NetApp public and hybrid cloud solutions

NetApp public and hybrid cloud solutions

NetApp
November 11, 2026

This PDF was generated from https://docs.netapp.com/us-en/netapp-solutions-cloud/index.html on November 11, 2026. Always check docs.netapp.com for the latest.

Table of Contents

NetApp public and hybrid cloud solutions . . . 1
What's new in NetApp public and hybrid cloud solutions . . . 2
18 August 2025 . . . 2
NetApp Hybrid Multicloud with VMware Solutions . . . 3
VMware for public cloud . . . 3
Overview of NetApp Hybrid Multicloud with VMware . . . 3
NetApp Solutions for VMware in Hyperscalers . . . 6
Supported Configurations for NetApp Hybrid Multicloud with VMware . . . 9
VMware in the hyperscalers configuration . . . 10
NetApp storage in public clouds . . . 34
Summary and Conclusion: Why NetApp Hybrid Multicloud with VMware . . . 132
VMware hybrid cloud use cases . . . 134
Use Cases for NetApp Hybrid Multicloud with VMware . . . 134
NetApp Solutions for VMware in Hyperscalers . . . 135
NetApp Solutions for Amazon VMware Managed Cloud (VMC) . . . 135
NetApp Solutions for Azure VMware Solution (AVS) . . . 136
NetApp Solutions for Google Cloud VMware Engine (GCVE) . . . 136
NetApp for AWS/VMC . . . 137
NetApp Capabilities for AWS VMC . . . 137
Protect workloads on AWS/VMC . . . 138
Migrate workloads on AWS/VMC . . . 251
NetApp for AWS / EVS . . . 266
Mount FSx ONTAP as external datastore for EVS . . . 266
Protect FSx ONTAP datastore for EVS . . . 290
NetApp for Azure / AVS . . . 323
NetApp Capabilities for Azure AVS . . . 323
Protect workloads on Azure / AVS . . . 325
Migrate workloads on Azure/AVS . . . 391
NetApp for GCP/GCVE . . . 405
NetApp Capabilities for Google Cloud Platform GCVE . . . 405
Protect workloads on GCP/GCVE . . . 406
Migrate workloads on GCP/GCVE . . . 437
Secure data with Google Cloud NetApp Volumes in Google Cloud . . . 454
NetApp Disaster Recovery . . . 517
3-2-1 Data Protection for VMware with SnapCenter Plug-in and NetApp Backup and Recovery to protect VMware workloads . . . 517
NetApp Disaster Recovery . . . 550
NetApp hybrid multicloud with Red Hat OpenShift container workloads . . . 591
Get started . . . 591
NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads . . . 591
Value propositions of NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads . . . 593
Supported Solutions of NetApp Hybrid Multicloud for Red Hat OpenShift Container workloads . . . 594
Versions of various components used in the solution validation . . . 597

Supported NetApp Storage integrations with Red Hat OpenShift Containers . . . 598
Additional Documentation . . . 601
Openshift for on-premises . . . 602
NetApp Solution with Red Hat OpenShift Container platform workloads on VMware . . . 602
Deploy and configure the Red Hat OpenShift Container platform on VMware . . . 603
Data protection using Trident Protect . . . 605
Hybrid cloud with self-managed components . . . 606
NetApp Solution with Red Hat OpenShift Container platform workloads in Hybrid Cloud . . . 606
Deploy and configure the Red Hat OpenShift Container platform on AWS . . . 608
Deploy and configure the Red Hat OpenShift Container Platform on Google Cloud . . . 610
Deploy and configure the Red Hat OpenShift Container platform on Azure . . . 612
Data protection using Trident Protect . . . 615
Hybrid cloud with provider-managed components . . . 615
NetApp Solution with Managed Red Hat OpenShift Container platform workloads . . . 615
Deploy and configure the Managed Red Hat OpenShift Container platform on AWS . . . 616
Deploy and configure OpenShift Dedicated on Google Cloud with Google Cloud NetApp Volumes . . . 619
Data protection . . . 621
Additional NetApp Hybrid Multicloud solutions for Red Hat OpenShift workloads . . . 621
Protect data using Trident Protect . . . 622
Data protection for Container Apps in OpenShift Container Platform using Trident Protect . . . 622
Protect VMs in Red Hat OpenShift Virtualization using Trident Protect . . . 634
Protect container apps using third-party tools . . . 654
Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP) . . . 654
Installation of OpenShift API for Data Protection (OADP) Operator . . . 656
Creating on-demand backup for Apps in OpenShift Container Platform . . . 665
Migrate an App from one cluster to another . . . 668
Restore an App from a backup . . . 672
Deleting backups and restores in using Velero . . . 679
Watch videos and demos on hybrid multicloud solutions with NetApp . . . 681
Watch videos on hybrid multicloud solutions for Amazon Web Services with NetApp . . . 681
Guest connected storage . . . 681
TCO and supplemental datastores . . . 681
Migration workflows (HCX) . . . 681
Watch videos on hybrid multicloud solutions for Azure VMware cloud with NetApp . . . 681
Supplemental datastores . . . 681
Disaster recovery . . . 681
Migration workflows (HCX) . . . 682
Watch videos on hybrid multicloud solutions for Red Hat OpenShift with NetApp . . . 682
Snapshots and restore . . . 682
Storage integration . . . 682
Failover and fail-back . . . 682
Solution automation . . . 682
Read cloud solution blogs from NetApp experts . . . 683
VMware integration with AWS and Amazon FSx for ONTAP . . . 683

Azure VMware solution . . . 683
Google Cloud VMware Engine . . . 683
Data protection and management . . . 683
NetApp for OpenShift container workloads . . . 683
VMware Sovereign Cloud . . . 683
Legal notices . . . 684
Copyright . . . 684
Trademarks . . . 684
Patents . . . 684
Privacy policy . . . 684

NetApp public and hybrid cloud solutions

What’s new in NetApp public and hybrid cloud solutions

Learn what’s new in public and hybrid cloud solutions.

18 August 2025

The NetApp Solutions site is now the NetApp solutions family, which includes the following sites:

- NetApp artificial intelligence solutions
- NetApp container solutions
- NetApp data management solutions
- NetApp database solutions
- NetApp public and hybrid cloud solutions
- NetApp solutions for SAP
- NetApp virtualization solutions

NetApp Hybrid Multicloud with VMware Solutions

VMware for public cloud

Overview of NetApp Hybrid Multicloud with VMware

Most IT organizations follow the hybrid cloud-first approach. These organizations are in a transformation phase and customers are evaluating their current IT landscape and then migrating their workloads to the cloud based on the assessment and discovery exercise.

The factors for customers migrating to the cloud can include elasticity and burst, data center exit, data center consolidation, end-of-life scenarios, mergers, acquisitions, and so on. The reason for this migration can vary based on each organization and their respective business priorities. When moving to the hybrid cloud, choosing the right storage in the cloud is very important in order to unleash the power of cloud deployment and elasticity.

VMware Cloud options in Public Cloud

This section describes how each of the cloud providers support a VMware Software Defined Data Center (SDDC) and/or VMware Cloud Foundation (VCF) stack within their respective public cloud offerings.

Azure VMware Solution

Azure VMware Solution is a hybrid cloud service that allows for fully functioning VMware SDDCs within the Microsoft Azure public cloud. Azure VMware Solution is a first-party solution fully managed and supported by Microsoft, verified by VMware leveraging Azure infrastructure. This means that when Azure VMware Solution is deployed, customer's get VMware's ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data center facilities and proximity to the rich ecosystem of native Azure services and solutions.

VMware Cloud on AWS

VMware Cloud on AWS brings VMware's enterprise-class SDDC software to the AWS Cloud with optimized access to native AWS services. Powered by VMware Cloud Foundation, VMware Cloud on AWS integrates VMware's compute, storage, and network virtualization products (VMware vSphere, VMware vSAN, and VMware NSX) along with VMware vCenter Server management, optimized to run on dedicated, elastic, bare-metal AWS infrastructure.

Google Cloud VMware Engine

Google Cloud VMware Engine is an infrastructure-as-a-service (IaaS) offering built on Google Cloud's highly performant scalable infrastructure and VMware Cloud Foundation stack – VMware vSphere, vCenter, vSAN, and NSX-T. This service enables a fast path to the cloud, seamlessly migrating or extending existing VMware workloads from on-premises environments to Google Cloud Platform without the cost, effort ,or risk of rearchitecting applications or retooling operations. It is a service sold and supported by Google, working closely with VMware.



SDDC private cloud and NetApp Cloud Volumes colocation provides the best performance with minimal network latency.

Did you know?

Regardless of the cloud used, when a VMware SDDC is deployed, the initial cluster includes the following products:

- VMware ESXi hosts for compute virtualization with a vCenter Server appliance for management
- VMware vSAN hyper-converged storage incorporating the physical storage assets of each ESXi host
- VMware NSX for virtual networking and security with an NSX Manager cluster for management

Storage configuration

For customers planning to host storage-intensive workloads and scale out on any cloud-hosted VMware solution, the default hyper-converged infrastructure dictates that the expansion should be on both the compute and storage resources.

By integrating with NetApp Cloud Volumes, such as Azure NetApp Files, Amazon FSx ONTAP, Cloud Volumes ONTAP (available in all three major hyperscalers), and Google Cloud NetApp Volumes for Google Cloud, customers now have options to independently scale their storage separately, and only add compute nodes to the SDDC cluster as needed.

Notes:

- VMware does not recommend unbalanced cluster configurations, hence expanding storage means adding more hosts, which implies more TCO.
- Only one vSAN environment is possible. Therefore, all storage traffic will compete directly with production workloads.
- There is no option to provide multiple performance tiers to align application requirements, performance, and cost.
- It is very easy to reach the limits of storage capacity of vSAN built on top of the cluster hosts. Use NetApp Cloud Volumes to scale storage to either host active datasets or tier cooler data to persistent storage.

Azure NetApp Files, Amazon FSx ONTAP, Cloud Volumes ONTAP (available in all three major hyperscalers), and Google Cloud NetApp Volumes for Google Cloud can be used in conjunction with guest VMs. This hybrid storage architecture consists of a vSAN datastore that holds the guest operating system and application binary data. The application data is attached to the VM through a guest-based iSCSI initiator or the NFS/SMB mounts that communicate directly with Amazon FSx ONTAP, Cloud Volume ONTAP, Azure NetApp Files and Google Cloud NetApp Volumes for Google Cloud respectively. This configuration allows you to easily overcome challenges with storage capacity as with vSAN, the available free space depends on the slack space and storage policies used.

Let’s consider a three-node SDDC cluster on VMware Cloud on AWS:

- The total raw capacity for a three-node SDDC = 31.1TB (roughly 10TB for each node).
- The slack space to be maintained before additional hosts are added = 25% = (.25 x 31.1TB) = 7.7TB.
- The usable raw capacity after slack space deduction = 23.4TB
- The effective free space available depends on the storage policy applied.

  For example:

  - RAID 0 = effective free space = 23.4TB (usable raw capacity/1)
  - RAID 1 = effective free space = 11.7TB (usable raw capacity/2)

- RAID 5 = effective free space = 17.5TB (usable raw capacity/1.33)

Thus, using NetApp Cloud Volumes as guest-connected storage would help in expanding the storage and optimizing the TCO while meeting the performance and data protection requirements.

In-guest storage was the only available option at the time this document was written.

Points to Remember

- In hybrid storage models, place tier 1 or high priority workloads on vSAN datastore to address any specific latency requirements because they are part of the host itself and within proximity. Use in-guest mechanisms for any workload VMs for which transactional latencies are acceptable.
- Use NetApp SnapMirror technology to replicate the workload data from the on-premises ONTAP system to Cloud Volumes ONTAP or Amazon FSx ONTAP to ease migration using block-level mechanisms. This does not apply to Azure NetApp Files and Google Cloud NetApp Volumes. For migrating data to Azure NetApp Files or Google Cloud NetApp Volumes, use NetApp XCP, NetApp Copy and Sync, rysnc or robocopy depending on the file protocol used.
- Testing shows 2-4ms additional latency while accessing storage from the respective SDDCs. Factor this additional latency into the application requirements when mapping the storage.
- For mounting guest-connected storage during test failover and actual failover, make sure iSCSI initiators are reconfigured, DNS is updated for SMB shares, and NFS mount points are updated in fstab.
- Make sure that in-guest Microsoft Multipath I/O (MPIO), firewall, and disk timeout registry settings are configured properly inside the VM.

This applies to guest connected storage only.

Benefits of NetApp cloud storage

NetApp cloud storage offers the following benefits:

- Improves compute-to-storage density by scaling storage independently of compute.
- Allows you to reduce the host count, thus reducing the overall TCO.
- Compute node failure does not impact storage performance.
- The volume reshaping and dynamic service-level capability of Azure NetApp Files allows you to optimize cost by sizing for steady-state workloads, and thus preventing over provisioning.
- The storage efficiencies, cloud tiering, and instance-type modification capabilities of Cloud Volumes ONTAP allow optimal ways of adding and scaling storage.
- Prevents over provisioning storage resources are added only when needed.
- Efficient Snapshot copies and clones allow you to rapidly create copies without any performance impact.
- Helps address ransomware attacks by using quick recovery from Snapshot copies.
- Provides efficient incremental block transfer-based regional disaster recovery and integrated backup block level across regions provides better RPO and RTOs.

Assumptions

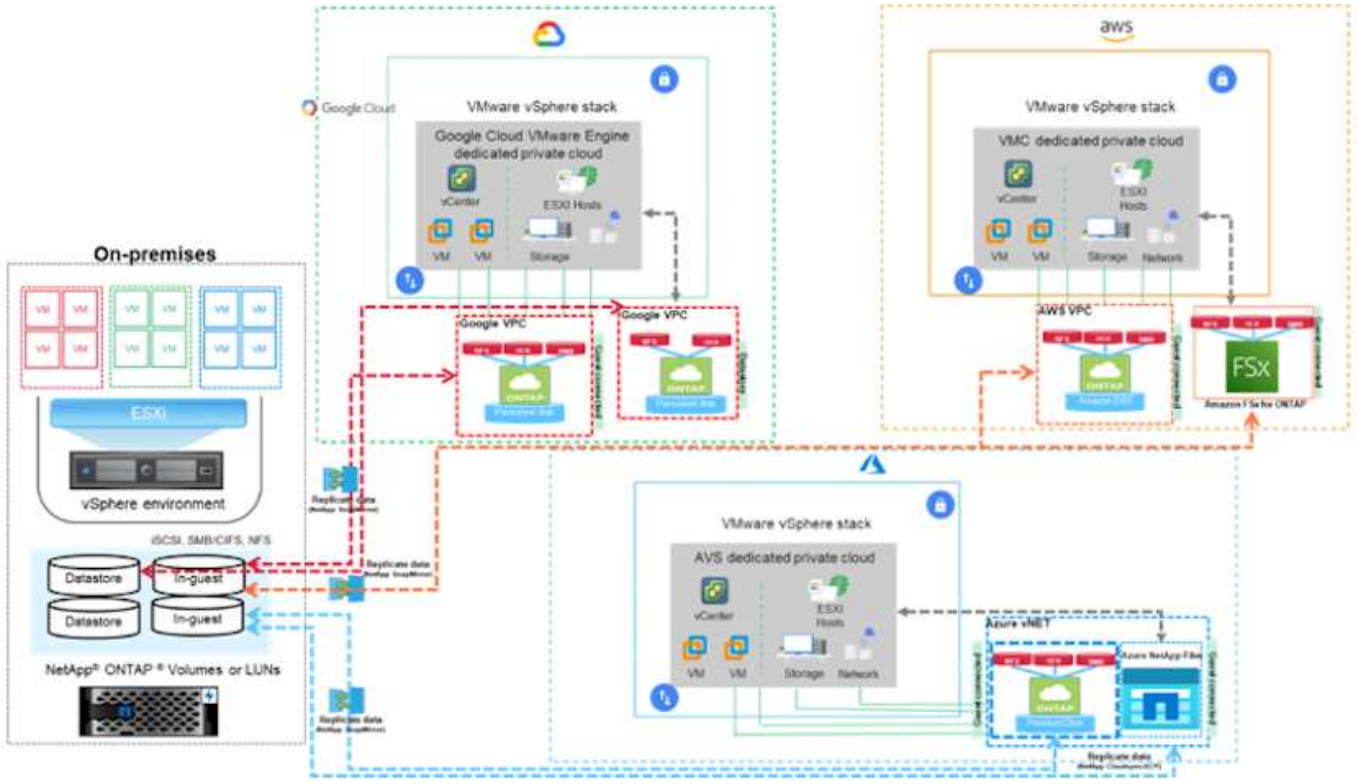
- SnapMirror technology or other relevant data migration mechanisms are enabled. There are many connectivity options, from on-premises to any hyperscaler cloud. Use the appropriate path and work with the relevant networking teams.

- In-guest storage was the only available option at the time this document was written.

> Engage NetApp solution architects and respective hyperscaler cloud architects for planning and sizing of storage and the required number of hosts. NetApp recommends identifying the storage performance requirements before using the Cloud Volumes ONTAP sizer to finalize the storage instance type or the appropriate service level with the right throughput.

Detailed architecture

From a high-level perspective, this architecture (shown in the figure below) covers how to achieve hybrid Multicloud connectivity and app portability across multiple cloud providers using NetApp Cloud Volumes ONTAP, Google Cloud NetApp Volumes for Google Cloud and Azure NetApp Files as an additional in-guest storage option.



NetApp Solutions for VMware in Hyperscalers

Learn more about the capabilities that NetApp brings to the three (3) primary hyperscalers - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Pick your cloud and let NetApp do the rest!



To see the capabilities for a specific hyperscaler, click on the appropriate tab for that hyperscaler.

Jump to the section for the desired content by selecting from the following options:

- VMware in the Hyperscalers Configuration
- NetApp Storage Options
- NetApp / VMware Cloud Solutions

VMware in the Hyperscalers Configuration

As with on-premises, planning a cloud based virtualization environment is critical for a successful production-ready environment for creating VMs and migration.

AWS / VMC

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.

In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to AWS VMC.

The setup process can be broken down into the following steps:

- Deploy and Configure VMware Cloud for AWS
- Connect VMware Cloud to FSx ONTAP

View the detailed configuration steps for VMC.

Azure / AVS

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.

In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to Azure VMware Solution.

The setup process can be broken down into the following steps:

- Register the resource provider and create a private cloud
- Connect to a new or existing ExpressRoute virtual network gateway
- Validate the network connectivity and access the private cloud

View the detailed configuration steps for AVS.

GCP / GCVE

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.

In-guest storage is the only supported method of connecting Cloud Volumes ONTAP and Google Cloud NetApp Volumes to GCVE.

The setup process can be broken down into the following steps:

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed configuration steps for GCVE.

NetApp Storage Options

NetApp storage can be utilized in several ways - either as guest connected or as a supplemental NFS datastore - within each of the 3 major hyperscalers.

Please visit Supported NetApp Storage Options for more information.

AWS / VMC

AWS supports NetApp storage in the following configurations:

- FSx ONTAP as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- FSx ONTAP as a supplemental NFS datastore

View the detailed guest connect storage options for VMC.
View the detailed supplemental NFS datastore options for VMC.

Azure / AVS

Azure supports NetApp storage in the following configurations:

- Azure NetApp Files (ANF) as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- Azure NetApp Files (ANF) as a supplemental NFS datastore

View the detailed guest connect storage options for AVS.
View the detailed supplemental NFS datastore options for AVS.

GCP / GCVE

Google Cloud supports NetApp storage in the following configurations:

- Cloud Volumes ONTAP (CVO) as guest connected storage
- Google Cloud NetApp Volumes (NetApp Volumes) as guest connected storage
- Google Cloud NetApp Volumes (NetApp Volumes) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.
View the detailed supplemental NFS datastore options for GCVE.

Read more about Google Cloud NetApp Volumes datastore support for Google Cloud VMware Engine (NetApp blog) or How to use Google Cloud NetApp Volumes as datastores for Google Cloud VMware Engine (Google blog)

NetApp / VMware Cloud Solutions

With NetApp and VMware cloud solutions, many use cases are simple to deploy in your hyperscaler of choice. VMware defines the primary cloud workload use-cases as:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

AWS / VMC

Browse the NetApp solutions for AWS / VMC

Azure / AVS

Browse the NetApp solutions for Azure / AVS

GCP / GCVE

Browse the NetApp solutions for Google Cloud Platform (GCP) / GCVE

Supported Configurations for NetApp Hybrid Multicloud with VMware

Understanding the combinations for NetApp storage support in the major hyperscalers.

	Guest Connected	Supplemental NFS Datastore
AWS	CVO FSx ONTAP Details	FSx ONTAP Details
Azure	CVO ANF Details	ANF Details

GCP	CVO NetApp Volumes Details	NetApp Volumes Details

VMware in the hyperscalers configuration

Configuring the virtualization environment in the cloud provider

Details for how to configure the virtualization environment in each of the supported hyperscalers are covered here.

AWS / VMC

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.

> In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to AWS VMC.

The setup process can be broken down into the following steps:

- Deploy and Configure VMware Cloud for AWS
- Connect VMware Cloud to FSx ONTAP

View the detailed configuration steps for VMC.

Azure / AVS

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.

> In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to Azure VMware Solution.

The setup process can be broken down into the following steps:

- Register the resource provider and create a private cloud
- Connect to a new or existing ExpressRoute virtual network gateway
- Validate the network connectivity and access the private cloud

View the detailed configuration steps for AVS.

GCP / GCVE

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.

> In-guest storage is the only supported method of connecting Cloud Volumes ONTAP and Google Cloud NetApp Volumes to GCVE.

The setup process can be broken down into the following steps:

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed configuration steps for GCVE.

Deploy and configure the Virtualization Environment on AWS

As with on-premises, planning VMware Cloud on AWS is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.

> In-guest storage is currently the only supported method of connecting Cloud Volumes ONTAP (CVO) to AWS VMC.

The setup process can be broken down into the following steps:

▼ **Deploy and configure VMware Cloud for AWS**

VMware Cloud on AWS provides for a cloud native experience for VMware based workloads in the AWS ecosystem. Each VMware Software-Defined Data Center (SDDC) runs in an Amazon Virtual Private Cloud (VPC) and provides a full VMware stack (including vCenter Server), NSX-T software-defined networking, vSAN software-defined storage, and one or more ESXi hosts that provide compute and storage resources to your workloads.

This section describes how to set up and manage VMware Cloud on AWS and use it in combination with Amazon FSx ONTAP and/or Cloud Volumes ONTAP on AWS with in-guest storage.

> In-guest storage is currently the only supported method of connecting Cloud Volumes ONTAP (CVO) to AWS VMC.

The setup process can be broken down into three parts:

▼ **Register for an AWS Account**

Register for an Amazon Web Services Account.

You need an AWS account to get started, assuming there isn't one created already. New or existing, you need administrative privileges in the account for many steps in this procedure. See this link for more information regarding AWS credentials.

▼ **Register for a My VMware Account**

Register for a My VMware account.

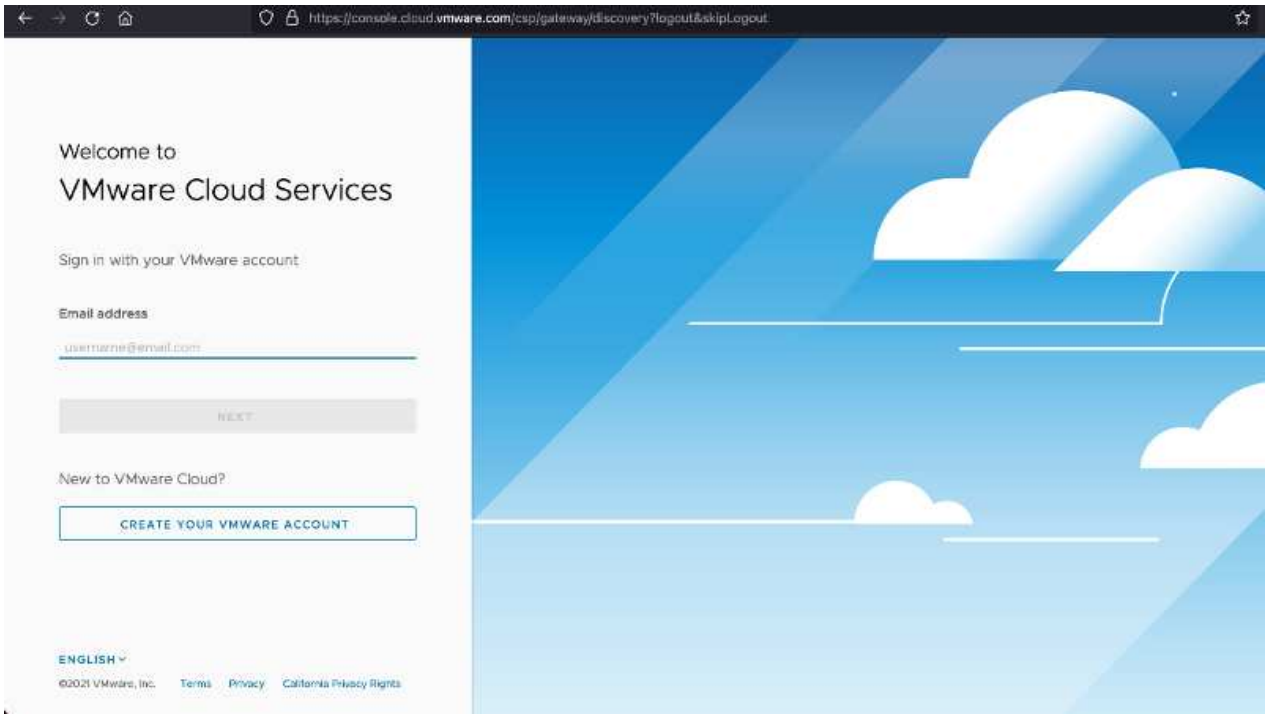
For access to VMware's cloud portfolio (including VMware Cloud on AWS), you need a VMware customer account or a My VMware account. If you have not already done so, create a VMware account here.

▼ **Provision SDDC in VMware Cloud**

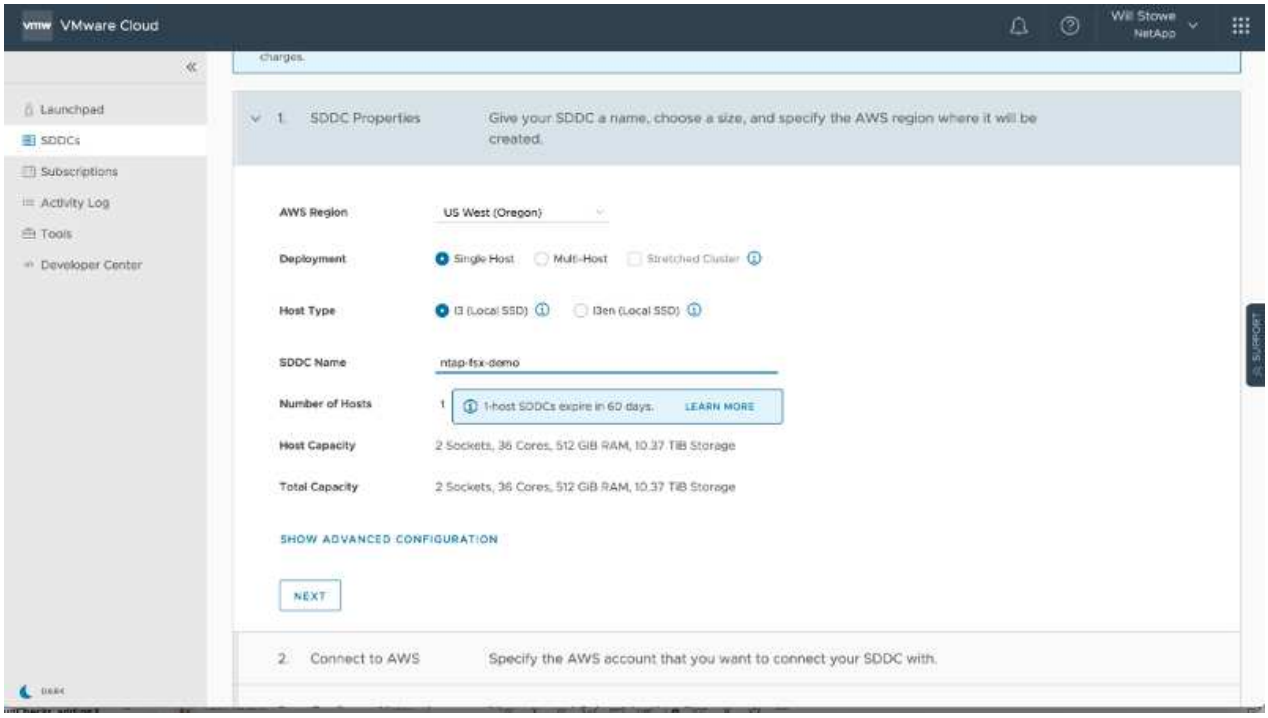
After the VMware account is configured and proper sizing is performed, deploying a Software-Defined Data Center is the obvious next step for using the VMware Cloud on AWS service. To create an SDDC, pick an AWS region to host it, give the SDDC a name, and specify how many ESXi hosts you want the SDDC to contain. If you don't already have an AWS account, you can still create a starter configuration

SDDC that contains a single ESXi host.

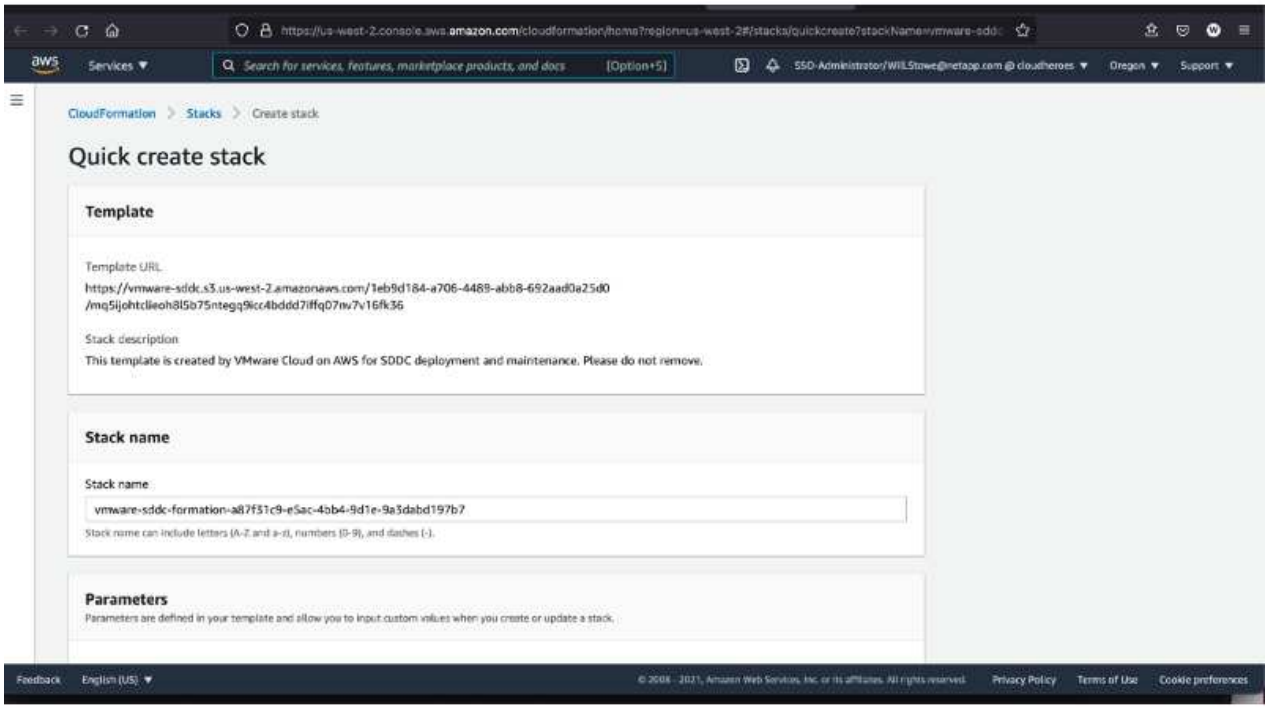
1. Log into the VMware Cloud Console using your existing or newly created VMware credentials.

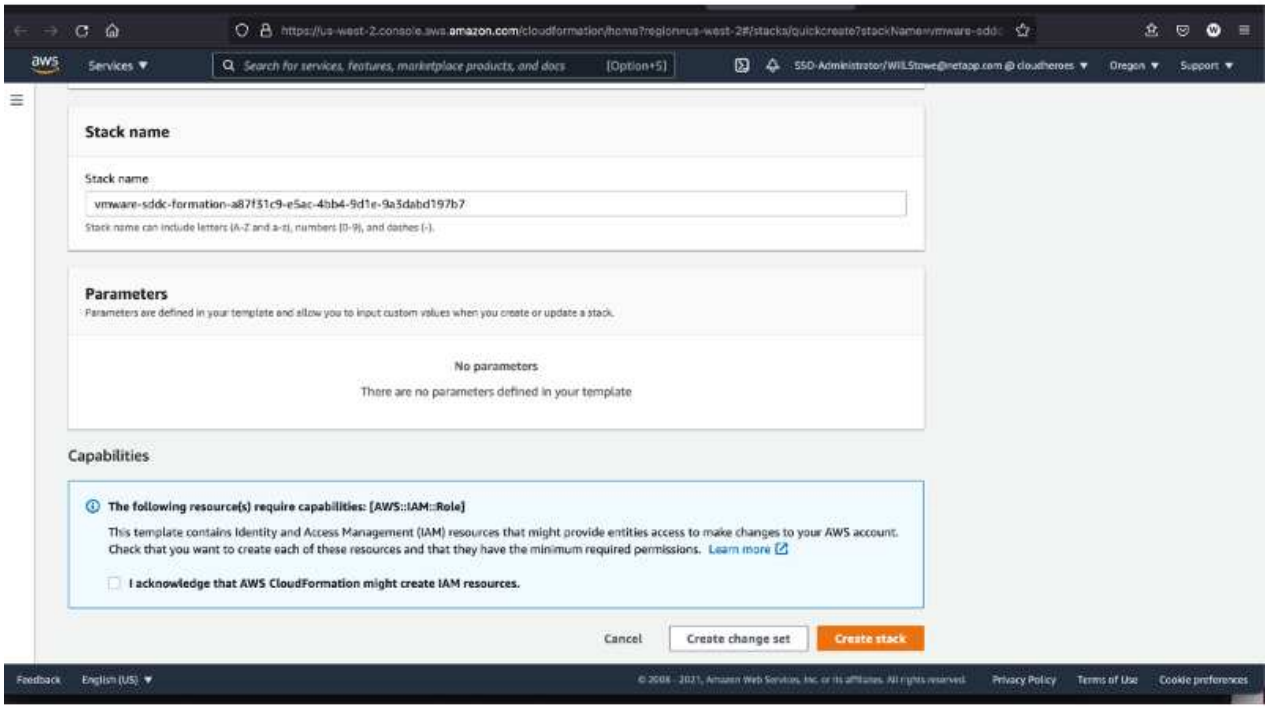


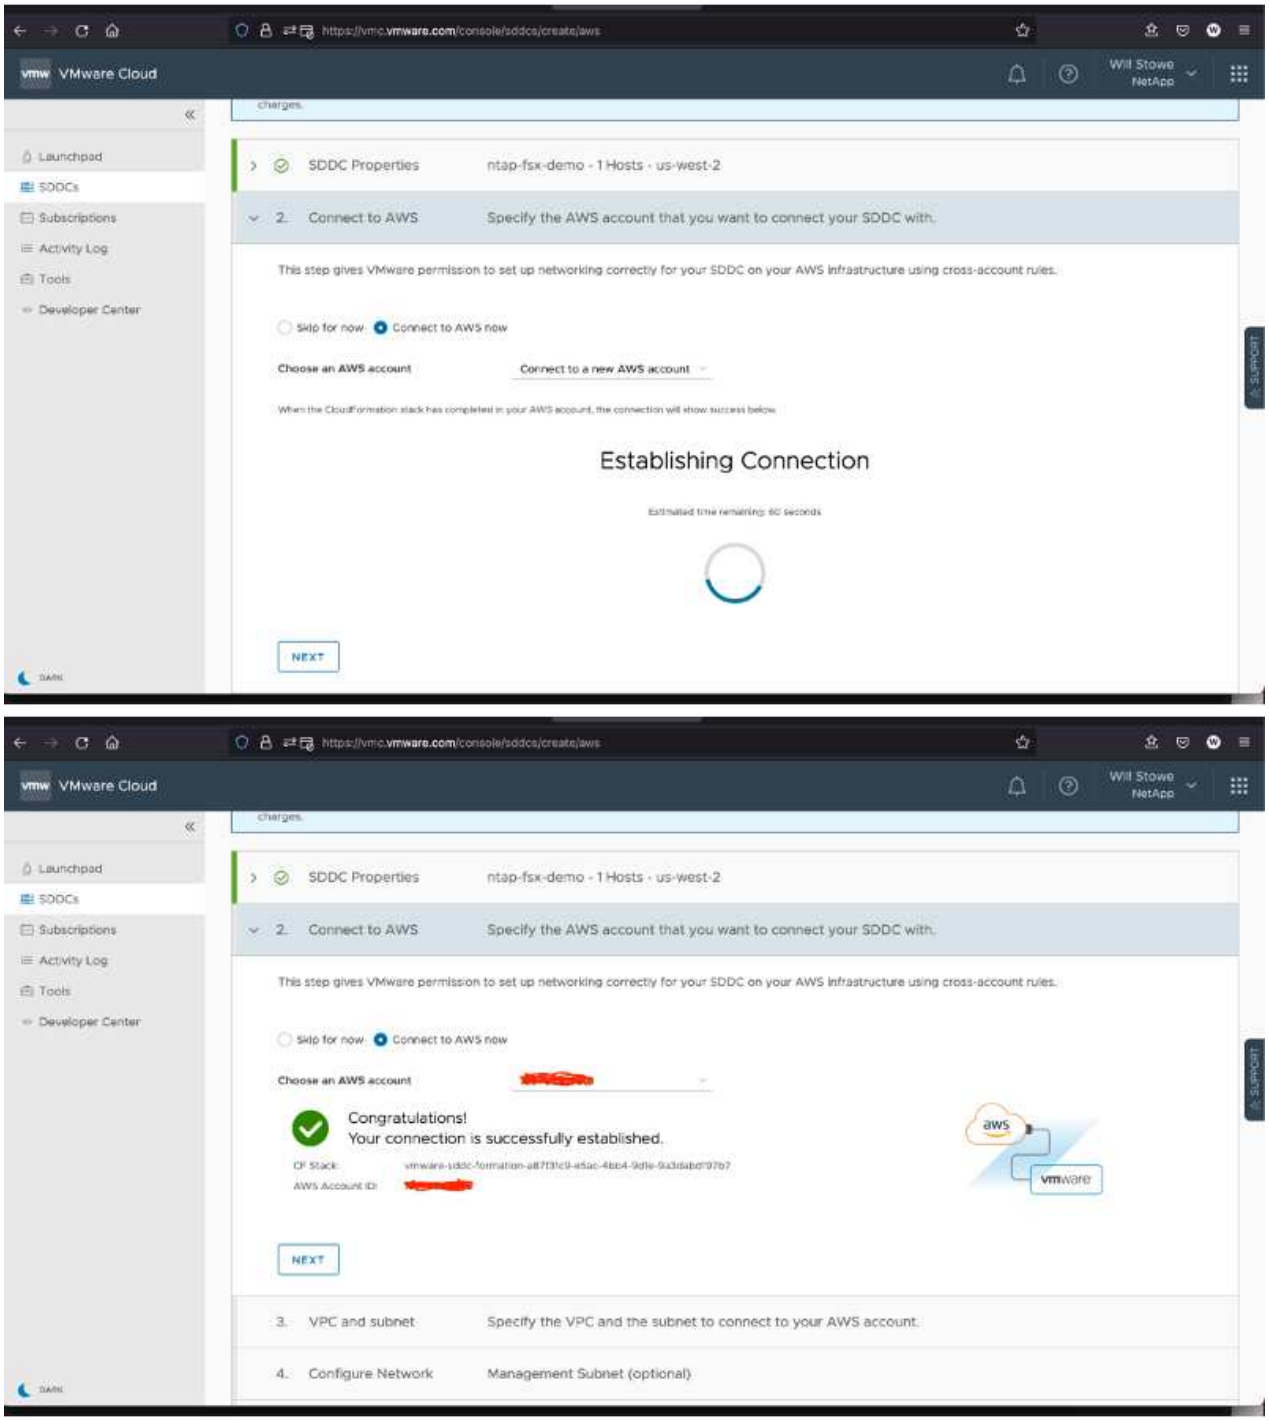
2. Configure the AWS region, deployment, and host type and the SDDC name:



3. Connect to the desired AWS account and execute the AWS Cloud Formation stack.

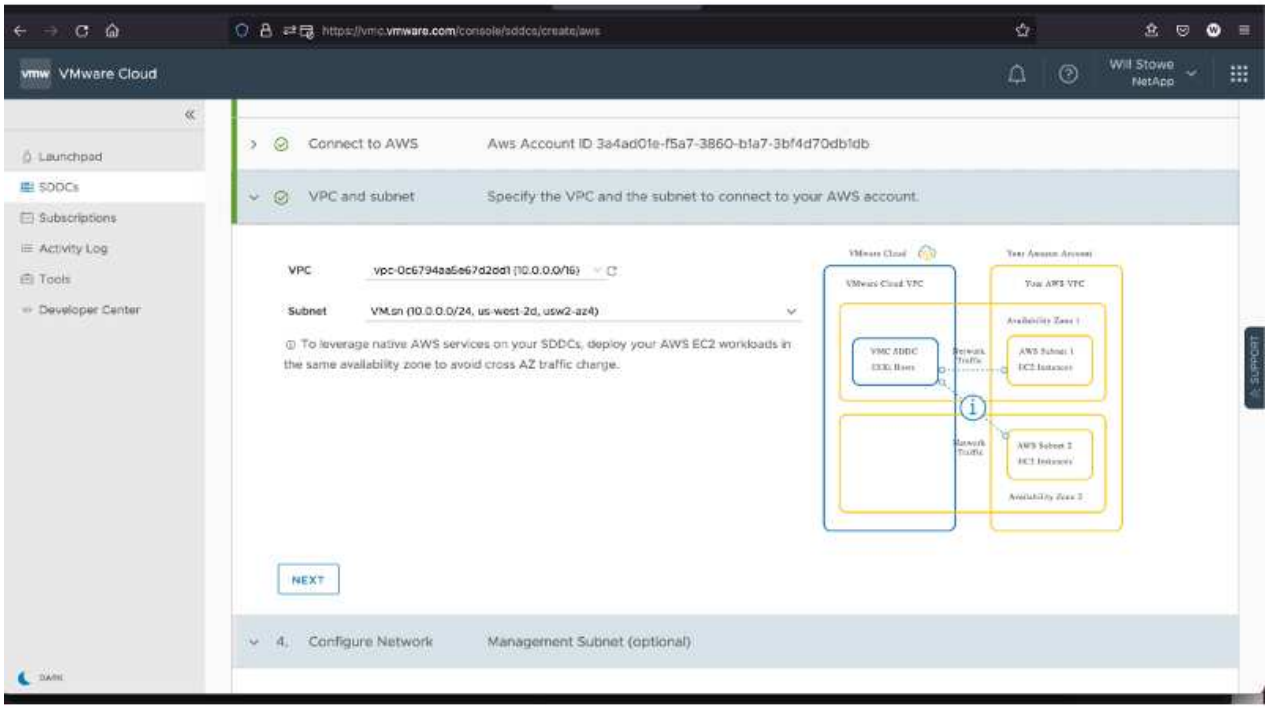

Services
Search for services, features, marketplace products, and docs [Option+S]
Oregon
Support
CloudFormation > Stacks > Create stack
Quick create stack
Template
Template URL
https://vmware-sddc.s3.us-west-2.amazonaws.com/1eb9d184-a706-4489-abb8-692aad0a25d0/mq5ljohtclieoh8l5b75ntegq9icc4bddd7iffq07nv7v16fk36
Stack description
This template is created by VMware Cloud on AWS for SDDC deployment and maintenance. Please do not remove.
Stack name
Stack name
vmware-sddc-formation-a87f31c9-e5ac-4bb4-9d1e-9a3dabd197b7
Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).
Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.
Feedback
English (US)
Privacy Policy
Terms of Use
Cookie preferences


Services
Search for services, features, marketplace products, and docs [Option+S]
Oregon
Support
Stack name
Stack name
vmware-sddc-formation-a87f31c9-e5ac-4bb4-9d1e-9a3dabd197b7
Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).
Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.
No parameters
There are no parameters defined in your template
Capabilities
The following resource(s) require capabilities: [AWS::IAM::Role]
This template contains Identity and Access Management (IAM) resources that might provide entities access to make changes to your AWS account. Check that you want to create each of these resources and that they have the minimum required permissions. Learn more
I acknowledge that AWS CloudFormation might create IAM resources.
Cancel
Create change set
Create stack
Feedback
English (US)
Privacy Policy
Terms of Use
Cookie preferences

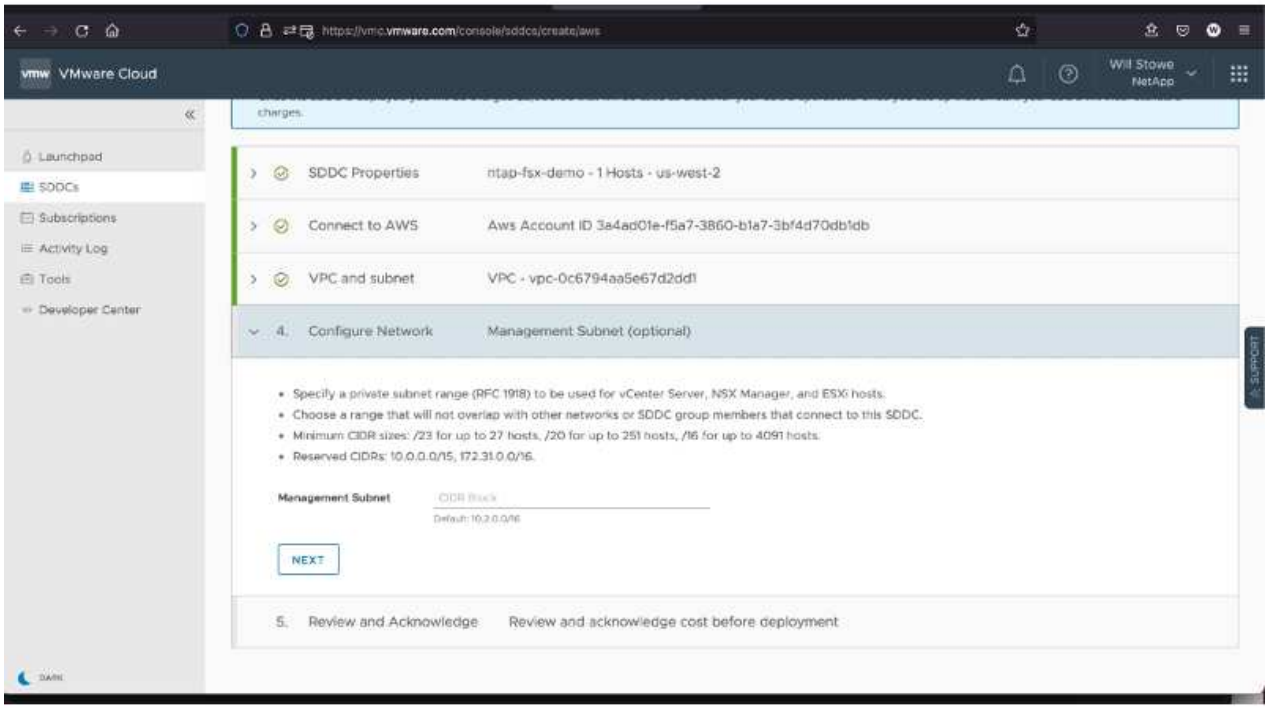


Single-host configuration is used in this validation.

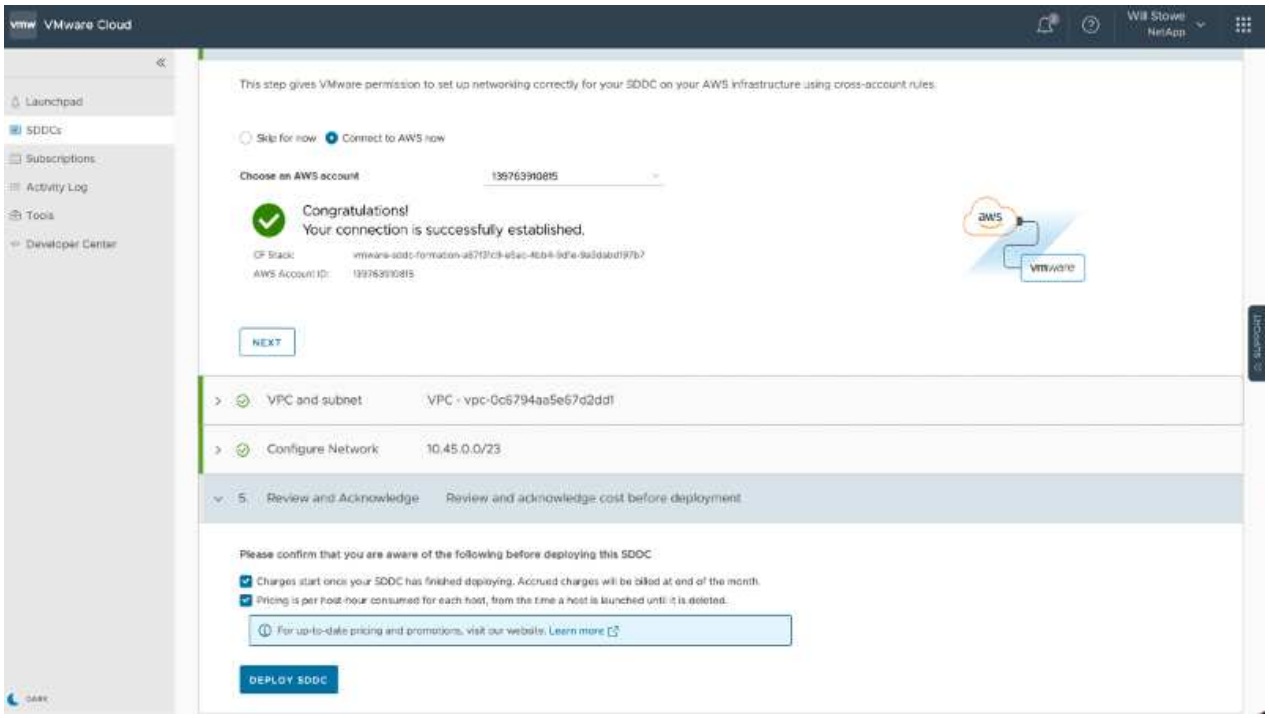
4. Select the desired AWS VPC to connect the VMC environment with.



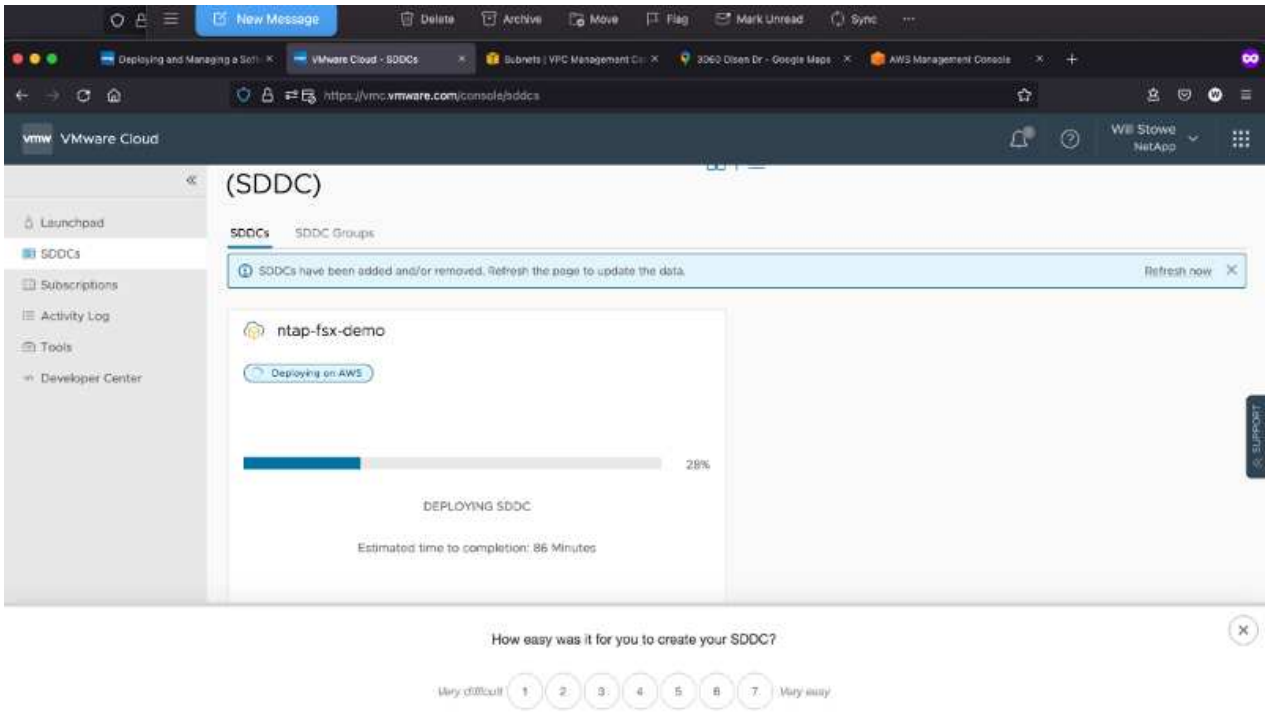
5. Configure the VMC Management Subnet; this subnet contains VMC-managed services like vCenter, NSX, and so on. Do not choose an overlapping address space with any other networks that need connectivity to the SDDC environment. Finally, follow the recommendations for CIDR size notated below.



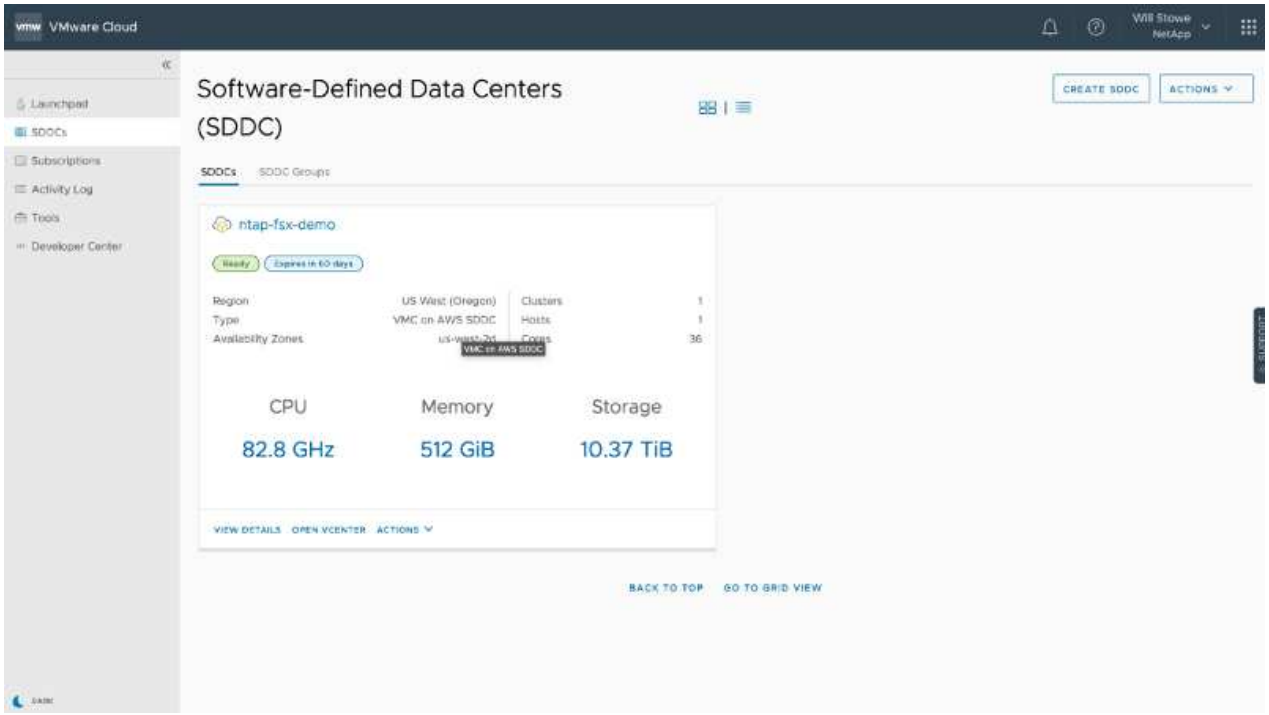
6. Review and acknowledge the SDDC configuration, and then click deploy the SDDC.



The deployment process typically takes approximately two hours to complete.



7. After completion, the SDDC is ready for use.

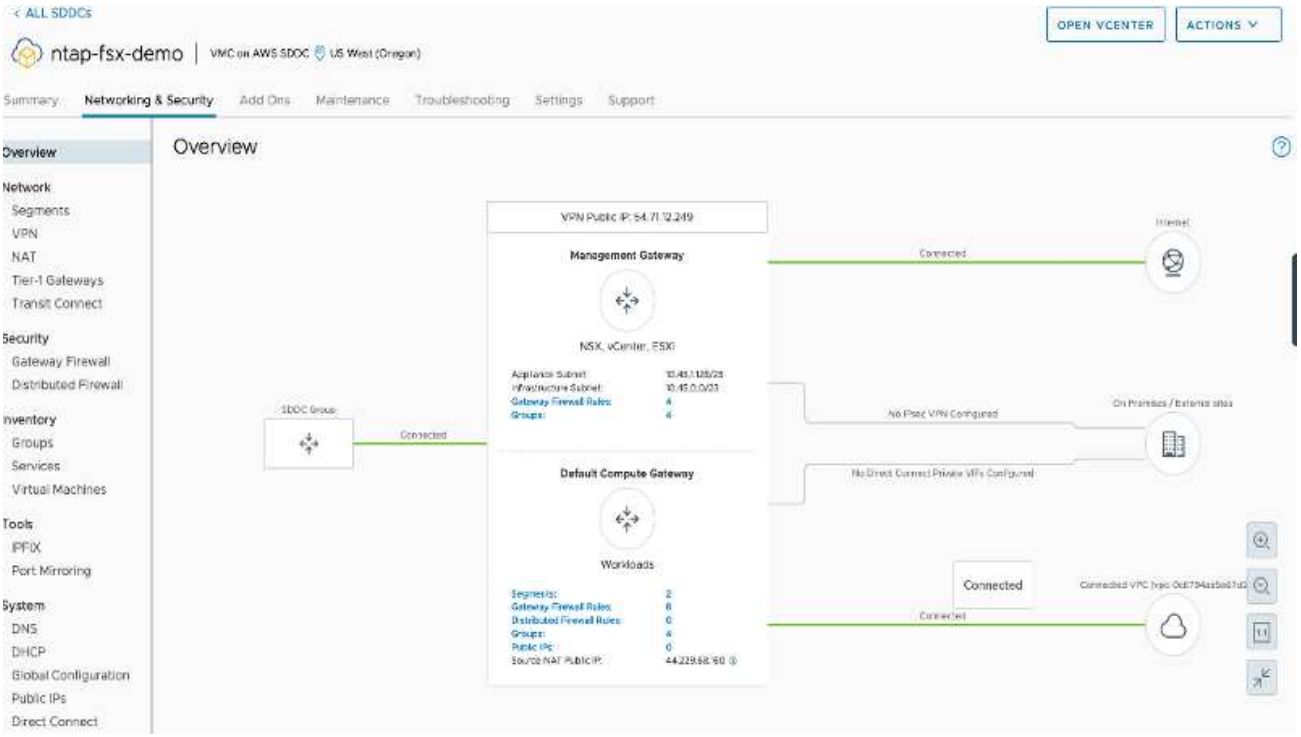


For a step-by-step guide on SDDC deployment, see Deploy an SDDC from the VMC Console.

▼ **Connect VMware Cloud to FSx ONTAP**

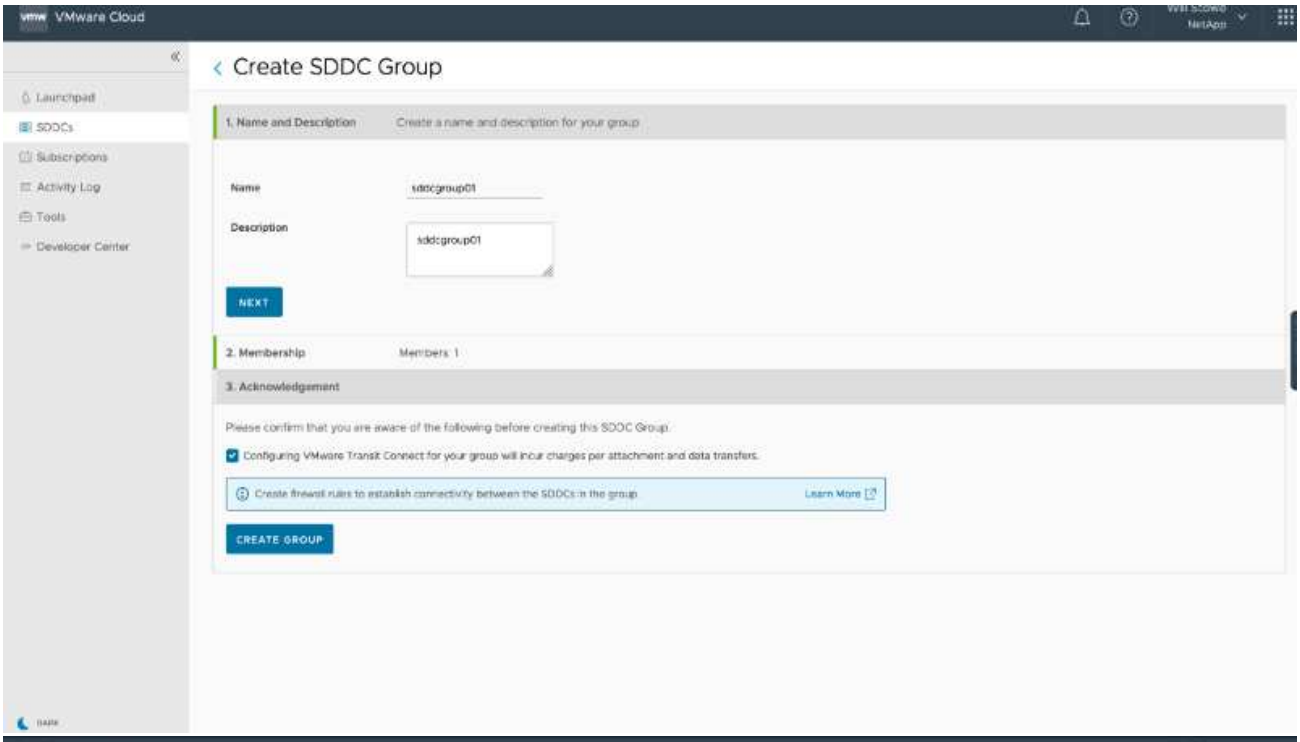
To connect VMware Cloud to FSx ONTAP, complete the following steps:

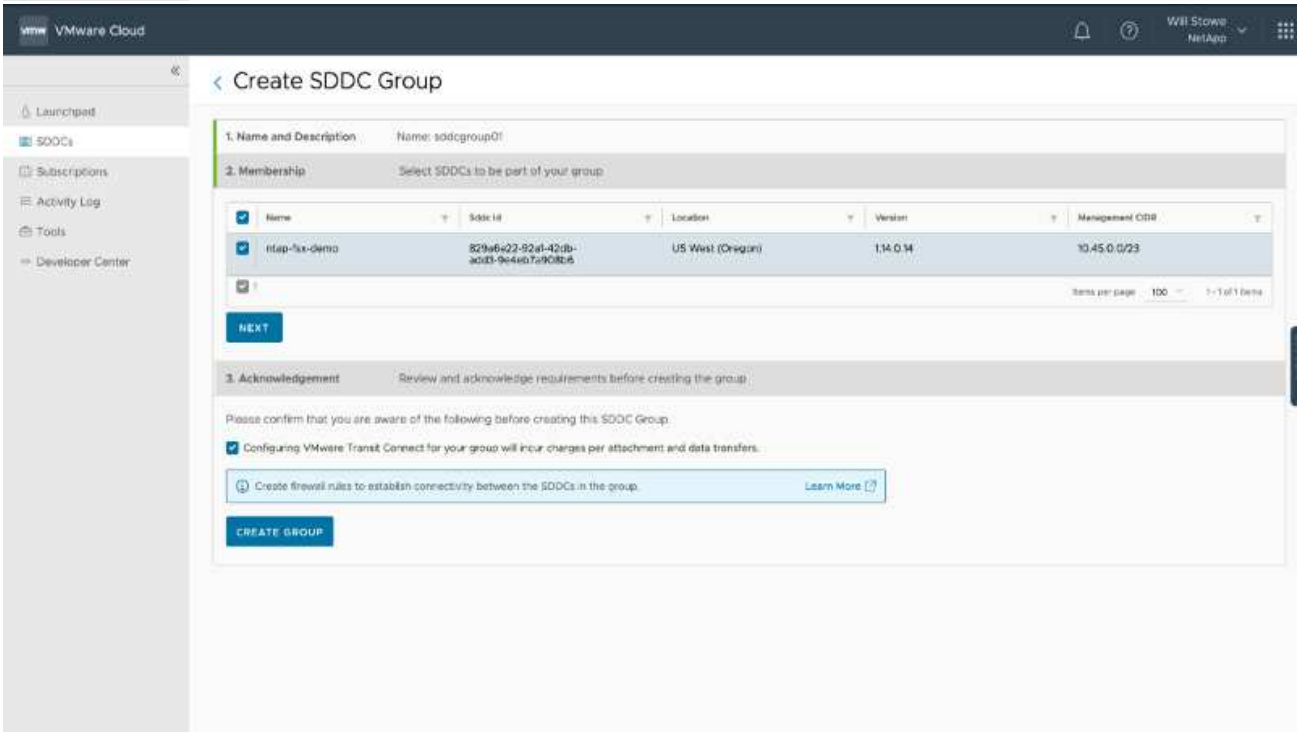
1. With VMware Cloud deployment completed and connected to AWS VPC, you must deploy Amazon FSx ONTAP into a new VPC rather than the original connected VPC (see the screenshot below). FSx (NFS and SMB floating IPs) is not accessible if it is deployed in the connected VPC. Keep in mind that ISCSI endpoints like Cloud Volumes ONTAP work just fine from the connected VPC.

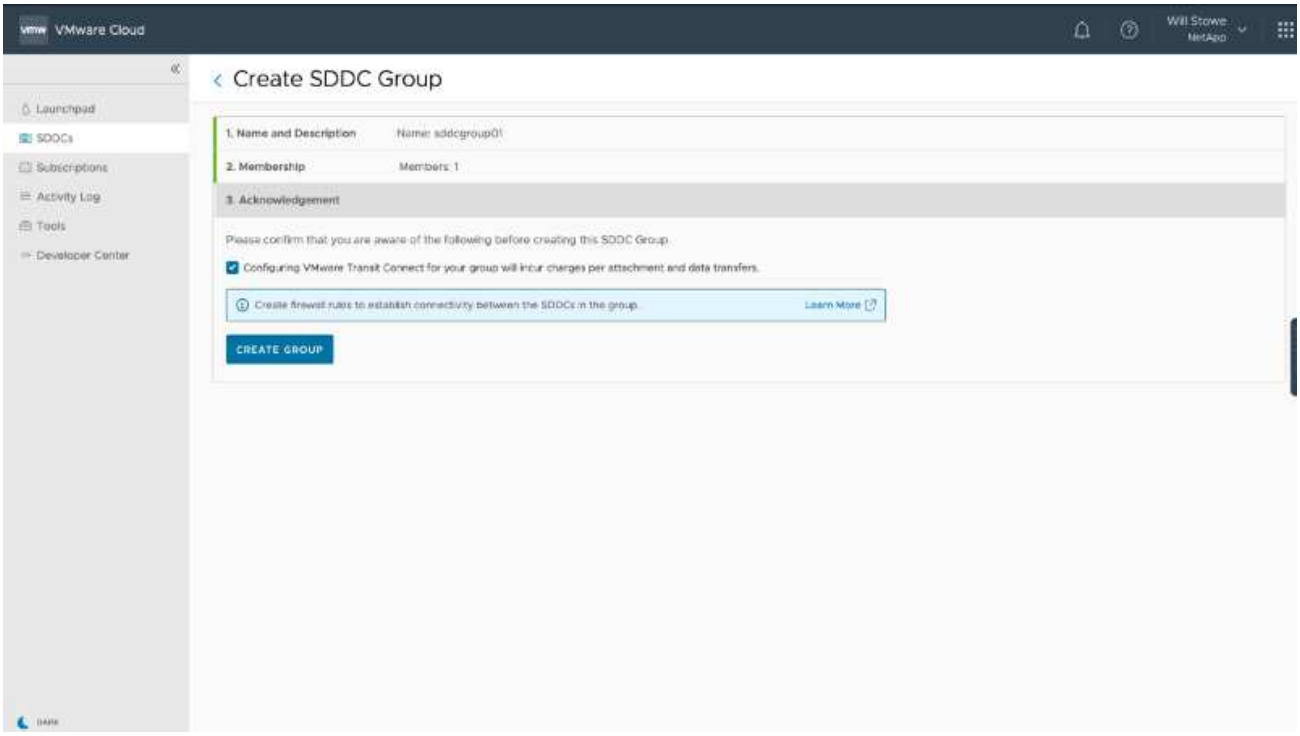


2. Deploy an additional VPC in the same region, and then deploy Amazon FSx ONTAP into the new VPC.

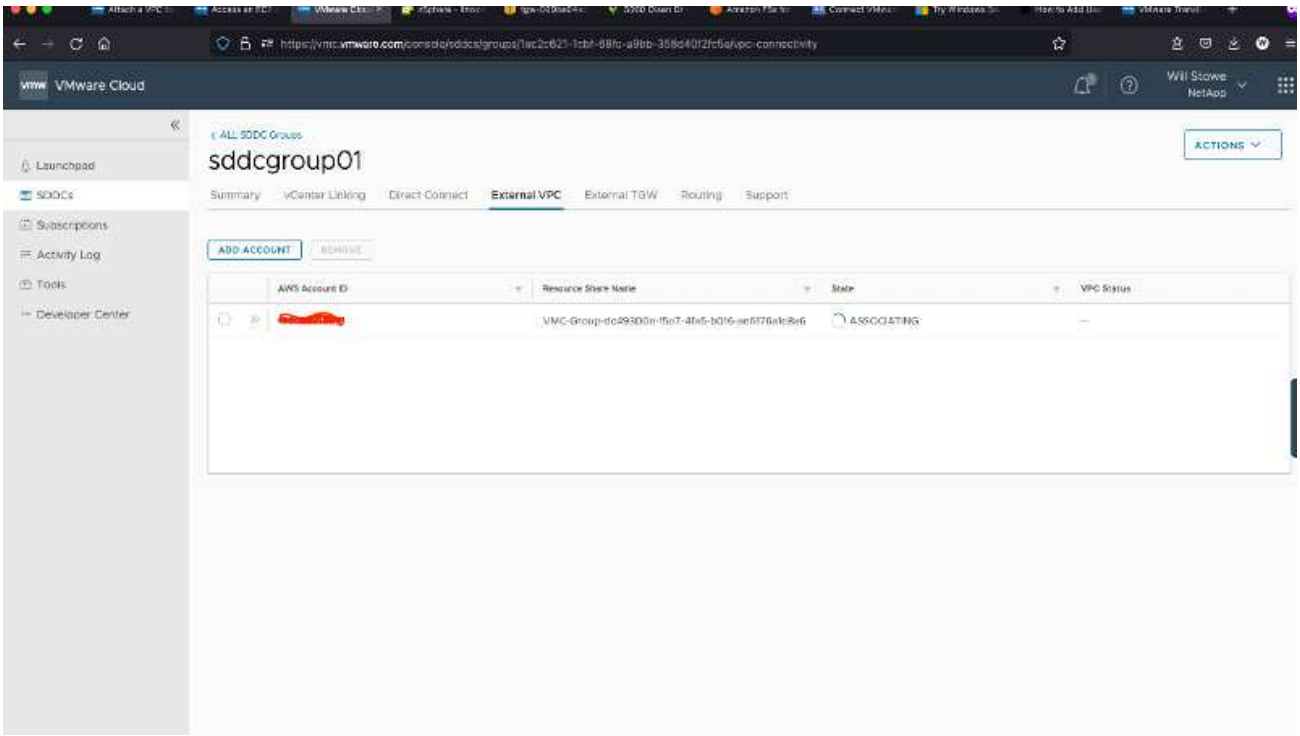
   Configuration of an SDDC group in the VMware Cloud console enables the networking configuration options required to connect to the new VPC where FSx is deployed. In step 3, verify that "Configuring VMware Transit Connect for your group will incur charges per attachment and data transfers" is checked, and then choose Create Group. The process can take a few minutes to complete.

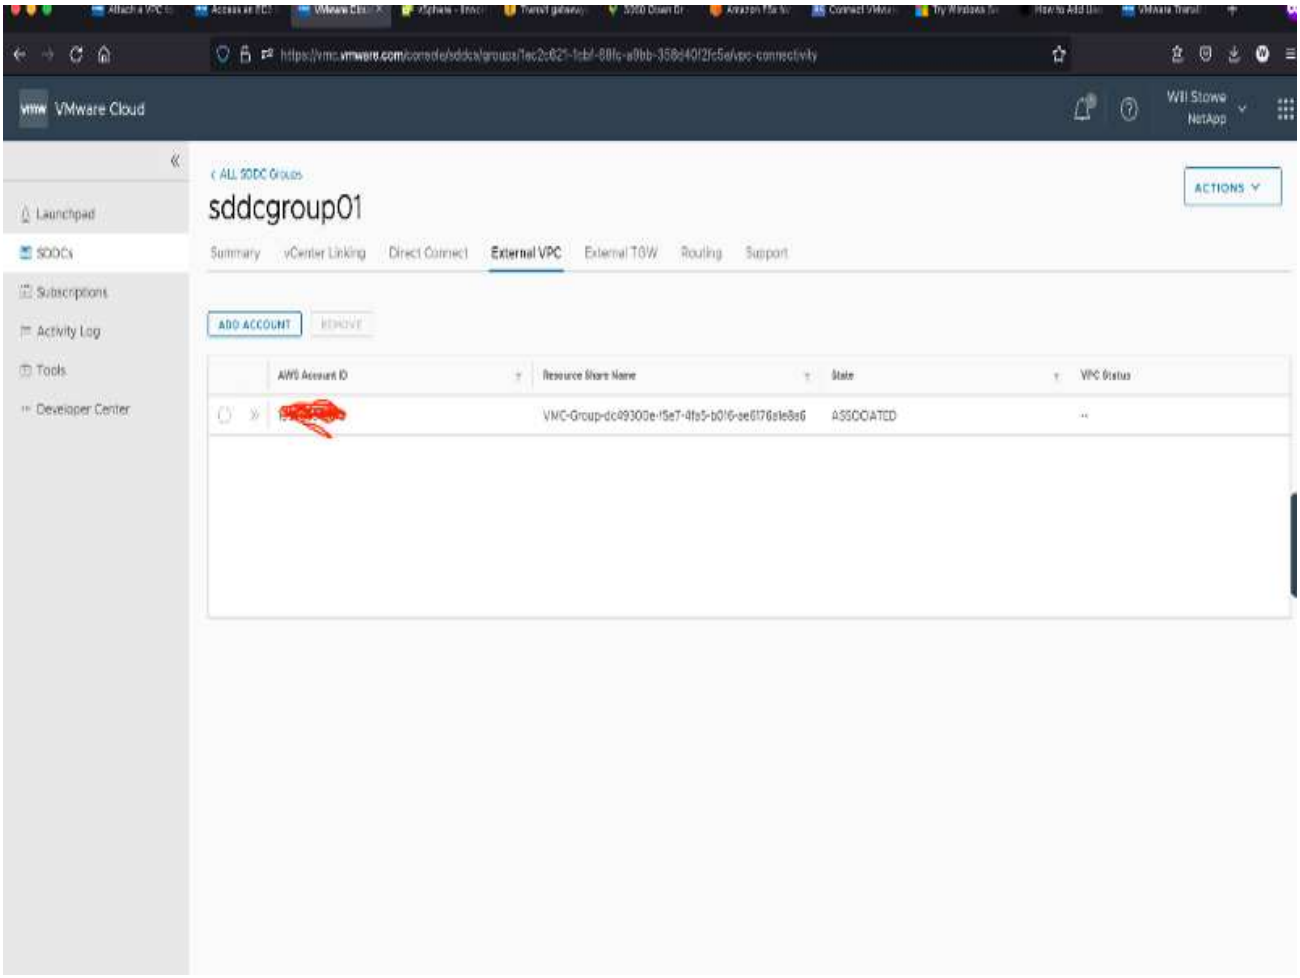




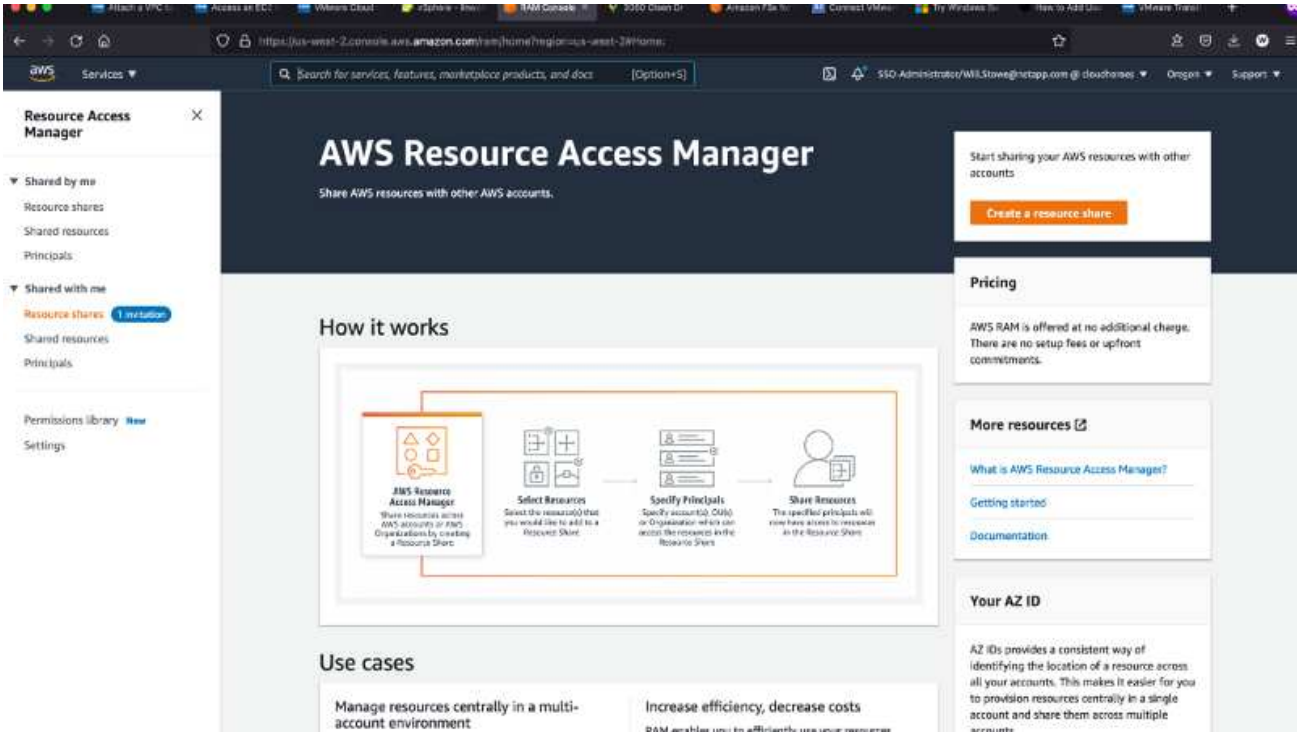


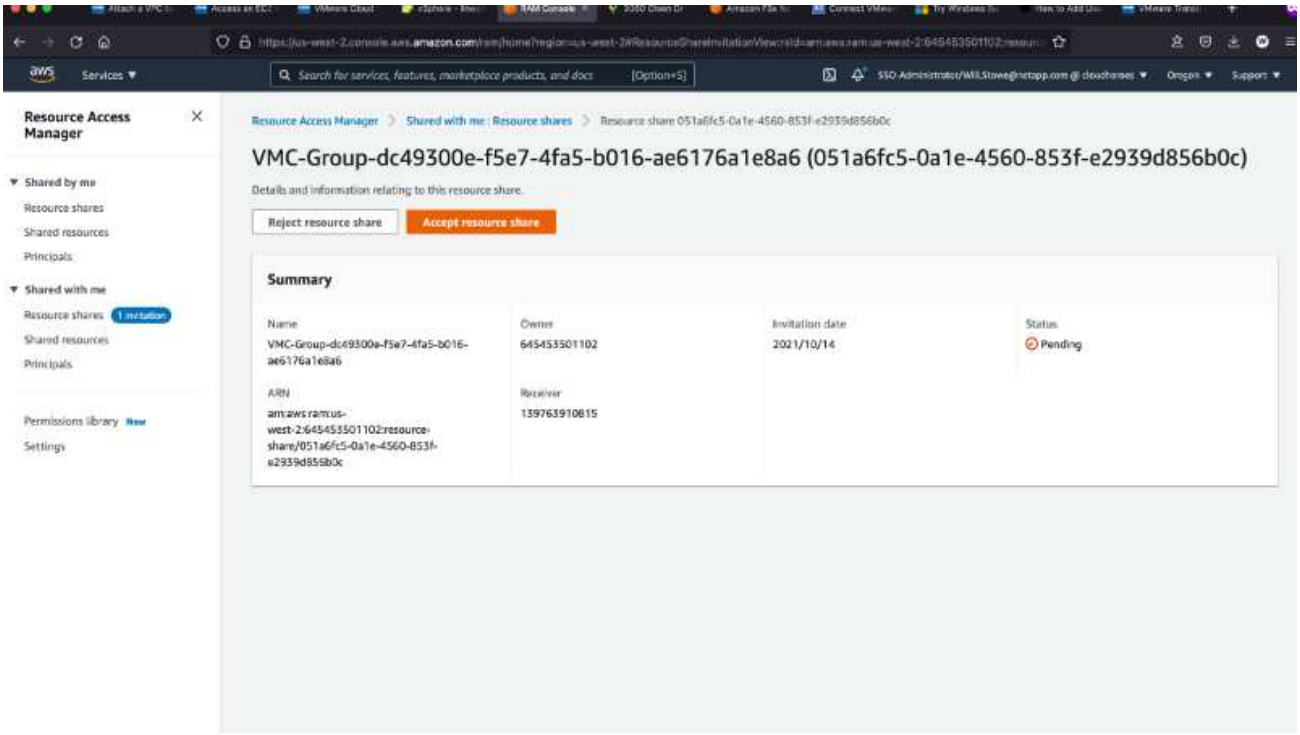
3. Attach the newly created VPC to the just created SDDC group. Select the External VPC tab and follow the instructions for attaching an External VPC to the group. This process can take 10 to 15 minutes to complete.



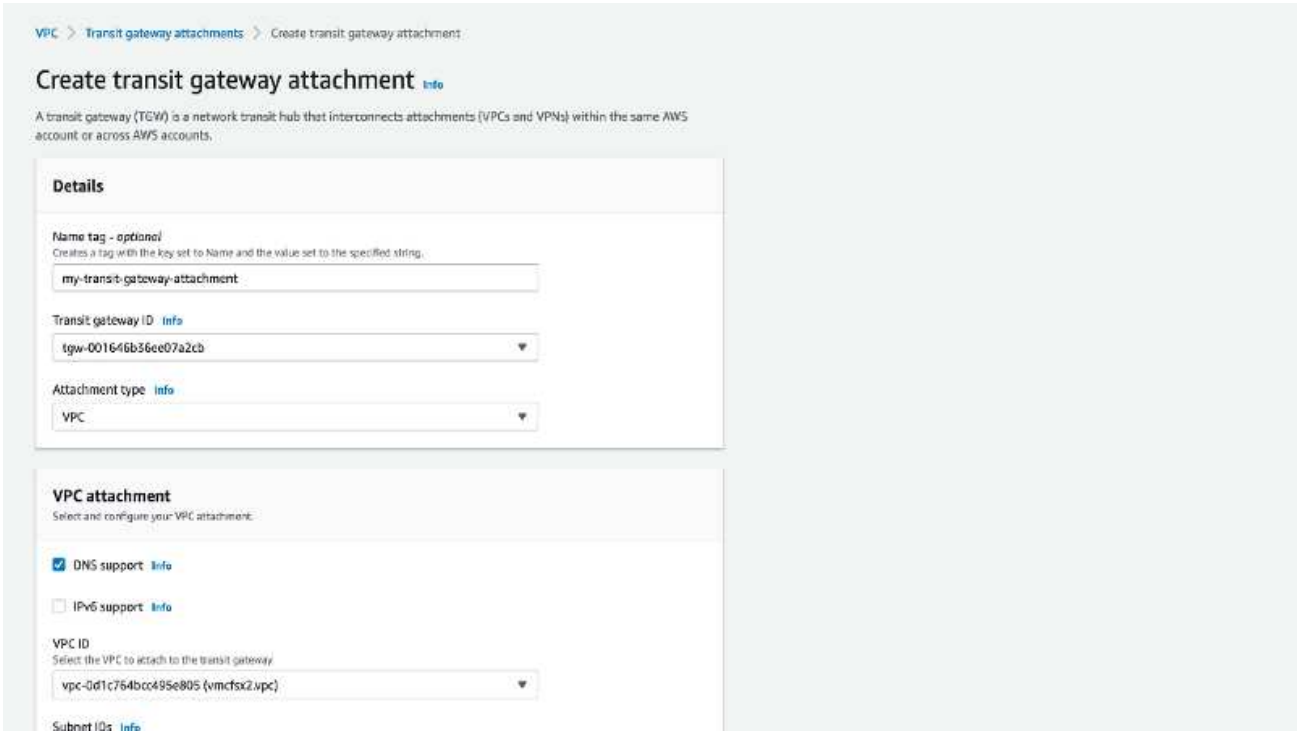


4. As part of the external VPC process, you are prompted through the AWS console to a new shared resource via the Resource Access Manager. The shared resource is the AWS Transit Gateway managed by VMware Transit Connect.

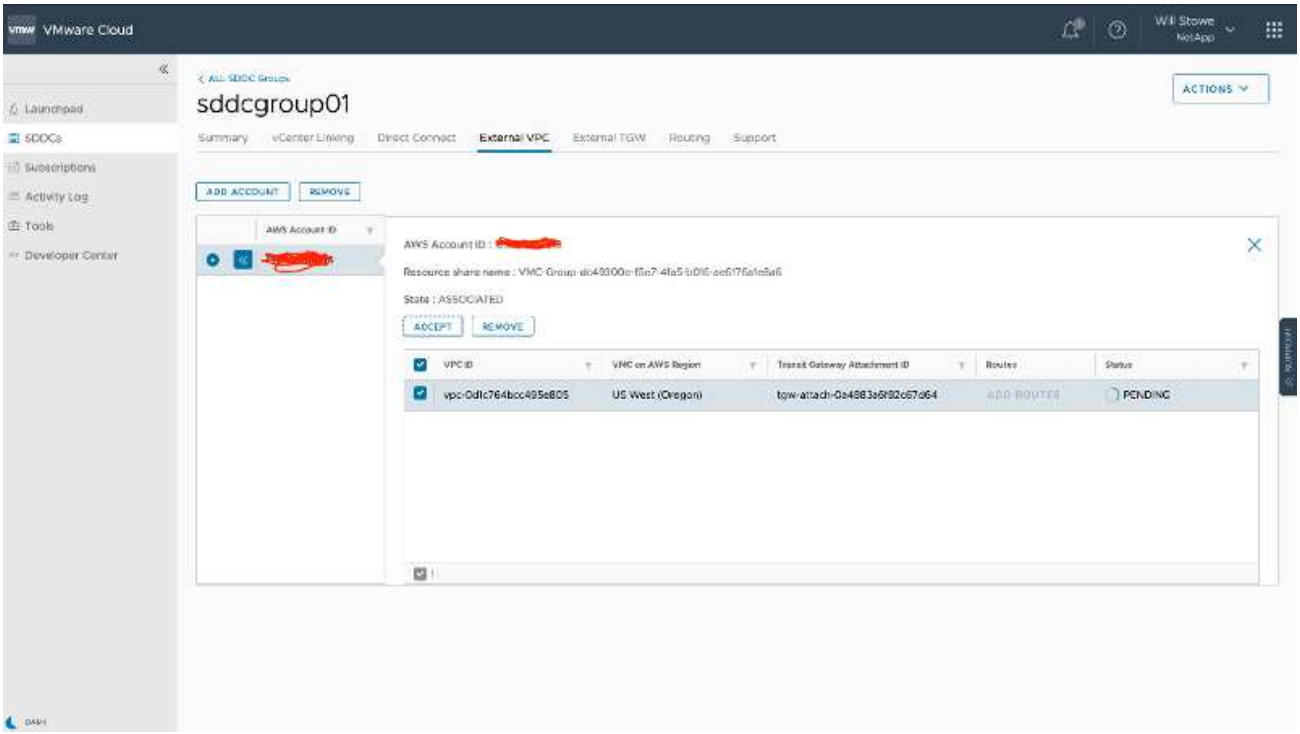




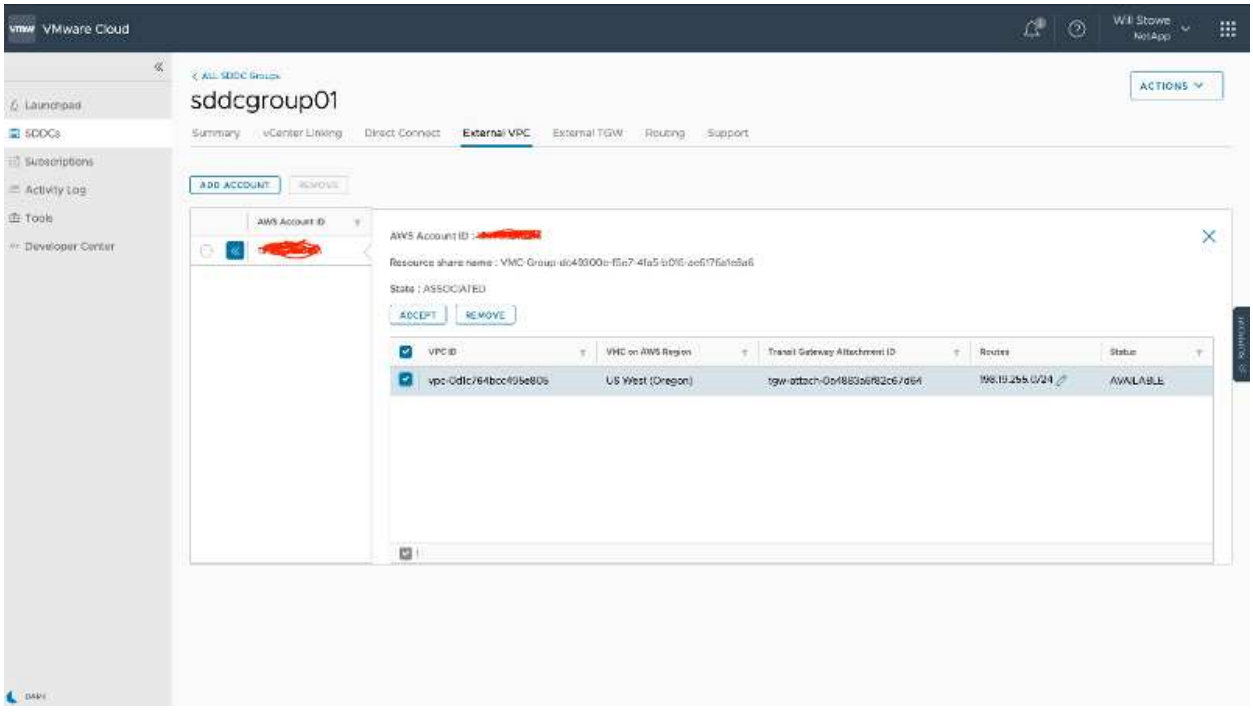
5. Create the Transit Gateway Attachment.



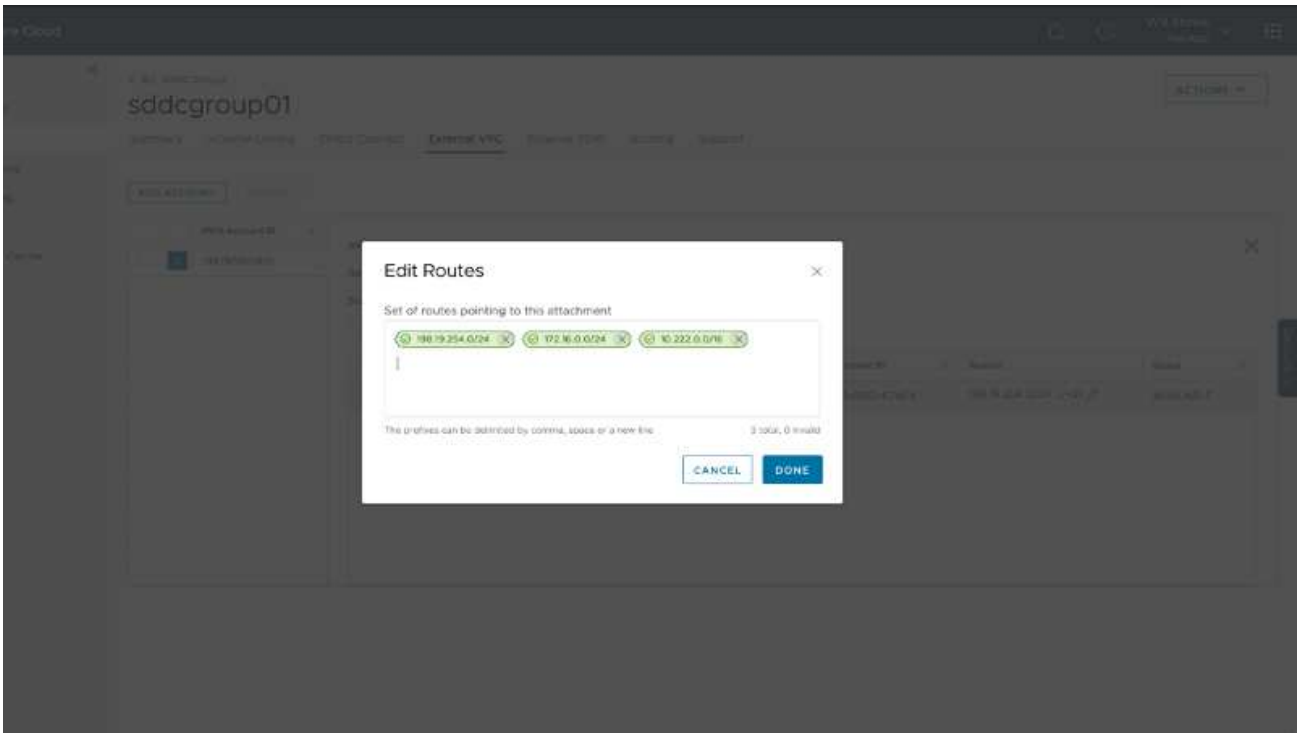
6. Back on the VMC Console, Accept the VPC attachment. This process can take approximately 10 minutes to complete.



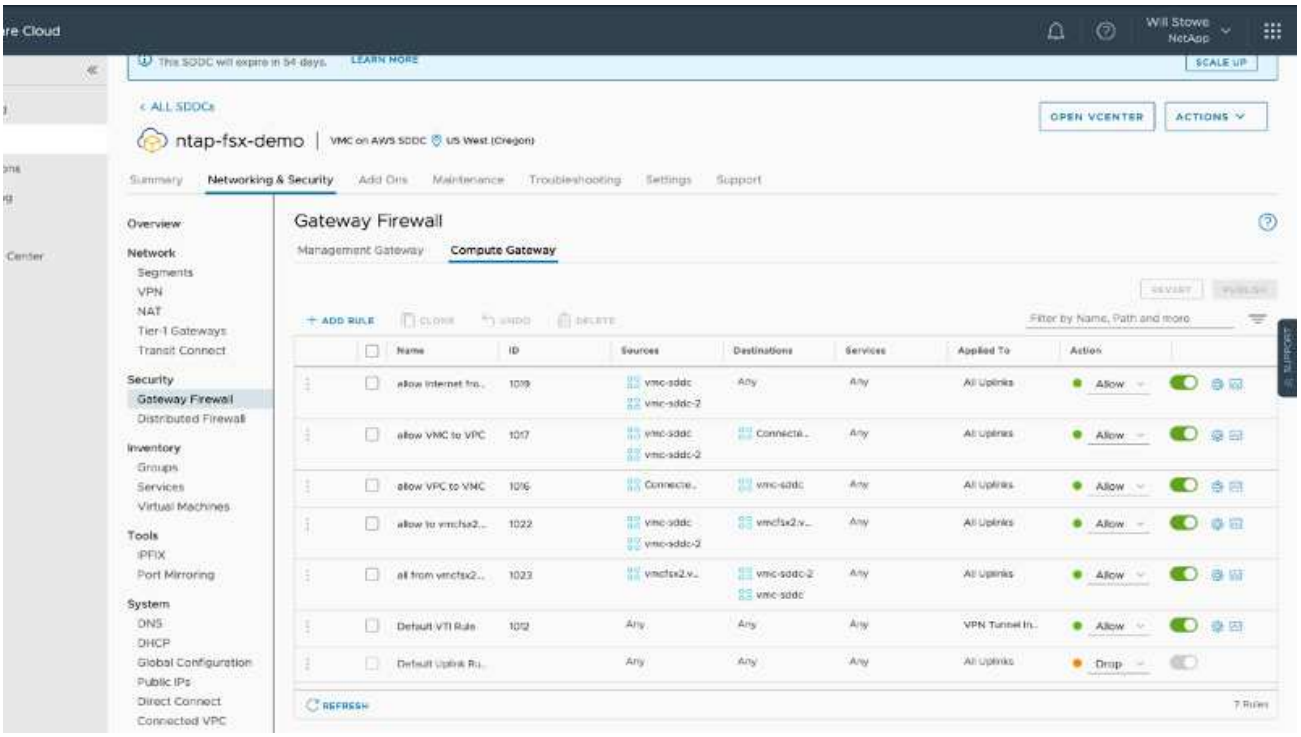
7. While in the External VPC tab, click the edit icon in the Routes column and add in the following required routes:
   - A route for the floating IP range for Amazon FSx ONTAP floating IPs.
   - A route for the floating IP range for Cloud Volumes ONTAP (if applicable).
   - A route for the newly created external VPC address space.



8. Finally, allow bidirectional traffic firewall rules for access to FSx/CVO. Follow these detailed steps for compute gateway firewall rules for SDDC workload connectivity.



9. After the firewall groups are configured for both the Management and Compute gateway, the vCenter can be accessed as follows:



The next step is to verify that Amazon FSx ONTAP or Cloud Volumes ONTAP is configured depending on your requirements and that the volumes are provisioned to offload storage components from vSAN to optimize the deployment.

Deploy and configure the Virtualization Environment on Azure

As with on-premises, planning Azure VMware Solution is critical for a successful production-ready environment for creating VMs and migration.

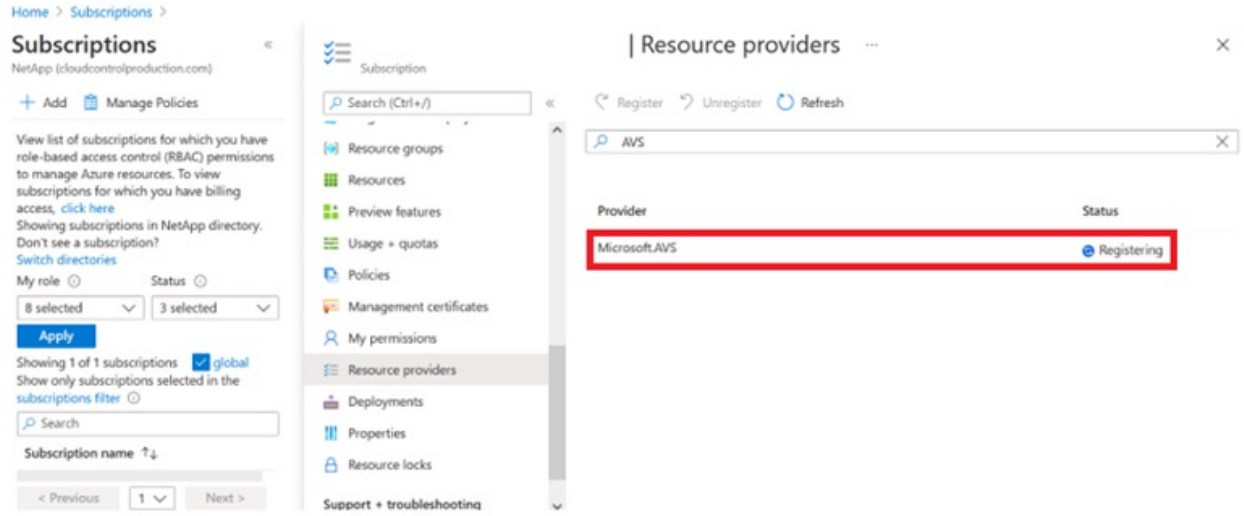
This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.

The setup process can be broken down into the following steps:

▼ **Register the resource provider and create a private cloud**

To use Azure VMware Solution, first register the resource provider within the identified subscription:

1. Sign in to the Azure portal.
2. On the Azure portal menu, select All Services.
3. In the All Services dialog box, enter the subscription and then select Subscriptions.
4. To view, select the subscription from the subscription list.
5. Select Resource Providers and enter Microsoft.AVS into the search.
6. If the resource provider is not registered, select Register.

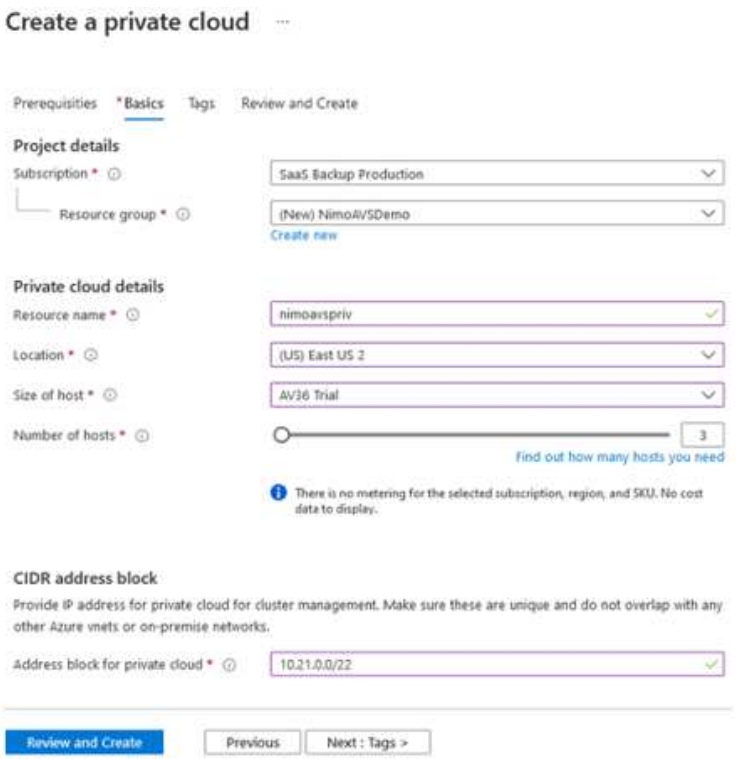




7. After the resource provider is registered, create an Azure VMware Solution private cloud by using the Azure portal.
8. Sign in to the Azure portal.
9. Select Create a New Resource.
10. In the Search the Marketplace text box, enter Azure VMware Solution and select it from the results.
11. On the Azure VMware Solution page, select Create.
12. From the Basics tab, enter the values in the fields and select Review + Create.

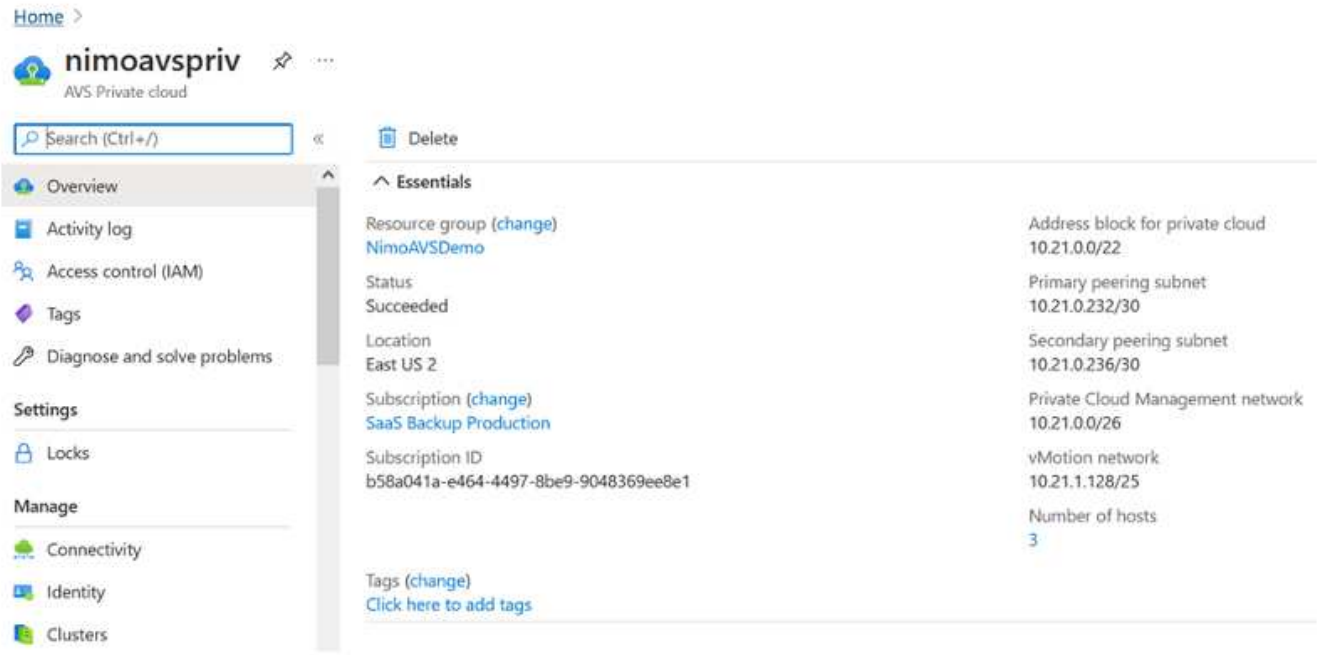
Notes:

- For a quick start, gather the required information during the planning phase.
- Select an existing resource group or create a new resource group for the private cloud. A resource group is a logical container in which the Azure resources are deployed and managed.
- Make sure the CIDR address is unique and does not overlap with other Azure Virtual Networks or on-premises networks. The CIDR represents the private cloud management network and is used for the cluster management services, such as vCenter Server and NSX-T Manager. NetApp recommends using a /22 address space. In this example, 10.21.0.0/22 is used.



The provisioning process takes approximately 4–5 hours. After the process is complete, verify that the deployment was successful by accessing the private cloud from the Azure portal. A status of Succeeded is displayed when the deployment is complete.

An Azure VMware Solution private cloud requires an Azure Virtual Network. Because Azure VMware Solution doesn't support on-premises vCenter, additional steps are required to integrate with an existing on-premises environment. Setting up an ExpressRoute circuit and a virtual network gateway is also required. While waiting for the cluster provisioning to complete, create a new virtual network or use an existing one to connect to Azure VMware Solution.



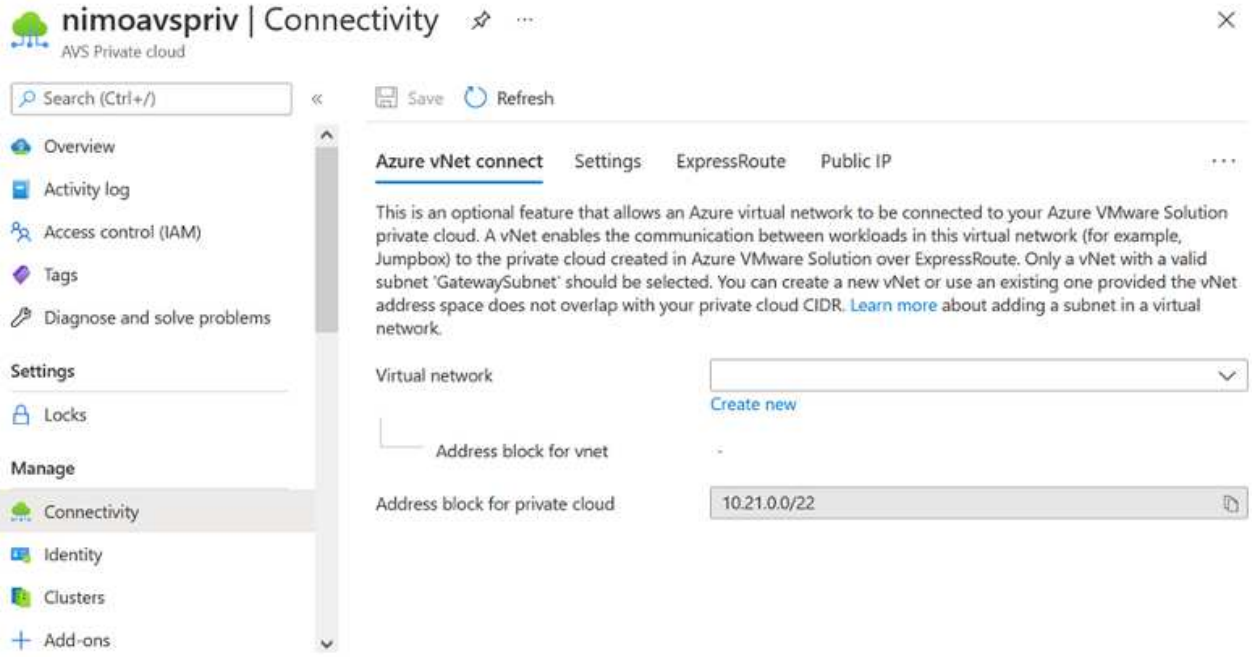
▼ **Connect to a new or existing ExpressRoute virtual network gateway**

To create a new Azure Virtual Network (VNet), select the Azure VNet Connect tab. Alternatively, you can create one manually from the Azure portal by using the Create Virtual Network wizard:

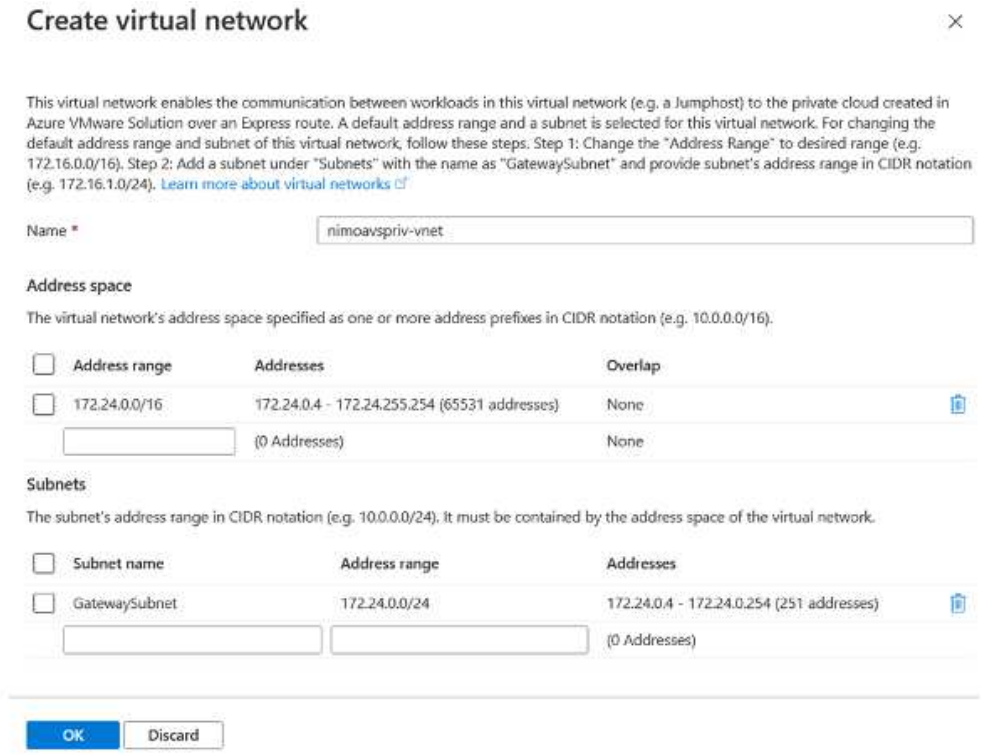
1. Go to Azure VMware Solution private cloud and access Connectivity under the Manage option.
2. Select Azure VNet Connect.
3. To create a new VNet, select the Create New option.

   This feature allows a VNet to be connected to the Azure VMware Solution private cloud. The VNet enables communication between workloads in this virtual network by automatically creating required components (for example, jump box, shared services such as Azure NetApp Files, and Cloud Volume ONTAP) to the private cloud created in Azure VMware Solution over ExpressRoute.

   Note: The VNet address space should not overlap with the private cloud CIDR.



4. Provide or update the information for the new VNet and select OK.



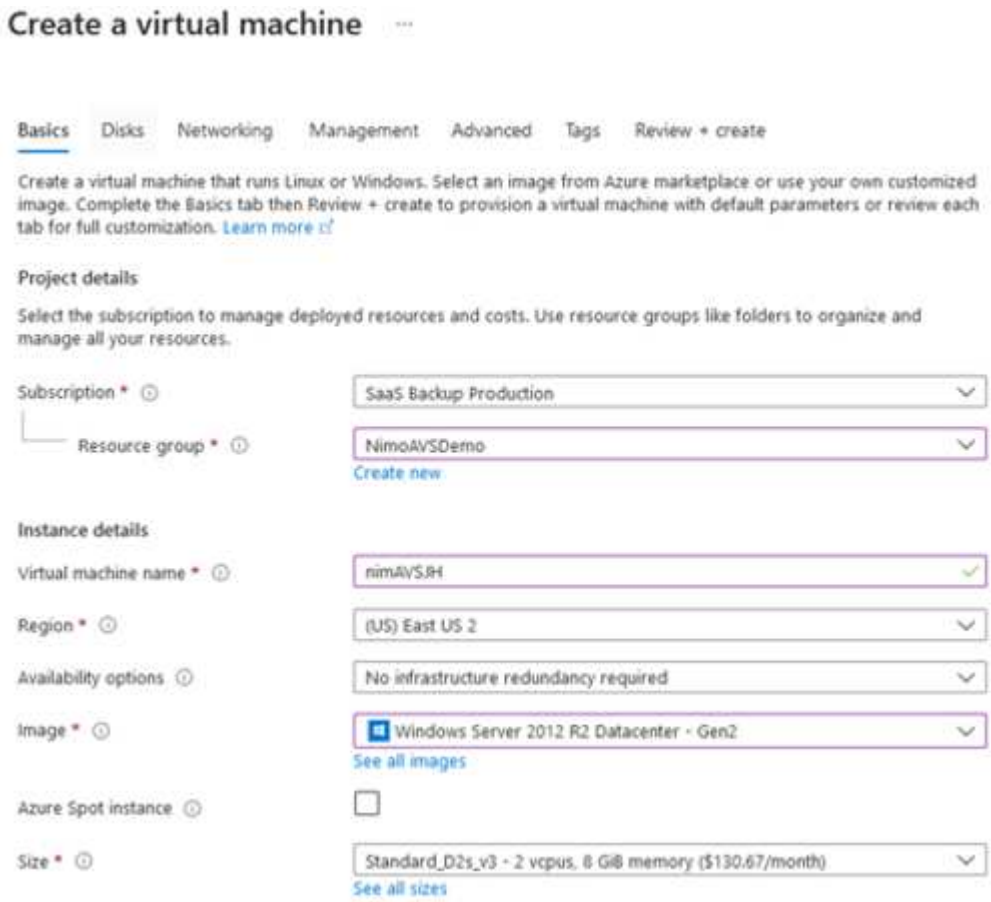
The VNet with the provided address range and gateway subnet is created in the designated subscription and resource group.

If you create a VNet manually, create a virtual network gateway with the appropriate SKU and ExpressRoute as the gateway type. After the deployment is complete, connect the ExpressRoute connection to the virtual network gateway containing Azure VMware Solution private cloud using the authorization key. For more information, see Configure networking for

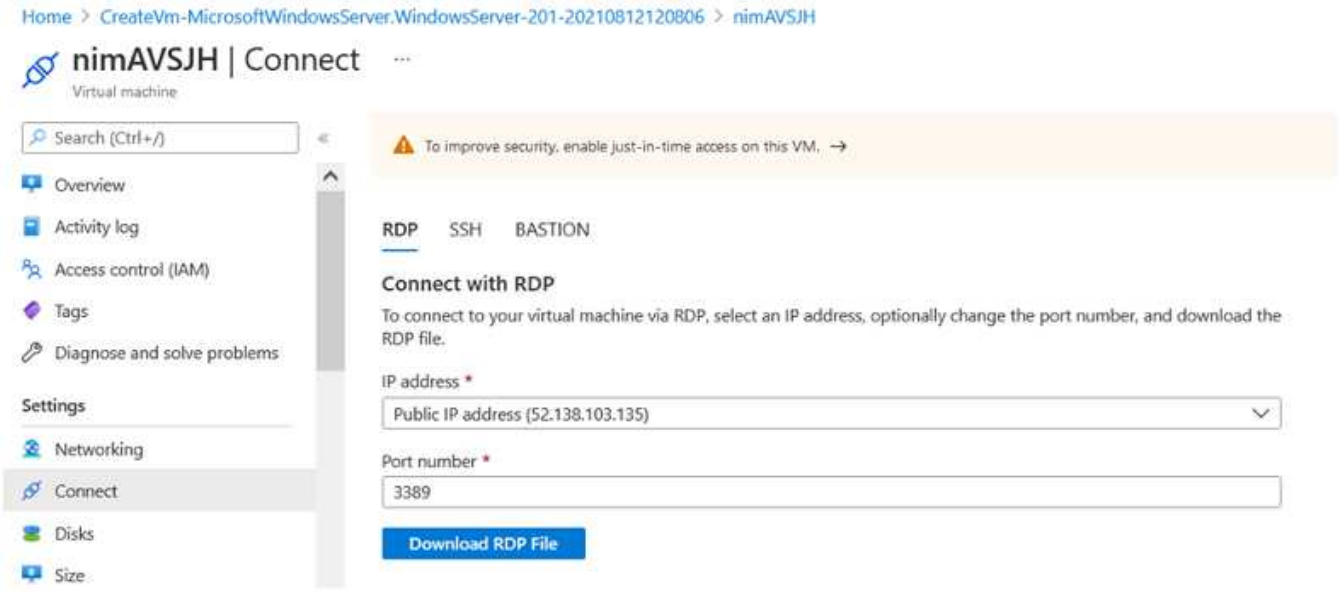
your VMware private cloud in Azure.

▼ **Validate the network connect and access to Azure VMware Solution private cloud**

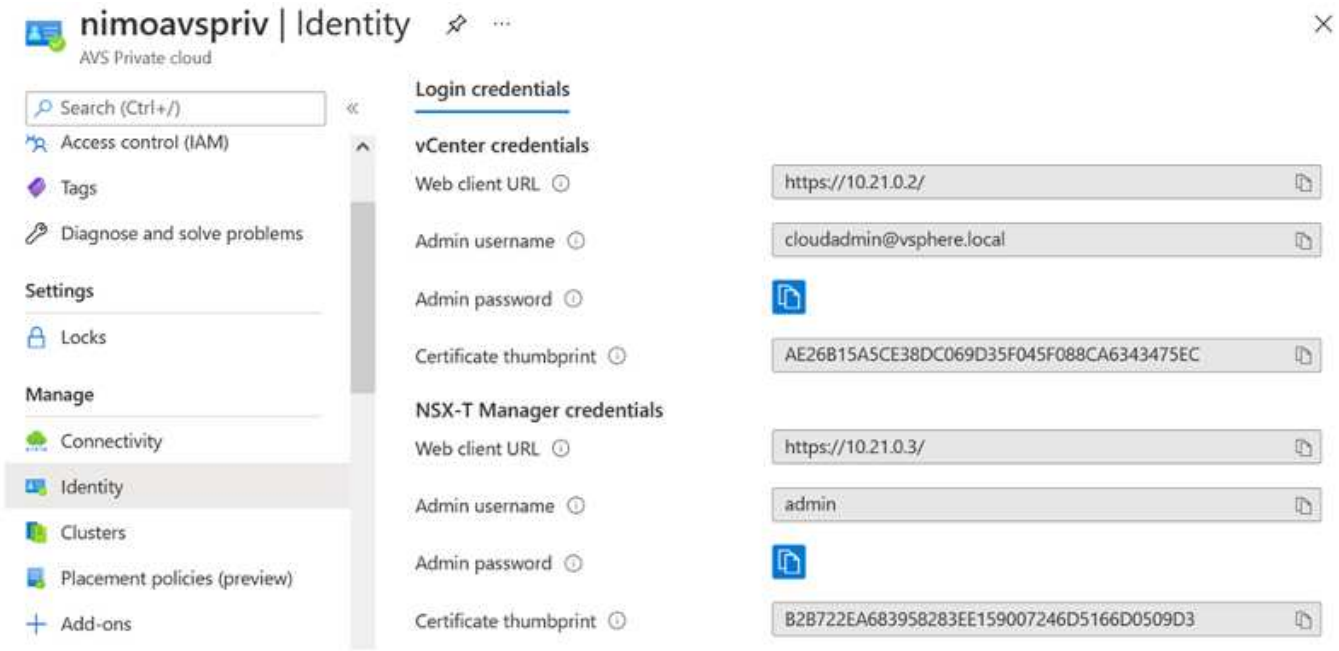
Azure VMware Solution does not allow you to manage a private cloud with on-premises VMware vCenter. Instead, jump host is required to connect to the Azure VMware Solution vCenter instance. Create a jump host in the designated resource group and sign in to the Azure VMware Solution vCenter. This jump host should be a Windows VM on the same virtual network that was created for connectivity and should provide access to both vCenter and the NSX Manager.



After the virtual machine is provisioned, use the Connect option to access RDP.



Sign in to vCenter from this newly created jump host virtual machine by using the cloud admin user . To access the credentials, go to the Azure portal and navigate to Identity (under the Manage option within the private cloud). The URLs and user credentials for the private cloud vCenter and NSX-T Manager can be copied from here.

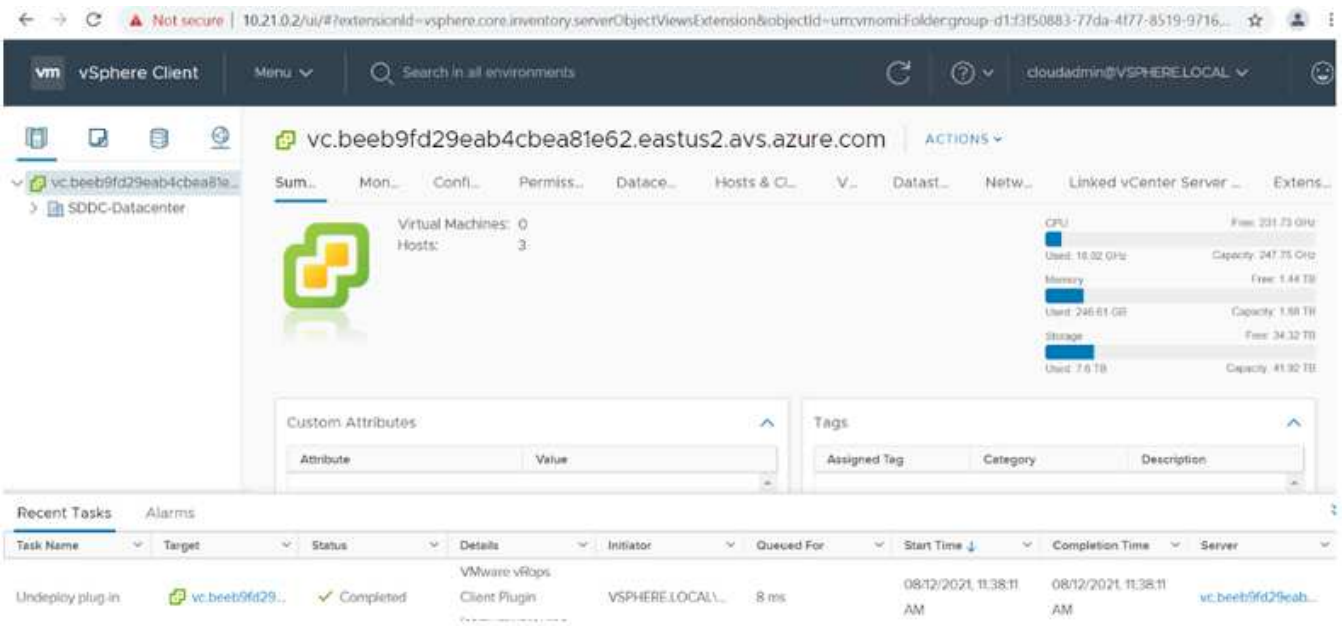


In the Windows virtual machine, open a browser and navigate to the vCenter web client URL (`"https://10.21.0.2/"`) and use the admin user name as **cloudadmin@vsphere.local** and paste the copied password. Similarly, NSX-T manager can also be accessed using the web client URL (`"https://10.21.0.3/"`) and use the admin user name and paste the copied password to create new segments or modify the existing tier gateways.



The web client URLs are different for each SDDC provisioned.





The Azure VMware Solution SDDC is now deployed and configured. Leverage ExpressRoute Global Reach to connect the on-premises environment to Azure VMware Solution private cloud. For more information, see Peer on-premises environments to Azure VMware Solution.

Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP)

As with on-premises, planning Google Cloud VMware Engine (GCVE) is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.

The setup process can be broken down into the following steps:

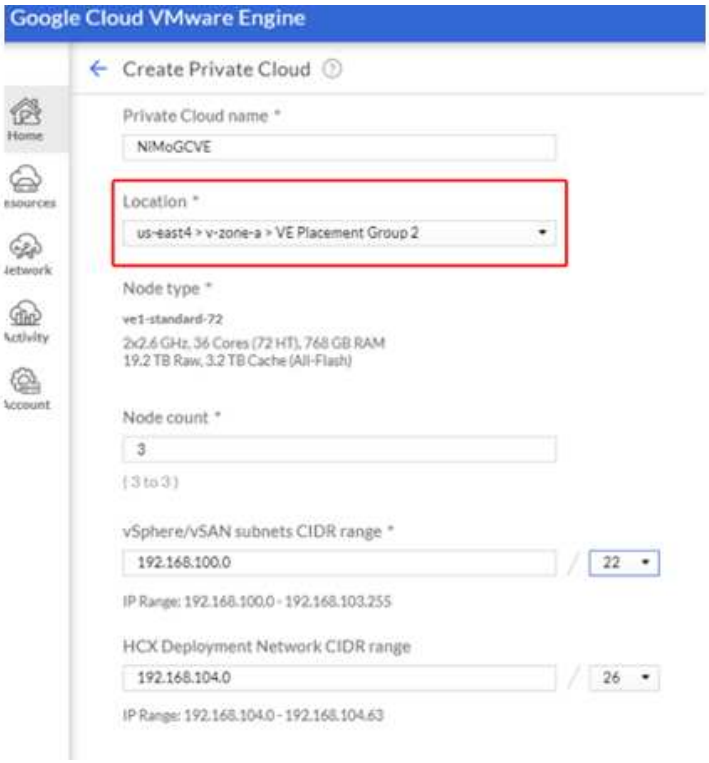
▼ Deploy and configure GCVE

To configure a GCVE environment on GCP, login to the GCP console and access the VMware Engine portal.

Click on the "New Private Cloud" button and enter the desired configuration for the GCVE Private Cloud. On "Location", make sure to deploy the private cloud in the same Region/Zone where NetApp Volumes/CVO is deployed, to ensure the best performance and lowest latency.

Pre-requisites:

- Setup VMware Engine Service Admin IAM role
- Enable VMWare Engine API access and node quota
- Make sure that the CIDR range doesn't overlap with any of your on-premises or cloud subnets. The CIDR range must be /27 or higher.



Note: Private cloud creation can take between 30 minutes to 2 hours.

▼ Enable Private Access to GCVE

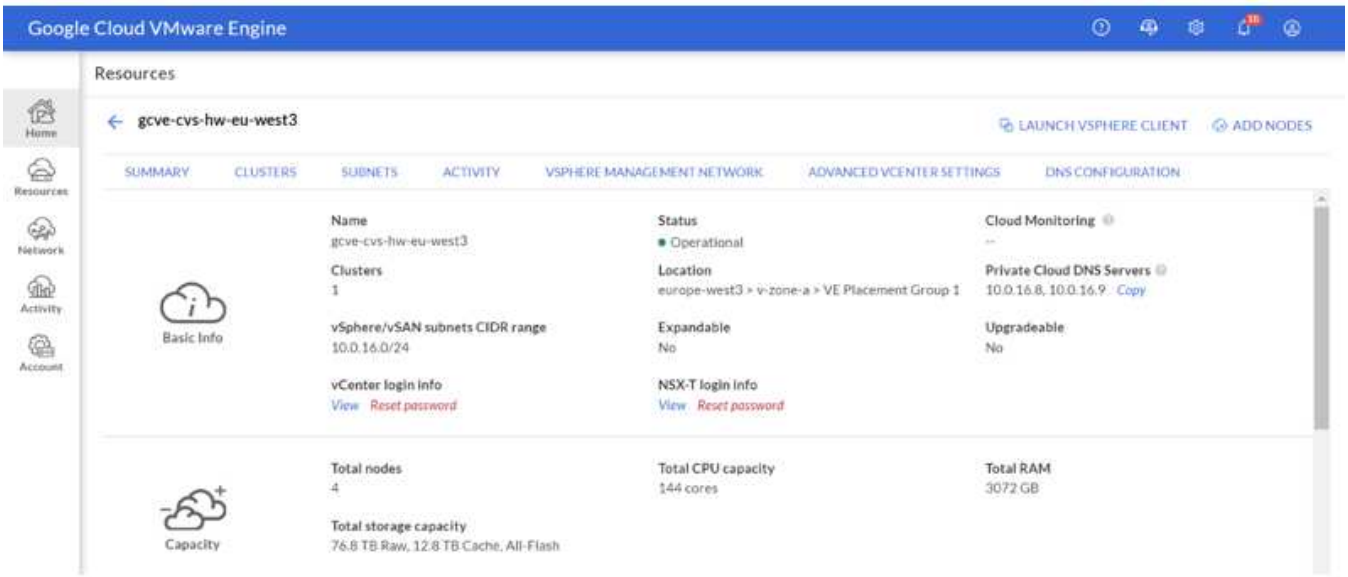
Once the Private Cloud is provisioned, configure private access to the Private Cloud for high-throughput and low-latency data-path connection.

This will ensure that the VPC network where Cloud Volumes ONTAP instances are running is able to communicate with the GCVE Private Cloud. To do so, follow the GCP documentation. For the Cloud Volume Service, establish a connection between VMware Engine and Google Cloud NetApp Volumes by performing a one-time peering between the tenant host projects. For detailed steps, follow this link.

Tenant P	Service	Region	Routing Mode	Peered Project ID	Peered VPC	VPC Peering Sta...	Region Status
ke841388caa56b...	VPC Network	europe-west3	Global	cv-performance-te...	cloud-volumes-vpc	Active	Connected
jbd729510b3ebbf...	NetApp CVS	europe-west3	Global	y2b6c17202af6dc...	netapp-tenant-vpc	Active	Connected

Sign in to vcenter using the CloudOwner@gve.local user. To access the credentials, go to the VMware

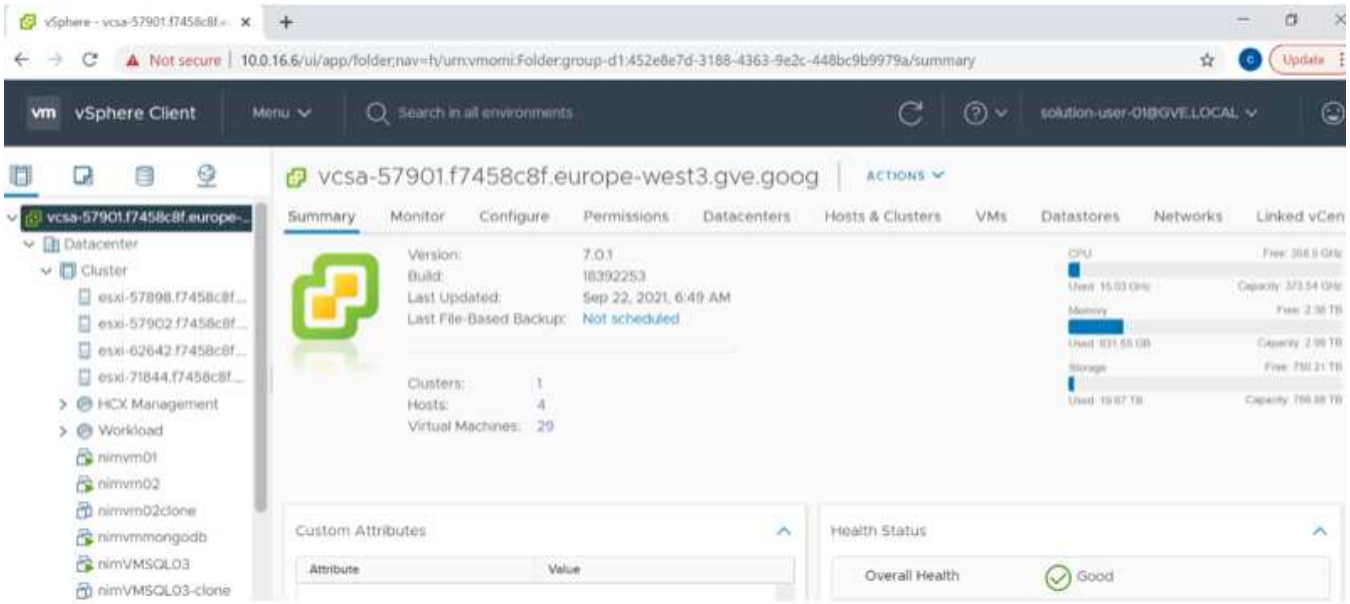
Engine portal, Go to Resources, and select the appropriate private cloud. In the Basic info section, click the View link for either vCenter login info (vCenter Server, HCX Manager) or NSX-T login info (NSX Manager).



In a Windows virtual machine, open a browser and navigate to the vCenter web client URL ("`https://10.0.16.6/`") and use the admin user name as CloudOwner@gve.local and paste the copied password. Similarly, NSX-T manager can also be accessed using the web client URL ("`https://10.0.16.11/`") and use the admin user name and paste the copied password to create new segments or modify the existing tier gateways.

For connecting from an on-premises network to VMware Engine private cloud, leverage cloud VPN or Cloud Interconnect for appropriate connectivity and make sure the required ports are open. For detailed steps, follow this link.





▼ **Deploy Google Cloud NetApp Volumes supplemental datastore to GCVE**

Refer Procedure to deploy supplemental NFS datastore with NetApp Volumes to GCVE

NetApp storage in public clouds

NetApp Storage options for Public Cloud Providers

Explore the options for NetApp as storage in the three major hyperscalers.

AWS / VMC

AWS supports NetApp storage in the following configurations:

- FSx ONTAP as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- FSx ONTAP as a supplemental NFS datastore

View the detailed guest connect storage options for VMC.
View the detailed supplemental NFS datastore options for VMC.

Azure / AVS

Azure supports NetApp storage in the following configurations:

- Azure NetApp Files (ANF) as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- Azure NetApp Files (ANF) as a supplemental NFS datastore

View the detailed guest connect storage options for AVS.
View the detailed supplemental NFS datastore options for AVS.

GCP / GCVE

Google Cloud supports NetApp storage in the following configurations:

- Cloud Volumes ONTAP (CVO) as guest connected storage
- Google Cloud NetApp Volumes (NetApp Volumes) as guest connected storage
- Google Cloud NetApp Volumes (NetApp Volumes) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.
View the detailed supplemental NFS datastore options for GCVE.

Read more about Google Cloud NetApp Volumes datastore support for Google Cloud VMware Engine (NetApp blog) or How to use Google Cloud NetApp Volumes as datastores for Google Cloud VMware Engine (Google blog)

Amazon Web Services: Options for using NetApp storage

NetApp storage can be attached to the Amazon Web Services as either guest connected or supplemental storage.

Amazon FSx for NetApp ONTAP (FSx ONTAP) as a supplemental NFS datastore

Amazon FSx ONTAP provides excellent options to deploy and manage application workloads along with file services while reducing the TCO by making data requirements seamless to the application layer. Whatever the use case, choose VMware Cloud on AWS along with Amazon FSx ONTAP for rapid realization of cloud benefits, consistent infrastructure, and operations from on-premises to AWS, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect storage.

For more information, please visit:

- FSx ONTAP as a Supplemental NFS Datastore: Overview
- Amazon FSx for ONTAP as a Supplement datastore

Amazon FSx for NetApp ONTAP as guest connected storage

Amazon FSx ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service.

FSx ONTAP provides feature-rich, fast, and flexible shared file storage that's broadly accessible from Linux, Windows, and macOS compute instances running in AWS or on premises. FSx ONTAP offers high-performance solid state drive (SSD) storage with sub-millisecond latencies. With FSx ONTAP, you can achieve SSD levels of performance for your workload while paying for SSD storage for only a small fraction of your data.

Managing your data with FSx ONTAP is easier because you can snapshot, clone, and replicate your files with the click of a button. In addition, FSx ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

FSx ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx ONTAP supports popular data security and antivirus applications.

For more information, please visit FSx ONTAP as Guest Connected Storage

Cloud Volumes ONTAP (CVO) as guest connected storage

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud.

For more information, please visit CVO as Guest Connected Storage.

TR-4938: Mount Amazon FSx ONTAP as a NFS datastore with VMware Cloud on AWS

This document outlines how to mount Amazon FSx ONTAP as a NFS datastore with VMware Cloud on AWS.

Introduction

Every successful organization is on a path of transformation and modernization. As part of this process, companies typically use their existing VMware investments to leverage cloud benefits and exploring how to migrate, burst, extend, and provide disaster recovery for processes as seamlessly as possible. Customers migrating to the cloud must evaluate the use cases for elasticity and burst, data-center exit, data-center consolidation, end-of-life scenarios, mergers, acquisitions, and so on.

Although VMware Cloud on AWS is the preferred option for the majority of the customers because it delivers unique hybrid capabilities to a customer, limited native storage options have restricted its usefulness for organizations with storage-heavy workloads. Because storage is directly tied to hosts, the only way to scale storage is to add more hosts, which can increase costs by 35-40% or more for storage intensive workloads. These workloads need additional storage and segregated performance, not additional horsepower, but that means paying for additional hosts. This is where the recent integration of FSx ONTAP comes in handy for storage and performance intensive workloads with VMware Cloud on AWS.

Let's consider the following scenario: a customer requires eight hosts for horsepower (vCPU/vMem), but they also have a substantial requirement for storage. Based on their assessment, they require 16 hosts to meet storage requirements. This increases the overall TCO because they must buy all that additional horsepower when all they really need is more storage. This is applicable for any use case, including migration, disaster recovery, bursting, dev/test, and so on.

This document walks you through the steps necessary to provision and attach FSx ONTAP as a NFS datastore for VMware Cloud on AWS.

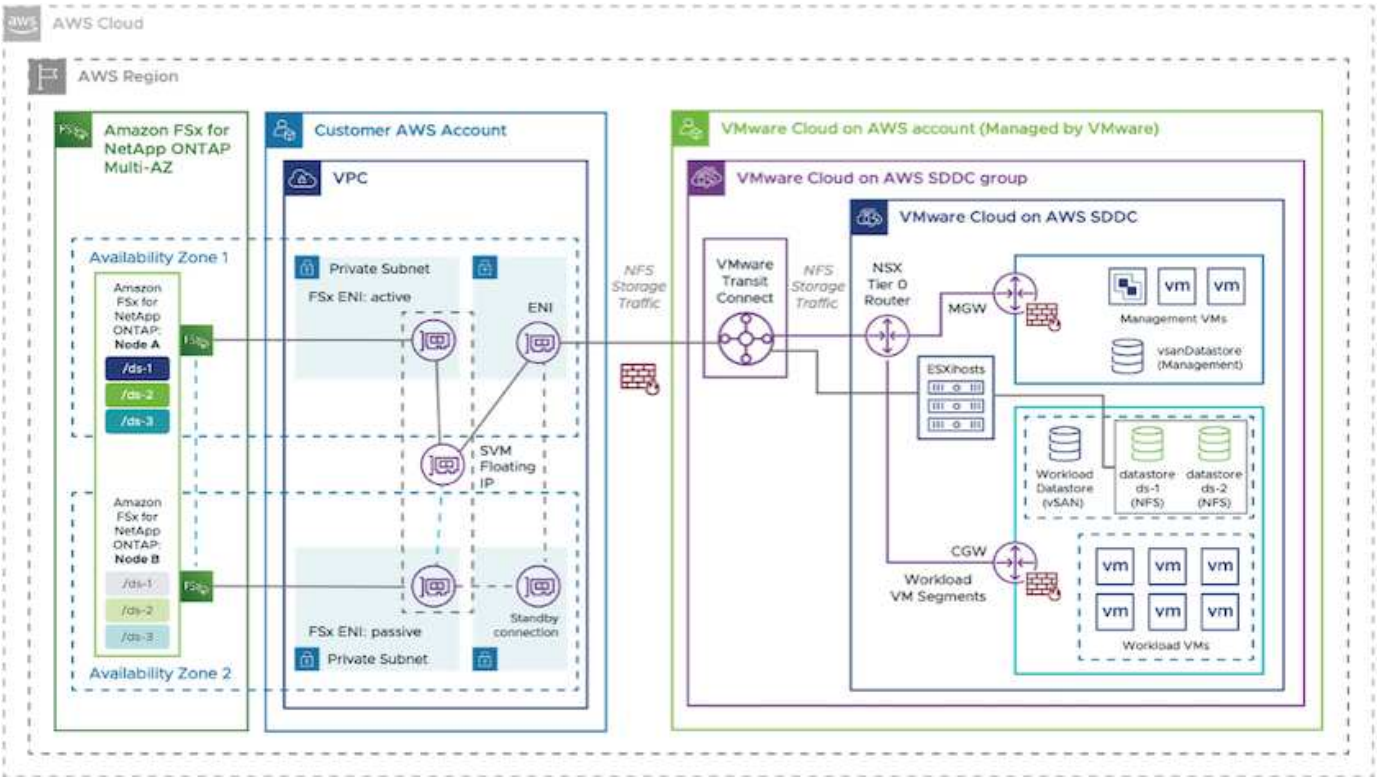


This solution is also available from VMware. Please visit the VMware Cloud on AWS Documentation for more information.

Connectivity options

VMware Cloud on AWS supports both multi-AZ and single-AZ deployments of FSx ONTAP.

This section describes the high-level connectivity architecture along with the steps needed to implement the solution to expand the storage in a SDDC cluster without the need for adding additional hosts.



The high-level deployment steps are as follows:

1. Create Amazon FSx ONTAP in a new designated VPC.
2. Create an SDDC group.
3. Create VMware Transit Connect and a TGW attachment.
4. Configure routing (AWS VPC and SDDC) and security groups.
5. Attach an NFS volume as a datastore to the SDDC cluster.

Before you provision and attach FSx ONTAP as a NFS datastore, you must first set up a VMware on Cloud SDDC environment or get an existing SDDC upgraded to v1.20 or above. For more information, see the Getting Started With VMware Cloud on AWS.

FSx ONTAP is not currently supported with stretched clusters.

Conclusion

This document covers the steps necessary to configure Amazon FSx ONTAP with VMware cloud on AWS. Amazon FSx ONTAP provides excellent options to deploy and manage application workloads along with file services while reducing the TCO by making data requirements seamless to the application layer. Whatever the use case, choose VMware Cloud on AWS along with Amazon FSx ONTAP for rapid realization of cloud benefits, consistent infrastructure, and operations from on-premises to AWS, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures

used to connect storage. Remember, it is just the position of the data that changed along with new names; the tools and processes all remain the same, and Amazon FSx ONTAP helps to optimize the overall deployment.

To learn more about this process, feel free to follow the detailed walkthrough video.

Amazon FSx ONTAP VMware Cloud

Creating a Supplemental NFS Datastore in AWS

After VMware Cloud is ready and connected to AWS VPC, you must deploy Amazon FSx ONTAP into a newly designated VPC rather than the original connected or existing default VPC.

To start, deploy an additional VPC in the same region and availability zone where SDDC resides, and then deploy Amazon FSx ONTAP into the new VPC. Configuration of an SDDC group in the VMware Cloud console enables the networking configuration options required to connect to the newly designated VPC where FSx ONTAP will be deployed.

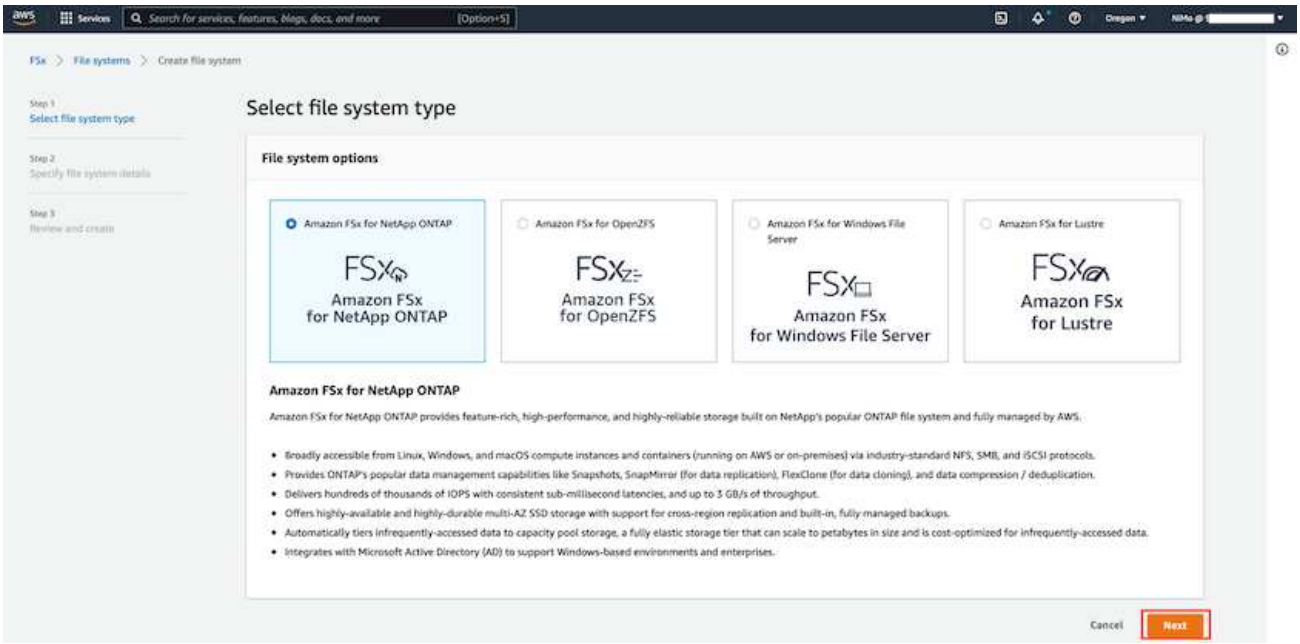
ⓘ Deploy FSx ONTAP in the same Availability Zone as VMware Cloud on AWS SDDC.

ⓘ You cannot deploy FSx ONTAP in the Connected VPC. Instead, you must deploy it in a new, designated VPC and then connect the VPC to a VMware Managed Transit Gateway (vTGW) via SDDC groups.

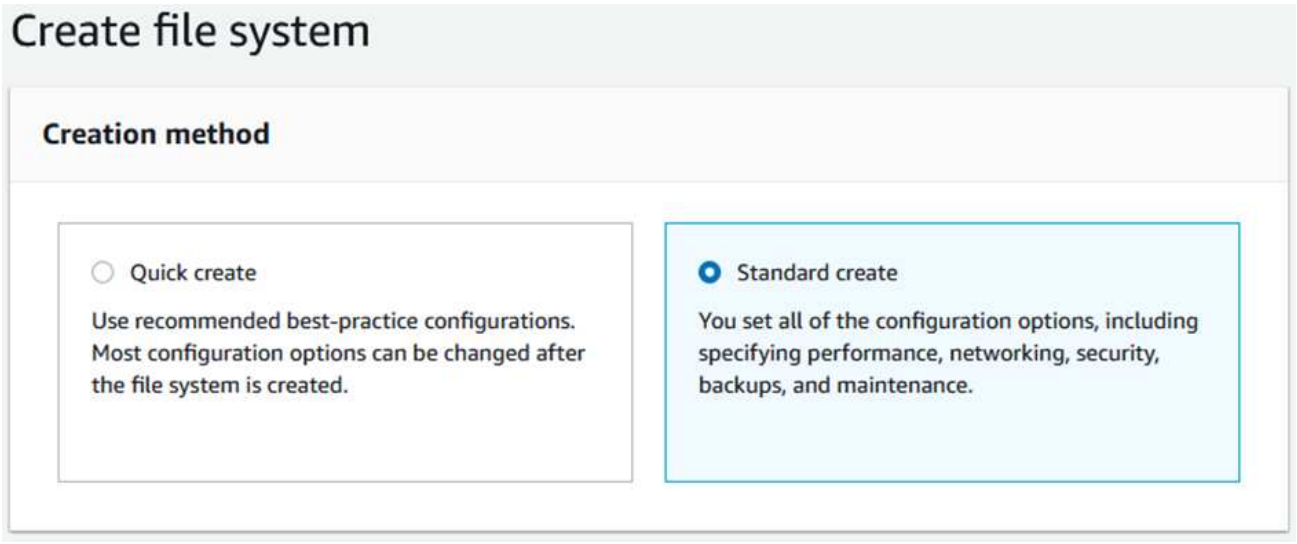
▼ **Step 1: Create Amazon FSx ONTAP in a new, designated VPC**

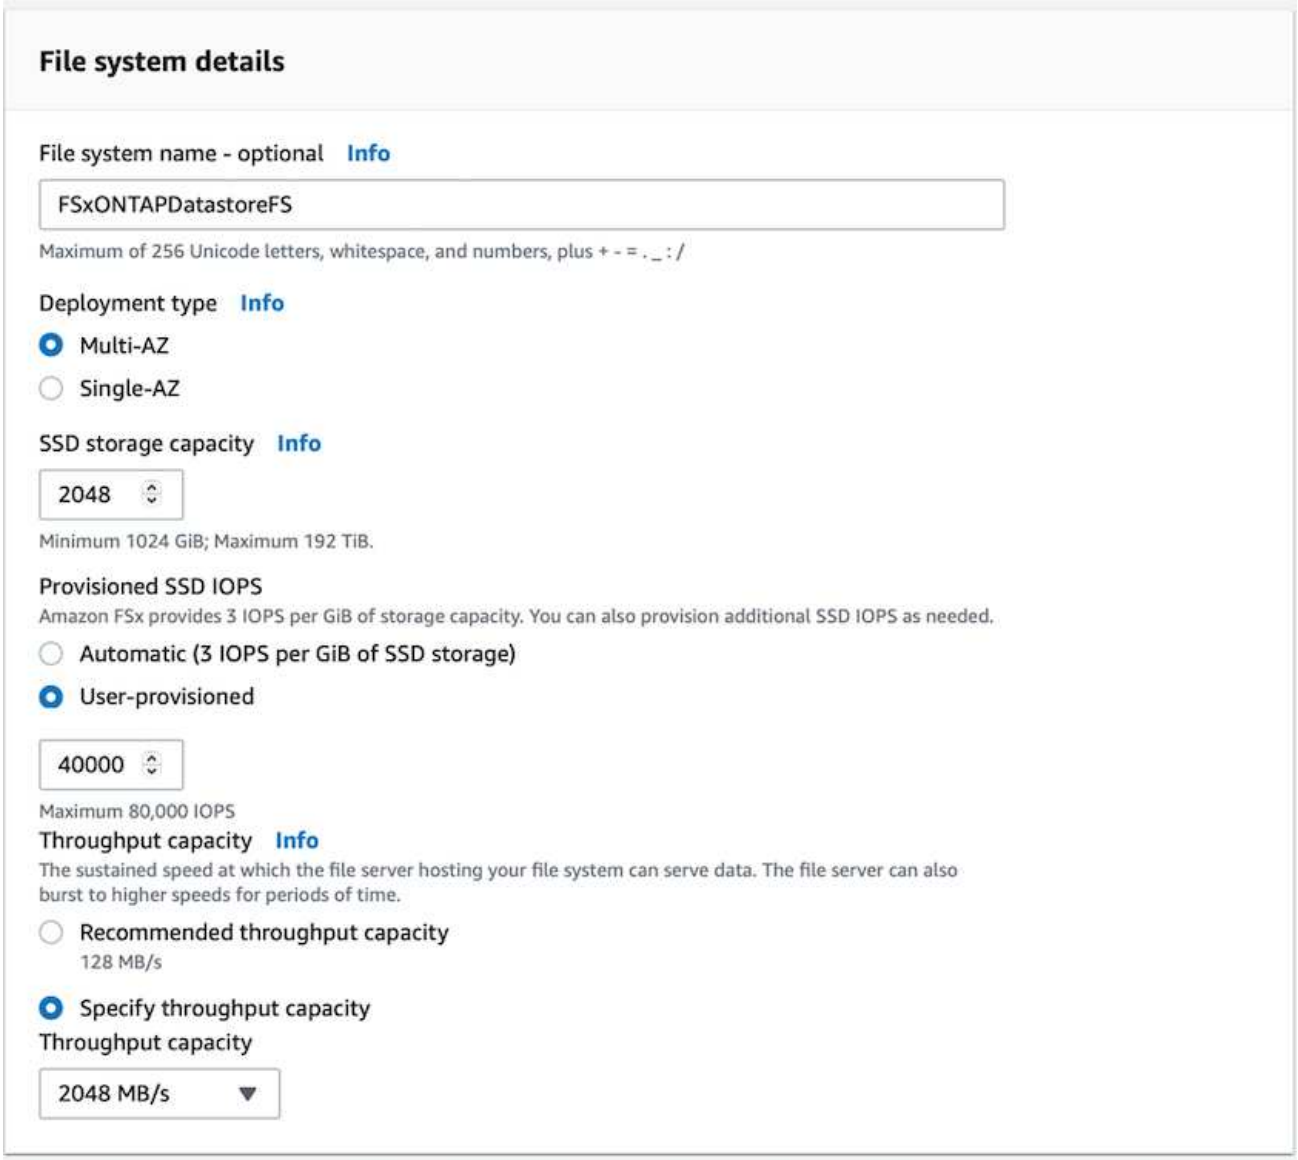
To create and mount the Amazon FSx ONTAP file system, complete the following steps:

1. Open the Amazon FSx console at `https://console.aws.amazon.com/fsx/` and choose **Create file system** to start the **File System Creation** wizard.
2. On the Select File System Type page, select **Amazon FSx ONTAP** and then click **Next**. The **Create File System** page appears.



3. For the creation method, choose **Standard Create**.





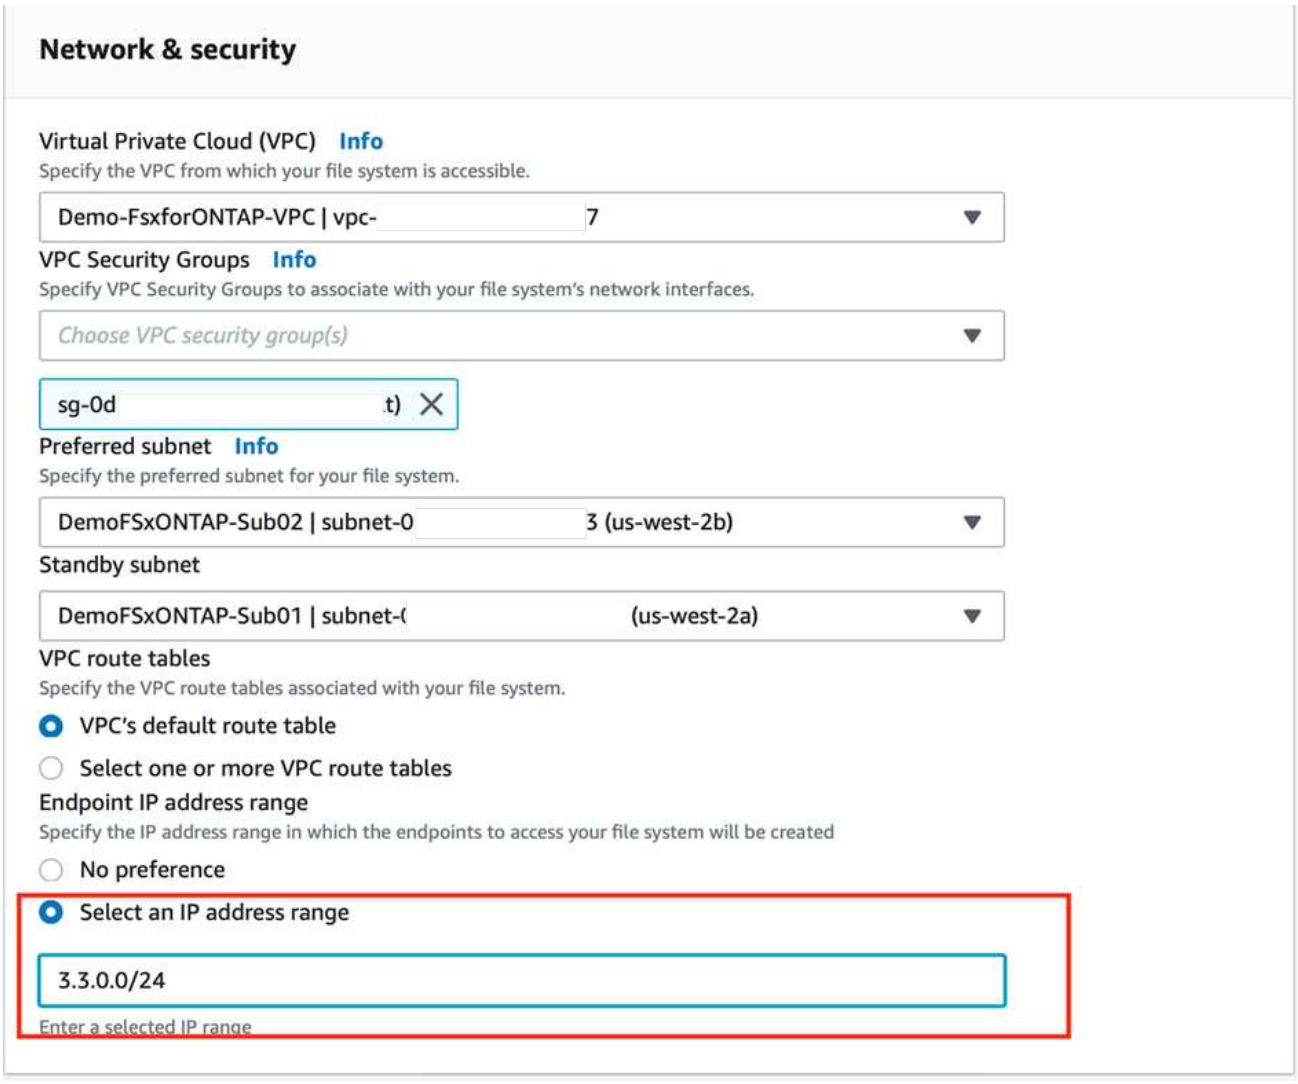
The datastore sizes vary quite a bit from customer to customer. Although the

recommended number of virtual machines per NFS datastore is subjective, many factors determine the optimum number of VMs that can be placed on each datastore. Although most administrators only consider capacity, the amount of concurrent I/O being sent to the VMDKs is one of the most important factors for overall performance. Use performance statistics from on-premises to size the datastore volumes accordingly.

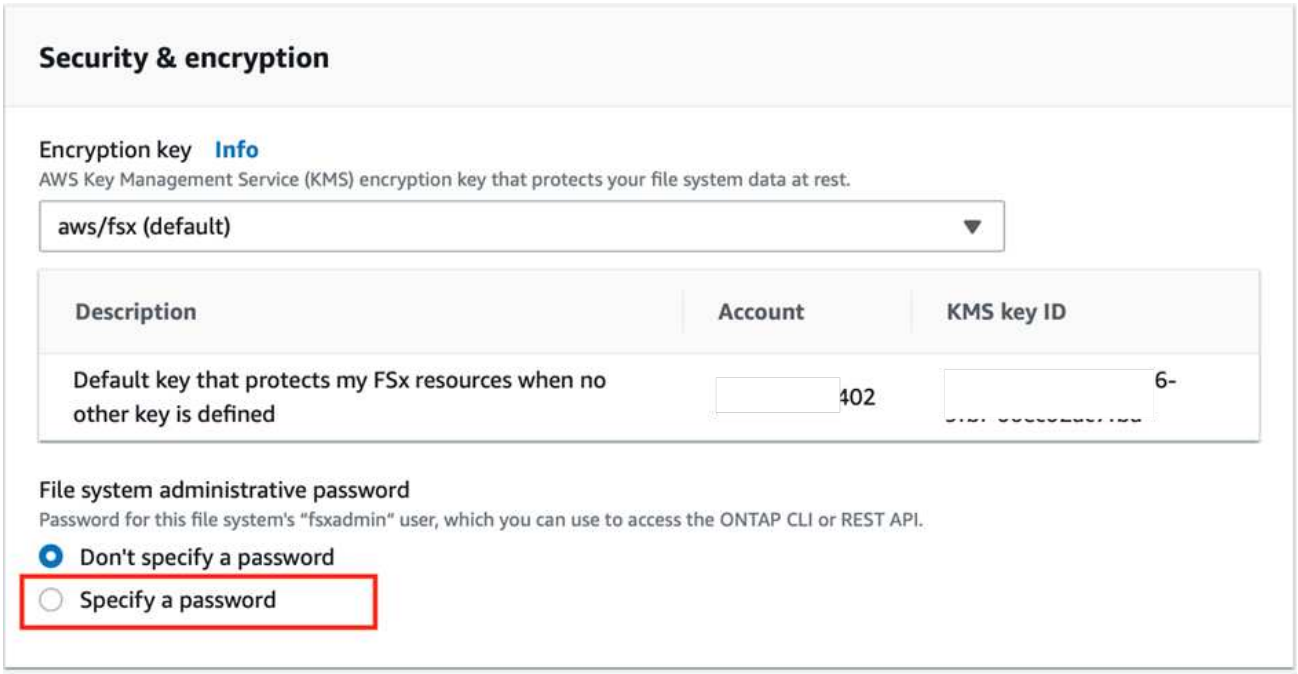
4. In the **Networking** section for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, Demo- FSxforONTAP-VPC is selected from the dropdown menu.

Make sure this is a new, designated VPC and not the connected VPC.

By default, FSx ONTAP uses 198.19.0.0/16 as the default endpoint IP address range for the file system. Make sure that the Endpoint IP address range does not conflict with the VMC on the AWS SDDC, associated VPC subnets and on-premises infrastructure. If you are unsure, use a non-overlapping range with no conflicts.

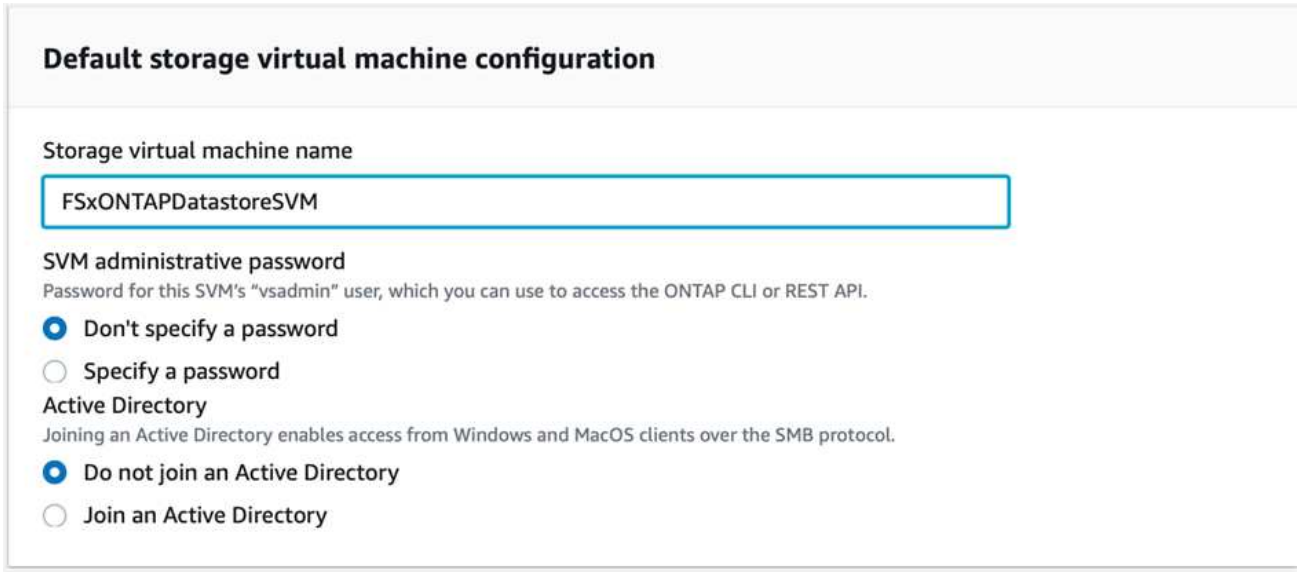


5. In the **Security & Encryption** section for the encryption key, choose the AWS Key Management Service (AWS KMS) encryption key that protects the file system's data at rest. For the **File System Administrative Password**, enter a secure password for the fsxadmin user.



6. In the **Default Storage Virtual Machine Configuration** section, specify the name of the SVM.

   As of GA, four NFS datastores are supported.



7. In the **Default Volume Configuration** section, specify the volume name and size required for datastore and click **Next**. This should be an NFSv3 volume. For **Storage Efficiency**, choose **Enabled** to turn on the ONTAP storage efficiency features (compression, deduplication, and compaction). After creation, use the shell to modify the volume parameters using ***volume modify*** as follows:

Setting	Configuration
Volume guarantee (Space Guarantee Style)	None (thin provisioned) – set by default
fractional_reserve (fractional-reserve)	0% – set by default
snap_reserve (percent-snapshot-space)	0%

Setting	Configuration
Autosize (autosize-mode)	grow_shrink
Storage efficiency	Enabled – set by default
Autodelete	volume / oldest_first
Volume Tiering Policy	Snapshot only – set by default
try_first	Autogrow
Snapshot policy	None

Use the following SSH command to create and modify volumes:

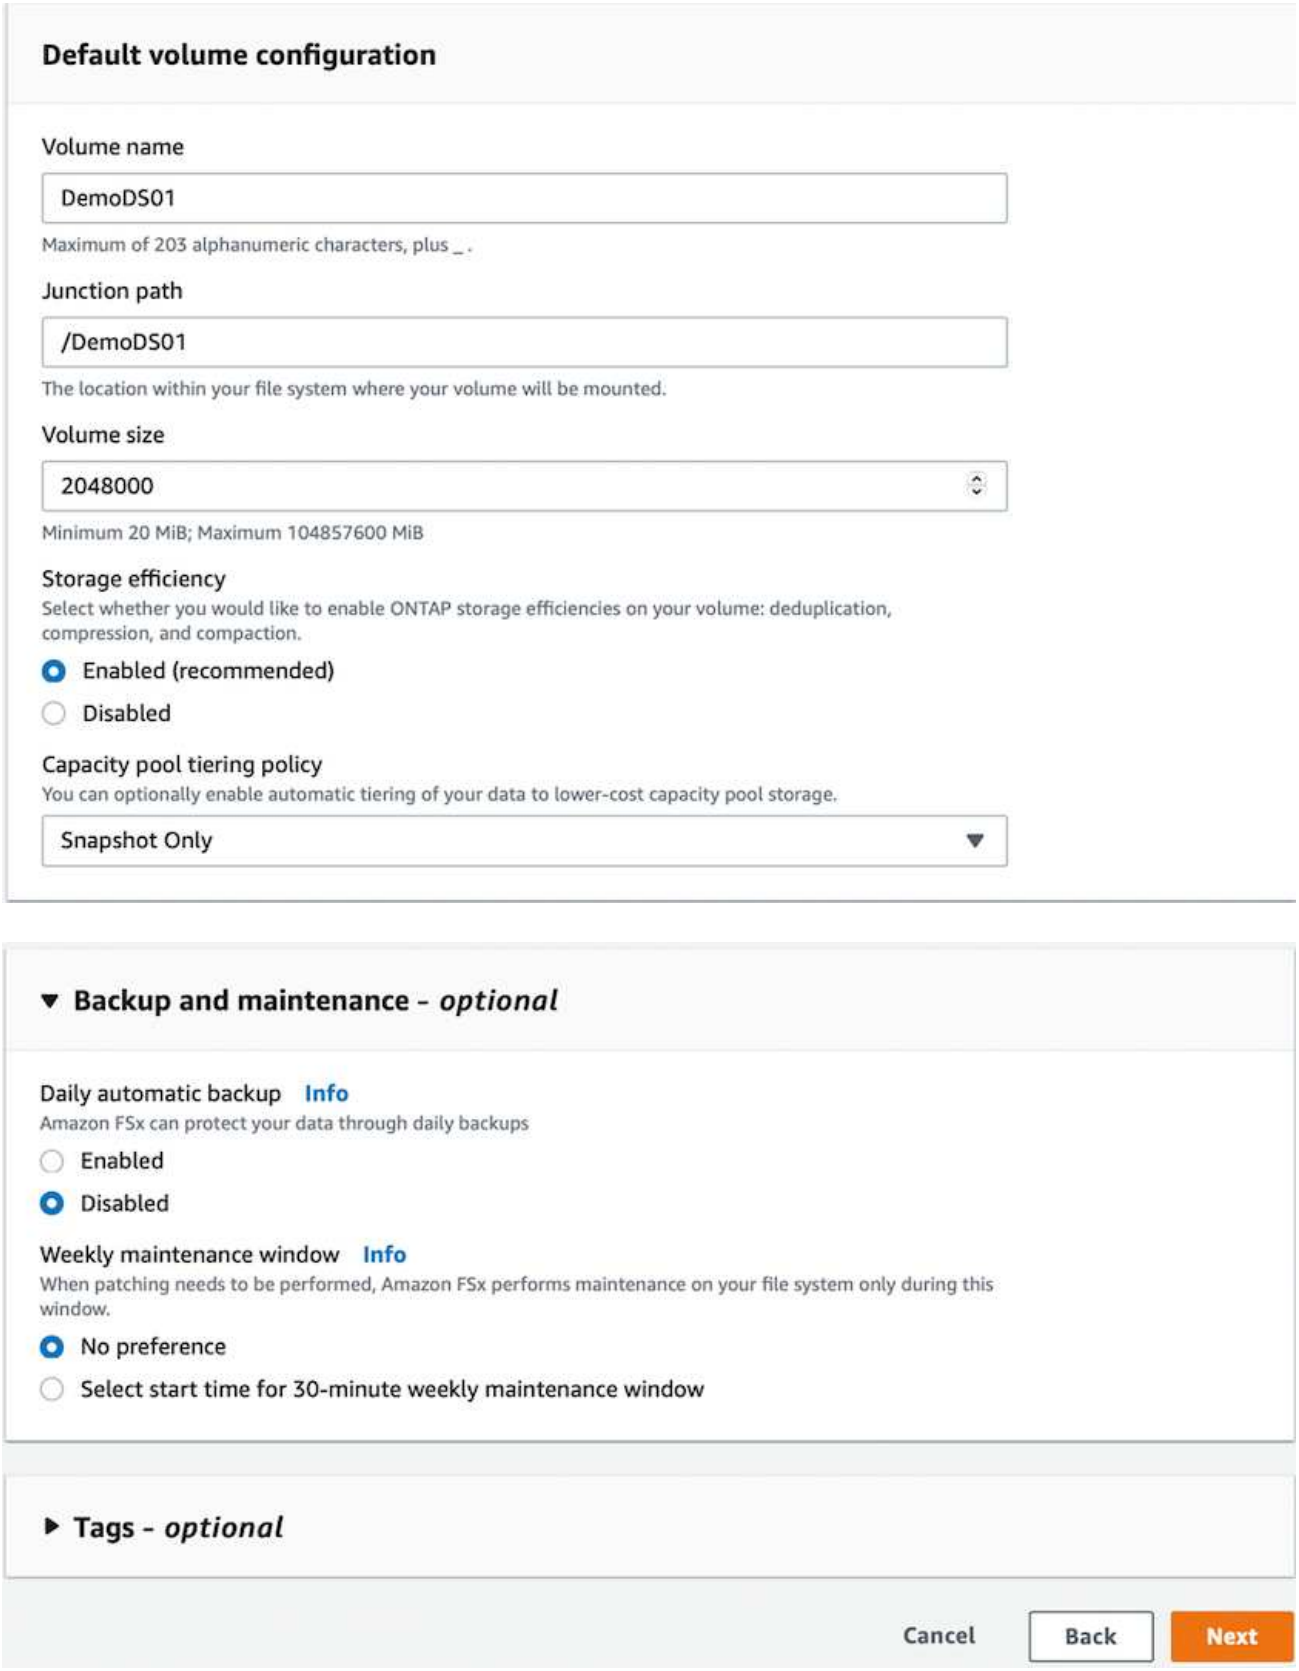
Command to create new datastore volume from shell:

```
volume create -vserver FSxONTAPDatastoreSVM -volume DemoDS002
-aggregate aggr1 -size 1024GB -state online -tiering-policy snapshot-
only -percent-snapshot-space 0 -autosize-mode grow -snapshot-policy
none -junction-path /DemoDS002
```

Note: The volumes created via shell will take few minutes to show up in the AWS Console.

Command to modify volume parameters which are not set by default:

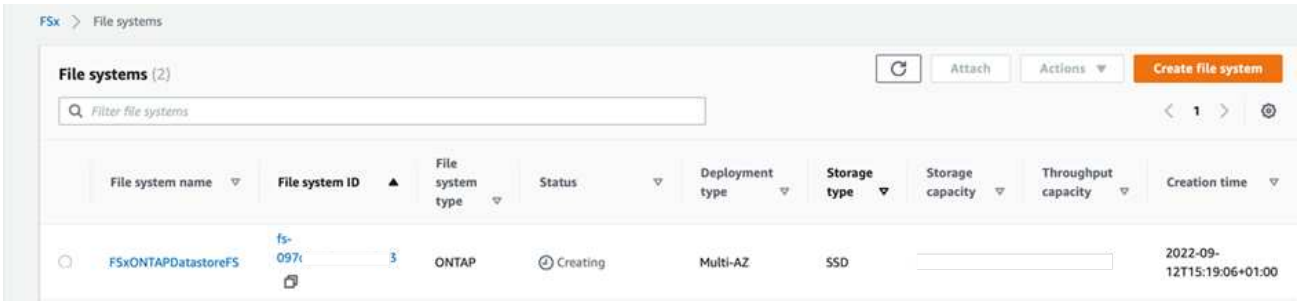
```
volume modify -vserver FSxONTAPDatastoreSVM -volume DemoDS002
-fractional-reserve 0
volume modify -vserver FSxONTAPDatastoreSVM -volume DemoDS002 -space
-mgmt-try-first vol_grow
volume modify -vserver FSxONTAPDatastoreSVM -volume DemoDS002
-autosize-mode grow
```

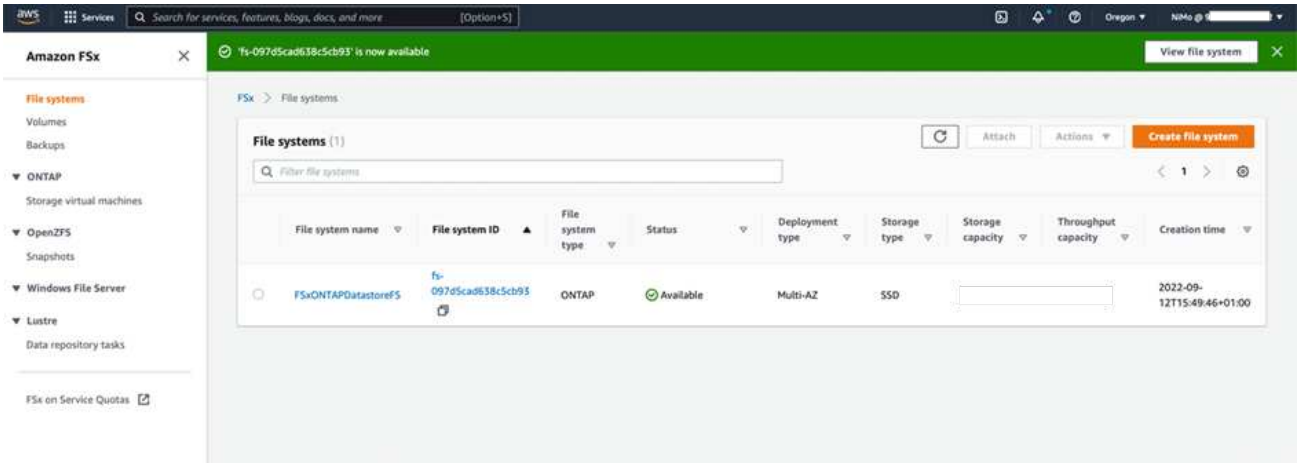


During initial migration scenario, the default snapshot policy can cause datastore capacity full issues. To overcome it, modify the snapshot policy to suit the needs.

8. Review the file system configuration shown on the **Create File System** page.

9. Click **Create File System**.





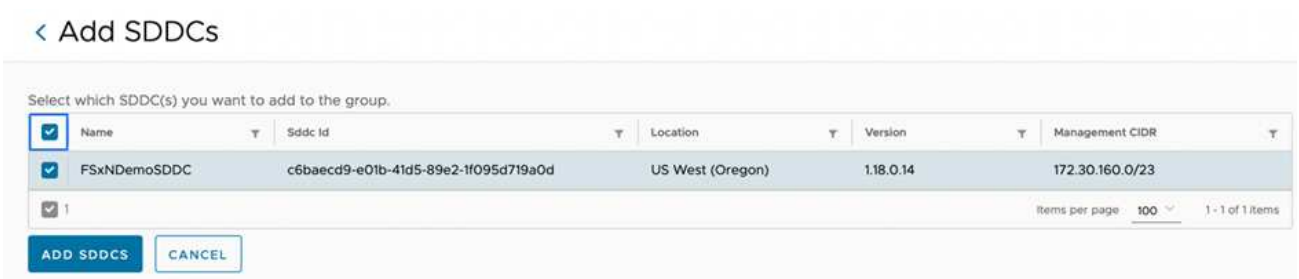
Repeat the previous steps to create more storage virtual machines or file systems and the datastore volumes according to the capacity and performance requirements.

To learn about Amazon FSx ONTAP performance, see Amazon FSx ONTAP performance.

▼ Step 2: Create SDDC group

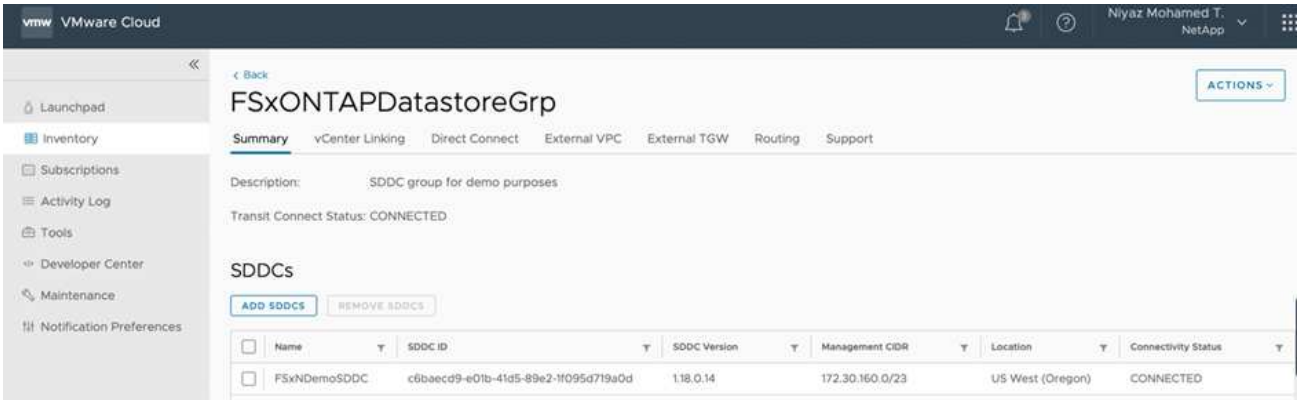
After the file systems and SVMs have been created, use VMware Console to create an SDDC group and to configure VMware Transit Connect. To do so, complete the following steps and remember that you must navigate between the VMware Cloud Console and the AWS Console.

1. Log into the VMC Console at `https://vmc.vmware.com`.
2. On the **Inventory** page, click **SDDC Groups**.
3. On the **SDDC Groups** tab, click **ACTIONS** and select **Create SDDC Group**. For demo purposes, the SDDC group is called `FSxONTAPDatastoreGrp`.
4. On the Membership grid, select the SDDCs to include as group members.



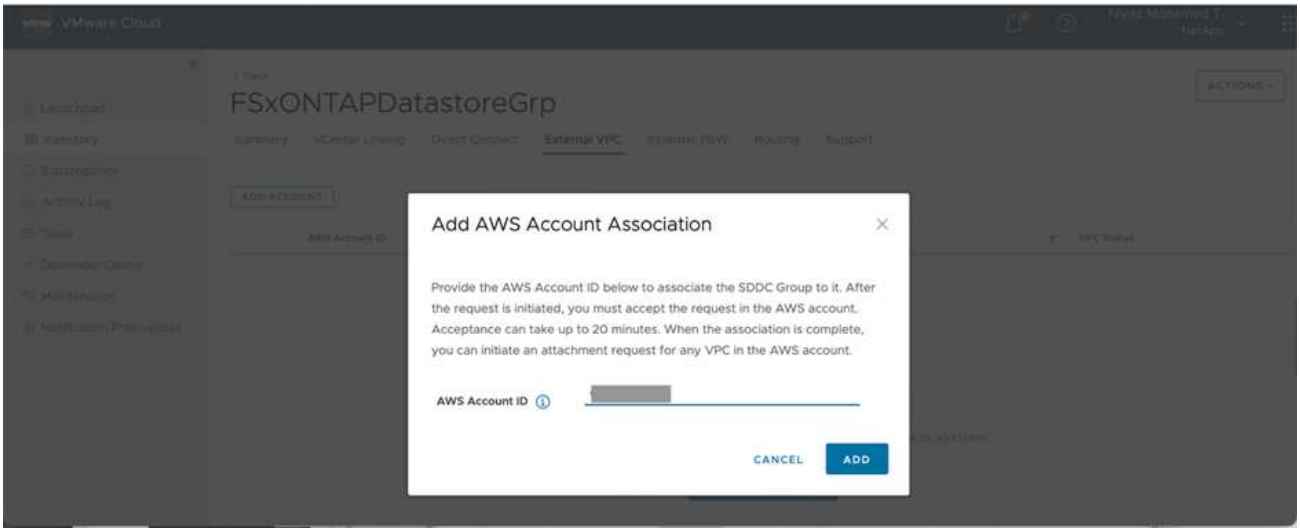
5. Verify that "Configuring VMware Transit Connect for your group will incur charges per attachment and

data transfers" is checked, then select **Create Group**. The process can take a few minutes to complete.

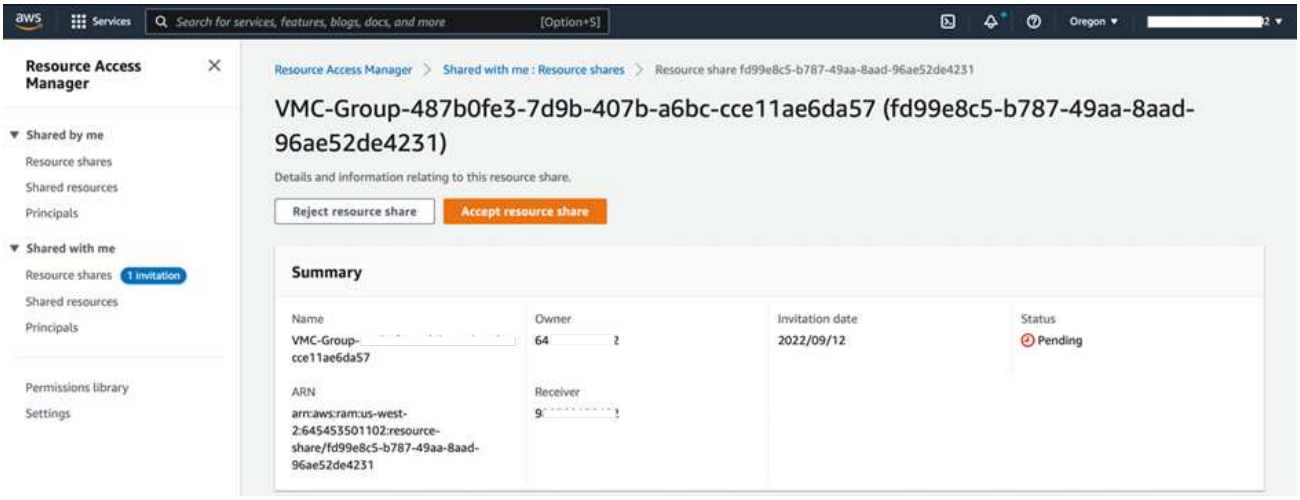


▼ **Step 3: Configure VMware Transit connect**

1. Attach the newly created designated VPC to the SDDC group. Select the **External VPC** tab and follow the instructions for attaching an External VPC to the group. This process can take 10-15 minutes to complete.

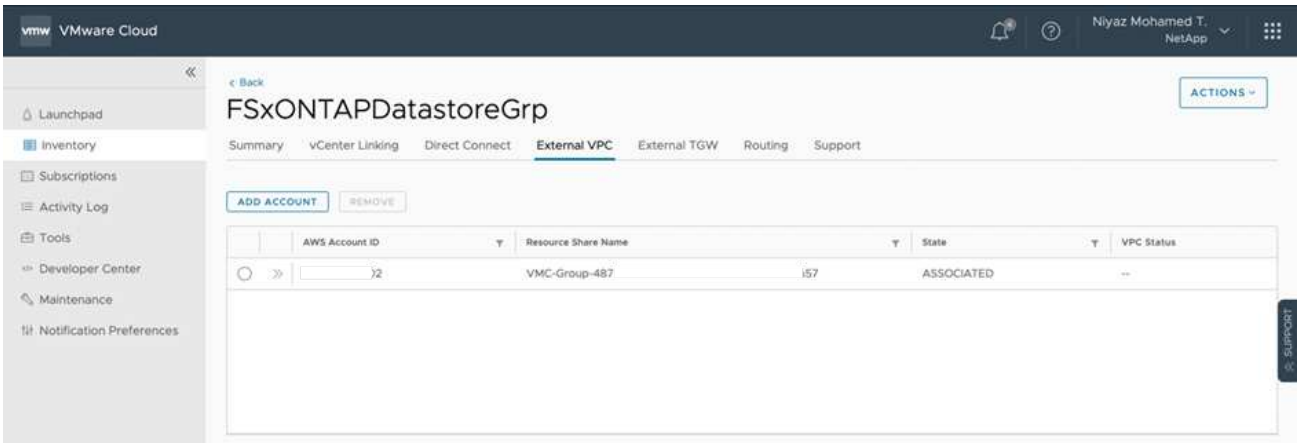


2. Click **Add Account**.
   a. Provide the AWS account that was used to provision the FSx ONTAP file system.
   b. Click **Add**.
3. Back in the AWS console, log into the same AWS account and navigate to the **Resource Access Manager** service page. There is a button for you to accept the resource share.



As part of the external VPC process, you'll be prompted via the AWS console to a new shared resource via the Resource Access Manager. The shared resource is the AWS Transit Gateway managed by VMware Transit Connect.

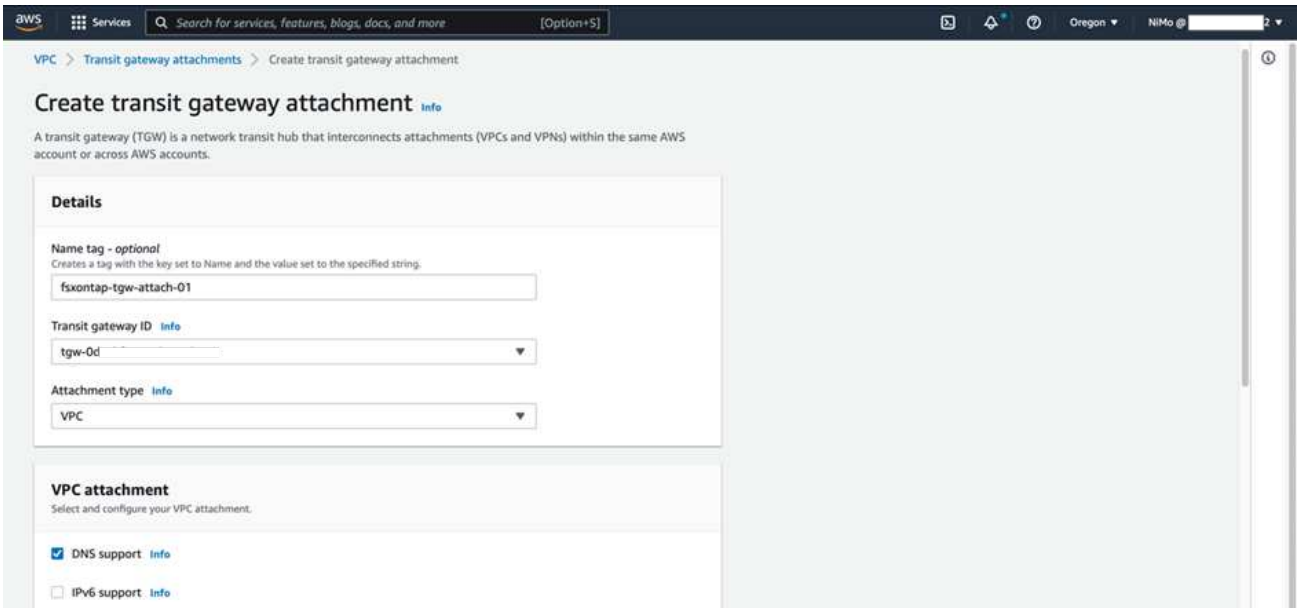
4. Click **Accept resource share**.



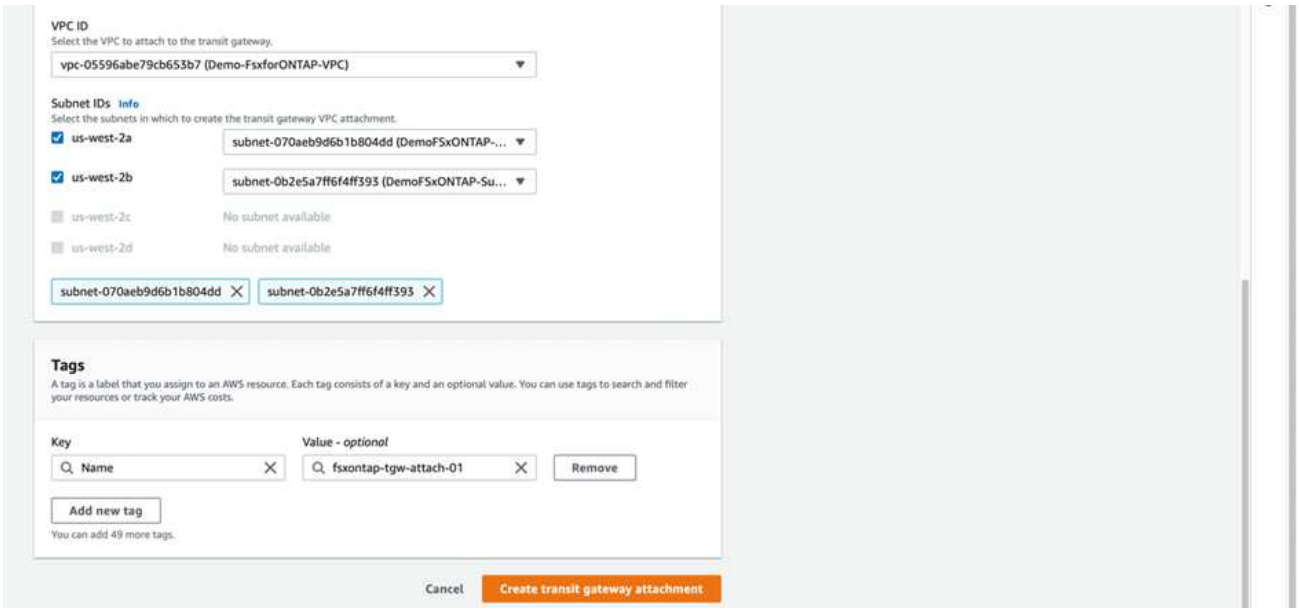
5. Back in the VMC Console, you now see that the External VPC is in an associated state. This can take several minutes to appear.

▼ Step 4: Create transit gateway attachment

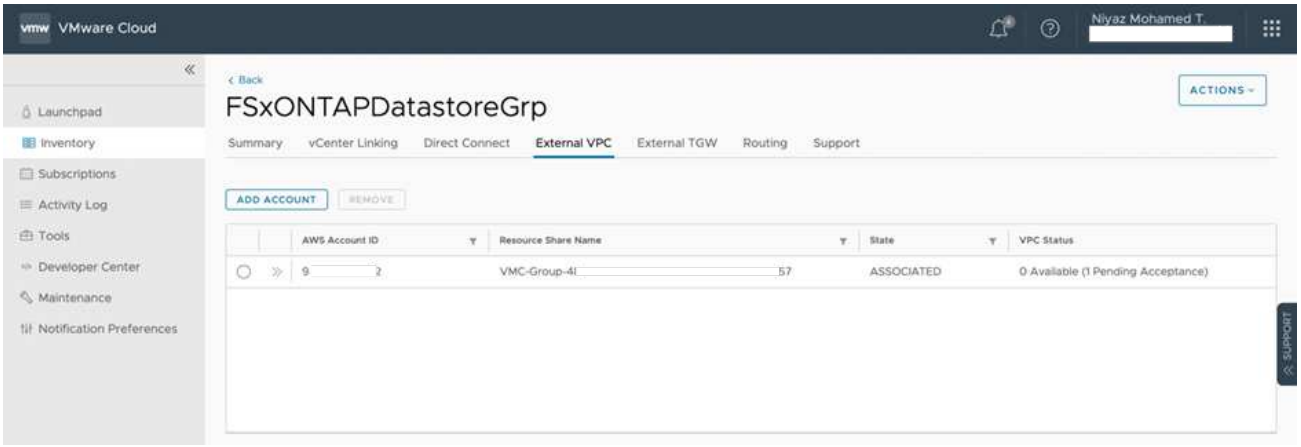
1. In the AWS Console, go to the VPC service page and navigate to the VPC that was used for provisioning the FSx file system. Here you create a transit gateway attachment by clicking **Transit Gateway Attachment** on the navigation pane on the right.
2. Under **VPC Attachment**, make sure that DNS Support is checked and select the VPC in which FSx ONTAP was deployed.

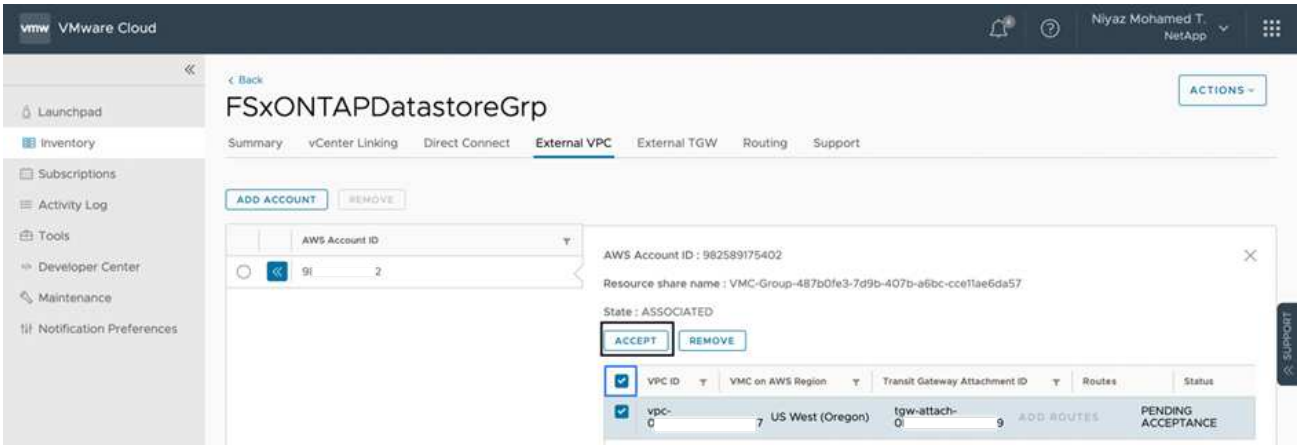


3. Click **Create transit gateway attachment**.



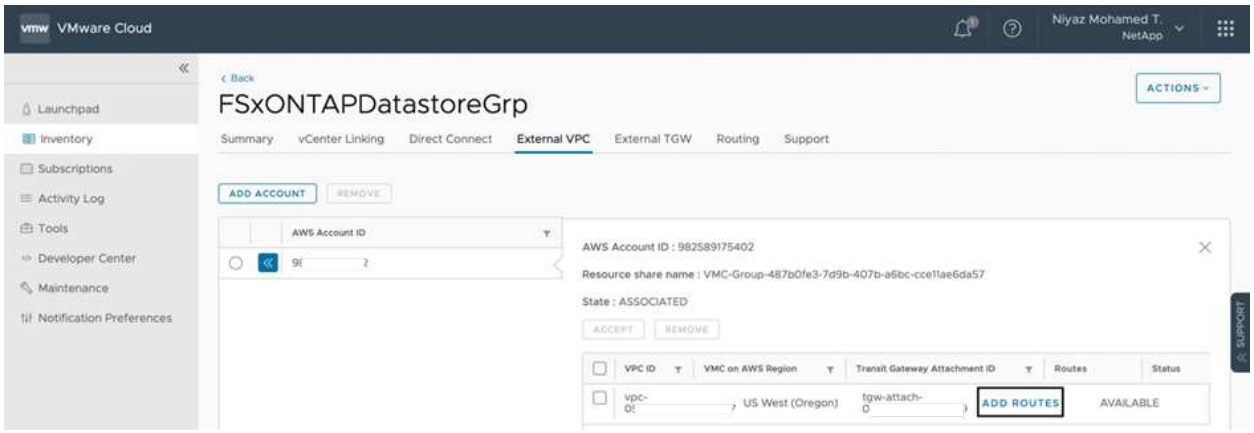
4. Back in the VMware Cloud Console, navigate back to SDDC Group > External VPC tab. Select the AWS account ID used for FSx and click the VPC and click **Accept**.

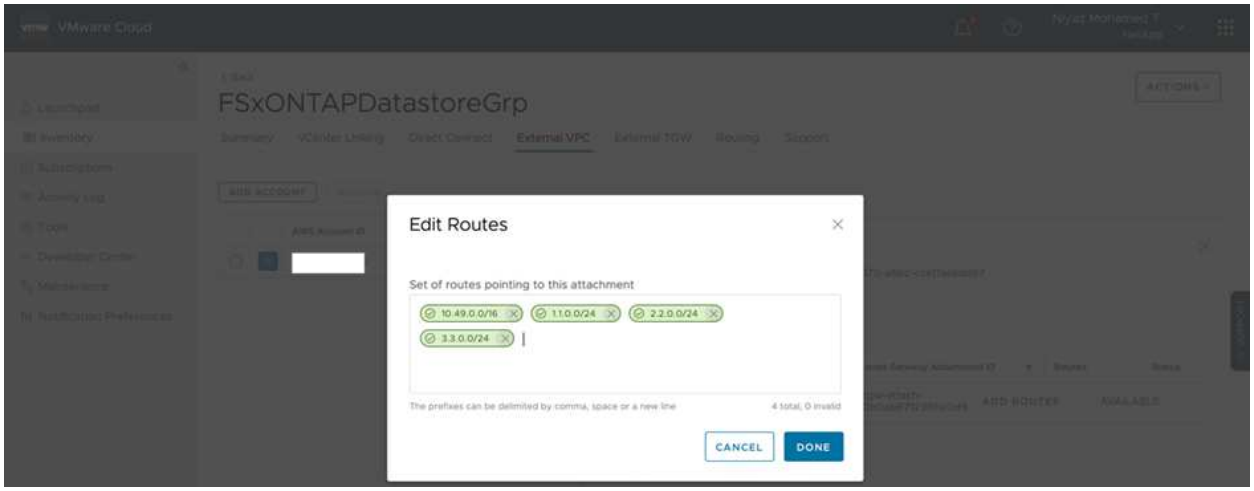




This option may take a several minutes to appear.

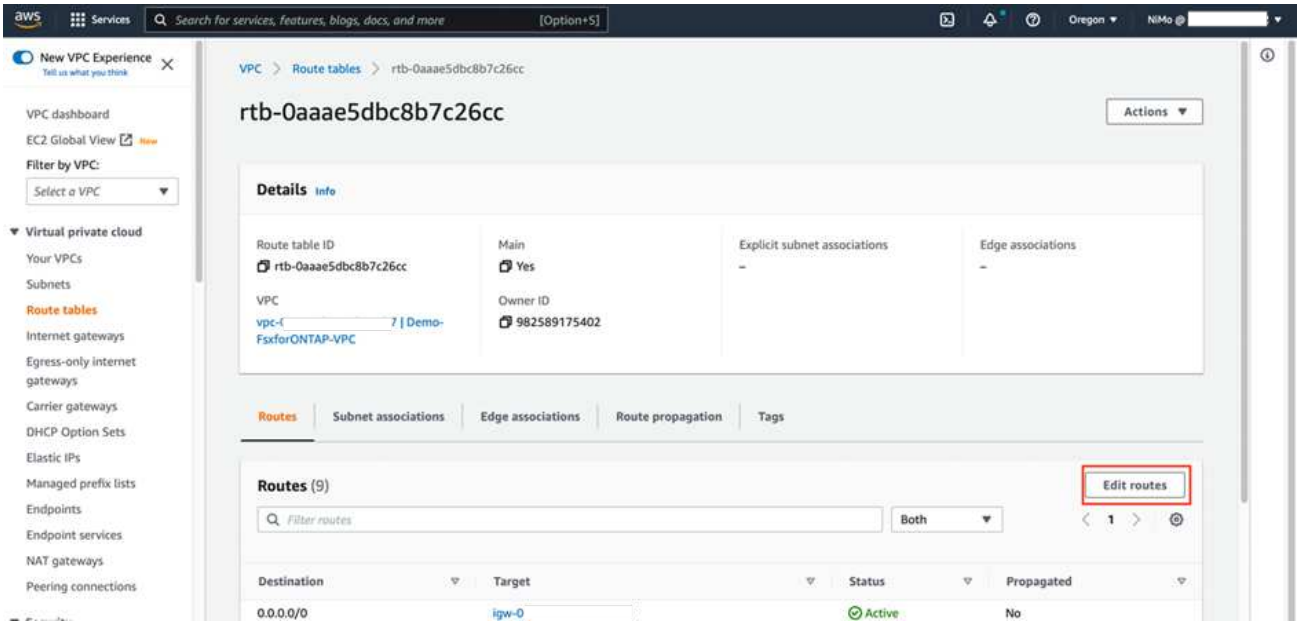
5. Then in the **External VPC** tab in the **Routes** column, click the **Add Routes** option and add in the required routes:
   - A route for the floating IP range for Amazon FSx ONTAP floating IPs.
   - A route for the newly created external VPC address space.



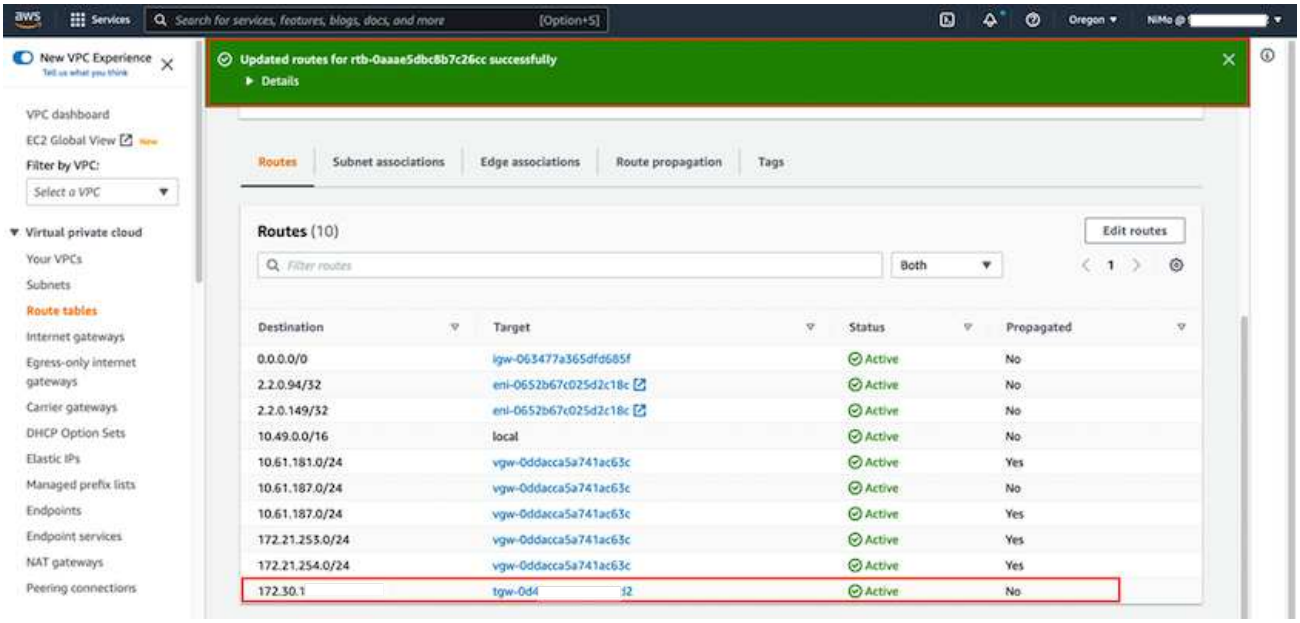


▼ **Step 5: Configure routing (AWS VPC and SDDC) and security groups**

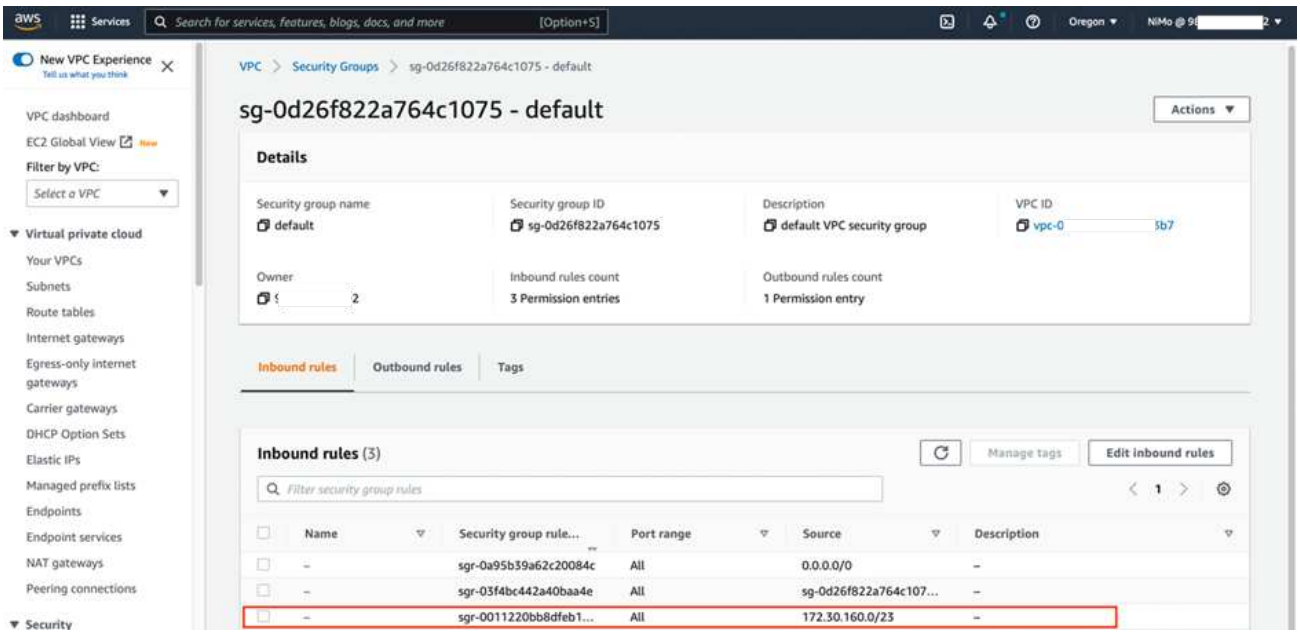
1. In the AWS Console, create the route back to the SDDC by locating the VPC in the VPC service page and select the **main** route table for the VPC.
2. Browse to the route table in the lower panel and click **Edit routes**.



3. In the **Edit routes** panel, click **Add route** and enter the CIDR for the SDDC infrastructure by selecting **Transit Gateway**, and the associated TGW ID. Click **Save changes**.



4. Next step is to verify that the security group in the associated VPC is updated with the correct inbound rules for the SDDC Group CIDR.
5. Update the inbound rule with the CIDR block of the SDDC infrastructure.



Verify that the VPC (where FSx ONTAP resides) route table is updated to avoid connectivity issues.

Update the security group to accept NFS traffic.

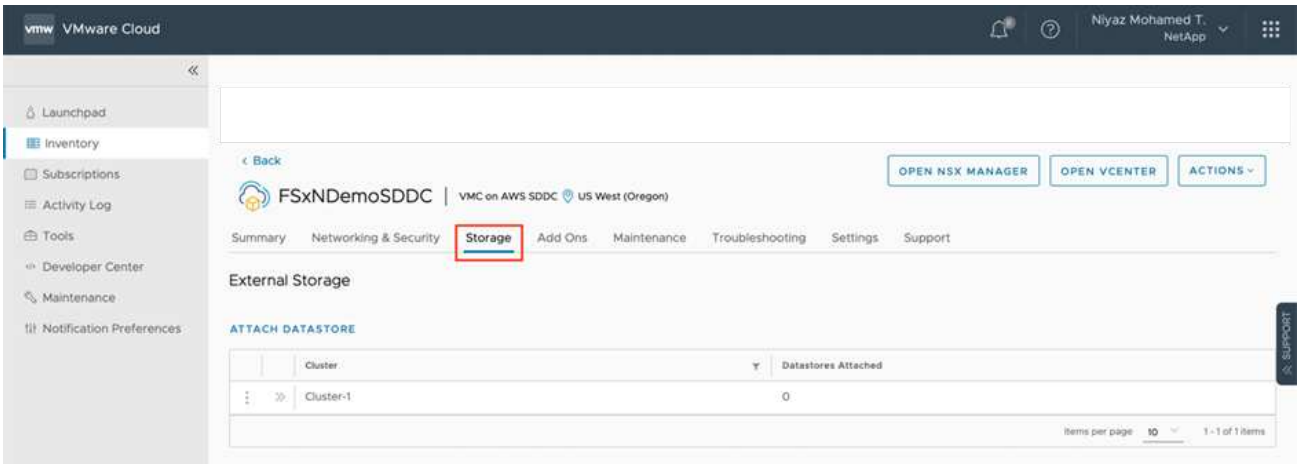
This is the final step in preparing the connectivity to the appropriate SDDC. With the file system configured, routes added, and security groups updated, it's time to mount the datastore(s).

▼ Step 6: Attach NFS volume as a datastore to SDDC cluster

After the file system is provisioned and the connectivity is in place, access VMware Cloud Console to mount

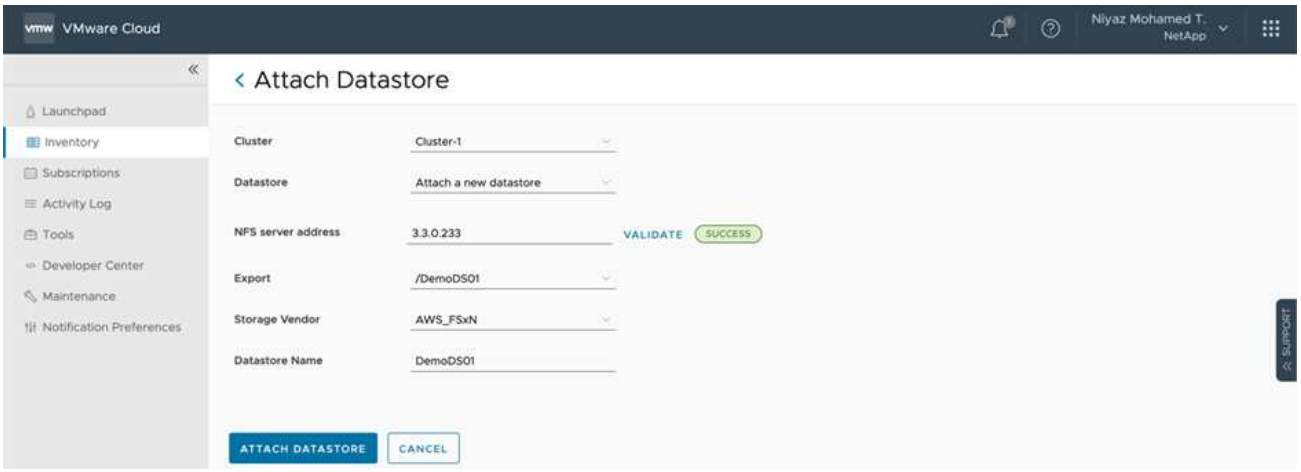
the NFS datastore.

1. In the VMC Console, open the **Storage** tab of the SDDC.

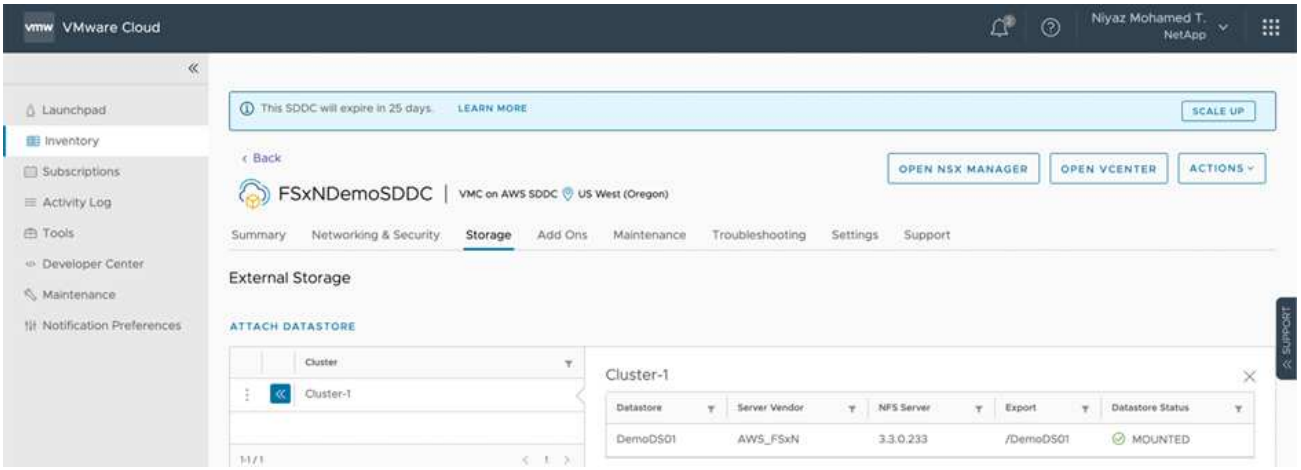


2. Click **ATTACH DATASTORE** and fill in the required values.

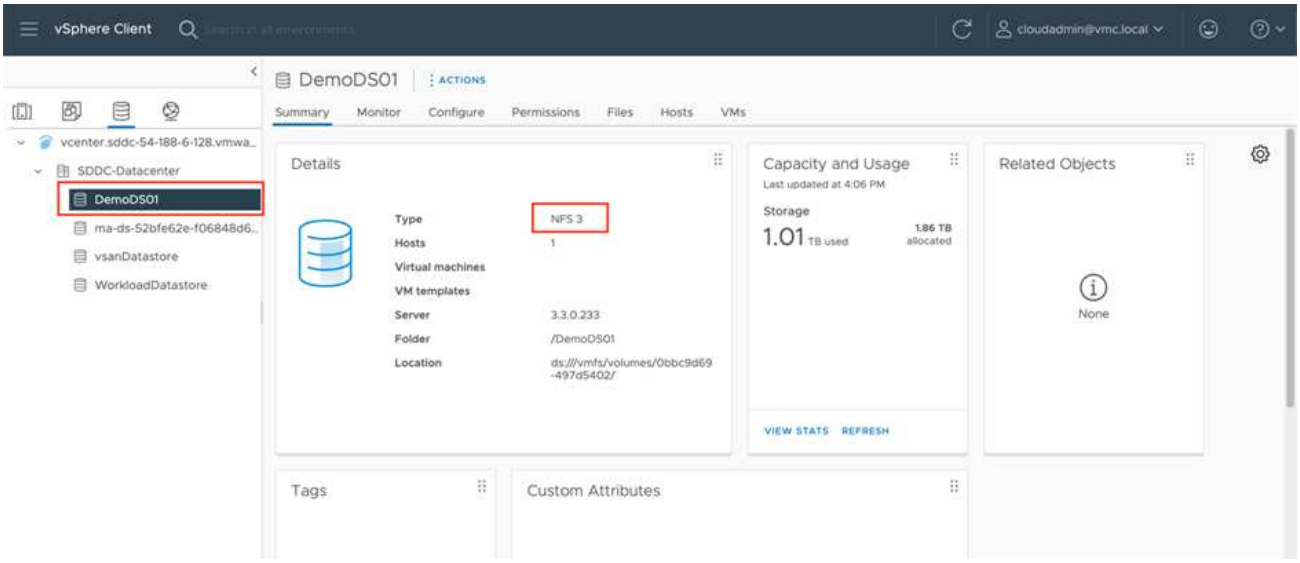
> NFS server address is the NFS IP address which can be found under the FSx > Storage virtual machines tab > Endpoints within AWS console.



3. Click **ATTACH DATASTORE** to attach the datastore to the cluster.



4. Validate the NFS datastore by accessing vCenter as shown below:



NetApp Guest Connected Storage Options for AWS

AWS supports guest connected NetApp storage with the native FSx service (FSx ONTAP) or with Cloud Volumes ONTAP (CVO).

FSx ONTAP

Amazon FSx ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service.

FSx ONTAP provides feature-rich, fast, and flexible shared file storage that's broadly accessible from Linux, Windows, and macOS compute instances running in AWS or on premises. FSx ONTAP offers high-performance solid state drive (SSD) storage with submillisecond latencies. With FSx ONTAP, you can achieve SSD levels of performance for your workload while paying for SSD storage for only a small fraction of your data.

Managing your data with FSx ONTAP is easier because you can snapshot, clone, and replicate your files with the click of a button. In addition, FSx ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

FSx ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx ONTAP supports popular data security and antivirus applications.

FSx ONTAP as guest connected storage

▼ **Configure Amazon FSx ONTAP with VMware Cloud on AWS**

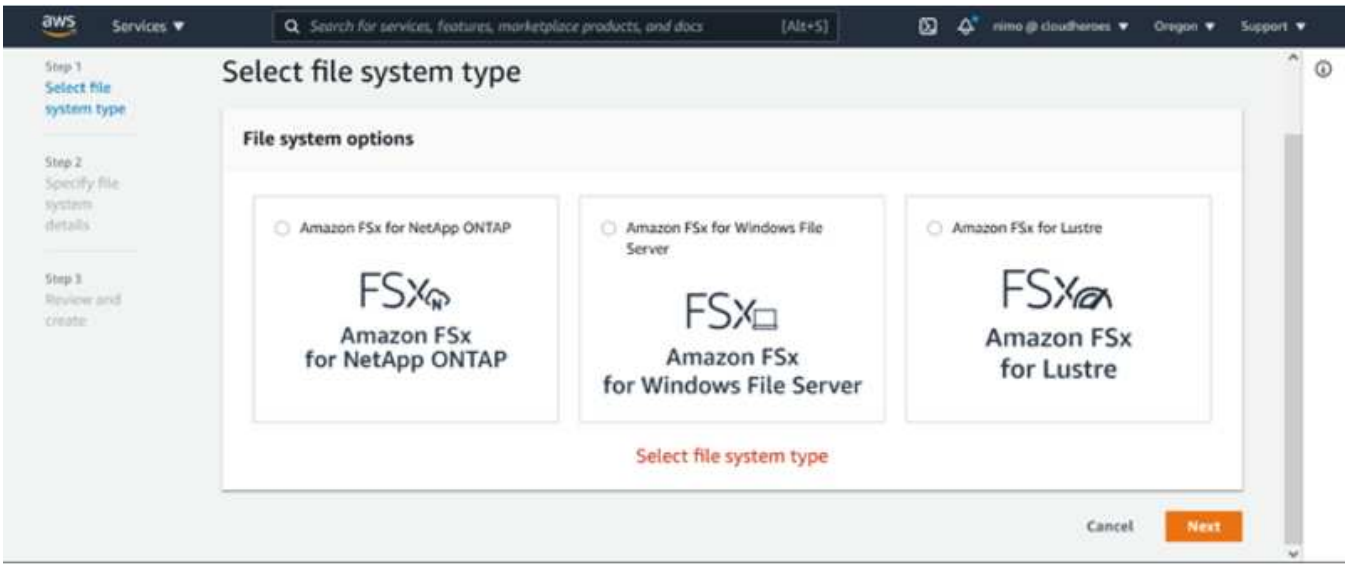
Amazon FSx ONTAP files shares and LUNs can be mounted from VMs that are created within the VMware SDDC environment at VMware Cloud at AWS. The volumes can also be mounted on the Linux client and mapped on the Windows client using the NFS or SMB protocol, and LUNS can be accessed on Linux or Windows clients as block devices when mounted over iSCSI. Amazon FSx for the NetApp ONTAP file system can be set up quickly with the following steps.

Amazon FSx ONTAP and VMware Cloud on AWS must be in the same availability zone to achieve better performance and avoid data transfer charges between availability zones.

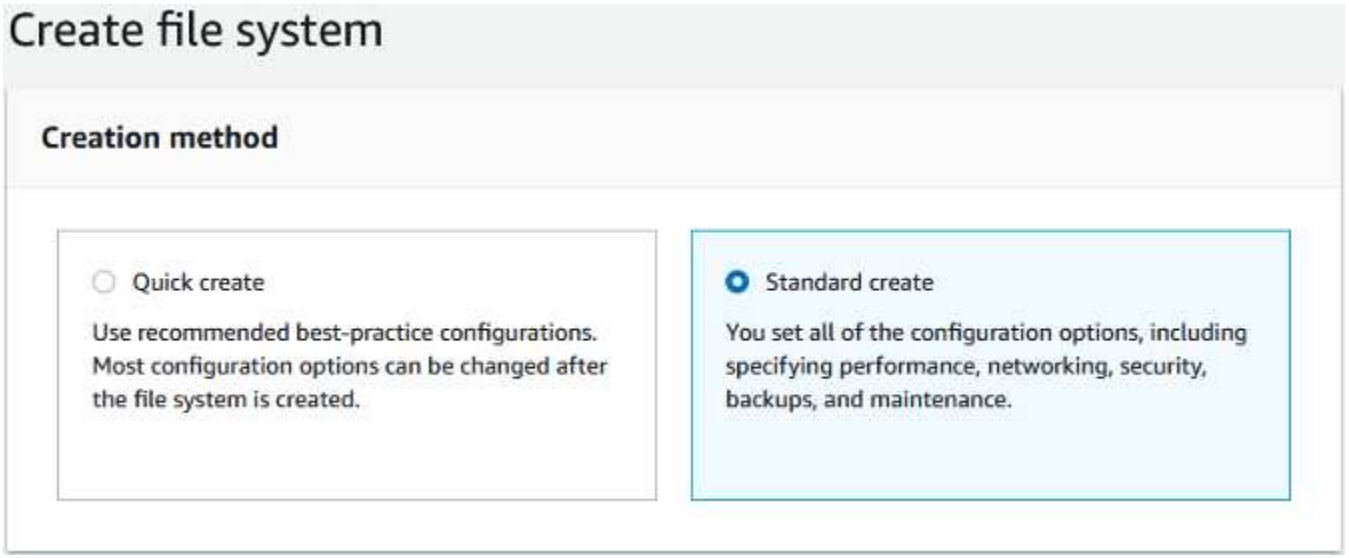
▼ **Create and mount Amazon FSx ONTAP volumes**

To create and mount Amazon FSx ONTAP file system, complete the following steps:

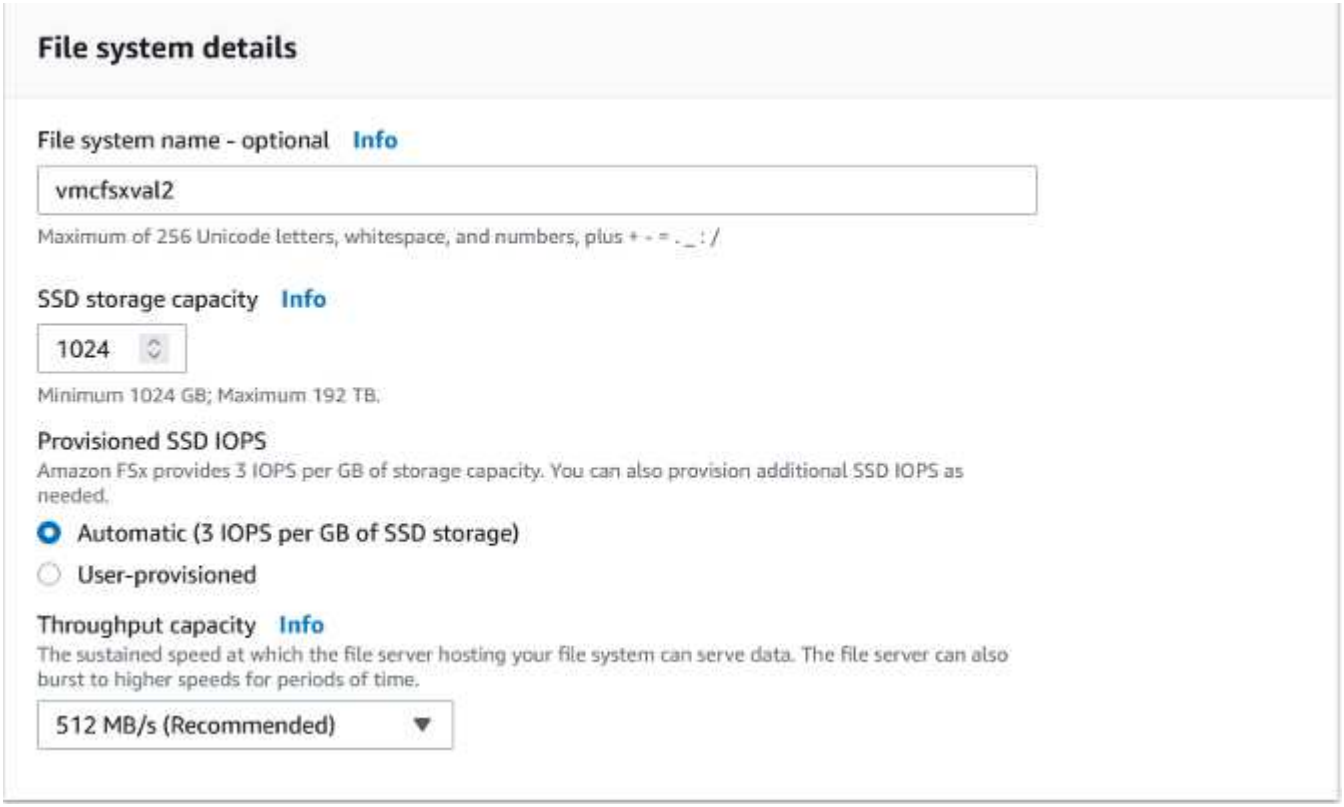
1. Open the Amazon FSx console and choose Create file system to start the file system creation wizard.
2. On the Select File System Type page, choose Amazon FSx ONTAP, and then choose Next. The Create File System page appears.



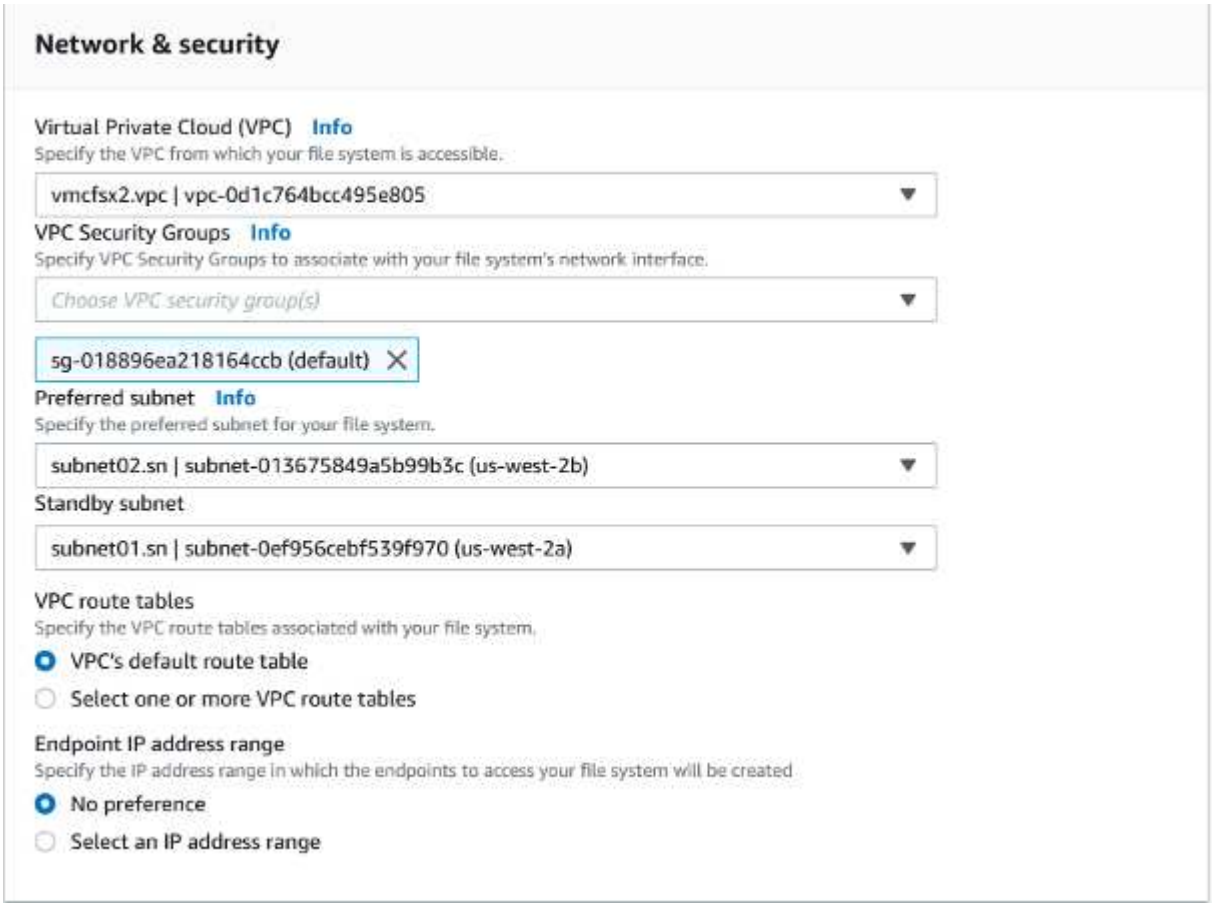
1. In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.



1. For the creation method, choose Standard Create. You can also choose Quick Create, but this document uses the Standard create option.

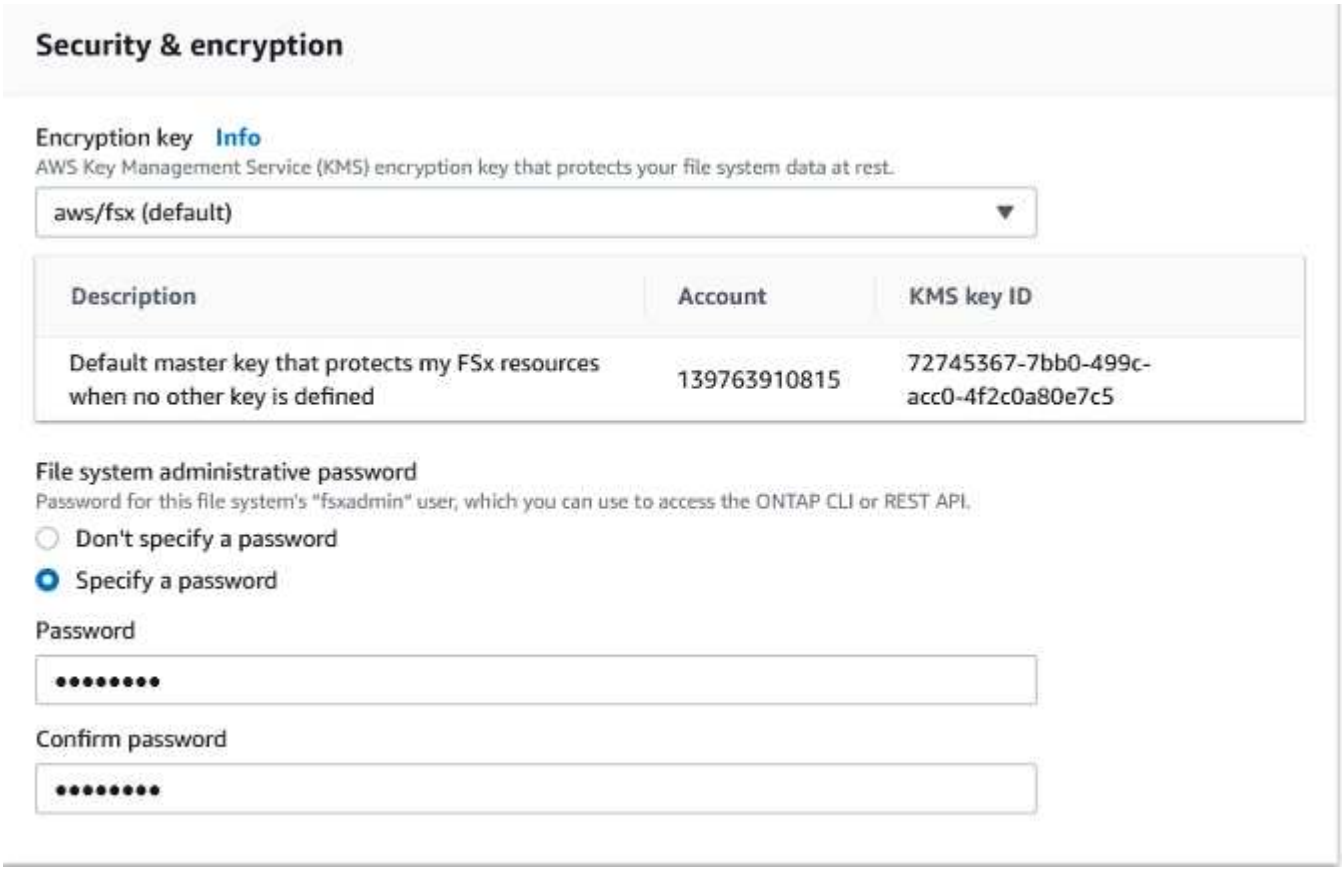


1. In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

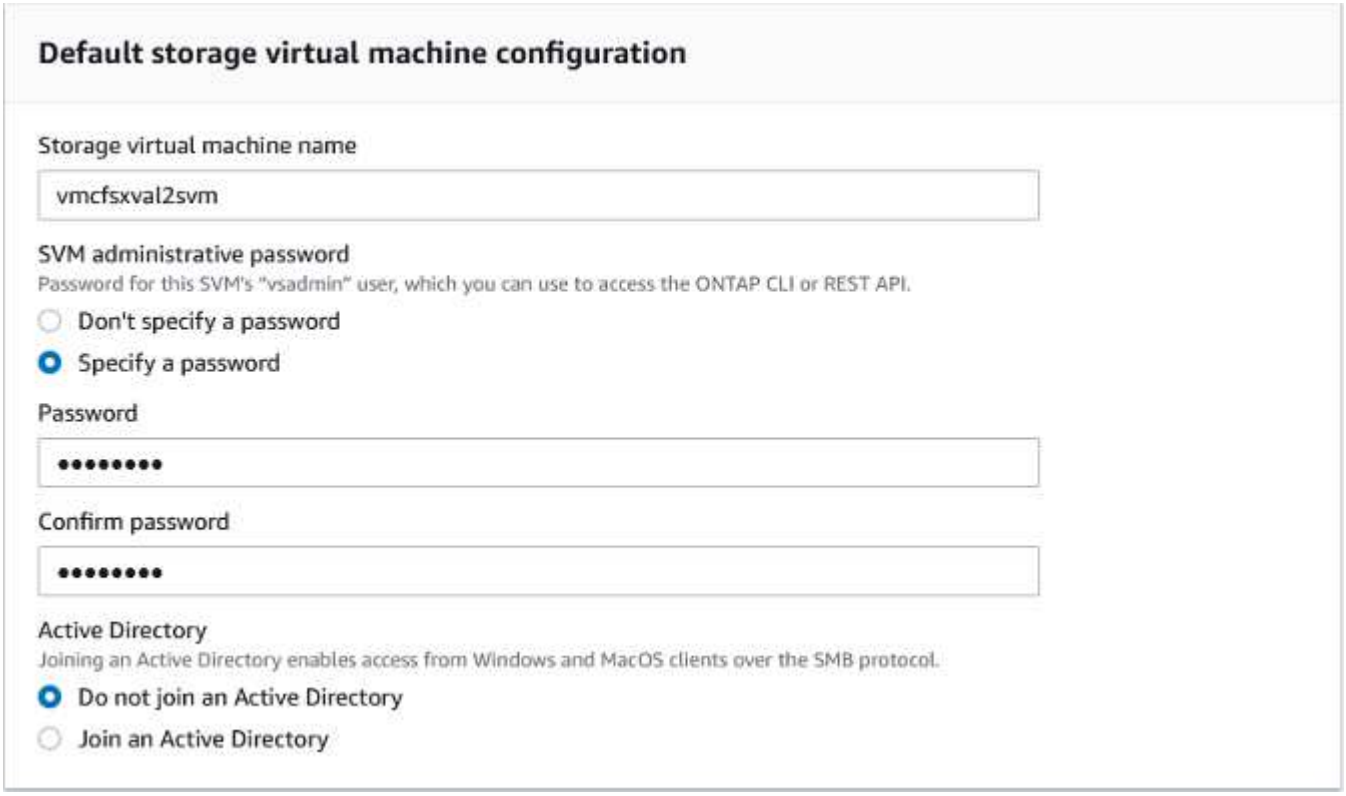


In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

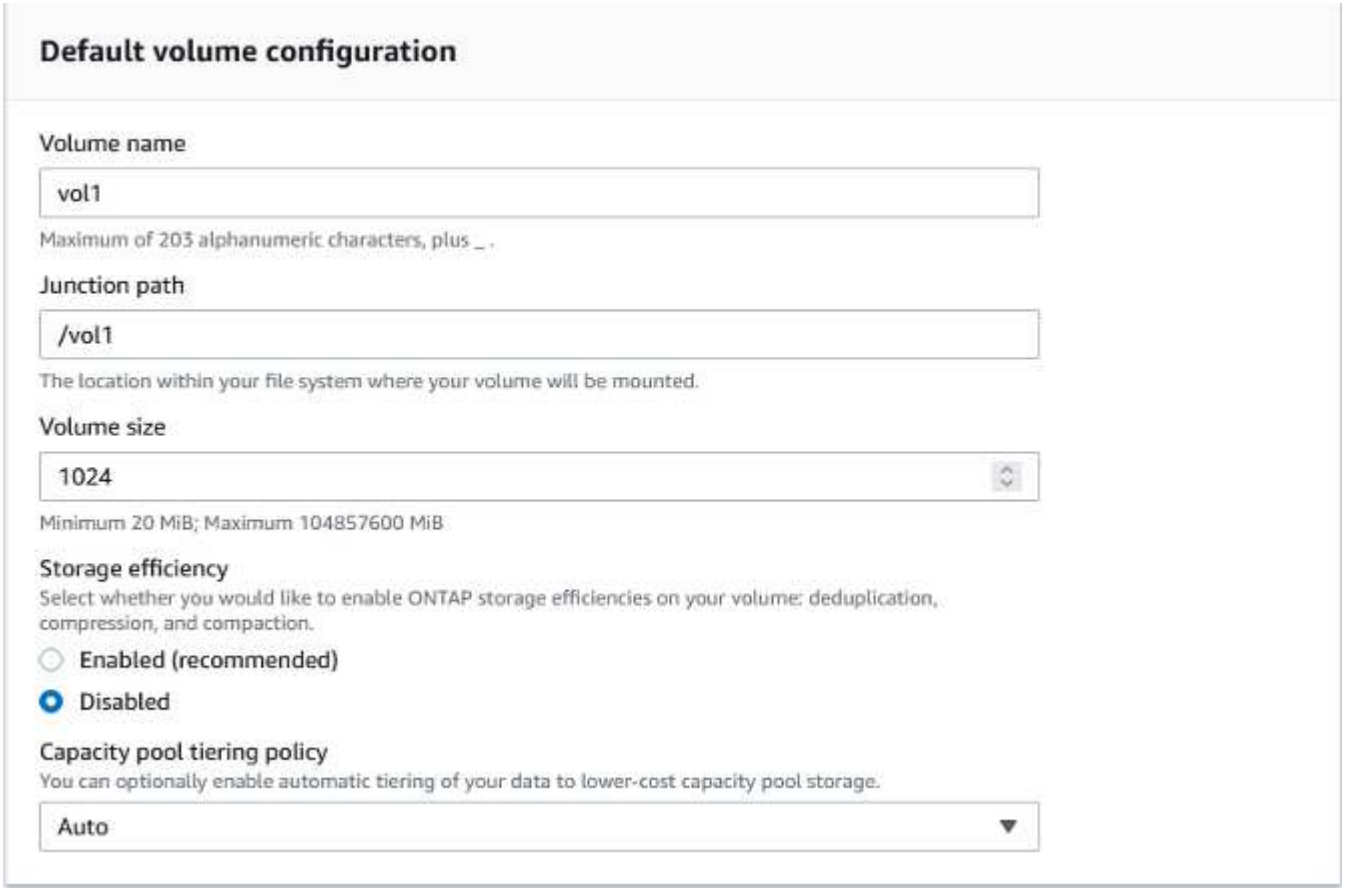
1. In the Security & Encryption section, for the Encryption Key, choose the AWS Key Management Service (AWS KMS) encryption key that protects the file system's data at rest. For the File System Administrative Password, enter a secure password for the fsxadmin user.



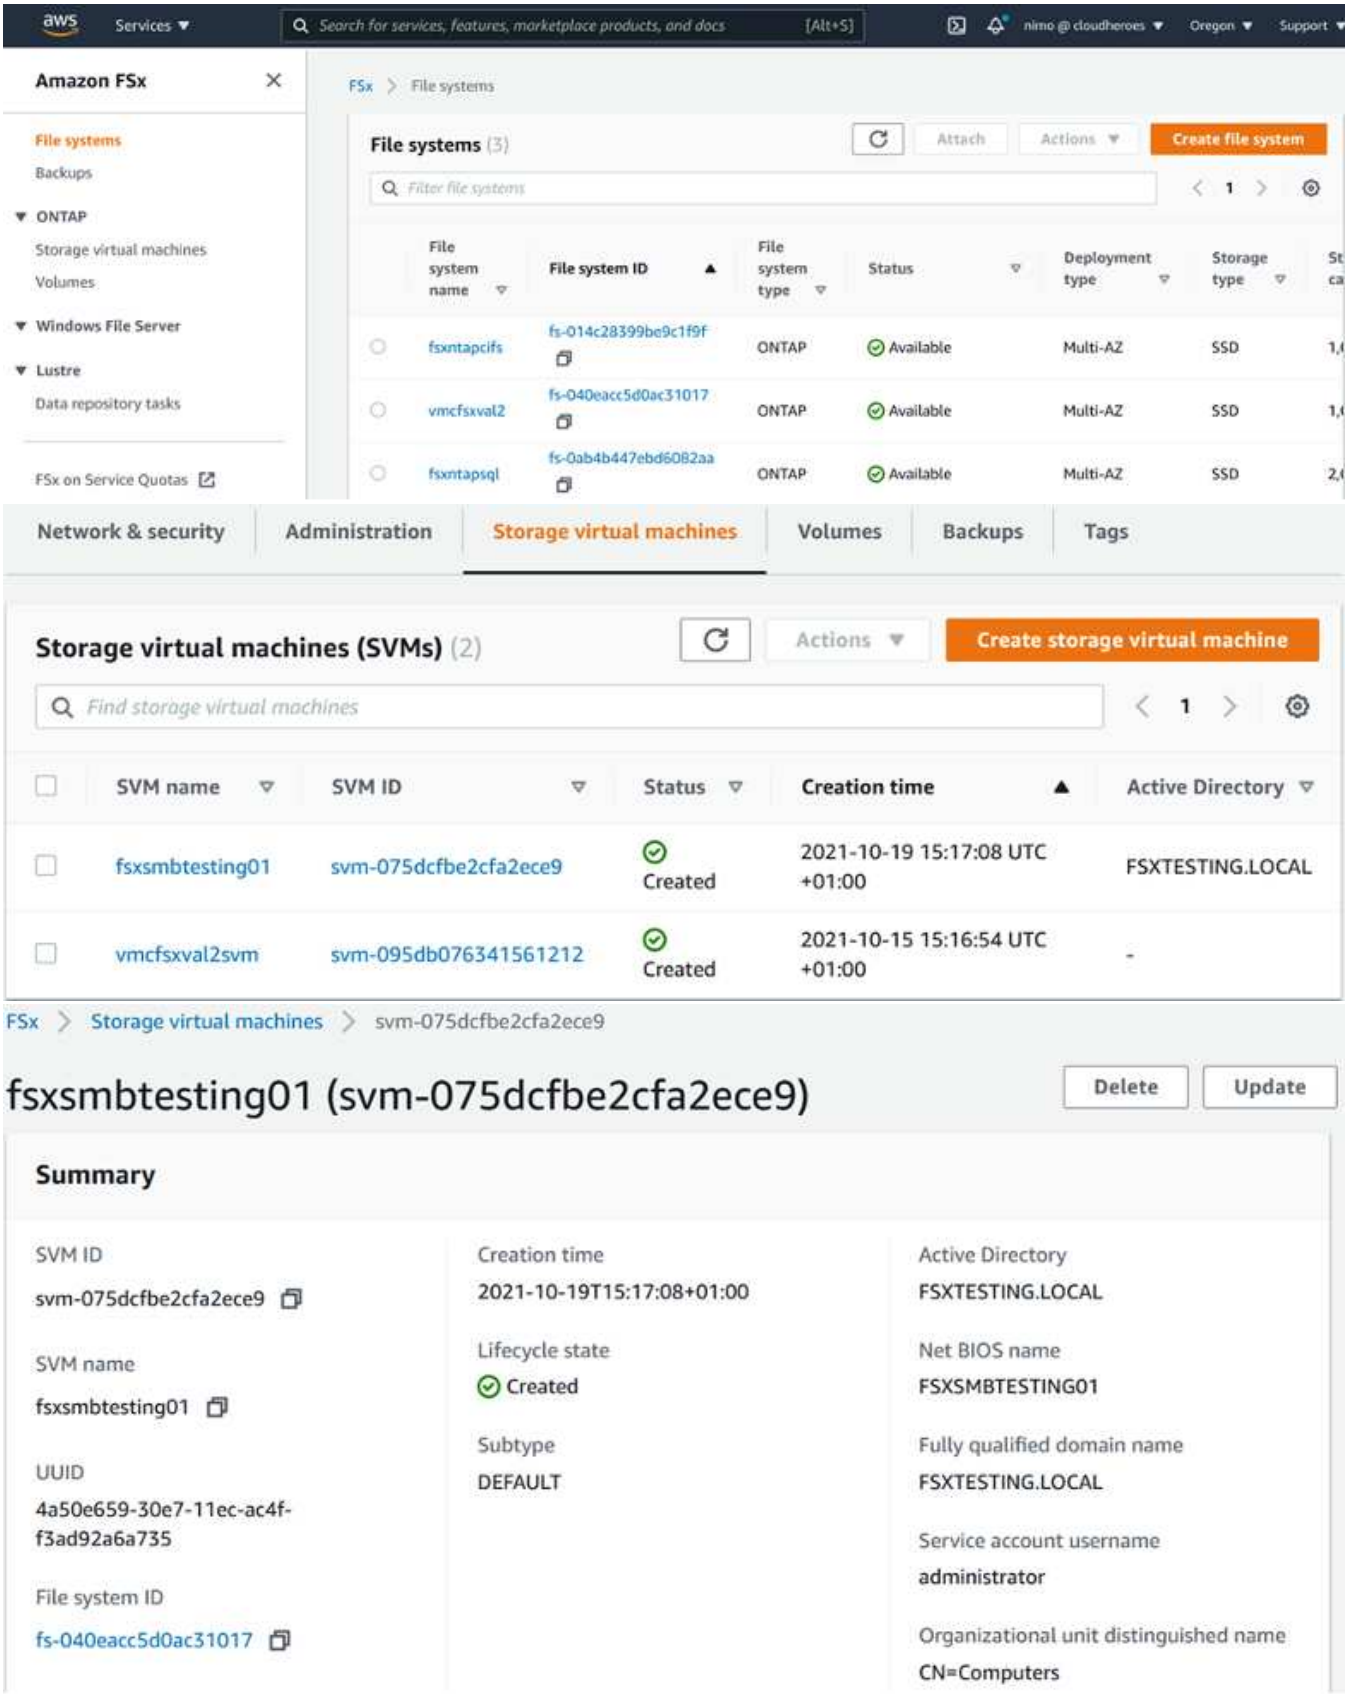
1. In virtual machine and specify the password to use with vsadmin for administering ONTAP using REST APIs or the CLI. If no password is specified, a fsxadmin user can be used for administering the SVM. In the Active Directory section, make sure to join Active Directory to the SVM for provisioning SMB shares. In the Default Storage Virtual Machine Configuration section, provide a name for the storage in this validation, SMB shares are provisioned using a self-managed Active Directory domain.



1. In the Default Volume Configuration section, specify the volume name and size. This is an NFS volume. For Storage Efficiency, choose Enabled to turn on the ONTAP storage efficiency features (compression, deduplication, and compaction) or Disabled to turn them off.



1. Review the file system configuration shown on the Create File System page.
2. Click Create File System.

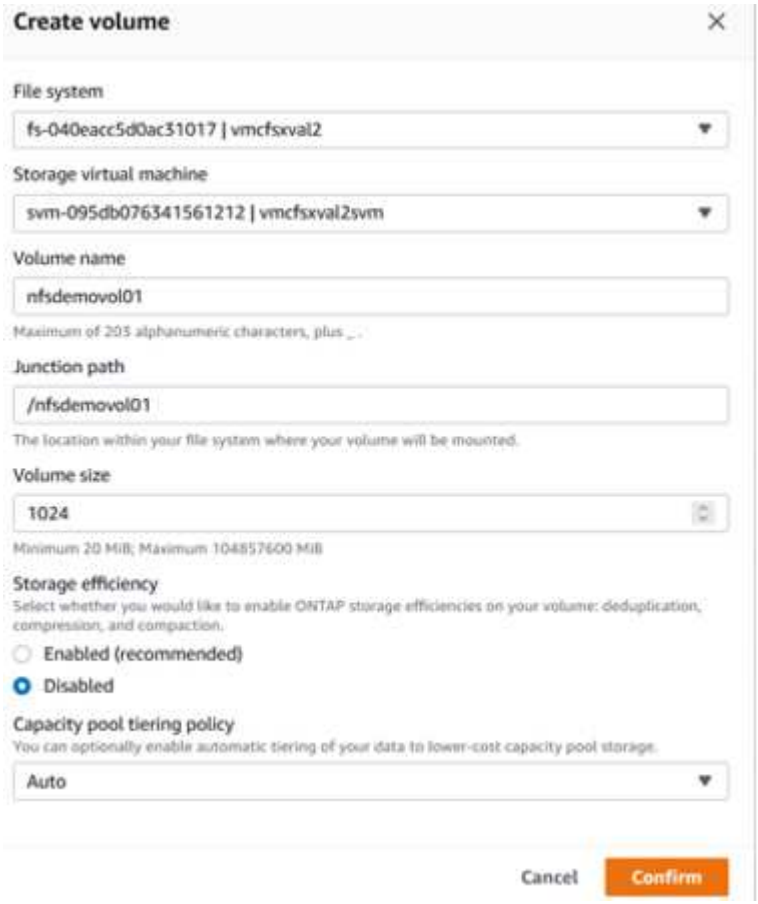


For more detailed information, see Getting started with Amazon FSx ONTAP.

After the file system is created as above, create the volume with the required size and protocol.

1. Open the Amazon FSx console.
2. In the left navigation pane, choose File systems, and then choose the ONTAP file system that you want to create a volume for.
3. Select the Volumes tab.
4. Select the Create Volume tab.
5. The Create Volume dialog box appears.

For demo purposes, an NFS volume is created in this section that can be easily mounted on VMs running on VMware cloud on AWS. nfsdemovol01 is created as depicted below:



▼ Mount FSx ONTAP volume on Linux client

To mount the FSx ONTAP volume created in the previous step. from the Linux VMs within VMC on AWS SDDC, complete the following steps:

1. Connect to the designated Linux instance.
2. Open a terminal on the instance using Secure Shell (SSH) and log in with the appropriate credentials.
3. Make a directory for the volume's mount point with the following command:

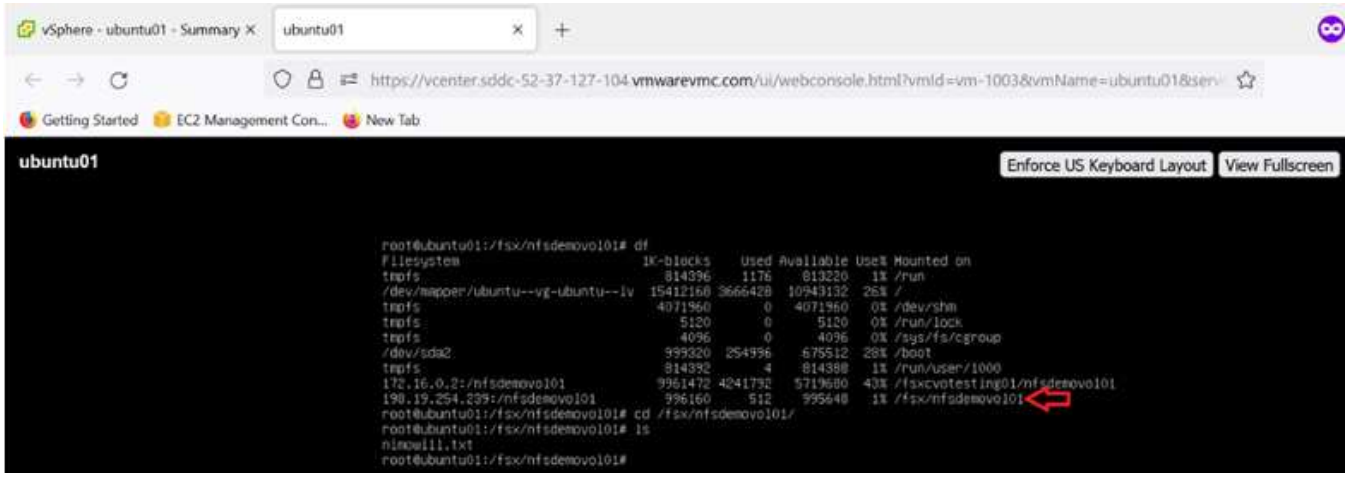
```
$ sudo mkdir /fsx/nfsdemovol01
```

4. Mount the Amazon FSx ONTAP NFS volume to the directory that is created in the previous step.

```
sudo mount -t nfs nfsvers=4.1,198.19.254.239:/nfsdemovol01
/fsx/nfsdemovol01
```

```
root@ubuntu01:/fsx/nfsdemovol01# mount -t nfs 198.19.254.239:/nfsdemovol01 /fsx/nfsdemovol01
```

1. Once executed, run the df command to validate the mount.



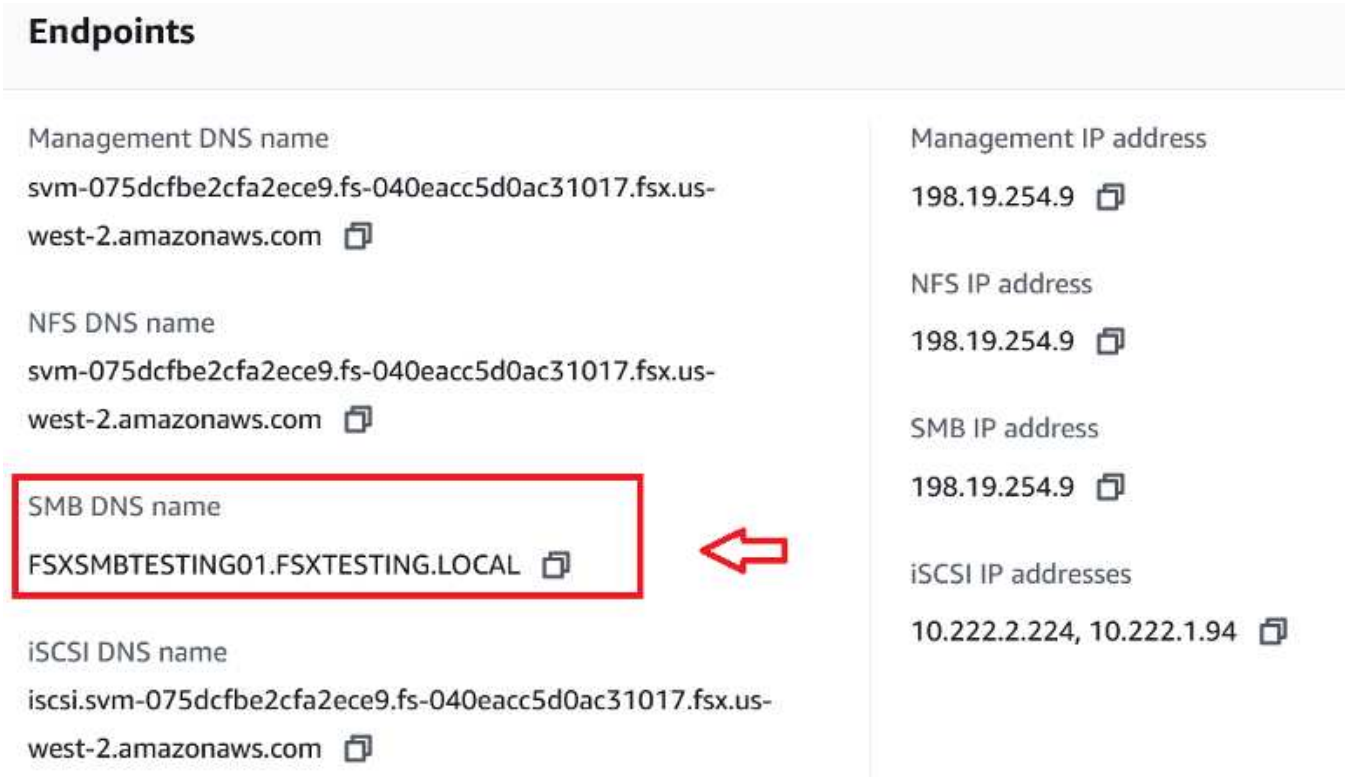
Mount FSx ONTAP volume on Linux client

▼ **Attach FSx ONTAP volumes to Microsoft Windows clients**

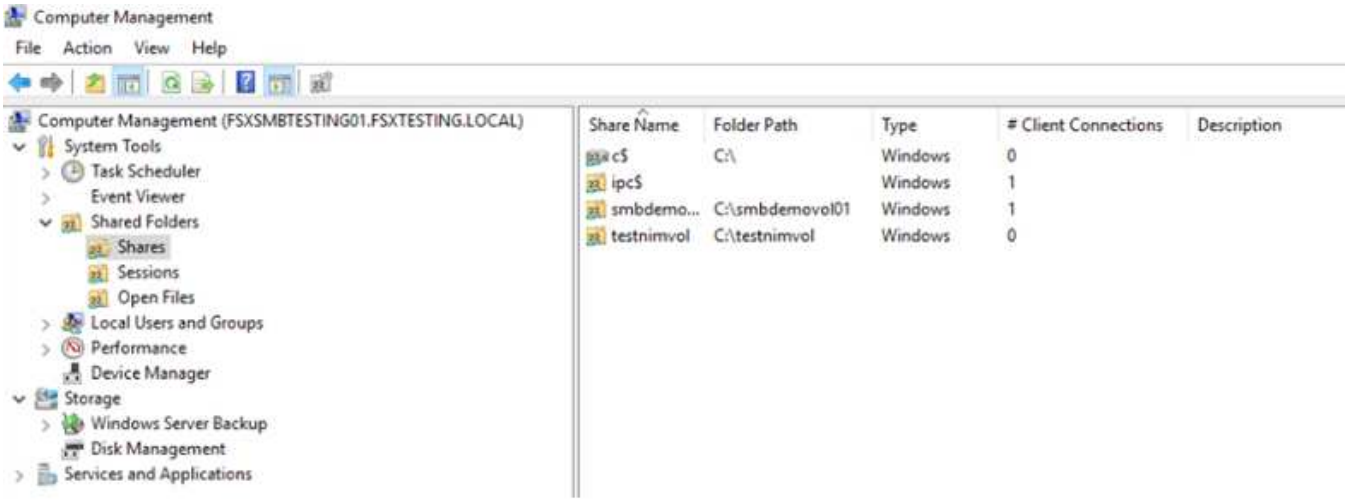
To manage and map file shares on an Amazon FSx file system, the Shared Folders GUI must be used.

1. Open the Start menu and run fsmgmt.msc using Run As Administrator. Doing this opens the Shared Folders GUI tool.
2. Click Action > All tasks and choose Connect to Another Computer.
3. For Another Computer, enter the DNS name for the storage virtual machine (SVM). For example, FSXSMBTESTING01.FSXTESTING.LOCAL is used in this example.

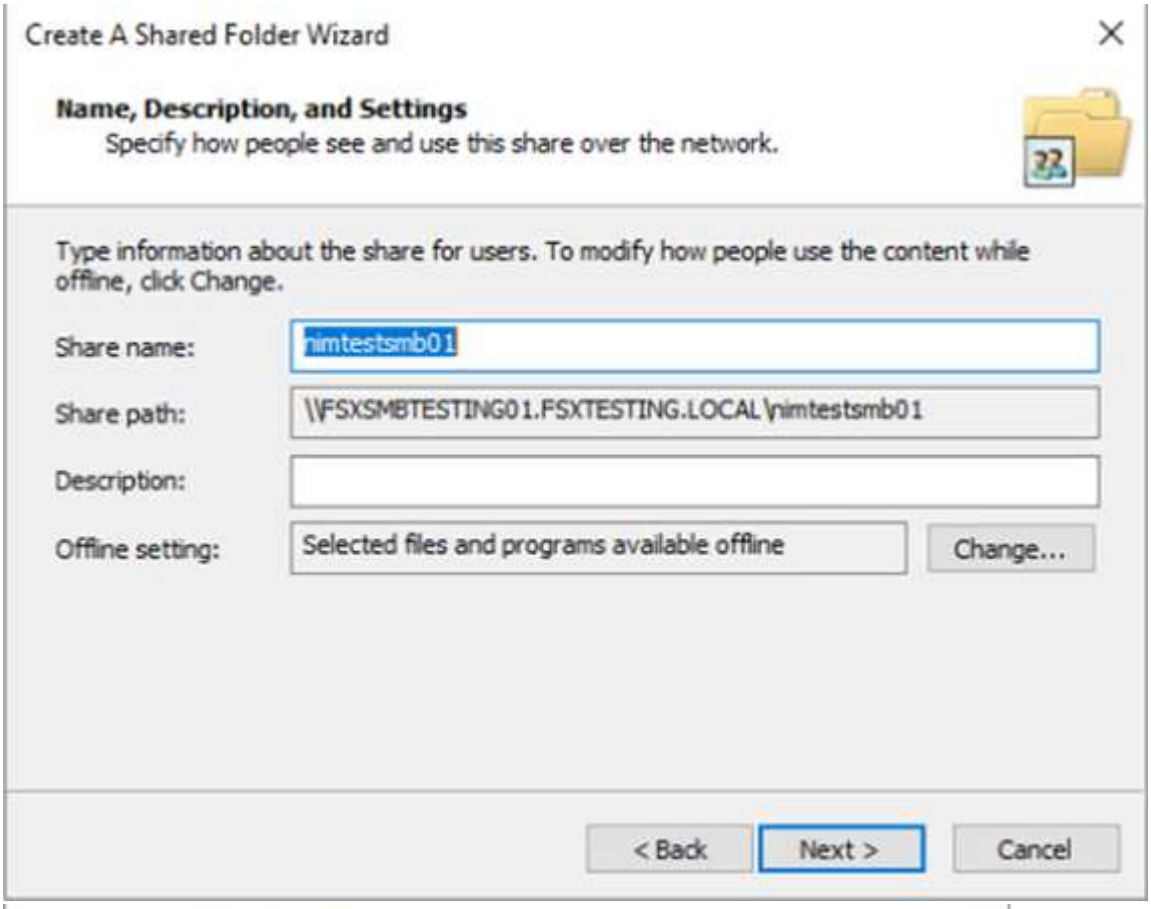
Tp find the SVM's DNS name on the Amazon FSx console, choose Storage Virtual Machines, choose SVM, and then scroll down to Endpoints to find the SMB DNS name. Click OK. The Amazon FSx file system appears in the list for the Shared Folders.



1. In the Shared Folders tool, choose Shares in the left pane to see the active shares for the Amazon FSx file system.



1. Now choose a new share and complete the Create a Shared Folder wizard.

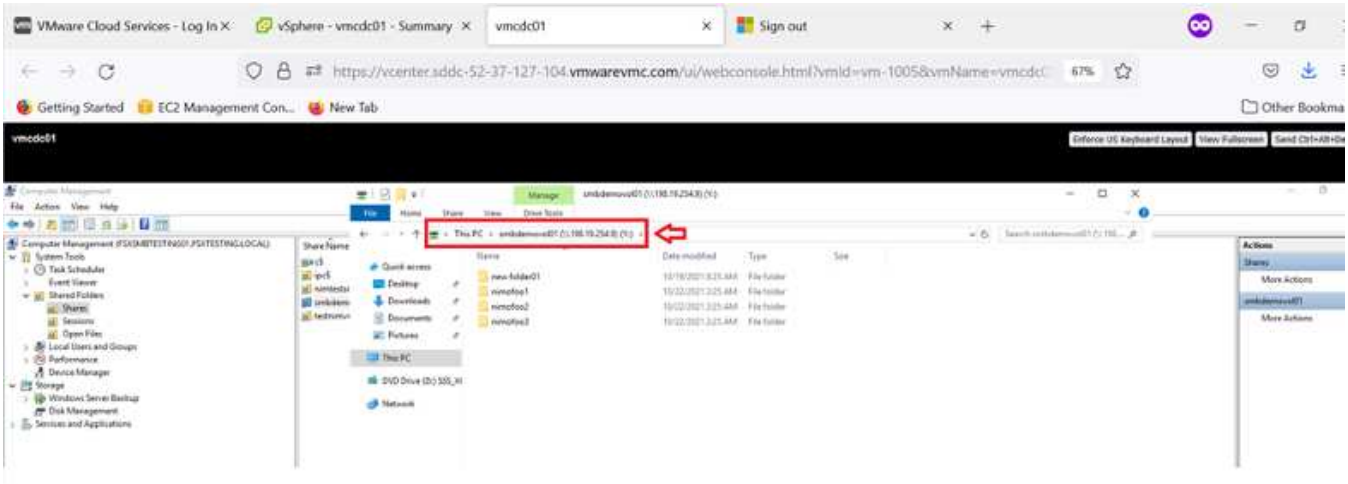




To learn more about creating and managing SMB shares on an Amazon FSx file system, see Creating SMB Shares.

1. After connectivity is in place, the SMB share can be attached and used for application data. To

accomplish this, Copy the share path and use the Map Network Drive option to mount the volume on the VM running on VMware Cloud on the AWS SDDC.



▼ **Connect a FSx ONTAP LUN to a host using iSCSI**

Connect a FSx ONTAP LUN to a host using iSCSI

iSCSI traffic for FSx traverses the VMware Transit Connect/AWS Transit Gateway via the routes provided in the previous section. To configure a LUN in Amazon FSx ONTAP, follow the documentation found here.

On Linux clients, make sure that the iSCSI daemon is running. After the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration with Ubuntu (as an example) here.

In this paper, connecting the iSCSI LUN to a Windows host is depicted:

▼ **Provision a LUN in FSx ONTAP:**

1. Access the NetApp ONTAP CLI using the management port of the FSx for the ONTAP file system.
2. Create the LUNs with the required size as indicated by the sizing output.

```
FsxId040eacc5d0ac31017::> lun create -vserver vmcfsxval2svm -volume
nimfsxscsivol -lun nimofsxlun01 -size 5gb -ostype windows -space
-reserve enabled
```

In this example, we created a LUN of size 5g (5368709120).

1. Create the necessary igroups to control which hosts have access to specific LUNs.

```
FsxId040eacc5d0ac31017::> igroup create -vserver vmcfsxval2svm -igroup
winIG -protocol iscsi -ostype windows -initiator iqn.1991-
05.com.microsoft:vmcdc01.fsxtesting.local

FsxId040eacc5d0ac31017::> igroup show

Vserver   Igroup       Protocol OS Type  Initiators
```

```
--------- ------------ -------- --------
-----------------------------------

vmcfsxval2svm

          ubuntu01     iscsi    linux    iqn.2021-
10.com.ubuntu:01:initiator01

vmcfsxval2svm

          winIG        iscsi    windows  iqn.1991-
05.com.microsoft:vmcdc01.fsxtesting.local
```

Two entries were displayed.

1. Map the LUNs to igroups using the following command:

```
FsxId040eacc5d0ac31017::> lun map -vserver vmcfsxval2svm -path
/vol/nimfsxscsivol/nimofsxlun01 -igroup winIG

FsxId040eacc5d0ac31017::> lun show

Vserver   Path                            State   Mapped   Type
Size

--------- ------------------------------- ------- -------- --------
--------

vmcfsxval2svm

          /vol/blocktest01/lun01          online  mapped   linux
5GB

vmcfsxval2svm

          /vol/nimfsxscsivol/nimofsxlun01 online  mapped   windows
5GB
```

Two entries were displayed.

1. Connect the newly provisioned LUN to a Windows VM:

To connect the new LUN tor a Windows host residing on VMware cloud on AWS SDDC, complete the following steps:

a. RDP to the Windows VM hosted on the VMware Cloud on AWS SDDC.

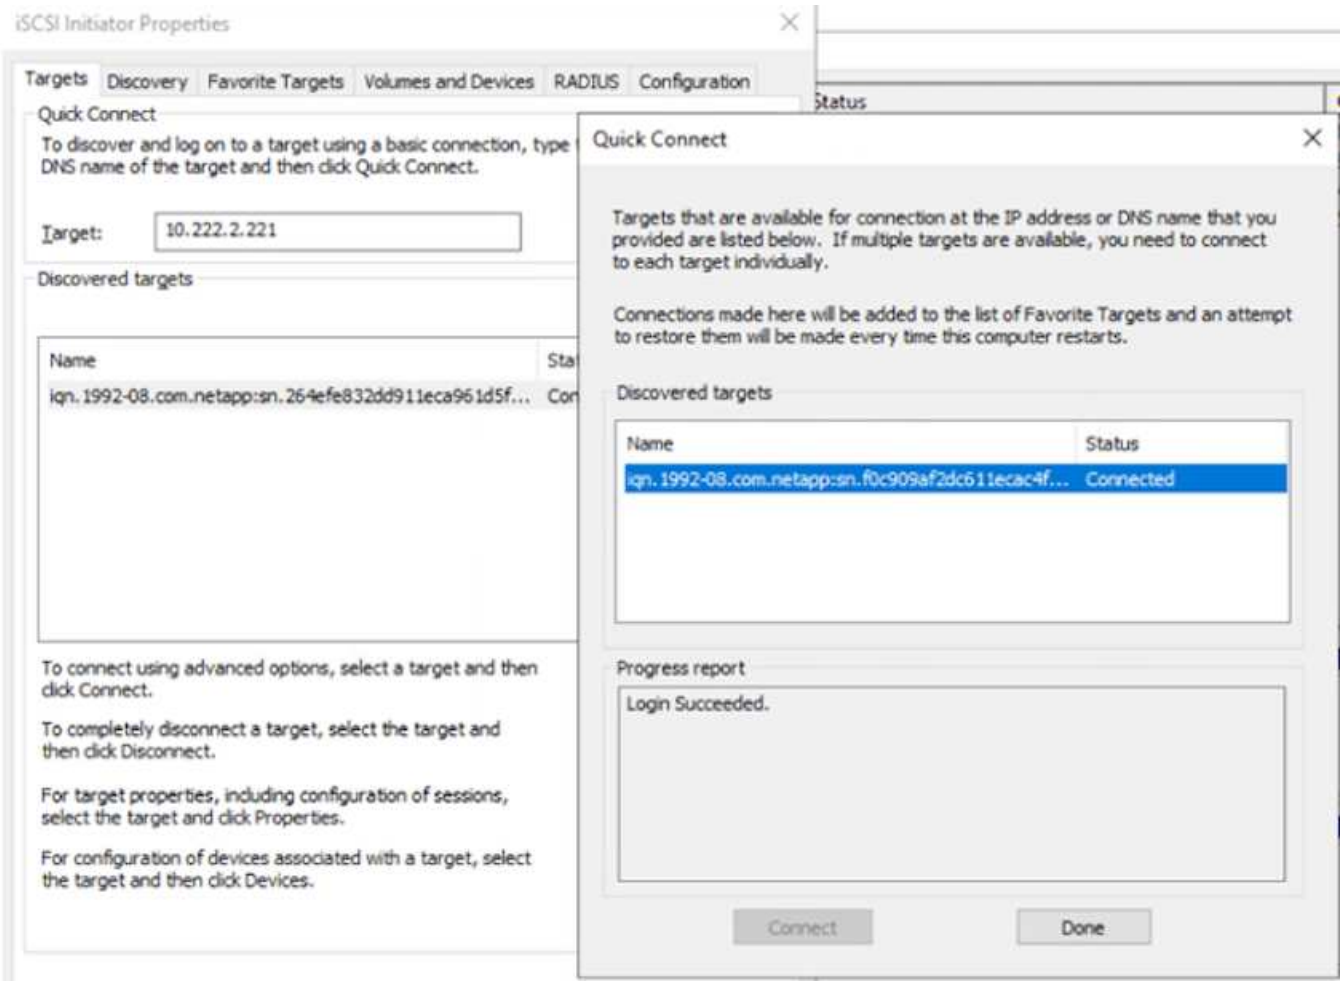
b. Navigate to Server Manager > Dashboard > Tools > iSCSI Initiator to open the iSCSI Initiator Properties dialog box.

c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.

d. From the Targets tab, select the target discovered and then click Log On or Connect.

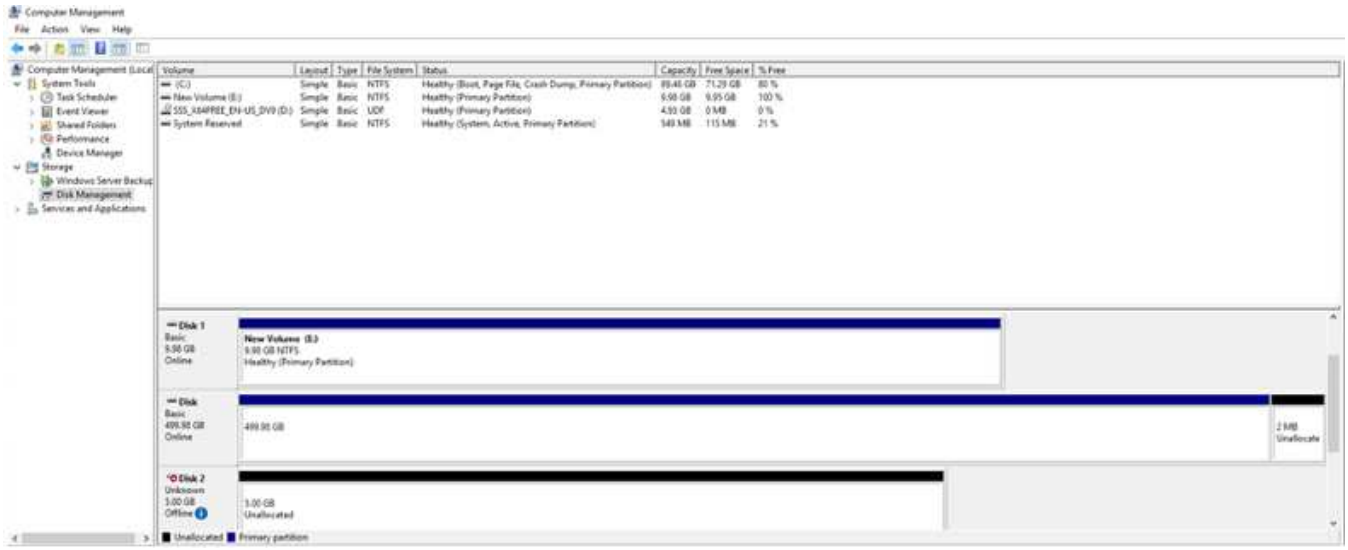
e. Select Enable Multipath, and then select "Automatically Restore This Connection When the Computer Starts" or "Add This Connection to the List of Favorite Targets". Click Advanced.

The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



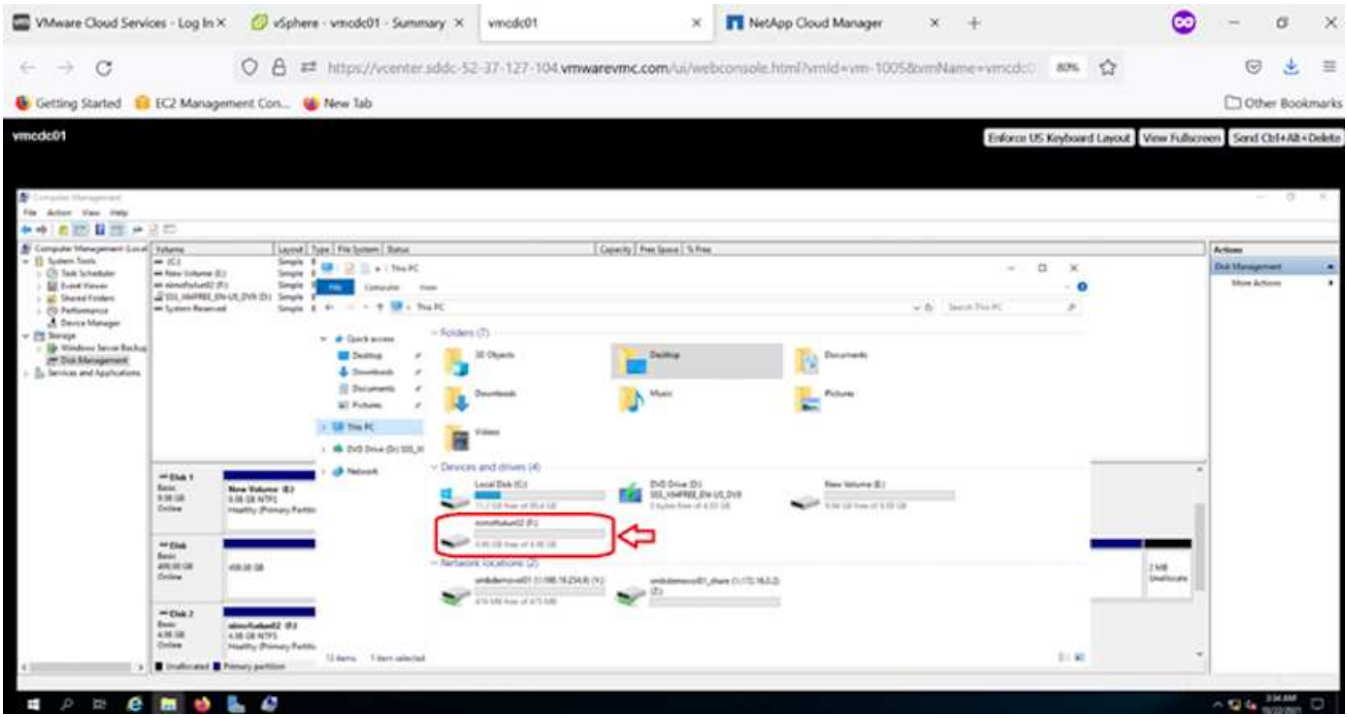
LUNs on the storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
2. Expand the Storage node in the navigation tree.
3. Click Disk Management.
4. Click Action > Rescan Disks.



When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN and, optionally, format the LUN with a file system by completing the following steps:

1. Start Windows Disk Management.
2. Right-click the LUN, and then select the required disk or partition type.
3. Follow the instructions in the wizard. In this example, drive F: is mounted.



Cloud Volumes ONTAP (CVO)

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to

manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

Cloud Volumes ONTAP (CVO) as guest connected storage

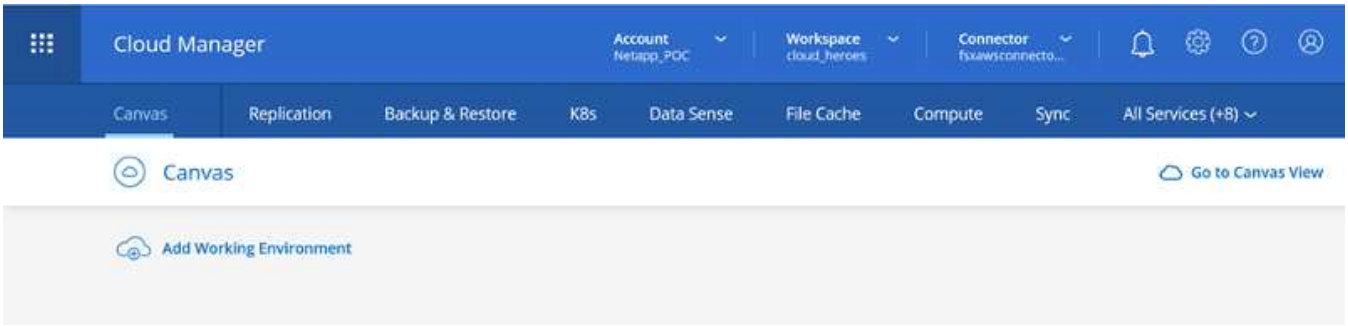
▼ **Deploy new Cloud Volumes ONTAP instance in AWS (do it yourself)**

Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the VMware Cloud on AWS SDDC environment. The volumes can also be mounted on native AWS VM Linux Windows clients, and LUNS can be accessed on Linux or Windows clients as block devices when mounted over iSCSI because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be set up in a few simple steps.

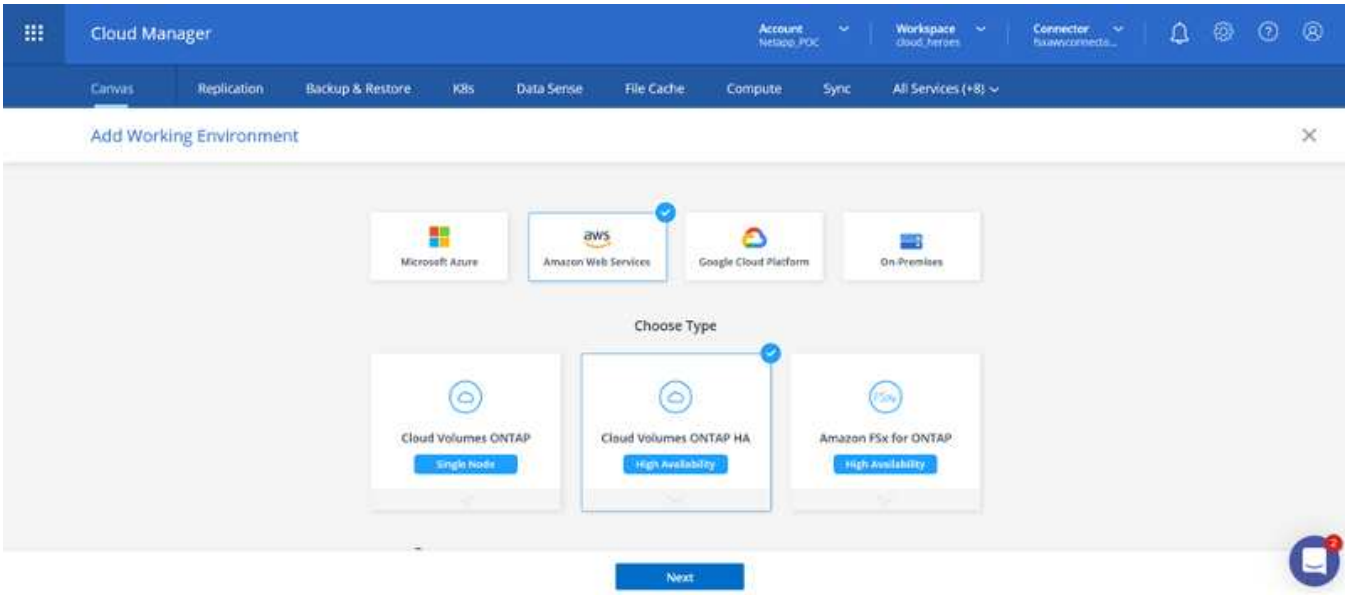
To replicate volumes from an on-premises environment to the cloud for disaster recovery or migration purposes, establish network connectivity to AWS, either using a site-to-site VPN or DirectConnect. Replicating data from on-premises to Cloud Volumes ONTAP is outside the scope of this document. To replicate data between on-premises and Cloud Volumes ONTAP systems, see Setting up data replication between systems.

Use the Cloud Volumes ONTAP sizer to accurately size the Cloud Volumes ONTAP instances. Also, monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

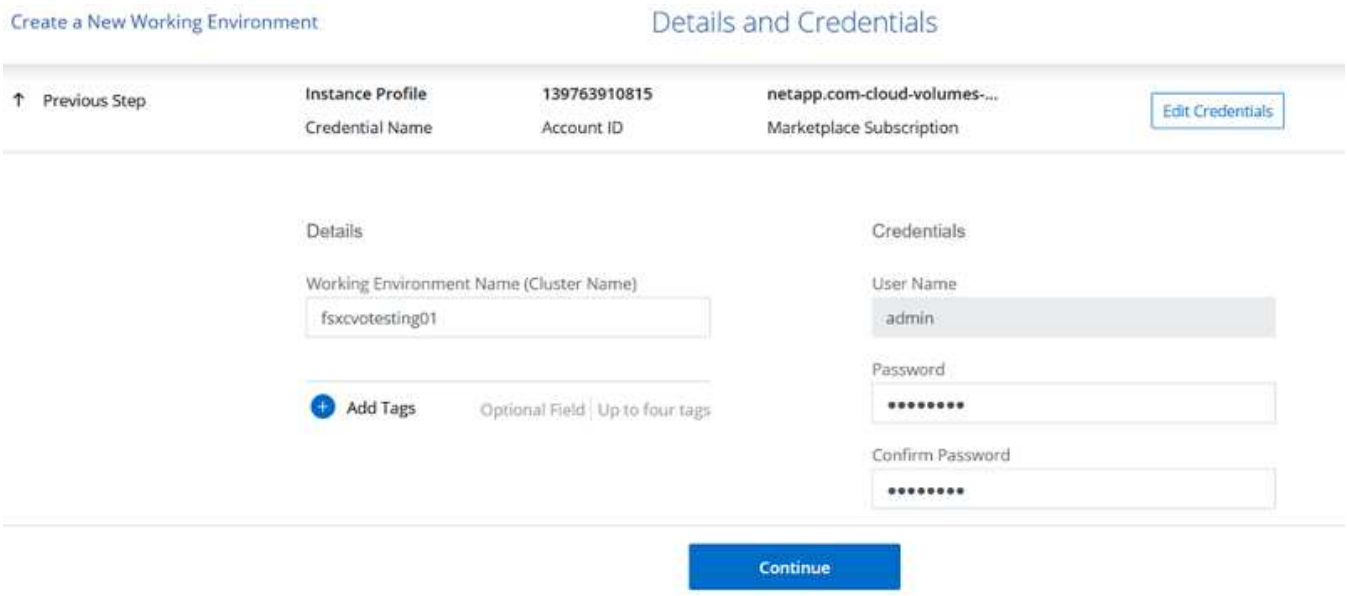
1. Log into NetApp Cloud Central; the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.



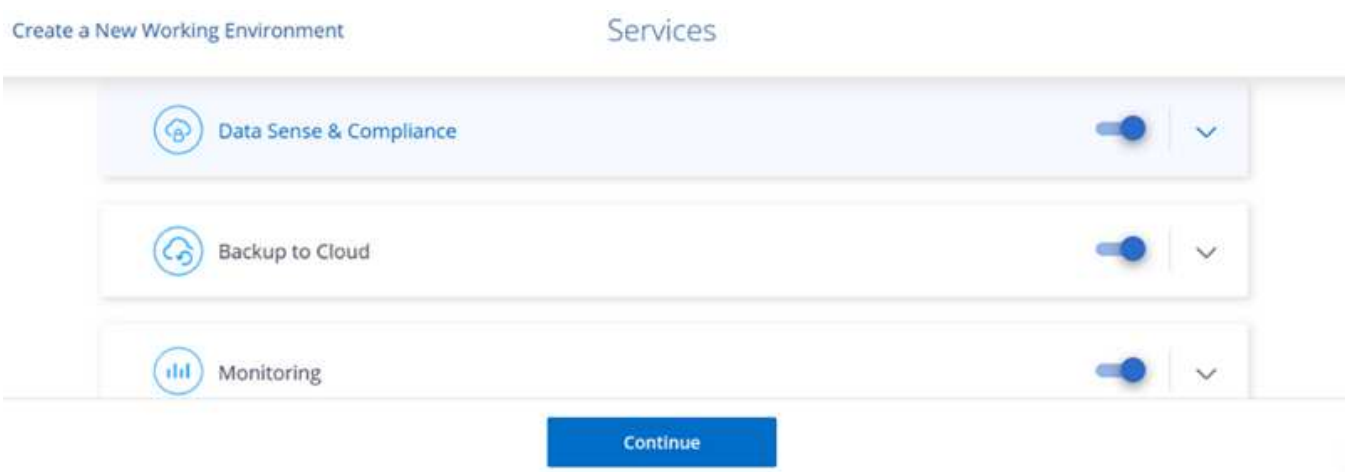
1. On the Cloud Manager home page, click Add a Working Environment and then select AWS as the cloud and the type of the system configuration.



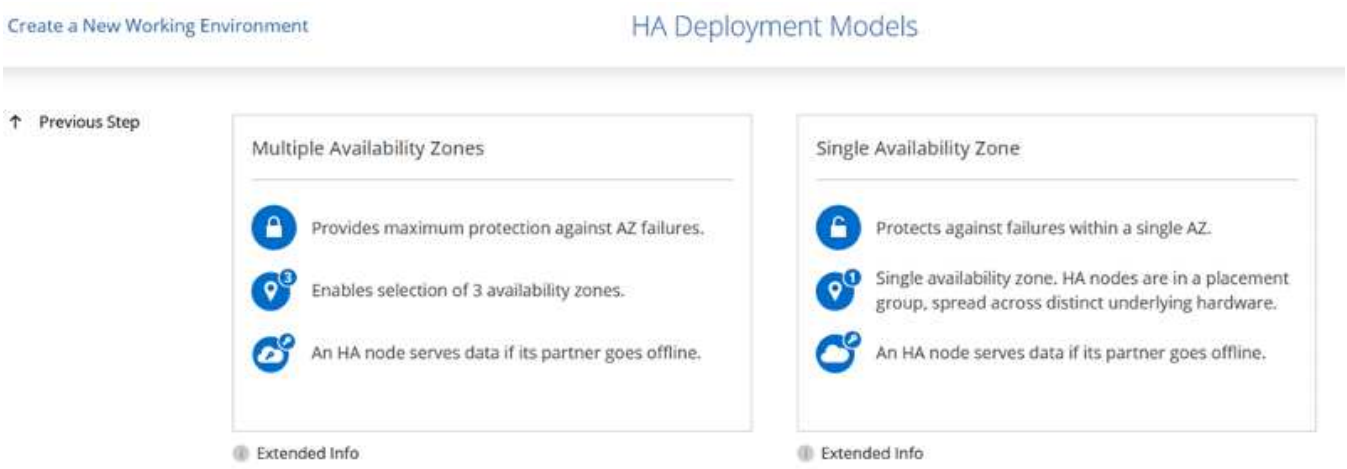
1. Provide the details of the environment to be created including the environment name and admin credentials. Click Continue.



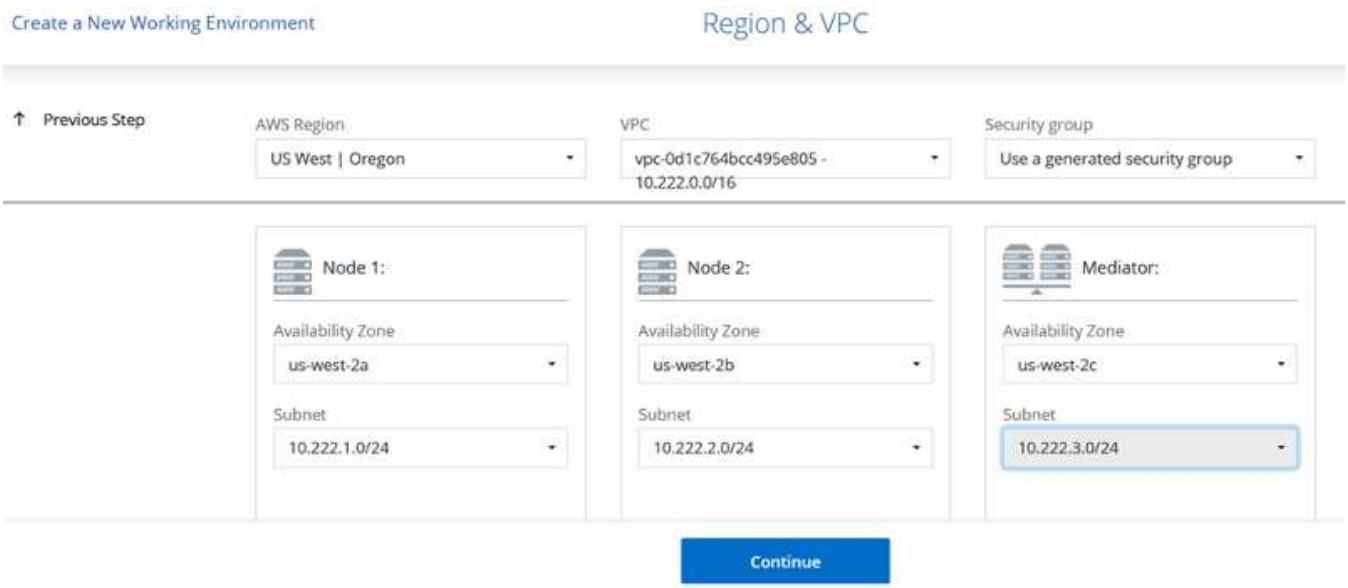
1. Select the add-on services for Cloud Volumes ONTAP deployment, including NetApp Data Classification, NetApp Backup and Recovery, and Data Infrastructure Insights. Click Continue.



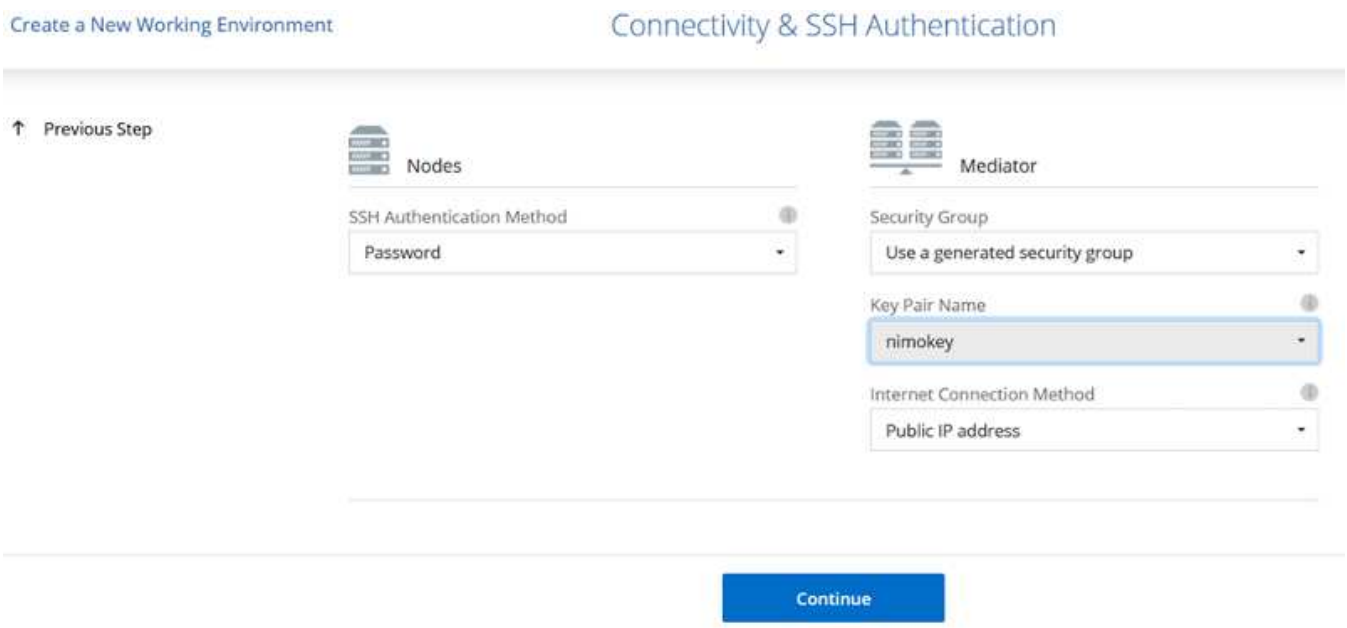
1. On the HA Deployment Models page, choose the Multiple Availability Zones configuration.



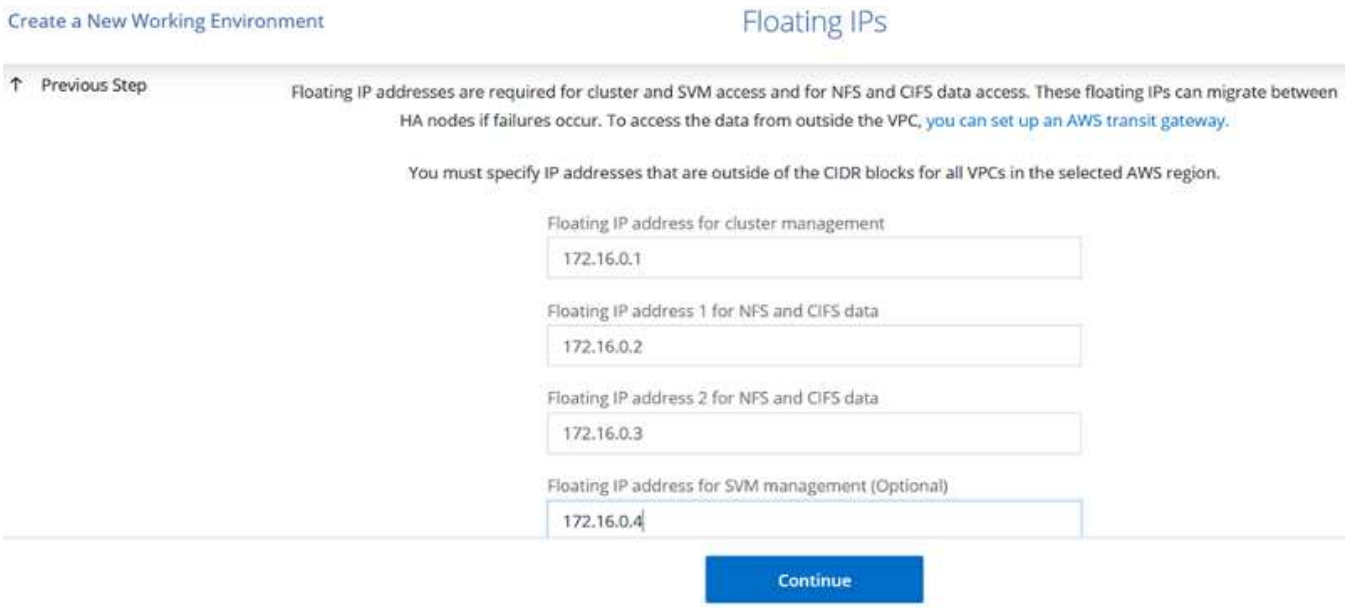
1. On the Region & VPC page, enter the network information and then click Continue.



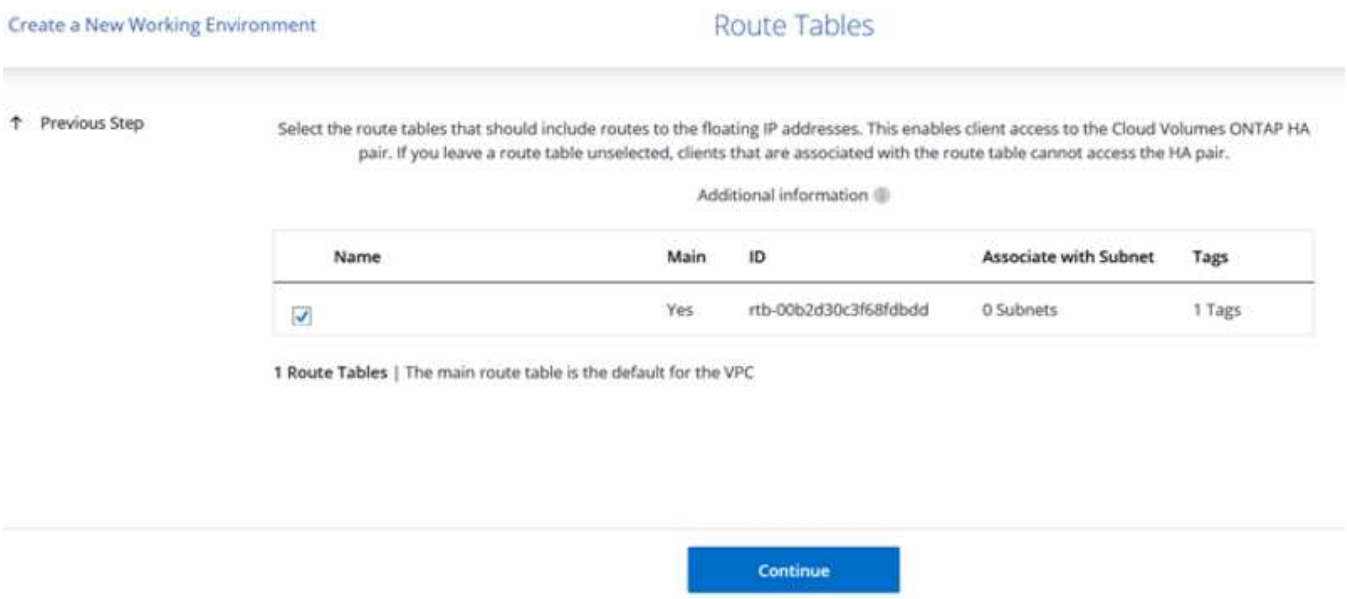
1. On the Connectivity and SSH Authentication page, choose connection methods for the HA pair and the mediator.



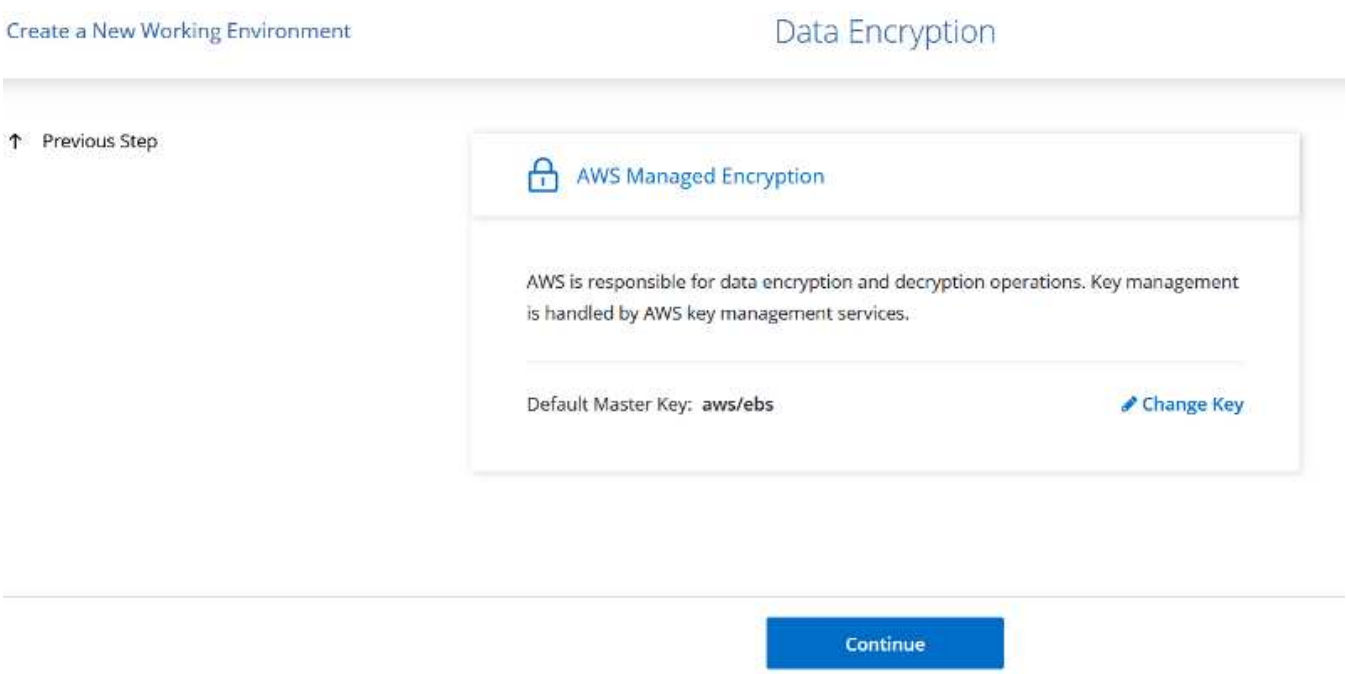
1. Specify the floating IP addresses and then click Continue.



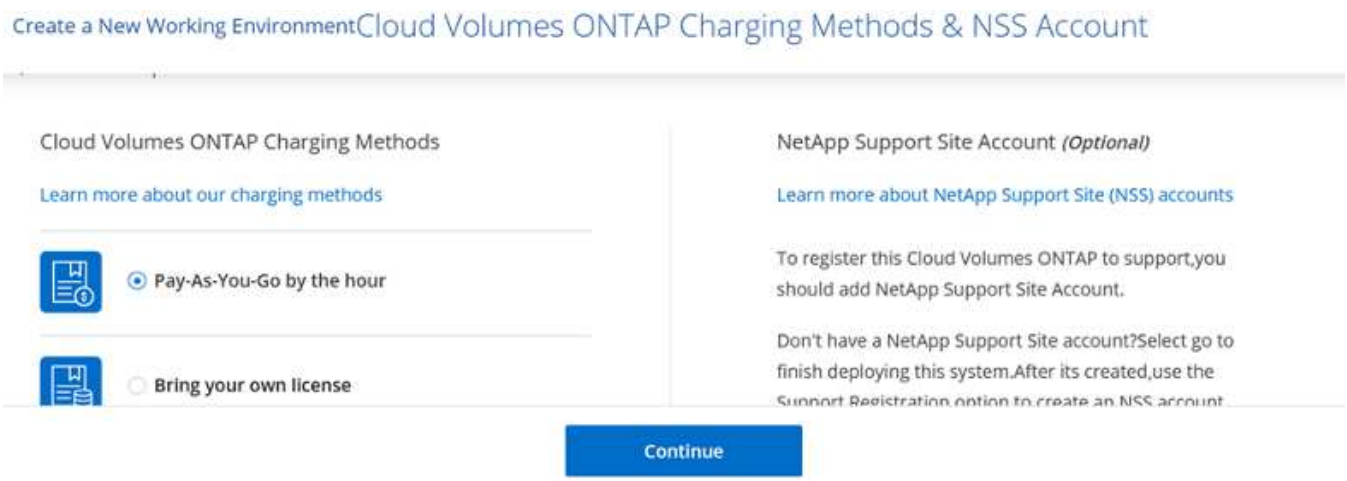
1. Select the appropriate route tables to include routes to the floating IP addresses and then click Continue.



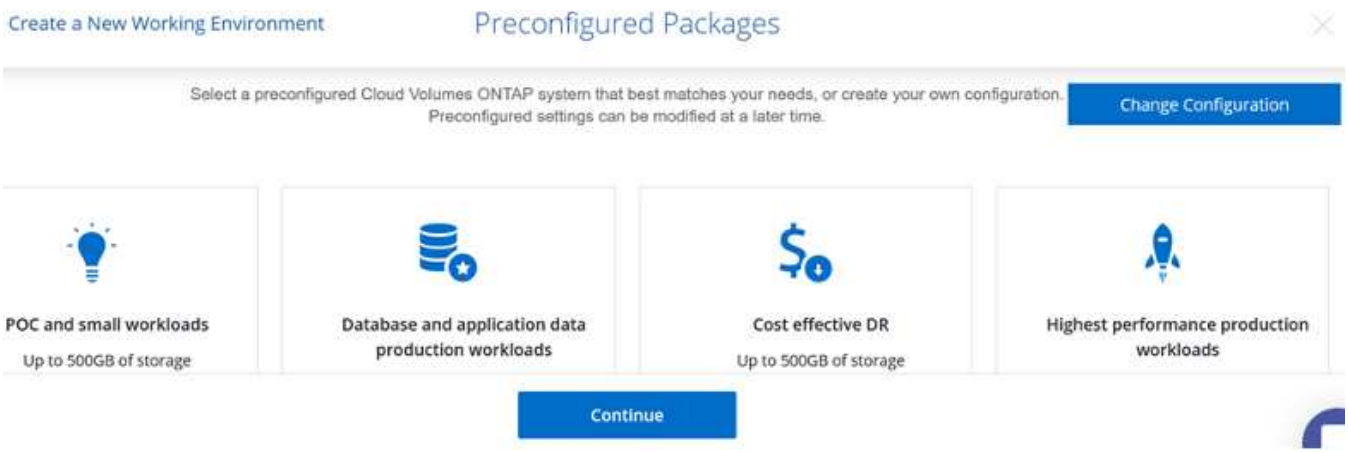
1. On the Data Encryption page, choose AWS-managed encryption.



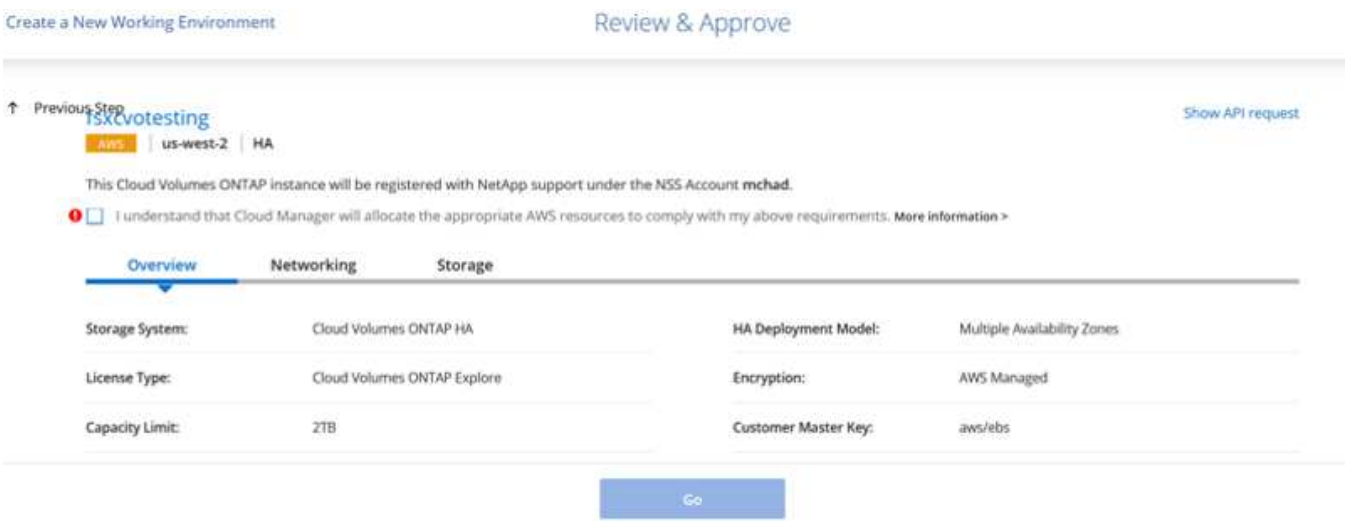
1. Select the license option: Pay-As-You-Go or BYOL for using an existing license. In this example, the Pay-As-You-Go option is used.



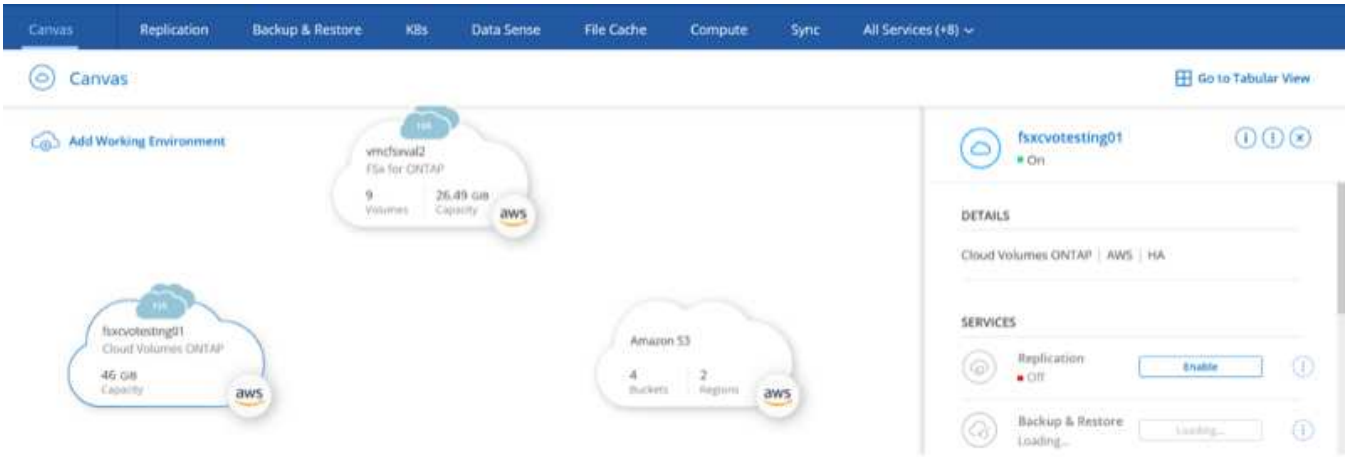
1. Select between several preconfigured packages available based on the type of workload to be deployed on the VMs running on the VMware cloud on AWS SDDC.



1. On the Review & Approve page, review and confirm the selections.To create the Cloud Volumes ONTAP instance, click Go.

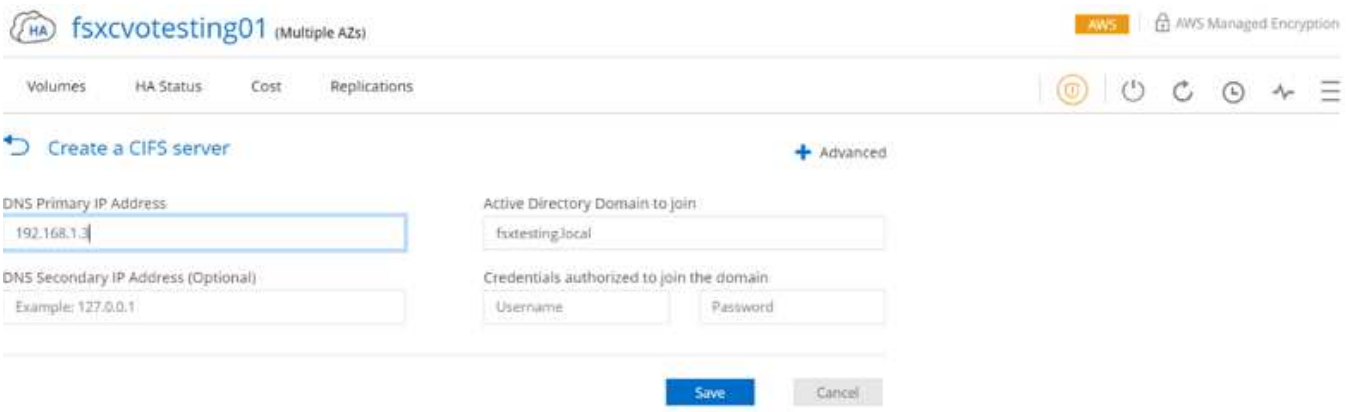


1. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.

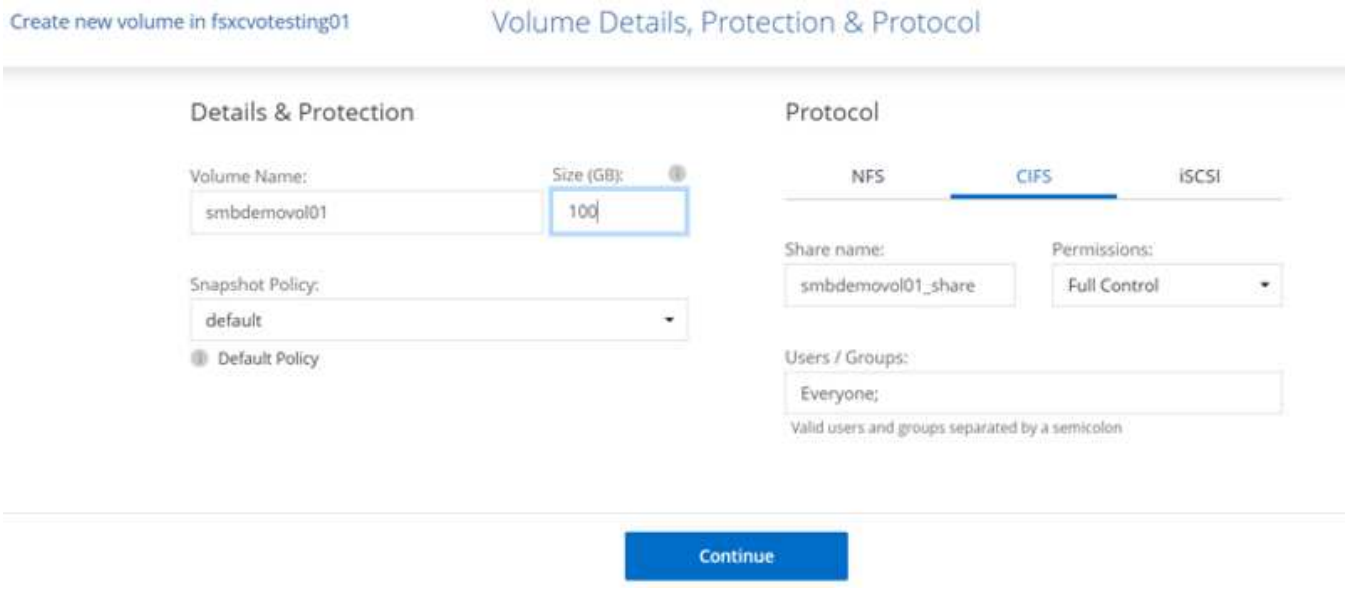


▼ **Additional configurations for SMB volumes**

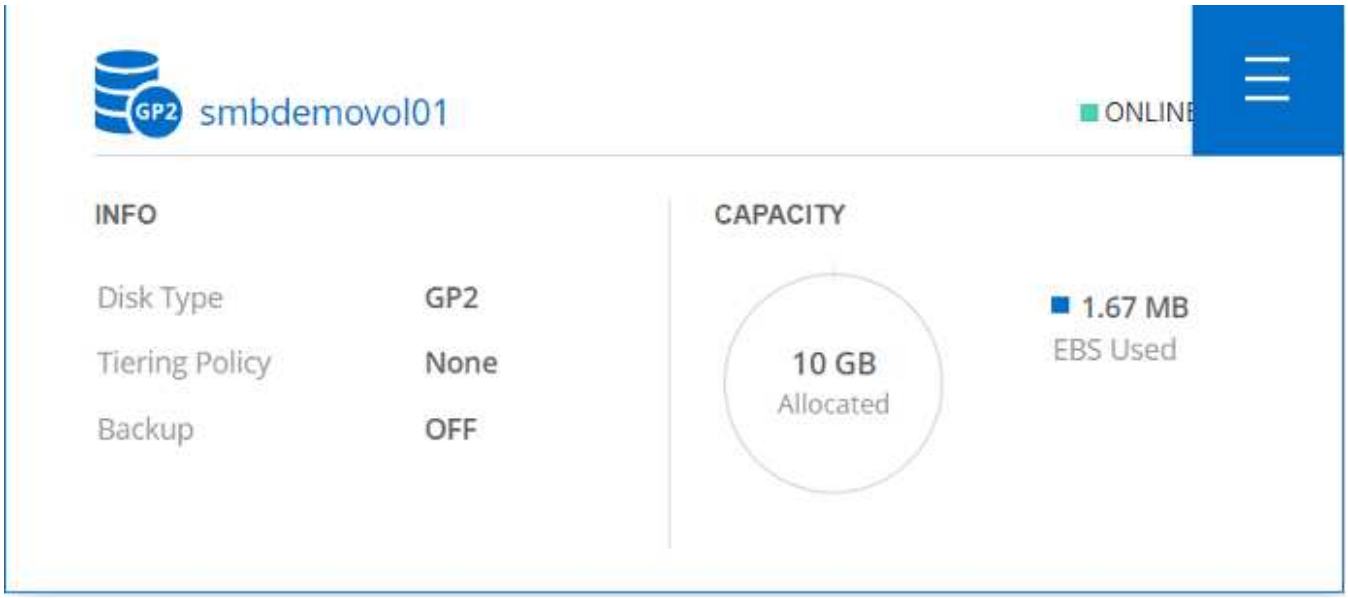
1. After the working environment is ready, make sure the CIFS server is configured with the appropriate DNS and Active Directory configuration parameters. This step is required before you can create the SMB volume.



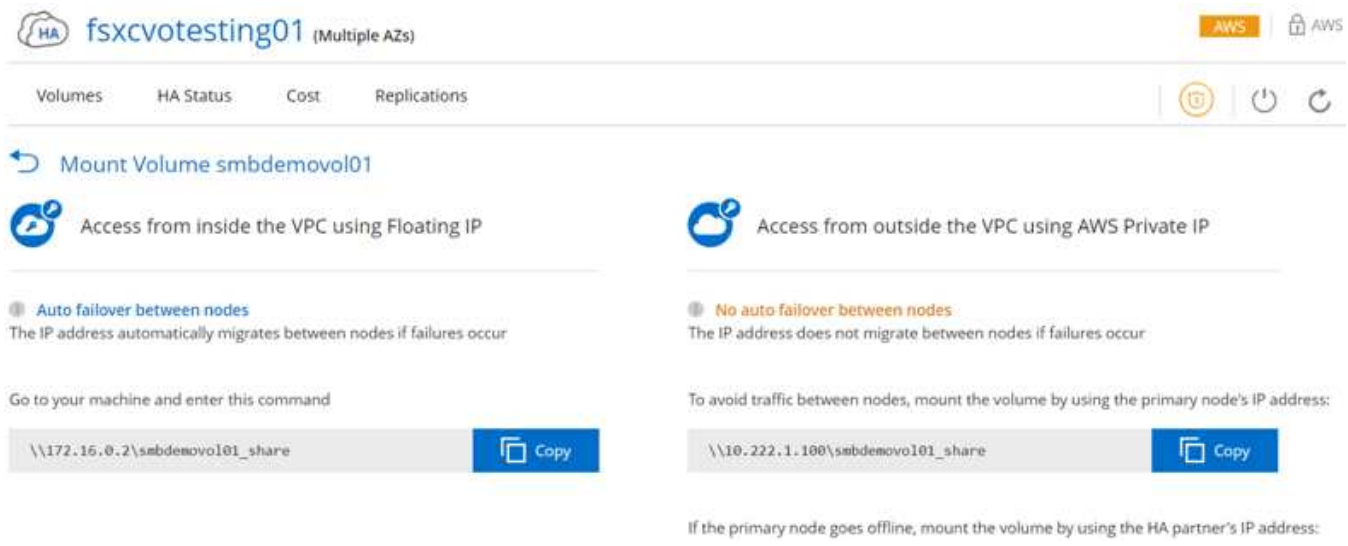
1. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected as the protocol.

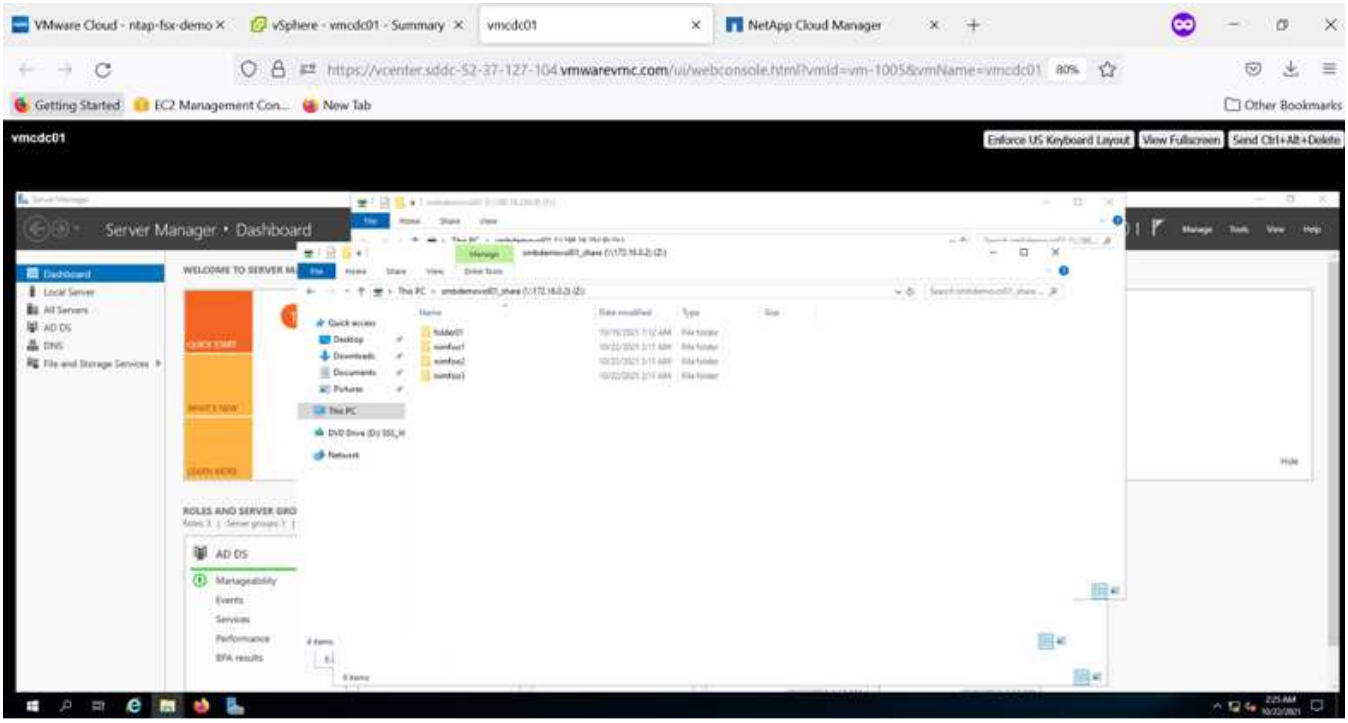


1. After the volume is provisioned, it is availabe under the Volumes pane. Because a CIFS share is provisioned, you should give your users or groups permission to the files and folders and verify that those users can access the share and create a file.



1. After the volume is created, use the mount command to connect to the share from the VM running on the VMware Cloud in AWS SDDC hosts.
2. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on the VMware Cloud in AWS SDDC.

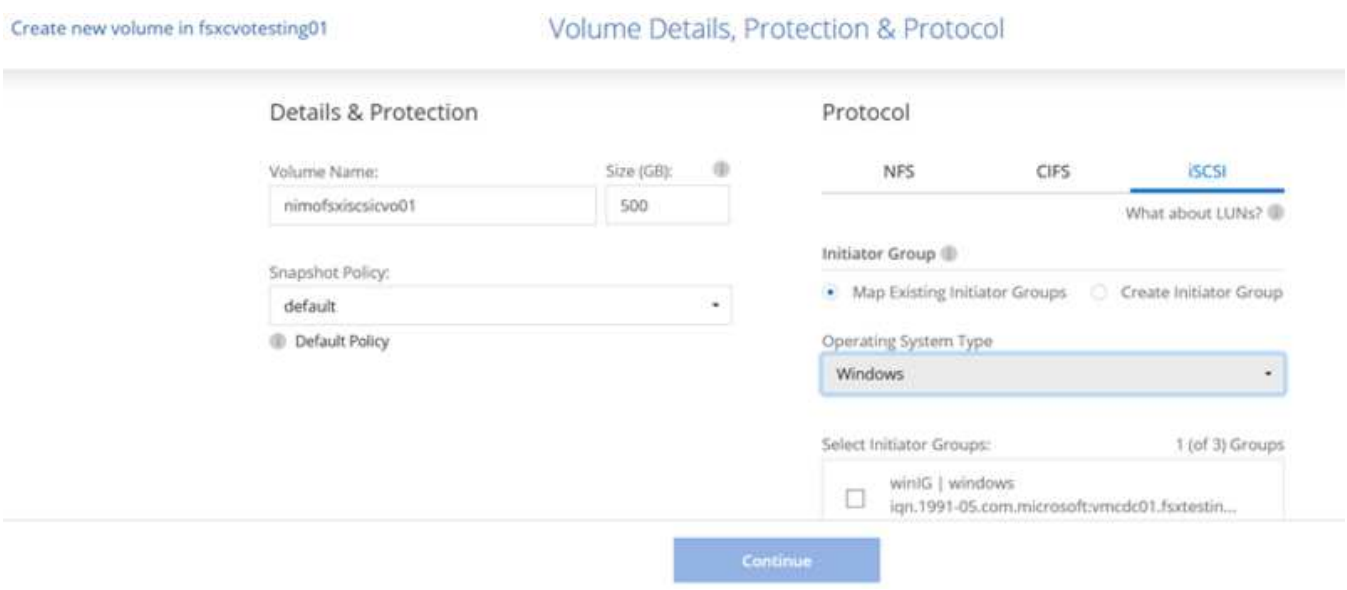


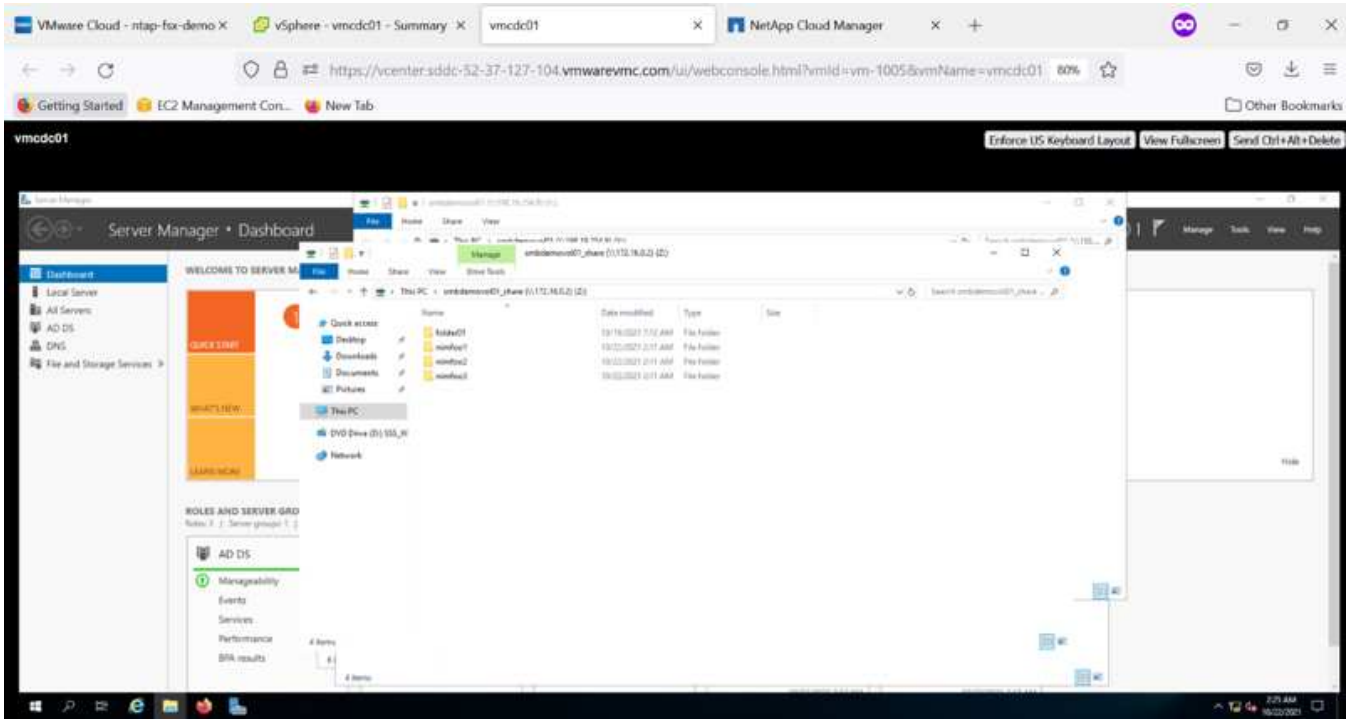


▼ Connect the LUN to a host

To connect the Cloud Volumes ONTAP LUN to a host, complete the following steps:

1. On the Cloud Manager Canvas page, double-click the Cloud Volumes ONTAP working environment to create and manage volumes.
2. Click Add Volume > New Volume, select iSCSI, and click Create Initiator Group. Click Continue.





1. After the volume is provisioned, select the volume, and then click Target IQN. To copy the iSCSI Qualified Name (IQN), click Copy. Set up an iSCSI connection from the host to the LUN.

To accomplish the same for the host residing on the VMware Cloud on AWS SDDC, complete the following steps:

a. RDP to the VM hosted on VMware cloud on AWS.

b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.

c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.

d. From the Targets tab, select the target discovered and then click Log On or Connect.

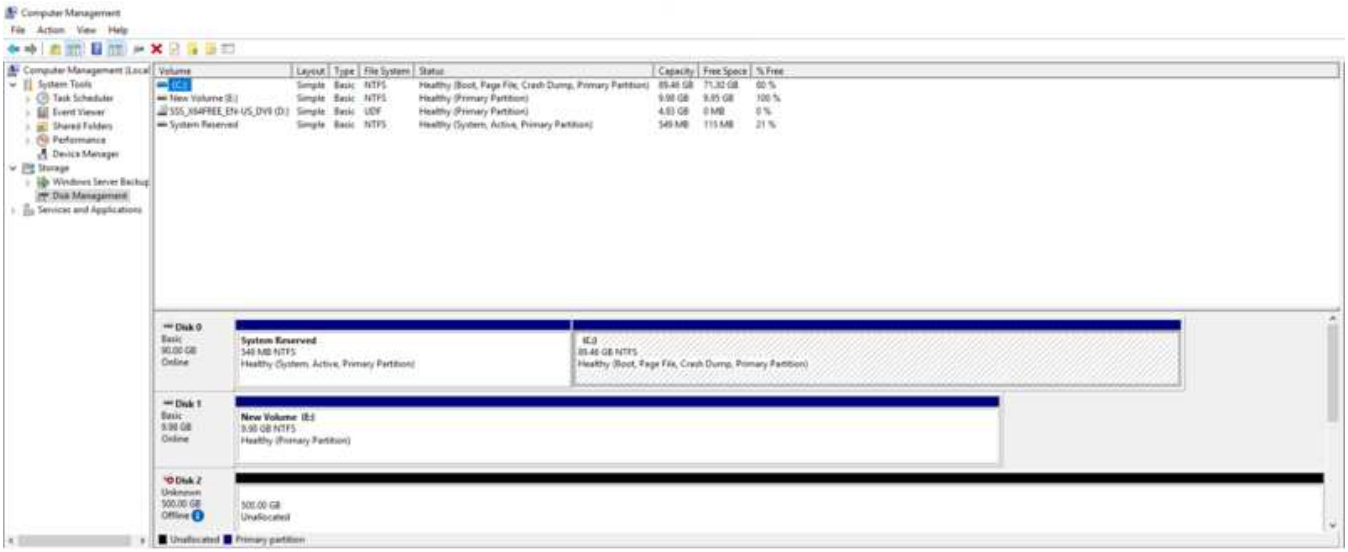
e. Select Enable Multipath, and then select Automatically Restore This Connection When the Computer Starts or Add This Connection to the List of Favorite Targets. Click Advanced.

The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



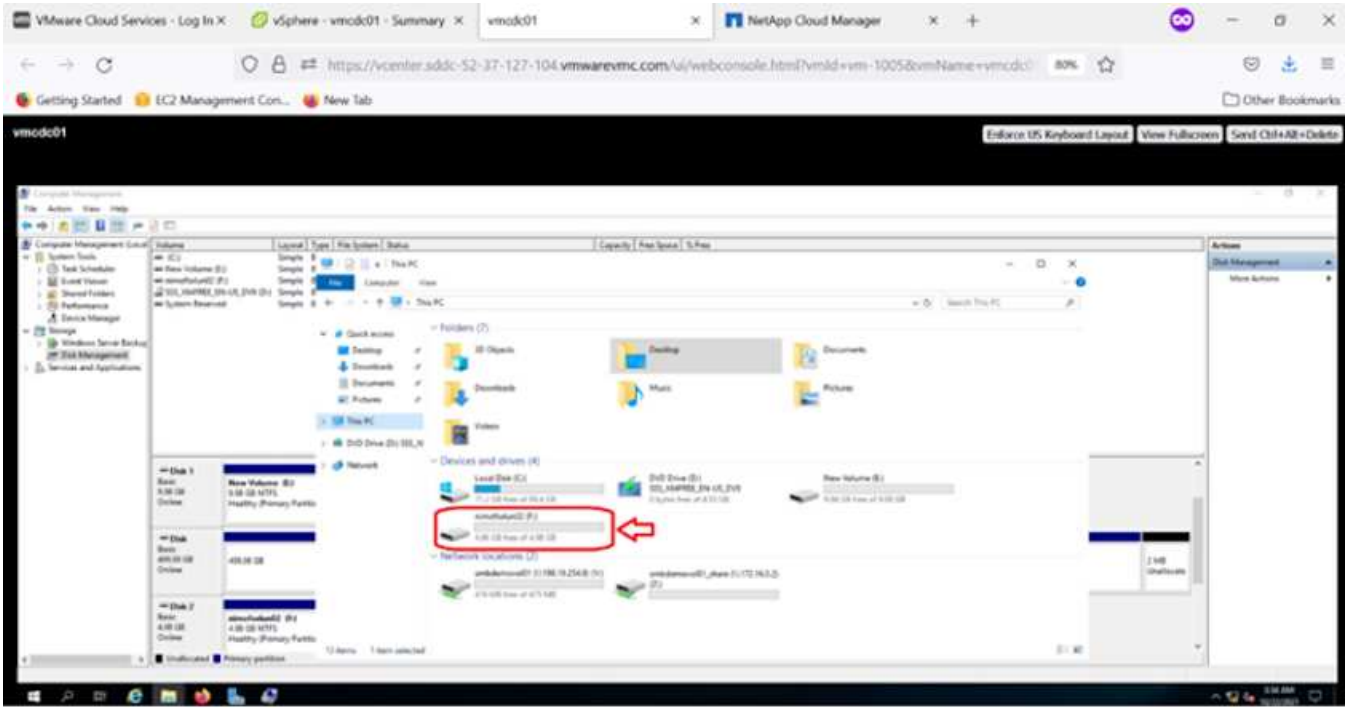
LUNs from the SVM appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
2. Expand the Storage node in the navigation tree.
3. Click Disk Management.
4. Click Action > Rescan Disks.



When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

1. Start Windows Disk Management.
2. Right-click the LUN, and then select the required disk or partition type.
3. Follow the instructions in the wizard. In this example, drive F: is mounted.



On the Linux clients, ensure the iSCSI daemon is running. After the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration for your Linux distribution. For example, Ubuntu iSCSI configuration can be found here. To verify, run lsblk cmd from the shell.

▼ **Mount Cloud Volumes ONTAP NFS volume on Linux client**

To mount the Cloud Volumes ONTAP (DIY) file system from VMs within VMC on AWS SDDC, complete the following steps:

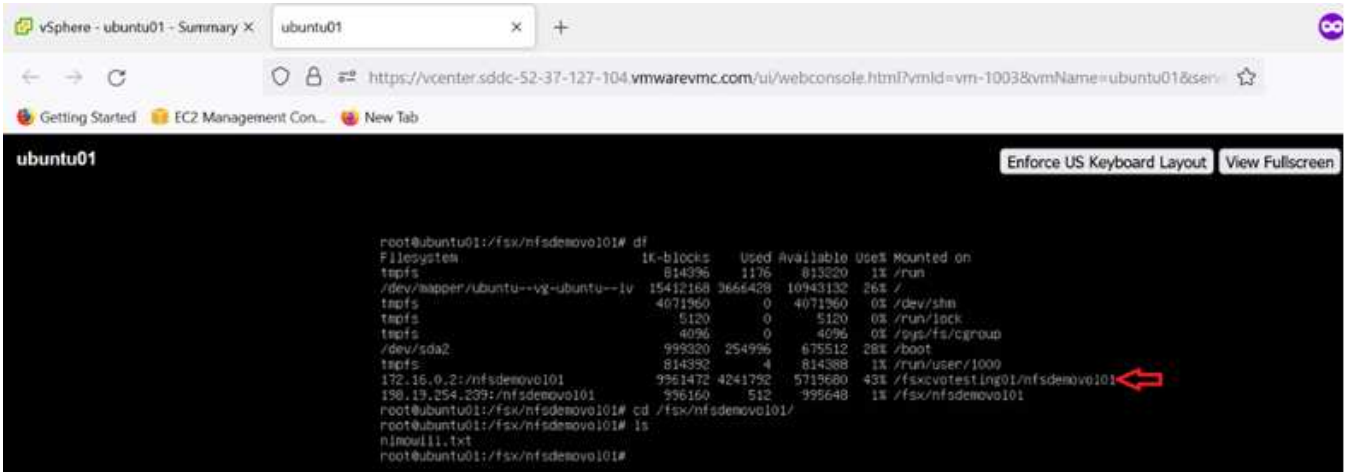
1. Connect to the designated Linux instance.
2. Open a terminal on the instance using secure shell (SSH) and log in with the appropriate credentials.
3. Make a directory for the volume's mount point with the following command.

```
$ sudo mkdir /fsxcvotesting01/nfsdemovol01
```

4. Mount the Amazon FSx ONTAP NFS volume to the directory that is created in the previous step.

```
sudo mount -t nfs nfsvers=4.1,172.16.0.2:/nfsdemovol01
/fsxcvotesting01/nfsdemovol01
```

```
root@ubuntu01:/fsx# mount -t nfs 172.16.0.2:/nfsdemovol01 /fsxcvotesting01/nfsdemovol01_
```



Azure Virtualization Service: Options for using NetApp storage

NetApp storage can be attached to the Azure VMware Service as either guest connected or supplemental storage.

Azure NetApp Files (ANF) as a supplemental NFS datastore

NFS datastore support was introduced with ESXi version 3 in on-premises deployments, which greatly extended vSphere's storage capabilities.

Running vSphere on NFS is a widely adopted option for virtualization deployments on-premises because it offers strong performance and stability. If you have significant network-attached storage (NAS) in an on-premises datacenter, you should consider deploying an Azure VMware solution SDDC in Azure with Azure NetApp File datastores to overcome capacity and performance challenges.

Azure NetApp Files is built on industry-leading, highly available NetApp ONTAP data management software. Microsoft Azure services are grouped into three categories: foundational, mainstream, and specialized. Azure NetApp Files is in the specialized category and is backed by hardware already deployed in many regions. With built-in high-availability (HA), Azure NetApp Files protects your data from most outages and offers you an industry-leading SLA of 99.99%^ uptime.

For more information on ANF as a supplemental NFS datastore, please visit:

- ANF as a Supplemental NFS Datastore: Overview
- Supplemental NFS Datastore Option in Azure

Azure NetApp Files (ANF) as guest connected storage

Azure NetApp Files brings enterprise-grade data management and storage to Azure so you can manage your workloads and applications with ease. Migrate your workloads to the cloud and run them without sacrificing performance.

Azure NetApp Files removes obstacles, so you can move all of your file-based applications to the cloud. For the first time, you do not have to re-architect your applications, and you get persistent storage for your applications without complexity.

Because the service is delivered through the Microsoft Azure Portal, users experience a fully managed service as part of their Microsoft enterprise Agreement. World-class support, managed by Microsoft, gives you complete peace of mind. This single solution enables you to quickly and easily add multiprotocol workloads.

you can build and deploy both Windows and Linux file-based applications, even for legacy environments.

For more information, please visit ANF as Guest Connected Storage.

Cloud Volumes ONTAP (CVO) as guest connected storage

Cloud Volumes ONTAP (CVO) is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

For more information, please visit CVO as Guest Connected Storage.

Overview of ANF Datastore Solutions

Every successful organization is on a path of transformation and modernization. As part of this process, companies typically use their existing VMware investments while leveraging cloud benefits and exploring how to make migration, burst, extend, and disaster recovery processes as seamless as possible. Customers migrating to the cloud must evaluate the issues of elasticity and burst, data center exit, data center consolidation, end- of- life scenarios, mergers, acquisitions, and so on. The approach adopted by each organization can vary based on their respective business priorities. When choosing cloud-based operations, selecting a low- cost model with appropriate performance and minimal hindrance is a critical goal. Along with choosing the right platform, storage and workflow orchestration is particularly important to unleash the power of cloud deployment and elasticity.

Use Cases

Although the Azure VMware solution delivers unique hybrid capabilities to a customer, limited native storage options have restricted its usefulness for organizations with storage-heavy workloads. Because storage is directly tied to hosts, the only way to scale storage is to add more hosts, which can increase costs by 35-40% or more for storage intensive workloads. These workloads need additional storage, not additional horsepower, but that means paying for additional hosts.

Let's consider the following scenario; a customer requires six hosts for horsepower (vCPU/vMem), but they also have a substantial requirement for storage. Based on their assessment, they require 12 hosts to meet storage requirements. This increases the overall TCO because they must buy all that additional horsepower when all they really need is more storage. This is applicable for any use case, including migration, disaster recovery, bursting, dev/test, and so on.

Another common use case for Azure VMware Solution is disaster recovery (DR). Most organizations do not

have a fool- proof DR strategy, or they might struggle to justify running a ghost datacenter just for DR. Administrators might explore zero- footprint DR options with a pilot- light cluster or an on-demand cluster. They could then scale the storage without adding additional hosts, potentially an attractive option.

So, to summarize, the use cases can be classified in two ways:

- Scaling storage capacity using ANF datastores
- Using ANF datastores as a disaster recovery target for a cost- optimized recovery workflow from on-premises or within Azure regions between the software-defined datacenters (SDDCs).This guide provides insight into using Azure NetApp Files to provide optimized storage for datastores (currently in public preview) along with best-in-class data protection and DR capabilities in an Azure VMware solution, which enables you to offload storage capacity from vSAN storage.



Contact NetApp or Microsoft solution architects in your region for additional information on using ANF datastores.

VMware Cloud options in Azure

Azure VMware Solution

The Azure VMware Solution (AVS) is a hybrid cloud service that provides fully functioning VMware SDDCs within a Microsoft Azure public cloud. AVS is a first-party solution fully managed and supported by Microsoft and verified by VMware that uses Azure infrastructure. Therefore, customers get VMware ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data center facilities, and proximity to the rich ecosystem of native Azure services and solutions. A combination of Azure VMware Solution SDDC and Azure NetApp Files provides the best performance with minimal network latency.

Regardless of the cloud used, when a VMware SDDC is deployed, the initial cluster includes the following components:

- VMware ESXi hosts for compute virtualization with a vCenter server appliance for management.
- VMware vSAN hyper-converged storage incorporating the physical storage assets of each ESXi host.
- VMware NSX for virtual networking and security with an NSX Manager cluster for management.

Conclusion

Whether you are targeting all-cloud or hybrid cloud, Azure NetApp files provide excellent options to deploy and manage the application workloads along with file services while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Azure VMware Solution along with Azure NetApp Files for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bi-directional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage. Remember, it is just the position of the data that changed along with new names; the tools and processes all remain the same, and Azure NetApp Files helps in optimizing the overall deployment.

Takeaways

The key points of this document include:

- You can now use Azure NetApp Files as a datastore on AVS SDDC.
- Boost the application response times and deliver higher availability to provide access workload data when and where it is needed.

- Simplify the overall complexity of the vSAN storage with simple and instant resizing capabilities.
- Guaranteed performance for mission-critical workloads using dynamic reshaping capabilities.
- If Azure VMware Solution Cloud is the destination, Azure NetApp Files is the right storage solution for optimized deployment.

Where to find additional information

To learn more about the information described in this document, refer to the following website links:

- Azure VMware Solution documentation

  https://docs.microsoft.com/en-us/azure/azure-vmware/

- Azure NetApp Files documentation

  https://docs.microsoft.com/en-us/azure/azure-netapp-files/

- Attach Azure NetApp Files datastores to Azure VMware Solution hosts (Preview)

  https://docs.microsoft.com/en-us/azure/azure-vmware/attach-azure-netapp-files-to-azure-vmware-solution-hosts?tabs=azure-portal/

Creating a Supplemental NFS Datastore in Azure

NFS datastore support was introduced with ESXi version 3 in on-premises deployments, which greatly extended vSphere's storage capabilities.

