Virtualization

NetApp Solutions

NetApp
July 31, 2024

This PDF was generated from https://docs.netapp.com/us-en/netapp-solutions/virtualization/virt-overview.html on July 31, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Virtualization . . . 1
A Quick Introduction to Virtualization . . . 1
NetApp Solutions for Virtualization with VMware by Broadcom . . . 3
NetApp Hyper-V Virtualization Solutions . . . 288
NetApp OpenShift Virtualization Solutions . . . 319

Virtualization

A Quick Introduction to Virtualization

The prevalence of virtualization has driven business outcomes for decades. From the early days of mainframes to the present era of storage-intensive use cases and the needs for AI-ready systems, virtualization has become an integral component of every customer environment.

Overview

Traditionally, datacenters consisted of physical servers, networking and storage, each dedicated to specific tasks. This approach created silos of resources that were often greatly underutilized. At the core, virtualization involves decoupling physical resources from the requirements and functions of the organization. This is achieved by creating virtual representations of the three main infrastructure components: **network**, **compute**, and **storage**. The deployment of virtual infrastructure enables organizations to achieve increased utilization of the underlying physical resources through these logical constructs.

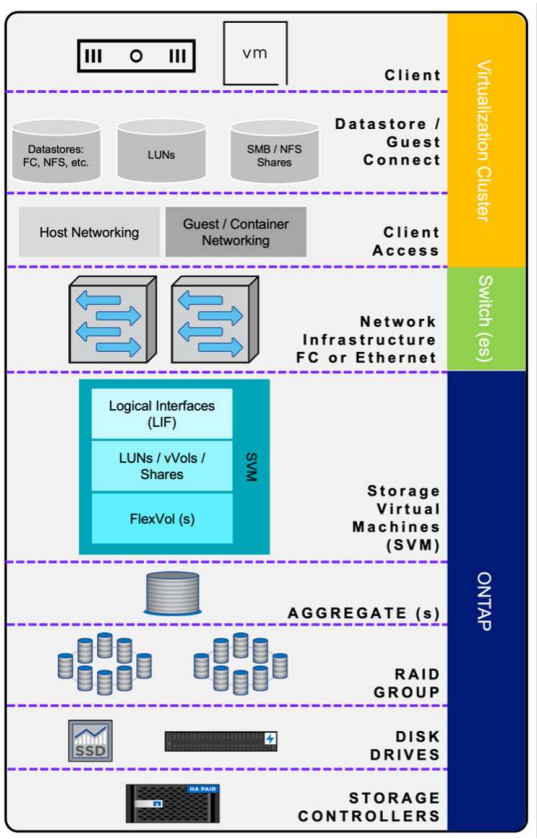

Virtualization is a case where the whole is greater than the sum of the parts through automation, policy-driven workflows and scalability. By improving efficiency, flexibility and while lowering TCO of IT Infrastructure, virtualization empowers organizations to optimize their technological resources to a greater extent. A virtual infrastructure powered by NetApp inherits key benefits of ONTAP:

- Provisioning, backup, and data protection with industry leading plug-ins and integrations
- Enterprise grade storage efficiencies

- Multiprotocol support for varying workload needs
- Policy-driven storage configuration and adaptive QoS
- Support for applications and workloads whether on-premises or the public clouds with one storage operating system

To learn more about NetApp ONTAP, see the following:

- ONTAP: The Leading Unified Storage Operating System

Network virtualization

Network virtualization refers to the abstraction of the physical network components (switches, routers, firewalls, interfaces, etc.) into logical constructs. These logical constructs, operating independently of the underlying physical infrastructure, deliver secure communication between clients, storage and other components in the network. Network virtualization is key to running the business as it allows for the sharing of resources while simultaneously allowing for the restriction of network traffic according to policies.

Network virtualization can combine multiple physical networks into one virtual fabric, or it can facilitate dividing a physical network into separate, discreet virtual networks. Multiples networks can be created and customized to meet specific IT requirements. Often network virtualization refers to Ethernet-based use cases but in many instances virtual fabrics can be configured dependent upon the switch manufacture's capabilities. Whether employing virtual LANs or virtual SANs, organizations achieve greater operational efficiency and overall improved network performance through network virtualization.

For more information on networking virtualization, see the following:

- Cisco virtual networking
- VMware virtual networking

Compute virtualization

Compute or server virtualization is perhaps the most well-known form of virtualization. With compute virtualization, hypervisors mimic the functions of physical servers allowing, operation teams to run multiple virtual machines on a single physical node. With compute virtualization, resources such as sever memory and CPU are shared. This sharing allows for an oversubscription of the underlying resources to the degree that is acceptable for the workloads and applications deployed.

With compute virtualization, each virtual machine has its own operating system and installed applications and resources; functioning independently of each other. Among the numerous advantages with compute virtualization include increased server utilization, reduced hardware expenditures, simplified management using the hypervisor's user interface (UI), and improved disaster recovery functionality. Additionally, with hypervisor plug-ins, storage administration, backups and protection relationships can be configured to further simplify operational tasks.

For more information on compute virtualization, see the following:

- VMware vSphere
- Red Hat OpenShift Virtualization
- Microsoft Hyper-V

Storage virtualization

Much like the network and compute virtualization, storage virtualization is important to a modern datacenter. NetApp ONTAP facilitates storage virtualization through Storage Virtual Machines (SVMs) which serve data to clients and hosts. SVMs are logical entities that allow for storage resources to not be tied to physical media. SVMs can be deployed based on workload type, application needs and organization groups for access.

There are multiple types of SVMs which aide in data access, administration and cluster and system level tasks. Data SVMs serve data to clients and hosts from one or more volumes, through one or more network logical interfaces (LIFs). These volumes and LIFs are logical constructs and are mapped through to storage aggregates and physical or logical network ports. This logical data access allows for the mobility of volumes or LIFs during maintenance scenarios or resource rebalancing much like compute virtual machines.

For more information on storage virtualization, see the following:

- ONTAP Storage Virtualization Overview

Closing

The components of virtual infrastructure described here: network, compute and storage provide same the same functionality as typical physical resources but through software. The allocation of virtual resources over physical resources accelerates time to value and allows for policy-driven configuration of resources. Pairing ONTAP with compute and network virtualization allows for clients and hosts to access resources through a software-defined virtual infrastructure.

NetApp Solutions for Virtualization with VMware by Broadcom

VMware Cloud Foundation

VMware Cloud Foundation (VCF) is an integrated software defined data center (SDDC) platform that provides a complete stack of software-defined infrastructure for running enterprise applications in a hybrid cloud environment. It combines compute, storage, networking, and management capabilities into a unified platform, offering a consistent operational experience across private and public clouds.

Author: Josh Powell

VMware Cloud Foundation with NetApp All-Flash SAN Arrays

This document provides information on storage options available for VMware Cloud Foundation using the NetApp All-Flash SAN Array. Supported storage options are covered with specific instruction for deploying iSCSI datastores as supplemental storage for management domains and both vVol (iSCSI) and NVMe/TCP datastores as supplemental datastores for workload domains. Also covered is data protection of VMs and datastores using SnapCenter for VMware vSphere.

Use Cases

Use cases covered in this documentation:

- Storage options for customers seeking uniform environments across both private and public clouds.

- Automated solution for deploying virtual infrastructure for workload domains.
- Scalable storage solution tailored to meet evolving needs, even when not aligned directly with compute resource requirements.
- Deploy supplemental storage to management and VI workload domains using ONTAP Tools for VMware vSphere.
- Protect VMs and datastores using the SnapCenter Plug-in for VMware vSphere.

Audience

This solution is intended for the following people:

- Solution architects looking for more flexible storage options for VMware environments that are designed to maximize TCO.
- Solution architects looking for VCF storage options that provide data protection and disaster recovery options with the major cloud providers.
- Storage administrators wanting specific instruction on how to configure VCF with principal and supplemental storage.
- Storage administrators wanting specific instruction on how to protect VMs and datastores residing on ONTAP storage.

Technology Overview

The VCF with NetApp ASA solution is comprised of the following major components:

VMware Cloud Foundation

VMware Cloud Foundation extends VMware's vSphere hypervisor offerings by combining key components such as SDDC Manager, vSphere, vSAN, NSX, and VMware Aria Suite to create a software-defined datacenter.

The VCF solution supports both native Kubernetes and virtual machine-based workloads. Key services such as VMware vSphere, VMware vSAN, VMware NSX-T Data Center, and VMware Aria Cloud Management are integral components of the VCF package. When combined, these services establish a software-defined infrastructure capable of efficiently managing compute, storage, networking, security, and cloud management.

VCF is comprised of a single management domain and up to 24 VI workload domains that each represent a unit of application-ready infrastructure. A workload domain is comprised of one or more vSphere clusters managed by a single vCenter instance.

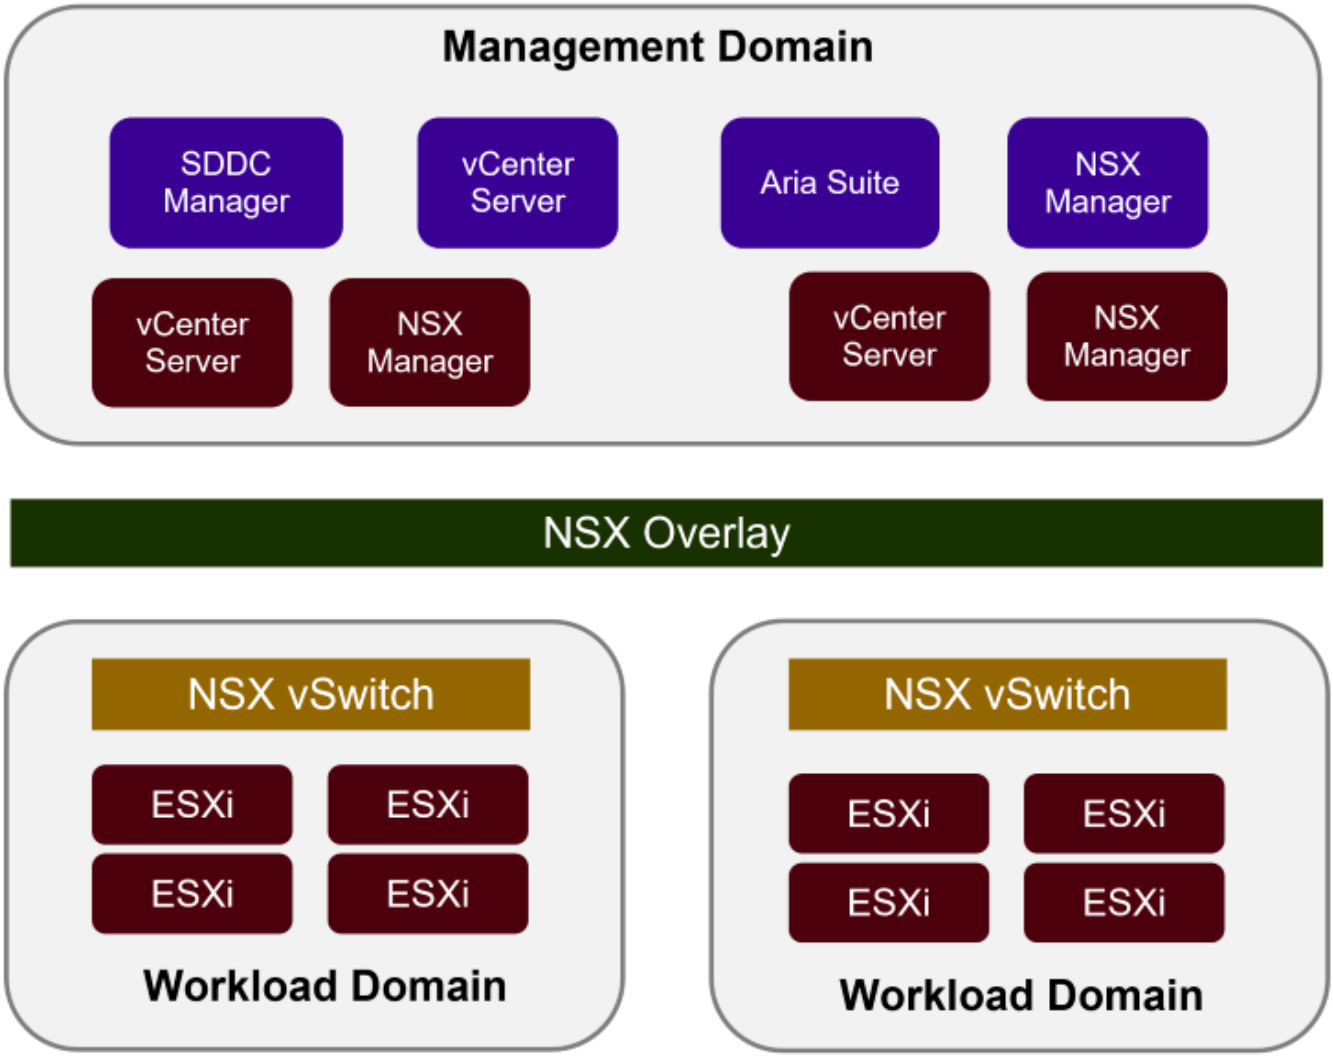

For more information on VCF architecture and planning, refer to Architecture Models and Workload Domain Types in VMware Cloud Foundation.

VCF Storage Options

VMware divides storage options for VCF into **principal** and **supplemental** storage. The VCF management domain must use vSAN as its principal storage. However, there are many supplemental storage options for the management domain and both principal and supplemental storage options available for VI workload domains.

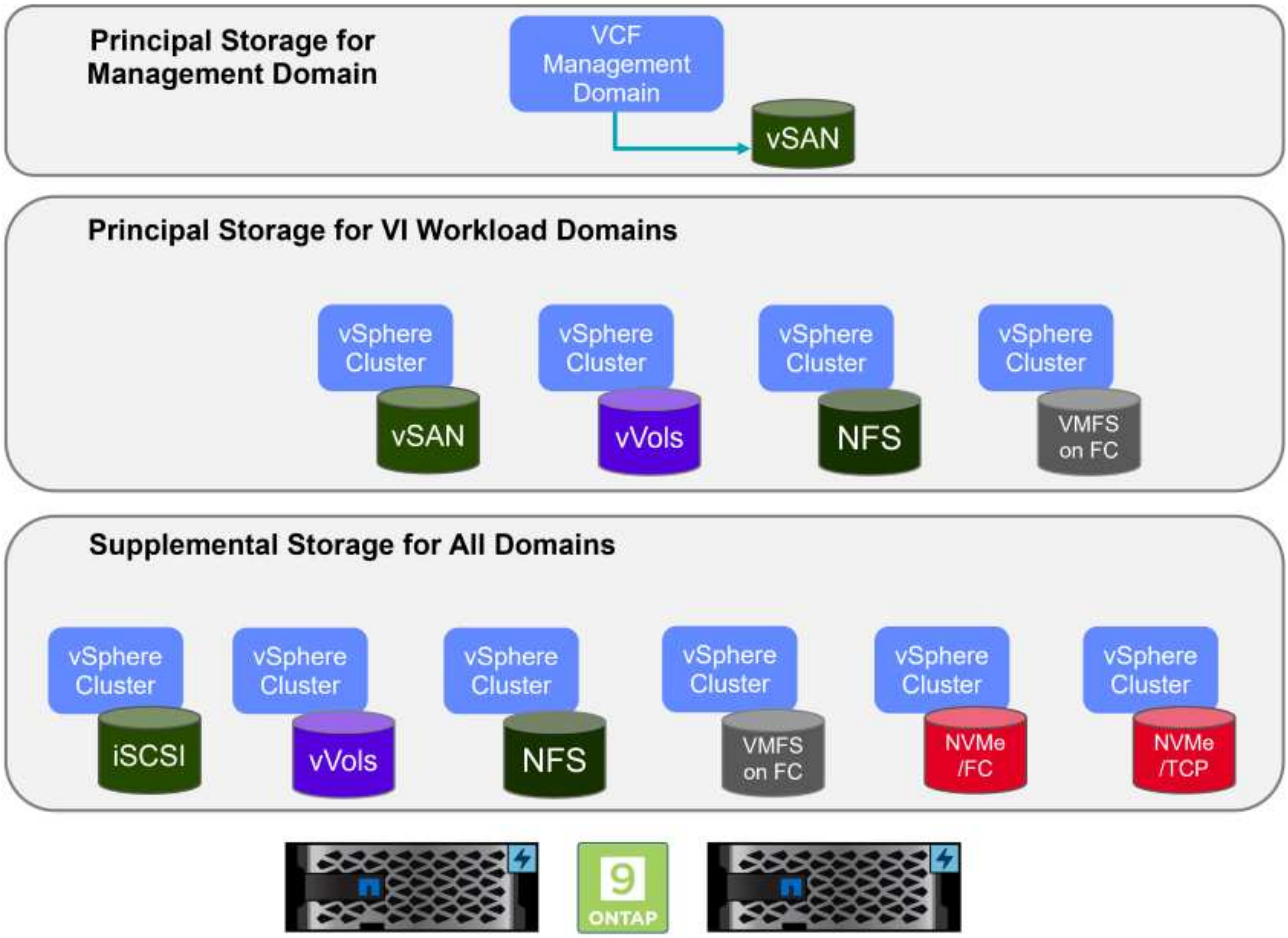

Principal Storage for Workload Domains
Principal storage refers to any type of storage that can be directly connected to a VI workload domain during the setup process within SDDC Manager. Principal storage is deployed with SDDC manager as part of cluster creation orchestration and is the first datastore configured for a workload domain. It includes vSAN, vVols (VMFS), NFS and VMFS on Fibre Channel.

Supplemental Storage for Management and Workload Domains
Supplemental storage is the storage type that can be added to the management or workload domains at any time after the cluster has been created. Supplemental storage represents the widest range of supported storage options, all of which are supported on NetApp ASA arrays. Supplemental storage can be deployed using ONTAP Tools for VMware vSphere for most storage protocol types.

Additional documentation resources for VMware Cloud Foundation:
* VMware Cloud Foundation Documentation
* Supported Storage Types for VMware Cloud Foundation
* Managing Storage in VMware Cloud Foundation

NetApp All-Flash SAN Arrays

The NetApp All-Flash SAN Array (ASA) is a high-performance storage solution designed to meet the demanding requirements of modern data centers. It combines the speed and reliability of flash storage with NetApp's advanced data management features to deliver exceptional performance, scalability, and data protection.

The ASA lineup is comprised of both A-Series and C-Series models.

The NetApp A-Series all-NVMe flash arrays are designed for high-performance workloads, offering ultra-low latency and high resiliency, making them suitable for mission-critical applications.

C-Series QLC flash arrays are aimed at higher-capacity use cases, delivering the speed of flash with the economy of hybrid flash.

For detailed information see the NetApp ASA landing page.

Storage Protocol Support

The ASA supports all standard SAN protocols including, iSCSI, Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and NVME over fabrics.

iSCSI - NetApp ASA provides robust support for iSCSI, allowing block-level access to storage devices over IP networks. It offers seamless integration with iSCSI initiators, enabling efficient provisioning and management of iSCSI LUNs. ONTAP's advanced features, such as multi-pathing, CHAP authentication, and ALUA support.

For design guidance on iSCSI configurations refer to the SAN Configuration reference documentation.

Fibre Channel - NetApp ASA offers comprehensive support for Fibre Channel (FC), a high-speed network technology commonly used in storage area networks (SANs). ONTAP seamlessly integrates with FC infrastructure, providing reliable and efficient block-level access to storage devices. It offers features like zoning, multi-pathing, and fabric login (FLOGI) to optimize performance, enhance security, and ensure seamless connectivity in FC environments.

For design guidance on Fibre Channel configurations refer to the SAN Configuration reference documentation.

NVMe over Fabrics - NetApp ONTAP and ASA support NVMe over fabrics. NVMe/FC enables the use of

NVMe storage devices over Fibre Channel infrastructure, and NVMe/TCP over storage IP networks.

For design guidance on NVMe refer to NVMe configuration, support and limitations

Active-active technology

NetApp All-Flash SAN Arrays allows for active-active paths through both controllers, eliminating the need for the host operating system to wait for an active path to fail before activating the alternative path. This means that the host can utilize all available paths on all controllers, ensuring active paths are always present regardless of whether the system is in a steady state or undergoing a controller failover operation.

Furthermore, the NetApp ASA offers a distinctive feature that greatly enhances the speed of SAN failover. Each controller continuously replicates essential LUN metadata to its partner. As a result, each controller is prepared to take over data serving responsibilities in the event of a sudden failure of its partner. This readiness is possible because the controller already possesses the necessary information to start utilizing the drives that were previously managed by the failed controller.

With active-active pathing, both planned and unplanned takeovers have IO resumption times of 2-3 seconds.

For more information see TR-4968, NetApp All-SAS Array – Data Availability and Integrity with the NetApp ASA.

Storage guarantees

NetApp offers a unique set of storage guarantees with NetApp All-flash SAN Arrays. The unique benefits include:

Storage efficiency guarantee: Achieve high performance while minimizing storage cost with the Storage Efficiency Guarantee. 4:1 for SAN workloads.

6 Nines (99.9999%) data availability guarantee: Guarantees remediation for unplanned downtime in excess of 31.56 seconds per year.

Ransomware recovery guarantee: Guaranteed data recovery in the event of a ransomware attack.

See the NetApp ASA product portal for more information.

NetApp ONTAP Tools for VMware vSphere

ONTAP Tools for VMware vSphere allows administrators to manage NetApp storage directly from within the vSphere Client. ONTAP Tools allows you to deploy and manage datastores, as well as provision vVol datastores.

ONTAP Tools allows mapping of datastores to storage capability profiles which determine a set of storage system attributes. This allows the creation of datastores with specific attributes such as storage performance and QoS.

ONTAP Tools also includes a **VMware vSphere APIs for Storage Awareness (VASA) Provider** for ONTAP storage systems, which enables the provisioning of VMware Virtual Volumes (vVols) datastores, creation and use of storage capability profiles, compliance verification, and performance monitoring.

For more information on NetApp ONTAP tools see the ONTAP tools for VMware vSphere Documentation page.

SnapCenter Plug-in for VMware vSphere

The SnapCenter Plug-in for VMware vSphere (SCV) is a software solution from NetApp that offers comprehensive data protection for VMware vSphere environments. It is designed to simplify and streamline the process of protecting and managing virtual machines (VMs) and datastores. SCV uses storage based snapshot and replication to secondary arrays to meet lower recovery time objectives.

The SnapCenter Plug-in for VMware vSphere provides the following capabilities in a unified interface, integrated with the vSphere client:

Policy-Based Snapshots - SnapCenter allows you to define policies for creating and managing application-consistent snapshots of virtual machines (VMs) in VMware vSphere.

Automation - Automated snapshot creation and management based on defined policies help ensure consistent and efficient data protection.

VM-Level Protection - Granular protection at the VM level allows for efficient management and recovery of individual virtual machines.

Storage Efficiency Features - Integration with NetApp storage technologies provides storage efficiency features like deduplication and compression for snapshots, minimizing storage requirements.

The SnapCenter Plug-in orchestrates the quiescing of virtual machines in conjunction with hardware-based snapshots on NetApp storage arrays. SnapMirror technology is utilized to replicate copies of backups to secondary storage systems including in the cloud.

For more information refer to the SnapCenter Plug-in for VMware vSphere documentation.

BlueXP integration enables 3-2-1 backup strategies that extend copies of data to object storage in the cloud.

For more information on 3-2-1 backup strategies with BlueXP visit 3-2-1 Data Protection for VMware with SnapCenter Plug-in and BlueXP backup and recovery for VMs.

Solution Overview

The scenarios presented in this documentation will demonstrate how to use ONTAP storage systems as supplemental storage for management and workload domains. In addition, the SnapCenter Plug-in for VMware vSphere is used to protect VMs and datastores.

Scenarios covered in this documentation:

- **Use Ontap Tools to deploy iSCSI datastores in a VCF management domain**. Click **here** for deployment steps.
- **Use Ontap Tools to deploy vVols (iSCSI) datastores in a VI workload domain**. Click **here** for deployment steps.
- **Configure NVMe over TCP datastores for use in a VI workload domain**. Click **here** for deployment steps.
- **Deploy and use the SnapCenter Plug-in for VMware vSphere to protect and restore VMs in a VI workload domain**. Click **here** for deployment steps.

In this scenario we will demonstrate how to deploy and use ONTAP Tools for VMware vSphere (OTV) to configure an iSCSI datastore for a VCF management domain.

Author: Josh Powell

Use ONTAP Tools to configure supplemental storage for VCF Management Domains

Scenario Overview

This scenario covers the following high level steps:

- Create a storage virtual machine (SVM) with logical interfaces (LIFs) for iSCSI traffic.
- Create distributed port groups for iSCSI networks on the VCF management domain.
- Create vmkernel adapters for iSCSI on the ESXi hosts for the VCF management domain.
- Deploy ONTAP Tools on the VCF management domain.
- Create a new VMFS datastore on the VCF management domain.

Prerequisites

This scenario requires the following components and configurations:

- An ONTAP ASA storage system with physical data ports on ethernet switches dedicated to storage traffic.
- VCF management domain deployment is complete and the vSphere client is accessible.

NetApp recommends fully redundant network designs for iSCSI. The following diagram illustrates an example of a redundant configuration, providing fault tolerance for storage systems, switches, networks adapters and host systems. Refer to the NetApp SAN configuration reference for additional information.

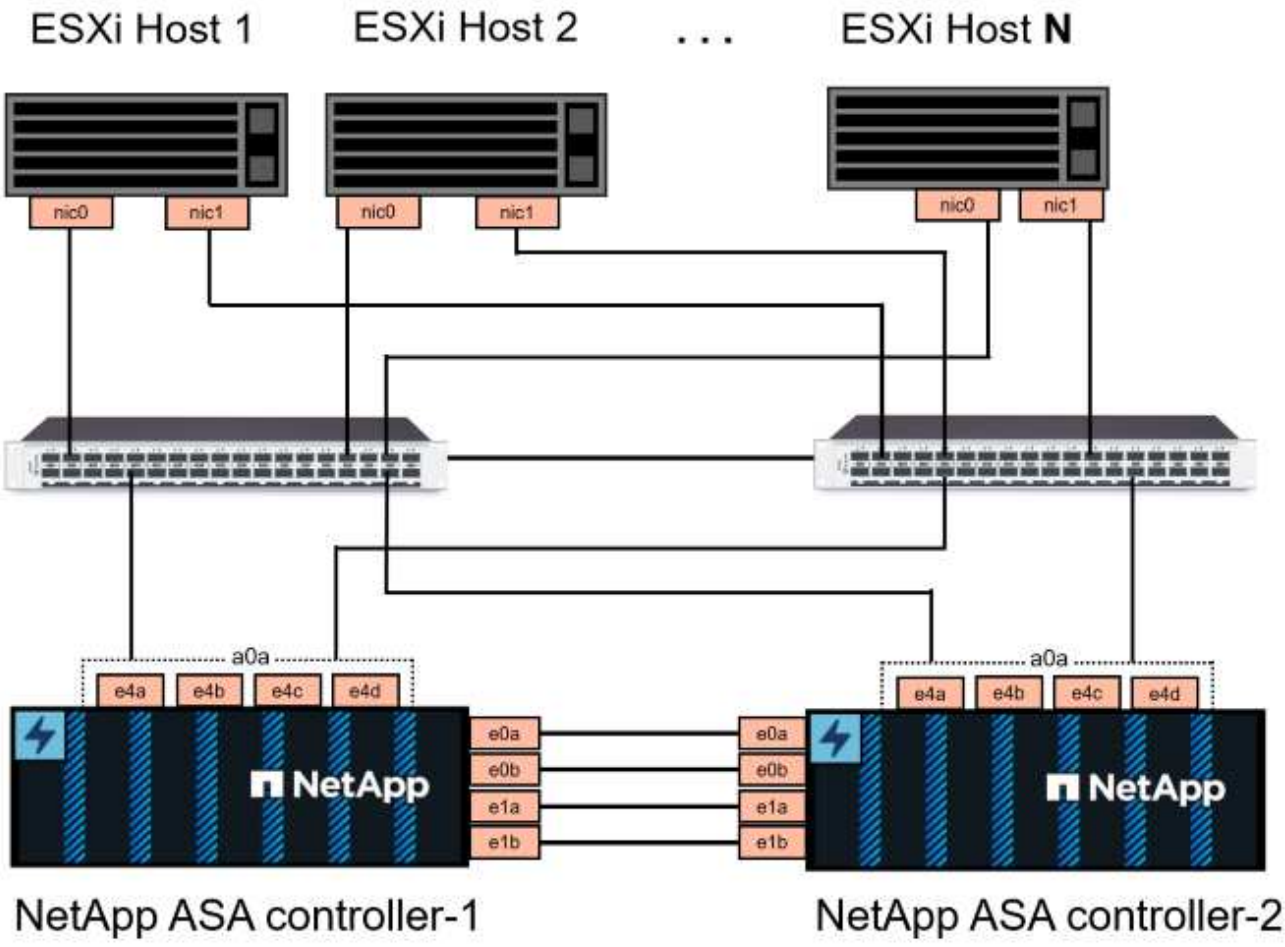

For multipathing and failover across multiple paths, NetApp recommends having a minimum of two LIFs per storage node in separate ethernet networks for all SVMs in iSCSI configurations.

This documentation demonstrates the process of creating a new SVM and specifying the IP address information to create multiple LIFs for iSCSI traffic. To add new LIFs to an existing SVM refer to Create a LIF (network interface).

For additional information on using VMFS iSCSI datastores with VMware refer to vSphere VMFS Datastore - iSCSI Storage backend with ONTAP.

In situations where multiple VMkernel adapters are configured on the same IP network, it is recommended to use software iSCSI port binding on the ESXi hosts to ensure that load balancing across the adapters occurs. Refer to KB article Considerations for using software iSCSI port binding in ESX/ESXi (2038869).

Deployment Steps

To deploy ONTAP Tools and use it to create a VMFS datastore on the VCF management domain, complete the following steps:

Create SVM and LIFs on ONTAP storage system

The following step is is performed in ONTAP System Manager.

Create the storage VM and LIFs

Complete the following steps to create an SVM together with multiple LIFs for iSCSI traffic.

1. From ONTAP System Manager navigate to **Storage VMs** in the left-hand menu and click on **+ Add** to start.

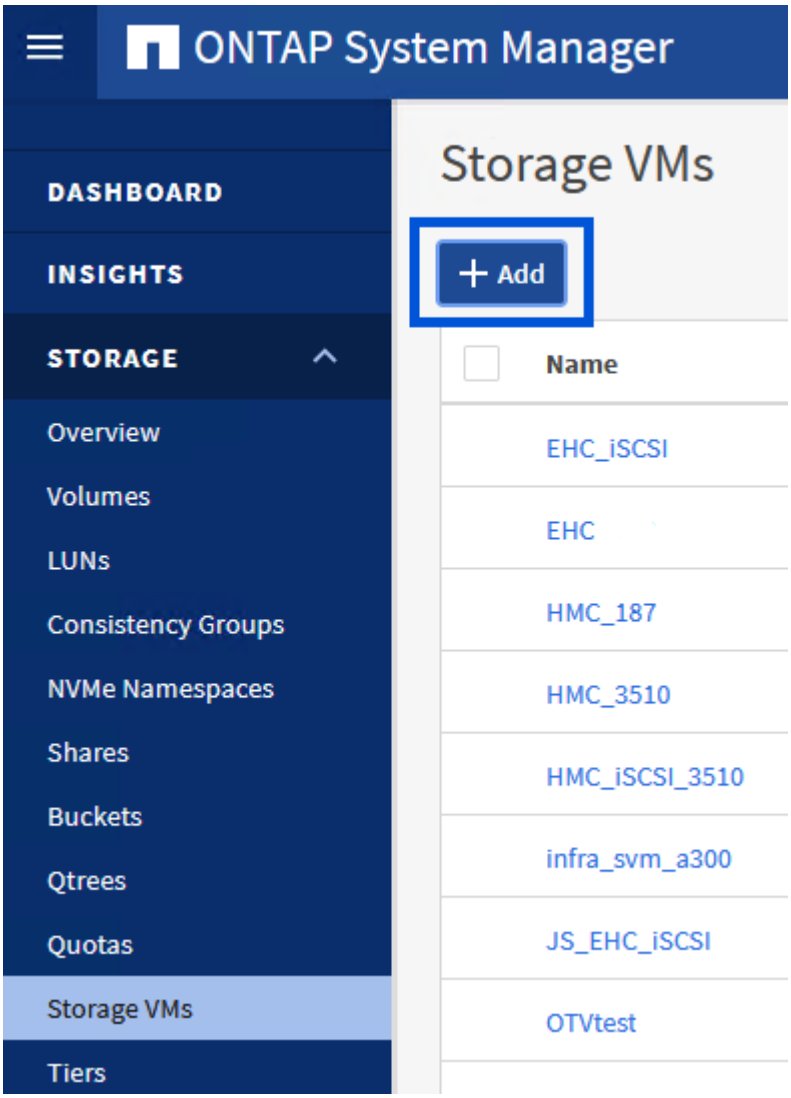

2. In the **Add Storage VM** wizard provide a **Name** for the SVM, select the **IP Space** and then, under **Access Protocol, click on the *iSCSI** tab and check the box to **Enable iSCSI**.

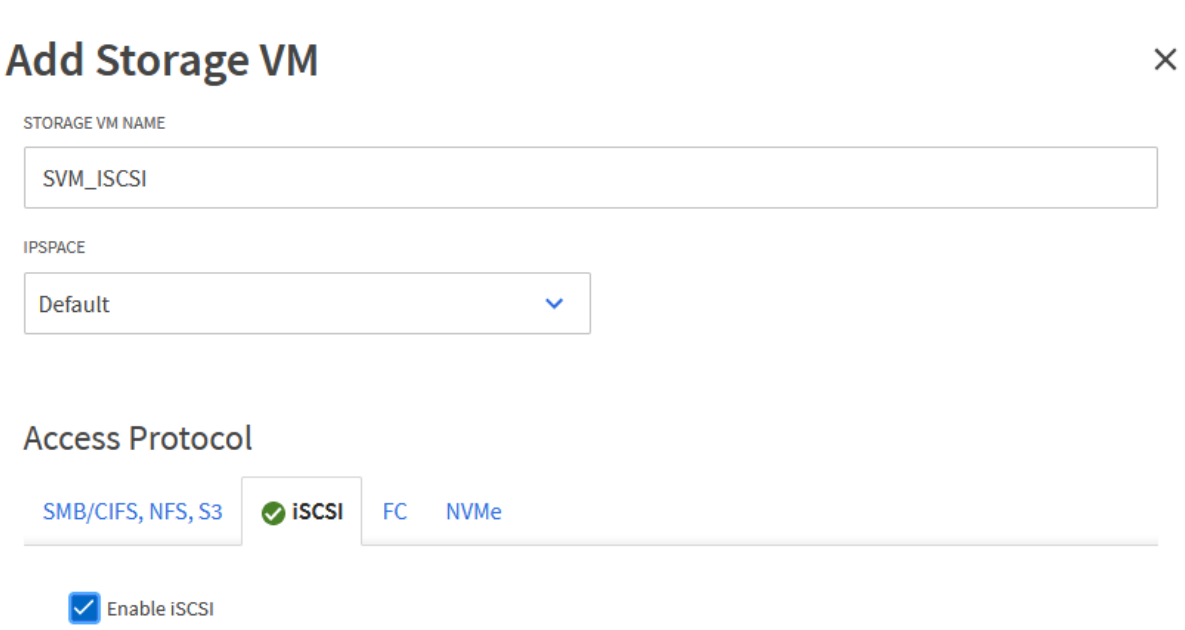

3. In the **Network Interface** section fill in the **IP address**, **Subnet Mask**, and **Broadcast Domain and Port** for the first LIF. For subsequent LIFs the checkbox may be enabled to use common settings across all remaining LIFs or use separate settings.

> For multipathing and failover across multiple paths, NetApp recommends having a minimum of two LIFs per storage node in separate Ethernet networks for all SVMs in iSCSI configurations.

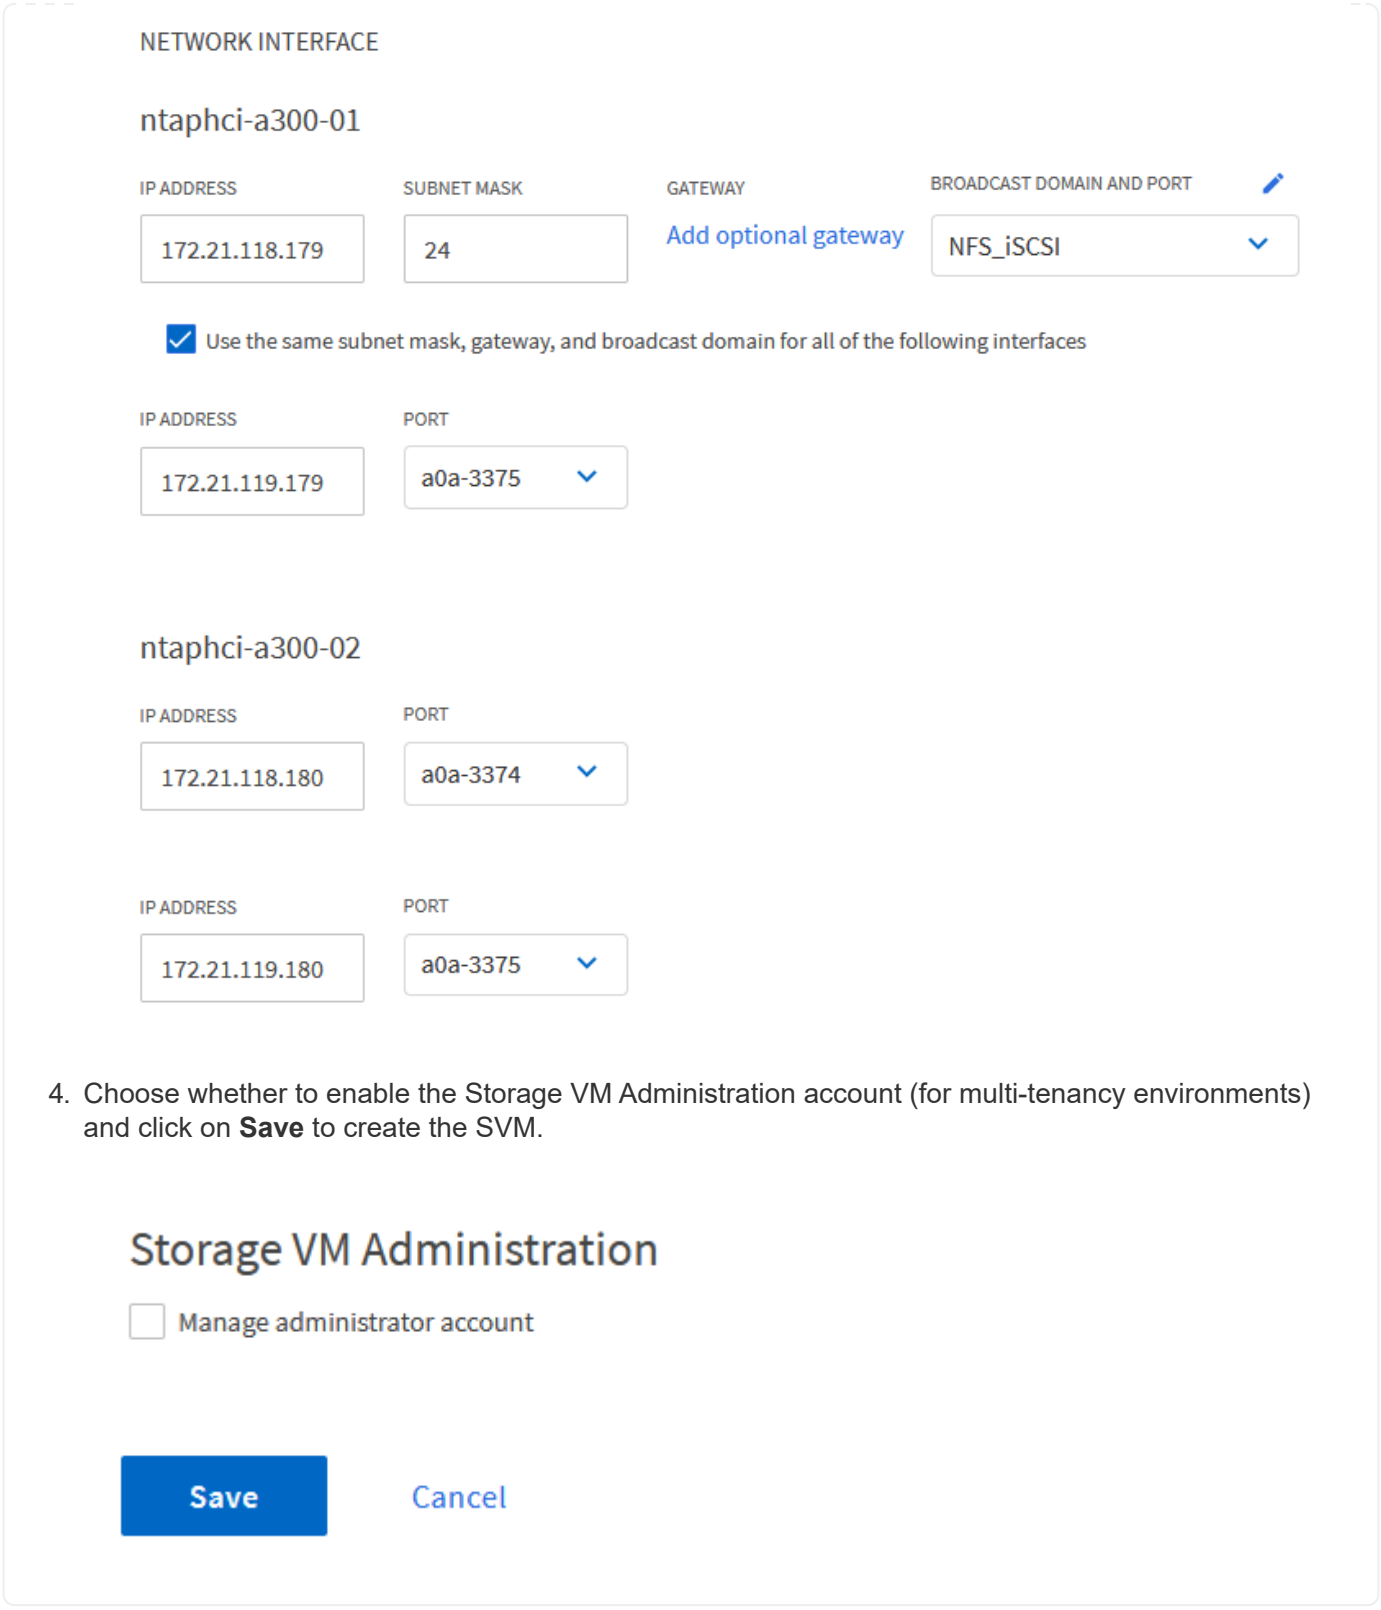

4. Choose whether to enable the Storage VM Administration account (for multi-tenancy environments) and click on **Save** to create the SVM.

Set up networking for iSCSI on ESXi hosts

The following steps are performed on the VCF management domain cluster using the vSphere client.

Create Distributed Port Groups for iSCSI traffic

Complete the following to create a new distributed port group for each iSCSI network:

1. From the vSphere client for the management domain cluster, navigate to **Inventory > Networking**. Navigate to the existing Distributed Switch and choose the action to create **New Distributed Port Group…**.

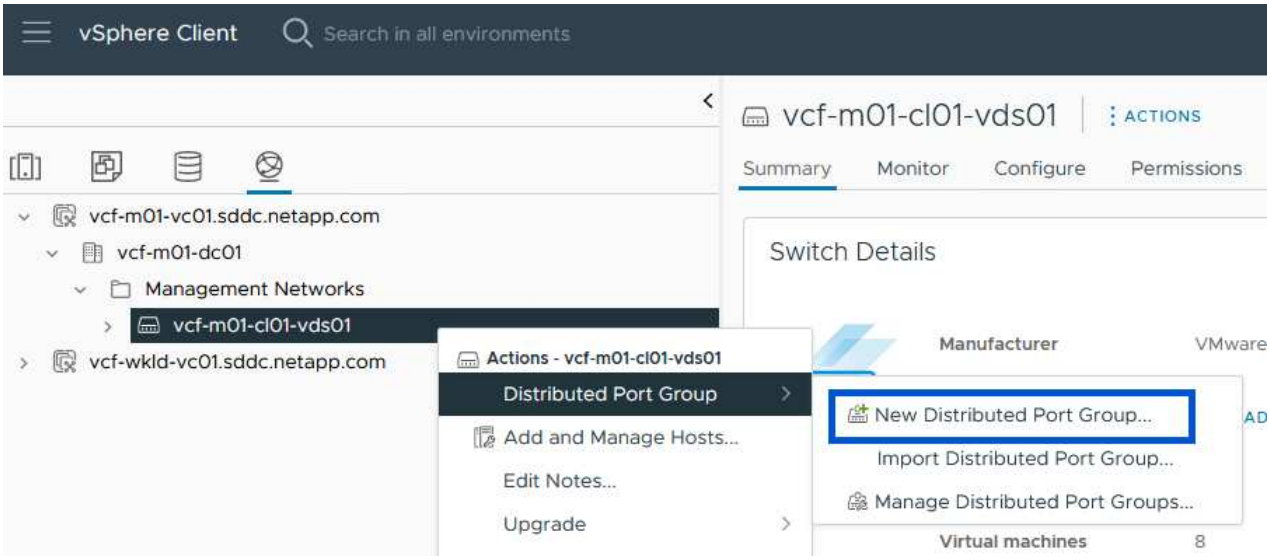

2. In the **New Distributed Port Group** wizard fill in a name for the new port group and click on **Next** to continue.
3. On the **Configure settings** page fill out all settings. If VLANs are being used be sure to provide the correct VLAN ID. Click on **Next** to continue.

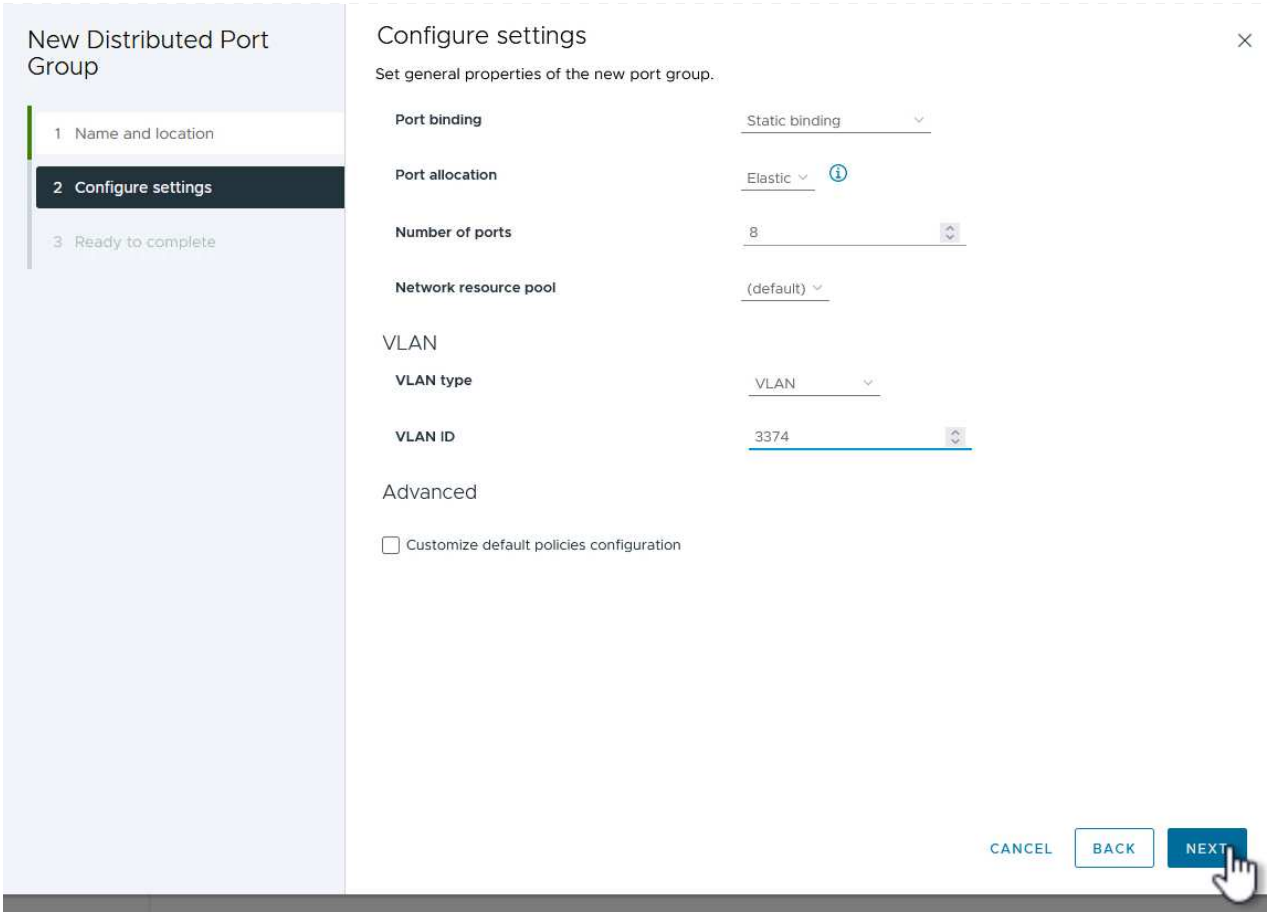

4. On the **Ready to complete** page, review the changes and click on **Finish** to create the new distributed port group.
5. Repeat this process to create a distributed port group for the second iSCSI network being used and ensure you have input the correct **VLAN ID**.
6. Once both port groups have been created, navigate to the first port group and select the action to **Edit settings…**.

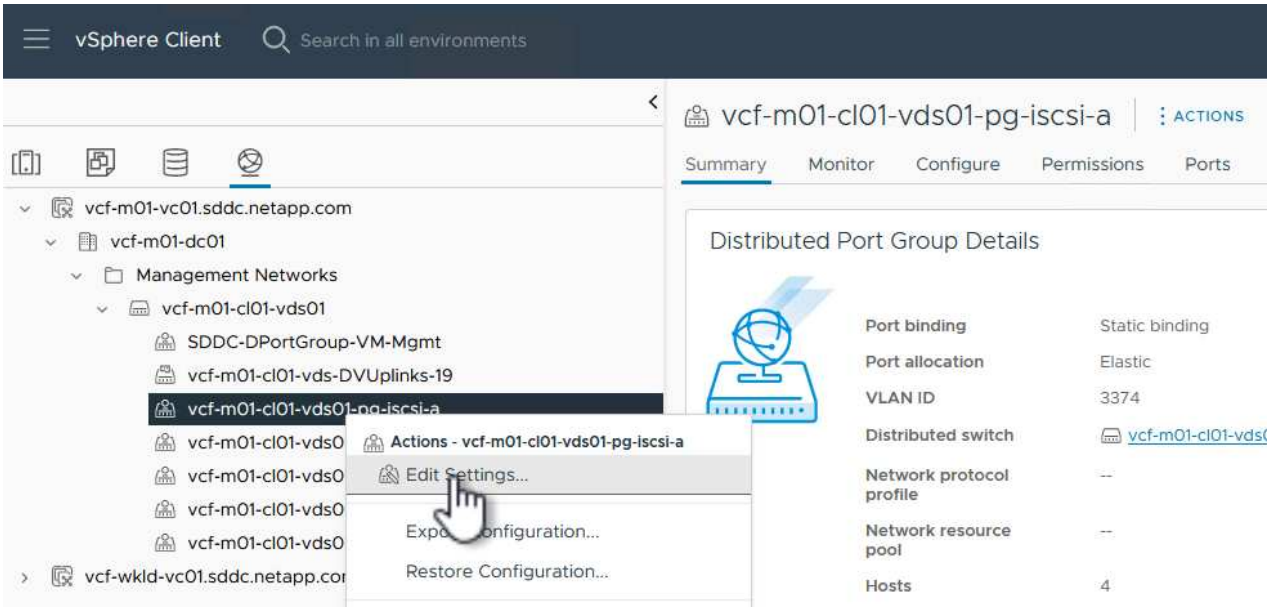

7. On **Distributed Port Group - Edit Settings** page, navigate to **Teaming and failover** in the left-hand menu and click on **uplink2** to move it down to **Unused uplinks**.

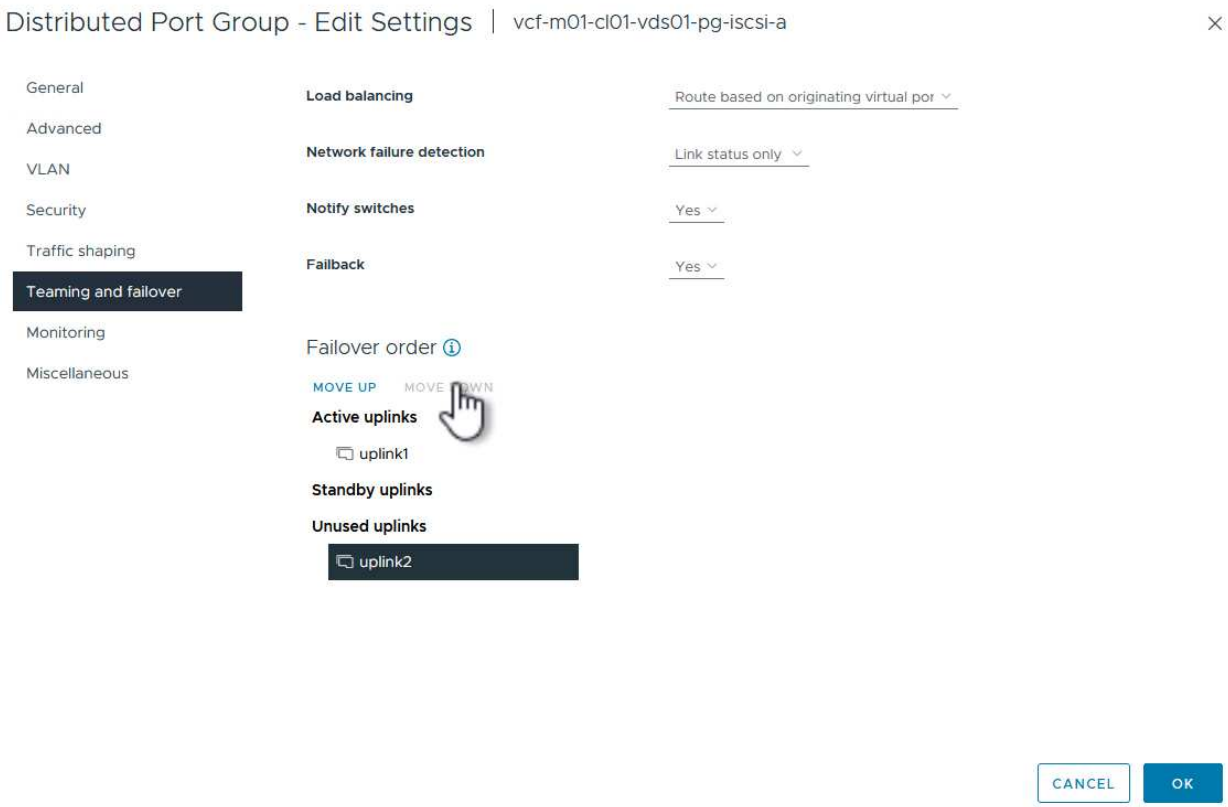

8. Repeat this step for the second iSCSI port group. However, this time move **uplink1** down to **Unused uplinks**.

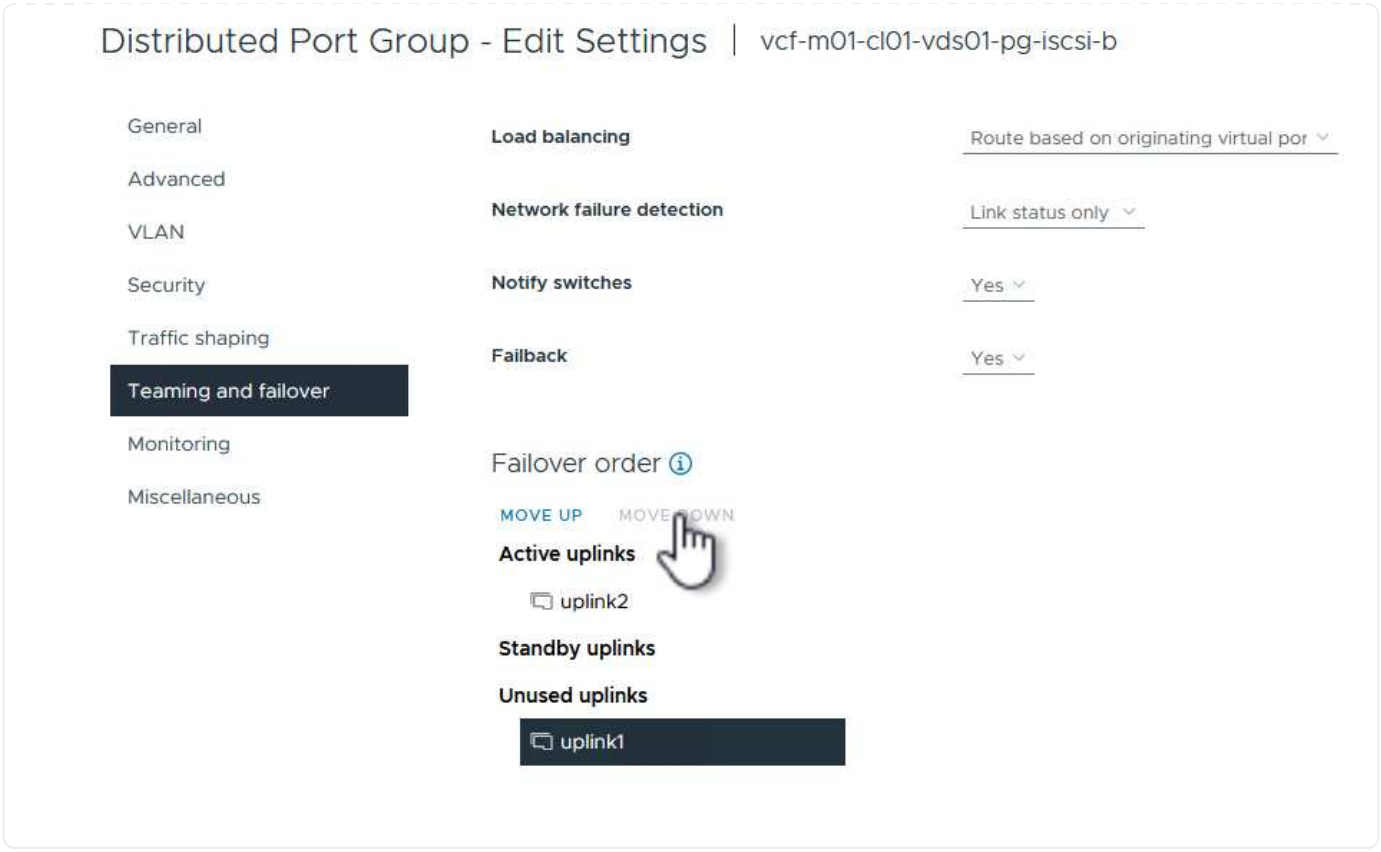
Distributed Port Group - Edit Settings | vcf-m01-cl01-vds01-pg-iscsi-b
General
Advanced
VLAN
Security
Traffic shaping
Teaming and failover
Monitoring
Miscellaneous
Load balancing
Route based on originating virtual por
Network failure detection
Link status only
Notify switches
Yes
Failback
Yes
Failover order
MOVE UP
MOVE DOWN
Active uplinks
uplink2
Standby uplinks
Unused uplinks
uplink1

Create VMkernel adapters on each ESXi host

Repeat this process on each ESXi host in the management domain.

1. From the vSphere client navigate to one of the ESXi hosts in the management domain inventory. From the **Configure** tab select **VMkernel adapters** and click on **Add Networking…** to start.

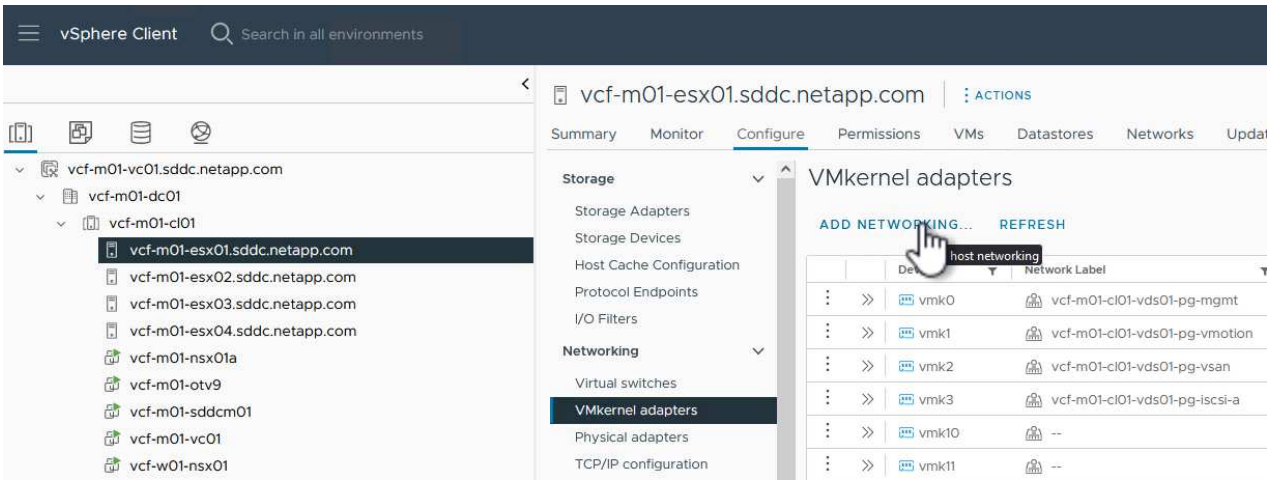

2. On the **Select connection type** window choose **VMkernel Network Adapter** and click on **Next** to continue.

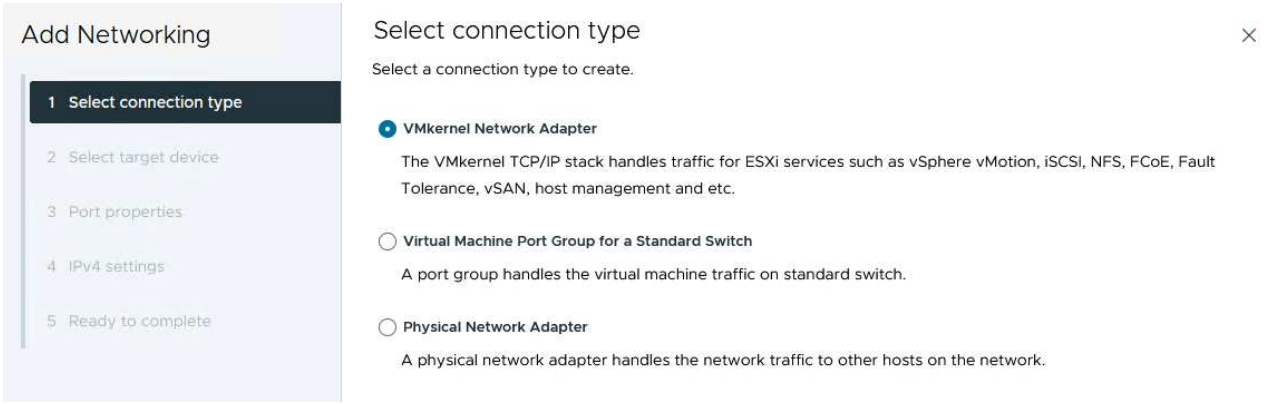

3. On the **Select target device** page, choose one of the distributed port groups for iSCSI that was created previously.

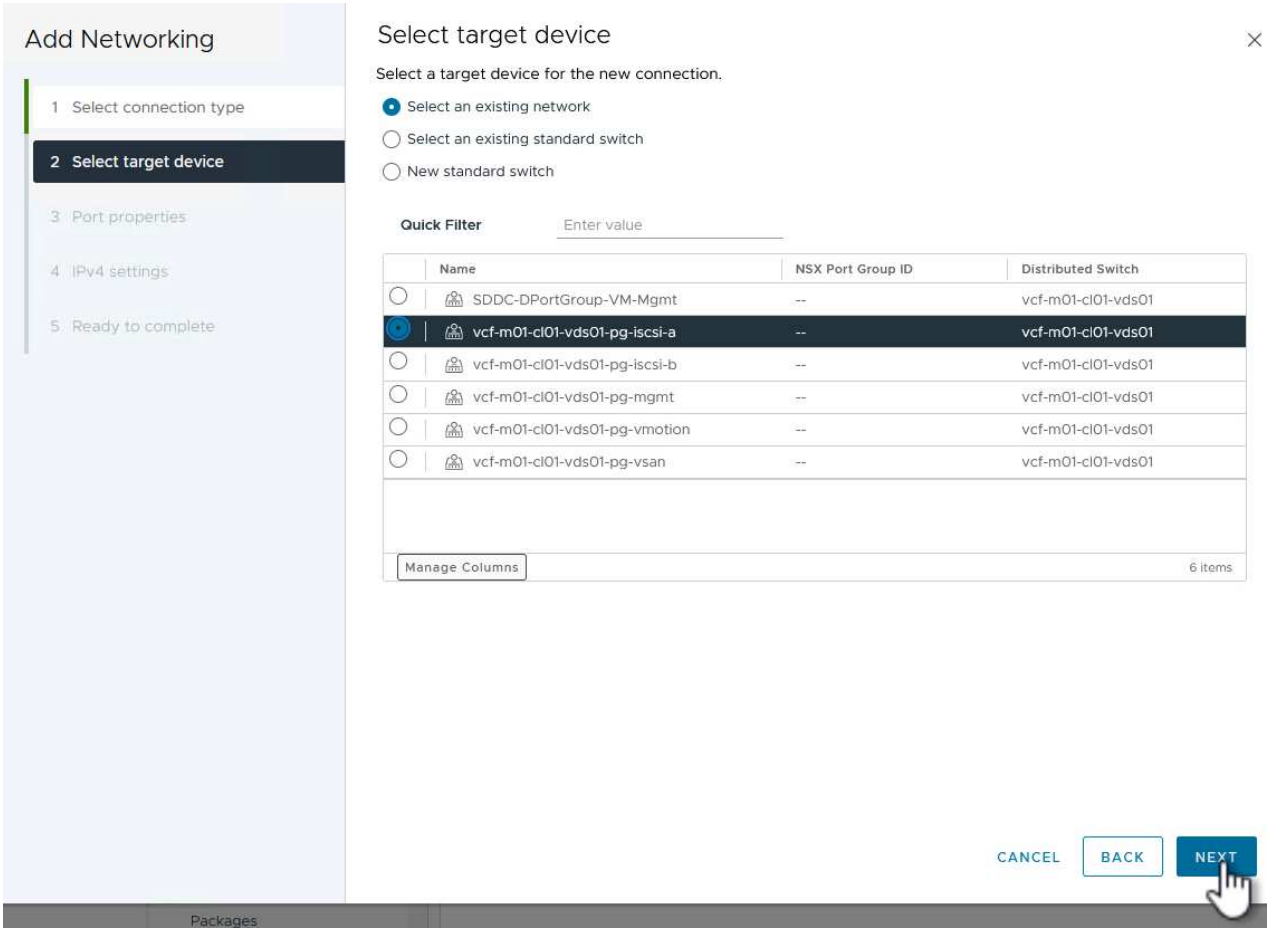

4. On the **Port properties** page keep the defaults and click on **Next** to continue.

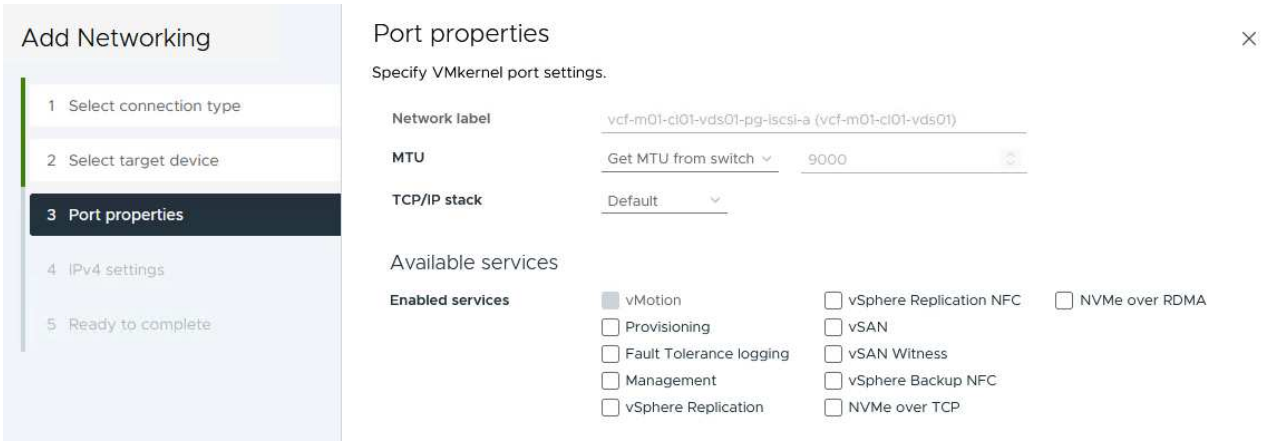

5. On the **IPv4 settings** page fill in the **IP address**, **Subnet mask**, and provide a new Gateway IP address (only if required). Click on **Next** to continue.