Running vSphere on NFS is a widely adopted option for virtualization deployments on-premises because it offers strong performance and stability. If you have significant network-attached storage (NAS) in an on-premises datacenter, you should consider deploying an Azure VMware solution SDDC in Azure with Azure NetApp File datastores to overcome capacity and performance challenges.

Azure NetApp Files is built on industry-leading, highly available NetApp ONTAP data management software. Microsoft Azure services are grouped into three categories: foundational, mainstream, and specialized. Azure NetApp Files is in the specialized category and is backed by hardware already deployed in many regions. With built-in high-availability (HA), Azure NetApp Files protects your data from most outages and offers you an industry-leading SLA of 99.99% uptime.

Before the introduction of Azure NetApp Files datastore capability, scale-out operation for customers planning to host performance and storage- intensive workloads required the expansion of both compute and storage.

Keep in mind the following issues:

- Unbalanced cluster configurations are not recommended in an SDDC cluster. Therefore expanding storage means adding more hosts, which implies more TCO.
- Only one vSAN environment is possible. Therefore, all storage traffic competes directly with production workloads.
- There is no option for providing multiple performance tiers to align application requirements, performance, and cost.
- It is easy to reach the limits of storage capacity for vSAN built on top of cluster hosts.By integrating Azure-native platform-as-a-service (PaaS) offerings like Azure NetApp Files as a datastore, customers have the option to independently scale their storage separately and only add compute nodes to the SDDC cluster as

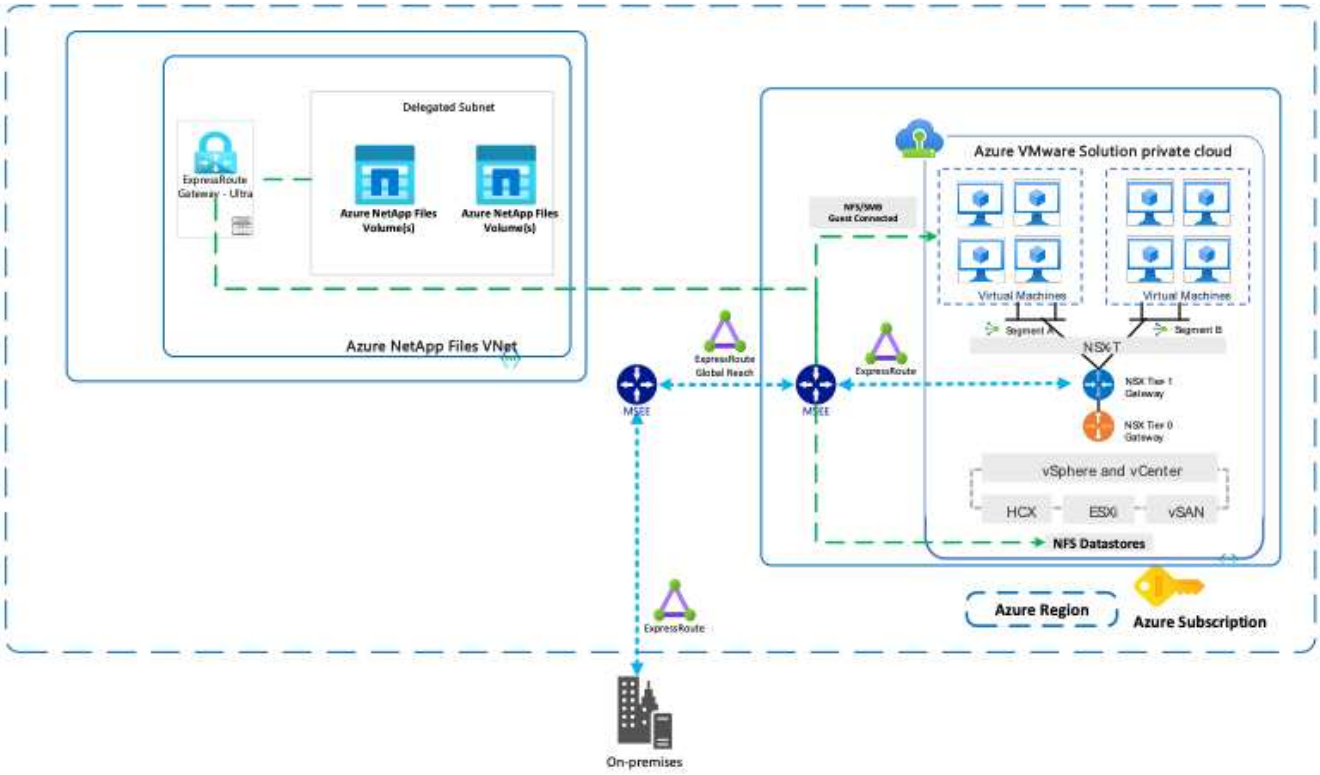
needed. This capability overcomes the challenges mentioned above.

Azure NetApp Files also enables you to deploy multiple datastores, which helps to mimic an on-premises deployment model by placing virtual machines in the appropriate datastore and by assigning the required service level to meet workload performance requirements. With the unique capability of multi-protocol support, guest storage is an additional option for database workloads like SQL and Oracle while also using supplemental NFS datastore capability to house the remaining VMDKs. Apart from this, native snapshot capability enables you to perform quick backup and granular restores.

Contact Azure and NetApp Solution Architects for planning and sizing of storage and determining the required number of hosts. NetApp recommends identifying storage performance requirements before finalizing the datastore layout for test, POC, and production deployments.

Detailed architecture

From a high-level perspective, this architecture describes how to achieve hybrid-cloud connectivity and app portability across on-premises environments and Azure. It also describes using Azure NetApp Files as a supplemental NFS datastore and as an in-guest storage option for guest virtual machines hosted on the Azure VMware solution.



Sizing

The most important aspect in migration or disaster recovery is to determine the correct size for the target environment. It is very important to understand how many nodes are required to accommodate a lift- and- shift exercise from on-premises to the Azure VMware Solution.

For sizing, use historic data from the on-premises environment using RVTools (preferred) or other tools like Live Optics or Azure Migrate. RVTools is an ideal tool to capture the vCPU, vMem, vDisk and all required information, including powered on or off VMs, to characterize the target environment.

To run RVtools, complete the following steps:

1. Download and install RVTools.
2. Run RVTools, enter the required info to connect to your on-premises vCenter Server, and press Login.
3. Export the inventory to an Excel spreadsheet.
4. Edit the spreadsheet and remove any VMs that are not ideal candidates from the vInfo tab.This approach provides a clear output about storage requirements that can be used to right- size the Azure VMware SDDC cluster with the required number of hosts.

> Guest VMs used with in-guest storage must be calculated separately; however, Azure NetApp Files can easily cover the additional storage capacity, thus keeping the overall TCO low.

Deploying and configuring Azure VMware Solution

Like on-premises, planning an Azure VMware solution is critical for a successful production- ready environment for creating virtual machines and migration.

This section describes how to set up and manage AVS for use in combination with Azure NetApp Files as a datastore with in-guest storage as well.

The setup process can be broken down into three parts:

- Register the resource provider and create a private cloud.
- Connect to a new or existing ExpressRoute virtual network gateway.
- Validate network connectivity and access the private cloud. Refer to this link for a step-by-step walkthrough of the Azure VMware solution SDDC provisioning process.

Configure Azure NetApp Files with Azure VMware Solution

The new integration between Azure NetApp Files enables you to create NFS datastores via the Azure VMware Solution resource provider APIs/CLI with Azure NetApp Files volumes and mount the datastores on your clusters of choice in a private cloud. Apart from housing the VM and App VMDKs, Azure NetApp file volumes can also be mounted from VMs that are created in the Azure VMware Solution SDDC environment. The volumes can be mounted on the Linux client and mapped on a Windows client, because Azure NetApp Files supports Server Message Block (SMB) and Network File System (NFS) protocols.

> For optimal performance, deploy the Azure NetApp Files in the same availability zone as the private cloud. Colocation with the Express route fastpath provides the best performance, with minimal network latency.

To attach an Azure NetApp File volume as the VMware datastore of an Azure VMware Solution private cloud, make sure the following prerequisites are met.

▼ Prerequisites

1. Use az login and validate the subscription is registered to CloudSanExperience feature in the Microsoft.AVS namespace.

```
az login –tenant xcvxcvxc- vxcv- xcvx- cvxc- vxcvxcvxcv
az feature show --name "CloudSanExperience" --namespace "Microsoft.AVS"
```

a. If it is not registered, then register it.

```
az feature register --name "CloudSanExperience" --namespace
"Microsoft.AVS"
```

Registration can take approximately 15 minutes to complete.

a. To check the status of registration, run the following command.

```
az feature show --name "CloudSanExperience" --namespace "Microsoft.AVS"
--query properties.state
```

a. If the registration is stuck in an intermediate state for longer than 15 minutes, unregister and then reregister the flag.

```
az feature unregister --name "CloudSanExperience" --namespace
"Microsoft.AVS"
az feature register --name "CloudSanExperience" --namespace
"Microsoft.AVS"
```

a. Verify that the subscription is registered to the AnfDatastoreExperience feature in the Microsoft.AVS namespace.

```
az feature show --name "AnfDatastoreExperience" --namespace
"Microsoft.AVS" --query properties.state
```

1. Verify that the vmware extension is installed.

```
az extension show --name vmware
```

a. If the extension is already installed, verify that the version is 3.0.0. If an older version is installed, update the extension.

```
az extension update --name vmware
```

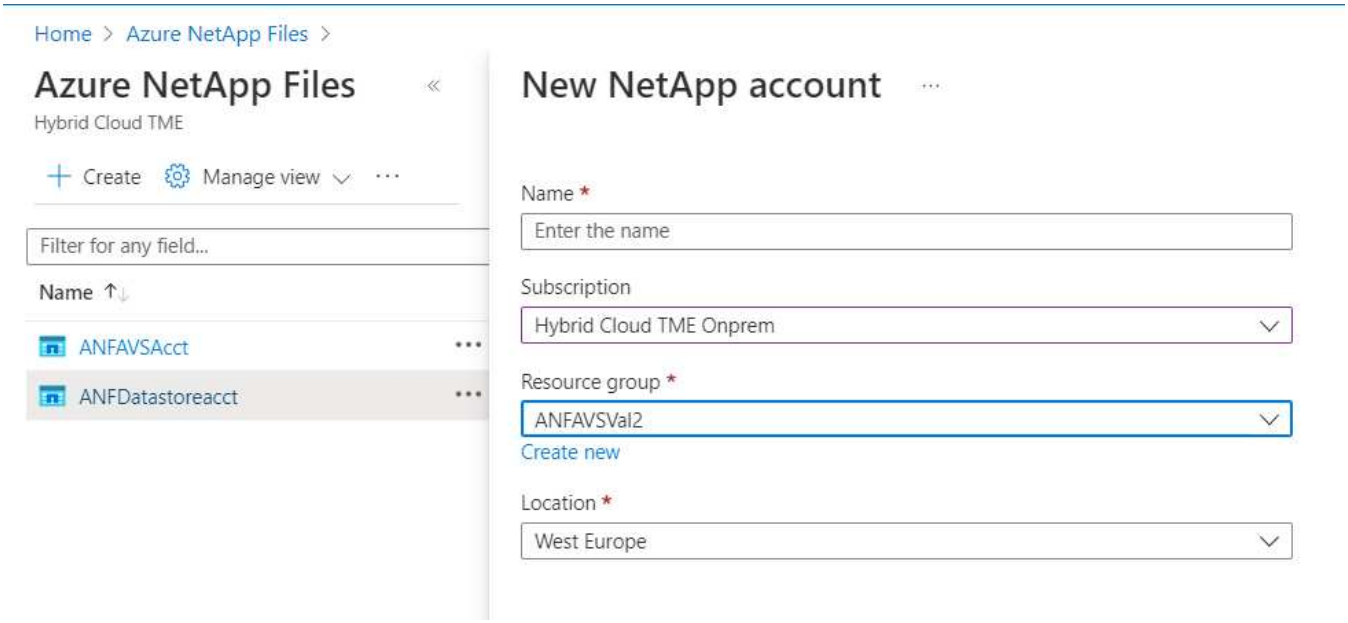
a. If the extension is not already installed, install it.

```
az extension add --name vmware
```

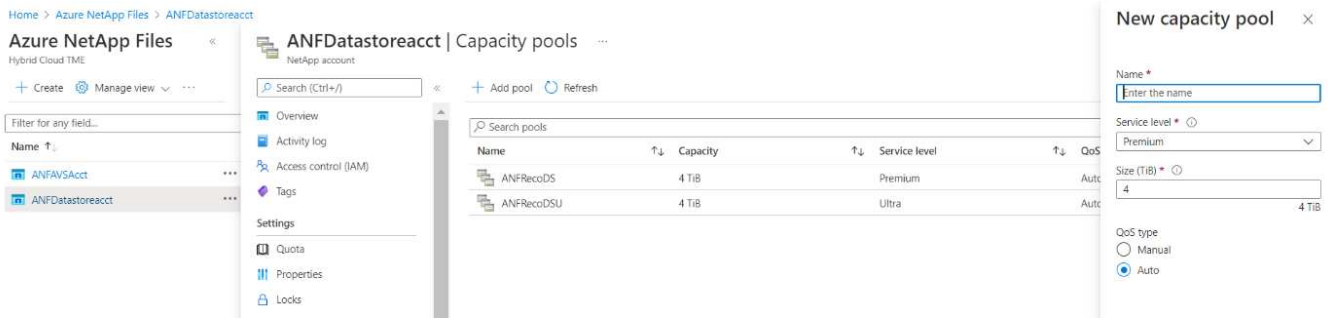
▼ Create and mount Azure NetApp Files volumes

1. Log into the Azure Portal and access Azure NetApp Files. Verify access to the Azure NetApp Files service and register the Azure NetApp Files Resource Provider by using the `az provider register`

`--namespace Microsoft.NetApp –wait` command. After registration, create a NetApp account. Refer to this link for detailed steps.



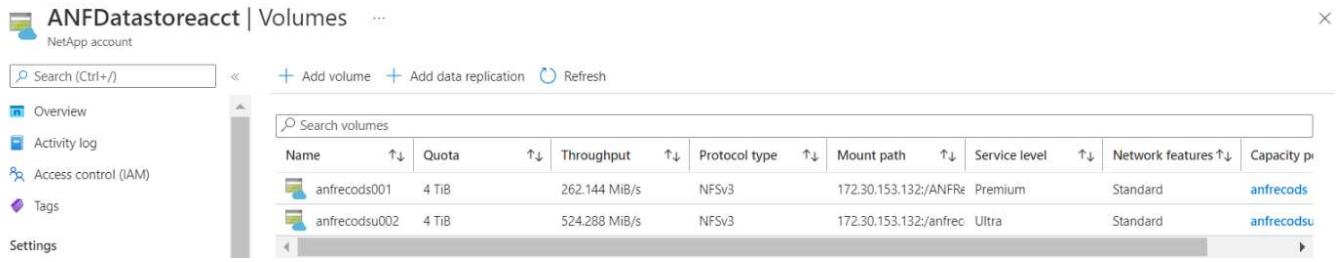
1. After a NetApp account is created, set up capacity pools with the required service level and size. For detailed information, refer to this link.



Points to Remember

- NFSv3 is supported for datastores on Azure NetApp Files.
- Use the Premium or standard tier for capacity bound workloads and Ultra tier for performance bound workloads where necessary while complementing default vSAN storage.

1. Configure a delegated subnet for Azure NetApp Files and specify this subnet when creating volumes. For detailed steps to create a delegated subnet, refer to this link.
2. Add an NFS volume for the datastore using the Volumes blade under the capacity pools blade.



To learn about Azure NetApp Files volume performance by size or quota, see Performance considerations for Azure NetApp Files.

▼ **Add Azure NetApp files datastore to private cloud**

Azure NetApp Files volume can be attached to your private cloud using Azure Portal. Follow this link from Microsoft for step by step approach of using Azure portal to mount an Azure NetApp files datastore.

To add an Azure NetApp files datastore to a private cloud, complete the following steps:

1. After the required features are registered, attach an NFS datastore to the Azure VMware Solution private cloud cluster by running the appropriate command.
2. Create a datastore using an existing ANF volume in the Azure VMware Solution private cloud cluster.

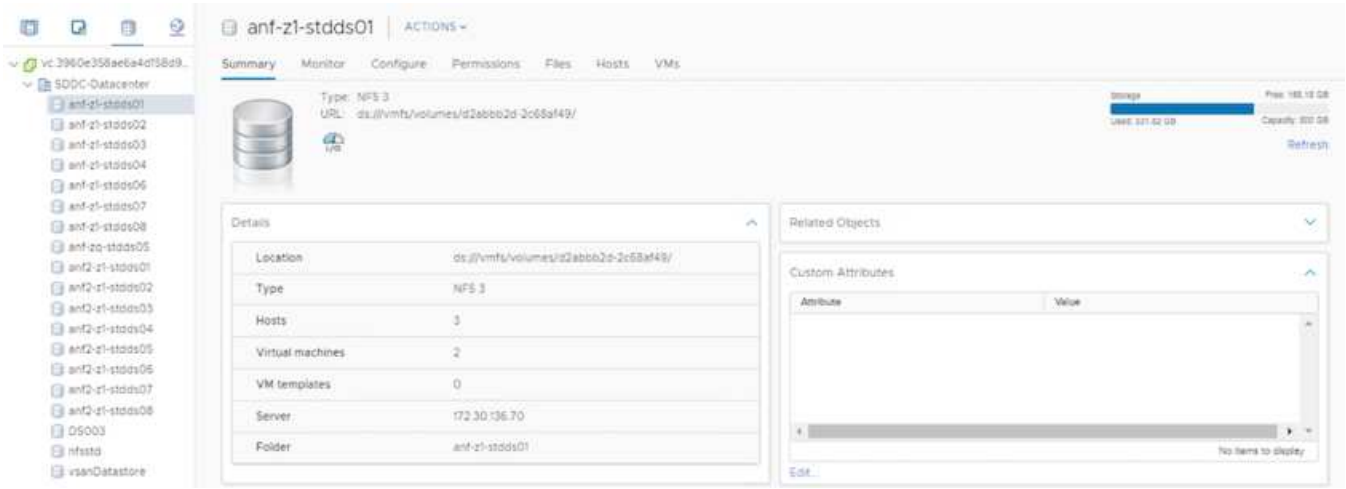
```
C:\Users\niyaz>az vmware datastore netapp-volume create --name
ANFRecoDSU002 --resource-group anfavsval2 --cluster Cluster-1 --private
-cloud ANFDataClus --volume-id /subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.NetApp/netApp
Accounts/anfdatastoreacct/capacityPools/anfrecodsu/volumes/anfrecodsU002
{
  "diskPoolVolume": null,
  "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.AVS/privateCl
ouds/ANFDataClus/clusters/Cluster-1/datastores/ANFRecoDSU002",
  "name": "ANFRecoDSU002",
  "netAppVolume": {
    "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.NetApp/netApp
Accounts/anfdatastoreacct/capacityPools/anfrecodsu/volumes/anfrecodsU002
",
    "resourceGroup": "anfavsval2"
  },
  "provisioningState": "Succeeded",
  "resourceGroup": "anfavsval2",
  "type": "Microsoft.AVS/privateClouds/clusters/datastores"
}

. List all the datastores in a private cloud cluster.
```

C:\Users\niyaz>az vmware datastore list --resource-group anfavsval2 --cluster Cluster-1 --private-cloud ANFDataClus
[
{
"diskPoolVolume": null,
"id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.AVS/privateClouds/ANFDataClus/clusters/Cl

uster-1/datastores/ANFRecoDS001",
"name": "ANFRecoDS001",
"netAppVolume": {
"id": "/subscriptions/0efa2dfb-917c-4497-b56a-b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.NetApp/netAppAccounts/anfdatastoreacct/capacityPools/anfrecods/volumes/ANFRecoDS001",
"resourceGroup": "anfavsval2"
},
"provisioningState": "Succeeded",
"resourceGroup": "anfavsval2",
"type": "Microsoft.AVS/privateClouds/clusters/datastores"
},
{
"diskPoolVolume": null,
"id": "/subscriptions/0efa2dfb-917c-4497-b56a-b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.AVS/privateClouds/ANFDataClus/clusters/Cluster-1/datastores/ANFRecoDSU002",
"name": "ANFRecoDSU002",
"netAppVolume": {
"id": "/subscriptions/0efa2dfb-917c-4497-b56a-b3f4eadb8111/resourceGroups/anfavsval2/providers/Microsoft.NetApp/netAppAccounts/anfdatastoreacct/capacityPools/anfrecodsu/volumes/anfrecodsU002",
"resourceGroup": "anfavsval2"
},
"provisioningState": "Succeeded",
"resourceGroup": "anfavsval2",
"type": "Microsoft.AVS/privateClouds/clusters/datastores"
}
]

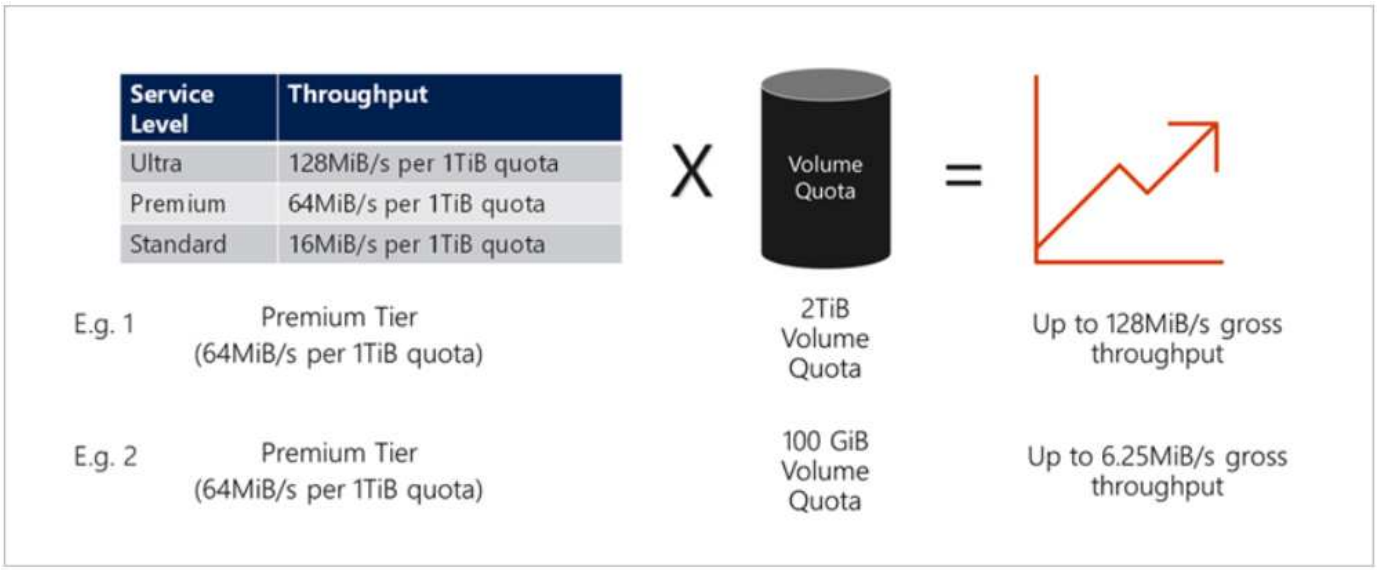
1. After the necessary connectivity is in place, the volumes are mounted as a datastore.



Sizing and performance optimization

Azure NetApp Files supports three service levels: Standard (16MBps per terabyte), Premium (64MBps per terabyte), and Ultra (128MBps per terabyte). Provisioning the right volume size is important for optimal performance of the database workload. With Azure NetApp Files, volume performance and the throughput limit are determined based on the following factors:

- The service level of the capacity pool to which the volume belongs
- The quota assigned to the volume
- The quality of service (QoS) type (auto or manual) of the capacity pool



For more information, see Service levels for Azure NetApp Files.

Refer to this link from Microsoft for detailed performance benchmarks that can be used during a sizing exercise.

Points to Remember
• Use the Premium or Standard tier for datastore volumes for optimal capacity and performance. If performance is required, then Ultra tier can be used. • For guest mount requirements, use Premium or Ultra tier and for file share requirements for guest VMs, use Standard or Premium tier volumes.

Performance considerations

It is important to understand that with NFS version 3 there is only one active pipe for the connection between the ESXi host and a single storage target. This means that although there might be alternate connections available for failover, the bandwidth for a single datastore and the underlying storage are limited to what a single connection can provide.

To leverage more available bandwidth with Azure NetApp Files volumes, an ESXi host must have multiple connections to the storage targets. To address this issue, you can configure multiple datastores, with each datastore using separate connections between the ESXi host and the storage.

For higher bandwidth, as a best practice create multiple datastores using multiple ANF volumes, create VMDKs, and stripe the logical volumes across VMDKs.

Refer to this link from Microsoft for detailed performance benchmarks that can be used during a sizing exercise.

Points to Remember

- Azure VMware solution allows eight NFS datastores by default. This can be increased via a support request.
- Leverage ER fastpath along with Ultra SKU for higher bandwidth and lower latency.
  More information
- With the "Basic" network features in Azure NetApp files, the connectivity from Azure VMware Solution is bound by the bandwidth of the ExpressRoute circuit and the ExpressRoute Gateway.
- For Azure NetApp Files volumes with "Standard" network features, ExpressRoute FastPath is supported. When enabled, FastPath sends network traffic directly to Azure NetApp Files volumes, bypassing the gateway providing higher bandwidth and lower latency.

Increasing the size of the datastore

Volume reshaping and dynamic service level changes are completely transparent to the SDDC. In Azure NetApp Files, these capabilities provide continuous performance, capacity, and cost optimizations. Increase the size of NFS datastores by resizing the volume from Azure Portal or by using the CLI. After you are done, access vCenter, go to the datastore tab, right-click the appropriate datastore, and select Refresh Capacity Information. This approach can be used to increase the datastore capacity and to increase the performance of the datastore in a dynamic fashion with no downtime. This process is also completely transparent to applications.

Points to remember

- Volume reshaping and dynamic service level capability allow you to optimize cost by sizing for steady-state workloads and thus avoid overprovisioning.
- VAAI is not enabled.

Workloads

▼ **Migration**

One of the most common use cases is migration. Use VMware HCX or vMotion to move on-premises VMs. Alternatively, you can use Rivermeadow to migrate VMs to Azure NetApp Files datastores.

▼ **Data Protection**

Backing up VMs and quickly recovering them are among the great strengths of ANF datastores. Use Snapshot copies to make quick copies of your VM or datastore without affecting performance, and then send them to Azure storage for longer-term data protection or to a secondary region using cross region replication for disaster recovery purposes. This approach minimizes storage space and network bandwidth by only storing changed information.

Use Azure NetApp Files Snapshot copies for general protection, and use application tools to protect transactional data such as SQL Server or Oracle residing on the guest VMs. These Snapshot copies are different from VMware (consistency) snapshots and are suitable for longer term protection.



With ANF datastores, the Restore to New Volume option can be used to clone an entire datastore volume, and the restored volume can be mounted as another datastore to hosts within AVS SDDC. After a datastore is mounted, VMs inside it can be registered, reconfigured, and customized as if they were individually cloned VMs.

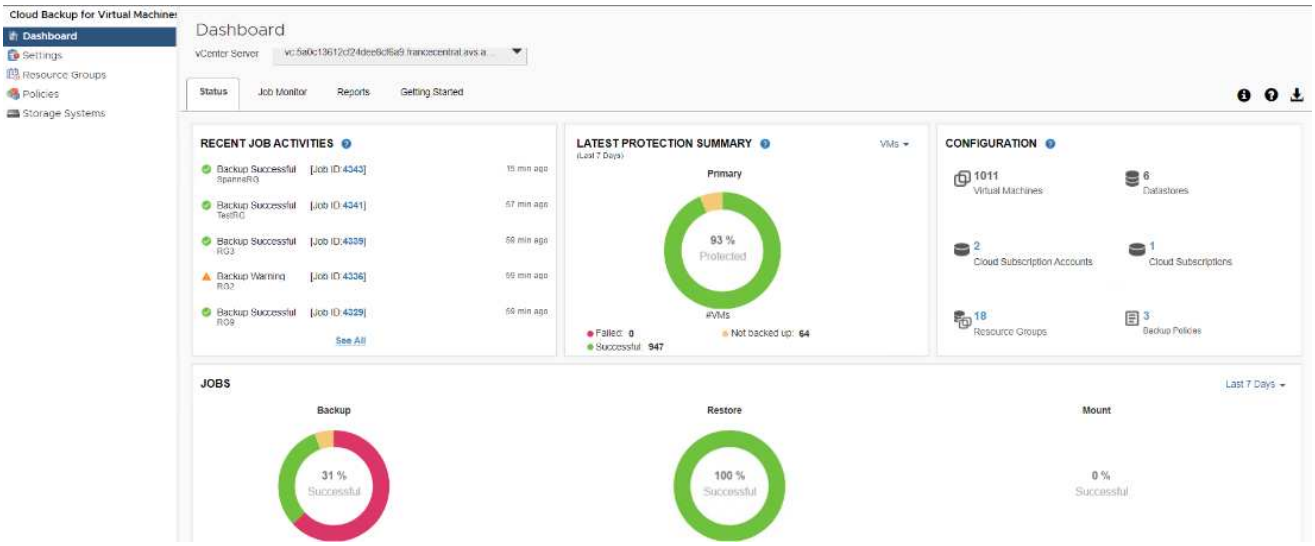
▼ **NetApp Backup and Recovery for Virtual Machines**

NetApp Backup and Recovery for Virtual Machines provides a vSphere web client GUI on vCenter to protect Azure VMware Solution virtual machines and Azure NetApp files datastores via backup policies. These policies can define schedule, retention, and other capabilities. The NetApp Backup and Recovery for Virtual Machine functionality can be deployed by using the Run command.

The setup and protection policies can be installed by completing the following steps:

1. Install NetApp Backup and Recovery for Virtual Machine in Azure VMware Solution private cloud using the Run command.
2. Add cloud subscription credentials (client and secret value), and then add a cloud subscription account (NetApp account and associated resource group) that contains the resources that you would like to protect.
3. Create one or more backup policies that manage the retention, frequency, and other settings for resource group backups.
4. Create a container to add one or more resources that need to be protected with backup policies.
5. In the event of a failure, restore the entire VM or specific individual VMDKs to the same location.

With Azure NetApp Files Snapshot technology, backups and restores are very fast.



▼ **Disaster Recovery with Azure NetApp Files, JetStream DR, and Azure VMware Solution**

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). Using the VMware VAIO framework, on-premises VMware workloads can be replicated to Azure Blob storage and recovered, enabling minimal or close to no data loss and near-zero RTO. JetStream DR can be used to seamlessly recover the workloads replicated from on-premises to AVS and specifically to Azure NetApp Files. It enables cost-effective disaster recovery by using minimal resources at the DR site and cost-effective cloud storage. JetStream DR automates recovery to ANF datastores via Azure Blob Storage. JetStream DR recovers independent VMs or groups of related VMs into recovery site infrastructure according to network mapping and provides point-in-time recovery for ransomware protection.

DR solution with ANF, JetStream and AVS.

NetApp Guest Connected Storage Options for Azure

Azure supports guest connected NetApp storage with the native Azure NetApp Files

(ANF) service or with Cloud Volumes ONTAP (CVO).

Azure NetApp Files (ANF)

Azure netApp Files brings enterprise-grade data management and storage to Azure so you can manage your workloads and applications with ease. Migrate your workloads to the cloud and run them without sacrificing performance.

Azure netApp Files removes obstacles, so you can move all of your file-based applications to the cloud. For the first time, you do not
have to re-architect your applications, and you get persistent storage for your applications without complexity.

Because the service is delivered through the Microsoft Azure Portal, users experience a fully managed service as part of their Microsoft enterprise Agreement. World-class support, managed by Microsoft, gives you complete peace of mind. This single solution enables you to quickly and easily add multiprotocol workloads. you can build and deploy both Windows and Linux file-based applications, even for legacy environments.

Azure NetApp Files (ANF) as guest connected storage

▼ **Configure Azure NetApp Files with Azure VMware Solution (AVS)**

Azure NetApp Files shares can be mounted from VMs that are created in the Azure VMware Solution SDDC environment. The volumes can also be mounted on the Linux client and mapped on the Windows client because Azure NetApp Files supports SMB and NFS protocols. Azure NetApp Files volumes can be set up in five simple steps.

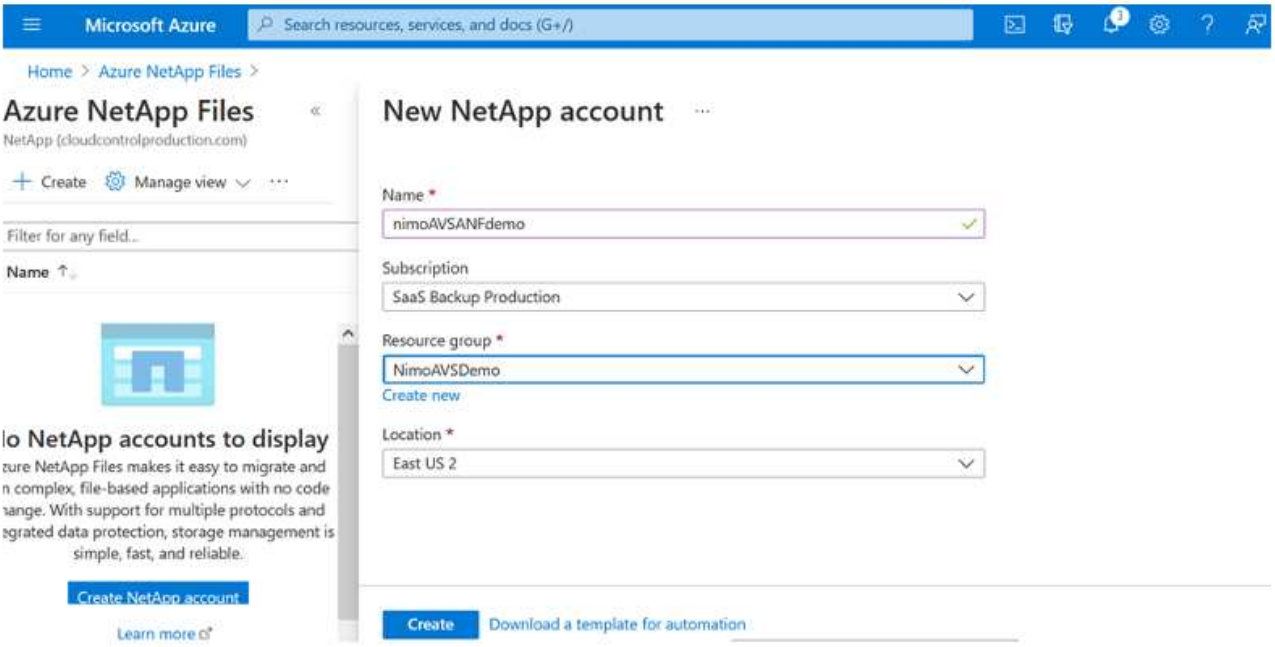
Azure NetApp Files and Azure VMware Solution must be in the same Azure region.

▼ **Create and mount Azure NetApp Files volumes**

To create and mount Azure NetApp Files volumes, complete the following steps:

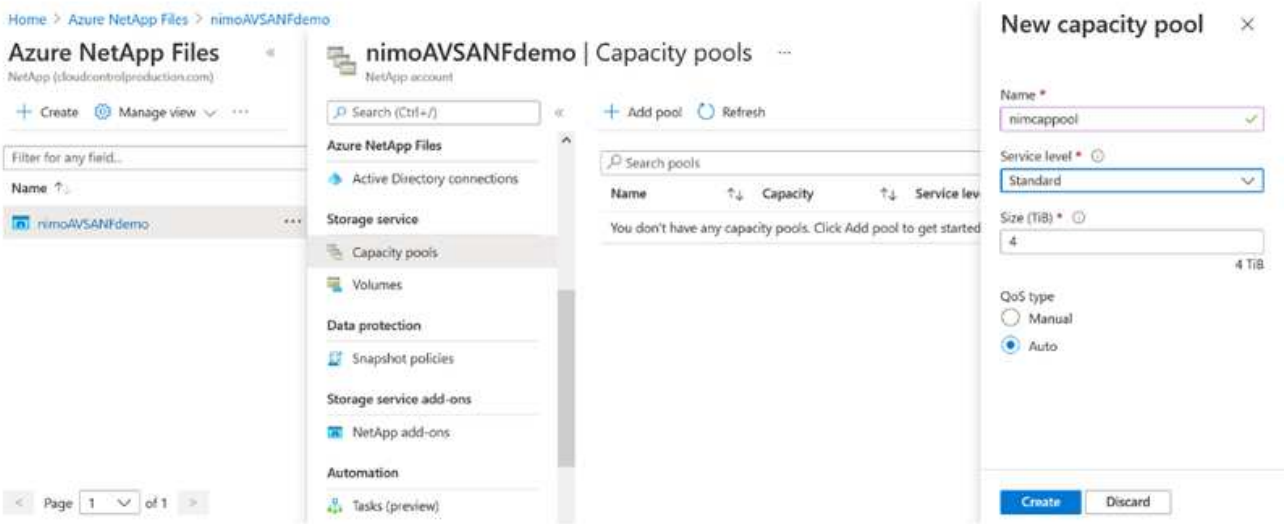
1. Log in to the Azure Portal and access Azure NetApp Files. Verify access to the Azure NetApp Files service and register the Azure NetApp Files Resource Provider by using the *az provider register --namespace Microsoft.NetApp –wait* command. After registration is complete, create a NetApp account.

   For detailed steps, see Azure NetApp Files shares. This page will guide you through the step-by-step process.

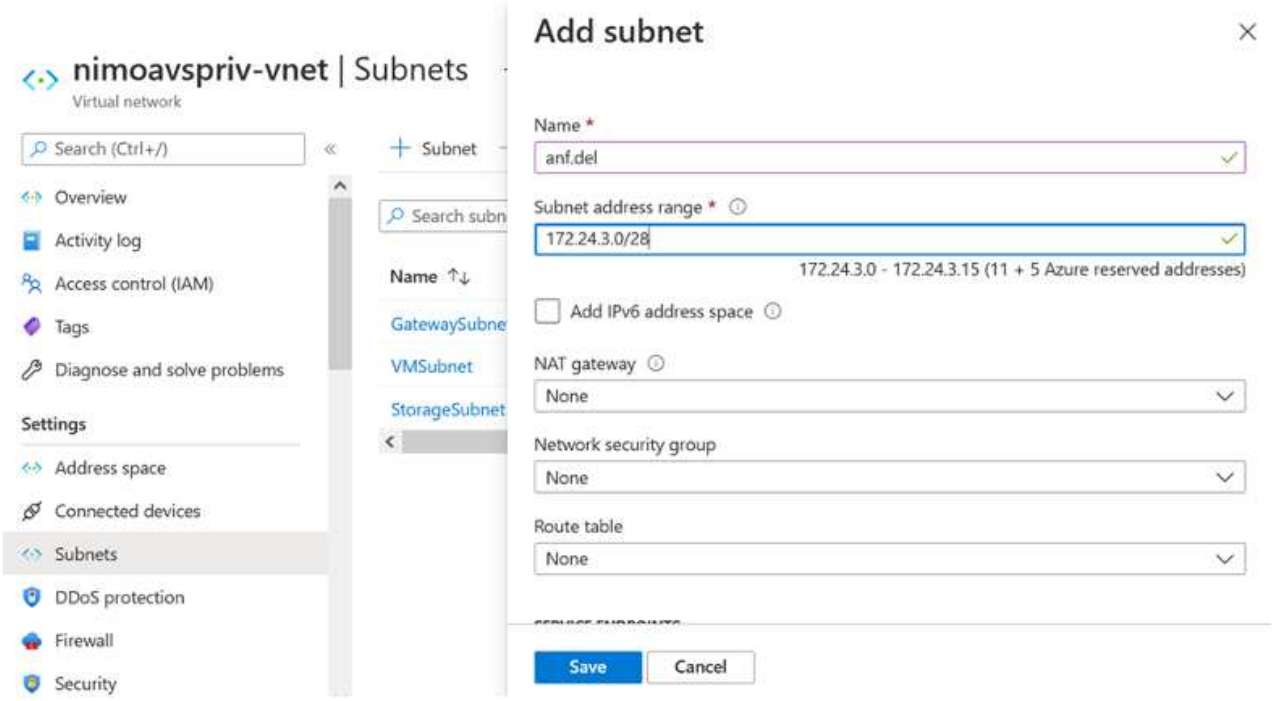


2. After the NetApp account is created, set up the capacity pools with the required service level and size.

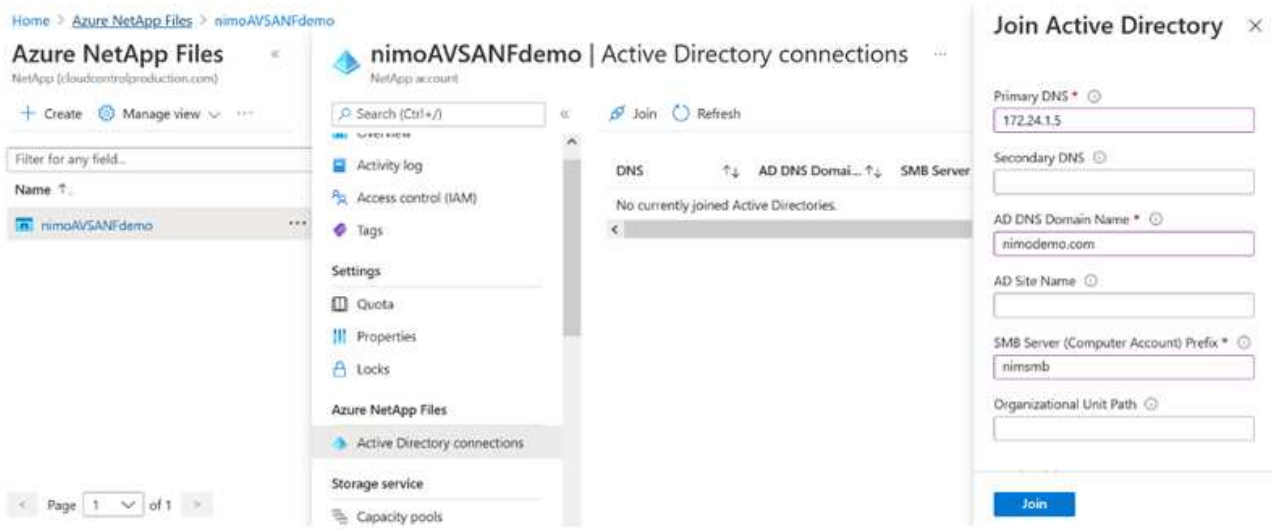
   For more information, see Set up a capacity pool.



3. Configure the delegated subnet for Azure NetApp Files and specify this subnet while creating the volumes. For detailed steps to create delegated subnet, see Delegate a subnet to Azure NetApp Files.

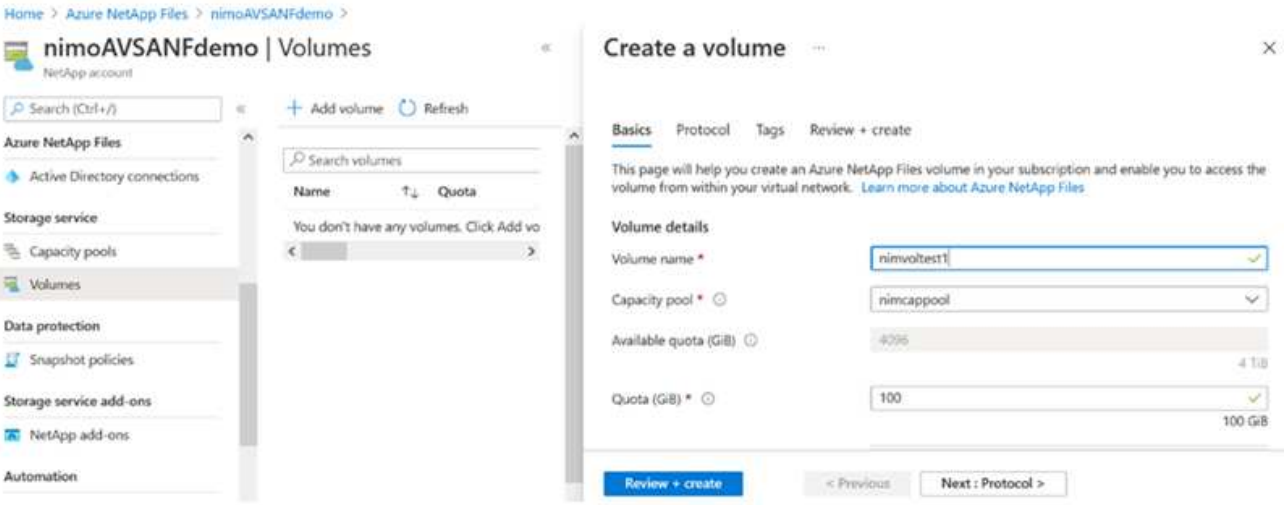


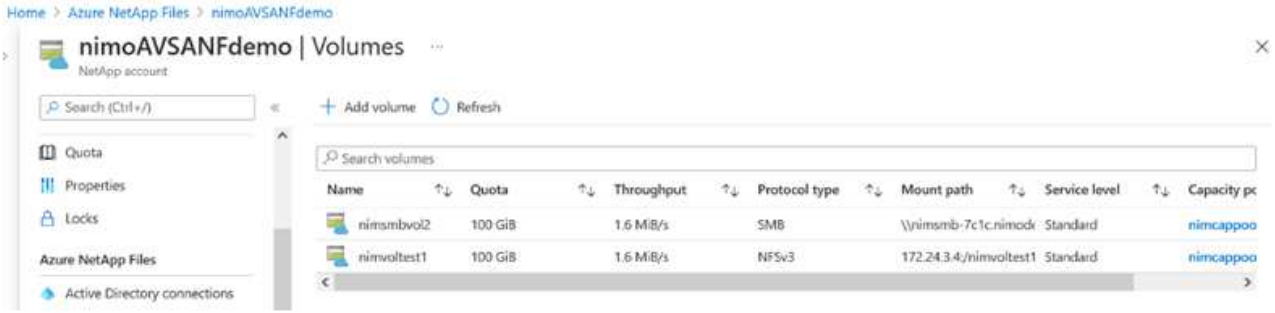
4. Add an SMB volume by using the Volumes blade under the Capacity Pools blade. Make sure the Active Directory connector is configured prior to creating the SMB volume.



5. Click Review + Create to create the SMB volume.

   If the application is SQL Server, then enable the SMB continuous availability.

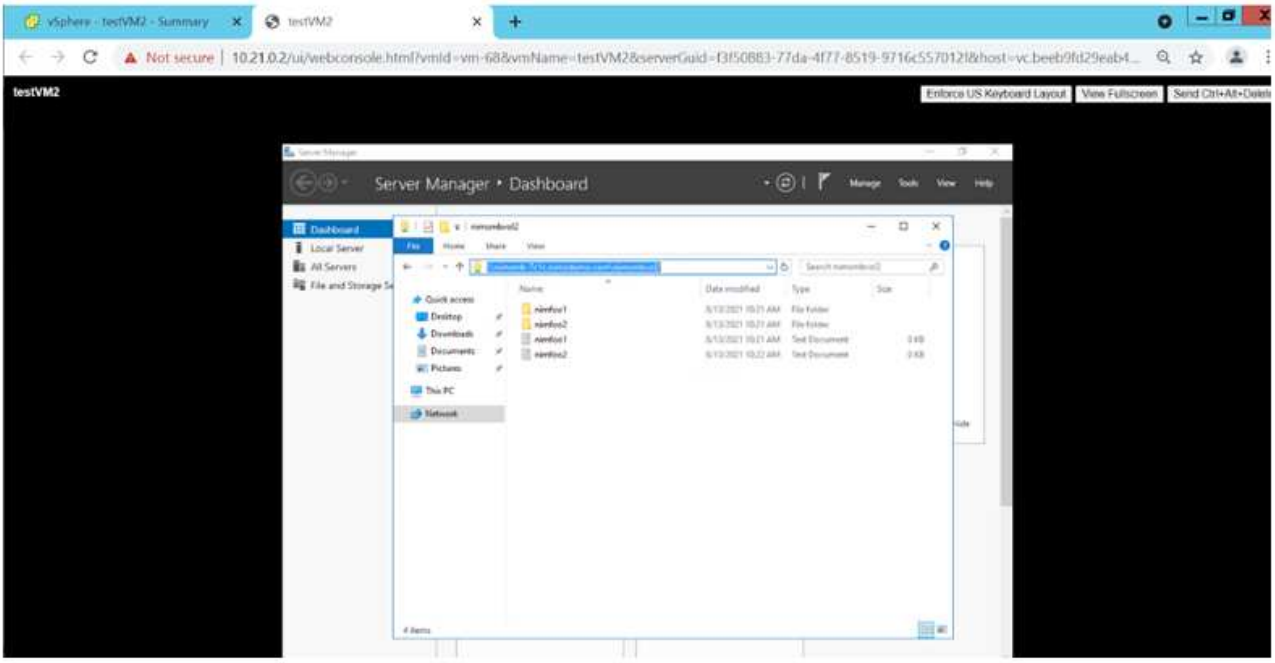


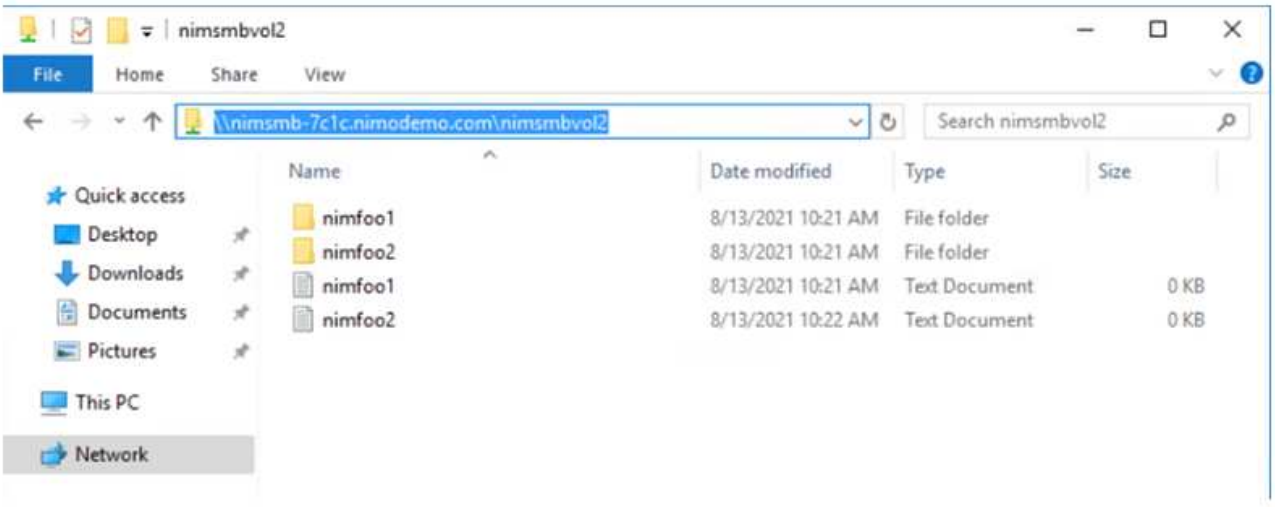


To learn more about Azure NetApp Files volume performance by size or quota, see Performance considerations for Azure NetApp Files.

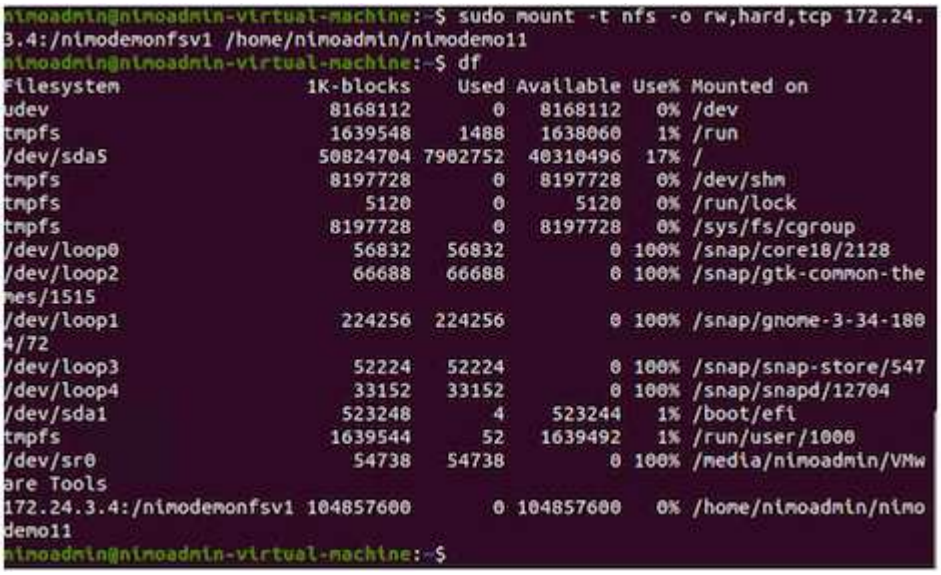
6. After the connectivity is in place, the volume can be mounted and used for application data.

   To accomplish this, from the Azure portal, click the Volumes blade, and then select the volume to mount and access the mount instructions. Copy the path and use the Map Network Drive option to mount the volume on the VM running on Azure VMware Solution SDDC.





7. To mount NFS volumes on Linux VMs running on Azure VMware Solution SDDC, use this same process. Use volume reshaping or dynamic service level capability to meet the workload demands.



For more information, see Dynamically change the service level of a volume.

Cloud Volumes ONTAP (CVO)

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

Cloud Volumes ONTAP (CVO) as guest connected storage

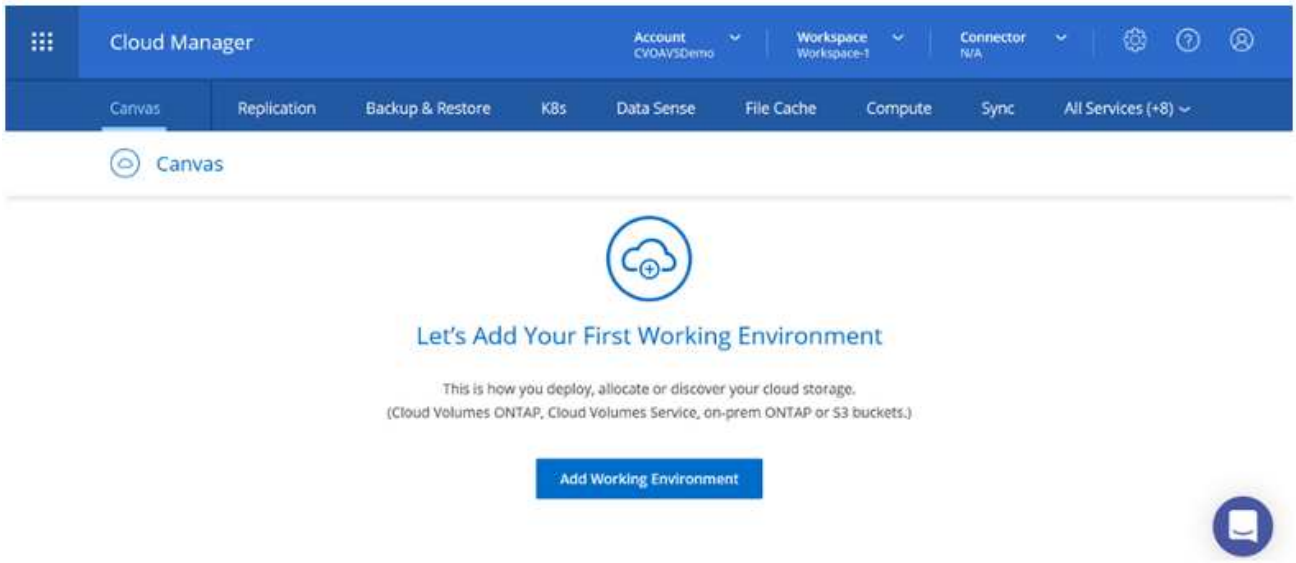
▼ **Deploy new Cloud Volumes ONTAP in Azure**

Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the Azure VMware Solution SDDC environment. The volumes can also be mounted on the Linux client and on Windows client because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be set up in a few simple steps.

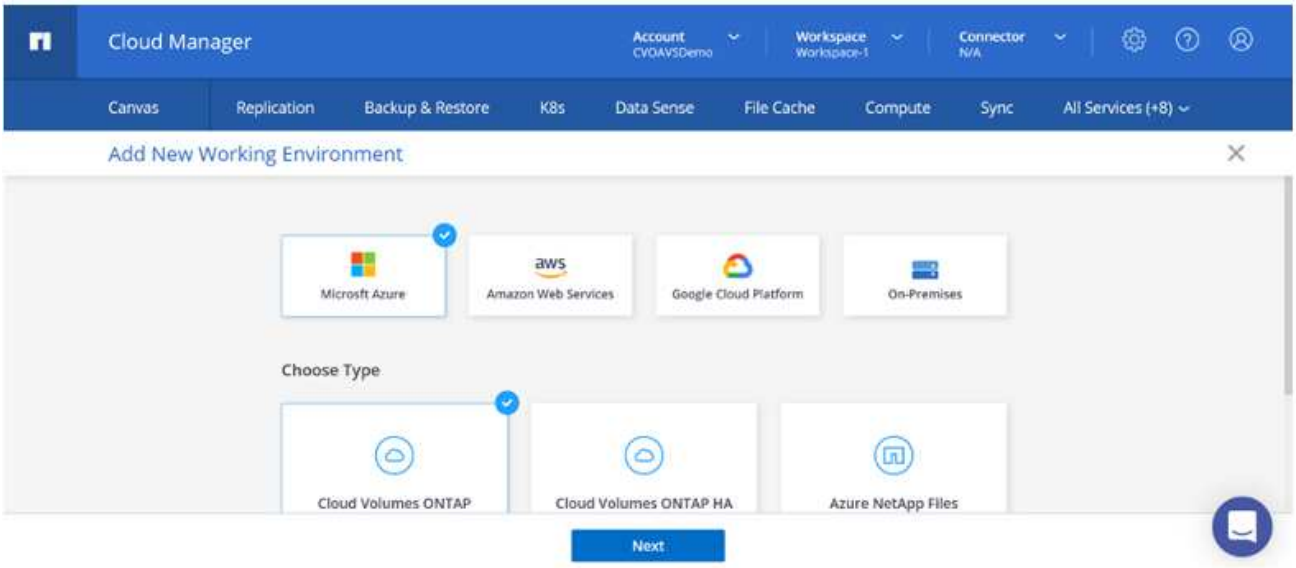
To replicate volumes from an on-premises environment to the cloud for disaster recovery or migration purposes, establish network connectivity to Azure, either using a site-to-site VPN or ExpressRoute. Replicating data from on-premises to Cloud Volumes ONTAP is outside the scope of this document. To replicate data between on-premises and Cloud Volumes ONTAP systems, see Setting up data replication between systems.

ⓘ Use Cloud Volumes ONTAP sizer to accurately size the Cloud Volumes ONTAP instances. Also monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

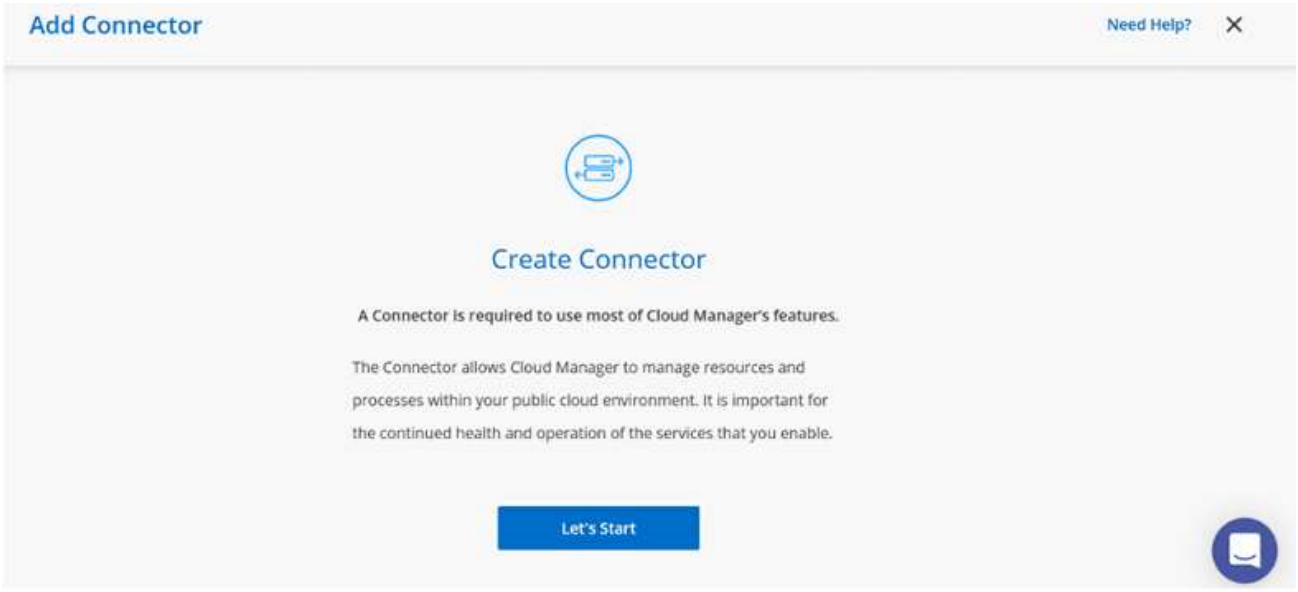
1. Log in to NetApp Cloud Central—the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.



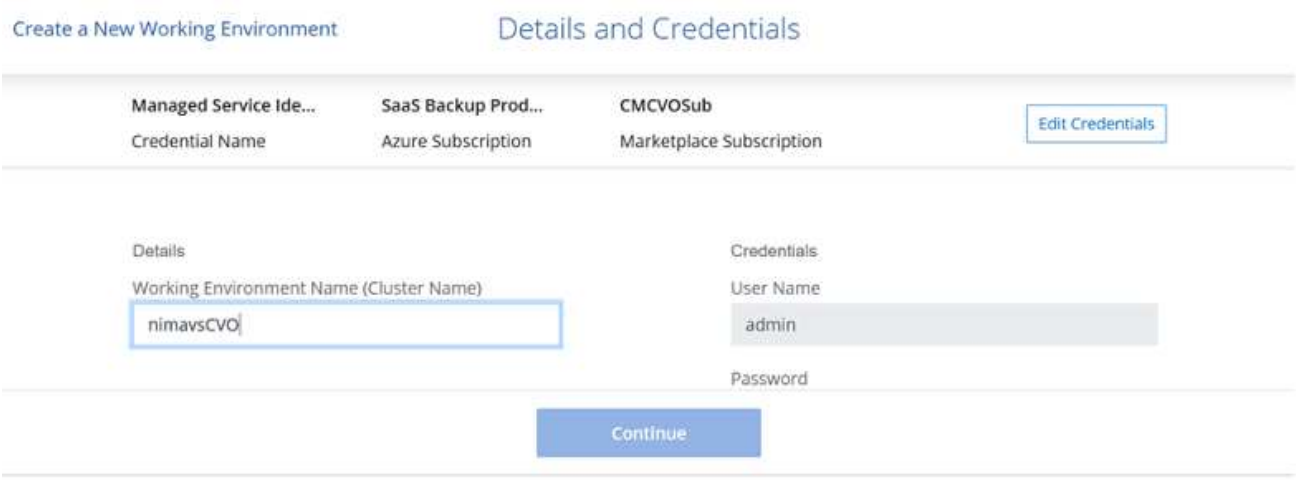
2. On the Cloud Manager home page, click Add a Working Environment and then select Microsoft Azure as the cloud and the type of the system configuration.



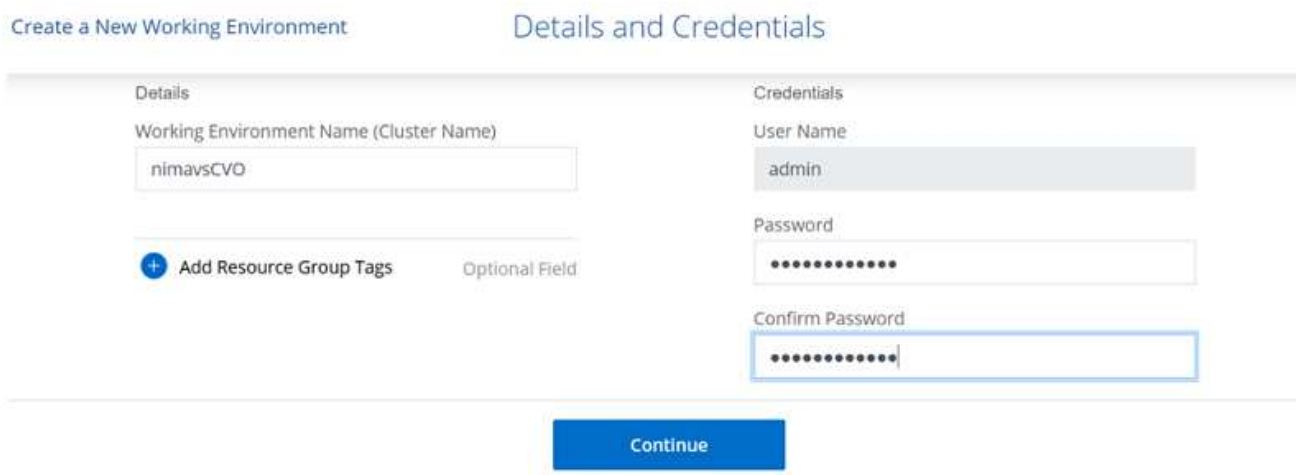
3. When creating the first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to deploy a Connector.



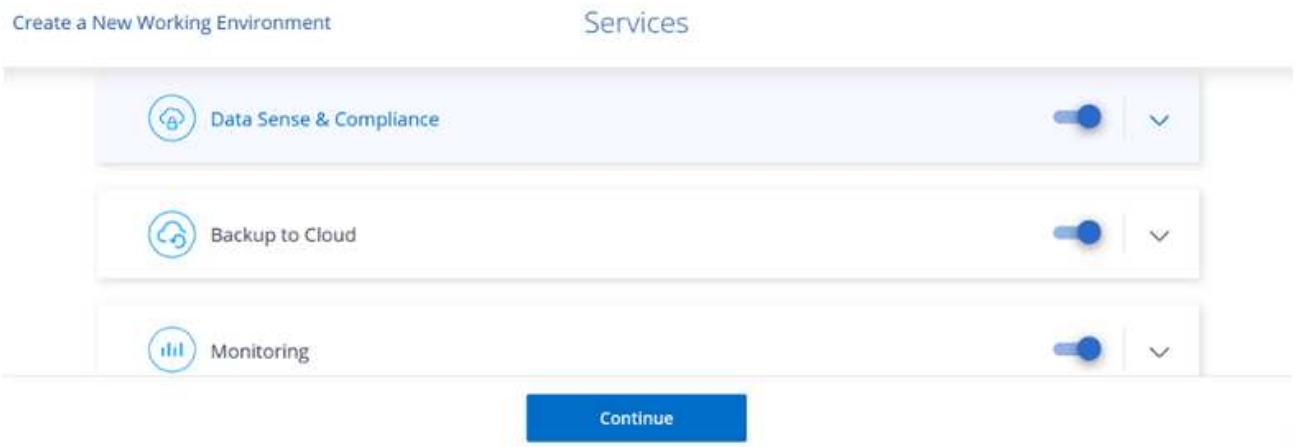
4. After the connector is created, update the Details and Credentials fields.



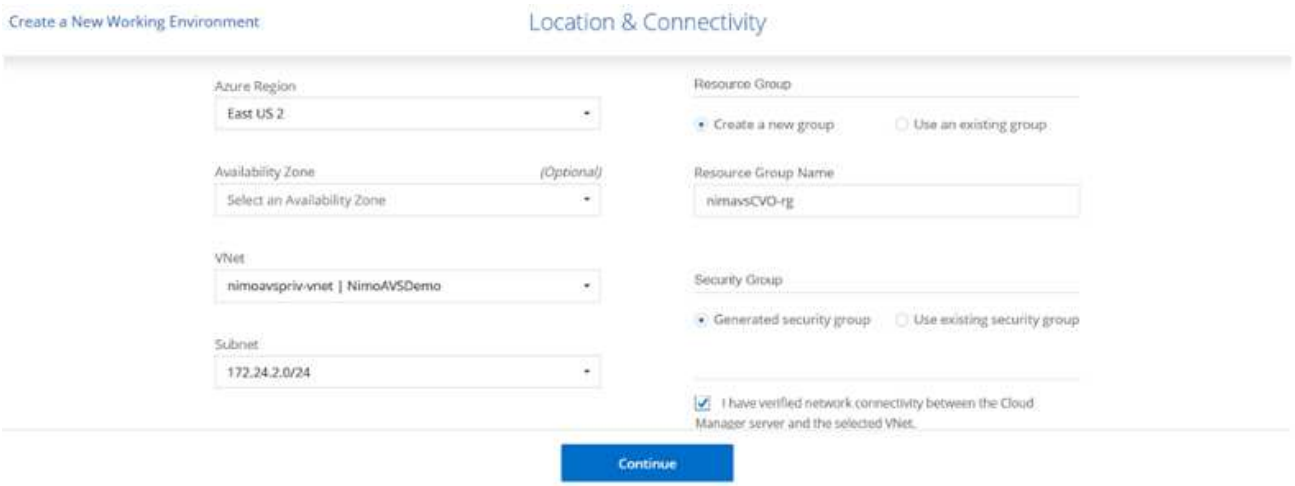
5. Provide the details of the environment to be created including the environment name and admin credentials. Add resource group tags for the Azure environment as an optional parameter. After you are done, click Continue.



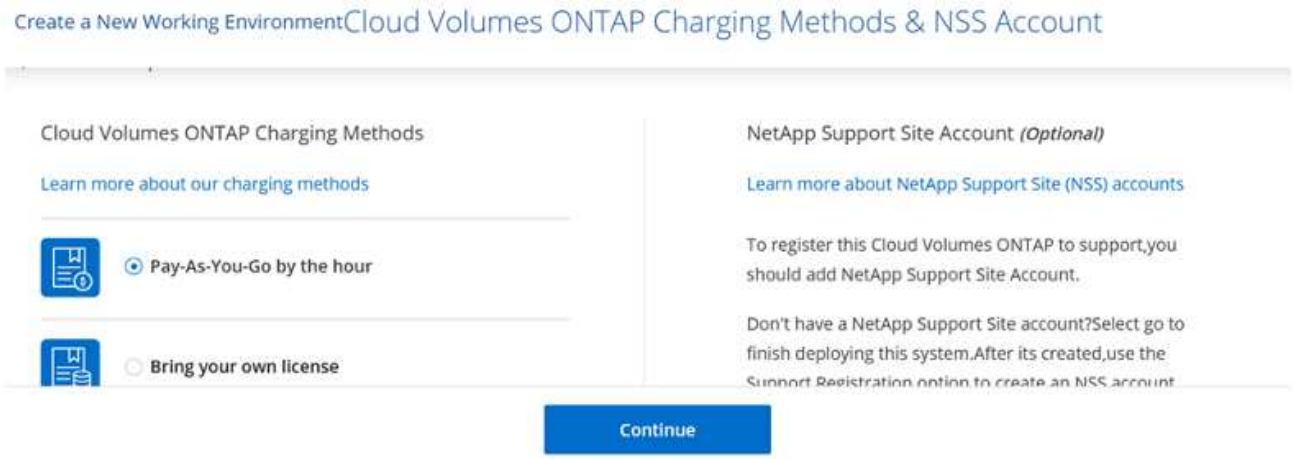
6. Select the add-on services for Cloud Volumes ONTAP deployment, including NetApp Data Classification, NetApp Backup and Recovery, and Data Infrastructure Insights. Select the services and then click Continue.



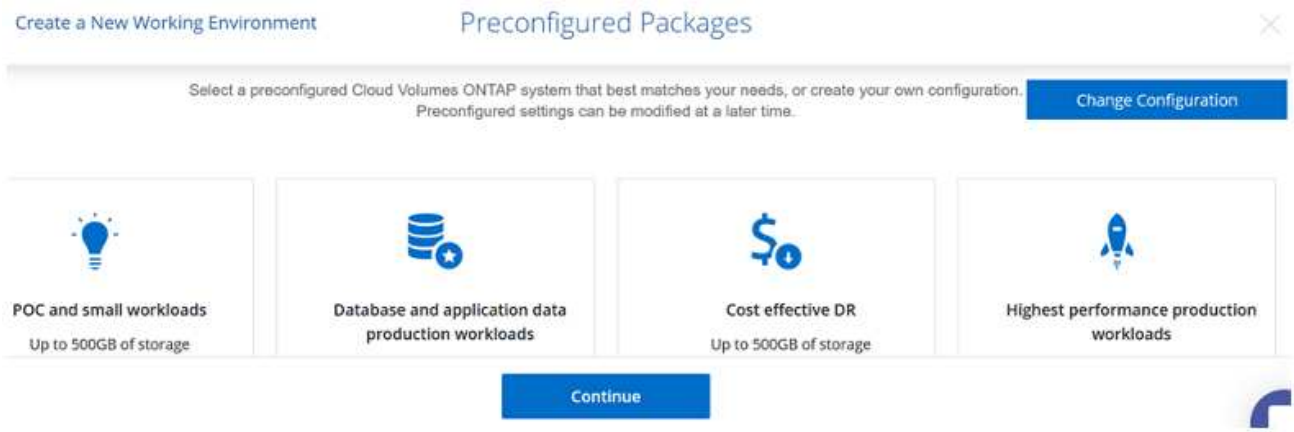
7. Configure the Azure location and connectivity. Select the Azure Region, resource group, VNet, and subnet to be used.



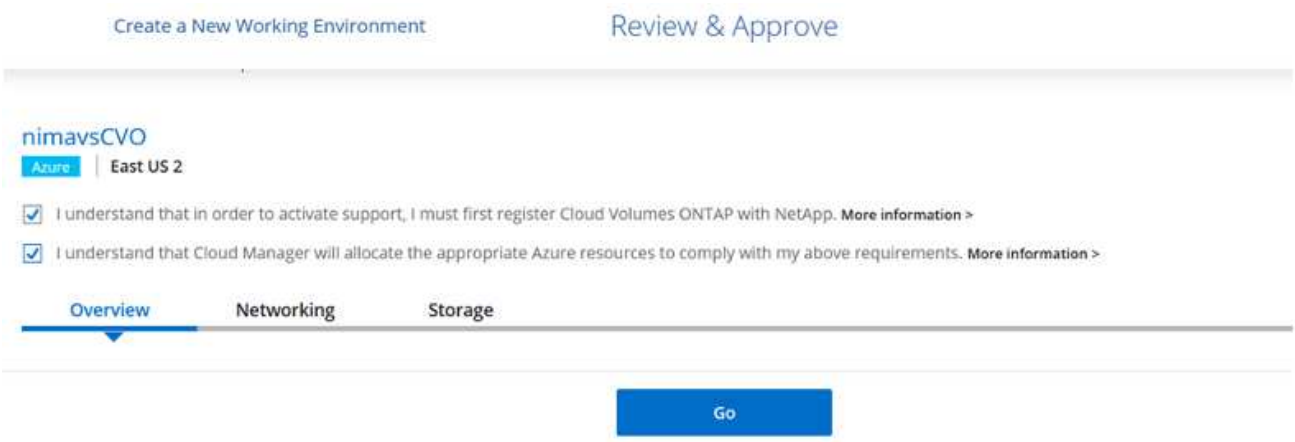
8. Select the license option: Pay-As-You-Go or BYOL for using existing license. In this example, Pay-As-You-Go option is used.



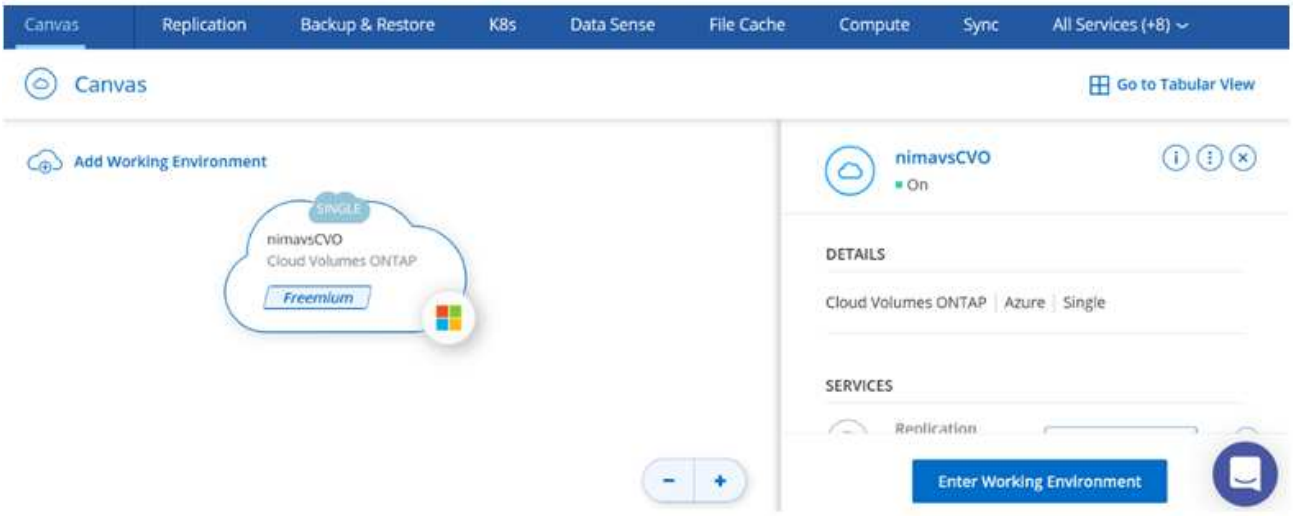
9. Select between several preconfigured packages available for the various types of workloads.



10. Accept the two agreements regarding activating support and allocation of Azure resources.To create the Cloud Volumes ONTAP instance, click Go.



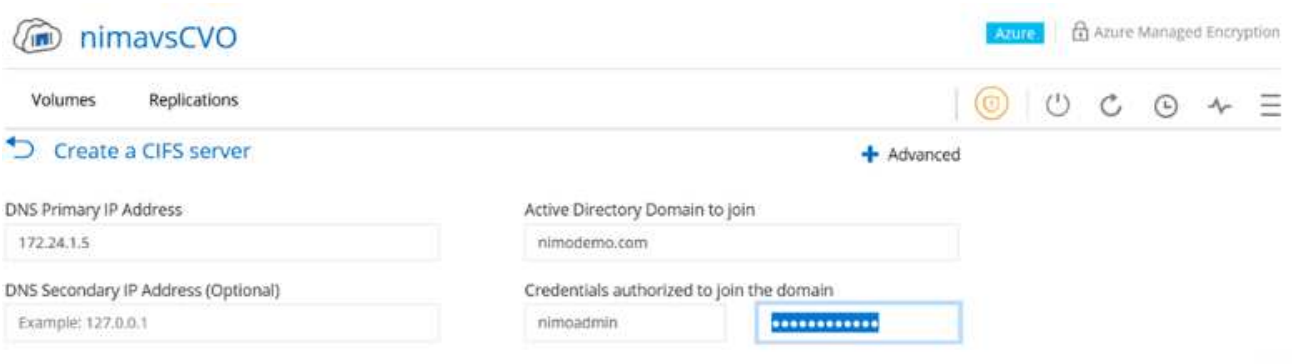
11. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.



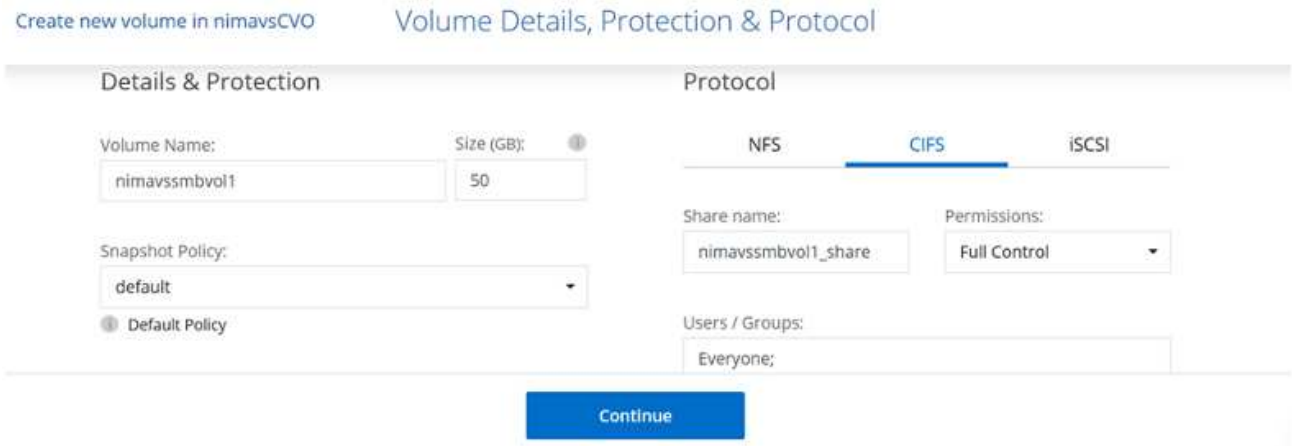
▼ Additional configurations for SMB volumes

1. After the working environment is ready, make sure the CIFS server is configured with the appropriate

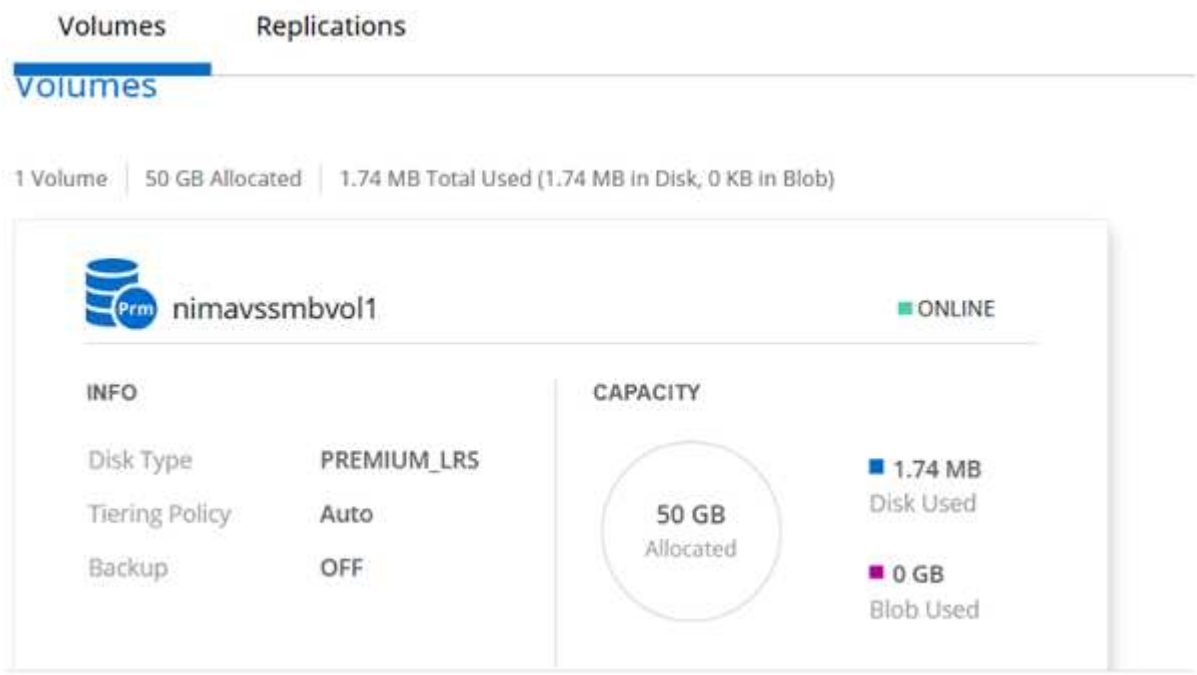
DNS and Active Directory configuration parameters. This step is required before you can create the SMB volume.



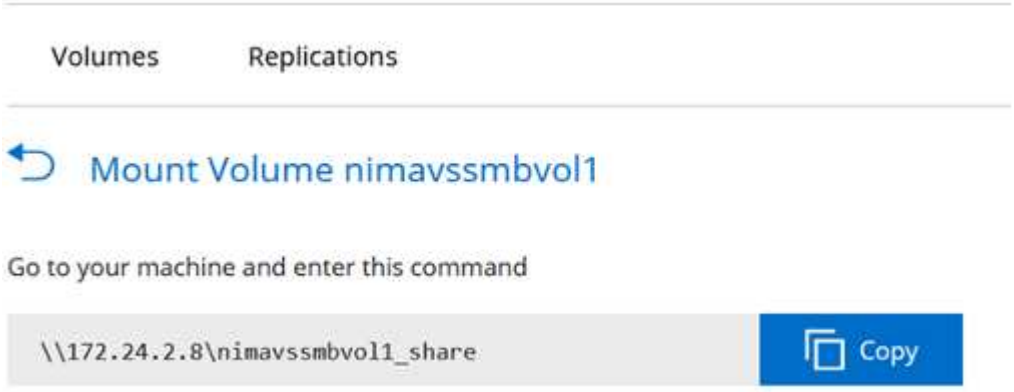
2. Creating the SMB volume is an easy process. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected as the protocol.

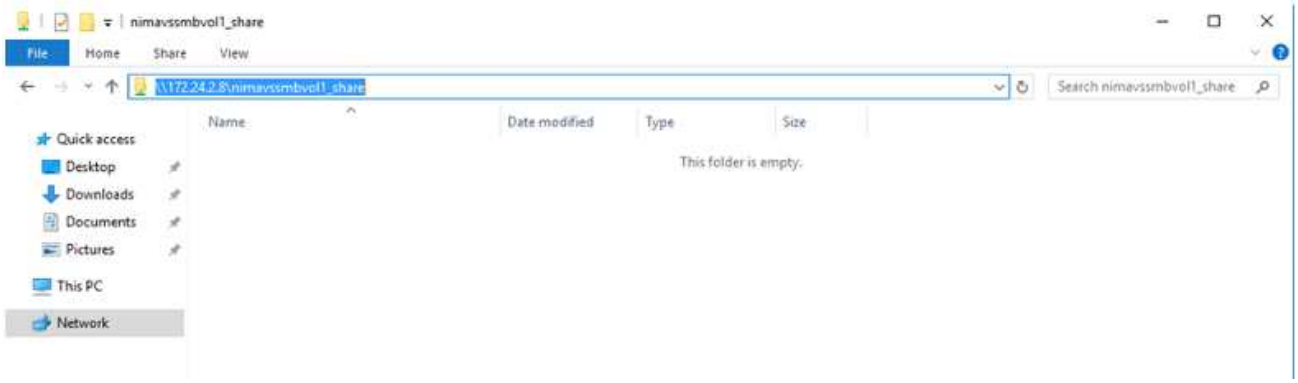


3. After the volume is provisioned, it will be availabe under the Volumes pane. Because a CIFS share is provisioned, give your users or groups permission to the files and folders and verify that those users can access the share and create a file. This step is not required if the volume is replicated from an on-premises environment because the file and folder permissions are all retained as part of SnapMirror replication.



4. After the volume is created, use the mount command to connect to the share from the VM running on the Azure VMware Solution SDDC hosts.
5. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on Azure VMware Solution SDDC.

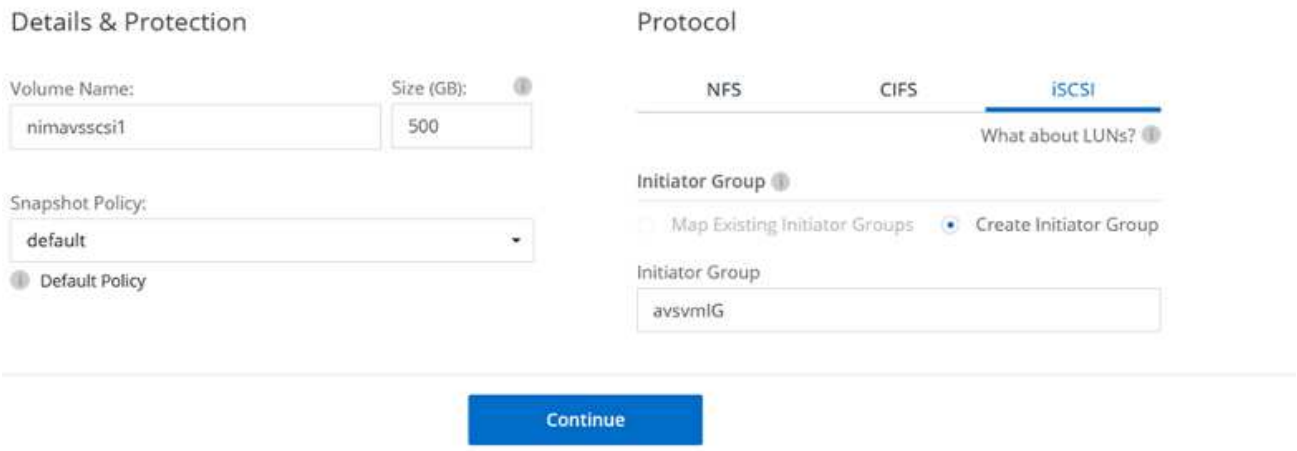




▼ **Connect the LUN to a host**

To connect the LUN to a host, complete the following steps:

1. On the Canvas page, double-click the Cloud Volumes ONTAP working environment to create and manage volumes.
2. Click Add Volume > New Volume and select iSCSI and click Create Initiator Group. Click Continue.



3. After the volume is provisioned, select the volume, and then click Target IQN. To copy the iSCSI Qualified Name (IQN), click Copy. Set up an iSCSI connection from the host to the LUN.

   To accomplish the same for the host residing on Azure VMware Solution SDDC:

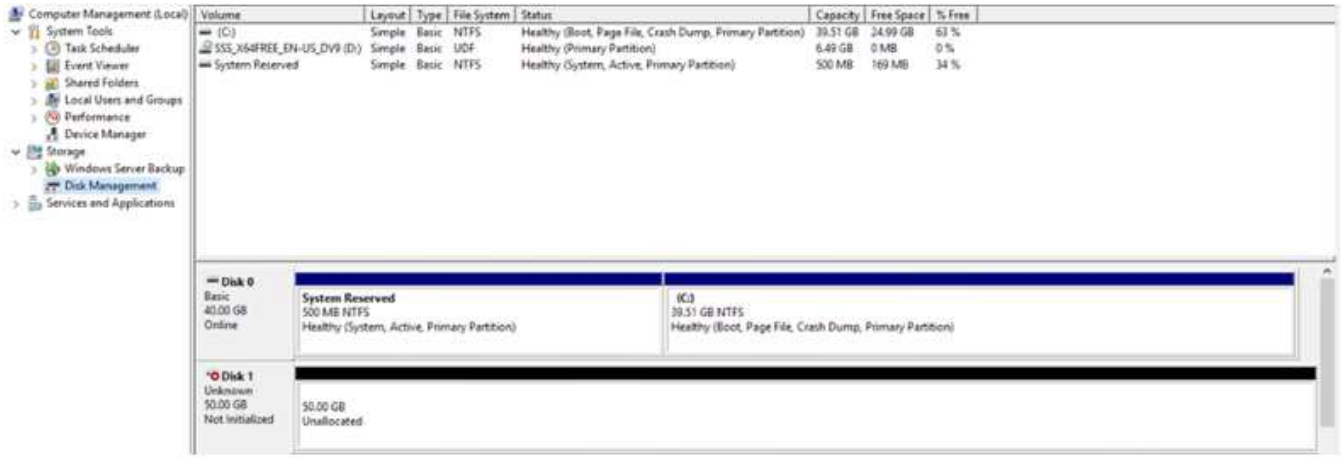
   a. RDP to the VM hosted on Azure VMware Solution SDDC.
   b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.
   c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.
   d. From the Targets tab, select the target discovered and then click Log on or Connect.
   e. Select Enable multipath, and then select Automatically Restore This Connection When the Computer Starts or Add This Connection to the List of Favorite Targets. Click Advanced.

      Note: The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



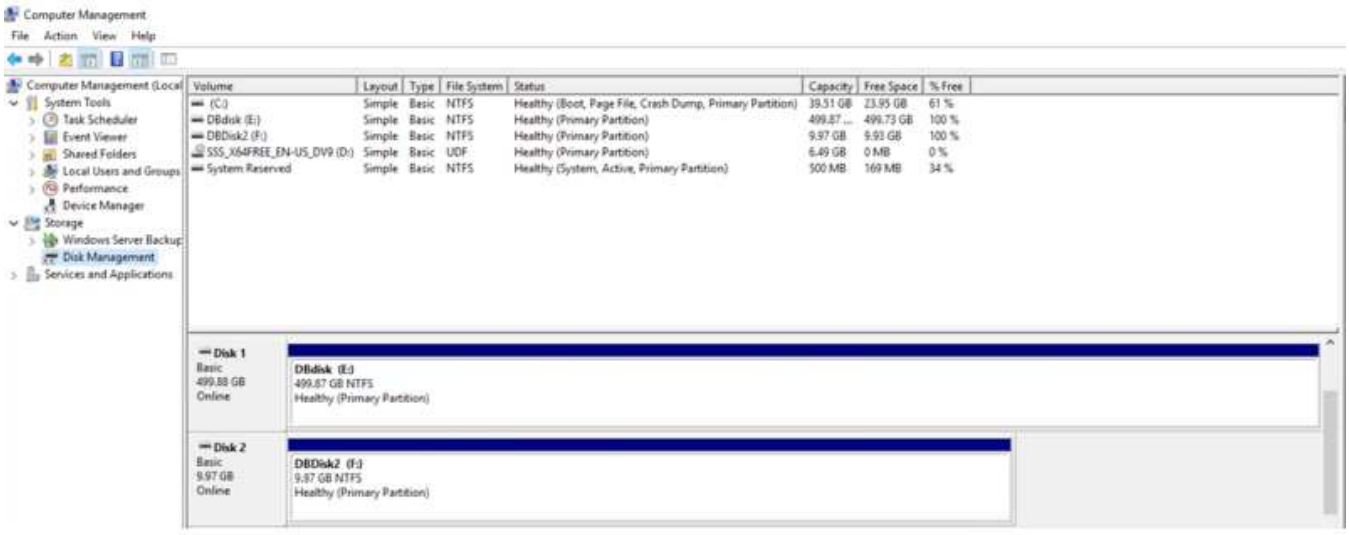
LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

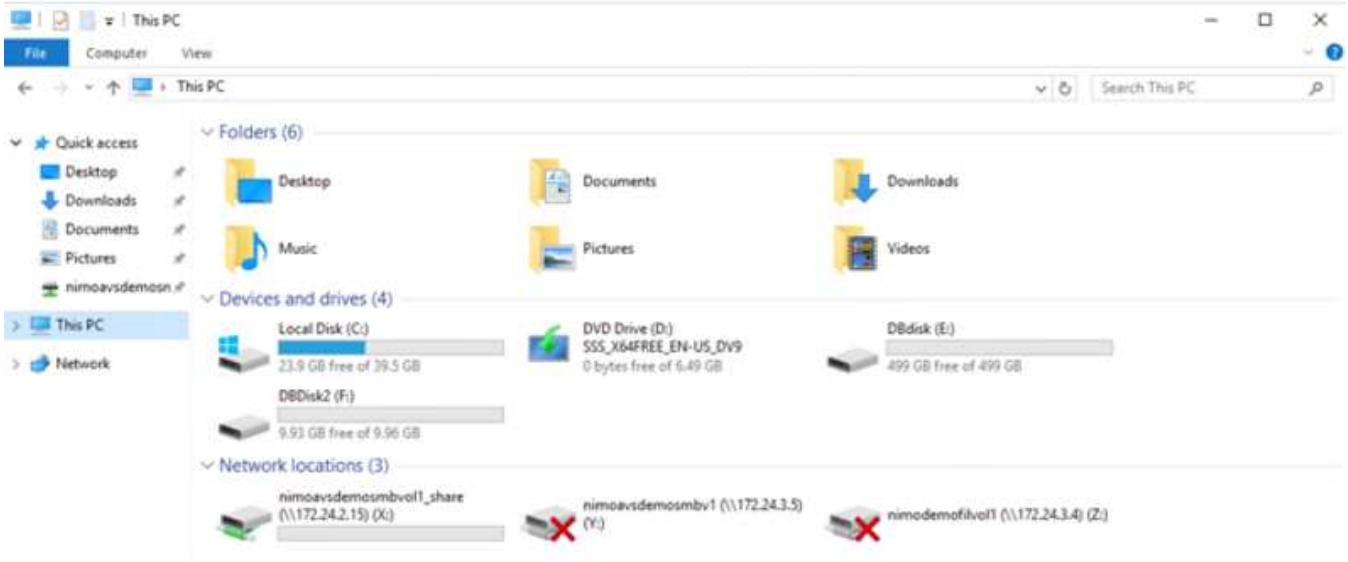
1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
2. Expand the Storage node in the navigation tree.
3. Click Disk Management.
4. Click Action > Rescan Disks.



When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

1. Start Windows Disk Management.
2. Right-click the LUN, and then select the required disk or partition type.
3. Follow the instructions in the wizard. In this example, drive E: is mounted





Google Cloud VMware Engine: Options for using NetApp storage

NetApp storage can be attached to the Google Cloud Virtualization Engine as either guest connected or supplemental storage.

Google Cloud NetApp Volumes (NetApp Volumes) as a supplemental NFS Datastore

Customers that requires additional storage capacity on their Google Cloud VMware Engine (GCVE) environment can utilize Google Cloud NetApp Volumes to mount as supplemental NFS datastore. Storing data on Google Cloud NetApp Volumes allows customers to replicate between regions to protect from diaster.

For more information, please visit Google Cloud NetApp Volumes (NetApp Volumes) as a supplemental NFS Datastore

NetApp CVO as Guest Connected Storage

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud.

For more information, please visit NetApp CVO as Guest Connected Storage

Google Cloud NetApp Volumes (NetApp Volumes) as Guest Connected Storage

Google Cloud NetApp Volumes shares can be mounted from VMs that are created in the VMware Engine environment. The volumes can also be mounted on the Linux client and mapped on the Windows client because Google Cloud NetApp Volumes supports SMB and NFS protocols. Google Cloud NetApp Volumes volumes can be set up in simple steps.

Google Cloud NetApp Volumes and Google Cloud VMware Engine private cloud must be in the same region.

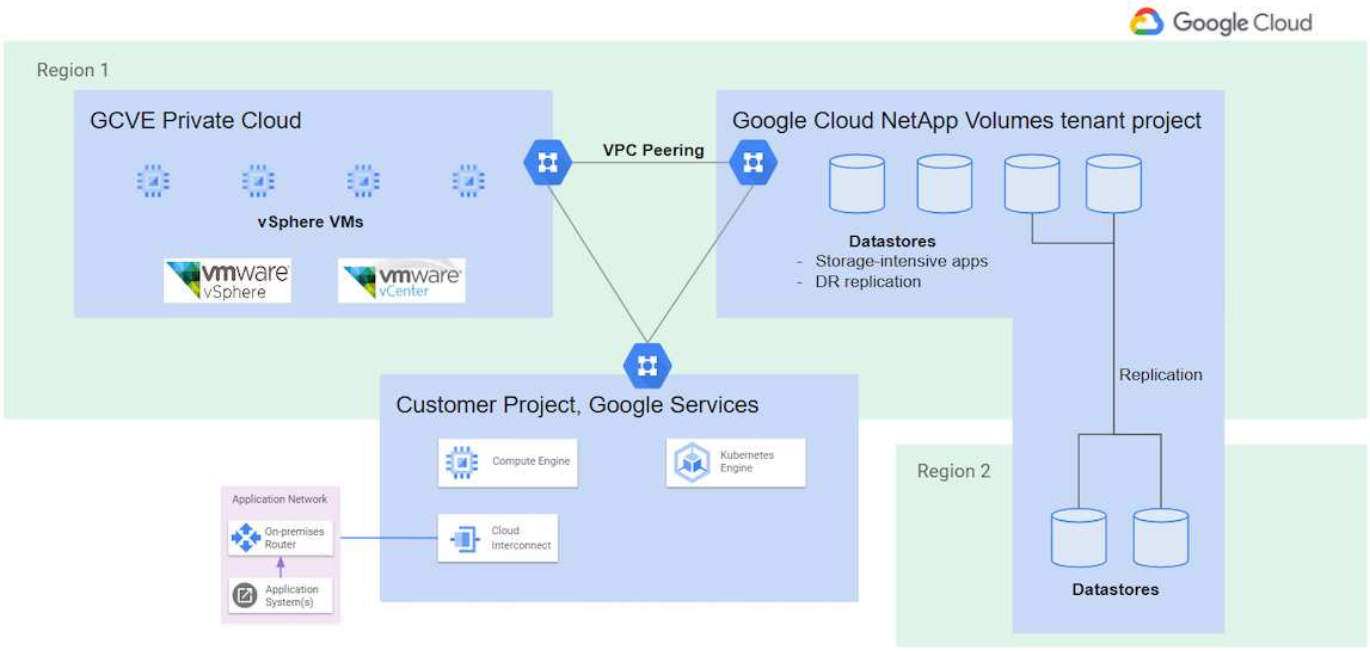
For more information, please visit Google Cloud NetApp Volumes (NetApp Volumes) as Guest Connected Storage

Google Cloud VMware Engine Supplemental NFS Datastore with Google Cloud NetApp Volumes

Customers can expand storage capacity on Google Cloud VMware Engine using NFS supplemental datastore with Google Cloud NetApp Volumes.

Overview

Customers that requires additional storage capacity on their Google Cloud VMware Engine (GCVE) environment can utilize Netapp Cloud Volume Service to mount as supplemental NFS datastore.
Storing data on Google Cloud NetApp Volumes allows customers to replicate between regions to protect from diaster.



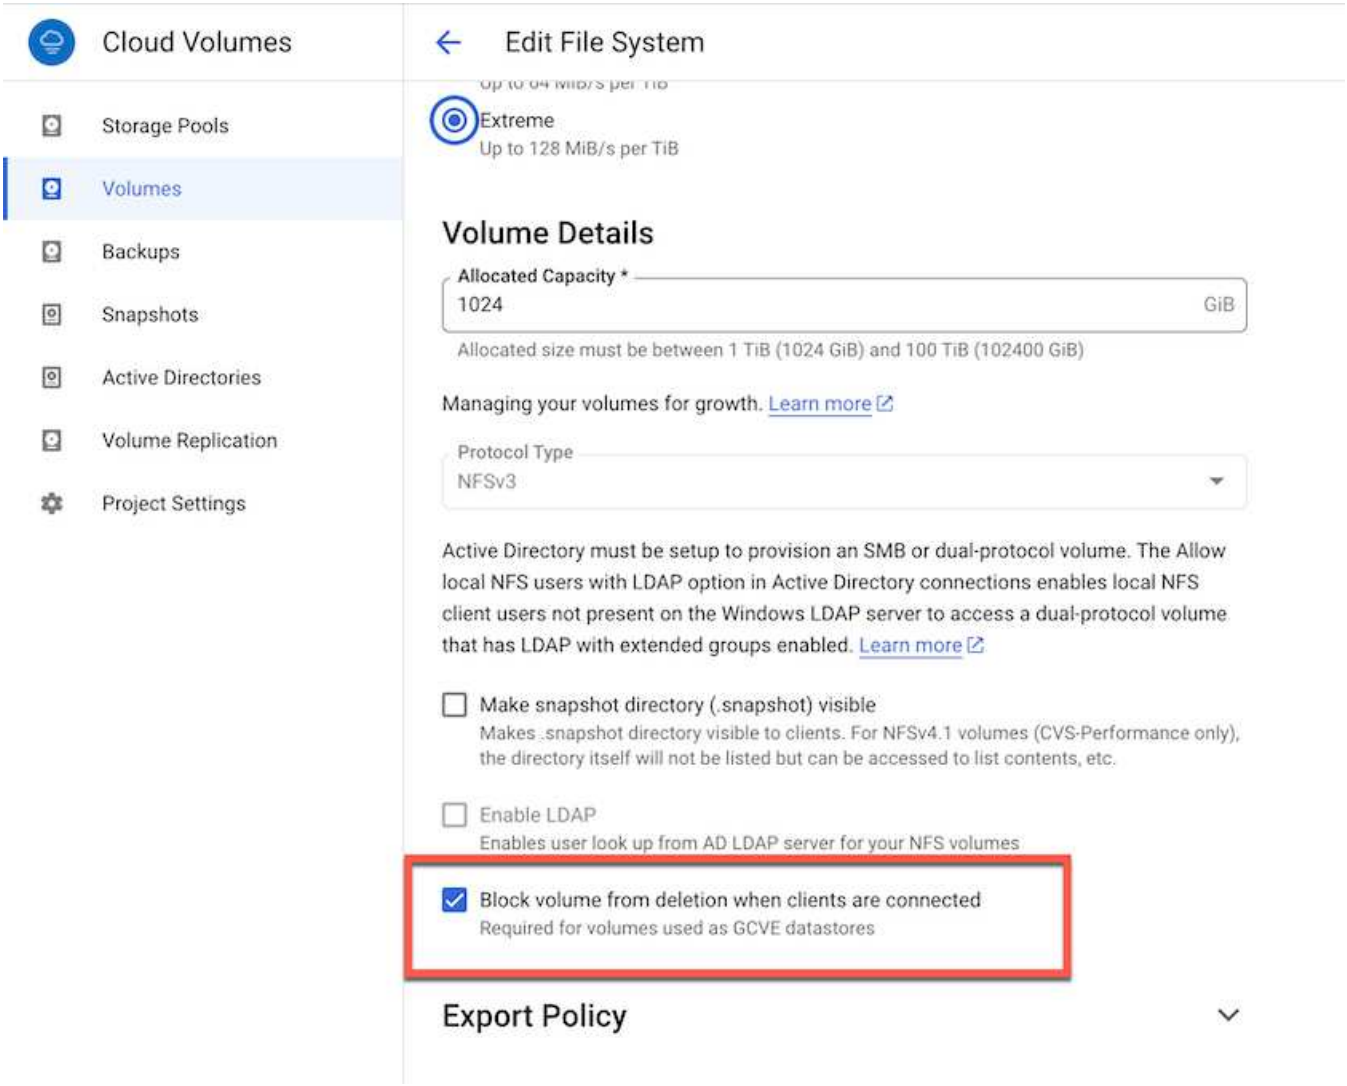
Deployment steps to mount NFS datastore from Google Cloud NetApp Volumes (NetApp Volumes) on GCVE

▼ Provision NetApp Volumes-Performance Volume

The Google Cloud NetApp Volumes volume can be either provisioned by
Using Google Cloud Console
Using Google Cloud CLI or API

▼ Mark that NetApp Volumes volume as non-deletable

To avoid accidental deletion of volume while VM is running, ensure the volume is marked as non-deletable as shown in screenshot below.



For more info, please refer Creating NFS Volume documentation.

▼ **Ensure Private Connection on GCVE exists for NetApp Volumes Tenant VPC.**

To mount NFS Datastore, there should be a private connection exists between GCVE and NetApp Volumes project.
For more info, please refer How to setup Private Service Access

▼ **Mount NFS datastore**

For instructions on how to mount NFS datastore on GCVE, please refer How to create NFS datastore with NetApp Volumes

As vSphere hosts are managed by Google, you don't have access to install NFS vSphere API for Array Integration (VAAI) vSphere Installation Bundle (VIB).
If you need support for Virtual Volumes (vVol), please let us know.
If you like to use Jumbo Frames, please refer Maximum supported MTU sizes on GCP

Savings with Google Cloud NetApp Volumes

To learn more about your potential saving with Google Cloud NetApp Volumes for your storage demands on GCVE, please check NetApp ROI Calculator

Reference Links

- Google Blog - How to use NetApp Volumes as datastores for Google Cloud VMware Engine
- NetApp Blog - A better way to migrate your storage-rich apps to Google Cloud

NetApp Storage Options for GCP

GCP supports guest connected NetApp storage with Cloud Volumes ONTAP (CVO) or Google Cloud NetApp Volumes (NetApp Volumes).

Cloud Volumes ONTAP (CVO)

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

Cloud Volumes ONTAP (CVO) as guest connected storage

▼ **Deploy Cloud Volumes ONTAP in Google Cloud (Do It Yourself)**

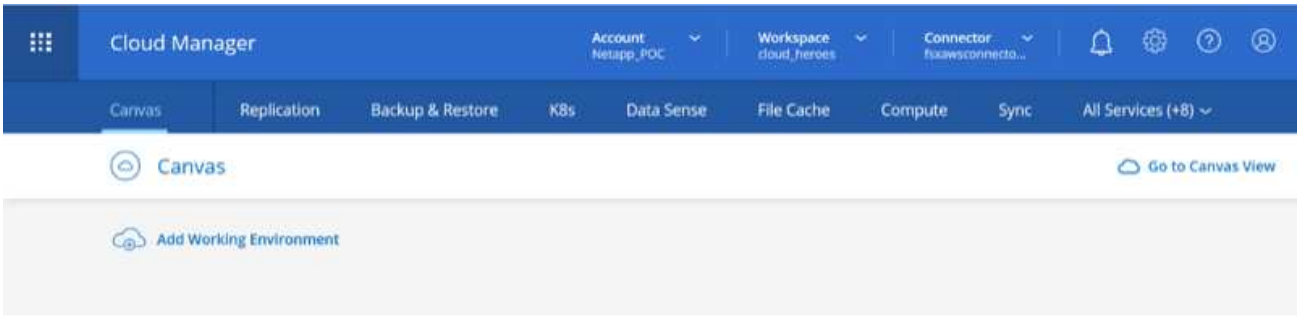
Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the GCVE private cloud environment. The volumes can also be mounted on the Linux client and on Windows client and LUNS can be accessed on Linux or Windows clients as block devices when mounted over iSCSI because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be set up in a few simple steps.

To replicate volumes from an on-premises environment to the cloud for disaster recovery or migration purposes, establish network connectivity to Google Cloud, either using a site-to-site VPN or Cloud Interconnect. Replicating data from on-premises to Cloud Volumes ONTAP is outside the scope of this document. To replicate data between on-premises and Cloud Volumes ONTAP systems, see xref:./vmware/Setting up data replication between systems.

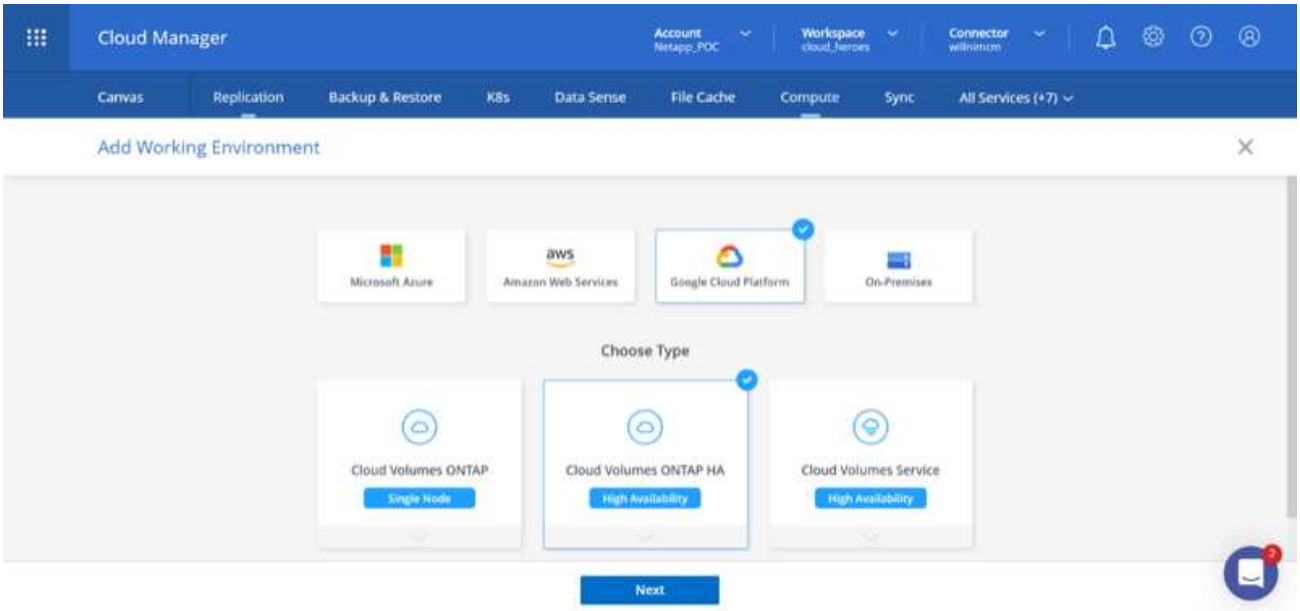


Use Cloud Volumes ONTAP sizer to accurately size the Cloud Volumes ONTAP instances. Also monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

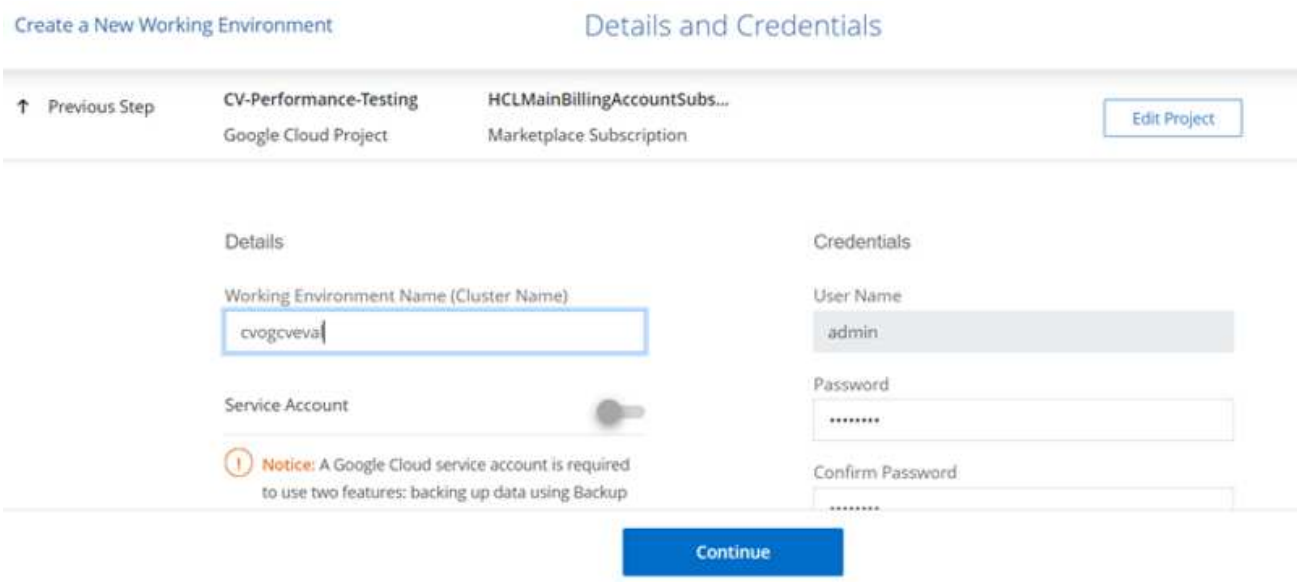
1. Log in to NetApp Cloud Central—the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.



2. On the Cloud Manager Canvas tab, click Add a Working Environment and then select Google Cloud Platform as the cloud and the type of the system configuration. Then, click Next.



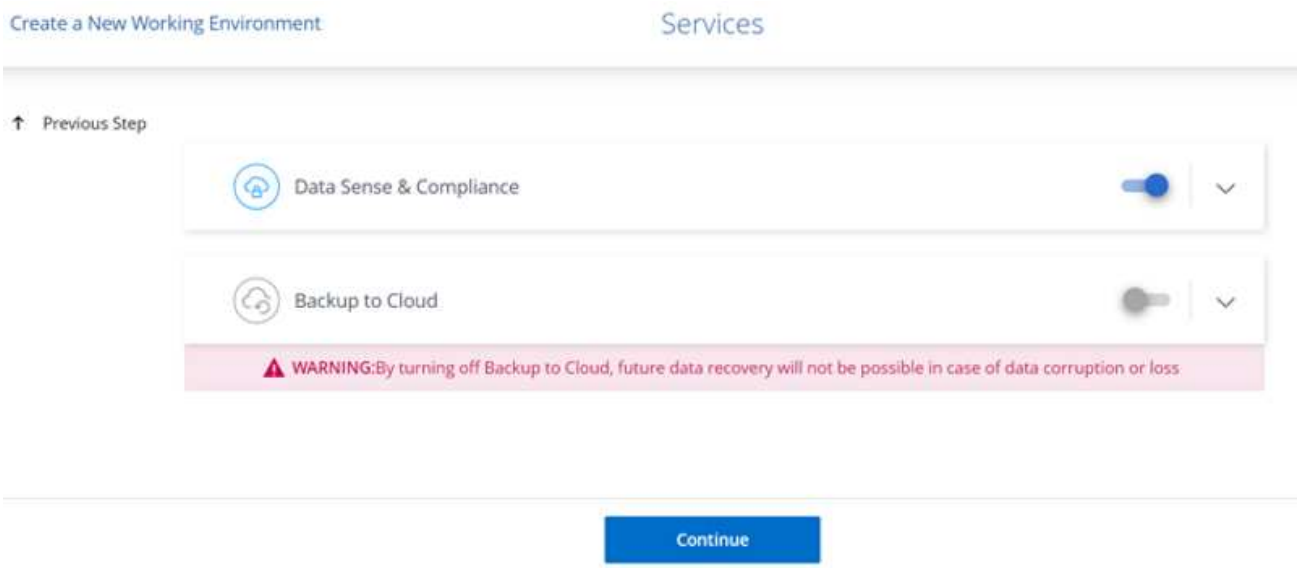
3. Provide the details of the environment to be created including the environment name and admin credentials. After you are done, click Continue.



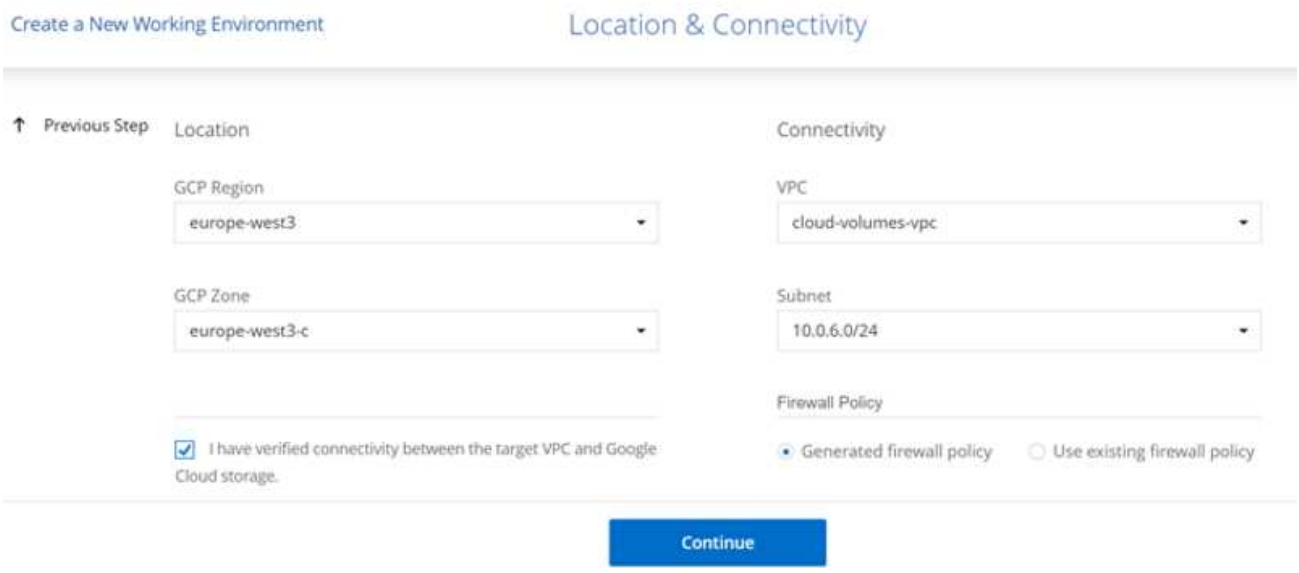
4. Select or deselect the add-on services for Cloud Volumes ONTAP deployment, including Data Sense &

Compliance or Backup to Cloud. Then, click Continue.

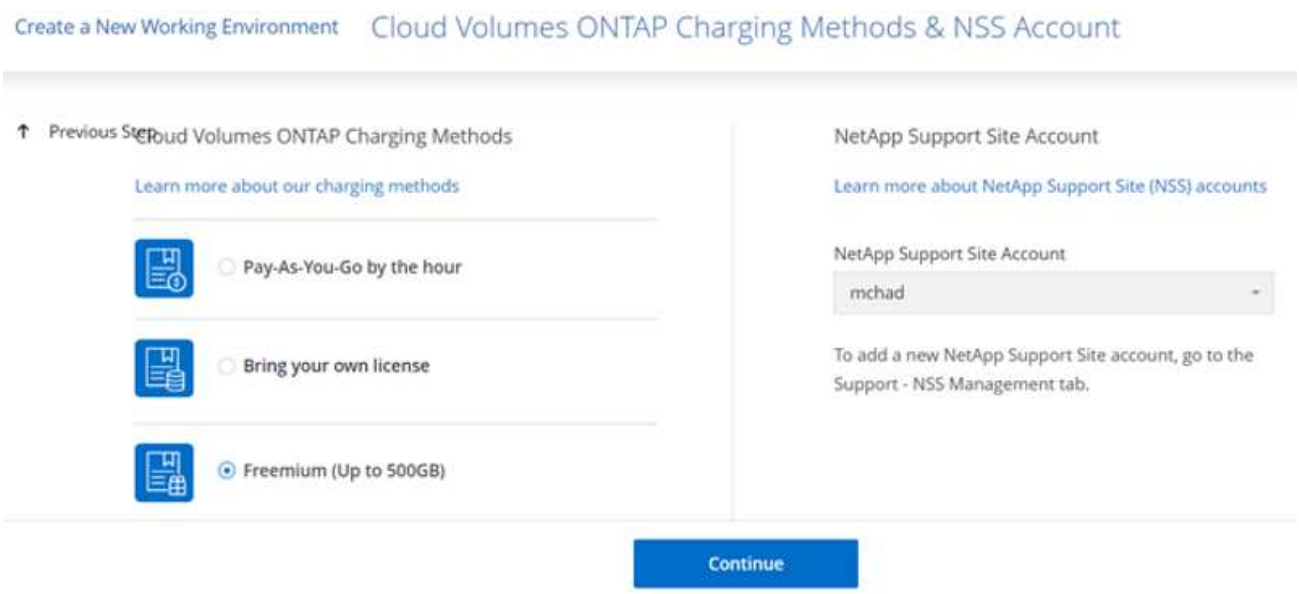
HINT: A verification pop-up message will be displayed when deactivating add-on services.
Add-on services can be added/removed after CVO deployment, consider to deselect them if not needed from the beginning to avoid costs.



5. Select a location, choose a firewall policy, and select the checkbox to confirm network connectivity to Google Cloud storage.

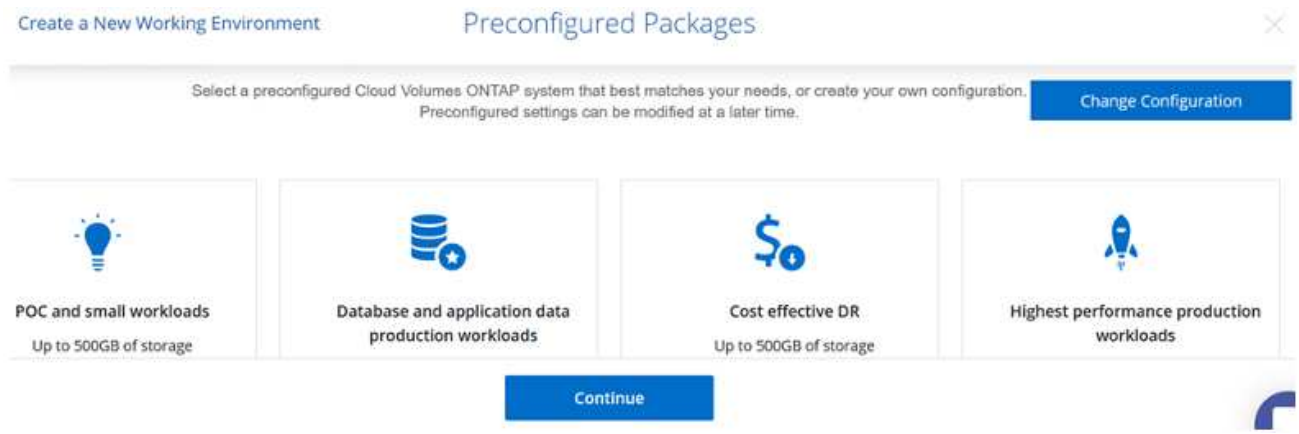


6. Select the license option: Pay-As-You-Go or BYOL for using existing license. In this example, Freemium option is used. Then, click on Continue.

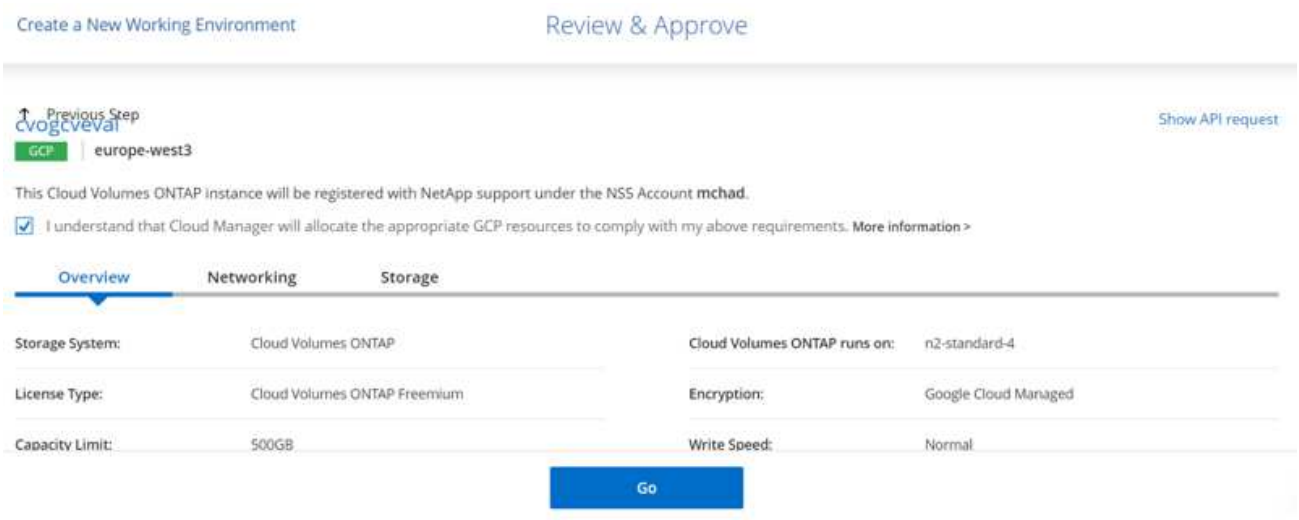


7. Select between several preconfigured packages available based on the type of workload that will be deployed on the VMs running on VMware cloud on AWS SDDC.

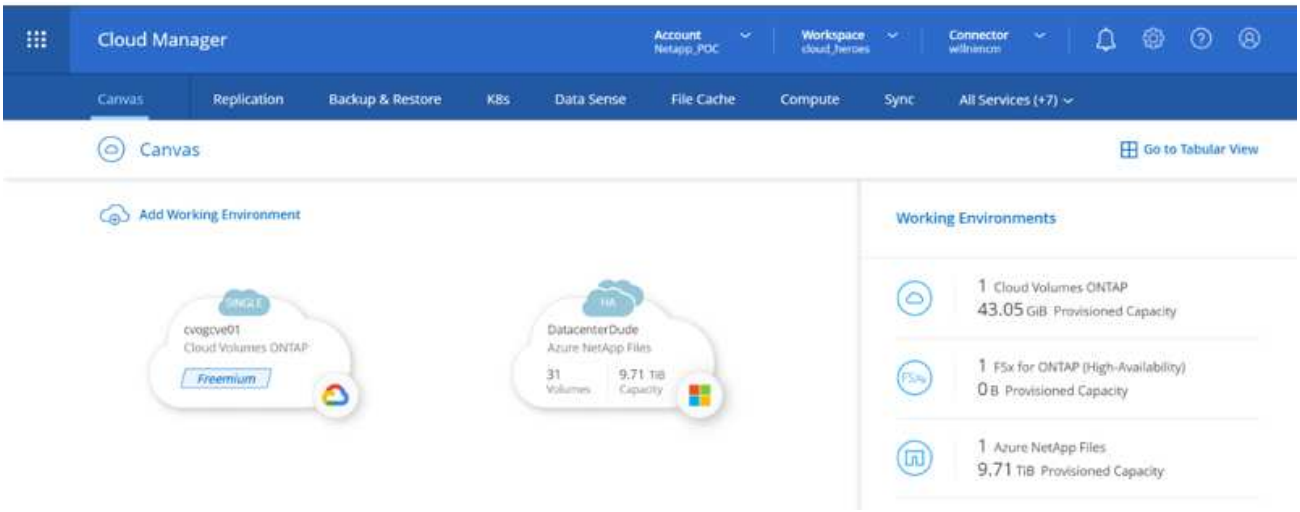
   HINT: Hoover your mouse over the tiles for details or customize CVO components and ONTAP version by clicking on Change Configuration.



8. On the Review & Approve page, review and confirm the selections.To create the Cloud Volumes ONTAP instance, click Go.



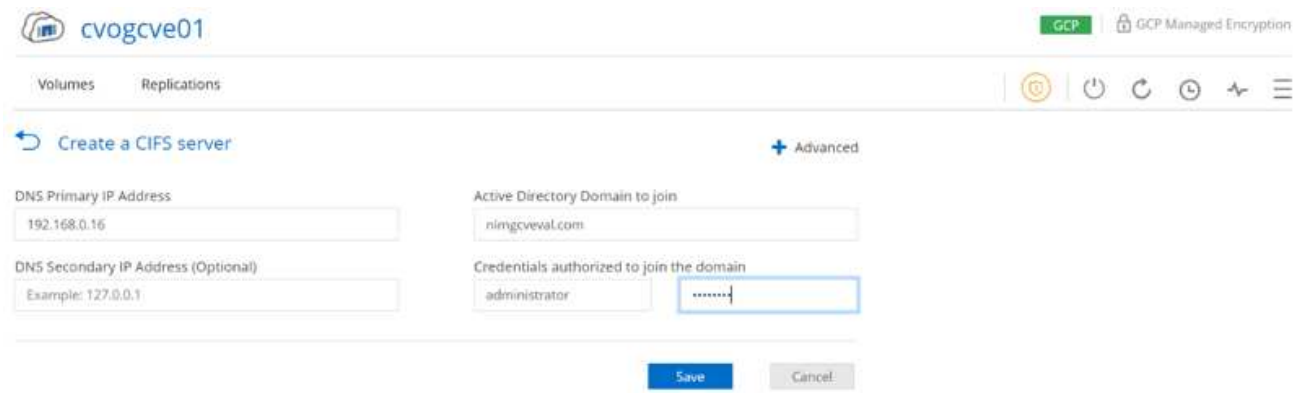
9. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.



▼ Additional configurations for SMB volumes

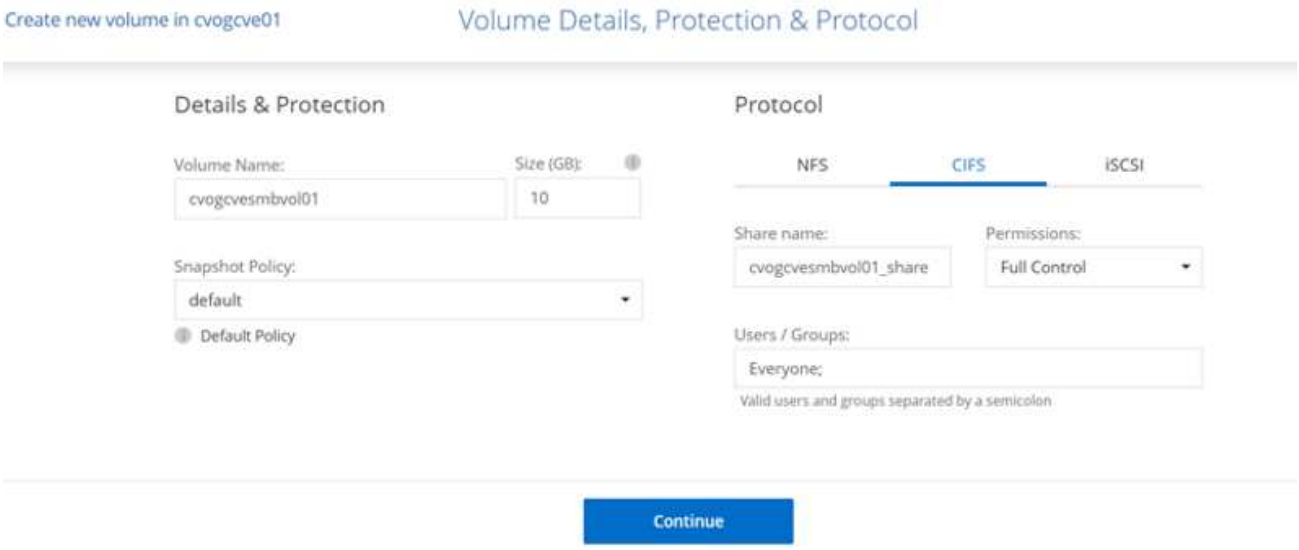
1. After the working environment is ready, make sure the CIFS server is configured with the appropriate DNS and Active Directory configuration parameters. This step is required before you can create the SMB volume.

   HINT: Click on the Menu Icon (º), select Advanced to display more options and select CIFS setup.



2. Creating the SMB volume is an easy process. At Canvas, double-click the Cloud Volumes ONTAP

working environment to create and manage volumes and click on the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, CIFS/SMB is selected as the protocol.



3. After the volume is provisioned, it will be availabe under the Volumes pane. Because a CIFS share is provisioned, give your users or groups permission to the files and folders and verify that those users can access the share and create a file. This step is not required if the volume is replicated from an on-premises environment because the file and folder permissions are all retained as part of SnapMirror replication.

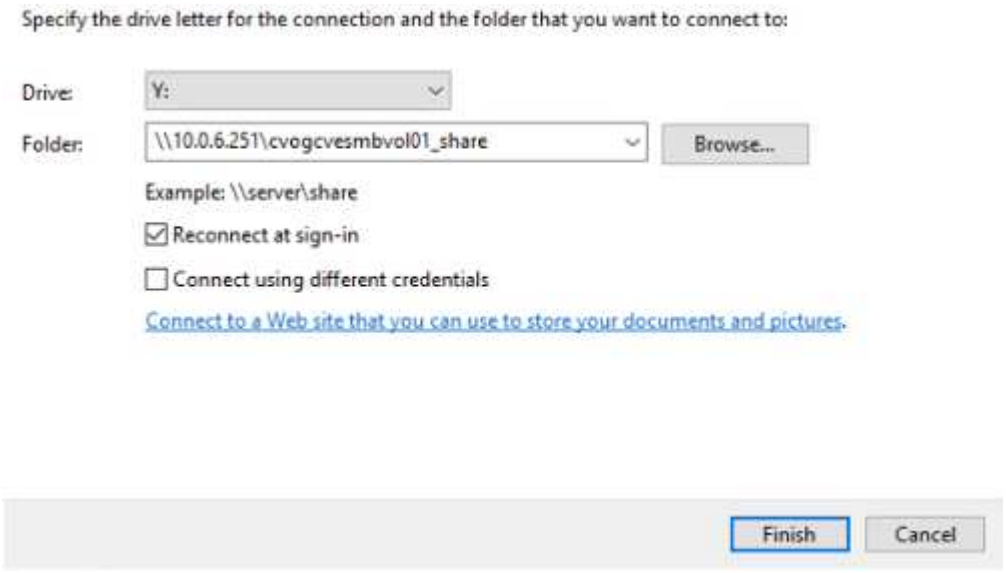
   HINT: Click on the volume menu (º) to display its options.



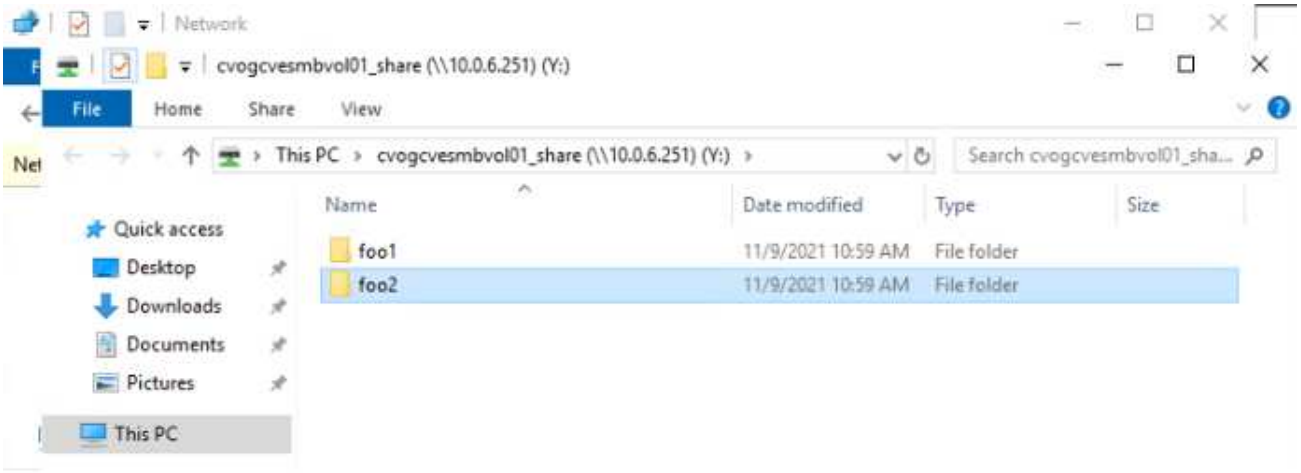
4. After the volume is created, use the mount command to display the volume connection instructions, then connect to the share from the VMs on Google Cloud VMware Engine.



5. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on the Google Cloud VMware Engine.



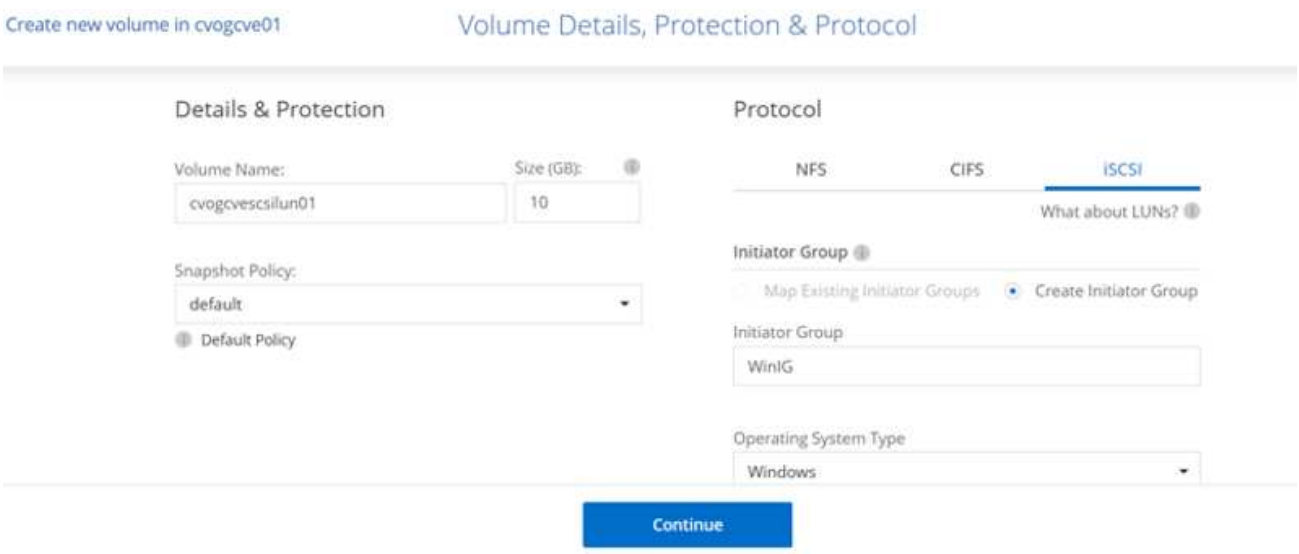
Once mapped, it can be easily accessed, and the NTFS permissions can be set accordingly.

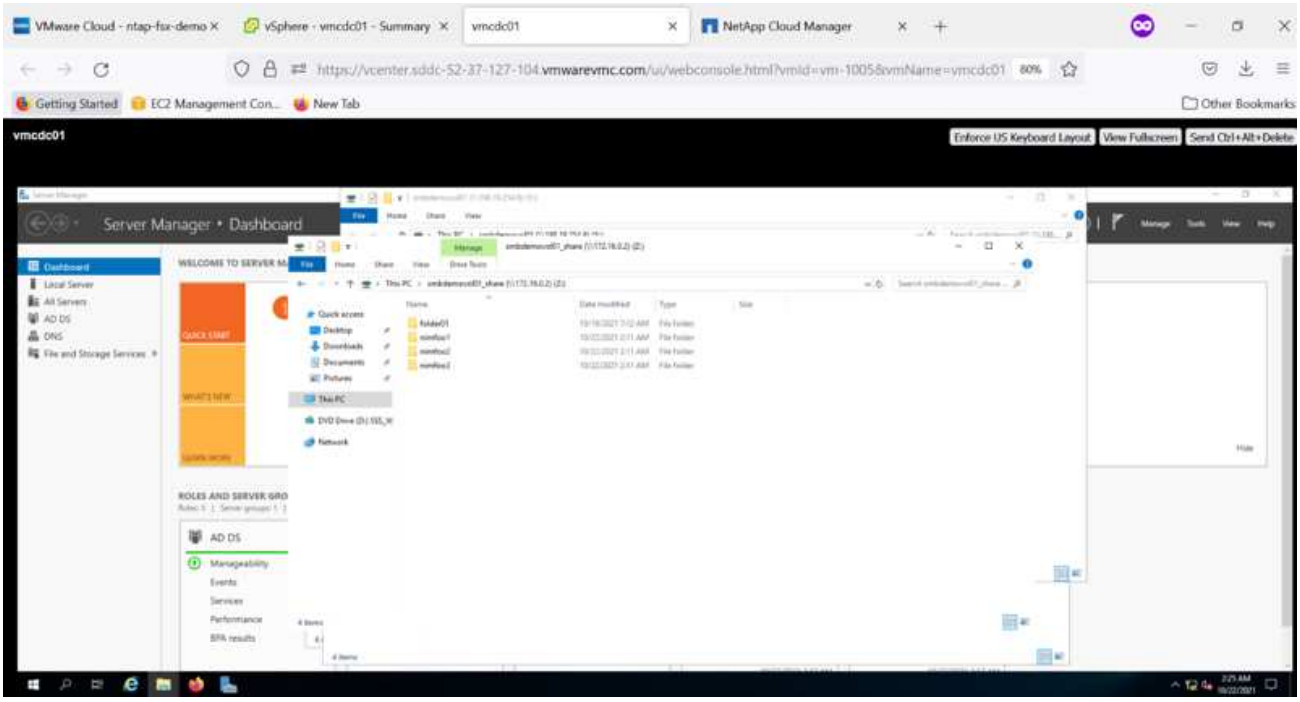


▼ **Connect the LUN on Cloud Volumes ONTAP to a host**

To connect the cloud volumes ONTAP LUN to a host, complete the following steps:

1. On the Canvas page, double-click the Cloud Volumes ONTAP working environment to create and manage volumes.
2. Click Add Volume > New Volume and select iSCSI and click Create Initiator Group. Click Continue.





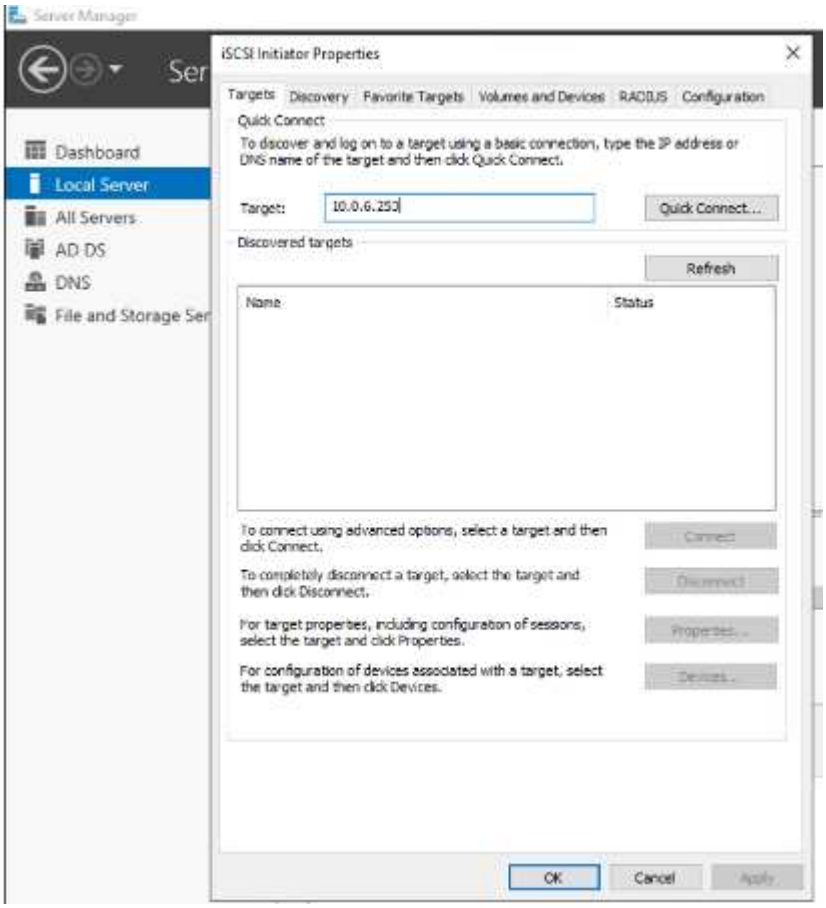
3. After the volume is provisioned, select the volume menu (º), and then click Target iQN. To copy the iSCSI Qualified Name (iQN), click Copy. Set up an iSCSI connection from the host to the LUN.

To accomplish the same for the host residing on Google Cloud VMware Engine:

a. RDP to the VM hosted on Google Cloud VMware Engine.
b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.
c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.
d. From the Targets tab, select the target discovered and then click Log on or Connect.
e. Select Enable multipath, and then select Automatically Restore This Connection When the Computer

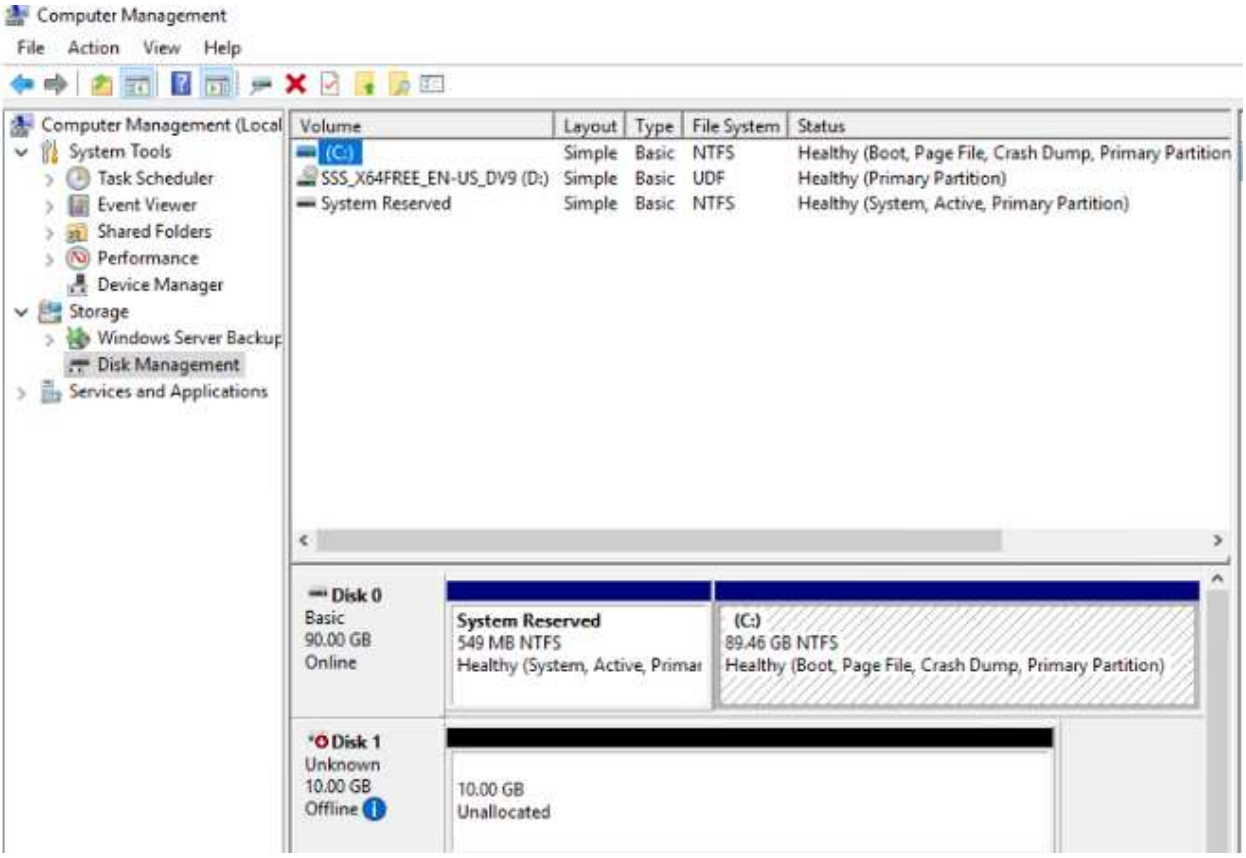
Starts or Add This Connection to the List of Favorite Targets. Click Advanced.

> The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



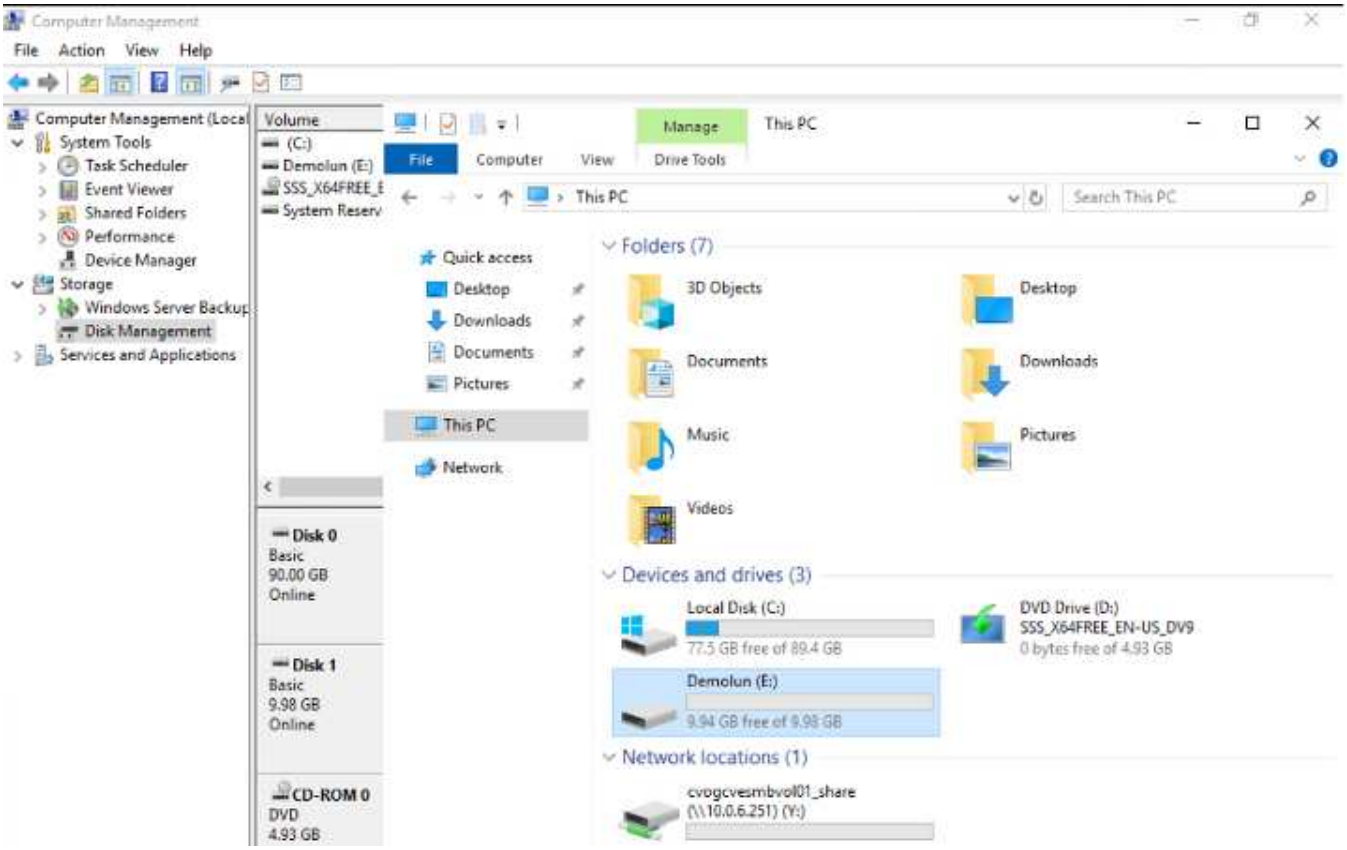
LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
2. Expand the Storage node in the navigation tree.
3. Click Disk Management.
4. Click Action > Rescan Disks.



When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

5. Start Windows Disk Management.
6. Right-click the LUN, and then select the required disk or partition type.
7. Follow the instructions in the wizard. In this example, drive F: is mounted.



On the Linux clients, ensure the iSCSI daemon is running. Once the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration with Ubuntu as an example here. To verify, run lsblk cmd from the shell.

```
nlyaz@nlmubu01:~$ lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
loop0    7:0    0 55.4M  1 loop /snap/core18/2128
loop1    7:1    0  219M  1 loop /snap/gnome-3-34-1804/72
loop2    7:2    0 65.1M  1 loop /snap/gtk-common-themes/1515
loop3    7:3    0   51M  1 loop /snap/snap-store/547
loop4    7:4    0 32.3M  1 loop /snap/snapd/12704
loop5    7:5    0 32.5M  1 loop /snap/snapd/13640
loop6    7:6    0 55.5M  1 loop /snap/core18/2246
loop7    7:7    0    4K  1 loop /snap/bare/5
loop8    7:8    0 65.2M  1 loop /snap/gtk-common-themes/1519
sda      8:0    0   16G  0 disk
├─sda1   8:1    0  512M  0 part /boot/efi
├─sda2   8:2    0    1K  0 part
└─sda5   8:5    0 15.5G  0 part /
sdb      8:16   0    1G  0 disk
```

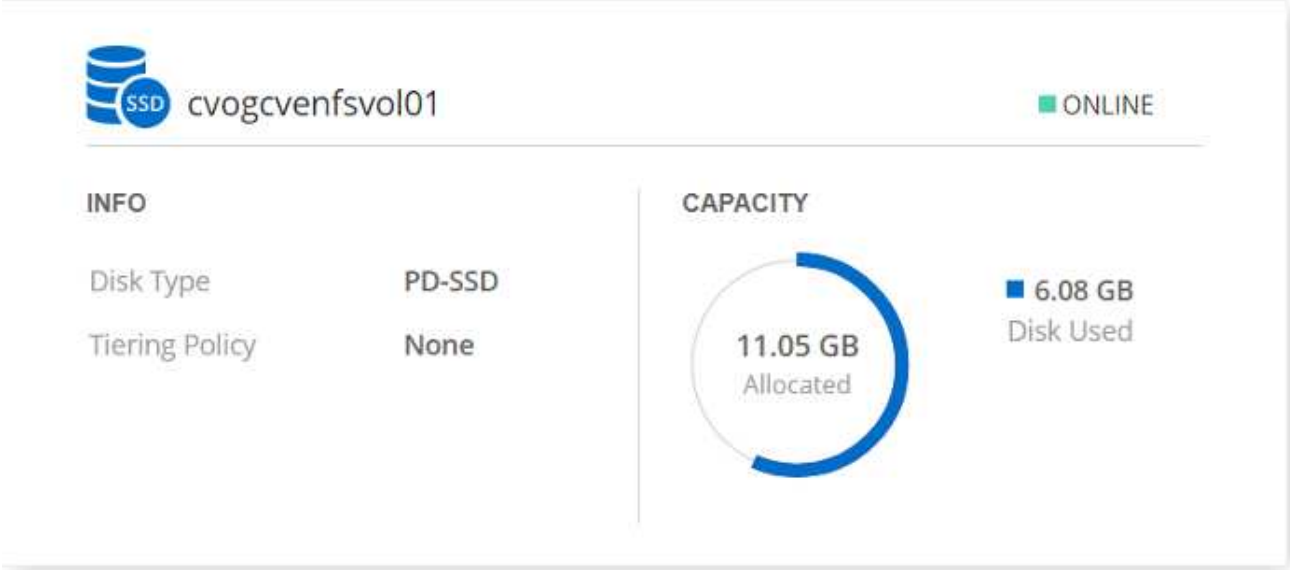
```
niyaz@nimubu01:~$ df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            1.9G     0  1.9G   0% /dev
tmpfs           394M  1.5M  392M   1% /run
/dev/sda5        16G  7.6G  6.9G  53% /
tmpfs           2.0G     0  2.0G   0% /dev/shm
tmpfs           5.0M     0  5.0M   0% /run/lock
tmpfs           2.0G     0  2.0G   0% /sys/fs/cgroup
/dev/loop1      219M  219M     0 100% /snap/gnome-3-34-1804/72
/dev/loop2       66M   66M     0 100% /snap/gtk-common-themes/1515
/dev/loop3       51M   51M     0 100% /snap/snap-store/547
/dev/loop0       56M   56M     0 100% /snap/core18/2128
/dev/loop4       33M   33M     0 100% /snap/snapd/12704
/dev/sda1       511M  4.0K  511M   1% /boot/efi
tmpfs           394M   64K  394M   1% /run/user/1000
/dev/loop5       33M   33M     0 100% /snap/snapd/13640
/dev/loop6       56M   56M     0 100% /snap/core18/2246
/dev/loop7      128K  128K     0 100% /snap/bare/5
/dev/loop8       66M   66M     0 100% /snap/gtk-common-themes/1519
/dev/sdb        976M  2.6M  907M   1% /mnt
```

▼ **Mount Cloud Volumes ONTAP NFS volume on Linux client**

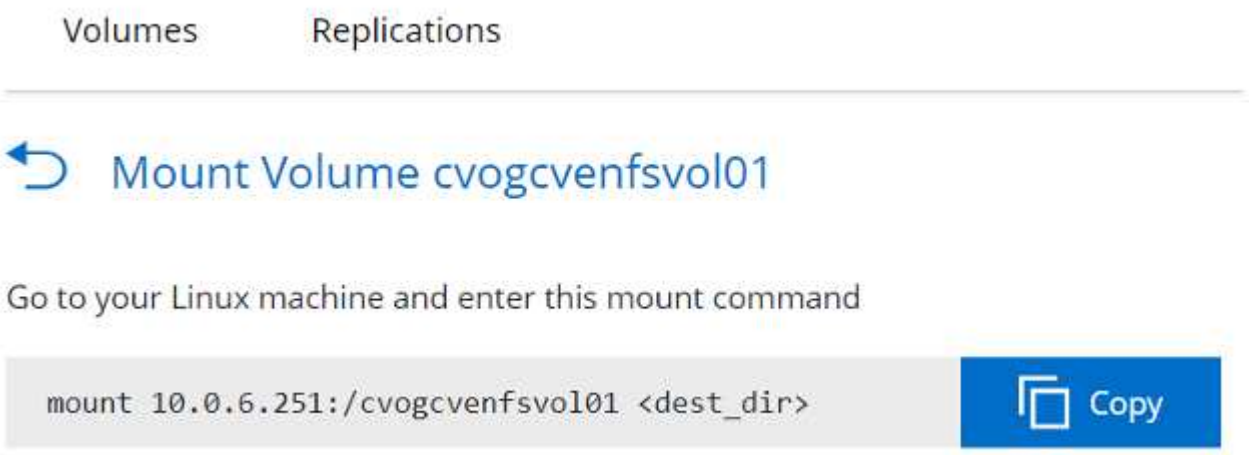
To mount the Cloud Volumes ONTAP (DIY) file system from VMs within Google Cloud VMware Engine, follow the below steps:

Provision the volume following the below steps

1. In the Volumes tab, click Create New Volume.
2. On the Create New Volume page, select a volume type:



3. In the Volumes tab, place your mouse cursor over the volume, select the menu icon (º), and then click Mount Command.



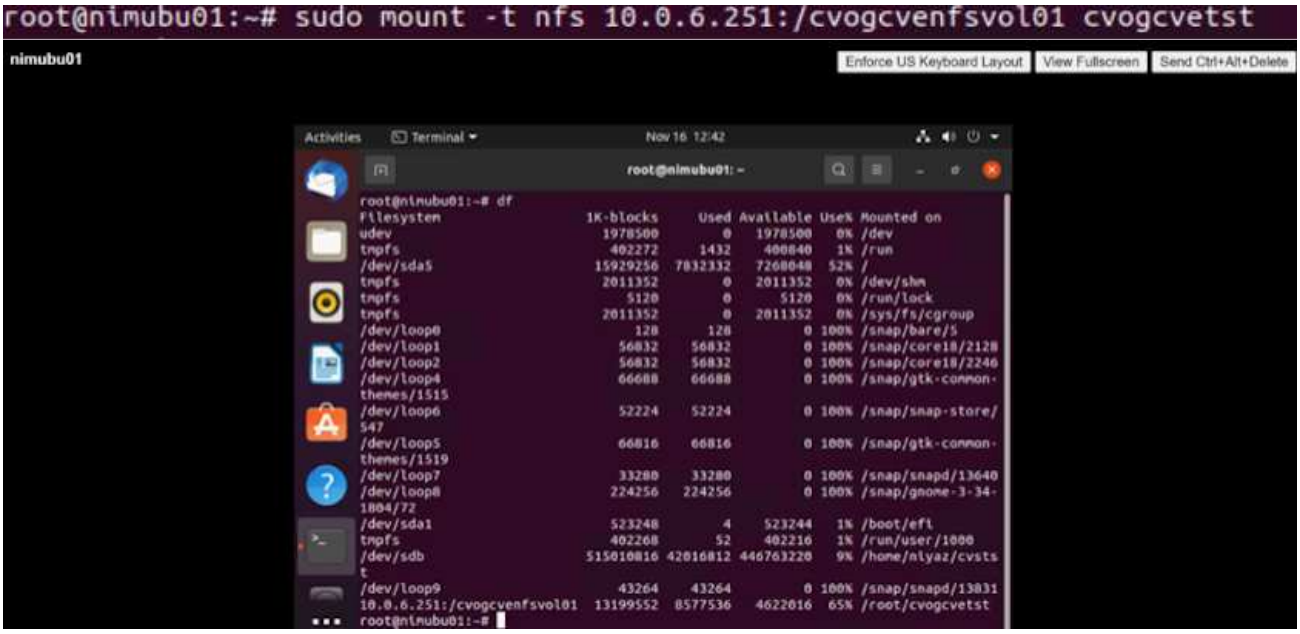
4. Click Copy.
5. Connect to the designated Linux instance.
6. Open a terminal on the instance using secure shell (SSH) and log in with the appropriate credentials.
7. Make a directory for the volume's mount point with the following command.

```
$ sudo mkdir /cvogcvetst
```

```
root@nimubu01:~# sudo mkdir cvogcvetst
```

8. Mount the Cloud Volumes ONTAP NFS volume to the directory that is created in the previous step.

```
sudo mount 10.0.6.251:/cvogcvenfsvol01 /cvogcvetst
```



Google Cloud NetApp Volumes (NetApp Volumes)

Google Cloud NetApp Volumes (NetApp Volumes) is a complete portfolio of data services to deliver advanced cloud solutions. NetApp Volumes supports multiple file access protocols for major cloud providers (NFS and SMB support).

Other benefits and features include: data protection and restore with Snapshot; special features to replicate, sync and migrate data destinations on-prem or in the cloud; and consistent high performance at the level of a dedicated flash storage system.

Google Cloud NetApp Volumes (NetApp Volumes) as guest connected storage

▼ Configure NetApp Volumes with VMware Engine

Google Cloud NetApp Volumes shares can be mounted from VMs that are created in the VMware Engine environment. The volumes can also be mounted on the Linux client and mapped on the Windows client because Google Cloud NetApp Volumes supports SMB and NFS protocols. Google Cloud NetApp Volumes volumes can be set up in simple steps.

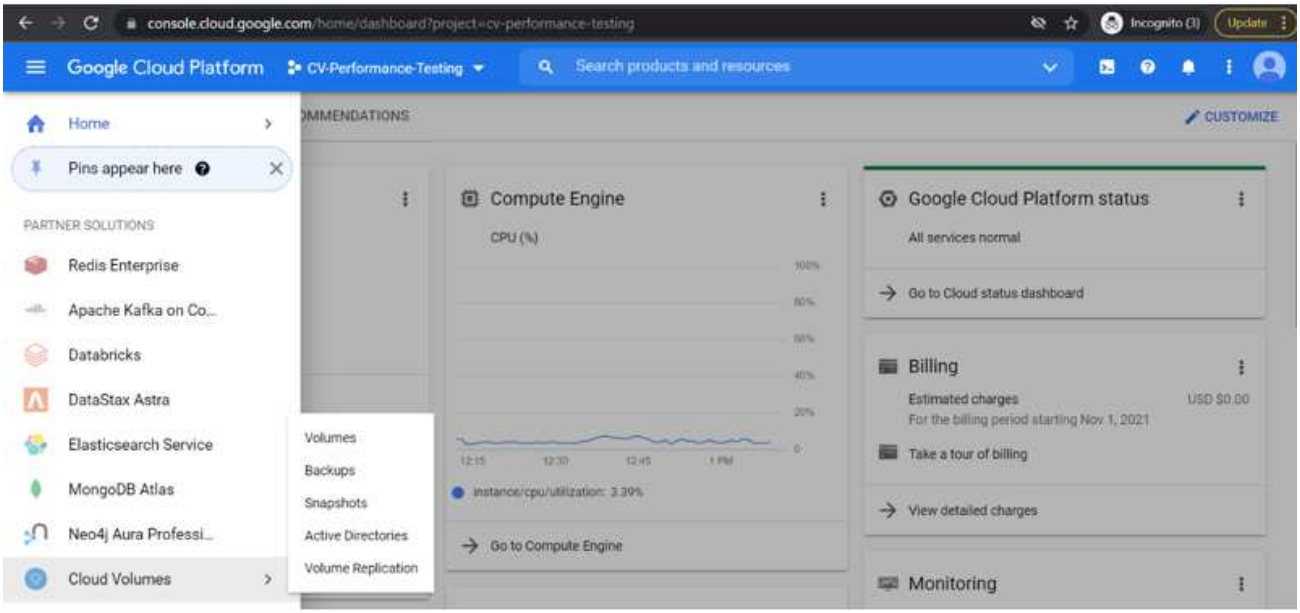
Google Cloud NetApp Volumes and Google Cloud VMware Engine private cloud must be in the same region.

To purchase, enable and configure Google Cloud NetApp Volumes for Google Cloud from the Google Cloud Marketplace, follow this detailed guide.

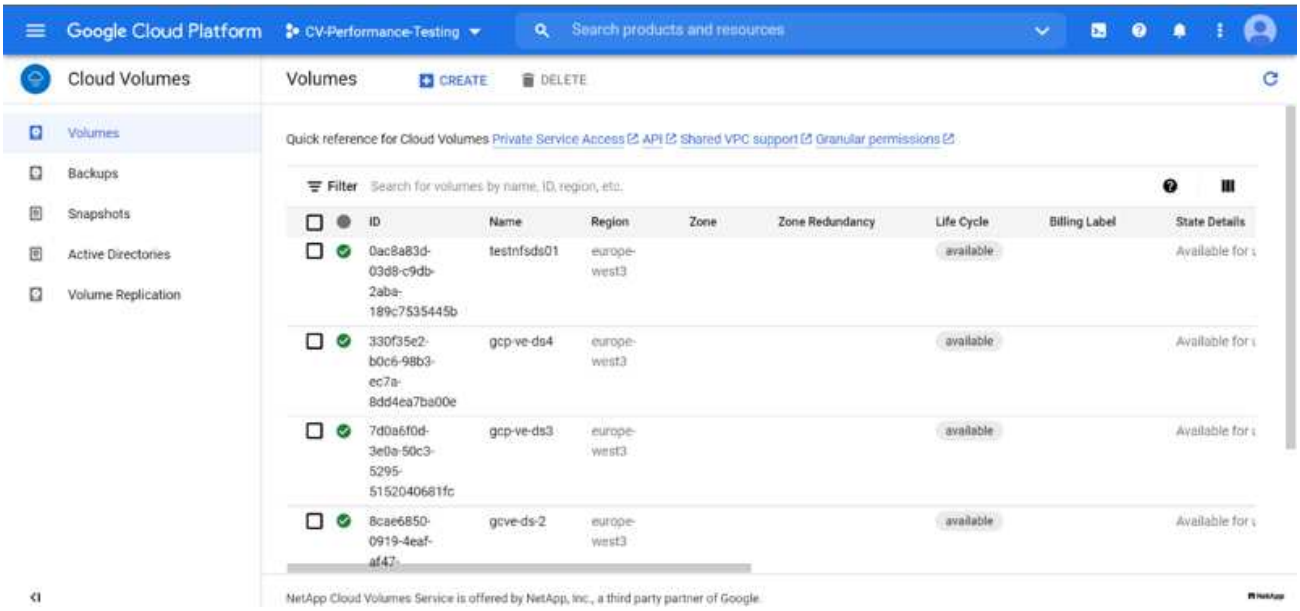
▼ Create a NetApp Volumes NFS volume to GCVE private cloud

To create and mount NFS volumes, complete the following steps:

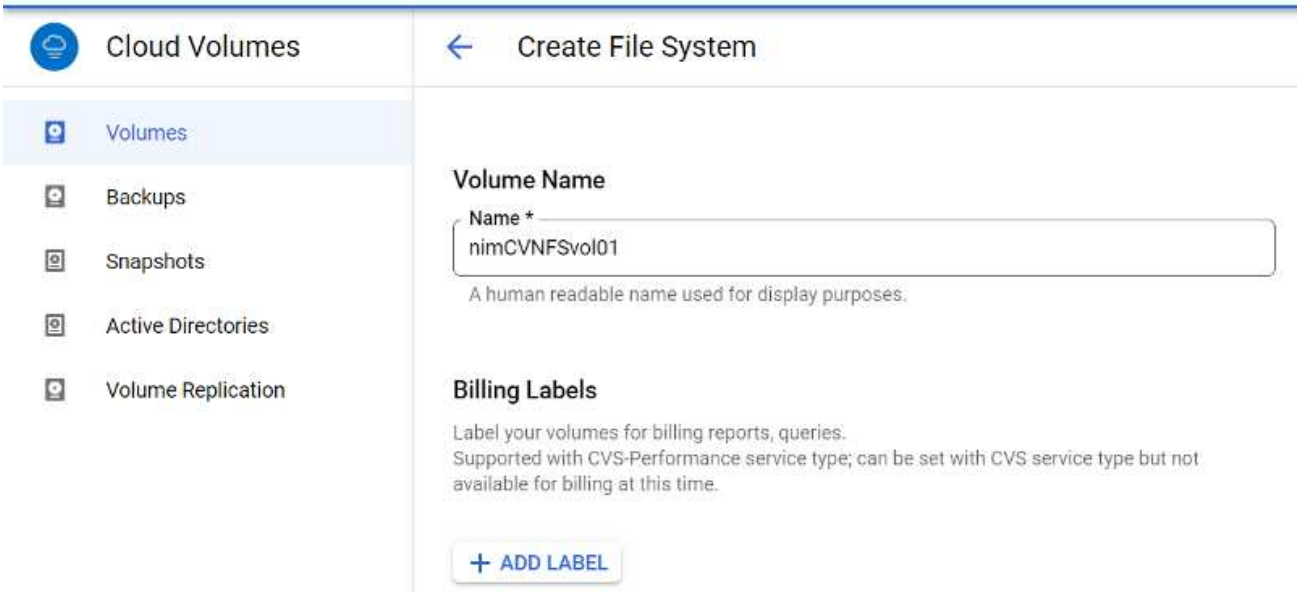
1. Access Cloud Volumes from Partner Solutions within the Google cloud console.



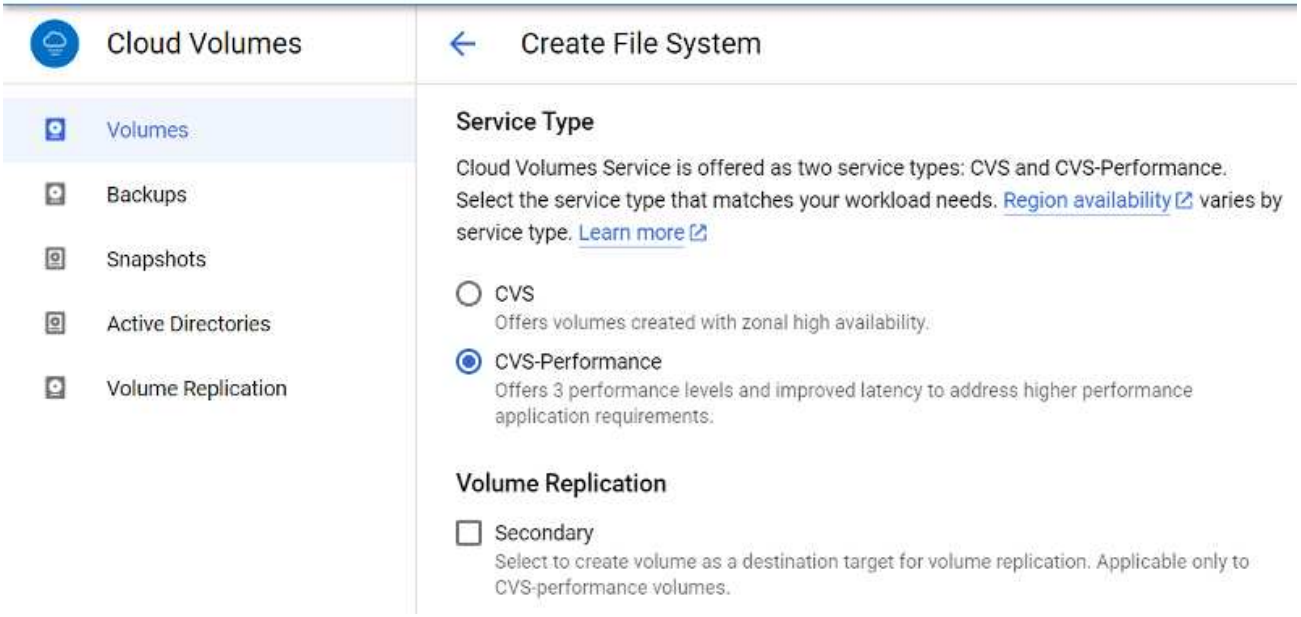
2. In the Cloud Volumes Console, go to the Volumes page and click Create.



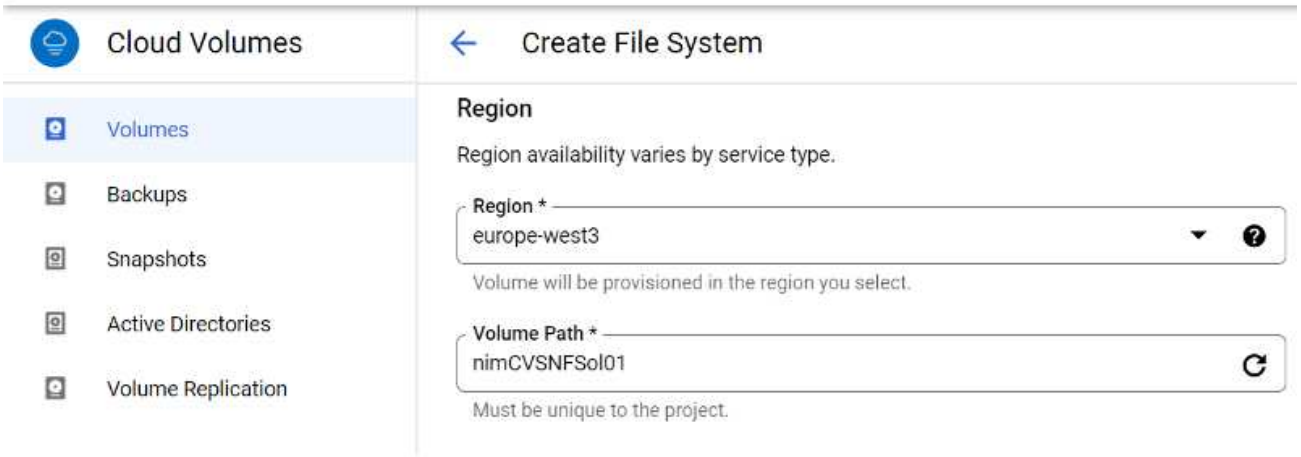
3. On the Create File System page, specify the volume name and billing labels as required for chargeback mechanisms.



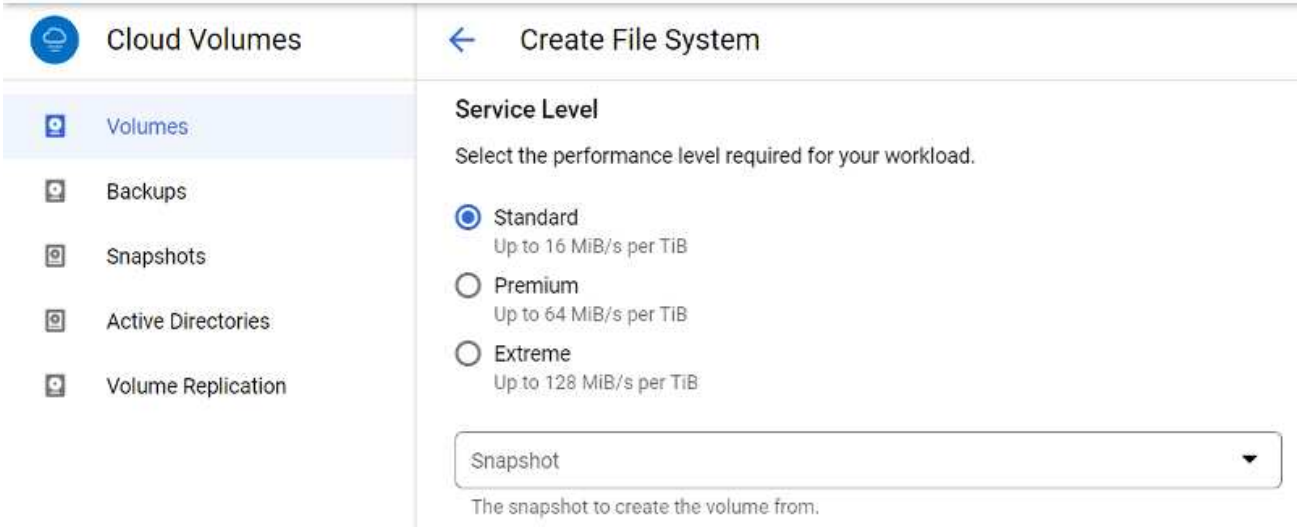
4. Select the appropriate service. For GCVE, choose NetApp Volumes-Performance and desired service level for improved latency and higher performance based on the application workload requirements.



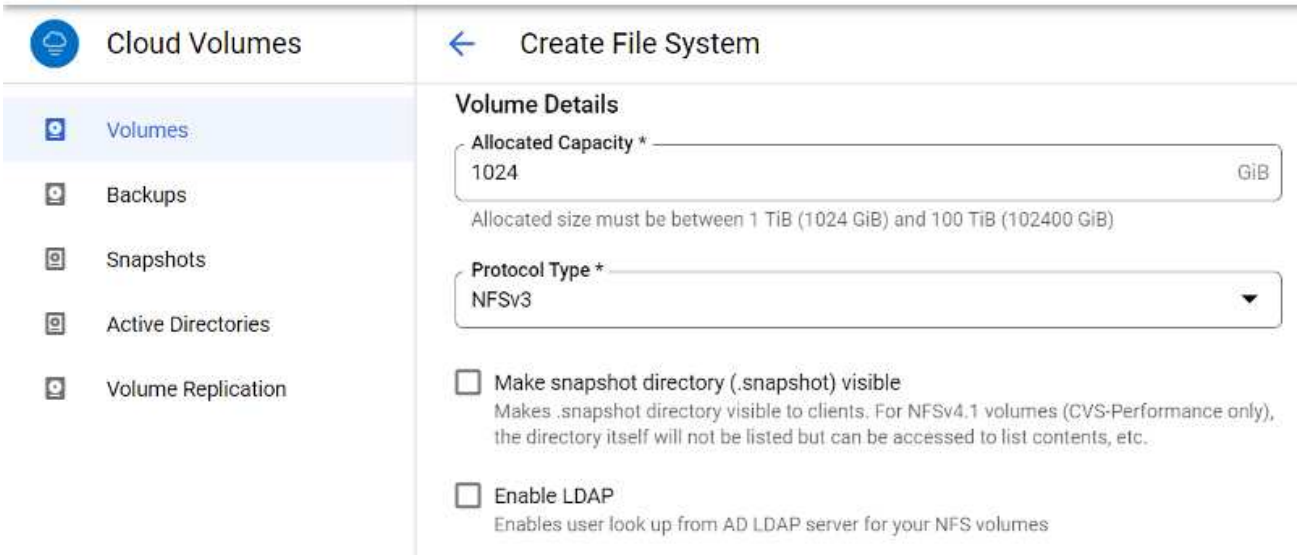
5. Specify the Google Cloud region for the volume and volume path (The volume path must be unique across all of cloud volumes in the project)



6. Select the level of performance for the volume.

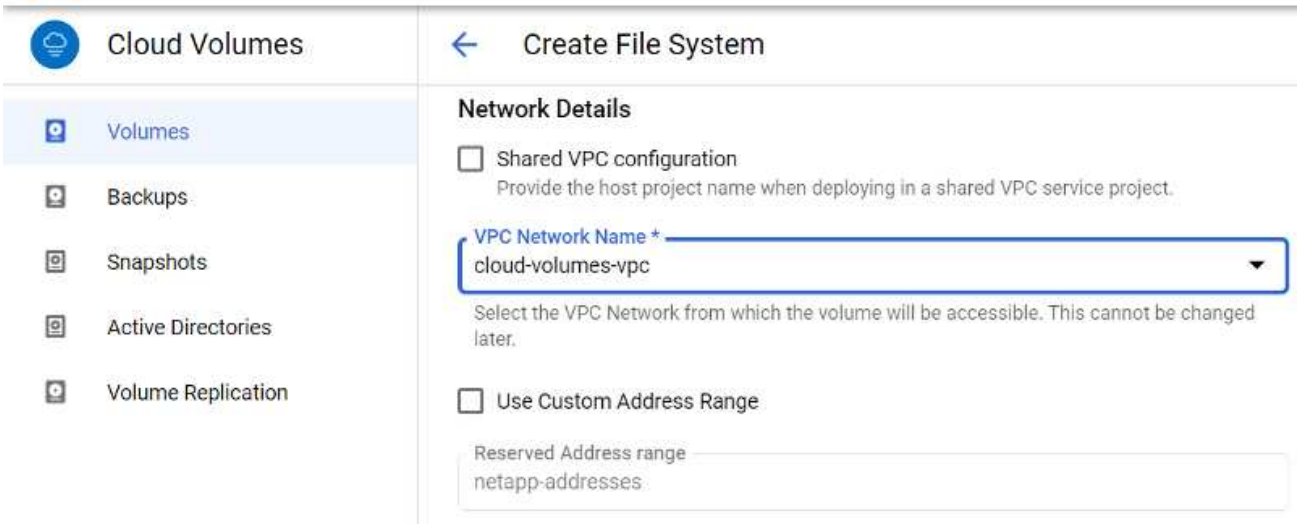


7. Specify the size of the volume and the protocol type. In this testing, NFSv3 is used.



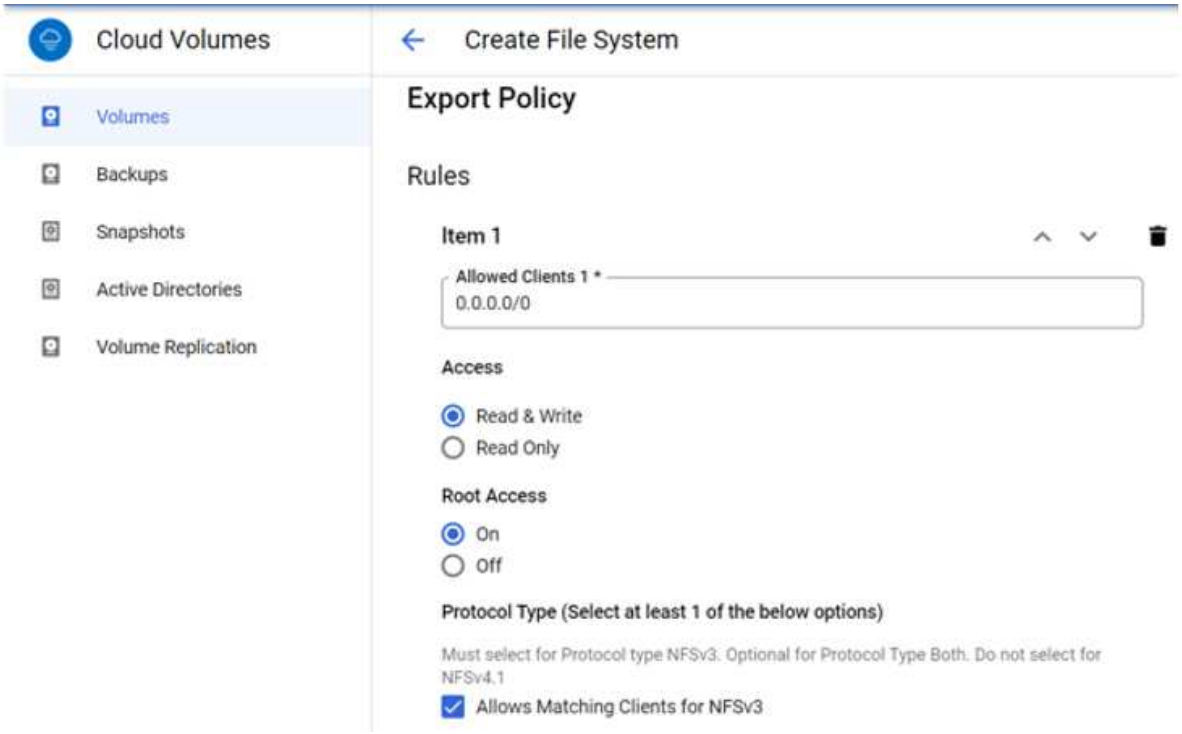
8. In this step, select the VPC Network from which the volume will be accessible. Ensure VPC peering is in place.

   HINT: If VPC peering has not been done, a pop-up button will be displayed to guide you through the peering commands. Open a Cloud Shell session and execute the appropriate commands to peer your VPC with Google Cloud NetApp Volumes producer. In case you decide to prepare VPC peering in beforehand, refer to these instructions.



9. Manage the Export policy rules by adding the appropriate rules and Select the checkbox for the corresponding NFS version.

   Note: Access to NFS volumes won't be possible unless an export policy is added.



10. Click Save to create the volume.



▼ **Mounting NFS exports to VMs running on VMware Engine**

Before preparing to mount the NFS volume, ensure the peering status of private connection is listed as Active. Once status is Active, use the mount command.

To mount an NFS volume, do the following:

1. In the Cloud Console, go to Cloud Volumes > Volumes.
2. Go to the Volumes page
3. Click the NFS volume for which you want to mount NFS exports.
4. Scroll to the right, under Show More, click Mount Instructions.

To perform the mounting process from within the guest OS of the VMware VM, follow the below steps:

1. Use SSH client and SSH to the virtual machine.
2. Install the nfs client on the instance.
   a. On Red Hat Enterprise Linux or SuSE Linux instance:

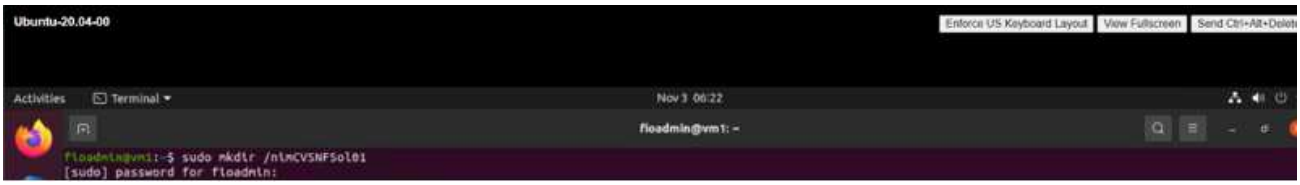
```
sudo yum install -y nfs-utils
```

   b. On an Ubuntu or Debian instance:

```
sudo apt-get install nfs-common
```

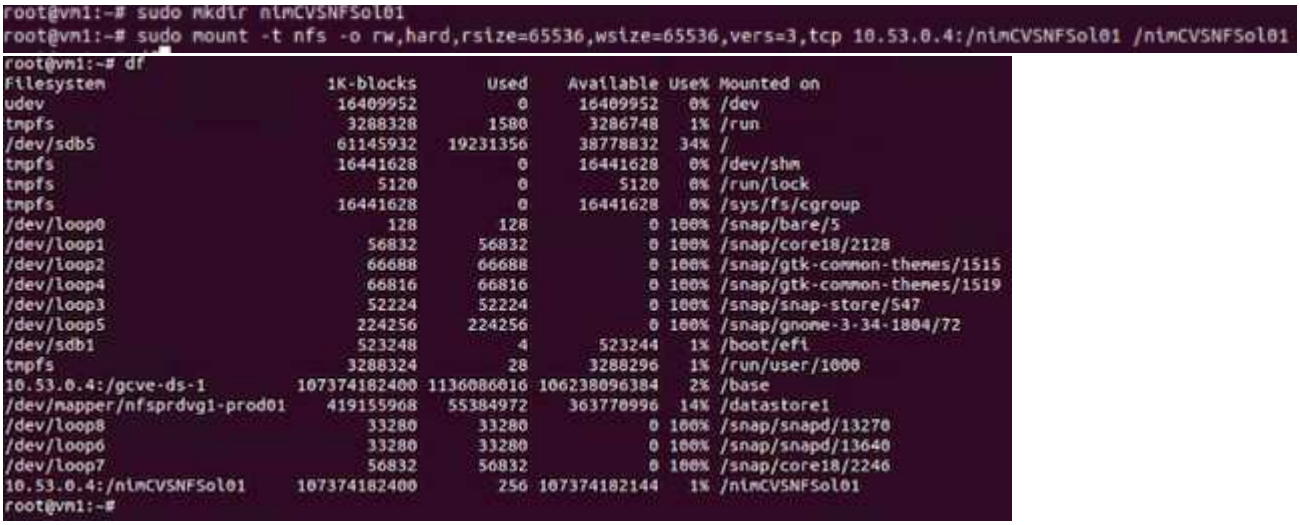
3. Create a new directory on the instance, such as "/nimCVSNFSol01":

```
sudo mkdir /nimCVSNFSol01
```



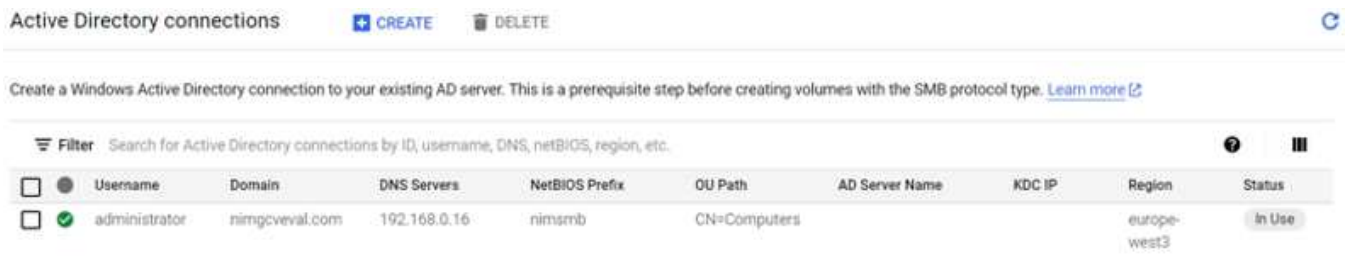
4. Mount the volume using the appropriate command. Example command from the lab is below:

```
sudo mount -t nfs -o rw,hard,rsize=65536,wsize=65536,vers=3,tcp
10.53.0.4:/nimCVSNFSol01 /nimCVSNFSol01
```



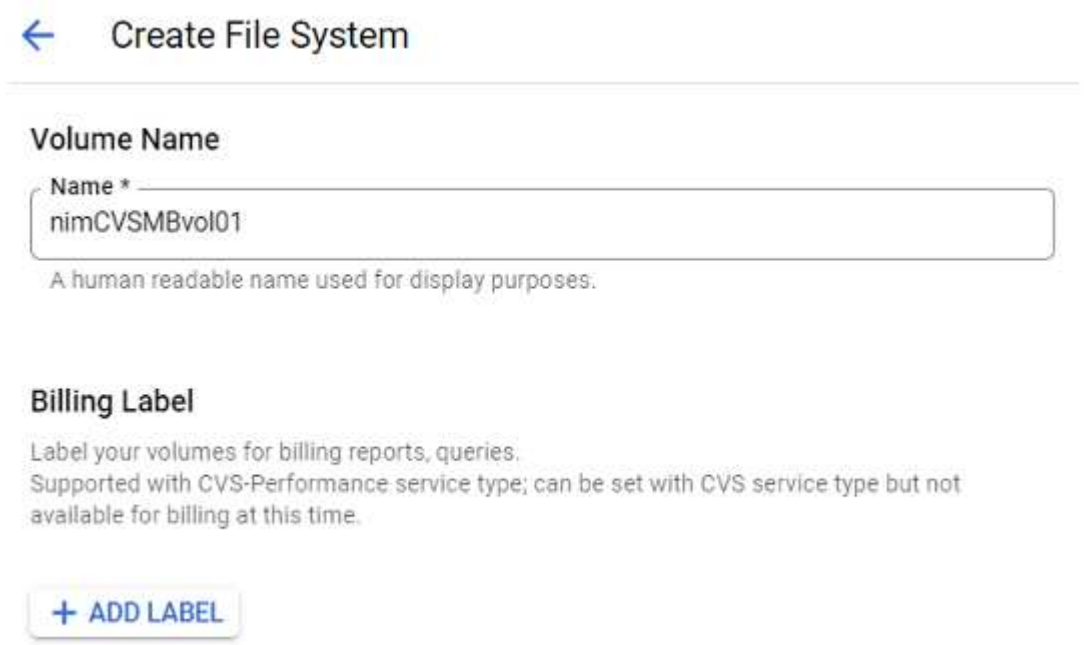
▼ **Creating and Mounting SMB Share to VMs running on VMware Engine**

For SMB volumes, make sure the Active Directory connections is configured prior to creating the SMB volume.

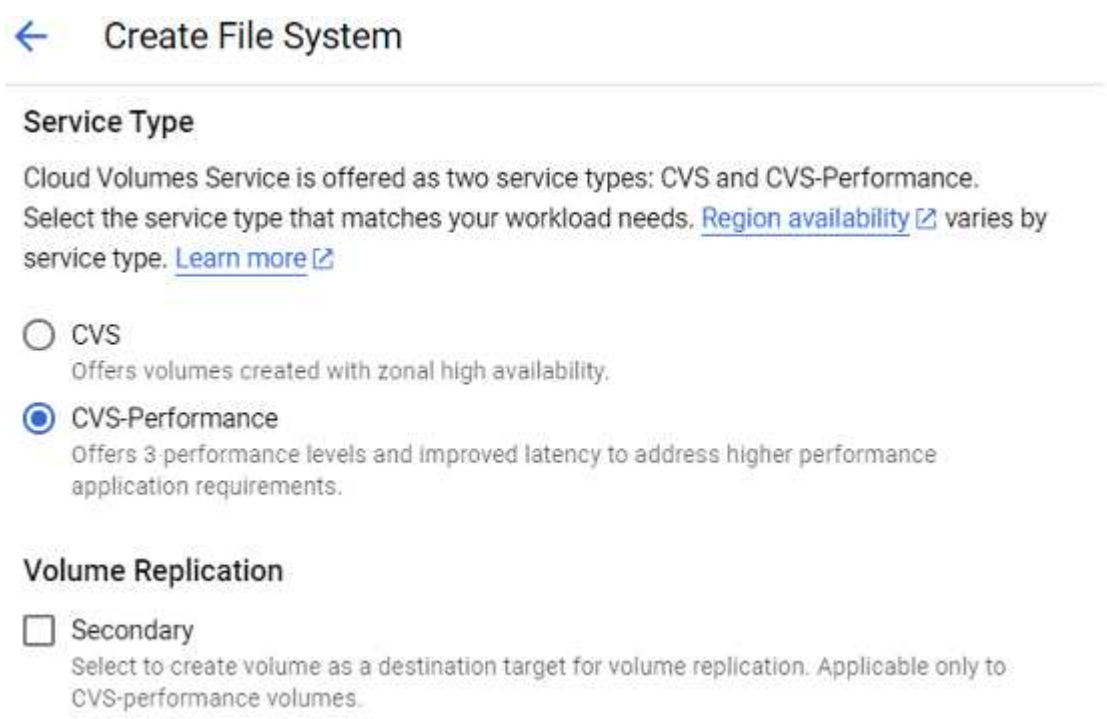


Once the AD connection is in place, create the volume with the desired service level. The steps are like creating NFS volume except selecting the appropriate protocol.

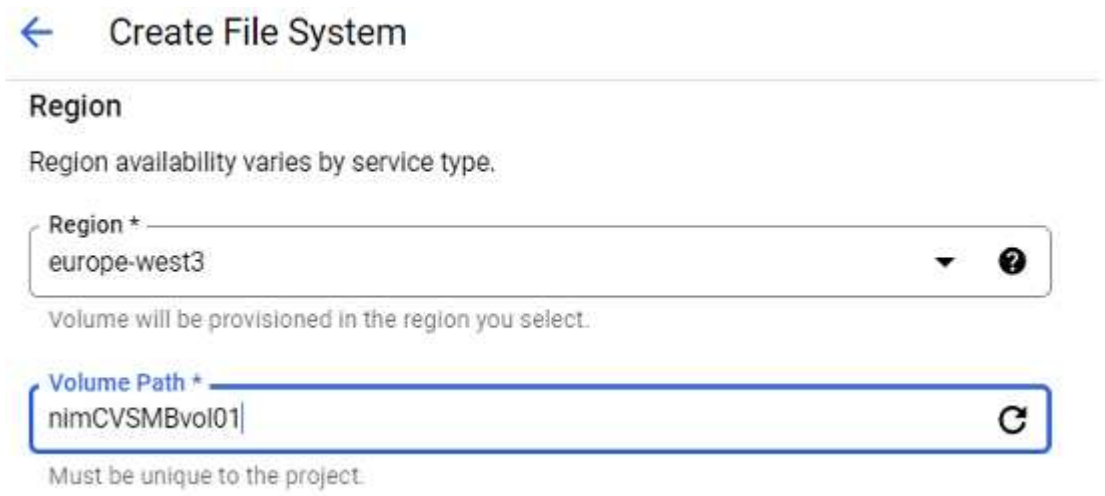
1. In the Cloud Volumes Console, go to the Volumes page and click Create.
2. On the Create File System page, specify the volume name and billing labels as required for chargeback mechanisms.



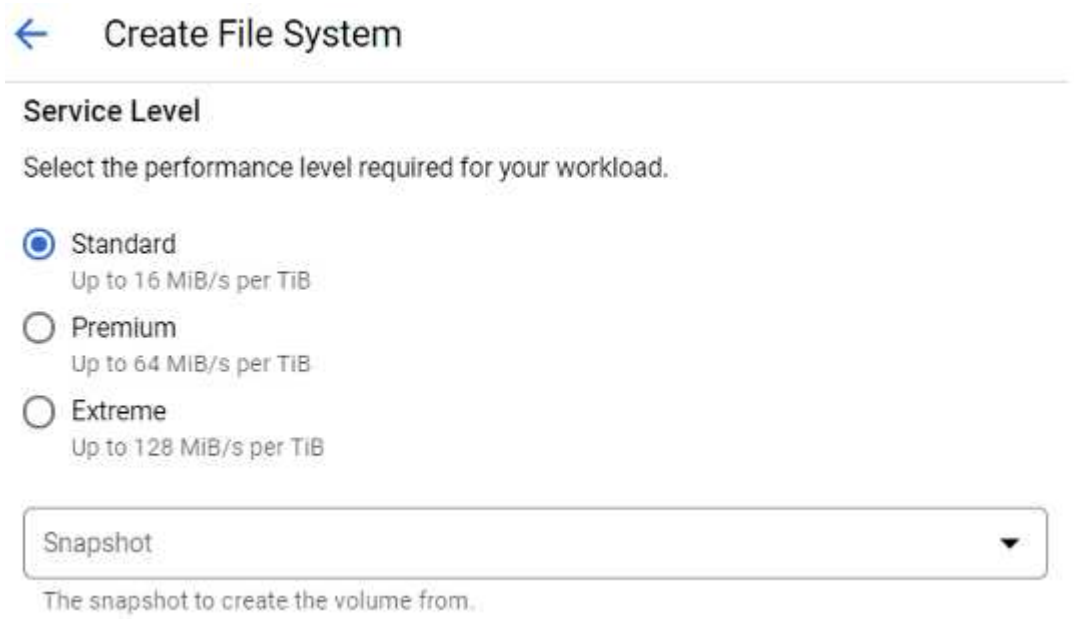
3. Select the appropriate service. For GCVE, choose NetApp Volumes-Performance and desired service level for improved latency and higher performance based on the workload requirements.



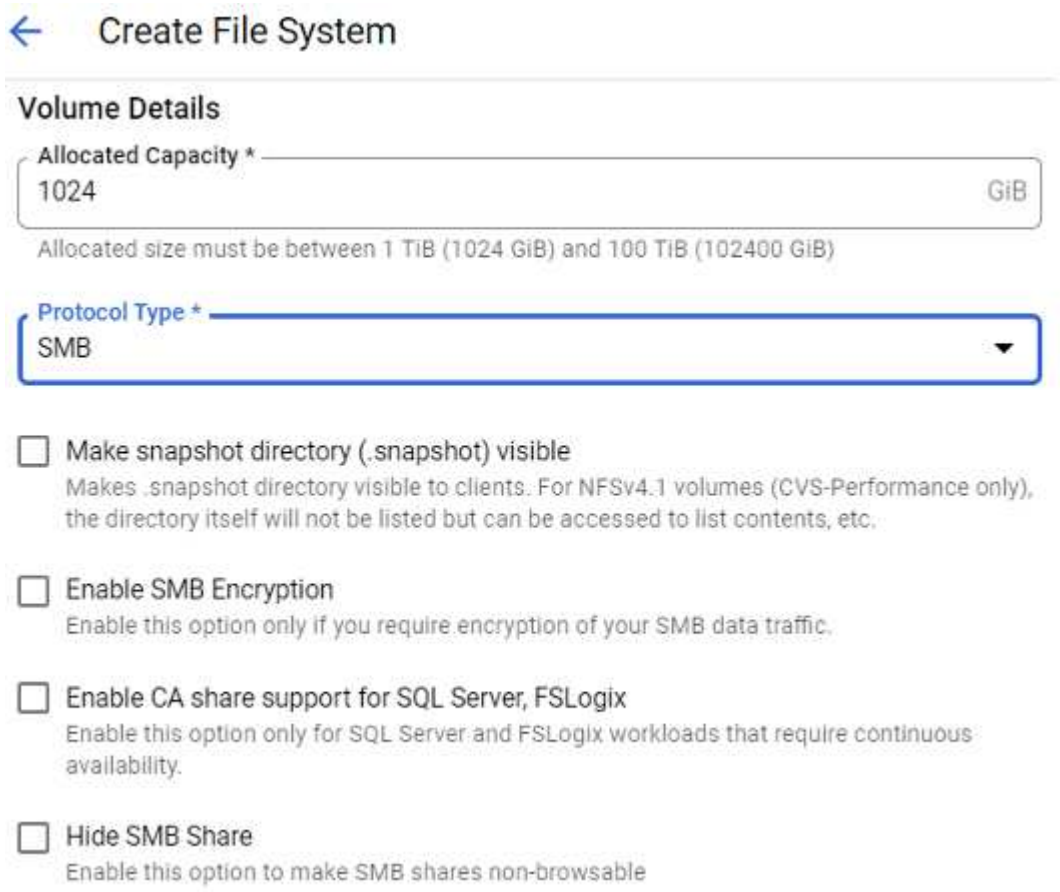
4. Specify the Google Cloud region for the volume and volume path (The volume path must be unique across all of cloud volumes in the project)



5. Select the level of performance for the volume.

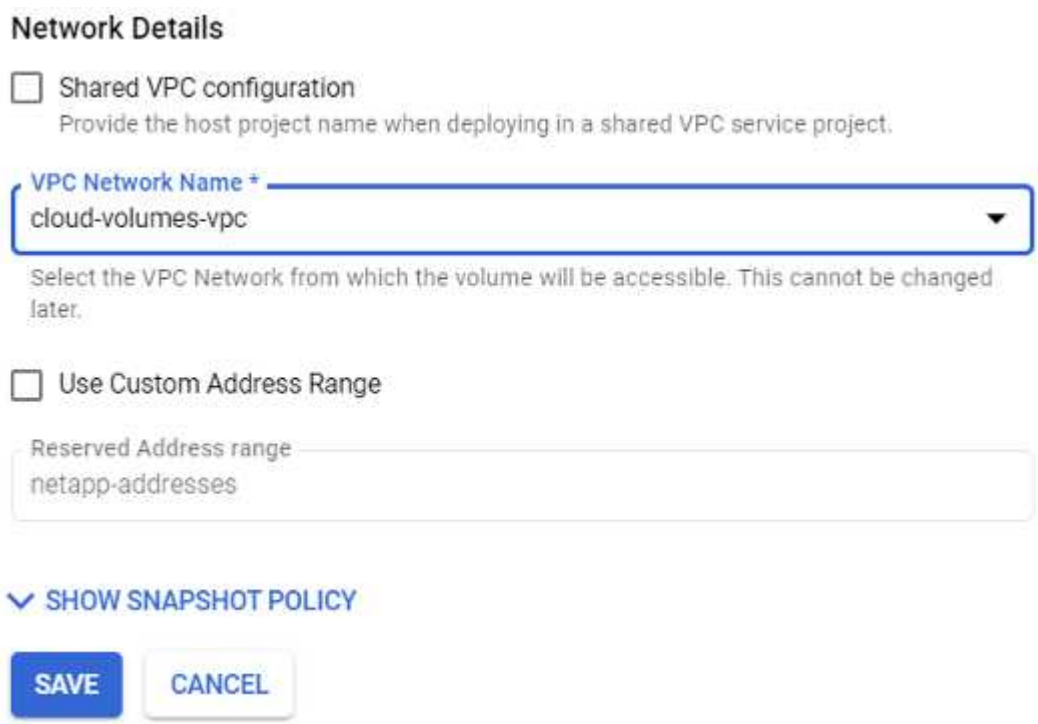


6. Specify the size of the volume and the protocol type. In this testing, SMB is used.



7. In this step, select the VPC Network from which the volume will be accessible. Ensure VPC peering is in place.

   HINT: If VPC peering has not been done, a pop-up button will be displayed to guide you through the peering commands. Open a Cloud Shell session and execute the appropriate commands to peer your VPC with Google Cloud NetApp Volumes producer. In case you decide to prepare VPC peering in beforehand, refer to these instructions.



8. Click Save to create the volume.



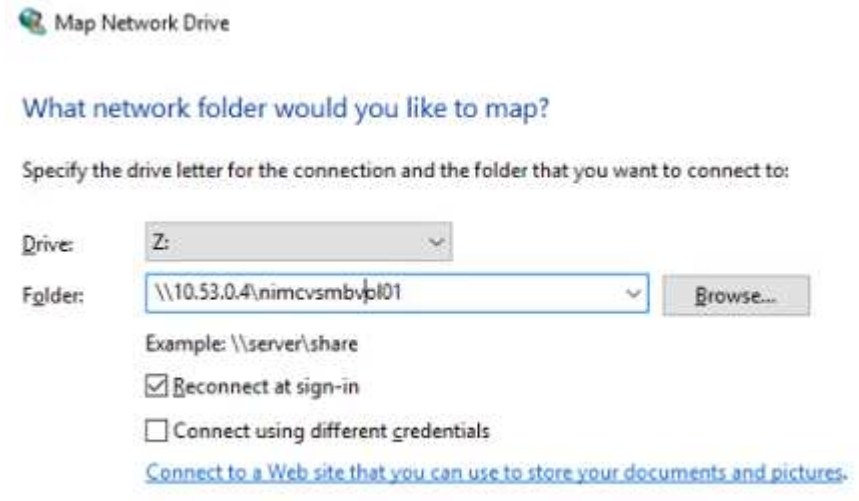
To mount the SMB volume, do the following:

1. In the Cloud Console, go to Cloud Volumes > Volumes.
2. Go to the Volumes page
3. Click the SMB volume for which you want to map an SMB share.
4. Scroll to the right, under Show More, click Mount Instructions.

To perform the mounting process from within the Windows guest OS of the VMware VM, follow the below steps:

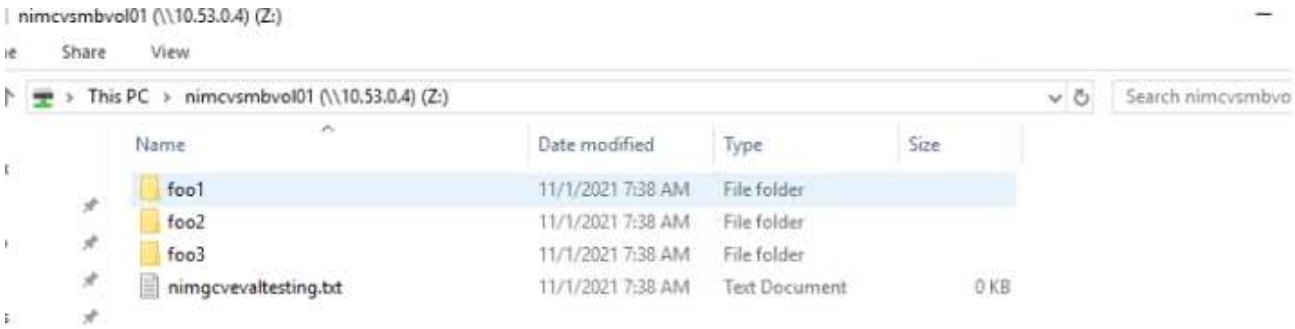
1. Click the Start button and then click on Computer.
2. Click Map Network Drive.
3. In the Drive list, click any available drive letter.
4. In the folder box, type:

```
\\nimsmb-3830.nimgcveval.com\nimCVSMBvol01
```



To connect every time you log on to your computer, select the Reconnect at sign-in check box.

5. Click Finish.



Summary and Conclusion: Why NetApp Hybrid Multicloud with VMware

NetApp Cloud Volumes along with VMware solutions for the major hyperscalers provides great potential for organizations looking to leverage hybrid cloud. The rest of this section provides the use cases that show integrating NetApp Cloud Volumes enables true hybrid Multicloud capabilities.

Use case #1: Optimizing storage

When performing a sizing exercise using RVtools output, it is always evident that the horsepower (vCPU/vMem) scale is parallel with storage. Many times, organizations find themselves in a situation where the storage space requires drives the size of the cluster well beyond what is needed for horsepower.

By integrating NetApp Cloud Volumes, organizations can realize a vSphere-based cloud solution with a simple migration approach, with no re-platforming, no IP changes, and no architectural changes. Additionally, this optimization enables you to scale the storage footprint while keeping the host count to least amount required in vSphere, but no change to the storage hierarchy, security, or files made available. This allows you to optimize the deployment and reduce the overall TCO by 35–45%. This integration also enables you to scale storage from warm storage to production-level performance in seconds.

Use case #2: Cloud migration

Organizations are under pressure to migrate applications from on-premises data centers to the Public Cloud for multiple reasons: an upcoming lease expiration; a finance directive to move from capital expenditure (capex) spending to operational expenditures (opex) spending; or simply a top-down mandate to move everything to the cloud.

When speed is critical, only a streamlined migration approach is feasible because re-platforming and refactoring applications to adapt to the cloud's particular IaaS platform is slow and expensive, often taking months. By combining NetApp Cloud Volumes with the bandwidth-efficient SnapMirror replication for guest-connected storage (including RDMs in conjunction with application-consistent Snapshot copies and HCX, cloud specific migration (e.g. Azure Migrate), or third-party products for replicating VMs), this transition is even easier than relying on time-consuming I/O filters mechanisms.

Use case #3: Data center expansion

When a data center reaches capacity limits due to seasonal demand spikes or just steady organic growth, moving to the cloud-hosted VMware along with NetApp Cloud Volumes is an easy solution. Leveraging NetApp Cloud Volumes allows storage creation, replication, and expansion very easily by providing high availability across availability zones and dynamic scaling capabilities. Leveraging NetApp Cloud Volumes helps in minimizing host cluster capacity by overcoming the need for stretch clusters.

Use case #4: Disaster recovery to the cloud

In a traditional approach, if a disaster occurs, the VMs replicated to the cloud would require conversion to the cloud's own hypervisor platform before they could be restored – not a task to be handled during a crisis.

By using NetApp Cloud Volumes for guest-connected storage using SnapCenter and SnapMirror replication from on-premises along with public cloud virtualization solutions, a better approach for disaster recovery can be devised allowing VM replicas to be recovered on fully consistent VMware SDDC infrastructure along with cloud specific recovery tools (e.g. Azure Site Recovery) or equivalent third-party tools such as Veeam. This approach also enables you to perform disaster recovery drills and recovery from ransomware quickly. This also enables you to scale to full production for testing or during a disaster by adding hosts on-demand.

Use case #5: Application modernization

After applications are in the public cloud, organizations will want to take advantage of the hundreds of powerful cloud services to modernize and extend them. With the use of NetApp Cloud Volumes, modernization is an easy process because the application data is not locked into vSAN and allows data mobility for a wide range of use cases, including Kubernetes.

Conclusion

Whether you are targeting an all-cloud or hybrid cloud, NetApp Cloud Volumes provides excellent options to deploy and manage the application workloads along with file services and block protocols while reducing the TCO by making the data requirements seamless to the application layer.

Whatever the use case, choose your favorite cloud/hyperscaler together with NetApp Cloud Volumes for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance.

It is the same familiar process and procedures that are used to connect the storage. Remember, it is just the position of the data that changed with new names; the tools and processes all remain the same and NetApp Cloud Volumes helps in optimizing the overall deployment.

VMware hybrid cloud use cases

Use Cases for NetApp Hybrid Multicloud with VMware

An overview of the use cases of importance to IT organization when planning hybrid-cloud or cloud-first deployments.

Popular Use Cases

Use cases include:

- Disaster recovery,
- Hosting workloads during data center maintenance, * quick burst in which additional resources are required beyond what's provisioned in the local data center,
- VMware site expansion,
- Fast migration to the cloud,
- Dev/test, and
- Modernization of apps leveraging cloud supplemental technologies.

Throughout this documentation, cloud workload references will be detailed using the VMware use-cases. These use-cases are:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

Inside the IT Journey

Most organizations are on a journey to transformation and modernization. As part of this process, companies are trying use their existing VMware investments while leveraging cloud benefits and exploring ways to make the migration process as seamless as possible. This approach would make their modernization efforts very easy because the data is already in the cloud.

The easiest answer to this scenario is VMware offerings in each hyperscaler. Like NetApp Cloud Volumes, VMware provides a way to move or extend on-premises VMware environments to any cloud, allowing you to retain existing on-premises assets, skills, and tools while running workloads natively in the cloud. This reduces risk because there will be no service breaks or a need for IP changes and provides the IT team the ability to operate the way they do on-premises using existing skills and tools. This can lead to accelerated cloud migrations and a much smoother transition to a hybrid Multicloud architecture.

Understanding the Importance of Supplemental NFS Storage Options

While VMware in any cloud delivers unique hybrid capabilities to every customer, limited supplemental NFS storage options have restricted its usefulness for organizations with storage-heavy workloads. Because storage is directly tied to hosts, the only way to scale storage is to add more hosts—and that can increase costs by 35–40 percent or more for storage intensive workloads. These workloads just need additional storage, not additional horsepower. But that means paying for additional hosts.

Let's consider this scenario:

A customer requires just five hosts for CPU and memory, but has a lot of storage needs, and needs 12 hosts to

meet the storage requirement. This requirement ends up really tipping the financial scale by having to buy the additional horsepower, when they only need to increment the storage.

When you're planning cloud adoption and migrations, it's always important to evaluate the best approach and take the easiest path that reduces total investments. The most common and easiest approach for any application migration is rehosting (also known as lift and shift) where there is no virtual machine (VM) or data conversion. Using NetApp Cloud Volumes with VMware software-defined data center (SDDC), while complementing vSAN, provides an easy lift-and-shift option.

NetApp Solutions for VMware in Hyperscalers

Learn more about the solutions that NetApp brings the VMware environment in each of the hyperscalers - from migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

NetApp Solutions for VMware Environments

Whether you are operating in a hybrid-cloud model or a "cloud first" model, NetApp offers a wide variety of solutions to address the most common use cases for managing workloads in a cloud or hybrid-cloud model.

NetApp also offers solutions for storage provisioned as in-guest (guest connected) storage or as a supplemental NFS datastore in each of the hyperscalers. All solutions are categorized in unison with the VMware classification of cloud workloads. Those classifications include:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

For more details on the solutions available for each hyperscaler, please visit:

- Solutions for AWS / VMC
- Solutions for Azure / AVS
- Solutions for GCP / GCVE

NetApp Solutions for Amazon VMware Managed Cloud (VMC)

Learn more about the solutions that NetApp brings to AWS.

VMware defines the cloud workloads into one of three categories:

- Protect (including both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

Browse the available solutions in the following sections.

Protect

- Disaster Recovery with VMC on AWS (guest connected)
- Veeam Backup & Restore in VMC with FSx ONTAP

- Disaster Recovery (DRO) with FSx ONTAP and VMC
- Using Veeam Replication and FSx ONTAP for Disaster recovery to VMware Cloud on AWS

Migrate

- Migrate Workloads to FSx ONTAP datastore using VMware HCX

Extend

COMING SOON!!

NetApp Solutions for Azure VMware Solution (AVS)

Learn more about the solutions that NetApp brings to Azure.

VMware defines the cloud workloads into one of three categories:

- Protect (including both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

Browse the available solutions in the following sections.

Protect

- Disaster Recovery with ANF and JetStream (supplemental NFS datastore)
- Disaster Recovery with ANF and CVO (guest connected storage)
- Disaster Recovery (DRO) with ANF and AVS
- Using Veeam Replication and Azure NetApp Files datastore for disaster recovery to Azure VMware Solution

Migrate

- Migrate Workloads to Azure NetApp Files datastore using VMware HCX

Extend

COMING SOON!!

NetApp Solutions for Google Cloud VMware Engine (GCVE)

Learn more about the solutions that NetApp brings to GCP.

VMware defines the cloud workloads into one of three categories:

- Protect (including both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

Browse the available solutions in the following sections.

Protect

- Application Disaster Recovery with SnapCenter, Cloud Volumes ONTAP and Veeam Replication
- Application Consistent Disaster Recovery with NetApp SnapCenter and Veeam Replication to NetApp Volumes on GCVE

Migrate

- Workload Migration using VMware HCX to Google Cloud NetApp Volumes NFS datastore
- VM Replication using Veeam to Google Cloud NetApp Volumes NFS datastore

Extend

COMING SOON!!

NetApp for AWS/VMC

NetApp Capabilities for AWS VMC

Learn more about the capabilities that NetApp brings to the AWS VMware Cloud (VMC) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- Configuring VMC in AWS
- NetApp Storage Options for VMC
- NetApp / VMware Cloud Solutions

Configuring VMC in AWS

As with on-premises, planning a cloud based virtualization environment is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to AWS VMC.

The setup process can be broken down into the following steps:

- Deploy and Configure VMware Cloud for AWS
- Connect VMware Cloud to FSx ONTAP

View the detailed configuration steps for VMC.

NetApp Storage Options for VMC

NetApp storage can be utilized in several ways - either as guess connected or as a supplemental NFS datastore - within AWS VMC.

Please visit Supported NetApp Storage Options for more information.

AWS supports NetApp storage in the following configurations:

- FSx ONTAP as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- FSx ONTAP as a supplemental NFS datastore

View the detailed guest connect storage options for VMC.
View the detailed supplemental NFS datastore options for VMC.

Solution Use Cases

With NetApp and VMware cloud solutions, many use cases are simple to deploy in your AWS VMC. Use cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Extend
- Migrate

Browse the NetApp solutions for AWS VMC

Protect workloads on AWS/VMC

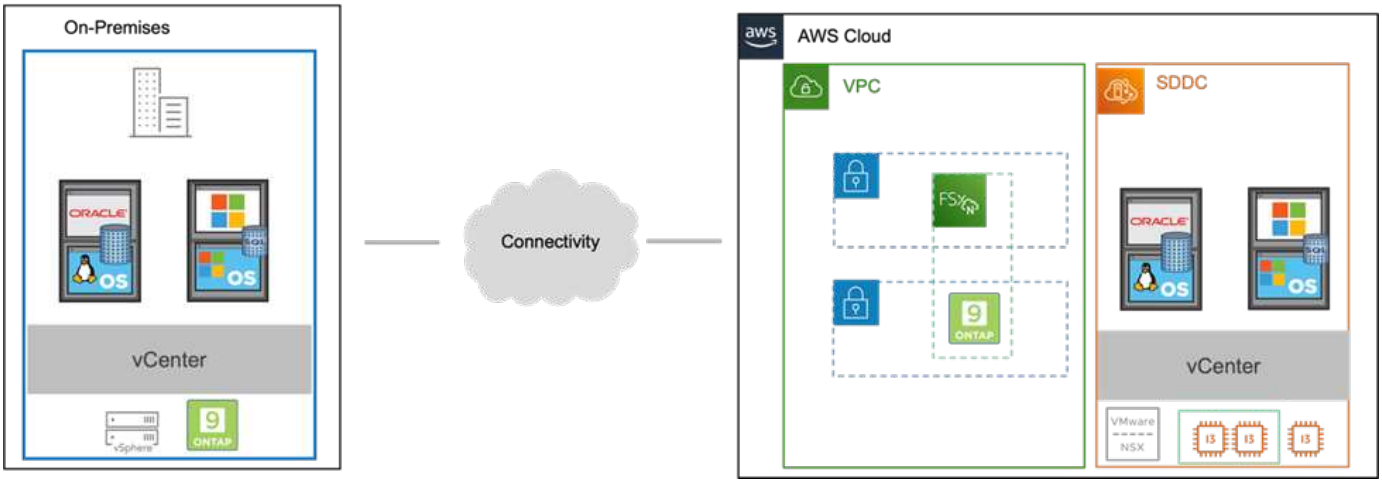
TR-4931: Disaster Recovery with VMware Cloud on Amazon Web Services and Guest Connect

A proven disaster recovery (DR) environment and plan is critical for organizations to ensure that business-critical applications can be rapidly restored in the event of a major outage. This solution focuses on demonstrating DR use cases with a focus on VMware and NetApp technologies, both on-premises and with VMware Cloud on AWS.

Overview

NetApp has a long history of integration with VMware as evidenced by the tens of thousands of customers that have chosen NetApp as their storage partner for their virtualized environment. This integration continues with guest-connected options in the cloud and recent integrations with NFS datastores as well. This solution focuses on the use case commonly referred to as guest-connected storage.

In guest-connected storage, the guest VMDK is deployed on a VMware-provisioned datastore, and application data is housed on iSCSI or NFS and mapped directly to the VM. Oracle and MS SQL applications are used to demonstrate a DR scenario, as shown in the following figure.



Assumptions, pre-requisites and component overview

Before deploying this solution, review the overview of the components, the required pre-requisites to deploy the solution and assumptions made in documenting this solution.

DR Solution Requirements, Pre-requisites and Planning

Performing DR with SnapCenter

In this solution, SnapCenter provides application-consistent snapshots for SQL Server and Oracle application data. This configuration, together with SnapMirror technology, provides high-speed data replication between our on-premises AFF and FSx ONTAP cluster. Additionally, Veeam Backup & Replication provides backup and restore capabilities for our virtual machines.

In this section, we cover the configuration of SnapCenter, SnapMirror, and Veeam for both backup and restore.

The following sections cover configuration and the steps needed to complete a failover at the secondary site:

Configure SnapMirror relationships and retention schedules

SnapCenter can update SnapMirror relationships within the primary storage system (primary > mirror) and to secondary storage systems (primary > vault) for the purpose of long-term archiving and retention. To do so, you must establish and initialize a data replication relationship between a destination volume and a source volume using SnapMirror.

The source and destination ONTAP systems must be in networks that are peered using Amazon VPC peering, a transit gateway, AWS Direct Connect, or an AWS VPN.

The following steps are required for setting up SnapMirror relationships between an on-premises ONTAP system and FSx ONTAP:

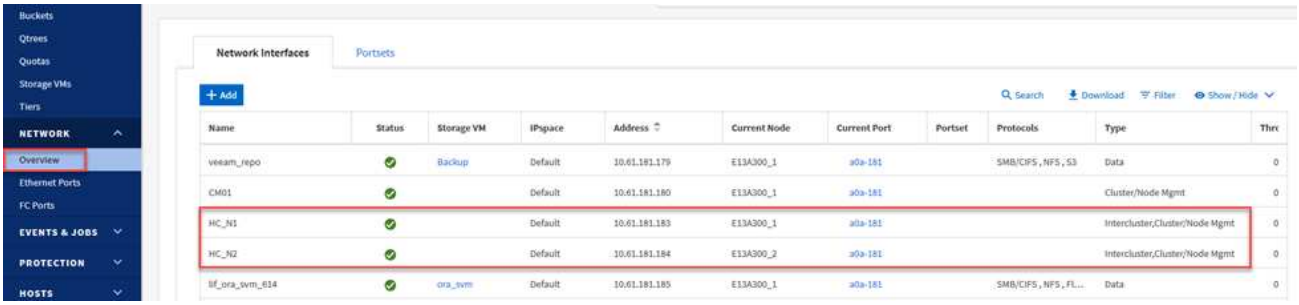
Refer to the FSx ONTAP – ONTAP User Guide for more information on creating SnapMirror relationships with FSx.

▼ **Record the source and destination Intercluster logical interfaces**

For the source ONTAP system residing on-premises, you can retrieve the inter-cluster LIF information from System Manager or from the CLI.

1. In ONTAP System Manager, navigate to the Network Overview page and retrieve the IP addresses of

Type: Intercluster that are configured to communicate with the AWS VPC where FSx is installed.



2. To retrieve the Intercluster IP addresses for FSx, log into the CLI and run the following command:

```
FSx-Dest::> network interface show -role intercluster
```

```
FsxId0ae40e08acc0dea67::> network interface show -role intercluster
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
FsxId0ae40e08acc0dea67
            inter_1      up/up    172.30.15.42/25    FsxId0ae40e08acc0dea67-01
                                                                   e0e     true
            inter_2      up/up    172.30.14.28/26    FsxId0ae40e08acc0dea67-02
                                                                   e0e     true
2 entries were displayed.
```

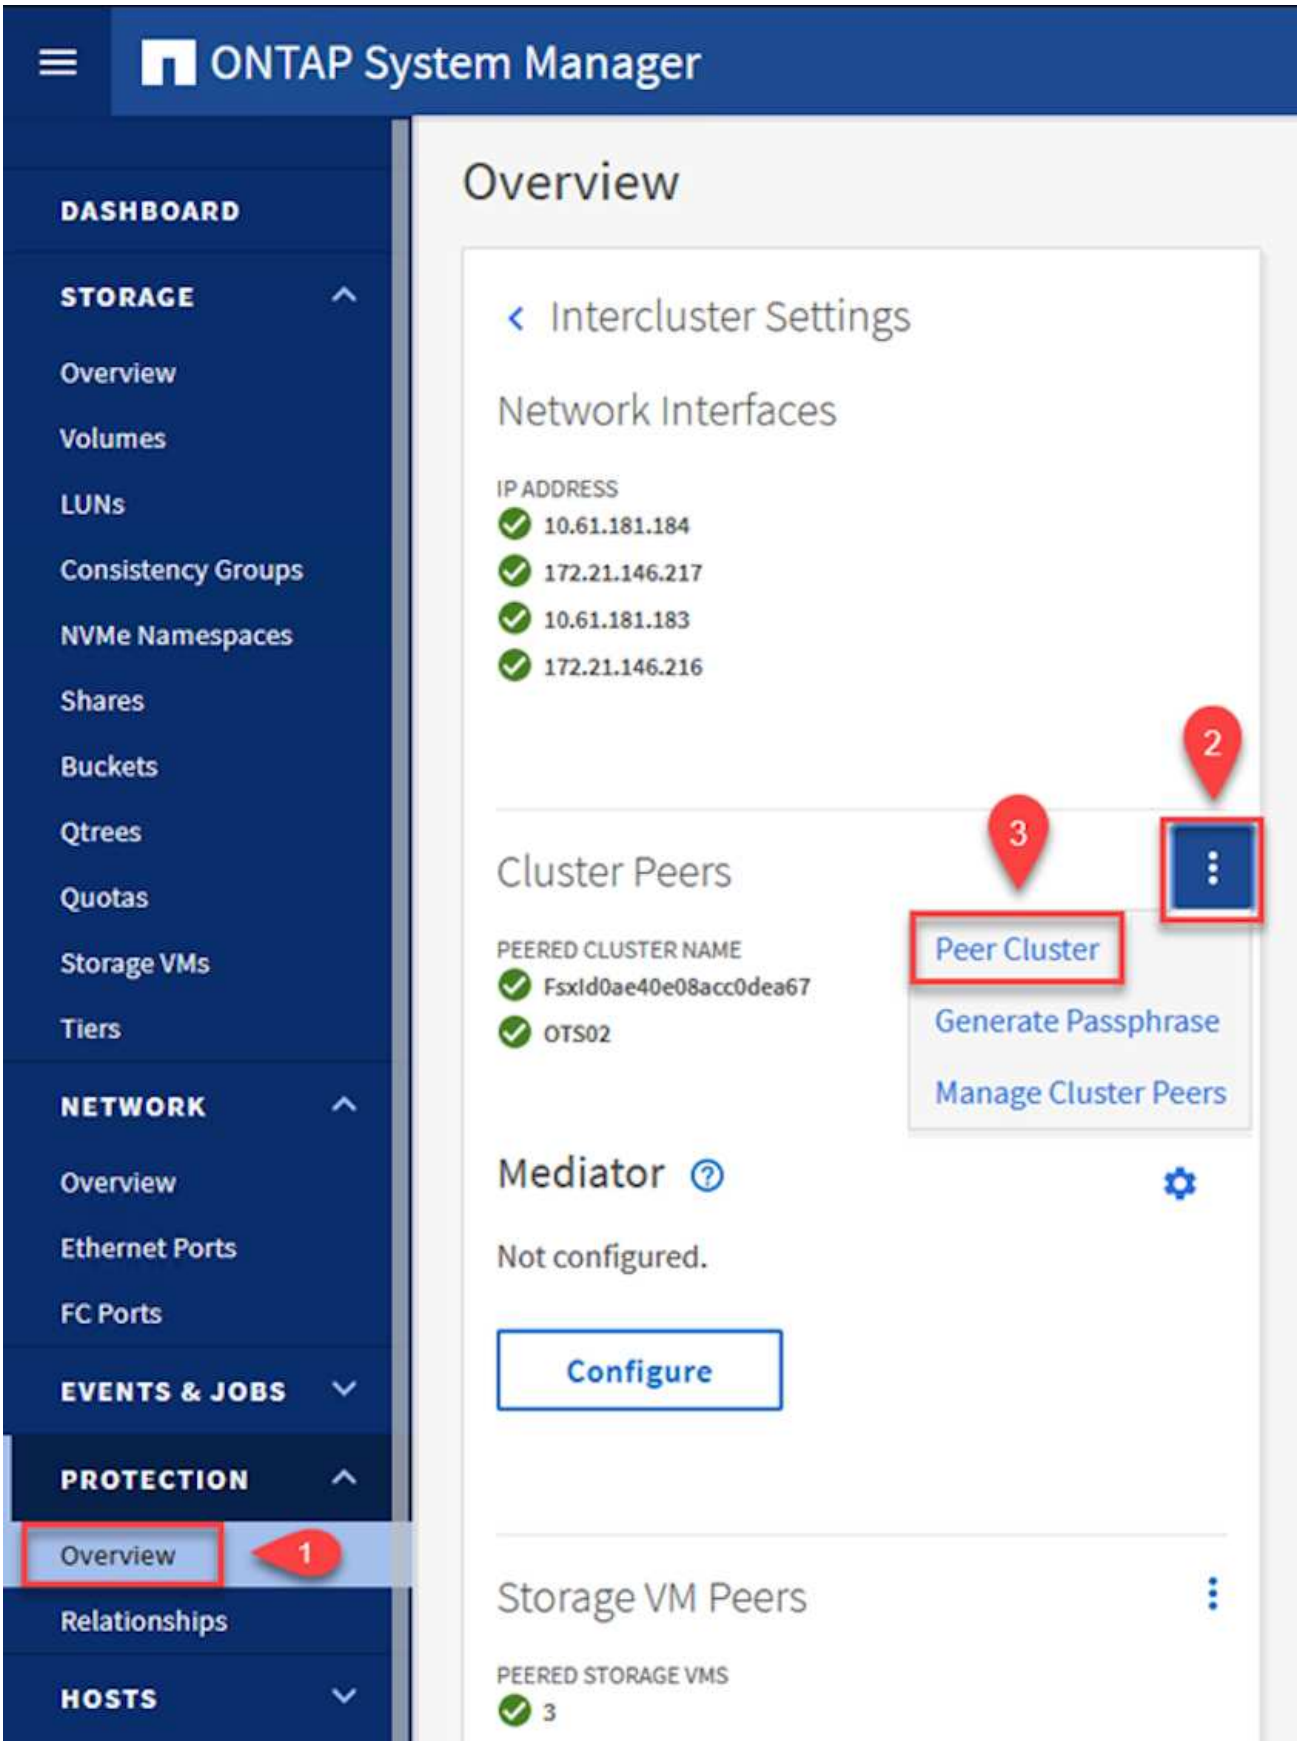
▼ **Establish cluster peering between ONTAP and FSx**

To establish cluster peering between ONTAP clusters, a unique passphrase entered at the initiating ONTAP cluster must be confirmed in the other peer cluster.

1. Set up peering on the destination FSx cluster using the `cluster peer create` command. When prompted, enter a unique passphrase that is used later on the source cluster to finalize the creation process.

```
FSx-Dest::> cluster peer create -address-family ipv4 -peer-addrs
source_intercluster_1, source_intercluster_2
Enter the passphrase:
Confirm the passphrase:
```

2. At the source cluster, you can establish the cluster peer relationship using either ONTAP System Manager or the CLI. From ONTAP System Manager, navigate to Protection > Overview and select Peer Cluster.



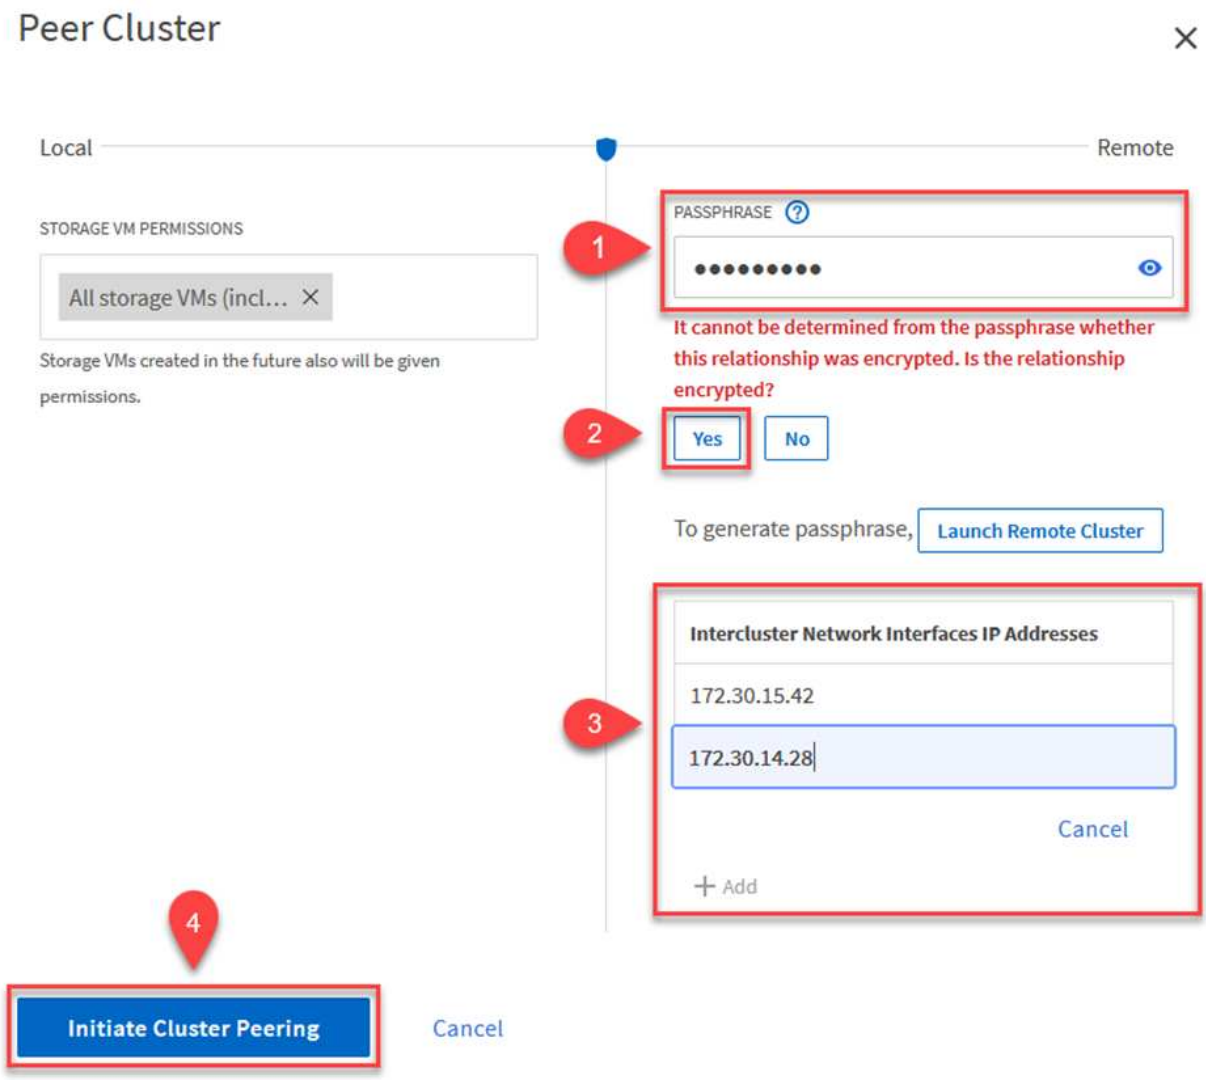
3. In the Peer Cluster dialog box, fill out the required information:
   a. Enter the passphrase that was used to establish the peer cluster relationship on the destination FSx

cluster.

b. Select `Yes` to establish an encrypted relationship.

c. Enter the intercluster LIF IP address(es) of the destination FSx cluster.

d. Click Initiate Cluster Peering to finalize the process.



4. Verify the status of the cluster peer relationship from the FSx cluster with the following command:

```
FSx-Dest::> cluster peer show
```

```
FsxId0ae40e08acc0dea67::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability   Authentication
------------------------- --------------------- -------------- --------------
E13A300                   1-80-000011           Available      ok
```

▼ **Establish SVM peering relationship**

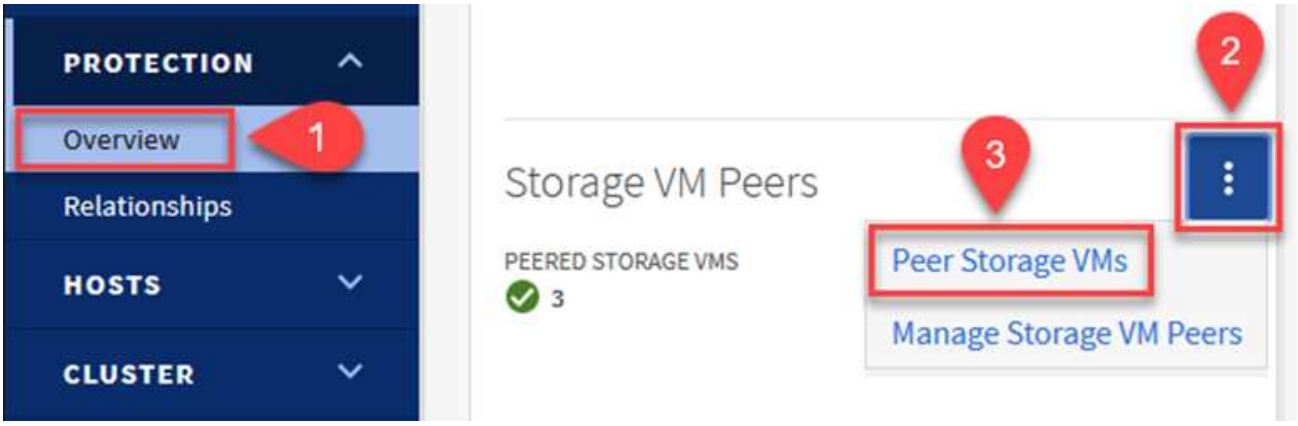
The next step is to set up an SVM relationship between the destination and source storage virtual machines

that contain the volumes that will be in SnapMirror relationships.

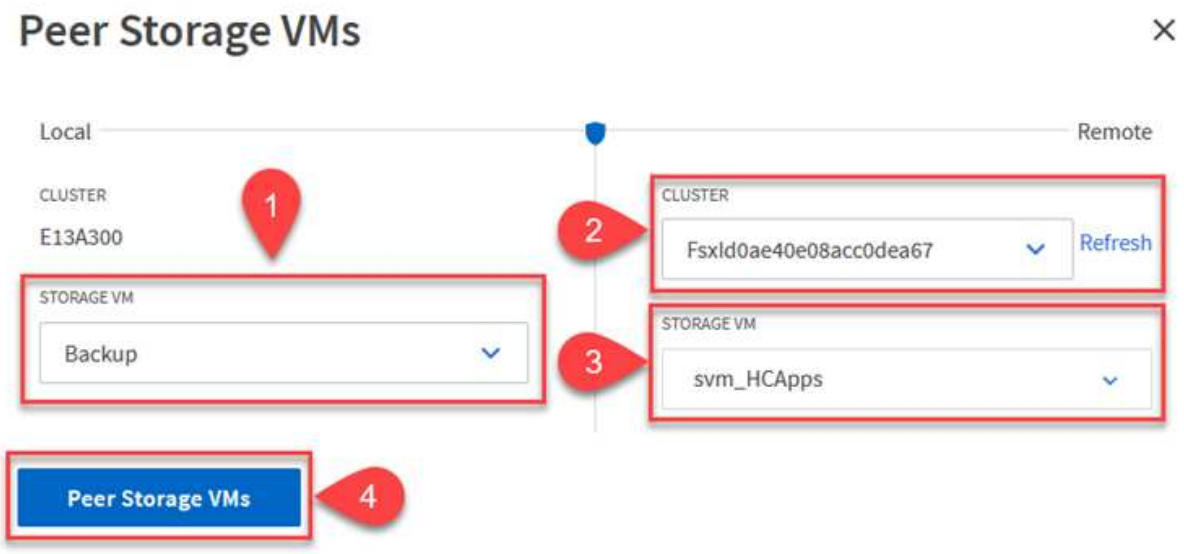
1. From the source FSx cluster, use the following command from the CLI to create the SVM peer relationship:

```
FSx-Dest::> vserver peer create -vserver DestSVM -peer-vserver Backup
-peer-cluster OnPremSourceSVM -applications snapmirror
```

2. From the source ONTAP cluster, accept the peering relationship with either ONTAP System Manager or the CLI.
3. From ONTAP System Manager, go to Protection > Overview and select Peer Storage VMs under Storage VM Peers.



4. In the Peer Storage VM's dialog box, fill out the required fields:
   - The source storage VM
   - The destination cluster
   - The destination storage VM



5. Click Peer Storage VMs to complete the SVM peering process.

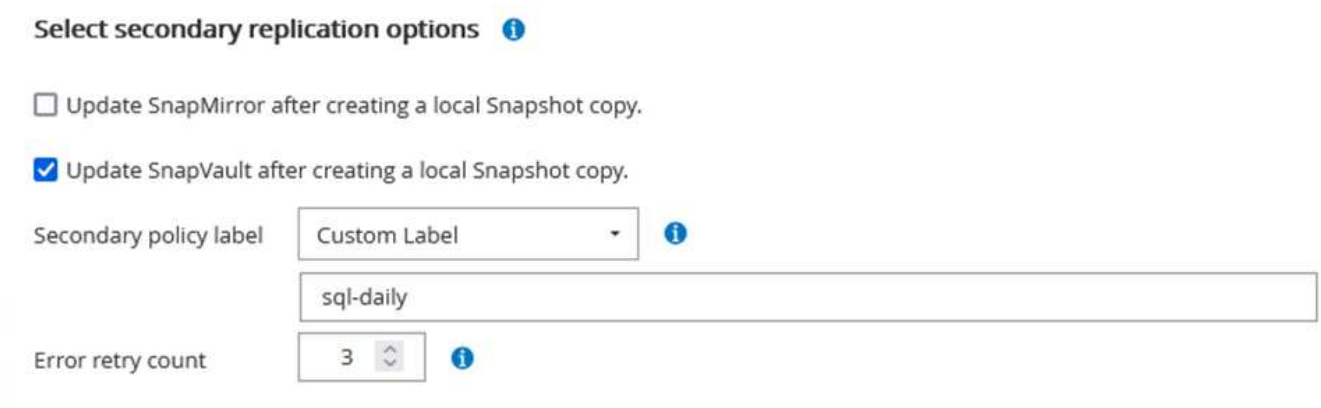
▼ Create a snapshot retention policy

SnapCenter manages retention schedules for backups that exist as snapshot copies on the primary storage system. This is established when creating a policy in SnapCenter. SnapCenter does not manage retention policies for backups that are retained on secondary storage systems. These policies are managed separately through a SnapMirror policy created on the secondary FSx cluster and associated with the destination volumes that are in a SnapMirror relationship with the source volume.

When creating a SnapCenter policy, you have the option to specify a secondary policy label that is added to the SnapMirror label of each snapshot generated when a SnapCenter backup is taken.

> ⓘ On the secondary storage, these labels are matched to policy rules associated with the destination volume for the purpose of enforcing retention of snapshots.

The following example shows a SnapMirror label that is present on all snapshots generated as part of a policy used for daily backups of our SQL Server database and log volumes.



For more information on creating SnapCenter policies for a SQL Server database, see the SnapCenter documentation.

You must first create a SnapMirror policy with rules that dictate the number of snapshot copies to retain.

1. Create the SnapMirror Policy on the FSx cluster.

```
FSx-Dest::> snapmirror policy create -vserver DestSVM -policy
PolicyName -type mirror-vault -restart always
```

2. Add rules to the policy with SnapMirror labels that match the secondary policy labels specified in the SnapCenter policies.

```
FSx-Dest::> snapmirror policy add-rule -vserver DestSVM -policy
PolicyName -snapmirror-label SnapMirrorLabelName -keep
#ofSnapshotsToRetain
```

The following script provides an example of a rule that could be added to a policy:

```
FSx-Dest::> snapmirror policy add-rule -vserver sql_svm_dest -policy
Async_SnapCenter_SQL -snapmirror-label sql-ondemand -keep 15
```

Create additional rules for each SnapMirror label and the number of snapshots to be retained (retention period).

▼ **Create destination volumes**

To create a destination volume on FSx that will be the recipient of snapshot copies from our source volumes, run the following command on FSx ONTAP:

```
FSx-Dest::> volume create -vserver DestSVM -volume DestVolName
-aggregate DestAggrName -size VolSize -type DP
```

▼ **Create the SnapMirror relationships between source and destination volumes**

To create a SnapMirror relationship between a source and destination volume, run the following command on FSx ONTAP:

```
FSx-Dest::> snapmirror create -source-path
OnPremSourceSVM:OnPremSourceVol -destination-path DestSVM:DestVol -type
XDP -policy PolicyName
```

▼ **Initialize the SnapMirror relationships**

Initialize the SnapMirror relationship. This process initiates a new snapshot generated from the source volume and copies it to the destination volume.

```
FSx-Dest::> snapmirror initialize -destination-path DestSVM:DestVol
```

Deploy and configure Windows SnapCenter server on-premises.

▼ **Deploy Windows SnapCenter Server on premises**

This solution uses NetApp SnapCenter to take application-consistent backups of SQL Server and Oracle databases. In conjunction with Veeam Backup & Replication for backing up virtual machine VMDKs, this provides a comprehensive disaster recovery solution for on-premises and cloud-based datacenters.

SnapCenter software is available from the NetApp support site and can be installed on Microsoft Windows systems that reside either in a domain or workgroup. A detailed planning guide and installation instructions can be found at the NetApp Documentation Center.

The SnapCenter software can be obtained at this link.

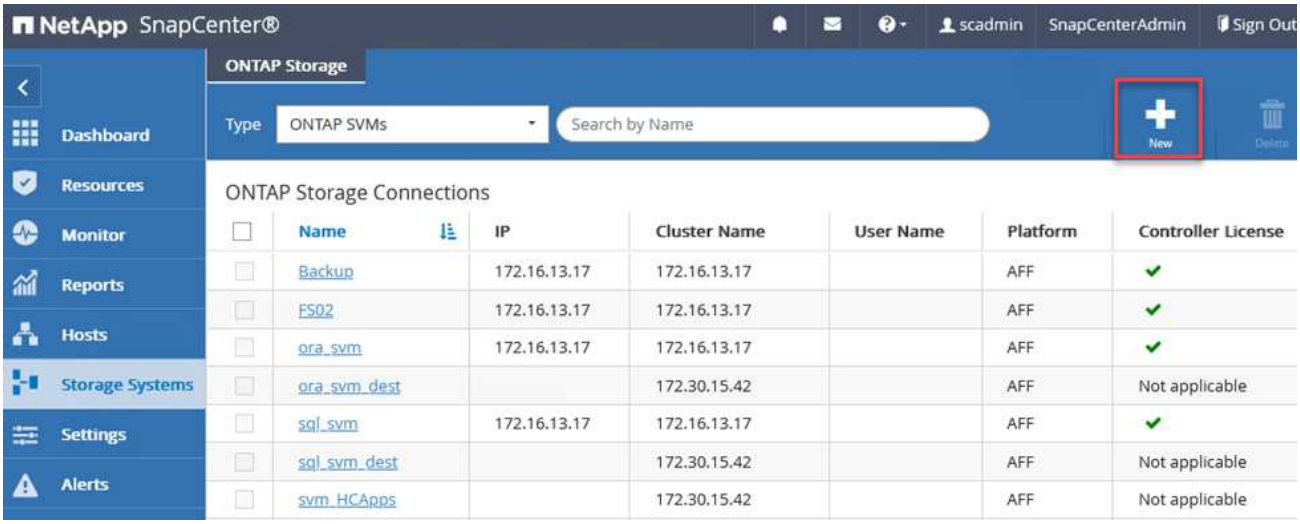
After it is installed, you can access the SnapCenter console from a web browser using *https://Virtual_Cluster_IP_or_FQDN:8146*.

After you log into the console, you must configure SnapCenter for backup SQL Server and Oracle databases.

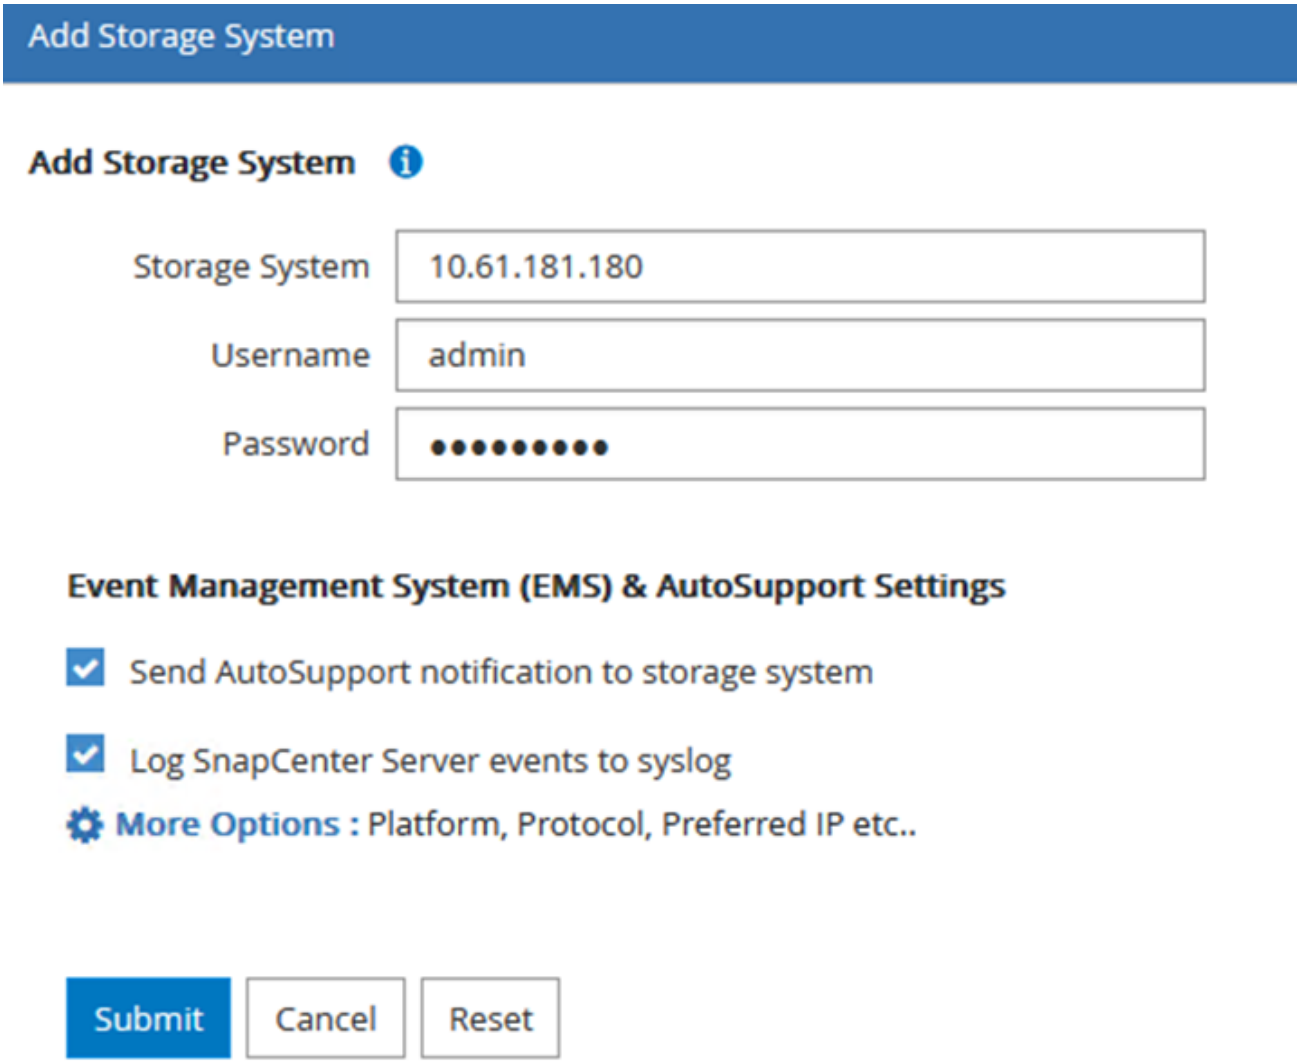
▼ **Add storage controllers to SnapCenter**

To add storage controllers to SnapCenter, complete the following steps:

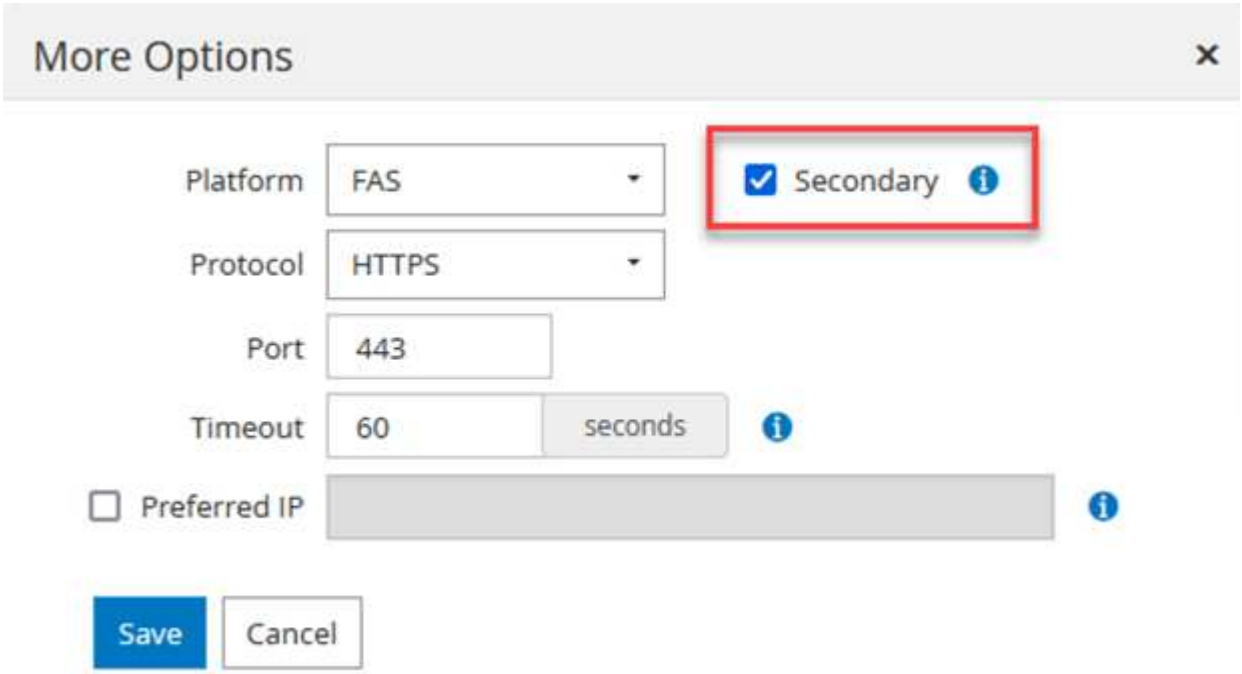
1. From the left menu, select Storage Systems and then click New to begin the process of adding your storage controllers to SnapCenter.



2. In the Add Storage System dialog box, add the management IP address for the local on-premises ONTAP cluster and the username and password. Then click Submit to begin discovery of the storage system.



3. Repeat this process to add the FSx ONTAP system to SnapCenter. In this case, select More Options at the bottom of the Add Storage System window and click the check box for Secondary to designate the FSx system as the secondary storage system updated with SnapMirror copies or our primary backup snapshots.

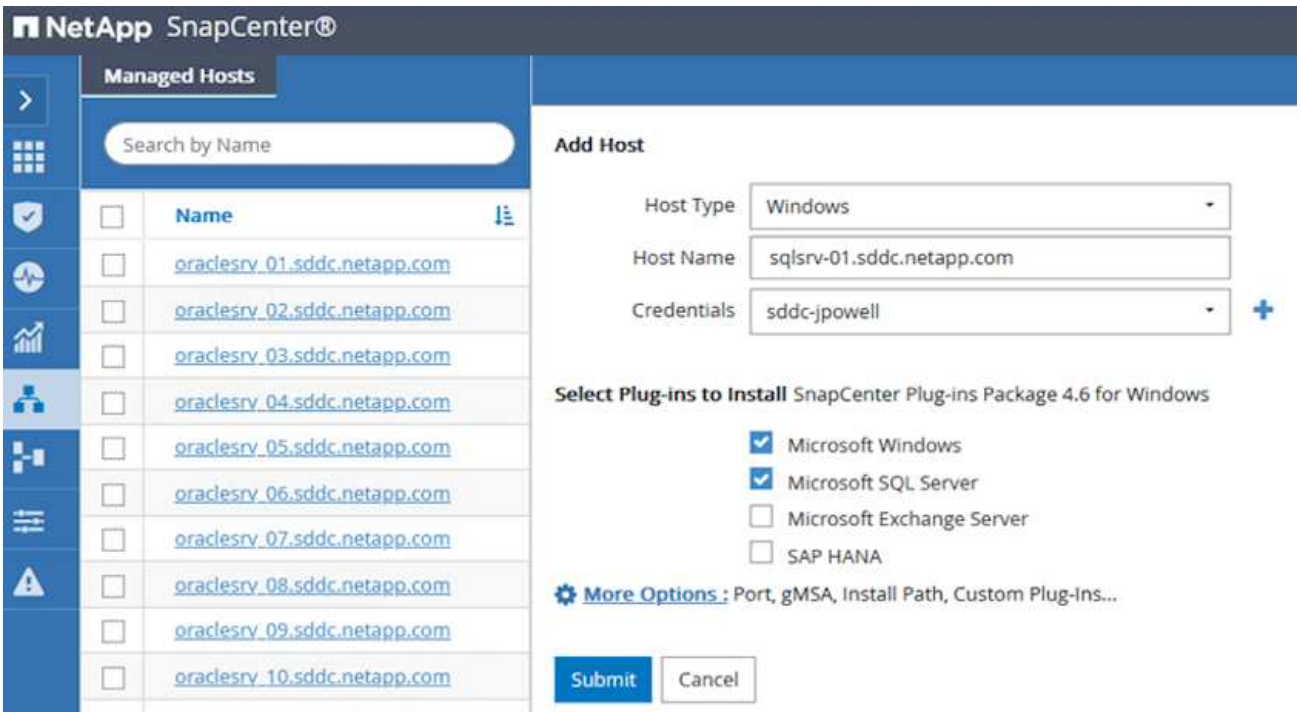


For more information related to adding storage systems to SnapCenter, see the documentation at this link.

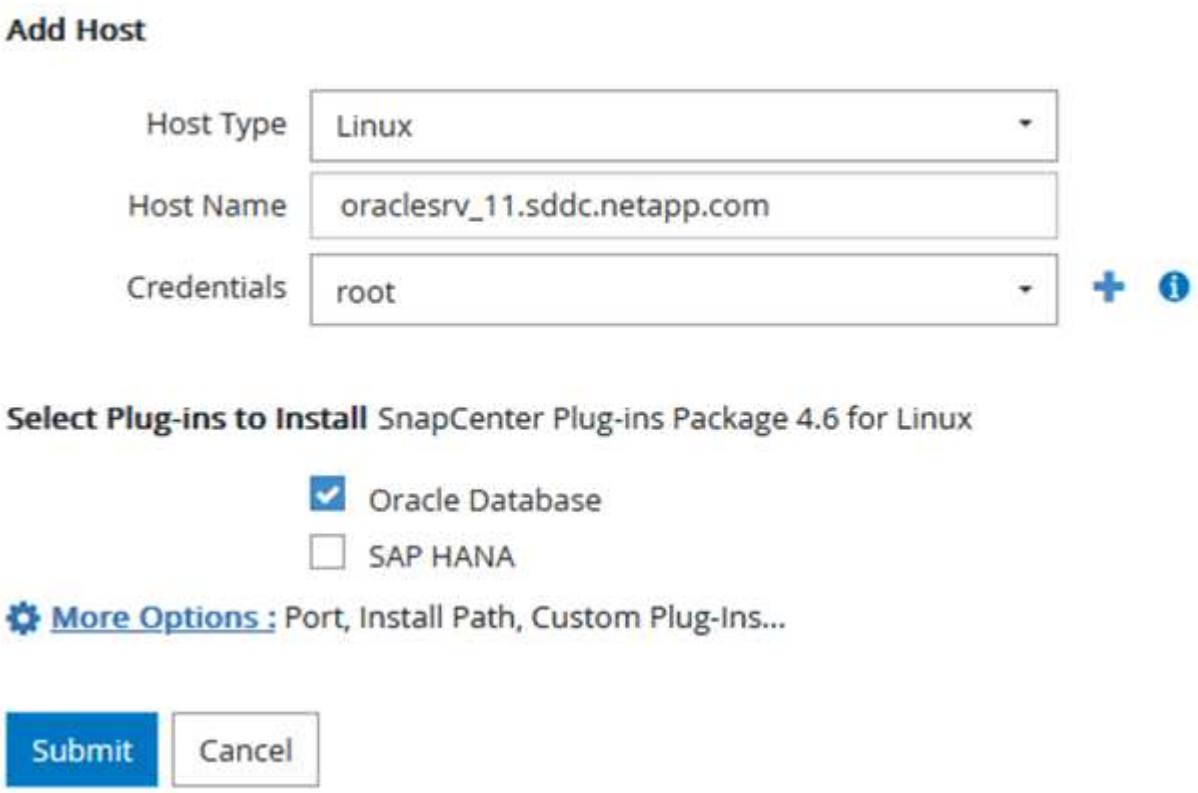
▼ **Add hosts to SnapCenter**

The next step is adding host application servers to SnapCenter. The process is similar for both SQL Server and Oracle.

1. From the left menu, select Hosts and then click Add to begin the process of adding storage controllers to SnapCenter.
2. In the Add Hosts window, add the Host Type, Hostname, and the host system Credentials. Select the plug-in type. For SQL Server, select the Microsoft Windows and Microsoft SQL Server plug-in.



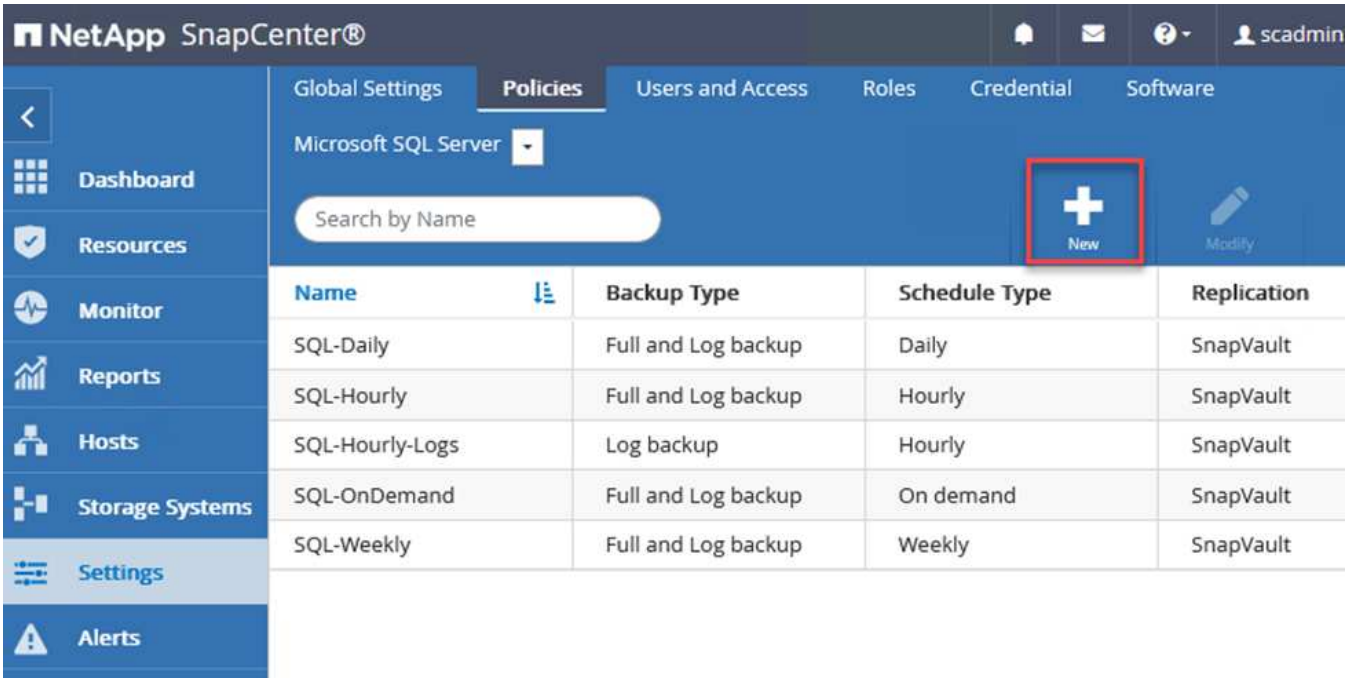
3. For Oracle, fill out the required fields in the Add Host dialog box and select the check box for the Oracle Database plug-in. Then click Submit to begin the discovery process and to add the host to SnapCenter.



▼ **Create SnapCenter policies**

Policies establish the specific rules to be followed for a backup job. They include, but are not limited to, the backup schedule, replication type, and how SnapCenter handles backing up and truncating transaction logs.

You can access policies in the Settings section of the SnapCenter web client.



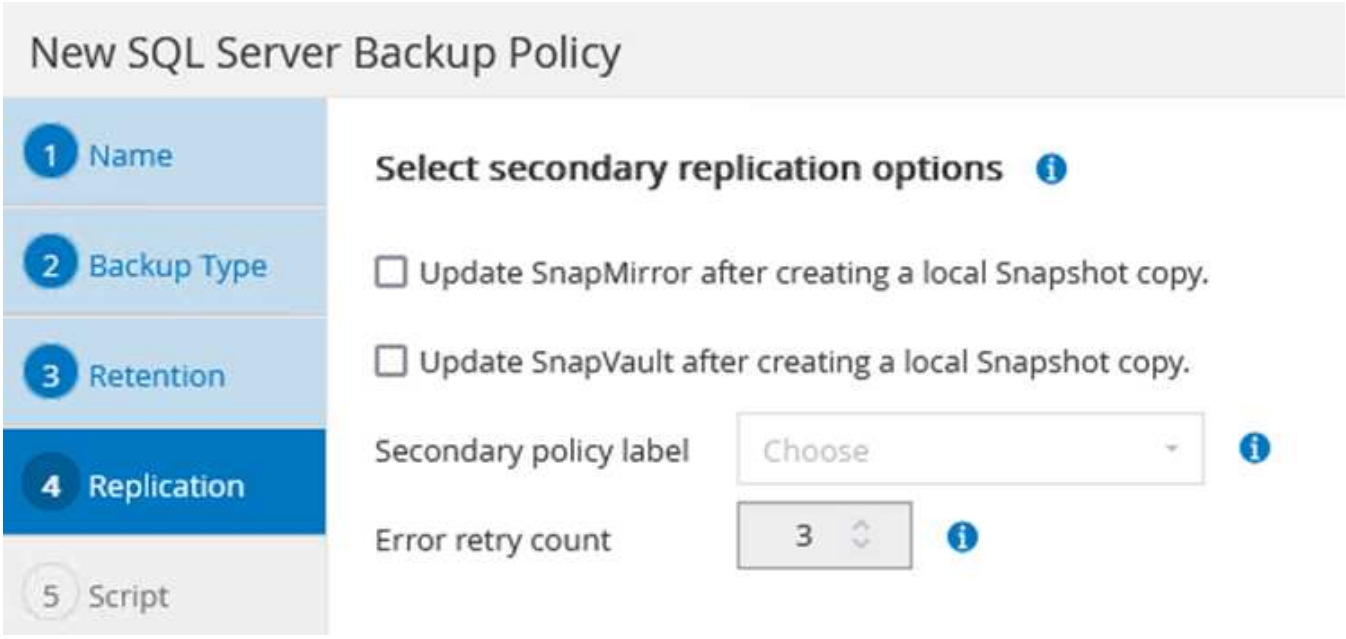
For complete information on creating policies for SQL Server backups, see the SnapCenter documentation.

For complete information on creating policies for Oracle backups, see the SnapCenter documentation.

Notes:

- As you progress through the policy creation wizard, take special note of the Replication section. In this section you stipulate the types of secondary SnapMirror copies that you want taken during the backups process.
- The "Update SnapMirror after creating a local Snapshot copy" setting refers to updating a SnapMirror relationship when that relationship exists between two storage virtual machines residing on the same cluster.
- The "Update SnapVault after creating a local SnapShot copy" setting is used to update a SnapMirror relationship that exists between two separate cluster and between an on-premises ONTAP system and Cloud Volumes ONTAP or FSx ONTAP.

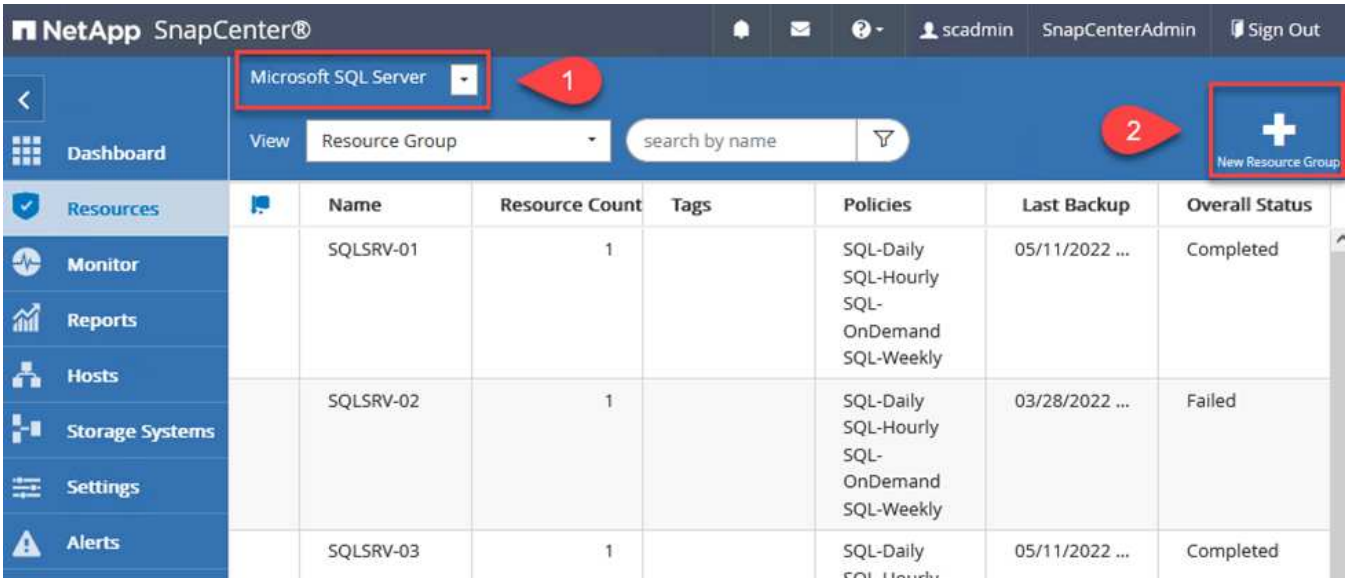
The following image shows the preceding options and how they look in the backup policy wizard.



▼ **Create SnapCenter Resource Groups**

Resource Groups allow you to select the database resources you want to include in your backups and the policies followed for those resources.

1. Go to the Resources section in the left-hand menu.
2. At the top of the window, select the resource type to work with (In this case Microsoft SQL Server) and then click New Resource Group.



The SnapCenter documentation covers step-by-step details for creating Resource Groups for both SQL Server and Oracle databases.

For backing up SQL resources, follow this link.

For Backing up Oracle resources, follow this link.

Deploy and configure Veeam Backup Server

Veeam Backup & Replication software is used in the solution to back up our application virtual machines and archive a copy of the backups to an Amazon S3 bucket using a Veeam scale-out backup repository (SOBR). Veeam is deployed on a Windows server in this solution. For specific guidance on deploying Veeam, see the Veeam help Center Technical documentation.

▼ Configure Veeam scale-out backup repository

After you deploy and license the software, you can create a scale-out backup repository (SOBR) as target storage for backup jobs. You should also include an S3 bucket as a backup of VM data offsite for disaster recovery.

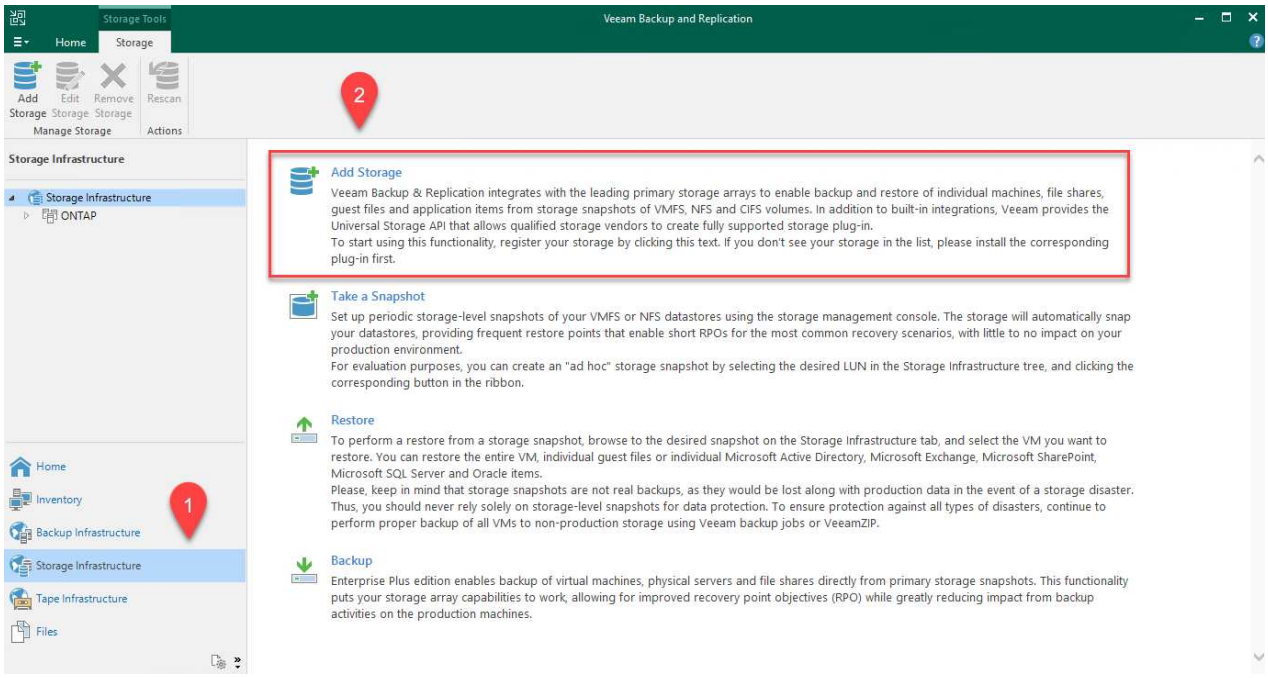
See the following prerequisites before getting started.

1. Create an SMB file share on your on-premises ONTAP system as the target storage for backups.
2. Create an Amazon S3 bucket to include in the SOBR. This is a repository for the offsite backups.

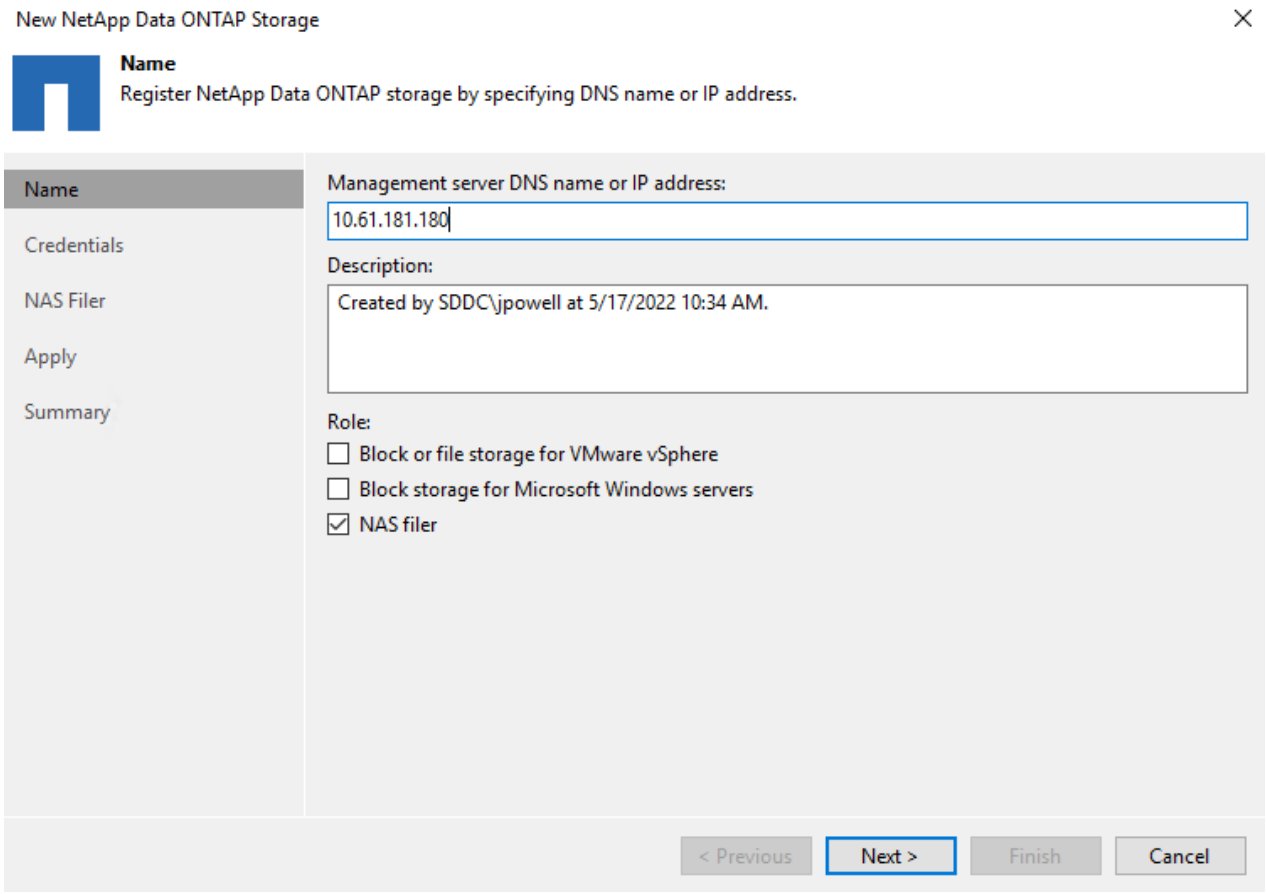
▼ Add ONTAP Storage to Veeam

First, add the ONTAP storage cluster and associated SMB/NFS filesystem as storage infrastructure in Veeam.

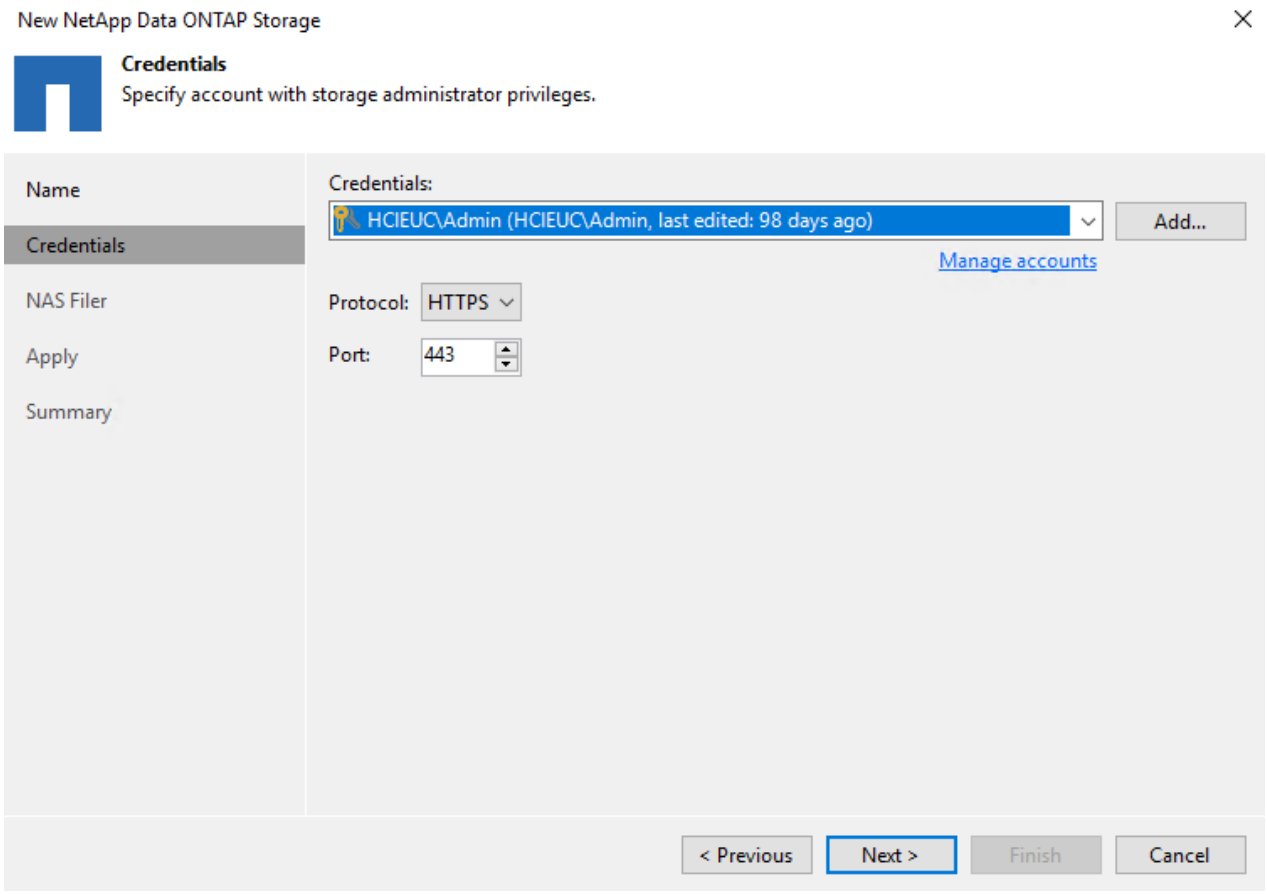
1. Open the Veeam console and log in. Navigate to Storage Infrastructure and then select Add Storage.



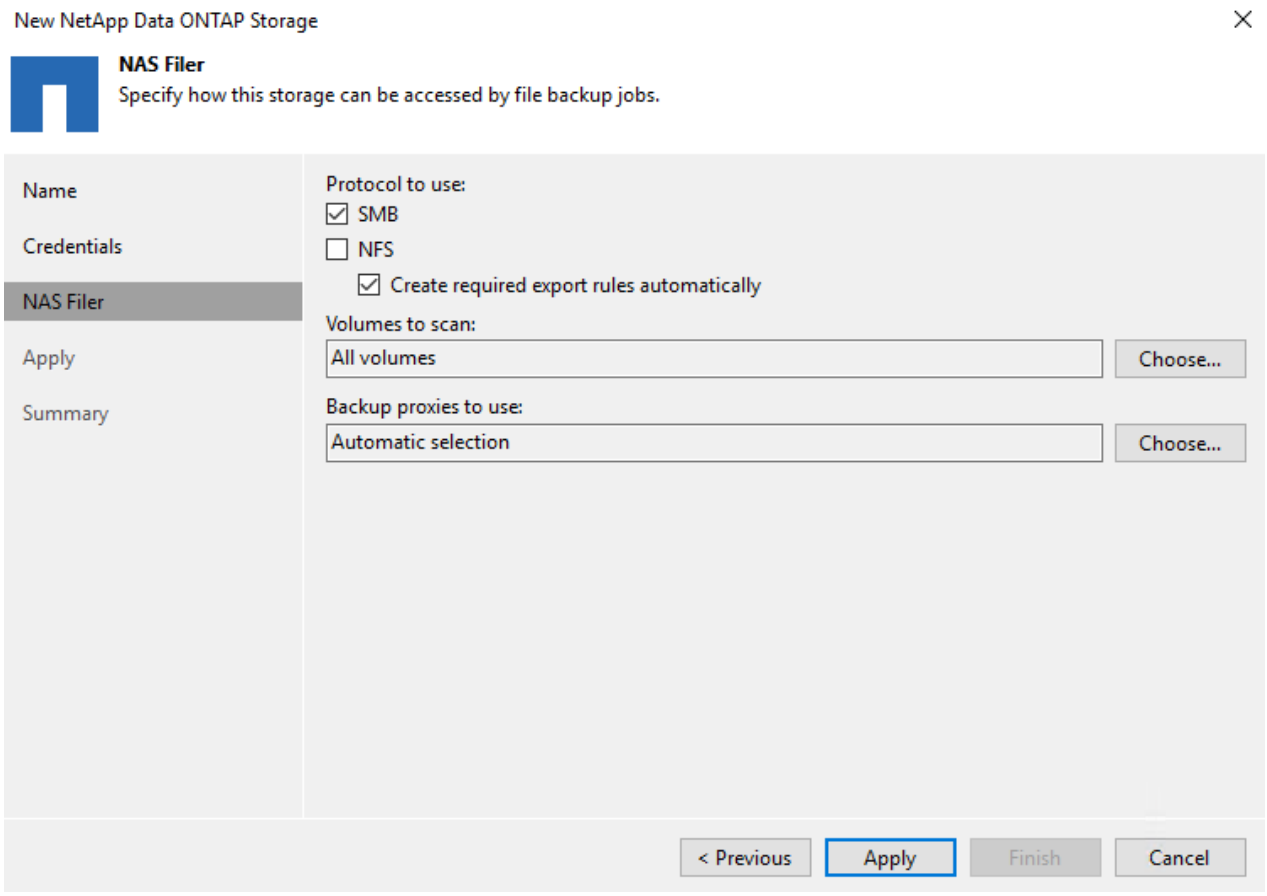
2. In the Add Storage wizard, select NetApp as the storage vendor and then select Data ONTAP.
3. Enter the management IP address and check the NAS Filer box. Click Next.



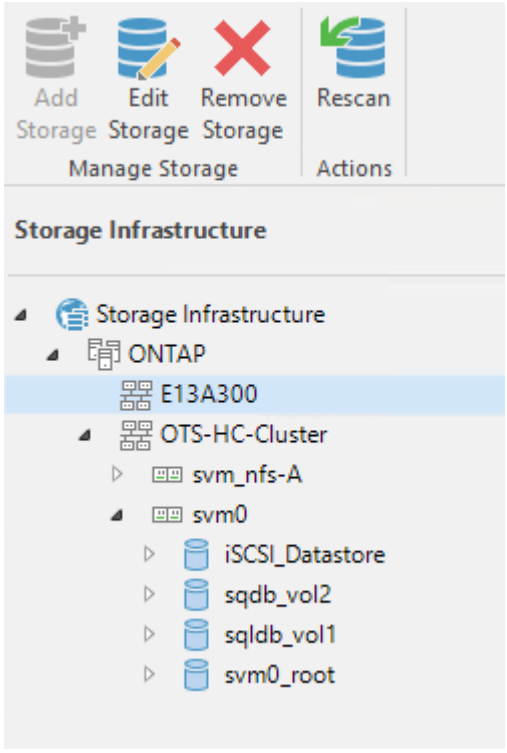
4. Add your credentials to access the ONTAP cluster.



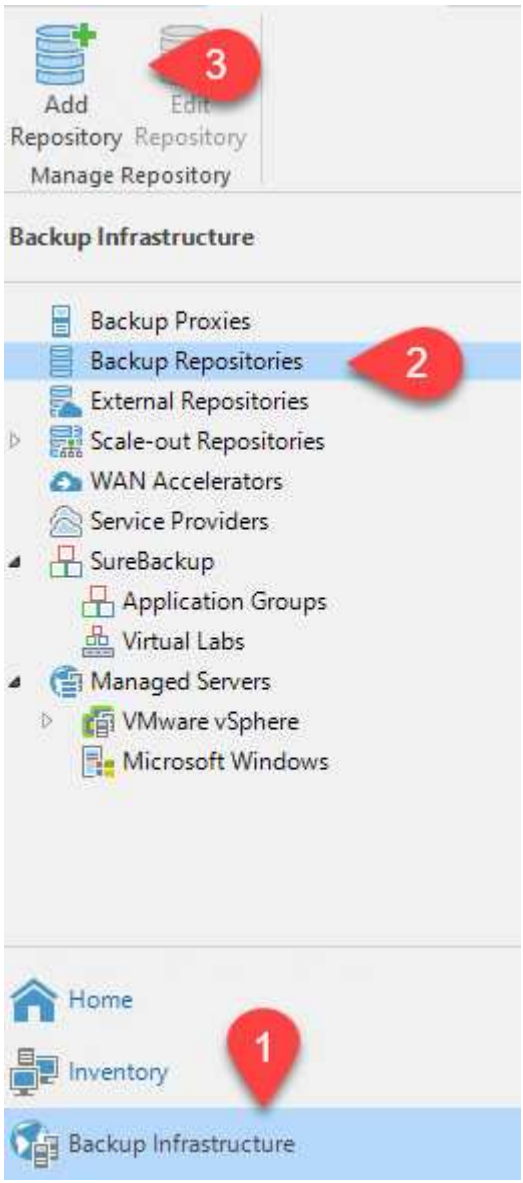
5. On the NAS Filer page choose the desired protocols to scan and select Next.



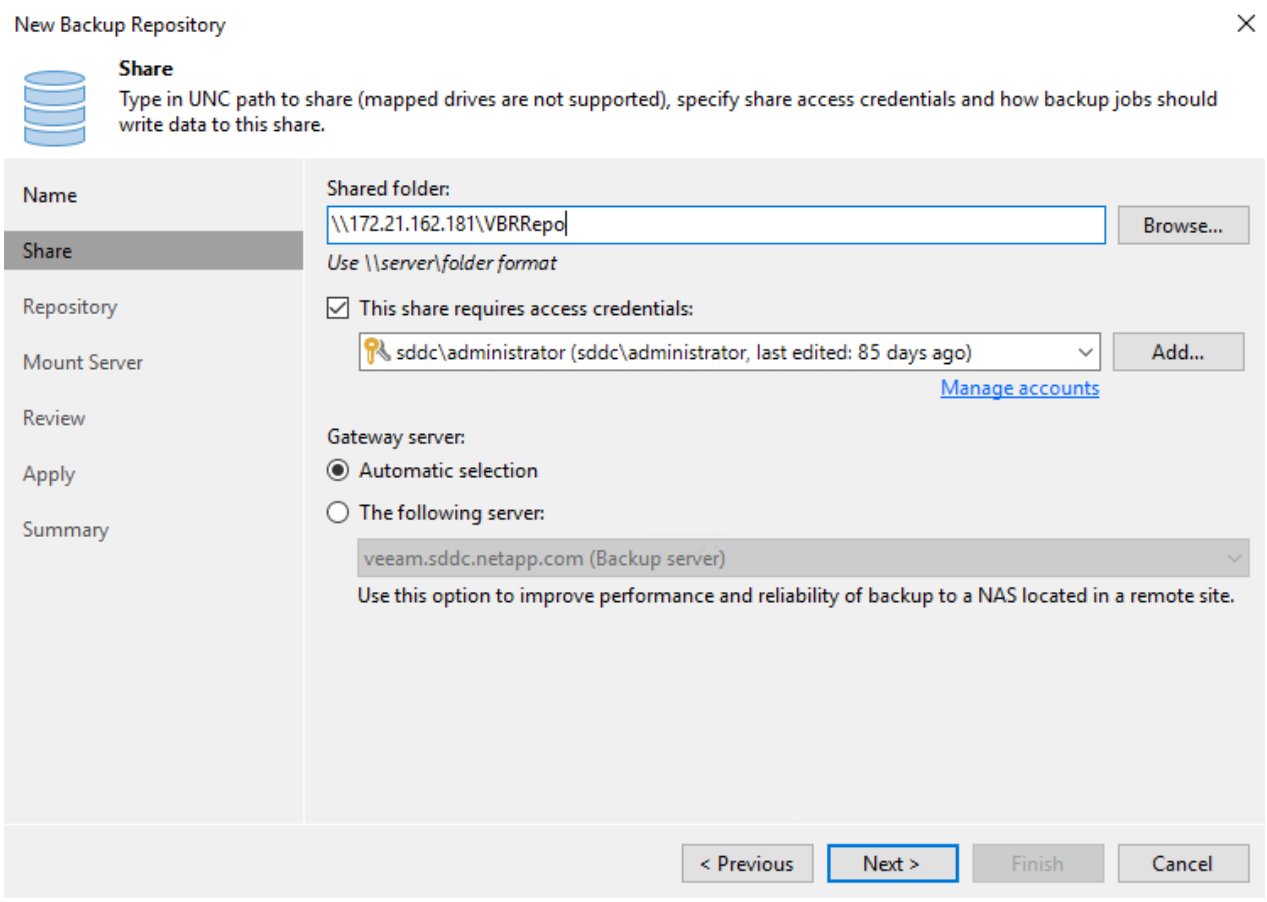
6. Complete the Apply and Summary pages of the wizard and click Finish to begin the storage discovery process. After the scan completes, the ONTAP cluster is added along with the NAS filers as available resources.



7. Create a backup repository using the newly discovered NAS shares. From Backup Infrastructure, select Backup Repositories and click the Add Repository menu item.



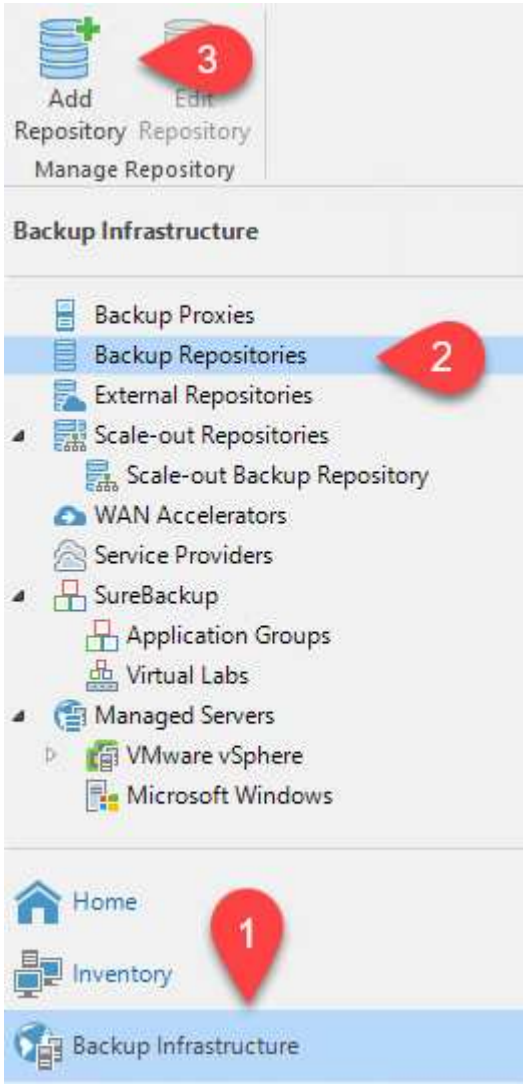
8. Follow all steps in the New Backup Repository Wizard to create the repository. For detailed information on creating Veeam Backup Repositories, see the Veeam documentation.



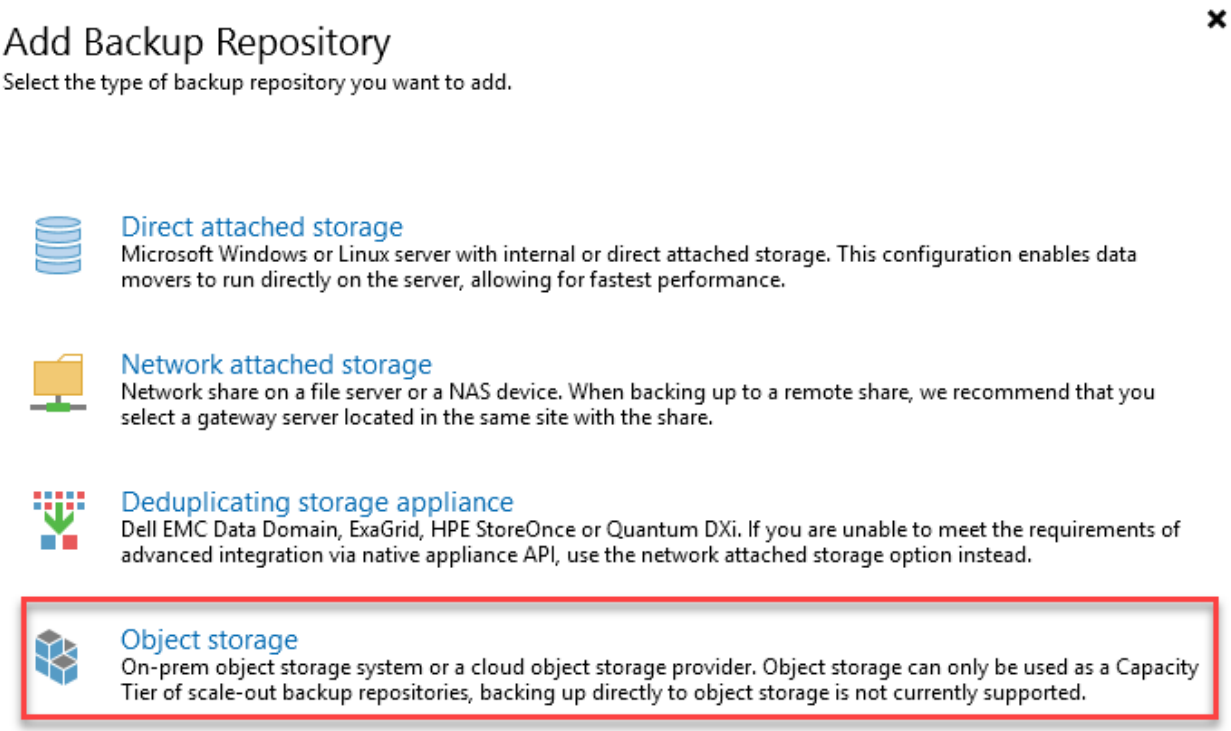
▼ Add the Amazon S3 bucket as a backup repository

The next step is to add the Amazon S3 storage as a backup repository.

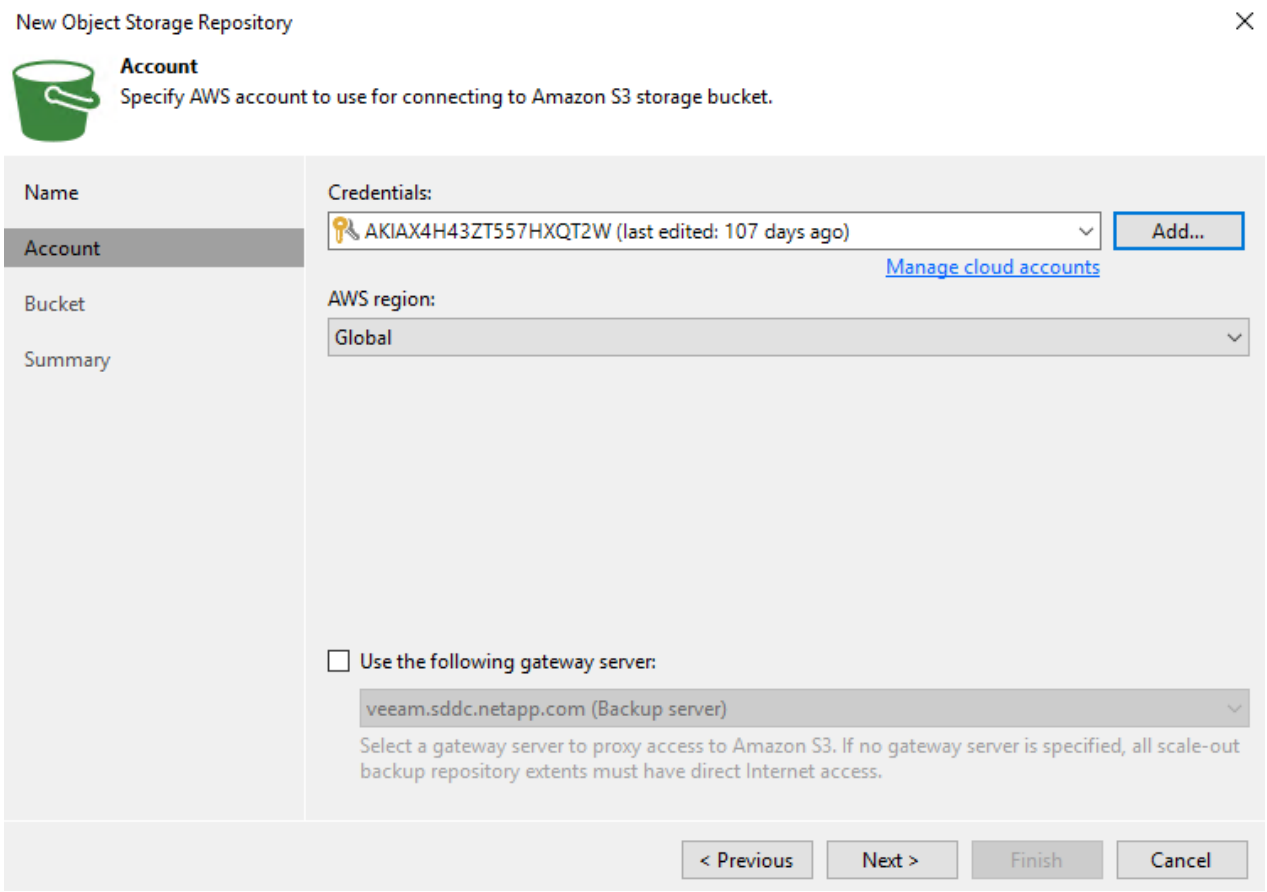
1. Navigate to Backup Infrastructure > Backup Repositories. Click Add Repository.



2. In the Add Backup Repository wizard, select Object Storage and then Amazon S3. This starts the New Object Storage Repository wizard.



3. Provide a name for your object storage repository and click Next.
4. In the next section, provide your credentials. You need an AWS Access Key and Secret Key.



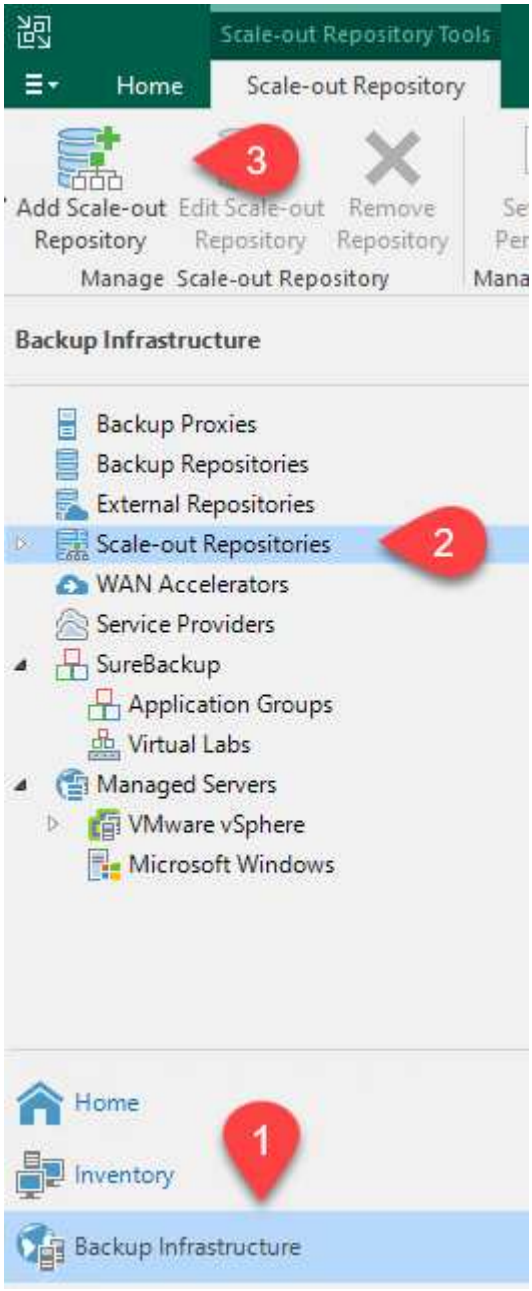
5. After the Amazon configuration loads, choose your datacenter, bucket, and folder and click Apply.

Finally, click Finish to close out the wizard.

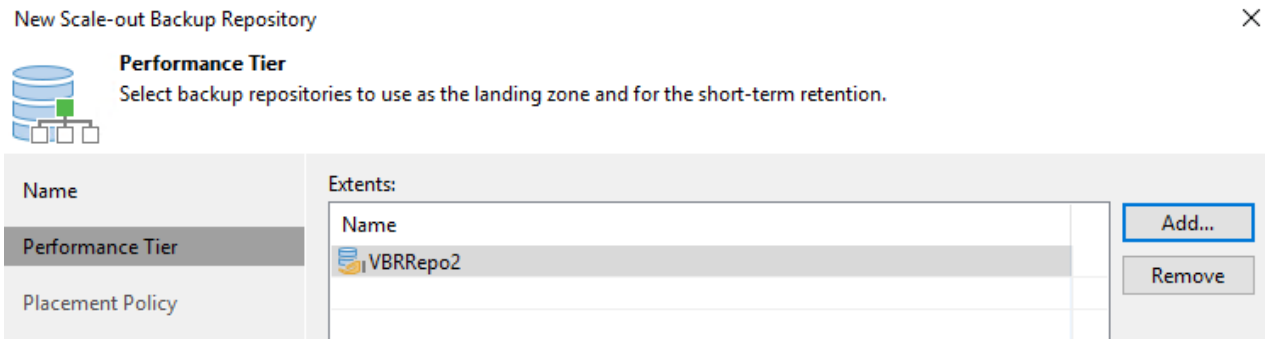
▼ **Create scale-out backup repository**

Now that we have added our storage repositories to Veeam, we can create the SOBR to automatically tier backup copies to our offsite Amazon S3 object storage for disaster recovery.

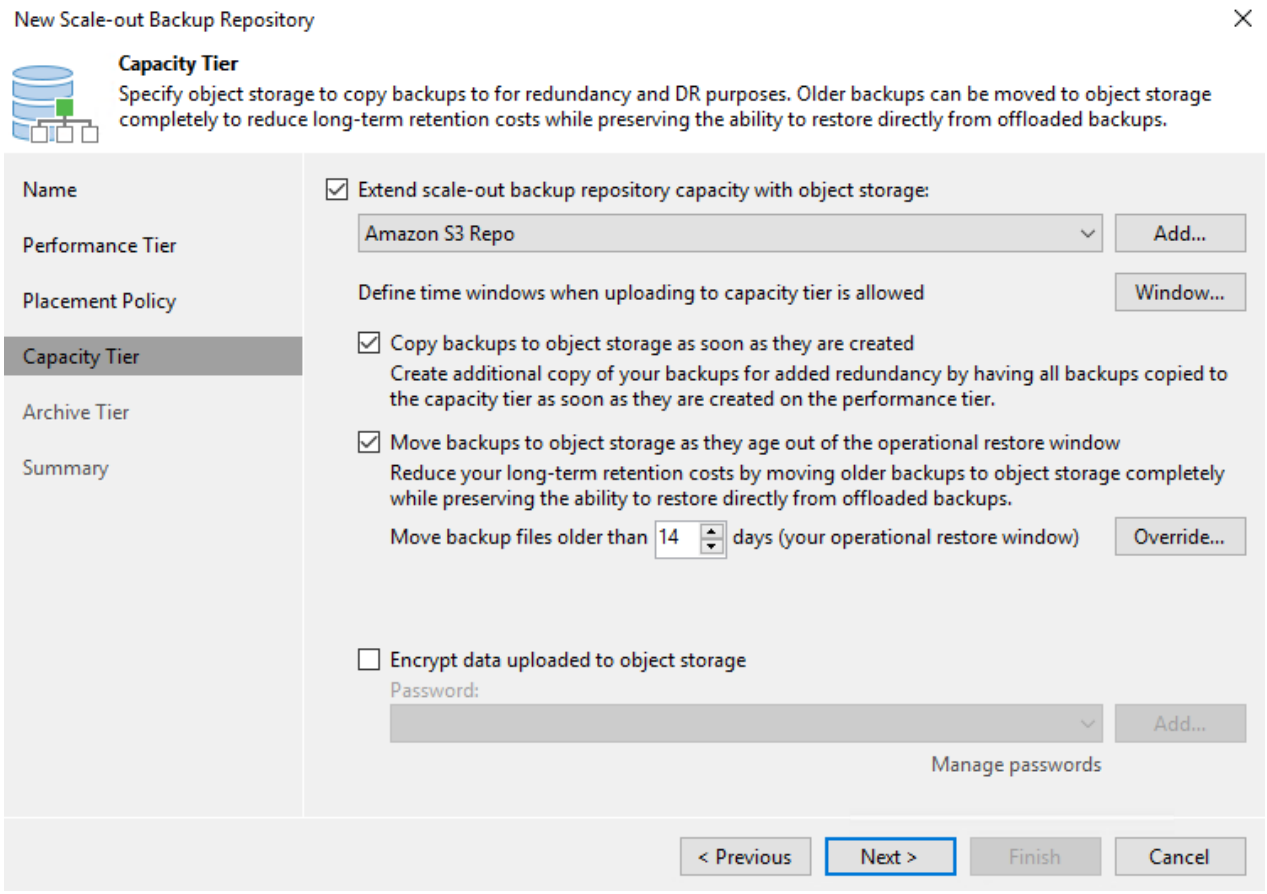
1. From Backup Infrastructure, select Scale-out Repositories and then click the Add Scale-out Repository menu item.



2. In the New Scale-out Backup Repository provide a name for the SOBR and click Next.
3. For the Performance Tier, choose the backup repository that contains the SMB share residing on your local ONTAP cluster.



4. For the Placement Policy, choose either Data Locality or Performance based your requirements. Select next.
5. For Capacity Tier we extend the SOBR with Amazon S3 object storage. For the purposes of disaster recovery, select Copy Backups to Object Storage as Soon as They are Created to ensure timely delivery of our secondary backups.



6. Finally, select Apply and Finish to finalize creation of the SOBR.

▼ Create the scale-out backup repository jobs

The final step to configuring Veeam is to create backup jobs using the newly created SOBR as the backup destination. Creating backup jobs is a normal part of any storage administrator's repertoire and we do not cover the detailed steps here. For more complete information on creating backup jobs in Veeam, see the Veeam Help Center Technical Documentation.

NetApp Backup and Recovery tools and configuration

To conduct a failover of application VMs and database volumes to VMware Cloud Volume services running in AWS, you must install and configure a running instance of both SnapCenter Server and Veeam Backup and Replication Server. After the failover is complete, you must also configure these tools to resume normal backup operations until a failback to the on-premises datacenter is planned and executed.

▼ **Deploy secondary Windows SnapCenter Server**

SnapCenter Server is deployed in the VMware Cloud SDDC or installed on an EC2 instance residing in a VPC with network connectivity to the VMware Cloud environment.

SnapCenter software is available from the NetApp support site and can be installed on Microsoft Windows systems that reside either in a domain or workgroup. A detailed planning guide and installation instructions can be found at the NetApp documentation center.

You can find the SnapCenter software at this link.

▼ **Configure secondary Windows SnapCenter Server**

To perform a restore of application data mirrored to FSx ONTAP, you must first perform a full restore of the on-premises SnapCenter database. After this process is complete, communication with the VMs is reestablished and application backups can now resume using FSx ONTAP as the primary storage.

To achieve this, you must complete the following items on the SnapCenter Server:

1. Configure the computer name to be identical to the original on-premises SnapCenter Server.
2. Configure networking to communicate with VMware Cloud and the FSx ONTAP instance.
3. Complete the procedure to restore the SnapCenter database.
4. Confirm that SnapCenter is in Disaster Recovery mode to make sure that FSx is now the primary storage for backups.
5. Confirm that communication is reestablished with the restored virtual machines.

▼ **Deploy secondary Veeam Backup & Replication server**

You can install the Veeam Backup & Replication server on a Windows server in the VMware Cloud on AWS or on an EC2 instance. For detailed implementation guidance, see the Veeam Help Center Technical Documentation.

▼ **Configure secondary Veeam Backup & Replication server**

To perform a restore of virtual machines that have been backed up to Amazon S3 storage, you must install the Veeam Server on a Windows server and configure it to communicate with VMware Cloud, FSx ONTAP, and the S3 bucket that contains the original backup repository. It must also have a new backup repository configured on FSx ONTAP to conduct new backups of the VMs after they are restored.

To perform this process, the following items must be completed:

1. Configure networking to communicate with VMware Cloud, FSx ONTAP, and the S3 bucket containing the original backup repository.
2. Configure an SMB share on FSx ONTAP to be a new backup repository.
3. Mount the original S3 bucket that was used as part of the scale-out backup repository on premises.
4. After restoring the VM, establish new backup jobs to protect SQL and Oracle VMs.

For more information on restoring VMs using Veeam, see the section "Restore Application VMs with Veeam

Full Restore".

SnapCenter database backup for disaster recovery

SnapCenter allows for the backup and recovery of its underlying MySQL database and configuration data for the purpose of recovering the SnapCenter server in the case of a disaster. For our solution, we recovered the SnapCenter database and configuration on an AWS EC2 instance residing in our VPC. For more information on disaster recovery of SnapCenter, see this link.

▼ SnapCenter backup prerequisites

The following prerequisites are required for SnapCenter backup:

- A volume and SMB share created on the on-premises ONTAP system to locate the backed-up database and configuration files.
- A SnapMirror relationship between the on-premises ONTAP system and FSx or CVO in the AWS account. This relationship is used for transporting the snapshot containing the backed-up SnapCenter database and configuration files.
- Windows Server installed in the cloud account, either on an EC2 instance or on a VM in the VMware Cloud SDDC.
- SnapCenter installed on the Windows EC2 instance or VM in VMware Cloud.

▼ SnapCenter backup and restore process summary

- Create a volume on the on-premises ONTAP system for hosting the backup db and config files.
- Set up a SnapMirror relationship between on-premises and FSx/CVO.
- Mount the SMB share.
- Retrieve the Swagger authorization token for performing API tasks.
- Start the db restore process.
- Use the xcopy utility to copy the db and config file local directory to the SMB share.
- On FSx, create a clone of the ONTAP volume (copied via SnapMirror from on-premises).
- Mount the SMB share from FSx to EC2/VMware Cloud.
- Copy the restore directory from the SMB share to a local directory.
- Run the SQL Server restore process from Swagger.

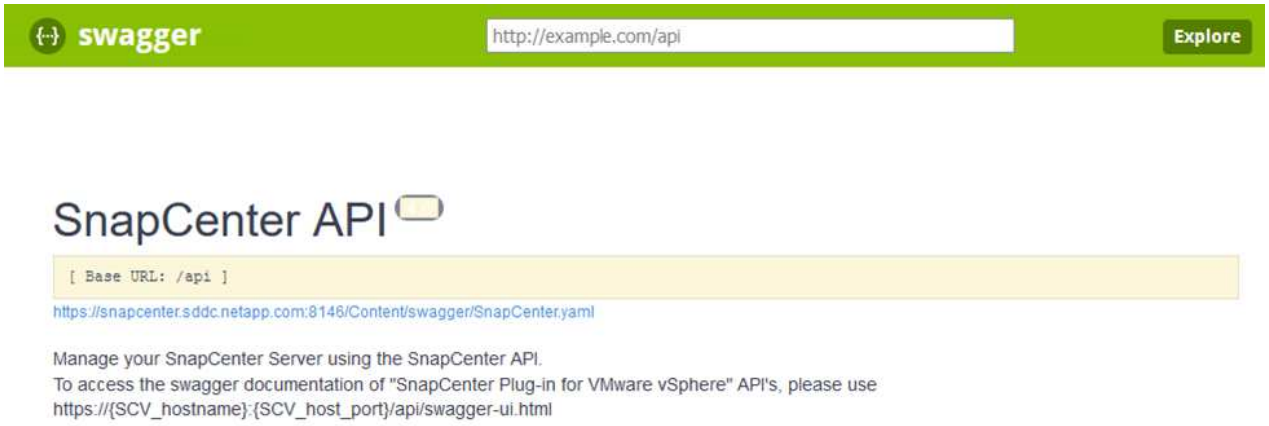
▼ Back up the SnapCenter database and configuration

SnapCenter provides a web client interface for executing REST API commands. For information on accessing the REST APIs through Swagger, see the SnapCenter documentation at this link.

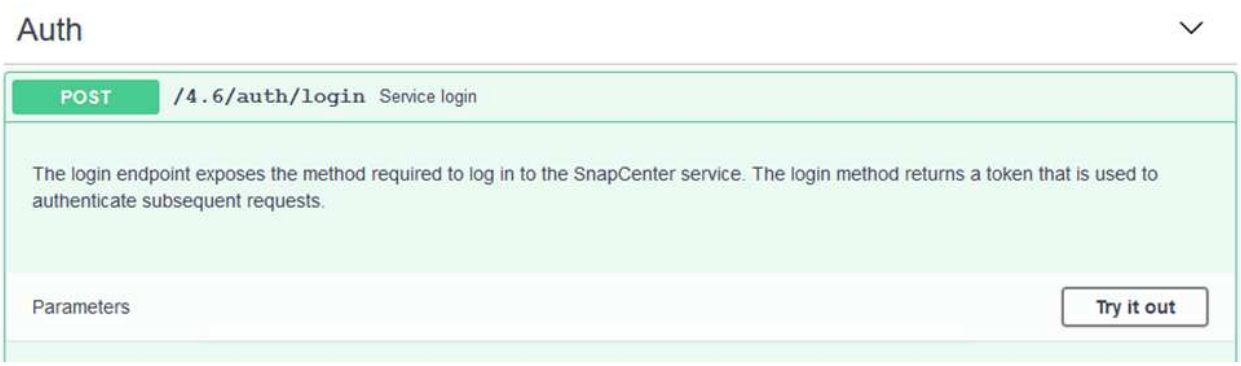
▼ Log into Swagger and obtain authorization token

After you have navigated to the Swagger page, you must retrieve an authorization token to initiate the database restore process.

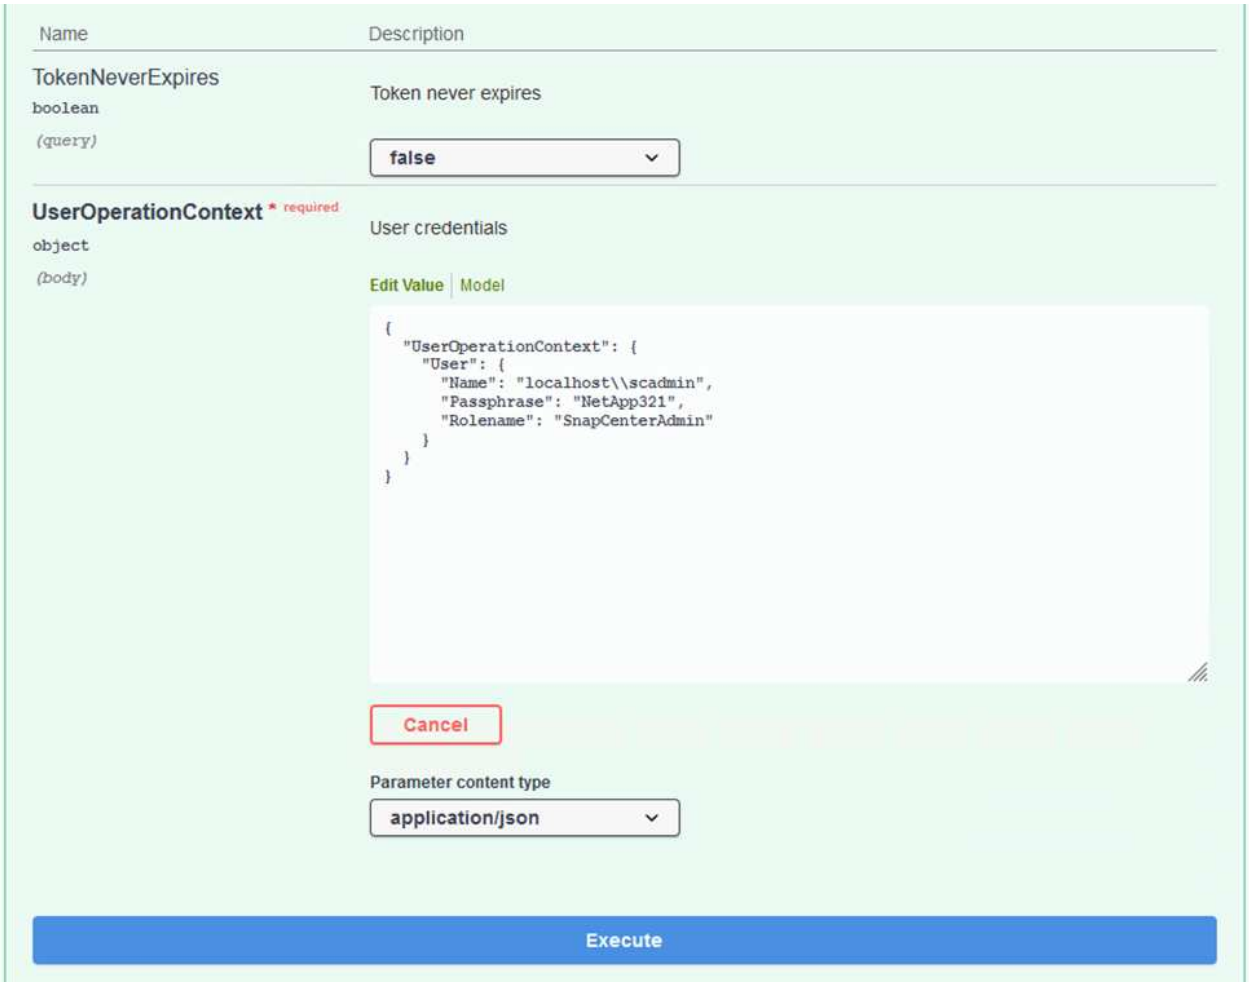
1. Access the SnapCenter Swagger API web page at *https://<SnapCenter Server IP>:8146/swagger/*.



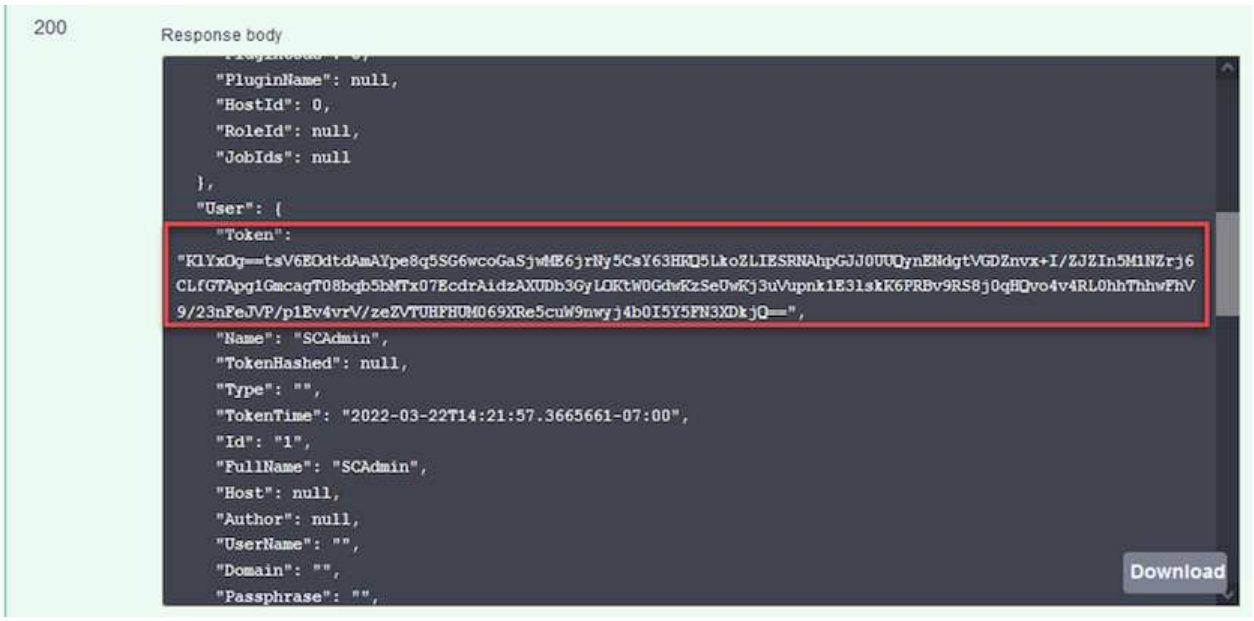
2. Expand the Auth section and click Try it Out.



3. In the UserOperationContext area, fill in the SnapCenter credentials and role and click Execute.



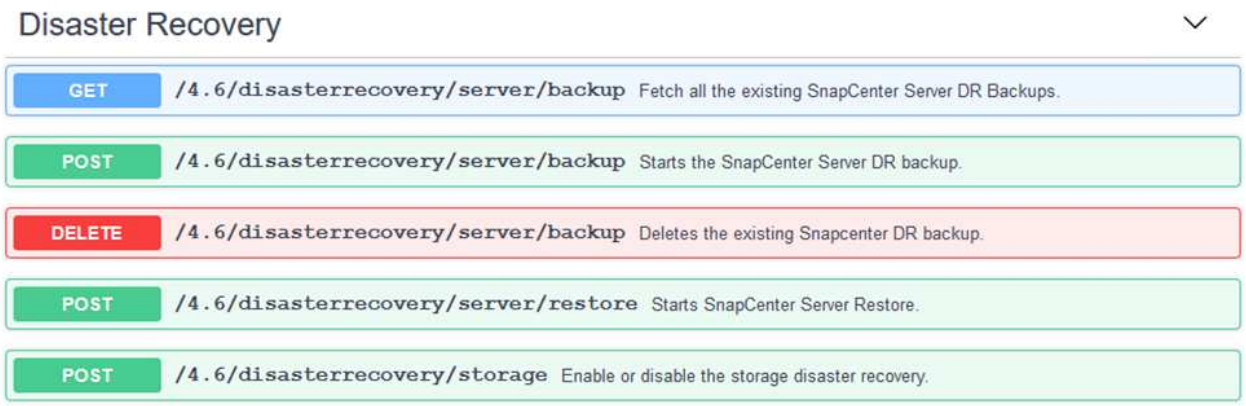
4. In the Response body below, you can see the token. Copy the token text for authentication when executing the backup process.



▼ **Perform a SnapCenter database backup**

Next go to the Disaster Recovery area on the Swagger page to begin the SnapCenter backup process.

1. Expand the Disaster Recovery area by clicking it.

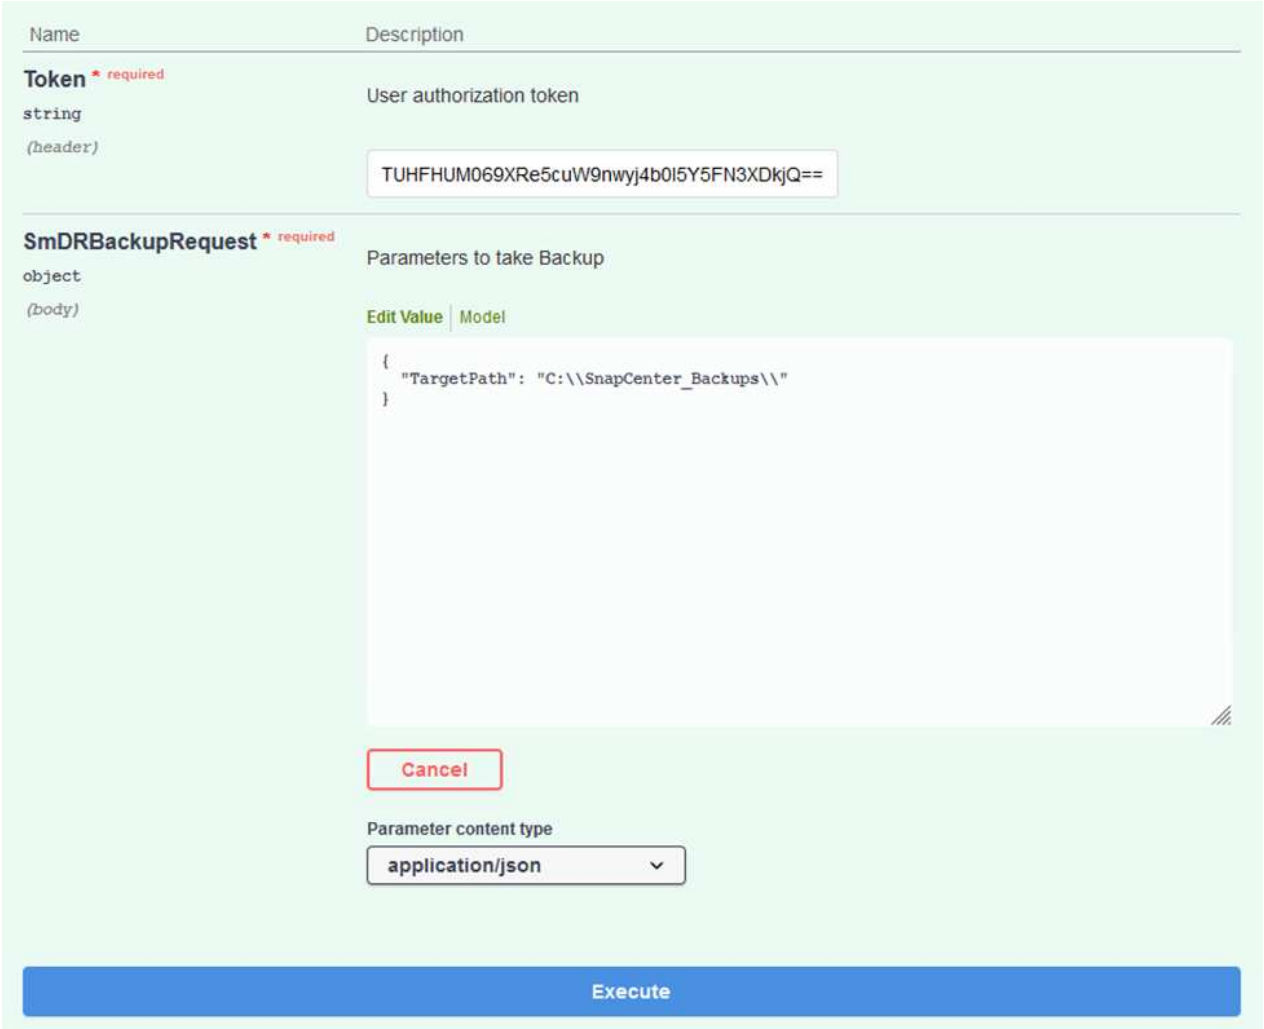


2. Expand the `/4.6/disasterrecovery/server/backup` section and click Try it Out.



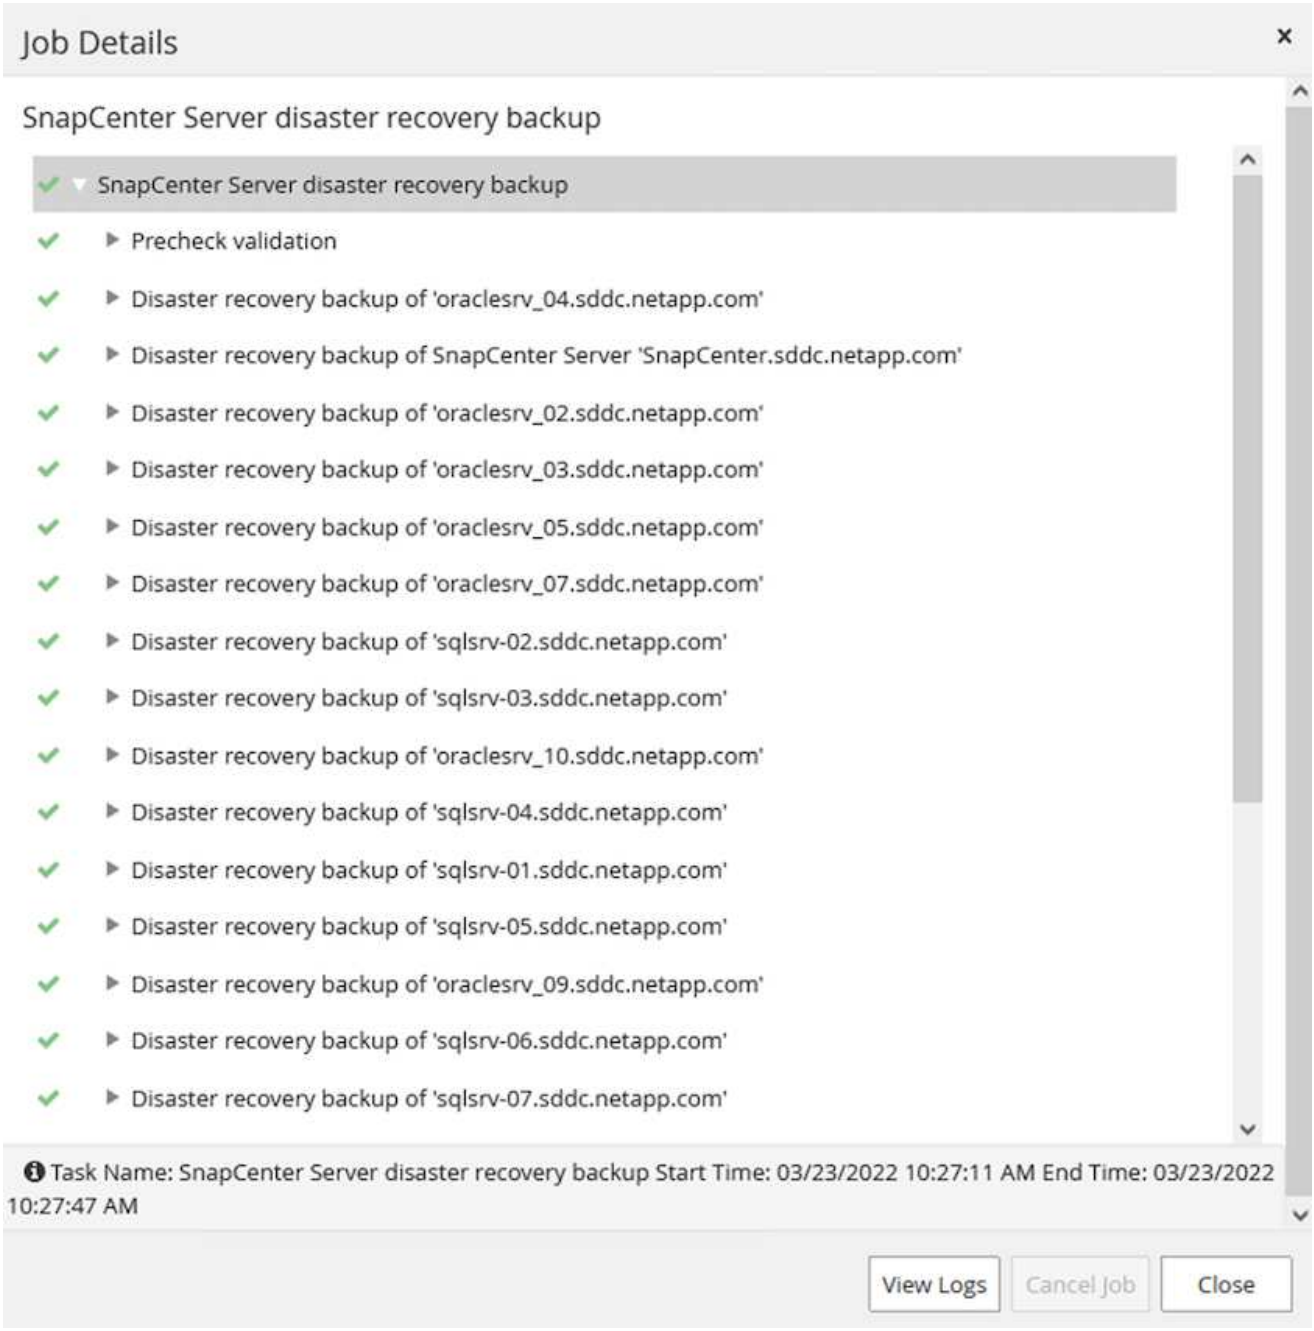
3. In the SmDRBackupRequest section, add the correct local target path and select Execute to start the backup of the SnapCenter database and configuration.

The backup process does not allow backing up directly to an NFS or CIFS file share.



▼ **Monitor the backup job from SnapCenter**

Log into SnapCenter to review log files when starting the database restore process. Under the Monitor section, you can view the details of the SnapCenter server disaster recovery backup.



▼ **Use XCOPY utility to copy the database backup file to the SMB share**

Next you must move the backup from the local drive on the SnapCenter server to the CIFS share that is used to SnapMirror copy the data to the secondary location located on the FSx instance in AWS. Use xcopy with specific options that retain the permissions of the files.

Open a command prompt as Administrator. From the command prompt, enter the following commands:

```
xcopy  <Source_Path>  \\<Destination_Server_IP>\<Folder_Path> /O /X /E
/H /K
xcopy c:\SC_Backups\SnapCenter_DR \\10.61.181.185\snapcenter_dr /O /X
/E /H /K
```

Failover

▼ Disaster occurs at primary site

For a disaster that occurs at the primary on-premises datacenter, our scenario includes failover to a secondary site residing on Amazon Web Services infrastructure using VMware Cloud on AWS. We assume that the virtual machines and our on-premises ONTAP cluster are no longer accessible. In addition, both the SnapCenter and Veeam virtual machines are no longer accessible and must be rebuilt at our secondary site.

This section address failover of our infrastructure to the cloud, and we cover the following topics:

- SnapCenter database restore. After a new SnapCenter server has been established, restore the MySQL database and configuration files and toggle the database into disaster recovery mode in order to allow the secondary FSx storage to become the primary storage device.
- Restore the application virtual machines using Veeam Backup & Replication. Connect the S3 storage that contains the VM backups, import the backups, and restore them to VMware Cloud on AWS.
- Restore the SQL Server application data using SnapCenter.
- Restore the Oracle application data using SnapCenter.

▼ SnapCenter database restore process

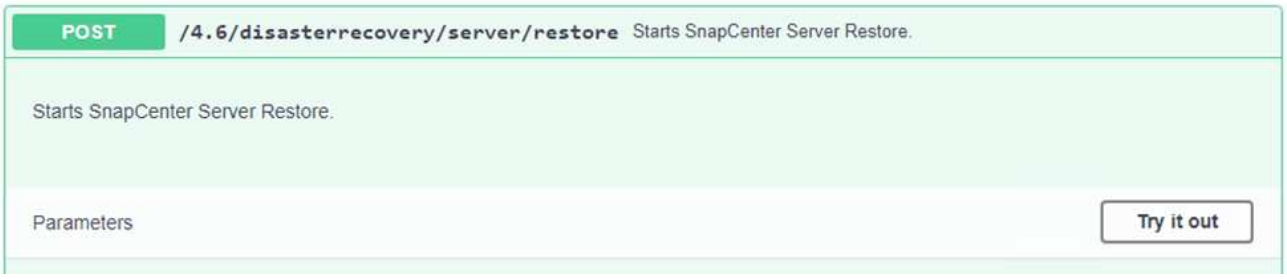
SnapCenter supports disaster recovery scenarios by allowing the backup and restore of its MySQL database and configuration files. This allows an administrator to maintain regular backups of the SnapCenter database at the on-premises datacenter and later restore that database to a secondary SnapCenter database.

To access the SnapCenter backup files on the remote SnapCenter server, complete the following steps:

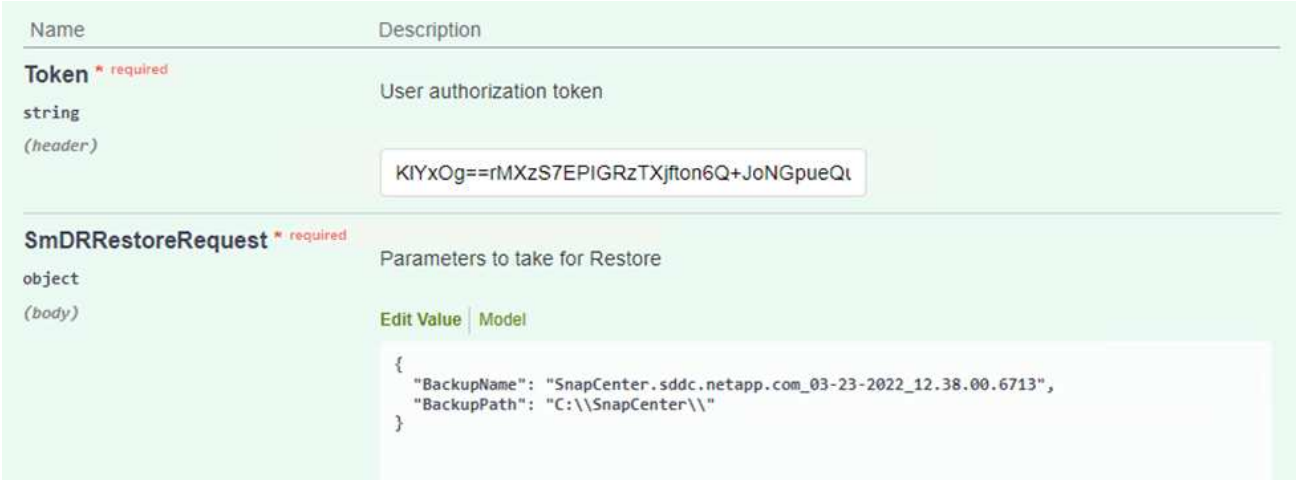
1. Break the SnapMirror relationship from the FSx cluster, which makes the volume read/write.
2. Create a CIFS server (if necessary) and create a CIFS share pointing to the junction path of the cloned volume.
3. Use xcopy to copy the backup files to a local directory on the secondary SnapCenter system.
4. Install SnapCenter v4.6.
5. Ensure that SnapCenter server has the same FQDN as the original server. This is required for the db restore to be successful.

To start the restore process, complete the following steps:

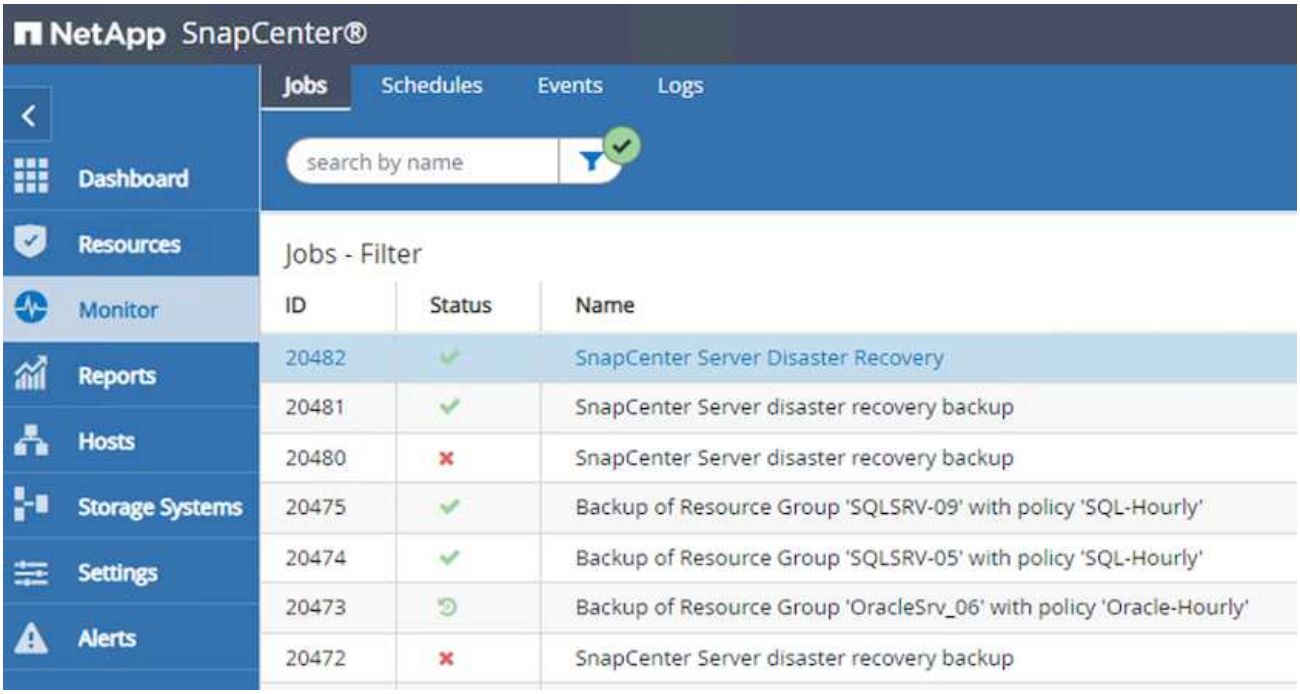
1. Navigate to the Swagger API web page for the secondary SnapCenter server and follow the previous instructions to obtain an authorization token.
2. Navigate to the Disaster Recovery section of the Swagger page, select `/4.6/disasterrecovery/server/restore`, and click Try it Out.

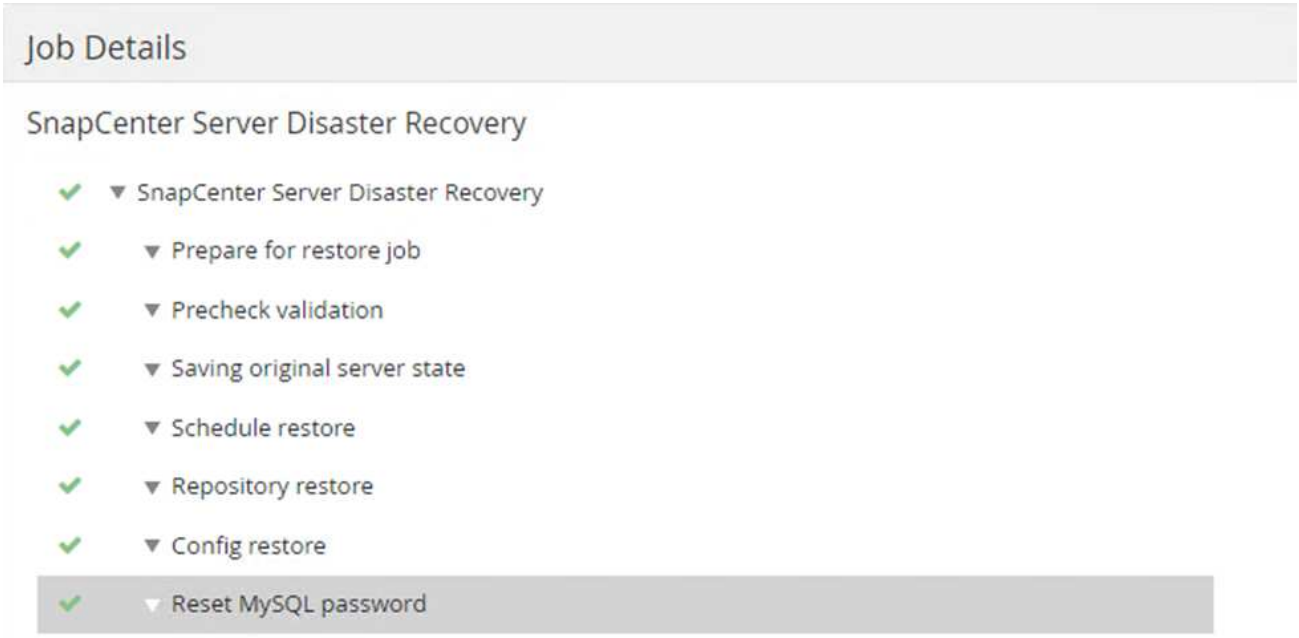


3. Paste in your authorization token and, in the SmDRResterRequest section, paste in the name of the backup and the local directory on the secondary SnapCenter server.

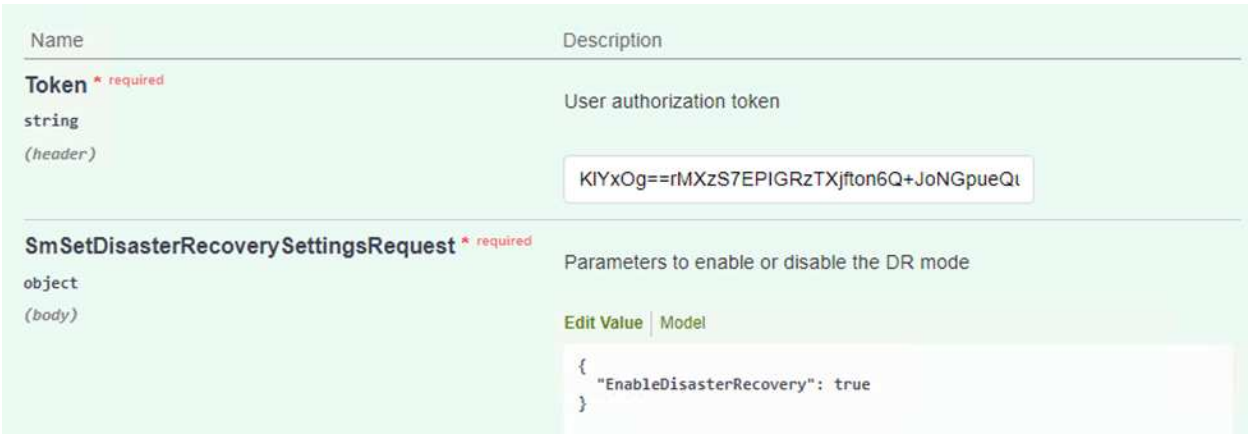


4. Select the Execute button to start the restore process.
5. From SnapCenter, navigate to the Monitor section to view the progress of the restore job.





6. To enable SQL Server restores from secondary storage, you must toggle the SnapCenter database into Disaster Recovery mode. This is performed as a separate operation and initiated on the Swagger API web page.
   a. Navigate to the Disaster Recovery section and click `/4.6/disasterrecovery/storage`.
   b. Paste in the user authorization token.
   c. In the SmSetDisasterRecoverySettingsRequest section, change `EnableDisasterRecover` to `true`.
   d. Click Execute to enable disaster recovery mode for SQL Server.



See comments regarding additional procedures.

Restore application VMs with Veeam full restore

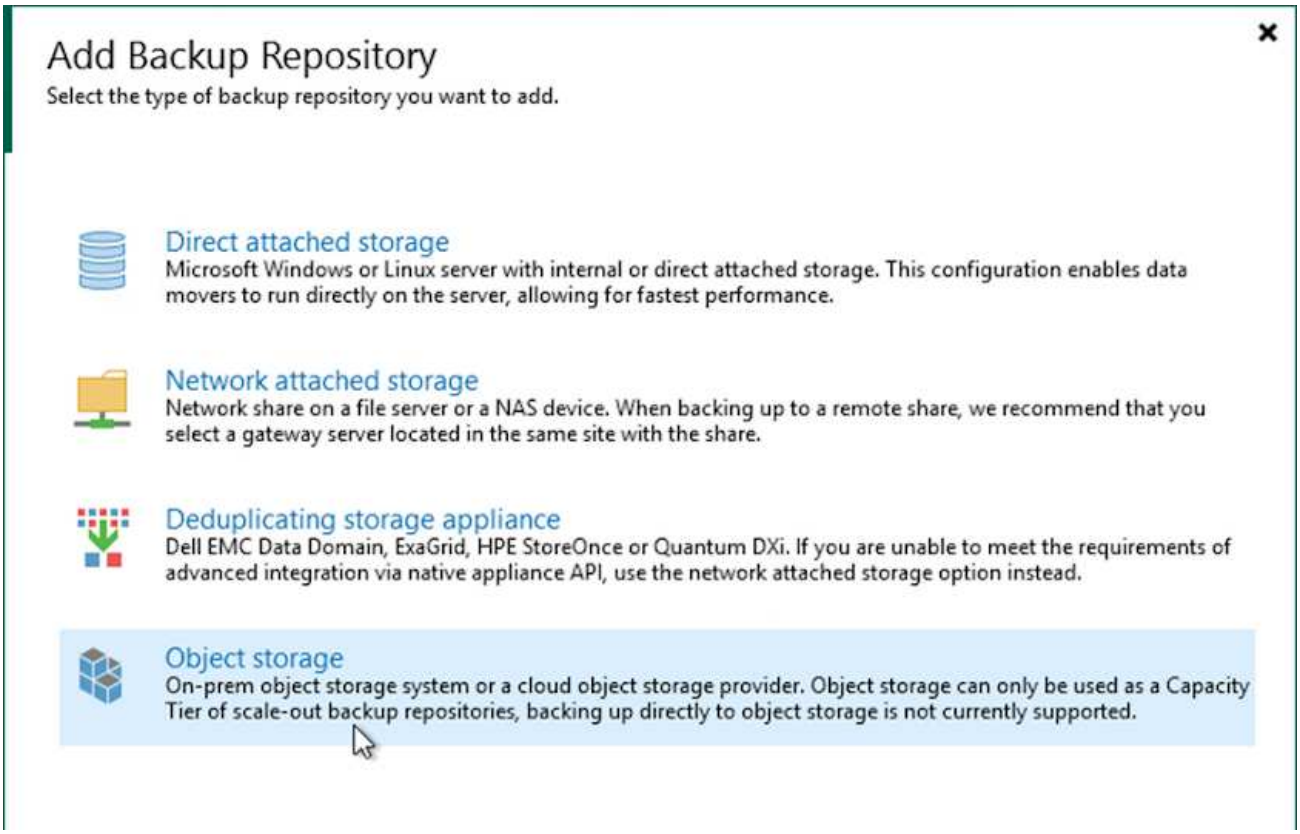
▼ Create a backup repository and import backups from S3

From the secondary Veeam server, import the backups from S3 storage and restore the SQL Server and Oracle VMs to your VMware Cloud cluster.

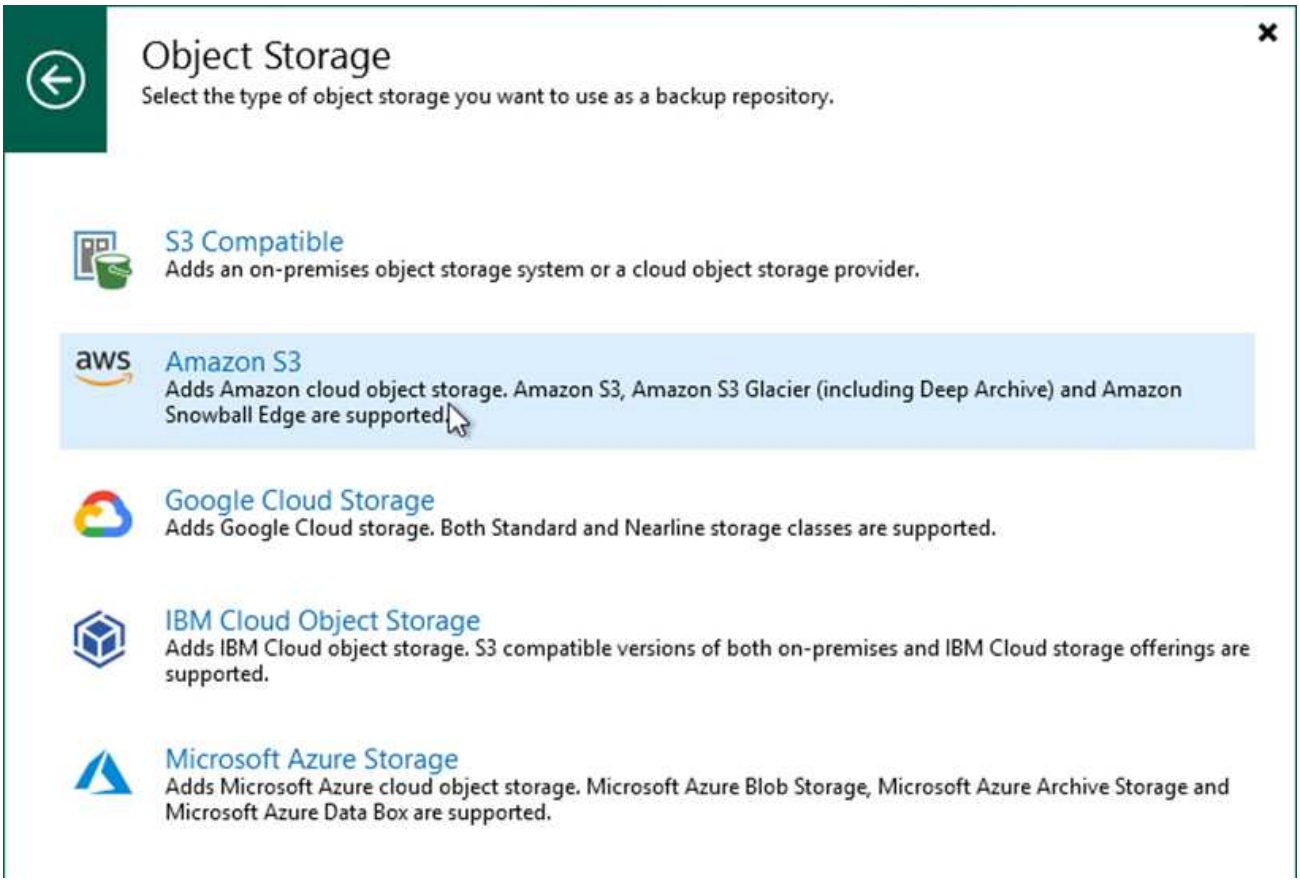
To import the backups from the S3 object that was part of the on-premises scale-out backup repository,

complete the following steps:

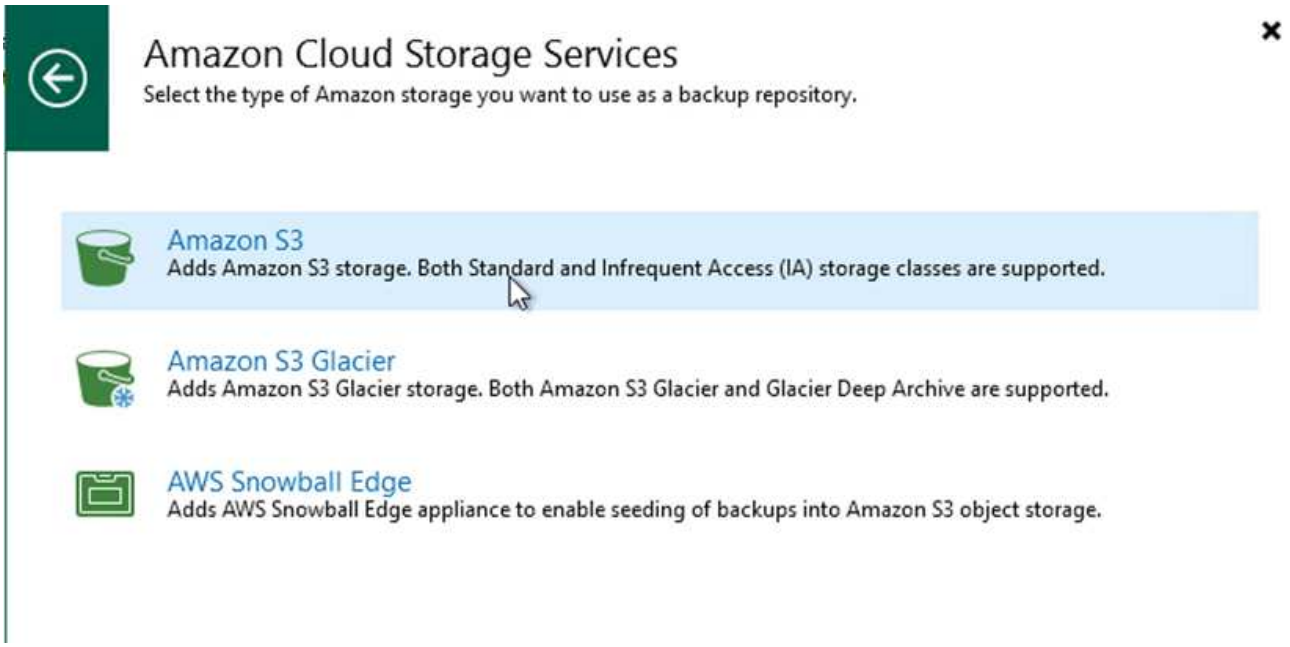
1. Go to Backup Repositories and click Add Repository in the top menu to launch the Add Backup Repository wizard. On the first page of the wizard, select Object Storage as the backup repository type.



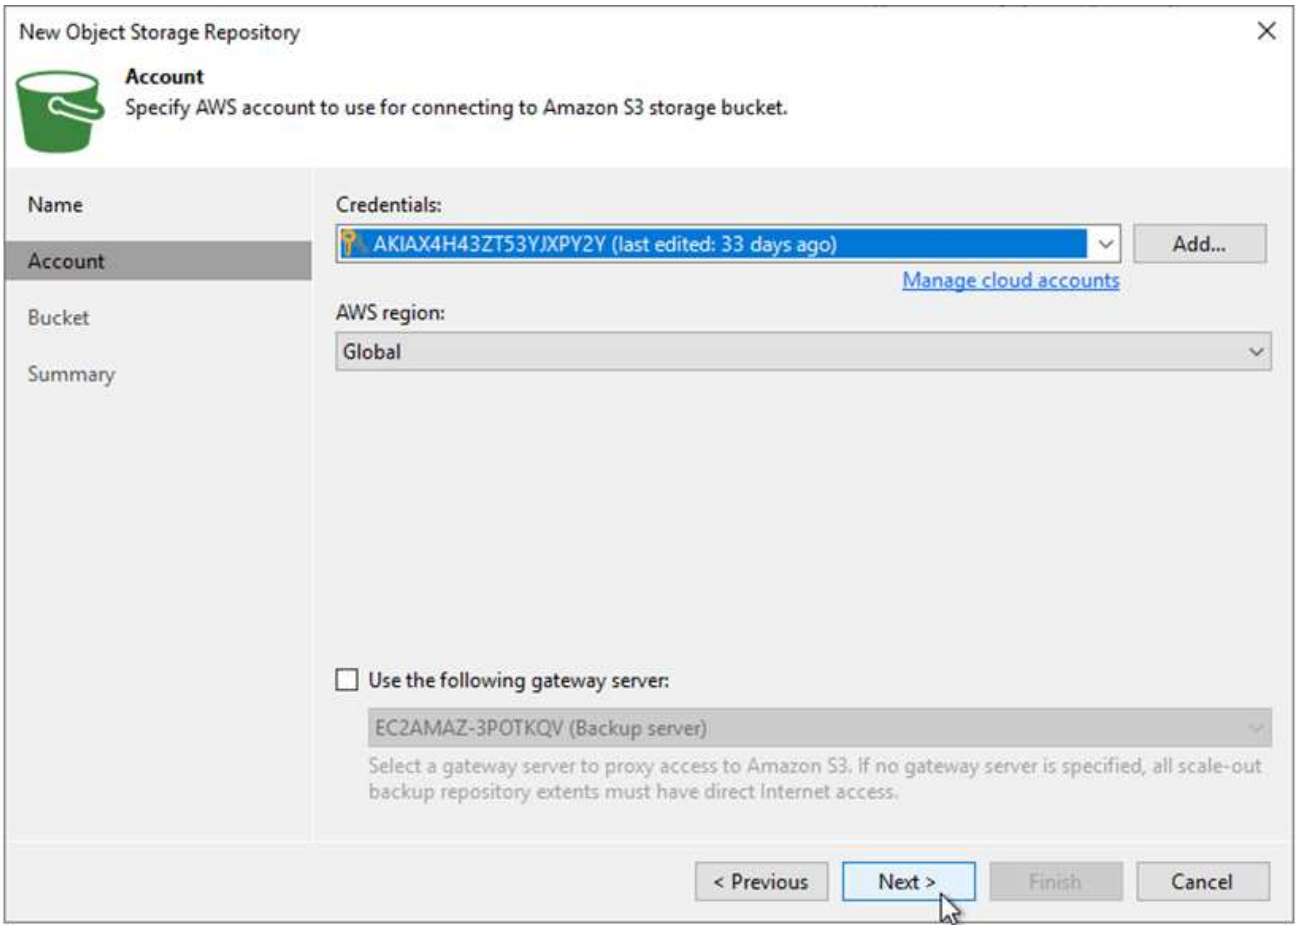
2. Select Amazon S3 as the Object Storage type.



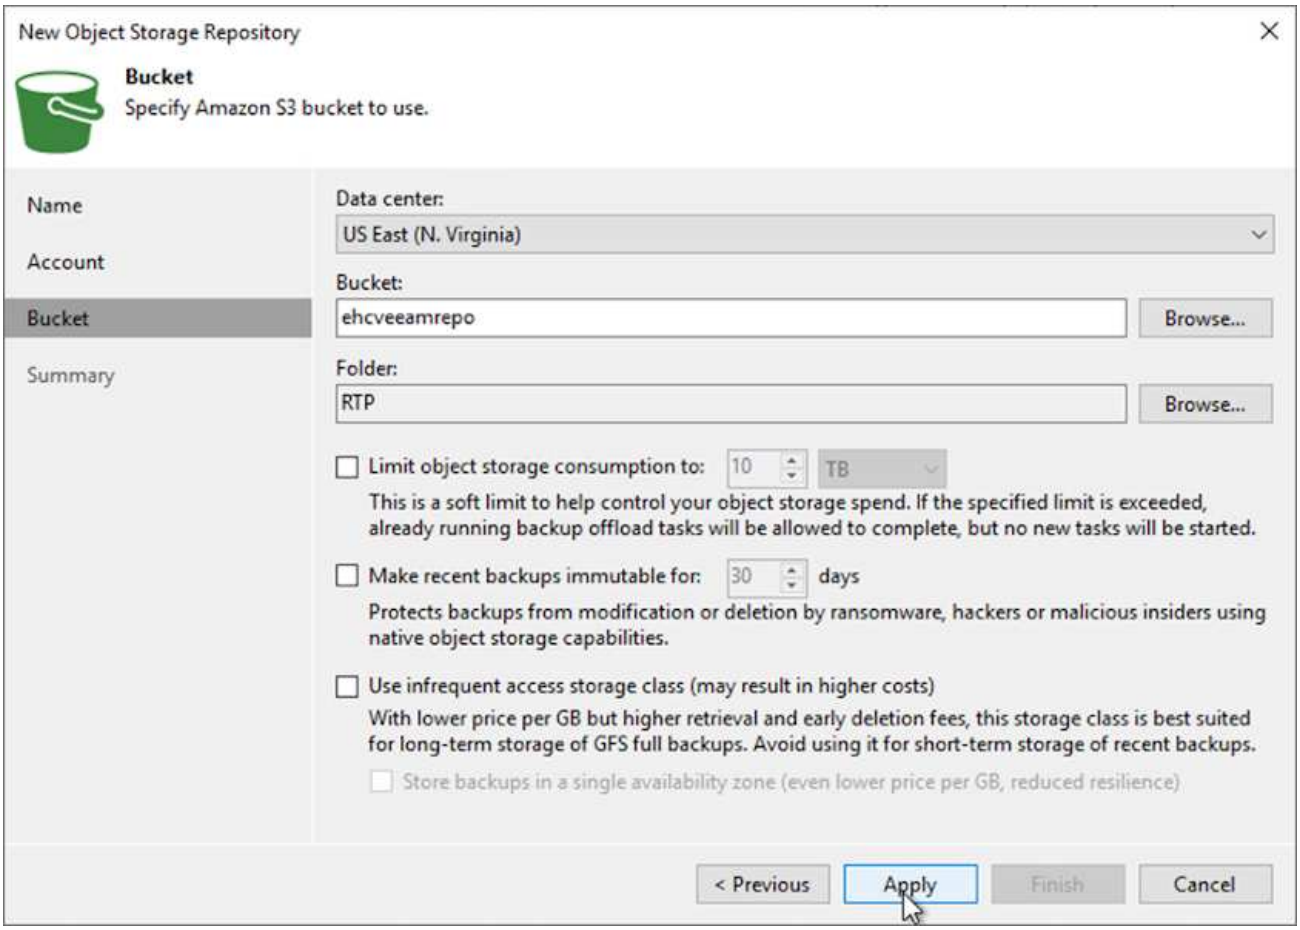
3. From the list of Amazon Cloud Storage Services, select Amazon S3.



4. Select your pre-entered credentials from the drop-down list or add a new credential for accessing the cloud storage resource. Click Next to continue.



5. On the Bucket page, enter the data center, bucket, folder, and any desired options. Click Apply.

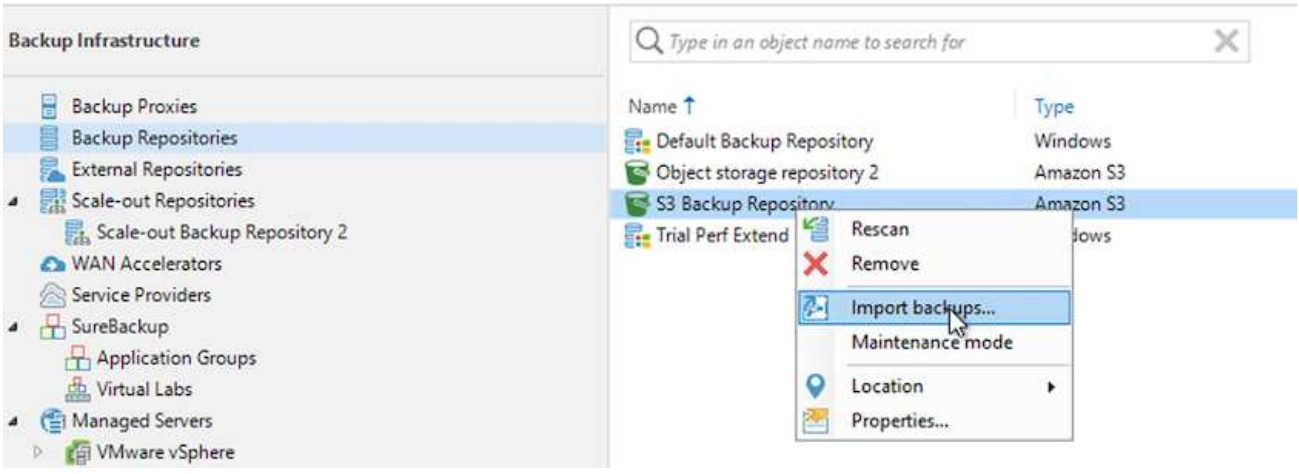


6. Finally, select Finish to complete the process and add the repository.

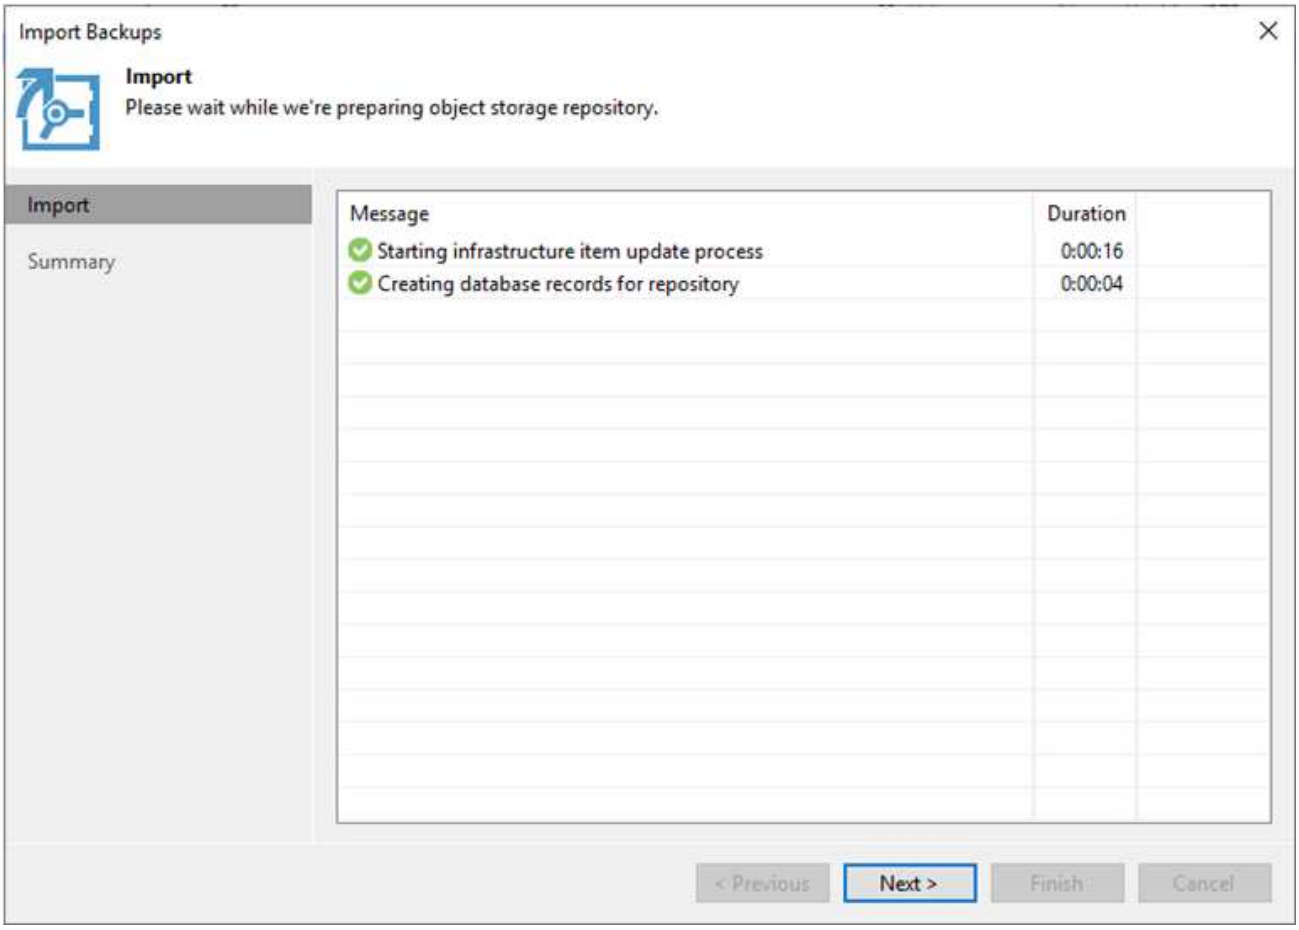
▼ **Import backups from S3 object storage**

To import the backups from the S3 repository that was added in the previous section, complete the following steps.

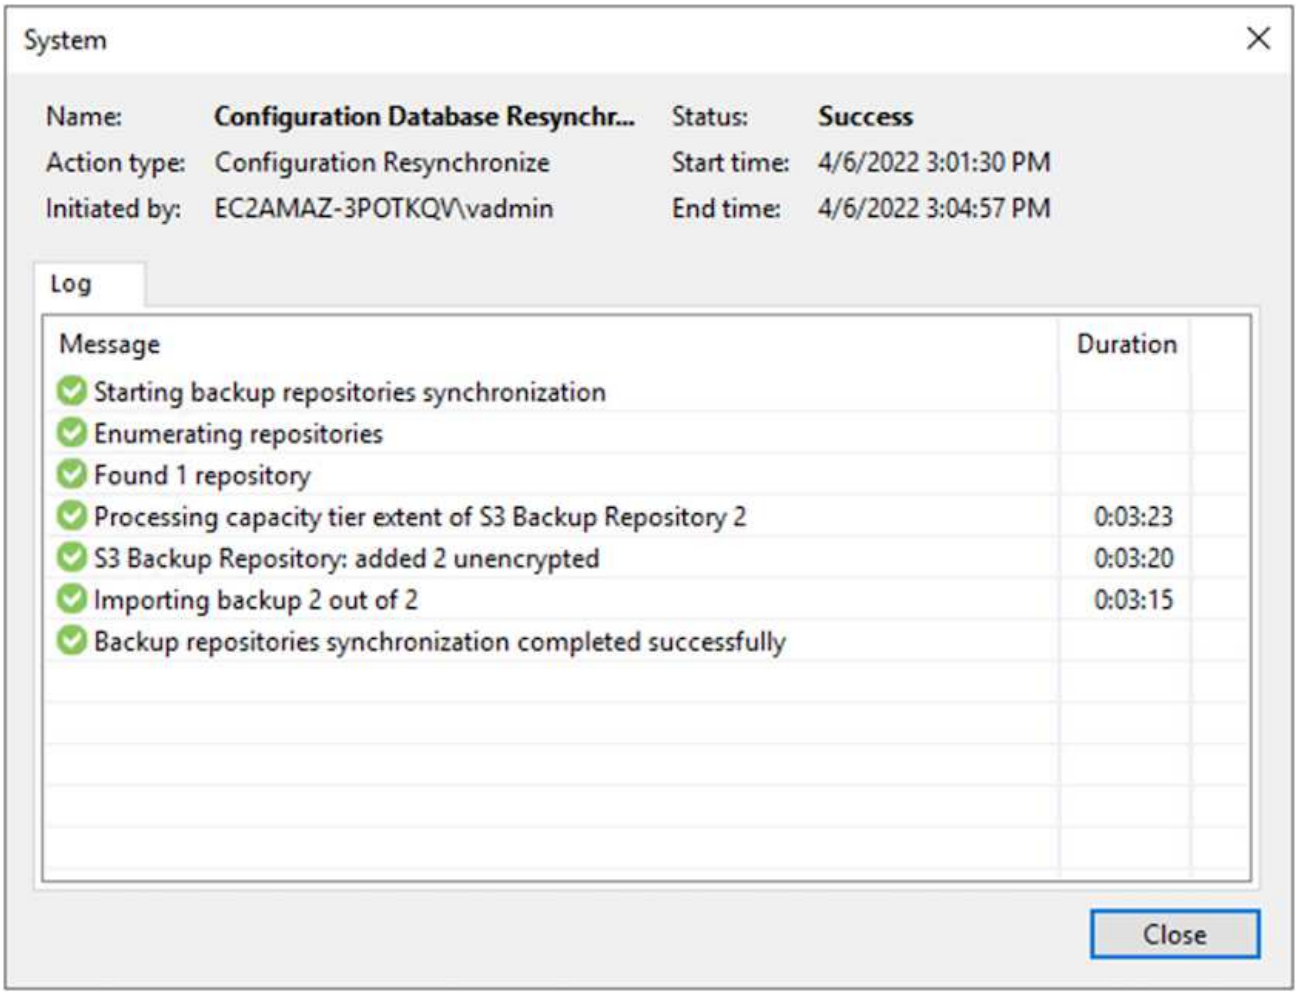
1. From the S3 backup repository, select Import Backups to launch the Import Backups wizard.



2. After the database records for the import have been created, select Next and then Finish at the summary screen to start the import process.



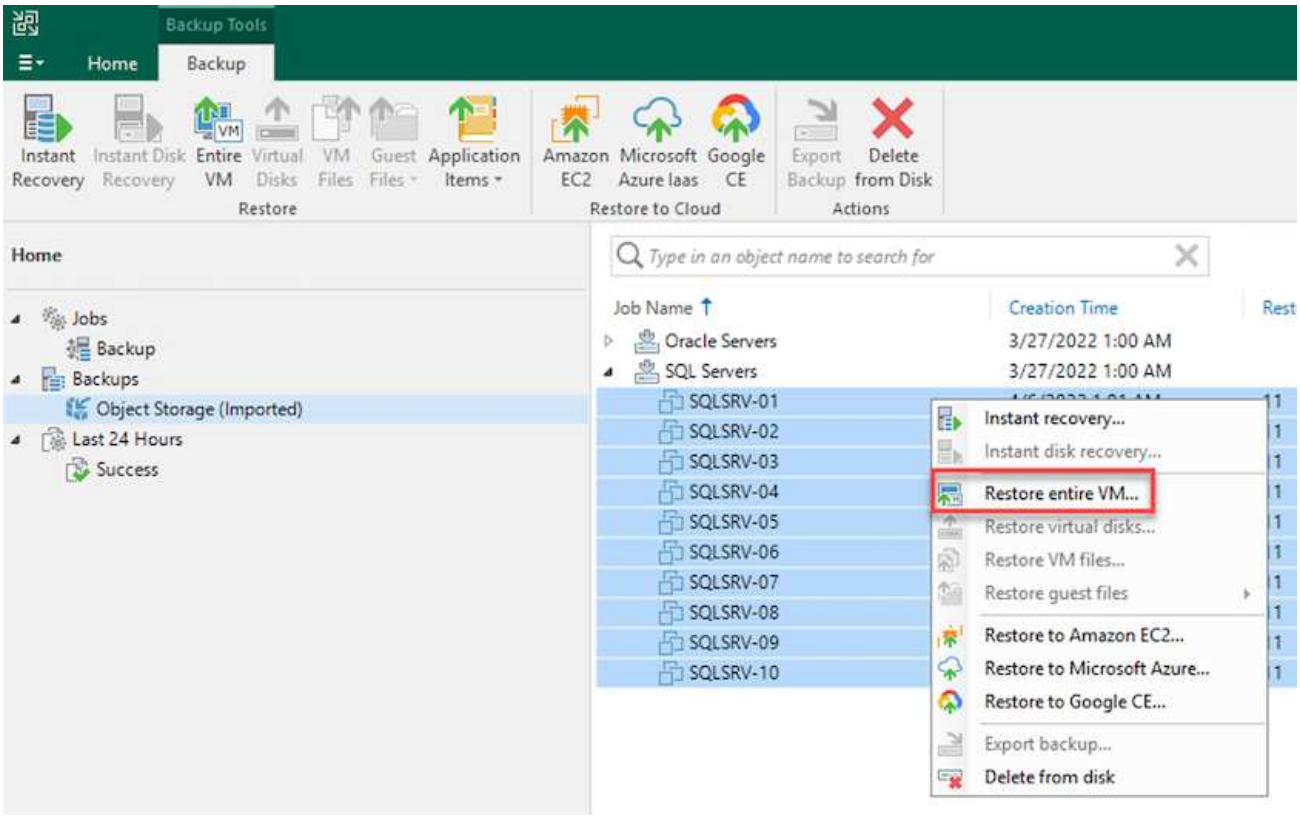
3. After the import is complete, you can restore VMs into the VMware Cloud cluster.



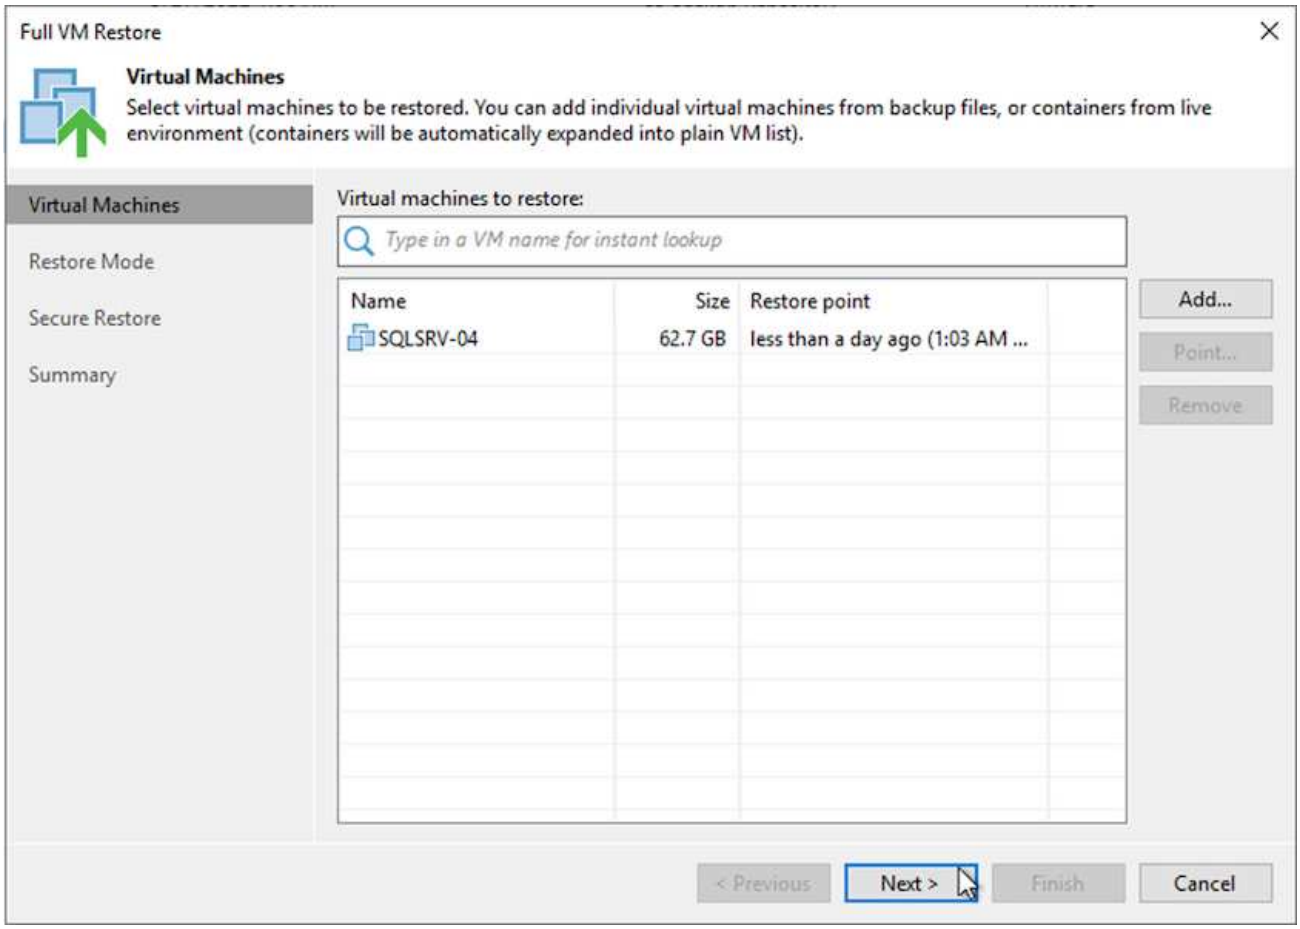
▼ **Restore application VMs with Veeam full restore to VMware Cloud**

To restore SQL and Oracle virtual machines to the VMware Cloud on AWS workload domain/cluster, complete the following steps.

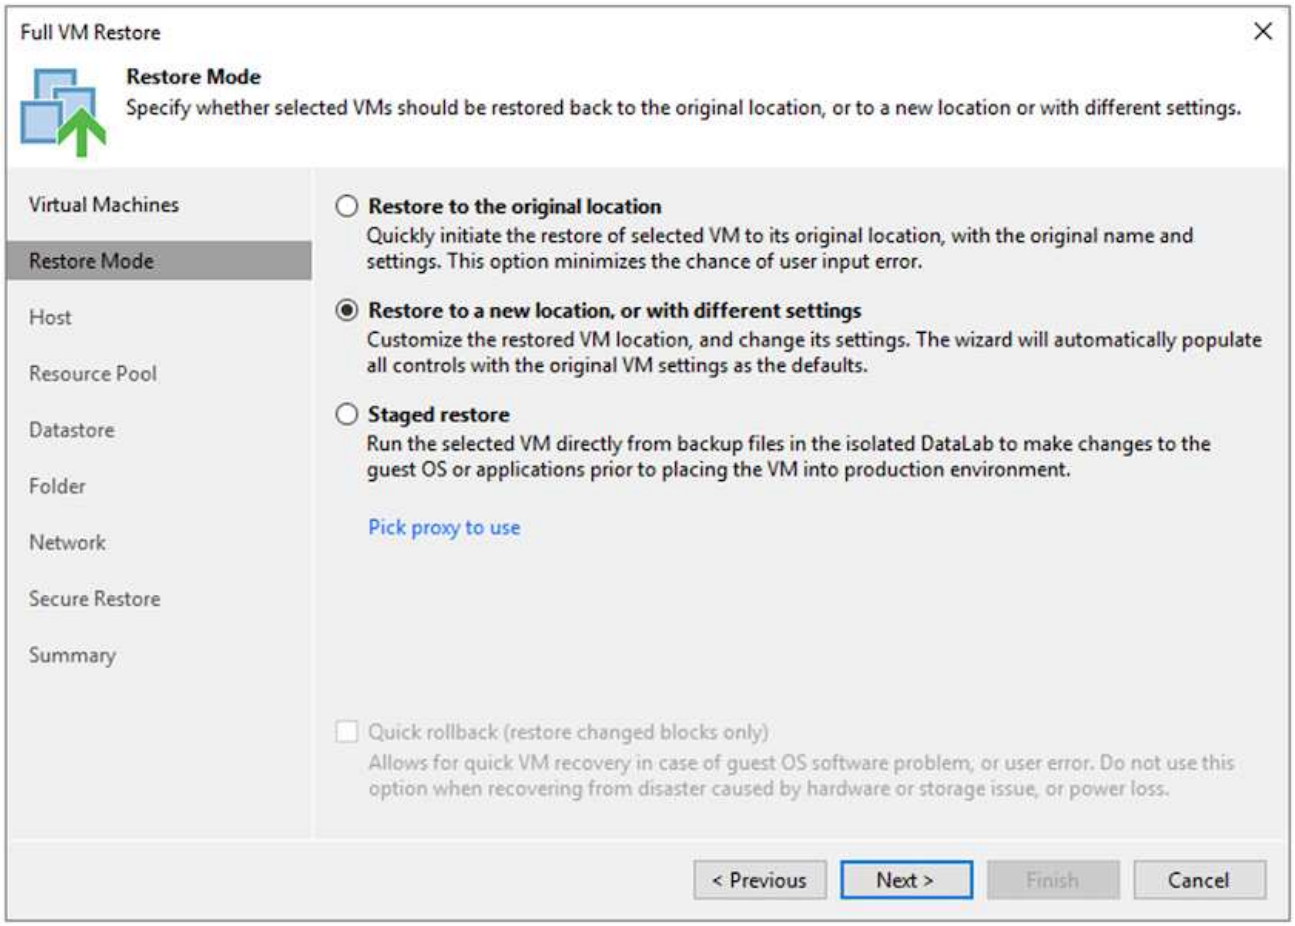
1. From the Veeam Home page, select the object storage containing the imported backups, select the VMs to restore, and then right click and select Restore Entire VM.



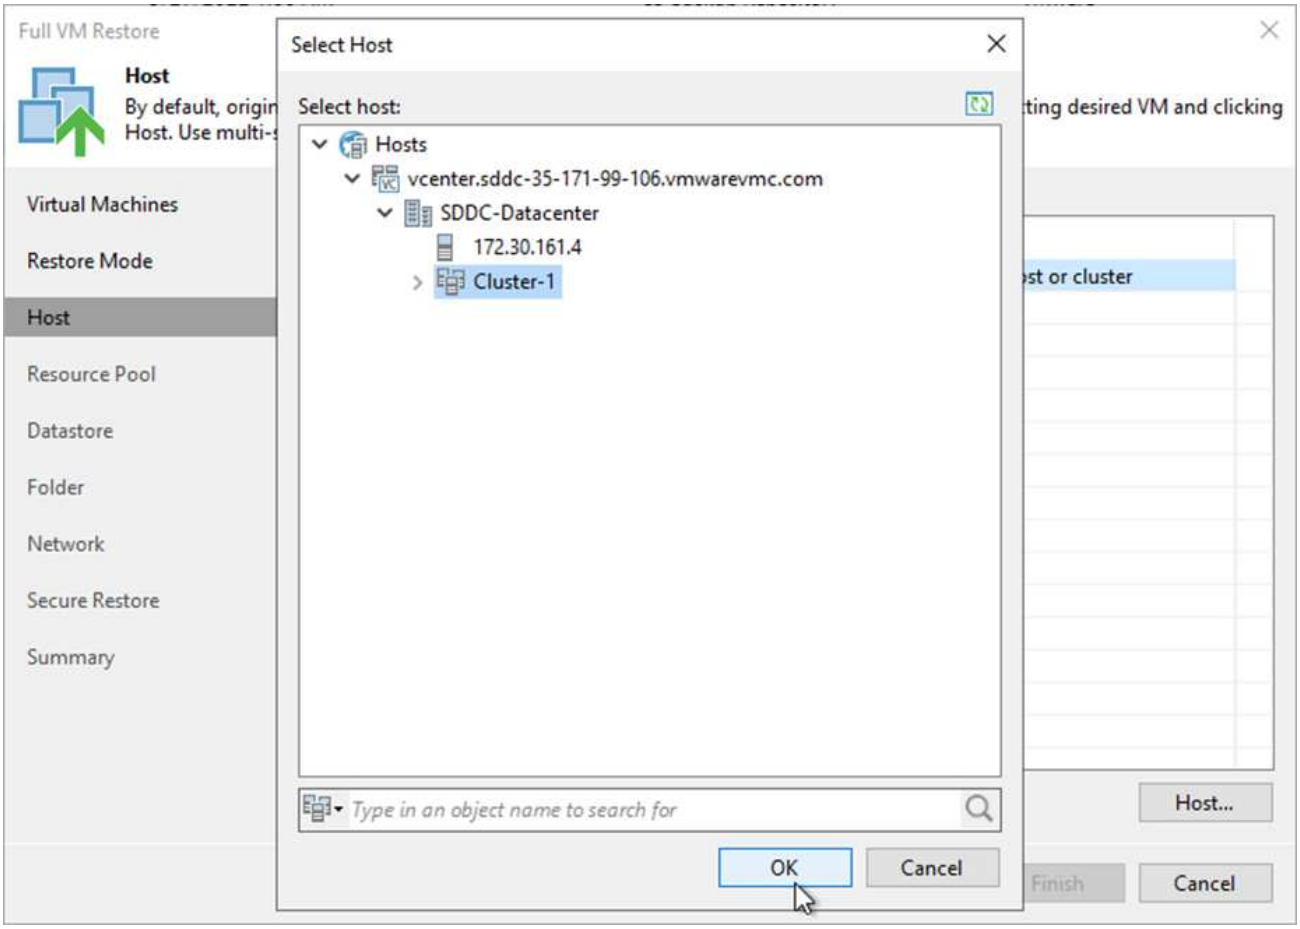
2. On the first page of the Full VM Restore wizard, modify the VMs to backup if desired and select Next.



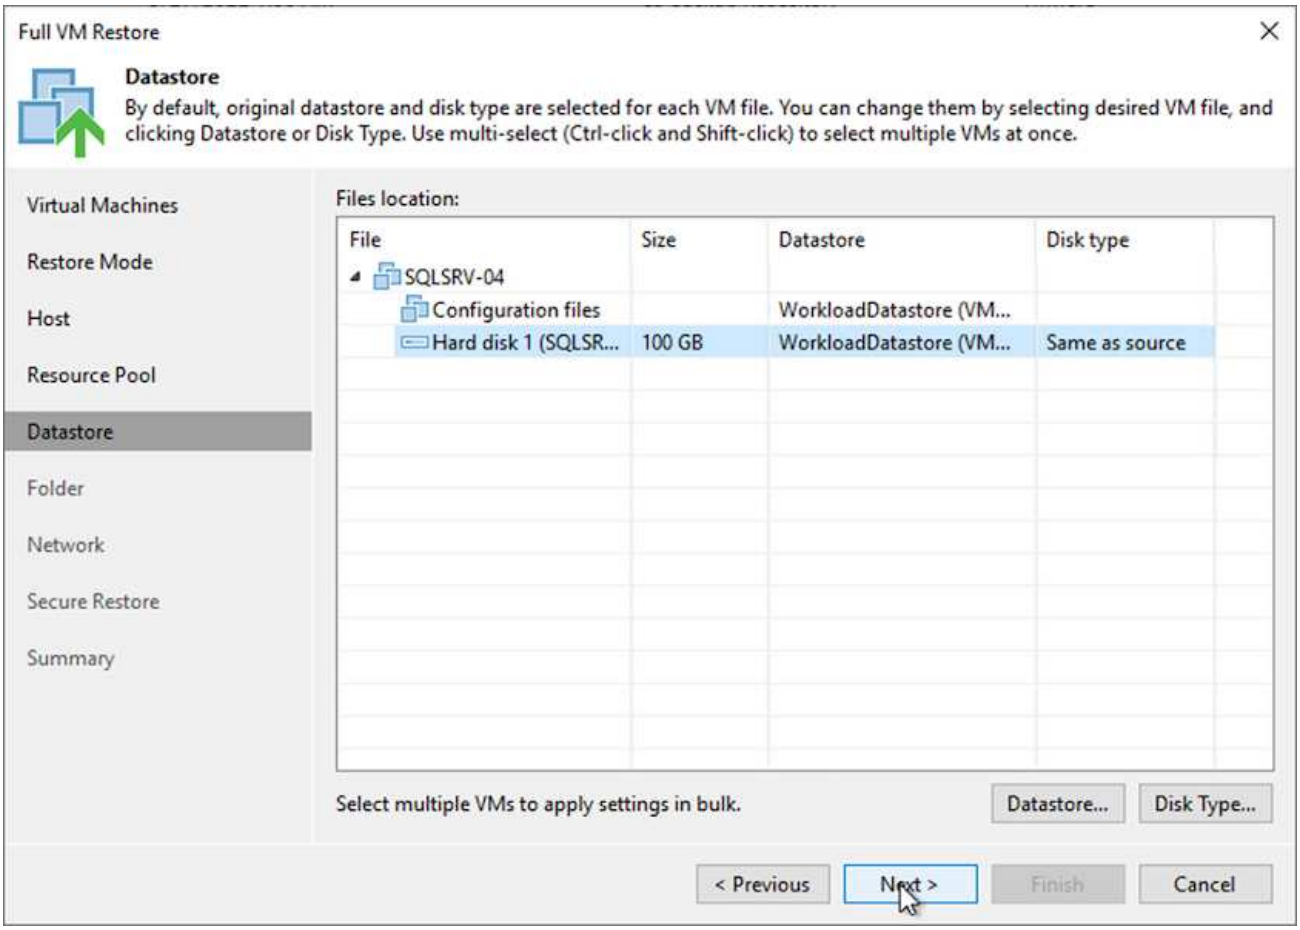
3. On the Restore Mode page, select Restore to a New Location, or with Different Settings.



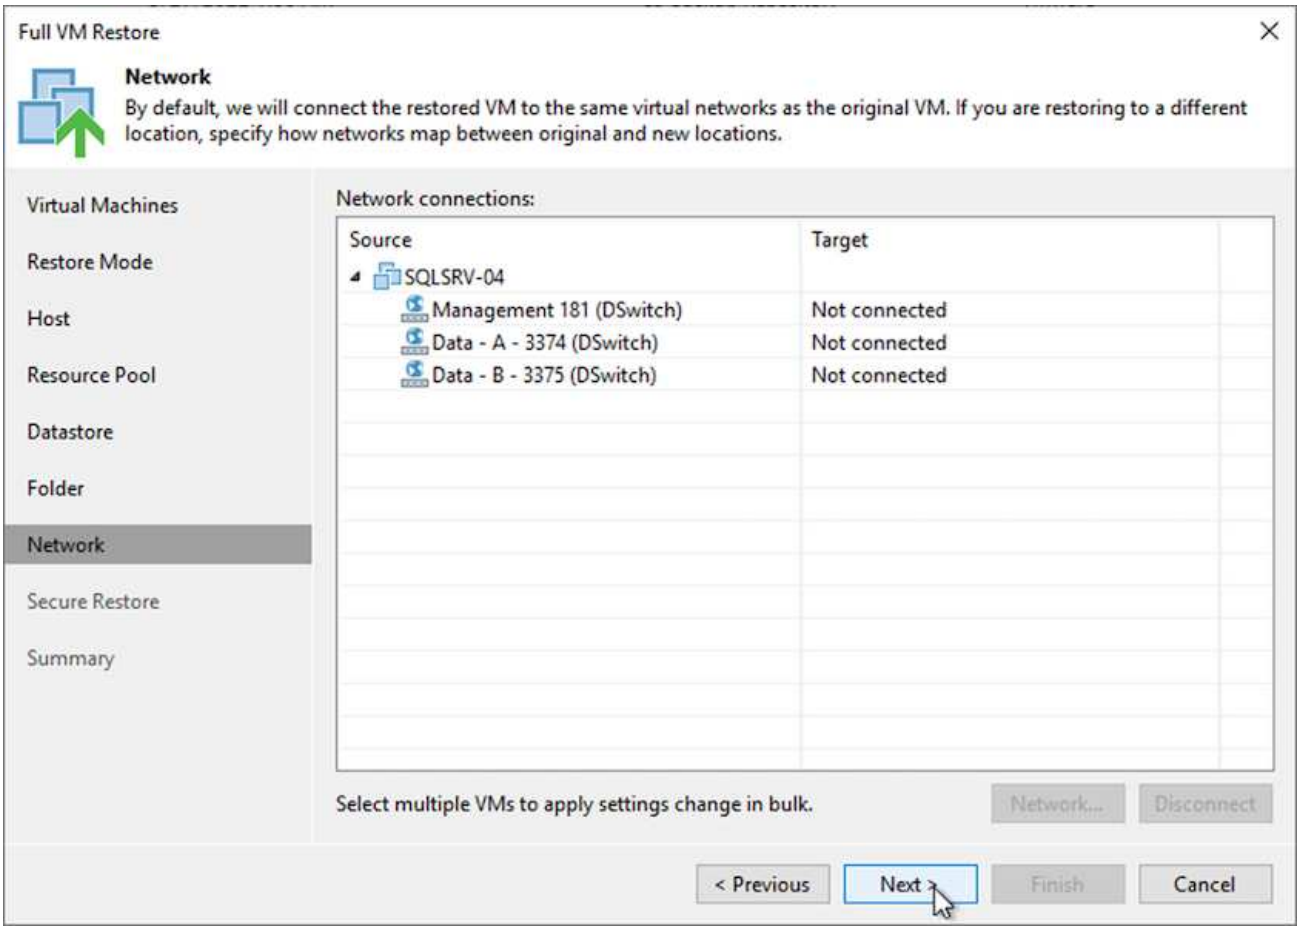
4. On the host page, select the Target ESXi host or cluster to restore the VM to.

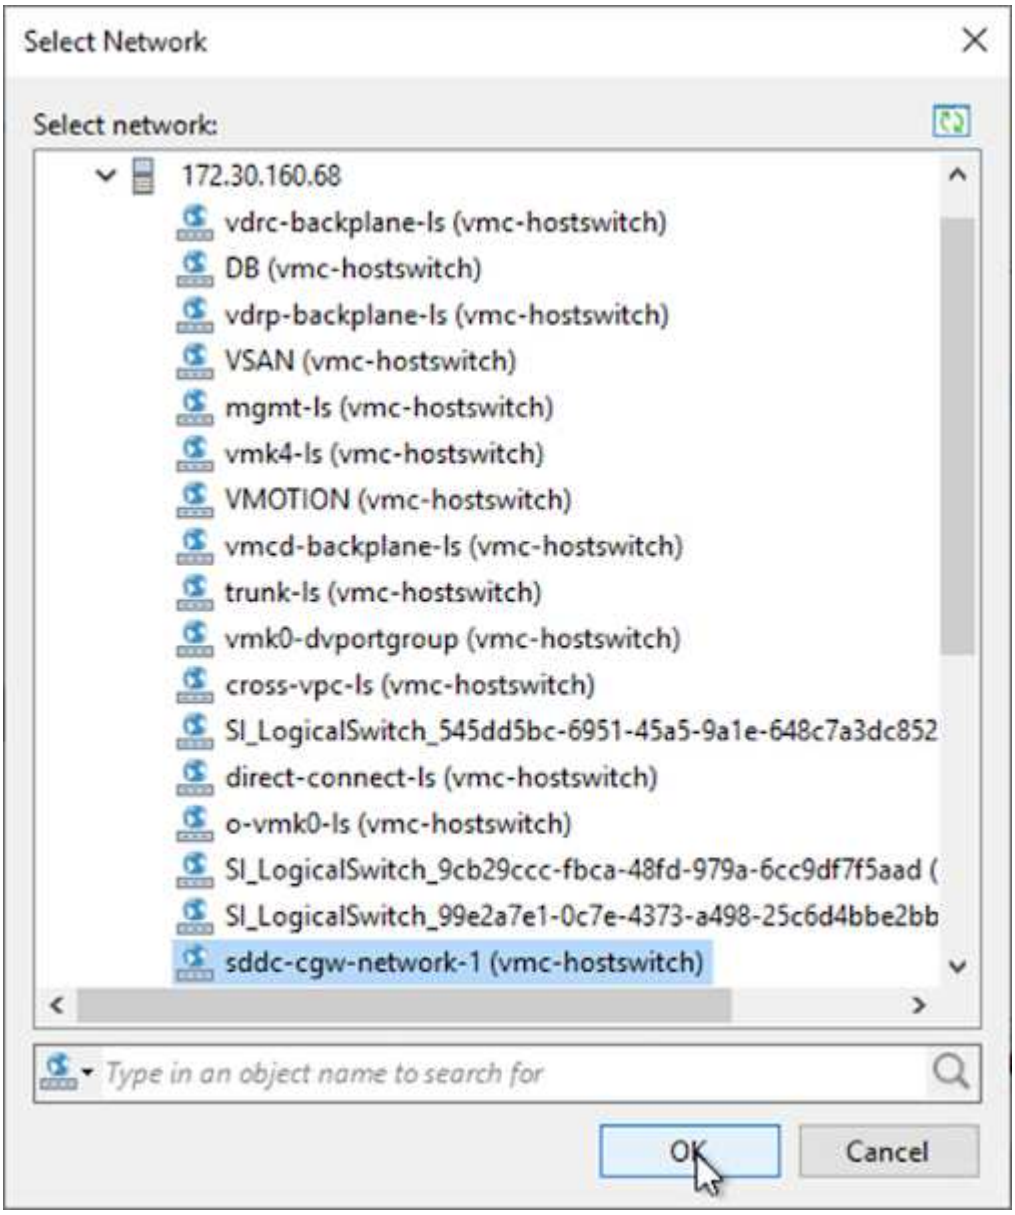


5. On the Datastores page, select the target datastore location for both the configuration files and hard disk.



6. On the Network page, map the original networks on the VM to the networks in the new target location.


Full VM Restore
Network
By default, we will connect the restored VM to the same virtual networks as the original VM. If you are restoring to a different location, specify how networks map between original and new locations.
Virtual Machines
Restore Mode
Host
Resource Pool
Datastore
Folder
Network
Secure Restore
Summary
Network connections:
Source
Target
SQLSRV-04
Management 181 (DSwitch)
Not connected
Data - A - 3374 (DSwitch)
Not connected
Data - B - 3375 (DSwitch)
Not connected
Select multiple VMs to apply settings change in bulk.
Network...
Disconnect
< Previous
Next >
Finish
Cancel



7. Select whether to scan the restored VM for malware, review the summary page, and click Finish to start the restore.

Restore SQL Server application data

The following process provides instructions on how to recover a SQL Server in VMware Cloud Services in AWS in the event of a disaster that renders the on-premises site inoperable.

The following prerequisites are assumed to be complete in order to continue with the recovery steps:

1. The Windows Server VM has been restored to the VMware Cloud SDDC using Veeam Full Restore.
2. A secondary SnapCenter server has been established and SnapCenter database restore and configuration has been completed using the steps outlined in the section "SnapCenter backup and restore process summary."

▼ VM: Post restore configuration for SQL Server VM

After the restore of the VM is complete, you must configure networking and other items in preparation for rediscovering the host VM within SnapCenter.

1. Assign new IP addresses for Management and iSCSI or NFS.
2. Join the host to the Windows domain.
3. Add the hostnames to DNS or to the hosts file on the SnapCenter server.

If the SnapCenter plug-in was deployed using domain credentials different than the current domain, you must change the Log On account for the Plug-in for Windows Service on the SQL Server VM. After changing the Log On account, restart the SnapCenter SMCore, Plug-in for Windows, and Plug-in for SQL Server services.

To automatically rediscover the restored VMs in SnapCenter, the FQDN must be identical to the VM that was originally added to the SnapCenter on premises.

▼ **Configure FSx storage for SQL Server restore**

To accomplish the disaster recovery restore process for a SQL Server VM, you must break the existing SnapMirror relationship from the FSx cluster and grant access to the volume. To do so, complete the following steps.

1. To break the existing SnapMirror relationship for the SQL Server database and log volumes, run the following command from the FSx CLI:

```
FSx-Dest::> snapmirror break -destination-path DestSVM:DestVolName
```

2. Grant access to the LUN by creating an initiator group containing the iSCSI IQN of the SQL Server Windows VM:

```
FSx-Dest::> igroup create -vserver DestSVM -igroup igroupName
-protocol iSCSI -ostype windows -initiator IQN
```

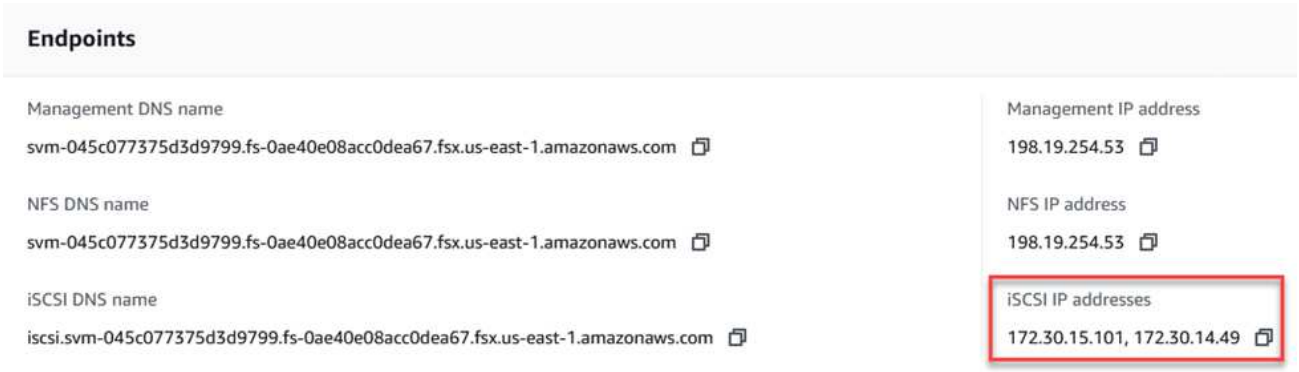
3. Finally, map the LUNs to the initiator group that you just created:

```
FSx-Dest::> lun mapping create -vserver DestSVM -path LUNPath igroup
igroupName
```

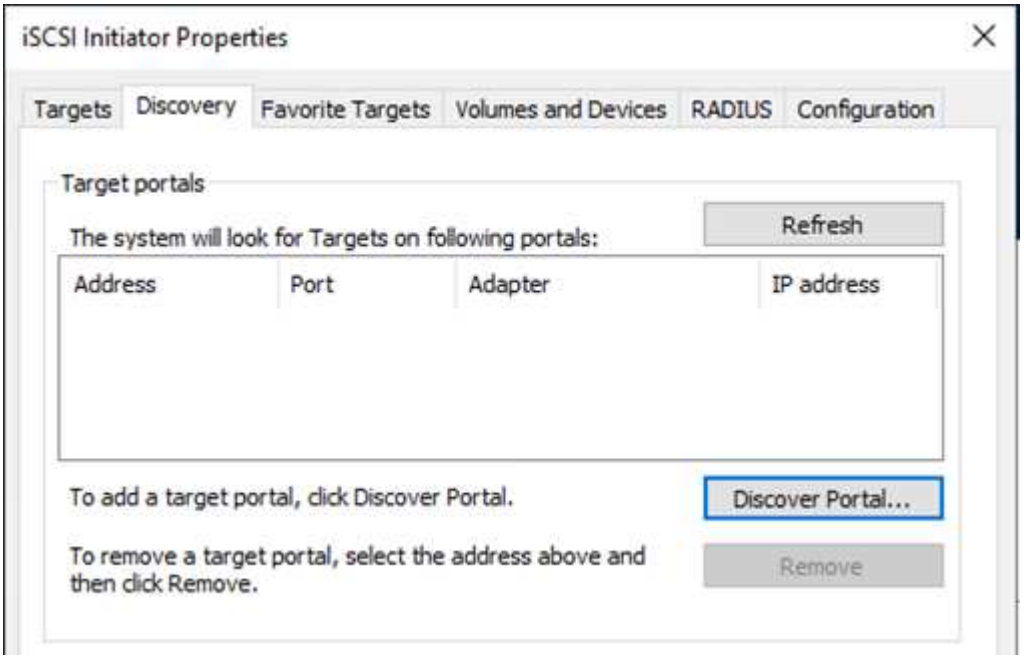
4. To find the path name, run the `lun show` command.

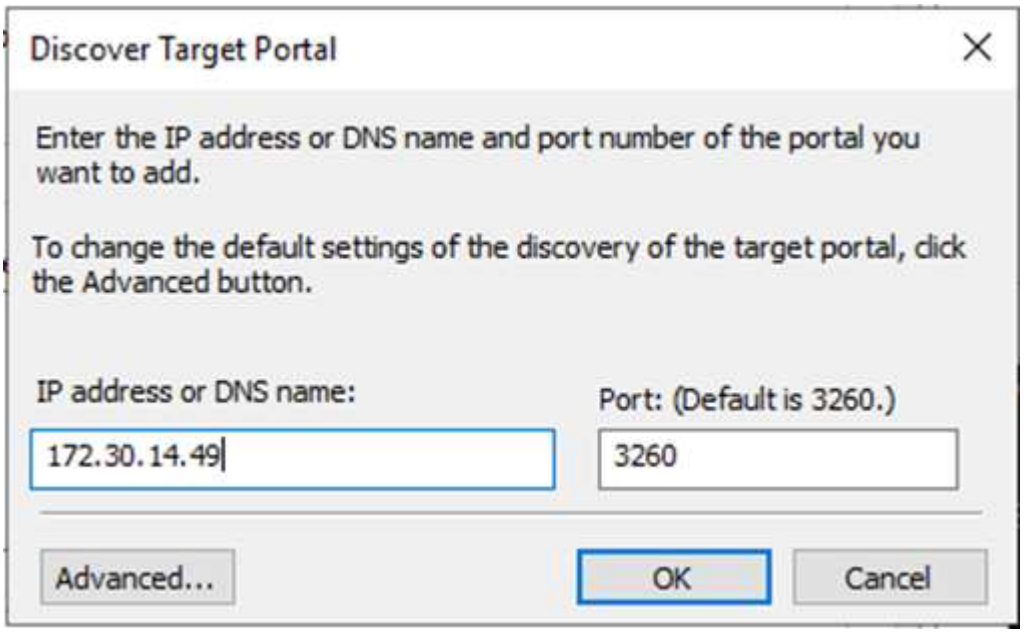
▼ **Set up the Windows VM for iSCSI access and discover the file systems**

1. From the SQL Server VM, set up your iSCSI network adapter to communicate on the VMware Port Group that has been established with connectivity to the iSCSI target interfaces on your FSx instance.
2. Open the iSCSI Initiator Properties utility and clear out the old connectivity settings on the Discovery, Favorite Targets, and Targets tabs.
3. Locate the IP address(es) for accessing the iSCSI logical interface on the FSx instance/cluster. This can be found in the AWS console under Amazon FSx > ONTAP > Storage Virtual Machines.



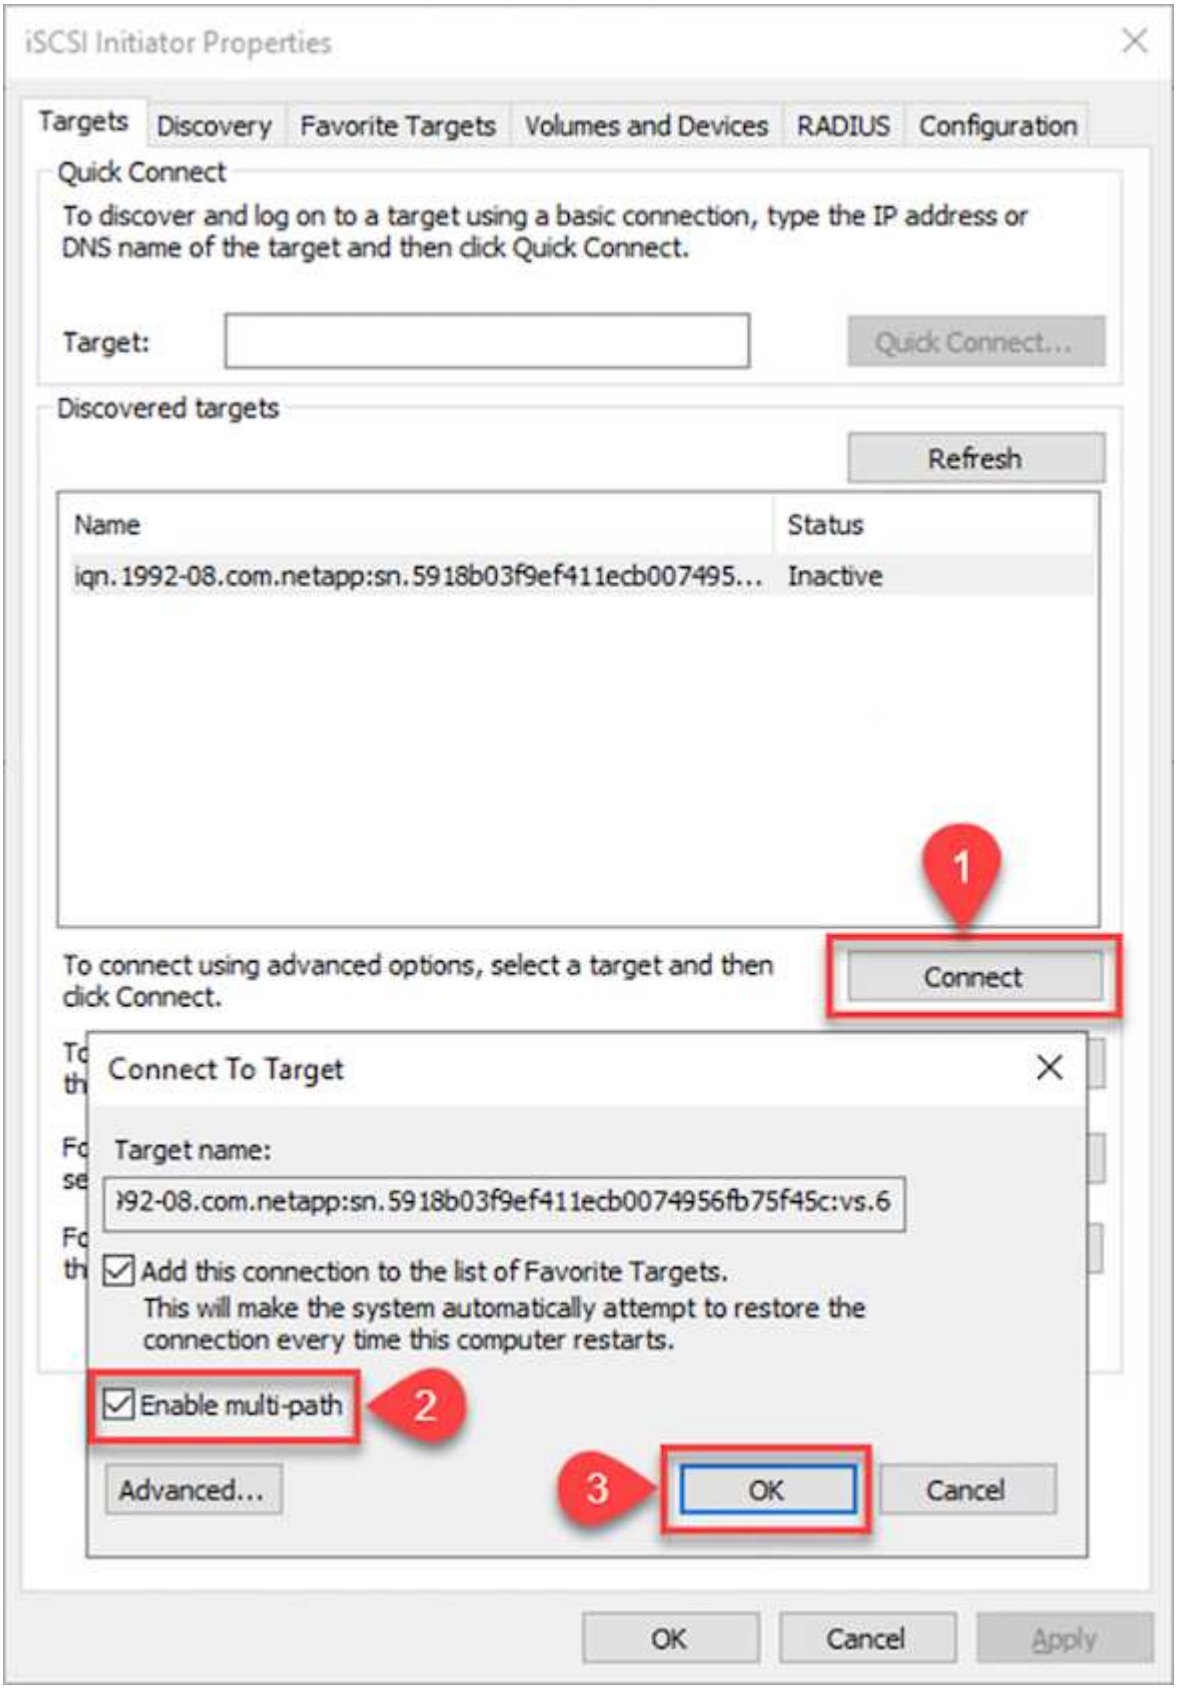
4. From the Discovery tab, click Discover Portal and enter the IP addresses for your FSx iSCSI targets.



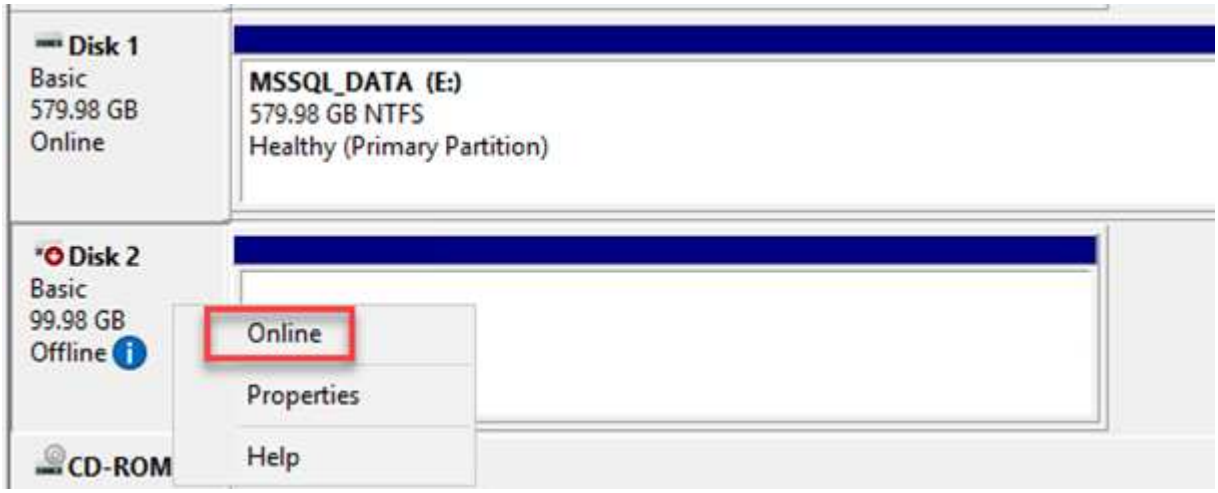


5. On the Target tab, click Connect, select Enable Multi-Path if appropriate for your configuration and then

click OK to connect to the target.

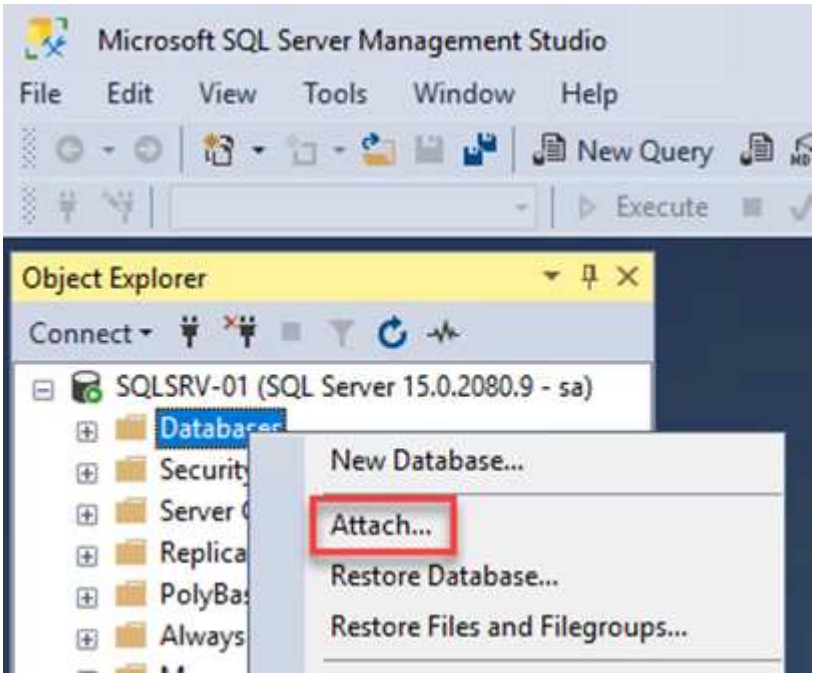


6. Open the Computer Management utility and bring the disks online. Verify that they retain the same drive letters that they previously held.

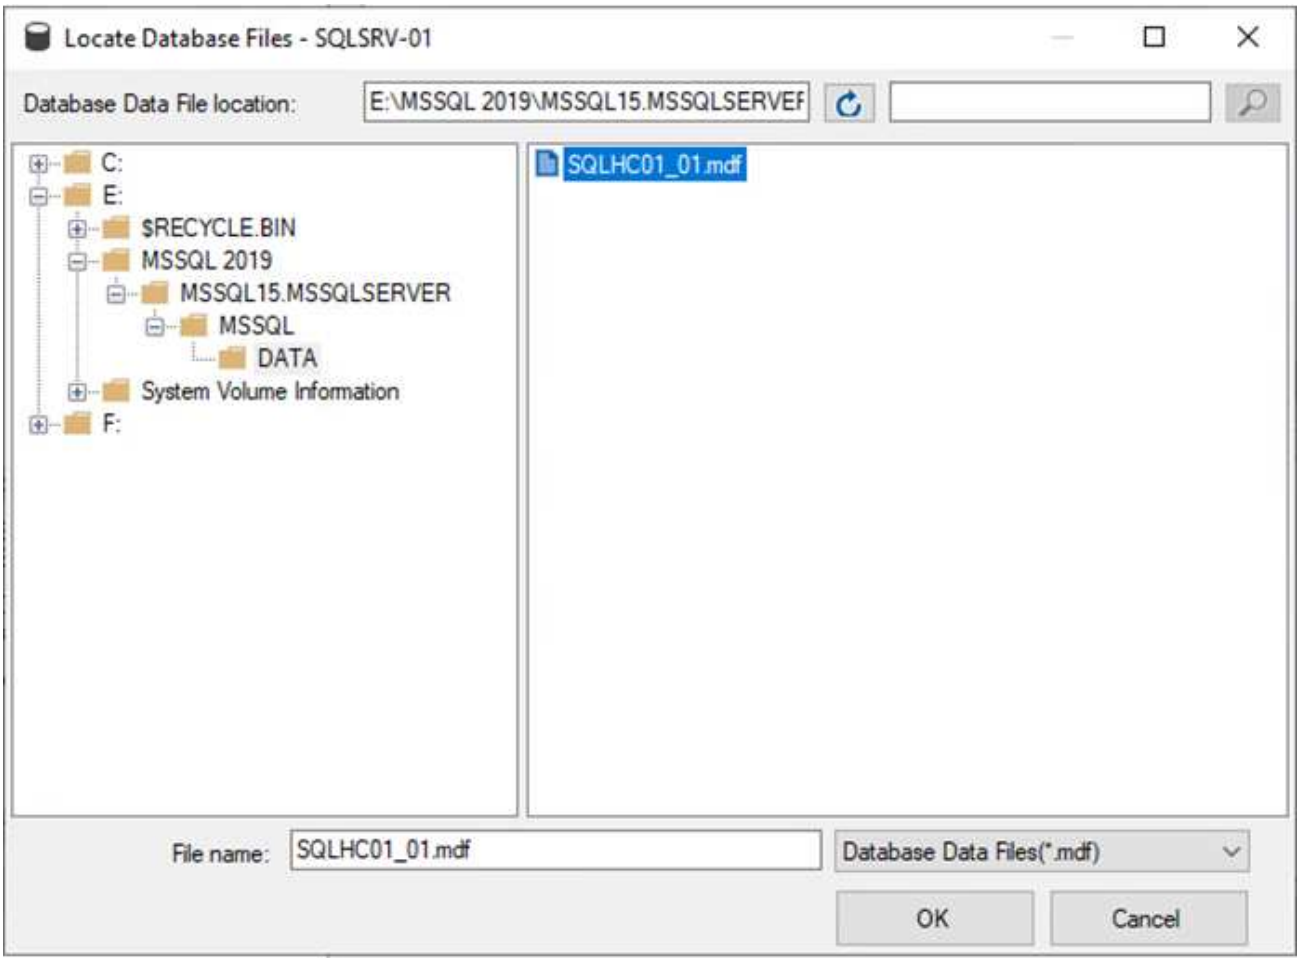


▼ **Attach the SQL Server databases**

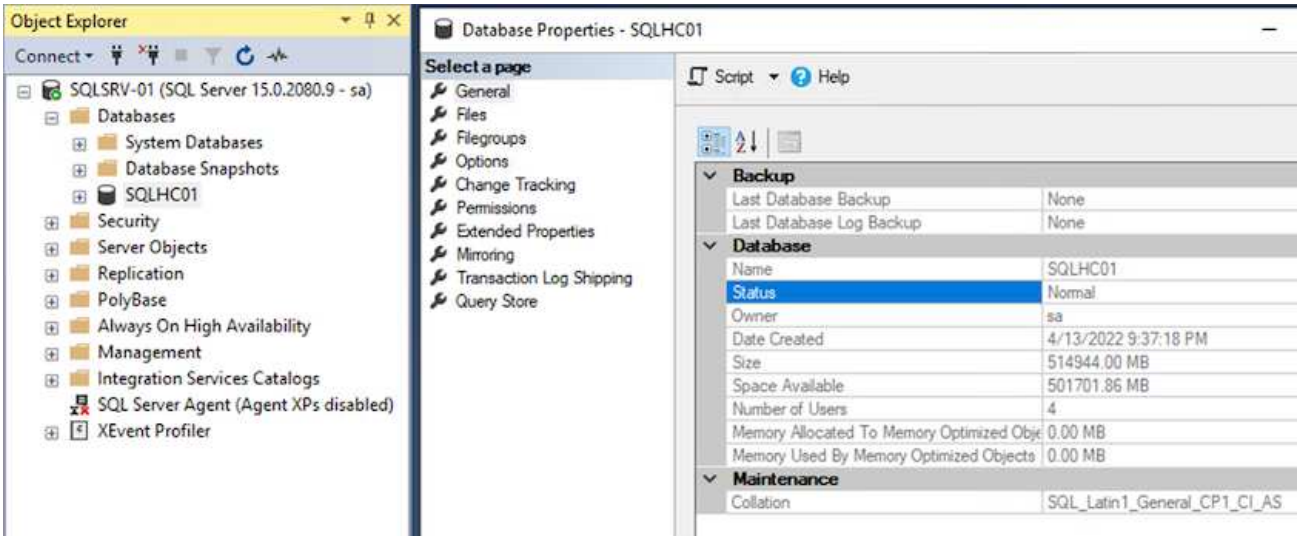
1. From the SQL Server VM, open Microsoft SQL Server Management Studio and select Attach to start the process of connecting to the database.



2. Click Add and navigate to the folder containing the SQL Server primary database file, select it, and click OK.



3. If the transaction logs are on a separate drive, choose the folder that contains the transaction log.
4. When finished, click OK to attach the database.



▼ Confirm SnapCenter communication with SQL Server Plug-in

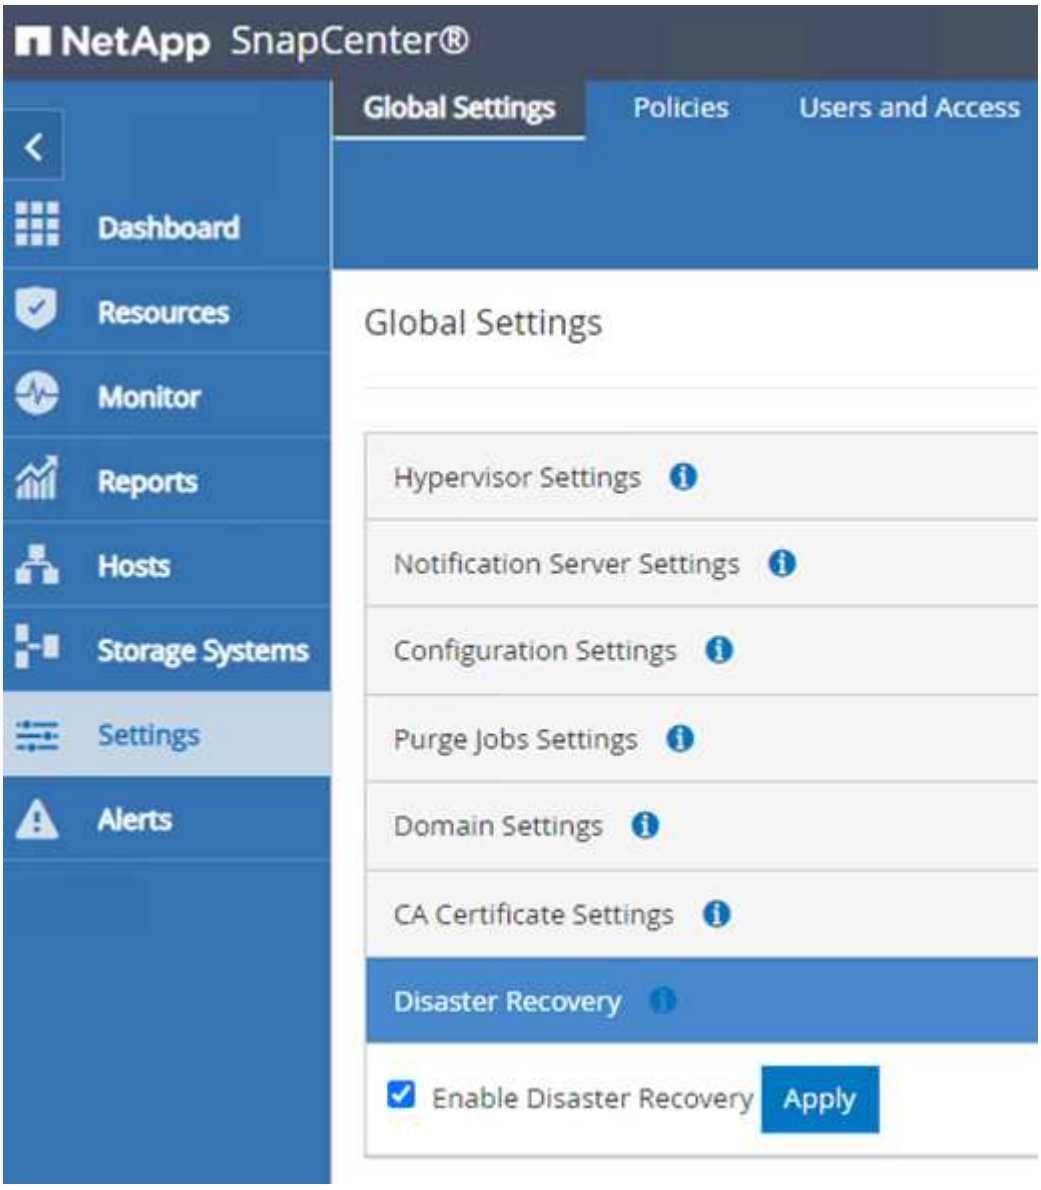
With the SnapCenter database restored to its previous state, it automatically rediscovers the SQL Server hosts. For this to work correctly, keep in mind the following prerequisites:

- SnapCenter must be placed in Disaster Recover mode. This can be accomplished through the Swagger

API or in Global Settings under Disaster Recovery.

- The FQDN of the SQL Server must be identical to the instance that was running in the on-premises datacenter.
- The original SnapMirror relationship must be broken.
- The LUNs containing the database must be mounted to the SQL Server instance and the database attached.

To confirm that SnapCenter is in Disaster Recovery mode, navigate to Settings from within the SnapCenter web client. Go to the Global Settings tab and then click Disaster Recovery. Make sure that the Enable Disaster Recovery checkbox is enabled.



Restore Oracle application data

The following process provides instructions on how to recover Oracle application data in VMware Cloud Services in AWS in the event of a disaster that renders the on-premises site inoperable.

Complete the following prerequisites to continue with the recovery steps:

1. The Oracle Linux server VM has been restored to the VMware Cloud SDDC using Veeam Full Restore.
2. A secondary SnapCenter server has been established and the SnapCenter database and configuration files have been restored using the steps outlined in this section "SnapCenter backup and restore process summary."

▼ **Configure FSx for Oracle restore – Break the SnapMirror relationship**

To make the secondary storage volumes hosted on the FSx ONTAP instance accessible to the Oracle servers, you must first break the existing SnapMirror relationship.

1. After logging into the FSx CLI, run the following command to view the volumes filtered by the correct name.

```
FSx-Dest::> volume show -volume VolumeName*
```

```
FsxId0ae40e08acc0dea67::> volume show -volume oraclesrv_03*
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
ora_svm_dest
          oraclesrv_03_u01_dest
                       aggr1        online     DP        100GB    93.12GB    6%
ora_svm_dest
          oraclesrv_03_u02_dest
                       aggr1        online     DP        200GB    34.98GB   82%
ora_svm_dest
          oraclesrv_03_u03_dest
                       aggr1        online     DP        150GB    33.37GB   77%
3 entries were displayed.

FsxId0ae40e08acc0dea67::>
```

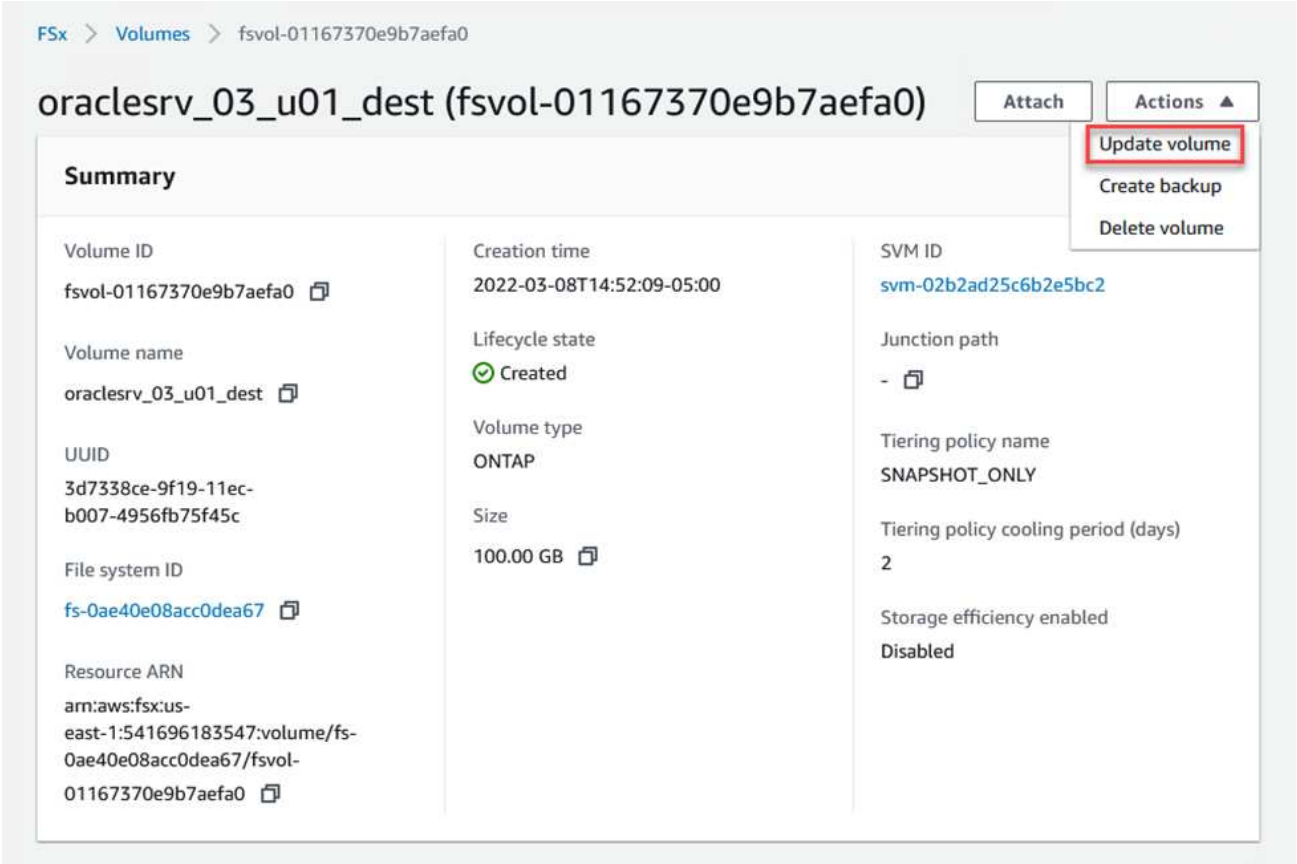
2. Run the following command to break the existing SnapMirror relationships.

```
FSx-Dest::> snapmirror break -destination-path DestSVM:DestVolName
```

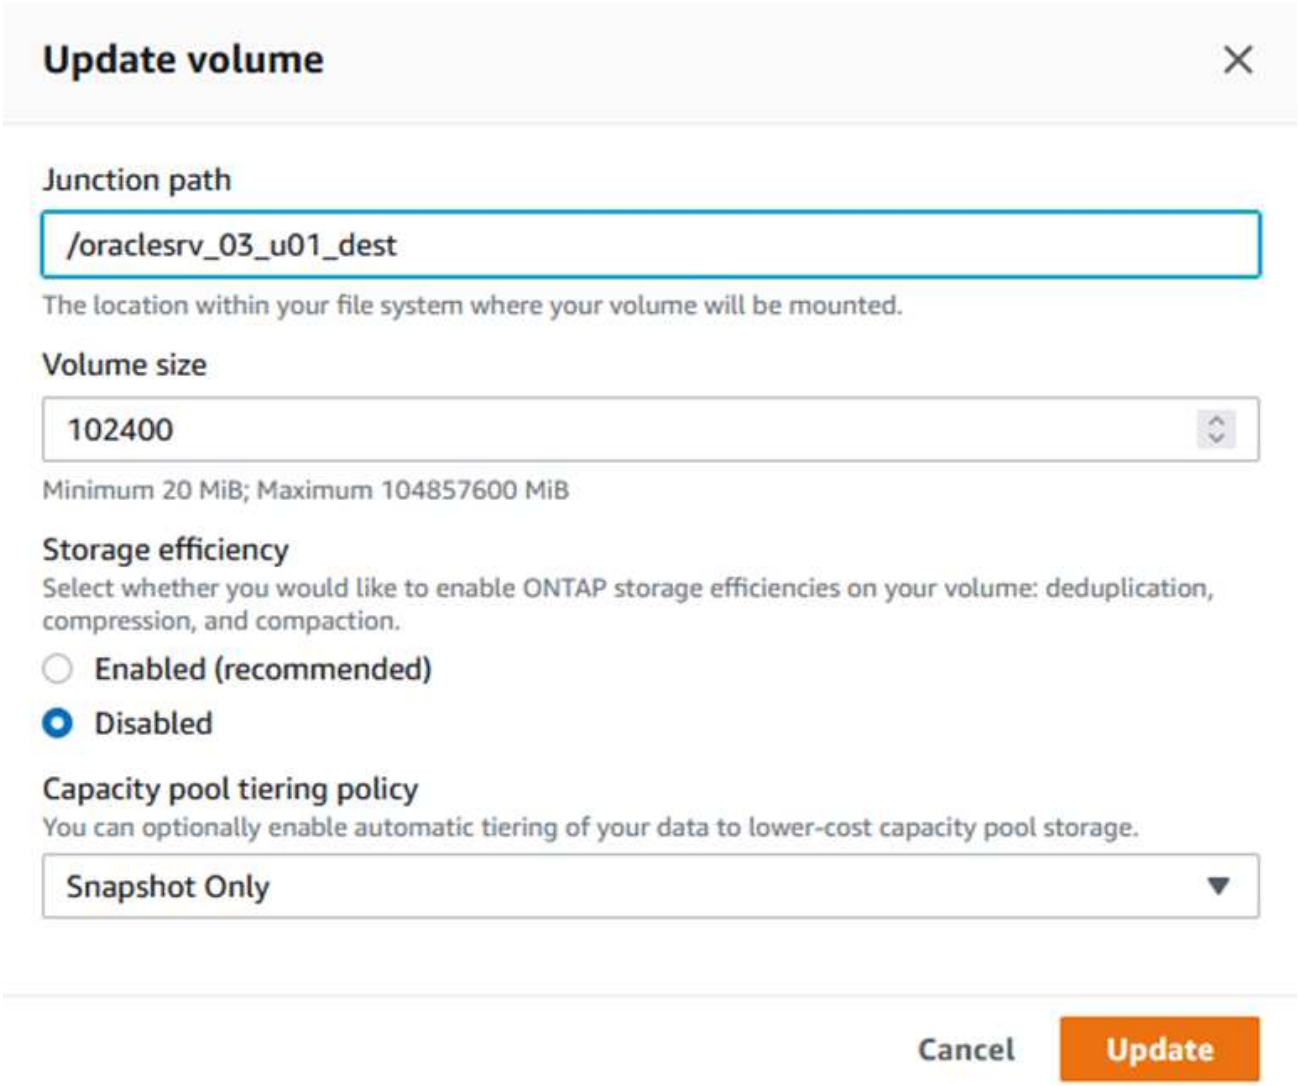
```
FsxId0ae40e08acc0dea67::> snapmirror break -destination-path ora_svm_dest:oraclesrv_03_u02_dest
Operation succeeded: snapmirror break for destination "ora_svm_dest:oraclesrv_03_u02_dest".

FsxId0ae40e08acc0dea67::> snapmirror break -destination-path ora_svm_dest:oraclesrv_03_u03_dest
Operation succeeded: snapmirror break for destination "ora_svm_dest:oraclesrv_03_u03_dest".
```

3. Update the junction-path in the Amazon FSx web client:



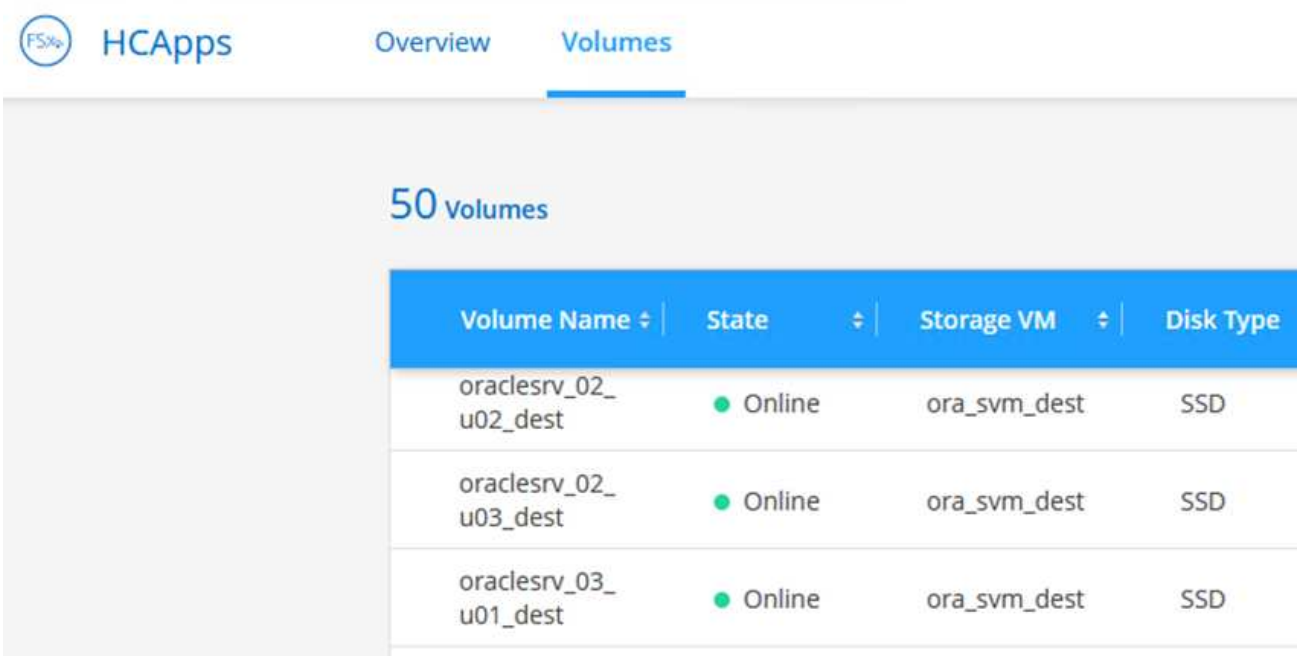
4. Add the junction path name and click Update. Specify this junction path when mounting the NFS volume from the Oracle server.



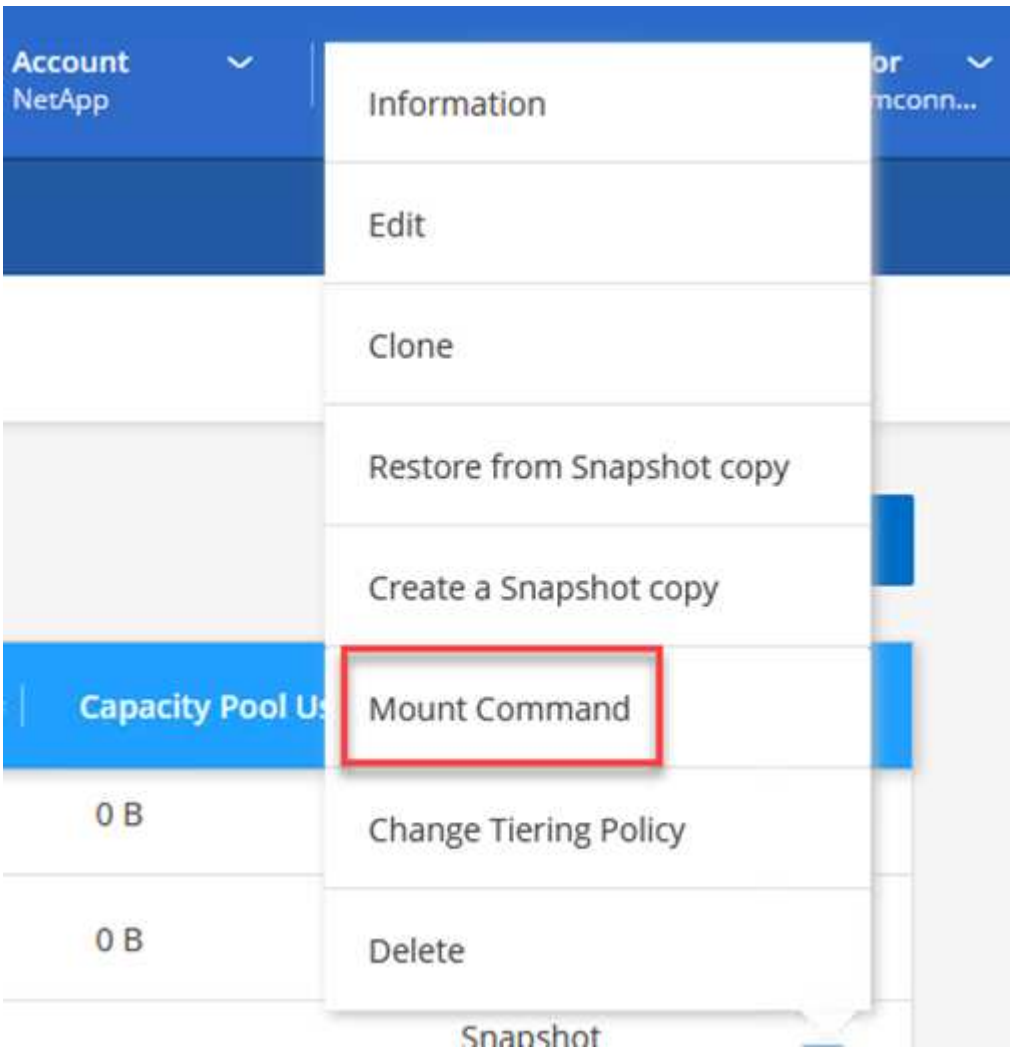
▼ **Mount NFS volumes on Oracle Server**

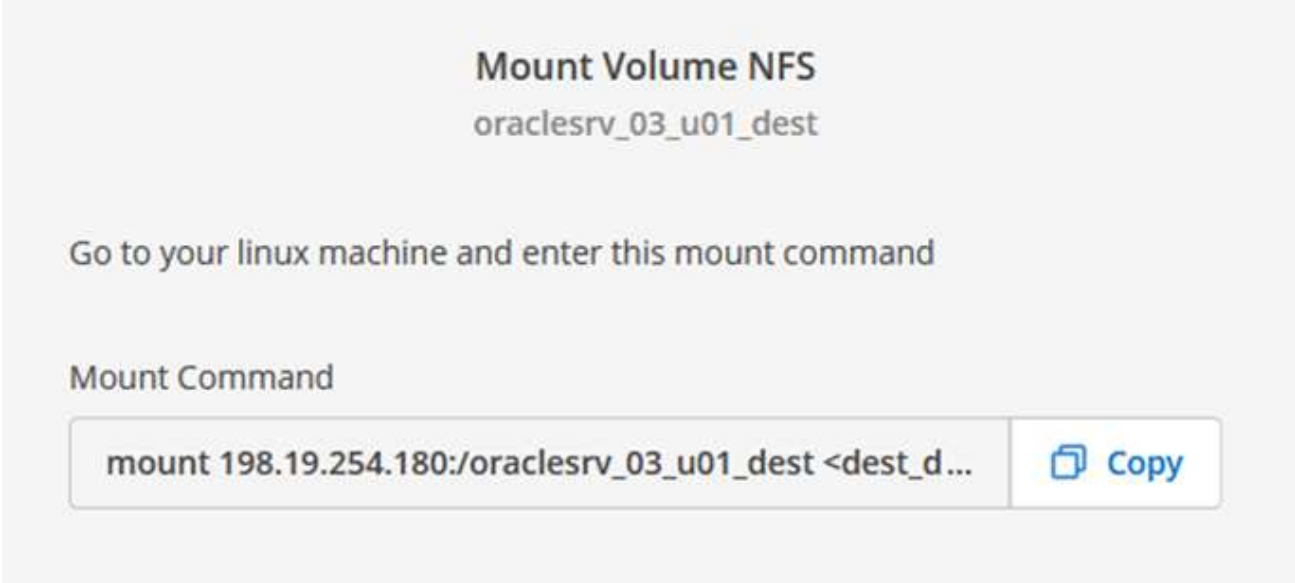
In Cloud Manager, you can obtain the mount command with the correct NFS LIF IP address for mounting the NFS volumes that contain the Oracle database files and logs.

1. In Cloud Manager, access the list of volumes for your FSx cluster.



2. From the action menu, select Mount Command to view and copy the mount command to be used on our Oracle Linux server.





3. Mount the NFS file system to the Oracle Linux Server. The directories for mounting the NFS share already exist on the Oracle Linux host.
4. From the Oracle Linux server, use the mount command to mount the NFS volumes.

```
FSx-Dest::> mount -t oracle_server_ip:/junction-path
```

Repeat this step for each volume associated with the Oracle databases.

> To make the NFS mount persistent upon rebooting, edit the `/etc/fstab` file to include the mount commands.

5. Reboot the Oracle server. The Oracle databases should start up normally and be available for use.

Failback

Upon successful completion of the failover process outlined in this solution, SnapCenter and Veeam resume their backup functions running in AWS, and FSx ONTAP is now designated as primary storage with no existing SnapMirror relationships with the original on-premises datacenter. After normal function has resumed on premises, you can use a process identical to the one outlined in this documentation to mirror data back to the on-premises ONTAP storage system.

As is also outlined in this documentation, you can configure SnapCenter to mirror the application data volumes from FSx ONTAP to an ONTAP storage system residing on premises. Similarly, you can configure Veeam to replicate backup copies to Amazon S3 using a scale-out backup repository so that those backups are accessible to a Veeam backup server residing at the on-premises datacenter.

Failback is outside the scope of this documentation, but failback differs little from the detailed process outlined here.

Conclusion

The use case presented in this documentation focuses on proven disaster recovery technologies that highlight the integration between NetApp and VMware. NetApp ONTAP storage systems provide proven data-mirroring technologies that allow organizations to design disaster recovery solutions that span on-premises and ONTAP

technologies residing with the leading cloud providers.

FSx ONTAP on AWS is one such solution that allows for seamless integration with SnapCenter and SyncMirror for replicating application data to the cloud. Veeam Backup & Replication is another well-known technology that integrates well with NetApp ONTAP storage systems and can provide failover to vSphere- native storage.

This solution presented a disaster recovery solution using guest connect storage from an ONTAP system hosting SQL Server and Oracle application data. SnapCenter with SnapMirror provides an easy-to-manage solution for protecting application volumes on ONTAP systems and replicating them to FSx or CVO residing in the cloud. SnapCenter is a DR-enabled solution for failing over all application data to VMware Cloud on AWS.

Veeam Backup & Restore in VMware Cloud, with Amazon FSx ONTAP

Veeam Backup & Replication is an effective and reliable solution for protecting data in VMware Cloud. This solution demonstrates the proper setup and configuration for using Veeam Backup and Replication to backup and restore application VMs residing on FSx ONTAP NFS datastores in VMware Cloud.

Overview

VMware Cloud (in AWS) supports the use of NFS datastores as supplemental storage, and FSx ONTAP is a secure solution for customers who need to store large amounts of data for their cloud applications that can scale independent of the number of ESXi hosts in the SDDC cluster. This integrated AWS storage service offers highly efficient storage with all of the traditional NetApp ONTAP capabilities.

Use Cases

This solution addresses the following use cases:

- Backup and restore of Windows and Linux virtual machines hosted in VMC using FSx ONTAP as a backup repository.
- Backup and restore of Microsoft SQL Server application data using FSx ONTAP as a backup repository.
- Backup and restore of Oracle application data using FSx ONTAP as a backup repository.

NFS Datastores Using Amazon FSx ONTAP

All virtual machines in this solution reside on FSx ONTAP supplemental NFS datastores. Using FSx ONTAP as a supplemental NFS datastore has several benefits. For example, it allows you to:

- Create a scalable and highly available file system in the cloud without the need for complex setup and management.
- Integrate with your existing VMware environment, allowing you to use familiar tools and processes to manage your cloud resources.
- Benefit from the advanced data management features provided by ONTAP, such as snapshots and replication, to protect your data and ensure its availability.

▼ Solution Deployment Overview

This list provides the high level steps necessary to configure Veeam Backup & Replication, execute backup and restore jobs using FSx ONTAP as a backup repository, and perform restores of SQL Server and Oracle VMs and databases:

1. Create the FSx ONTAP file system to be used as iSCSI backup repository for Veeam Backup & Replication.
2. Deploy Veeam Proxy to distribute backup workloads and mount iSCSI backup repositories hosted on FSx ONTAP.
3. Configure Veeam Backup Jobs to backup SQL Server, Oracle, Linux and Windows virtual machines.
4. Restore SQL Server virtual machines and individual databases.
5. Restore Oracle virtual machines and individual databases.

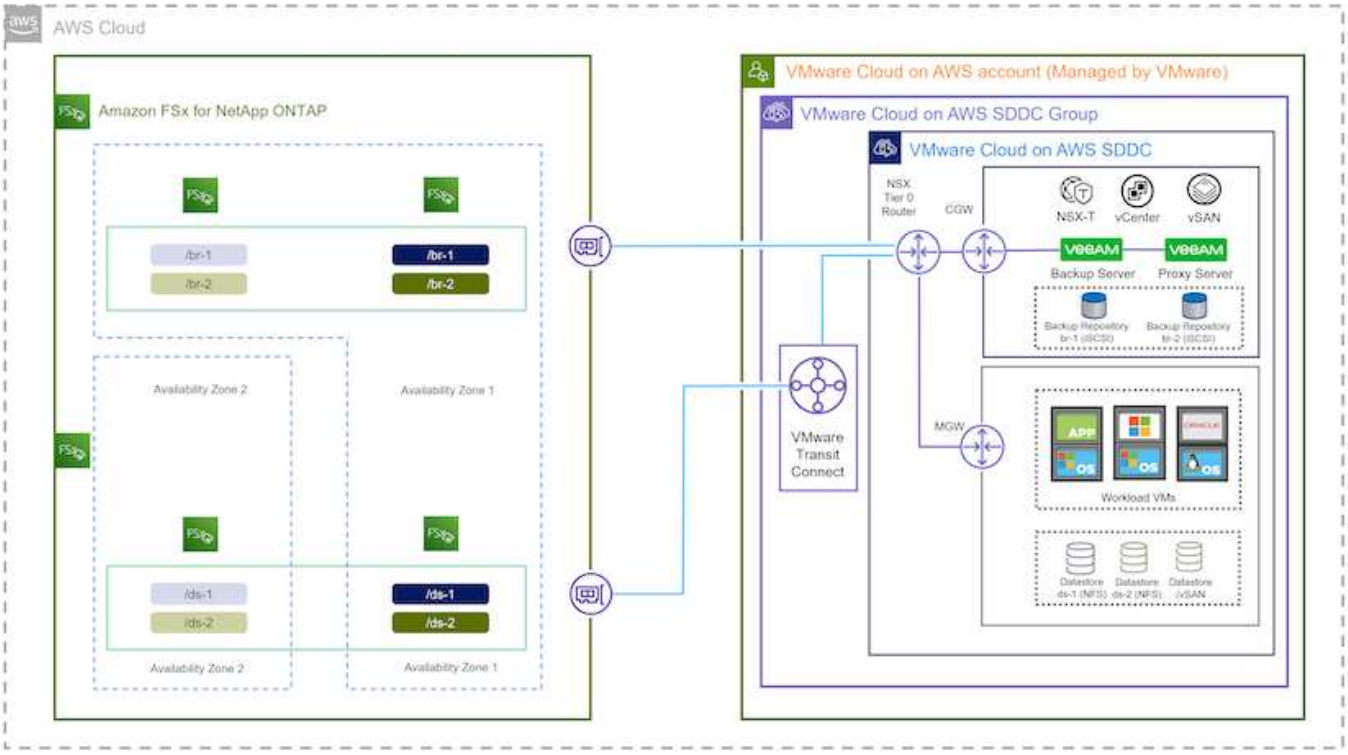
▼ **Prerequisites**

The purpose of this solution is to demonstrate data protection of virtual machines running in VMware Cloud and located on NFS Datastores hosted by FSx ONTAP. This solution assumes the following components are configured and ready for use:

1. FSx ONTAP filesystem with one or more NFS datastores connected to VMware Cloud.
2. Microsoft Windows Server VM with Veeam Backup & Replication software installed.
   - vCenter server has been discovered by the Veeam Backup & Replication server using their IP address or fully qualified domain name.
3. Microsoft Windows Server VM to be installed with Veeam Backup Proxy components during the solution deployment.
4. Microsoft SQL Server VMs with VMDKs and application data residing on FSx ONTAP NFS datastores. For this solution we had two SQL databases on two separate VMDKs.
   - Note: As a best practice database and transaction log files are placed on separate drives as this will improve performance and reliability. This is in part due to the fact that transaction logs are written sequentially, whereas database files are written randomly.
5. Oracle Database VMs with VMDKs and application data residing on FSx ONTAP NFS datastores.
6. Linux and Windows file server VMs with VMDKs residing on FSx ONTAP NFS datastores.
7. Veeam requires specific TCP ports for communication between servers and components in the backup environment. On Veeam backup infrastructure components, the required firewall rules are automatically created.
   For a full listing of the network port requirements refer to the Ports section of the Veeam Backup and Replication User Guide for VMware vSphere.

High Level Architecture

The testing / validation of this solution was performed in a lab that may or may not match the final deployment environment. For more information, please refer to the following sections.



▼ **Hardware / Software Components**

The purpose of this solution is to demonstrate data protection of virtual machines running in VMware Cloud and located on NFS Datastores hosted by FSx ONTAP. This solution assumes the following components are already configured and ready for use:

- Microsoft Windows VM's located on an FSx ONTAP NFS Datastore
- Linux (CentOS) VM's located on an FSx ONTAP NFS Datastore
- Microsoft SQL Server VM's located on an FSx ONTAP NFS Datastore
  - Two databases hosted on separate VMDK's
- Oracle VM's located on an FSx ONTAP NFS Datastore

Solution Deployment

In this solution we provide detailed instructions for deploying and validating a solution utilizing Veeam Backup and Replication software to perform backup and recovery of SQL Server, Oracle, and Windows and Linux file server virtual machines in a VMware Cloud SDDC on AWS. The Virtual Machines in this solution reside on a supplemental NFS datastore hosted by FSx ONTAP. In addition, a separate FSx ONTAP file system is used to host iSCSI volumes that will be used for Veeam backup repositories.

We will go over FSx ONTAP file system creation, mounting iSCSI volumes to be used as backup repositories, creating and running backup jobs, and performing VM and database restores.

For detailed information on FSx ONTAP refer to the FSx ONTAP User Guide.

For detailed information on Veeam Backup and Replication refer to the Veeam Help Center Technical Documentation site.

For considerations and limitations when using Veeam Backup and Replication with VMware Cloud on AWS, refer to VMware Cloud on AWS and VMware Cloud on Dell EMC Support. Considerations and Limitations.

Deploy Veeam Proxy server

A Veeam proxy server is a component of the Veeam Backup & Replication software that acts as an intermediary between the source and the backup or replication target. The proxy server helps to optimize and accelerate data transfer during backup jobs by processing data locally and can use different Transport Modes to access data using VMware vStorage APIs for Data Protection or through direct storage access.

When choosing a Veeam proxy server design it is important to consider the number of concurrent tasks and the transport mode or type of storage access desired.

For sizing the number of proxy servers, and for their system requirements, refer to the Veeam VMware vSphere Best Practice Guide.

The Veeam Data Mover is a component of the Veeam Proxy Server and utilizes a Transport Mode as a method for obtaining VM data from the source and transferring it to the target. The transport mode is specified during the configuration of the backup job. It is possible to increase the efficiency backups from NFS datastores by using direct storage access.

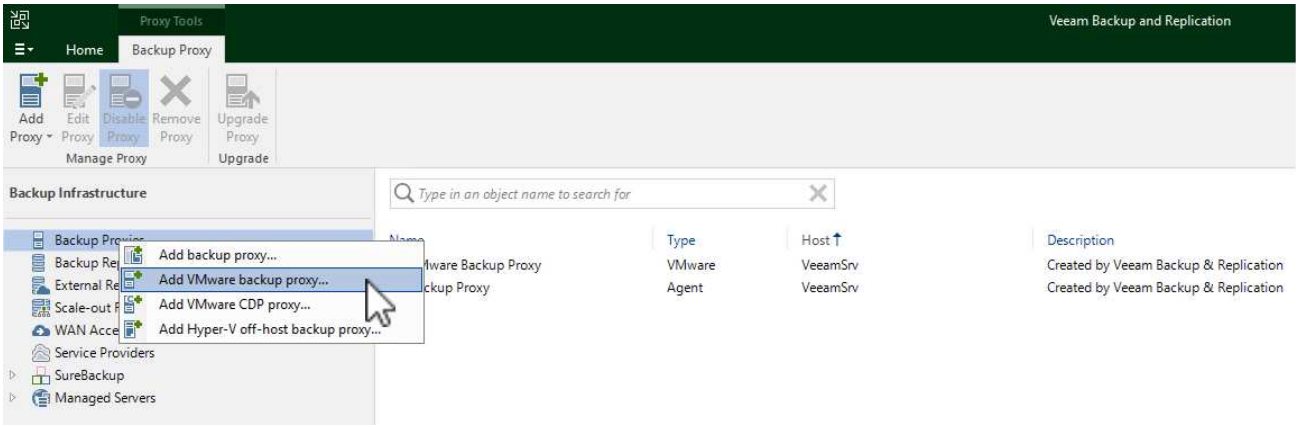
For more information on Transport Modes refer to the Veeam Backup and Replication User Guide for VMware vSphere.

In the following step we cover deployment of the Veeam Proxy Server on a Windows VM in the VMware Cloud SDDC.

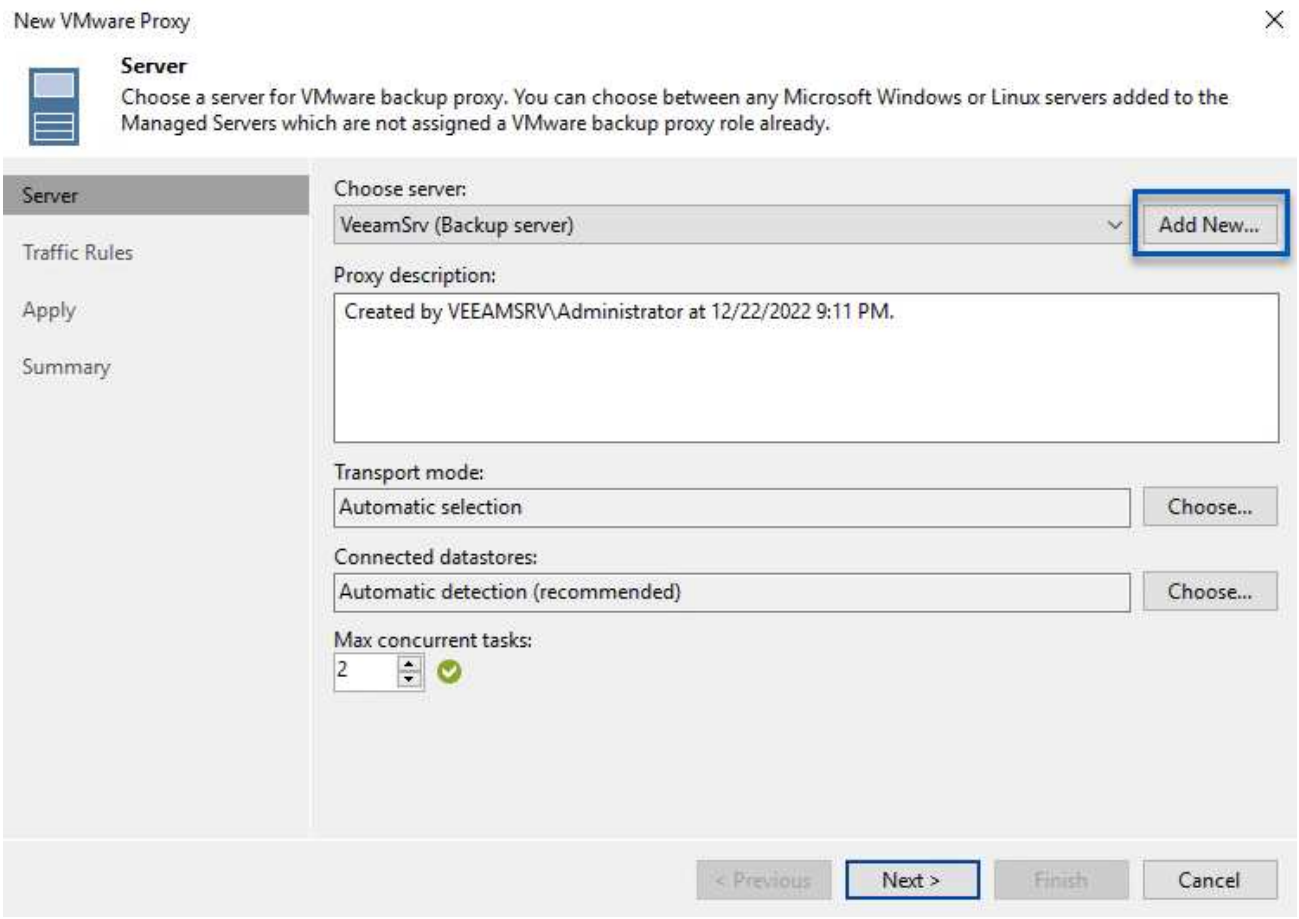
▼ **Deploy Veeam Proxy to distribute backup workloads**

In this step the Veeam Proxy is deployed to an existing Windows VM. This allows backup jobs to be distributed between the primary Veeam Backup Server and the Veeam Proxy.

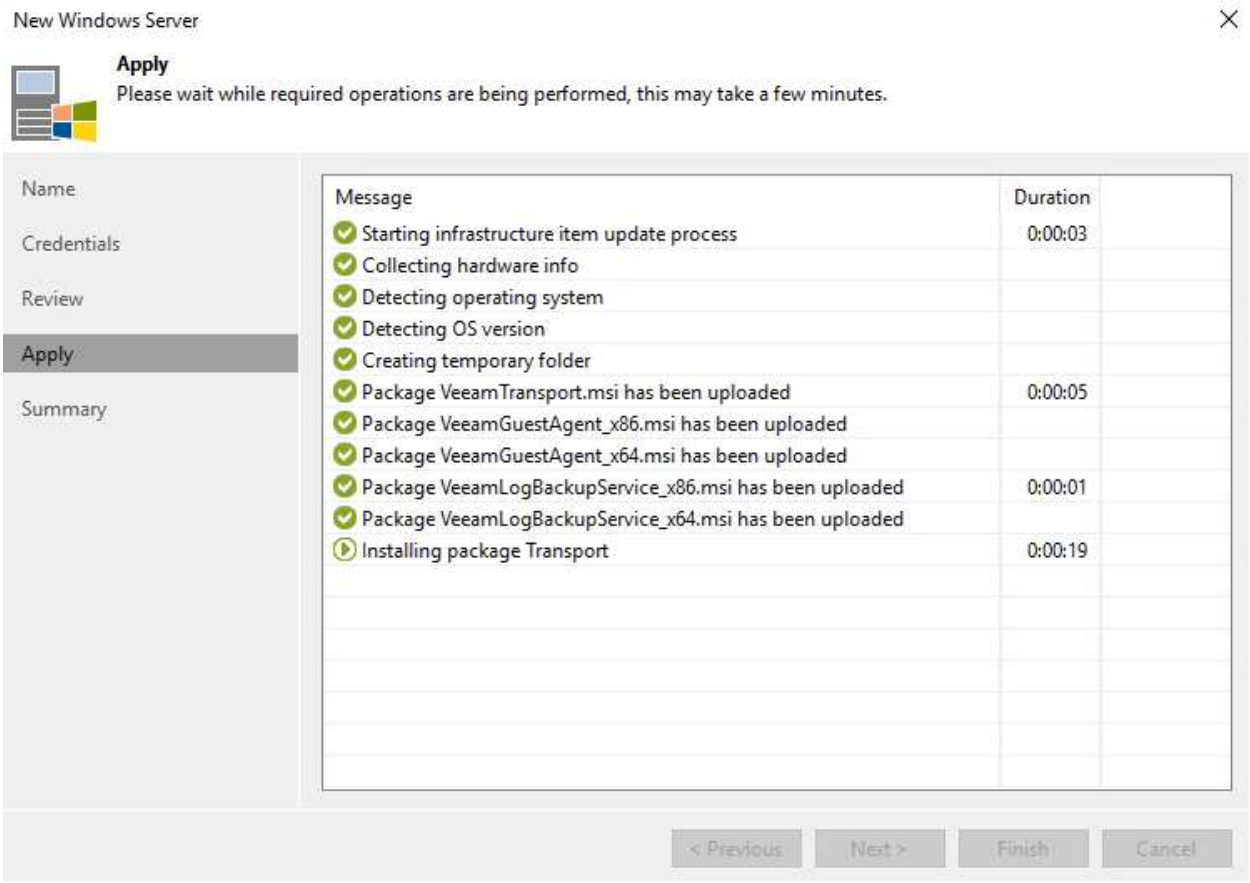
1. On the Veeam Backup and Replication server, open the administration console and select **Backup Infrastructure** in the lower left menu.
2. Right click on **Backup Proxies** and click on **Add VMware backup proxy…** to open the wizard.



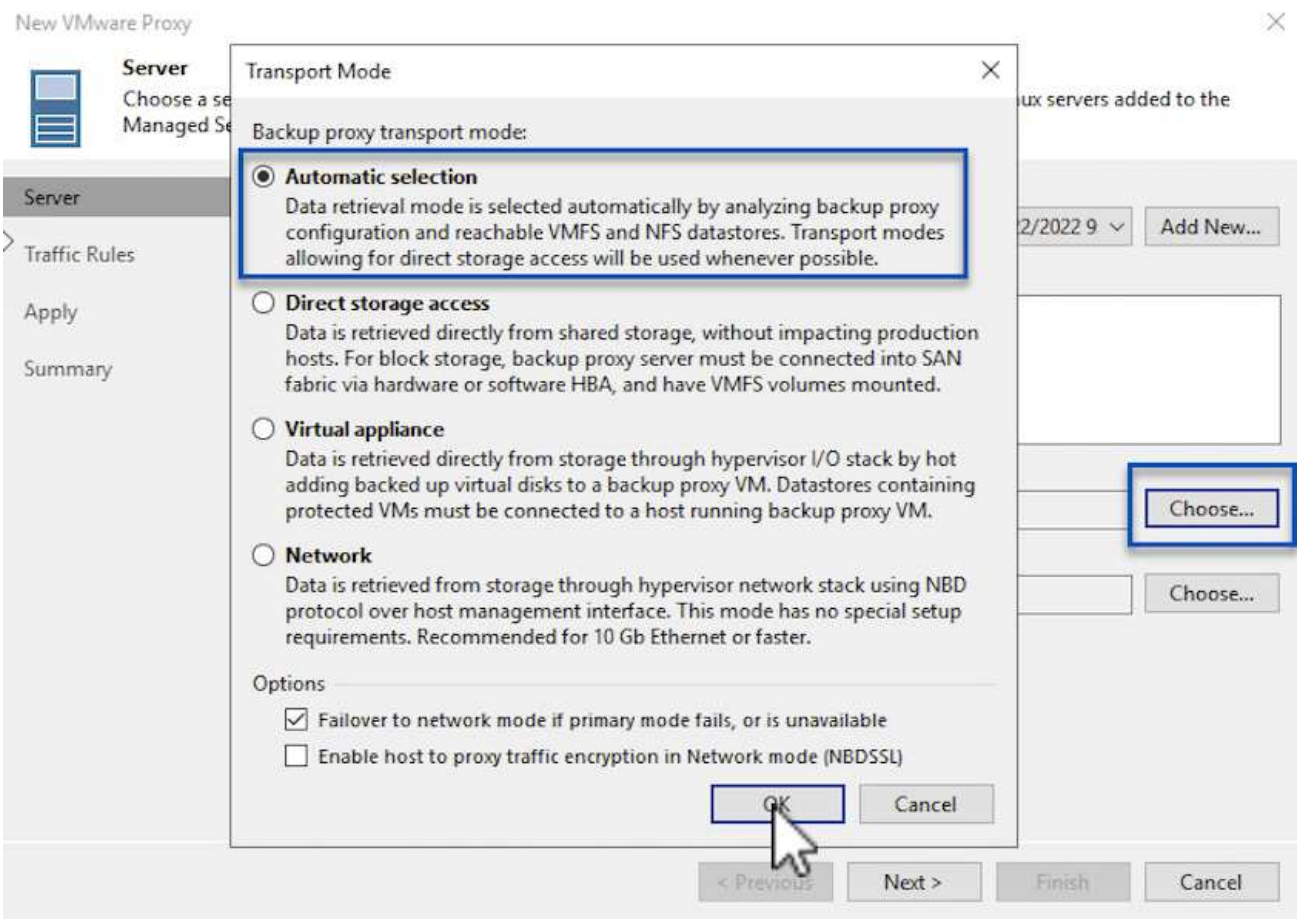
3. In the **Add VMware Proxy** wizard click the **Add New…** button to add a new proxy server.



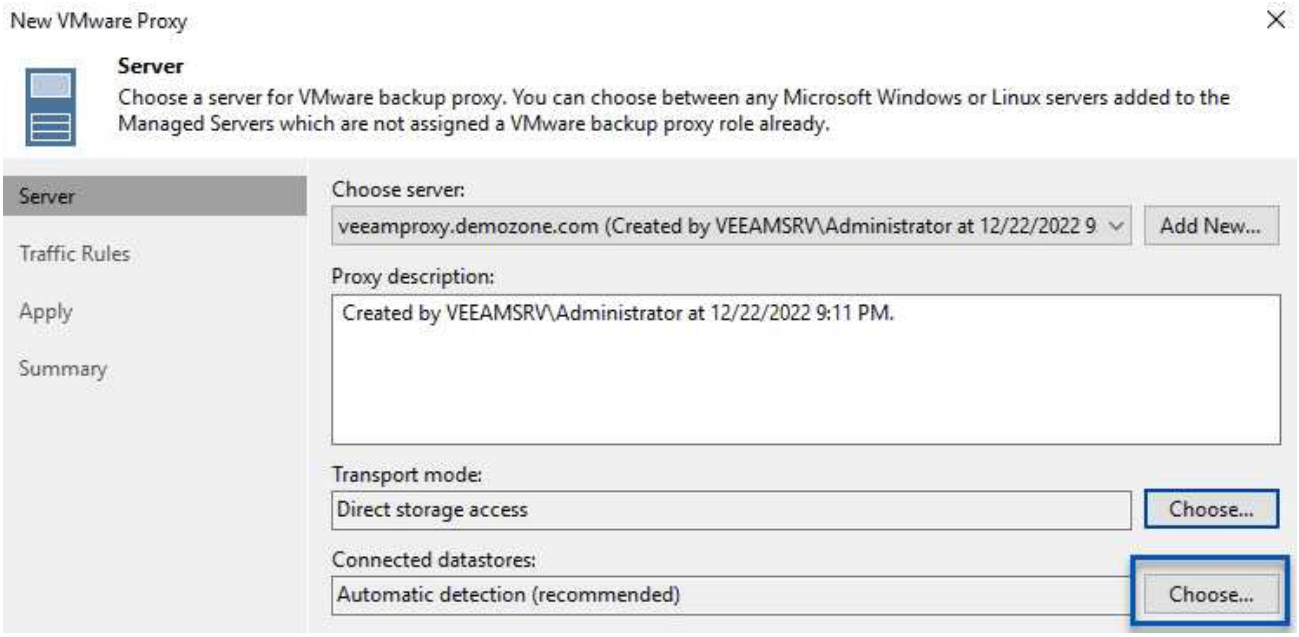
4. Select to add Microsoft Windows and follow the prompts to add the server:
   - Fill out the DNS name or IP address
   - Select an account to use for Credentials on the new system or add new credentials
   - Review the components to be installed and then click on **Apply** to begin the deployment

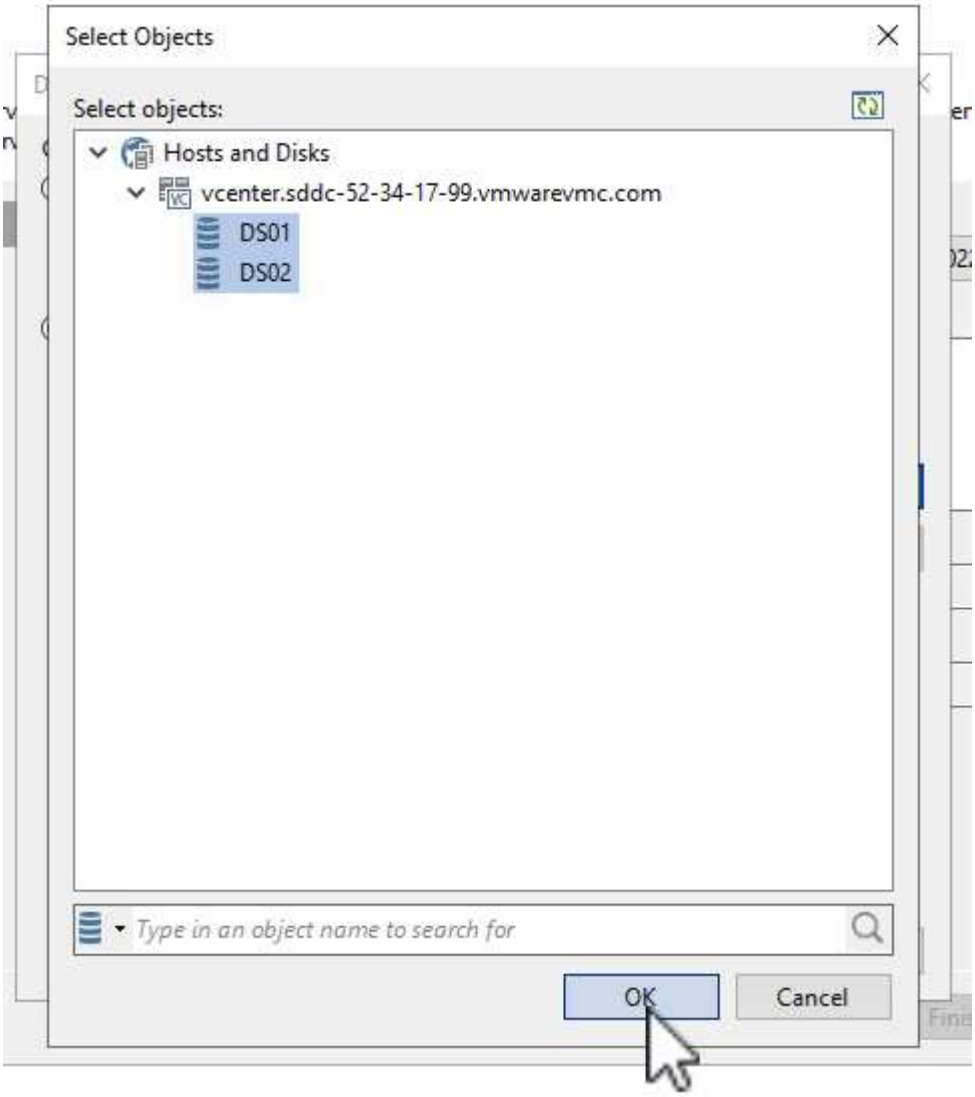


5. Back in the **New VMware Proxy** wizard, choose a Transport Mode. In our case we chose **Automatic Selection**.

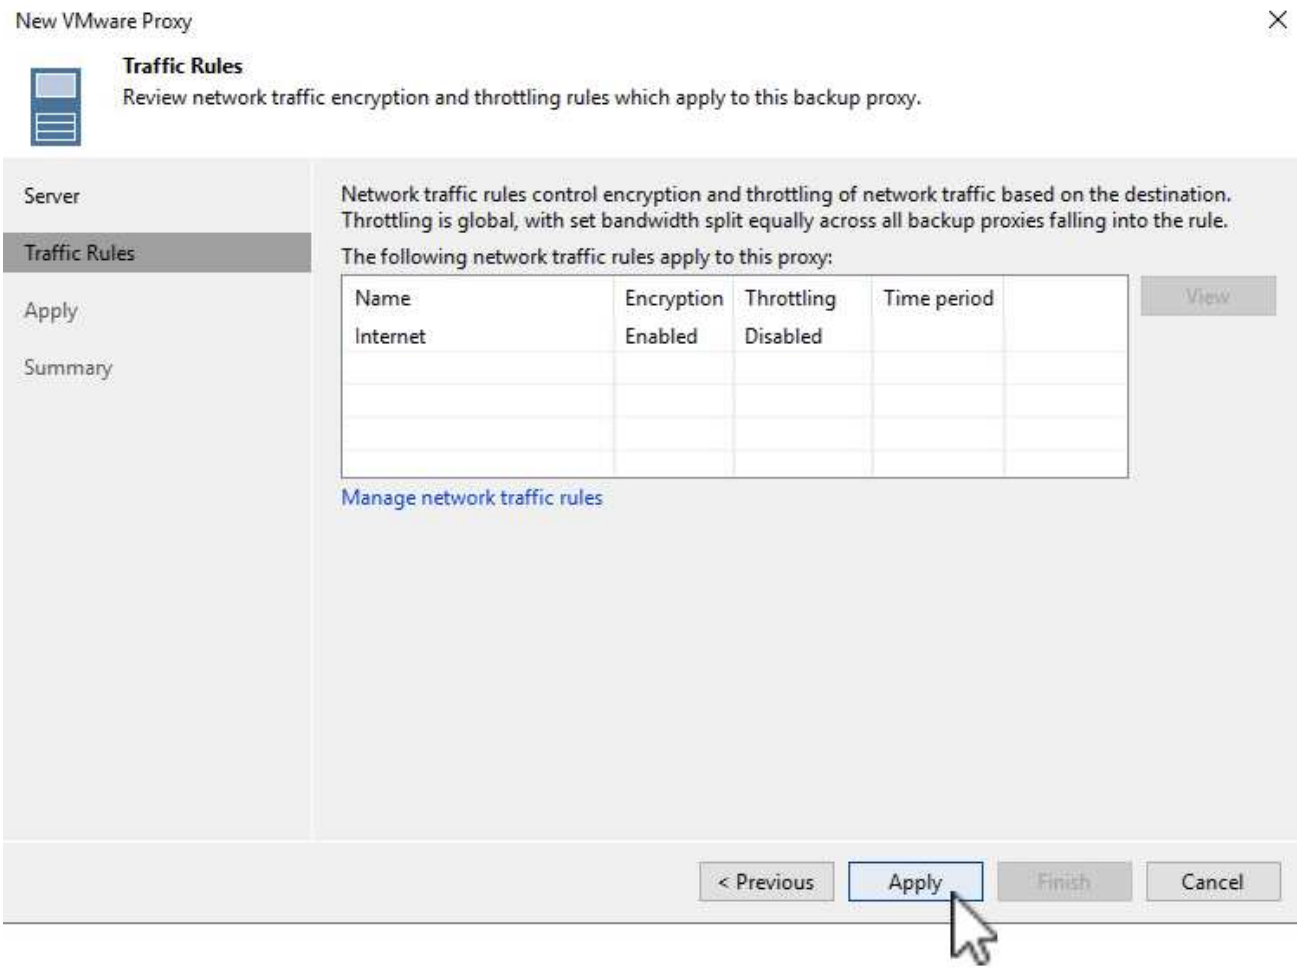


6. Select the Connected datastores that you want the VMware Proxy to have direct access to.





7. Configure and apply any specific network traffic rules such as encryption or throttling that are desired. When complete click on the **Apply** button to complete the deployment.



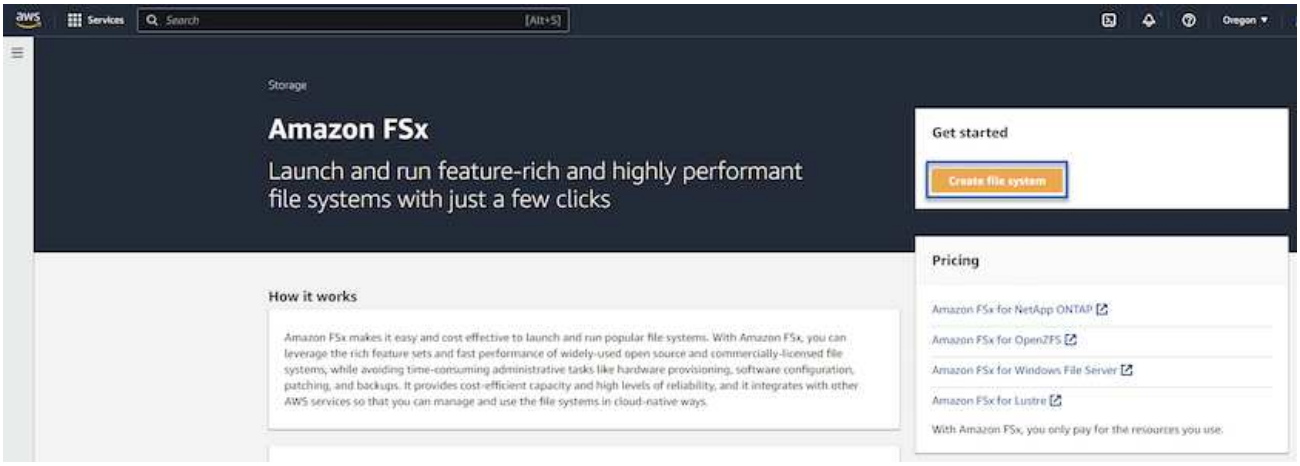
Configure storage and Backup Repositories

The primary Veeam Backup server and Veeam Proxy server have access to a backup repository in the form of direct connected storage. In this section we cover creating an FSx ONTAP file system, mounting iSCSI LUNs to the Veeam servers and creating Backup Repositories.

▼ Create FSx ONTAP file system

Create an FSx ONTAP file system that will be used to host the iSCSI volumes for the Veeam Backup Repositories.

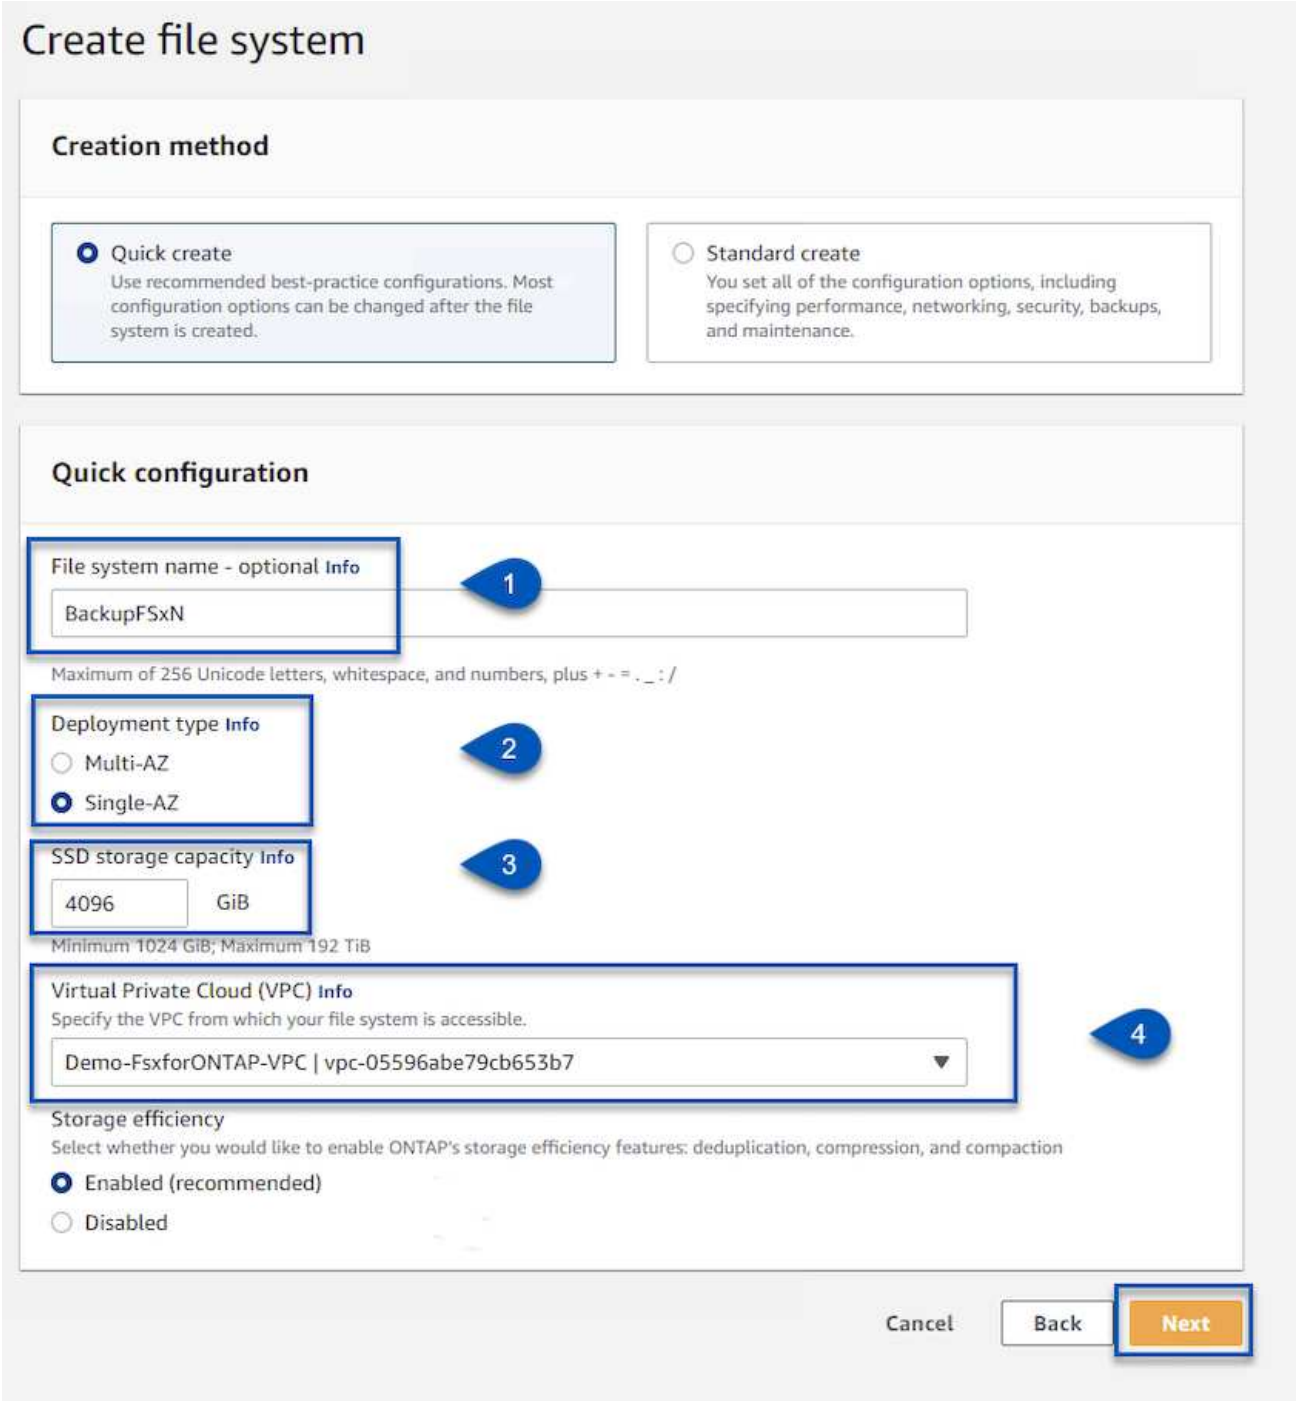
1. In the AWS console, Go to FSx and then **Create file system**



2. Select **Amazon FSx ONTAP** and then **Next** to continue.



3. Fill in the file system name, deployment type, SSD storage capacity and the VPC in which the FSx ONTAP cluster will reside. This must be a VPC configured to communicate with the virtual machine network in VMware Cloud. Click on **Next**.



4. Review the deployment steps and click on **Create File System** to begin the file system creation process.

▼ Configure and mount iSCSI LUNs

Create and configure the iSCSI LUNs on FSx ONTAP and mount to the Veeam backup and proxy servers. These LUNs will later be used to create Veeam backup repositories.

ⓘ Creating an iSCSI LUN on FSx ONTAP is a multi-step process. The first step of creating the volumes can be accomplished in the Amazon FSx Console or with the NetApp ONTAP CLI.

ⓘ For more information on using FSx ONTAP, see the FSx ONTAP User Guide.

1. From the NetApp ONTAP CLI create the initial volumes using the following command:

```
FSx-Backup::> volume create -vserver svm_name -volume vol_name
-aggregate aggregate_name -size vol_size -type RW
```

2. Create LUNs using the volumes created in the previous step:

```
FSx-Backup::> lun create -vserver svm_name -path
/vol/vol_name/lun_name -size size -ostype windows -space-allocation
enabled
```

3. Grant access to the LUNs by creating an initiator group containing the iSCSI IQN of the Veeam backup and proxy servers:

```
FSx-Backup::> igroup create -vserver svm_name -igroup igroup_name
-protocol iSCSI -ostype windows -initiator IQN
```

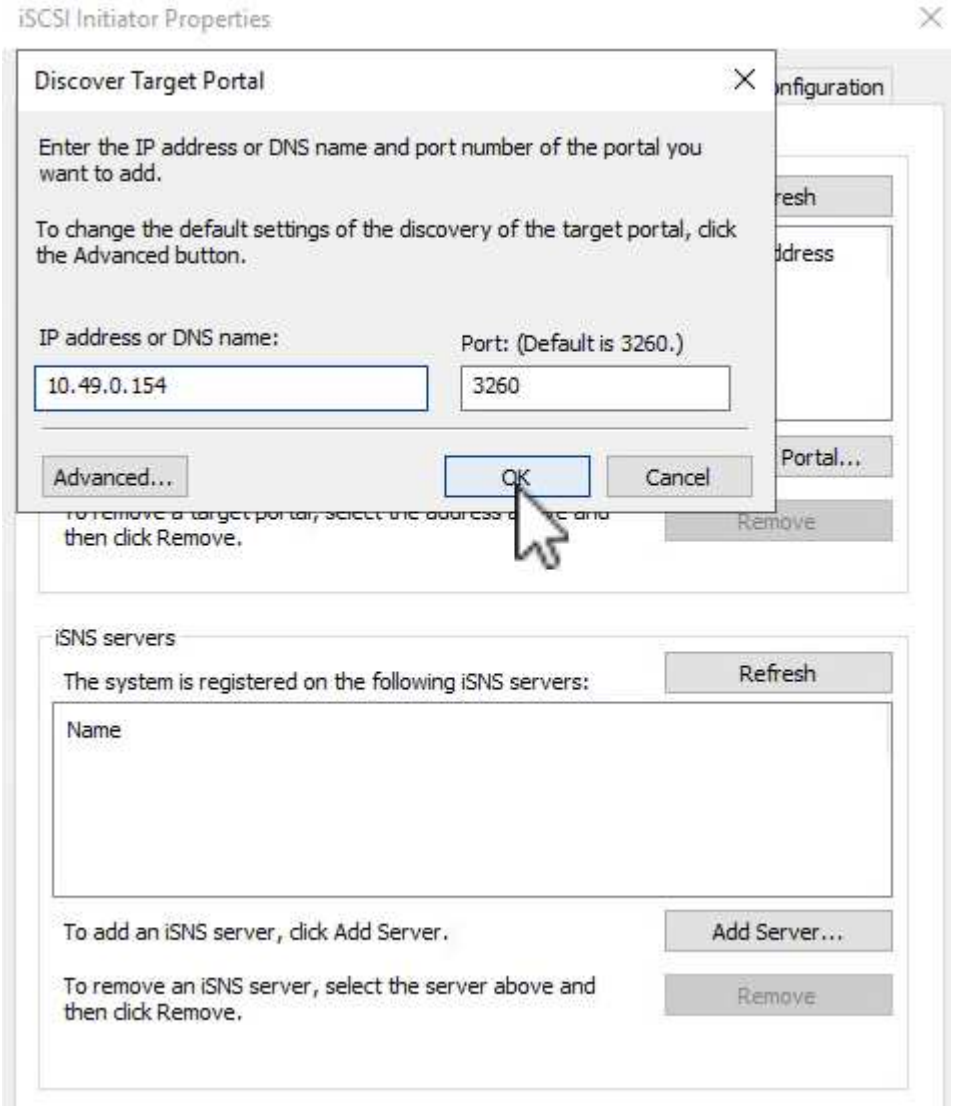


To complete the preceding step you will need to first retrieve the IQN from the iSCSI initiator properties on the Windows servers.

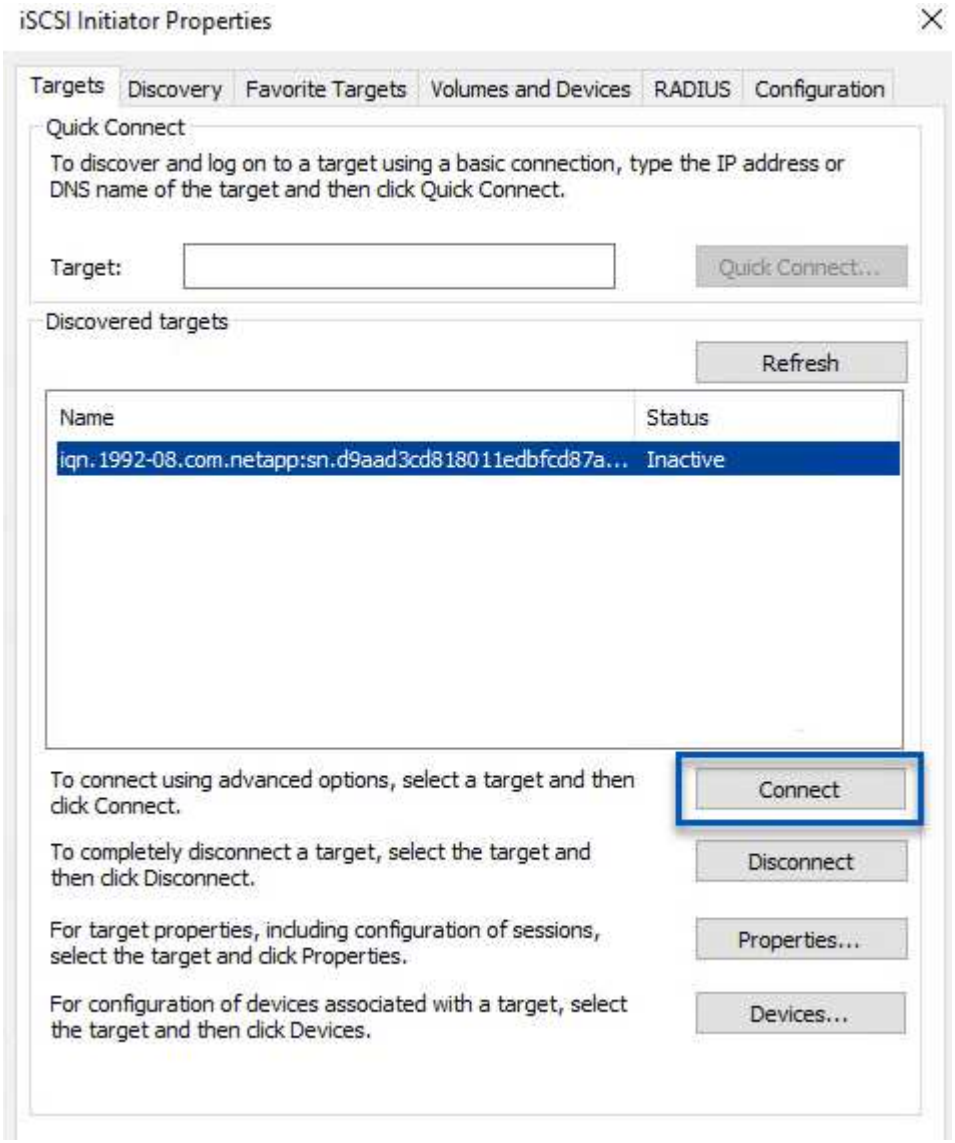
4. Finally, map the LUNs to the initiator group that you just created:

```
FSx-Backup::> lun mapping create -vserver svm_name -path
/vol/vol_name/lun_name igroup igroup_name
```

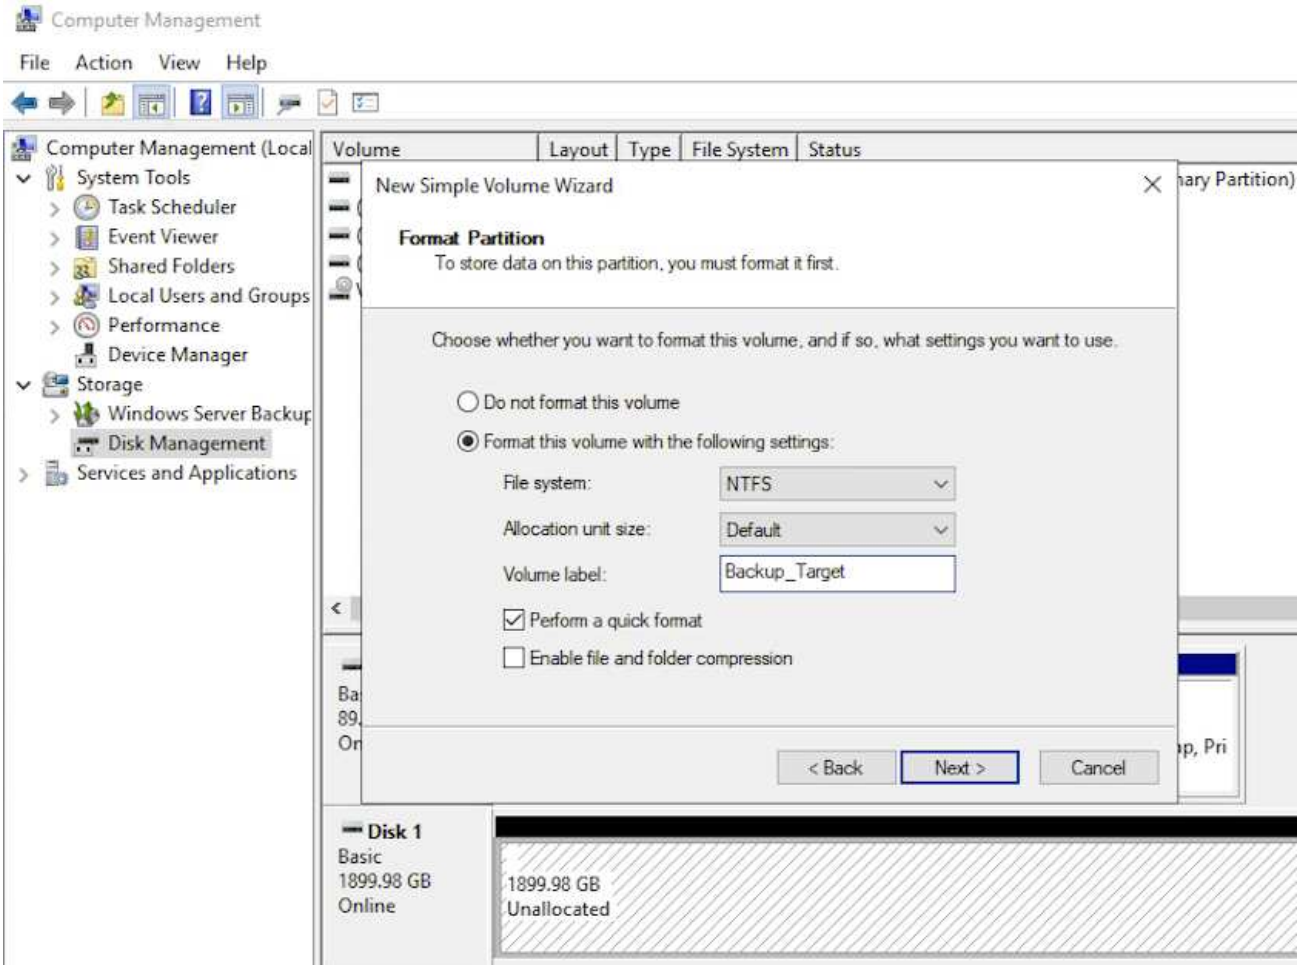
5. To mount the iSCSI LUNs, log into the Veeam Backup & Replication Server and open iSCSI Initiator Properties. Go to the **Discover** tab and enter the iSCSI target IP address.



6. On the **Targets** tab, highlight the inactive LUN and click on **Connect**. Check the **Enable multi-path** box and click on **OK** to connect to the LUN.



7. In the Disk Management utility initialize the new LUN and create a volume with the desired name and drive letter. Check the **Enable multi-path** box and click on **OK** to connect to the LUN.

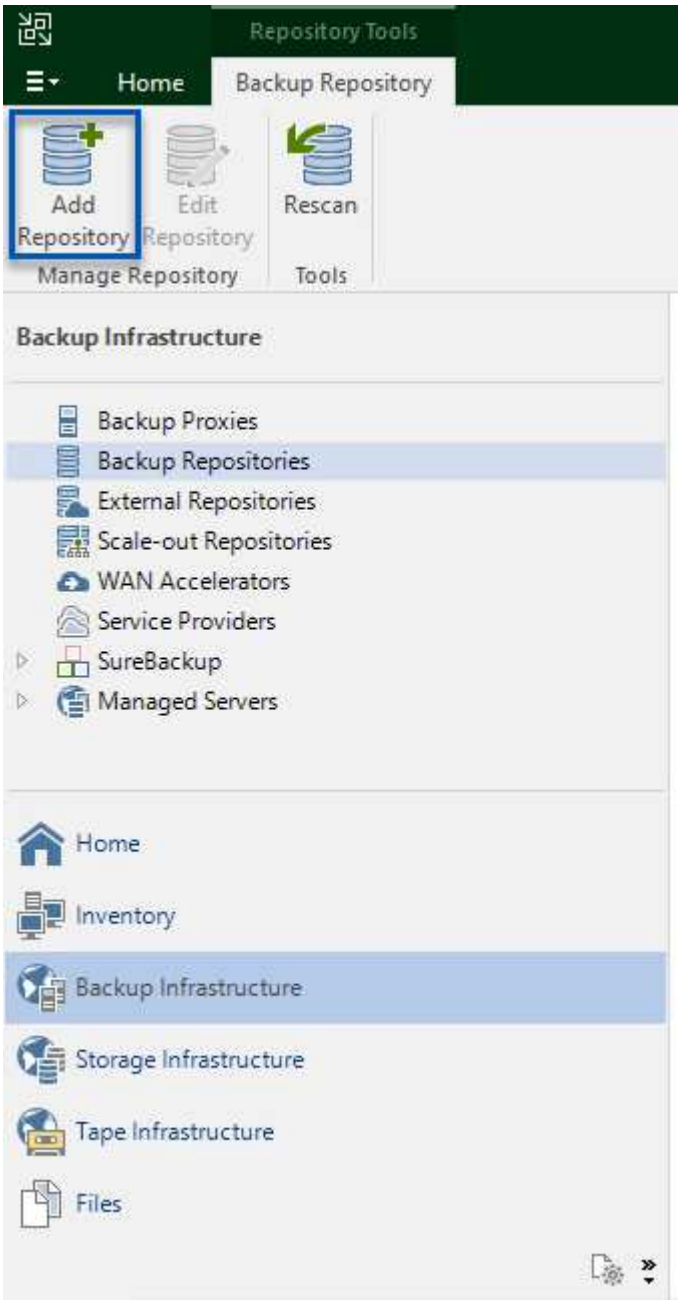


8. Repeat these steps to mount the iSCSI volumes on the Veeam Proxy server.

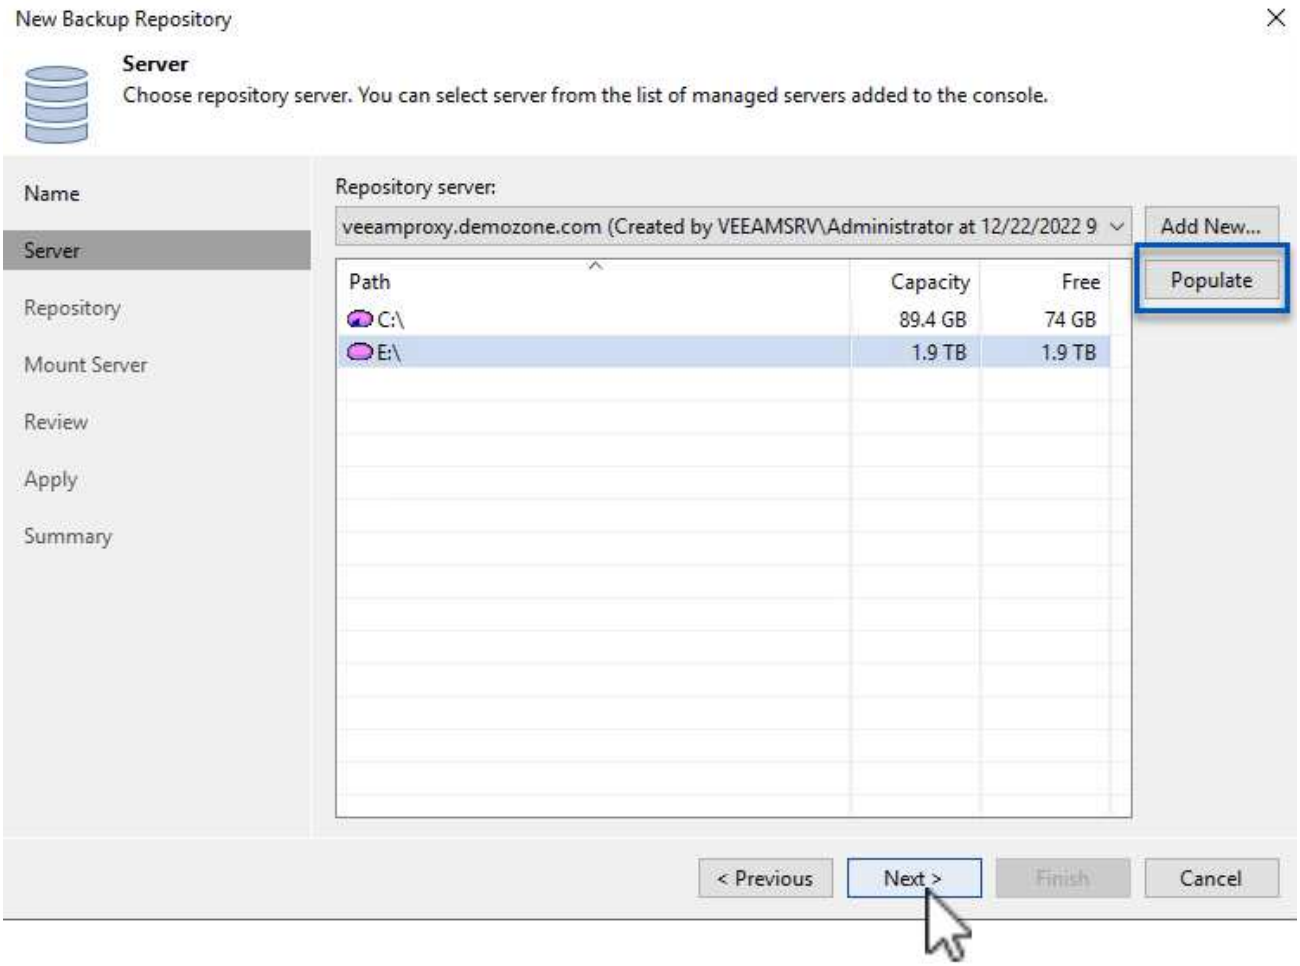
▼ **Create Veeam Backup Repositories**

In the Veeam Backup and Replication console, create backup repositories for the Veeam Backup and Veeam Proxy servers. These repositories will be used as backup targets for the virtual machines backups.

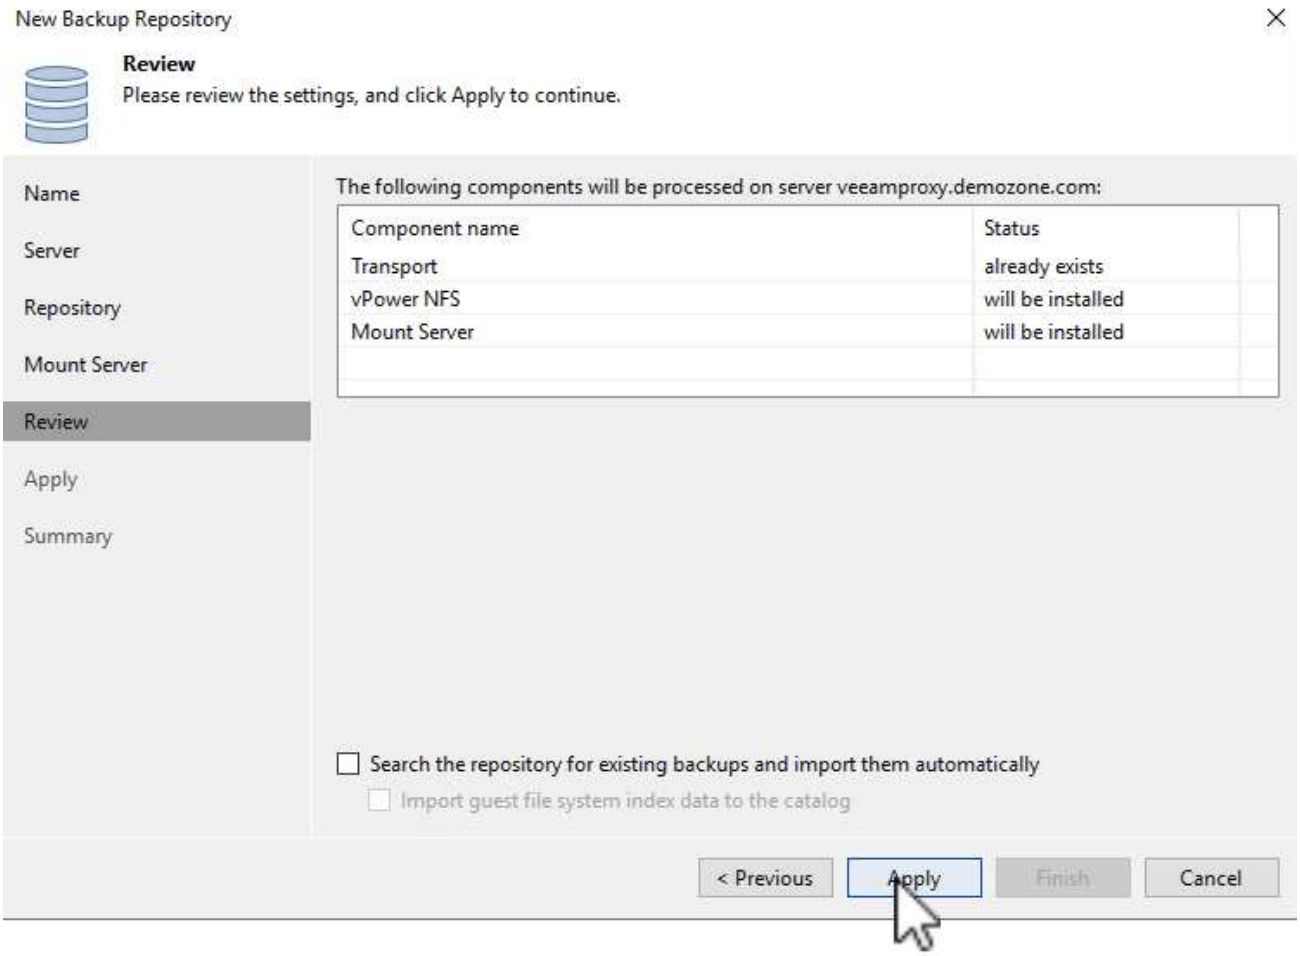
1. In the Veeam Backup and Replication console click on **Backup Infrastructure** in the lower left and then select **Add Repository**



2. In the New Backup Repository wizard, enter a name for the repository and then select the server from the drop-down list and click on the **Populate** button to choose the NTFS volume that will be used.



3. On the next page choose a Mount server that will be used to mount backups to when performing advanced restores. By default this is the same server that has the repository storage connected.
4. Review your selections and click on **Apply** to start the backup repository creation.



5. Repeat these steps for any additional proxy servers.

Configure Veeam backup jobs

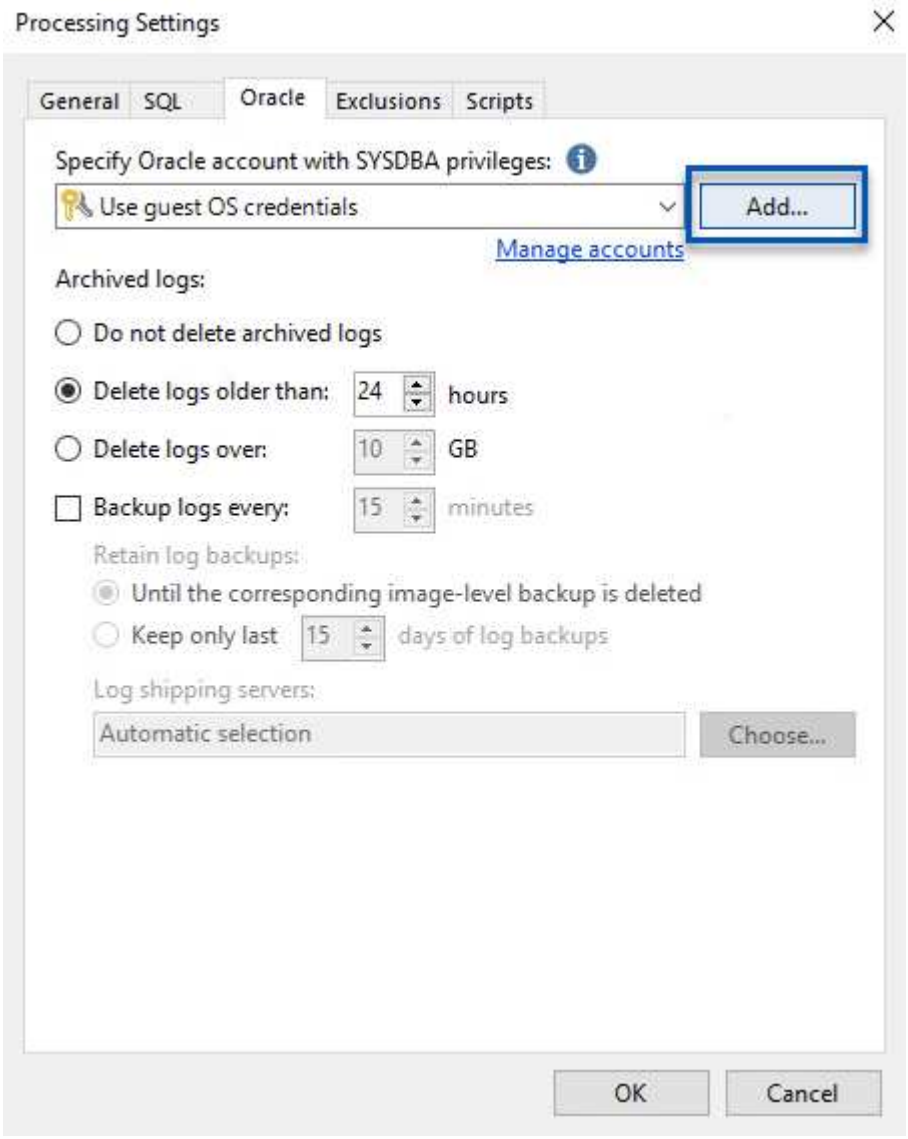
Backup jobs should be created utilizing the the Backup Repositories in the previous section. Creating backup jobs is a normal part of any storage administrator's repertoire and we do not cover all of the steps here. For more complete information on creating backup jobs in Veeam, see the Veeam Help Center Technical Documentation.

In this solution separate backup jobs were created for:

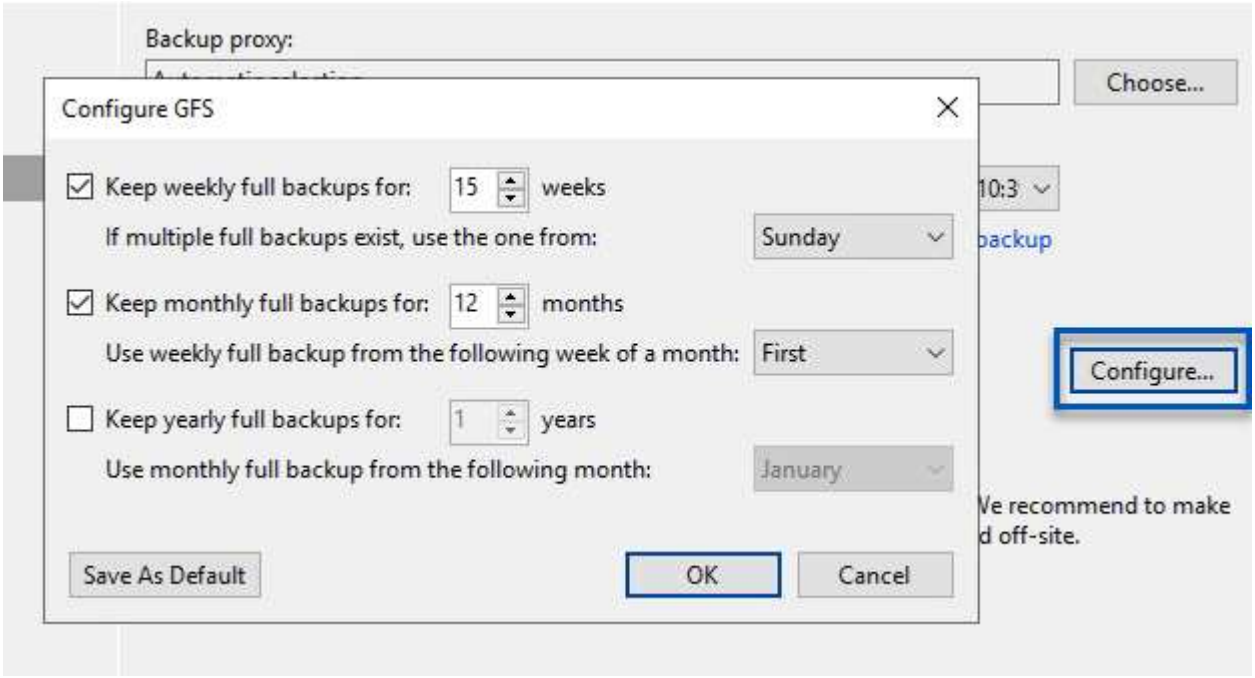
- Microsoft Windows SQL Servers
- Oracle database servers
- Windows file servers
- Linux file servers

▼ General considerations when configuring Veeam backup jobs

1. Enable application-aware processing to create consistent backups and perform transaction log processing.
2. After enabling application-aware processing add the correct credentials with admin privileges to the application as this may be different than the guest OS credentials.



3. To manage the retention policy for the backup check the **Keep certain full backups longer for archival purposes** and click the **Configure…** button to configure the policy.



Restore Application VMs with Veeam full restore

Performing a full restore with Veeam is the first step in performing an application restore. We validated that full restores of our VMs powered on and all services were running normally.

Restoring servers is a normal part of any storage administrator's repertoire and we do not cover all of the steps here. For more complete information on performing full restores in Veeam, see the Veeam Help Center Technical Documentation.

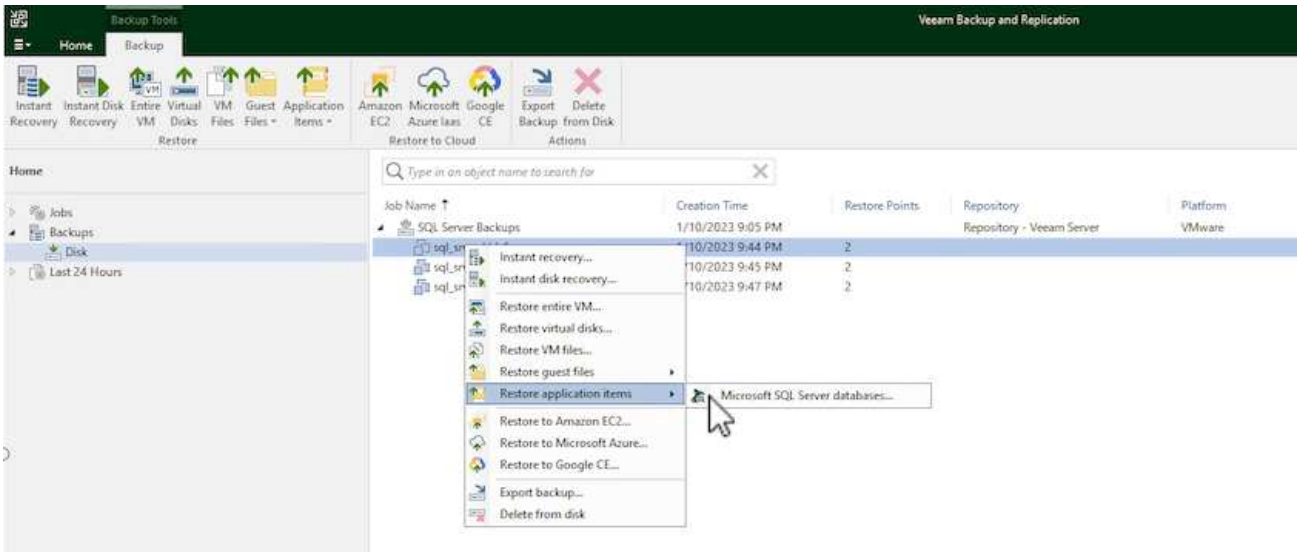
Restore SQL Server databases

Veeam Backup & Replication provides several options for restoring SQL Server databases. For this validation we used the Veeam Explorer for SQL Server with Instant Recovery to execute restores of our SQL Server databases. SQL Server Instant Recovery is a feature that allows you to quickly restore SQL Server databases without having to wait for a full database restore. This rapid recovery process minimizes downtime and ensures business continuity. Here's how it works:

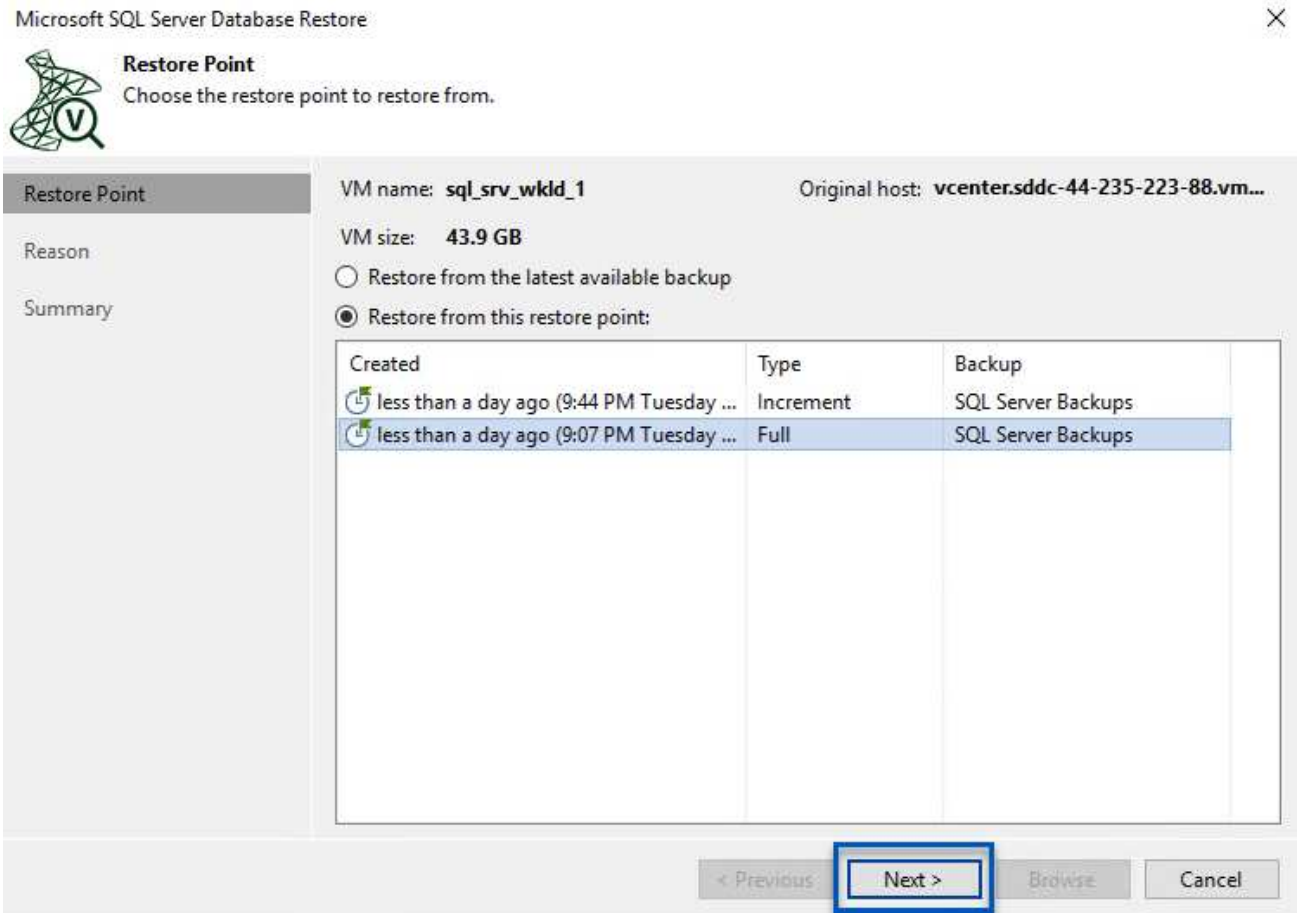
- Veeam Explorer **mounts the backup** containing the SQL Server database to be restored.
- The software **publishes the database** directly from the mounted files, making it accessible as a temporary database on the target SQL Server instance.
- While the temporary database is in use, Veeam Explorer **redirects user queries** to this database, ensuring that users can continue to access and work with the data.
- In the background, Veeam **performs a full database restore**, transferring data from the temporary database to the original database location.
- Once the full database restore is complete, Veeam Explorer **switches user queries back to the original** database and removes the temporary database.

▼ Restore SQL Server database with Veeam Explorer Instant Recovery

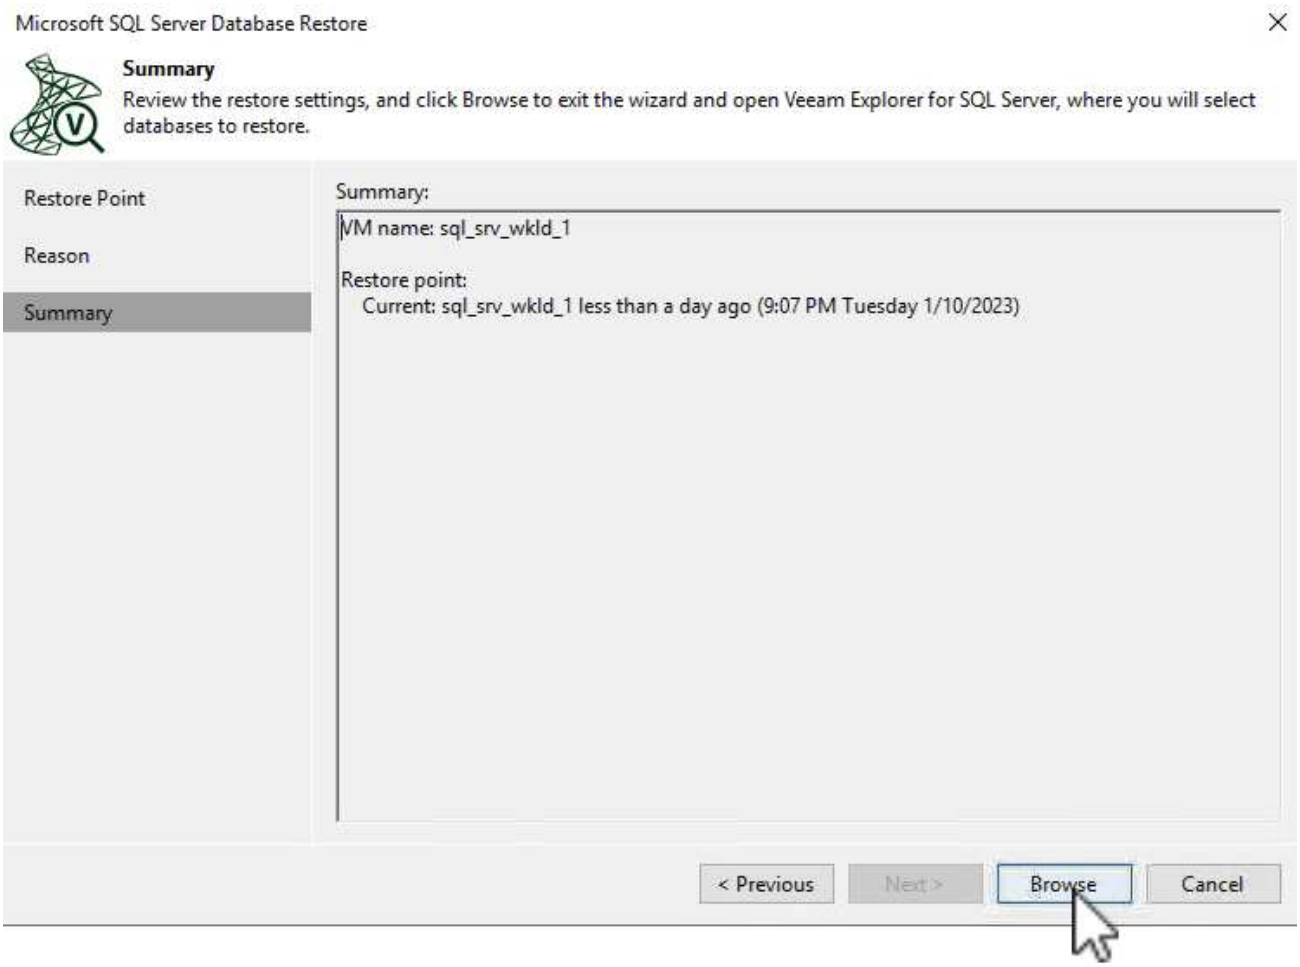
1. In the Veeam Backup and Replication console, navigate to the list of SQL Server backups, right click on a server and select **Restore application items** and then **Microsoft SQL Server databases…**.



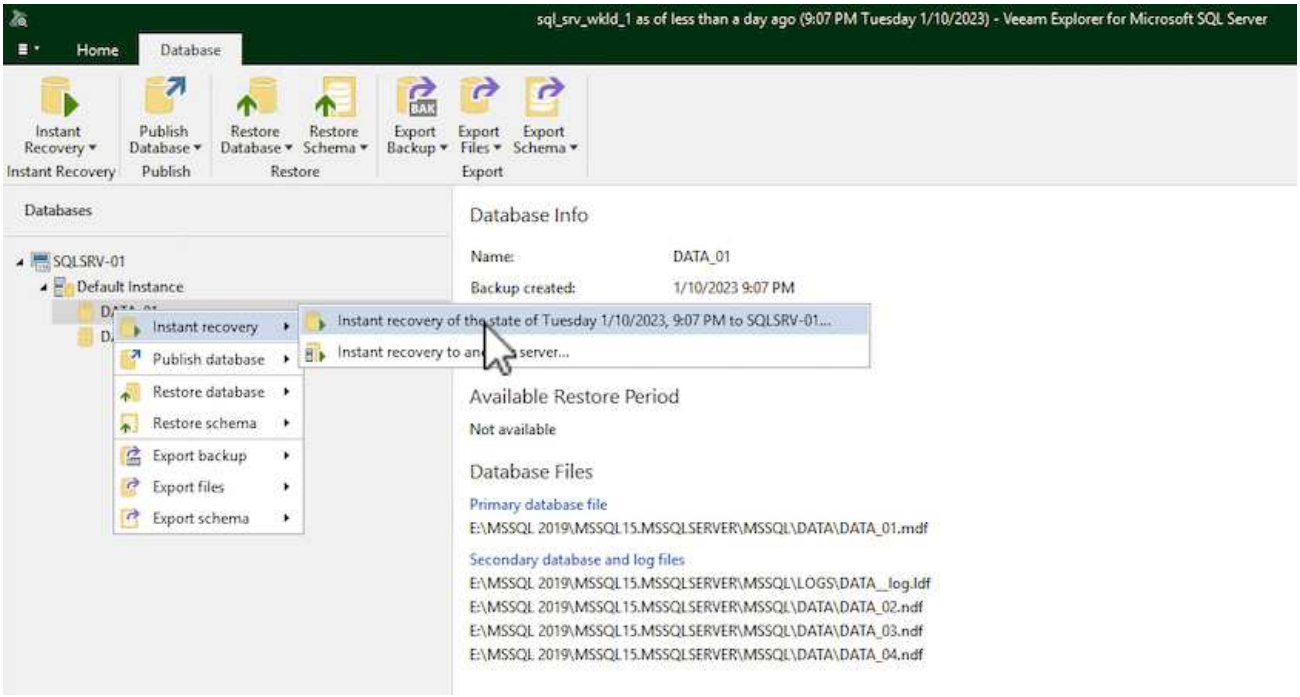
2. In the Microsoft SQL Server Database Restore Wizard select a restore point from the list and click on **Next**.



3. Enter a **Restore reason** if desired and then, on the Summary page, click on the **Browse** button to launch Veeam Explorer for Microsoft SQL Server.

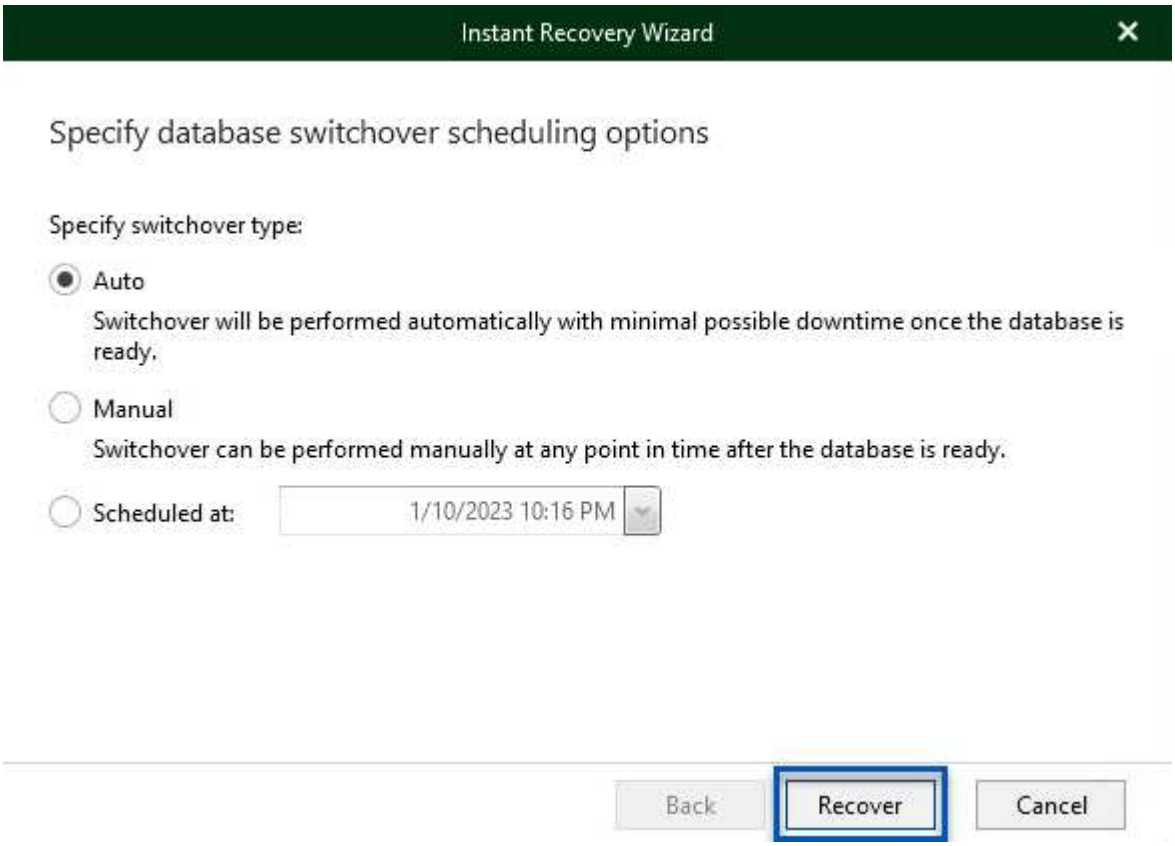


4. In Veeam Explorer expand the list of database instances, right click and select **Instant recovery** and then the specific restore point to recover to.

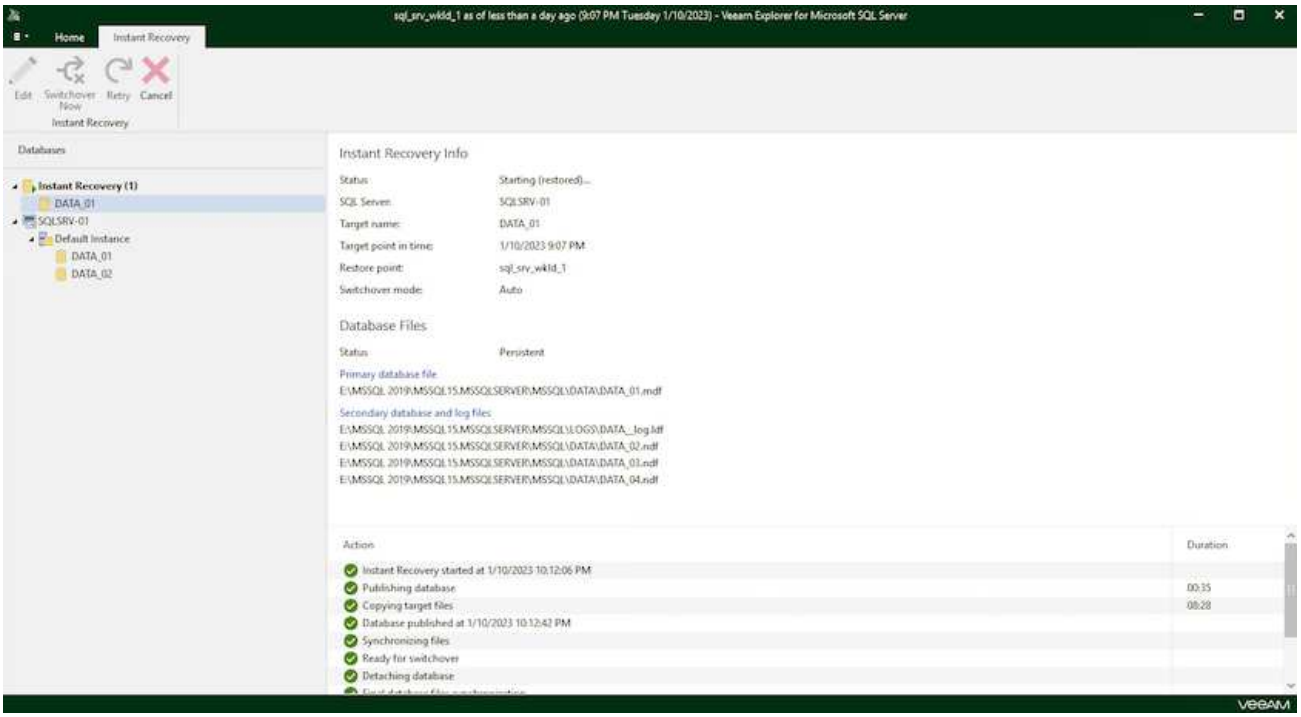


5. In the Instant Recovery Wizard specify the switchover type. This can either be automatically with

minimal downtime, manually, or at a specified time. Then click the **Recover** button to begin the restore process.



6. The recovery process can be monitored from Veeam Explorer.



For more detailed information on performing SQL Server restore operations with Veeam Explorer refer to the Microsoft SQL Server section in the Veeam Explorers User Guide.

Restore Oracle databases with Veeam Explorer

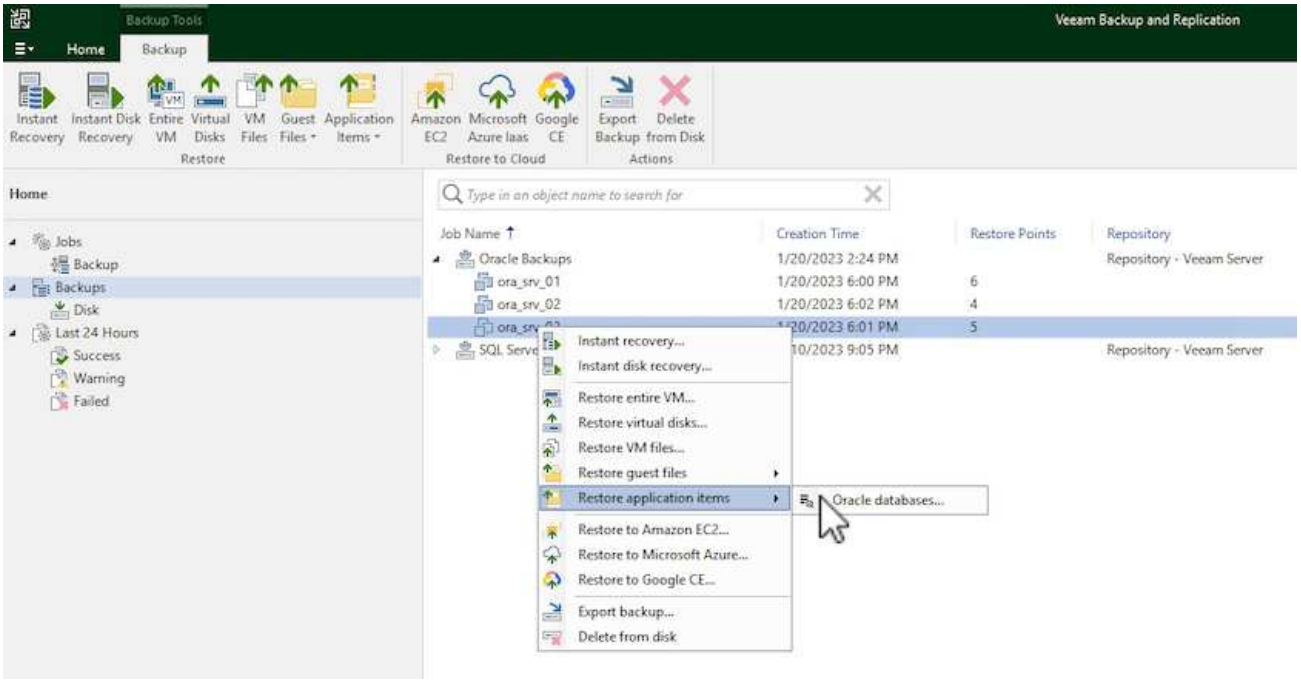
Veeam Explorer for Oracle database provides the ability to perform a standard Oracle database restore or an uninterrupted restore using Instant Recovery. It also supports publishing databases for fast access, recovery of Data Guard databases and restores from RMAN backups.

For more detailed information on performing Oracle database restore operations with Veeam Explorer refer to the Oracle section in the Veeam Explorers User Guide.

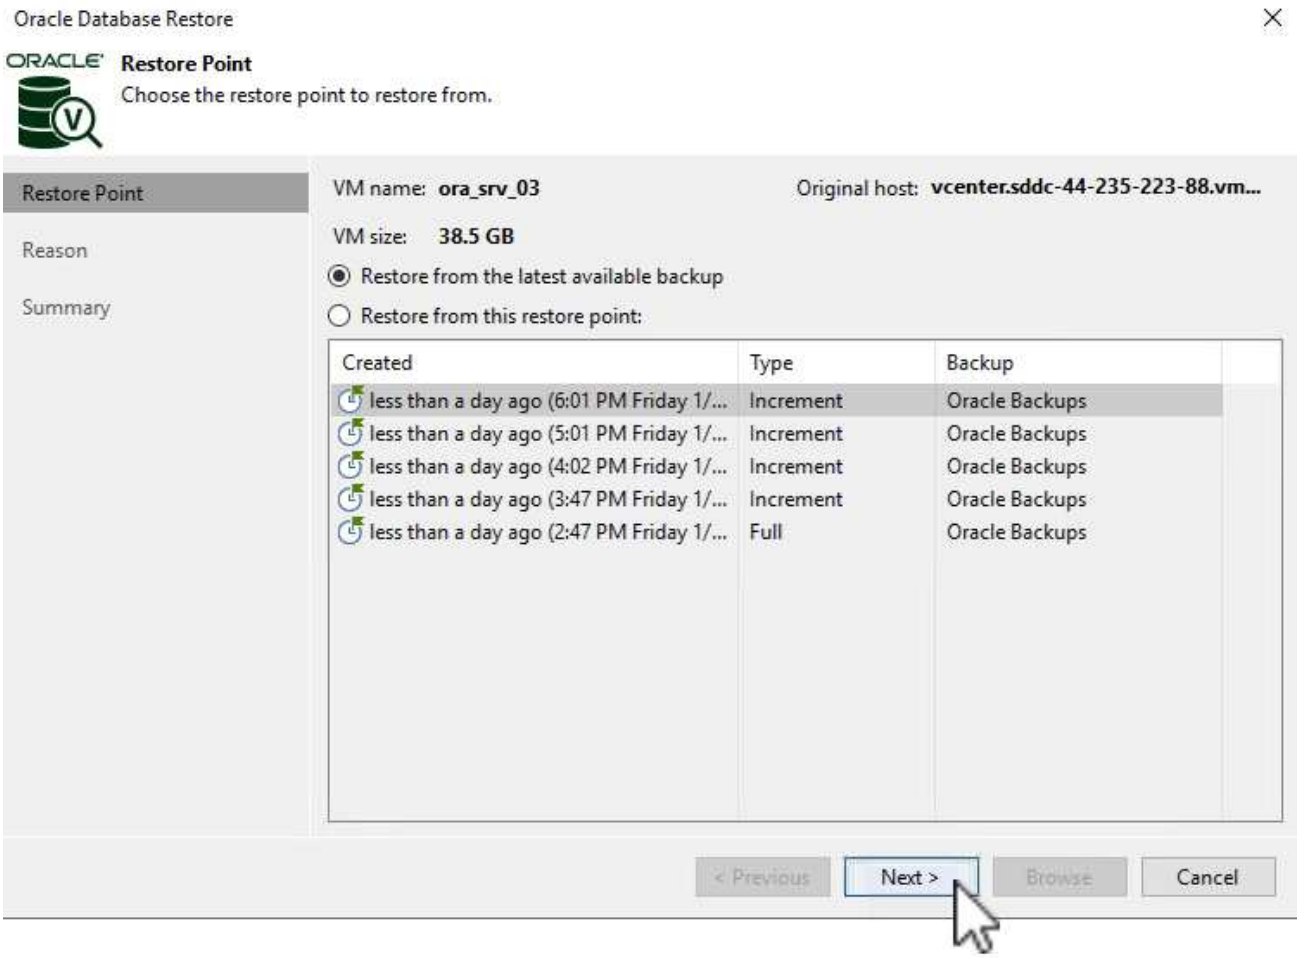
▼ **Restore Oracle database with Veeam Explorer**

In this section an Oracle database restore to a different server is covered using Veeam Explorer.

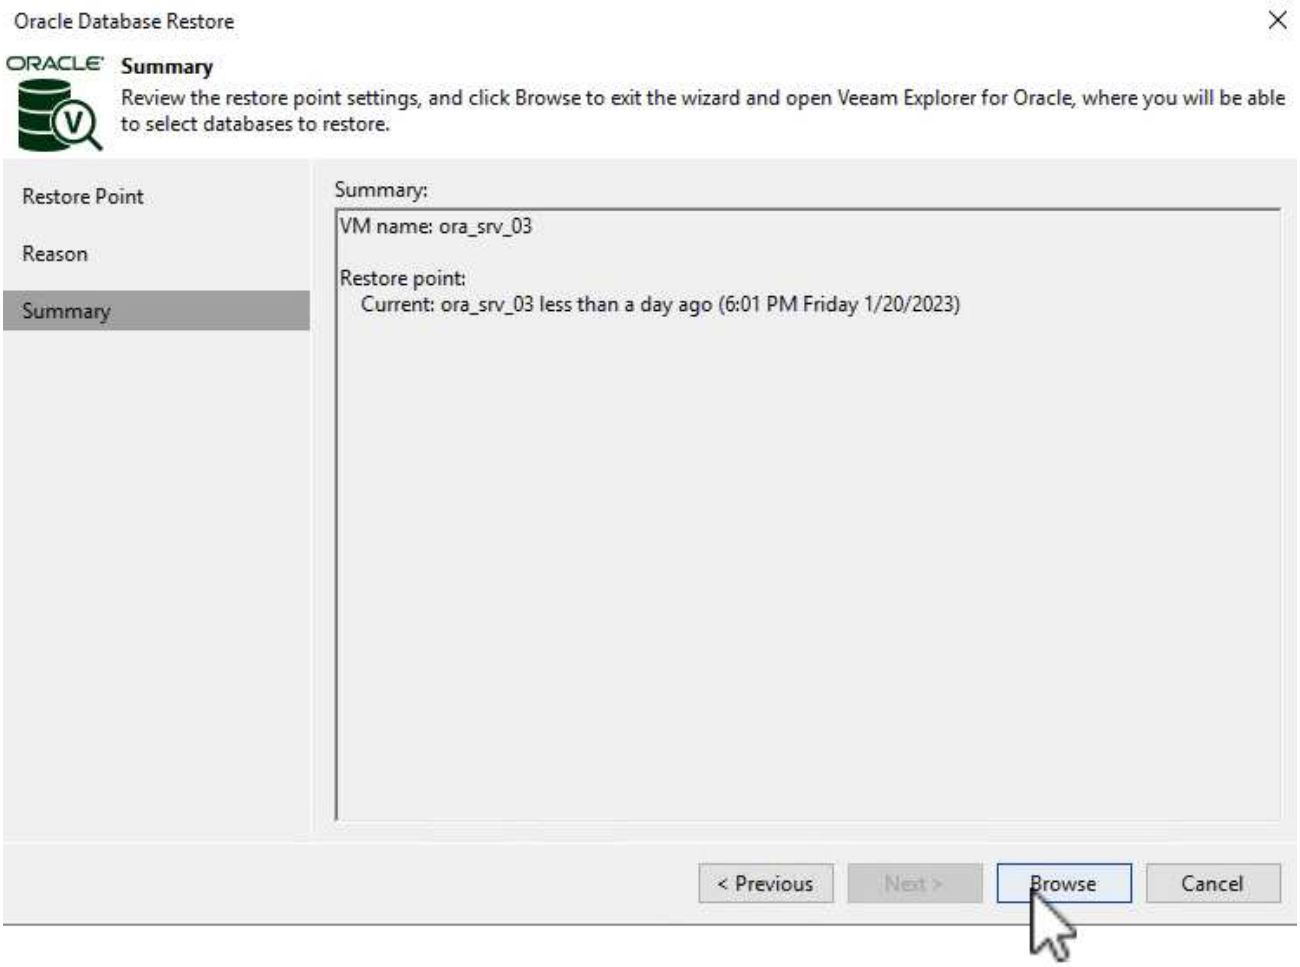
1. In the Veeam Backup and Replication console, navigate to the list of Oracle backups, right click on a server and select **Restore application items** and then **Oracle databases…**.



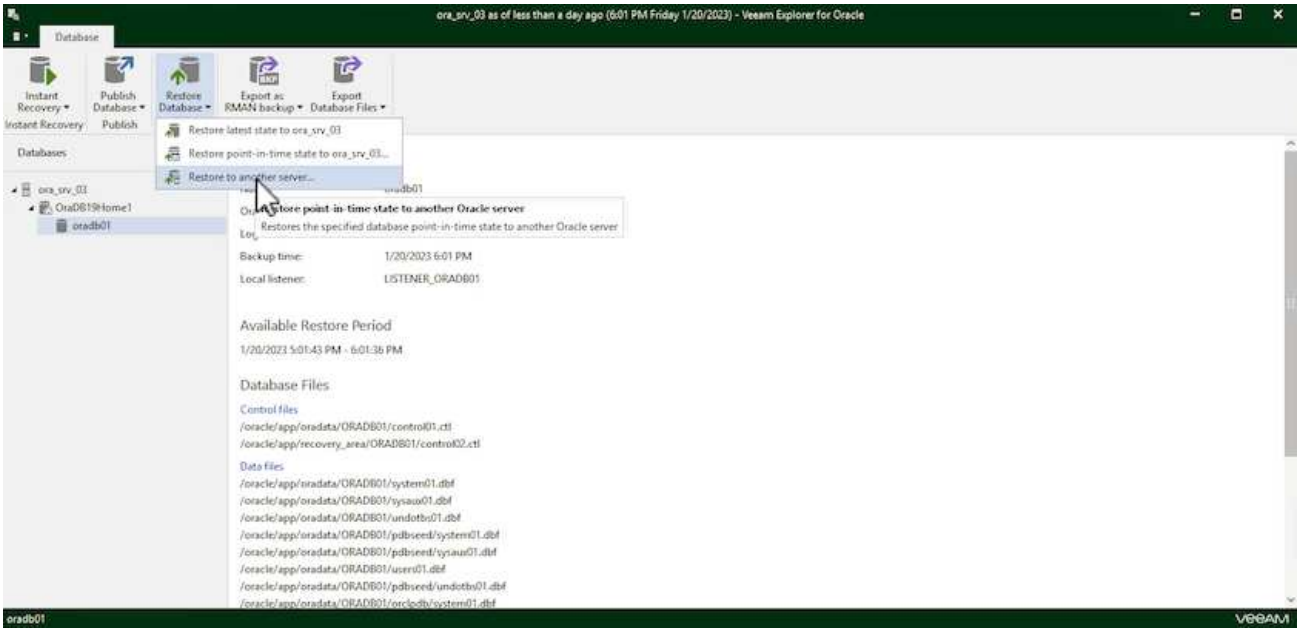
2. In the Oracle Database Restore Wizard select a restore point from the list and click on **Next**.



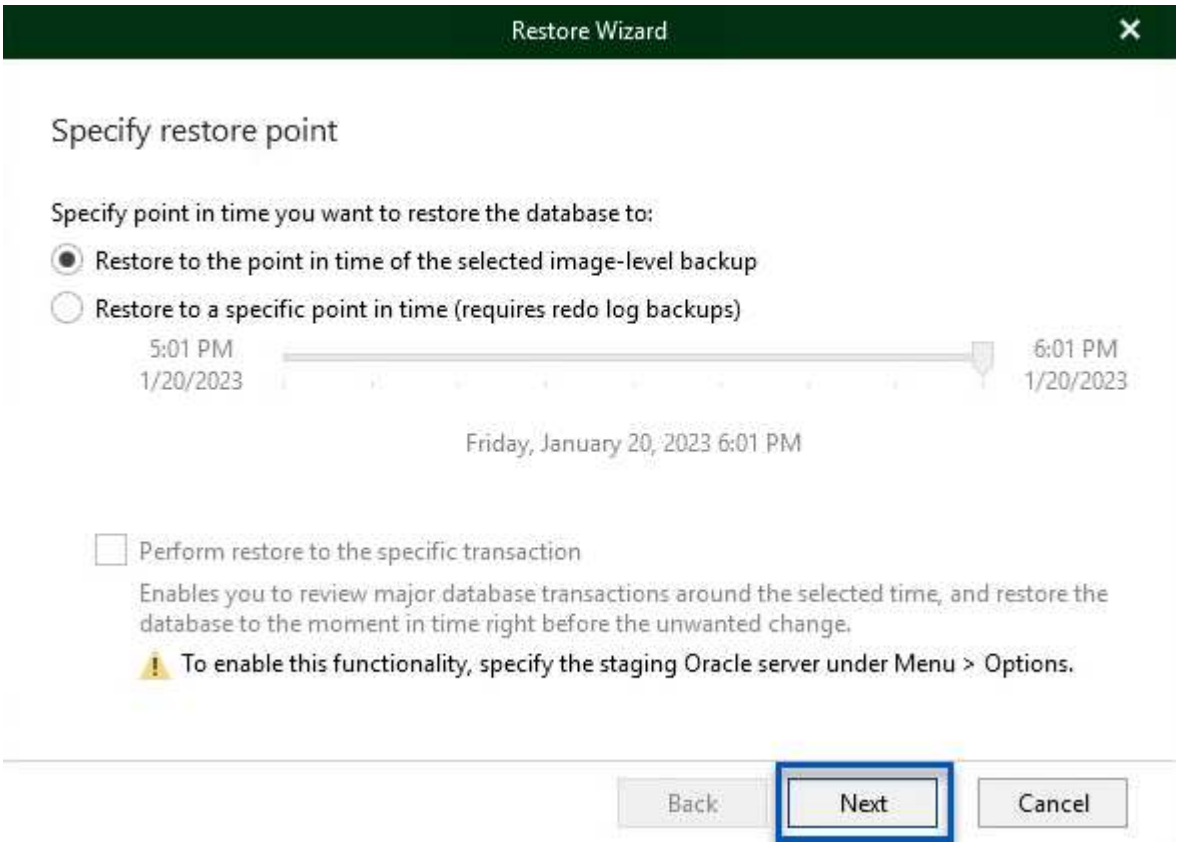
3. Enter a **Restore reason** if desired and then, on the Summary page, click on the **Browse** button to launch Veeam Explorer for Oracle.



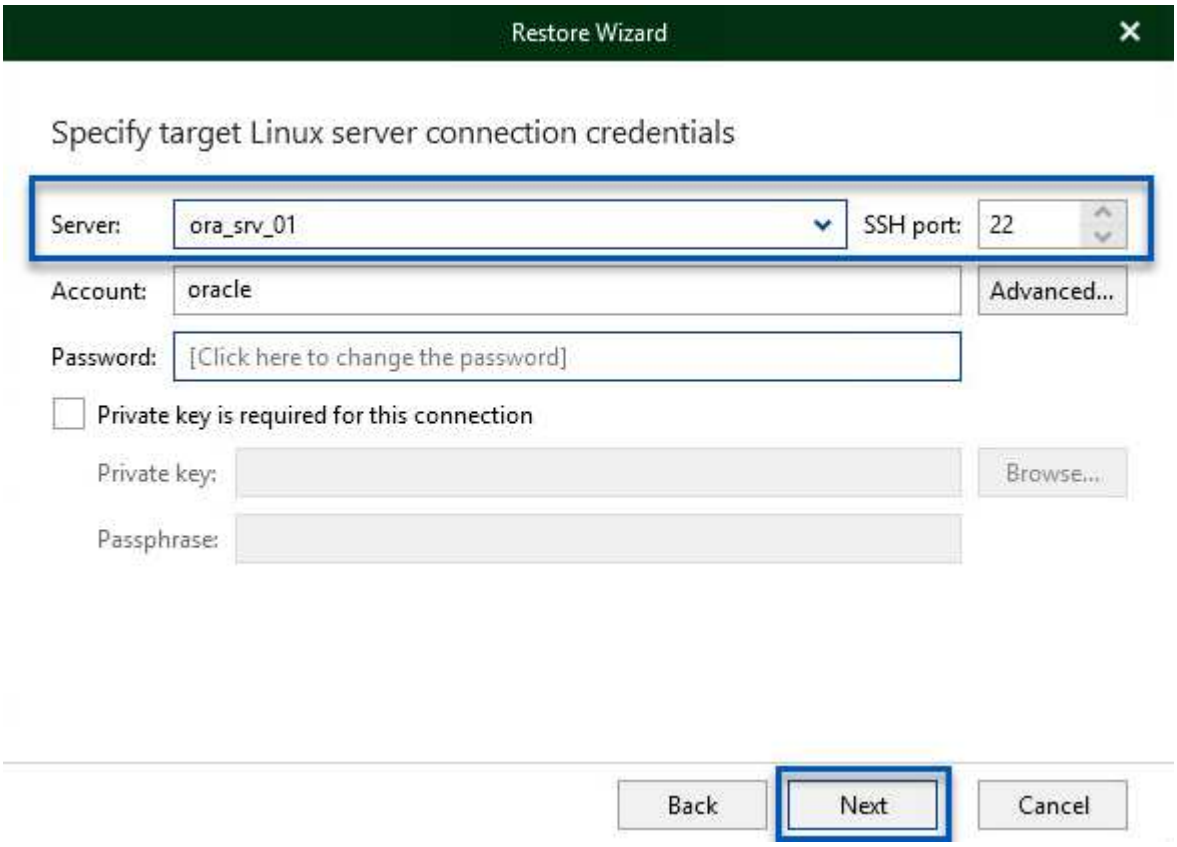
4. In Veeam Explorer expand the list of database instances, click on the database to be restored and then from the **Restore Database** drop-down menu at the top select **Restore to another server…**.



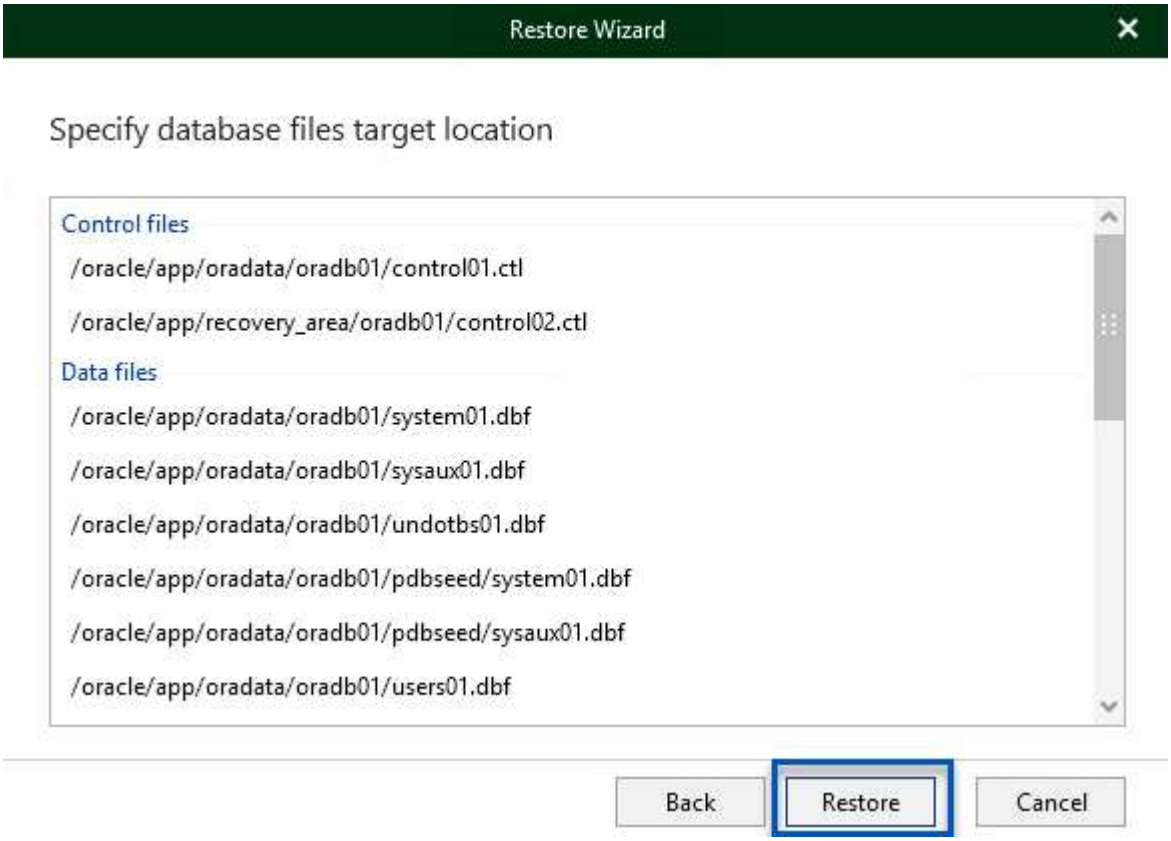
5. In the Restore Wizard specify the restore point to restore from and click **Next**.



6. Specify the target server the database will be restored to and the account credentials and click **Next**.



7. Finally, specify the database files target location and click the **Restore** button to start the restore process.

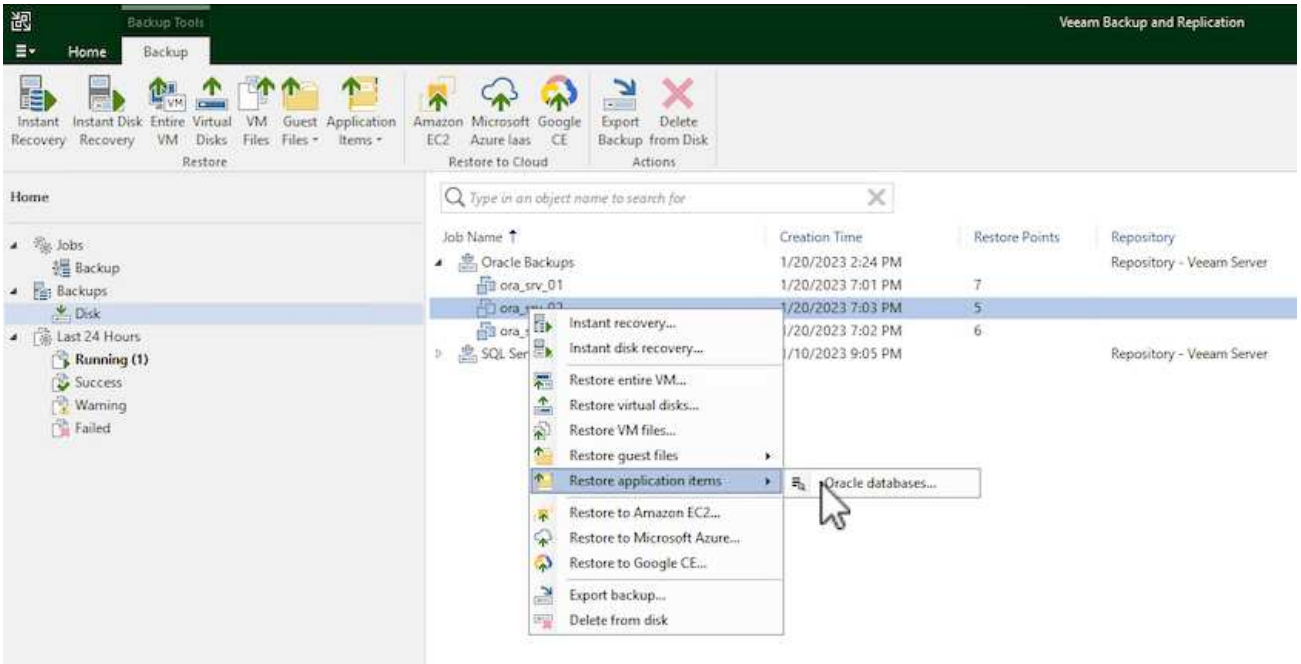


8. Once the database recovery is complete check that the Oracle database starts properly on the server.

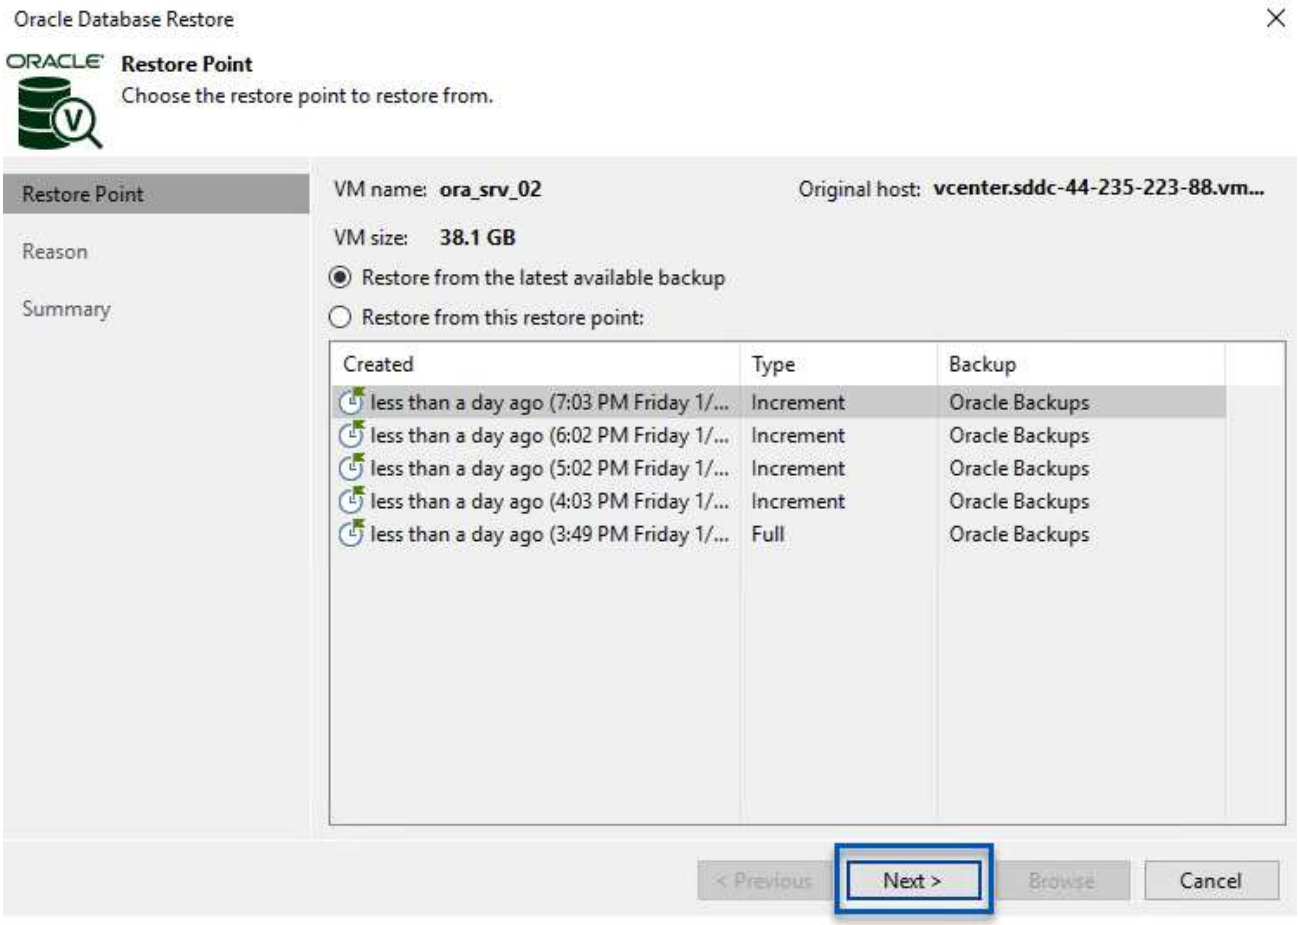
▼ **Publish Oracle database to alternate server**

In this section a database is published to an alternate server for fast access without launching a full restore.

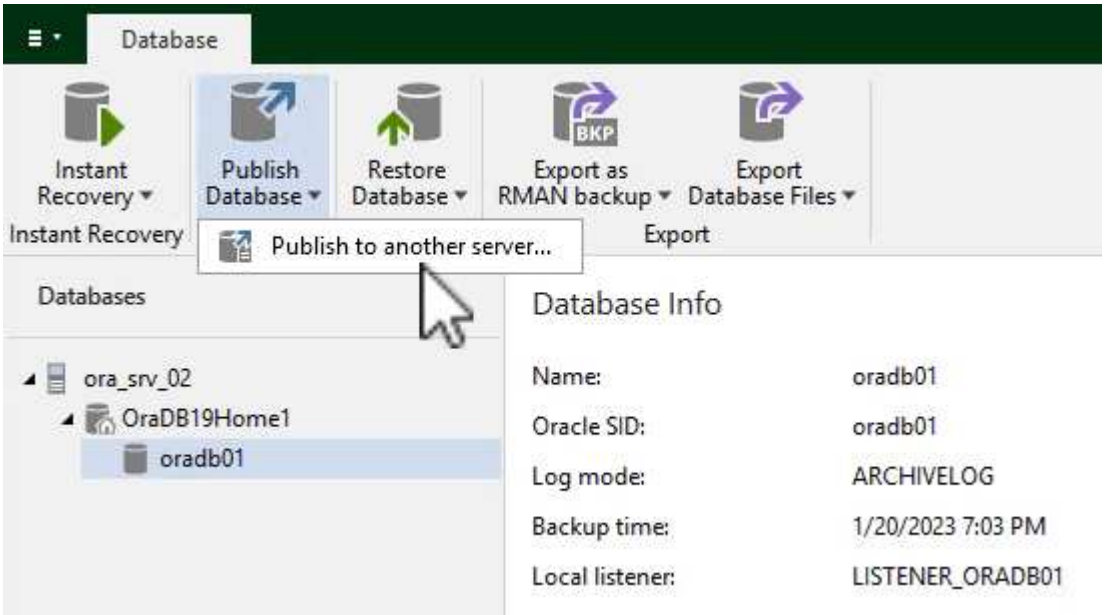
1. In the Veeam Backup and Replication console, navigate to the list of Oracle backups, right click on a server and select **Restore application items** and then **Oracle databases…**.



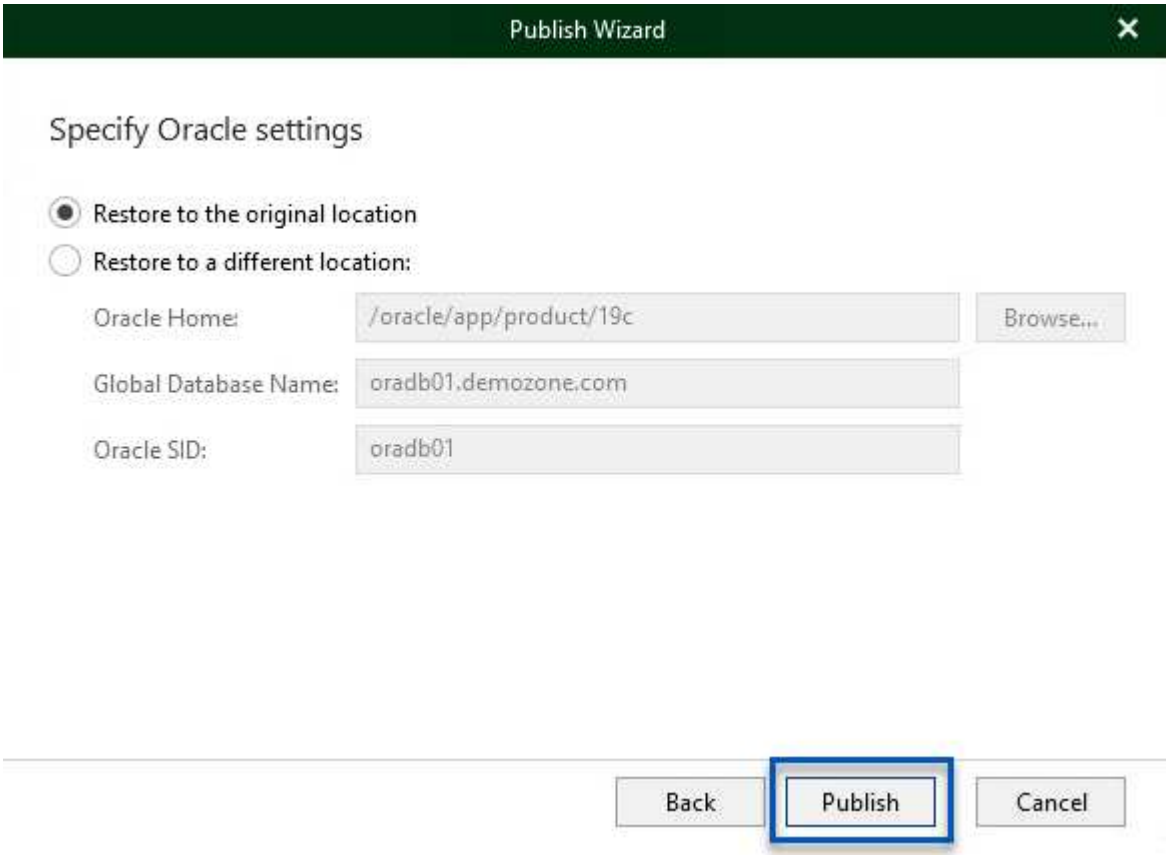
2. In the Oracle Database Restore Wizard select a restore point from the list and click on **Next**.



3. Enter a **Restore reason** if desired and then, on the Summary page, click on the **Browse** button to launch Veeam Explorer for Oracle.
4. In Veeam Explorer expand the list of database instances, click on the database to be restored and then from the **Publish Database** drop-down menu at the top select **Publish to another server…**.



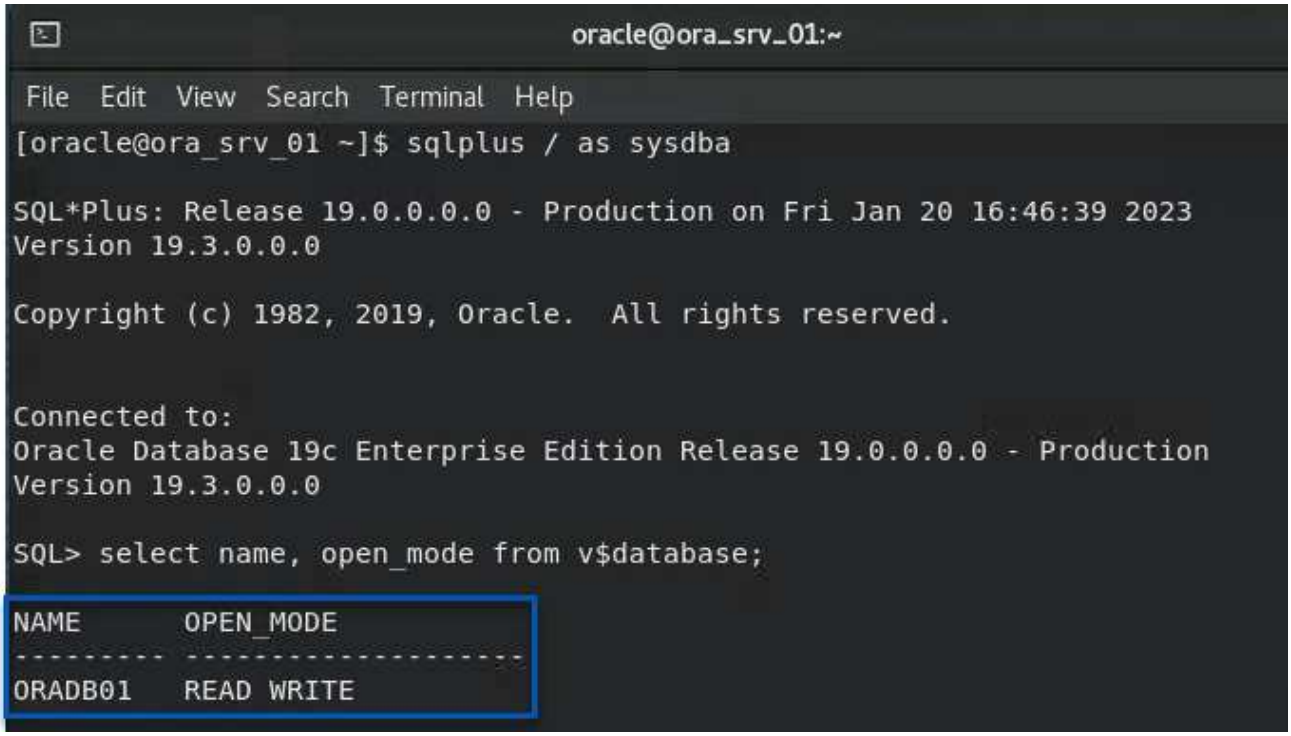
5. In the Publish wizard, specify the restore point at which to publish the database from and click **Next**.
6. Finally, specify the target linux file system location and click on **Publish** to begin the restore process.



7. Once the publish has completed log into the target server and run the following commands to ensure the database is running:

```
oracle@ora_srv_01> sqlplus / as sysdba
```

```
SQL> select name, open_mode from v$database;
```



Conclusion

VMware Cloud is a powerful platform for running business-critical applications and storing sensitive data. A secure data protection solution is essential for businesses that rely on VMware Cloud to ensure business continuity and help protect against cyber threats and data loss. By choosing a reliable and robust data protection solution, businesses can be confident that their critical data is safe and secure, no matter what.

The use case presented in this documentation focuses on proven data protection technologies that highlight the integration between NetApp, VMware, and Veeam. FSx ONTAP is supported as supplemental NFS datastores for VMware Cloud in AWS and is used for all virtual machine and application data. Veeam Backup & Replication is a comprehensive data protection solution designed to help businesses improve, automate, and streamline their backup and recovery processes. Veeam is used in conjunction with iSCSI backup target volumes, hosted on FSx ONTAP, to provide a secure and easy to manage data protection solution for application data residing in VMware Cloud.

Additional Information

To learn more about the technologies presented in this solution refer to the following additional information.

- FSx ONTAP User Guide
- Veeam Help Center Technical Documentation
- VMware Cloud on AWS Support. Considerations and Limitations

TR-4955: Disaster Recovery with FSx ONTAP and VMC (AWS VMware Cloud)

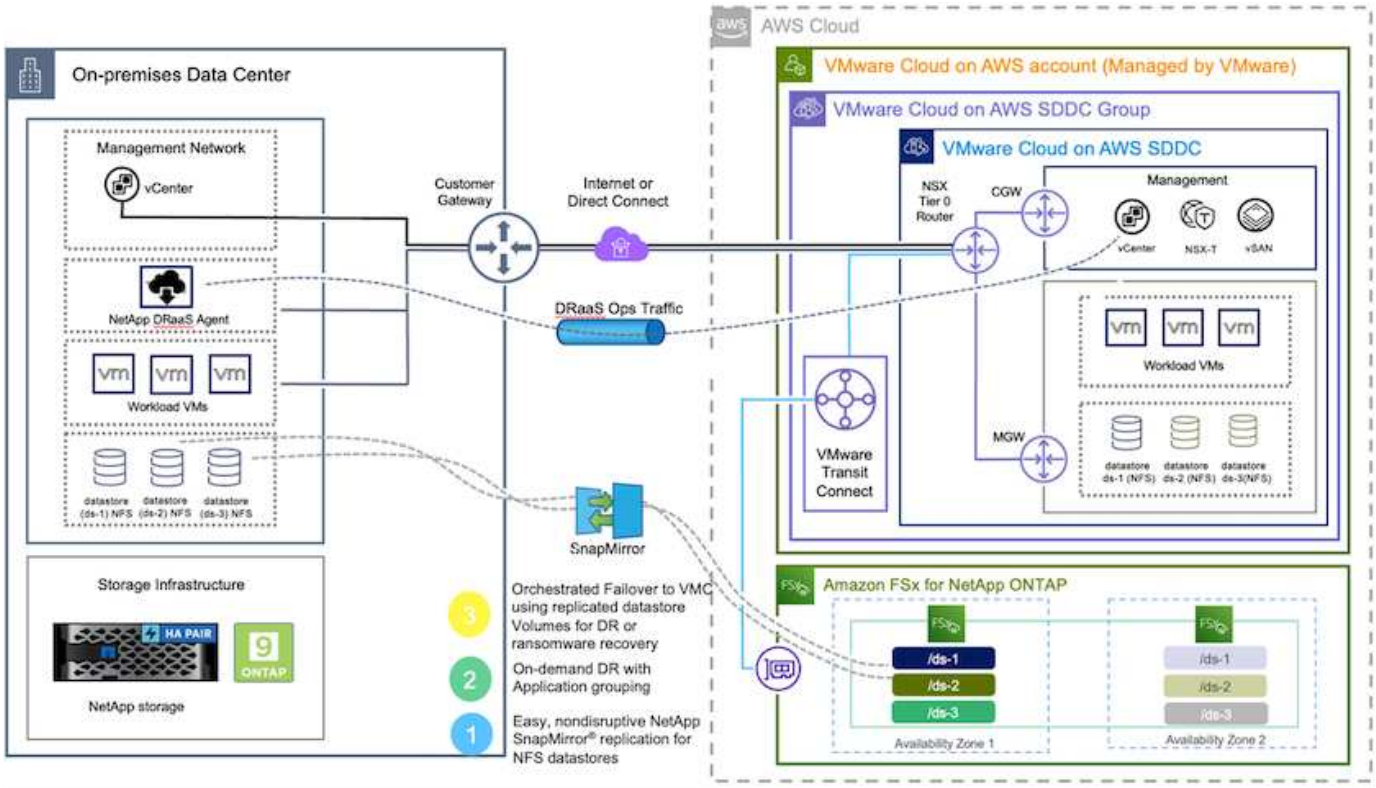
Disaster Recovery Orchestrator (DRO; a scripted solution with UI) can be used to seamlessly recover workloads replicated from on-premises to FSx ONTAP. DRO automates the recovery from the SnapMirror level, through VM registration to VMC, to network mappings directly on NSX-T. This feature is included with all VMC environments.

Niyaz Mohamed, NetApp

Overview

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). With NetApp SnapMirror technology, on-premises VMware workloads can be replicated to FSx ONTAP running in AWS.

Disaster Recovery Orchestrator (DRO; a scripted solution with UI) can be used to seamlessly recover workloads replicated from on-premises to FSx ONTAP. DRO automates the recovery from the SnapMirror level, through VM registration to VMC, to network mappings directly on NSX-T. This feature is included with all VMC environments.



Getting started

Deploy and configure VMware Cloud on AWS

VMware Cloud on AWS provides a cloud-native experience for VMware-based workloads in the AWS ecosystem. Each VMware Software-Defined Data Center (SDDC) runs in an Amazon Virtual Private Cloud (VPC) and provides a full VMware stack (including vCenter Server), NSX-T software-defined networking, vSAN software-defined storage, and one or more ESXi hosts that provide compute and storage resources to the workloads. To configure a VMC environment on AWS, follow the steps at this link. A pilot-light cluster can also be used for DR purposes.

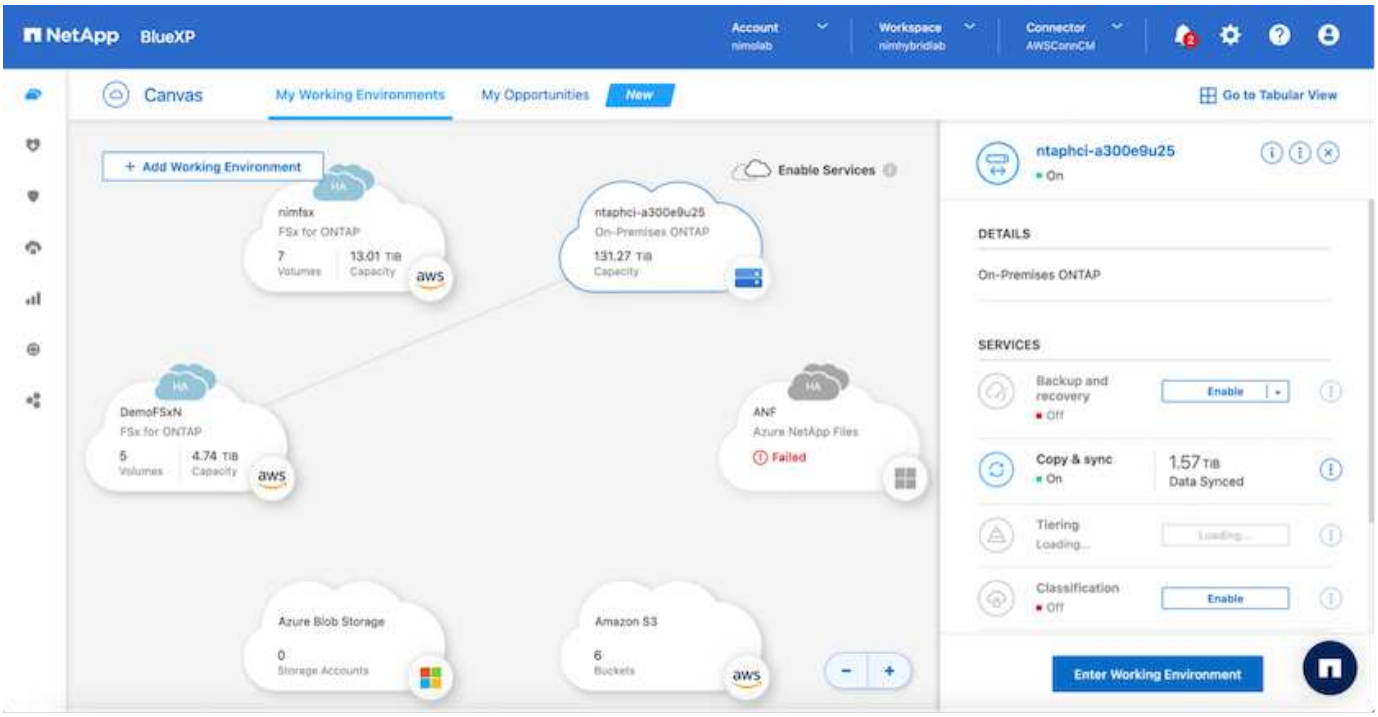
> In the initial release, DRO supports an existing pilot-light cluster. On-demand SDDC creation will be available in an upcoming release.

Provision and configure FSx ONTAP

Amazon FSx ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on the popular NetApp ONTAP file system. Follow the steps at this link to provision and configure FSx ONTAP.

Deploy and configure SnapMirror to FSx ONTAP

The next step is to use NetApp Console and discover the provisioned FSx ONTAP on AWS instance and replicate the desired datastore volumes from an on-premises environment to FSx ONTAP with the appropriate frequency and NetApp Snapshot copy retention:



Follow the steps in this link to configure DRO. You can also use the NetApp ONTAP CLI to schedule replication following this link.

A SnapMirror relationship is a prerequisite and must be created beforehand.

DRO installation

To get started with DRO, use the Ubuntu operating system on a designated EC2 instance or virtual machine to make sure you meet the prerequisites. Then install the package.

Prerequisites

- Make sure that connectivity to the source and destination vCenter and storage systems exists.
- DNS resolution should be in place if you are using DNS names. Otherwise, you should use IP addresses for the vCenter and storage systems.
- Create a user with root permissions. You can also use sudo with an EC2 instance.

OS requirements

- Ubuntu 20.04 (LTS) with minimum of 2GB and 4 vCPUs
- The following packages must be installed on the designated agent VM:
  - Docker
  - Docker-compose

- Jq

Change permissions on `docker.sock`: `sudo chmod 666 /var/run/docker.sock`.

The `deploy.sh` script executes all the required prerequisites.

Install the package

1. Download the installation package on the designated virtual machine:

```
git clone https://github.com/NetApp/DRO-AWS.git
```

The agent can be installed on-premises or within an AWS VPC.

2. Unzip the package, run the deployment script, and enter the host IP (for example, 10.10.10.10).

```
tar xvf DRO-prereq.tar
```

3. Navigate to the directory and run the deploy script as follows:

```
sudo sh deploy.sh
```

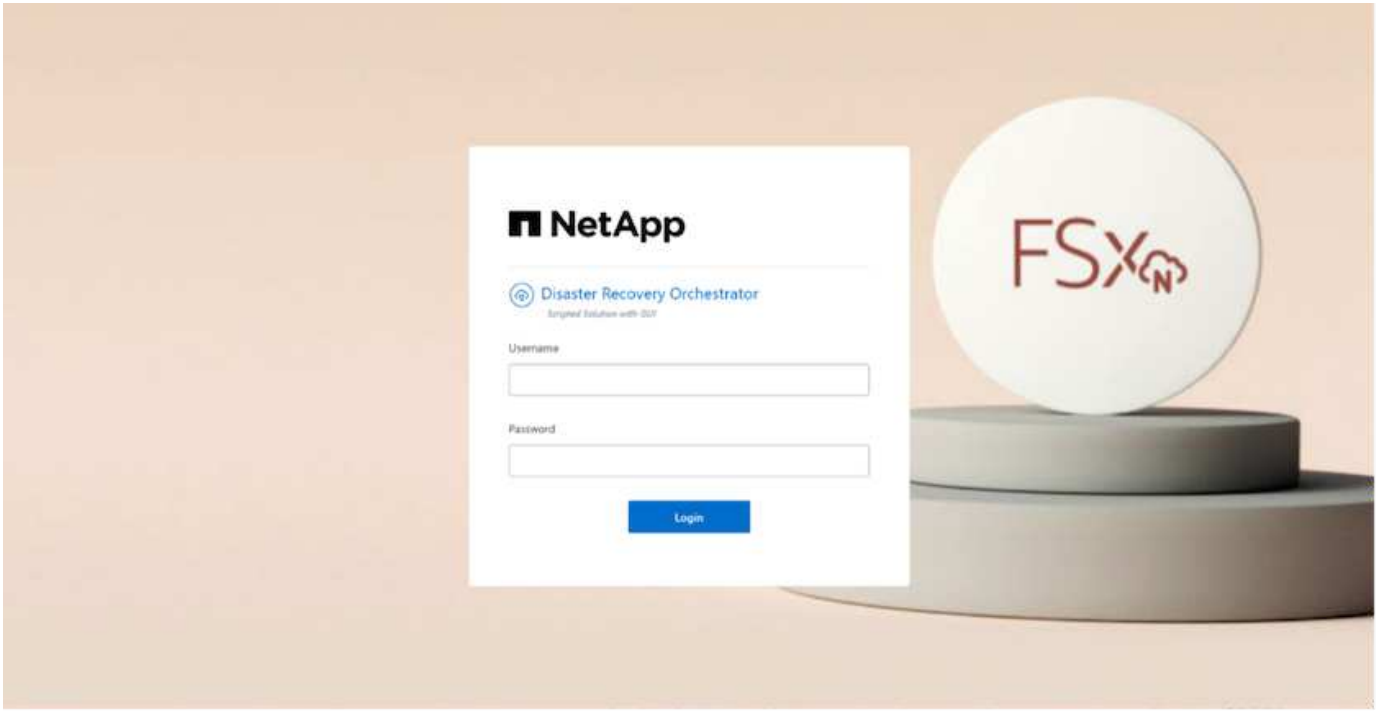
4. Access the UI using:

```
https://<host-ip-address>
```

with the following default credentials:

```
Username: admin
Password: admin
```

The password can be changed using the "Change Password" option.



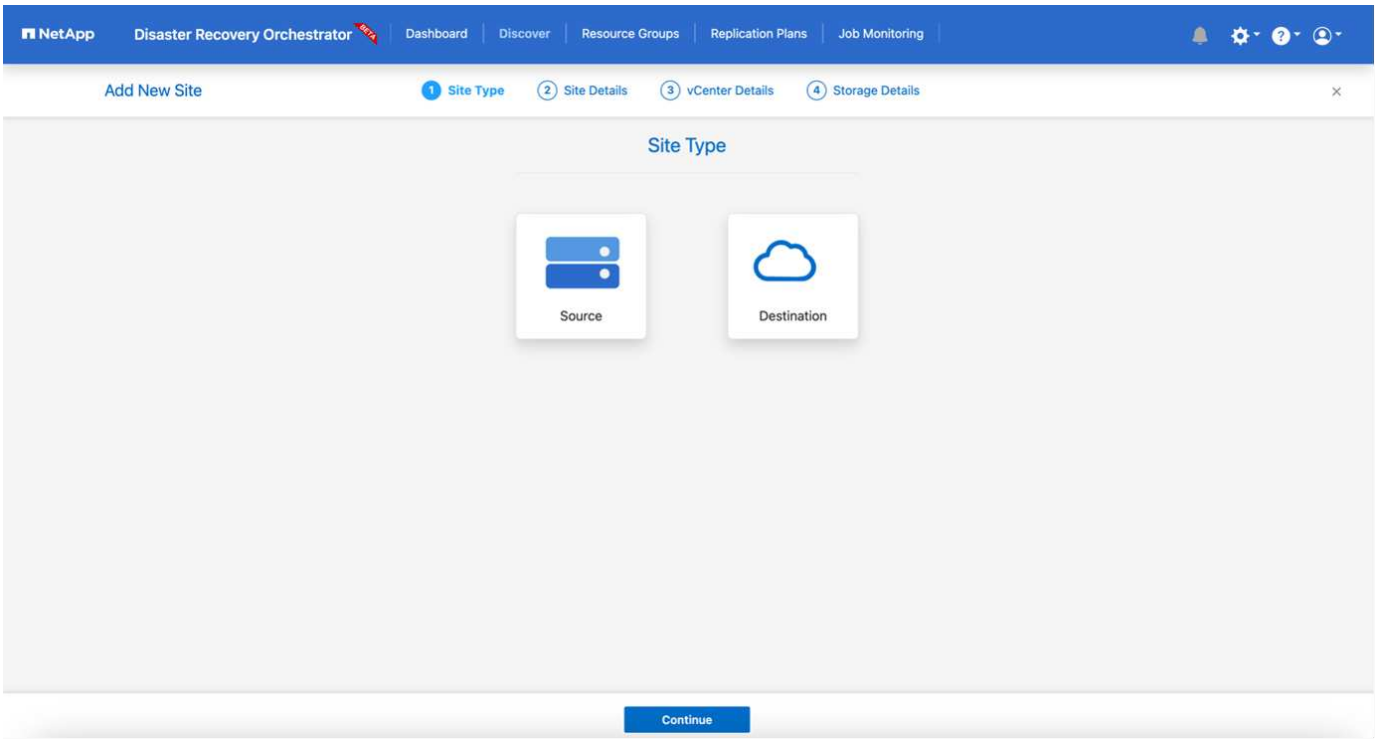
DRO configuration

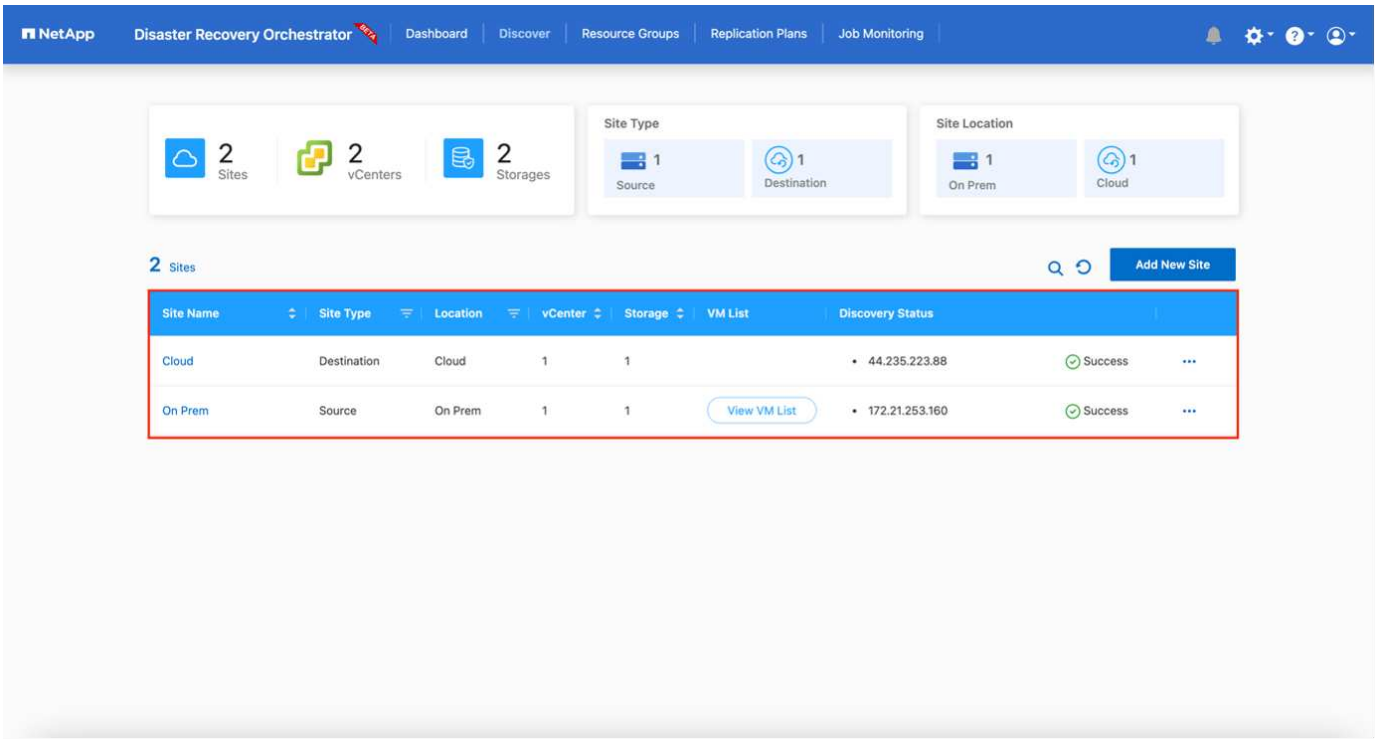
After FSx ONTAP and VMC have been configured properly, you can begin configuring DRO to automate the recovery of on-premises workloads to VMC by using the read-only SnapMirror copies on FSx ONTAP.

NetApp recommends deploying the DRO agent in AWS and also to the same VPC where FSx ONTAP is deployed (it can be peer connected too), so that the DRO agent can communicate through the network with your on-premises components as well as with the FSx ONTAP and VMC resources.

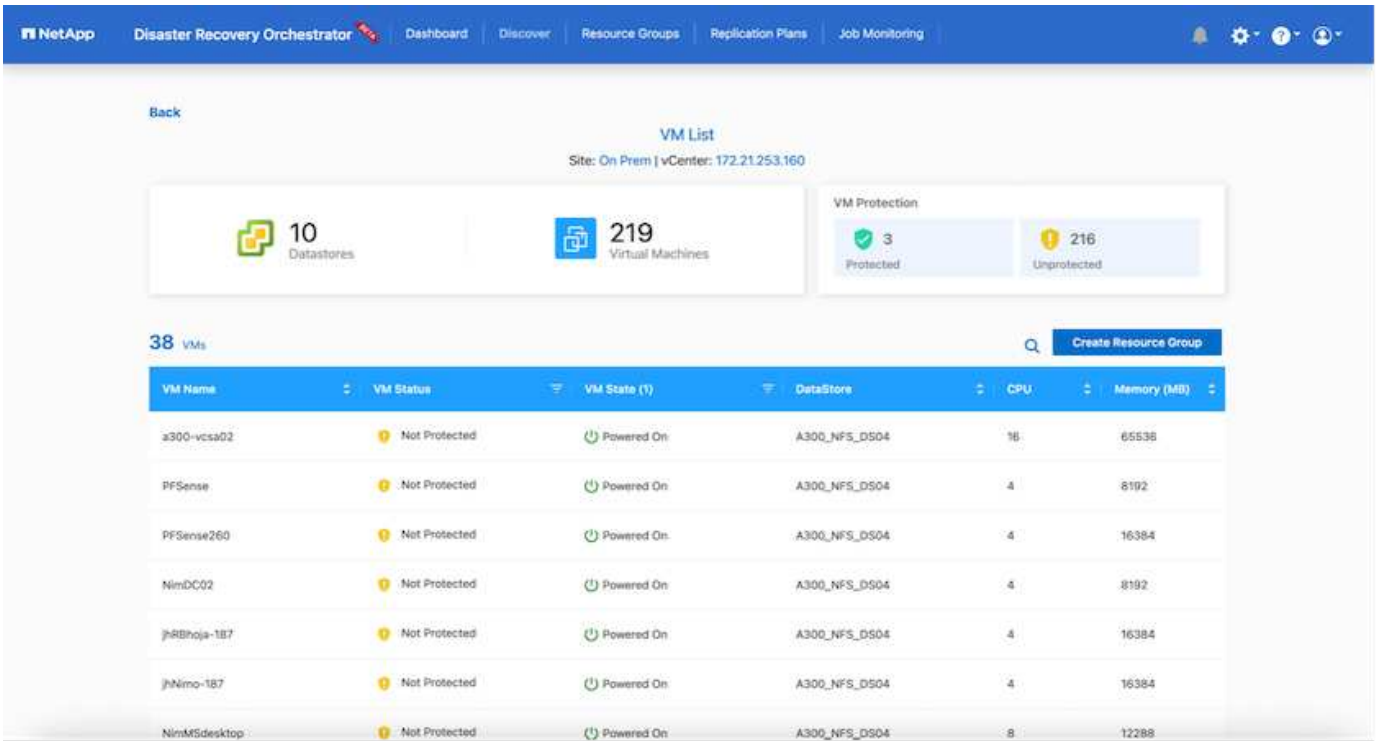
The first step is to discover and add the on-premises and cloud resources (both vCenter and storage) to DRO. Open DRO in a supported browser and use the default username and password (admin/admin) and Add Sites. Sites can also be added using the Discover option. Add the following platforms:

- On-premises
  - On-premises vCenter
  - ONTAP storage system
- Cloud
  - VMC vCenter
  - FSx ONTAP





Once added, DRO performs automatic discovery and displays the VMs that have corresponding SnapMirror replicas from the source storage to FSx ONTAP. DRO automatically detects the networks and portgroups used by the VMs and populates them.



The next step is to group the required VMs into functional groups to serve as resource groups.

Resource groupings

After the platforms have been added, you can group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.

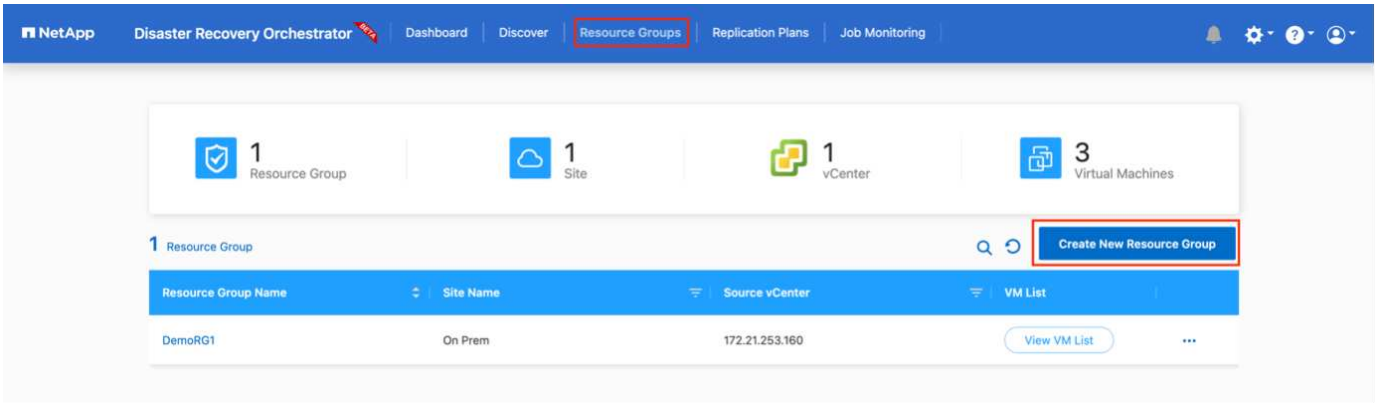
To start creating resource groups, complete the following steps:

1. Access **Resource Groups**, and click **Create New Resource Group**.
2. Under **New resource group**, select the source site from the dropdown and click **Create**.
3. Provide **Resource Group Details** and click **Continue**.
4. Select the appropriate VMs using the search option.
5. Select the boot order and boot delay (secs) for the selected VMs. Set the order of the power-on sequence by selecting each VM and setting up the priority for it. Three is the default value for all VMs.

   Options are as follows:

   1 – The first virtual machine to power on
   3 – Default
   5 – The last virtual machine to power on

6. Click **Create Resource Group**.

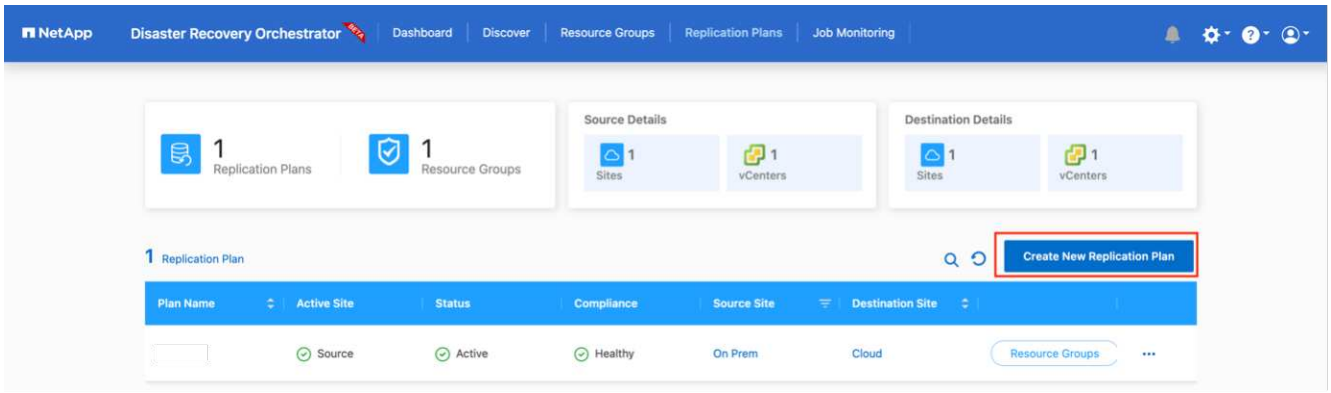


Replication plans

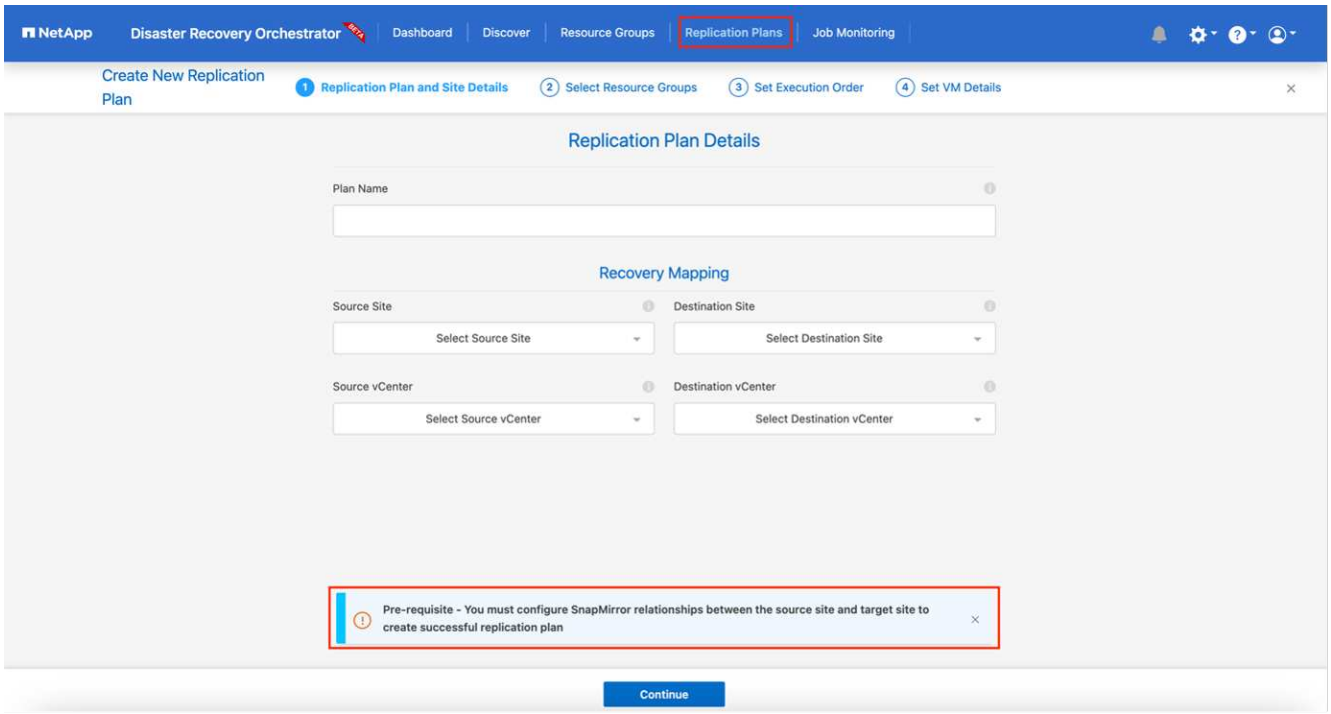
You need a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down and pick the resource groups to be included in this plan, along with the grouping of how applications should be restored and powered on (for example, domain controllers, then tier-1, then tier-2, and so on). Such plans are sometimes also called blueprints. To define the recovery plan, navigate to the **Replication Plan** tab and click **New Replication Plan**.

To start creating a replication plan, complete the following steps:

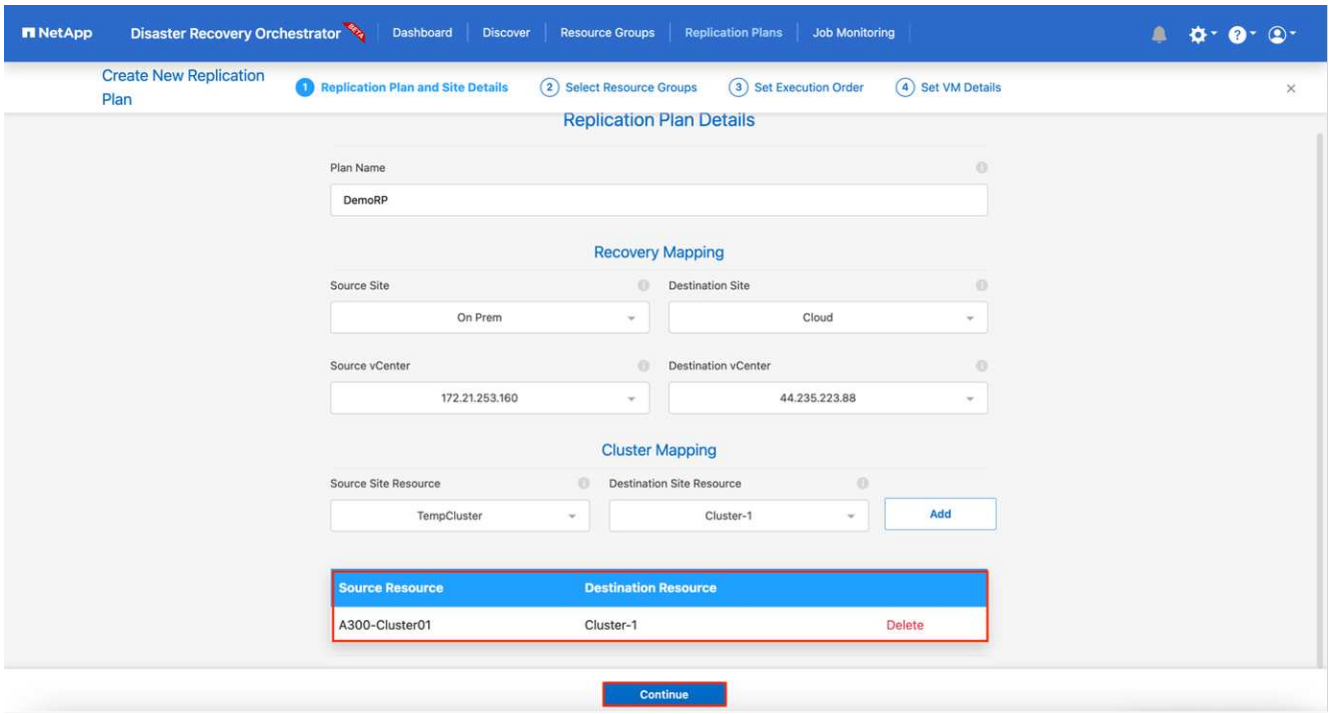
1. Access **Replication Plans**, and click **Create New Replication Plan**.



2. Under **New Replication Plan**, provide a name for the plan and add recovery mappings by selecting the source site, associated vCenter, destination site, and associated vCenter.

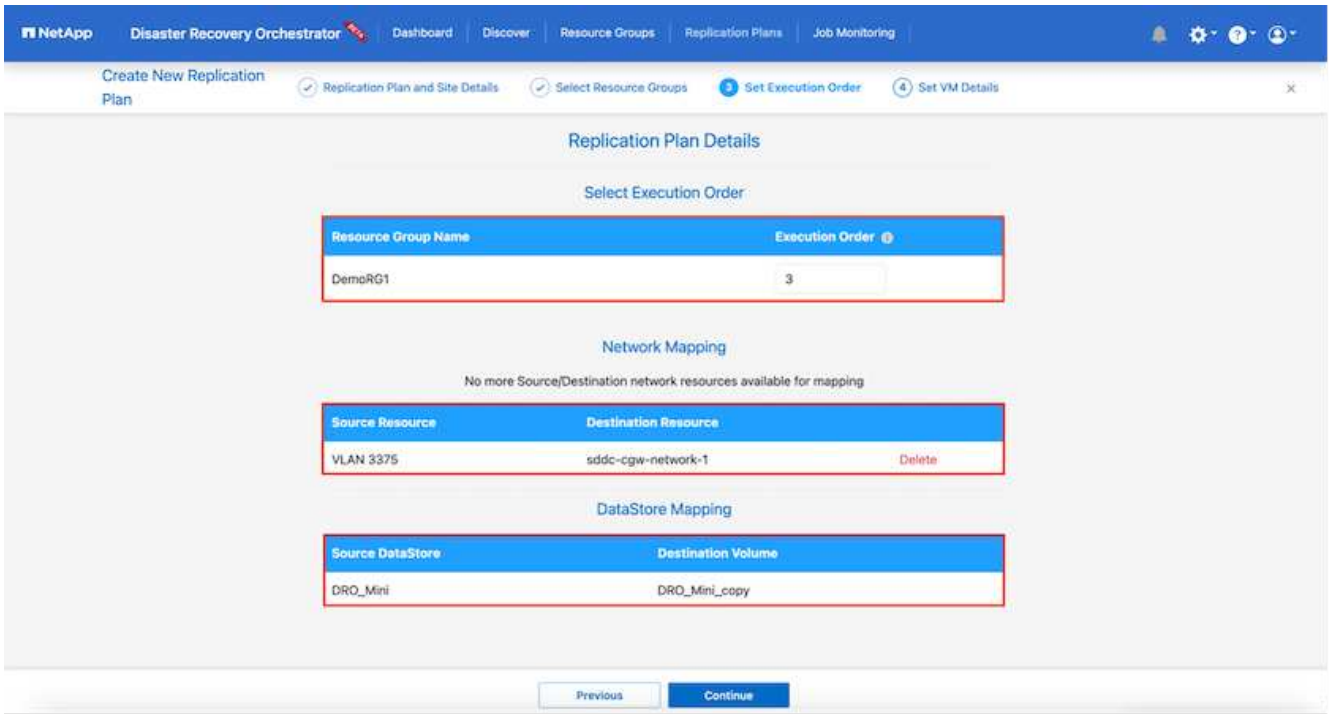


3. After Recovery mapping is completed, select the cluster mapping.

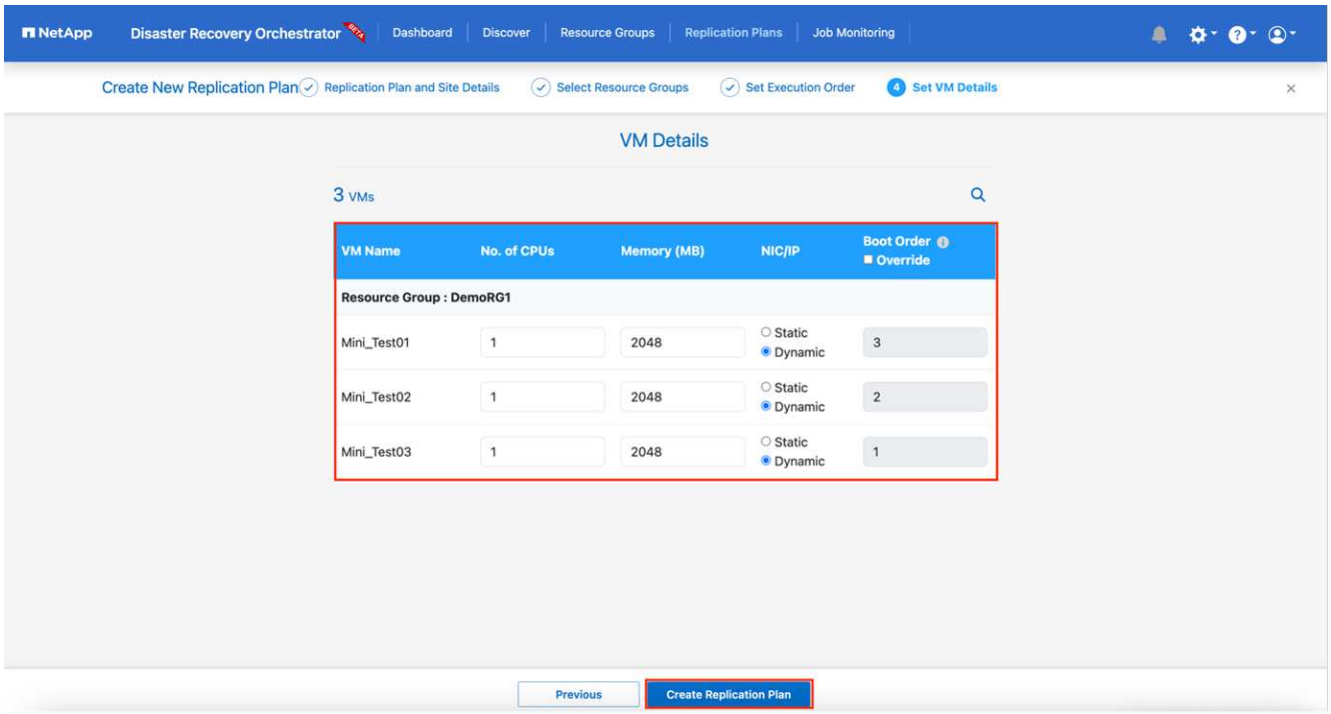


4. Select **Resource Group Details** and click **Continue**.
5. Set the execution order for the resource group. This option enables you to select the sequence of operations when multiple resource groups exist.
6. After you are done, select the network mapping to the appropriate segment. The segments should already be provisioned within VMC, so select the appropriate segment to map the VM.
7. Based on the selection of VMs, datastore mappings are automatically selected.

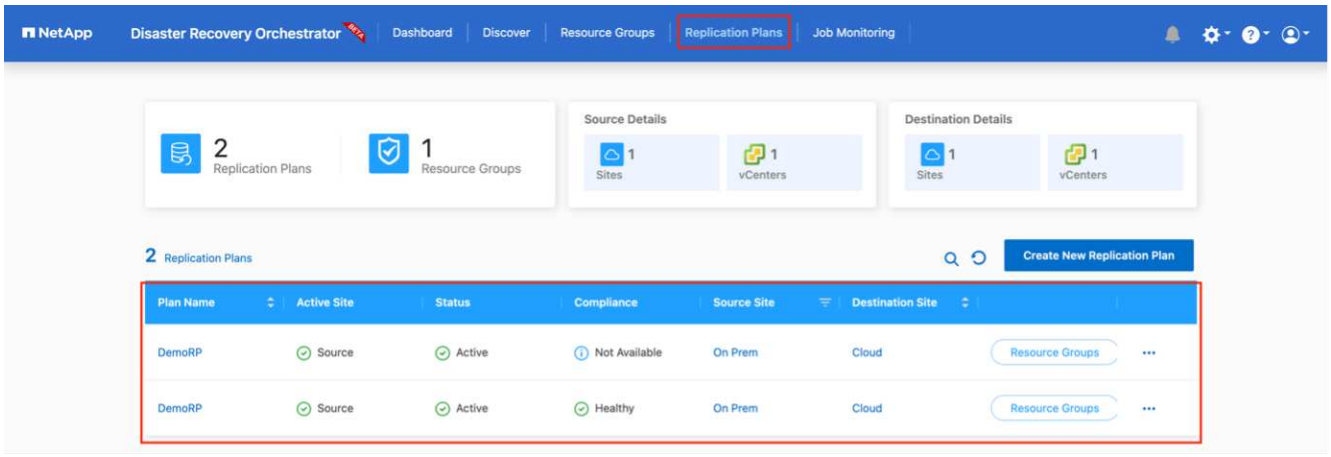
SnapMirror is at the volume level. Therefore, all VMs are replicated to the replication destination. Make sure to select all VMs that are part of the datastore. If they are not selected, only the VMs that are part of the replication plan are processed.



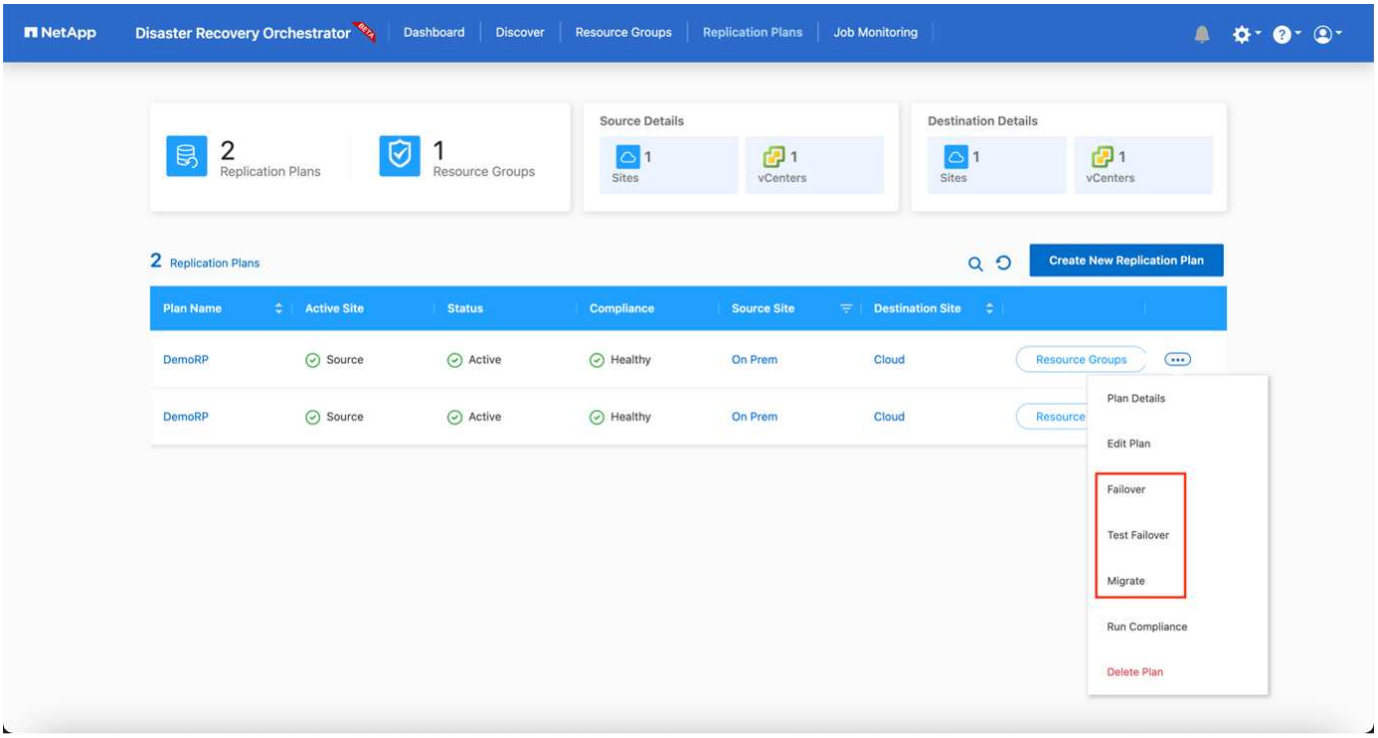
8. Under the VM details, you can optionally resize the VM's CPU and RAM parameters; this can be very helpful when recovering large environments to smaller target clusters or for conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, you can modify the boot order and boot delay (seconds) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if there are any changes required from those selected during the resource-group boot-order selection. By default, the boot order selected during resource-group selection is used; however, any modifications can be performed at this stage.

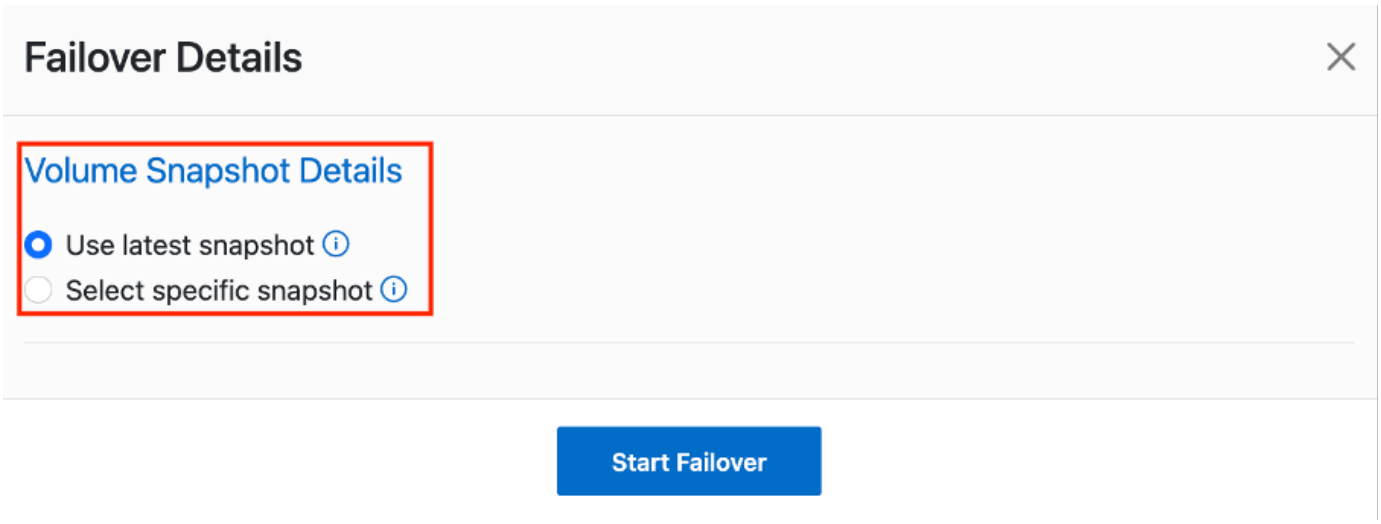


9. Click **Create Replication Plan**.

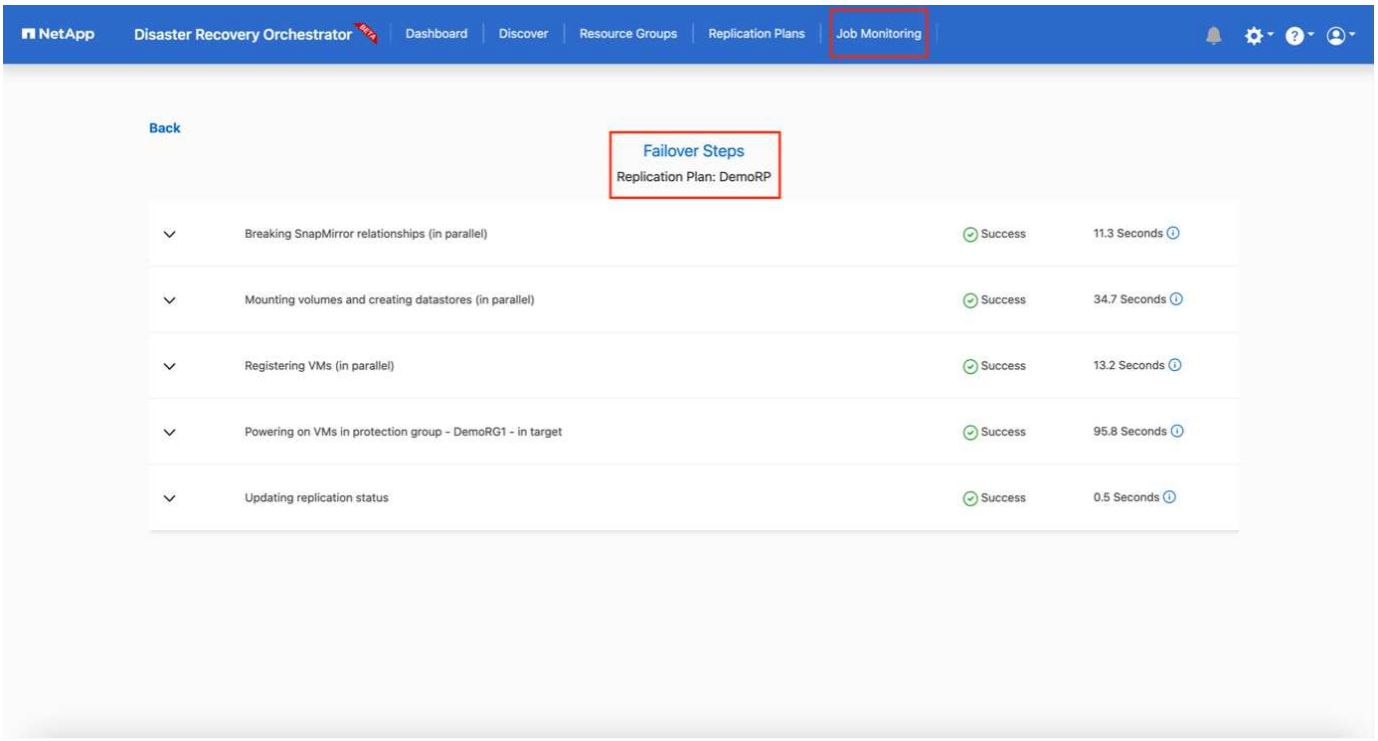


After the replication plan is created, the failover option, the test-failover option, or the migrate option can be exercised depending on the requirements. During the failover and test-failover options, the most recent SnapMirror Snapshot copy is used, or a specific Snapshot copy can be selected from a point-in-time Snapshot copy (per the retention policy of SnapMirror). The point-in-time option can be very helpful if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available points in time. To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test failover**.

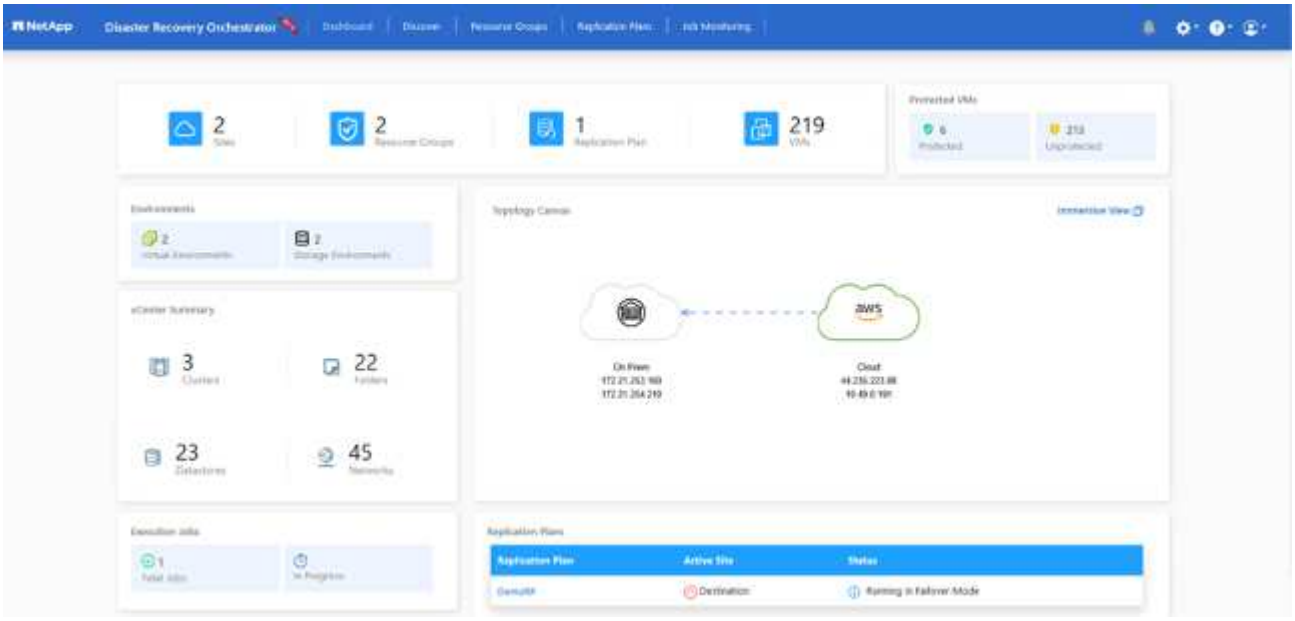




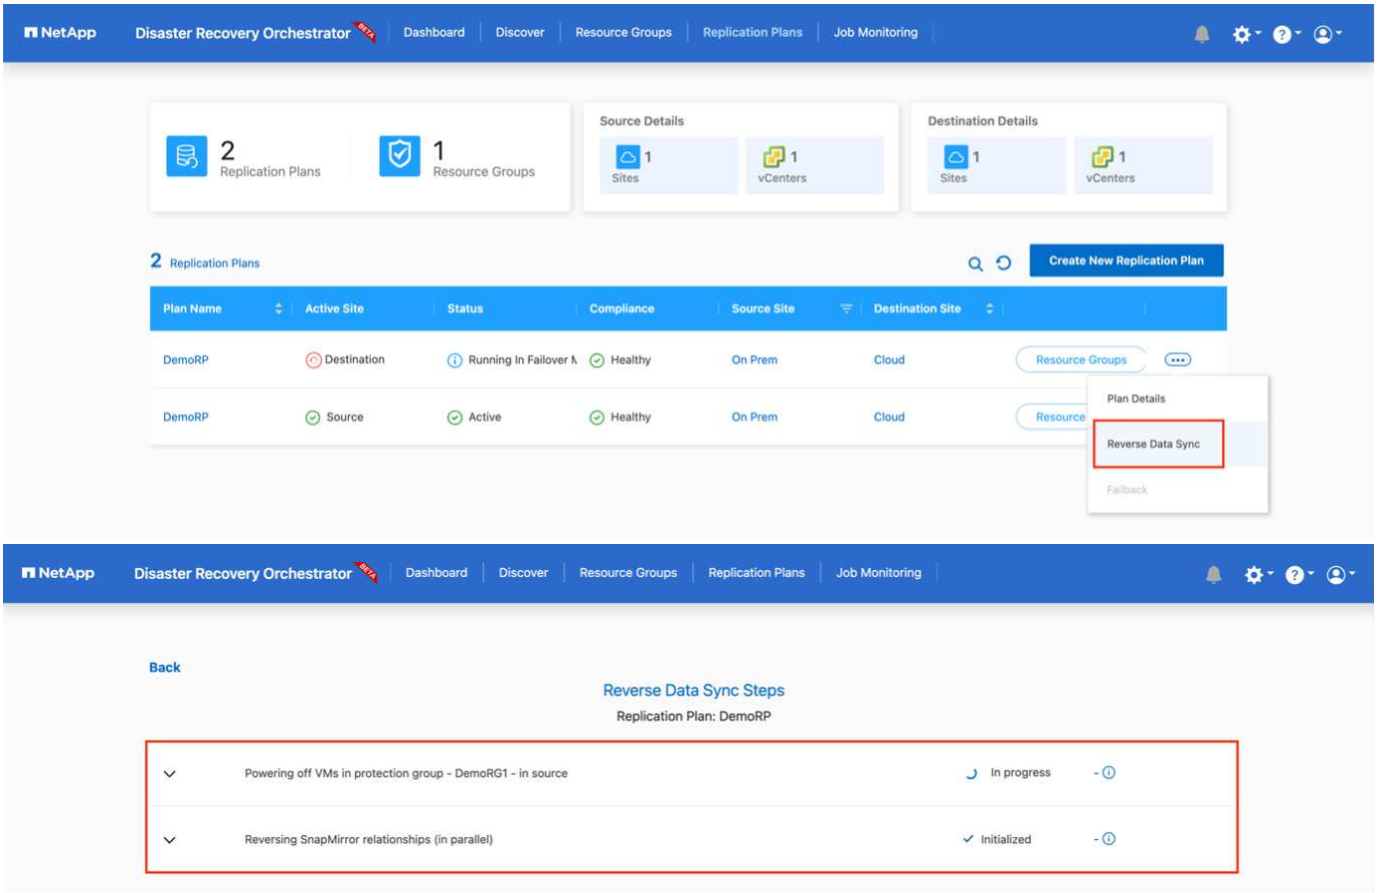
The replication plan can be monitored in the task menu:



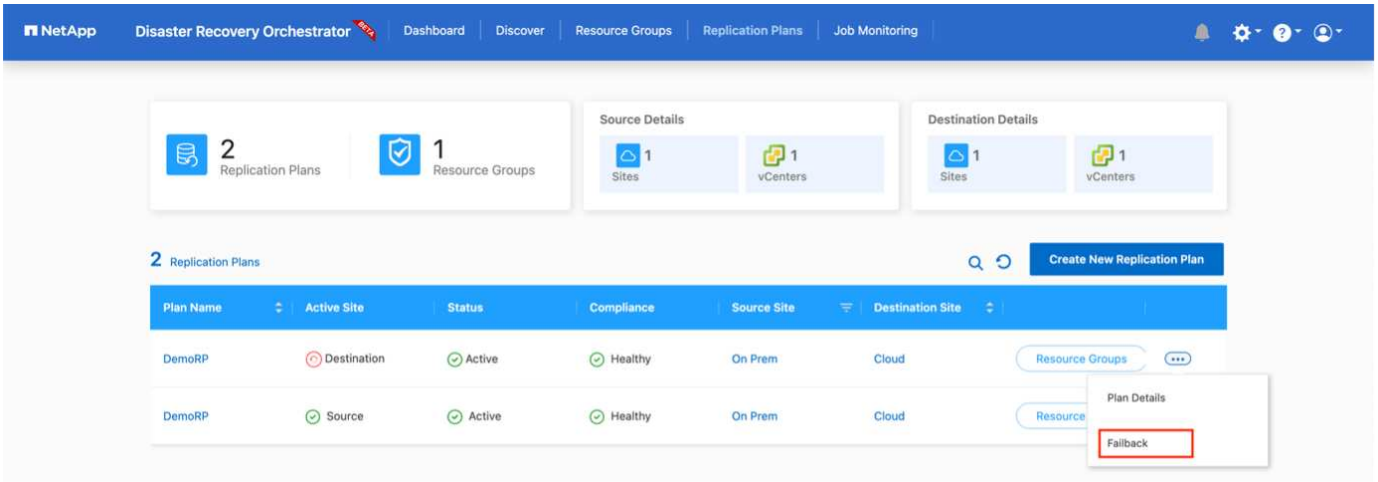
After failover is triggered, the recovered items can be seen in the VMC vCenter (VMs, networks, datastores). By default, the VMs are recovered to the Workload folder.

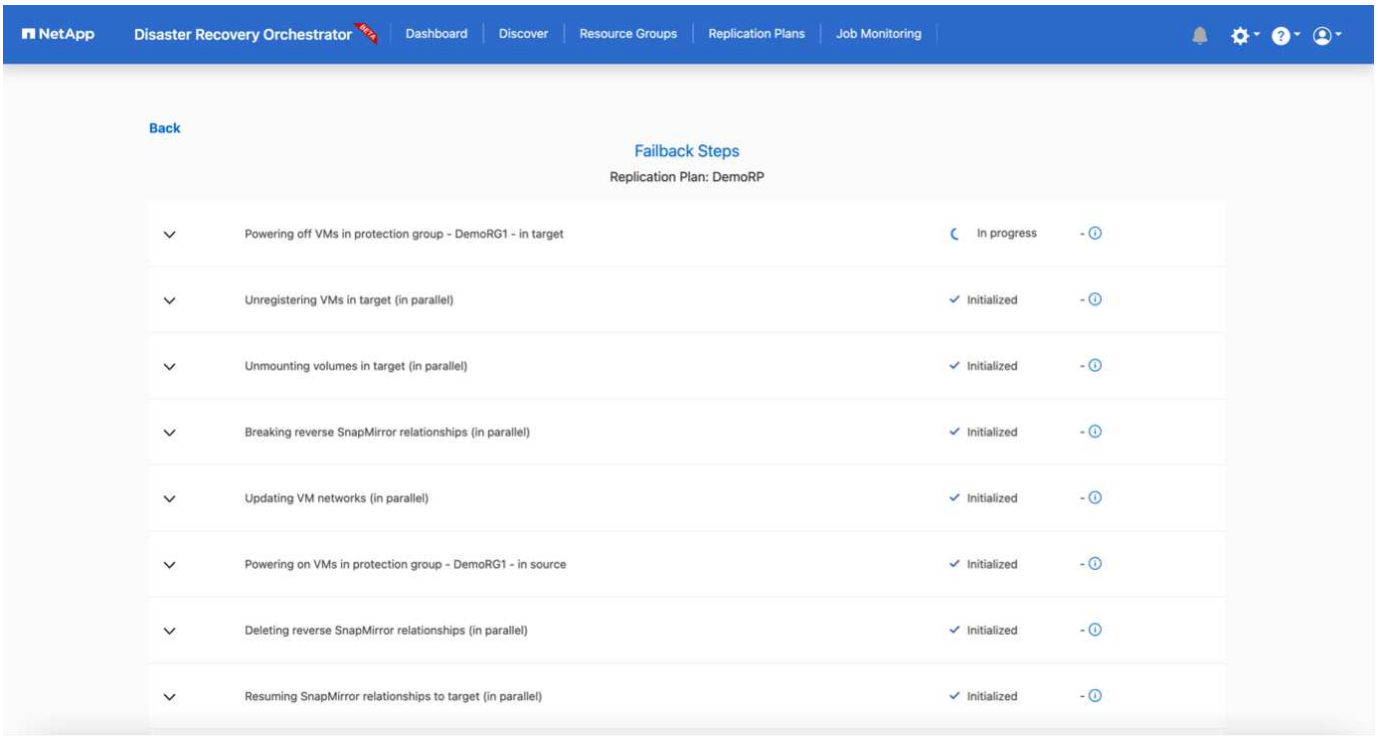


Failback can be triggered at the replication-plan level. For a test failover, the tear-down option can be used to roll back the changes and remove the FlexClone relationship. Failback related to failover is a two-step process. Select the replication plan and select **Reverse data sync**.

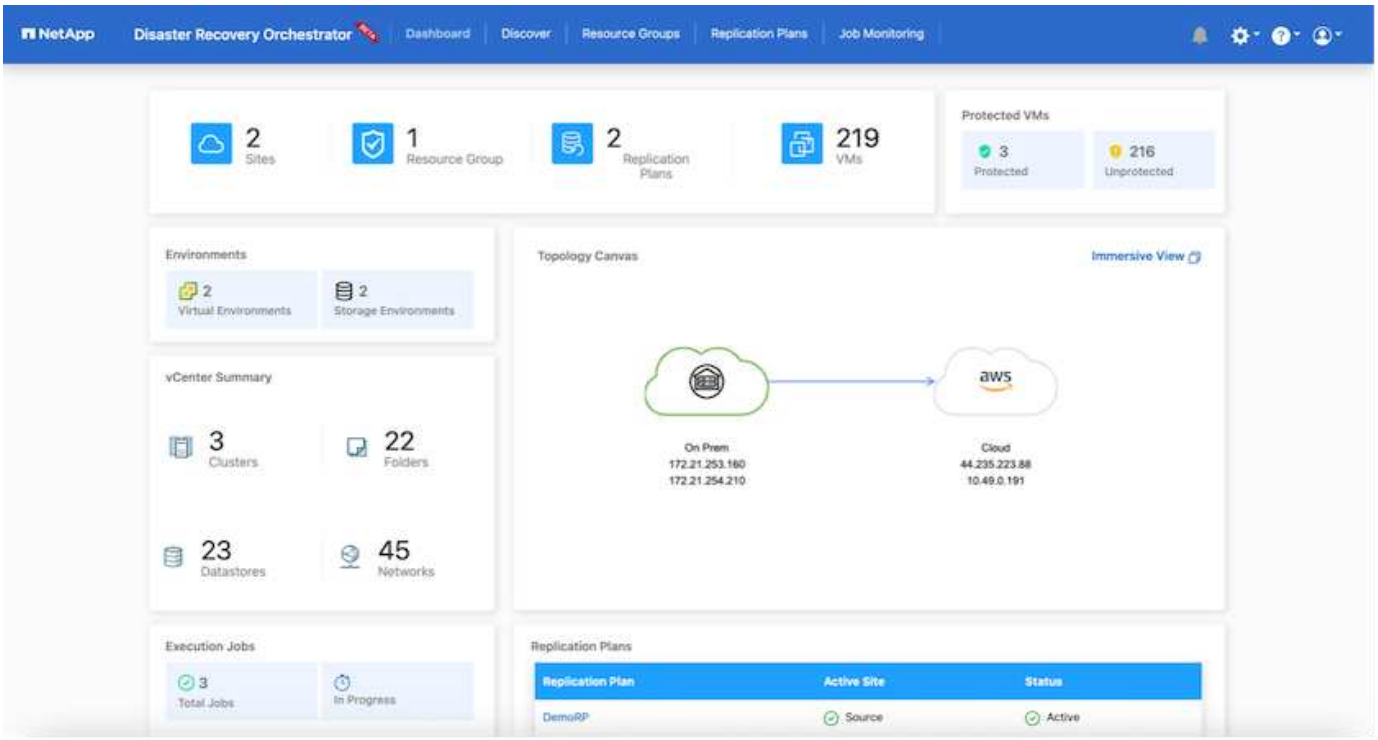


Once completed, you can trigger failback to move back to original production site.





From NetApp Console, we can see that replication health has broken off for the appropriate volumes (those that were mapped to VMC as read-write volumes). During test failover, DRO does not map the destination or replica volume. Instead, it makes a FlexClone copy of the required SnapMirror (or Snapshot) instance and exposes the FlexClone instance, which does not consume additional physical capacity for FSx ONTAP. This process makes sure that the volume is not modified and replica jobs can continue even during DR tests or triage workflows. Additionally, this process makes sure that, if errors occur or corrupted data is recovered, the recovery can be cleaned up without the risk of the replica being destroyed.



Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to pin-point where the safe point of return is and, once that is determined, to protect recovered workloads from reoccurring attacks from, for example, sleeping malware or vulnerable applications.

DRO addresses these concerns by enabling you to recover your system from any available point in time. You can also recover workloads to functional and yet isolated networks so that applications can function and communicate with each other in a location where they are not exposed to north-south traffic. This gives your security team a safe place to conduct forensics and make sure there is no hidden or sleeping malware.

Benefits

- Use of the efficient and resilient SnapMirror replication.
- Recovery to any available point in time with Snapshot copy retention.
- Full automation of all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery with ONTAP FlexClone technology using a method that doesn't change the replicated volume.
  - Avoids risk of data corruption for volumes or Snapshot copies.
  - Avoids replication interruptions during DR test workflows.
  - Potential use of DR data with cloud computing resources for workflows beyond DR such as DevTest, security testing, patch or upgrade testing, and remediation testing.
- CPU and RAM optimization to help lower cloud costs by allowing recovery to smaller compute clusters.

Using Veeam Replication and FSx ONTAP for Disaster recovery to VMware Cloud on AWS

Amazon FSx ONTAP integration with VMware Cloud on AWS is an AWS-managed external NFS datastore built on NetApp's ONTAP file system that can be attached to a cluster in the SDDC. It provides customers with flexible, high-performance virtualized storage infrastructure that scales independently of compute resources.

Overview

For those customers looking to use VMware Cloud on AWS SDDC as the disaster recovery target, FSx ONTAP datastores can be used to replicate data from on-premises using any validated third-party solution that provides VM replication capability. By adding FSx ONTAP datastore, it will enable cost optimised deployment than building VMware cloud on AWS SDDC with enormous amount of ESXi hosts just to accommodate the storage.

This approach also helps customers to use pilot light cluster in VMC along with FSx ONTAP datastores to host the VM replicas. The same process can also be extended as a migration option to VMware Cloud on AWS by gracefully failing over the replication plan.

Problem Statement

This document describes how to use FSx ONTAP datastore and Veeam Backup and replication to set up disaster recovery for on-premises VMware VMs to VMware Cloud on AWS using the VM replication functionality.

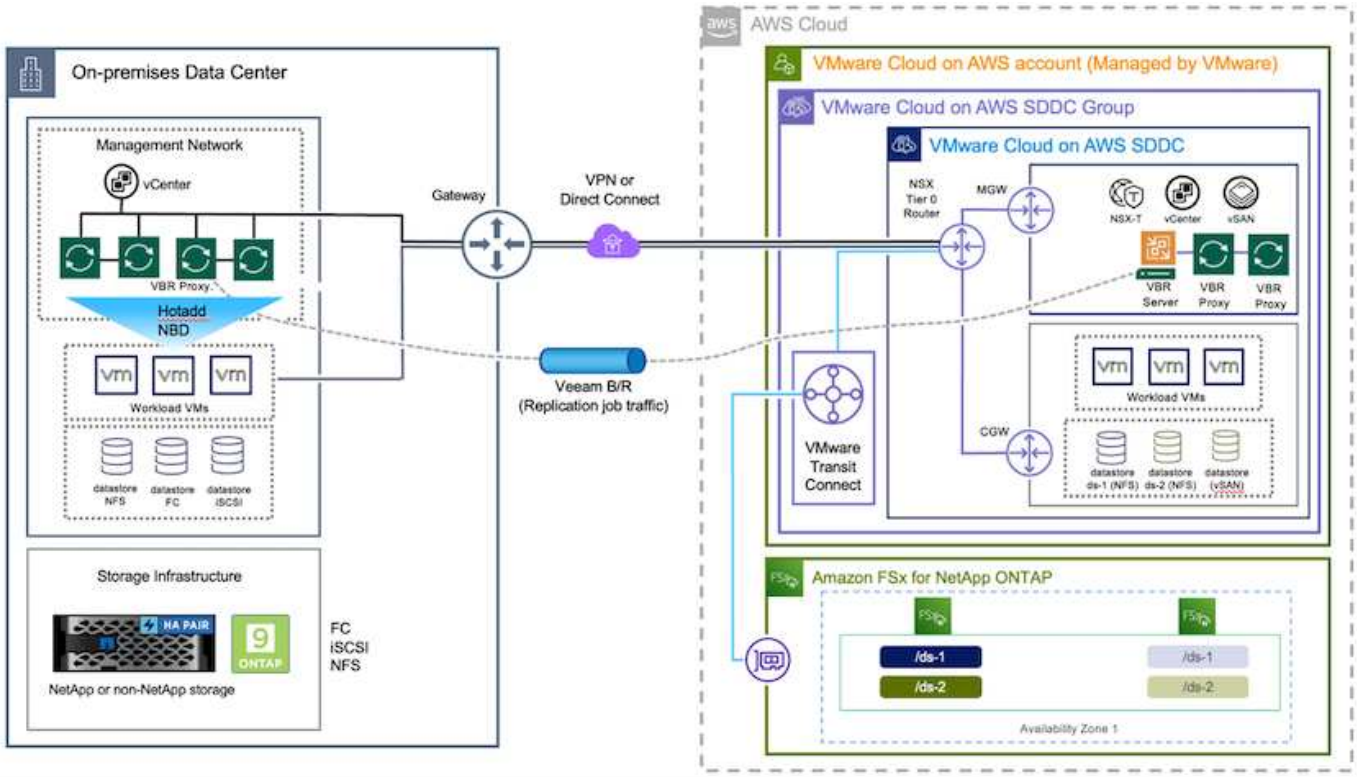
Veeam Backup & Replication allows onsite and remote replication for disaster recovery (DR). When virtual

machines are replicated, Veeam Backup & Replication creates an exact copy of the VMs in the native VMware vSphere format on the target VMware Cloud on AWS SDDC cluster and keeps the copy synchronized with the original VM.

Replication provides the best recovery time objective (RTO) values as there is a copy of a VM in the ready-to-start state. This replication mechanism ensures that the workloads can quickly start in VMware Cloud on AWS SDDC in case of a disaster event. The Veeam Backup & Replication software also optimizes traffic transmission for replication over WAN and slow connections. In addition, it also filters out duplicate data blocks, zero data blocks, swap files and excluded VM guest OS files, and compresses the replica traffic.

To prevent replication jobs from consuming the entire network bandwidth, WAN accelerators and network throttling rules can be put in place. The replication process in Veeam Backup & Replication is job driven which means replication is performed by configuring replication jobs. In case of a disaster event, failover can be triggered to recover the VMs by failing over to its replica copy.

When failover is performed, a replicated VM takes over the role of the original VM. Fail over can be performed to the latest state of a replica or to any of its good known restore points. This enables ransomware recovery or isolated testing as needed. In Veeam Backup & Replication, failover and failback are temporary intermediate step that should be further finalized. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.



Solution Deployment

High level steps

1. Veeam Backup and Replication software is running in on-premises environment with appropriate network connectivity.
2. Configure VMware Cloud on AWS, see the VMware Cloud Tech Zone article VMware Cloud on AWS integration with Amazon FSx ONTAP Deployment Guide to deploy, configure VMware Cloud on AWS SDDC and FSx ONTAP as NFS datastore. (A pilot-light environment set up with a minimal configuration

can be used for DR purposes. VMs will fail over to this cluster in the event of an incident, and additional nodes can be added).

3. Set up replication jobs to create VM replicas using Veeam Backup and Replication.
4. Create failover plan and perform failover.
5. Switch back to production VMs once the disaster event is complete and primary site is Up.

Pre-requisites for Veeam VM Replication to VMC and FSx ONTAP datastores

1. Ensure Veeam Backup & Replication backup VM is connected to the source vCenter as well as the target VMware cloud on AWS SDDC clusters.
2. The backup server must be able to resolve short names and connect to source and target vCenters.
3. The target FSx ONTAP datastore must have enough free space to store VMDKs of replicated VMs

For additional information, refer to "Considerations and Limitations" covered here.

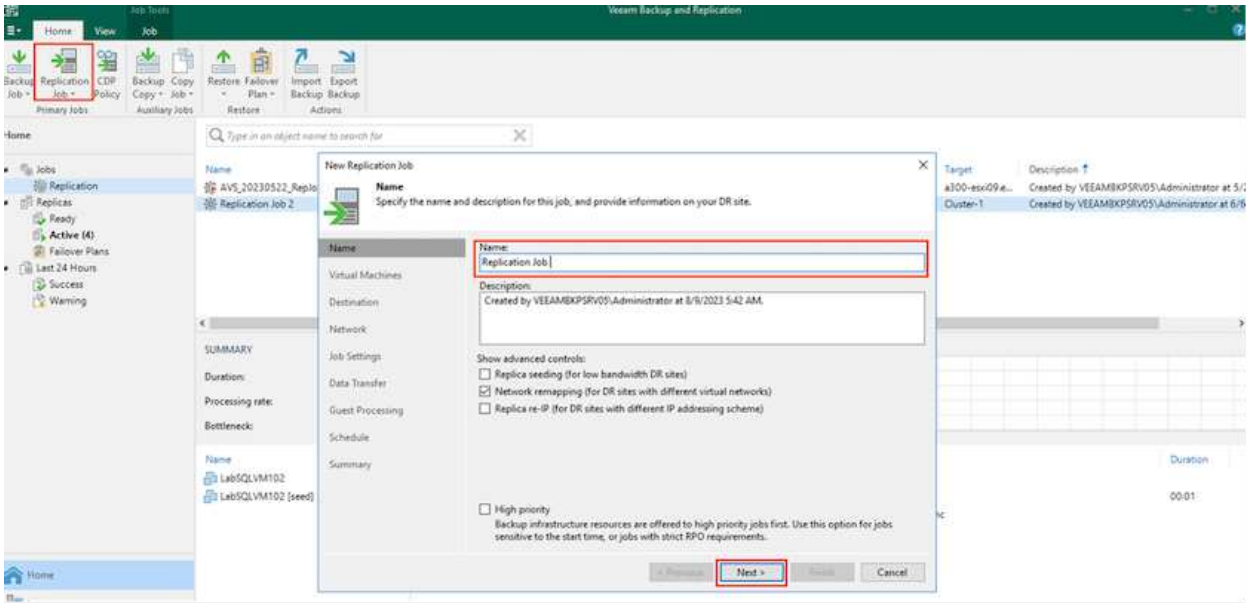
Deployment Details

▼ Step 1: Replicate VMs

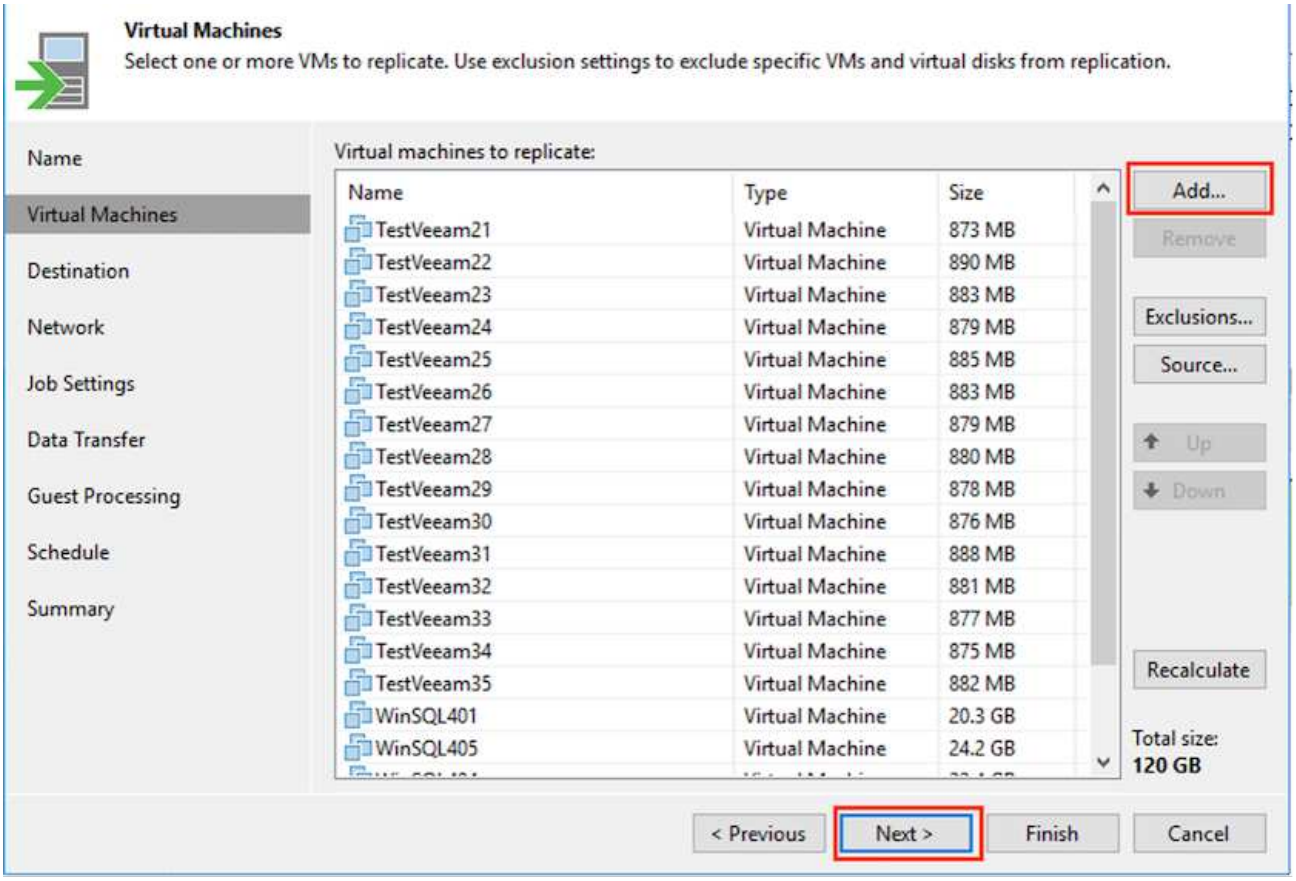
Veeam Backup & Replication leverages VMware vSphere snapshot capabilities and during replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot is the point-in-time copy of a VM that includes virtual disks, system state, configuration and so on. Veeam Backup & Replication uses the snapshot as a source of data for replication.

To replicate VMs, follow the below steps:

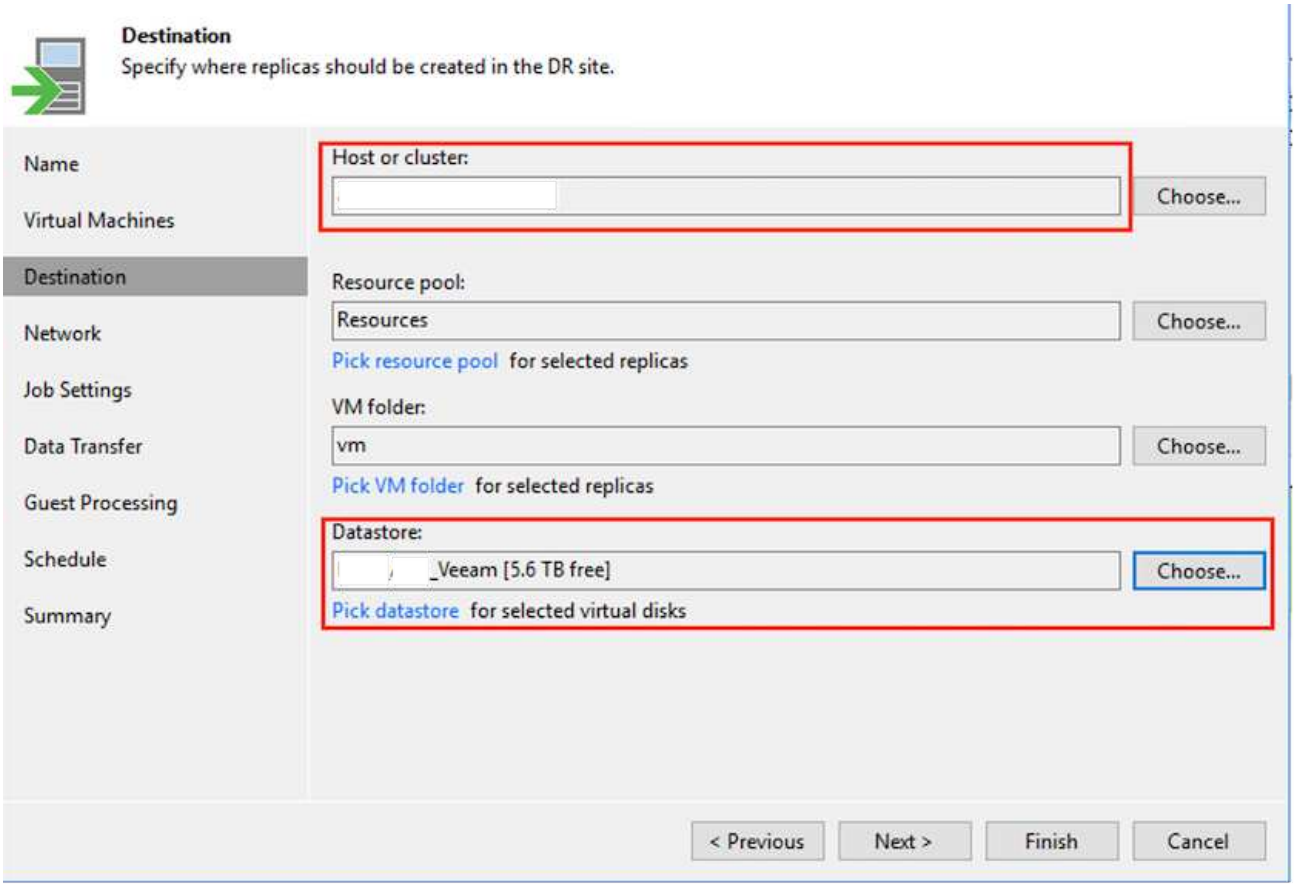
1. Open the Veeam Backup & Replication Console.
2. On the Home view, select Replication Job > Virtual machine > VMware vSphere.
3. Specify a job name and select the appropriate advanced control checkbox. Click Next.
    - Select the Replica seeding check box if connectivity between on-premises and AWS has restricted bandwidth.
    - Select the Network remapping (for AWS VMC sites with different networks) check box if segments on VMware Cloud on AWS SDDC do not match that of on-premises site networks.
    - If the IP addressing scheme in on-premises production site differs from the scheme in the AWS VMC site, select the Replica re-IP (for DR sites with different IP addressing scheme) check box.



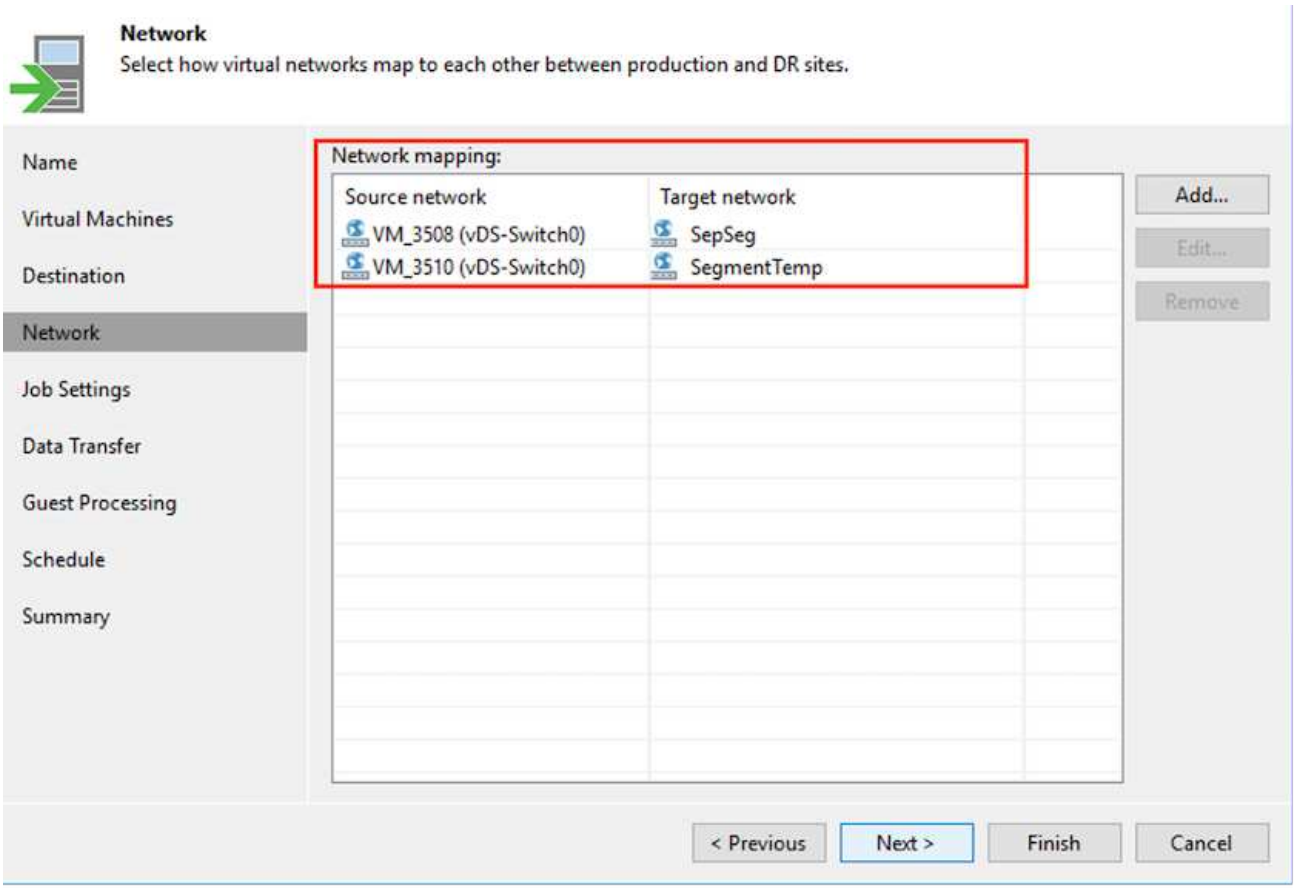
4. Select the VMs that needs to be replicated to FSx ONTAP datastore attached to VMware Cloud on AWS SDDC in the **Virtual Machines** step. The Virtual machines can be placed on vSAN to fill the available vSAN datastore capacity. In a pilot light cluster, the usable capacity of a 3-node cluster will be limited. The rest of the data can be replicated to FSx ONTAP datastores. Click **Add**, then in the **Add Object** window select the necessary VMs or VM containers and click **Add**. Click **Next**.



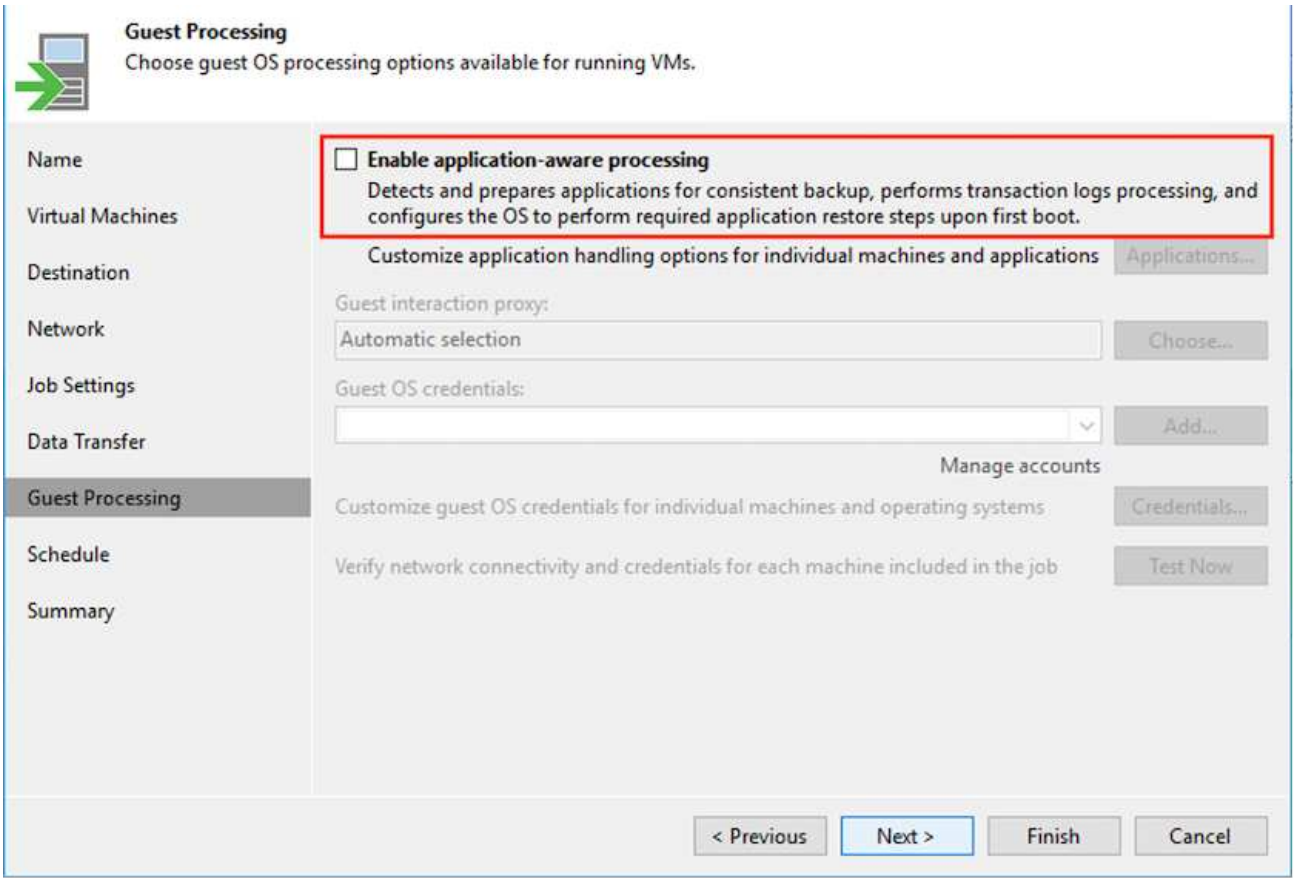
5. After that, select the destination as VMware Cloud on AWS SDDC cluster / host and the appropriate resource pool, VM folder and FSx ONTAP datastore for VM replicas. Then Click **Next**.



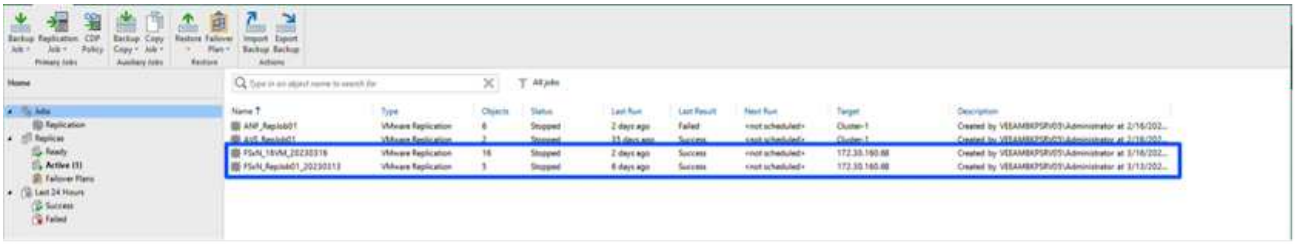
6. In the next step, create the mapping between source and destination virtual network as needed.



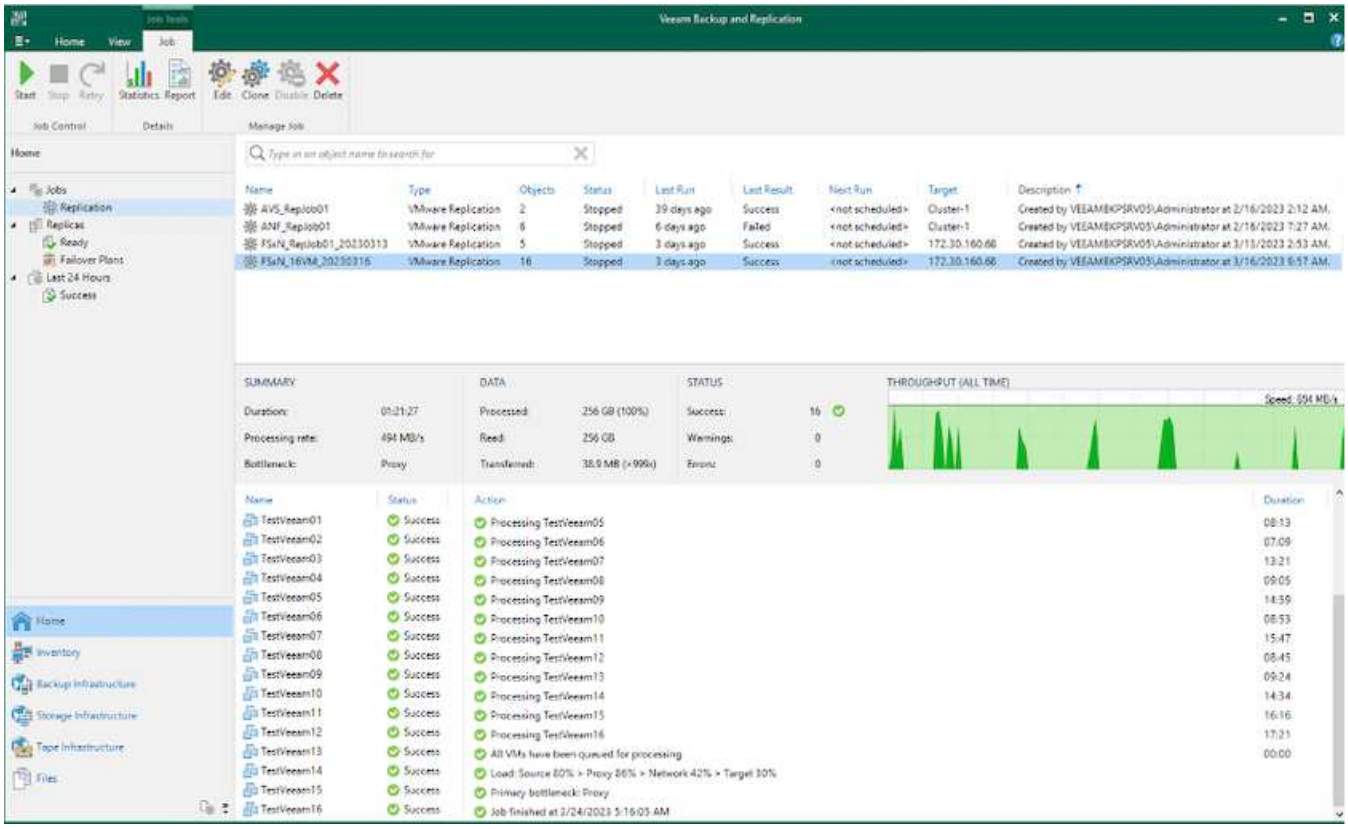
7. In the **Job Settings** step, specify the backup repository that will store metadata for VM replicas, retention policy and so on.
8. Update the **Source** and **Target** proxy servers in the **Data Transfer** step and leave **Automatic** selection (default) and keep **Direct** option selected and click **Next**.
9. At the **Guest Processing** step, select **Enable application-aware processing** option as needed. Click **Next**.



10. Choose the replication schedule to run the replication job to run on a regular basis.
11. At the **Summary** step of the wizard, review details of the replication job. To start the job right after the wizard is closed, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.



Once the replication job starts, the VMs with the suffix specified will be populated on the destination VMC SDDC cluster / host.



For additional information for Veeam replication, refer to How Replication Works.

▼ Step 2: Create a failover plan

When the initial replication or seeding is complete, create the failover plan. Failover plan helps in performing failover for dependent VMs one by one or as a group automatically. Failover plan is the blueprint for the order in which the VMs are processed including the boot delays. The failover plan also helps to ensure that critical dependant VMs are already running.

To create the plan, navigate to the new sub section called Replicas and select Failover Plan. Choose the appropriate VMs. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas.

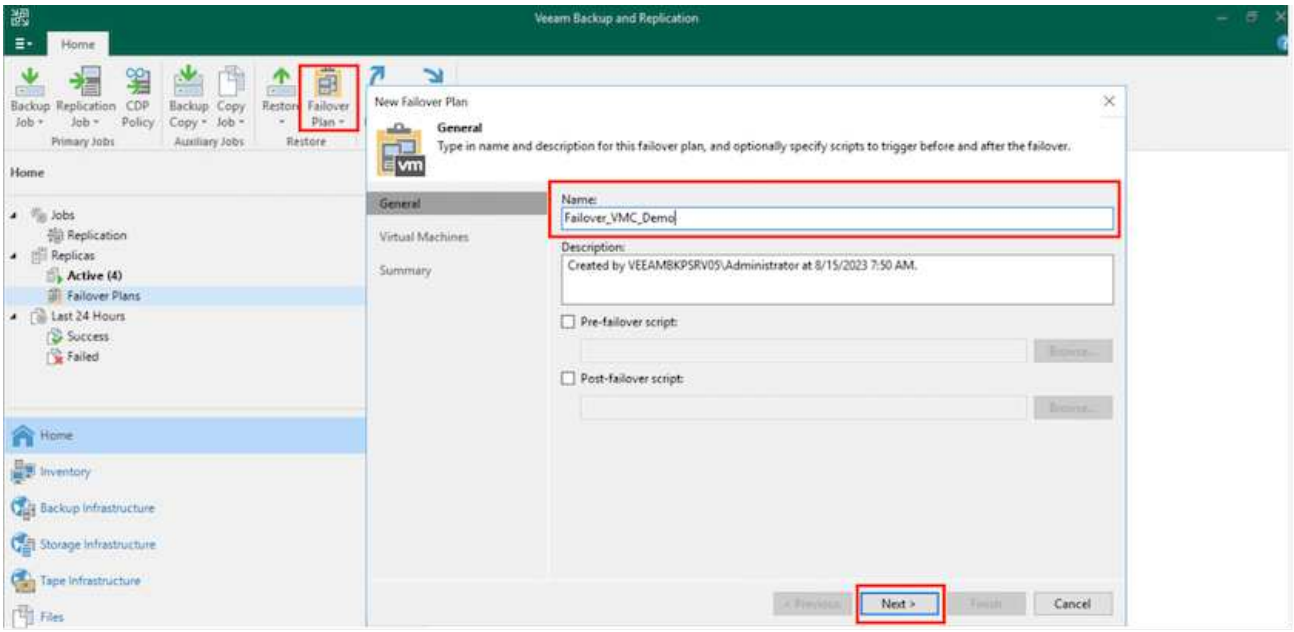
ⓘ The failover plan can only be added once the initial replication is complete and the VM replicas are in Ready state.

ⓘ The maximum number of VMs that can be started simultaneously when running a failover plan is 10.

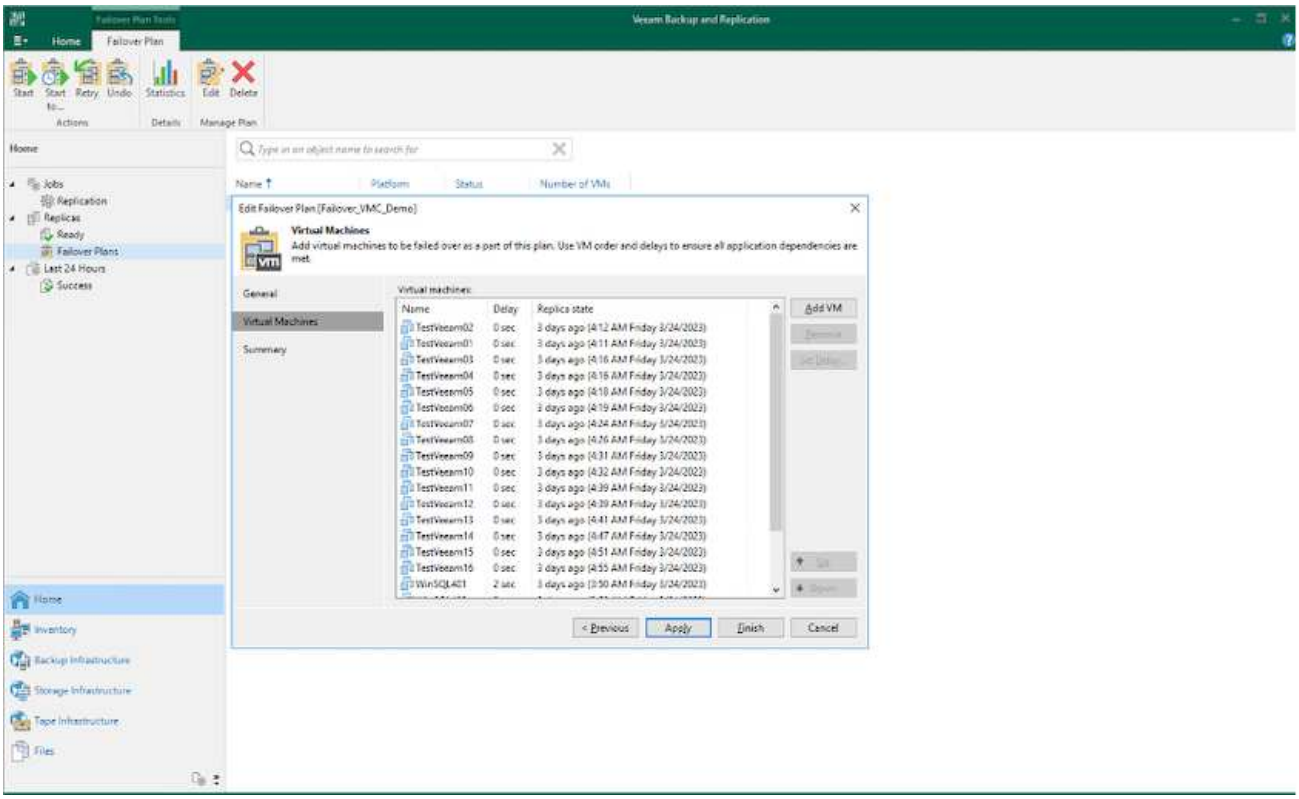
ⓘ During the failover process, the source VMs will not be powered off.

To create the **Failover Plan**, do the following:

1. On the Home view, select **Failover Plan > VMware vSphere**.
2. Next, provide a name and a description to the plan. Pre and Post-failover script can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.