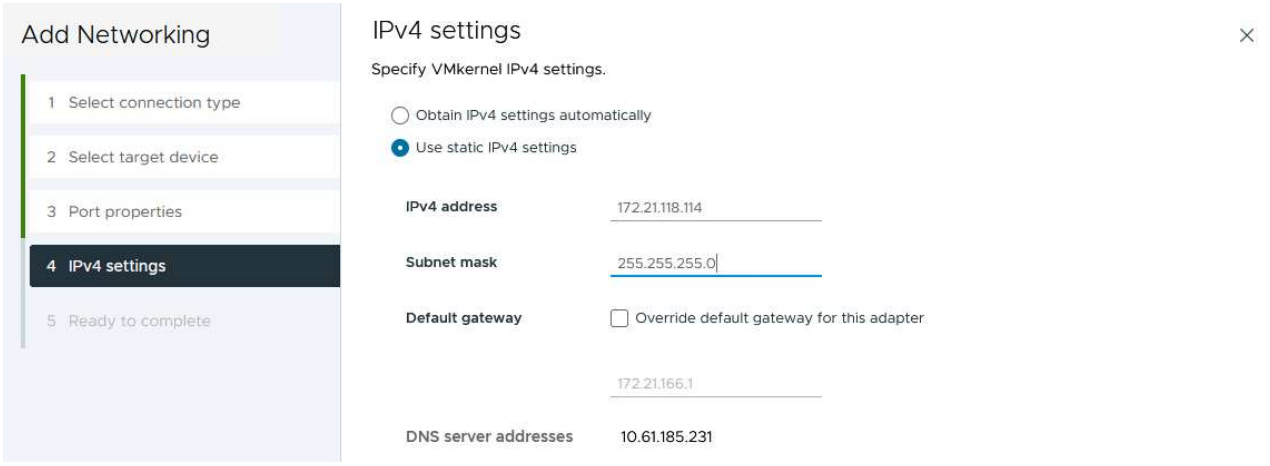

6. Review the your selections on the **Ready to complete** page and click on **Finish** to create the VMkernel adapter.

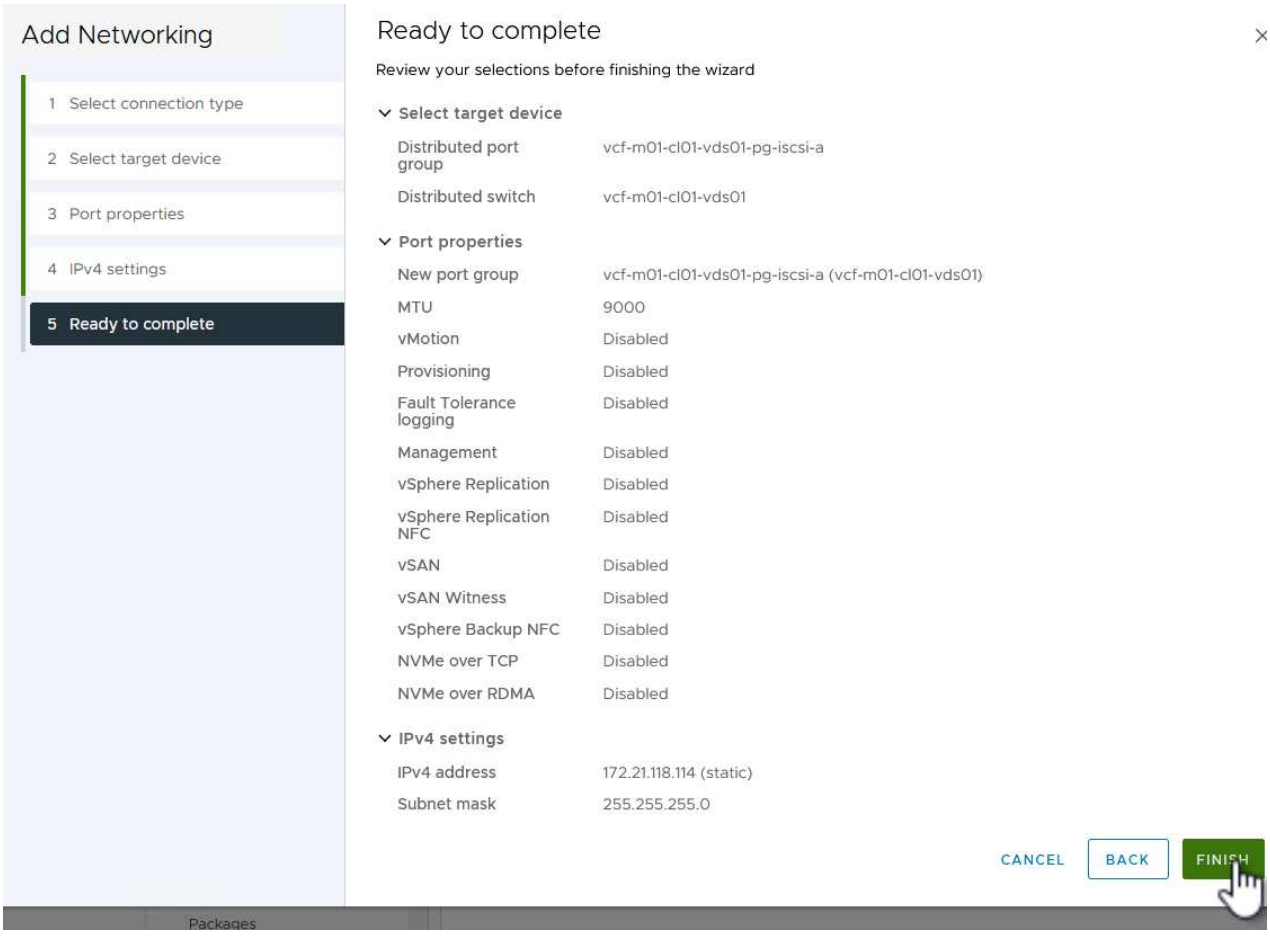

7. Repeat this process to create a VMkernel adapter for the second iSCSI network.

Deploy and use ONTAP Tools to configure storage

The following steps are performed on the VCF management domain cluster using the vSphere client and involve deploying OTV, creating a VMFS iSCSI datastore, and migrating management VM's to the new datastore.

Deploy ONTAP tools for VMware vSphere

ONTAP tools for VMware vSphere (OTV) is deployed as a VM appliance and provides an integrated vCenter UI for managing ONTAP storage.

Complete the following to Deploy ONTAP tools for VMware vSphere:

1. Obtain the ONTAP tools OVA image from the NetApp Support site and download to a local folder.
2. Log into the vCenter appliance for the VCF management domain.
3. From the vCenter appliance interface right-click on the management cluster and select **Deploy OVF Template…**

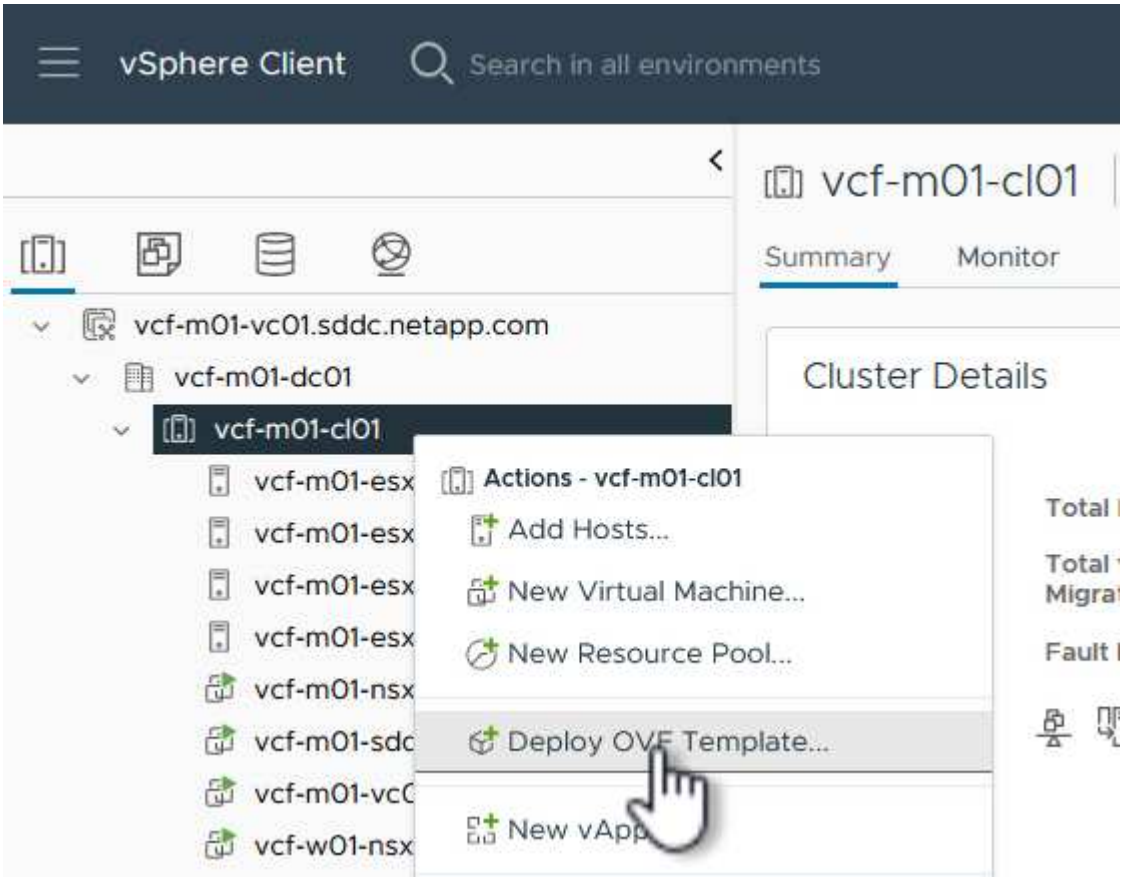

4. In the **Deploy OVF Template** wizard click the **Local file** radio button and select the ONTAP tools OVA file downloaded in the previous step.

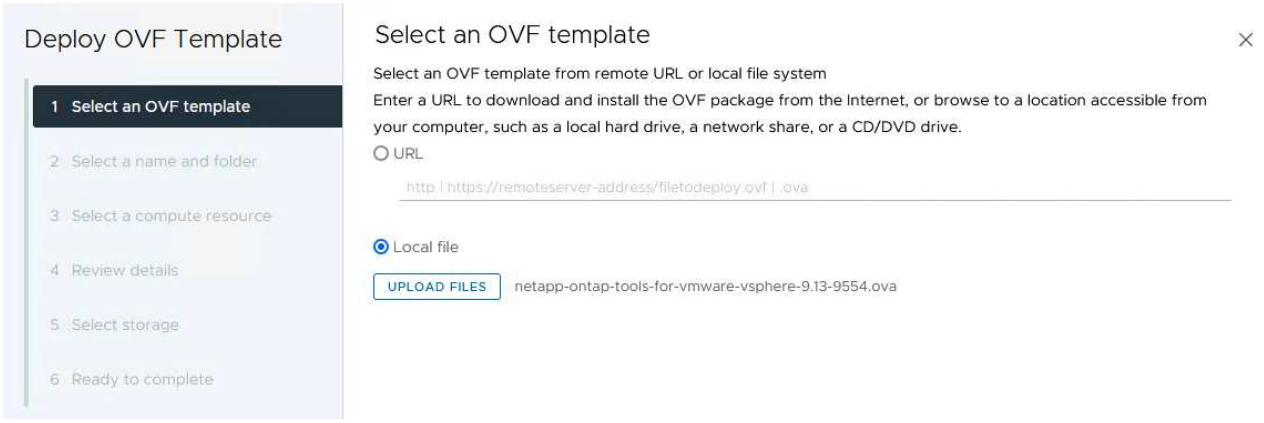

5. For steps 2 through 5 of the wizard select a name and folder for the VM, select the compute resource, review the details, and accept the license agreement.
6. For the storage location of the configuration and disk files, select the vSAN datastore of the VCF management domain cluster.

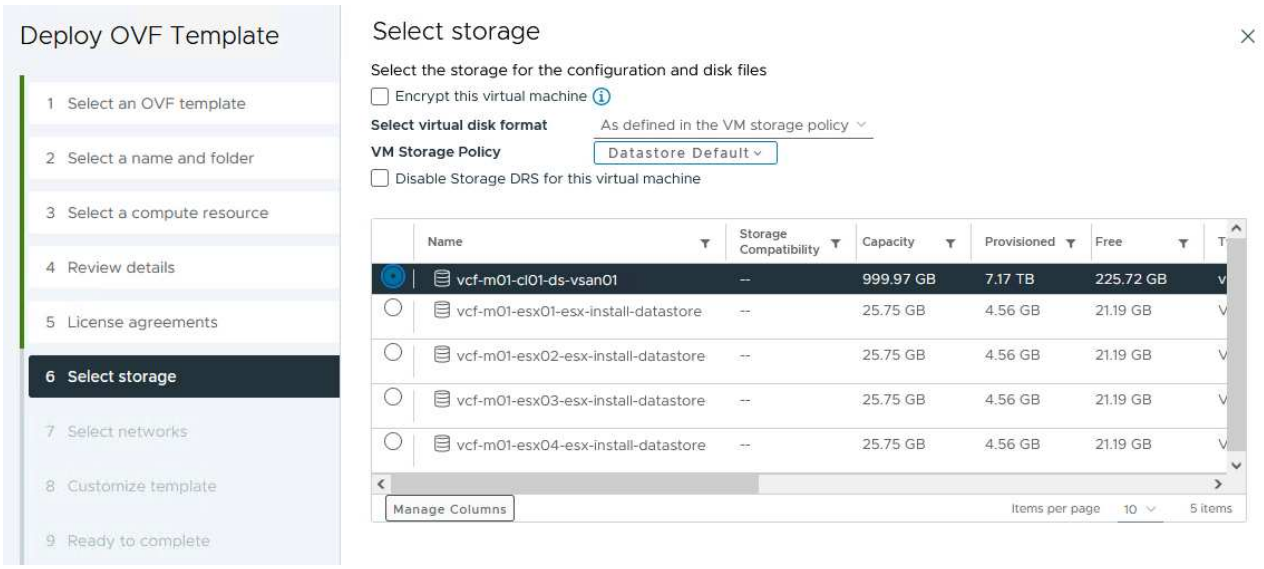

7. On the Select network page select the network used for management traffic.

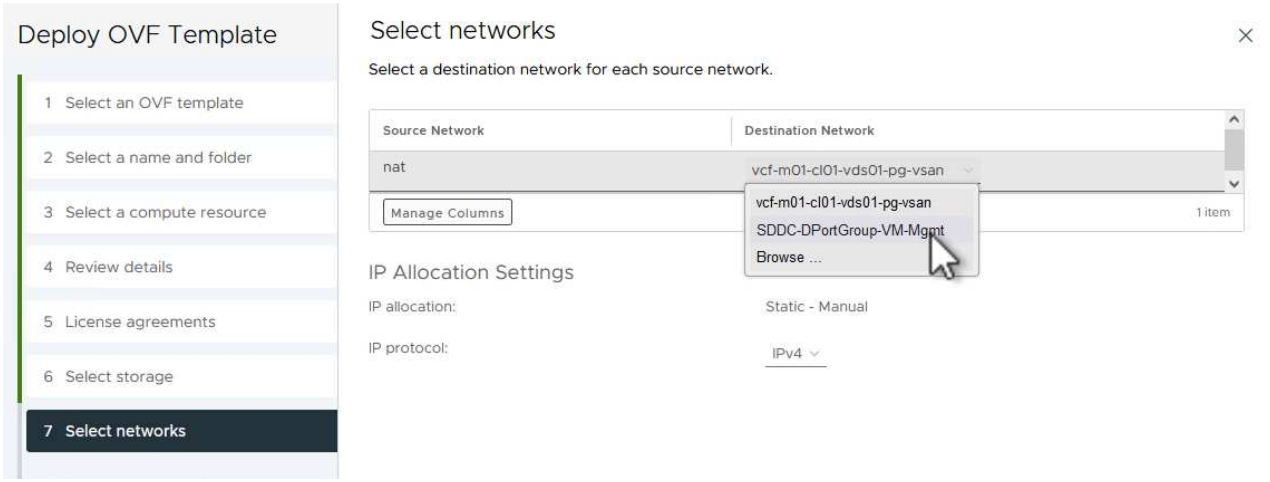

8. On the Customize template page fill out all required information:
   - Password to be used for administrative access to OTV.
   - NTP server IP address.
   - OTV maintenance account password.
   - OTV Derby DB password.
   - Do not check the box to **Enable VMware Cloud Foundation (VCF)**. VCF mode is not required for deploying supplemental storage.
   - FQDN or IP address of the vCenter appliance and provide credentials for vCenter.
   - Provide the required network properties fields.

     Click on **Next** to continue.

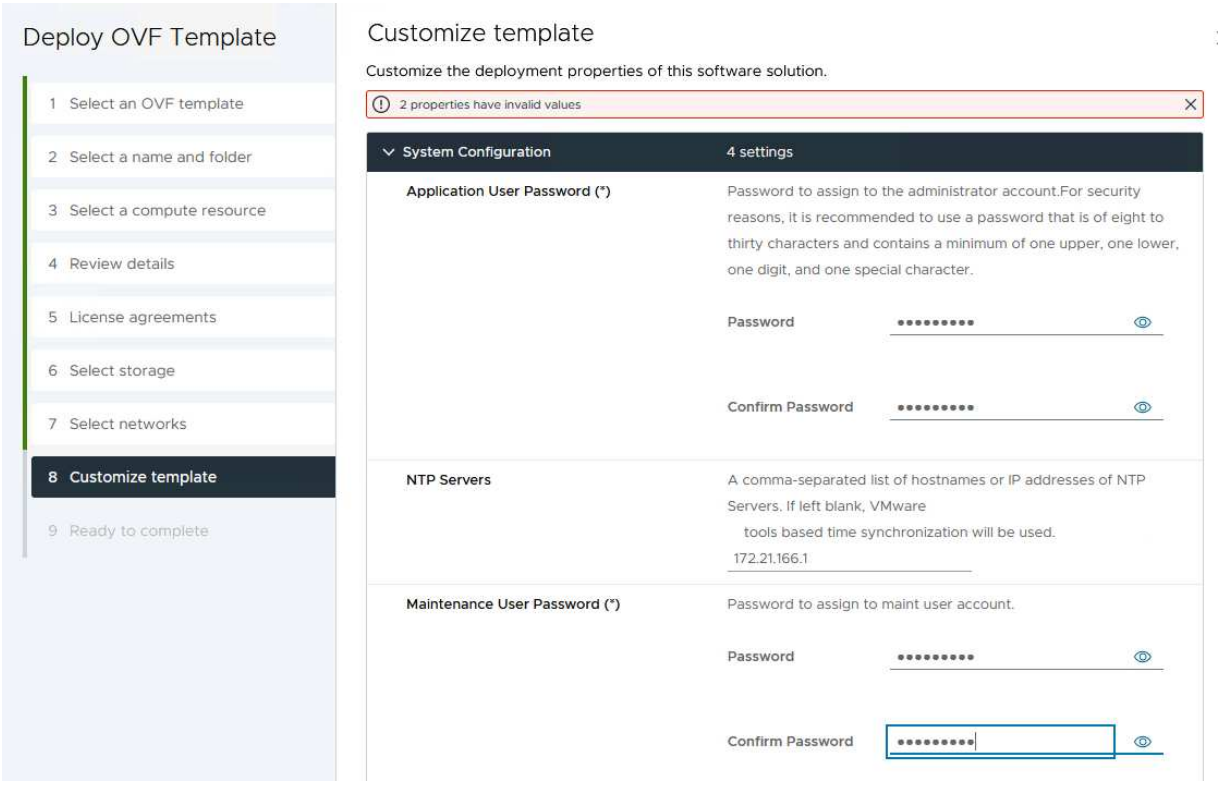

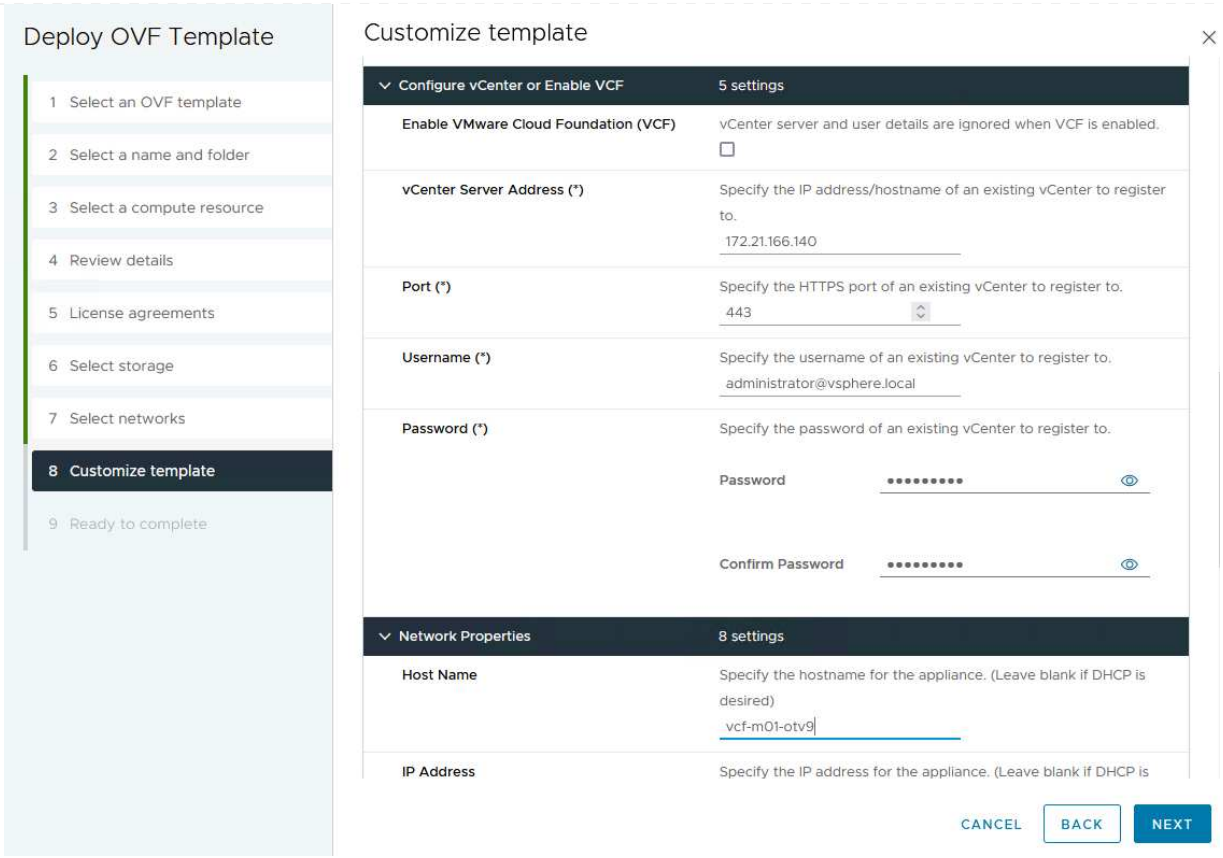

9. Review all information on the Ready to complete page and the click Finish to begin deploying the OTV appliance.

Configure a VMFS iSCSI datastore on Management Domain using OTV

Complete the following to use OTV to configure a VMFS iSCSI datastore as supplemental storage on the management domain:

1. In the vSphere client navigate to the main menu and select **NetApp ONTAP Tools**.

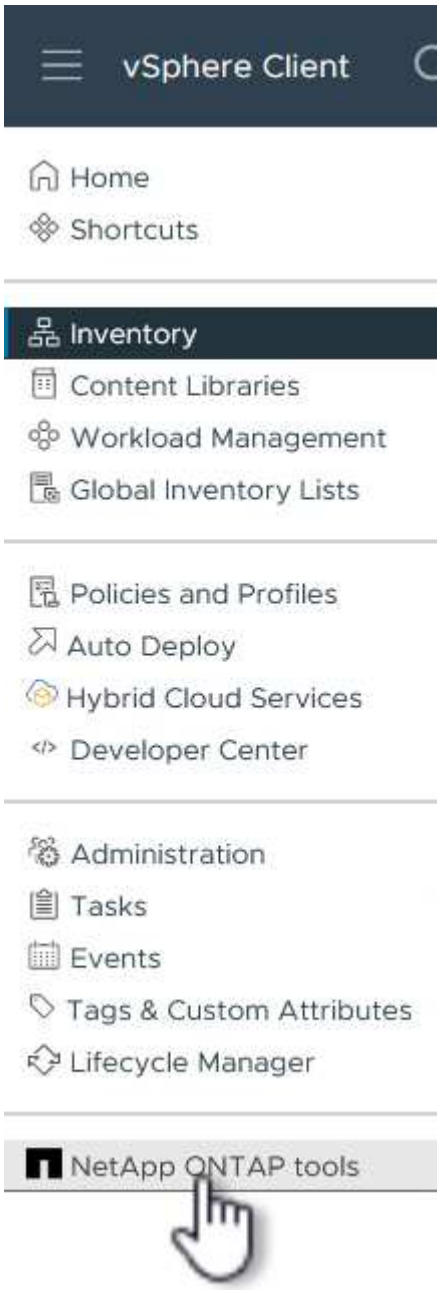

2. Once in **ONTAP Tools**, from the Getting Started page (or from **Storage Systems**), click on **Add** to add a new storage system.

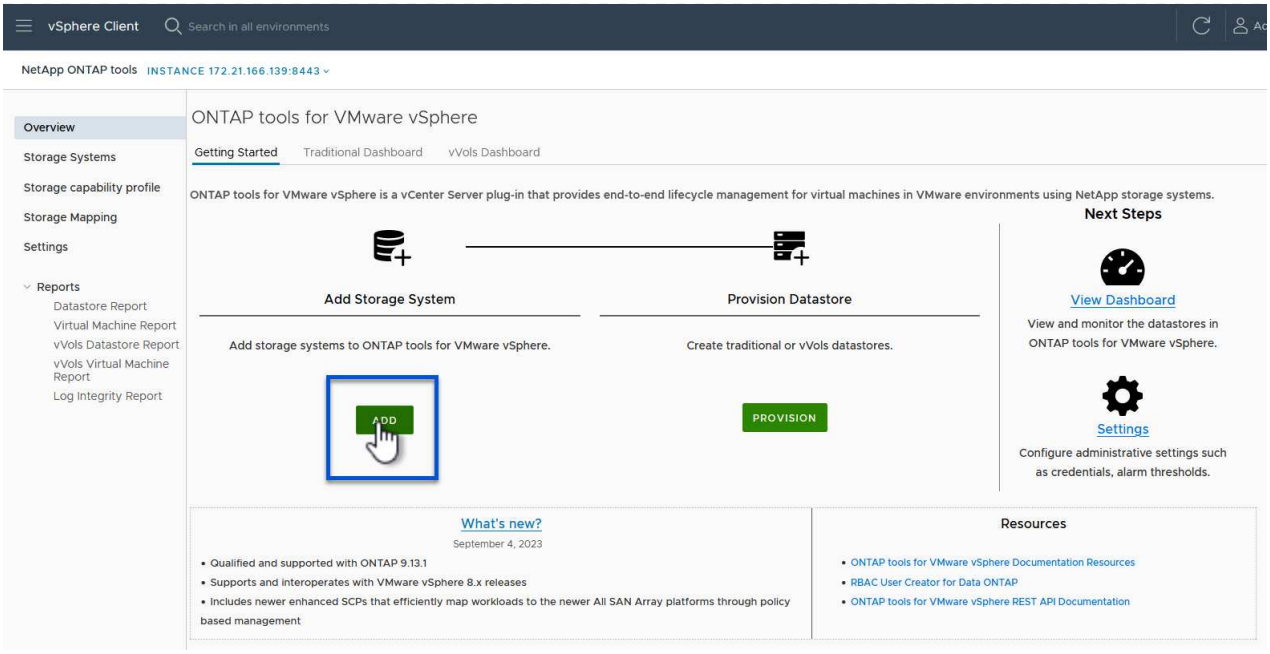

3. Provide the IP address and credentials of the ONTAP storage system and click on **Add**.

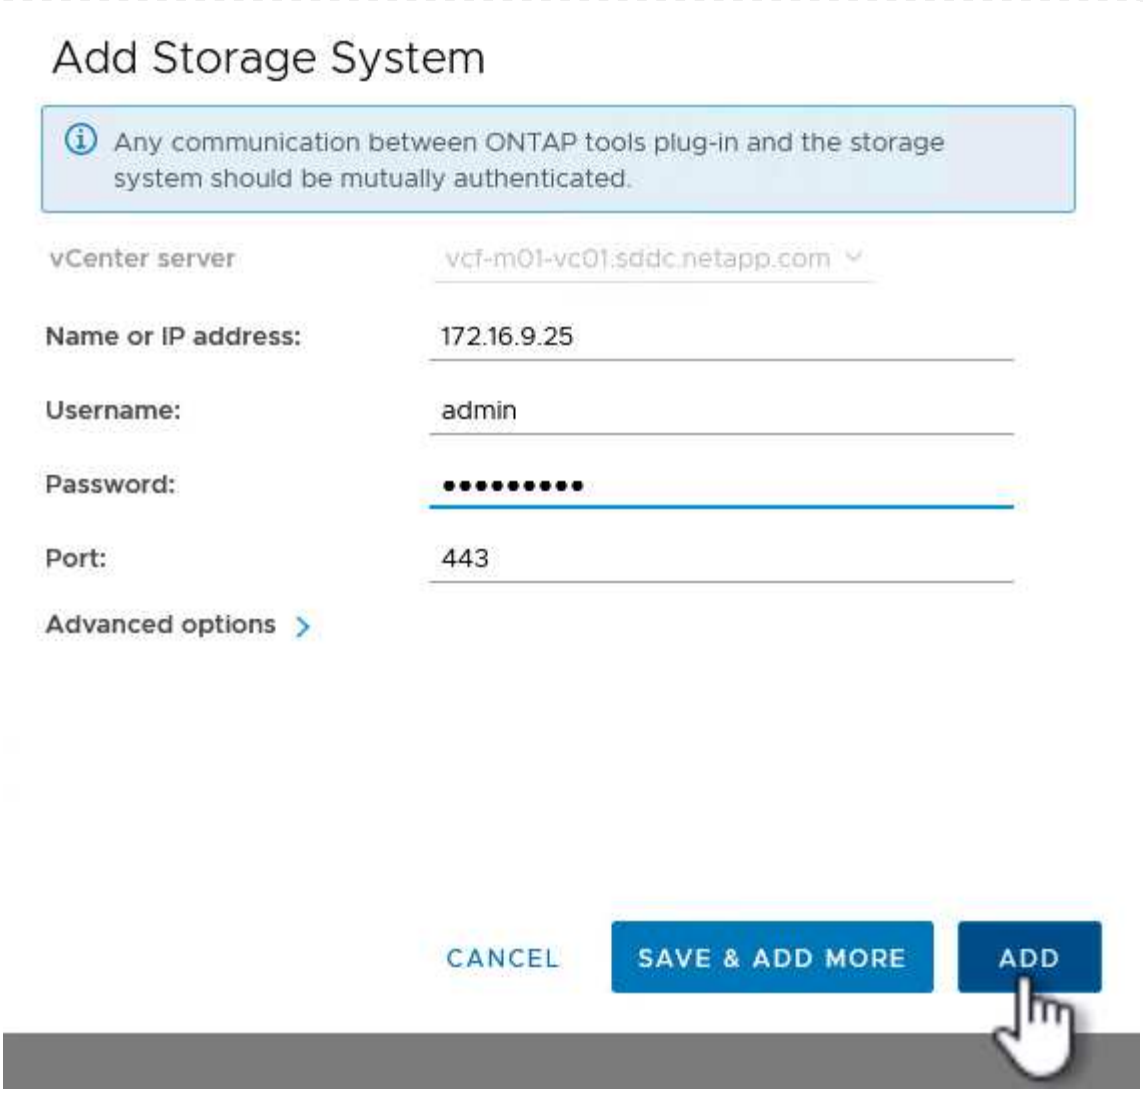

4. Click on **Yes** to authorize the cluster certificate and add the storage system.

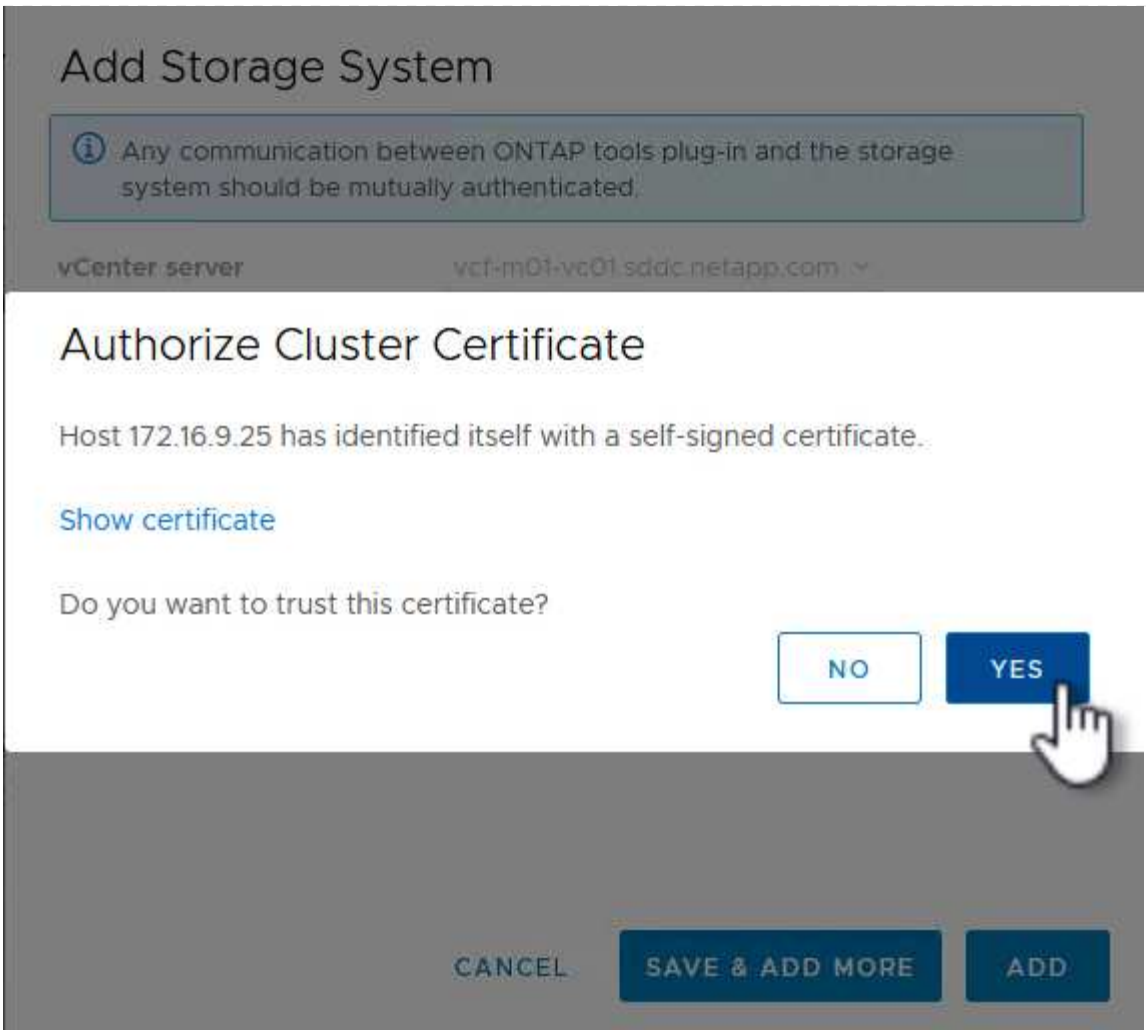
Add Storage System
Any communication between ONTAP tools plug-in and the storage system should be mutually authenticated.
vCenter server
vcf-m01-vc01.sddc.netapp.com
Authorize Cluster Certificate
Host 172.16.9.25 has identified itself with a self-signed certificate.
Show certificate
Do you want to trust this certificate?
NO
YES
CANCEL
SAVE & ADD MORE
ADD

Migrate management VM's to iSCSI Datastore

In cases where it is preferred to use ONTAP storage to protect the VCF management VM's vMotion can be use to migrate the VM's to the newly created iSCSI datastore.

Complete the following steps to migrate the VCF management VM's to the iSCSI datastore.

1. From the vSphere Client navigate to the management domain cluster and click on the **VMs** tab.
2. Select the VMs to be migrated to the iSCSI datastore, right click and select **Migrate..**.

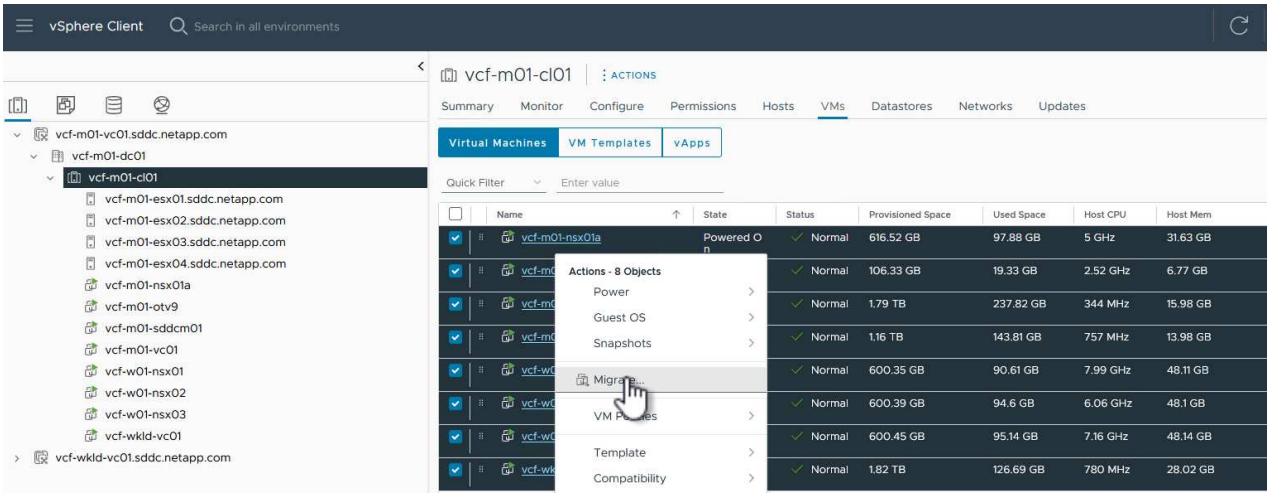

3. In the **Virtual Machines - Migrate** wizard, select **Change storage only** as the migration type and click on **Next** to continue.

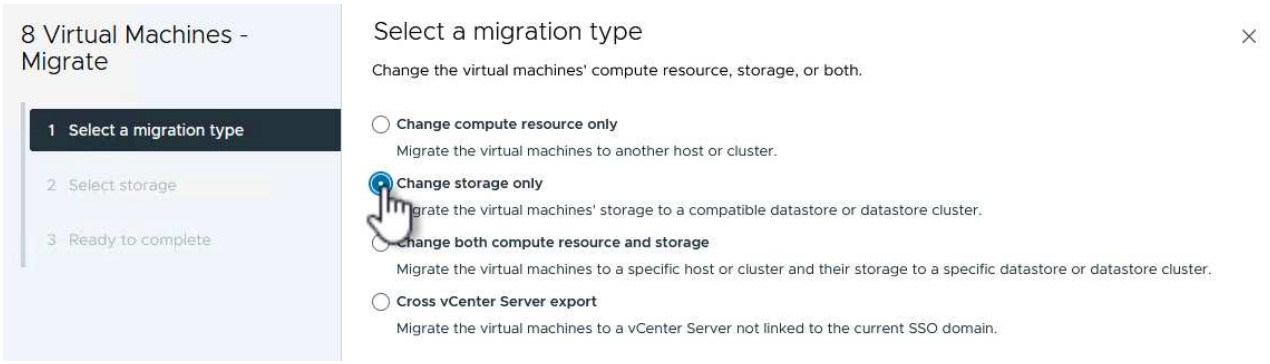

4. On the **Select storage** page, select the iSCSi datastore and select **Next** to continue.

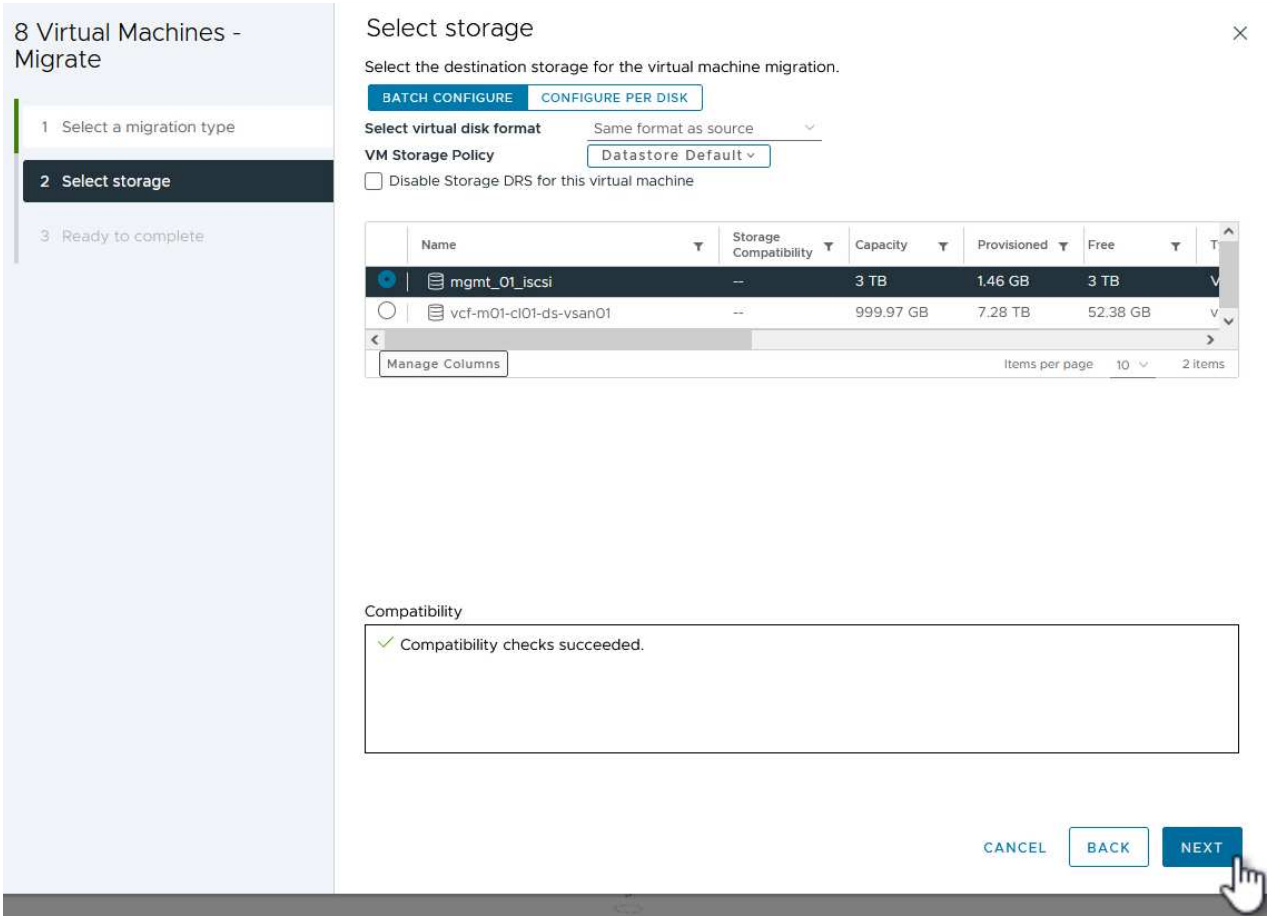

5. Review the selections and click on **Finish** to start the migration.
6. The relocation status can be viewed from the **Recent Tasks** pane.

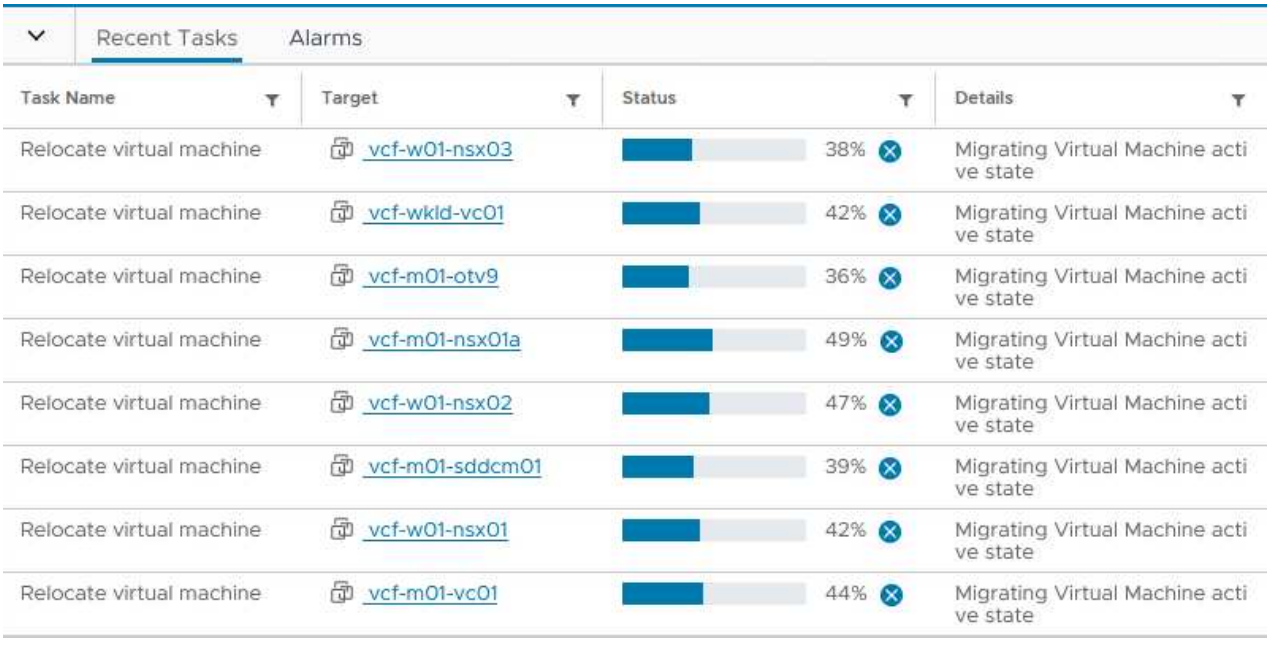

Additional information

For information on configuring ONTAP storage systems refer to the ONTAP 9 Documentation center.

For information on configuring VCF refer to VMware Cloud Foundation Documentation.

Video demo for this solution

iSCSI Datastores as Supplemental Storage for VCF Management Domains

In this scenario we will demonstrate how to deploy and use ONTAP Tools for VMware vSphere (OTV) to configure a **vVols datastore** for a VCF workload domain.

iSCSI is used as the storage protocol for the vVols datastore.

Author: Josh Powell

Use ONTAP Tools to configure supplemental storage (vVols) for VCF Workload Domains

Scenario Overview

This scenario covers the following high level steps:

- Create a storage virtual machine (SVM) with logical interfaces (LIFs) for iSCSI traffic.
- Create distributed port groups for iSCSI networks on the VI workload domain.
- Create vmkernel adapters for iSCSI on the ESXi hosts for the VI workload domain.
- Deploy ONTAP Tools on the VI workload domain.
- Create a new vVols datastore on the VI workload domain.

Prerequisites

This scenario requires the following components and configurations:

- An ONTAP ASA storage system with physical data ports on ethernet switches dedicated to storage traffic.
- VCF management domain deployment is complete and the vSphere client is accessible.
- A VI workload domain has been previously deployed.

NetApp recommends fully redundant network designs for iSCSI. The following diagram illustrates an example of a redundant configuration, providing fault tolerance for storage systems, switches, networks adapters and host systems. Refer to the NetApp SAN configuration reference for additional information.

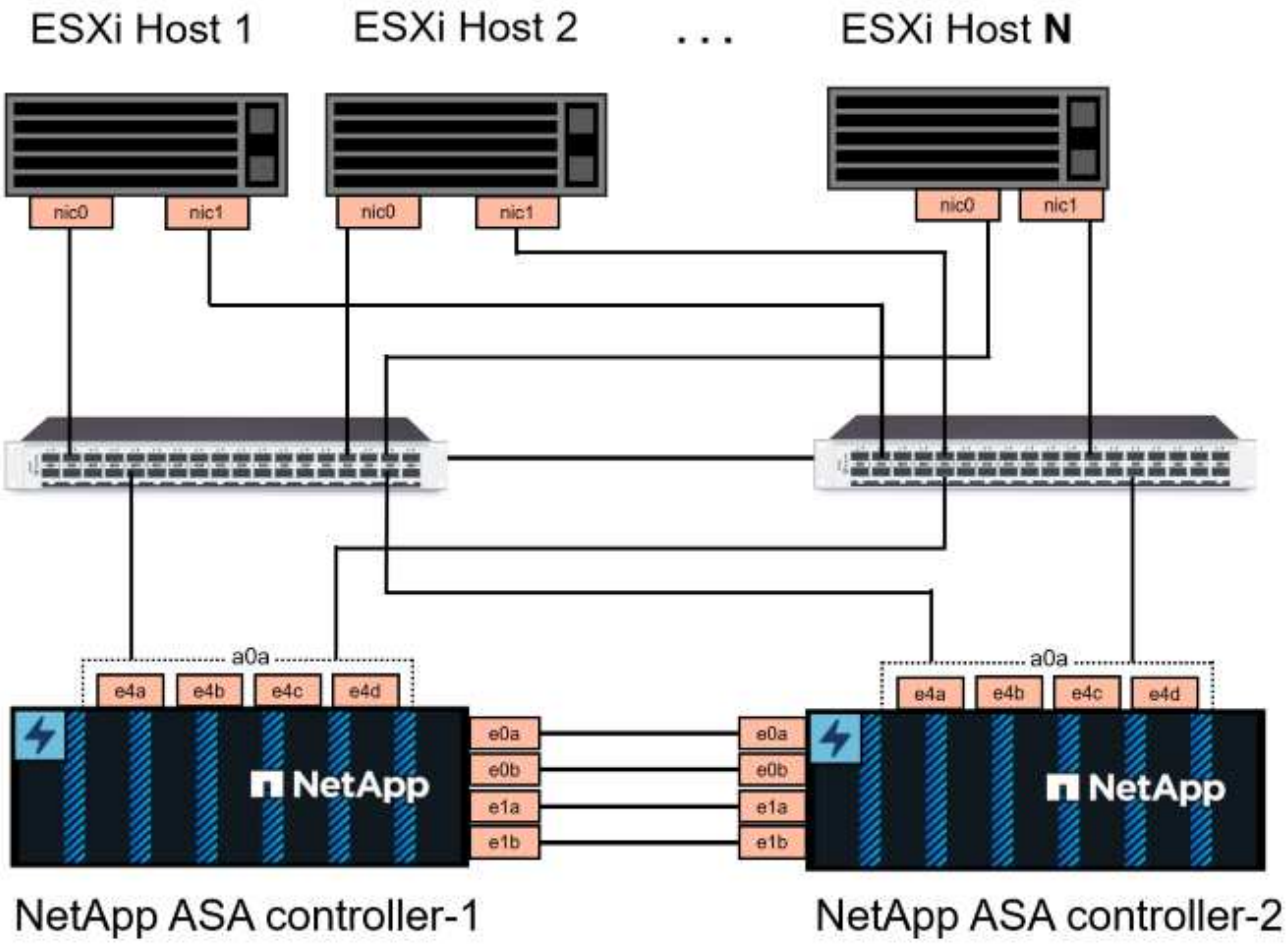

For multipathing and failover across multiple paths, NetApp recommends having a minimum of two LIFs per storage node in separate ethernet networks for all SVMs in iSCSI configurations.

This documentation demonstrates the process of creating a new SVM and specifying the IP address information to create multiple LIFs for iSCSI traffic. To add new LIFs to an existing SVM refer to Create a LIF (network interface).

In situations where multiple VMkernel adapters are configured on the same IP network, it is recommended to use software iSCSI port binding on the ESXi hosts to ensure that load balancing across the adapters occurs. Refer to KB article Considerations for using software iSCSI port binding in ESX/ESXi (2038869).

For additional information on using VMFS iSCSI datastores with VMware refer to vSphere VMFS Datastore - iSCSI Storage backend with ONTAP.

Deployment Steps

To deploy ONTAP Tools and use it to create a vVols datastore on the VCF management domain, complete the following steps:

Create SVM and LIFs on ONTAP storage system

The following step is performed in ONTAP System Manager.

Create the storage VM and LIFs

Complete the following steps to create an SVM together with multiple LIFs for iSCSI traffic.

1. From ONTAP System Manager navigate to **Storage VMs** in the left-hand menu and click on **+ Add** to start.

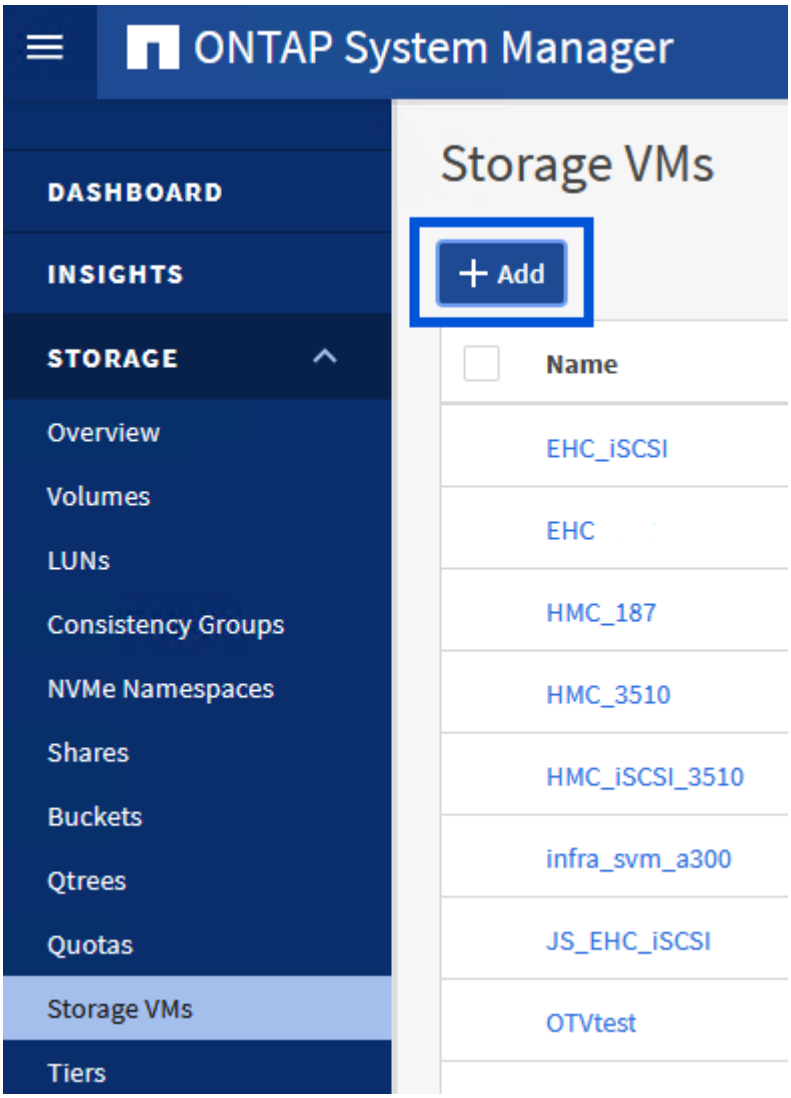

2. In the **Add Storage VM** wizard provide a **Name** for the SVM, select the **IP Space** and then, under **Access Protocol**, click on the **iSCSI** tab and check the box to **Enable iSCSI**.

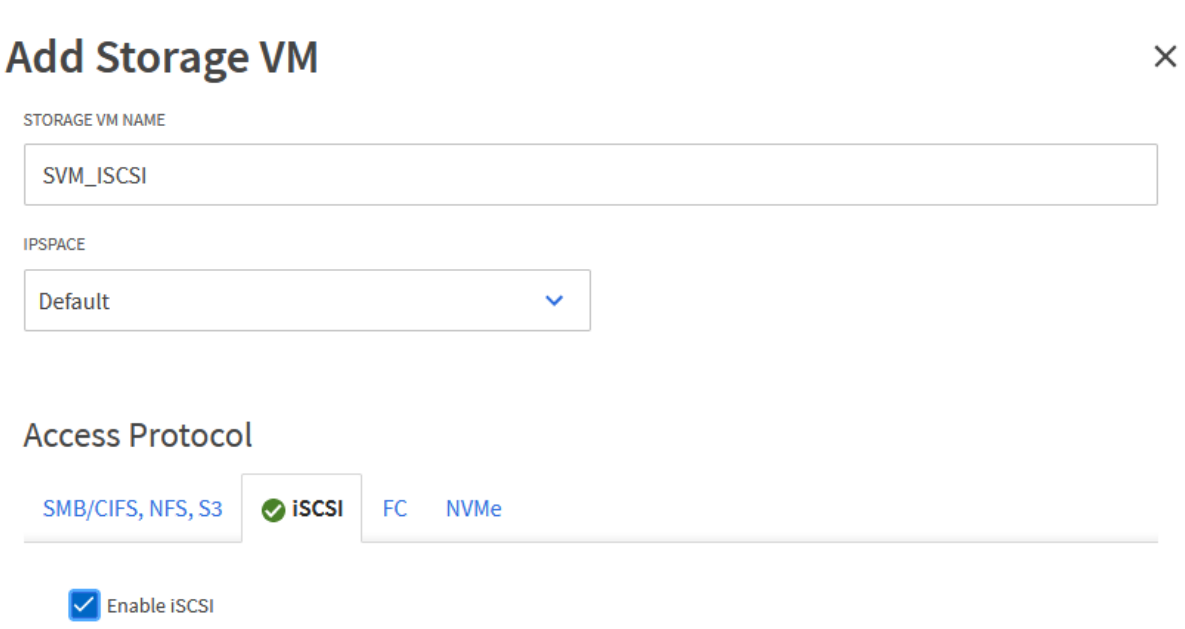

3. In the **Network Interface** section fill in the **IP address**, **Subnet Mask**, and **Broadcast Domain and Port** for the first LIF. For subsequent LIFs the checkbox may be enabled to use common settings across all remaining LIFs or use separate settings.

> For multipathing and failover across multiple paths, NetApp recommends having a minimum of two LIFs per storage node in separate Ethernet networks for all SVMs in iSCSI configurations.

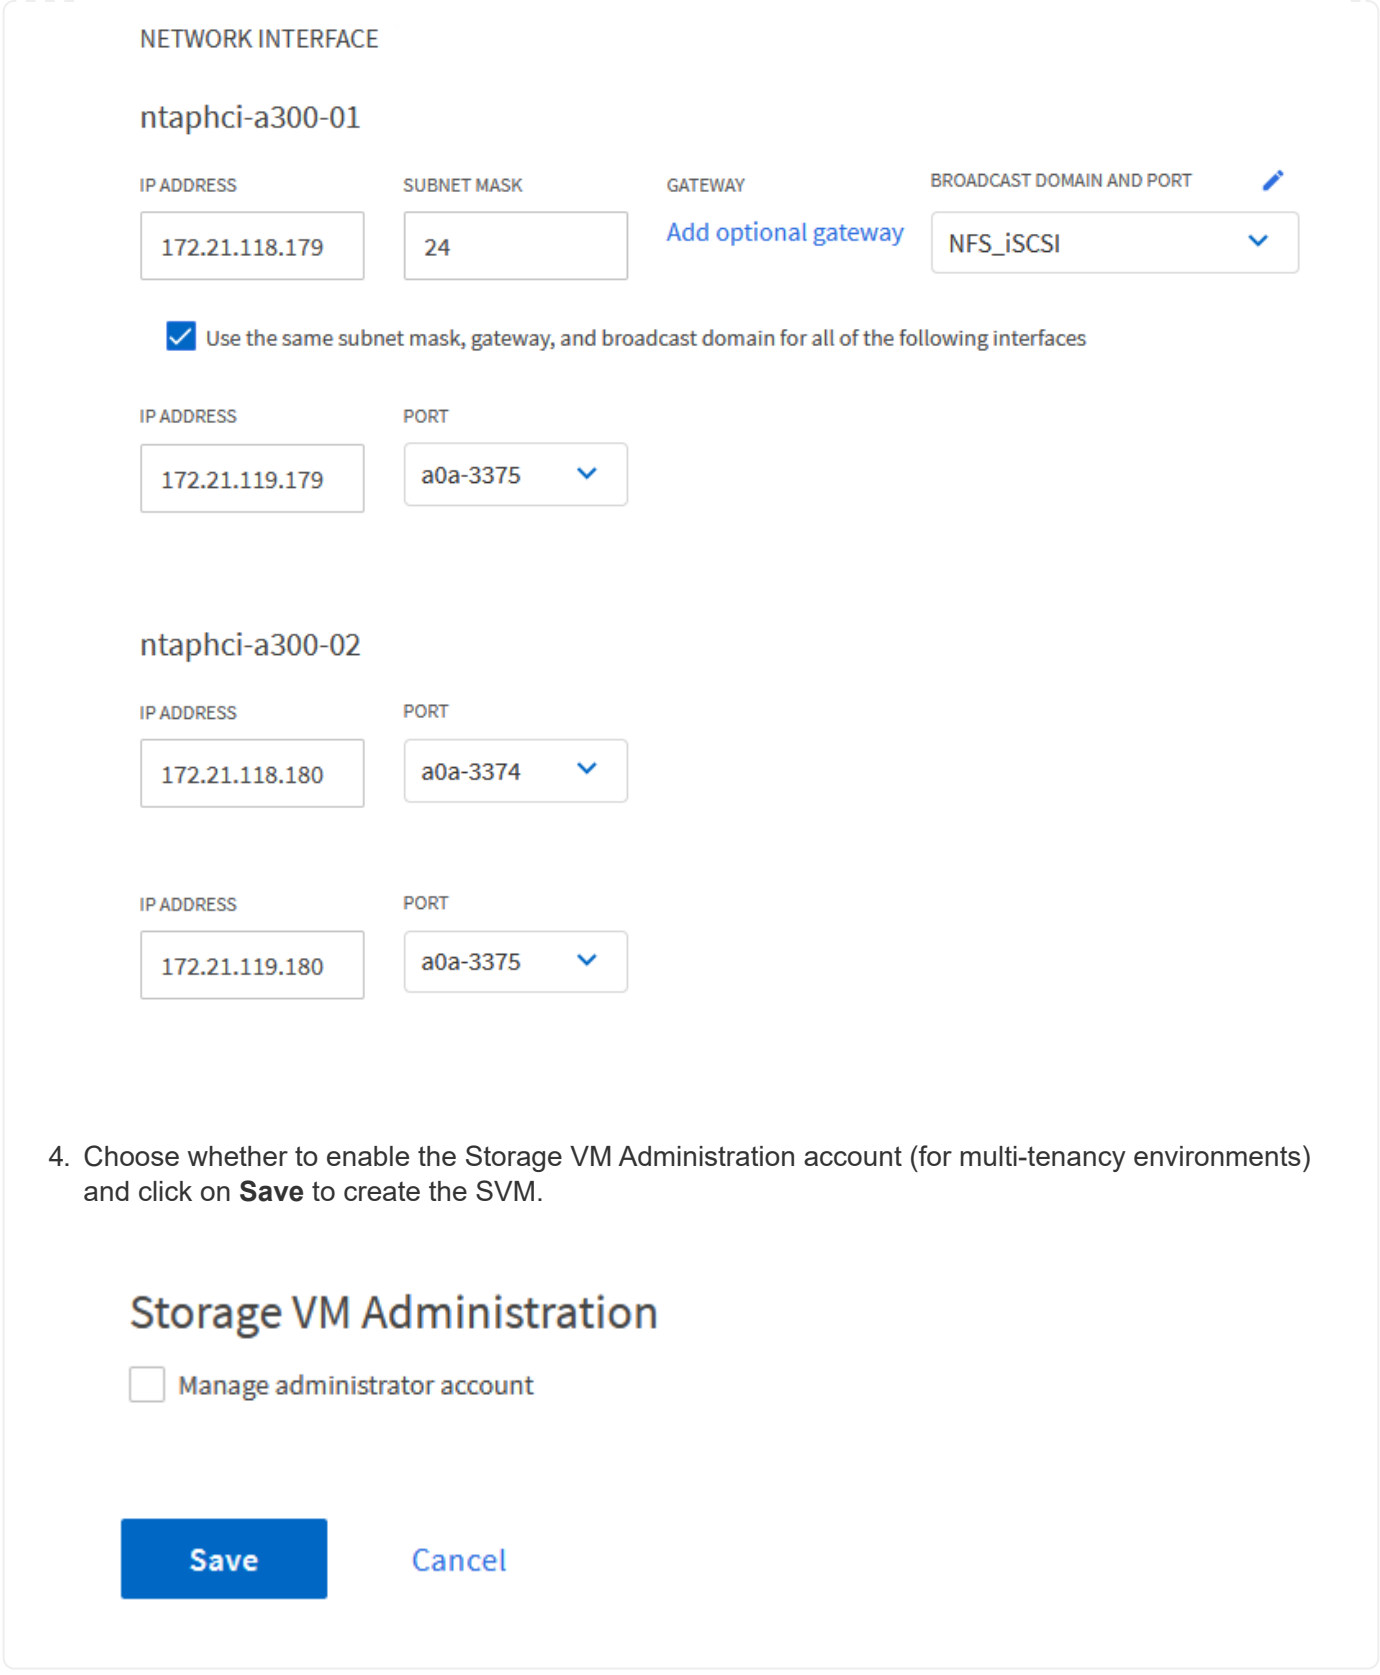

4. Choose whether to enable the Storage VM Administration account (for multi-tenancy environments) and click on **Save** to create the SVM.

Set up networking for iSCSI on ESXi hosts

The following steps are performed on the VI Workload Domain cluster using the vSphere client. In this case vCenter Single Sign-On is being used so the vSphere client is common across the management and workload domains.

Create Distributed Port Groups for iSCSI traffic

Complete the following to create a new distributed port group for each iSCSI network:

1. From the vSphere client , navigate to **Inventory > Networking** for the workload domain. Navigate to the existing Distributed Switch and choose the action to create **New Distributed Port Group…**.

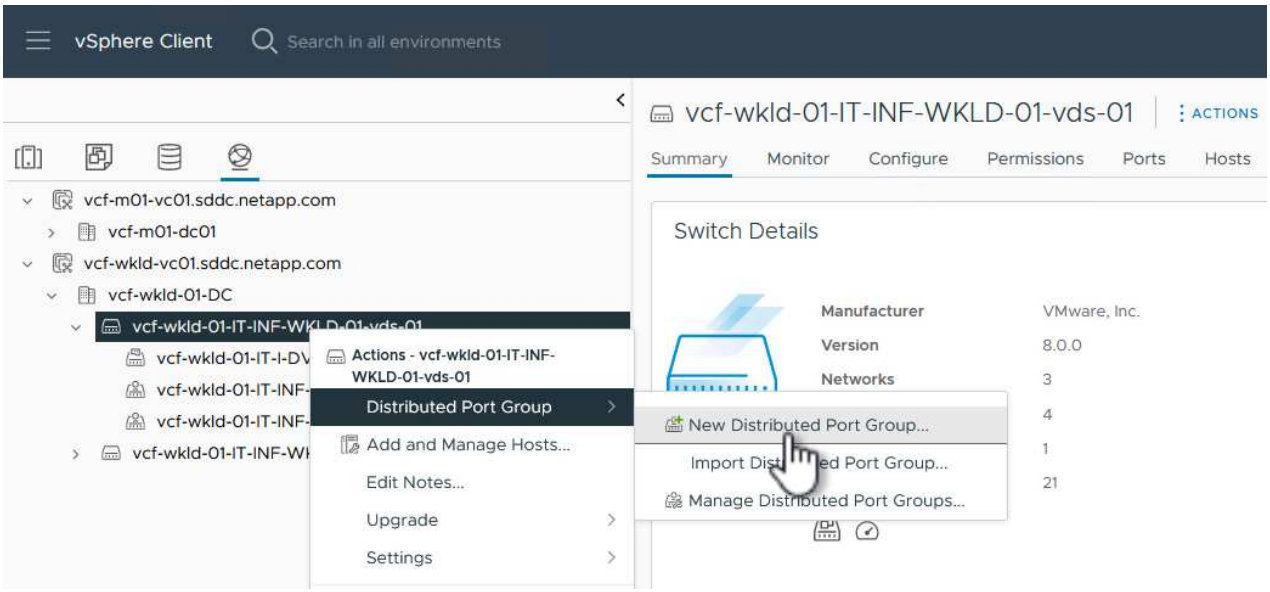

2. In the **New Distributed Port Group** wizard fill in a name for the new port group and click on **Next** to continue.
3. On the **Configure settings** page fill out all settings. If VLANs are being used be sure to provide the correct VLAN ID. Click on **Next** to continue.

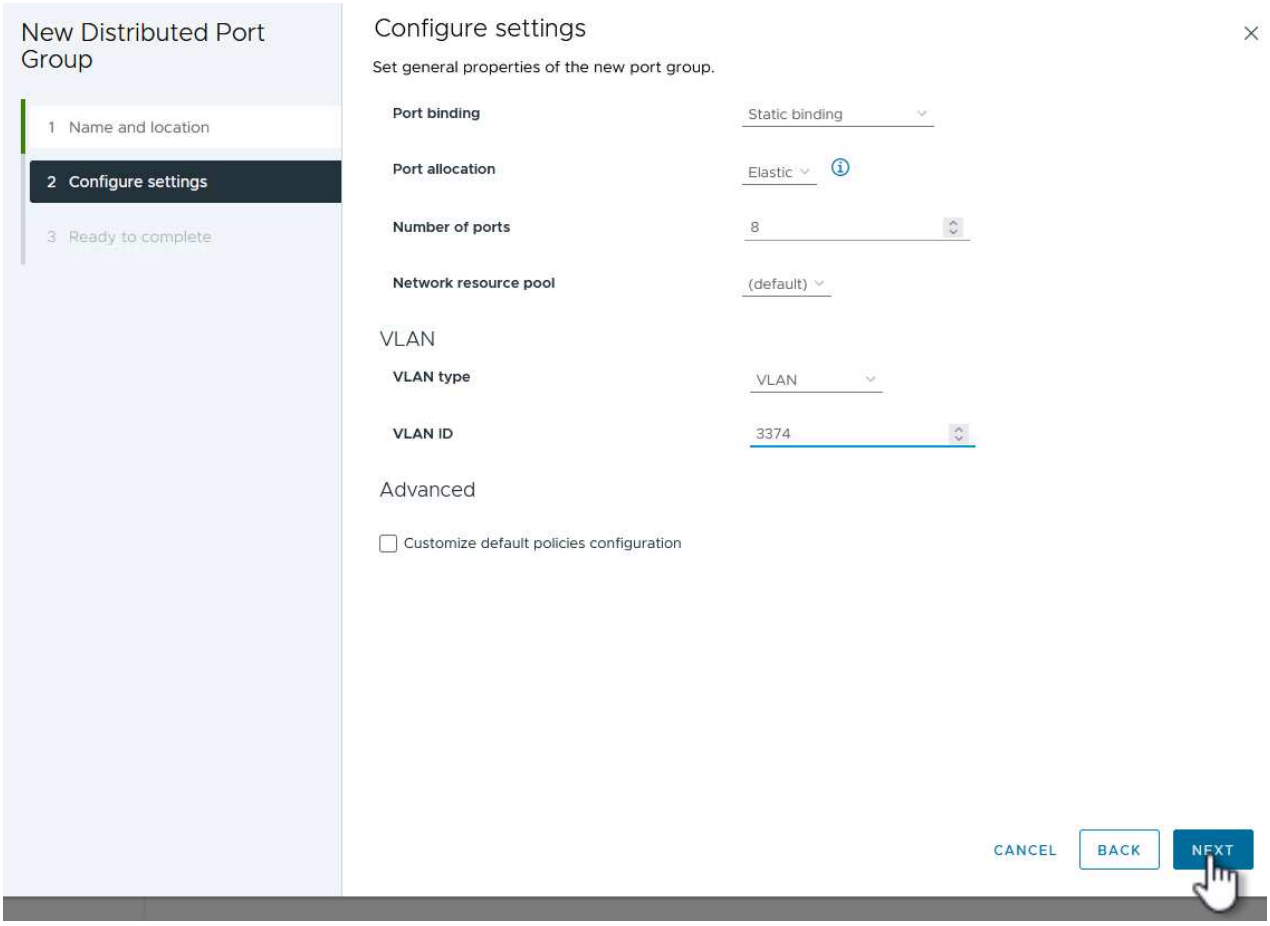

4. On the **Ready to complete** page, review the changes and click on **Finish** to create the new distributed port group.
5. Repeat this process to create a distributed port group for the second iSCSI network being used and ensure you have input the correct **VLAN ID**.
6. Once both port groups have been created, navigate to the first port group and select the action to **Edit settings…**.

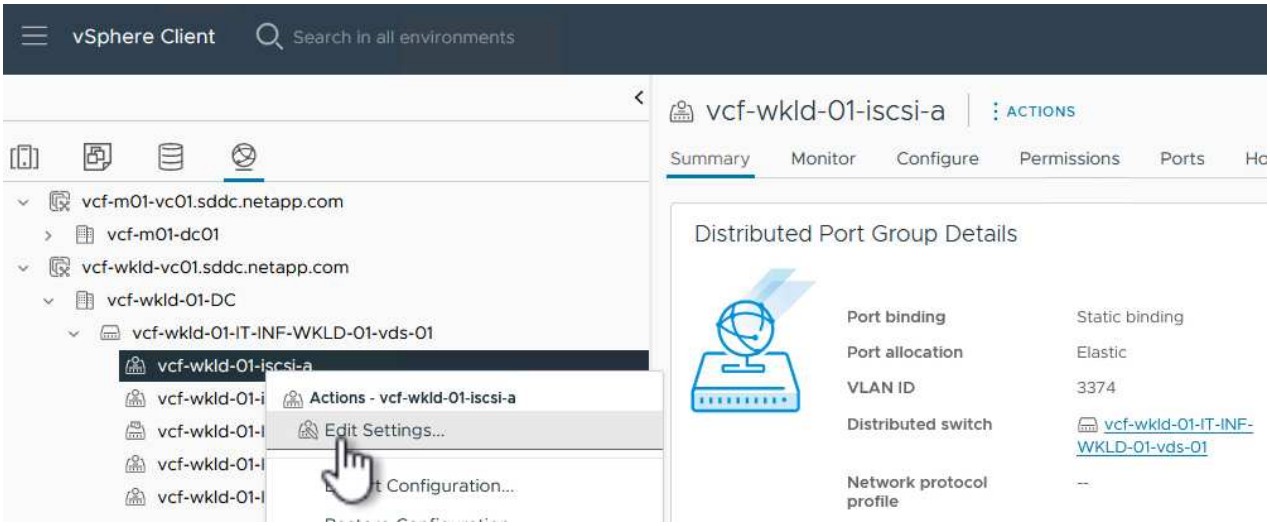

7. On **Distributed Port Group - Edit Settings** page, navigate to **Teaming and failover** in the left-hand menu and click on **uplink2** to move it down to **Unused uplinks**.