3. Add the VMs to the plan and modify the VM boot order and boot delays to meet the application dependencies.



For additional information for creating replication jobs, refer Creating Replication Jobs.

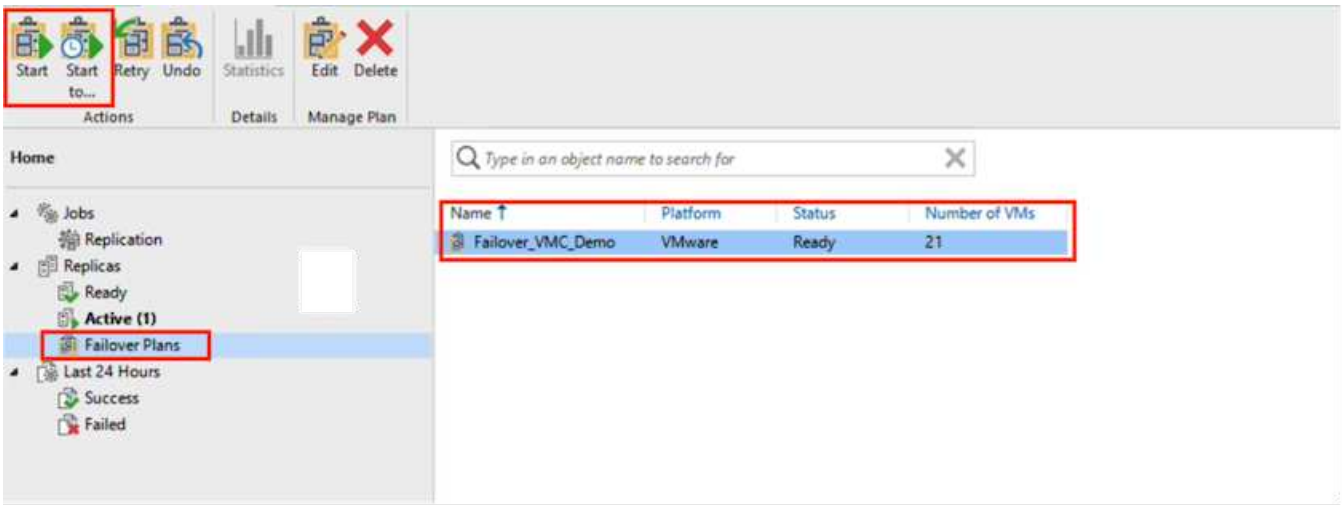
▼ **Step 3: Run the failover plan**

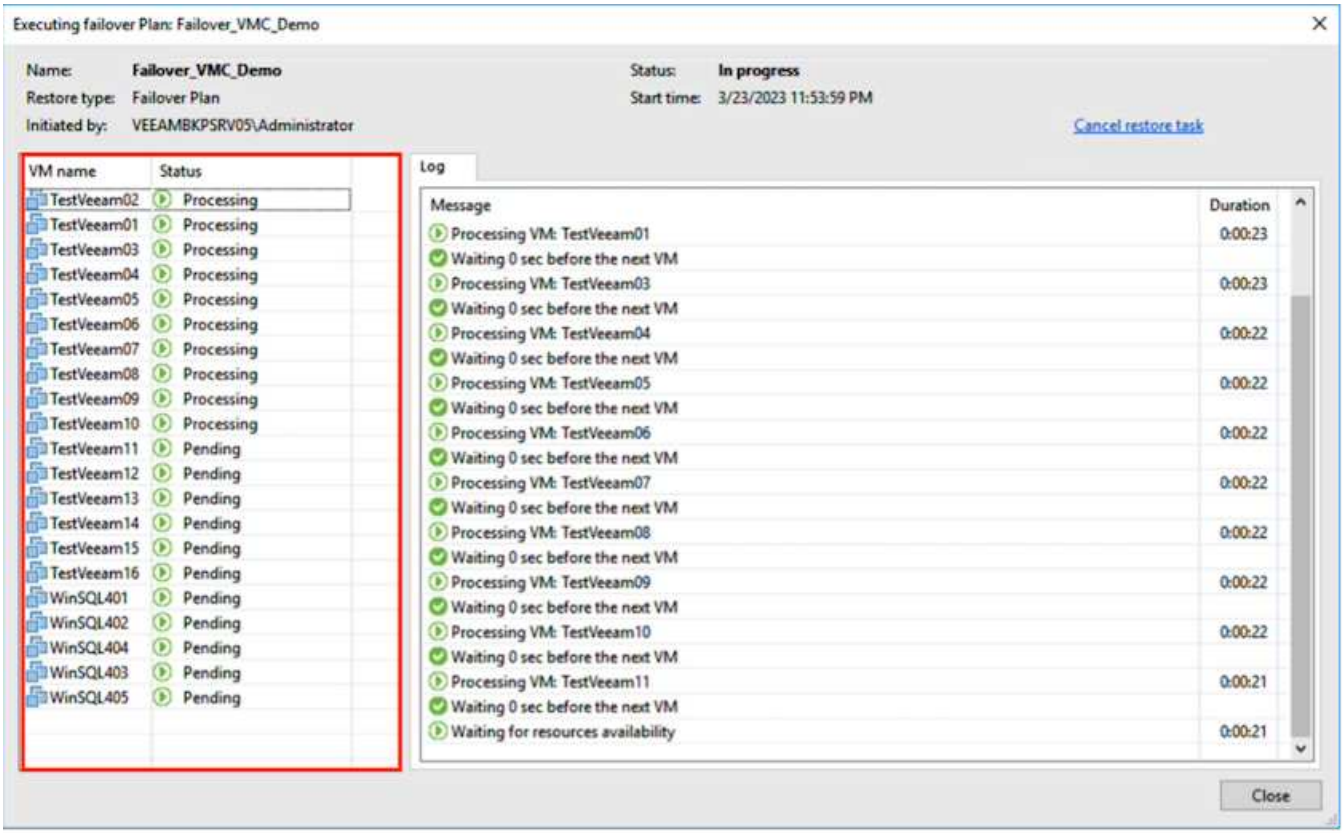
During failover, the source VM in the production site is switched over to its replica at the disaster recovery site. As part of the failover process, Veeam Backup & Replication restores the VM replica to the required restore point and moves all I/O activities from the source VM to its replica. Replicas can be used not only in case of a disaster, but also to simulate DR drills. During failover simulation, the source VM remains running. Once all the necessary tests have been conducted, you can undo the failover and return to normal

operations.

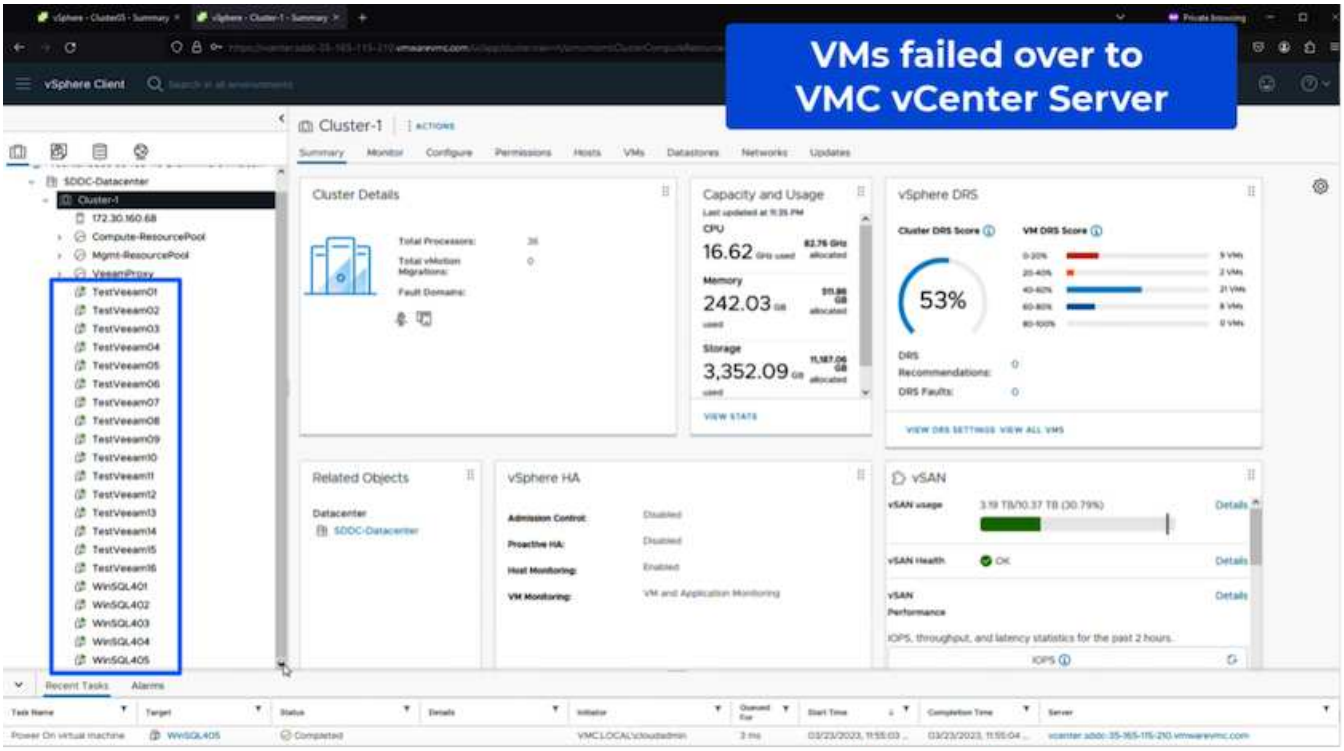
Make sure network segmentation is in place to avoid IP conflicts during DR drills.

To start the failover plan, simply click in **Failover Plans** tab and right click on the failover plan. Select **Start**. This will failover using the latest restore points of VM replicas. To fail over to specific restore points of VM replicas, select **Start to**.

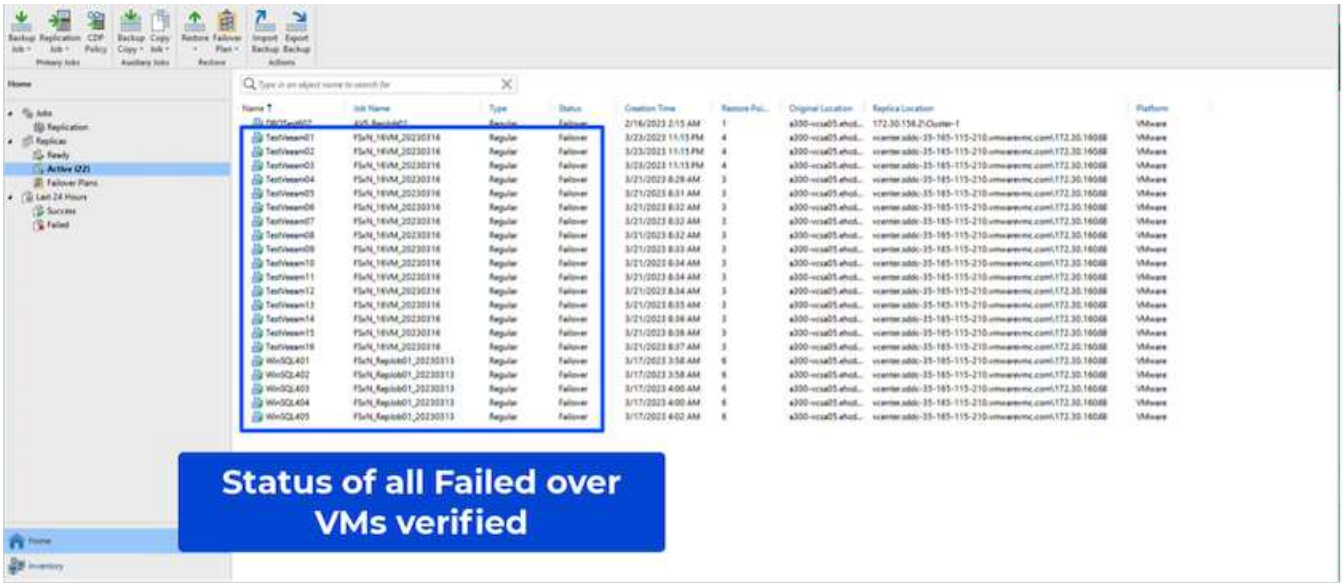




The state of the VM replica changes from Ready to Failover and VMs will start on the destination VMware Cloud on AWS SDDC cluster / host.



Once the failover is complete, the status of the VMs will change to "Failover".



Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the Ready state.

For detailed information about failover plans, refer to Failover Plans.

▼ **Step 4: Failback to the Production site**

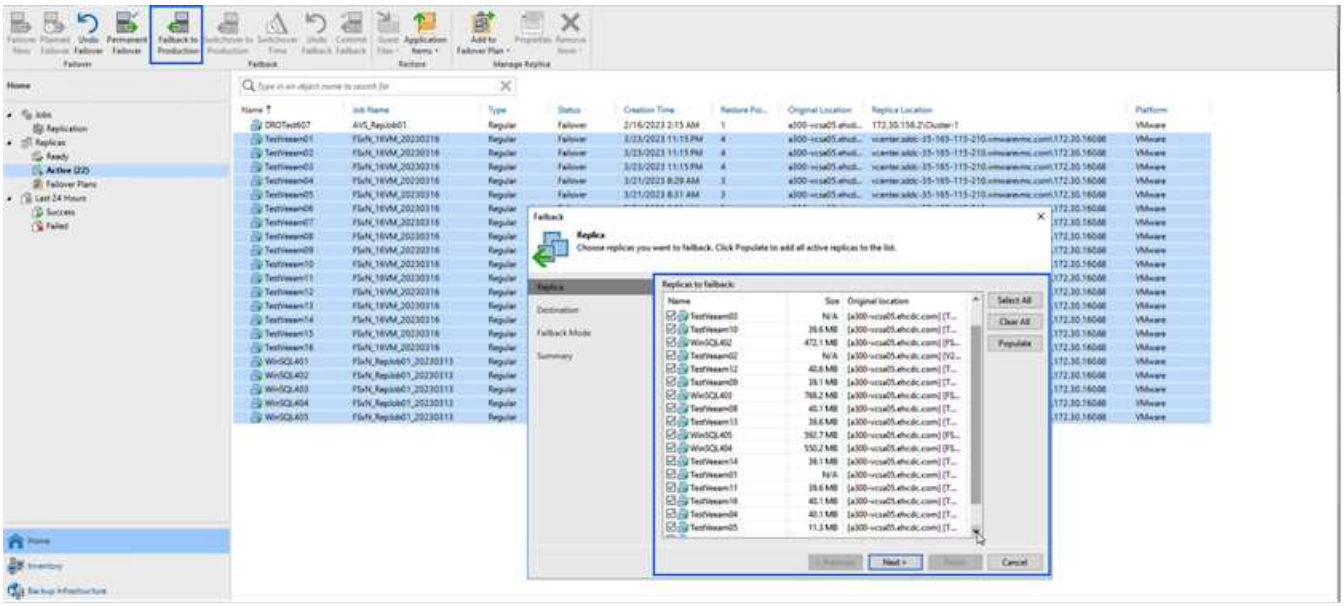
When the failover plan is running, it is considered as an intermediate step and needs to be finalized based on the requirement. The options include the following:

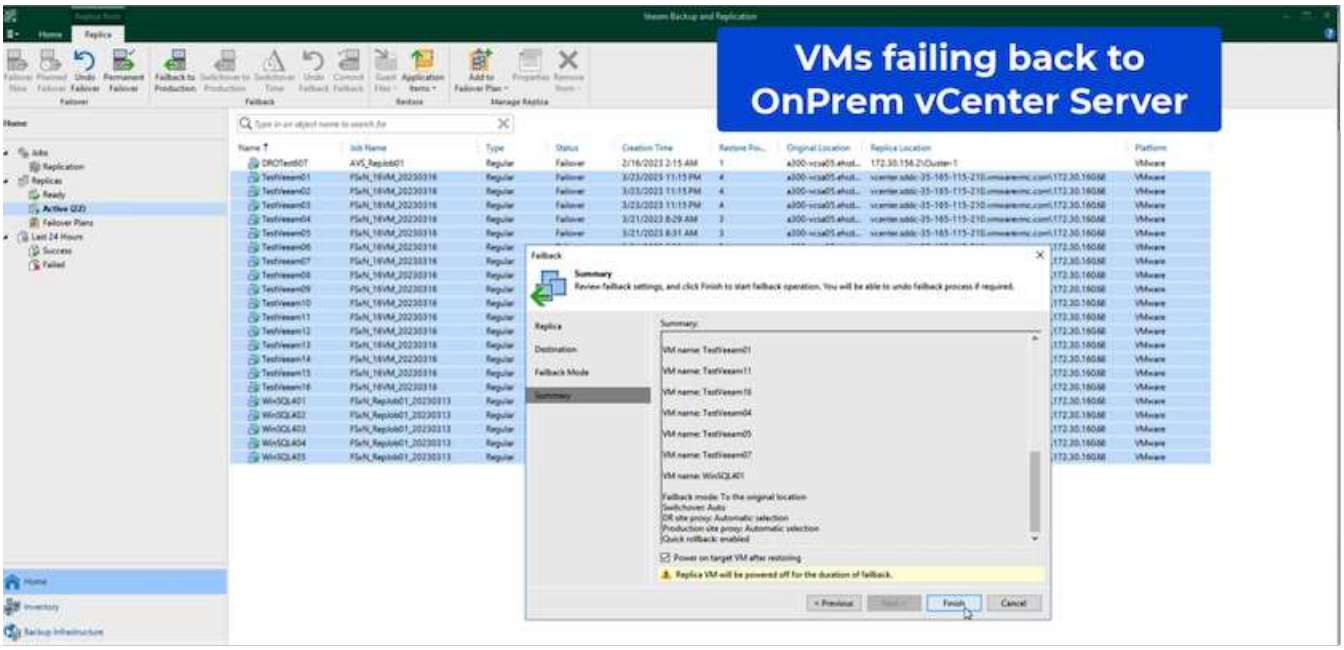
- **Failback to production** - switch back to the original VM and transfer all changes that took place while the VM replica was running to the original VM.

When you perform failback, changes are only transferred but not published. Choose **Commit failback** (once the original VM is confirmed to work as expected) or **Undo failback** to get back to the VM replica If the original VM is not working as expected.

- **Undo failover** - switch back to the original VM and discard all changes made to the VM replica while it was running.
- **Permanent Failover** - permanently switch from the original VM to a VM replica and use this replica as the original VM.

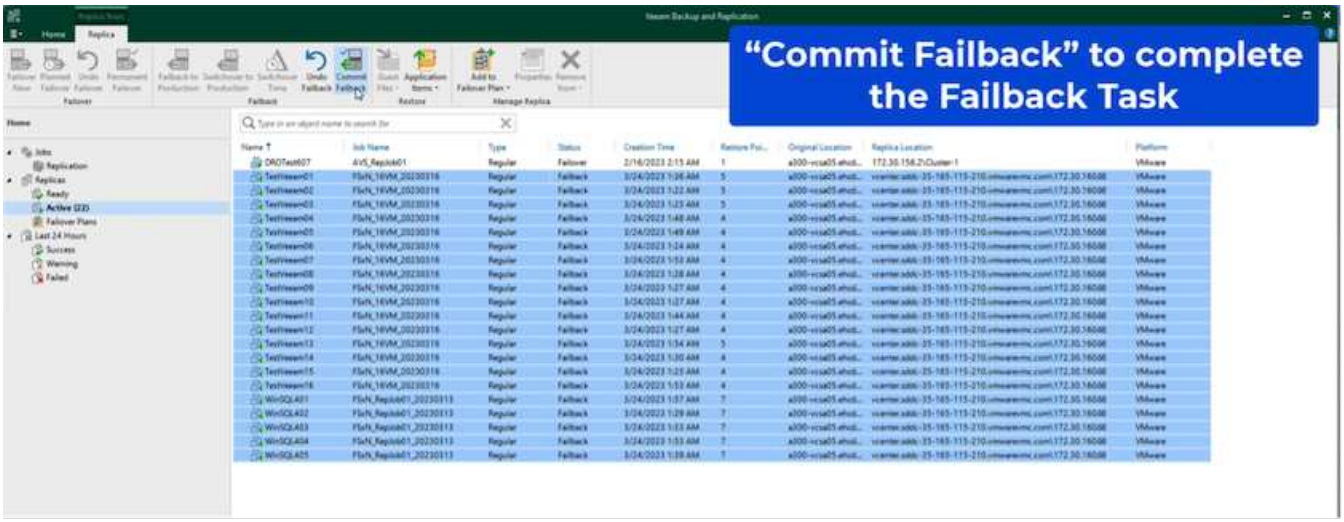
In this demo, Failback to production was chosen. Failback to the original VM was selected during the Destination step of the wizard and "Power on VM after restoring" check box was enabled.

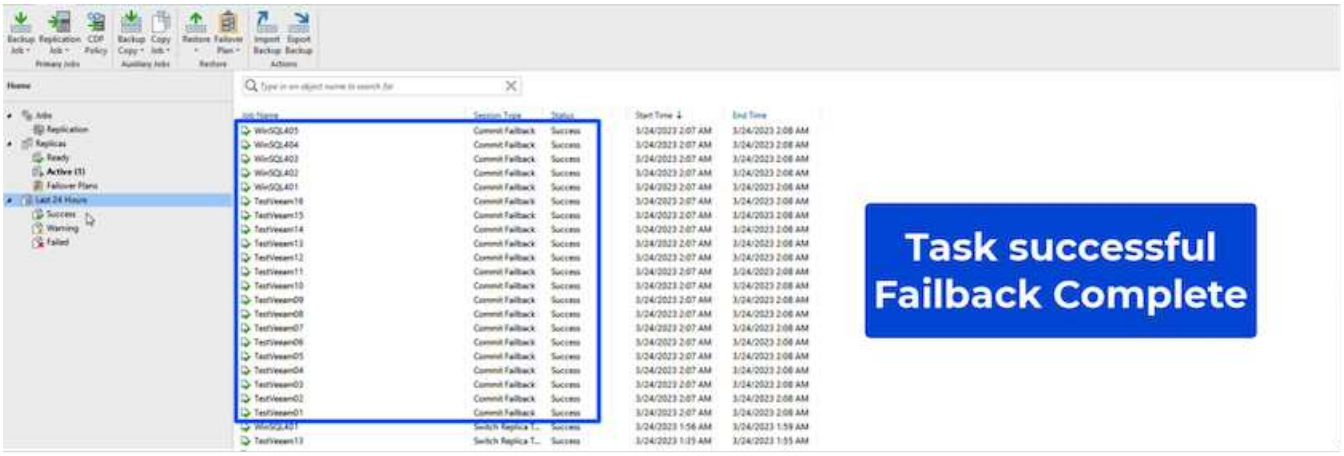




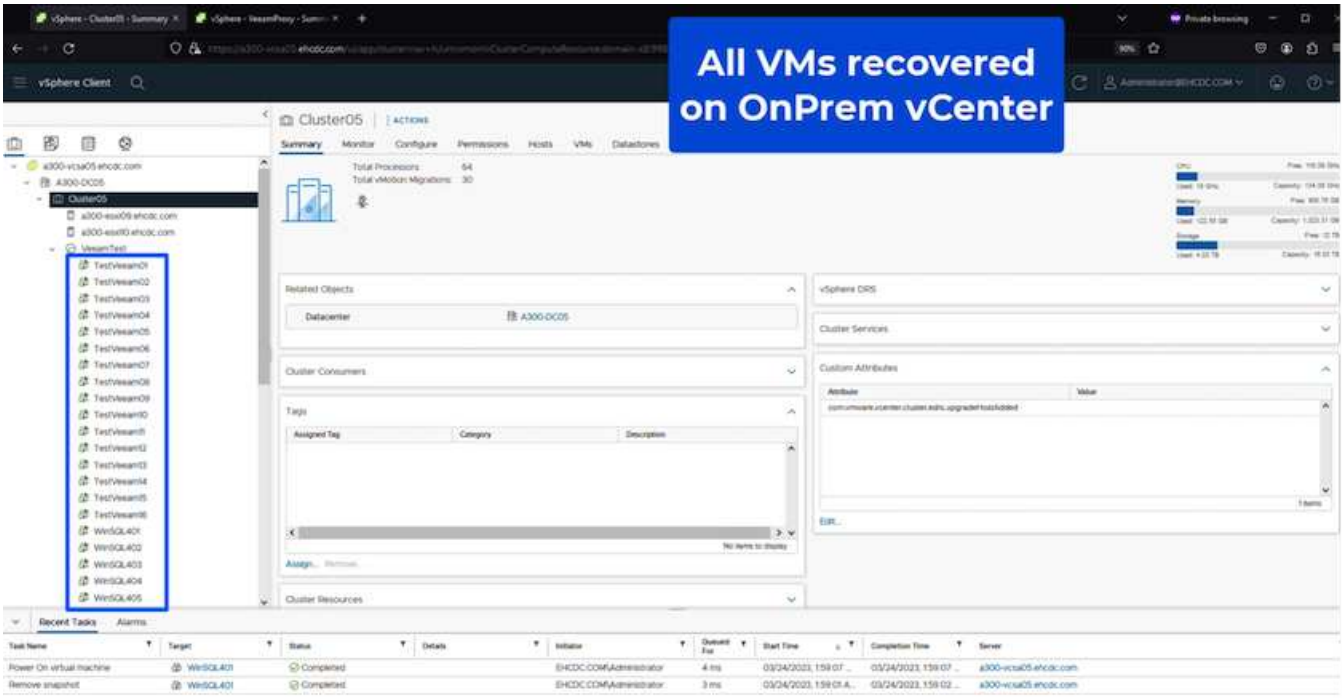
Failback commit is one of the ways to finalize failback operation. When failback is committed, it confirms that the changes sent to the VM which is failed back (the production VM) are working as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

For detailed information about the failback process, refer Veeam documentation for Failover and Failback for replication.





After failback to production is successful, the VMs are all restored back to the original production site.



Conclusion

FSx ONTAP datastore capability enables Veeam or any validated third-party tool to provide low-cost DR solution using Pilot light cluster and without standing up large number of hosts in the cluster just to accommodate the VM replica copy. This provides a powerful solution to handle a tailored, customized disaster recovery plan and also allows to reuse existing backup products in house to meet the DR needs, thus enabling cloud-based disaster recovery by exiting DR datacentres on-premises. Failover can be done as planned failover or failover with a click of a button when disaster occurs, and decision is made to activate the DR site.

To learn more about this process, feel free to follow the detailed walkthrough video.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=15fed205-8614-4ef7-b2d0-b061015e925a

Migrate workloads on AWS/VMC

TR 4942: Migrate Workloads to FSx ONTAP datastore using VMware HCX

A common use case for VMware Cloud (VMC) on Amazon Web Services (AWS), with its supplemental NFS datastore on Amazon FSx ONTAP, is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration methods to move on-premises virtual machines (VMs) and their data, running on any VMware supported datastores, to VMC datastores, which includes supplemental NFS datastores on FSx ONTAP.

Overview: Migrating virtual machines with VMware HCX, FSx ONTAP supplemental datastores, and VMware Cloud

VMware HCX is primarily a mobility platform that is designed to simplify workload migration, workload rebalancing, and business continuity across clouds. It is included as part of VMware Cloud on AWS and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for deploying and configuring VMware HCX, including all its main components, on-premises and on the cloud data center side, which enables various VM migration

mechanisms.

For more information, see VMware HCX User Guide and Install Checklist B - HCX with a VMware Cloud on AWS SDDC Destination Environment.

▼ High-level steps

This list provides the high-level steps to install and configure VMware HCX:

1. Activate HCX for the VMC software-defined data center (SDDC) through VMware Cloud Services Console.
2. Download and deploy the HCX Connector OVA installer in the on-premises vCenter Server.
3. Activate HCX with a license key.
4. Pair on-premises VMware HCX Connector with VMC HCX Cloud Manager.
5. Configure the network profile, compute profile, and service mesh.
6. (Optional) Perform Network Extension to extend the network and avoid re-IP.
7. Validate the appliance status and ensure that migration is possible.
8. Migrate the VM workloads.

▼ Prerequisites

Before you begin, make sure the following prerequisites are met. For more information, see Preparing for Installation. After the prerequisites are in place, including connectivity, configure and activate HCX by generating a license key from the VMware HCX Console at VMC. After HCX is activated, the vCenter Plug-in is deployed and can be accessed by using the vCenter Console for management.

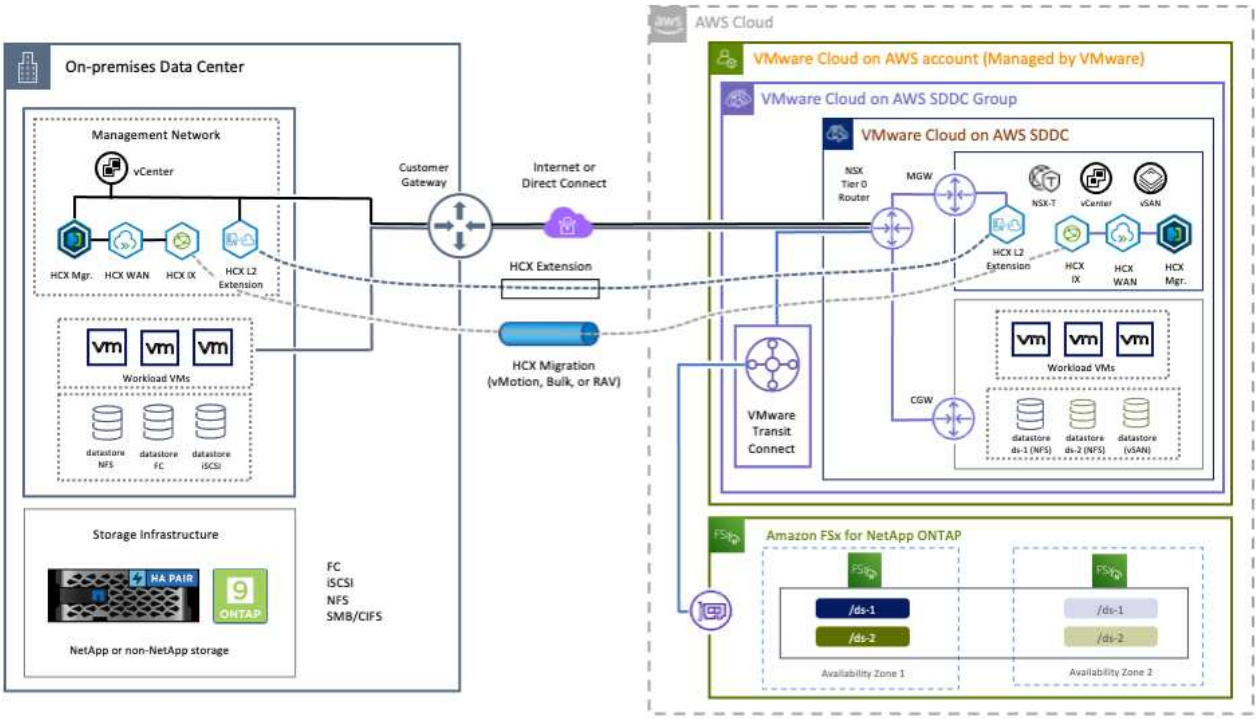
The following installation steps must be completed before proceeding with HCX activation and deployment:

1. Use an existing VMC SDDC or create a new SDDC following this NetApp link or this VMware link.
2. The network path from the on-premises vCenter environment to the VMC SDDC must support migration of VMs by using vMotion.
3. Make sure the required firewall rules and ports are allowed for vMotion traffic between the on-premises vCenter Server and the SDDC vCenter.
4. The FSx ONTAP NFS volume should be mounted as a supplemental datastore in the VMC SDDC. To attach the NFS datastores to the appropriate cluster, follow the steps outlined in this NetApp link or this VMware link.

▼ High Level Architecture

For testing purposes, the on-premises lab environment used for this validation was connected through a site-to-site VPN to AWS VPC, which allowed on-premises connectivity to AWS and to VMware cloud SDDC through External transit gateway. HCX migration and network extension traffic flows over the internet between on-premises and VMware cloud destination SDDC. This architecture can be modified to use Direct Connect private virtual interfaces.

The following image depicts the high-level architecture.



Solution Deployment

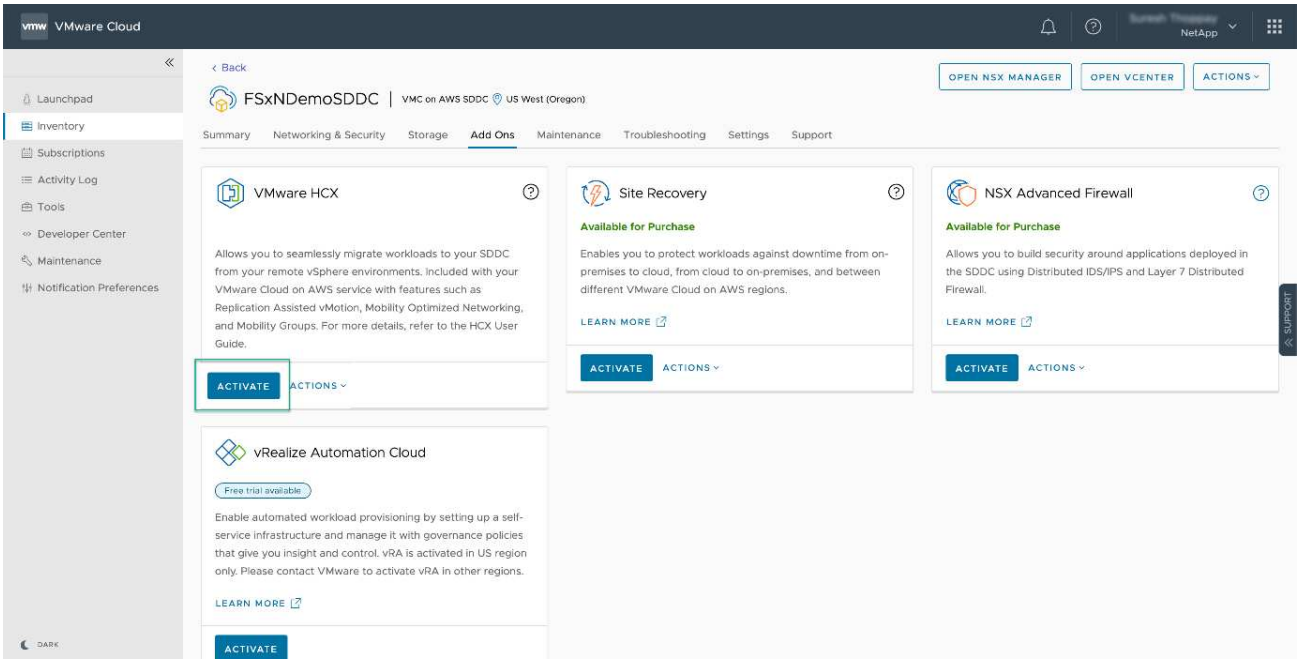
Follow the series of steps to complete the deployment of this solution:

▼ **Step 1: Activate HCX through VMC SDDC using the Add-ons option**

To perform the installation, complete the following steps:

1. Log in to the VMC Console at vmc.vmware.com and access Inventory.
2. To select the appropriate SDDC and access Add- ons, click View Details on SDDC and select the Add Ons tab.
3. Click Activate for VMware HCX.

This step takes up to 25 minutes to complete.

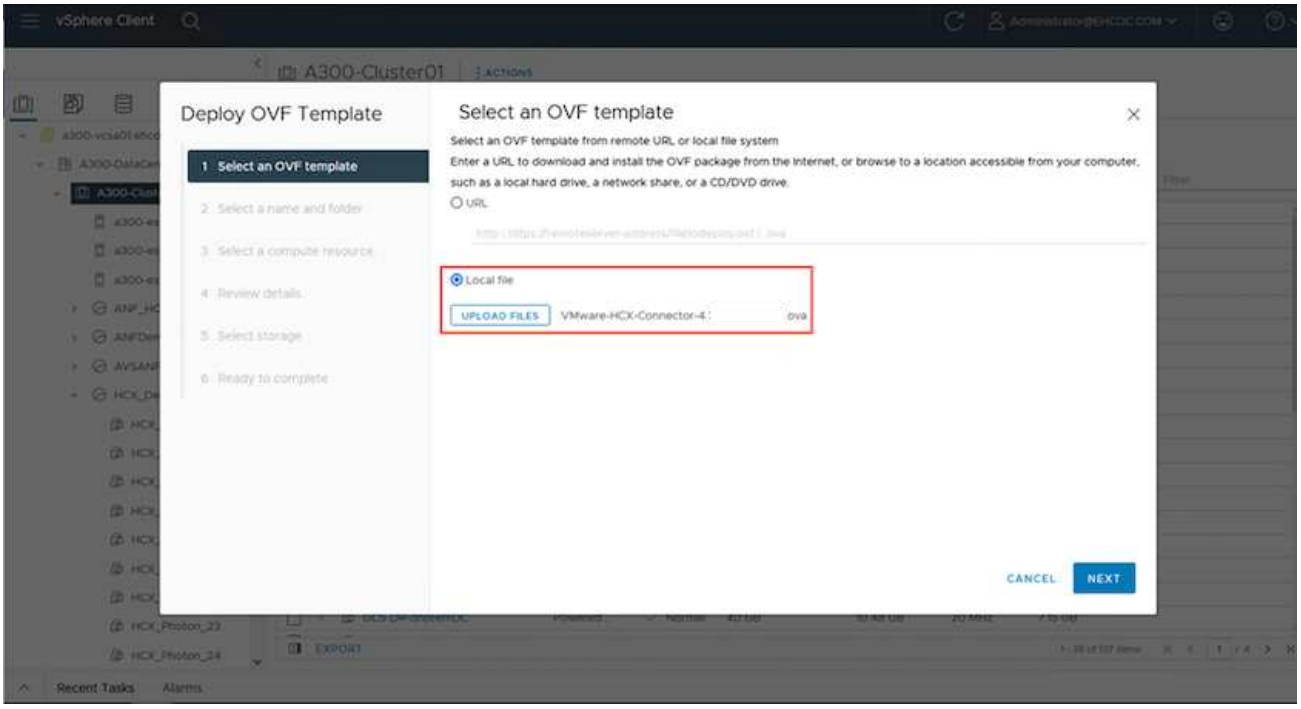


4. After the deployment is complete, validate the deployment by confirming that HCX Manager and its associated plug-ins are available in vCenter Console.
5. Create the appropriate Management Gateway firewalls to open the ports necessary to access HCX Cloud Manager.HCX Cloud Manager is now ready for HCX operations.

▼ Step 2: Deploy the installer OVA in the on-premises vCenter Server

For the on-premises Connector to communicate with the HCX Manager in VMC, make sure that the appropriate firewall ports are open in the on-premises environment.

1. From the VMC Console, navigate to the HCX Dashboard, go to Administration, and select the Systems Update tab. Click Request a Download Link for the HCX Connector OVA image.
2. With the HCX Connector downloaded, deploy the OVA in the on-premises vCenter Server. Right- click vSphere Cluster and select the Deploy OVF Template option.

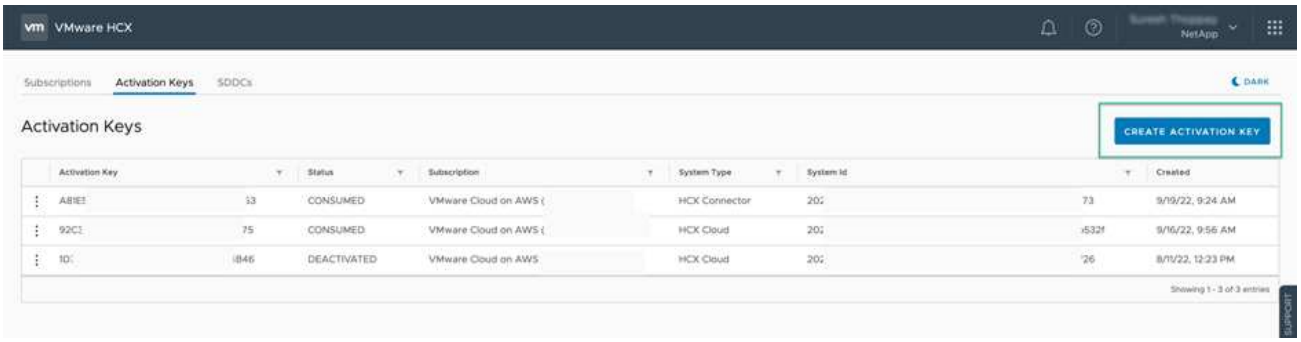


3. Enter the required information in the Deploy OVF Template wizard, click Next and then Finish to deploy the VMware HCX Connector OVA.
4. Power on the virtual appliance manually.For step- by- step instructions, go to VMware HCX User Guide.

▼ Step 3: Activate HCX Connector with the license key

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the VMware HCX Console at VMC and input the license during the VMware HCX Connector setup.

1. From the VMware Cloud Console, go to Inventory, select the SDDC, and click View Details. From the Add Ons tab, in the VMware HCX tile, click Open HCX.
2. From the Activation Keys tab, click Create Activation Key. Select the System Type as HCX Connector and click Confirm to generate the key. Copy the activation key.



> A separate key is required for each HCX Connector deployed on-premises.

3. Log in to the on-premises VMware HCX Connector at "`https://hcxconnectorIP:9443`" using administrator credentials.

> Use the password defined during the OVA deployment.

4. In the Licensing section, enter the activation key copied from step 2 and click Activate.

   The on-premises HCX Connector must have internet access for the activation to complete successfully.

5. Under Datacenter Location, provide the desired location for installing the VMware HCX Manager on-premises. Click Continue.
6. Under System Name, update the name and click Continue.
7. Select Yes and then Continue.
8. Under Connect Your vCenter, provide the IP address or fully qualified domain name (FQDN) and the credentials for the vCenter Server and click Continue.

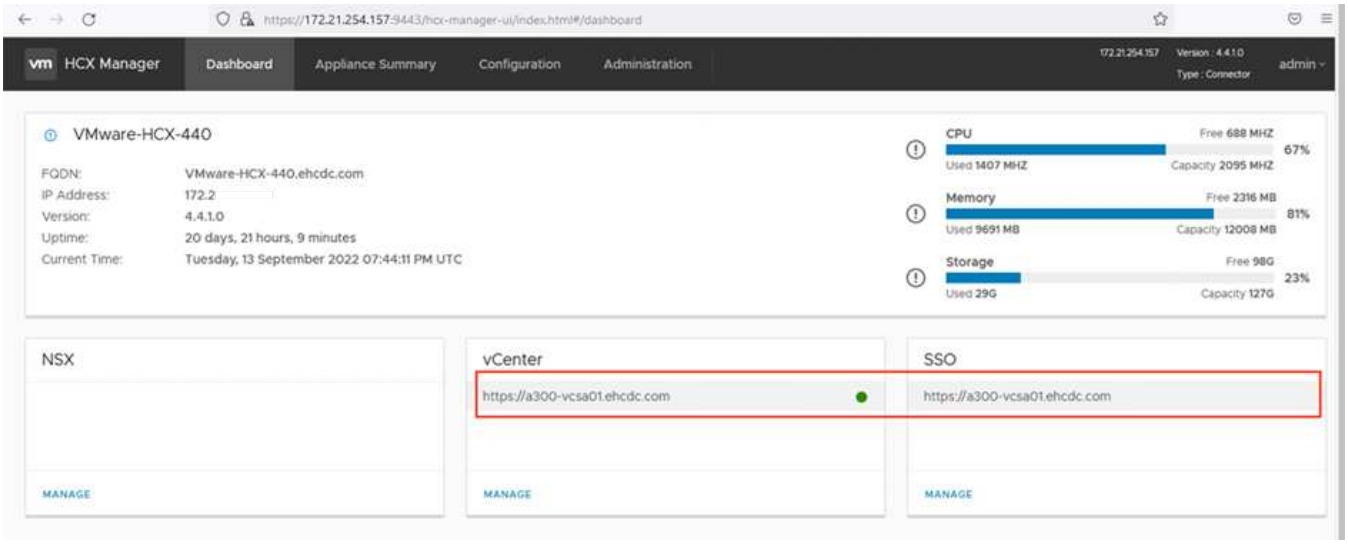
   Use the FQDN to avoid communication issues later.

9. Under Configure SSO/PSC, provide the Platform Services Controller's FQDN or IP address and click Continue.

   Enter the vCenter Server's IP address or FQDN.

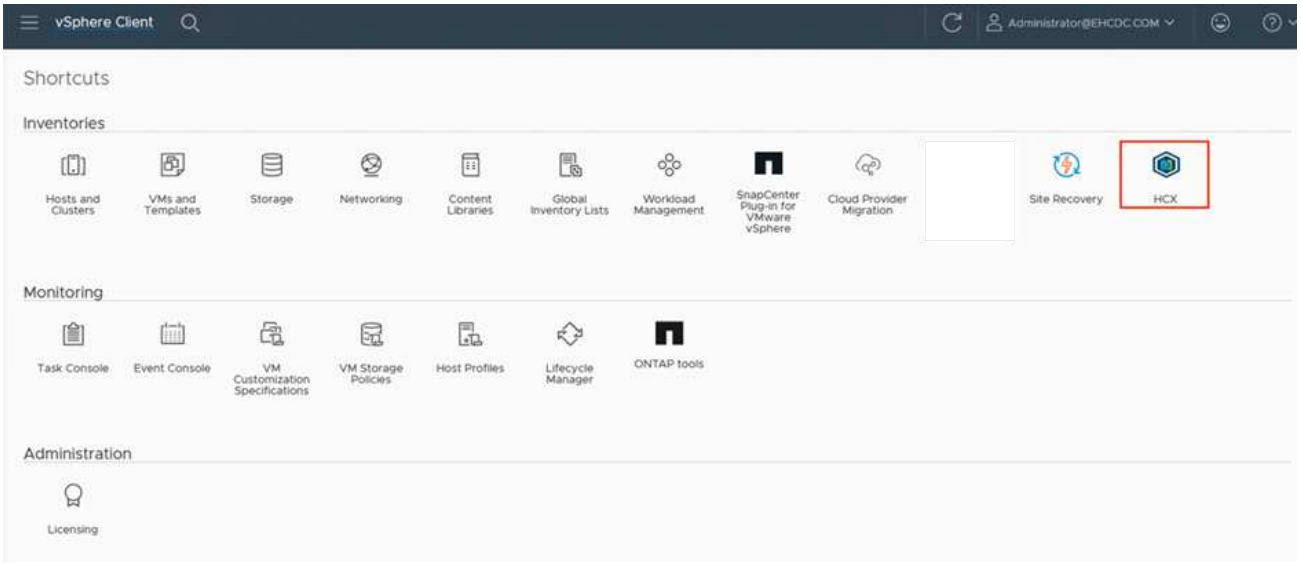
10. Verify that the information is entered correctly and click Restart.
11. After complete, the vCenter Server is displayed as green. Both the vCenter Server and SSO must have the correct configuration parameters, which should be the same as the previous page.

    This process should take approximately 10–20 minutes and for the plug-in to be added to the vCenter Server.

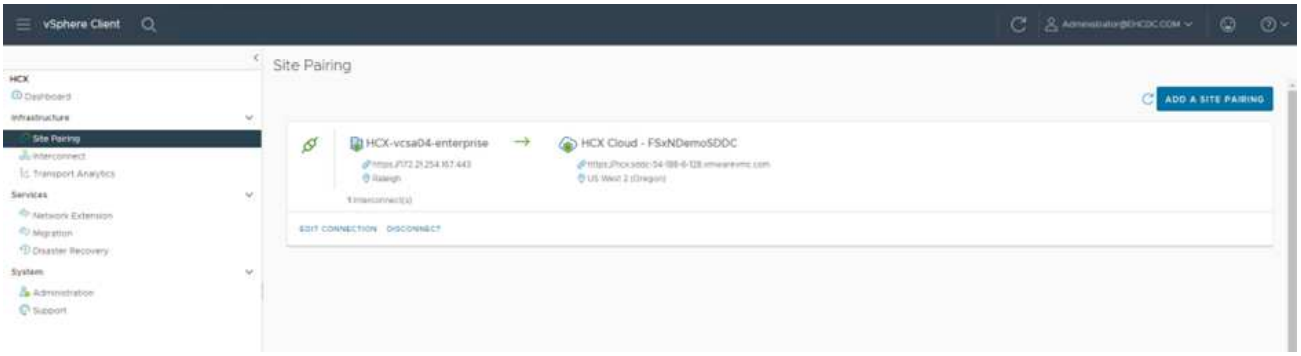


▼ Step 4: Pair on-premises VMware HCX Connector with VMC HCX Cloud Manager

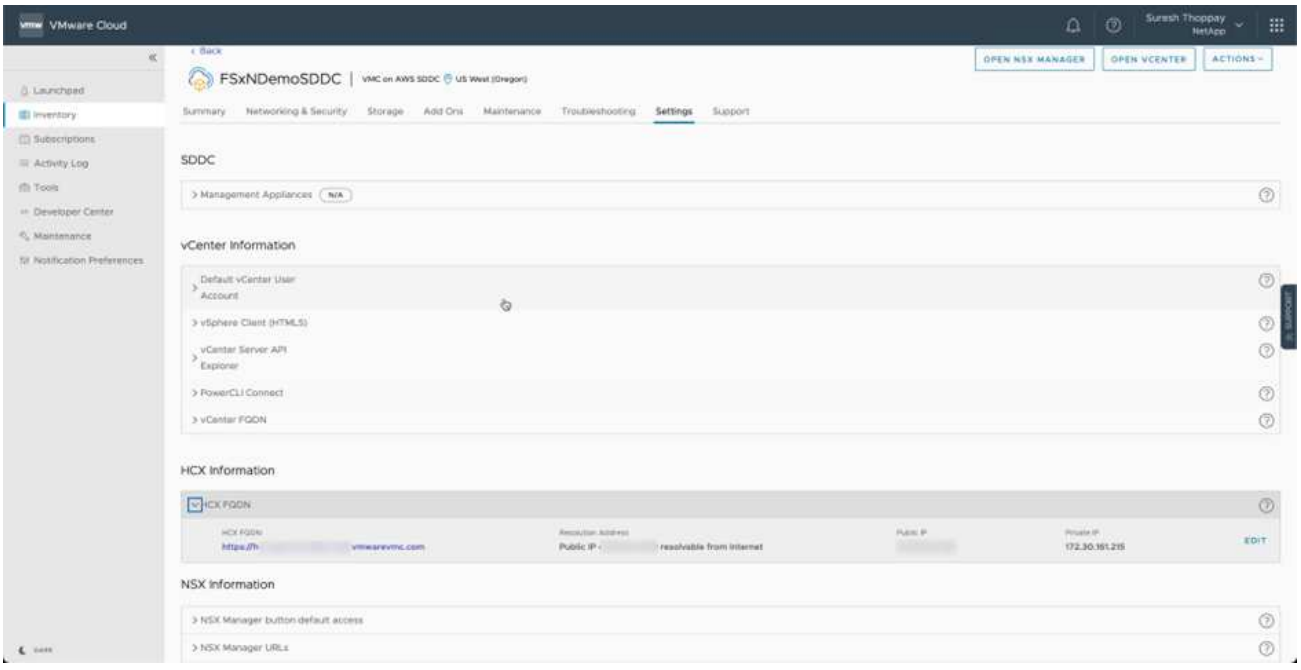
1. To create a site pair between the on-premises vCenter Server and the VMC SDDC, log in to the on-premises vCenter Server and access the HCX vSphere Web Client Plug- in.

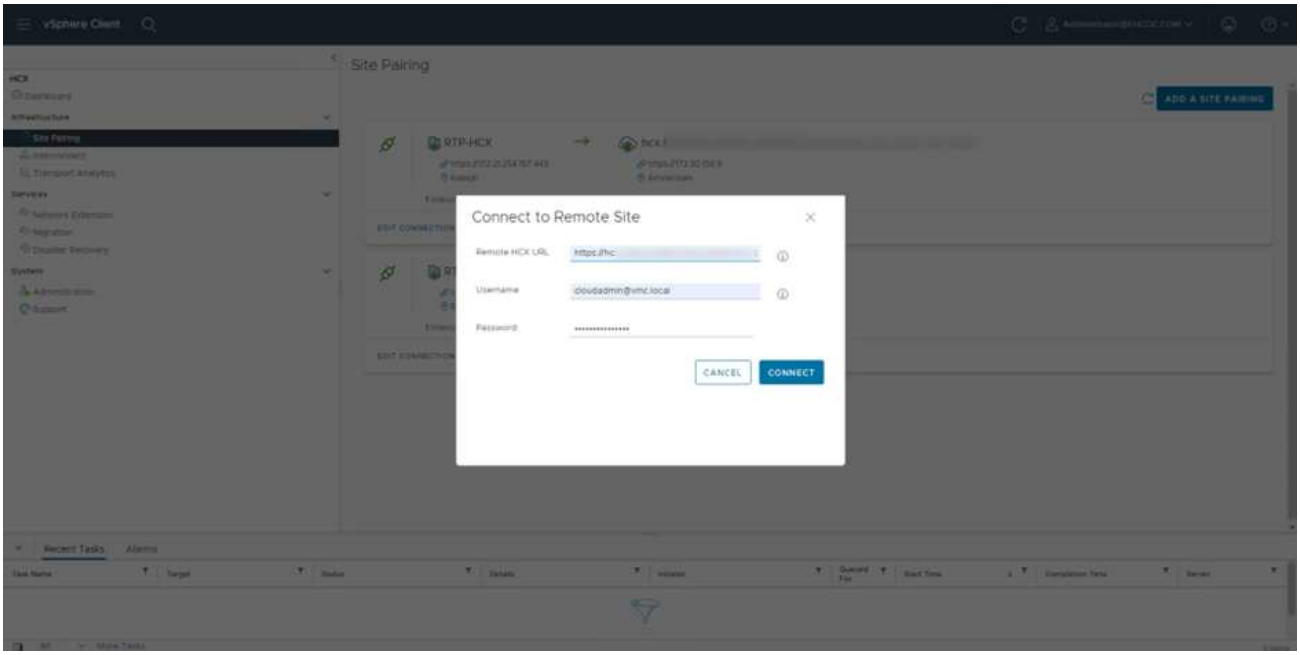


2. Under Infrastructure, click Add a Site Pairing. To authenticate the remote site, enter the VMC HCX Cloud Manager URL or IP address and the credentials for the CloudAdmin role.



HCX information can be retrieved from the SDDC Settings page.





3. To initiate the site pairing, click Connect.

> VMware HCX Connector must be able to communicate with the HCX Cloud Manager IP over port 443.

4. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.

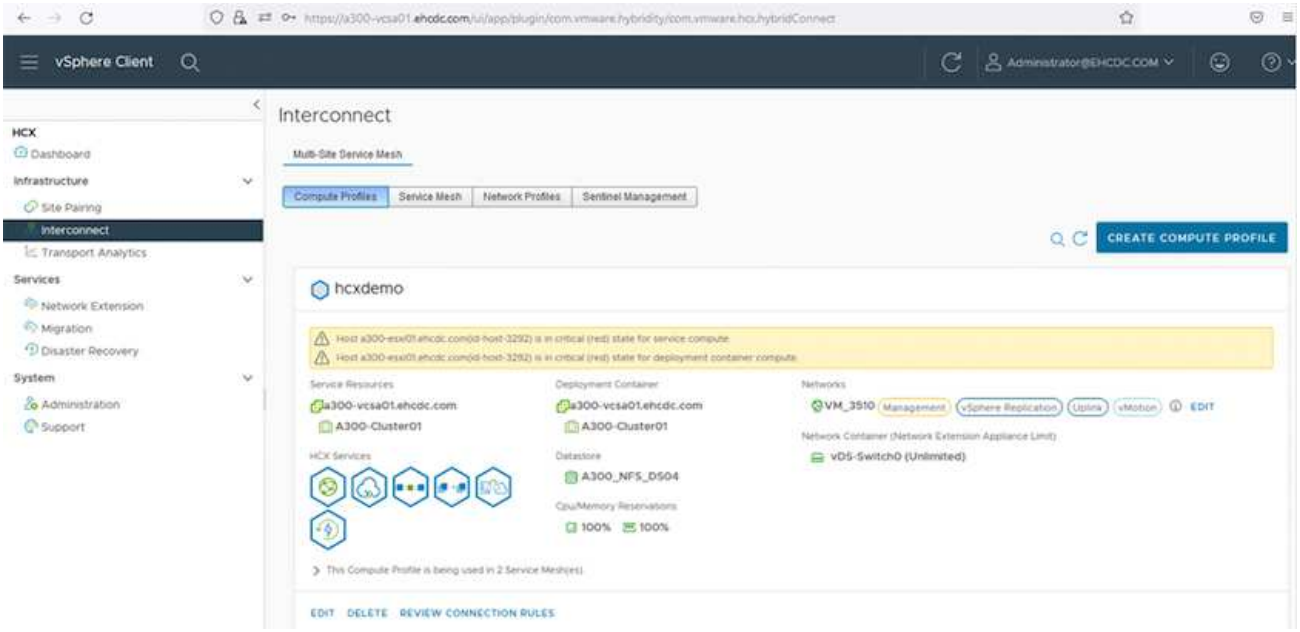
▼ **Step 5: Configure the network profile, compute profile, and service mesh**

The VMware HCX Interconnect (HCX-IX) appliance provides secure tunnel capabilities over the internet and private connections to the target site that enable replication and vMotion-based capabilities. The interconnect provides encryption, traffic engineering, and an SD-WAN. To create the HCI-IX Interconnect Appliance, complete the following steps:

1. Under Infrastructure, select Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile.

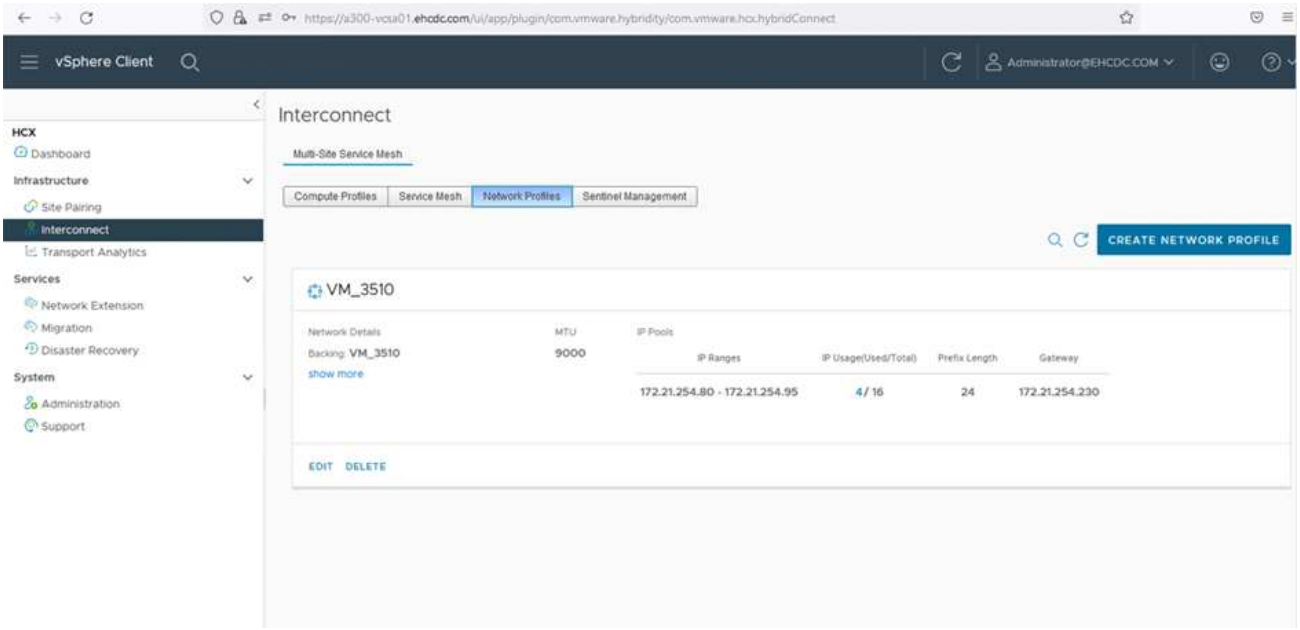
> Compute profiles contain the compute, storage, and network deployment parameters required to deploy an interconnect virtual appliance. They also specify which portion of the VMware data center will be accessible to the HCX service.

For detailed instructions, see Creating a Compute Profile.



2. After the compute profile is created, create the network profile by selecting Multi-Site Service Mesh > Network Profiles > Create Network Profile.
3. The network profile defines a range of IP address and networks that will be used by HCX for its virtual appliances.

> This will require two or more IP address. These IP addresses will be assigned from the management network to virtual appliances.



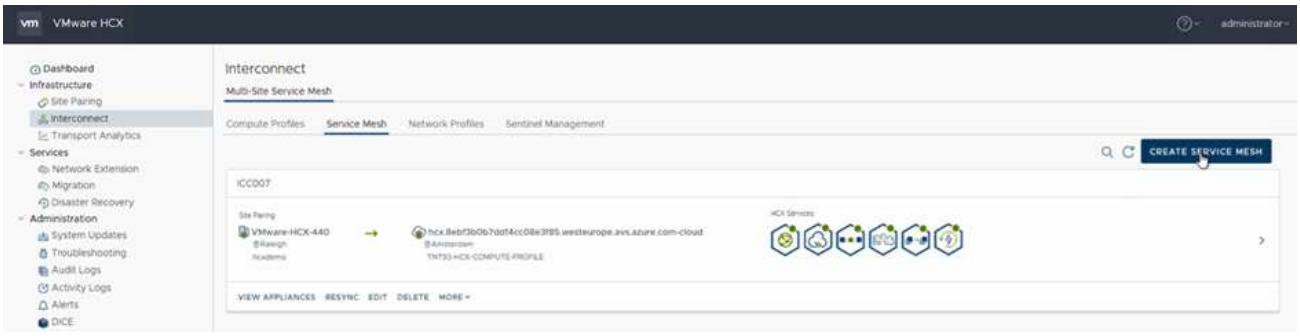
For detailed instructions, see Creating a Network Profile.

> If you are connecting with an SD-WAN over the internet, you have to reserve public IPs under the Networking and Security section.

4. To create a service mesh, select the Service Mesh tab within the Interconnect option and select on-

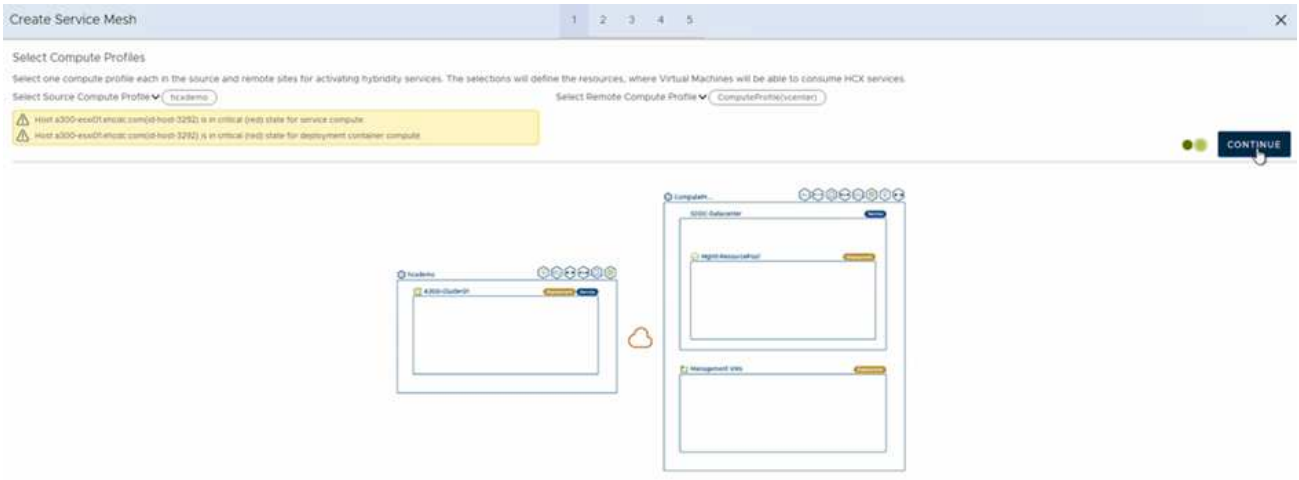
premises and VMC SDDC sites.

The service mesh establishes a local and remote compute and network profile pair.

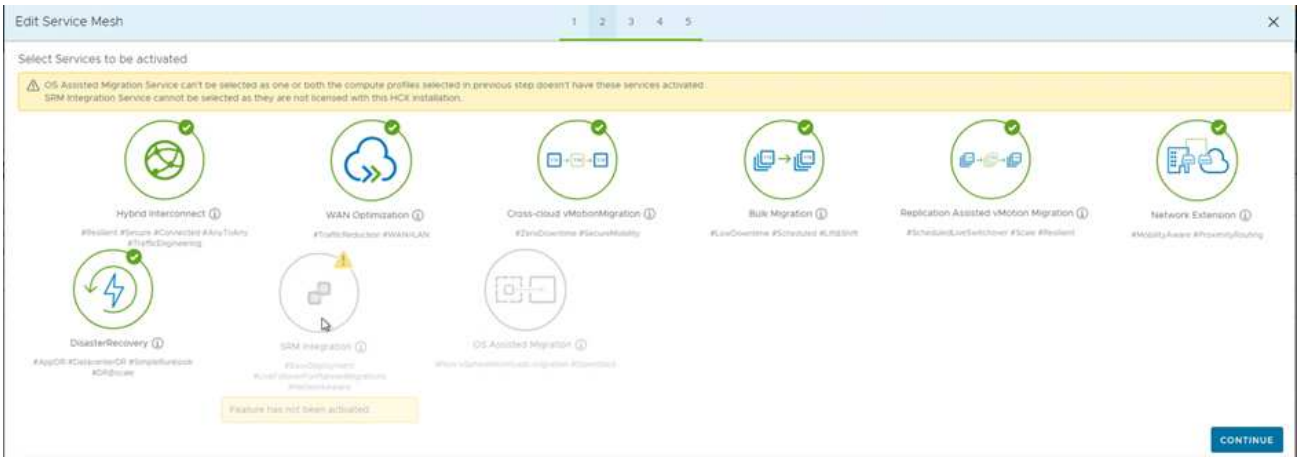


Part of this process involves deploying HCX appliances that will be automatically configured on both the source and target sites, creating a secure transport fabric.

5. Select the source and remote compute profiles and click Continue.

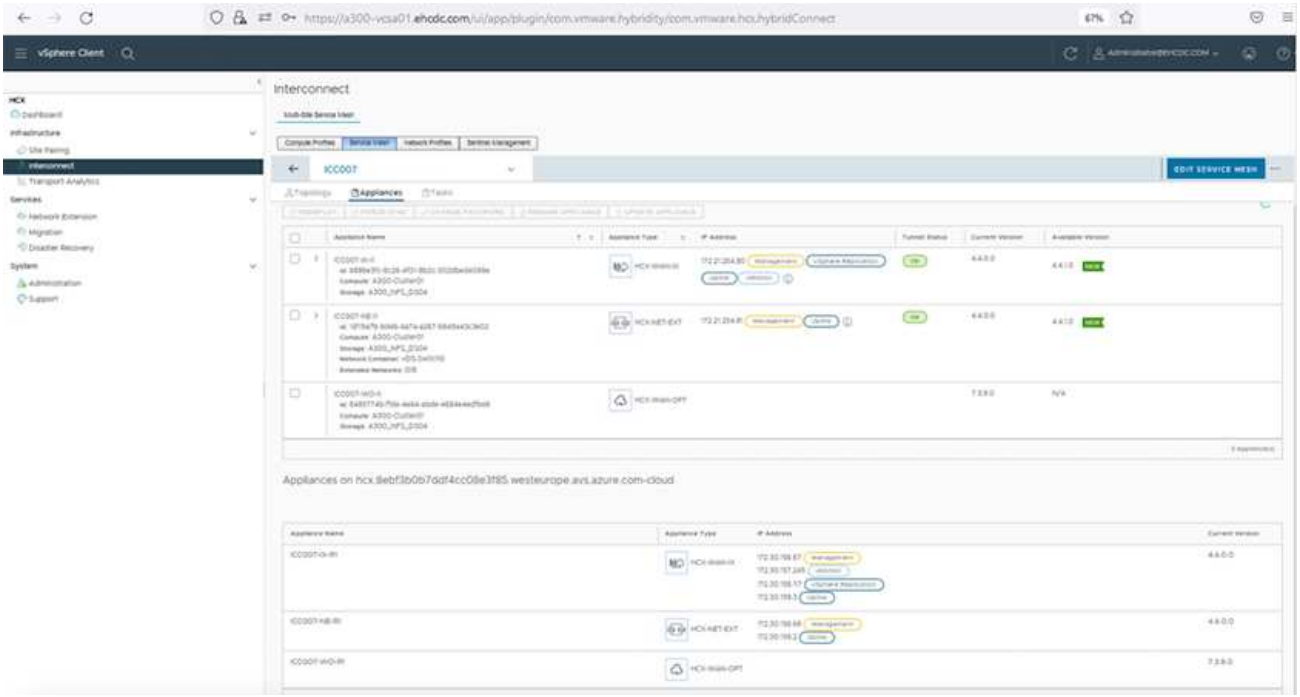


6. Select the service to be activated and click Continue.



An HCX Enterprise license is required for Replication Assisted vMotion Migration, SRM Integration, and OS Assisted Migration.

7. Create a name for the service mesh and click Finish to begin the creation process. The deployment should take approximately 30 minutes to complete. After the service mesh is configured, the virtual infrastructure and networking required to migrate the workload VMs has been created.



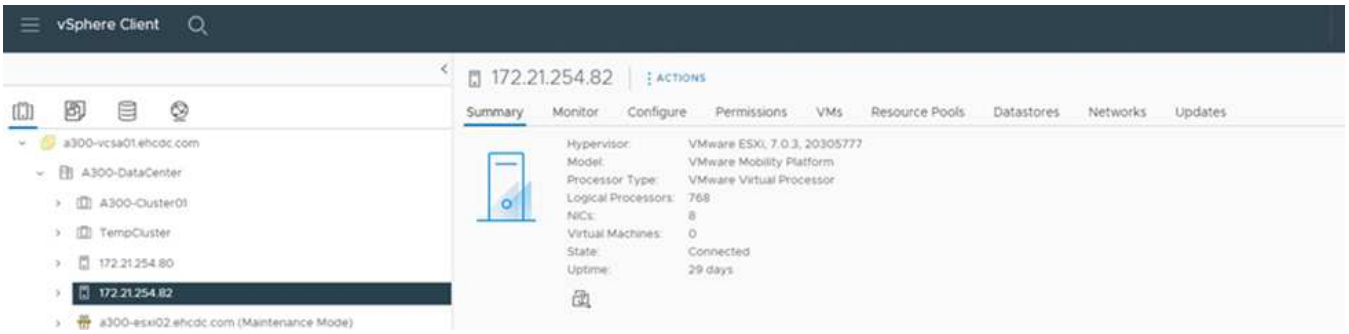
▼ Step 6: Migrating Workloads

HCX provides bidirectional migration services between two or more distinct environments such as on-premises and VMC SDDCs. Application workloads can be migrated to and from HCX activated sites using a variety of migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with HCX Enterprise edition).

To learn more about available HCX migration technologies, see VMware HCX Migration Types

The HCX-IX appliance uses the Mobility Agent service to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations.

The HCX-IX appliance adds the Mobility Agent service as a host object in the vCenter Server. The processor, memory, storage and networking resources displayed on this object do not represent actual consumption on the physical hypervisor hosting the IX appliance.



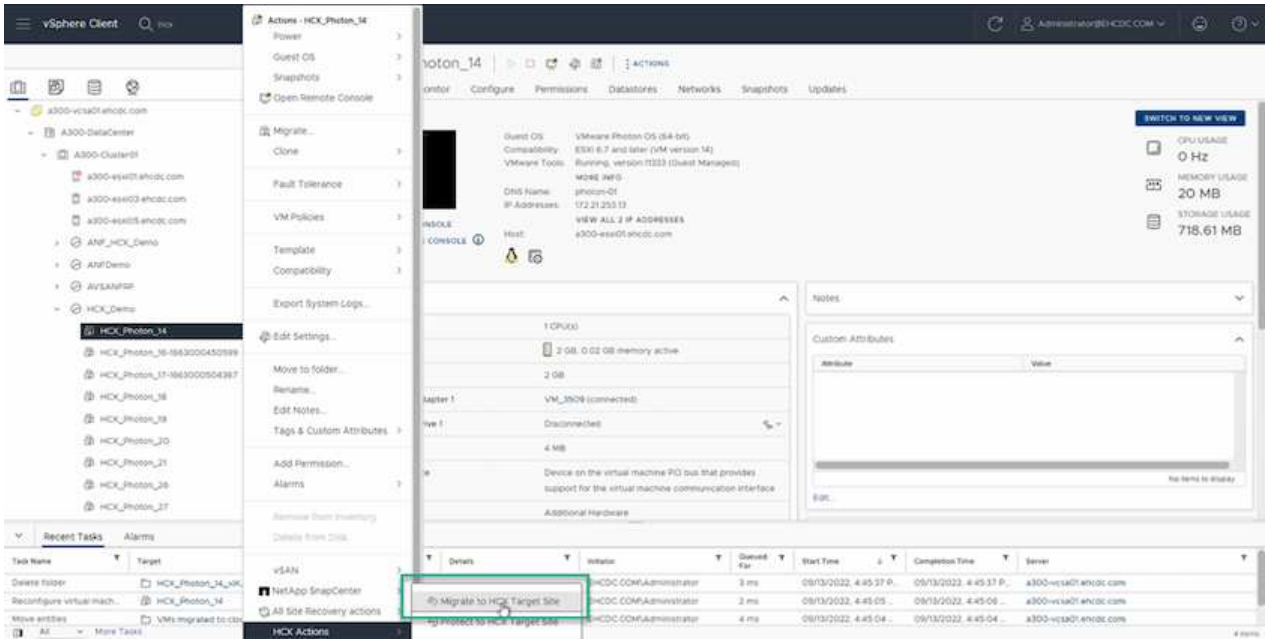
▼ VMware HCX vMotion

This section describes the HCX vMotion mechanism. This migration technology uses the VMware

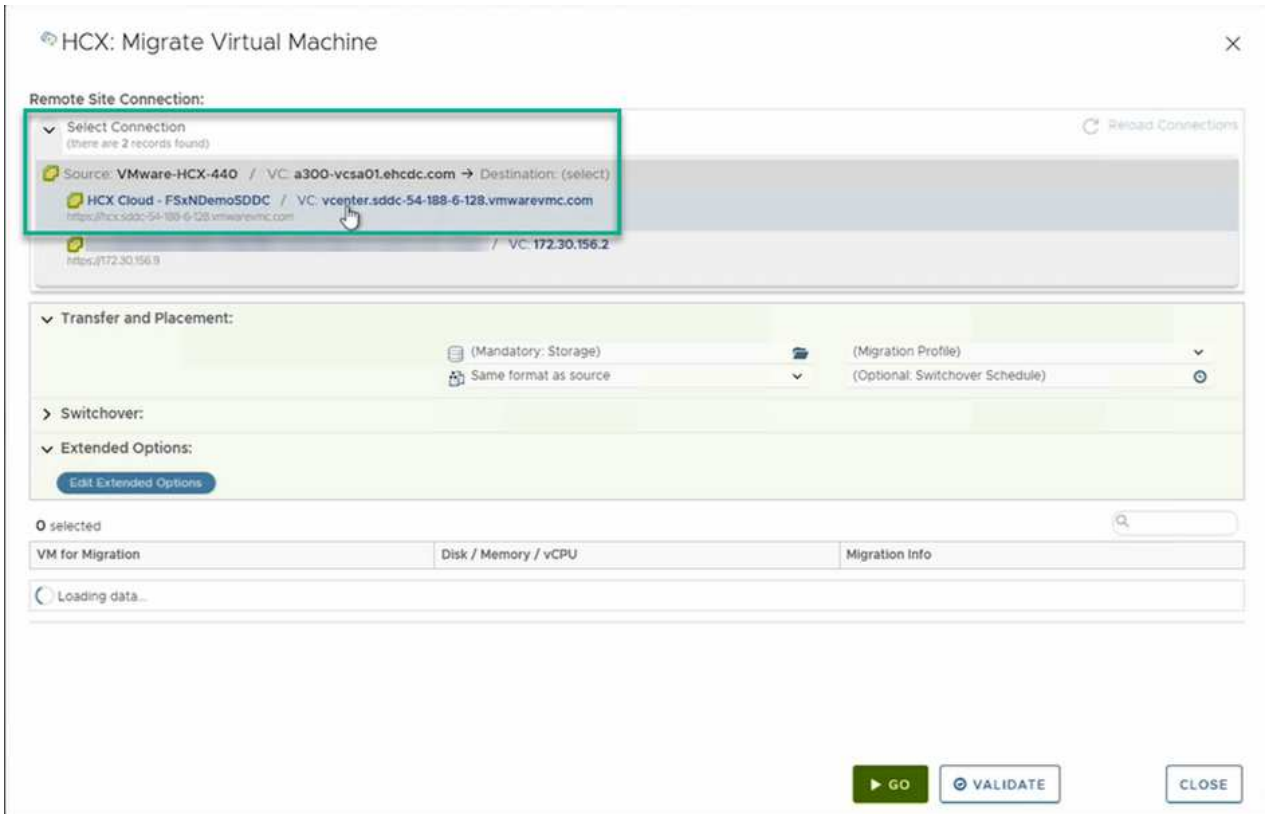
vMotion protocol to migrate a VM to VMC SDDC. The vMotion migration option is used for migrating the VM state of a single VM at a time. There is no service interruption during this migration method.

Network Extension should be in place (for the port group in which the VM is attached) in order to migrate the VM without the need to make an IP address change.

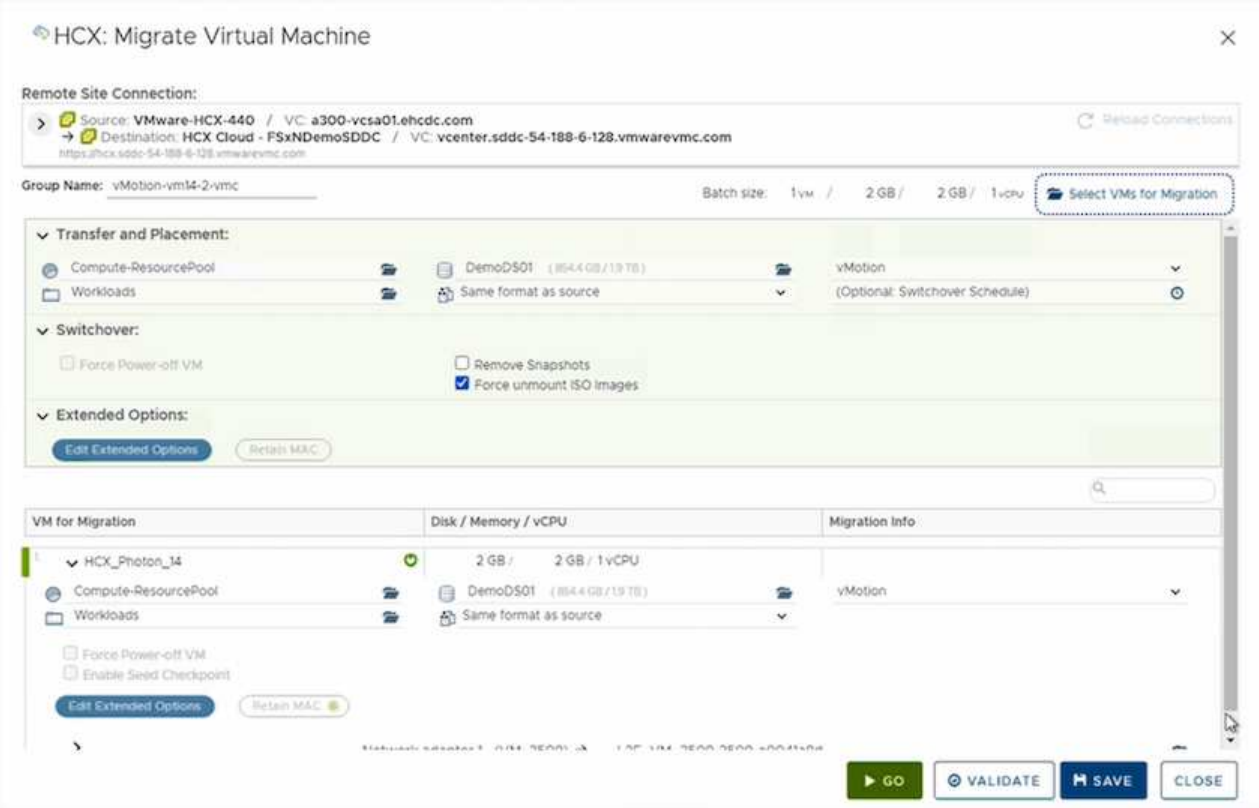
1. From the on-premises vSphere client, go to Inventory, right- click on the VM to be migrated, and select HCX Actions > Migrate to HCX Target Site.



2. In the Migrate Virtual Machine wizard, select the Remote Site Connection (target VMC SDDC).



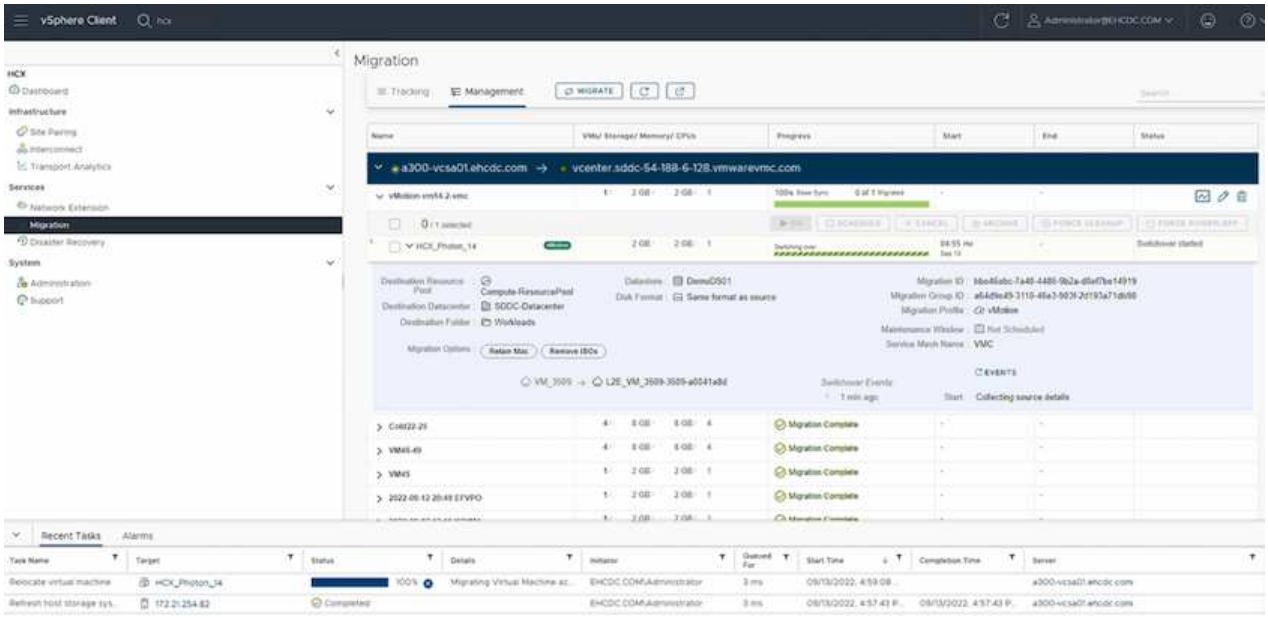
3. Add a group name and under Transfer and Placement, update the mandatory fields (Cluster, Storage, and Destination Network), Click Validate.



4. After the validation checks are complete, click Go to initiate the migration.

> The vMotion transfer captures the VM active memory, its execution state, its IP address, and its MAC address. For more information about the requirements and limitations of HCX vMotion, see Understanding VMware HCX vMotion and Cold Migration.

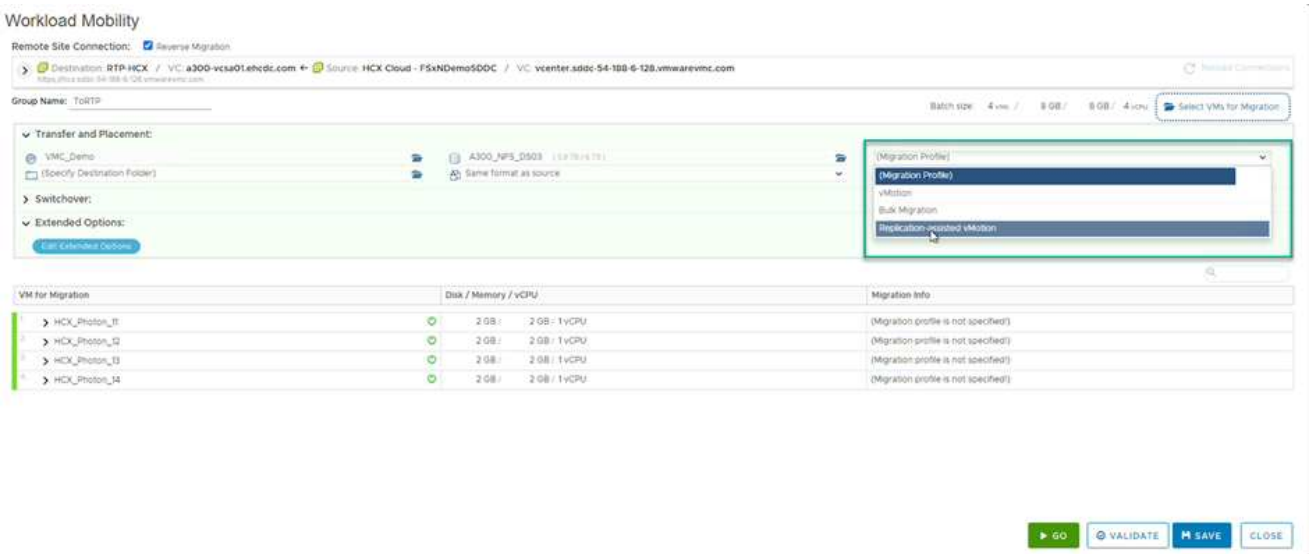
5. You can monitor the progress and completion of the vMotion from the HCX > Migration dashboard.



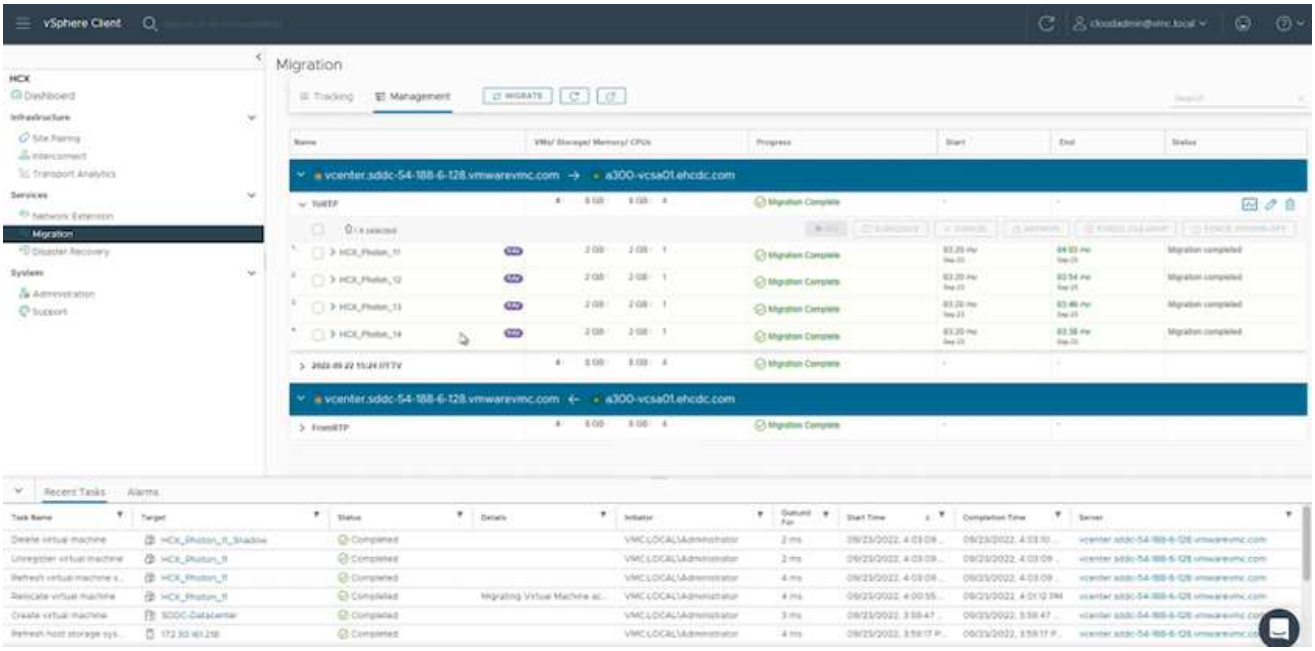
▼ **VMware Replication Assisted vMotion**

As you might have noticed from VMware documentation, VMware HCX Replication Assisted vMotion (RAV) combines the benefits of bulk migration and vMotion. Bulk migration uses vSphere Replication to migrate multiple VMs in parallel—the VM gets rebooted during switchover. HCX vMotion migrates with no downtime, but it is performed serially one VM at a time in a replication group. RAV replicates the VM in parallel and keeps it in sync until the switchover window. During the switchover process, it migrates one VM at a time with no downtime for the VM.

The following screenshot show the migration profile as Replication Assisted vMotion.



The duration of the replication might be longer compared to the vMotion of a small number of VMs. With RAV, only sync the deltas and include the memory contents. The following is a screenshot of the migration status—it shows how the start time of the migration is the same and the end time is different for each VM.



For additional information about the HCX migration options and on how to migrate workloads from on-premises to VMware Cloud on AWS using HCX, see the VMware HCX User Guide.

VMware HCX vMotion requires 100Mbps or higher throughput capability.

The target VMC FSx ONTAP datastore must have sufficient space to accommodate the migration.

Conclusion

Whether you are targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in on-premises, Amazon FSx ONTAP along with HCX provide excellent options to deploy and migrate the workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose VMC along with FSx ONTAP datastore for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere replication, VMware vMotion or even NFC copy.

Takeaways

The key points of this document include:

- You can now use Amazon FSx ONTAP as a datastore with VMC SDDC.
- You can easily migrate data from any on-premises datacenter to VMC running with FSx ONTAP datastore
- You can easily grow and shrink the FSx ONTAP datastore to meet the capacity and performance requirements during migration activity.

Where to find additional information

To learn more about the information described in this document, refer to the following website links:

- VMware Cloud documentation

https://docs.vmware.com/en/VMware-Cloud-on-AWS/

- Amazon FSx ONTAP documentation

  https://docs.aws.amazon.com/fsx/latest/ONTAPGuide

  VMware HCX User Guide

- https://techdocs.broadcom.com/us/en/vmware-cis/hcx/vmware-hcx/4-10/vmware-hcx-user-guide-4-10.html

NetApp for AWS / EVS

Mount FSx ONTAP as external datastore for EVS

How to mount Amazon FSx for NetApp ONTAP as an external datastore for Amazon Elastic VMware Service

This document will present detailed steps for creating an FSx for ONTAP file system and attaching it for use as either an Internet Small Computer System Interface (iSCSI) or Network File System (NFS) datastore for use with Amazon EVS.

Introduction

Organizations with existing on-premises VMware investments looking for efficient and cost-effective ways to modernize their deployments by migrating, bursting, extending, or protecting (backup, disaster recovery) copies on AWS have faced several challenges. With the release of Amazon Elastic VMware Service (Amazon EVS) there is now a native AWS service that makes it possible to run VMware Cloud Foundation (VCF) directly in an Amazon Virtual Private Cloud (VPC).

Amazon EVS makes it possible to extend existing VCF deployments to AWS, burst into AWS to expand performance during peak usage, and entirely run VCF workloads on AWS to facilitate modernizations, data center exits, consolidations, and other use cases.

Although Amazon EVS provides strong benefits, the default hyper-converged architecture couples compute and storage, preventing independent scale of one or the other as business needs dictate. In today's data-centric world, storage demands in many organizations are higher than those for compute. However, with the VMware hyper-converged architecture, storage resources are coupled with compute resources. That means scaling storage can only be done by adding compute hosts that will never be used. It's a situation that can raise costs considerably, especially in storage-heavy workloads.

To solve this issue, AWS and NetApp offer Amazon FSx for NetApp ONTAP (FSx for ONTAP) as the external storage service for Amazon EVS. FSx for ONTAP is an AWS-native storage service that delivers the high performance, data resilience, and cost-efficiency features of NetApp ONTAP software. By using FSx for ONTAP as external storage, Amazon EVS users can scale compute and storage resources independently, allowing for optimal total cost of ownership (TCO).

Overview

This document will present detailed steps for creating an FSx for ONTAP file system and attaching it for use as either an Internet Small Computer System Interface (iSCSI) or Network File System (NFS) datastore for use with Amazon EVS.

Purpose of this Guide

This guide provides instructions on how to set up an FSx for ONTAP datastore for Amazon EVS, as well as best practices, supportability requirements, sizing considerations, and other information to help you plan, design, and implement the integration.

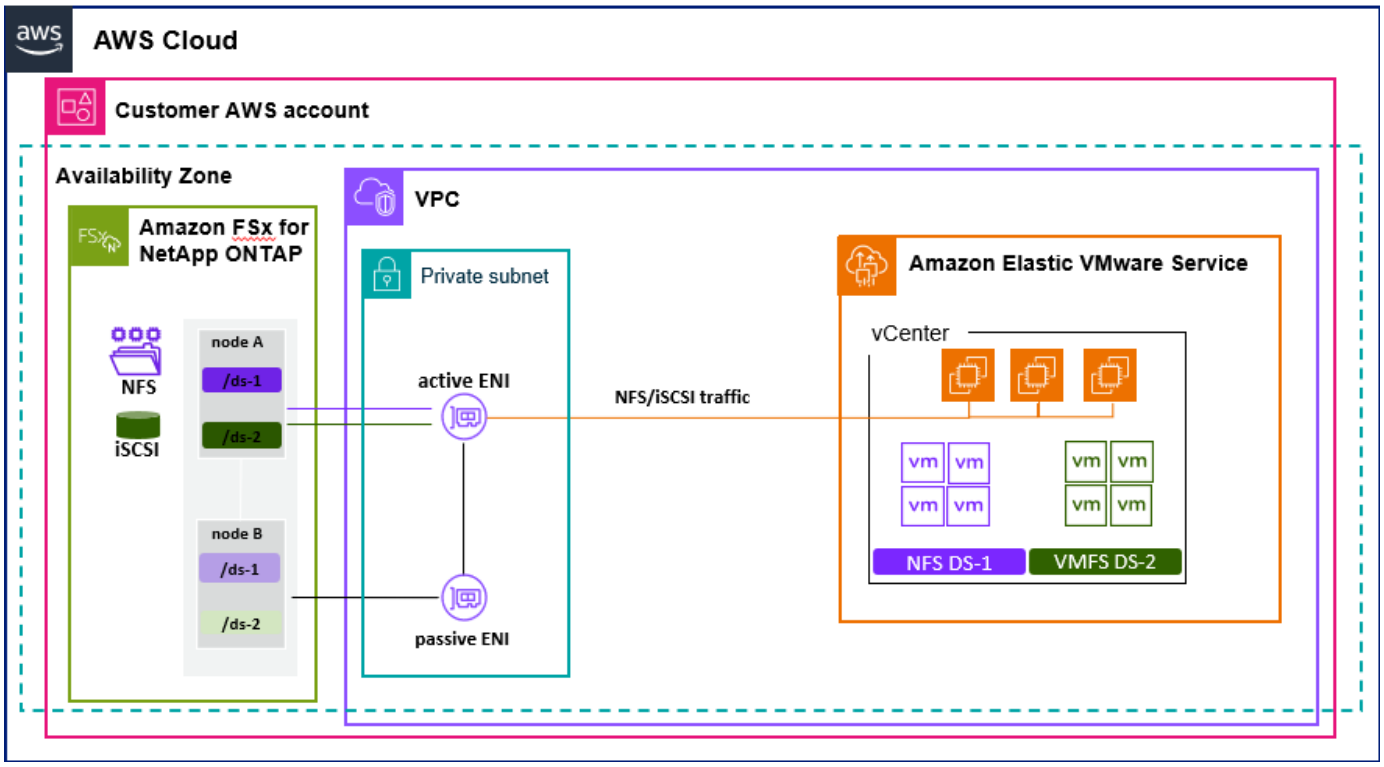
Audience

This tutorial is intended for Cloud Architects and Cloud Administrators familiar with VMware vSphere, Amazon EVS, the AWS Console, and FSx for ONTAP.

High-level architecture

In this section, we will give an overview of the solution architecture for using FSx for ONTAP as a datastore for Amazon EVS.

Below you can see a high-level view of the architecture of the FSx for ONTAP datastore with an Amazon EVS environment within an AWS VPC.



Deployment procedure for using FSx for ONTAP as an NFS datastore

This section will show the steps on how to deploy FSx for ONTAP as an NFS datastore. The steps for deploying an iSCSI datastore are detailed in the next section.

High-level deployment steps

Deploying the FSx for ONTAP file system for use as an NFS datastore for Amazon EVS and establishing connectivity between the two services will include the following steps:

- Setting up an FSx for ONTAP file system in the same VPC where the Amazon EVS VCF environment is located.

- Validating the security groups and routing table.
- Attaching an FSx for ONTAP storage volume for use as an NFS datastore for Amazon EVS.

Prerequisites

The following prerequisites should be in place before you begin the process to set up an FSx for ONTAP as an NFS datastore for Amazon EVS:

- An Amazon EVS virtualization environment. For more information on how to do that, see the documentation on Getting started with Amazon EVS guide.
- Create Single-Availability Zone (AZ) FSx for ONTAP with appropriate SSD storage capacity, provisioned IOPS and throughput.
- For optimal performance, place FSx for ONTAP file system within the same VPC as the Amazon EVS instance. Note that VPC Peering is also supported for the case where FSx for ONTAP is deployed in a different VPC.
- Based on the performance requirements, select the correct IOPS and throughput needed for a Single-AZ FSx for ONTAP file system. See performance specification details for FSx for ONTAP to understand the throughput and IOPS each available option provides.



For this guide we will deploy the FSx for ONTAP file system in the same subnet as the service access subnet.

Once the Amazon EVS virtualization environment has been deployed in the VPC, you can proceed with setting up the Single-AZ FSx for ONTAP file system that will be used as an NFS datastore.

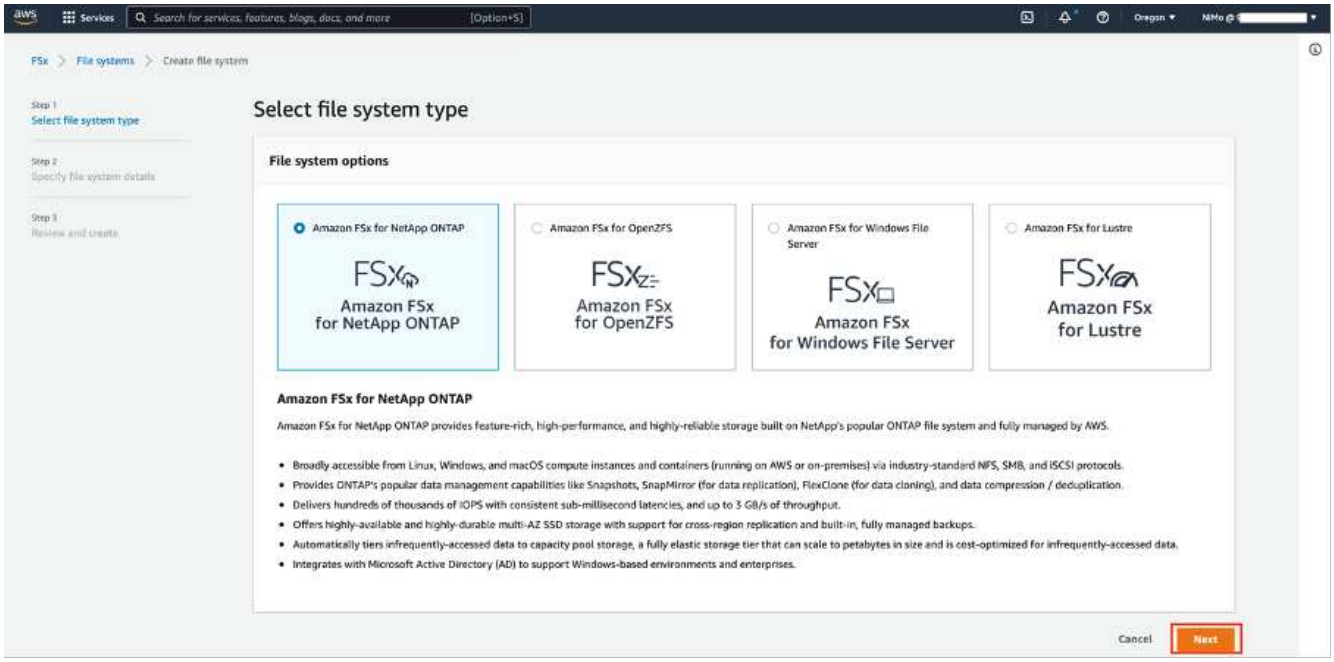


Your FSx for ONTAP file system should be deployed in the same AZ as the Amazon EVS SDDC.

Create an FSx for ONTAP file system

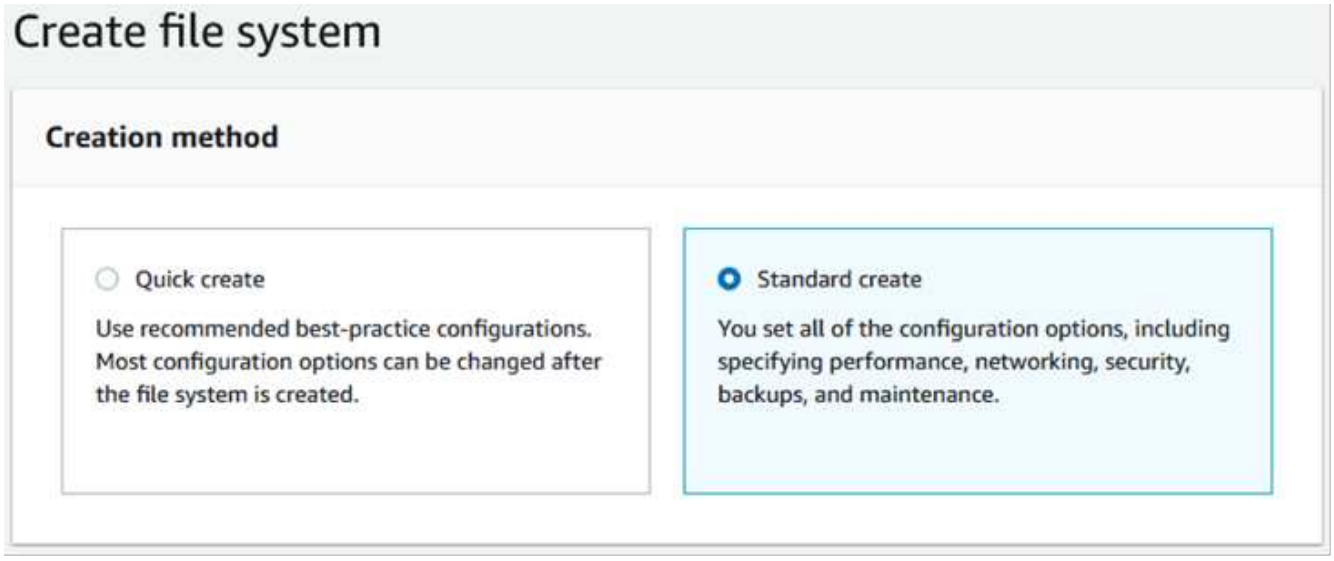
The following steps will show you how to provision and mount an FSx for ONTAP file system:

1. Go to the Amazon FSx Console at `https://console.aws.amazon.com/fsx/` using your internet browser and sign in with your credentials.
2. In the console, select Create file system. This will begin the setup wizard to create an FSx for ONTAP file system.
3. First, select the file system type. From the file system options available, select "Amazon FSx for NetApp ONTAP" and proceed by clicking the "Next" button.

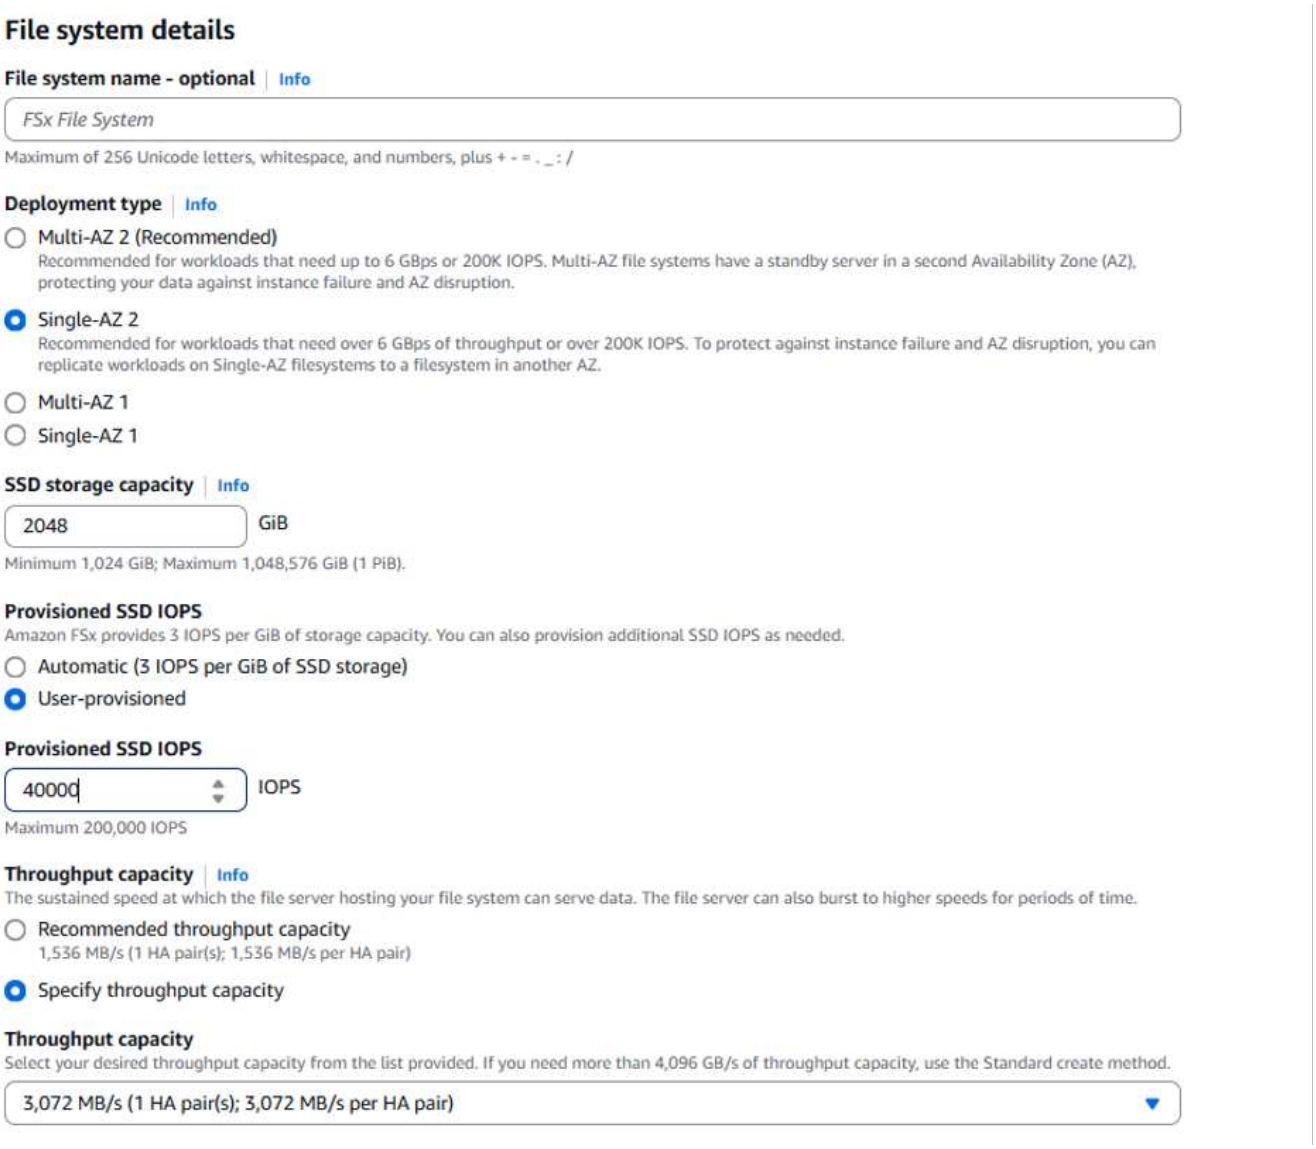


4. You will be presented with a choice between the Quick create and Standard create options. Quick create offers predefined configurations and Standard create lets you define your own configurations.

   Select the "Standard create" option.

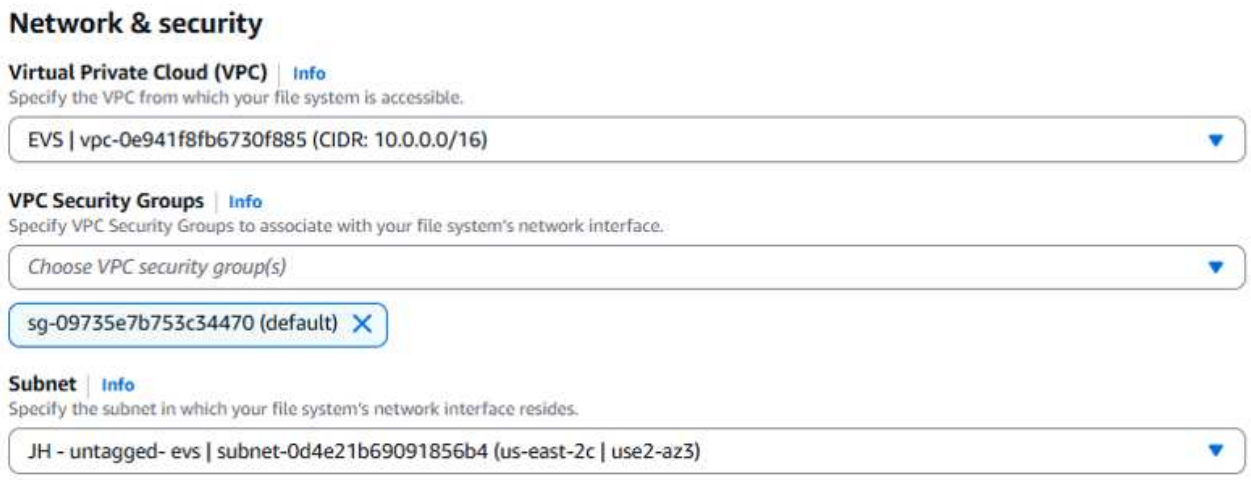


   You will now be in the Standard create file system setup wizard. In the file system details page, choose and enter the same information as shown in the image below:



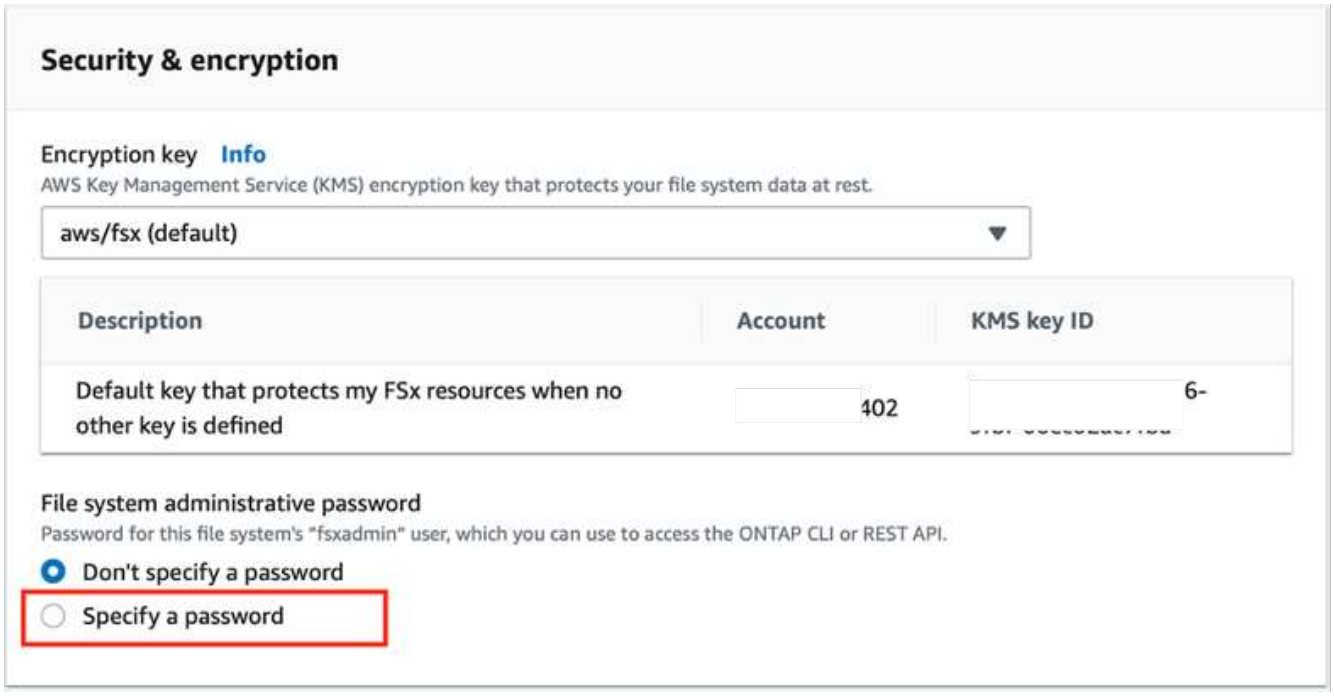
5. Scroll down to where you find a section labelled "Network & security." In this section you will set the details for the subnet and VPC. For the purposes of this guide, we selected "EVS-VPC" from the drop-down menu, and selected service access subnet (untagged VLAN) as the subnet for the EVS virtual environment.

> ⓘ The subnet can be any untagged private subnet in the VPC. It can also be in a different VPC/subnet connected via VPC Peering, as long as proper routing has been configured.

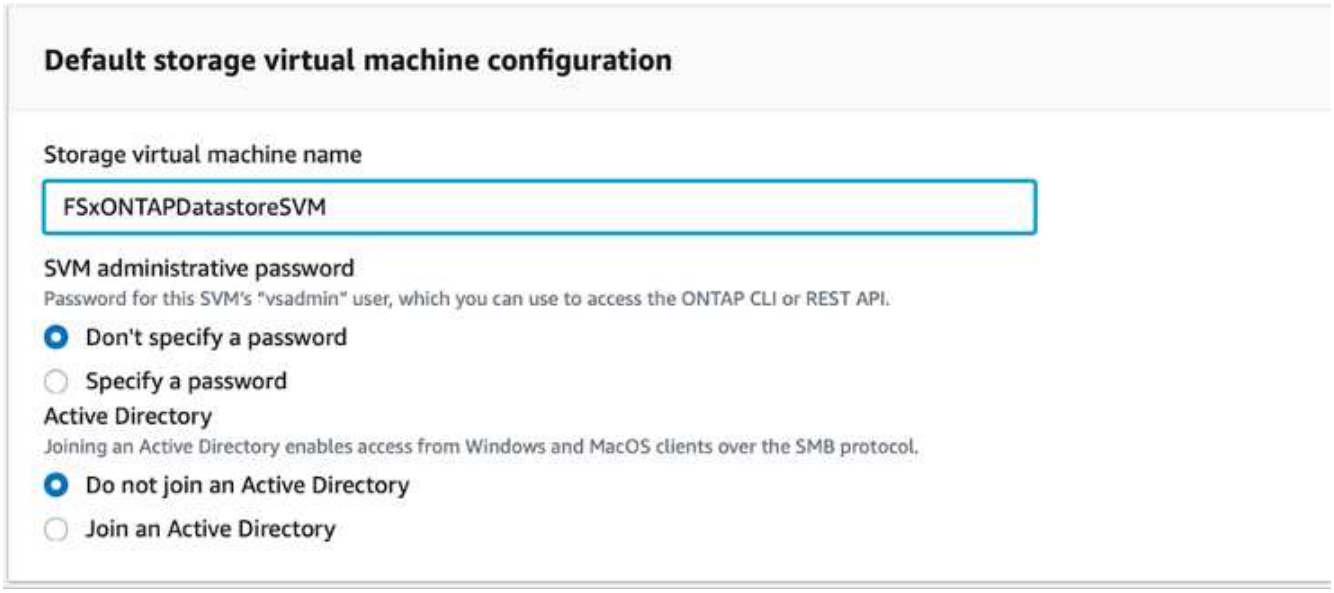


6. Scrolling down, you'll find the "Security & encryption" section. Here, select the encryption key from the drop-down menu. These keys are provided by AWS Key Management Service (AWS KMS) and will protect the at-rest data in your FSx for ONTAP file system.

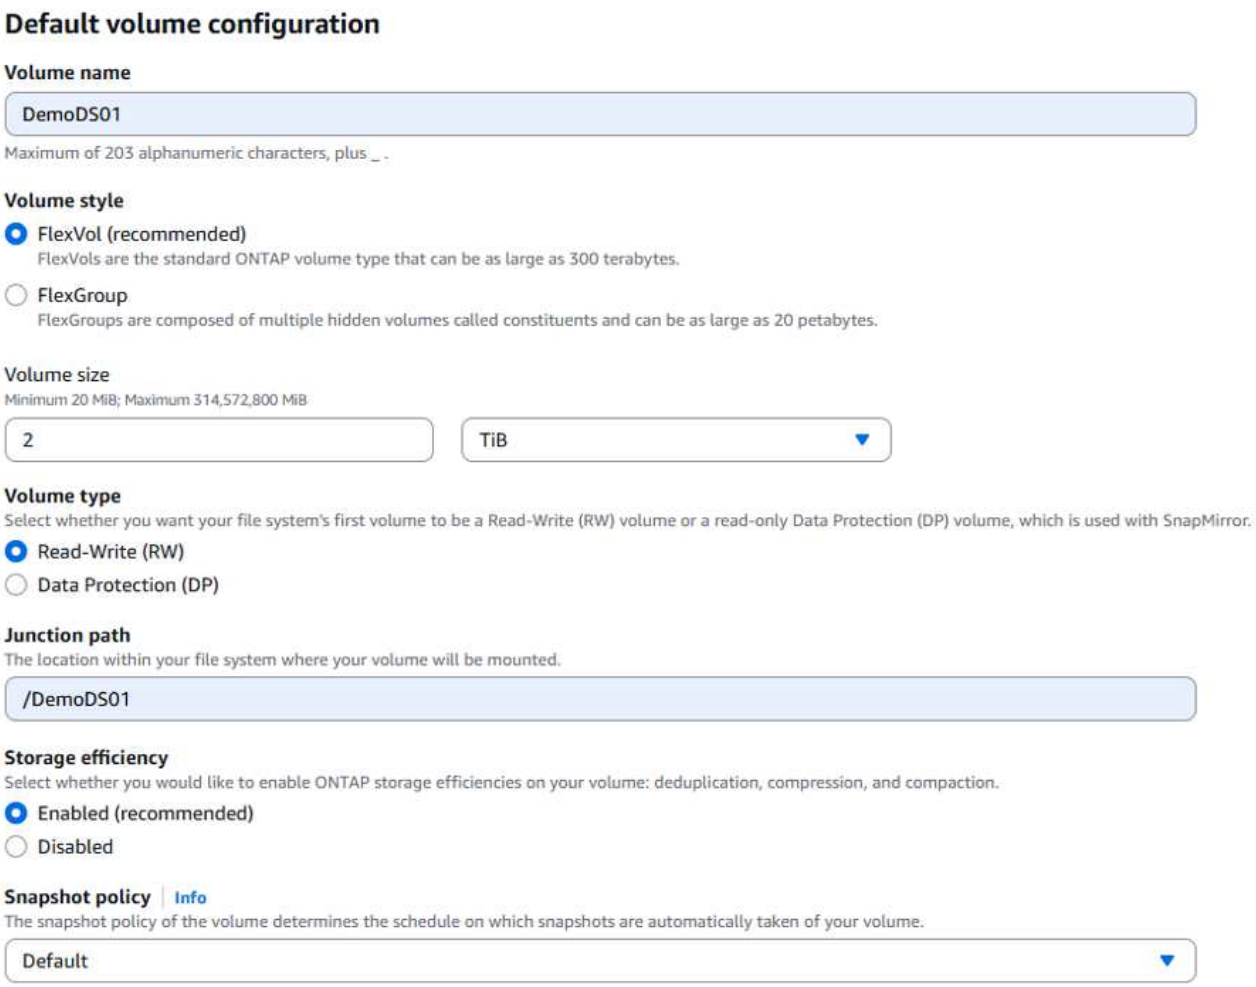
   Next, in the "File system administrative password" setting, specify a unique password for the fsxadmin user who will be administering the file system.



7. Next, you'll scroll down to a section titled "Default storage virtual machine configuration." Here, you will need to provide a name for the storage virtual machine (SVM).

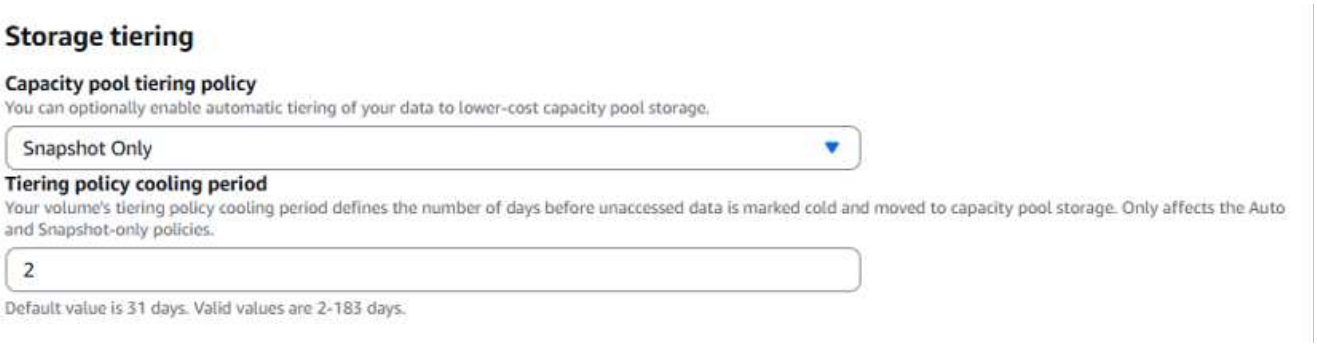


8. Next, scroll down to the section titled "Default volume configuration." Here you will provide a name for the volume and set the volume size. The volume protocol should be NFS v3 (enabled by default). Click the option to enable storage efficiency features (compression, deduplication, and compaction) to reduce your overall storage usage and costs.



9. For the storage tiering option, select "Snapshot Only" from the drop-down menu and set a tiering policy

cooling period. The default is 31 days, but for this example we set a two-day cooling period.

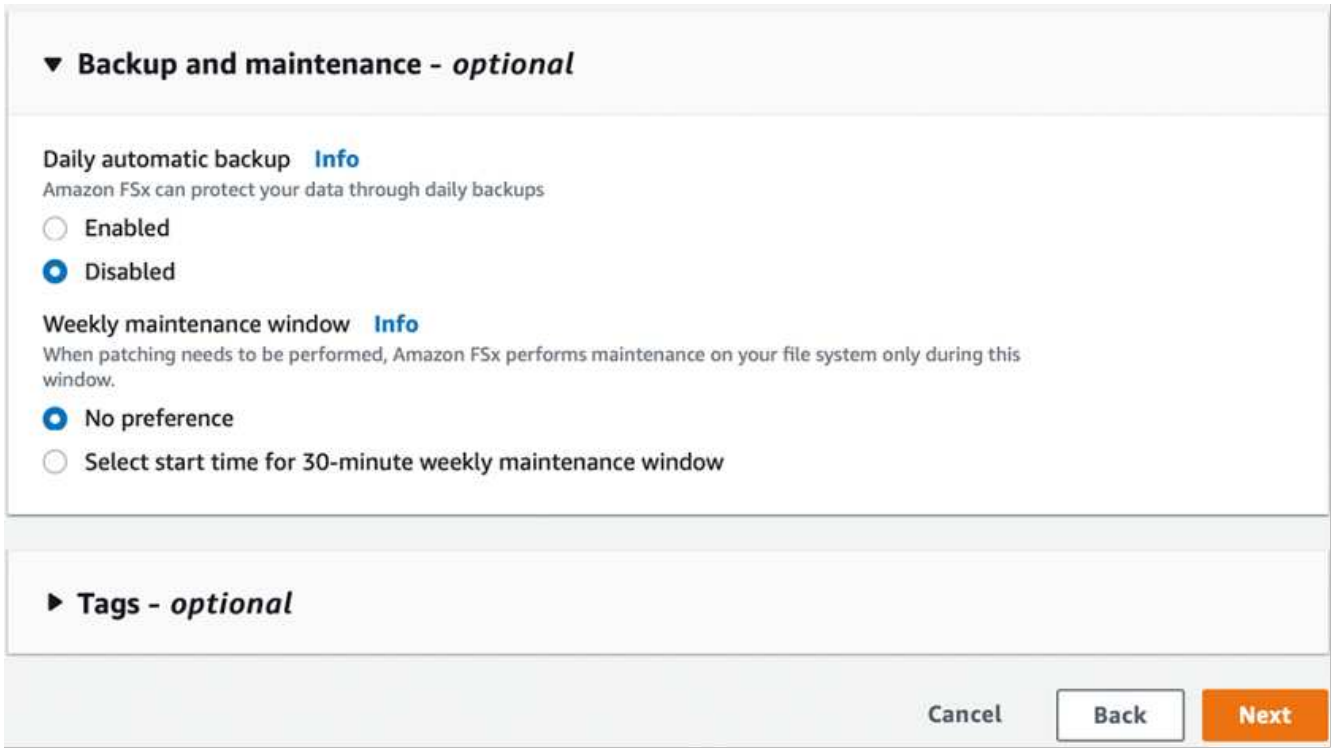


10. Once you've created the volume, modify the following parameters using the ONTAP command line interface (ONTAP CLI):

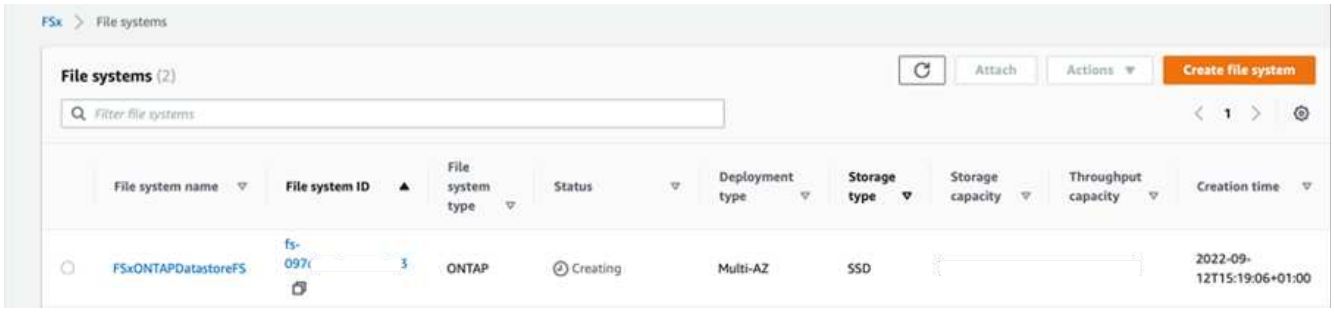
Setting	Configuration
Volume guarantee	None (thin provisioned)
fractional_reserve	0%
snap_reserve	0%
Autosize	grow_shrink
Storage efficiency	Enabled
Volume Tiering Policy	Snapshot only
Access time update (-atime-update)	False
Minimum readahead (-min-readahead)	False
Snapshot policy	None

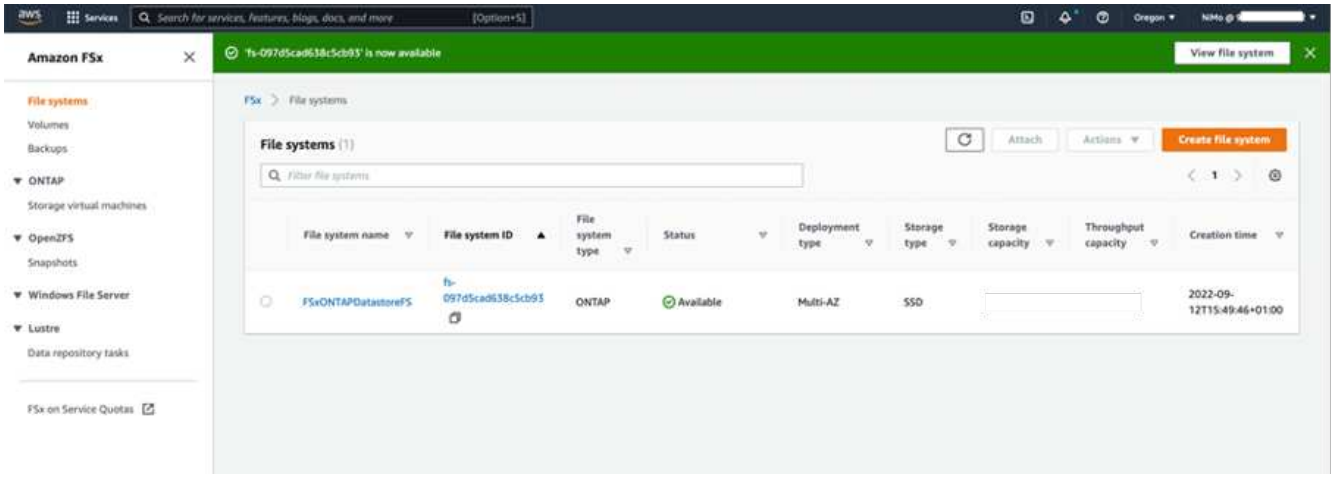
11. In the optional "Backup and maintenance" step, disable daily automatic backups and select "No preference" for the weekly maintenance window.

    Click "Next" when you are done.



12. You will now be provided with a review of the entire file system configuration that you set up following the previous steps. Check that all the parameters are correct, then click the "Create file system" button to complete the set up.





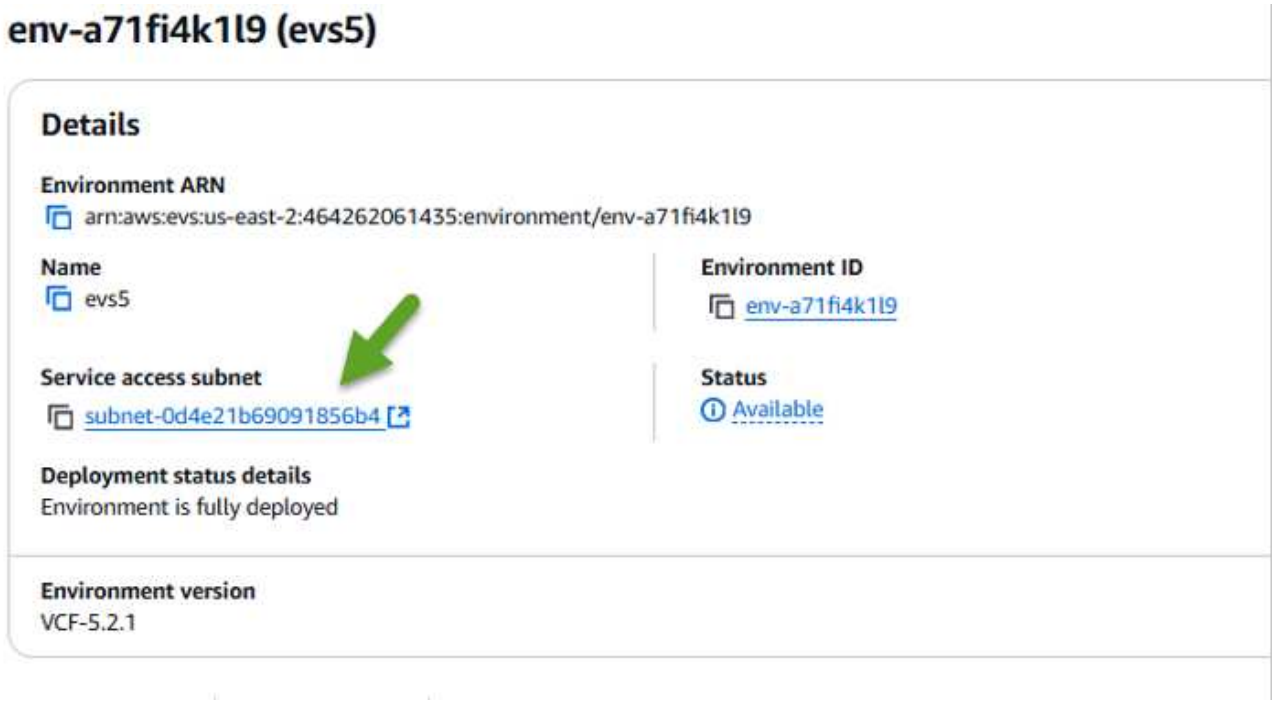
To create additional FSx for ONTAP file systems, you can follow the steps detailed above, adjusting the values that affect performance and disk capacity as needed. To learn more

about performance options for FSx for ONTAP, see this documentation page.

Validating VPC and virtualization environment routing and security groups

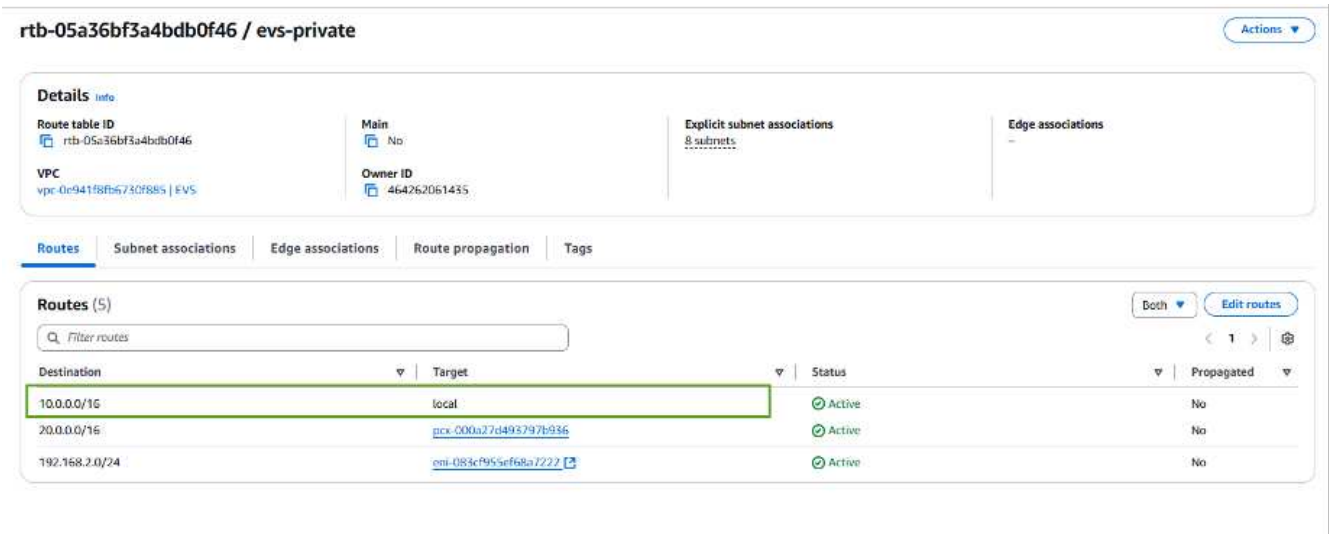
Validating the VPC and SDDC routing and security groups is important to make sure the Amazon EVS virtualization environment and FSx for ONTAP file system components can work together properly.

1. Browse to the Elastic VMware Service > Virtualization Environments > [your Amazon EVS environment] and select the Service access subnet, as indicated by the green arrow below:



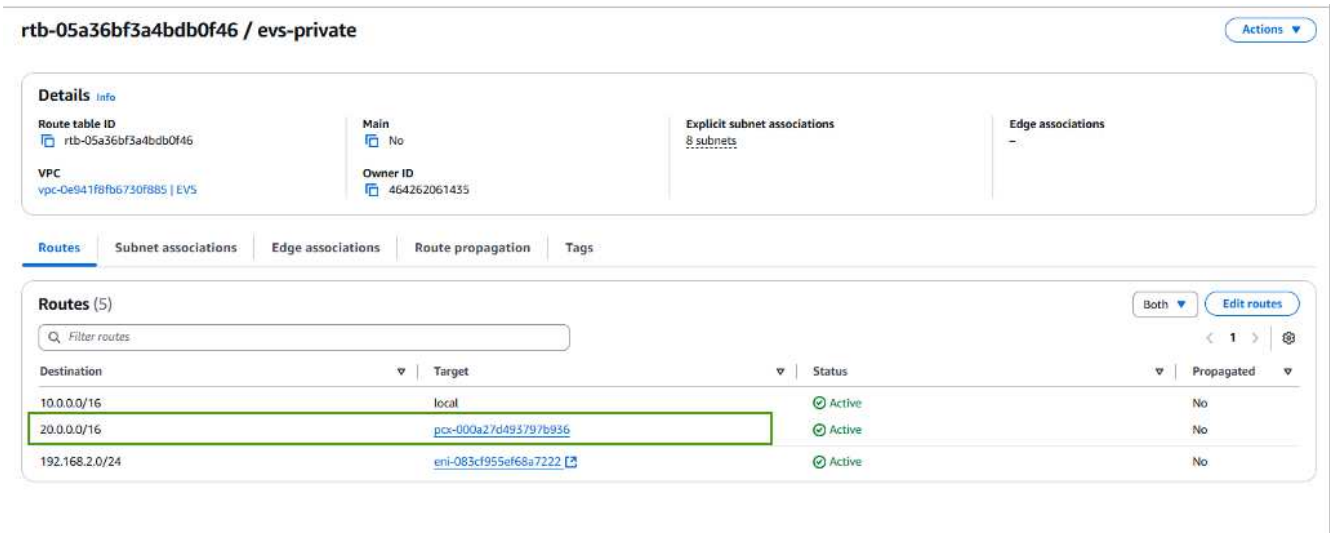
2. This will open the subnet panel. Find the Routes table on this panel and click it.

   In this example the FSxN ONTAP file system is in the same VPC as the Amazon EVS VCF environment, so the default routes are sufficient.



   If FSx for ONTAP file system is in a different VPC, then check that the routes to that VPC have been added. In this example, routes for VPC Peering to the VPC where the FSx for ONTAP system resides have

been added.

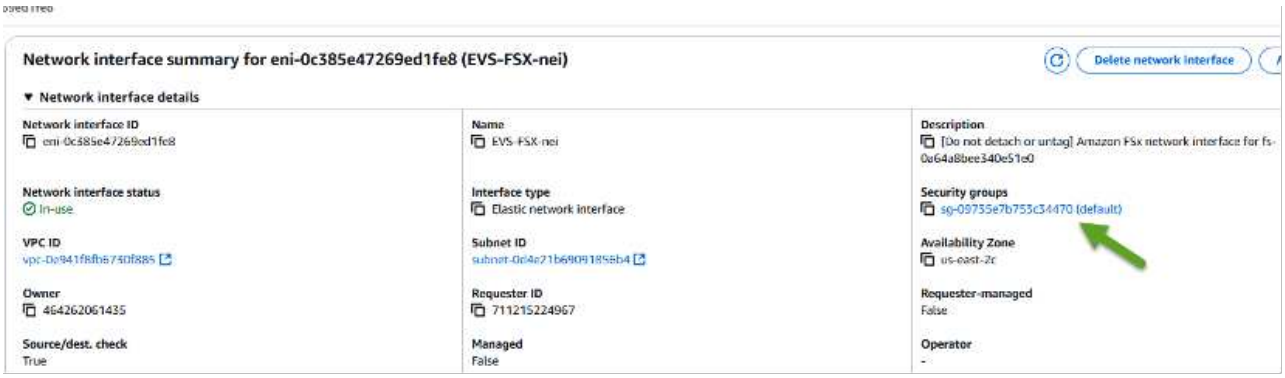


3. Next, you'll need to check that the security group connected to the FSx for ONTAP file system (i.e. the one connected to the elastic network interface (ENI)—allows access to NFS ports).

   For a full list of NFS and iSCSI protocol ports, refer to the link:https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/limit-access-security-groups.html]File System Access Control with Amazon VPC documentation].

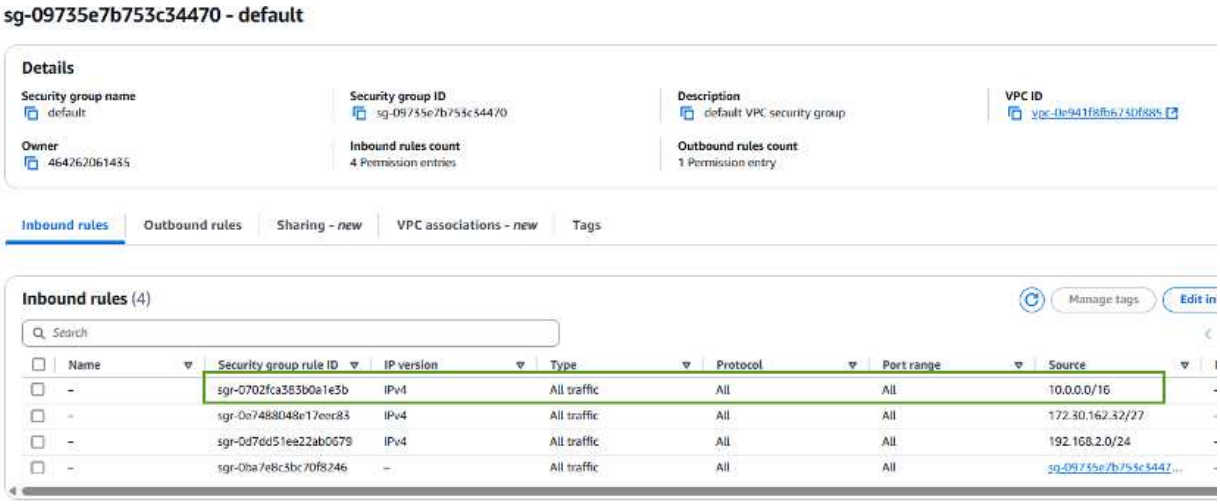
   To verify your FSx for ONTAP file system is associated with the security group:

   a. In the AWS Console, navigate to FSx > File Systems > [your FSx for ONTAP file system].
   b. In the Network & Security tab, click the ENI.
   c. Find the Security group listed under the Network interface details. associated with the ENI (indicated by the green arrow in the image below).



   d. Validate the ports. In this example shown in the following screen, all traffic is allowed. However, you can limit those to NFS ports only. Here is the list of the ports:
      - Port 111 TCP: Portmapper, used to negotiate which ports are used in NFS requests.
      - Port 635 TCP: Mountd, which is for receiving incoming mount requests.
      - Port 2049 TCP: NFS, which handles NFS traffic.
      - Port 4045 TCP: Network Lock Manager (NLM), which deals with lock requests.
      - Port 4046 TCP: Network Status Monitor (NSM), which sends NFS clients notifications when the

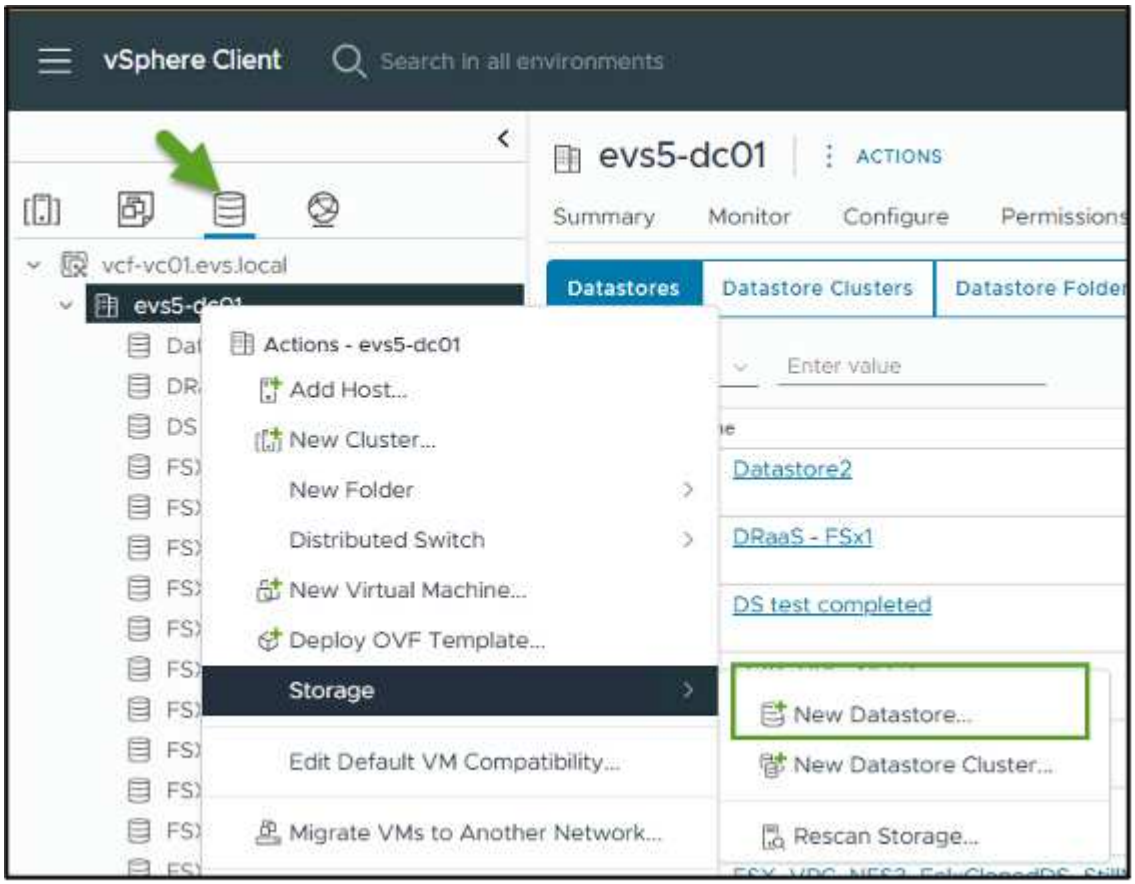
server reboots for lock management.



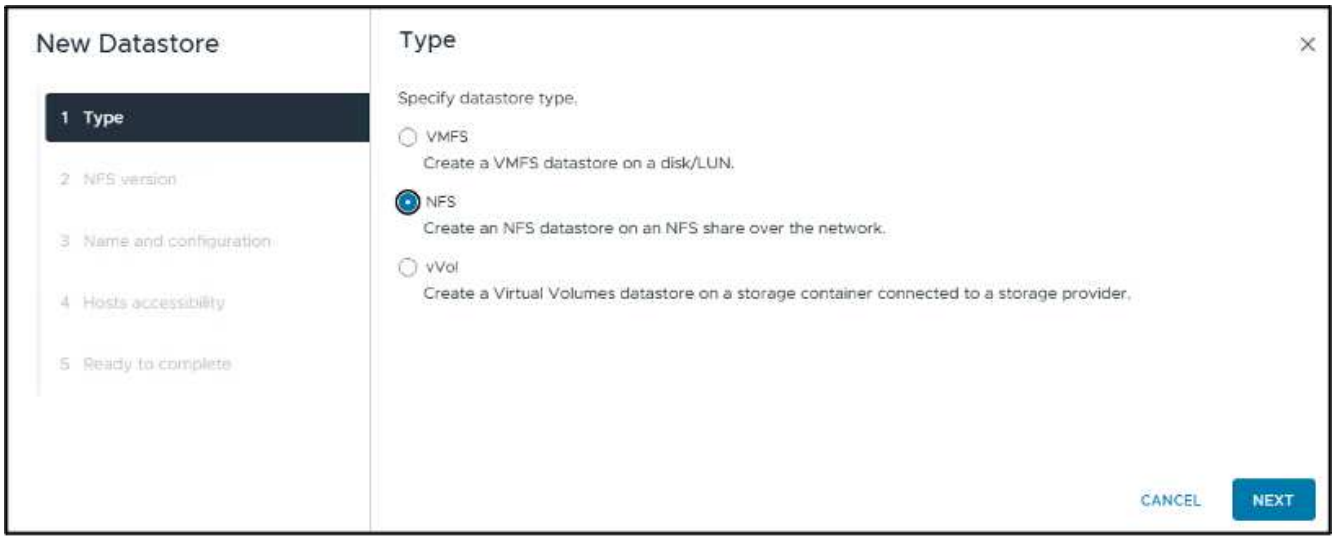
Attaching an FSx for ONTAP NFS volume to Amazon EVS for use as a datastore

Now that your FSx for ONTAP file system has been provisioned and all the proper connectivity parameters have been validated, it is time to attach an FSx for ONTAP storage volume to Amazon EVS.
The following steps will show how to access vCenter and mount the FSx for ONTAP volume as an NFS datastore:

1. In the vSphere client, go to the "Datastore" tab. Find the Datacenter navigate to Storage > New Datastore, as shown below:

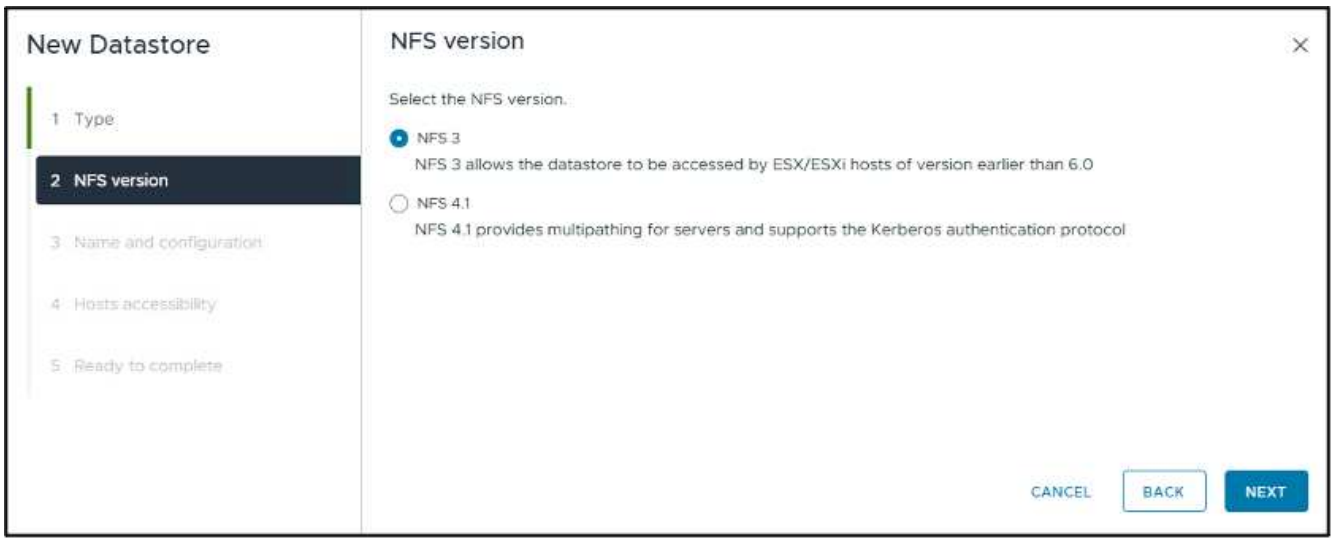


2. You will now be in the New Datastore wizard. In the "Type" step, you will choose a datastore type. Select "NFS" and click "NEXT" to proceed.



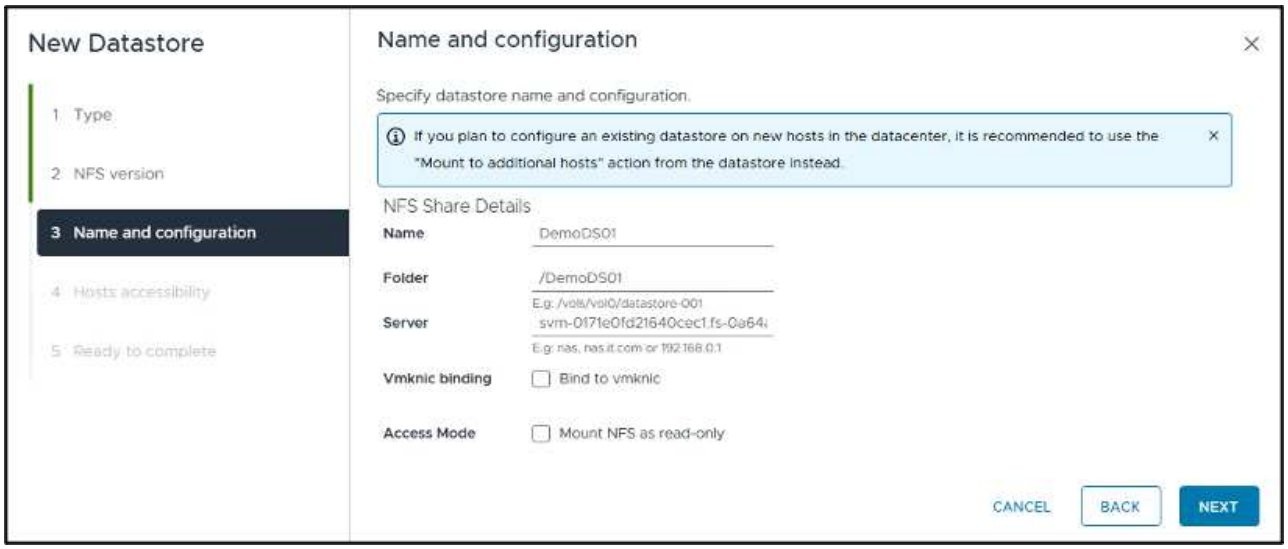
3. In the "NFS version" step, you will select an NFS version.

   For the purposes of this example, we will choose "NFS 3" but note that NFS v4.1 can also be used.

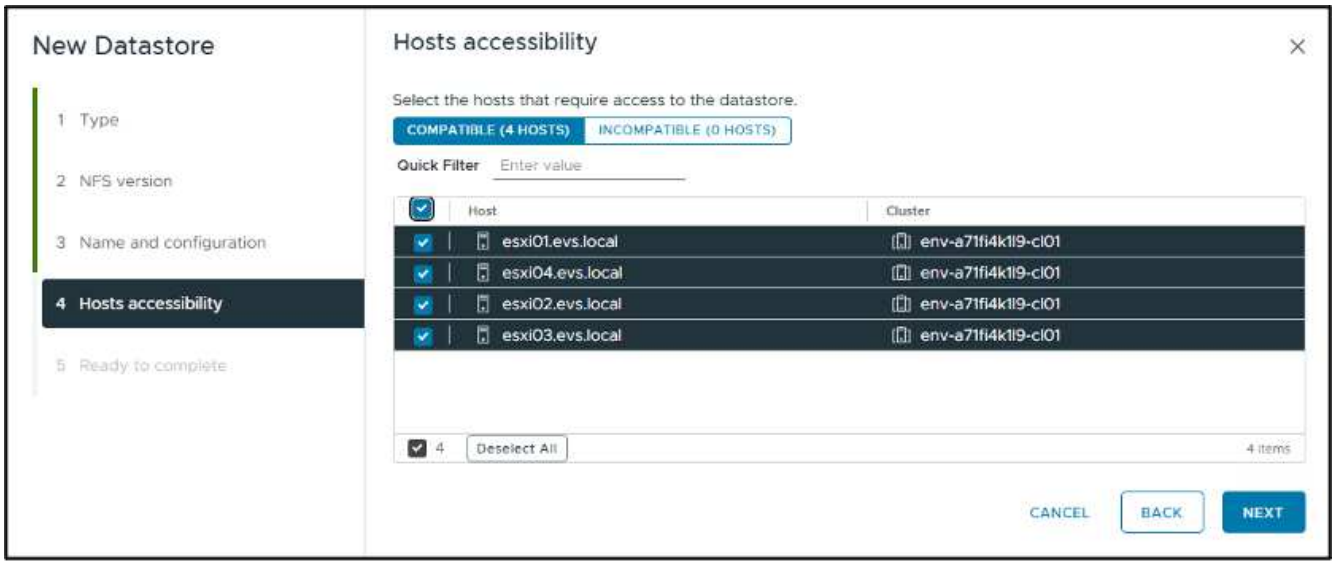


4. In the "Name and configuration" step, you will:
   a. Name the datastore.
   b. Name the folder. Use the junction path to your NFS volume.
   c. Name the server. Use either the NFS DNS name of your SVM or the IP address of the NFS.

      Click "NEXT" to continue.



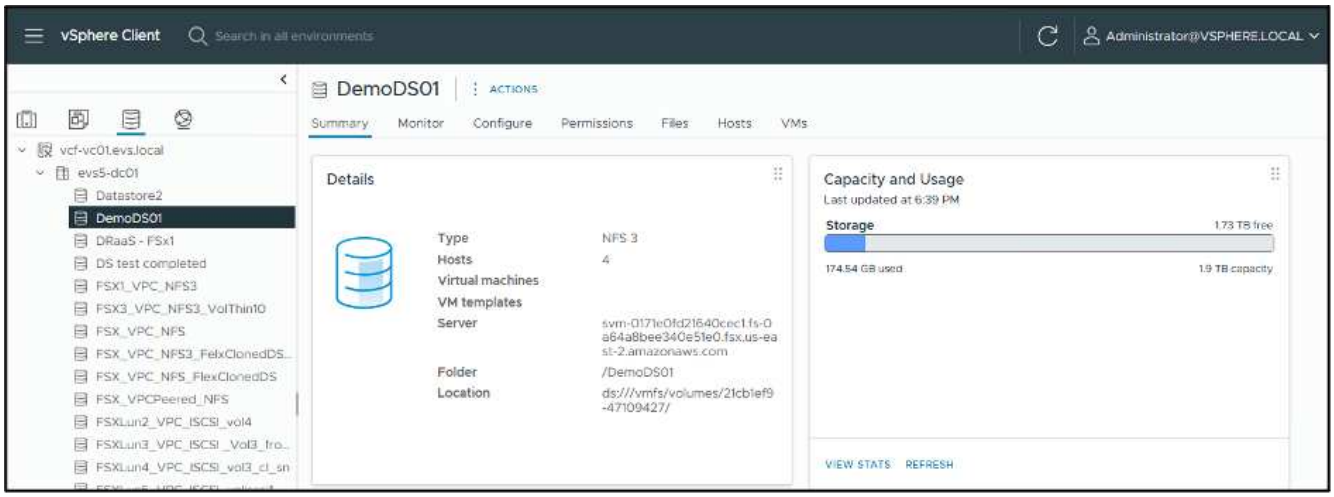
5. In the "Host accessibility" step, select all the hosts that require access to the datastore, then click "NEXT" to continue.



6. In the "Ready to complete" step, review the data and click "FINISH" to complete the set up.



7. Validate that the NFS datastore has been connected by accessing vCenter, as shown below:



Deployment procedure for using FSx for ONTAP as an iSCSI datastore

This section covers the steps to configure an iSCSI datastore using FSx for ONTAP.

Prerequisites

Before you provision and attach FSx for ONTAP as an iSCSI datastore, you must first deploy an Amazon EVS virtualization environment. For more information, see the documentation on Getting started with Amazon EVS guide. This section also assumes you have deployed an FSx for ONTAP file system, please see the Detailed Approach section above for appropriate steps.

High-level deployment steps

Deploying the FSx for ONTAP file system for use as an iSCSI datastore for Amazon EVS and establishing connectivity between the two services will include the following steps:

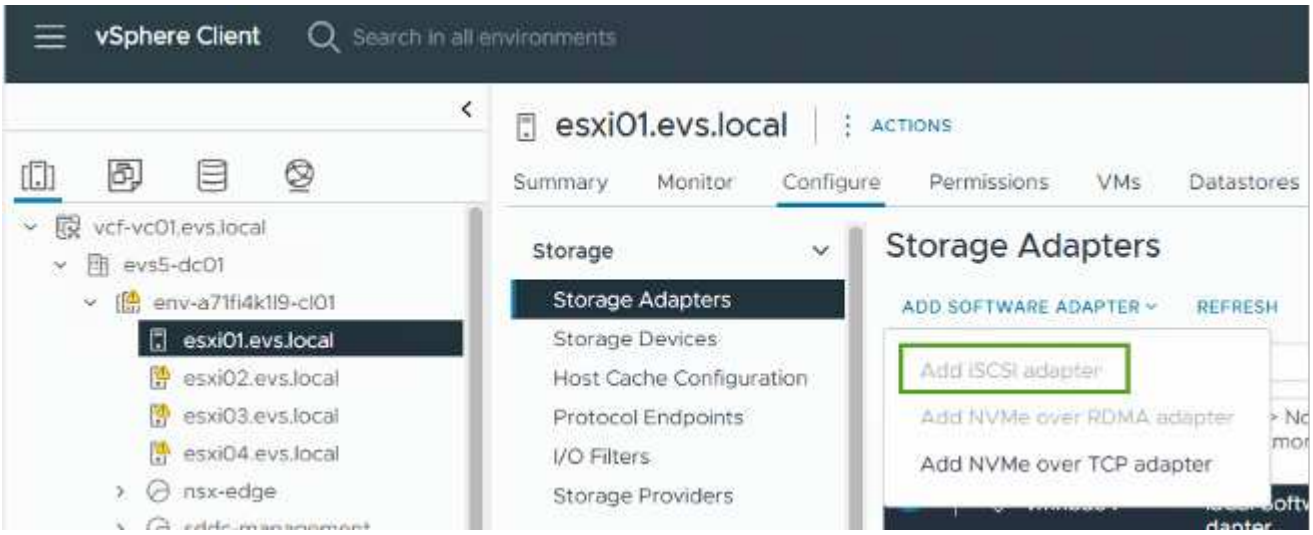
- Setting up an iSCSI network adapter for each ESXi host.
- Creating and configuring the FSx for ONTAP file system, volume, LUN, etc.

- Connecting the ESXi hosts to the FSx for ONTAP file system server.
- Creating an iSCSI datastore virtual machine file system (VMFS).

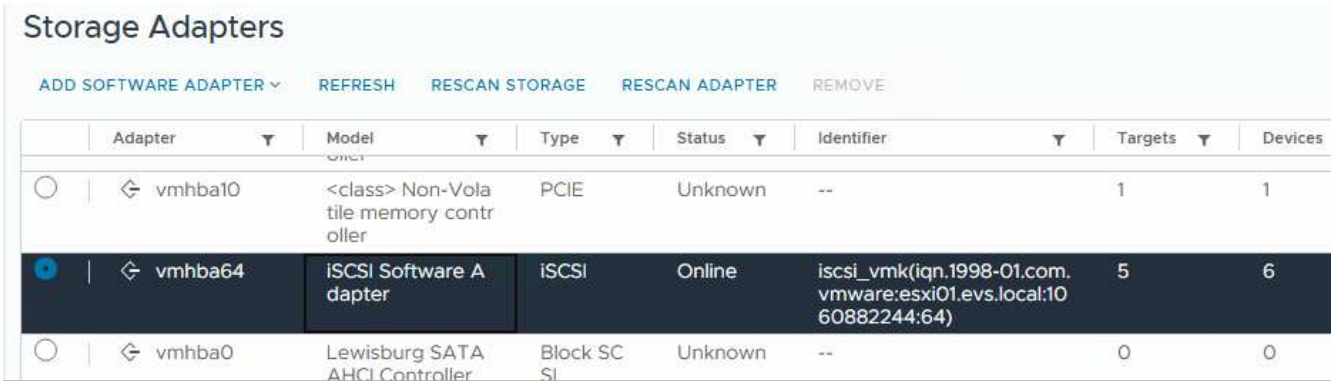
Add iSCSI network adapter

The following steps need to be performed on each ESXi host in your environment.

1. Enter the vSphere client and select the ESXi host. Under the Configure tab, choose "Add Software Adapter" then select "Add iSCSI adapter" from the drop-down menu, as shown here:

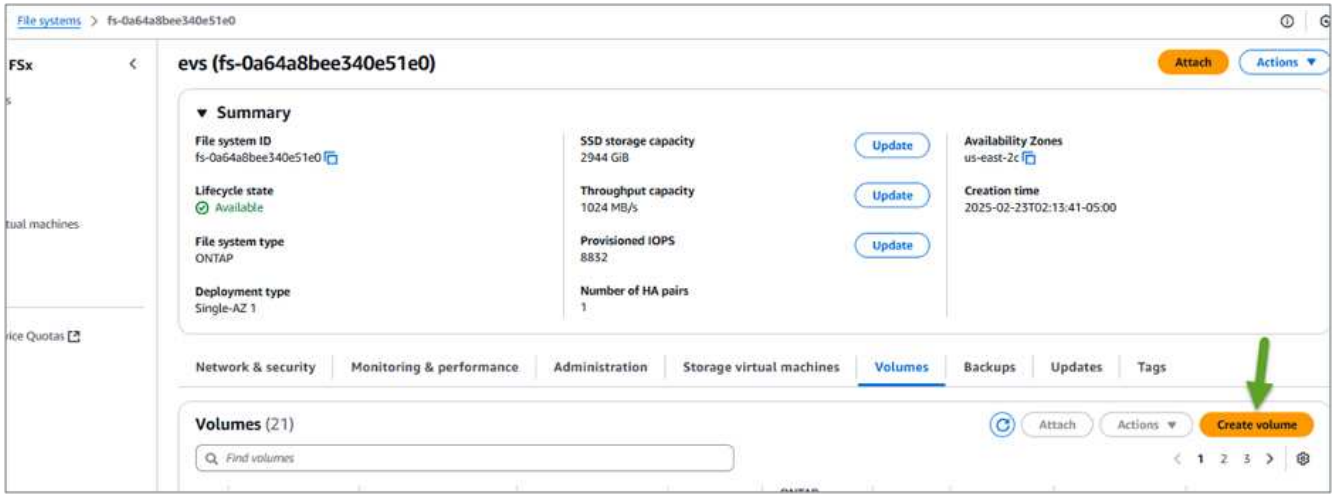


2. Once the installation is complete, you'll be on the Storage Adapters screen. Select "Rescan Adapter" from the menu bar. Record the IQN of the iSCSI adapter: iqn.xxxx. This value will be used in the LUN mapping step later.



Create an FSx for ONTAP volume

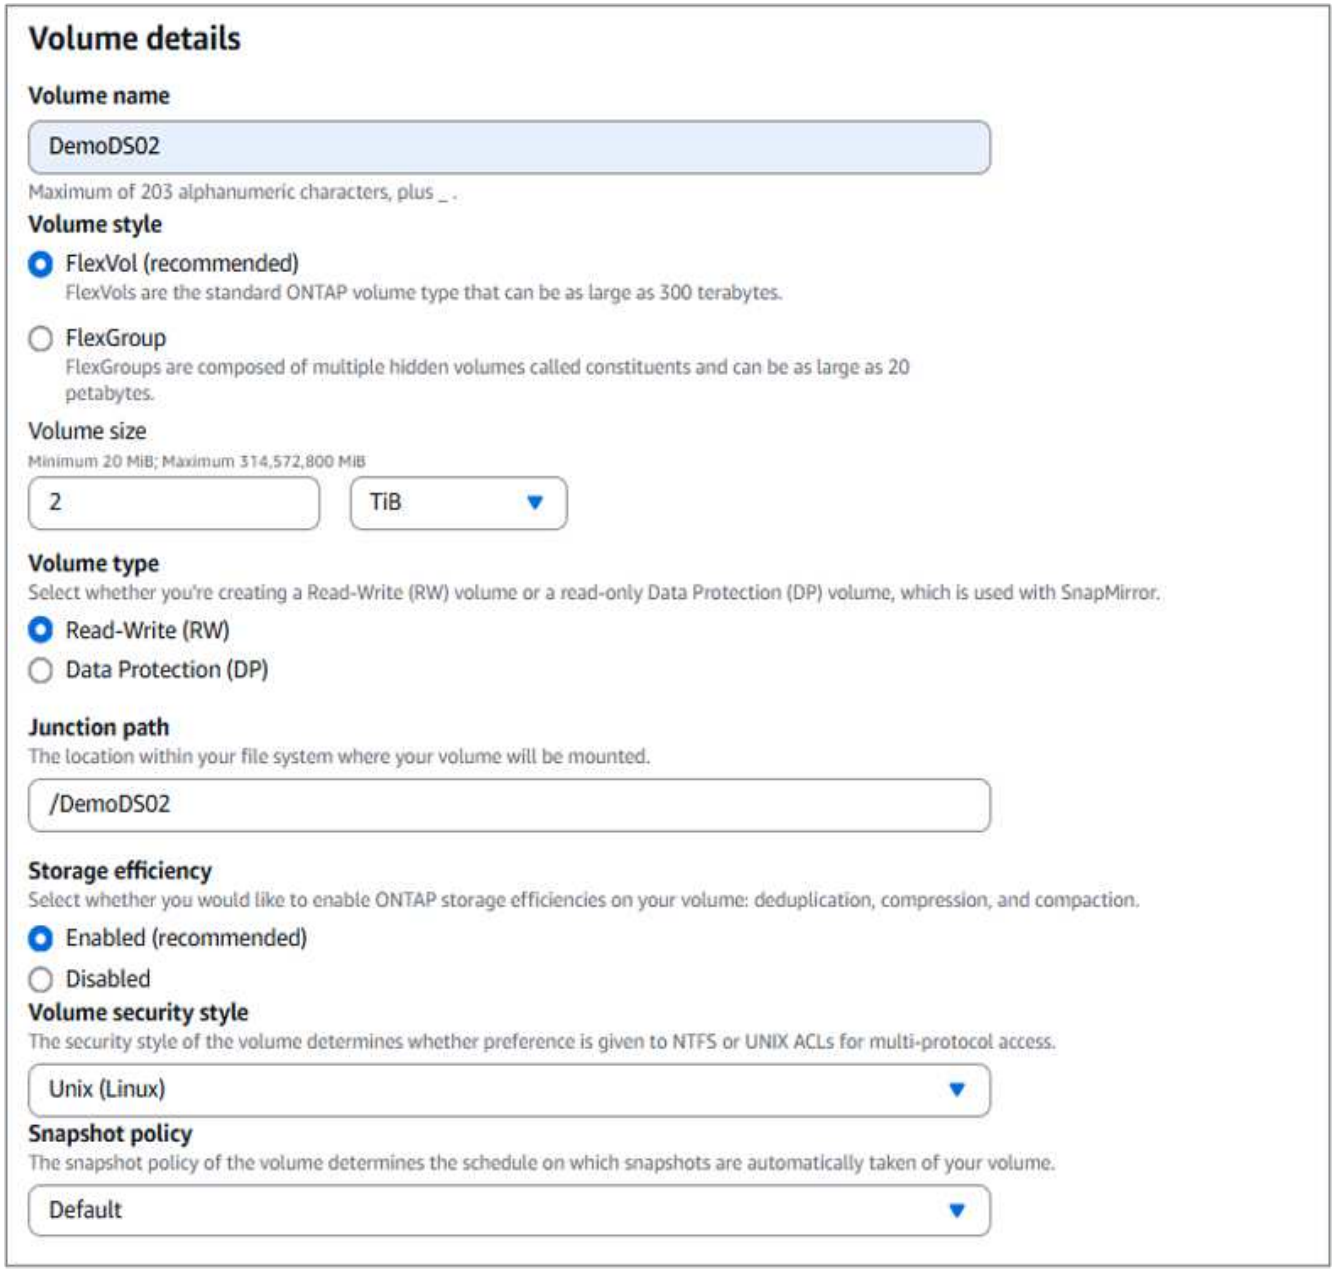
1. Go to the AWS Console and enter the page for your FSx for ONTAP file system and click "Create volume".



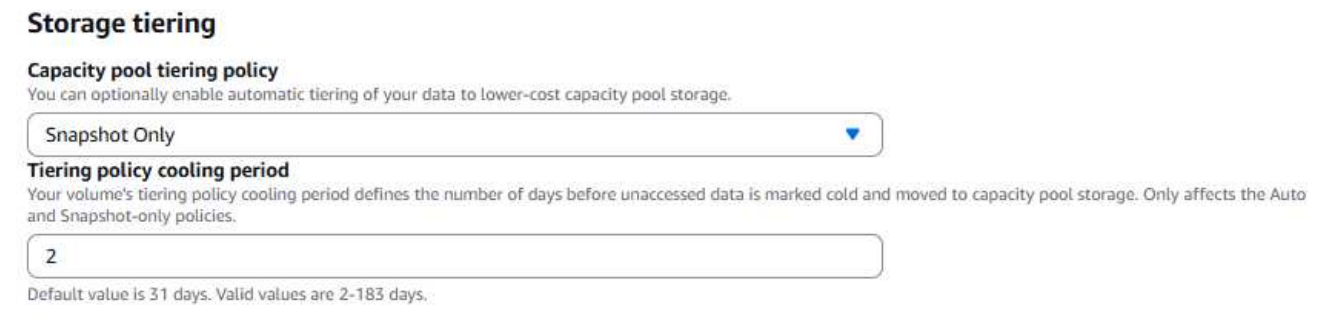
2. Select the SVM where you want to create the volume.



3. On the "Volume details" page, enter the details as shown in the screenshot below. When you are done, click "Create" to set up the volume. Choose "Enabled" to turn on the ONTAP storage efficiency features (compression, deduplication, and compaction).



4. For the storage tiering option, select "Snapshot Only" from the drop-down menu and set a tiering policy cooling period. The default is 31 days, but for this example we set a two-day cooling period.



5. Once the volume has been created, use the ONTAP CLI to set the following values:

Setting	Configuration
Volume guarantee	None (thin provisioned)
fractional_reserve	0%
snap_reserve	0%
Autosize	grow_shrink
Storage efficiency	Enabled
Volume Tiering Policy	Snapshot only
Access time update (-atime-update)	False
Minimum readahead (-min-readahead)	False
Snapshot policy	None
LUN space reservation	Disabled
LUN space allocation	Enabled

Create a LUN, igroup, and mapping the LUN

In this section we show how to create the LUN which will appear as a disk in vSphere. We will also show how to create an initiator group (igroup), which is the list of the initiators that have access to the LUN. The igroup will allow access to the LUN from the vSphere ESXi hosts. Finally, we will show how to map the igroup to the LUN.

1. To access the ONTAP CLI, establish an SSH session on the management port of the FSx for ONTAP file system or SVM by running the following command:

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

2. Create the LUN. Be sure to set space-allocation and space-reservation as enabled. The size should be up to 90% of the volume size:

```
Fsxid::> lun create -vserver fsx -path /vol/DemoDS02/lun_name -size 1.8T
-ostype vmware -space-allocation enabled -space-reservation disabled
```

3. Create the igroup, and add the IQN of all ESXi hosts within the cluster.

```
Fsxid::> lun igroup create -vserver fsx -igroup evs1 -protocol iscsi
-ostype vmware -initiator iqn.1998-
01.com.vmware:esxi01.evs.local:1060882244:64,iqn.1998-
01.com.vmware:esxi02.evs.local:1911302492:64,iqn.1998-
01.com.vmware:esxi03.evs.local:2069609753:64,iqn.1998-
01.com.vmware:esxi04.evs.local:1165297648:64
```

4. Map the LUN with the following code:

```
Fsxid::> lun mapping create -vserver fsx -path /vol/DemoDS02/lun1
-igroup evs1 -lun-id 2
```

5. Review the LUN mapping and verify that the state was changed to "mapped." This is a necessary double check to make sure the LUN is in a mapped state:

```
Fsxid::> lun show
Vserver   Path                            State   Mapped   Type
Size
--------- ------------------------------- ------- -------- --------
--------
fsx       /vol/DemoDS02/lun1              online  mapped   vmware
1.80TB
```

Obtain iSCSI configuration information from the SVM

1. Run iSCSI show and record the Target Name. This will be used in the vSphere discovery of the FSx for ONTAP disks (Note: This target name is unique to your SVM):

```
Fsxid::> iscsi show -vserver fsx
                 Vserver: fsx
             Target Name: iqn.1992-
08.com.netapp:sn.c083e374f1b711ef9d9f8123a3ebcf46:vs.3
            Target Alias: fsx
   Administrative Status: up
```

2. Record the iSCSI IP addresses:

```
Fsxid::> network interface show -vserver fsx -data-protocol iscsi
-fields address
vserver lif       address
------- -------   -----------
fsx     iscsi_1   10.0.10.134
fsx     iscsi_2   10.0.10.227
```

Discover the FSx for ONTAP iSCSI server

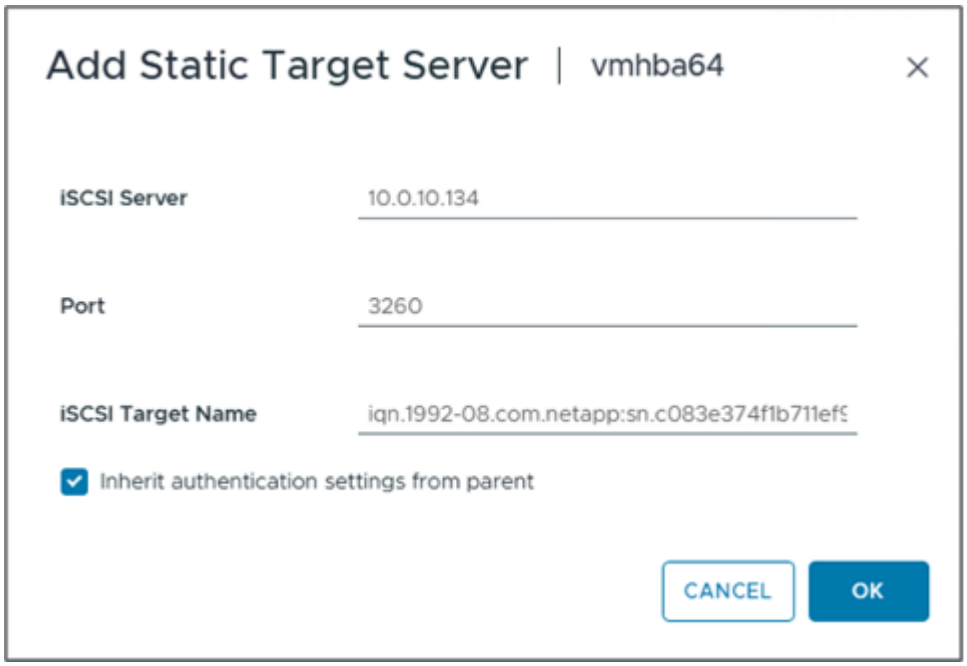
Now that we have the LUN mapped, we can discover the FSx for ONTAP iSCSI server for the SVM. Note that for each ESXi host that exists in your SDDC you will need to repeat the steps listed here.

1. First, make sure the security group linked to the FSx for ONTAP file system (i.e., the one connected to the ENI) allows iSCSI ports.

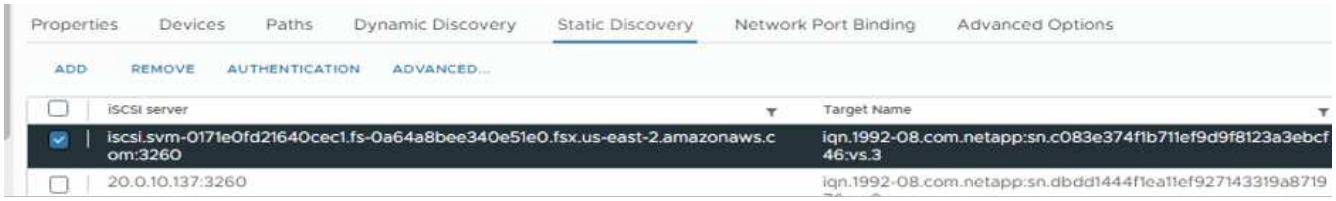
   For a full list of iSCSI protocol ports and how to apply them refer to File System Access Control with

Amazon VPC.

2. In the vSphere Client, go to ESXi Host > Storage Adapters > Static Discovery and click "Add."
3. Enter the iSCSI Server IP address from above (Port is 3260). The iSCSI target name is the IQN from the iSCSI show command. Click "OK" to proceed.



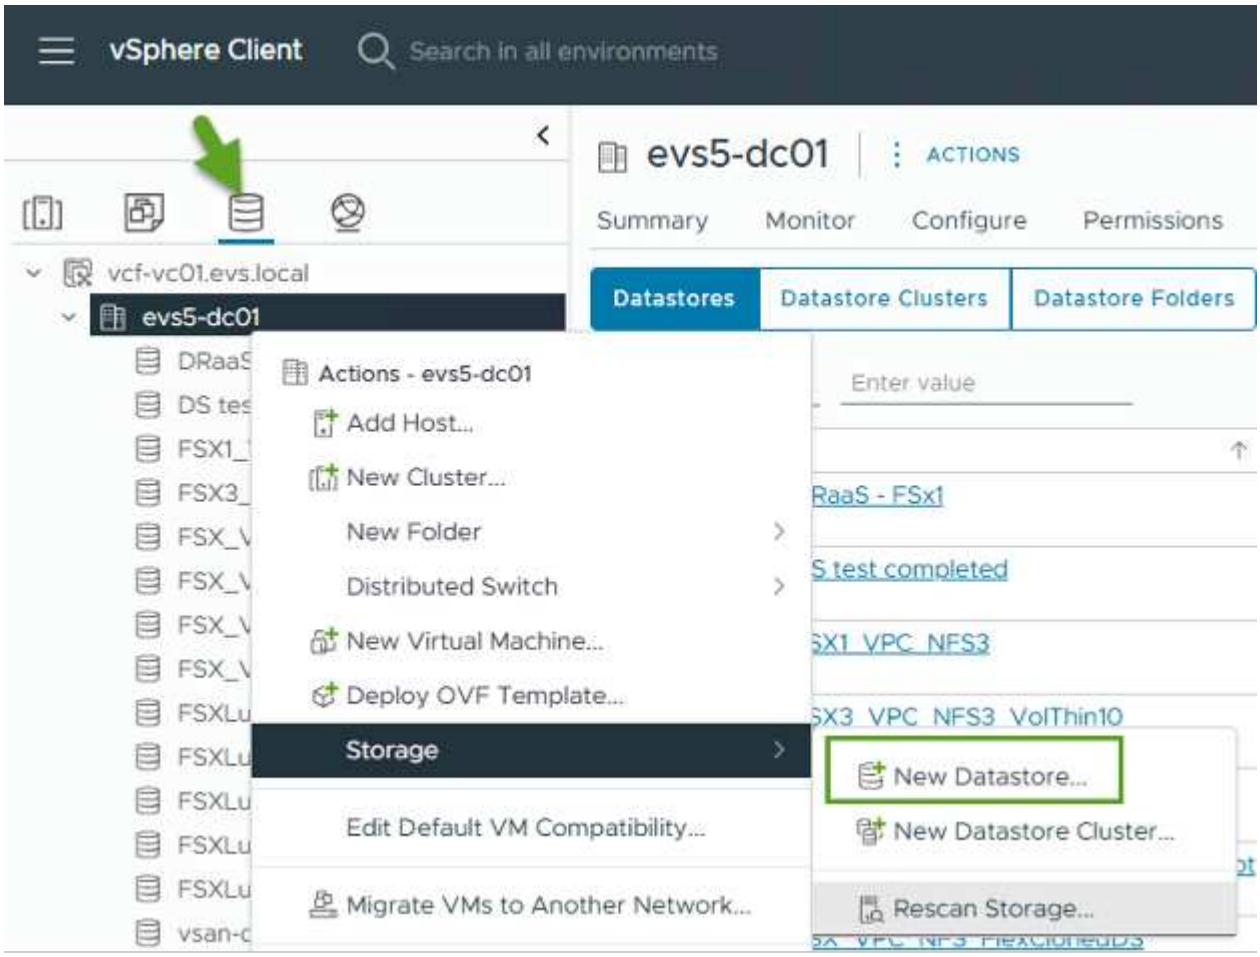
4. The wizard will close and you will be on the Datastore Static Discovery screen. In the table on this page you will be able to verify that the target has been discovered.



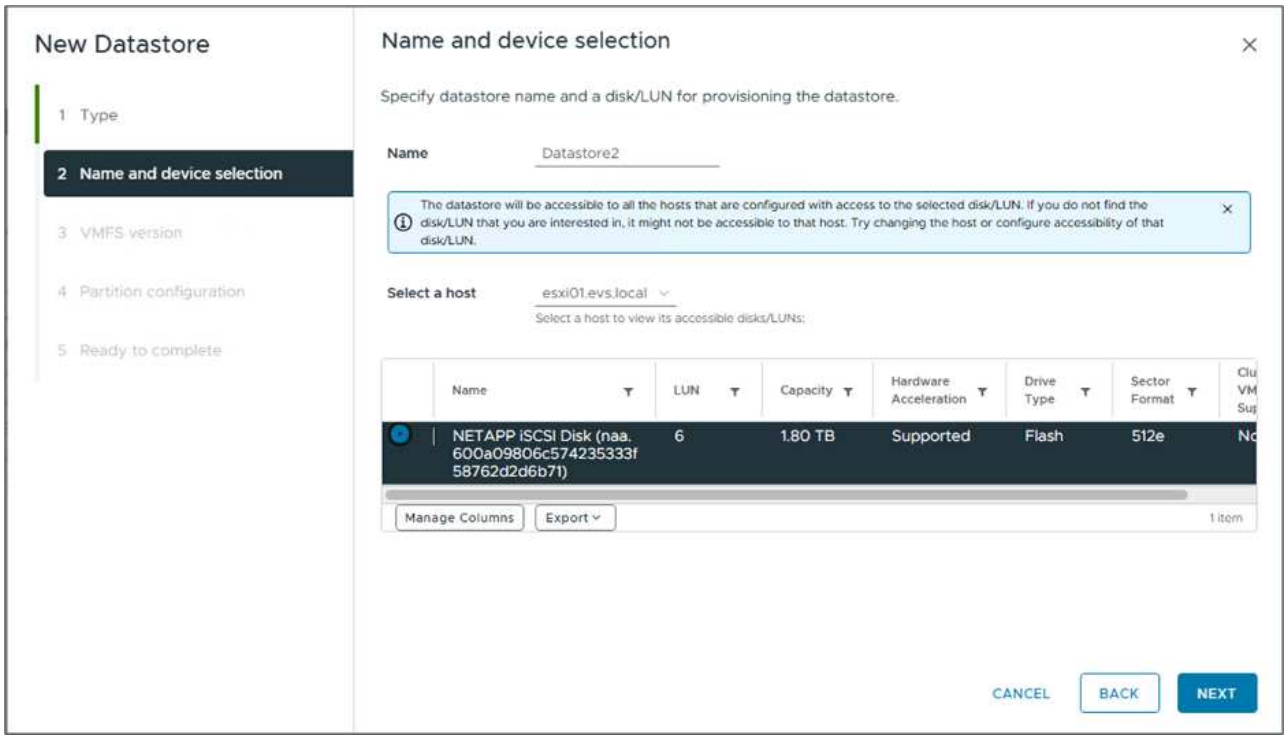
Create an iSCSI datastore

Now that we have discovered the iSCSI server, we can create an iSCSI datastore.

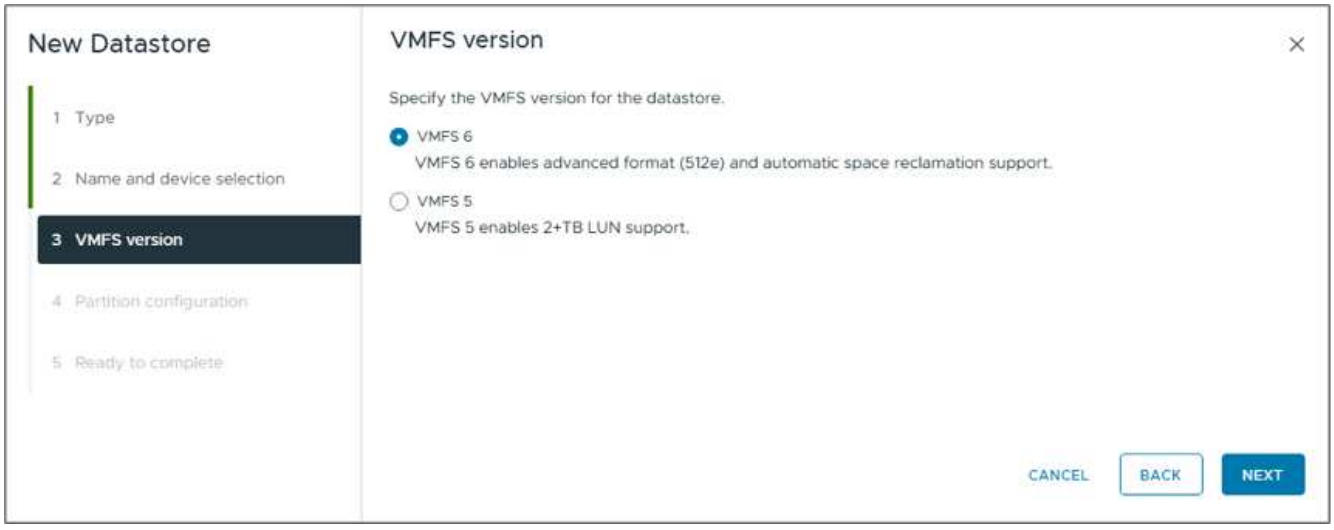
1. In the vSphere client, go to the Datastore tab, select the SDDC where you want to deploy the datastore. Right click and select the Storage icon (indicated by the green arrow in the screenshot below), then select "New Datastore" from the drop-down menu:



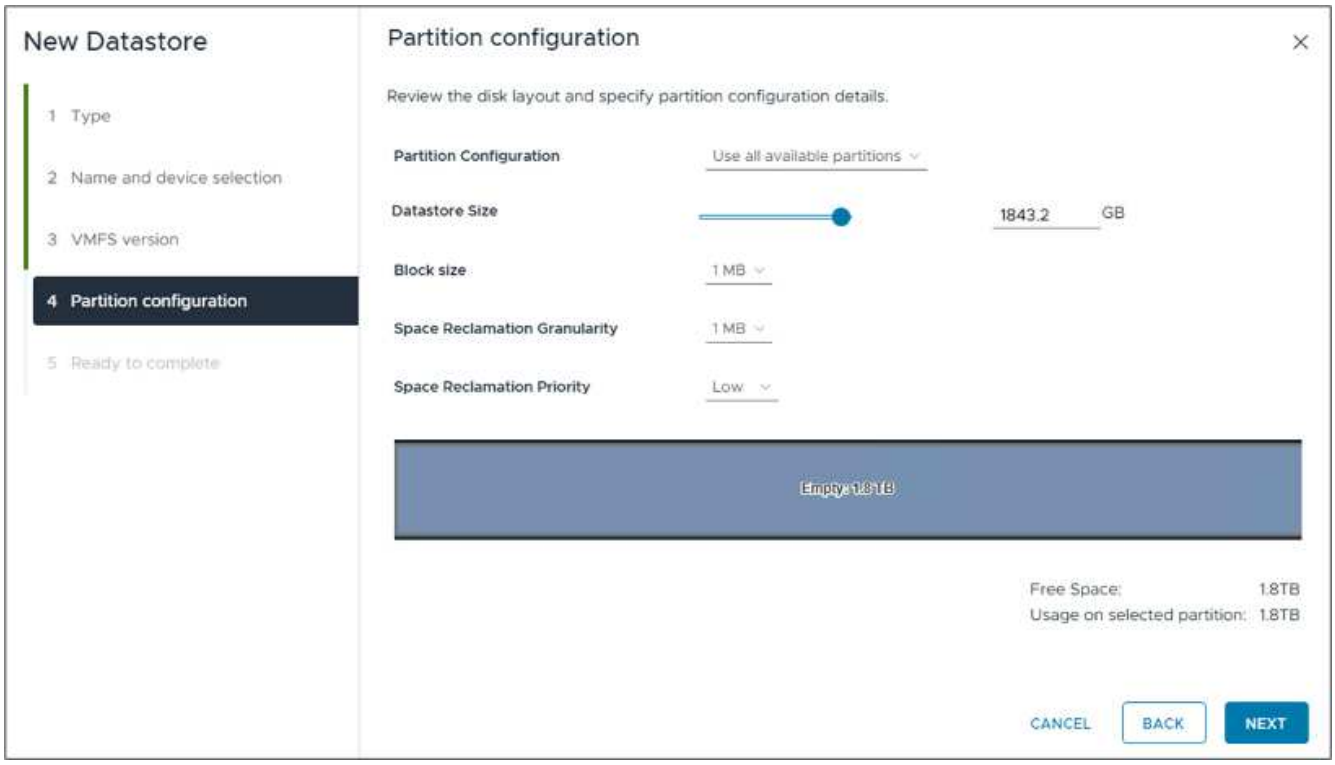
2. You will now be in the New Datastore wizard. In the "Type" step, select the VMFS option.
3. In the "Name and device selection" step:
   a. Provide a name for your datastore.
   b. Select which ESXi host you want to connect to the datastore.
   c. Select the discovered disk (LUN) and click "Next."



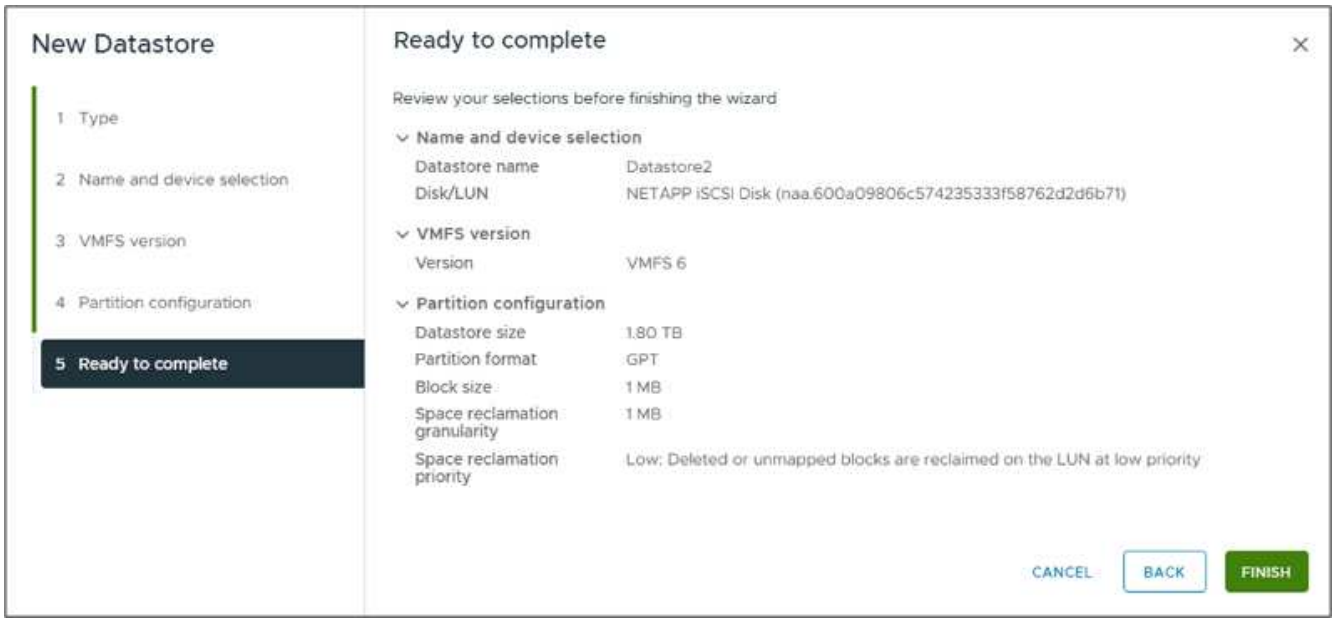
4. In the "VMFS version" step, select "VMFS 6."



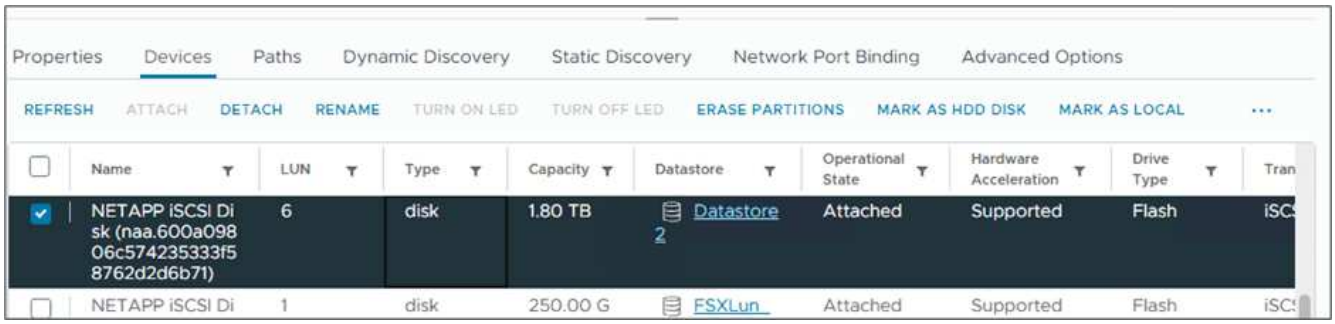
5. In the "Partition configuration" step, leave the default settings as is, including "Use all available partitions" option. Click "Next" to proceed.



6. In the "Ready to complete" step, make sure the settings are correct. When you are done, click "FINISH" to complete the set up.



7. Return to the Devices page and verify that the datastore has been attached.



Conclusion

This document has provided directions on how to set up FSx for ONTAP as an external storage for Amazon EVS for use with both the NFS and iSCSI protocols.

Using FSx for ONTAP as external storage for Amazon EVS provides clear operational and business values. It provides better data resilience, protection, replication, and cost efficiency through ONTAP technology, and keeps TCO optimized thanks to decoupling the VMware hyper-converged architecture. With FSx for ONTAP, Amazon EVS users can scale storage without adding compute resources, and without inflating costs. That opens up the full potential for elasticity and scale presented by the AWS Cloud.

To learn more, read the Amazon what's new announcement: Amazon EVS now integrates with FSx for ONTAP, and the NetApp blog: Amazon EVS announces support for FSx for ONTAP.

Protect FSx ONTAP datastore for EVS

How to protect Amazon FSx for NetApp ONTAP datastores with Amazon Elastic VMware Service using NetApp SnapCenter Plug-in for VMware

This guide will demonstrate a simple deployment approach for protecting FSx for ONTAP Network File System (NFS) and iSCSI Amazon EVS datastores and associated VMs using the NetApp SnapCenter® Plug-in for VMware vSphere (SCV) backup and recovery solution for VMs.

Introduction

Amazon FSx for NetApp ONTAP (FSx for ONTAP)-based supplemental datastores for Amazon Elastic VMware Service (Amazon EVS) offer an excellent solution for running both performance-demanding workloads—such as leveraging nConnect, Amazon Virtual Private Cloud (Amazon VPC) peering, and using multiple network connections per datastore mount—and storage-intensive workloads.

This capability enables customers to overcome the limitations of local instance storage provided by a virtual storage area network (vSAN), scale storage independently of compute, and lower overall total cost of ownership. Beyond total cost of ownership (TCO) benefits, FSx for ONTAP delivers Amazon EVS with enterprise-grade features familiar to customers of on-premises ONTAP®, helping optimize and streamline deployments. Using FSx for ONTAP provides access to NetApp® Snapshot™, SnapMirror®, and SnapVault® technologies that support back up and restoration operations for virtual machines (VMs), datastores, and virtual machine disks (VMDKs). One key feature is Snapshot-based protection.

Overview

As organizations scale and consolidate environments, the need to protect those workloads and the underlying storage becomes increasingly important. SCV can be used to back up datastores and virtual machines directly

from the Amazon EVS vSphere client.

Purpose of this Guide

This guide provides instructions on how to protect an FSx for ONTAP datastore for Amazon EVS using SCV. Example scenarios of protecting VMs and databases are provided as well as restore.

Audience

This tutorial is intended for Cloud Architects and Cloud Administrators familiar with VMware vSphere, Amazon EVS, the AWS Console, and FSx for ONTAP.

High-level steps for deployment and configuration

Below are the high-level steps for deploying and configuring SCV to protect your VMs and datastores on Amazon EVS:

1. Download the SCV Open Virtual Appliance (OVA).
2. Log in with the vSphere Client credentials to Amazon EVS.
3. Deploy an Open Virtualization Format (OVF) template to start the VMware deployment wizard.
4. To access the plug-in, select SnapCenter Plug-in for VMware from the menu.
5. Add storage.
6. Create backup policies.
7. Create resource groups.
8. Backup resource groups.
9. Restore the entire virtual machine or particular virtual disk or guest files.

Setup procedure for SnapCenter Plug-in for VMware

This section covers the steps to setup SnapCenter Plug-in for VMware.

Prerequisites

The following prerequisites are required before you get started:

1. For registering the vCenter server while deploying the OVA, specify the fully qualified domain name (FQDN) or the IP address of the vCenter.
2. For the vCenter credentials, use an account with minimum required privileges
3. Open firewall ports for vCenter Server access.
4. Update the Security Group of the Amazon VPC where FSx for ONTAP is running to allow SCV access to FSx for ONTAP.



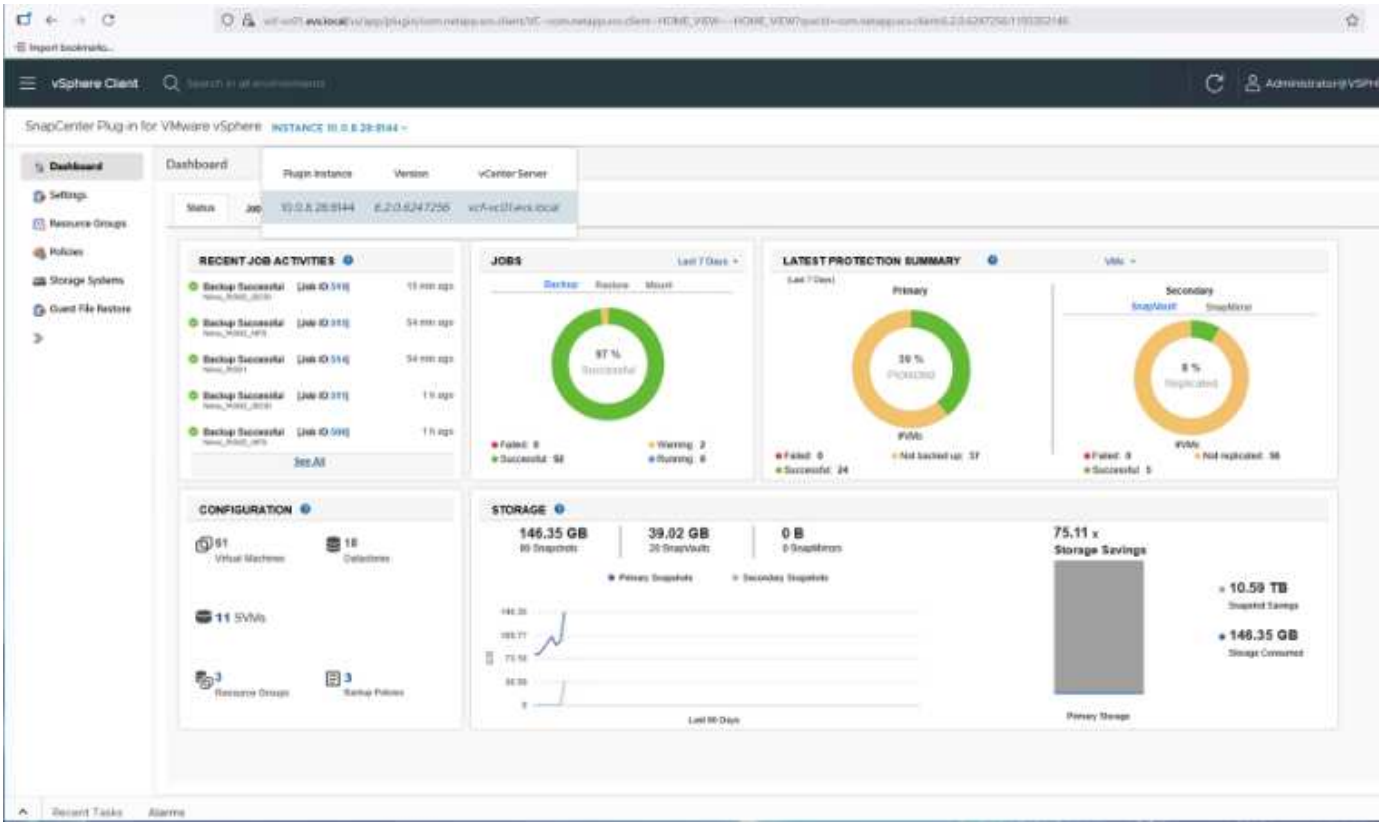
SCV is available free of charge.

Setting up the SnapCenter Plug-in for VMware

To protect VMs and NFS/iSCSI-based datastores hosting them, SCV must be deployed. It’s a simple OVF import.

The steps to deploy are as follows:

1. Download the OVA from NetApp Support Site.
2. Log in to the Amazon EVS vCenter.
3. Within vCenter, right-click any inventory object such as a data center, folder, cluster, or host and select "Deploy OVF template".
4. Select the correct settings, including storage, segment and customize the template to update the vCenter and its credentials, create SCV credentials and IP details. Once reviewed, select "Finish".
5. Wait for the OVF import and deployment tasks to complete.
6. Once SCV is successfully deployed, it will be registered within vCenter.
7. To access the plug-in, navigate to the left sidebar of the vCenter web client page, select "SnapCenter Plug-in for VMware".



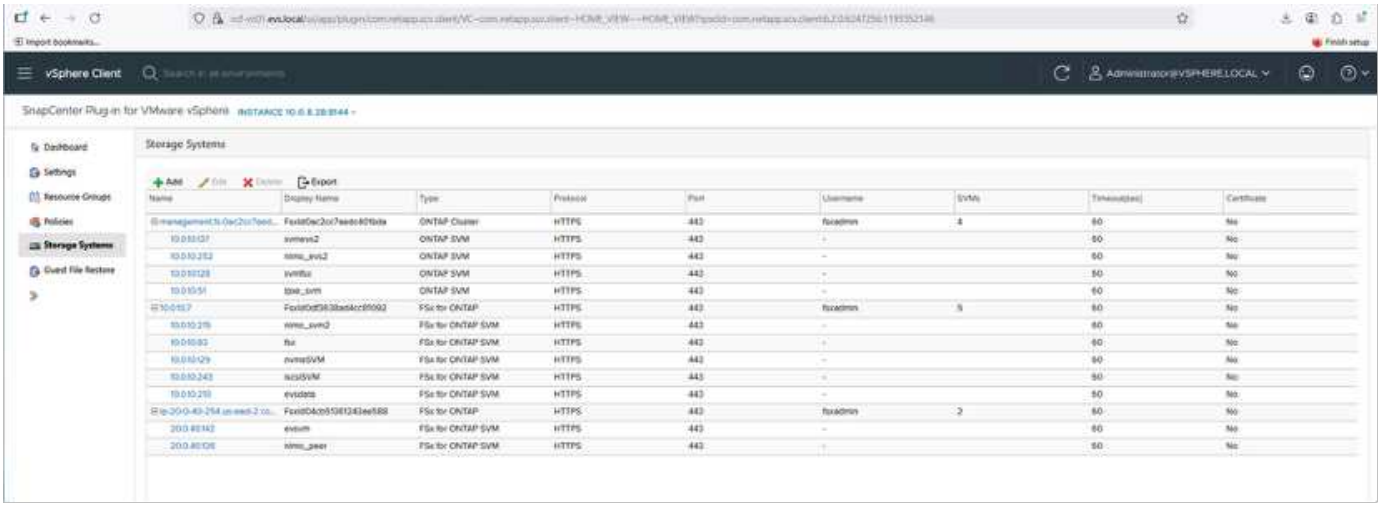
Configuration procedure for SnapCenter Plug-in for VMware

This section covers the steps to configure SnapCenter Plug-in for VMware. The following sections will detail how to add storage, create a backup policy, and how to create resource groups.

Adding the storage system

The next step in protecting workloads with SVC is to add the storage system. Adding storage allows SCV to detect and manage back up and restore operations within vCenter. Include the FSx for NetApp ONTAP management endpoint or the storage virtual machine (SVM) administration endpoint IP as a storage system to enable VM backup and restore.

1. From the left navigation menu of the EVS vCenter, select "SnapCenter Plug-in for VMware".
2. Select "Storage System".
3. Select "Add" to add the FSx for ONTAP details.
4. Use the "Credentials as the Authentication" method and enter FSxadmin as the username and password. Then select "Add" to save the settings.



All SVMs are automatically added when the file system is added.

If the SCV version is 6.2, go to the SCV support page and download the SCV 6.2P1 release.

Create a backup policy

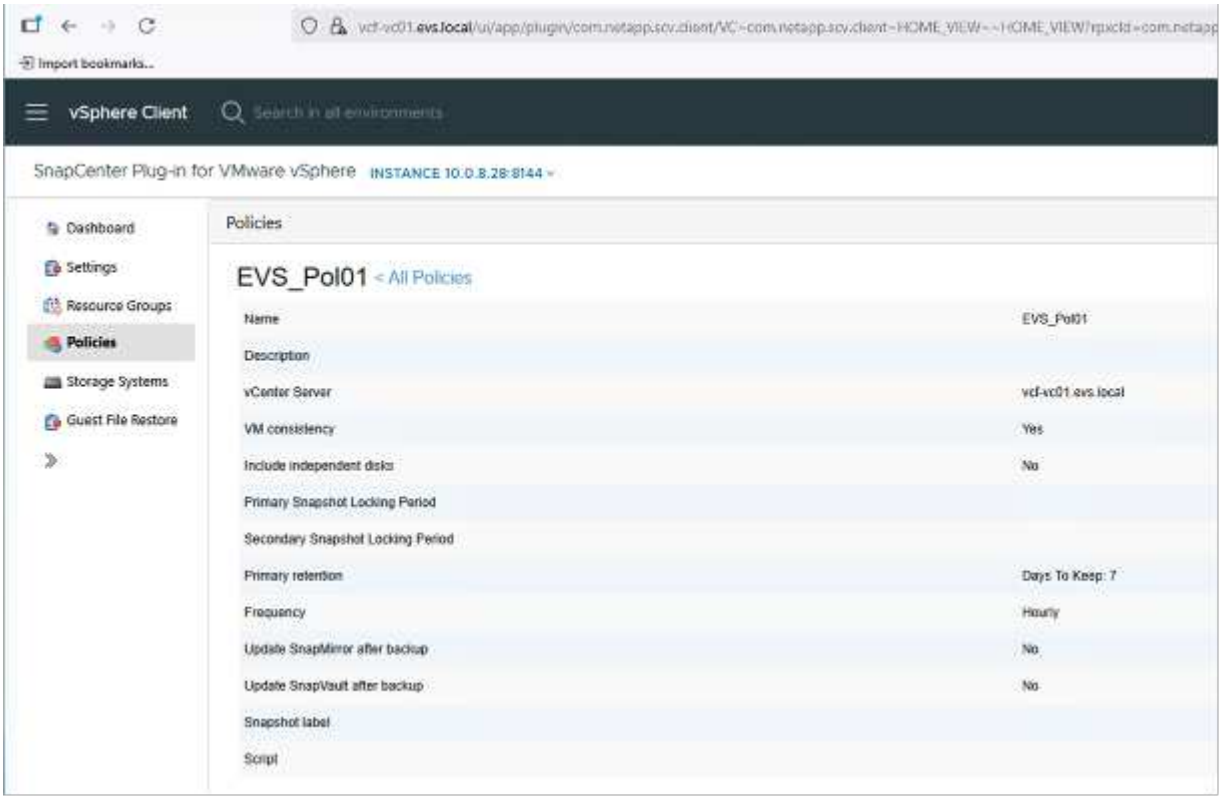
A robust backup strategy should define key factors such as the back up schedule, the data to be backed up, and the retention period for backups. Snapshot copies can be scheduled on an hourly or daily basis to protect entire datastores. Scheduling backups this way not only secures the datastores but also supports back up and restore operations for the VMs and VMDKs contained within those datastores.

Before initiating backups for VMs and datastores, create a backup policy and a resource group. The backup policy specifies parameters such as the backup schedule and retention settings.

Here are the steps to create a backup policy:

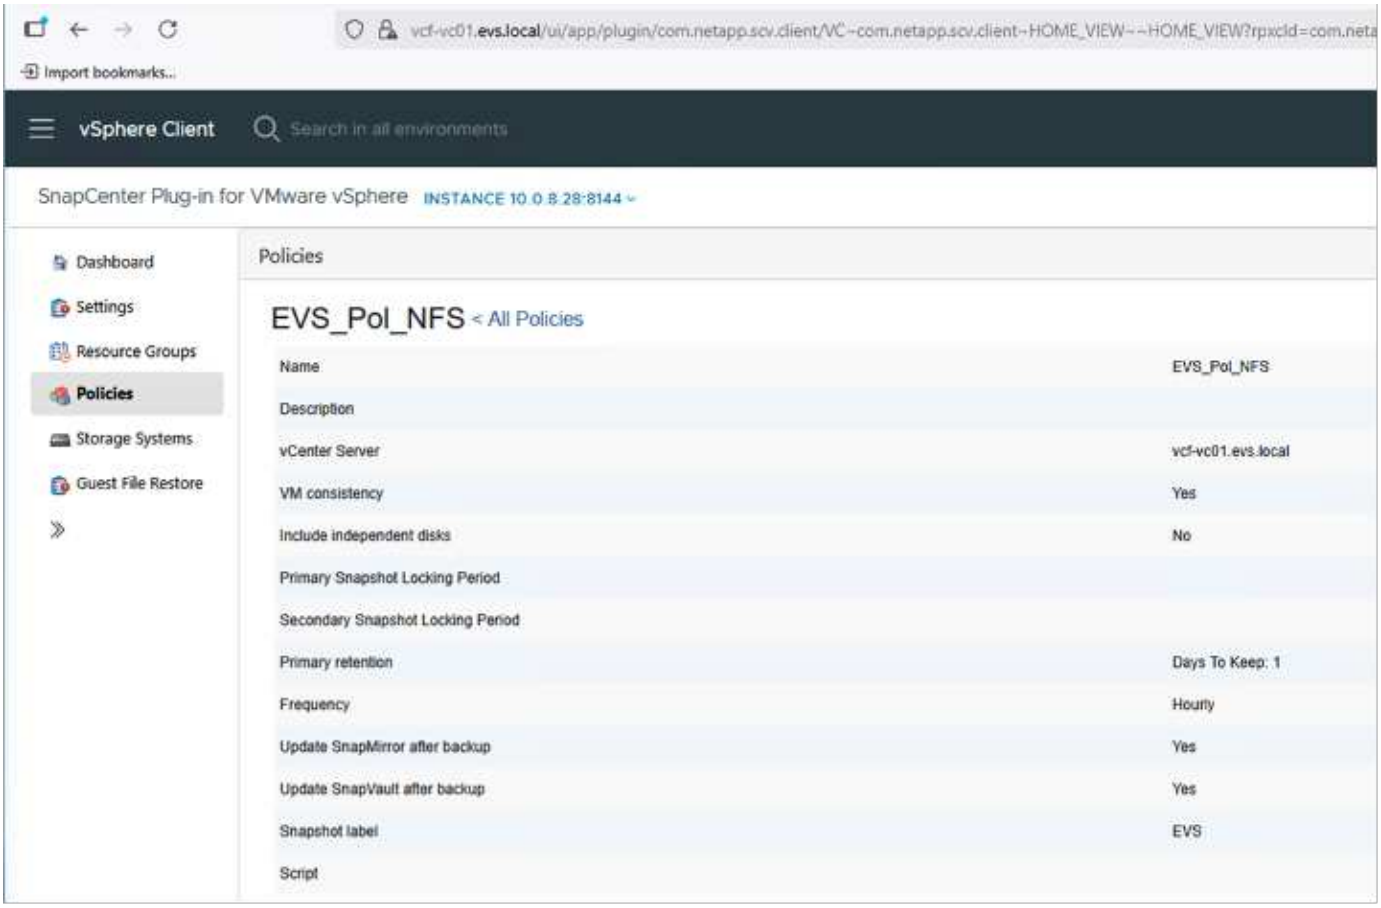
1. Go to the left-side navigation menu in SCV and select "Policies".
2. On the Policies page, select "Create" to start the wizard.
3. On the "New Backup Policy page", enter a name for the policy.
4. Specify the retention period, the back up frequency, and replication settings.
5. Select the "Locking Period" checkbox to enable Snapshot locking. You can select the primary and secondary Snapshot-locking periods as Days/Months/Years.

Regardless of the retention value defined in the SnapMirror policy, the secondary Snapshot copy remains intact until the specified secondary snapshot locking period expires.



To enable VM-consistent backups, VMware tools must be installed and running. When the VM consistency box is checked, the VMs are first quiesced, then VMware performs a VM-consistent snapshot (excluding memory), and then SCV performs its backup operation, and then VM operations are resumed.

To replicate Snapshot copies to a secondary mirror or vault file system, the relationships must be configured beforehand.



Once the policy is created, the next step is to create the resource group which will define the appropriate NFS or iSCSI datastores and VMs that should be backed up.

Create a resource group

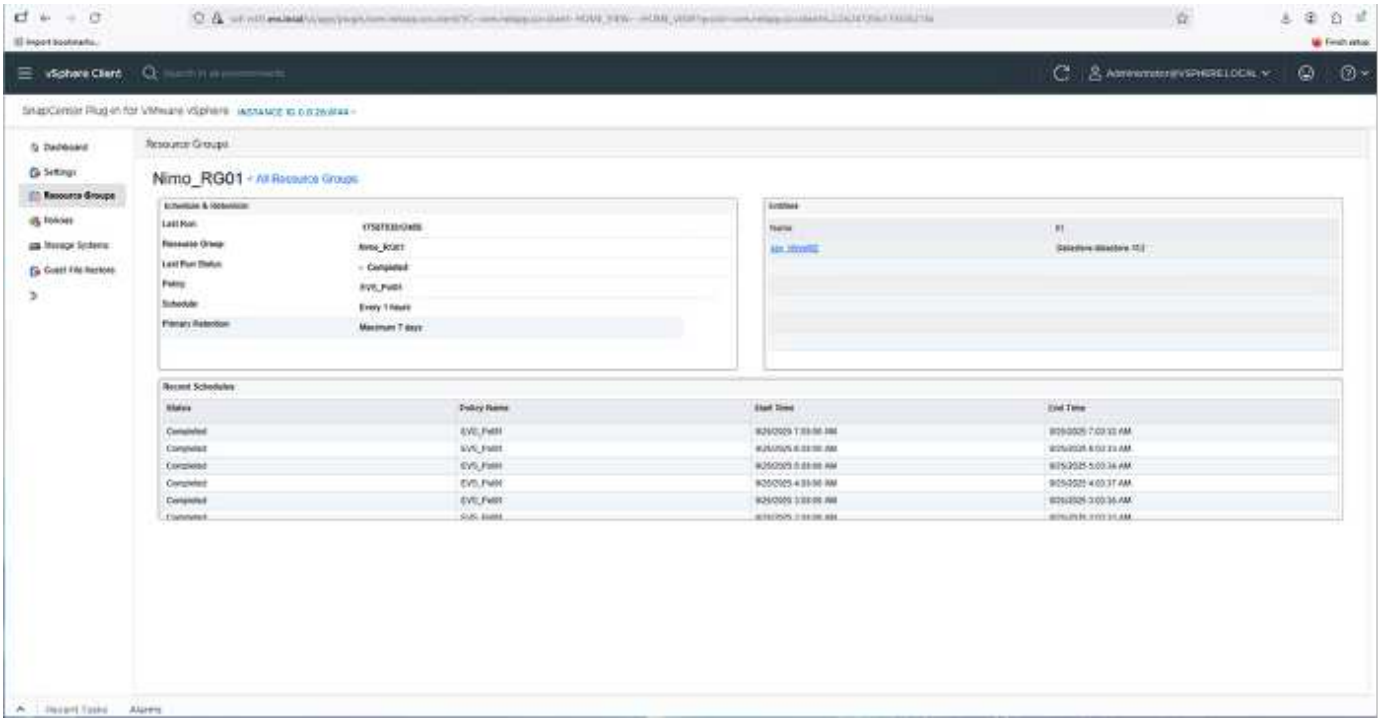
A resource group serves as a container for VMs and datastores that require protection. Resources can be added to or removed from a resource group at any time, providing flexibility in managing backup configurations.

You can create a resource group using the following steps:

1. Go to the left-side navigation pane in SCV and select "Resource Groups".
2. On the Resource Groups page, select "Create" to start the resource group creation wizard.

> ⓘ Another option to create resource groups is by selecting an individual VM or datastore and creating a resource group that contains only that resource.

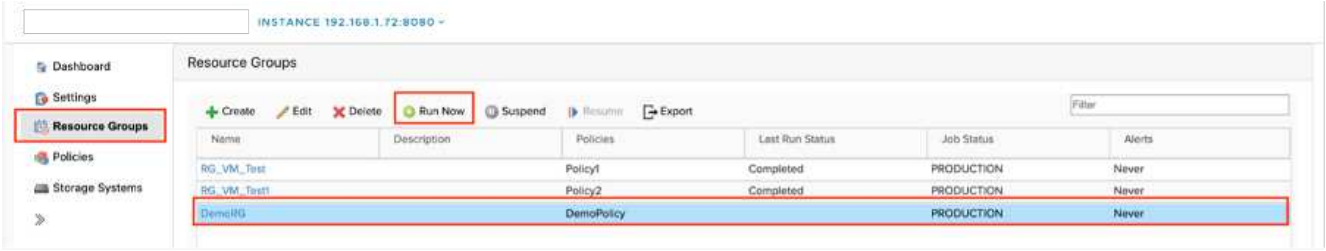
3. On the "Resources" page, select the scope (VMs or datastores) and the data center.
4. On the "Spanning disks" page, select the option for "Virtual Machines with multiple VMDKs across multiple datastores".
5. Next, associate a backup policy. Select an existing policy or create a new back up policy.
6. On the "Schedules" page, configure the backup schedule for each selected policy.
7. Once you make the appropriate selections, select "Finish". This will create a new resource group and add to the resource group list.



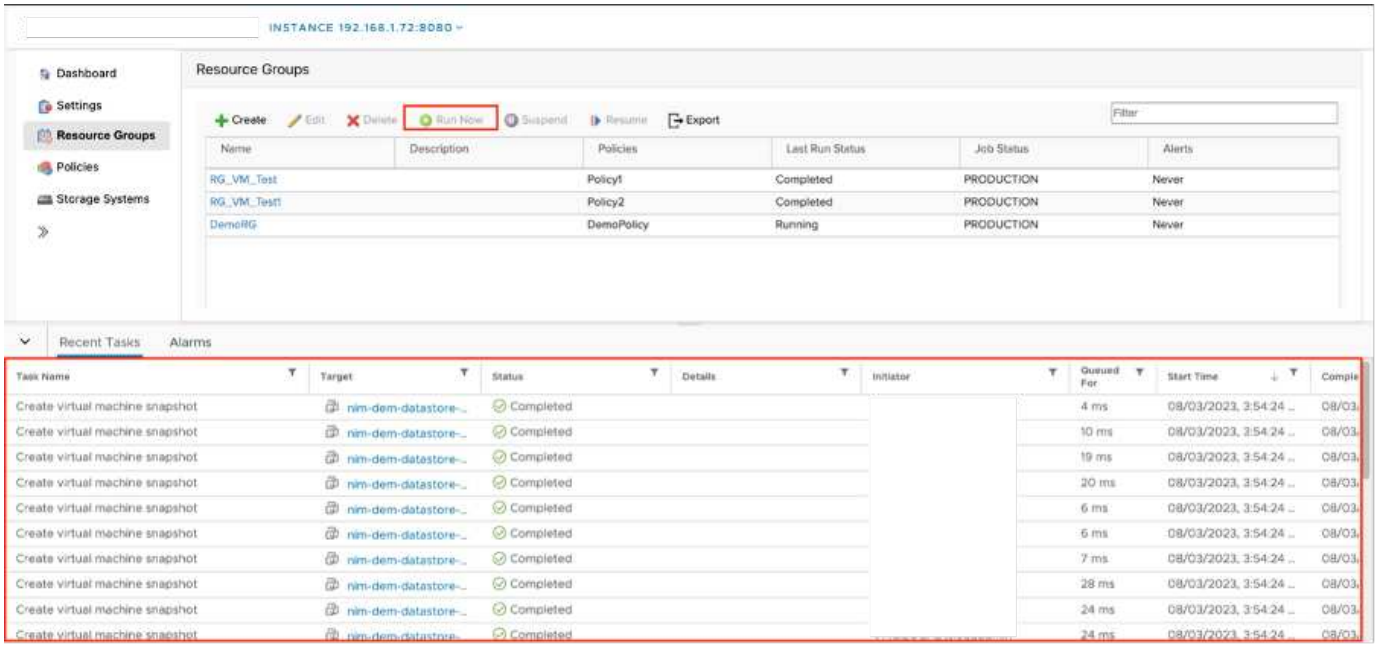
Backing up resource groups

Back up operations are executed on all resources defined within a resource group. If a resource group has a policy attached and a schedule configured, back ups occur automatically according to the schedule.

1. In the left navigation of the vCenter web client page, select "SnapCenter Plug-in for VMware" then "Resource Groups".
2. In the Resource Groups page, find the designated resource group and select "Run Now" to start the manual back up.



3. If the resource group has multiple policies configured, select the policy for the back up operation in the "Backup Now" dialog box.
4. Select "OK" to initiate the backup.

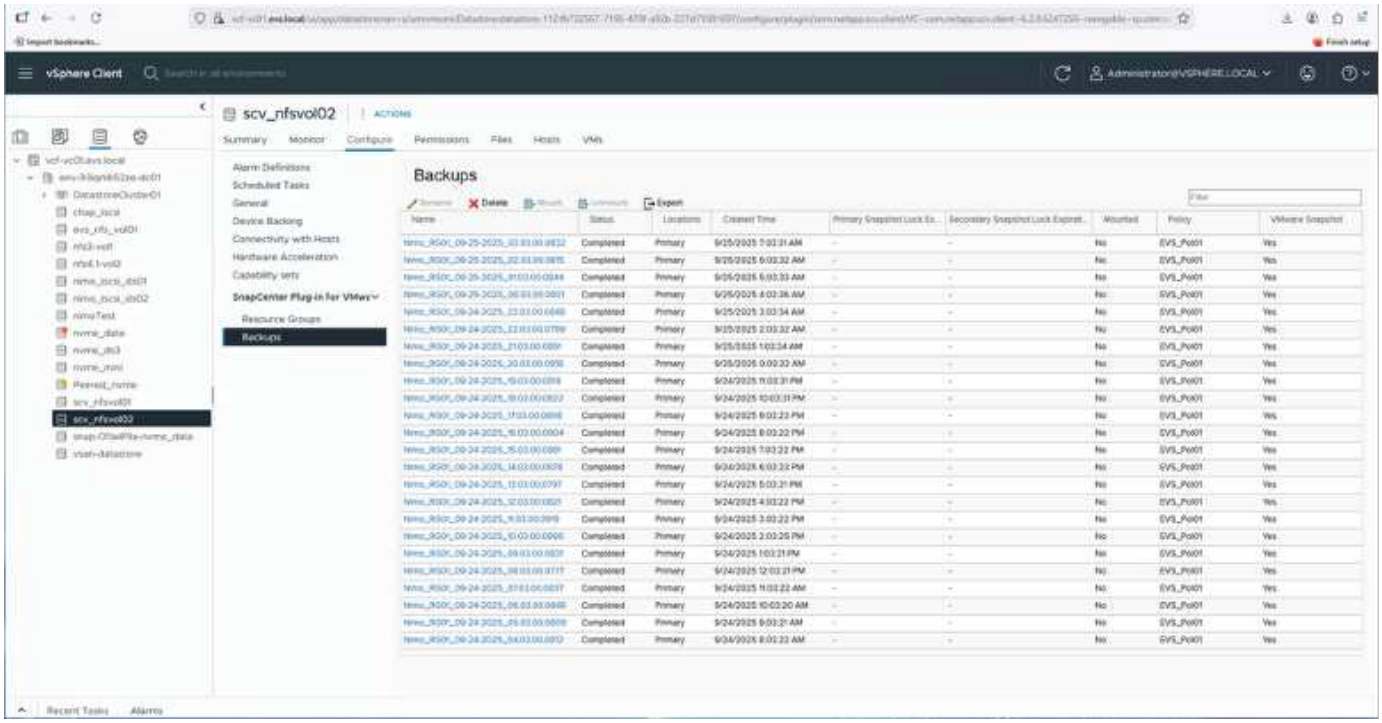


Monitor the operation progress by selecting "Recent Tasks" at the bottom of the window or, for more details, by selecting "Job Monitor" on the dashboard.

The process for creating back up policies, configuring resource groups, and performing back up operations is consistent across different storage protocols, including NFS and iSCSI.

Backups are performed according to the schedule defined in the policy.

After the back up completes, it can be viewed at the datastore level or for individual VMs by selecting "SnapCenter Plug-in for VMware vSphere" from the drop-down list and then choosing "Restore".



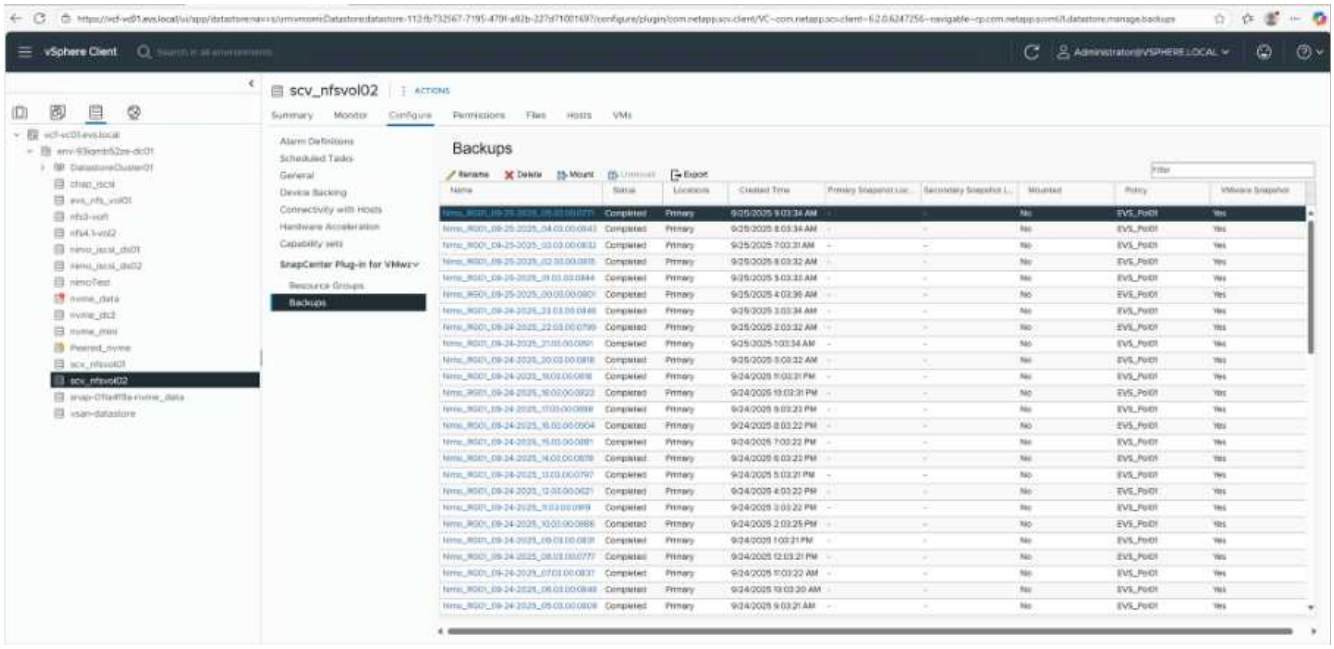
Restore procedure for VMs protected by SnapCenter Plug-in for VMware

This section covers the steps to restore VMs from backups. SCV can be used to restore VMs within vCenter, allowing them to be recovered to the original datastore on the original ESXi host (overwriting existing content), to restore deleted or renamed VMs from a backup copy, to a different datastore on the original ESXi host, to the original datastore on a different ESXi host managed by the same vCenter, or to a different datastore on another ESXi host within the same vCenter environment.

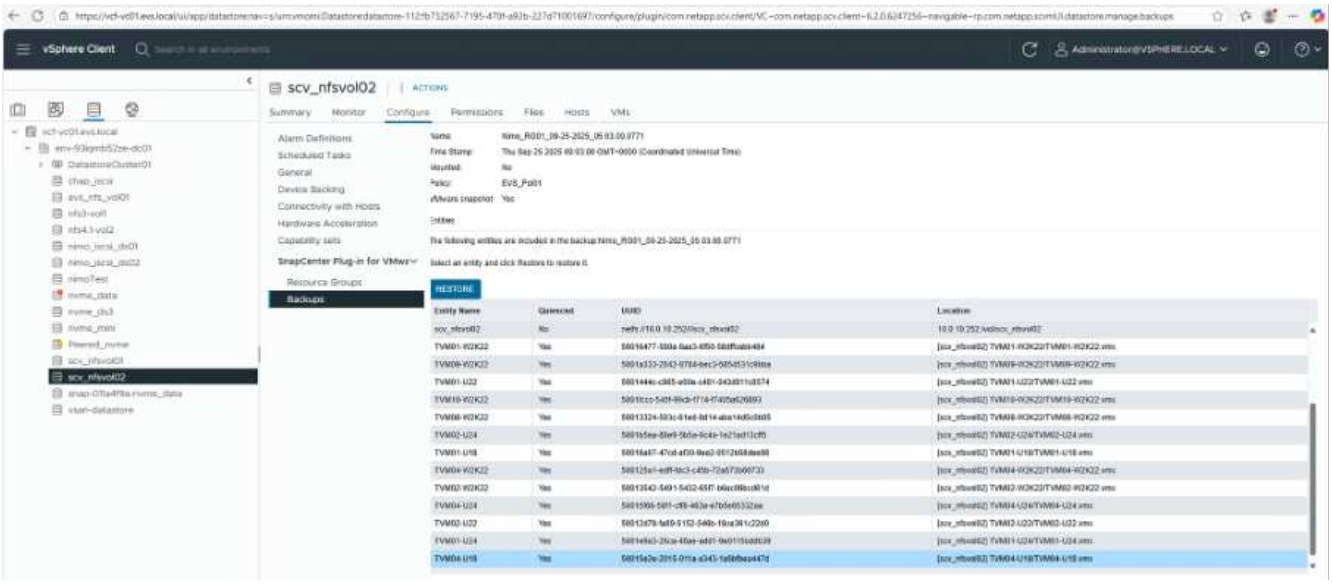
Steps to perform restore operation

To perform the restore operation, follow these steps:

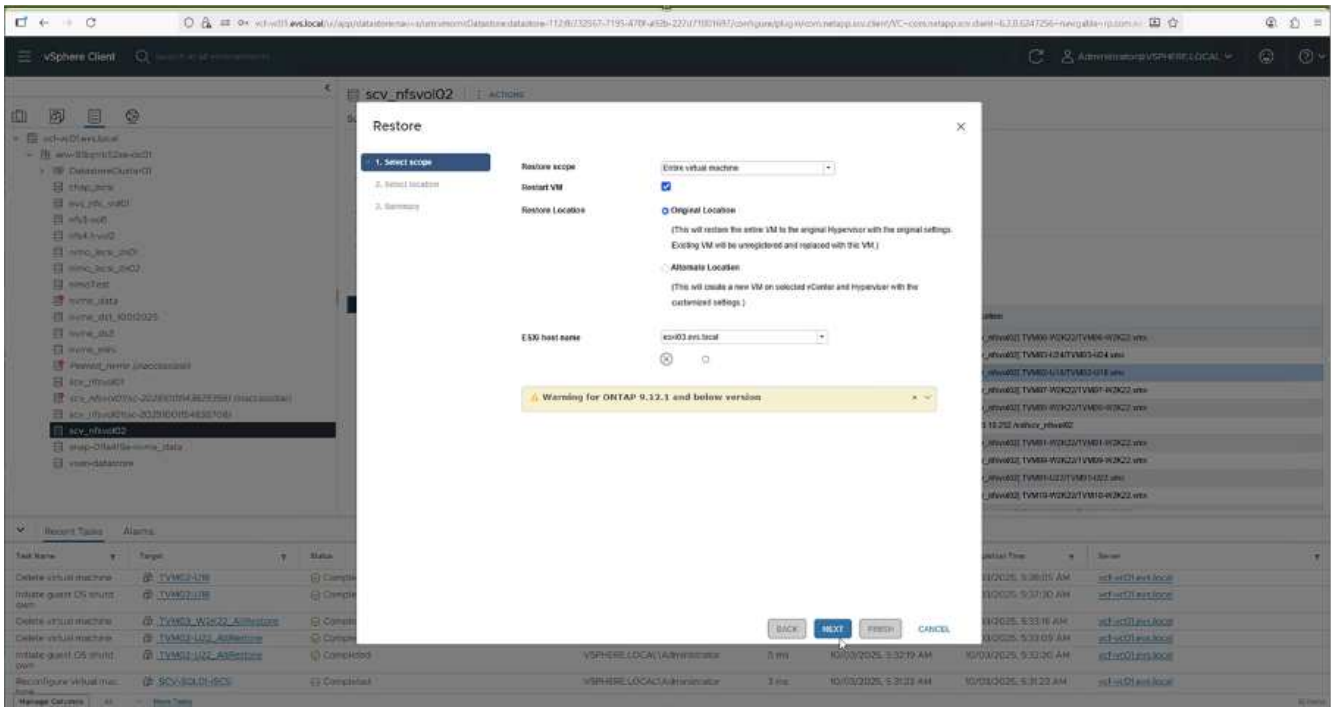
1. In the EVS vSphere web client GUI, go to **Inventory › Datastores** and select the desired datastore. Then navigate to **Configure › Backups** under the "SnapCenter Plug-in for VMware vSphere" section.



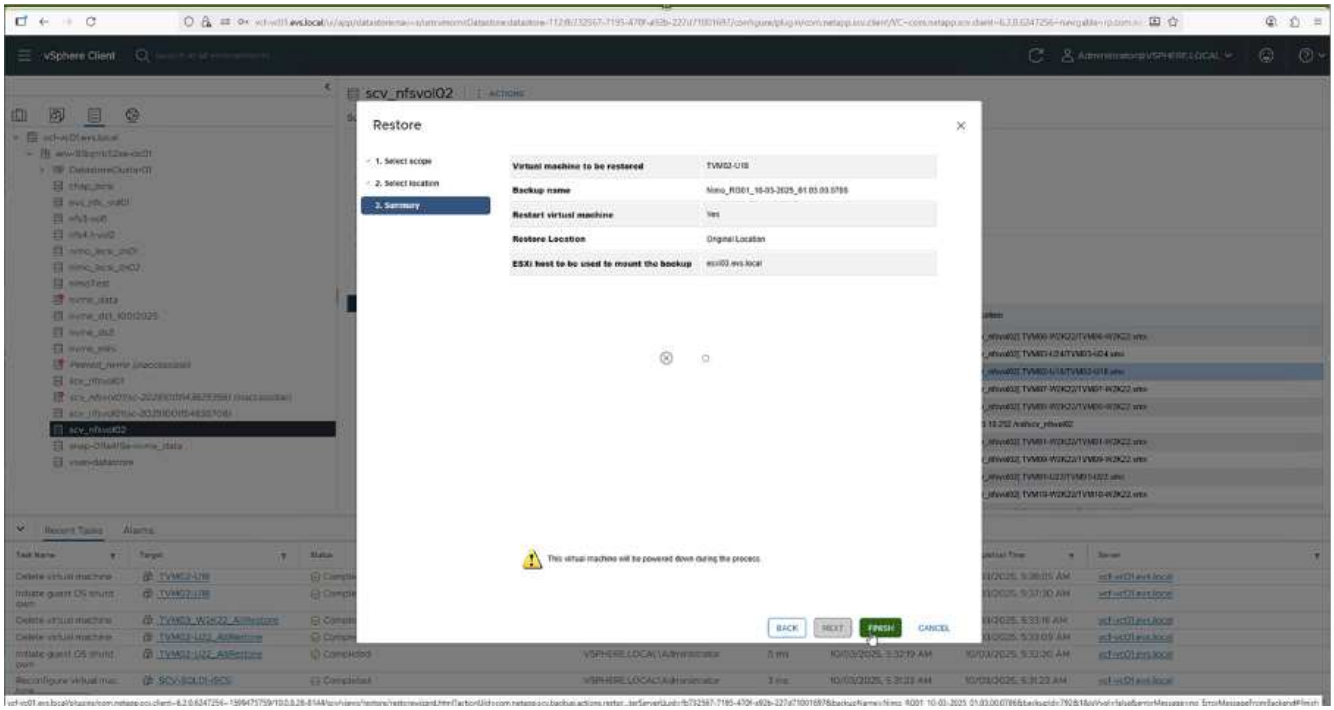
2. Double-click a backup to view all VMs included in that backup. Find the VM to be restored in the list and select "Restore".



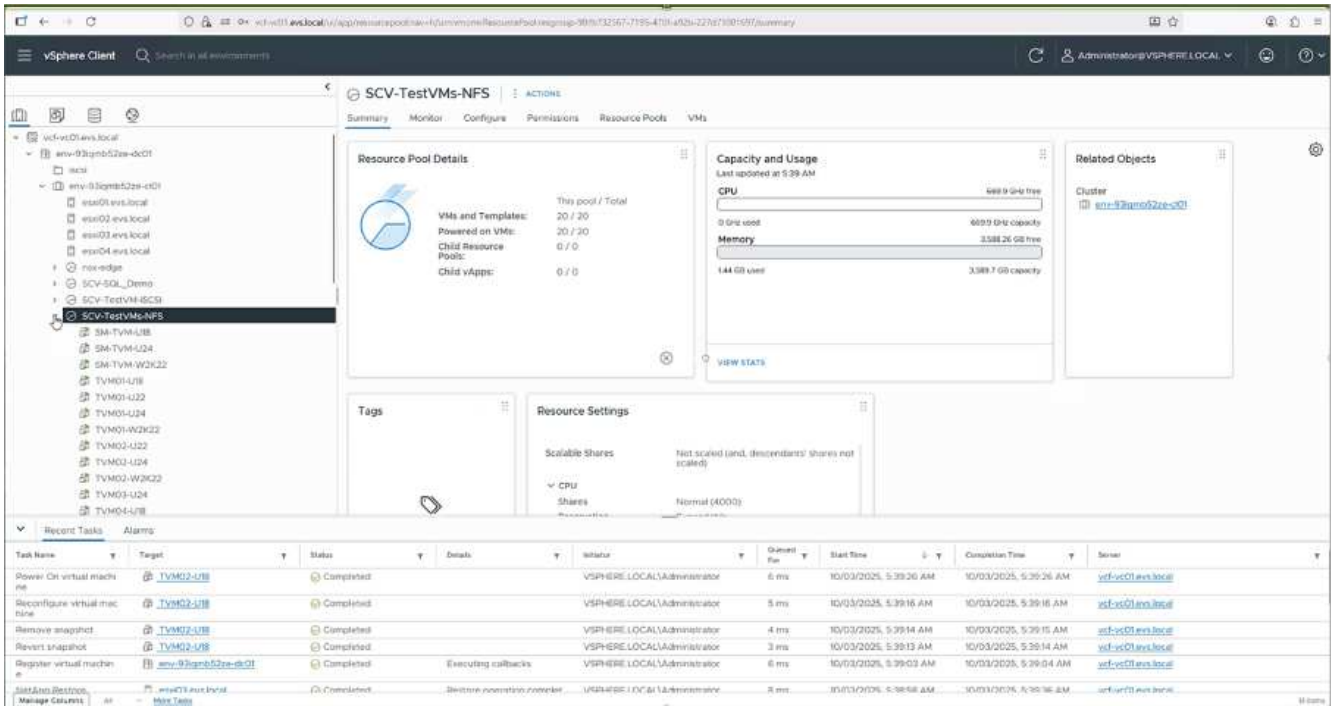
3. On the "Select Scope page", select "Entire Virtual Machine" in the "Restore scope" field, then select "Restore location" and enter the destination ESXi information where the backup should be mounted.
4. Enable the "Restart VM" checkbox if the VM needs to be powered on after the restore operation.
5. On the "Select Location page", select a primary location.



6. Review the "Summary" page and then select "Finish".



7. Monitor the operation progress by selecting "Recent Tasks" at the bottom of the screen.



Although the VMs are restored, they're not automatically added to their former resource groups. If you want to protect those VMs, add them manually to the appropriate resource groups.

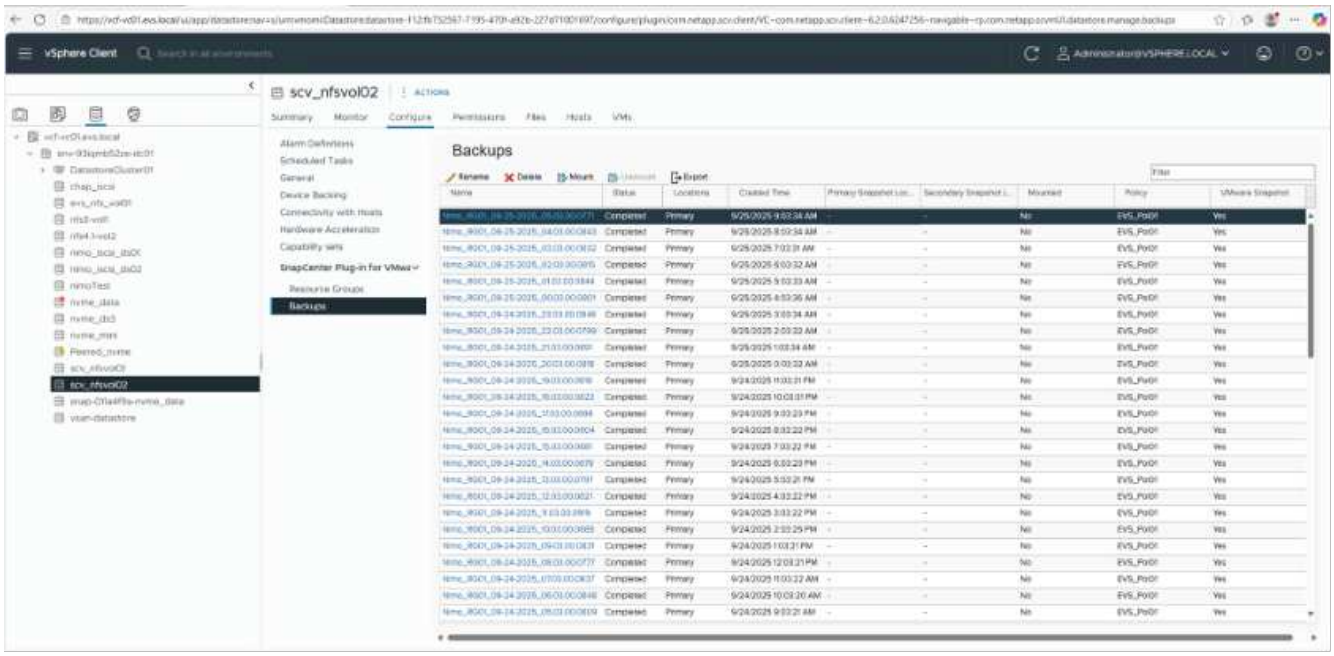
The restore operation can also be triggered from individual VM level.

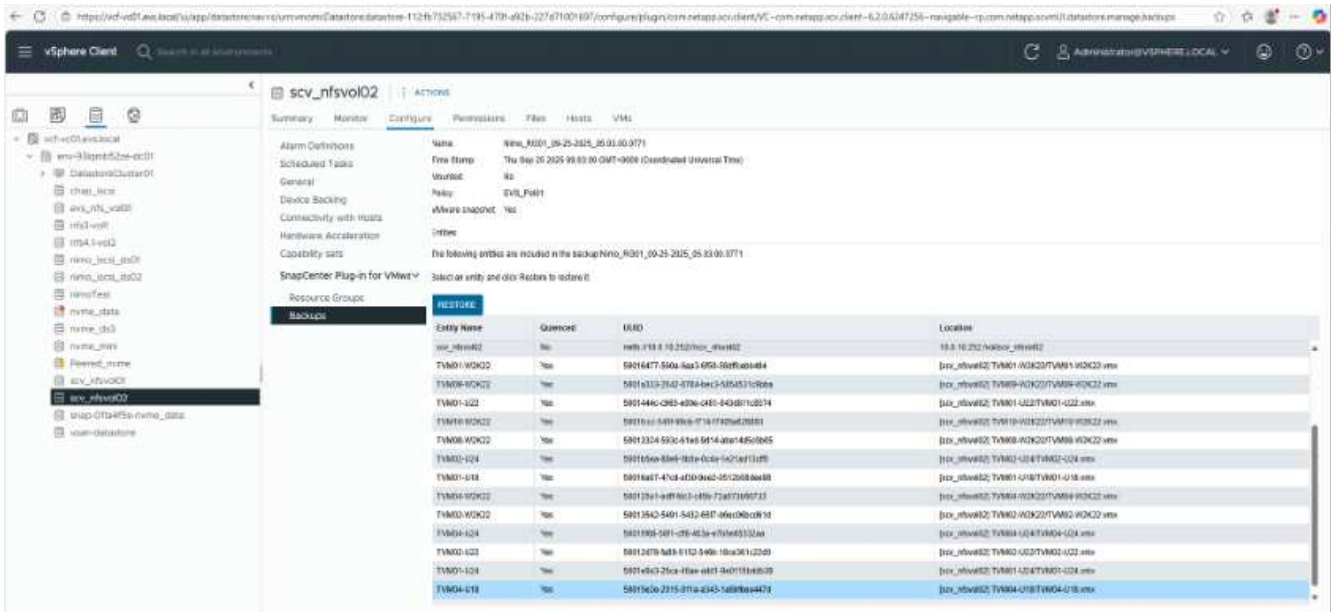
Steps to restore a deleted VM

Restoring a deleted VM is a simple process with SCV.

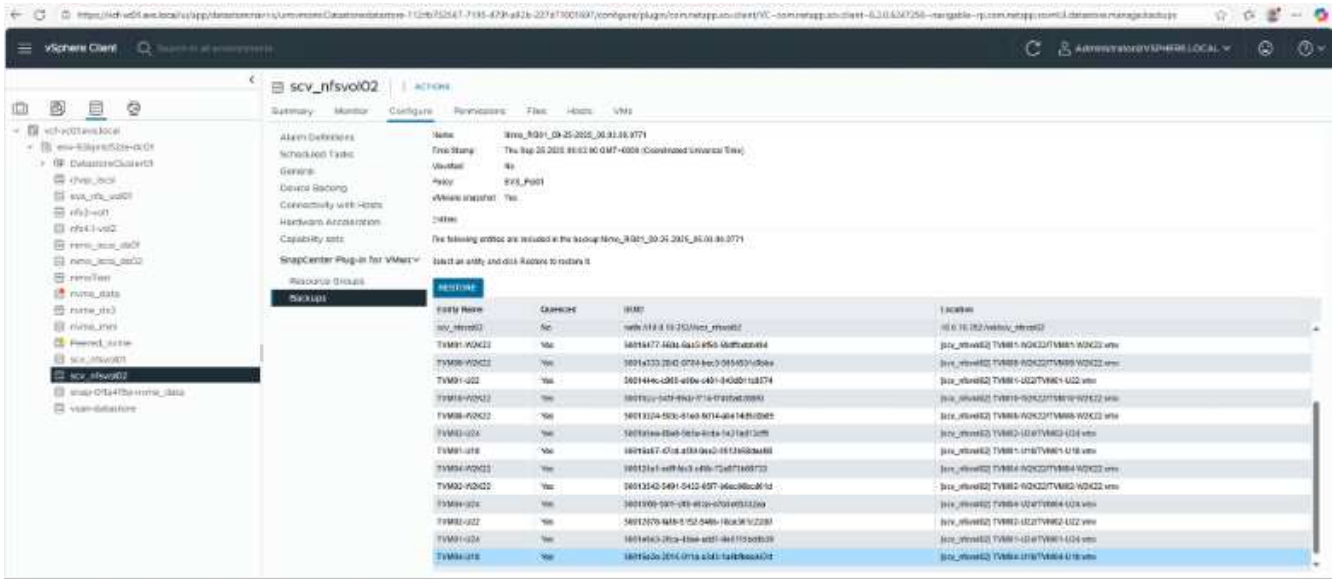
The restore operation for a deleted VM can be performed from the datastore level.

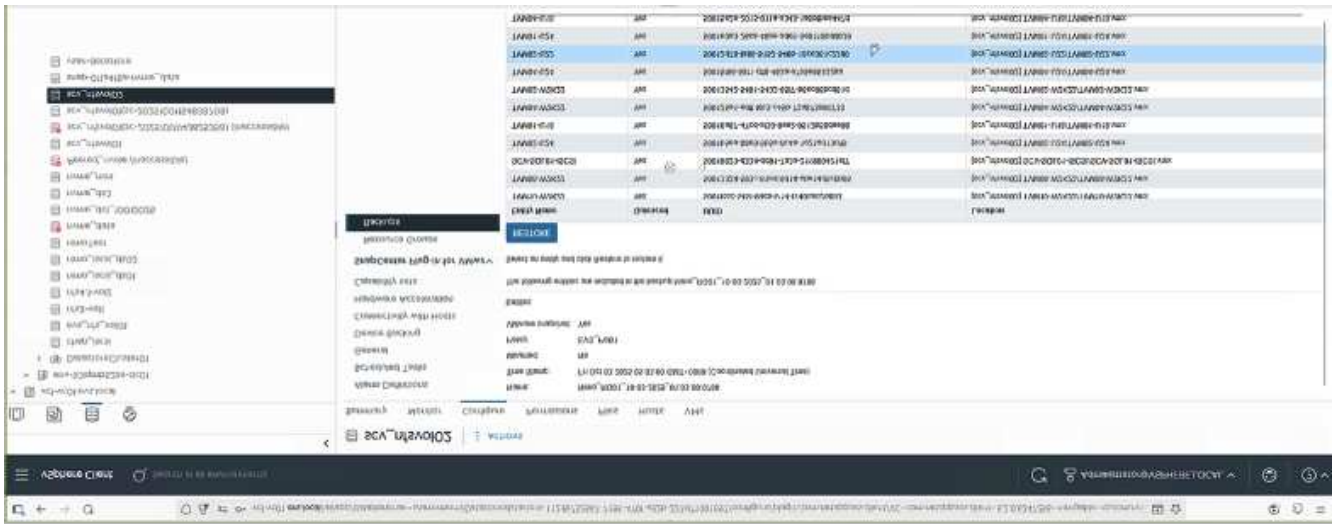
1. Go to respective Datastore > Configure > Backups. There, select the deleted VM and select "Restore".

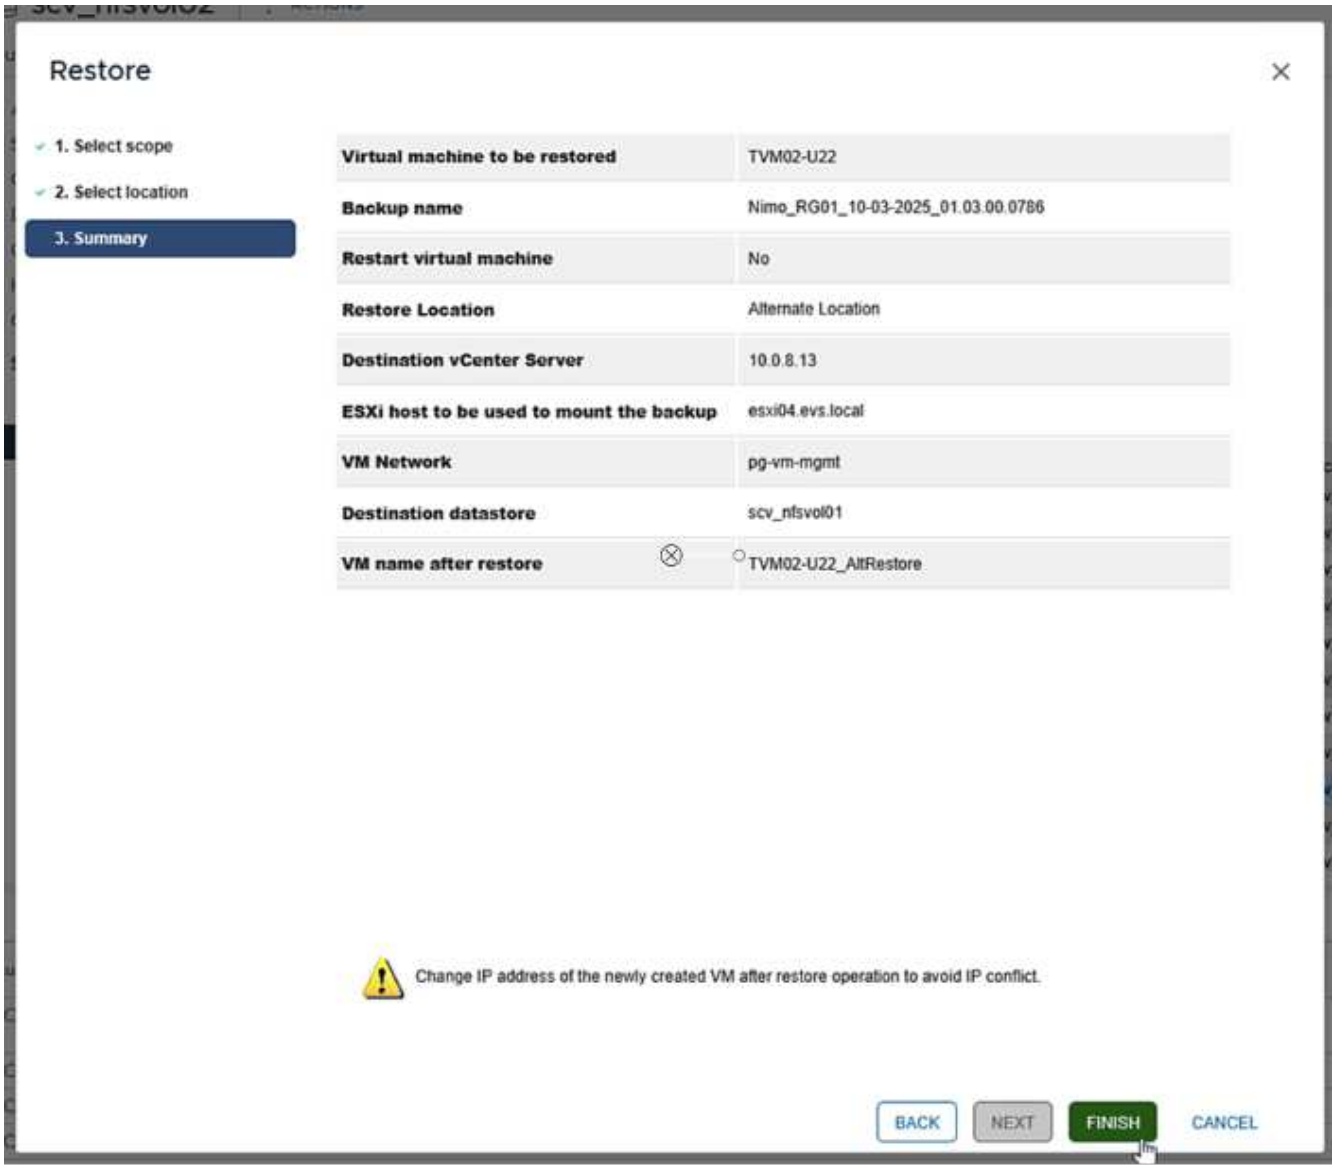


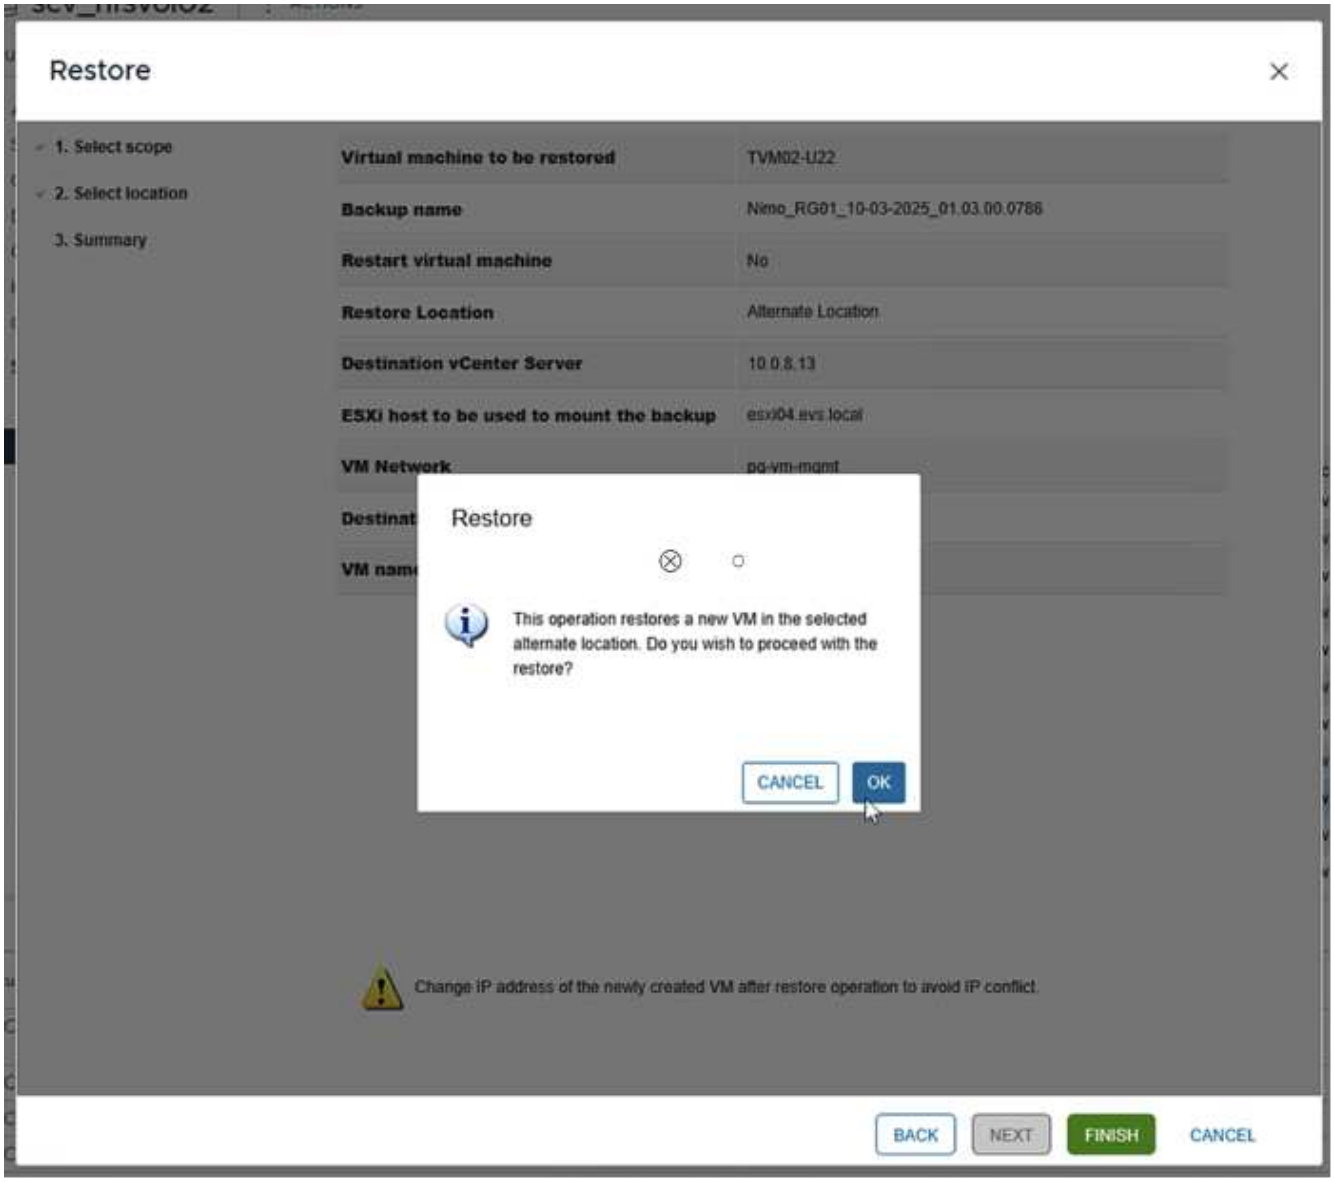


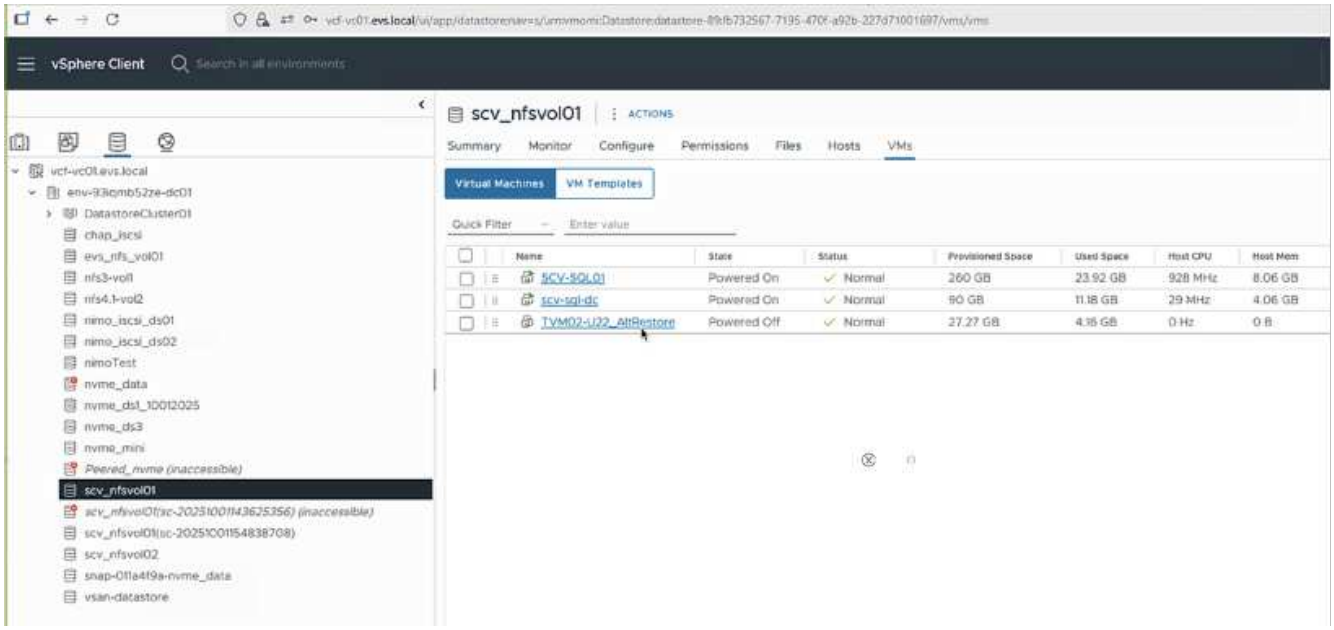
2. The screenshot flow that follows covers the restoration of a VM to an alternate ESXi host and different datastore.


vSphere Client
scv_nfsvol02
Summary
Monitor
Configure
Permissions
Files
Hosts
VMs
SnapCenter Plug-in for VMware
Resource Groups
Backups
RESTORE




Restore
1. Select scope
2. Select location
3. Summary
Virtual machine to be restored
TVM02-U22
Backup name
Nimo_RG01_10-03-2025_01.03.00.0786
Restart virtual machine
No
Restore Location
Alternate Location
Destination vCenter Server
10.0.8.13
ESXi host to be used to mount the backup
esxi04.evs.local
VM Network
pg-vm-mgmt
Destination datastore
scv_nfsvol01
VM name after restore
TVM02-U22_AltRestore
Change IP address of the newly created VM after restore operation to avoid IP conflict.
BACK
NEXT
FINISH
CANCEL


Restore
1. Select scope
2. Select location
3. Summary
Virtual machine to be restored
TVM02-U22
Backup name
Nimo_RG01_10-03-2025_01.03.00.0786
Restart virtual machine
No
Restore Location
Alternate Location
Destination vCenter Server
10.0.8.13
ESXi host to be used to mount the backup
esxi04.evs.local
VM Network
pg-vm-mgmt
Restore
This operation restores a new VM in the selected alternate location. Do you wish to proceed with the restore?
CANCEL
OK
Change IP address of the newly created VM after restore operation to avoid IP conflict.
BACK
NEXT
FINISH
CANCEL

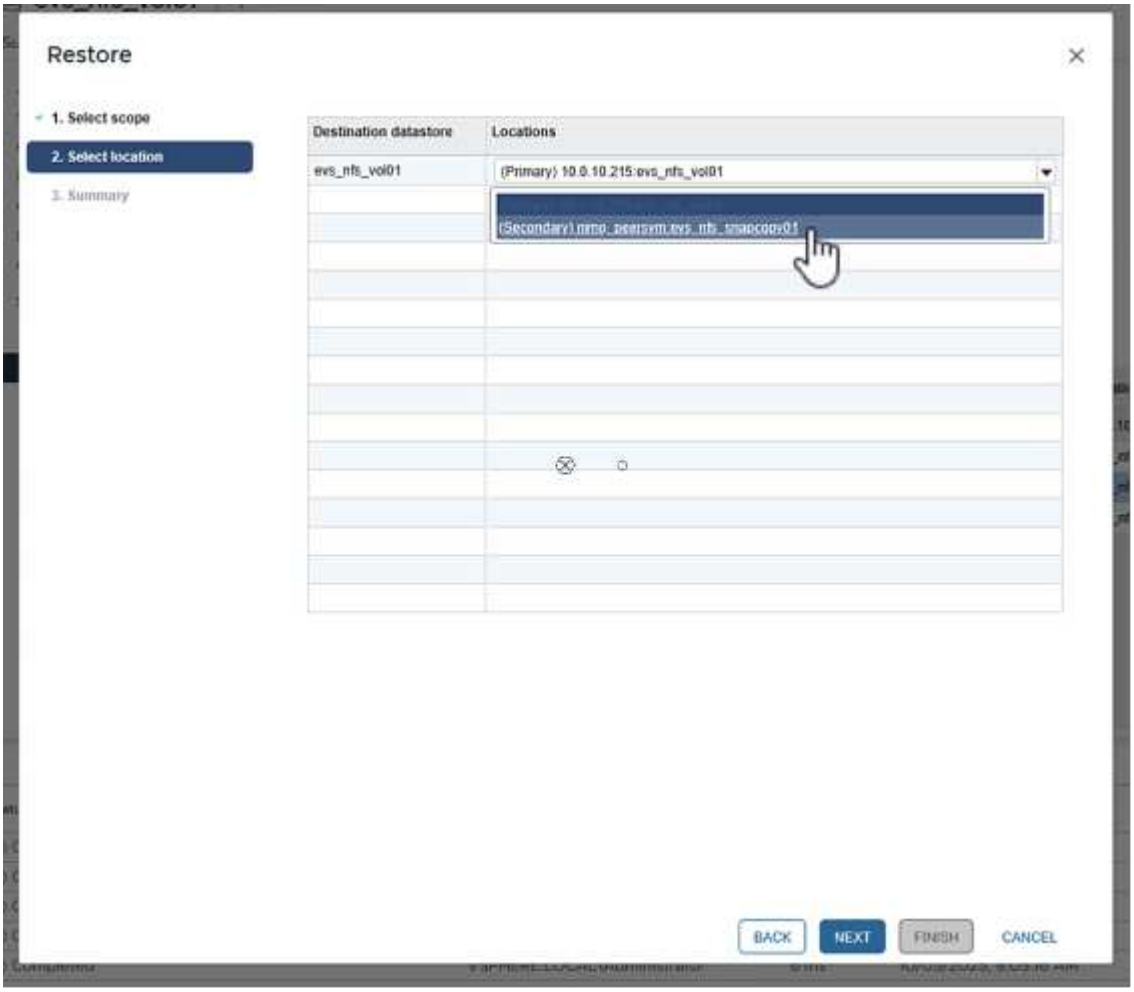

vSphere Client
scv_nfsvol01
ACTIONS
Summary
Monitor
Configure
Permissions
Files
Hosts
VMs
Virtual Machines
VM Templates
Quick Filter
Enter value
Name
State
Status
Provisioned Space
Used Space
Host CPU
Host Mem
SCV-SQL01
Powered On
Normal
260 GB
23.92 GB
928 MHz
8.06 GB
scv-sql-dc
Powered On
Normal
90 GB
11.18 GB
29 MHz
4.06 GB
TVM02-U22_AltRestore
Powered Off
Normal
27.27 GB
4.16 GB
0 Hz
0 B
vcf-vc01.evs.local
env-93ojmb52ze-dc01
DatastoreCluster01
chap_iscsi
evs_nfs_vol01
nfs3-vol1
nfs4.1-vol2
nimo_iscsi_ds01
nimo_iscsi_ds02
nimoTest
nvme_data
nvme_ds1_10012025
nvme_ds3
nvme_mini
Peered_nvme (inaccessible)
scv_nfsvol01
scv_nfsvol01(sc-20251001154838708)
scv_nfsvol02
snap-011a4f9a-nvme_data
vsan-datastore

The destination datastore must have adequate free space and the vCenter specified in the "Restore to Alternate Location"field must be DNS-resolvable.

For better restore efficiency in NFS environments, turn on the VMware vStorage API for Array Integration (VAAI).

Restores can also be performed from secondary locations, such as a SnapMirror destination.

+



Restore procedure for guest files and folders protected by SnapCenter Plug-in for VMware

The following section shows how to restore guest files and folders from VMDKs.

SnapCenter supports file and folder restoration from VMDKs only when the guest operating system is Windows.

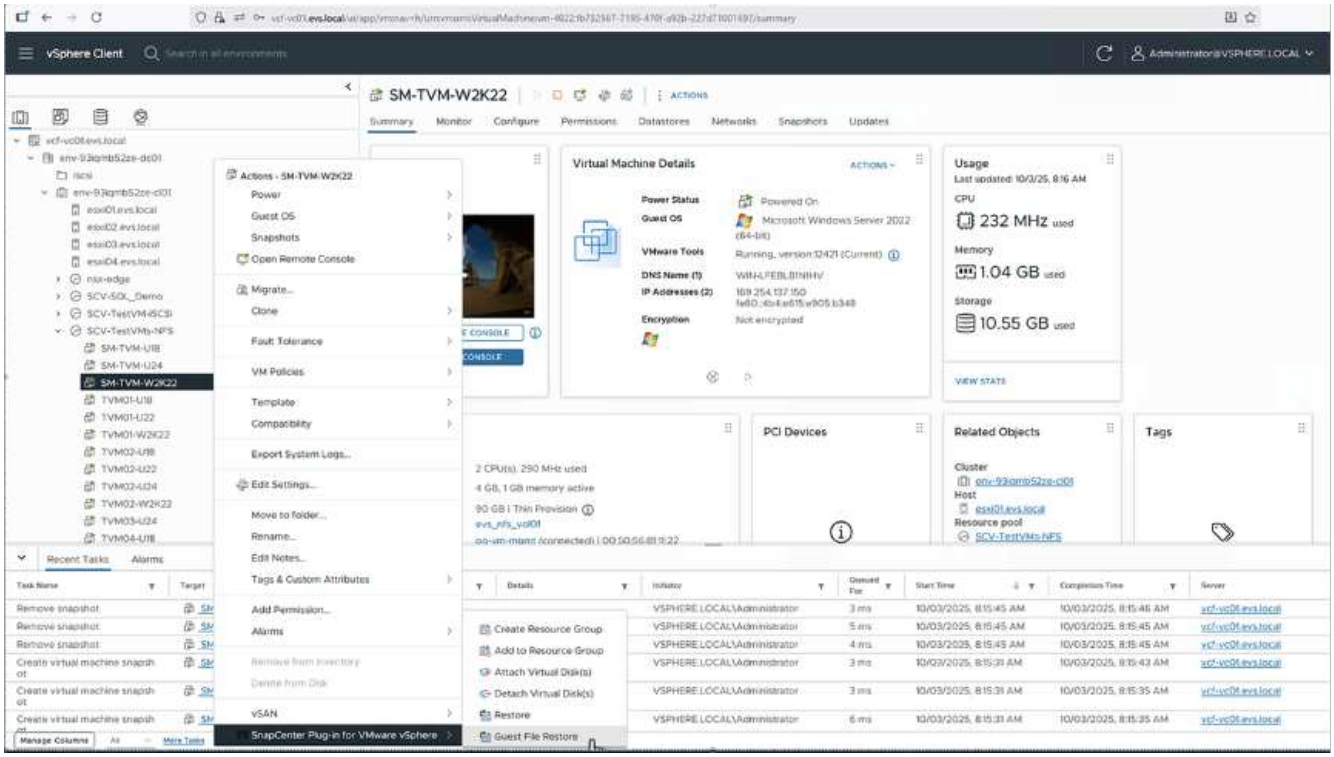
Considerations for guest file and folder restoration with SCV:

- VMware tools must be running on the VM.
- The attached VMDK is accessible for 24 hours by default.
- Only one restore or attach operation can be performed on a virtual machine at a time.

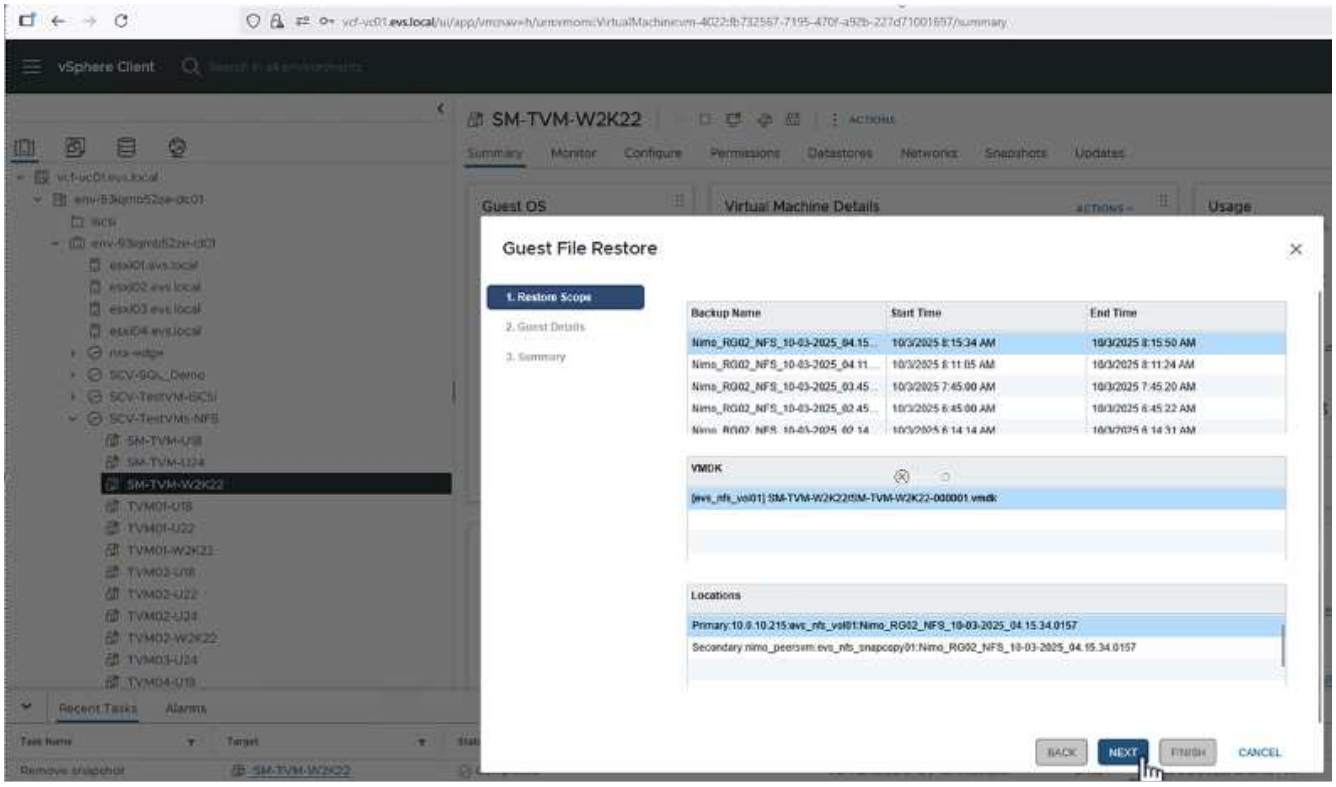
- System, hidden, and encrypted files can be restored; however, their original attributes will not be preserved.
- Avoid restoring encrypted files to encrypted folders or overwriting system files.

Steps for the guest file and folder restoration:

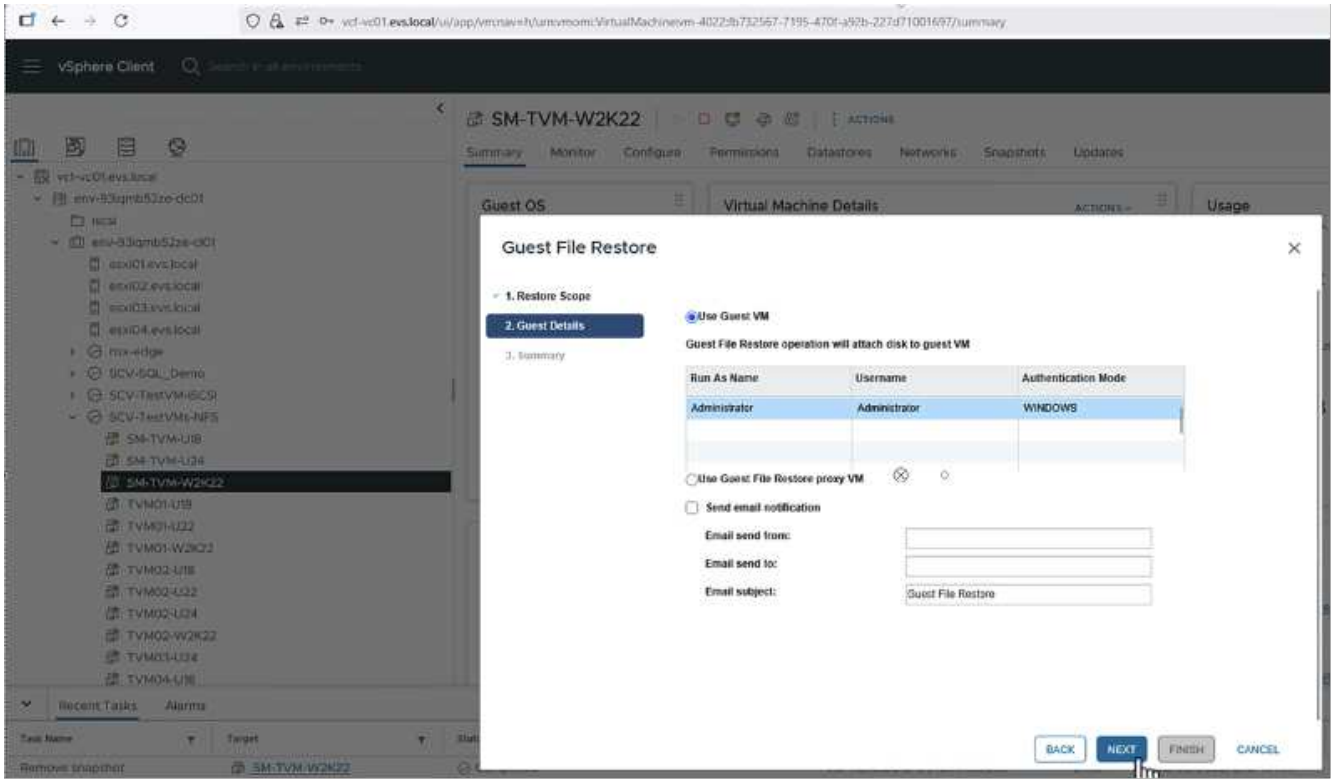
1. Right-click the target VM and select **NetApp SnapCenter › Guest File Restore**.



2. Choose the backup containing the required VMDK. Specify the VMDK and its location (primary or secondary).



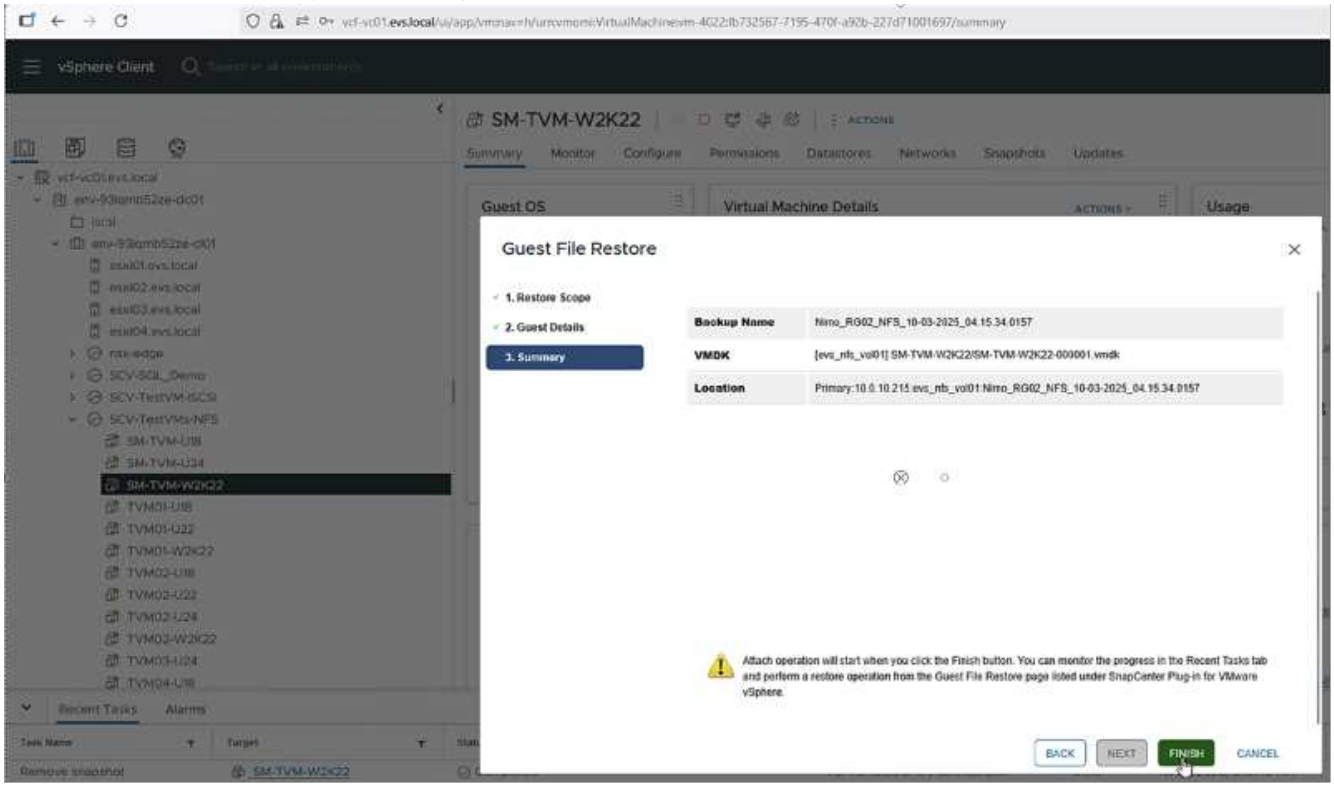
3. Choose "Disk Attachment Method".

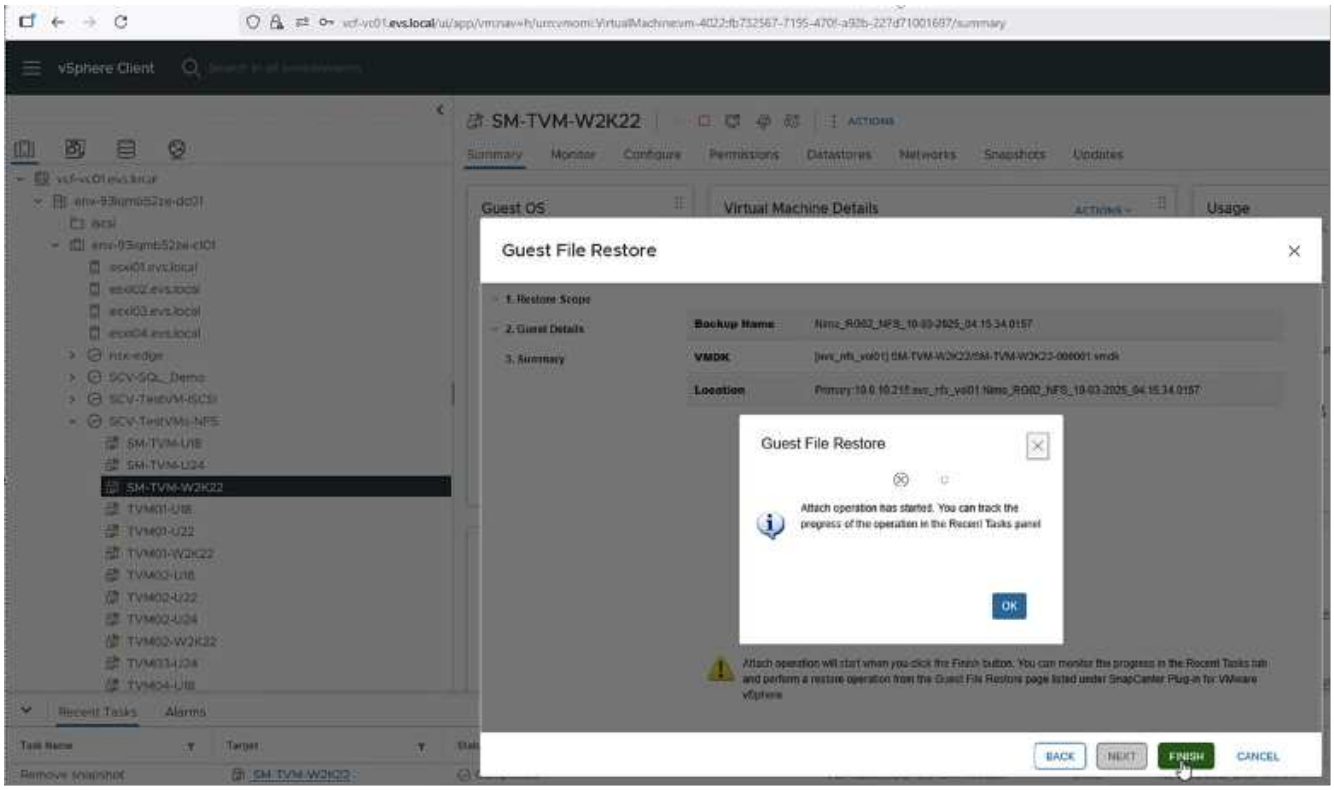


- Guest VM Option: Attach the disk to the original VM (requires pre-configured credentials). This is the option used in this demonstration.
- Proxy VM Option: Attach the disk to a designated proxy VM (must be configured in advance).

4. Wait for the attach operation to finish before browsing and restoring files. Once the operation completes,

SnapCenter automatically creates a guest file restore session.





5. Browse to the VMDK in the Guest File Restore session, select the required files or folders, and restore them to the specified location.

Protection procedures for databases running on VMware VMDKs with SnapCenter for SQL (SCSQL)

The following section shows how to protect databases running on VMware VMDKs with SnapCenter for SQL (SCSQL).

Setting up SCSQL

The set up for using SCSQL to protect databases on VMDKs are as follows:

1. Install and configure SnapCenter Server.
2. Log in to SnapCenter.
3. Configure the SnapCenter environment:
   - Add or assign storage system connections.
   - Create credentials.
4. For Microsoft SQL Server databases on VMware VMDKs:
   - Deploy SCV (already covered in the setup section).
   - Register the plug-in with SnapCenter.
5. Add SQL server hosts and install plug-ins:
   - Configure the SQL plug-ins.
   - Configure the protection plan and schedules for database and logs.

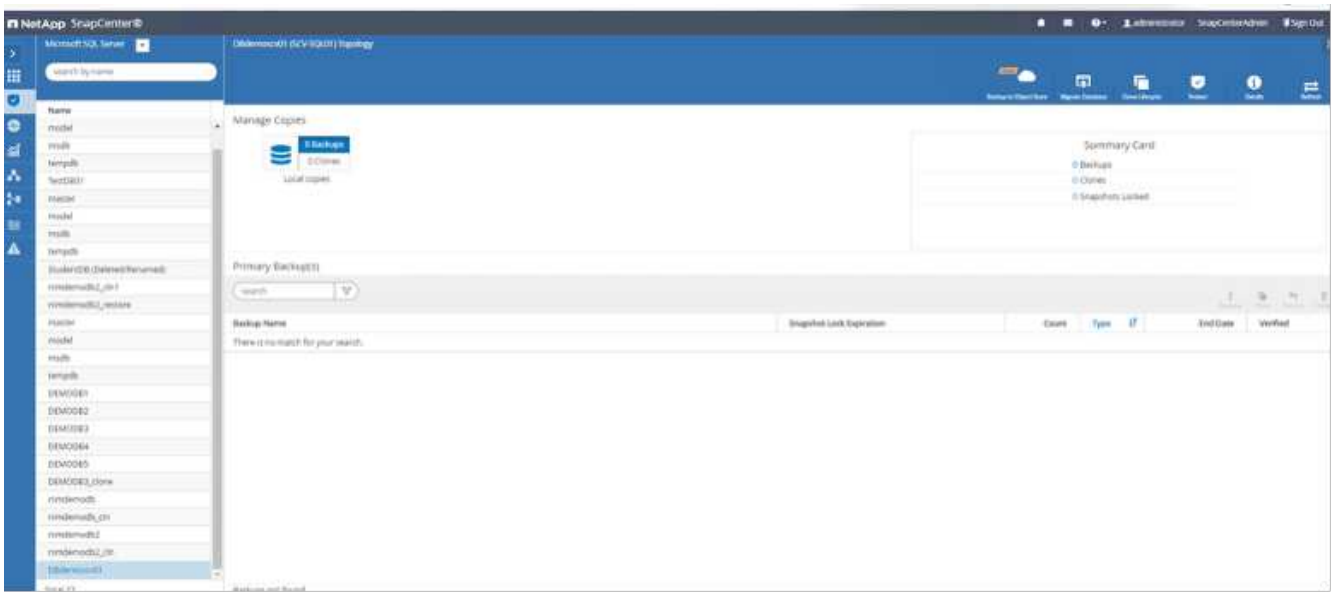


Setting up a backup policy

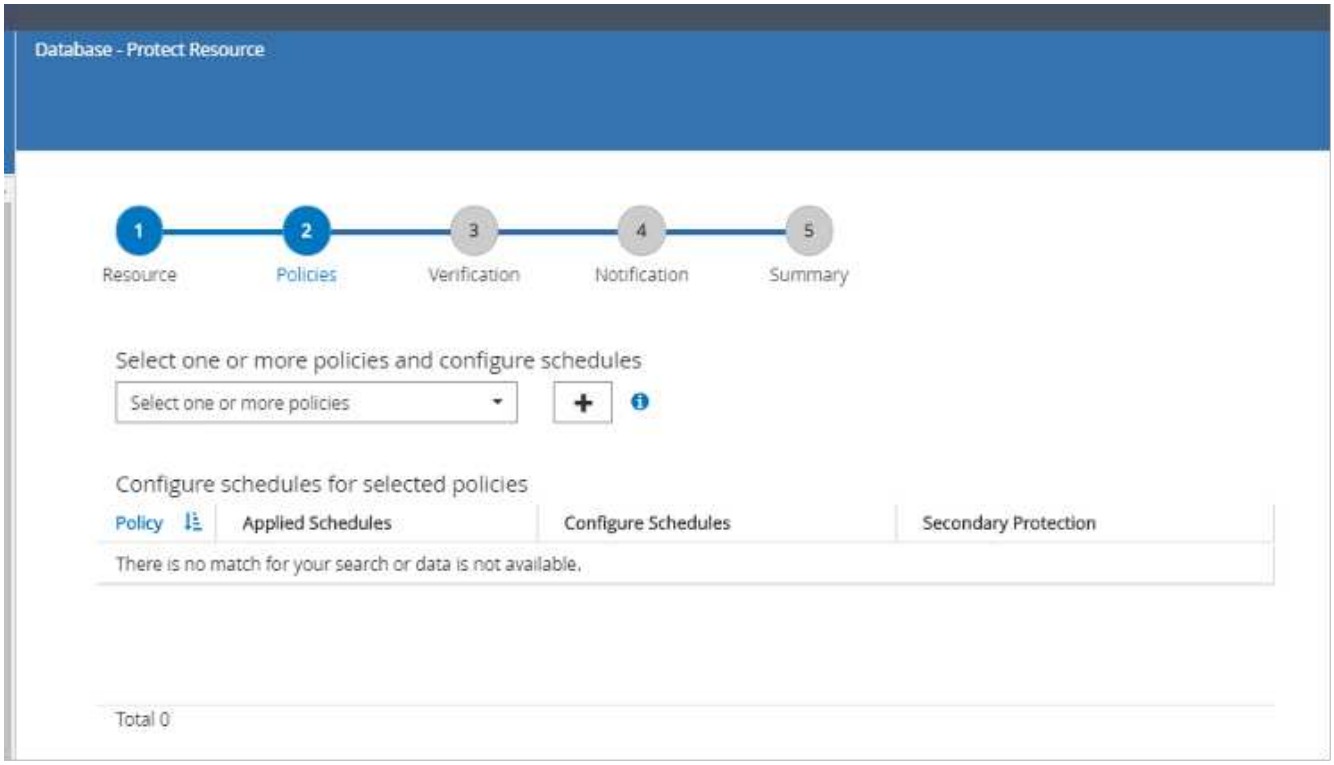
The following steps will show how to set up a backup policy.

1. Select the database from the resource tab and select "Protect" from the upper-right corner.

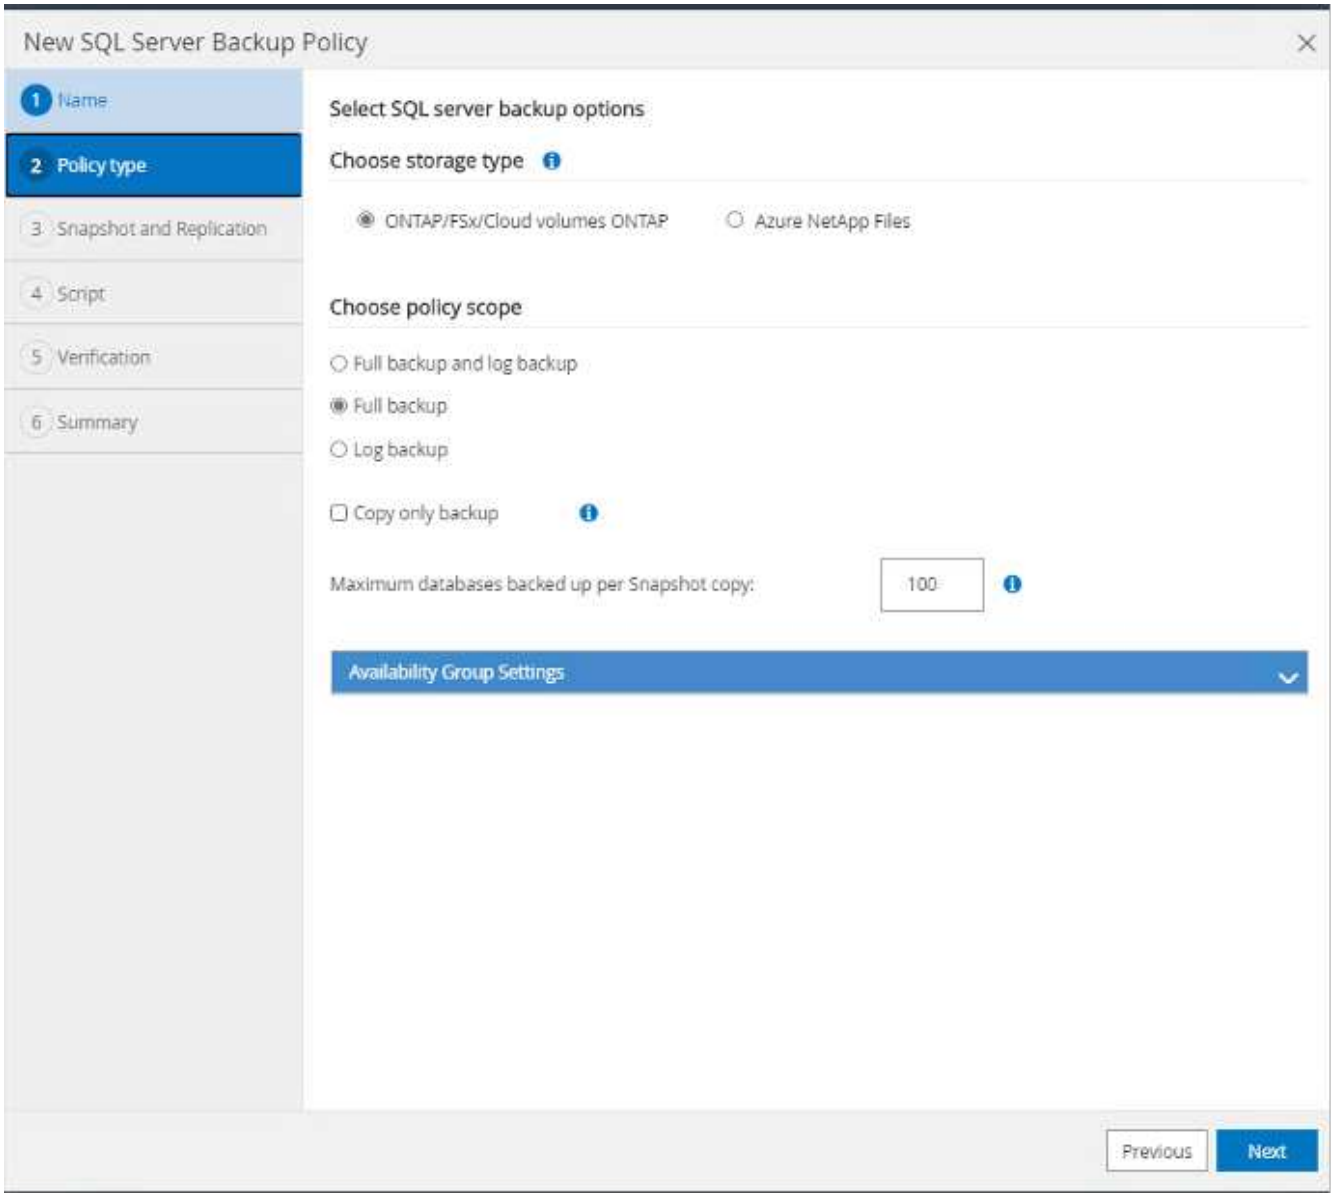
   Validations are performed against a SQL Server database running on the guest VM.



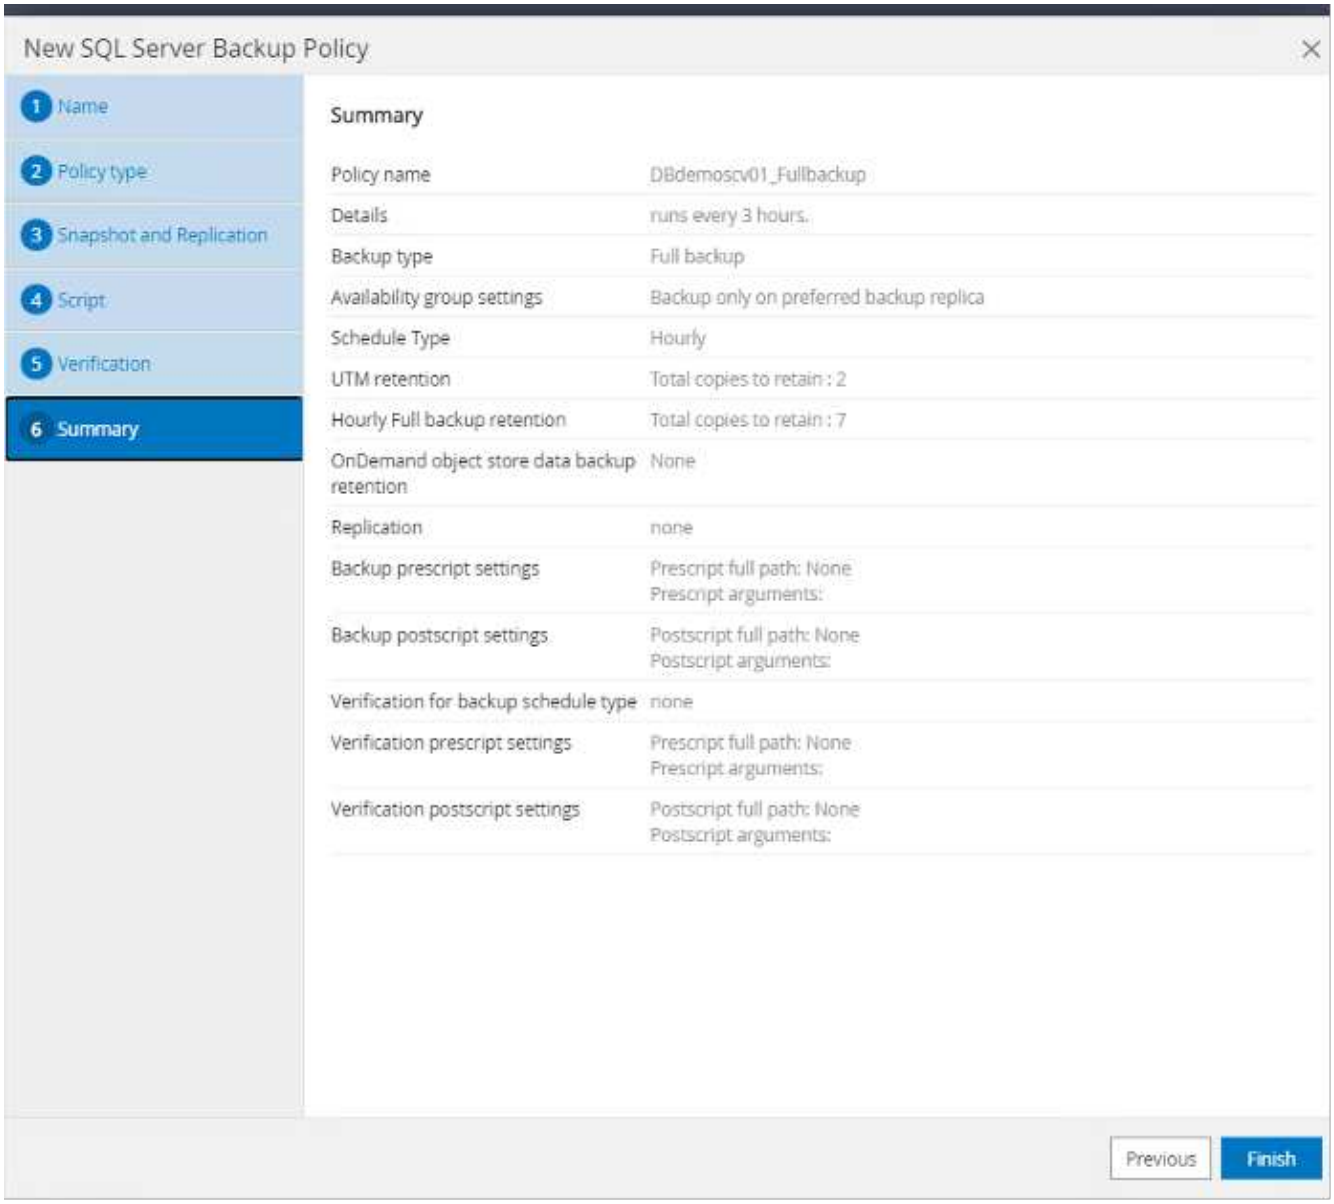
2. Click on the plus sign icon to configure a SQL Server backup policy.



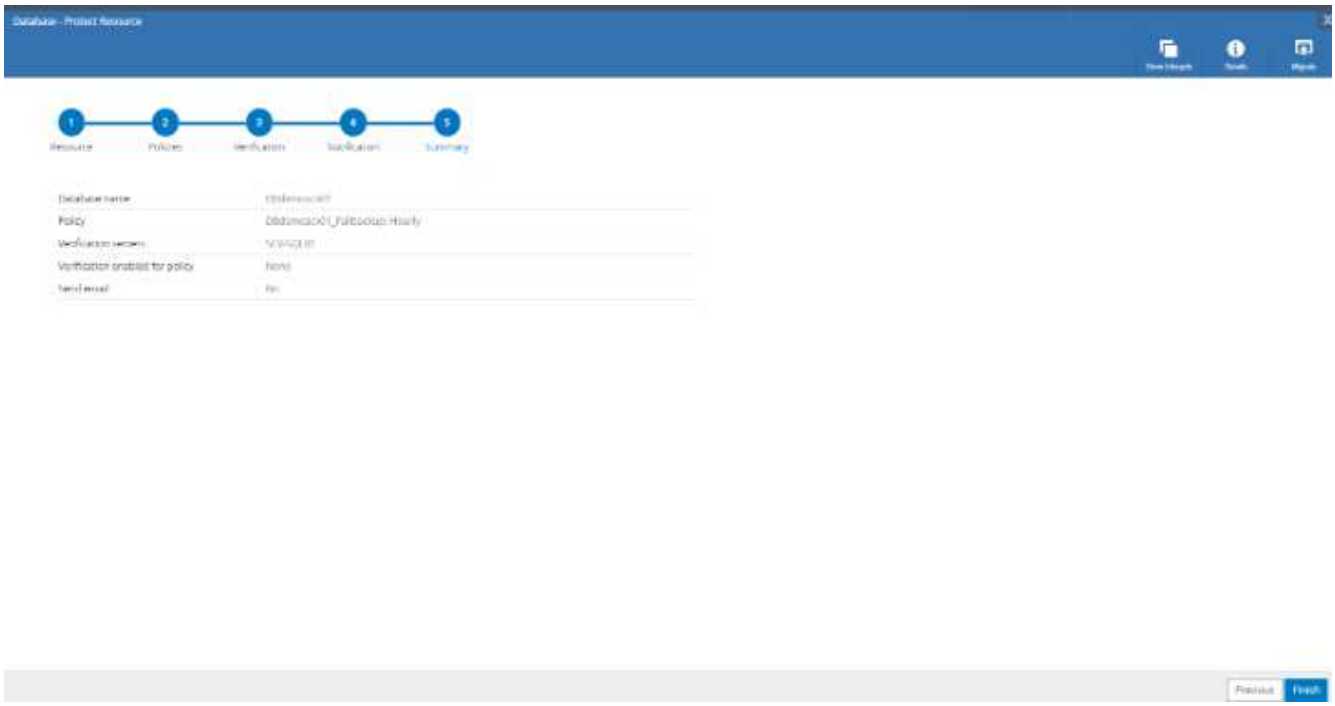
3. In this demo, "Full backup" is selected. Select "Next" to continue.



4. On the Summary page, select "Finish".



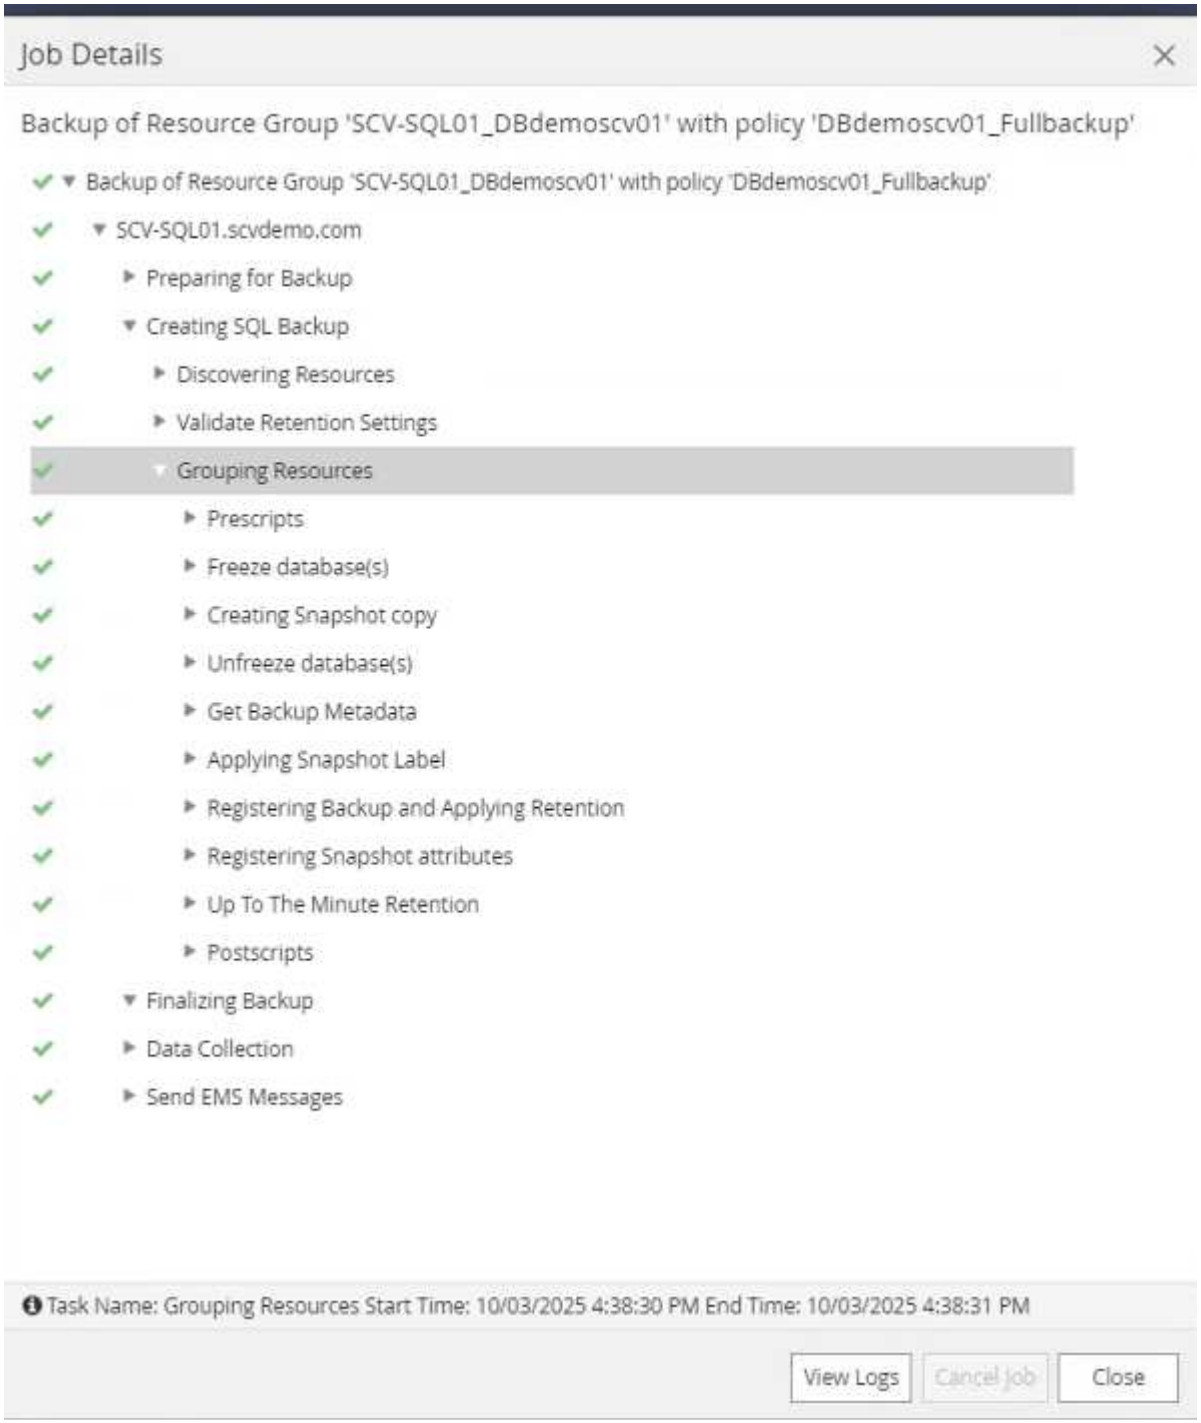
5. Proceed to "Finish".



Monitoring a backup

Once you have set up a backup policy, you can monitor your back up operations as they take place. You can do this either by clicking on Backup now to trigger an on-demand back up operation or by waiting for your scheduled backup policy to kick in.

1. Monitor the job progress using **Monitor › Jobs**.
2. Select the backup job, and then click "Details" to view the job details as below.

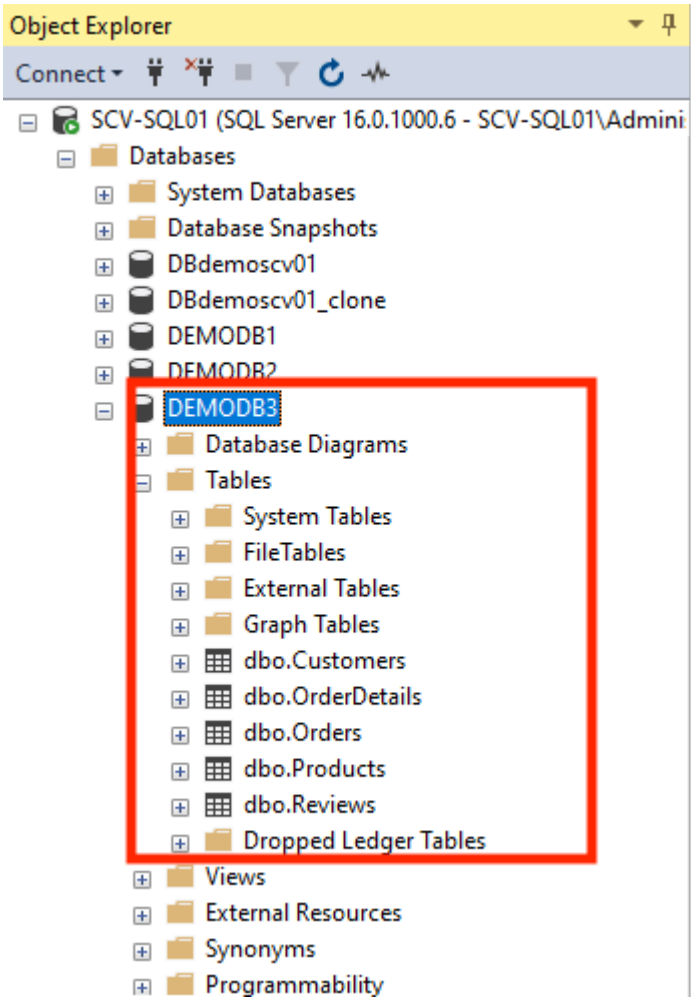


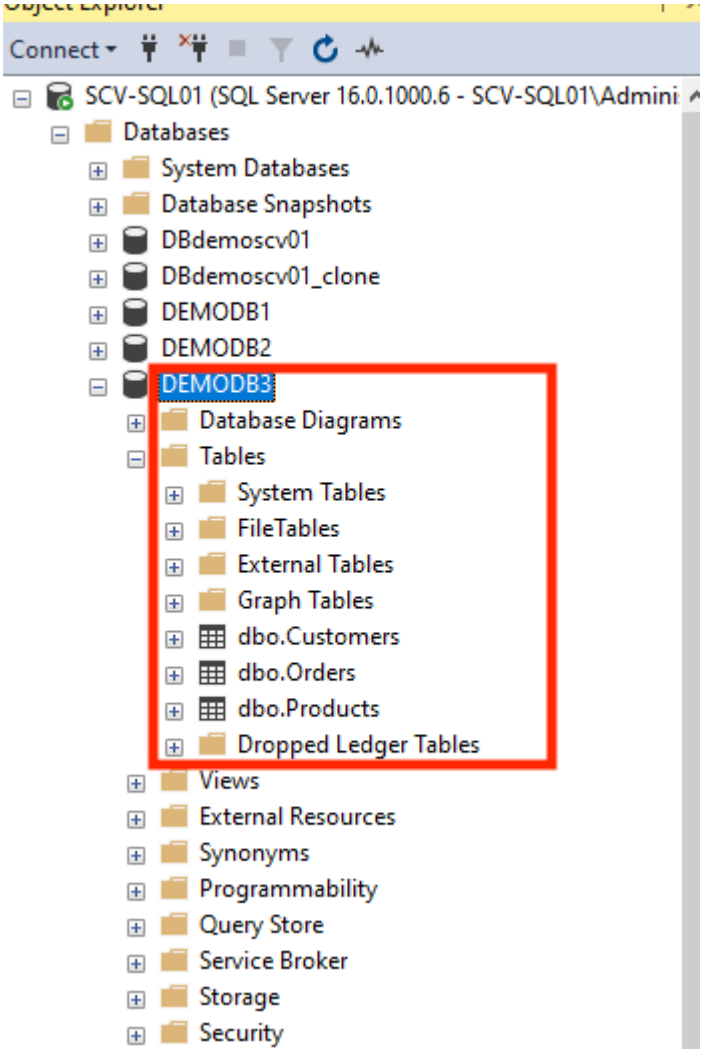
Restore of databases residing on the VMDKs

The database on the VMDK can be restored to a specific point in time.

We can simulate the importance of this by deleting some tables in your database which will then be corrected by performing a point-in-time restore.

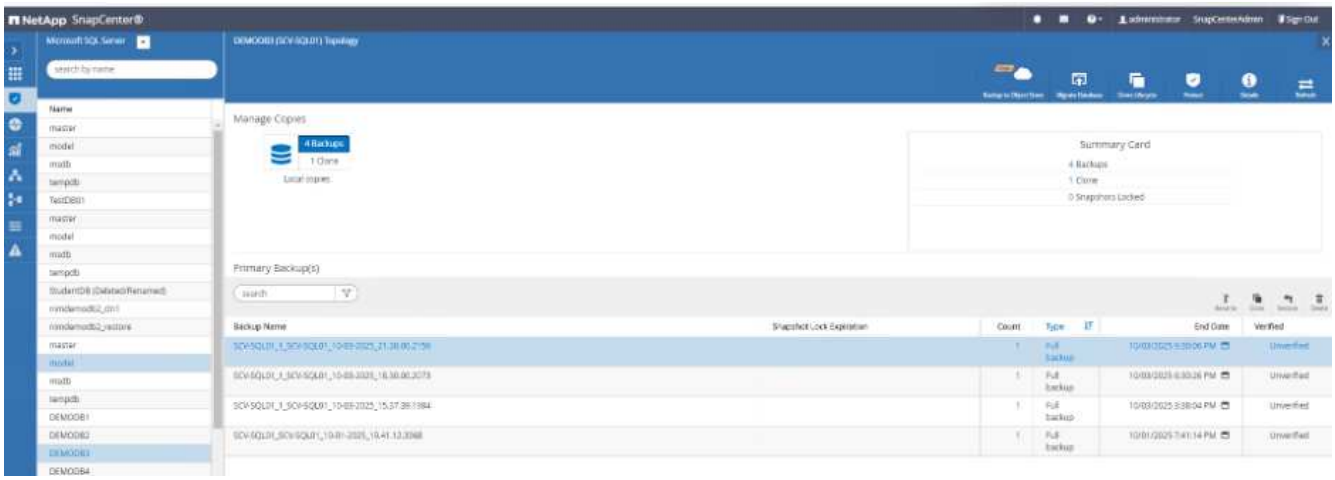
Below you can see that several tables from our database shown in the first screenshot have been deleted in the second screenshot:


Object Explorer
Connect
SCV-SQL01 (SQL Server 16.0.1000.6 - SCV-SQL01\Admini
Databases
System Databases
Database Snapshots
DBdemoscv01
DBdemoscv01_clone
DEMODB1
DEMODB2
DEMODB3
Database Diagrams
Tables
System Tables
FileTables
External Tables
Graph Tables
dbo.Customers
dbo.OrderDetails
dbo.Orders
dbo.Products
dbo.Reviews
Dropped Ledger Tables
Views
External Resources
Synonyms
Programmability

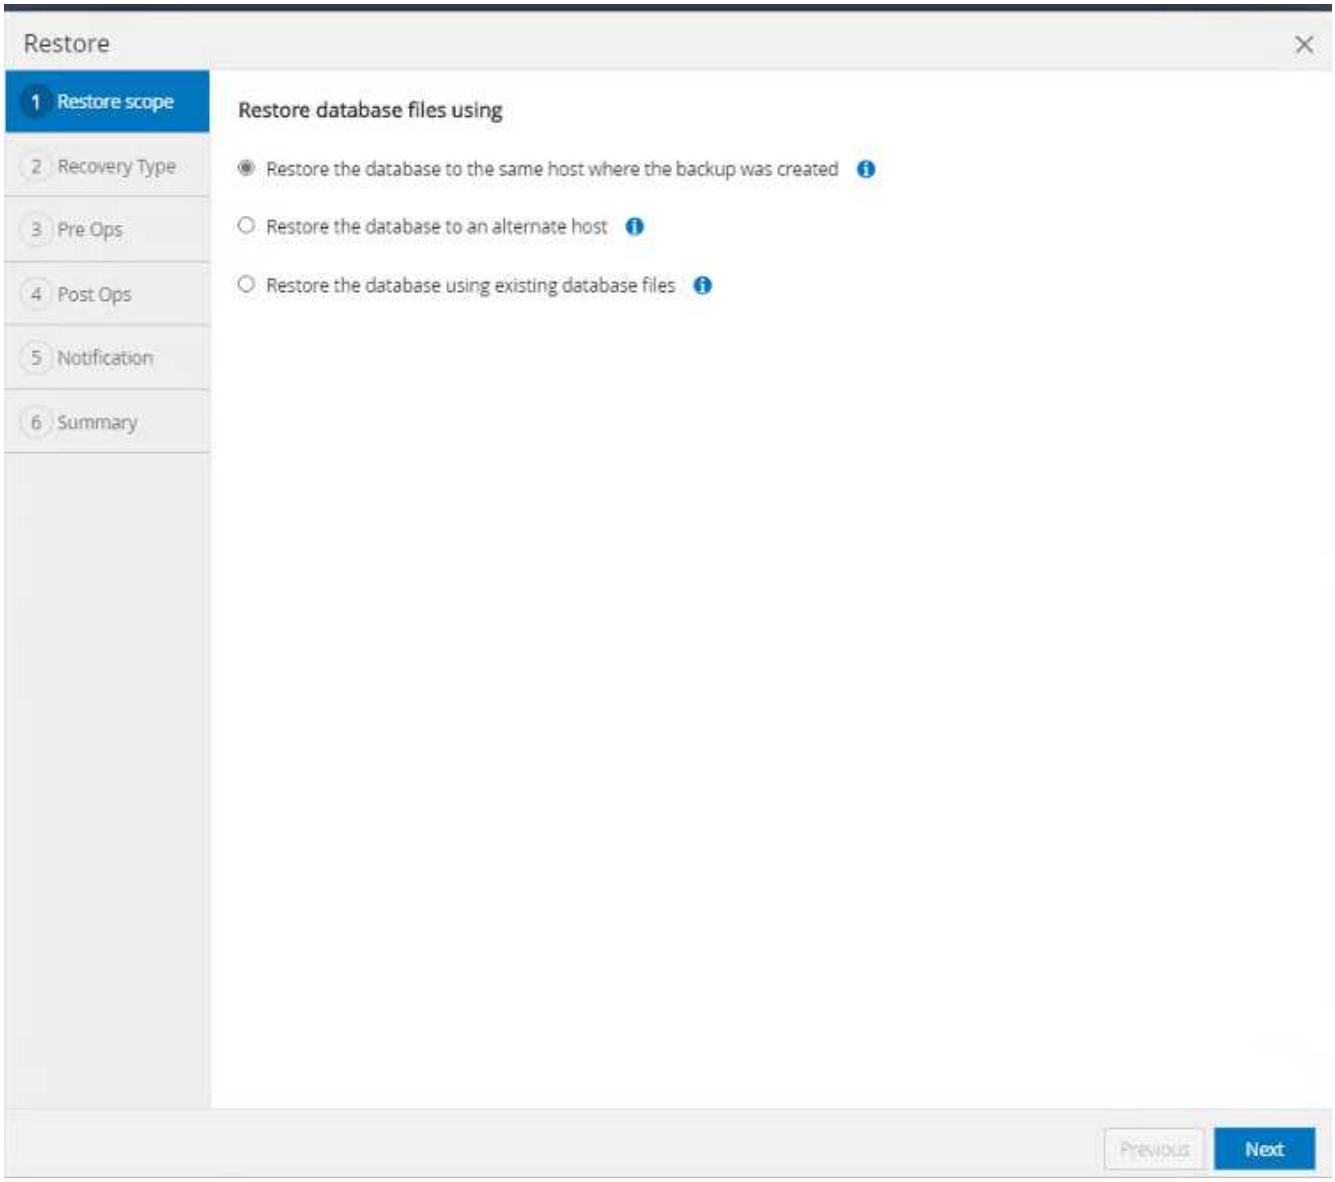


The following steps show how to carry out a restoration to recover those deleted tables:

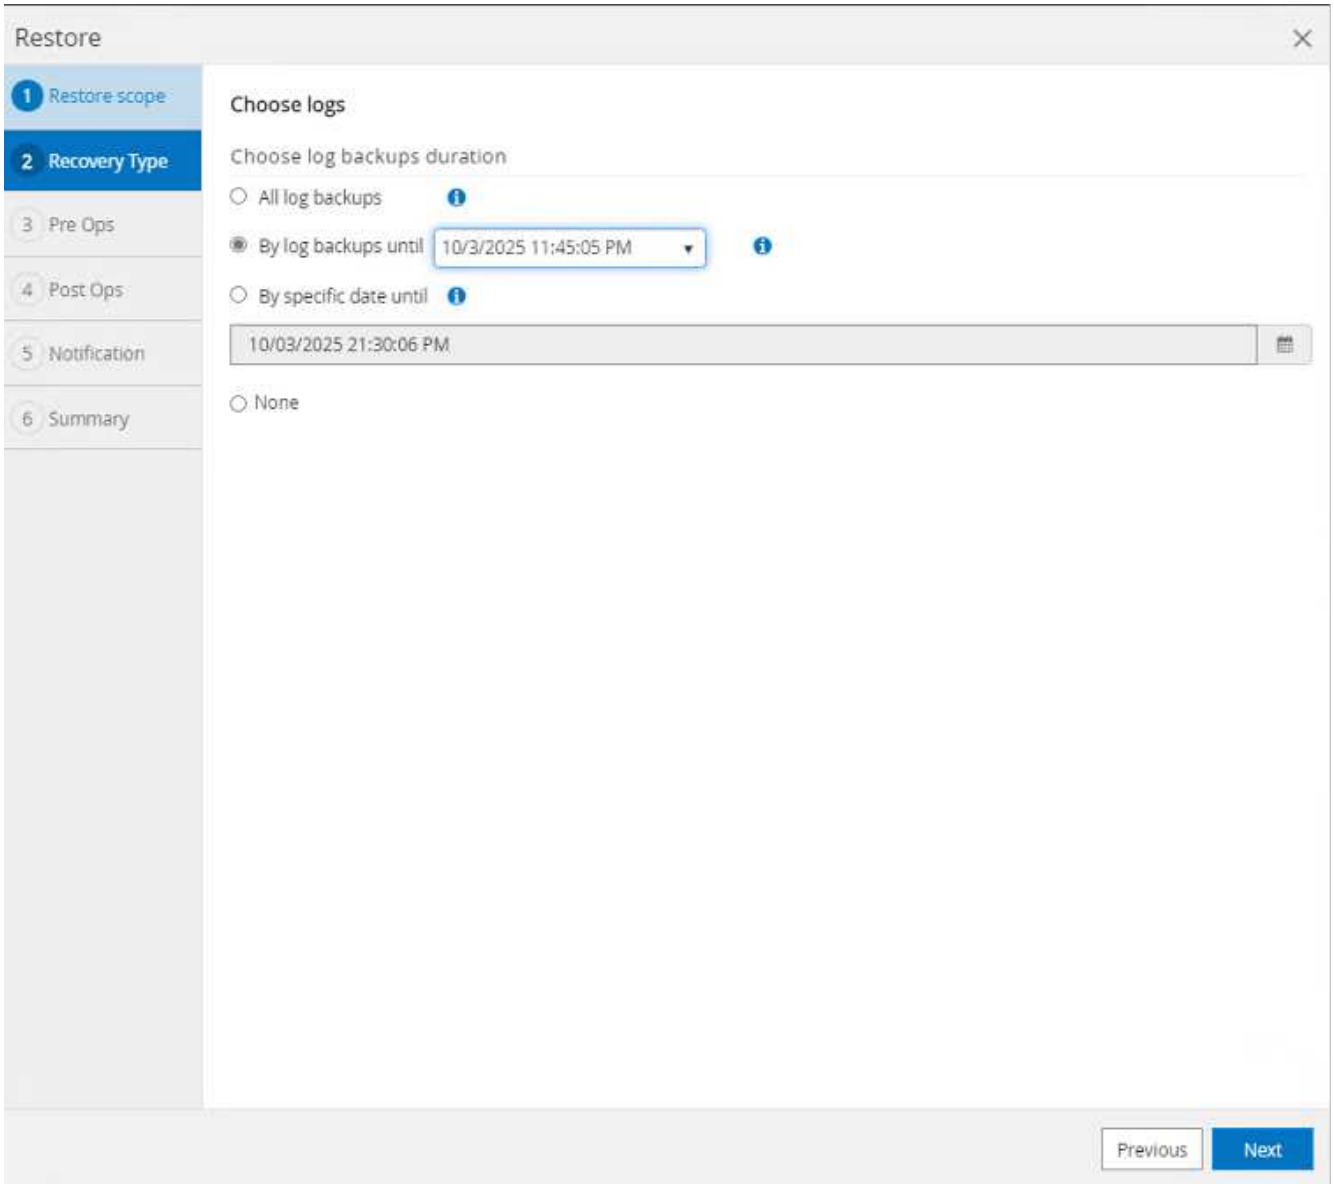
1. Select the database from the resource group.



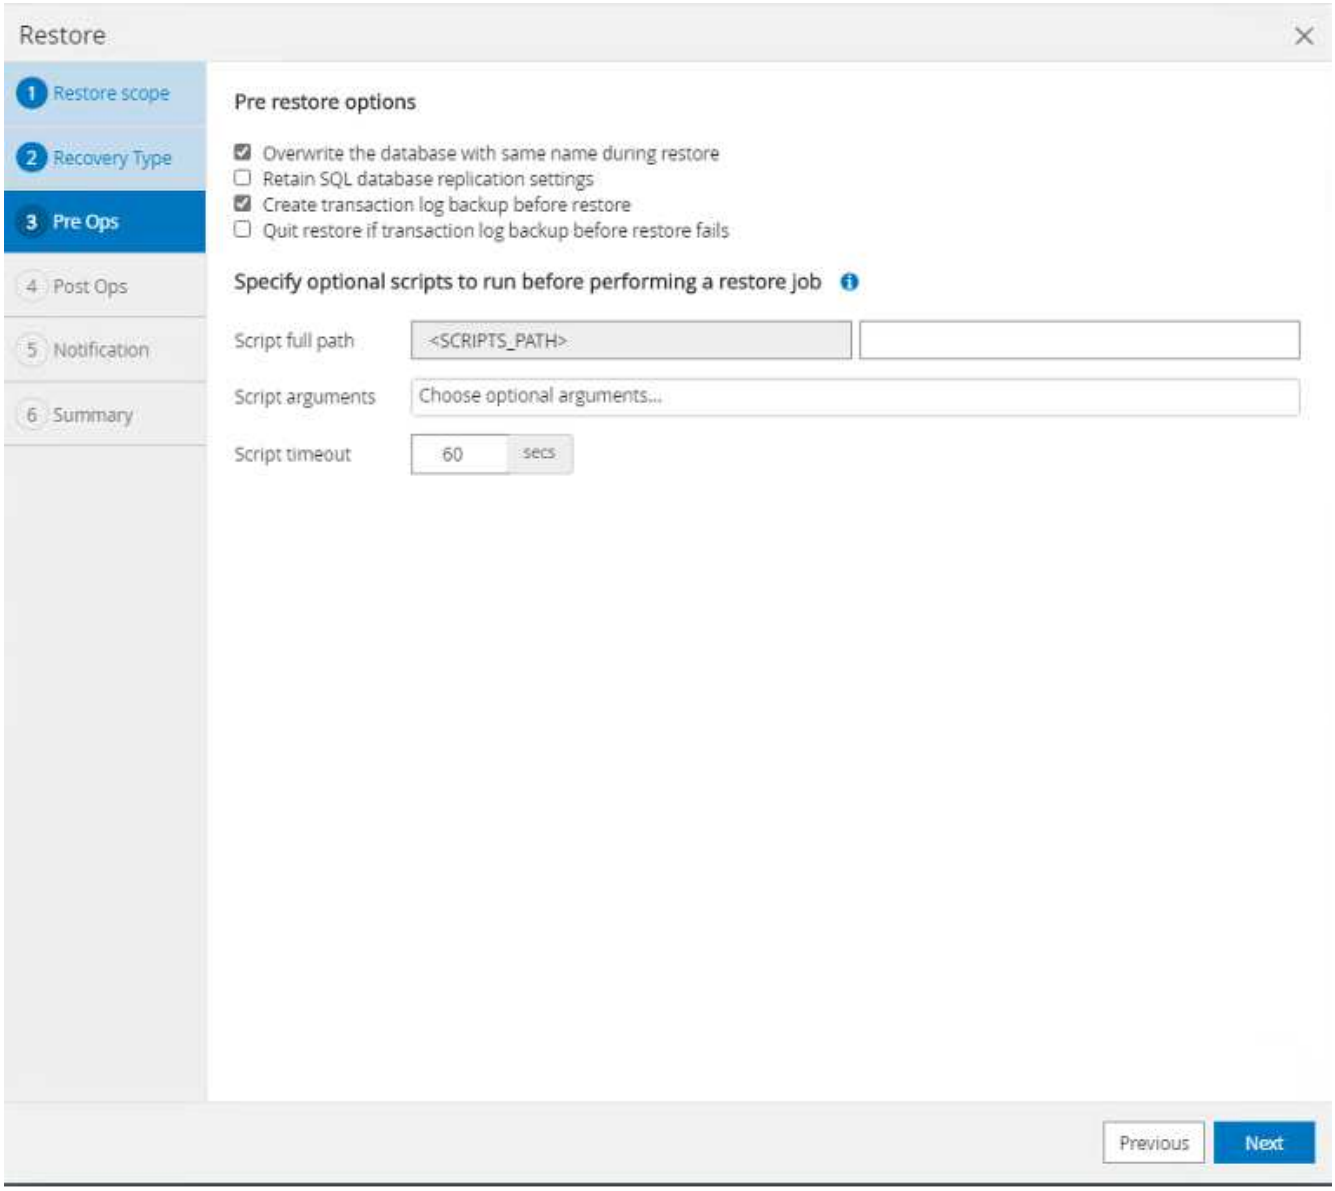
2. In the Restore page, select the option "Restore the database to the same host where the backup was created".



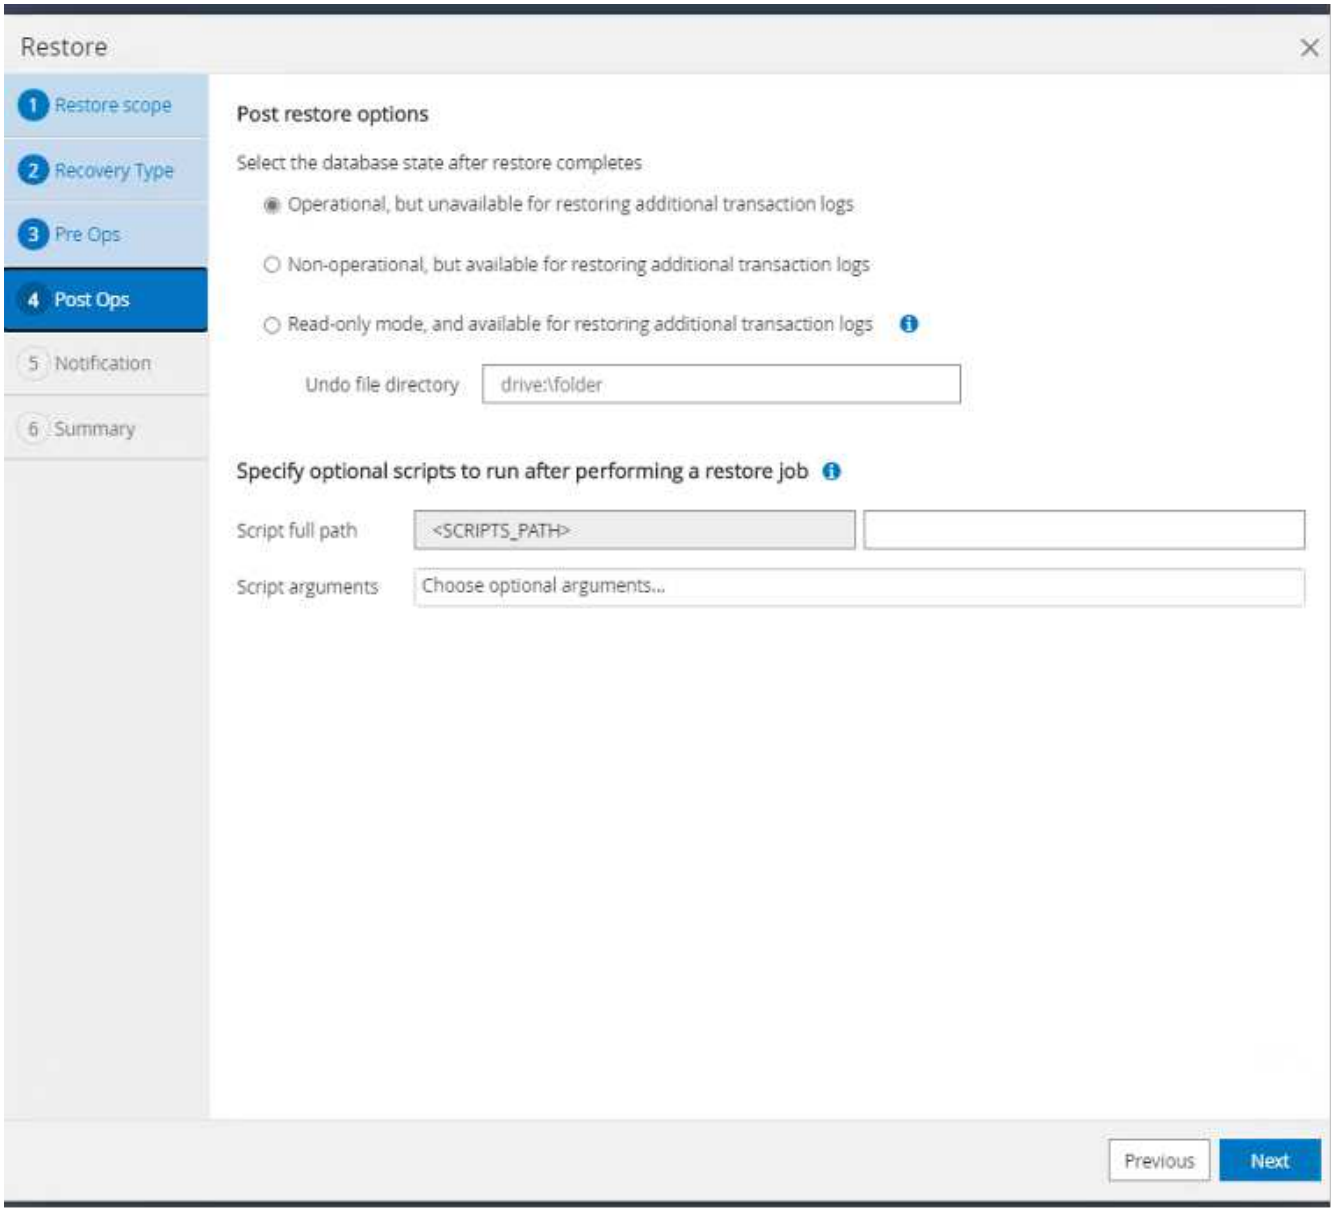
3. Select the option to use a log backup to a point in time.



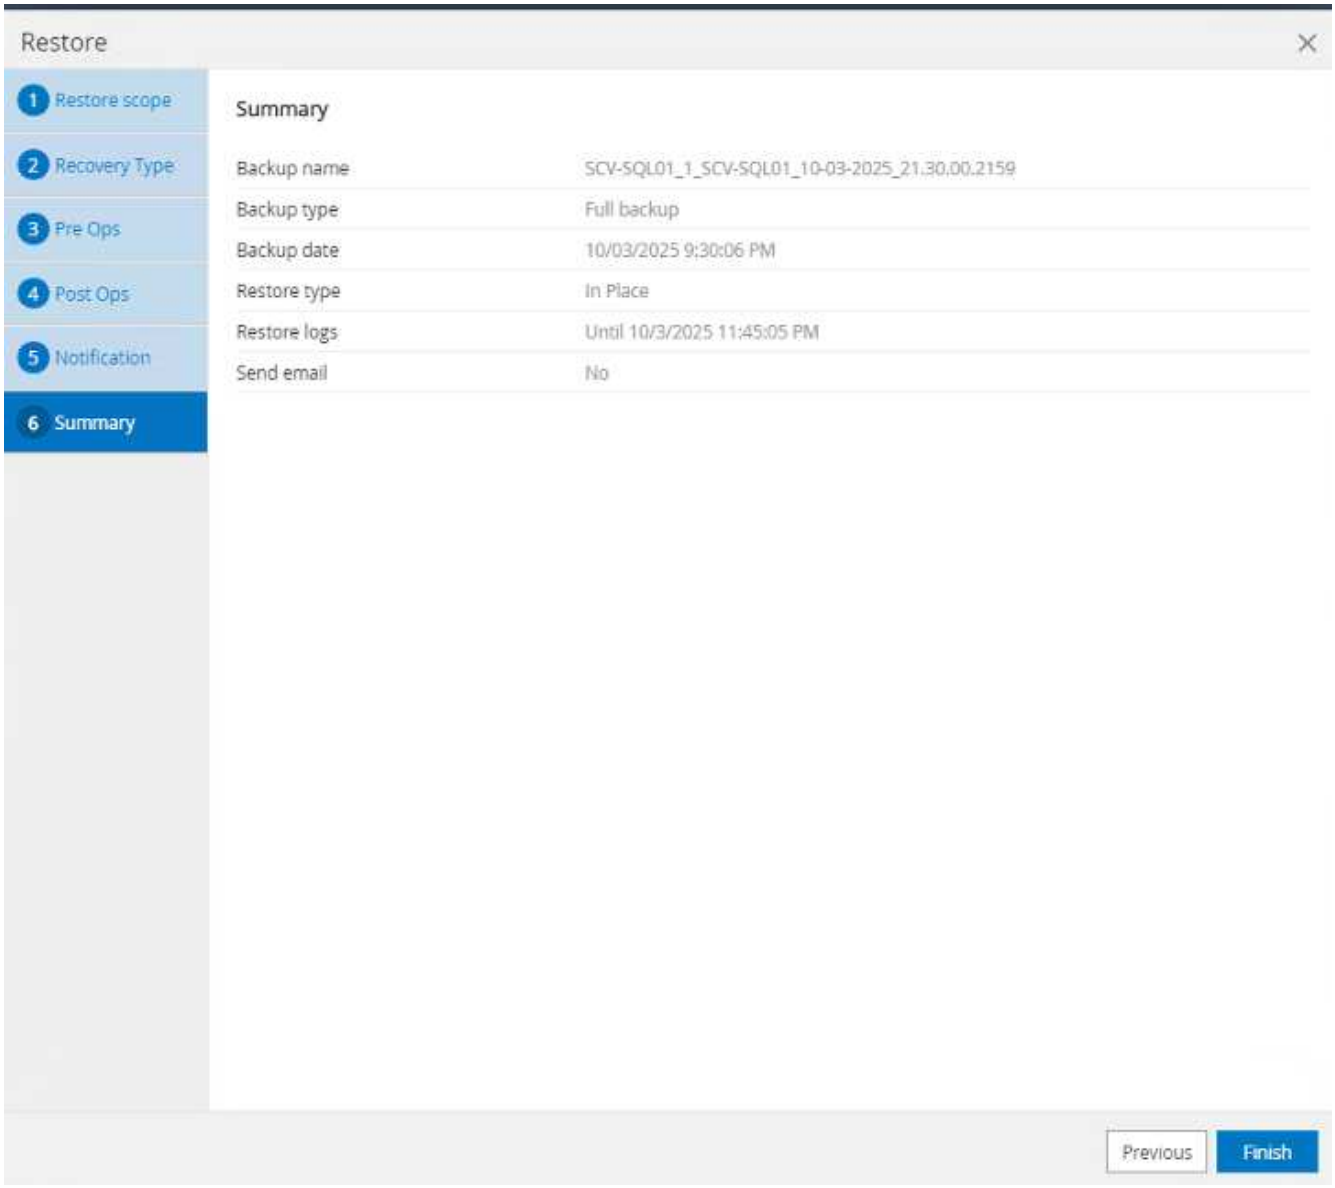
4. Select "Overwrite database with same name during restore, Create transaction log backup before restore", and select "Next".



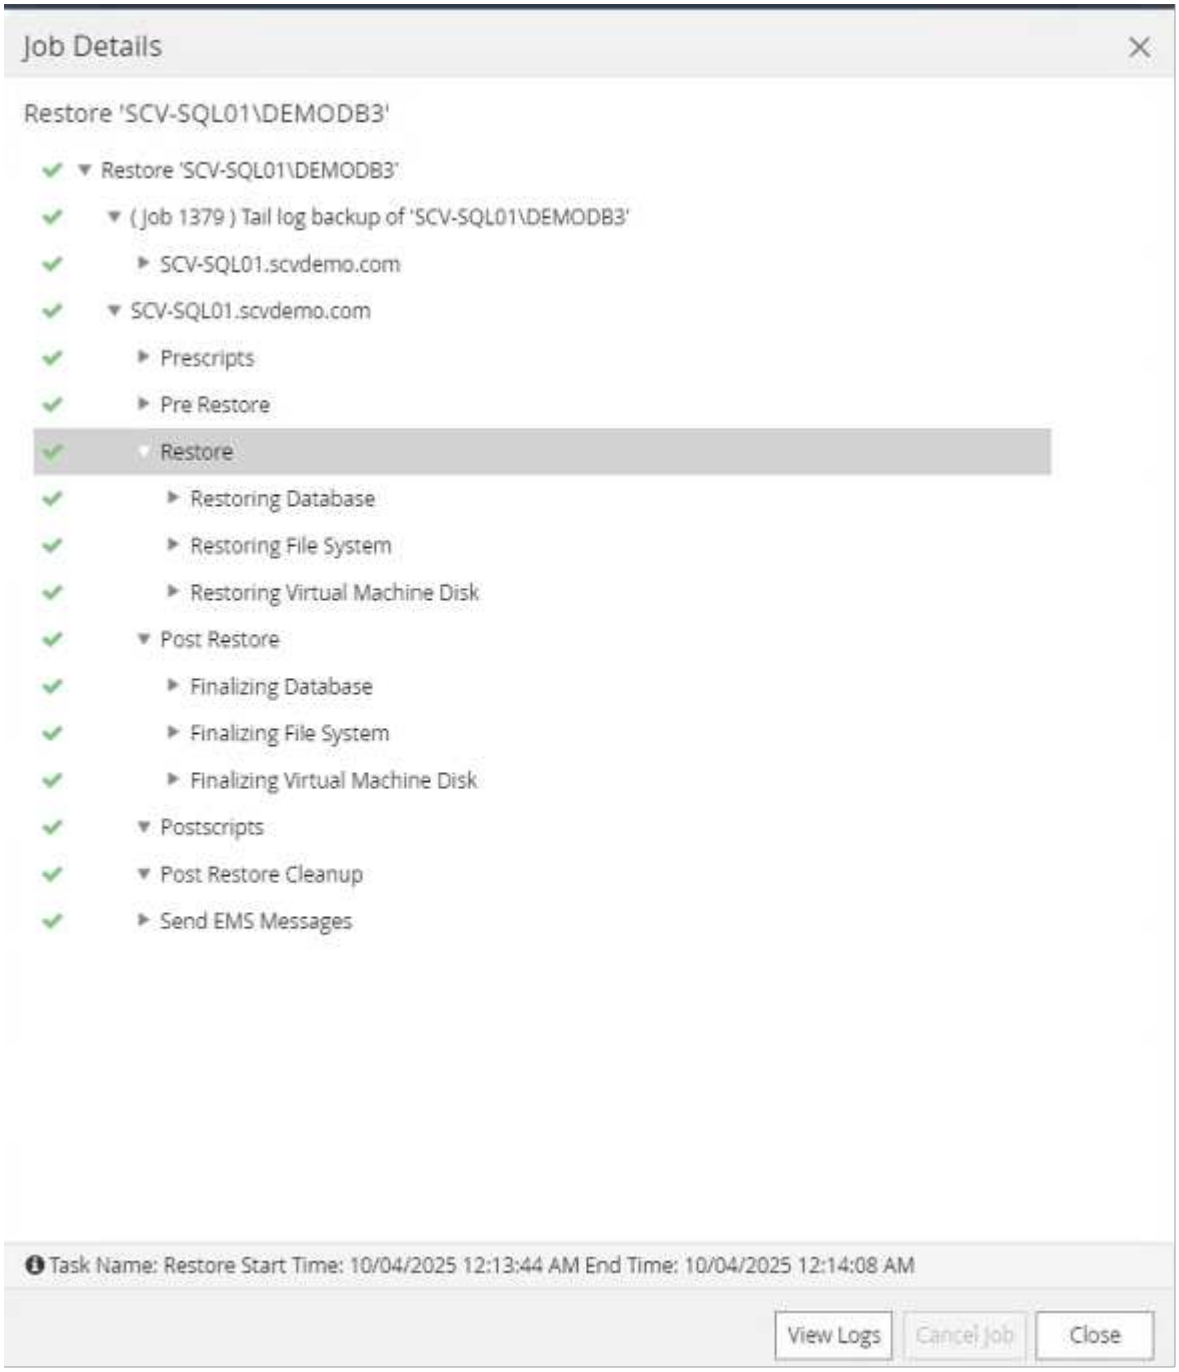
5. Select the post-restore state you wish the database to assume from the list options and select "Next".



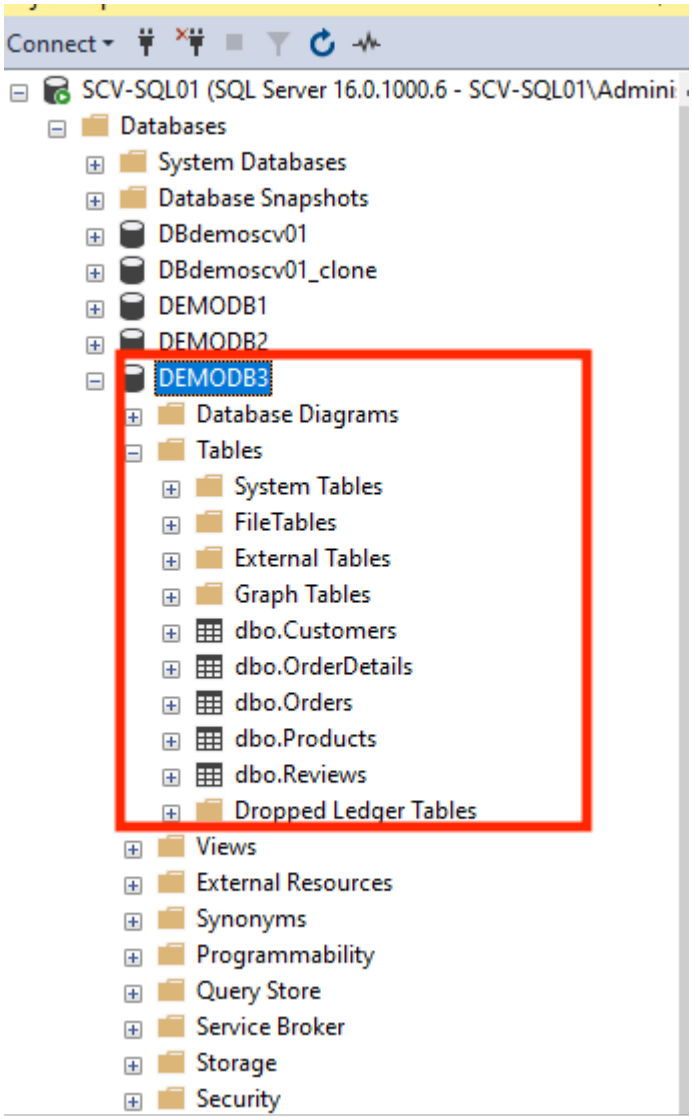
6. On the summary page, select "Finish".



7. Monitor the job to completion.



8. Verify that the deleted tables have been successfully restored.



Conclusion

In this document, we provided detailed step by step guidance of deploying, configuring, and using SCV and SnapCenter Plug-in for SQL Server (SCSQL) while using FSx for ONTAP as the supplemental datastore with Amazon EVS clusters.

Conclusion

To summarize, when using FSx for ONTAP to optimize TCO for an Amazon EVS deployment, SCV offers a simple and efficient method for backing up VMs. It can back up and restore VMs in a seamless and quick manner since it relies on Snapshot copies, which take literally seconds to complete.

Try the SnapCenter plug-in for VMware now for free.

NetApp for Azure / AVS

NetApp Capabilities for Azure AVS

Learn more about the capabilities that NetApp brings to the Azure VMware Solution

(AVS) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- Configuring AVS in Azure
- NetApp Storage Options for AVS
- NetApp / VMware Cloud Solutions

Configuring AVS in Azure

As with on-premises, planning a cloud based virtualization environment is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to Azure VMware Solution.

The setup process can be broken down into the following steps:

- Register the resource provider and create a private cloud
- Connect to a new or existing ExpressRoute virtual network gateway
- Validate the network connectivity and access the private cloud

View the detailed configuration steps for AVS.

NetApp Storage Options for AVS

NetApp storage can be utilized in several ways - either as guess connected or as a supplemental NFS datastore - within Azure AVS.

Please visit Supported NetApp Storage Options for more information.

Azure supports NetApp storage in the following configurations:

- Azure NetApp Files (ANF) as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- Azure NetApp Files (ANF) as a supplemental NFS datastore

View the detailed guest connect storage options for AVS.
View the detailed supplemental NFS datastore options for AVS.

Solution Use Cases

With NetApp and VMware cloud solutions, many use cases are simple to deploy in Azure AVS. se cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)

- Extend
- Migrate

Browse the NetApp solutions for Azure AVS

Protect workloads on Azure / AVS

Disaster Recovery with ANF and JetStream

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). Using the VMware VAIO framework, on-premises VMware workloads can be replicated to Azure Blob storage and recovered, enabling minimal or close to no data loss and near-zero RTO.

JetStream DR can be used to seamlessly recover the workloads replicated from on-premises to AVS and specifically to Azure NetApp Files. It enables cost-effective disaster recovery by using minimal resources at the DR site and cost-effective cloud storage. JetStream DR automates recovery to ANF datastores via Azure Blob Storage. JetStream DR recovers independent VMs or groups of related VMs into recovery site infrastructure according to network mapping and provides point-in-time recovery for ransomware protection.

This document provides an understanding of the JetStream DR principles of operations and its main components.

▼ Solution deployment overview

1. Install JetStream DR software in the on-premises data center.
    a. Download the JetStream DR software bundle from Azure Marketplace (ZIP) and deploy the JetStream DR MSA (OVA) in the designated cluster.
    b. Configure the cluster with the I/O filter package (install JetStream VIB).
    c. Provision Azure Blob (Azure Storage Account) in the same region as the DR AVS cluster.
    d. Deploy DRVA appliances and assign replication log volumes (VMDK from existing datastore or shared iSCSI storage).
    e. Create protected domains (groups of related VMs) and assign DRVAs and Azure Blob Storage/ANF.
    f. Start protection.
2. Install JetStream DR software in the Azure VMware Solution private cloud.
    a. Use the Run command to install and configure JetStream DR.
    b. Add the same Azure Blob container and discover domains using the Scan Domains option.
    c. Deploy required DRVA appliances.
    d. Create replication log volumes using available vSAN or ANF datastores.
    e. Import protected domains and configure RocVA (recovery VA) to use ANF datastore for VM placements.
    f. Select the appropriate failover option and start continuous rehydration for near-zero RTO domains or VMs.
3. During a disaster event, trigger failover to Azure NetApp Files datastores in the designated AVS DR site.
4. Invoke failback to the protected site after the protected site has been recovered.Before starting, make

sure that the prerequisites are met as indicated in this link and also run the Bandwidth Testing Tool (BWT) provided by JetStream Software to evaluate the potential performance of Azure Blob storage and its replication bandwidth when used with JetStream DR software. After the pre-requisites, including connectivity, are in place, set up and subscribe to JetStream DR for AVS from the Azure Marketplace. After the software bundle is downloaded, proceed with the installation process described above.

When planning and starting protection for a large number of VMs (for example, 100+), use the Capacity Planning Tool (CPT) from the JetStream DR Automation Toolkit. Provide a list of VMs to be protected together with their RTO and recovery group preferences, and then run CPT.

CPT performs the following functions:

- Combining VMs into protection domains according to their RTO.
- Defining the optimal number of DRVAs and their resources.
- Estimating required replication bandwidth.
- Identifying replication log volume characteristics (capacity, bandwidth, and so on).
- Estimating required object storage capacity, and more.

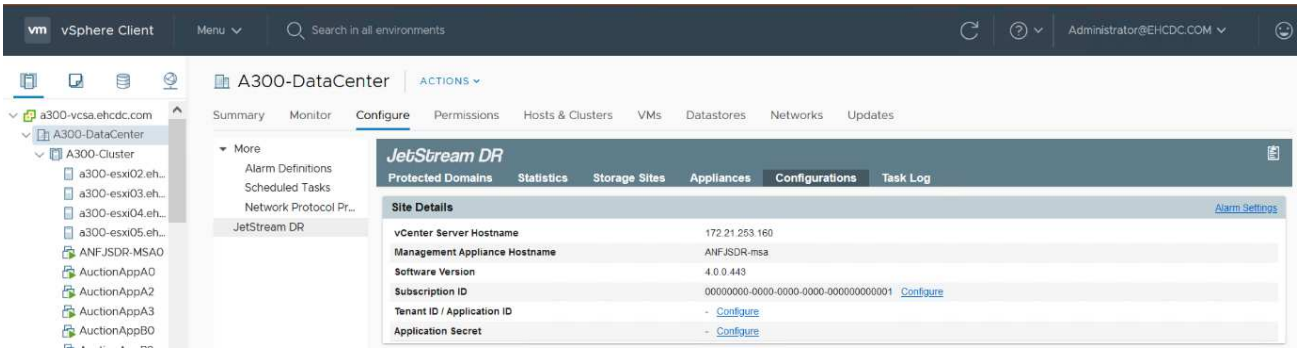
The number and content of domains prescribed depend upon various VM characteristics such as average IOPS, total capacity, priority (which defines failover order), RTO, and others.

Install JetStream DR in On-Premises Datacenter

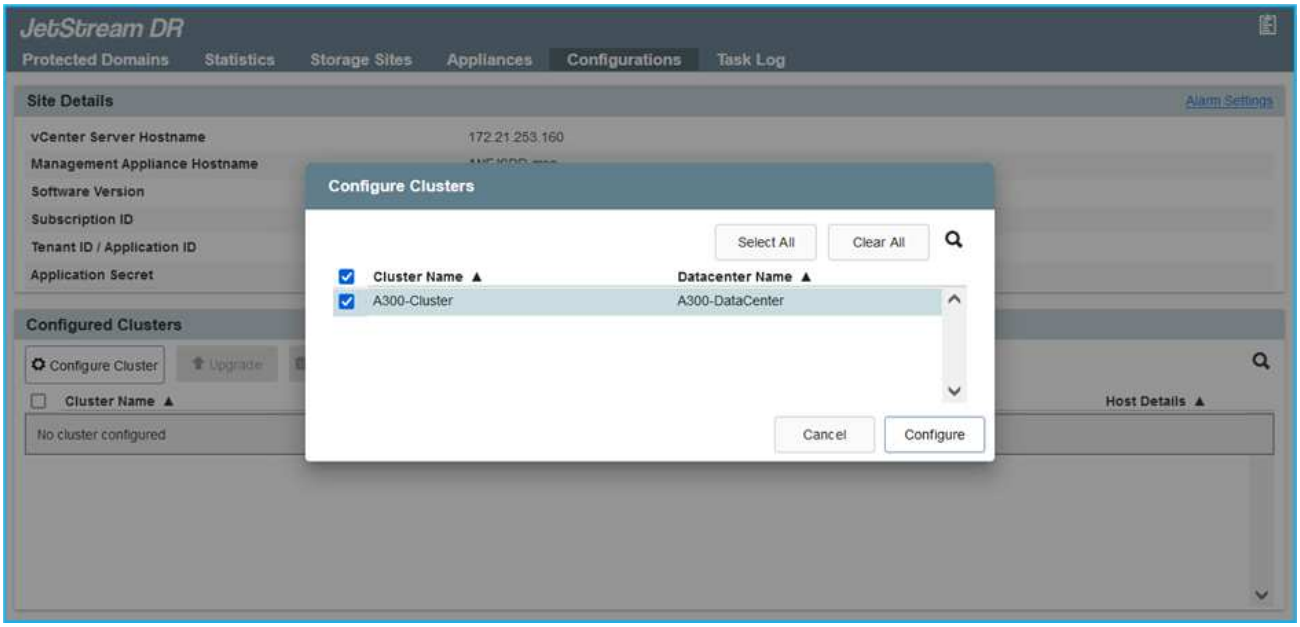
JetStream DR software consists of three major components: JetStream DR Management Server Virtual Appliance (MSA), DR Virtual Appliance (DRVA), and host components (I/O Filter packages). MSA is used to install and configure host components on the compute cluster and then to administer JetStream DR software. The following list provides a high-level description of the installation process:

▼ How to install JetStream DR for on-premises

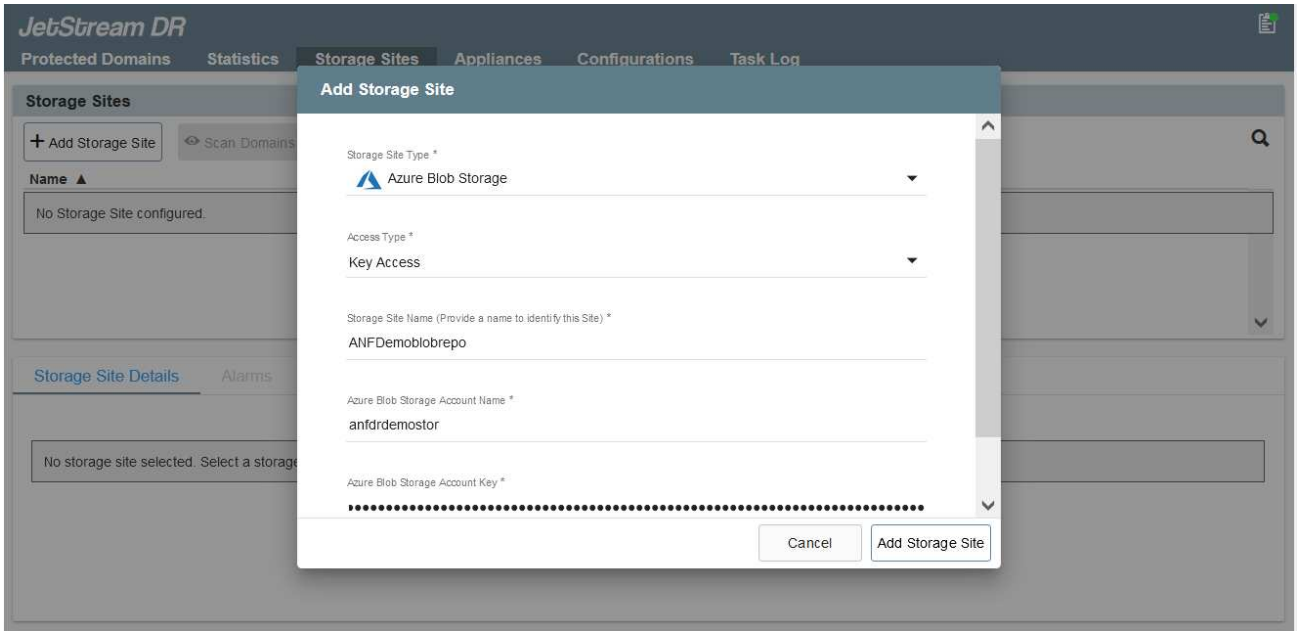
1. Check prerequisites.
2. Run the Capacity Planning Tool for resource and configuration recommendations (optional but recommended for proof-of-concept trials).
3. Deploy the JetStream DR MSA to a vSphere host in the designated cluster.
4. Launch the MSA using its DNS name in a browser.
5. Register the vCenter server with the MSA.To perform the installation, complete the following detailed steps:
6. After JetStream DR MSA has been deployed and the vCenter Server has been registered, access the JetStream DR plug-in using the vSphere Web Client. This can be done by navigating to Datacenter > Configure > JetStream DR.



7. From the JetStream DR interface, select the appropriate cluster.



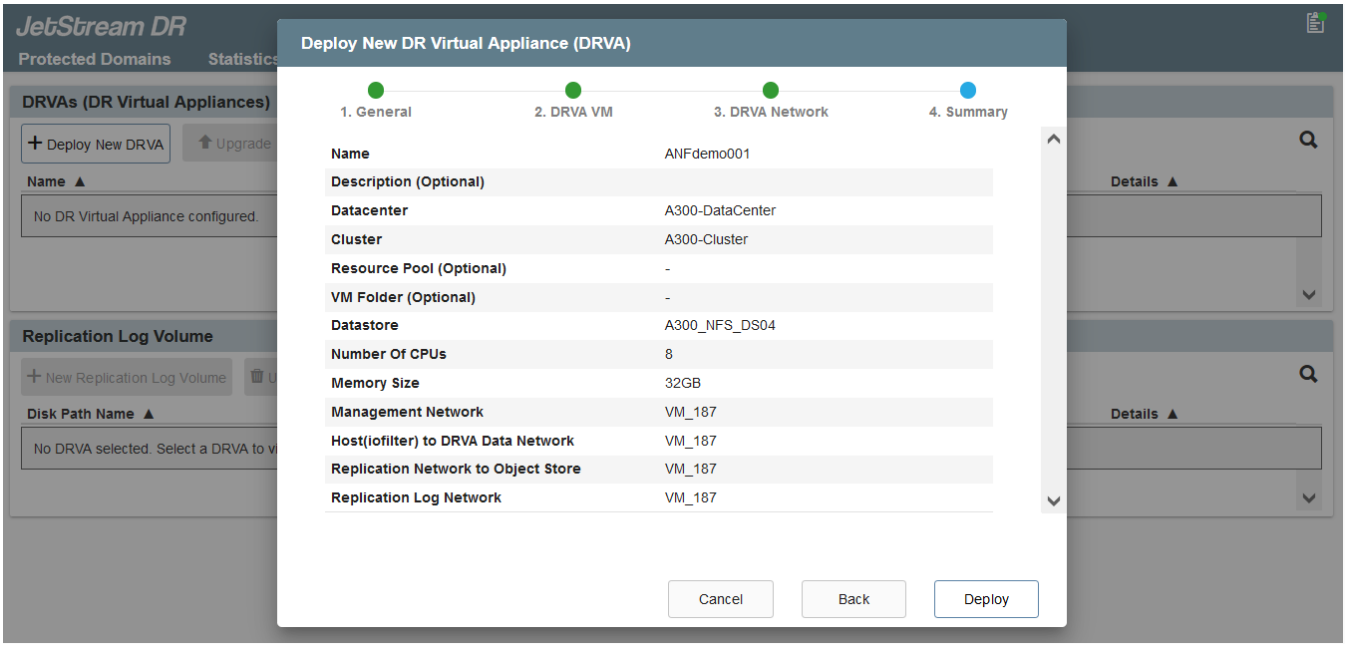
8. Configure the cluster with the I/O filter package.



9. Add Azure Blob Storage located at the recovery site.
10. Deploy a DR Virtual Appliance (DRVA) from the Appliances tab.

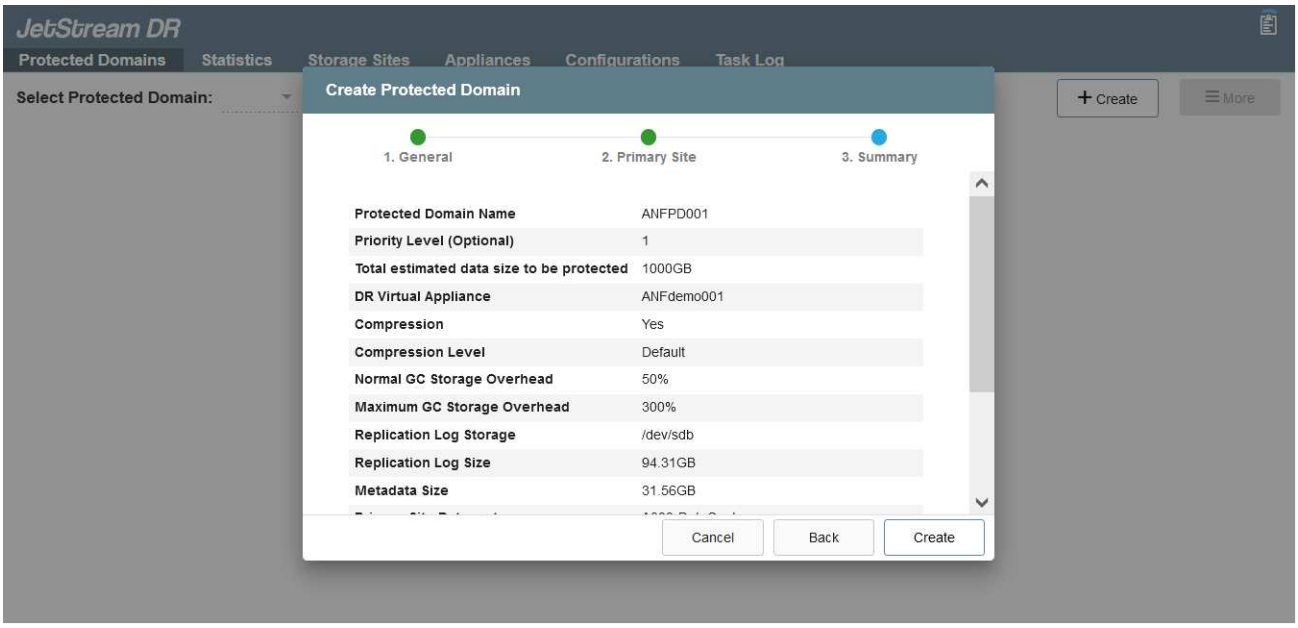
> DRVAs can be automatically created by CPT, but for POC trials we recommend configuring and running the DR cycle manually (start protection > failover > failback).

The JetStream DRVA is a virtual appliance that facilitates key functions in the data replication process. A protected cluster must contain at least one DRVA, and typically one DRVA is configured per host. Each DRVA can manage multiple protected domains.



In this example, four DRVA's were created for 80 virtual machines.

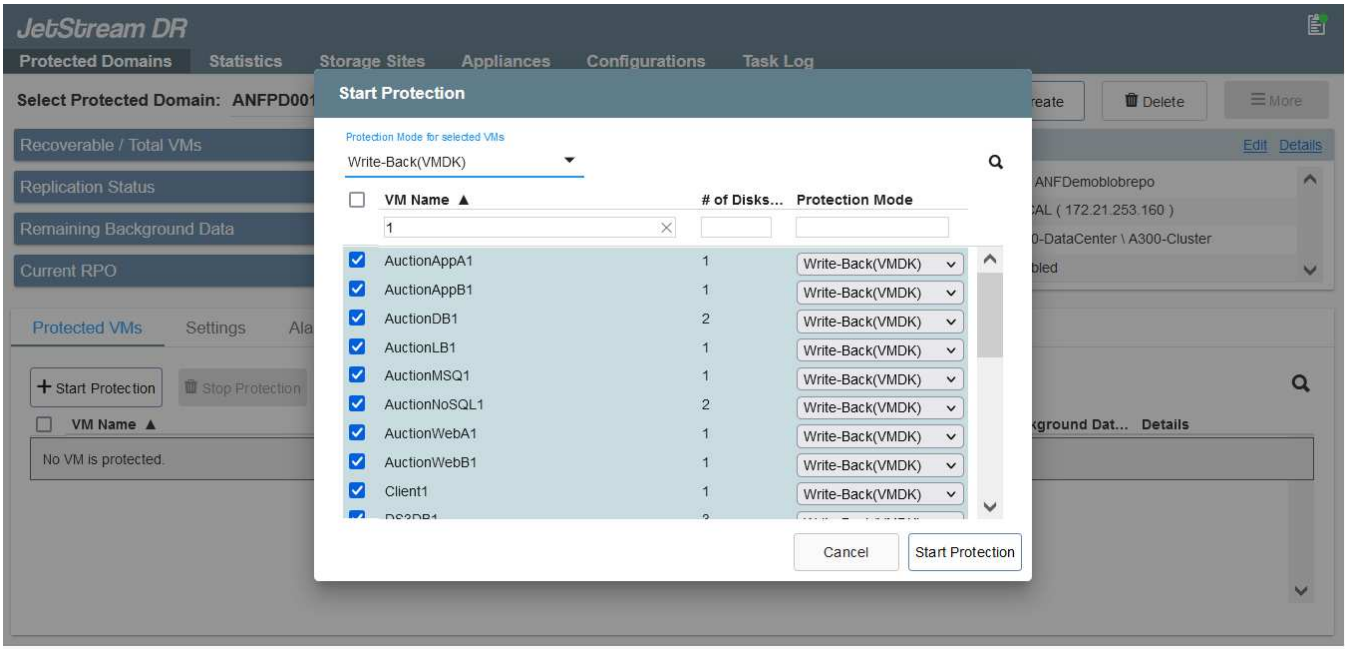
1. Create replication log volumes for each DRVA using VMDK from the datastores available or independent shared iSCSI storage pools.
2. From the Protected Domains tab, create the required number of protected domains using information about the Azure Blob Storage site, DRVA instance, and replication log. A protected domain defines a specific VM or set of VMs within the cluster that are protected together and assigned a priority order for failover/failback operations.



3. Select VMs you want to protect and start VM protection of the protected domain. This begins data replication to the designated Blob Store.

Verify that the same protection mode is used for all VMs in a protected domain.

Write- Back(VMDK) mode can offer higher performance.



Verify that replication log volumes are placed on high performance storage.

Failover run books can be configured to group the VMs (called Recovery Group), set boot order sequence, and modify the CPU/memory settings along with IP configurations.

Install JetStream DR for AVS in an Azure VMware Solution private cloud using the Run command

A best practice for a recovery site (AVS) is to create a three-node pilot-light cluster in advance. This allows the recovery site infrastructure to be preconfigured, including the following items:

- Destination networking segments, firewalls, services like DHCP and DNS, and so on.
- Installation of JetStream DR for AVS
- Configuration of ANF volumes as datastores, and moreJetStream DR supports near-zero RTO mode for mission- critical domains. For these domains, destination storage should be preinstalled. ANF is a recommended storage type in this case.

Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.

Depending on the SLA and RTO requirements, continuous failover or regular (standard) failover mode can be used. For near-zero RTO, continuous rehydration should be started at the recovery site.

▼ How to install JetStream DR for AVS in a private cloud

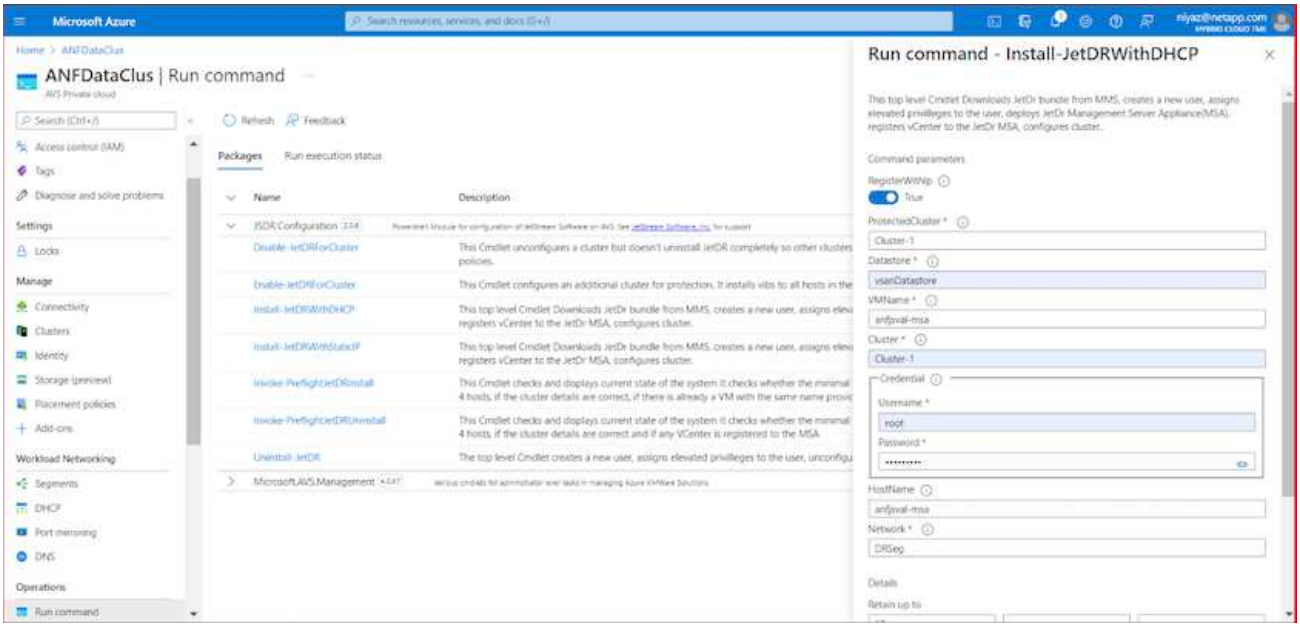
To install JetStream DR for AVS on an Azure VMware Solution private cloud, complete the following steps:

1. From the Azure portal, go to the Azure VMware solution, select the private cloud, and select Run command > Packages > JSDR.Configuration.

The default CloudAdmin user in Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. Azure VMware Solution enables simplified and

automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

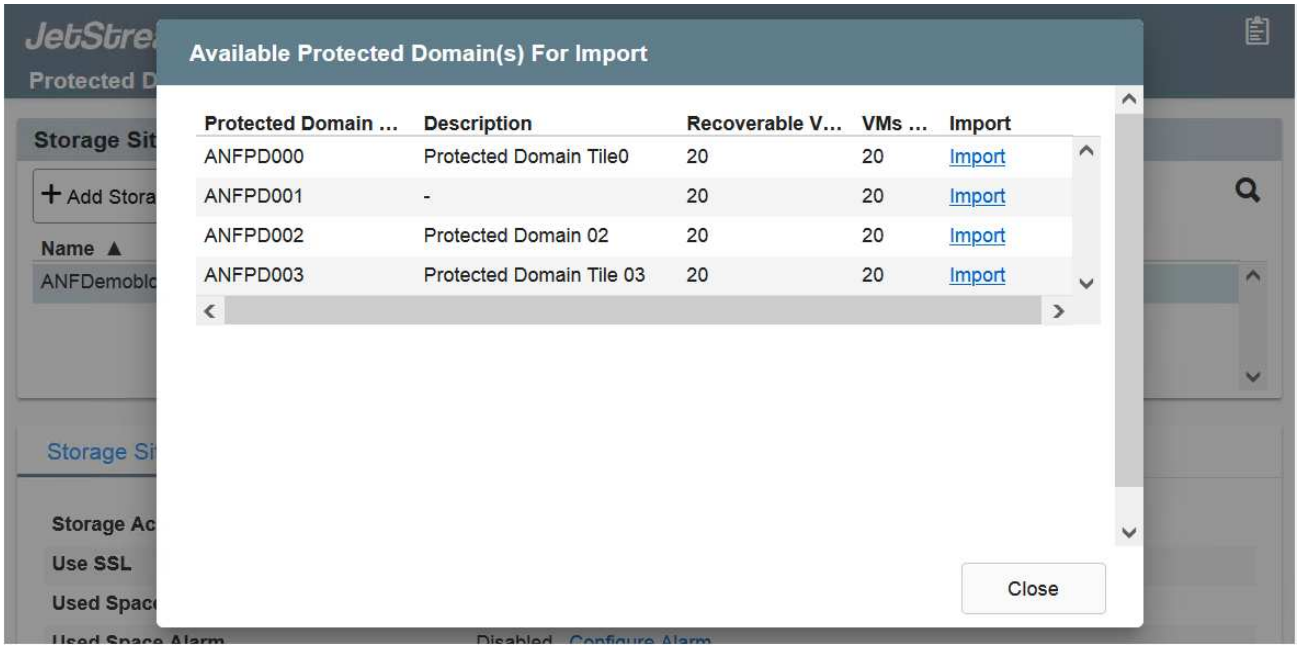
The following screenshot shows installation using a DHCP-based IP address.



2. After JetStream DR for AVS installation is complete, refresh the browser. To access the JetStream DR UI, go to SDDC Datacenter > Configure > JetStream DR.



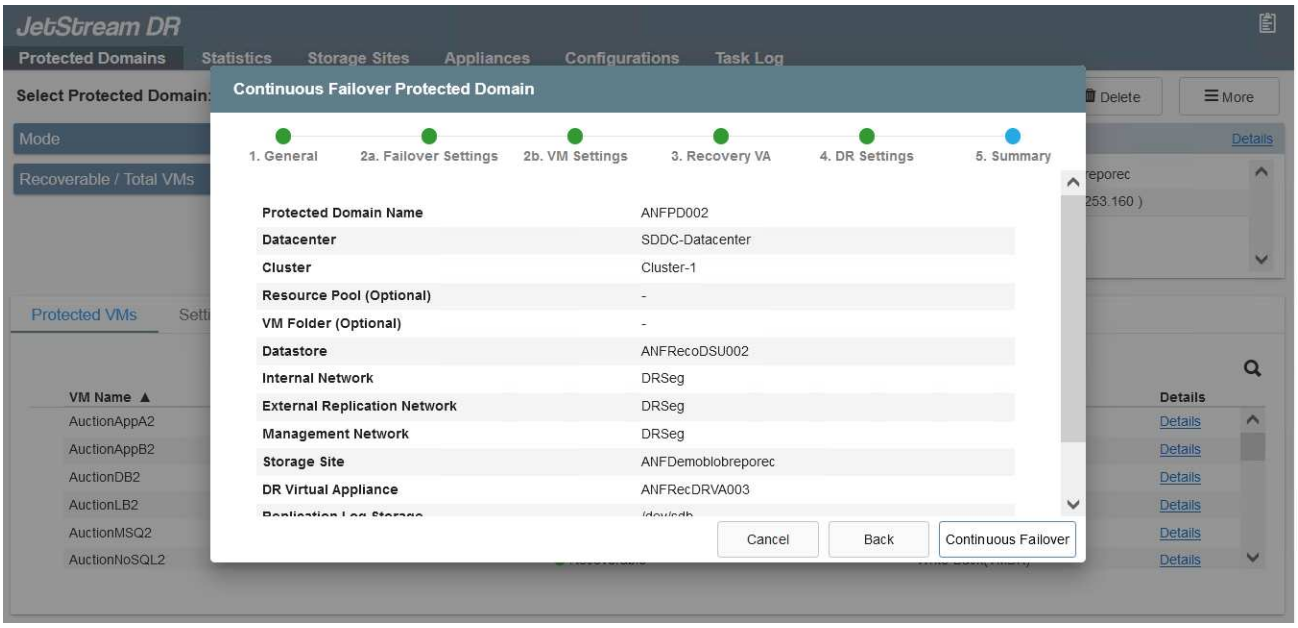
3. From the JetStream DR interface, add the Azure Blob Storage account that was used to protect the on-premises cluster as a storage site and then run the Scan Domains option.



4. After the protected domains are imported, deploy DRVA appliances. In this example, continuous rehydration is started manually from the recovery site using the JetStream DR UI.

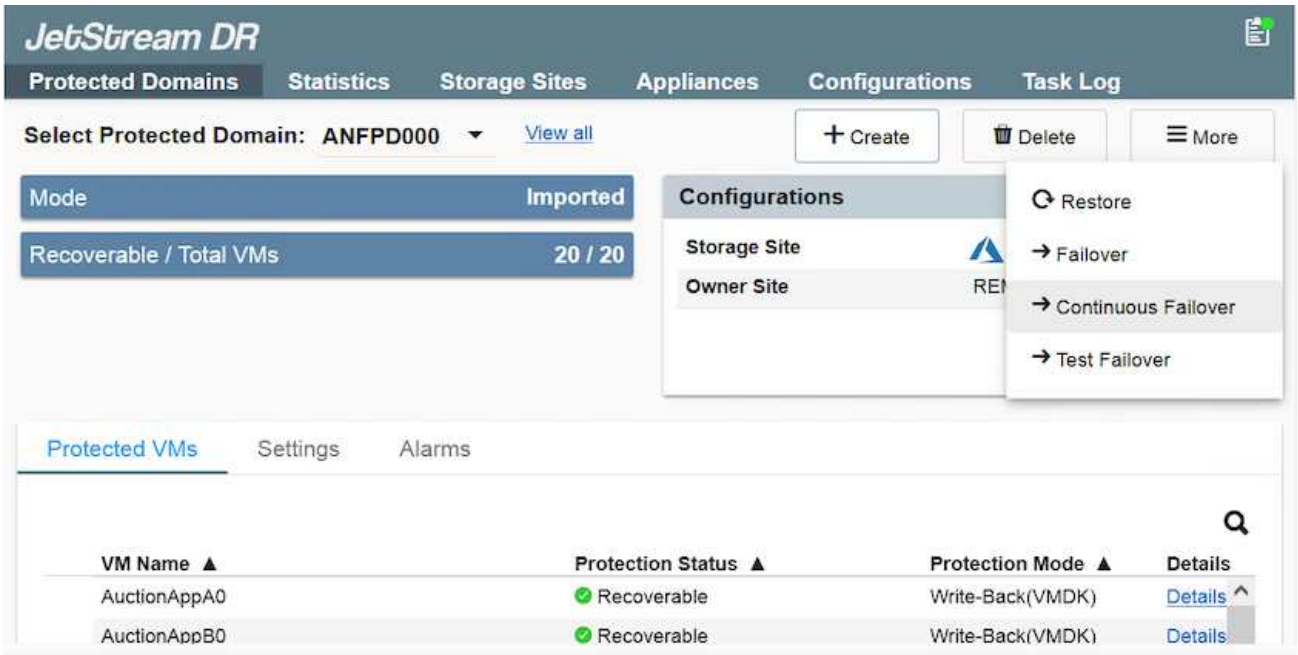
> These steps can also be automated using CPT created plans.

5. Create replication log volumes using available vSAN or ANF datastores.
6. Import the protected domains and configure the Recovery VA to use the ANF datastore for VM placements.



> Make sure that DHCP is enabled on the selected segment and enough IPs are available. Dynamic IPs are temporarily used while domains are recovering. Each recovering VM (including continuous rehydration) requires an individual dynamic IP. After recovery is complete, the IP is released and can be reused.

7. Select the appropriate failover option (continuous failover or failover). In this example, continuous rehydration (continuous failover) is selected.



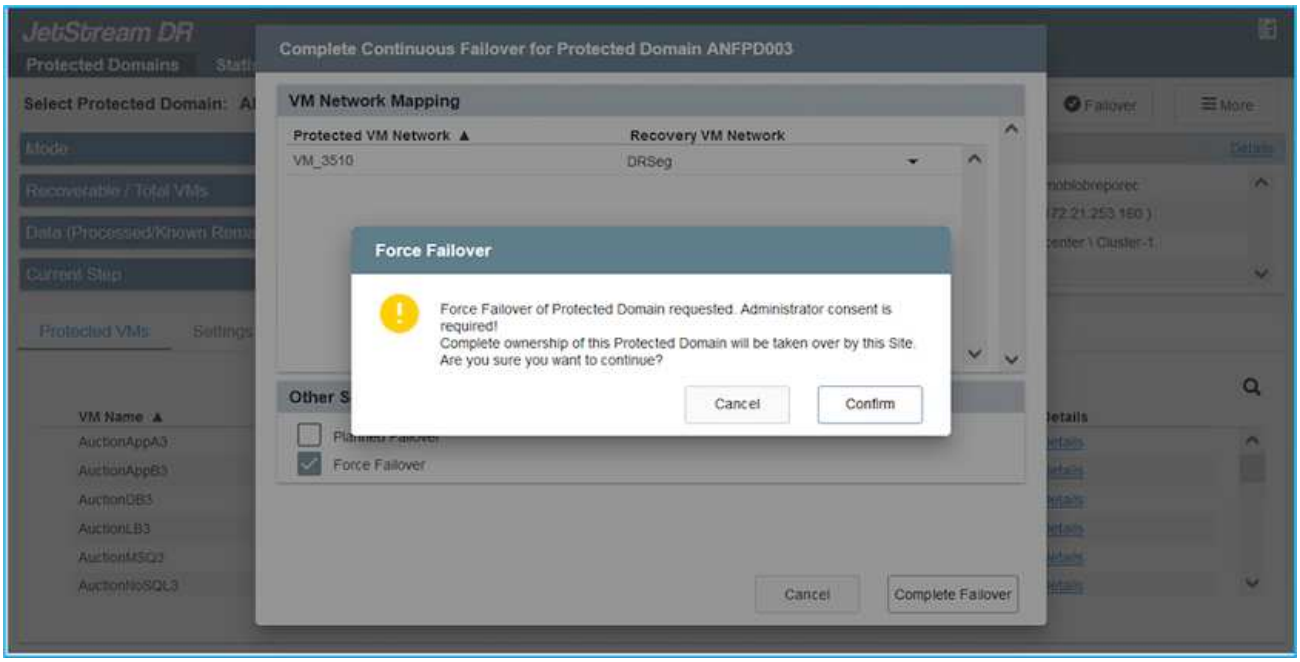
Performing Failover / Failback

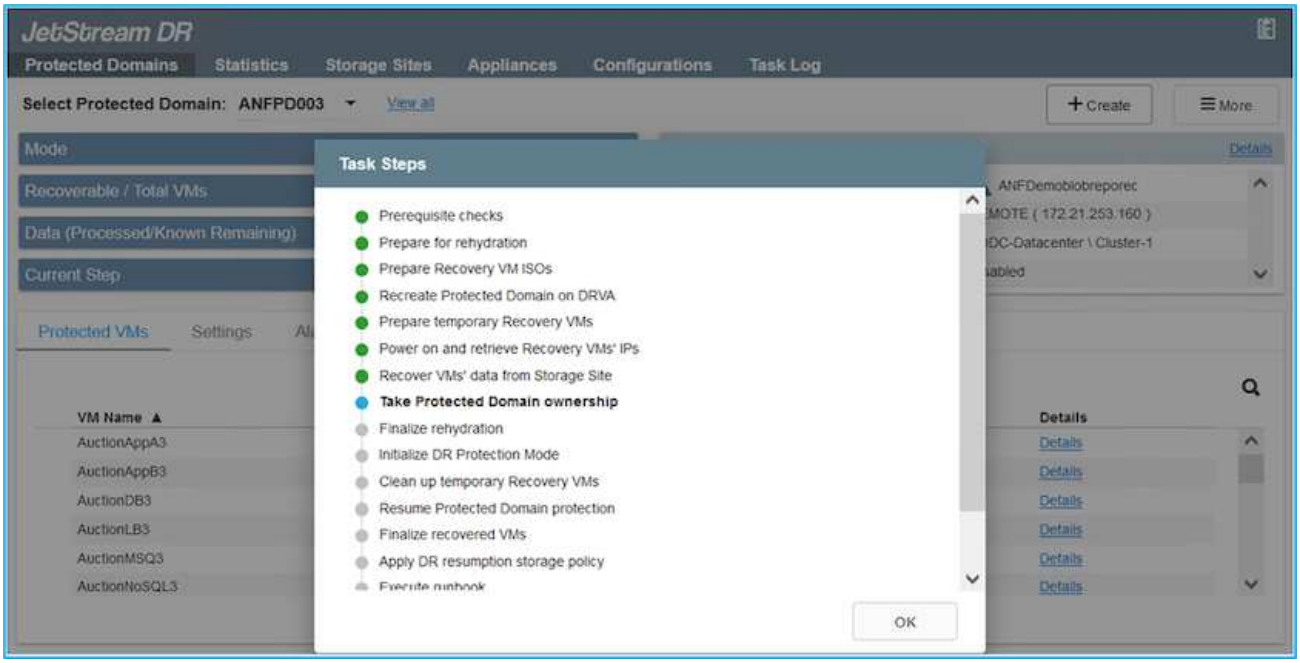
▼ How to perform a Failover / Failback

1. After a disaster occurs in the protected cluster of the on-premises environment (partial or full failure), trigger the failover.

CPT can be used to execute the failover plan to recover the VMs from Azure Blob Storage into the AVS cluster recovery site.

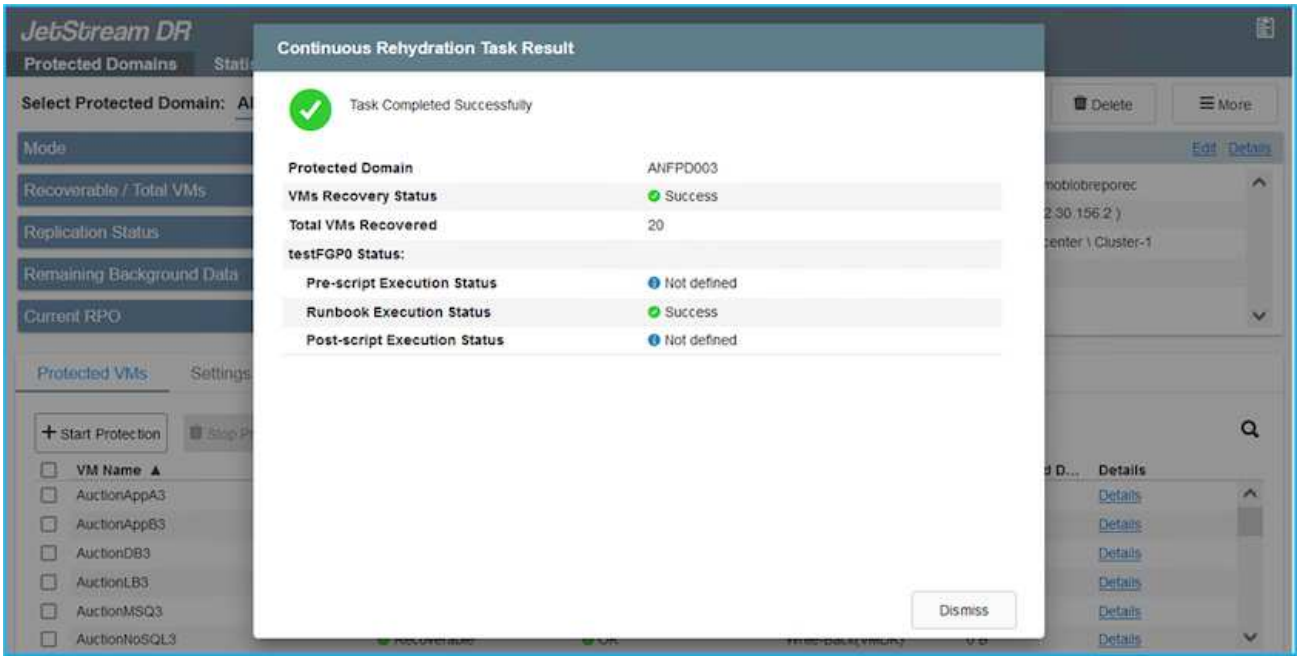
After failover (for continuous or standard rehydration) when the protected VMs have been started in AVS, protection is automatically resumed and JetStream DR continues to replicate their data into the appropriate/original containers in Azure Blob Storage.





The task bar shows progress of failover activities.

2. When the task is complete, access the recovered VMs and business continues as normal.

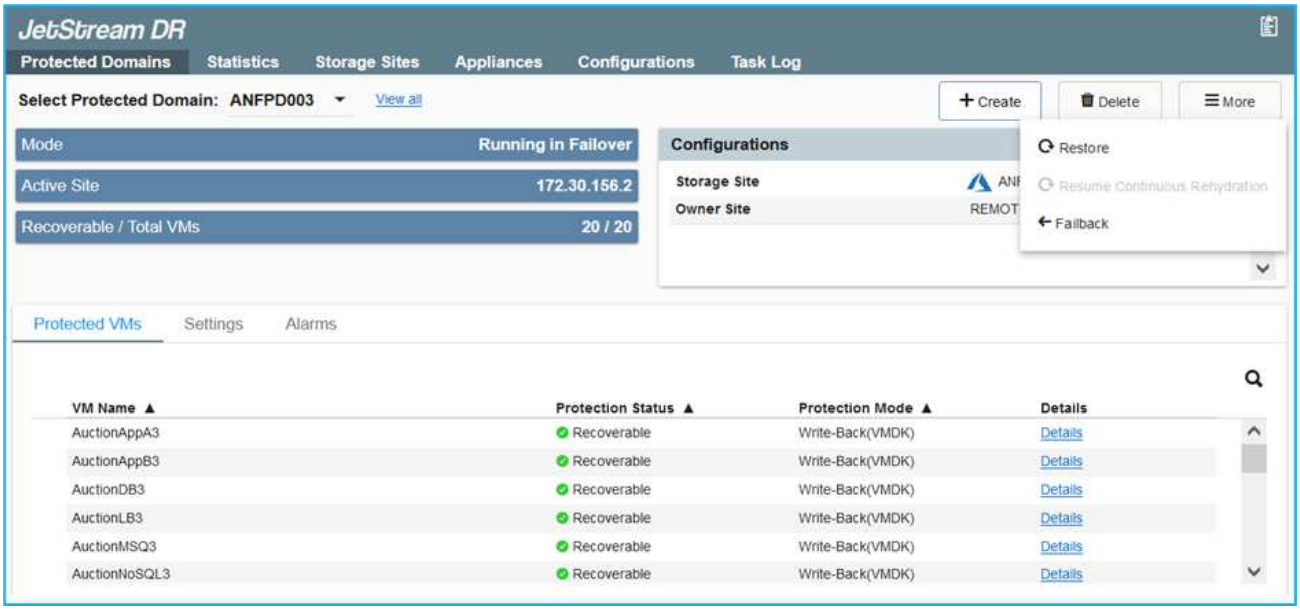


After the primary site is up and running again, failback can be performed. VM protection is resumed and data consistency should be checked.

3. Restore the on-premises environment. Depending upon the type of disaster incident, it might be necessary to restore and/or verify the configuration of the protected cluster. If necessary, JetStream DR software might need to be reinstalled.

> Note: The `recovery_utility_prepare_failback` script provided in the Automation Toolkit can be used to help clean the original protected site of any obsolete VMs, domain information, and so on.

4. Access the restored on-premises environment, go to the Jetstream DR UI, and select the appropriate protected domain. After the protected site is ready for failback, select the Failback option in the UI.



> The CPT generated failback plan can also be used to initiate the return of the VMs and their

data from the object store back to the original VMware environment.

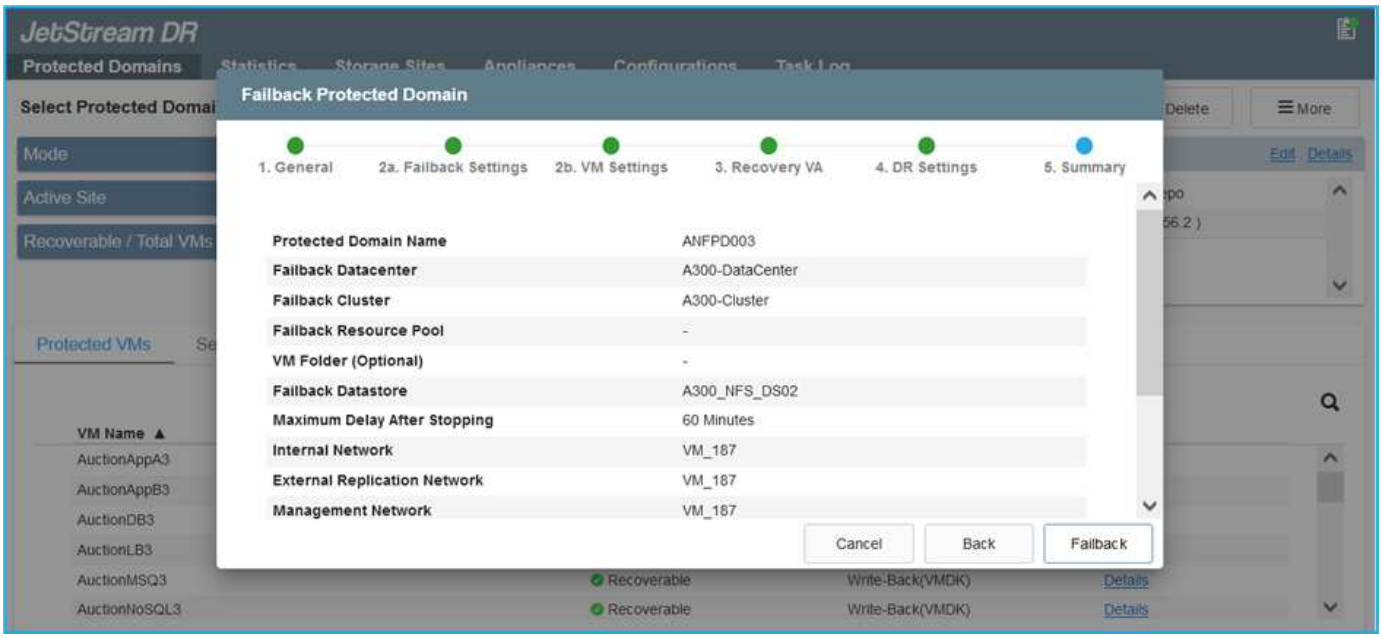
Specify the maximum delay after pausing VMs in the recovery site and restarting in the protected site. This time includes completing replication after stopping failover VMs, the time to clean recovery site, and the time to recreate VMs in protected site. The NetApp recommended value is 10 minutes.

Complete the failback process, and then confirm the resumption of VM protection and data consistency.

Ransomeware Recovery

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to determine the safe point of return and, once determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (from sleeping malware or through vulnerable applications).

JetStream DR for AVS together with Azure NetApp Files datastores can address these concerns by allowing organizations to recover from available points in time, so that workloads are recovered to a functional, isolated network if required. Recovery allows applications to function and communicate with each other while not exposing them to north- south traffic, thereby giving security teams a safe place to perform forensics and other necessary remediation.



Disaster Recovery with CVO and AVS (guest-connected storage)

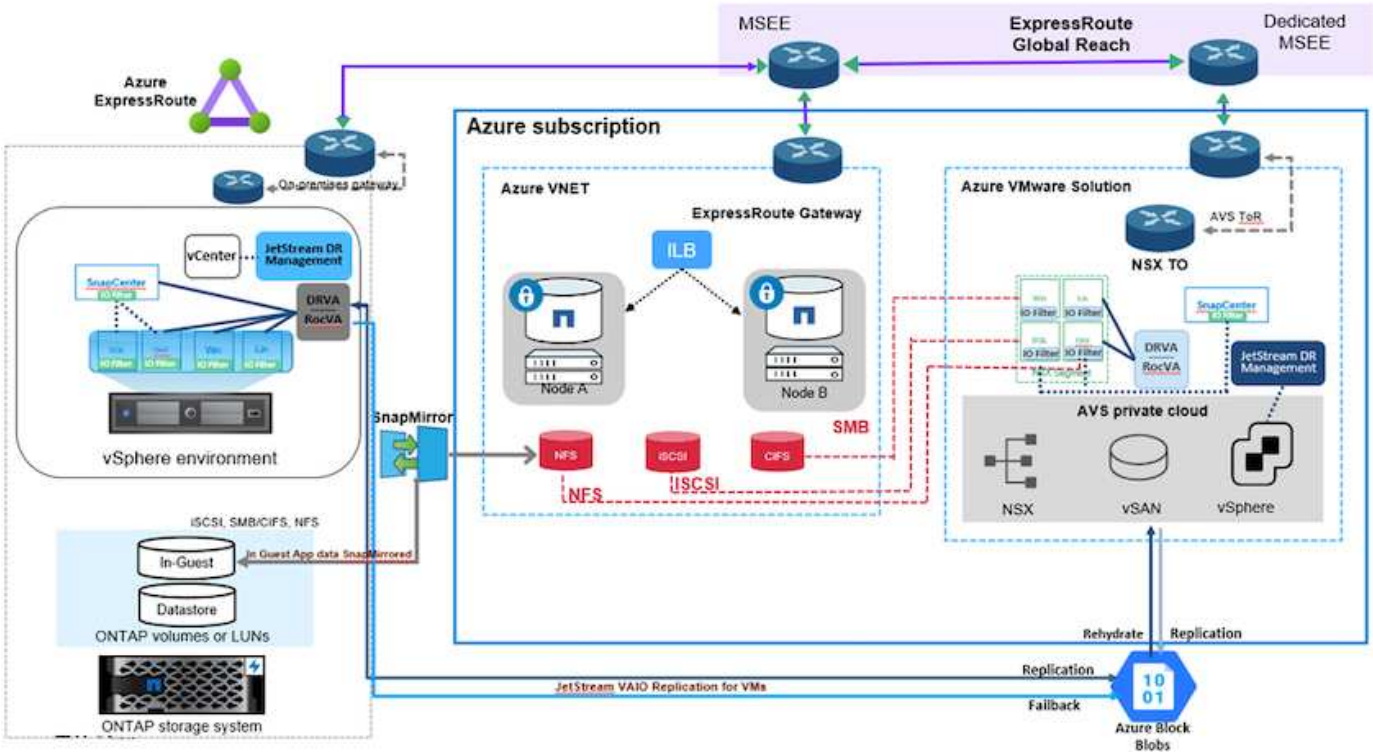
Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Azure.

Overview

```
This covers application data; however, what about the actual VMs
themselves. Disaster recovery should cover all dependent components,
```

including virtual machines, VMDKs, application data, and more. To accomplish this, SnapMirror along with Jetstream can be used to seamlessly recover workloads replicated from on-premises to Cloud Volumes ONTAP while using vSAN storage for VM VMDKs.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, JetStream, and the Azure VMware Solution (AVS).



Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we assume that the on-premises environment is using SnapCenter for application-consistent backups.

> This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Azure virtual network, use the express route global reach or a virtual WAN with a VPN gateway. Segments should be created based on the on-premises vLAN design.

> There are multiple options for connecting on-premises datacenters to Azure, which prevents us from outlining a specific workflow in this document. Refer to the Azure documentation for the appropriate on-premises-to-Azure connectivity method.

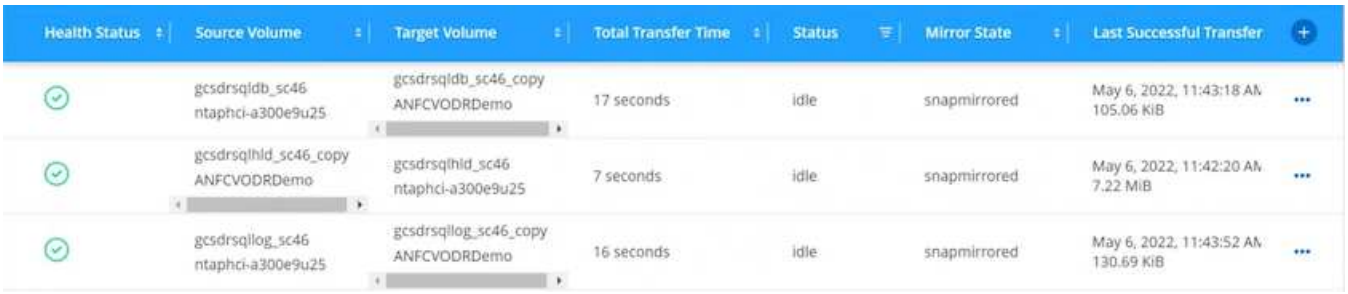
Solution Deployment Overview

1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
2. Provision Cloud Volumes ONTAP with the correct instance size using Cloud manager within the appropriate subscription and virtual network.
   a. Configure SnapMirror for the relevant application volumes.
   b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
3. Install the JetStream DR software in the on-premises data center and start protection for virtual machines.
4. Install JetStream DR software in the Azure VMware Solution private cloud.
5. During a disaster event, break the SnapMirror relationship using Cloud Manager and trigger failover of virtual machines to Azure NetApp Files or to vSAN datastores in the designated AVS DR site.
   a. Reconnect the ISCSI LUNs and NFS mounts for the application VMs.
6. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

Deployment Details

▼ Configure CVO on Azure and replicate volumes to CVO

The first step is to configure Cloud Volumes ONTAP on Azure (Link) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.



Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
✓	gcsdrsqldb_sc46 ntaphci-a300e9u25	gcsdrsqldb_sc46_copy ANFCVODRDemo	17 seconds	idle	snapmirrored	May 6, 2022, 11:43:18 AM 105.06 KiB
✓	gcsdrsqlhld_sc46_copy ANFCVODRDemo	gcsdrsqlhld_sc46 ntaphci-a300e9u25	7 seconds	idle	snapmirrored	May 6, 2022, 11:42:20 AM 7.22 MiB
✓	gcsdrsqllog_sc46 ntaphci-a300e9u25	gcsdrsqllog_sc46_copy ANFCVODRDemo	16 seconds	idle	snapmirrored	May 6, 2022, 11:43:52 AM 130.69 KiB

▼ Configure AVS hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the Azure VMware solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

The decision to deploy an AVS cluster is primarily based on the RPO/RTO requirements. With the Azure VMware solution, the SDDC can be provisioned just in time in preparation for either testing or an actual disaster event. An SDDC deployed just in time saves on ESXi host costs when you are not dealing with a disaster. However, this form of deployment affects the RTO by a few of hours while SDDC is being provisioned.

The most common deployed option is to have SDDC running in an always-on, pilot-light mode of operation. This option provides a small footprint of three hosts that are always available, and it also speeds up recovery operations by providing a running baseline for simulation activities and compliance checks, thus avoiding the risk of operational drift between the production and DR sites. The pilot-light cluster can be scaled up quickly to the desired level when needed to handle an actual DR event.

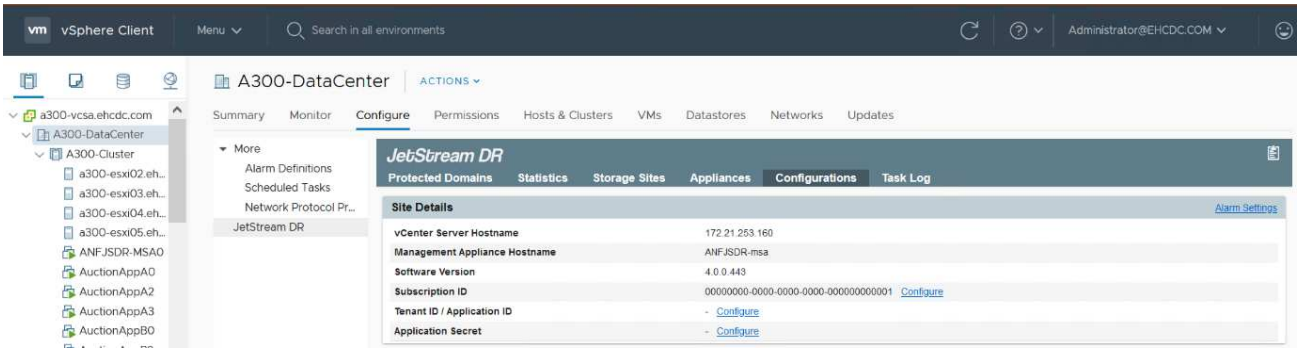
To configure AVS SDDC (be it on-demand or in pilot-light mode), see Deploy and configure the Virtualization Environment on Azure. As a prerequisite, verify that the guest VMs residing on the AVS hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and AVS have been configured properly, begin configuring Jetstream to automate the recovery of on-premises workloads to AVS (VMs with application VMDKs and VMs with in-guest storage) by using the VAIO mechanism and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

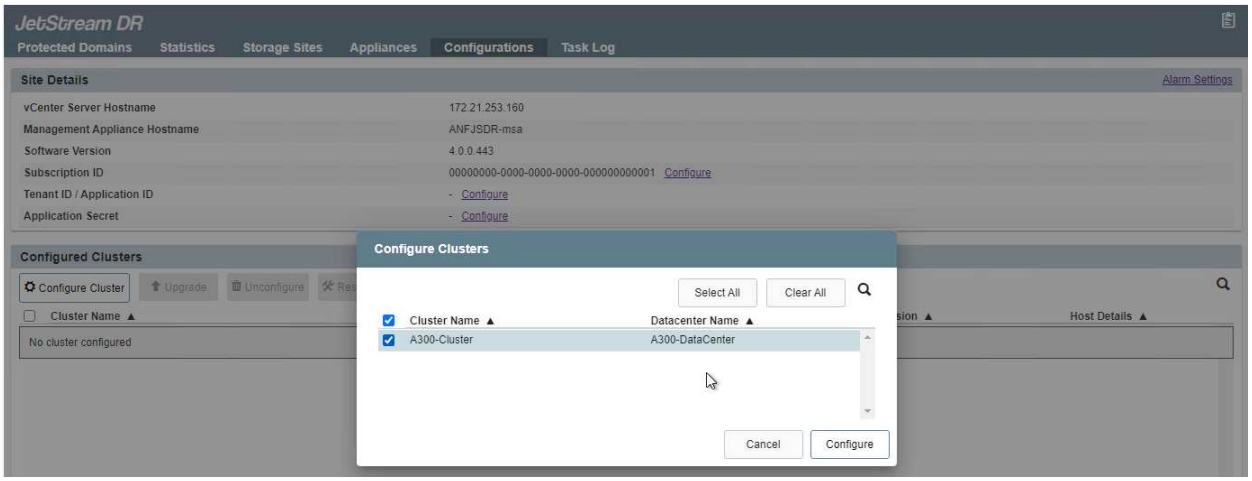
▼ Install JetStream DR in on-premises datacenter

JetStream DR software consists of three major components: the JetStream DR Management Server Virtual Appliance (MSA), the DR Virtual Appliance (DRVA), and host components (I/O filter packages). The MSA is used to install and configure host components on the compute cluster and then to administer JetStream DR software. The installation process is as follows:

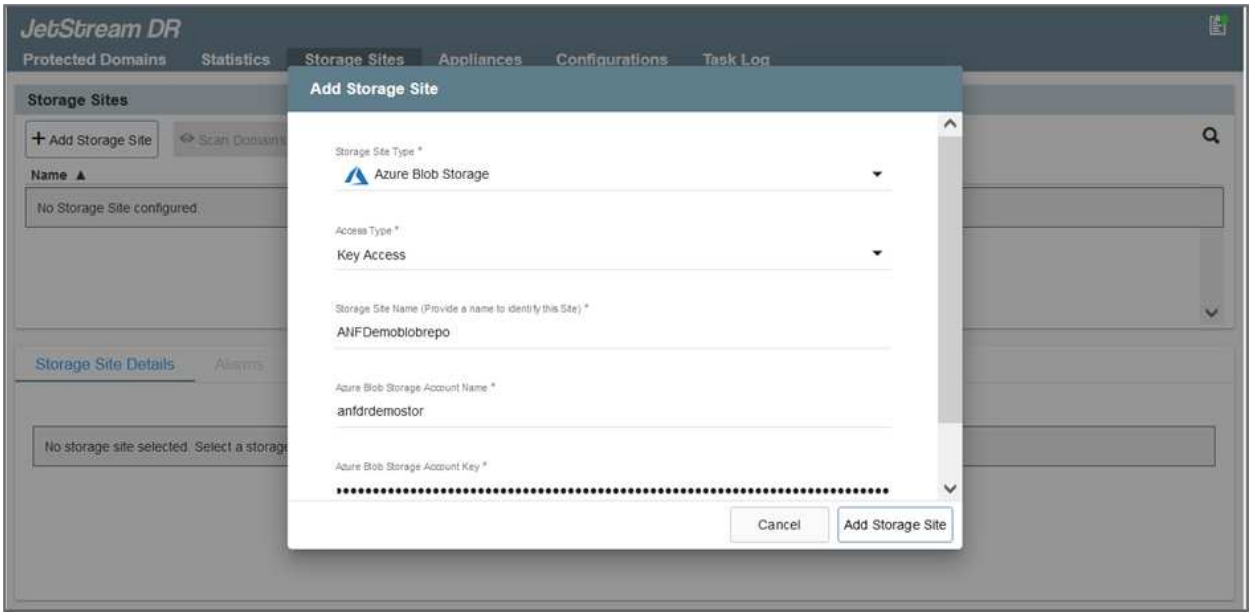
1. Check the prerequisites.
2. Run the Capacity Planning Tool for resource and configuration recommendations.
3. Deploy the JetStream DR MSA to each vSphere host in the designated cluster.
4. Launch the MSA using its DNS name in a browser.
5. Register the vCenter server with the MSA.
6. After JetStream DR MSA has been deployed and the vCenter Server has been registered, navigate to the JetStream DR plug-in with the vSphere Web Client. This can be done by navigating to Datacenter > Configure > JetStream DR.



7. From the JetStream DR interface, complete the following tasks:
   a. Configure the cluster with the I/O filter package.



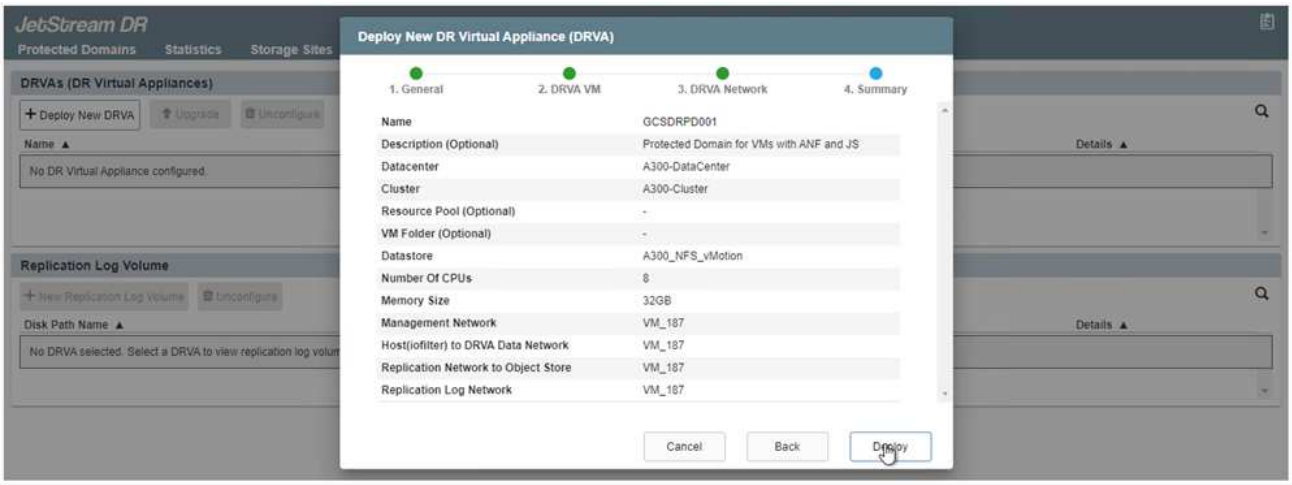
b. Add the Azure Blob storage located at the recovery site.



8. Deploy the required number of DR Virtual Appliances (DRVAs) from the Appliances tab.

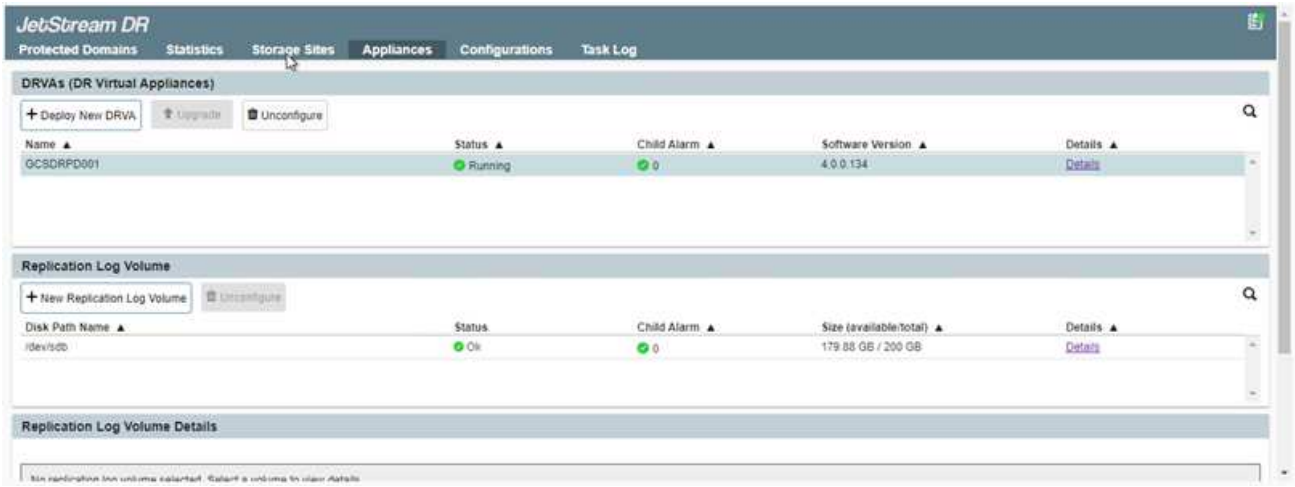
Use the capacity planning tool to estimate the number of DRVAs required.



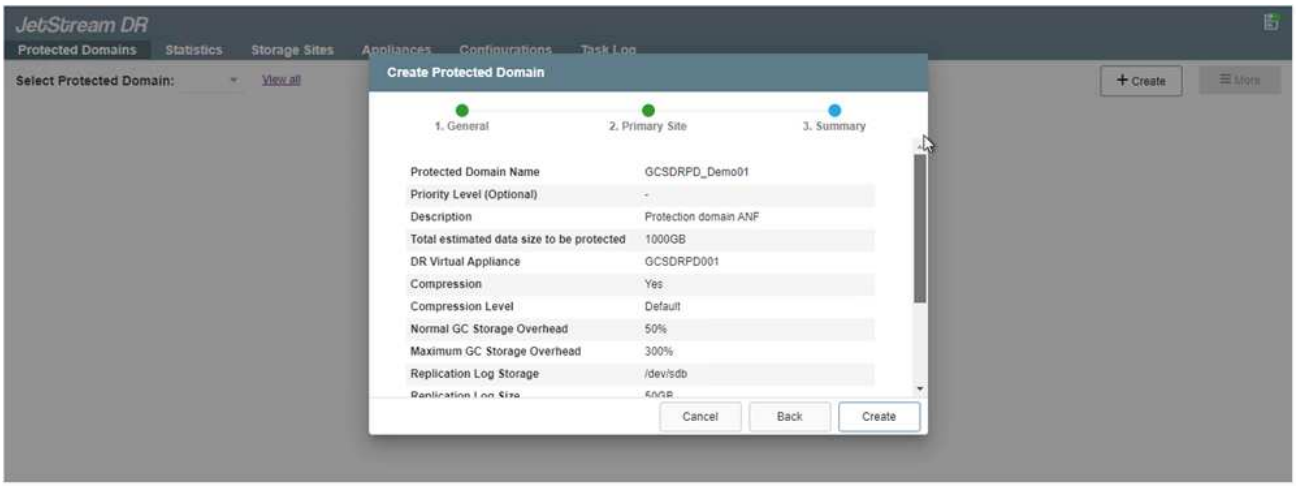


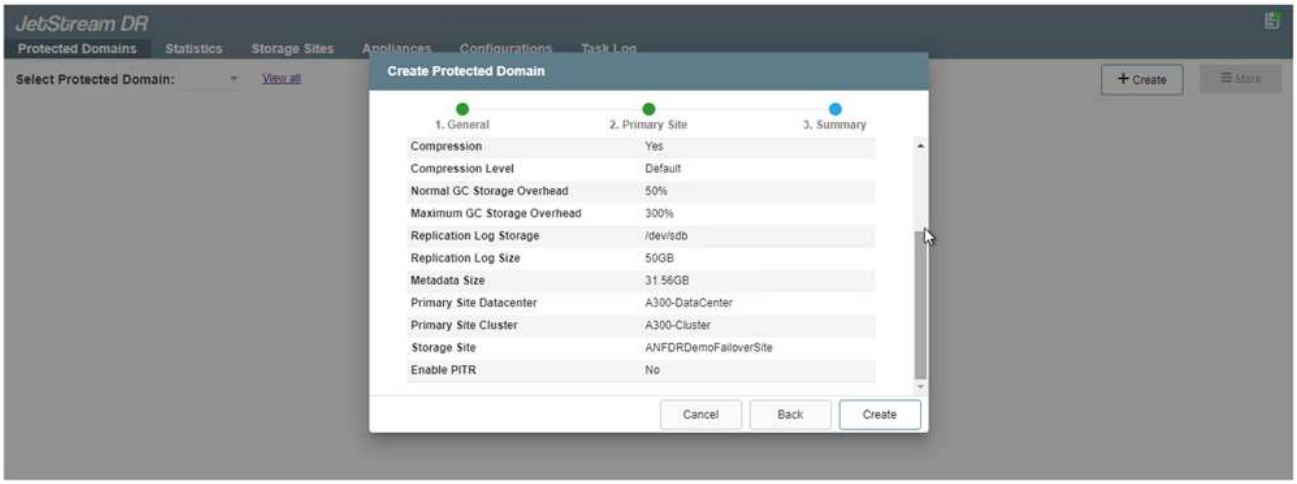
9. Create replication log volumes for each DRVA using the VMDK from the datastores available or the

independent shared iSCSI storage pool.



10. From the Protected Domains tab, create the required number of protected domains using information about the Azure Blob Storage site, the DRVA instance, and the replication log. A protected domain defines a specific VM or set of application VMs within the cluster that are protected together and assigned a priority order for failover/failback operations.



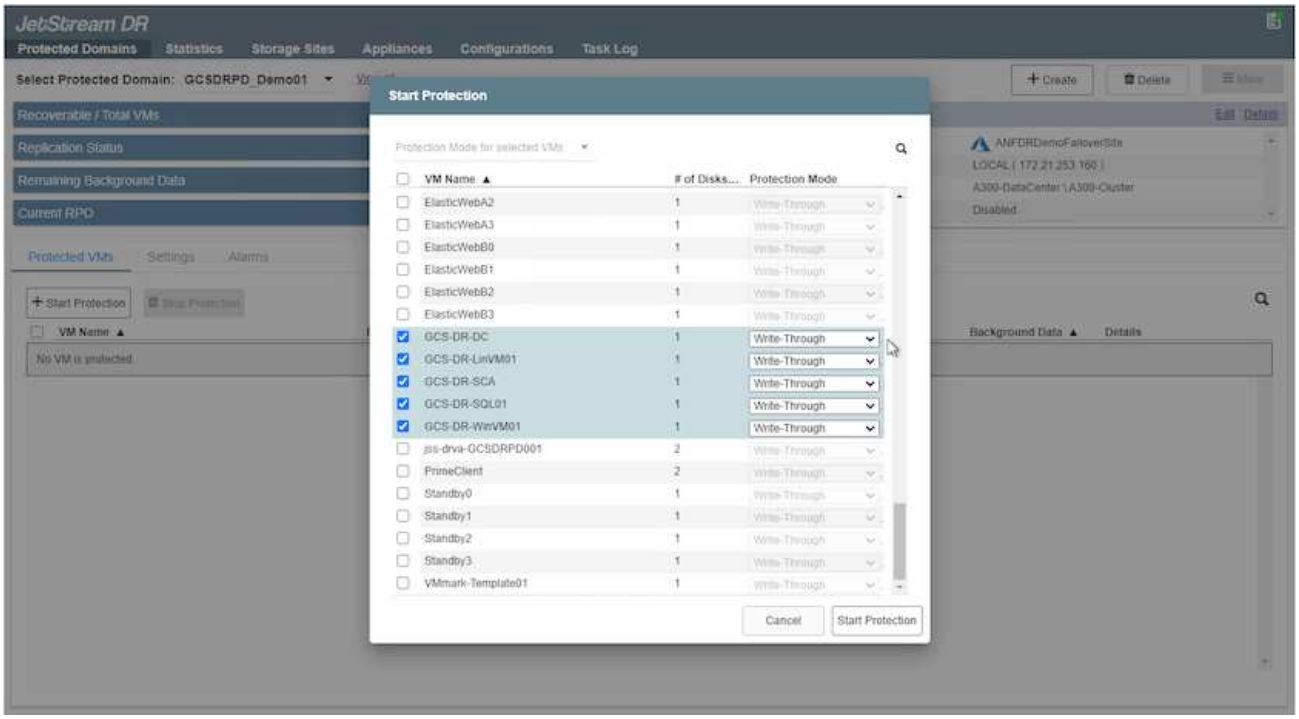


11. Select the VMs to be protected and group the VMs into applications groups based on dependency. Application definitions allow you to group sets of VMs into logical groups that contain their boot orders,

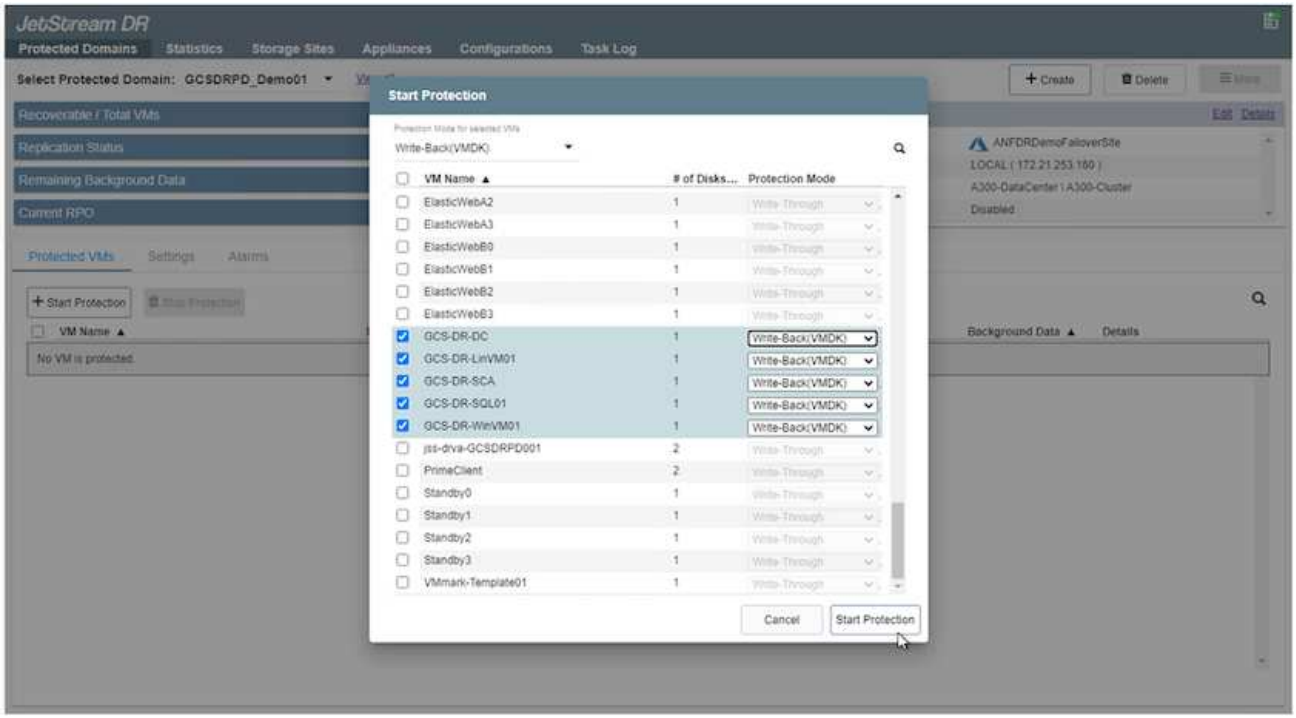
boot delays, and optional application validations that can be executed upon recovery.

Make sure that the same protection mode is used for all VMs in a protected domain.

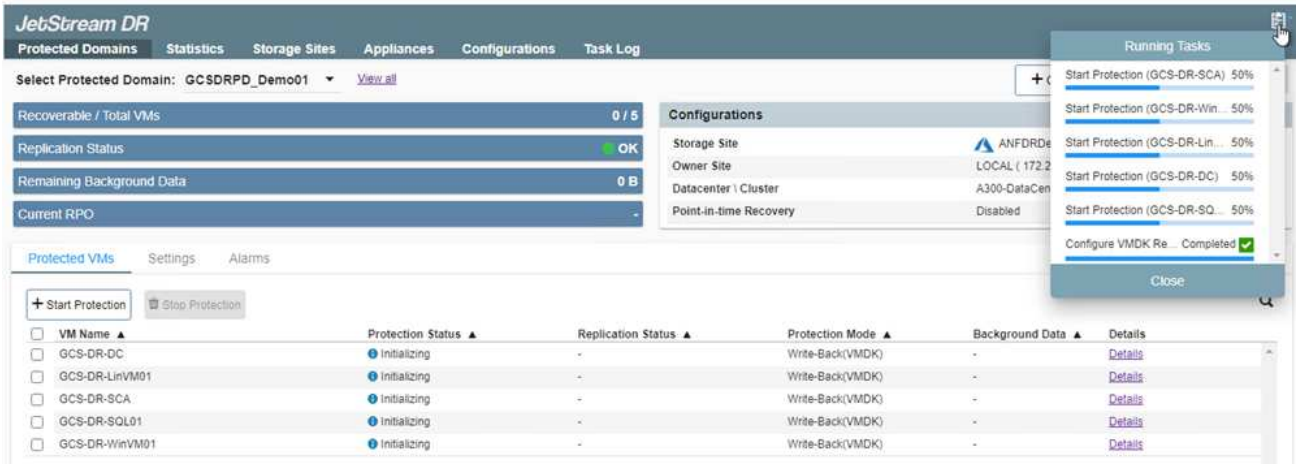
Write-Back(VMDK) mode offers higher performance.



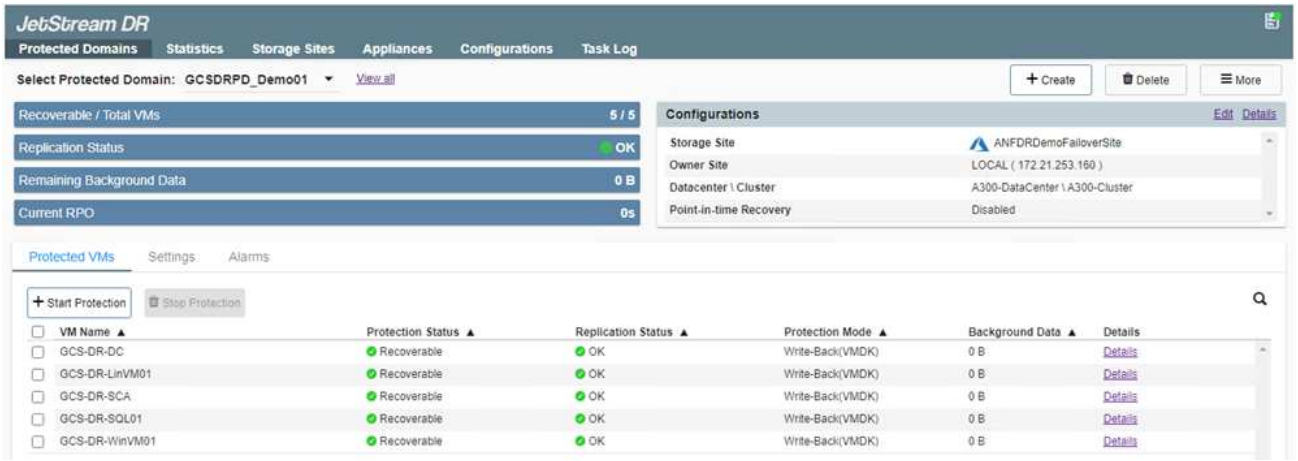
12. Make sure that replication log volumes are placed on high- performance storage.



13. After you are done, click Start Protection for the protected domain. This starts data replication for the selected VMs to the designated Blob store.



14. After replication is completed, the VM protection status is marked as Recoverable.



> Failover runbooks can be configured to group the VMs (called a recovery group), set the boot order sequence, and modify the CPU/memory settings along with the IP configurations.

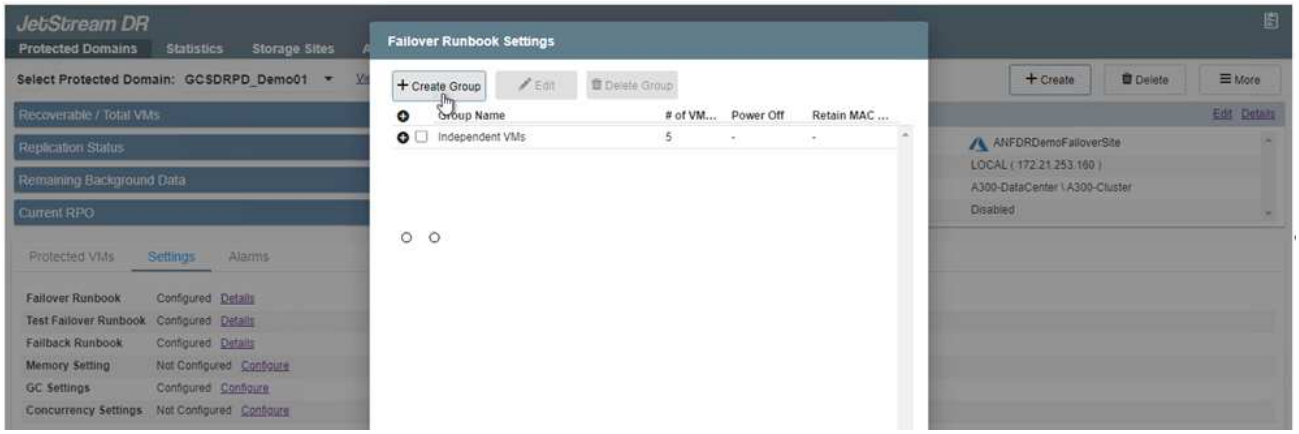
15. Click Settings and then click the runbook Configure link to configure the runbook group.



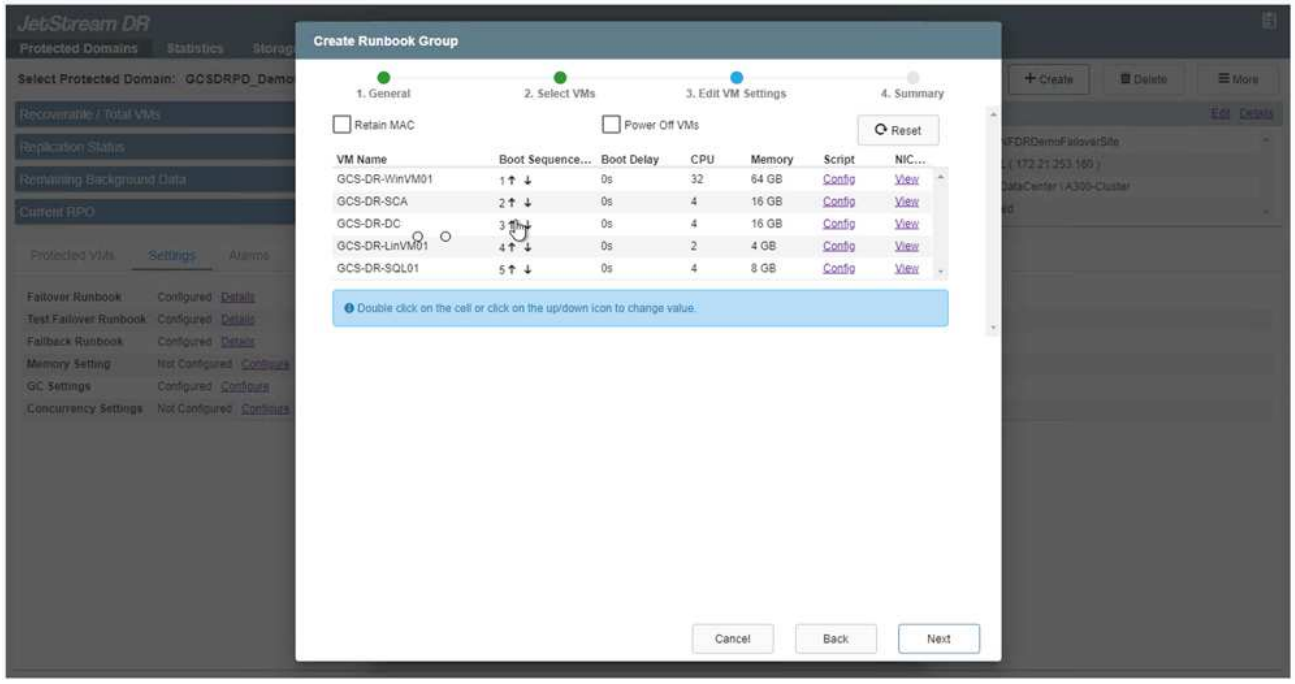
16. Click the Create Group button to begin creating a new runbook group.

> If needed, in the lower portion of the screen, apply custom pre-scripts and post-scripts to

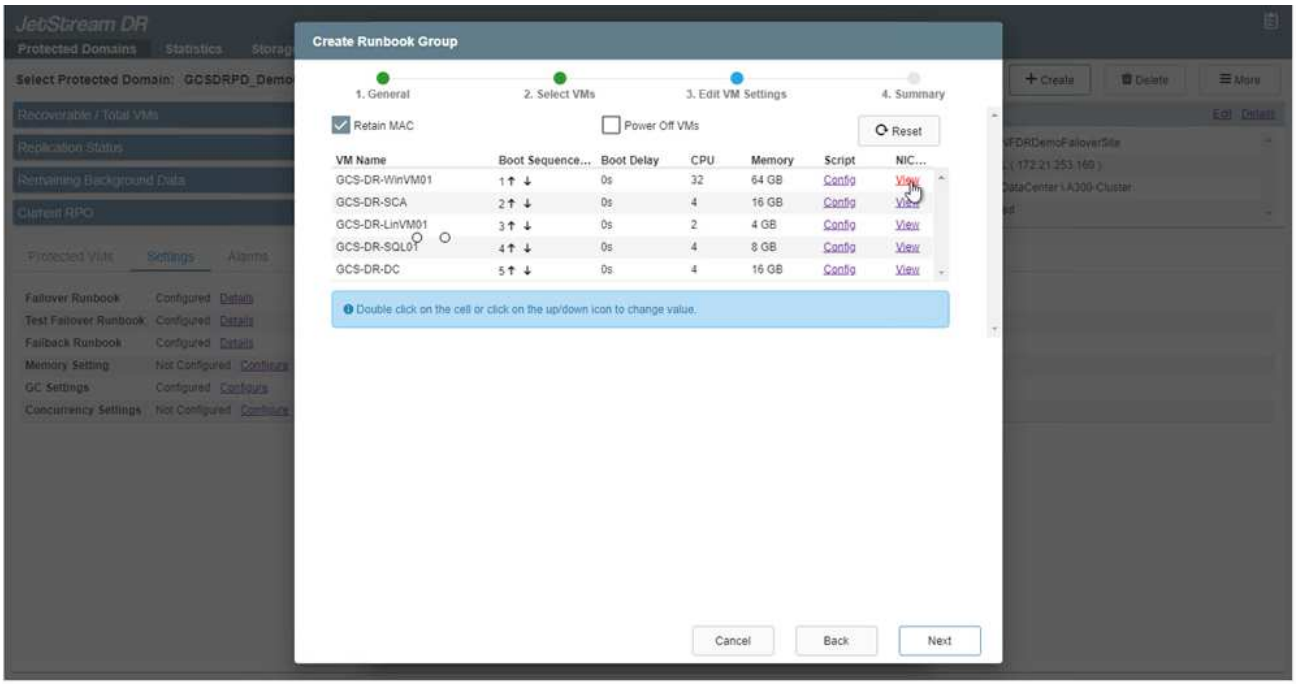
automatically run prior to and following operation of the runbook group. Make sure that the Runbook scripts are residing on the management server.



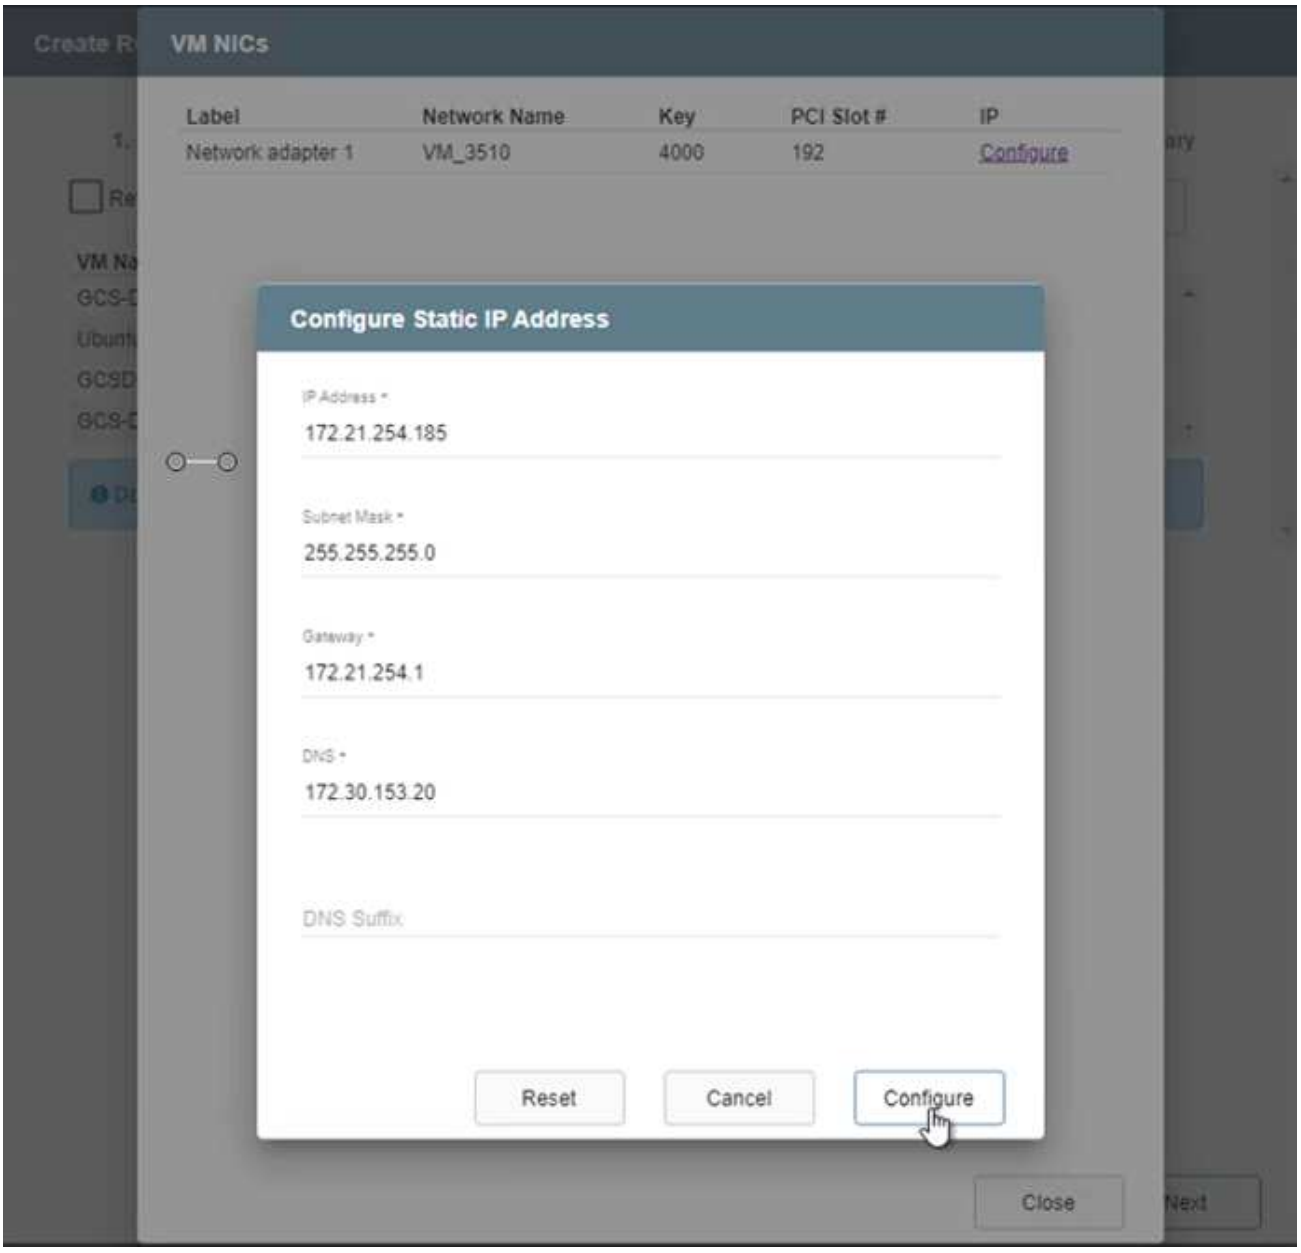
17. Edit the VM settings as required. Specify the parameters for recovering the VMs, including the boot sequence, the boot delay (specified in seconds), the number of CPUs, and the amount of memory to allocate. Change the boot sequence of the VMs by clicking the up or down arrows. Options are also provided to Retain MAC.

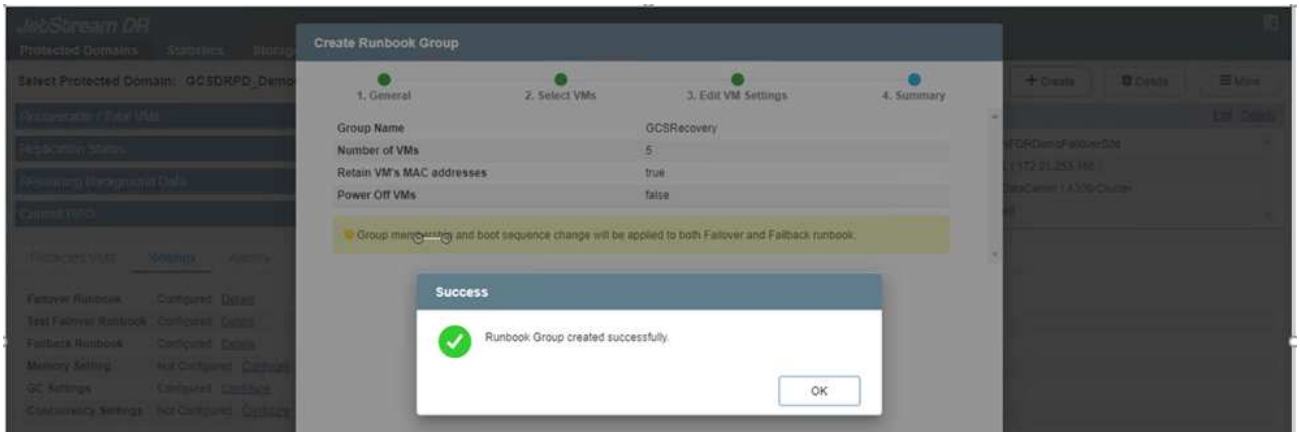


18. Static IP addresses can be manually configured for the individual VMs of the group. Click the NIC View link of a VM to manually configure its IP address settings.



19. Click the Configure button to save NIC settings for the respective VMs.





The status of both the failover and failback runbooks is now listed as Configured. Failover and failback runbook groups are created in pairs using the same initial group of VMs and settings. If necessary, the settings of any runbook group can be individually customized by clicking its respective Details link and making changes.

▼ Install JetStream DR for AVS in private cloud

A best practice for a recovery site (AVS) is to create a three-node pilot-light cluster in advance. This allows the recovery site infrastructure to be preconfigured, including the following:

- Destination networking segments, firewalls, services like DHCP and DNS, and so on
- Installation of JetStream DR for AVS
- Configuration of ANF volumes as datastores and more

JetStream DR supports a near-zero RTO mode for mission-critical domains. For these domains, destination storage should be preinstalled. ANF is a recommended storage type in this case.

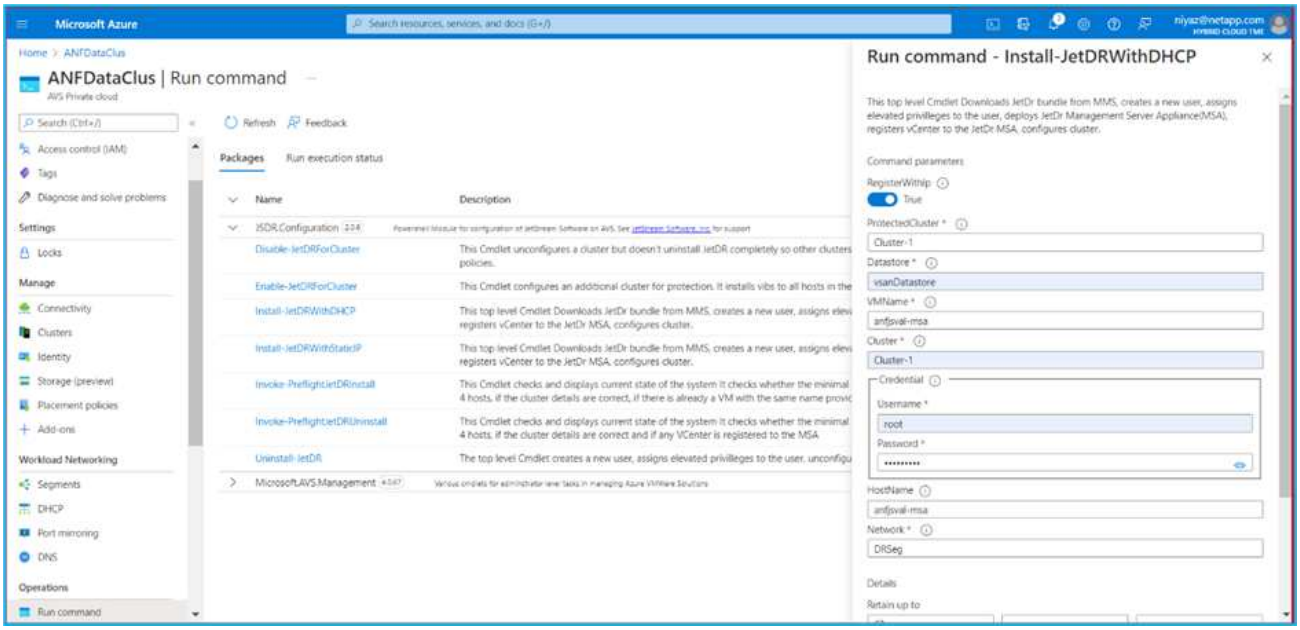
ⓘ Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.

ⓘ Depending on the SLA and RTO requirements, you can use continuous failover or regular (standard) failover mode. For near-zero RTO, you should start continuous rehydration at the recovery site.

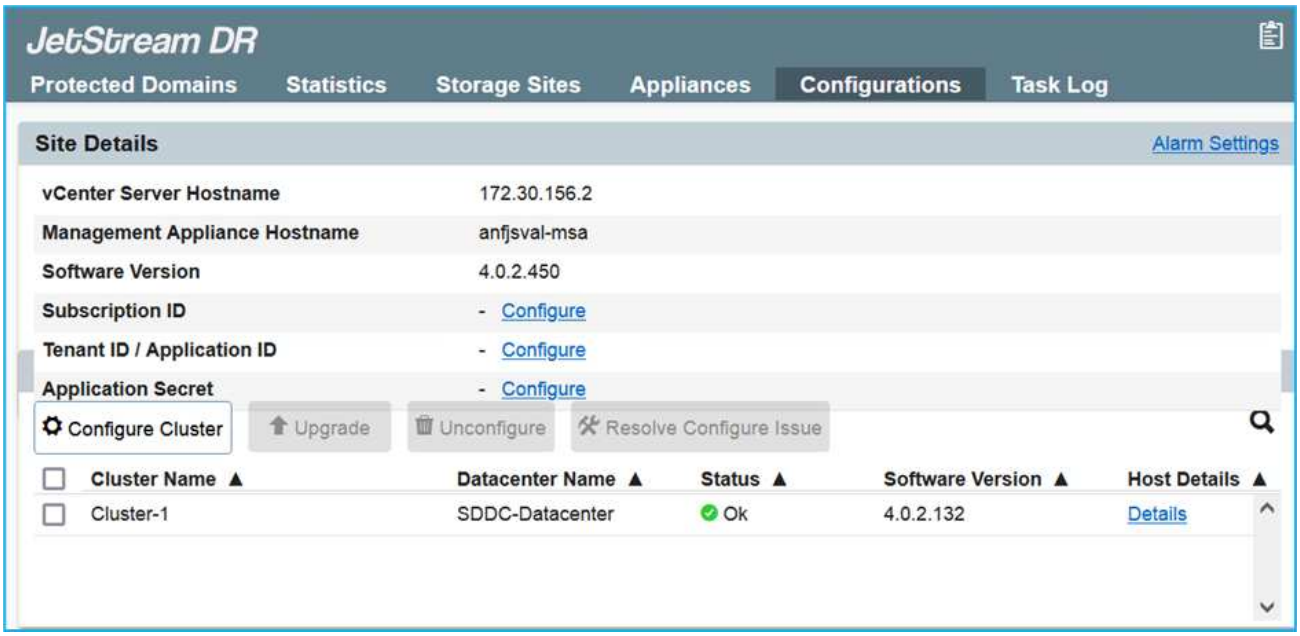
1. To install JetStream DR for AVS on an Azure VMware Solution private cloud, use the Run command. From the Azure portal, go to Azure VMware solution, select the private cloud, and select Run command > Packages > JSDR.Configuration.

   ⓘ The default CloudAdmin user of the Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. The Azure VMware Solution enables simplified and automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

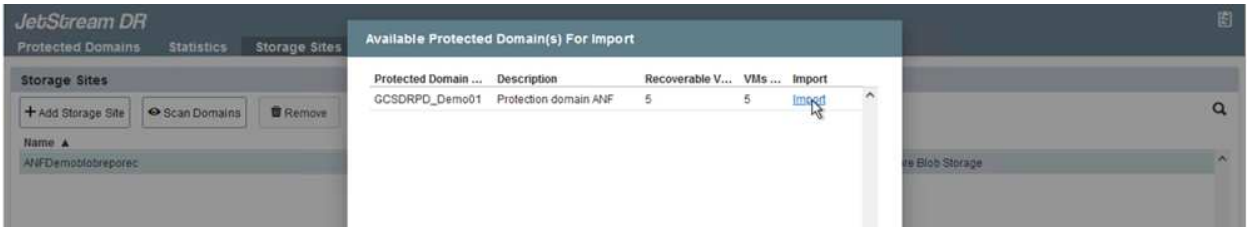
   The following screenshot shows installation using a DHCP-based IP address.



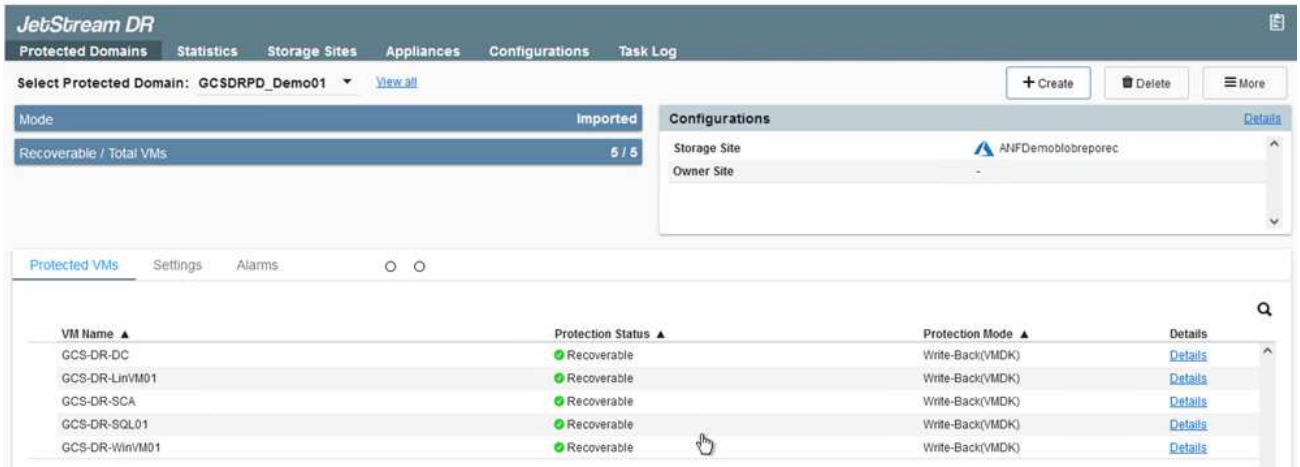
2. After JetStream DR for AVS installation is complete, refresh the browser. To access the JetStream DR UI, go to SDDC Datacenter > Configure > JetStream DR.



3. From the JetStream DR interface, complete the following tasks:
   a. Add the Azure Blob Storage account that was used to protect the on-premises cluster as a storage site and then run the Scan Domains option.
   b. In the pop-up dialog window that appears, select the protected domain to import and then click its Import link.



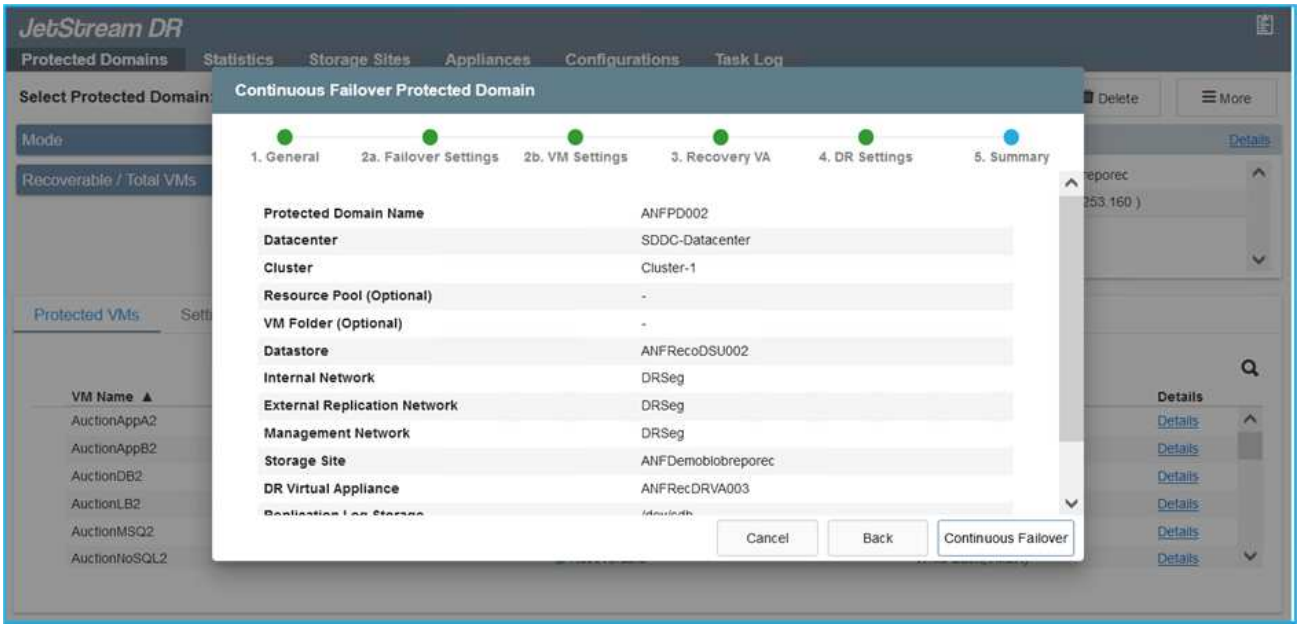
4. The domain is imported for recovery. Go to the Protected Domains tab and verify that the intended domain has been selected or choose the desired one from the Select Protected Domain menu. A list of the recoverable VMs in the protected domain is displayed.



5. After the protected domains are imported, deploy DRVA appliances.

These steps can also be automated using CPT- created plans.

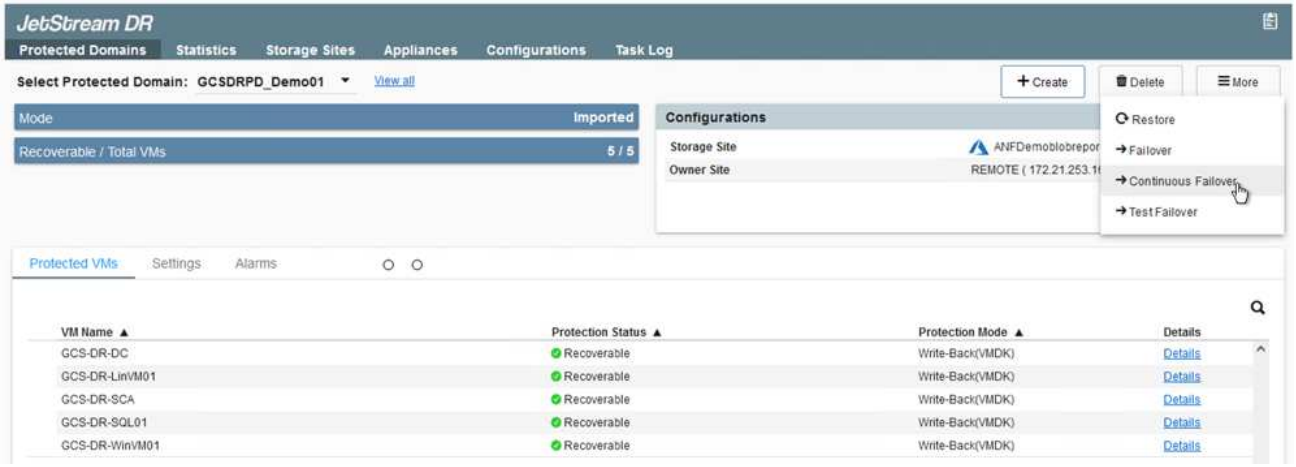
6. Create replication log volumes using available vSAN or ANF datastores.
7. Import the protected domains and configure the recovery VA to use an ANF datastore for VM placements.



Make sure that DHCP is enabled on the selected segment and that enough IPs are available. Dynamic IPs are temporarily used while domains are recovering. Each recovering VM (including continuous rehydration) requires an individual dynamic IP. After recovery is complete, the IP is released and can be reused.

8. Select the appropriate failover option (continuous failover or failover). In this example, continuous rehydration (continuous failover) is selected.

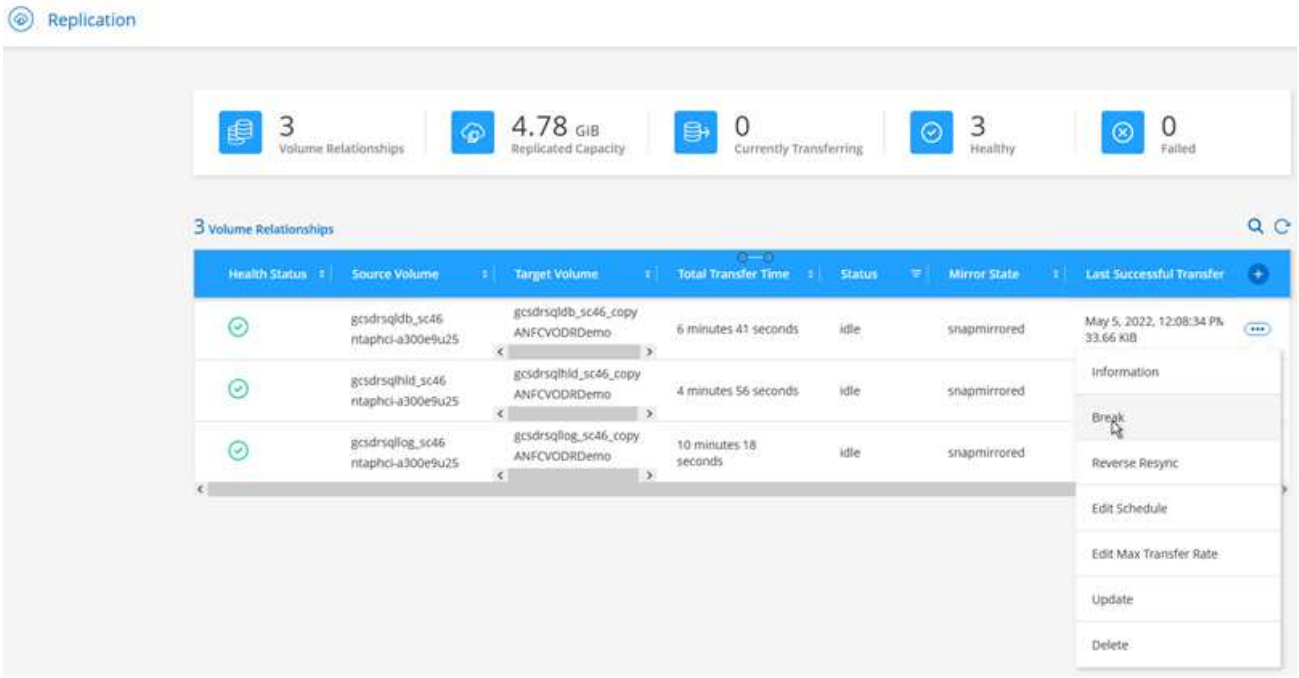
Although Continuous Failover and Failover modes differ on when configuration is performed, both failover modes are configured using the same steps. Failover steps are configured and performed together in response to a disaster event. Continuous failover can be configured at any time and then allowed to run in the background during normal system operation. After a disaster event has occurred, continuous failover is completed to immediately transfer ownership of the protected VMs to the recovery site (near-zero RTO).

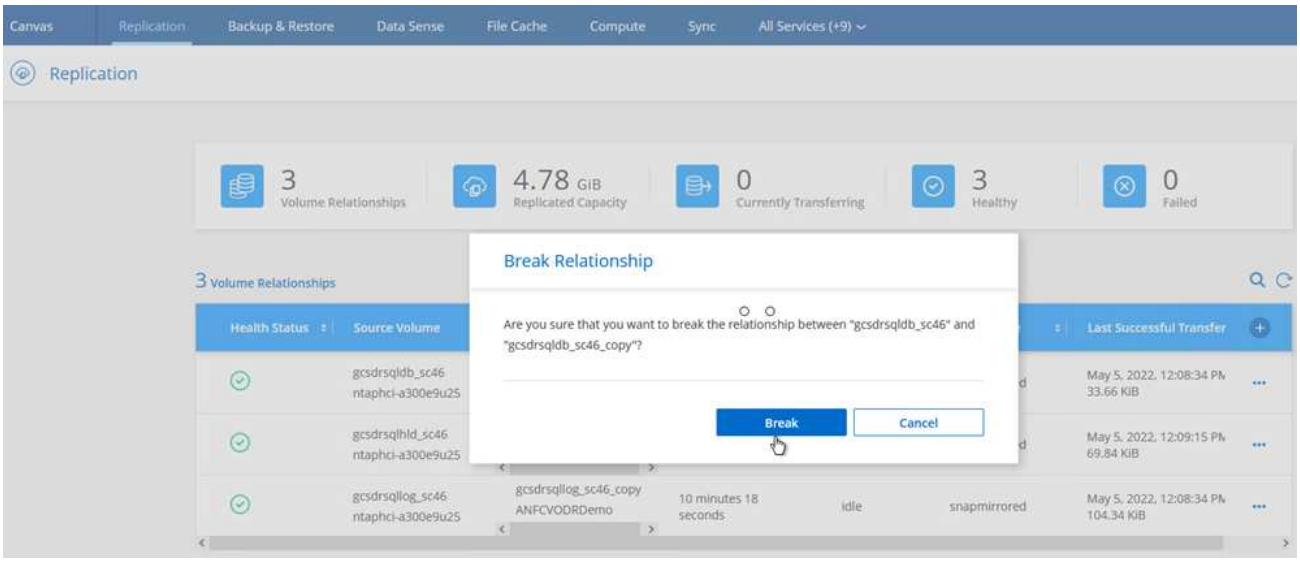


The continuous failover process begins, and its progress can be monitored from the UI. Clicking the blue icon in the Current Step section exposes a pop-up window showing details of the current step of the failover process.

▼ **Failover and Failback**

1. After a disaster occurs in the protected cluster of the on-premises environment (partial or complete failure), you can trigger the failover for VMs using Jetstream after breaking the SnapMirror relationship for the respective application volumes.

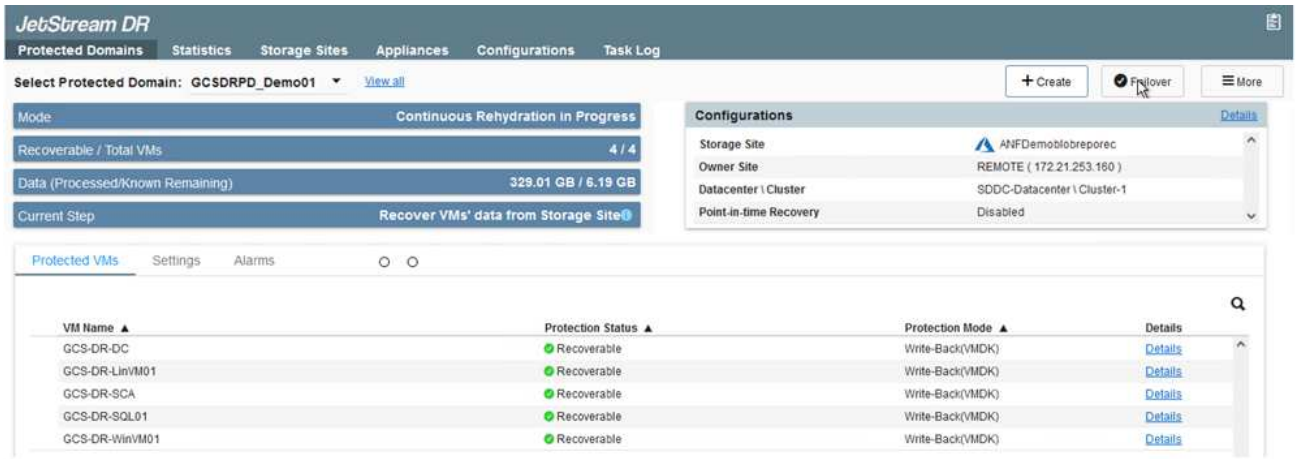


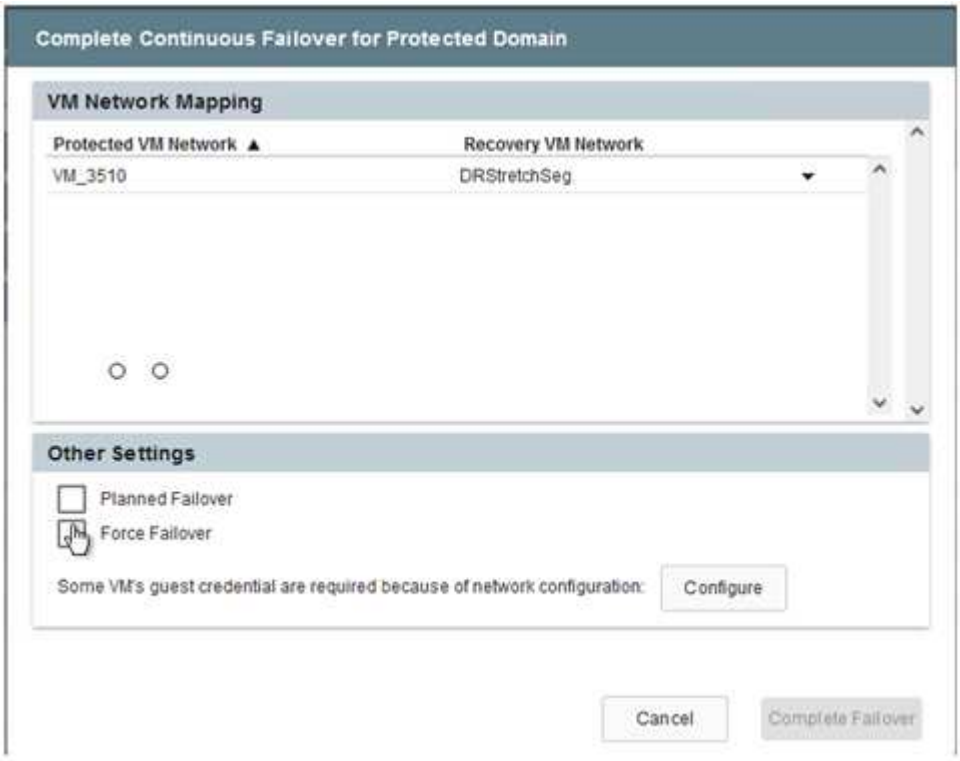


This step can easily be automated to facilitate the recovery process.

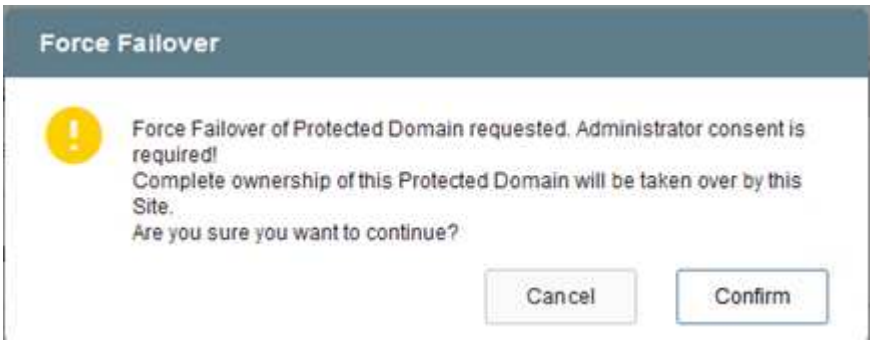
2. Access the Jetstream UI on AVS SDDC (destination side) and trigger the failover option to complete failover. The task bar shows progress for failover activities.

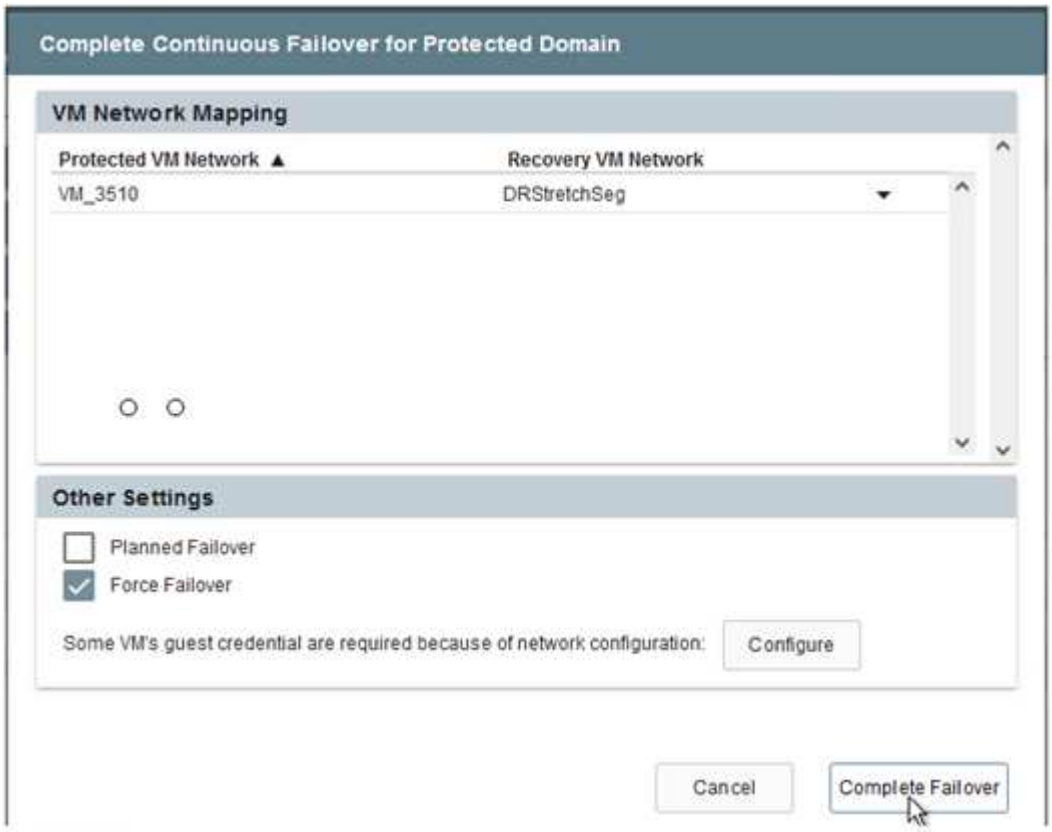
   In the dialog window that appears when completing failover, the failover task can be specified as planned or assumed to be forced.





Forced failover assumes the primary site is no longer accessible and ownership of the protected domain should be directly assumed by the recovery site.





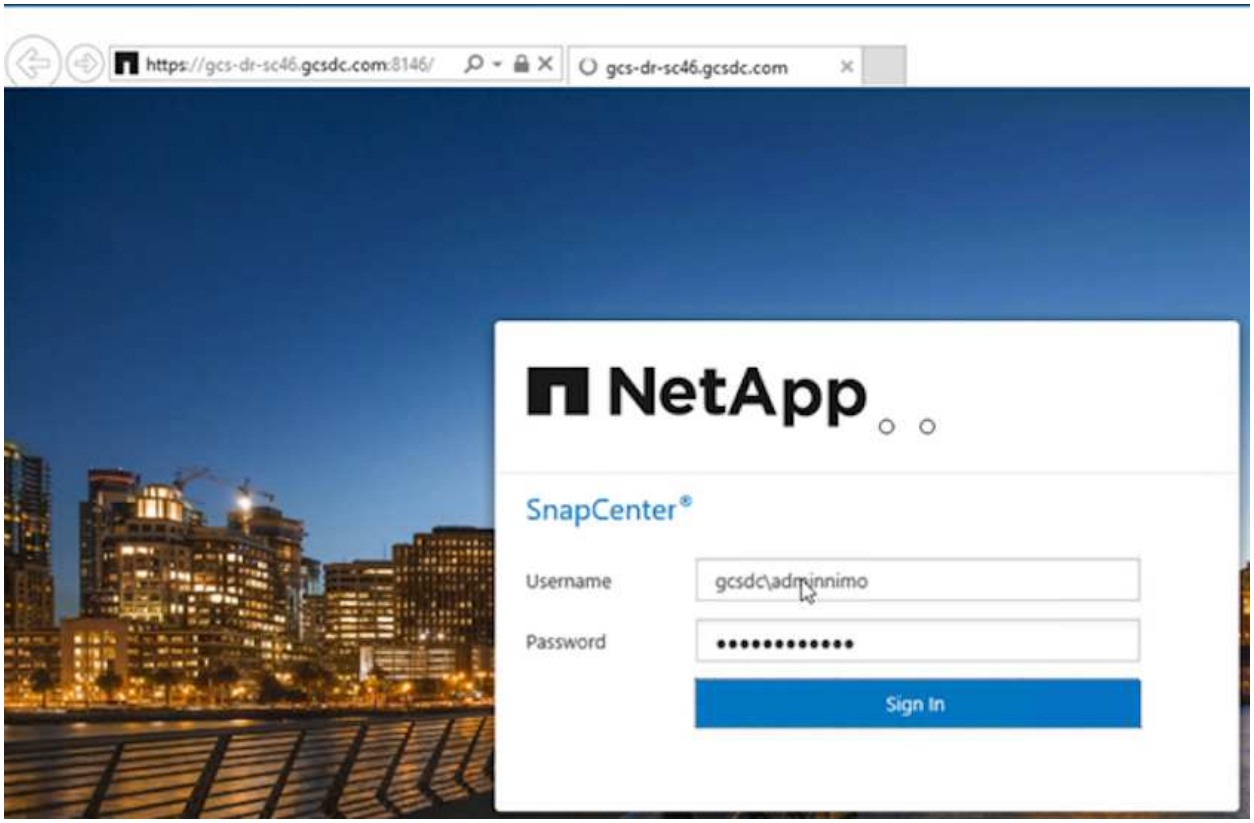
3. After continuous failover is complete, a message appears confirming completion of the task. When the task is complete, access the recovered VMs to configure ISCSI or NFS sessions.

The failover mode changes to Running in Failover and the VM status is Recoverable. All the VMs of the protected domain are now running at the recovery site in the state specified by the failover runbook settings.

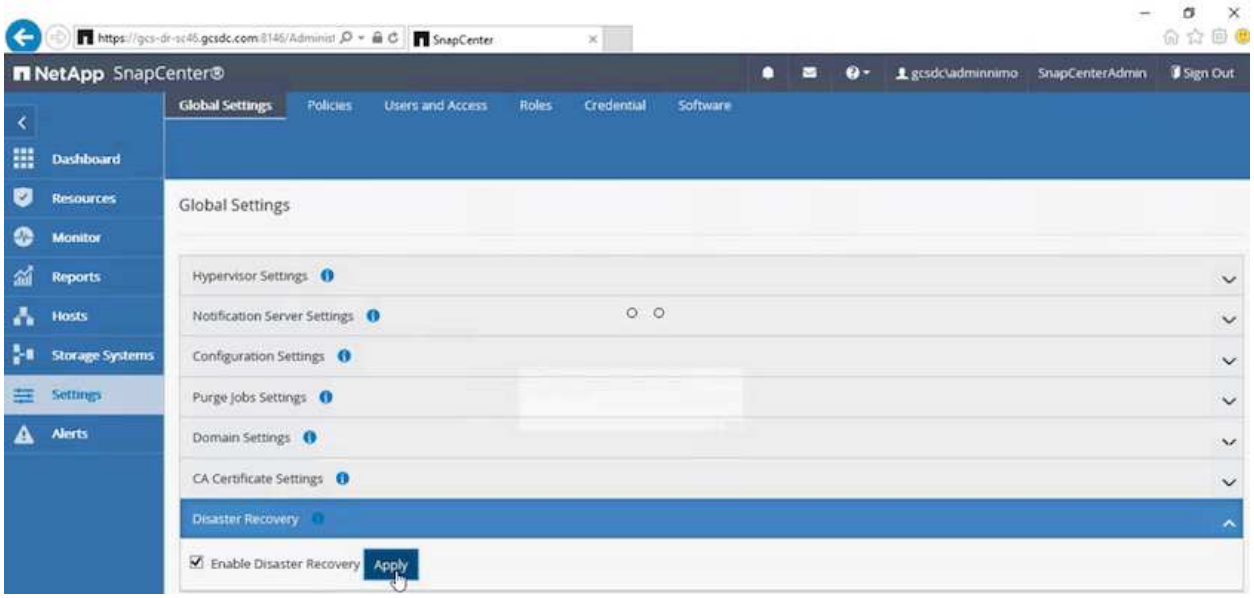
To verify the failover configuration and infrastructure, JetStream DR can be operated in test mode (Test Failover option) to observe the recovery of virtual machines and their data from the object store into a test recovery environment. When a failover procedure is executed in test mode, its operation resembles an actual failover process.



4. After the virtual machines are recovered, use storage disaster recovery for in-guest storage. To demonstrate this process, SQL server is used in this example.
5. Log into the recovered SnapCenter VM on AVS SDDC and enable DR mode.
   a. Access the SnapCenter UI using the browserN.



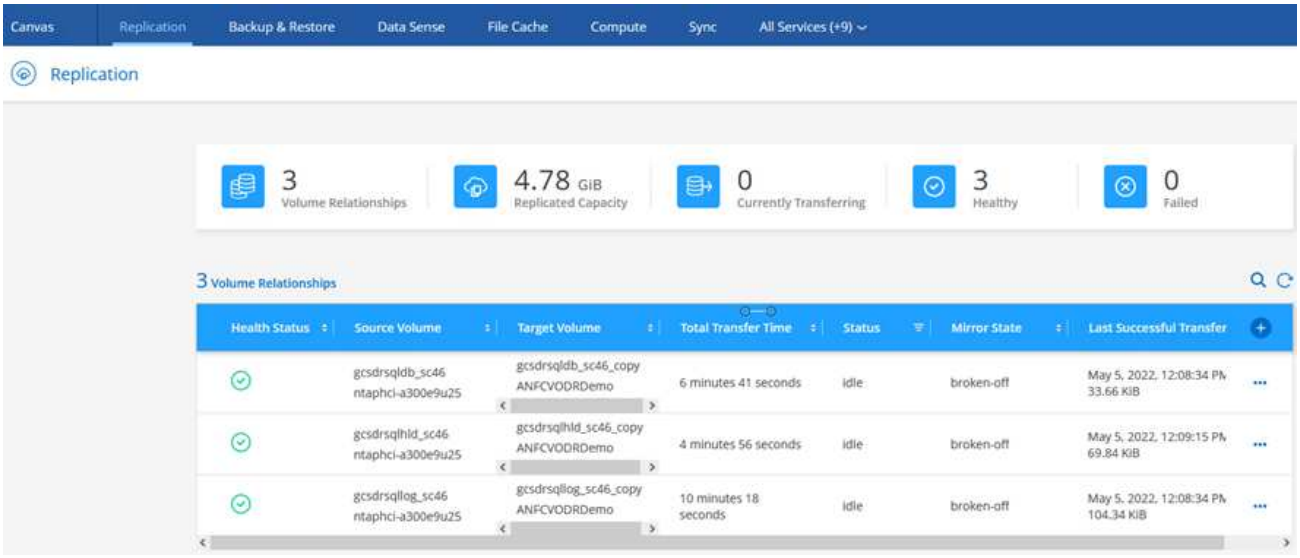
   b. In the Settings page, navigate to Settings > Global Settings > Disaster Recovery.
   c. Select Enable Disaster Recovery.
   d. Click Apply.



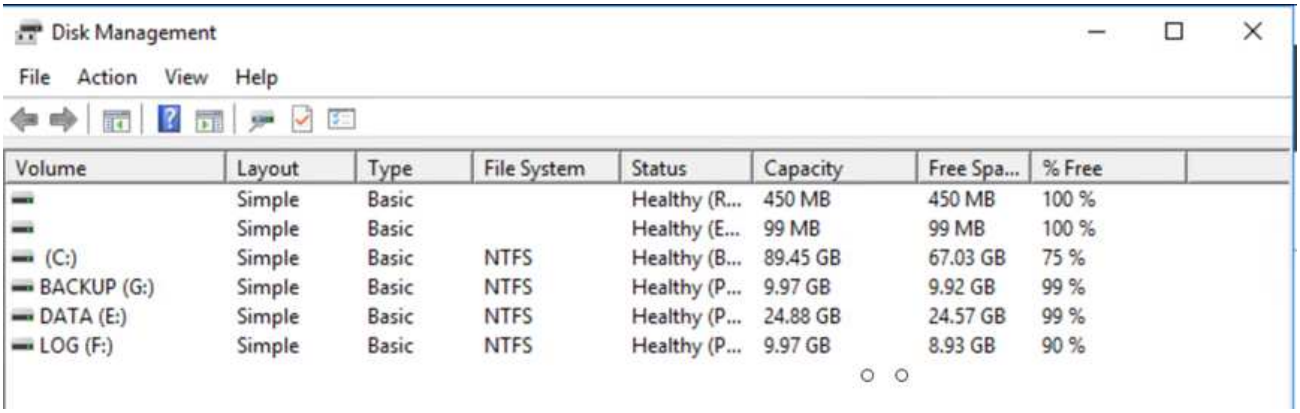
e. Verify whether the DR job is enabled by clicking Monitor > Jobs.

> NetApp SnapCenter 4.6 or later should be used for storage disaster recovery. For previous versions, application-consistent snapshots (replicated using SnapMirror) should be used and manual recovery should be executed in case previous backups must be recovered in the disaster recovery site.

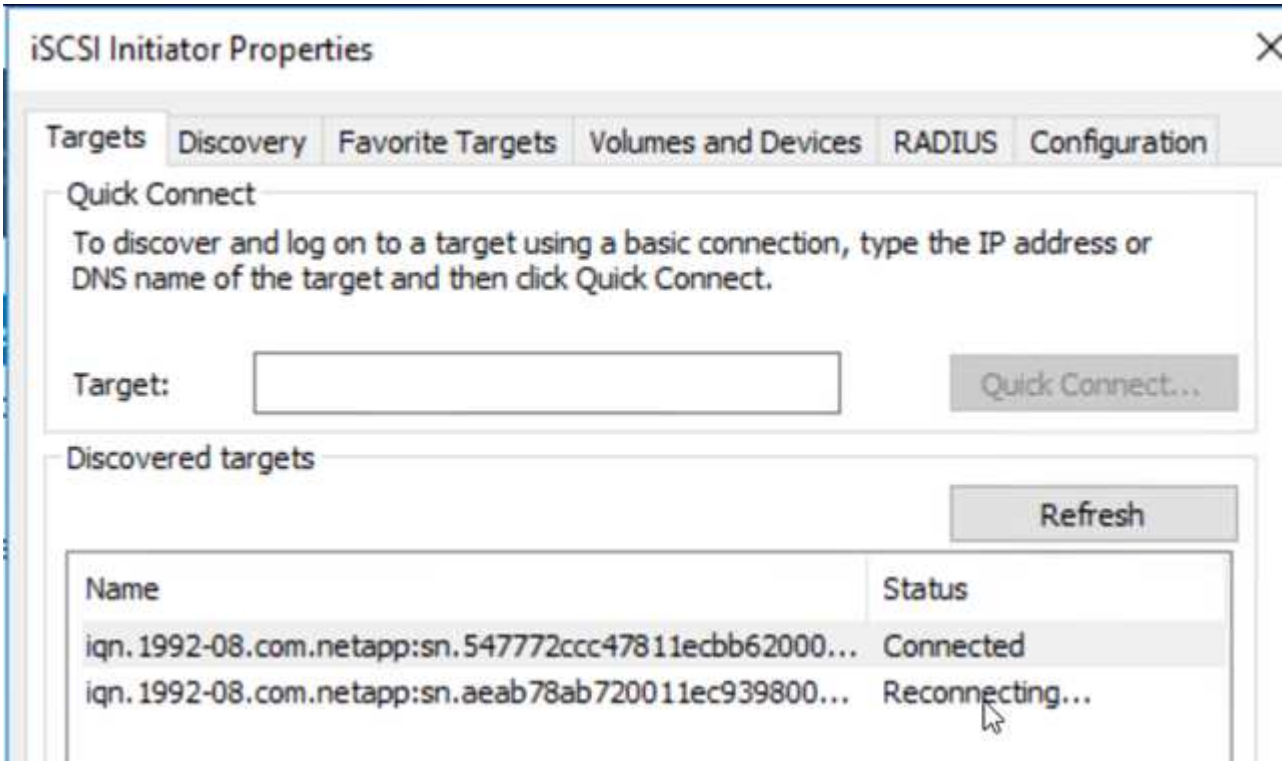
6. Make sure that the SnapMirror relationship is broken.



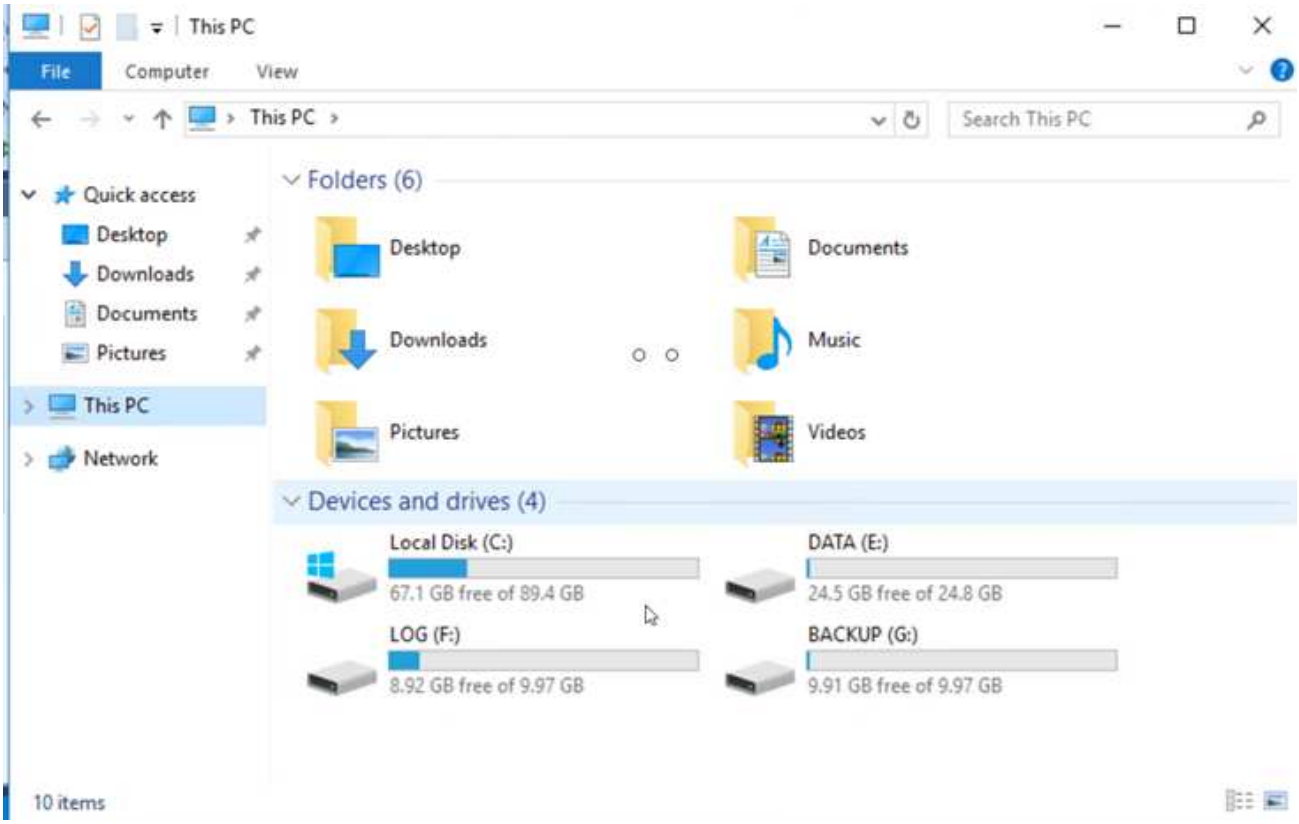
7. Attach the LUN from Cloud Volumes ONTAP to the recovered SQL guest VM with same drive letters.



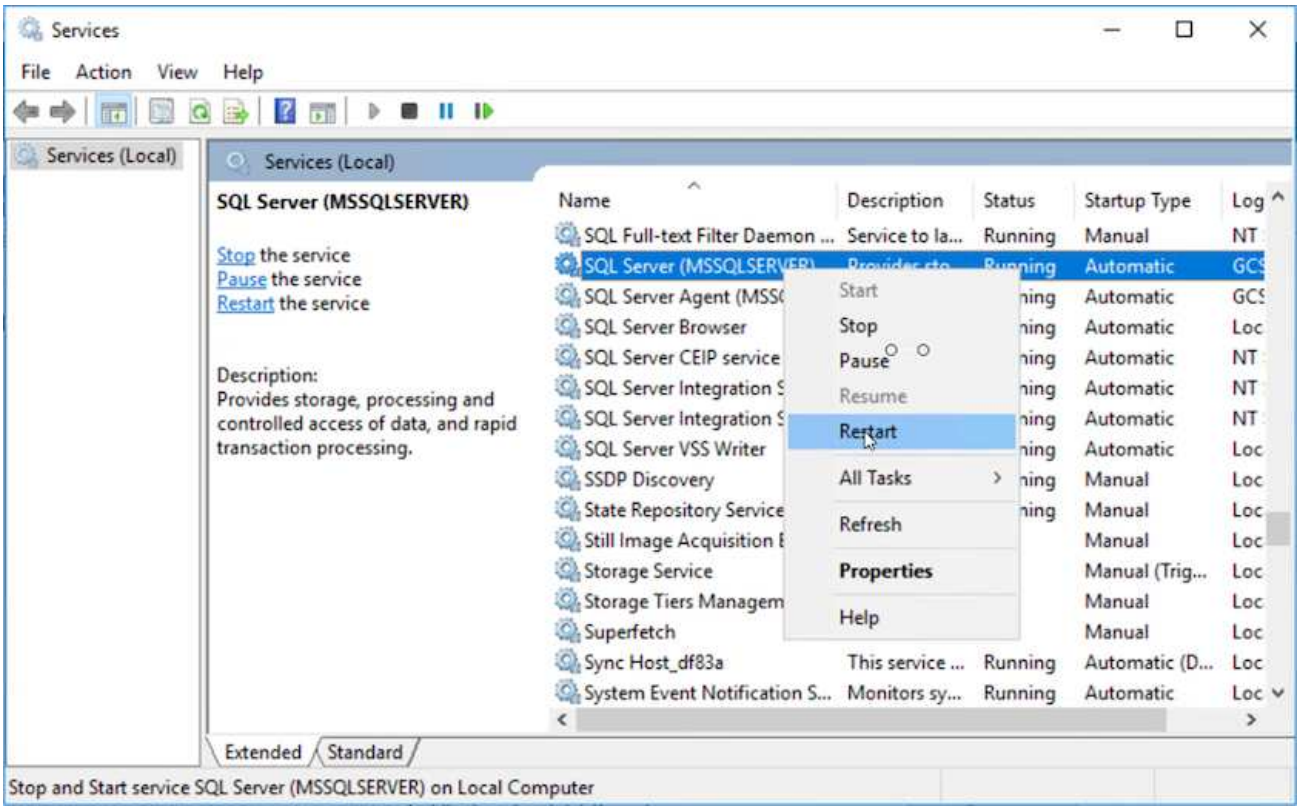
8. Open iSCSI Initiator, clear the previous disconnected session and add the new target along with multipath for the replicated Cloud Volumes ONTAP volumes.



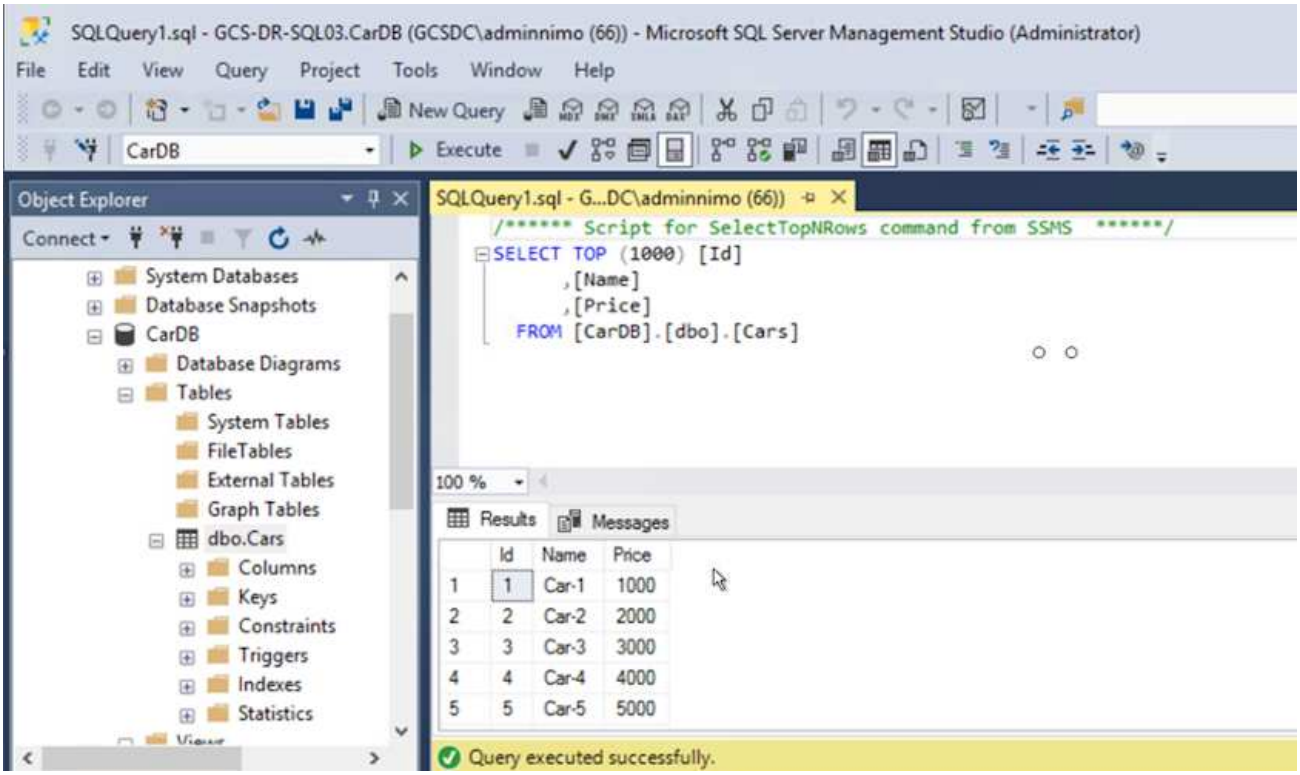
9. Make sure that all the disks are connected using the same drive letters that were used prior to DR.



10. Restart the MSSQL server service.



11. Make sure that the SQL resources are back online.



> In the case of NFS, attach the volumes using the mount command and update the `/etc/fstab` entries.

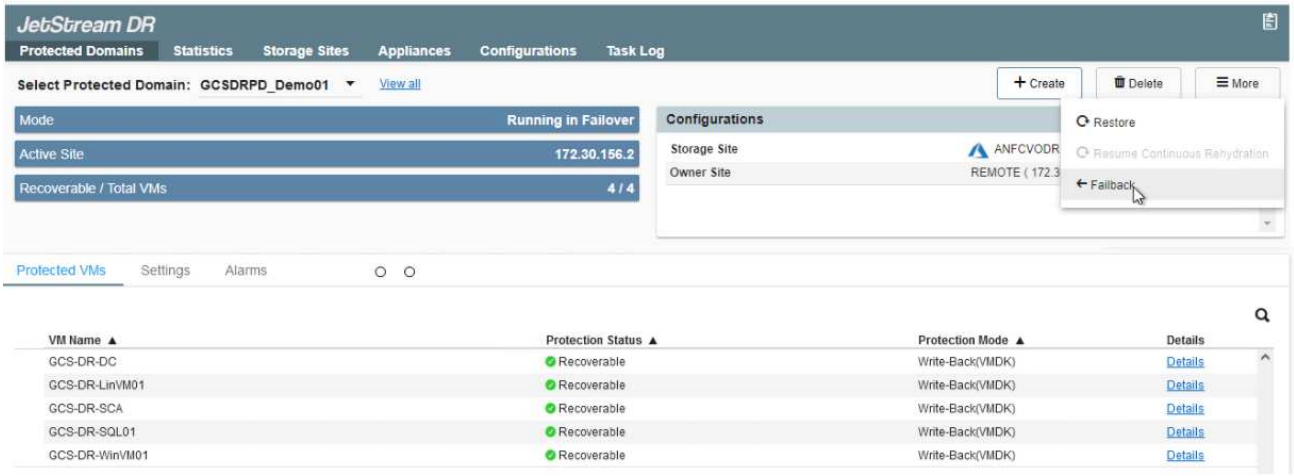
At this point, operations can be run and business continues normally.

On the NSX-T end, a separate dedicated tier-1 gateway can be created for simulating failover scenarios. This ensures that all workloads can communicate with each other but that no traffic can route in or out of the environment, so that any triage, containment, or hardening tasks can be performed without risk of cross-contamination. This operation is outside of the scope of this document, but it can easily be achieved for simulating isolation.

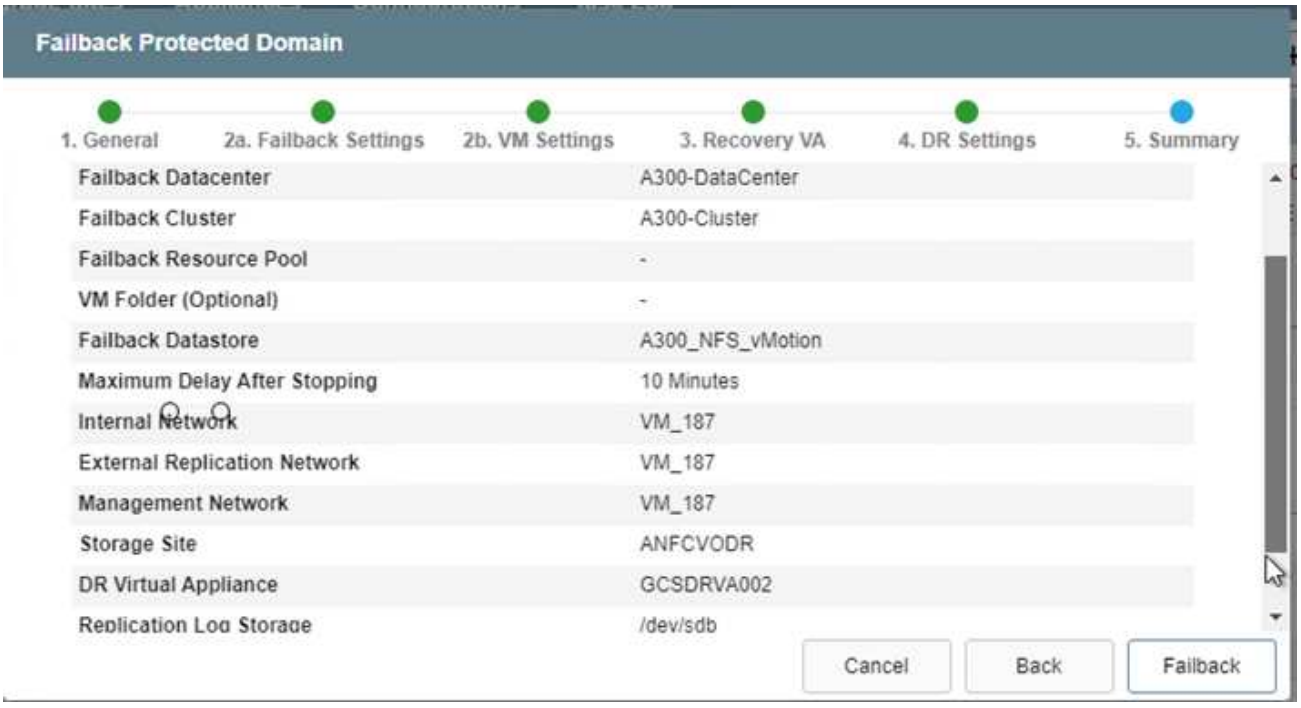
After the primary site is up and running again, you can perform failback. VM protection is resumed by Jetstream and the SnapMirror relationship must be reversed.

1. Restore the on-premises environment. Depending on the type of disaster incident, it might be necessary to restore and/or verify the configuration of the protected cluster. If necessary, JetStream DR software might need to be reinstalled.
2. Access the restored on-premises environment, go to the Jetstream DR UI, and select the appropriate protected domain. After the protected site is ready for failback, select the Failback option in the UI.

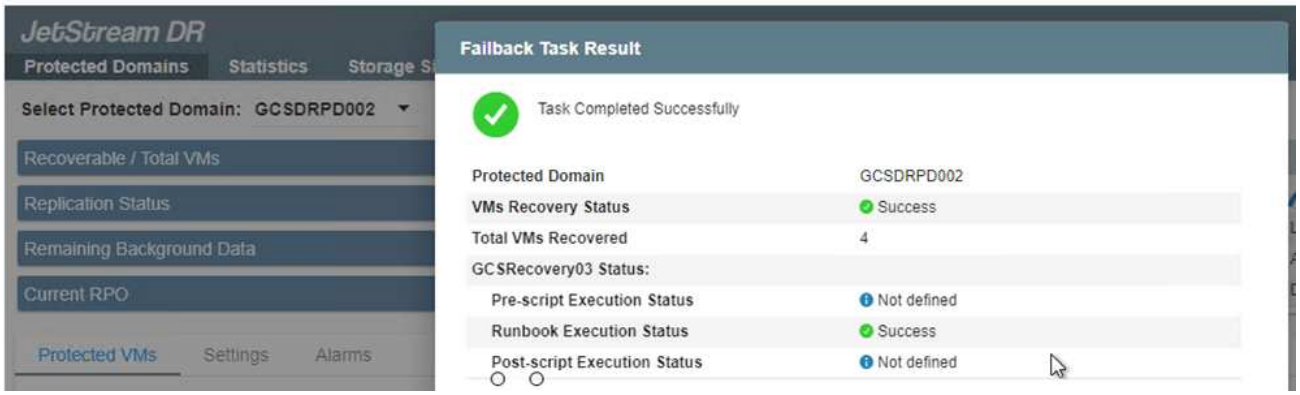
The CPT-generated failback plan can also be used to initiate the return of the VMs and their data from the object store back to the original VMware environment.



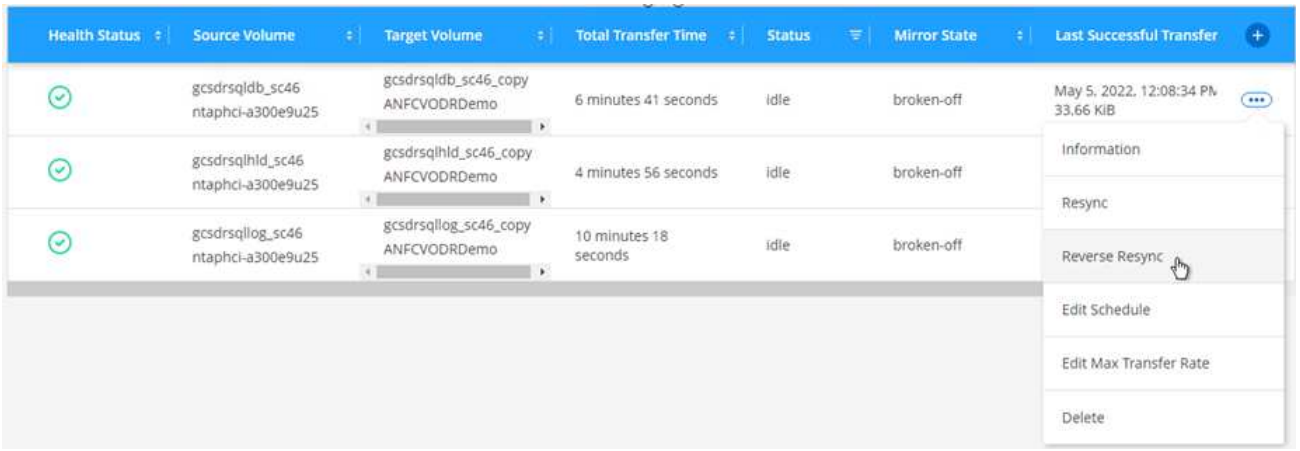
Specify the maximum delay after pausing the VMs in the recovery site and restarting them in the protected site. The time need to complete this process includes the completion of replication after stopping failover VMs, the time needed to clean the recovery site, and the time needed to recreate VMs in the protected site. NetApp recommends 10 minutes.

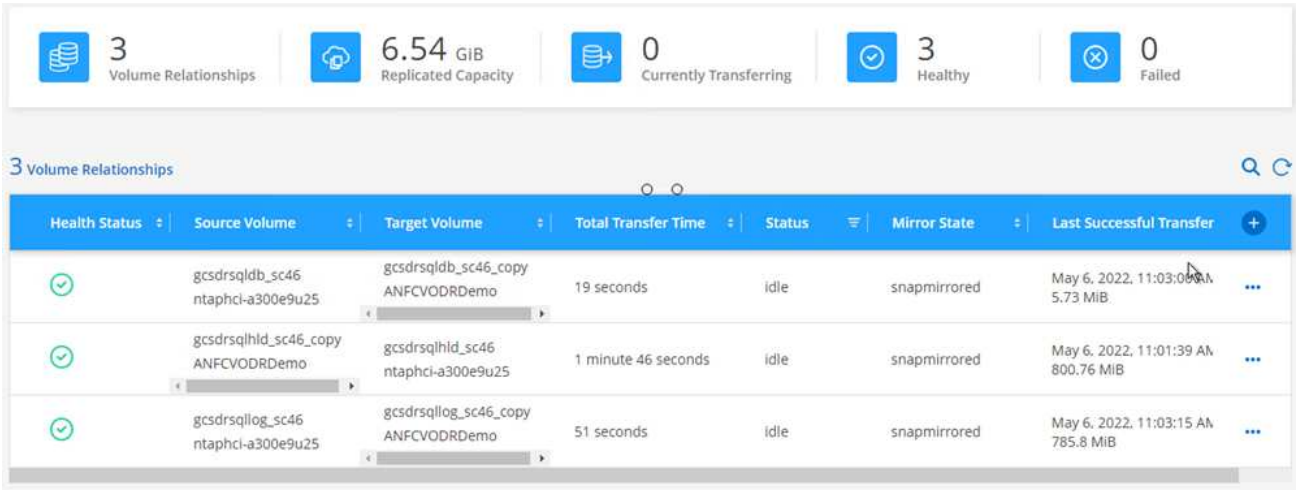


3. Complete the failback process and then confirm the resumption of VM protection and data consistency.

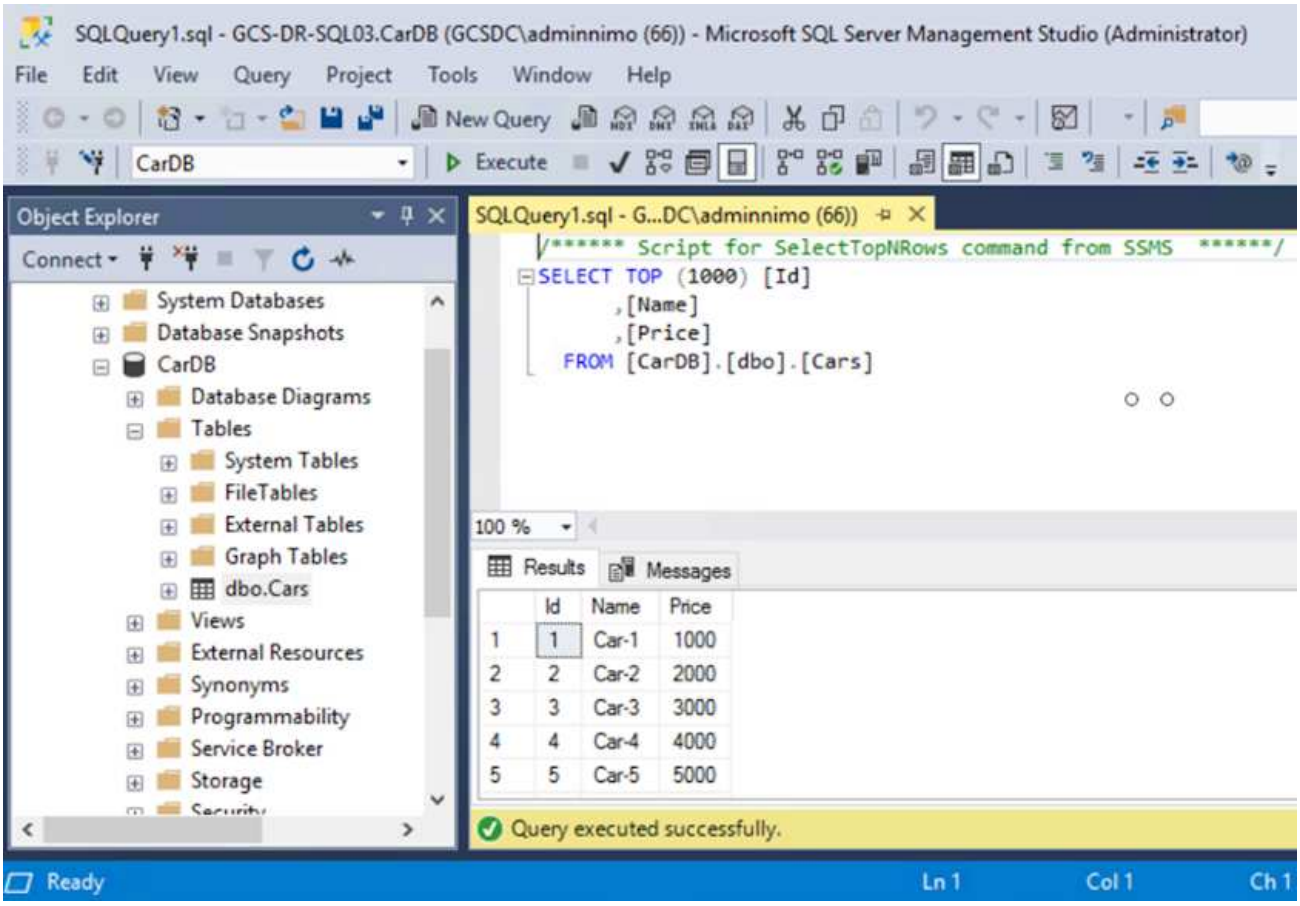


4. After the VMs are recovered, disconnect the secondary storage from the host and connect to the primary storage.





5. Restart the MSSQL server service.
6. Verify that the SQL resources are back online.



> To failback to the primary storage, make sure that the relationship direction remains the same as it was before the failover by performing a reverse resync operation.

> To retain the roles of primary and secondary storage after the reverse resync operation, perform the reverse resync operation again.

This process is applicable to other applications like Oracle, similar database flavors, and any other

applications using guest-connected storage.

As always, test the steps involved for recovering the critical workloads before porting them into production.

Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
  - This avoids the risk of data corruption for volumes and snapshots.
  - Avoids replication interruptions during DR test workflows.
  - Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- CPU and RAM optimization can help lower cloud costs by enabling recovery to smaller compute clusters.

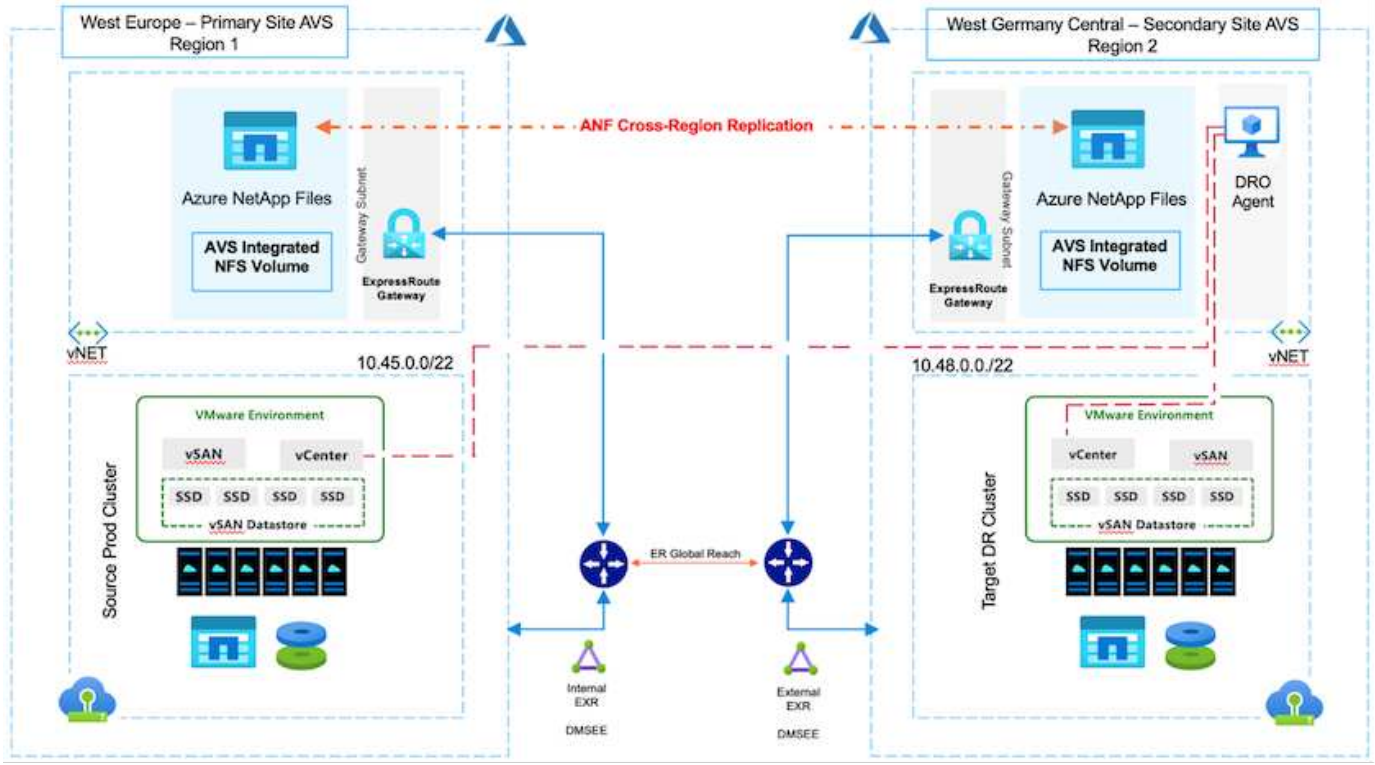
TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS)

Disaster recovery using block-level replication between regions within the cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware).

Overview

With Azure NetApp files (ANF) cross-region volume replication, VMware workloads running on an Azure VMware Solution (AVS) SDDC site using Azure NetApp files volumes as an NFS datastore on the primary AVS site can be replicated to a designated secondary AVS site in the target recovery region.

Disaster Recovery Orchestrator (DRO) (a scripted solution with a UI) can be used to seamlessly recover workloads replicated from one AVS SDDC to another. DRO automates recovery by breaking replication peering and then mounting the destination volume as a datastore, through VM registration to AVS, to network mappings directly on NSX-T (included with all AVS private clouds).



Prerequisites and general recommendations

- Verify that you have enabled cross-region replication by creating replication peering. See Create volume replication for Azure NetApp Files.
- You must configure ExpressRoute Global Reach between the source and target Azure VMware Solution private clouds.
- You must have a service principal that can access resources.
- The following topology is supported: primary AVS site to secondary AVS site.
- Configure the replication schedule for each volume appropriately based on business needs and the data-change rate.

Cascading and fan- in and fan- out topologies are not supported.

Getting started

Deploy Azure VMware Solution

The Azure VMware Solution (AVS) is a hybrid cloud service that provides fully functional VMware SDDCs within a Microsoft Azure public cloud. AVS is a first-party solution fully managed and supported by Microsoft and verified by VMware that uses Azure infrastructure. Therefore, customers get VMware ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data- center facilities, and proximity to the rich ecosystem of native Azure services and solutions. A combination of Azure VMware Solution SDDC and Azure NetApp Files provides the best performance with minimal network latency.

To configure an AVS private cloud on Azure, follow the steps in this link for NetApp documentation and in this link for Microsoft documentation. A pilot- light environment set up with a minimal configuration can be used for DR purposes. This setup only contains core components to support critical applications, and it can scale out and spawn more hosts to take the bulk of the load if a failover occurs.

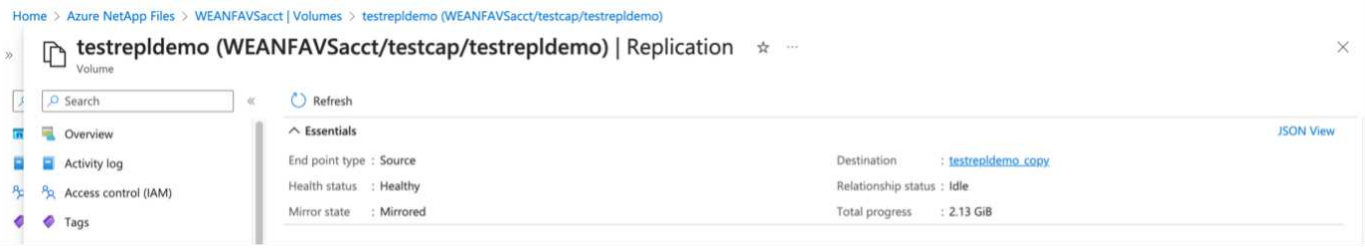
In the initial release, DRO supports an existing AVS SDDC cluster. On-demand SDDC creation will be available in an upcoming release.

Provision and configure Azure NetApp Files

Azure NetApp Files is a high-performance, enterprise-class, metered file- storage service. Follow the steps in this link to provision and configure Azure NetApp Files as a NFS datastore to optimize AVS private cloud deployments.

Create volume replication for Azure NetApp Files-powered datastore volumes

The first step is to set up cross- region replication for the desired datastore volumes from the AVS primary site to the AVS secondary site with the appropriate frequencies and retentions.



Follow the steps in this link to set up cross-region replication by creating replication peering. The service level for the destination capacity pool can match that of the source capacity pool. However, for this specific use case, you can select the standard service level and then modify the service level in the event of a real disaster or DR simulations.

A cross- region replication relationship is a prerequisite and must be created beforehand.

DRO installation

To get started with DRO, use the Ubuntu operating system on the designated Azure virtual machine and make sure you meet the prerequisites. Then install the package.

Prerequisites:

- Service principal that can access resources.
- Make sure that appropriate connectivity exists to the source and destination SDDC and Azure NetApp Files instances.
- DNS resolution should be in place if you are using DNS names. Otherwise, use IP addresses for vCenter.

OS requirements:

- Ubuntu Focal 20.04 (LTS)The following packages must be installed on the designated agent virtual machine:
- Docker
- Docker- compose
- JqChange `docker.sock` to this new permission: `sudo chmod 666 /var/run/docker.sock`.

The `deploy.sh` script executes all required prerequisites.

The steps are as follows:

1. Download the installation package on the designated virtual machine:

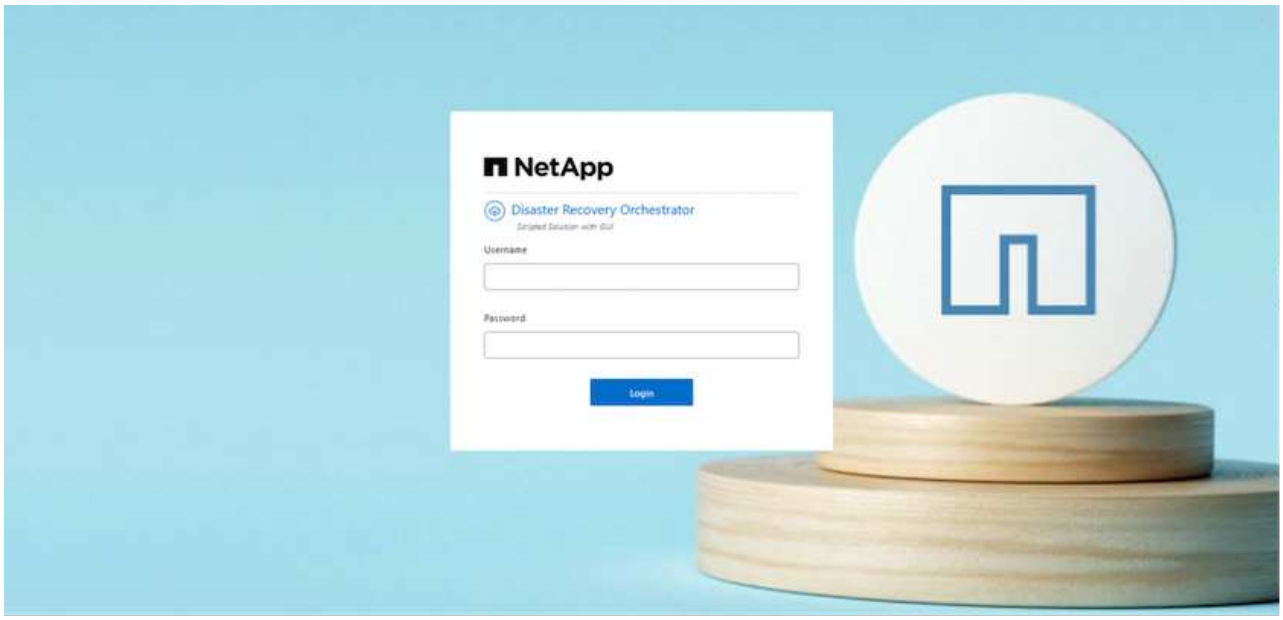
```
git clone https://github.com/NetApp/DRO-Azure.git
```

> The agent must be installed in the secondary AVS site region or in the primary AVS site region in a separate AZ than the SDDC.

2. Unzip the package, run the deployment script, and enter the host IP (for example, `10.10.10.10`).

```
tar xvf draas_package.tar
Navigate to the directory and run the deploy script as below:
sudo sh deploy.sh
```

3. Access the UI using the following credentials:
   - Username: `admin`
   - Password: `admin`



DRO configuration

After Azure NetApp Files and AVS have been configured properly, you can begin configuring DRO to automate the recovery of workloads from the primary AVS site to the secondary AVS site. NetApp recommends deploying the DRO agent in the secondary AVS site and configuring the ExpressRoute gateway connection so that the DRO agent can communicate via the network with the appropriate AVS and Azure NetApp Files components.

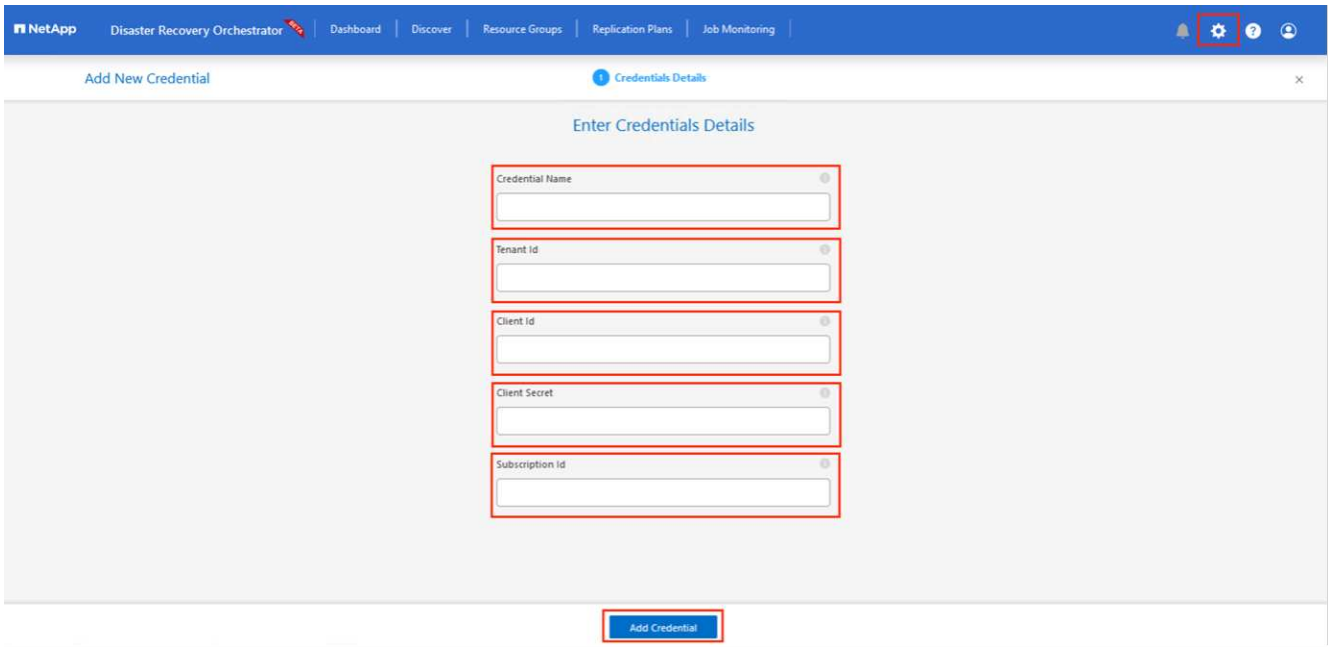
The first step is to Add credentials. DRO requires permission to discover Azure NetApp Files and the Azure VMware Solution. You can grant the required permissions to an Azure account by creating and setting up an Azure Active Directory (AD) application and by obtaining the Azure credentials that DRO needs. You must bind

the service principal to your Azure subscription and assign it a custom role that has the relevant required permissions. When you add source and destination environments, you are prompted to select the credentials associated with the service principal. You need to add these credentials to DRO before you can click Add New Site.

To perform this operation, complete the following steps:

1. Open DRO in a supported browser and use the default username and password (`admin/admin`). The password can be reset after the first login using the Change Password option.
2. In the upper right of the DRO console, click the **Settings** icon, and select **Credentials**.
3. Click Add New Credential and follow the steps in the wizard.
4. To define the credentials, enter information about the Azure Active Directory service principal that grants the required permissions:
   - Credential name
   - Tenant ID
   - Client ID
   - Client secret
   - Subscription ID

     You should have captured this information when you created the AD application.
5. Confirm the details about the new credentials and click Add Credential.

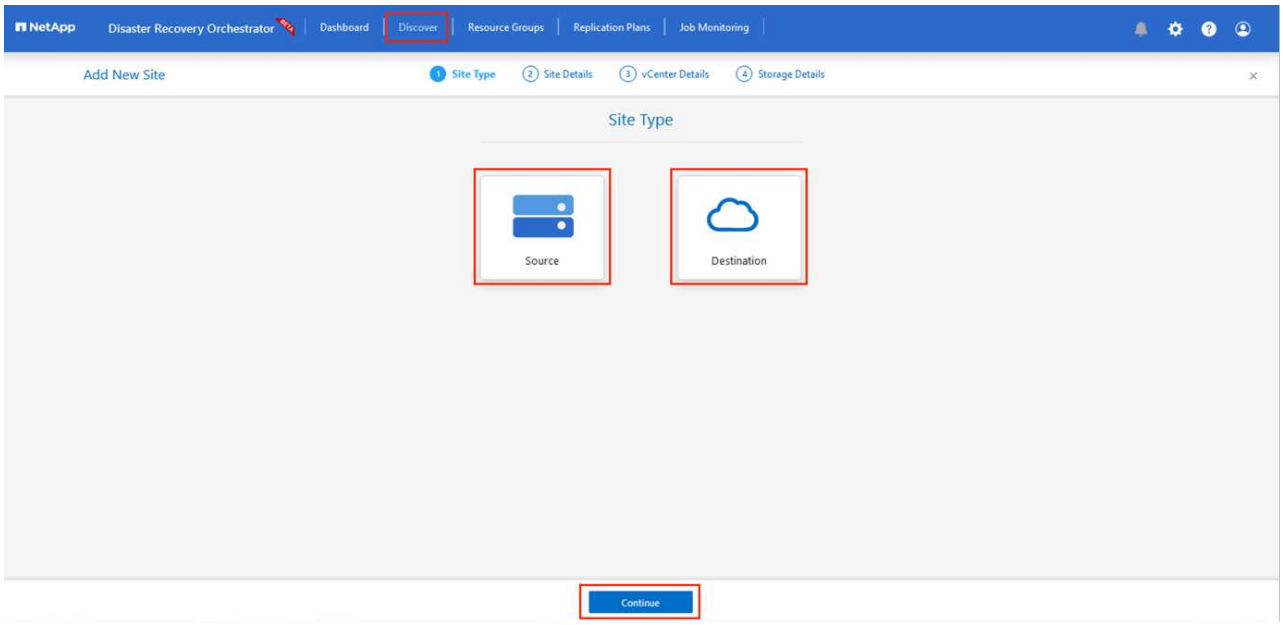


After you add the credentials, it's time to discover and add the primary and secondary AVS sites (both vCenter and the Azure NetApp files storage account) to DRO. To add the source and destination site, complete the following steps:

6. Go to the **Discover** tab.
7. Click **Add New Site**.
8. Add the following primary AVS site (designated as **Source** in the console).

- SDDC vCenter
- Azure NetApp Files storage account

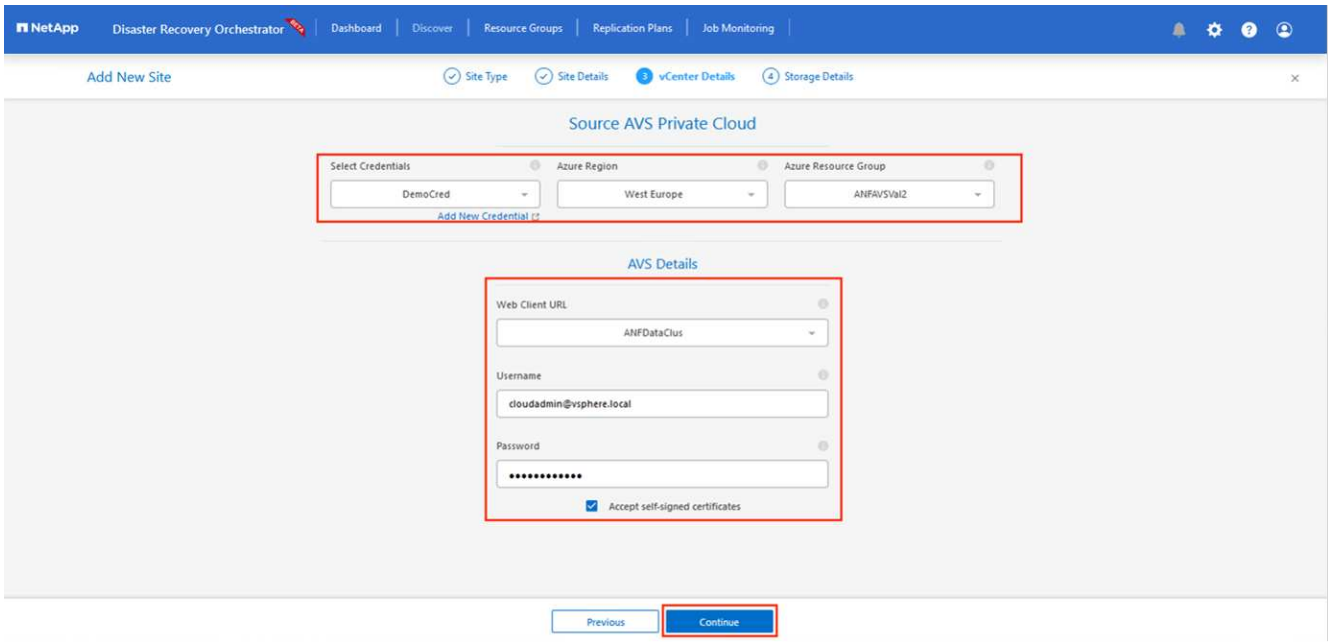
9. Add the following secondary AVS site (designated as **Destination** in the console).
   - SDDC vCenter
   - Azure NetApp Files storage account



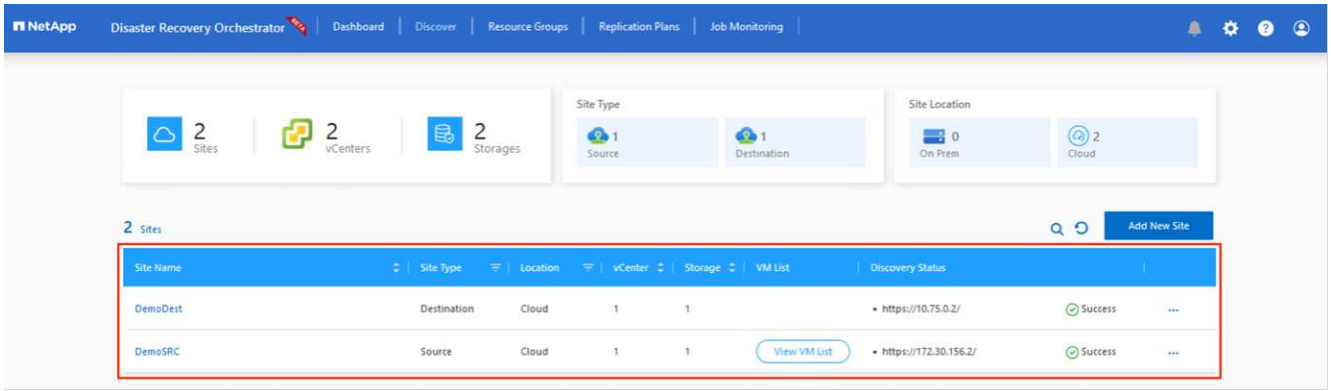
10. Add site details by clicking **Source,** entering a friendly site name, and select the connector. Then click **Continue**.

    For demonstration purposes, adding a source site is covered in this document.

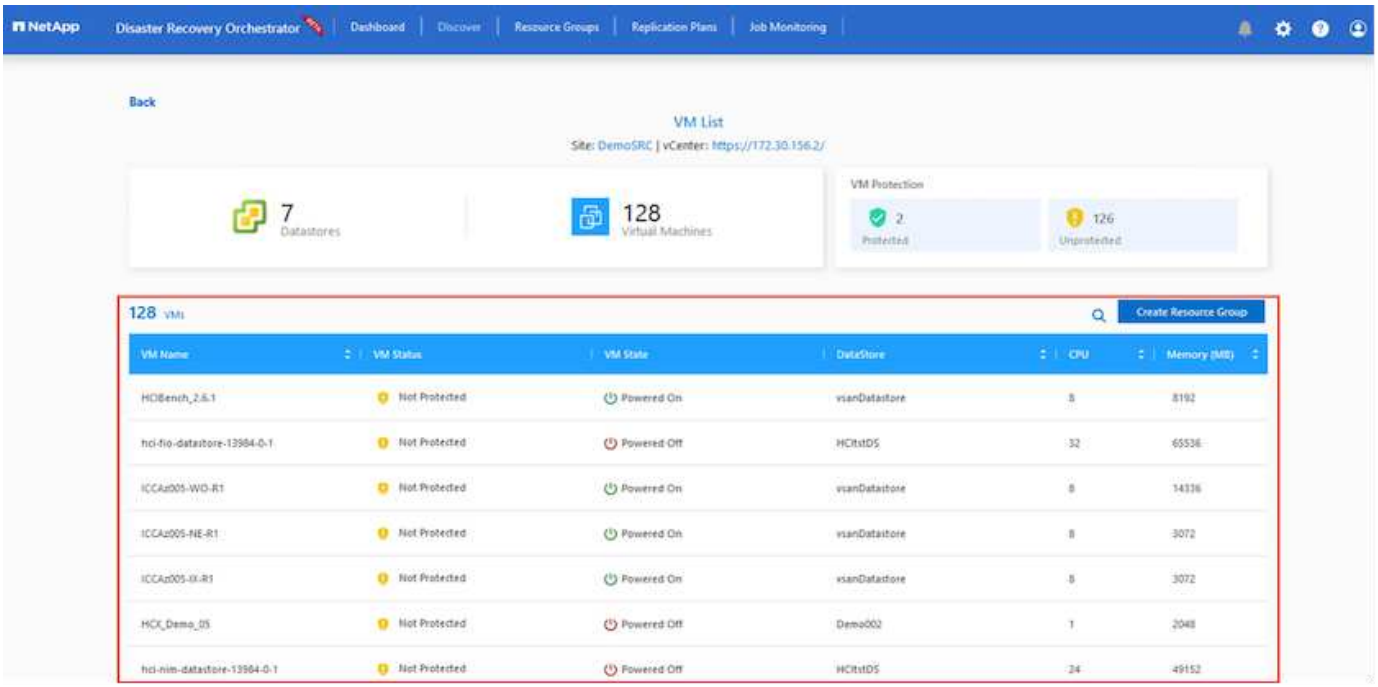
11. Update the vCenter details. To do this, select the credentials, Azure region, and resource group from the dropdown for the primary AVS SDDC.
12. DRO lists all the available SDDCs within the region. Select the designated private cloud URL from the dropdown.
13. Enter the `cloudadmin@vsphere.local` user credentials. This can be accessed from Azure Portal. Follow the steps mentioned in this link. Once done, click **Continue**.



14. Select the Source Storge details (ANF) by selecting the Azure Resource group and NetApp account.
15. Click **Create Site**.



Once added, DRO performs automatic discovery and displays the VMs that have corresponding cross- region replicas from the source site to the destination site. DRO automatically detects the networks and segments used by the VMs and populates them.



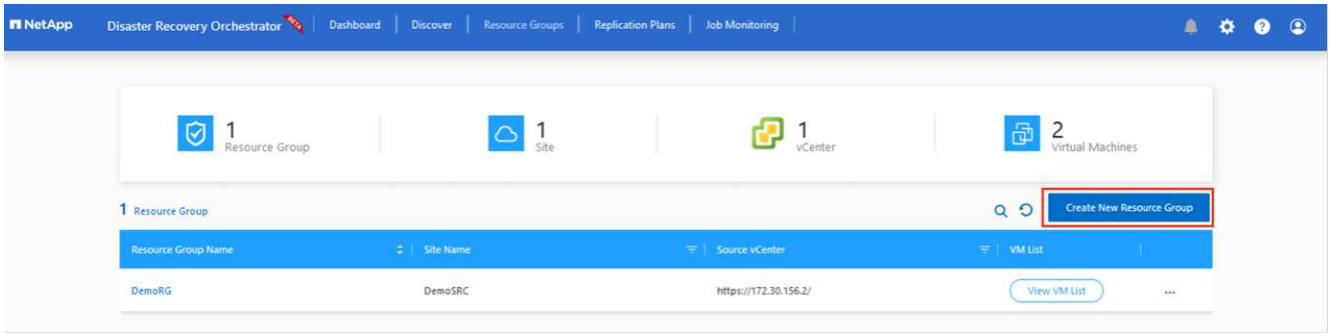
The next step is to group the required VMs into their functional groups as resource groups.

Resource groupings

After the platforms have been added, group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.

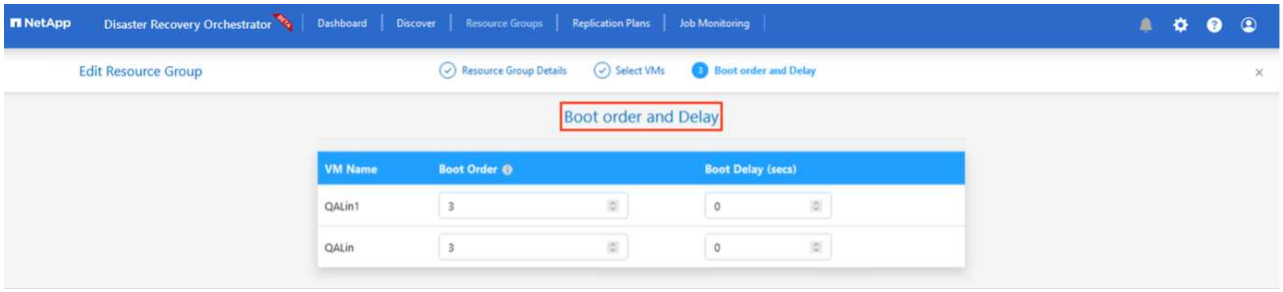
To start creating resource groups, click the **Create New Resource Group** menu item.

1. Access **Resource Grou*ps and click *Create New Resource Group**.

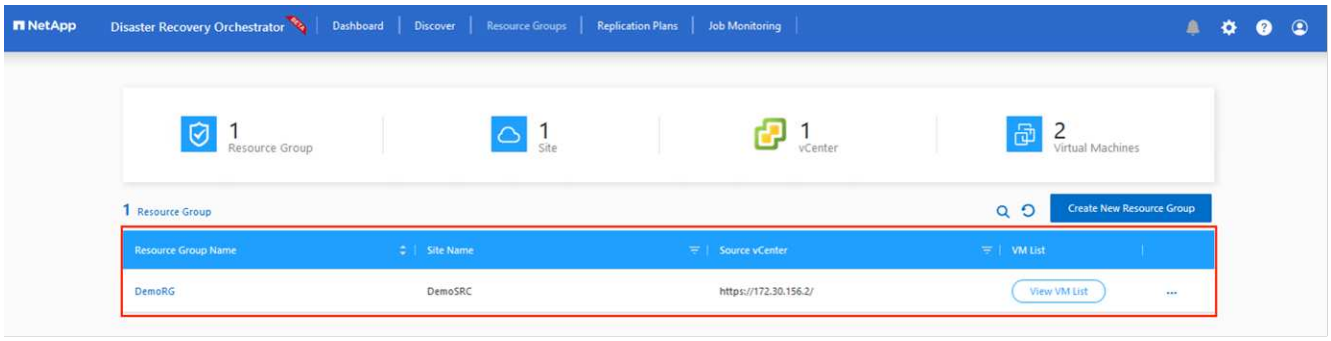


2. Under New Resource Group, select the source site from the dropdown and click **Create**.
3. Provide the resource group details and click **Continue**.
4. Select appropriate VMs using the search option.
5. Select the **Boot Order** and **Boot Delay** (secs) for all the selected VMs. Set the order of the power- on sequence by selecting each virtual machine and setting up the priority for it. The default value for all virtual machines is 3. The options are as follows:
   - The first virtual machine to power on
   - Default

- The last virtual machine to power on



6. Click **Create Resource Group**.

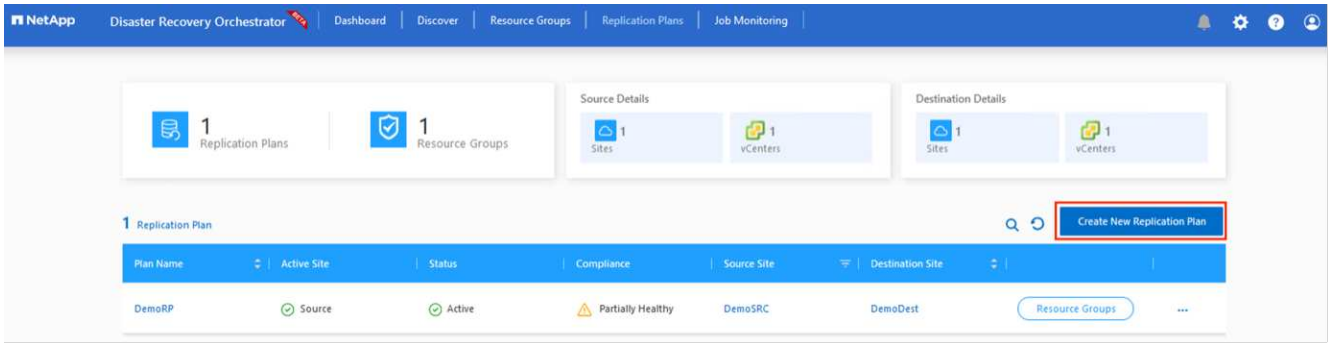


Replication plans

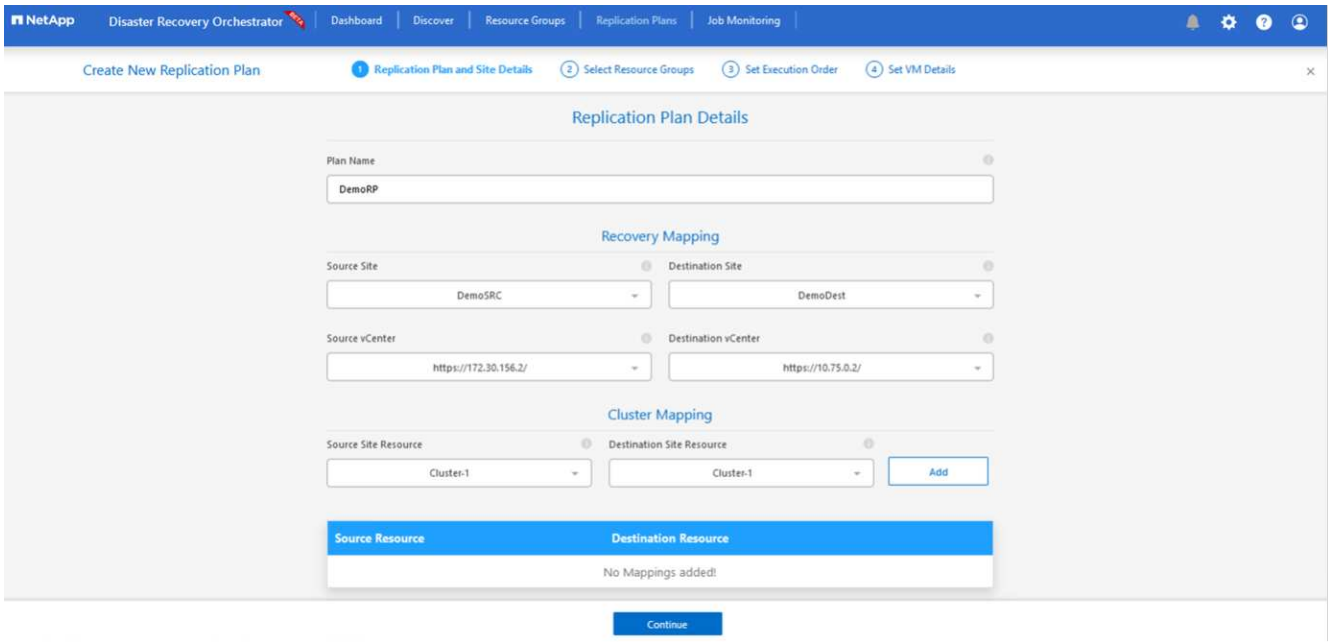
You must have a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down, pick the resource groups to be included in this plan, and also include the grouping of how applications should be restored and powered on (for example, domain controllers, tier-1, tier-2, and so on). Plans are often called blueprints as well. To define the recovery plan, navigate to the Replication Plan tab, and click **New Replication Plan**.

To start creating a replication plan, complete the following steps:

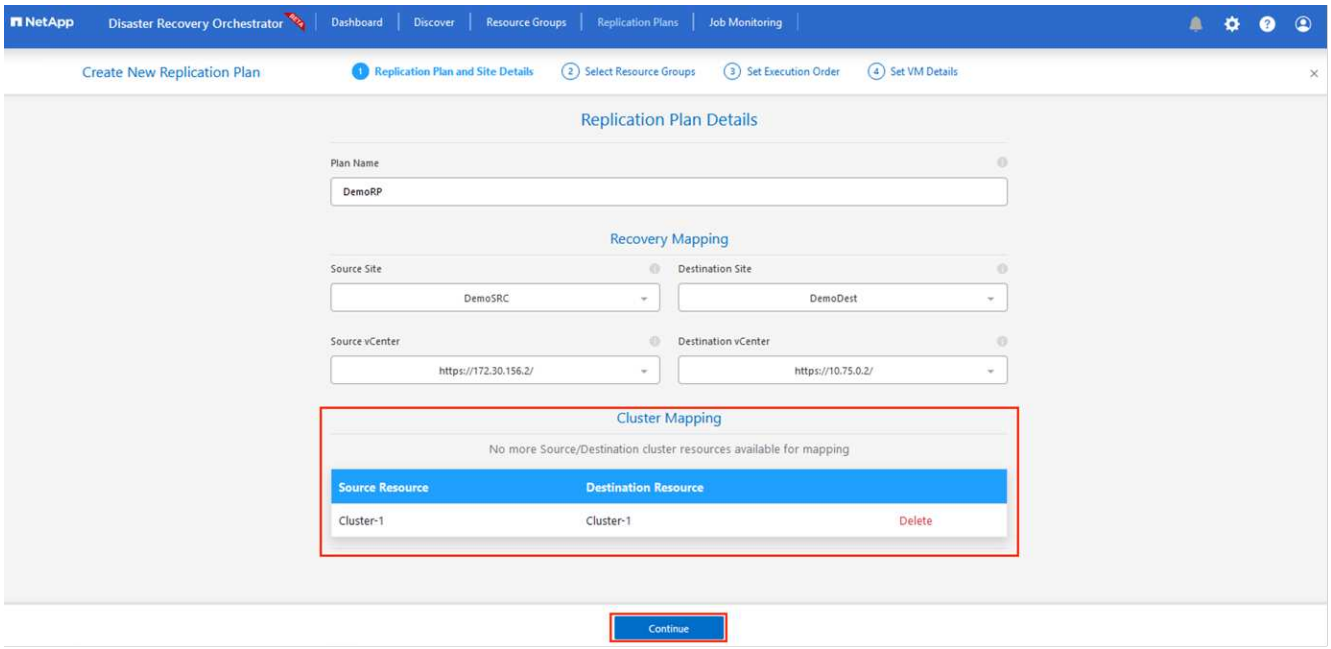
1. Navigate to **Replication Plans** and click **Create New Replication Plan**.



2. On the **New Replication Plan**, provide a name for the plan and add recovery mappings by selecting the Source Site, associated vCenter, Destination Site, and associated vCenter.

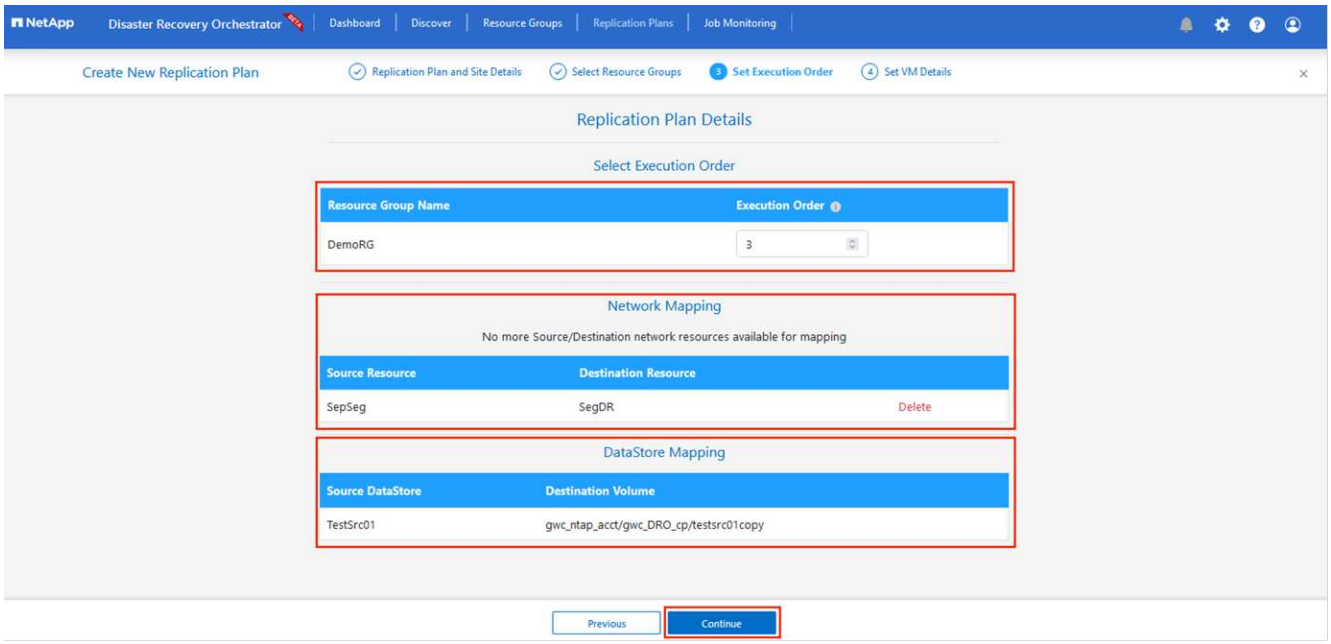


3. After recovery mapping is complete, select the **Cluster Mapping**.

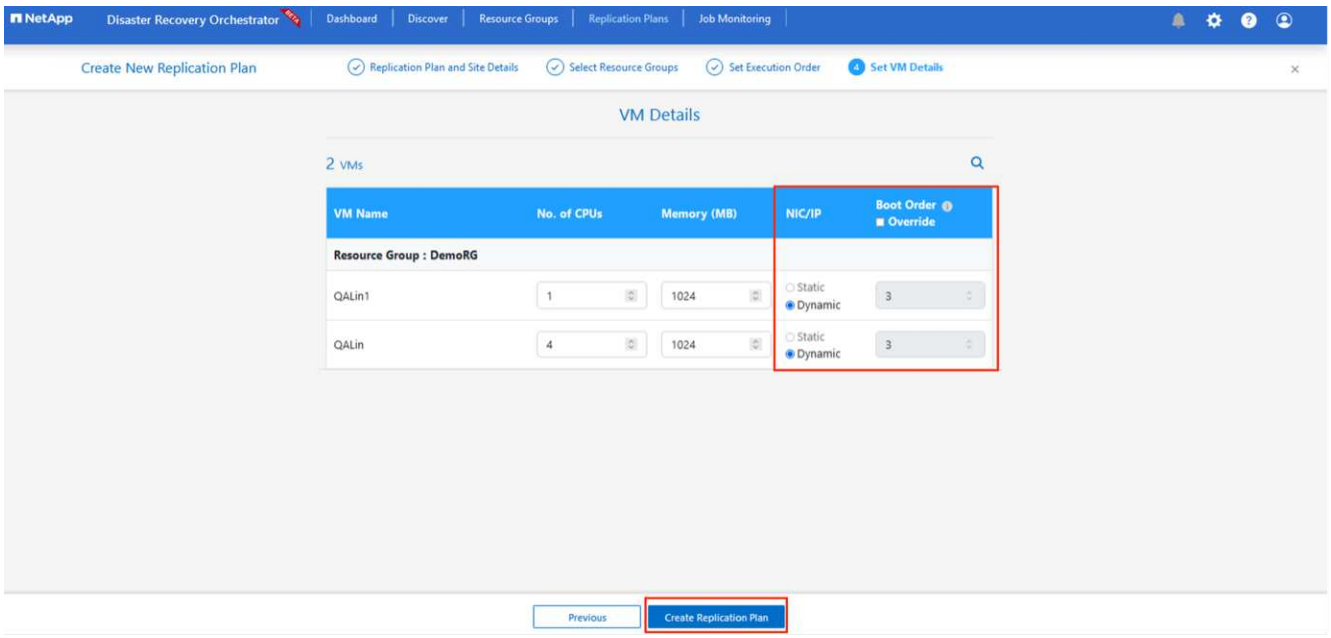


4. Select **Resource Group Details** and click **Continue**.
5. Set the execution order for the resource group. This option enables you to select the sequence of operations when multiple resource groups exist.
6. Once done, set network mapping to the appropriate segment. The segments should already be provisioned on the secondary AVS cluster, and, to map the VMs to those, select the appropriate segment.
7. Datastore mappings are automatically selected based on the selection of VMs.

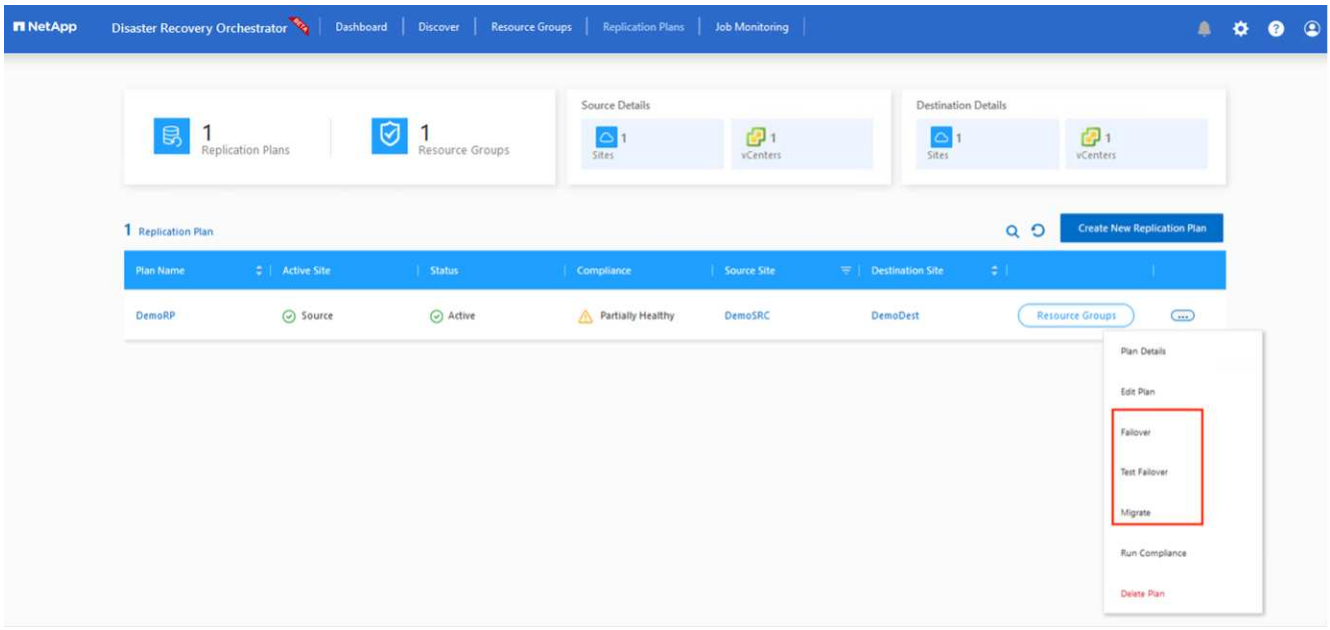
> Cross- region replication (CRR) is at the volume level. Therefore, all VMs residing on the respective volume are replicated to the CRR destination. Make sure to select all VMs that are part of the datastore, because only virtual machines that are part of the replication plan are processed.



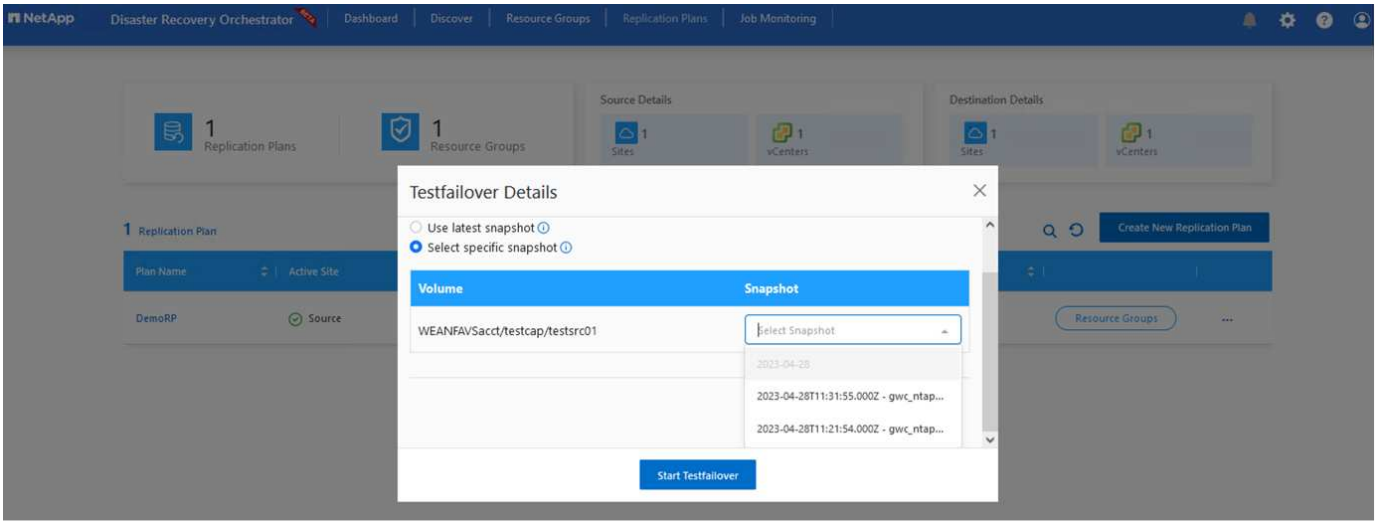
8. Under VM details, you can optionally resize the VMs CPU and RAM parameters. This can be very helpful when you are recovering large environments to smaller target clusters or when you are conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, modify the boot order and boot delay (secs) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if any changes are required from what you selected during resource- group boot-order selection. By default, the boot order selected during resource- group selection is used, however any modifications can be performed at this stage.



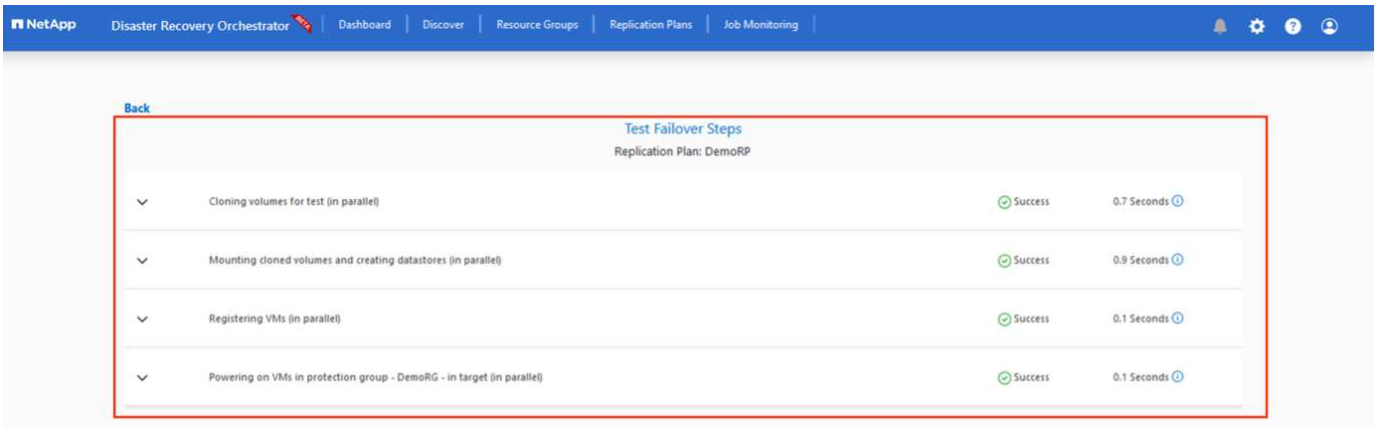
9. Click **Create Replication Plan**.After the replication plan is created, you can exercise the failover, test failover, or migrate options depending on your requirements.



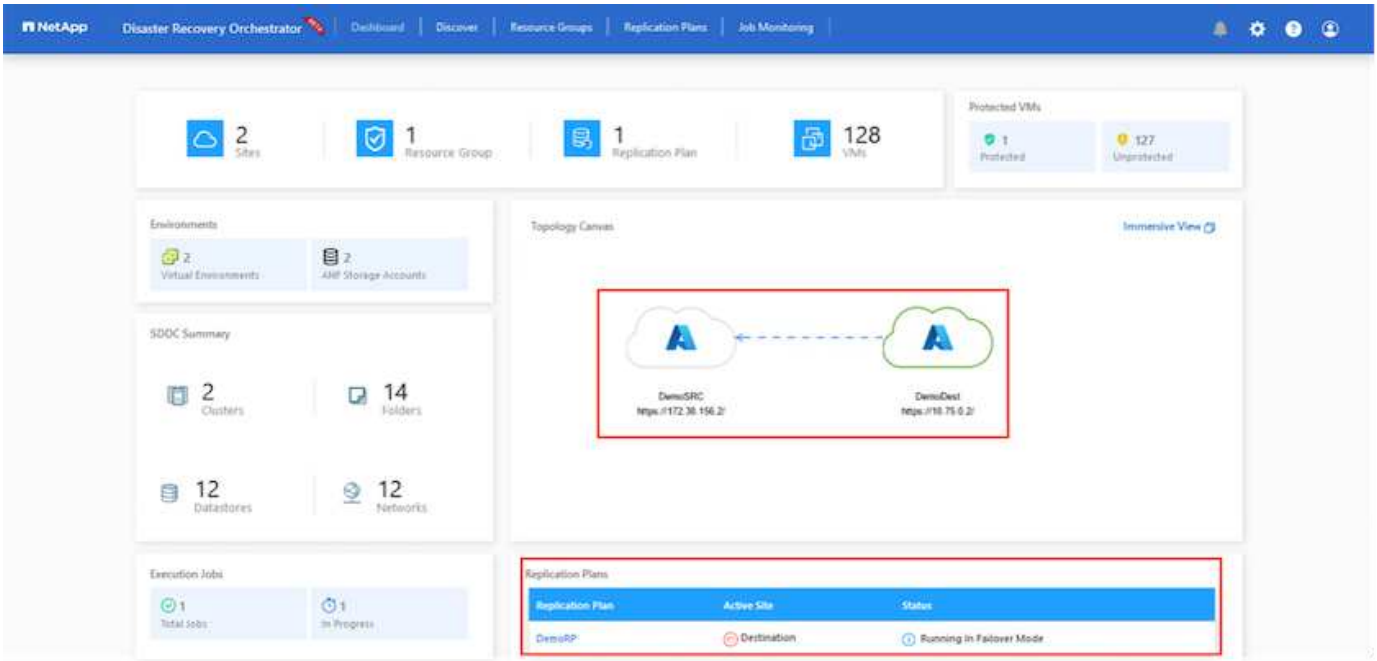
During the failover and test failover options, the most recent snapshot is used, or a specific snapshot can be selected from a point-in-time snapshot. The point-in-time option can be very beneficial if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available time points.



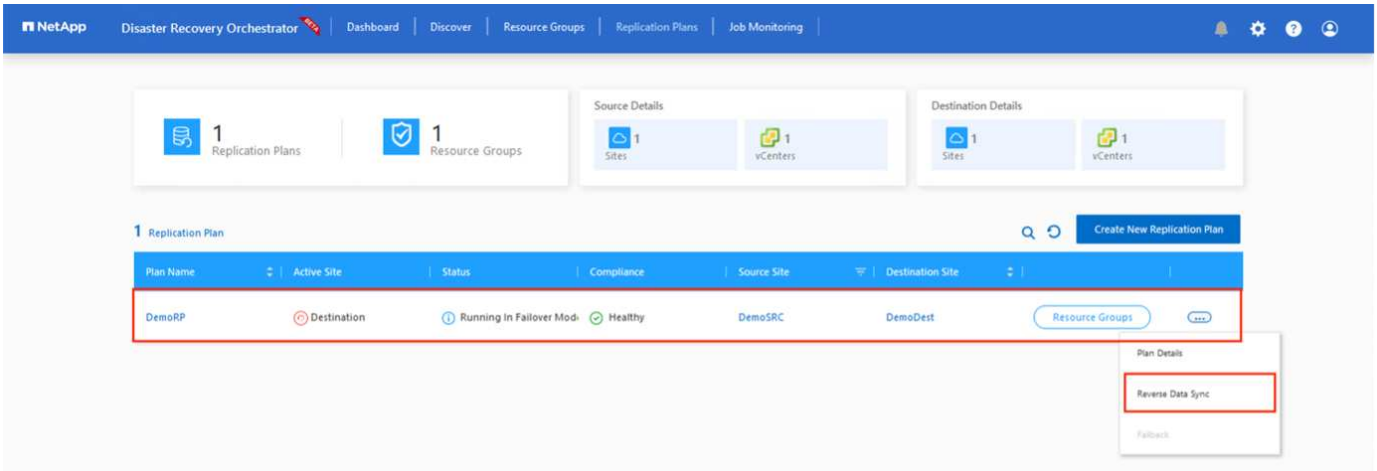
To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test Failover**. You can monitor the replication plan in the task menu.



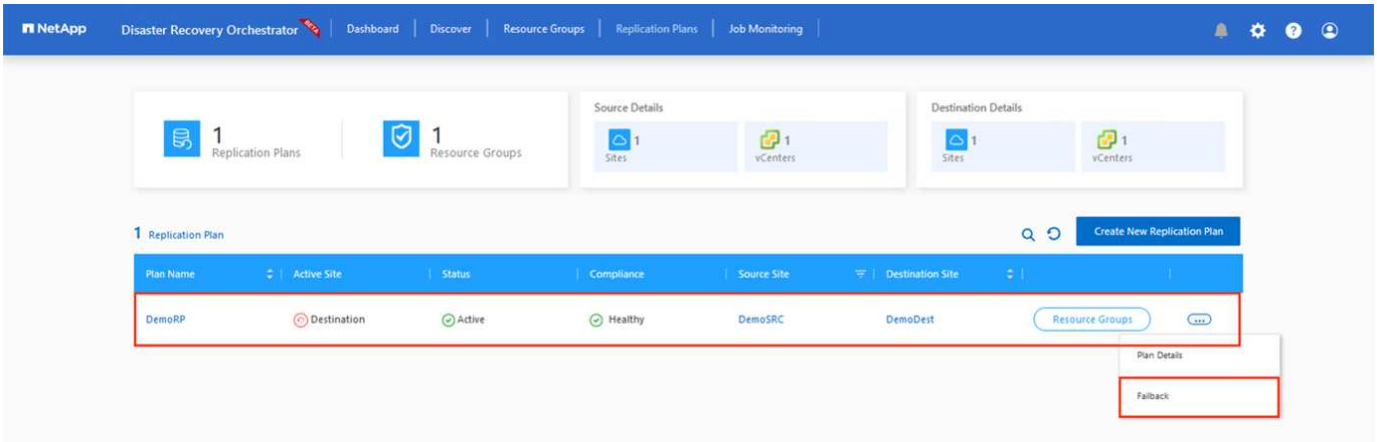
After failover is triggered, the recovered items can be seen in the secondary site AVS SDDC vCenter (VMs, networks, and datastores). By default, the VMs are recovered to Workload folder.

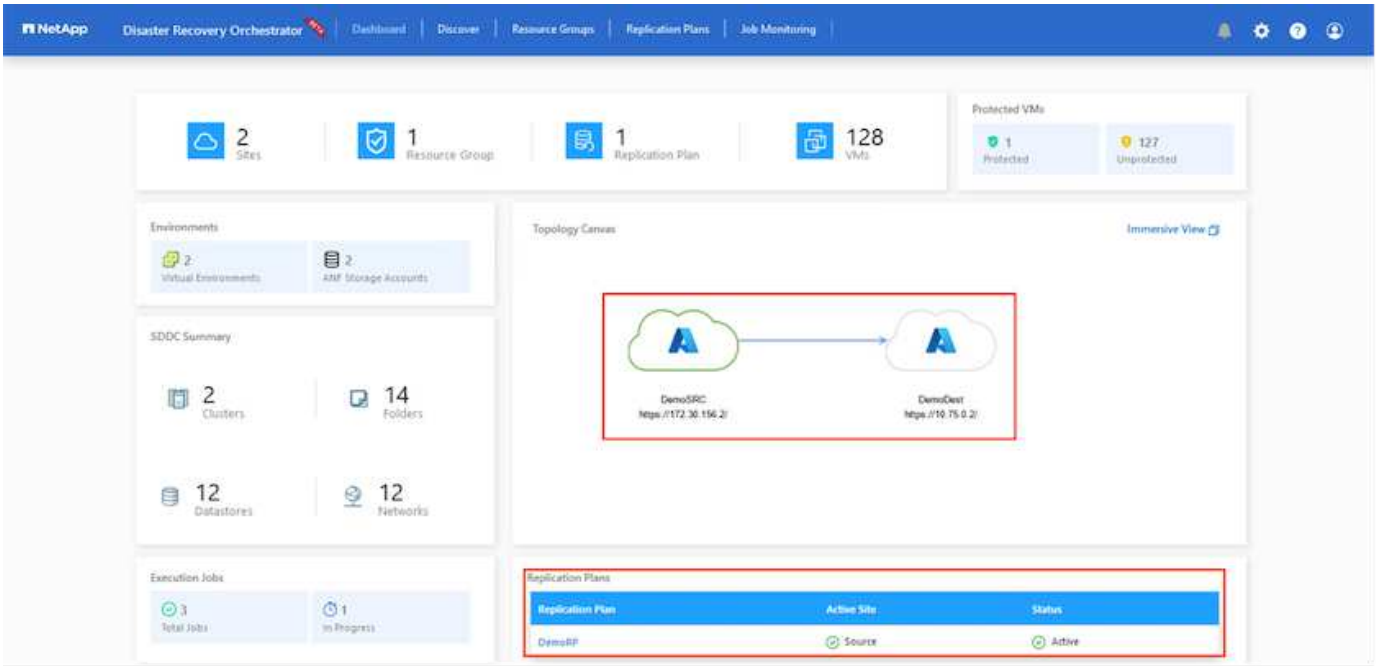


Failback can be triggered at the replication plan level. In case of test failover, the tear down option can be used to roll back the changes and remove the newly created volume. Failbacks related to failover are a two- step process. Select the replication plan and select **Reverse Data sync**.



After this step is complete, trigger failback to move back to the primary AVS site.





From the Azure portal, we can see that the replication health has been broken off for the appropriate volumes that were mapped to the secondary site AVS SDDC as read/write volumes. During test failover, DRO does not map the destination or replica volume. Instead, it creates a new volume of the required cross- region replication snapshot and exposes the volume as a datastore, which consumes additional physical capacity from the capacity pool and ensures that the source volume is not modified. Notably, replication jobs can continue during DR tests or triage workflows. Additionally, this process makes sure that the recovery can be cleaned up without the risk of the replica being destroyed if errors occur or corrupted data is recovered.

Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be difficult for IT organizations to pinpoint what the safe point of return is, and, once that's determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (for example, from sleeping malware or through vulnerable applications).

DRO addresses these concerns by allowing organizations to recover from any available point-in-time. Workloads are then recovered to functional and yet isolated networks, so that applications can function and communicate with each other but are not exposed to any north- south traffic. This process gives security teams

a safe place to conduct forensics and identify any hidden or sleeping malware.

Conclusion

The Azure NetApp Files and Azure VMware disaster recovery solution provide you with the following benefits:

- Leverage efficient and resilient Azure NetApp Files cross- region replication.
- Recover to any available point-in-time with snapshot retention.
- Fully automate all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery leverages the "Create new volumes from the most recent snapshots" process, which doesn't manipulate the replicated volume.
- Avoid any risk of data corruption on the volumes or snapshots.
- Avoid replication interruptions during DR test workflows.
- Leverage DR data and cloud compute resources for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- CPU and RAM optimization can help lower cloud costs by allowing recovery to smaller compute clusters.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Create volume replication for Azure NetApp Files

  https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-create-peering

- Cross-region replication of Azure NetApp Files volumes

  https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-introduction#service-level-objectives

- Azure VMware Solution

  https://learn.microsoft.com/en-us/azure/azure-vmware/introduction

- Deploy and configure the Virtualization Environment on Azure

  Setup AVS on Azure

- Deploy and configure Azure VMware Solution

  https://learn.microsoft.com/en-us/azure/azure-vmware/deploy-azure-vmware-solution?tabs=azure-portal

Using Veeam Replication and Azure NetApp Files datastore for disaster recovery to Azure VMware Solution

Azure NetApp Files (ANF) datastores decouples storage from compute and unlocks the flexibility needed for any organisation to take their workloads to the cloud. It provides customers with flexible, high-performance storage infrastructure that scales independently of compute resources. Azure NetApp Files datastore's simplifies and

optimizes the deployment alongside Azure VMware Solution (AVS) as a disaster recovery site for on premises VMWare environments.

Overview

Azure NetApp Files (ANF) volume based NFS datastores can be used to replicate data from on-premises using any validated third-party solution that provides VM replication capability. By adding Azure NetApp Files datastores, it will enable cost optimised deployment vs building an Azure VMware Solution SDDC with enormous amount of ESXi hosts to accommodate the storage. This approach is called a "Pilot Light Cluster". A pilot light cluster is a minimal AVS host configuration (3 x AVS nodes) along with Azure NetApp Files Datastore capacity.

The objective is to maintain a low-cost infrastructure with all the core components to handle a failover. A pilot light cluster can scale out and provision more AVS hosts if a failover does occur. And once the failover is complete and normal operations are restored, the pilot light cluster can scale back down to low-cost mode of operations.

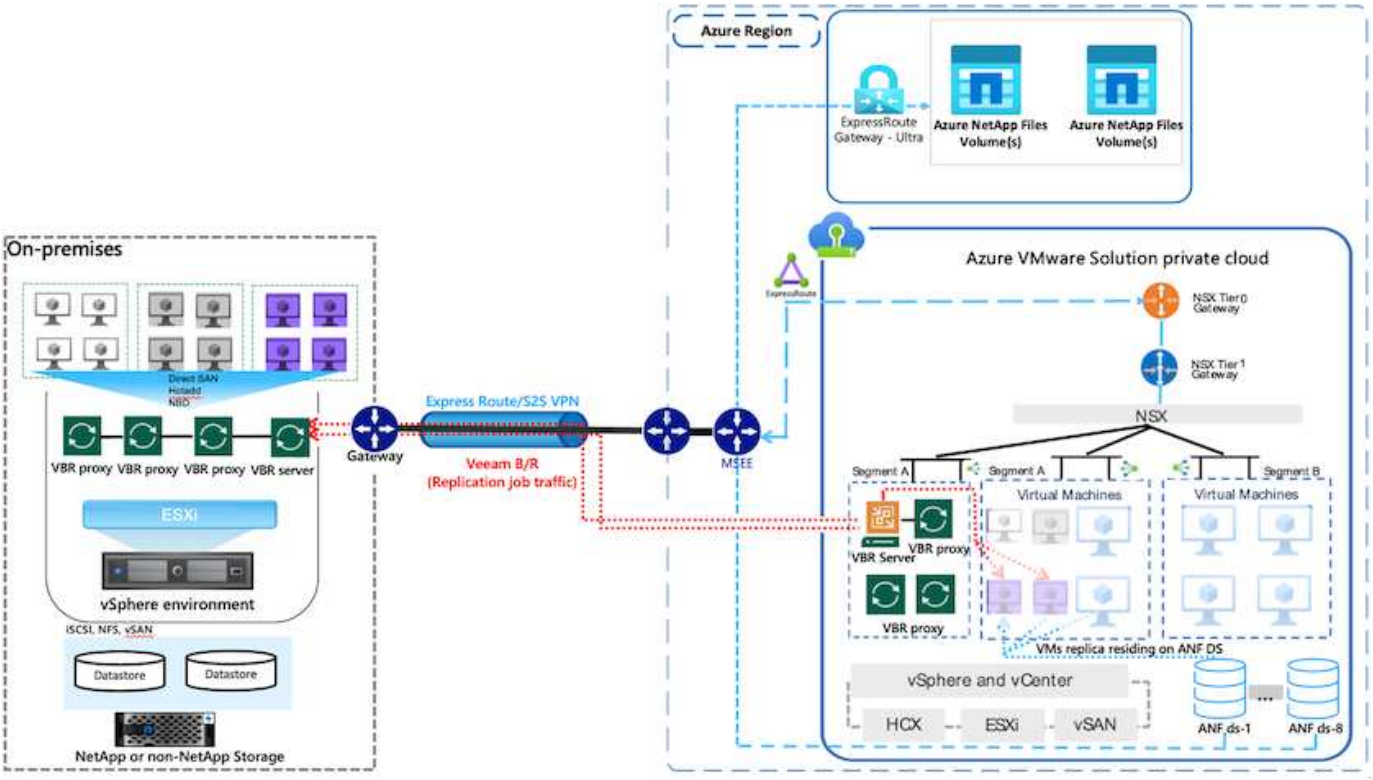
Purposes of this document

This article describes how to use Azure NetApp Files datastore with Veeam Backup and replication to set up disaster recovery for on-premises VMware VMs to (AVS) using the Veeam VM replication software functionality.

Veeam Backup & Replication is a backup and replication application for virtual environments. When virtual machines are replicated, Veeam Backup & Replication is replicated from on AVS, the software will create an exact copy of the VMs in the native VMware vSphere format on the target AVS SDDC cluster. Veeam Backup & Replication will keep the copy synchronized with the original VM. Replication provides the best recovery time objective (RTO) as there is a mounted copy of a VM at the DR site in a ready-to-start state.

This replication mechanism ensures that the workloads can quickly start in a AVS SDDC in the case of a disaster event. The Veeam Backup & Replication software also optimizes traffic transmission for replication over WAN and slow connections. In addition, it also filters out duplicate data blocks, zero data blocks, swap files, and "excluded VM guest OS files". The software will also compress the replica traffic. To prevent replication jobs from consuming the entire network bandwidth, WAN accelerators and network throttling rules can be utilized.

The replication process in Veeam Backup & Replication is job driven which means replication is performed by configuring replication jobs. In the case of a disaster event, failover can be triggered to recover the VMs by failing over to its replica copy. When failover is performed, a replicated VM takes over the role of the original VM. Failover can be performed to the latest state of a replica or to any of its good known restore points. This enables ransomware recovery or isolated testing as needed. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.



Solution Deployment

High level steps

1. Veeam Backup and Replication software is running in an on-premises environment with appropriate network connectivity.
2. Deploy Azure VMware Solution (AVS) private cloud and attach Azure NetApp Files datastores to Azure VMware Solution hosts.

   A pilot-light environment set up with a minimal configuration can be used for DR purposes. VMs will fail over to this cluster in the event of an incident, and additional nodes can be added).

3. Set up replication job to create VM replicas using Veeam Backup and Replication.
4. Create failover plan and perform failover.
5. Switch back to production VMs once the disaster event is complete and primary site is Up.

Pre-requisites for Veeam VM Replication to AVS and ANF datastores

1. Ensure the Veeam Backup & Replication backup VM is connected to the source as well as the target AVS SDDC clusters.
2. The backup server must be able to resolve short names and connect to source and target vCenters.
3. The target Azure NetApp Files datastore must have enough free space to store VMDKs of replicated VMs.

For additional information, refer to "Considerations and Limitations" covered here.

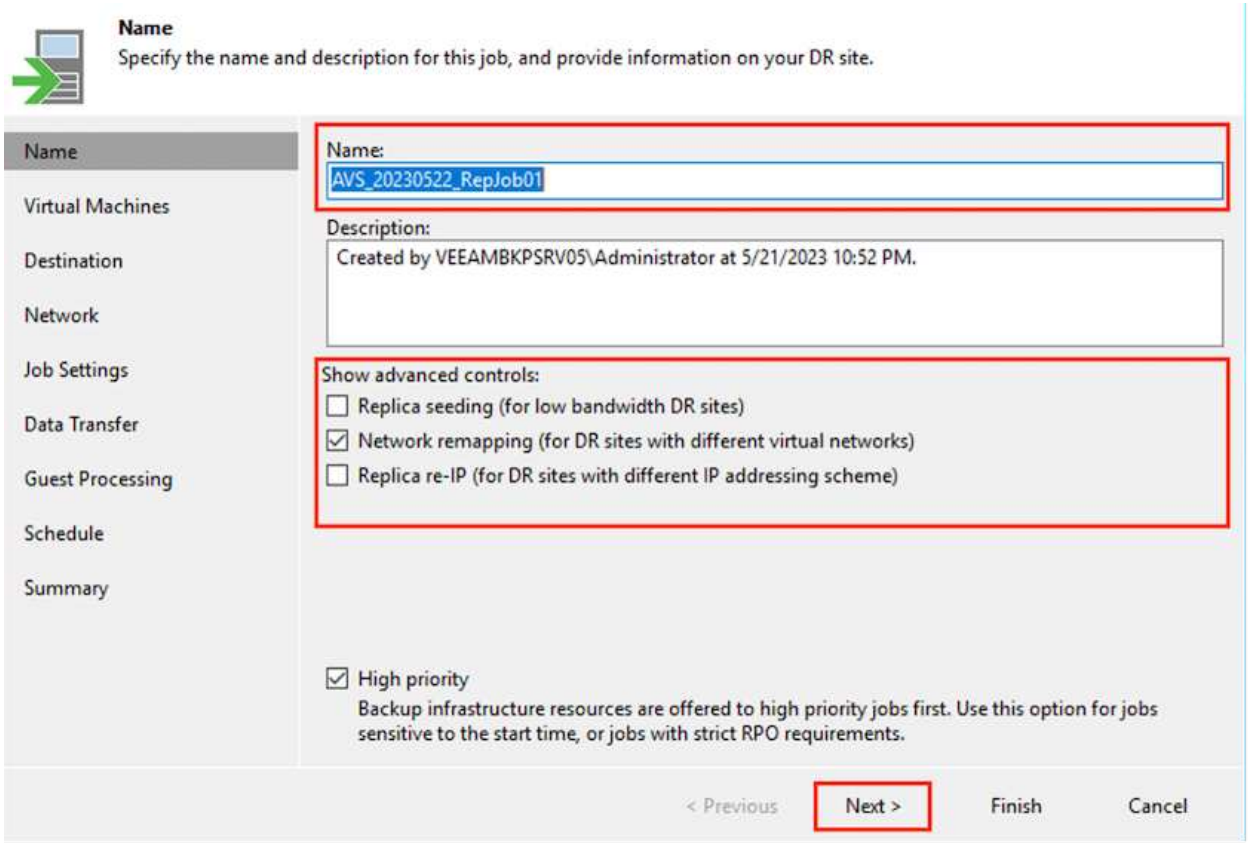
Deployment Details

▼ Step 1: Replicate VMs

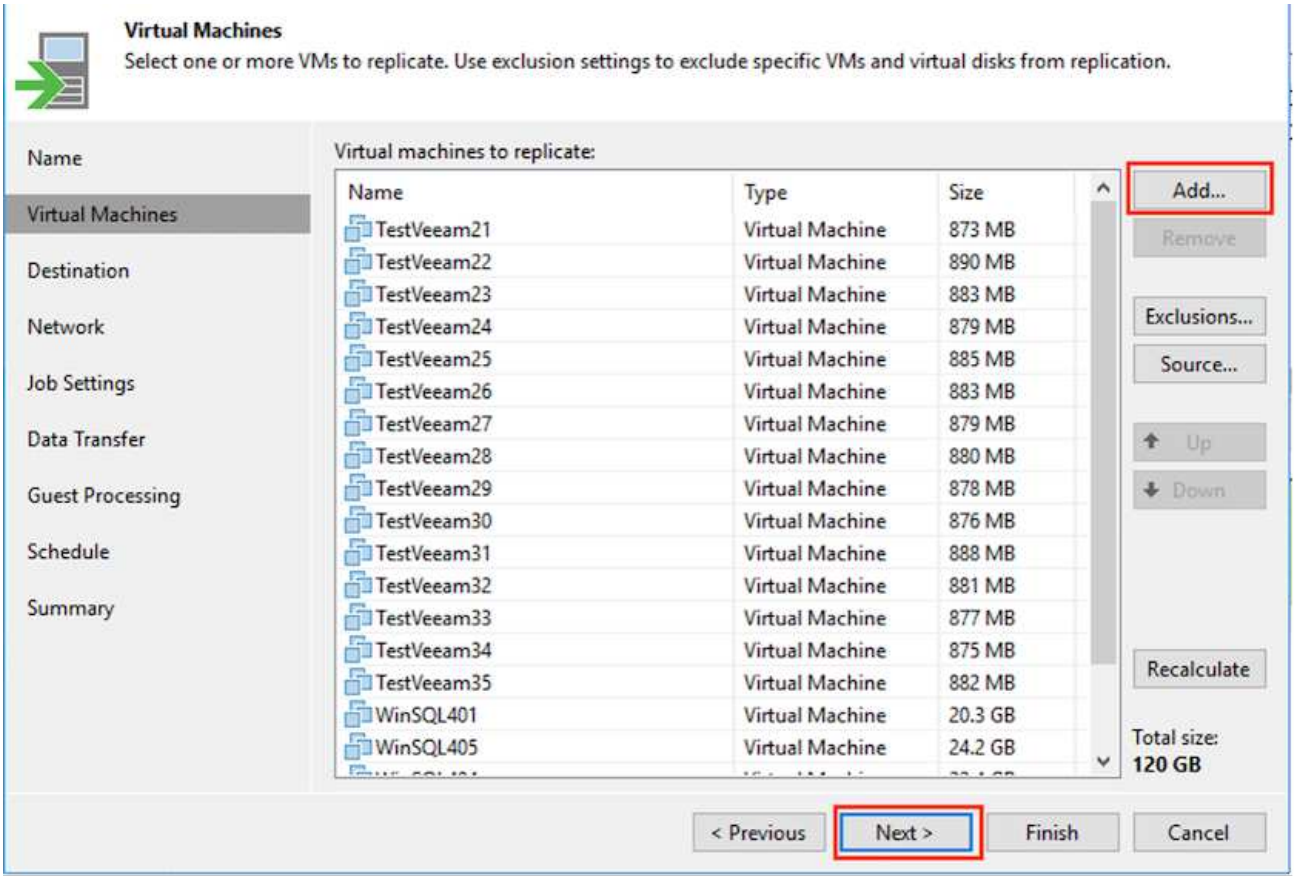
Veeam Backup & Replication leverages VMware vSphere snapshot capabilities/During replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot is the point-in-time copy of a VM that includes virtual disks, system state, configuration and metadata. Veeam Backup & Replication uses the snapshot as a source of data for replication.

To replicate VMs, follow the below steps:

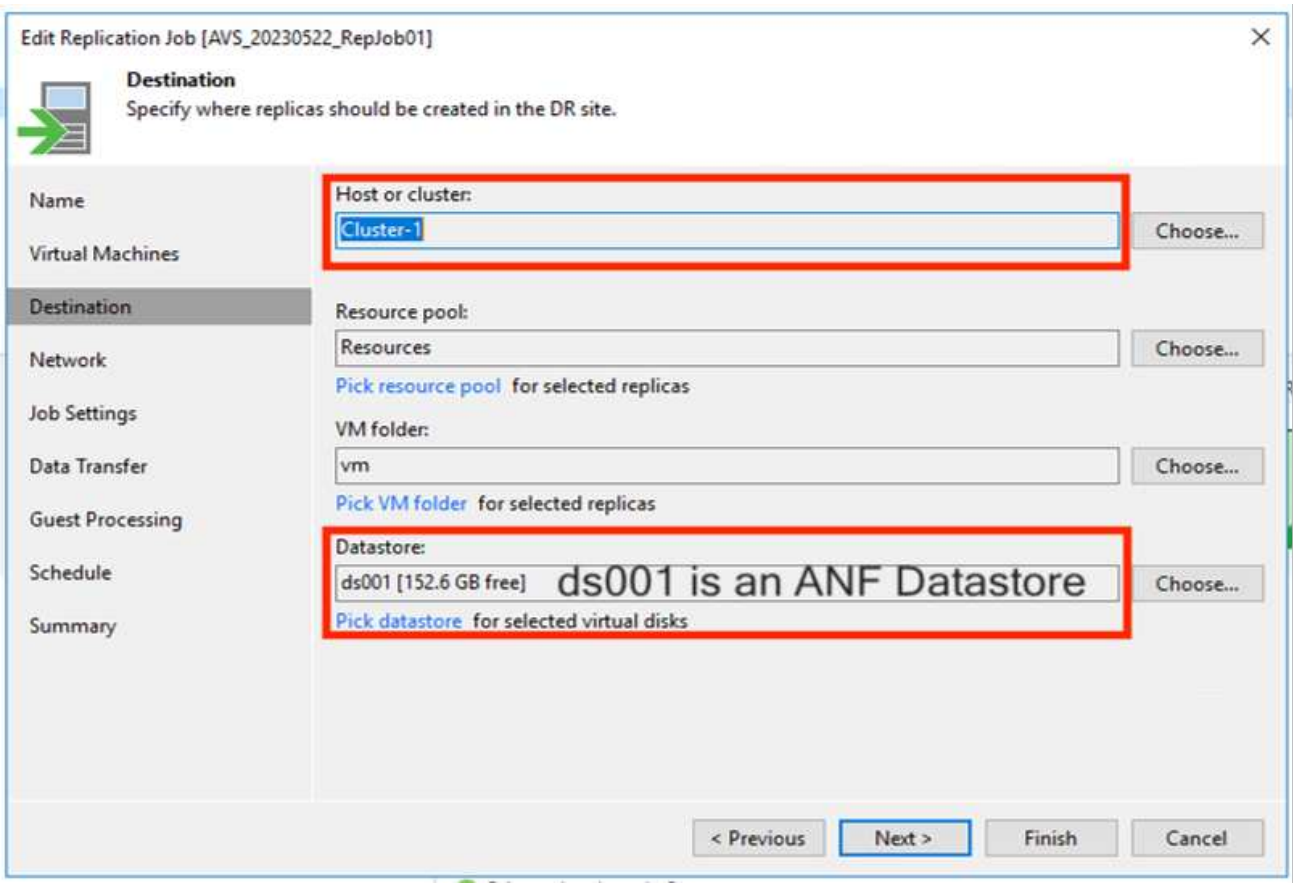
1. Open the Veeam Backup & Replication Console.
2. On the Home view. Right click the jobs node and select Replication Job > Virtual machine.
3. Specify a job name and select the appropriate advanced control checkbox. Click Next.
   - Select the Replica seeding check box if connectivity between on-premises and Azure has restricted bandwidth.
     *Select the Network remapping (for AVS SDDC sites with different networks) check box if segments on Azure VMware Solution SDDC do not match that of on-premises site networks.
   - If the IP addressing scheme in on-premises production site differs from the scheme in the target AVS site, select the Replica re-IP (for DR sites with different IP addressing scheme) check box.



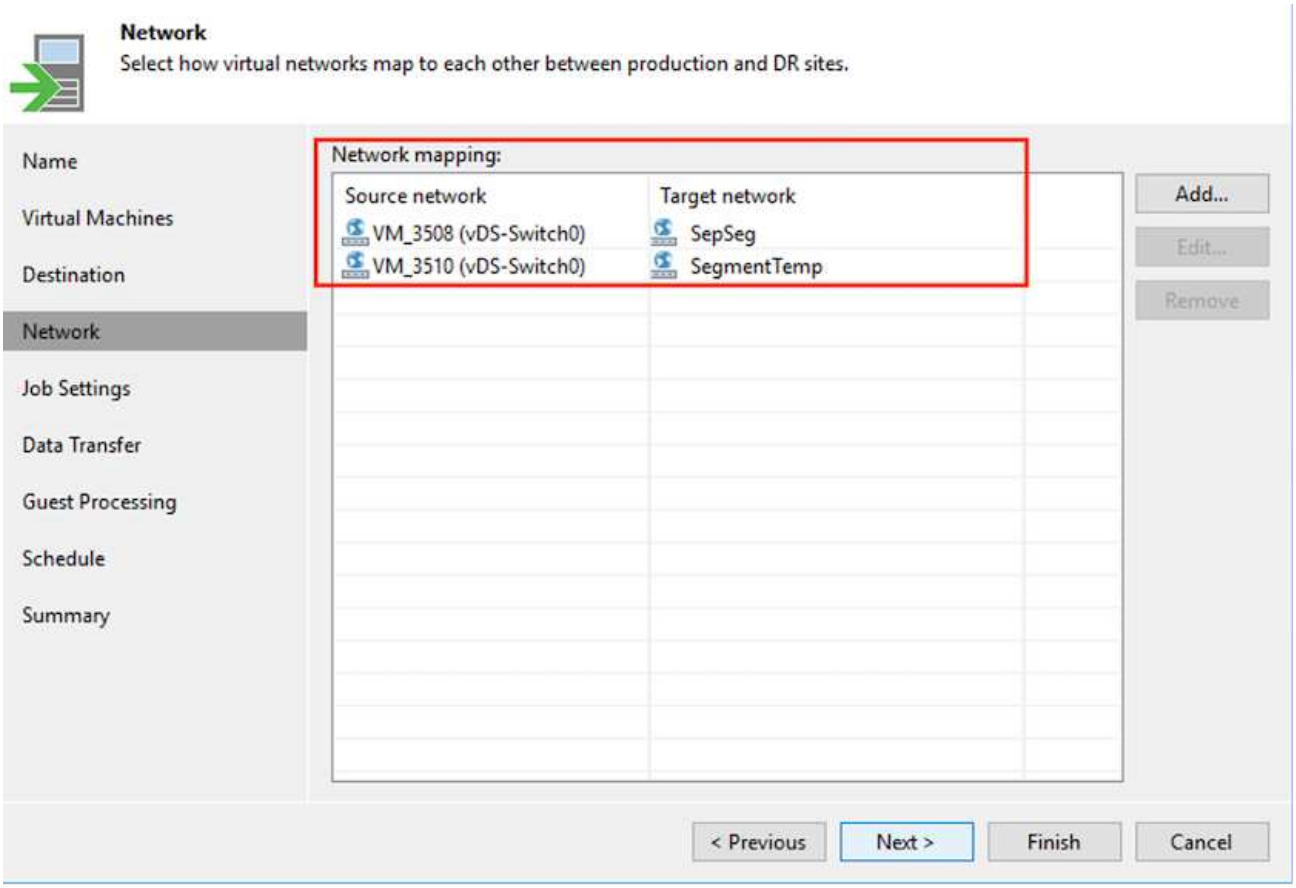
4. Select the VMs to be replicated to Azure NetApp Files datastore attached to a Azure VMware Solution SDDC in the **Virtual** Machines* step. The Virtual machines can be placed on vSAN to fill the available vSAN datastore capacity. In a pilot light cluster, the usable capacity of a 3-node cluster will be limited. The rest of the data can be easily placed on Azure NetApp Files datastores so that the VMs can recovered, and cluster can be expanded to meet the CPU/mem requirements. Click **Add**, then in the **Add Object** window select the necessary VMs or VM containers and click **Add**. Click **Next**.



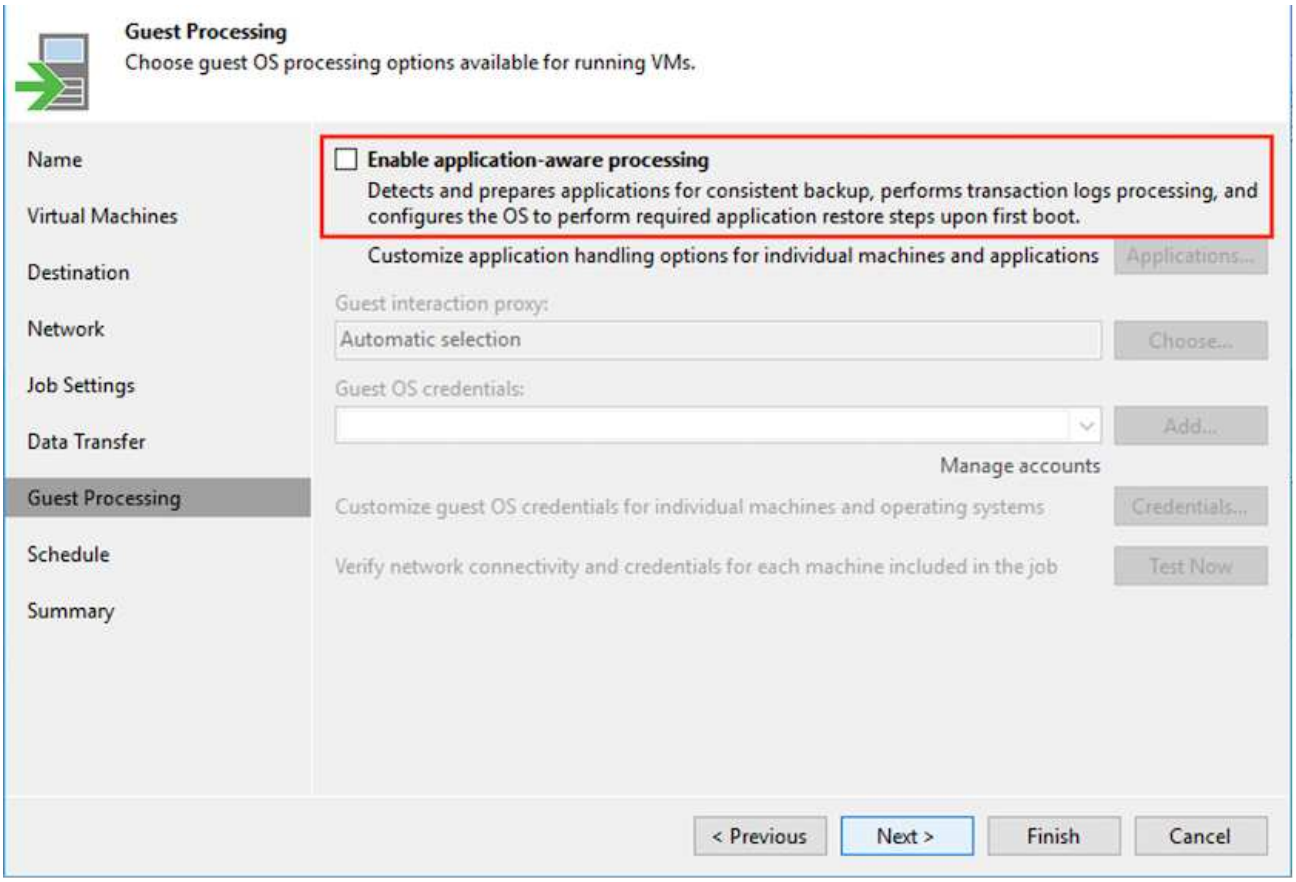
5. After that, select the destination as Azure VMware Solution SDDC cluster / host and the appropriate resource pool, VM folder and FSx ONTAP datastore for VM replicas. Then click **Next**.



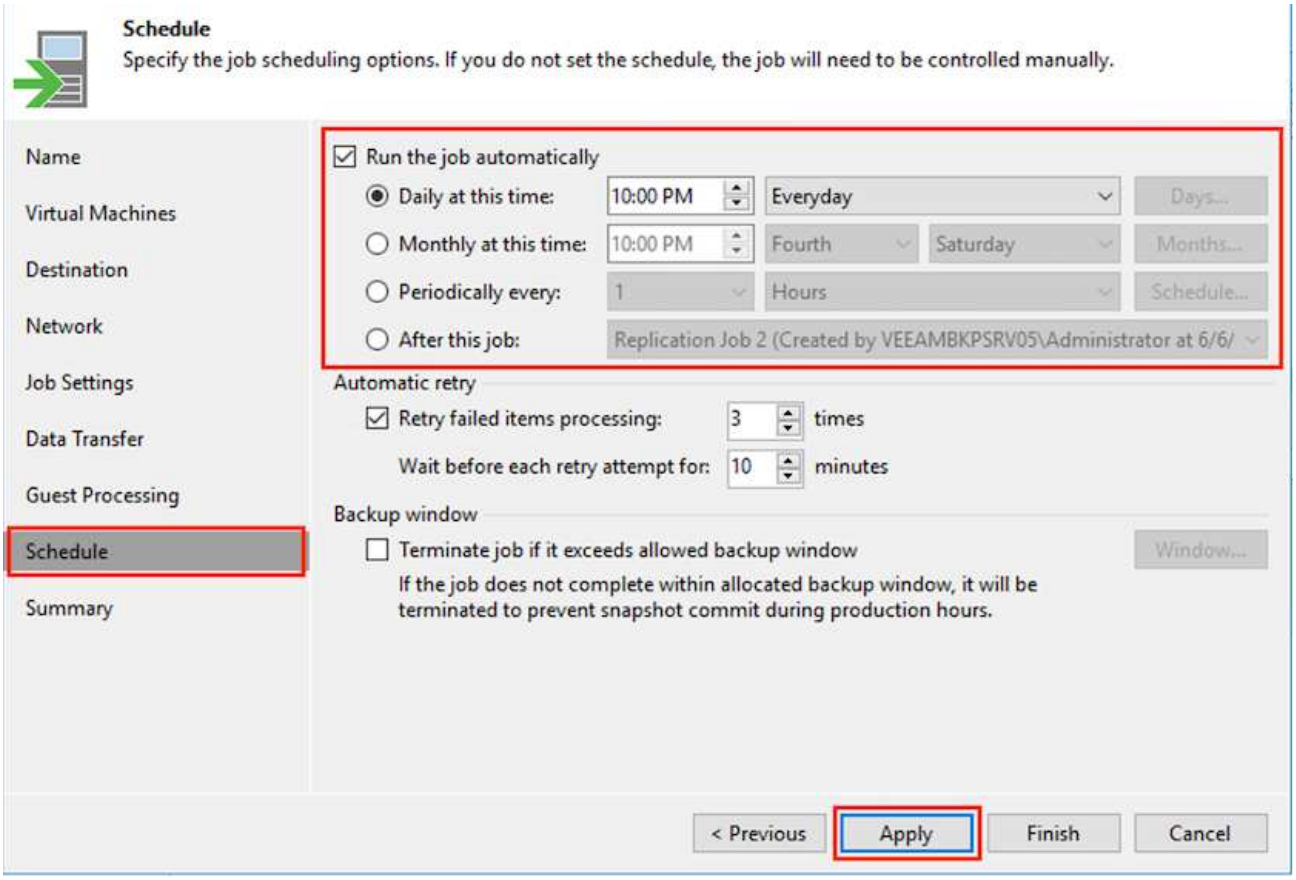
6. In the next step, create the mapping between source and destination virtual network as needed.



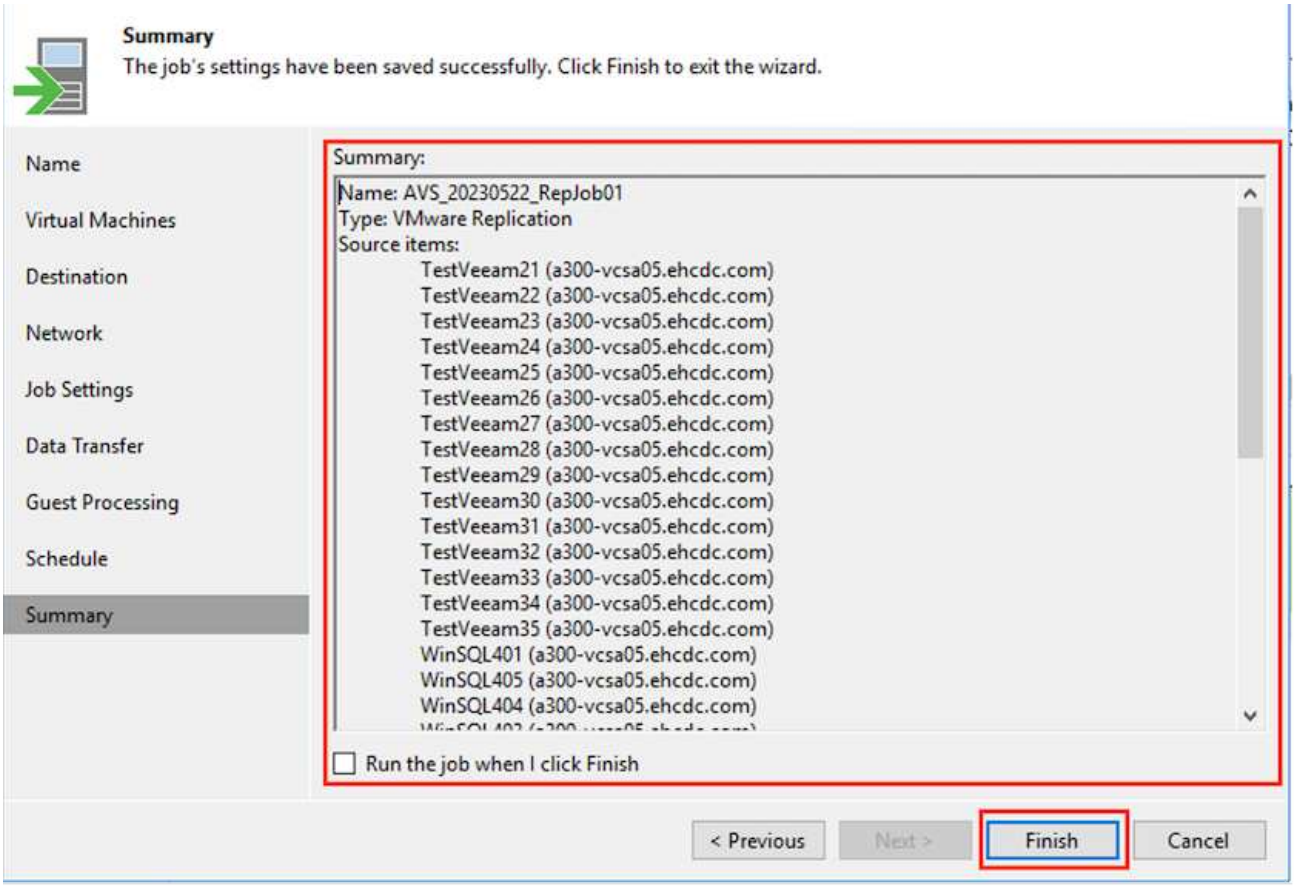
7. In the **Job Settings** step, specify the backup repository that will store metadata for VM replicas, retention policy and so on.
8. Update the **Source** and **Target** proxy servers in the **Data Transfer** step and leave **Automatic** selection (default) and keep **Direct** option selected and click **Next**.
9. At the **Guest Processing** step, select **Enable application-aware processing** option as needed. Click **Next**.



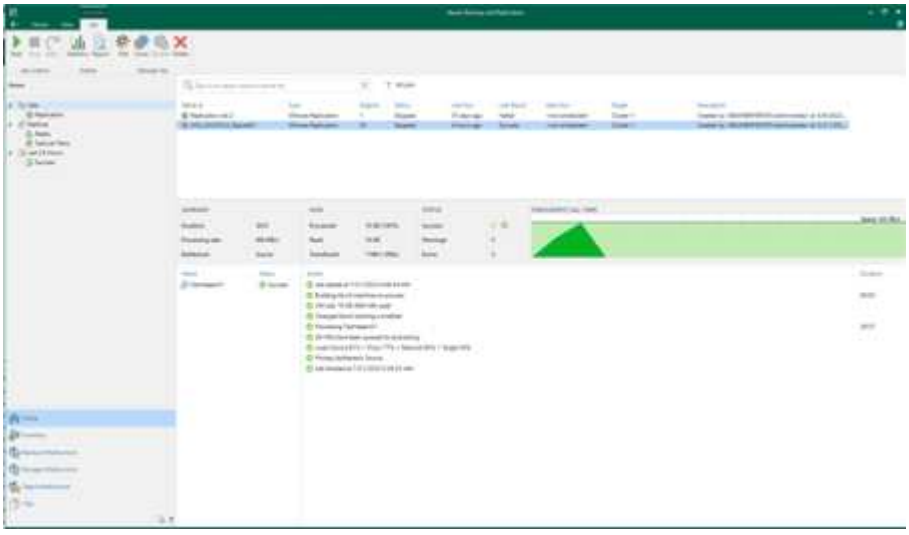
10. Choose the replication schedule to run the replication job to run on a regular basis.



11. At the **Summary** step of the wizard, review details of the replication job. To start the job right after the wizard is closed, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.



Once the replication job starts, the VMs with the suffix specified will be populated on the destination AVS SDDC cluster / host.



For additional information for Veeam replication, refer How Replication Works

▼ **Step 2: Create a failover plan**

When the initial replication or seeding is complete, create the failover plan. Failover plan helps in performing failover for dependent VMs one by one or as a group automatically. Failover plan is the blueprint for the order in which the VMs are processed including the boot delays. The failover plan also helps to ensure that critical dependant VMs are already running.

To create the plan, navigate to the new sub section called **Replicas** and select **Failover Plan**. Choose the appropriate VMs. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas.

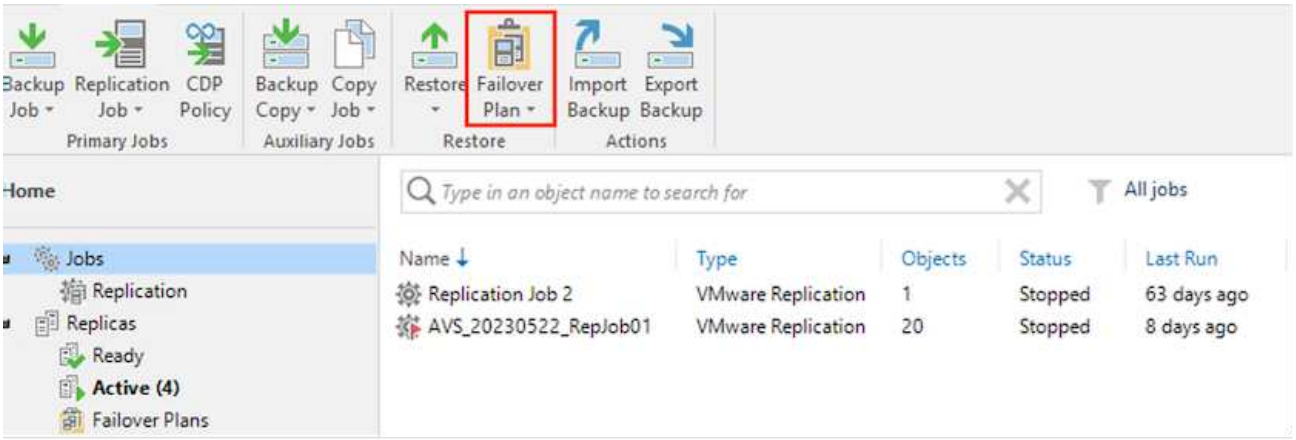
The failover plan can only be added once the initial replication is complete and the VM replicas are in Ready state.

The maximum number of VMs that can be started simultaneously when running a failover plan is 10

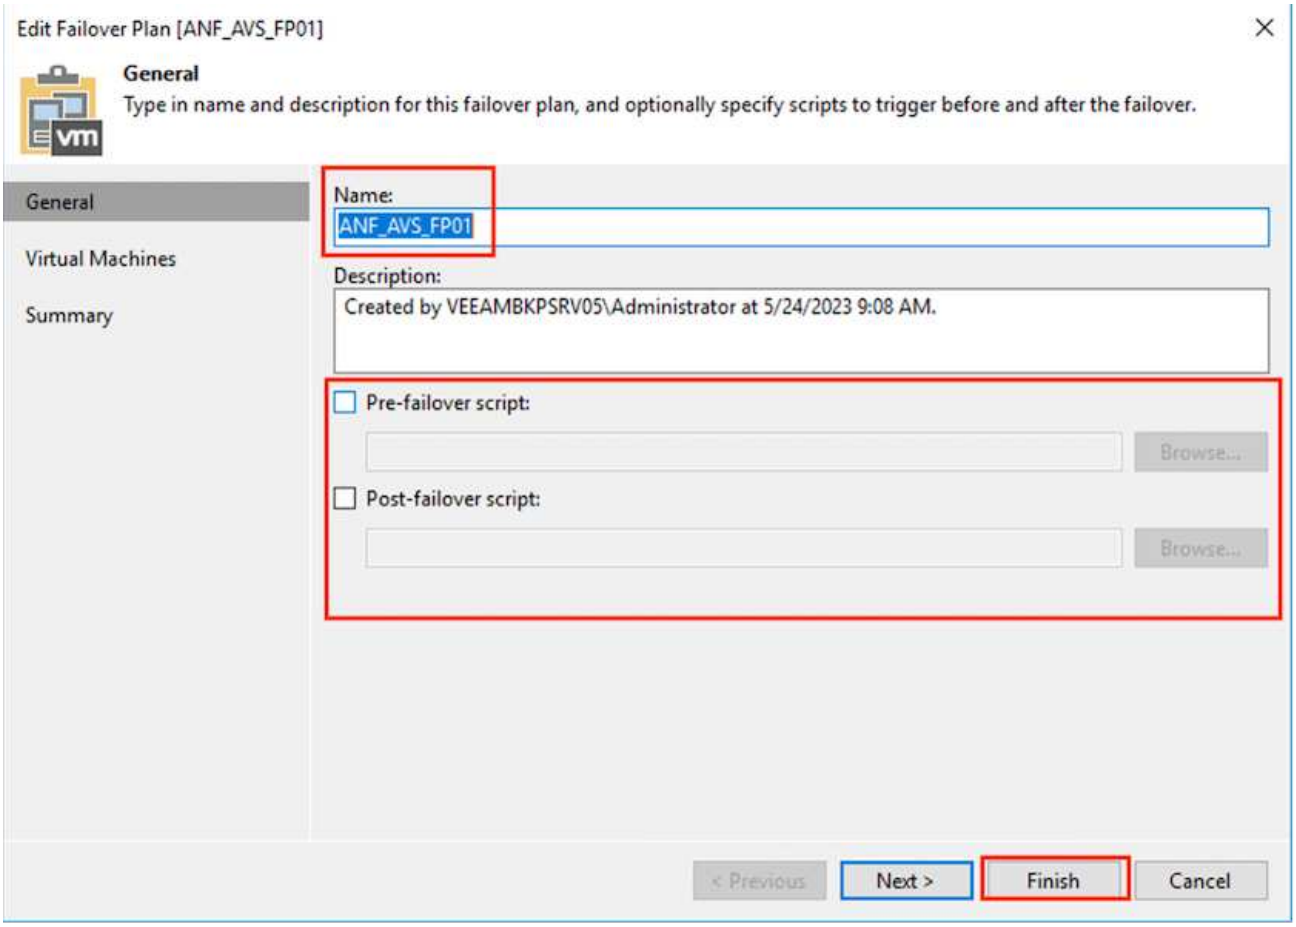
During the failover process, the source VMs will not be powered off

To create the **Failover Plan**, do the following:

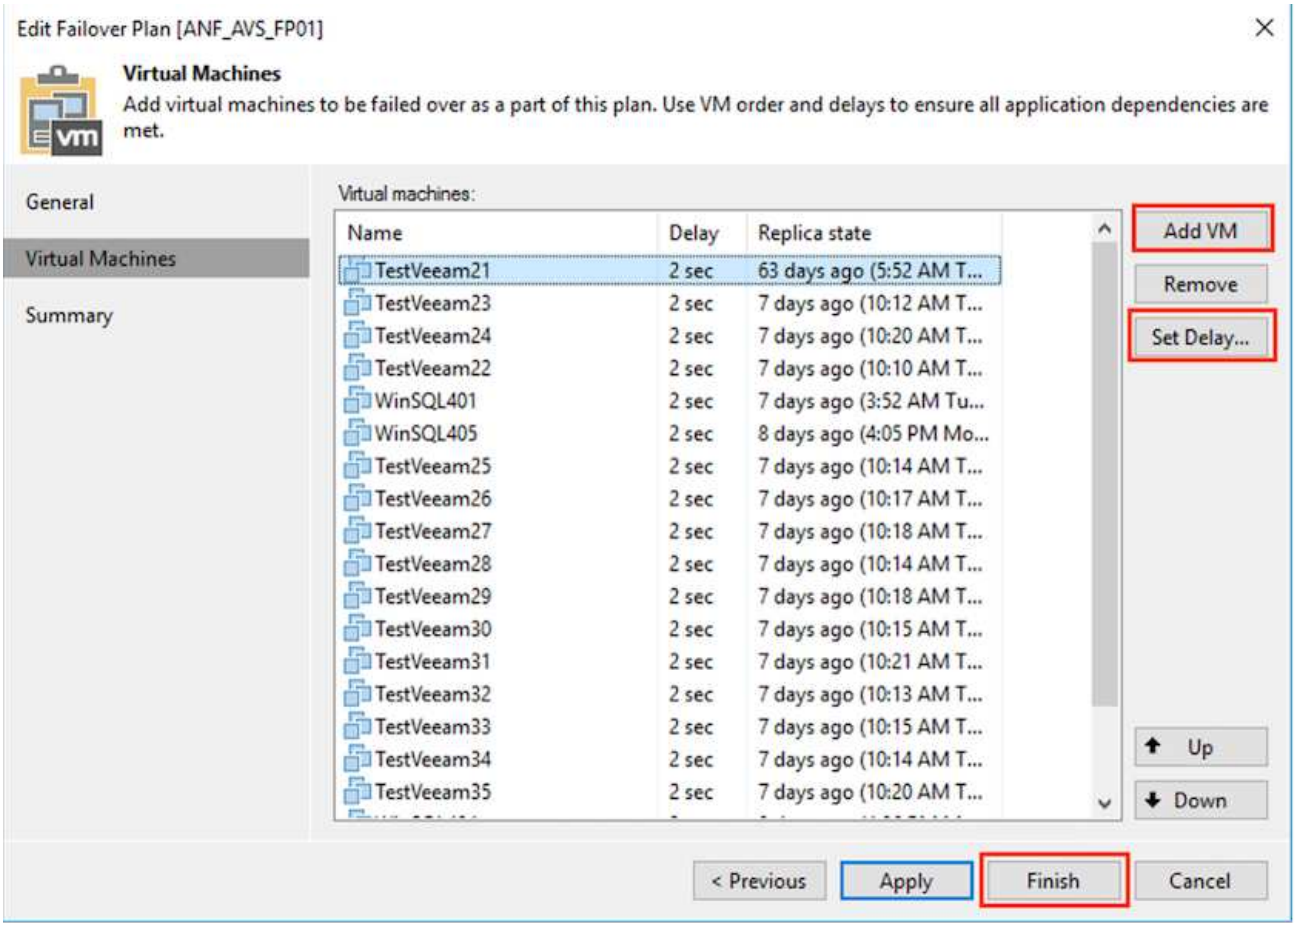
1. On the Home view. Right click the Replicas node and select Failover Plans > Failover Plan > VMware vSphere.



2. Next provide a name and a description to the plan. Pre and Post-failover script can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.



3. Add the VMs to the plan and modify the VM boot order and boot delays to meet the application dependencies.



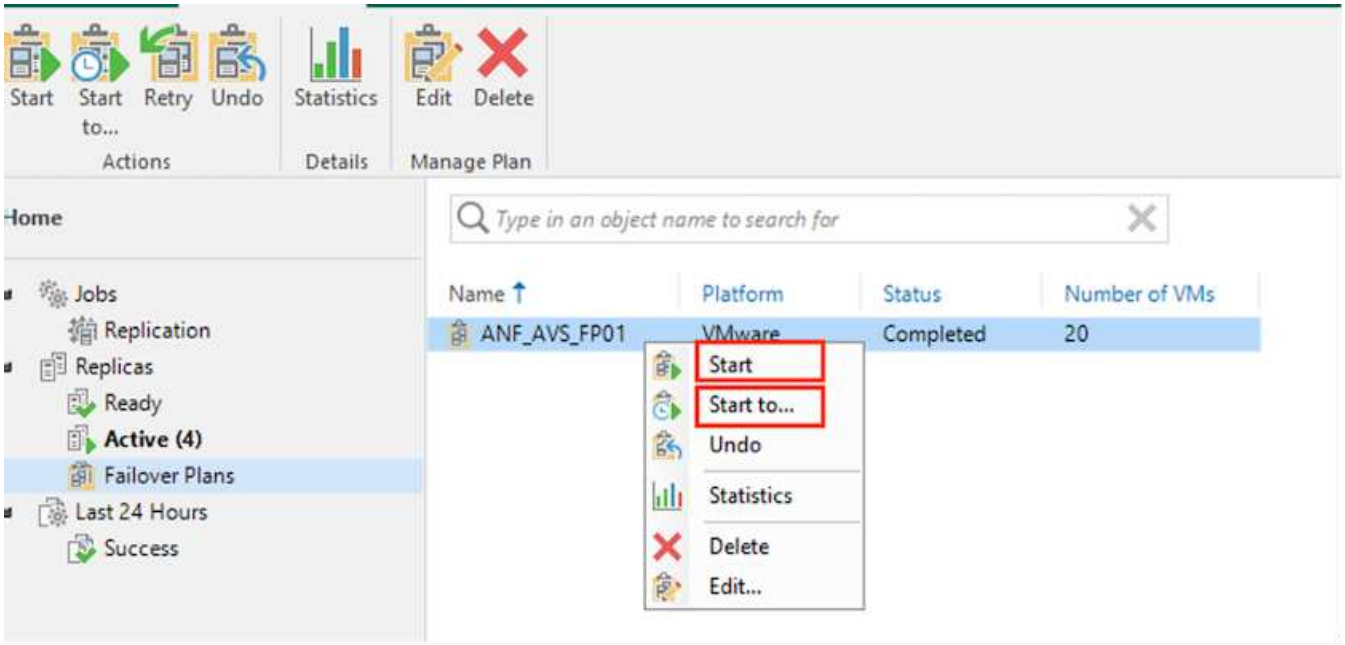
For additional information for creating replication jobs, refer Creating Replication Jobs.

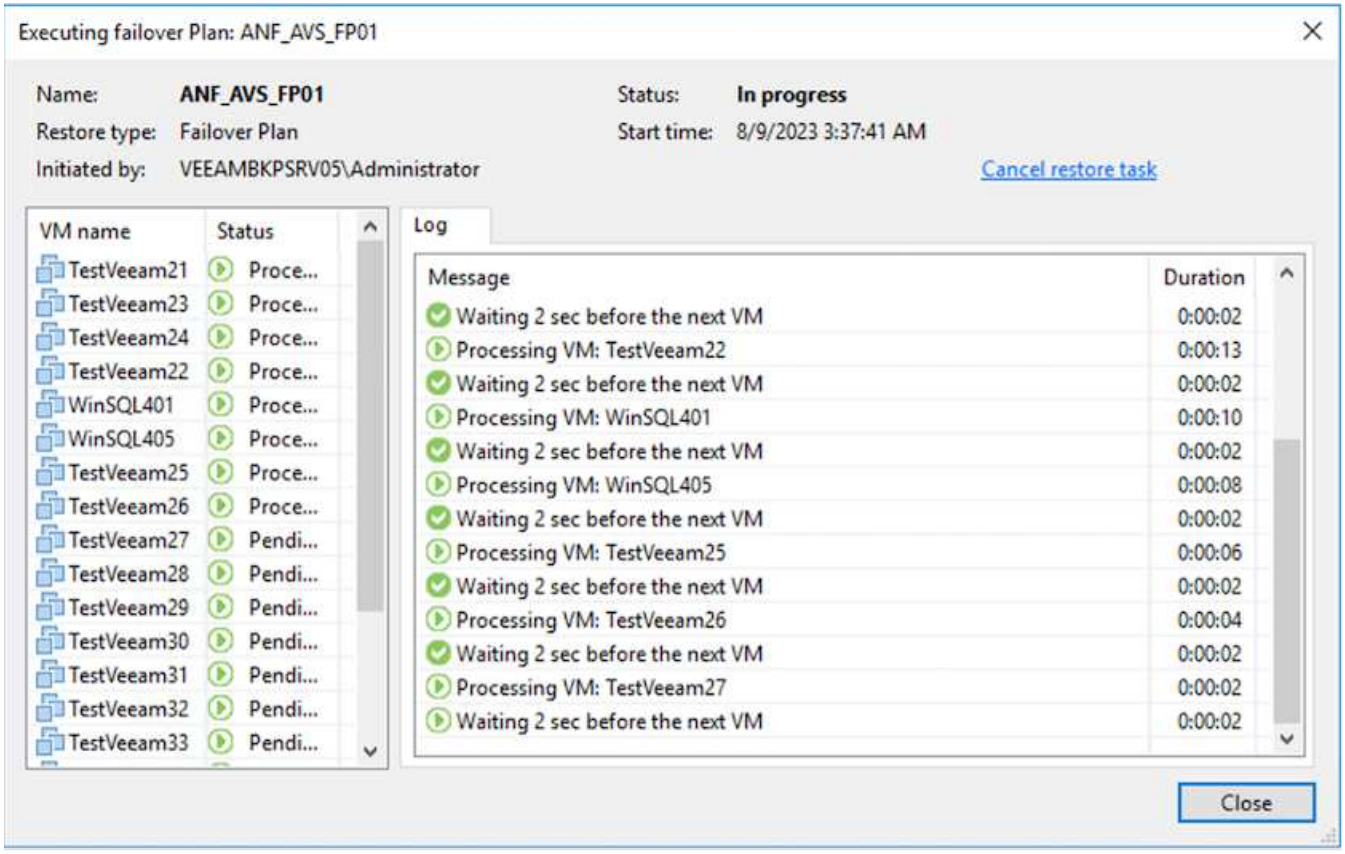
▼ **Step 3: Run the failover plan**

During failover, the source VM in the production site is switched over to its replica at the disaster recovery site. As part of the failover process, Veeam Backup & Replication restores the VM replica to the required restore point and moves all I/O activities from the source VM to its replica. Replicas can be used not only in case of a disaster, but also to simulate DR drills. During failover simulation, the source VM remains running. Once all the necessary tests have been conducted, you can undo the failover and return to normal operations.

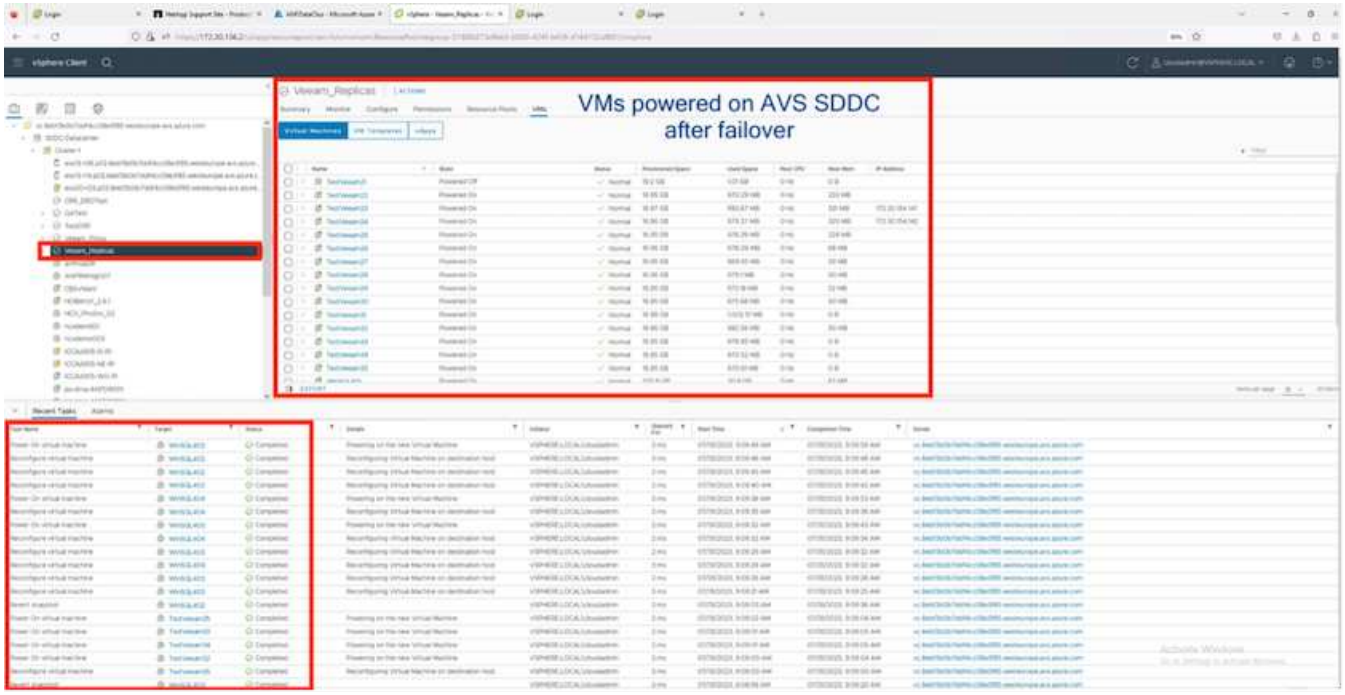
ⓘ Make sure network segmentation is in place to avoid IP conflicts during failover.

To start the failover plan, simply click in **Failover Plans** tab and right click on your failover plan. Select ***Start**. This will failover using the latest restore points of VM replicas. To fail over to specific restore points of VM replicas, select **Start to**.

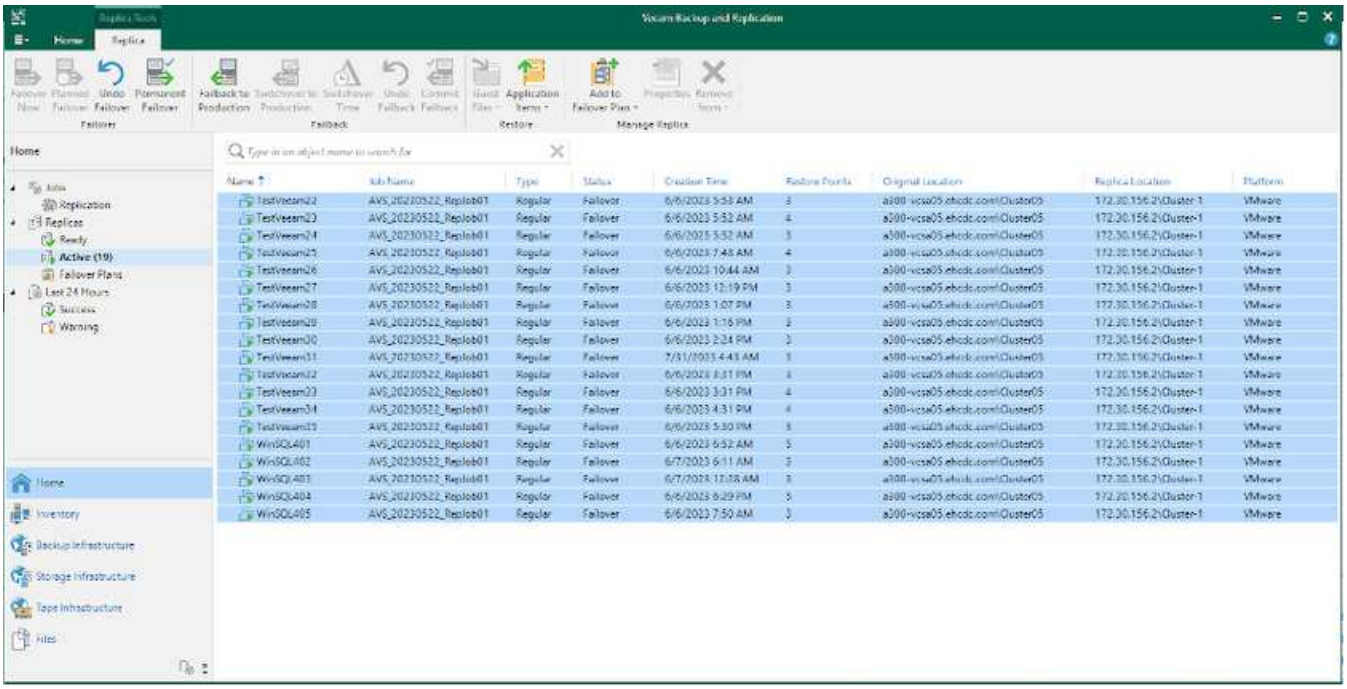




The state of the VM replica changes from Ready to Failover and VMs will start on the destination Azure VMware Solution (AVS) SDDC cluster / host.



Once the failover is complete, the status of the VMs will change to "Failover".



Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the Ready state.

For detailed information about failover plans, refer Failover Plans.

▼ **Step 4: Failback to the Production site**

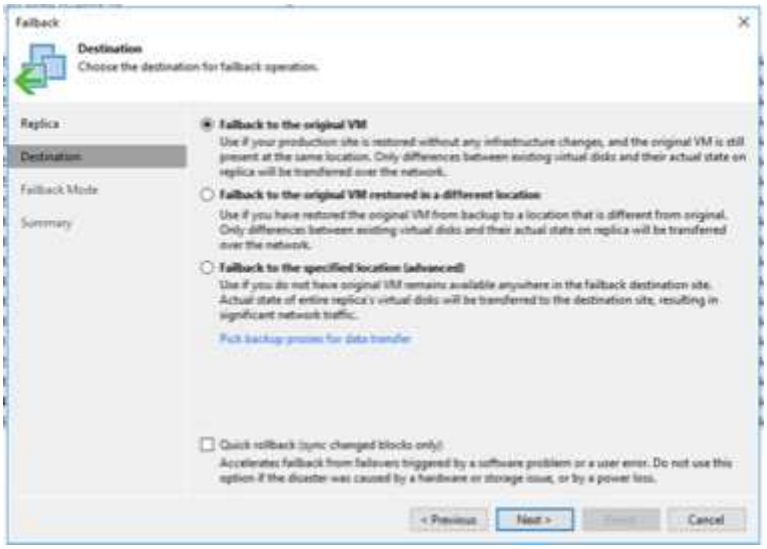
When the failover plan is running, it is considered as an intermediate step and needs to be finalized based on the requirement. The options include the following:

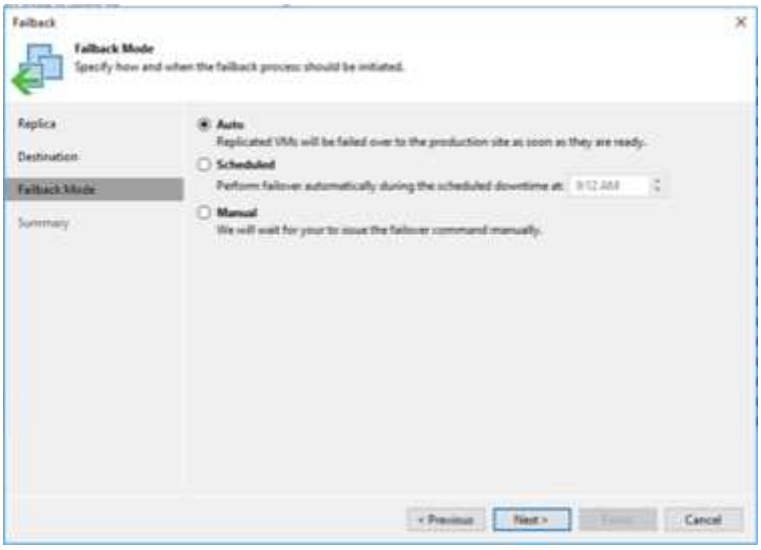
- **Failback to production** - switch back to the original VM and transfer all changes that took place while the VM replica was running to the original VM.

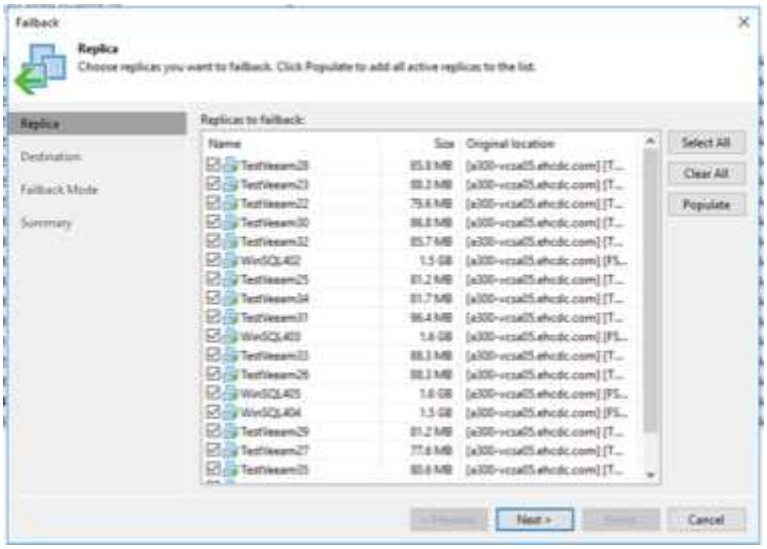
> When you perform failback, changes are only transferred but not published. Choose **Commit failback** (once the original VM is confirmed to work as expected) or Undo failback to get back to the VM replica If the original VM is not working as expected.

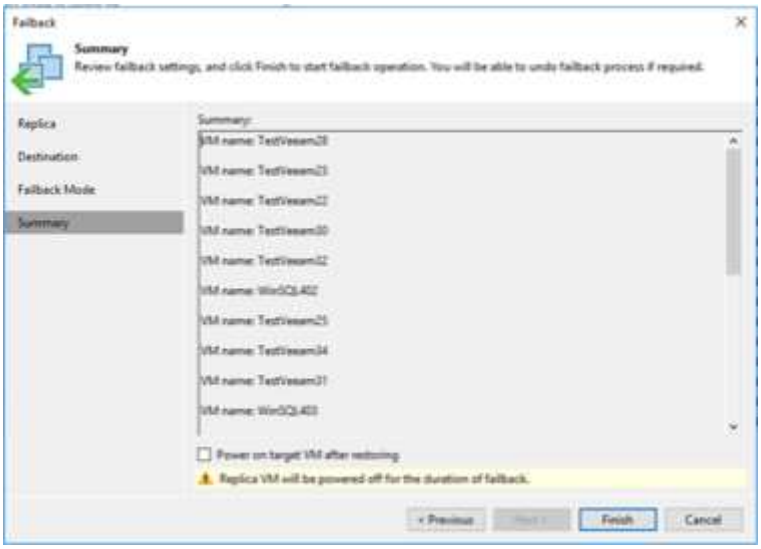
- **Undo failover** - switch back to the original VM and discard all changes made to the VM replica while it was running.
- **Permanent Failover** - permanently switch from the original VM to a VM replica and use this replica as the original VM.

In this demo, Failback to production was chosen. Failback to the original VM was selected during the Destination step of the wizard and "Power on VM after restoring" check box was enabled.



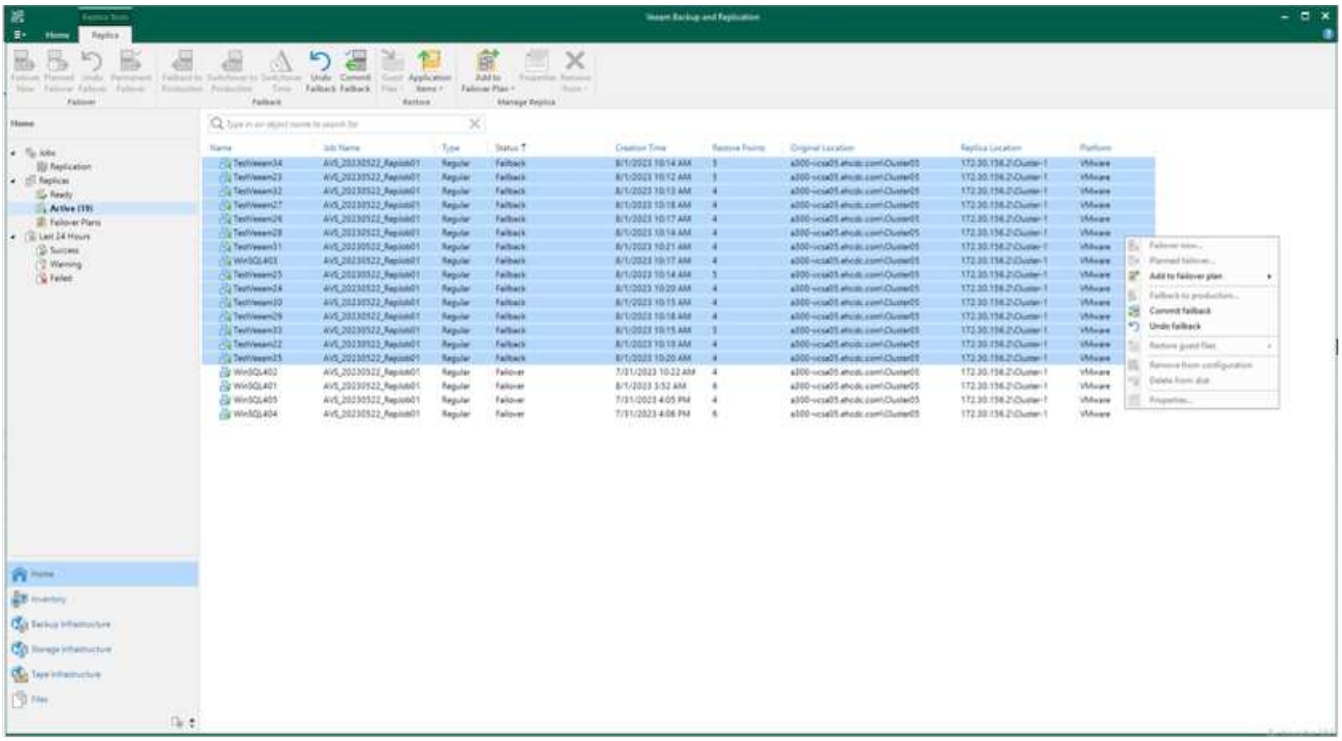




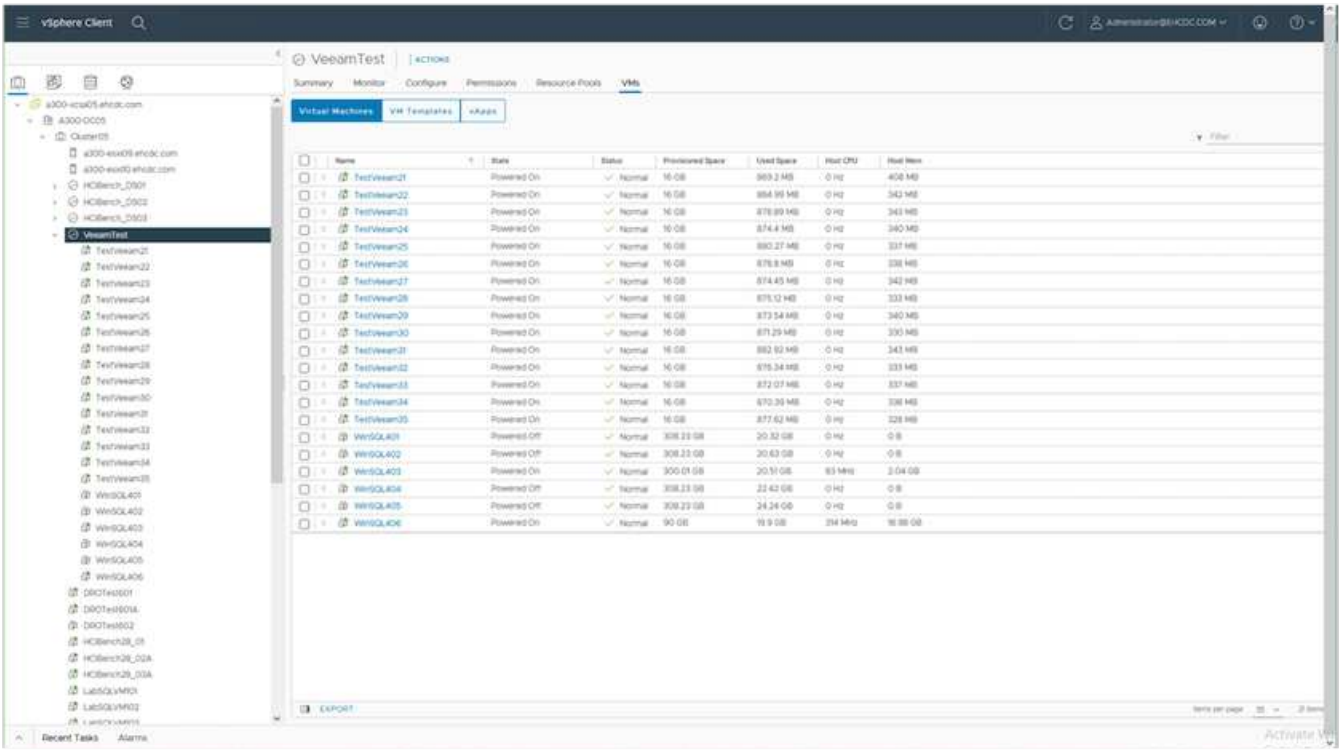


Failback commit is one of the ways to finalize failback operation. When failback is committed, it confirms that the changes sent to the VM which is failed back (the production VM) are working as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

For detailed information about the failback process, refer Veeam documentation for Failover and Failback for replication.



After failback to production is successful, the VMs are all restored back to the original production site.



Conclusion

Azure NetApp Files datastore capability enables Veeam or any validated third-party tool to provide a low-cost DR solution by leveraging Pilot light clusters instead of standing up a large cluster only to accommodate VM replicas. This provides an efficacious way to handle a tailored, customized disaster recovery plan and to reuse existing backup products in house for DR, enabling cloud-based disaster recovery by exiting on-premises DR datacenters. It is possible to failover by clicking a button in case of disaster or to failover automatically if a

disaster occurs.

To learn more about this process, feel free to follow the detailed walkthrough video.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=2855e0d5-97e7-430f-944a-b061015e9278

Migrate workloads on Azure/AVS

TR-4940: Migrate workloads to Azure NetApp Files datastore using VMware HCX - Quickstart guide

One of the most common use cases for the Azure VMware Solution and Azure NetApp Files datastore is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration mechanisms to move on-premises virtual machines (VMs) and its data to Azure NetApp Files datastores.

Overview: Migrating virtual machines with VMware HCX, Azure NetApp Files datastores, and Azure VMware solution

VMware HCX is primarily a migration platform that is designed to simplify application migration, workload rebalancing, and even business continuity across clouds. It is included as part of Azure VMware Solution Private Cloud and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for provisioning Azure NetApp Files datastore followed by downloading, deploying, and configuring VMware HCX, including all its main components in on-premises and the Azure VMware Solution side including Interconnect, Network Extension, and WAN optimization for enabling various VM migration mechanisms.



VMware HCX works with any datastore type as the migration is at the VM level. Hence this document is applicable to existing NetApp customers and non-NetApp customers who are planning to deploy Azure NetApp Files with Azure VMware Solution for a cost-effective VMware cloud deployment.

▼ **High-level steps**

This list provides the high-level steps necessary to install and configure HCX Cloud Manager on the Azure cloud side and install HCX Connector on-premises:

1. Install HCX through the Azure portal.
2. Download and deploy the HCX Connector Open Virtualization Appliance (OVA) installer in the on-premises VMware vCenter Server.
3. Activate HCX with the license key.
4. Pair the on-premises VMware HCX Connector with Azure VMware Solution HCX Cloud Manager.
5. Configure the network profile, compute profile, and service mesh.
6. (Optional) Perform network extension to avoid re-IP during migrations.
7. Validate the appliance status and ensure that migration is possible.
8. Migrate the VM workloads.

▼ **Prerequisites**

Before you begin, make sure the following prerequisites are met. For more information, see this link. After the prerequisites, including connectivity, are in place, configure and activate HCX by generating the license

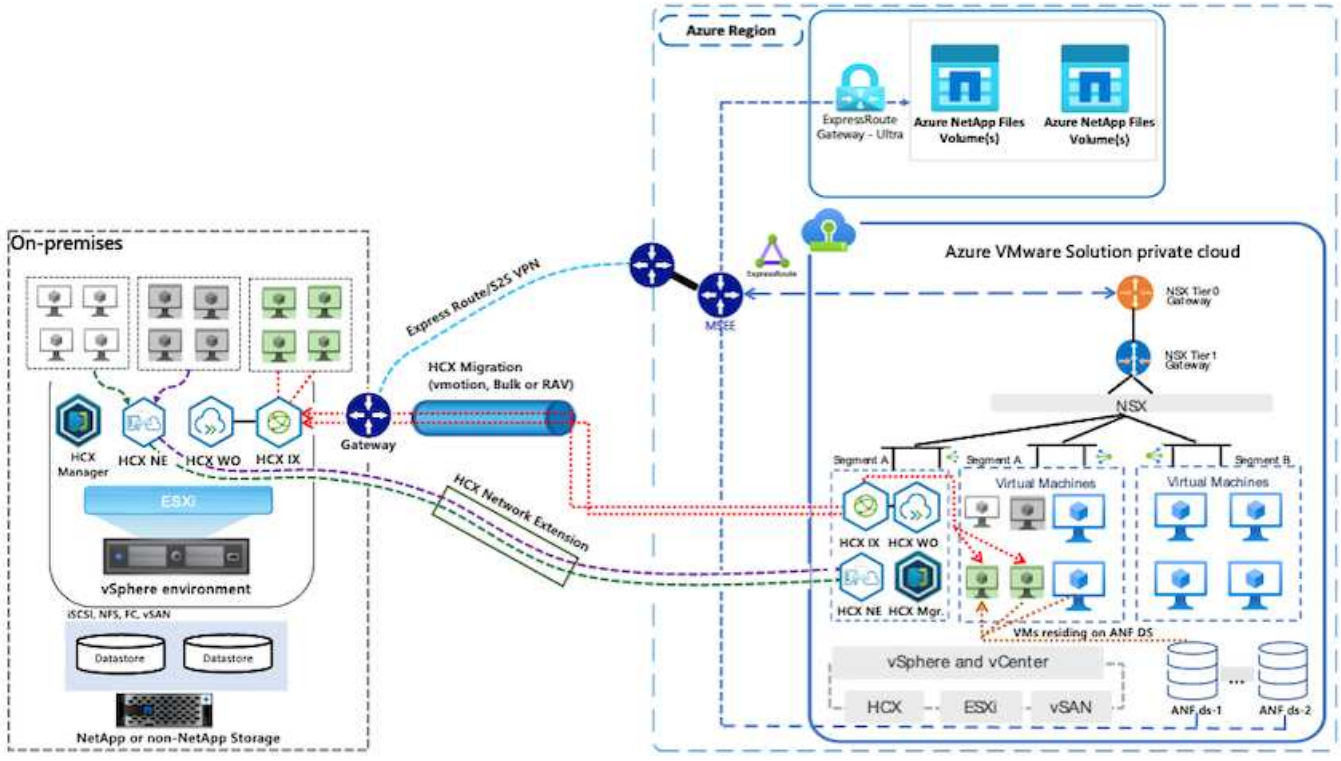
key from the Azure VMware Solution portal. After the OVA installer is downloaded, proceed with the installation process as described below.

> HCX advanced is the default option and VMware HCX Enterprise edition is also available through a support ticket and supported at no additional cost.

- Use an existing Azure VMware solution software-defined data center (SDDC) or create a private cloud by using this NetApp link or this Microsoft link.
- Migration of VMs and associated data from the on-premises VMware vSphere- enabled data center requires network connectivity from the data center to the SDDC environment. Before migrating workloads, set up a site-to-site VPN or Express route global reach connection between the on-premises environment and the respective private cloud.
- The network path from on-premises VMware vCenter Server environment to the Azure VMware Solution private cloud must support the migration of VMs by using vMotion.
- Make sure the required firewall rules and ports are allowed for vMotion traffic between the on-premises vCenter Server and SDDC vCenter. On the private cloud, routing on the vMotion network is configured by default.
- Azure NetApp Files NFS volume should be mounted as a datastore in Azure VMware Solution. Follow the steps detailed in this link to attach Azure NetApp Files datastores to Azure VMware Solutions hosts.

▼ High Level Architecture

For testing purposes, the lab environment from on-premises used for this validation was connected through a site-to-site VPN, which allows on-premises connectivity to Azure VMware Solution.



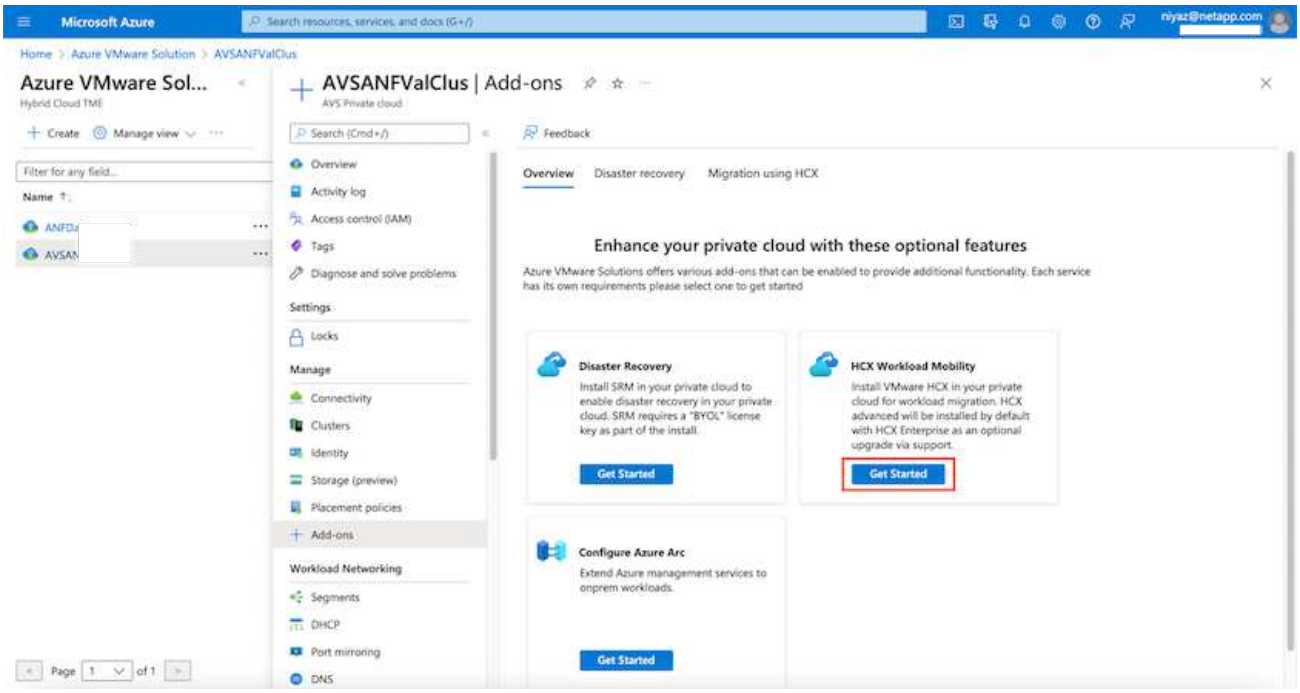
Solution Deployment

Follow the series of steps to complete the deployment of this solution:

▼ Step 1: Install HCX through Azure Portal using the Add-ons option

To perform the installation, complete the following steps:

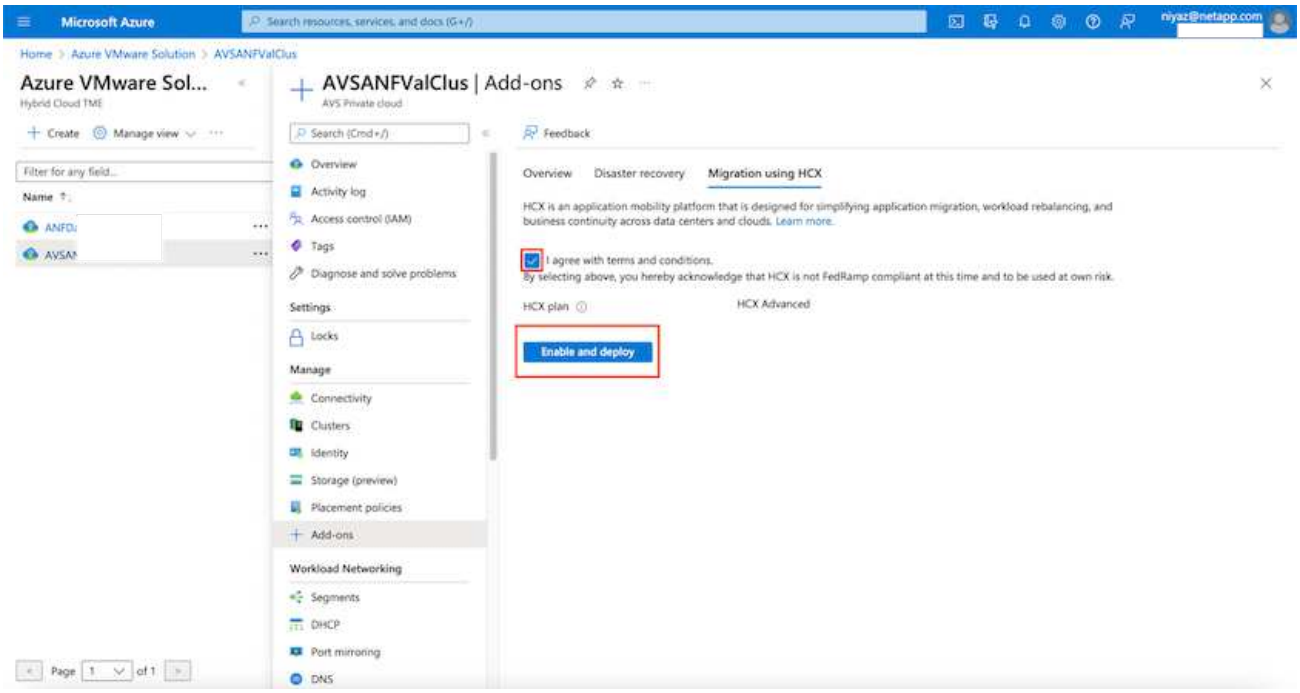
1. Log in to the Azure Portal and access the Azure VMware Solution private cloud.
2. Select the appropriate private cloud and access Add-ons. This can be done by navigating to **Manage > Add-ons**.
3. In the HCX Workload Mobility section, click **Get Started**.



4. Select the **I Agree with Terms and Conditions** option and click **Enable and Deploy**.

> The default deployment is HCX Advanced. Open a support request to enable the Enterprise edition.

> The deployment takes approximately 25 to 30 minutes.



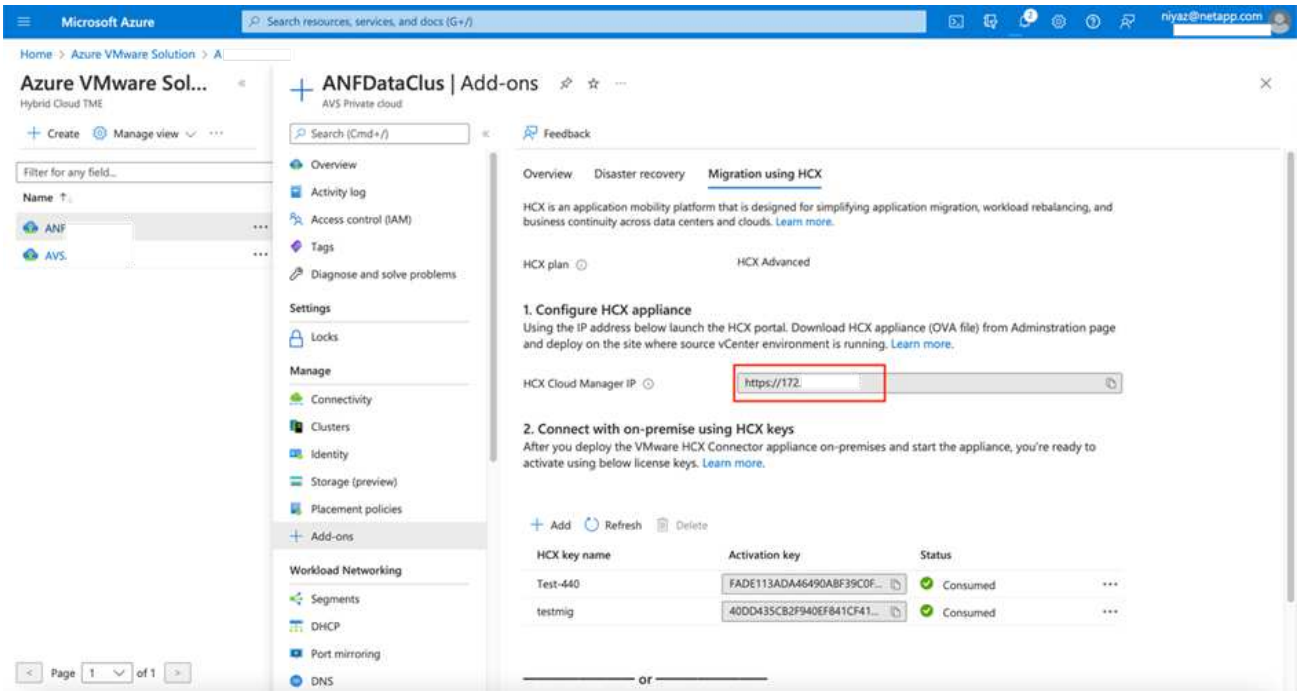
▼ **Step 2: Deploy the installer OVA in the on-premises vCenter Server**

For the on-premises Connector to connect to the HCX Manager in Azure VMware Solution, make sure the appropriate firewall ports are open in the on-premises environment.

To download and install HCX Connector in the on-premises vCenter Server, complete the following steps:

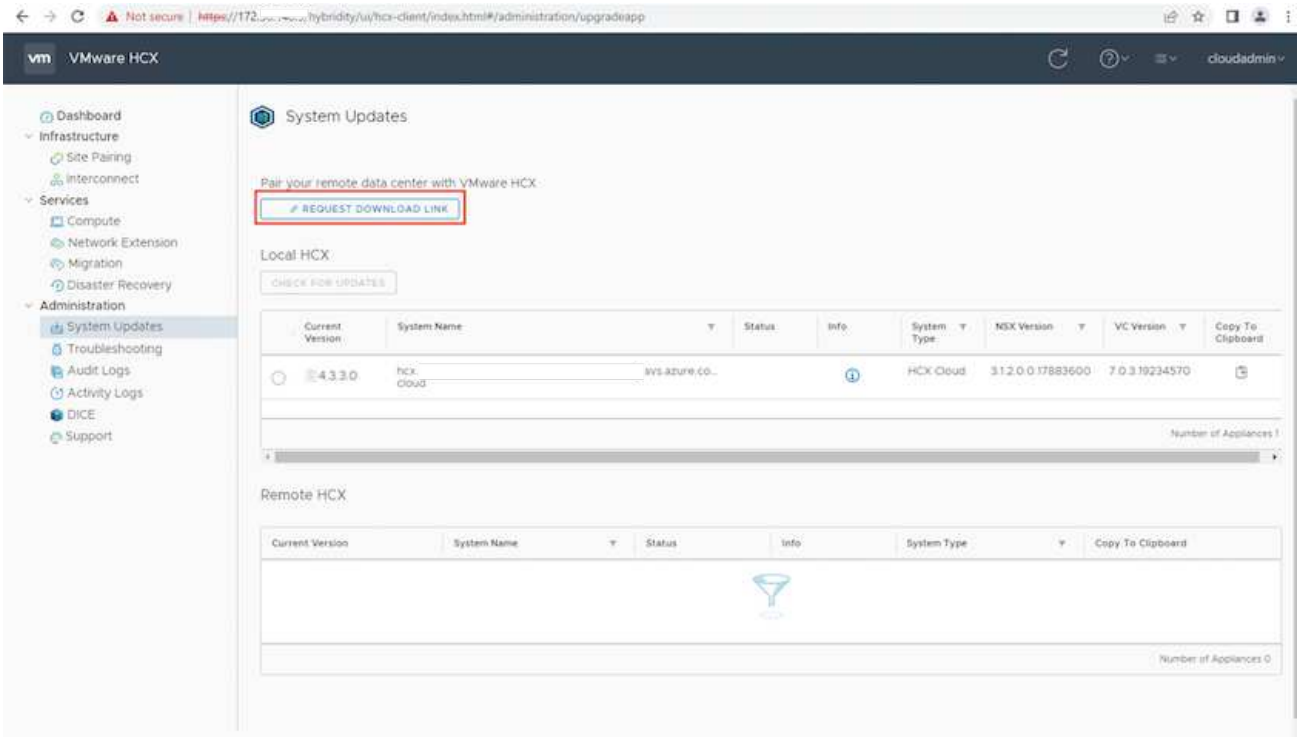
1. From the Azure portal, go to the Azure VMware Solution, select the private cloud, and select **Manage > Add-ons > Migration** using HCX and copy the HCX Cloud Manager portal to download the OVA file.

Use the default CloudAdmin user credentials to access the HCX portal.

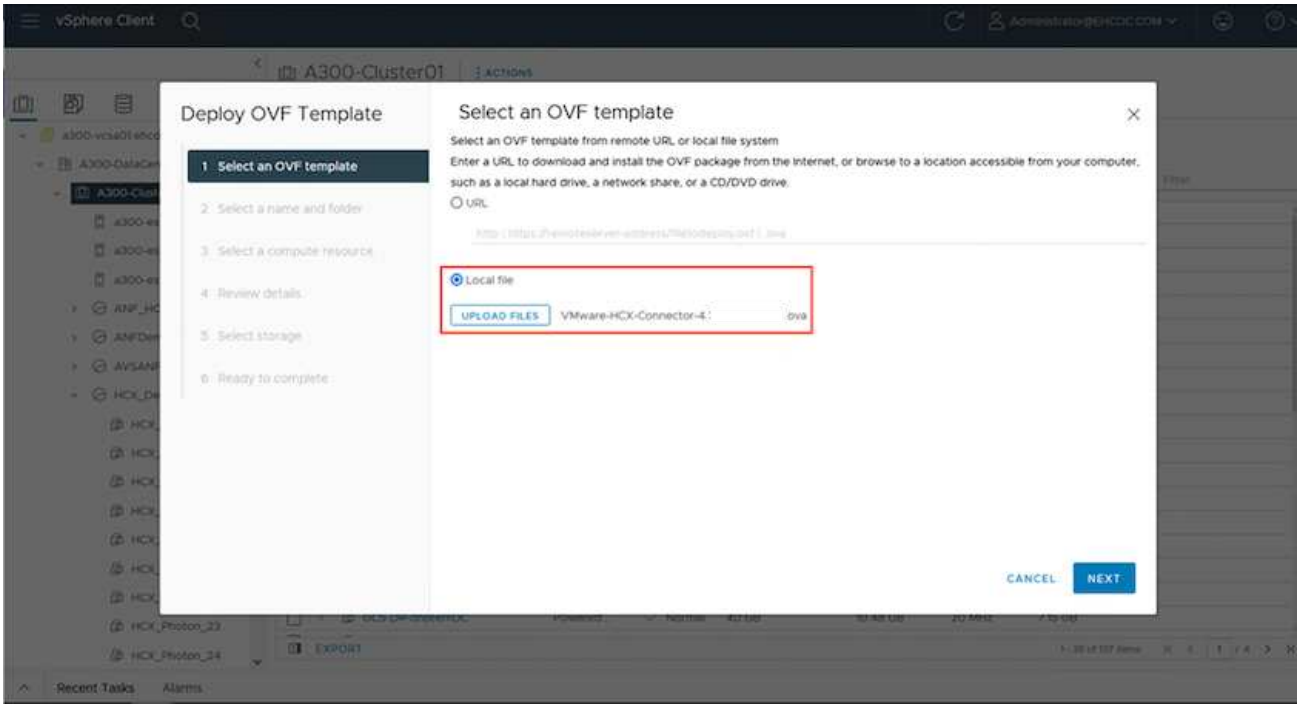


2. After you access the HCX portal with cloudadmin@vsphere.local using the jumphost, navigate to **Administration > System Updates** and click **Request Download Link**.

> Either download or copy the link to the OVA and paste it into a browser to begin the download process of the VMware HCX Connector OVA file to deploy on the on-premises vCenter Server.



3. After the OVA is downloaded, deploy it on to the on-premises VMware vSphere environment by using the **Deploy OVF Template** option.



4. Enter all the required information for the OVA deployment, click **Next**, and then click **Finish** to deploy the VMware HCX connector OVA.

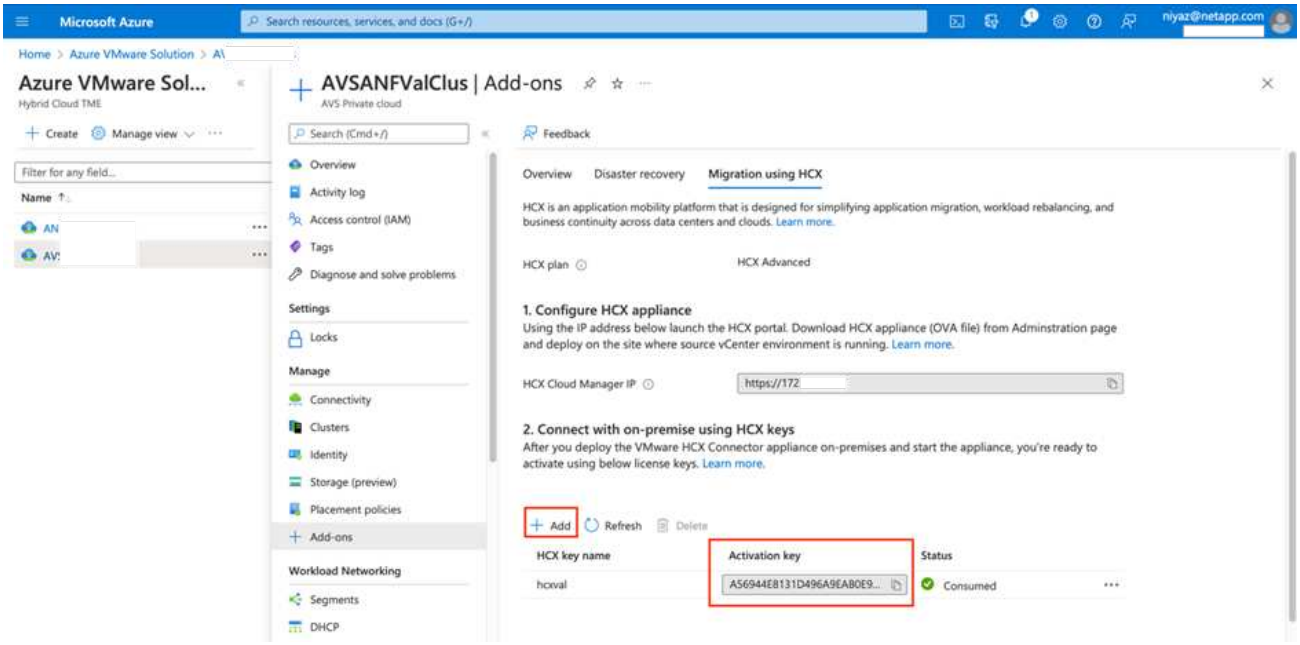
   Power on the virtual appliance manually.

For step-by-step instructions, see the VMware HCX User Guide.

▼ Step 3: Activate HCX Connector with the license key

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the Azure VMware Solution portal and activate it in VMware HCX Manager.

1. From the Azure portal, go to the Azure VMware Solution, select the private cloud, and select **Manage > Add-ons > Migration using HCX**.
2. Under **Connect with on-premise Using HCX keys**, click **Add** and copy the activation key.



   A separate key is required for each on-premises HCX Connector that is deployed.

3. Log into the on-premises VMware HCX Manager at "`https://hcxmanagerIP:9443`" using administrator credentials.

   Use the password defined during the OVA deployment.

4. In the licensing, enter the key copied from step 3 and click **Activate**.

   The on-premises HCX Connector should have internet access.

5. Under **Datacenter Location**, provide the nearest location for installing the VMware HCX Manager on-premises. Click **Continue**.
6. Under **System Name**, update the name and click **Continue**.
7. Click **Yes, Continue**.

8. Under **Connect your vCenter**, provide the fully qualified domain name (FQDN) or IP address of vCenter Server and the appropriate credentials and click **Continue**.

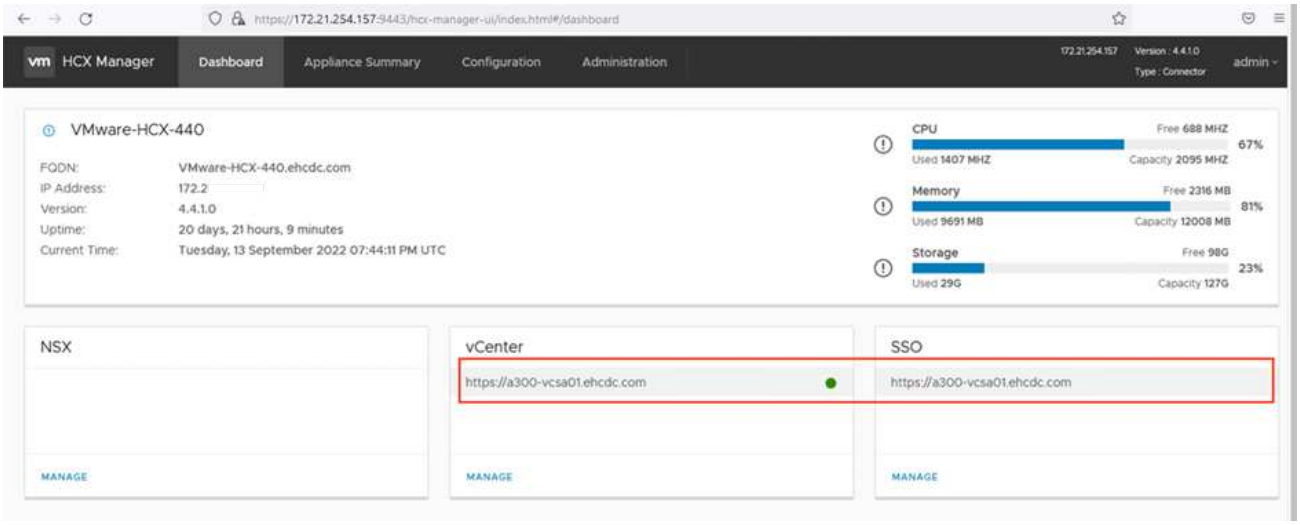
   Use the FQDN to avoid connectivity issues later.

9. Under **Configure SSO/PSC**, provide the Platform Services Controller's FQDN or IP address and click **Continue**.

   Enter the VMware vCenter Server FQDN or IP address.

10. Verify that the information entered is correct and click **Restart**.
11. After the services restart, vCenter Server is displayed as green on the page that appears. Both vCenter Server and SSO must have the appropriate configuration parameters, which should be the same as the previous page.

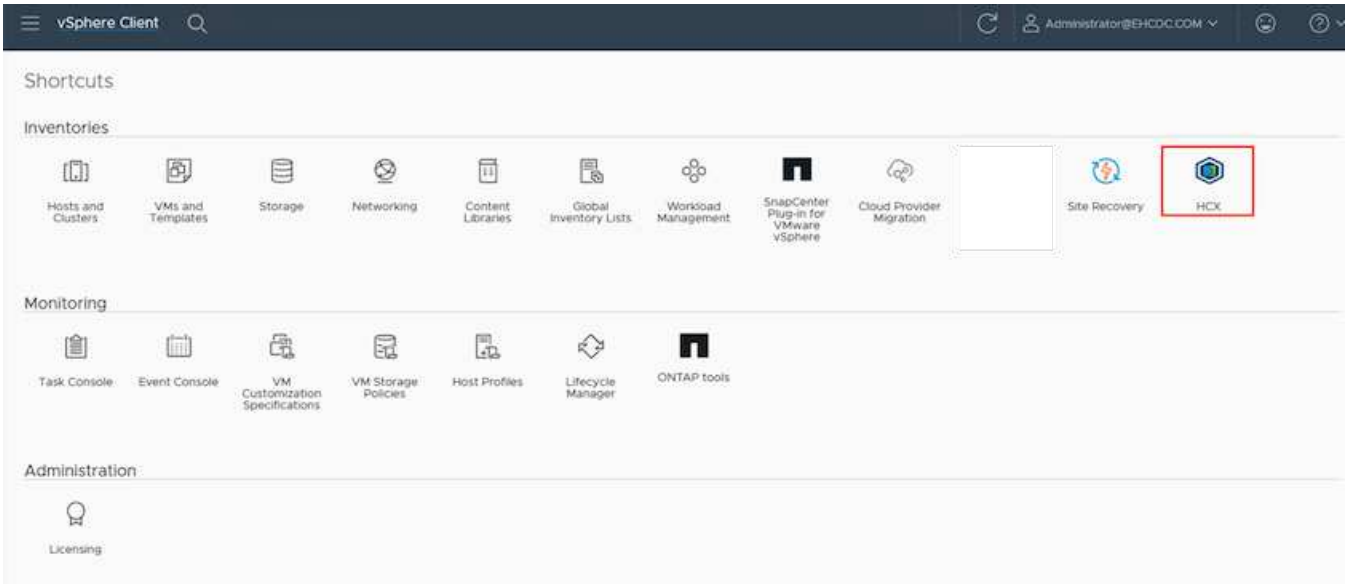
    This process should take approximately 10 to 20 minutes and for the plug-in to be added to the vCenter Server.



▼ Step 4: Pair on-premises VMware HCX Connector with Azure VMware Solution HCX Cloud Manager

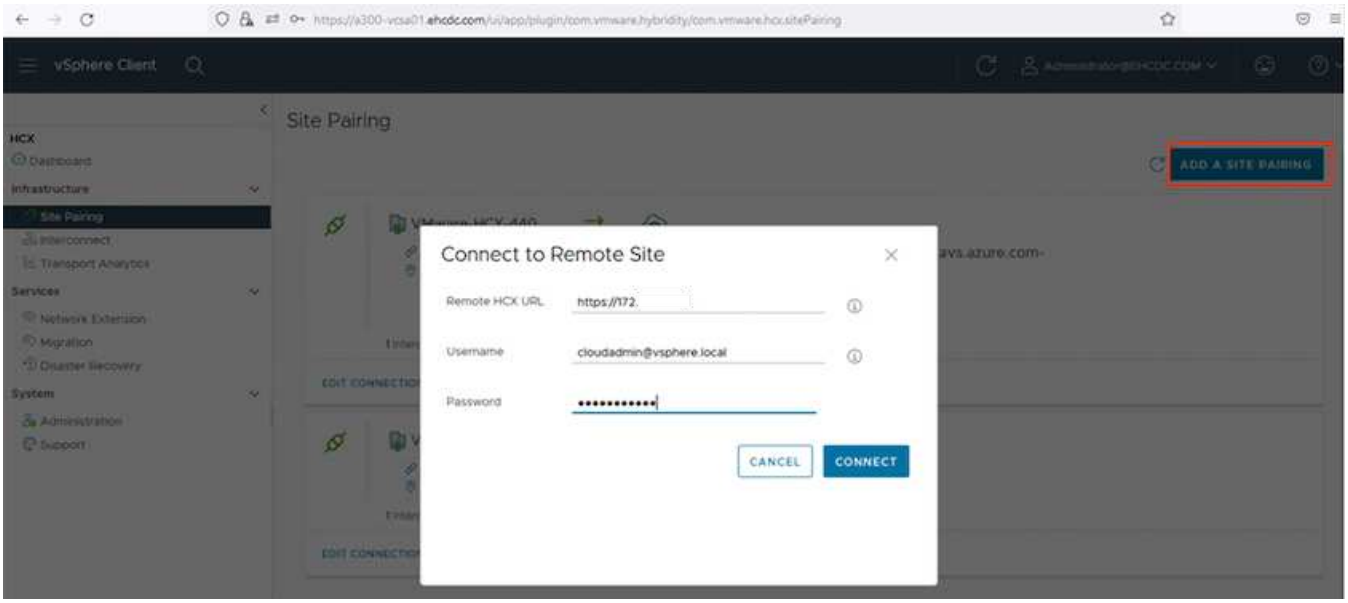
After HCX Connector is installed in both on-premises and Azure VMware Solution, configure the on-premises VMware HCX Connector for Azure VMware Solution private cloud by adding the pairing. To configure the site pairing, complete the following steps:

1. To create a site pair between the on-premises vCenter environment and Azure VMware Solution SDDC, log in to the on-premises vCenter Server and access the new HCX vSphere Web Client plug-in.



1. Under Infrastructure, click **Add a Site Pairing**.

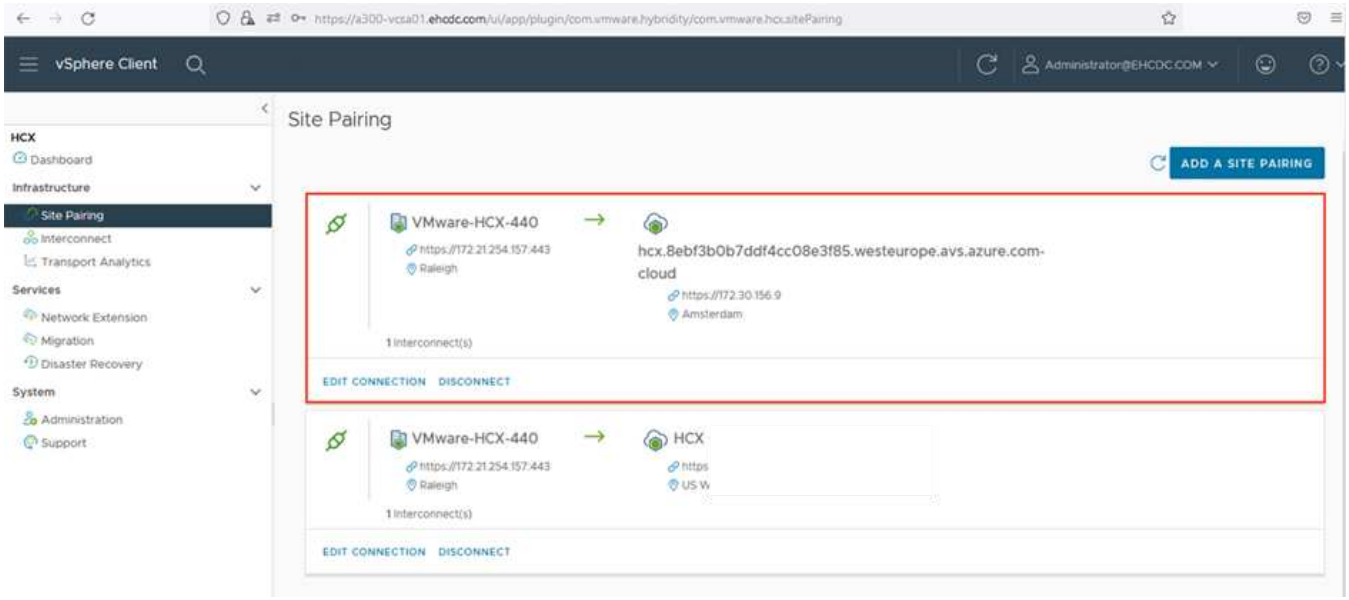
Enter the Azure VMware Solution HCX Cloud Manager URL or IP address and the credentials for CloudAdmin role for accessing the private cloud.



1. Click **Connect**.

VMware HCX Connector must be able to route to HCX Cloud Manager IP over port 443.

1. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.