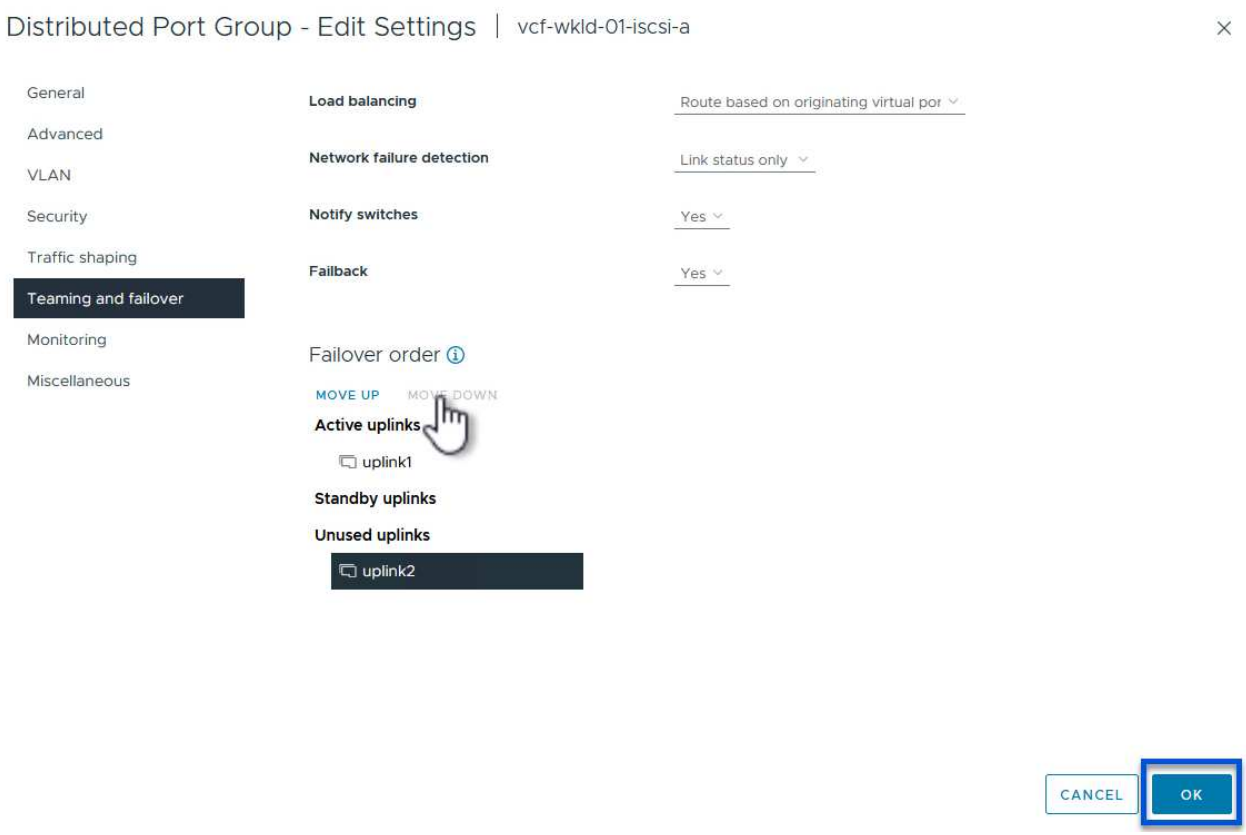

8. Repeat this step for the second iSCSI port group. However, this time move **uplink1** down to **Unused uplinks**.

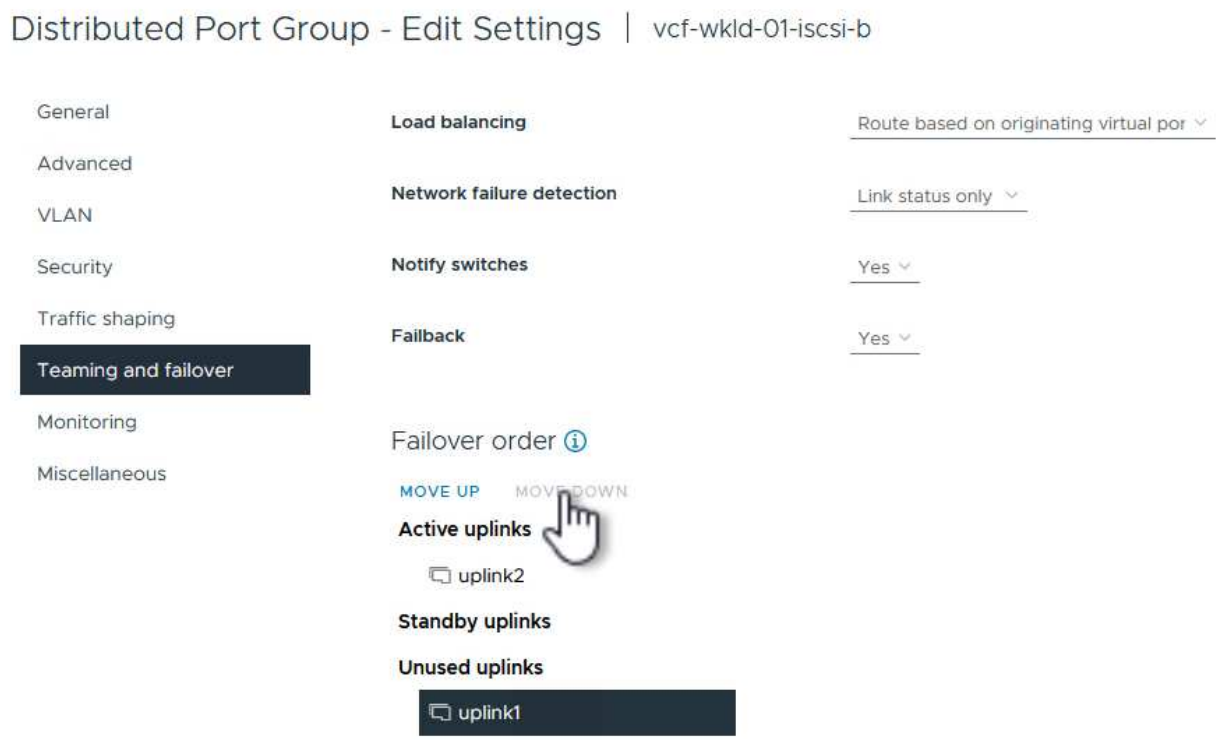

Create VMkernel adapters on each ESXi host

Repeat this process on each ESXi host in the workload domain.

1. From the vSphere client navigate to one of the ESXi hosts in the workload domain inventory. From the **Configure** tab select **VMkernel adapters** and click on **Add Networking…** to start.

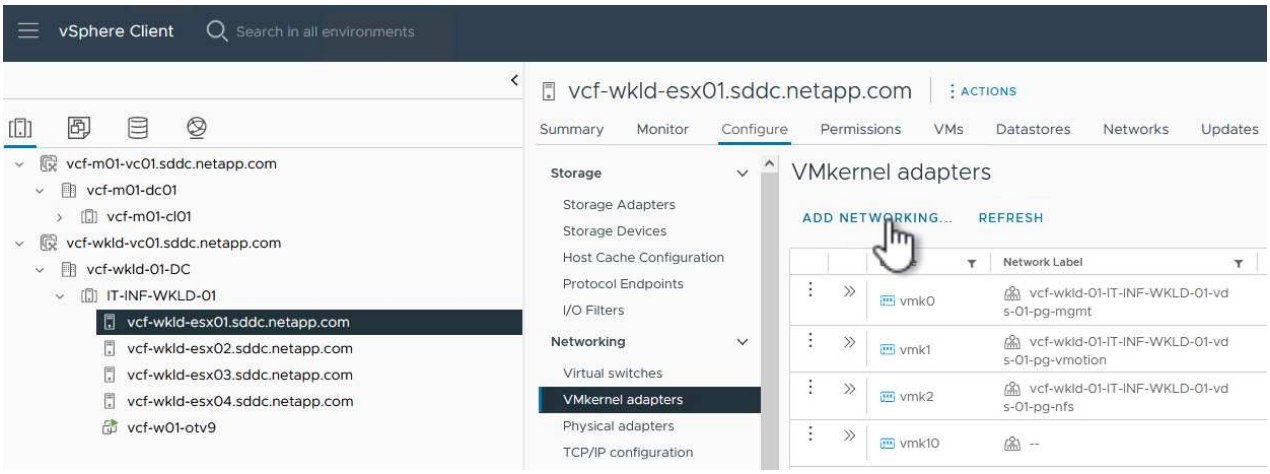

2. On the **Select connection type** window choose **VMkernel Network Adapter** and click on **Next** to continue.

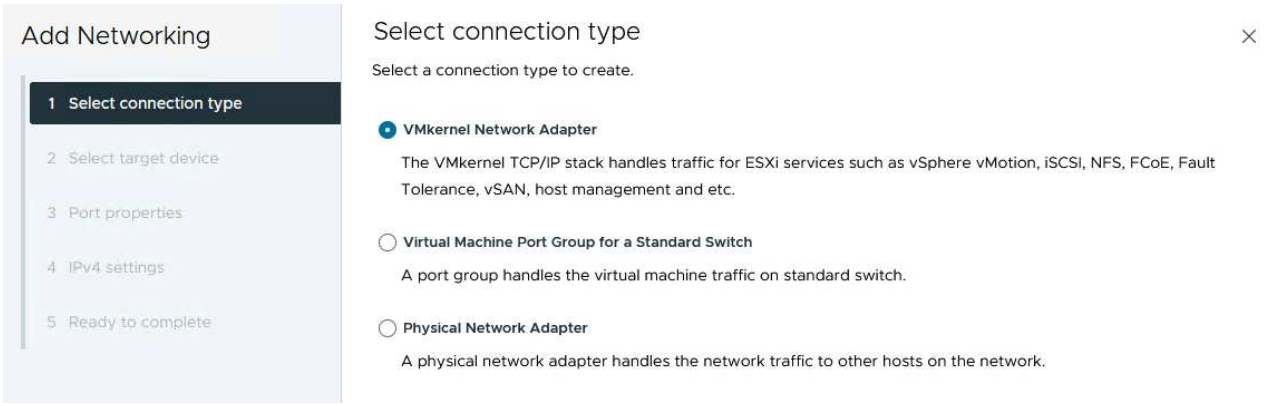

3. On the **Select target device** page, choose one of the distributed port groups for iSCSI that was created previously.

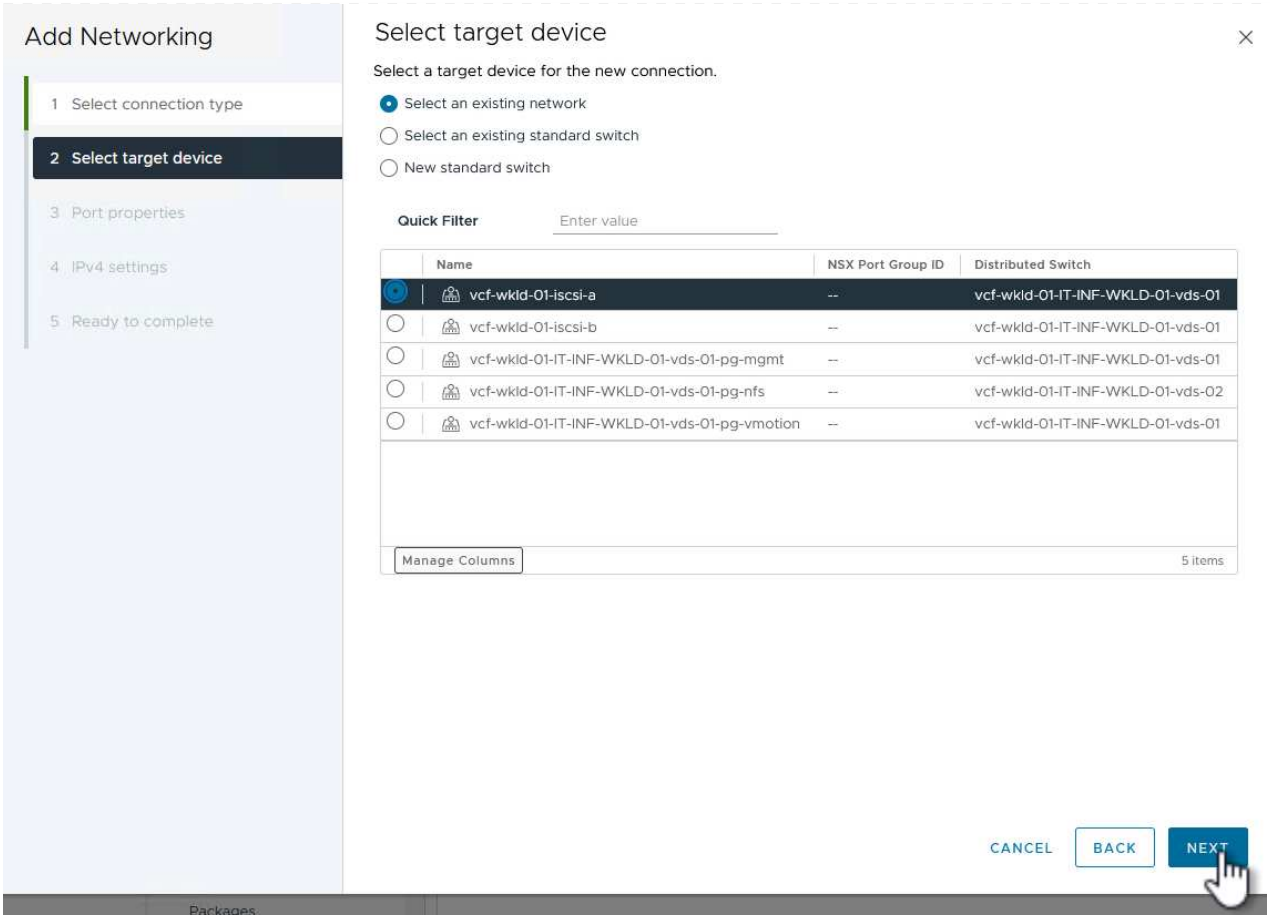

4. On the **Port properties** page keep the defaults and click on **Next** to continue.

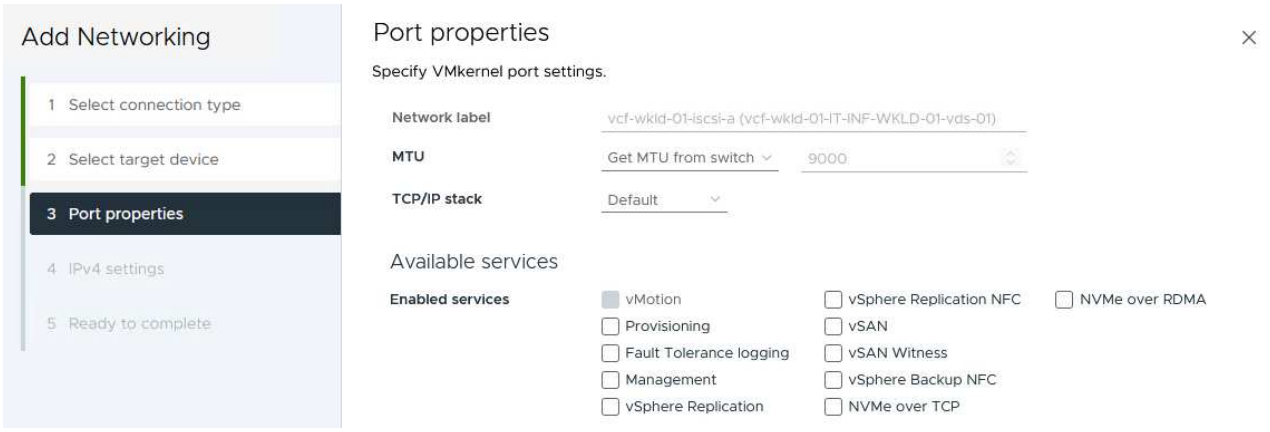

5. On the **IPv4 settings** page fill in the **IP address**, **Subnet mask**, and provide a new Gateway IP address (only if required). Click on **Next** to continue.

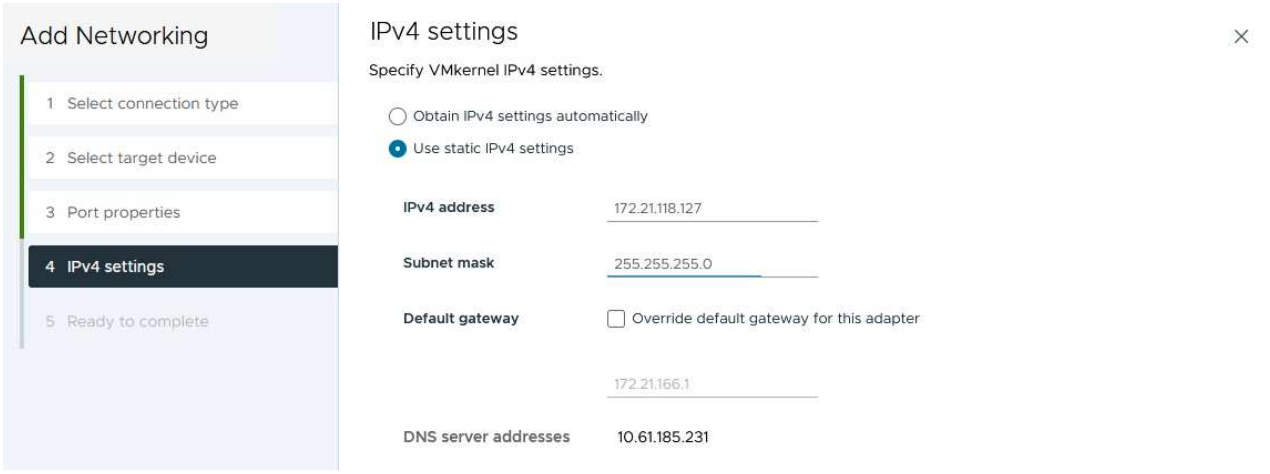

6. Review the your selections on the **Ready to complete** page and click on **Finish** to create the VMkernel adapter.

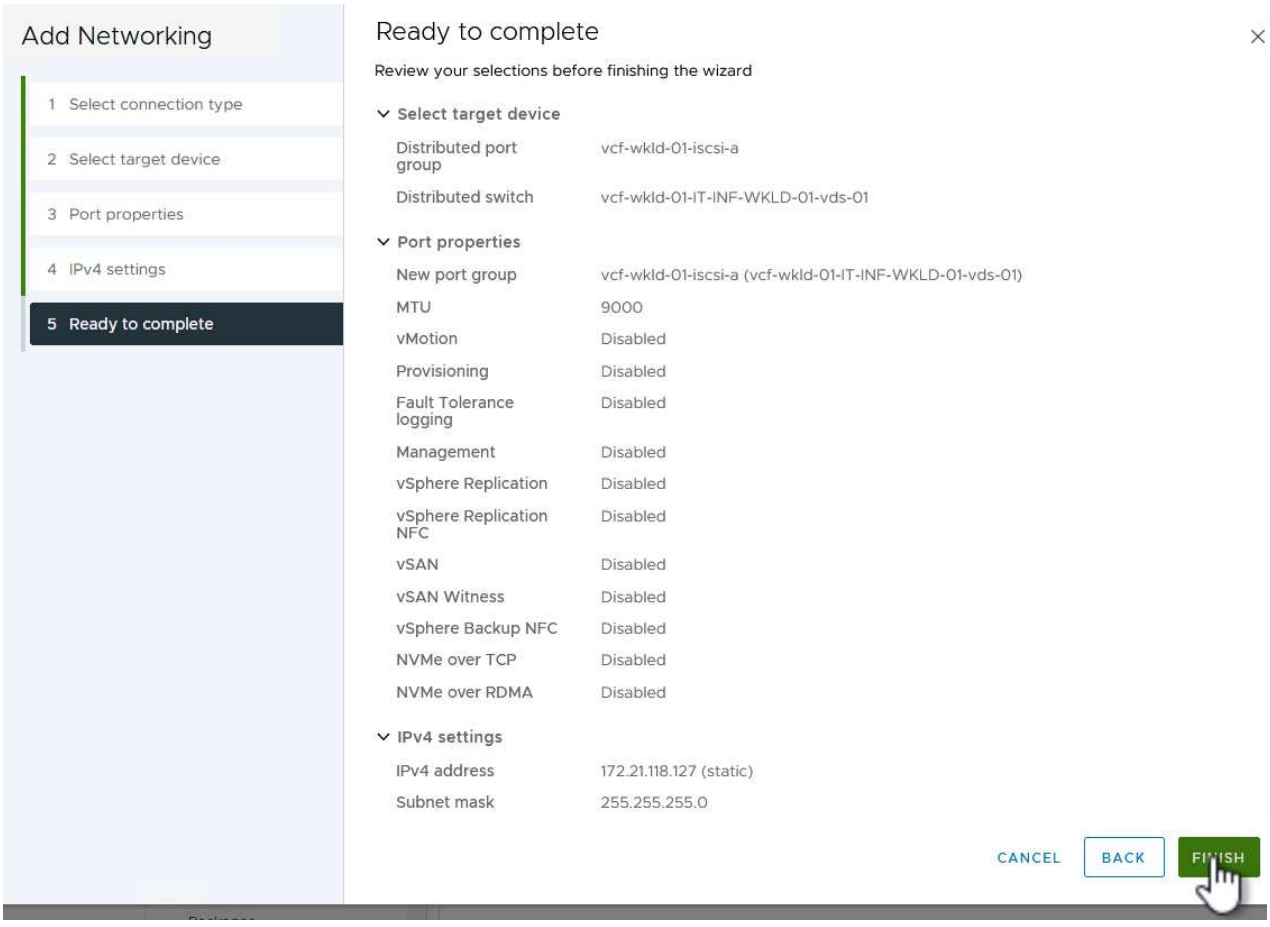

7. Repeat this process to create a VMkernel adapter for the second iSCSI network.

Deploy and use ONTAP Tools to configure storage

The following steps are performed on the VCF management domain cluster using the vSphere client and involve deploying OTV, creating a vVols iSCSI datastore, and migrating management VM's to the new datastore.

For VI workload domains, OTV is installed to the VCF Management Cluster but registered with the vCenter associated with the VI workload domain.

For additional information on deploying and using ONTAP Tools in a multiple vCenter environment refer to Requirements for registering ONTAP tools in multiple vCenter Servers environment.

Deploy ONTAP tools for VMware vSphere

ONTAP tools for VMware vSphere (OTV) is deployed as a VM appliance and provides an integrated vCenter UI for managing ONTAP storage.

Complete the following to Deploy ONTAP tools for VMware vSphere:

1. Obtain the ONTAP tools OVA image from the NetApp Support site and download to a local folder.
2. Log into the vCenter appliance for the VCF management domain.
3. From the vCenter appliance interface right-click on the management cluster and select **Deploy OVF Template…**

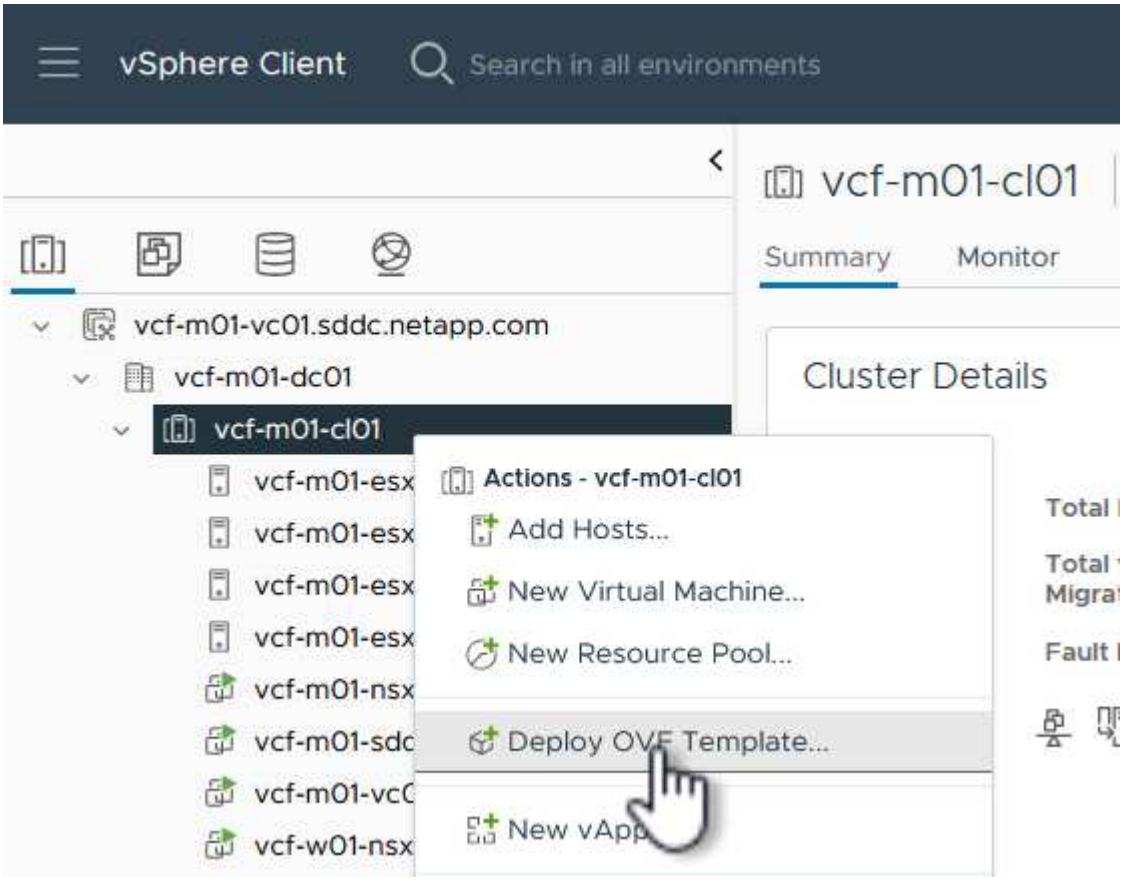

4. In the **Deploy OVF Template** wizard click the **Local file** radio button and select the ONTAP tools OVA file downloaded in the previous step.

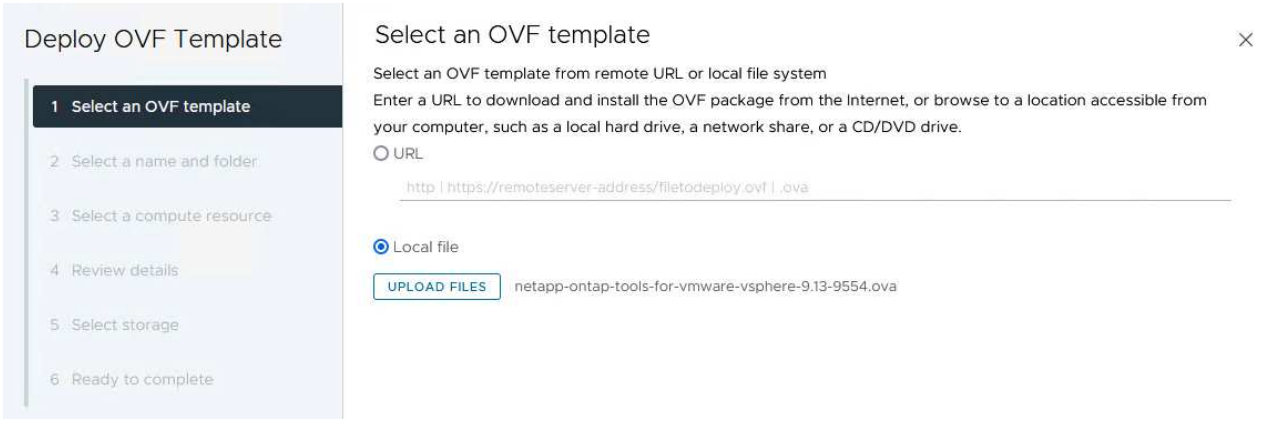

5. For steps 2 through 5 of the wizard select a name and folder for the VM, select the compute resource, review the details, and accept the license agreement.
6. For the storage location of the configuration and disk files, select the vSAN datastore of the VCF management domain cluster.

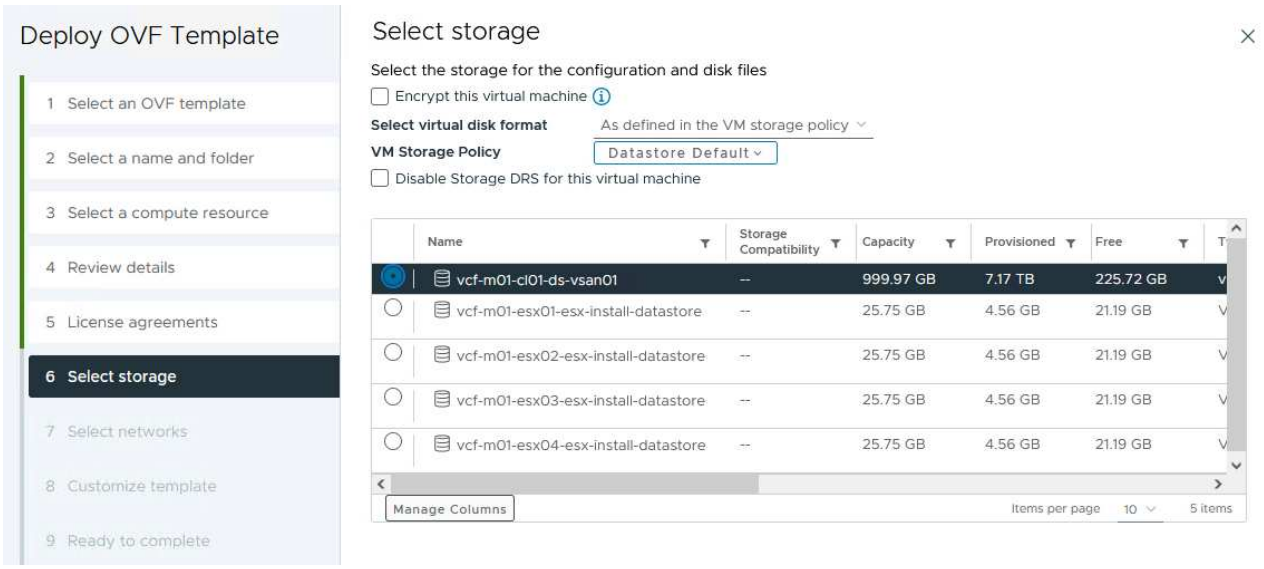

7. On the Select network page select the network used for management traffic.

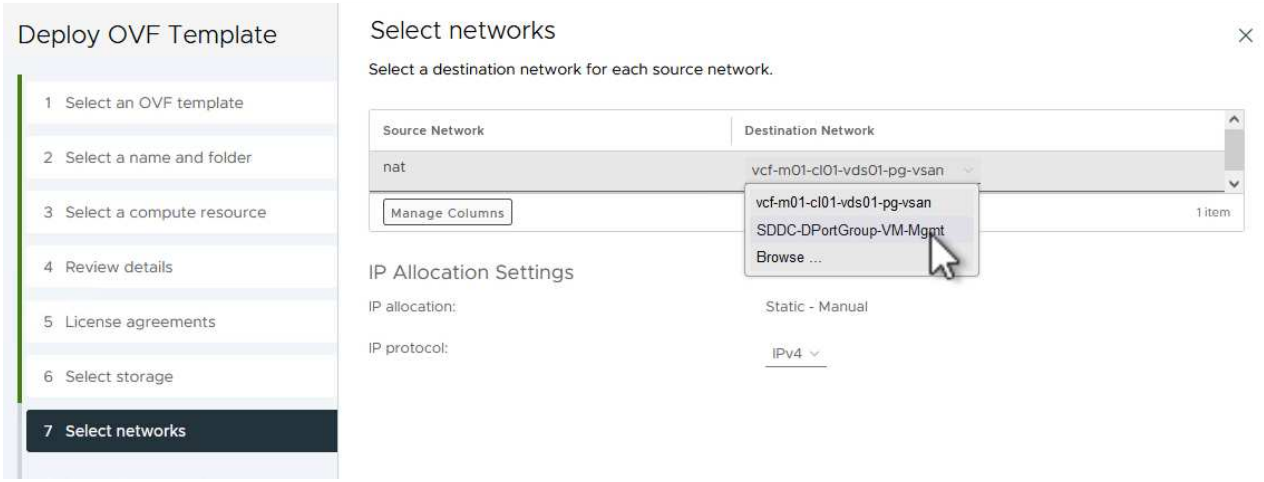

8. On the Customize template page fill out all required information:
   - Password to be used for administrative access to OTV.
   - NTP server IP address.
   - OTV maintenance account password.
   - OTV Derby DB password.
   - Do not check the box to **Enable VMware Cloud Foundation (VCF)**. VCF mode is not required for deploying supplemental storage.
   - FQDN or IP address of the vCenter appliance for the **VI Workload Domain**
   - Credentials for the vCenter appliance of the **VI Workload Domain**
   - Provide the required network properties fields.

     Click on **Next** to continue.

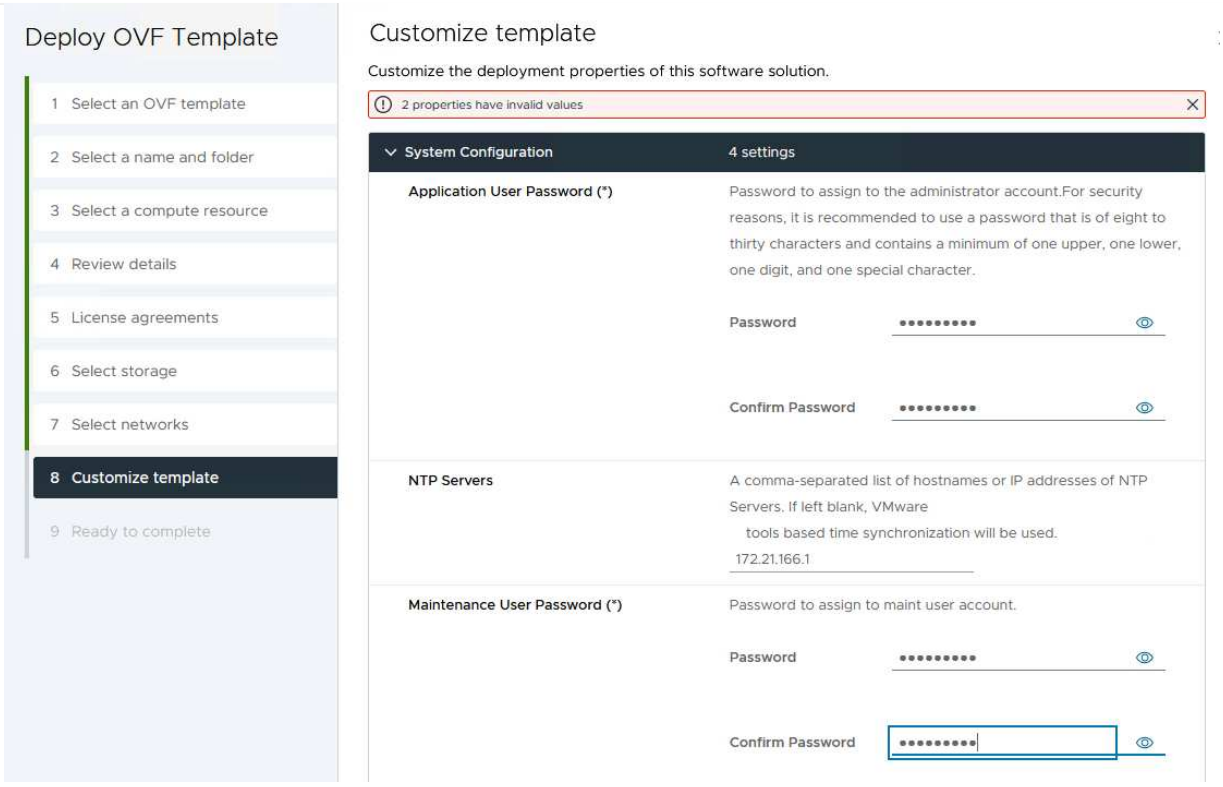

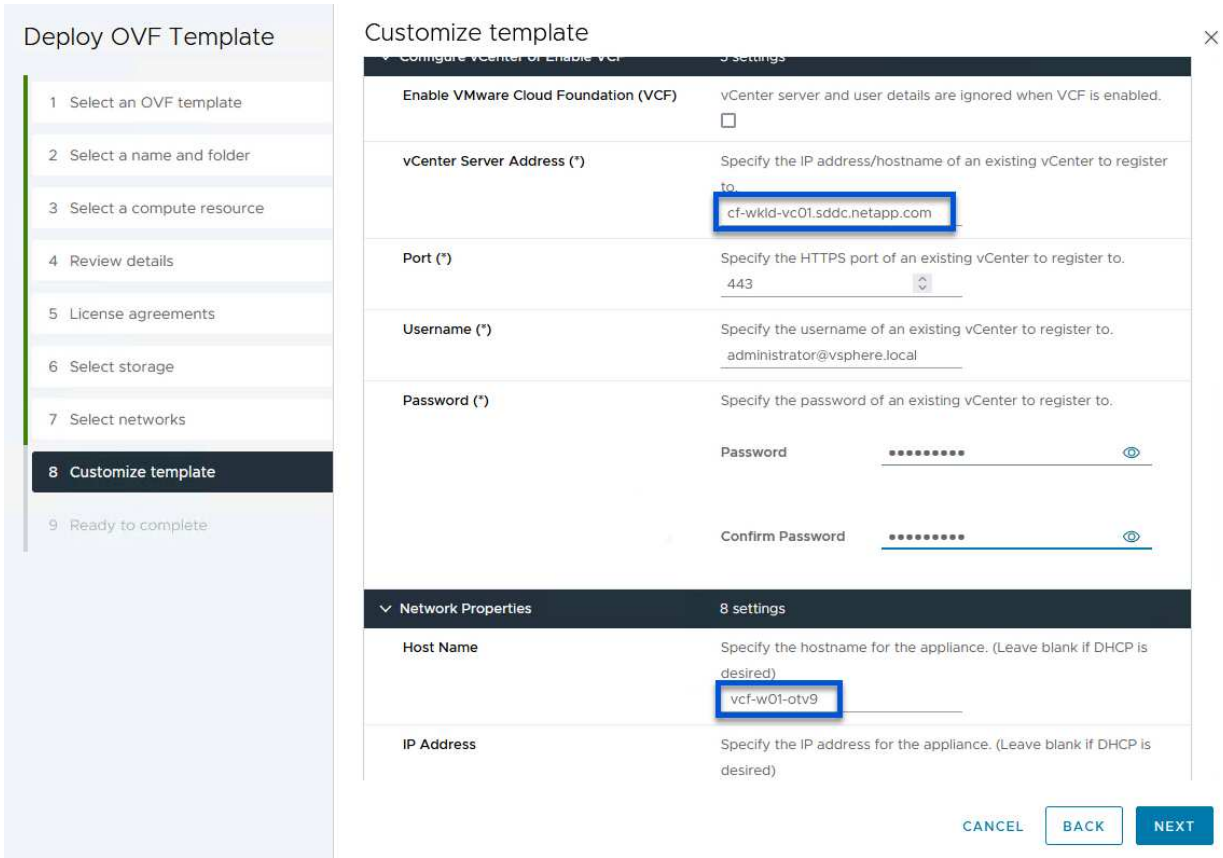

9. Review all information on the Ready to complete page and the click Finish to begin deploying the OTV appliance.

Add a storage system to ONTAP Tools.

1. Access NetApp ONTAP Tools by selecting it from the main menu in the vSphere client.

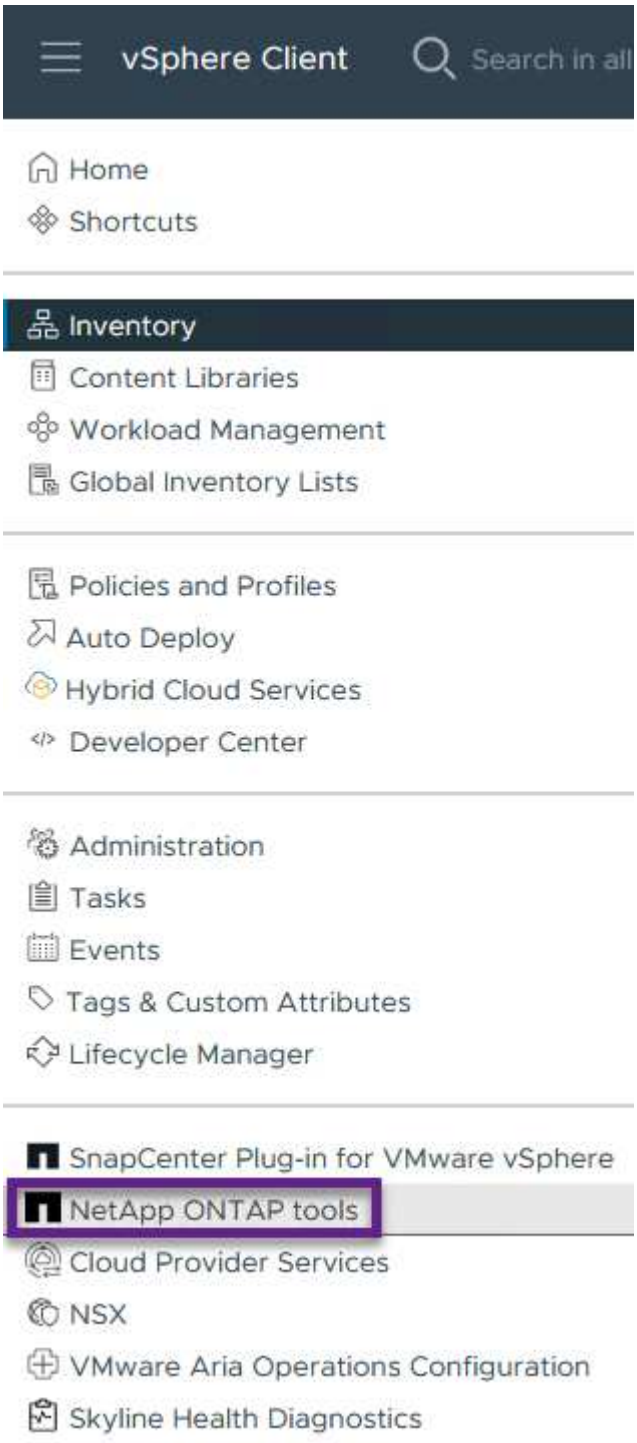

2. From the **INSTANCE** drop down menu in the ONTAP Tool interface, select the OTV instance associated with the workload domain to be managed.

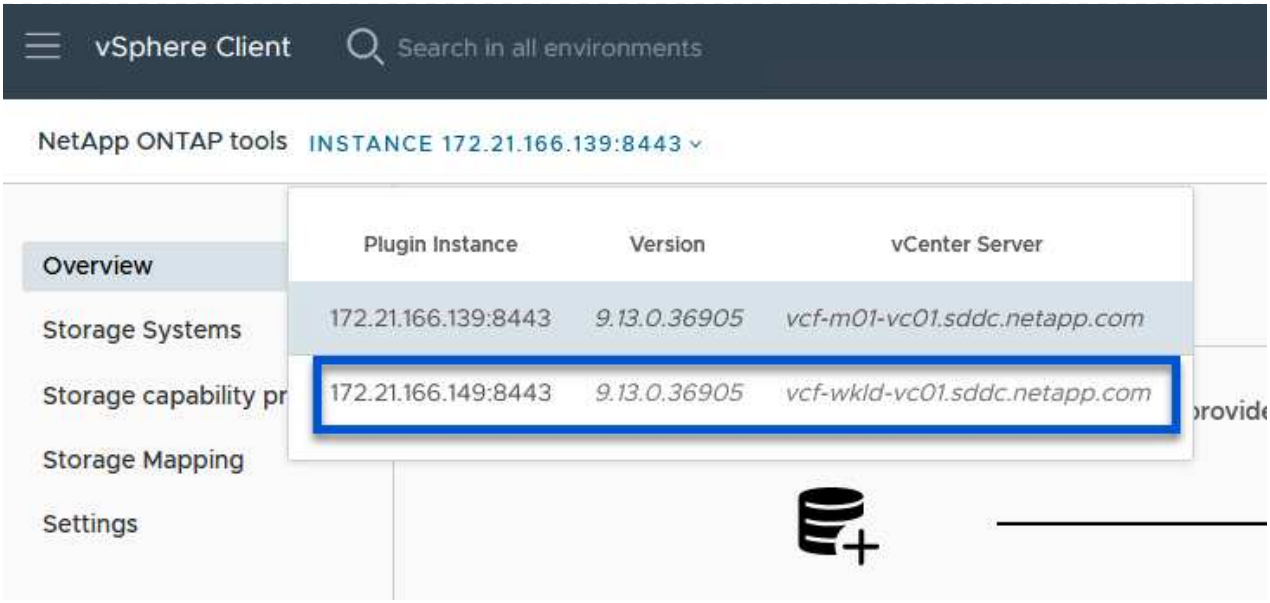

3. In ONTAP Tools select **Storage Systems** from the left hand menu and then press **Add**.

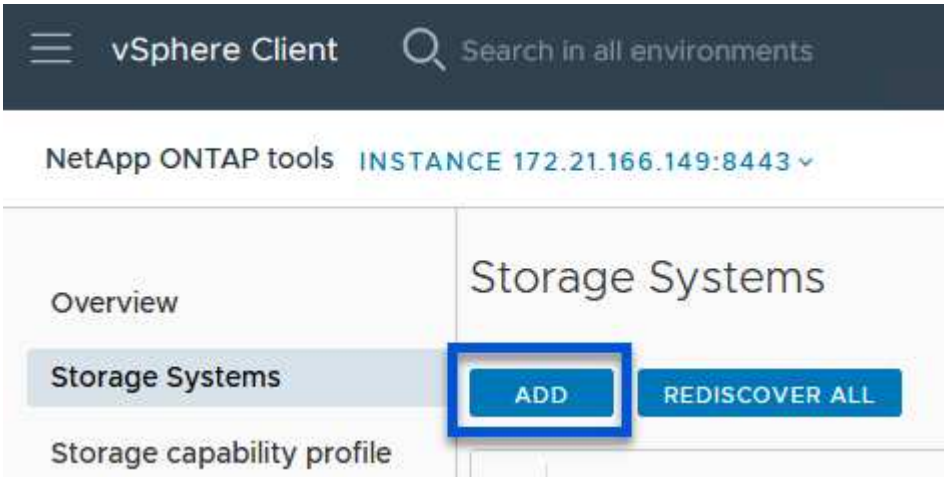

4. Fill out the IP Address, credentials of the storage system and the port number. Click on **Add** to start the discovery process.

> vVol requires ONTAP cluster credentials rather than SVM credentials. For more information refer to Add storage systems In the ONTAP Tools documentation.

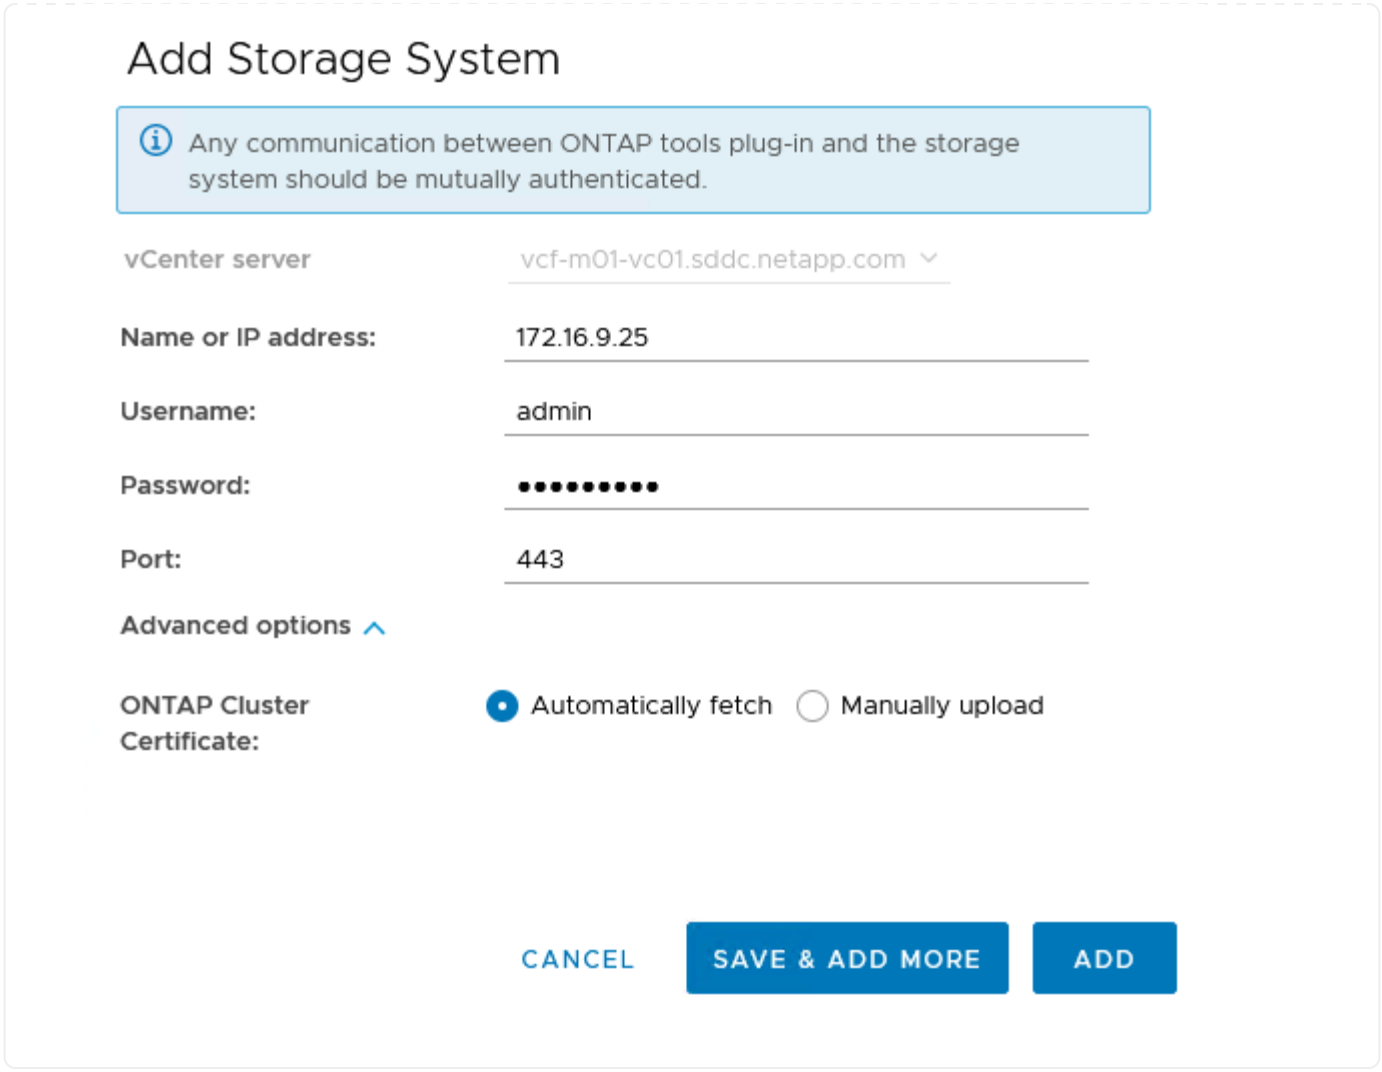
Add Storage System
Any communication between ONTAP tools plug-in and the storage system should be mutually authenticated.
vCenter server
vcf-m01-vc01.sddc.netapp.com
Name or IP address:
172.16.9.25
Username:
admin
Password:
Port:
443
Advanced options
ONTAP Cluster Certificate:
Automatically fetch
Manually upload
CANCEL
SAVE & ADD MORE
ADD

Create a storage capability profile in ONTAP Tools

Storage capability profiles describe the features provided by a storage array or storage system. They include quality of service definitions and are used to select storage systems that meet the parameters defined in the profile. One of the provided profiles can be used or new ones can be created.

To create a storage capability profile in ONTAP Tools complete the following steps:

1. In ONTAP Tools select **Storage capability profile** from the left-hand menu and then press **Create**.

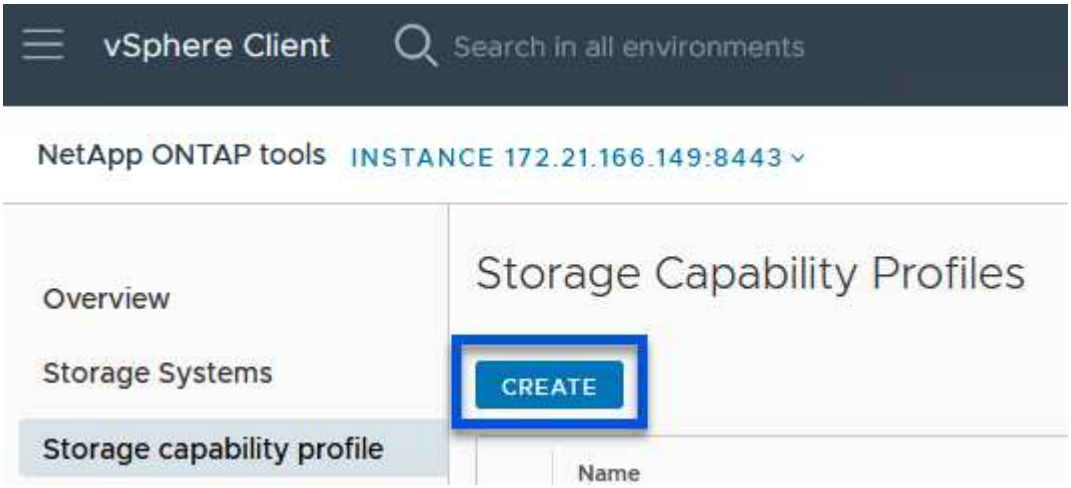

2. In the **Create Storage Capability profile** wizard provide a name and description of the profile and click on **Next**.

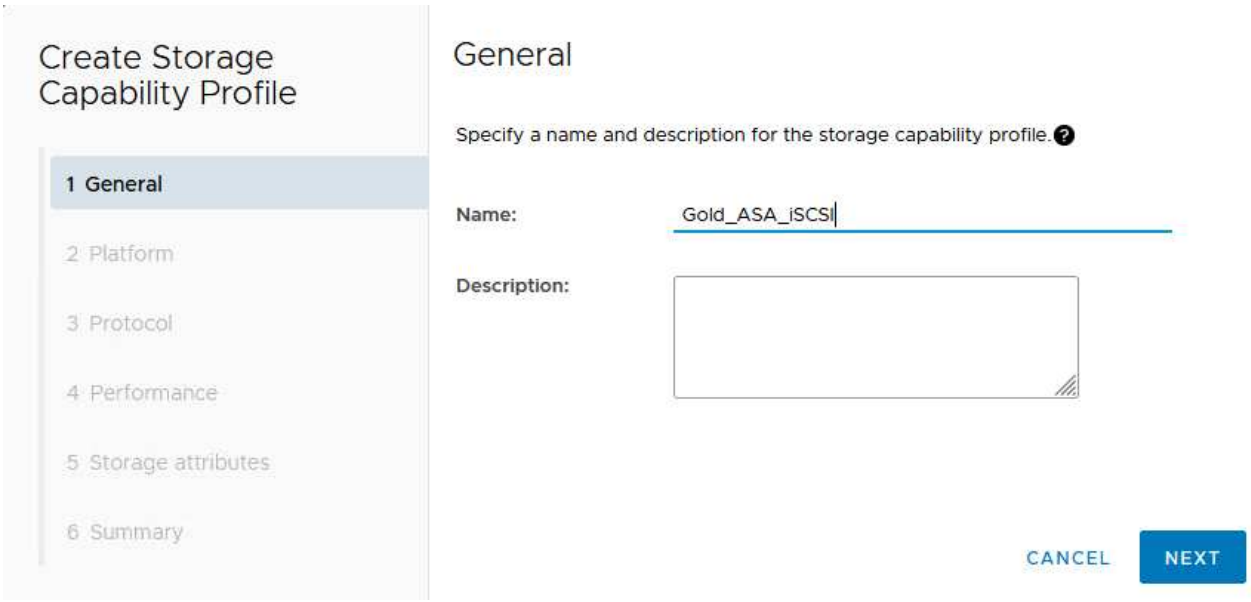

3. Select the platform type and to specify the storage system is to be an All-Flash SAN Array set **Asymmetric** to false.

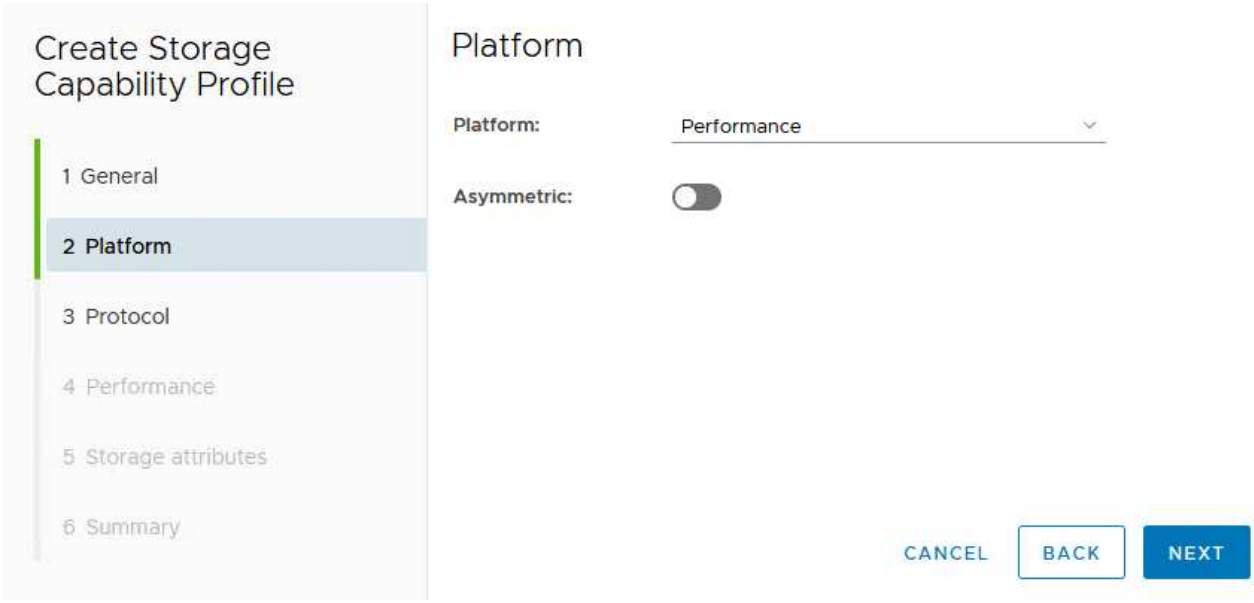

4. Next, select choice of protocol or **Any** to allow all possible protocols. Click **Next** to continue.

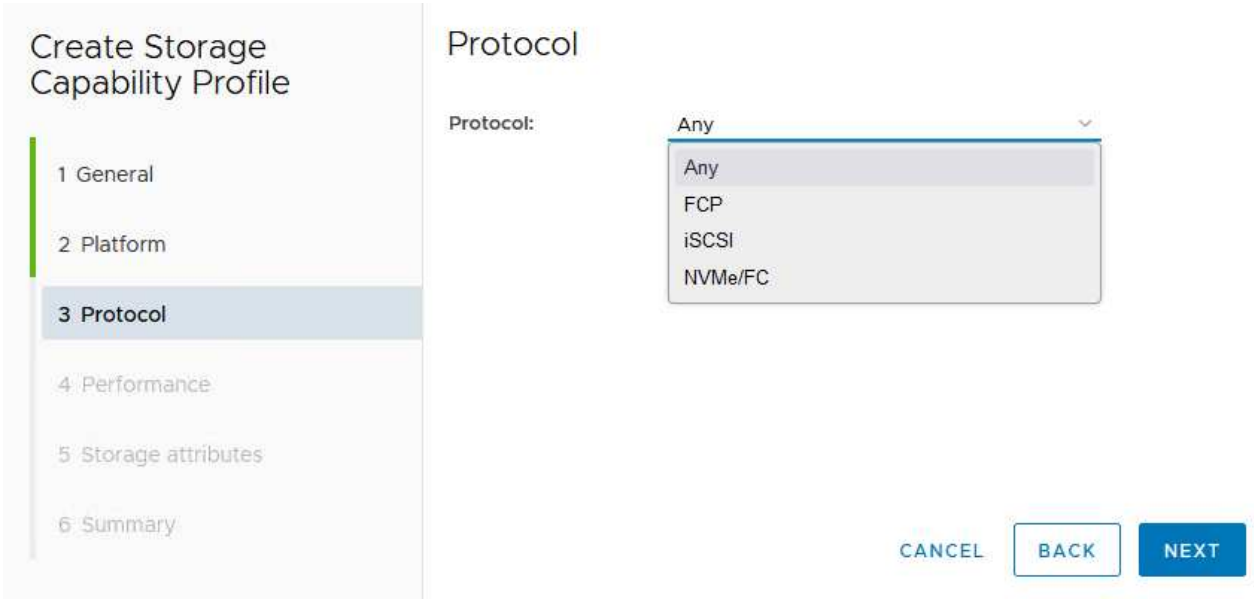

5. The **performance** page allows setting of quality of service in form of minimum and maximum IOPs allowed.

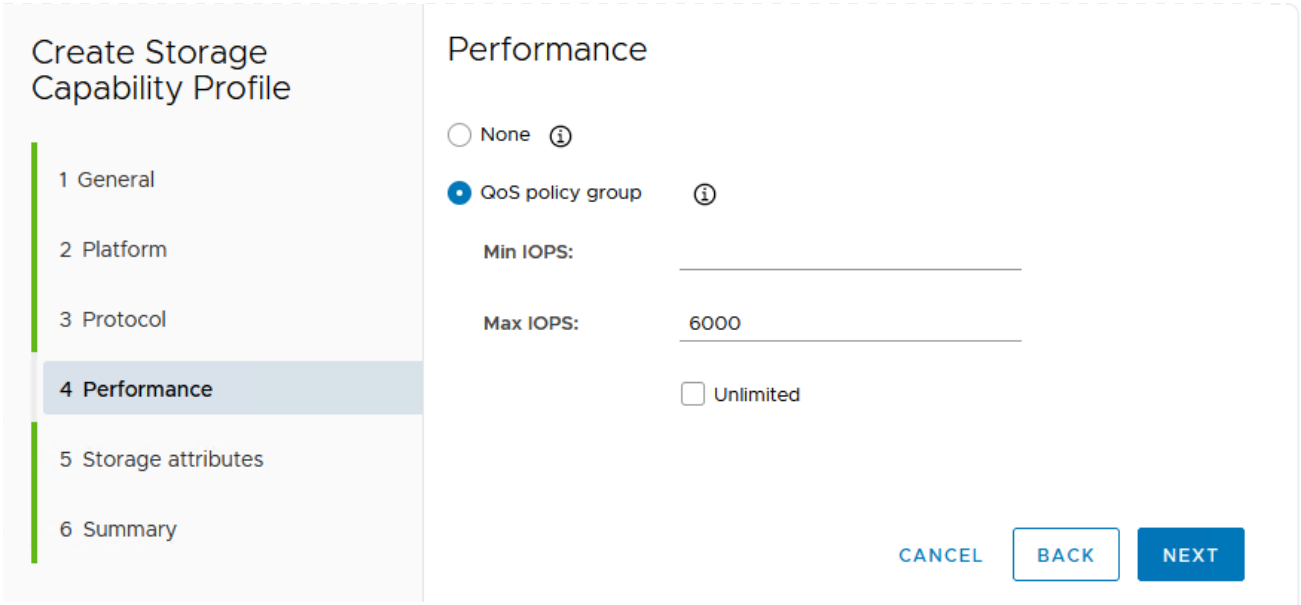

6. Complete the **storage attributes** page selecting storage efficiency, space reservation, encryption and any tiering policy as needed.

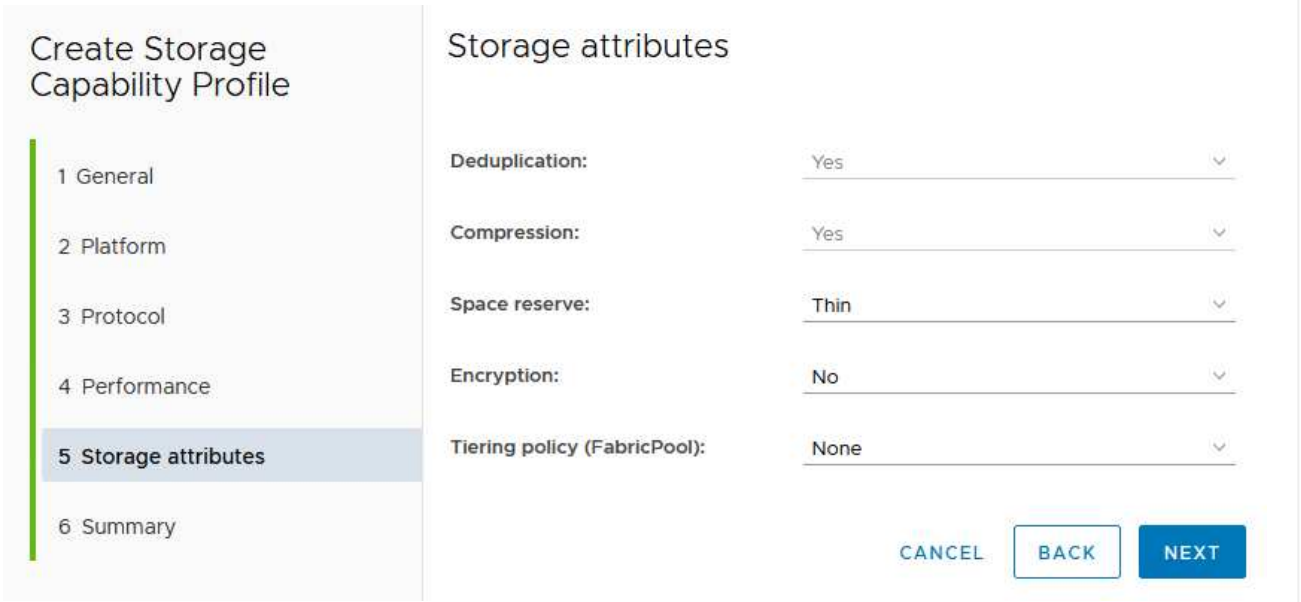

7. Finally, review the summary and click on Finish to create the profile.

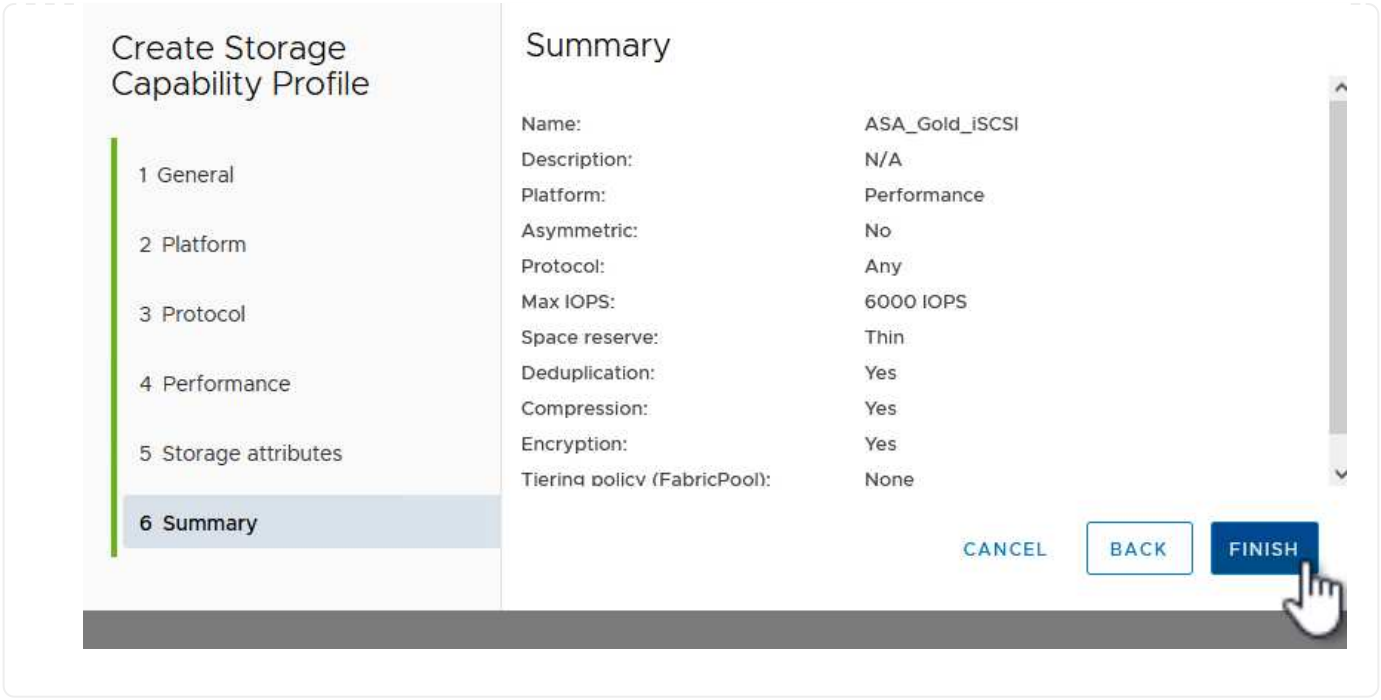
Create Storage Capability Profile
1 General
2 Platform
3 Protocol
4 Performance
5 Storage attributes
6 Summary
Summary
Name: ASA_Gold_iSCSI
Description: N/A
Platform: Performance
Asymmetric: No
Protocol: Any
Max IOPS: 6000 IOPS
Space reserve: Thin
Deduplication: Yes
Compression: Yes
Encryption: Yes
Tiering policy (FabricPool): None
CANCEL
BACK
FINISH

Create a vVols datastore in ONTAP Tools

To create a vVols datastore in ONTAP Tools complete the following steps:

1. In ONTAP Tools select **Overview** and from the **Getting Started** tab click on **Provision** to start the wizard.

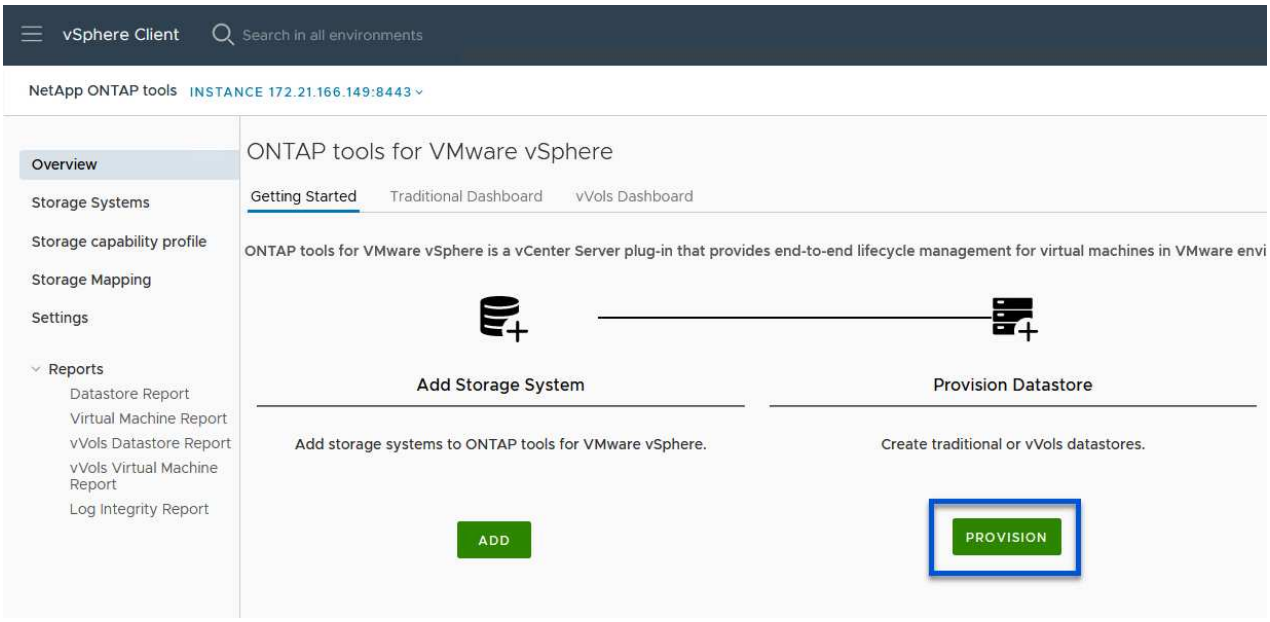

2. On the **General** page of the New Datastore wizard select the vSphere datacenter or cluster destination. Select **vVols** as the datastore type, fill out a name for the datastore, and select **iSCSI** as the protocol. Click on **Next** to continue.