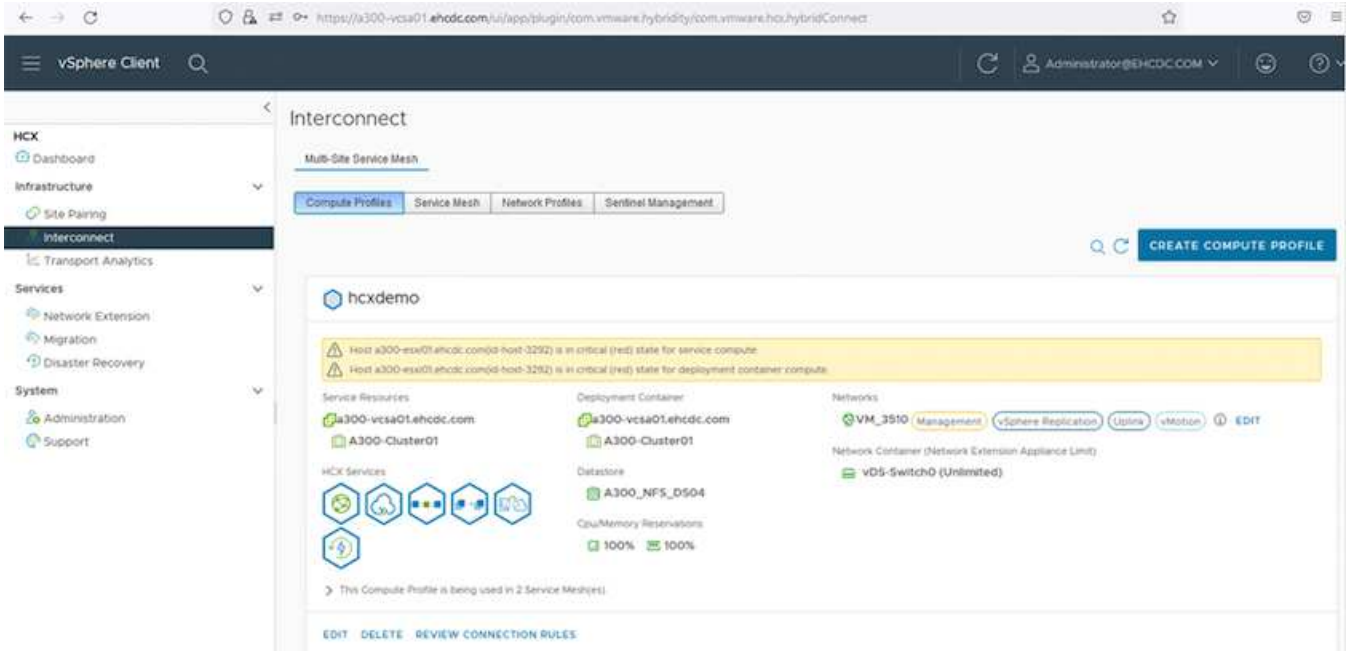


▼ **Step 5: Configure the network profile, compute profile, and service mesh**

The VMware HCX Interconnect service appliance provides replication and vMotion-based migration capabilities over the internet and private connections to the target site. The interconnect provides encryption, traffic engineering, and VM mobility. To create an Interconnect service appliance, complete the followings steps:

1. Under Infrastructure, select **Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile**.

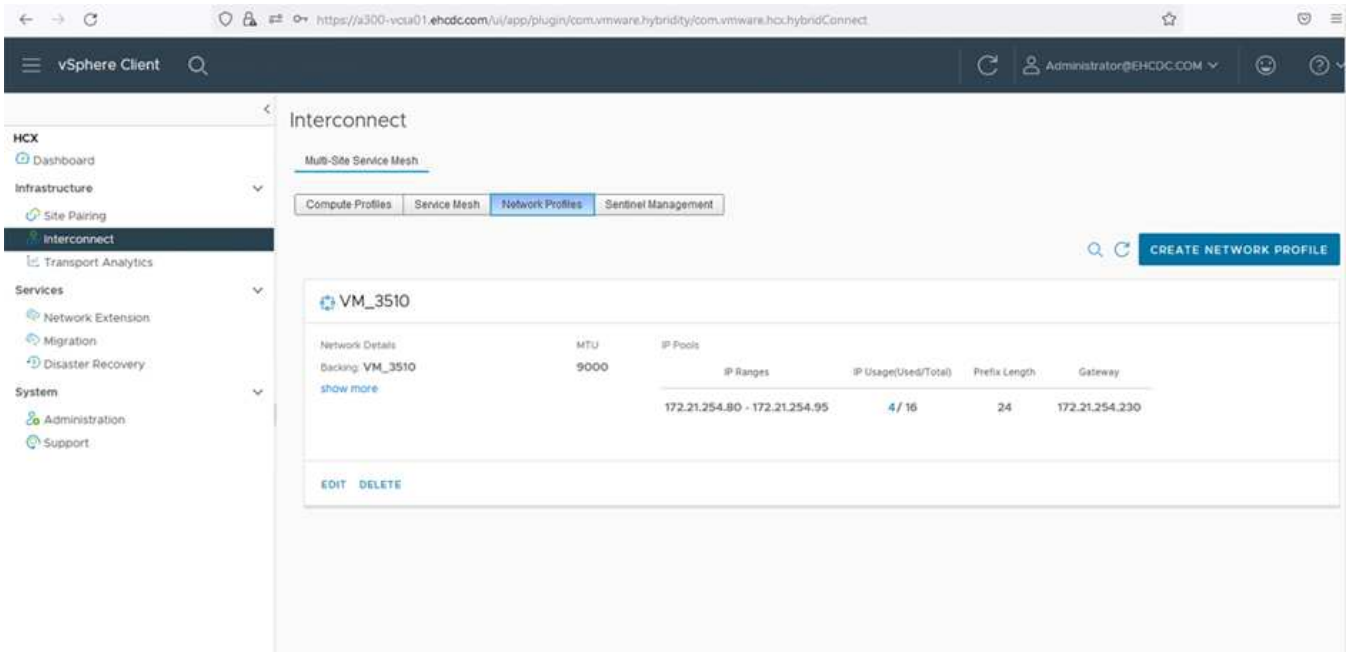
The compute profiles define the deployment parameters including the appliances that are deployed and which portion of the VMware data center are accessible to HCX service.



1. After the compute profile is created, create the network profiles by selecting **Multi-Site Service Mesh > Network Profiles > Create Network Profile**.

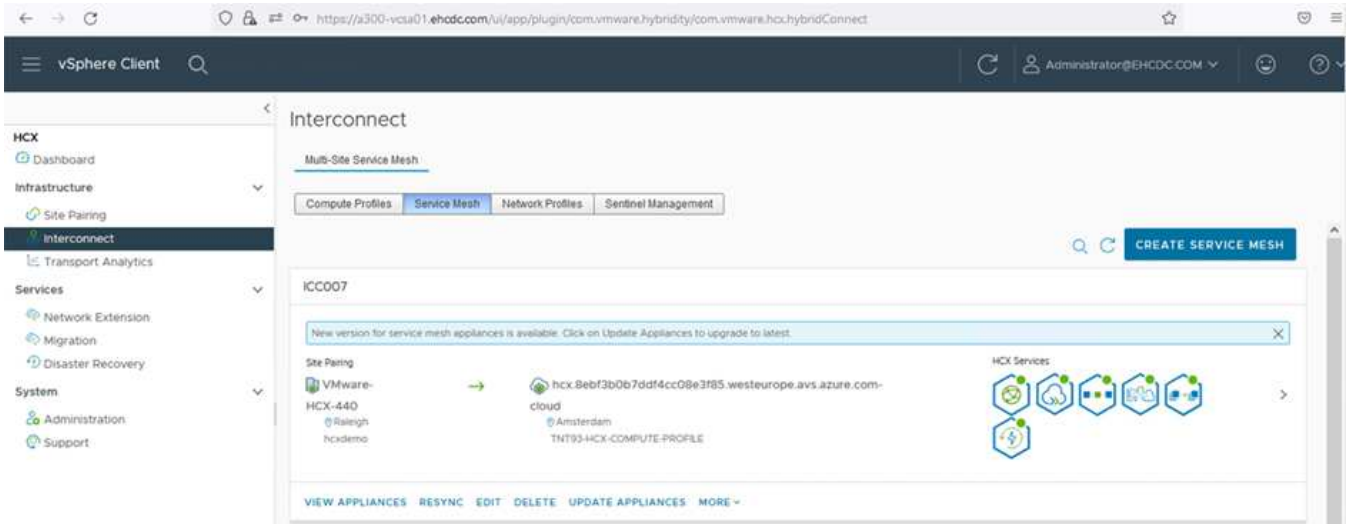
The network profile defines a range of IP address and networks that are used by HCX for its virtual appliances.

This step requires two or more IP addresses. These IP addresses are assigned from the management network to the Interconnect Appliances.



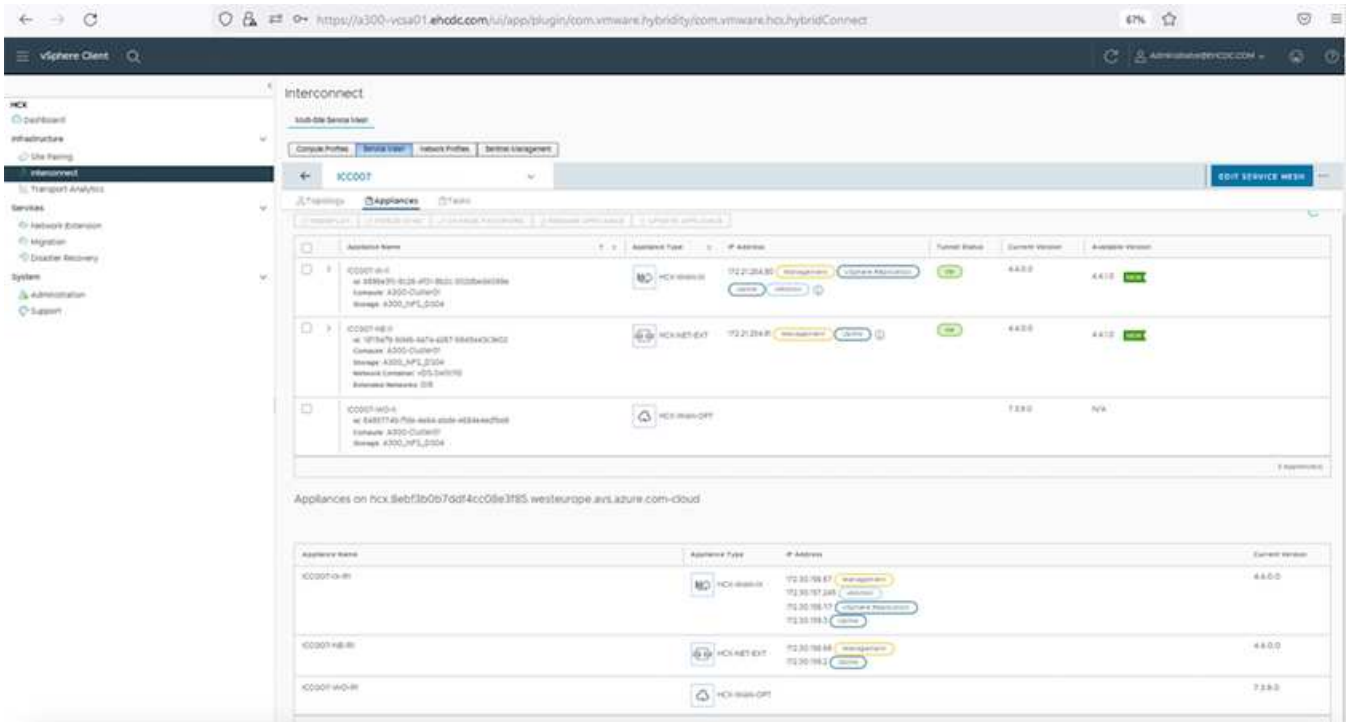
1. At this time, the compute and network profiles have been successfully created.
2. Create the Service Mesh by selecting the **Service Mesh** tab within the **Interconnect** option and select the on-premises and Azure SDDC sites.
3. The Service Mesh specifies a local and remote compute and network profile pair.

As part of this process, the HCX appliances are deployed and automatically configured on both the source and target sites in order to create a secure transport fabric.



1. This is the final step of configuration. This should take close to 30 minutes to complete the deployment. After the service mesh is configured, the environment is ready with the IPsec tunnels successfully

created to migrate the workload VMs.



▼ **Step 6: Migrate workloads**

Workloads can be migrated bidirectionally between on-premises and Azure SDDCs using various VMware HCX migration technologies. VMs can be moved to and from VMware HCX-activated entities using multiple migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with the HCX Enterprise edition).

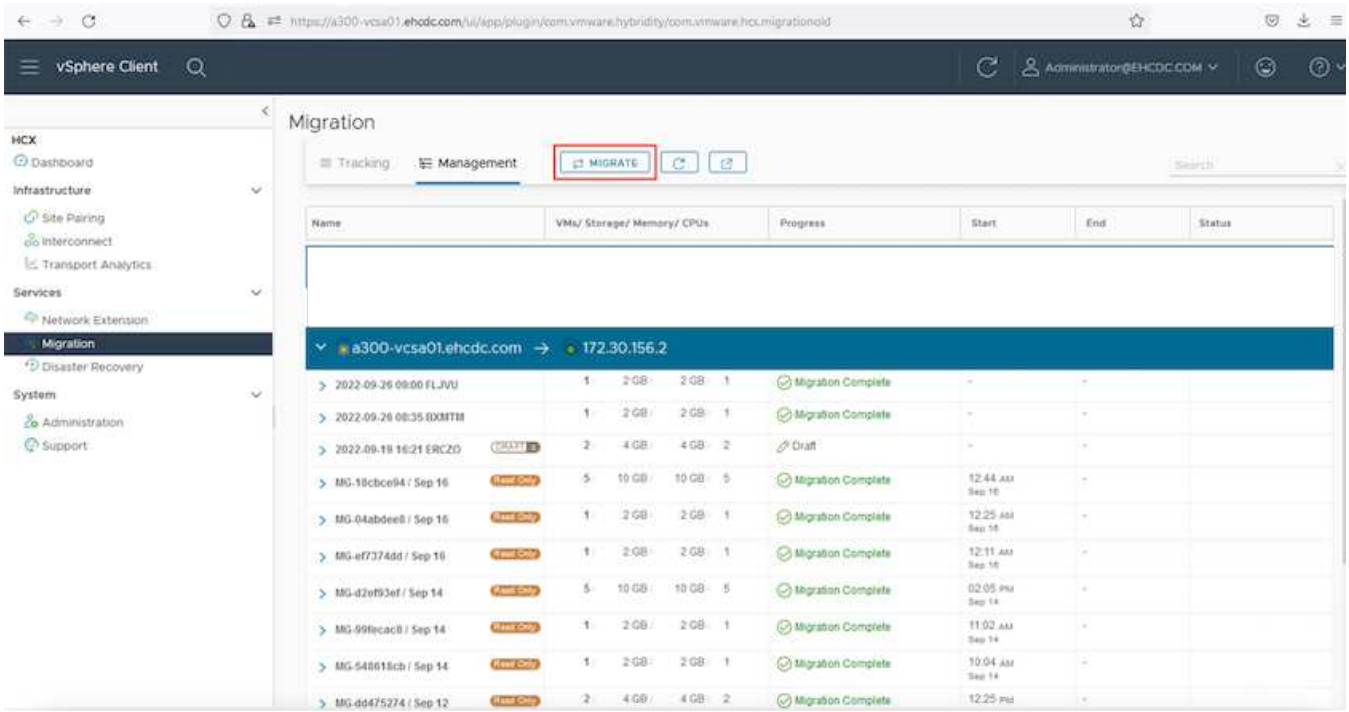
To learn more about various HCX migration mechanisms, see VMware HCX Migration Types.

Bulk migration

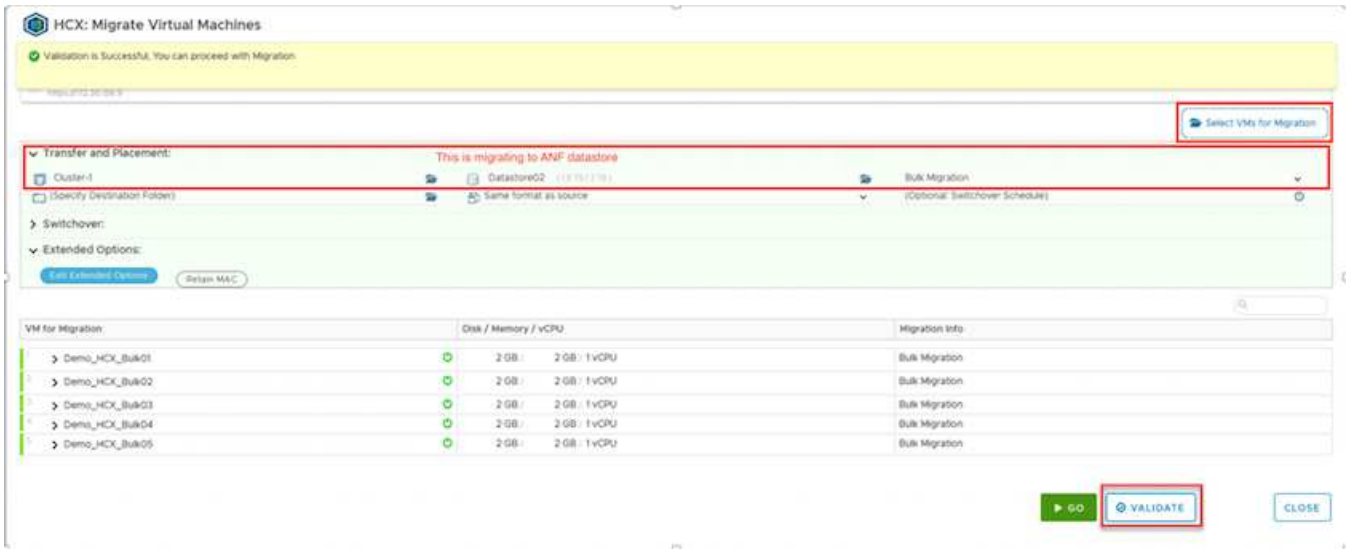
This section details the bulk migration mechanism. During a bulk migration, the bulk migration capability of HCX uses vSphere Replication to migrate disk files while recreating the VM on the destination vSphere HCX instance.

To initiate bulk VM migrations, complete the following steps:

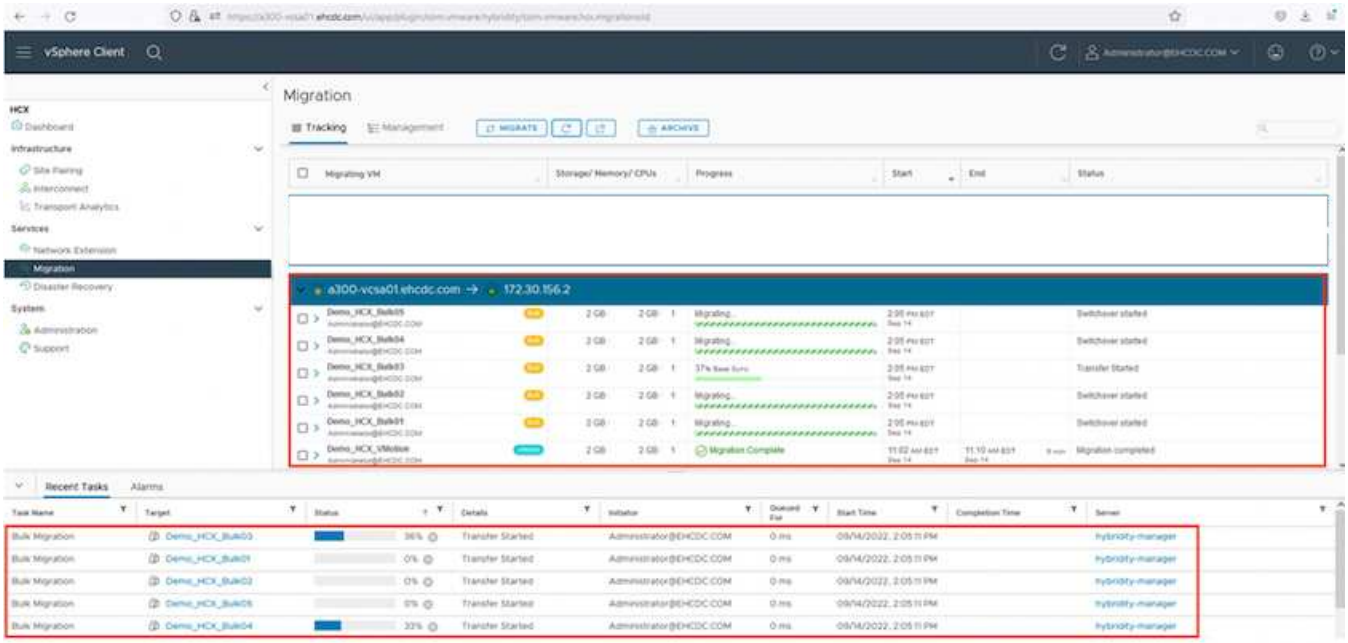
1. Access the **Migrate** tab under **Services > Migration**.



1. Under **Remote Site Connection**, select the remote site connection and select the source and destination. In this example, the destination is Azure VMware Solution SDDC HCX endpoint.
2. Click **Select VMs for Migration**. This provides a list of all the on-premises VMs. Select the VMs based on the match:value expression and click **Add**.
3. In the **Transfer and Placement** section, update the mandatory fields (**Cluster**, **Storage**, **Destination**, and **Network**), including the migration profile, and click **Validate**.

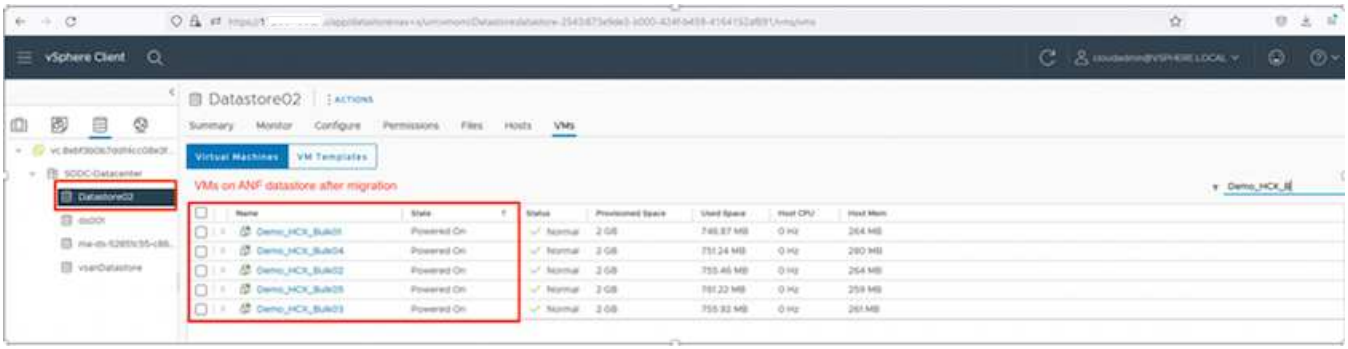


1. After the validation checks are complete, click **Go** to initiate the migration.



During this migration, a placeholder disk is created on the specified Azure NetApp Files datastore within the target vCenter to enable replication of the source VM disk's data to the placeholder disks. HBR is triggered for a full sync to the target, and after the baseline is complete, an incremental sync is performed based on the recovery point objective (RPO) cycle. After the full/incremental sync is complete, switchover is triggered automatically unless a specific schedule is set.

1. After the migration is complete, validate the same by accessing the destination SDDC vCenter.

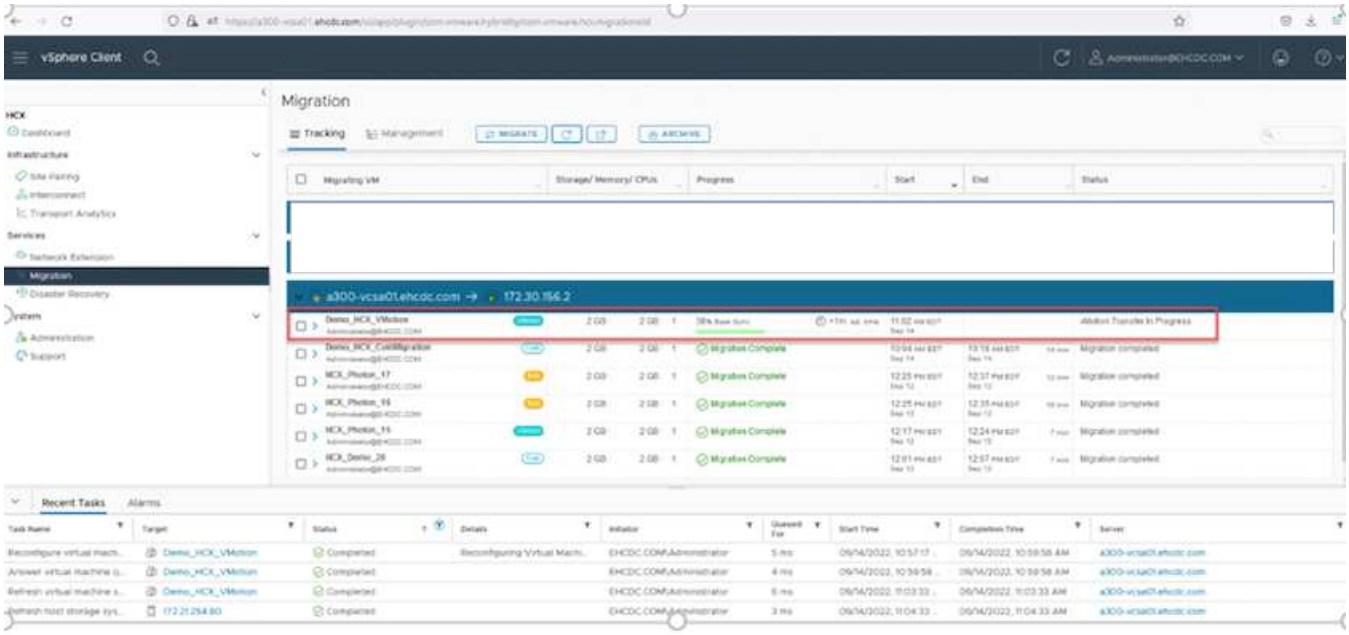


For additional and detailed information about various migration options and on how to migrate workloads from on-premises to Azure VMware Solution using HCX, see VMware HCX Migration Considerations.

To learn more about this process, feel free to watch the following video:

Workload Migration using HCX

Here is a screenshot of HCX vMotion option.



To learn more about this process, feel free to watch the following video:

HCX vMotion

Make sure sufficient bandwidth is available to handle the migration.

The target ANF datastore should have sufficient space to handle the migration.

Conclusion

Whether you're targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in on-premises, Azure NetApp Files and HCX provide excellent options to deploy and migrate the application workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Azure VMware Solution along with Azure NetApp Files for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere Replication, VMware vMotion, or even network file copy (NFC).

Takeaways

The key points of this document include:

- You can now use Azure NetApp Files as a datastore on Azure VMware Solution SDDC.
- You can easily migrate data from on-premises to Azure NetApp Files datastore.
- You can easily grow and shrink the Azure NetApp Files datastore to meet the capacity and performance requirements during migration activity.

Where to find additional information

To learn more about the information described in this document, refer to the following website links:

- Azure VMware Solution documentation

https://docs.microsoft.com/en-us/azure/azure-vmware/

- Azure NetApp Files documentation

https://docs.microsoft.com/en-us/azure/azure-netapp-files/

- VMware HCX migration considerations

https://learn.microsoft.com/en-us/azure/azure-vmware/architecture-migrate

NetApp for GCP/GCVE

NetApp Capabilities for Google Cloud Platform GCVE

Learn more about the capabilities that NetApp brings to the Google Cloud Platform (GCP) Google Cloud VMware Engine (GCVE) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- Configuring GCVE in GCP
- NetApp Storage Options for GCVE
- NetApp / VMware Cloud Solutions

Configuring GCVE in GCP

As with on-premises, planning a cloud based virtualization environment is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP and Google Cloud NetApp Volumes to GCVE.

The setup process can be broken down into the following steps:

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed configuration steps for GCVE.

NetApp Storage Options for GCVE

NetApp storage can be utilized in several ways - either as guess connected or as a supplemental NFS datastore - within GCP GCVE.

Please visit Supported NetApp Storage Options for more information.

Google Cloud supports NetApp storage in the following configurations:

- Cloud Volumes ONTAP (CVO) as guest connected storage
- Google Cloud NetApp Volumes (NetApp Volumes) as guest connected storage
- Google Cloud NetApp Volumes (NetApp Volumes) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.
View the detailed supplemental NFS datastore options for GCVE.

Read more about Google Cloud NetApp Volumes datastore support for Google Cloud VMware Engine (NetApp blog) or How to use Google Cloud NetApp Volumes as datastores for Google Cloud VMware Engine (Google blog)

Solution Use Cases

With NetApp and VMware cloud solutions, many use cases are simple to deploy in Azure AVS. se cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Extend
- Migrate

Browse the NetApp solutions for Google Cloud GCVE

Protect workloads on GCP/GCVE

Application Consistent Disaster Recovery with NetApp SnapCenter and Veeam Replication

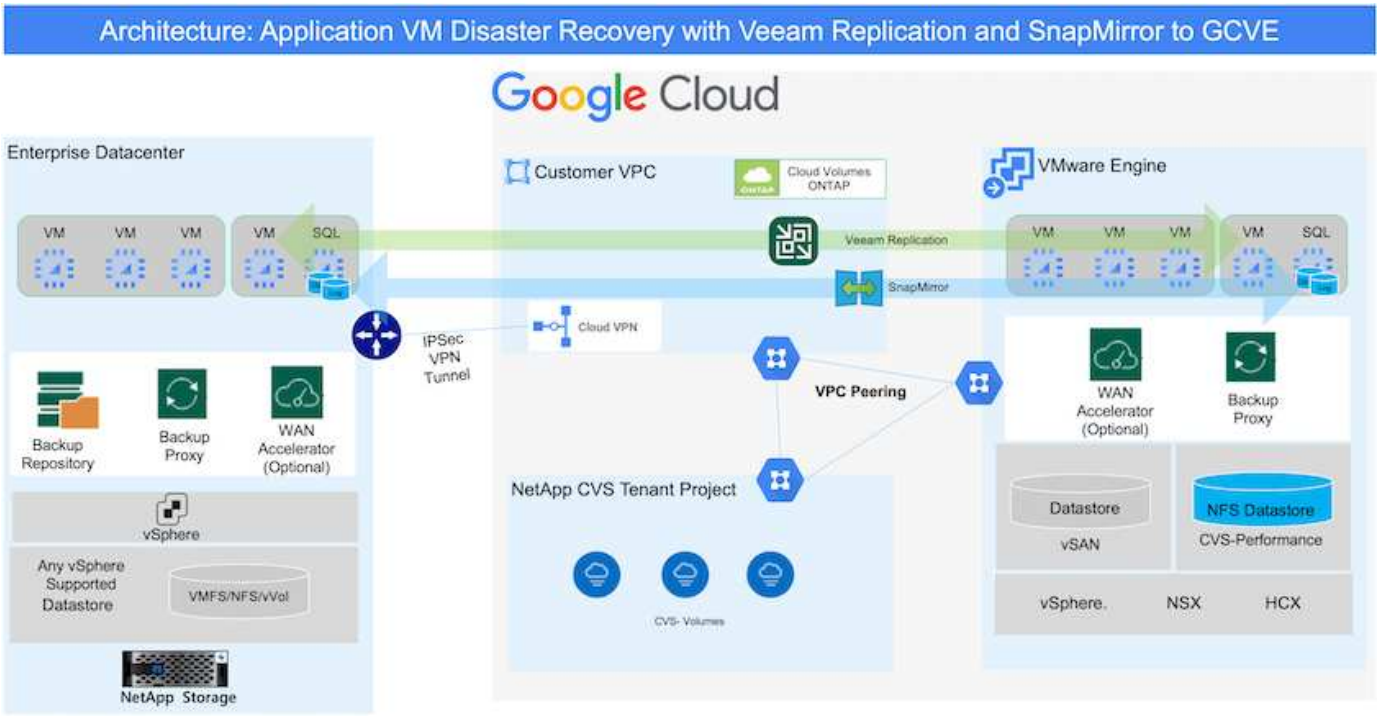
Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Google Cloud.

Overview

Many customers are looking for an effective disaster recovery solution for their application VMs hosted on VMware vSphere. Many of them use their existing backup solution to perform recovery during diaster.
Many times that solution increase the RTO and doesn’t meet their expectations. To reduce the RPO and RTO, Veeam VM replication can be utilized even from on-prem to GCVE as long as network connectivity and environemnt with appropriate permissions are available.
NOTE: Veeam VM Replication doesn’t protect VM guest connected storage devices like iSCSI or NFS mounts inside the guest VM. Need to protect those seperately.

For application consistent replication for SQL VM and to reduce the RTO, we used SnapCenter to orchestrate snapmirror operations of SQL database and log volumes.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, Veeam, and the Google Cloud VMware Engine (GCVE).



Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we assume that the on-premises environment is using SnapCenter for application-consistent backups.

> This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Google Cloud network, use the connectivity options like dedicated interconnect or Cloud VPN. Segments should be created based on the on-premises VLAN design.

> There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google connectivity method.

Deploying the DR Solution

Solution Deployment Overview

1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
2. Provision Cloud Volumes ONTAP with the correct instance size using NetApp Console within the appropriate subscription and virtual network.
   a. Configure SnapMirror for the relevant application volumes.
   b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
3. Install the Veeam software and start replicating virtual machines to Google Cloud VMware Engine instance.
4. During a disaster event, break the SnapMirror relationship using NetApp Console and trigger failover of virtual machines with Veeam.

a. Reconnect the ISCSI LUNs and NFS mounts for the application VMs.

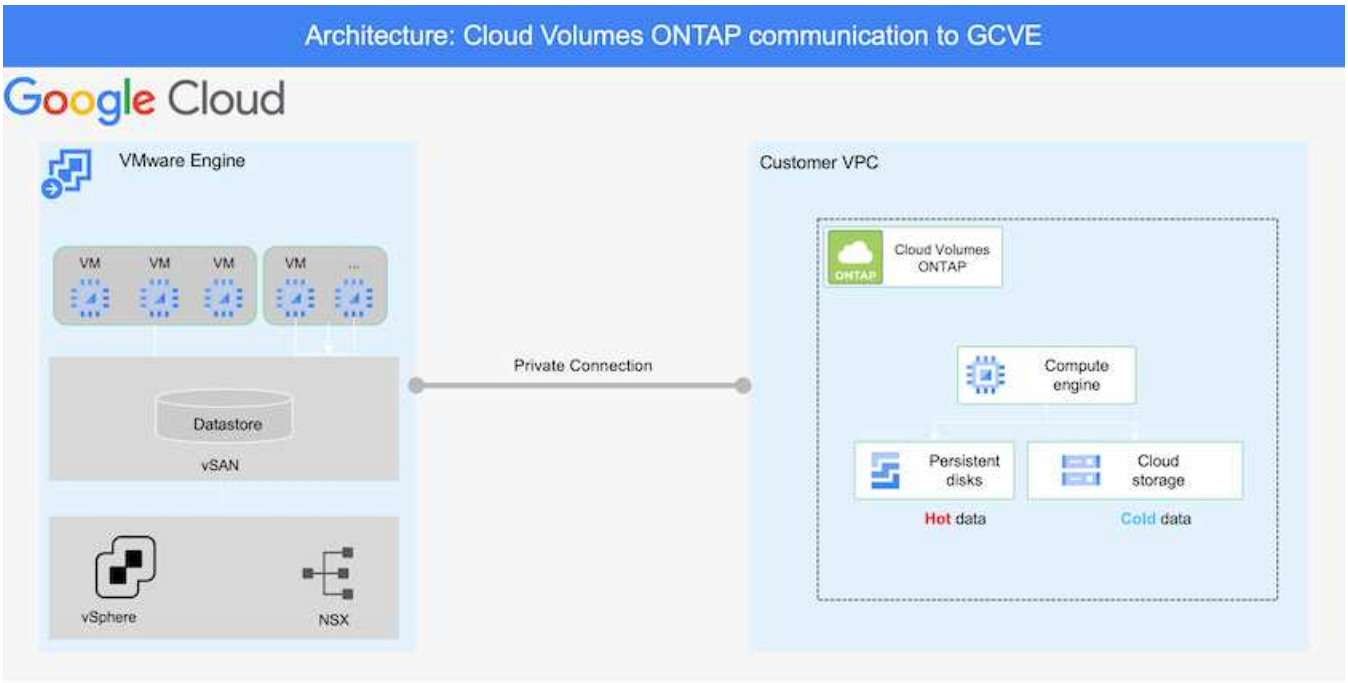
b. Bring up applications online.

5. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

Deployment Details

▼ Configure CVO on Google Cloud and replicate volumes to CVO

The first step is to configure Cloud Volumes ONTAP on Google Cloud (cvo) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.



For sample step-by-step instructions on setting up SnapCenter and replicating the data, Refer to Setup Replication with SnapCenter

Review of SQL VM protection with SnapCenter

▼ Configure GCVE hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the GCVE solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

Google Cloud NetApp Volumes for NFS Datastore and Cloud Volumes ONTAP for SQL databases and log can be deployed to any VPC and GCVE should have private connection to that VPC to mount NFS datastore and have VM connect to iSCSI LUNs.

To configure GCVE SDDC, see Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP). As a prerequisite, verify that the guest VMs residing on the GCVE hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and GCVE have been configured properly, begin configuring Veeam to automate the recovery of on-premises workloads to GCVE (VMs with application VMDKs and VMs with in-

guest storage) by using the Veeam Replication feature and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

▼ **Install Veeam Components**

Based on deployment scenario, the Veeam backup server, backup repository and backup proxy that needs to be deployed. For this use case, there is no need to deploy object store for Veeam and Scale-out repository also not required.
Refer to the Veeam documentation for the installation procedure
For additional information, please refer Migration with Veeam Replication

▼ **Setup VM Replication with Veeam**

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. Setup vSphere VM Replication Job At the Guest Processing step of wizard, select disable application processing as we will be utilizing SnapCenter for application aware backup and recovery.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=8b7e4a9b-7de1-4d48-a8e2-b01200f00692

▼ **Failover of Microsoft SQL Server VM**

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=9762dc99-081b-41a2-ac68-b01200f00ac0

Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
  - This avoids the risk of data corruption for volumes and snapshots.
  - Avoids replication interruptions during DR test workflows.
  - Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- Veeam Replication allows changing VM IP addresses on DR site.

Application Disaster Recovery with SnapCenter, Cloud Volumes ONTAP and Veeam Replication

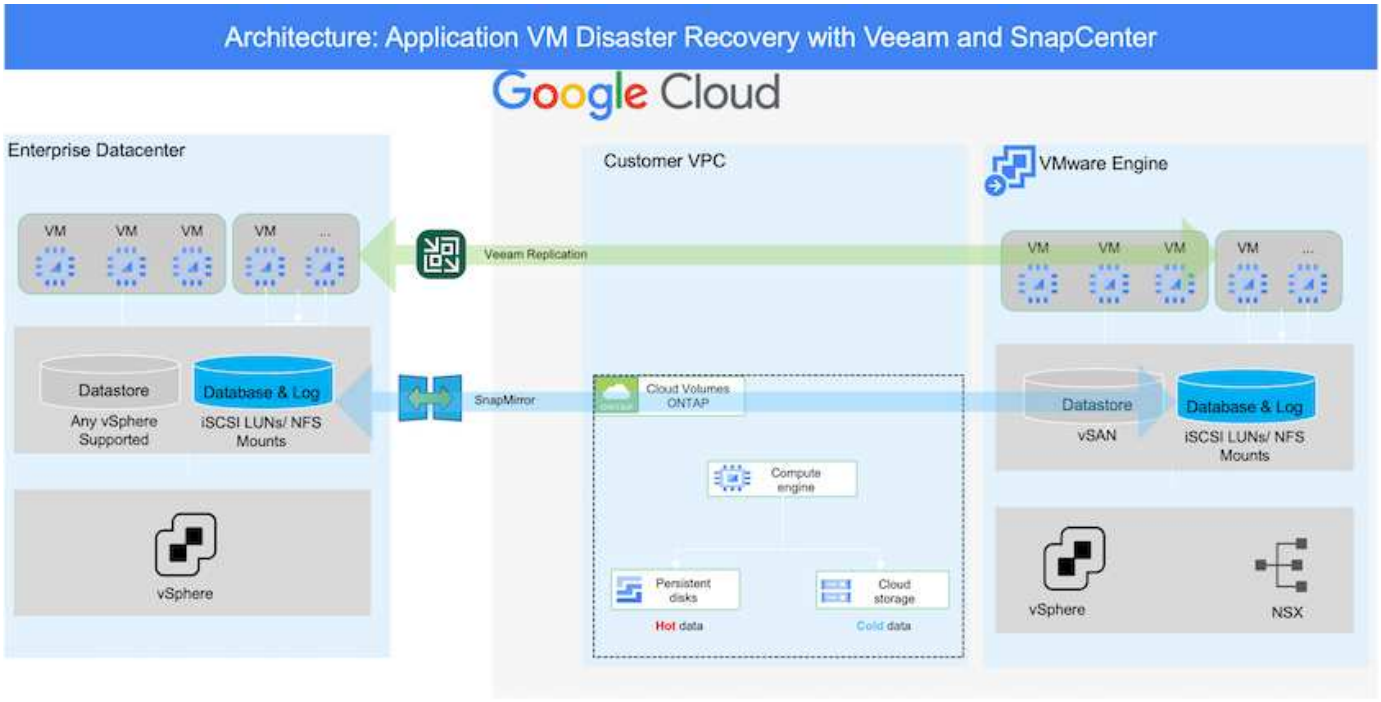
Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Google Cloud.

Overview

This covers application data; however, what about the actual VMs themselves. Disaster recovery should cover all dependent components, including virtual machines, VMDKs, application data, and more. To accomplish this, SnapMirror along with Veeam can be used to seamlessly recover workloads replicated from on-premises to Cloud Volumes ONTAP while using vSAN storage for VM VMDKs.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses

NetApp SnapMirror, Veeam, and the Google Cloud VMware Engine (GCVE).



Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we assume that the on-premises environment is using SnapCenter for application-consistent backups.

> This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Google Cloud network, use the connectivity options like dedicated interconnect or Cloud VPN. Segments should be created based on the on-premises VLAN design.

> There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google connectivity method.

Deploying the DR Solution

Solution Deployment Overview

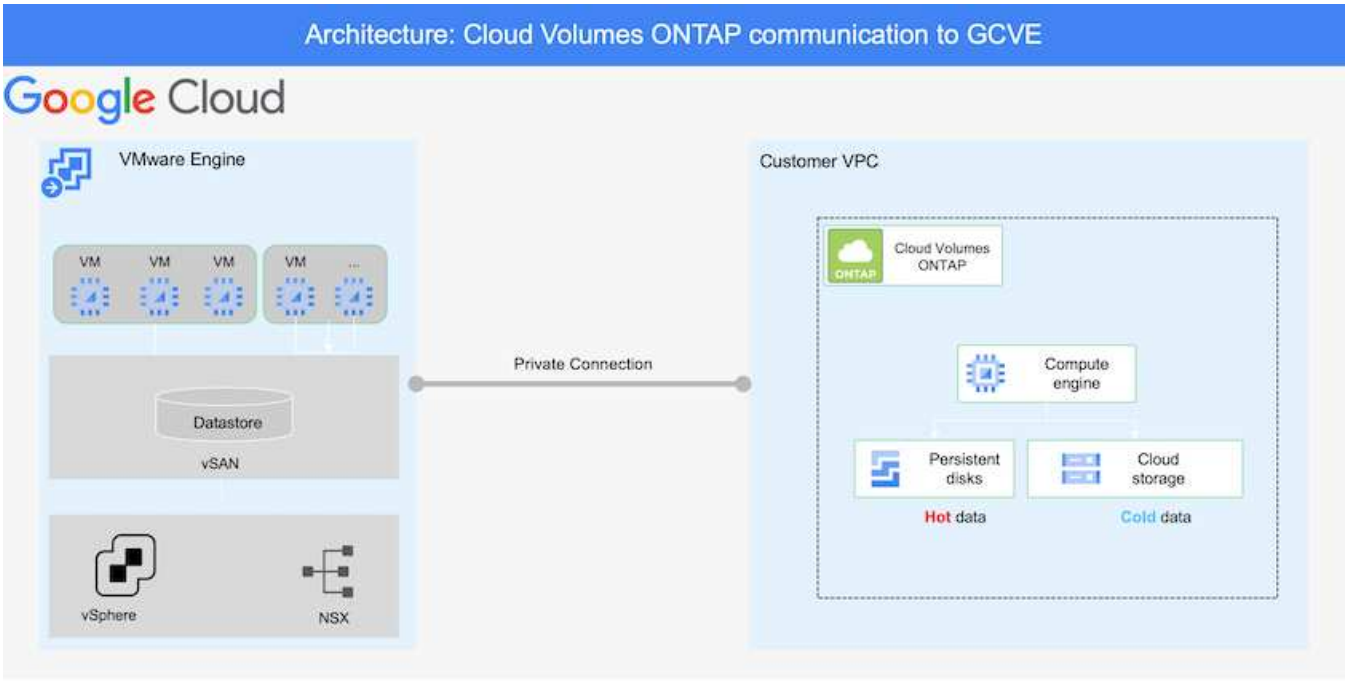
1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
2. Provision Cloud Volumes ONTAP with the correct instance size using Cloud manager within the appropriate subscription and virtual network.
   a. Configure SnapMirror for the relevant application volumes.
   b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
3. Install the Veeam software and start replicating virtual machines to Google Cloud VMware Engine instance.

4. During a disaster event, break the SnapMirror relationship using Cloud Manager and trigger failover of virtual machines with Veeam.
   a. Reconnect the ISCSI LUNs and NFS mounts for the application VMs.
   b. Bring up applications online.
5. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

Deployment Details

▼ Configure CVO on Google Cloud and replicate volumes to CVO

The first step is to configure Cloud Volumes ONTAP on Google Cloud (cvo) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.



For sample step-by-step instructions on setting up SnapCenter and replicating the data, Refer to Setup Replication with SnapCenter

Setup Replication with SnapCenter

▼ Configure GCVE hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the GCVE solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

Cloud Volumes ONTAP can be deployed to any VPC and GCVE should have private connection to that VPC to have VM connect to iSCSI LUNs.

To configure GCVE SDDC, see Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP). As a prerequisite, verify that the guest VMs residing on the GCVE hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and GCVE have been configured properly, begin configuring Veeam to

automate the recovery of on-premises workloads to GCVE (VMs with application VMDKs and VMs with in-guest storage) by using the Veeam Replication feature and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

▼ **Install Veeam Components**

Based on deployment scenario, the Veeam backup server, backup repository and backup proxy that needs to be deployed. For this use case, there is no need to deploy object store for Veeam and Scale-out repository also not required.
Refer to the Veeam documentation for the installation procedure

▼ **Setup VM Replication with Veeam**

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. Setup vSphere VM Replication Job At the Guest Processing step of wizard, select disable application processing as we will be utilizing SnapCenter for application aware backup and recovery.

Setup vSphere VM Replication Job

▼ **Failover of Microsoft SQL Server VM**

Failover of Microsoft SQL Server VM

Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
  - This avoids the risk of data corruption for volumes and snapshots.
  - Avoids replication interruptions during DR test workflows.
  - Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- Veeam Replication allows changing VM IP addresses on DR site.

Using Veeam Replication and Google Cloud NetApp Volumes datastore for disaster recovery to Google Cloud VMware Engine

A comprehensive disaster recovery plan is critical for businesses in times of crisis. Many organizations leverage cloud computing for daily operations and disaster recovery. This proactive approach can reduce or eliminate expensive business disruptions.

This article describes how to use Veeam Backup & Replication to set up disaster recovery for on-premises VMware VMs to Google Cloud VMware Engine (GCVE) with Google Cloud NetApp Volumes (NetApp Volumes).

Overview

Google Cloud NetApp Volumes is a storage service from Google and NetApp that is available for Google Cloud. NetApp Volumes service provides high performance NFS/SMB storage. VMware certified NetApp Volumes NFS storage can be used as an external datastore for ESXi hosts in GCVE. Users are required to make a peering connection between their GCVE private cloud and NetApp Volumes project. There are no

network charges resulting from storage access within a region. Users can create NetApp Volumes volumes in the Google Cloud console and enable deletion protection before mounting volumes as datastores to their ESXi hosts.

NetApp Volumes based NFS datastores can be used to replicate data from on-premises using any validated third-party solution that provides VM replication capability. By adding NetApp Volumes datastores, it enables cost optimized deployment instead of building a Google Cloud VMware Engine (GCVE) based SDDC with a large number of ESXi hosts to accommodate the storage. This approach is called a "Pilot Light Cluster". A pilot light cluster is a minimal GCVE host configuration (3 x GCVE ESXi hosts) along with NetApp Volumes datastores capacity to allow for independent scaling to meet capacity requirements.

The objective is to sustain a cost-effective infrastructure with just the core components to manage a failover. A pilot light cluster can expand and add more GCVE hosts in the event of a failover. Once the failover is resolved and normal operations resume, the pilot light cluster can reduce its scale, returning to a low-cost operational mode.

Purposes of this document

This article describes how to use a Google Cloud NetApp Volumes datastore with Veeam Backup & Replication to set up disaster recovery for on-premises VMware VMs to GCVE using the Veeam VM replication software functionality.

Veeam Backup & Replication is a backup and replication application for virtual environments. When virtual machines are replicated, Veeam Backup & Replication will create an exact copy of the VMs in the native VMware vSphere format on the target GCVE SDDC cluster. Veeam Backup & Replication will keep the copy synchronized with the original VM. Replication provides the best recovery time objective (RTO) as there is a mounted copy of a VM at the DR site in a ready-to-start state.

This replication mechanism ensures that the workloads can quickly start in GCVE in the case of a disaster event. The Veeam Backup & Replication software also optimizes traffic transmission for replication over WAN and slow connections. In addition, it also filters out duplicate data blocks, zero data blocks, swap files, and "excluded VM guest OS files". The software will also compress the replica traffic. To prevent replication jobs from consuming the entire network bandwidth, WAN accelerators and network throttling rules can be utilized.

The replication process in Veeam Backup & Replication is job driven which means replication is performed by configuring replication jobs. In the case of a disaster event, failover can be triggered to recover the VMs by failing over to its replica copy. When failover is performed, a replicated VM takes over the role of the original VM. Failover can be performed to the latest state of a replica or to any of its known good restore points. This enables ransomware recovery or isolated testing as needed. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.

Solution Overview

This solution covers the following high level steps:

1. Create an NFS volume using Google Cloud NetApp Volumes
2. Follow GCP process to create a GCVE datastore from the NetApp Volumes NFS volume.
3. Set up a replication job to create VM replicas using Veeam Backup & Replication.
4. Create a failover plan and perform failover.
5. Switch back to production VMs once the disaster event is complete and the primary site is up.

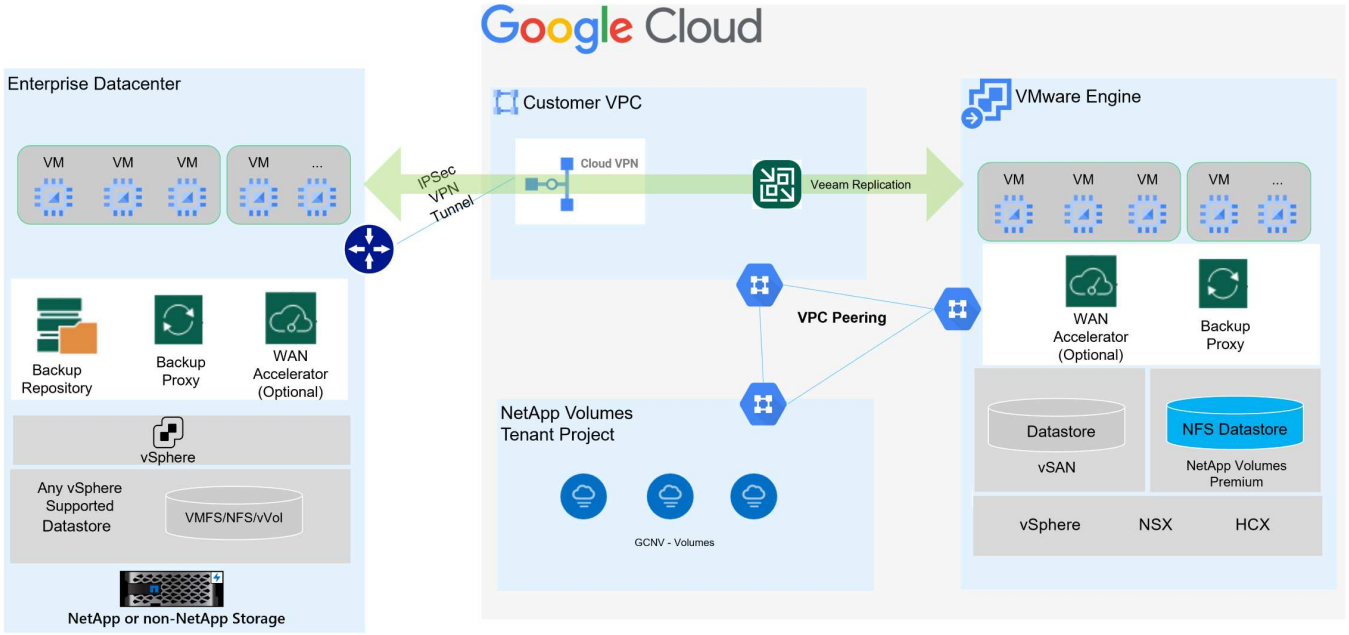


When creating a volume in NetApp Volumes, for use as a GCVE datastore, only NFS v3 is supported.

For more information on using NetApp Volumes NFS volumes as datastores for GCVE, check out Using NFS volume as vSphere datastore hosted by Google Cloud NetApp Volumes.

Architecture

The following diagram shows the architecture of the solution presented in this documentation. A recommended best practice is to have a Veeam Backup & Replication server located at both the on-premises site and in the GCVE SDDC. Backup and recovery is performed and managed by the Veeam server on-premises, and replication is managed by the Veeam server in the GCVE SDDC. This architecture provides the highest availability when a failure occurs in the primary datacenter.



Pre-requisites for Veeam Replication to GCVE and NetApp Volumes datastores

This solution requires the following components and configurations:

1. NetApp Volumes has a Storage Pool available with enough free capacity to accommodate the NFS volume to be created.
2. Veeam Backup and Replication software is running in an on-premises environment with appropriate network connectivity.
3. Ensure the Veeam Backup & Replication backup VM is connected to the source as well as the target GCVE SDDC clusters.
4. Ensure the Veeam Backup & Replication backup VM is connected to the Veeam Proxy server VMs at both the source and target GCVE clusters.
5. The backup server must be able to resolve short names and connect to source and target vCenters.

Users are required to make a peering connection between their GCVE private cloud and NetApp Volumes project using the VPC Network peering or Private connections pages within the VMware Engine Cloud console UI.

Veeam requires a GCVE solution user account with elevated privileges when adding the GCVE vCenter server to the Veeam Backup and Replication inventory. For more information refer to the Google Cloud Platform (GCP) documentation, Elevating VMware Engine privileges.

For additional information refer to Considerations and Limitations in the Veeam Backup & Replication documentation.

Deployment steps

The following sections outline the deployment steps to create and mount an NFS datastore using Google Cloud NetApp Volumes, and use Veeam Backup and Replication to implement a full disaster recovery solution between an on-premises datacenter and Google Cloud VMware Engine.

Create NetApp Volumes NFS volume and datastore for GCVE

Refer to Using NFS volume as vSphere datastore hosted by Google Cloud NetApp Volumes for an overview of how to Google Cloud NetApp Volumes as a datastore for GCVE.

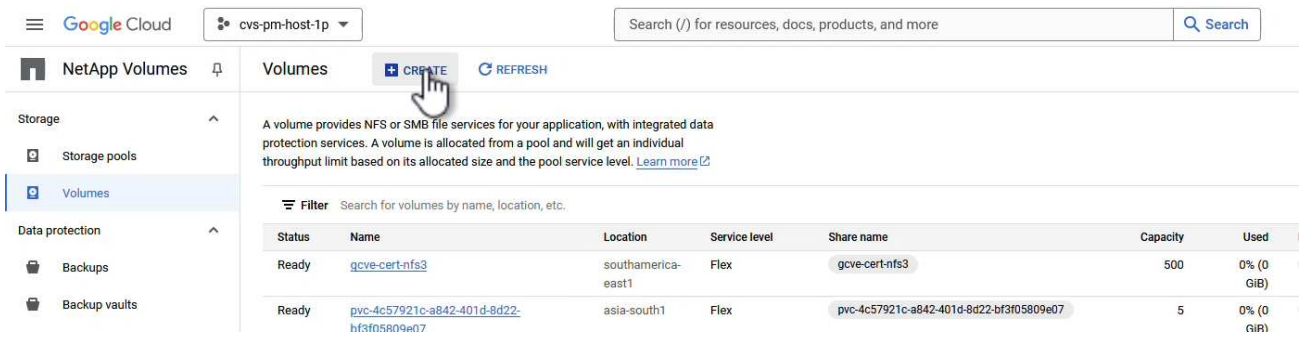
Complete the following steps to create and use an NFS datastore for GCVE using NetApp Volumes:

▼ **Create NetApp Volumes NFS volume**

Google Cloud NetApp Volumes is accessed from the Google Cloud Platform (GCP) console.

Refer to Create a volume in the Google Cloud NetApp Volumes documentation for detailed information on this step.

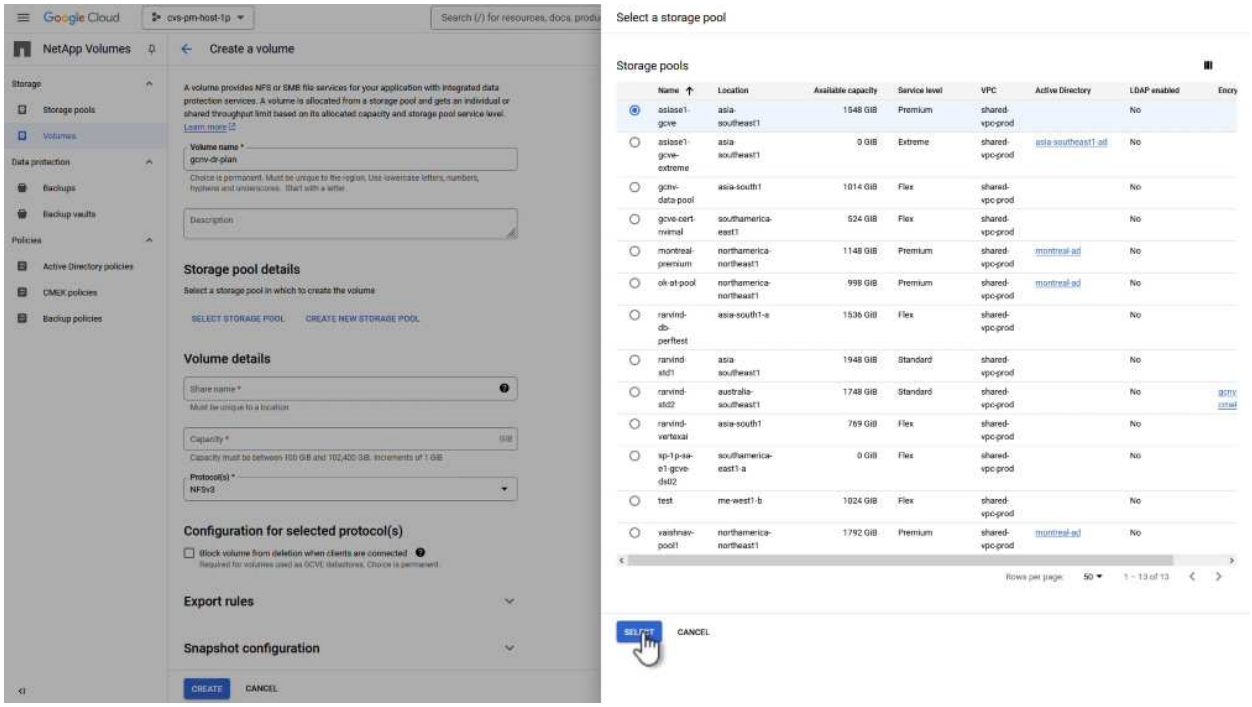
1. In a web browser, navigate to https://console.cloud.google.com/ and log into your GCP console. Search for **NetApp Volumes** to get started.
2. In the **NetApp Volumes** management interface, click on **Create** to get started creating an NFS volume.

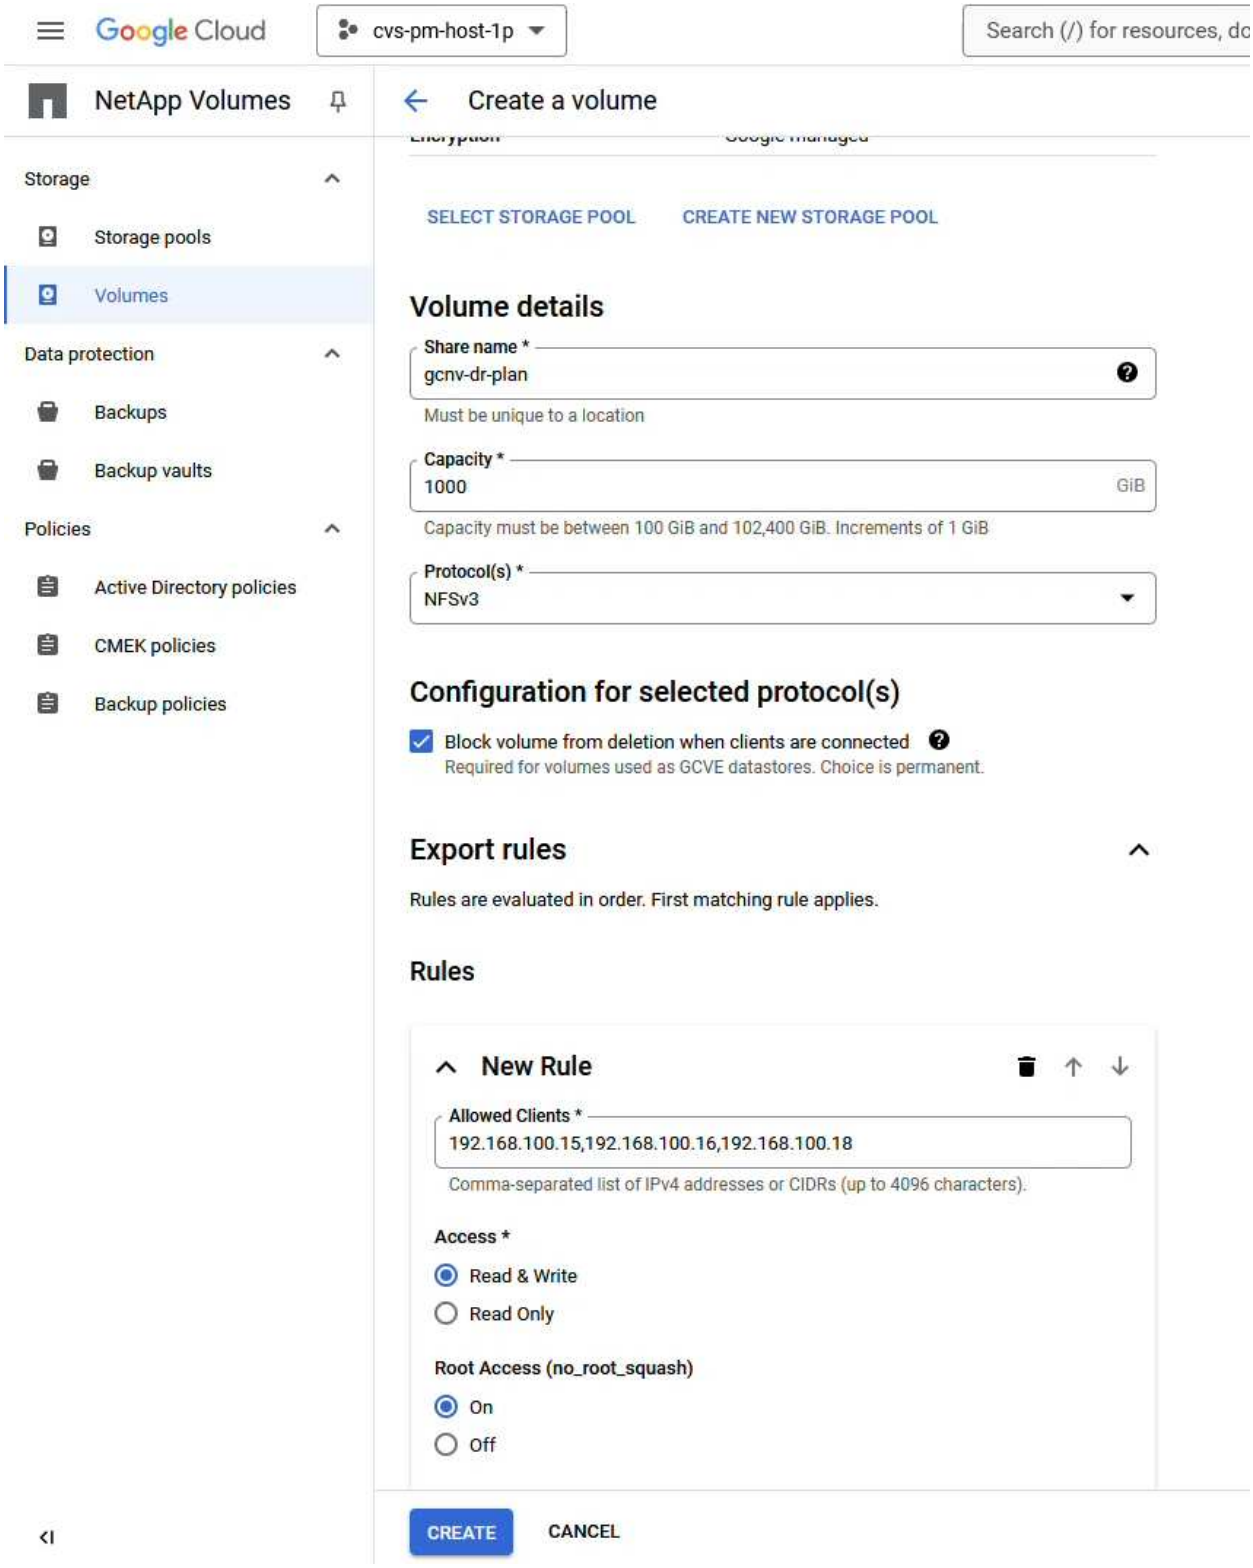


3. In the **Create a volume** wizard, fill out all required information:
   - A name for the volume.
   - The Storage Pool on which to create the volume.
   - A share name used when mounting the NFS volume.
   - The capacity of the volume in GiB.
   - The storage protocol to be used.
   - Check the box to **Block volume from deletion when clients are connected** (required by GCVE when mounting as a datastore).
   - The export rules for accessing the volume. This is the IP addresses of the ESXi adapters on the NFS network.
   - A snapshot schedule used to protect the volume using local snapshots.

◦ Optionally, choose to backup the volume and/or create labels for the volume.

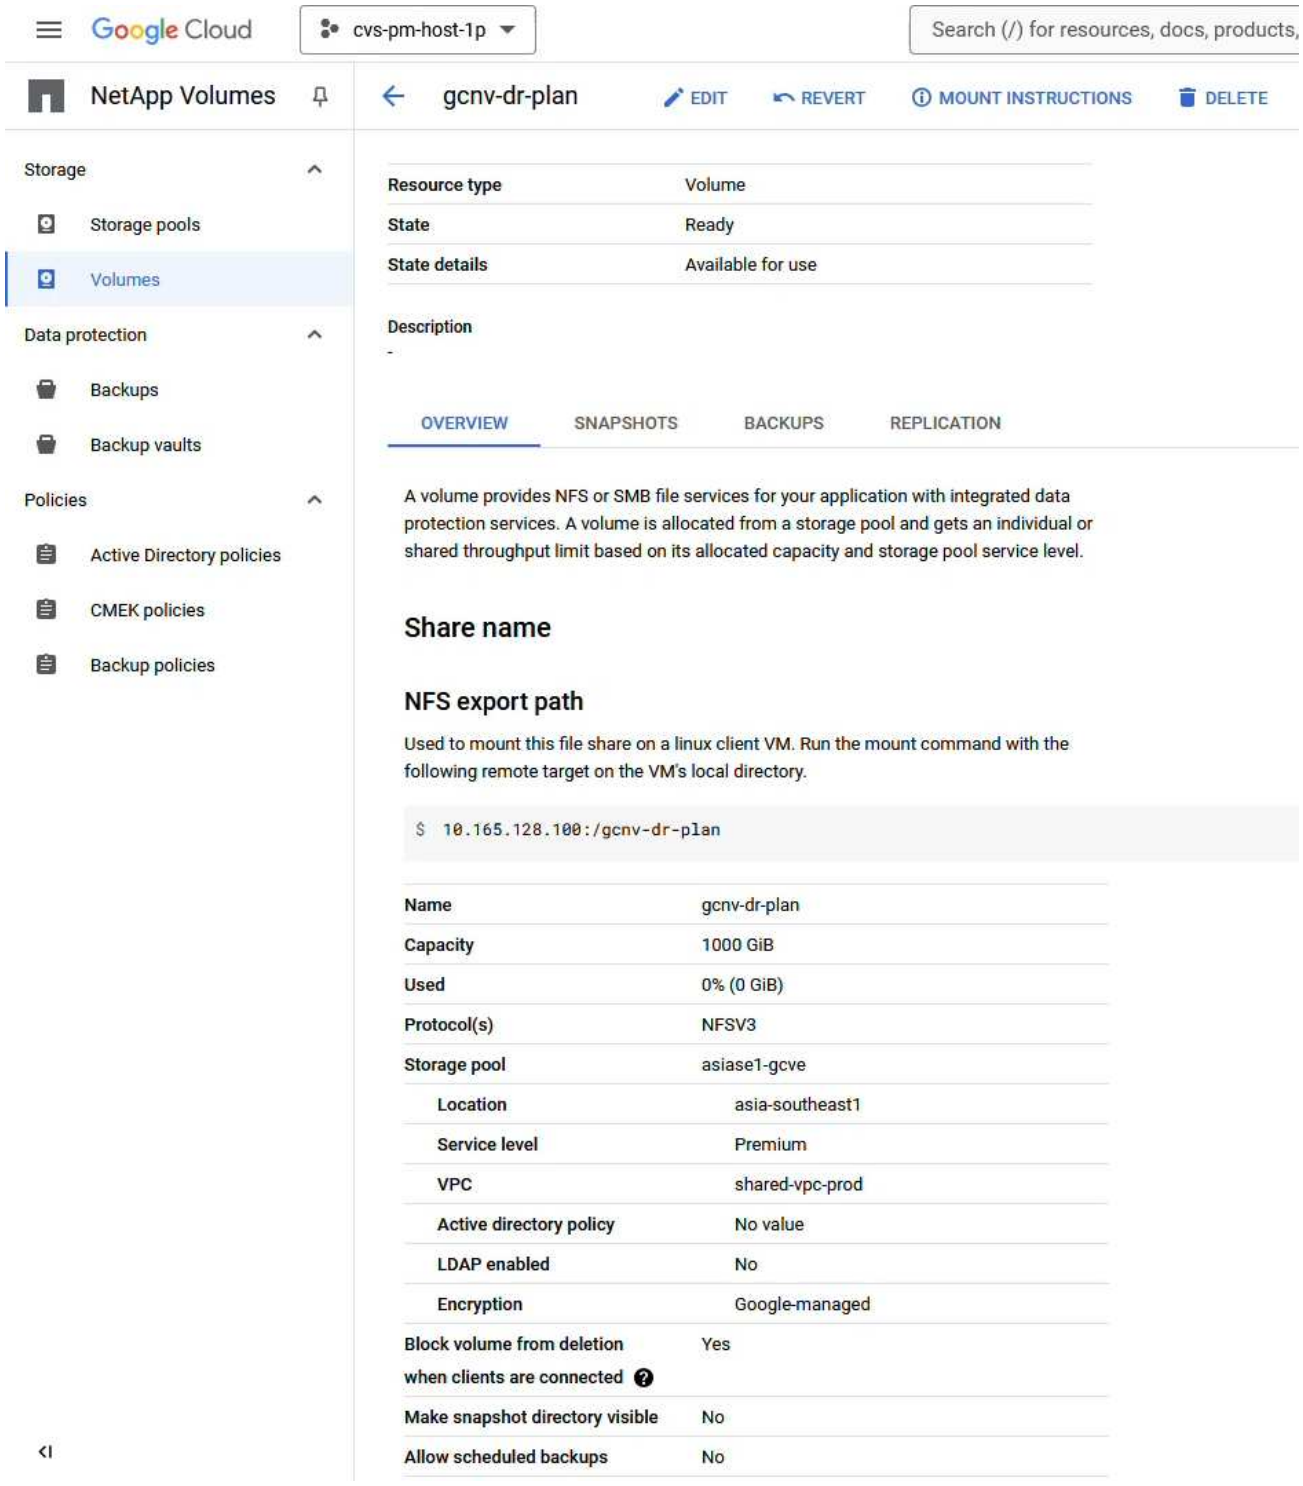
When creating a volume in NetApp Volumes, for use as a GCVE datastore, only NFS v3 is supported.





Click on **Create** to finish creating the volume.

4. Once the volume is created, the NFS export path required to mount the volume can be viewed from the volume's properties page.



▼ **Mount the NFS datastore in GCVE**

At the time of this writing the process to mount a datastore in GCVE requires opening a GCP support ticket to have the volume mounted as an NFS datastore.

Refer to Using NFS volume as vSphere datastore hosted by Google Cloud NetApp Volumes for more information.

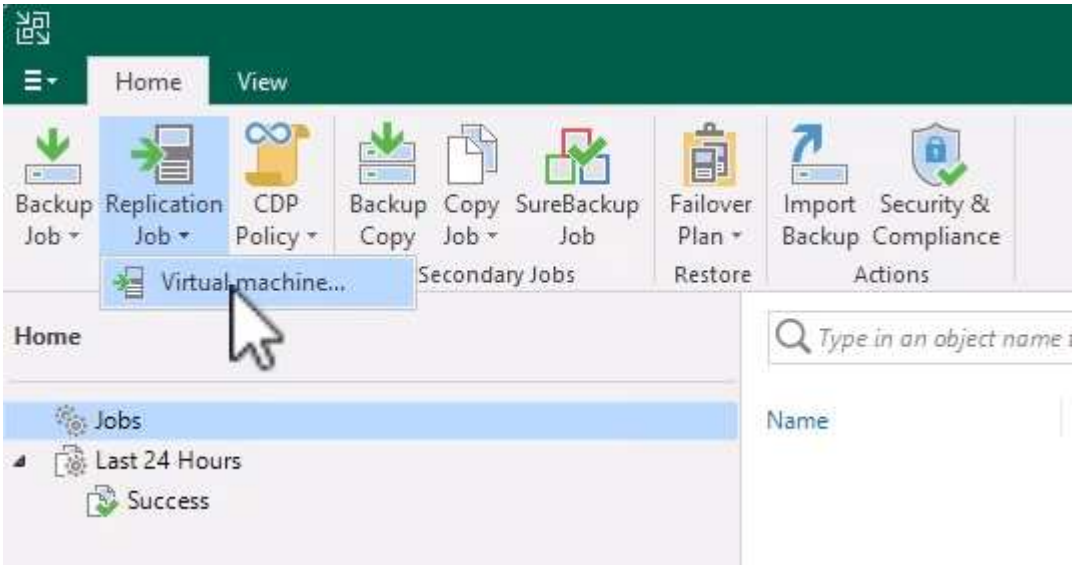
Replicate VMs to GCVE and execute failover plan, and failback

▼ **Replicate VMs to NFS datastore in GCVE**

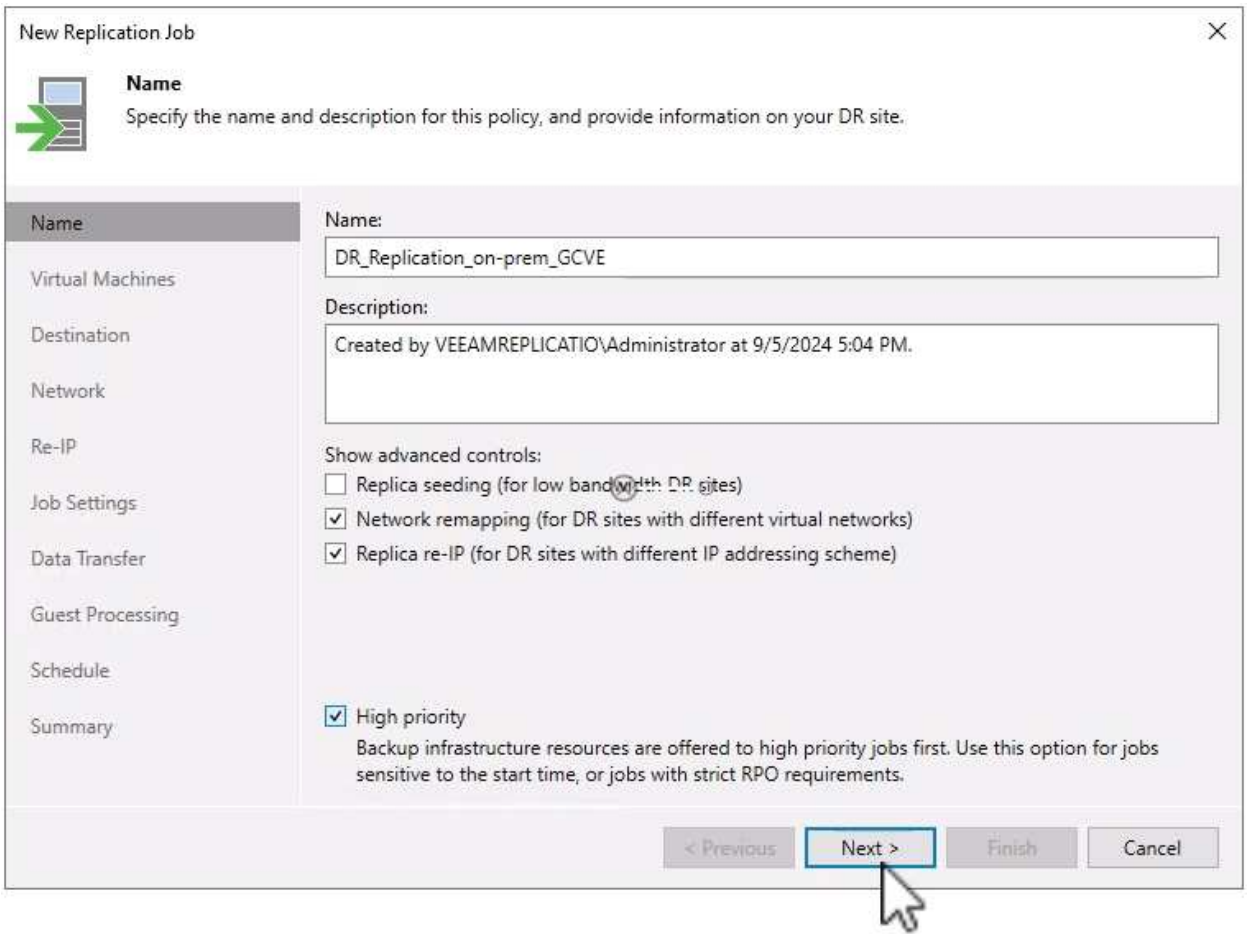
Veeam Backup & Replication leverages VMware vSphere snapshot capabilities during replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot is the point-in-time copy of a VM that includes virtual disks, system state, configuration and metadata. Veeam Backup & Replication uses the snapshot as a source of data for replication.

To replicate VMs, complete the following steps:

1. Open the Veeam Backup & Replication Console.
2. On the **Home** tab, click on **Replication Job > Virtual machine…**

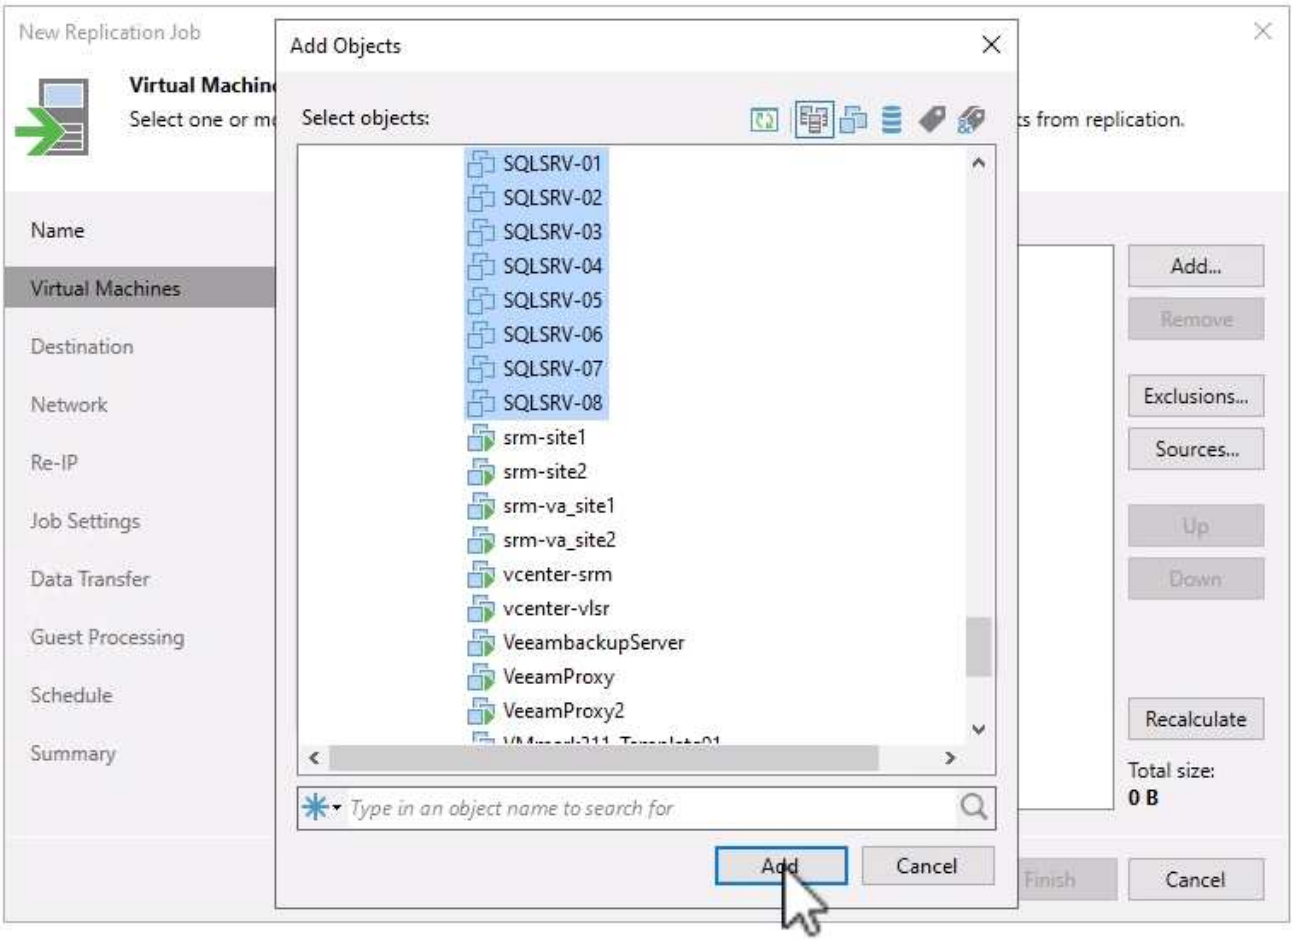


3. On the **Name** page of the **New Replication Job** wizard, specify a job name and select the appropriate advanced control checkboxes.
   - Select the Replica seeding check box if connectivity between on-premises and GCP has restricted bandwidth.
   - Select the Network remapping (for GCVE SDDC sites with different networks) check box if segments on the GCVE SDDC do not match that of the on-premises site networks.
   - Select the Replica re-IP (for DR sites with different IP addressing scheme) check box if the IP addressing scheme in the on-premises production site differs from the scheme in the target GCVE site.

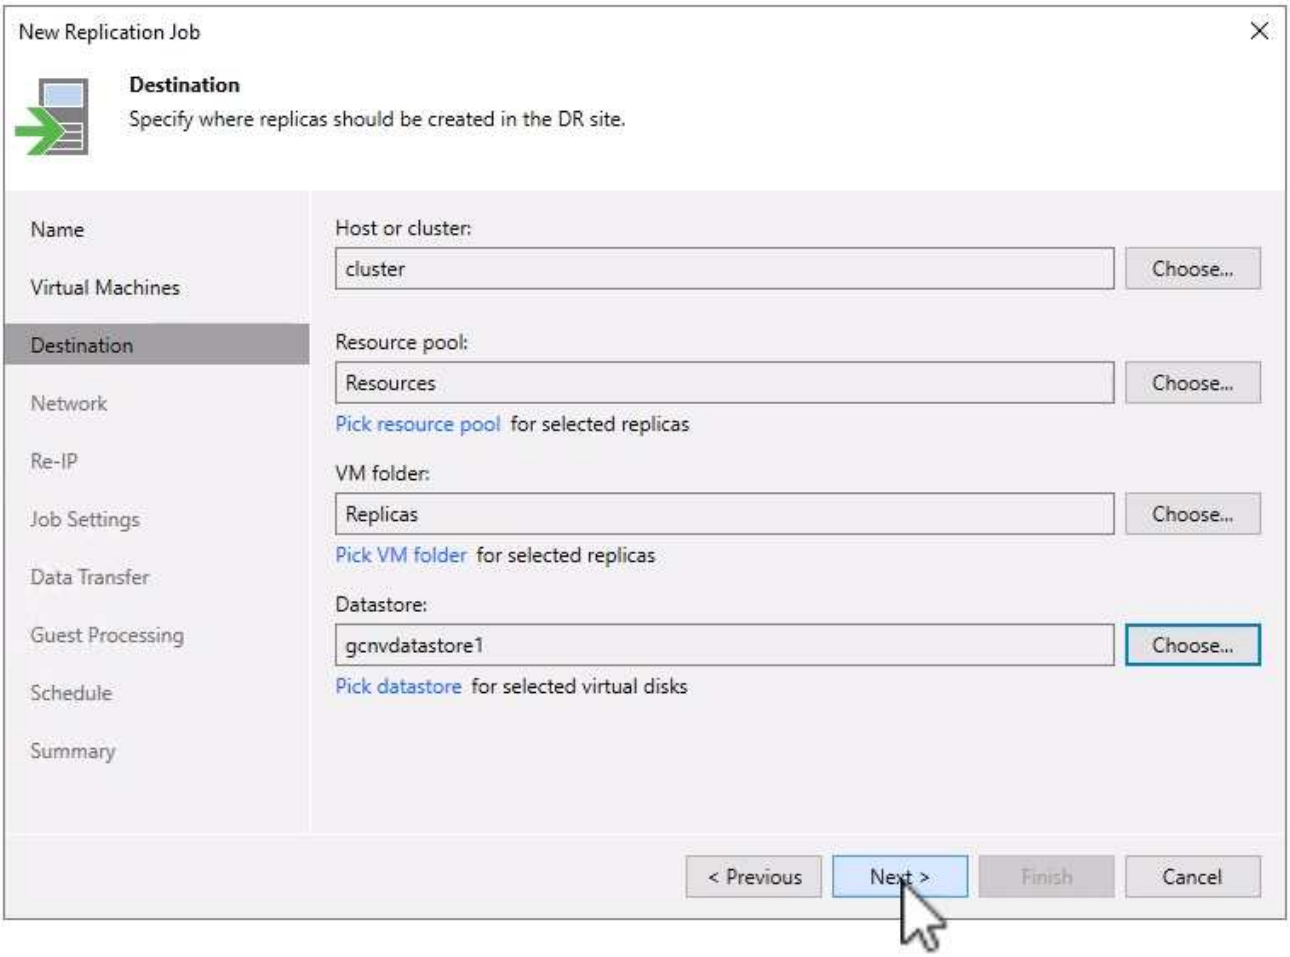


4. On the **Virtual Machines** page, select the VMs to be replicated to the NetApp Volumes datastore attached to a GCVE SDDC. Click **Add**, then in the **Add Object** window select the necessary VMs or VM containers and click **Add**. Click **Next**.

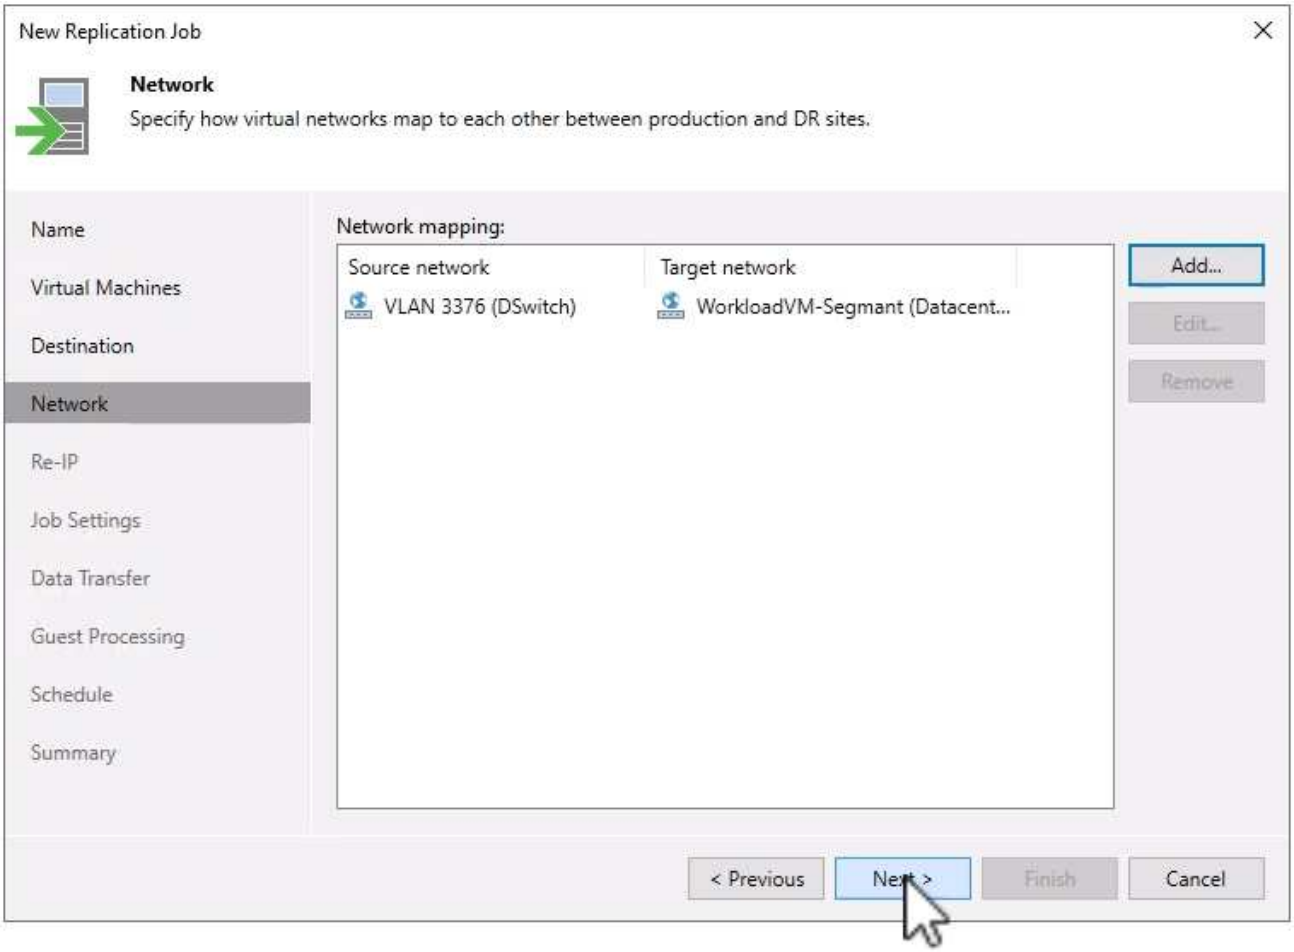
> The Virtual machines can be placed on vSAN to fill the available vSAN datastore capacity. In a pilot light cluster, the usable capacity of a 3-node vSAN cluster will be limited. The rest of the data can be easily placed on Google Cloud NetApp Volumes datastores so that the VMs can be recovered, and the cluster can later be expanded to meet the CPU/mem requirements.



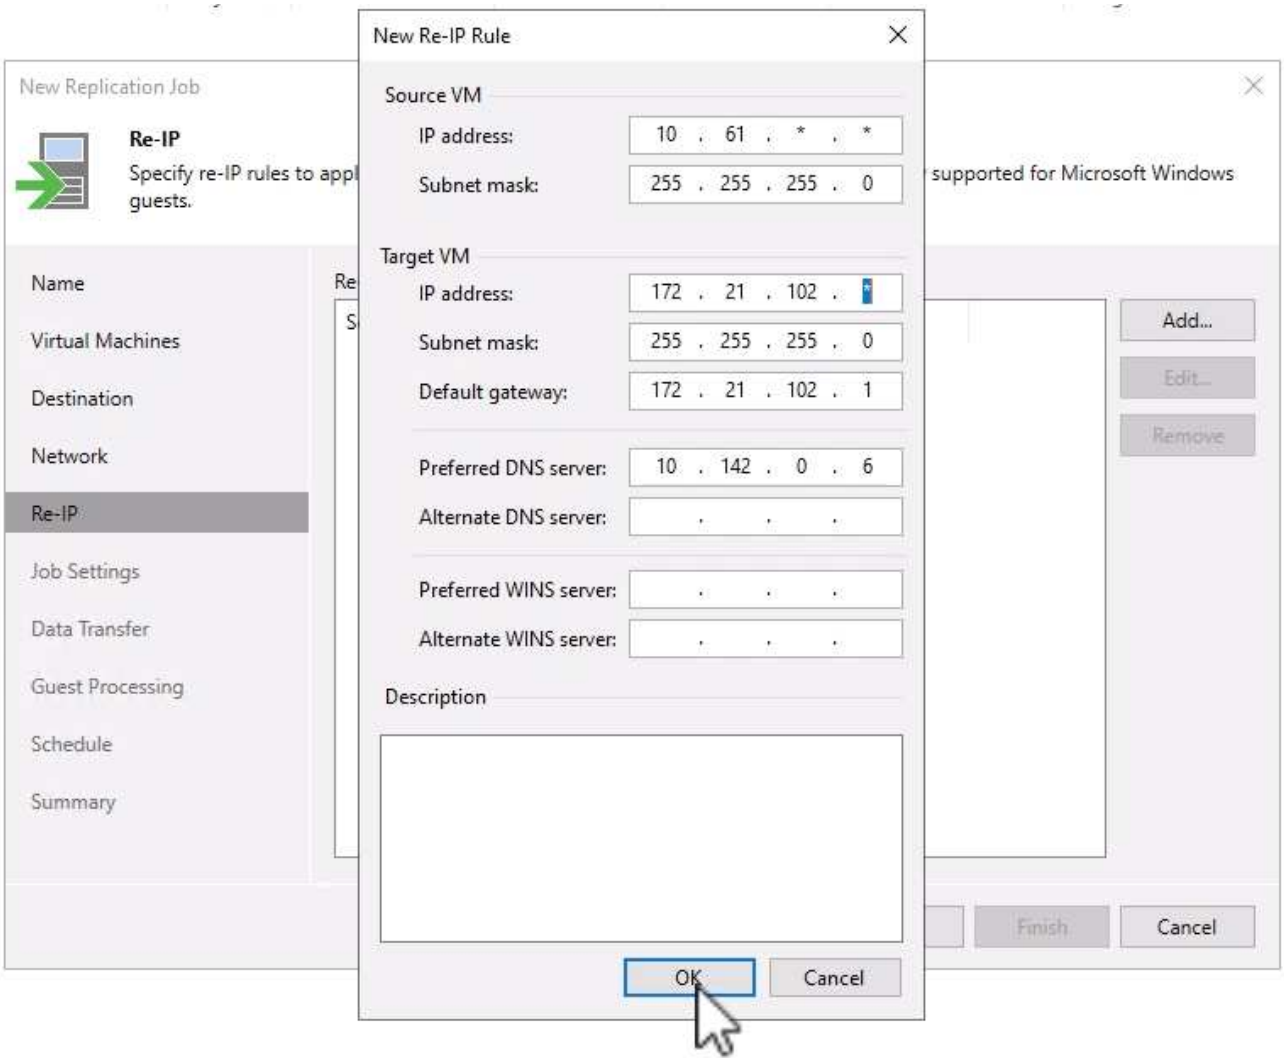
5. On the **Destination** page, select the destination as the GCVE SDDC cluster / hosts and the appropriate resource pool, VM folder and NetApp Volumes datastore for the VM replicas. Click **Next** to continue.



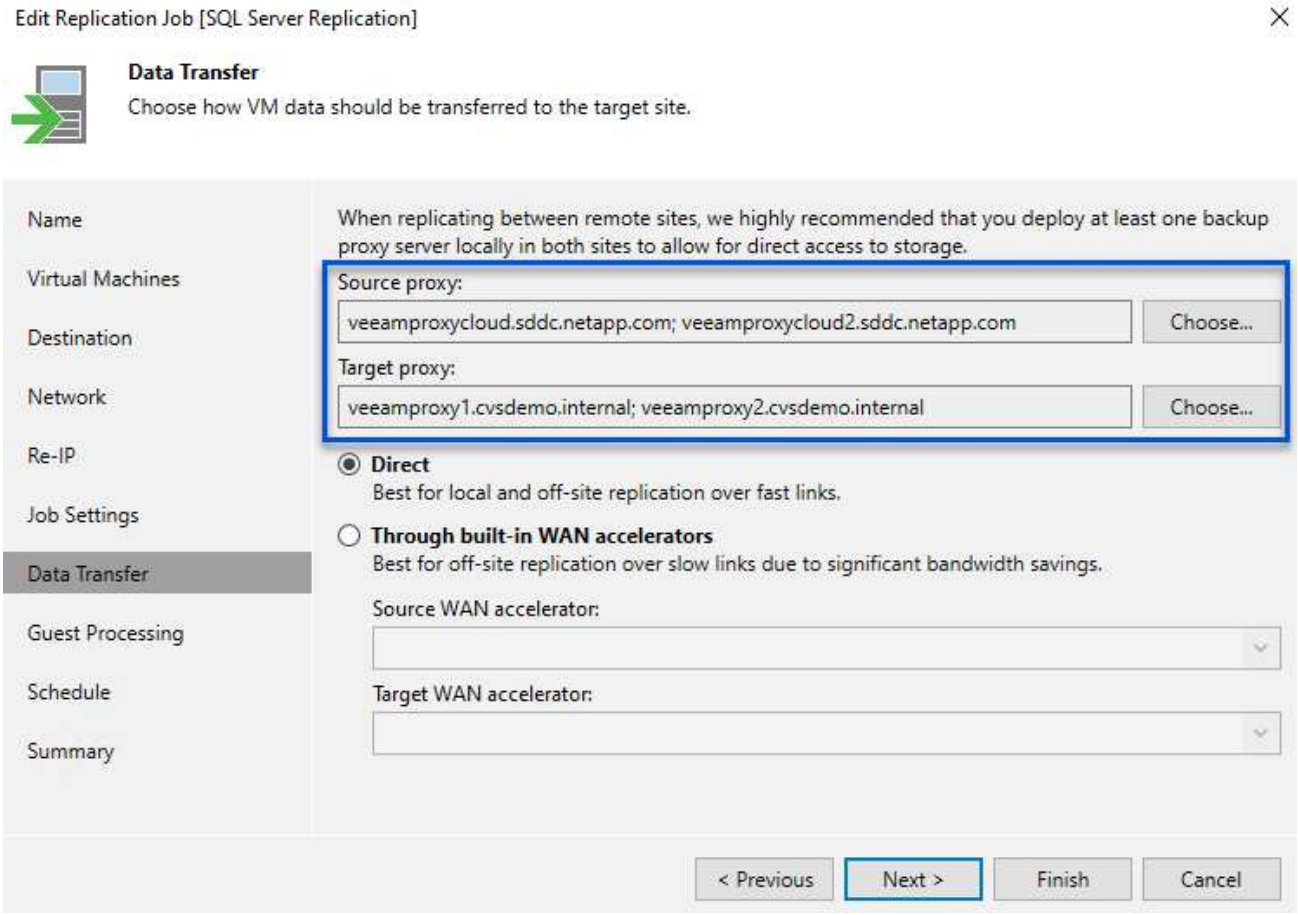
6. On the **Network** page, create the mapping between source and target virtual networks as needed. Click **Next** to continue.



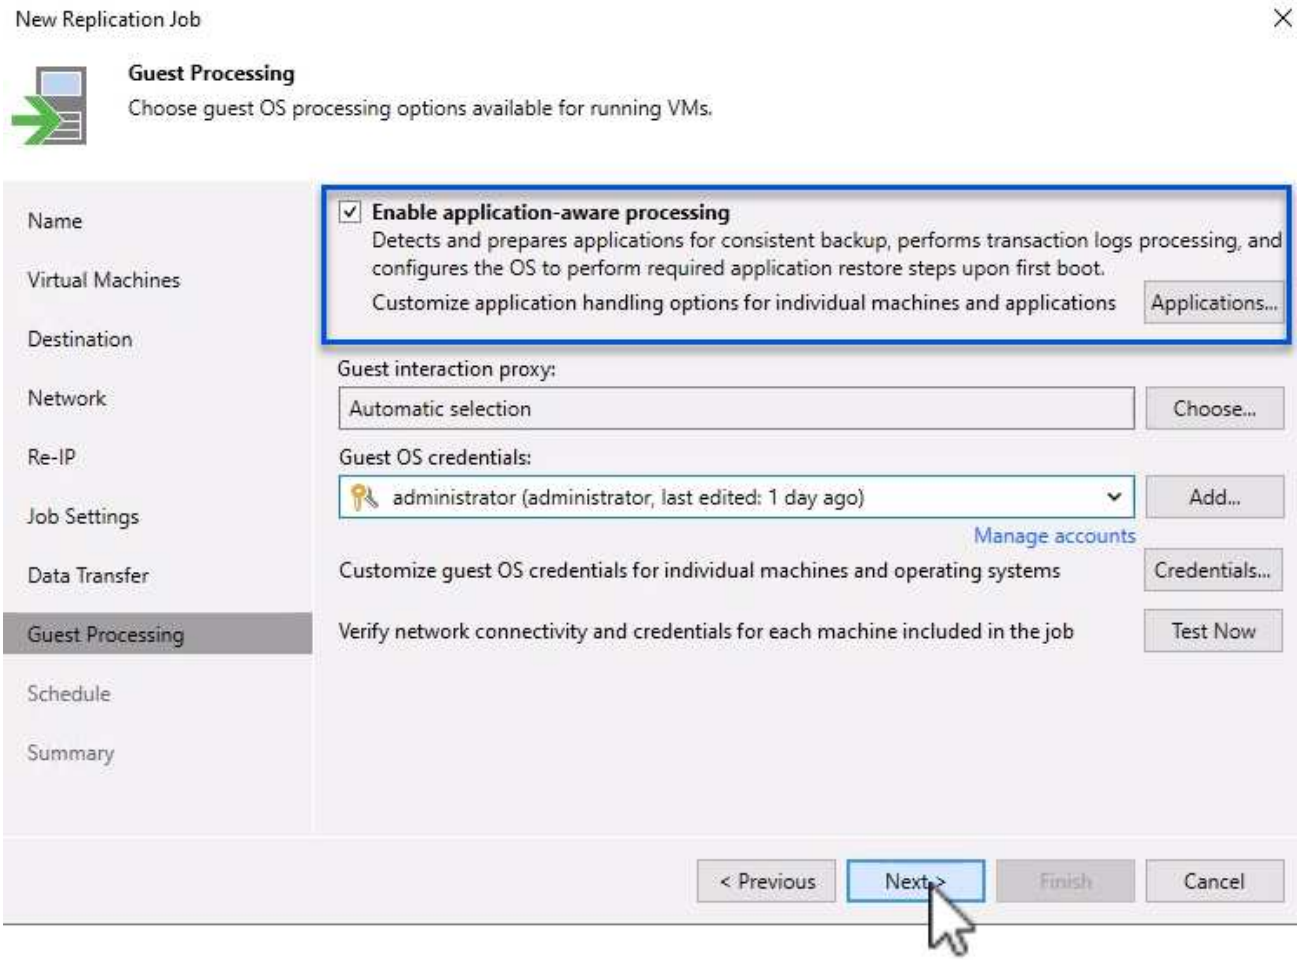
7. On the **Re-IP** page, click on the **Add…** button to add a new re-ip rule. Fill out the source and target VM ip ranges to specify the networking that will be applied to the source VM's in the case of a failover. Use asterisks to specify a range of addresses is indicated for that octet. Click **Next** to continue.



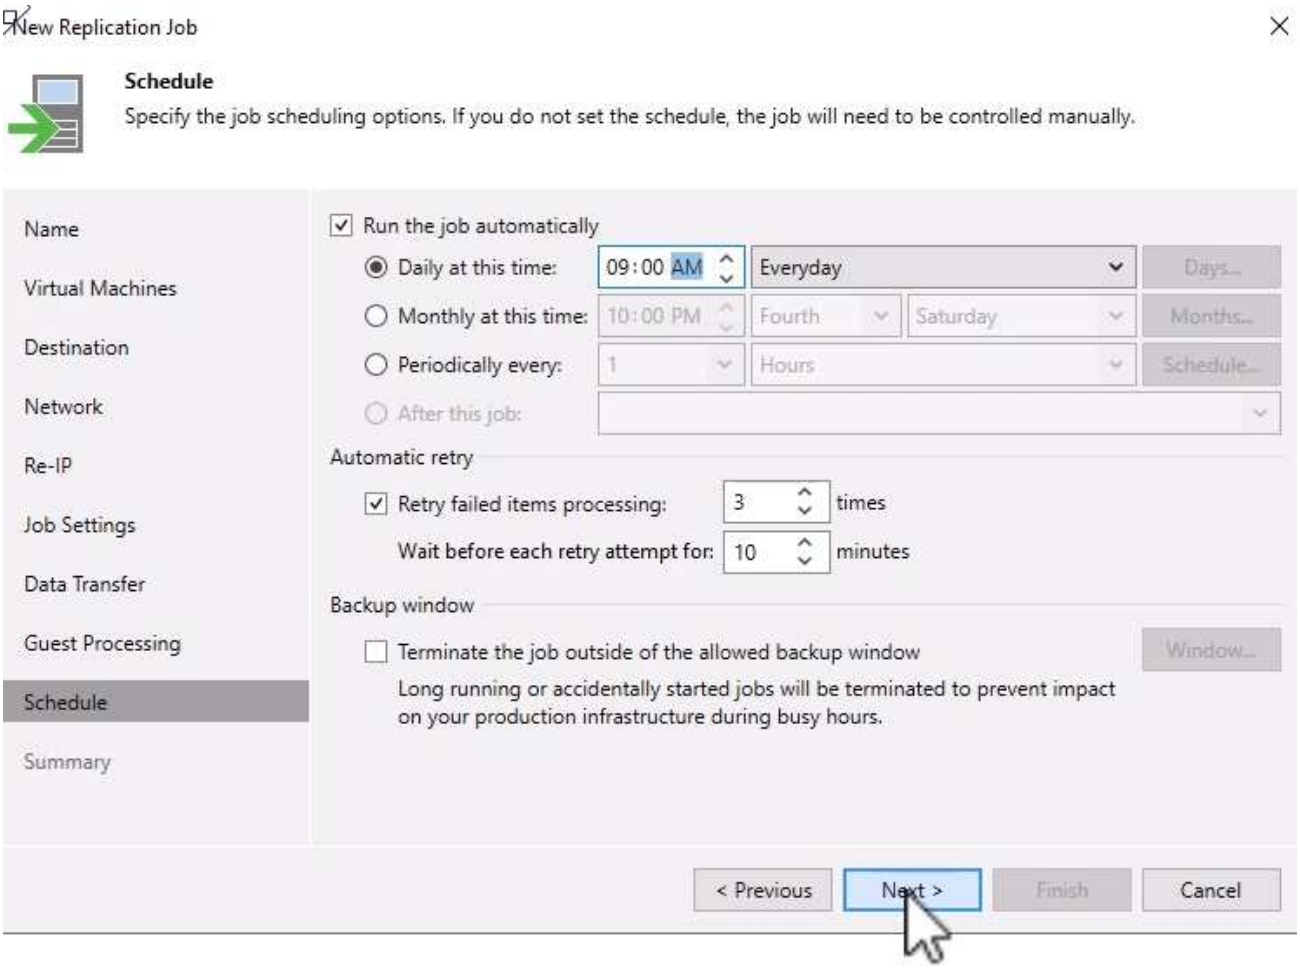
8. On the **Job Settings** page, specify the backup repository that will store metadata for VM replicas, the retention policy and select the button at the bottom for **Advanced…** button at the bottom for additional job settings. Click **Next** to continue.
9. On the **Data Transfer**, select the proxy servers that reside at the source and targets sites, and keep the Direct option selected. WAN accelerators can also be selected here, if configured. Click **Next** to continue.



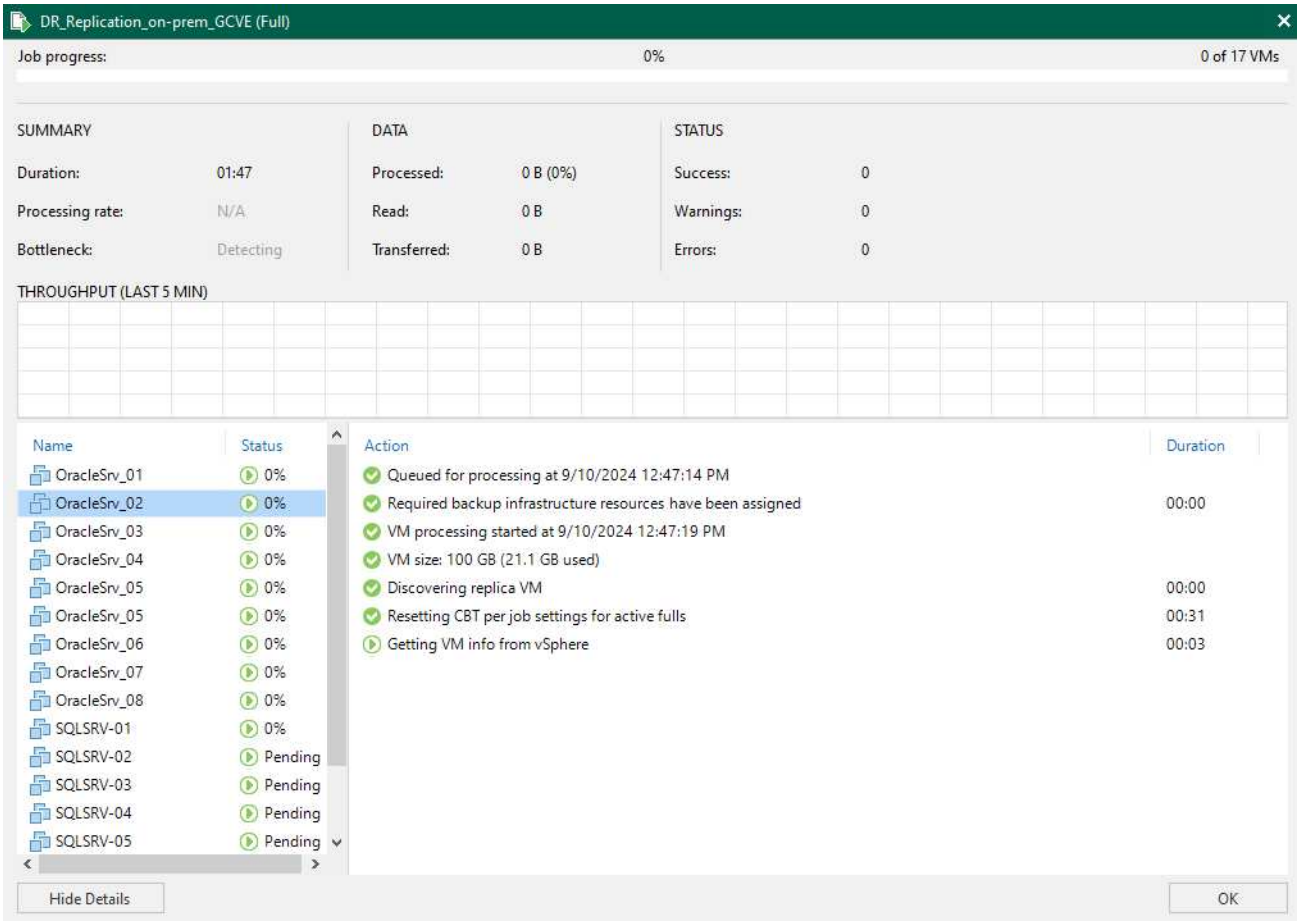
10. On the **Guest Processing** page, check the box for **Enable application-aware processing** as needed and select the **Guest OS credentials**. Click **Next** to continue.



11. On the **Schedule** page, define the times and frequency at which the replication job will run. Click **Next** to continue.



12. Finally, review the job setting on the **Summary** page. Check the box to **Run the job when I click Finish**, and click on **Finish** to complete creating the replication job.
13. Once run, the replication job can be viewed in the job status window.



For additional information on Veeam replication, refer to How Replication Works

▼ **Create a failover plan**

When the initial replication or seeding is complete, create the failover plan. Failover plan helps in performing failover for dependent VMs one by one or as a group automatically. Failover plan is the blueprint for the order in which the VMs are processed including the boot delays. The failover plan also helps to ensure that critical dependent VMs are already running.

After completing the initial replication or seeding, create a failover plan. This plan serves as a strategic blueprint for orchestrating the failover of dependent VMs, either individually or as a group. It defines the processing order of VMs, incorporates necessary boot delays, and ensures that critical dependent VMs are operational before others. By implementing a well-structured failover plan, organizations can streamline their disaster recovery process, minimizing downtime and maintaining the integrity of interdependent systems during a failover event.

When creating the plan, Veeam Backup & Replication automatically identifies and uses the most recent restore points to initiate the VM replicas.

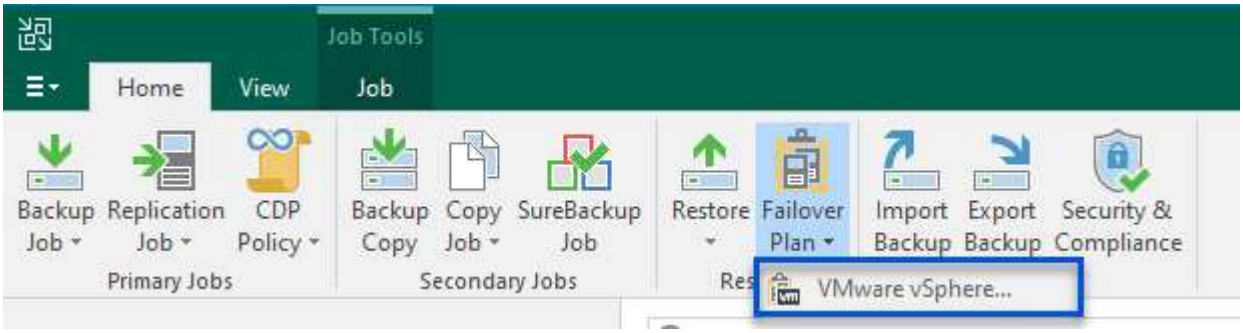
ⓘ The failover plan can only be created once the initial replication is complete and the VM replicas are in Ready state.

ⓘ The maximum number of VMs that can be started simultaneously when running a failover plan is 10.

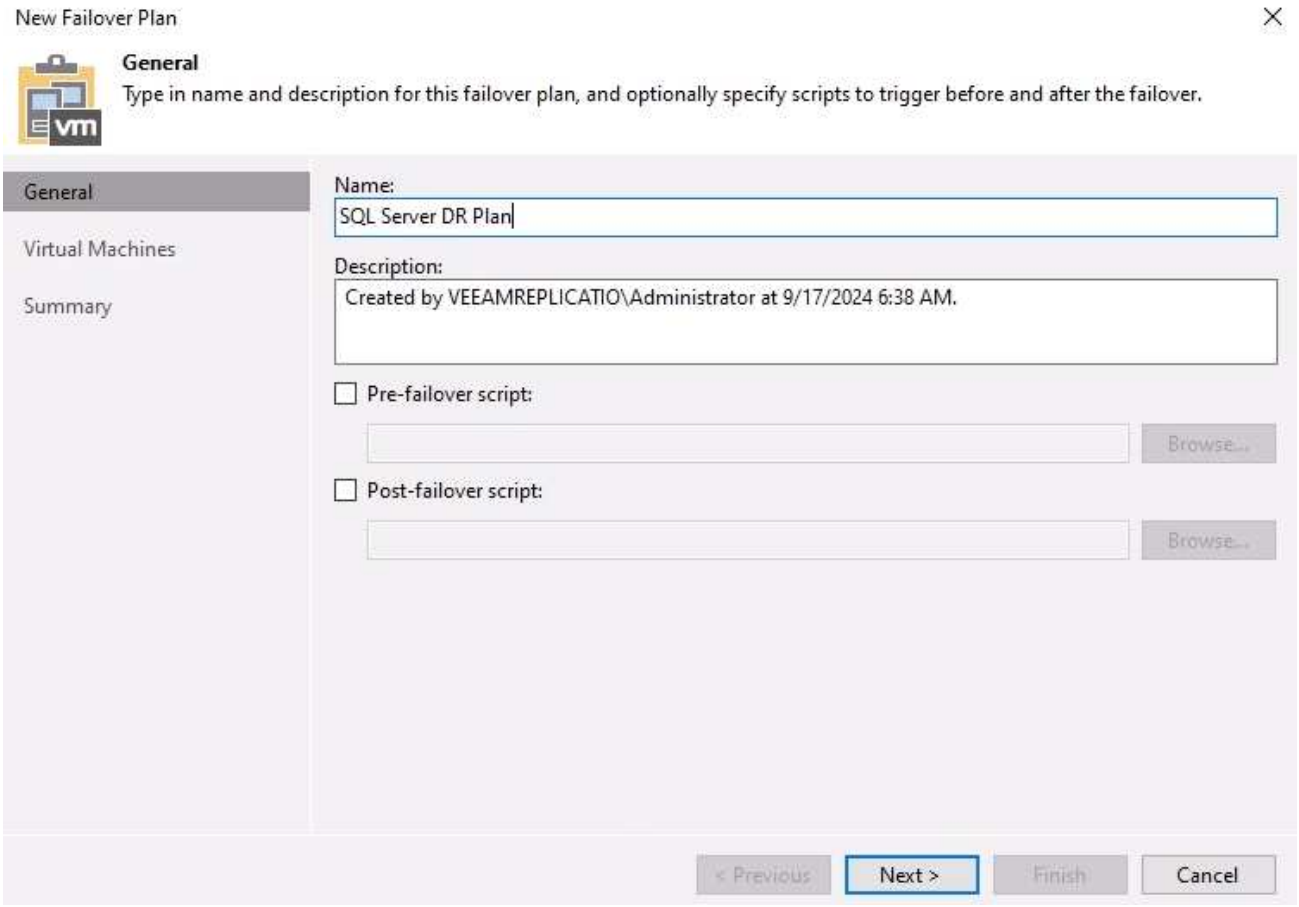
ⓘ During the failover process, the source VMs will not be powered off.

To create the **Failover Plan**, complete the following steps:

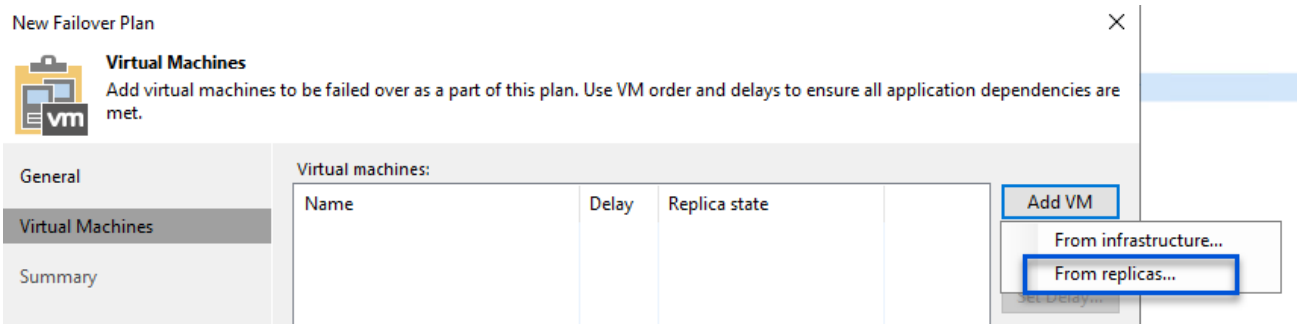
1. On the **Home** view, Click on the **Failover Plan** button in the **Restore** section. In the drop down, select **VMware vSphere…**

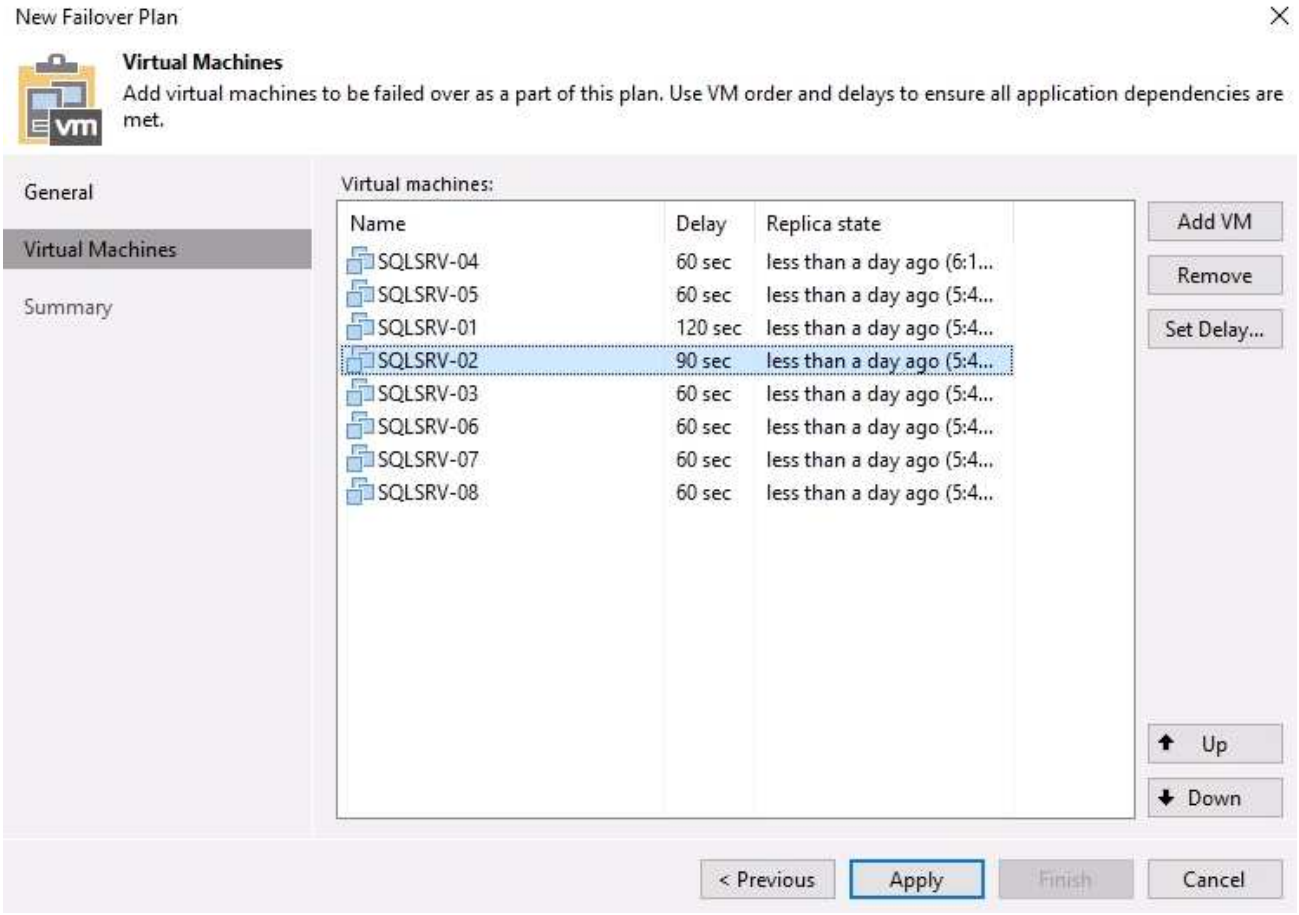


2. On the **General** page of the **New Failover Plan** wizard, provide a name and a description to the plan. Pre and Post-failover scripts can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.



3. On the **Virtual Machines** page, click the button to **Add VM** and select **From replicas…**. Choose the VMs to be part of the failover plan, and then modify the VM boot order and any required boot delays to meet application dependencies.





Click on **Apply** to continue.

4. Finally review all of the failover plan settings and click on **Finish** to create the failover plan.

For additional information on creating replication jobs, refer to Creating Replication Jobs.

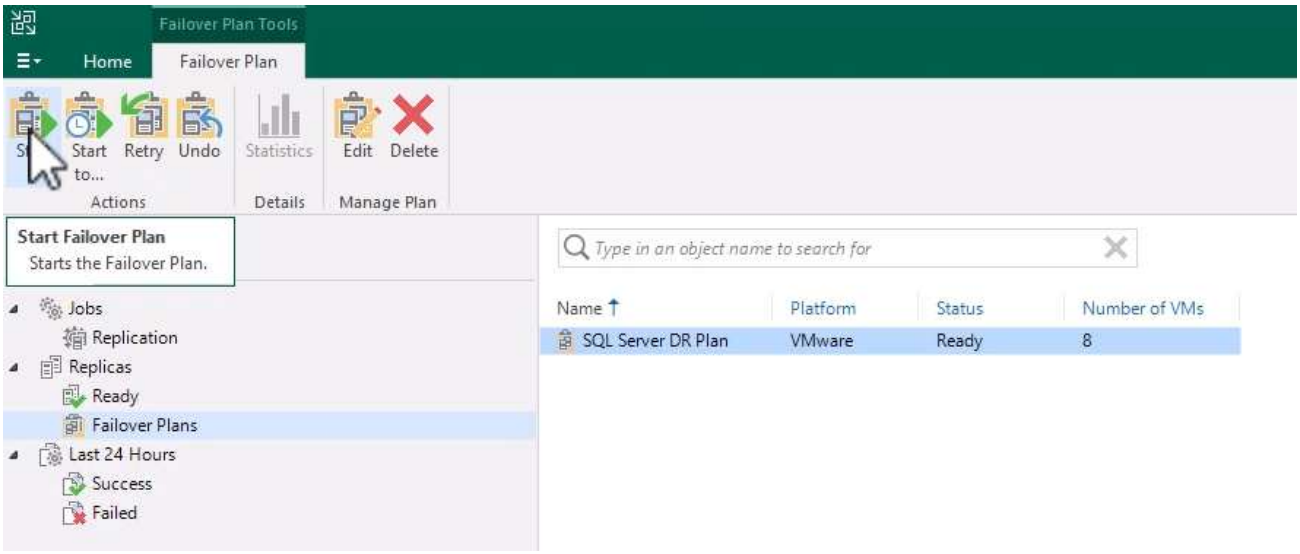
▼ **Run the failover plan**

During failover, the source VM in the production site switches over to its replica at the disaster recovery site. As part of the process, Veeam Backup & Replication restores the VM replica to the required restore point and transfers all I/O activities from the source VM to its replica. Replicas serve not only for actual disasters but also for simulating DR drills. In failover simulation, the source VM continues running. Upon completion of necessary tests, the failover can be undone, returning operations to normal.

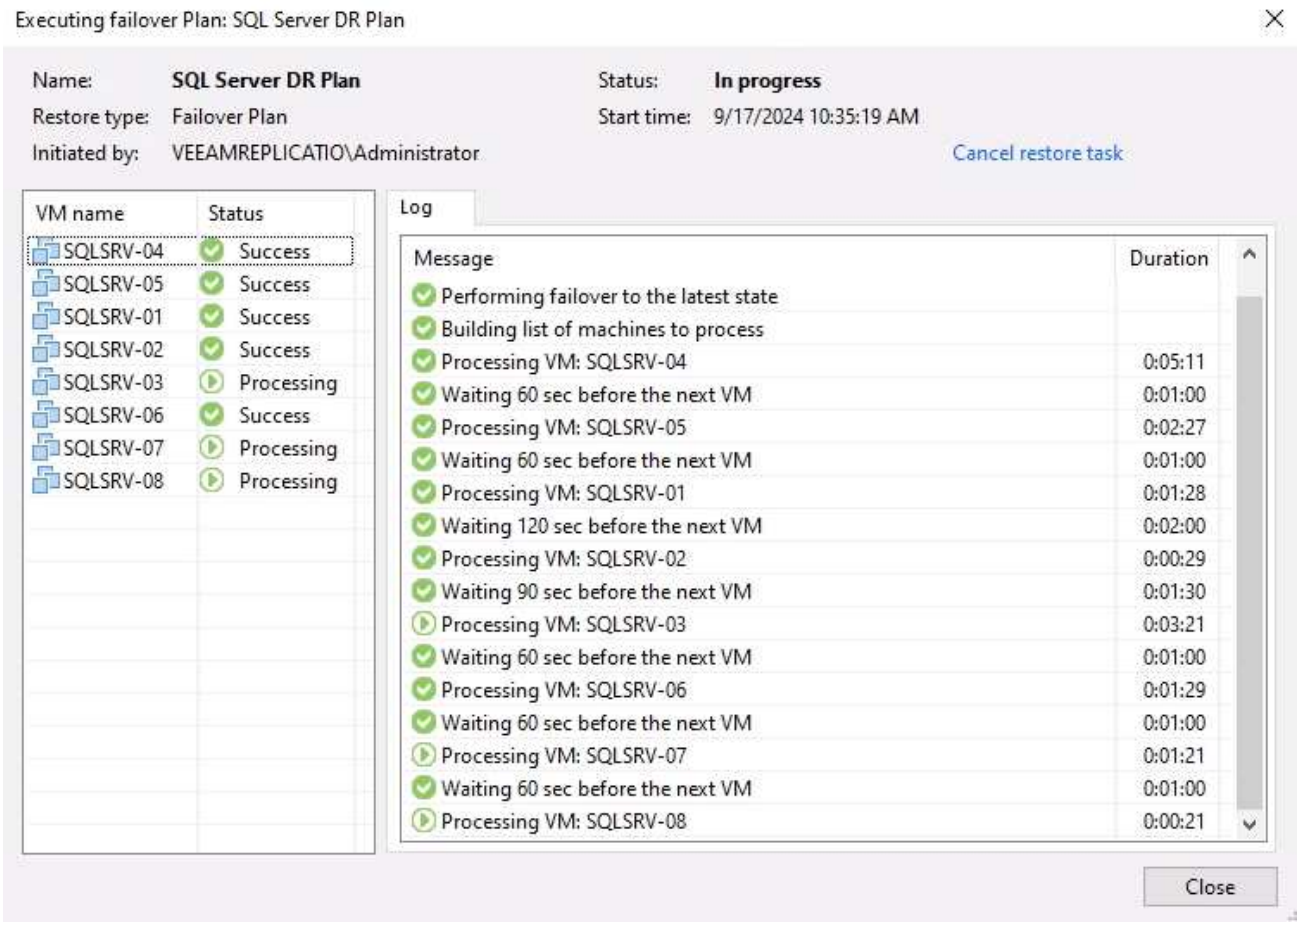
Make sure network segmentation is in place to avoid IP conflicts during failover.

Complete the follow steps to start the failover plan:

1. To get started, in the **Home** view, click on **Replicas > Failover Plans** in the left-hand menu and then on the **Start** button. Alternately, the **Start to…** button can be used to failover to a prior restore point.



2. Monitor the progress of the failover in the **Executing failover plan** window.



Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the Ready state.

For detailed information about failover plans, refer Failover Plans.

▼ **Failback to the production site**

Conducting a failover is considered an intermediate step and needs to be finalized based on the requirement. The options include the following:

- **Failback to production** - Revert to the original VM and synchronize all modifications made during the replica's active period back to the source VM.

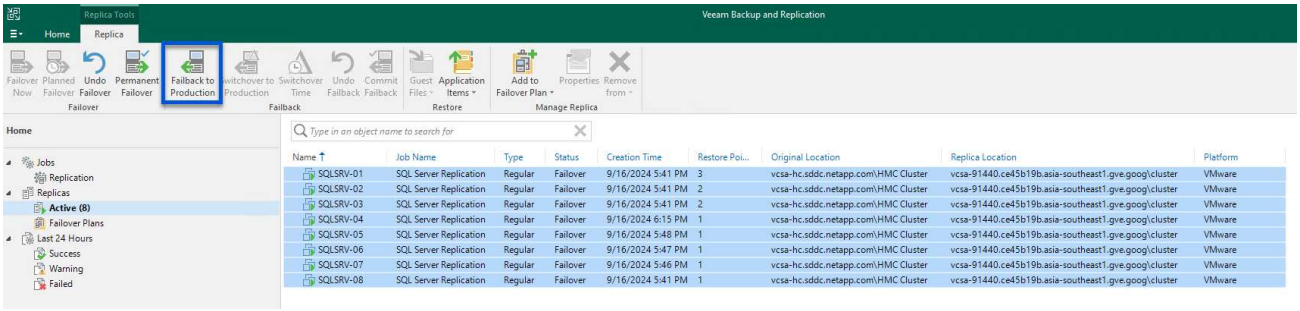
During failback, changes are transferred but not immediately applied. Select **Commit failback** once the original VM's functionality is verified. Alternatively, choose **Undo failback** to revert to the VM replica if the original VM exhibits unexpected behavior.

- **Undo failover** - Revert to the original VM, discarding all changes made to the VM replica during its operational period.
- **Permanent Failover** - Permanently switch from the original VM to its replica, establishing the replica as the new primary VM for ongoing operations.

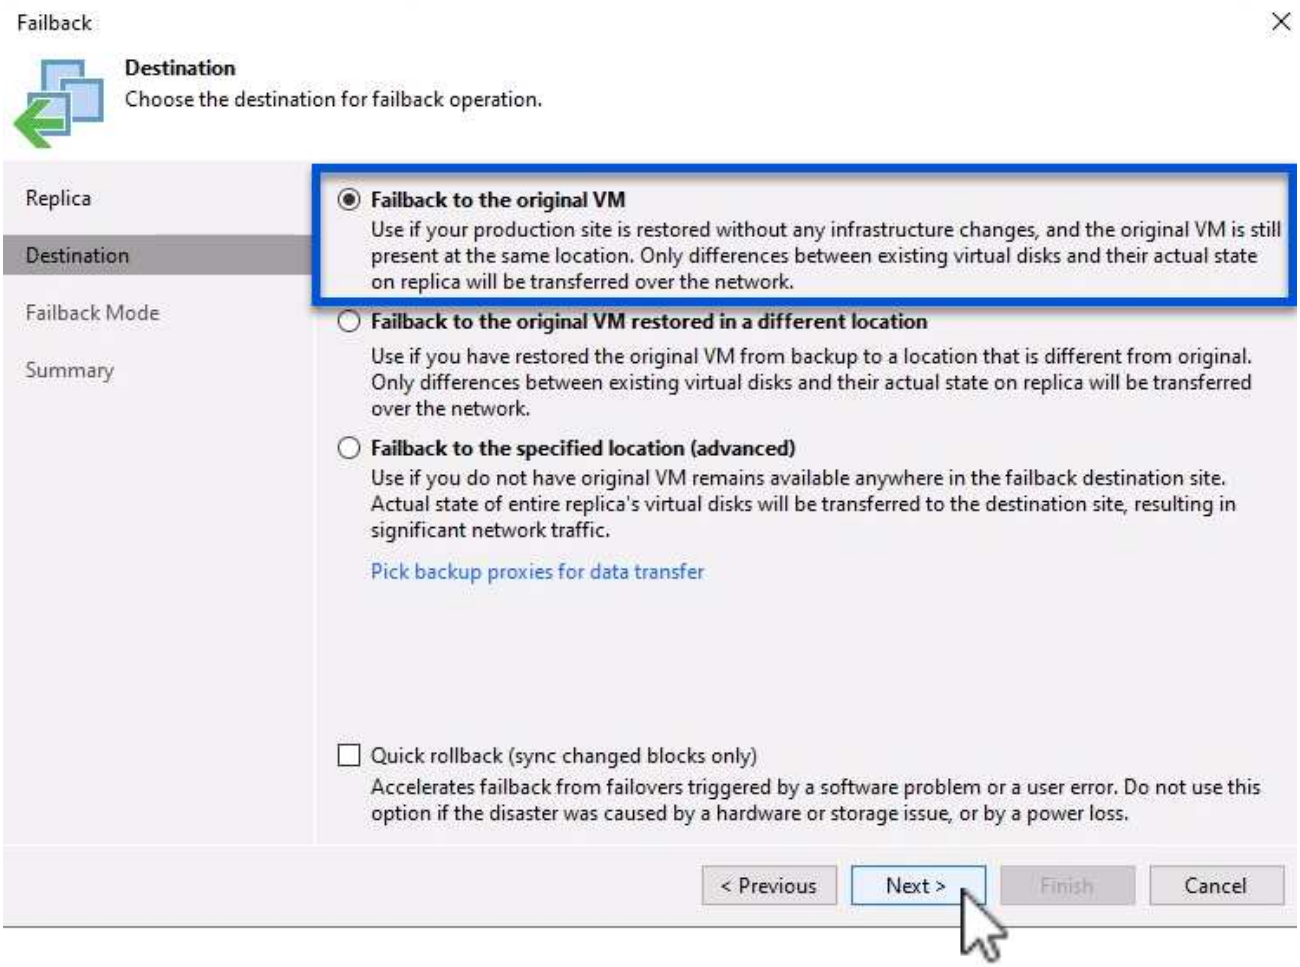
In this scenario, the "Failback to production" option was selected.

Complete the following steps to perform a failback to the production site:

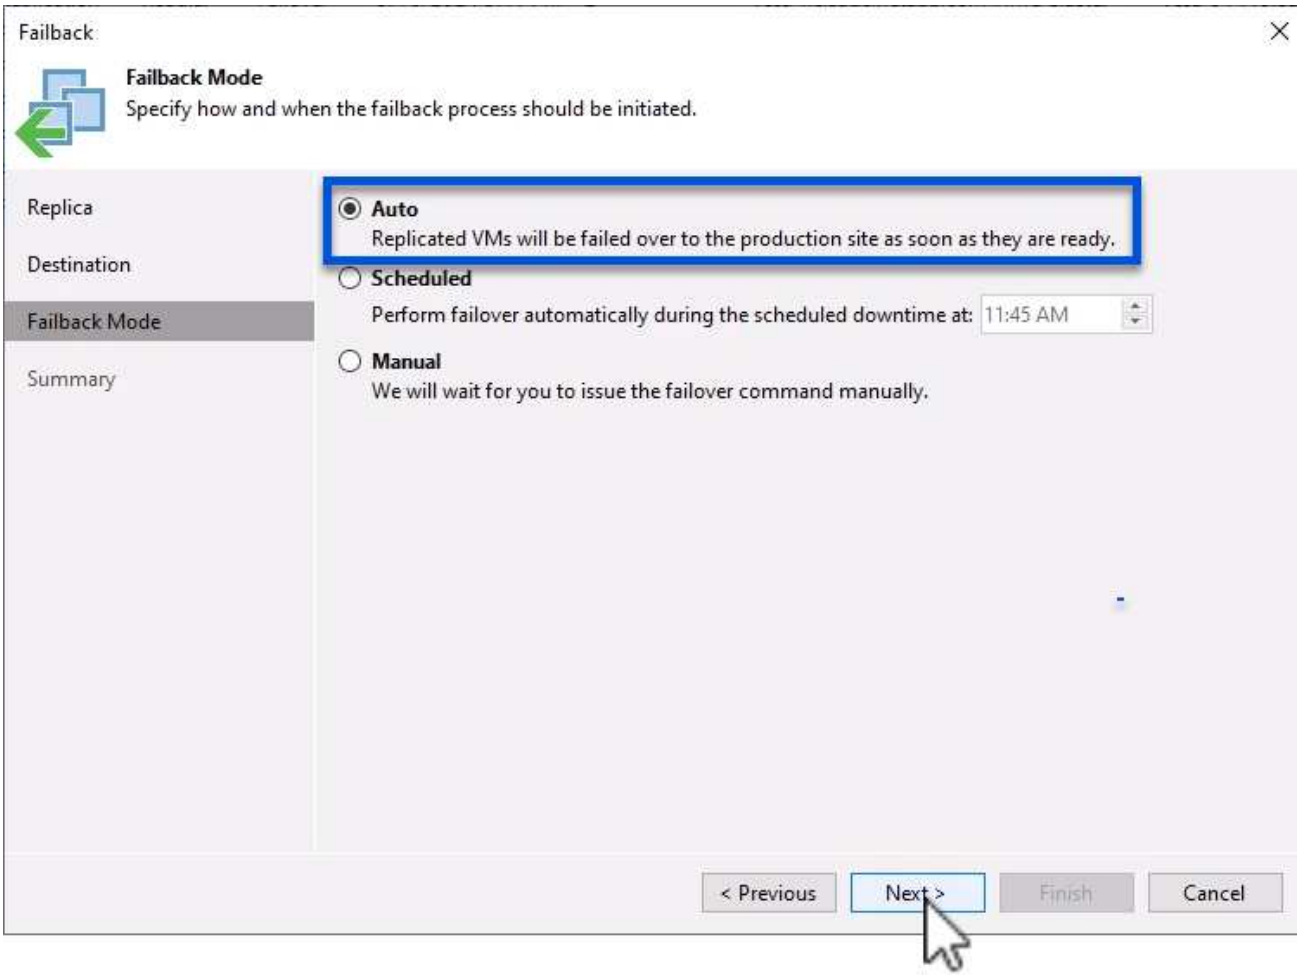
1. From the **Home** view, click on **Replicas > Active** in the left-hand menu. Select the VMs to be included and click on the **Failback to Production** button in the top menu.



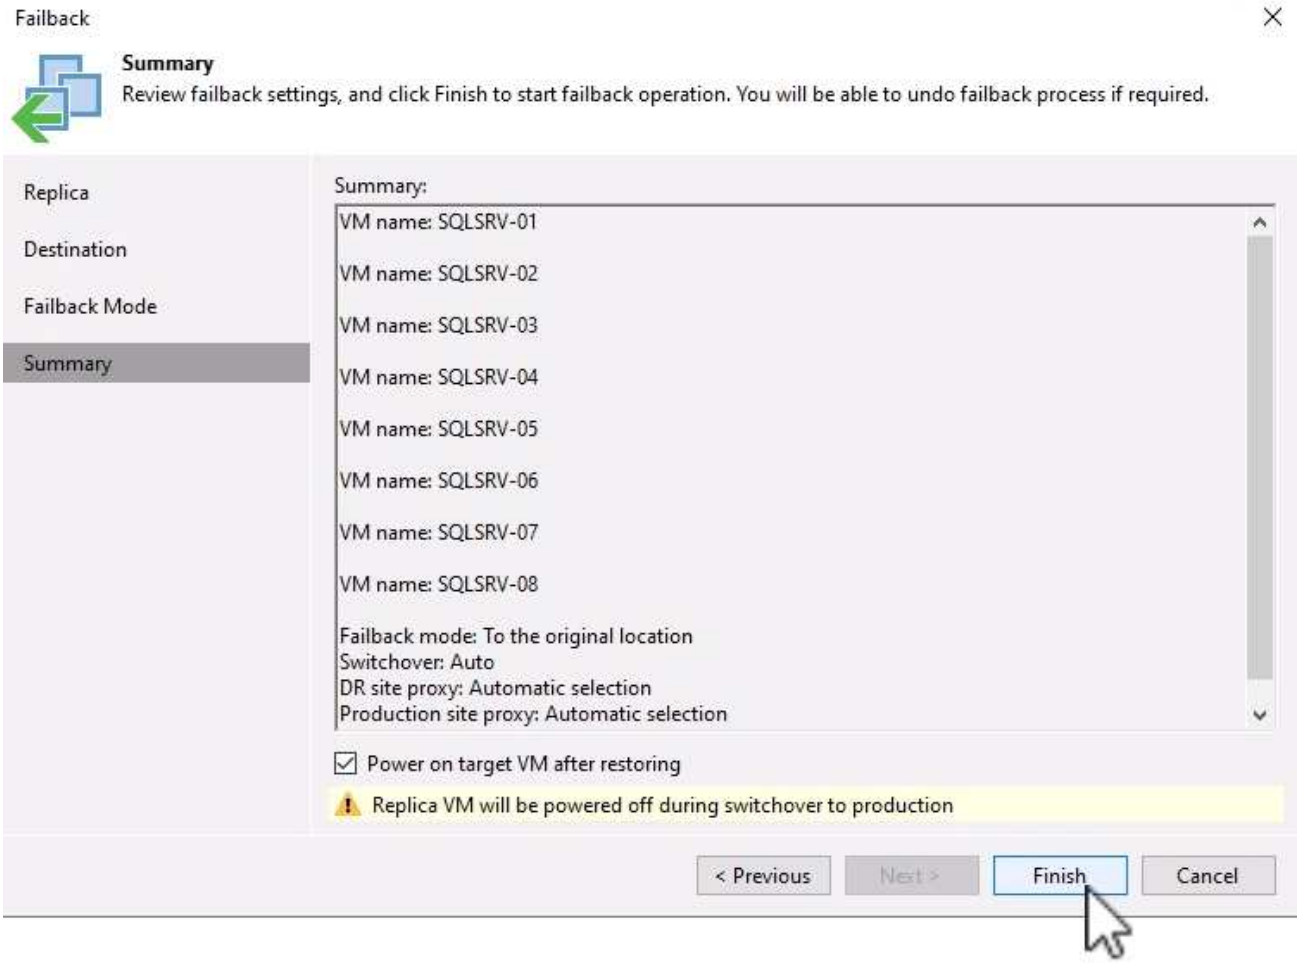
2. On the **Replica** page of the **Failback** wizard, select the replicas to include in the failback job.
3. On the **Destination** page, select **Failback to the original VM** and click on **Next** to continue.



4. On the **Failback Mode** page, select **Auto** to start the failback as soon as possible.

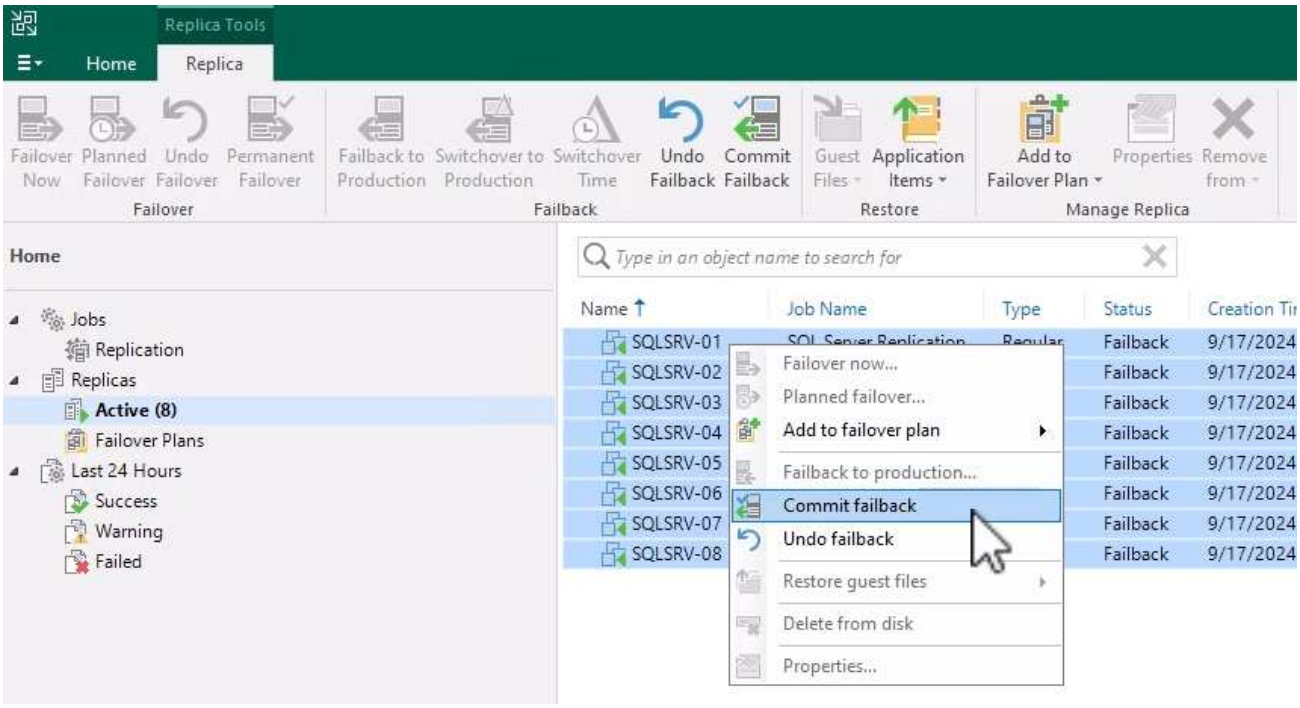


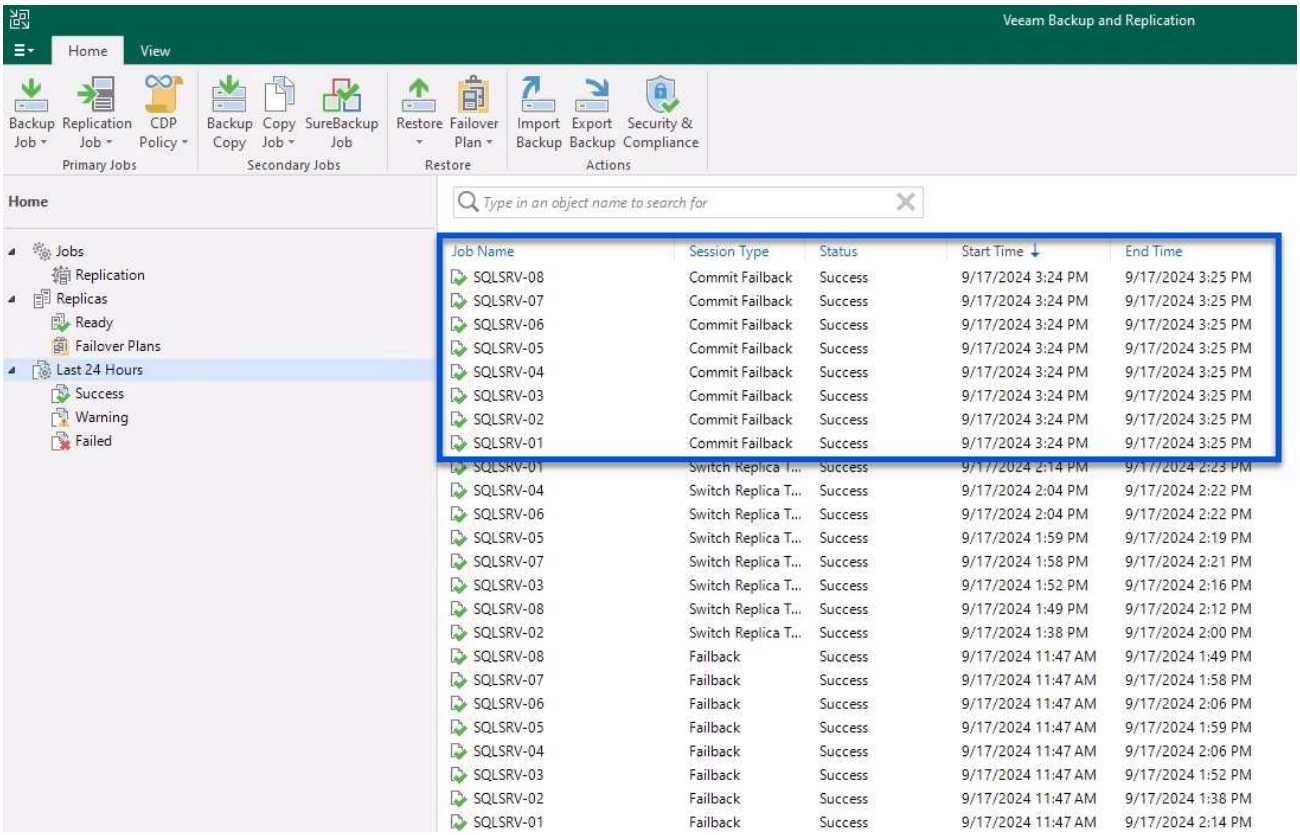
5. On the **Summary** page, choose whether to **Power on target VM after restoring** and then click on Finish to start the failback job.



Failback commit finalizes the failback operation, confirming the successful integration of changes to the production VM. Upon commit, Veeam Backup & Replication resumes regular replication activities for the restored production VM. This changes the status of the restored replica from *Failback* to *Ready*.

1. To commit failback, navigate to **Replicas > Active**, select the VMs to be committed, right click and select **Commit failback**.





After failback to production is successful, the VMs are all restored back to the original production site.

For detailed information about the failback process, refer Veeam documentation for Failover and Failback for replication.

Conclusion

Google Cloud NetApp Volumes datastore functionality empowers Veeam and other validated third-party tools to deliver cost-effective disaster recovery (DR) solutions. By utilizing Pilot light clusters instead of large, dedicated clusters for VM replicas, organizations can significantly reduce expenses. This approach enables tailored DR strategies that leverage existing in-house backup solutions for cloud-based disaster recovery, eliminating the need for additional on-premises datacenters. In the event of a disaster, failover can be initiated with a single click or configured to occur automatically, ensuring business continuity with minimal downtime.

To learn more about this process, feel free to follow the detailed walkthrough video.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=b2fb8597-c3fe-49e2-8a84-b1f10118db6d

Migrate workloads on GCP/GCVE

Migrate workloads to Google Cloud NetApp Volumes datastore on Google Cloud VMware Engine using VMware HCX - Quickstart guide

One of the most common use cases for the Google Cloud VMware Engine and Cloud Volume Service datastore is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration mechanisms to move on-premises virtual machines (VMs) and its data to Cloud Volume Service NFS datastores.

Overview: Migrating virtual machines with VMware HCX, Google Cloud NetApp Volumes datastores, and Google Cloud VMware Engine (GCVE)

VMware HCX is primarily a migration platform that is designed to simplify application migration, workload rebalancing, and even business continuity across clouds. It is included as part of Google Cloud VMware Engine Private Cloud and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for provisioning Cloud Volume Service datastore followed by downloading, deploying, and configuring VMware HCX, including all its main components in on-premises and the Google Cloud VMware Engine side including Interconnect, Network Extension, and WAN optimization for enabling various VM migration mechanisms.

> ⓘ VMware HCX works with any datastore type as the migration is at the VM level. Hence this document is applicable to existing NetApp customers and non-NetApp customers who are planning to deploy Cloud Volume Service with Google Cloud VMware Engine for a cost-effective VMware cloud deployment.

▼ **High-level steps**

This list provides the high-level steps necessary to pair & Migrate the VMs to HCX Cloud Manager on the Google Cloud VMware Engine side from HCX Connector on-premises:

1. Prepare HCX through the Google VMware Engine portal.
2. Download and deploy the HCX Connector Open Virtualization Appliance (OVA) installer in the on-premises VMware vCenter Server.
3. Activate HCX with the license key.
4. Pair the on-premises VMware HCX Connector with Google Cloud VMware Engine HCX Cloud Manager.
5. Configure the network profile, compute profile, and service mesh.

6. (Optional) Perform network extension to avoid re-IP during migrations.
7. Validate the appliance status and ensure that migration is possible.
8. Migrate the VM workloads.

▼ Prerequisites

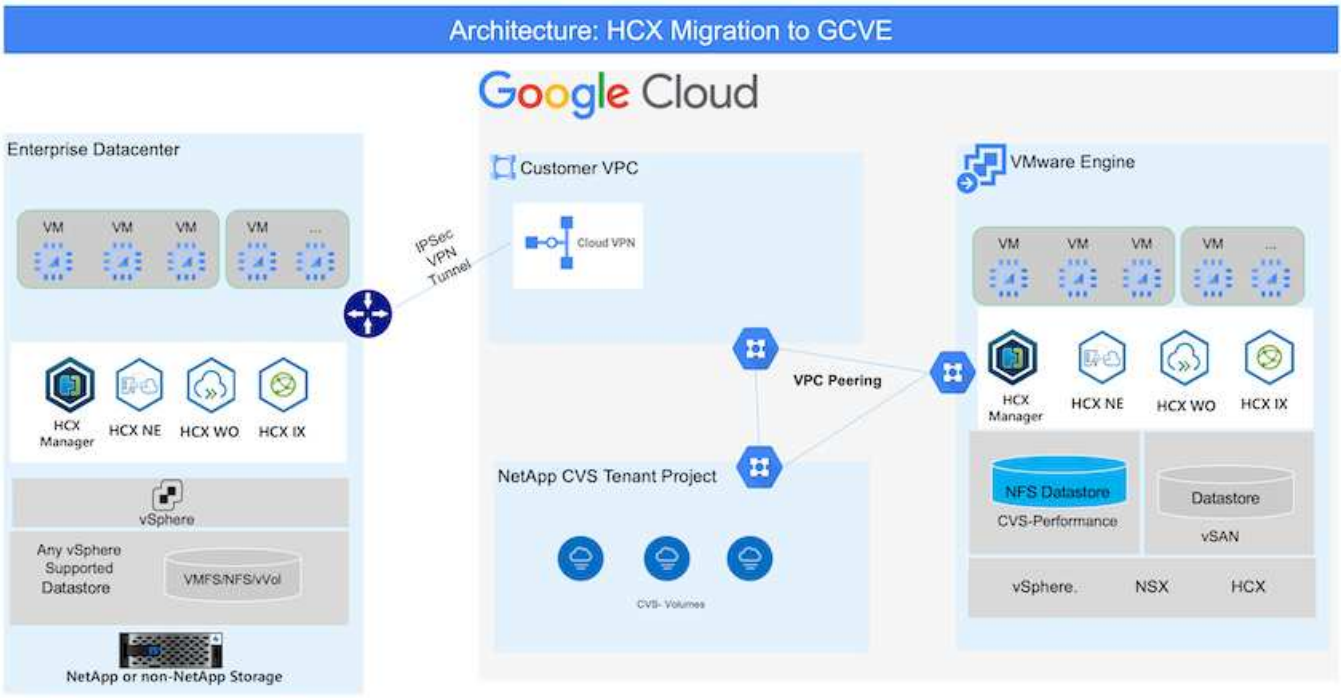
Before you begin, make sure the following prerequisites are met. For more information, see this link. After the prerequisites, including connectivity, are in place, download HCX license key from the Google Cloud VMware Engine portal. After the OVA installer is downloaded, proceed with the installation process as described below.

> ⓘ HCX advanced is the default option and VMware HCX Enterprise edition is also available through a support ticket and supported at no additional cost. Refer this link

- Use an existing Google Cloud VMware Engine software-defined data center (SDDC) or create a private cloud by using this NetApp link or this Google link.
- Migration of VMs and associated data from the on-premises VMware vSphere- enabled data center requires network connectivity from the data center to the SDDC environment. Before migrating workloads, set up a Cloud VPN or Cloud Interconnect connection between the on-premises environment and the respective private cloud.
- The network path from on-premises VMware vCenter Server environment to the Google Cloud VMware Engine private cloud must support the migration of VMs by using vMotion.
- Make sure the required firewall rules and ports are allowed for vMotion traffic between the on-premises vCenter Server and SDDC vCenter.
- Cloud Volume Service NFS volume should be mounted as a datastore in Google Cloud VMware Engine. Follow the steps detailed in this link to attach Cloud Volume Service datastores to Google Cloud VMware Engines hosts.

▼ High Level Architecture

For testing purposes, the lab environment from on-premises used for this validation was connected through a Cloud VPN, which allows on-premises connectivity to Google Cloud VPC.



For more details regarding using VMware HCX with Google, please refer to VMware link

Solution Deployment

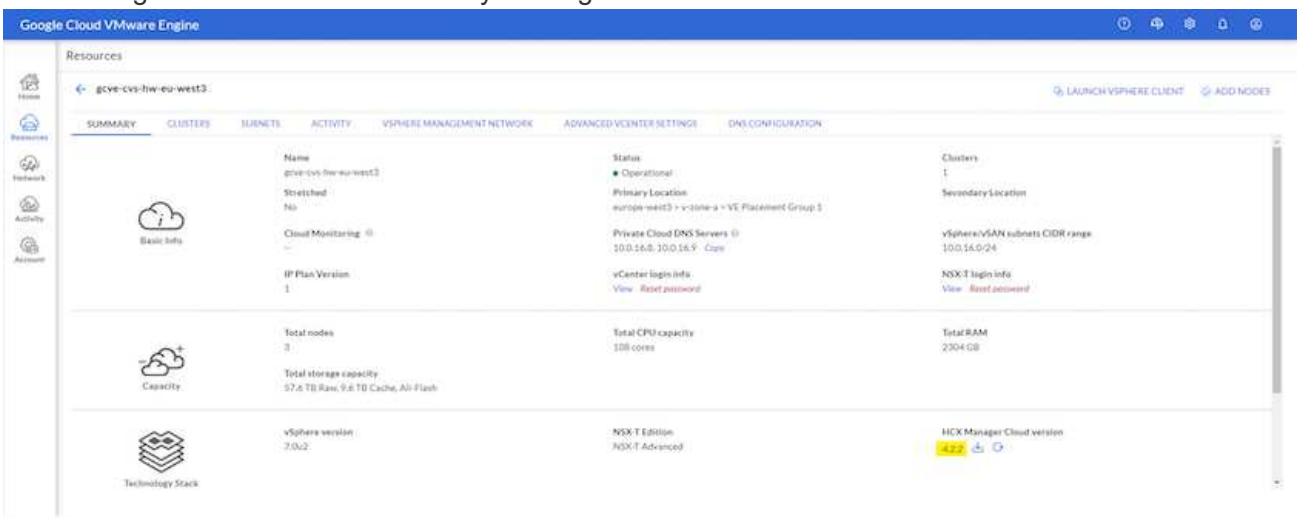
Follow the series of steps to complete the deployment of this solution:

▼ **Step 1: Prepare HCX through the Google VMware Engine Portal**

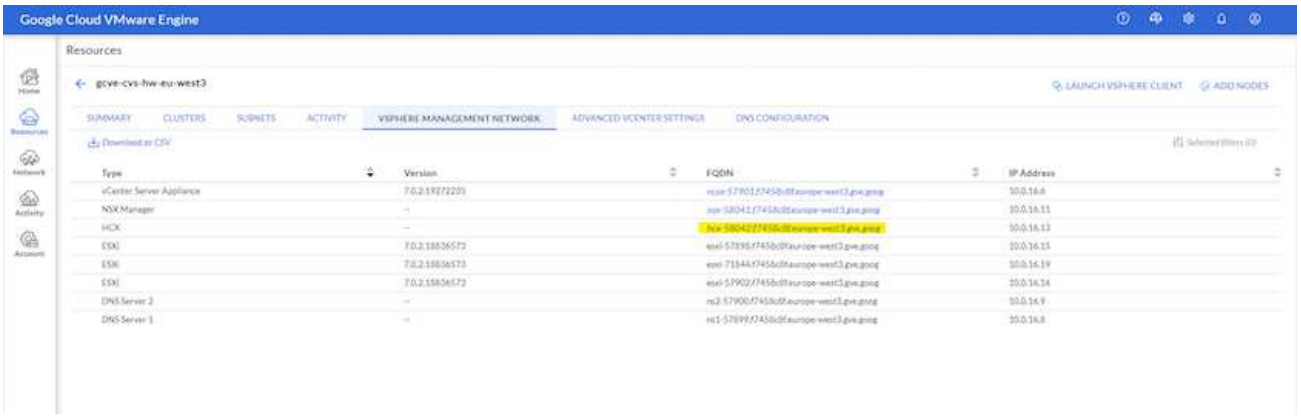
HCX Cloud Manager component automatically gets installed as you provision private cloud with VMware Engine. To prepare for site pairing, complete the following steps:

1. Log in to the Google VMware Engine Portal and sign-in to the HCX Cloud Manager.

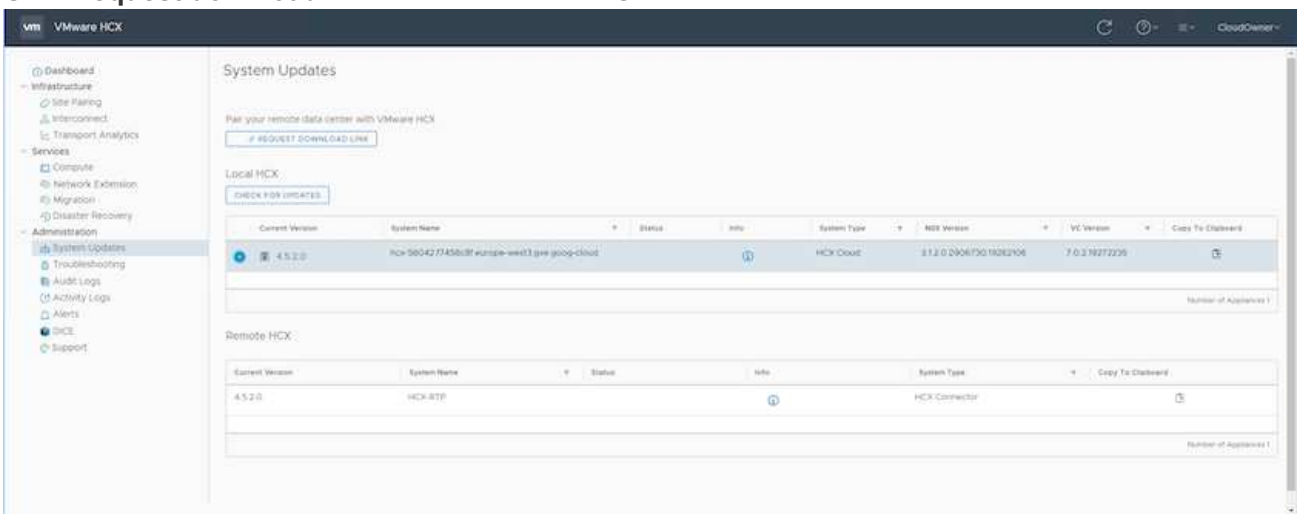
   You can login to HCX Console either by clicking on the HCX version link



   or clicking on HCX FQDN under vSphere Management Network tab.



2. In HCX Cloud Manager, go to **Administration > System Updates**.
3. Click **Request download link** and download the OVA file.



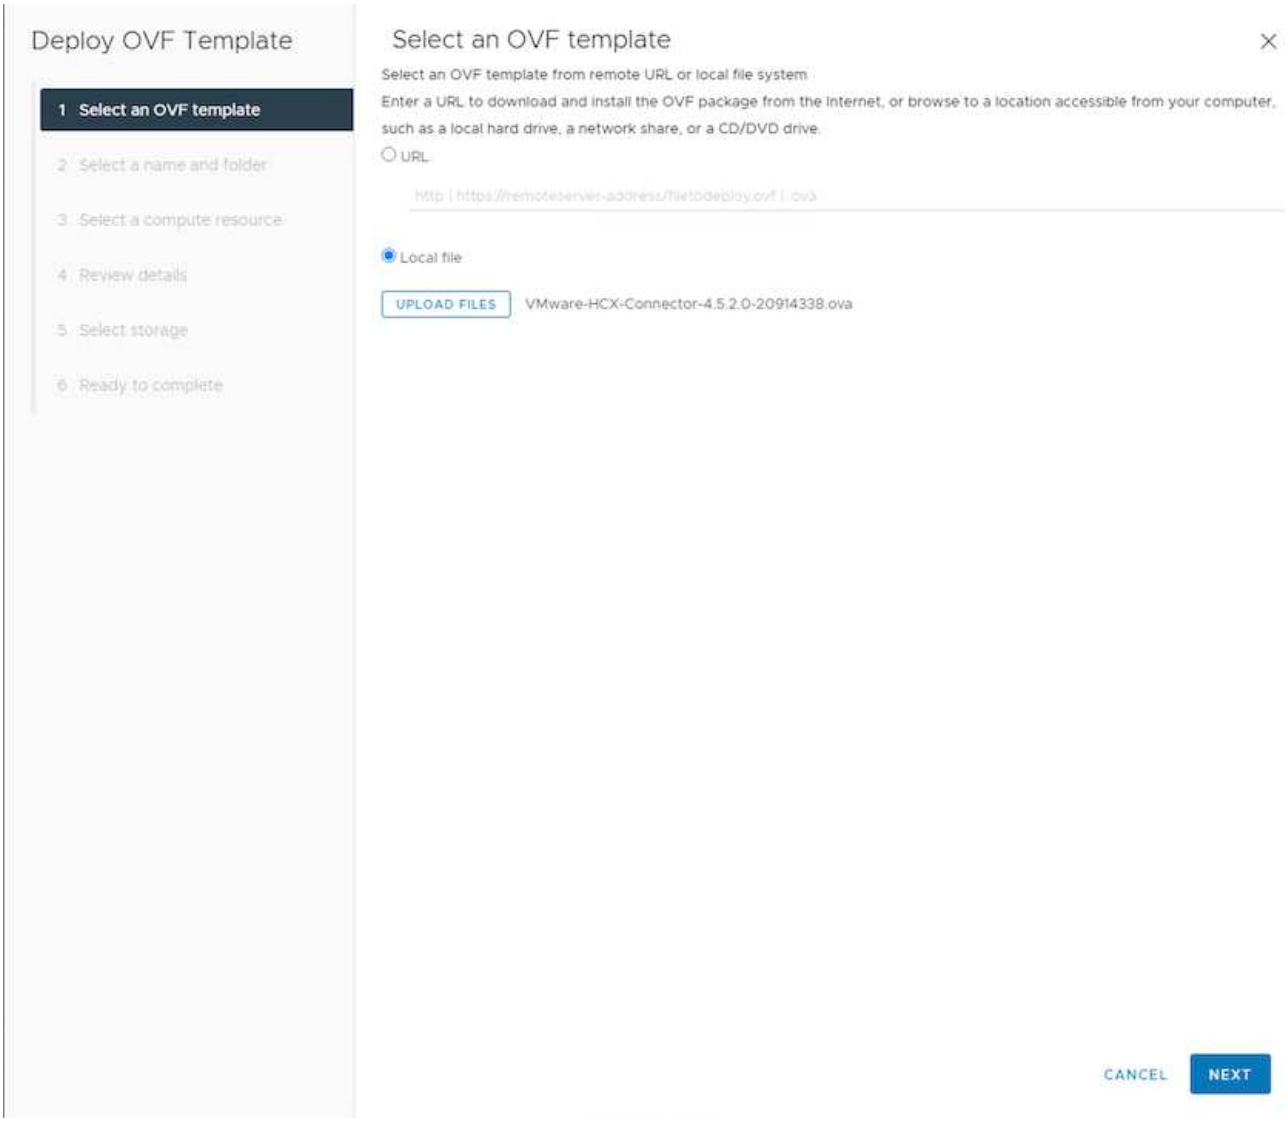
4. Update HCX Cloud Manager to the latest version available from the HCX Cloud Manager UI.

▼ **Step 2: Deploy the installer OVA in the on-premises vCenter Server**

For the on-premises Connector to connect to the HCX Manager in Google Cloud VMware Engine, make sure the appropriate firewall ports are open in the on-premises environment.

To download and install HCX Connector in the on-premises vCenter Server, complete the following steps:

1. Have the ova downloaded from the HCX Console on Google Cloud VMware Engine as stated in previous step.
2. After the OVA is downloaded, deploy it on to the on-premises VMware vSphere environment by using the **Deploy OVF Template** option.



3. Enter all the required information for the OVA deployment, click **Next**, and then click **Finish** to deploy the VMware HCX connector OVA.

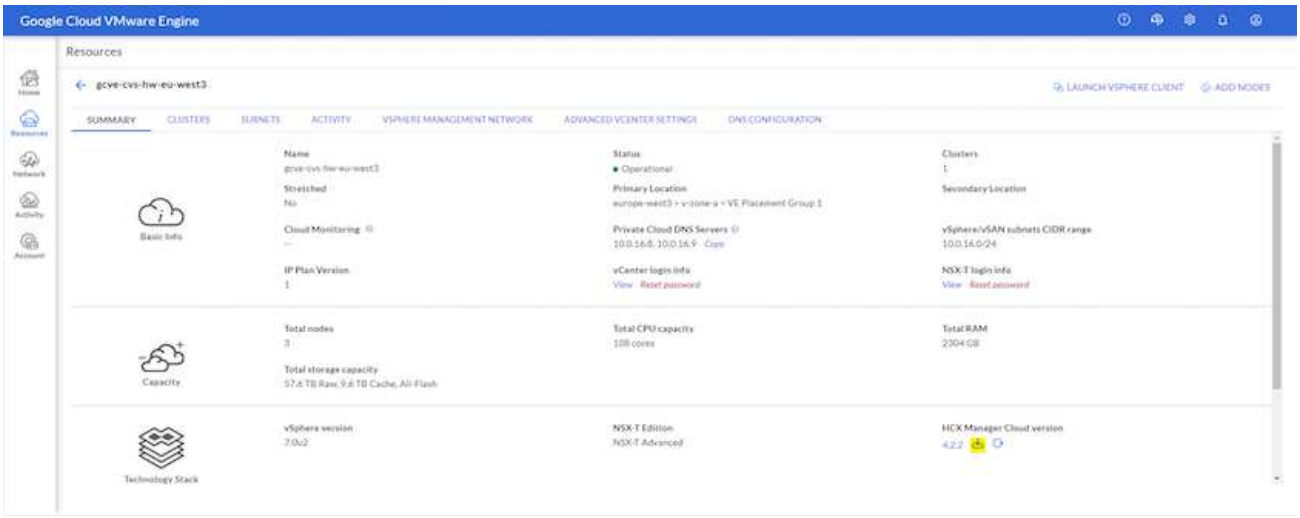
   Power on the virtual appliance manually.

For step-by-step instructions, see the Google HCX Documentation.

▼ **Step 3: Activate HCX Connector with the license key**

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the Google Cloud VMware Engine portal and activate it in VMware HCX Manager.

1. From the VMware Engine portal, Click on Resources, select the private cloud, and **click on download icon under HCX Manager Cloud Version**.



Open Downloaded file and copy the License Key String.

2. Log into the on-premises VMware HCX Manager at `"https://hcxmanagerIP:9443"` using administrator credentials.

   Use the hcxmanagerIP and password defined during the OVA deployment.

3. In the licensing, enter the key copied from step 3 and click **Activate**.

   The on-premises HCX Connector should have internet access.

4. Under **Datacenter Location**, provide the nearest location for installing the VMware HCX Manager on-premises. Click **Continue**.
5. Under **System Name**, update the name and click **Continue**.
6. Click **Yes, Continue**.
7. Under **Connect your vCenter**, provide the fully qualified domain name (FQDN) or IP address of vCenter Server and the appropriate credentials and click **Continue**.

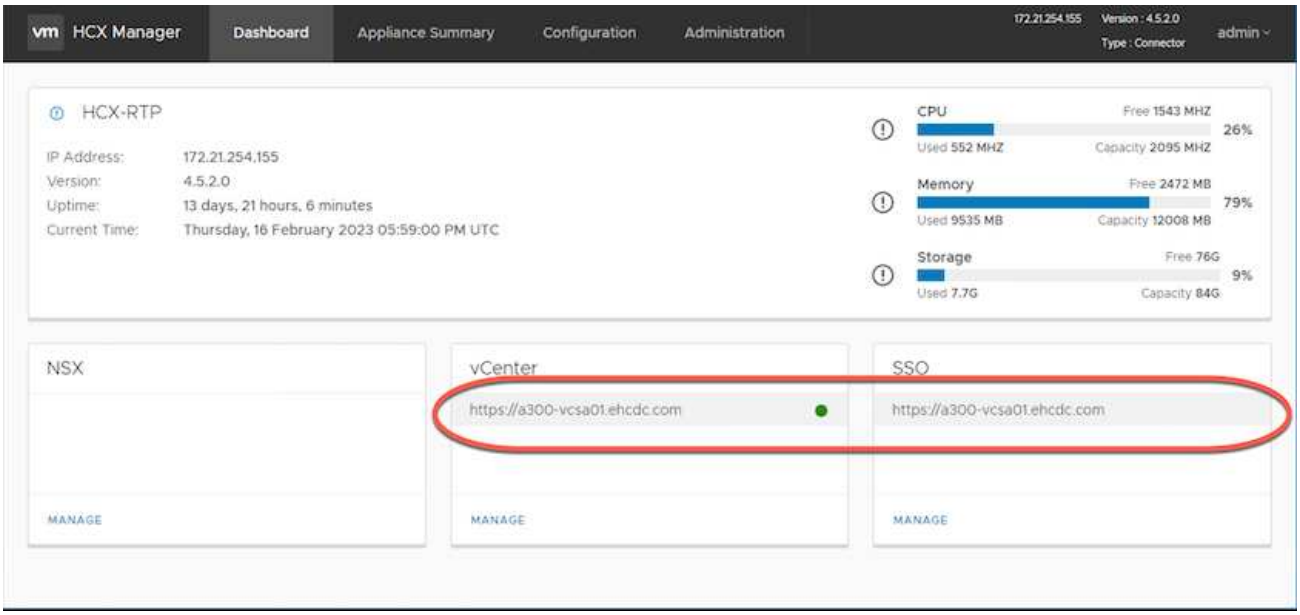
   Use the FQDN to avoid connectivity issues later.

8. Under **Configure SSO/PSC**, provide the Platform Services Controller's(PSC) FQDN or IP address and click **Continue**.

   For Embedded PSC, Enter the VMware vCenter Server FQDN or IP address.

9. Verify that the information entered is correct and click **Restart**.
10. After the services restart, vCenter Server is displayed as green on the page that appears. Both vCenter Server and SSO must have the appropriate configuration parameters, which should be the same as the previous page.

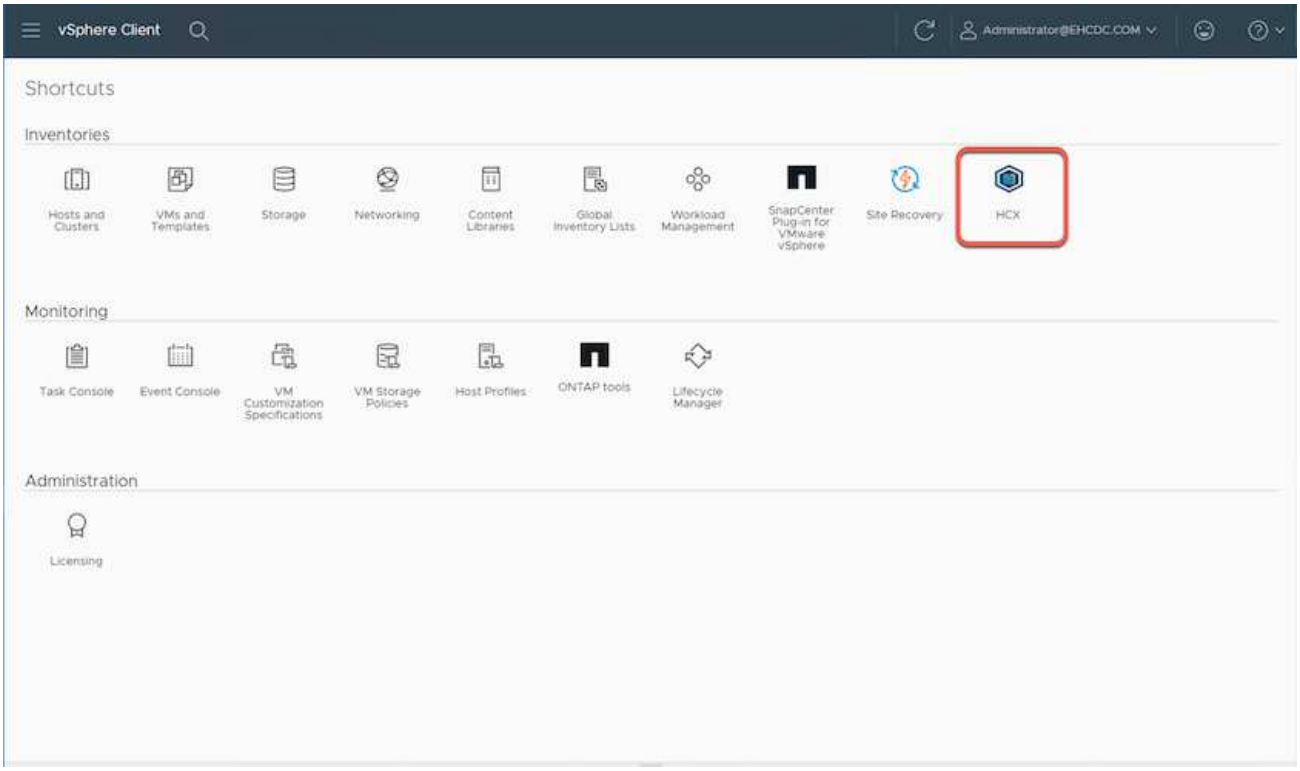
    This process should take approximately 10 to 20 minutes and for the plug-in to be added to the vCenter Server.



▼ **Step 4: Pair on-premises VMware HCX Connector with Google Cloud VMware Engine HCX Cloud Manager**

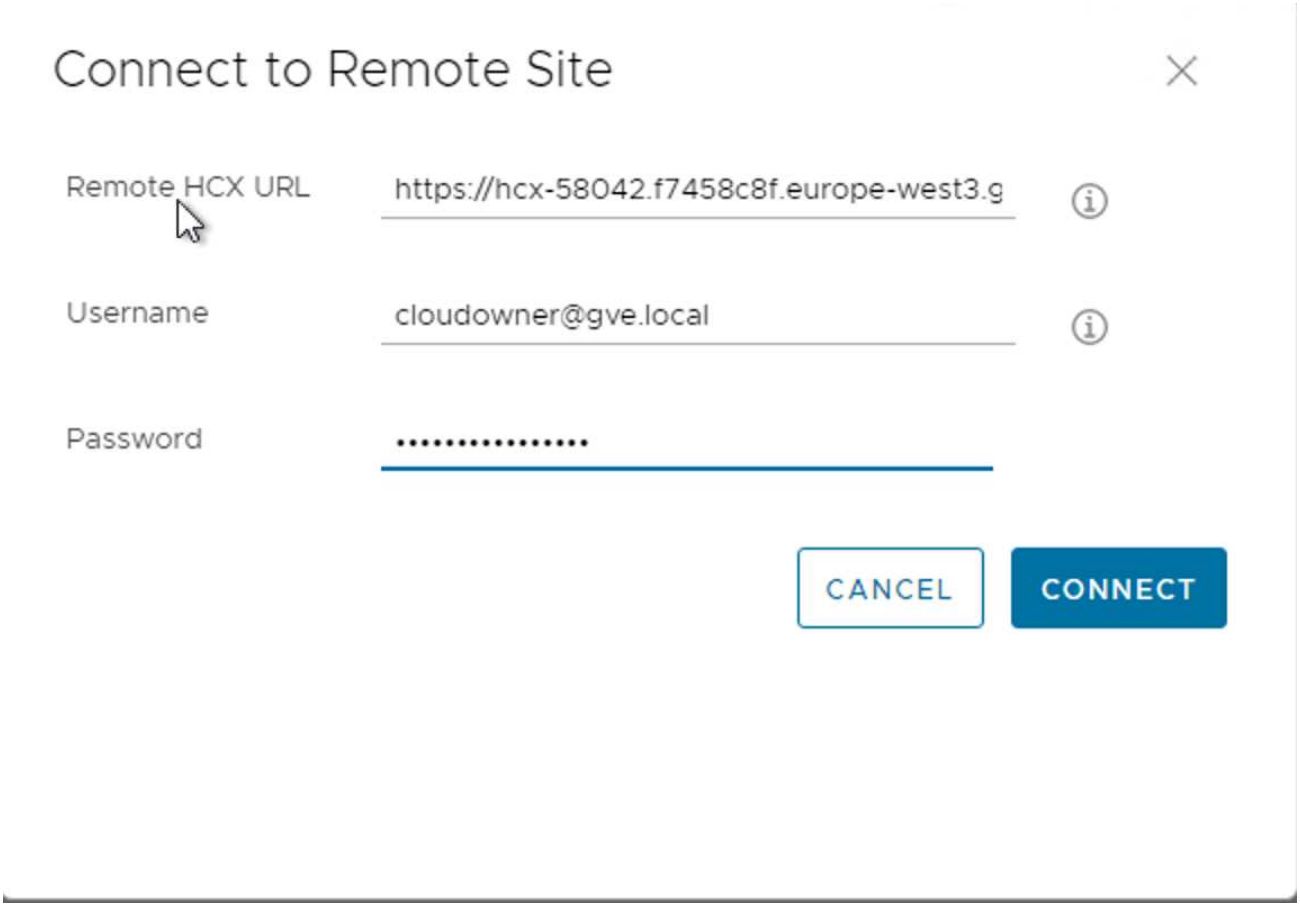
After HCX Connector is deployed and configured on on-premises vCenter, establish connection to Cloud Manager by adding the pairing. To configure the site pairing, complete the following steps:

1. To create a site pair between the on-premises vCenter environment and Google Cloud VMware Engine SDDC, log in to the on-premises vCenter Server and access the new HCX vSphere Web Client plug-in.



2. Under Infrastructure, click **Add a Site Pairing**.

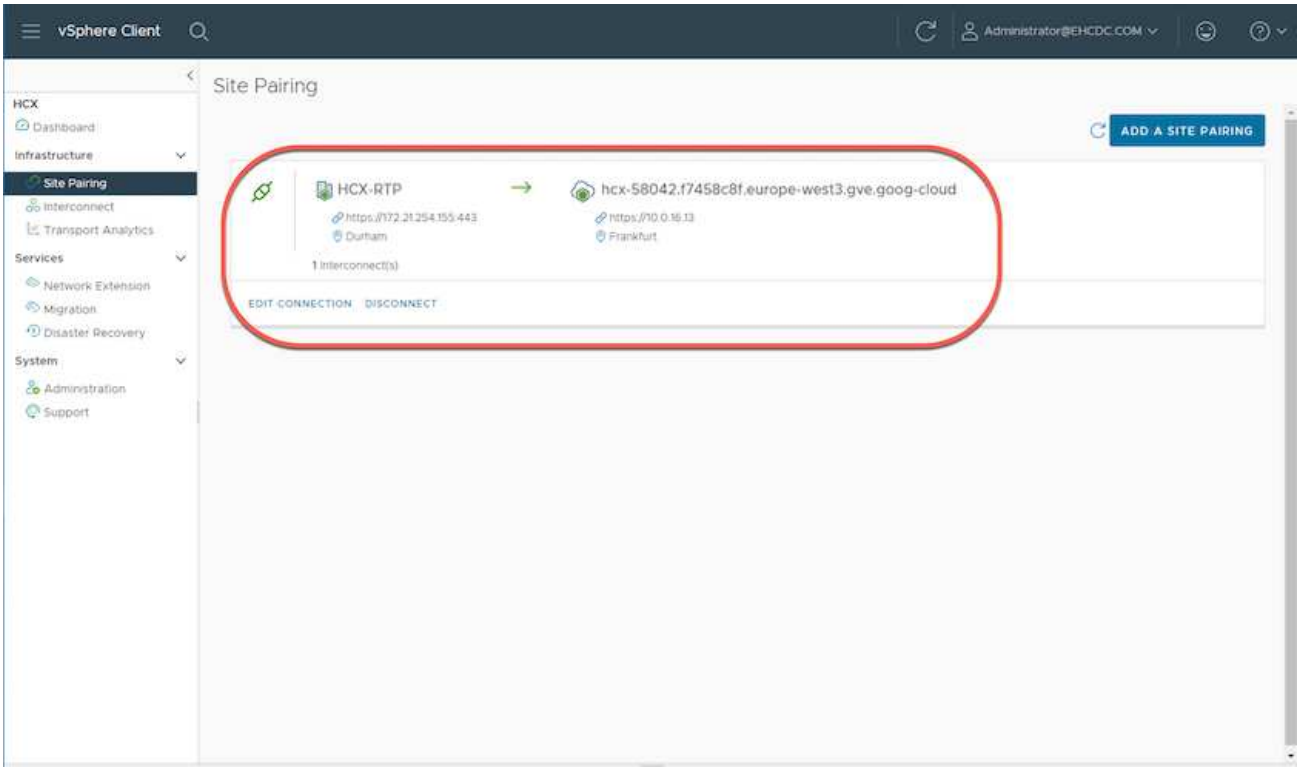
Enter the Google Cloud VMware Engine HCX Cloud Manager URL or IP address and the credentials for user with Cloud-Owner-Role privileges for accessing the private cloud.



3. Click **Connect**.

   VMware HCX Connector must be able to route to HCX Cloud Manager IP over port 443.

4. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.

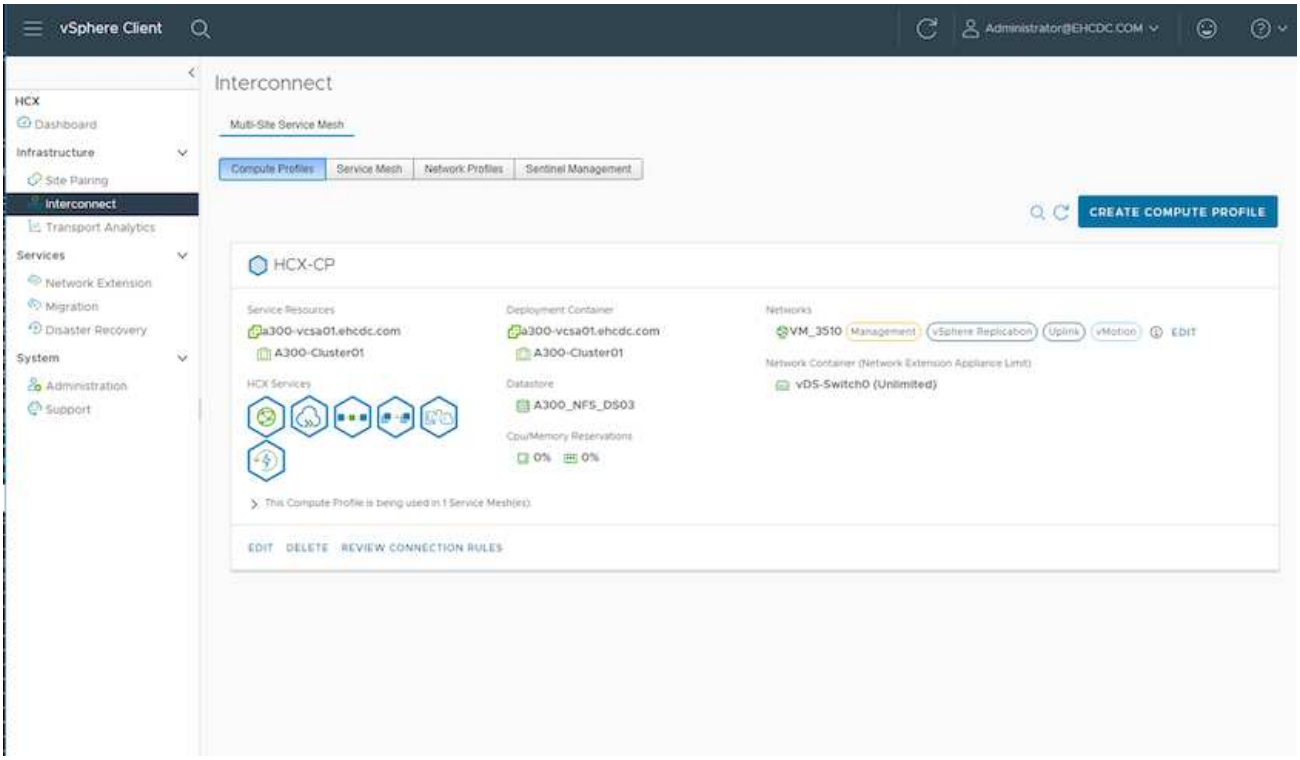


▼ **Step 5: Configure the network profile, compute profile, and service mesh**

The VMware HCX Interconnect service appliance provides replication and vMotion-based migration capabilities over the internet and private connections to the target site. The interconnect provides encryption, traffic engineering, and VM mobility. To create an Interconnect service appliance, complete the followings steps:

1. Under Infrastructure, select **Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile**.

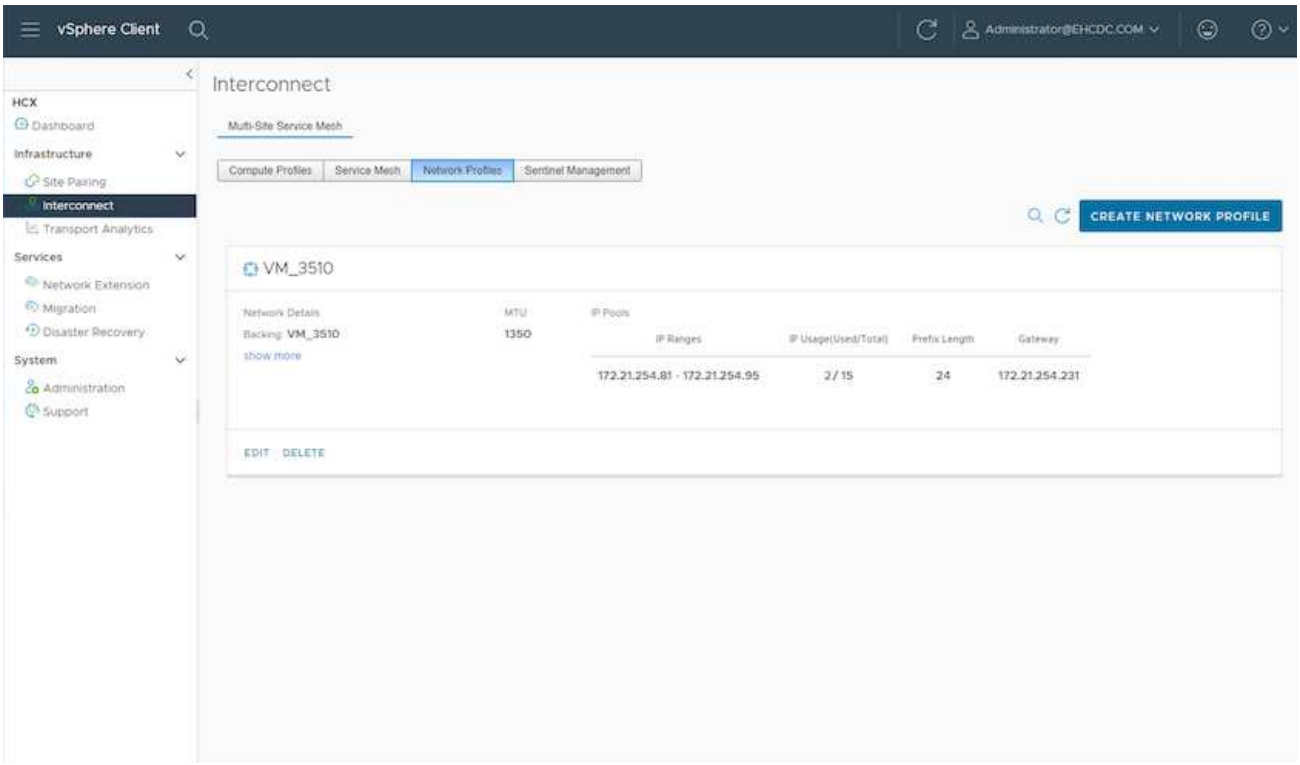
> The compute profiles define the deployment parameters including the appliances that are deployed and which portion of the VMware data center are accessible to HCX service.



2. After the compute profile is created, create the network profiles by selecting **Multi-Site Service Mesh > Network Profiles > Create Network Profile**.

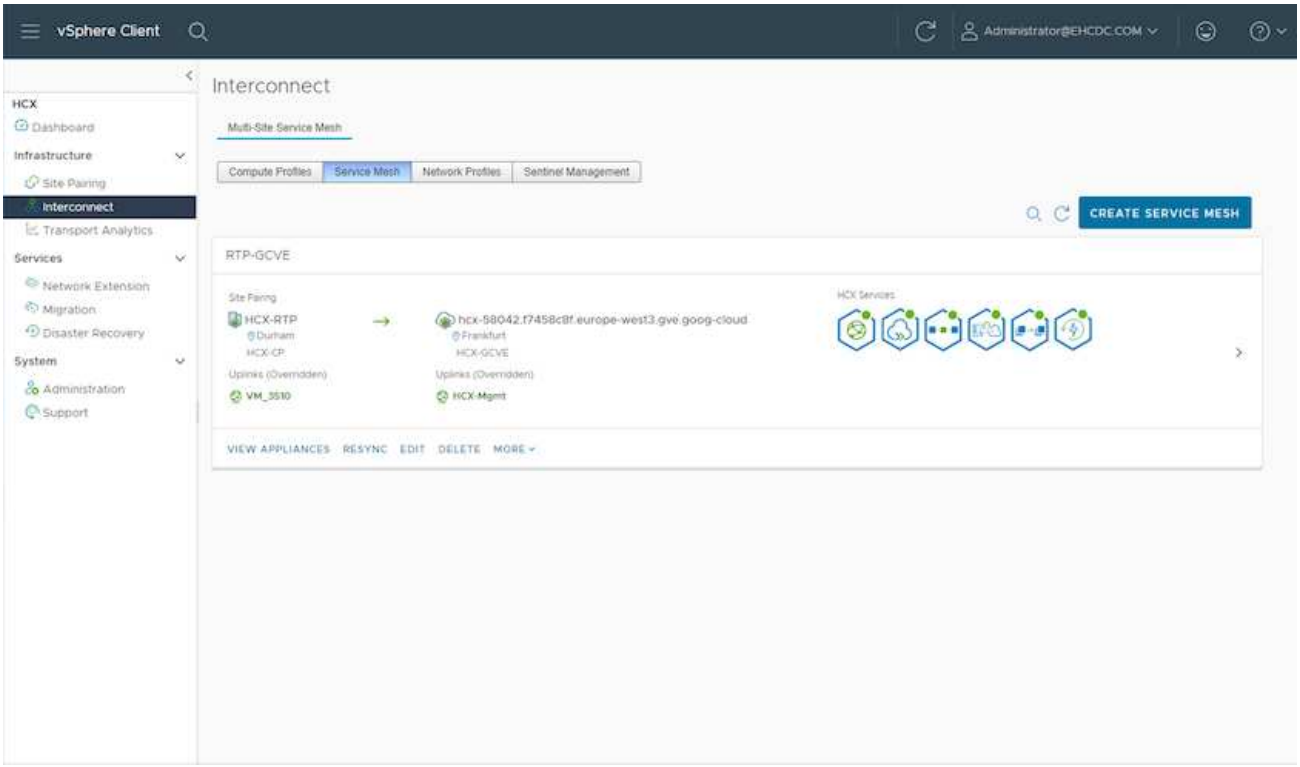
   The network profile defines a range of IP address and networks that are used by HCX for its virtual appliances.

   This step requires two or more IP addresses. These IP addresses are assigned from the management network to the Interconnect Appliances.

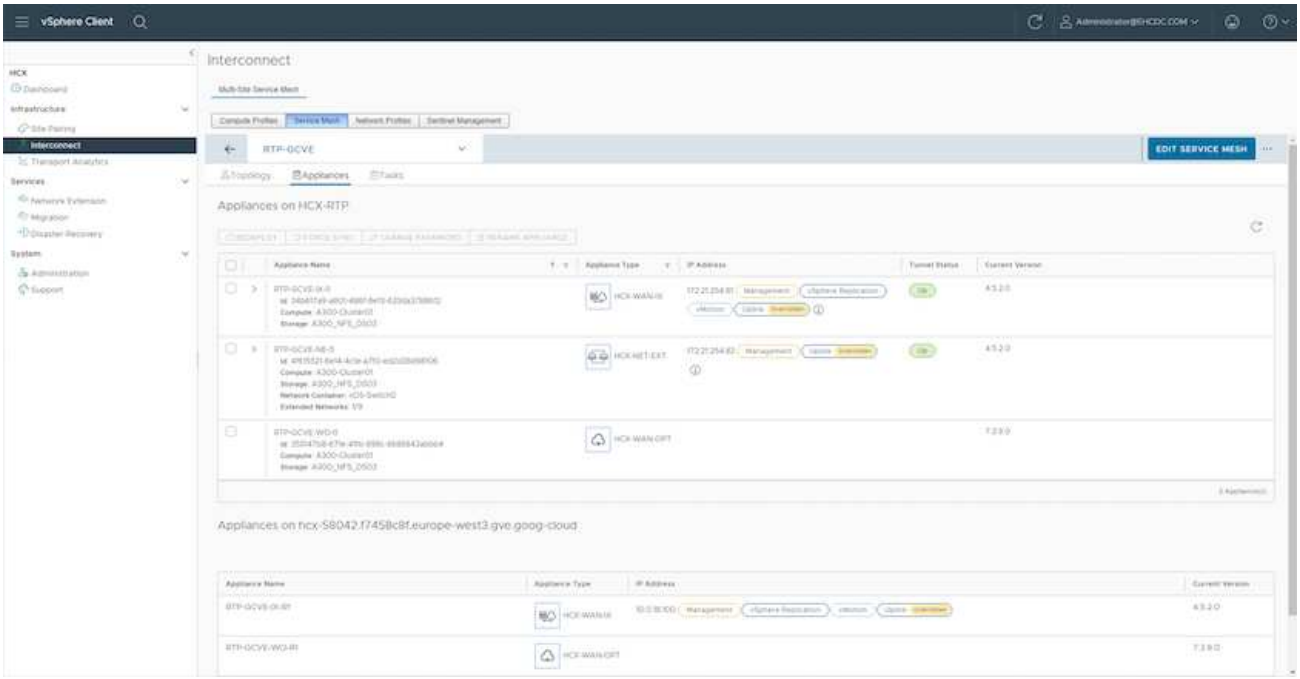


3. At this time, the compute and network profiles have been successfully created.
4. Create the Service Mesh by selecting the **Service Mesh** tab within the **Interconnect** option and select the on-premises and GCVE SDDC sites.
5. The Service Mesh specifies a local and remote compute and network profile pair.

> As part of this process, the HCX appliances are deployed and automatically configured on both the source and target sites in order to create a secure transport fabric.



6. This is the final step of configuration. This should take close to 30 minutes to complete the deployment. After the service mesh is configured, the environment is ready with the IPsec tunnels successfully created to migrate the workload VMs.



▼ Step 6: Migrate workloads

Workloads can be migrated bidirectionally between on-premises and GCVE SDDCs using various VMware HCX migration technologies. VMs can be moved to and from VMware HCX-activated entities using multiple migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with the HCX Enterprise edition).

To learn more about various HCX migration mechanisms, see Migrating VMware VMs using VMware HCX documentation.

The HCX-IX appliance uses the Mobility Agent service to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations.

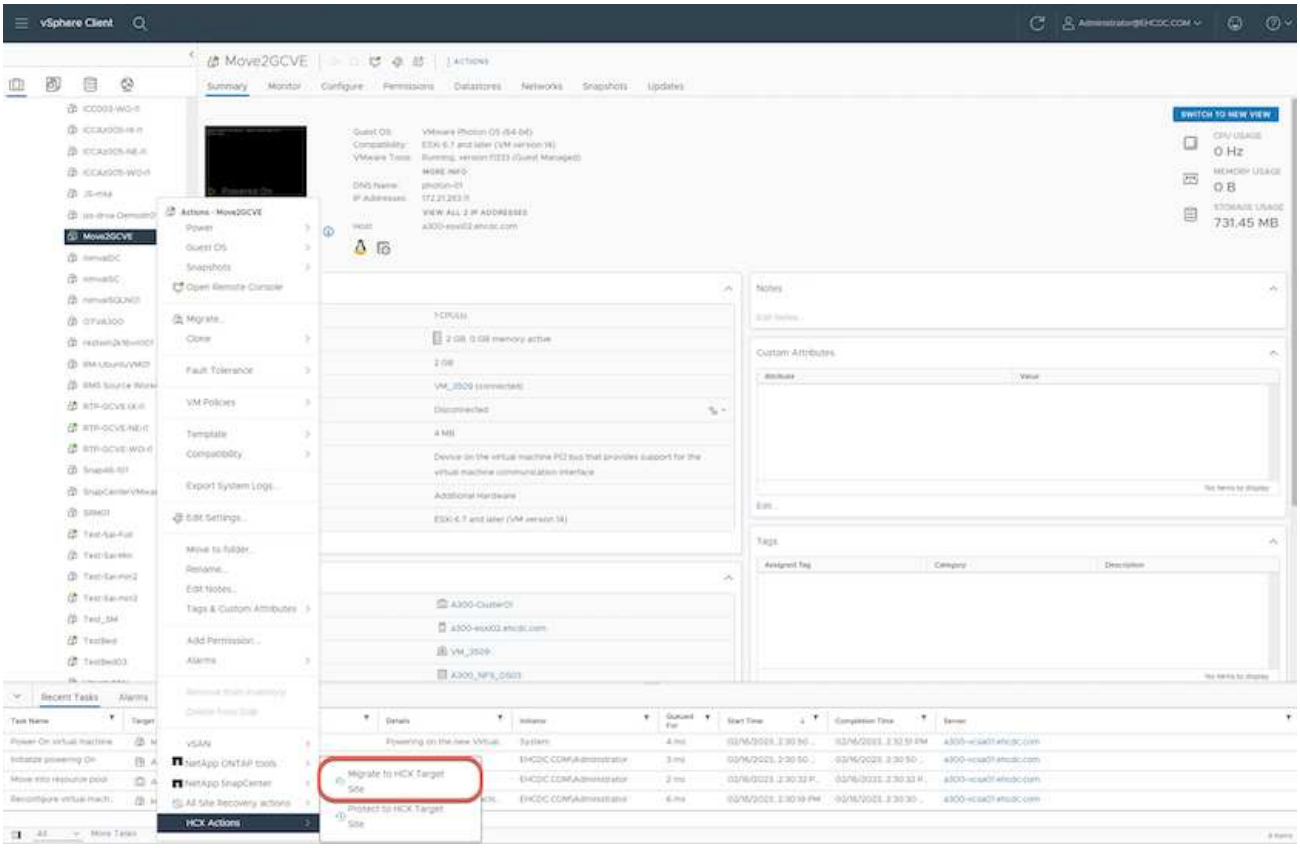
> ⓘ The HCX-IX appliance adds the Mobility Agent service as a host object in the vCenter Server. The processor, memory, storage and networking resources displayed on this object do not represent actual consumption on the physical hypervisor hosting the IX appliance.

HCX vMotion

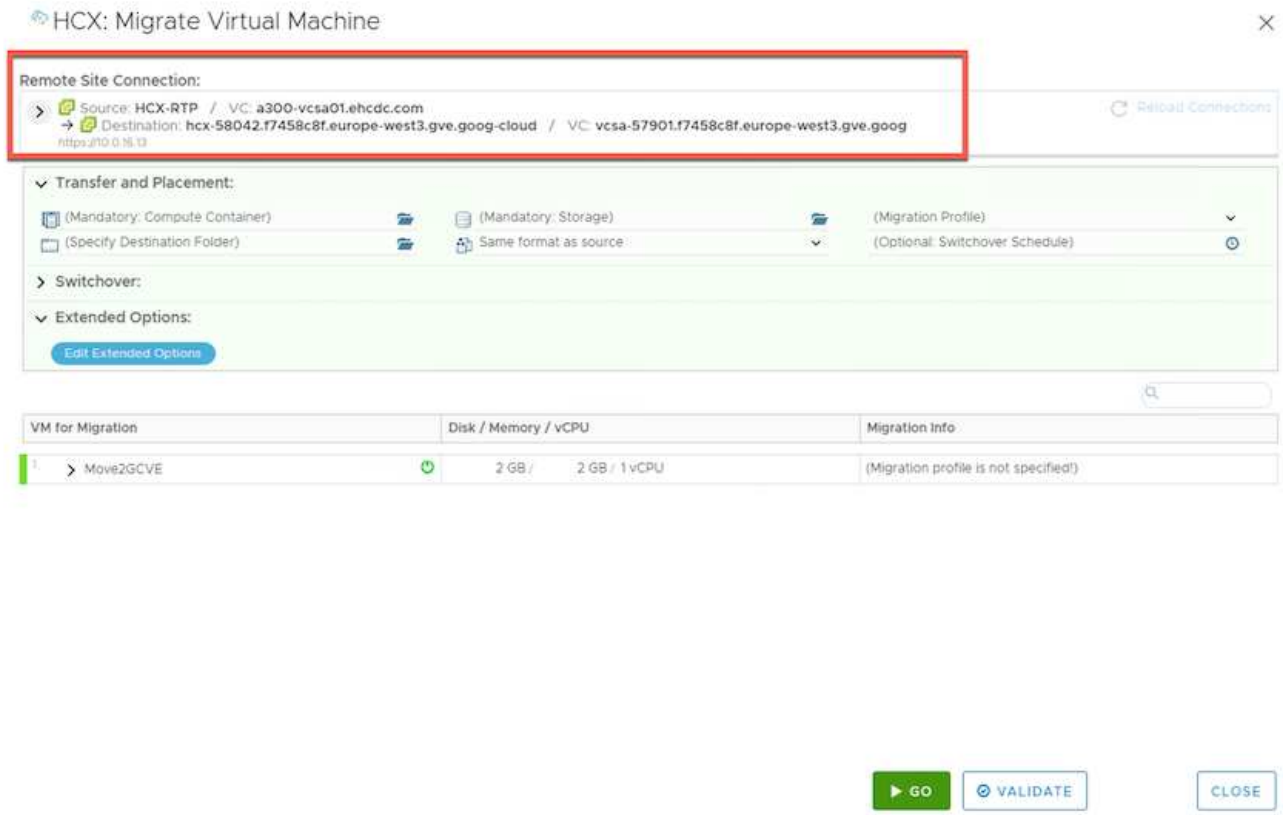
This section describes the HCX vMotion mechanism. This migration technology uses the VMware vMotion protocol to migrate a VM to GCVE. The vMotion migration option is used for migrating the VM state of a single VM at a time. There is no service interruption during this migration method.

> ⓘ Network Extension should be in place (for the port group in which the VM is attached) in order to migrate the VM without the need to make an IP address change.

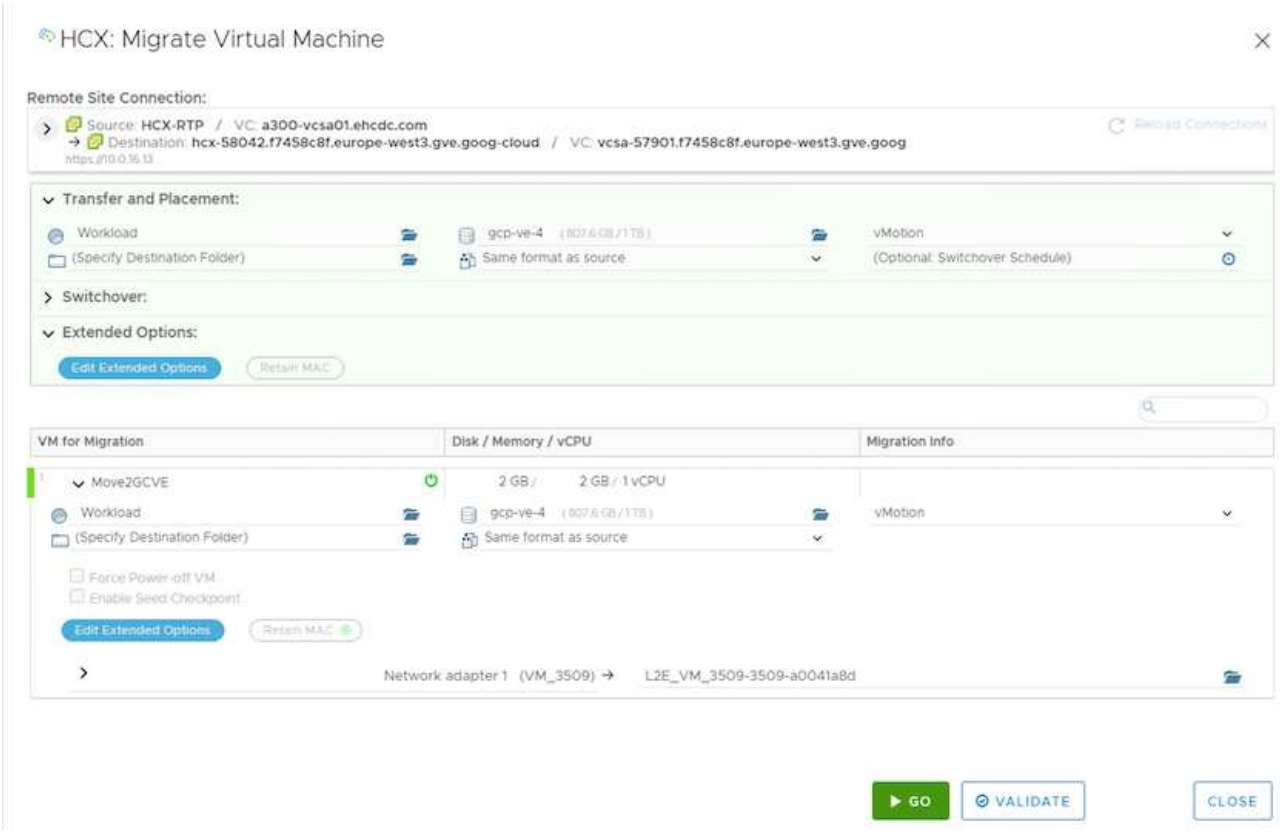
1. From the on-premises vSphere client, go to Inventory, right- click on the VM to be migrated, and select HCX Actions > Migrate to HCX Target Site.



2. In the Migrate Virtual Machine wizard, select the Remote Site Connection (target GCVE).



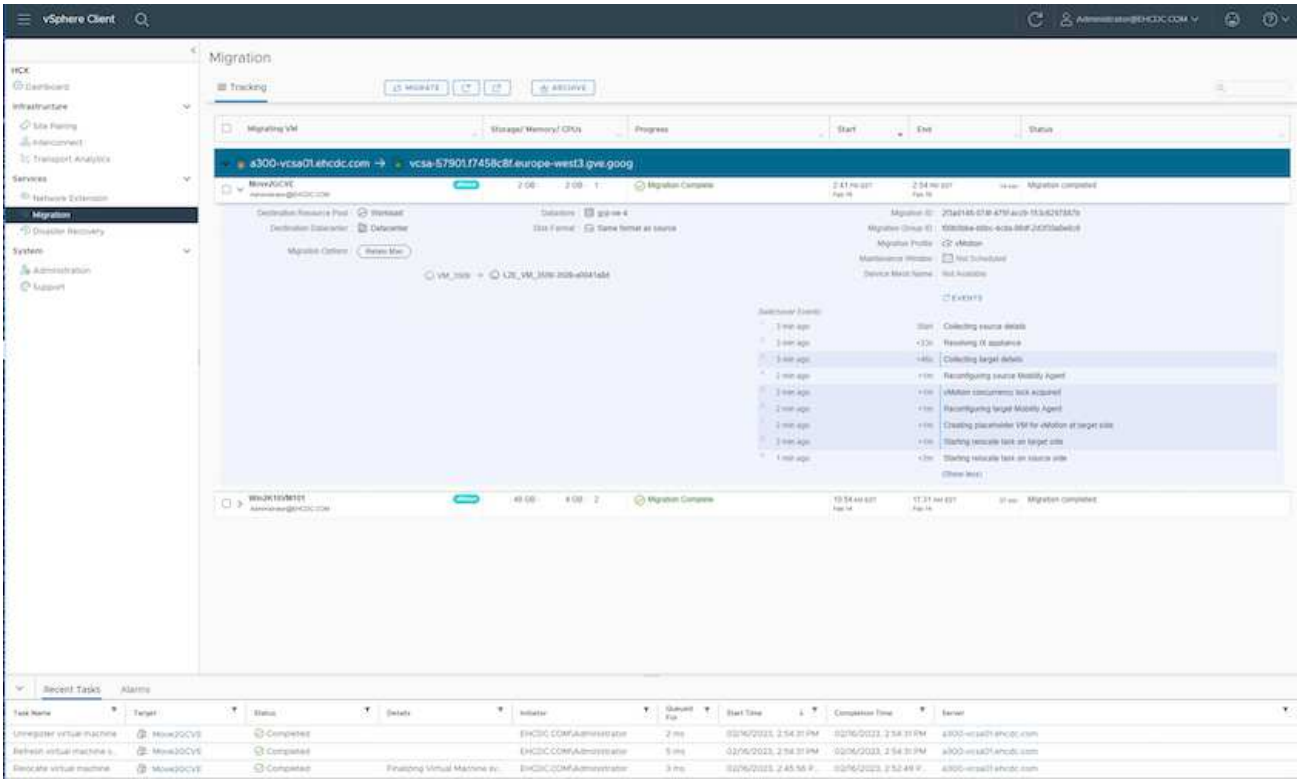
3. Update the mandatory fields (Cluster, Storage, and Destination Network), Click Validate.



4. After the validation checks are complete, click Go to initiate the migration.

> The vMotion transfer captures the VM active memory, its execution state, its IP address, and its MAC address. For more information about the requirements and limitations of HCX vMotion, see Understanding VMware HCX vMotion and Cold Migration.

5. You can monitor the progress and completion of the vMotion from the HCX > Migration dashboard.



The target Google Cloud NetApp Volumes (NetApp Volumes) NFS datastore should have sufficient space to handle the migration.

Conclusion

Whether you're targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in on-premises, Cloud Volume Service and HCX provide excellent options to deploy and migrate the application workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Google Cloud VMware Engine along with Cloud Volume Service for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere Replication, VMware vMotion, or even network file copy (NFC).

Takeaways

The key points of this document include:

- You can now use Cloud Volume Service as a datastore on Google Cloud VMware Engine SDDC.
- You can easily migrate data from on-premises to Cloud Volume Service datastore.
- You can easily grow and shrink the Cloud Volume Service datastore to meet the capacity and performance requirements during migration activity.

Videos from Google and VMware for reference

▼ **From Google**

- Deploy HCX Connector with GCVE
- Configure HCX ServiceMesh with GCVE

- Migrate VM with HCX to GCVE

▼ **From VMware**

- HCX Connector deployment for GCVE
- HCX ServiceMesh configuration for GCVE
- HCX Workload Migration to GCVE

Where to find additional information

To learn more about the information described in this document, refer to the following website links:

- Google Cloud VMware Engine documentation

  https://cloud.google.com/vmware-engine/docs/overview

- Cloud Volume Service documentation

  https://docs.cloud.google.com/netapp/volumes/docs/discover/overview

- VMware HCX User Guide

  https://docs.vmware.com/en/VMware-HCX/index.html

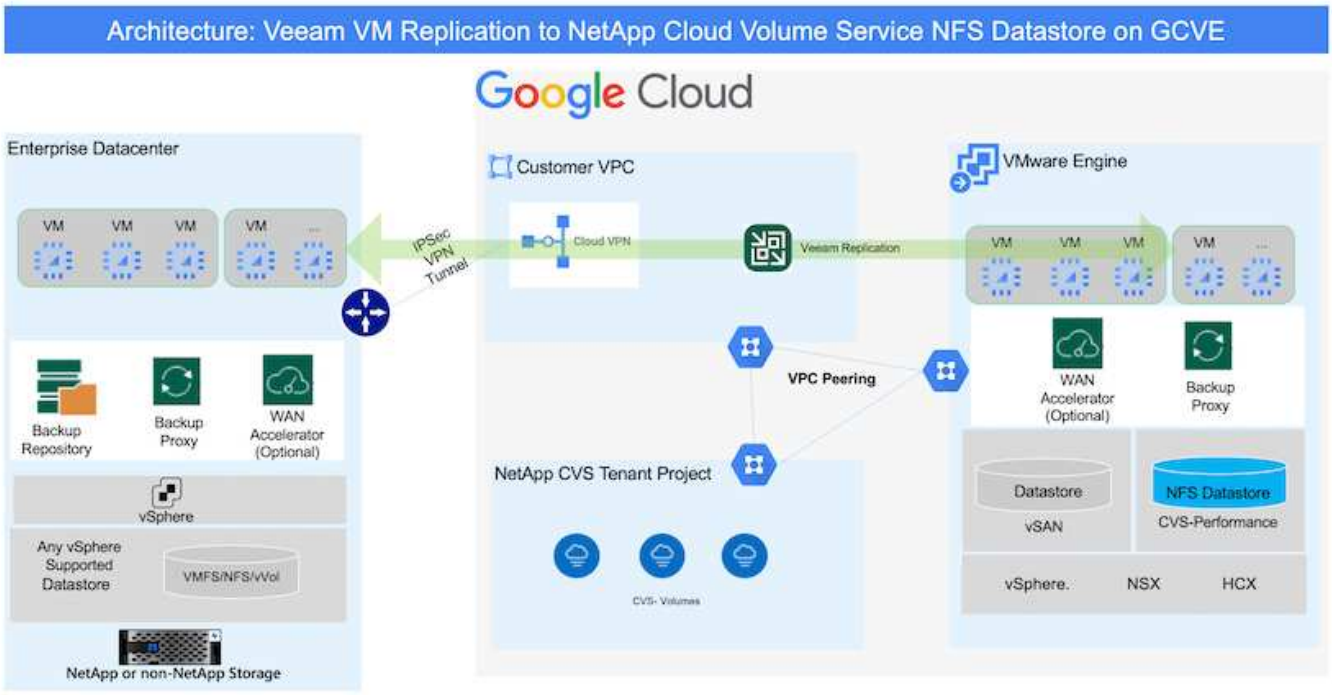
VM Migration to Google Cloud NetApp Volumes NFS Datastore on Google Cloud VMware Engine using Veeam Replication feature

Customers who currently use Veeam for their data protection requirements continue using that solution to migrate the workloads to GCVE and enjoy the benefits of Google Cloud NetApp Volumes NFS Datastores.

Overview

VM Workloads running on VMware vSphere can be migrated to Google Cloud VMware Engine (GCVE) utilizing Veeam Replication feature.

This document provides a step-by-step approach for setting up and performing VM migration that uses Google Cloud NetApp Volumes, Veeam, and the Google Cloud VMware Engine (GCVE).



Assumptions

This document assumes you have either Google Cloud VPN or Cloud Interconnect or other networking option in place to establish network connectivity from existing vSphere servers to Google Cloud VMware Engine.

> There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document.
> Refer to the Google Cloud documentation for the appropriate on-premises-to-Google connectivity method.

Deploying the Migration Solution

Solution Deployment Overview

1. Make sure NFS datastore from Google Cloud NetApp Volumes is mounted on GCVE vCenter.
2. Ensure Veeam Backup Recovery is deployed on existing VMware vSphere environment
3. Create Replication Job to start replicating virtual machines to Google Cloud VMware Engine instance.
4. Perform Failover of Veeam Replication Job.
5. Perform Permanent Failover on Veeam.

Deployment Details

Make sure NFS datastore from Google Cloud NetApp Volumes is mounted on GCVE vCenter

Login to GCVE vCenter and ensure NFS datastore with sufficient space is available.
If not, Please refer Mount NetApp Volumes as NFS datastore on GCVE

Ensure Veeam Backup Recovery is deployed on existing VMware vSphere environment

Please refer Veeam Replication Components documentation to install required components.

Create Replication Job to start replicating virtual machines to Google Cloud VMware Engine instance.

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. Setup vSphere VM Replication Job
Here is a video explaining how to
Configure Replication Job.

> Replica VM can have different IP from the source VM and can also be connected to different port group. For more details, check the video above.

Perform Failover of Veeam Replication Job

To Migrate VMs, perform Perform Failover

Perform Permanent Failover on Veeam.

To treat GCVE as your new source environment, perform Permanent Failover

Benefits of this solution

- Existing Veeam backup infrastructure can be utilized for migration.
- Veeam Replication allows changing VM IP addresses on target site.
- Has ability to remap existing data replicated outside of Veeam (like replicated data from NetApp System)
- Has ability to specify different network portgroup on target site.
- Can specify the order of VMs to power on.
- Utilizes VMware Change Block Tracking to minimize the amount of data to send across WAN.
- Capability to execute pre and post scripts for replication.
- Capability to execute pre and post scripts for snapshots.

Secure data with Google Cloud NetApp Volumes in Google Cloud

TR-4918: Security overview - Google Cloud NetApp Volumes in Google Cloud

Oliver Krause, Justin Parisi, NetApp

Security, particularly in the cloud where infrastructure is outside of the control of storage administrators, is paramount to trusting your data to service offerings provided by cloud providers. This document is an overview of the security offerings that Google Cloud NetApp Volumes provides in Google Cloud.

Intended audience

This document's intended audience includes, but is not limited to, the following roles:

- Cloud providers
- Storage administrators
- Storage architects
- Field resources

- Business decision makers

If you have questions about the content of this technical report, see the section "Contact us."

Abbreviation	Definition
NetApp Volumes-SW	Google Cloud NetApp Volumes, Service Type NetApp Volumes
NetApp Volumes-Performance	Cloud Volume Service, Service Type NetApp Volumes-Performance
PSA	

How Google Cloud NetApp Volumes in Google Cloud secures your data

Google Cloud NetApp Volumes in Google Cloud provides a multitude of ways to natively secure your data.

Secure architecture and tenancy model

Google Cloud NetApp Volumes provides a secure architecture in Google Cloud by segmenting the service management (control plane) and the data access (data plane) across different endpoints so that neither can impact the other (see the section "Google Cloud NetApp Volumes architecture"). It uses Google's private services access (PSA) framework to provide the service. This framework distinguishes between the service producer, which is provided and operated by NetApp, and the service consumer, which is a Virtual Private Cloud (VPC) in a customer project, hosting the clients that want to access Google Cloud NetApp Volumes file shares.

In this architecture, tenants (see the section "Tenancy model") are defined as Google Cloud projects that are completely isolated from each other unless explicitly connected by the user. Tenants allow complete isolation of data volumes, external name services, and other essential pieces of the solution from other tenants using the Google Cloud NetApp Volumes volume platform. Because the Google Cloud NetApp Volumes platform is connected through VPC peering, that isolation applies to it also. You can enable sharing of Google Cloud NetApp Volumes volumes between multiple projects by using a shared-VPC (see the section "Shared VPCs"). You can apply access controls to SMB shares and NFS exports to limit who or what can view or modify datasets.

Strong identity management for the control plane

In the control plane where Google Cloud NetApp Volumes configuration takes place, identity management is managed by using Identity Access Management (IAM). IAM is a standard service that enables you to control authentication (logins) and authorization (permissions) to Google Cloud project instances. All configuration is performed with Google Cloud NetApp Volumes APIs over a secure HTTPS transport using TLS 1.2 encryption, and authentication is performed by using JWT tokens for added security. The Google console UI for Google Cloud NetApp Volumes translates user input into Google Cloud NetApp Volumes API calls.

Security hardening - Limiting attack surfaces

Part of effective security is limiting the number of attack surfaces available in a service. Attack surfaces can include a variety of things, including data at-rest, in-flight transfers, logins, and the datasets themselves.

A managed service removes some of the attack surfaces inherently in its design. Infrastructure management, as described in the section "Service operation," is handled by a dedicated team and is automated to reduce the number of times a human actually touches configurations, which helps reduce the number of intentional and

unintentional errors. Networking is fenced off so that only necessary services can access one another. Encryption is baked into the data storage and only the data plane needs security attention from Google Cloud NetApp Volumes administrators. By hiding most of the management behind an API interface, security is achieved by limiting the attack surfaces.

Zero Trust model

Historically, IT security philosophy has been to trust but verify, and manifested as relying solely on external mechanisms (such as firewalls and intrusion detection systems) to mitigate threats. However, attacks and breaches evolved to bypass the verification in environments through phishing, social engineering, insider threats and other methods that provide the verification to enter networks and wreak havoc.

Zero Trust has become a new methodology in security, with the current mantra being "trust nothing while still verifying everything." Therefore, nothing is allowed access by default. This mantra is enforced in a variety of ways, including standard firewalls and intrusion detection systems (IDS) and also with the following methods:

- Strong authentication methods (such as AES-encrypted Kerberos or JWT tokens)
- Single strong sources of identities (such as Windows Active Directory, Lightweight Directory Access Protocol (LDAP), and Google IAM)
- Network segmentation and secure multitenancy (only tenants are allowed access by default)
- Granular access controls with Least Privileged Access policies
- Small exclusive lists of dedicated, trusted administrators with digital audit and paper trails

Google Cloud NetApp Volumes running in Google Cloud adheres to the Zero Trust model by implementing the "trust nothing, verify everything" stance.

Encryption

Encrypt data at-rest (see the section "Data encryption at rest") by using XTS-AES-256 ciphers with NetApp Volume Encryption (NVE) and in-flight with "SMB encryption" or NFS Kerberos 5p support. Rest easy knowing cross-region replication transfers are protected by TLS 1.2 encryption (see the section "Cross-region replication"). In addition, Google networking also provides encrypted communications (see the section "Data encryption in transit") for an added layer of protection against attacks. For more information about transport encryption, see the section "Google Cloud network".

Data protection and backups

Security isn't just about the prevention of attacks. It is also about how we recover from attacks if or when they occur. This strategy includes data protection and backups. Google Cloud NetApp Volumes provides methods to replicate to other regions in case of outages (see the section "Cross-region replication") or if a dataset is affected by a ransomware attack. It can also perform asynchronous backups of data to locations outside of the Google Cloud NetApp Volumes instance by using Google Cloud NetApp Volumes backup. With regular backups, mitigation of security events can take less time and save money and angst for administrators.

Fast ransomware mitigation with industry leading Snapshot copies

In addition to data protection and backups, Google Cloud NetApp Volumes provides support for immutable Snapshot copies (see the section "Immutable Snapshot copies") of volumes that allow recovery from ransomware attacks (see the section "Service operation") within seconds of discovering the issue and with minimal disruption. Recovery time and effects depend on the Snapshot schedule, but you can create Snapshot copies that provide as little as one-hour deltas in ransomware attacks. Snapshot copies have a negligible effect on performance and capacity usage and are a low-risk, high-reward approach to protecting your datasets.

Security considerations and attack surfaces

The first step in understanding how to secure your data is identifying the risks and potential attack surfaces.

These include (but are not limited to) the following:

- Administration and logins
- Data at rest
- Data in flight
- Network and firewalls
- Ransomware, malware, and viruses

Understanding attack surfaces can help you to better secure your environments. Google Cloud NetApp Volumes in Google Cloud already considers many of these topics and implements security functionality by default, without any administrative interaction.

Ensuring secure logins

When securing your critical infrastructure components, it is imperative to make sure that only approved users can log in and manage your environments. If bad actors breach your administrative credentials, then they have the keys to the castle and can do anything they want—change configurations, delete volumes and backups, create backdoors, or disable Snapshot schedules.

Google Cloud NetApp Volumes for Google Cloud provides protection against unauthorized administrative logins through the obfuscation of storage as a service (StaaS). Google Cloud NetApp Volumes is completely maintained by the cloud provider with no availability to login externally. All setup and configuration operations are fully automated, so a human administrator never has to interact with the systems except in very rare circumstances.

If login is required, Google Cloud NetApp Volumes in Google Cloud secures logins by maintaining a very short list of trusted administrators that have access to log in to the systems. This gatekeeping helps reduce the number of potential bad actors with access. Additionally, the Google Cloud networking hides the systems behind layers of network security and exposes only what is needed to the outside world. For information about the Google Cloud, Google Cloud NetApp Volumes architecture, see the section "Google Cloud NetApp Volumes architecture."

Cluster administration and upgrades

Two areas with potential security risks include cluster administration (what happens if a bad actor has admin access) and upgrades (what happens if a software image is compromised).

Storage administration protection

Storage provided as a service removes the added risk of exposure to administrators by removing that access to end users outside of the cloud data center. Instead, the only configuration done is for the data access plane by customers. Each tenant manages their own volumes, and no tenant can reach other Google Cloud NetApp Volumes instances. The service is managed by automation, with a very small list of trusted administrators given access to the systems through the processes covered in the section "Service operation."

The NetApp Volumes-Performance service type offers cross-region replication as an option to provide data protection to a different region in the event of a region failure. In those cases, Google Cloud NetApp Volumes can be failed over to the unaffected region to maintain data access.

Service upgrades

Updates help protect vulnerable systems. Each update provides security enhancements and bug fixes that minimize attack surfaces. Software updates are downloaded from centralized repositories and are validated before the updates are allowed to verify that official images are used and that the upgrades are not compromised by bad actors.

With Google Cloud NetApp Volumes, updates are handled by the cloud provider teams, which removes risk exposure for administrator teams by providing experts well versed in configuration and upgrades that have automated and fully tested the process. Upgrades are nondisruptive, and Google Cloud NetApp Volumes maintains the latest updates for best overall results.

For information about the administrator team that performs these service upgrades, see the section "Service operation."

Securing data at-rest

Data-at-rest encryption is important to protect sensitive data in the event of a disk that is stolen, returned, or repurposed. Data in Google Cloud NetApp Volumes is protected at rest by using software-based encryption.

- Google-generated keys are used for NetApp Volumes-SW.
- For NetApp Volumes-Performance, the per-volume keys are stored in a key manager built into Google Cloud NetApp Volumes, which uses NetApp ONTAP CryptoMod to generate AES-256 encryption keys. CryptoMod is listed on the CMVP FIPS 140-2 validated modules list. See FIPS 140-2 Cert #4144.

Starting in November 2021, preview Customer-managed Encryption (CMEK) functionality was made available for NetApp Volumes-Performance. This functionality allows you to encrypt the per-volume keys with per-project, per-region master-keys that are hosted in Google Key Management Service (KMS). KMS enables you to attach external key managers.

For details about how to configure KMS for NetApp Volumes-Performance, see the Google Cloud NetApp Volumes documentation.

For more information about architecture, see the section "Google Cloud NetApp Volumes architecture."

Securing data in-flight

In addition to securing data at rest, you must also be able to secure data when it is in flight between the Google Cloud NetApp Volumes instance and a client or replication target. Google Cloud NetApp Volumes provides encryption for in-flight data over NAS protocols by using encryption methods such as SMB encryption using Kerberos, the signing/sealing of packets, and NFS Kerberos 5p for end-to-end encryption of data transfers.

Replication of Google Cloud NetApp Volumes volumes uses TLS 1.2, which takes advantage of AES-GCM encryption methods.

Most insecure in-flight protocols such as telnet, NDMP, and so on are disabled by default. DNS, however, is not encrypted by Google Cloud NetApp Volumes (no DNS Sec support) and should be encrypted by using external network encryption when possible. See the section "Data encryption in transit" for more information about securing data in-flight.

For information about NAS protocol encryption, see the section "NAS protocols."

Users and groups for NAS permissions

Part of securing your data in the cloud involves proper user and group authentication, where the users

accessing the data are verified as real users in the environment and the groups contain valid users. These users and groups provide initial share and export access, as well as permission validation for files and folders in the storage system.

Google Cloud NetApp Volumes uses standard Active Directory-based Windows user and group authentication for SMB shares and Windows-style permissions. The service can also leverage UNIX identity providers such as LDAP for UNIX users and groups for NFS exports, NFSv4 ID validation, Kerberos authentication, and NFSv4 ACLs.

> Currently only Active Directory LDAP is supported with Google Cloud NetApp Volumes for LDAP functionality.

Detection, prevention and mitigation of ransomware, malware, and viruses

Ransomware, malware, and viruses are a persistent threat to administrators, and detection, prevention, and mitigation of those threats are always top of mind for enterprise organizations. A single ransomware event on a critical dataset can potentially cost millions of dollars, so it is beneficial to do what you can to minimize the risk.

Although Google Cloud NetApp Volumes currently doesn't include native detection or prevention measures, such as antivirus protection or automatic ransomware detection, there are ways to quickly recover from a ransomware event by enabling regular Snapshot schedules. Snapshot copies are immutable and read only pointers to changed blocks in the file system, are near instantaneous, have minimal impact on performance, and only use up space when data is changed or deleted. You can set schedules for Snapshot copies to match your desired acceptable recovery point objective (RPO)/recovery time objective (RTO) and can keep up to 1,024 Snapshot copies per volume.

Snapshot support is included at no additional cost (beyond data storage charges for changed blocks/data retained by Snapshot copies) with Google Cloud NetApp Volumes and, in the event of a ransomware attack, can be used to roll back to a Snapshot copy before the attack occurred. Snapshot restores take just seconds to complete, and you then can get back to serving data as normal. For more information, see The NetApp Solution for Ransomware.

Preventing ransomware from affecting your business requires a multilayered approach that includes one or more of the following:

- Endpoint protection
- Protection against external threats through network firewalls
- Detection of data anomalies
- Multiple backups (onsite and offsite) of critical datasets
- Regular restore tests of backups
- Immutable read-only NetApp Snapshot copies
- Multifactor authentication for critical infrastructure
- Security audits of system logins

This list is far from exhaustive but is a good blueprint to follow when dealing with the potential of ransomware attacks. Google Cloud NetApp Volumes in Google Cloud provides several ways to protect against ransomware events and reduce their effects.

Immutable Snapshot copies

Google Cloud NetApp Volumes natively provides immutable read-only Snapshot copies that are taken on a

customizable schedule for quick point-in-time recovery in the event of data deletion or if an entire volume has been victimized by a ransomware attack. Snapshot restores to previous good Snapshot copies are fast and minimize data loss based on the retention period of your Snapshot schedules and RTO/RPO. The performance effect with Snapshot technology is negligible.

Because Snapshot copies in Google Cloud NetApp Volumes are read-only, they cannot be infected by ransomware unless the ransomware has proliferated into the dataset unnoticed and Snapshot copies have been taken of the data infected by ransomware. This is why you must also consider ransomware detection based on data anomalies. Google Cloud NetApp Volumes does not currently provide detection natively, but you can use external monitoring software.

Backups and restores

Google Cloud NetApp Volumes provides standard NAS client backup capabilities (such as backups over NFS or SMB).

- NetApp Volumes-Performance offers cross-region volume replication to other NetApp Volumes-Performance volumes. For more information, see volume replication in the Google Cloud NetApp Volumes documentation.
- NetApp Volumes-SW offers service-native volume backup/restore capabilities. For more information, see cloud backup in the Google Cloud NetApp Volumes documentation.

Volume replication provides an exact copy of the source volume for fast failover in the case of a disaster, including ransomware events.

Cross-region replication

NetApp Volumes-Performance enables you to securely replicate volumes across Google Cloud regions for data protection and archive use cases by using TLS1.2 AES 256 GCM encryption on a NetApp-controlled backend service network using specific interfaces used for replication running on Google's network. A primary (source) volume contains the active production data and replicates to a secondary (destination) volume to provide an exact replica of the primary dataset.

Initial replication transfers all blocks, but updates only transmit the changed blocks in a primary volume. For instance, if a 1TB database that resides on a primary volume is replicated to the secondary volume, then 1TB of space is transferred on the initial replication. If that database has a few hundred rows (hypothetically, a few MB) that change between the initialization and the next update, only the blocks with the changed rows are replicated to the secondary (a few MB). This helps to make sure that the transfer times remain low and keeps replication charges down.

All permissions on files and folders are replicated to the secondary volume, but share access permissions (such as export policies and rules or SMB shares and share ACLs) must be handled separately. In the case of a site failover, the destination site should leverage the same name services and Active Directory domain connections to provide consistent handling of user and group identities and permissions. You can use a secondary volume as a failover target in the event of a disaster by breaking the replication relationship, which converts the secondary volume to read-write.

Volume replicas are read-only, which provides an immutable copy of data offsite for quick recovery of data in instances where a virus has infected data or ransomware has encrypted the primary dataset. Read-only data won't be encrypted, but, if the primary volume is affected and replication occurs, the infected blocks also replicate. You can use older, non-affected Snapshot copies to recover, but SLAs might fall out of range of the promised RTO/RPO depending on how quickly an attack is detected.

In addition, you can prevent malicious administrative actions, such as volume deletions, Snapshot deletions, or

Snapshot schedule changes, with cross-region replication (CRR) management in Google Cloud. This is done by creating custom roles that separate volume administrators, who can delete source volumes but not break mirrors and therefore cannot delete destination volumes, from CRR administrators, who cannot perform any volume operations. See Security Considerations in the Google Cloud NetApp Volumes documentation for permissions allowed by each administrator group.

Google Cloud NetApp Volumes backup

Although Google Cloud NetApp Volumes provides high data durability, external events can cause data loss. In the event of a security event such as a virus or ransomware, backups and restores become critical for resumption of data access in a timely manner. An administrator might accidentally delete a Google Cloud NetApp Volumes volume. Or users simply want to retain backup versions of their data for many months and keeping the extra Snapshot copy space inside the volume becomes a cost challenge. Although Snapshot copies should be the preferred way to keep backup versions for the last few weeks to restore lost data from them, they are sitting inside the volume and are lost if the volume goes away.

For all these reasons, Google Cloud NetApp Volumes offers backup services through Google Cloud NetApp Volumes backup.

Google Cloud NetApp Volumes backup generates a copy of the volume on Google Cloud Storage (GCS). It only backs up the actual data stored within the volume, not the free space. It works as incremental forever, meaning it transfers the volume content once and from there on continues backing up changed data only. Compared to classical backup concepts with multiple full backups, it saves large amounts of backup storage, reducing cost. Because the monthly price of backup space is lower compared to a volume, it is an ideal place to keep backup versions longer.

Users can use a Google Cloud NetApp Volumes backup to restore any backup version to the same or a different volume within the same region. If the source volume is deleted, the backup data is retained and needs to be managed (for example, deleted) independently.

Google Cloud NetApp Volumes backup is built into Google Cloud NetApp Volumes as option. Users can decide which volumes to protect by activating Google Cloud NetApp Volumes backup on a per-volume basis. See the Google Cloud NetApp Volumes backup documentation for information about backups, the number of maximum backup versions supported, scheduling, and pricing.

All backup data of a project is stored within a GCS bucket, which is managed by the service and not visible to the user. Each project uses a different bucket. Currently, the buckets are in same region as the Google Cloud NetApp Volumes volumes, but more options are being discussed. Consult the documentation for the latest status.

Data transport from a Google Cloud NetApp Volumes bucket to GCS uses service-internal Google networks with HTTPS and TLS1.2. Data is encrypted at-rest with Google-managed keys.

To manage Google Cloud NetApp Volumes backup (creating, deleting, and restoring backups), a user must have the roles/netapp.admin role.

Overview

Part of trusting a cloud solution is understanding the architecture and how it is secured. This section calls out different aspects of the Google Cloud NetApp Volumes architecture in Google to help alleviate potential concerns about how data is secured, as well as call out areas where additional configuration steps might be required to obtain the most secure deployment.

The general architecture of Google Cloud NetApp Volumes can be broken down into two main components: the control plane and the data plane.

Control plane

The control plane in Google Cloud NetApp Volumes is the backend infrastructure managed by Google Cloud NetApp Volumes administrators and NetApp native automation software. This plane is completely transparent to end users and includes networking, storage hardware, software updates, and so on to help deliver value to a cloud-resident solution such as Google Cloud NetApp Volumes.

Data plane

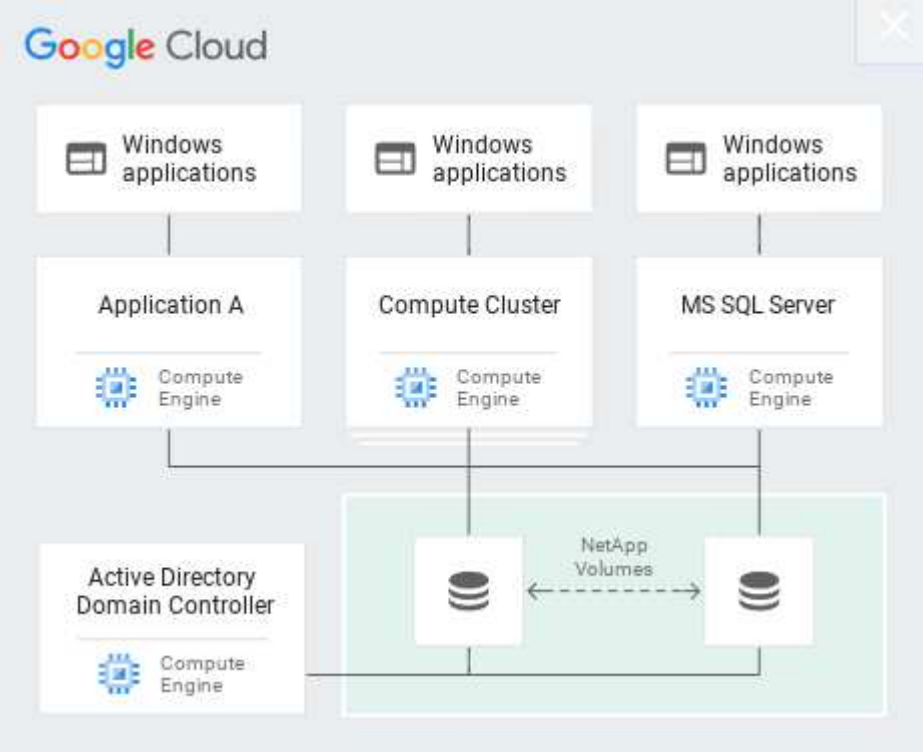
The data plane in Google Cloud NetApp Volumes includes the actual data volumes and the overall Google Cloud NetApp Volumes configuration (such as access control, Kerberos authentication, and so on). The data plane is entirely under the control of the end users and the consumers of the Google Cloud NetApp Volumes platform.

There are distinct differences in how each plane is secured and managed. The following sections cover these differences, starting with a Google Cloud NetApp Volumes architecture overview.

Google Cloud NetApp Volumes architecture

In a manner similar to other Google Cloud native services such as CloudSQL, Google Cloud VMware Engine (GCVE), and FileStore, Google Cloud NetApp Volumes uses Google PSA to deliver the service. In PSA, services are built inside a service producer project, which uses VPC network peering to connect to the service consumer. The service producer is provided and operated by NetApp, and the service consumer is a VPC in a customer project, hosting the clients that want to access Google Cloud NetApp Volumes file shares.

The following figure, referenced from the Google Cloud NetApp Volumes documentation, reviews a high-level view.



The part above the dotted line shows the control plane of the service, which controls the volume lifecycle. The part below the dotted line shows the data plane. The left blue box depicts the user VPC (service consumer), the right blue box is the service producer provided by NetApp. Both are connected through VPC peering.

Tenancy model

In Google Cloud NetApp Volumes, individual projects are considered unique tenants. This means that manipulation of volumes, Snapshot copies, and so on are performed on a per- project basis. In other words, all volumes are owned by the project that they were created in and only that project can manage and access the data inside of them by default. This is considered the control plane view of the service.

Shared VPCs

On the data plane view, Google Cloud NetApp Volumes can connect to a shared VPC. You can create volumes in the hosting project or in one of the service projects connected to the shared VPC. All projects (host or service) connected to that shared VPC are able to reach the volumes at the network layer (TCP/IP). Because all clients with network connectivity on the shared- VPC can potentially access the data through NAS protocols, access control on the individual volume (such as user/group access control lists (ACLs) and hostnames/IP addresses for NFS exports) must be used to control who can access the data.

You can connect Google Cloud NetApp Volumes to up to five VPCs per customer project. On the control plane, the project enables you to manage all created volumes, no matter which VPC they are connected to. On the data plane, VPCs are isolated from one another, and each volume can only be connected to one VPC.

Access to the individual volumes is controlled by protocol specific (NFS/SMB) access control mechanisms.

In other words, on the network layer, all project s connected to the shared VPC are able to see the volume, while, on the management side, the control plane only allows the owner project to see the volume.

VPC Service Controls

VPC Service Controls establish an access control perimeter around Google Cloud services that are attached to the internet and are accessible worldwide. These services provide access control through user identities but cannot restrict which network location requests originate from. VPC Service Controls close that gap by introducing the capabilities to restrict access to defined networks.

The Google Cloud NetApp Volumes data plane is not connected to the external internet but to private VPCs with well-defined network boundaries (perimeters). Within that network, each volume uses protocol-specific access control. Any external network connectivity is explicitly created by Google Cloud project administrators. The control plane, however, does not provide the same protections as the data plane and can be accessed by anyone from anywhere with valid credentials (JWT tokens).

In short, the Google Cloud NetApp Volumes data plane provides the capability of network access control, without the requirement to support VPC Service Controls and does not explicitly use VPC Service Controls.

Packet sniffing/trace considerations

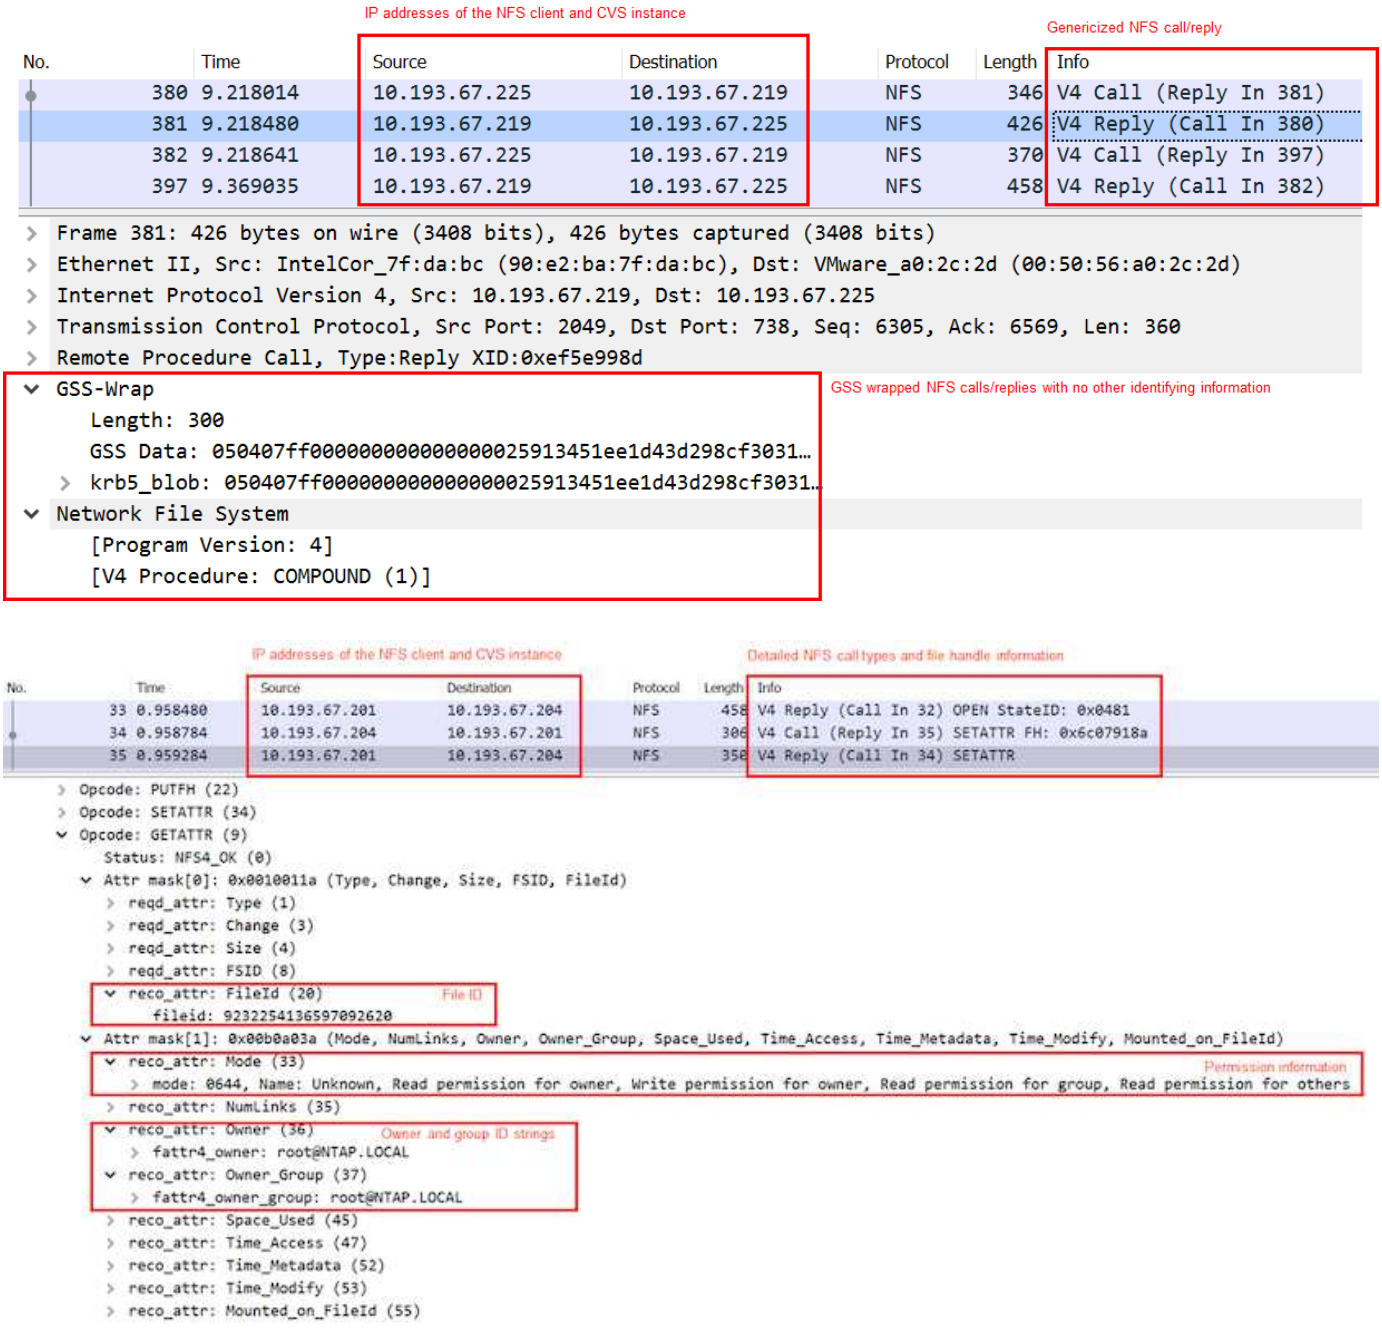
Packet captures can be useful for troubleshooting network issues or other problems (such as NAS permissions, LDAP connectivity, and so on), but can also be used maliciously to gain information about network IP addresses, MAC addresses, user and group names, and what level of security is being used on endpoints. Because of the way Google Cloud networking, VPCs, and firewall rules are configured, unwanted access to network packets should be difficult to obtain without user login credentials or JWT tokens into the cloud instances. Packet captures are only possible on endpoints (such as virtual machines (VMs)) and only possible on endpoints internal to the VPC unless a shared VPC and/or external network tunnel/IP forwarding is in use to explicitly allow external traffic to endpoints. There is no way to sniff traffic outside of the clients.

When shared VPCs are used, in-flight encryption with NFS Kerberos and/or SMB encryption can mask much of the information gleaned from traces. However, some traffic is still sent in plaintext, such as DNS and LDAP queries. The following figure shows a packet capture from a plaintext LDAP query originating from Google Cloud NetApp Volumes and the potential identifying information that is exposed. LDAP queries in Google Cloud NetApp Volumes currently do not support encryption or LDAP over SSL. NetApp Volumes-Performance support LDAP signing, if requested by Active Directory. NetApp Volumes-SW does not support LDAP signing.



unixUserPassword is queried by LDAP and is not sent in plaintext but instead in a salted hash. By default, Windows LDAP does not populate the unixUserPassword fields. This field is only required if you need to leverage Windows LDAP for interactive logins through LDAP to clients. Google Cloud NetApp Volumes does not support interactive LDAP logins to the instances.

The following figure shows a packet capture from an NFS Kerberos conversation next to a capture of NFS over AUTH_SYS. Note how the information available in a trace differs between the two and how enabling in-flight encryption offers greater overall security for NAS traffic.



VM network interfaces

One trick attackers might attempt is to add a new network interface card (NIC) to a VM in promiscuous mode (port mirroring) or enable promiscuous mode on an existing NIC in order to sniff all traffic. In Google Cloud, adding a new NIC requires a VM to be shut down entirely, which creates alerts, so attackers cannot do this

unnoticed.

In addition, NICs cannot be set to promiscuous mode at all and will trigger alerts in Google Cloud.

Control plane architecture

All management actions to Google Cloud NetApp Volumes are done through API. Google Cloud NetApp Volumes management integrated into the GCP Cloud Console also uses the Google Cloud NetApp Volumes API.

Identity and Access Management

Identity and Access Management (IAM) is a standard service that enables you to control authentication (logins) and authorization (permissions) to Google Cloud project instances. Google IAM provides a full audit trail of permissions authorization and removal. Currently Google Cloud NetApp Volumes does not provide control plane auditing.

Authorization/permission overview

IAM offers built-in, granular permissions for Google Cloud NetApp Volumes. You can find a complete list of granular permissions here.

IAM also offers two predefined roles called `netappcloudvolumes.admin` and `netappcloudvolumes.viewer`. These roles can be assigned to specific users or service accounts.

Assign appropriate roles and permission to allow IAM users to manage Google Cloud NetApp Volumes.

Examples for using granular permissions include the following:

- Build a custom role with only get/list/create/update permissions so that users cannot delete volumes.
- Use a custom role with only `snapshot.*` permissions to create a service account that is used to build application- consistent Snapshot integration.
- Build a custom role to delegate `volumereplication.*` to specific users.

Service accounts

To make Google Cloud NetApp Volumes API calls through scripts or Terraform, you must create a service account with the `roles/netappcloudvolumes.admin` role. You can use this service account to generate the JWT tokens required to authenticate Google Cloud NetApp Volumes API requests in two different ways:

- Generate a JSON key and use Google APIs to derive a JWT token from it. This is the simplest approach, but it involves manual secrets (the JSON key) management.
- Use Service account impersonation with `roles/iam.serviceAccountTokenCreator`. The code (script, Terraform, and so on.) runs with Application Default Credentials and impersonates the service account to gain its permissions. This approach reflects Google security best practices.

See Creating your service account and private key in the Google cloud documentation for more information.

Google Cloud NetApp Volumes API

Google Cloud NetApp Volumes API uses a REST-based API by using HTTPS (TLSv1.2) as the underlying network transport. You can find the latest API definition and information about how to use the API at Cloud

Volumes APIs in the Google cloud documentation.

The API endpoint is operated and secured by NetApp using standard HTTPS (TLSv1.2) functionality.

JWT tokens

Authentication to the API is performed with JWT bearer tokens (RFC-7519). Valid JWT tokens must be obtained by using Google Cloud IAM authentication. This must be done by fetching a token from IAM by providing a service account JSON key.

Audit logging

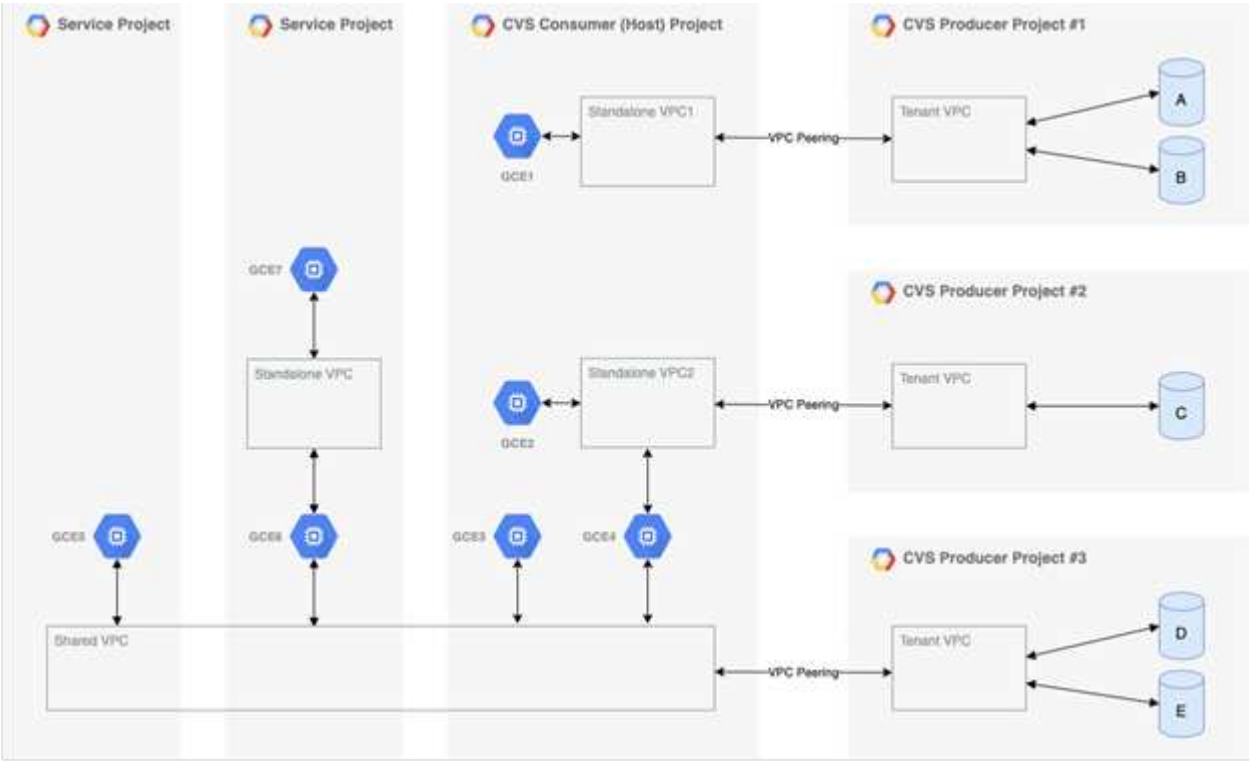
Currently, no user-accessible control plane audit logs are available.

Data plane architecture

Google Cloud NetApp Volumes for Google Cloud leverages the Google Cloud private services access framework. In this framework, users can connect to the Google Cloud NetApp Volumes. This framework uses Service Networking and VPC peering constructs like other Google Cloud services, ensuring complete isolation between tenants.

For an architecture overview of Google Cloud NetApp Volumes for Google Cloud, see Architecture for Google Cloud NetApp Volumes.

User VPCs (standalone or shared) are peered to VPCs within Google Cloud NetApp Volumes managed tenant projects, which hosts the volumes.



The preceding figure shows a project (the NetApp Volumes consumer project in the middle) with three VPC networks connected to Google Cloud NetApp Volumes and multiple Compute Engine VMs (GCE1-7) sharing volumes:

- VPC1 allows GCE1 to access volumes A and B.
- VPC2 allows GCE2 and GCE4 to access volume C.
- The third VPC network is a shared VPC, shared with two service projects. It allows GCE3, GCE4, GCE5, and GCE6 to access volumes D and E. Shared VPC networks are only supported for volumes of the NetApp Volumes-Performance service type.

GCE7 cannot access any volume.

Data can be encrypted both in-transit (using Kerberos and/or SMB encryption) and at-rest in Google Cloud NetApp Volumes.

Data encryption in transit

Data in transit can be encrypted at the NAS protocol layer, and the Google Cloud network itself is encrypted, as described in the following sections.

Google Cloud network

Google Cloud encrypts traffic on the network level as described in Encryption in transit in the Google documentation. As mentioned in the section "Google Cloud NetApp Volumes architecture," Google Cloud NetApp Volumes is delivered out of a NetApp-controlled PSA producer project.

In case of NetApp Volumes-SW, the producer tenant runs Google VMs to provide the service. Traffic between user VMs and Google Cloud NetApp Volumes VMs is encrypted automatically by Google.

Although the data path for NetApp Volumes-Performance isn't fully encrypted on the network layer, NetApp and Google use a combination of IEEE 802.1AE encryption (MACSec), encapsulation (data encryption), and physically restricted networks to protect data in transit between the Google Cloud NetApp Volumes NetApp Volumes-Performance service type and Google Cloud.

NAS protocols

NFS and SMB NAS protocols provide optional transport encryption at the protocol layer.

SMB encryption

SMB encryption provides end-to-end encryption of SMB data and protects data from eavesdropping occurrences on untrusted networks. You can enable encryption for both the client/server data connection (only available to SMB3.x capable clients) and the server/domain controller authentication.

When SMB encryption is enabled, clients that do not support encryption cannot access the share.

Google Cloud NetApp Volumes supports RC4-HMAC, AES-128-CTS-HMAC-SHA1, and AES-256-CTS-HMAC-SHA1 security ciphers for SMB encryption. SMB negotiates to the highest supported encryption type by the server.

NFSv4.1 Kerberos

For NFSv4.1, NetApp Volumes-Performance offers Kerberos authentication as described in RFC7530. You can enable Kerberos on a per-volume basis.

The current strongest available encryption type for Kerberos is AES-256-CTS-HMAC-SHA1. Google Cloud NetApp Volumes supports AES-256-CTS-HMAC-SHA1, AES-128-CTS-HMAC-SHA1, DES3, and DES for

NFS. It also supports ARCFOUR-HMAC (RC4) for CIFS/SMB traffic, but not for NFS.

Kerberos provides three different security levels for NFS mounts that offer choices for how strong the Kerberos security should be.

As per RedHat's Common Mount Options documentation:

```
sec=krb5 uses Kerberos V5 instead of local UNIX UIDs and GIDs to
authenticate users.
sec=krb5i uses Kerberos V5 for user authentication and performs integrity
checking of NFS operations using secure checksums to prevent data
tampering.
sec=krb5p uses Kerberos V5 for user authentication, integrity checking,
and encrypts NFS traffic to prevent traffic sniffing. This is the most
secure setting, but it also involves the most performance overhead.
```

As a general rule, the more the Kerberos security level has to do, the worse the performance is, as the client and server spend time encrypting and decrypting NFS operations for each packet sent. Many clients and NFS servers provide support for AES-NI offloading to the CPUs for a better overall experience, but the performance impact of Kerberos 5p (full end-to-end encryption) is significantly greater than the impact of Kerberos 5 (user authentication).

The following table shows differences in what each level does for security and performance.

Security level	Security	Performance
NFSv3—sys	• Least secure; plain text with numeric user IDs/group IDs • Able to view UID, GID, client IP addresses, export paths, file names, permissions in packet captures	• Best for most cases
NFSv4.x—sys	• More secure than NFSv3 (client IDs, name string/domain string matching) but still plain text • Able to view UID, GID, client IP addresses, name strings, domain IDs, export paths, file names, permissions in packet captures	• Good for sequential workloads (such as VMs, databases, large files) • Bad with high file count/high metadata (30-50% worse)

Security level	Security	Performance
NFS—krb5	• Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper • User requesting access to mount needs a valid Kerberos ticket (either through username/password or manual key tab exchange); ticket expires after a specified time period and user must reauthenticate for access • No encryption for NFS operations or ancillary protocols like mount/portmapper/nlm (can see export paths, IP addresses, file handles, permissions, file names, atime/mtime in packet captures)	• Best in most cases for Kerberos; worse than AUTH_SYS
NFS—krb5i	• Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper • User requesting access to mount needs a valid Kerberos ticket (either via username/password or manual key tab exchange); ticket expires after a specified time period and user must reauthenticate for access • No encryption for NFS operations or ancillary protocols like mount/portmapper/nlm (can see export paths, IP addresses, file handles, permissions, file names, atime/mtime in packet captures) • Kerberos GSS checksum is added to every packet to ensure nothing intercepts the packets. If checksums match, conversation is allowed.	• Better than krb5p because the NFS payload is not encrypted; only added overhead compared to krb5 is the integrity checksum. Performance of krb5i won’t be much worse than krb5 but will see some degradation.

Security level	Security	Performance
NFS – krb5p	• Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper • User requesting access to mount needs a valid Kerberos ticket (either via username/password or manual keytab exchange); ticket expires after specified time period and user must reauthenticate for access • All of the NFS packet payloads are encrypted with the GSS wrapper (cannot see file handles, permissions, file names, atime/mtime in packet captures). • Includes integrity check. • NFS operation type is visible (FSINFO, ACCESS, GETATTR, and so on). • Ancillary protocols (mount, portmap, nlm, and so on) are not encrypted - (can see export paths, IP addresses)	• Worst performance of the security levels; krb5p has to encrypt/decrypt more. • Better performance than krb5p with NFSv4.x for high file count workloads.

In Google Cloud NetApp Volumes, a configured Active Directory server is used as Kerberos server and LDAP server (to lookup user identities from an RFC2307 compatible schema). No other Kerberos or LDAP servers are supported. NetApp highly recommends that you use LDAP for identity management in Google Cloud NetApp Volumes. For information on how NFS Kerberos is shown in packet captures, see the section "Packet sniffing/trace considerations."

Data encryption at rest

All volumes in Google Cloud NetApp Volumes are encrypted-at-rest using AES-256 encryption, which means all user data written to media is encrypted and can only be decrypted with a per-volume key.

- For NetApp Volumes-SW, Google-generated keys are used.
- For NetApp Volumes-Performance, the per-volume keys are stored in a key manager built into the Google Cloud NetApp Volumes.

Starting in November 2021, preview customer-managed encryption keys (CMEK) functionality was made available. This enables you to encrypt the per-volume keys with a per-project, per-region master key that is hosted in Google Key Management Service (KMS). KMS enables you to attach external key managers.

For information about configuring KMS for NetApp Volumes-Performance, see Setting up customer-managed

encryption keys.

Firewall

Google Cloud NetApp Volumes exposes multiple TCP ports to serve NFS and SMB shares:

- Ports required for NFS access
- Ports required for SMB access

Additionally, SMB, NFS with LDAP including Kerberos, and dual-protocol configurations require access to a Windows Active Directory domain. Active Directory connections must be configured on a per-region basis. Active Directory Domain controllers (DC) are identified by using DNS-based DC discovery using the specified DNS servers. Any of the DCs returned are used. The list of eligible DCs can be limited by specifying an Active Directory site.

Google Cloud NetApp Volumes reaches out with IP addresses from the CIDR range allocated with the `gcloud compute address` command while on-boarding the Google Cloud NetApp Volumes. You can use this CIDR as source addresses to configure inbound firewalls to your Active Directory domain controllers.

Active Directory Domain Controllers must expose ports to the Google Cloud NetApp Volumes CIDRs as mentioned here.

NAS protocols overview

NAS protocols include NFS (v3 and v4.1) and SMB/CIFS (2.x and 3.x). These protocols are how NetApp Volumes allows shared access to data across multiple NAS clients. In addition, Google Cloud NetApp Volumes can provide access to NFS and SMB/CIFS clients simultaneously (dual-protocol) while honoring all of the identity and permission settings on files and folders in the NAS shares. To maintain the highest possible data transfer security, Google Cloud NetApp Volumes supports protocol encryption in flight using SMB encryption and NFS Kerberos 5p.

Dual-protocol is available with NetApp Volumes-Performance only.

Basics of NAS protocols

NAS protocols are ways for multiple clients on a network to access the same data on a storage system, such as Google Cloud NetApp Volumes on GCP. NFS and SMB are the defined NAS protocols and operate on a client/server basis where Google Cloud NetApp Volumes acts as the server. Clients send access, read, and write requests to the server, and the server is responsible for coordinating the locking mechanisms for files, storing permissions and handling identity and authentication requests.

For example, the following general process is followed if a NAS client wants to create a new file in a folder.

1. The client asks the server for information about the directory (permissions, owner, group, file ID, available space, and so on); the server responds with the information if the requesting client and user have the necessary permissions on the parent folder.

2. If the permissions on the directory allow access, the client then asks the server if the file name being created already exists in the file system. If the file name is already in use, creation fails. If the file name does not exist, the server lets the client know it can proceed.
3. The client issues a call to the server to create the file with the directory handle and file name and sets the access and modified times. The server issues a unique file ID to the file to make sure that no other files are created with the same file ID.
4. The client sends a call to check file attributes before the WRITE operation. If permissions allow it, the client then writes the new file. If locking is used by the protocol/application, the client asks the server for a lock to prevent other clients from accessing the file while locked to prevent data corruption.

NFS

NFS is a distributed file system protocol that is an open IETF standard defined in Request for Comments (RFC) that allows anyone to implement the protocol.

Volumes in Google Cloud NetApp Volumes are shared out to NFS clients by exporting a path that is accessible to a client or set of clients. Permissions to mount these exports are defined by export policies and rules, which are configurable by Google Cloud NetApp Volumes administrators.

The NetApp NFS implementation is considered a gold standard for the protocol and is used in countless enterprise NAS environments. The following sections cover NFS and specific security features available in Google Cloud NetApp Volumes and how they are implemented.

Default local UNIX users and groups

Google Cloud NetApp Volumes contains several default UNIX users and groups for various basic functionalities. These users and groups cannot currently be modified or deleted. New local users and groups cannot currently be added to Google Cloud NetApp Volumes. UNIX users and groups outside of the default users and groups need to be provided by an external LDAP name service.

The following table shows the default users and groups and their corresponding numeric IDs. NetApp recommends not creating new users or groups in LDAP or on the local clients that re-use these numeric IDs.

Default users: numeric IDs	Default groups: numeric IDs
• root:0 • pcuser:65534 • nobody:65535	• root:0 • daemon:1 • pcuser:65534 • nobody:65535

When using NFSv4.1, the root user might display as nobody when running directory listing commands on NFS clients. This is due to the client's ID domain mapping configuration. See the section called NFSv4.1 and the nobody user/group for details on this issue and how to resolve it.

The root user

In Linux, the root account has access to all commands, files, and folders in a Linux-based file system. Because of the power of this account, security best practices often require the root user to be disabled or restricted in some fashion. In NFS exports, the power a root user has over the files and folders can be controlled in Google Cloud NetApp Volumes through export policies and rules and a concept known as root squash.

Root squashing ensures that the root user accessing an NFS mount is squashed to the anonymous numeric

user 65534 (see the section "The anonymous user") and is currently only available when using NetApp Volumes-Performance by selecting Off for root access during export policy rule creation. If the root user is squashed to the anonymous user, it no longer has access to run chown or setuid/setgid commands (the sticky bit) on files or folders in the NFS mount, and files or folders created by the root user show the anon UID as the owner/group. In addition, NFSv4 ACLs cannot be modified by the root user. However, the root user still has access to chmod and deleted files that it does not have explicit permissions for. If you want to limit access to a root user's file and folder permissions, consider using a volume with NTFS ACLs, creating a Windows user named `root`, and applying the desired permissions to the files or folders.

The anonymous user

The anonymous (anon) user ID specifies a UNIX user ID or username that is mapped to client requests that arrive without valid NFS credentials. This can include the root user when root squashing is used. The anon user in Google Cloud NetApp Volumes is 65534.

This UID is normally associated with the username `nobody` or `nfsnobody` in Linux environments. Google Cloud NetApp Volumes also uses 65534 as the local UNIX user` pcuser` (see the section "Default local UNIX users and groups"), which is also the default fallback user for Windows to UNIX name mappings when no valid matching UNIX user can be found in LDAP.

Because of the differences in usernames across Linux and Google Cloud NetApp Volumes for UID 65534, the name string for users mapped to 65534 might not match when using NFSv4.1. As a result, you might see `nobody` as the user on some files and folders. See the section "NFSv4.1 and the nobody user/group" for information about this issue and how to resolve it.

Access control/exports

Initial export/share access for NFS mounts is controlled through host- based export policy rules contained within an export policy. A host IP, host name, subnet, netgroup, or domain is defined to allow access to mount the NFS share and the level of access allowed to the host. Export policy rule configuration options depend on the Google Cloud NetApp Volumes level.

For NetApp Volumes-SW, the following options are available for export-policy configuration:

- **Client match.** Comma-separated list of IP addresses, comma-separated list of hostnames, subnets, netgroups, domain names.
- **RO/RW access rules.** Select read/write or read only to control level of access to export.NetApp Volumes-Performance provides the following options:
- **Client match.** Comma-separated list of IP addresses, comma-separated list of hostnames, subnets, netgroups, domain names.
- **RO/RW access rules.** Select read/write or read only to control level of access to export.
- **Root access (on/off).** Configures root squash (see the section "The root user" for details).
- **Protocol type.** This limits access to the NFS mount to a specific protocol version. When specifying both NFSv3 and NFSv4.1 for the volume, either leave both blank or check both boxes.
- **Kerberos security level (when Enable Kerberos is selected).** Provides the options of krb5, krb5i, and/or krb5p for read-only or read-write access.

Change ownership (chown) and change group (chgrp)

NFS on Google Cloud NetApp Volumes only allows the root user to run chown/chgrp on files and folders. Other users see an `Operation not permitted` error— even on files they own. If you use root squash (as covered in the section "The root user"), the root is squashed to a nonroot user and is not allowed access to

chown and chgrp. There are currently no workarounds in Google Cloud NetApp Volumes to allow chown and chgrp for non-root users. If ownership changes are required, consider using dual protocol volumes and set the security style to NTFS to control permissions from the Windows side.

Permission management

Google Cloud NetApp Volumes supports both mode bits (such as 644, 777, and so on for rwx) and NFSv4.1 ACLs to control permissions on NFS clients for volumes that use the UNIX security style. Standard permission management is used for these (such as chmod, chown, or nfs4_setfacl) and work with any Linux client that supports them.

Additionally, when using dual protocol volumes set to NTFS, NFS clients can leverage Google Cloud NetApp Volumes name mapping to Windows users, which then are used to resolve the NTFS permissions. This requires an LDAP connection to Google Cloud NetApp Volumes to provide numeric-ID-to- username translations because Google Cloud NetApp Volumes requires a valid UNIX username to map properly to a Windows username.

Providing granular ACLs for NFSv3

Mode bit permissions cover only owner, group, and everyone else in the semantics—meaning that there are no granular user access controls in place for basic NFSv3. Google Cloud NetApp Volumes does not support POSIX ACLs, nor extended attributes (such as chattr), so granular ACLs are only possible in the following scenarios with NFSv3:

- NTFS security style volumes (CIFS server required) with valid UNIX to Windows user mappings.
- NFSv4.1 ACLs applied using an admin client mounting NFSv4.1 to apply ACLs.

Both methods require an LDAP connection for UNIX identity management and a valid UNIX user and group information populated (see the section "LDAP") and are only available with NetApp Volumes-Performance instances. To use NTFS security style volumes with NFS, you must use dual-protocol (SMB and NFSv3) or dual-protocol (SMB and NFSv4.1), even if no SMB connections are made. To use NFSv4.1 ACLs with NFSv3 mounts, you must select `Both (NFSv3/NFSv4.1)` as the protocol type.

Regular UNIX mode bits don’t provide the same level of granularity in permissions that NTFS or NFSv4.x ACLs provide. The following table compares the permission granularity between NFSv3 mode bits and NFSv4.1 ACLs. For information about NFSv4.1 ACLs, see nfs4_acl - NFSv4 Access Control Lists.

NFSv3 mode bits	NFSv4.1 ACLs
• Set user ID on execution • Set group ID on execution • Save swapped text (not defined in POSIX) • Read permission for owner • Write permission for owner • Execute permission for owner on a file; or look up (search) permission for owner in directory • Read permission for group • Write permission for group • Execute permission for group on a file; or look up (search) permission for group in directory • Read permission for others • Write permission for others • Execute permission for others on a file; or look up (search) permission for others in directory	Access control entry (ACE) types (Allow/Deny/Audit) * Inheritance flags * directory-inherit * file-inherit * no-propagate-inherit * inherit-only Permissions * read-data (files) / list-directory (directories) * write-data (files) / create-file (directories) * append-data (files) / create-subdirectory (directories) * execute (files) / change-directory (directories) * delete * delete-child * read-attributes * write-attributes * read-named-attributes * write-named-attributes * read-ACL * write-ACL * write-owner * Synchronize

Finally, NFS group membership (in both NFSv3 and NFSV4.x) is limited to a default maximum of 16 for AUTH_SYS as per the RPC packet limits. NFS Kerberos provides up to 32 groups and NFSv4 ACLs remove the limitation by way of granular user and group ACLs (up to 1024 entries per ACE).

Additionally, Google Cloud NetApp Volumes provides extended group support to extend the maximum supported groups up to 32. This requires an LDAP connection to an LDAP server that contains valid UNIX user and group identities. For more information about configuring this, see Creating and managing NFS volumes in the Google documentation.

NFSv3 user and group IDs

NFSv3 user and group IDs come across the wire as numeric IDs rather than names. Google Cloud NetApp Volumes does no username resolution for these numeric IDs with NFSv3, with UNIX security style volumes using just mode bits. When NFSv4.1 ACLs are present, a numeric ID lookup and/or name string lookup is needed to resolve the ACL properly—even when using NFSv3. With NTFS security style volumes, Google Cloud NetApp Volumes must resolve a numeric ID to a valid UNIX user and then map to a valid Windows user to negotiate access rights.

Security limitations of NFSv3 user and group IDs

With NFSv3, the client and server never have to confirm that the user attempting a read or write with a numeric ID is a valid user; it is just implicitly trusted. This opens the file system up to potential breaches simply by spoofing any numeric ID. To prevent security holes like this, there are a few options available to Google Cloud NetApp Volumes.

- Implementing Kerberos for NFS forces users to authenticate with a username and password or keytab file to get a Kerberos ticket to allow access into a mount. Kerberos is available with NetApp Volumes-Performance instances and only with NFSv4.1.
- Limiting the list of hosts in your export policy rules limits which NFSv3 clients have access to the Google

Cloud NetApp Volumes volume.

- Using dual-protocol volumes and applying NTFS ACLs to the volume forces NFSv3 clients to resolve numeric IDs to valid UNIX usernames to authenticate properly to access mounts. This requires enabling LDAP and configuring UNIX user and group identities.
- Squashing the root user limits the damage a root user can do to an NFS mount but does not completely remove risk. For more information, see the section "The root user."

Ultimately, NFS security is limited to what the protocol version you are using offers. NFSv3, while more performant in general than NFSv4.1, does not provide the same level of security.

NFSv4.1

NFSv4.1 provides greater security and reliability as compared to NFSv3, for the following reasons:

- Integrated locking through a lease-based mechanism
- Stateful sessions
- All NFS functionality over a single port (2049)
- TCP only
- ID domain mapping
- Kerberos integration (NFSv3 can use Kerberos, but only for NFS, not for ancillary protocols such as NLM)

NFSv4.1 dependencies

Because of the additionally security features in NFSv4.1, there are some external dependencies involved that were not needed to use NFSv3 (similar to how SMB requires dependencies such as Active Directory).

NFSv4.1 ACLs

Google Cloud NetApp Volumes offers support for NFSv4.x ACLs, which deliver distinct advantages over normal POSIX-style permissions, such as the following:

- Granular control of user access to files and directories
- Better NFS security
- Improved interoperability with CIFS/SMB
- Removal of the NFS limitation of 16 groups per user with AUTH_SYS security
- ACLs bypass the need for group ID (GID) resolution, which effectively removes the GID limitNFSv4.1 ACLs are controlled from NFS clients—not from Google Cloud NetApp Volumes. To use NFSv4.1 ACLs, be sure your client's software version supports them and the proper NFS utilities are installed.

Compatibility between NFSv4.1 ACLs and SMB clients

NFSv4 ACLs are different from Windows file-level ACLs (NTFS ACLs) but carry similar functionality. However, in multiprotocol NAS environments, if NFSv4.1 ACLs are present and you are using dual-protocol access (NFS and SMB on the same datasets), clients using SMB2.0 and later won't be able to view or manage ACLs from Windows security tabs.

How NFSv4.1 ACLs work

For reference, the following terms are defined:

- **Access control list (ACL).** A list of permissions entries.
- **Access control entry (ACE).** A permission entry in the list.

When a client sets an NFSv4.1 ACL on a file during a SETATTR operation, Google Cloud NetApp Volumes sets that ACL on the object, replacing any existing ACL. If there is no ACL on a file, then the mode permissions on the file are calculated from OWNER@, GROUP@, and EVERYONE@. If there are any existing SUID/SGID/STICKY bits on the file, they are not affected.

When a client gets an NFSv4.1 ACL on a file during the course of a GETATTR operation, Google Cloud NetApp Volumes reads the NFSv4.1 ACL associated with the object, constructs a list of ACEs, and returns the list to the client. If the file has an NT ACL or mode bits, then an ACL is constructed from mode bits and is returned to the client.

Access is denied if a DENY ACE is present in the ACL; access is granted if an ALLOW ACE exists. However, access is also denied if neither of the ACEs is present in the ACL.

A security descriptor consists of a security ACL (SACL) and a discretionary ACL (DACL). When NFSv4.1 interoperates with CIFS/SMB, the DACL is one-to-one mapped with NFSv4 and CIFS. The DACL consists of the ALLOW and the DENY ACEs.

If a basic `chmod` is run on a file or folder with NFSv4.1 ACLs set, existing user and group ACLs are preserved, but the default OWNER@, GROUP@, EVERYONE@ ACLs are modified.

A client using NFSv4.1 ACLs can set and view ACLs for files and directories on the system. When a new file or subdirectory is created in a directory that has an ACL, that object inherits all ACEs in the ACL that have been tagged with the appropriate inheritance flags.

If a file or directory has an NFSv4.1 ACL, that ACL is used to control access no matter which protocol is used to access the file or directory.

Files and directories inherit ACEs from NFSv4 ACLs on parent directories (possibly with appropriate modifications) as long as the ACEs have been tagged with the correct inheritance flags.

When a file or directory is created as the result of an NFSv4 request, the ACL on the resulting file or directory depends on whether the file creation request includes an ACL or only standard UNIX file access permissions. The ACL also depends on whether the parent directory has an ACL.

- If the request includes an ACL, that ACL is used.
- If the request includes only standard UNIX file access permissions and the parent directory does not have an ACL, the client file mode is used to set standard UNIX file access permissions.
- If the request includes only standard UNIX file access permissions and the parent directory has a noninheritable ACL, a default ACL based on the mode bits passed into the request is set on the new object.
- If the request includes only standard UNIX file access permissions but the parent directory has an ACL, the ACEs in the parent directory's ACL are inherited by the new file or directory as long as the ACEs have been tagged with the appropriate inheritance flags.

ACE permissions

NFSv4.1 ACLs permissions uses a series of upper- and lower-case letter values (such as `rxtncy`) to control access. For more information about these letter values, see HOW TO: Use NFSv4 ACL.

NFSv4.1 ACL behavior with umask and ACL inheritance

NFSv4 ACLs provide the ability to offer ACL inheritance. ACL inheritance means that files or folders created beneath objects with NFSv4.1 ACLs set can inherit the ACLs based on the configuration of the ACL inheritance flag.

Umask is used to control the permission level at which files and folders are created in a directory without administrator interaction. By default, Google Cloud NetApp Volumes allows umask to override inherited ACLs, which is expected behavior as per RFC 5661.

ACL formatting

NFSv4.1 ACLs have specific formatting. The following example is an ACE set on a file:

```
A::ldapuser@domain.netapp.com:rwatTnNcCy
```

The preceding example follows the ACL format guidelines of:

```
type:flags:principal:permissions
```

A type of `A` means "allow." The inherit flags are not set in this case, because the principal is not a group and does not include inheritance. Also, because the ACE is not an AUDIT entry, there is no need to set the audit flags. For more information about NFSv4.1 ACLs, see http://linux.die.net/man/5/nfs4_acl.

If the NFSv4.1 ACL is not set properly (or a name string cannot be resolved by the client and server), the ACL might not behave as expected, or the ACL change might fail to apply and throw an error.

Sample errors include:

```
Failed setxattr operation: Invalid argument
Scanning ACE string 'A:: user@rwaDxtTnNcCy' failed.
```

Explicit DENY

NFSv4.1 permissions can include explicit DENY attributes for OWNER, GROUP, and EVERYONE. That is because NFSv4.1 ACLs are default-deny, which means that if an ACL is not explicitly granted by an ACE, then it is denied. Explicit DENY attributes override any ACCESS ACEs, explicit or not.

DENY ACEs are set with an attribute tag of `D`.

In the example below, GROUP@ is allowed all read and execute permissions, but denied all write access.

```
sh-4.1$ nfs4_getfacl /mixed
A::ldapuser@domain.netapp.com:ratTnNcCy
A::OWNER@:rwaDxtTnNcCy
D::OWNER@:
A:g:GROUP@:rxtncy
```

```
D:g:GROUP@:waDTC
A::EVERYONE@:rxtncy
D::EVERYONE@:waDTC
```

DENY ACEs should be avoided whenever possible because they can be confusing and complicated; ALLOW ACLs that are not explicitly defined are implicitly denied. When DENY ACEs are set, users might be denied access when they expect to be granted access.

The preceding set of ACEs is equivalent to 755 in mode bits, which means:

- The owner has full rights.
- Groups have read only.
- Others have read only.

However, even if permissions are adjusted to the 775 equivalent, access can be denied because of the explicit DENY set on EVERYONE.

NFSv4.1 ID domain mapping dependencies

NFSv4.1 leverages ID domain mapping logic as a security layer to help verify that a user attempting access to an NFSv4.1 mount is indeed who they claim to be. In these cases, the username and group name coming from the NFSv4.1 client appends a name string and sends it to the Google Cloud NetApp Volumes instance. If that username/group name and ID string combination does not match, then the user and/or group is squashed to the default nobody user specified in the `/etc/idmapd.conf` file on the client.

This ID string is a requirement for proper permission adherence, especially when NFSv4.1 ACLs and/or Kerberos are in use. As a result, name service server dependencies such as LDAP servers are necessary to ensure consistency across clients and Google Cloud NetApp Volumes for proper user and group name identity resolution.

Google Cloud NetApp Volumes uses a static default ID domain name value of `defaultv4iddomain.com`. NFS clients default to the DNS domain name for its ID domain name settings, but you can manually adjust the ID domain name in `/etc/idmapd.conf`.

If LDAP is enabled in Google Cloud NetApp Volumes, then Google Cloud NetApp Volumes automates the NFS ID domain to change to what is configured for the search domain in DNS and clients won't need to be modified unless they use different DNS domain search names.

When Google Cloud NetApp Volumes can resolve a username or group name in local files or LDAP, the domain string is used and non-matching domain IDs squash to nobody. If Google Cloud NetApp Volumes cannot find a username or group name in local files or LDAP, the numeric ID value is used and the NFS client resolves the name properly (this is similar to NFSv3 behavior).

Without changing the client's NFSv4.1 ID domain to match what the Google Cloud NetApp Volumes volume is using, you see the following behavior:

- UNIX users and groups with local entries in Google Cloud NetApp Volumes (such as root, as defined in local UNIX users and groups) are squashed to the nobody value.
- UNIX users and groups with entries in LDAP (if Google Cloud NetApp Volumes is configured to use LDAP) squashes to nobody if DNS domains are different between NFS clients and Google Cloud NetApp Volumes.

- UNIX users and groups with no local entries or LDAP entries use the numeric ID value and resolve to the name specified on the NFS client. If no name exists on the client, only the numeric ID is shown.

The following shows the results of the preceding scenario:

```
# ls -la /mnt/home/prof1/nfs4/
total 8
drwxr-xr-x 2 nobody nobody 4096 Feb  3 12:07 .
drwxrwxrwx 7 root   root   4096 Feb  3 12:06 ..
-rw-r--r-- 1   9835   9835    0 Feb  3 12:07 client-user-no-name
-rw-r--r-- 1 nobody nobody    0 Feb  3 12:07 ldap-user-file
-rw-r--r-- 1 nobody nobody    0 Feb  3 12:06 root-user-file
```

When the client and server ID domains match, this is how the same file listing looks:

```
# ls -la
total 8
drwxr-xr-x 2 root   root         4096 Feb  3 12:07 .
drwxrwxrwx 7 root   root         4096 Feb  3 12:06 ..
-rw-r--r-- 1   9835         9835    0 Feb  3 12:07 client-user-no-name
-rw-r--r-- 1 apache apache-group    0 Feb  3 12:07 ldap-user-file
-rw-r--r-- 1 root   root            0 Feb  3 12:06 root-user-file
```

For more information about this issue and how to resolve it, see the section "NFSv4.1 and the nobody user/group."

Kerberos dependencies

If you plan to use Kerberos with NFS, you must have the following with Google Cloud NetApp Volumes:

- Active Directory domain for Kerberos Distribution Center services (KDC)
- Active Directory domain with user and group attributes populated with UNIX information for LDAP functionality (NFS Kerberos in Google Cloud NetApp Volumes requires a user SPN to UNIX user mapping for proper functionality.)
- LDAP enabled on the Google Cloud NetApp Volumes instance
- Active Directory domain for DNS services

NFSv4.1 and the nobody user/group

One of the most common issues seen with an NFSv4.1 configuration is when a file or folder is shown in a listing using `ls` as being owned by the `user:group` combination of `nobody:nobody`.

For example:

```
sh-4.2$ ls -la | grep prof1-file
```

```
-rw-r--r-- 1 nobody nobody    0 Apr 24 13:25 prof1-file
```

And the numeric ID is `99`.

```
sh-4.2$ ls -lan | grep prof1-file
-rw-r--r-- 1 99 99    0 Apr 24 13:25 prof1-file
```

In some instances, the file might show the correct owner but `nobody` as the group.

```
sh-4.2$ ls -la | grep newfile1
-rw-r--r-- 1 prof1  nobody    0 Oct  9  2019 newfile1
```

Who is nobody?

The `nobody` user in NFSv4.1 is different from the `nfsnobody` user. You can view how an NFS client sees each user by running the `id` command:

```
# id nobody
uid=99(nobody) gid=99(nobody) groups=99(nobody)
# id nfsnobody
uid=65534(nfsnobody) gid=65534(nfsnobody) groups=65534(nfsnobody)
```

With NFSv4.1, the `nobody` user is the default user defined by the `idmapd.conf` file and can be defined as any user you want to use.

```
# cat /etc/idmapd.conf | grep nobody
#Nobody-User = nobody
#Nobody-Group = nobody
```

Why does this happen?

Because security through name string mapping is a key tenet of NFSv4.1 operations, the default behavior when a name string does not match properly is to squash that user to one that won't normally have any access to files and folders owned by users and groups.

When you see `nobody` for the user and/or group in file listings, this generally means something in NFSv4.1 is misconfigured. Case sensitivity can come into play here.

For example, if user1@CVSDEMO.LOCAL (uid 1234, gid 1234) is accessing an export, then Google Cloud NetApp Volumes must be able to find user1@CVSDEMO.LOCAL (uid 1234, gid 1234). If the user in Google Cloud NetApp Volumes is USER1@CVSDEMO.LOCAL, then it won't match (uppercase USER1 versus lowercase user1). In many cases, you can see the following in the messages file on the client:

```
May 19 13:14:29 centos7 nfsidmap[17481]: nss_getpwnam: name
```

```
'root@defaultv4iddomain.com' does not map into domain 'CVSDEMO.LOCAL'
May 19 13:15:05 centos7 nfsidmap[17534]: nss_getpwnam: name 'nobody' does
not map into domain 'CVSDEMO.LOCAL'
```

The client and server must both agree that a user is indeed who they are claiming to be, so you must check the following to ensure that the user that the client sees has the same information as the user that Google Cloud NetApp Volumes sees.

- **NFSv4.x ID domain.** Client: `idmapd.conf` file; Google Cloud NetApp Volumes uses `defaultv4iddomain.com` and cannot be changed manually. If using LDAP with NFSv4.1, Google Cloud NetApp Volumes changes the ID domain to what the DNS search domain is using, which is the same as the AD domain.
- **User name and numeric IDs.** This determines where the client is looking for user names and leverages the name service switch configuration—client: `nsswitch.conf` and/or local passwd and group files; Google Cloud NetApp Volumes does not allow modifications to this but automatically adds LDAP to the configuration when it is enabled.
- **Group name and numeric IDs.** This determines where the client is looking for group names and leverages the name service switch configuration—client: `nsswitch.conf` and/or local passwd and group files; Google Cloud NetApp Volumes does not allow modifications to this but automatically adds LDAP to the configuration when it is enabled.

In almost all cases, if you see `nobody` in user and group listings from clients, the issue is user or group name domain ID translation between Google Cloud NetApp Volumes and the NFS client. To avoid this scenario, use LDAP to resolve user and group information between clients and Google Cloud NetApp Volumes.

Viewing name ID strings for NFSv4.1 on clients

If you are using NFSv4.1, there is a name-string mapping that takes place during NFS operations, as previously described.

In addition to using `/var/log/messages` to find an issue with NFSv4 IDs, you can use the nfsidmap -l command on the NFS client to view which usernames have properly mapped to the NFSv4 domain.

For example, this is output of the command after a user that can be found by the client and Google Cloud NetApp Volumes accesses an NFSv4.x mount:

```
# nfsidmap -l
4 .id_resolver keys found:
  gid:daemon@CVSDEMO.LOCAL
  uid:nfs4@CVSDEMO.LOCAL
  gid:root@CVSDEMO.LOCAL
  uid:root@CVSDEMO.LOCAL
```

When a user that does not map properly into the NFSv4.1 ID domain (in this case, `netapp-user`) tries to access the same mount and touches a file, they are assigned `nobody:nobody`, as expected.

```
# su netapp-user
sh-4.2$ id
```

```
uid=482600012(netapp-user), 2000(secondary)
sh-4.2$ cd /mnt/nfs4/
sh-4.2$ touch newfile
sh-4.2$ ls -la
total 16
drwxrwxrwx  5 root   root   4096 Jan 14 17:13 .
drwxr-xr-x. 8 root   root     81 Jan 14 10:02 ..
-rw-r--r--  1 nobody nobody    0 Jan 14 17:13 newfile
drwxrwxrwx  2 root   root   4096 Jan 13 13:20 qtree1
drwxrwxrwx  2 root   root   4096 Jan 13 13:13 qtree2
drwxr-xr-x  2 nfs4   daemon 4096 Jan 11 14:30 testdir
```

The `nfsidmap -l` output shows the user `pcuser` in the display but not `netapp-user`; this is the anonymous user in our export-policy rule (`65534`).

```
# nfsidmap -l
6 .id_resolver keys found:
  gid:pcuser@CVSDEMO.LOCAL
  uid:pcuser@CVSDEMO.LOCAL
  gid:daemon@CVSDEMO.LOCAL
  uid:nfs4@CVSDEMO.LOCAL
  gid:root@CVSDEMO.LOCAL
  uid:root@CVSDEMO.LOCAL
```

SMB

SMB is a network file sharing protocol developed by Microsoft that provides centralized user/group authentication, permissions, locking, and file sharing to multiple SMB clients over an Ethernet network. Files and folders are presented to clients by way of shares, which can be configured with a variety of share properties and offers access control through share-level permissions. SMB can be presented to any client that offers support for the protocol, including Windows, Apple, and Linux clients.

Google Cloud NetApp Volumes provides support for the SMB 2.1 and 3.x versions of the protocol.

Access control/SMB shares

- When a Windows username requests access to the Google Cloud NetApp Volumes volume, Google Cloud NetApp Volumes looks for a UNIX username using the methods configured by Google Cloud NetApp Volumes administrators.
- If an external UNIX identity provider (LDAP) is configured and Windows/UNIX usernames are identical, then Windows usernames will map 1:1 to UNIX usernames without any additional configuration needed. When LDAP is enabled, Active Directory is used to host those UNIX attributes for user and group objects.
- If Windows names and UNIX names do not match identically, then LDAP must be configured to allow Google Cloud NetApp Volumes to use the LDAP name mapping configuration (see the section "Using LDAP for asymmetric name mapping").

- If LDAP is not in use, then Windows SMB users map to a default local UNIX user named `pcuser` in Google Cloud NetApp Volumes. This means files written in Windows by users that map to the `pcuser` show UNIX ownership as `pcuser` in multiprotocol NAS environments. `pcuser` here is effectively the `nobody` user in Linux environments (UID 65534).

In deployments with SMB only, the `pcuser` mapping still occurs, but it won't matter, because Windows user and group ownership is correctly displayed and NFS access to the SMB-only volume is not allowed. In addition, SMB-only volumes do not support conversion to NFS or dual-protocol volumes after they are created.

Windows leverages Kerberos for username authentication with the Active Directory domain controllers, which requires a username/password exchange with the AD DCs, which is external to the Google Cloud NetApp Volumes instance. Kerberos authentication is used when the `\\SERVERNAME` UNC path is used by the SMB clients and the following is true:

- DNS A/AAAA entry exists for SERVERNAME
- A valid SPN for SMB/CIFS access exists for SERVERNAME

When a Google Cloud NetApp Volumes SMB volume is created, the machine account name is created as defined in the section "How Google Cloud NetApp Volumes shows up in Active Directory." That machine account name also becomes the SMB share access path because Google Cloud NetApp Volumes leverages Dynamic DNS (DDNS) to create the necessary A/AAAA and PTR entries in DNS and the necessary SPN entries on the machine account principal.

> For PTR entries to be created, the reverse lookup zone for the Google Cloud NetApp Volumes instance IP address must exist on the DNS server.

For example, this Google Cloud NetApp Volumes volume uses the following UNC share path: `\\cvs-east-433d.cvsdemo.local`.

In Active Directory, these are the Google Cloud NetApp Volumes-generated SPN entries:

```
PS C:\> setspn /L CVS-EAST-433D
Registered ServicePrincipalNames for CN=CVS-EAST-433D,CN=Computers,DC=cvsdemo,DC=local:
        HOST/cvs-east-433d.cvsdemo.local
        HOST/CVS-EAST-433D
```

This is the DNS forward/reverse lookup result:

```
PS C:\> nslookup NetApp Volumes-EAST-433D
Server:  activedirectory. region. lab. internal
Address:  10. xx.0. xx
Name:    NetApp Volumes-EAST-433D.cvsdemo.local
Address:  10. xxx.0. x
PS C:\> nslookup 10. xxx.0. x
Server:  activedirectory.region.lab.internal
Address:  10.xx.0.xx
Name:    NetApp Volumes-EAST-433D.CVSDEMO.LOCAL
Address:  10. xxx.0. x
```

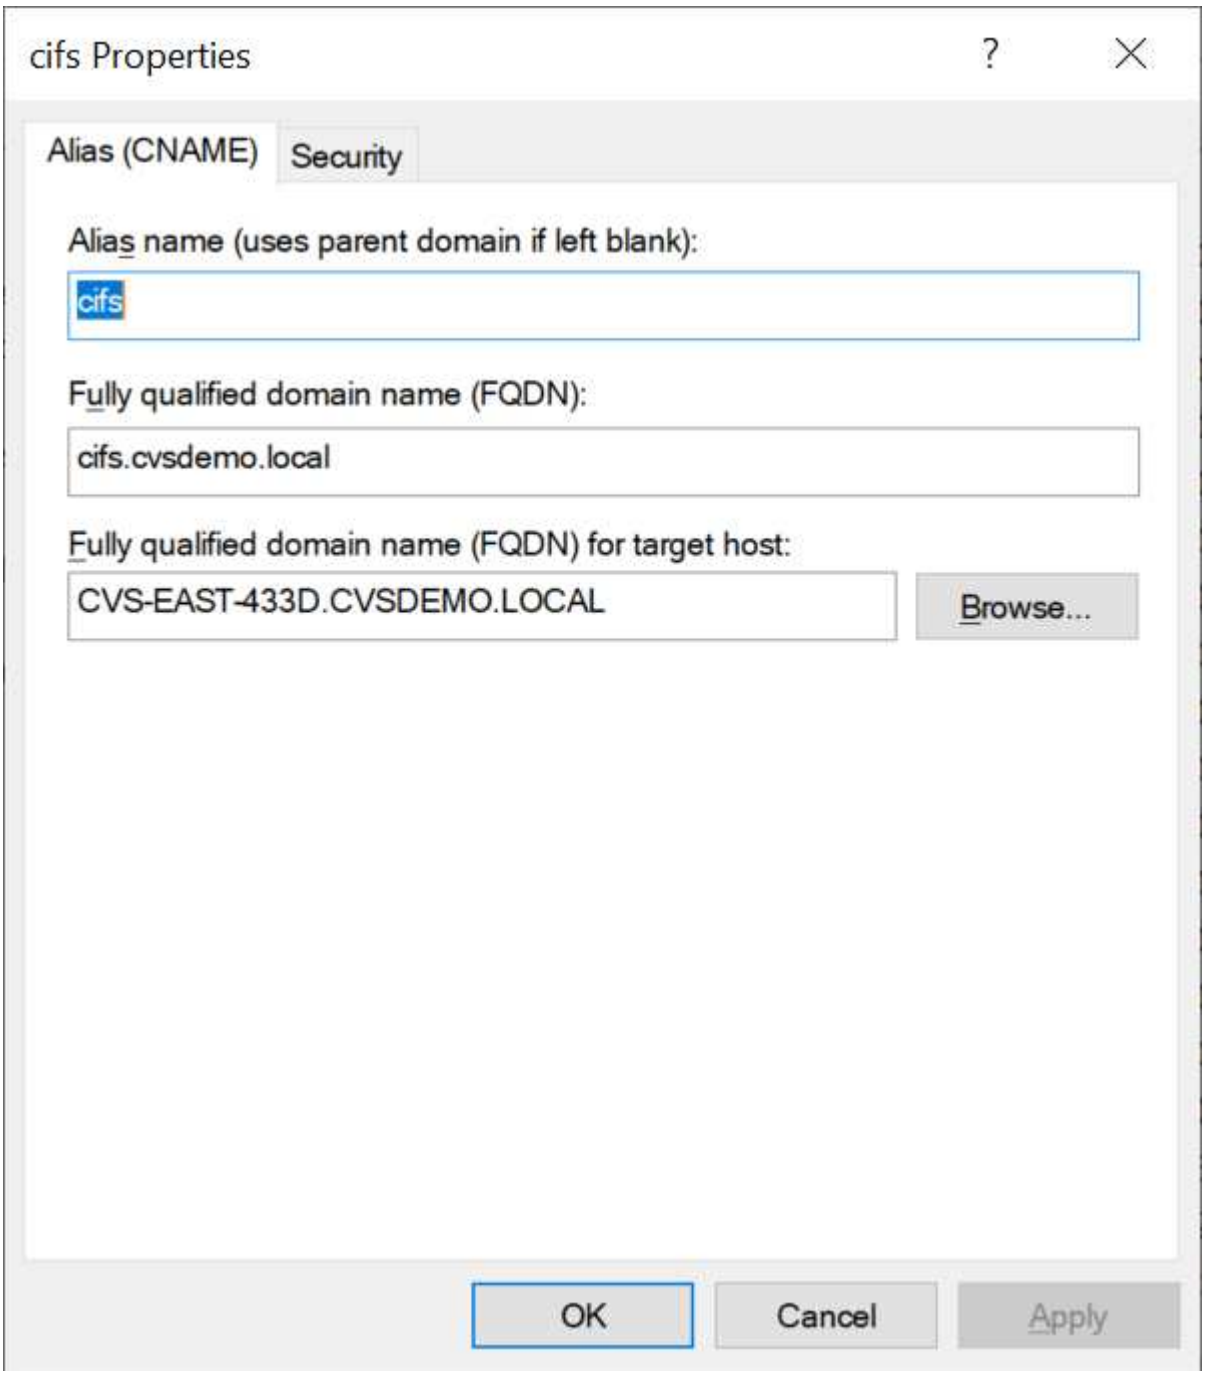
Optionally, more access control can be applied by enabling/requiring SMB encryption for SMB shares in Google Cloud NetApp Volumes. If SMB encryption isn't supported by one of the endpoints, then access is not

allowed.

Using SMB name aliases

In some cases, it might be a security concern for end users to know the machine account name in use for Google Cloud NetApp Volumes. In other cases, you might simply want to provide a simpler access path to your end users. In those cases, you can create SMB aliases.

If you want to create aliases for the SMB share path, you can leverage what is known as a CNAME record in DNS. For example, if you want to use the name `\\CIFS` to access shares instead of `\\cvs-east-433d.cvsdemo.local`, but you still want to use Kerberos authentication, a CNAME in DNS that points to the existing A/AAAA record and an additional SPN added to the existing machine account provides Kerberos access.



This is the resulting DNS forward lookup result after adding a CNAME:

```
PS C:\> nslookup cifs
Server:  ok-activedirectory.us-east4-a.c.cv-solution-architect-
lab.internal
Address:  10. xx.0. xx
Name:    NetApp Volumes-EAST-433D.cvsdemo.local
Address:  10. xxx.0. x
Aliases:  cifs.cvsdemo.local
```

This is the resulting SPN query after adding new SPNs:

```
PS C:\> setspn /L CVS-EAST-433D
Registered ServicePrincipalNames for CN=CVS-EAST-433D,CN=Computers,DC=cvsdemo,DC=local:
        cifs/cifs.cvsdemo.local
        cifs/cifs
        HOST/cvs-east-433d.cvsdemo.local
        HOST/CVS-EAST-433D
```

In a packet capture, we can see the Session Setup Request using the SPN tied to the CNAME.

```
431 4.156722                                    SMB2     308 Negotiate Protocol Response
432 4.156785                                    SMB2     232 Negotiate Protocol Request
434 4.158108                                    SMB2     374 Negotiate Protocol Response
435 4.160977                                    SMB2    1978 Session Setup Request
437 4.166224                                    SMB2     322 Session Setup Response
438 4.166891                                    SMB2     152 Tree Connect Request Tree: \\cifs\IPC$
439 4.168063                                    SMB2     138 Tree Connect Response
```

```
realm: CVSDEMO.LOCAL
sname
   name-type: kRB5-NT-SRV-INST (2)
   sname-string: 2 items
      SNameString: cifs
      SNameString: cifs
enc-part
   etype: eTYPE-ARCFOUR-HMAC-MD5 (23)
```

SMB authentication dialects

Google Cloud NetApp Volumes supports the following dialects for SMB authentication:

- LM
- NTLM
- NTLMv2
- Kerberos

Kerberos authentication for SMB share access is the most secure level of authentication you can use. With AES and SMB encryption enabled, the security level is further increased.

Google Cloud NetApp Volumes also supports backward compatibility for LM and NTLM authentication. When Kerberos is misconfigured (such as when creating SMB aliases), share access falls back to weaker authentication methods (such as NTLMv2). Because these mechanisms are less secure, they are disabled in some Active Directory environments. If weaker authentication methods are disabled and Kerberos is not

configured properly, share access fails because there is no valid authentication method to fall back to.

For information about configuring/viewing your supported authentication levels in Active Directory, see Network security: LAN Manager authentication level.

Permission models

NTFS/File permissions

NTFS permissions are the permissions applied to files and folders in file systems adhering to NTFS logic. You can apply NTFS permissions in `Basic` or `Advanced` and can be set to `Allow` or `Deny` for access control.

Basic permissions include the following:

- Full Control
- Modify
- Read & Execute
- Read
- Write

When you set permissions for a user or group, referred to as an ACE, it resides in an ACL. NTFS permissions use the same read/write/execute basics as UNIX mode bits, but they can also extend to more granular and extended access controls (also known as Special Permissions), such as Take Ownership, Create Folders/Append Data, Write Attributes, and more.

Standard UNIX mode bits do not provide the same level of granularity as NTFS permissions (such as being able to set permissions for individual user and group objects in an ACL or setting extended attributes). However, NFSv4.1 ACLs do provide the same functionality as NTFS ACLs.

NTFS permissions are more specific than share permissions and can be used in conjunction with share permissions. With NTFS permission structures, the most restrictive applies. As such, explicit denials to a user or group overrides even Full Control when defining access rights.

NTFS permissions are controlled from Windows SMB clients.

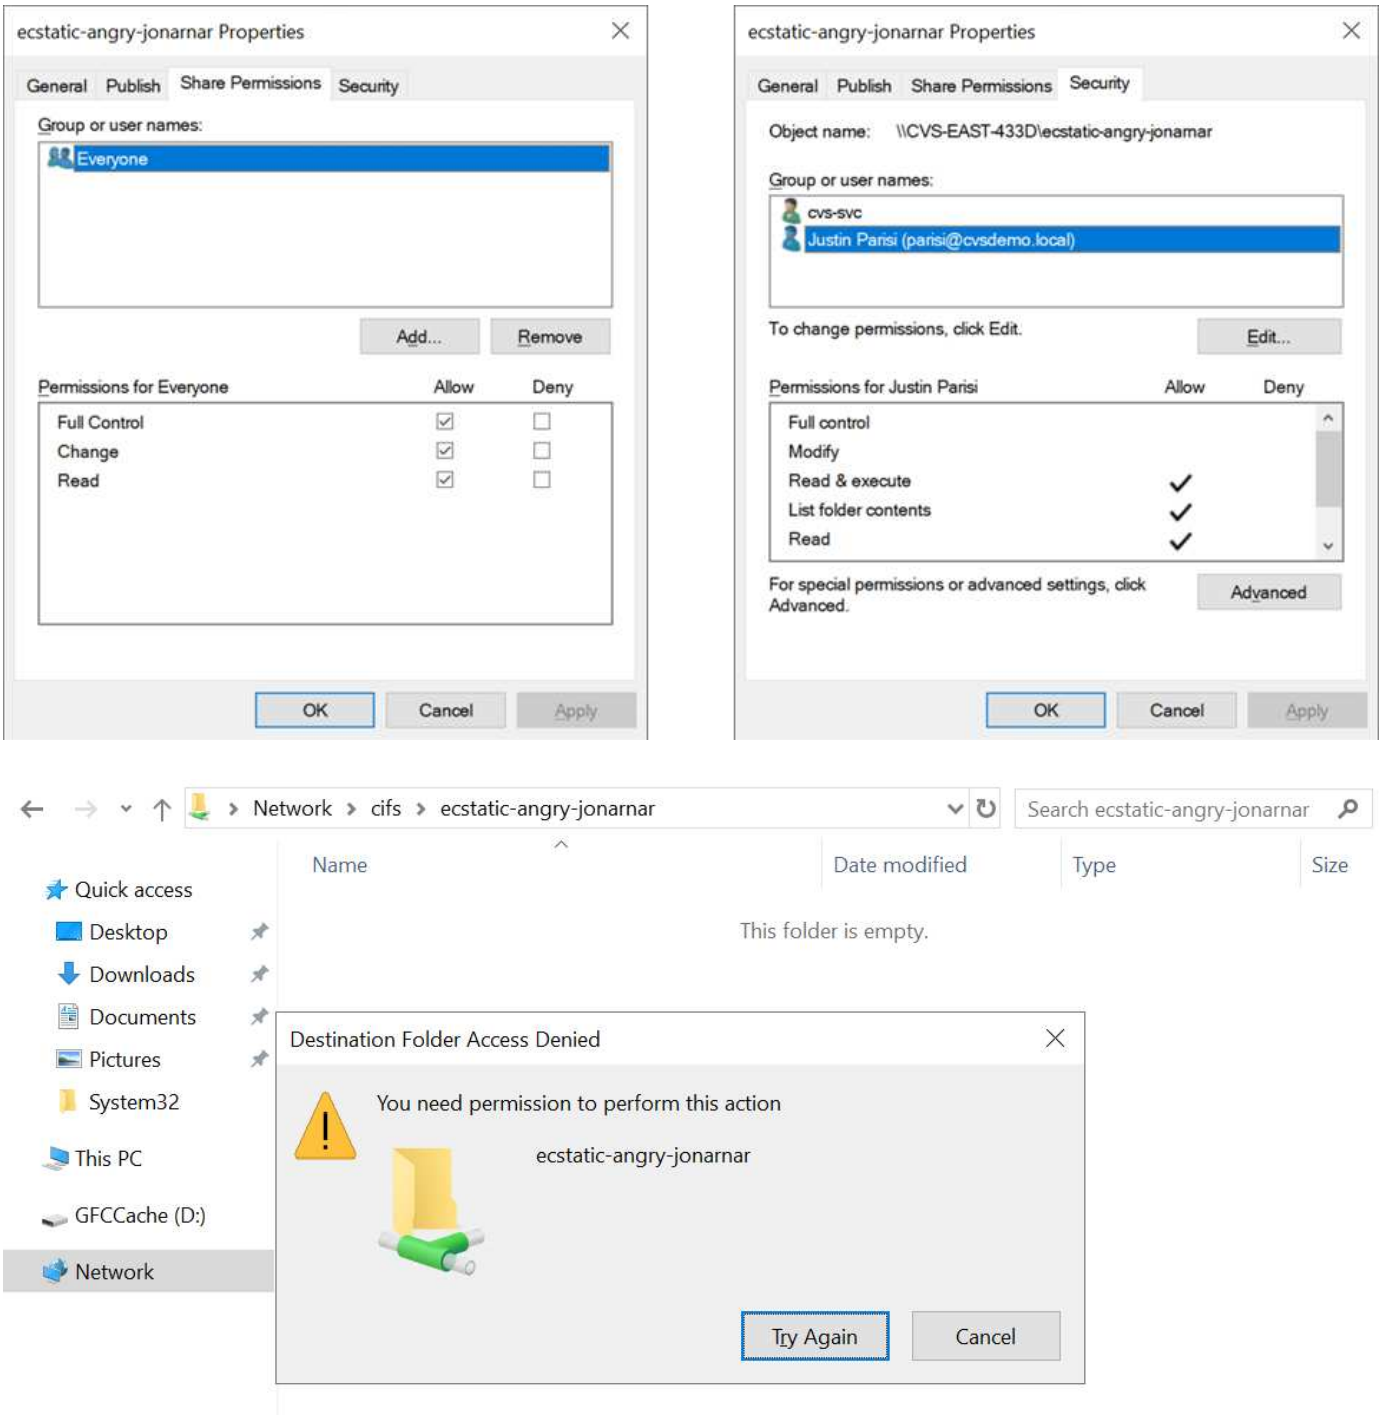
Share permissions

Share permissions are more general than NTFS permissions (Read/Change/Full Control only) and control the initial entry into an SMB share—similar to how NFS export policy rules work.

Although NFS export policy rules control access through host-based information such as IP addresses or host names, SMB share permissions can control access by using user and group ACEs in a share ACL. You can set share ACLs either from the Windows client or from the Google Cloud NetApp Volumes management UI.

By default, share ACLs and initial volume ACLs include Everyone with Full Control. The file ACLs should be changed but share permissions are overruled by the file permissions on objects in the share.

For instance, if a user is only allowed Read access to the Google Cloud NetApp Volumes volume file ACL, they are denied access to create files and folders even though the share ACL is set to Everyone with Full Control, as shown in the following figure.



For best security results, do the following:

- Remove Everyone from the share and file ACLs and instead set share access for users or groups.
- Use groups for access control instead of individual users for ease of management and faster removal/addition of users to share ACLs through group management.
- Allow less restrictive, more general share access to the ACEs on the share permissions and lock down access to users and groups with file permissions for more granular access control.
- Avoid general use of explicit deny ACLs, because they override allow ACLs. Limit use of explicit deny ACLs for users or groups that need to be restricted from access to a file system quickly.
- Make sure that you pay attention to the ACL inheritance settings when modifying permissions; setting the

inheritance flag at the top level of a directory or volume with high file counts means that each file below that directory or volume has inherited permissions added to it, which can create unwanted behavior such as unintended access/denial and long churn of permission modification as each file is adjusted.

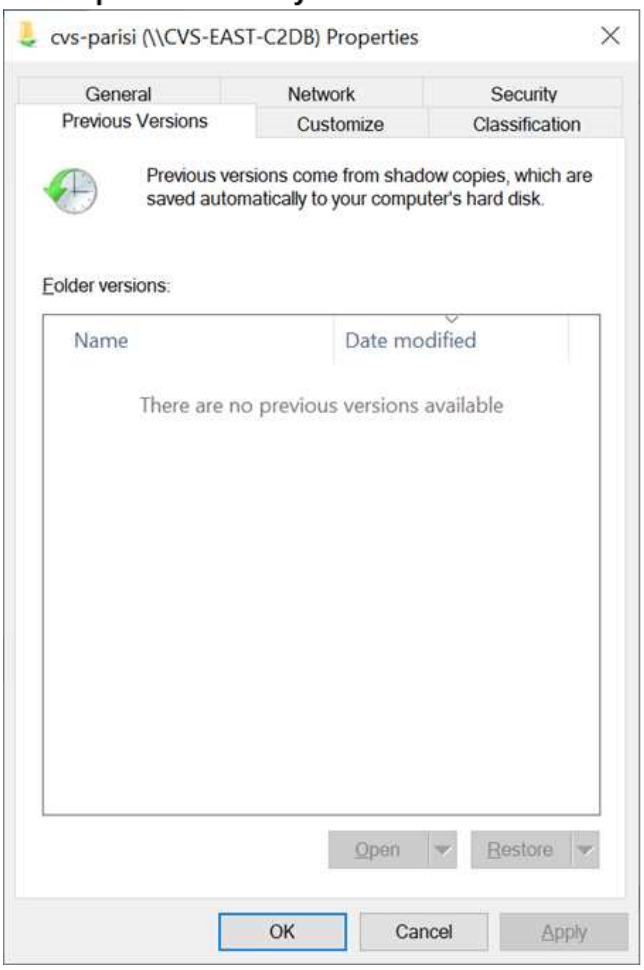
SMB share security features

When you first create a volume with SMB access in Google Cloud NetApp Volumes, you are presented with a series of choices for securing that volume.

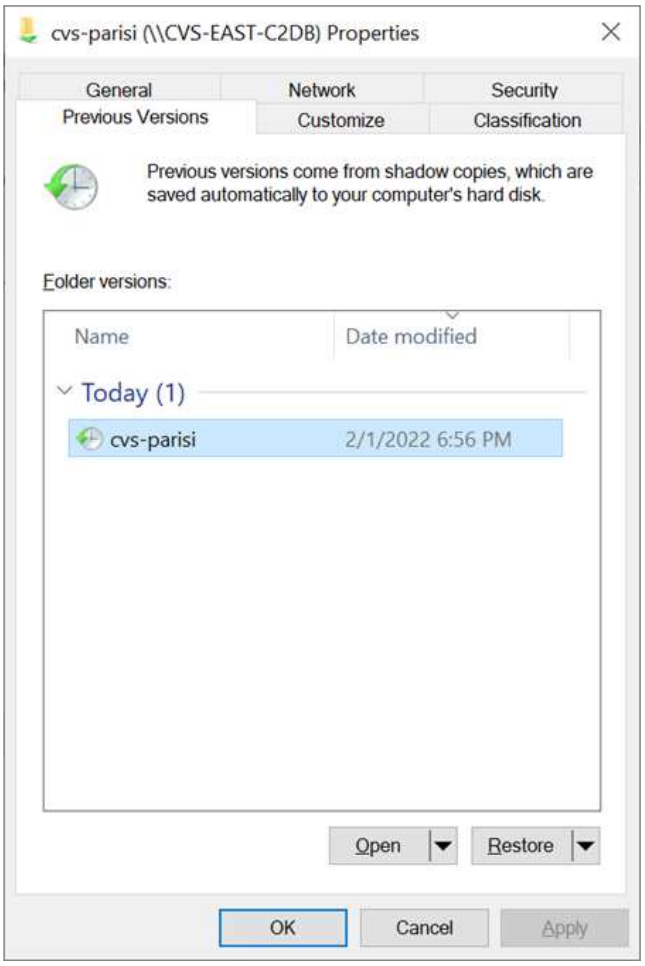
Some of these choices depend on the Google Cloud NetApp Volumes level (Performance or Software) and choices include:

- **Make snapshot directory visible (available for both NetApp Volumes-Performance and NetApp Volumes-SW).** This option controls whether or not SMB clients can access the Snapshot directory in an SMB share (`\\server\share\~snapshot` and/or Previous Versions tab). The default setting is Not Checked, which means that the volume defaults to hiding and disallowing access to the `~snapshot` directory, and no Snapshot copies appear in the Previous Versions tab for the volume.

Snapshot directory visible: Unchecked



Snapshot directory visible: Checked



Hiding Snapshot copies from end users might be desired for security reasons, performance reasons (hiding these folders from AV scans) or preference. Google Cloud NetApp Volumes Snapshots are read- only, so even if these Snapshots are visible, end users cannot delete or modify files in the Snapshot directory. File permissions on the files or folders at the time the Snapshot copy was taken apply. If a file or folder's permissions change between Snapshot copies, then the changes also apply to the files or folders in the Snapshot directory. Users and groups can gain access to these files or folders based on permissions. While

deletes or modifications of files in the Snapshot directory are not possible, it is possible to copy files or folders out of the Snapshot directory.

- **Enable SMB encryption (available for both NetApp Volumes-Performance and NetApp Volumes-SW).** SMB encryption is disabled on the SMB share by default (unchecked). Checking the box enables SMB encryption, which means traffic between the SMB client and server is encrypted in-flight with the highest supported encryption levels negotiated. Google Cloud NetApp Volumes supports up to AES-256 encryption for SMB. Enabling SMB encryption does carry a performance penalty that might or might not be noticeable to your SMB clients—roughly in the 10-20% range. NetApp strongly encourages testing to see if that performance penalty is acceptable.
- **Hide SMB share (available for both NetApp Volumes-Performance and NetApp Volumes-SW).** Setting this option hides the SMB share path from normal browsing. This means that clients that do not know the share path cannot see the shares when accessing the default UNC path (such as `\\NetApp Volumes-SMB`). When the checkbox is selected, only clients that explicitly know the SMB share path or have the share path defined by a Group Policy Object can access it (security through obfuscation).
- **Enable access-based enumeration (ABE) (NetApp Volumes-SW only).** This is similar to hiding the SMB share, except the shares or files are only hidden from users or groups that do not have permissions to access the objects. For instance, if Windows user `joe` is not allowed at least Read access through the permissions, then the Windows user `joe` cannot see the SMB share or files at all. This is disabled by default, and you can enable it by selecting the checkbox. For more information on ABE, see the NetApp Knowledge Base article How does Access Based Enumeration (ABE) work?
- **Enable Continuously Available (CA) share support (NetApp Volumes-Performance only).** Continuously Available SMB shares provide a way to minimize application disruptions during failover events by replicating lock states across nodes in the Google Cloud NetApp Volumes backend system. This is not a security feature, but it does offer better overall resiliency. Currently, only SQL Server and FSLogix applications are supported for this functionality.

Default hidden shares

When an SMB server is created in Google Cloud NetApp Volumes, there are hidden administrative shares (using the $ naming convention) that are created in addition to the data volume SMB share. These include C$ (namespace access) and IPC$ (sharing named pipes for communication between programs, such as the remote procedure calls (RPC) used for Microsoft Management Console (MMC) access).

The IPC$ share contains no share ACLs and cannot be modified—it is strictly used for RPC calls and Windows disallows anonymous access to these shares by default.

The C$ share allows BUILTIN/Administrators access by default, but Google Cloud NetApp Volumes automation removes the share ACL and does not allow access to anyone because access to the C$ share allows visibility into all mounted volumes in the Google Cloud NetApp Volumes file systems. As a result, attempts to navigate to `\\SERVER\C$` fail.

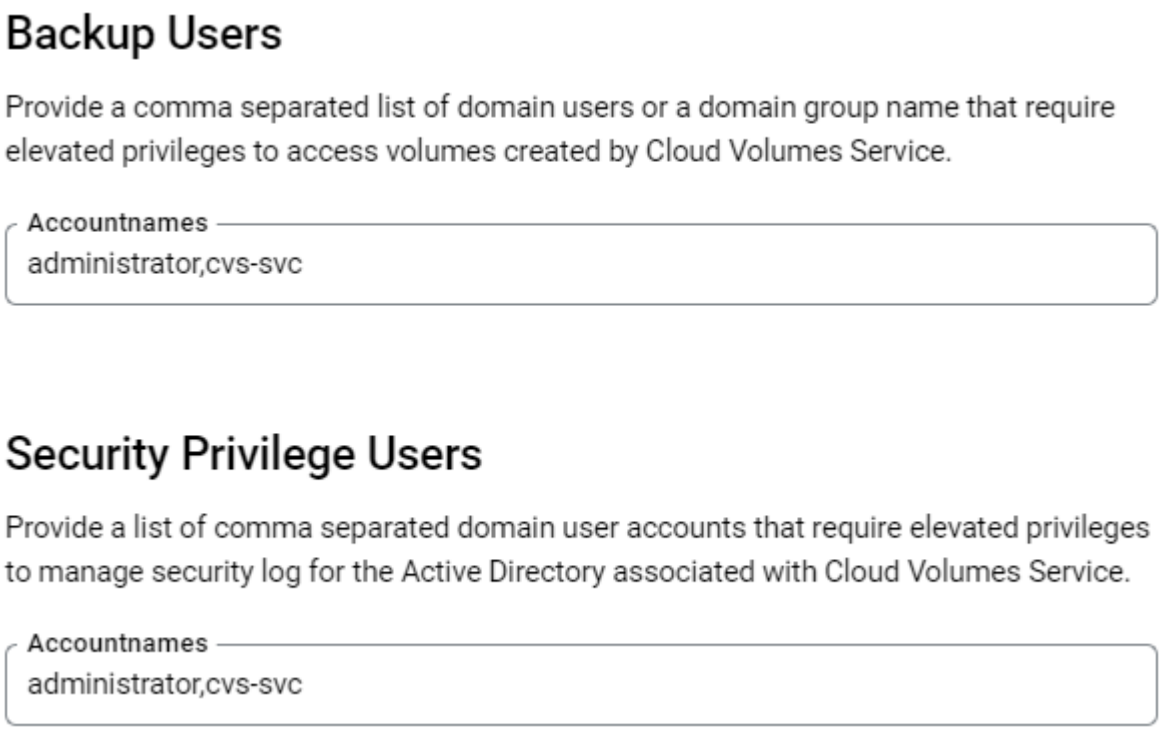
Accounts with local/BUILTIN administrator/backup rights

Google Cloud NetApp Volumes SMB servers maintain similar functionality to regular Windows SMB servers in that there are local groups (such as BUILTIN\Administrators) that apply access rights to select domain users and groups.

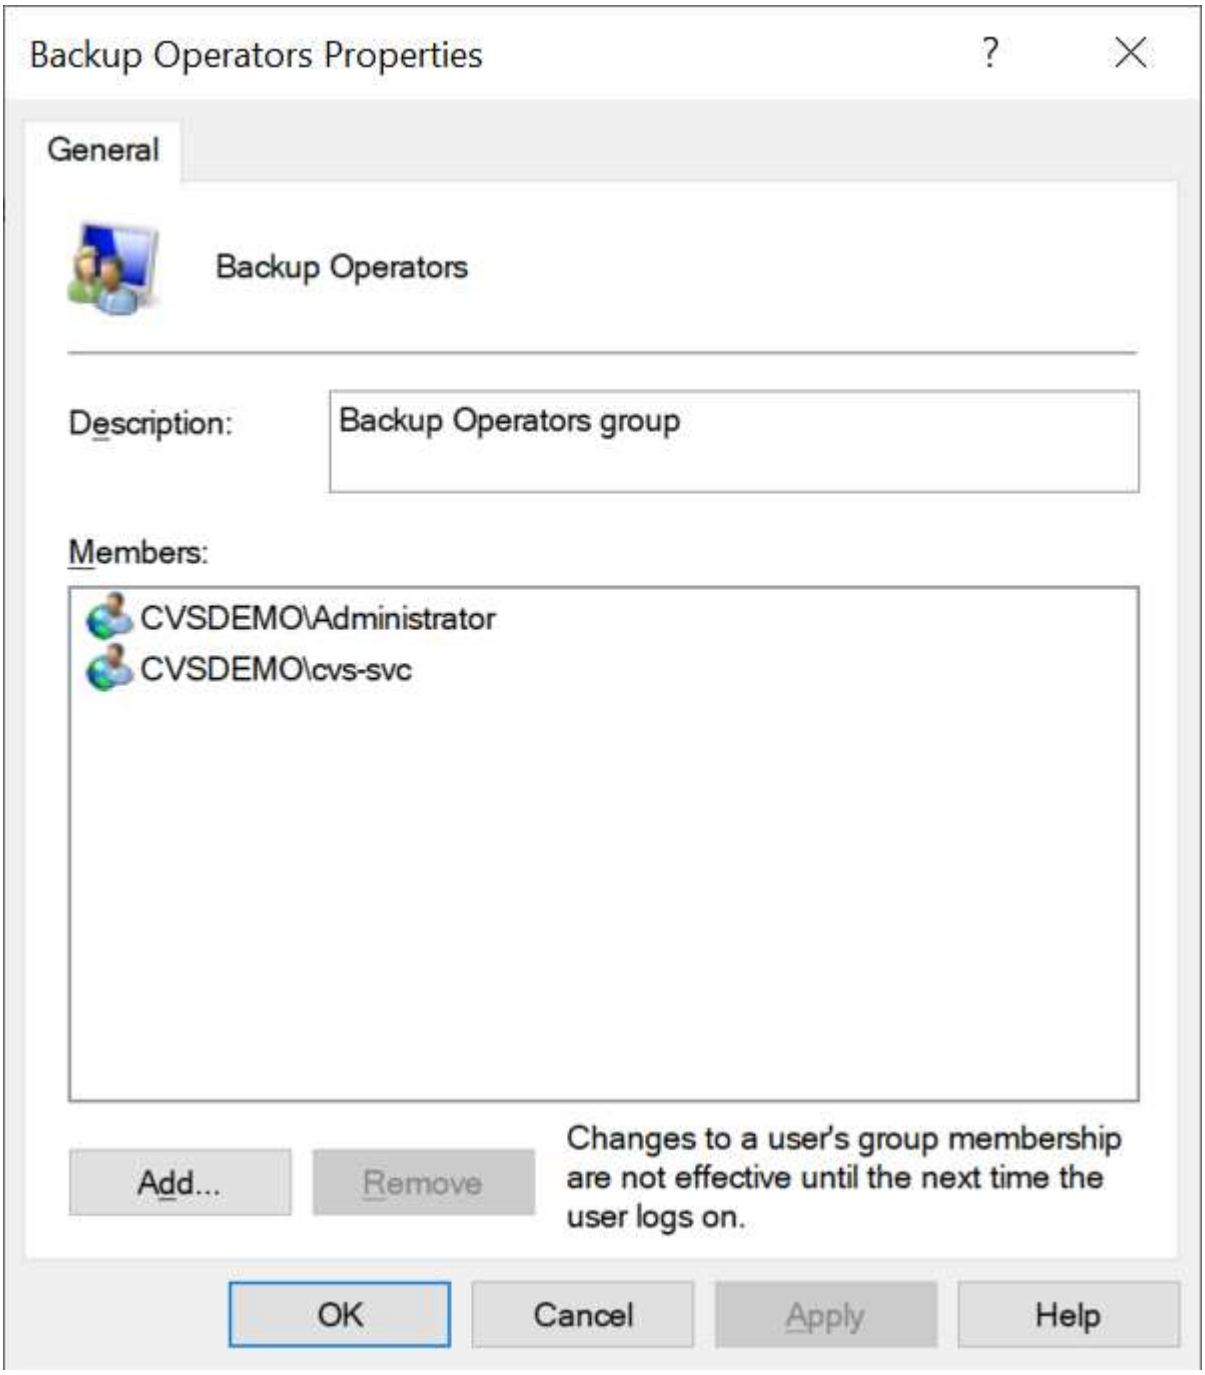
When you specify a user to be added to Backup Users, the user is added to the BUILTIN\Backup Operators group in the Google Cloud NetApp Volumes instance that uses that Active Directory connection, which then gets the SeBackupPrivilege and SeRestorePrivilege.

When you add a user to Security Privilege Users, the user is given the SeSecurityPrivilege, which is useful in

some application use cases, such as SQL Server on SMB shares.



You can view Google Cloud NetApp Volumes local group memberships through the MMC with the proper privileges. The following figure shows users that have been added by using the Google Cloud NetApp Volumes console.

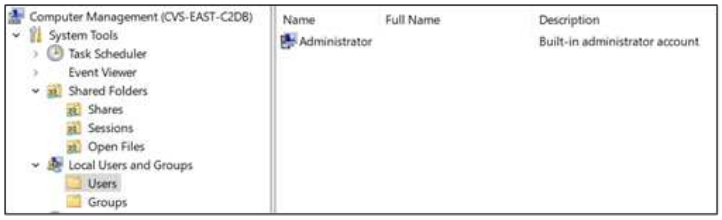


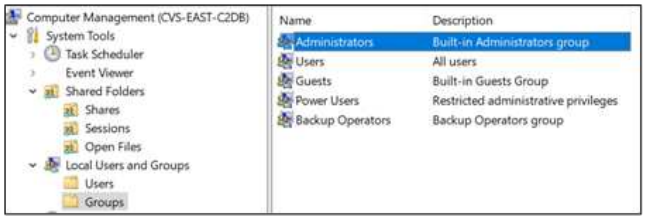
The following table shows the list of default BUILTIN groups and what users/groups are added by default.

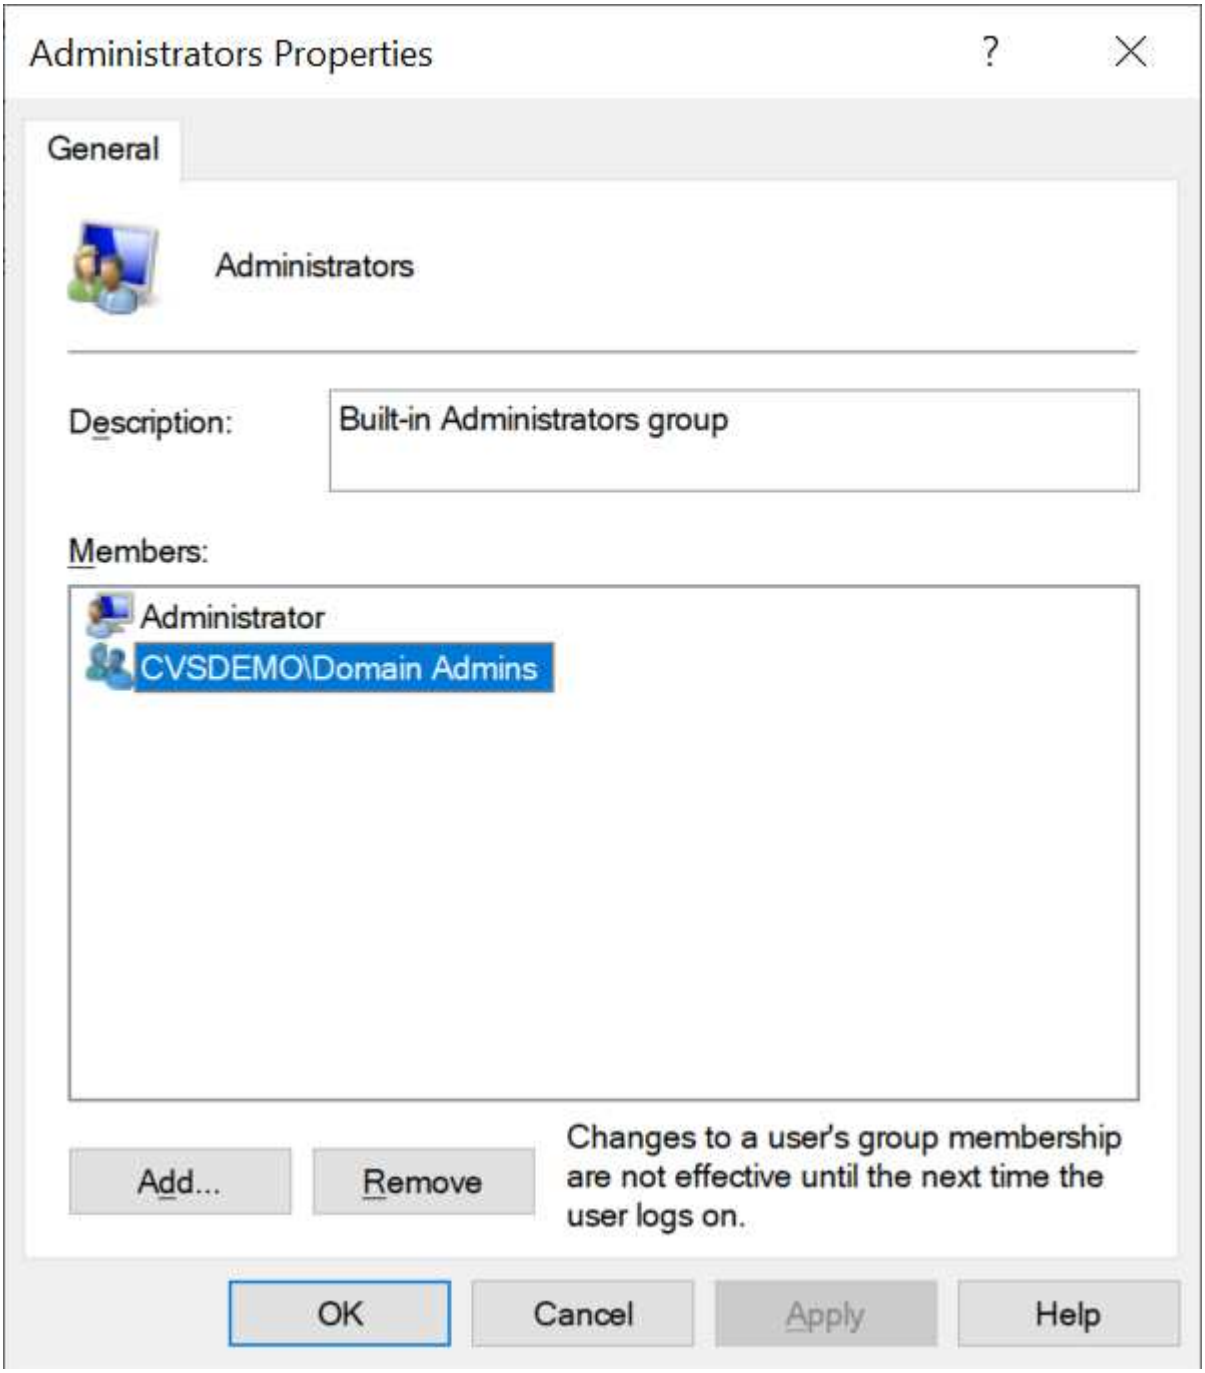
Local/BUILTIN group	Default members
BUILTIN\Administrators*	DOMAIN\Domain Admins
BUILTIN\Backup Operators*	None
BUILTIN\Guests	DOMAIN\Domain Guests
BUILTIN\Power Users	None
BUILTIN\Domain Users	DOMAIN\Domain Users

*Group membership controlled in Google Cloud NetApp Volumes Active Directory connection configuration.

You can view local users and groups (and group members) in the MMC window, but you cannot add or delete objects or change group memberships from this console. By default, only the Domain Admins group and Administrator are added to the BUILTIN\Administrators group in Google Cloud NetApp Volumes. Currently, you cannot modify this.







MMC/Computer Management access

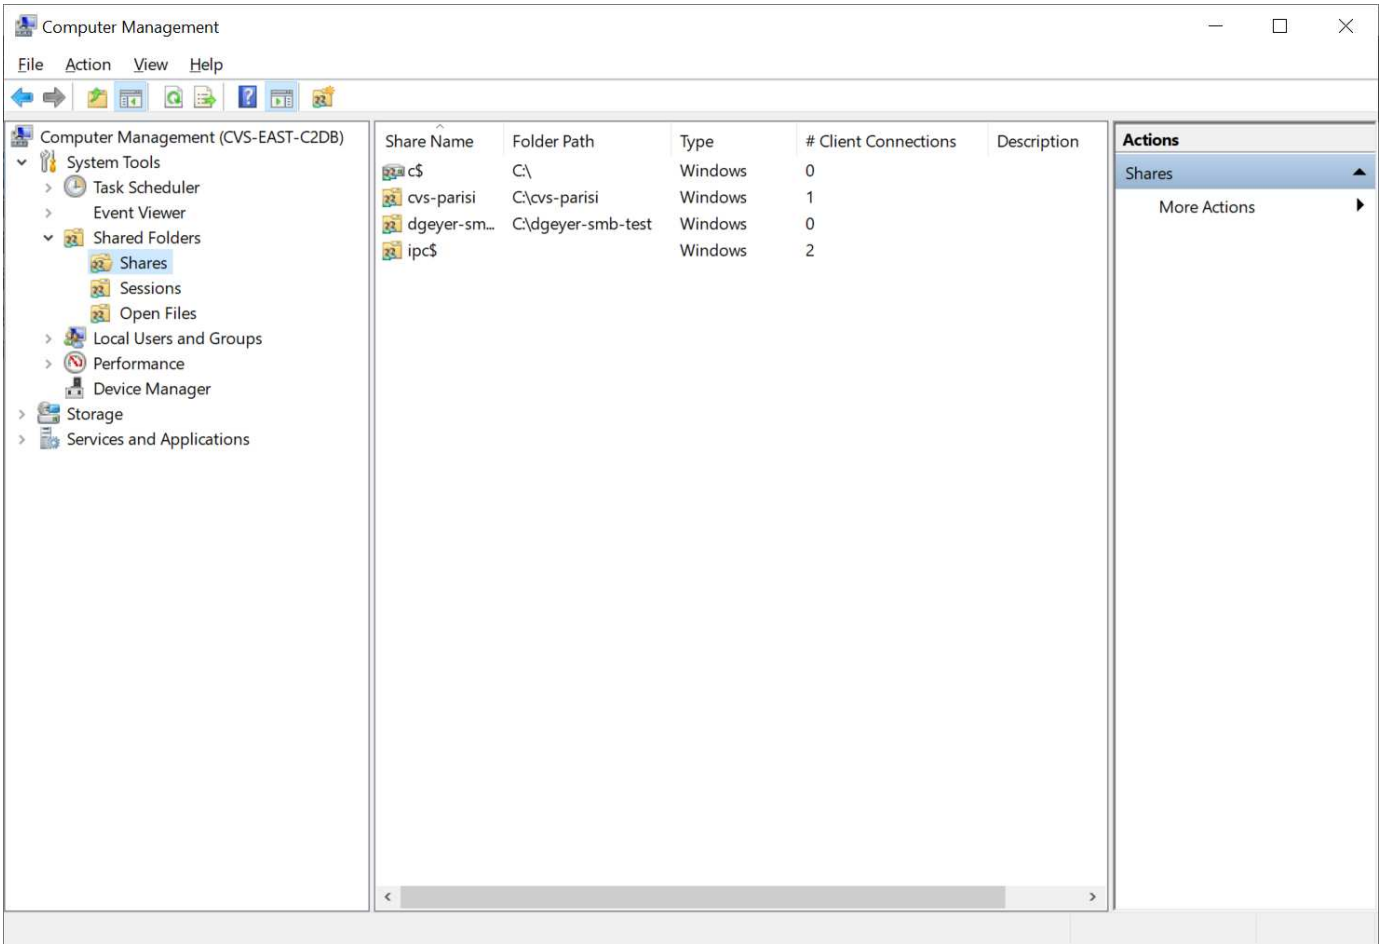
SMB access in Google Cloud NetApp Volumes provides connectivity to the Computer Management MMC, which allows you to view shares, manage share ACLs, ands view/manage SMB sessions and open files.

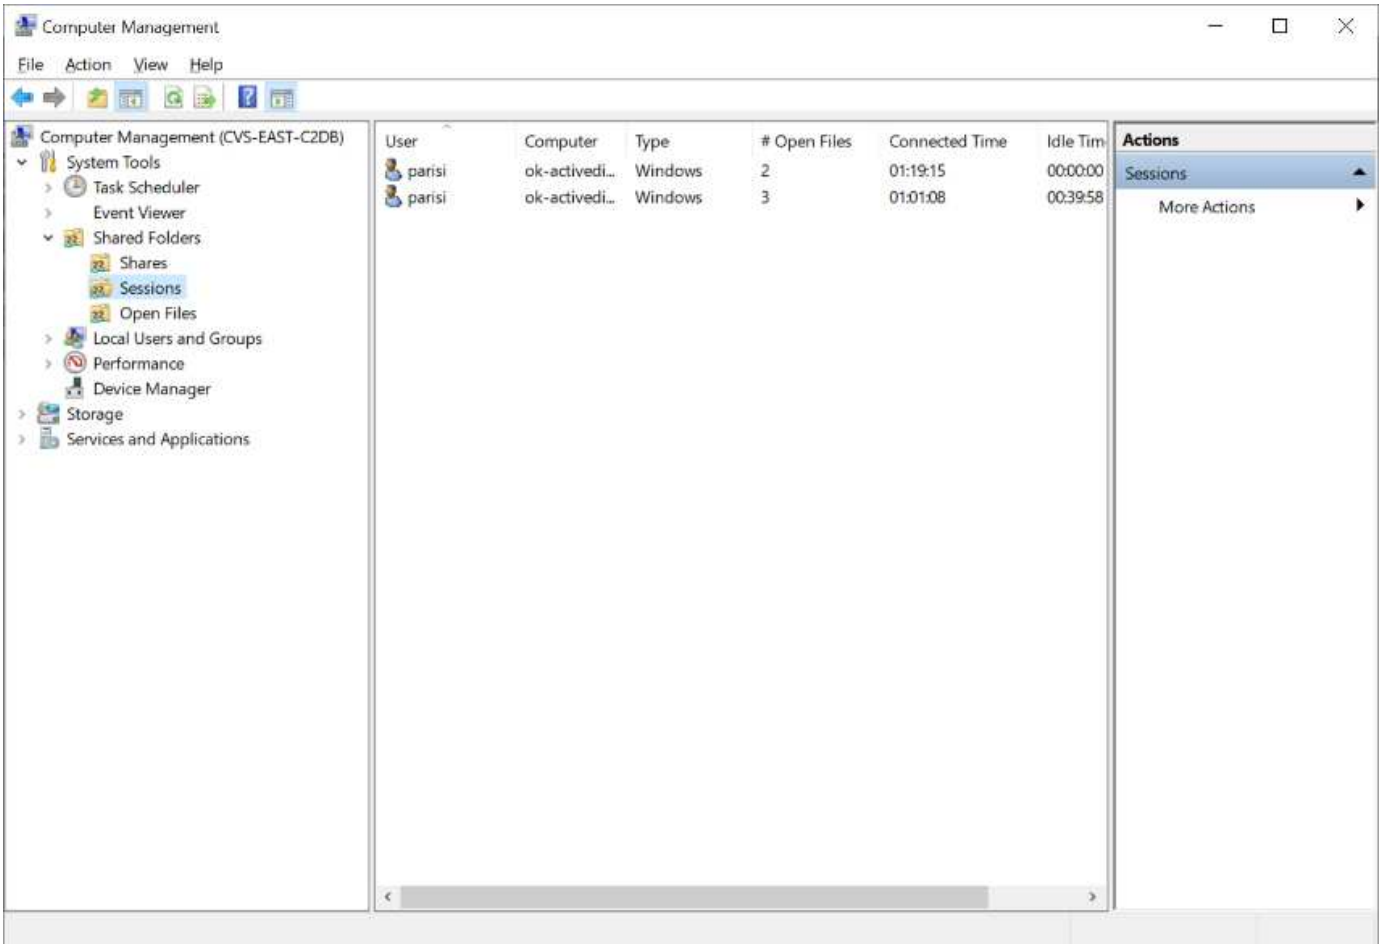
To use the MMC to view SMB shares and sessions in Google Cloud NetApp Volumes, the user logged in currently must be a domain administrator. Other users are allowed access to view or manage the SMB server from MMC and receive a You Do Not Have Permissions dialog box when attempting to view shares or sessions on the Google Cloud NetApp Volumes SMB instance.

To connect to the SMB server, open Computer Management, right click Computer Management and then select Connect To Another Computer. This opens the Select Computer dialog box where you can enter the SMB server name (found in the Google Cloud NetApp Volumes volume information).

When you view SMB shares with the proper permissions, you see all available shares in the Google Cloud NetApp Volumes instance that share the Active Directory connection. To control this behavior, set the Hide SMB Shares option on the Google Cloud NetApp Volumes volume instance.

Remember, only one Active Directory connection is allowed per region.





The following table shows a list of supported/unsupported functionality for the MMC.

Supported functions	Unsupported functions
• View shares • View active SMB sessions • View open files • View local users and groups • View local group memberships • Enumerate the list of sessions, files, and tree connections in the system • Close open files in the system • Close open sessions • Create/manage shares	• Creating new local users/groups • Managing/viewing existing local user/groups • View events or performance logs • Managing storage • Managing services and applications

SMB server security information

The SMB server in Google Cloud NetApp Volumes uses a series of options that define security policies for SMB connections, including things such as Kerberos clock skew, ticket age, encryption, and more.

The following table contains a list of those options, what they do, the default configurations, and if they can be modified with Google Cloud NetApp Volumes. Some options do not apply to Google Cloud NetApp Volumes.

Security option	What it does	Default value	Can change?
Maximum Kerberos Clock Skew (minutes)	Maximum time skew between Google Cloud NetApp Volumes and domain controllers. If the time skew exceeds 5 minutes, Kerberos authentication fails. This is set to the Active Directory default value.	5	No
Kerberos Ticket Lifetime (hours)	Maximum time a Kerberos ticket remains valid before requiring a renewal. If no renewal occurs before the 10 hours, you must obtain a new ticket. Google Cloud NetApp Volumes performs these renewals automatically. 10 hours is the Active Directory default value.	10	No
Maximum Kerberos Ticket Renewal (days)	Maximum number of days that a Kerberos ticket can be renewed before a new authorization request is needed. Google Cloud NetApp Volumes automatically renews tickets for SMB connections. Seven days is the Active Directory default value.	7	No
Kerberos KDC Connection Timeout (secs)	The number of seconds before a KDC connection times out.	3	No
Require Signing for Incoming SMB Traffic	Setting to require signing for SMB traffic. If set to true, clients that do not support signing fail connectivity.	False	
Require Password Complexity for Local User Accounts	Used for passwords on local SMB users. Google Cloud NetApp Volumes does not support local user creation, so this option does not apply to Google Cloud NetApp Volumes.	True	No

Security option	What it does	Default value	Can change?
Use start_tls for Active Directory LDAP Connections	Used to enable start TLS connections for Active Directory LDAP. Google Cloud NetApp Volumes does not currently support enabling this.	False	No
Is AES-128 and AES-256 Encryption for Kerberos Enabled	This controls whether AES encryption is used for Active Directory connections and is controlled with the Enable AES Encryption for Active Directory Authentication option when creating/modifying the Active Directory connection.	False	Yes
LM Compatibility Level	Level of supported authentication dialects for Active Directory connections. See the section "SMB authentication dialects" for more information.	ntlmv2-krb	No
Require SMB Encryption for Incoming CIFS Traffic	Requires SMB encryption for all shares. This is not used by Google Cloud NetApp Volumes; instead, set encryption on a per-volume basis (see the section "SMB share security features").	False	No
Client Session Security	Sets signing and/or sealing for LDAP communication. This is not currently set in Google Cloud NetApp Volumes but might be needed in future releases to address . Remediation for LDAP authentication issues due to the Windows patch is covered in the section "LDAP channel binding.".	None	No
SMB2 enable for DC connections	Uses SMB2 for DC connections. Enabled by default.	System-default	No

Security option	What it does	Default value	Can change?
LDAP Referral Chasing	When using multiple LDAP servers, referral chasing allows the client to refer to other LDAP servers in the list when an entry is not found in the first server. This is currently not supported by Google Cloud NetApp Volumes.	False	No
Use LDAPS for Secure Active Directory Connections	Enables the use of LDAP over SSL. Currently not supported by Google Cloud NetApp Volumes.	False	No
Encryption is required for DC Connection	Requires encryption for successful DC connections. Disabled by default in Google Cloud NetApp Volumes.	False	No

Dual-protocol/multiprotocol

Google Cloud NetApp Volumes offers the ability to share the same datasets to both SMB and NFS clients while maintaining proper access permissions (dual-protocol). This is done by coordinating identity mapping between protocols and using a centralized backend LDAP server to provide the UNIX identities to Google Cloud NetApp Volumes. You can use Windows Active Directory to provide both Windows and UNIX users for ease of use.

Access control

- **Share access controls.** Determine which clients and/or user and groups can access a NAS share. For NFS, export policies and rules control client access to exports. NFS exports are managed from the Google Cloud NetApp Volumes instance. SMB makes use of CIFS/SMB shares and share ACLs to provide more granular control at the user and group level. You can only configure share-level ACLs from SMB clients by using MMC/Computer Management with an account that has administrator rights on the Google Cloud NetApp Volumes instance (see the section "Accounts with local/BUILTIN administrator/backup rights.").
- **File access controls.** Control permissions at a file or folder level and are always managed from the NAS client. NFS clients can make use of traditional mode bits (rwx) or NFSv4 ACLs. SMB clients leverage NTFS permissions.

The access control for volumes that serve data to both NFS and SMB depends on the protocol in use. For information on permissions with dual protocol, see the section "Permission model."

User mapping

When a client accesses a volume, Google Cloud NetApp Volumes attempts to map the incoming user to a valid user in the opposite direction. This is necessary for proper access to be determined across protocols and to ensure that the user requesting access is indeed who they claim to be.

For example, if a Windows user named `joe` attempts access to a volume with UNIX permissions through SMB, then Google Cloud NetApp Volumes performs a search to find a corresponding UNIX user named `joe`. If one exists, then files that are written to an SMB share as Windows user `joe` appears as UNIX user `joe` from NFS clients.

Alternately, if a UNIX user named `joe` attempts access to a Google Cloud NetApp Volumes volume with Windows permissions, then the UNIX user must be able to map to a valid Windows user. Otherwise, access to the volume is denied.

Currently, only Active Directory is supported for external UNIX identity management with LDAP. For more information about configuring access to this service, see Creating an AD connection.

Permission model

When using dual-protocol setups, Google Cloud NetApp Volumes makes use of security styles for volumes to determine the type of ACL. These security styles are set based on which NAS protocol is specified, or in the case of dual protocol, is a choice made at the time of Google Cloud NetApp Volumes volume creation.

- If you are only using NFS, Google Cloud NetApp Volumes volumes use UNIX permissions.
- If you are only using SMB, Google Cloud NetApp Volumes volumes use NTFS permissions.

If you are creating a dual-protocol volume, you can choose the ACL style at volume creation. This decision should be made based on the desired permissions management. If your users manage permissions from Windows/SMB clients, select NTFS. If your users prefer using NFS clients and chmod/chown, use UNIX security styles.

Considerations for creating Active Directory connections

Google Cloud NetApp Volumes provides the ability to connect your Google Cloud NetApp Volumes instance to an external Active Directory server for identity management for both SMB and UNIX users. Creating an Active Directory connection is required to use SMB in Google Cloud NetApp Volumes.

The configuration for this provides several options that require some consideration for security. The external Active Directory server can be an on-premises instance or cloud native. If you are using an on-premises Active Directory server, don't expose the domain to the external network (such as with a DMZ or an external IP address). Instead, use secure private tunnels or VPNs, one-way forest trusts, or dedicated network connections to the on-premises networks with Private Google Access. See the Google Cloud documentation for more information about best practices using Active Directory in Google Cloud.

> NetApp Volumes-SW requires Active Directory servers to be located in the same region. If a DC connection is attempted in NetApp Volumes-SW to another region, the attempt fails. When using NetApp Volumes-SW, be sure to create Active Directory sites that include the Active Directory DCs and then specify sites in Google Cloud NetApp Volumes to avoid cross-region DC connection attempts.

Active Directory credentials

When SMB or LDAP for NFS is enabled, Google Cloud NetApp Volumes interacts with the Active Directory controllers to create a machine account object to use for authentication. This is no different from how a Windows SMB client joins a domain and requires the same access rights to Organizational Units (OUs) in Active Directory.

In many cases, security groups do not allow the use of a Windows administrator account on external servers such as Google Cloud NetApp Volumes. In some cases, the Windows Administrator user is disabled entirely as a security best practice.

Permissions needed to create SMB machine accounts

To add Google Cloud NetApp Volumes machine objects to an Active Directory, an account that either has administrative rights to the domain or has delegated permissions to create and modify machine account objects to a specified OU is required. You can do this with the Delegation of Control Wizard in Active Directory by creating a custom task that provides a user access to creation/deletion of computer objects with the following access permissions provided:

- Read/Write
- Create/Delete All Child Objects
- Read/Write All Properties
- Change/Reset Password

Doing this automatically adds a security ACL for the defined user to the OU in Active Directory and minimizes the access to the Active Directory environment. After a user has been delegated, that username and password can be provided as Active Directory Credentials in this window.



The username and password that is passed to the Active Directory domain leverages Kerberos encryption during the machine account object query and creation for added security.

Active Directory connection details

The Active Directory Connection Details provide guidance for administrators to give specific Active Directory schema information for machine account placement, such as the following:

- **Active Directory Connection Type.** Used to specify whether the Active Directory connection in a region is used for volumes of either Google Cloud NetApp Volumes or NetApp Volumes-Performance service type. If this is set incorrectly on an existing connection, it might not work properly when used or edited.
- **Domain.** The Active Directory domain name.
- **Site.** Limits Active Directory servers to a specific site for security and performance considerations. This is necessary when multiple Active Directory servers span regions because Google Cloud NetApp Volumes does not currently support allowing Active Directory authentication requests to Active Directory servers in a different region than the Google Cloud NetApp Volumes instance. (For instance, the Active Directory domain controller is in a region that only NetApp Volumes-Performance supports but you want an SMB share in a NetApp Volumes-SW instance.)
- **DNS servers.** DNS servers to use in name lookups.
- **NetBIOS name (optional).** If desired, the NetBIOS name for the server. This what is used when new machine accounts are created using the Active Directory connection. For instance, if the NetBIOS name is set to NetApp Volumes-EAST then the machine account names will be NetApp Volumes-EAST-{1234}. See the section "How Google Cloud NetApp Volumes shows up in Active Directory" for more information.
- **Organizational Unit (OU).** The specific OU to create the computer account. This is useful if you're delegating control to a user for machine accounts to a specific OU.
- **AES Encryption.** You can also check or uncheck the Enable AES Encryption for AD Authentication checkbox. Enabling AES encryption for Active Directory authentication provides extra security for Google Cloud NetApp Volumes to Active Directory communication during user and group lookups. Before enabling this option, check with your domain administrator to confirm that the Active Directory domain controllers

support AES authentication.

> By default, most Windows servers do not disable weaker ciphers (such as DES or RC4-HMAC), but if you choose to disable weaker ciphers, confirm Google Cloud NetApp Volumes Active Directory connection has been configured to enable AES. Otherwise, authentication failures occur. Enabling AES encryption doesn't disable weaker ciphers but instead adds support for AES ciphers to the Google Cloud NetApp Volumes SMB machine account.

Kerberos realm details

This option does not apply to SMB servers. Rather, it is used when configuring NFS Kerberos for the Google Cloud NetApp Volumes system. When these details are populated, the NFS Kerberos realm is configured (similar to a krb5.conf file on Linux) and is used when NFS Kerberos is specified on the Google Cloud NetApp Volumes volume creation, as the Active Directory connection acts as the NFS Kerberos Distribution Center (KDC).

> Non-Windows KDCs are currently unsupported for use with Google Cloud NetApp Volumes.

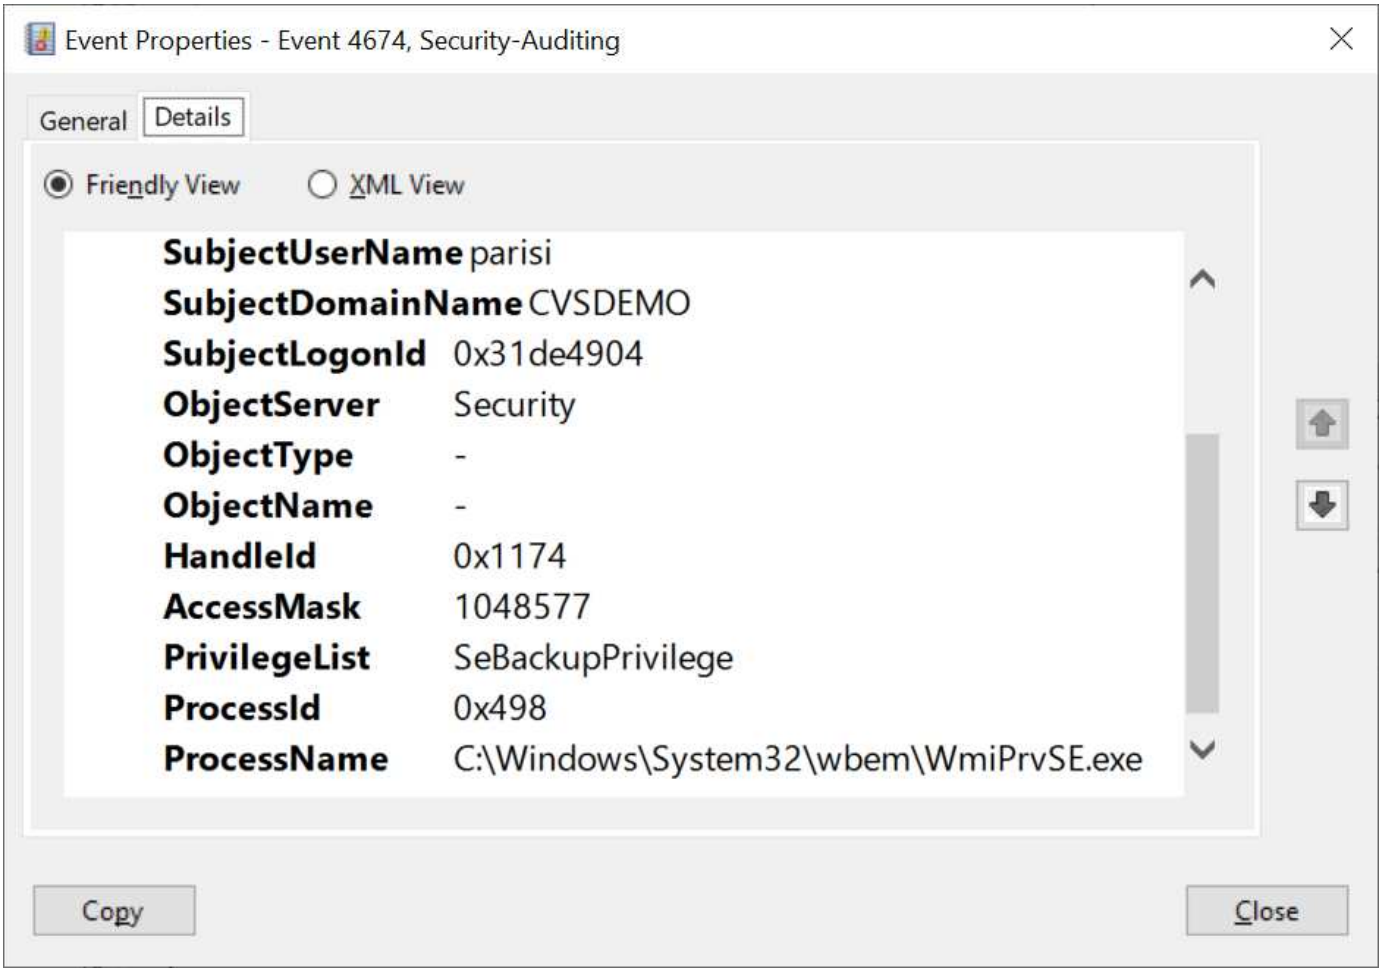
Region

A region enables you to specify the location where the Active Directory connection resides. This region must be the same region as the Google Cloud NetApp Volumes volume.

- **Local NFS Users with LDAP.** In this section, there is also an option to Allow Local NFS Users with LDAP. This option must be left unselected if you want to extend your UNIX user group membership support beyond the 16-group limitation of NFS (extended groups). However, using extended groups requires a configured LDAP server for UNIX identities. If you don't have an LDAP server, leave this option unselected. If you have an LDAP server and want to also use local UNIX users (such as root), select this option.

Backup users

This option enables you to specify Windows users that have backup permissions to the Google Cloud NetApp Volumes volume. Backup privileges (SeBackupPrivilege) are necessary for some applications to properly backup and restore data in NAS volumes. This user has a high level of access to data in the volume, so you should consider enabling auditing of that user access. After it is enabled, audit events display in Event Viewer > Windows Logs > Security.



Security privilege users

This option enables you to specify Windows users that have security modification permissions to the Google Cloud NetApp Volumes volume. Security privileges (SeSecurityPrivilege) are necessary for some applications (such as SQL Server) to properly set permissions during installation. This privilege is needed to manage the security log. Although this privilege is not as powerful as SeBackupPrivilege, NetApp recommends auditing user access of users with this privilege level if needed.

For more information, see Special privileges assigned to new logon.

How Google Cloud NetApp Volumes shows up in Active Directory

Google Cloud NetApp Volumes shows up in Active Directory as a normal machine account object. The naming conventions are as follows.

- CIFS/SMB and NFS Kerberos create separate machine account objects.
- NFS with LDAP enabled creates a machine account in Active Directory for Kerberos LDAP binds.
- Dual protocol volumes with LDAP share the CIFS/SMB machine account for LDAP and SMB.
- CIFS/SMB machine accounts use a naming convention of NAME-1234 (random four digit ID with hyphen appended to <10 character name) for the machine account. You can define NAME by the NetBIOS name setting on the Active Directory connection (see the section "Active Directory connection details").
- NFS Kerberos uses NFS-NAME-1234 as the naming convention (up to 15 characters). If more than 15 characters are used, the name is NFS-TRUNCATED-NAME-1234.

- NFS-only NetApp Volumes-Performance instances with LDAP enabled create an SMB machine account for binding to the LDAP server with the same naming convention as CIFS/SMB instances.
- When an SMB machine account is created, default hidden admin shares (see the section "Default hidden shares") are also created (c$, admin$, ipc$), but those shares have no ACLs assigned and are inaccessible.
- The machine account objects are placed in CN=Computers by default, but a you can specify a different OU when necessary. See the section "Permissions needed to create SMB machine accounts" for information about what access rights are needed to add/remove machine account objects for Google Cloud NetApp Volumes.

When Google Cloud NetApp Volumes adds the SMB machine account to Active Directory, the following fields are populated:

- cn (with the specified SMB server name)
- dNSHostName (with SMBserver.domain.com)
- msDS-SupportedEncryptionTypes (Allows DES_CBC_MD5, RC4_HMAC_MD5 if AES encryption is not enabled; if AES encryption is enabled, DES_CBC_MD5, RC4_HMAC_MD5, AES128_CTS_HMAC_SHA1_96, AES256_CTS_HMAC_SHA1_96 are allowed for Kerberos ticket exchange with the machine account for SMB)
- name (with the SMB server name)
- sAMAccountName (with SMBserver$)
- servicePrincipalName (with host/smbserver.domain.com and host/smbserver SPNs for Kerberos)

If you want to disable weaker Kerberos encryption types (enctype) on the machine account, you can change the msDS-SupportedEncryptionTypes value on the machine account to one of the values in the following table to allow AES only.

msDS-SupportedEncryptionTypes value	Enctype enabled
2	DES_CBC_MD5
4	RC4_HMAC
8	AES128_CTS_HMAC_SHA1_96 only
16	AES256_CTS_HMAC_SHA1_96 only
24	AES128_CTS_HMAC_SHA1_96 and AES256_CTS_HMAC_SHA1_96
30	DES_CBC_MD5, RC4_HMAC, AES128_CTS_HMAC_SHA1_96 and AES256_CTS_HMAC_SHA1_96

To enable AES encryption for SMB machine accounts, click Enable AES Encryption for AD Authentication when creating the Active Directory connection.

To enable AES encryption for NFS Kerberos, see the Google Cloud NetApp Volumes documentation.

Other NAS Infrastructure service dependencies (KDC, LDAP, and DNS)

When using Google Cloud NetApp Volumes for NAS shares, there might be external dependencies required for proper functionality. These dependencies are in play under

specific circumstances. The following table shows various configuration options and what, if any, dependencies are required.

Configuration	Dependencies required
NFSv3 only	None
NFSv3 Kerberos only	Windows Active Directory: * KDC * DNS * LDAP
NFSv4.1 only	Client ID mapping configuration (/etc/idmap.conf)
NFSv4.1 Kerberos only	• Client ID mapping configuration (/etc/idmap.conf) • Windows Active Directory: KDC DNS LDAP
SMB only	Active Directory: * KDC * DNS
Multiprotocol NAS (NFS and SMB)	• Client ID mapping configuration (NFSv4.1 only; /etc/idmap.conf) • Windows Active Directory: KDC DNS LDAP

Kerberos keytab rotation/password resets for machine account objects

With SMB machine accounts, Google Cloud NetApp Volumes schedules periodic password resets for the SMB machine account. These password resets occur using Kerberos encryption and operate on a schedule of every fourth Sunday at a random time between 11PM and 1AM. These password resets change the Kerberos key versions, rotate the keytabs stored on the Google Cloud NetApp Volumes system, and help maintain a greater level of security for SMB servers running in Google Cloud NetApp Volumes. Machine account passwords are randomized and are not known to administrators.

For NFS Kerberos machine accounts, password resets take place only when a new keytab is created/exchanged with the KDC. Currently, this is not possible to do in Google Cloud NetApp Volumes.

Network ports for use with LDAP and Kerberos

When using LDAP and Kerberos, you should determine the network ports in use by these services. You can find a complete list of ports in use by Google Cloud NetApp Volumes in the Google Cloud NetApp Volumes documentation on security considerations.

LDAP

Google Cloud NetApp Volumes acts as an LDAP client and uses standard LDAP search queries for user and group lookups for UNIX identities. LDAP is necessary if you intend to use users and groups outside the standard default users provided by Google Cloud NetApp Volumes. LDAP is also necessary if you plan on using NFS Kerberos with user principals (such as user1@domain.com). Currently, only LDAP using Microsoft

Active Directory is supported.

To use Active Directory as a UNIX LDAP server, you must populate the necessary UNIX attributes on users and groups you intend to use for UNIX identities. Google Cloud NetApp Volumes uses a default LDAP schema template that queries attributes based on RFC-2307-bis. As a result, the following table shows the bare minimum necessary Active Directory attributes to populate for users and groups and what each attribute is used for.

For more information about setting LDAP attributes in Active Directory, see Managing dual-protocol access.

Attribute	What it does
uid*	Specifies the UNIX user name
uidNumber*	Specifies the UNIX user's numeric ID
gidNumber*	Specifies the UNIX user's primary group numeric ID
objectClass*	Specifies what type of object is being used; Google Cloud NetApp Volumes requires "user" to be included in the list of object classes (is included in most Active Directory deployments by default).
name	General information about the account (real name, phone number, and so on—also known as gecos)
unixUserPassword	No need to set this; not used in UNIX identity lookups for NAS authentication. Setting this puts the configured unixUserPassword value in plaintext.
unixHomeDirectory	Defines path to UNIX home directories when a user authenticates against LDAP from a Linux client. Set this if you want to use LDAP for UNIX home directory functionality.
loginShell	Defines path to the bash/profile shell for Linux clients when a user authenticates against LDAP.

*Denotes attribute is required for proper functionality with Google Cloud NetApp Volumes. Remaining attributes are for client-side use only.

Attribute	What it does
cn*	Specifies the UNIX group name. When using Active Directory for LDAP, this is set when the object is first created, but it can be changed later. This name cannot be the same as other objects. For instance, if your UNIX user named user1 belongs to a group named user1 on your Linux client, Windows doesn't allow two objects with the same cn attribute. To work around this, rename the Windows user to a unique name (such as user1-UNIX); LDAP in Google Cloud NetApp Volumes uses the uid attribute for UNIX user names.
gidNumber*	Specifies the UNIX group numeric ID.

Attribute	What it does
objectClass*	Specifies what type of object is being used; Google Cloud NetApp Volumes requires group to be included in the list of object classes (this attribute is included in most Active Directory deployments by default).
memberUid	Specifies which UNIX users are members of the UNIX group. With Active Directory LDAP in Google Cloud NetApp Volumes, this field is not necessary. The Google Cloud NetApp Volumes LDAP schema uses the Member field for group memberships.
Member*	Required for group memberships/secondary UNIX groups. This field is populated by adding Windows users to Windows groups. However, if the Windows groups don't have UNIX attributes populated, they are not included in the UNIX user's group membership lists. Any groups that need to be available in NFS must populate the required UNIX group attributes listed in this table.

*Denotes attribute is required for proper functionality with Google Cloud NetApp Volumes. Remaining attributes are for client-side use only.

LDAP bind information

To query users in LDAP, Google Cloud NetApp Volumes must bind (login) to the LDAP service. This login has read-only permissions and is used to query LDAP UNIX attributes for directory lookups. Currently, LDAP binds are possible only by using an SMB machine account.

You can only enable LDAP for `NetApp Volumes-Performance` instances and use it for NFSv3, NFSv4.1, or dual-protocol volumes. An Active Directory connection must be established in the same region as the Google Cloud NetApp Volumes volume for successful deployment of the LDAP-enabled volume.

When LDAP is enabled, the following occurs in specific scenarios.

- If only NFSv3 or NFSv4.1 is used for the Google Cloud NetApp Volumes project, then a new machine account is created in the Active Directory domain controller, and the LDAP client in Google Cloud NetApp Volumes binds to Active Directory by using the machine account credentials. No SMB shares are created for the NFS volume and default hidden administrative shares (see the section "Default hidden shares") have share ACLs removed.
- If dual-protocol volumes are used for the Google Cloud NetApp Volumes project, then only the single machine account created for SMB access is used to bind the LDAP client in Google Cloud NetApp Volumes to Active Directory. No additional machine accounts are created.
- If dedicated SMB volumes are created separately (either before or after NFS volumes with LDAP are enabled), then the machine account for LDAP binds is shared with the SMB machine account.
- If NFS Kerberos is also enabled, two machine accounts are created—one for SMB shares and/or LDAP binds and one for NFS Kerberos authentication.

LDAP queries

Although LDAP binds are encrypted, LDAP queries are passed over the wire in plaintext by using the common LDAP port 389. This well-known port cannot currently be changed in Google Cloud NetApp Volumes. As a

result, someone with access to packet sniffing in the network can see user and group names, numeric IDs, and group memberships.

However, Google Cloud VMs cannot sniff other VM's unicast traffic. Only VMs actively participating in LDAP traffic (that is, being able to bind) can see traffic from the LDAP server. For more information about packet sniffing in Google Cloud NetApp Volumes, see the section "Packet sniffing/trace considerations."

LDAP client configuration defaults

When LDAP is enabled in a Google Cloud NetApp Volumes instance, an LDAP client configuration is created with specific configuration details by default. In some cases, options either do not apply to Google Cloud NetApp Volumes (not supported) or are not configurable.

LDAP client option	What it does	Default value	Can change?
LDAP Server List	Sets LDAP server names or IP addresses to use for queries. This is not used for Google Cloud NetApp Volumes. Instead, Active Directory Domain is used to define LDAP servers.	Not set	No
Active Directory Domain	Sets the Active Directory Domain to use for LDAP queries. Google Cloud NetApp Volumes leverages SRV records for LDAP in DNS to find LDAP servers in the domain.	Set to the Active Directory domain specified in the Active Directory connection.	No
Preferred Active Directory Servers	Sets the preferred Active Directory servers to use for LDAP. Not supported by Google Cloud NetApp Volumes. Instead, use Active Directory sites to control LDAP server selection.	Not set.	No
Bind using SMB Server Credentials	Binds to LDAP by using the SMB machine account. Currently, the only supported LDAP bind method in Google Cloud NetApp Volumes.	True	No
Schema Template	The schema template used for LDAP queries.	MS-AD-BIS	No

LDAP client option	What it does	Default value	Can change?
LDAP Server Port	The port number used for LDAP queries. Google Cloud NetApp Volumes currently uses only the standard LDAP port 389. LDAPS/port 636 is not currently supported.	389	No
Is LDAPS Enabled	Controls whether LDAP over Secure Sockets Layer (SSL) is used for queries and binds. Currently not supported by Google Cloud NetApp Volumes.	False	No
Query Timeout (sec)	Timeout for queries. If queries take longer than the specified value, queries fail.	3	No
Minimum Bind Authentication Level	The minimum supported bind level. Because Google Cloud NetApp Volumes uses machine accounts for LDAP binds and Active Directory does not support anonymous binds by default, this option does not come into play for security.	Anonymous	No
Bind DN	The user/distinguished name (DN) used for binds when simple bind is used. Google Cloud NetApp Volumes uses machine accounts for LDAP binds and does not currently support simple bind authentication.	Not set	No
Base DN	The base DN used for LDAP searches.	The Windows domain use for the Active Directory connection, in DN format (that is, DC=domain, DC=local).	No
Base search scope	The search scope for base DN searches. Values can include base, onelevel, or subtree. Google Cloud NetApp Volumes only supports subtree searches.	Subtree	No

LDAP client option	What it does	Default value	Can change?
User DN	Defines the DN where user searches start for LDAP queries. Currently not supported for Google Cloud NetApp Volumes, so all user searches start at the base DN.	Not set	No
User search scope	The search scope for user DN searches. Values can include base, onelevel, or subtree. Google Cloud NetApp Volumes does not support setting the user search scope.	Subtree	No
Group DN	Defines the DN where group searches start for LDAP queries. Currently not supported for Google Cloud NetApp Volumes, so all group searches start at the base DN.	Not set	No
Group search scope	The search scope for group DN searches. Values can include base, onelevel, or subtree. Google Cloud NetApp Volumes does not support setting the group search scope.	Subtree	No
Netgroup DN	Defines the DN where netgroup searches start for LDAP queries. Currently not supported for Google Cloud NetApp Volumes, so all netgroup searches start at the base DN.	Not set	No
Netgroup search scope	The search scope for netgroup DN searches. Values can include base, onelevel, or subtree. Google Cloud NetApp Volumes does not support setting the netgroup search scope.	Subtree	No

LDAP client option	What it does	Default value	Can change?
Use start_tls over LDAP	Leverages Start TLS for certificate based LDAP connections over port 389. Currently not supported by Google Cloud NetApp Volumes.	False	No
Enable netgroup-by-host lookup	Enables netgroup lookups by hostname rather than expanding netgroups to list all members. Currently not supported by Google Cloud NetApp Volumes.	False	No
Netgroup-by-host DN	Defines the DN where netgroup-by-host searches start for LDAP queries. Netgroup-by-host is currently not supported for Google Cloud NetApp Volumes.	Not set	No
Netgroup-by-host search scope	The search scope for netgroup-by-host DN searches. Values can include base, onelevel or subtree. Netgroup-by-host is currently not supported for Google Cloud NetApp Volumes.	Subtree	No
Client session security	Defines what level of session security is used by LDAP (sign, seal, or none). LDAP signing is supported by NetApp Volumes-Performance, if requested by Active Directory. NetApp Volumes-SW does not support LDAP signing. For both service types, sealing is currently not supported.	None	No
LDAP referral chasing	When using multiple LDAP servers, referral chasing allows the client to refer to other LDAP servers in the list when an entry is not found in the first server. This is currently not supported by Google Cloud NetApp Volumes.	False	No

LDAP client option	What it does	Default value	Can change?
Group membership filter	Provides a custom LDAP search filter to be used when looking up group membership from an LDAP server. Not currently supported with Google Cloud NetApp Volumes.	Not set	No

Using LDAP for asymmetric name mapping

Google Cloud NetApp Volumes, by default, maps Windows users and UNIX users with identical usernames bidirectionally without special configuration. As long as Google Cloud NetApp Volumes can find a valid UNIX user (with LDAP), then 1:1 name mapping occurs. For instance, if Windows user `johnsmith` is used, then, if Google Cloud NetApp Volumes can find a UNIX user named `johnsmith` in LDAP, name mapping succeeds for that user, all files/folders created by `johnsmith` show the correct user ownership, and all ACLs affecting `johnsmith` are honored regardless of the NAS protocol in use. This is known as symmetric name mapping.

Asymmetric name mapping is when the Windows user and UNIX user identity don't match. For instance, if Windows user `johnsmith` has a UNIX identity of `jsmith`, Google Cloud NetApp Volumes needs a way to be told about the variation. Because Google Cloud NetApp Volumes currently doesn't support creation of static name mapping rules, LDAP must be used to look up the identity of the users for both Windows and UNIX identities to ensure proper ownership of files and folders and expected permissions.

By default, Google Cloud NetApp Volumes includes `LDAP` in the ns-switch of the instance for the name map database, so that to provide name mapping functionality by using LDAP for asymmetric names, you only need to modify some of the user/group attributes to reflect what Google Cloud NetApp Volumes looks for.

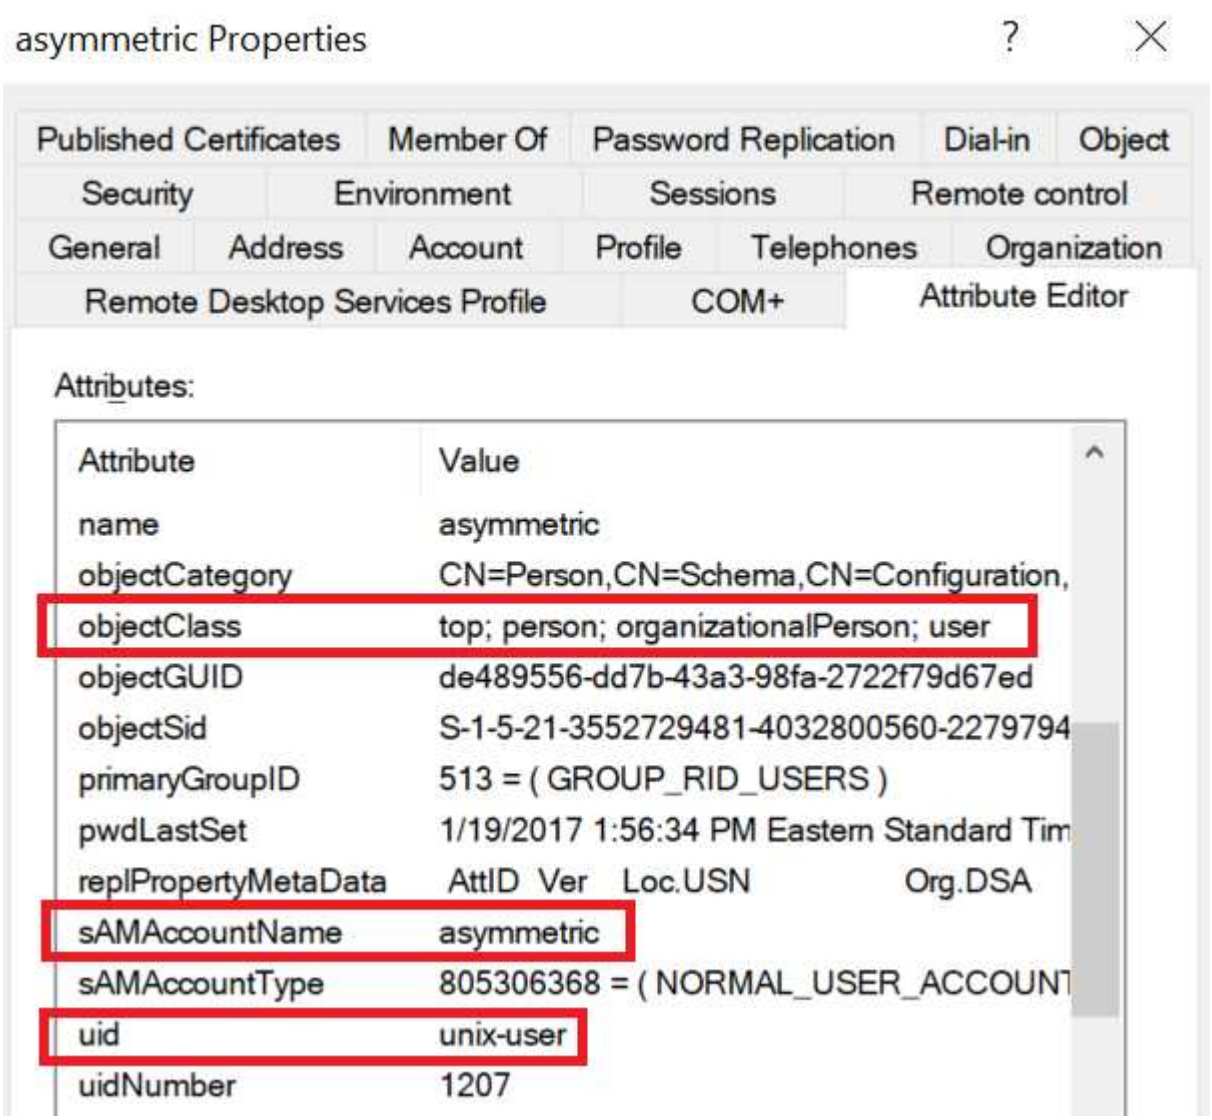
The following table shows what attributes must be populated in LDAP for asymmetric name mapping functionality. In most cases, Active Directory is already configured to do this.

Google Cloud NetApp Volumes attribute	What it does	Value used by Google Cloud NetApp Volumes for name mapping
Windows to UNIX objectClass	Specifies the type of object being used. (That is, user, group, posixAccount, and so on)	Must include user (can contain multiple other values, if desired.)
Windows to UNIX attribute	that defines the Windows username at creation. Google Cloud NetApp Volumes uses this for Windows to UNIX lookups.	No change needed here; sAMAccountName is the same as the Windows login name.
UID	Defines the UNIX username.	Desired UNIX username.

Google Cloud NetApp Volumes currently does not use domain prefixes in LDAP lookups, so multiple domain LDAP environments do not function properly with LDAP namemap lookups.

The following example shows a user with the Windows name `asymmetric`, the UNIX name `unix-user`, and the behavior it follows when writing files from both SMB and NFS.

The following figure shows how LDAP attributes look from the Windows server.



From an NFS client, you can query the UNIX name but not the Windows name:

```
# id unix-user
uid=1207(unix-user) gid=1220(sharedgroup) groups=1220(sharedgroup)
# id asymmetric
id: asymmetric: no such user
```

When a file is written from NFS as `unix-user`, the following is the result from the NFS client:

```
sh-4.2$ pwd
/mnt/home/ntfssh-4.2$ touch unix-user-file
sh-4.2$ ls -la | grep unix-user
-rwx------  1 unix-user sharedgroup     0 Feb 28 12:37 unix-user-nfs
sh-4.2$ id
uid=1207(unix-user) gid=1220(sharedgroup) groups=1220(sharedgroup)
```

From a Windows client, you can see that the owner of the file is set to the proper Windows user:

```
PS C:\ > Get-Acl \\demo\home\ntfs\unix-user-nfs | select Owner
Owner
-----
NTAP\asymmetric
```

Conversely, files created by the Windows user `asymmetric` from an SMB client show the proper UNIX owner, as shown in the following text.

SMB:

```
PS Z:\ntfs> echo TEXT > asymmetric-user-smb.txt
```

NFS:

```
sh-4.2$ ls -la | grep asymmetric-user-smb.txt
-rwx------  1 unix-user         sharedgroup   14 Feb 28 12:43 asymmetric-
user-smb.txt
sh-4.2$ cat asymmetric-user-smb.txt
TEXT
```

LDAP channel binding

Because of a vulnerability with Windows Active Directory domain controllers, Microsoft Security Advisory ADV190023 changes how DCs allow LDAP binds.

The impact for Google Cloud NetApp Volumes is the same as for any LDAP client. Google Cloud NetApp Volumes does not currently support channel binding. Because Google Cloud NetApp Volumes supports LDAP signing by default through negotiation, LDAP channel binding should not be an issue. If you do have issues binding to LDAP with channel binding enabled, follow the remediation steps in ADV190023 to allow LDAP binds from Google Cloud NetApp Volumes to succeed.

DNS

Active Directory and Kerberos both have dependencies on DNS for host name to IP/IP to host name resolution. DNS requires port 53 to be open. Google Cloud NetApp Volumes does not make any modifications to DNS records, nor does it currently support the use of dynamic DNS on network interfaces.

You can configure Active Directory DNS to restrict which servers can update DNS records. For more information, see Secure Windows DNS.

Note that resources within a Google project default to using Google Cloud DNS, which isn't connected with Active Directory DNS. Clients using Cloud DNS cannot resolve UNC paths returned by Google Cloud NetApp Volumes. Windows clients joined to the Active Directory domain are configured to use Active Directory DNS and can resolve such UNC paths.

To join a client to Active Directory, you must configure its DNS configuration to use Active Directory DNS.

Optionally, you can configure Cloud DNS to forward requests to Active Directory DNS. See Why can't my client resolve the SMB NetBIOS name? for more information.

Google Cloud NetApp Volumes does not currently support DNSSEC and DNS queries are performed in plaintext.

File access auditing

Currently not supported for Google Cloud NetApp Volumes.

Antivirus protection

You must perform antivirus scanning in Google Cloud NetApp Volumes at the client to a NAS share. There is currently no native antivirus integration with Google Cloud NetApp Volumes.

Service operation

The Google Cloud NetApp Volumes team manages the backend services in Google Cloud and uses multiple strategies to secure the platform and prevent unwanted access.

Each customer gets their own unique subnet that has access fenced off from other customers by default, and every tenant in Google Cloud NetApp Volumes gets their own namespace and VLAN for total data isolation. After a user is authenticated, the Service Delivery Engine (SDE) can only read configuration data specific to that tenant.

Physical security

With proper preapproval, only onsite engineers and NetApp-badged Field Support Engineers (FSEs) have access to the cage and racks for physical work. Storage and network management is not permitted. Only these onsite resources are able to perform hardware maintenance tasks.

For onsite engineers, a ticket is raised for the statement of work (SOW) that includes the rack ID and device location (RU) and all other details are included in the ticket. For NetApp FSEs, a site visitation ticket must be raised with the COLO and the ticket includes the visitor's details, date, and time for auditing purposes. The SOW for the FSE is communicated internally to NetApp.

Operations team

The operations team for Google Cloud NetApp Volumes consists of Production Engineering and a Site Reliability Engineer (SRE) for Google Cloud NetApp Volumes and NetApp Field Support Engineers and Partners for hardware. All operations team members are accredited for work in Google Cloud and detailed records of work are maintained for every ticket raised. In addition, there is a stringent change control and approval process in place to ensure each decision is appropriately scrutinized.

The SRE team manages the control plane and how the data is routed from UI requests to backend hardware and software in Google Cloud NetApp Volumes. The SRE team also manages system resources, such as volume and inode maximums. SREs are not allowed to interact with or have access to customer data. SREs also provide coordination with Return Material Authorizations (RMAs), such as new disk or memory replacement requests for the backend hardware.

Customer responsibilities

Customers of Google Cloud NetApp Volumes manage their organization's Active Directory and user role management as well as the volume and data operations. Customers can have administrative roles and can

delegate permissions to other end users within the same Google Cloud project using the two predefined roles that NetApp and Google Cloud provide (Administrator and Viewer).

The administrator can peer any VPC within the customer project to Google Cloud NetApp Volumes that the customer determines to be appropriate. It is the responsibility of the customer to manage access to their Google Cloud marketplace subscription and to manage the VPCs that have access to the data plane.

Malicious SRE protection

One concern that could arise is how does Google Cloud NetApp Volumes protect against scenarios in which there is a malicious SRE or when SRE credentials have been compromised?

Access to the production environment is with a limited number of SRE individuals only. Administrative privileges are further restricted to a handful of experienced administrators. All actions performed by anyone in the Google Cloud NetApp Volumes production environment are logged and any anomalies to the baseline or suspicious activities are detected by our security information and event management (SIEM) threat intelligence platform. As a result, malicious actions can be tracked and mitigated before too much damage is done to the Google Cloud NetApp Volumes backend.

Volume life cycle

Google Cloud NetApp Volumes manages only the objects within the service—not the data within the volumes. Only clients accessing the volumes can manage the data, the ACLs, file owners, and so on. The data in these volumes is encrypted at rest and access is limited to tenants of the Google Cloud NetApp Volumes instance.

The volume lifecycle for Google Cloud NetApp Volumes is create-update-delete. Volumes retain Snapshot copies of volumes until the volumes are deleted, and only validated Google Cloud NetApp Volumes administrators can delete volumes in Google Cloud NetApp Volumes. When a volume deletion is requested by an administrator, an additional step of entering the volume name is required to verify the deletion. After a volume is deleted, the volume is gone and cannot be recovered.

In cases where a Google Cloud NetApp Volumes contract is terminated, NetApp marks volumes for deletion after a specific time period. Before that time period expires, you can recover volumes at the customer's request.

Certifications

Google Cloud NetApp Volumes for Google Cloud is currently certified to ISO/IEC 27001:2013 and ISO/IEC 27018:2019 standards. The service also recently received its SOC2 Type I attestation report. For information about the NetApp commitment to data security and privacy, see Compliance: Data security and data privacy.

GDPR

Our commitments to privacy and compliance with GDPR are available in a number of our customer contracts, such as our Customer Data Processing Addendum, which includes the Standard Contractual Clauses provided by the European Commission. We also make these commitments in our Privacy Policy, backed by the core values set out in our corporate Code of Conduct.

Additional information and contact information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Google Cloud documentation for Google Cloud NetApp Volumes

https://docs.cloud.google.com/netapp/volumes/docs/discover/overview

- Google private service access

  https://cloud.google.com/vpc/docs/private-services-access?hl=en_US

- NetApp product documentation

  https://www.netapp.com/support-and-training/documentation/

- Cryptographic Validation Module Program—NetApp CryptoMod

  https://csrc.nist.gov/projects/cryptographic-module-validation-program/certificate/4144

- The NetApp Solution for Ransomware

  https://www.netapp.com/pdf.html?item=/media/16716-sb-3938pdf.pdf&v=202093745

- TR-4616: NFS Kerberos in ONTAP

  https://www.netapp.com/pdf.html?item=/media/19384-tr-4616.pdf

Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com. Include TECHNICAL REPORT 4918 in the subject line.

NetApp Disaster Recovery

3-2-1 Data Protection for VMware with SnapCenter Plug-in and NetApp Backup and Recovery to protect VMware workloads

The 3-2-1 backup strategy is an industry accepted data protection method, providing a comprehensive approach to safeguarding valuable data. This strategy is reliable and ensures that even if some unexpected disaster strikes, there will still be a copy of the data available.

Overview

The strategy is comprised of three fundamental rules:

1. Keep at least three copies of your data. This ensures that even if one copy is lost or corrupted, you still have at least two remaining copies to fall back on.
2. Store two backup copies on different storage media or devices. Diversifying storage media helps protect against device-specific or media-specific failures. If one device gets damaged or one type of media fails, the other backup copy remains unaffected.
3. Finally, ensure that at least one backup copy is offsite. Offsite storage serves as a fail-safe against localized disasters like fires or floods that could render onsite copies unusable.

This solution document covers a 3-2-1 backups solution using SnapCenter Plug-in for VMware vSphere (SCV) to create primary and secondary backups of our on-premises virtual machines and NetApp Backup and

Recovery for virtual machines to backup a copy of our data to cloud storage or StorageGRID.

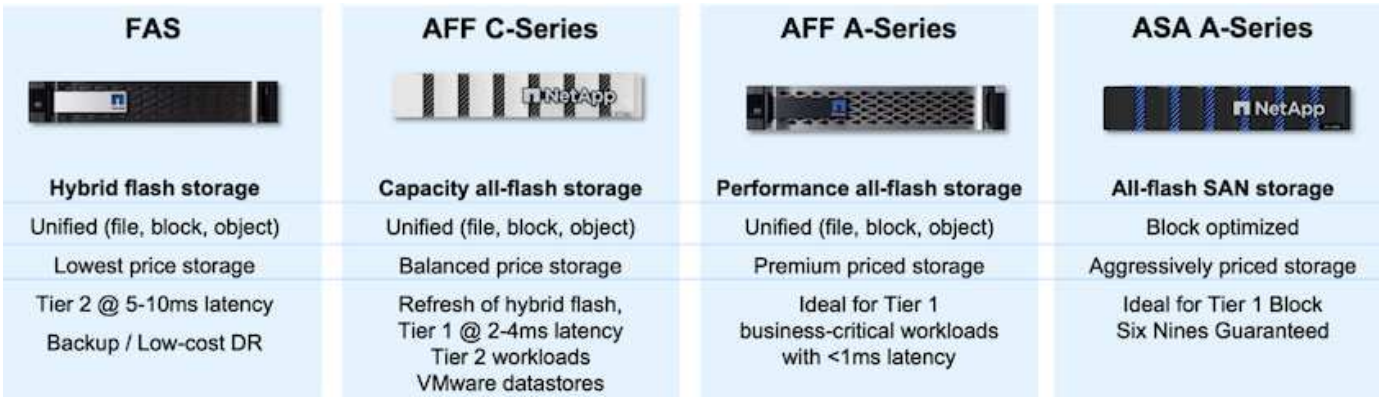
Use Cases

This solution addresses the following use cases:

- Backup and restore of on-premises virtual machines and datastores using using SnapCenter Plug-in for VMware vSphere.
- Backup and restore of on-premises virtual machines and datastores, hosted on ONTAP clusters, and backed up to object storage using NetApp Backup and Recovery for virtual machines.

NetApp ONTAP Data Storage

ONTAP is NetApp's industry leading storage solution that offers unified storage whether you access over SAN or NAS protocols. The 3-2-1 backup strategy ensures on-premises data is protected on more than one media type and NetApp offers platforms ranging from high-speed flash to lower-cost media.



For more information on all of NetApp's hardware platform's check out NetApp Data Storage.

SnapCenter Plug-in for VMware vSphere

The SnapCenter Plugin for VMware vSphere is a data protection offering which is tightly integrated with VMware vSphere and allows easy management of backup and restores for virtual machines. As part of that solution, SnapMirror provides a fast and reliable method to create a second immutable backup copy of virtual machine data on a secondary ONTAP storage cluster. With this architecture in place, virtual machine restore operations can easily be initiated from either the primary or secondary backup locations.

SCV is deployed as a linux virtual appliance using an OVA file. The plug-in now uses a remote plug-in architecture. The remote plug-in runs outside of the vCenter server and is hosted on the SCV virtual appliance.

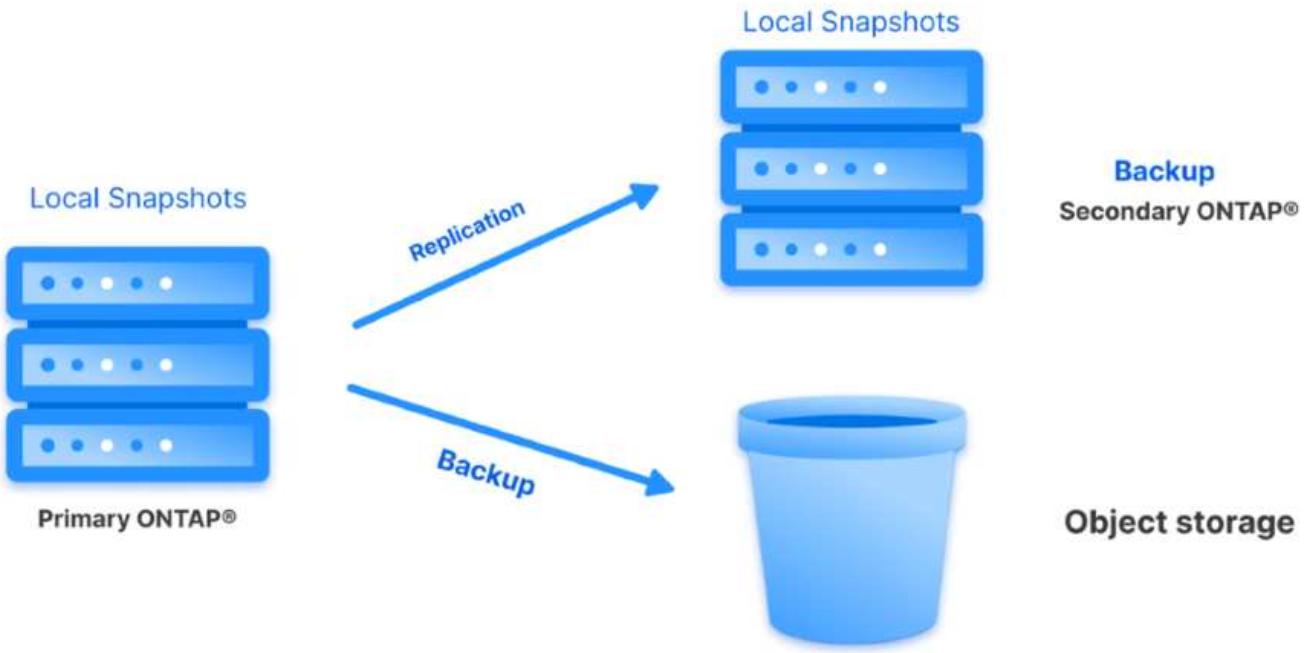
For detailed information on SCV refer to SnapCenter Plug-in for VMware vSphere documentation.

NetApp Backup and Recovery to protect VMware Workloads

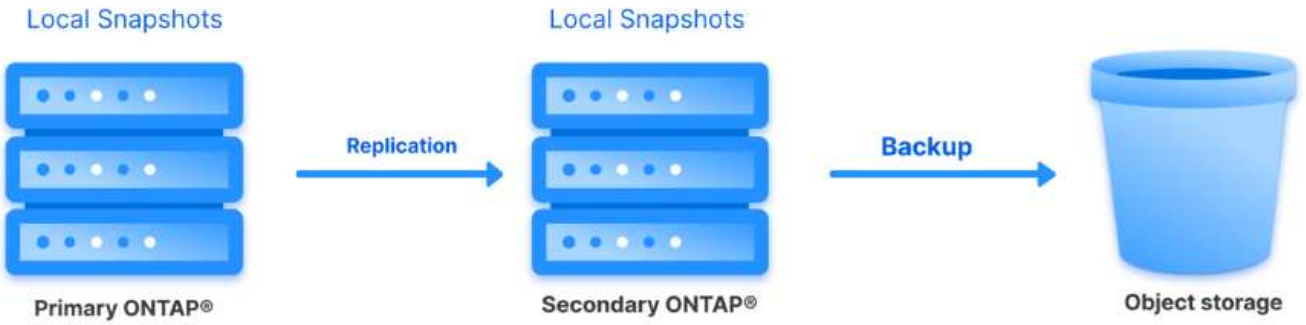
NetApp Backup and Recovery is a cloud based tool for data management that provides a single control plane for a wide range of backup and recovery operations across both on-premises and cloud environments. Part of the NetApp Backup and Recovery for VM suite is a feature that integrates with the SnapCenter Plugin for VMware vSphere (on-premises) to extend a copy of the data to object storage in the cloud. This establishes a third copy of the data offsite that is sourced from the primary or secondary storage backups. NetApp Backup and Recovery for VMs makes it easy to set up storage policies that transfer copies of your data from either of these two on-prem locations.

Choosing between the primary and secondary backups as the source in NetApp Backup and Recovery will result in one of two topologies being implemented:

Fan-out Topology – When a backup is initiated by the SnapCenter Plug-in for VMware vSphere, a local snapshot is immediately taken. SCV then initiates a SnapMirror operation that replicates the most recent snapshot to the Secondary ONTAP cluster. In NetApp Backup and Recovery, a policy specifies the primary ONTAP cluster as the source for a snapshot copy of the data to be transferred to object storage in your cloud provider of choice.



Cascading Topology – Creating the primary and secondary data copies using SCV is identical to the fan-out topology mentioned above. However, this time a policy is created in NetApp Backup and Recovery specifying that the backup to object storage will originate from the secondary ONTAP cluster.



NetApp Backup and Recovery can create backup copies of on-premises ONTAP snapshots to AWS Glacier, Azure Blob, and GCP Archive storage.



AWS Glacier and Deep Glacier



Azure Blob Archive



GCP Archive Storage

In addition, you can use NetApp StorageGRID as the object storage backup target. For more on StorageGRID refer to the StorageGRID landing page.

▼ Solution Deployment Overview

This list provides the high level steps necessary to configure this solution and execute backup & restore operations from SCV and NetApp Console:

1. Configure SnapMirror relationship between the ONTAP clusters to be used for primary and secondary data copies.
2. Configure SnapCenter Plug-In for VMware vSphere.
   a. Add Storage Systems
   b. Create backup policies
   c. Create resource groups
   d. Run backup first backup jobs
3. Configure NetApp Backup and Recovery for virtual machines
   a. Add working environment
   b. Discover SCV and vCenter appliances
   c. Create backup policies
   d. Activate backups
4. Restore virtual machines from primary and secondary storage using SCV.
5. Restore virtual machines from object storage using NetApp Console.

▼ Prerequisites

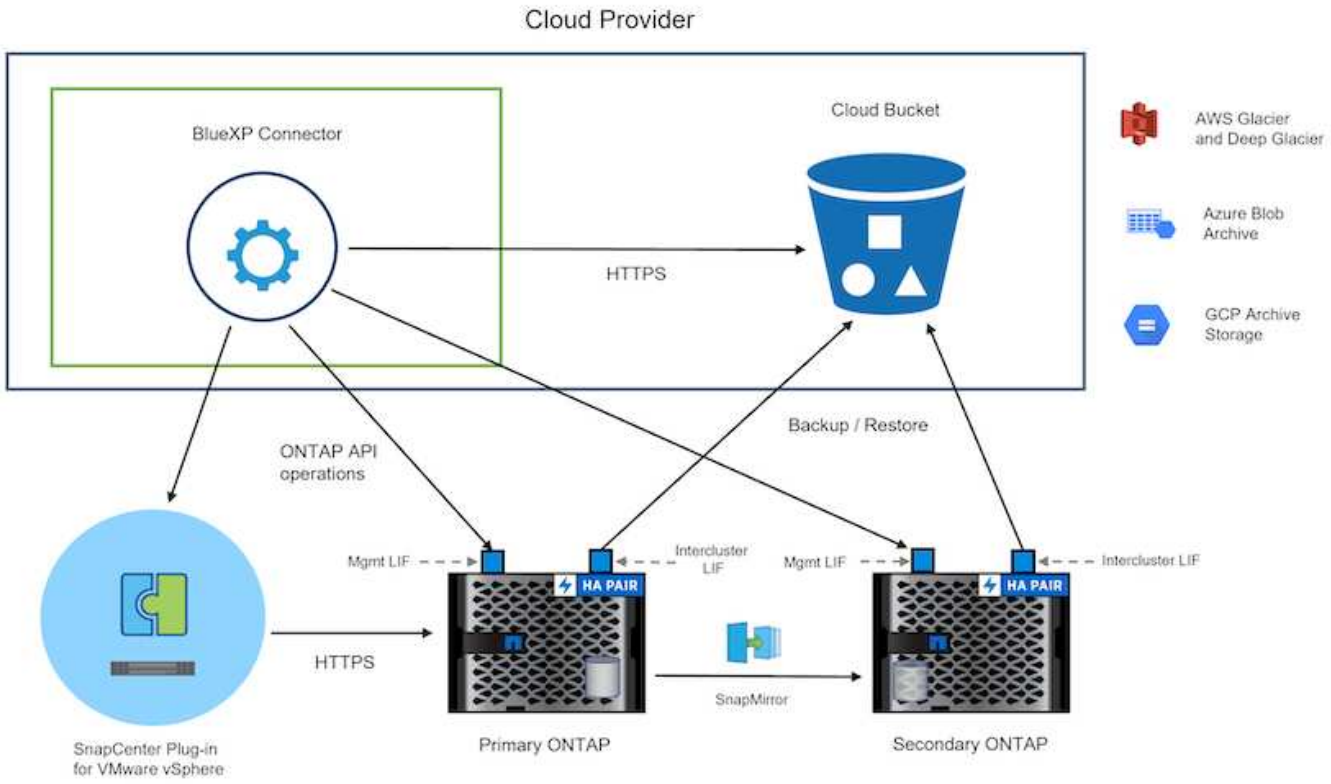
The purpose of this solution is to demonstrate data protection of virtual machines running in VMware vSphere and located on NFS Datastores hosted by NetApp ONTAP. This solution assumes the following components are configured and ready for use:

1. ONTAP storage cluster with NFS or VMFS datastores connected to VMware vSphere. Both NFS and VMFS datastores are supported. NFS datastores were utilized for this solution.
2. Secondary ONTAP storage cluster with SnapMirror relationships established for volumes used for NFS datastores.
3. Console agent installed for cloud provider used for object storage backups.

4. Virtual machines to be backed are on NFS datastores residing on the primary ONTAP storage cluster.
5. Network connectivity between the Console agent and on-premises ONTAP storage cluster management interfaces.
6. Network connectivity between the Console agent and on-premises SCV appliance VM and between the Console agent and vCenter.
7. Network connectivity between the on-premises ONTAP intercluster LIFs and the object storage service.
8. DNS configured for management SVM on primary and secondary ONTAP storage clusters. For more information refer to Configure DNS for host-name resolution.

High Level Architecture

The testing / validation of this solution was performed in a lab that may or may not match the final deployment environment.



Solution Deployment

In this solution, we provide detailed instructions for deploying and validating a solution that utilizes SnapCenter Plug-in for VMware vSphere, along with Backup and Recovery, to perform the backup and recovery of Windows and Linux virtual machines within a VMware vSphere cluster located in an on-premises data center. The virtual machines in this setup are stored on NFS datastores hosted by an ONTAP A300 storage cluster. Additionally, a separate ONTAP A300 storage cluster serves as a secondary destination for volumes replicated using SnapMirror. Furthermore, object storage hosted on Amazon Web Services and Azure Blob were employed as targets for a third copy of the data.

We will go over creating SnapMirror relationships for secondary copies of our backups managed by SCV and configuration of backup jobs in both SCV and NetApp Console.

For detailed information on SnapCenter Plug-in for VMware vSphere refer to the SnapCenter Plug-in for

VMware vSphere documentation.

For detailed information on Backup and Recovery refer to the Backup and Recovery documentation.

Establish SnapMirror relationships between ONTAP Clusters

SnapCenter Plug-in for VMware vSphere uses ONTAP SnapMirror technology to manage the transport of secondary SnapMirror and/or SnapVault copies to a secondary ONTAP Cluster.

SCV backup policies have the option of using SnapMirror or SnapVault relationships. The primary difference is that when using the SnapMirror option, the retention schedule configured for backups in the policy will be the same at the primary and secondary locations. SnapVault is designed for archiving and when using this option a separate retention schedule can be established with the SnapMirror relationship for the snapshot copies on the secondary ONTAP storage cluster.

Setting up SnapMirror relationships can be done in NetApp Console where many of the steps are automated, or it can be done using System Manager and the ONTAP CLI. All of these methods are discussed below.

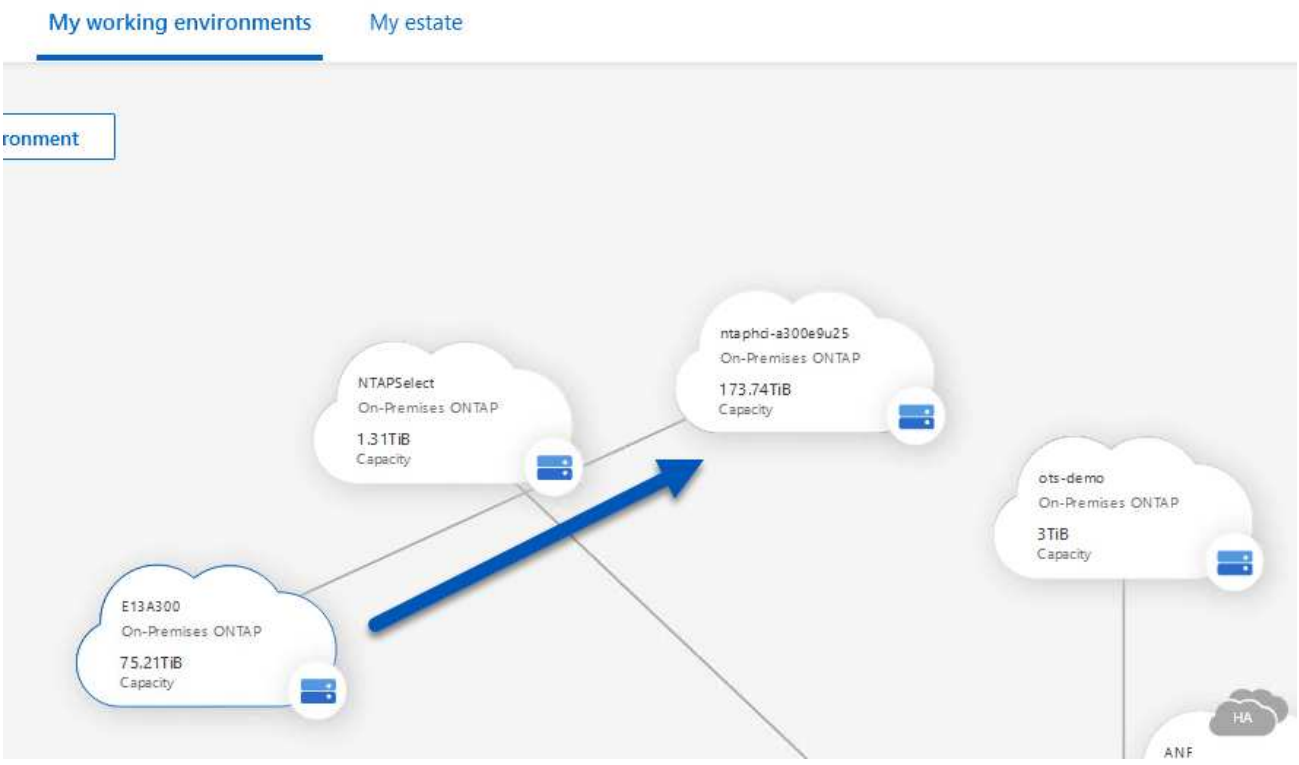
Establish SnapMirror relationships with NetApp Console

The following steps must be completed from the NetApp Console:

▼ **Replication setup for primary and secondary ONTAP storage systems**

Begin by logging into the NetApp Console and navigating to the Systems.

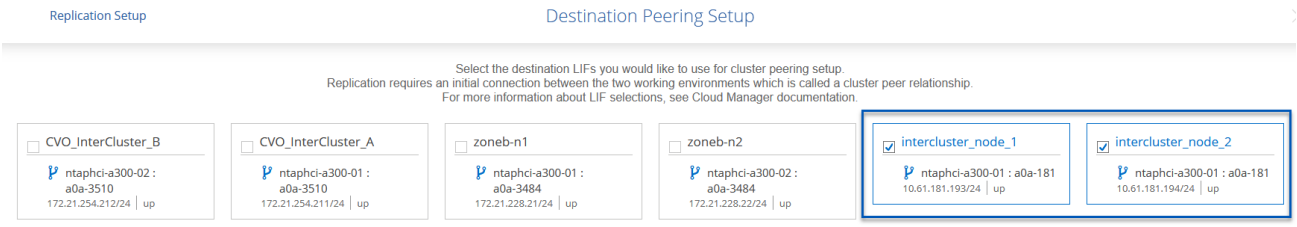
1. Drag and drop the source (primary) ONTAP storage system onto the destination (secondary) ONTAP storage system.



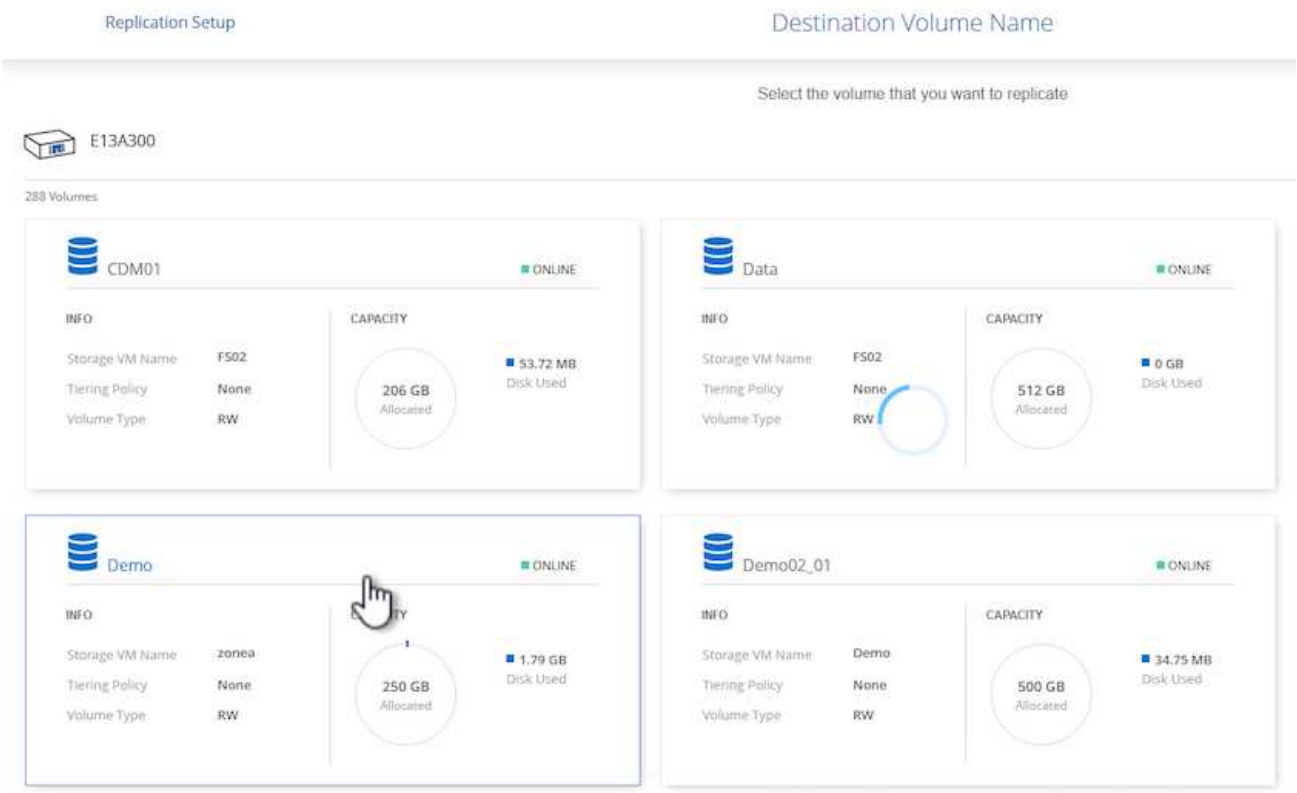
2. From the menu that appears select **Replication**.



3. On the **Destination Peering Setup** page select the destination Intercluster LIFs to be used for the connection between storage systems.

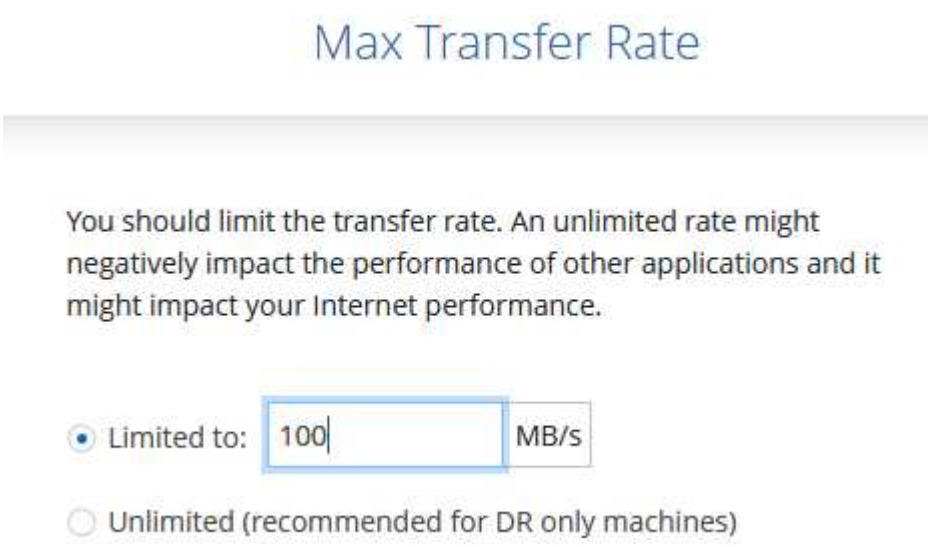


4. On the **Destination Volume Name** page, first select the source volume and then fill out the destination volume name and select the destination SVM and aggregate. Click on **Next** to continue.

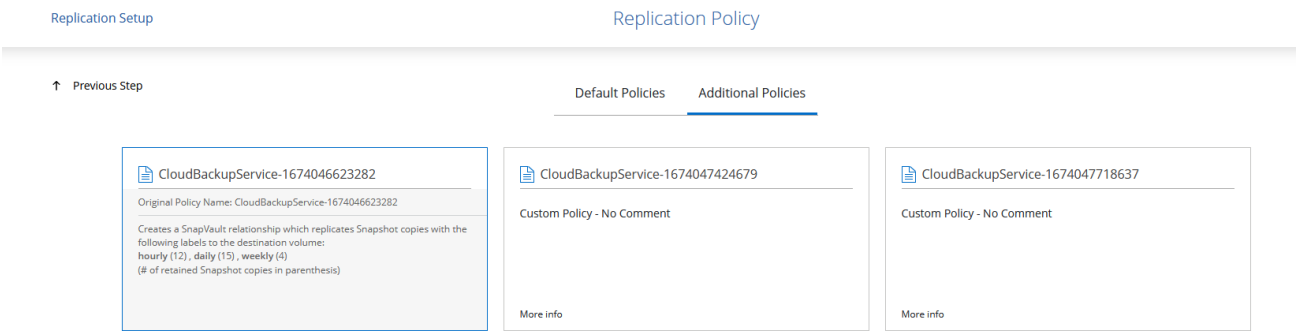




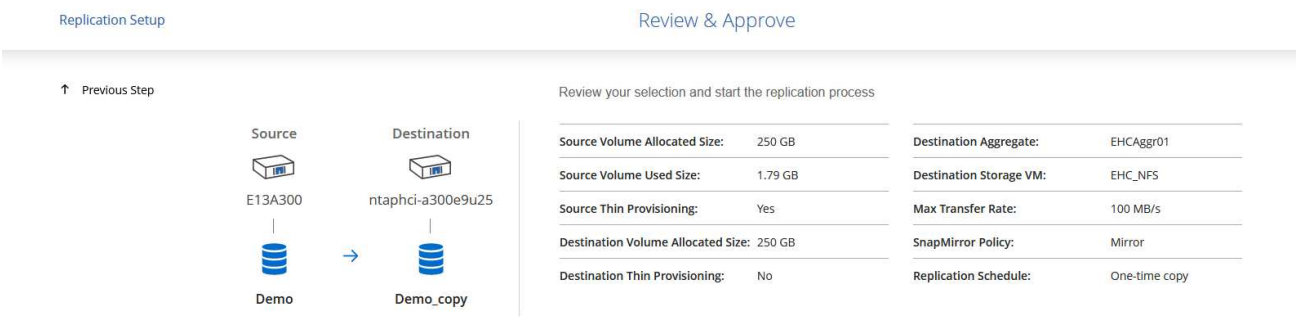
5. Choose the max transfer rate for replication to occur at.



6. Choose the policy that will determine the retention schedule for secondary backups. This policy can be created beforehand (see the manual process below in the **Create a snapshot retention policy** step) or can be changed after the fact if desired.



7. Finally, review all information and click on the **Go** button to start the replication setup process.



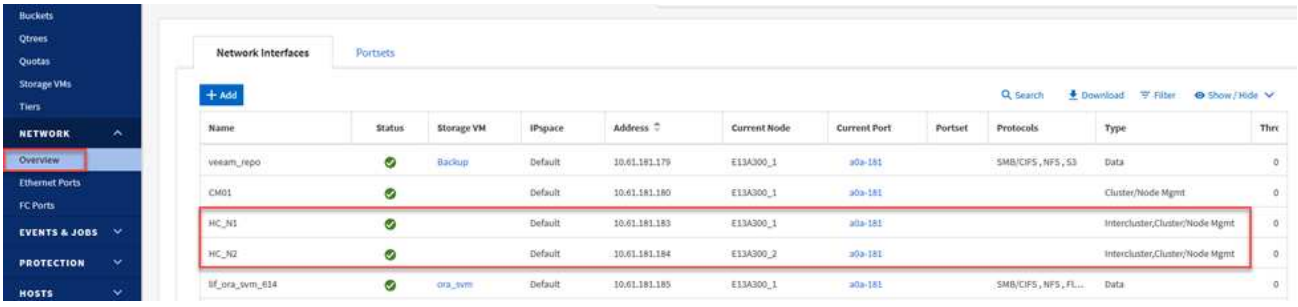
Establish SnapMirror relationships with System Manager and ONTAP CLI

All required steps for establishing SnapMirror relationships can be accomplished with System Manager or the ONTAP CLI. The following section provides detailed information for both methods:

▼ Record the source and destination Intercluster logical interfaces

For the source and destination ONTAP clusters, you can retrieve the inter-cluster LIF information from System Manager or from the CLI.

1. In ONTAP System Manager, navigate to the Network Overview page and retrieve the IP addresses of Type: Intercluster that are configured to communicate with the AWS VPC where FSx is installed.



2. To retrieve the Intercluster IP addresses using the CLI run the following command:

```
ONTAP-Dest::> network interface show -role intercluster
```

▼ Establish cluster peering between ONTAP clusters

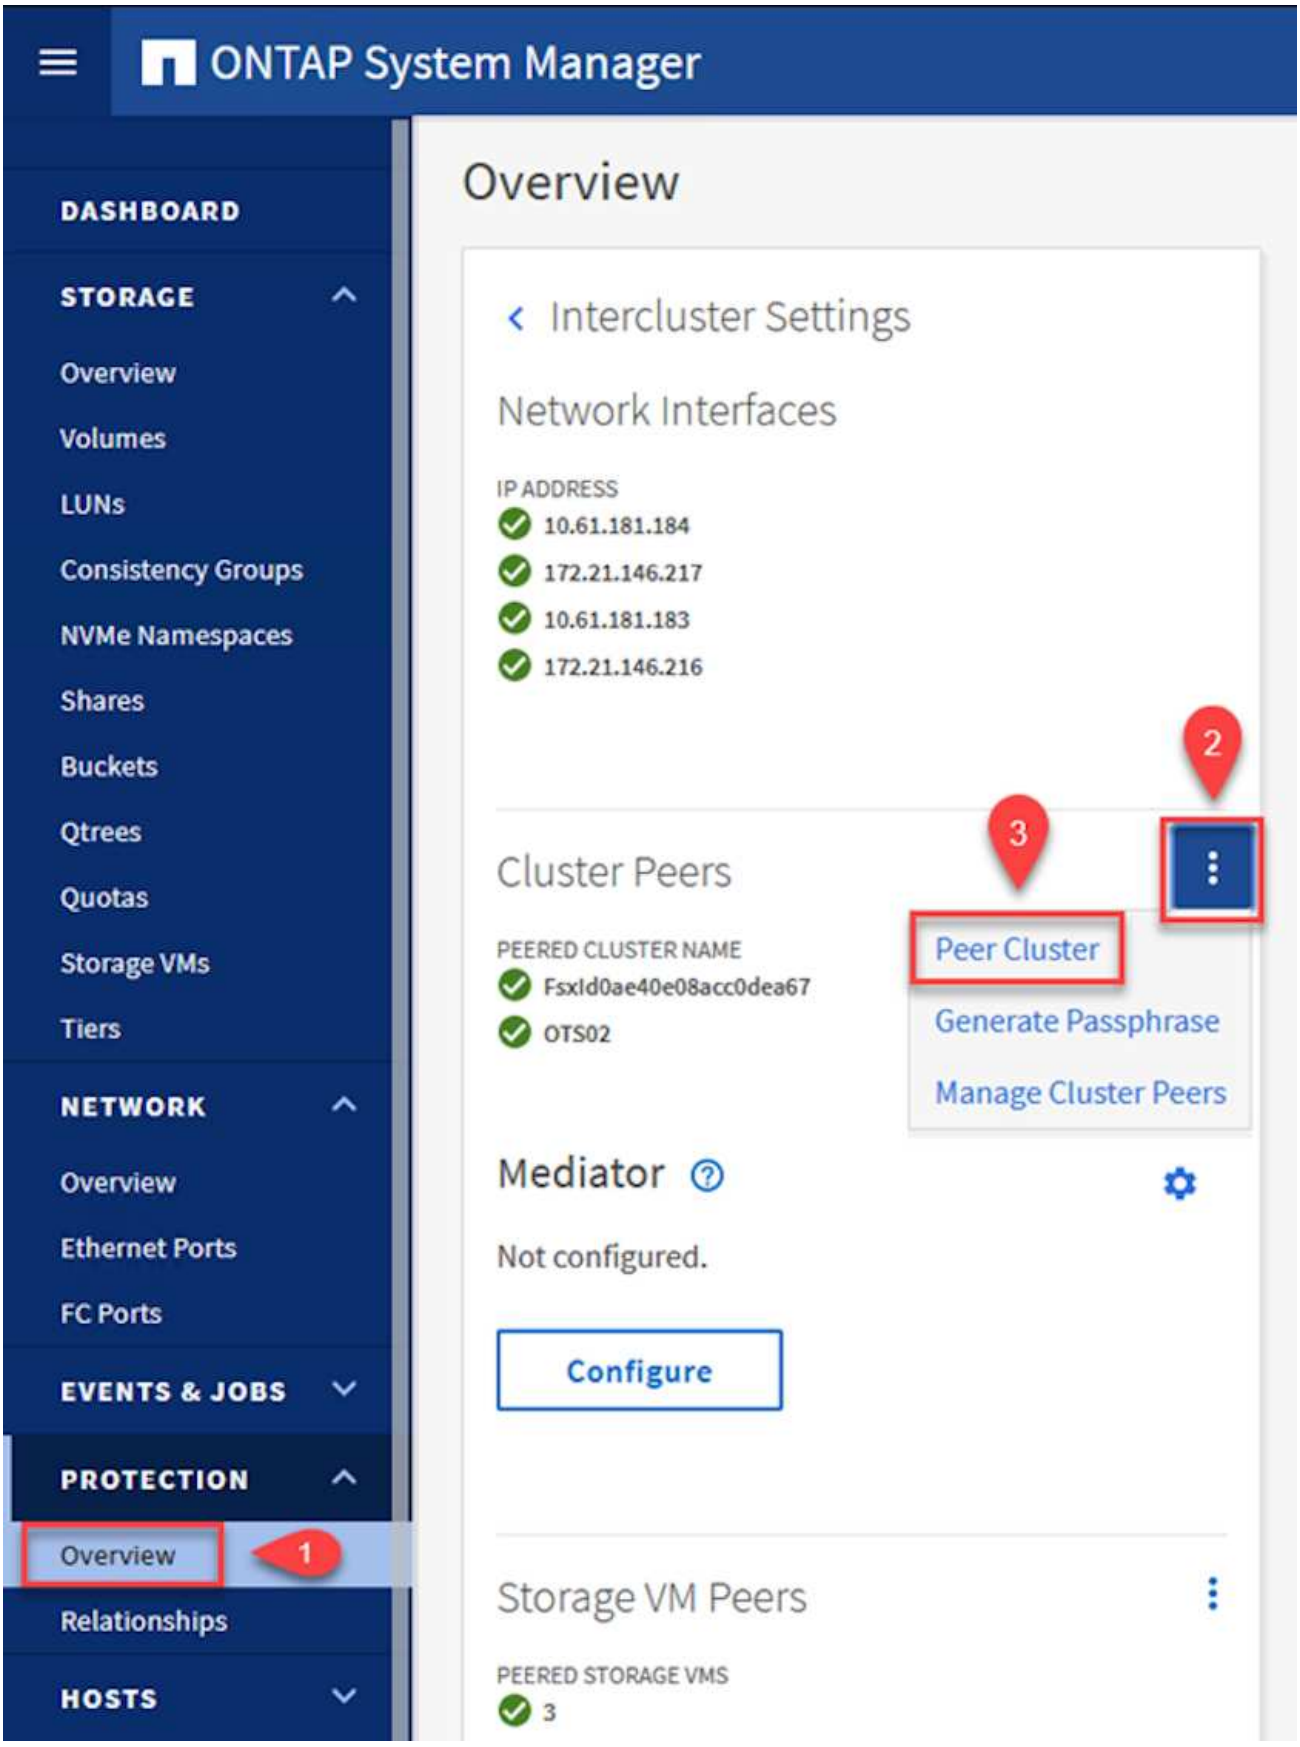
To establish cluster peering between ONTAP clusters, a unique passphrase entered at the initiating ONTAP

cluster must be confirmed in the other peer cluster.

1. Set up peering on the destination ONTAP cluster using the `cluster peer create` command. When prompted, enter a unique passphrase that is used later on the source cluster to finalize the creation process.

```
ONTAP-Dest::> cluster peer create -address-family ipv4 -peer-addrs
source_intercluster_1, source_intercluster_2
Enter the passphrase:
Confirm the passphrase:
```

2. At the source cluster, you can establish the cluster peer relationship using either ONTAP System Manager or the CLI. From ONTAP System Manager, navigate to Protection > Overview and select Peer Cluster.



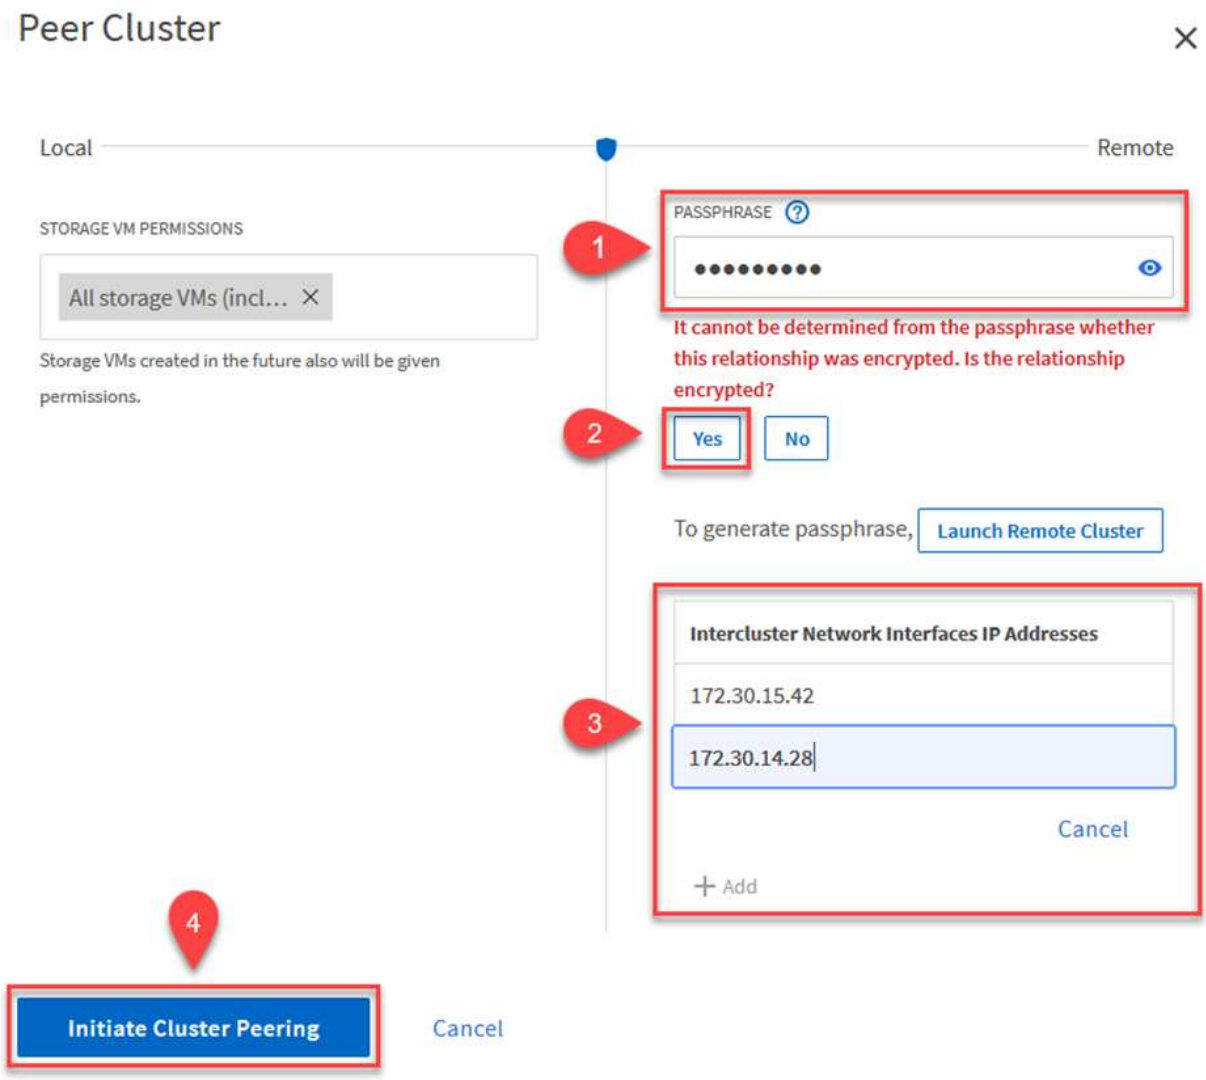
3. In the Peer Cluster dialog box, fill out the required information:
   a. Enter the passphrase that was used to establish the peer cluster relationship on the destination

ONTAP cluster.

b. Select `Yes` to establish an encrypted relationship.

c. Enter the intercluster LIF IP address(es) of the destination ONTAP cluster.

d. Click Initiate Cluster Peering to finalize the process.



4. Verify the status of the cluster peer relationship from the destination ONTAP cluster with the following command:

```
ONTAP-Dest::> cluster peer show
```

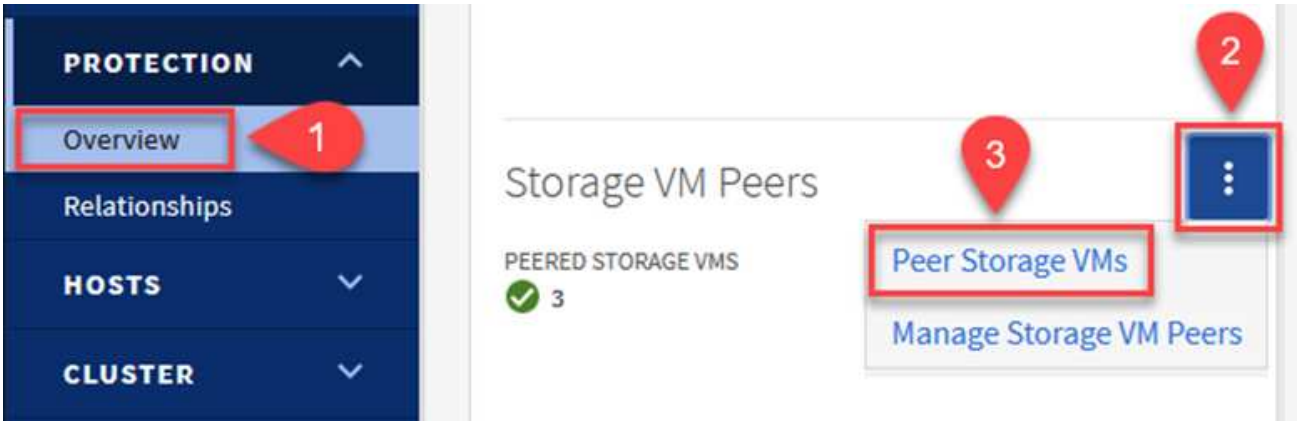
▼ **Establish SVM peering relationship**

The next step is to set up an SVM relationship between the destination and source storage virtual machines that contain the volumes that will be in SnapMirror relationships.

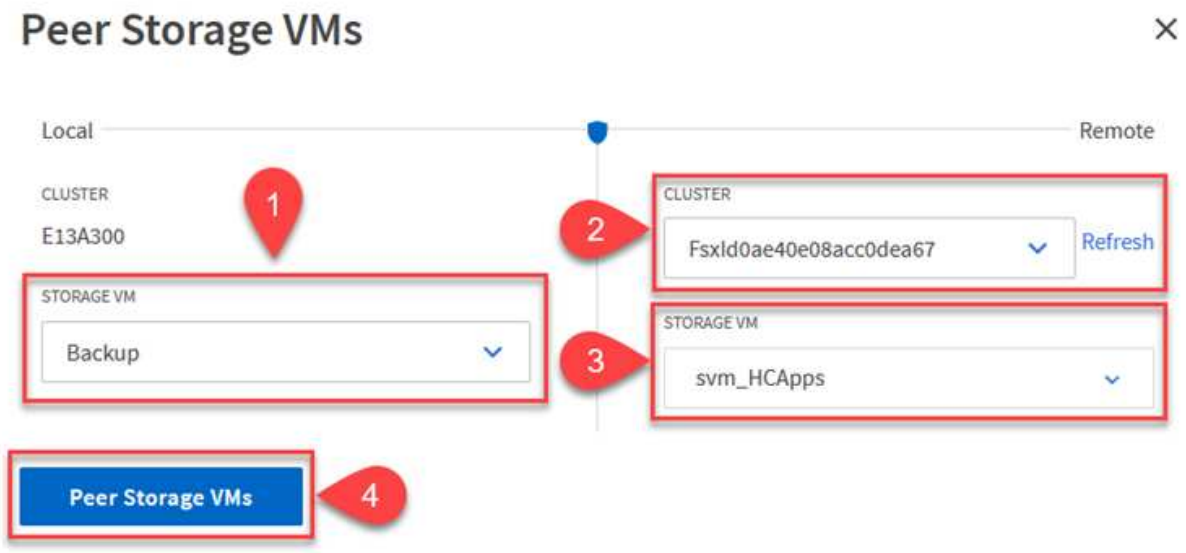
1. From the destination ONTAP cluster, use the following command from the CLI to create the SVM peer relationship:

```
ONTAP-Dest::> vserver peer create -vserver DestSVM -peer-vserver
Backup -peer-cluster OnPremSourceSVM -applications snapmirror
```

2. From the source ONTAP cluster, accept the peering relationship with either ONTAP System Manager or the CLI.
3. From ONTAP System Manager, go to Protection > Overview and select Peer Storage VMs under Storage VM Peers.



4. In the Peer Storage VM's dialog box, fill out the required fields:
   - The source storage VM
   - The destination cluster
   - The destination storage VM



5. Click Peer Storage VMs to complete the SVM peering process.

▼ **Create a snapshot retention policy**

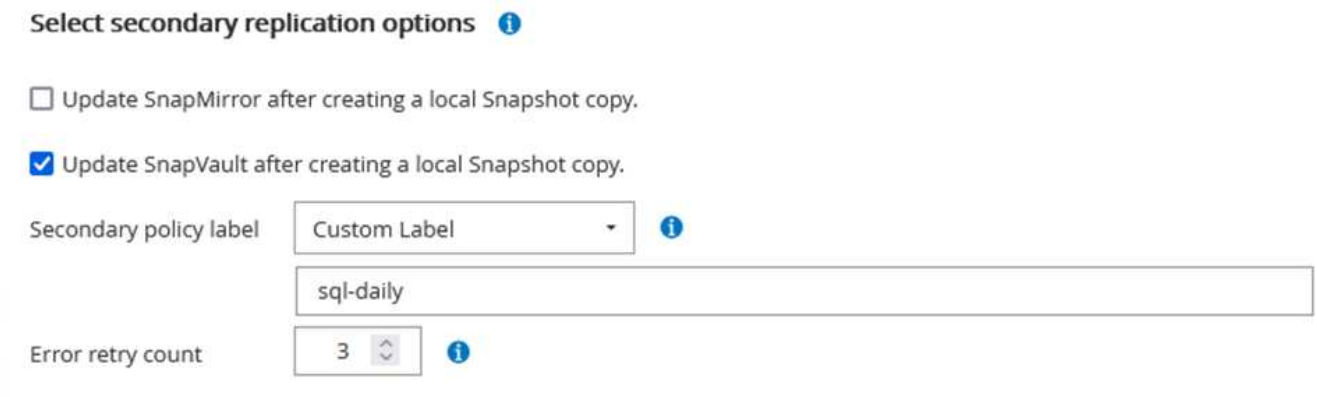
SnapCenter manages retention schedules for backups that exist as snapshot copies on the primary storage system. This is established when creating a policy in SnapCenter. SnapCenter does not manage retention policies for backups that are retained on secondary storage systems. These policies are managed

separately through a SnapMirror policy created on the secondary FSx cluster and associated with the destination volumes that are in a SnapMirror relationship with the source volume.

When creating a SnapCenter policy, you have the option to specify a secondary policy label that is added to the SnapMirror label of each snapshot generated when a SnapCenter backup is taken.

> On the secondary storage, these labels are matched to policy rules associated with the destination volume for the purpose of enforcing retention of snapshots.

The following example shows a SnapMirror label that is present on all snapshots generated as part of a policy used for daily backups of our SQL Server database and log volumes.



For more information on creating SnapCenter policies for a SQL Server database, see the SnapCenter documentation.

You must first create a SnapMirror policy with rules that dictate the number of snapshot copies to retain.

1. Create the SnapMirror Policy on the FSx cluster.

```
ONTAP-Dest::> snapmirror policy create -vserver DestSVM -policy
PolicyName -type mirror-vault -restart always
```

2. Add rules to the policy with SnapMirror labels that match the secondary policy labels specified in the SnapCenter policies.

```
ONTAP-Dest::> snapmirror policy add-rule -vserver DestSVM -policy
PolicyName -snapmirror-label SnapMirrorLabelName -keep
#ofSnapshotsToRetain
```

The following script provides an example of a rule that could be added to a policy:

```
ONTAP-Dest::> snapmirror policy add-rule -vserver sql_svm_dest -policy
Async_SnapCenter_SQL -snapmirror-label sql-ondemand -keep 15
```

Create additional rules for each SnapMirror label and the number of snapshots to be retained (retention period).

▼ **Create destination volumes**

To create a destination volume on ONTAP that will be the recipient of snapshot copies from our source volumes, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> volume create -vserver DestSVM -volume DestVolName
-aggregate DestAggrName -size VolSize -type DP
```

▼ **Create the SnapMirror relationships between source and destination volumes**

To create a SnapMirror relationship between a source and destination volume, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> snapmirror create -source-path
OnPremSourceSVM:OnPremSourceVol -destination-path DestSVM:DestVol -type
XDP -policy PolicyName
```

▼ **Initialize the SnapMirror relationships**

Initialize the SnapMirror relationship. This process initiates a new snapshot generated from the source volume and copies it to the destination volume.

To create a volume, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> snapmirror initialize -destination-path DestSVM:DestVol
```

Configure the SnapCenter Plug-in for VMware vSphere

Once installed, the SnapCenter Plug-in for VMware vSphere can be accessed from the vCenter Server Appliance Management interface. SCV will manage backups for the NFS datastores mounted to the ESXi hosts and that contain the Windows and Linux VMs.

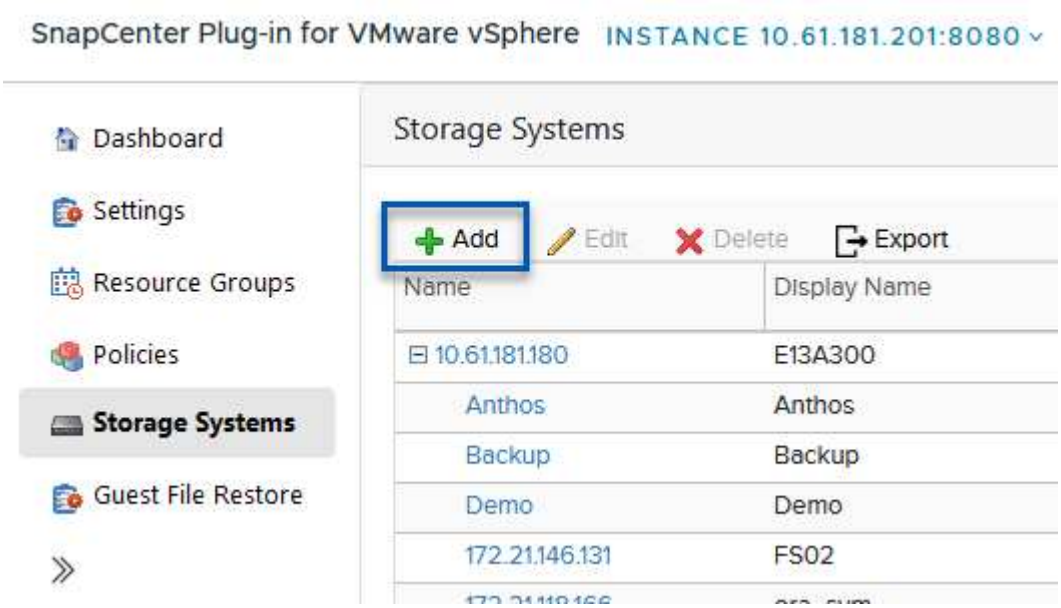
Review the Data protection workflow section of the SCV documentation for more information on the steps involved in configuring backups.

To configure backups of your virtual machines and datastores the following steps will need to be completed from the plug-in interface.

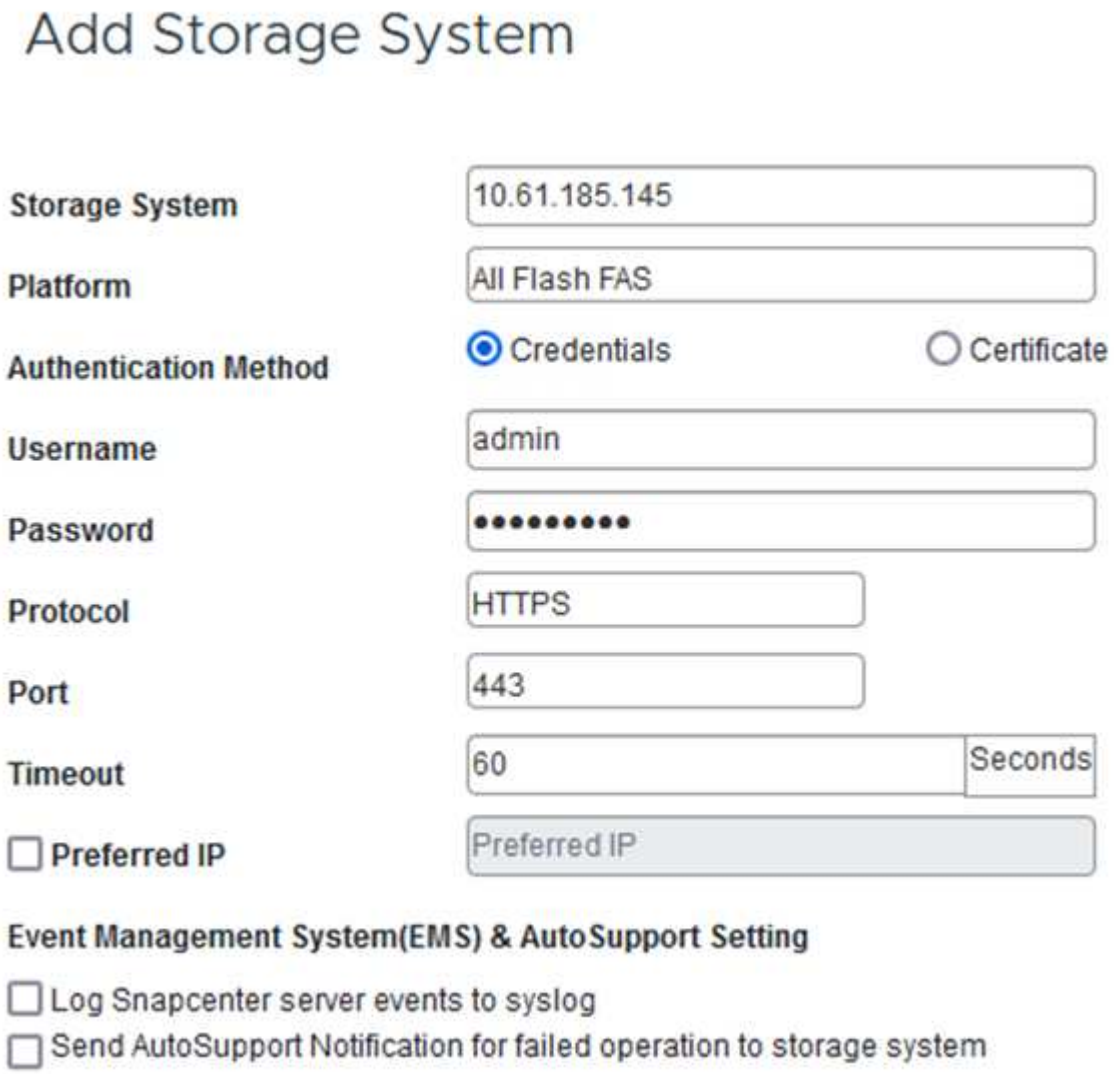
▼ **Discovery ONTAP storage systems**

Discover the ONTAP storage clusters to be used for both primary and secondary backups.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Storage Systems** in the left-hand menu and click on the **Add** button.



2. Fill out the credentials and platform type for the primary ONTAP storage system and click on **Add**.



3. Repeat this procedure for the secondary ONTAP storage system.

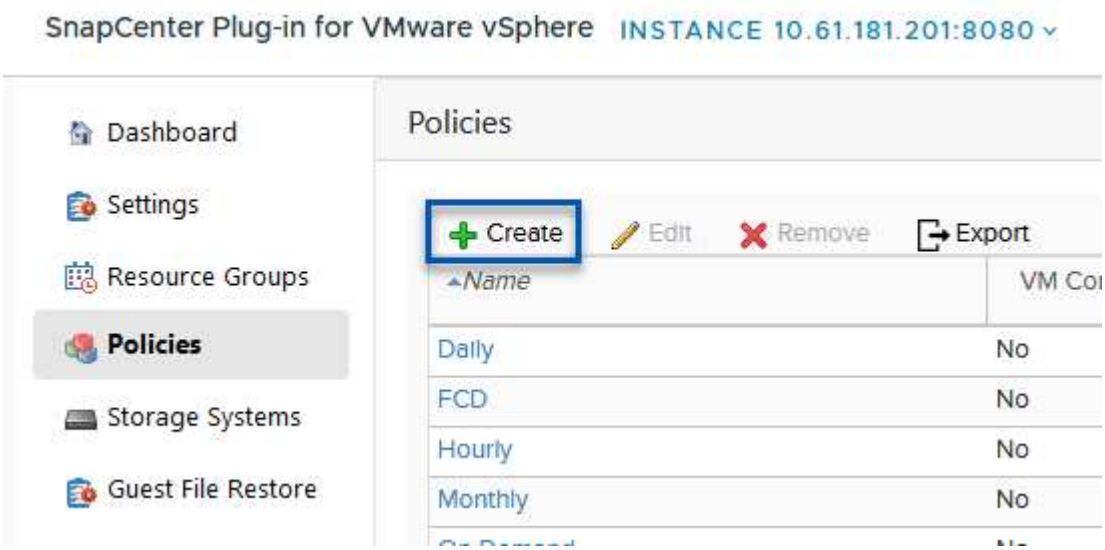
▼ **Create SCV backup policies**

Policies specify the retention period, frequency and replication options for the backups managed by SCV.

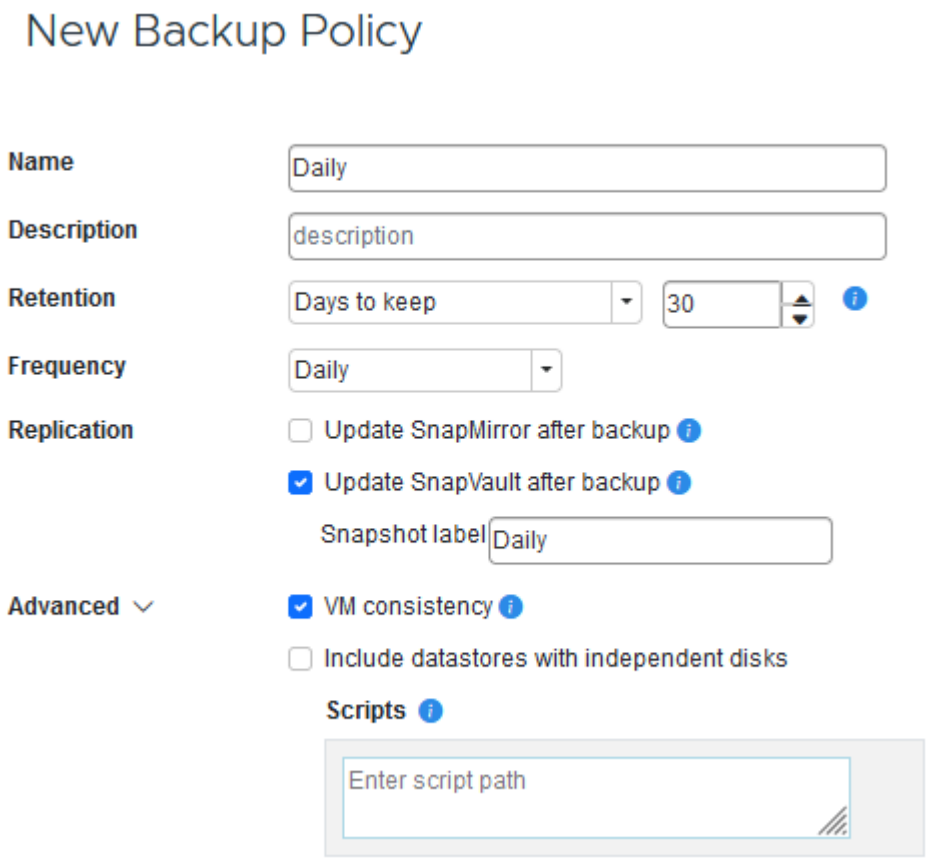
Review the Create backup policies for VMs and datastores section of the documentation for more information.

To create backup policies complete the following steps:

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Policies** in the left-hand menu and click on the **Create** button.



2. Specify a name for the policy, retention period, frequency and replication options, and snapshot label.



When creating a policy in the SnapCenter Plug-in you will see options for SnapMirror and SnapVault. If you choose SnapMirror, the retention schedule specified in the policy will be the same for both the primary and secondary snapshots. If you choose SnapVault, the retention schedule for the secondary snapshot will be based on a separate schedule implemented with the SnapMirror relationship. This is useful when you wish longer retention periods for secondary backups.

Snapshot labels are useful in that they can be used to enact policies with a specific retention period for the SnapVault copies replicated to the secondary ONTAP cluster. When SCV is used with NetApp Data Services, the Snapshot label field must either be blank or match the label specified in the backup policy.

3. Repeat the procedure for each policy required. For example, separate policies for daily, weekly, and monthly backups.

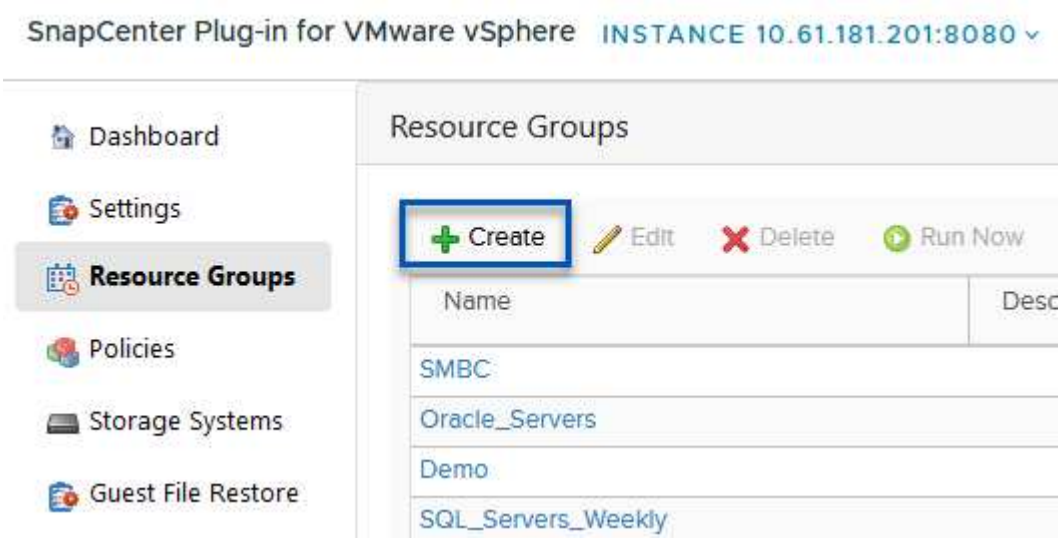
▼ **Create resource groups**

Resource groups contain the datastores and virtual machines to be included in a backup job, along with the associated policy and backup schedule.

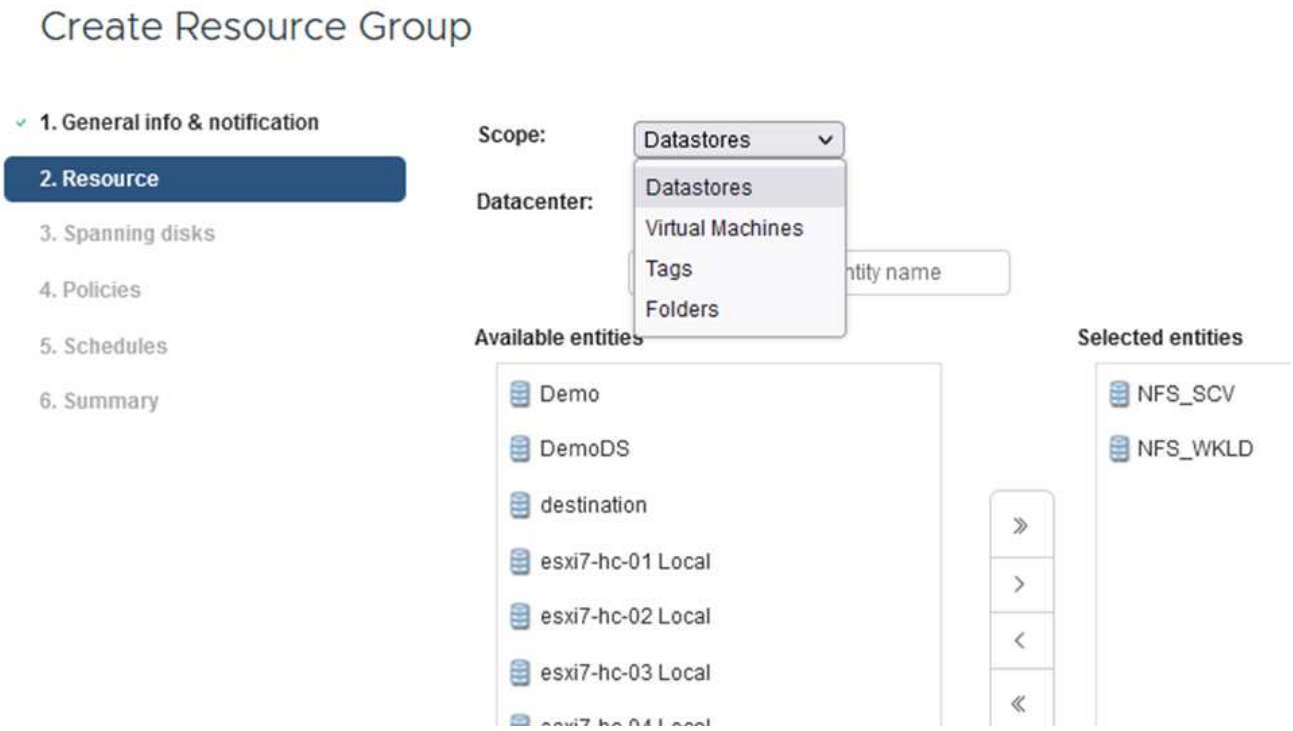
Review the Create resource groups section of the documentation for more information.

To create resource groups complete the following steps.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Resource Groups** in the left-hand menu and click on the **Create** button.

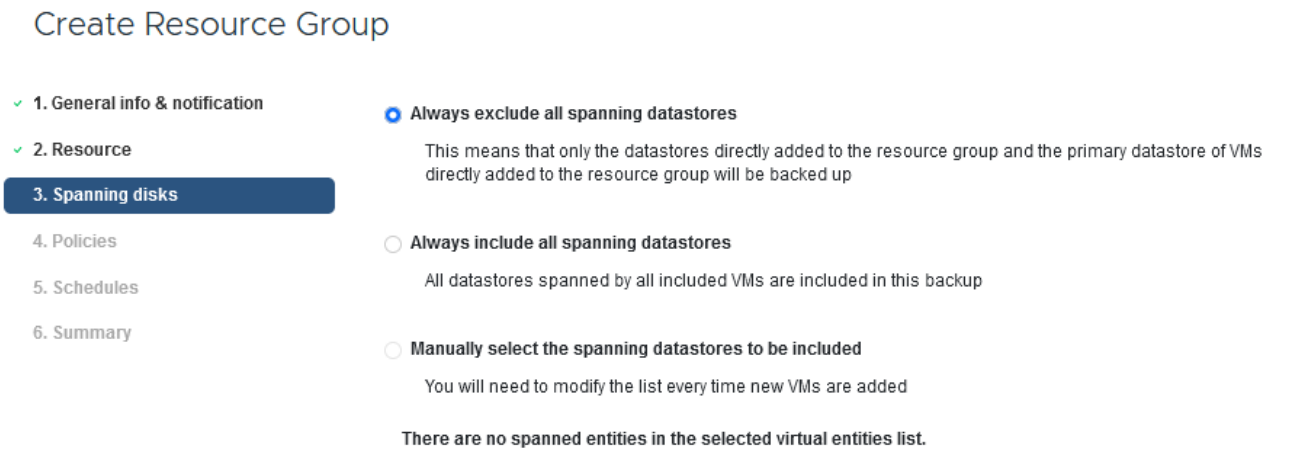


2. In the Create Resource Group wizard, enter a name and description for the group, as well as information required to receive notifications. Click on **Next**
3. On the next page select the datastores and virtual machines that wish to be included in the backup job and then click on **Next**.



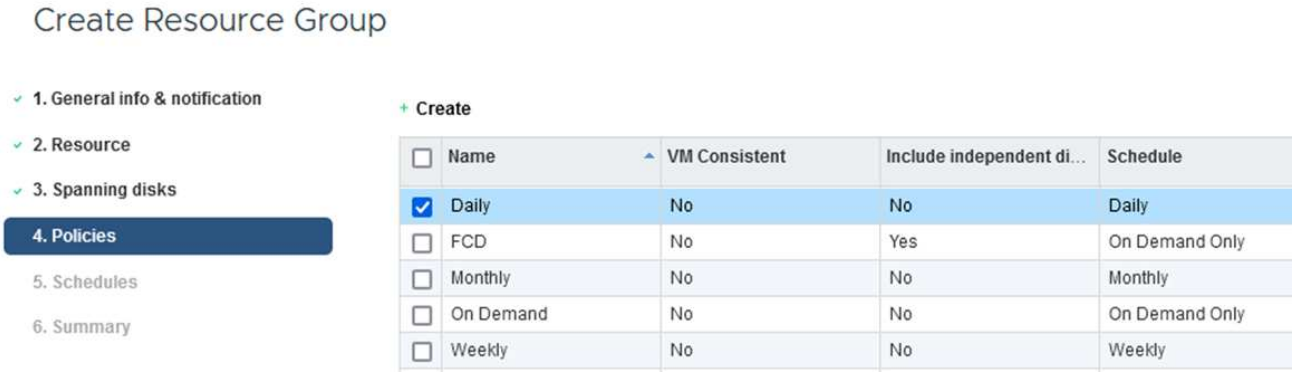
You have the option to select specific VMs or entire datastores. Regardless of which you choose, the entire volume (and datastore) is backed up since the backup is the result of taking a snapshot of the underlying volume. In most cases, it is easiest to choose the entire datastore. However, if you wish to limit the list of available VMs when restoring, you can choose only a subset of VMs for backup.

4. Choose options for spanning datastores for VMs with VMDKs that reside on multiple datastores and then click on **Next**.



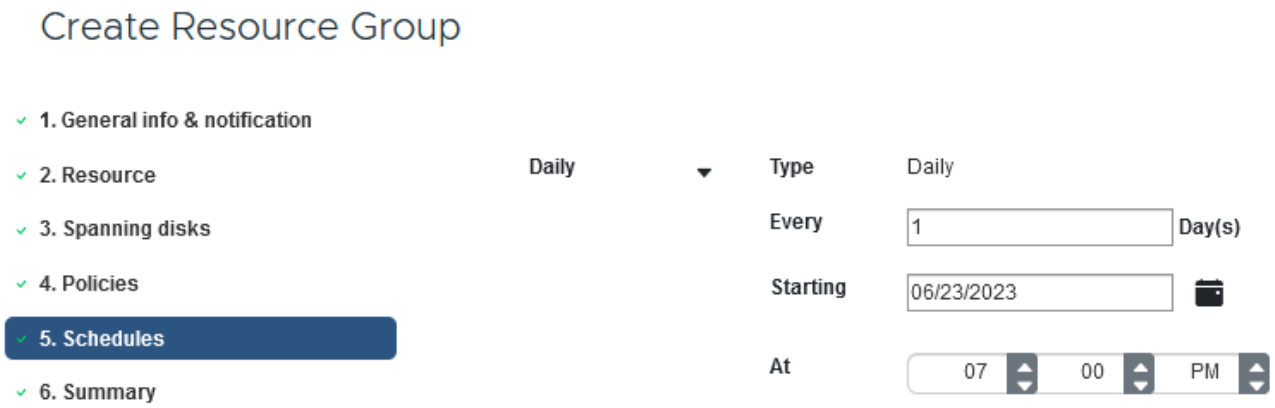
Backup and Recovery to protect VMware workloads does not currently support backing up VMs with VMDKs that span multiple datastores.

5. On the next page select the policies that will be associated with the resource group and click on **Next**.



When backing up SCV managed snapshots to object storage using Backup and Recovery, each resource group can only be associated with a single policy.

6. Select a schedule that will determine at what times the backups will run. Click on **Next**.

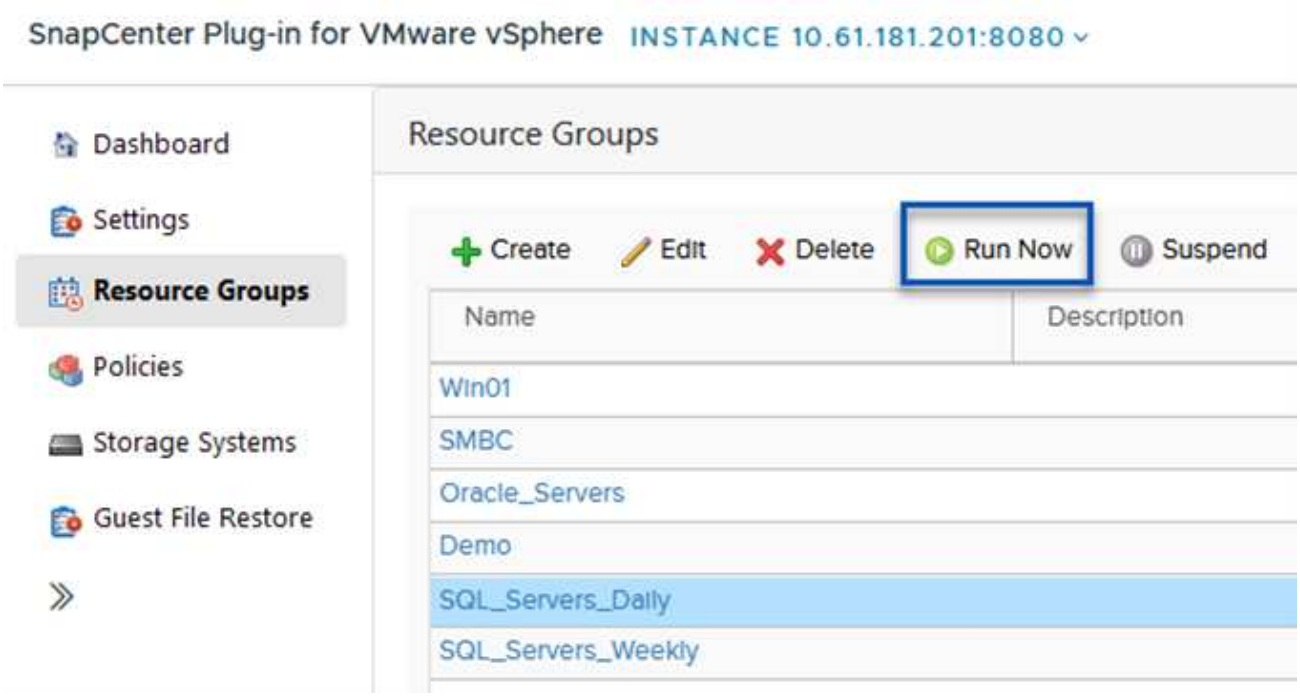


7. Finally, review the summary page and then on **Finish** to complete the resource group creation.

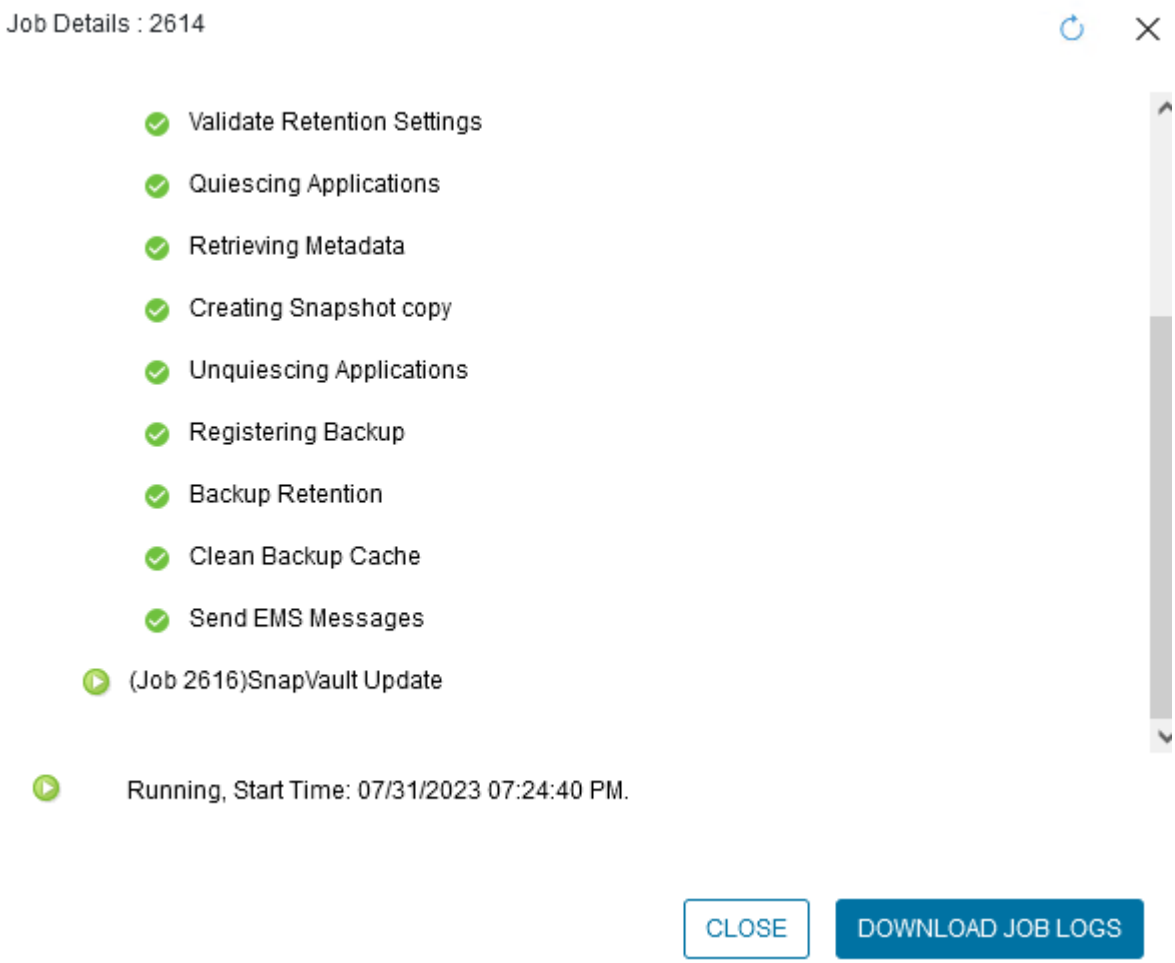
▼ **Run a backup job**

In this final step, run a backup job and monitor its progress. At least one backup job must be successfully completed in SCV before resources can be discovered from NetApp Console.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Resource Groups** in the left-hand menu.
2. To initiate a backup job, select the desired resource group and click the **Run Now** button.



3. To monitor the backup job, navigate to **Dashboard** on the left hand menu. Under **Recent Job Activities** click on the Job ID number to monitor the job progress.



Configure Backups to Object Storage in NetApp Console

For NetApp Console to manage the data infrastructure effectively, it requires the prior installation of a Console agent. The Console agent executes the actions involved in discovering resources and managing data operations.

For more information on the Console agent refer to Learn about Console agents in the NetApp Console documentation.

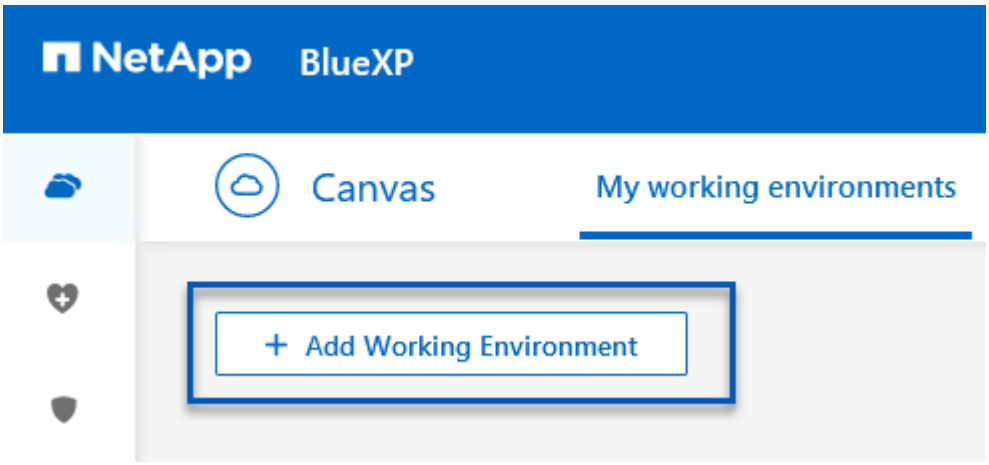
Once the agent is installed for the cloud provider being utilized, a graphic representation of the object storage will be viewable from the Systems.

To configure Backup and Recovery to backup data managed by SCV on-premises, complete the following steps:

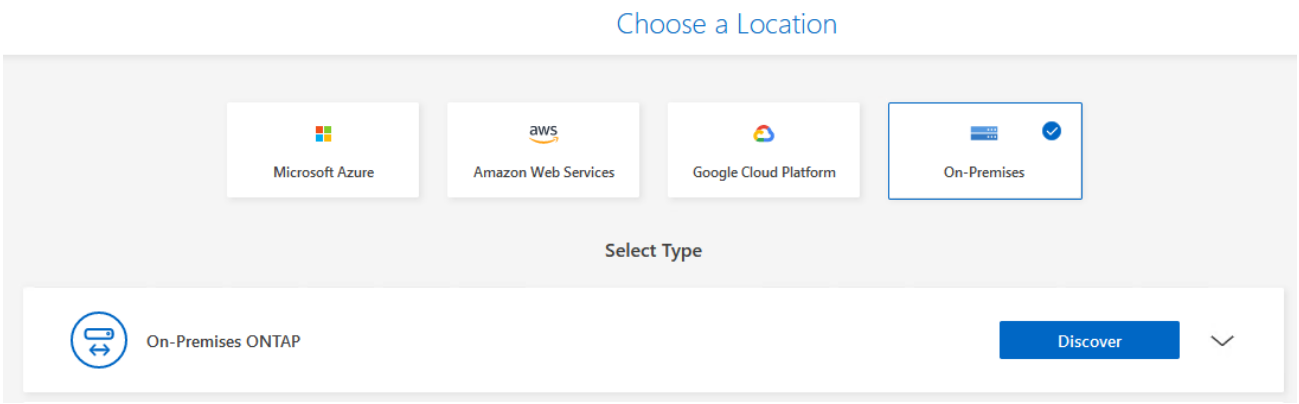
▼ **Add ONTAP storage system to Systems**

The first step is to add the on-premises ONTAP storage systems to NetApp Console.

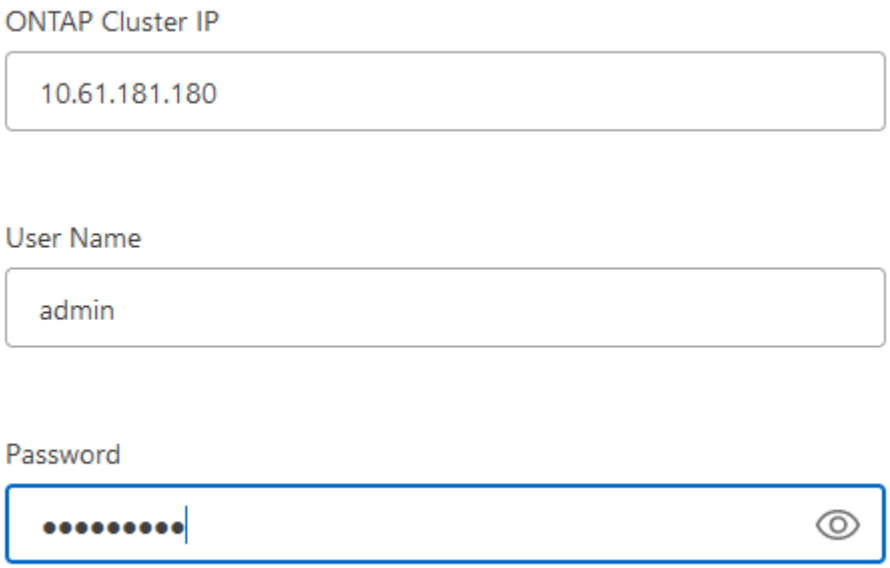
1. From the Systems Click **+ Add** to begin.



2. Select **On-Premises** from the choice of locations and then click on the **Discover** button.



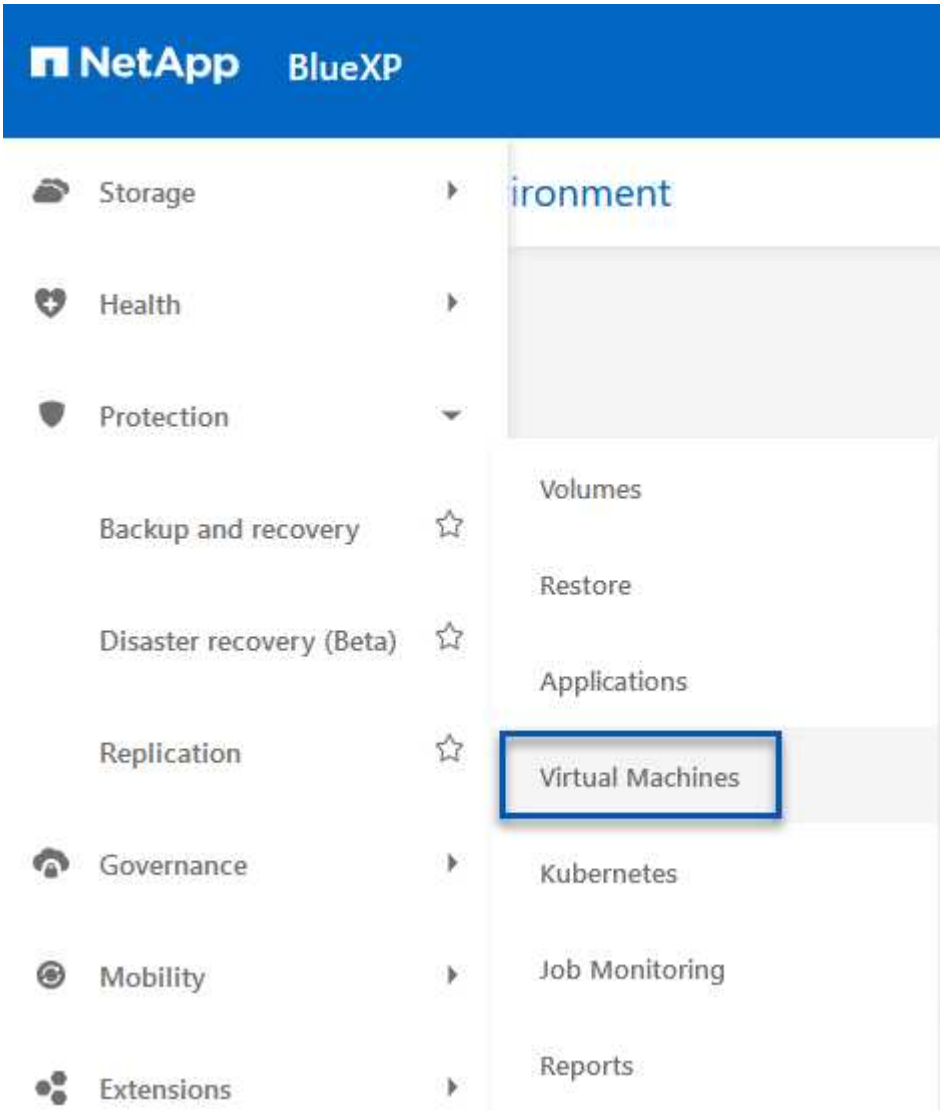
3. Fill out the credentials for the ONTAP storage system and click the **Discover** button to add the working environment.



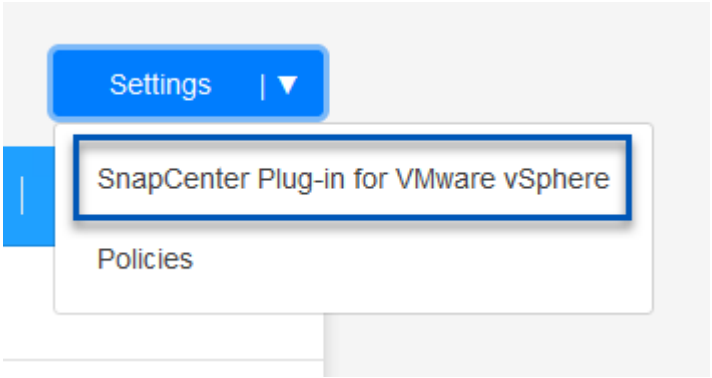
▼ **Discover on-premises SCV appliance and vCenter**

To discover the on-premises datastore and virtual machine resources, add info for the SCV data broker and credentials for the vCenter management appliance.

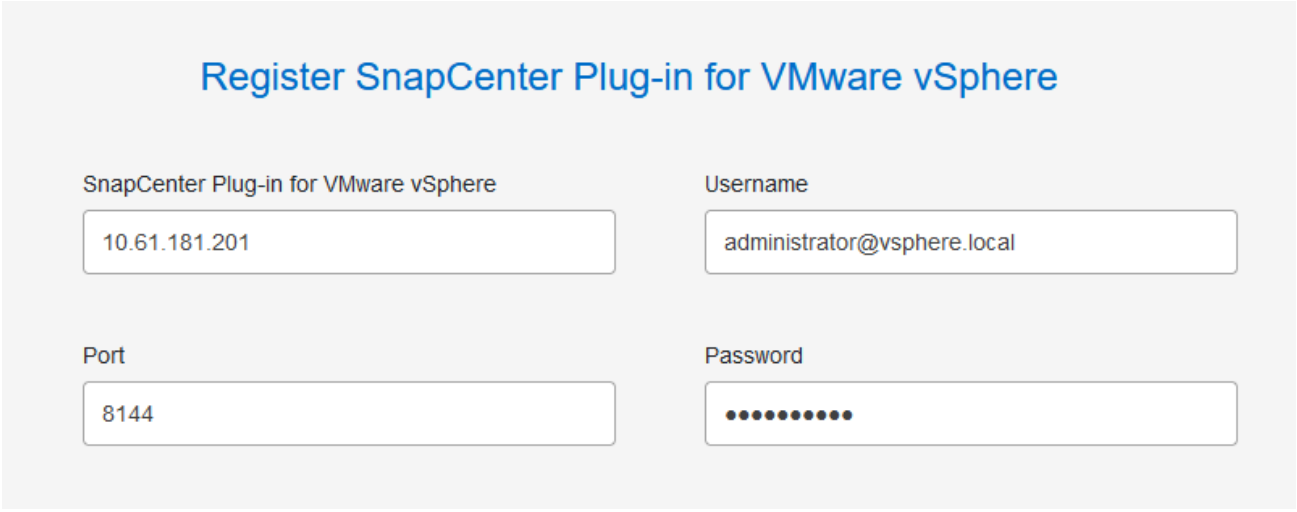
1. From the NetApp Console left-hand menu selection **Protection > Backup and recovery > Workloads >VMware**



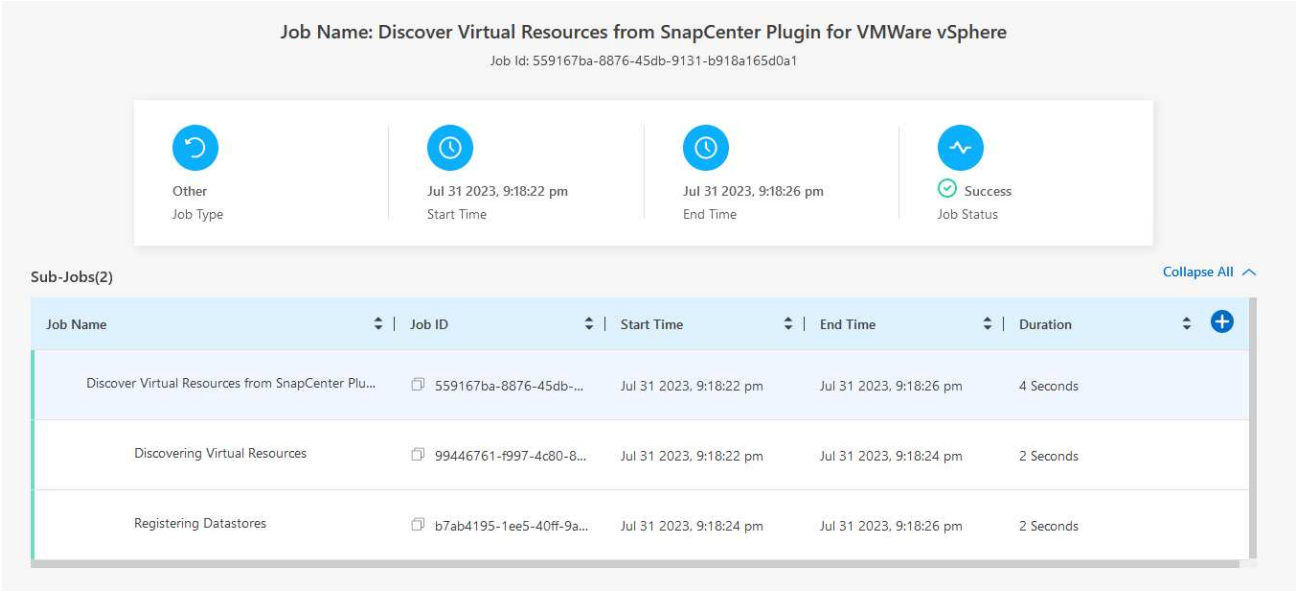
2. From the Virtual Machines main screen access the **Settings** drop down menu and select **SnapCenter Plug-in for VMware vSphere**.



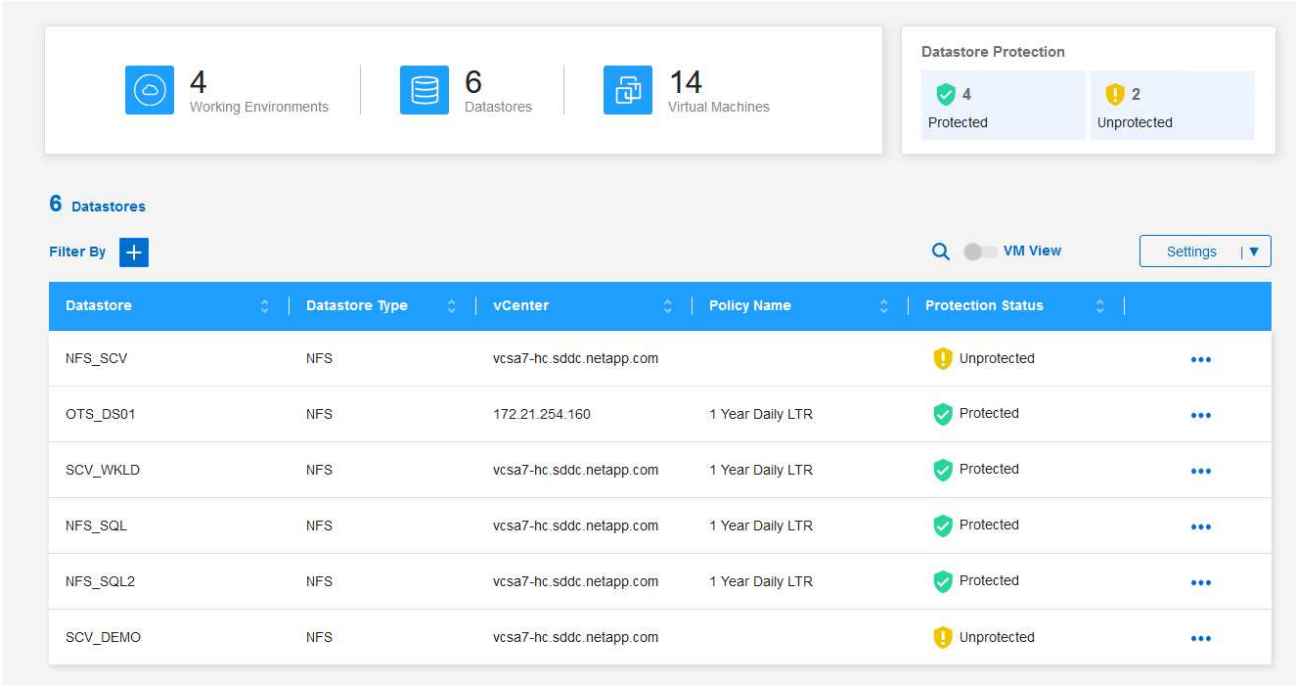
3. Click on the **Register** button and then enter the IP address and port number for the SnapCenter Plug-in appliance and the username and password for the vCenter management appliance. Click on the **Register** button to begin the discovery process.



4. The progress of jobs can be monitored from the Job Monitoring tab.



5. Once discovery is complete you will be able to view the datastores and virtual machines across all discovered SCV appliances.

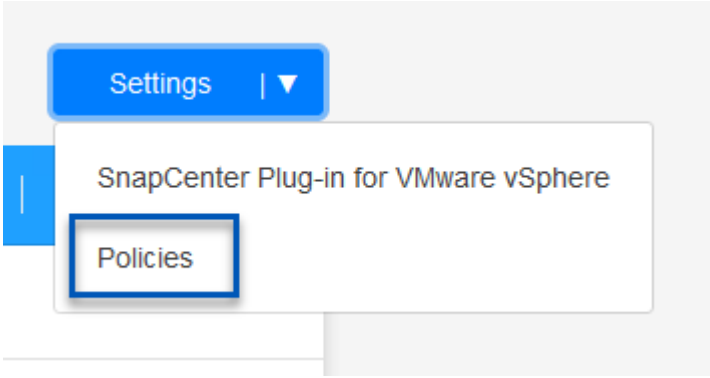


▼ Create backup policies

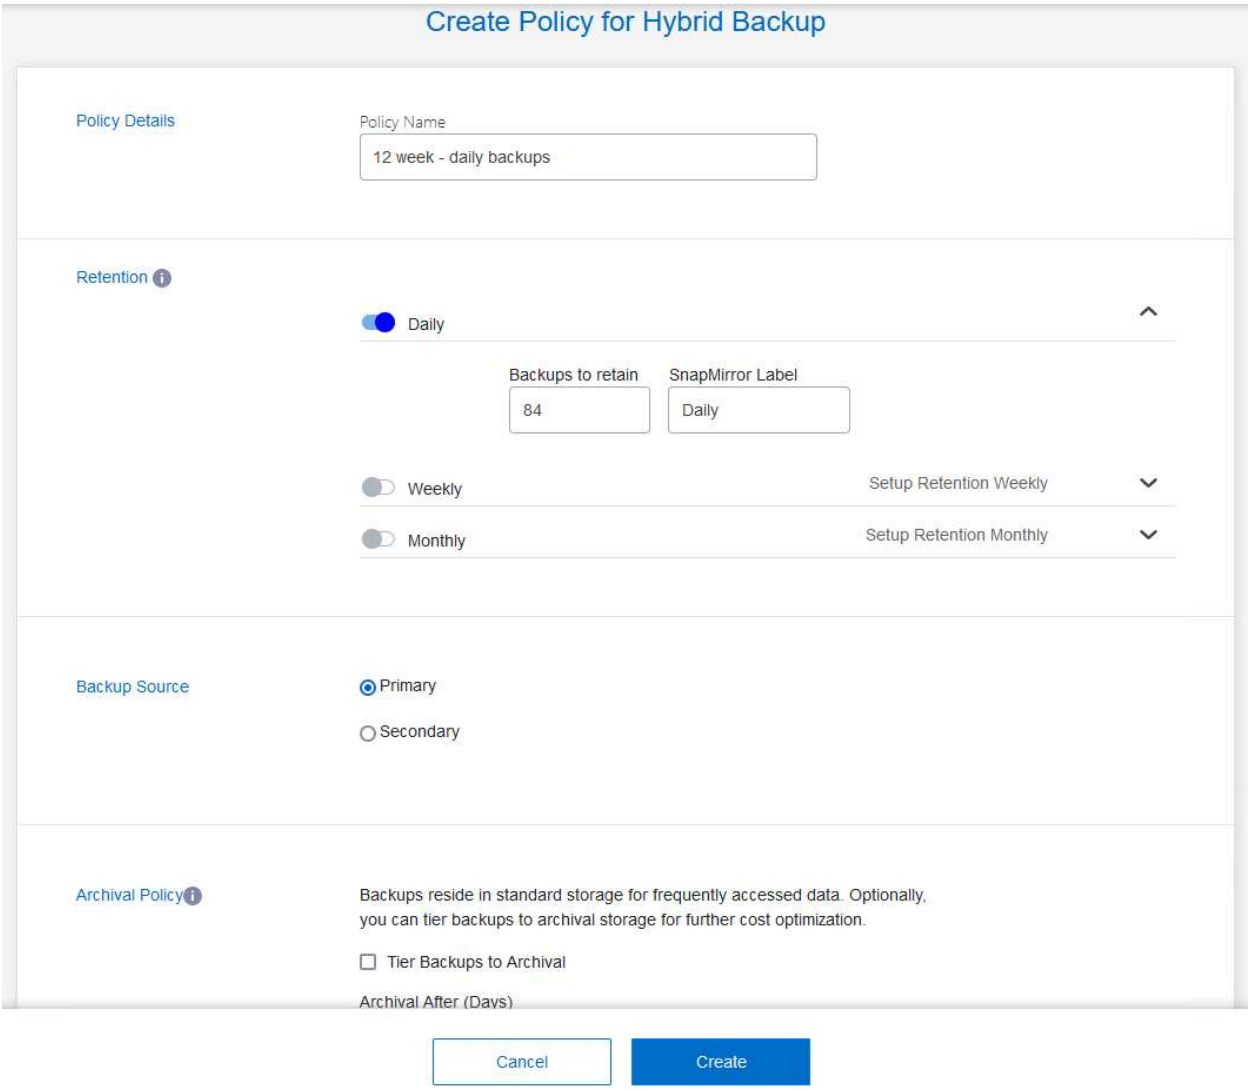
In NetApp Backup and Recovery for Virtual Machines, create policies to specify the retention period, backup source and the archival policy.

For more information on creating protection groups refer to Create a protection group.

1. From the NetApp Backup and Recovery for Virtual Machines main page, access the **Settings** drop down menu and select **Policies**.



2. Click on **Create Policy** to access the **Create Policy for Hybrid Backup** window.
   a. Add a name for the policy
   b. Select the desired retention period
   c. Select if backups will be sourced from the primary or secondary on-premises ONTAP storage system
   d. Optionally, specify after what period of time backups will be tiered to archival storage for additional cost savings.



The SnapMirror Label entered here is used to identify which backups to apply the policy too. The label name must match the label name in the corresponding on-premises SCV policy.

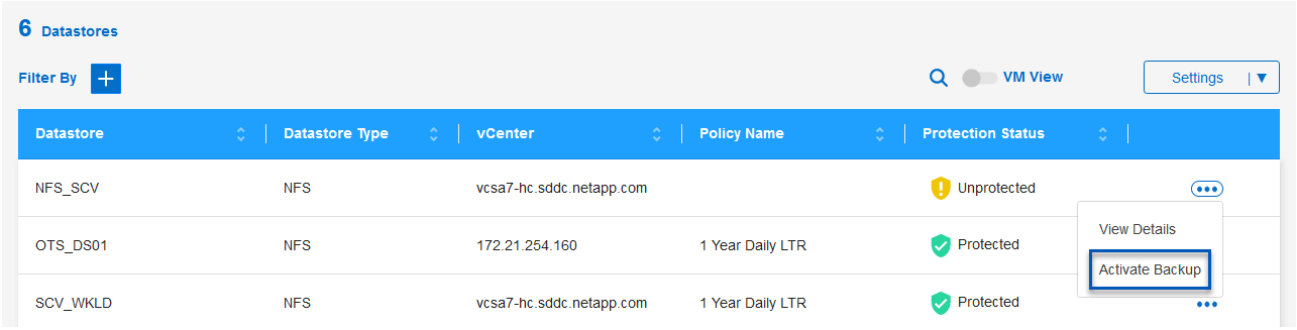
3. Click on **Create** to complete the policy creation.

▼ **Backup datastores to Amazon Web Services**

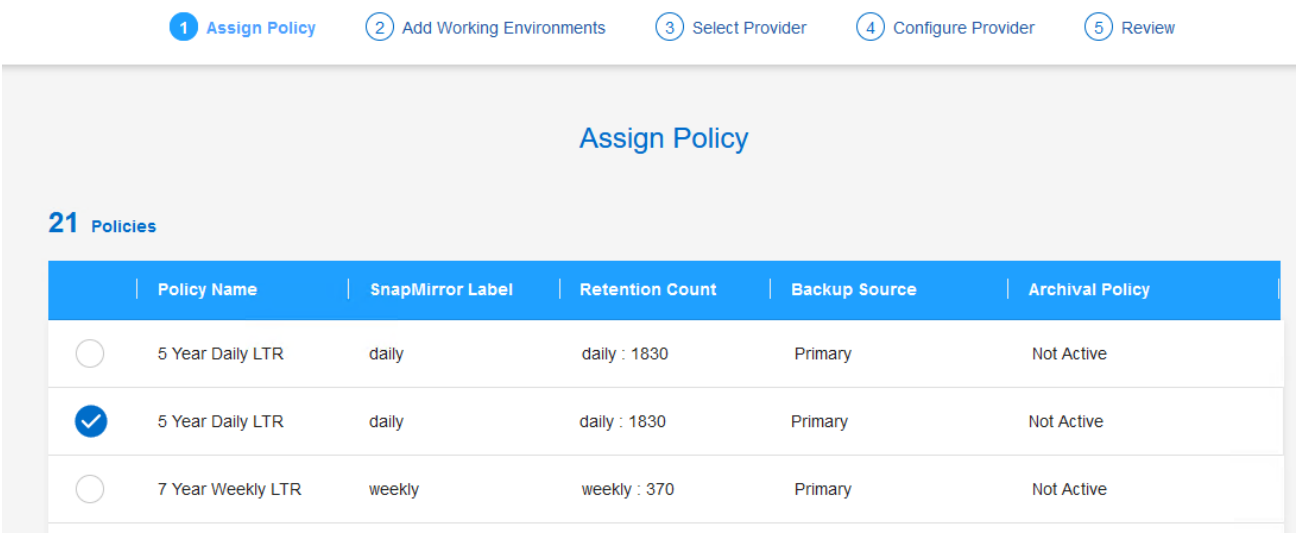
The final step is to activate data protection for the individual datastores and virtual machines. The following steps outline how to activate backups to AWS.

For more information refer to Back up VMware workloads.

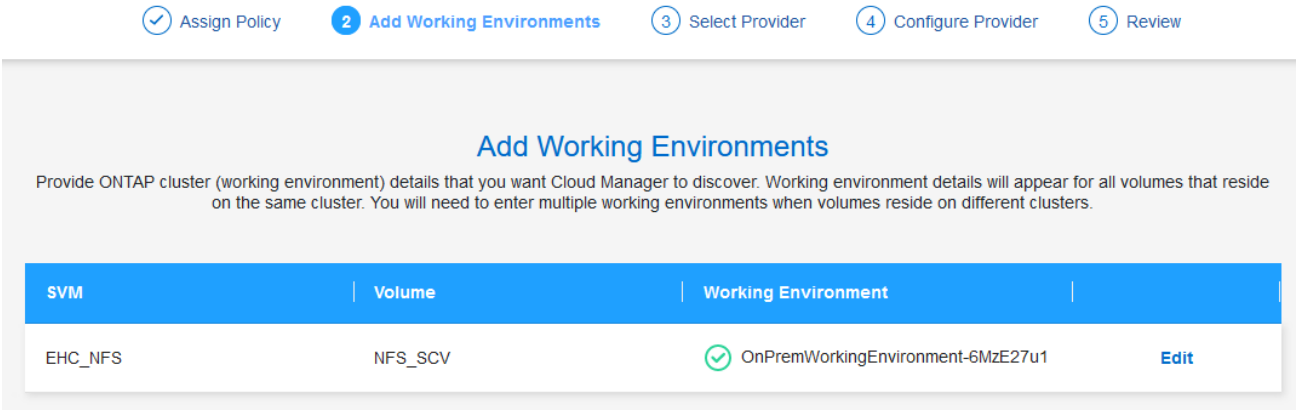
1. From the NetApp Backup and Recovery for Virtual Machines main page, access the settings drop down for the datastore to be backed up and select **Activate Backup**.



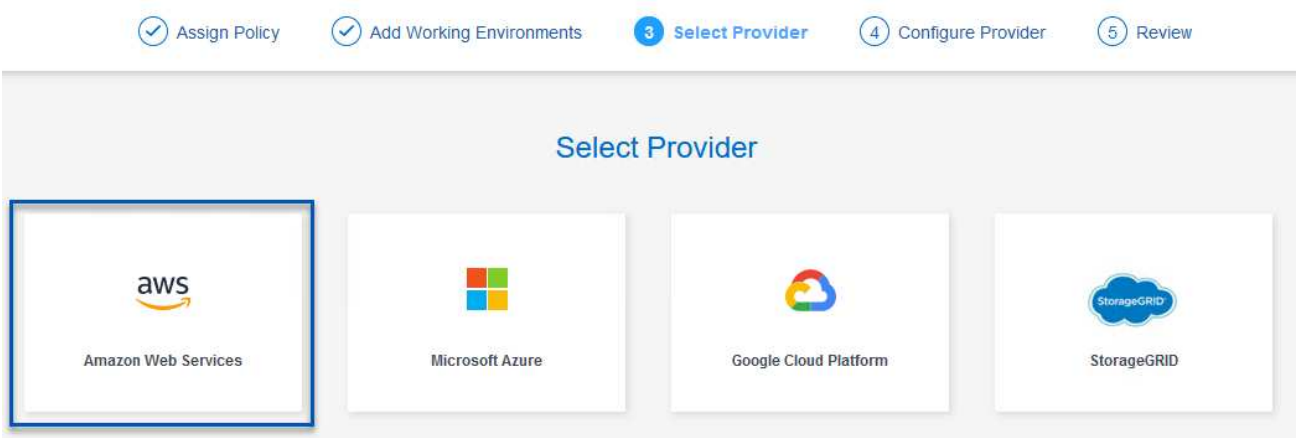
2. Assign the policy to be used for the data protection operation and click on **Next**.



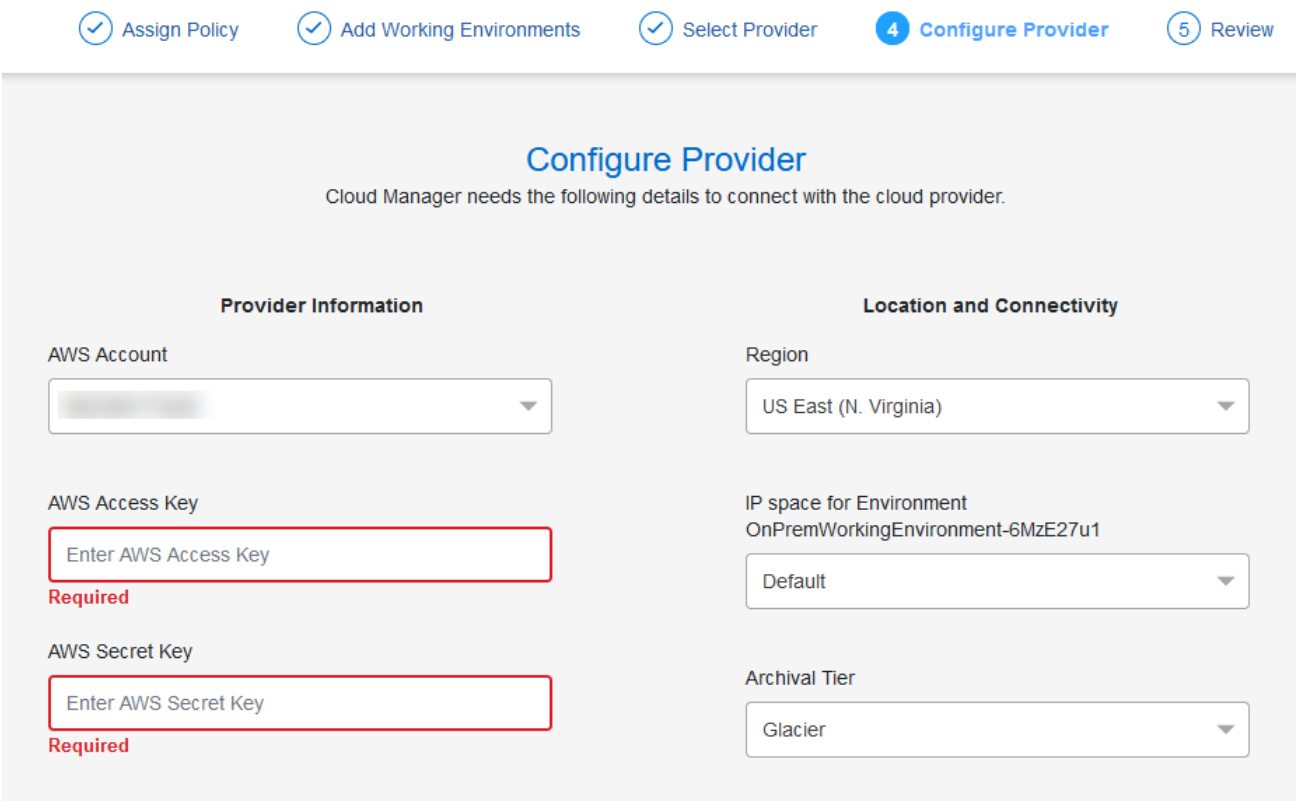
3. At the **Add Working Environments** page the datastore and working environment with a check mark should appear if the working environment has been previously discovered. If the working environment has not been previously discovered you can add it here. Click on **Next** to continue.



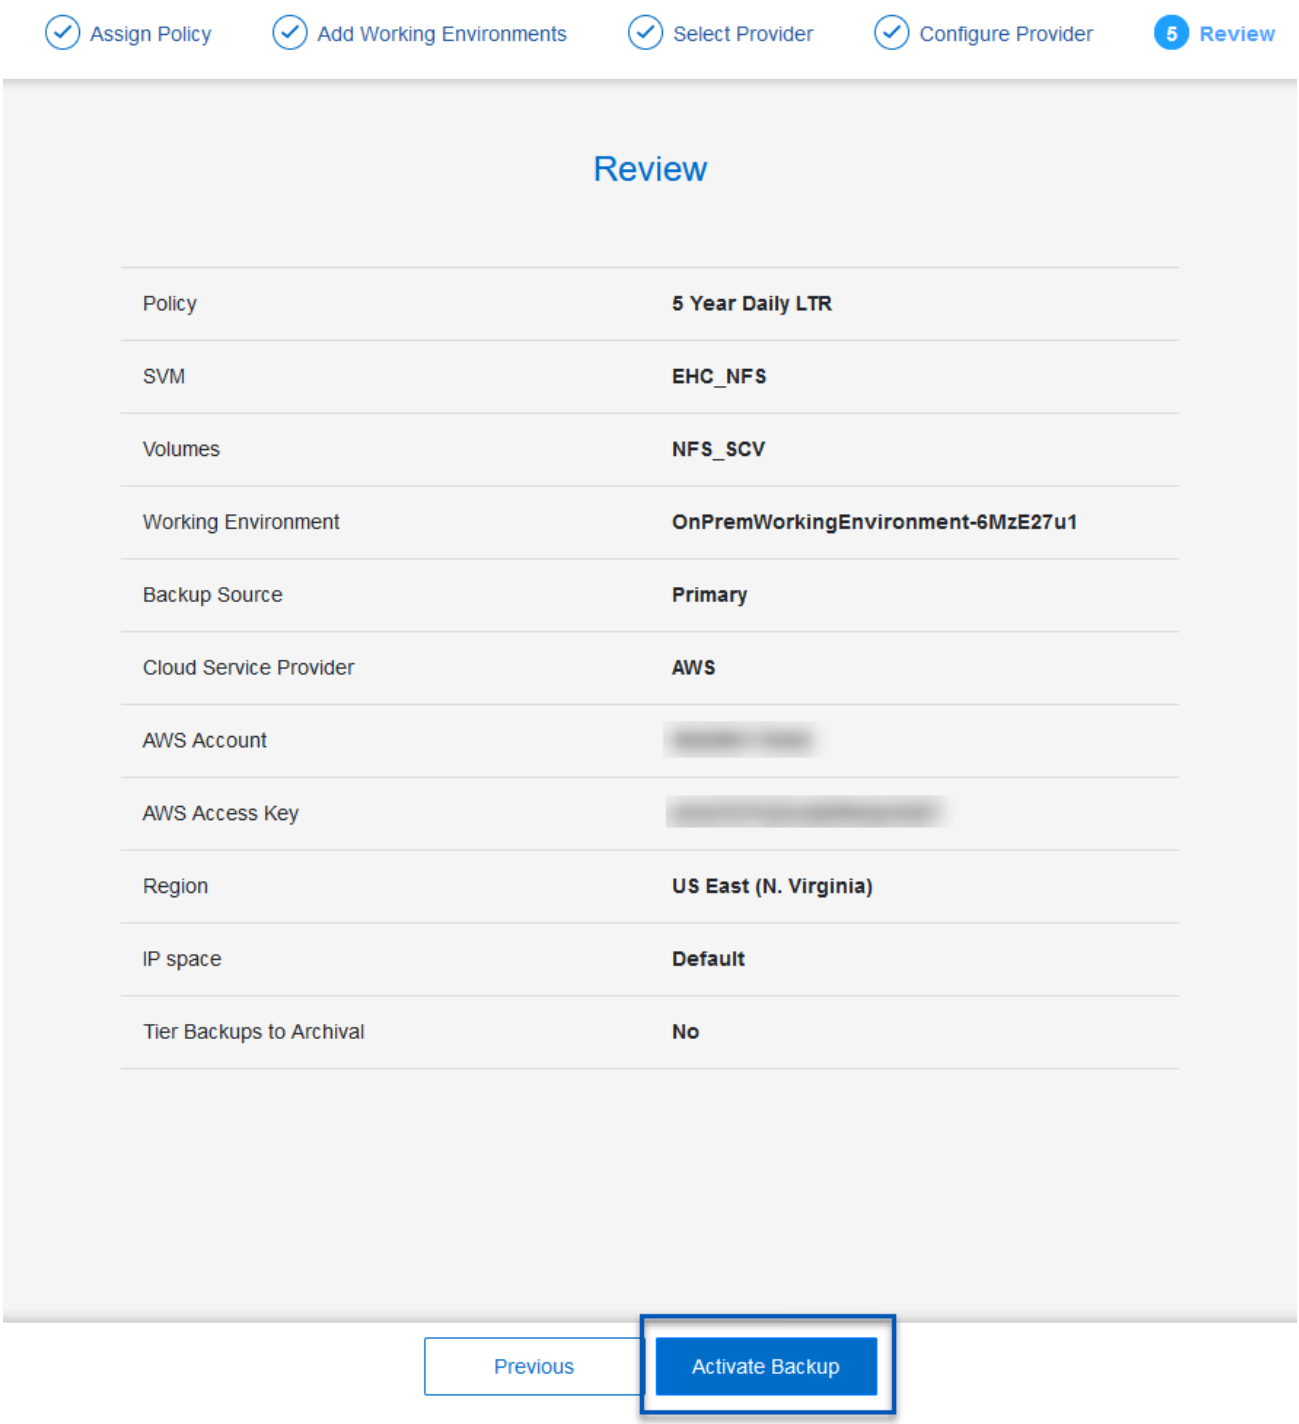
4. At the **Select Provider** page click on AWS and then click on the **Next** button to continue.



5. Fill out the provider specific credential information for AWS including the AWS access key and secret key, region, and archival tier to be used. Also, select the ONTAP IP space for the on-premises ONTAP storage system. Click on **Next**.



6. Finally, review the backup job details and click on the **Activate Backup** button to initiate data protection of the datastore.



At this point data transfer may not immediately begin. NetApp Backup and Recovery scans for any outstanding snapshots every hour and then transfers them to object storage.

Restoring Virtual Machines in the case of data loss

Ensuring the safeguarding of your data is only one aspect of comprehensive data protection. Equally crucial is the ability to promptly restore data from any location in the event of data loss or a ransomware attack. This capability is vital for maintaining seamless business operations and meeting recovery point objectives.

NetApp offers a highly adaptable 3-2-1 strategy, providing customized control over retention schedules at the

primary, secondary, and object storage locations. This strategy provides the flexibility to tailor data protection approaches to specific needs.

This section provides an overview of the data restoration process from both the SnapCenter Plug-in for VMware vSphere and NetApp Backup and Recovery for virtual machines.

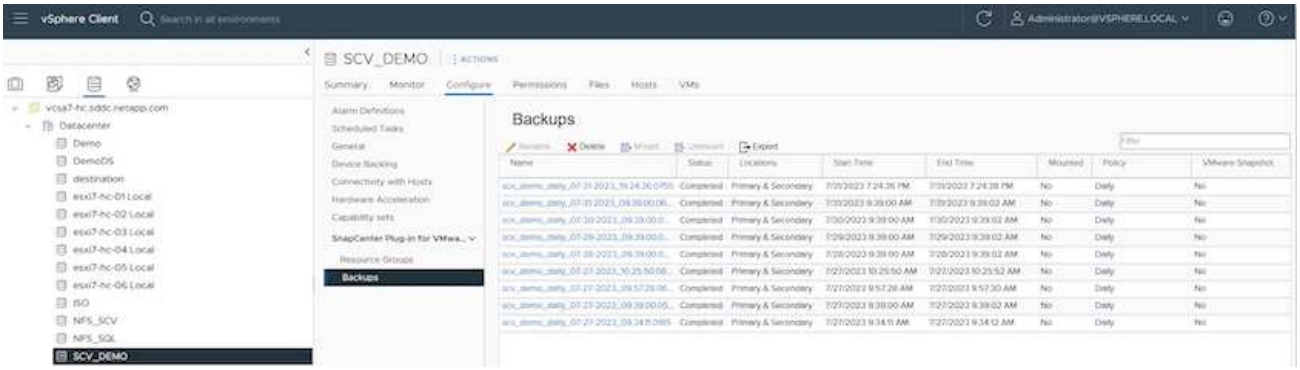
Restoring Virtual Machines from SnapCenter Plug-in for VMware vSphere

For this solution virtual machines were restored to original and alternate locations. Not all aspects of SCV's data restoration capabilities will be covered in this solution. For in depth information on all that SCV has to offer refer to the Restore VMs from backups in the product documentation.

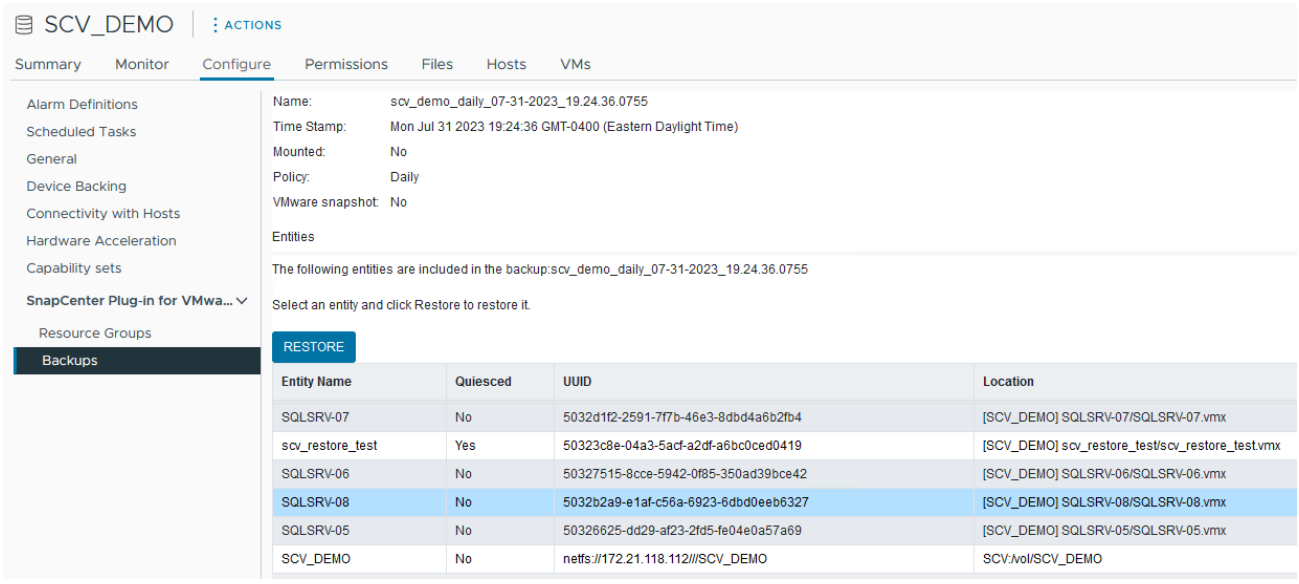
▼ **Restore virtual machines from SCV**

Complete the following steps to restore a virtual machine restore from primary or secondary storage.

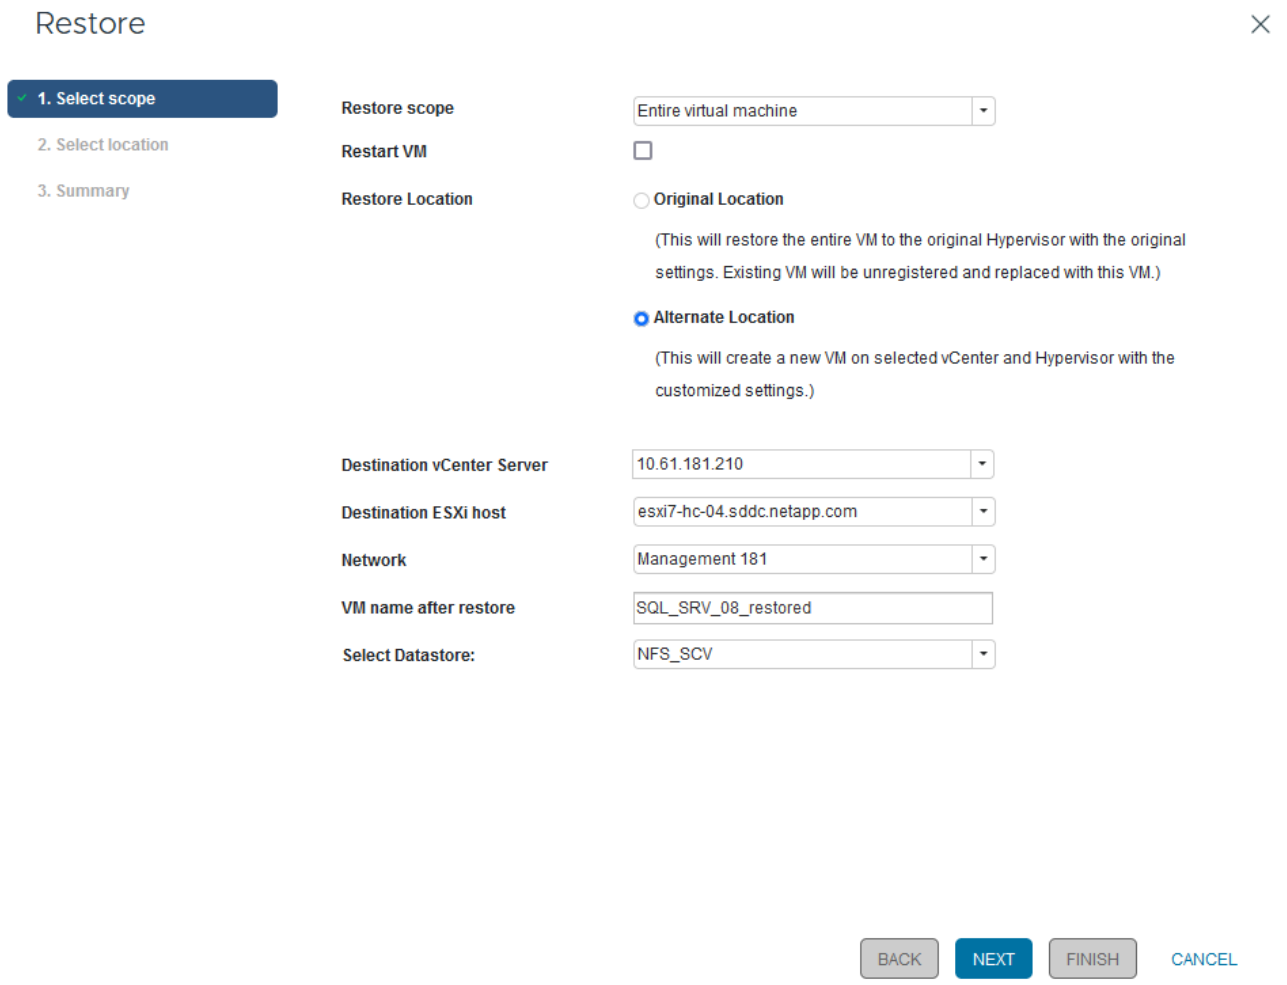
1. From the vCenter client navigate to **Inventory > Storage** and click on the datastore that contains the virtual machines you wish to restore.
2. From the **Configure** tab click on **Backups** to access the list of available backups.



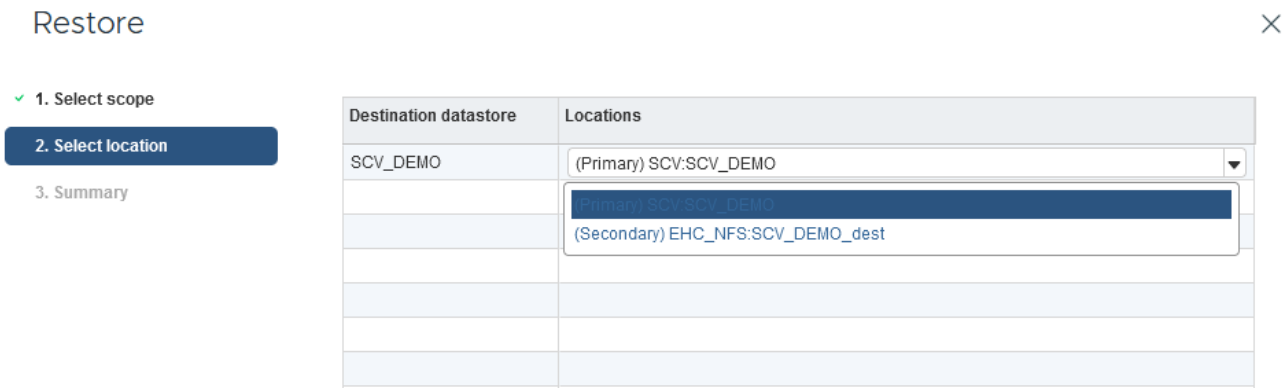
3. Click on a backup to access the list of VMs and then select a VM to restore. Click on **Restore**.



4. From the Restore wizard select to restore the entire virtual machine or a specific VMDK. Select to install to the original location or alternate location, provide VM name after restore, and destination datastore. Click **Next**.



5. Choose to backup from the primary or secondary storage location.



6. Finally, review a summary of the backup job and click on Finish to begin the restore process.

Restoring Virtual Machines from NetApp Backup and Recovery for Virtual Machines

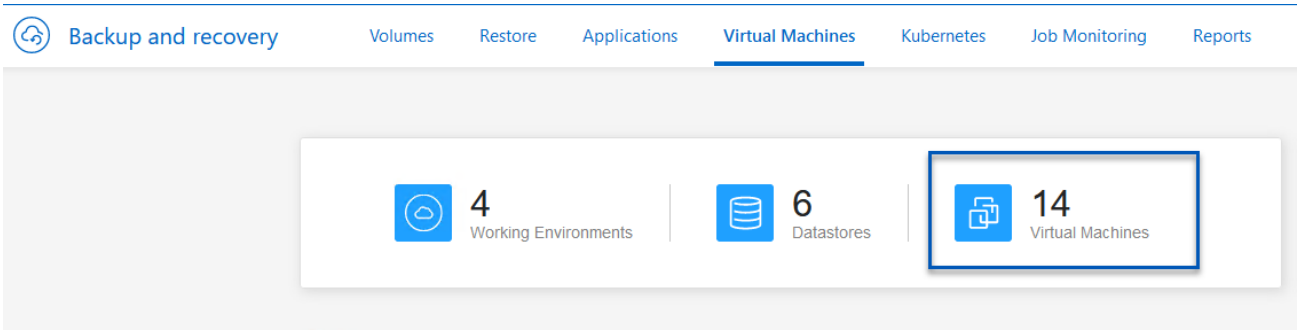
NetApp Backup and Recovery for Virtual Machines allows restores of virtual machines to their original location. Restore functions are accessed through the NetApp Backup and Recovery web console.

For more information refer to Restore virtual machines data from the cloud.

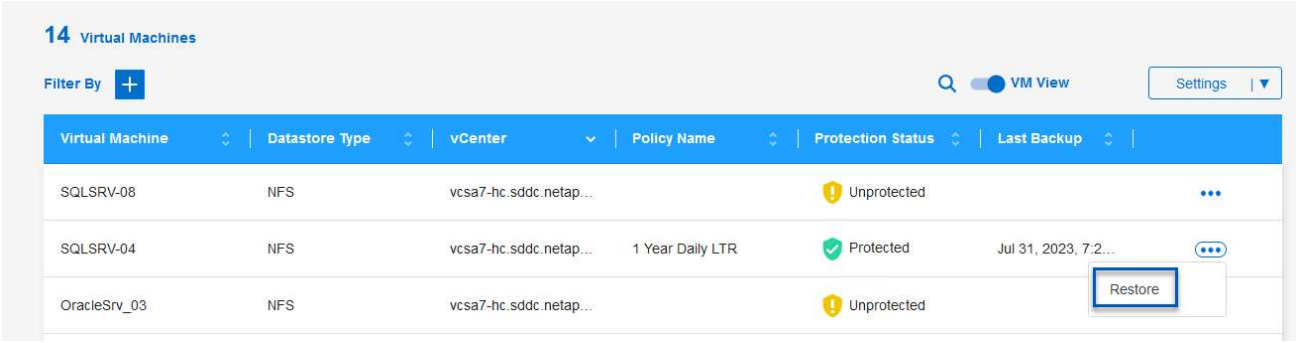
▼ **Restore virtual machines from NetApp Backup and Recovery**

To restore a virtual machine from NetApp Backup and Recovery, complete the following steps.

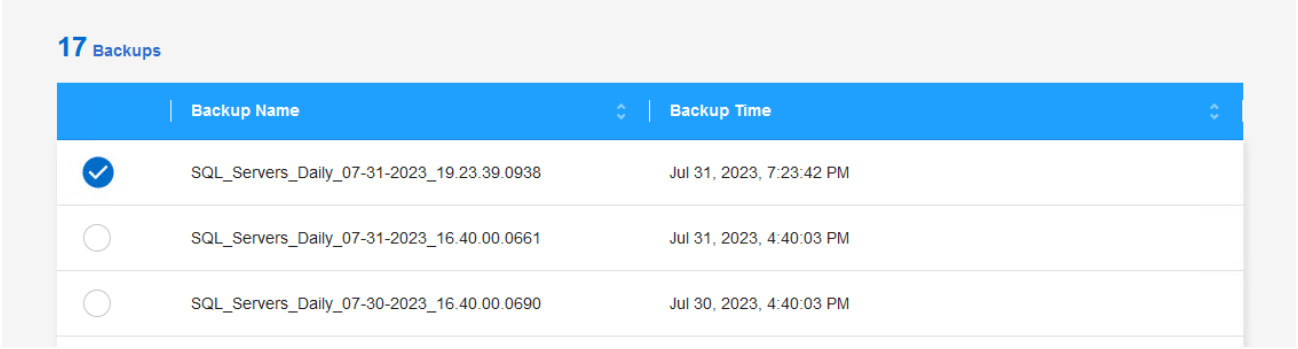
1. Navigate to **Protection > Backup and recovery > Virtual Machines** and click on Virtual Machines to view the list of virtual machines available to be restored.



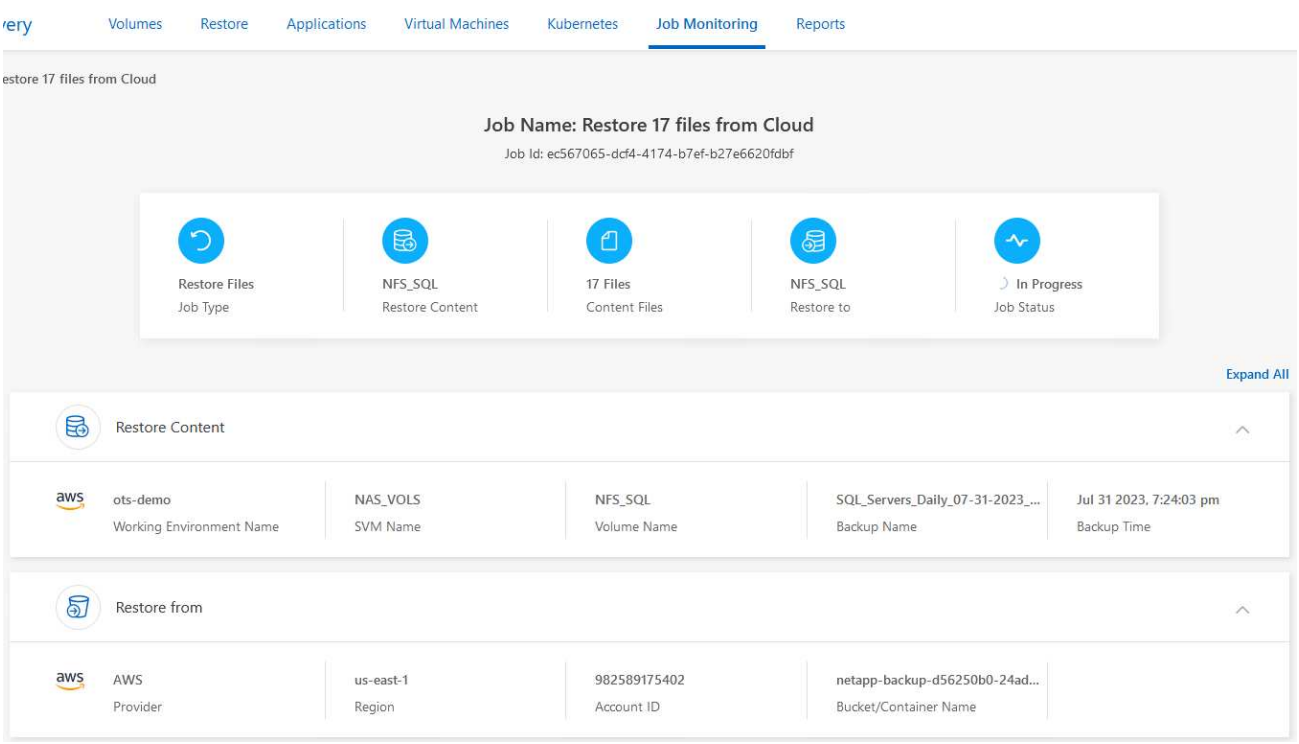
2. Access the settings drop down menu for the VM to be restored and select



3. Select the backup to restore from and click on **Next**.



4. Review a summary of the backup job and click on **Restore** to start the restore process.
5. Monitor the progress of the restore job from the **Job Monitoring** tab.



Conclusion

The 3-2-1 backup strategy, when implemented with SnapCenter Plug-in for VMware vSphere and NetApp Backup and Recovery for virtual machines, offers a robust, reliable, and cost-effective solution for data protection. This strategy not only ensures data redundancy and accessibility but also provides the flexibility of restoring data from any location and from both on-premises ONTAP storage systems and cloud based object storage.

The use case presented in this documentation focuses on proven data protection technologies that highlight the integration between NetApp, VMware, and the leading cloud providers. The SnapCenter Plug-in for VMware vSphere provides seamless integration with VMware vSphere, allowing for efficient and centralized management of data protection operations. This integration streamlines the backup and recovery processes for virtual machines, enabling easy scheduling, monitoring, and flexible restore operations within the VMware ecosystem. NetApp Backup and Recovery for virtual machines provides the one (1) in 3-2-1 by providing secure, air-gapped backups of virtual machine data to cloud based object storage. The intuitive interface and logical workflow provide a secure platform for long-term archival of critical data.

Additional Information

To learn more about the technologies presented in this solution refer to the following additional information.

- SnapCenter Plug-in for VMware vSphere documentation
- NetApp Console documentation

NetApp Disaster Recovery

Overview

Disaster Recovery is foremost in the minds of every VMware administrator. Because VMware encapsulates entire servers into a series of files that make up the virtual

machine; administrators take advantage of block storage-based techniques such as clones, snapshots and replicas to protect these VMs. ONTAP arrays offer built-in replication to transfer volume data, and therefore the virtual machines residing on the designated datastore LUNs, from one site to another. NetApp Disaster Recovery integrates with vSphere and automates the entire workflow for seamless failover and failback in the event of disaster. By combining storage replication with intelligent automation, administrators now have a manageable way to not only configure, automate, and test disaster recovery plans, but the means to easily run them in the case of a disaster.

Most time-consuming parts of a DR failover in a VMware vSphere environment is the execution of the steps necessary to inventory, register, reconfigure, and power up VMs at the DR site. An ideal solution has both a low RPO (as measured in minutes) and a low RTO (measured in minutes to hours). One factor that is often overlooked in a DR solution is the ability to test the DR solution efficiently on a periodic interval.

To architect a DR solution, keep the following factors in mind:

- The recovery time objective (RTO). The RTO is how quickly a business can recover from a disaster, or, more specifically, how long it takes to execute the recovery process to make business services available again.
- The recovery point objective (RPO). The RPO is how old the recovered data is after it has been made available, relative to the time that the disaster occurred.
- Scalability and adaptability. This factor includes the ability to grow storage resources incrementally as demand increases.

For more technical information on the available solutions, please see:

- NetApp Disaster Recovery for NFS Datastores
- NetApp Disaster Recovery for VMFS Datastores

NetApp Disaster Recovery for NFS Datastores

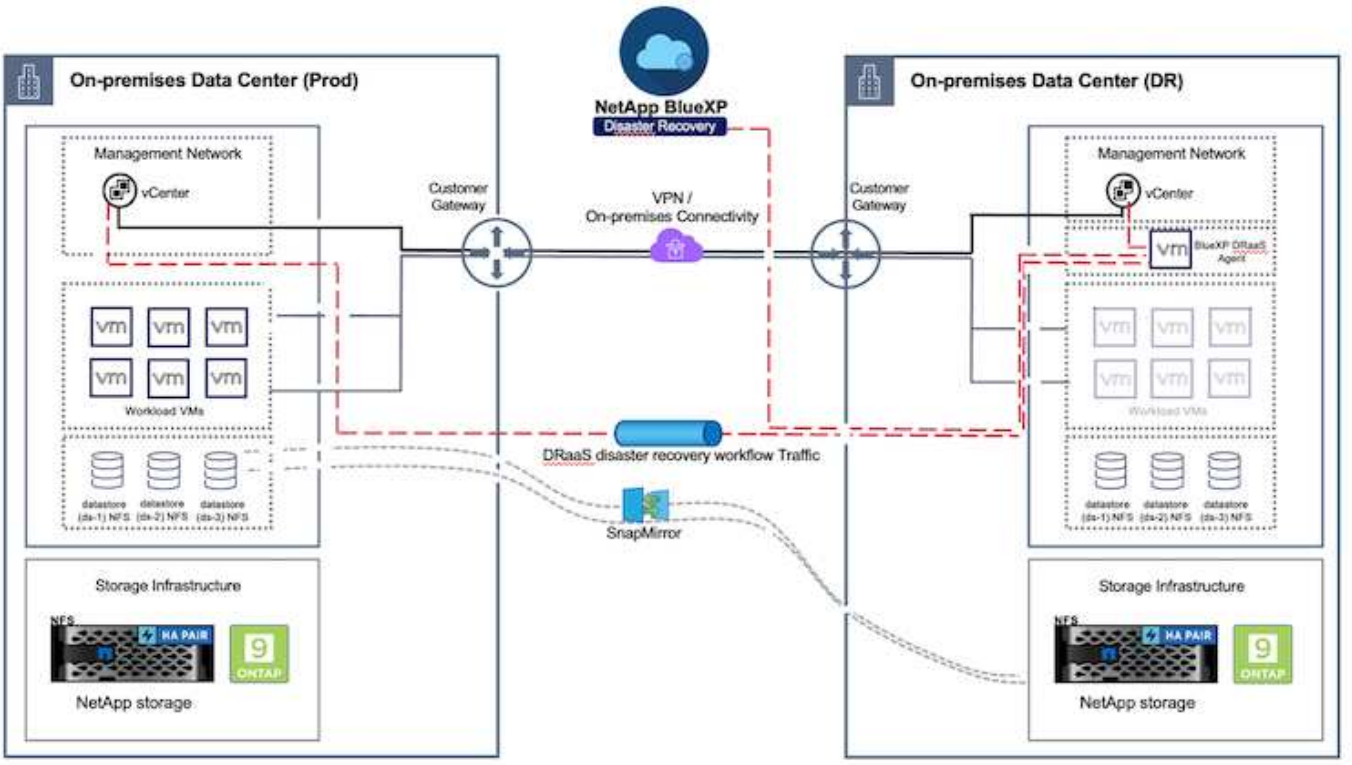
Implementing disaster recovery through block-level replication from the production site to the disaster recovery site is a resilient and cost-effective method for safeguarding workloads against site outages and data corruption events, such as ransomware attacks. Using NetApp SnapMirror replication, VMware workloads running on on-premises ONTAP systems with NFS datastore can be replicated to another ONTAP storage system located in a designated recovery datacenter where VMware is also deployed.

This section of the document describes the configuration of NetApp Disaster Recovery to set up disaster recovery for on-premises VMware VMs to another designated site. As part of this setup, Console agent, onboarding the ONTAP arrays, Defining Sites and associating VMware vCenter. In addition, this document details how to configure replication between sites and how to setup and test a recovery plan. The last section has instructions for performing a full site failover and how to failback when the primary site is recovered and bought online.

Utilizing the Disaster Recovery, part of NetApp Data Services, companies can easily discover their on-premises VMware vCenters and ONTAP storage. Organizations can then create resource groups, create a disaster recovery plan, associate it with resource groups, and test or execute failover and failback. SnapMirror provides storage-level block replication to keep the two sites up to date with incremental changes, resulting in

a Recovery Point Objective (RPO) of up to 5 minutes. Additionally, it is possible to simulate disaster recovery procedures without affecting production or incurring additional storage costs.

Disaster Recovery leverages ONTAP's FlexClone technology to create a space-efficient copy of the NFS datastore from the last replicated Snapshot on the disaster recovery site. After completing the disaster recovery test, customers can easily delete the test environment without impacting actual replicated production resources. In case of an actual failover, the Disaster Recovery orchestrates all the necessary steps to automatically bring up the protected virtual machines on the designated disaster recovery site with just a few clicks. The service will also reverse the SnapMirror relationship to the primary site and replicate any changes from the secondary to the primary for a failback operation, when needed. All these capabilities come at a fraction of the cost compared to other well-known alternatives.

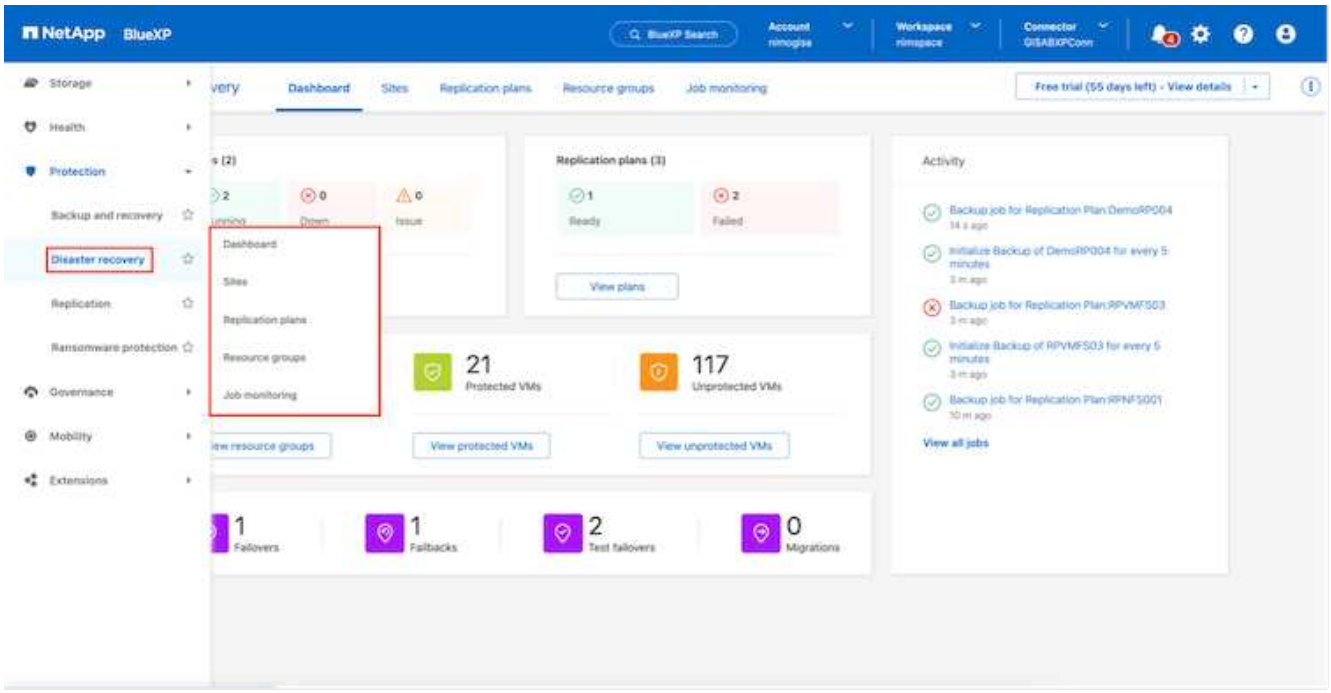


For up-to-date information and supported versions of vCenter Server, ESXi, and ONTAP for Disaster Recovery, refer to the Disaster Recovery supported configurations.

Getting started

To get started with NetApp Disaster Recovery, use NetApp console and then access the service.

1. Log in to Console.
2. From the Console left navigation, select Protection > Disaster Recovery.
3. The Disaster Recovery Dashboard appears.

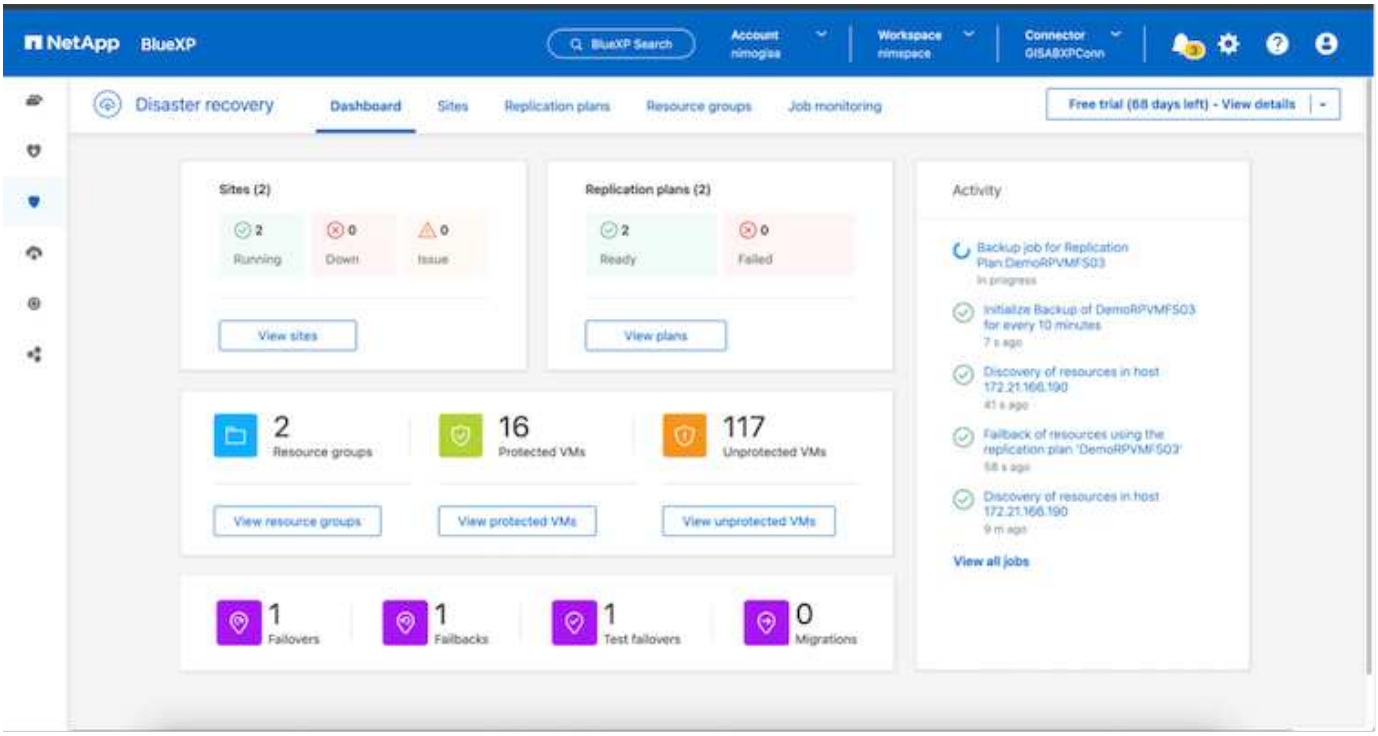


Before configuring Disaster Recovery plan, ensure the following pre-requisites are met:

- Console agent is deployed.
- Agent have connectivity to the source and destination vCenter and storage systems.
- ONTAP cluster with NFS service enabled SVM is available.
- On-premises NetApp storage systems hosting NFS datastores for VMware are added in Console.
- DNS resolution should be in place when using DNS names. Otherwise, use IP addresses for the vCenter.
- SnapMirror replication is configured for the designated NFS based datastore volumes.
- Make sure that the environment has supported versions of vCenter Server and ESXi servers.

Once the connectivity is established between the source and destination sites, proceed with configuration steps, which should take couple of clicks and about 3 to 5 minutes.

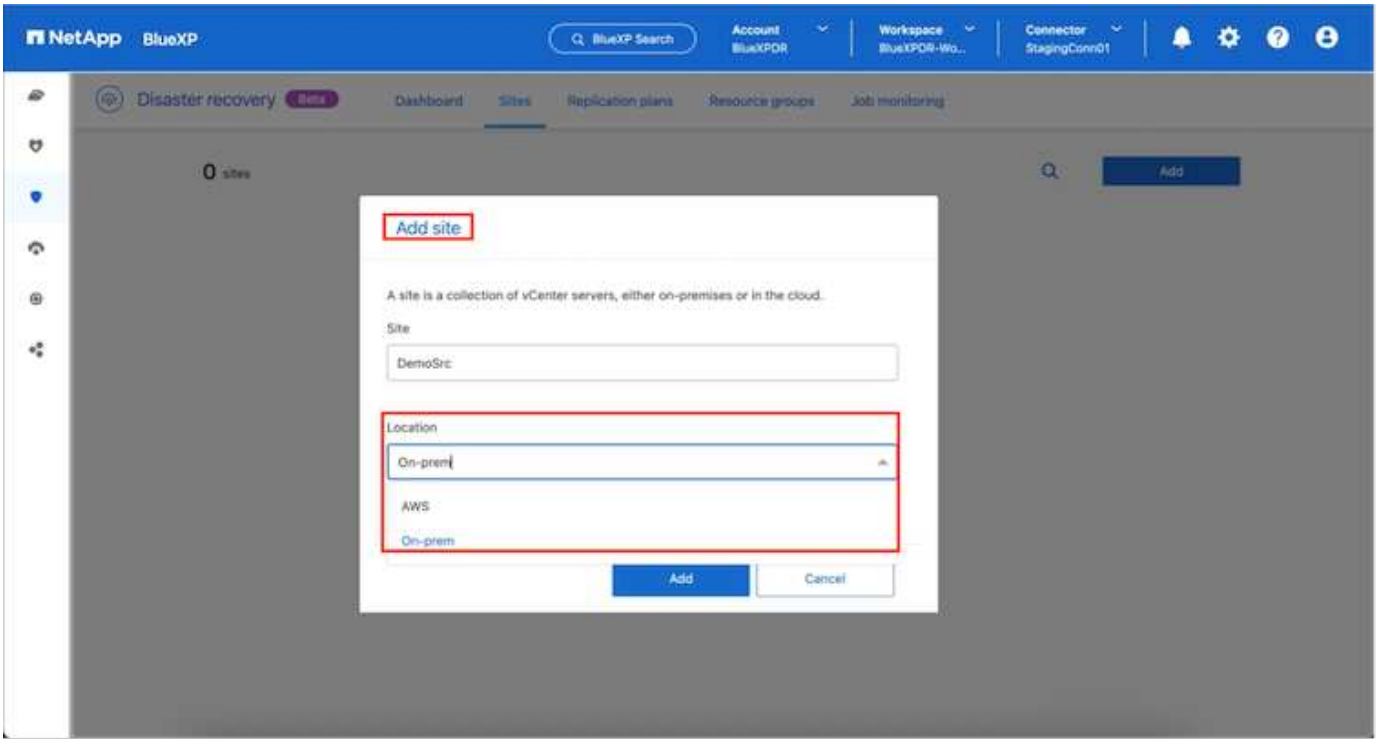
NetApp recommends deploying the Console agent in the destination site or in a third site, so that the Console agent can communicate through the network with source and destination resources.



NetApp Disaster Recovery configuration

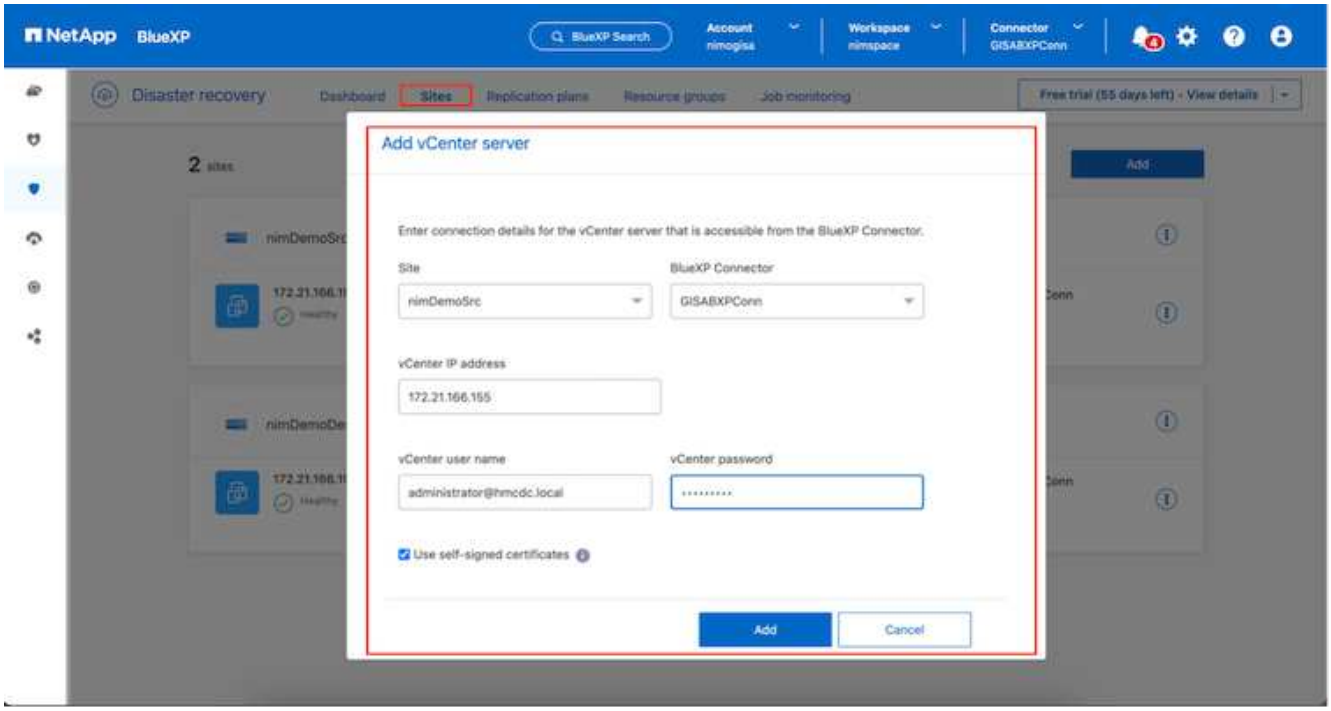
The first step in preparing for disaster recovery is to discover and add the on-premises vCenter and storage resources to NetApp Disaster Recovery.

Open Console and select **Protection > Disaster Recovery** from left navigation. Click **Sites**, provide name and pick the location type. Then, click **Add vCenter** button.

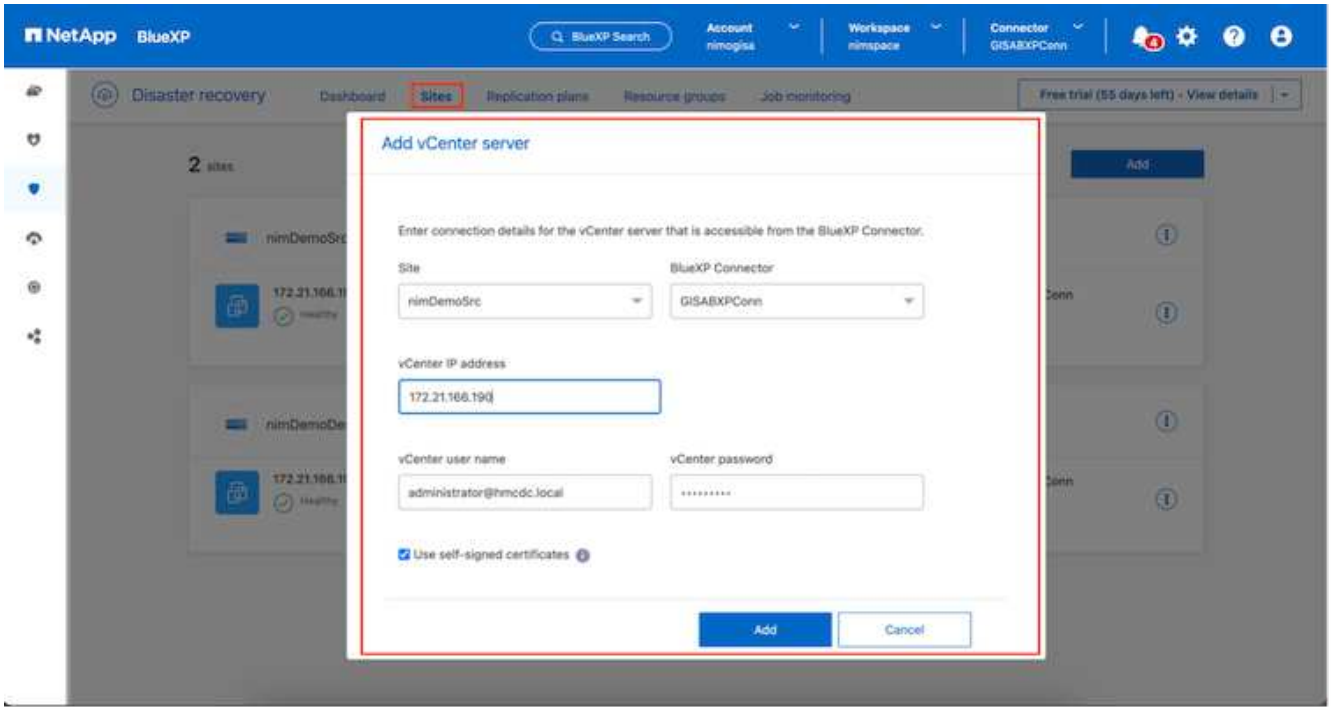


Add the following platforms:

- **Source**. On-premises vCenter.



- **Destination**. VMC SDDC vCenter.



Once the vCenters are added, automated discovery is triggered.

Configuring Storage replication between source site array and destination site array

SnapMirror provides data replication in a NetApp environment. Built on NetApp Snapshot technology, SnapMirror replication is extremely efficient because it replicates only the blocks that have been changed or added since the previous update. SnapMirror is easily configured by using either NetApp OnCommand System

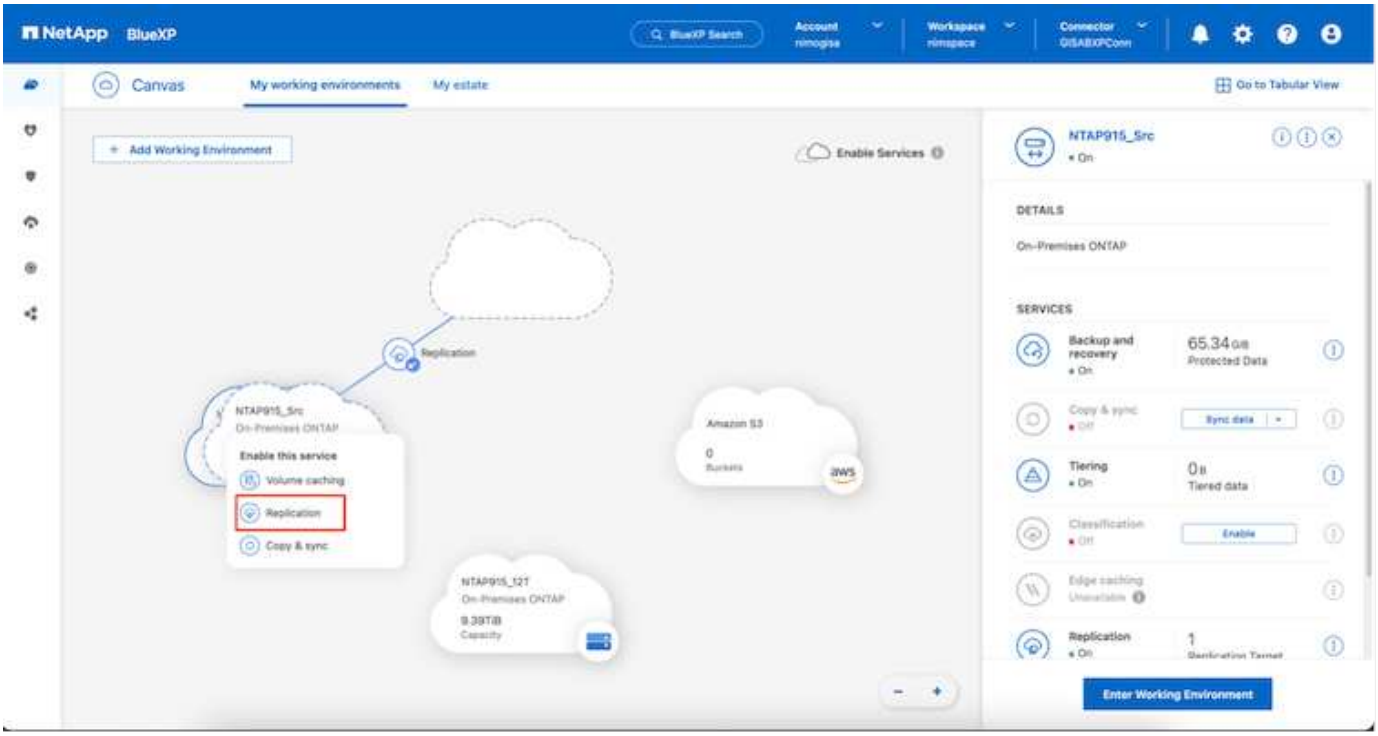
Manager or the ONTAP CLI. NetApp Data Services also creates the SnapMirror relationship provided cluster and SVM peering is configured beforehand.

For cases in which the primary storage is not completely lost, SnapMirror provides an efficient means of resynchronizing the primary and DR sites. SnapMirror can resynchronize the two sites, transferring only changed or new data back to the primary site from the DR site by simply reversing the SnapMirror relationships. This means replication plans in Disaster Recovery can be resynchronized in either direction after a failover without recopying the entire volume. If a relationship is resynchronized in the reverse direction, only new data that was written since the last successful synchronization of the Snapshot copy is sent back to the destination.

If SnapMirror relationship is already configured for the volume via CLI or System Manager, NetApp Data Services picks up the relationship and continues with the rest of the workflow operations.

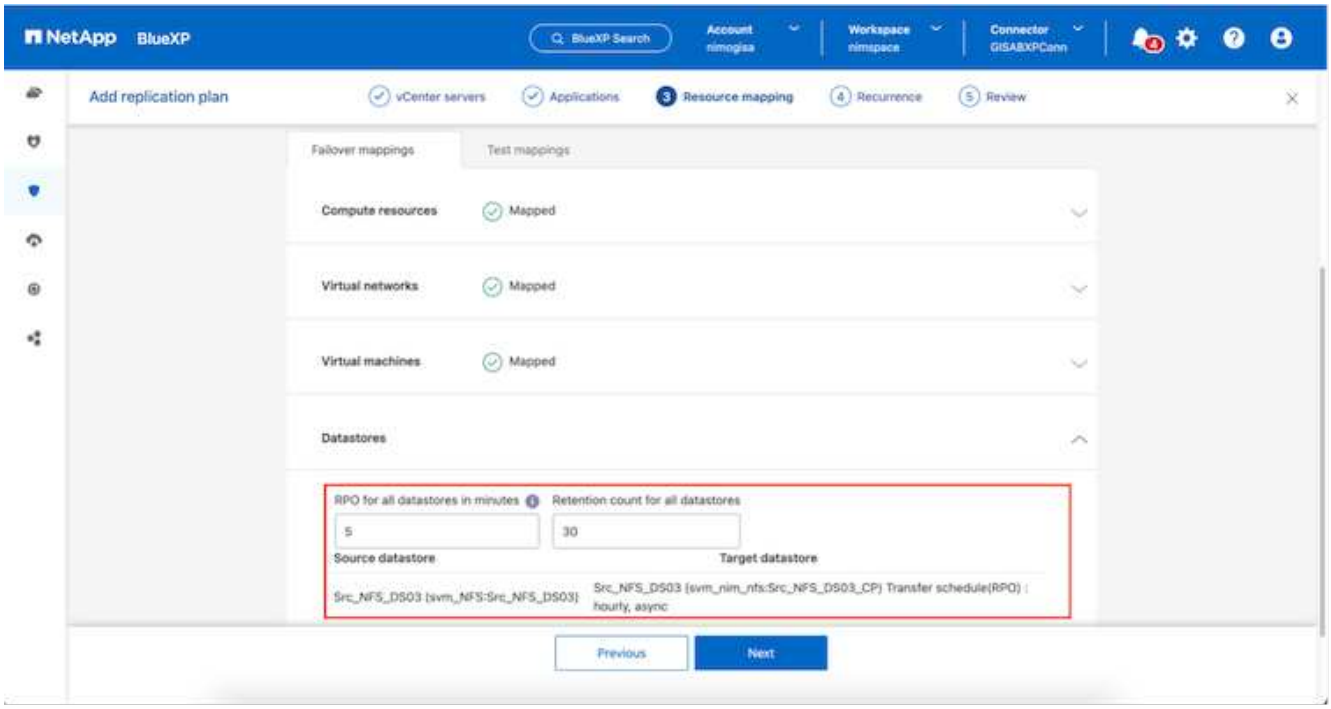
How to set it up for VMware Disaster Recovery

The process to create SnapMirror replication remains the same for any given application. The process can be manual or automated. The easiest way is to leverage NetApp Console to configure SnapMirror replication by using simple drag & drop of the source ONTAP system in the environment onto the destination to trigger the wizard that guides through the rest of the process.



Disaster Recovery can also automate the same provided the following two criteria's are met:

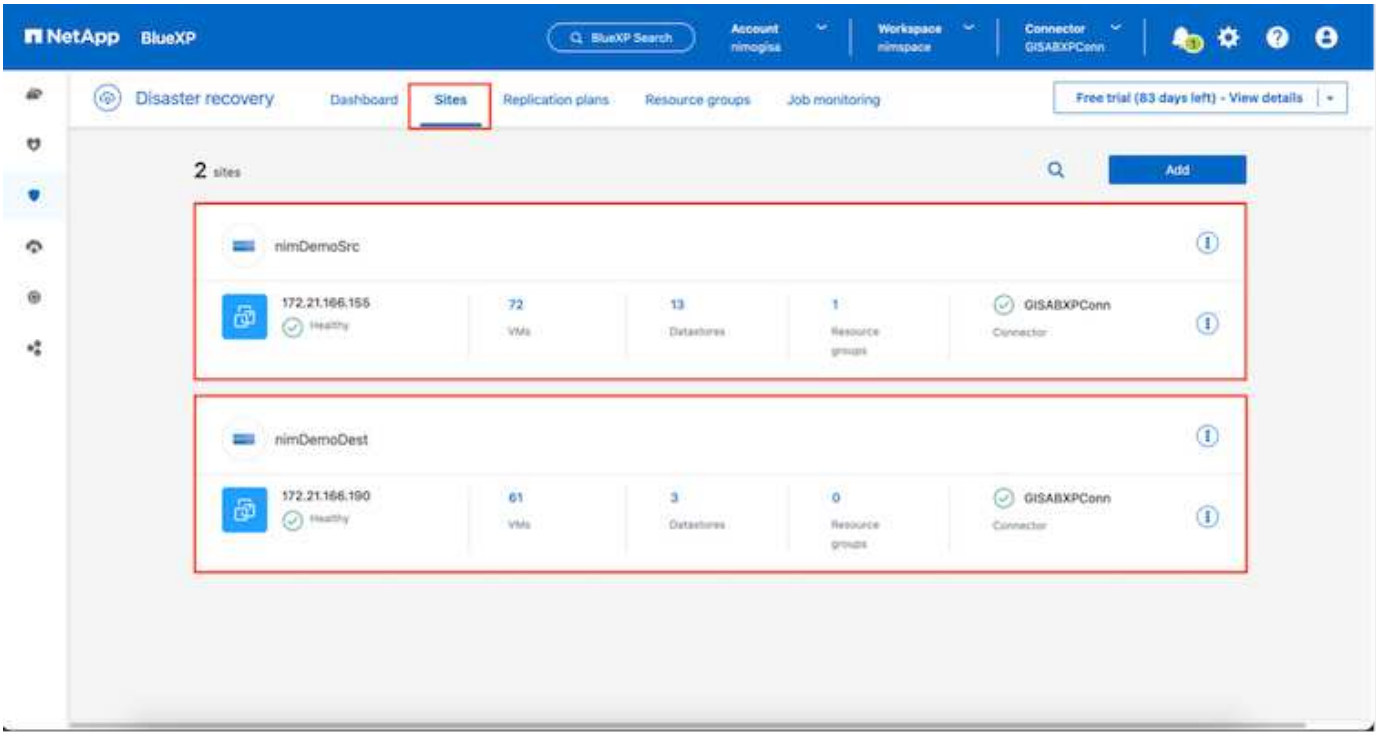
- Source and destination clusters have a peer relationship.
- Source SVM and destination SVM have a peer relationship.



If SnapMirror relationship is already configured for the volume via CLI, Disaster Recovery picks up the relationship and continues with the rest of the workflow operations.

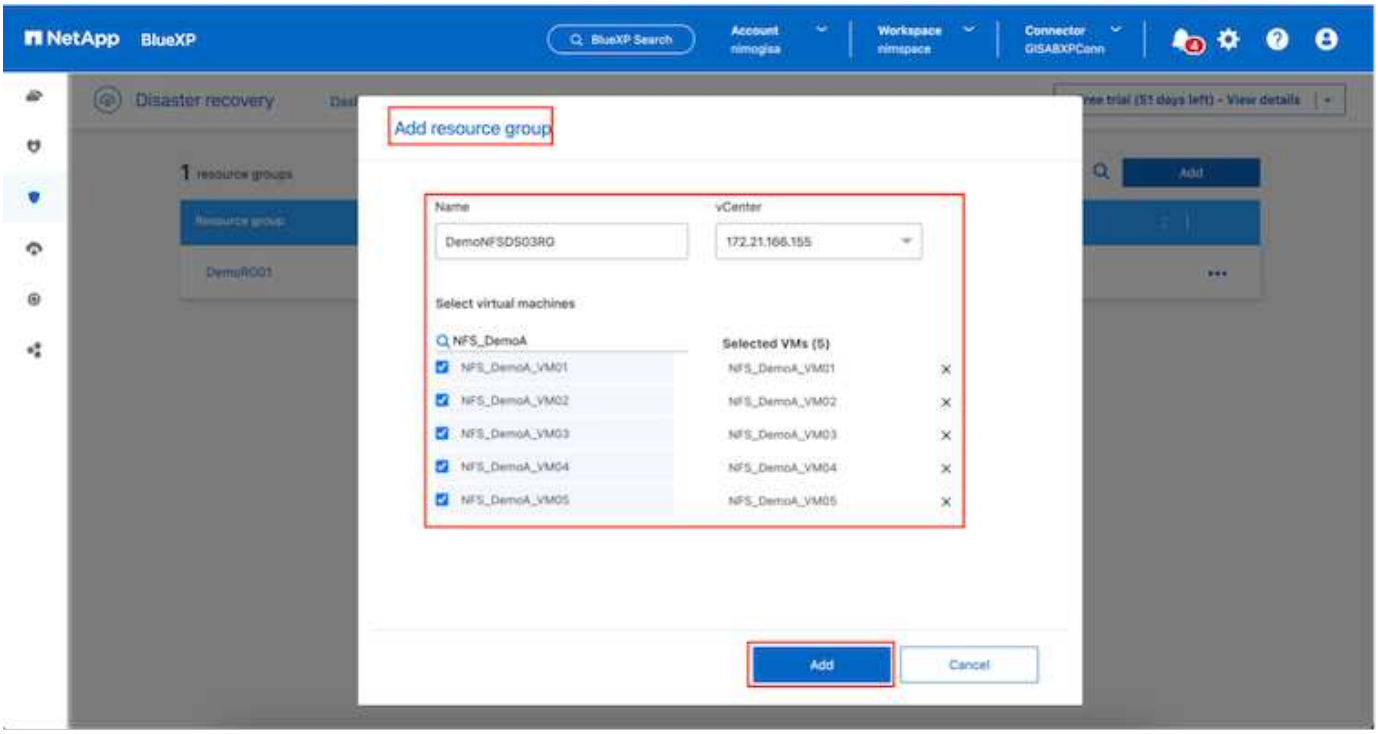
What can Disaster Recovery do for you?

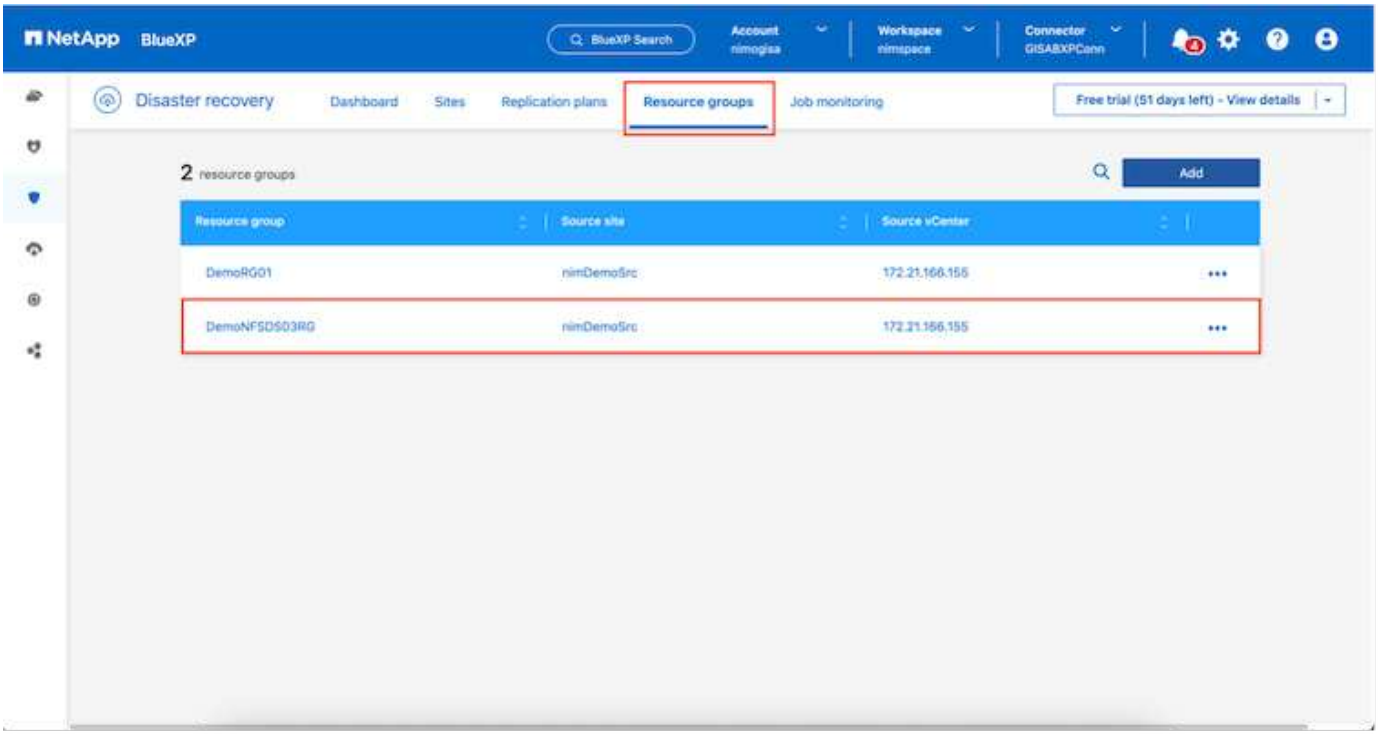
After the source and destination sites are added, Disaster Recovery performs automatic deep discovery and displays the VMs along with associated metadata. Disaster Recovery also automatically detects the networks and port groups used by the VMs and populates them.



After the sites have been added, VMs can be grouped into resource groups. Disaster Recovery resource

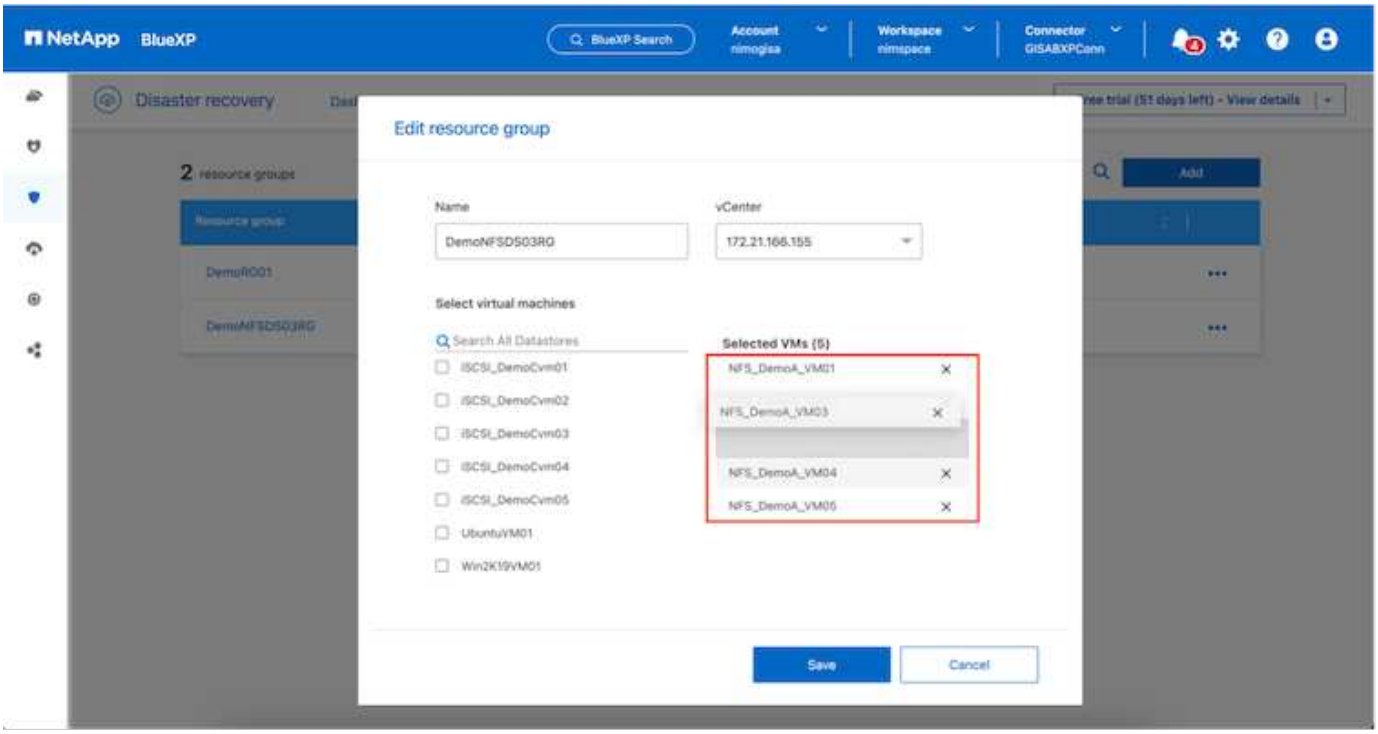
groups allow you to group a set of dependent VMs into logical groups that contain their boot orders and boot delays that can be executed upon recovery. To start creating resource groups, navigate to **Resource Groups** and click **Create New Resource Group**.



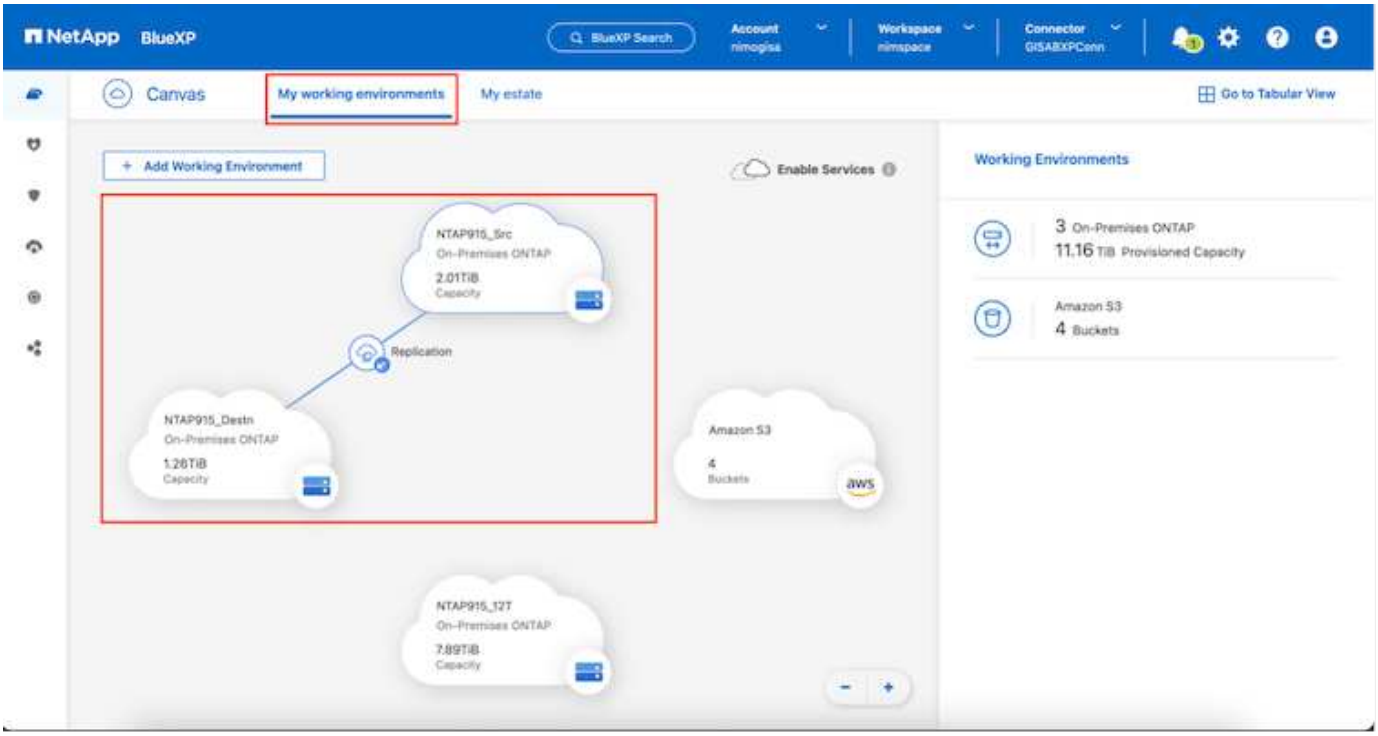


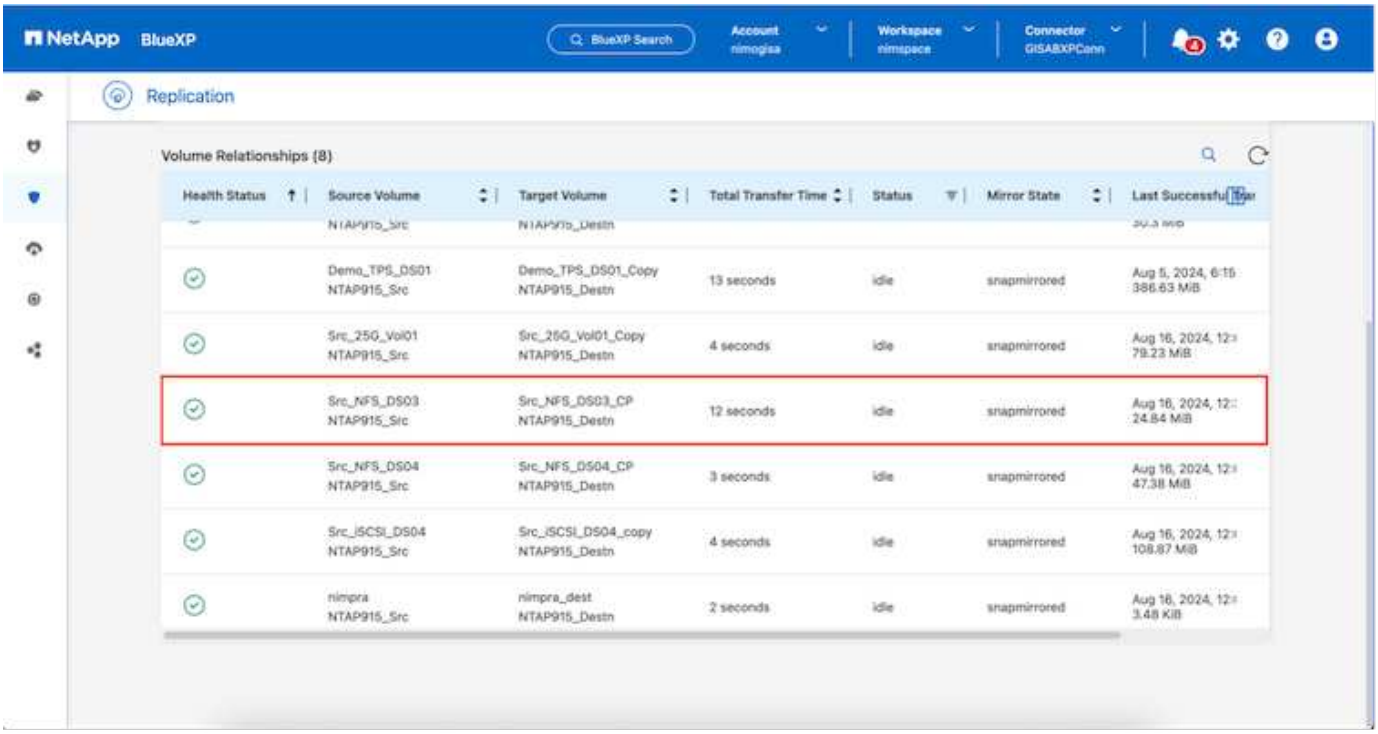
The resource group can also be created while creating a replication plan.

The boot order of the VMs can be defined or modified during the creation of resource groups by using simple drag and drop mechanism.



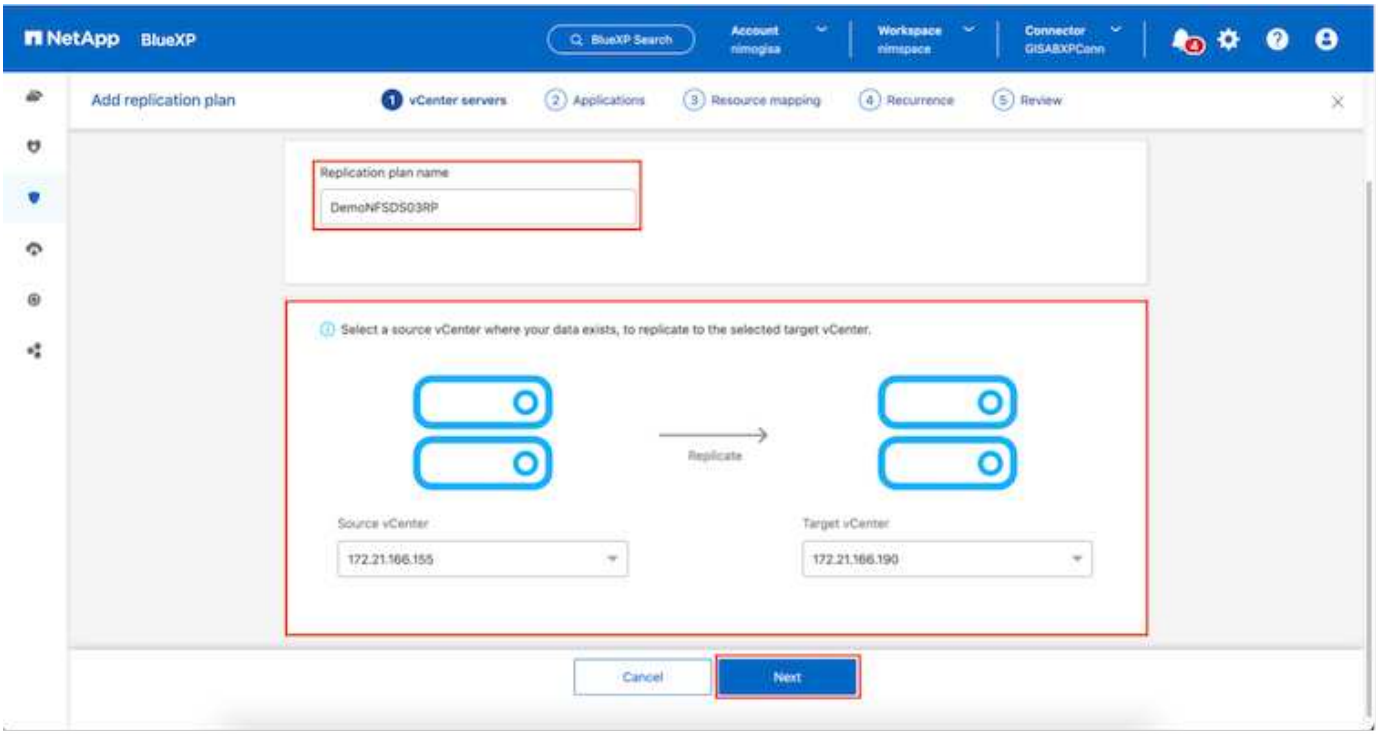
Once the resource groups are created, the next step is to create the execution blueprint or a plan to recover virtual machines and applications in the event of a disaster. As mentioned in the prerequisites, SnapMirror replication can be configured beforehand or DRaaS can configure it using the RPO and retention count specified during creation of the replication plan.





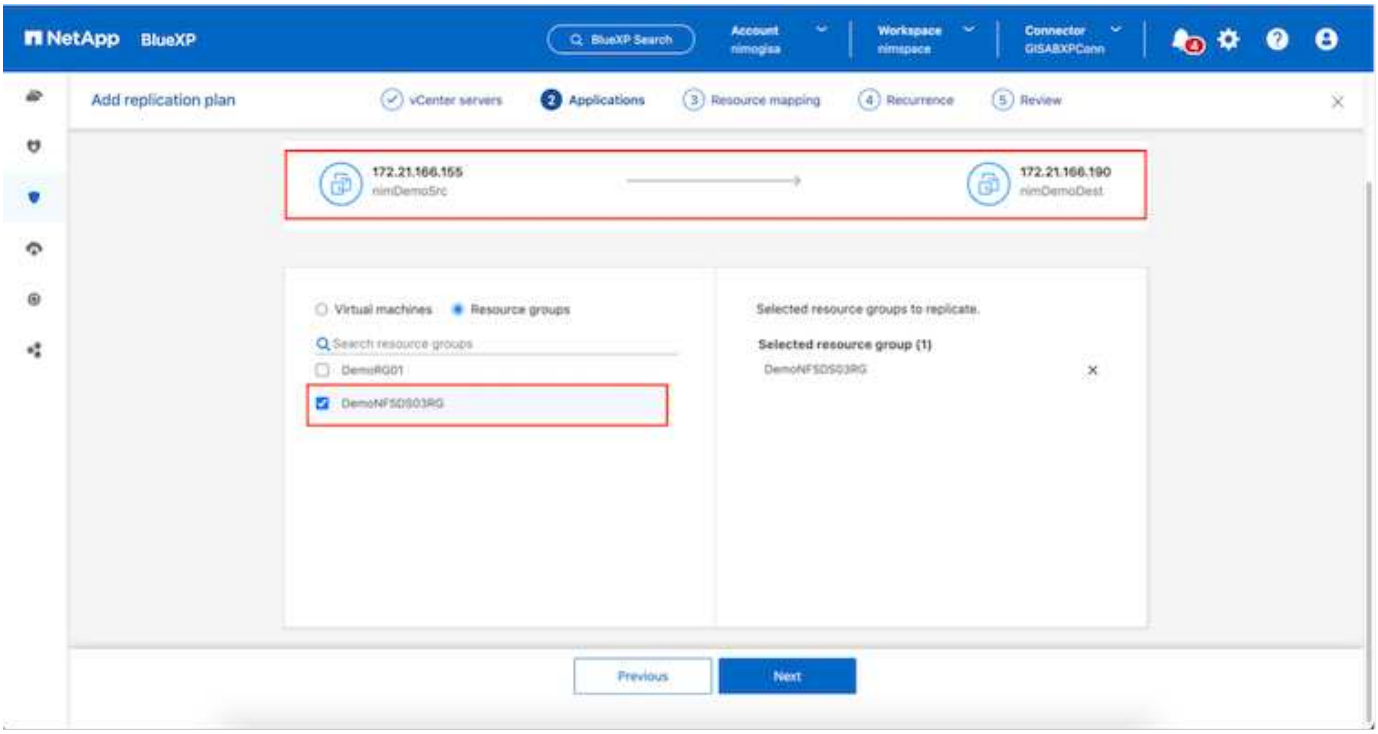
Configure the replication plan by selecting the source and destination vCenter platforms from the drop down and pick the resource groups to be included in the plan, along with the grouping of how applications should be restored and powered on and mapping of clusters and networks. To define the recovery plan, navigate to the **Replication Plan** tab and click **Add Plan**.

First, select the source vCenter and then select the destination vCenter.



The next step is to select existing resource groups. If no resource groups created, then the wizard helps to group the required virtual machines (basically create functional resource groups) based on the recovery objectives. This also helps define the operation sequence of how application virtual machines should be

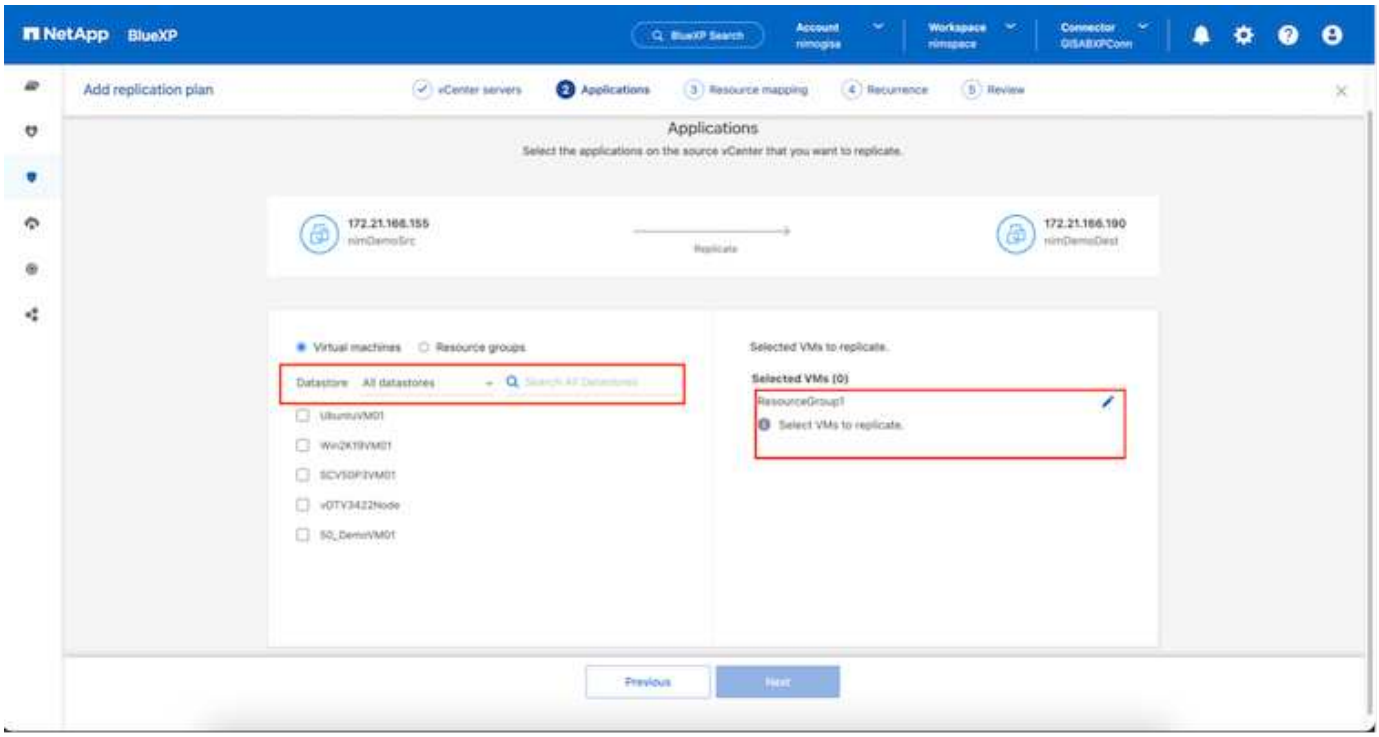
restored.



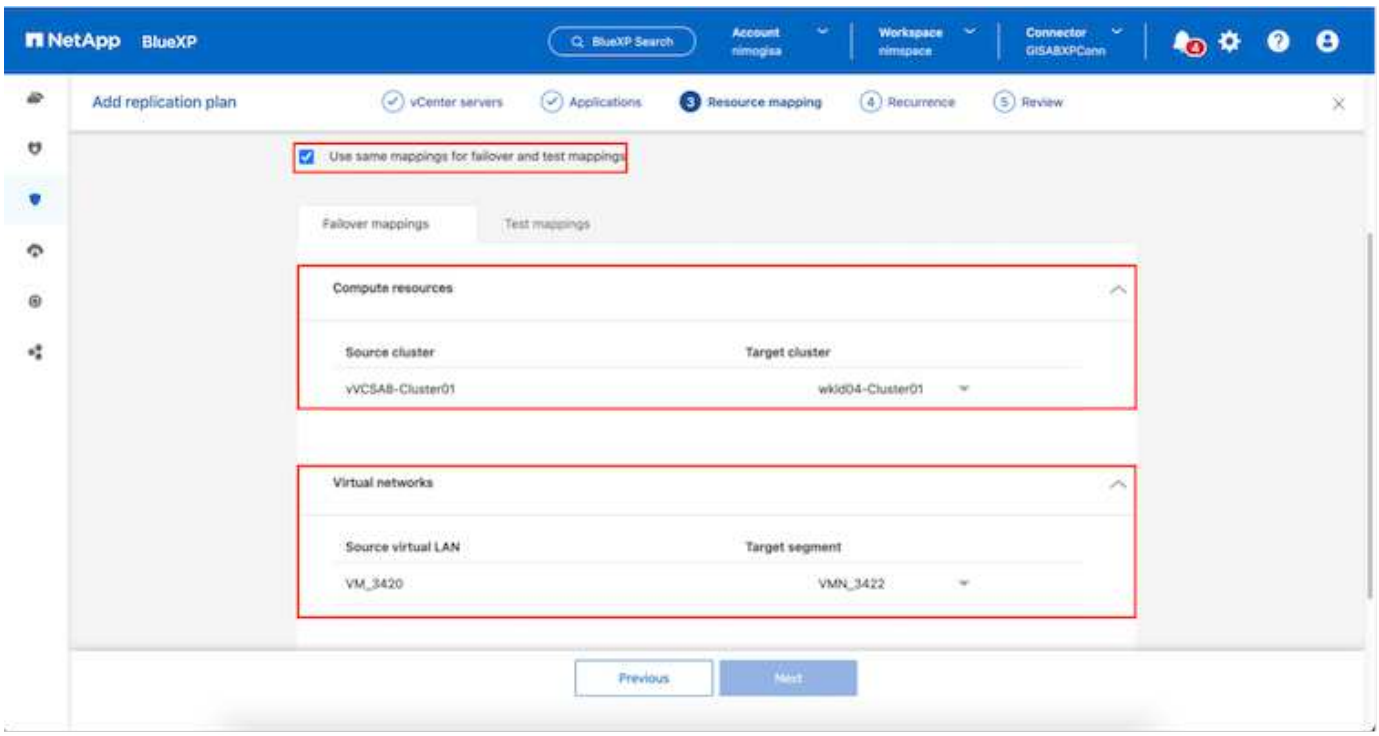
Resource group allows to set boot order using the drag and drop functionality. It can be used to easily modify the order in which the VMs would be powered on during the recovery process.

Each virtual machine within a resource group is started in sequence based on the order. Two resource groups are started in parallel.

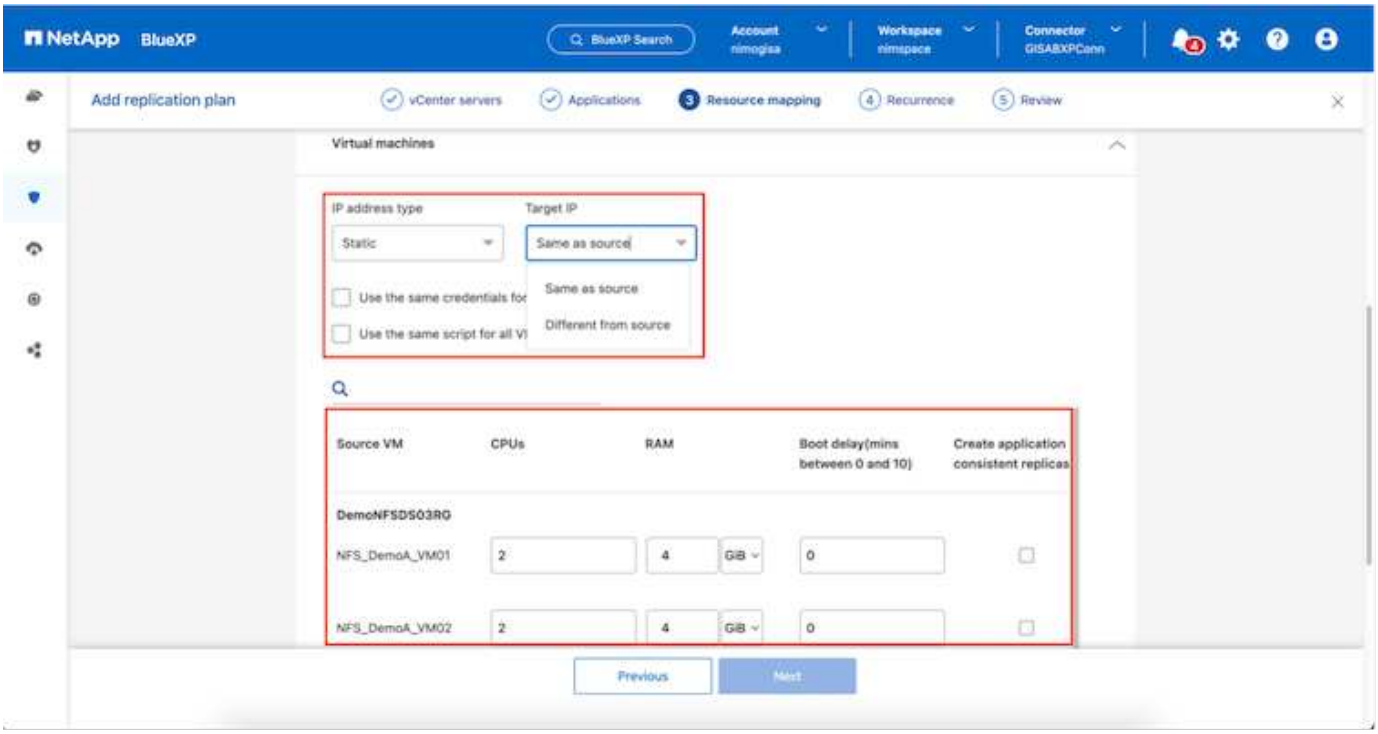
The below screenshot shows the option to filter virtual machines or specific datastores based on organizational requirements if resource groups are not created beforehand.



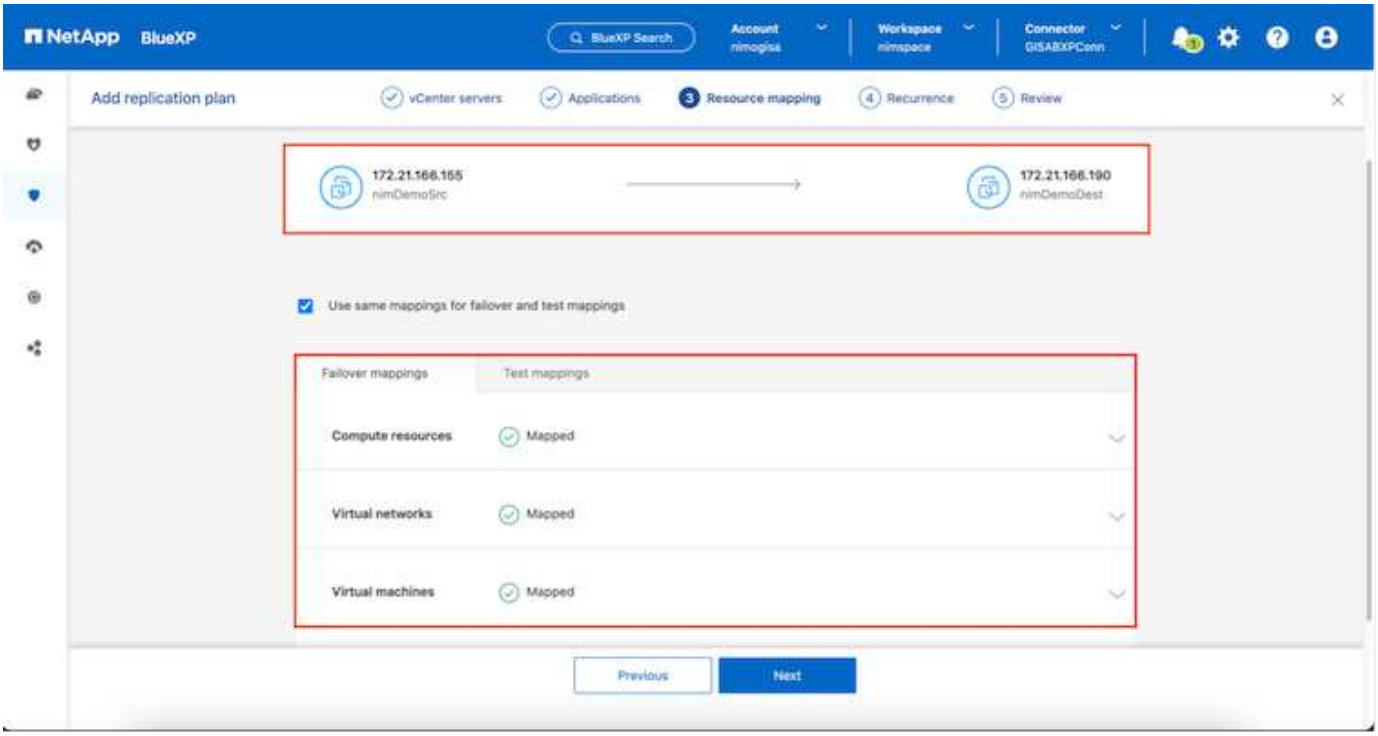
Once the resource groups are selected, create the failover mappings. In this step, specify how the resources from the source environment maps to the destination. This includes compute resources, virtual networks. IP customization, pre- and post-scripts, boot delays, application consistency and so on. For detailed information, refer to Create a replication plan.



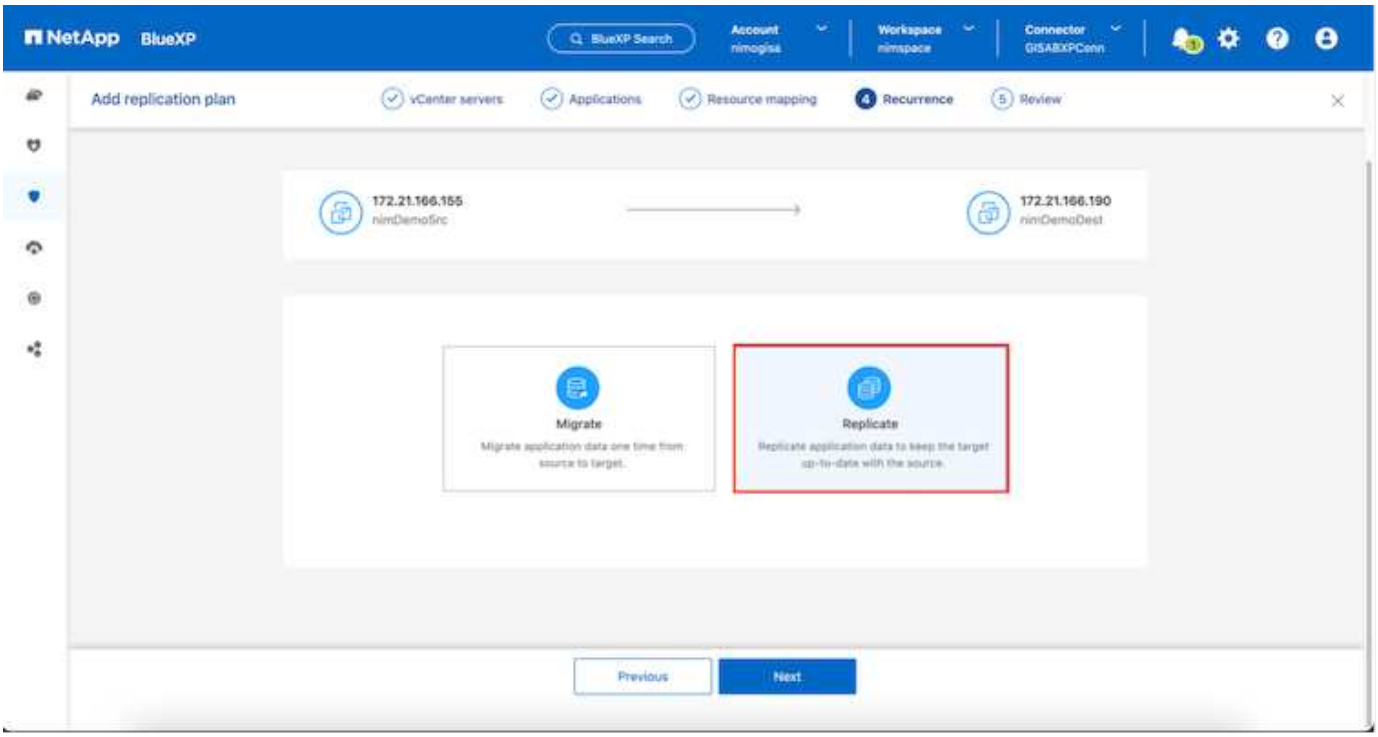
By default, same mapping parameters are used for both test and failover operations. To set different mappings for test environment, select the Test mapping option after unchecking the checkbox as shown below:



Once the resource mapping is complete, click Next.

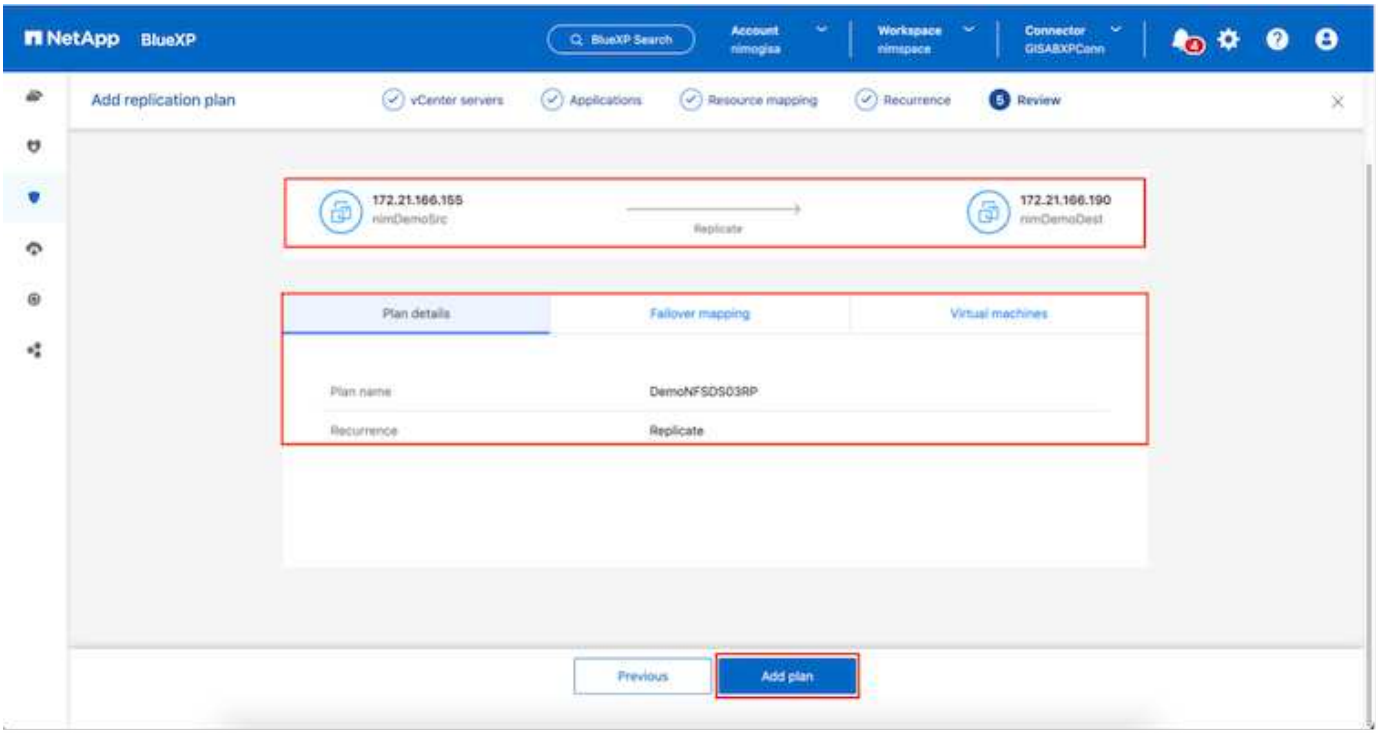


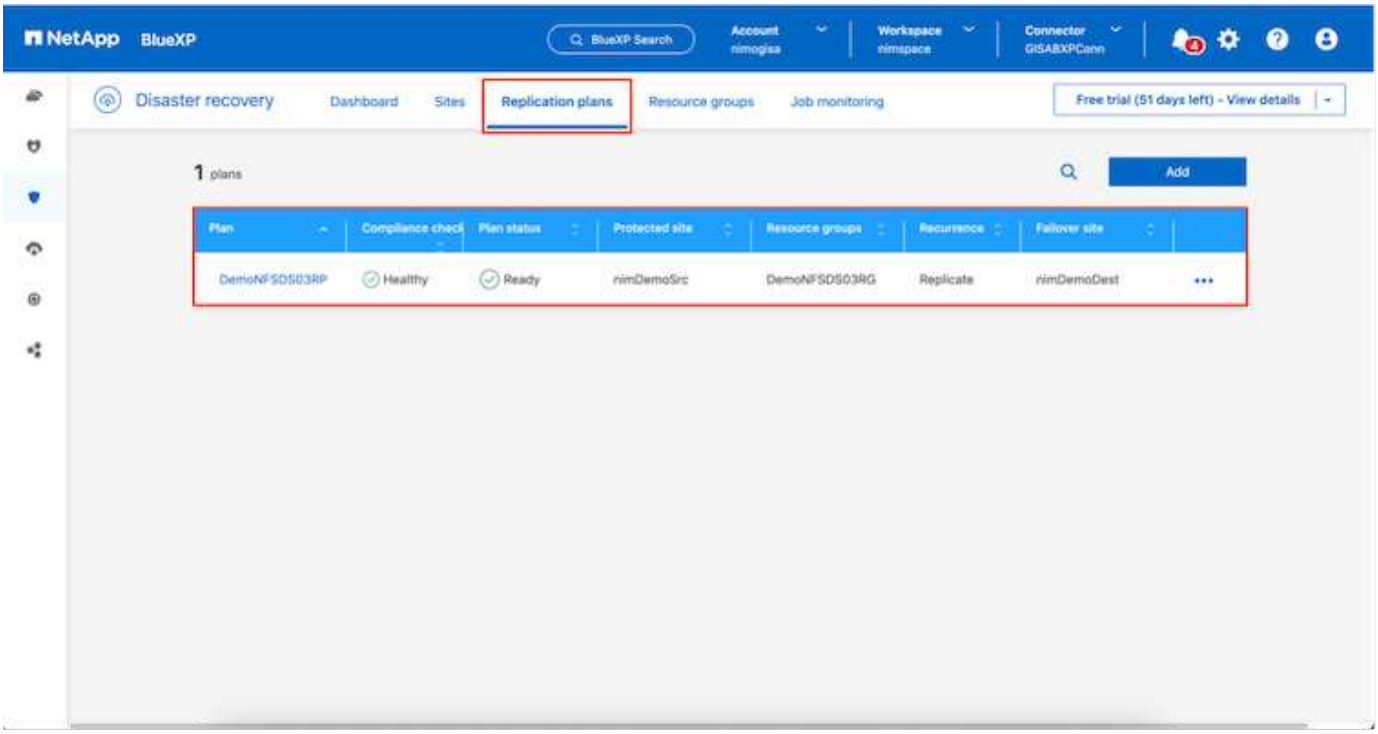
Select the recurrence type. In simple words, select Migrate (one time migration using failover) or recurring continuous replication option. In this walkthrough, Replicate option is selected.



Once done, review the created mappings and then click on **Add plan**.

VMs from different volumes and SVMs can be included in a replication plan. Depending on the VM placement (be it on same volume or separate volume within the same SVM, separate volumes on different SVMs), the Disaster Recovery creates a Consistency Group Snapshot.



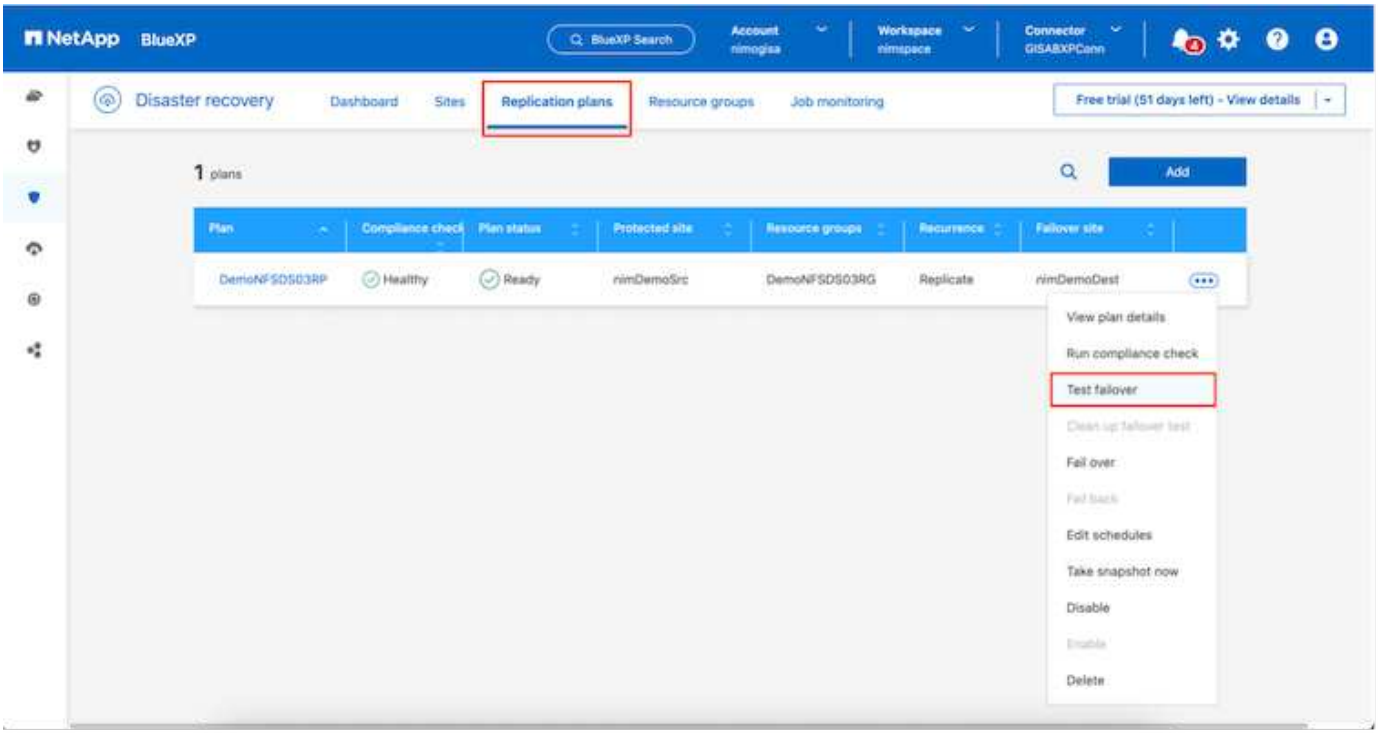


Disaster Recovery consists of the following workflows:

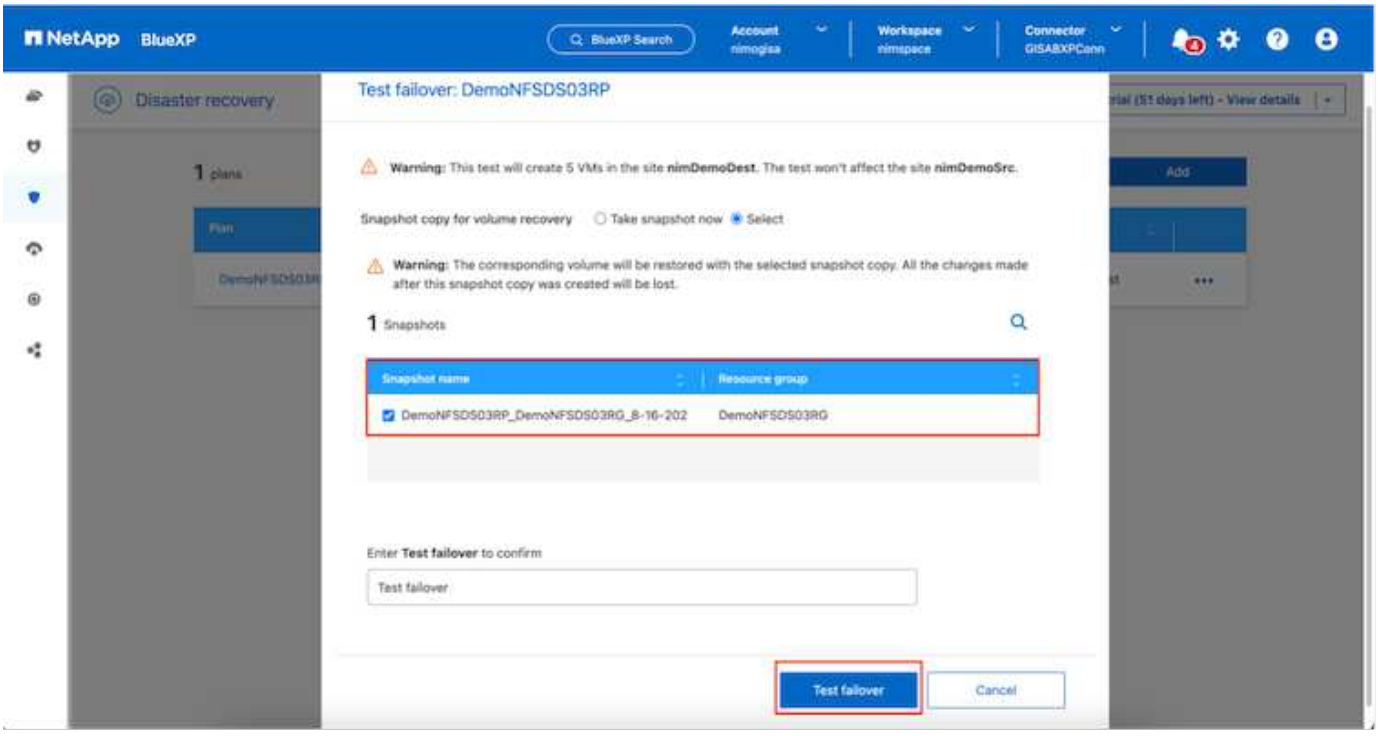
- Test failover (including periodic automated simulations)
- Cleanup failover test
- Failover
- Failback

Test failover

Test failover in Disaster Recovery is an operational procedure that allows VMware administrators to fully validate their recovery plans without disrupting their production environments. During a test failover, Disaster Recovery creates space-efficient NetApp FlexClone copies of the replicated NFS datastore volumes on the DR storage system. These FlexClone volumes are then mounted to the ESXi hosts at the DR site, and the VMs are powered on in an isolated test network. This process allows administrators to verify that their disaster recovery plans will function as intended in the event of an actual disaster.



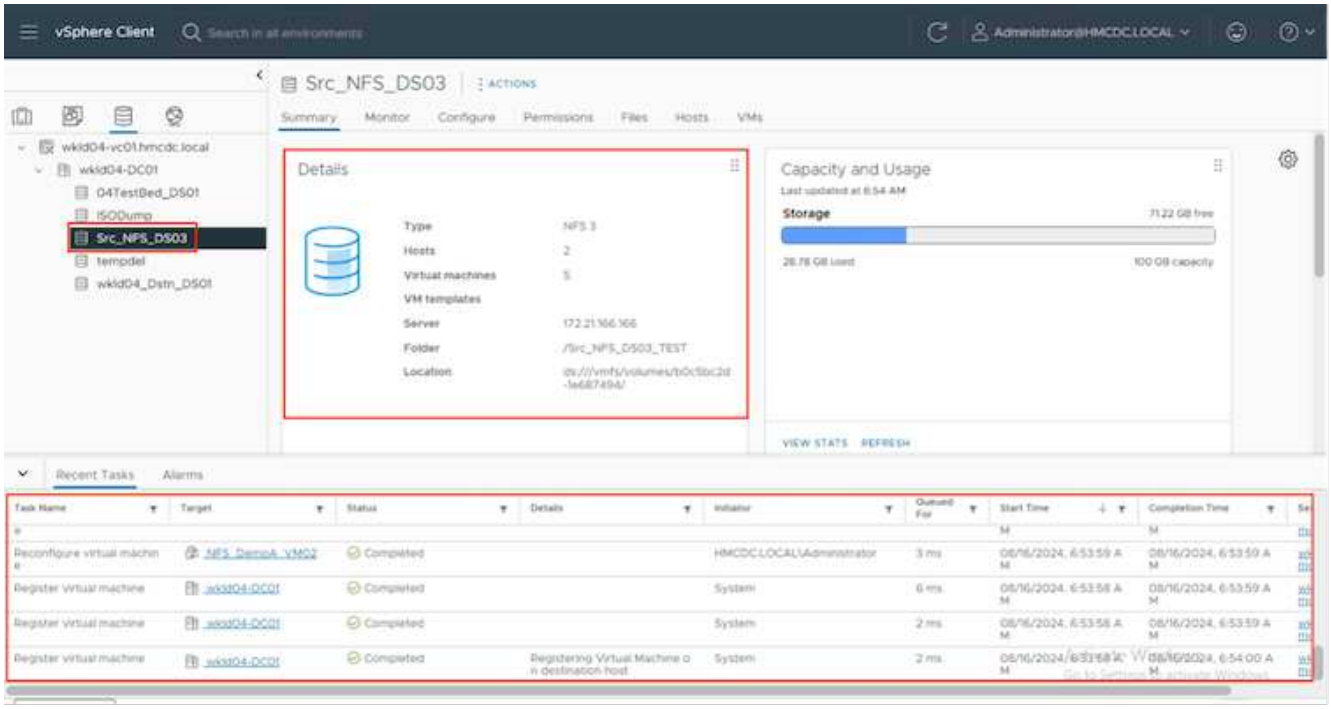
Disaster Recovery incorporates the ability to select the snapshot as an optional capability in the test failover operation. This capability allows the VMware administrator to verify that any changes that were recently made in the environment are replicated to the destination site and thus are present during the test. Such changes include patches to the VM guest operating system



When the VMware administrator runs a test failover operation, Disaster Recovery automates the following tasks:

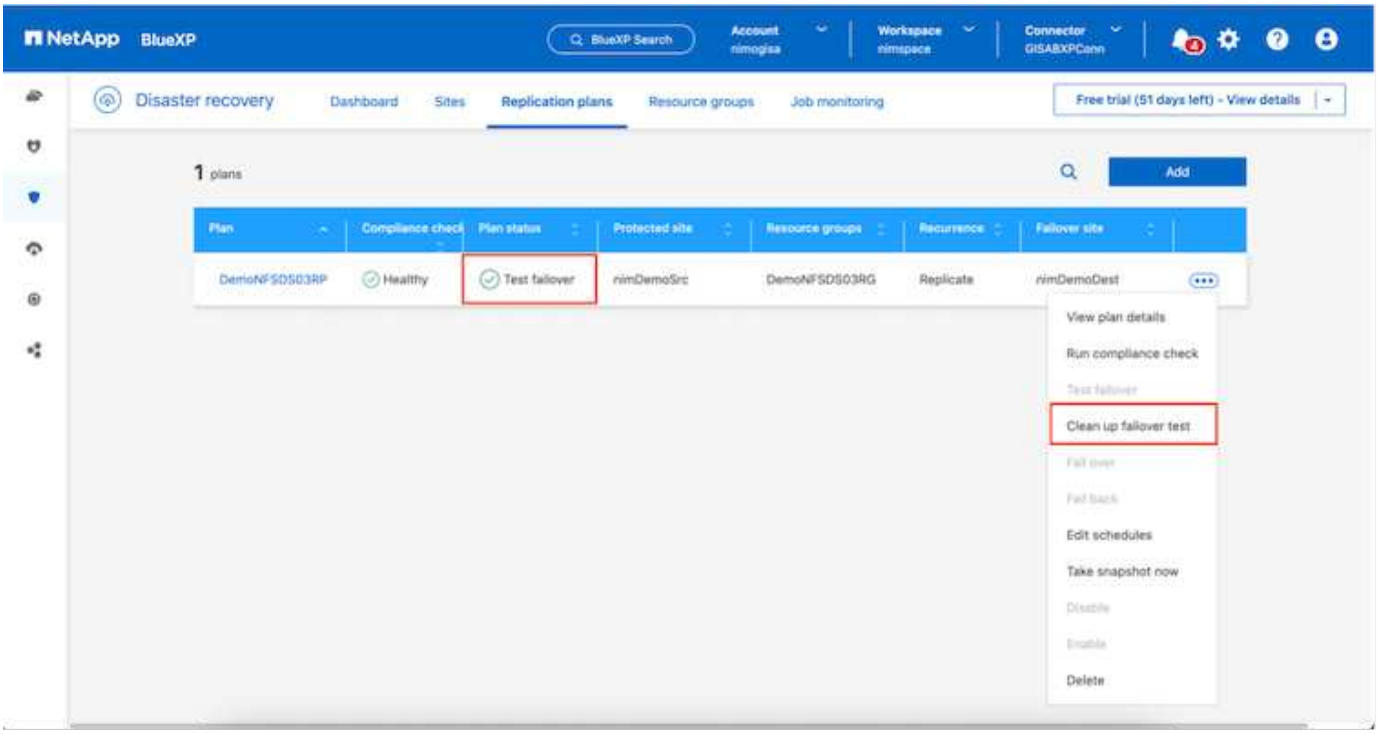
- Triggering SnapMirror relationships to update storage at the destination site with any recent changes that were made at the production site.

- Creating NetApp FlexClone volumes of the FlexVol volumes on the DR storage array.
- Connecting the NFS datastores in the FlexClone volumes to the ESXi hosts at the DR site.
- Connecting the VM network adapters to the test network specified during the mapping.
- Reconfiguring the VM guest operating system network settings as defined for the network at the DR site.
- Executing any custom commands that have been stored in the replication plan.
- Powering on the VMs in the order that is defined in the replication plan.



Cleanup failover test Operation

The cleanup failover test operation occurs after the replication plan test has been completed and the VMware administrator responds to the cleanup prompt.



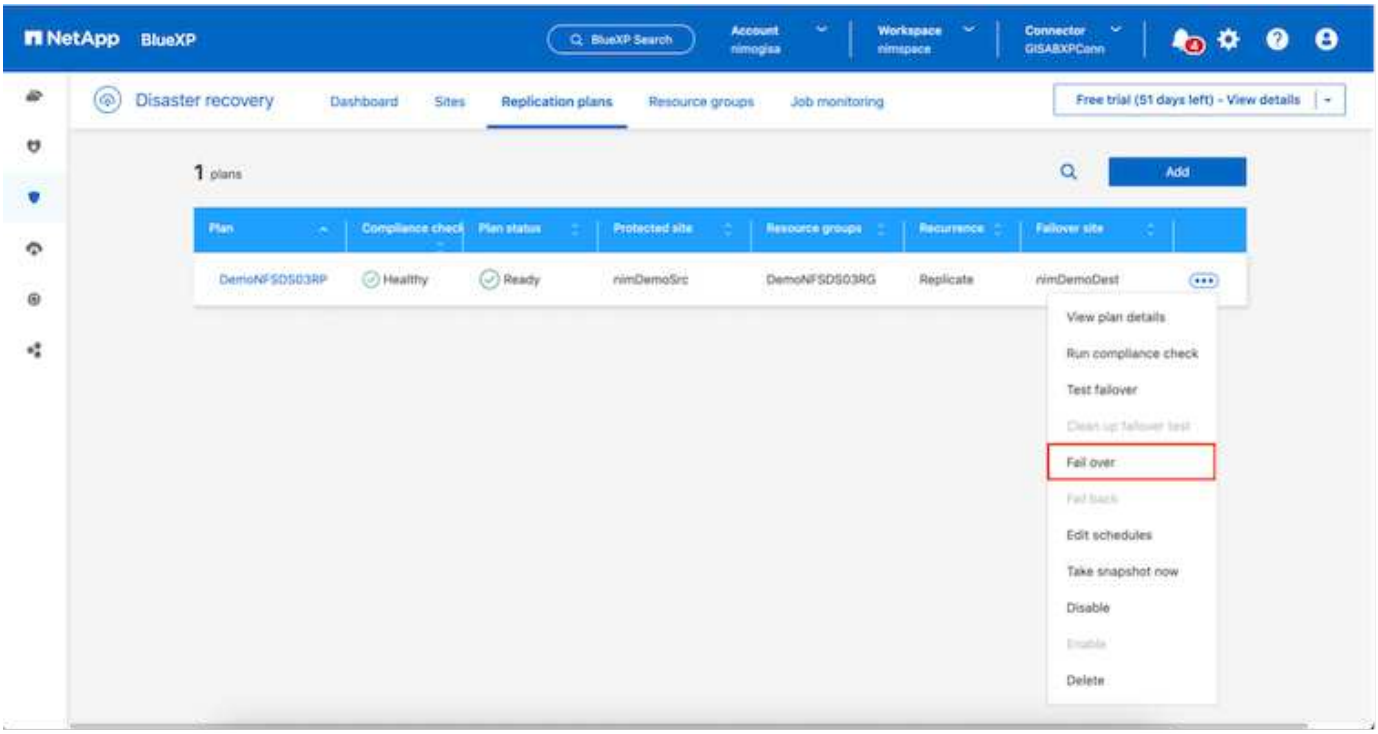
This action will reset the virtual machines (VMs) and the status of the replication plan to the ready state.

When the VMware administrator performs a recovery operation, Disaster Recovery completes the following process:

1. It powers off each recovered VM in the FlexClone copy that was used for testing.
2. It deletes the FlexClone volume that was used to present the recovered VMs during the test.

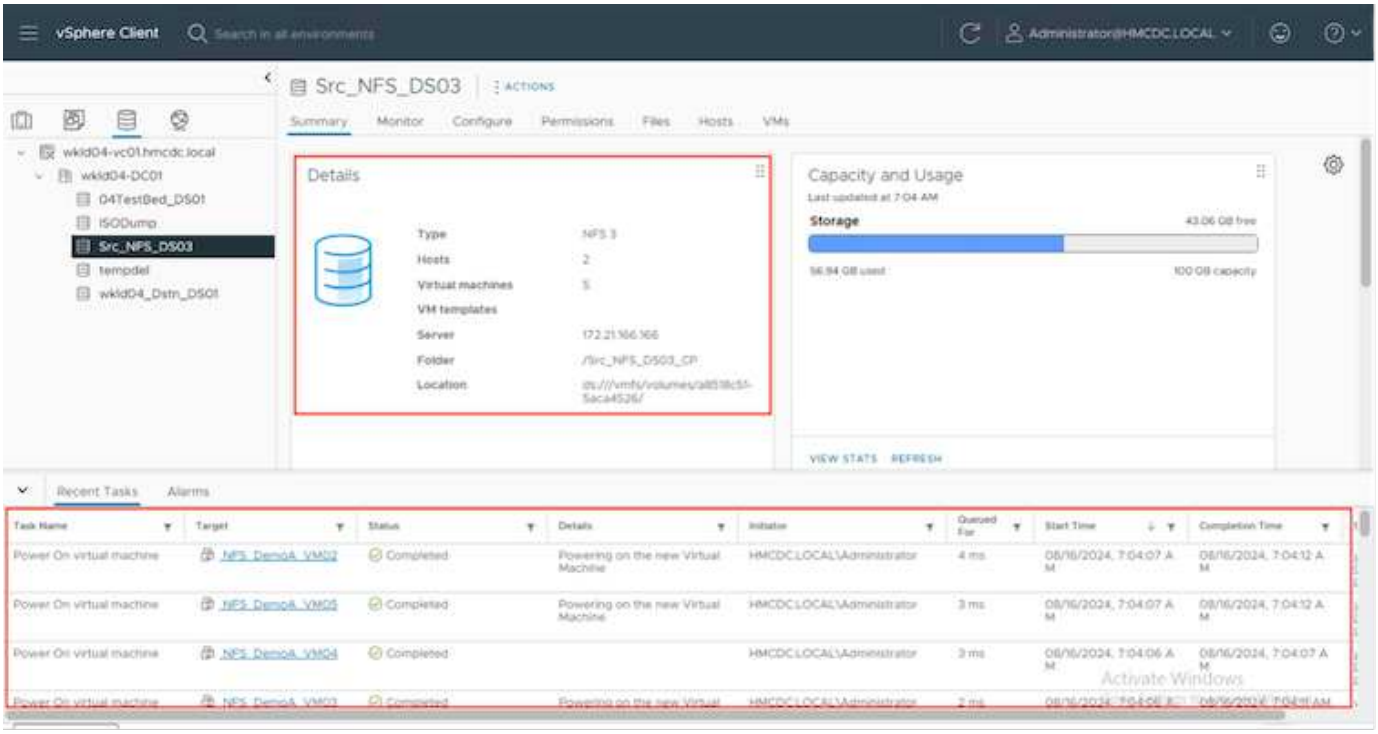
Planned Migration and Fail over

Disaster Recovery has two methods for performing a real failover: planned migration and fail over. The first method, planned migration, incorporates VM shutdown and storage replication synchronization into the process to recover or effectively move the VMs to the destination site. Planned migration requires access to the source site. The second method, failover, is an planned/unplanned failover in which the VMs are recovered at the destination site from the last storage replication interval that was able to complete. Depending on the RPO that was designed into the solution, some amount of data loss can be expected in the DR scenario.



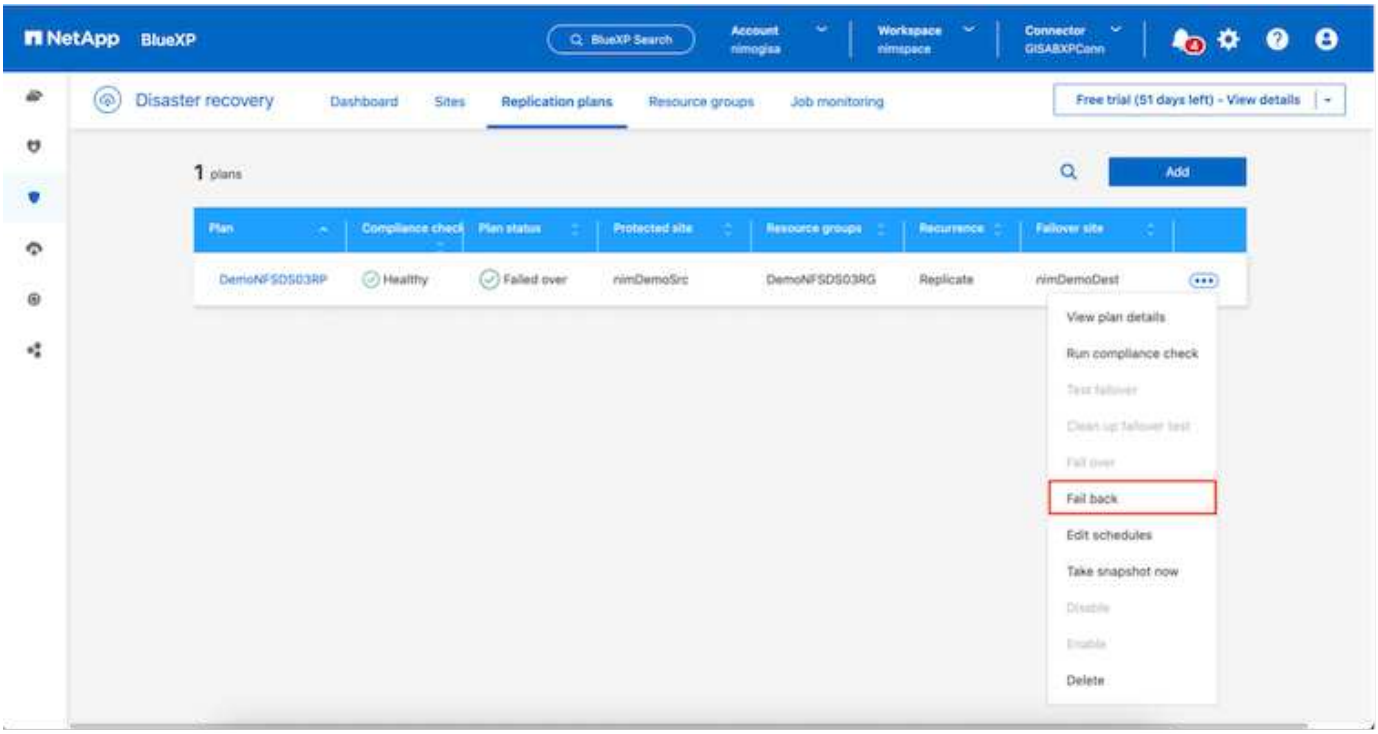
When the VMware administrator performs a failover operation, Disaster Recovery automates the following tasks:

- Break and fail over the NetApp SnapMirror relationships.
- Connect the replicated NFS datastores to the ESXi hosts at the DR site.
- Connect the VM network adapters to the appropriate destination site network.
- Reconfigure the VM guest operating system network settings as defined for the network at the destination site.
- Execute any custom commands (if any) that have been stored in the replication plan.
- Power on the VMs in the order that was defined in the replication plan.



Failback

A failback is an optional procedure that restores the original configuration of the source and destination sites after a recovery.



VMware administrators can configure and run a failback procedure when they are ready to restore services to the original source site.

NOTE: Disaster Recovery replicates (resyncs) any changes back to the original source virtual machine before reversing the replication direction. This process starts from a relationship that has completed failing over to a

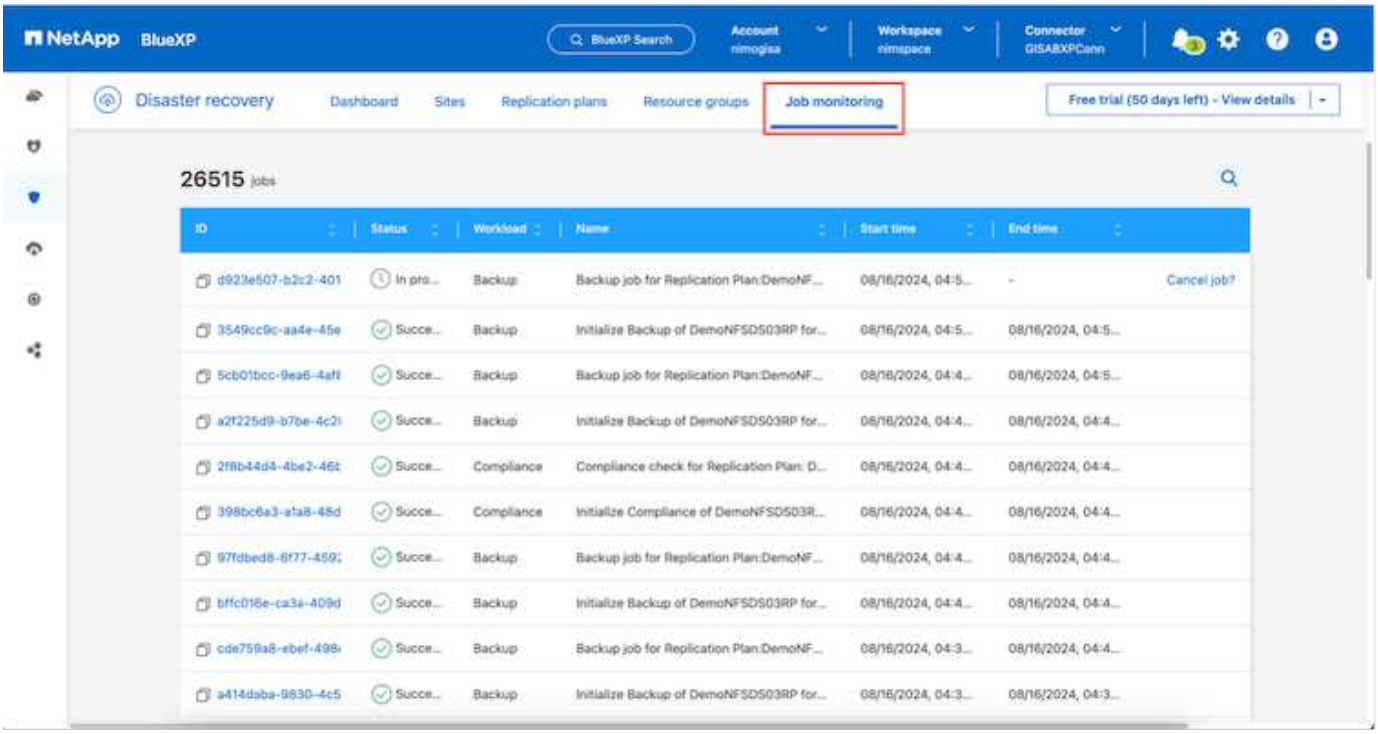
target and involves the following steps:

- Power off and unregister the virtual machines and volumes on the destination site are unmounted.
- Break the SnapMirror relationship on the original source is broken to make it read/write.
- Resynchronize the SnapMirror relationship to reverse the replication.
- Mount the volume on the source, power on and register the source virtual machines.

For more details about accessing and configuring Disaster Recovery, see the Learn about NetApp Disaster Recovery for VMware.

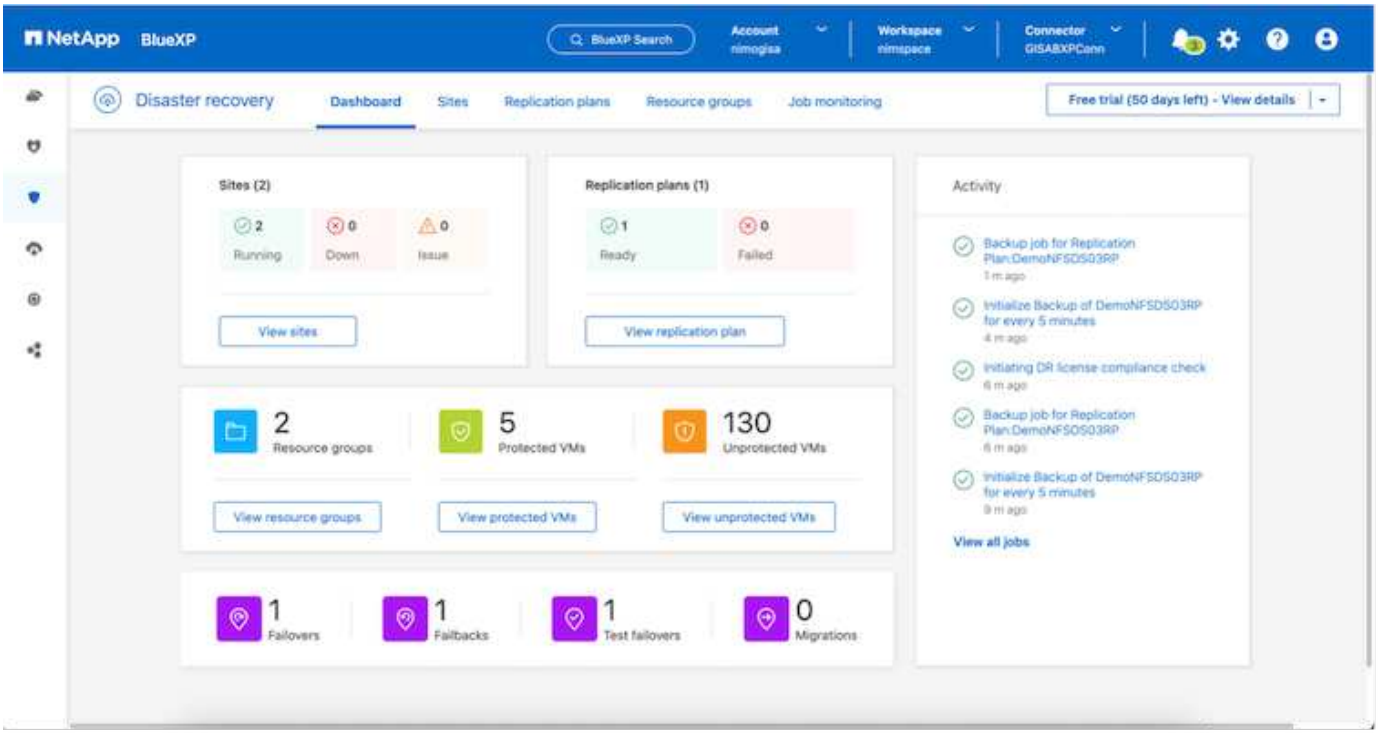
Monitoring and Dashboard

From Console or the ONTAP CLI, you can monitor the replication health status for the appropriate datastore volumes, and the status of a failover or test failover can be tracked via Job Monitoring.



If a job is currently in progress or queued, and you wish to stop it, there is an option to cancel it.

With the Disaster Recovery dashboard, confidently evaluate the status of disaster recovery sites and replication plans. This enables administrators to swiftly identify healthy, disconnected, or degraded sites and plans.



This provides a powerful solution to handle a tailored and customized disaster recovery plan. Failover can be done as planned failover or failover with a click of a button when disaster occurs and decision is made to activate the DR site.

To learn more about this process, feel free to follow the detailed walkthrough video or use the solution simulator.

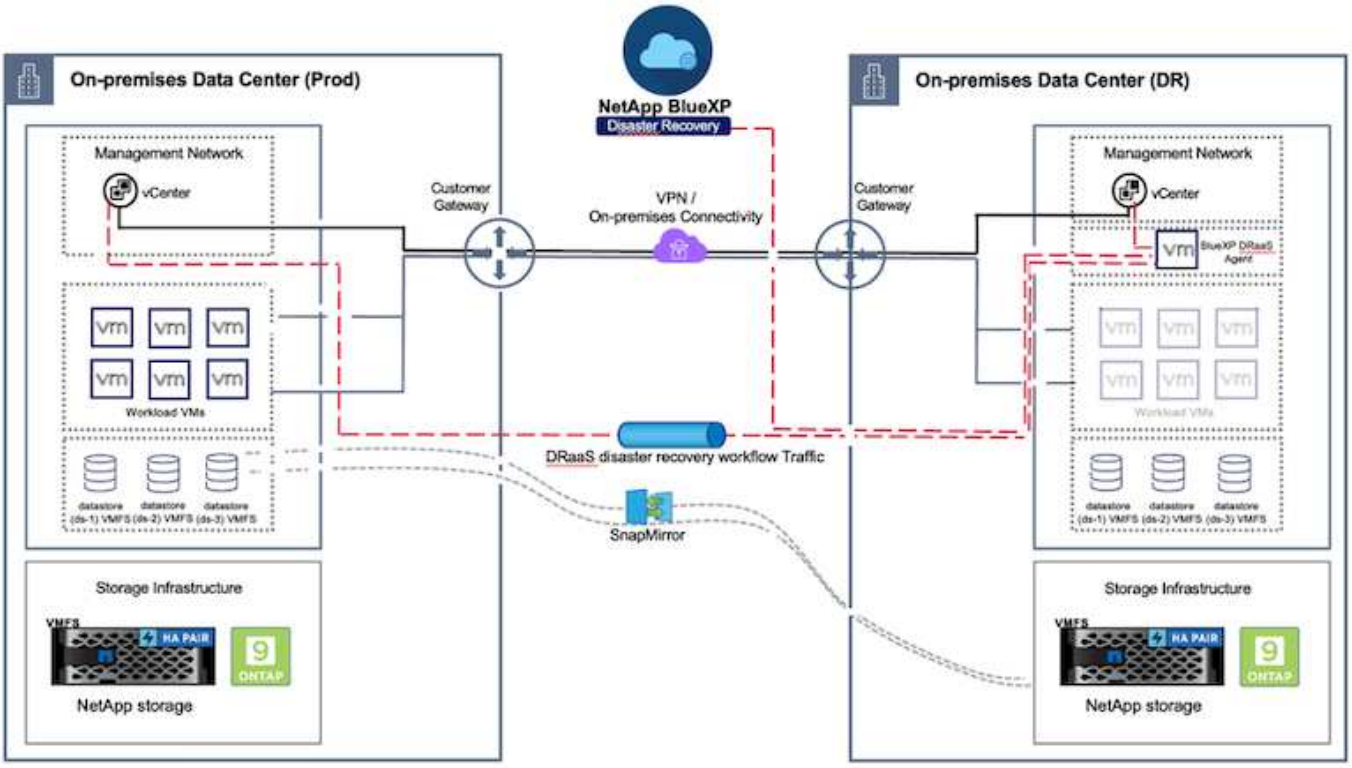
NetApp Disaster Recovery for VMFS Datastores

Disaster recovery using block-level replication from production site to disaster recovery site is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events, like ransomware attacks. With NetApp SnapMirror replication, VMware workloads running on-premises ONTAP systems using VMFS datastore can be replicated to another ONTAP storage system in a designated recovery datacenter where VMware resides

This section of the document describes the configuration of NetApp Diaster Recovery to set up disaster recovery for on-premises VMware VMs to another designated site. As part of this setup, Console agent, onboarding the ONTAP arrays, Defining Sites and associating VMware vCenter. In addition, this document details how to configure replication between sites and how to setup and test a recovery plan. The last section has instructions for performing a full site failover and how to failback when the primary site is recovered and bought online.

Utilizing the Disaster Recovery, part of NetApp Data Services,, customers can discover their on-premises VMware vCenters along with ONTAP storage, create resource groupings, create a disaster recovery plan, associate it with resource groups, and test or execute failover and failback. SnapMirror provides storage-level block replication to keep the two sites up to date with incremental changes, resulting in a RPO of up to 5 minutes. It is also possible to simulate DR procedures as a regular drill without impacting the production and replicated datastores or incurring additional storage costs. NetApp Data Services takes advantage of ONTAP's FlexClone technology to create a space-efficient copy of the VMFS datastore from the last replicated Snapshot on the DR site. Once the DR test is complete, customers can simply delete the test environment, again without

any impact to actual replicated production resources. When there is a need (planned or unplanned) for actual failover, with a few clicks, the Disaster Recovery will orchestrate all the steps needed to automatically bring up the protected virtual machines on designated disaster recovery site. The service will also reverse the SnapMirror relationship to the primary site and replicate any changes from secondary to primary for a failback operation, when needed. All of these can be achieved with a fraction of cost compared to other well-known alternatives.

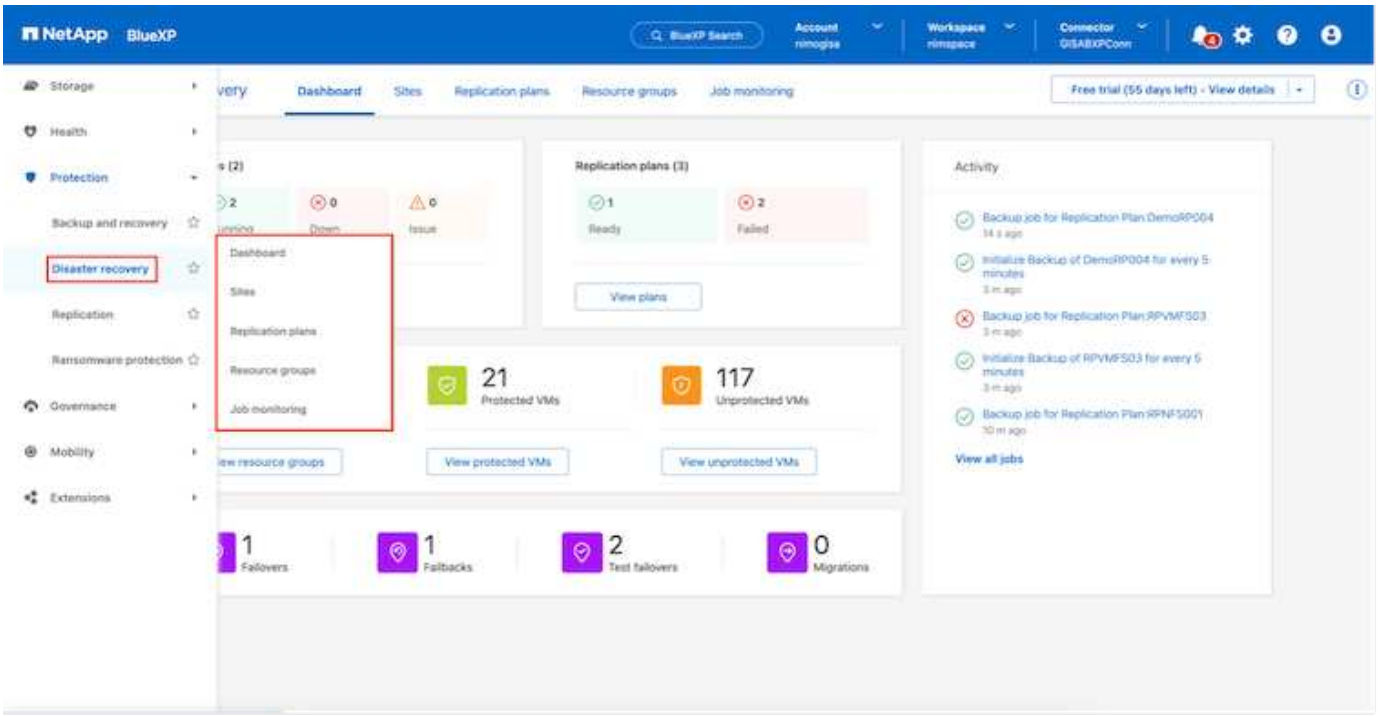


For up-to-date information and supported versions of vCenter Server, ESXi, and ONTAP for Disaster Recovery, refer to the Disaster Recovery supported configurations.

Getting started

To get started with Disaster Recovery, use NetApp Console and then access the service.

1. Log in to Console.
2. From the Console left navigation, select Protection > Disaster recovery.
3. The Disaster Recovery Dashboard appears.

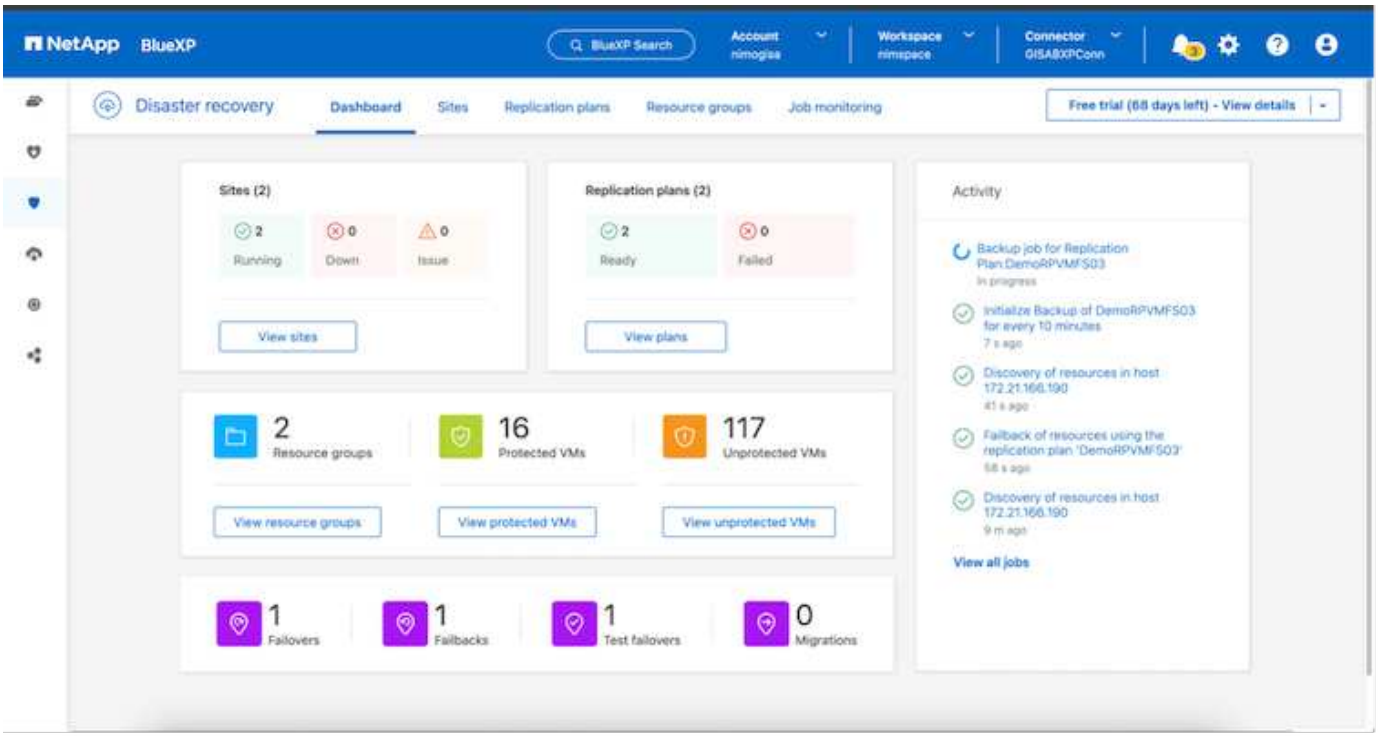


Before configuring disaster recovery plan, ensure the following pre-requisites are met:

- Console agent is deployed in AWS VPC.
- Console agent instance have connectivity to the source and destination vCenter and storage systems.
- On-premises NetApp storage systems hosting VMFS datastores for VMware are added in Console.
- DNS resolution should be in place when using DNS names. Otherwise, use IP addresses for the vCenter.
- SnapMirror replication is configured for the designated VMFS based datastore volumes.

Once the connectivity is established between the source and destination sites, proceed with configuration steps, which should take about 3 to 5 minutes.

NetApp recommends deploying the Console agent in the disaster recovery site or in a third site, so that the Console agent can communicate through the network with source and destination resources during real outages or natural disasters.

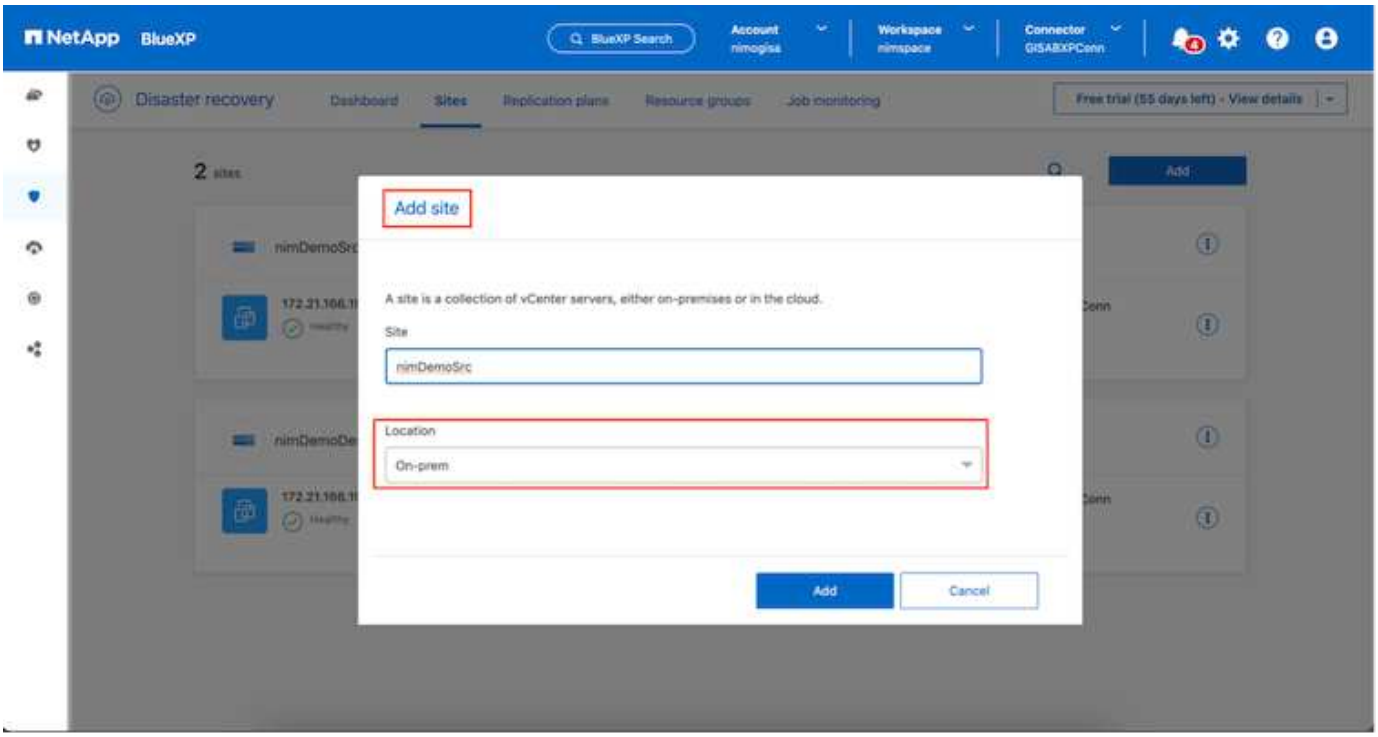


Support for on-premises to on-premises VMFS datastores is in technology preview while writing this document. The capability is supported with both FC and ISCSI protocol based VMFS datastores.

Disaster Recovery configuration

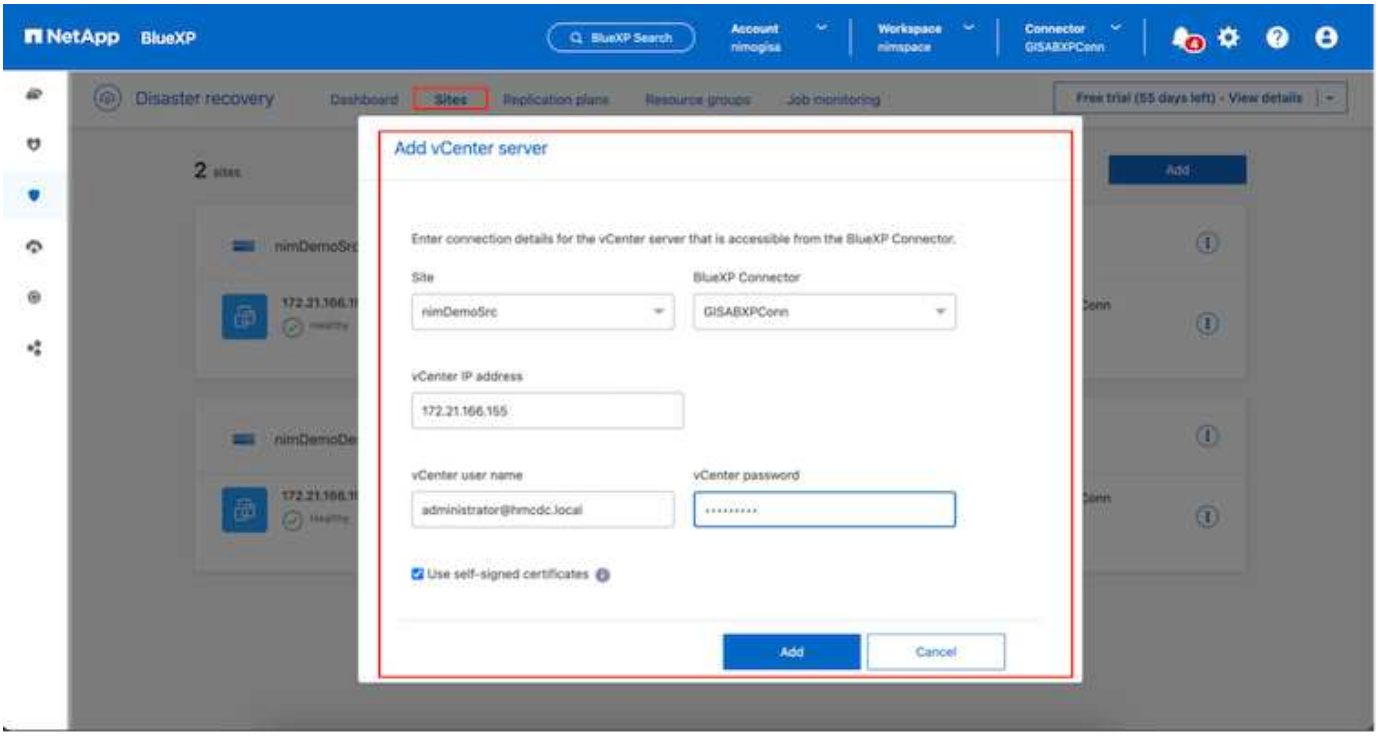
The first step in preparing for disaster recovery is to discover and add the on-premises vCenter and storage resources to Disaster Recovery.

Ensure the ONTAP storage systems are added to the working environment within the canvas. Open Console and select **Protection > Disaster Recovery** from left navigation. Select **Discover vCenter servers** or use top menu, Select **Sites > Add > Add vCenter**.

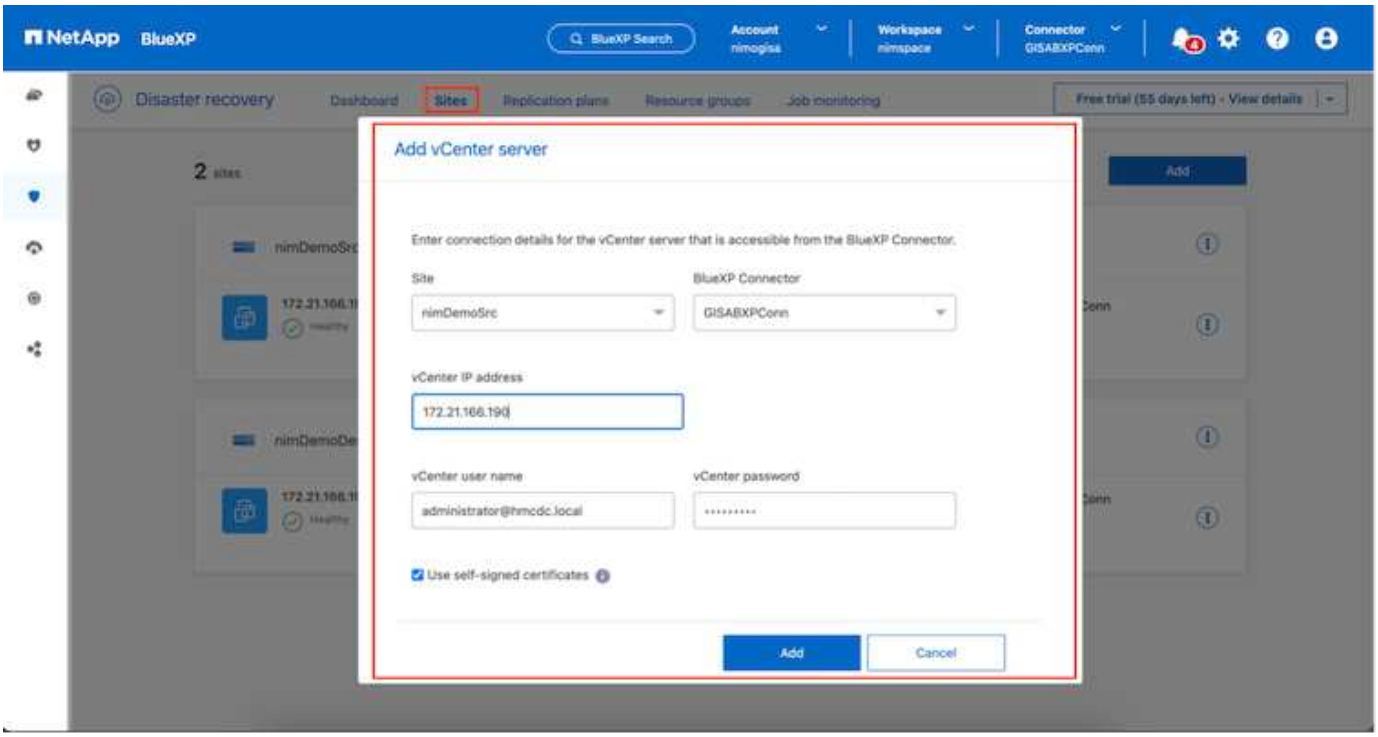


Add the following platforms:

- **Source**. On-premises vCenter.



- **Destination**. VMC SDDC vCenter.



Once the vCenters are added, automated discovery is triggered.

Configuring Storage replication between source and destination site

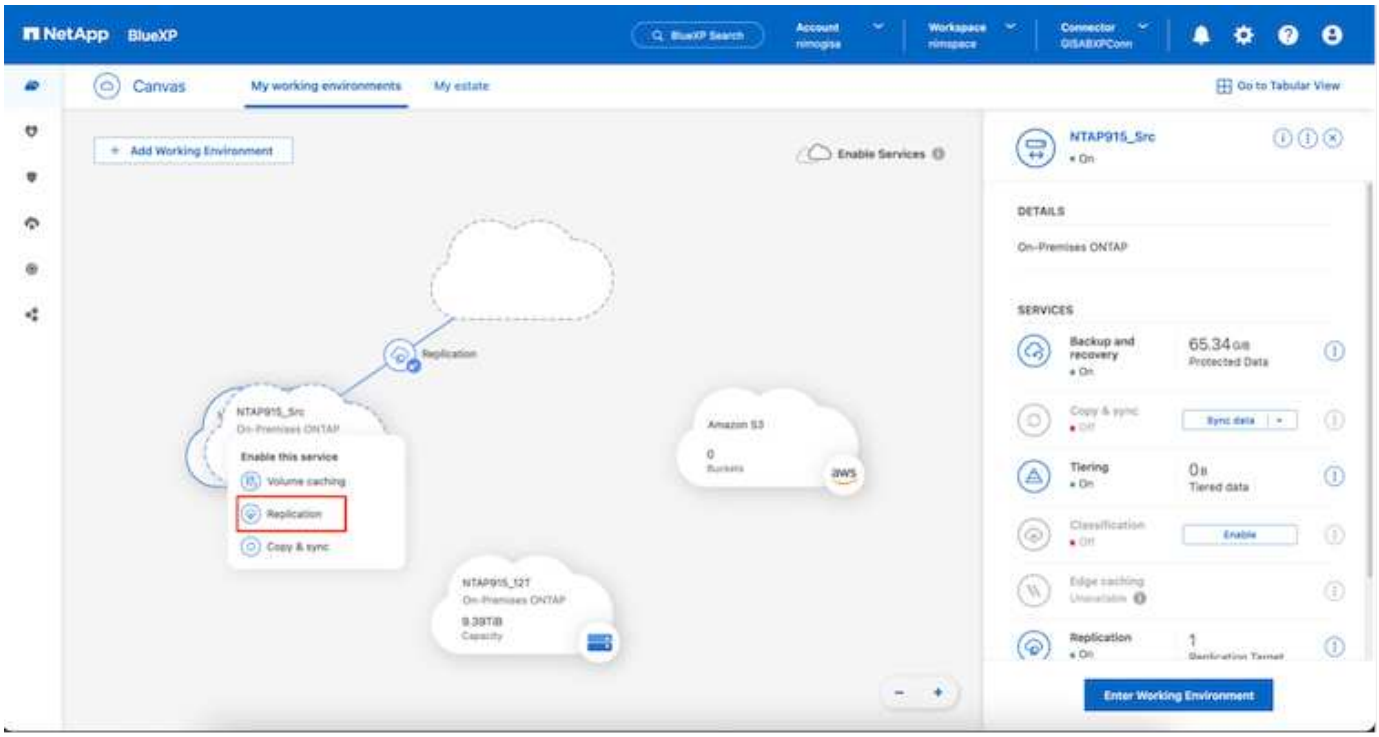
SnapMirror makes use of ONTAP snapshots to manage the transfer of data from one location to another. Initially, a full copy based on a snapshot of the source volume is copied over to the destination to perform a baseline synchronization. As data changes occur at the source, a new snapshot is created and compared to the baseline snapshot. The blocks found to have changed are then replicated to the destination, with the newer snapshot becoming the current baseline, or newest common snapshot. This enables the process to be repeated and incremental updates to be sent to the destination.

When a SnapMirror relationship has been established, the destination volume is in an online read-only state, and so is still accessible. SnapMirror works with physical blocks of storage, rather than at a file or other logical level. This means that the destination volume is an identical replica of the source, including snapshots, volume settings, etc. If ONTAP space efficiency features, such as data compression and data deduplication, are being used by the source volume, the replicated volume will retain these optimizations.

Breaking the SnapMirror relationship makes the destination volume writable and would typically be used to perform a failover when SnapMirror is being used to synchronize data to a DR environment. SnapMirror is sophisticated enough to allow the data changed at the failover site to be efficiently resynchronized back to the primary system, should it later come back online, and then allow for the original SnapMirror relationship to be re-established.

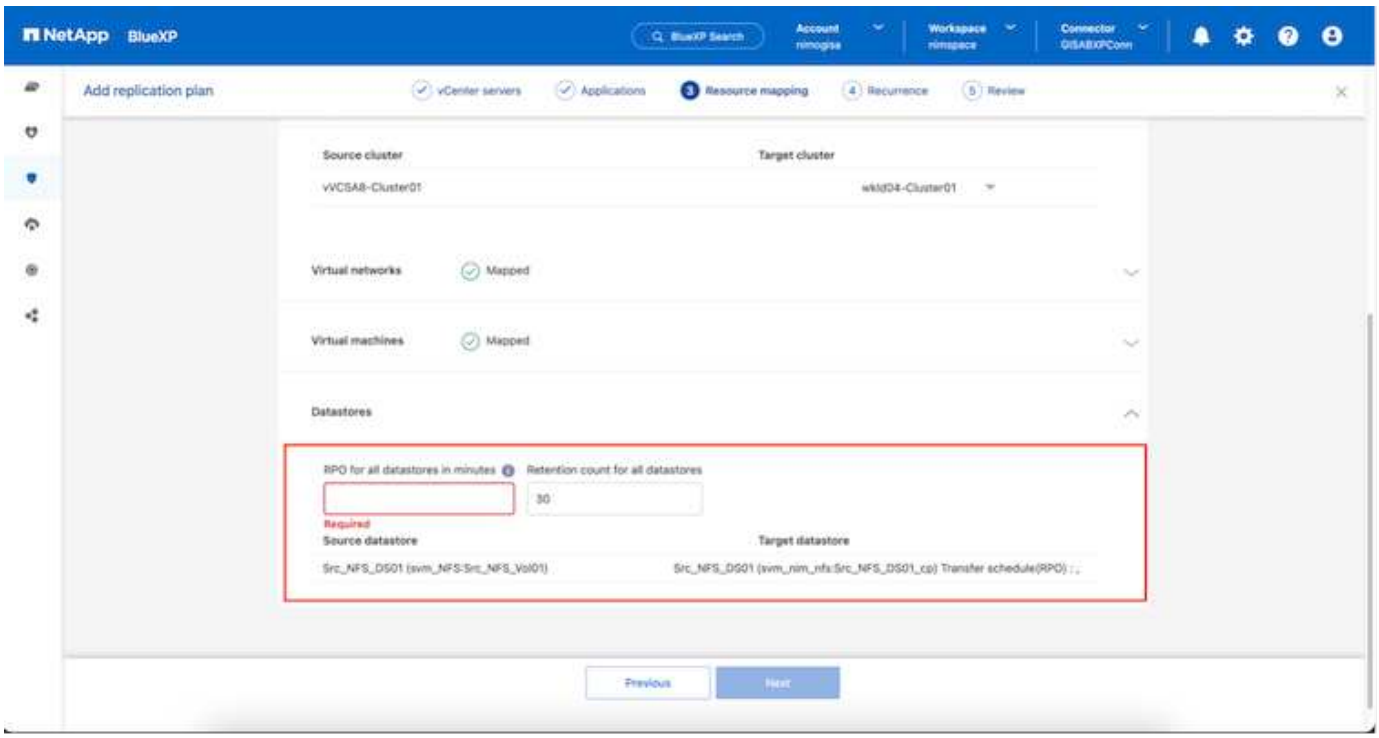
How to set it up for VMware Disaster Recovery

The process to create SnapMirror replication remains the same for any given application. The process can be manual or automated. The easiest way is to leverage Console to configure SnapMirror replication by using simple drag & drop of the source ONTAP system in the environment onto the destination to trigger the wizard that guides through the rest of the process.



NetApp Data Services can also automate the same provided the following two criteria's are met:

- Source and destination clusters have a peer relationship.
- Source SVM and destination SVM have a peer relationship.

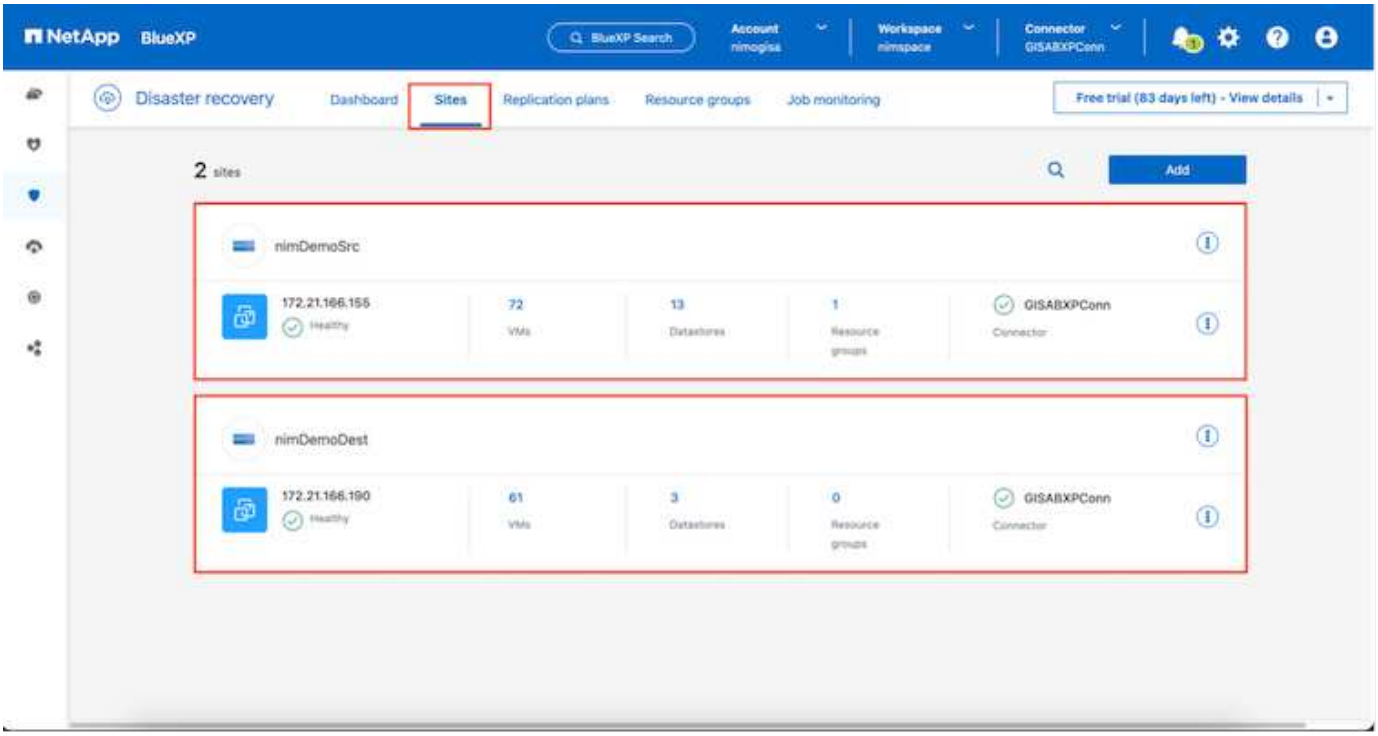


If SnapMirror relationship is already configured for the volume via CLI, Console picks up the relationship and continues with the rest of the workflow operations.

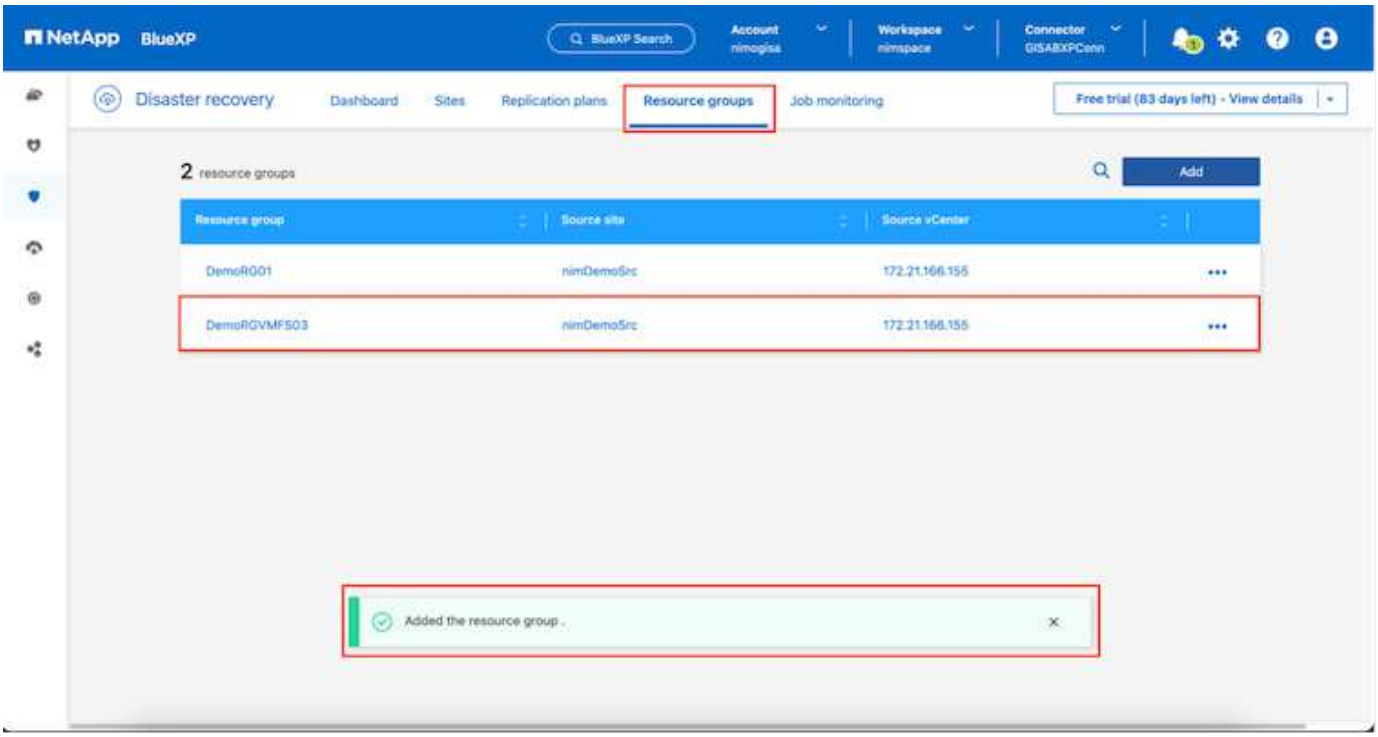
Apart from the above approaches, SnapMirror replication can also be created via ONTAP CLI or System Manager. Irrespective of the approach used to synchronize the data using SnapMirror, Disaster Recovery orchestrates the workflow for seamless and efficient disaster recovery operations.

What can Disaster Recovery do for you?

After the source and destination sites are added, Disaster Recovery performs automatic deep discovery and displays the VMs along with associated metadata. Disaster Recovery also automatically detects the networks and port groups used by the VMs and populates them.

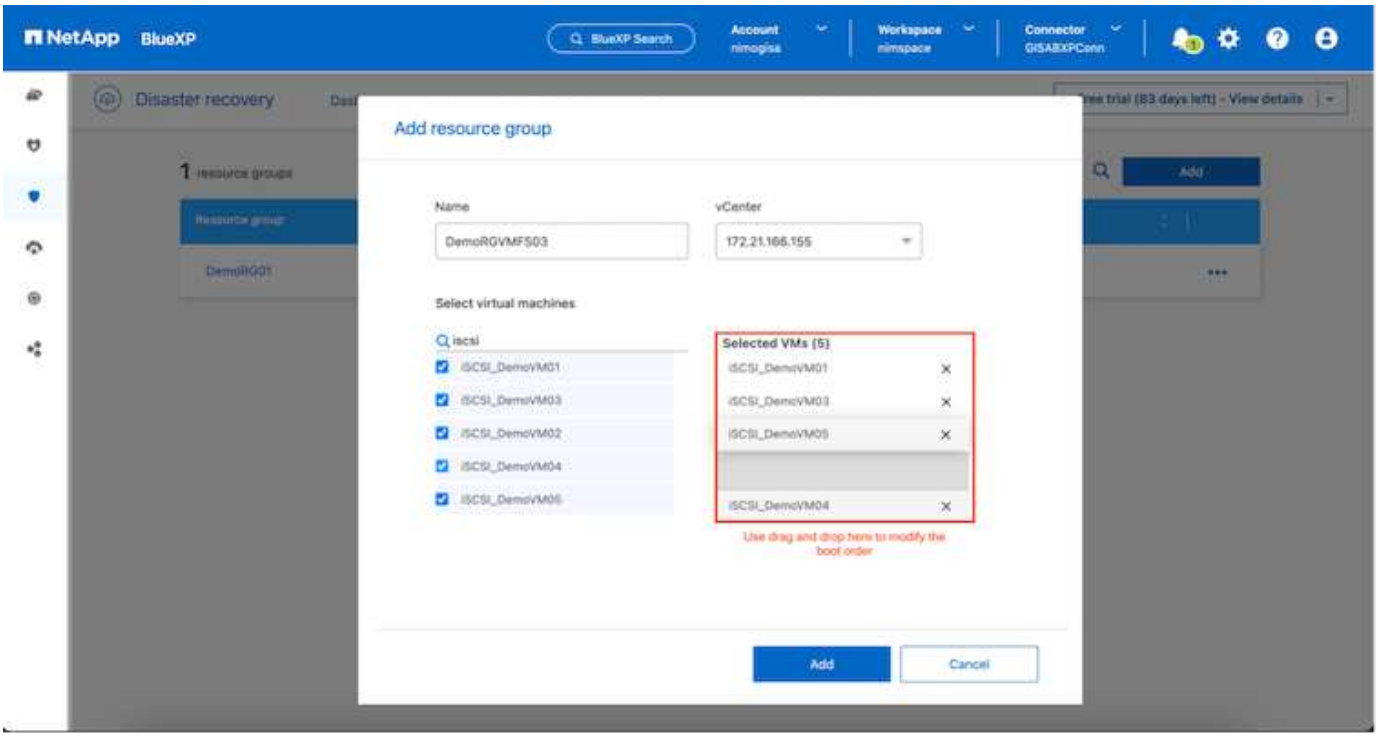


After the sites have been added, VMs can be grouped into resource groups. Disaster Recovery resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders and boot delays that can be executed upon recovery. To start creating resource groups, navigate to **Resource Groups** and click **Create New Resource Group**.

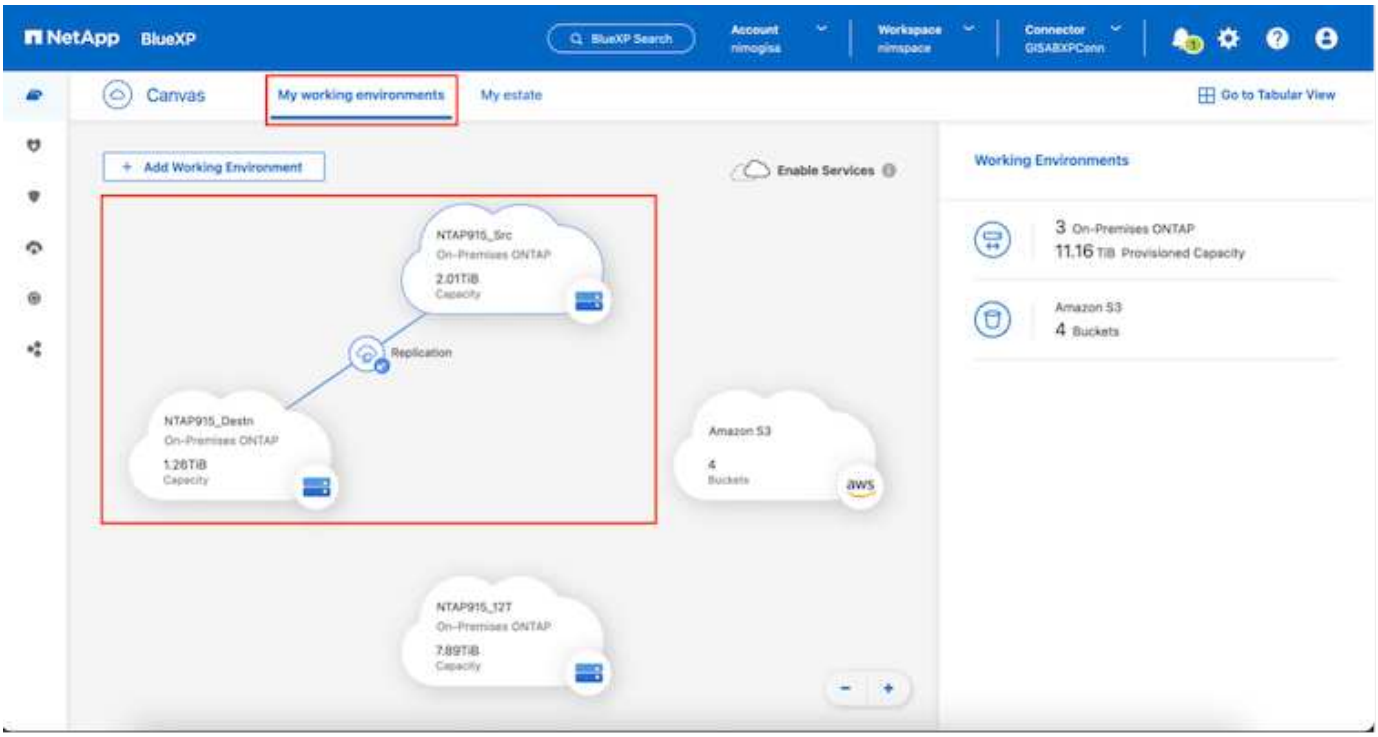


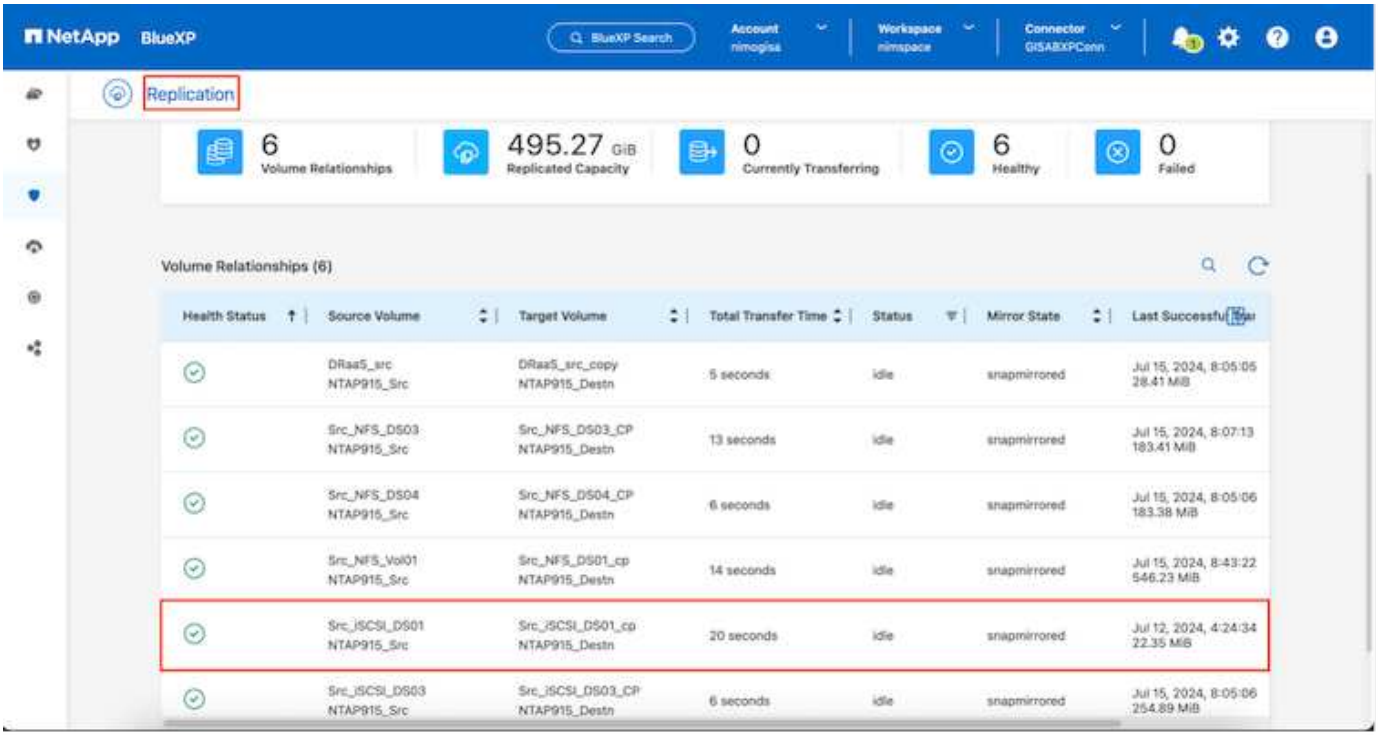
The resource group can also be created while creating a replication plan.

The boot order of the VMs can be defined or modified during the creation of resource groups by using simple drag and drop mechanism.



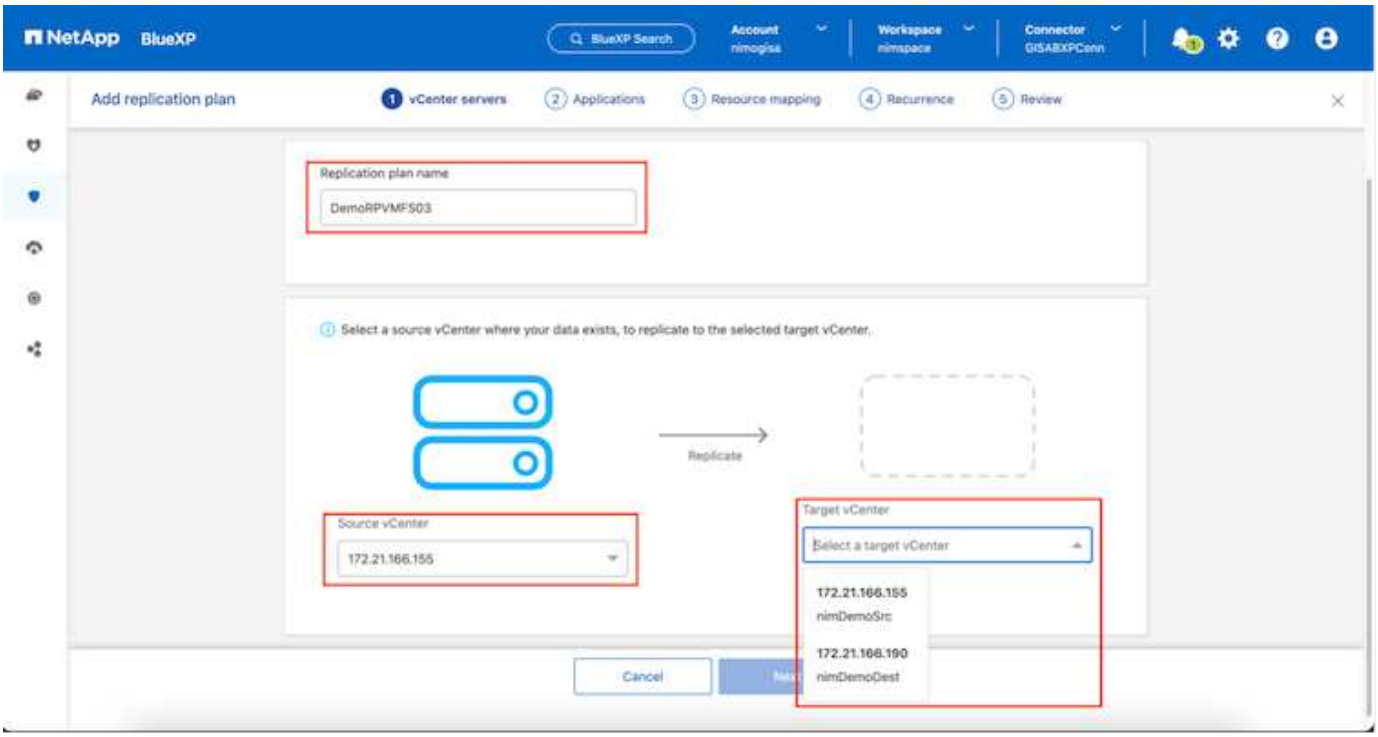
Once the resource groups are created, the next step is to create the execution blueprint or a plan to recover virtual machines and applications in the event of a disaster. As mentioned in the prerequisites, SnapMirror replication can be configured beforehand or DRaaS can configure it using the RPO and retention count specified during creation of the replication plan.



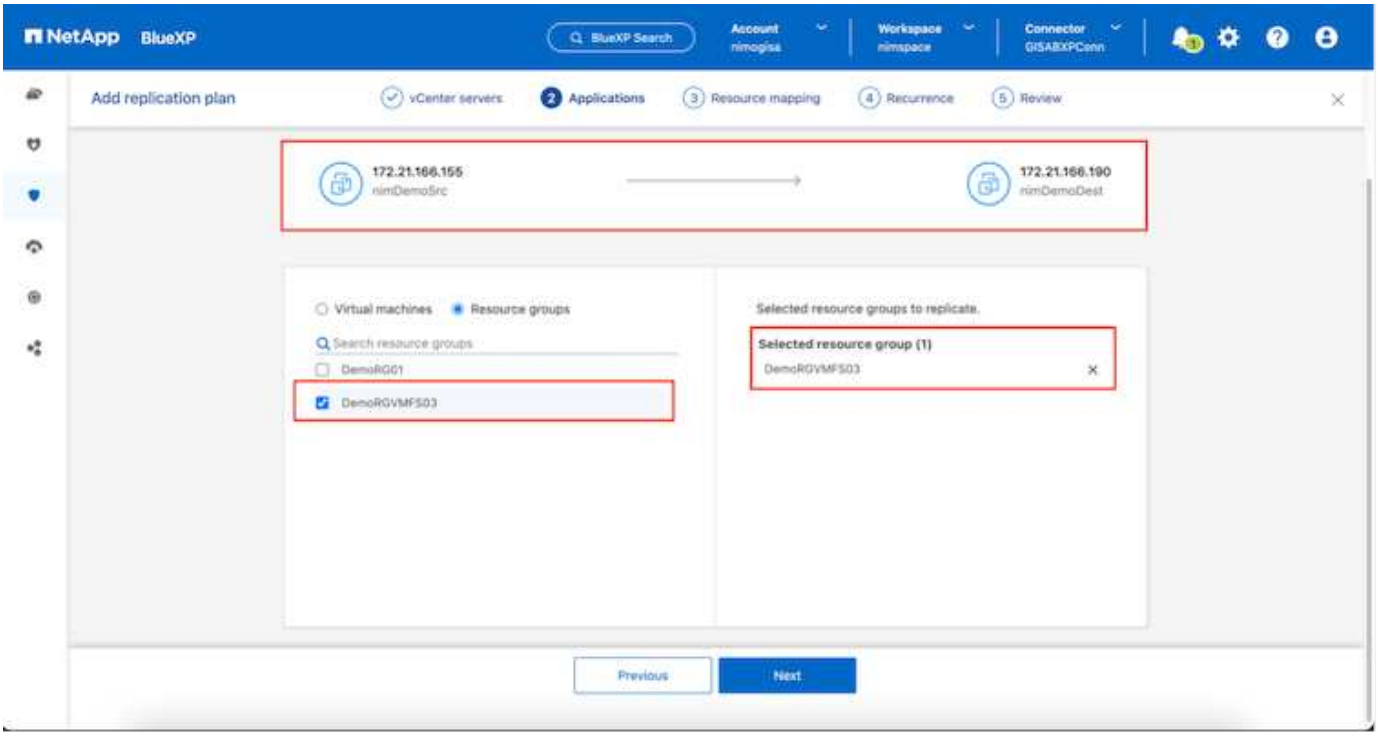


Configure the replication plan by selecting the source and destination vCenter platforms from the drop down and pick the resource groups to be included in the plan, along with the grouping of how applications should be restored and powered on and mapping of clusters and networks. To define the recovery plan, navigate to the **Replication Plan** tab and click **Add Plan**.

First, select the source vCenter and then select the destination vCenter.



The next step is to select existing resource groups. If no resource groups created, then the wizard helps to group the required virtual machines (basically create functional resource groups) based on the recovery objectives. This also helps define the operation sequence of how application virtual machines should be restored.

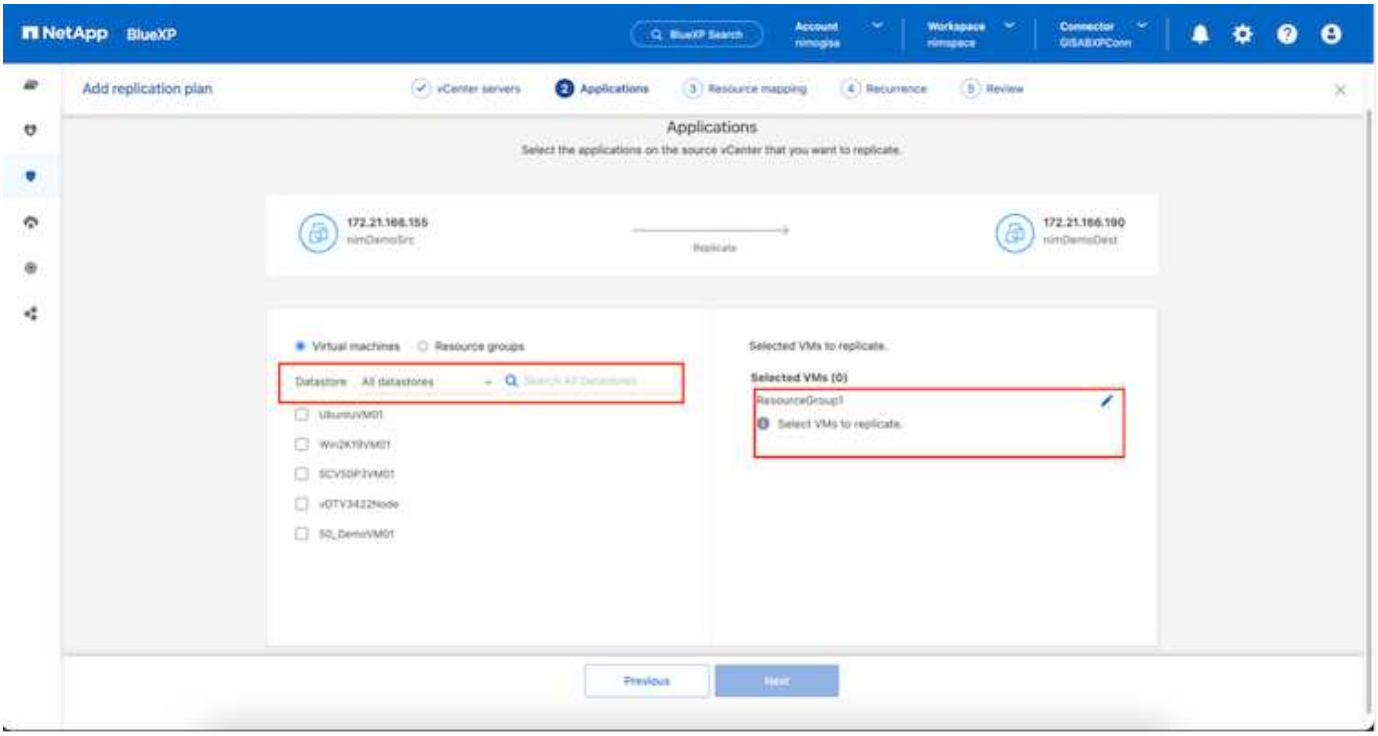


Resource group allows to set boot order using the drag and drop functionality. It can be used to easily modify the order in which the VMs would be powered on during the recovery process.

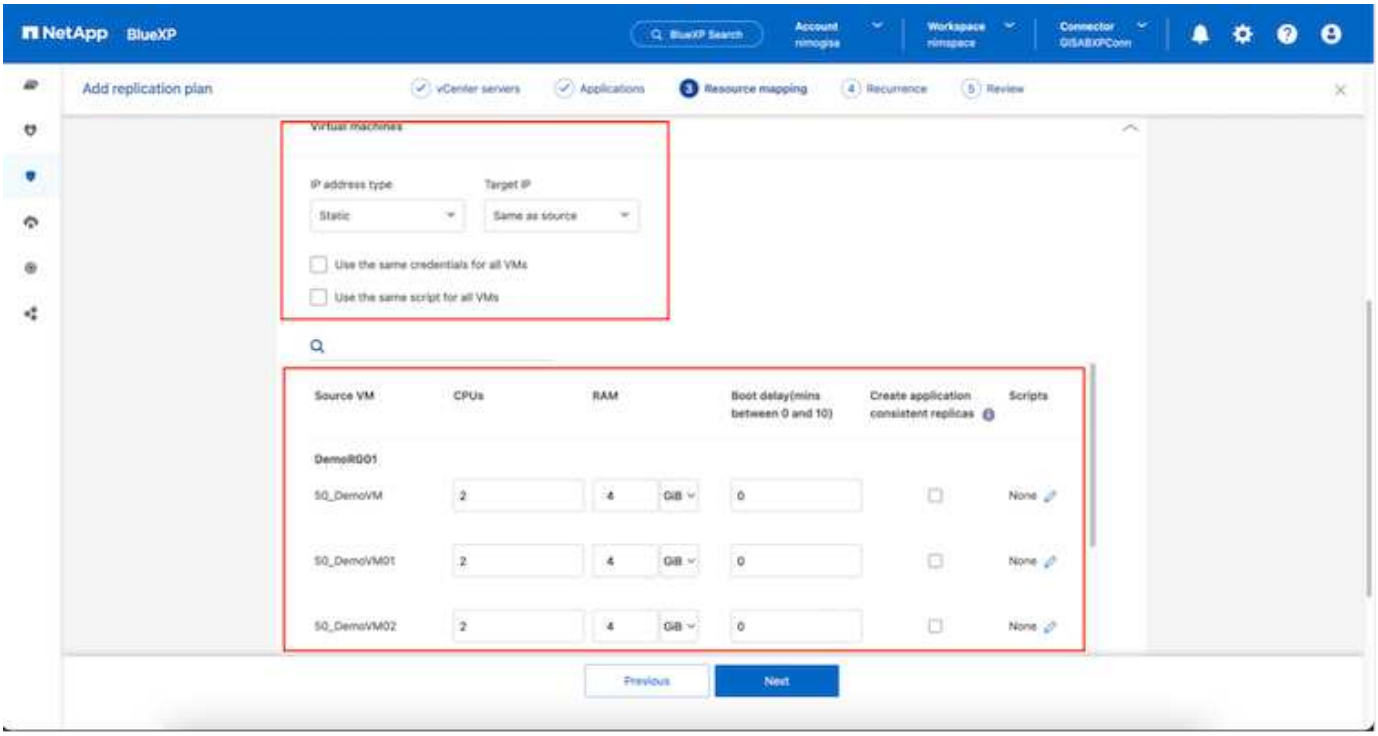
Each virtual machine within a resource group is started in sequence based on the order. Two

resource groups are started in parallel.

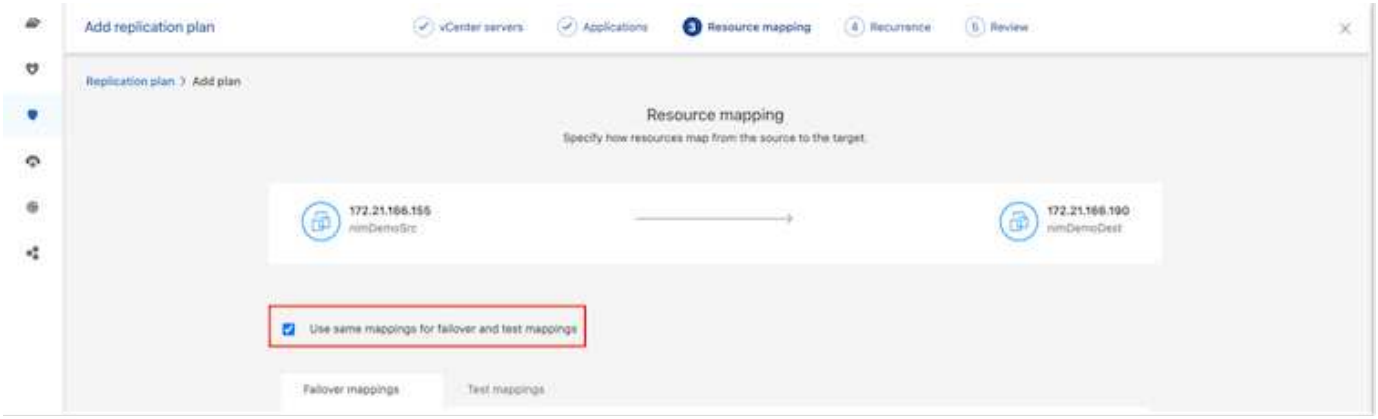
The below screenshot shows the option to filter virtual machines or specific datastores based on organizational requirements if resource groups are not created beforehand.



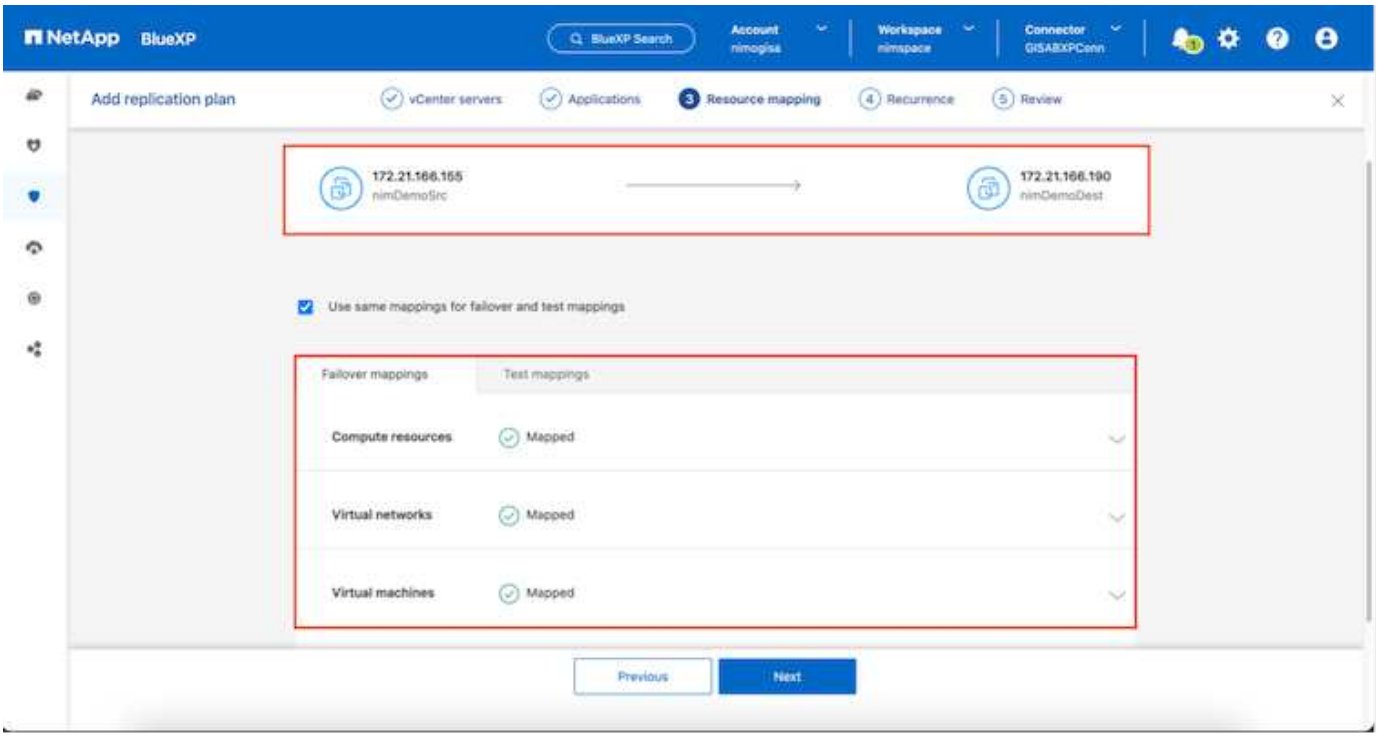
Once the resource groups are selected, create the failover mappings. In this step, specify how the resources from the source environment maps to the destination. This includes compute resources, virtual networks. IP customization, pre- and post-scripts, boot delays, application consistency and so on. For detailed information, refer to Create a replication plan.



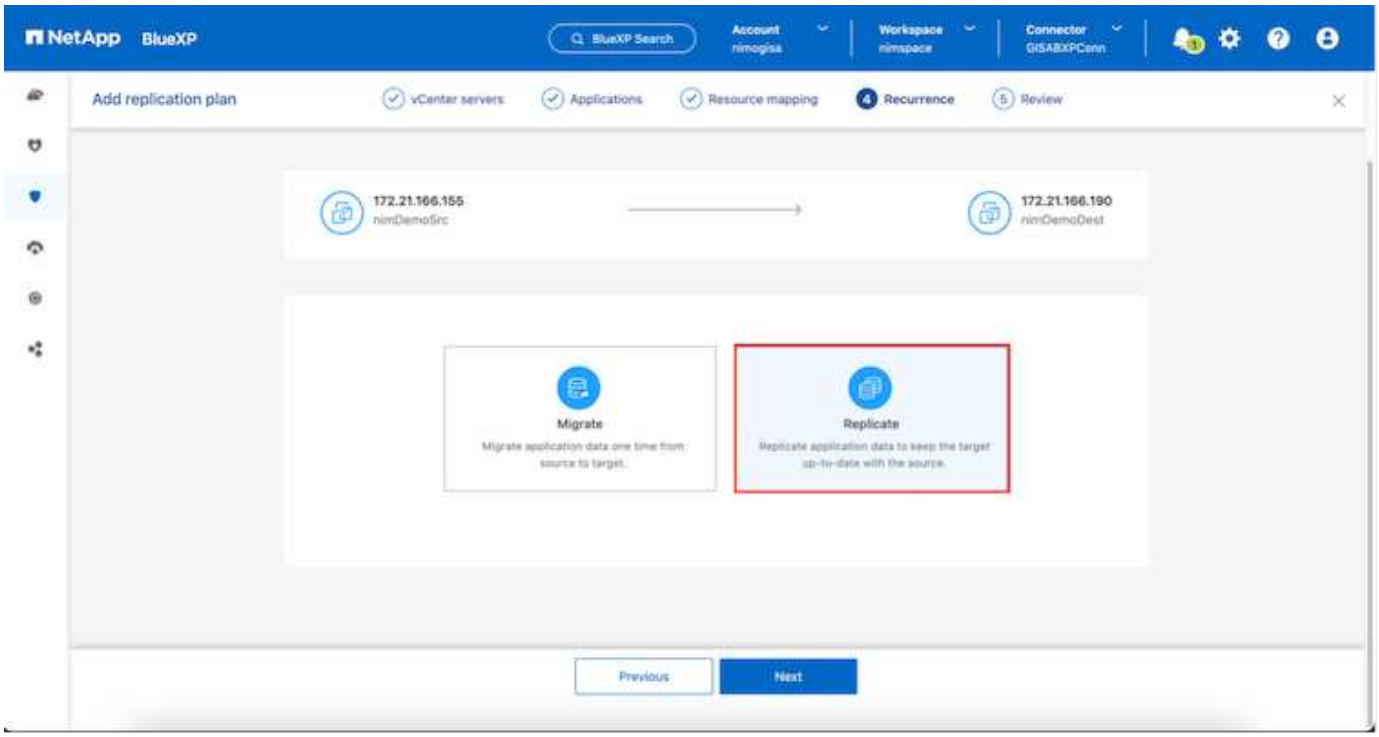
By default, same mapping parameters are used for both test and failover operations. To apply different mappings for test environment, select the Test mapping option after unchecking the checkbox as shown below:



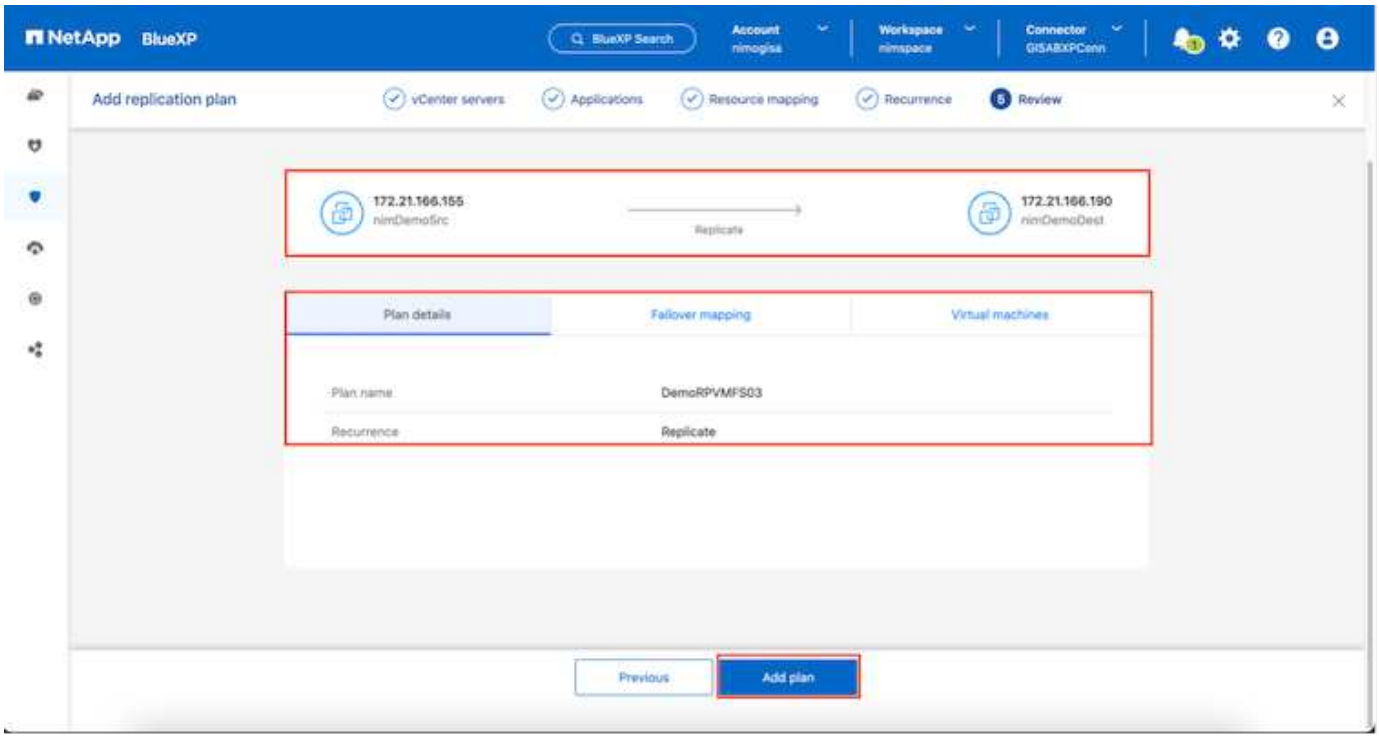
Once the resource mapping is complete, click Next.

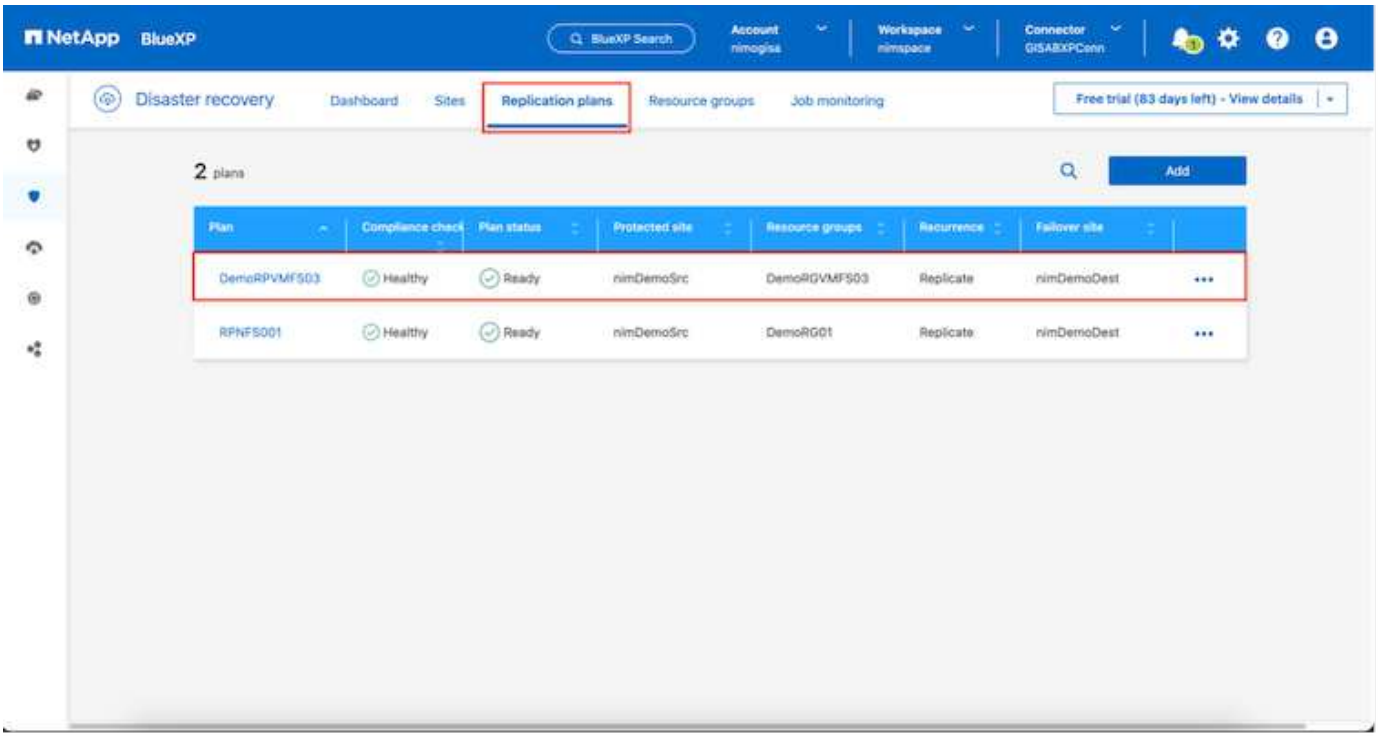


Select the recurrence type. In simple words, select Migrate (one time migration using failover) or recurring continuous replication option. In this walkthrough, Replicate option is selected.

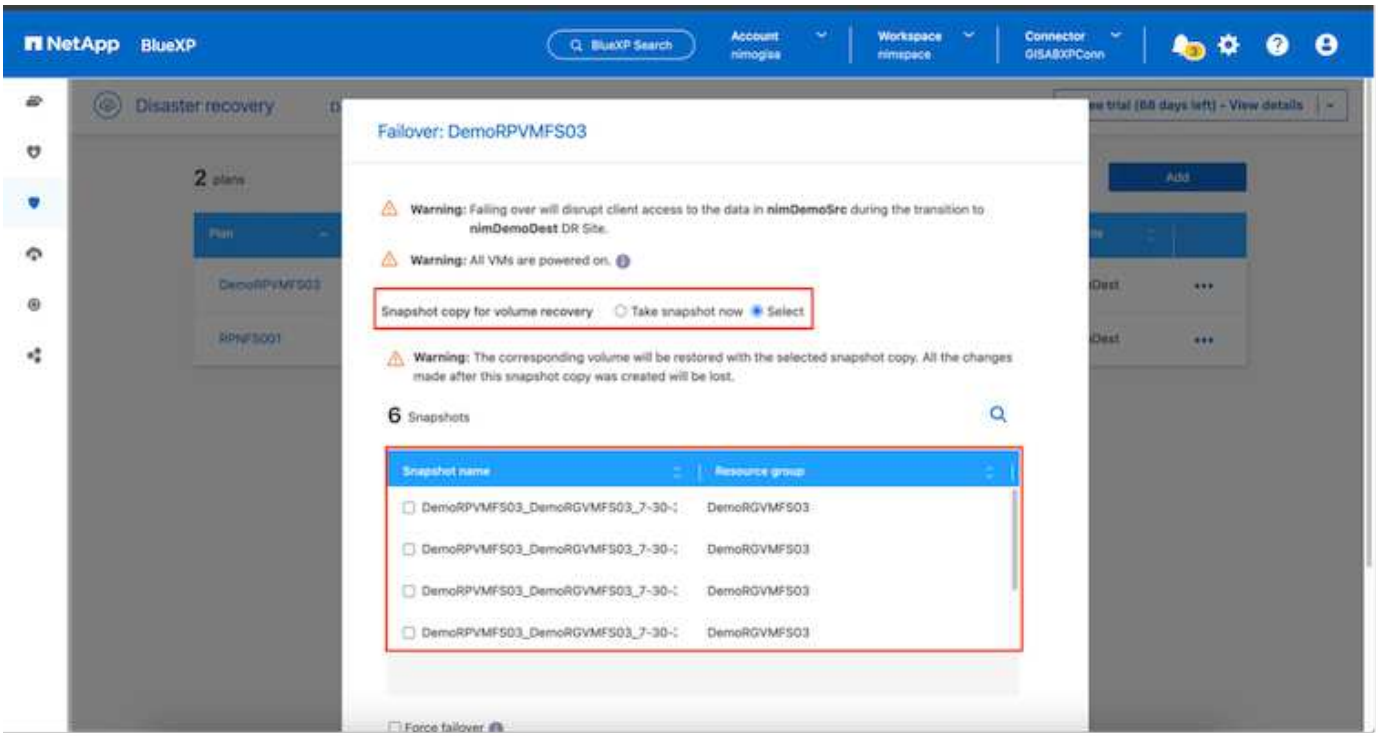


Once done, review the created mappings and then click on Add plan.

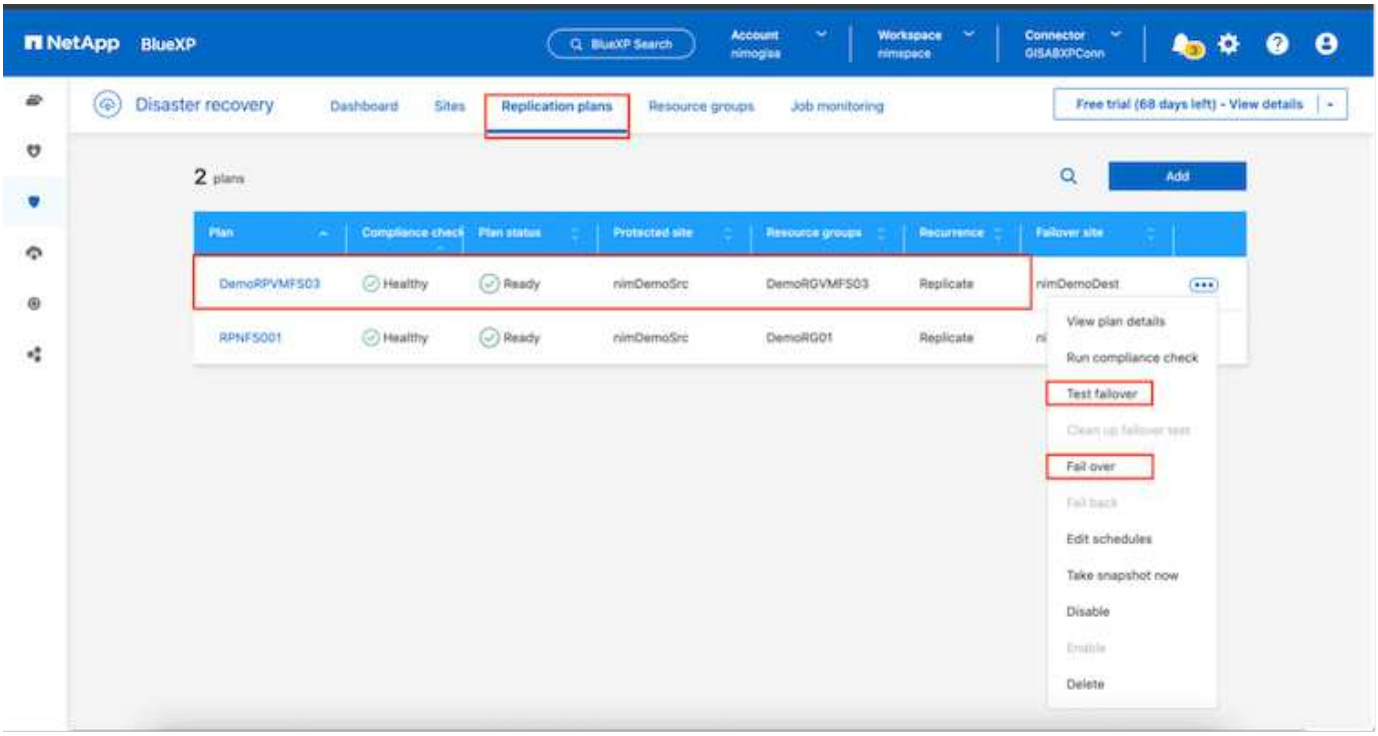




Once the replication plan is created, failover can be performed depending on the requirements by selecting the failover option, test-failover option, or the migrate option. Disaster Recovery ensures that the replication process is being executed according to the plan every 30 minutes. During the failover and test-failover options, you can use the most recent SnapMirror Snapshot copy, or you can select a specific Snapshot copy from a point-in-time Snapshot copy (per the retention policy of SnapMirror). The point-in-time option can be very helpful if there is a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. Disaster Recovery shows all available recovery points.



To trigger failover or test failover with the configuration specified in the replication plan, click on **Failover** or **Test failover**.

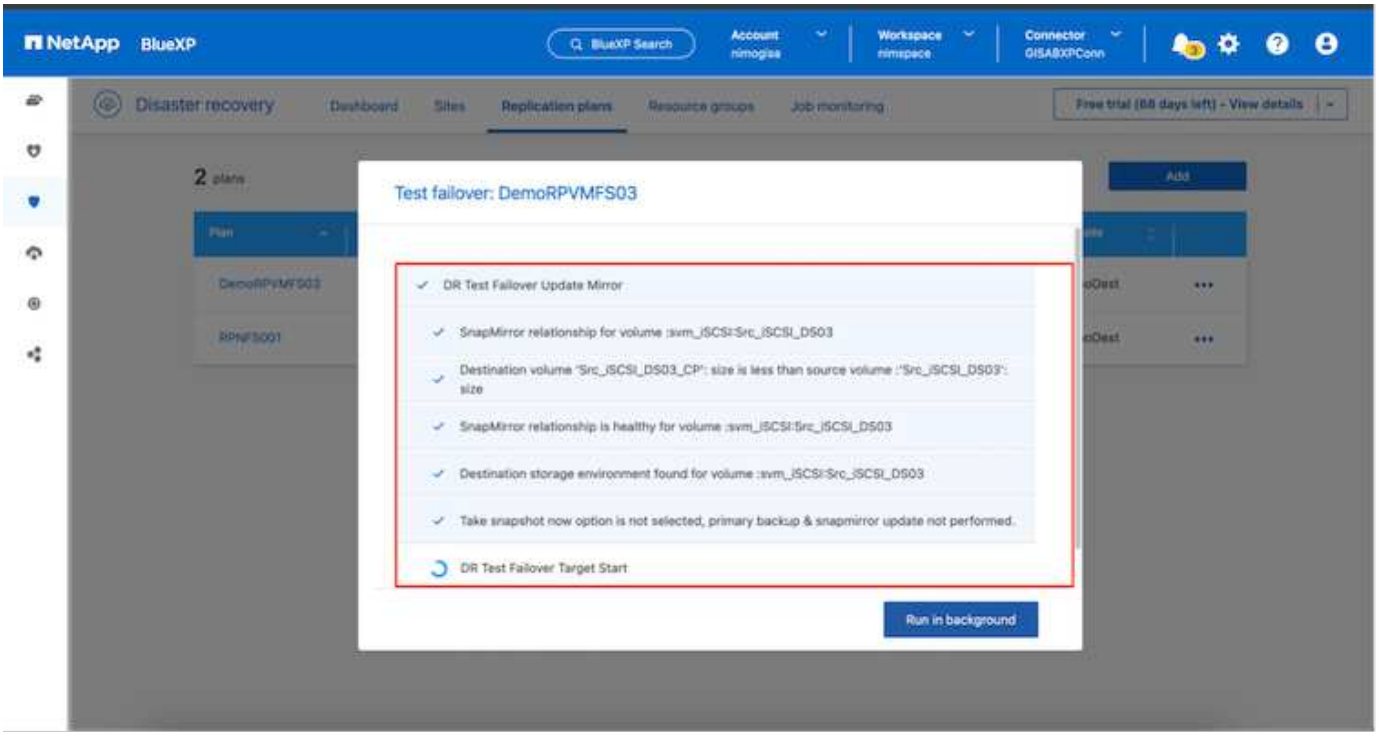


What happens during a failover or test failover operation?

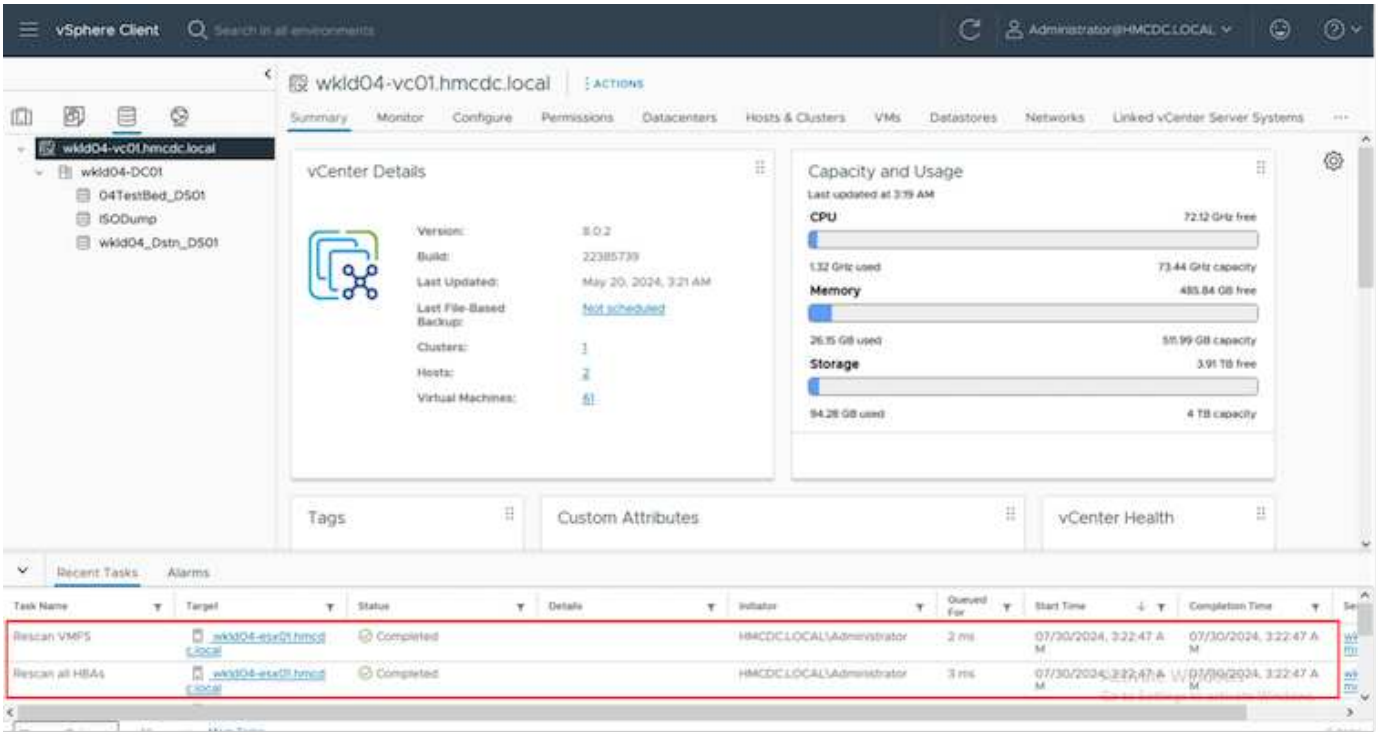
During a test failover operation, Disaster Recovery creates a FlexClone volume on the destination ONTAP storage system using the latest Snapshot copy or a selected snapshot of the destination volume.

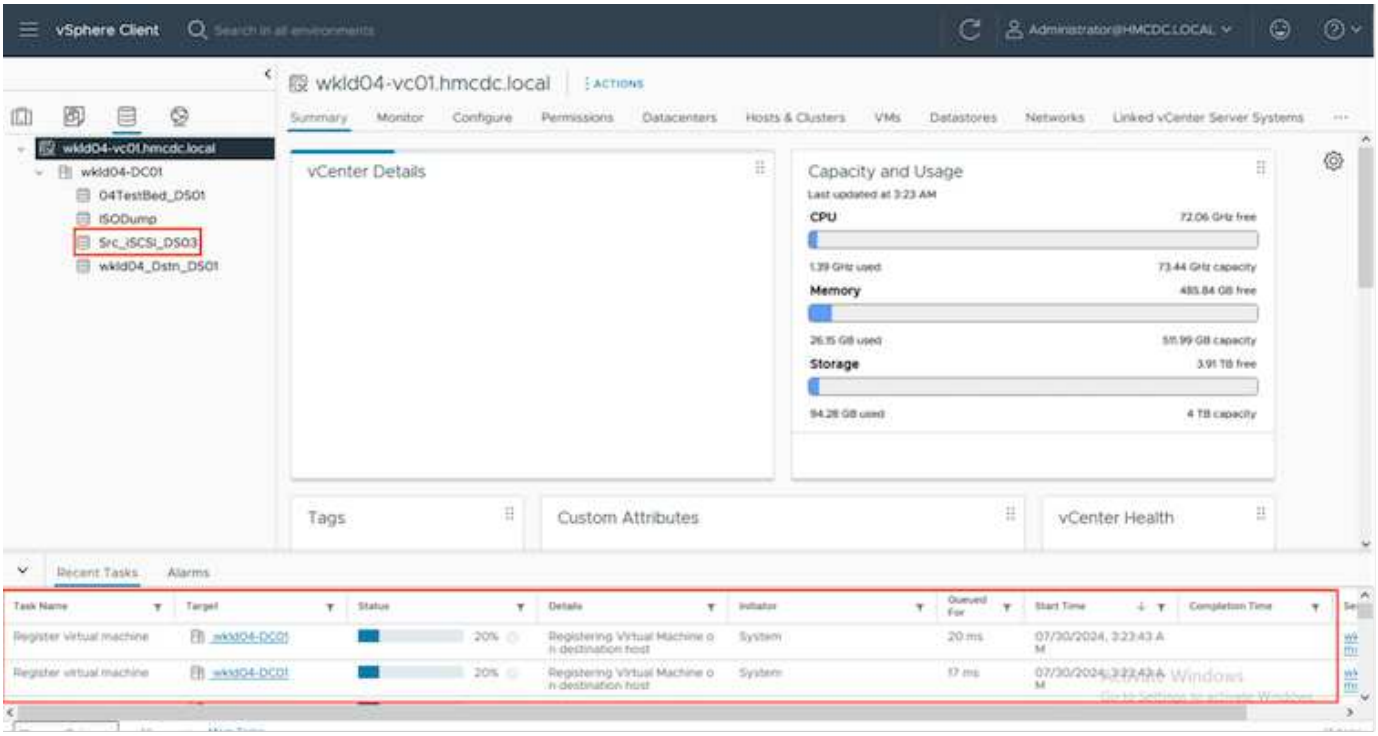
ⓘ A test failover operation creates a cloned volume on the destination ONTAP storage system.

ⓘ Running a test recovery operation does not affect the SnapMirror replication.



During the process, Disaster Recovery does not map the original target volume. Instead, it makes a new FlexClone volume from the selected Snapshot and a temporary datastore backing the FlexClone volume is mapped to the ESXi hosts.



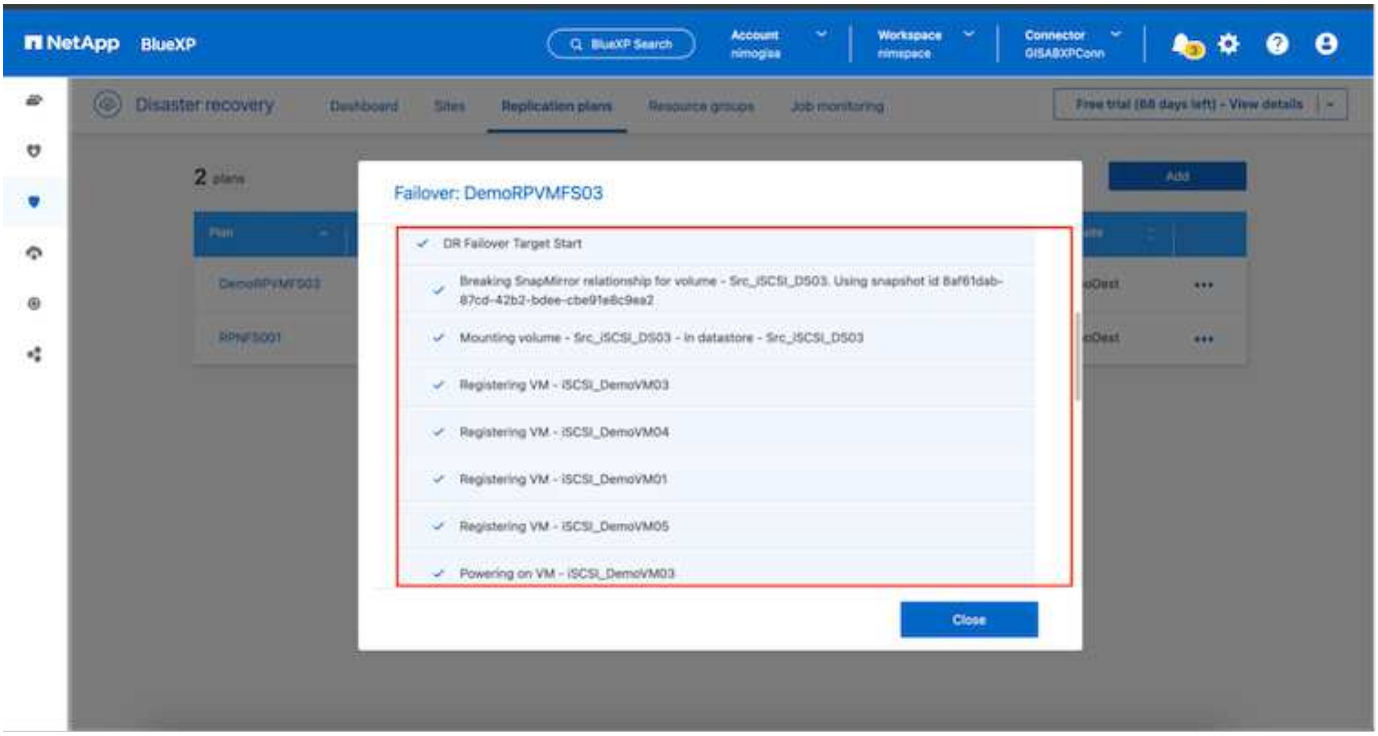


When the test failover operation completes, the cleanup operation can be triggered using **"Clean Up failover test"**. During this operation, Disaster Recovery destroys the FlexClone volume that was used in the operation.

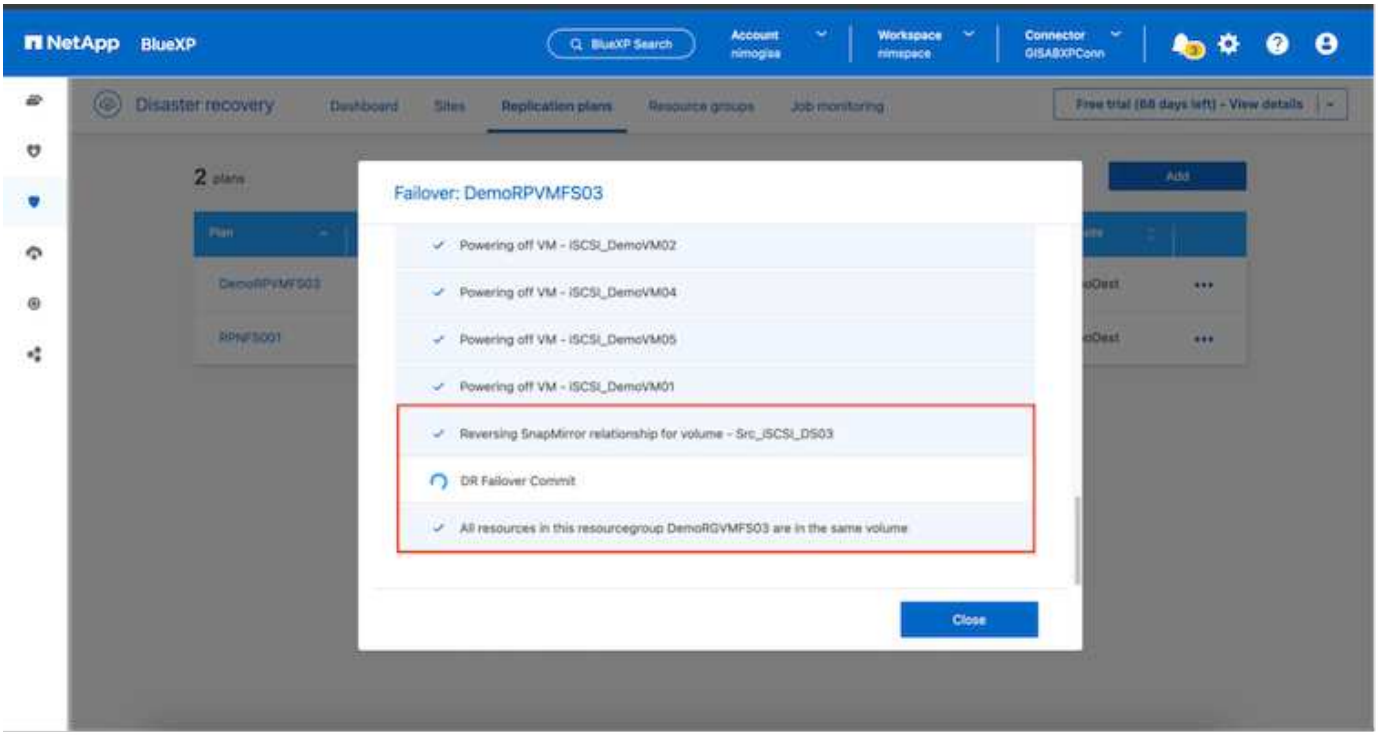
In the event of real disaster event occurs, Disaster Recovery performs the following steps:

1. Breaks the SnapMirror relationship between the sites.

2. Mounts the VMFS datastore volume after resignature for immediate use.
3. Register the VMs
4. Power on VMs



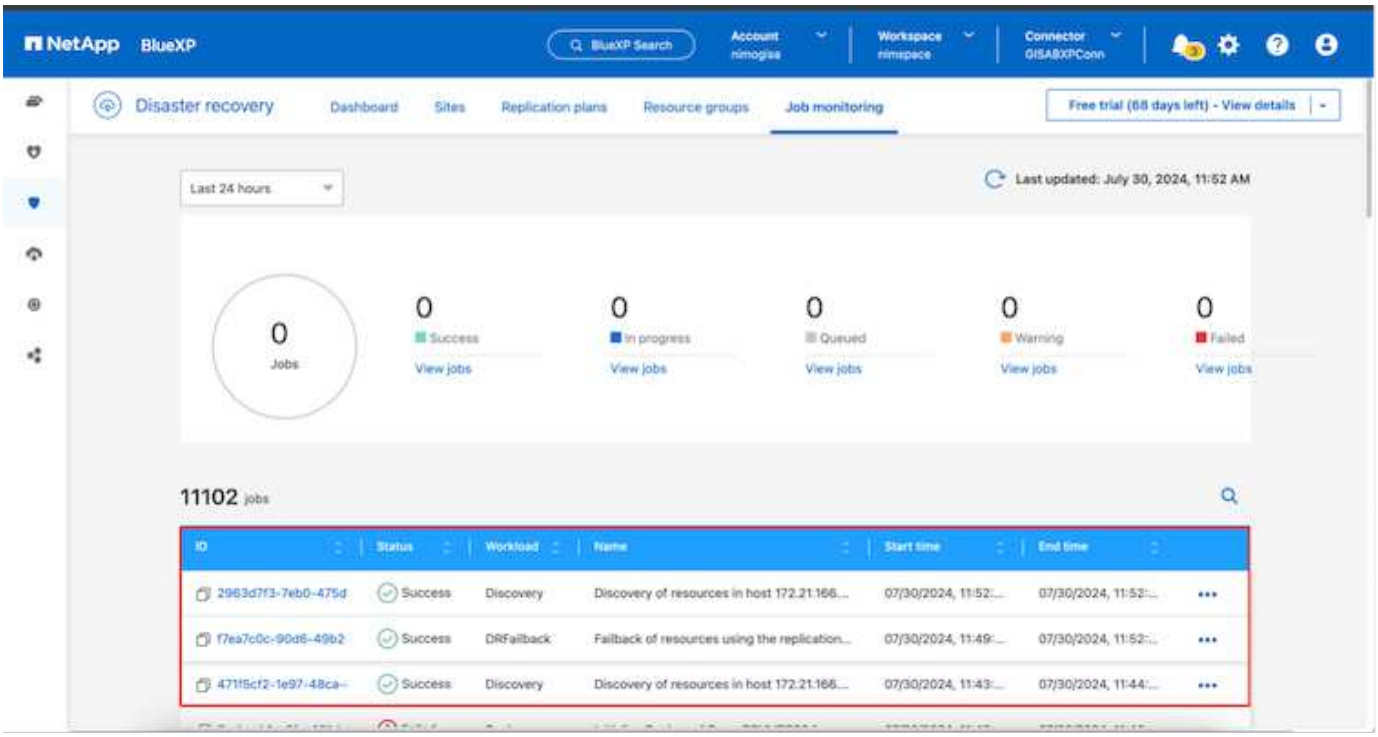
Once the primary site is up and running, Disaster Recovery enables reverse resync for SnapMirror and enables failback, which again can be performed with the click of a button.



And if migrate option is chosen, it is considered as a planned failover event. In this case, an additional step is triggered which is to shut down the virtual machines at the source site. The rest of the steps remains the same

as failover event.

From Console or the ONTAP CLI, you can monitor the replication health status for the appropriate datastore volumes, and the status of a failover or test failover can be tracked via Job Monitoring.



This provides a powerful solution to handle a tailored and customized disaster recovery plan. Failover can be done as planned failover or failover with a click of a button when disaster occurs and decision is made to activate the DR site.

To learn more about this process, feel free to follow the detailed walkthrough video or use the solution simulator.

NetApp hybrid multicloud with Red Hat OpenShift container workloads

Get started

NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

NetApp is seeing a significant increase in customers modernizing their legacy enterprise applications and building new applications using containers and orchestration platforms built around Kubernetes. Red Hat OpenShift Container Platform is one example that we see adopted by many of our customers.

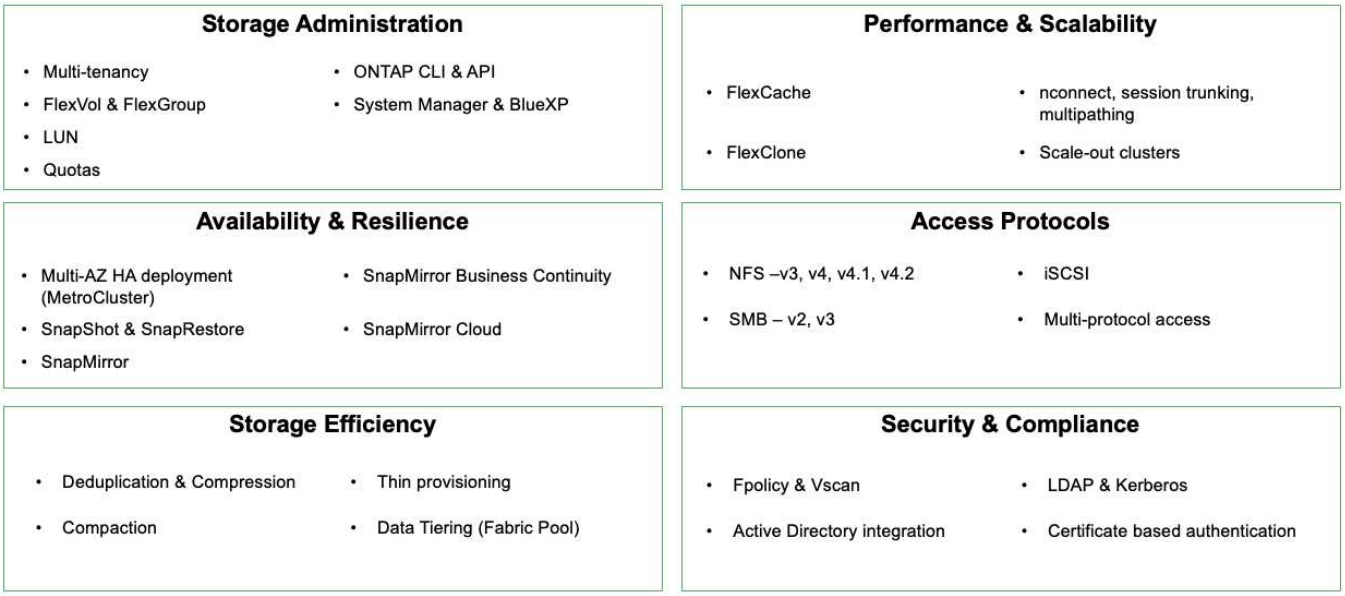
Overview

As more and more customers begin adopting containers within their enterprises, NetApp is perfectly positioned to help serve the persistent storage needs of their stateful applications and classic data management needs such as data protection, data security, and data migration. However, these needs are met using different strategies, tools, and methods.

NetApp ONTAP based storage options listed below, deliver security, data protection, reliability, and flexibility for containers and Kubernetes deployments.

- Self-managed storage in on-premises:
  - NetApp Fabric Attached Storage (FAS), NetApp All Flash FAS Arrays (AFF), NetApp All SAN Array (ASA) and ONTAP Select
- Provider-managed storage in on-premises:
  - NetApp Keystone provides Storage as a Service (STaaS)
- Self-managed storage in the cloud:
  - NetApp Cloud Volumes ONTAP(CVO) provide self managed storage in the hyperscalers
- Provider-managed storage in the cloud:
  - Amazon FSx for NetApp ONTAP, Azure NetApp Files, and Google Cloud NetApp Volumes provide file-based storage in the cloud.

ONTAP feature highlights

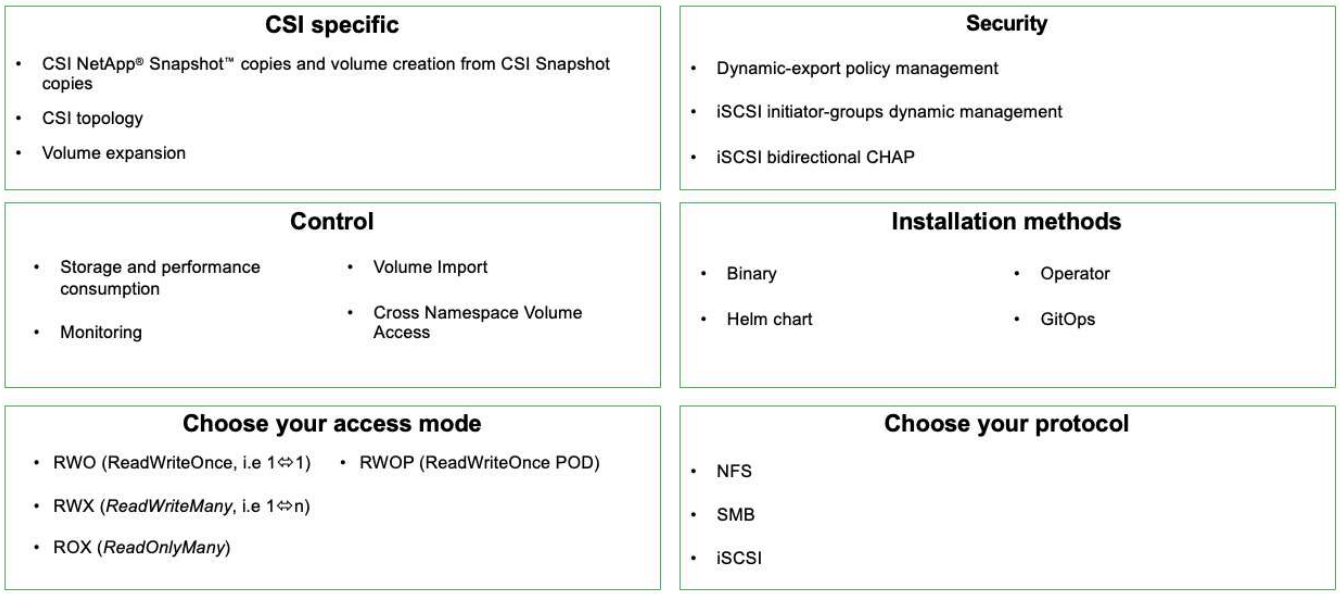


NetApp Console enables you to manage all of your storage and data assets from a single control plane/interface.

You can use NetApp Console to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files), to move, protect, and analyze data, and to control many on-prem and edge storage devices.

NetApp Trident is a CSI Compliant Storage Orchestrator that enable quick and easy consumption of persistent storage backed by a variety of the above-mentioned NetApp storage options. It is an open-source software maintained and supported by NetApp.

Astra Trident CSI feature highlights



Business critical container workloads need more than just persistent volumes. Their data management requirements require protection and migration of the application kubernetes objects as well.



Application data includes kubernetes objects in addition to the user data: Some examples are as follows:
- kubernetes objects such as pods specs, PVCs, deployments, services
- custom config objects such as config maps and secrets
- persistent data such as Snapshot copies, backups, clones
- custom resources such as CRs and CRDs

NetApp Trident Protect, available as a free add-on to NetApp Trident, delivers advanced application data management capabilities that enhance the functionality and availability of stateful Kubernetes applications backed by NetApp ONTAP storage systems. Trident Protect simplifies the management, protection, and movement of containerized workloads across both public clouds and on-premises environments. It also offers robust automation capabilities through its API and CLI.

This reference documentation provides validation of data protection of container-based applications, deployed on RedHat OpenShift platform using Trident Protect. In addition, the solution provides high-level details for the deployment and the use of Red Hat Advanced Cluster Management (ACM) for managing the container platforms. The document also highlights the details for the integration of NetApp storage with Red Hat OpenShift container platforms using Trident CSI provisioner.

Value propositions of NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

Most customers do not just start out building Kubernetes based environments without any existing infrastructure. Perhaps they are a traditional IT shop running most of their enterprise applications on virtual machines (in large VMware environments for example). Then they start building small container-based environments to satisfy the needs of their modern application development teams. These initiatives usually start small and begin to become more pervasive as the teams learn these new technologies and skills, and begin to recognize the many benefits of adopting them.
The good news for customers is that NetApp can serve the needs of both environments. This set of solutions for hybrid multicloud with Red Hat OpenShift will empower NetApp customers to adopt modern cloud technologies and services without having to overhaul their entire infrastructure and organization. Whether customer applications and data are hosted on-premises, in cloud, run on virtual machines, or on containers, NetApp can provide consistent data management, protection, security, and portability. With these new solutions, the same value NetApp has delivered in on-premises data center environments for decades will be available across the enterprise entire data horizon, without requiring significant investment to retool, acquire new skills, or build new teams. NetApp is positioned well to help customers solve these business challenges regardless of what phase of their cloud journey they are in.

NetApp Hybrid Multi-Cloud with Red Hat Openshift:

- Gives customers validated designs and practices which demonstrate the best ways for customers to manage, protect, secure, and migrate their data and applications when using Red Hat OpenShift with NetApp based storage solutions.

- Present best practices for customers running Red Hat OpenShift with NetApp storage in VMware environments, bare metal infrastructure, or a combination of both.
- Demonstrate strategies and options for both on-prem and cloud environments, as well as hybrid environments where both are used.

Supported Solutions of NetApp Hybrid Multicloud for Red Hat OpenShift Container workloads

The solution tests and validates Migration & Centralized Data Protection with OpenShift Container Platform (OCP), OpenShift Advanced Cluster Manager (ACM), NetApp ONTAP, NetApp Console and NetApp Trident Protect.

For this solution, the following scenarios are tested and validated by NetApp. The solution is separated into multiple scenarios based on the following characteristics:

- on-premises
- cloud
  - self-managed OpenShift clusters and self-managed NetApp storage
  - provider-managed OpenShift clusters and provider-managed NetApp storage

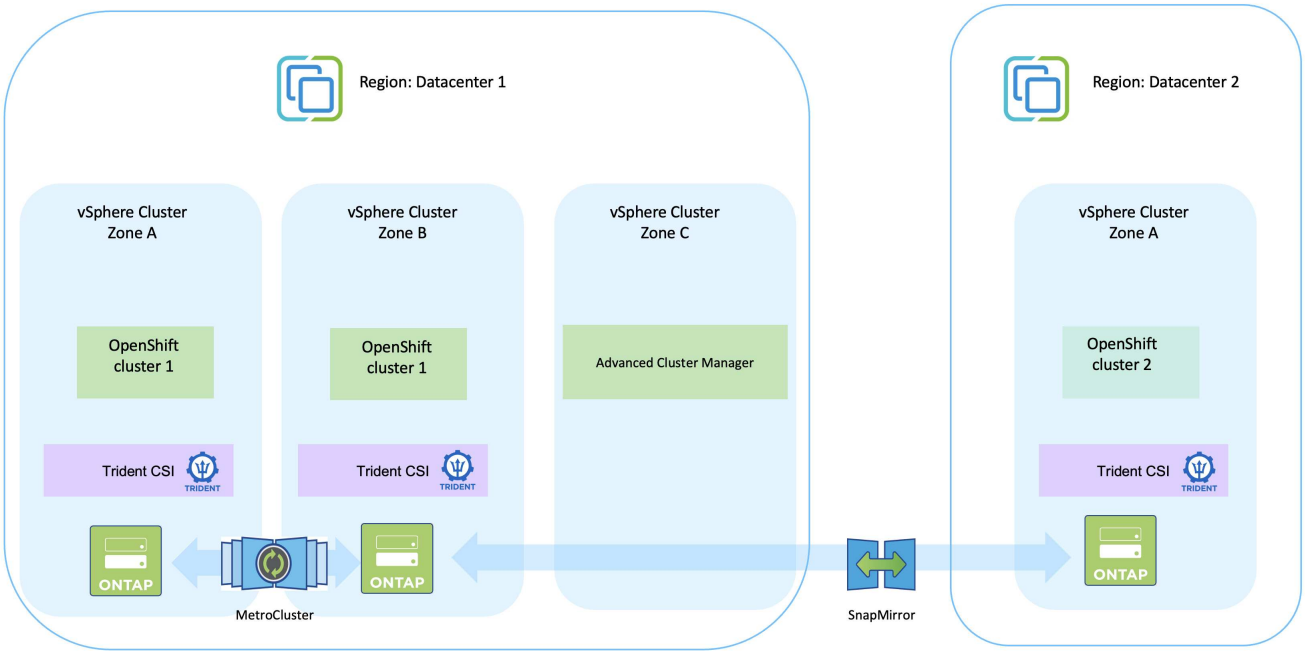
We will be building out additional solutions and use cases in the future.

Scenario 1: Data protection and migration within the on-premises environment using Trident Protect

On-premises: self-managed OpenShift clusters and self-managed NetApp storage

- Use Trident Protect to protect applications and support migration between the on-premises clusters.

Scenario 1

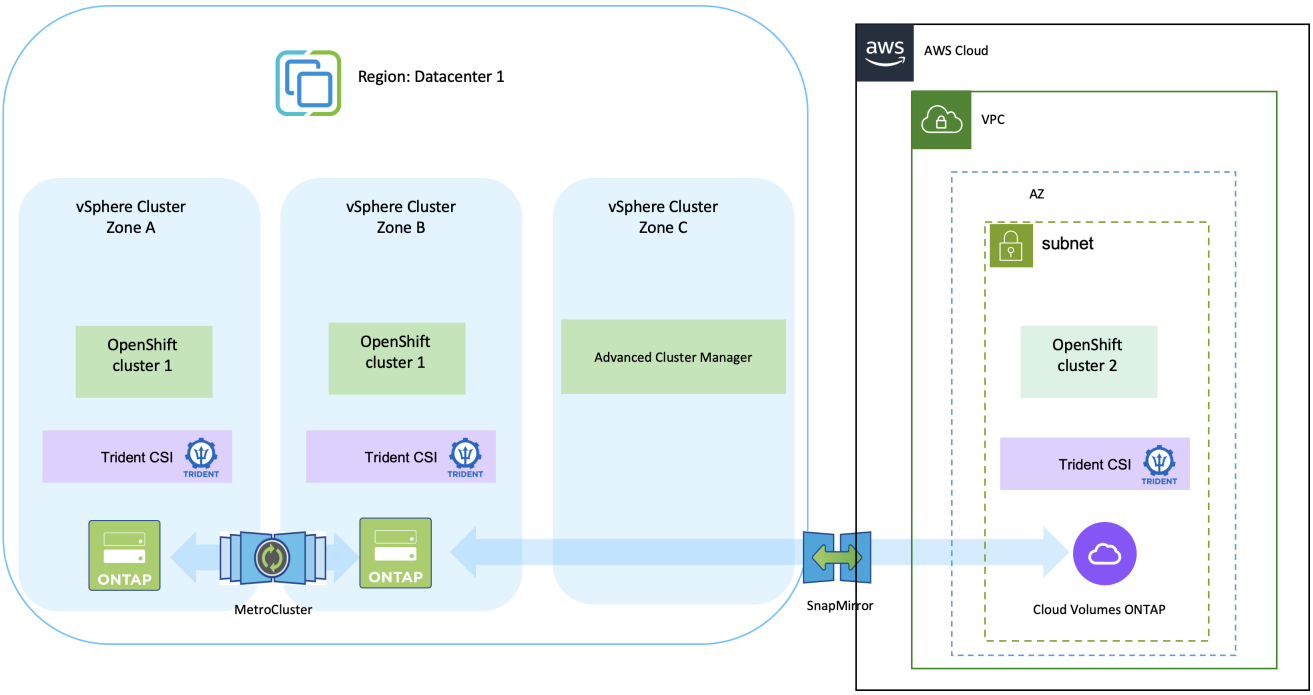


Scenario 2: Data protection and migration from the on-premises environment to AWS environment using Trident Protect

On-premises: Self-managed OpenShift cluster and self-managed storage
AWS Cloud: Self-managed OpenShift cluster and self-managed storage

- Use Trident Protect to protect applications and support migration between the on-premises clusters and the self managed OpenShift clusters in AWS.

Scenario 2

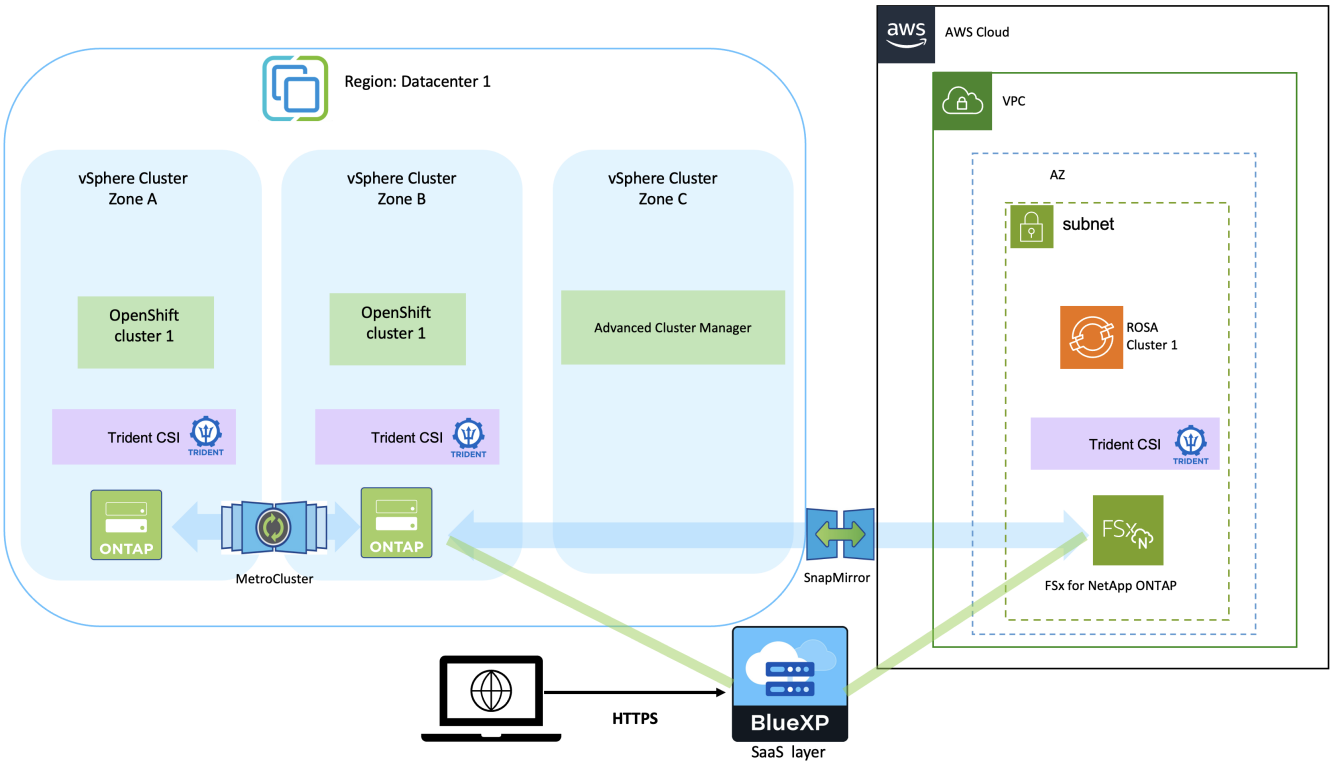


Scenario 3: Data protection and migration from the on-premises environment to AWS environment

On-premises: Self-managed OpenShift cluster and self-managed storage
AWS Cloud: Provider-managed OpenShift cluster (ROSA) and provider-managed storage (FSx ONTAP)

- Using NetApp Console, perform replication of persistent volumes (FSx ONTAP).
- Using OpenShift GitOps, recreate application metadata.

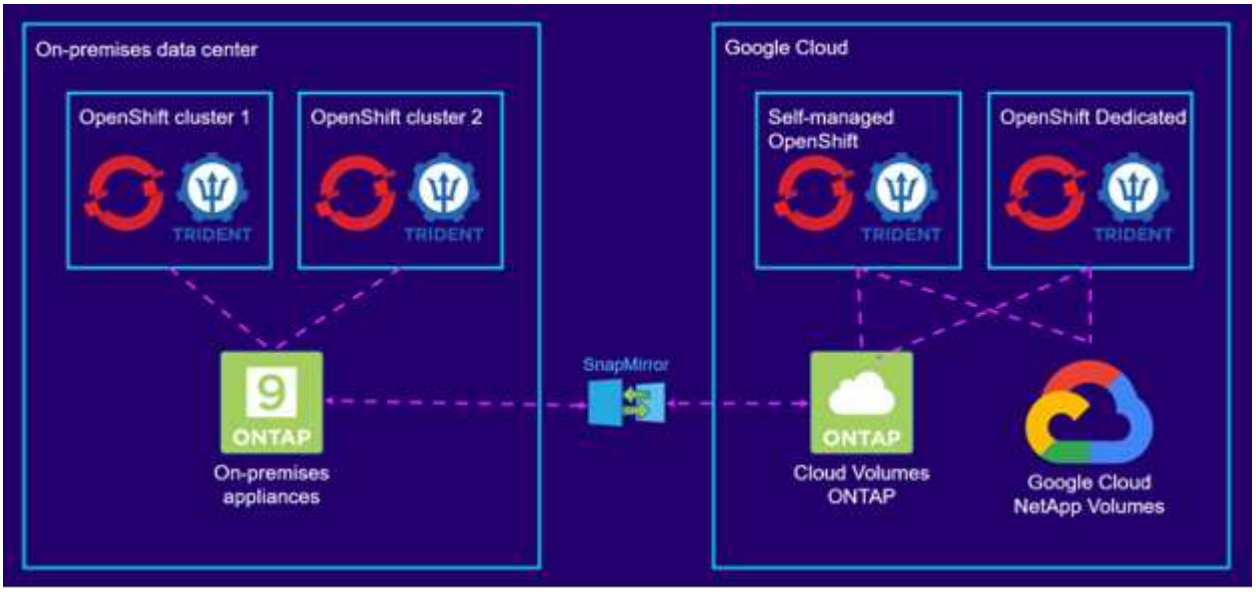
Scenario 3



Scenario 4: Data protection and migration from the on-premises environment to GCP environment using Trident Protect

On-premises: Self-managed OpenShift cluster and self-managed storage
Google Cloud: Self-managed OpenShift cluster and self-managed storage

- Use Trident Protect to protect applications and support migration between the on-premises and Google Cloud OpenShift clusters.

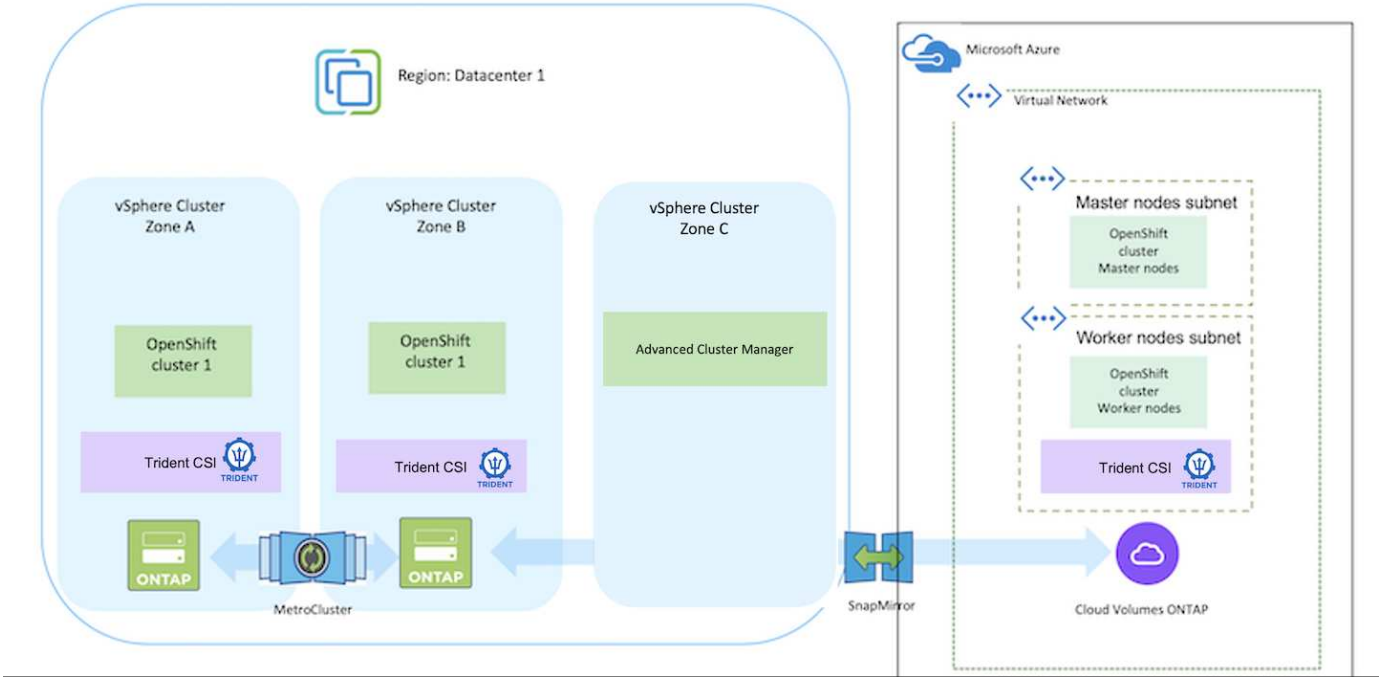


For considerations when using ONTAP in a MetroCluster configuration, refer here.

Scenario 5: Data protection and migration from the on-premises environment to Azure environment using Trident Protect

On-premises: Self-managed OpenShift cluster and self-managed storage
Azure Cloud: Self-managed OpenShift cluster and self-managed storage

- Use Trident Protect to protect applications and support migration between the on-premises and OpenShift clusters in Azure.



For considerations when using ONTAP in a MetroCluster configuration, refer here.

Versions of various components used in the solution validation

The solution tests and validates Migration & Centralized Data Protection with OpenShift container platform, OpenShift Advanced Cluster Manager, NetApp ONTAP, and NetApp Trident Protect.

Scenarios 1, 2 and 3 of the solution were validated using the versions as shown in the table below:

Component	Version
VMware	vSphere Client version 8.0.0.10200 VMware ESXi, 8.0.0, 20842819
Hub Cluster	OpenShift 4.11.34
Source and Destination Clusters	OpenShift 4.12.9 on-premises and in AWS
NetApp Trident	Trident Server and Client 23.04.0
NetApp Trident Protect	Trident Protect 22.11.0-82
NetApp ONTAP	ONTAP 9.12.1

AWS FSx ONTAP	Single AZ

Scenario 4 of the solution was validated using the versions as shown in the table below:

Component	Version
VMware	vSphere Client version 8.0.2.00000 VMware ESXi, 8.0.2, 22380479
Hub Cluster	OpenShift 4.13.13
Source and Destination Clusters	OpenShift 4.13.12 on-premises and in Google Cloud
NetApp Trident	Trident Server and Client 23.07.0
NetApp ONTAP	ONTAP 9.12.1
Cloud Volumes ONTAP	Single AZ, Single node,9.14.0

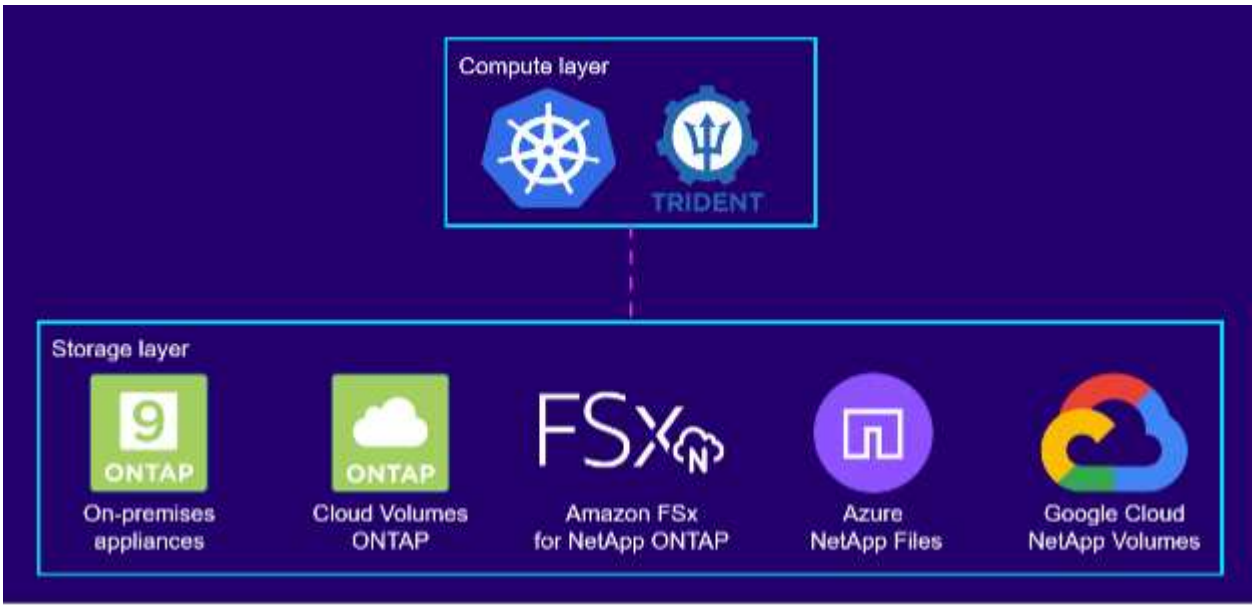
Scenario 5 of the solution was validated using the versions as shown in the table below:

Component	Version
VMware	vSphere Client version 8.0.2.00000 VMware ESXi, 8.0.2, 22380479
Source and Destination Clusters	OpenShift 4.13.25 on-premises and in Azure
NetApp Trident	Trident Server and Client 23.10.0
NetApp Trident Protect	Trident Protect 23.10
NetApp ONTAP	ONTAP 9.12.1
Cloud Volumes ONTAP	Single AZ, Single node,9.14.0

Supported NetApp Storage integrations with Red Hat OpenShift Containers

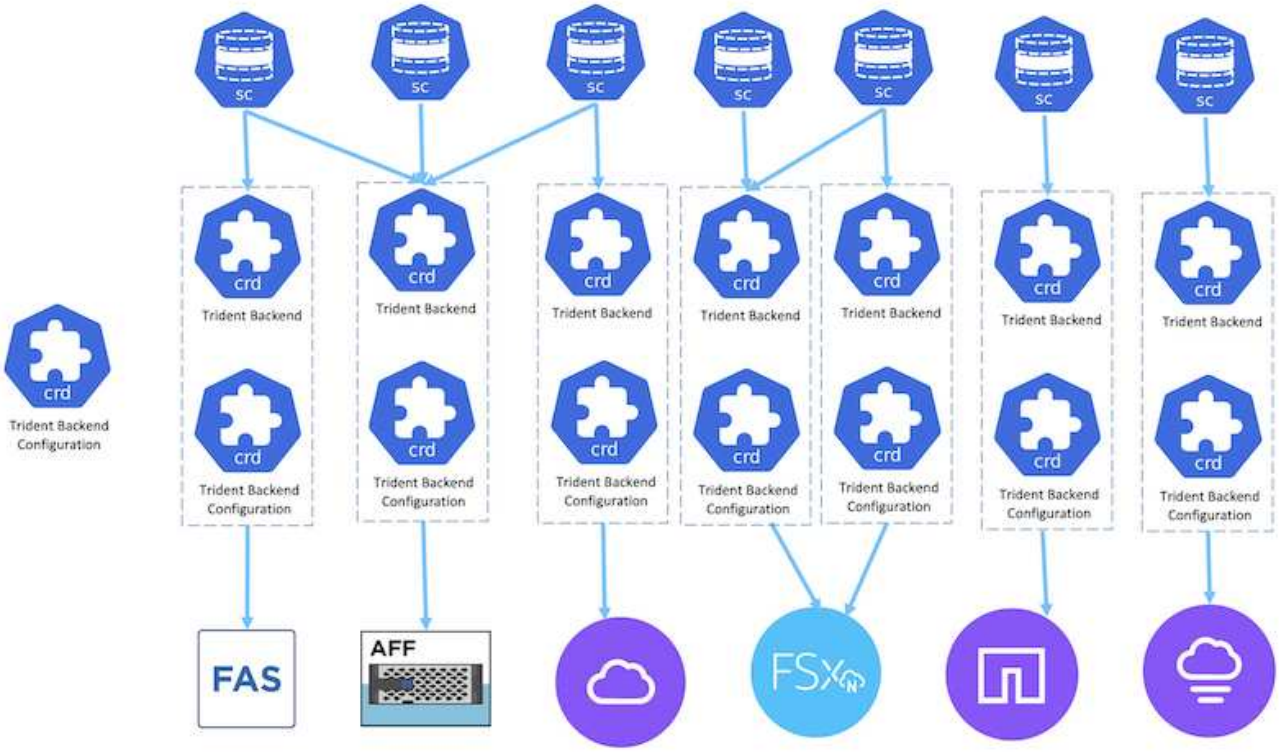
Whether the Red Hat OpenShift containers are running on VMware or in the hyperscalers, NetApp Trident can be used as the CSI provisioner for the various types of backend NetApp storage that it supports.

The following diagram depicts the various backend NetApp storage that can be integrated with OpenShift clusters using NetApp Trident.



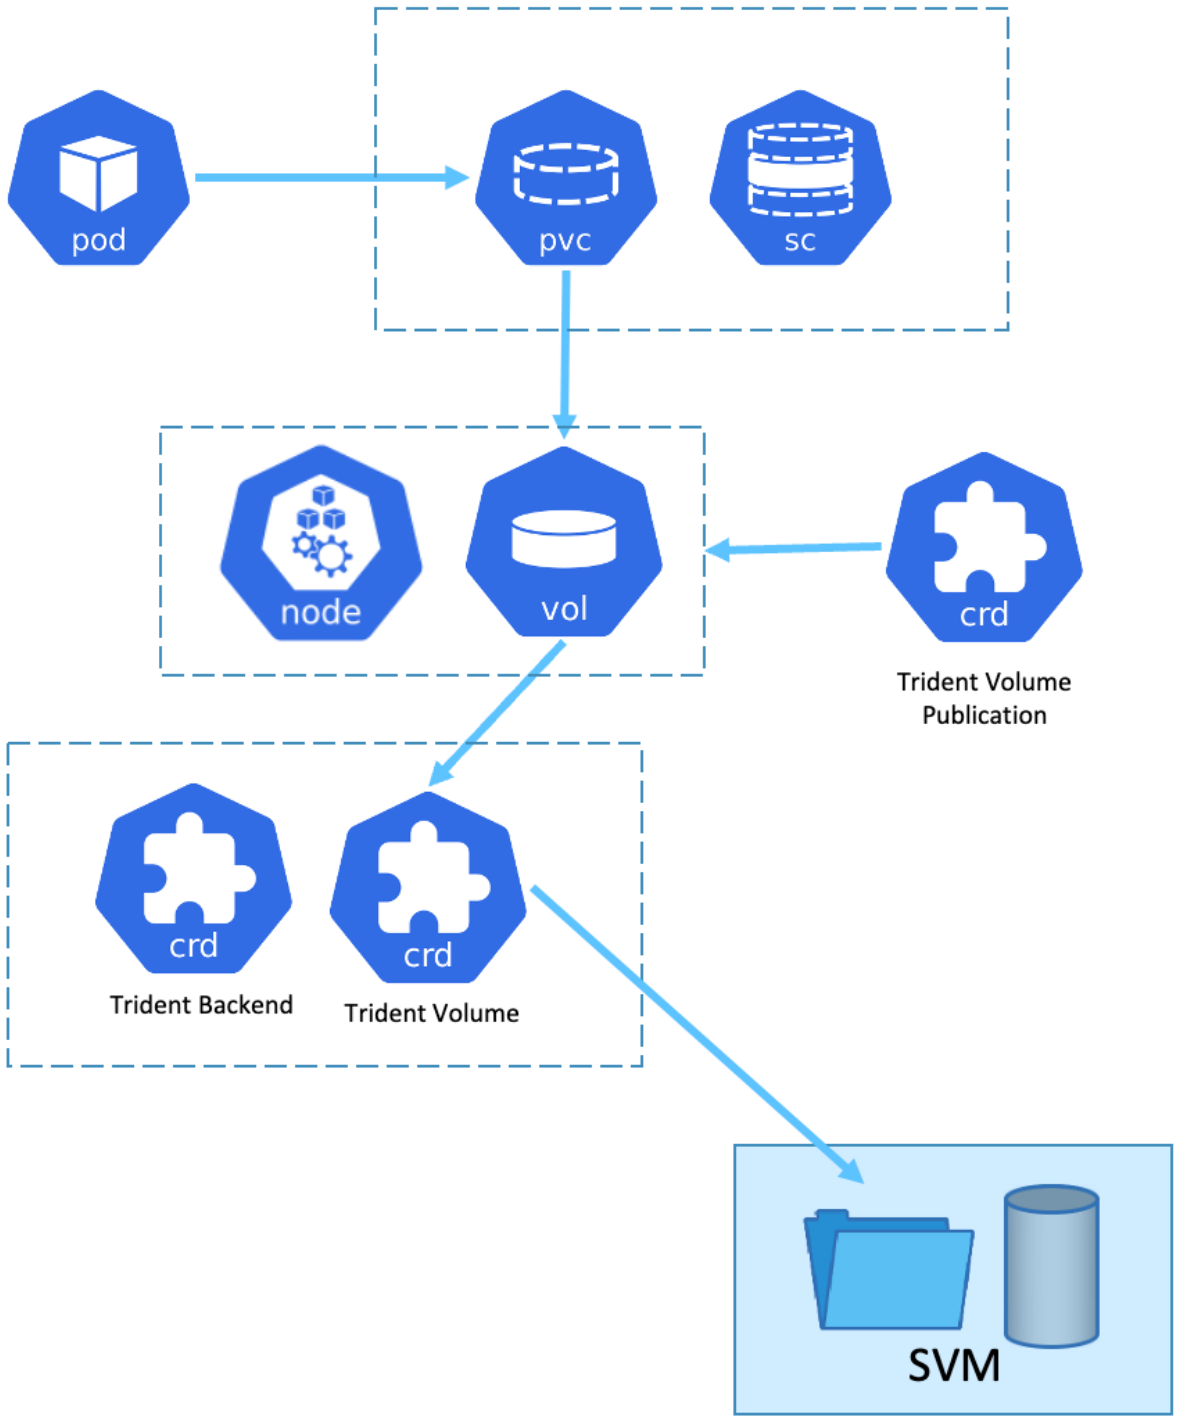
ONTAP Storage Virtual Machine (SVM) provides secure multi-tenancy. A Single OpenShift cluster can connect to single SVM or multiple SVMs or even to multiple ONTAP clusters. Storage class filters the backend storage based on parameters or by labels. Storage administrators define the parameters to connect to storage system using trident backend configuration. On successful connection establishment, it creates the trident backend and populates the information which the storage class can filter.

The relationship between the storageclass and backend is shown below.



Application owner requests persistent volume using storage class. The storage class filters the backend storage.

The relationship between the pod and backend storage is shown below.



Container Storage Interface (CSI) Options

On vSphere environments, customers can pick VMware CSI driver and/or Trident CSI to integrate with ONTAP. With VMware CSI, the persistent volumes are consumed as local SCSI disks, whereas with Trident, it is consumed with network.

As VMware CSI does not support RWX access modes with ONTAP, applications need to use Trident CSI if RWX mode is required. With FC based deployments, VMware CSI is preferred and SnapMirror Business Continuity (SMBC) provides zone level high availability.

VMware CSI supports

- Core Block based datastores (FC, FCoE, iSCSI, NVMeoF)
- Core File based datastores (NFS v3, v4)
- vVol datastores (block and file)

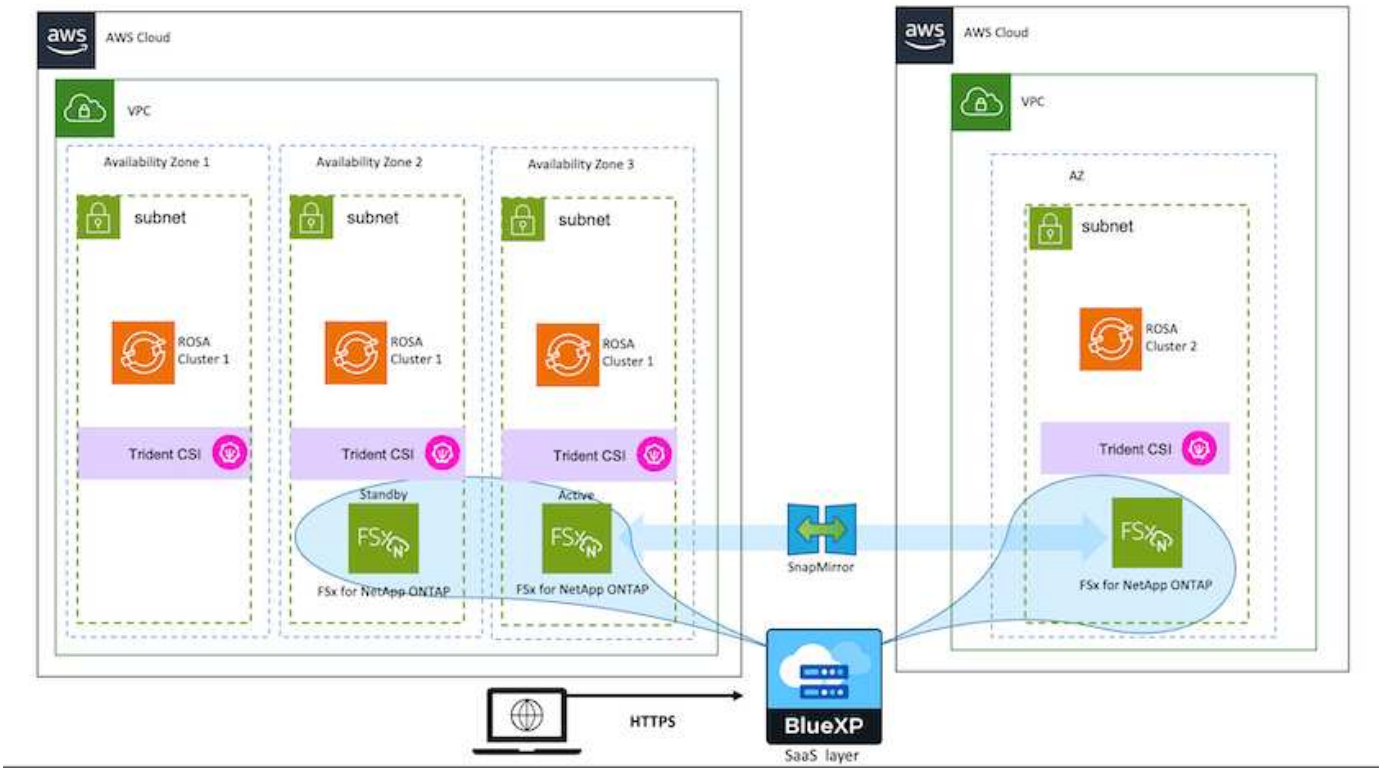
Trident has following drivers to support ONTAP

- ontap-san (dedicated volume)
- ontap-san-economy (shared volume)
- ontap-nas (dedicated volume)
- ontap-nas-economy (shared volume)
- ontap-nas-flexgroup (dedicated large scale volume)

For both VMware CSI and Trident CSI, ONTAP supports nconnect, session trunking, kerberos, etc. for NFS and multipathing, chap authentication, etc. for block protocols.

In AWS, Amazon FSx for NetApp ONTAP (FSx ONTAP) can be deployed in single Availability Zone (AZ) or in Multi AZ. For production workloads that requires high availability, multi-AZ provides zonal level fault tolerance and has better NVMe read cache compared to single AZ. For more info, check AWS performance guidelines.

To save cost on disaster recovery site, single AZ FSx ONTAP can be utilized.



For number of SVMs that are supported by FSx ONTAP, refer managing FSx ONTAP storage virtual machine

Additional Documentation

Here are some additional documentation for reference:

1. VMware vSphere
2. OpenShift Container platform product documentation
3. Installing OpenShift container platform clusters
4. Advanced cluster management product documentation
5. Creating a cluster using ACM
6. Deploying Red Hat Quay on OpenShift
7. Trident
8. ONTAP 9
9. NetApp Verda - Verda (Open source) has a collection of reference execution hooks for popular cloud native data-intensive applications
10. NetApp data services
11. Red Hat OpenShift Service on AWS classic architecture
12. Amazon FSx ONTAP
13. Installing OpenShift Clusters on GCP
14. Installing OpenShift Clusters on Azure

Openshift for on-premises

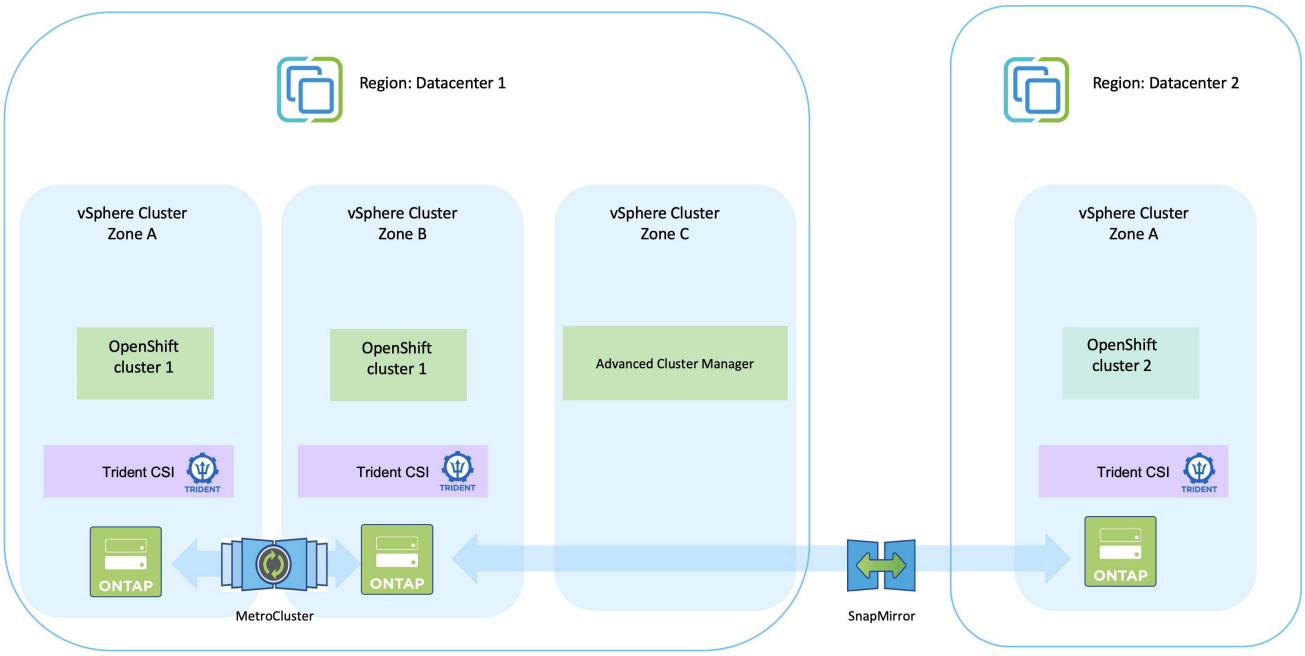
NetApp Solution with Red Hat OpenShift Container platform workloads on VMware

If customers have a need to run their modern containerized applications on infrastructure in their private data centers, they can do so. They should plan and deploy the Red Hat OpenShift container platform (OCP) for a successful production-ready environment for deploying their container workloads. Their OCP clusters can be deployed on VMware or bare metal.

NetApp ONTAP storage delivers data protection, reliability, and flexibility for container deployments. Trident serves as the dynamic storage provisioner to consume persistent ONTAP storage for customers' stateful applications. NetApp Trident Protect can be used for the many data management requirements of stateful applications such as data protection, migration, and business continuity.

With VMware vSphere, NetApp ONTAP tools provides a vCenter Plugin which can be utilized to provision datastores. Apply tags and use it with OpenShift for storing the node configuration and data. NVMe based storage provides lower latency and high performance.

Data protection and migration solution for OpenShift Container workloads using Trident Protect

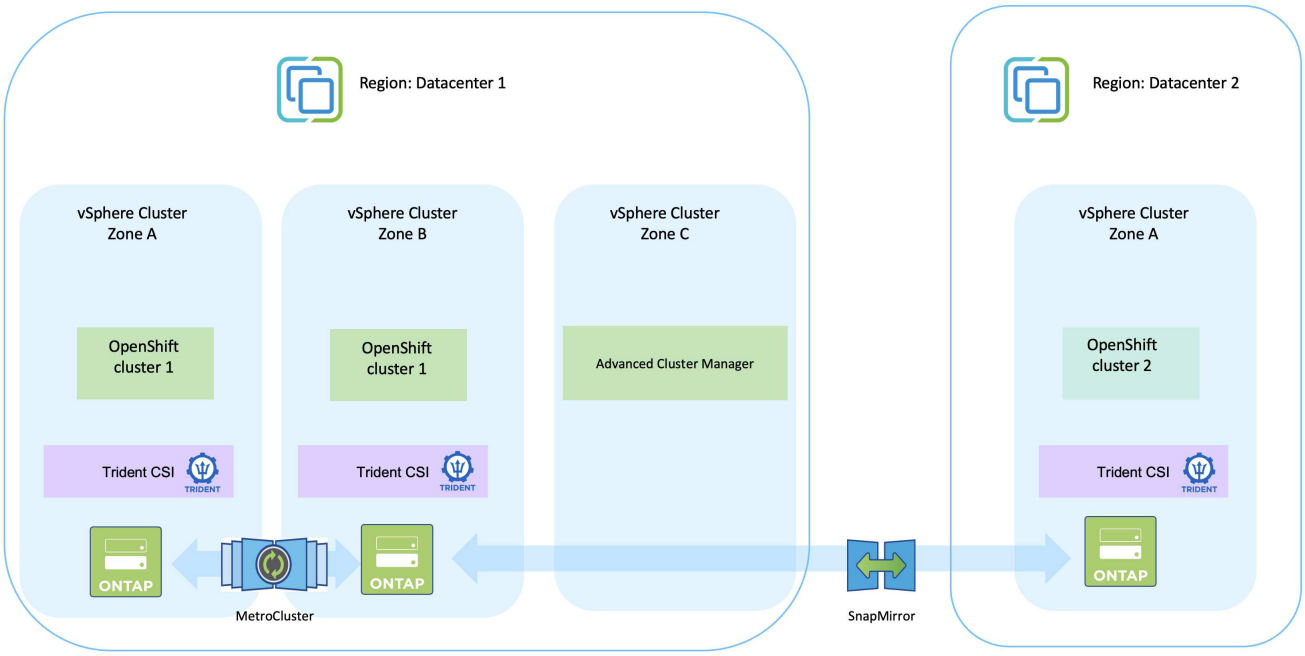


Deploy and configure the Red Hat OpenShift Container platform on VMware

This section describes a high-level workflow of how to set up and manage OpenShift clusters and manage stateful applications on them. It shows the use of NetApp ONTAP storage arrays with the help of Trident to provide persistent volumes.

> There are several ways of deploying Red Hat OpenShift Container platform clusters. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

Here is a diagram that depicts the clusters deployed on VMware in a data center.



The setup process can be broken down into the following steps:

▼ **Deploy and configure a CentOS VM**

- It is deployed in the VMware vSphere environment.
- This VM is used for deploying some components such as NetApp Trident and NetApp Trident Protect for the solution.
- A root user is configured on this VM during installation.

▼ **Deploy and configure an OpenShift Container Platform cluster on VMware vSphere (Hub Cluster)**

Refer to the instructions for the Assisted deployment method to deploy an OCP cluster.

Remember the following:
- Create ssh public and private key to provide to the installer. These keys will be used to login to the master and worker nodes if needed.
- Download the installer program from the assisted installer. This program is used to boot the VMs that you create in the VMware vSphere environment for the master and worker nodes.
- VMs should have the minimum CPU, memory, and hard disk requirement. (Refer to the vm create commands on this page for the master and the worker nodes which provide this information)
- The diskUUID should be enabled on all VMs.
- Create a minimum of 3 nodes for master and 3 nodes for worker.
- Once they are discovered by the installer, turn on the VMware vSphere integration toggle button.

▼ **Install Advanced Cluster Management on the Hub cluster**

This is installed using the Advanced Cluster Management Operator on the Hub Cluster.
Refer to the instructions here.

▼ **Install two additional OCP clusters (Source and Destination)**

- The additional clusters can be deployed using the ACM on the Hub Cluster.
- Refer to the instructions here.

▼ **Configure NetApp ONTAP storage**

- Install an ONTAP cluster with connectivity to the OCP VMs in VMWare environment.
- Create an SVM.
- Configure NAS data lif to access the storage in SVM.

▼ **Install NetApp Trident on the OCP clusters**

- Install NetApp Trident on all three clusters: Hub, source, and destination clusters
- Refer to the instructions here.
- Create a storage backend for ontap-nas .
- Create a storage class for ontap-nas.
- Refer to instructions here.

▼ **Deploy an Application on Source Cluster**

Use OpenShift GitOps to deploy an application. (eg. Postgres, Ghost)

The next step is to use the Trident Protect for Data protection and Data migration from the source to the destination cluster.
Refer here for instructions.

Data protection using Trident Protect

This page shows the data protection options for Red Hat OpenShift Container based applications running on VMware vSphere using Trident Protect.

As users take their journey of modernizing their applications with Red Hat OpenShift, a data protection strategy should be in place to protect them from accidental deletion or any other human errors. Often a protection strategy is also required for regulatory or compliance purposes to protect their data from a disaster.

The requirements of data protection varies from reverting back to a point in time copy to automatically failing over to a different fault domain without any human intervention. Many customers pick ONTAP as their preferred storage platform for their Kubernetes applications because of its rich features like multitenancy, multi-protocol, high performance and capacity offerings, replication and caching for multi-site locations, security and flexibility.

For details about protecting and migrating container workloads using Trident Protect, review the following section:
- Protect and migrate Container apps using Trident Protect

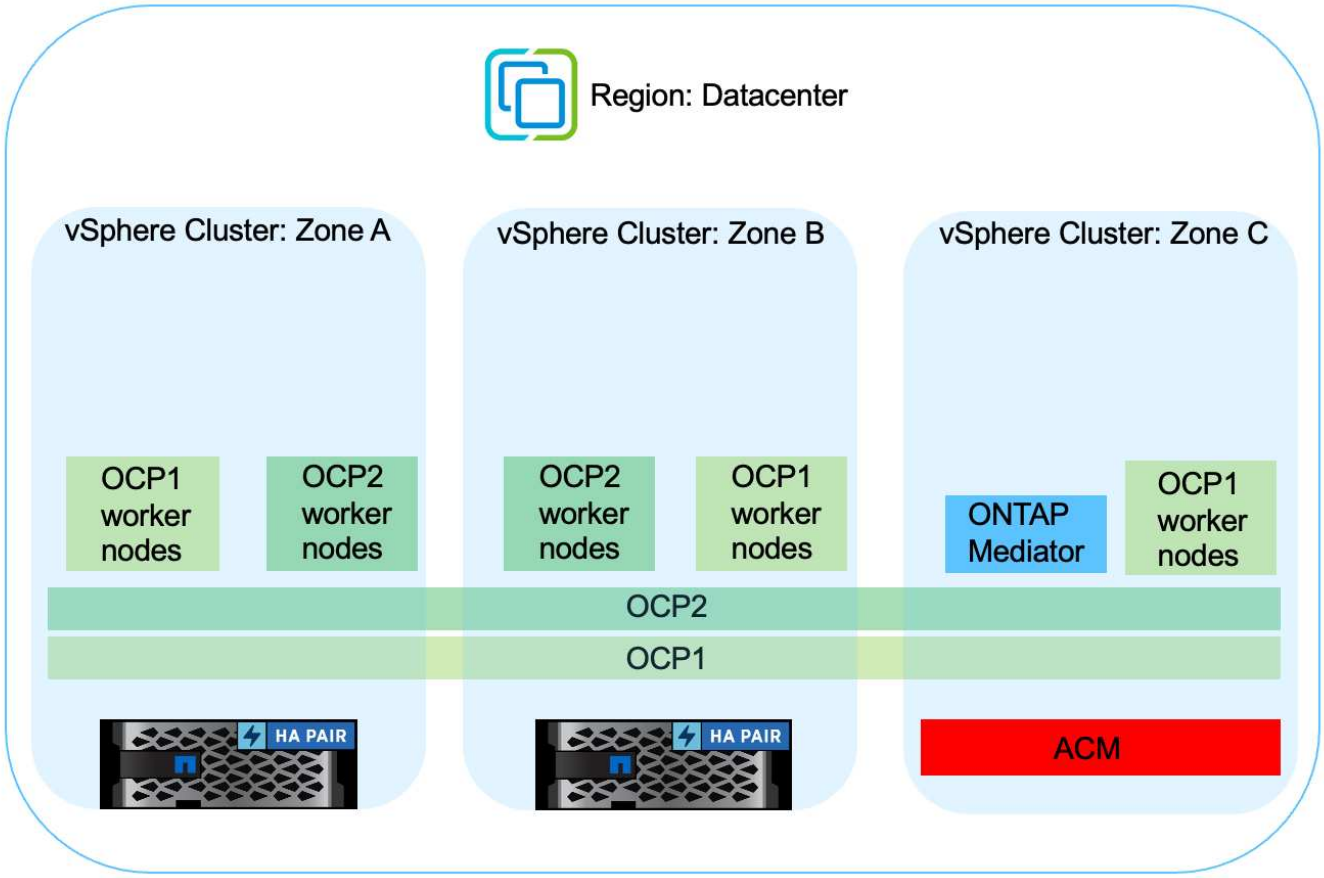
Business Continuity with MetroCluster

Most of our hardware platform for ONTAP has high availability features to protect from device failures avoiding the need to perform diaster recovery. But to protect from fire or any other disaster and to continue the business with zero RPO and low RTO, often a MetroCluster solution is used.

Customers who currently have an ONTAP system can extend to MetroCluster by adding supported ONTAP systems within the distance limitations for providing zone level disaster recovery.
Trident, the CSI (Container Storage Interface) supports NetApp ONTAP including MetroCluster configuration as well as other options like Cloud Volumes ONTAP, Azure NetApp Files, AWS FSx ONTAP, etc. Trident provides

five storage driver options for ONTAP and all are supported for MetroCluster configuration. Refer here for additional details about ONTAP storage drivers supported by Trident.

The MetroCluster solution requires layer 2 network extension or capability to access the same network address from both fault domains. Once MetroCluster configuration is in place, the solution is transparent to application owners as all the volumes in the MetroCluster svm are protected and get the benefits of SyncMirror (zero RPO).



For Trident Backend Configuration (TBC), do not specify the dataLIF and SVM when using MetroCluster configuration. Specify SVM management IP for managementLIF and use vsadmin role credentials.

Details on Trident Protect Data Protection features are available here.

Hybrid cloud with self-managed components

NetApp Solution with Red Hat OpenShift Container platform workloads in Hybrid Cloud

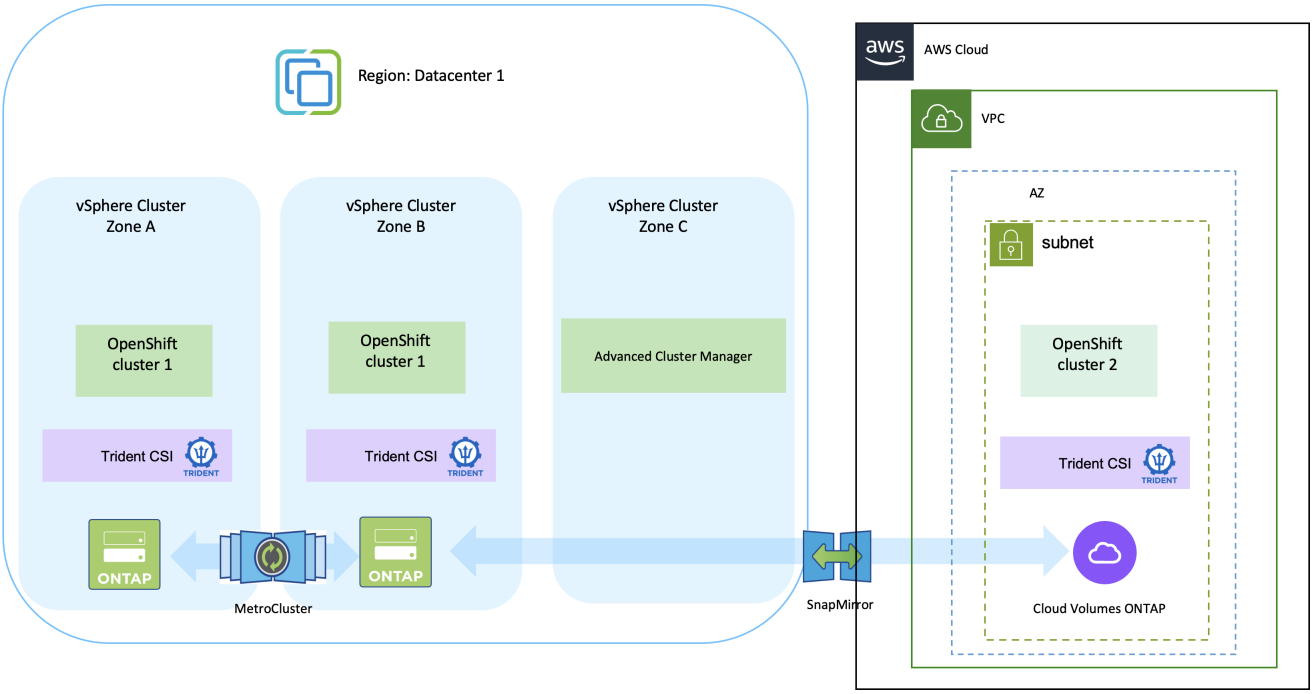
Customers may be at a point in their modernization journey when they are ready to move some select workloads or all workloads from their data centers to the cloud. They may choose to use self-managed OpenShift containers and self-managed NetApp storage in the cloud for various reasons. They should plan and deploy the Red Hat OpenShift container platform (OCP) in the cloud for a successful production-ready environment for

migrating their container workloads from their data centers. Their OCP clusters can be deployed on VMware or Bare Metal in their data centers and on AWS, Azure or Google Cloud in the cloud environment.

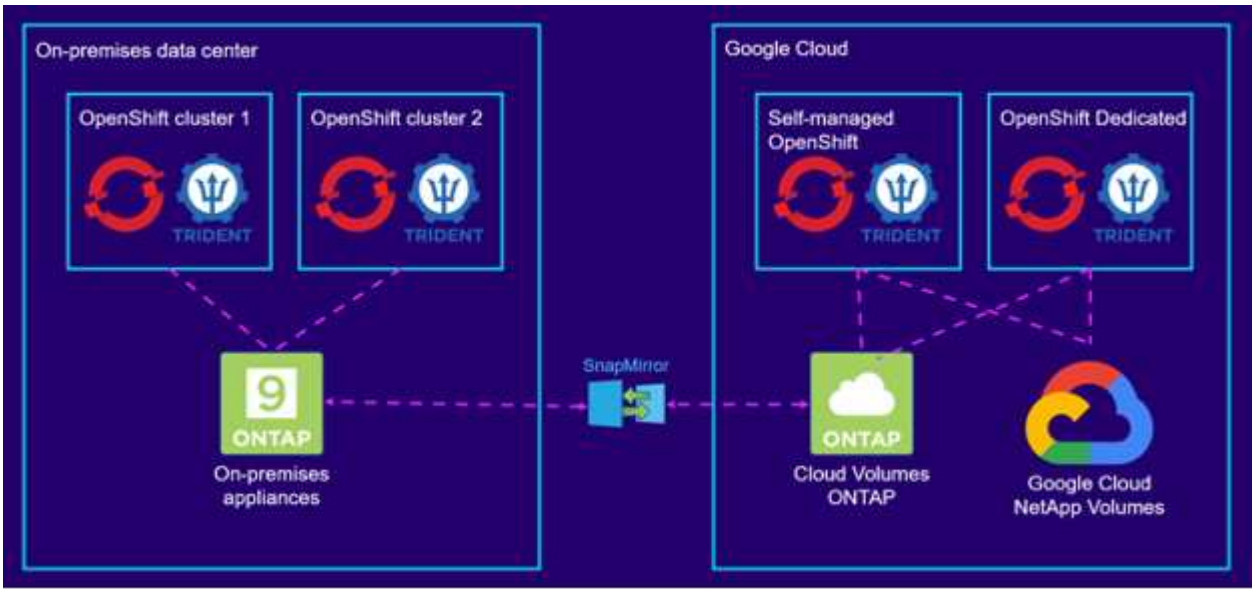
NetApp Cloud Volumes ONTAP storage delivers data protection, reliability, and flexibility for container deployments in AWS, Azure, and Google Cloud. Trident serves as the dynamic storage provisioner to consume the persistent Cloud Volumes ONTAP storage for customers' stateful applications. Trident Protect can be used for the data management requirements of stateful applications such as data protection, migration, and business continuity.

Data protection and migration solution for OpenShift Container workloads in a hybrid cloud using Trident Protect

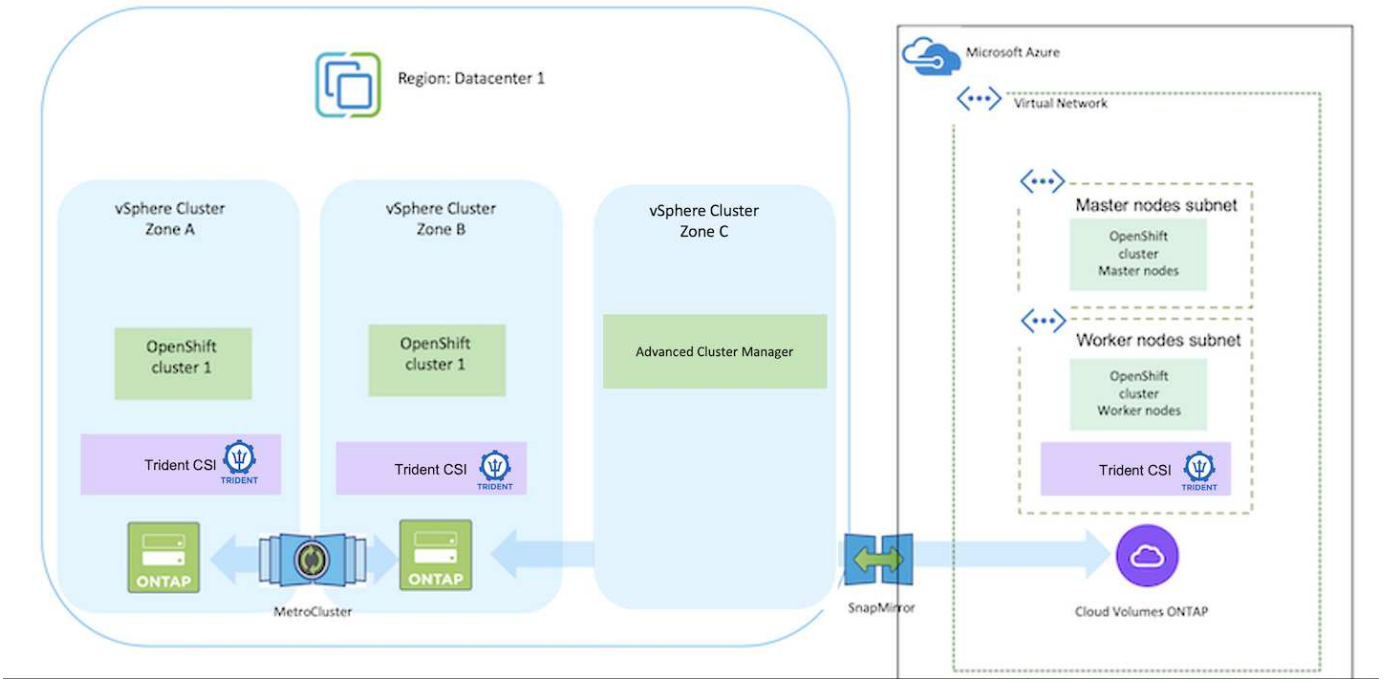
On-premises and AWS



On-premises and Google Cloud



On-premises and Azure Cloud



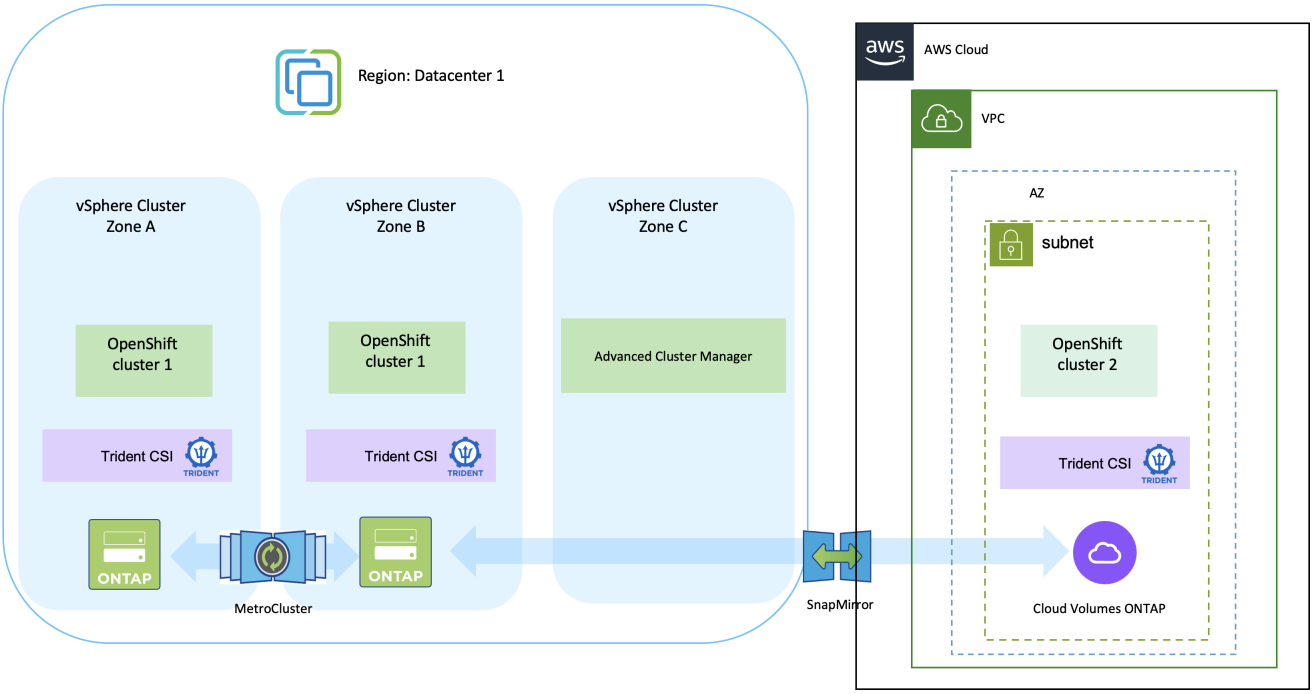
Deploy and configure the Red Hat OpenShift Container platform on AWS

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in AWS and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Trident to provide persistent volumes. Details are provided about the use of Trident Protect to perform data protection and migration activities for the stateful applications.

There are several ways of deploying Red Hat OpenShift Container platform clusters on AWS. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources

section.

Here is a diagram that depicts the clusters deployed on AWS and connected to the data center using a VPN.



The setup process can be broken down into the following steps:

▼ **Install an OCP cluster on AWS from the Advanced Cluster Management.**

- Create a VPC with a site-to-site VPN connection (using pfsense) to connect to the on-premises network.
- On-premises network has internet connectivity.
- Create 3 private subnets in 3 different AZs.
- Create a Route 53 private hosted zone and a DNS resolver for the VPC.

Create OpenShift Cluster on AWS from the Advanced Cluster Management (ACM) Wizard. Refer to instructions here.

You can also create the cluster in AWS from the OpenShift Hybrid Cloud console. Refer here for instructions.

When creating the cluster using the ACM, you have the ability to customize the installation by editing the yaml file after filling in the details in the form view. After the cluster is created, you can ssh login to the nodes of the cluster for troubleshooting or additional manual configuration. Use the ssh key you provided during installation and the username core to login.

▼ **Deploy Cloud Volumes ONTAP in AWS using NetApp Console.**

- Install the connector in on-premises VMware environment. Refer to instructions here.
- Deploy a CVO instance in AWS using the connector. Refer to instructions here.

The connector can also be installed in the cloud environment. Refer here for additional information.

▼ **Install Trident in the OCP Cluster**

- Deploy Trident Operator using Helm.
  Refer to instructions here
- Create a backend and a storage class. Refer to instructions here.

Using CSI Topology feature of Trident for multi-zone architectures

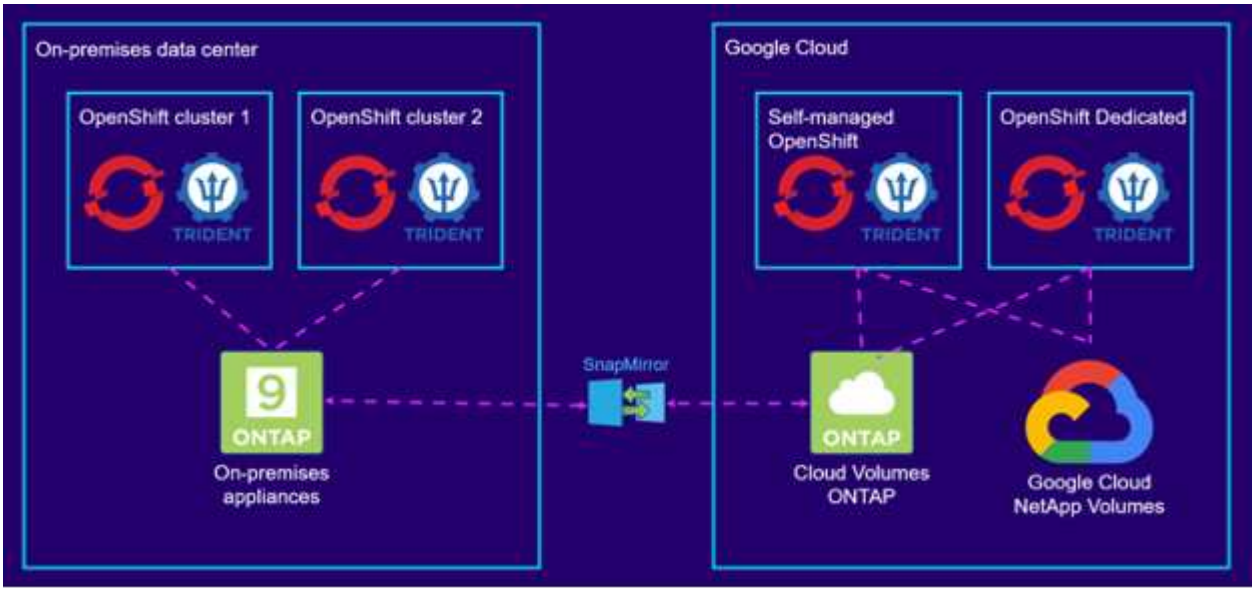
Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer here for additional details.

Kubernetes supports two volume binding modes:
- When ***VolumeBindingMode* is set to *Immediate*** (default), Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.
- When ***VolumeBindingMode* set to *WaitForFirstConsumer***, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.
Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass)
Refer here for additional details.

Deploy and configure the Red Hat OpenShift Container Platform on Google Cloud

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in GCP and deploy stateful applications on them. It shows the use of Google Cloud NetApp Volumes and NetApp Cloud Volumes ONTAP storage with the help of Trident to provide persistent volumes.

Here is a diagram that shows the clusters deployed on GCP and connected to the data center using a VPN.



There are several ways of deploying Red Hat OpenShift Container platform clusters in GCP. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

The setup process can be broken down into the following steps:

Install an OCP cluster on GCP from the CLI

- Ensure that you have met all the prerequisites stated here.
- For the VPN connectivity between on-premises and GCP, a pfsense VM was created and configured. For instructions, see here.
  - The remote gateway address in pfsense can be configured only after you have created a VPN gateway in Google Cloud Platform.
  - The remote network IP addresses for the Phase 2 can be configured only after the OpenShift cluster installation program runs and creates the infrastructure components for the cluster.
  - The VPN in Google Cloud can only be configured after the infrastructure components for the cluster are created by the installation program.
- Now install the OpenShift cluster on GCP.
  - Obtain the installation program and the pull secret and deploy the cluster following the steps provided in the documentation here.
  - The installation creates a VPC network in Google Cloud Platform. It also creates a private zone in Cloud DNS and adds A records.
    - Use the CIDR block address of the VPC network to configure the pfsense and establish the VPN connection. Ensure firewalls are setup correctly.
    - Add A records in the DNS of the on-premises environment using the IP address in the A records of the Google Cloud DNS.
  - The installation of the cluster completes and will provide a kubeconfig file and username and password to login to the console of the cluster.

Deploy Google Cloud NetApp Volumes

- Google Cloud NetApp Volumes can be added to your projects as outlined here.

Deploy Cloud Volumes ONTAP in GCP using NetApp Console

- Install a Console agent in Google Cloud. Refer to instructions here.
- Deploy a CVO instance in Google Cloud using the agent. Refer to instructions here.

Install Trident in the OCP Cluster in GCP

- There are many methods to deploy Trident as shown here.
- For this project, Trident was installed by deploying Trident Operator manually using the instructions here.
- Create backend and a storage classes. Refer to instructions here.

Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer here for additional details.

> Kubernetes supports two volume binding modes:
> - When ***VolumeBindingMode*** **is set to** ***Immediate*** (default), Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.
> - When ***VolumeBindingMode*** **set to** ***WaitForFirstConsumer***, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.
> Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass)
> Refer here for additional details.

Demonstration Video

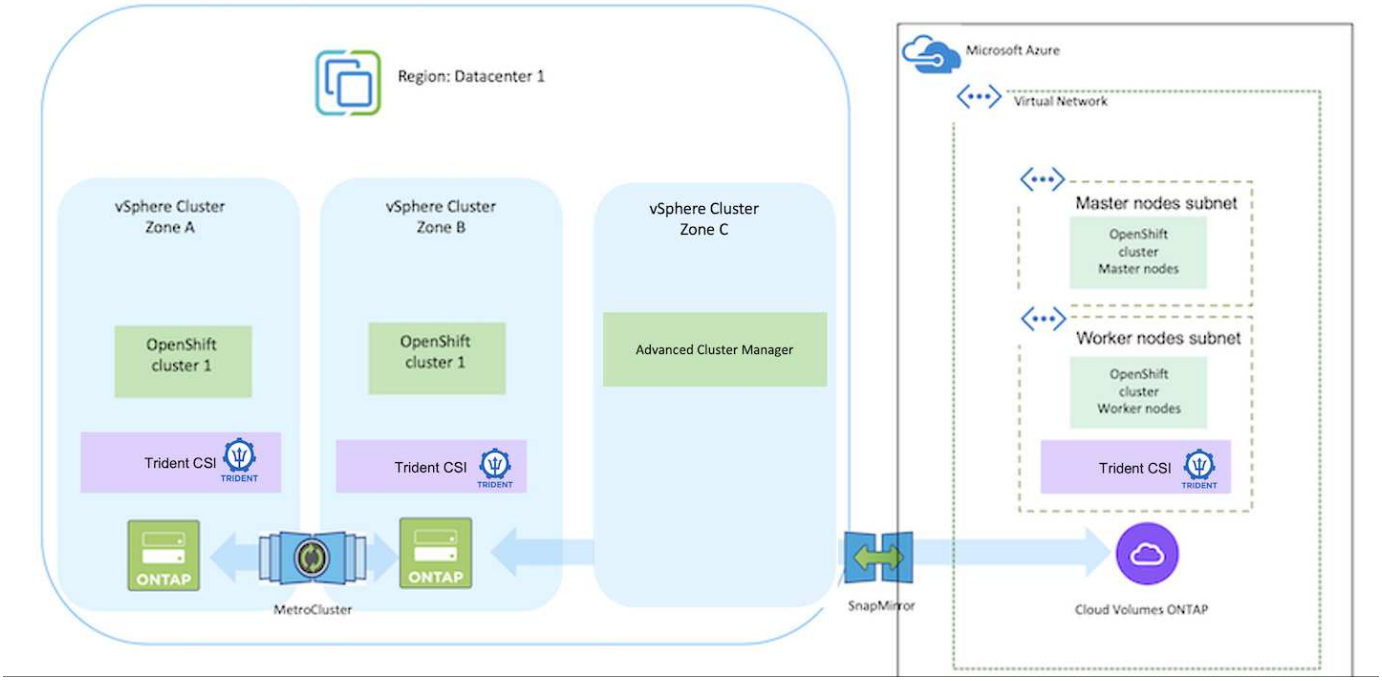
OpenShift Cluster installation on Google Cloud Platform

Importing OpenShift clusters into Trident Protect

Deploy and configure the Red Hat OpenShift Container platform on Azure

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in Azure and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Trident to provide persistent volumes. Details are provided about the use of Trident Protect to perform data protection and migration activities for the stateful applications.

Here is a diagram that shows the clusters deployed on Azure and connected to the data center using a VPN.



There are several ways of deploying Red Hat OpenShift Container platform clusters in Azure. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

The setup process can be broken down into the following steps:

▼ **Install an OCP cluster on Azure from the CLI.**

- Ensure that you have met all the prerequisites stated here.
- Create a VPN, subnets and network security groups and a private DNS zone. Create VPN gateway and site-to-site VPN Connection.
- For the VPN connectivity between on-premises and Azure, a pfsense VM was created and configured. For instructions, see here.
- Obtain the installation program and the pull secret and deploy the cluster following the steps provided in the documentation here.
- The installation of the cluster completes and will provide a kubeconfig file and username and password to login to the console of the cluster.

A sample install-config.yaml file is given below.