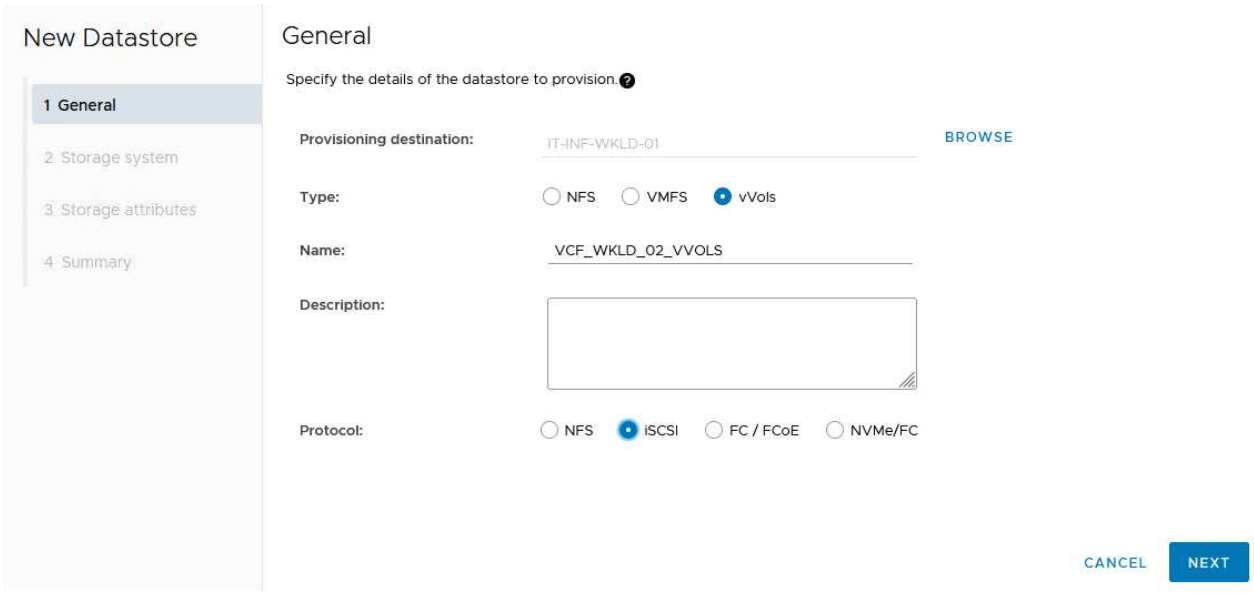

3. On the **Storage system** page select the select a storage capability profile, the storage system and SVM. Click on **Next** to continue.

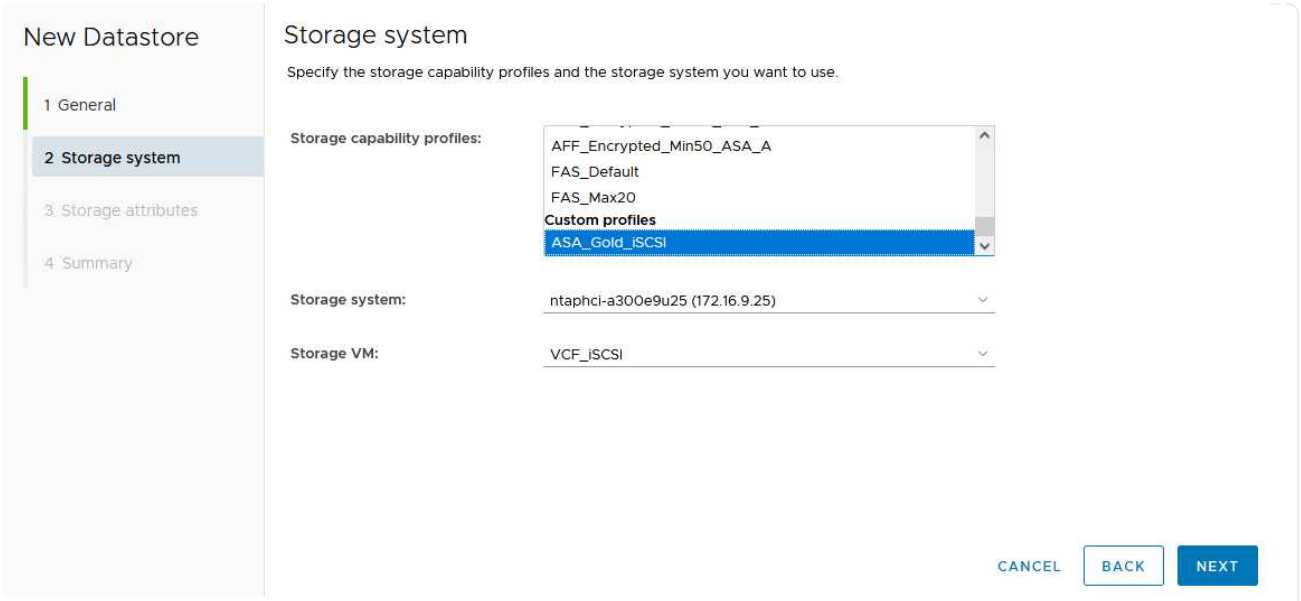

4. On the **Storage attributes** page select to create a new volume for the datastore and fill out the storage attributes of the volume to be created. Click on **Add** to create the volume and then **Next** to continue.

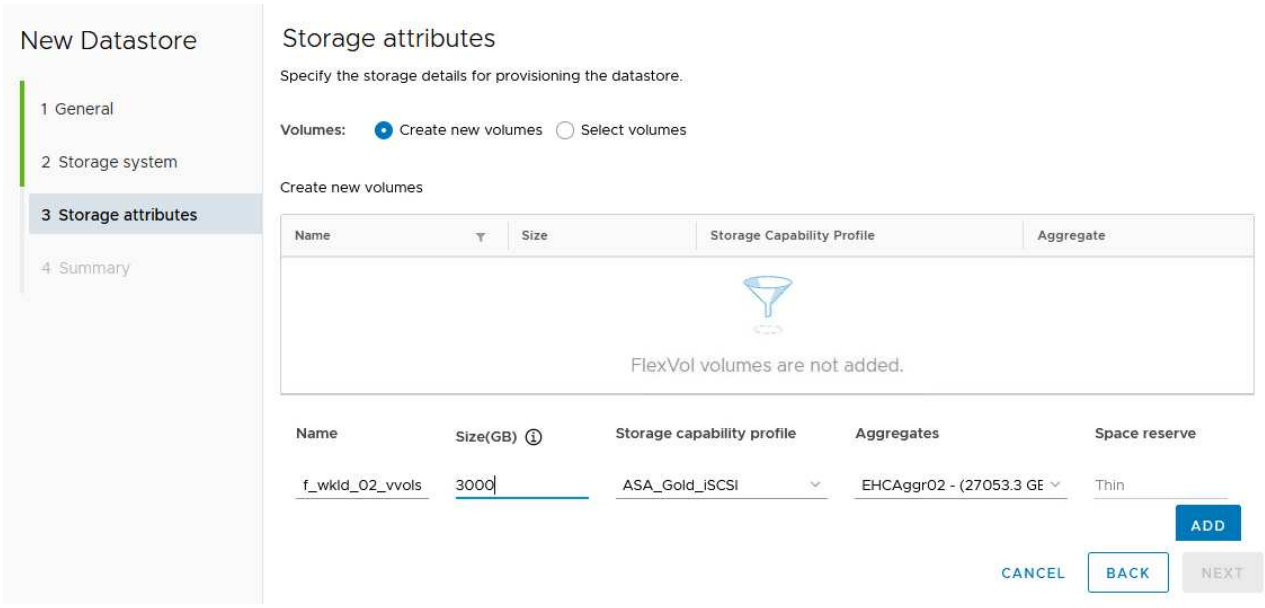

5. Finally, review the summary and click on **Finish** to start the vVol datastore creation process.

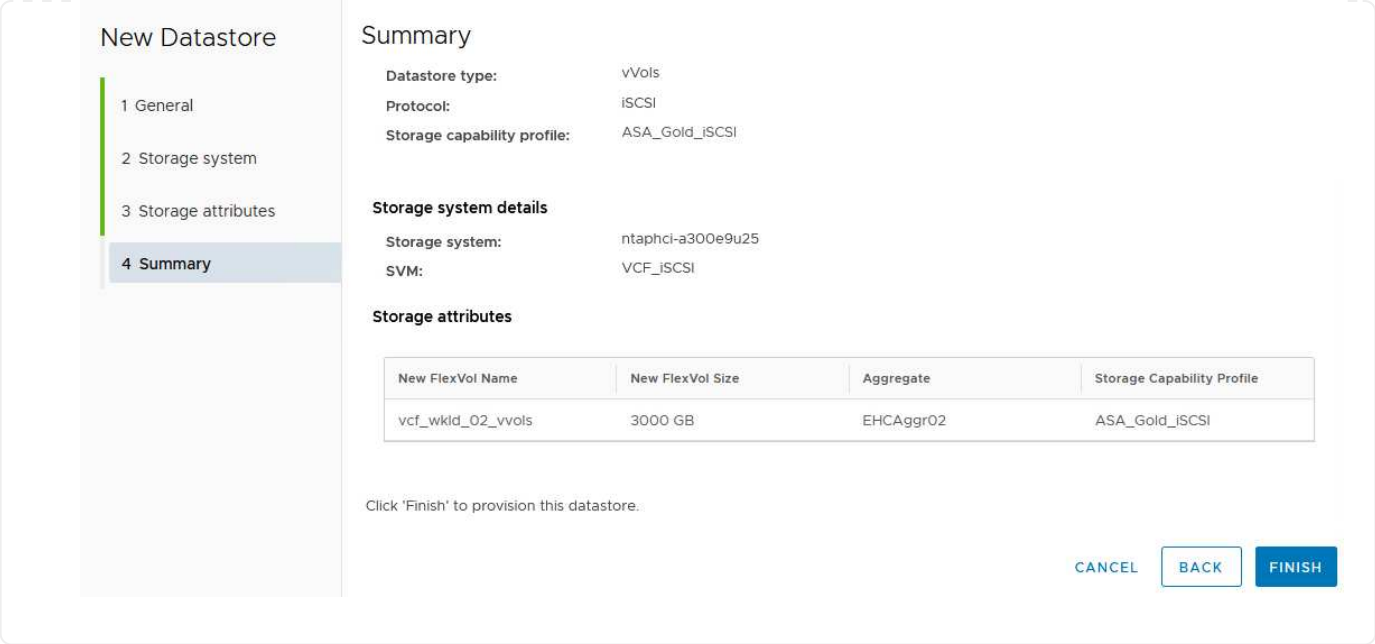

Additional information

For information on configuring ONTAP storage systems refer to the ONTAP 9 Documentation center.

For information on configuring VCF refer to VMware Cloud Foundation Documentation.

In this scenario we will demonstrate how to configure NVMe/TCP supplemental storage for a VCF workload domain.

Author: Josh Powell

Configure NVMe/TCP supplemental storage for VCF Workload Domains

Scenario Overview

This scenario covers the following high level steps:

- Create a storage virtual machine (SVM) with logical interfaces (LIFs) for NVMe/TCP traffic.
- Create distributed port groups for iSCSI networks on the VI workload domain.
- Create vmkernel adapters for iSCSI on the ESXi hosts for the VI workload domain.
- Add NVMe/TCP adapters on ESXi hosts.
- Deploy NVMe/TCP datastore.

Prerequisites

This scenario requires the following components and configurations:

- An ONTAP ASA storage system with physical data ports on ethernet switches dedicated to storage traffic.
- VCF management domain deployment is complete and the vSphere client is accessible.
- A VI workload domain has been previously deployed.

NetApp recommends fully redundant network designs for NVMe/TCP. The following diagram illustrates an

example of a redundant configuration, providing fault tolerance for storage systems, switches, networks adapters and host systems. Refer to the NetApp SAN configuration reference for additional information.

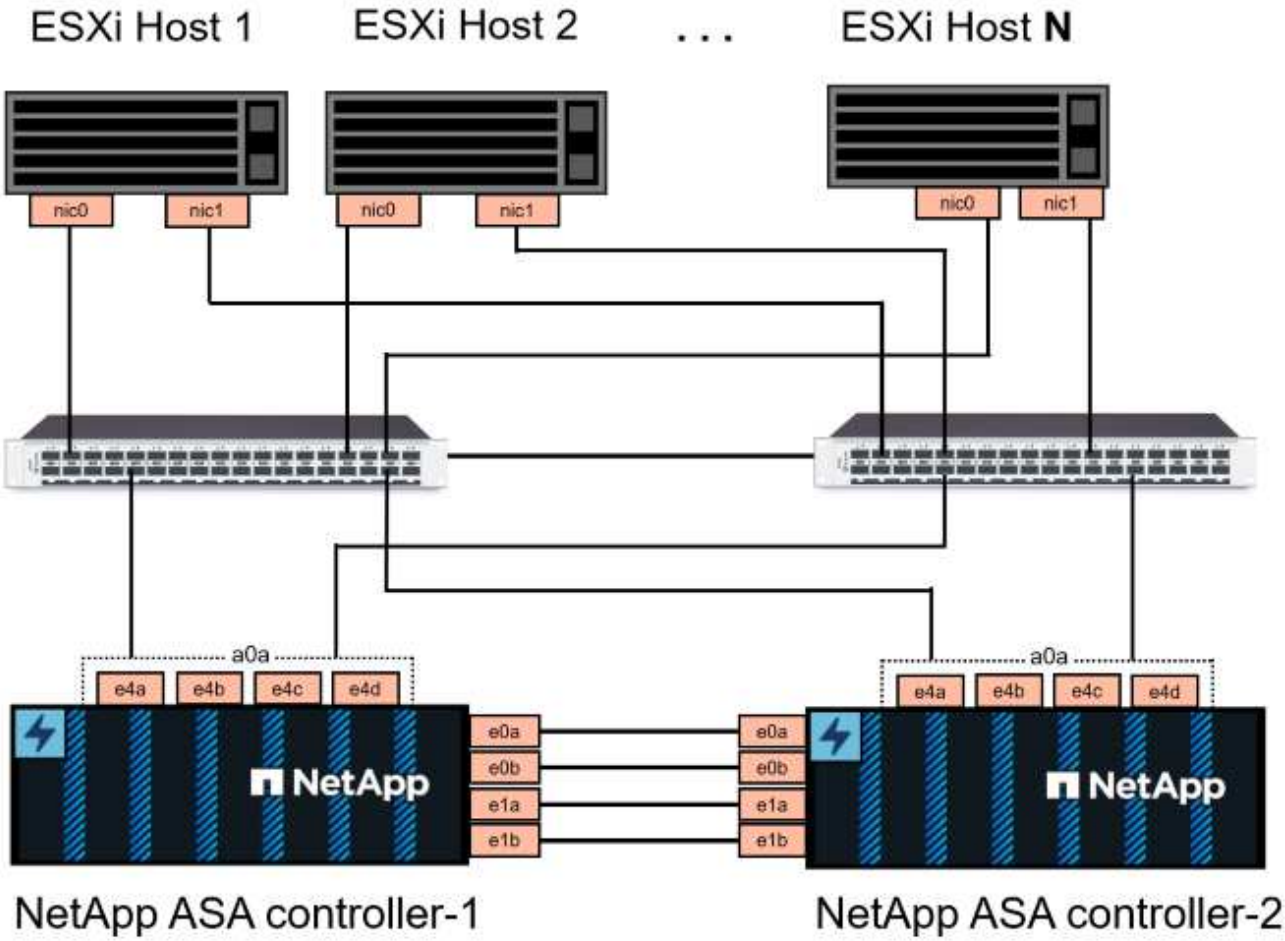

For multipathing and failover across multiple paths, NetApp recommends having a minimum of two LIFs per storage node in separate ethernet networks for all SVMs in NVMe/TCP configurations.

This documentation demonstrates the process of creating a new SVM and specifying the IP address information to create multiple LIFs for NVMe/TCP traffic. To add new LIFs to an existing SVM refer to Create a LIF (network interface).

For additional information on NVMe design considerations for ONTAP storage systems, refer to NVMe configuration, support and limitations.

Deployment Steps

To create a VMFS datastore on a VCF workload domain using NVMe/TCP, complete the following steps.

Create SVM, LIFs and NVMe Namespace on ONTAP storage system

The following step is performed in ONTAP System Manager.

Create the storage VM and LIFs

Complete the following steps to create an SVM together with multiple LIFs for NVMe/TCP traffic.

1. From ONTAP System Manager navigate to **Storage VMs** in the left-hand menu and click on **+ Add** to start.

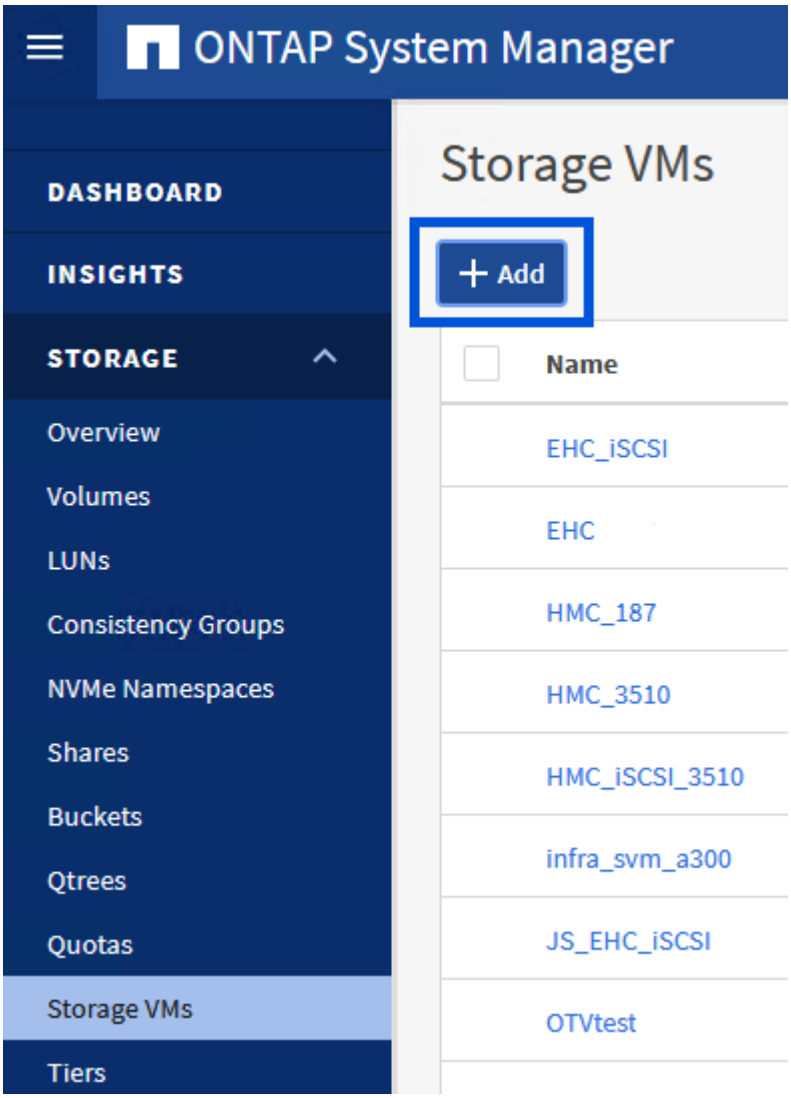

2. In the **Add Storage VM** wizard provide a **Name** for the SVM, select the **IP Space** and then, under **Access Protocol**, click on the **NVMe** tab and check the box to **Enable NVMe/TCP**.

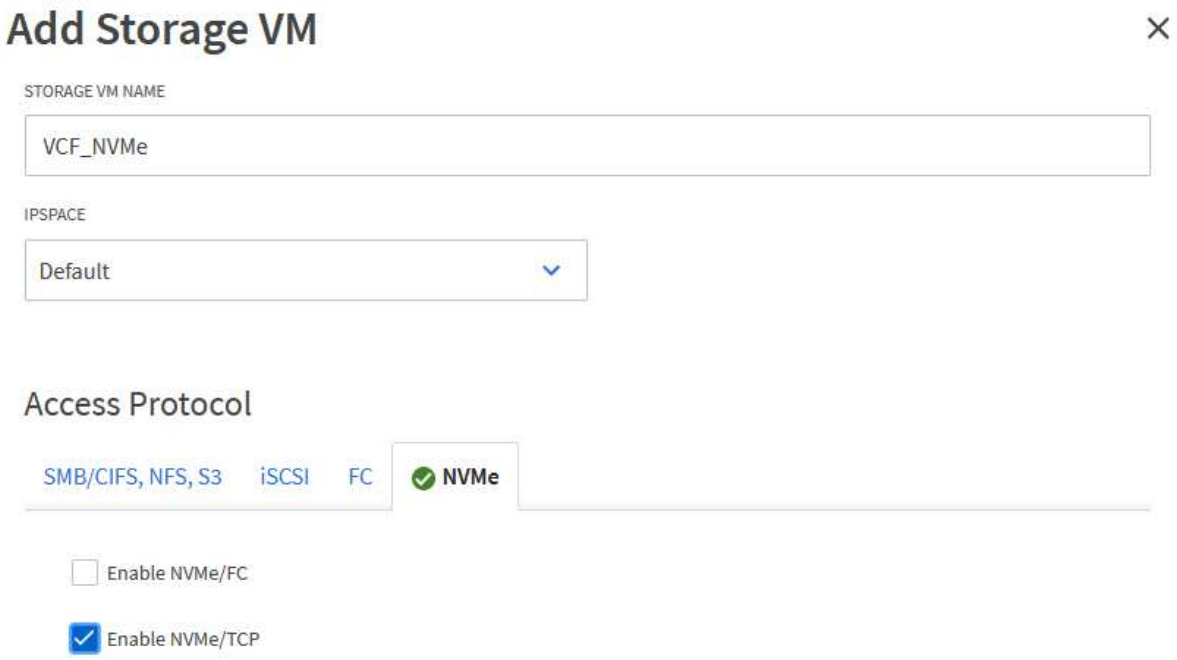

3. In the **Network Interface** section fill in the **IP address**, **Subnet Mask**, and **Broadcast Domain and Port** for the first LIF. For subsequent LIFs the checkbox may be enabled to use common settings across all remaining LIFs, or use separate settings.

> For multipathing and failover across multiple paths, NetApp recommends having a minimum of two LIFs per storage node in separate Ethernet networks for all SVMs in NVMe/TCP configurations.

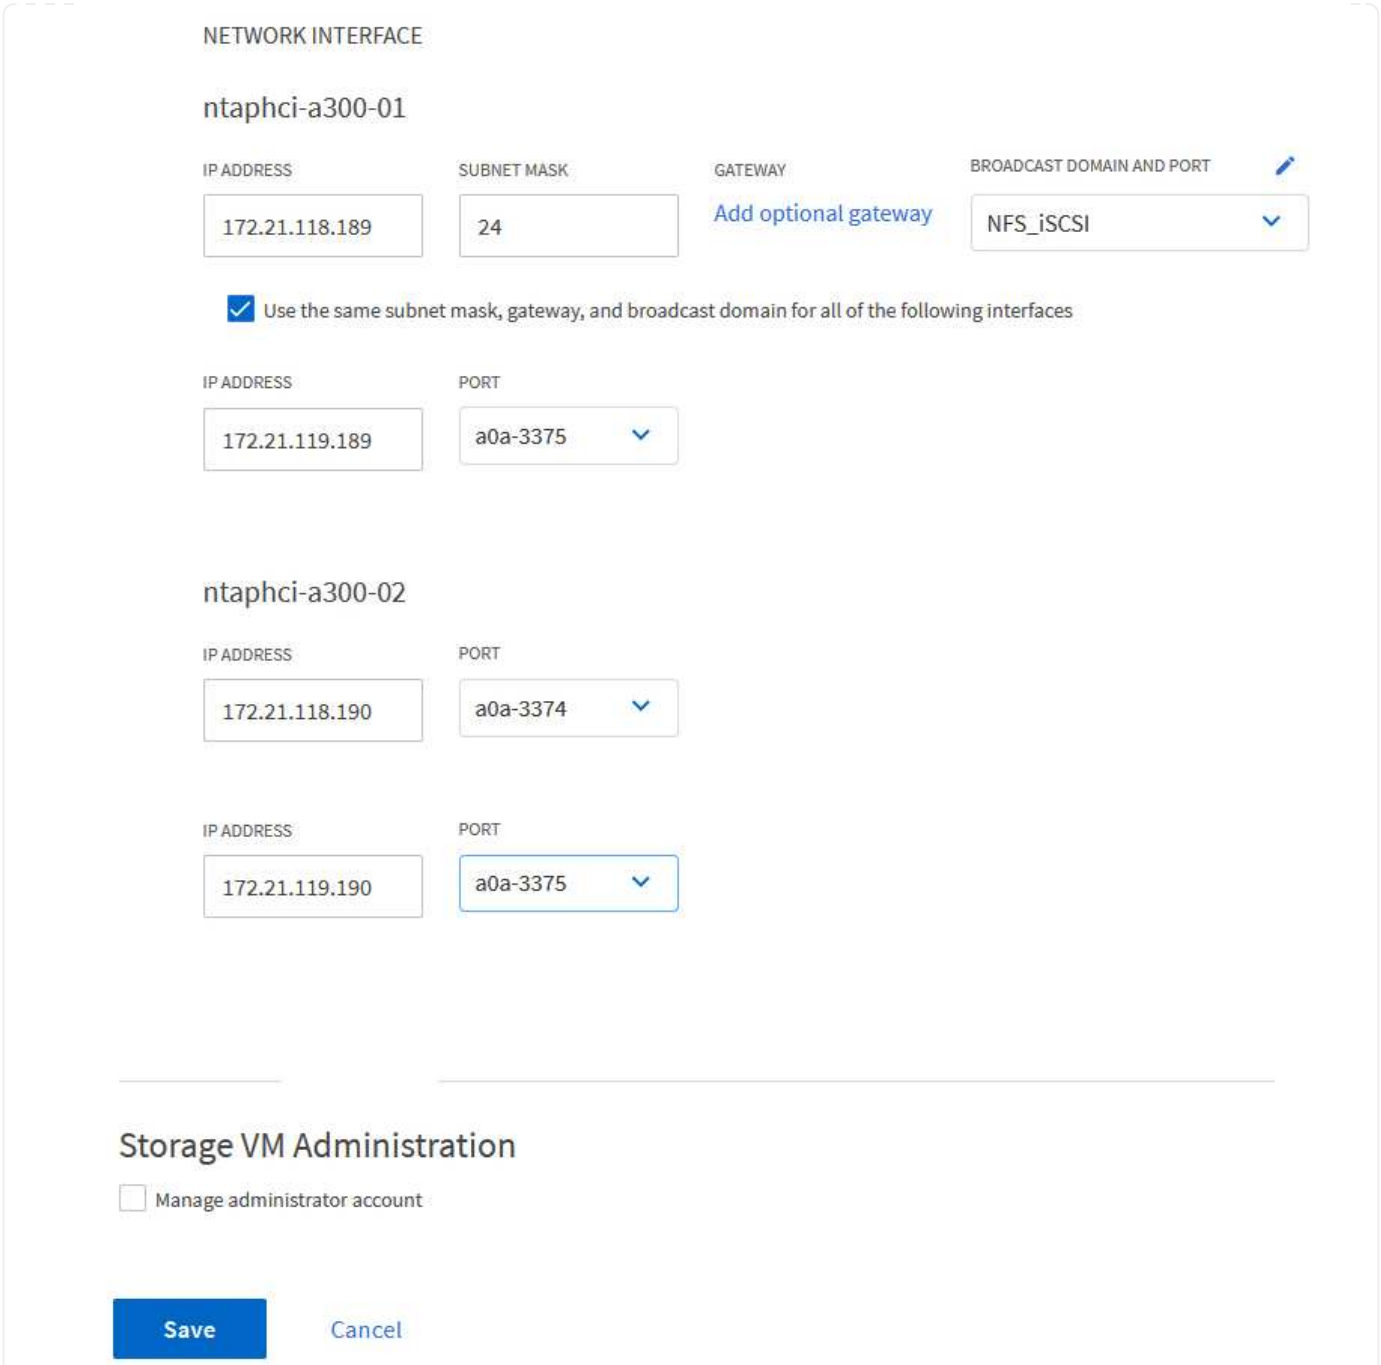

4. Choose whether to enable the Storage VM Administration account (for multi-tenancy environments) and click on **Save** to create the SVM.

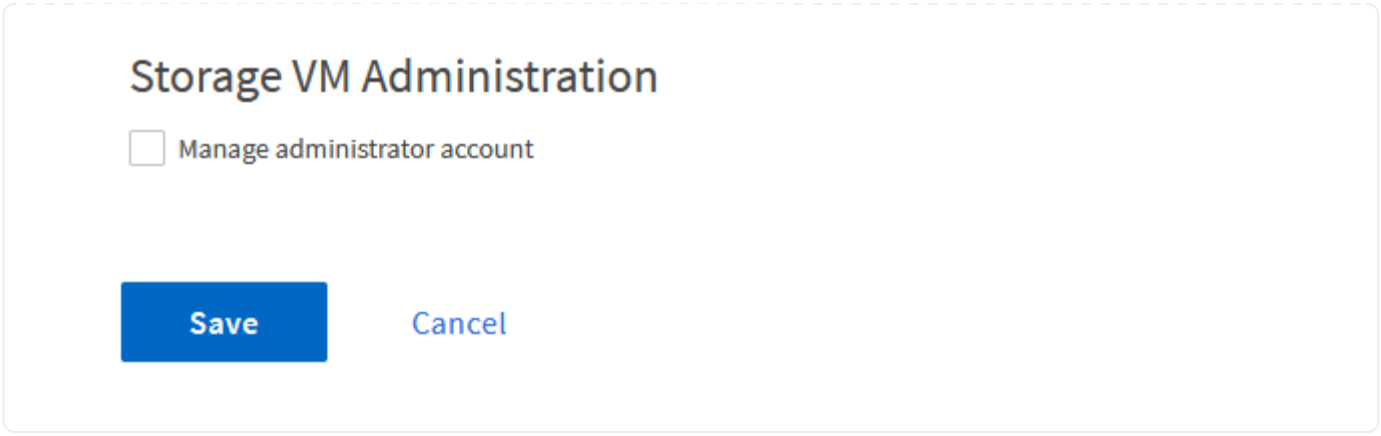
Storage VM Administration
Manage administrator account
Save
Cancel

Create the NVMe Namespace

NVMe namespaces are analogous to LUNs for iSCSi or FC. The NVMe Namespace must be created before a VMFS datastore can be deployed from the vSphere Client. To create the NVMe namespace, the NVMe Qualified Name (NQN) must first be obtained from each ESXi host in the cluster. The NQN is used by ONTAP to provide access control for the namespace.

Complete the following steps to create an NVMe Namespace:

1. Open an SSH session with an ESXi host in the cluster to obtain its NQN. Use the following command from the CLI:

   ```
   esxcli nvme info get
   ```

   An output similar to the following should be displayed:

   ```
   Host NQN: nqn.2014-08.com.netapp.sddc:nvme:vcf-wkld-esx01
   ```

2. Record the NQN for each ESXi host in the cluster
3. From ONTAP System Manager navigate to **NVMe Namespaces** in the left-hand menu and click on **+ Add** to start.

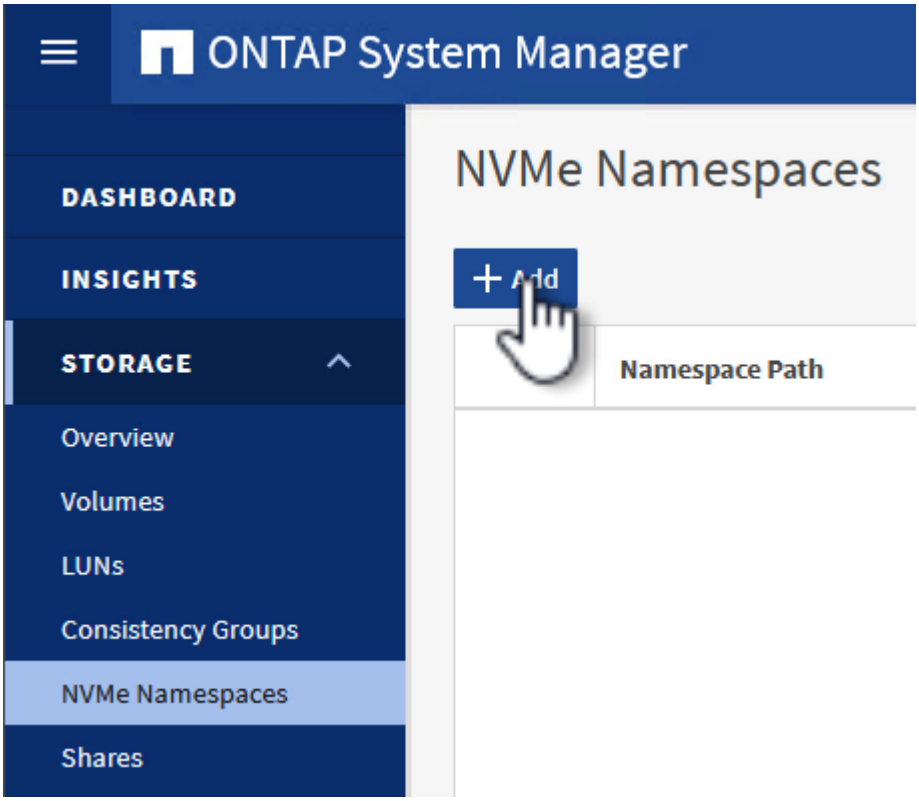

4. On the **Add NVMe Namespace** page, fill in a name prefix, the number of namespaces to create, the size of the namespace, and the host operating system that will be accessing the namespace. In the **Host NQN** section create a comma separated list of the NQN's previously collected from the ESXi

hosts that will be accessing the namespaces.

Click on **More Options** to configure additional items such as the snapshot protection policy. Finally, click on **Save** to create the NVMe Namespace.

+
image::vmware-vcf-asa-image93.png[Click +Add to create NVMe Namespace]

Set up networking and NVMe software adapters on ESXi hosts

The following steps are performed on the VI workload domain cluster using the vSphere client. In this case vCenter Single Sign-On is being used so the vSphere client is common to both the management and workload domains.

Create Distributed Port Groups for NVME/TCP traffic

Complete the following to create a new distributed port group for each NVMe/TCP network:

1. From the vSphere client , navigate to **Inventory > Networking** for the workload domain. Navigate to the existing Distributed Switch and choose the action to create **New Distributed Port Group…**.

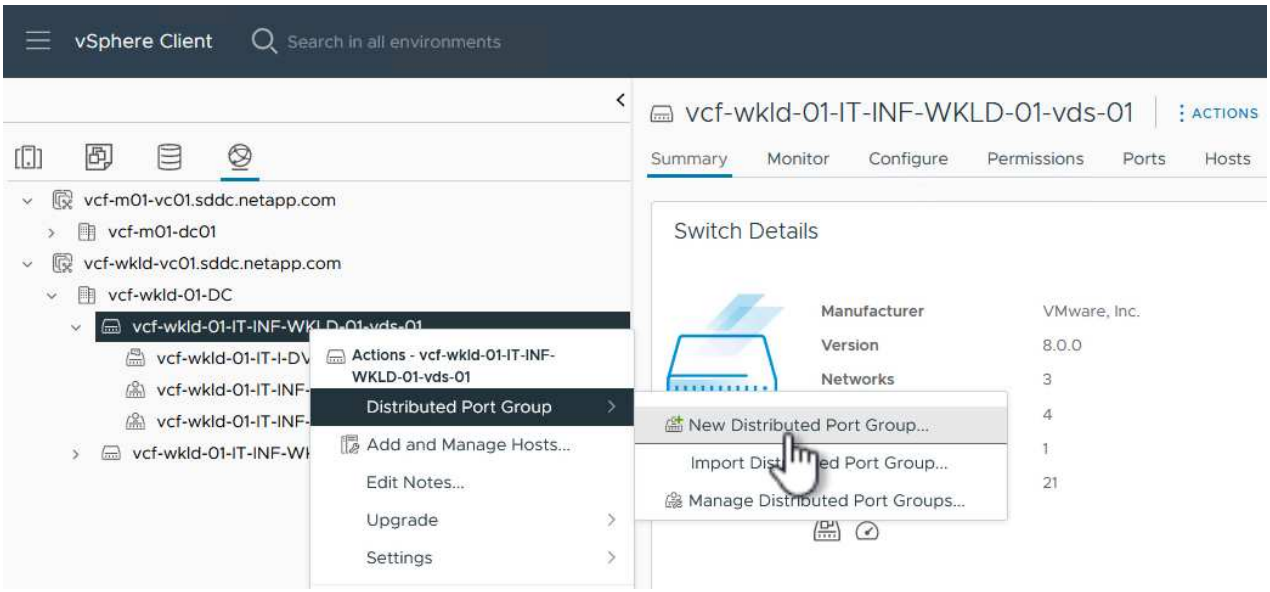

2. In the **New Distributed Port Group** wizard fill in a name for the new port group and click on **Next** to continue.
3. On the **Configure settings** page fill out all settings. If VLANs are being used be sure to provide the correct VLAN ID. Click on **Next** to continue.

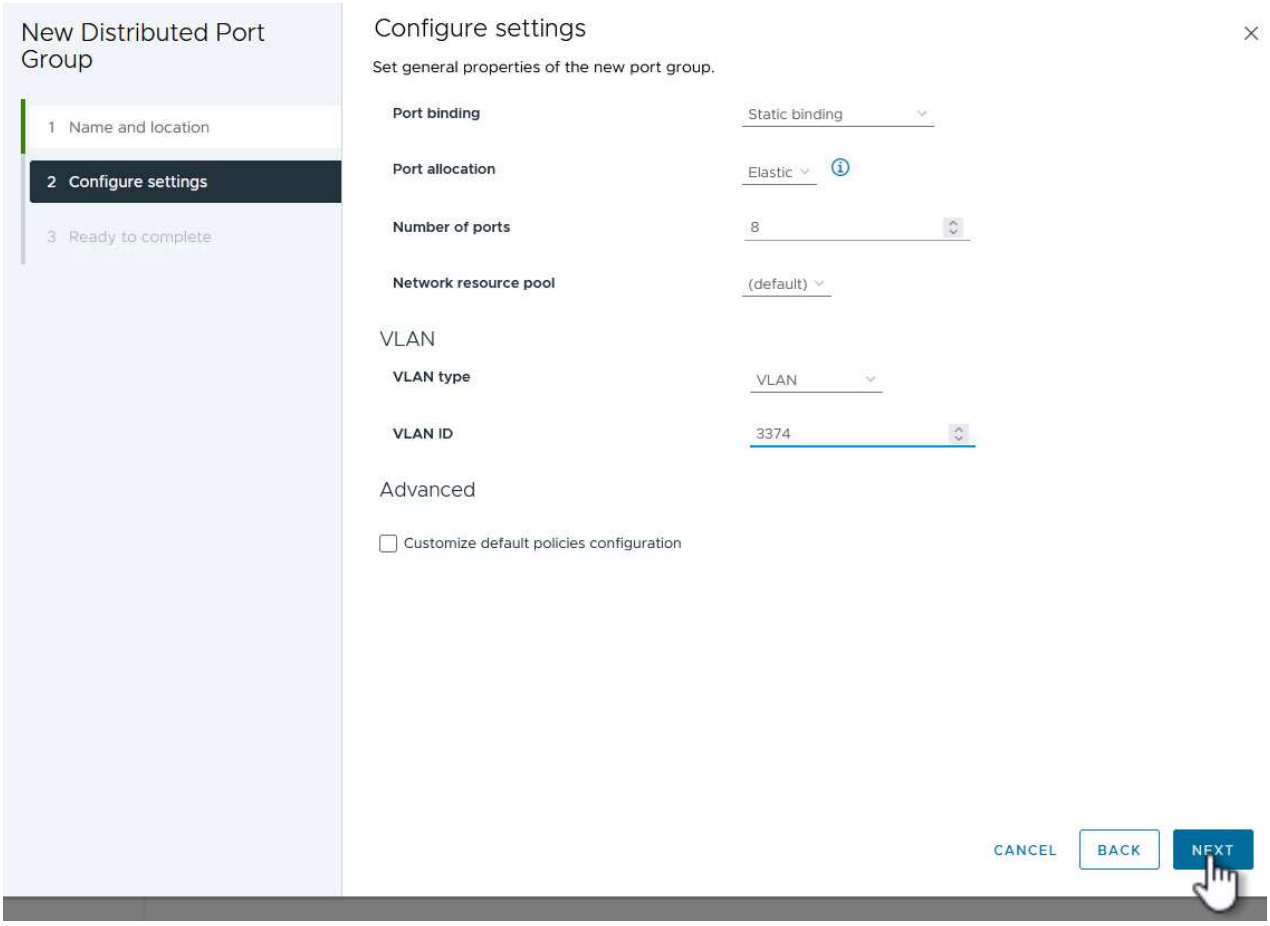

4. On the **Ready to complete** page, review the changes and click on **Finish** to create the new distributed port group.
5. Repeat this process to create a distributed port group for the second NVMe/TCP network being used and ensure you have input the correct **VLAN ID**.
6. Once both port groups have been created, navigate to the first port group and select the action to **Edit settings…**.

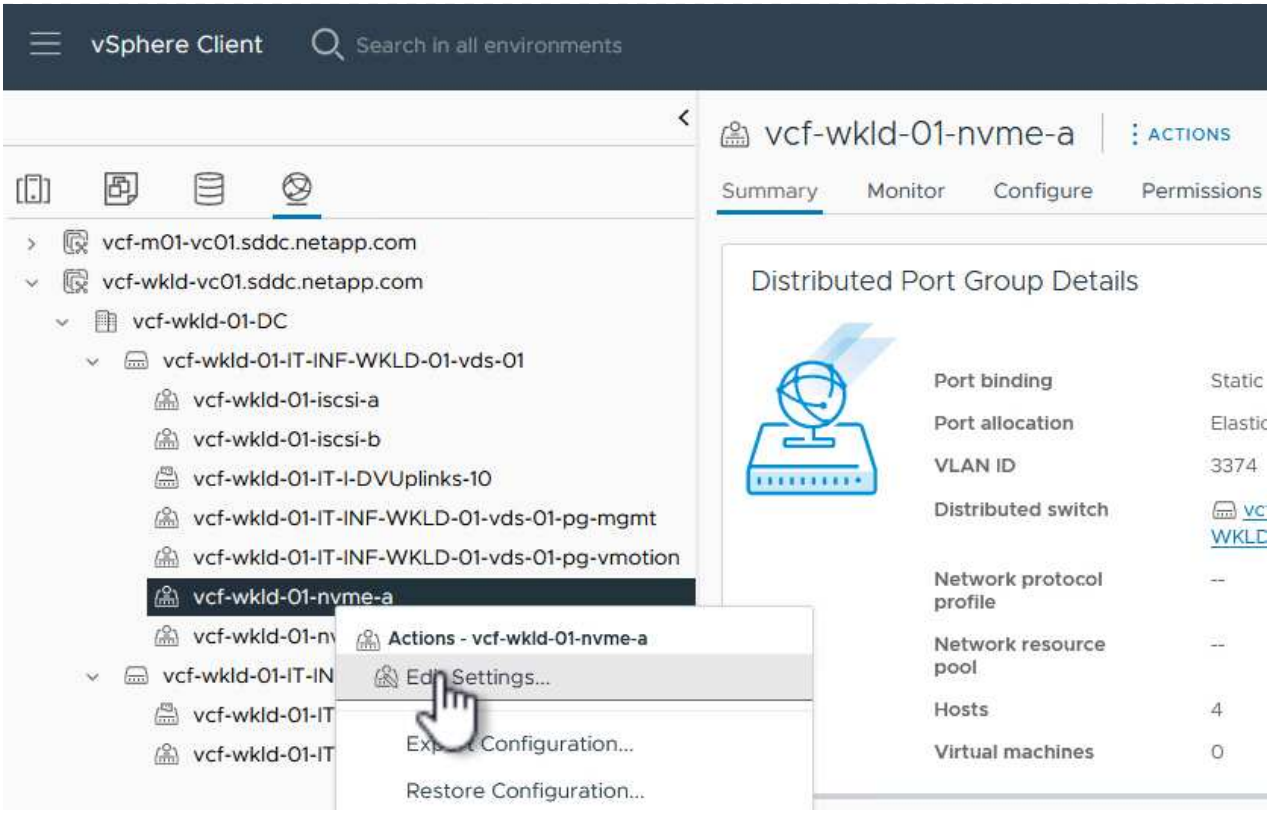

7. On **Distributed Port Group - Edit Settings** page, navigate to **Teaming and failover** in the left-hand menu and click on **uplink2** to move it down to **Unused uplinks**.

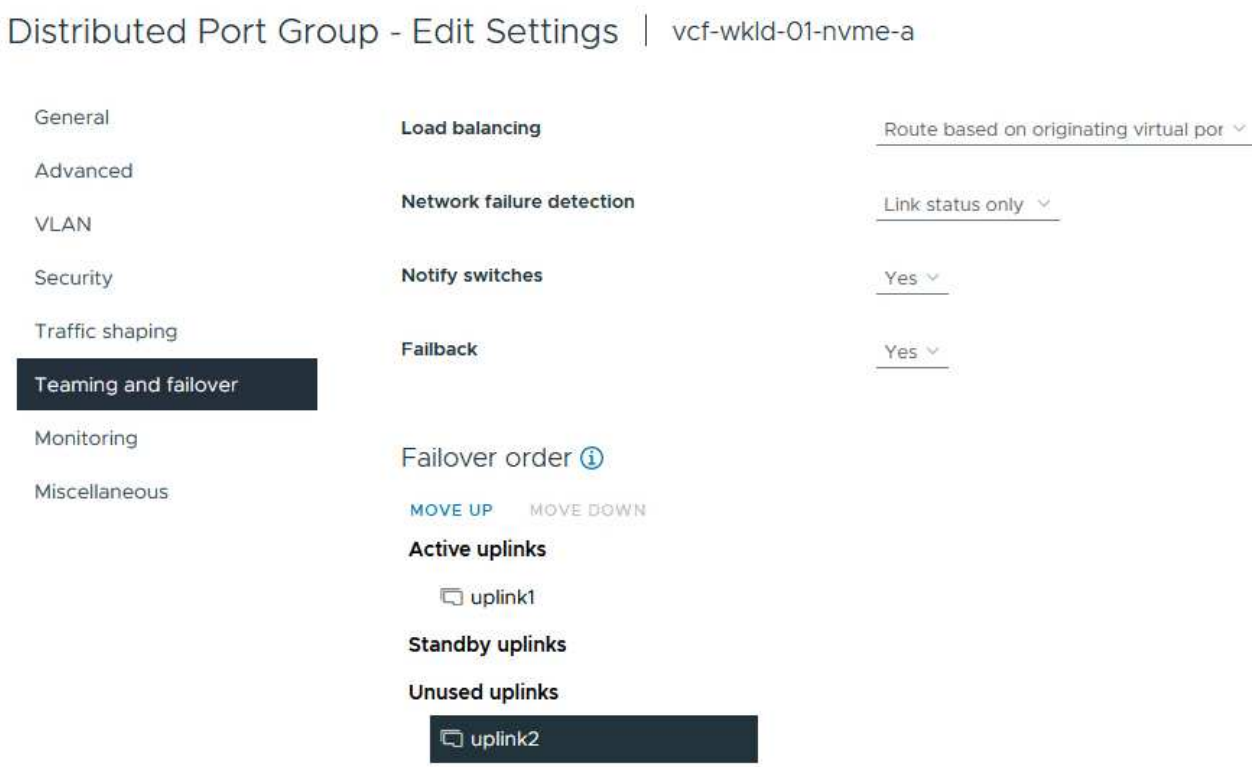

8. Repeat this step for the second NVMe/TCP port group. However, this time move **uplink1** down to

Unused uplinks.

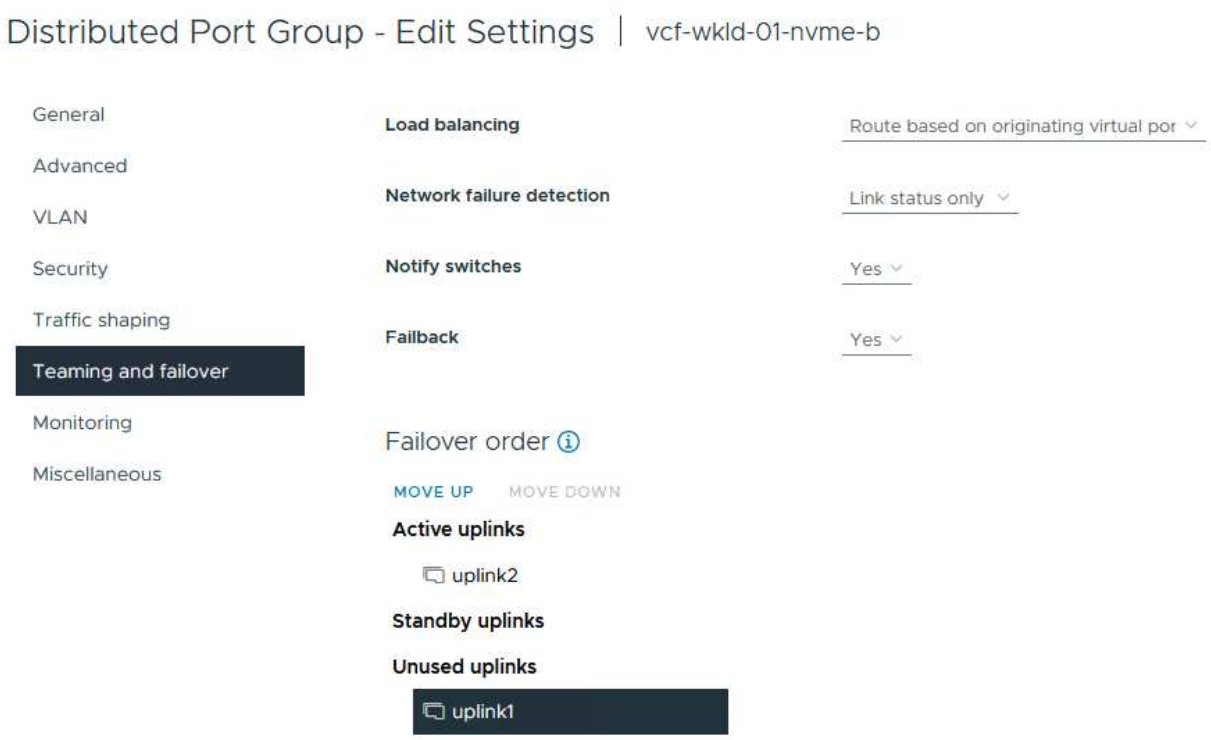

Create VMkernel adapters on each ESXi host

Repeat this process on each ESXi host in the workload domain.

1. From the vSphere client navigate to one of the ESXi hosts in the workload domain inventory. From the **Configure** tab select **VMkernel adapters** and click on **Add Networking…** to start.

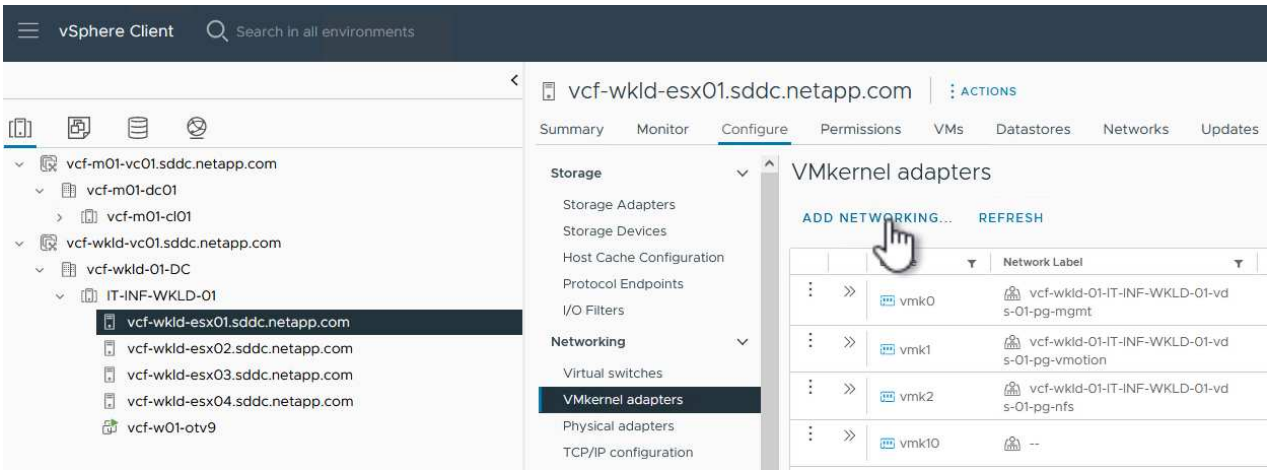

2. On the **Select connection type** window choose **VMkernel Network Adapter** and click on **Next** to continue.

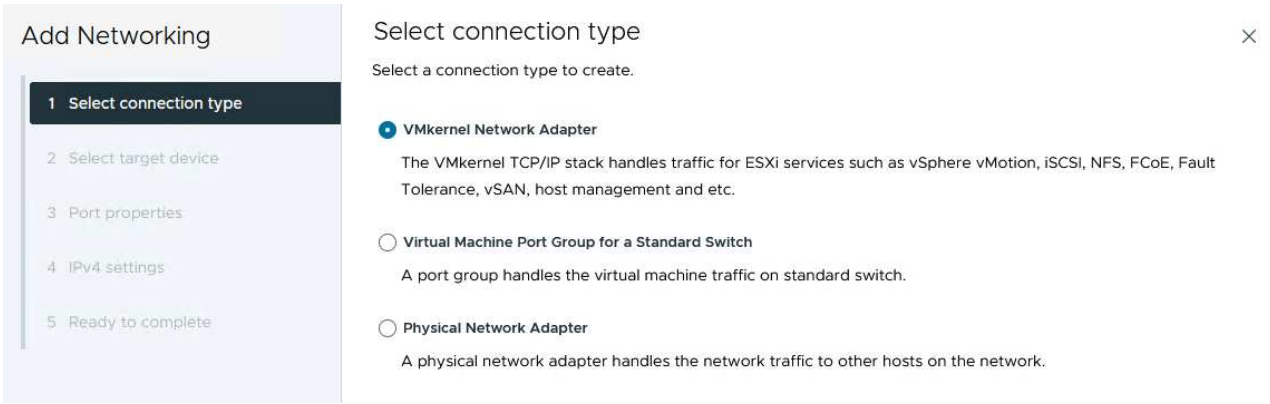

3. On the **Select target device** page, choose one of the distributed port groups for iSCSI that was created previously.

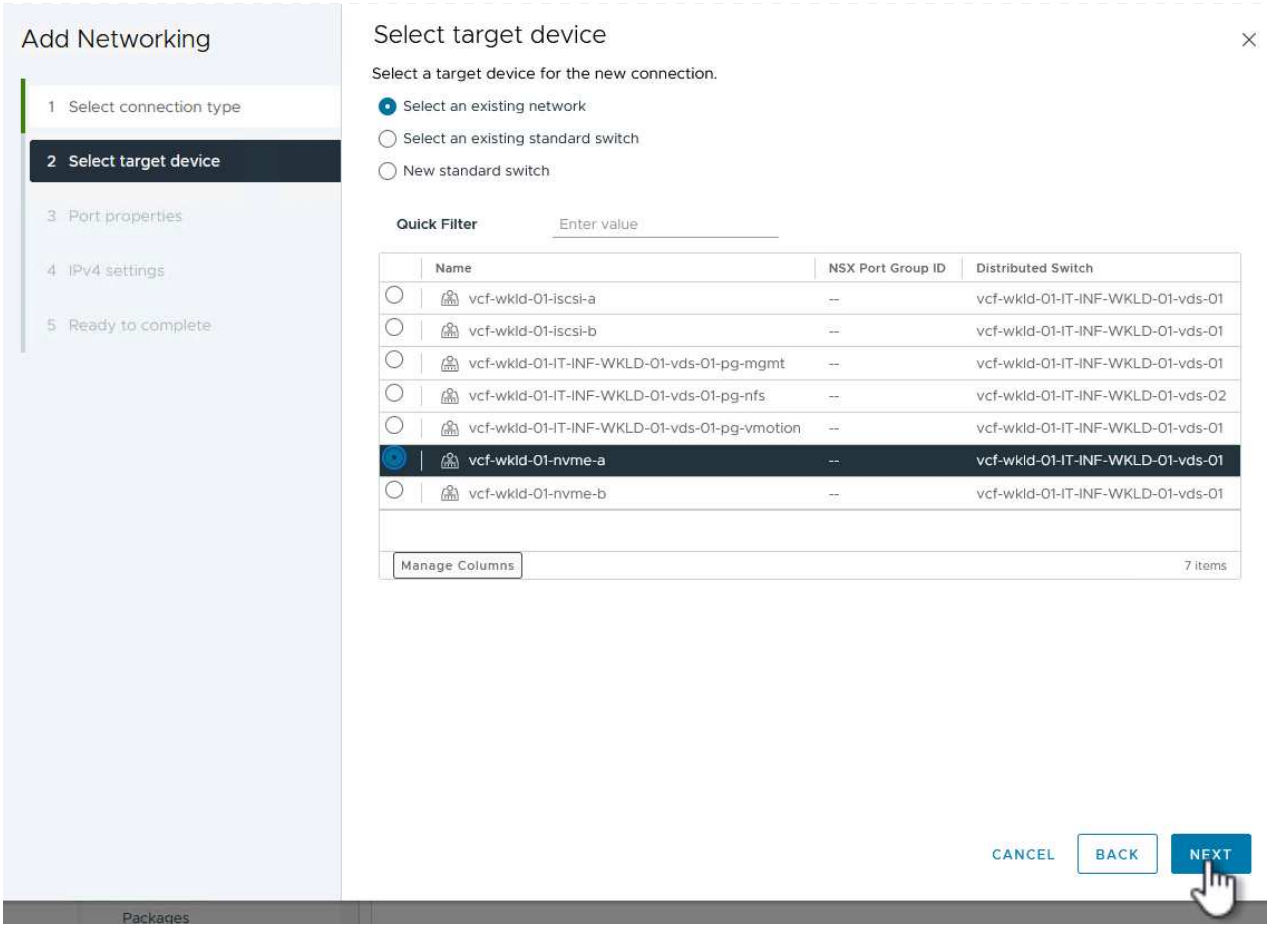

4. On the **Port properties** page click the box for **NVMe over TCP** and click on **Next** to continue.

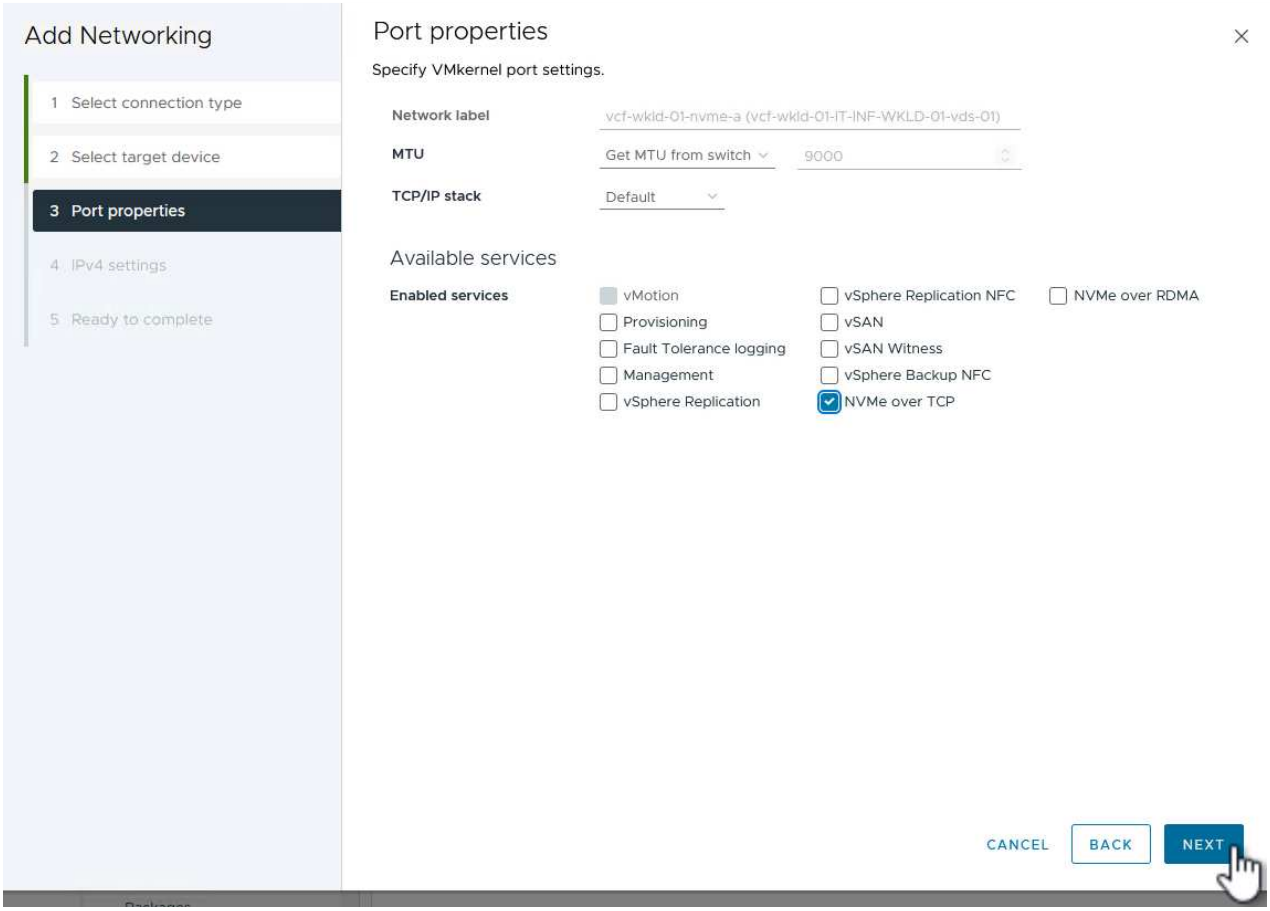

5. On the **IPv4 settings** page fill in the **IP address**, **Subnet mask**, and provide a new Gateway IP address (only if required). Click on **Next** to continue.

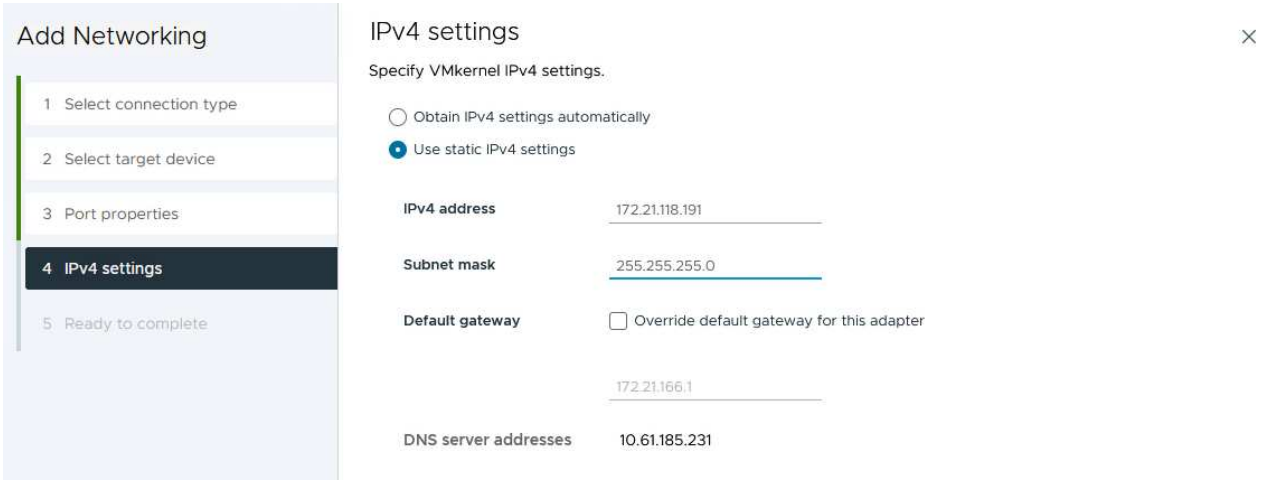

6. Review the your selections on the **Ready to complete** page and click on **Finish** to create the VMkernel adapter.

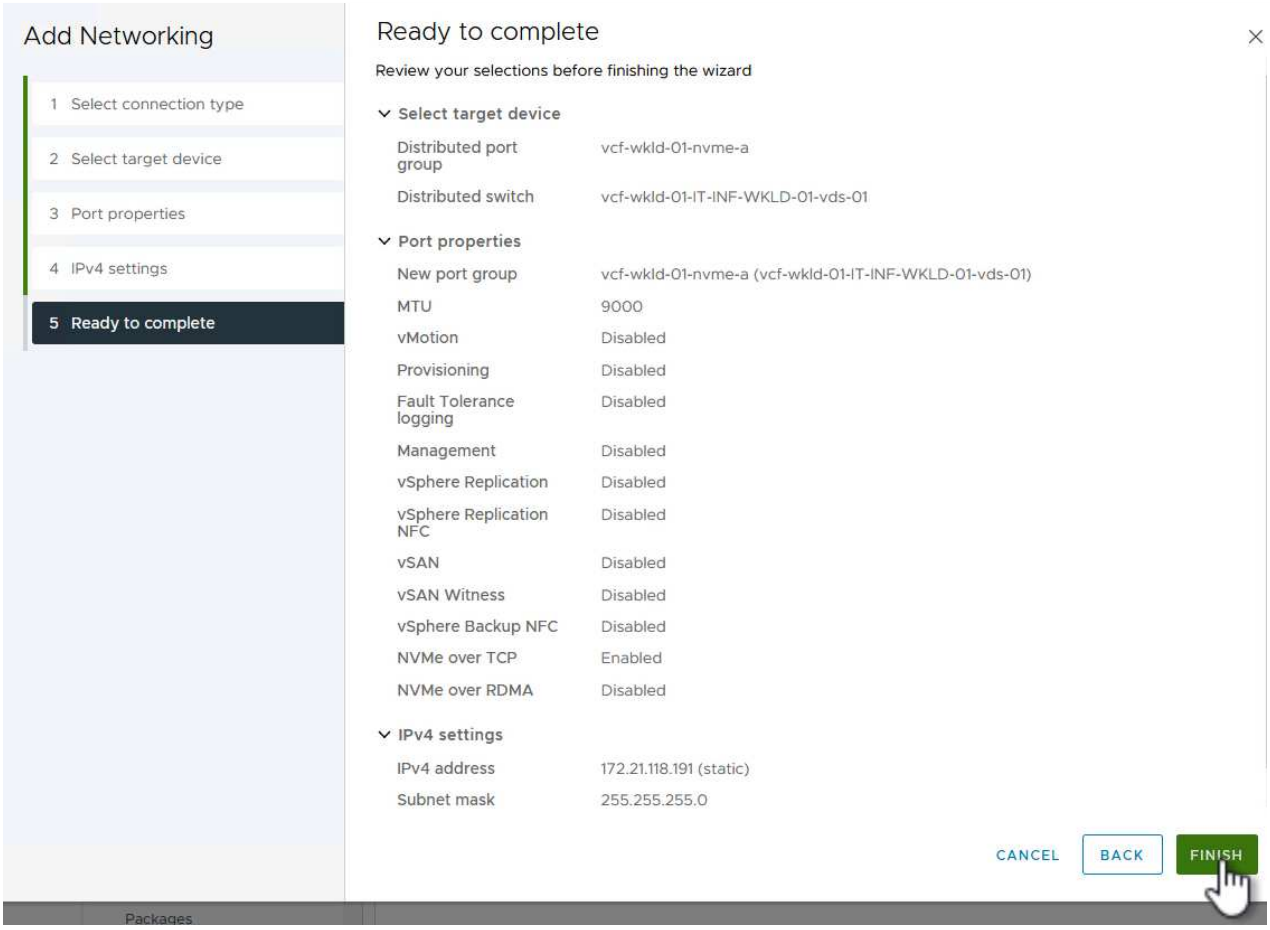

7. Repeat this process to create a VMkernel adapter for the second iSCSI network.

Add NVMe over TCP adapter

Each ESXi host in the workload domain cluster must have an NVMe over TCP software adapter installed for every established NVMe/TCP network dedicated to storage traffic.

To install NVMe over TCP adapters and discover the NVMe controllers, complete the following steps:

1. In the vSphere client navigate to one of the ESXi hosts in the workload domain cluster. From the **Configure** tab click on **Storage Adapters** in the menu and then, from the **Add Software Adapter** drop-down menu, select **Add NVMe over TCP adapter**.

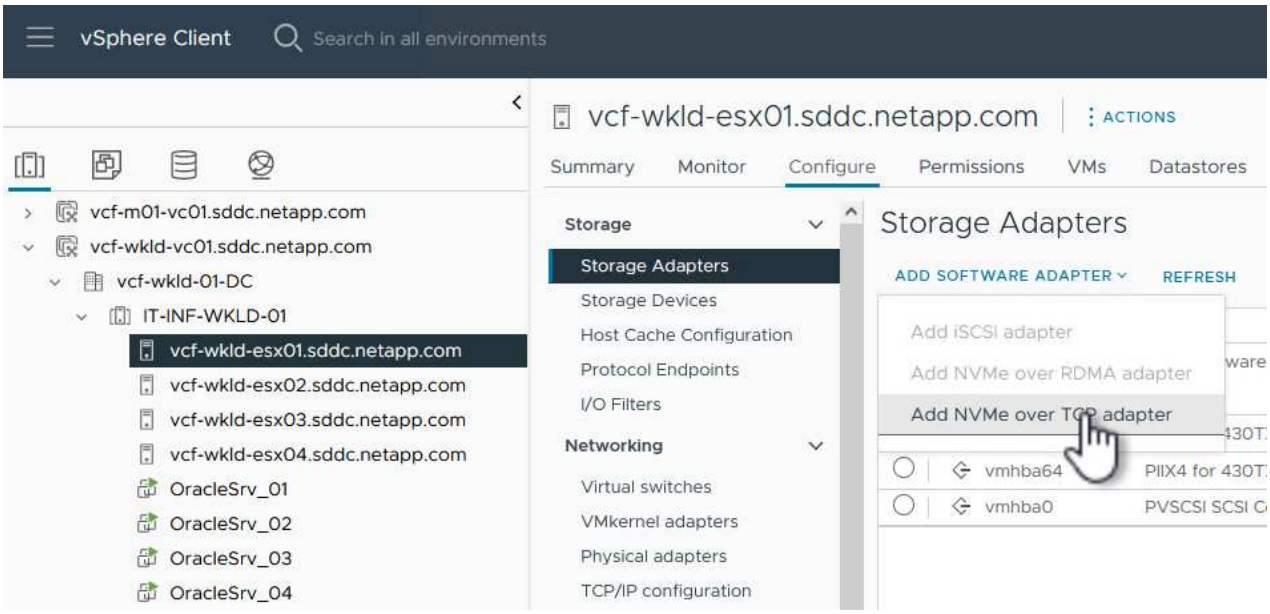

2. In the **Add Software NVMe over TCP adapter** window, access the **Physical Network Adapter** drop-down menu and select the correct physical network adapter on which to enable the NVMe adapter.