```
apiVersion: v1
baseDomain: sddc.netapp.com
compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
```

```
    azure:
      encryptionAtHost: false
      osDisk:
        diskSizeGB: 512
        diskType: "StandardSSD_LRS"
      type: Standard_D2s_v3
      ultraSSDCapability: Disabled
      #zones:
      #- "1"
      #- "2"
      #- "3"
  replicas: 3
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    azure:
      encryptionAtHost: false
      osDisk:
        diskSizeGB: 1024
        diskType: Premium_LRS
      type: Standard_D8s_v3
      ultraSSDCapability: Disabled
  replicas: 3
metadata:
  creationTimestamp: null
  name: azure-cluster
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  azure:
    baseDomainResourceGroupName: ocp-base-domain-rg
    cloudName: AzurePublicCloud
    computeSubnet: ocp-subnet2
    controlPlaneSubnet: ocp-subnet1
    defaultMachinePlatform:
      osDisk:
        diskSizeGB: 1024
```

```
        diskType: "StandardSSD_LRS"
      ultraSSDCapability: Disabled
    networkResourceGroupName: ocp-nc-us-rg
    #outboundType: UserDefinedRouting
    region: northcentralus
    resourceGroupName: ocp-cluster-ncusrg
    virtualNetwork: ocp_vnet_ncus
publish: Internal
pullSecret:
```

▼ **Deploy Cloud Volumes ONTAP in Azure using NetApp Console.**

- Install a NetApp Console agent in Azure. Refer to instructions here.
- Deploy a CVO instance in Azure using the agent. Refer to instructions here.

Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer here for additional details.

> Kubernetes supports two volume binding modes:
> - When ***VolumeBindingMode*** **is set to** ***Immediate*** (default), Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.
> - When ***VolumeBindingMode*** **set to** ***WaitForFirstConsumer***, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.
>
> Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass)
> Refer here for additional details.

Data protection using Trident Protect

For details about protecting and migrating container workloads using Trident Protect, review the following section:
- Protect and migrate Container apps using Trident Protect

Hybrid cloud with provider-managed components

NetApp Solution with Managed Red Hat OpenShift Container platform workloads

Customers may be "born in the cloud" or may be at a point in their modernization journey when they are ready to move some select workloads or all workloads from their data centers to the cloud. They may choose to use provider-managed OpenShift containers

and provider-managed NetApp storage in the cloud for running their workloads. They should plan and deploy the Managed Red Hat OpenShift container clusters in the cloud for a successful production-ready environment for their container workloads. NetApp provides fully managed storage offerings for Managed Red Hat solutions in all three leading public clouds.

Amazon FSx for NetApp ONTAP (FSx ONTAP)

FSx ONTAP delivers data protection, reliability, and flexibility for container deployments in AWS. Trident serves as the dynamic storage provisioner to consume the persistent FSx ONTAP storage for customers' stateful applications.

As ROSA can be deployed in HA mode with control plane nodes spread across multiple availability zones, FSx ONTAP can also be provisioned with Multi-AZ option which provides high availability and protect against AZ failures.

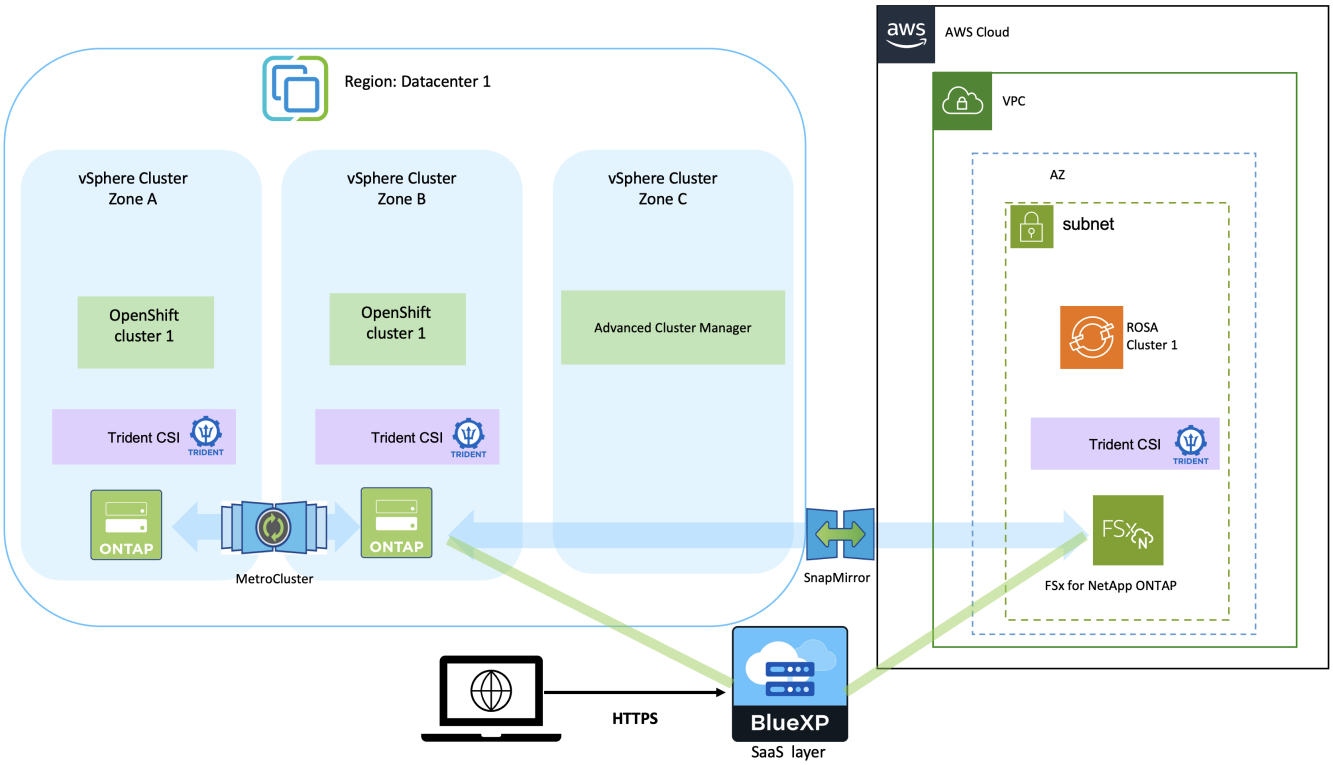
Google Cloud NetApp Volumes

Red Hat OpenShift Dedicated is a fully managed application platform that enables you to quickly build, deploy, and scale applications across the hybrid cloud. Google Cloud NetApp Volumes provides persistent volumes bringing the full suite of ONTAP’s enterprise data management capabilities to OpenShift deployments in Google Cloud.

Deploy and configure the Managed Red Hat OpenShift Container platform on AWS

This section describes a high-level workflow of setting up the Managed Red Hat OpenShift clusters on AWS(ROSA). It shows the use of managed Amazon FSx for NetApp ONTAP (FSx ONTAP) as the storage backend by Trident to provide persistent volumes. Details are provided about the deployment of FSx ONTAP on AWS using NetApp Console. Also, details are provided about the use of NetApp Console and OpenShift GitOps (Argo CD) to perform data protection and migration activities for the stateful applications on ROSA clusters.

Here is a diagram that depicts the ROSA clusters deployed on AWS and using FSx ONTAP as the backend storage.



This solution was verified by using two ROSA clusters in two VPCs in AWS. Each ROSA cluster was integrated with FSx ONTAP using Trident. There are several ways of deploying ROSA clusters and FSx ONTAP in AWS. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

The setup process can be broken down into the following steps:

Install ROSA clusters

- Create two VPCs and set up VPC peering connectivity between the VPCs.
- Refer here for instructions to install ROSA clusters.

Install FSx ONTAP

- Install FSx ONTAP on the VPCs from NetApp Console.
  Refer here for NetApp Console account creation and to get started.
  Refer here for installing FSx ONTAP.
  Refer here for creating a connector in AWS to manage the FSx ONTAP.
- Deploy FSx ONTAP using AWS.
  Refer here for deployment using AWS console.

Install Trident on ROSA clusters (using Helm chart)

- Use Helm chart to install Trident on ROSA clusters.
  Refer to the documentation xref:./openshift/ here.

Integration of FSx ONTAP with Trident for ROSA clusters

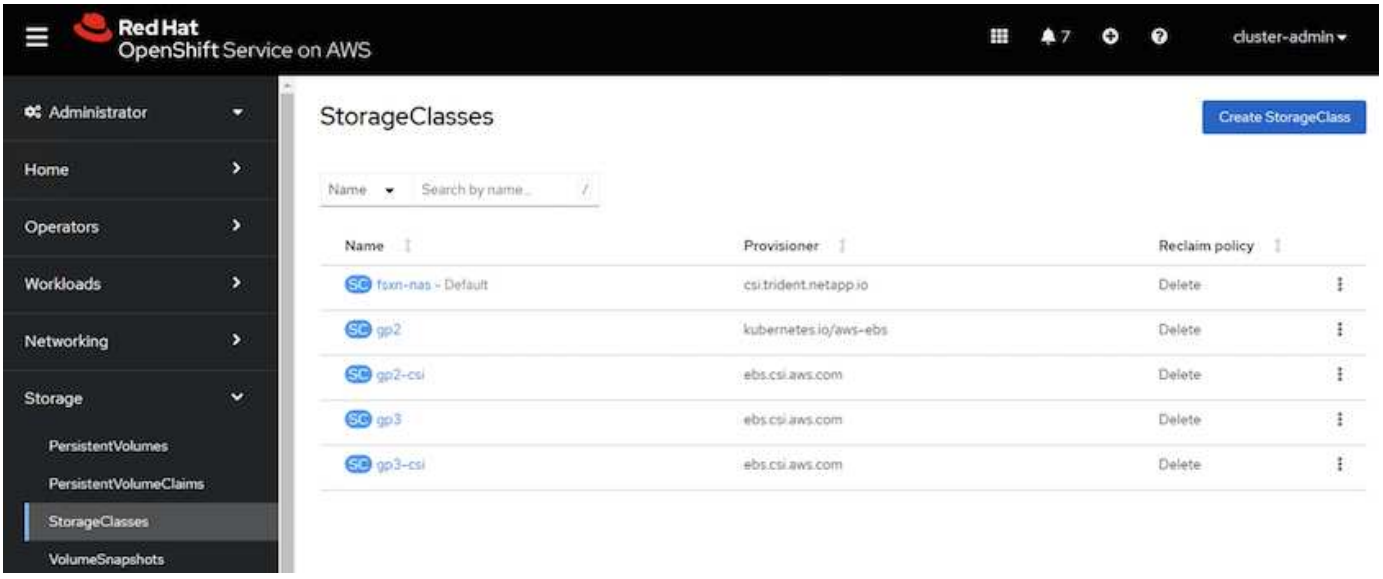
OpenShift GitOps can be utilized to deploy Trident CSI to all managed clusters as they get registered to ArgoCD using ApplicationSet.

```yaml
apiVersion: argoproj.io/v1alpha1
kind: ApplicationSet
metadata:
  name: trident-operator
spec:
  generators:
  - clusters: {}
    # selector:
    #   matchLabels:
    #     tridentversion: '23.04.0'
  template:
    metadata:
      name: '{{nameNormalized}}-trident'
    spec:
      destination:
        namespace: trident
        server: '{{server}}'
      source:
        repoURL: 'https://netapp.github.io/trident-helm-chart'
        targetRevision: 23.04.0
        chart: trident-operator
      project: default
      syncPolicy:
        syncOptions:
          - CreateNamespace=true
```



Create backend and storage classes using Trident (for FSx ONTAP)

- Refer here for details about creating backend and storage class.
- Make the storage class created for FsxN with Trident CSI as default from OpenShift Console. See screenshot below:



Deploy an application using OpenShift GitOps (Argo CD)

- Install OpenShift GitOps operator on the cluster. Refer to instructions here.
- SetUp a new Argo CD instance for the cluster. Refer to instructions here.

Open the console of Argo CD and deploy an app.

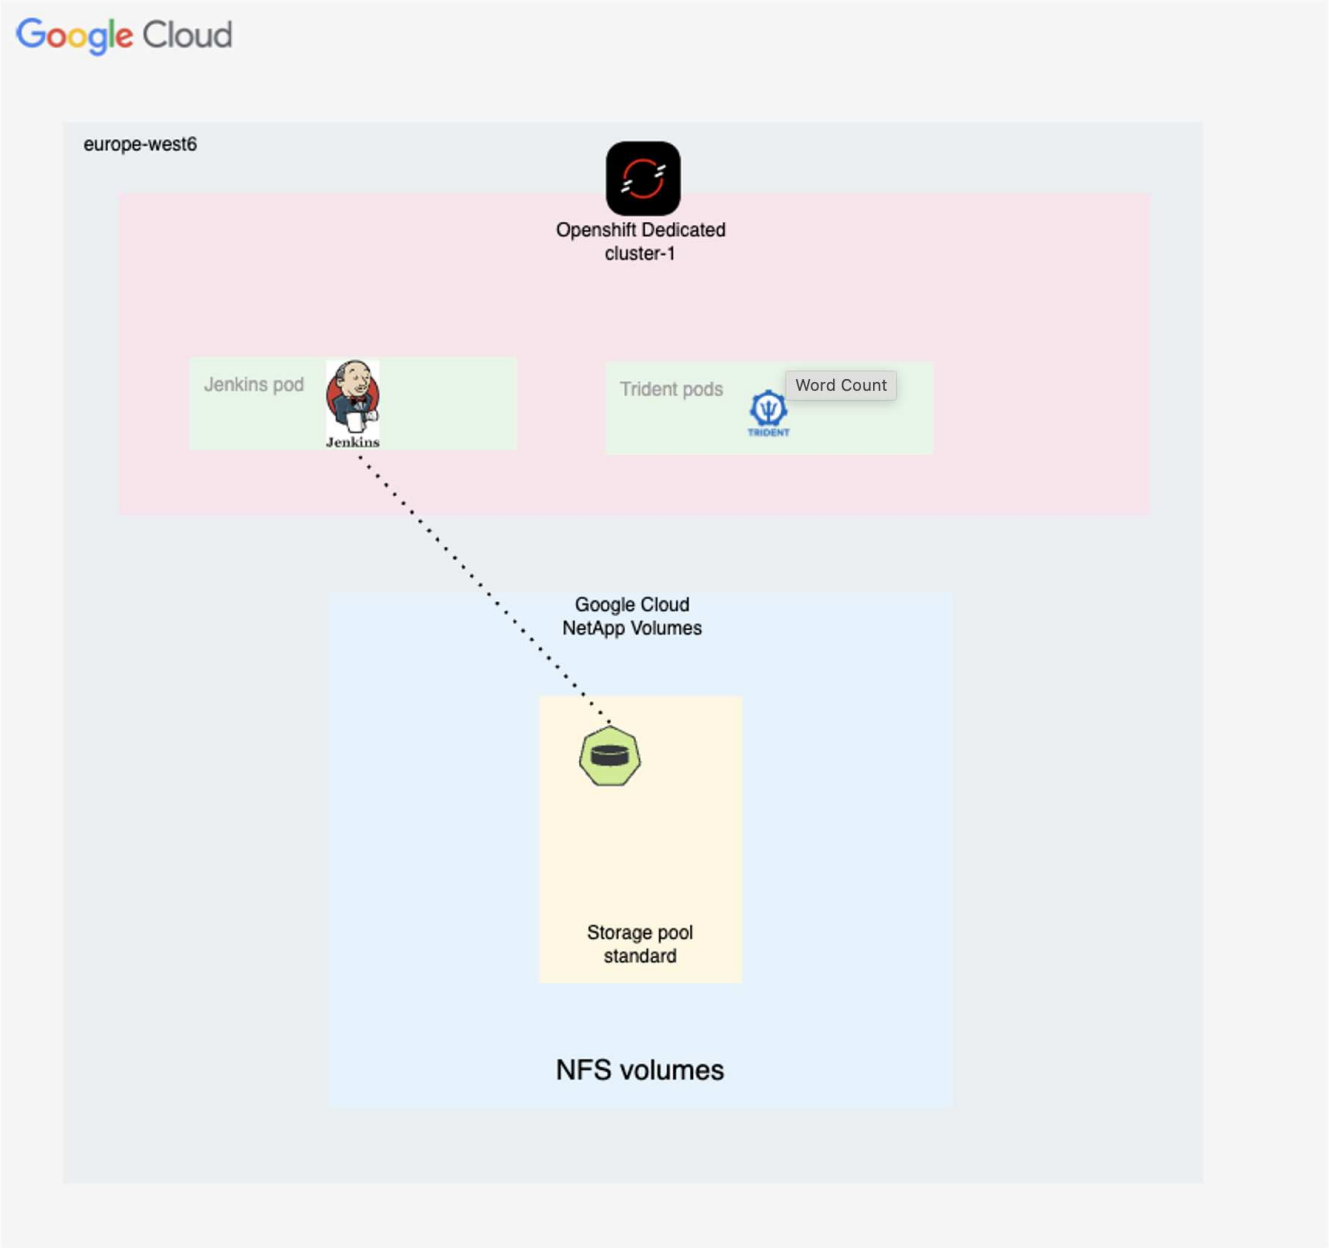
As an example, you can deploy a Jenkins App using Argo CD with a Helm Chart.
When creating the application, the following details were provided:
Project: default
cluster: 'https://kubernetes.default.svc' (without the quotes)
Namespace: Jenkins
The url for the Helm Chart: 'https://charts.bitnami.com/bitnami' (without the quotes)

Helm Parameters:
global.storageClass: fsxn-nas

Deploy and configure OpenShift Dedicated on Google Cloud with Google Cloud NetApp Volumes

This section describes a high-level workflow of setting up OpenShift Dedicated (OSD) clusters on the Google Cloud platform. It shows NetApp Trident using Google Cloud NetApp Volumes as the storage backend to provide persistent volumes for stateful applications running with Kubernetes.

Here is a diagram that depicts an OSD cluster deployed on Google Cloud and using NetApp Volumes as the backend storage.



The setup process can be broken down into the following steps:

Install OSD clusters in Google Cloud

- If you wish to use an existing VPC for the cluster, you must create the VPC, two subnets, a cloud router, and two GCP cloud NATs for the OSD cluster. Refer here for instructions.
- Refer here for instructions to install OSD clusters on GCP using the Customer Cloud Subscription (CCS) billing model. OSD is also included on Google Cloud Marketplace. A video showing how to install OSD using the Google Cloud Marketplace solution is located here.

Enable Google Cloud NetApp Volumes

- Refer here for information on setting up access to Google Cloud NetApp Volumes. Follow all the steps up to and including
- Create a storage pool. Refer here for information on how to set up a storage pool on Google Cloud NetApp Volumes. Volumes for the stateful Kubernetes applications running on OSD will be created within the storage pool.

Install Trident on OSD clusters (using Helm chart)

- Use a Helm chart to install Trident on OSD clusters. Refer here for instructions on how to install the Helm Chart.

Integration of NetApp Volumes with NetApp Trident for OSD clusters

Create backend and storage classes using Trident (for Google Cloud NetApp Volumes)

- Refer here for details about creating backend.
- If any of the current storage classes in kubernetes are marked as default, remove that annotation by editing the storage class.
- Create at least one storage class for NetApp volumes with the Trident CSI provisioner. Make exactly one of the storage classes the default using an annotation. This will allow a PVC to use this storage class when it is not explicitly called out in the PVC manifest. An example with the annotation is shown below.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: gcnv-standard-k8s
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
provisioner: csi.trident.netapp.io
parameters:
  backendType: "google-cloud-netapp-volumes"
  trident.netapp.io/nasType: "nfs"
allowVolumeExpansion: true
```

Deploy an application using OpenShift GitOps (Argo CD)

- Install OpenShift GitOps operator on the cluster. Refer to instructions here.
- SetUp a new Argo CD instance for the cluster. Refer to instructions here.

Open the console of Argo CD and deploy an app.
As an example, you can deploy a Jenkins App using Argo CD with a Helm Chart.
When creating the application, the following details were provided:
Project: default
cluster: 'https://kubernetes.default.svc' (without the quotes)
Namespace: Jenkins
The url for the Helm Chart: 'https://charts.bitnami.com/bitnami' (without the quotes)

Data protection

For details about protecting and migrating container workloads using Trident Protect, review the following section:
- Protect and migrate Container apps using Trident Protect

Additional NetApp Hybrid Multicloud solutions for Red Hat OpenShift workloads

See the following additional solutions:

- Red Hat OpenShift Container solution on-premises
- Red Hat OpenShift Service on AWS Container solution
- Red Hat OpenShift Virtualization solution on-premises
- Red Hat OpenShift Service on AWS Virtualization solution

Protect data using Trident Protect

Data protection for Container Apps in OpenShift Container Platform using Trident Protect

This section of the reference document provides details for creating snapshots and backups of Container Apps using Trident Protect. NetApp Trident Protect provides advanced application data management capabilities that enhance the functionality and availability of stateful Kubernetes applications backed by NetApp ONTAP storage systems and the NetApp Trident CSI storage provisioner.
Trident Protect creates application snapshots and backups. This means that snapshots and backups are created not only for application data in persistent volumes, but also for application metadata. The snapshots and backups created by Trident Protect can be stored in any of the following object storage and restored from them at a later point.

- AWS S3
- Azure Blob storage
- Google Cloud Storage
- ONTAP S3
- StorageGrid
- Any other S3-compatible storage

Trident Protect uses the Kubernetes model of role-based access control (RBAC).
By default, Trident Protect provides a single system namespace called trident-protect and its associated default service account. If you have an organization with many users or specific security needs, you can use the RBAC features of Trident Protect to gain more granular control over access to resources and namespaces.

Additional information on RBAC in Trident Protect can be found in the
Trident Protect documentation



The cluster administrator has access to resources in the default trident-protect namespace and can also access resources in all other namespaces. Users cannot create application data management custom resource (CRs) like Snapshot and Backup CRs in the trident-protect namespace. As a best practice, users will need to create those CRs in the application namespace.

Trident Protect can be installed using the instructions provided in the documentation here. This section will show the workflow for the data protection of container applications and restoration of the applications using Trident Protect.
1. Snapshot creation (on demand or scheduled)
2. Restore from Snapshot (restore to same and different namespace)
3. Backup creation

4. Restore from Backup

▼ **Prerequisite**

Prior to creating the snapshots and backups for an application, object storage must be configured in Trident Protect to store the snapshots and backups. This is done using the bucket CR. Only administrators can create a bucket CR and configure it.
The bucket CR is known as AppVault in Trident Protect. AppVault objects are the declarative Kubernetes workflow representation of a storage bucket. An AppVault CR contains the configurations necessary for a bucket to be used in protection operations, such as backups, snapshots, restore operations, and SnapMirror replication.

In this example, we will show the use of ONTAP S3 as object storage.
Here is the workflow for creating AppVault CR for ONTAP S3:
1. Create an S3 object store server in the SVM in ONTAP cluster.
2. Create a bucket in the object store server.
3. Create an S3 user in the SVM. Keep the access key and the secret key in a safe location.
4. In OpenShift, create a secret to store the ONTAP S3 credentials.
5. Create an AppVault object for ONTAP S3

Configure Trident Protect AppVault for ONTAP S3

Sample YAML file for configuring Trident Protect with ONTAP S3 as the AppVault

```
# alias tp='tridentctl-protect'

appvault-secret.yaml

apiVersion: v1
stringData:
  accessKeyID: "<access key id created for a user to access ONTAP S3
bucket>"
  secretAccessKey: "corresponding Secret Access Key"
#data:
# base 64 encoded values
#  accessKeyID: <base64 access key id created for a user to access ONTAP
S3 bucket>
#  secretAccessKey: <base 64  Secret Access Key>
kind: Secret
metadata:
  name: appvault-secret
  namespace: trident-protect
type: Opaque

appvault.yaml

apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: ontap-s3-appvault
```

```
  namespace: trident-protect
spec:
  providerConfig:
    azure:
      accountName: ""
      bucketName: ""
      endpoint: ""
    gcp:
      bucketName: ""
      projectID: ""
    s3:
      bucketName: <bucket-name for storing the snapshots and backups>
      endpoint: <endpoint IP for S3>
      secure: "false"
      skipCertValidation: "true"
  providerCredentials:
    accessKeyID:
      valueFromSecret:
        key: accessKeyID
        name: appvault-secret
    secretAccessKey:
      valueFromSecret:
        key: secretAccessKey
        name: appvault-secret
  providerType: OntapS3

# oc create -f appvault-secret.yaml -n trident-protect
# oc create -f appvault.yaml -n trident-protect
```

```
[root@localhost openshift-benchmark]#
[root@localhost openshift-benchmark]# tp get appvault -n trident-protect
+------------------+----------+-----------+------+-------+
|       NAME       | PROVIDER |   STATE   | AGE  | ERROR |
+------------------+----------+-----------+------+-------+
| ontap-s3-appvault | OntapS3  | Available | 2d2h |       |
+------------------+----------+-----------+------+-------+
[root@localhost openshift-benchmark]#
```

Sample yaml file for installing postgresql app

```
postgres.yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: postgres
```

```
spec:
  replicas: 1
  selector:
    matchLabels:
      app: postgres
  template:
    metadata:
      labels:
        app: postgres
    spec:
      containers:
      - name: postgres
        image: postgres:14
        env:
        - name: POSTGRES_USER
          #value: "myuser"
          value: "admin"
        - name: POSTGRES_PASSWORD
          #value: "mypassword"
          value: "adminpass"
        - name: POSTGRES_DB
          value: "mydb"
        - name: PGDATA
          value: "/var/lib/postgresql/data/pgdata"
        ports:
        - containerPort: 5432
        volumeMounts:
        - name: postgres-storage
          mountPath: /var/lib/postgresql/data
      volumes:
      - name: postgres-storage
        persistentVolumeClaim:
          claimName: postgres-pvc
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: postgres-pvc
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 5Gi
---
apiVersion: v1
```

```
kind: Service
metadata:
  name: postgres
spec:
  selector:
    app: postgres
  ports:
  - protocol: TCP
    port: 5432
    targetPort: 5432
  type: ClusterIP

Now create the Trident Protect application CR for the postgres app.
Include the objects in the namespace postgres and create it in the
postgres namespace.
# tp create app postgres-app --namespaces postgres -n postgres
```

```
[root@localhost RedHat]# tp get app -n postgres
+--------------+------------+-------+-----+
|     NAME     | NAMESPACES | STATE | AGE |
+--------------+------------+-------+-----+
| postgres-app | postgres   | Ready | 24s |
+--------------+------------+-------+-----+
[root@localhost RedHat]#
```

▼ **Create Snapshots**

Creating an on-demand snapshot

```
# tp create snapshot postgres-snap1 --app postgres-app --appvault ontap-
s3-appvault -n postgres
Snapshot "postgres-snap1" created.
```

```
[root@localhost RedHat]# tp get snapshot -n postgres
+-----------------+--------------+-----------+-----+-------+
|      NAME       |   APP REF    |   STATE   | AGE | ERROR |
+-----------------+--------------+-----------+-----+-------+
| postgres-snap1 | postgres-app | Completed | 19s |       |
+-----------------+--------------+-----------+-----+-------+
[root@localhost RedHat]#
```

```
[root@localhost DataProtection]# oc get all,pvc,volumesnapshot -n postgres
Warning: apps.openshift.io/v1 DeploymentConfig is deprecated in v4.14+, unavailable in v4.10000+
Warning: kubevirt.io/v1 VirtualMachineInstancePresets is now deprecated and will be removed in v2.
NAME                            READY   STATUS    RESTARTS   AGE
pod/postgres-cd9d6ccb-jfx49     1/1     Running   0          3h47m

NAME               TYPE        CLUSTER-IP       EXTERNAL-IP   PORT(S)    AGE
service/postgres   ClusterIP   172.30.132.112   <none>        5432/TCP   3h47m

NAME                       READY   UP-TO-DATE   AVAILABLE   AGE
deployment.apps/postgres   1/1     1            1           3h47m

NAME                                  DESIRED   CURRENT   READY   AGE
replicaset.apps/postgres-cd9d6ccb     1         1         1       3h47m

NAME                                                STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAG
ECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/data-postgres-postgresql-0    Bound    pvc-9f89514e-3f2c-41ad-b7a3-792cea503f00   8Gi        RWO            sc-zon
ea-nas   <unset>                 4h40m
persistentvolumeclaim/postgres-pvc                  Bound    pvc-951a9910-9edb-48ae-898a-1aed9aa25dc3   5Gi        RWO            sc-zon
ea-nas   <unset>                 3h47m

NAME                                                                                                                                READYTO
USE   SOURCEPVC                  SOURCESNAPSHOTCONTENT   RESTORESIZE   SNAPSHOTCLASS           SNAPSHOTCONTENT
           CREATIONTIME   AGE
volumesnapshot.snapshot.storage.k8s.io/snapshot-2e94d04c-c8ea-446a-8d47-64b0bee67107-pvc-951a9910-9edb-48ae-898a-1aed9aa25dc3   true
      postgres-pvc                                      53676Ki       trident-snapshotclass   snapcontent-796ea7f8-59a0-493e-bbd8-3a
e76fe9036c   13m            13m
volumesnapshot.snapshot.storage.k8s.io/snapshot-2e94d04c-c8ea-446a-8d47-64b0bee67107-pvc-9f89514e-3f2c-41ad-b7a3-792cea503f00   true
      data-postgres-postgresql-0                        368Ki         trident-snapshotclass   snapcontent-86a464d4-ffd8-4279-9cf7-88
88a097c001   13m            13m
```

Creating a Schedule

Using the following command, Snapshots will be created daily at 15:33 and two snapshots and backups will be retained.

```
# tp create schedule schedule1 --app postgres-app --appvault ontap-s3-
appvault --backup-retention 2 --snapshot-retention 2 --granularity Daily
--hour 15 --minute 33 --data-mover Restic -n postgres
Schedule "schedule1" created.
```

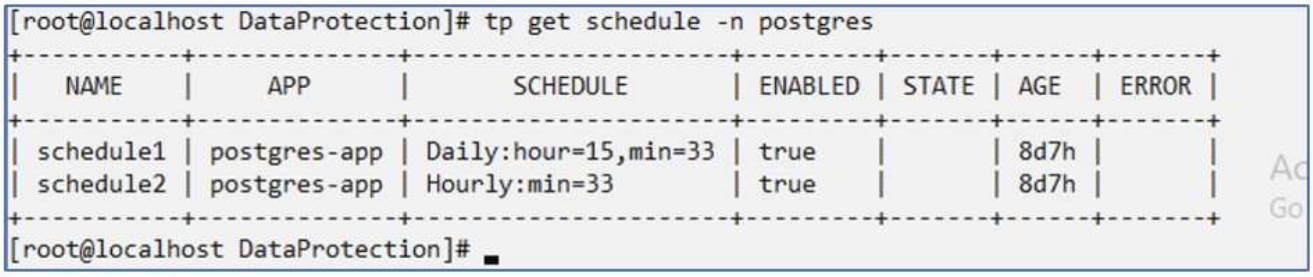
```
[root@localhost DataProtection]# tp get schedule -n postgres
+-----------+--------------+----------------------+---------+-------+-----+-------+
|   NAME    |     APP      |       SCHEDULE       | ENABLED | STATE | AGE | ERROR |
+-----------+--------------+----------------------+---------+-------+-----+-------+
| schedule1 | postgres-app | Daily:hour=15,min=33 | true    |       | 17s |       |
+-----------+--------------+----------------------+---------+-------+-----+-------+
[root@localhost DataProtection]#
```

Creating a Schedule using yaml

```
# tp create schedule schedule2 --app postgres-app --appvault ontap-s3-
appvault --backup-retention 2 --snapshot-retention 2 --granularity Daily
--hour 15 --minute 33 --data-mover Restic -n postgres --dry-run >
hourly-snapshotschedule.yaml

cat hourly-snapshotschedule.yaml

apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  creationTimestamp: null
  name: schedule2
  namespace: postgres
spec:
  appVaultRef: ontap-s3-appvault
  applicationRef: postgres-app
  backupRetention: "2"
  dataMover: Restic
  dayOfMonth: ""
  dayOfWeek: ""
  enabled: true
  granularity: Hourly
  #hour: "15"
  minute: "33"
  recurrenceRule: ""
  snapshotRetention: "2"
status: {}
```



You can see snapshots created on this schedule.

```
[root@localhost DataProtection]# tp get snap -n postgres
+-----------------------------+--------------+-----------+-------+-------+
|            NAME             |   APP REF    |   STATE   |  AGE  | ERROR |
+-----------------------------+--------------+-----------+-------+-------+
| hourly-3f1ee-20250214183300 | postgres-app | Completed | 19s   |       |
| postgres-snap1              | postgres-app | Completed | 1h25m |       |
+-----------------------------+--------------+-----------+-------+-------+
[root@localhost DataProtection]#
```

Volume snapshots are also created.

```
[root@localhost DataProtection]# oc get volumesnapshots -n postgres
NAME                                                                                                  READYTOUSE   SOURCEPVC
                                CREATIONTIME   AGE
snapshot-2e94d04c-c8ea-446a-8d47-64b0bee67107-pvc-951a9910-9edb-48ae-898a-1aed9aa25dc3   true         postgres-pvc
-59a0-493e-bbd8-3ae76fe9036c   114m           114m
snapshot-2e94d04c-c8ea-446a-8d47-64b0bee67107-pvc-9f89514e-3f2c-41ad-b7a3-792cea503f00   true         data-postgres-postgresql-0
-ffd8-4279-9cf7-8888a097c001   114m           114m
snapshot-ce75a274-ecb2-48c9-a0a5-94c10f8e6cb1-pvc-951a9910-9edb-48ae-898a-1aed9aa25dc3   true         postgres-pvc
-7adc-4042-a8c9-7606d1103ead   30m            30m
snapshot-ce75a274-ecb2-48c9-a0a5-94c10f8e6cb1-pvc-9f89514e-3f2c-41ad-b7a3-792cea503f00   true         data-postgres-postgresql-0
-3a82-43f6-9868-dcadd2cc04e2   30m            30m
```

▼ Delete the Application to simulate loss of application

```
# oc delete deployment/postgres -n postgres
# oc get pod,pvc -n postgres
No resources found in postgres namespace.
```

▼ Restore from Snapshot to the same namespace

```
# tp create sir postgres-sir --snapshot postgres/hourly-3f1ee-
20250214183300 -n postgres
SnapshotInplaceRestore "postgres-sir" created.
```

```
[root@localhost DataProtection]# tp get sir -n postgres
+--------------+--------------------+-----------+-------+-------+
|     NAME     |      APPVAULT      |   STATE   |  AGE  | ERROR |
+--------------+--------------------+-----------+-------+-------+
| postgres-sir | ontap-s3-appvault | Completed | 2m39s |       |
+--------------+--------------------+-----------+-------+-------+
```

Application and its PVCis restored to the same namespace.

```
[root@localhost DataProtection]# oc get pods,pvc -n postgres
NAME                     READY   STATUS    RESTARTS   AGE
pod/postgres-cd9d6ccb-x85tg   1/1     Running   0          98s

NAME                                          STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/data-postgres-postgresql-0   Bound    pvc-b2cd67fd-fe4d-49b1-9e06-a53bf7be575e   8Gi        RWO            sc-zonea-nas   <unset>                 105s
persistentvolumeclaim/postgres-pvc                 Bound    pvc-2d549395-0cc6-4529-b2b9-7361bfb14fa8   5Gi        RWO            sc-zonea-nas   <unset>                 105s
[root@localhost DataProtection]#
```

▼ **Restore from Snapshot to a different namespace**

```
# tp create snapshotrestore postgres-restore --snapshot postgres/hourly-3f1ee-20250214183300 --namespace-mapping postgres:postgres-restore -n postgres-restore
SnapshotRestore "postgres-restore" created.
```

```
[root@localhost DataProtection]# tp get snapshotrestore -n postgres-restore
+------------------+--------------------+-----------+-------+-------+
|       NAME       |     APPVAULT       |   STATE   |  AGE  | ERROR |
+------------------+--------------------+-----------+-------+-------+
| postgres-restore | ontap-s3-appvault | Completed | 1m15s |       |
+------------------+--------------------+-----------+-------+-------+
```

You can see that the application has been restored to a new namespace.

```
[root@localhost DataProtection]#
[root@localhost DataProtection]# oc get pods,pvc -n postgres-restore
NAME                          READY   STATUS    RESTARTS   AGE
pod/postgres-cd9d6ccb-pfxw9   1/1     Running   0          8d

NAME                                                STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
persistentvolumeclaim/data-postgres-postgresql-0    Bound    pvc-3dff4e42-828d-49f9-84e9-4daf75479292   8Gi        RWO            sc-zonea-nas
persistentvolumeclaim/postgres-pvc                  Bound    pvc-f57321b2-f21f-4eb7-8f86-4a17f915318b   5Gi        RWO            sc-zonea-nas
[root@localhost DataProtection]#
```

▼ **Create Backups**

Creating an on-demand Backup

```
# tp create backup postgres-backup1 --app postgres-app --appvault ontap-s3-appvault -n postgres
Backup "postgres-backup1" created.
```

```
[root@localhost DataProtection]# tp get backup -n postgres
+-----------------------------+--------------+-----------+--------+-------+
|            NAME             |   APP REF    |   STATE   |  AGE   | ERROR |
+-----------------------------+--------------+-----------+--------+-------+
| backup1                     | postgres-app | Completed | 5d12h  |       |
| daily-feac1-20250222153300  | postgres-app | Completed | 1d10h  |       |
| daily-feac1-20250223153300  | postgres-app | Completed | 10h36m |       |
| hourly-3f1ee-20250224003300 | postgres-app | Completed | 1h36m  |       |
| hourly-3f1ee-20250224013300 | postgres-app | Completed | 36m27s |       |
| postgres-backup1            | postgres-app | Completed | 6m19s  |       |
+-----------------------------+--------------+-----------+--------+-------+
```

Creating Schedule for Backup

The daily and the hourly backups in the list above are created from the schedule set up previously.

```
# tp create schedule schedule1 --app postgres-app --appvault ontap-s3-
```

```
appvault --backup-retention 2 --snapshot-retention 2 --granularity Daily
--hour 15 --minute 33 --data-mover Restic -n postgres
Schedule "schedule1" created.
```

```
[root@localhost DataProtection]# tp get schedule -n postgres
+-----------+--------------+----------------------+---------+-------+------+-------+
|   NAME    |     APP      |       SCHEDULE       | ENABLED | STATE | AGE  | ERROR |
+-----------+--------------+----------------------+---------+-------+------+-------+
| schedule1 | postgres-app | Daily:hour=15,min=33 | true    |       | 9d8h |       |
| schedule2 | postgres-app | Hourly:min=33        | true    |       | 9d8h |       |
+-----------+--------------+----------------------+---------+-------+------+-------+
```

▼ **Restore from backup**

Delete the application and PVCs to simulate a data loss.

```
[root@localhost DataProtection]# oc get pods -n postgres
NAME                        READY   STATUS    RESTARTS   AGE
postgres-cd9d6ccb-dftkt     1/1     Running   0          11s
[root@localhost DataProtection]# oc get deployment -n postgres
NAME       READY   UP-TO-DATE   AVAILABLE   AGE
postgres   1/1     1            1           20s
[root@localhost DataProtection]# oc delete deployment/postgres -n postgres
deployment.apps "postgres" deleted
[root@localhost DataProtection]# oc get pods -n postgres
No resources found in postgres namespace.
[root@localhost DataProtection]# oc get pvc -n postgres
NAME                          STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   VOLUMEATTRIBUTE
SCLASS   AGE
data-postgres-postgresql-0    Bound    pvc-b2cd67fd-fe4d-49b1-9e06-a53bf7be575e   8Gi        RWO            sc-zonea-nas   <unset>
         5d13h
postgres-pvc                  Bound    pvc-2d549395-0cc6-4529-b2b9-7361bfb14fa8   5Gi        RWO            sc-zonea-nas   <unset>
         5d13h
[root@localhost DataProtection]# oc delete pvc/data-postgres-postgresql-0  -n postgres
persistentvolumeclaim "data-postgres-postgresql-0" deleted
[root@localhost DataProtection]# oc delete pvc/postgres-pvc -n postgres
persistentvolumeclaim "postgres-pvc" deleted
[root@localhost DataProtection]# oc get pvc -n postgres
No resources found in postgres namespace.
[root@localhost DataProtection]#
```

Activate Windows
Go to Settings to activate Windows.

Restore to same namespace

#tp create bir postgres-bir --backup postgres/hourly-3f1ee-20250224023300 -n postgres
BackupInplaceRestore "postgres-bir" created.

```
[root@localhost DataProtection]# tp get bir -n postgres
+--------------+--------------------+-----------+-------+-------+
|     NAME     |      APPVAULT      |   STATE   |  AGE  | ERROR |
+--------------+--------------------+-----------+-------+-------+
| postgres-bir | ontap-s3-appvault | Completed | 2m19s |       |
+--------------+--------------------+-----------+-------+-------+
```

The application and the PVCs are restored in the same namespace.

```
[root@localhost DataProtection]# oc get pods -n postgres
NAME                        READY   STATUS    RESTARTS   AGE
postgres-cd9d6ccb-t857w     1/1     Running   0          10m
[root@localhost DataProtection]# oc get pvc -n postgres
NAME                          STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   VOLUMEATTRIBUT
ESCLASS   AGE
data-postgres-postgresql-0    Bound    pvc-0a849c19-16fe-466f-9733-85e82a8b1677   8Gi        RWO            sc-zonea-nas   <unset>
          10m
postgres-pvc                  Bound    pvc-ded304ea-02d4-4225-b606-63007666ad66   5Gi        RWO            sc-zonea-nas   <unset>
          10m
```

Restore to a different namespace

Create a new namespace.

Restore from a backup to the new namespace.

```
[root@localhost DataProtection]# oc create ns postgres-restore-from-backup
namespace/postgres-restore-from-backup created
[root@localhost DataProtection]# tp create backuprestore postgres-restore-from-backup --backup postgres/postgres-backup1 --namespace-map
ping postgres:postgres-restore-from-backup -n postgres-restore-from-backup
BackupRestore "postgres-restore-from-backup" created.
[root@localhost DataProtection]# oc get backuprestore -n postgres-restore-from-backup
NAME                           STATE     ERROR   AGE
postgres-restore-from-backup   Running           37s
[root@localhost DataProtection]# oc get backuprestore -n postgres-restore-from-backup
NAME                           STATE     ERROR   AGE
postgres-restore-from-backup   Running           56s
[root@localhost DataProtection]# oc get backuprestore -n postgres-restore-from-backup
NAME                           STATE       ERROR   AGE
postgres-restore-from-backup   Completed           2m52s
[root@localhost DataProtection]# oc get pods -n postgres-restore-from-backup
NAME                        READY   STATUS    RESTARTS   AGE
postgres-cd9d6ccb-p659p     1/1     Running   0          2m9s
[root@localhost DataProtection]# oc get pvc -n postgres-restore-from-backup
NAME                          STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   VOLUMEATTRIBUT
ESCLASS   AGE
data-postgres-postgresql-0    Bound    pvc-36df7399-95da-4c67-a621-af9434015bdb   8Gi        RWO            sc-zonea-nas   <unset>
          2m18s
postgres-pvc                  Bound    pvc-633de3aa-a4f9-4f3b-93cc-e91afbd4fe02   5Gi        RWO            sc-zonea-nas   <unset>
          2m18s
[root@localhost DataProtection]#
```

Activate Windows
Go to Settings to activate Windows.

▼ Migrate Applications

To clone or migrate an application to a different cluster (perform a cross-cluster clone), create a backup on the source cluster, and then restore the backup to a different cluster. Make sure that Trident Protect is installed on the destination cluster.

On the source cluster, perform the steps as shown in the image below:

```
[root@localhost DataProtection]# tp create backup postgres-backup-cluster1 --app postgres-app --appvault ontap-s3-appvault -n postgres
Backup "postgres-backup-cluster1" created.
[root@localhost DataProtection]# tp get backup -n postgres
+-----------------------------+--------------+-----------+--------+-------+
|            NAME             |   APP REF    |   STATE   |  AGE   | ERROR |
+-----------------------------+--------------+-----------+--------+-------+
| backup1                     | postgres-app | Completed | 5d14h  |       |
| daily-feac1-20250222153300  | postgres-app | Completed | 1d12h  |       |
| daily-feac1-20250223153300  | postgres-app | Completed | 12h18m |       |
| hourly-3f1ee-20250224023300 | postgres-app | Completed | 1h18m  |       |
| hourly-3f1ee-20250224033300 | postgres-app | Completed | 18m26s |       |
| postgres-backup-cluster1    | postgres-app | Running   | 35s    |       |
| postgres-backup1            | postgres-app | Completed | 1h48m  |       |
+-----------------------------+--------------+-----------+--------+-------+
[root@localhost DataProtection]# tp get backup -n postgres
+-----------------------------+--------------+-----------+--------+-------+
|            NAME             |   APP REF    |   STATE   |  AGE   | ERROR |
+-----------------------------+--------------+-----------+--------+-------+
| backup1                     | postgres-app | Completed | 5d14h  |       |
| daily-feac1-20250222153300  | postgres-app | Completed | 1d12h  |       |
| daily-feac1-20250223153300  | postgres-app | Completed | 12h19m |       |
| hourly-3f1ee-20250224023300 | postgres-app | Completed | 1h19m  |       |
| hourly-3f1ee-20250224033300 | postgres-app | Completed | 19m41s |       |
| postgres-backup-cluster1    | postgres-app | Completed | 1m50s  |       |
| postgres-backup1            | postgres-app | Completed | 1h49m  |       |
+-----------------------------+--------------+-----------+--------+-------+
```

Activate Windows
Go to Settings to activate Windows.

From the source cluster, switch context to the destination cluster.
Then, ensure that the AppVault is accessible from the destination cluster context and get the AppVault contents from the destination cluster.

```
[root@localhost DataProtection]# kubectl config use-context default/api-bm-cluster5-min-ocpv-sddc-netapp-com:6443/kube:admin
Switched to context "default/api-bm-cluster5-min-ocpv-sddc-netapp-com:6443/kube:admin".
[root@localhost DataProtection]# tp get appvault -n trident-protect
+-------------------+----------+-----------+------+-------+
|       NAME        | PROVIDER |   STATE   | AGE  | ERROR |
+-------------------+----------+-----------+------+-------+
| ontap-s3-appvault | OntapS3  | Available | 3d6h |       |
+-------------------+----------+-----------+------+-------+
[root@localhost DataProtection]# tp get appvaultcontent ontap-s3-appvault --show-resources backup --show-paths
+---------------+--------------+--------+----------------------------+--------------------------+-----------------------------------
----------------------------------------------------------------------------------+
|    CLUSTER    |     APP      |  TYPE  |            NAME            |        TIMESTAMP         |
                      PATH                                                        |
+---------------+--------------+--------+----------------------------+--------------------------+-----------------------------------
----------------------------------------------------------------------------------+
| ocp-cluster11 | bbox         | backup | bboxbackup1                | 2025-01-17 15:57:49 (UTC) | bbox_68b4f05f-d5f1-4304-8c69-ad67514
fe393/backups/bboxbackup1_3960c945-8ee0-42fe-945c-c57bb0a3af6f                                  |
| ocp-cluster11 | postgres-app | backup | backup1                    | 2025-02-18 13:31:50 (UTC) | postgres-app_4d798ed5-cfa8-49ff-a5b6
-c5e2d89aeb89/backups/backup1_28e1bd9a-9b04-4412-8b96-811f9b62e2e3                              |
| ocp-cluster11 | postgres-app | backup | daily-feac1-20250222153300 | 2025-02-22 15:34:44 (UTC) | postgres-app_4d798ed5-cfa8-49ff-a5b6
-c5e2d89aeb89/backups/daily-feac1-20250222153300_23d1386b-09f0-456f-aa49-a5865fd48abd  |
| ocp-cluster11 | postgres-app | backup | daily-feac1-20250223153300 | 2025-02-23 15:34:42 (UTC) | postgres-app_4d798ed5-cfa8-49ff-a5b6
-c5e2d89aeb89/backups/daily-feac1-20250223153300_c492a4d1-38a9-4472-9684-4705c12a206d  |
| ocp-cluster11 | postgres-app | backup | hourly-3f1ee-20250224033300 | 2025-02-24 03:34:44 (UTC) | postgres-app_4d798ed5-cfa8-49ff-a5b6
-c5e2d89aeb89/backups/hourly-3f1ee-20250224033300_3d09ab81-f6fe-47fa-a699-28006160cdbc |
| ocp-cluster11 | postgres-app | backup | hourly-3f1ee-20250224043300 | 2025-02-24 04:34:47 (UTC) | postgres-app_4d798ed5-cfa8-49ff-a5b6
-c5e2d89aeb89/backups/hourly-3f1ee-20250224043300_66805e4f-7631-48a6-98f7-d34bb8626031 |
| ocp-cluster11 | postgres-app | backup | postgres-backup-cluster1   | 2025-02-24 03:52:36 (UTC) | postgres-app_4d798ed5-cfa8-49ff-a5b6
-c5e2d89aeb89/backups/postgres-backup-cluster1_ec0ed3f3-5500-4e72-afa8-117a04a0b1c3    |
```

Use the backup path from the list and create a backuprestore CR object as shown in the command below.

```
# tp create backuprestore backup-restore-cluster2 --namespace-mapping
postgres:postgres --appvault ontap-s3-appvault --path postgres-
app_4d798ed5-cfa8-49ff-a5b6-c5e2d89aeb89/backups/postgres-backup-
cluster1_ec0ed3f3-5500-4e72-afa8-117a04a0b1c3 -n postgres
BackupRestore "backup-restore-cluster2" created.
```

```
[root@localhost DataProtection]# tp get backuprestore -n postgres
+--------------------------+--------------------+-----------+--------+-------+
|           NAME           |      APPVAULT      |   STATE   |  AGE   | ERROR |
+--------------------------+--------------------+-----------+--------+-------+
| backup-restore-cluster2 | ontap-s3-appvault | Completed | 12m41s |       |
+--------------------------+--------------------+-----------+--------+-------+
```

You can now see that the application pods and the pvcs are created in the destination cluster.

```
[root@localhost DataProtection]# oc get pods -n postgres
NAME                     READY   STATUS    RESTARTS   AGE
postgres-cd9d6ccb-2lvcq   1/1     Running   0          13m
[root@localhost DataProtection]# oc get pvc -n postgres
NAME                         STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   VOLUMEATTRIBUT
ESCLASS   AGE
data-postgres-postgresql-0   Bound    pvc-872a5182-601b-4848-b410-fef368337d07   8Gi        RWO            sc-zoneb-san   <unset>
          13m
postgres-pvc                 Bound    pvc-caf9fa71-76a8-4645-9bb5-2ed72e72948b   5Gi        RWO            sc-zoneb-san   <unset>
          13m
[root@localhost DataProtection]#
```

Activate Windows
Go to Settings to activate Windows.

Protect VMs in Red Hat OpenShift Virtualization using Trident Protect

Protect VMs in OpenShift Virtualization using snapshots and backups. This procedure includes creating an AppVault using ONTAP S3 object storage, configuring Trident Protect to capture VM data, including Kubernetes resource objects, persistent volumes, and internal images, and restoring the data when necessary.

Virtual machines in the OpenShift Virtualization environment are containerized applications that run in the worker nodes of your OpenShift Container platform. It is important to protect the VM metadata as well as the persistent disks of the VMs, so that when they are lost or corrupted, you can recover them.

The persistent disks of the OpenShift Virtualization VMs can be backed by ONTAP storage integrated to the OpenShift Cluster using Trident CSI. In this section we use Trident Protect to create snapshots and backups of VMs including their data volumes to ONTAP object storage.

We then restore from a snapshot or a backup when needed.

Trident Protect enables snapshots, backups, restores, and disaster recovery of applications and VMs on an OpenShift cluster. For OpenShift Virtualization VMs, data that can be protected with Trident Protect include Kubernetes resource objects associated with the VMs, persistent volumes, and internal images.

The following are the versions of the various components used for the examples in this section

- OpenShift Cluster 4.17
- OpenShift Virtualization installed via OpenShift Virtualization Operator provided by Red Hat
- Trident 25.02
- Trident Protect 25.02
- ONTAP 9.16

Create AppVault for object storage

▼ Create AppVault

Prior to creating the snapshots and backups for an application or a VM, object storage must be configured in Trident Protect to store the snapshots and backups. This is done using the bucket CR. Only administrators can create a bucket CR and configure it.
The bucket CR is known as AppVault in Trident Protect. AppVault objects are the declarative Kubernetes workflow representation of a storage bucket. An AppVault CR contains the configurations necessary for a bucket to be used in protection operations, such as backups, snapshots, restore operations, and SnapMirror replication.

In this example, we will show the use of ONTAP S3 as object storage.
Here is the workflow for creating AppVault CR for ONTAP S3:
1. Create an S3 object store server in the SVM in ONTAP Cluster.
2. Create a bucket in the object store server.
3. Create an S3 user in the SVM. Keep the access key and the secret key in a safe location.
4. In OpenShift, create a secret to store the ONTAP S3 credentials.
5. Create an AppVault object for ONTAP S3

Configure Trident Protect AppVault for ONTAP S3

```
# alias tp='tridentctl-protect'
```

```
# cat appvault-secret.yaml
apiVersion: v1
stringData:
  accessKeyID: "<access key of S3>"
  secretAccessKey: "<secret access key of S3>"
# you can also provide base 64 encoded values instead of string values
#data:
# base 64 encoded values
#  accessKeyID: < base 64 encoded access key>
#  secretAccessKey: <base 64 encoded secretAccess key>
kind: Secret
metadata:
  name: appvault-secret
  namespace: trident-protect
type: Opaque

# cat appvault.yaml
apiVersion: protect.trident.netapp.io/v1
kind: AppVault
metadata:
  name: ontap-s3-appvault
  namespace: trident-protect
spec:
  providerConfig:
    azure:
      accountName: ""
      bucketName: ""
      endpoint: ""
    gcp:
      bucketName: ""
      projectID: ""
    s3:
      bucketName: trident-protect
      endpoint: <lif for S3 access>
      secure: "false"
      skipCertValidation: "true"
  providerCredentials:
    accessKeyID:
      valueFromSecret:
        key: accessKeyID
        name: appvault-secret
    secretAccessKey:
      valueFromSecret:
        key: secretAccessKey
        name: appvault-secret
```

```
  providerType: OntapS3

# oc create -f appvault-secret.yaml -n trident-protect
# oc create -f appvault.yaml -n trident-protect
```

```
[root@localhost VM-DataProtection]# tp get appvault
+-------------------+----------+-----------+-------+---------+-------+
|       NAME        | PROVIDER |   STATE   | ERROR | MESSAGE |  AGE  |
+-------------------+----------+-----------+-------+---------+-------+
| ontap-s3-appvault | OntapS3  | Available |       |         | 8d17h |
+-------------------+----------+-----------+-------+---------+-------+
[root@localhost VM-DataProtection]#
```

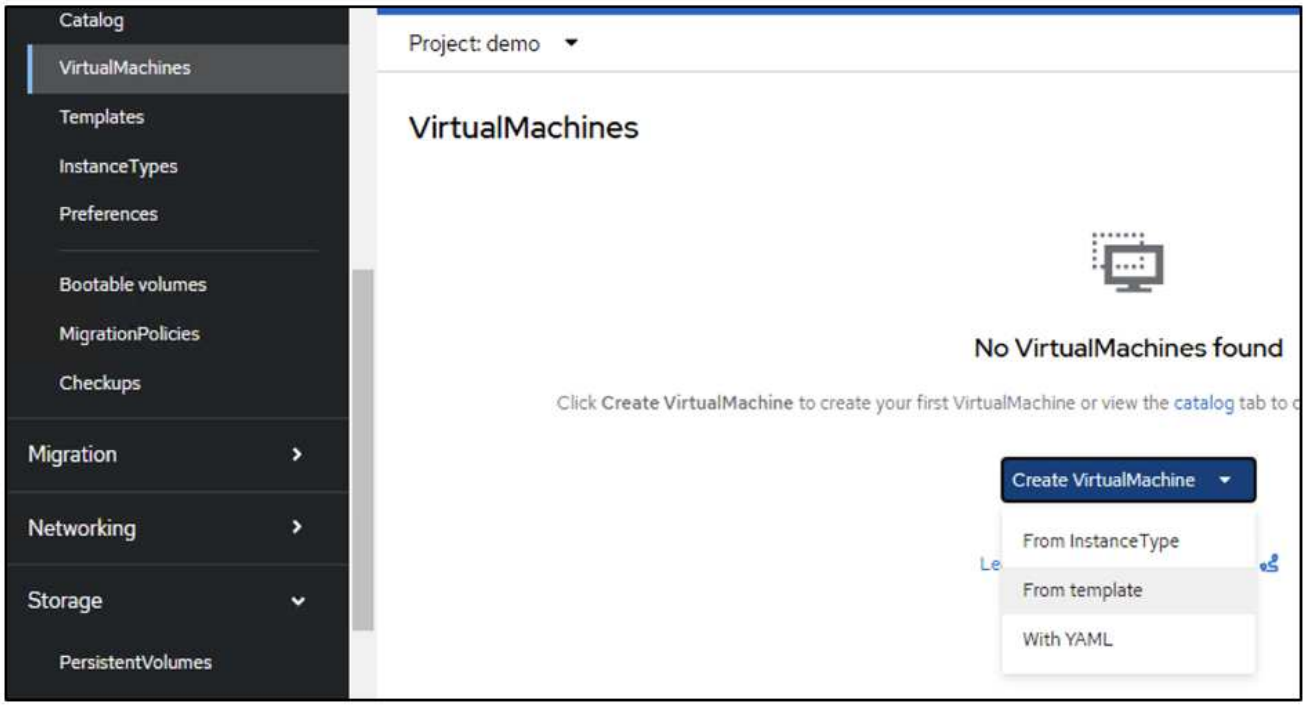
Create a VM in OpenShift Virtualization

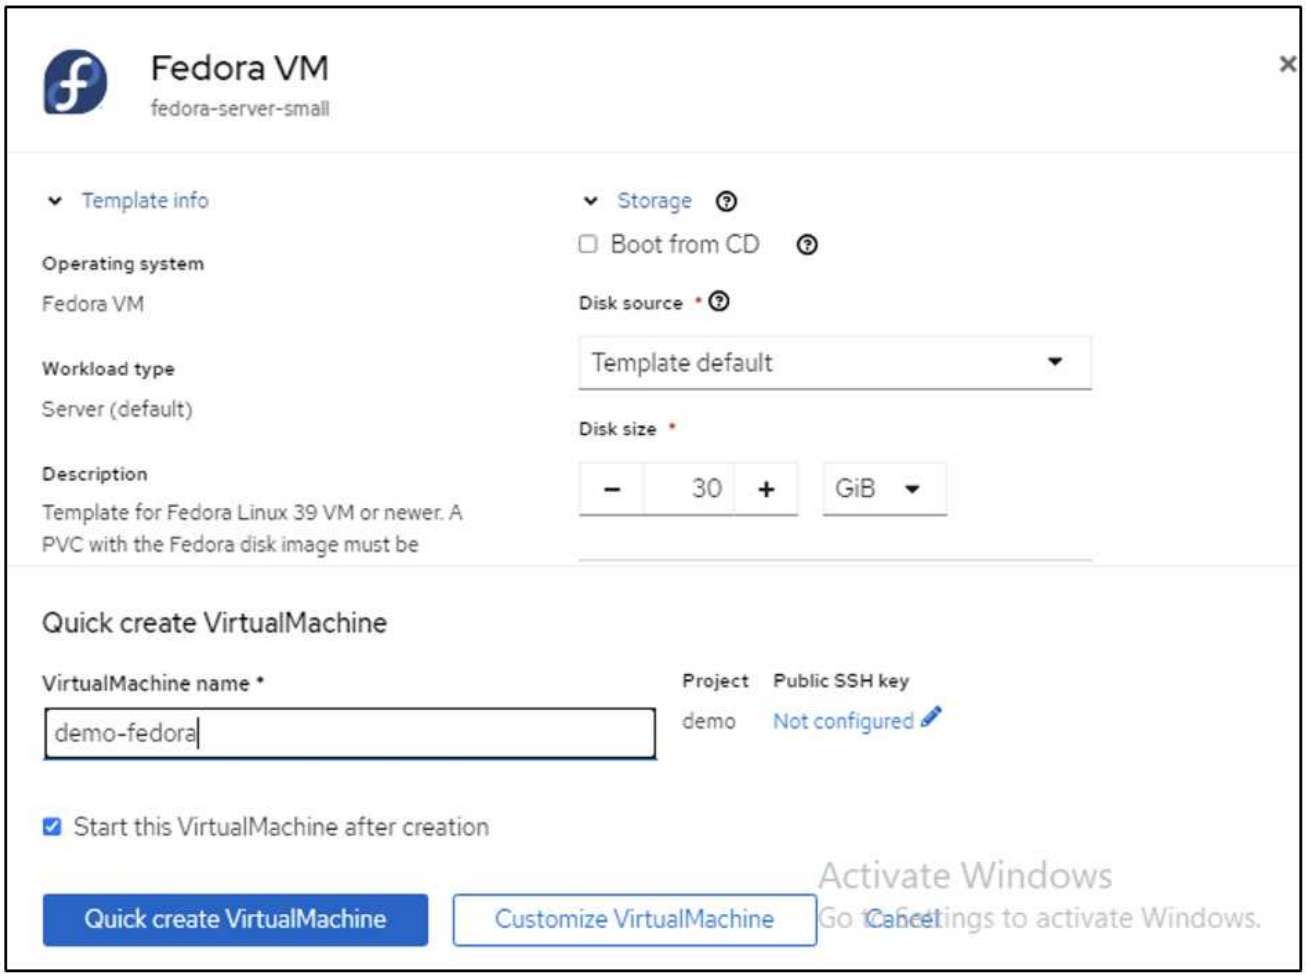
▼ **Create a VM in OpenShift Virtualization**

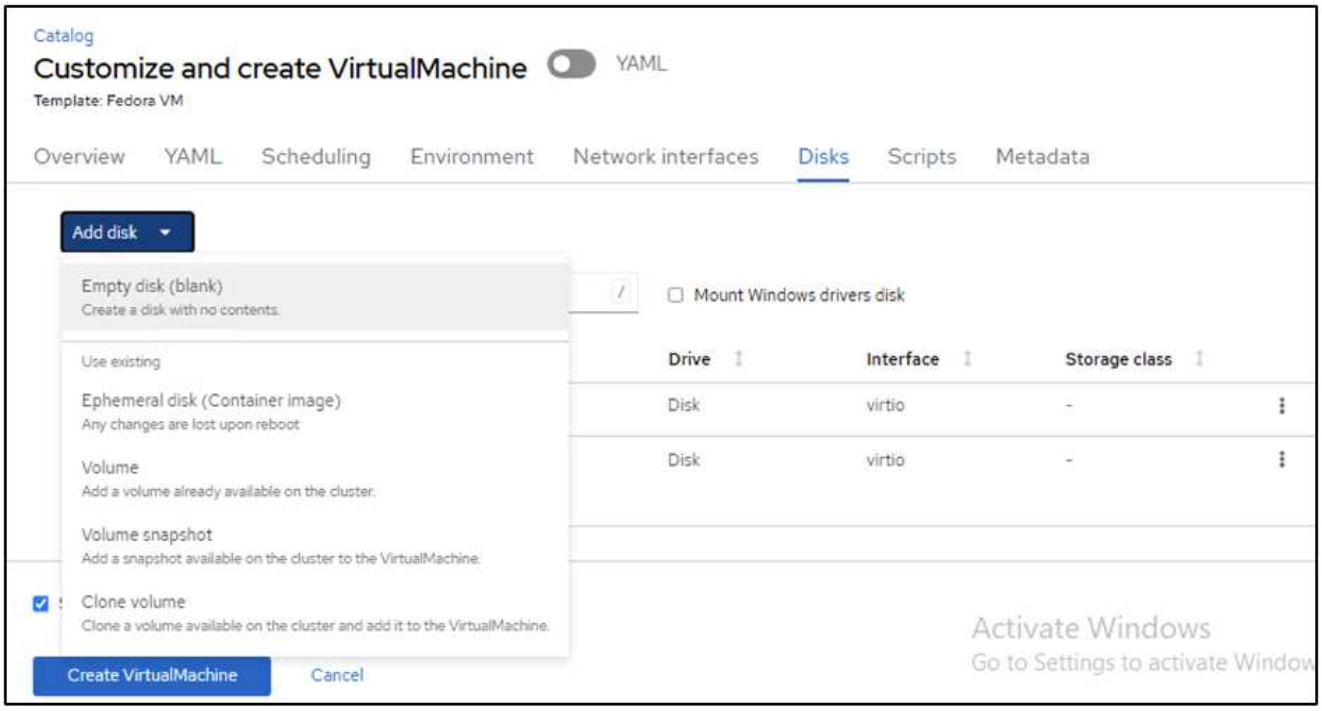
The following screen shots show the creation of the VM (demo-fedora in namespace demo) from the console using the template. The root disk chooses the default storage class automatically, so, verify that the default storage class is set appropriately. In this setup, the default storage class is **sc-zonea-san**. Ensure that when you create the additional disk, you choose the storage class sc-zonea-san and check the "**Apply optimized storage settings**" checkbox. This will set the Access modes to RWX and Volume Mode to Block.

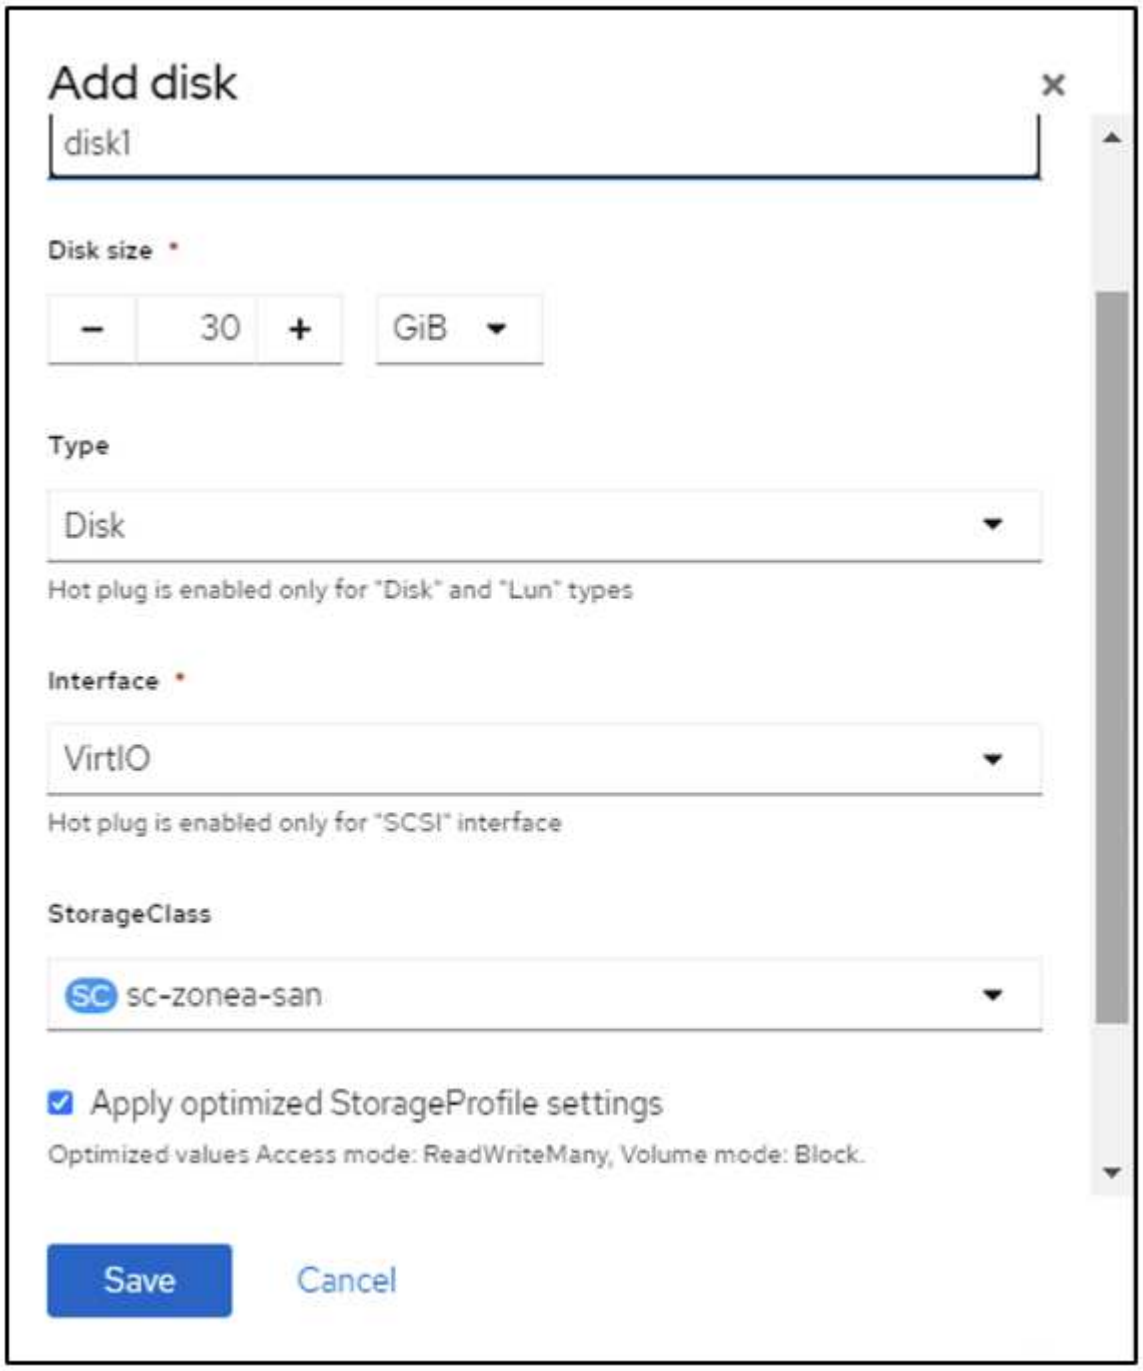
ⓘ Trident supports RWX access mode in Block Volume mode for SAN (iSCSI, NVMe/TCP and FC). (It is the default access mode for NAS). RWX Access mode is required if you need to do Live migration of the VMs at a later point.

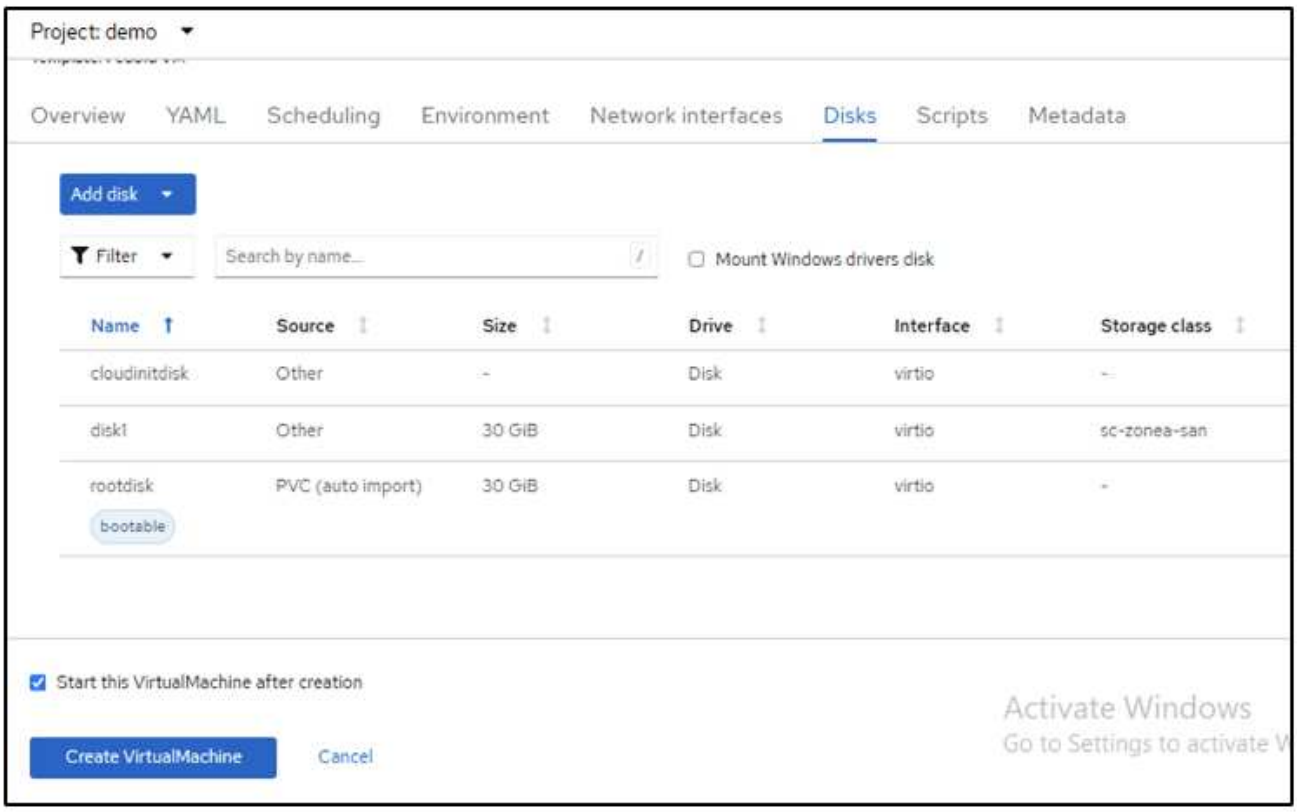



Catalog
VirtualMachines
Templates
InstanceTypes
Preferences
Bootable volumes
MigrationPolicies
Checkups
Migration
Networking
Storage
PersistentVolumes
Project: demo
VirtualMachines
No VirtualMachines found
Click Create VirtualMachine to create your first VirtualMachine or view the catalog tab to
Create VirtualMachine
From InstanceType
From template
With YAML


Fedora VM
fedora-server-small
Template info
Operating system
Fedora VM
Workload type
Server (default)
Description
Template for Fedora Linux 39 VM or newer. A PVC with the Fedora disk image must be
Storage
Boot from CD
Disk source *
Template default
Disk size *
30
GiB
Quick create VirtualMachine
VirtualMachine name *
demo-fedora
Project
demo
Public SSH key
Not configured
Start this VirtualMachine after creation
Activate Windows
Go to Settings to activate Windows.
Quick create VirtualMachine
Customize VirtualMachine
Cancel


Catalog
Customize and create VirtualMachine
YAML
Template: Fedora VM
Overview
YAML
Scheduling
Environment
Network interfaces
Disks
Scripts
Metadata
Add disk
Empty disk (blank)
Create a disk with no contents.
Use existing
Ephemeral disk (Container image)
Any changes are lost upon reboot
Volume
Add a volume already available on the cluster.
Volume snapshot
Add a snapshot available on the cluster to the VirtualMachine.
Clone volume
Clone a volume available on the cluster and add it to the VirtualMachine.
Mount Windows drivers disk
Drive
Interface
Storage class
Disk
virtio
-
Disk
virtio
-
Create VirtualMachine
Cancel
Activate Windows
Go to Settings to activate Window


Add disk
disk1
Disk size *
30
GiB
Type
Disk
Hot plug is enabled only for "Disk" and "Lun" types
Interface *
VirtIO
Hot plug is enabled only for "SCSI" interface
StorageClass
SC sc-zonea-san
Apply optimized StorageProfile settings
Optimized values Access mode: ReadWriteMany, Volume mode: Block.
Save
Cancel




```
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo
NAME                                   AGE     STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   2m54s   Running   True

NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-fedora-69cts   1/1     Running   0          110s

NAME                                                 STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
persistentvolumeclaim/demo-fedora                    Bound    pvc-406d9d57-4a5e-4123-93c1-13070b7b0140   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-311f2f81-1d25-4a9a-b0cb-836b4e702f04   30Gi       RWX            sc-zonea-san
[root@localhost VM-DataProtection]#
```

Create an app

▼ Create App

Create a Trident Protect app for the VM

In the example, the demo namespace has one VM and all resources of the namespace are included when creating the app.

```
# alias tp='tridentctl-protect'
# tp create app demo-vm --namespaces demo -n demo --dry-run > app.yaml

# cat app.yaml
apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  creationTimestamp: null
  name: demo-vm
```

```
  namespace: demo
spec:
  includedNamespaces:
  - namespace: demo
# oc create -f app.yaml -n demo
```

```
[root@localhost VM-DataProtection]# tp get app -n demo
+---------+------------+-------+-----+
|  NAME   | NAMESPACES | STATE | AGE |
+---------+------------+-------+-----+
| demo-vm | demo       | Ready | 45s |
+---------+------------+-------+-----+
[root@localhost VM-DataProtection]#
```

Protect the app by creating a backup

▼ **Create Backups**

Create an On-demand Backup

Create a backup for the app (demo-vm) created previously that includes all resources in the demo namespace. Provide the appvault name where the backups will be stored.

```
# tp create backup demo-vm-backup-on-demand --app demo-vm --appvault
ontap-s3-appvault -n demo
Backup "demo-vm-backup-on-demand" created.
```

```
[root@localhost VM-DataProtection]# tp get backup -n demo
+--------------------------+---------+----------------+-----------+-------+--------+
|           NAME           |   APP   | RECLAIM POLICY |   STATE   | ERROR |  AGE   |
+--------------------------+---------+----------------+-----------+-------+--------+
| demo-vm-backup-on-demand | demo-vm | Retain         | Completed |       | 12m53s |
+--------------------------+---------+----------------+-----------+-------+--------+
[root@localhost VM-DataProtection]#
```

Create Backups on a Schedule

Create a schedule for the backups specifying the granularity and the number of backups to retain.

```
# tp create schedule backup-schedule1 --app demo-vm --appvault ontap-s3-
appvault --granularity Hourly --minute 45 --backup-retention 1 -n demo
--dry-run>backup-schedule-demo-vm.yaml
schedule.protect.trident.netapp.io/backup-schedule1 created
```

```
#cat backup-schedule-demo-vm.yaml
apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  creationTimestamp: null
  name: backup-schedule1
  namespace: demo
spec:
  appVaultRef: ontap-s3-appvault
  applicationRef: demo-vm
  backupRetention: "1"
  dayOfMonth: ""
  dayOfWeek: ""
  enabled: true
  granularity: Hourly
  hour: ""
  minute: "45"
  recurrenceRule: ""
  snapshotRetention: "0"
status: {}
# oc create -f backup-schedule-demo-vm.yaml -n demo
```

```
[root@localhost VM-DataProtection]# tp get schedule -n demo
+------------------+---------+---------------+---------+-------+-------+-----+
|       NAME       |   APP   |   SCHEDULE    | ENABLED | STATE | ERROR | AGE |
+------------------+---------+---------------+---------+-------+-------+-----+
| backup-schedule1 | demo-vm | Hourly:min=45 | true    |       |       | 9s  |
+------------------+---------+---------------+---------+-------+-------+-----+
```

```
[root@localhost VM-DataProtection]# tp get backups -n demo
+-----------------------------+---------+----------------+-----------+-------+--------+
|            NAME             |   APP   | RECLAIM POLICY |   STATE   | ERROR |  AGE   |
+-----------------------------+---------+----------------+-----------+-------+--------+
| demo-vm-backup-on-demand    | demo-vm | Retain         | Completed |       | 44m4s  |
| hourly-4c094-20250312144500 | demo-vm | Retain         | Completed |       | 20m34s |
+-----------------------------+---------+----------------+-----------+-------+--------+
[root@localhost VM-DataProtection]#
```

Restore from Backup

▼ **Restore from Backups**

Restore the VM to the same namespace

In the example the backup demo-vm-backup-on-demand contains the backup with the demo-app for the fedora VM.

First, delete the VM and ensure that the PVCs, pod and the VM objects are deleted from the namespace

"demo"

```
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo
NAME                                  AGE   STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   59m   Running   True

NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-fedora-69cts   1/1     Running   0          58m

NAME                                                  STATUS   VOLUME                                     CAPACITY
CLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-fedora                     Bound    pvc-406d9d57-4a5e-4123-93c1-13070b7b0140   30Gi
a-san   <unset>                 59m
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-311f2f81-1d25-4a9a-b0cb-836b4e702f04   30Gi
a-san   <unset>                 59m
[root@localhost VM-DataProtection]#
[root@localhost VM-DataProtection]# oc delete vm demo-fedora -n demo
virtualmachine.kubevirt.io "demo-fedora" deleted
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo
No resources found in demo namespace.
[root@localhost VM-DataProtection]#
```

Now, create a backup-in-place restore object.

```
# tp create bir demo-fedora-restore --backup demo/demo-vm-backup-on-
demand -n demo --dry-run>vm-demo-bir.yaml

# cat vm-demo-bir.yaml
apiVersion: protect.trident.netapp.io/v1
kind: BackupInplaceRestore
metadata:
  annotations:
    protect.trident.netapp.io/max-parallel-restore-jobs: "25"
  creationTimestamp: null
  name: demo-fedora-restore
  namespace: demo
spec:
  appArchivePath: demo-vm_cc8adc7a-0c28-460b-a32f-
0a7b3d353e13/backups/demo-vm-backup-on-demand_f6af3513-9739-480e-88c7-
4cca45808a80
  appVaultRef: ontap-s3-appvault
  resourceFilter: {}
status:
  postRestoreExecHooksRunResults: null
  state: ""

# oc create -f vm-demo-bir.yaml -n demo
backupinplacerestore.protect.trident.netapp.io/demo-fedora-restore
created
```

```
[root@localhost VM-DataProtection]# tp get bir -n demo
+----------------------+--------------------+-----------+-------+--------+
|         NAME         |      APPVAULT      |   STATE   | ERROR |  AGE   |
+----------------------+--------------------+-----------+-------+--------+
| demo-fedora-restore | ontap-s3-appvault | Completed |       | 28m17s |
+----------------------+--------------------+-----------+-------+--------+
[root@localhost VM-DataProtection]#
```

Verify that the VM, pods and PVCs are restored

```
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo
NAME                                        AGE    STATUS    READY
virtualmachine.kubevirt.io/demo-fedora      116s   Running   True

NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-fedora-9kfxh  1/1     Running   0          116s

NAME                                                  STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
persistentvolumeclaim/demo-fedora                     Bound    pvc-6f69a62c-285c-4980-b0dd-6c85baccf346   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87 Bound    pvc-81d82d82-7aca-40fc-8f8f-6e99246e63f8   30Gi       RWX            sc-zonea-san
[root@localhost VM-DataProtection]#
```

Restore the VM to a different namespace

First create a new namespace to which you want to restore the app to, in this example demo2. Then create a backup restore object

```
# tp create br demo2-fedora-restore --backup demo/hourly-4c094-
20250312154500 --namespace-mapping demo:demo2 -n demo2 --dry-run>vm
-demo2-br.yaml

# cat vm-demo2-br.yaml
apiVersion: protect.trident.netapp.io/v1
kind: BackupRestore
metadata:
  annotations:
    protect.trident.netapp.io/max-parallel-restore-jobs: "25"
  creationTimestamp: null
  name: demo2-fedora-restore
  namespace: demo2
spec:
  appArchivePath: demo-vm_cc8adc7a-0c28-460b-a32f-
0a7b3d353e13/backups/hourly-4c094-20250312154500_aaa14543-a3fa-41f1-
a04c-44b1664d0f81
  appVaultRef: ontap-s3-appvault
  namespaceMapping:
  - destination: demo2
    source: demo
  resourceFilter: {}
status:
```

```
  conditions: null
  postRestoreExecHooksRunResults: null
  state: ""
# oc create -f vm-demo2-br.yaml -n demo2
```

```
[root@localhost VM-DataProtection]# tp get br -n demo2
+----------------------+--------------------+-----------+-------+--------+
|         NAME         |     APPVAULT       |   STATE   | ERROR |  AGE   |
+----------------------+--------------------+-----------+-------+--------+
| demo2-fedora-restore | ontap-s3-appvault | Completed |       | 38m52s |
+----------------------+--------------------+-----------+-------+--------+
```

Verify that the VM, pods and pvcs are created in the new namespace demo2.

```
[root@localhost VM-DataProtection]#
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo2
NAME                                  AGE    STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   5m8s   Running   True

NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-fedora-c7xc6   1/1     Running   0          5m7s

NAME                                                 STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
persistentvolumeclaim/demo-fedora                    Bound    pvc-4d278ae2-76cc-46f1-bbf8-071ae75e4a82   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-9b96d09c-7226-4ffc-829b-2cee88e7a117   30Gi       RWX            sc-zonea-san
[root@localhost VM-DataProtection]#
```

Protect the app using Snapshots

▼ Create Snapshots

Create an on-demand snapshot

Create a snapshot for the app and specify the appvault where it needs to be stored.

```
# tp create snapshot demo-vm-snapshot-ondemand --app demo-vm --appvault
ontap-s3-appvault -n demo --dry-run
# cat demo-vm-snapshot-on-demand.yaml
apiVersion: protect.trident.netapp.io/v1
kind: Snapshot
metadata:
  creationTimestamp: null
  name: demo-vm-snapshot-ondemand
  namespace: demo
spec:
  appVaultRef: ontap-s3-appvault
  applicationRef: demo-vm
  completionTimeout: 0s
  volumeSnapshotsCreatedTimeout: 0s
  volumeSnapshotsReadyToUseTimeout: 0s
status:
  conditions: null
```

```
  postSnapshotExecHooksRunResults: null
  preSnapshotExecHooksRunResults: null
  state: ""

# oc create -f demo-vm-snapshot-on-demand.yaml
snapshot.protect.trident.netapp.io/demo-vm-snapshot-ondemand created
```

```
[root@localhost VM-DataProtection]#
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo2
NAME                                   AGE    STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   5m8s   Running   True

NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-fedora-c7xc6   1/1     Running   0          5m7s

NAME                                                STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
persistentvolumeclaim/demo-fedora                   Bound    pvc-4d278ae2-76cc-46f1-bbf8-071ae75e4a82   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-9b96d09c-7226-4ffc-829b-2cee88e7a117   30Gi       RWX            sc-zonea-san
[root@localhost VM-DataProtection]#
```

Create a schedule for snapshots

Create schedule for the snapshots. Specify the granularity and the number of snapshots to be retained.

```
# tp create Schedule snapshot-schedule1 --app demo-vm --appvault ontap-
s3-appvault --granularity Hourly --minute 50 --snapshot-retention 1 -n
demo --dry-run>snapshot-schedule-demo-vm.yaml

# cat snapshot-schedule-demo-vm.yaml
apiVersion: protect.trident.netapp.io/v1
kind: Schedule
metadata:
  creationTimestamp: null
  name: snapshot-schedule1
  namespace: demo
spec:
  appVaultRef: ontap-s3-appvault
  applicationRef: demo-vm
  backupRetention: "0"
  dayOfMonth: ""
  dayOfWeek: ""
  enabled: true
  granularity: Hourly
  hour: ""
  minute: "50"
  recurrenceRule: ""
  snapshotRetention: "1"
status: {}

# oc create -f snapshot-schedule-demo-vm.yaml
```

```
schedule.protect.trident.netapp.io/snapshot-schedule1 created
```

```
[root@localhost VM-DataProtection]# tp get schedule -n demo
+--------------------+---------+---------------+---------+-------+-------+-------+
|        NAME        |   APP   |   SCHEDULE    | ENABLED | STATE | ERROR |  AGE  |
+--------------------+---------+---------------+---------+-------+-------+-------+
| backup-schedule1   | demo-vm | Hourly:min=45 | true    |       |       | 5d23h |
| snapshot-schedule1 | demo-vm | Hourly:min=50 | true    |       |       | 12s   |
+--------------------+---------+---------------+---------+-------+-------+-------+
[root@localhost VM-DataProtection]#
```

```
[root@localhost VM-DataProtection]# tp get snapshots -n demo
+-----------------------------------------+---------+----------------+-----------+-------+--------+
|                  NAME                   |   APP   | RECLAIM POLICY |   STATE   | ERROR |  AGE   |
+-----------------------------------------+---------+----------------+-----------+-------+--------+
| backup-39b67e1c-f875-4045-93df-78634bae9dfb | demo-vm | Delete         | Completed |       | 6m29s  |
| demo-vm-snapshot-ondemand               | demo-vm | Delete         | Completed |       | 21m30s |
| hourly-51839-20250318135000             | demo-vm | Delete         | Completed |       | 1m29s  |
+-----------------------------------------+---------+----------------+-----------+-------+--------+
[root@localhost VM-DataProtection]#
```

Restore from Snapshot

▼ **Restore from Snapshot**

Restore the VM from the snapshot to the same namespace

Delete the VM demo-fedora from the demo2 namespace.

```
[root@localhost RedHat]# oc get vm,pvc -n demo
NAME                                  AGE   STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   28h   Running   True

NAME                                                    STATUS   VOLUME
ORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-fedora                       Bound    pvc-e0d5f79d-dff9-450d-be0e-90ab6880b7af
-zonea-san   <unset>                 28h
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-e6f7890a-70c7-4538-9035-5e2e9379511d
-zonea-san   <unset>                 28h
[root@localhost RedHat]# oc delete virtualmachine.kubevirt.io/demo-fedora -n demo
virtualmachine.kubevirt.io "demo-fedora" deleted
[root@localhost RedHat]#
```

Create a snapshot-in-place-restore object from the snapshot of the VM.

```
# tp create sir demo-fedora-restore-from-snapshot --snapshot demo/demo-
vm-snapshot-ondemand -n demo --dry-run>vm-demo-sir.yaml

# cat vm-demo-sir.yaml
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotInplaceRestore
metadata:
  creationTimestamp: null
  name: demo-fedora-restore-from-snapshot
  namespace: demo
```

```
spec:
  appArchivePath: demo-vm_cc8adc7a-0c28-460b-a32f-
0a7b3d353e13/snapshots/20250318132959_demo-vm-snapshot-
ondemand_e3025972-30c0-4940-828a-47c276d7b034
  appVaultRef: ontap-s3-appvault
  resourceFilter: {}
status:
  conditions: null
  postRestoreExecHooksRunResults: null
  state: ""

# oc create -f vm-demo-sir.yaml
snapshotinplacerestore.protect.trident.netapp.io/demo-fedora-restore-
from-snapshot created
```

```
[root@localhost VM-DataProtection]# tp get sir -n demo
+-----------------------------------+-------------------+-----------+-------+--------+
|               NAME                |     APPVAULT      |   STATE   | ERROR |  AGE   |
+-----------------------------------+-------------------+-----------+-------+--------+
| demo-fedora-restore-from-snapshot | ontap-s3-appvault | Completed |       | 58m17s |
+-----------------------------------+-------------------+-----------+-------+--------+
[root@localhost VM-DataProtection]#
```

Verify that the VM and its PVCs are created in the demo namespace.

```
[root@localhost RedHat]# oc get vm,pvc -n demo
NAME                                  AGE     STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   5m17s   Running   True

NAME                                                STATUS   VOLUME
persistentvolumeclaim/demo-fedora                   Bound    pvc-e2f418b0-1b97-40fc-9cb8-943b370d85bc
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-db085154-079f-45ad-9e62-9656e913d01c
[root@localhost RedHat]#
[root@localhost RedHat]#
```

Restore the VM from the snapshot to a different namespace

Delete the VM in the demo2 namespace previously restored from the backup.

```
[root@localhost VM-DataProtection]# oc get vm,pods,pvc -n demo2
NAME                                  AGE     STATUS    READY
virtualmachine.kubevirt.io/demo-centos   3m12s   Running   True
virtualmachine.kubevirt.io/demo-fedora   3m11s   Running   True

NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-centos-w77rr   1/1     Running   0          3m11s
pod/virt-launcher-demo-fedora-wwdtc   1/1     Running   0          3m11s

NAME                                                                 STATUS   VOLUME                                     CAPACITY   ACCESS MODES
  STORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-centos                                    Bound    pvc-7aafd5e2-d4cf-4af6-a259-c68e016ec6cd   30Gi       RWX
  sc-zonea-san   <unset>                 3m26s
persistentvolumeclaim/demo-fedora                                    Bound    pvc-c1f8145a-56ea-42c5-abb4-6457b6853e1c   30Gi       RWX
  sc-zonea-san   <unset>                 3m25s
persistentvolumeclaim/dv-demo-centos-lavender-tortoise-34   Bound    pvc-e82b6d98-6762-4473-9bb8-7e98efcb7987   30Gi       RWX
  sc-zonea-san   <unset>                 3m20s
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87       Bound    pvc-ee13cfbc-2921-4129-a214-a0f6b10335cd   30Gi       RWX
  sc-zonea-san   <unset>                 3m16s
[root@localhost VM-DataProtection]# oc delete virtualmachine.kubevirt.io/demo-fedora -n demo2
virtualmachine.kubevirt.io "demo-fedora" deleted
```

Create the snapshot restore object from the snapshot and provide the namespace mapping.

```
# tp create sr demo2-fedora-restore-from-snapshot --snapshot demo/demo-
vm-snapshot-ondemand --namespace-mapping demo:demo2 -n demo2 --dry
-run>vm-demo2-sr.yaml

# cat vm-demo2-sr.yaml
apiVersion: protect.trident.netapp.io/v1
kind: SnapshotRestore
metadata:
  creationTimestamp: null
  name: demo2-fedora-restore-from-snapshot
  namespace: demo2
spec:
  appArchivePath: demo-vm_cc8adc7a-0c28-460b-a32f-
0a7b3d353e13/snapshots/20250318132959_demo-vm-snapshot-
ondemand_e3025972-30c0-4940-828a-47c276d7b034
  appVaultRef: ontap-s3-appvault
  namespaceMapping:
  - destination: demo2
    source: demo
  resourceFilter: {}
status:
  postRestoreExecHooksRunResults: null
  state: ""

# oc create -f vm-demo2-sr.yaml
snapshotrestore.protect.trident.netapp.io/demo2-fedora-restore-from-
snapshot created
```

```
[root@localhost VM-DataProtection]# tp get sr -n demo2
+-----------------------------------+--------------------+-----------+-------+--------+
|                NAME               |     APPVAULT       |   STATE   | ERROR |  AGE   |
+-----------------------------------+--------------------+-----------+-------+--------+
| demo2-fedora-restore-from-snapshot | ontap-s3-appvault | Completed |       | 15m22s |
+-----------------------------------+--------------------+-----------+-------+--------+
```

Verify that the VM and its PVCs are restored in the new namespace demo2.

```
[root@localhost RedHat]# oc get vm,pvc -n demo2
NAME                                      AGE   STATUS    READY
virtualmachine.kubevirt.io/demo-fedora    29h   Running   True

NAME                                                    STATUS   VOLUME
ORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-fedora                       Bound    pvc-35dcd9b2-4fca-486c-af9e-596bc5b4dc15
-zonea-san   <unset>                 29h
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-575a3111-382f-4933-a778-0089fa1ea2af
-zonea-san   <unset>                 29h
[root@localhost RedHat]#
```

Restore a Specific VM

▼ Selecting specific VMs in a namespace to create snapshots/backups and restore

In the previous example, we had a single VM within a namespace. By including the entire namespace in the backup, all resources associated with that VM were captured. In the following example, we add another VM to the same namespace and create an app just for this new VM by using a label selector.

Create a new VM (demo-centos vm) in the demo namespace

```
[root@localhost VM-DataProtection]# oc get vm,pod,pvc -n demo
NAME                                  AGE     STATUS    READY
virtualmachine.kubevirt.io/demo-centos   2m47s   Running   True
virtualmachine.kubevirt.io/demo-fedora   81m     Running   True

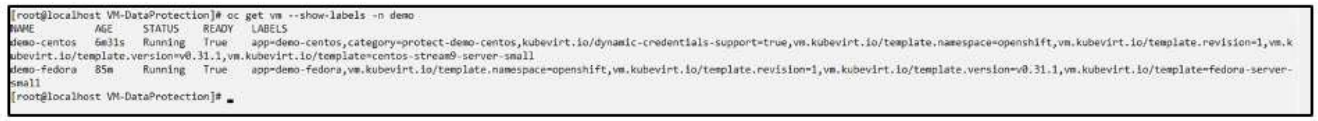
NAME                                 READY   STATUS    RESTARTS   AGE
pod/virt-launcher-demo-centos-2nq6g   1/1     Running   0          113s
pod/virt-launcher-demo-fedora-9kfxh   1/1     Running   0          81m

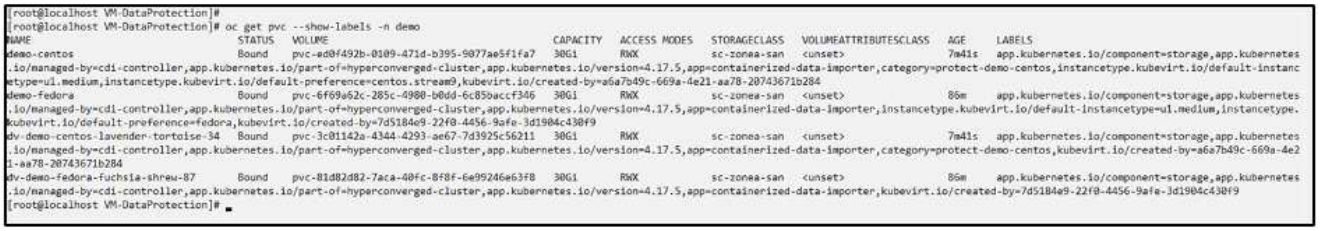
NAME                                                  STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
persistentvolumeclaim/demo-centos                     Bound    pvc-ed0f492b-0109-471d-b395-9077ae5f1fa7   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/demo-fedora                     Bound    pvc-6f69a62c-285c-4980-b0dd-6c85baccf346   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/dv-demo-centos-lavender-tortoise-34   Bound    pvc-3c01142a-4344-4293-ae67-7d3925c56211   30Gi       RWX            sc-zonea-san
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-81d82d82-7aca-40fc-8f8f-6e99246e63f8   30Gi       RWX            sc-zonea-san
[root@localhost VM-DataProtection]#
```

Label the demo-centos vm and its associated resources

```
[root@localhost VM-DataProtection]# oc label vm demo-centos category=protect-demo-centos -n demo
virtualmachine.kubevirt.io/demo-centos labeled
[root@localhost VM-DataProtection]# oc label pvc demo-centos category=protect-demo-centos -n demo
persistentvolumeclaim/demo-centos labeled
[root@localhost VM-DataProtection]# oc label pvc dv-demo-centos-lavender-tortoise-34 category=protect-demo-centos -n demo
persistentvolumeclaim/dv-demo-centos-lavender-tortoise-34 labeled
[root@localhost VM-DataProtection]#
```

Verify that the demo-centos vm and pvcs got the labels





Create an app for only a specific VM (demo-centos) using the label selector

```
# tp create app demo-centos-app --namespaces 'demo(category=protect-demo-centos)' -n demo --dry-run>demo-centos-app.yaml

# cat demo-centos-app.yaml

apiVersion: protect.trident.netapp.io/v1
kind: Application
metadata:
  creationTimestamp: null
  name: demo-centos-app
  namespace: demo
spec:
  includedNamespaces:
  - labelSelector:
      matchLabels:
        category: protect-demo-centos
    namespace: demo
status:
  conditions: null

# oc create -f demo-centos-app.yaml -n demo
application.protect.trident.netapp.io/demo-centos-app created
```

```
[root@localhost VM-DataProtection]# tp get app -n demo
+------------------+------------+-------+------+
|      NAME        | NAMESPACES | STATE | AGE  |
+------------------+------------+-------+------+
| demo-centos-app  | demo       | Ready | 56s  |
| demo-vm          | demo       | Ready | 4h6m |
+------------------+------------+-------+------+
```

The method of creating backups and snapshots on-demand and on a schedule is the same as shown previously.
Since the trident-protect app that is being used to create the snapshots or backups only contains the specific VM from the namespace, restoring from them only restores a specific VM.
A sample backup/restore operation is shown as an example below.

Create a backup of a specific VM in a namespace by using its corresponding app

In the previous steps, an app was created using label selectors to include only the centos vm in the demo namespace. Create a backup (on-demand backup, in this example) for this app.

```
# tp create backup demo-centos-backup-on-demand --app demo-centos-app
--appvault ontap-s3-appvault -n demo
Backup "demo-centos-backup-on-demand" created.
```

```
+------------------------------+-----------------+----------------+-----------+-------+---------+
|             NAME             |       APP       | RECLAIM POLICY |   STATE   | ERROR |   AGE   |
+------------------------------+-----------------+----------------+-----------+-------+---------+
| demo-centos-backup-on-demand | demo-centos-app | Retain         | Completed |       | 13m22s  |
| demo-vm-backup-on-demand     | demo-vm         | Retain         | Completed |       | 4h19m   |
| hourly-4c094-20250312174500  | demo-vm         | Retain         | Completed |       | 56m17s  |
+------------------------------+-----------------+----------------+-----------+-------+---------+
```

Restore a specific VM to the same namespace
The backup of a specific VM (centos) was created using the corresponding app.
If a backup-in-place-restore or a backup-restore is created from this, only this specific VM is restored.
Delete the Centos VM.

```
[root@localhost RedHat]# oc get vm,pvc -n demo
NAME                                    AGE     STATUS    READY
virtualmachine.kubevirt.io/demo-centos   4m27s   Running   True
virtualmachine.kubevirt.io/demo-fedora   4m27s   Running   True

NAME                                                                  STATUS   VOLUME
  STORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-centos                                     Bound    pvc-e8faeaf8-fc0c-4d92-96de-c83a335a7a17
  sc-zonea-san   <unset>                 4m33s
persistentvolumeclaim/demo-fedora                                     Bound    pvc-e2f418b0-1b97-40fc-9cb8-943b370d85bc
  sc-zonea-san   <unset>                 4m33s
persistentvolumeclaim/dv-demo-centos-lavender-tortoise-34             Bound    pvc-66eb7996-1420-4513-a67c-2824f08534da
  sc-zonea-san   <unset>                 4m33s
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87                 Bound    pvc-db085154-079f-45ad-9e62-9656e913d01c
  sc-zonea-san   <unset>                 4m32s
```

```
[root@localhost RedHat]# oc delete virtualmachine.kubevirt.io/demo-centos -n demo
virtualmachine.kubevirt.io "demo-centos" deleted
[root@localhost RedHat]# oc get vm,pvc -n demo
NAME                                  AGE     STATUS    READY
virtualmachine.kubevirt.io/demo-fedora   5m17s   Running   True

NAME                                                  STATUS   VOLUME
persistentvolumeclaim/demo-fedora                     Bound    pvc-e2f418b0-1b97-40fc-9cb8-943b370d85bc
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-db085154-079f-45ad-9e62-9656e913d01c
[root@localhost RedHat]#
```

Create a backup in-place restore from demo-centos-backup-on-demand and verify that the centos VM has been recreated.

```
#tp create bir demo-centos-restore --backup demo/demo-centos-backup-on-demand -n demo
BackupInplaceRestore "demo-centos-restore" created.
```

```
[root@localhost RedHat]# tp get bir -n demo
+----------------------+--------------------+-----------+-------+-------+
|         NAME         |      APPVAULT      |   STATE   | ERROR |  AGE  |
+----------------------+--------------------+-----------+-------+-------+
| demo-centos-restore  | ontap-s3-appvault  | Completed |       | 57m9s |
| demo-fedora-restore  | ontap-s3-appvault  | Completed |       | 7d5h  |
+----------------------+--------------------+-----------+-------+-------+
```

```
[root@localhost RedHat]# oc get vm,pvc -n demo
NAME                                  AGE   STATUS    READY
virtualmachine.kubevirt.io/demo-centos   29m   Running   True
virtualmachine.kubevirt.io/demo-fedora   85m   Running   True

NAME                                                  STATUS   VOLUME                                     C
  STORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-centos                     Bound    pvc-82954bf7-4a7e-4e0c-9a04-4fa152e1b0ef   1
  sc-zonea-san   <unset>                 29m
persistentvolumeclaim/demo-fedora                     Bound    pvc-e2f418b0-1b97-40fc-9cb8-943b370d85bc   1
  sc-zonea-san   <unset>                 85m
persistentvolumeclaim/dv-demo-centos-lavender-tortoise-34   Bound    pvc-2a8d4eb5-ed6d-4408-b85d-e218e9a5d4b0   1
  sc-zonea-san   <unset>                 29m
persistentvolumeclaim/dv-demo-fedora-fuchsia-shrew-87   Bound    pvc-db085154-079f-45ad-9e62-9656e913d01c   1
  sc-zonea-san   <unset>                 85m
[root@localhost RedHat]#
```

Restore a specific VM to a different namespace

Create a backup restore to a different namespace (demo3) from demo-centos-backup-on-demand and verify that the centos VM has been recreated.

```
# tp create br demo2-centos-restore --backup demo/demo-centos-backup-on-demand --namespace-mapping demo:demo3 -n demo3
BackupRestore "demo2-centos-restore" created.
```

```
[root@localhost RedHat]#
[root@localhost RedHat]# tp get br -n demo3
+----------------------+--------------------+-----------+-------+--------+
|         NAME         |     APPVAULT       |   STATE   | ERROR |  AGE   |
+----------------------+--------------------+-----------+-------+--------+
| demo2-centos-restore | ontap-s3-appvault | Completed |       | 52m57s |
+----------------------+--------------------+-----------+-------+--------+
[root@localhost RedHat]#
```

```
[root@localhost RedHat]#
[root@localhost RedHat]# oc get vm,pvc -n demo3
NAME                                    AGE   STATUS    READY
virtualmachine.kubevirt.io/demo-centos   19m   Running   True

NAME                                                          STATUS   VOLUME
  STORAGECLASS   VOLUMEATTRIBUTESCLASS   AGE
persistentvolumeclaim/demo-centos                             Bound    pvc-0a14e38f-07de-4e09-8f88-14a9a8bb45c2
  sc-zonea-san   <unset>                 19m
persistentvolumeclaim/dv-demo-centos-lavender-tortoise-34   Bound    pvc-d4f9cf2f-264c-4d02-94bf-0db28b427acc
  sc-zonea-san   <unset>                 19m
[root@localhost RedHat]#
```

Video Demonstration

The following video shows a demonstration of protecting a VM using Snapshots

Protecting a VM

Protect container apps using third-party tools

Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP)

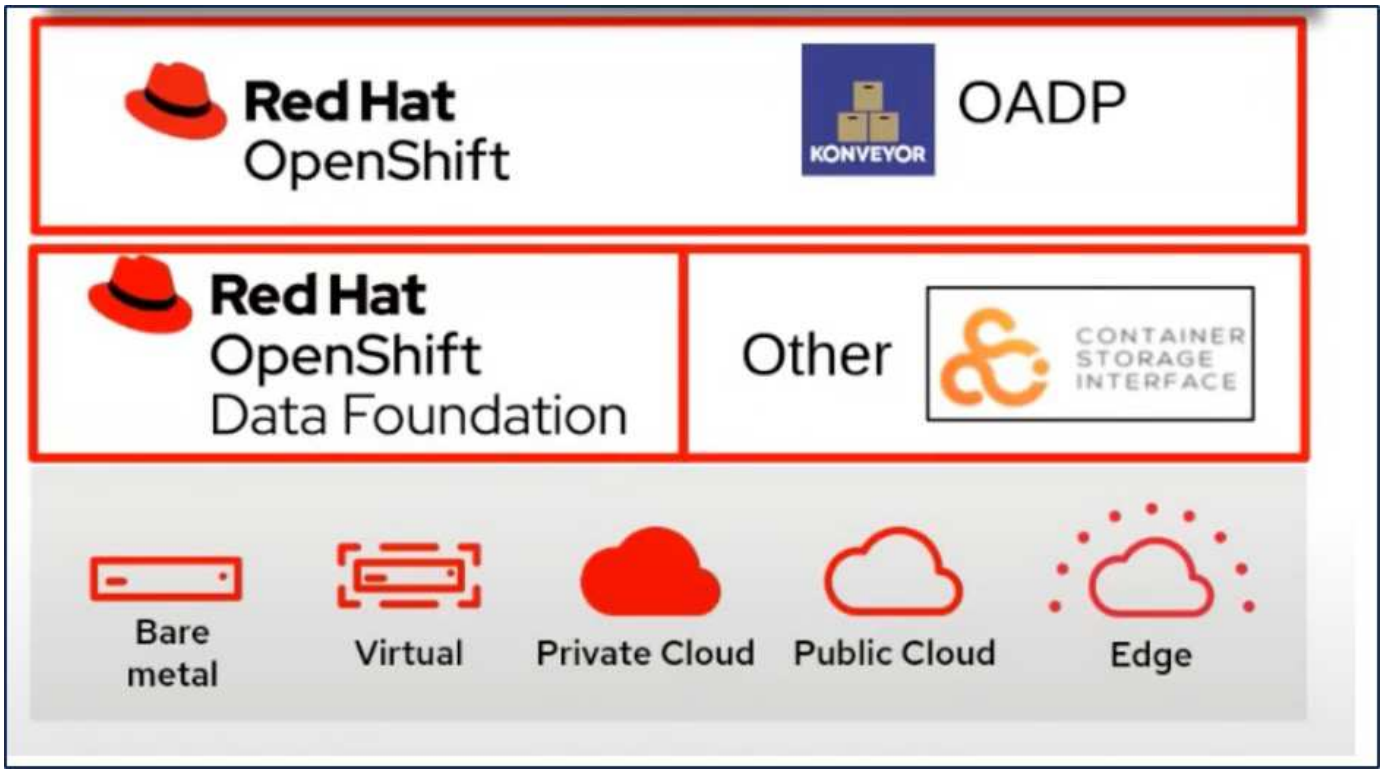
This section of the reference document provides details for creating backups of Container Apps using the OpenShift API for Data Protection (OADP) with Velero on NetApp ONTAP S3 or NetApp StorageGRID S3. The backups of namespace scoped resources including Persistent Volumes(PVs) of the app are created using CSI Trident Snapshots.

The persistent storage for container apps can be backed by ONTAP storage integrated to the OpenShift Cluster using Trident CSI. In this section we use OpenShift API for Data Protection (OADP) to perform backup of apps including its data volumes to

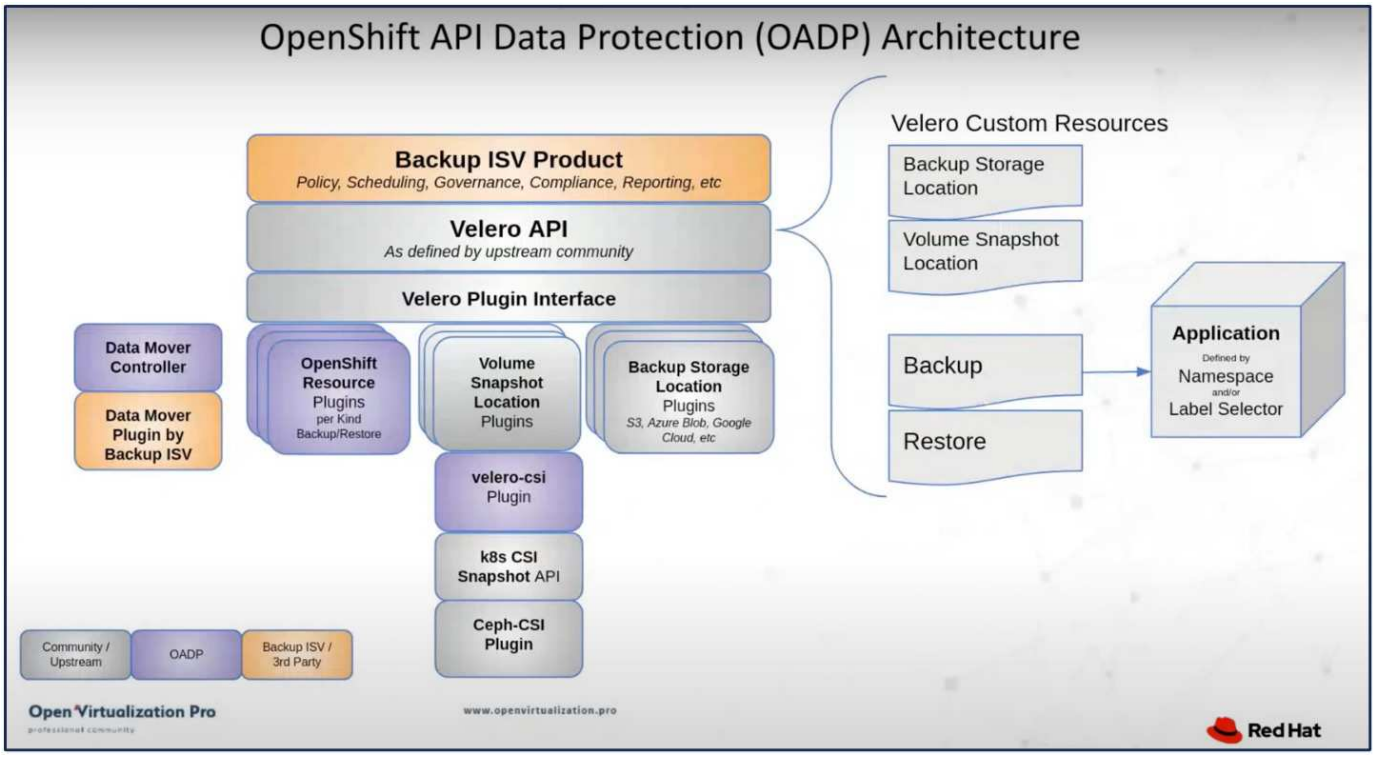
- ONTAP Object Storage
- StorageGrid

We then restore from the backup when needed. Please note that the app can be restored only to the cluster from where the backup was created.

OADP enables backup, restore, and disaster recovery of applications on an OpenShift cluster. Data that can be protected with OADP include Kubernetes resource objects, persistent volumes, and internal images.



Red Hat OpenShift has leveraged the solutions developed by the OpenSource communities for data protection. Velero is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes. To use Velero easily, OpenShift has developed the OADP operator and the Velero plugin to integrate with the CSI storage drivers. The core of the OADP APIs that are exposed are based on the Velero APIs. After installing the OADP operator and configuring it, the backup/restore operations that can be performed are based on the operations exposed by the Velero API.



OADP 1.3 is available from the operator hub of OpenShift cluster 4.12 and later. It has a built-in Data Mover

that can move CSI volume snapshots to a remote object store. This provides portability and durability by moving snapshots to an object storage location during backup. The snapshots are then available for restoration after disasters.

The following are the versions of the various components used for the examples in this section

- OpenShift Cluster 4.14
- OADP Operator 1.13 provided by Red Hat
- Velero CLI 1.13 for Linux
- Trident 24.02
- ONTAP 9.12
- postgresql installed using helm.

Trident CSI
OpenShift API for Data Protection (OADP)
Velero

Installation of OpenShift API for Data Protection (OADP) Operator

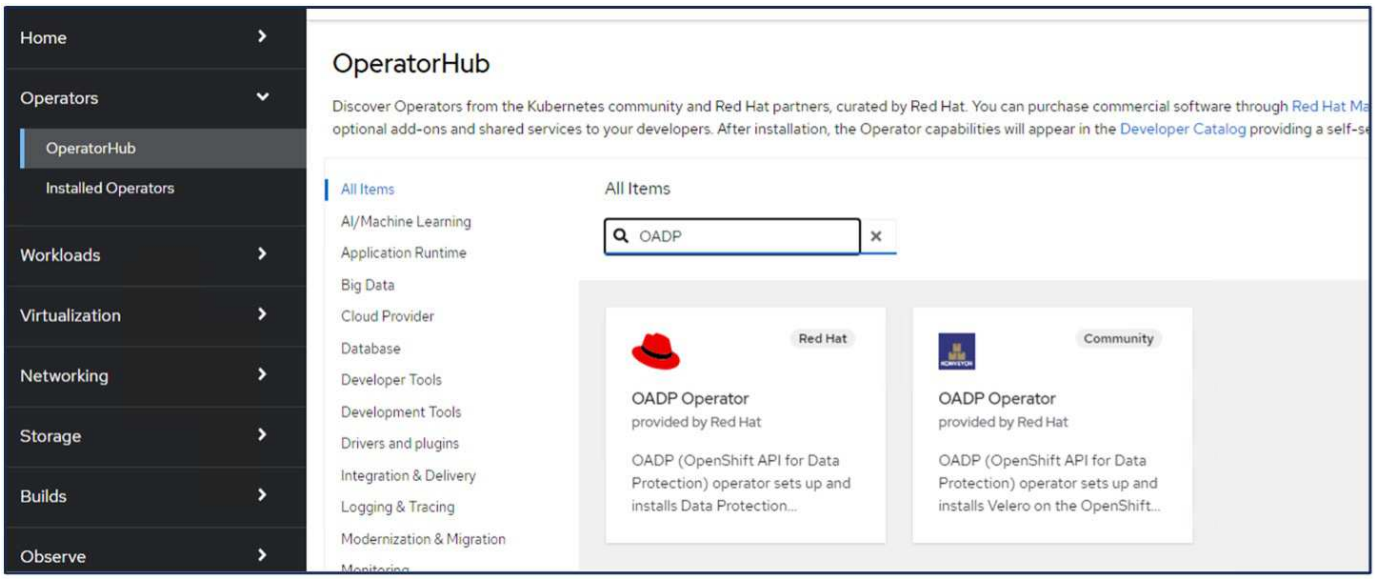
This section outlines the installation of OpenShift API for Data Protection (OADP) Operator.

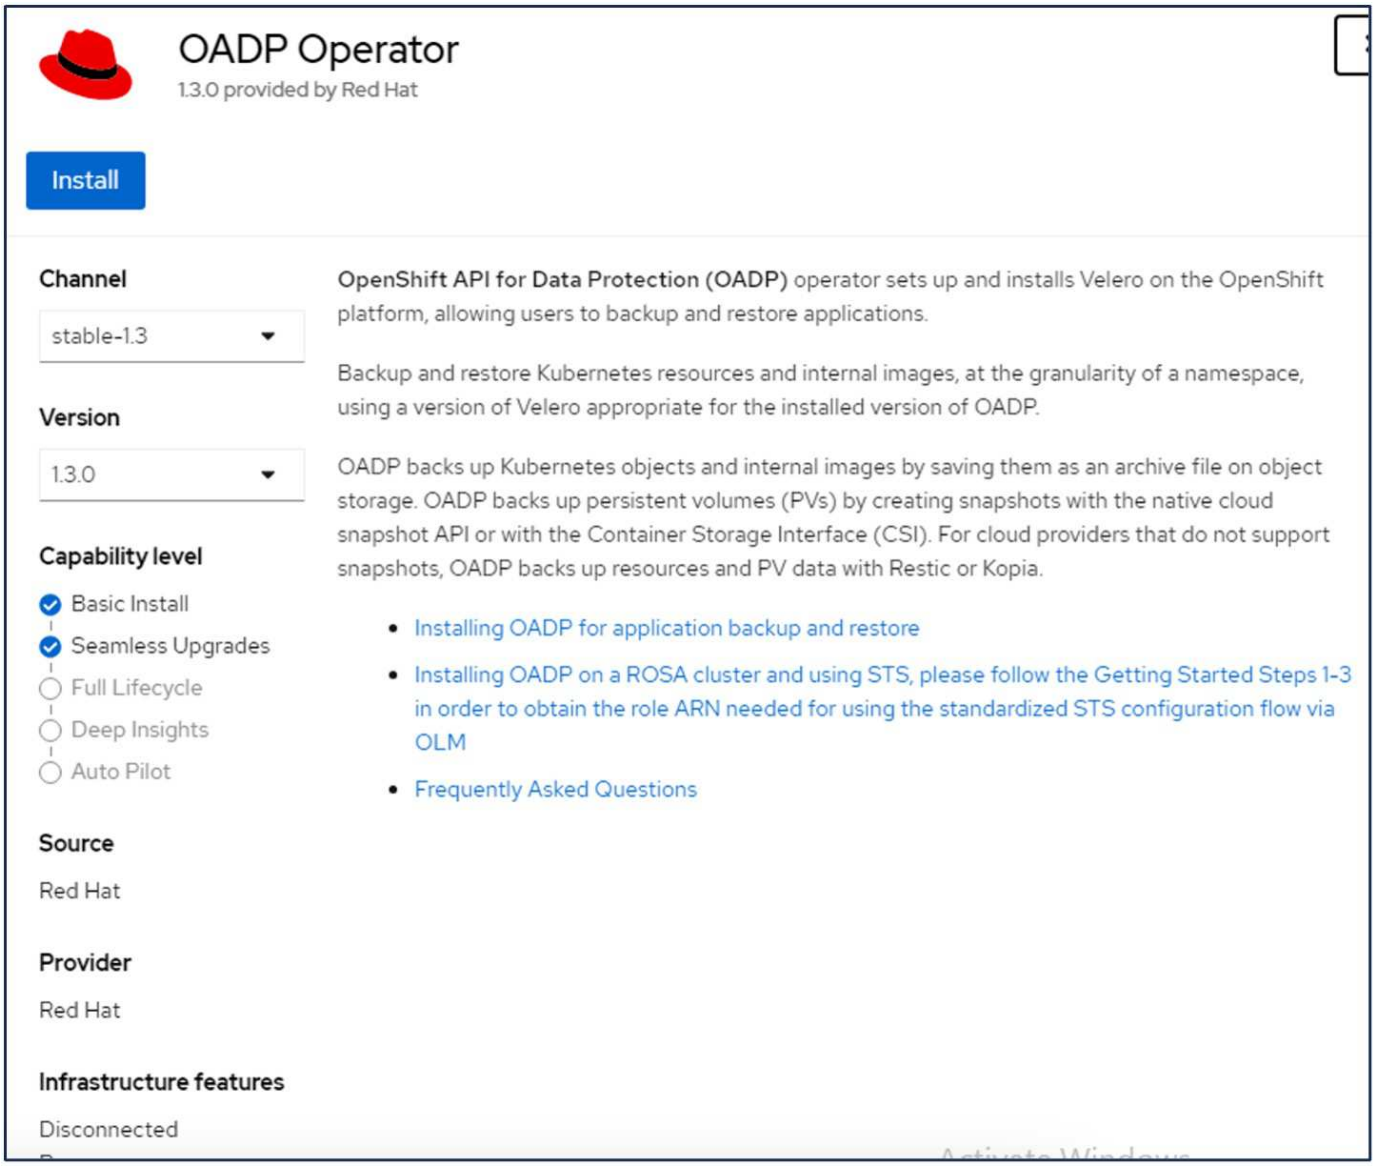
Prerequisites

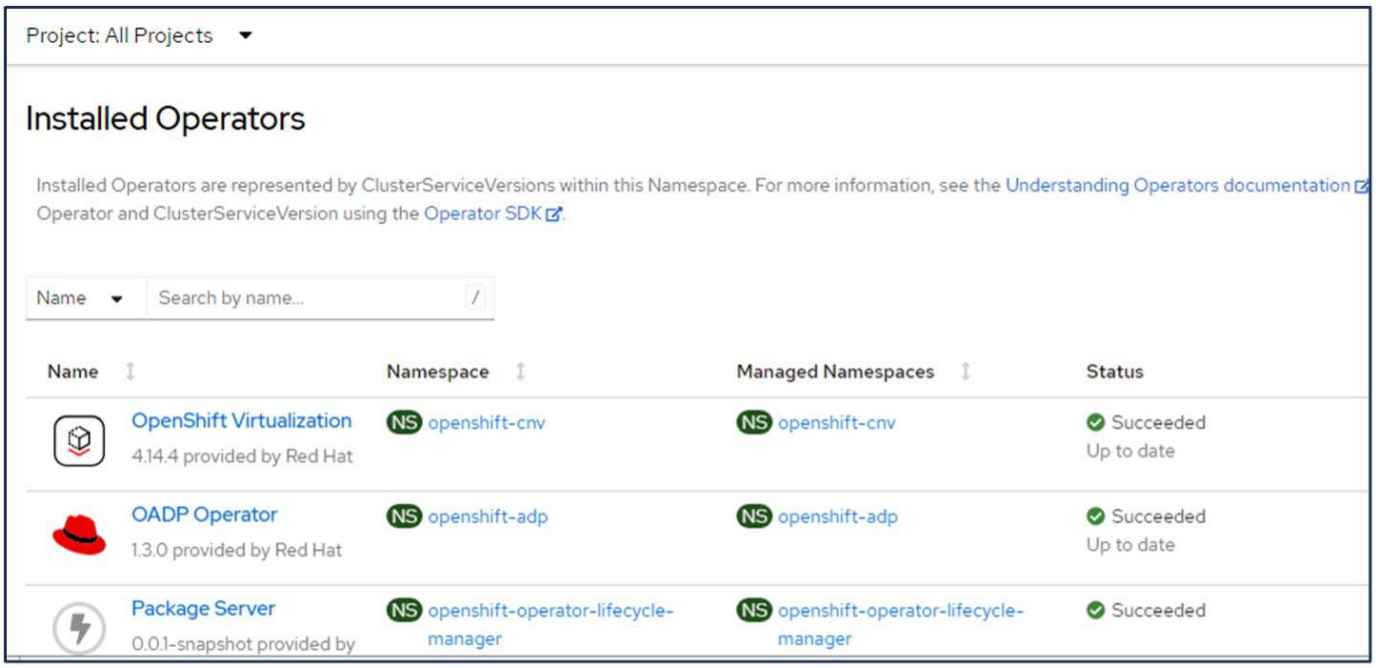
- A Red Hat OpenShift cluster (later than version 4.12) installed on bare-metal infrastructure with RHCOS worker nodes
- A NetApp ONTAP cluster integrated with the cluster using Trident
- A Trident backend configured with an SVM on ONTAP cluster
- A StorageClass configured on the OpenShift cluster with Trident as the provisioner
- Trident Snapshot class created on the cluster
- Cluster-admin access to Red Hat OpenShift cluster
- Admin access to NetApp ONTAP cluster
- An application eg. postgresql deployed on the cluster
- An admin workstation with tridentctl and oc tools installed and added to $PATH

Steps to install OADP Operator

1. Go to the Operator Hub of the cluster and select Red Hat OADP operator. In the Install page, use all the default selections and click install. On the next page, again use all the defaults and click Install. The OADP operator will be installed in the namespace openshift-adp.


Home
Operators
OperatorHub
Installed Operators
Workloads
Virtualization
Networking
Storage
Builds
Observe
OperatorHub
Discover Operators from the Kubernetes community and Red Hat partners, curated by Red Hat. You can purchase commercial software through Red Hat Ma
optional add-ons and shared services to your developers. After installation, the Operator capabilities will appear in the Developer Catalog providing a self-se
All Items
AI/Machine Learning
Application Runtime
Big Data
Cloud Provider
Database
Developer Tools
Development Tools
Drivers and plugins
Integration & Delivery
Logging & Tracing
Modernization & Migration
All Items
OADP
Red Hat
OADP Operator
provided by Red Hat
OADP (OpenShift API for Data Protection) operator sets up and installs Data Protection...
Community
OADP Operator
provided by Red Hat
OADP (OpenShift API for Data Protection) operator sets up and installs Velero on the OpenShift...


OADP Operator
1.3.0 provided by Red Hat
Install
Channel
stable-1.3
Version
1.3.0
Capability level
Basic Install
Seamless Upgrades
Full Lifecycle
Deep Insights
Auto Pilot
Source
Red Hat
Provider
Red Hat
Infrastructure features
Disconnected
OpenShift API for Data Protection (OADP) operator sets up and installs Velero on the OpenShift platform, allowing users to backup and restore applications.
Backup and restore Kubernetes resources and internal images, at the granularity of a namespace, using a version of Velero appropriate for the installed version of OADP.
OADP backs up Kubernetes objects and internal images by saving them as an archive file on object storage. OADP backs up persistent volumes (PVs) by creating snapshots with the native cloud snapshot API or with the Container Storage Interface (CSI). For cloud providers that do not support snapshots, OADP backs up resources and PV data with Restic or Kopia.
Installing OADP for application backup and restore
Installing OADP on a ROSA cluster and using STS, please follow the Getting Started Steps 1-3 in order to obtain the role ARN needed for using the standardized STS configuration flow via OLM
Frequently Asked Questions



Prerequisites for Velero configuration with Ontap S3 details

After the installation of the operator succeeds, configure the instance of Velero.
Velero can be configured to use S3 compatible Object Storage. Configure ONTAP S3 using the procedures shown in the Object Storage Management section of ONTAP documentation. You will need the following information from your ONTAP S3 configuration to integrate with Velero.

- A Logical Interface (LIF) that can be used to access S3
- User credentials to access S3 that includes the access key and the secret access key
- A bucket name in S3 for backups with access permissions for the user
- For secure access to the Object storage, TLS certificate should be installed on the Object Storage server.

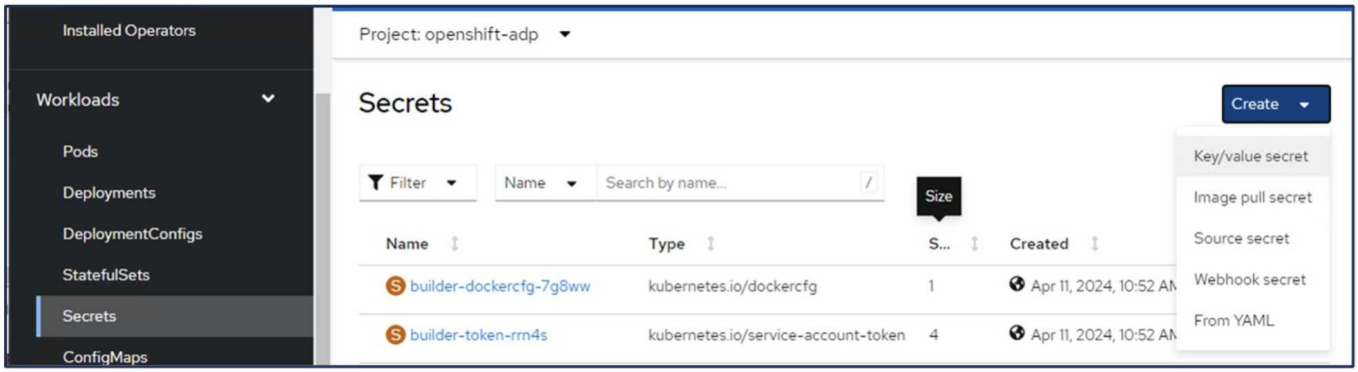
Prerequisites for Velero configuration with StorageGrid S3 details

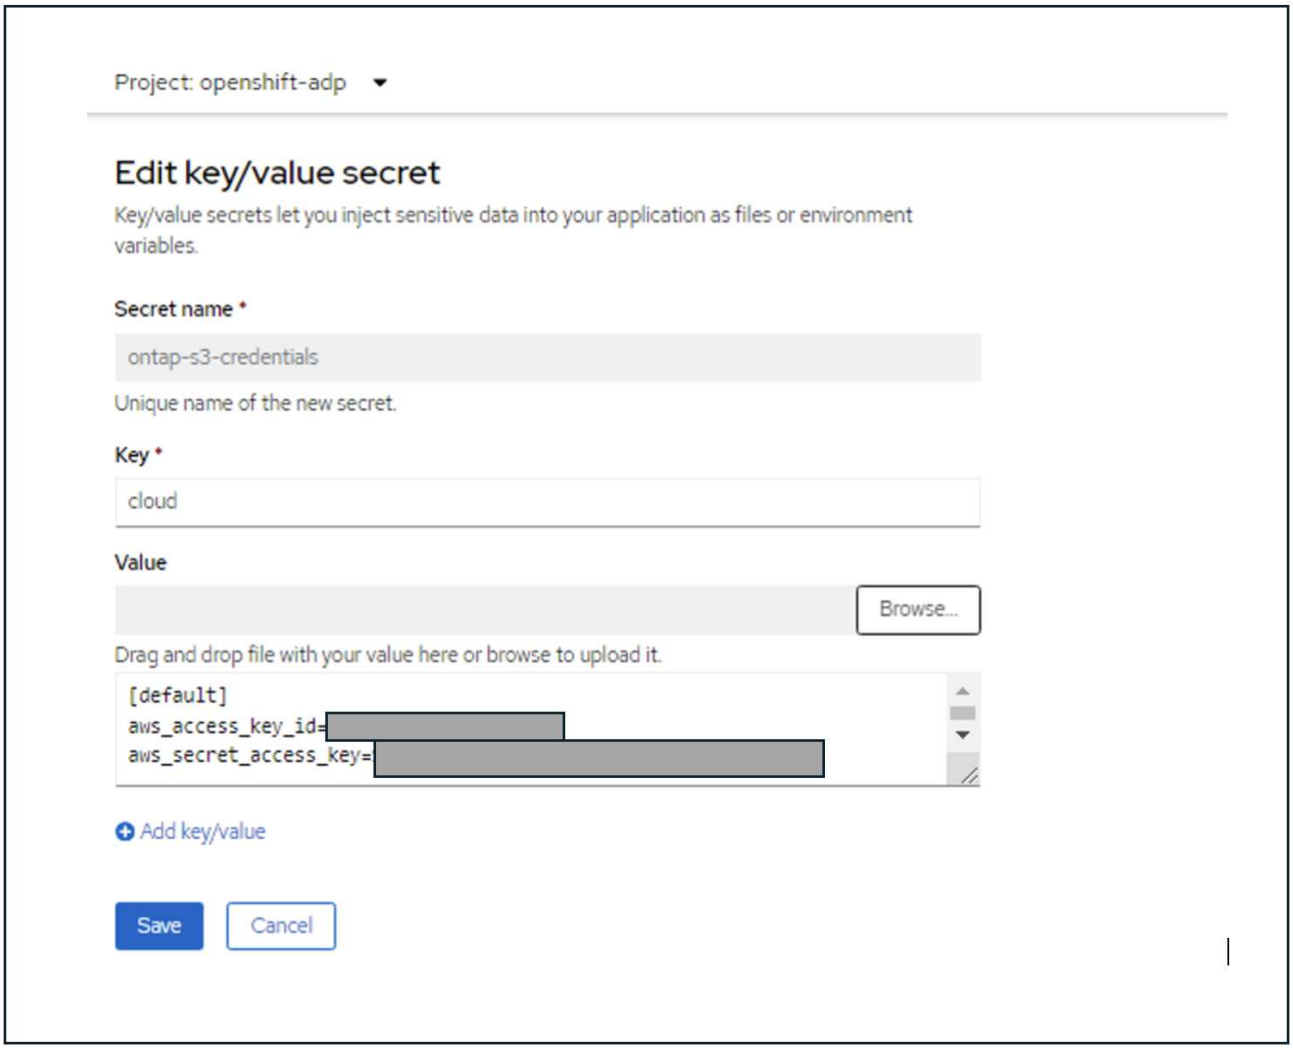
Velero can be configured to use S3 compatible Object Storage. You can configure StorageGrid S3 using the procedures shown in the StorageGrid documentation. You will need the following information from your StorageGrid S3 configuration to integrate with Velero.

- The endpoint that can be used to access S3
- User credentials to access S3 that includes the access key and the secret access key
- A bucket name in S3 for backups with access permissions for the user
- For secure access to the Object storage, TLS certificate should be installed on the Object Storage server.

Steps to configure Velero

- First, create a secret for an ONTAP S3 user credential or StorageGrid Tenant user credentials. This will be used to configure Velero later. You can create a secret from the CLI or from the web console.
  To create a secret from the web console, select Secrets, then click on Key/Value Secret. Provide the values for the credential name, key and the value as shown. Be sure to use the Access Key Id and Secret Access Key of your S3 user. Name the secret appropriately. In the sample below, a secret with ONTAP S3 user credentials named ontap-s3-credentials is created.





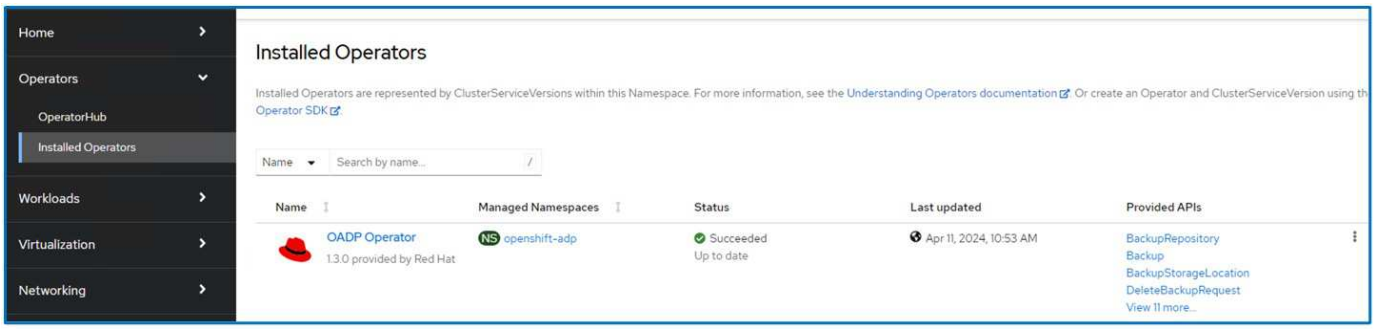
To create a secret named sg-s3-credentials from the CLI you can use the following command.

```
# oc create secret generic sg-s3-credentials --namespace openshift-adp --from-file
cloud=cloud-credentials.txt

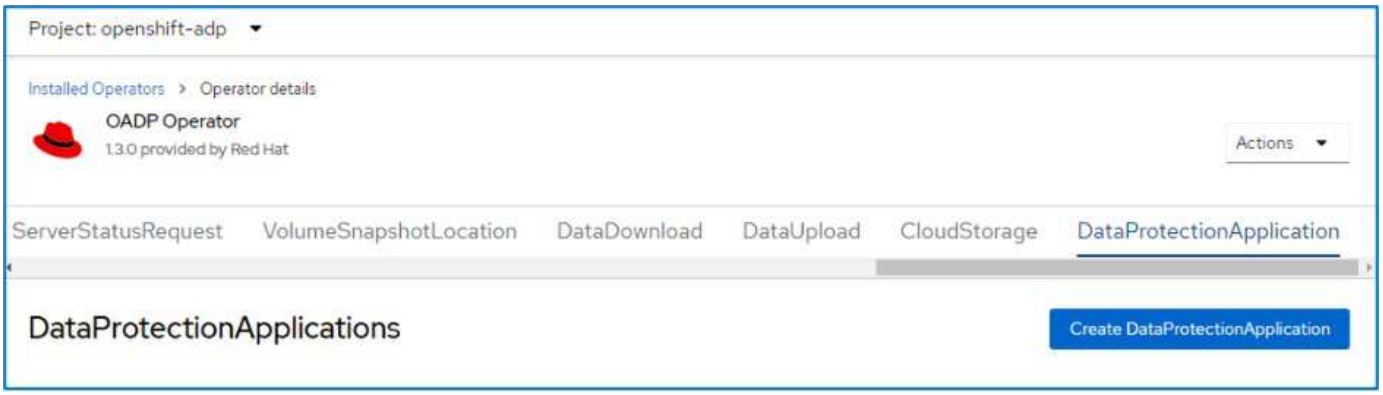
Where credentials.txt file contains the Access Key Id and the Secret Access Key of the S3
user in the following format:

[default]
aws_access_key_id=< Access Key ID of S3 user>
aws_secret_access_key=<Secret Access key of S3 user>
```

- Next, to configure Velero, select Installed Operators from the menu item under Operators, click on OADP operator, and then select the **DataProtectionApplication** tab.



Click on Create DataProtectionApplication. In the form view, provide a name for the DataProtection Application or use the default name.



Now go to the YAML view and replace the spec information as shown in the yaml file examples below.

Sample yaml file for configuring Velero with ONTAP S3 as the backupLocation

```
spec:
  backupLocations:
    - velero:
        config:
          insecureSkipTLSVerify: 'false' ->use this for https
```

```
communication with ONTAP S3
          profile: default
          region: us-east-1
          s3ForcePathStyle: 'true'  ->This allows use of IP in s3URL
          s3Url: 'https://10.61.181.161' ->Ensure TLS certificate for S3
is configured
        credential:
          key: cloud
          name: ontap-s3-credentials -> previously created secret
        default: true
        objectStorage:
          bucket: velero -> Your bucket name previously created in S3 for
backups
          prefix: container-demo-backup ->The folder that will be created
in the bucket
          caCert: <base64 encoded CA Certificate installed on ONTAP
Cluster with the SVM Scope where the bucker exists>
        provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
      #default Data Mover uses Kopia to move snapshots to Object Storage
    velero:
      defaultPlugins:
        - csi ->This plugin to use CSI snapshots
        - openshift
        - aws
        - kubevirt -> This plugin to use Velero with OIpenShift
Virtualization
```

Sample yaml file for configuring Velero with StorageGrid S3 as the backupLocation

```
spec:
  backupLocations:
    - velero:
        config:
          insecureSkipTLSVerify: 'true'
          profile: default
          region: us-east-1 ->region of your StorageGrid system
          s3ForcePathStyle: 'True'
          s3Url: 'https://172.21.254.25:10443' ->the IP used to access S3
        credential:
          key: cloud
          name: sg-s3-credentials ->secret created earlier
```

```
        default: true
        objectStorage:
          bucket: velero
          prefix: demobackup
        provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
        - aws
        - kubevirt
```

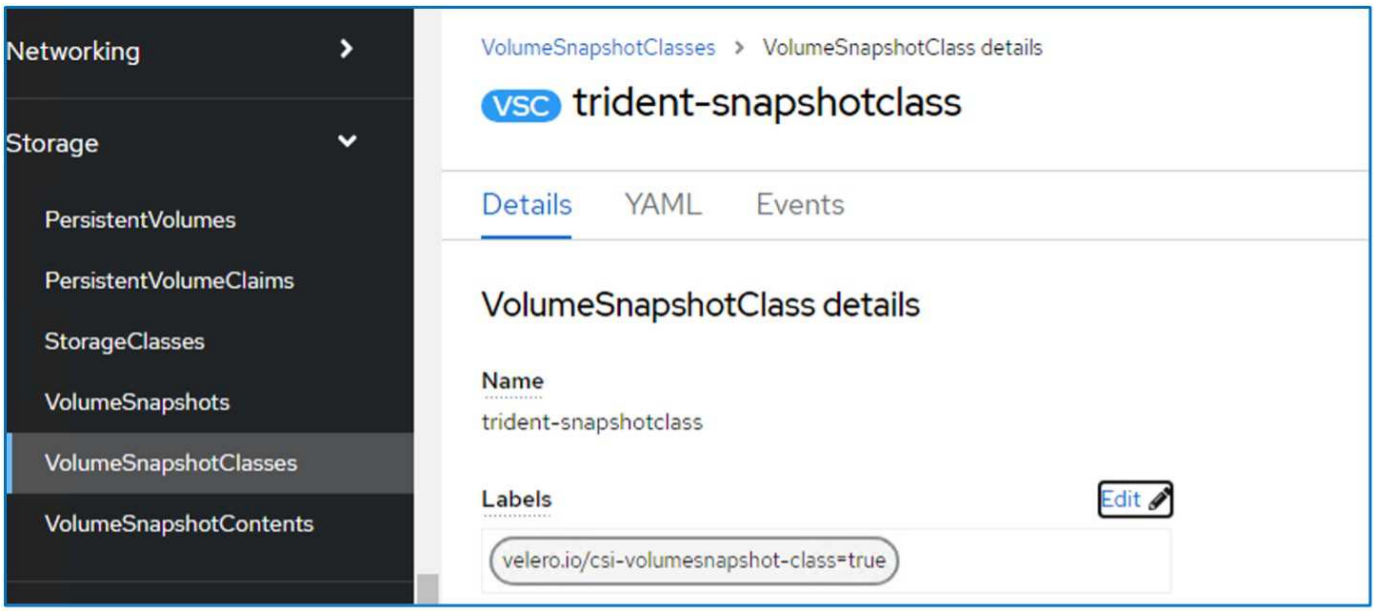
The spec section in the yaml file should be configured appropriately for the following parameters similar to the example above

backupLocations
ONTAP S3 or StorageGrid S3 (with its credentials and other information as shown in the yaml) is configured as the default BackupLocation for velero.

snapshotLocations
If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a VolumeSnapshotClass CR to register the CSI driver. In our example, you use Trident CSI and you have previously created VolumeSnapShotClass CR using the Trident CSI driver.

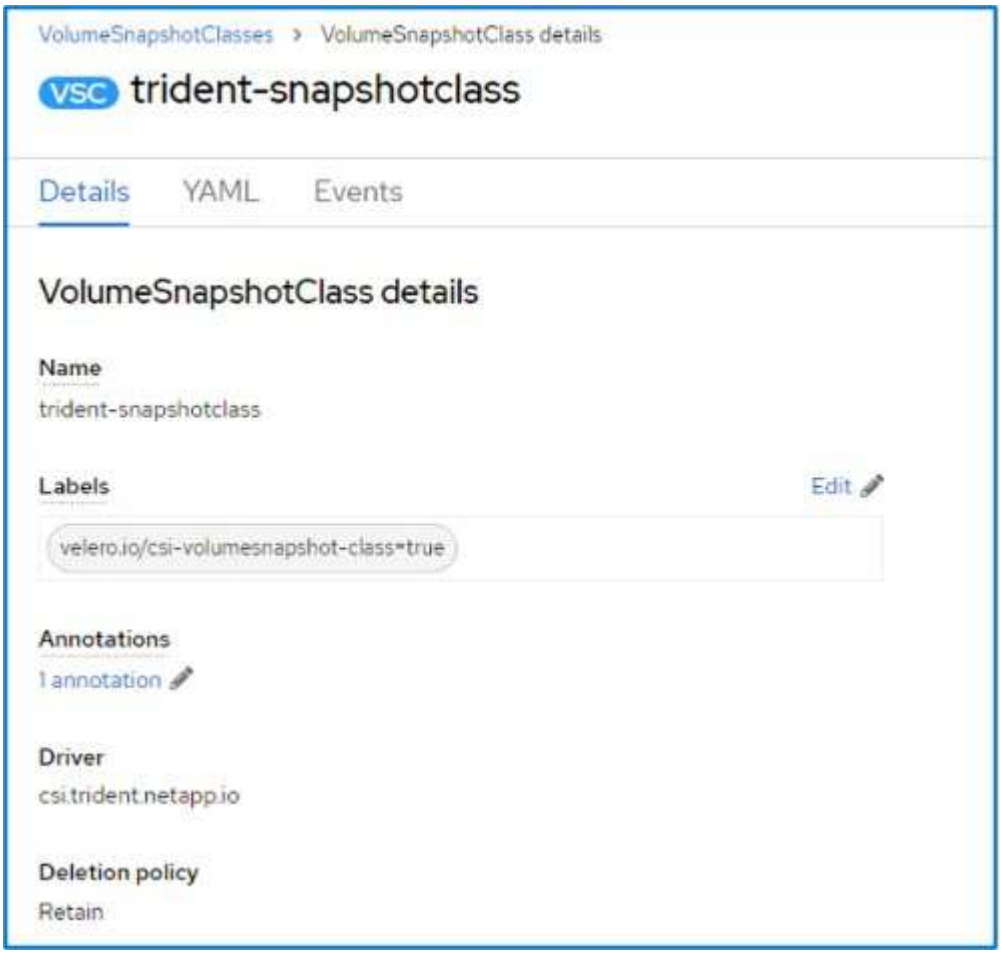
Enable CSI plugin
Add csi to the defaultPlugins for Velero to back up persistent volumes with CSI snapshots.
The Velero CSI plugins, to backup CSI backed PVCs, will choose the VolumeSnapshotClass in the cluster that has **velero.io/csi-volumesnapshot-class** label set on it. For this

- You must have the trident VolumeSnapshotClass created.
- Edit the label of the trident-snapshotclass and set it to **velero.io/csi-volumesnapshot-class=true** as shown below.

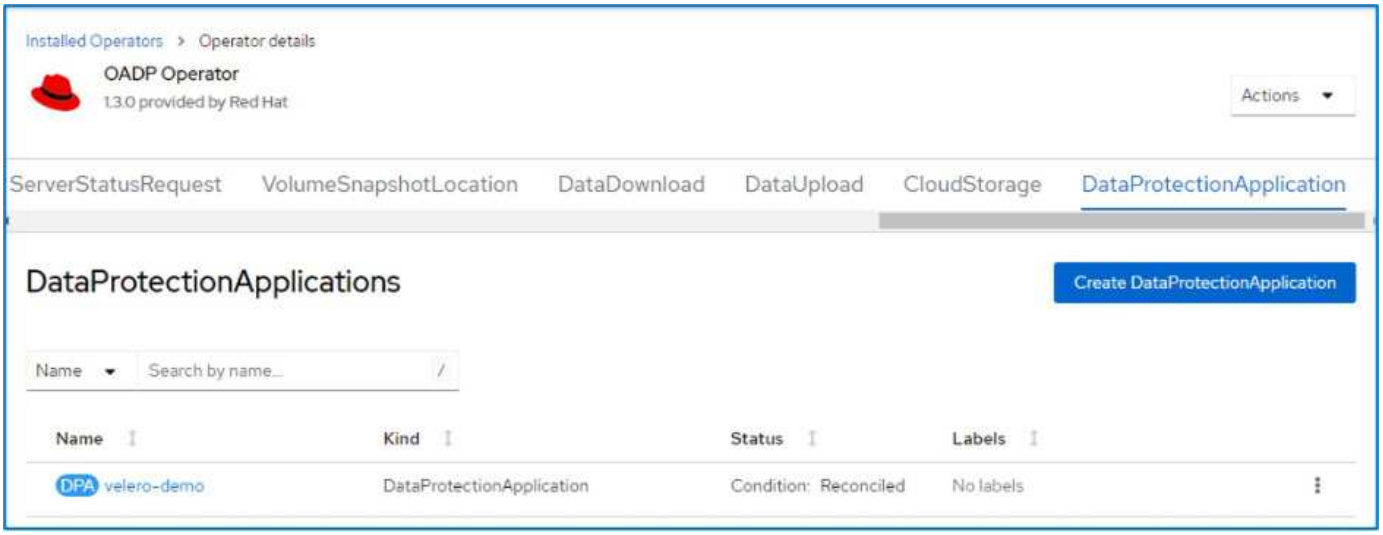


Ensure that the snapshots can persist even if the VolumeSnapshot objects are deleted. This can be done by setting the **deletionPolicy** to Retain. If not, deleting a namespace will completely lose all PVCs ever backed up in it.

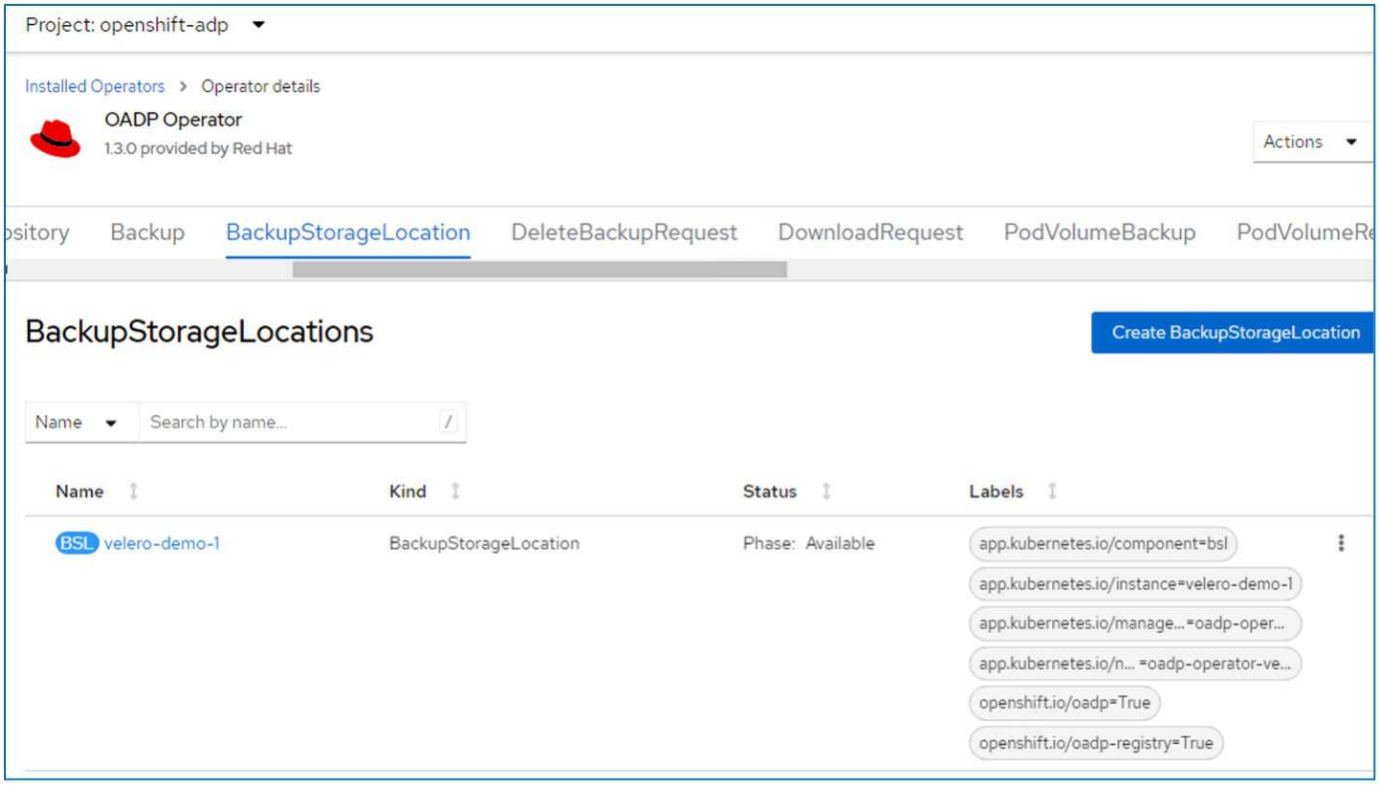
```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: trident-snapshotclass
driver: csi.trident.netapp.io
deletionPolicy: Retain
```



Ensure that the DataProtectionApplication is created and is in condition:Reconciled.



The OADP operator will create a corresponding BackupStorageLocation.This will be used when creating a backup.



Creating on-demand backup for Apps in OpenShift Container Platform

This section outlines how to create on-demand backup for VMs in OpenShift Virtualization.

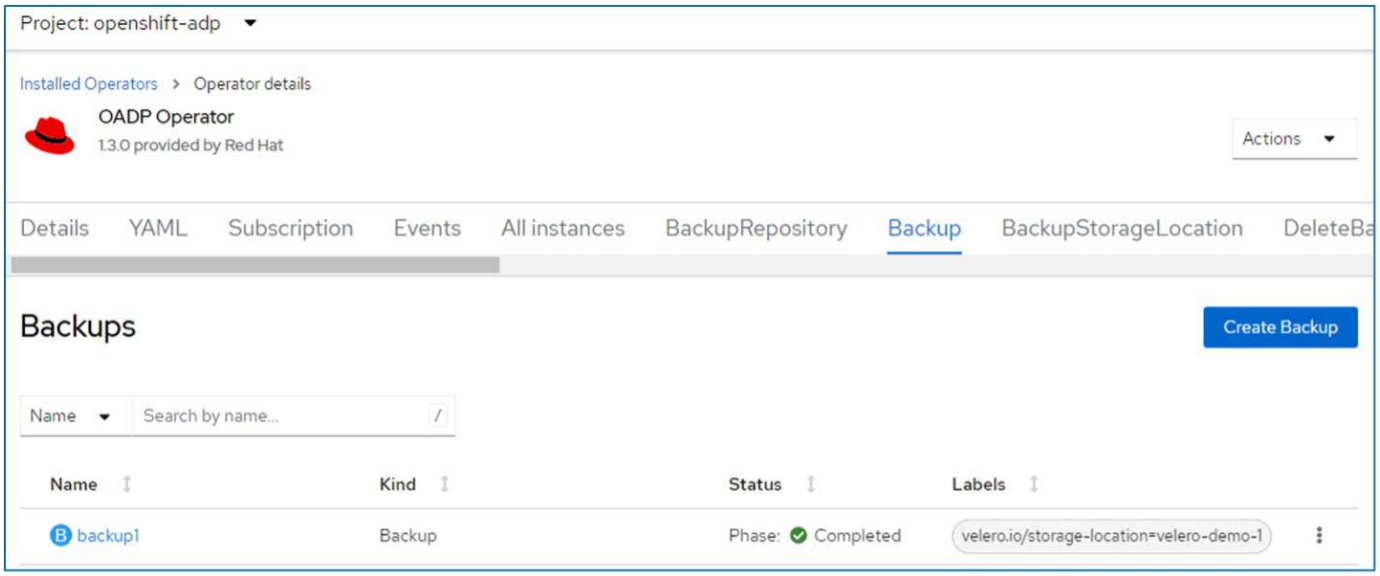
Steps to create a backup of an App

To create an on-demand backup of an app (app metadata and persistent volumes of the app), click on the **Backup** tab to create a Backup Custom Resource (CR). A sample yaml is provided to create the Backup CR. Using this yaml, the app and its persistent storage in the specified namespace will be backed up. Additional parameters can be set as shown in the documentation.

A snapshot of the persistent volumes and the app resources in the namespace specified will be created by the CSI. This snapshot will be stored in the backup location specified in the yaml. The backup will remain in the system for 30 days as specified in the ttl.

```
spec:
  csiSnapshotTimeout: 10m0s
  defaultVolumesToFsBackup: false
  includedNamespaces:
    - postgresql ->namespace of the app
  itemOperationTimeout: 4h0m0s
  snapshotMoveData: false
  storageLocation: velero-container-backup-ontap-1 -->this is the
backupStorageLocation previously created when Velero is configured.
  ttl: 720h0m0s
```

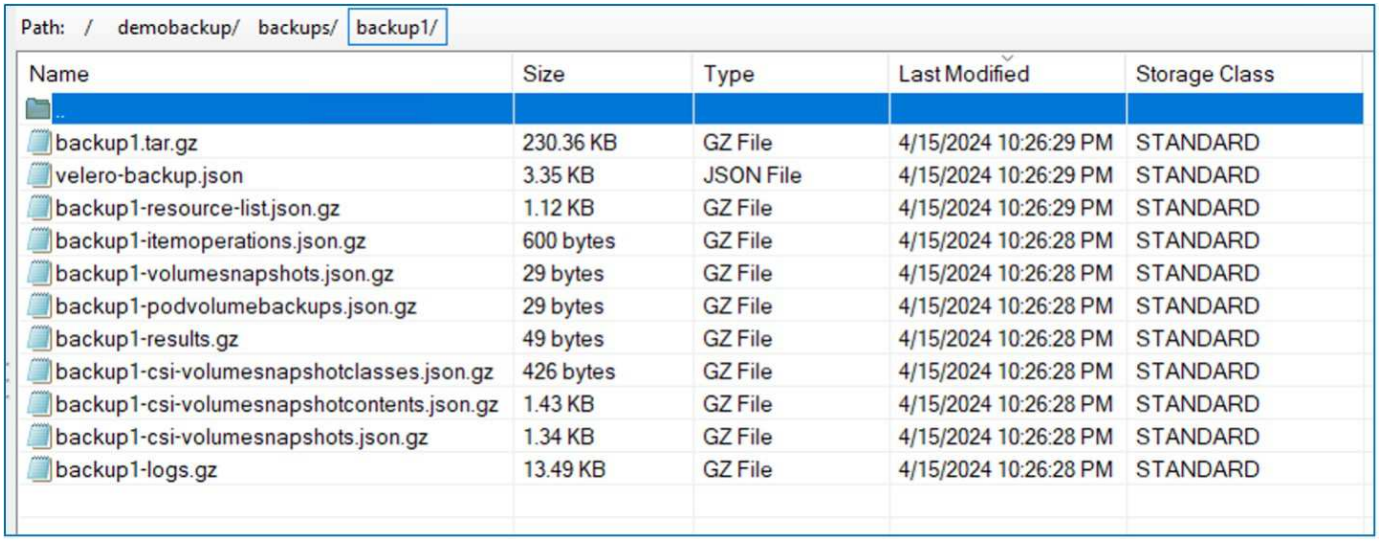
Once the backup completes, its Phase will show as completed.



You can inspect the backup in the Object storage with the help of an S3 browser application. The path of the backup shows up in the configured bucket with the prefix name (velero/container-demo-backup). You can see the contents of the backup includes the volume snapshots, logs, and other metadata of the application.



In StorageGrid, you can also use the S3 console that is available from the Tenant Manager to view the backup objects.



Creating scheduled backups for Apps

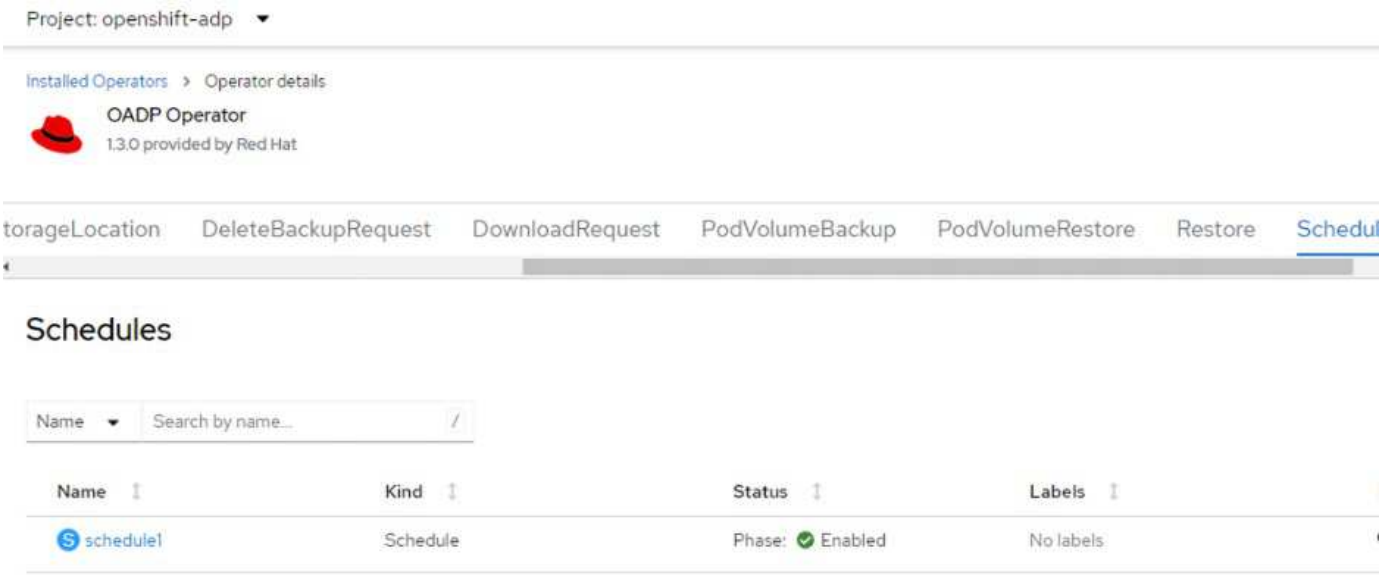
To create backups on a schedule, you need to create a Schedule CR.
The schedule is simply a Cron expression allowing you to specify the time at which you want to create the backup. A sample yaml to create a Schedule CR is shown below.

```
apiVersion: velero.io/v1
kind: Schedule
metadata:
```

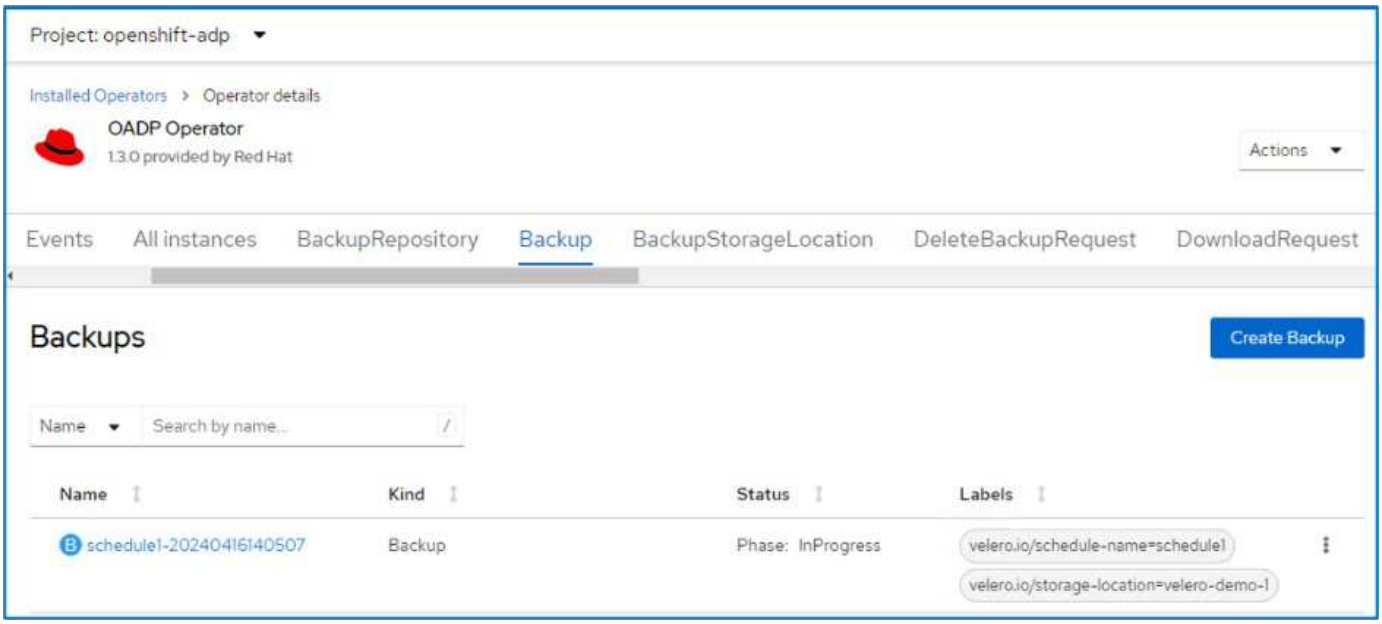
```
  name: schedule1
  namespace: openshift-adp
spec:
  schedule: 0 7 * * *
  template:
    includedNamespaces:
      - postgresql
    storageLocation: velero-container-backup-ontap-1
```

The Cron expression 0 7 * * * means a backup will be created at 7:00 every day.
The namespaces to be included in the backup and the storage location for the backup are also specified. So instead of a Backup CR, Schedule CR is used to create a backup at the specified time and frequency.

Once the schedule is created, it will be Enabled.



Backups will be created according to this schedule, and can be viewed from the Backup tab.



Migrate an App from one cluster to another

Velero's backup and restore capabilities make it a valuable tool for migrating your data between clusters. This section describes how to migrate apps(s) from one cluster to another by creating a backup of the app in Object storage from one cluster and then restoring the app from the same object storage to another cluster. .

▼ **Backup from first cluster**

Prerequisites on Cluster 1

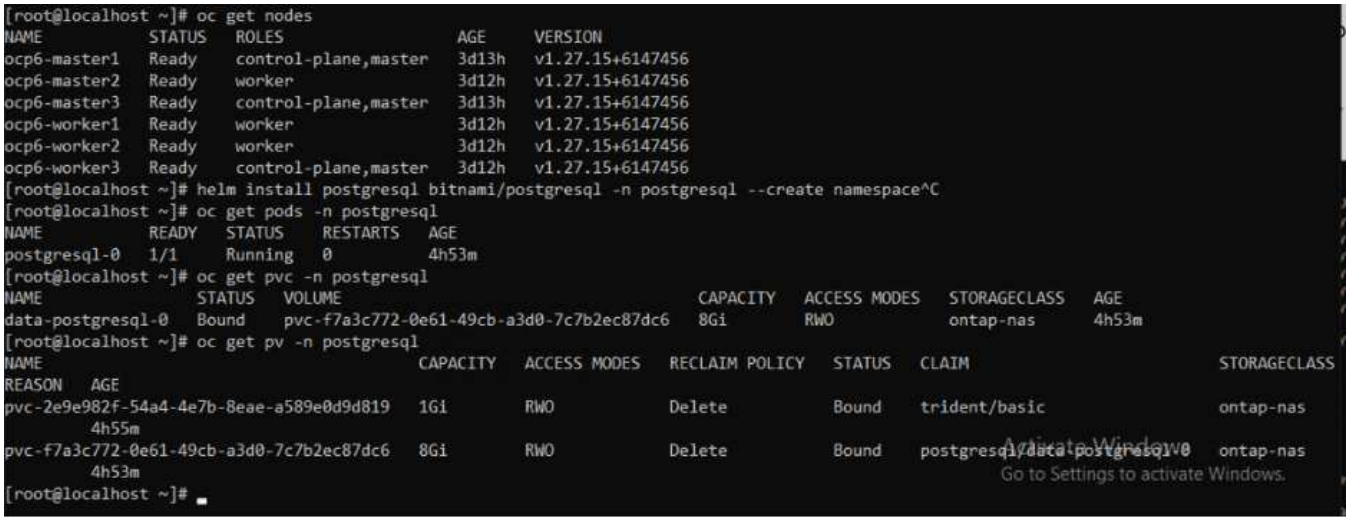
- Trident must be installed on the cluster.
- A trident backend and Storage class must be created.
- OADP operator must be installed on the cluster.
- The DataProtectionApplication should be configured.

Use the following spec to configure the DataProtectionApplication object.

```
spec:
  backupLocations:
    - velero:
        config:
          insecureSkipTLSVerify: 'false'
          profile: default
          region: us-east-1
          s3ForcePathStyle: 'true'
          s3Url: 'https://10.61.181.161'
        credential:
          key: cloud
          name: ontap-s3-credentials
        default: true
```

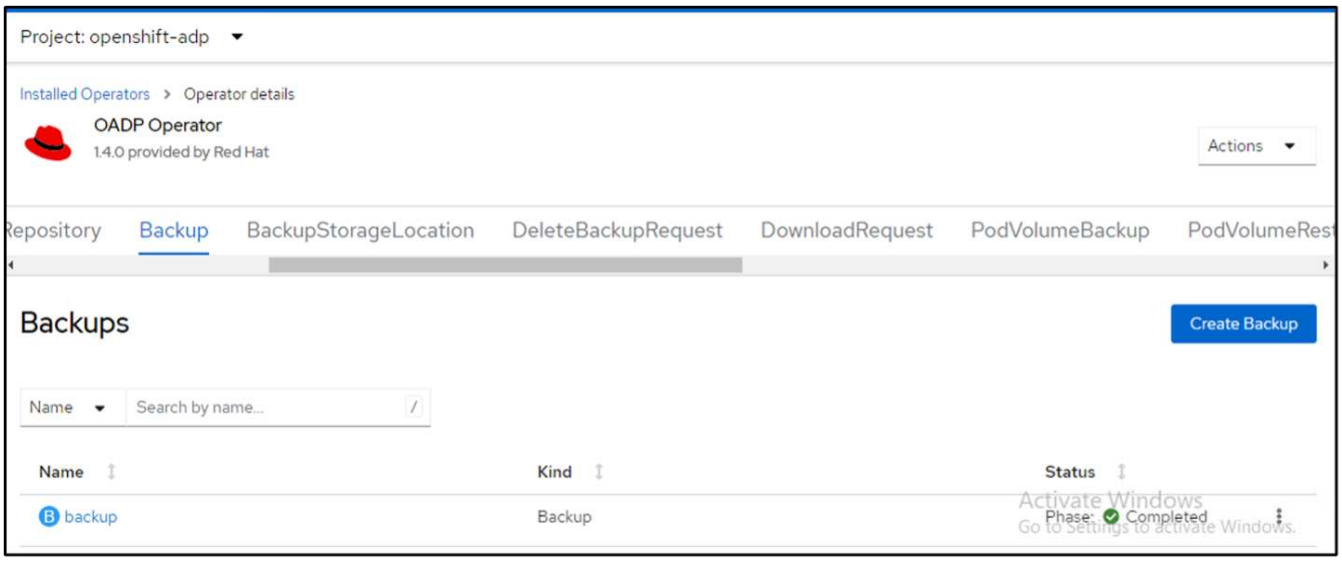
```
        objectStorage:
          bucket: velero
          caCert: <base-64 encoded tls certificate>
          prefix: container-backup
        provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
        - aws
        - kubevirt
```

- Create an application on the cluster and take a backup of this application.
  As an example, install a postgres application.



- Use the following spec for the backup CR:

```
spec:
  csiSnapshotTimeout: 10m0s
  defaultVolumesToFsBackup: false
  includedNamespaces:
    - postgresql
  itemOperationTimeout: 4h0m0s
  snapshotMoveData: true
  storageLocation: velero-sample-1
  ttl: 720h0m0s
```



You can click on the **All instances** tab to see the different objects being created and moving through different phases to finally come to the backup **completed** phase.

A backup of the resources in the namespace postgresql will be stored in the Object Storage location (ONTAP S3) specified in the backupLocation in the OADP spec.

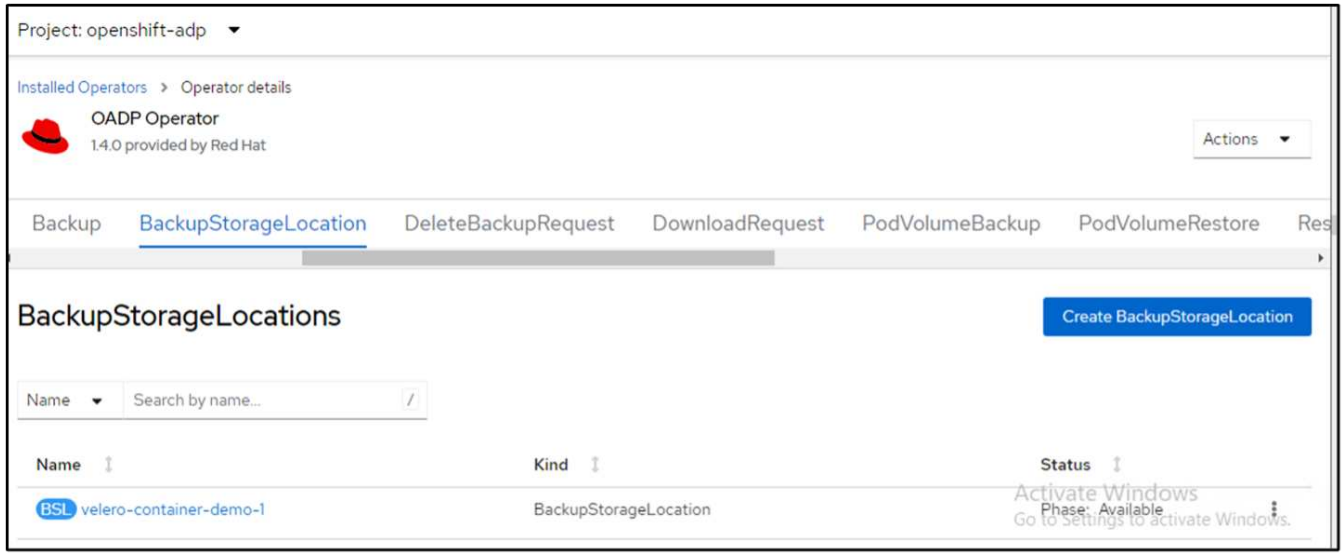
▼ **Restore to a second cluster**

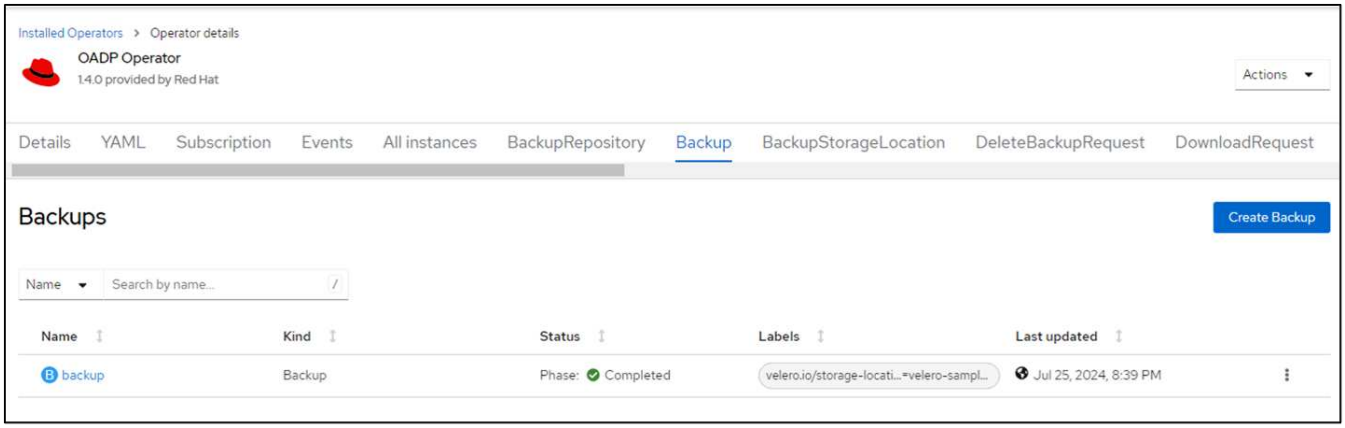
Prerequisites on Cluster 2

- Trident must be installed on cluster 2.
- The postgresql app must NOT be already installed in the postgresql namespace.
- OADP operator must be installed on cluster 2, and the BackupStorage Location must be pointing to the same object storage location where the backup was stored from the first cluster.
- The Backup CR must be visible from the second cluster.

```
[root@localhost ~]# oc get pods -n trident
NAME                                READY   STATUS    RESTARTS   AGE
trident-controller-6799cfb77f-8rzvk   6/6     Running   6          2d7h
trident-node-linux-7wvjz            2/2     Running   2          2d7h
trident-node-linux-8vvm2            2/2     Running   0          2d7h
trident-node-linux-bgs6f            2/2     Running   2          2d7h
trident-node-linux-njwb8            2/2     Running   0          2d7h
trident-node-linux-scqjl            2/2     Running   0          2d7h
trident-node-linux-swr69            2/2     Running   2          2d7h
trident-operator-b88b86fc8-7fk68    1/1     Running   1          2d7h
[root@localhost ~]#
```

```
[root@localhost ~]# oc get nodes
NAME           STATUS   ROLES                  AGE   VERSION
ocp7-master1   Ready    control-plane,master   3d    v1.27.15+6147456
ocp7-master2   Ready    control-plane,master   3d    v1.27.15+6147456
ocp7-master3   Ready    control-plane,master   3d    v1.27.15+6147456
ocp7-worker1   Ready    worker                 3d    v1.27.15+6147456
ocp7-worker2   Ready    worker                 3d    v1.27.15+6147456
ocp7-worker3   Ready    worker                 3d    v1.27.15+6147456
[root@localhost ~]# oc get pods -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]# oc get pvc -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]# oc get pv -n postgresql
NAME                                       CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM                 STORAGECLASS   REASON   AGE
pvc-c6660630-0cfe-484b-aaa3-5ada54c8b9a7   1Gi        RWO            Delete           Bound    trident/basic         ontap-nas               11m
pvc-edcc6551-81b0-40b4-8547-e9df70c1740d   10Gi       RWO            Delete           Bound    default/test-pvc      vsphere-sc              2d7
h
[root@localhost ~]#
```



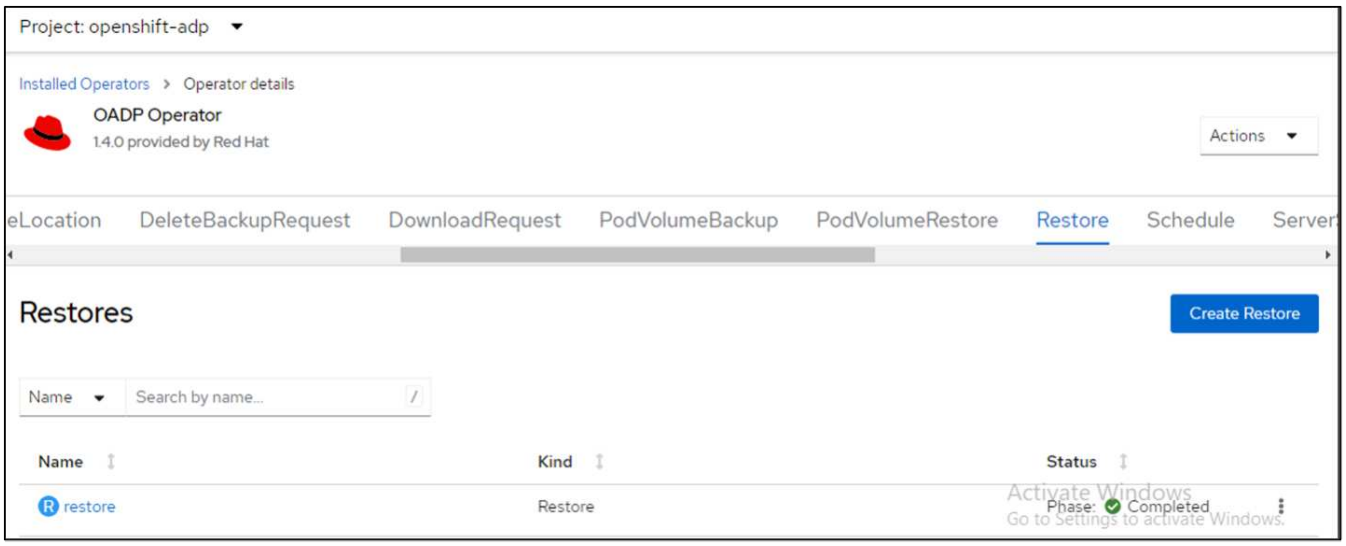


Restore the app on this cluster from the backup. Use the following yaml to create the Restore CR.

```
apiVersion: velero.io/v1
kind: Restore
apiVersion: velero.io/v1
metadata:
  name: restore
  namespace: openshift-adp
spec:
```

```
  backupName: backup
  restorePVs: true
```

When the restore is completed, you will see that the postgresql app is running on this cluster and is associated with the pvc and a corresponding pv. The state of the app is the same as when the backup was taken.




```
[root@localhost ~]# export KUBECONFIG=ocp-cluster7/kubeconfig-ocp-cluster7
[root@localhost ~]# oc get nodes
NAME          STATUS   ROLES                  AGE    VERSION
ocp7-master1  Ready    control-plane,master   3d3h   v1.27.15+6147456
ocp7-master2  Ready    control-plane,master   3d3h   v1.27.15+6147456
ocp7-master3  Ready    control-plane,master   3d3h   v1.27.15+6147456
ocp7-worker1  Ready    worker                 3d3h   v1.27.15+6147456
ocp7-worker2  Ready    worker                 3d3h   v1.27.15+6147456
ocp7-worker3  Ready    worker                 3d3h   v1.27.15+6147456
[root@localhost ~]# oc get pods -n postgresql
NAME           READY   STATUS    RESTARTS   AGE
postgresql-0   1/1     Running   0          31m
[root@localhost ~]# oc get pvc -n postgresql
NAME                 STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS   AGE
data-postgresql-0    Bound    pvc-ce7044e3-2ba5-4934-8bad-553fa7d35128   8Gi        RWO            ontap-nas      31m
[root@localhost ~]# oc get pv
NAME                                       CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM                           STORAGECLASS
REASON   AGE
pvc-c6660630-0cfe-484b-aaa3-5ada54c8b9a7   1Gi        RWO            Delete           Bound    trident/basic                   ontap-nas
         3h27m
pvc-ce7044e3-2ba5-4934-8bad-553fa7d35128   8Gi        RWO            Delete           Bound    postgresql/data-postgresql-0    ontap-nas
         31m
pvc-edcc6551-81b0-40b4-8547-e9df70c1740d   10Gi       RWO            Delete           Bound    default/test-pvc                vsphere-sc
         2d10h
[root@localhost ~]#
```

Activate Windows
Go to Settings to activate Windows.

Restore an App from a backup

This section describes how to restore apps(s) from a backup.

Prerequisites

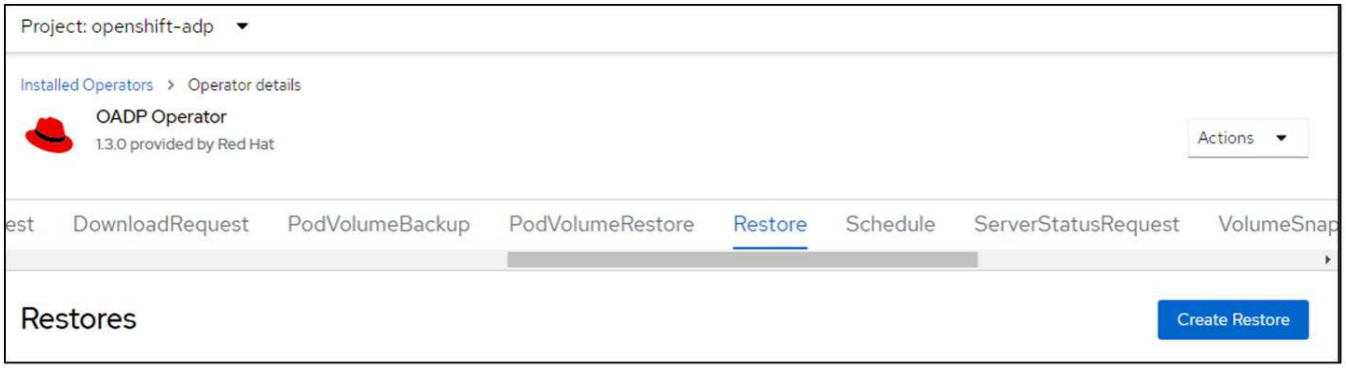
To restore from a backup, let us assume that the namespace where the app existed got accidentally deleted.

```
[root@localhost ~]# oc get pods -n postgresql
NAME           READY   STATUS    RESTARTS   AGE
postgresql-0   1/1     Running   0          102s
[root@localhost ~]# oc delete ns postgresql
namespace "postgresql" deleted

[root@localhost ~]#
[root@localhost ~]#
[root@localhost ~]# oc get pods -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]#
```

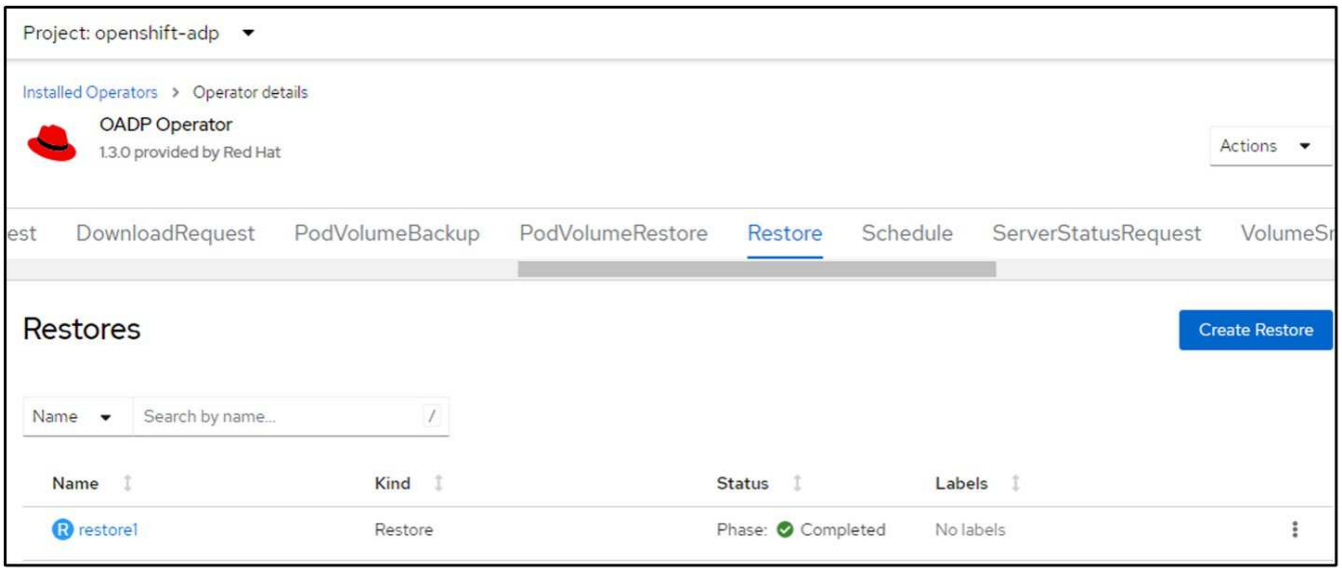
▼ **Restore to the same namespace**

To restore from the backup that we just created, we need to create a Restore Custom Resource (CR). We need to provide it a name, provide the name of the backup that we want to restore from and set the restorePVs to true. Additional parameters can be set as shown in the documentation. Click on Create button.




```
apiVersion: velero.io/v1
kind: Restore
apiVersion: velero.io/v1
metadata:
  name: restore
  namespace: openshift-adp
spec:
  backupName: backup-postgresql-ontaps3
  restorePVs: true
```

When the phase shows completed, you can see that the app has been restored to the state when the snapshot was taken. The app is restored to the same namespace.




```
[root@localhost ~]#
[root@localhost ~]# oc get pods -n postgresql
No resources found in postgresql namespace.
[root@localhost ~]# oc get pods -n postgresql
NAME           READY   STATUS              RESTARTS   AGE
postgresql-0   0/1     ContainerCreating   0          16s
[root@localhost ~]# oc get pods -n postgresql
NAME           READY   STATUS    RESTARTS   AGE
postgresql-0   0/1     Running   0          22s
[root@localhost ~]# oc get pods -n postgresql
NAME           READY   STATUS    RESTARTS   AGE
postgresql-0   0/1     Running   0          29s
[root@localhost ~]# oc get pods -n postgresql
NAME           READY   STATUS    RESTARTS   AGE
postgresql-0   1/1     Running   0          37s
[root@localhost ~]#
```

▼ **Restore to a different namespace**

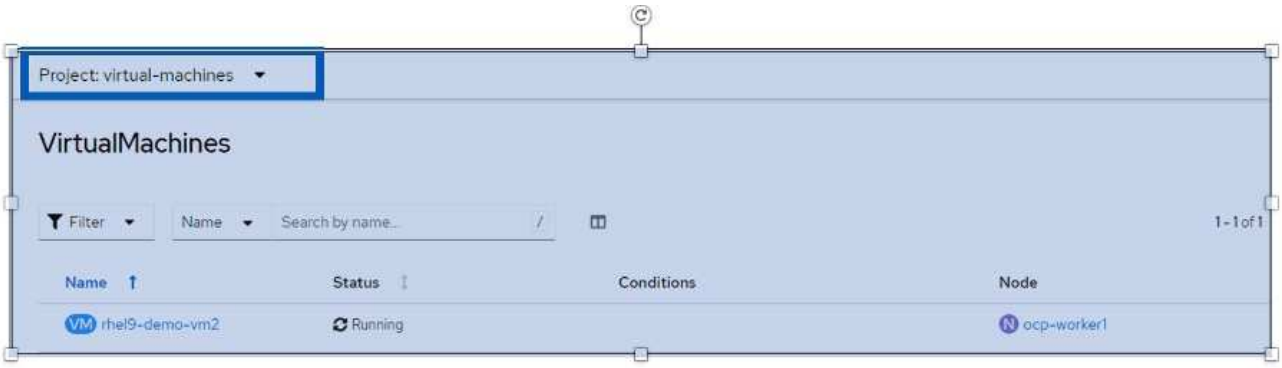
To restore the App to a different namespace, you can provide a namespaceMapping in the yaml definition of the Restore CR.

The following sample yaml file creates a Restore CR to restore an App and its persistent storage from the postgresql namespace, to the new namespace postgresql-restored.

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: restore-to-different-ns
  namespace: openshift-adp
spec:
  backupName: backup-postgresql-ontaps3
```

```
  restorePVs: true
  includedNamespaces:
  - postgresql
  namespaceMapping:
    postgresql: postgresql-restored
```

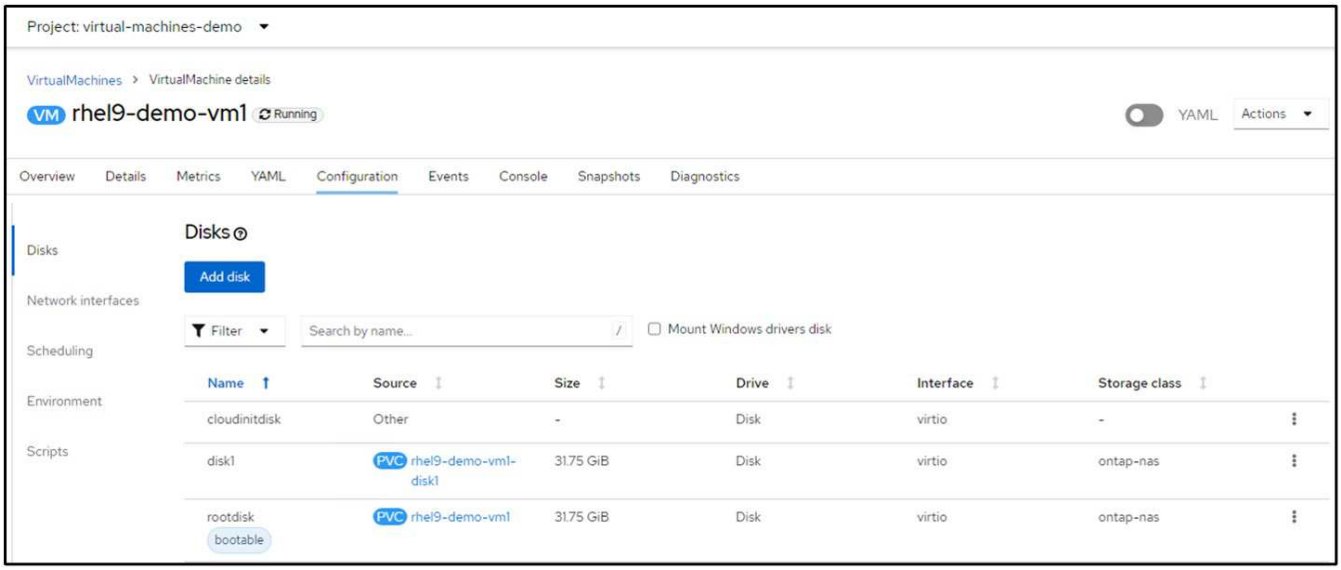
When the phase shows completed, you can see that the app has been restored to the state when the snapshot was taken. The App is restored to a different namespace as specified in the yaml.

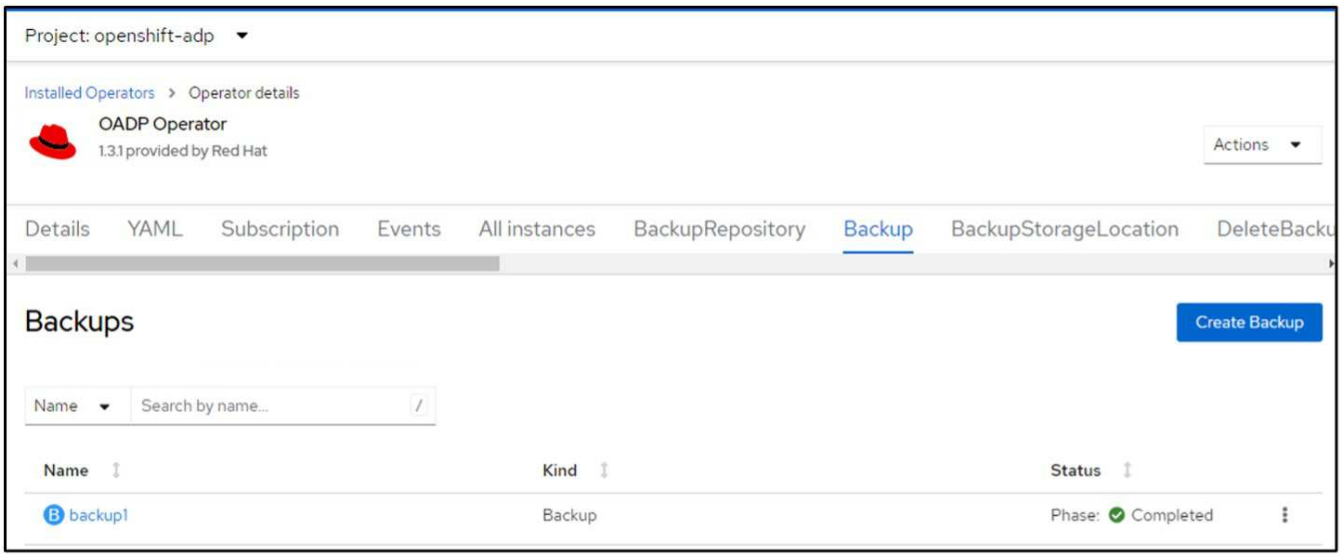


▼ **Restore to a different storage class**

Velero provides a generic ability to modify the resources during restore by specifying json patches. The json patches are applied to the resources before they are restored. The json patches are specified in a configmap and the configmap is referenced in the restore command. This feature enables you to restore using different storage class.

In the example below, the app, during deployment uses ontap-nas as the storage class for its persistent volumes. A backup of the app named backup-postgresql-ontaps3 is created.





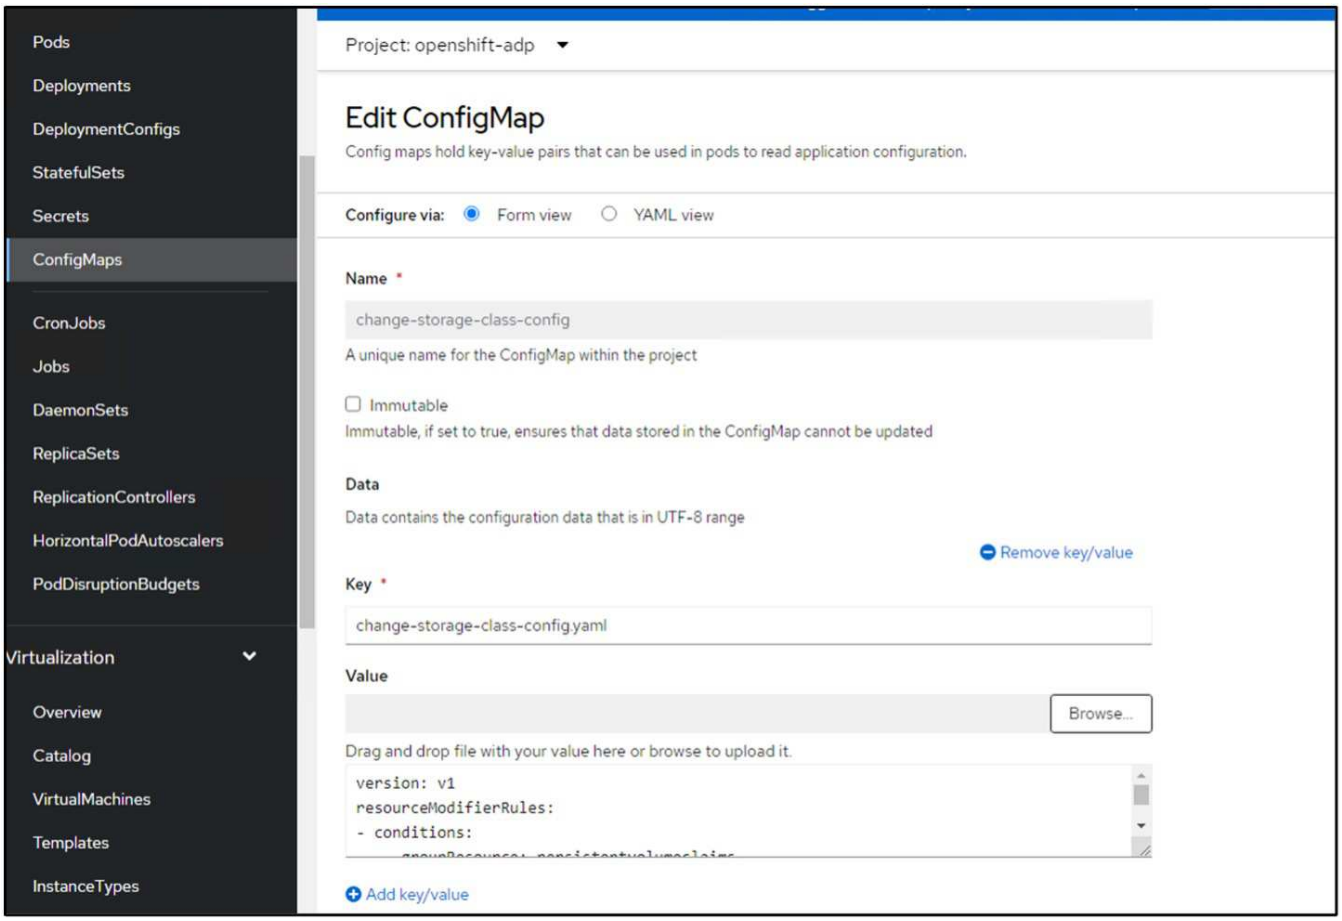
Simulate a loss of the app by uninstalling the app.

To restore the VM using a different storage class, for example, ontap-nas-eco storage class, you need to do the following two steps:

Step 1

Create a config map (console) in the openshift-adp namespace as follows:
Fill in the details as shown in the screenshot:
Select namespace : openshift-adp
Name: change-ontap-sc (can be any name)
Key: change-ontap-sc-config.yaml:
Value:

```
version: v1
resourceModifierRules:
- conditions:
     groupResource: persistentvolumeclaims
     resourceNameRegex: "data-postgresql*"
     namespaces:
     - postgresql
  patches:
  - operation: replace
    path: "/spec/storageClassName"
    value: "ontap-nas-eco"
```



The resulting config map object should look like this (CLI):

```
# kubectl describe cm/change-storage-class-config -n openshift-
adp
Name:         change-storage-class-config
Namespace:    openshift-adp
Labels:       velero.io/change-storage-class=RestoreItemAction
              velero.io/plugin-config=
Annotations:  <none>

Data
====
change-storage-class-config.yaml:
----
version: v1
resourceModifierRules:
- conditions:
     groupResource: persistentvolumeclaims
     resourceNameRegex: "^rhel*"
     namespaces:
     - virtual-machines-demo
  patches:
  - operation: replace
    path: "/spec/storageClassName"
    value: "ontap-nas-eco"


BinaryData
====

Events:  <none>
```

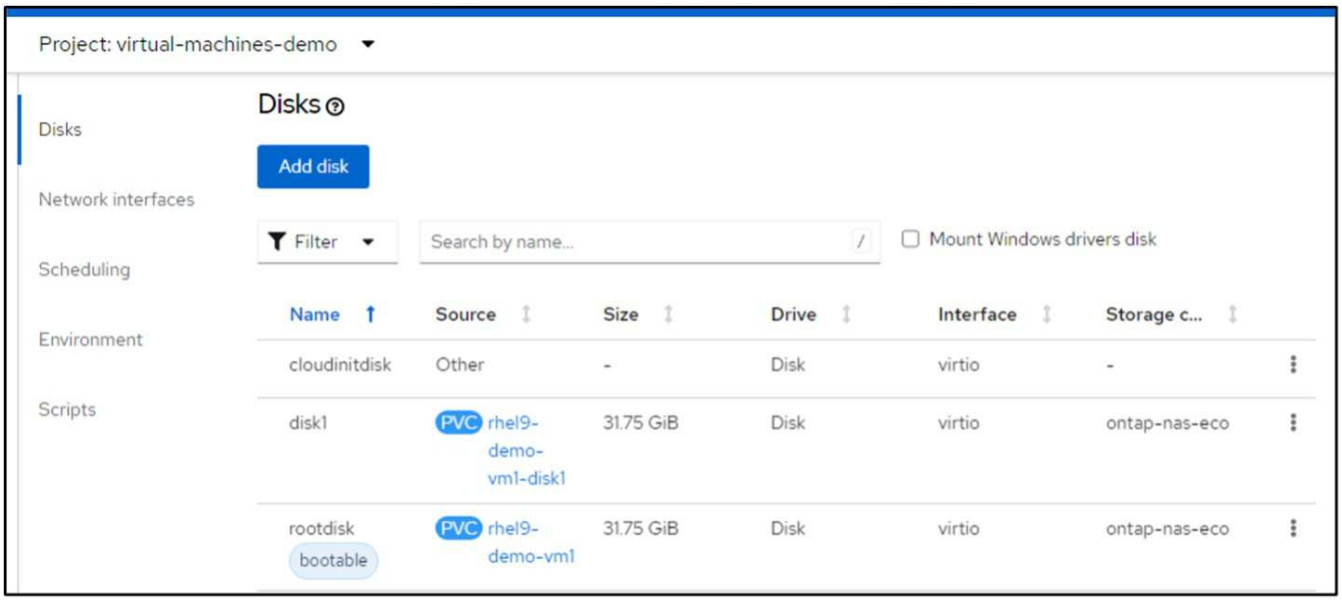
This config map will apply the resource modifier rule when the restore is created. A patch will be applied to replace the storage class name to ontap-nas-eco for all persistent volume claims starting with rhel.

Step 2

To restore the VM use the following command from the Velero CLI:

```
#velero restore create restore1 --from-backup backup1 --resource
-modifier-configmap change-storage-class-config -n openshift-adp
```

The app is restored in the same namespace with the persistent volume claims created using the storage class ontap-nas-eco.



Deleting backups and restores in using Velero

This section outlines how to delete backups and restores of Apps in OpenShift container platform using Velero.

List all backups

You can list all Backup CRs by using the OC CLI tool or the Velero CLI tool.
Download the Velero CLI as given in the instructions in the Velero documentation.

```
[root@localhost ~]# oc get backups -n openshift-adp
NAME                      AGE
backup-postgresql-ontaps3   23h
backup2                     26s
schedule1-20240717070005    6h42m
[root@localhost ~]# velero get backups -n openshift-adp
NAME                      STATUS      ERRORS   WARNINGS   CREATED                         EXPIRES   STORAGE LOCATION                SELECTOR
backup-postgresql-ontaps3   Completed   0        0          2024-07-16 10:01:08 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
backup2                     Completed   0        0          2024-07-17 09:42:32 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
schedule1-20240717070005    Completed   0        0          2024-07-17 03:00:05 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
[root@localhost ~]#
```

Deleting a backup

You can delete a Backup CR without deleting the Object Storage data by using the OC CLI tool. The backup will be removed from the CLI/Console output. However, since the corresponding backup is not removed from the object storage, it will re-appear in the CLI/console output.

```
[root@localhost ~]# oc delete backup backup2 -n openshift-adp
backup.velero.io "backup2" deleted
[root@localhost ~]# oc get backups -n openshift-adp
NAME                          AGE
backup-postgresql-ontaps3     23h
schedule1-20240717070005      6h49m
[root@localhost ~]# oc get backups -n openshift-adp
NAME                          AGE
backup-postgresql-ontaps3     23h
backup2                       24s
schedule1-20240717070005      6h50m
[root@localhost ~]#
```

If you want to delete the Backup CR AND the associated object storage data, you can do so by using the Velero CLI tool.

```
[root@localhost ~]# velero get backups -n openshift-adp
NAME                          STATUS      ERRORS   WARNINGS   CREATED                         EXPIRES   STORAGE LOCATION                 SELECTOR
backup-postgresql-ontaps3     Completed   0        0          2024-07-16 10:01:08 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
backup2                       Completed   0        0          2024-07-17 09:42:32 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
schedule1-20240717070005      Completed   0        0          2024-07-17 03:00:05 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
[root@localhost ~]# velero delete backup backup2 -n openshift-adp
Are you sure you want to continue (Y/N)? Y
Request to delete backup "backup2" submitted successfully.
The backup will be fully deleted after all associated data (disk snapshots, backup files, restores) are removed.
[root@localhost ~]# velero get backups -n openshift-adp
NAME                          STATUS      ERRORS   WARNINGS   CREATED                         EXPIRES   STORAGE LOCATION                 SELECTOR
backup-postgresql-ontaps3     Completed   0        0          2024-07-16 10:01:08 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
schedule1-20240717070005      Completed   0        0          2024-07-17 03:00:05 -0400 EDT   29d       velero-container-backup-ontap-1   <none>
[root@localhost ~]#
```

Deleting the Restore

You can delete the Restore CR Object by using either the OC CLI or the Velero CLI

```
[root@localhost ~]# velero get restore -n openshift-adp
NAME       BACKUP                    STATUS      STARTED                         COMPLETED                       ERRORS   WARNINGS   CREATED                         SELECTOR
restore    backup-postgresql-ontaps3   Completed   2024-07-16 14:59:22 -0400 EDT   2024-07-16 14:59:45 -0400 EDT   0        10         2024-07-16 14:59:22 -0400 EDT   <none>
restore1   backup-postgresql-ontaps3   Completed   2024-07-16 16:36:37 -0400 EDT   2024-07-16 16:36:59 -0400 EDT   0        9          2024-07-16 16:36:37 -0400 EDT   <none>
[root@localhost ~]# velero restore delete restore1 -n openshift-adp
Are you sure you want to continue (Y/N)? Y
Request to delete restore "restore1" submitted successfully.
The restore will be fully deleted after all associated data (restore files in object storage) are removed.
[root@localhost ~]# velero get restore -n openshift-adp
NAME       BACKUP                    STATUS      STARTED                         COMPLETED                       ERRORS   WARNINGS   CREATED                         SELECTOR
restore    backup-postgresql-ontaps3   Completed   2024-07-16 14:59:22 -0400 EDT   2024-07-16 14:59:45 -0400 EDT   0        10         2024-07-16 14:59:22 -0400 EDT   <none>
[root@localhost ~]#
[root@localhost ~]# oc delete restore restore -n openshift-adp
restore.velero.io "restore" deleted
[root@localhost ~]# oc get restore -n openshift-adp
No resources found in openshift-adp namespace.
[root@localhost ~]# velero get restore -n openshift-adp
[root@localhost ~]#
```

Activate Windows

Watch videos and demos on hybrid multicloud solutions with NetApp

Watch videos on hybrid multicloud solutions for Amazon Web Services with NetApp

Explore NetApp's comprehensive video library featuring hybrid multicloud solutions for AWS with NetApp. Watch videos on guest connected storage, TCO savings, supplemental datastores, and migration workflows for VMware Cloud on AWS and Amazon FSx for ONTAP, including Windows and Linux scenarios.

Guest connected storage

Windows Guest Connected Storage with FSx ONTAP using iSCSI

Linux Guest Connected Storage with FSx ONTAP using NFS

TCO and supplemental datastores

VMware Cloud on AWS TCO savings with Amazon FSx ONTAP

VMware Cloud on AWS supplemental datastore w/ Amazon FSx ONTAP

Migration workflows (HCX)

VMware HCX Deployment and Configuration Setup for VMC

vMotion Migration Demonstration with VMware HCX for VMC and FSx ONTAP

Cold Migration Demonstration with VMware HCX for VMC and FSx ONTAP

Watch videos on hybrid multicloud solutions for Azure VMware cloud with NetApp

Explore NetApp's comprehensive video library featuring hybrid multicloud solutions for Azure VMware Solution (AVS) with NetApp. This collection includes videos on supplemental datastores, disaster recovery, and migration workflows.

Supplemental datastores

Azure VMware Solution supplemental datastore overview with Azure NetApp Files

Disaster recovery

Azure VMware Solution DR with Cloud Volumes ONTAP, SnapCenter and JetStream

Migration workflows (HCX)

Cold Migration Demonstration with VMware HCX for AVS and ANF

vMotion Demonstration with VMware HCX for AVS and ANF

Bulk Migration Demonstration with VMware HCX for AVS and ANF

Watch videos on hybrid multicloud solutions for Red Hat OpenShift with NetApp

Explore NetApp's comprehensive video library featuring hybrid multicloud solutions for Red Hat OpenShift with NetApp. This collection includes demos on snapshots and restore for OpenShift on AWS (ROSA) with FSx ONTAP, integration with Trident, failover and fail-back for apps, and automated Oracle RAC deployment on FlexPod with Ansible.

Snapshots and restore

Snapshot/Restore for Applications on Red Hat OpenShift Service on AWS (ROSA)clusters with Amazon FSx ONTAP storage

Storage integration

Integration of FSx ONTAP with Trident

Failover and fail-back

Failover and Fail-back of apps on ROSA with FSx ONTAP

Solution automation



Automated Oracle 19c RAC Deployment on FlexPod with Ansible

Read cloud solution blogs from NetApp experts

Read the blogs that offer practical guidance and best practices for integrating ONTAP storage with NetApp cloud solutions.

VMware integration with AWS and Amazon FSx for ONTAP

- Migrate VMware VMs to AWS using Amazon EC2 and Amazon FSx ONTAP iSCSI protocol
- Amazon EVS now integrates with FSx for ONTAP
- Amazon EVS announces support for FSx for ONTAP
- Configure Hybrid Cloud with FSx ONTAP and VMware Cloud on AWS SDDC using VMware HCX

Azure VMware solution

- Optimize storage and data-intensive VMware workloads in any cloud
- VMware workloads to Azure VMware Solution? Read this first.

Google Cloud VMware Engine

- Google Cloud NetApp Volumes datastore support for Google Cloud VMware Engine (NetApp)
- How to use NetApp Volumes as datastores for Google Cloud VMware Engine (Google)
- NetApp Blog - A better way to migrate your storage-rich apps to Google Cloud

Data protection and management

- 3-2-1 Data Protection for VMware with SnapCenter Plug-in and NetApp Backup and Recovery for VMs
- Observability for the modern datacenter with NetApp Cloud Insights

NetApp for OpenShift container workloads

- NetApp ONTAP doubles as storage for logs in Red Hat OpenShift container platform

VMware Sovereign Cloud

- Introducing the new Sovereign Cloud initiative

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

https://www.netapp.com/company/legal/copyright/

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

https://www.netapp.com/company/legal/trademarks/

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.