3. Repeat this process for the second network assigned to NVMe over TCP traffic, assigning the correct physical adapter.
4. Select one of the newly installed NVMe over TCP adapters and, on the **Controllers** tab, select **Add Controller**.

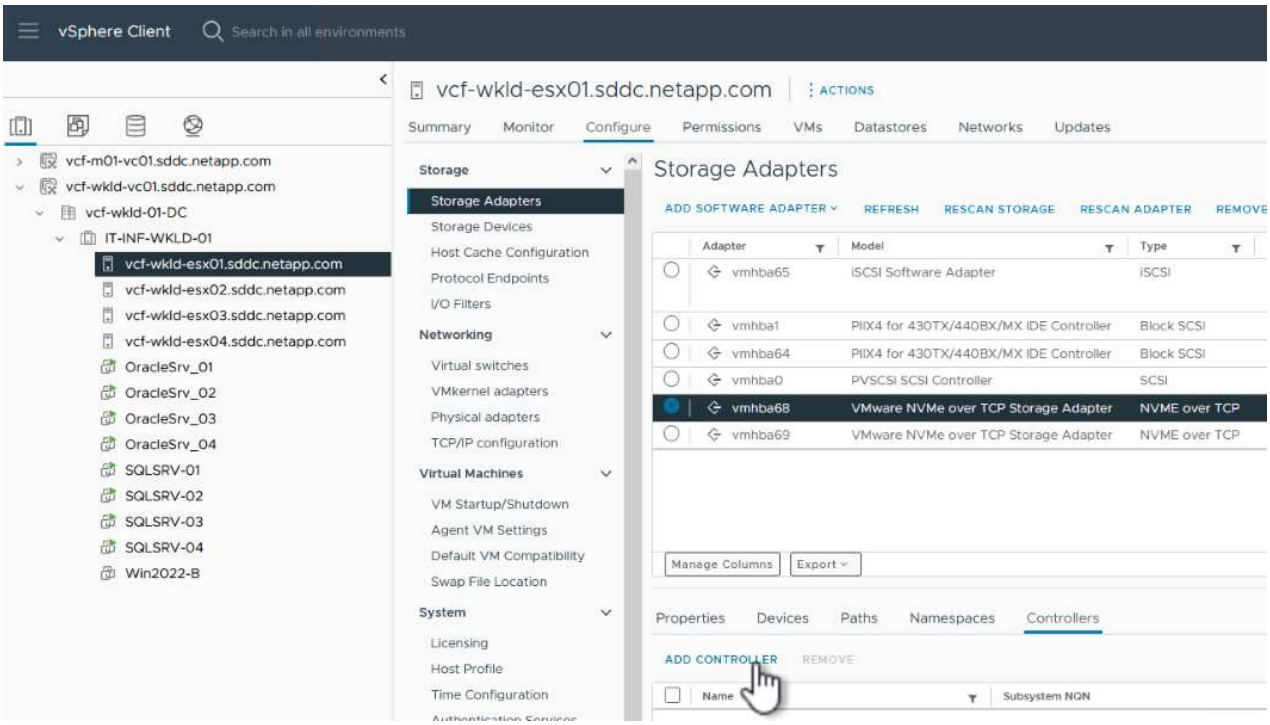

5. In the **Add controller** window, select the **Automatically** tab and complete the following steps.
   - Fill in an IP addresses for one of the SVM logical interfaces on the same network as the physical adapter assigned to this NVMe over TCP adapter.
   - Click on the **Discover Controllers** button.
   - From the list of discovered controllers, click the check box for the two controllers with network addresses aligned with this NVMe over TCP adapter.
   - Click on the **OK** button to add the selected controllers.

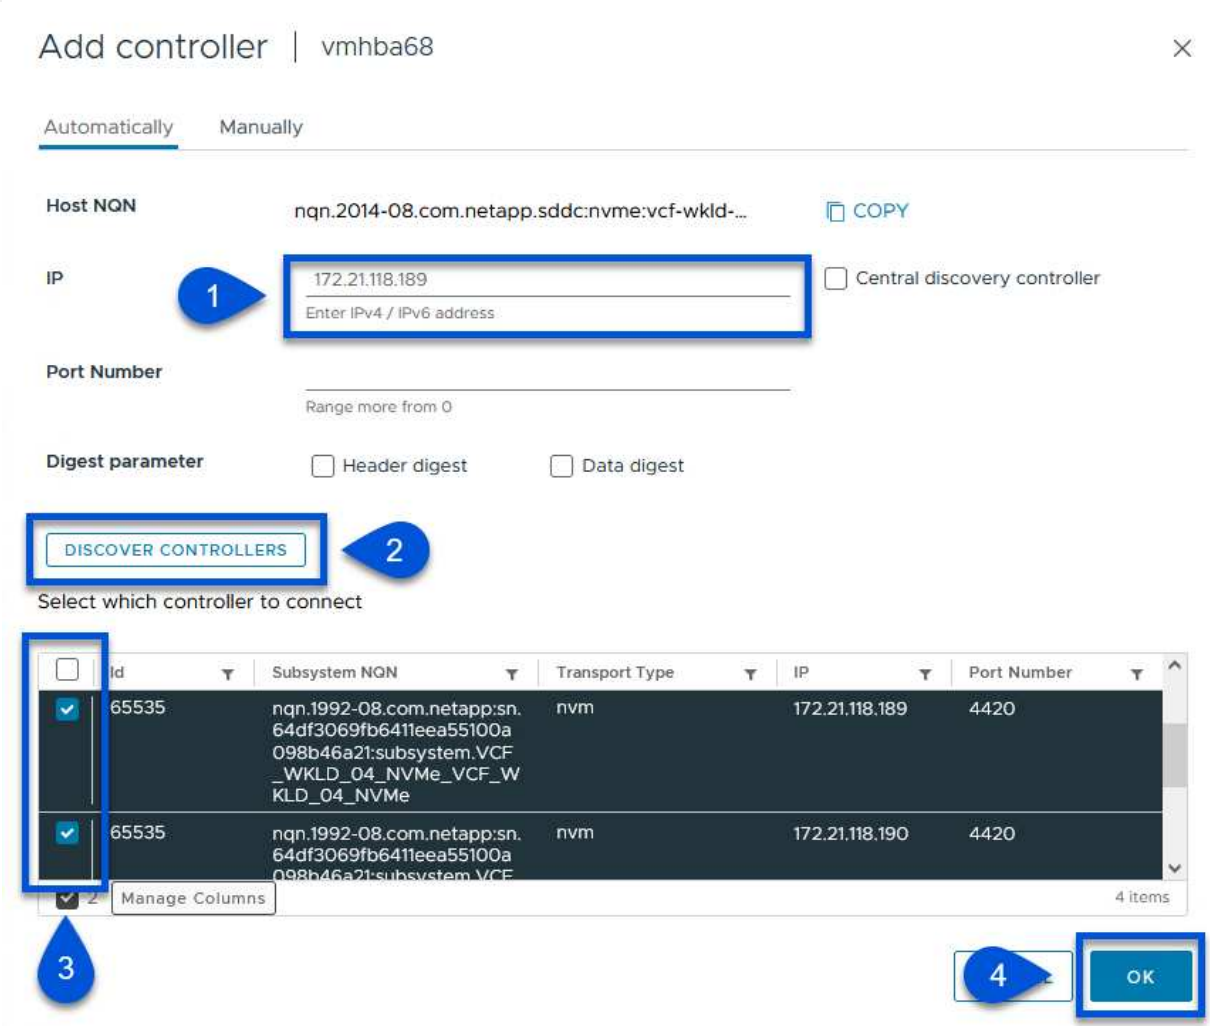

6. After a few seconds you should see the NVMe namespace appear on the Devices tab.

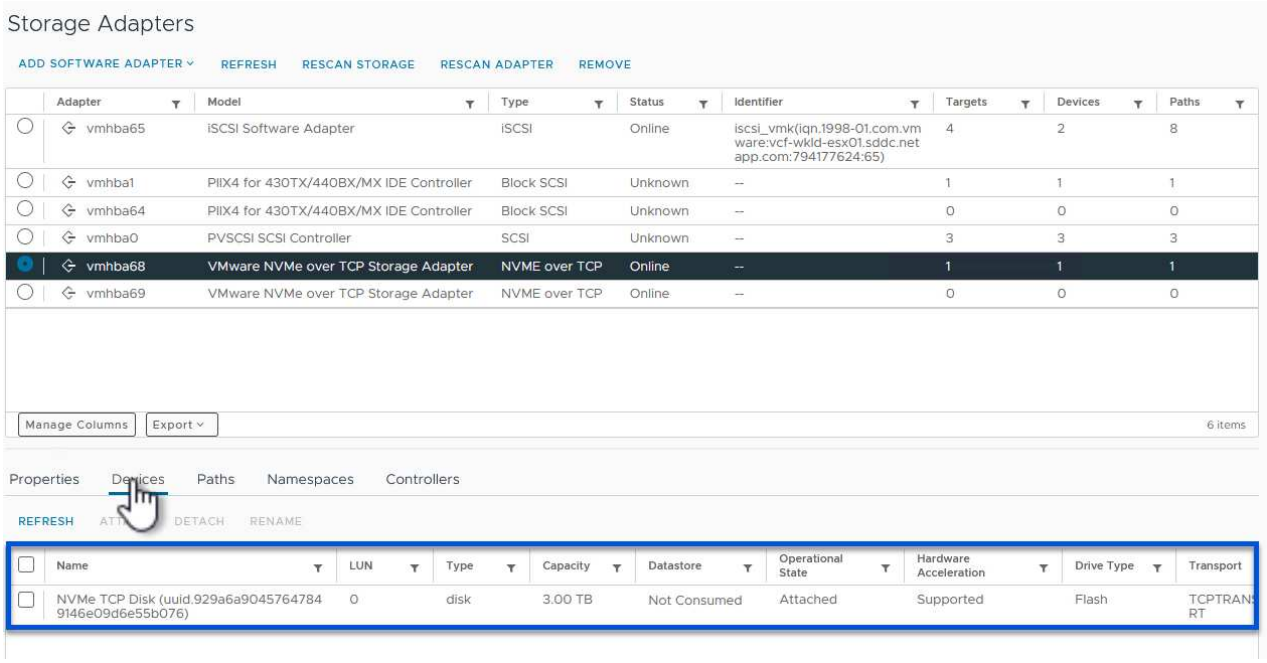

7. Repeat this procedure to create an NVMe over TCP adapter for the second network established for NVMe/TCP traffic.

Deploy NVMe over TCP datastore

To create a VMFS datastore on the NVMe namespace, complete the following steps:

1. In the vSphere client navigate to one of the ESXi hosts in the workload domain cluster. From the **Actions** menu select **Storage > New Datastore…**.

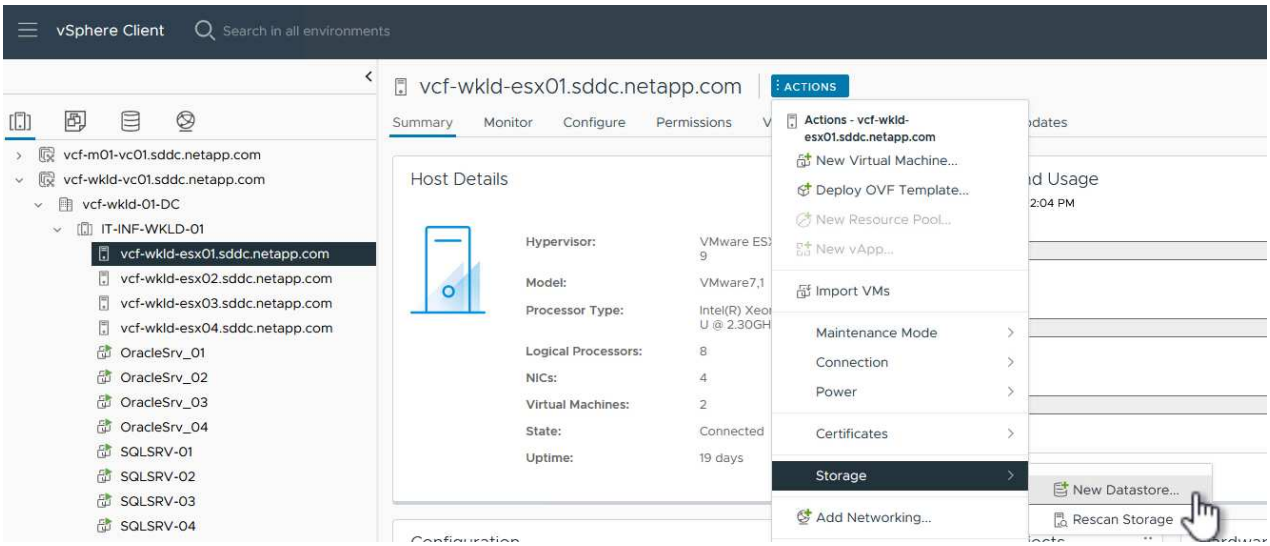

2. In the **New Datastore** wizard, select **VMFS** as the type. Click on **Next** to continue.
3. On the **Name and device selection** page, provide a name for the datastore and select the NVMe namespace from the list of available devices.

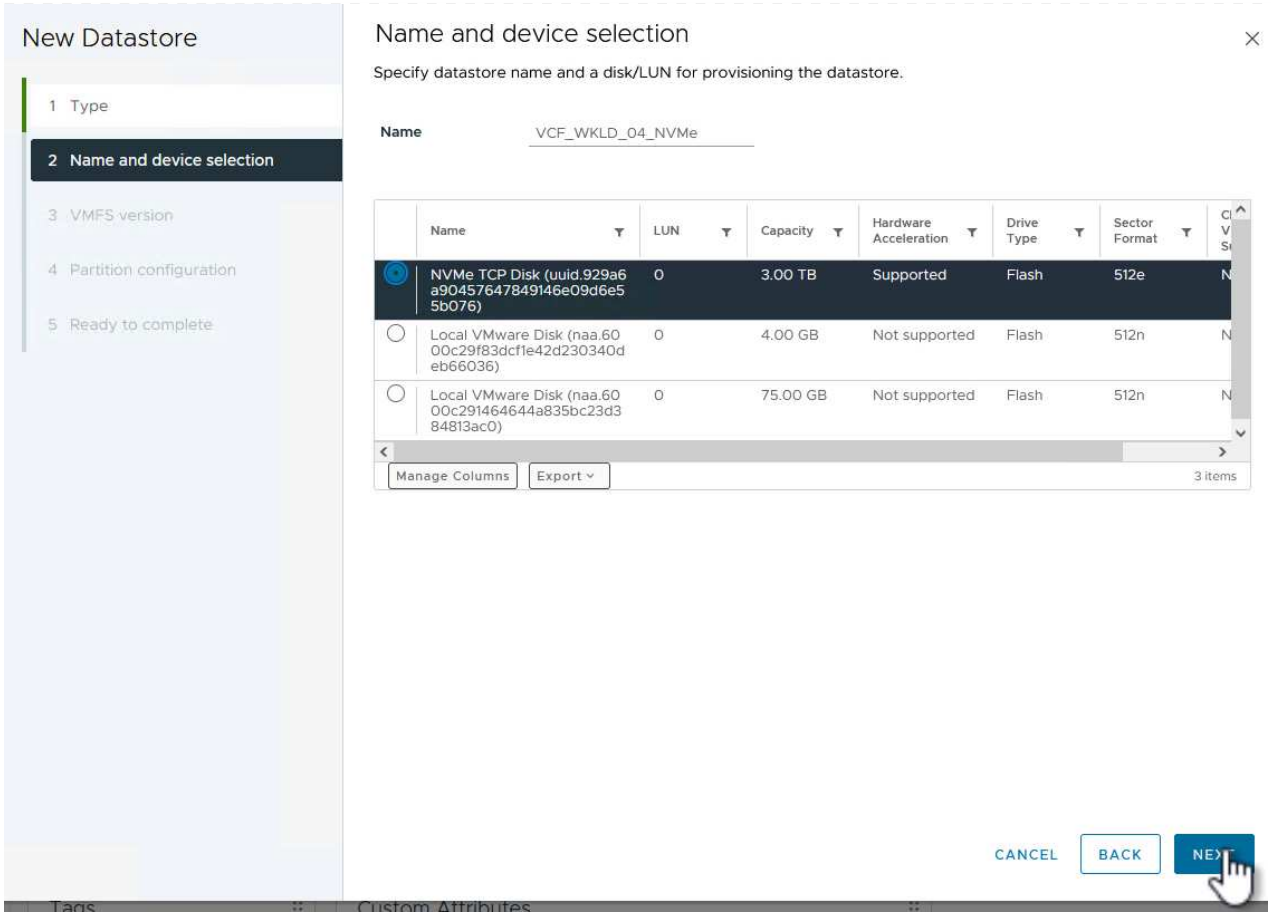

4. On the **VMFS version** page select the version of VMFS for the datastore.
5. On the **Partition configuration** page, make any desired changes to the default partition scheme. Click on **Next** to continue.

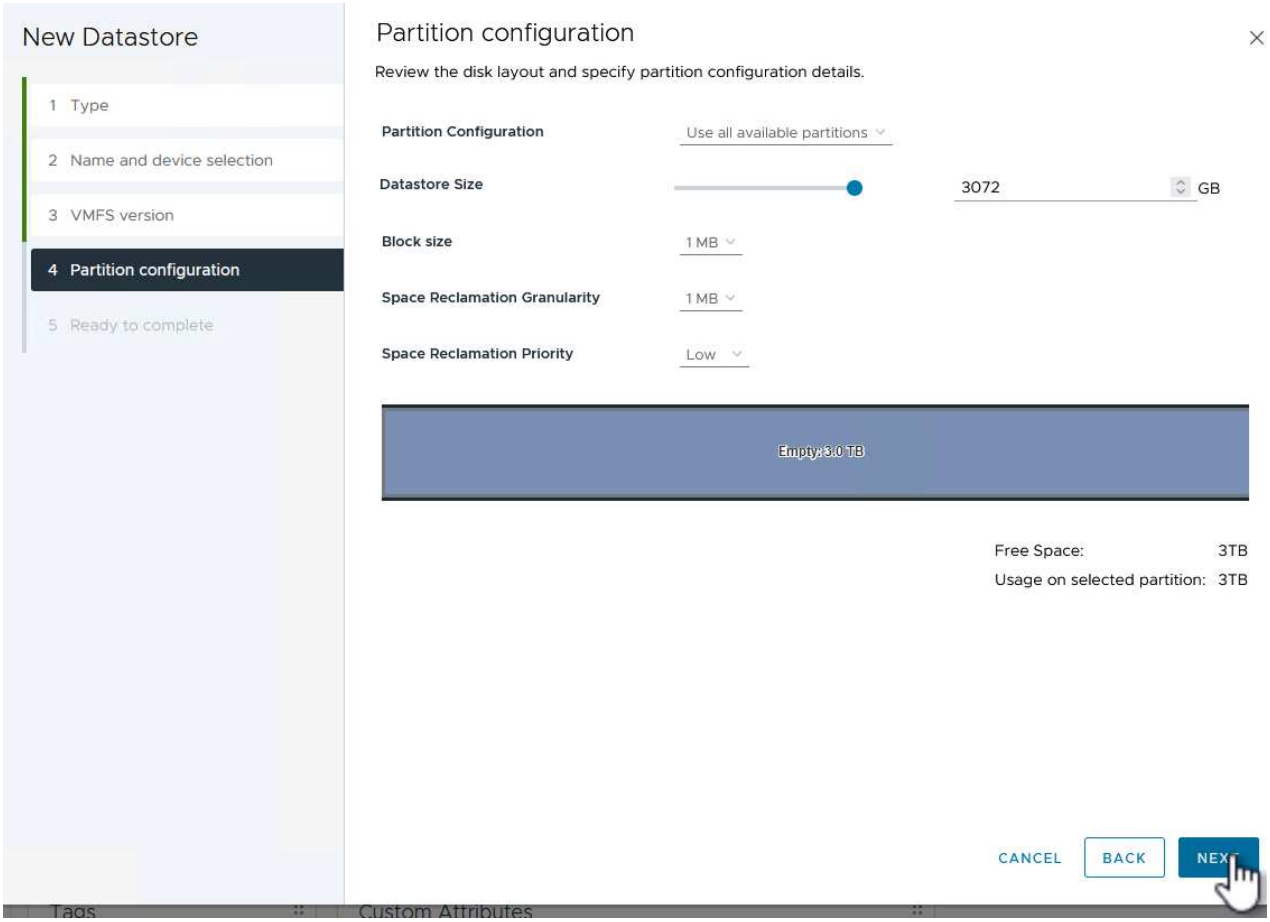

6. On the **Ready to complete** page, review the summary and click on **Finish** to create the datastore.
7. Navigate to the new datastore in inventory and click on the **Hosts** tab. If configured correctly, all ESXi hosts in the cluster should be listed and have access to the new datastore.

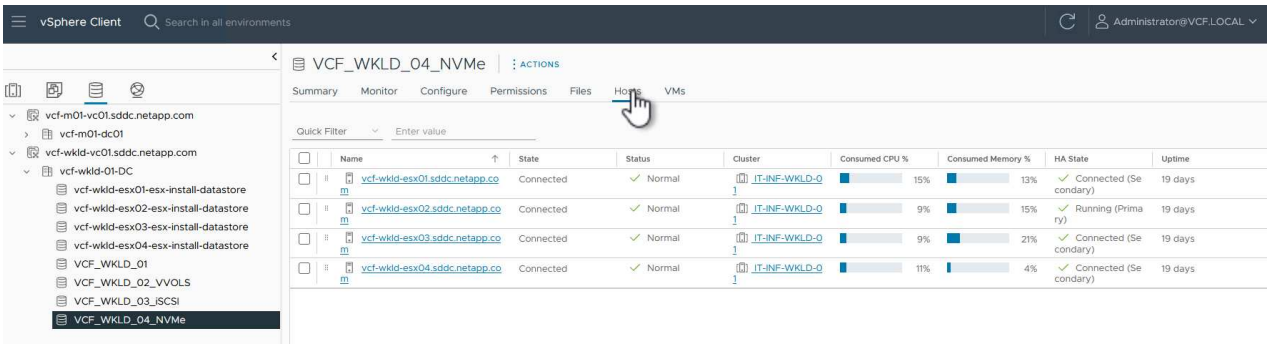

Additional information

For information on configuring ONTAP storage systems refer to the ONTAP 9 Documentation center.

For information on configuring VCF refer to VMware Cloud Foundation Documentation.

In this scenario we will demonstrate how to deploy and use the SnapCenter Plug-in for VMware vSphere (SCV) to backup and restore VM's and datastores on a VCF workload domain. SCV uses ONTAP snapshot technology to take fast and efficient backup copies of the ONTAP storage volumes hosting vSphere datastores. SnapMirror and SnapVault technology are used to create secondary backups on a separate storage system and with retention policies that mimic the original volume or can be independent of the original volume for longer term retention.

iSCSI is used as the storage protocol for the VMFS datastore in this solution.

Author: Josh Powell

Use SnapCenter Plug-in for VMware vSphere to protect VMs on VCF Workload Domains

Scenario Overview

This scenario covers the following high level steps:

- Deploy the SnapCenter Plug-in for VMware vSphere (SCV) on the VI workload domain.
- Add storage systems to SCV.
- Create backup policies in SCV.
- Create Resource Groups in SCV.
- Use SCV to backup datastores or specific VMs.
- Use SCV to restores VMs to an alternate location in the cluster.
- Use SCV to restores files to a windows file system.

Prerequisites

This scenario requires the following components and configurations:

- An ONTAP ASA storage system with iSCSI VMFS datastores allocated to the workload domain cluster.
- A secondary ONTAP storage system configured to received secondary backups using SnapMirror.
- VCF management domain deployment is complete and the vSphere client is accessible.
- A VI workload domain has been previously deployed.
- Virtual machines are present on the cluster SCV is designated to protect.

For information on configuring iSCSI VMFS datastores as supplemental storage refer to **iSCSI as supplemental storage for Management Domains** in this documentation. The process for using OTV to deploy datastores is identical for management and workload domains.

In addition to replicating backups taken with SCV to secondary storage, offsite copies of data can be made to object storage on one of the three (3) leading cloud providers using NetApp BlueXP backup and recovery for VMs. For more information refer to the solution 3-2-1 Data Protection for VMware with SnapCenter Plug-in and BlueXP backup and recovery for VMs.

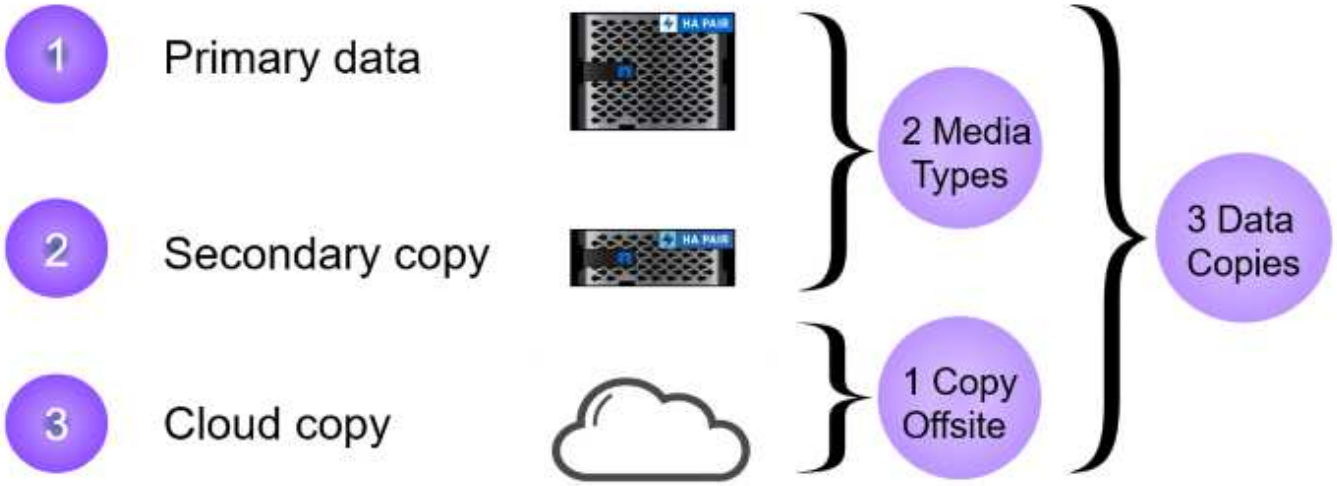

Deployment Steps

To deploy the SnapCenter Plug-in and use it to create backups, and restore VMs and datastores, complete the following steps:

Deploy and use SCV to protect data in a VI workload domain

Complete the following steps to deploy, configure, and use SCV to protect data in a VI workload domain:

Deploy the SnapCenter Plug-in for VMware vSphere

The SnapCenter Plug-in is hosted on the VCF management domain but registered to the vCenter for the VI workload domain. One SCV instance is required for each vCenter instance and, keep in mind that, a Workload domain can include multiple clusters managed by a single vCenter instance.

Complete the following steps from the vCenter client to deploy SCV to the VI workload domain:

1. Download the OVA file for the SCV deployment from the download area of the NetApp support site **HERE**.
2. From the management domain vCenter Client, select to **Deploy OVF Template…**.

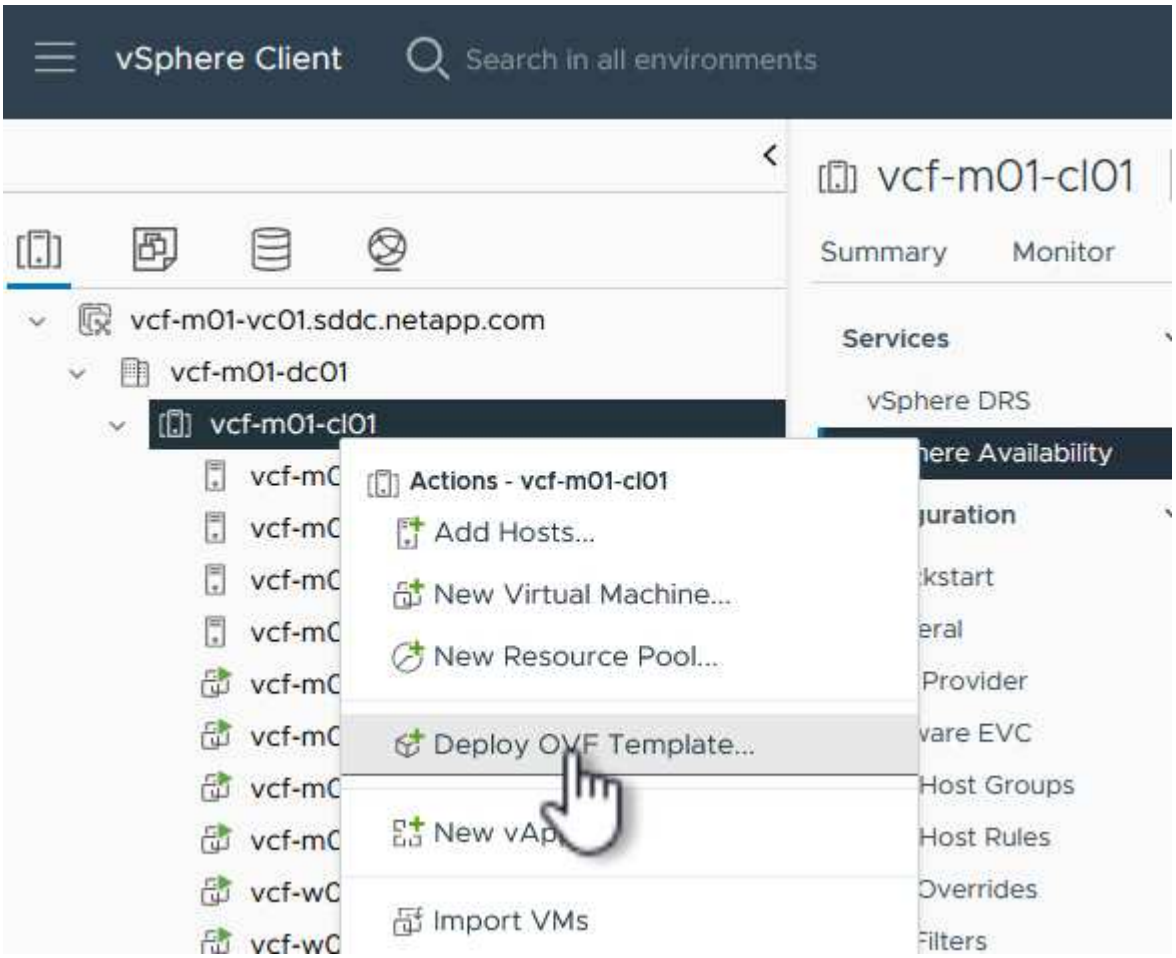

3. In the **Deploy OVF Template** wizard, click on the **Local file** radio button and then select to upload the previously downloaded OVF template. Click on **Next** to continue.

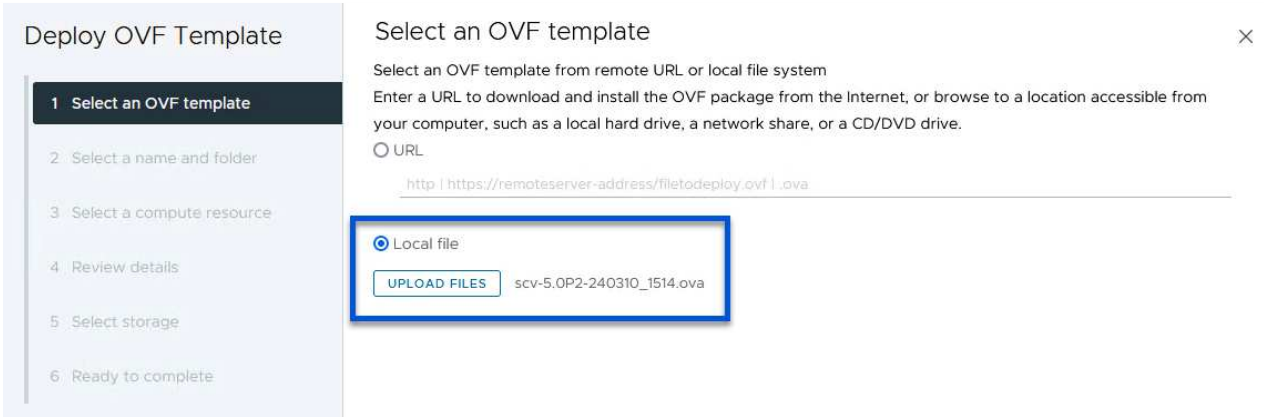

4. On the **Select name and folder** page, provide a name for the SCV data broker VM and a folder on the management domain. Click on **Next** to continue.
5. On the **Select a compute resource** page, select the management domain cluster or specific ESXi host within the cluster to install the VM to.
6. Review information pertaining to the OVF template on the **Review details** page and agree to the licensing terms on the **Licensing agreements** page.
7. On the **Select storage** page choose the datastore which the VM will be installed to and select the **virtual disk format** and **VM Storage Policy**. In this solution, the VM will be installed on an iSCSI VMFS datastore located on an ONTAP storage system, as previously deployed in a separate section of this documentation. Click on **Next** to continue.

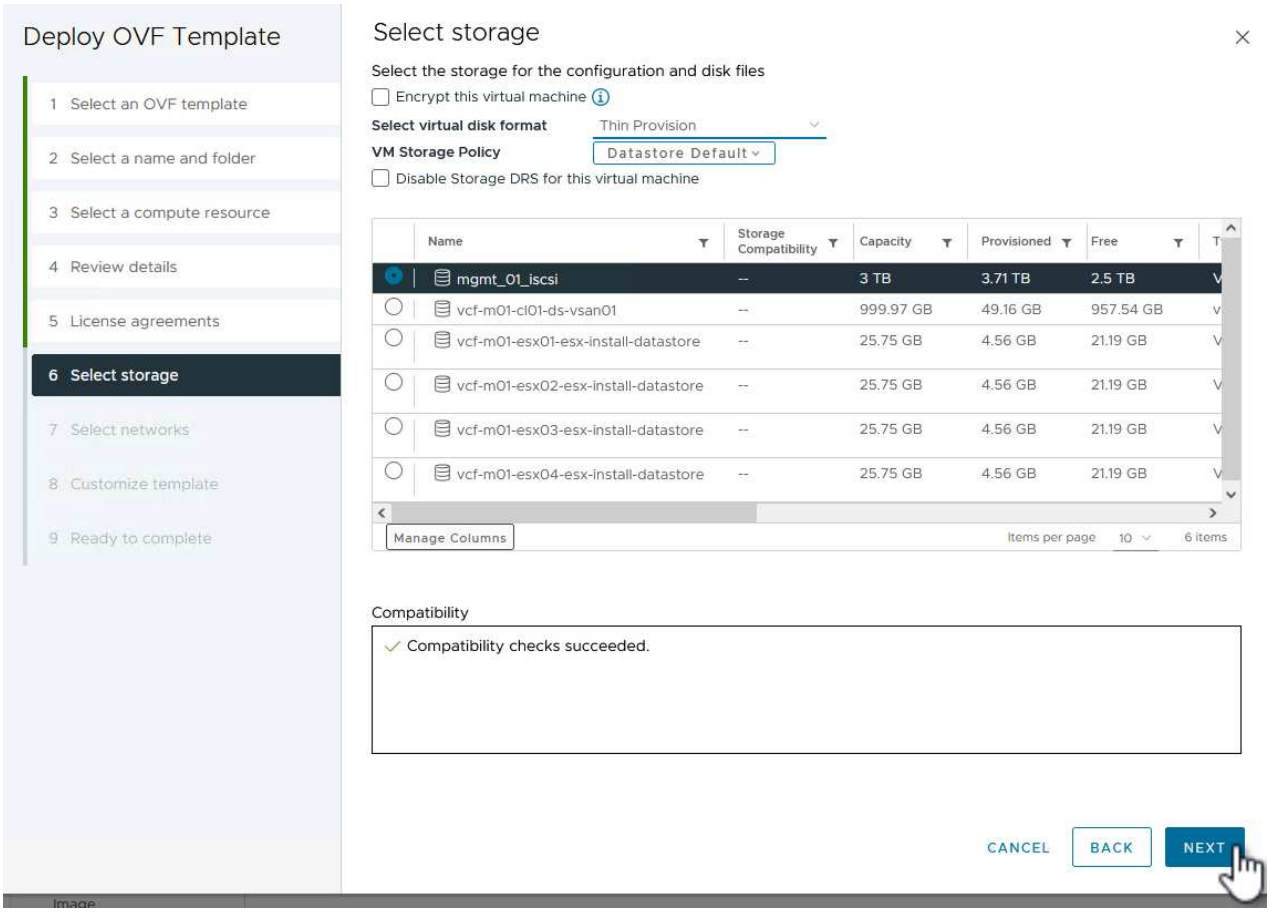

8. On the **Select network** page, select the management network that is able to communicate with the workload domain vCenter appliance and both the primary and secondary ONTAP storage systems.

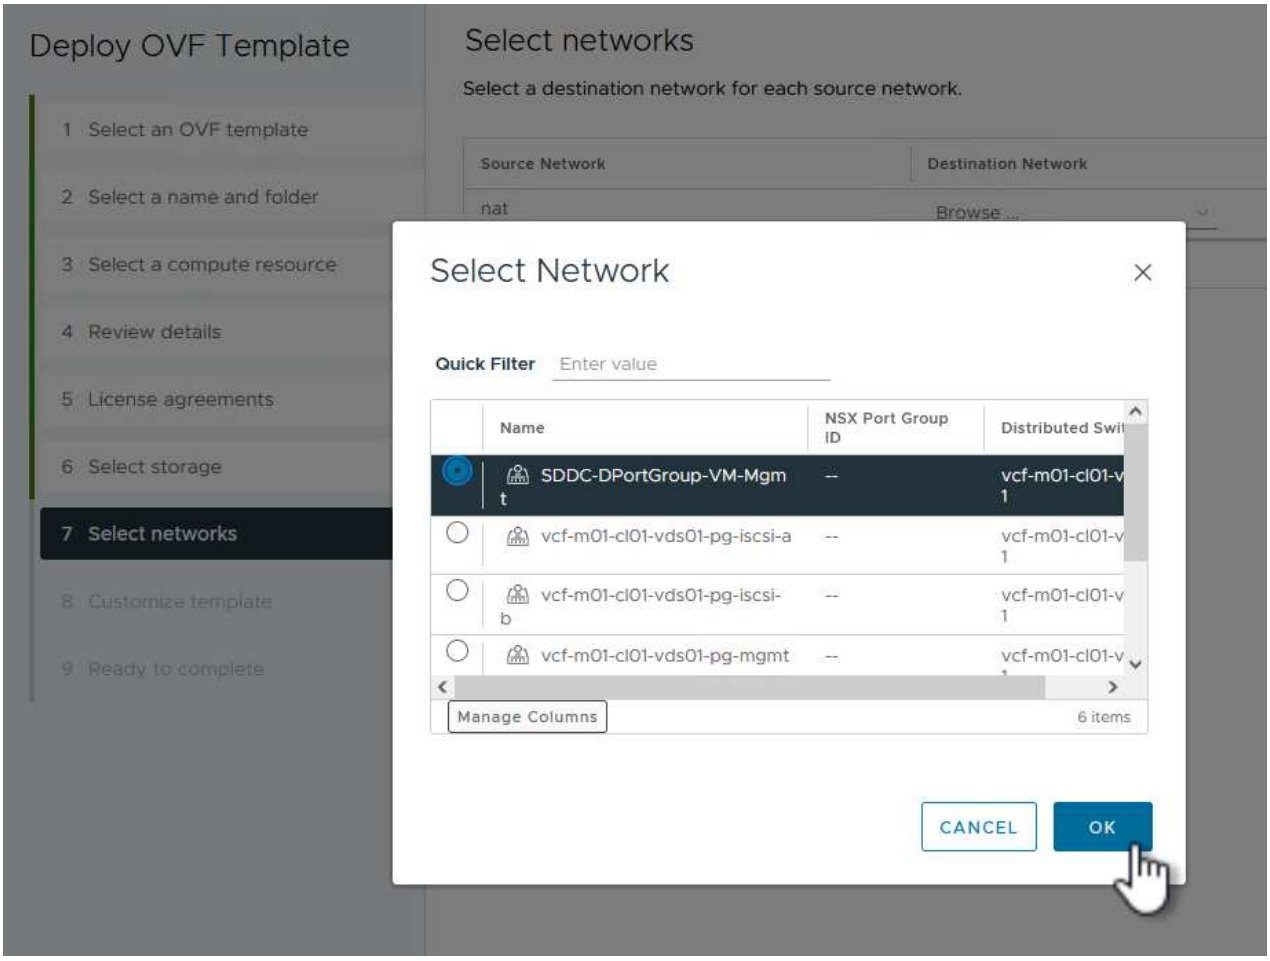

9. On the **Customize template** page fill out all information required for the deployment:
   - FQDN or IP, and credentials for the workload domain vCenter appliance.
   - Credentials for the SCV administrative account.
   - Credentials for the SCV maintenance account.
   - IPv4 Network Properties details (IPv6 can also be used).
   - Date and Time settings.

   Click on **Next** to continue.

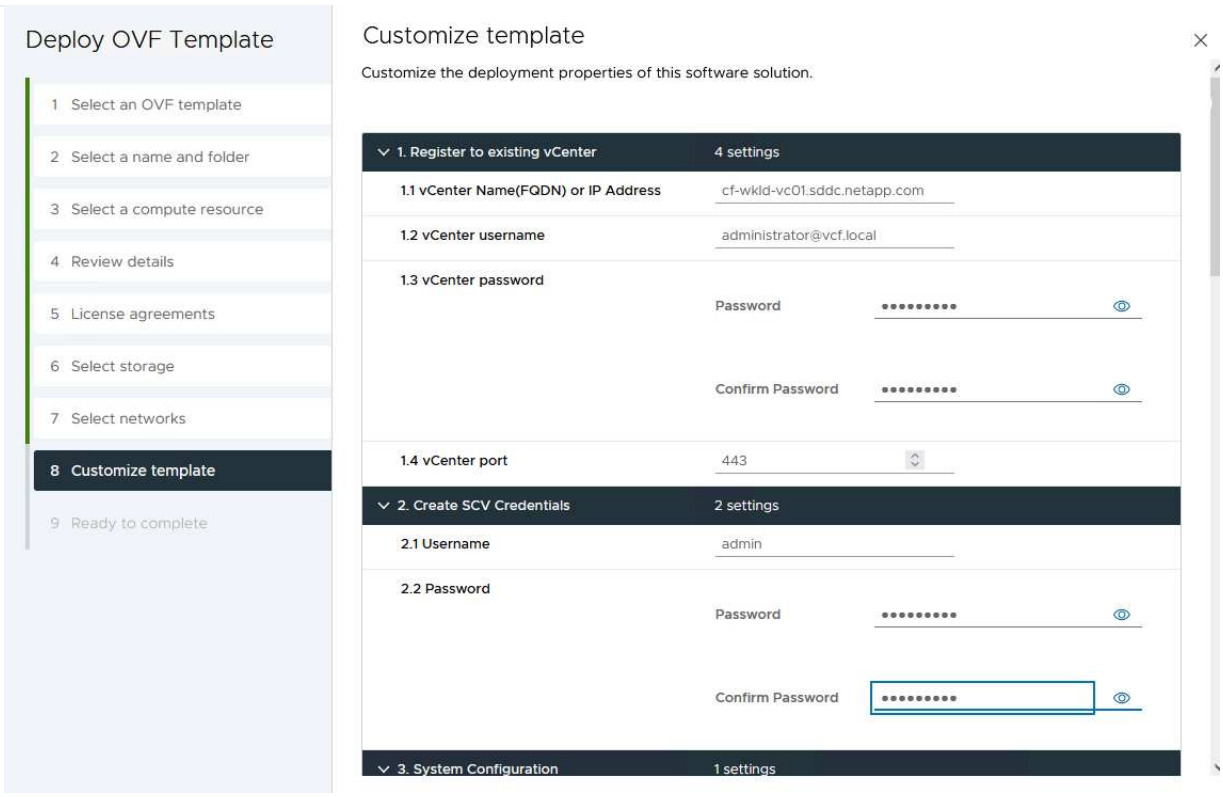
Deploy OVF Template
1 Select an OVF template
2 Select a name and folder
3 Select a compute resource
4 Review details
5 License agreements
6 Select storage
7 Select networks
8 Customize template
9 Ready to complete
Customize template
Customize the deployment properties of this software solution.
1. Register to existing vCenter
4 settings
1.1 vCenter Name(FQDN) or IP Address
cf-wkld-vc01.sddc.netapp.com
1.2 vCenter username
administrator@vcf.local
1.3 vCenter password
Password
Confirm Password
1.4 vCenter port
443
2. Create SCV Credentials
2 settings
2.1 Username
admin
2.2 Password
Password
Confirm Password
3. System Configuration
1 settings

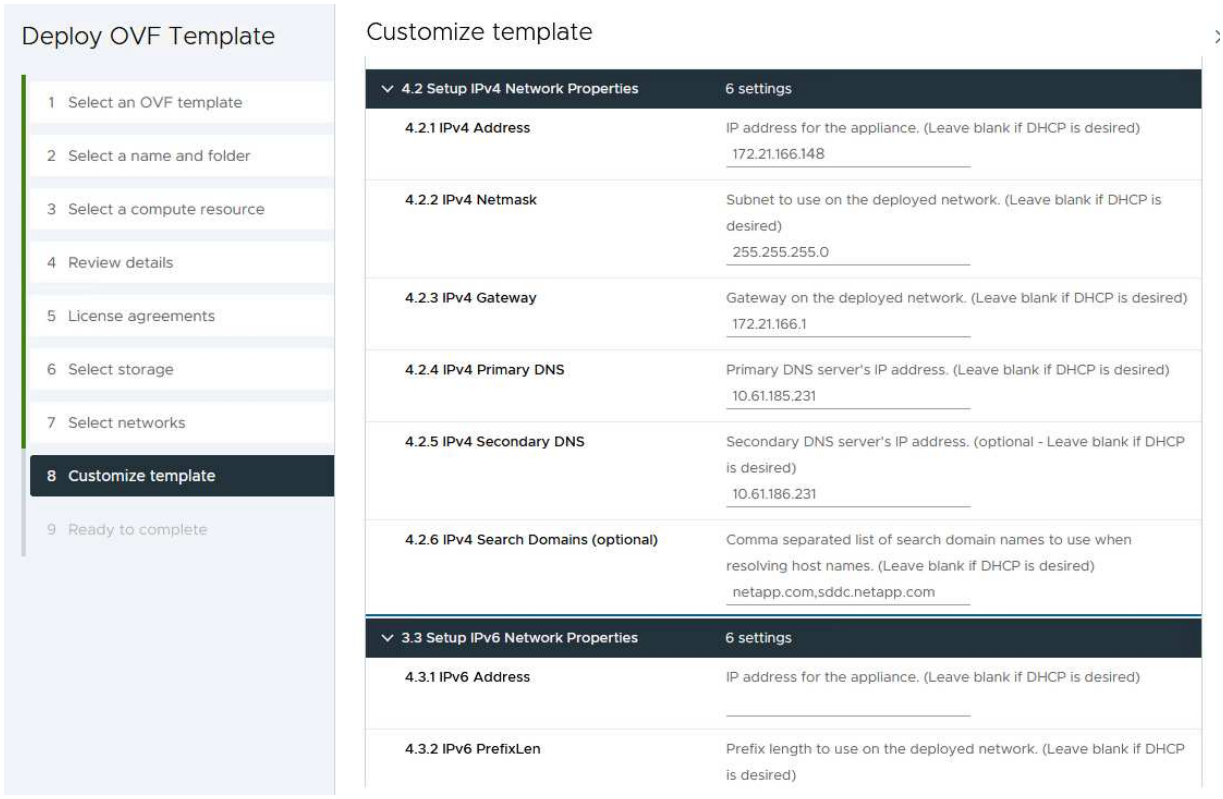
Deploy OVF Template
1 Select an OVF template
2 Select a name and folder
3 Select a compute resource
4 Review details
5 License agreements
6 Select storage
7 Select networks
8 Customize template
9 Ready to complete
Customize template
4.2 Setup IPv4 Network Properties
6 settings
4.2.1 IPv4 Address
IP address for the appliance. (Leave blank if DHCP is desired)
172.21.166.148
4.2.2 IPv4 Netmask
Subnet to use on the deployed network. (Leave blank if DHCP is desired)
255.255.255.0
4.2.3 IPv4 Gateway
Gateway on the deployed network. (Leave blank if DHCP is desired)
172.21.166.1
4.2.4 IPv4 Primary DNS
Primary DNS server's IP address. (Leave blank if DHCP is desired)
10.61.185.231
4.2.5 IPv4 Secondary DNS
Secondary DNS server's IP address. (optional - Leave blank if DHCP is desired)
10.61.186.231
4.2.6 IPv4 Search Domains (optional)
Comma separated list of search domain names to use when resolving host names. (Leave blank if DHCP is desired)
netapp.com,sddc.netapp.com
3.3 Setup IPv6 Network Properties
6 settings
4.3.1 IPv6 Address
IP address for the appliance. (Leave blank if DHCP is desired)
4.3.2 IPv6 PrefixLen
Prefix length to use on the deployed network. (Leave blank if DHCP is desired)

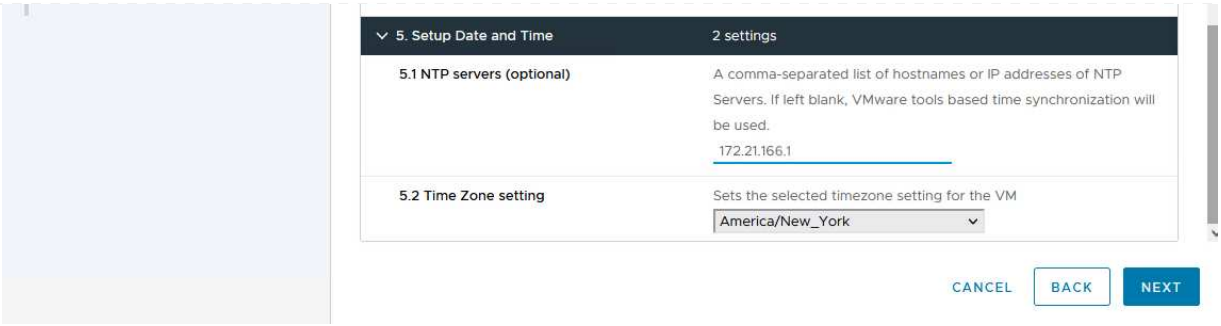

10. Finally, on the **Ready to complete page**, review all settings and click on Finish to start the deployment.

Add Storage Systems to SCV

Once the SnapCenter Plug-in is installed complete the following steps to add storage systems to SCV:

1. SCV can be accessed from the main menu in the vSphere Client.

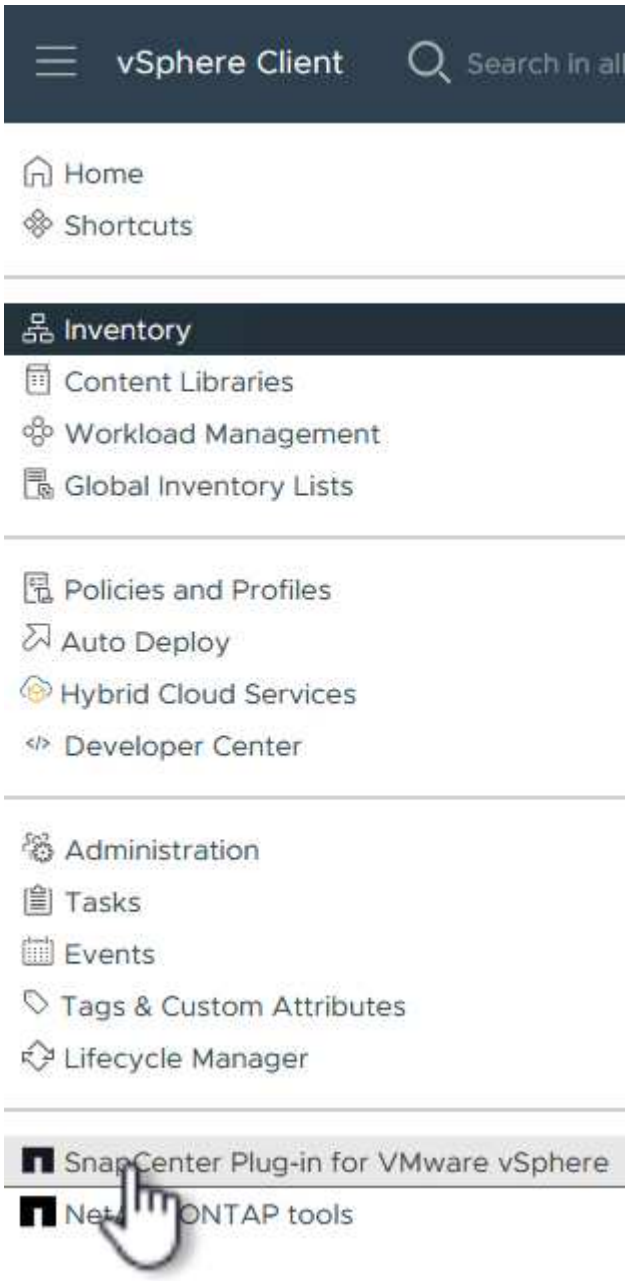

2. At the top of the SCV UI interface, select the correct SCV instance that matches the vSphere cluster to be protected.

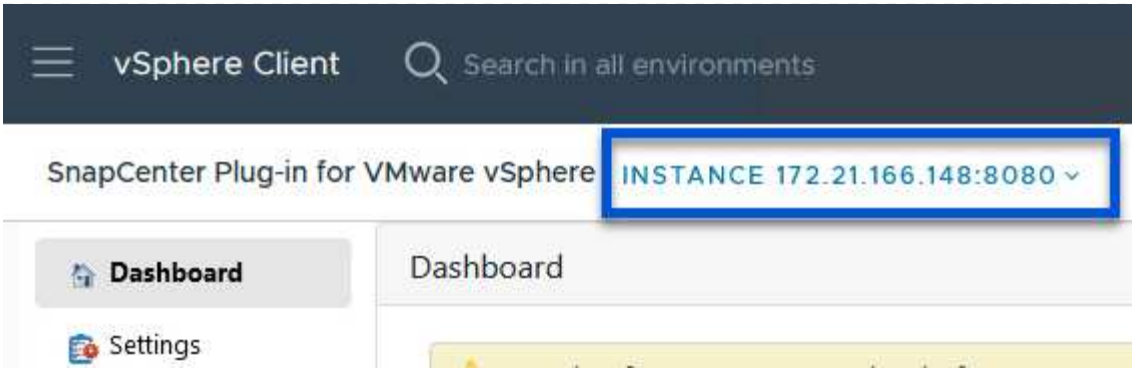

3. Navigate to **Storage Systems** in the left-hand menu and click on **Add** to get started.

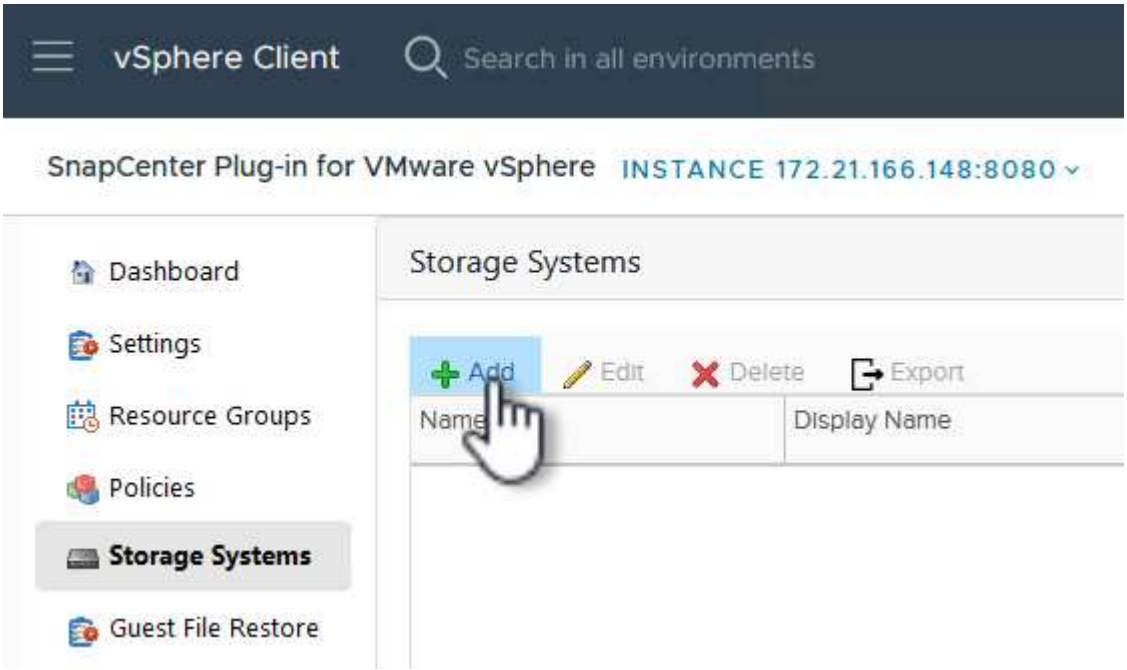

4. On the **Add Storage System** form, fill in the IP address and credentials of the ONTAP storage system to be added, and click on **Add** to complete the action.

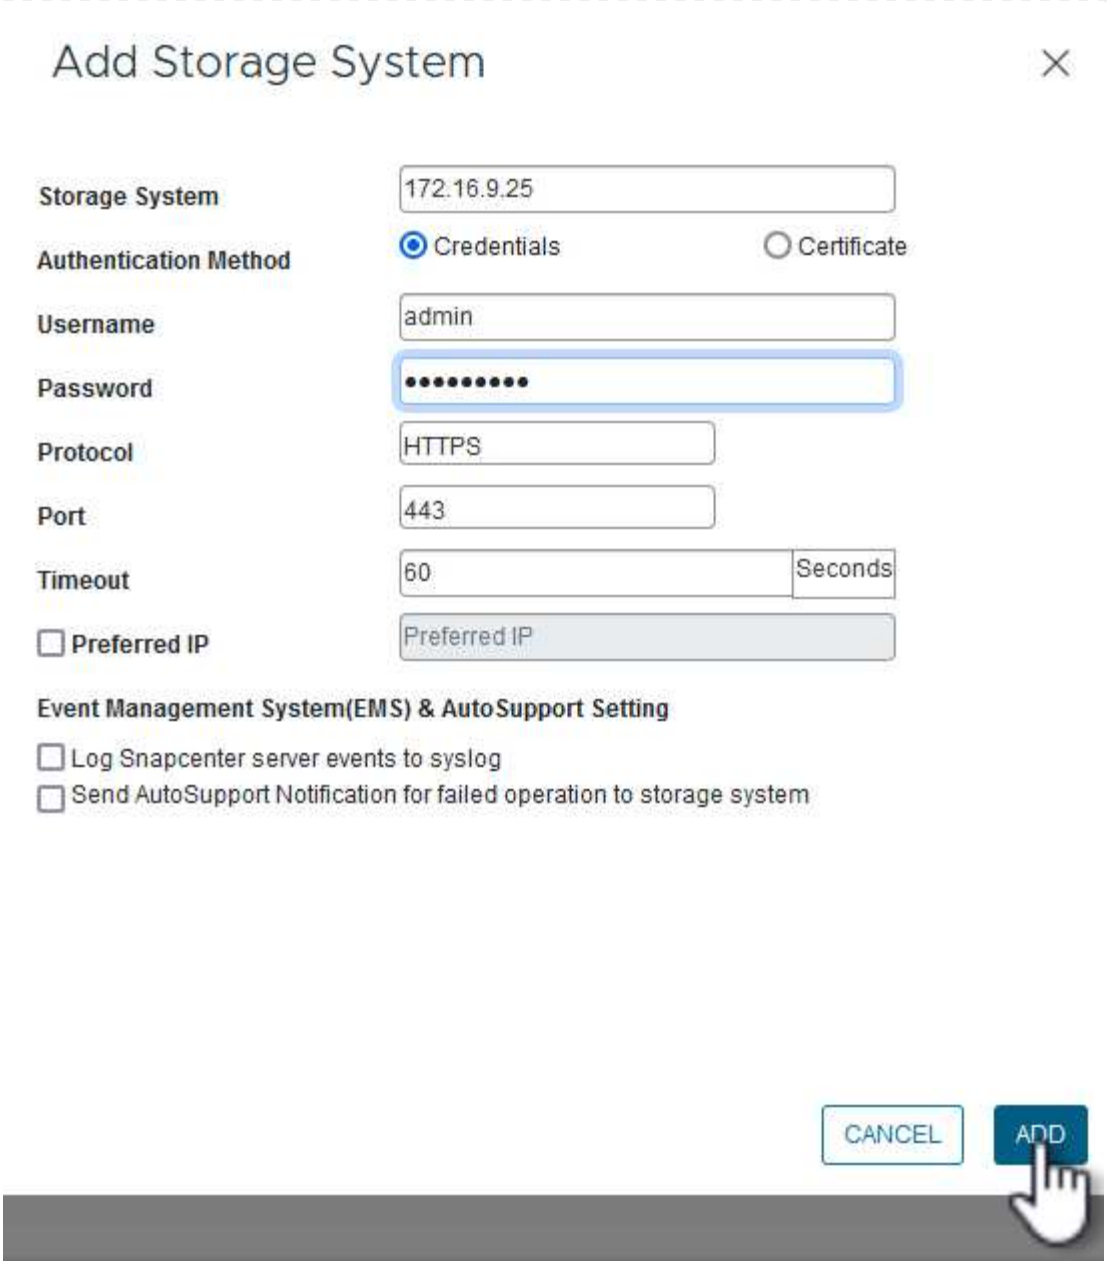

5. Repeat this procedure for any additional storage systems to be managed, including any systems to be used as secondary backup targets.

Configure backup policies in SCV

For more information on creating SCV backup policies refer to Create backup policies for VMs and datastores.

Complete the following steps to create a new backup policy:

1. From the left-hand menu select **Policies** and click on **Create** to begin.

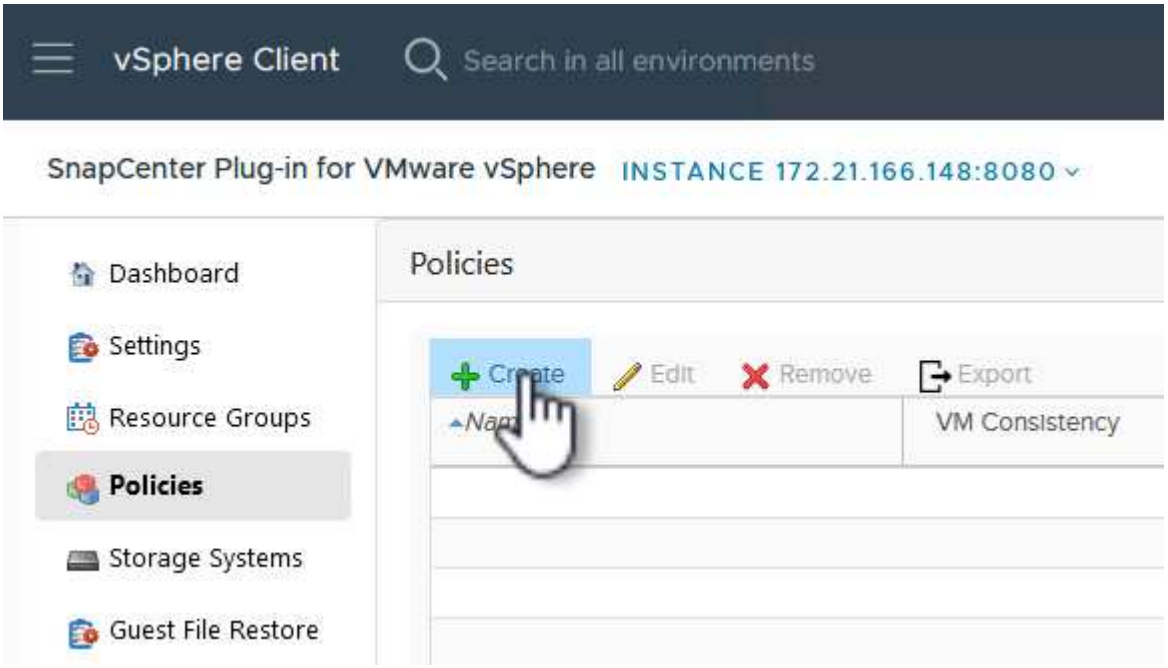

2. On the **New Backup Policy** form, provide a **Name** and **Description** for the policy, the **Frequency** at which the backups will take place, and the **Retention** period which specifies how long the backup is retained.

   Locking Period enables the ONTAP SnapLock feature to create tamper proof snapshots and allows configuration of the locking period.

   For **Replication** Select to update the underlying SnapMirror or SnapVault relationships for the ONTAP storage volume.

   > SnapMirror and SnapVault replication are similar in that they both utilize ONTAP SnapMirror technology to asynchronously replicate storage volumes to a secondary storage system for increased protection and security. For SnapMirror relationships, the retention schedule specified in the SCV backup policy will govern retention for both the primary and secondary volume. With SnapVault relationships, a separate retention schedule can be established on the secondary storage system for longer term or differing retention schedules. In this case the snapshot label is specified in the SCV backup policy and in the policy associated with the secondary volume, to identify which volumes to apply the independent retention schedule to.

   Choose any additional advanced options and click on **Add** to create the policy.

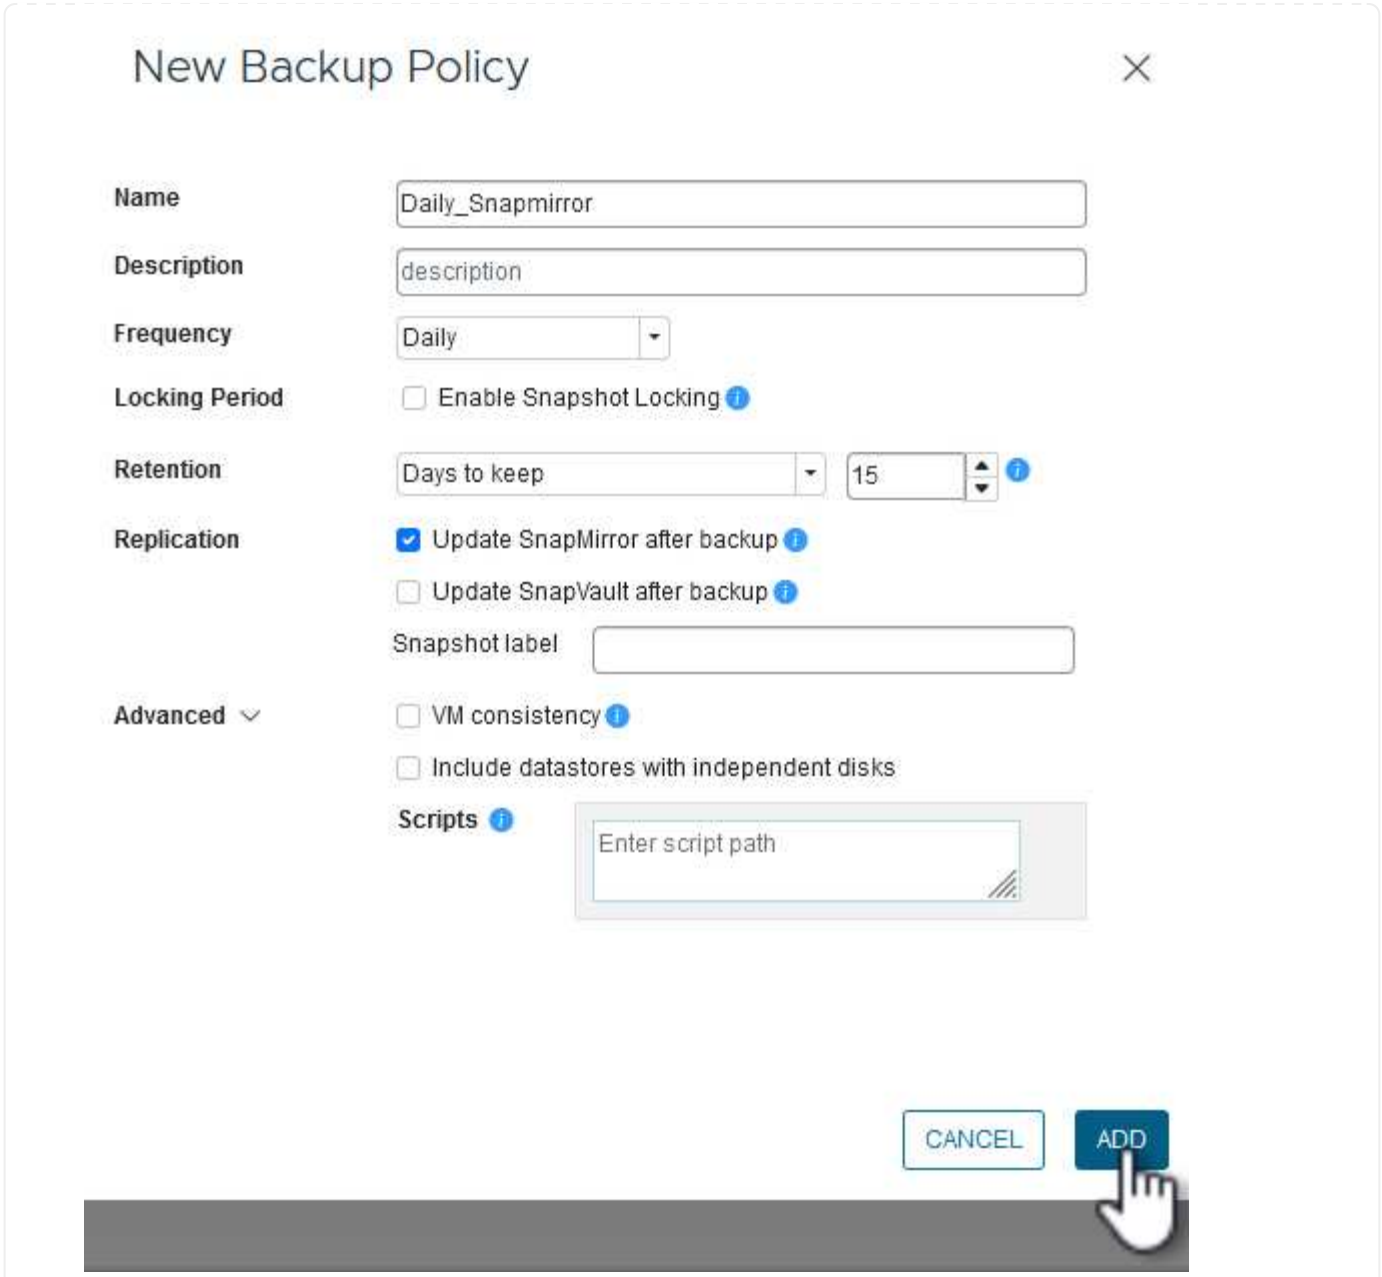
New Backup Policy
Name
Daily_Snapmirror
Description
description
Frequency
Daily
Locking Period
Enable Snapshot Locking
Retention
Days to keep
15
Replication
Update SnapMirror after backup
Update SnapVault after backup
Snapshot label
Advanced
VM consistency
Include datastores with independent disks
Scripts
Enter script path
CANCEL
ADD

Create resource groups in SCV

For more information on creating SCV Resource Groups refer to Create resource groups.

Complete the following steps to create a new resource group:

1. From the left-hand menu select **Resource Groups** and click on **Create** to begin.

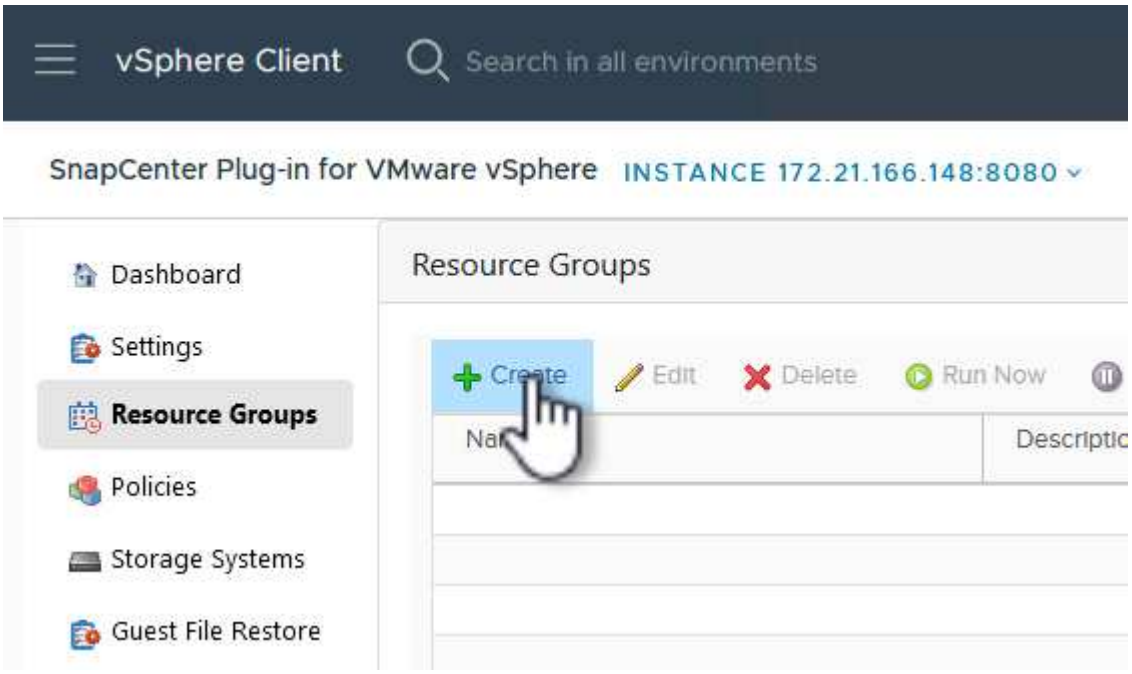

2. On the **General info & notification** page, provide a name for for the resource group, notification settings, and any additional options for the naming of the snapshots.
3. On the **Resource** page select the datastores and VM's to be protected in the resource group. Click on **Next** to continue.

Even when only specific VMs are selected, the entire datastore is always backed up. This is because ONTAP takes snapshots of the volume hosting the datastore. However, note that selecting only specific VMs for backup limits the ability to restore to only those VMs.

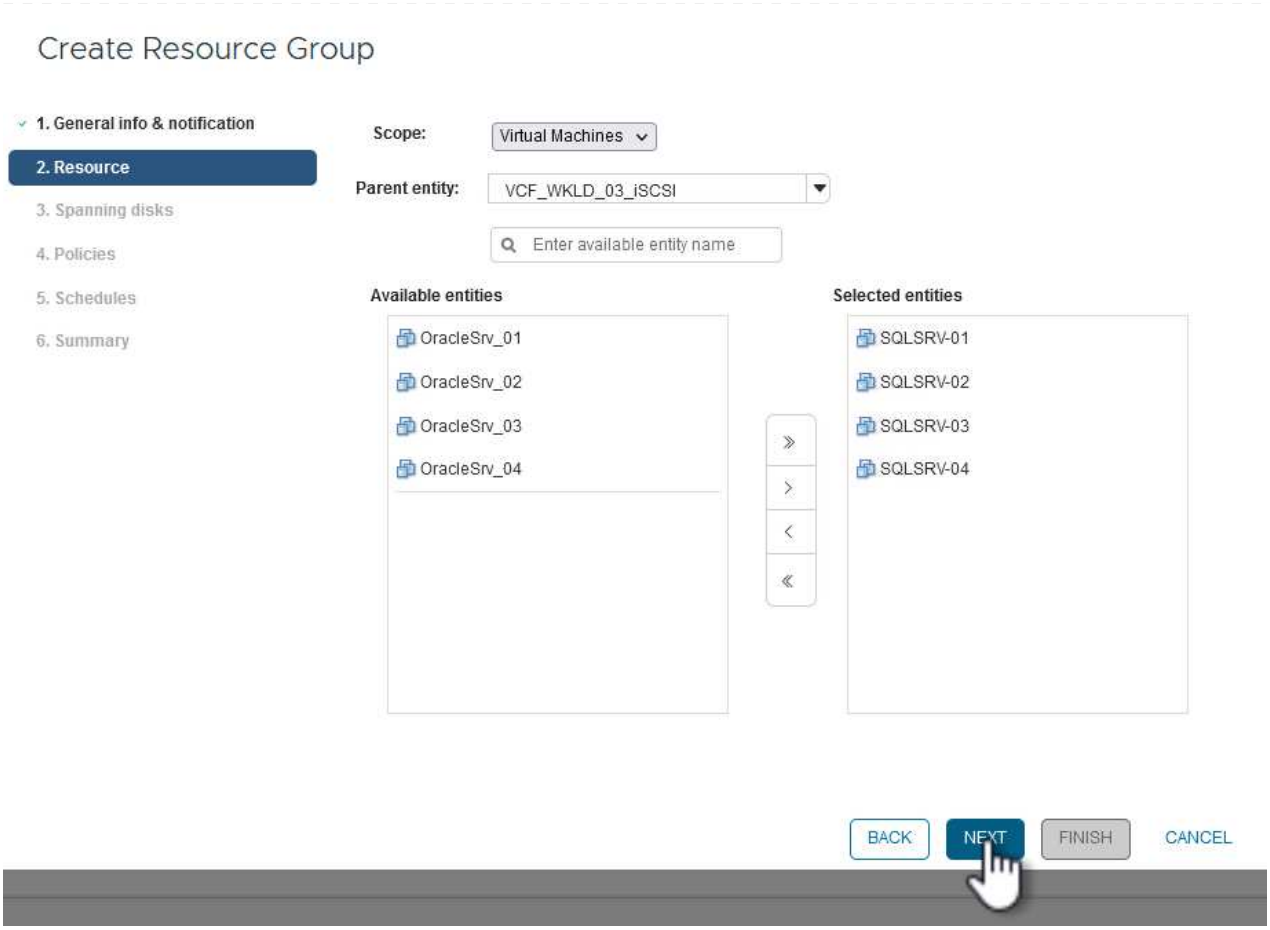

4. On the **Spanning disks** page select the option for how to handle VMs with VMDK's that span multiple datastores. Click on **Next** to continue.

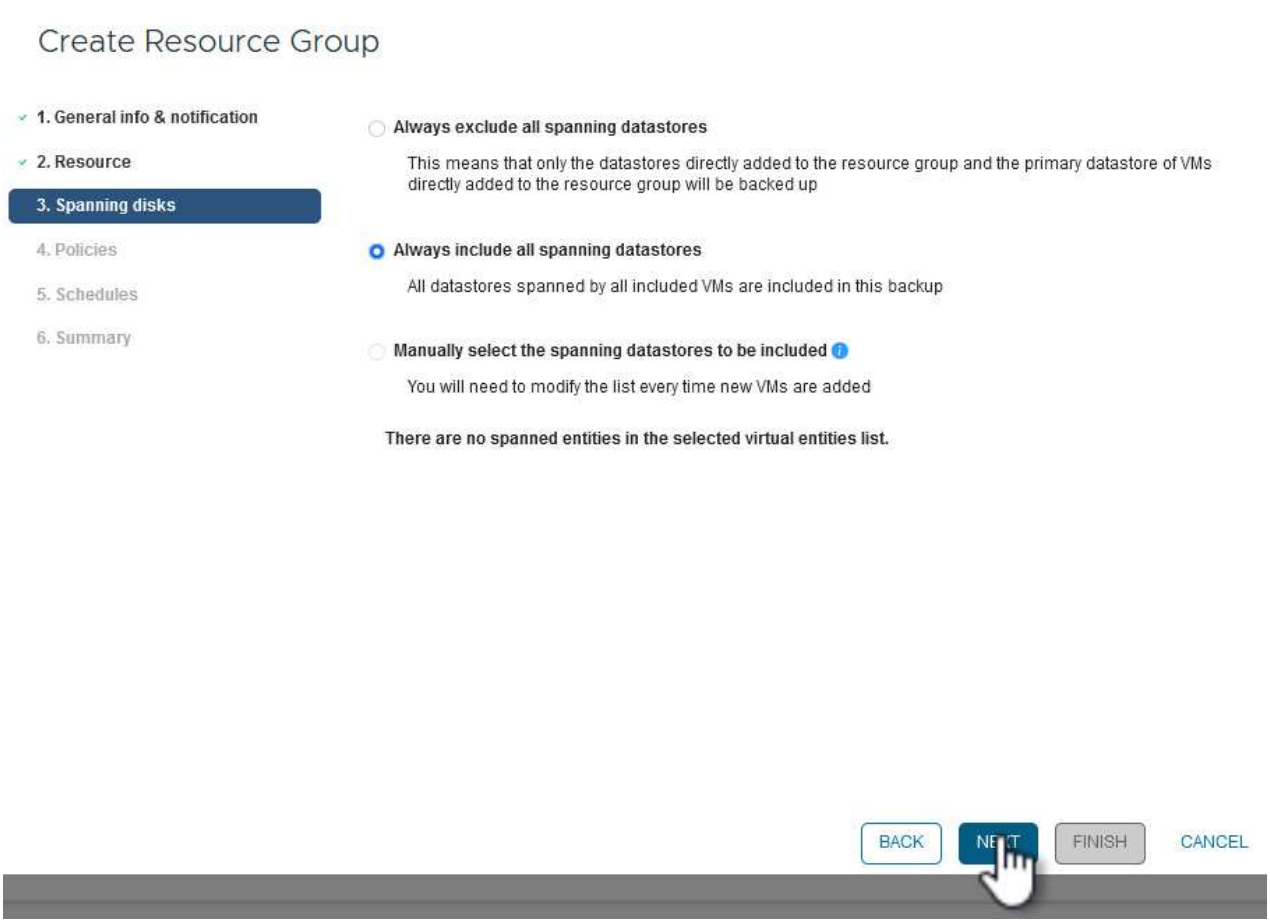

5. On the **Policies** page select a previously created policy or multiple policies that will be used with this resource group. Click on **Next** to continue.

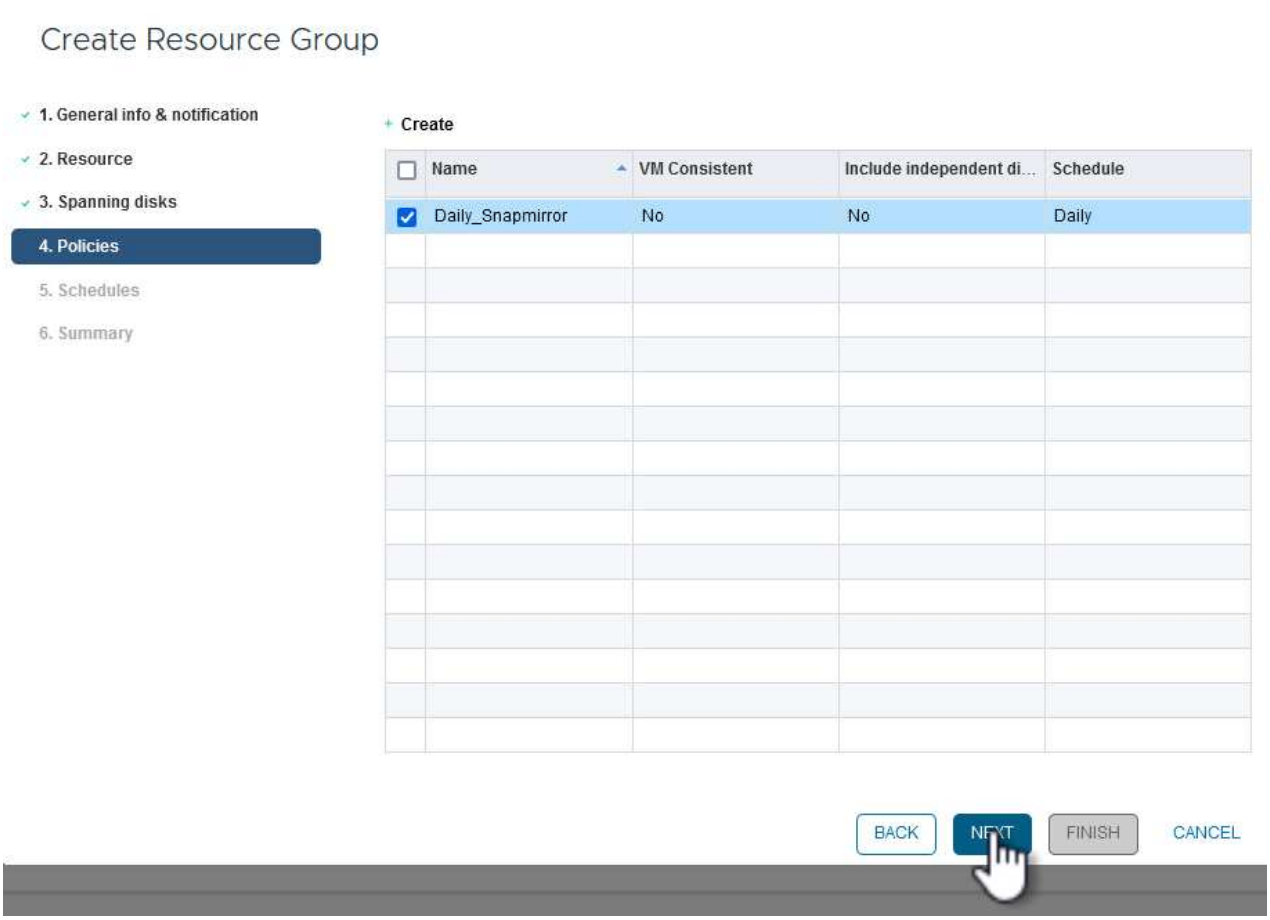

6. On the **Schedules** page establish for when the backup will run by configuring the recurrence and time of day. Click on **Next** to continue.

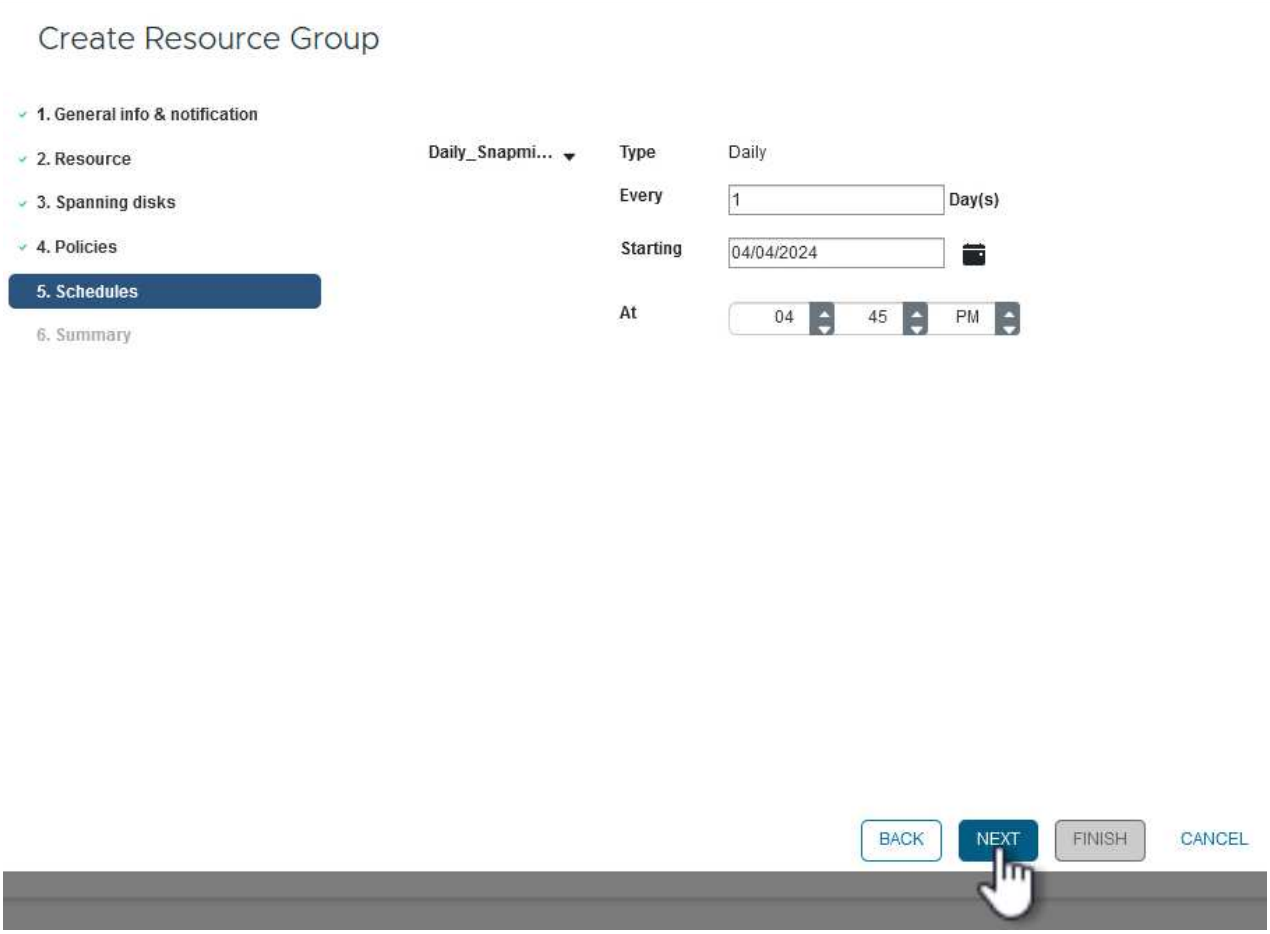

7. Finally review the **Summary** and click on **Finish** to create the resource group.

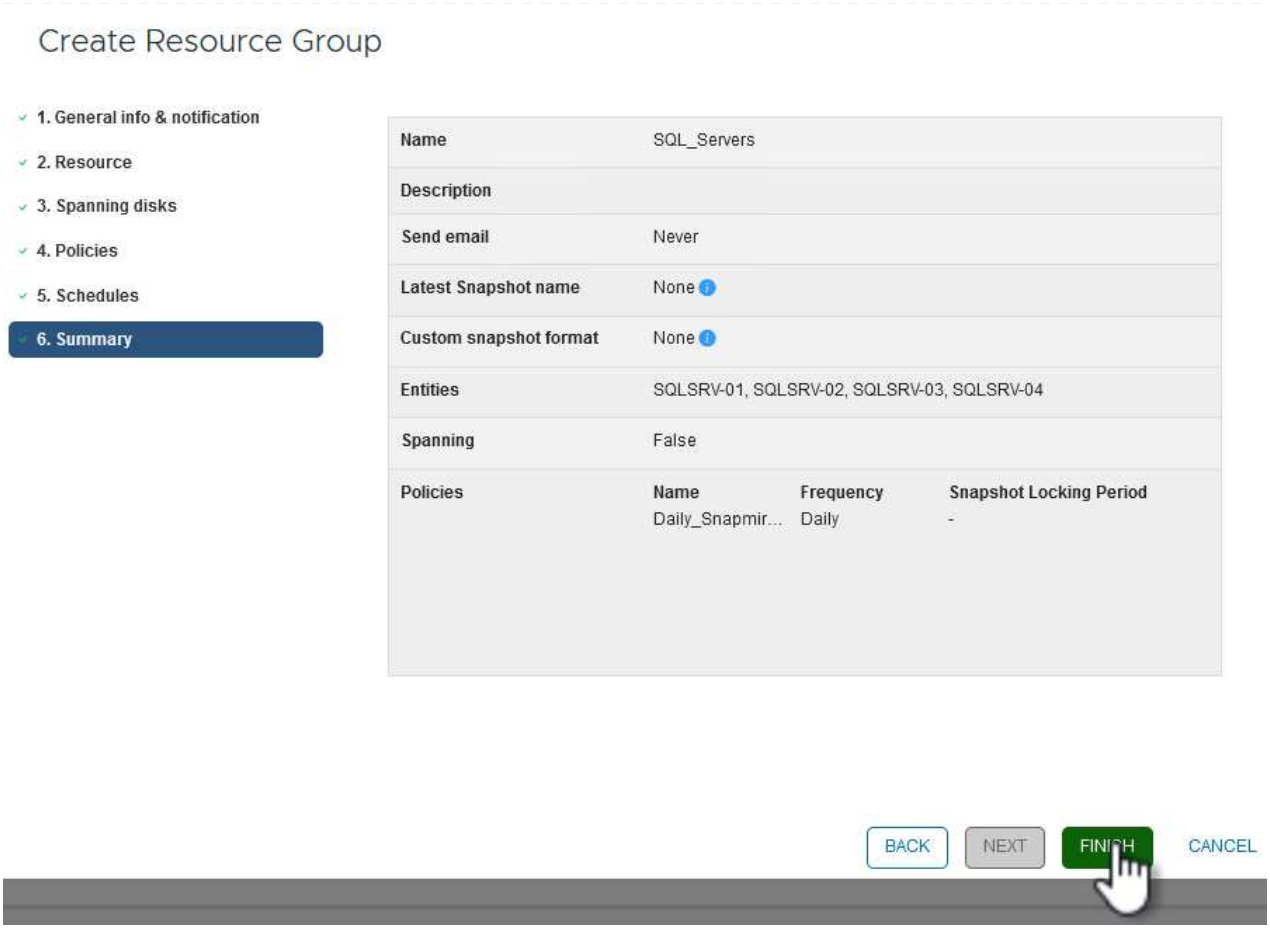

8. With the resource group created click on the **Run Now** button to run the first backup.

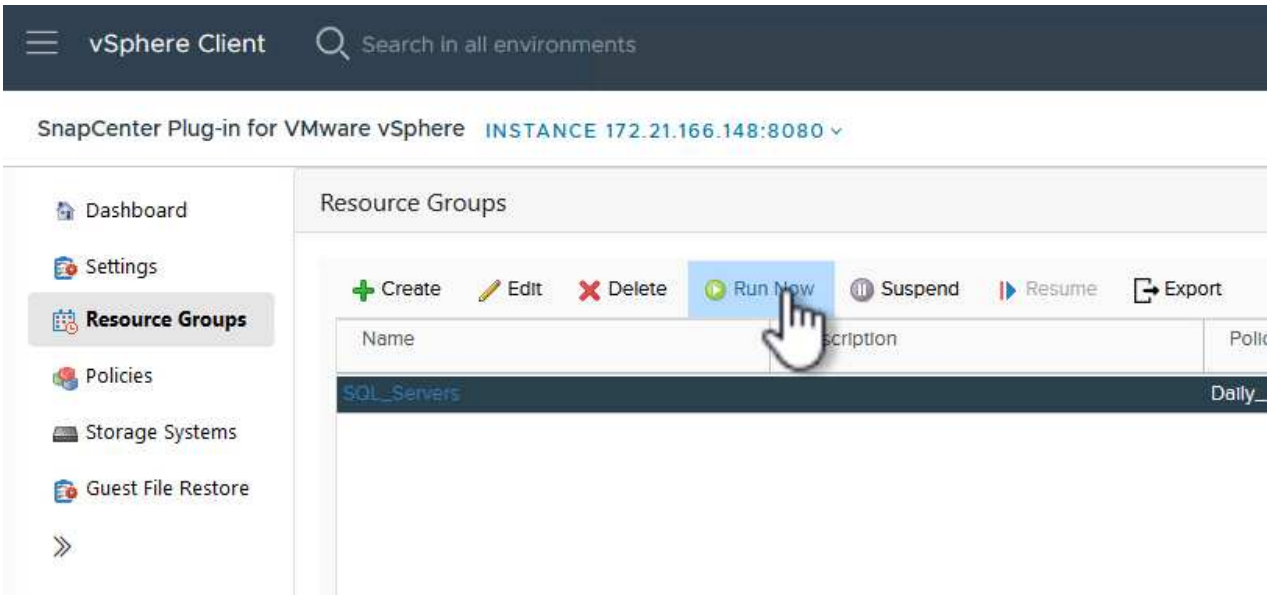

9. Navigate to the **Dashboard** and, under **Recent Job Activities** click on the number next to **Job ID** to open the job monitor and view the progress of the running job.

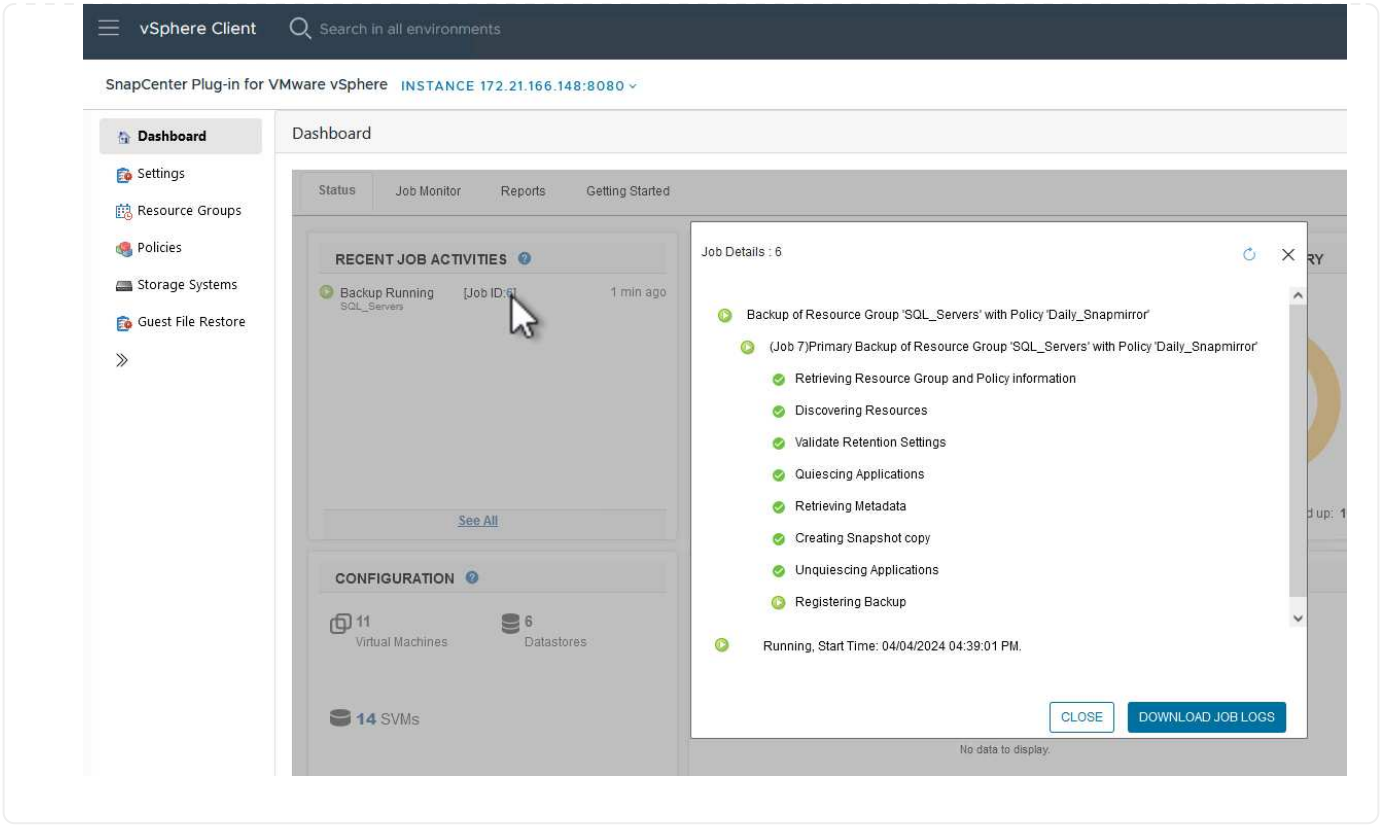

Use SCV to restore VMs, VMDKs and files

The SnapCenter Plug-in allows restores of VMs, VMDKs, files, and folders from primary or secondary backups.

VMs can be restored to the original host, or to an alternate host in the same vCenter Server, or to an alternate ESXi host managed by the same vCenter or any vCenter in linked mode.

vVol VMs can be restored to the original host.

VMDKs in traditional VMs can be restored to either the original or to an alternate datastore.

VMDKs in vVol VMs can be restored to the original datastore.

Individual files and folders in a guest file restore session can be restored, which attaches a backup copy of a virtual disk and then restores the selected files or folders.

Complete the following steps to restore VMs, VMDKs or individual folders.

Restore VMs using SnapCenter Plug-in

Complete the following steps to restore a VM with SCV:

1. Navigate to the VM to be restored in the vSphere client, right click and navigate to **SnapCenter Plug-in for VMware vSphere**. Select **Restore** from the sub-menu.

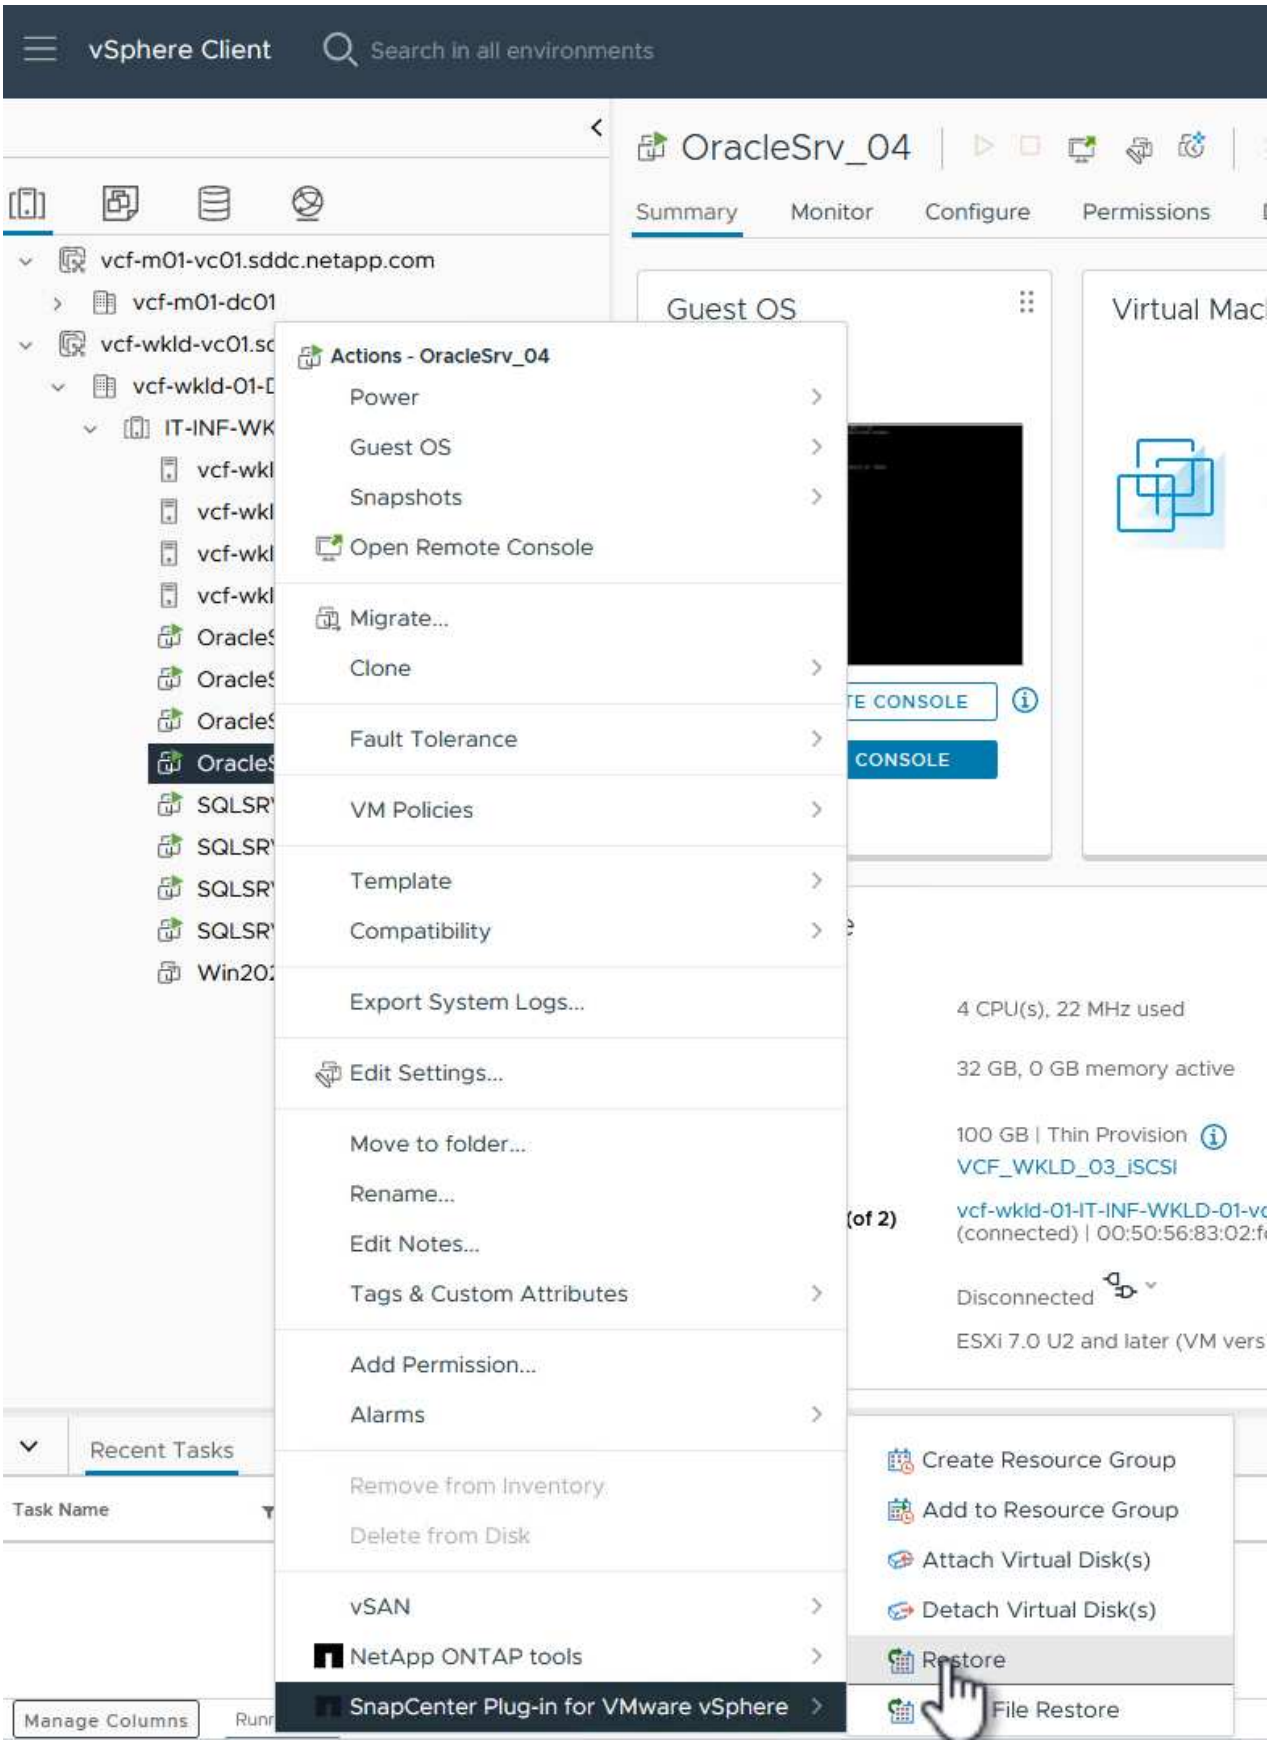
vSphere Client
Search in all environments
OracleSrv_04
Summary
Monitor
Configure
Permissions
Guest OS
Virtual Mac
vcf-m01-vc01.sddc.netapp.com
vcf-m01-dc01
vcf-wkld-vc01.sc
vcf-wkld-01-D
IT-INF-WK
vcf-wkl
vcf-wkl
vcf-wkl
vcf-wkl
Oracle
Oracle
Oracle
Oracle
SQLSR
SQLSR
SQLSR
SQLSR
Win20
Actions - OracleSrv_04
Power
Guest OS
Snapshots
Open Remote Console
Migrate...
Clone
Fault Tolerance
VM Policies
Template
Compatibility
Export System Logs...
Edit Settings...
Move to folder...
Rename...
Edit Notes...
Tags & Custom Attributes
Add Permission...
Alarms
Remove from Inventory
Delete from Disk
vSAN
NetApp ONTAP tools
SnapCenter Plug-in for VMware vSphere
TE CONSOLE
CONSOLE
4 CPU(s), 22 MHz used
32 GB, 0 GB memory active
100 GB | Thin Provision
VCF_WKLD_03_iSCSI
(of 2)
vcf-wkld-01-IT-INF-WKLD-01-vc
(connected) | 00:50:56:83:02:f
Disconnected
ESXi 7.0 U2 and later (VM vers
Recent Tasks
Task Name
Manage Columns
Create Resource Group
Add to Resource Group
Attach Virtual Disk(s)
Detach Virtual Disk(s)
Restore
File Restore

An alternative is to navigate to the datastore in inventory and then under the **Configure** tab go to **SnapCenter Plug-in for VMware vSphere > Backups**. From the chosen backup, select the VMs to be restored.

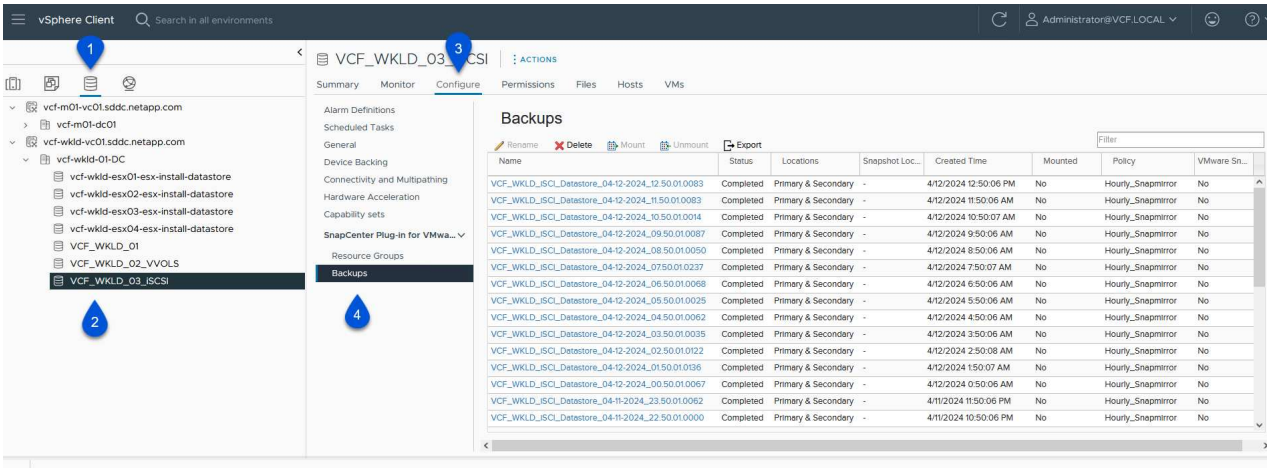

2. In the **Restore** wizard select the backup to be used. Click on **Next** to continue.

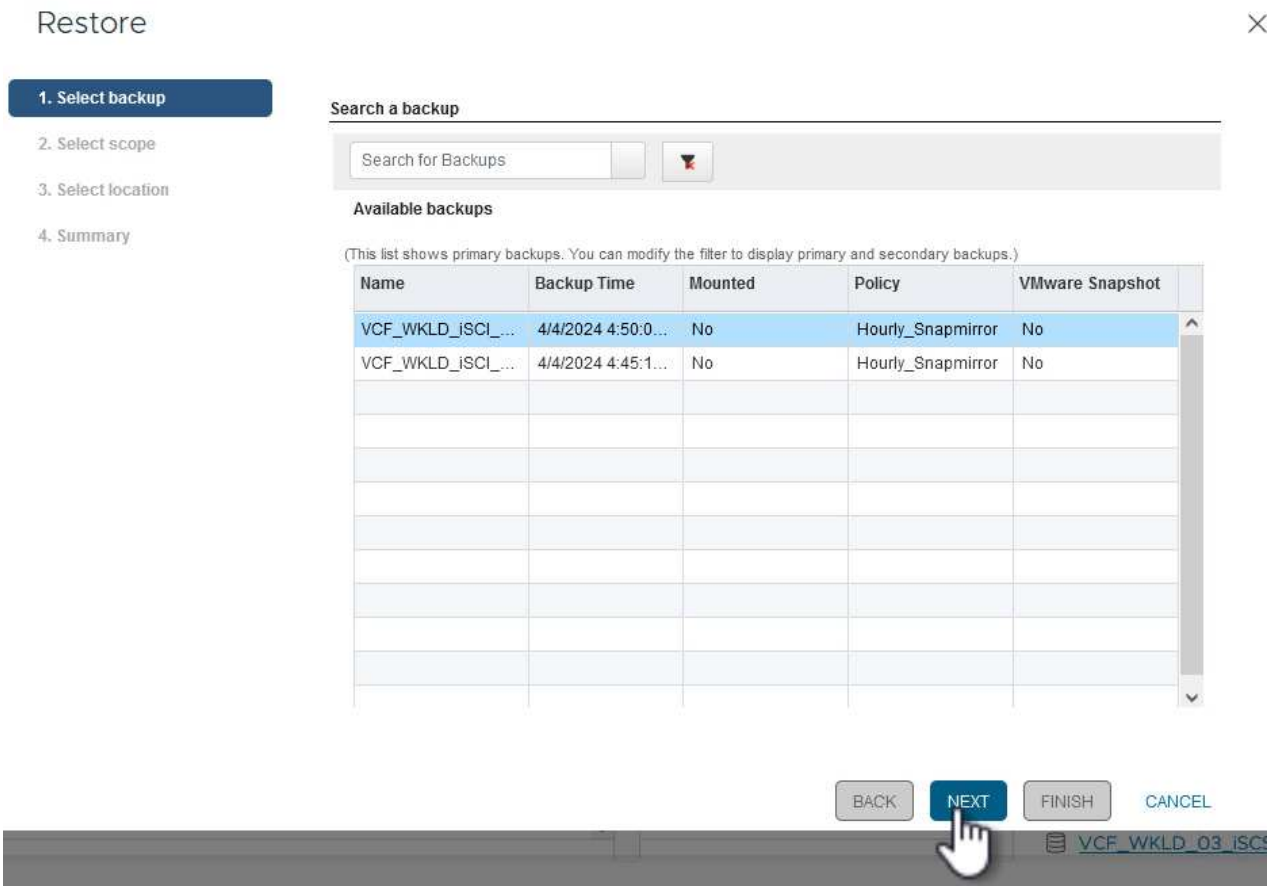

3. On the **Select scope** page fill out all required fields:
   - **Restore scope** - Select to restore the entire virtual machine.

- **Restart VM** - Choose whether to start the VM after the restore.
- **Restore Location** - Choose to restore to the orginal location or to an alternate location. When choosing alternate location select the options from each of the fields:
  - **Destination vCenter Server** - local vCenter or alternate vCenter in linked mode
  - **Destination ESXi host**
  - **Network**
  - **VM name after restore**
  - **Select datastore:**

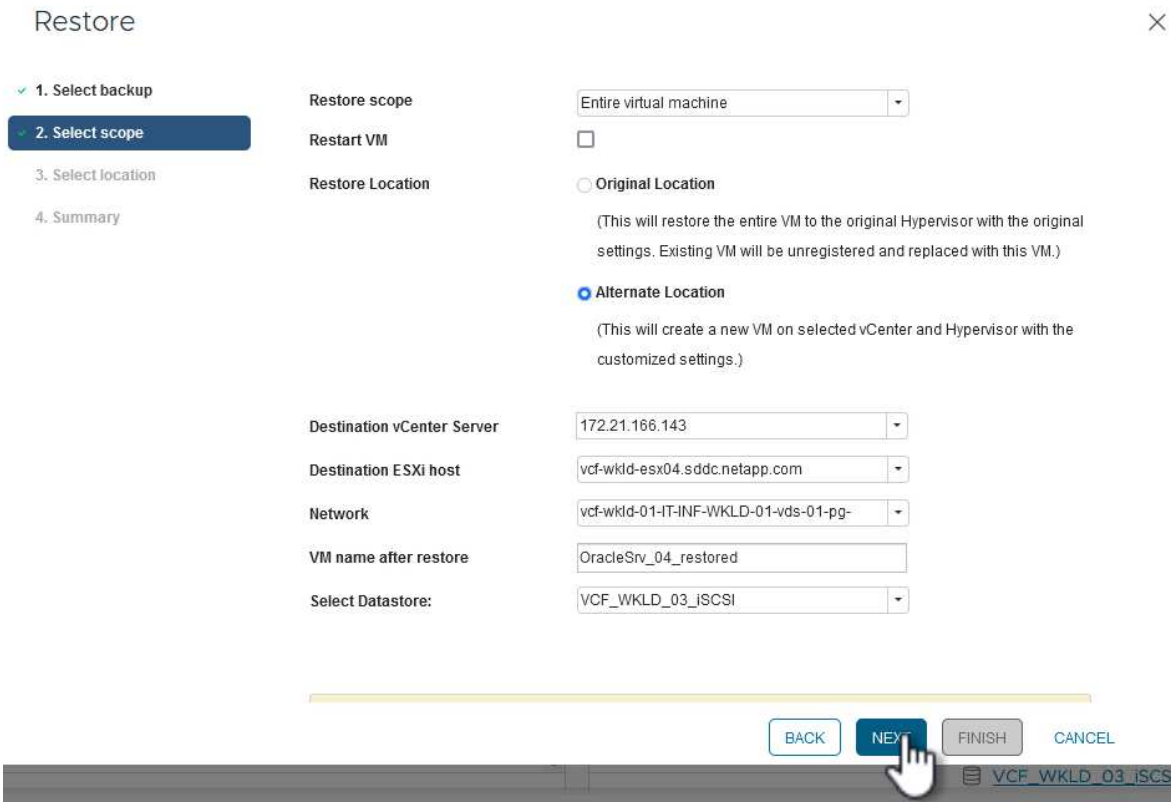

Click on **Next** to continue.

4. On the **Select location** page, choose to restore the VM from the primary or secondary ONTAP storage system. Click on **Next** to continue.

5. Finally, review the **Summary** and click on **Finish** to start the restore job.

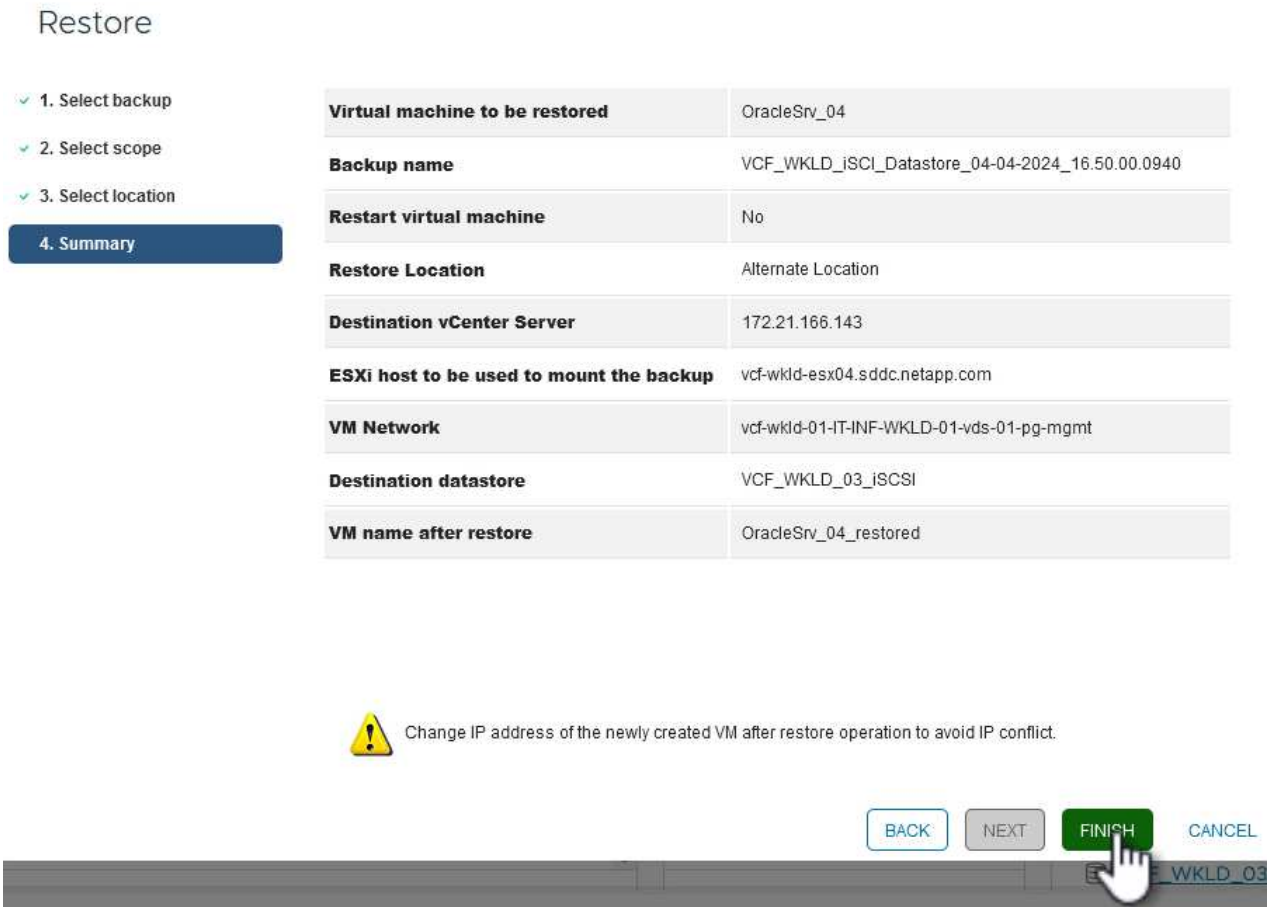

6. The restore job progress can be monitored from the **Recent Tasks** pane in the vSphere Client and from the job monitor in SCV.

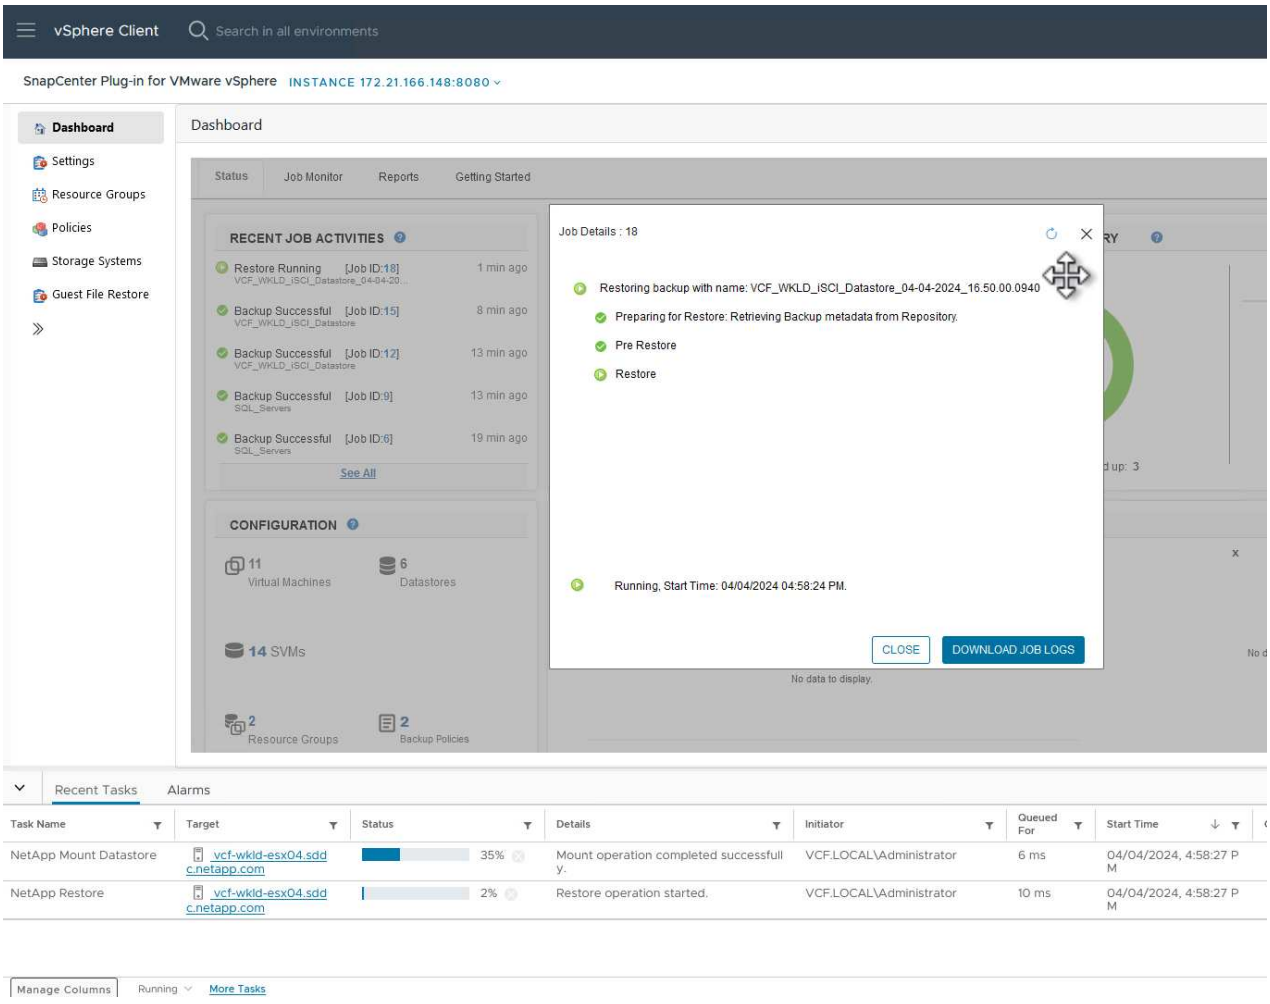
vSphere Client
Search in all environments
SnapCenter Plug-in for VMware vSphere
INSTANCE 172.21.166.148:8080
Dashboard
Settings
Resource Groups
Policies
Storage Systems
Guest File Restore
Dashboard
Status
Job Monitor
Reports
Getting Started
RECENT JOB ACTIVITIES
Restore Running [Job ID:18] 1 min ago
Backup Successful [Job ID:15] 8 min ago
Backup Successful [Job ID:12] 13 min ago
Backup Successful [Job ID:9] 13 min ago
Backup Successful [Job ID:6] 19 min ago
See All
CONFIGURATION
11 Virtual Machines
6 Datastores
14 SVMs
2 Resource Groups
2 Backup Policies
Job Details : 18
Restoring backup with name: VCF_WKLD_iSCI_Datastore_04-04-2024_16.50.00.0940
Preparing for Restore: Retrieving Backup metadata from Repository.
Pre Restore
Restore
Running, Start Time: 04/04/2024 04:58:24 PM.
CLOSE
DOWNLOAD JOB LOGS
No data to display.
Recent Tasks
Alarms
Task Name
Target
Status
Details
Initiator
Queued For
Start Time
NetApp Mount Datastore
vcf-wkld-esx04.sddc.netapp.com
35%
Mount operation completed successfully.
VCF.LOCAL\Administrator
6 ms
04/04/2024, 4:58:27 PM
NetApp Restore
vcf-wkld-esx04.sddc.netapp.com
2%
Restore operation started.
VCF.LOCAL\Administrator
10 ms
04/04/2024, 4:58:27 PM
Manage Columns
Running
More Tasks

Restore VMDKs using SnapCenter Plug-in

ONTAP Tools allows full restore of VMDK's to their original location or the ability to attach a VMDK as a new disk to a host system. In this scenario a VMDK will be attached to a Windows host in order to access the file system.

To attach a VMDK from a backup, complete the following steps:

1. In the vSphere Client navigate to a VM and, from the **Actions** menu, select **SnapCenter Plug-in for VMware vSphere > Attach Virtual Disk(s)**.

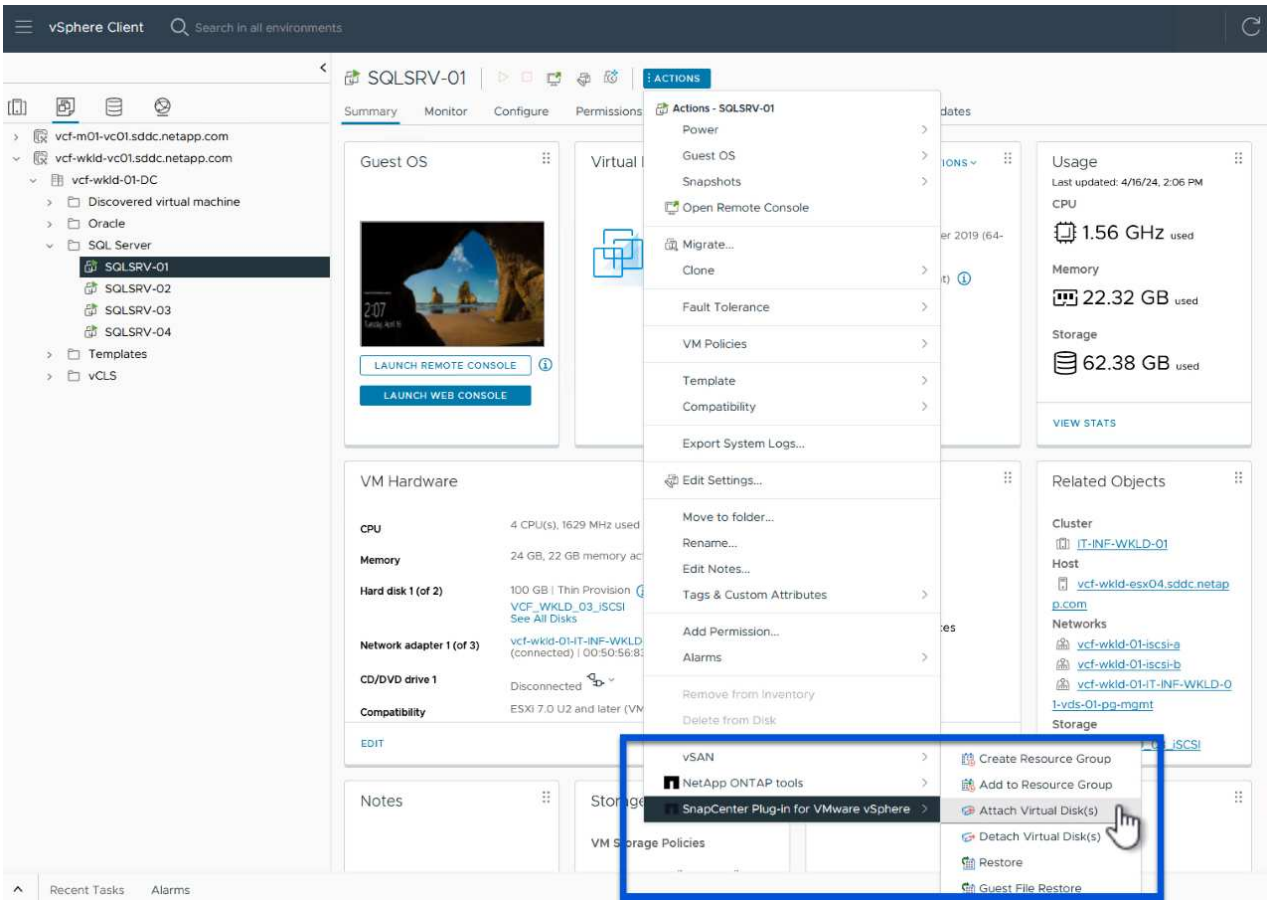

2. In the **Attach Virtual Disk(s)** wizard, select the backup instance to be used and the particular VMDK to be attached.

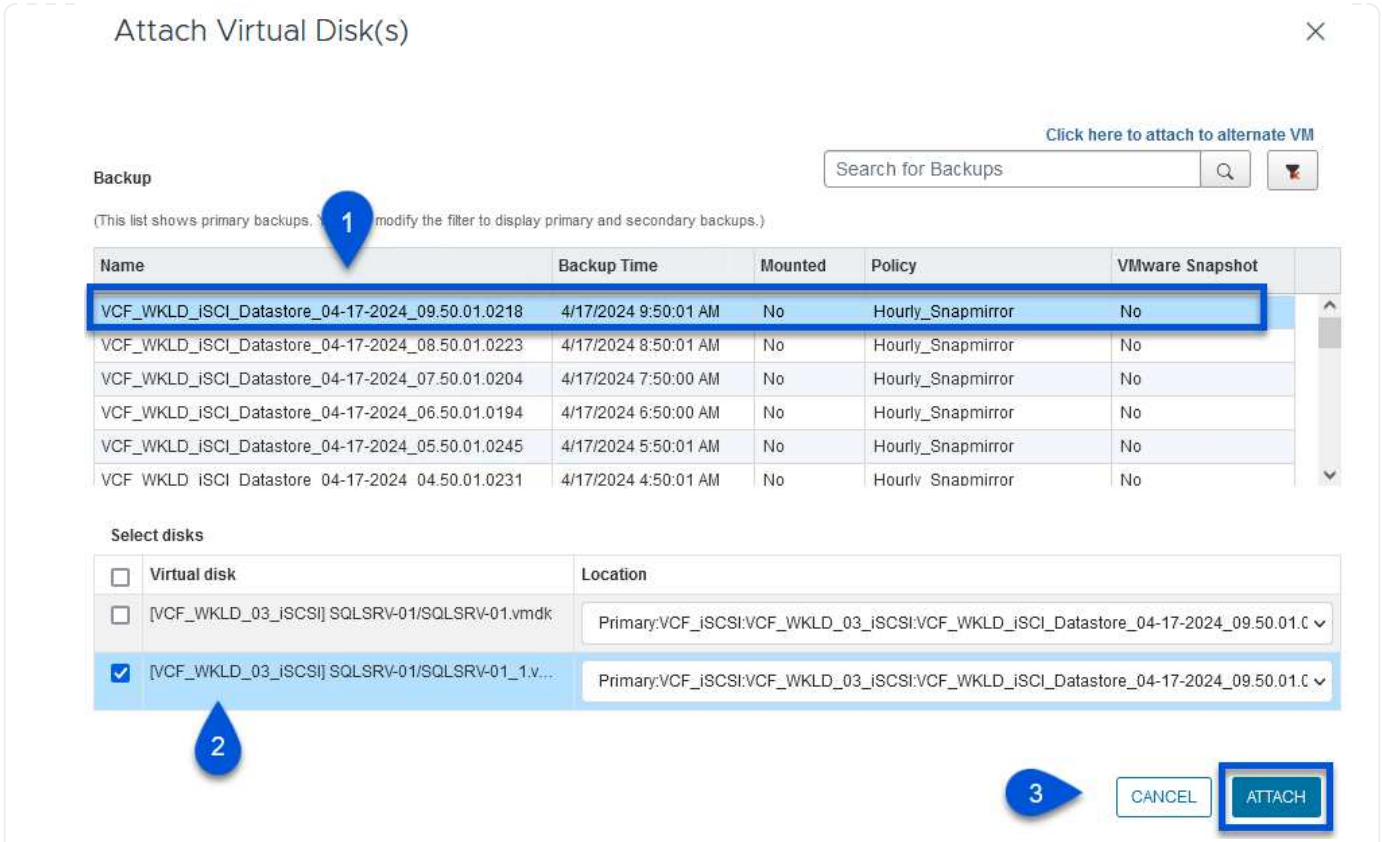

Filter options can be used to locate backups and to display backups from both primary and secondary storage systems.

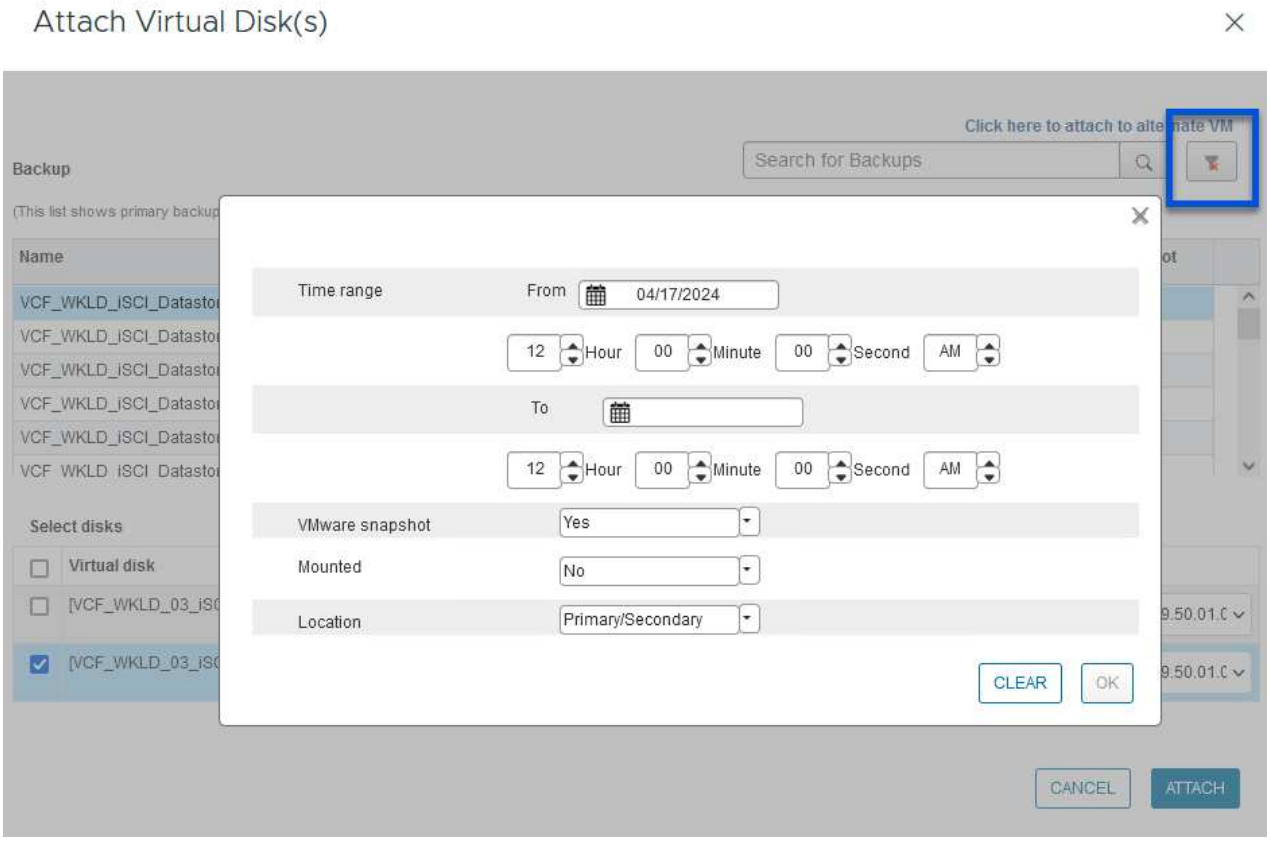

3. After selecting all options, click on the **Attach** button to begin the restore process and attached the VMDK to the host.
4. Once the attach procedure is complete the disk can be accessed from the OS of the host system. In this case SCV attached the disk with its NTFS file system to the E: drive of our Windows SQL Server and the SQL database files on the file system are accessible through File Explorer.

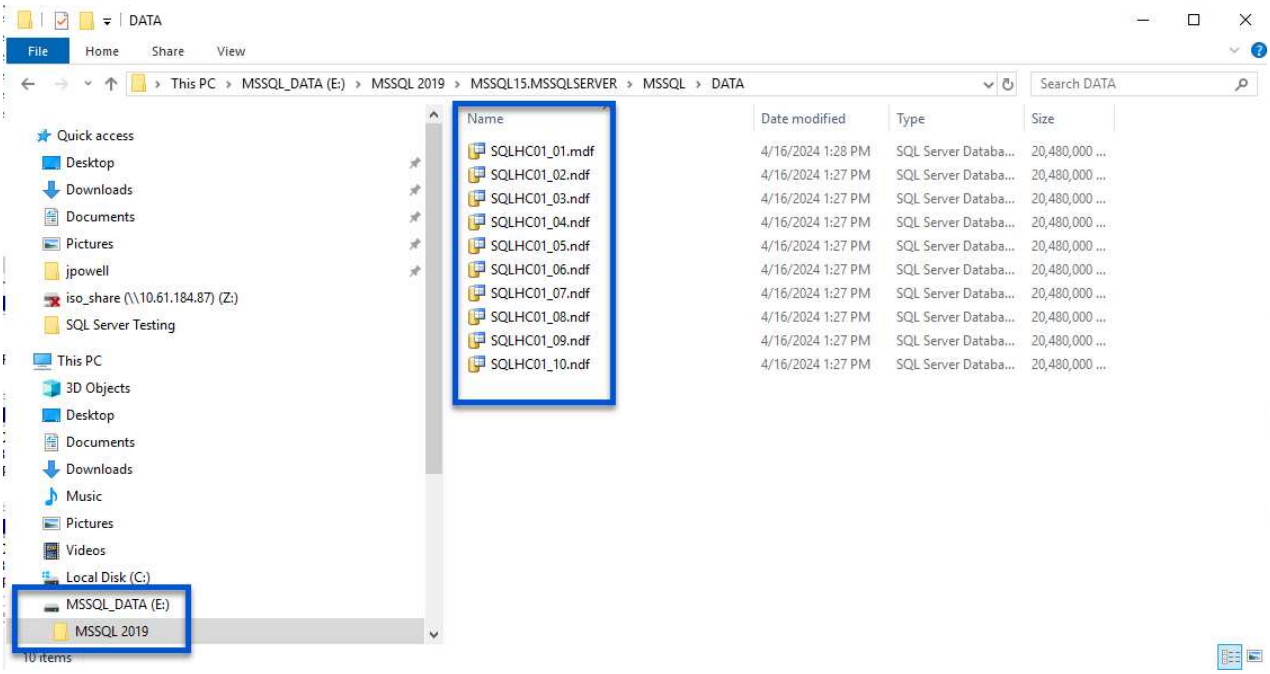

Guest File System Restore using SnapCenter Plug-in

ONTAP Tools features guest file system restores from a VMDK on Windows Server OSes. This is preformed centrally from the SnapCenter Plug-in interface.

For detailed information refer to Restore guest files and folders at the SCV documentation site.

To perform a guest file system restore for a Windows system, complete the following steps:

1. The first step is to create Run As credentials to provide access to the Windows host system. In the vSphere Client navigate to the CSV plug-in interface and click on **Guest File Restore** in the main menu.

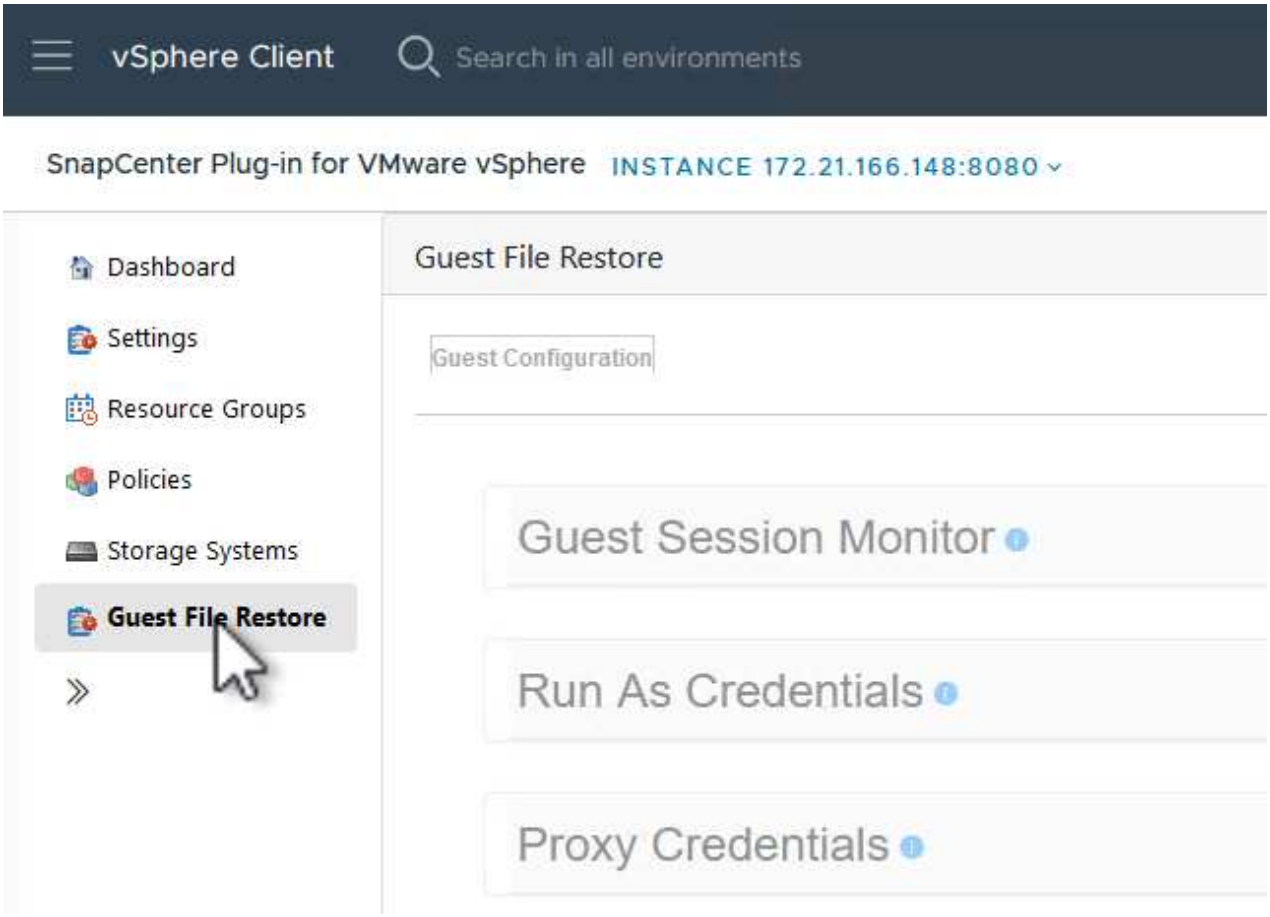

2. Under **Run As Credentials** click on the **+** icon to open the **Run As Credentials** window.
3. Fill in a name for the credentials record, an administrator username and password for the Windows system, and then click on the **Select VM** button to select an optional Proxy VM to be used for the restore.
image::vmware-vcf-asa-image85.png[Run as credentials window]

4. On the Proxy VM page provide a name for the VM and locate it by searching by ESXi host or by name. Once selected, click on **Save**.

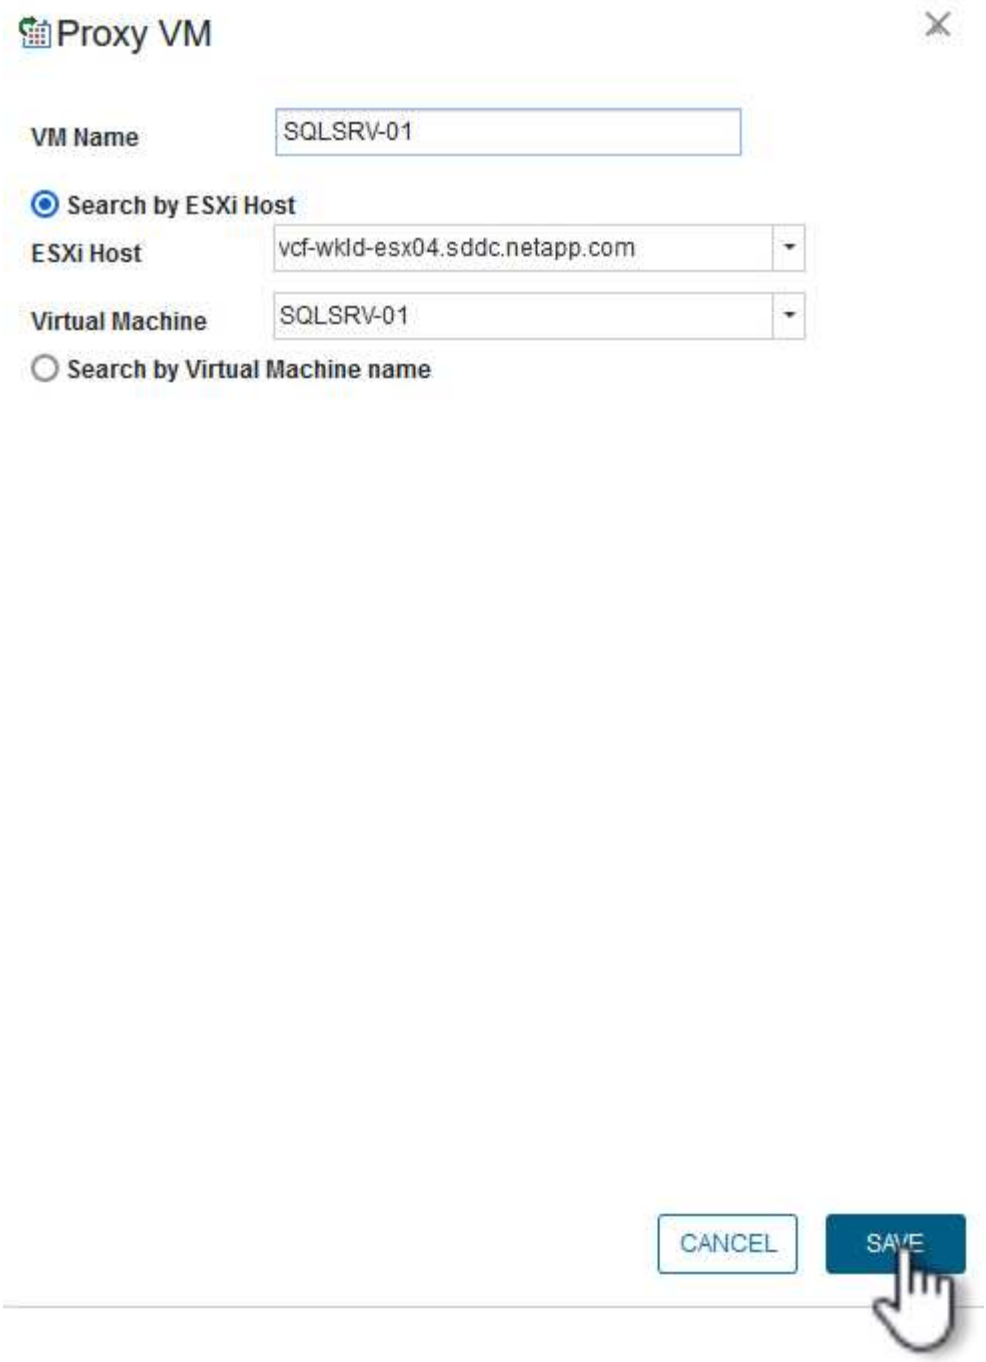

5. Click on **Save** again in the **Run As Credentials** window to complete saving the record.
6. Next, navigate to a VM in the inventory. From the **Actions** menu, or by right-clicking on the VM, select **SnapCenter Plug-in for VMware vSphere > Guest File Restore**.

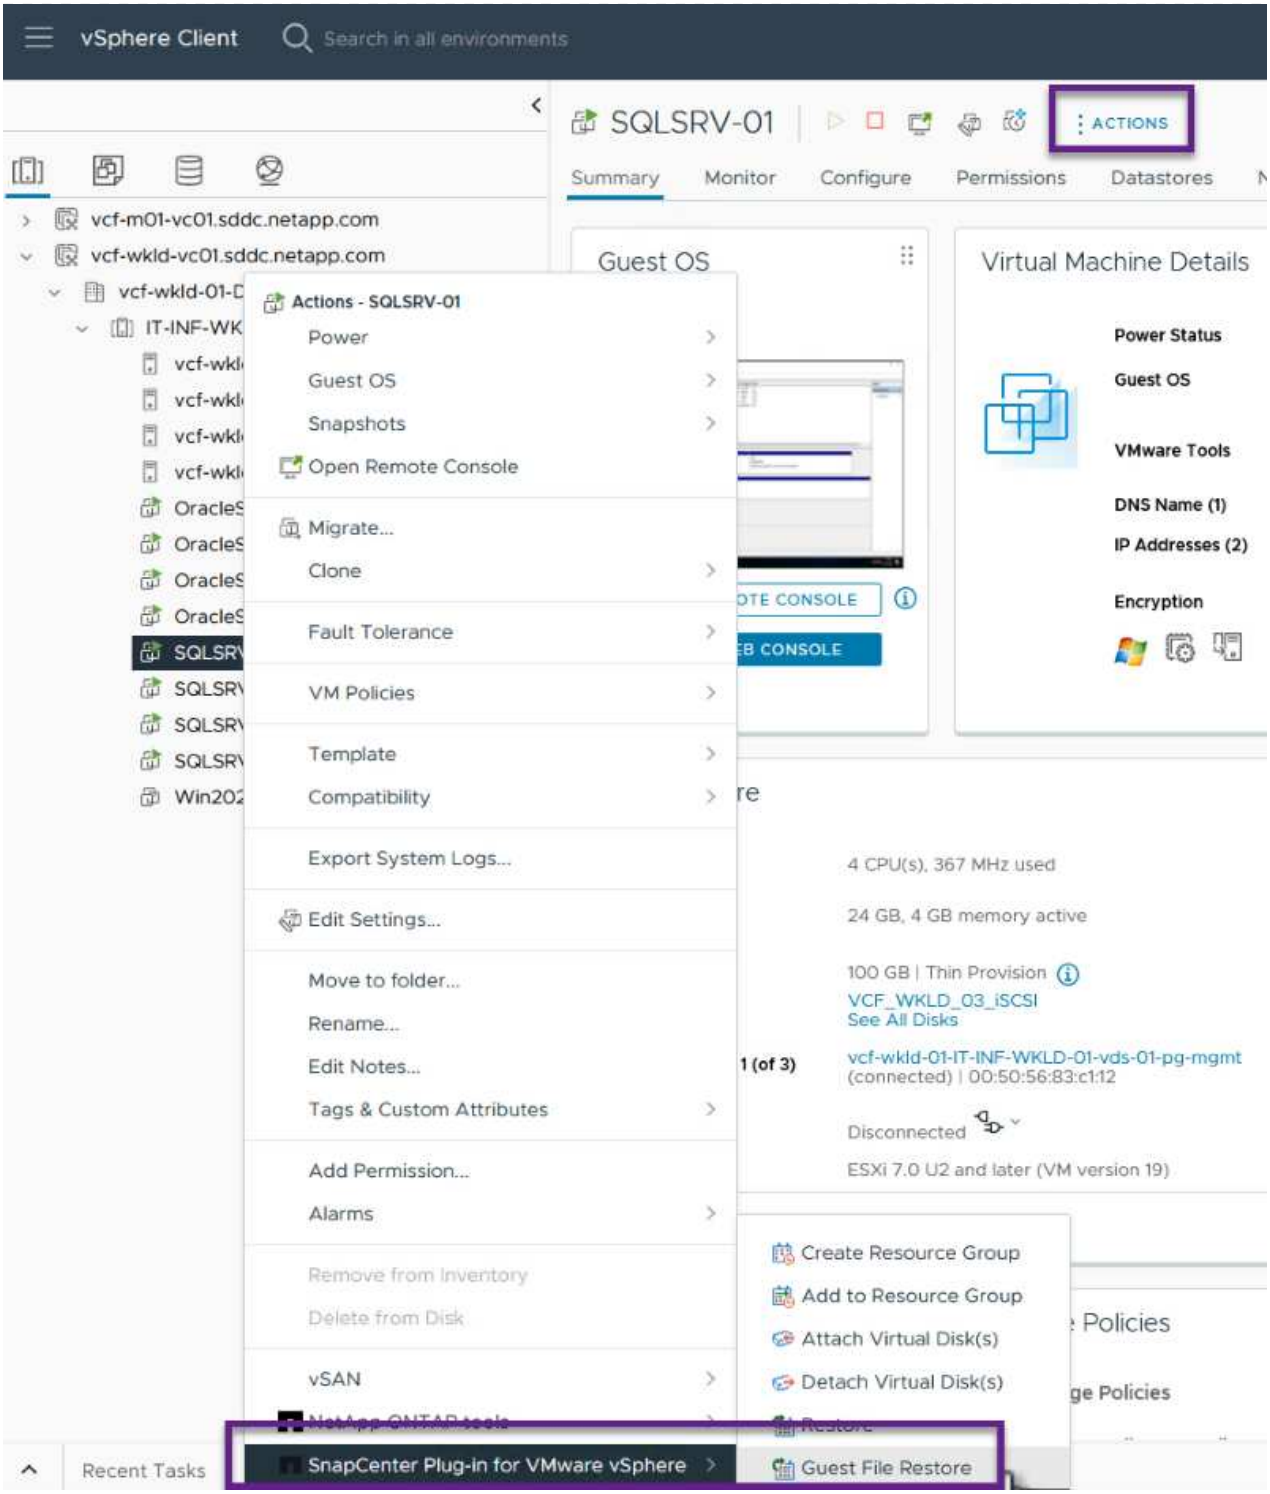

7. On the **Restore Scope** page of the **Guest File Restore** wizard, select the backup to restore from, the particular VMDK, and the location (primary or secondary) to restore the VMDK from. Click on **Next** to continue.

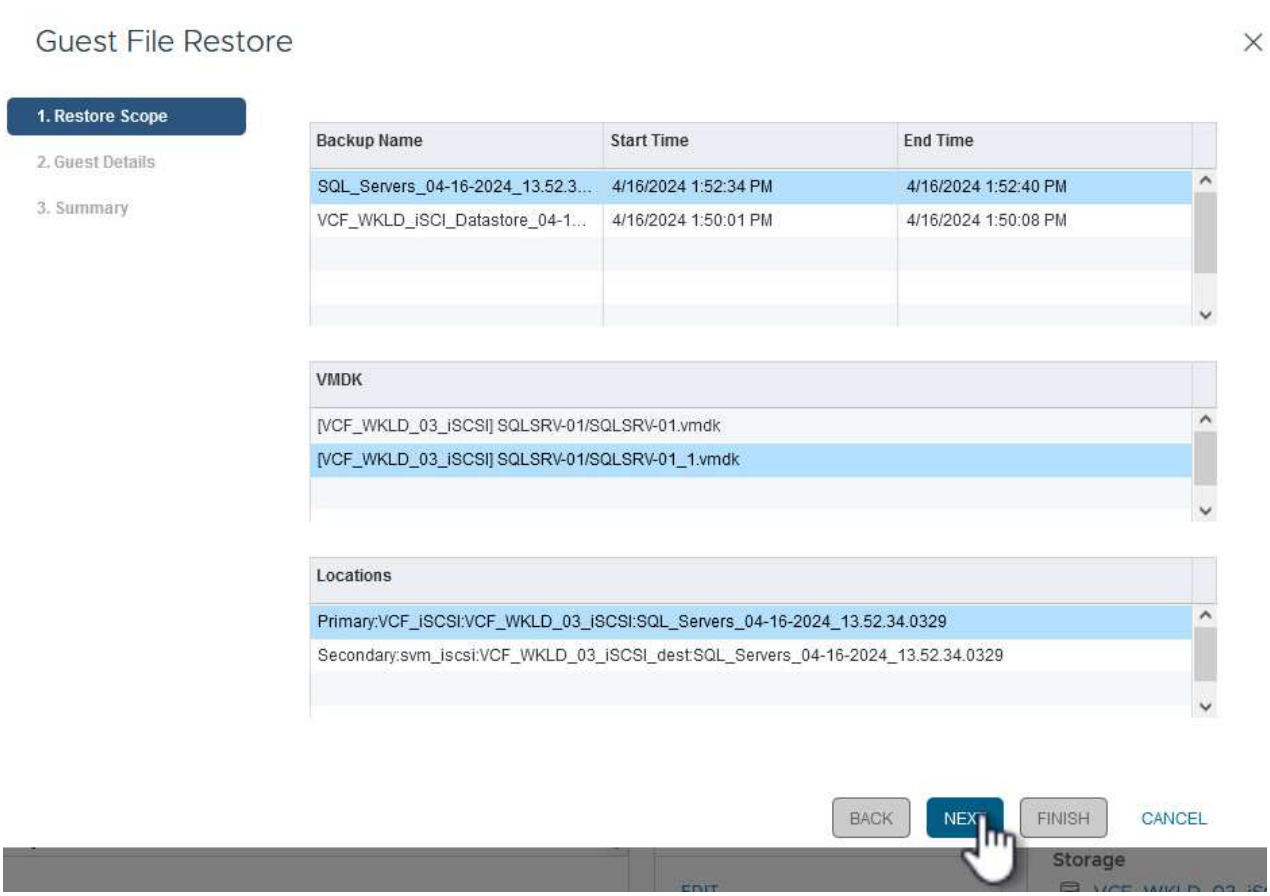

8. On the **Guest Details** page, select to use **Guest VM** or **Use Gues File Restore proxy VM** for the restore. Also, fill out email notification settings here if desired. Click on **Next** to continue.

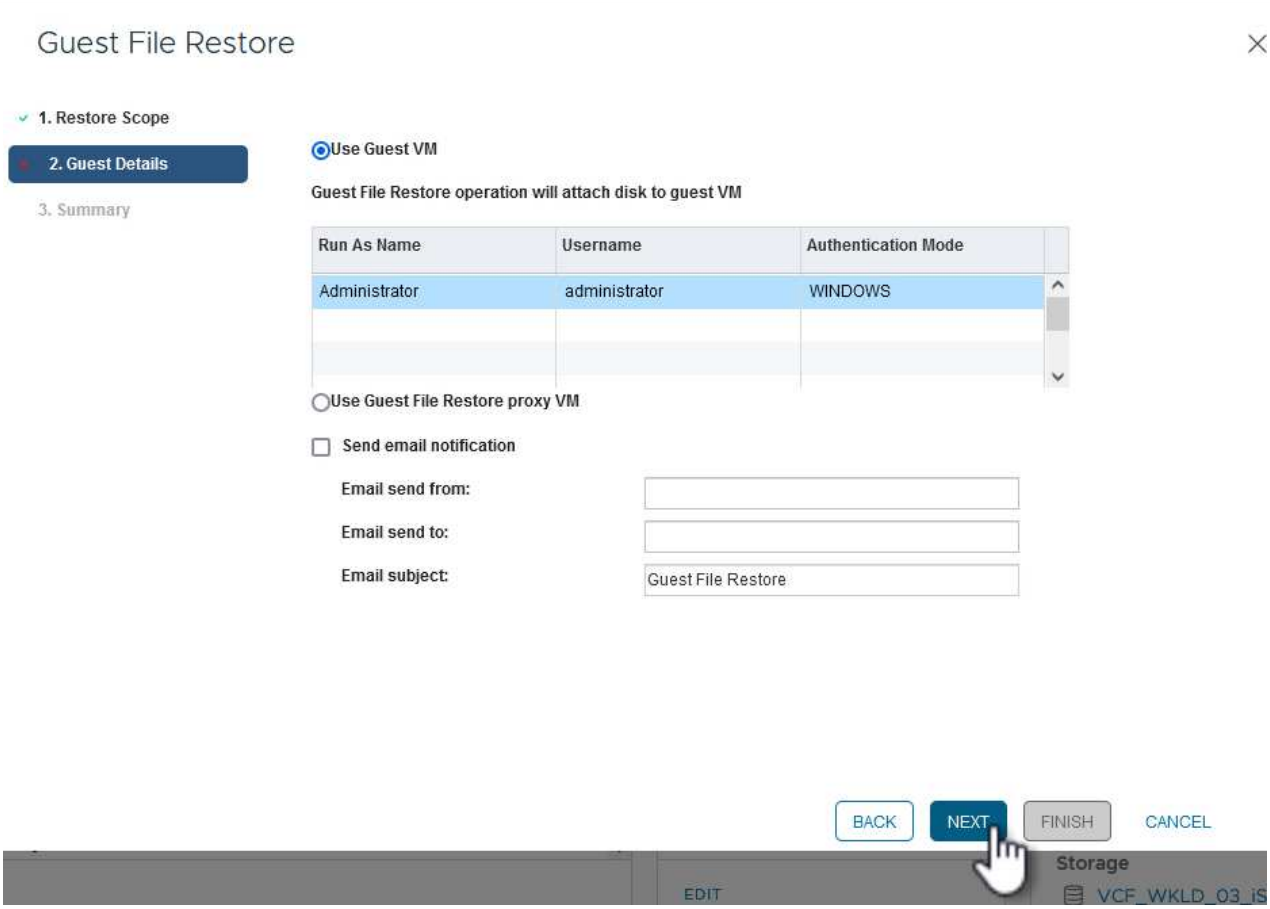

9. Finally, review the **Summary** page and click on **Finish** to begin the Guest File System Restore session.
10. Back in the SnapCenter Plug-in interface, navigate to **Guest File Restore** again and view the running session under **Guest Session Monitor**. Click on the icon under **Browse Files** to continue.

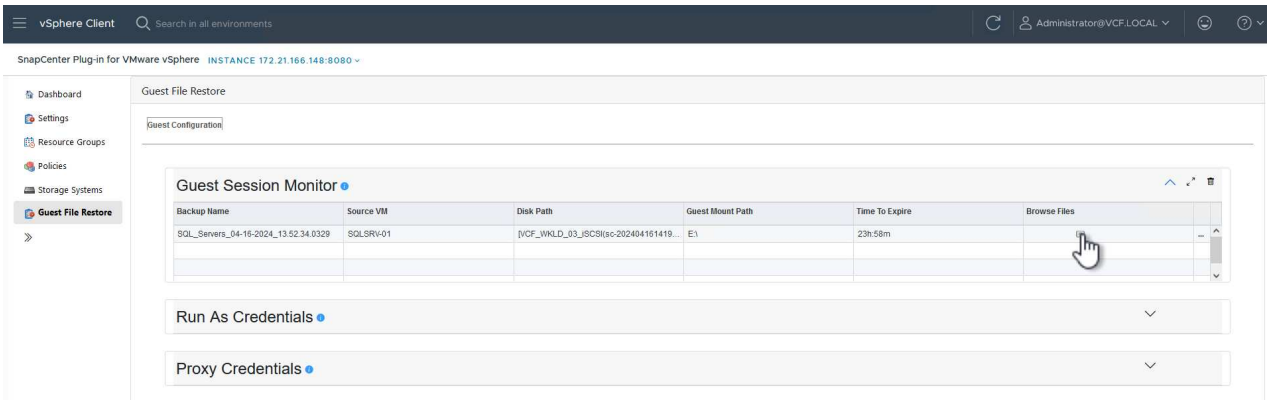

11. In the **Guest File Browse** wizard select the folder or files to restore and the file system location to restore them to. Finally, click on **Restore** to start the **Restore** process.

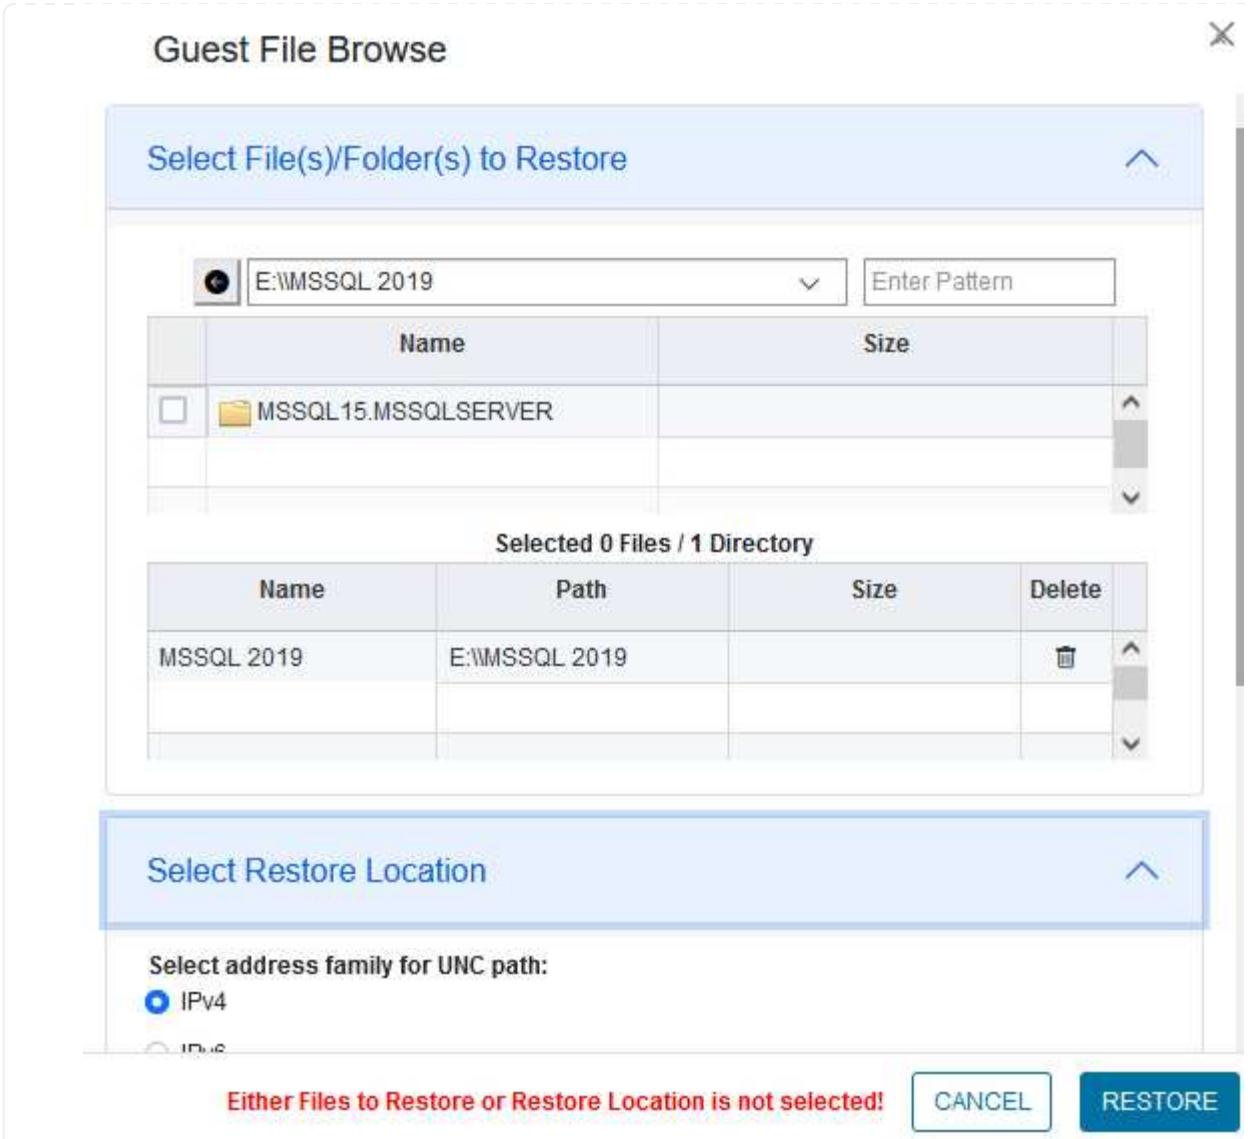
Guest File Browse
Select File(s)/Folder(s) to Restore
E:\MSSQL 2019
Enter Pattern
Name
Size
MSSQL15.MSSQLSERVER
Selected 0 Files / 1 Directory
Name
Path
Size
Delete
MSSQL 2019
E:\MSSQL 2019
Select Restore Location
Select address family for UNC path:
IPv4
Either Files to Restore or Restore Location is not selected!
CANCEL
RESTORE

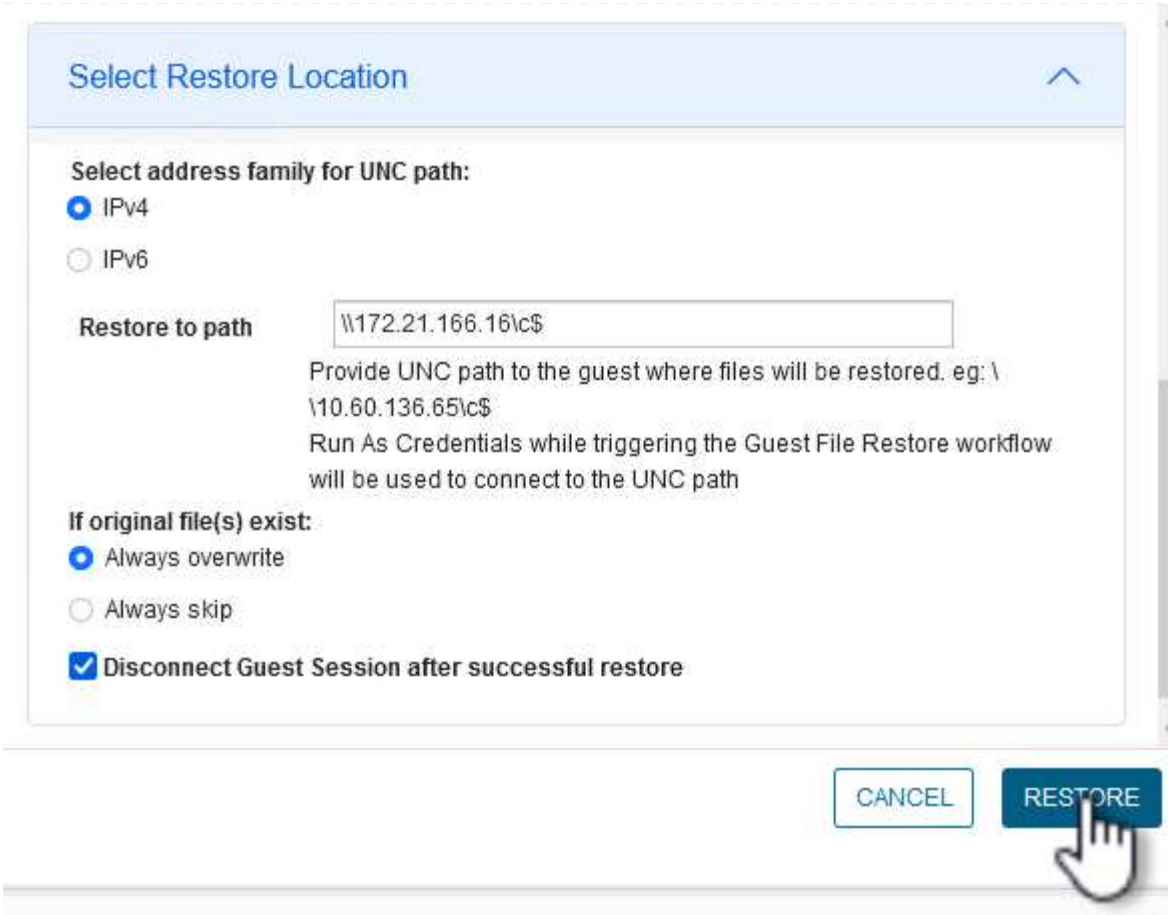

12. The restore job can be monitored from the vSphere Client task pane.

Additional information

For information on configuring VCF refer to VMware Cloud Foundation Documentation.

For information on configuring ONTAP storage systems refer to the ONTAP 9 Documentation center.

For information on using the SnapCenter Plug-in for VMware vSphere refer to the SnapCenter Plug-in for VMware vSphere documentation.

VMware Cloud Foundation (VCF) is an integrated software defined data center (SDDC) platform that provides a complete stack of software-defined infrastructure for running enterprise applications in a hybrid cloud environment. It combines compute, storage, networking, and management capabilities into a unified platform, offering a consistent operational experience across private and public clouds.

Author: Josh Powell, Ravi BCB

VMware Cloud Foundation with NetApp AFF Arrays

This document provides information on storage options available for VMware Cloud Foundation using the NetApp All-Flash AFF storage system. Supported storage options are covered with specific instruction for

creating workload domains with NFS and vVol datastores as principal storage as well as a range of supplemental storage options.

Use Cases

Use cases covered in this documentation:

- Storage options for customers seeking uniform environments across both private and public clouds.
- Automated solution for deploying virtual infrastructure for workload domains.
- Scalable storage solution tailored to meet evolving needs, even when not aligned directly with compute resource requirements.
- Deploy VCF VI Workload Domains using ONTAP as principal storage.
- Deploy supplemental storage to VI Workload Domains using ONTAP Tools for VMware vSphere.

Audience

This solution is intended for the following people:

- Solution architects looking for more flexible storage options for VMware environments that are designed to maximize TCO.
- Solution architects looking for VCF storage options that provide data protection and disaster recovery options with the major cloud providers.
- Storage administrators wanting to understand how to configure VCF with principal and supplemental storage.

Technology Overview

The VCF with NetApp AFF solution is comprised of the following major components:

VMware Cloud Foundation

VMware Cloud Foundation extends VMware's vSphere hypervisor offerings by combining key components such as SDDC Manager, vSphere, vSAN, NSX, and VMware Aria Suite to create a virtualized datacenter.

The VCF solution supports both native Kubernetes and virtual machine-based workloads. Key services such as VMware vSphere, VMware vSAN, VMware NSX-T Data Center, and VMware vRealize Cloud Management are integral components of the VCF package. When combined, these services establish a software-defined infrastructure capable of efficiently managing compute, storage, networking, security, and cloud management.

VCF is comprised of a single management domain and up to 24 VI Workload Domains that each represent a unit of application-ready infrastructure. A workload domain is comprised of one or more vSphere clusters managed by a single vCenter instance.

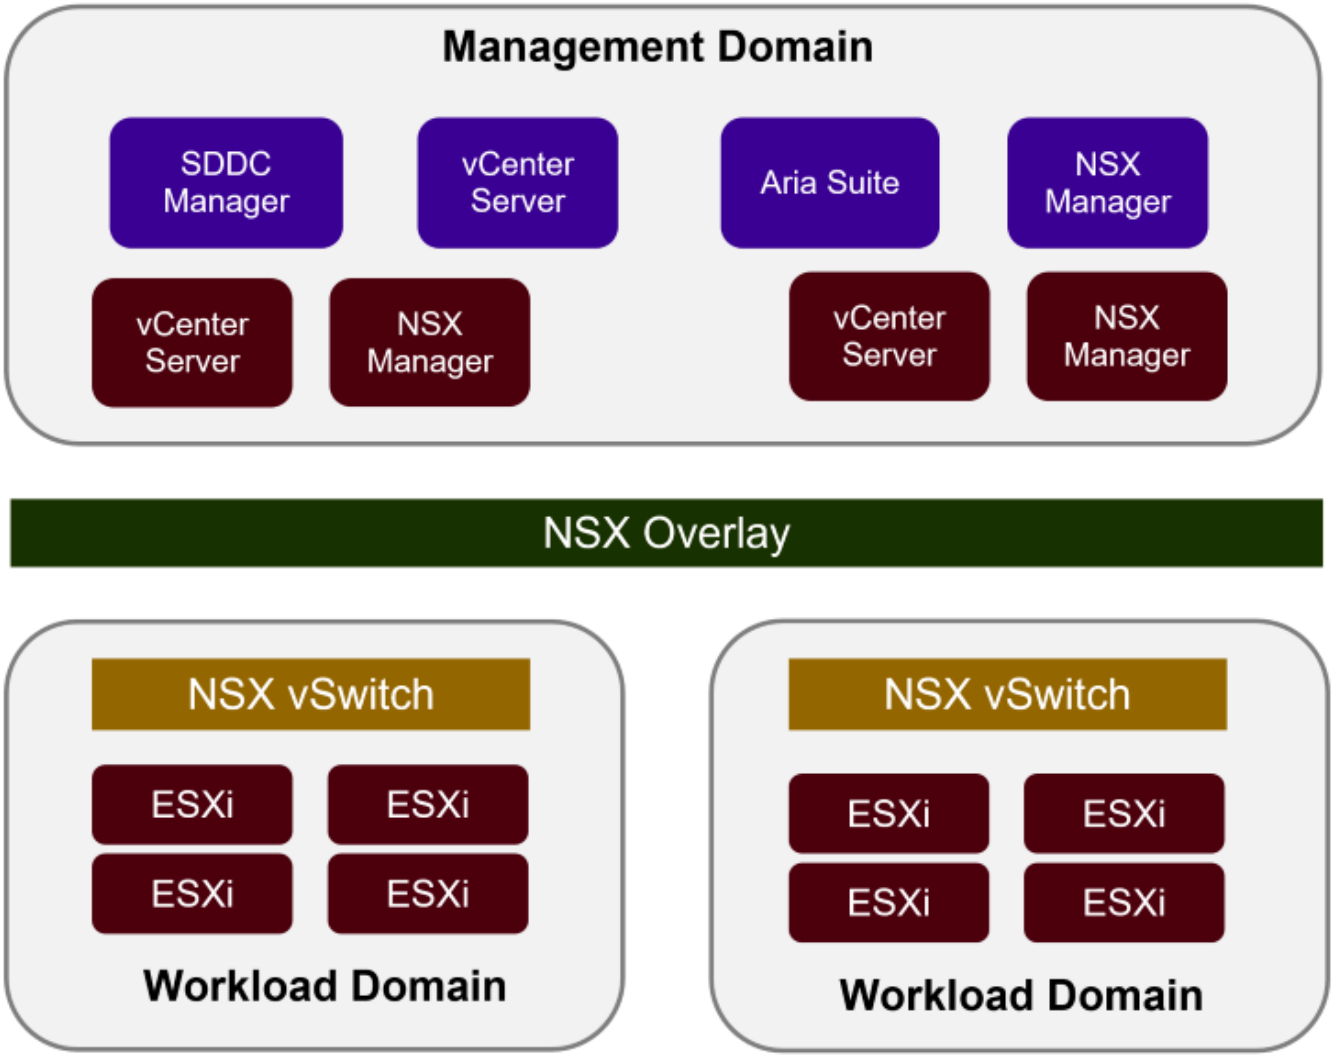

For more information on VCF architecture and planning, refer to Architecture Models and Workload Domain Types in VMware Cloud Foundation.

VCF Storage Options

VMware divides storage options for VCF into **principal** and **supplemental** storage. The VCF Management Domain must use vSAN as its principal storage. However, there are many supplemental storage options for the Management Domain and both principal and supplemental storage options available for VI Workload Domains.

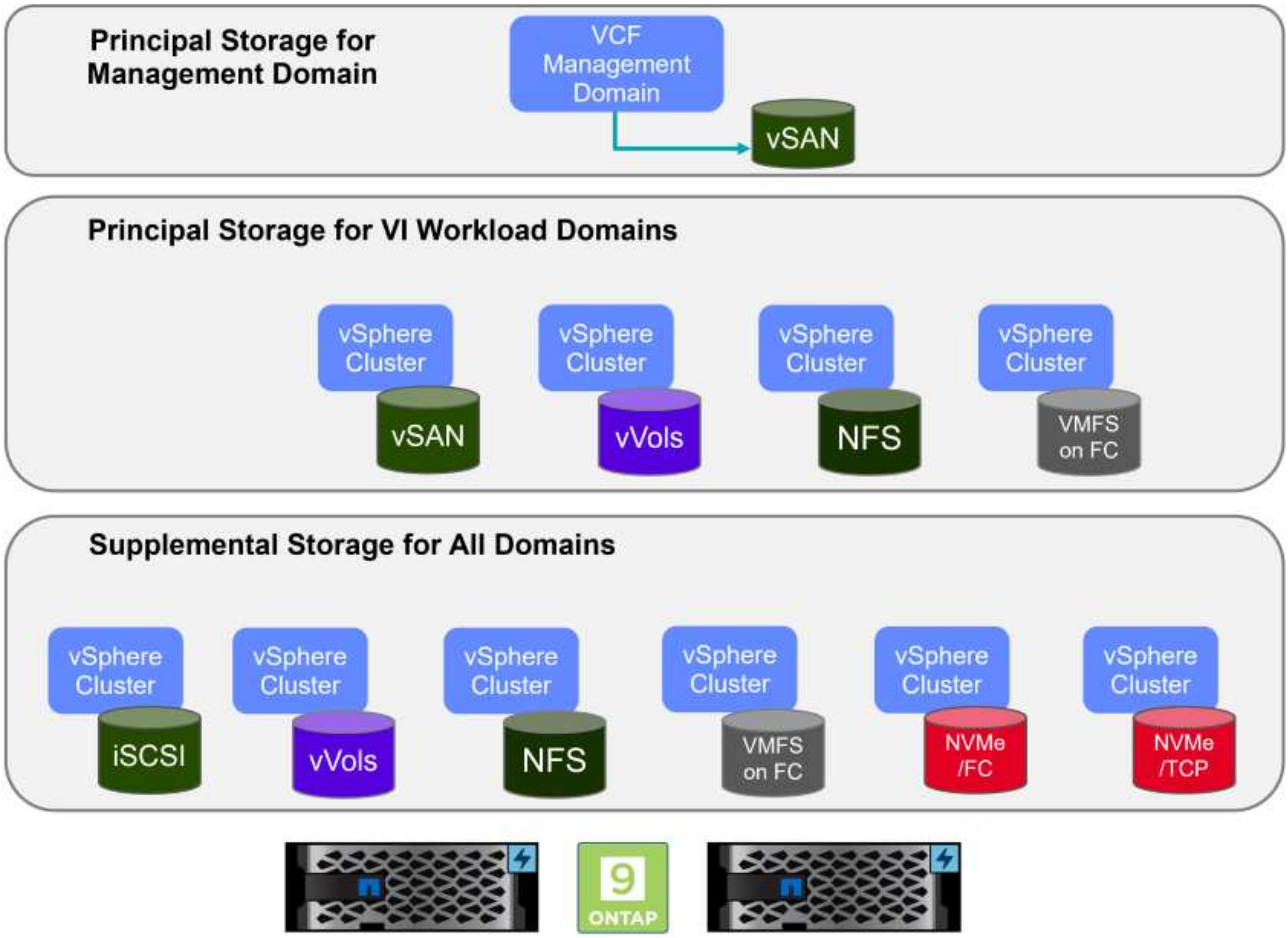

Principal Storage for Workload Domains
Principal Storage refers to any type of storage that can be directly connected to a VI Workload Domain during the setup process within SDDC Manager. Principal storage is the first datastore configured for a Workload Domain and includes vSAN, vVols (VMFS), NFS and VMFS on Fibre Channel.

Supplemental Storage for Management and Workload Domains
Supplemental storage is the storage type that can be added to the management or workload domains at any time after the cluster has been created. Supplemental storage represents the widest range of supported storage options, all of which are supported on NetApp AFF arrays.

Additional documentation resources for VMware Cloud Foundation:
* VMware Cloud Foundation Documentation
* Supported Storage Types for VMware Cloud Foundation
* Managing Storage in VMware Cloud Foundation

NetApp All-Flash Storage Arrays

NetApp AFF (All Flash FAS) arrays are high-performance storage solutions designed to leverage the speed and efficiency of flash technology. AFF arrays incorporate integrated data management features such as snapshot-based backups, replication, thin provisioning, and data protection capabilities.

NetApp AFF arrays utilize the ONTAP storage operating system, offering comprehensive storage protocol support for all storage options compatible with VCF, all within a unified architecture.

NetApp AFF storage arrays are available in the highest performing A-Series and a QLC flash-based C-Series. Both series use NVMe flash drives.

For more information on NetApp AFF A-Series storage arrays see the NetApp AFF A-Series landing page.

For more information on NetApp C-Series storage arrays see the NetApp AFF C-Series landing page.

NetApp ONTAP Tools for VMware vSphere

ONTAP Tools for VMware vSphere (OTV) allows administrators to manage NetApp storage directly from within the vSphere Client. ONTAP Tools allows you to deploy and manage datastores, as well as provision vVol datastores.

ONTAP Tools allows mapping of datastores to storage capability profiles which determine a set of storage system attributes. This allows the creation of datastores with specific attributes such as storage performance and QoS.

ONTAP Tools also includes a **VMware vSphere APIs for Storage Awareness (VASA) Provider** for ONTAP storage systems which enables the provisioning of VMware Virtual Volumes (vVols) datastores, creation and use of storage capability profiles, compliance verification, and performance monitoring.

For more information on NetApp ONTAP tools see the ONTAP tools for VMware vSphere Documentation page.

Solution Overview

In the scenarios presented in this documentation we will demonstrate how to use ONTAP storage systems as principal storage for VCF VI Workload Domain deployments. In addition, we will install and use ONTAP Tools for VMware vSphere to configure supplemental datastores for VI Workload Domains.

Scenarios covered in this documentation:

- **Configure and use an NFS datastore as principal storage during VI Workload Domain deployment.** Click **here** for deployment steps.
- **Install and demonstrate the use of ONTAP Tools to configure and mount NFS datastores as supplemental storage in VI Workload Domains.** Click **here** for deployment steps.

In this scenario we will demonstrate how to configure an NFS datastore as principal storage for the deployment of a VI Workload Domain in VCF. Where appropriate we will refer to external documentation for the steps that must be performed in VCF's SDDC Manager, and cover those steps that are specific to the storage configuration portion.

Author: Josh Powell, Ravi BCB

NFS as principal storage for VI Workload Domains

Scenario Overview

This scenario covers the following high level steps:

- Verify networking for the ONTAP storage virtual machine (SVM) and that a logical interface (LIF) is present to carry NFS traffic.

- Create an export policy to allow the ESXi hosts access to the NFS volume.
- Create an NFS volume on the ONTAP storage system.
- Create a Network Pool for NFS and vMotion traffic in SDDC Manager.
- Commission hosts in VCF for use in a VI Workload Domain.
- Deploy a VI Workload Domain in VCF using an NFS datastore as principal storage.
- Install NetApp NFS Plug-in for VMware VAAI

Prerequisites

This scenario requires the following components and configurations:

- NetApp AFF storage system with a storage virtual machine (SVM) configured to allow NFS traffic.
- Logical interface (LIF) has been created on the IP network that is to carry NFS traffic and is associated with the SVM.
- VCF management domain deployment is complete and the SDDC Manager interface is accessible.
- 4 x ESXi hosts configured for communication on the VCF management network.
- IP addresses reserved for vMotion and NFS storage traffic on the VLAN or network segment established for this purpose.

> When deploying a VI Workload Domain, VCF validates connectivity to the NFS Server. This is done using the management adapter on the ESXi hosts before any additional vmkernel adapter is added with the NFS IP address. Therefore, it is necessary to ensure that either 1) the management network is routable to the NFS Server, or 2) a LIF for the management network has been added to the SVM hosting the NFS datastore volume, to ensure that the validation can proceed.

For information on configuring ONTAP storage systems refer to the ONTAP 9 Documentation center.

For information on configuring VCF refer to VMware Cloud Foundation Documentation.

Deployment Steps

To deploy a VI Workload Domain with an NFS datastore as principal storage, complete the following steps:

Verify networking for ONTAP SVM

Verify that the required logical interfaces have been established for the network that will carry NFS traffic between the ONTAP storage cluster and VI Workload Domain.

1. From ONTAP System Manager navigate to **Storage VMs** in the left-hand menu and click on the SVM to be used for NFS traffic. On the **Overview** tab, under **NETWORK IP INTERFACES**, click on the numeric to the right of **NFS**. In the list verify that the required LIF IP addresses are listed.

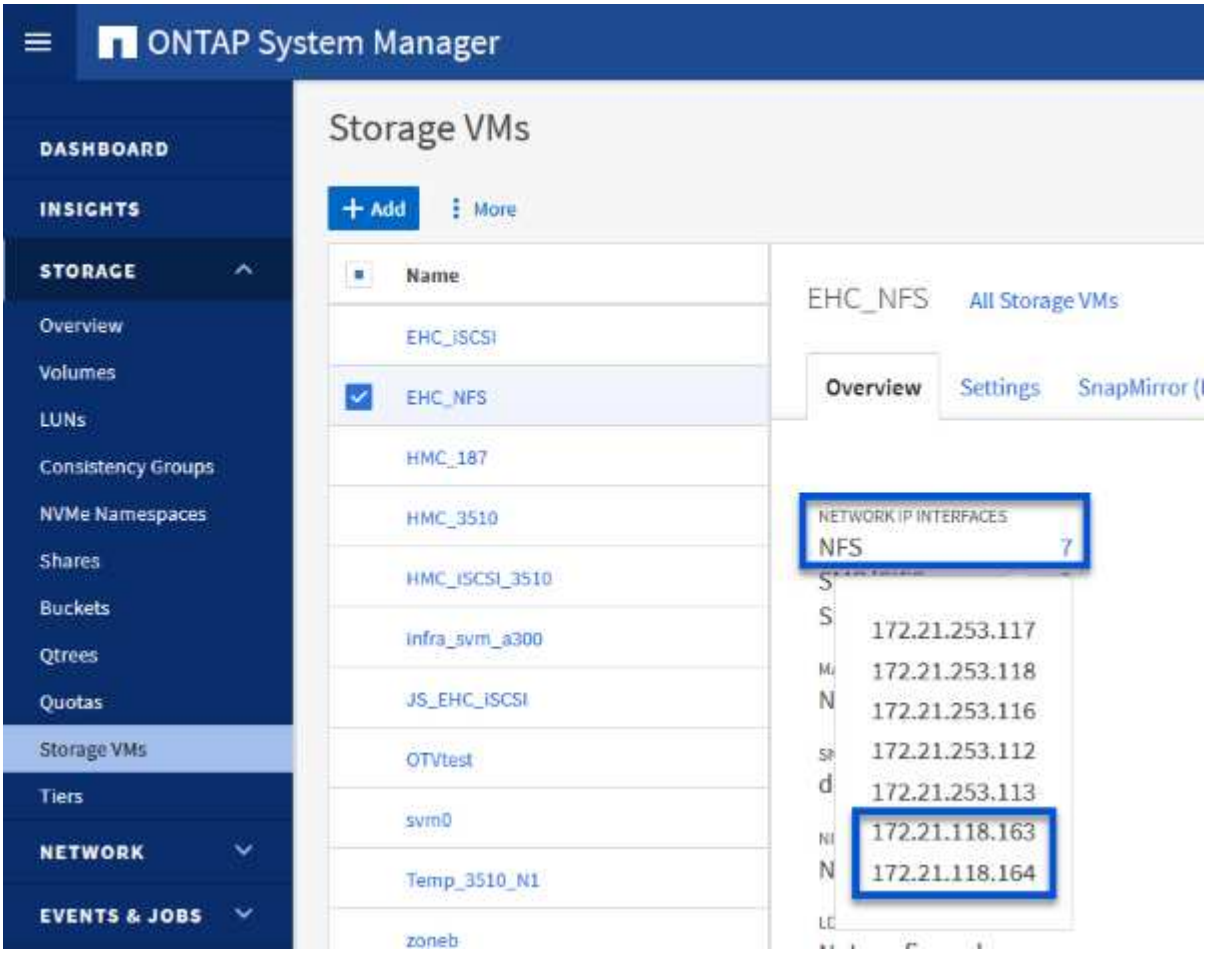

Alternately, verify the LIFs associated with an SVM from the ONTAP CLI with the following command:

```
network interface show -vserver <SVM_NAME>
```

1. Verify that the ESXi hosts can communicate to the ONTAP NFS Server. Log into the ESXi host via SSH and ping the SVM LIF:

```
vmkping <IP Address>
```

When deploying a VI Workload Domain, VCF validates connectivity to the NFS Server. This is done using the management adapter on the ESXi hosts before any additional vmkernel adapter is added with the NFS IP address. Therefore, it is necessary to ensure that either 1) the management network is routable to the NFS Server, or 2) a LIF for the management network has been added to the SVM hosting the NFS datastore volume, to ensure that the validation can proceed.

Create Export Policy for sharing NFS volume

Create an export policy in ONTAP System Manager to define access control for NFS volumes.

1. In ONTAP System Manager click on **Storage VMs** in the left-hand menu and select an SVM from the list.
2. On the **Settings** tab locate **Export Policies** and click on the arrow to access.

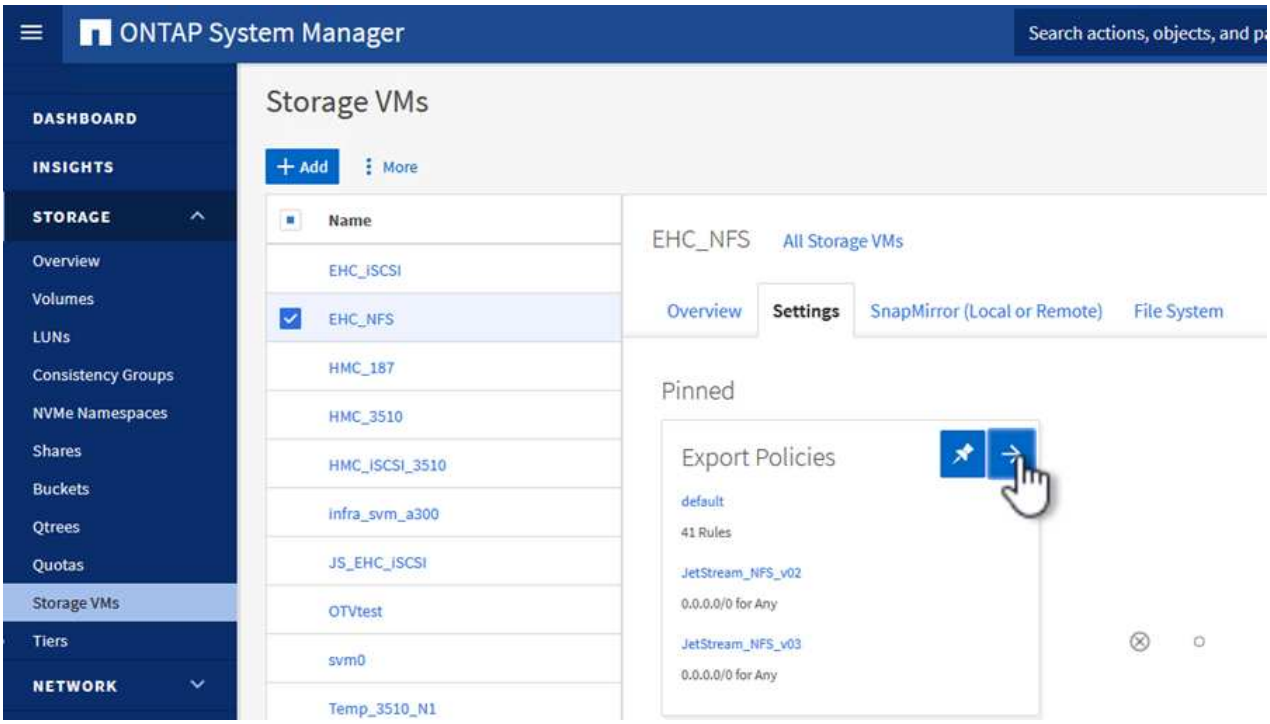

3. In the **New export policy** window add a name for the policy, click on the **Add new rules** button and then on the **+Add** button to begin adding a new rule.

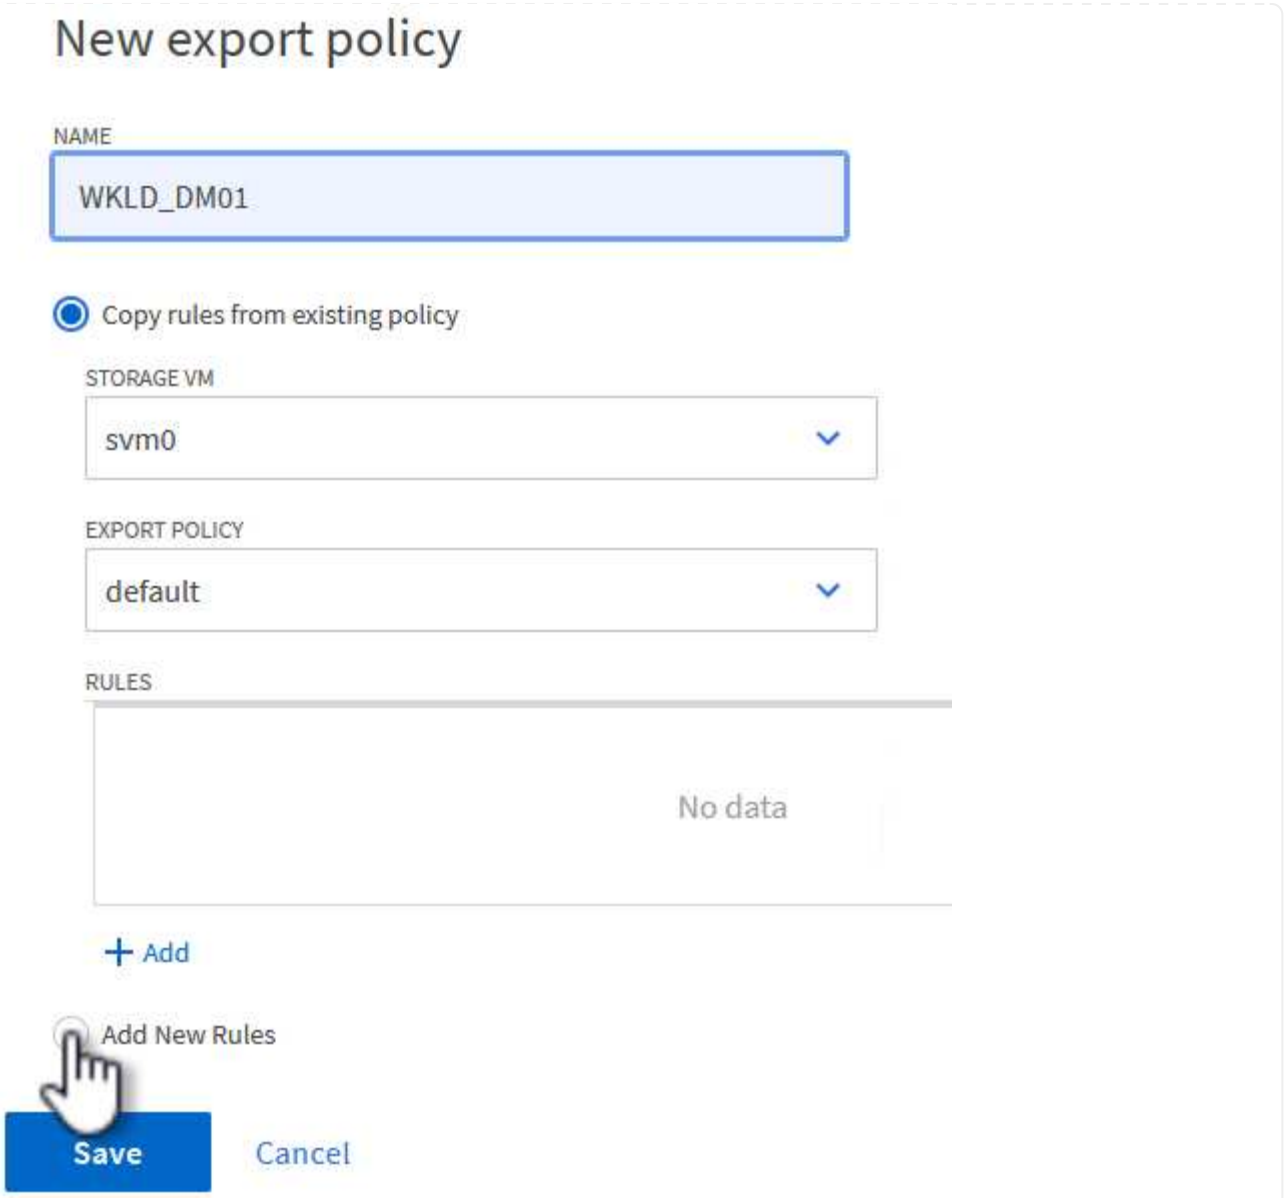

4. Fill in the IP Addresses, IP address range, or network that you wish to include in the rule. Uncheck the **SMB/Cifs** and **FlexCache** boxes and make selections for the access details below. Selecting the UNIX boxes is sufficient for ESXi host access.

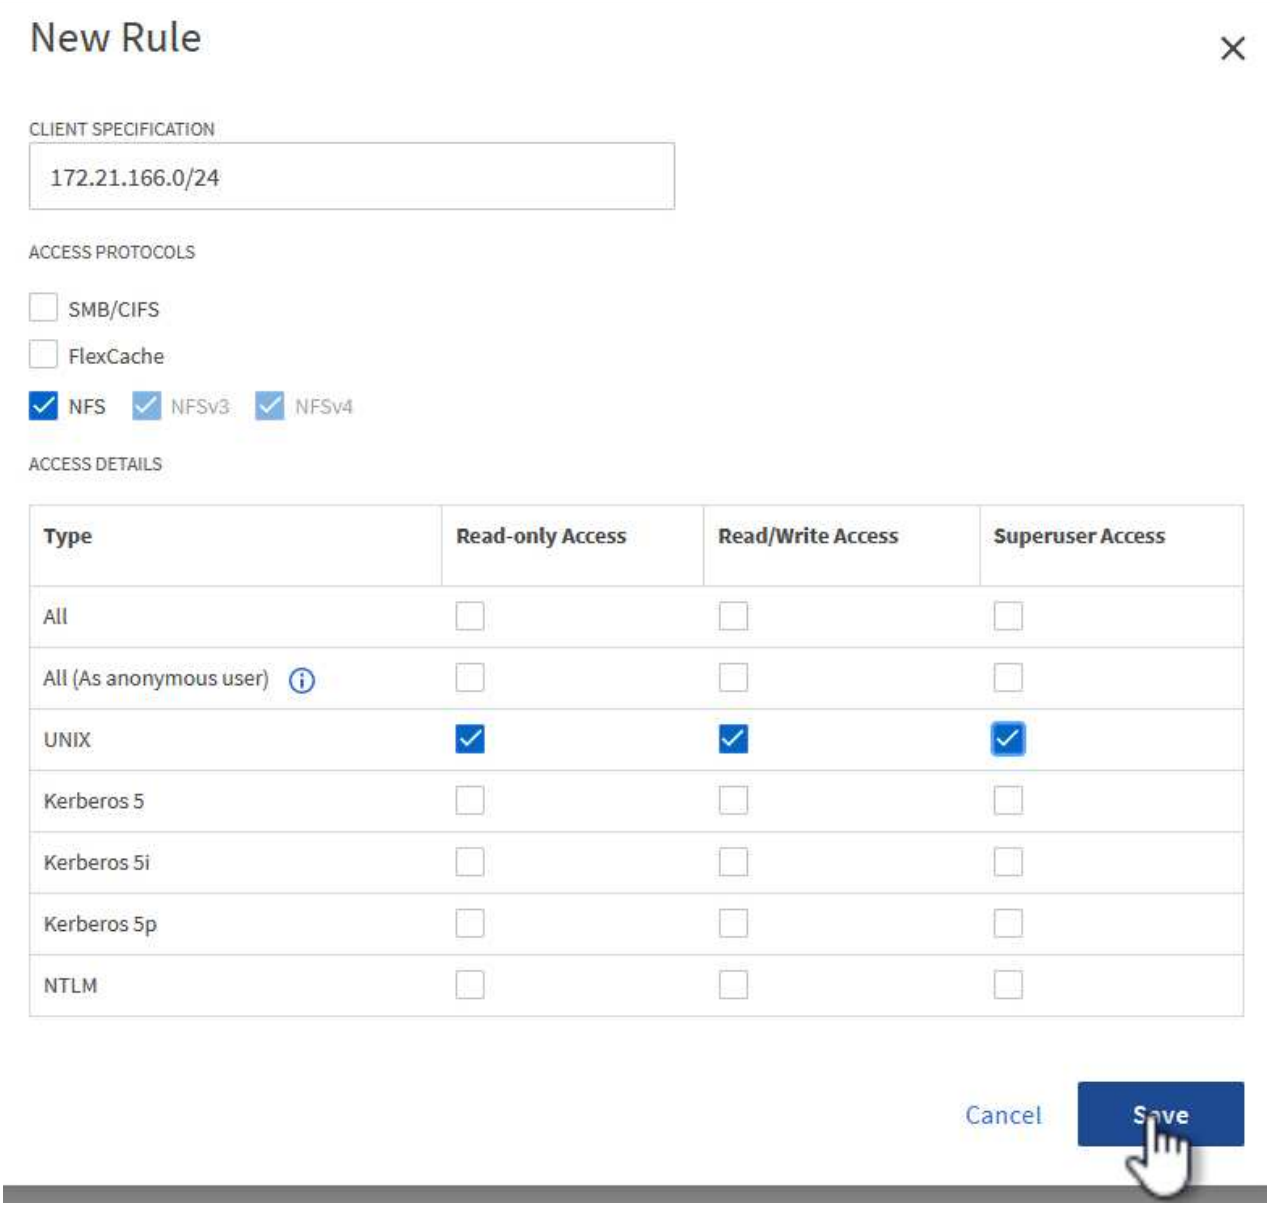

When deploying a VI Workload Domain, VCF validates connectivity to the NFS Server. This is done using the management adapter on the ESXi hosts before any additional vmkernel adapter is added with the NFS IP address. Therefore, it is necessary to ensure that the export policy includes the VCF management network in order to allow the validation to proceed.

5. Once all rules have been entered click on the **Save** button to save the new Export Policy.
6. Alternately, you can create export policies and rules in the ONTAP CLI. Refer to the steps for creating an export policy and adding rules in the ONTAP documentation.
   - Use the ONTAP CLI to Create an export policy.
   - Use the ONTAP CLI to Add a rule to an export policy.

Create NFS volume

Create an NFS volume on the ONTAP storage system to be used as a datastore in the Workload Domain deployment.

1. From ONTAP System Manager navigate to **Storage > Volumes** in the left-hand menu and click on **+Add** to create a new volume.

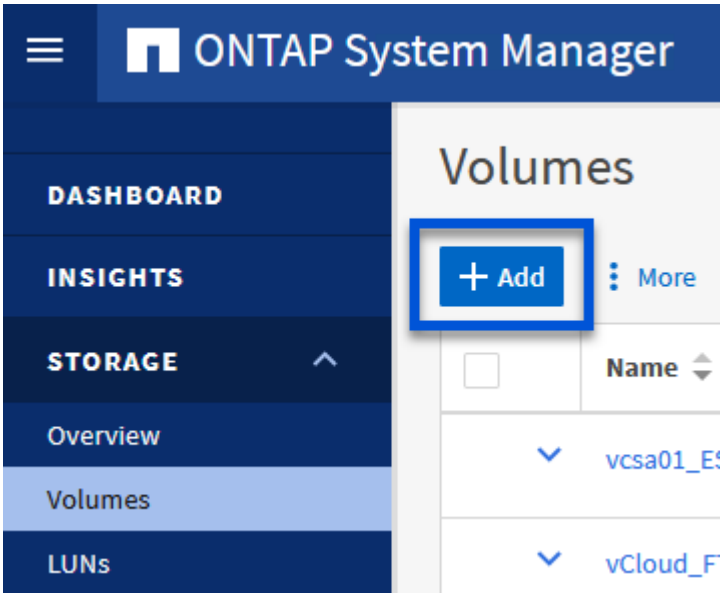

2. Add a name for the volume, fill out the desired capacity and selection the storage VM that will host the volume. Click on **More Options** to continue.

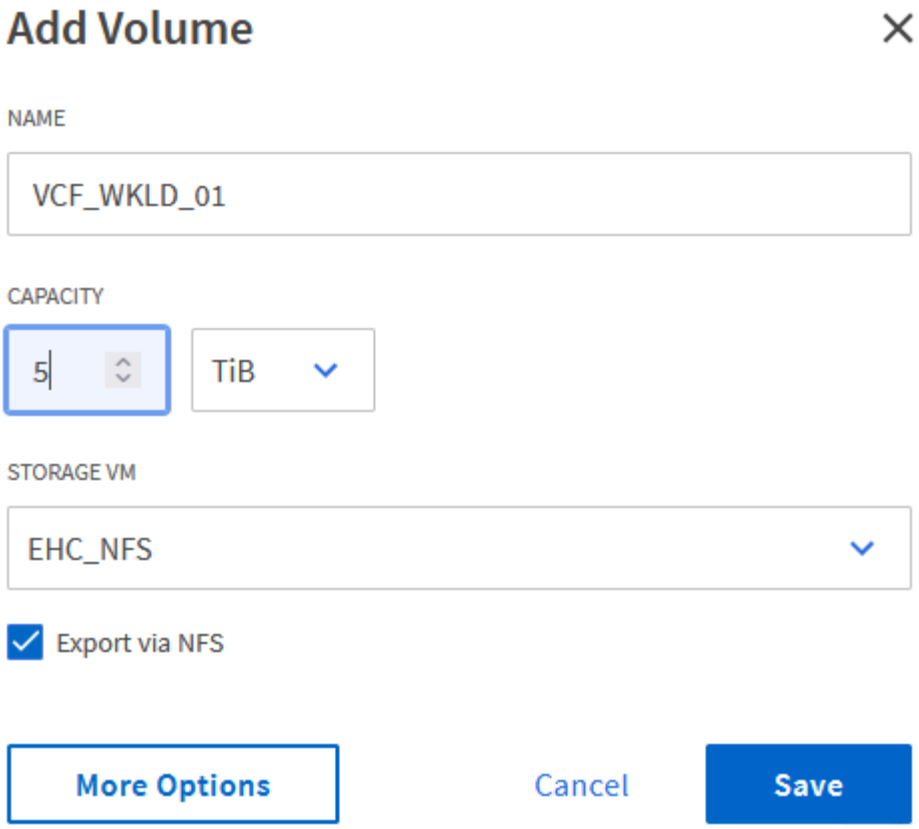

3. Under Access Permissions, select the Export Policy which includes the VCF management network or IP address and NFS network IP addresses that will be used for both validation of the NFS Server and NFS traffic.

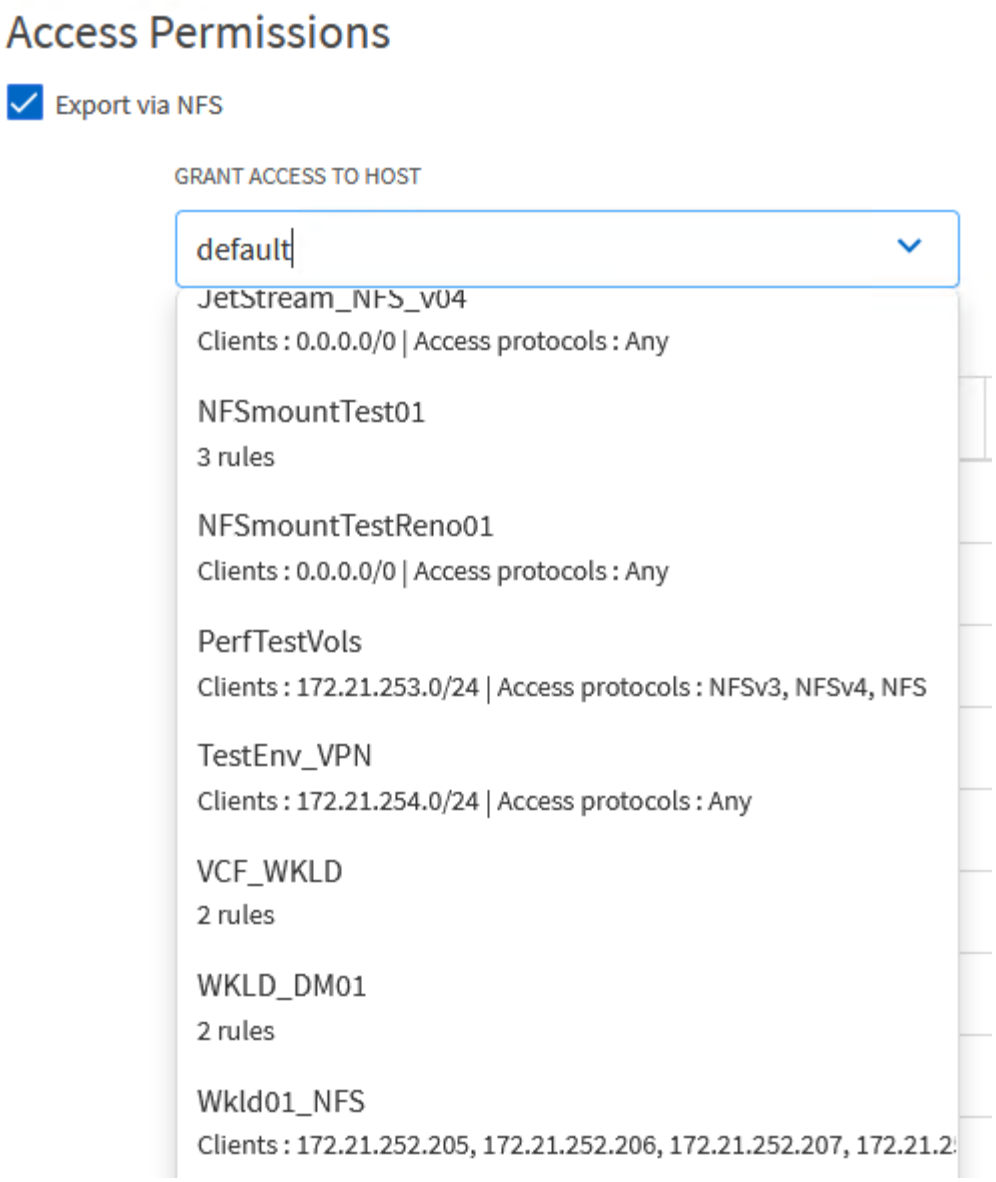

+

When deploying a VI Workload Domain, VCF validates connectivity to the NFS Server. This is done using the management adapter on the ESXi hosts before any additional vmkernel adapter is added with the NFS IP address. Therefore, it is necessary to ensure that either 1) the management network is routable to the NFS Server, or 2) a LIF for the management network has been added to the SVM hosting the NFS datastore volume, to ensure that the validation can proceed.

4. Alternately, ONTAP Volumes can be created in the ONTAP CLI. For more information refer to the lun create command in the ONTAP commands documentation.

Create Network Pool in SDDC Manager

ANetwork Pool must be created in SDDC Manager before commissioning the ESXi hosts, as preparation for deploying them in a VI Workload Domain. The Network Pool must include the network information and IP address range(s) for VMkernel adapters to be used for communication with the NFS server.

1. From the SDDC Manager web interface navigate to **Network Settings** in the left-hand menu and click on the **+ Create Network Pool** button.

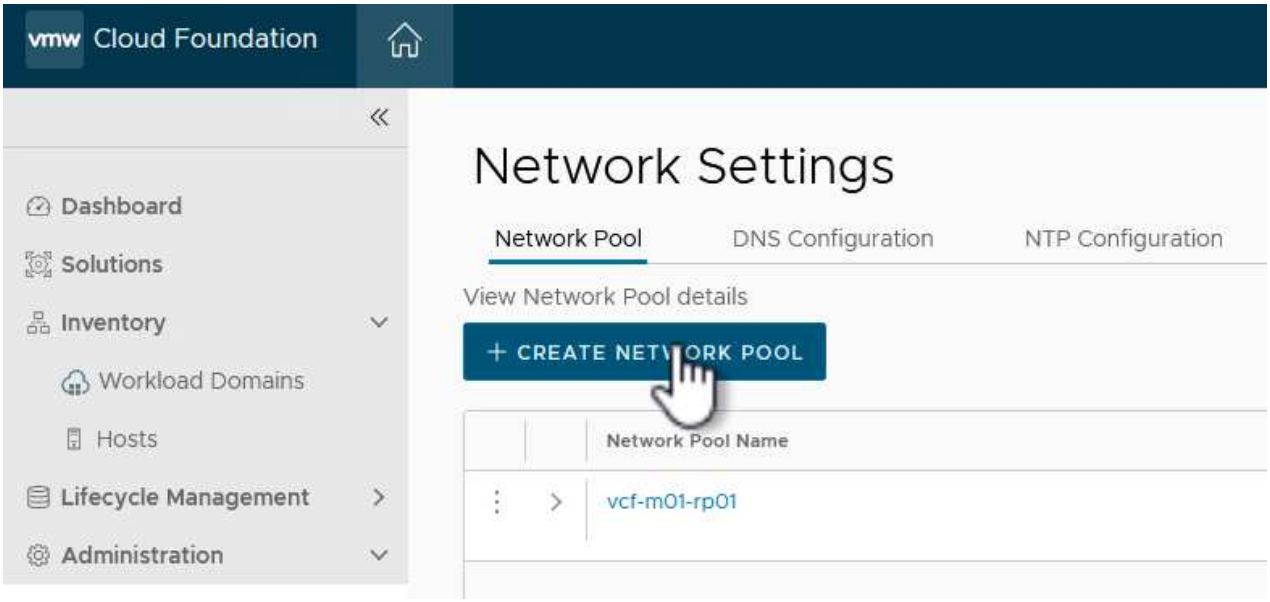

2. Fill out a name for the Network Pool, select the check box for NFS and fill out all networking details. Repeat this for the vMotion network information.

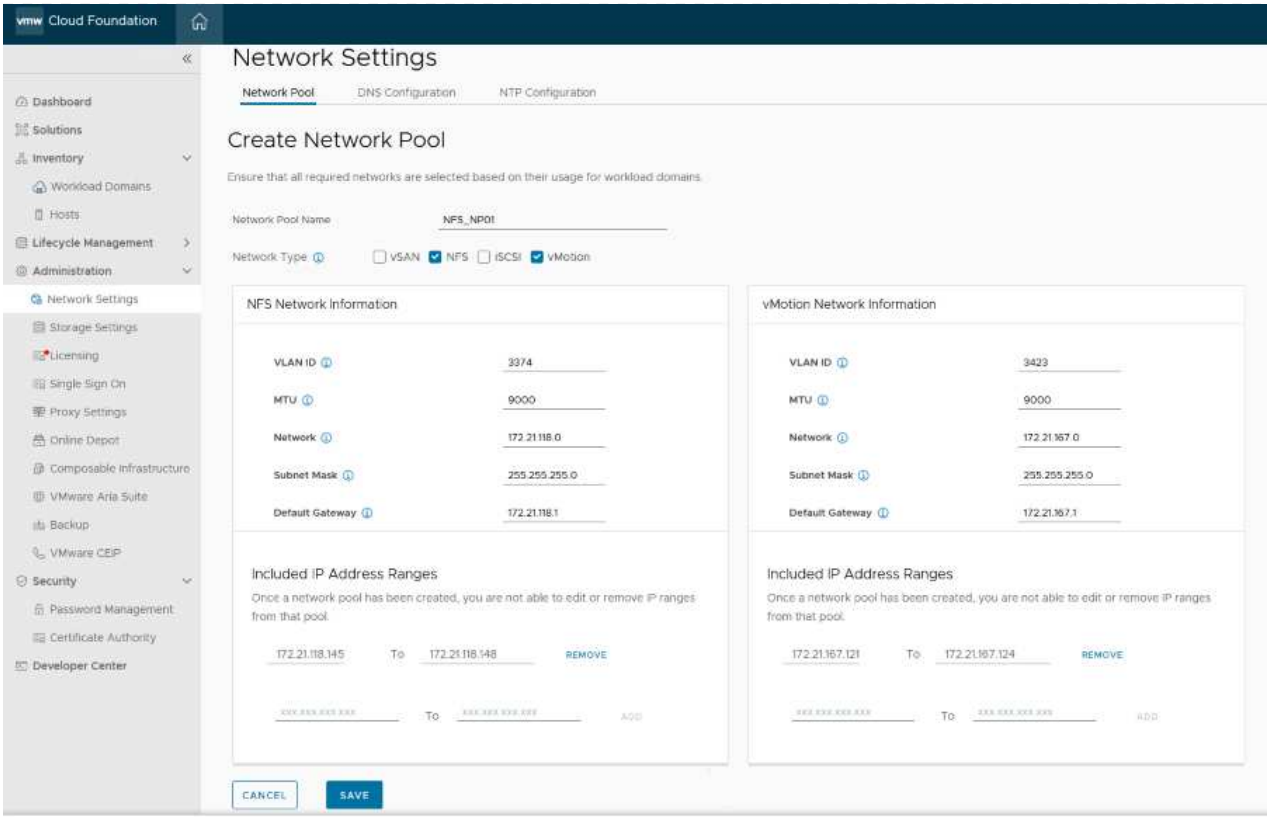

3. Click the **Save** button to complete creating the Network Pool.

Commission Hosts

Before ESXi hosts can be deployed as a workload domain they must be added to the SDDC Manager inventory. This involves providing the required information, passing validation and starting the commissioning process.

For more information see Commission Hosts in the VCF Administration Guide.

1. From the SDDC Manager interface navigate to **Hosts** in the left-hand menu and click on the **Commission Hosts** button.

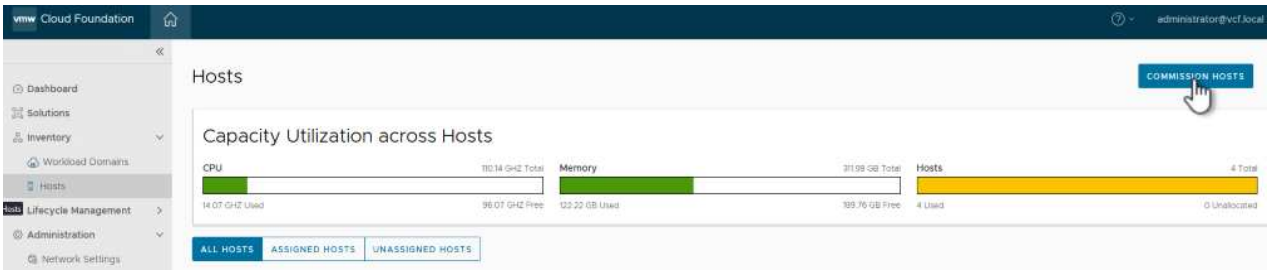

2. The first page is a prerequisite checklist. Double-check all prerequisites and select all checkboxes to proceed.

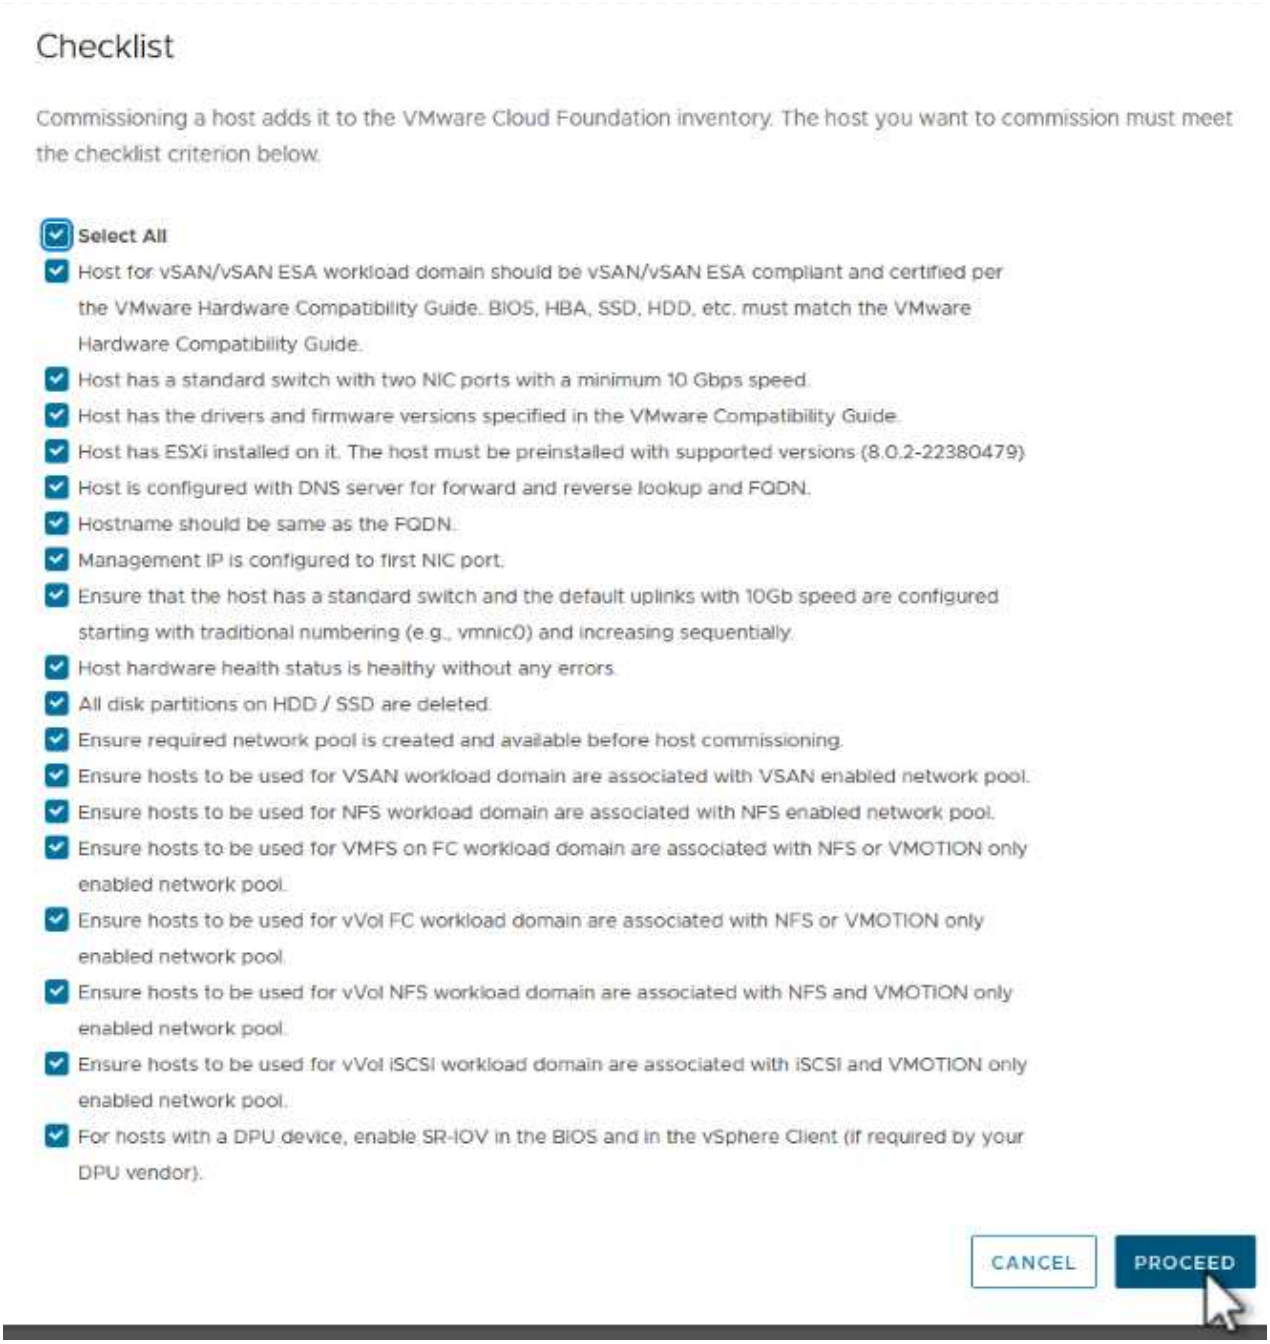

3. In the **Host Addition and Validation** window fill out the **Host FQDN**, **Storage Type**, The **Network Pool** name that includes the vMotion and NFS storage IP addresses to be used for the workload domain, and the credentials to access the ESXi host. Click on **Add** to add the host to the group of hosts to be validated.

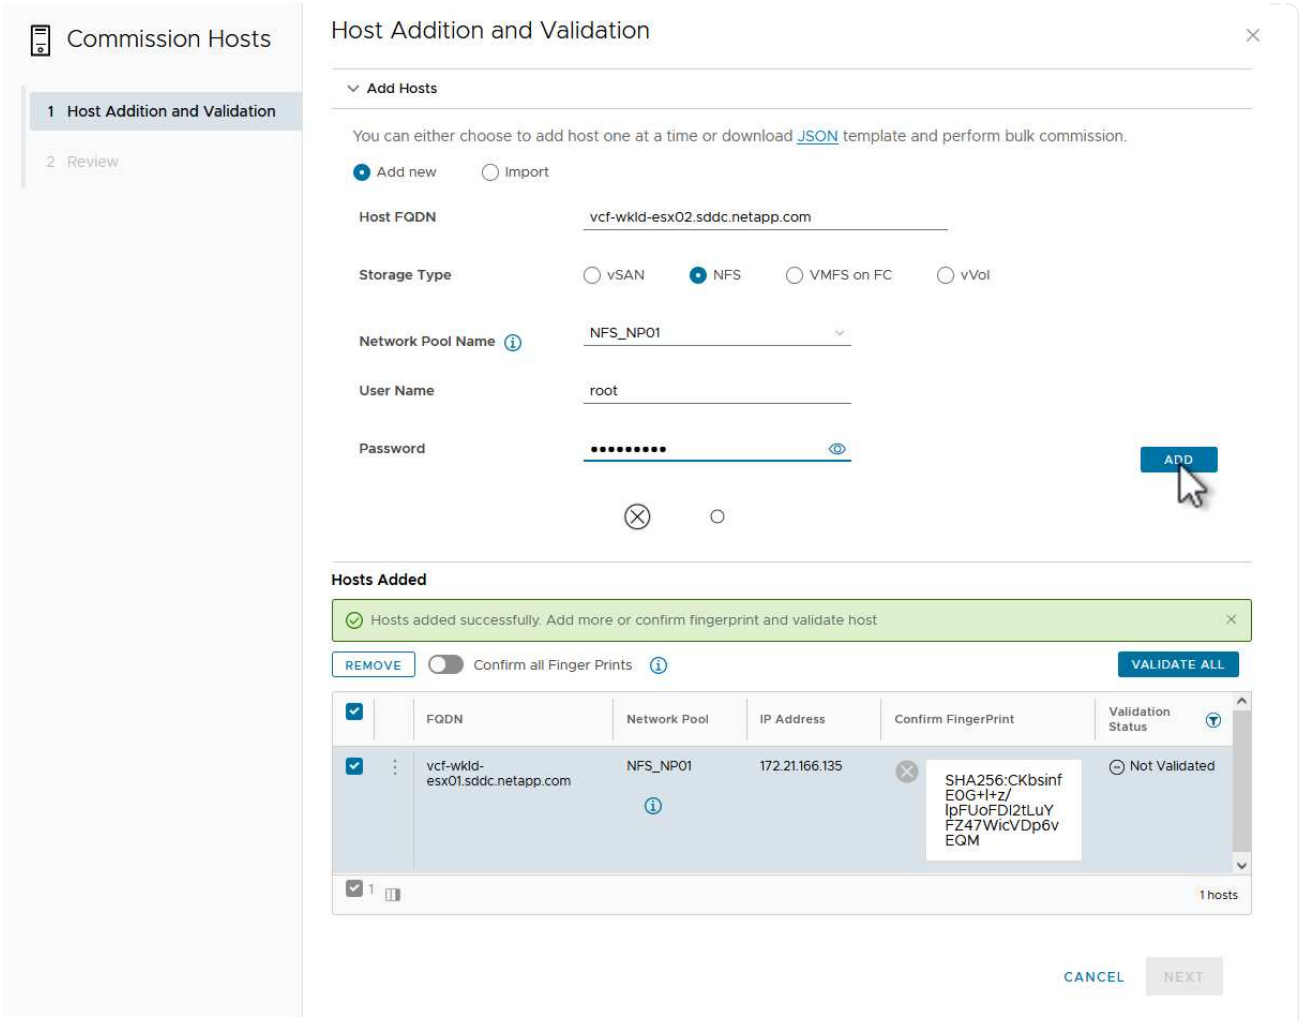

4. Once all hosts to be validated have been added, click on the **Validate All** button to continue.
5. Assuming all hosts are validated, click on **Next** to continue.

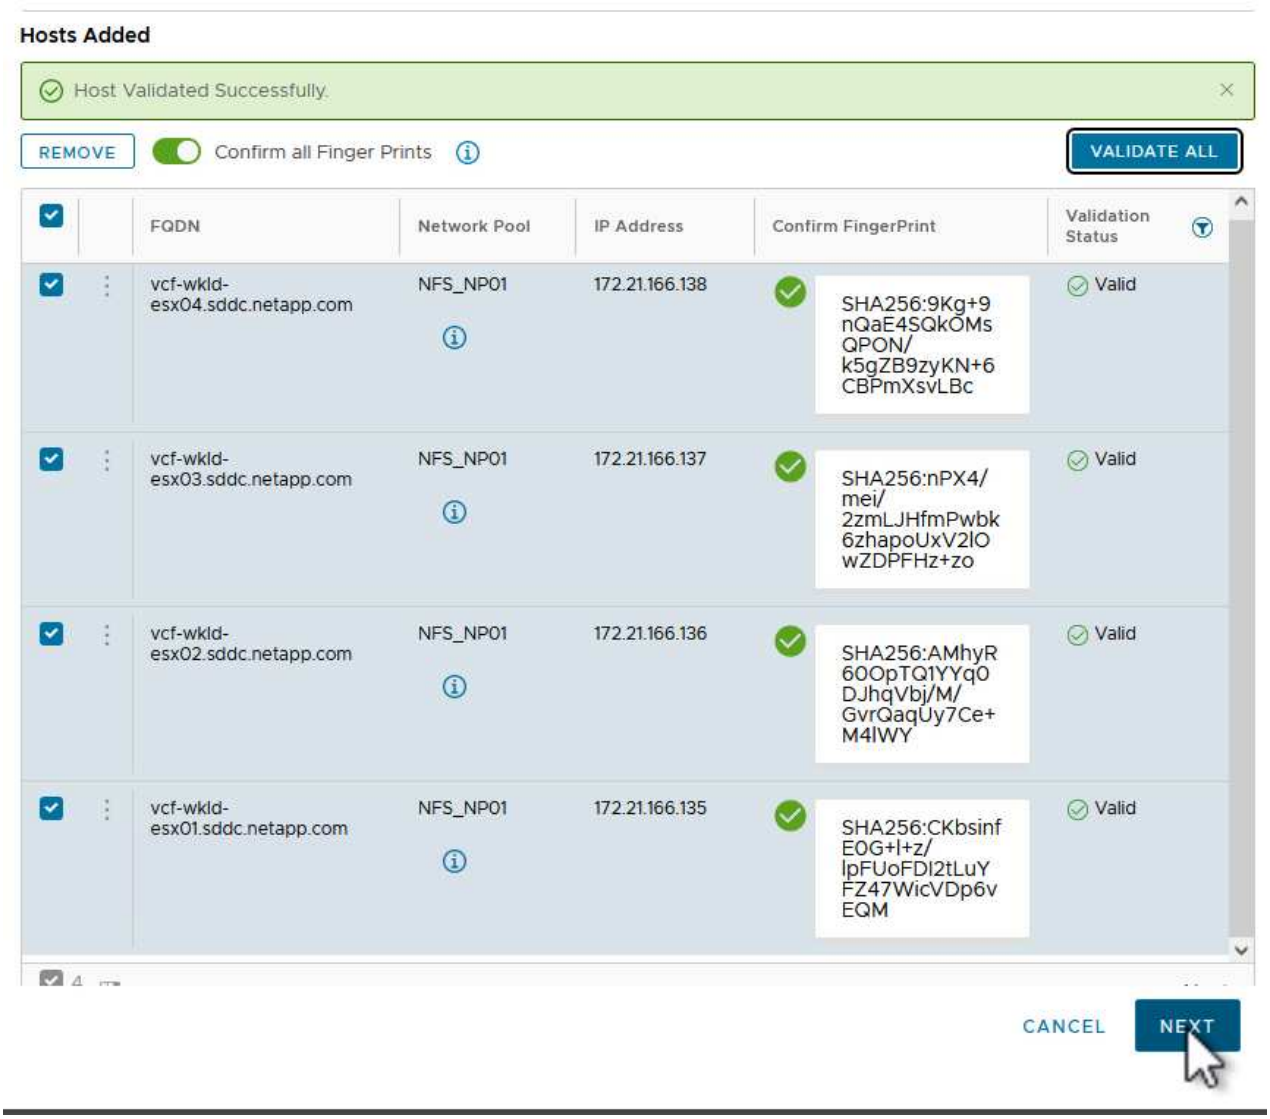

6. Review the list of hosts to be commissioned and click on the **Commission** button to start the process. Monitor the commissioning process from the Task pane in SDDC manager.

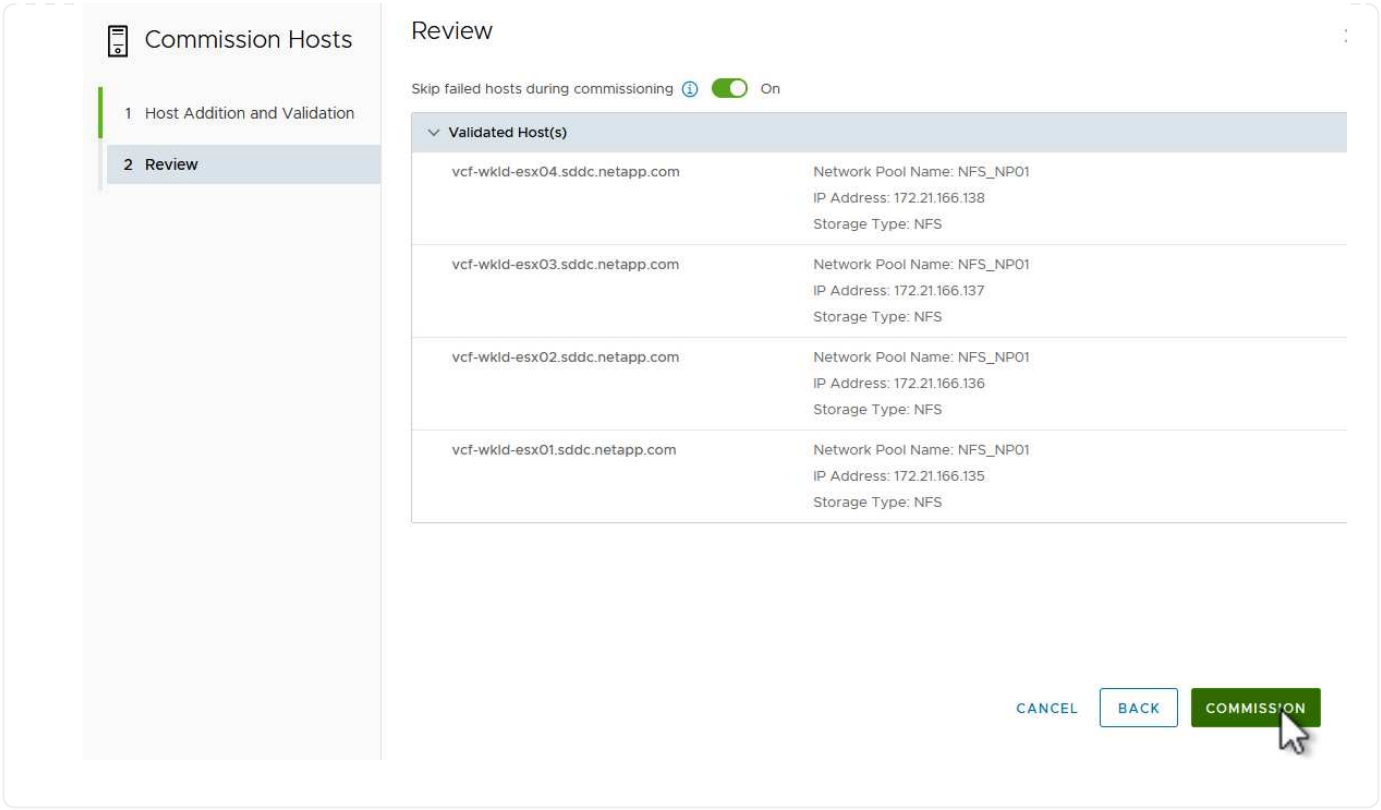
Commission Hosts
1 Host Addition and Validation
2 Review
Review
Skip failed hosts during commissioning
On
Validated Host(s)
vcf-wkld-esx04.sddc.netapp.com
Network Pool Name: NFS_NP01
IP Address: 172.21.166.138
Storage Type: NFS
vcf-wkld-esx03.sddc.netapp.com
Network Pool Name: NFS_NP01
IP Address: 172.21.166.137
Storage Type: NFS
vcf-wkld-esx02.sddc.netapp.com
Network Pool Name: NFS_NP01
IP Address: 172.21.166.136
Storage Type: NFS
vcf-wkld-esx01.sddc.netapp.com
Network Pool Name: NFS_NP01
IP Address: 172.21.166.135
Storage Type: NFS
CANCEL
BACK
COMMISSION

Deploy VI Workload Domain

Deploying VI workload domains is accomplished using the VCF Cloud Manager interface. Only the steps related to the storage configuration will be presented here.

For step-by-step instructions on deploying a VI workload domain refer to Deploy a VI Workload Domain Using the SDDC Manager UI.

1. From the SDDC Manager Dashboard click on **+ Workload Domain** in the upper right hand corner to create a new Workload Domain.

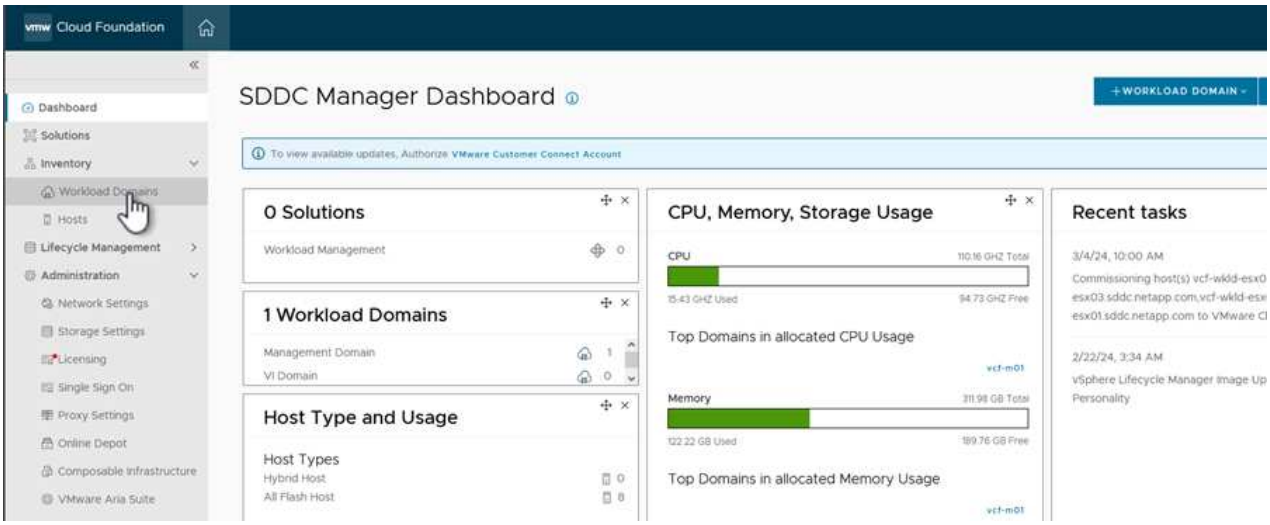

2. In the VI Configuration wizard fill out the sections for **General Info, Cluster, Compute, Networking**, and **Host Selection** as required.

For information on filling out the information required in the VI Configuration wizard refer to Deploy a VI Workload Domain Using the SDDC Manager UI.

+
image::vmware-vcf-aff-image13.png[VI Configuration Wizard]

1. In the NFS Storage section fill out the Datastore Name, the folder mount point of the NFS volume and the IP address of the ONTAP NFS storage VM LIF.

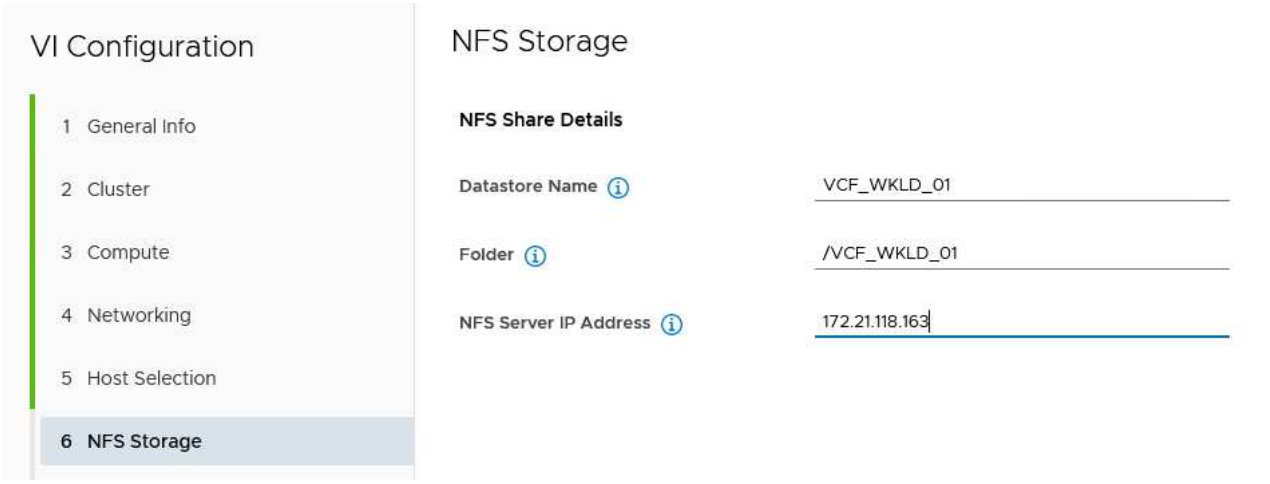

2. In the VI Configuration wizard complete the Switch Configuration and License steps, and then click on **Finish** to start the Workload Domain creation process.

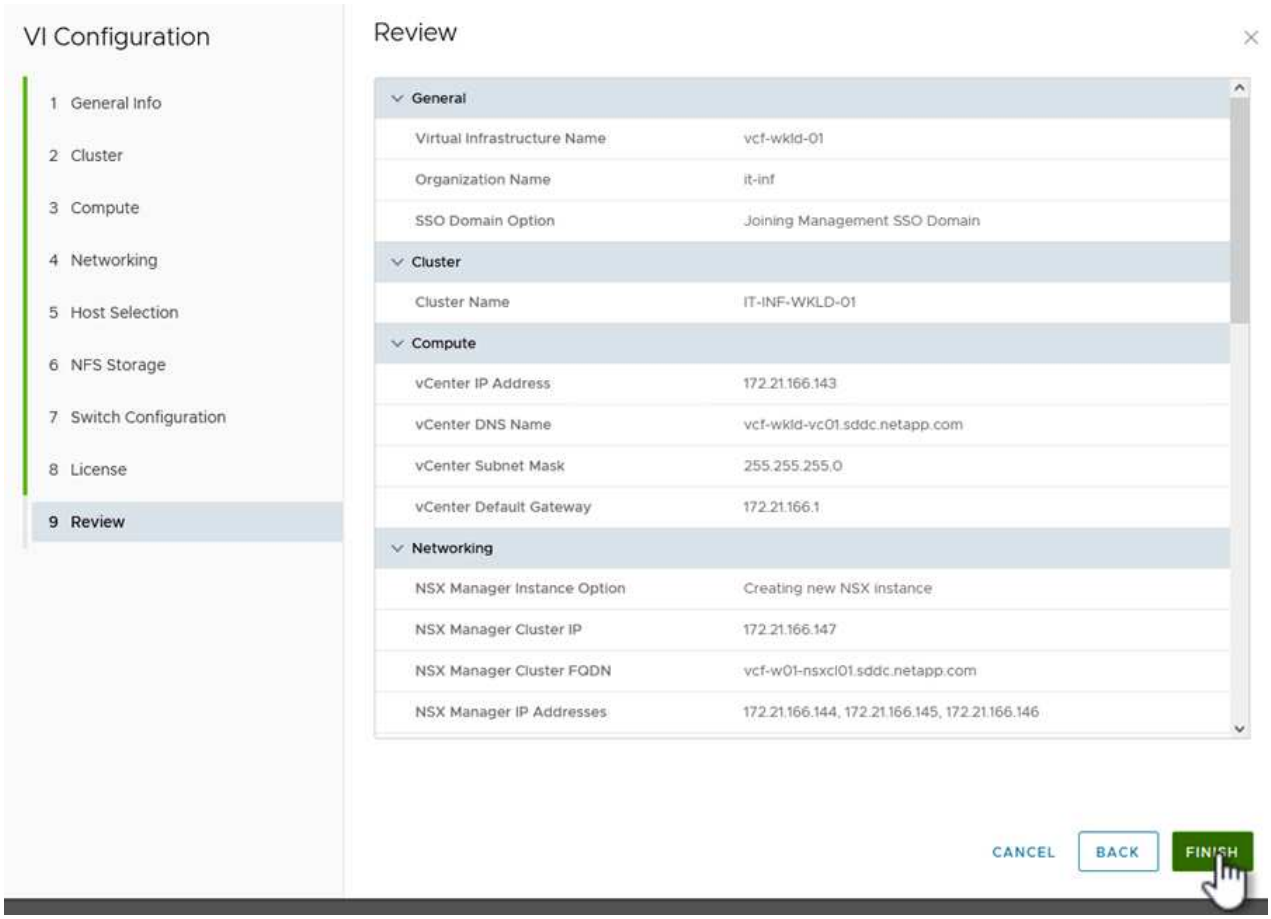

3. Monitor the process and resolve any validation issues that arise during the process.

Install NetApp NFS Plug-in for VMware VAAI

The NetApp NFS Plug-in for VMware VAAI integrates the VMware Virtual Disk Libraries installed on the ESXi host and provides higher performance cloning operations that finish faster. This is a recommended procedure when using ONTAP storage systems with VMware vSphere.

For step-by-step instructions on deploying the NetApp NFS Plug-in for VMware VAAI following the instructions at Install NetApp NFS Plug-in for VMware VAAI.

Video demo for this solution

NFS Datastores as Principal Storage for VCF Workload Domains

Migration of VMs

Migrate VMs to ONTAP Datastores

Author: Suresh Thoppay

VMware vSphere by Broadcom supports VMFS, NFS, and vVol datastores for hosting virtual machines. Customers have the option to create those datastores with hyper converged infrastructures or with centralized shared storage systems. Customers often see the value with hosting on ONTAP based storage systems to provide space efficient snapshots and clones of Virtual machines, flexiblity to choose various deployment models across the datacenters and clouds, operational efficiency with monitoring and alerting tools, security, governance and optional compliance tools to inspect VM data, etc,.

VMs hosted on ONTAP datastores can be protected using SnapCenter Plugin for VMware vSphere (SCV). SCV creates storage based snapshots and also replicates to remote ONTAP storage system. Restores can be performed either from Primary or Secondary storage systems.

Customers has flexibility to choose Cloud Insights or Aria Operations or combination of both or other third party tools that use ONTAP api to troubleshoot, performance monitoring, reporting and alert notification features.

Customers can easily provision datastore using ONTAP Tools vCenter Plug-in or its API and VMs can be migrated to ONTAP datastores even while it is powered on.

Some VMs which are deployed with external management tool like Aria Automation, Tanzu (or other Kubernetes flavors) are usually depends on VM storage policy. If migrating between the datastores within same VM storage policy, it should be of less impact for the applications. Check with Application owners to properly migrate those VMs to new datastore. vSphere 8 introduced vMotion notification to prepare application for the vMotion.

Network Requirements

VM migration with vMotion

It is assumed that dual storage network is already in place for the ONTAP datastore to provide connectivity, fault tolerance and performance boost.

Migration of VMs across the vSphere hosts are also handled by the VMKernel interface of the vSphere host. For hot migration (powered on VMs), VMKernel interface with vMotion enabled service is used and for cold migration (powered off VMs), VMKernel interface with Provisioning service enabled is consumed to move the data. If no valid interface was found, it will use the management interface to move the data which may not be desirable for certain use cases.

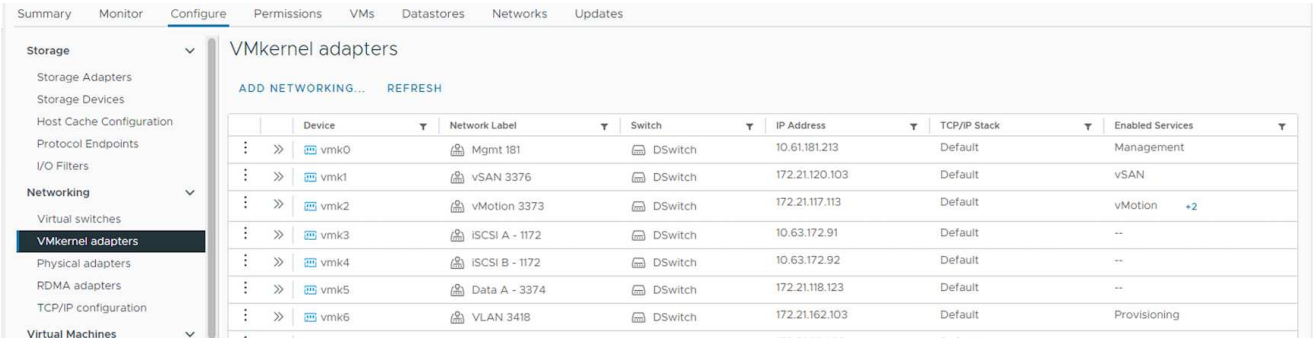

When you edit the VMKernel interface, here is the option to enable the required services.

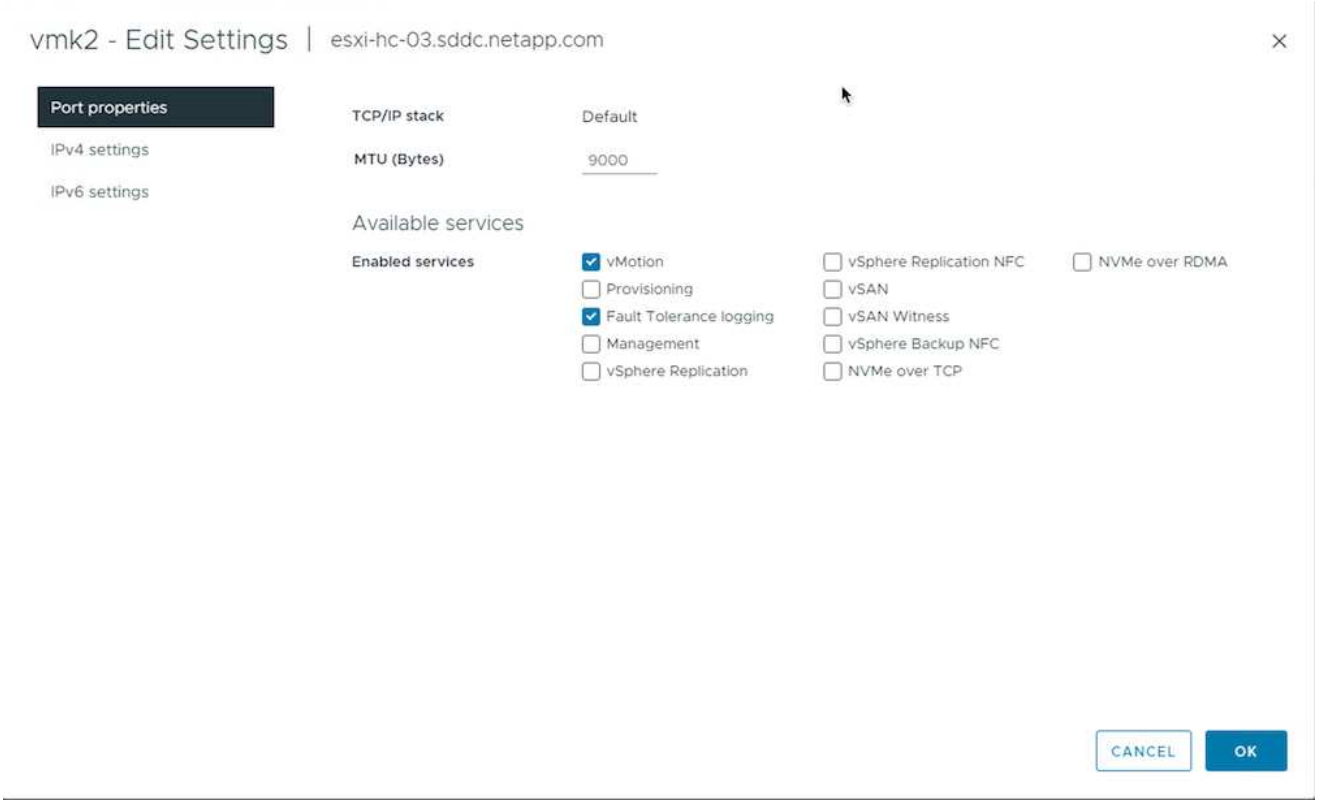

Ensure at least two high-speed active uplink nics are available for the portgroup used by vMotion and Provisioning VMkernel interfaces.

VM Migration Scenarios

vMotion is often used to migrate the VMs irrespective of its power state. Additional considerations and migration procedure for specific scenarios is available below.

Understand VM Conditions and Limitation of vSphere vMotion before proceeding with any VM migration options.

Migration of VMs from specific vSphere Datastore

Follow the procedure below to migrate VMs to new Datastore using UI.

1. With vSphere Web Client, select the Datastore from the storage inventory and click on VMs tab.

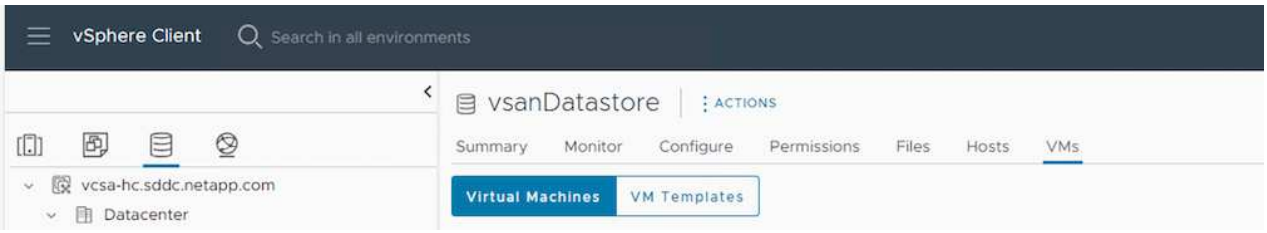

2. Select the VMs that needs to be migrated and right click to select Migrate option.

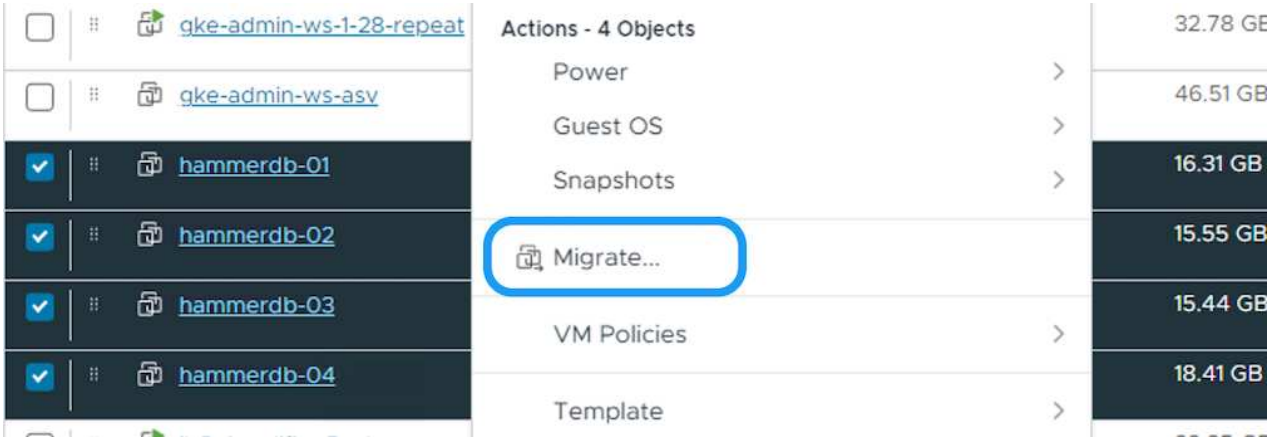

3. Choose option to change storage only, Click Next

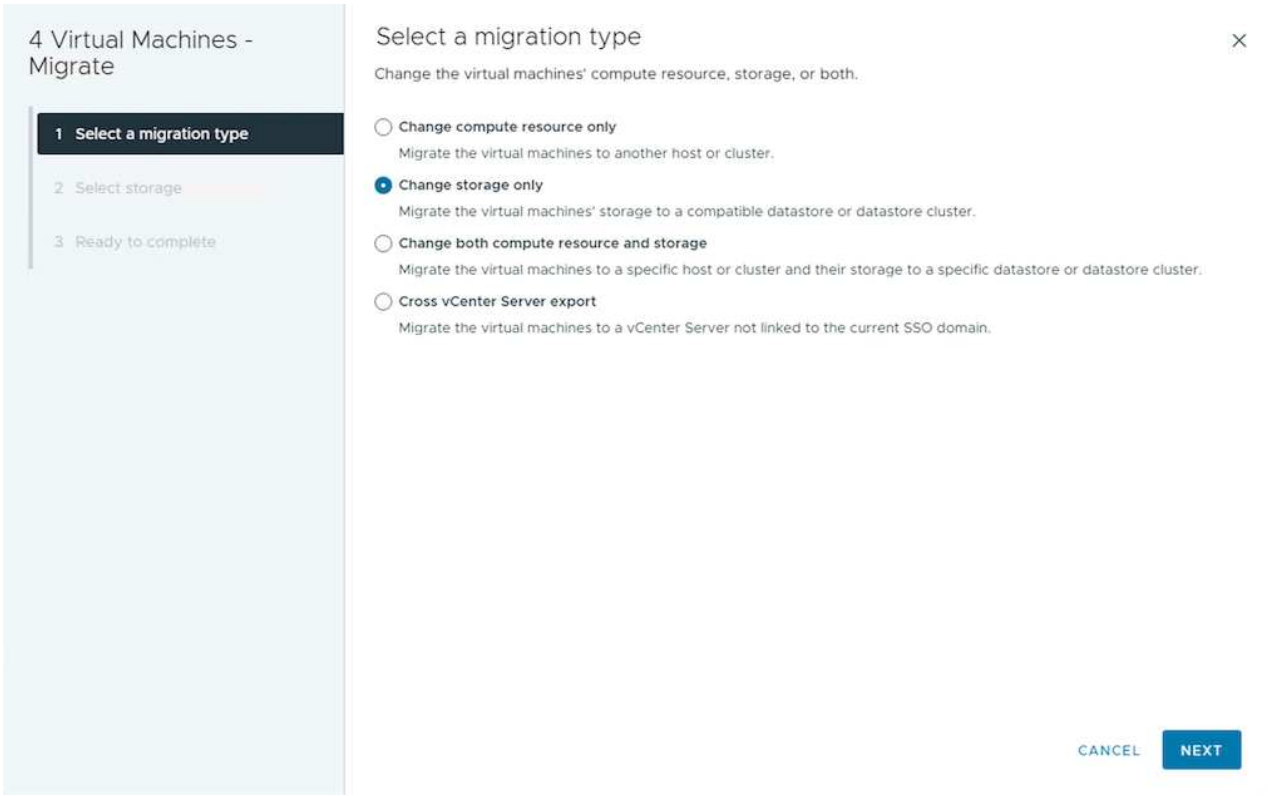

4. Select the desired VM Storage Policy and pick the datastore that is compatible. Click Next.

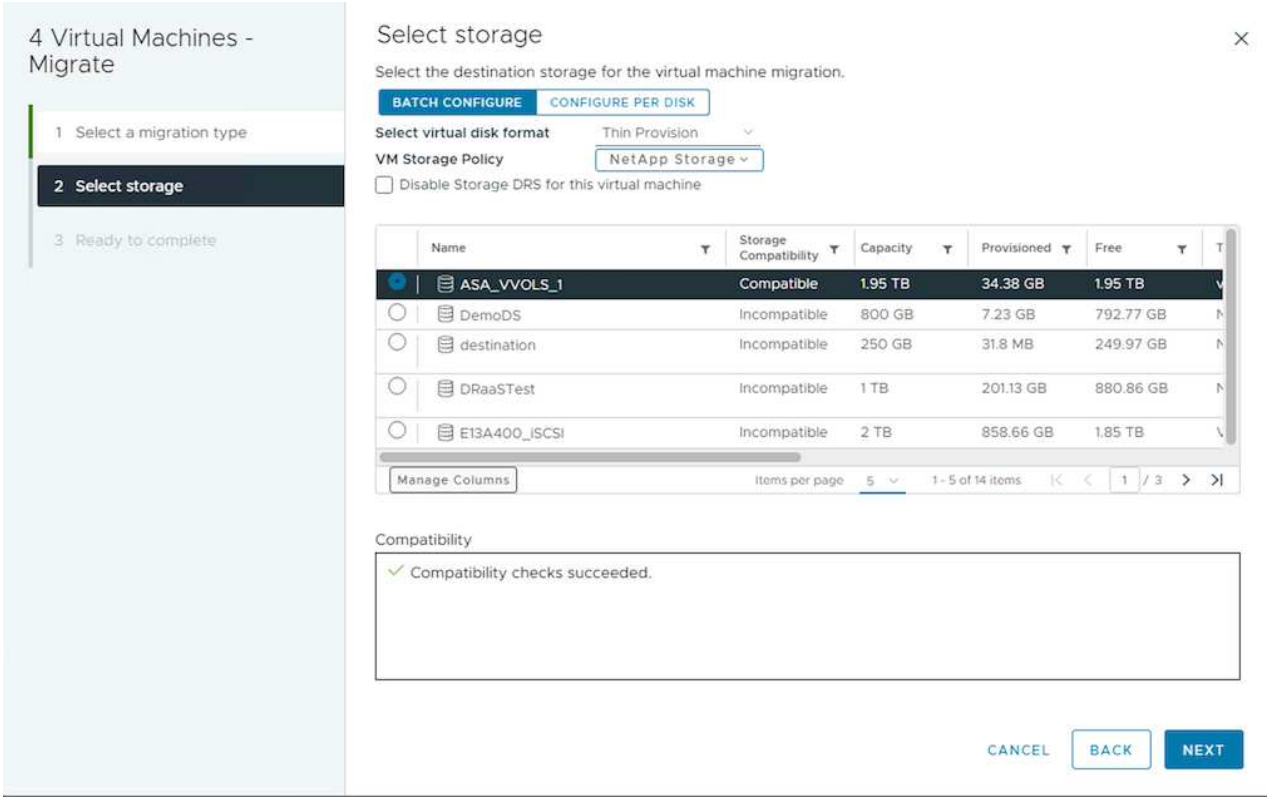

5. Review and click on Finish.

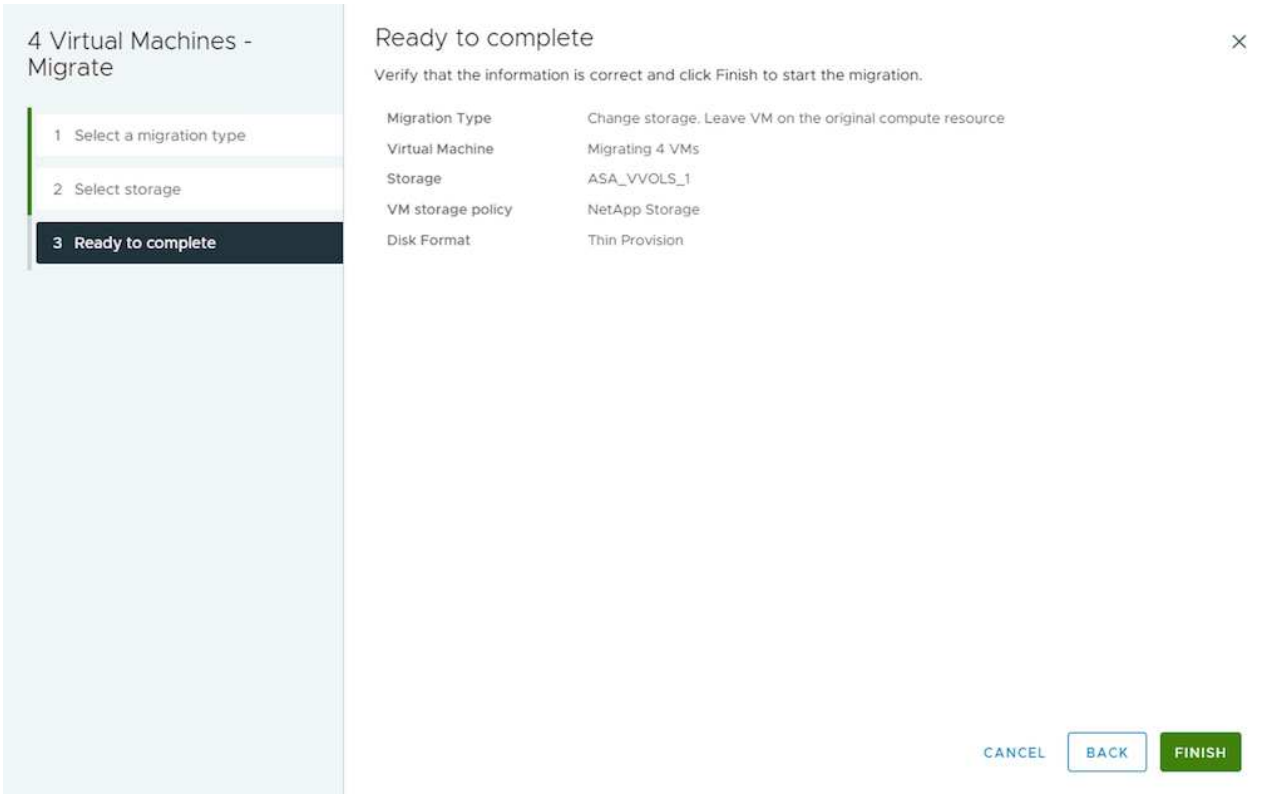

To migrate VMs using PowerCLI, here is the sample script.

```
#Authenticate to vCenter
Connect-VIServer -server vcsa.sddc.netapp.local -force

# Get all VMs with filter applied for a specific datastore
$vm = Get-DataStore 'vSanDatastore' | Get-VM Har*

#Gather VM Disk info
$vmdisk = $vm | Get-HardDisk

#Gather the desired Storage Policy to set for the VMs. Policy should be
available with valid datastores.
$storagepolicy = Get-SPBMStoragePolicy 'NetApp Storage'

#set VM Storage Policy for VM config and its data disks.
$vm, $vmdisk | Get-SPBMEntityConfiguration | Set-
SPBMEntityConfiguration -StoragePolicy $storagepolicy

#Migrate VMs to Datastore specified by Policy
$vm | Move-VM -Datastore (Get-SPBMCompatibleStorage -StoragePolicy
$storagepolicy)

#Ensure VM Storage Policy remains compliant.
$vm, $vmdisk | Get-SPBMEntityConfiguration
```

Migration of VMs in same vSphere cluster

Follow the procedure below to migrate VMs to new Datastore using UI.

1. With vSphere Web Client, select the Cluster from the Host and Cluster inventory and click on VMs tab.

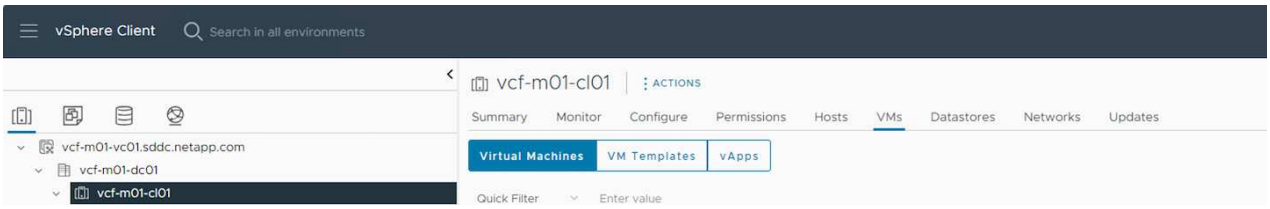

2. Select the VMs that needs to be migrated and right click to select Migrate option.

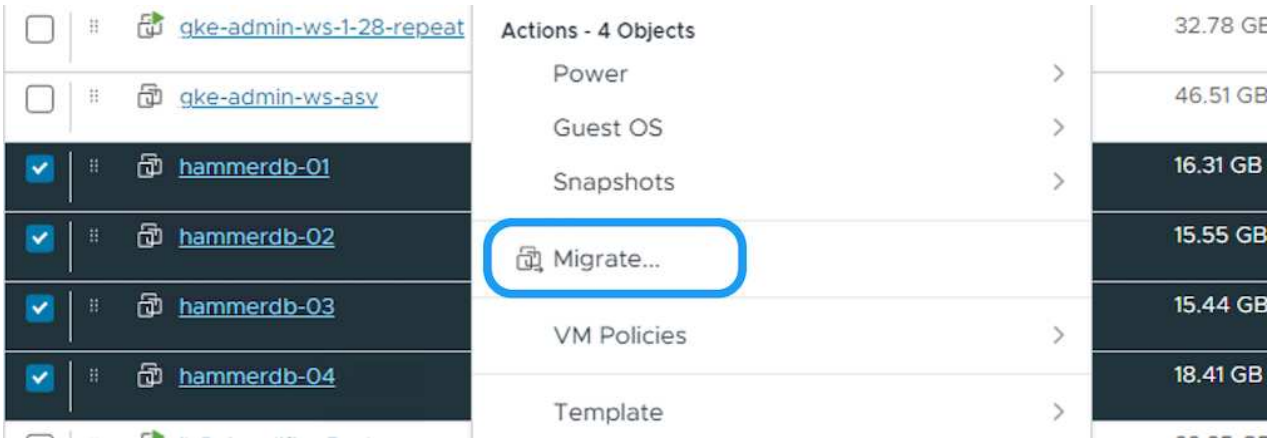

3. Choose option to change storage only, Click Next

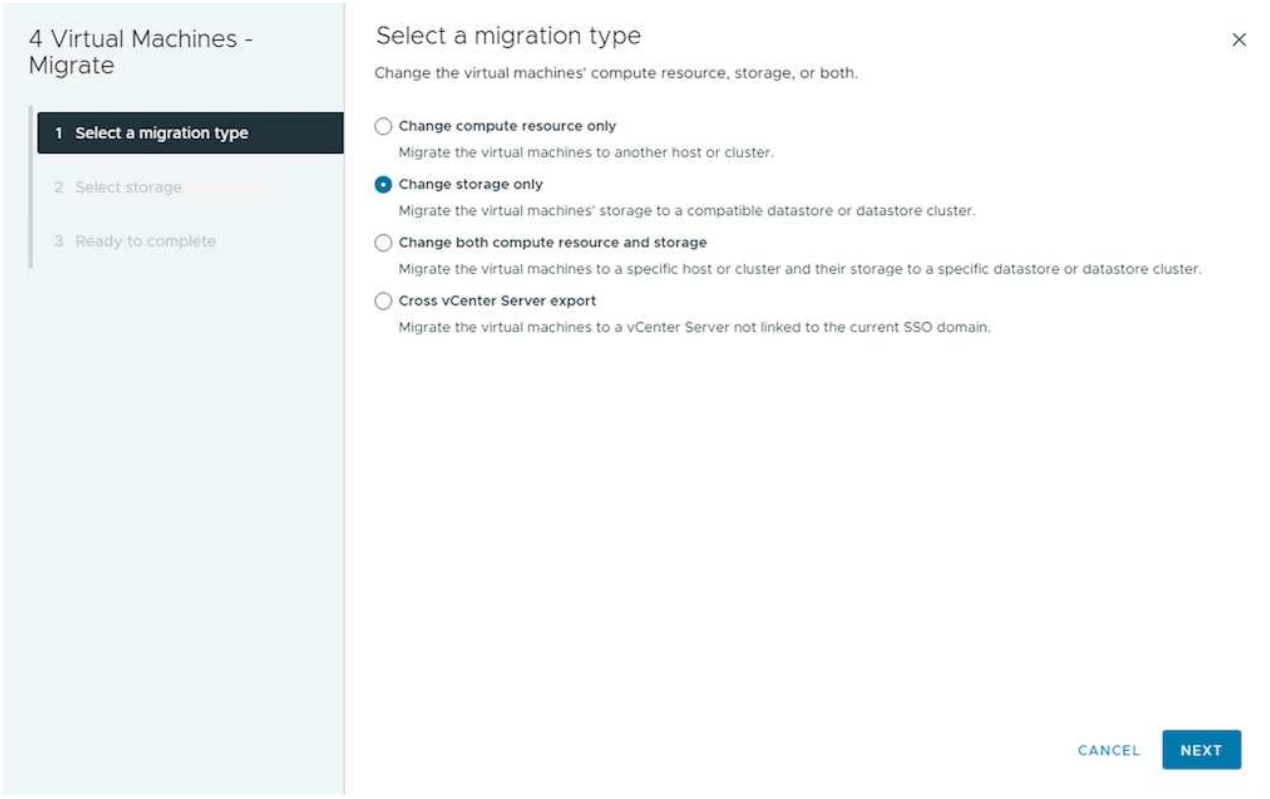

4. Select the desired VM Storage Policy and pick the datastore that is compatible. Click Next.

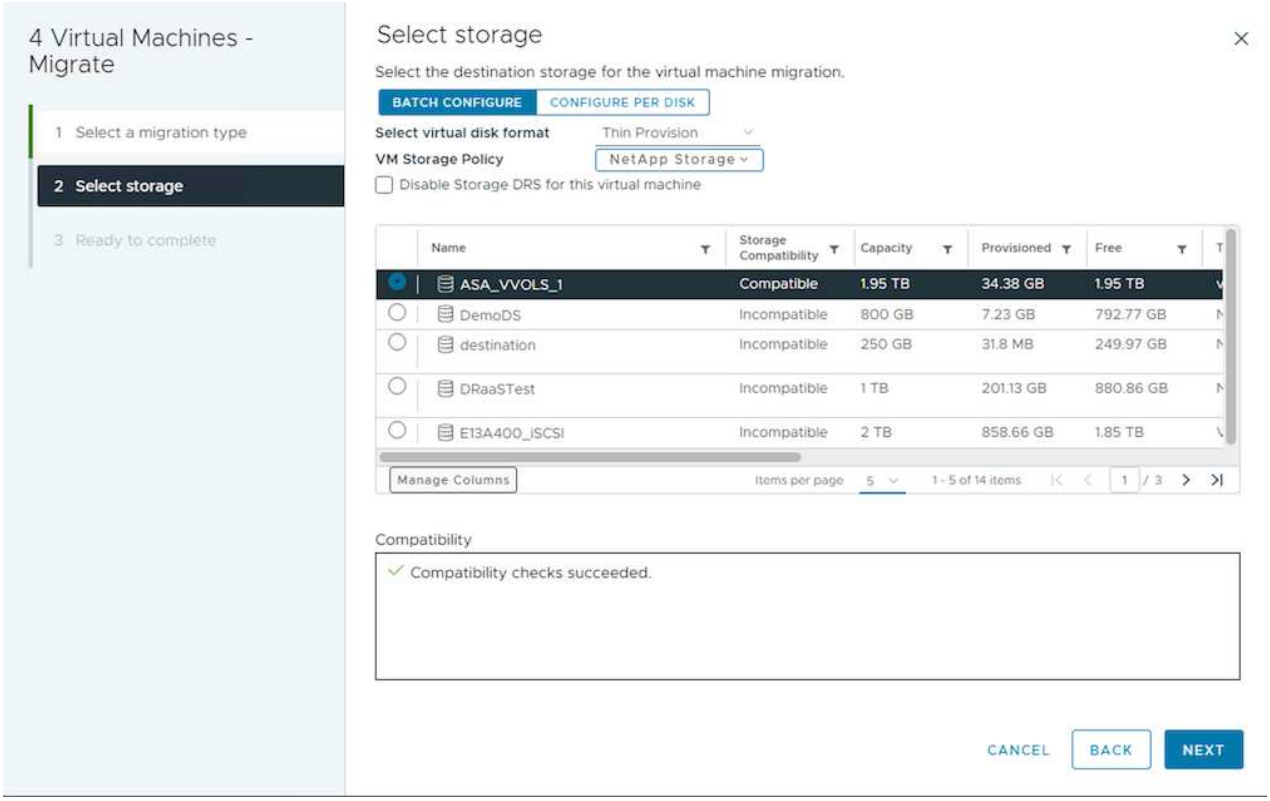

5. Review and click on Finish.

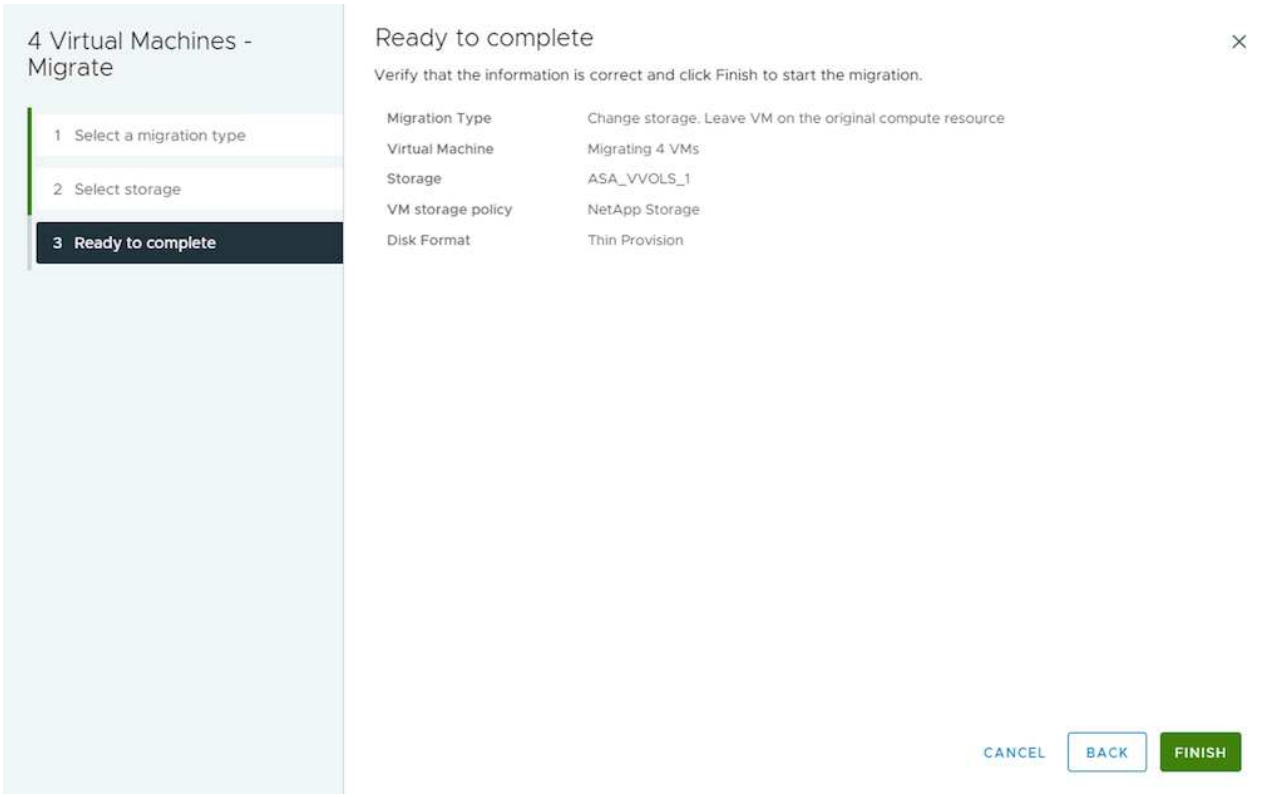

To migrate VMs using PowerCLI, here is the sample script.

```
#Authenticate to vCenter
Connect-VIServer -server vcsa.sddc.netapp.local -force

# Get all VMs with filter applied for a specific cluster
$vm = Get-Cluster 'vcf-m01-cl01' | Get-VM Aria*

#Gather VM Disk info
$vmdisk = $vm | Get-HardDisk

#Gather the desired Storage Policy to set for the VMs. Policy should be
available with valid datastores.
$storagepolicy = Get-SPBMStoragePolicy 'NetApp Storage'

#set VM Storage Policy for VM config and its data disks.
$vm, $vmdisk | Get-SPBMEntityConfiguration | Set-
SPBMEntityConfiguration -StoragePolicy $storagepolicy

#Migrate VMs to Datastore specified by Policy
$vm | Move-VM -Datastore (Get-SPBMCompatibleStorage -StoragePolicy
$storagepolicy)

#Ensure VM Storage Policy remains compliant.
$vm, $vmdisk | Get-SPBMEntityConfiguration
```

When Datastore Cluster is in use with fully automated storage DRS (Dynamic Resource Scheduling) and both (source & target) datastores are of same type (VMFS/NFS/vVol), Keep both datastores in same storage cluster and migrate VMs from source datastore by enabling maintenance mode on the source. Experience will be similar to how compute hosts are handled for maintenance.

Migration of VMs across multiple vSphere clusters

Refer CPU Compatibility and vSphere Enhanced vMotion Compatibility when source and target hosts are of different CPU family or model.

Follow the procedure below to migrate VMs to new Datastore using UI.

1. With vSphere Web Client, select the Cluster from the Host and Cluster inventory and click on VMs tab.

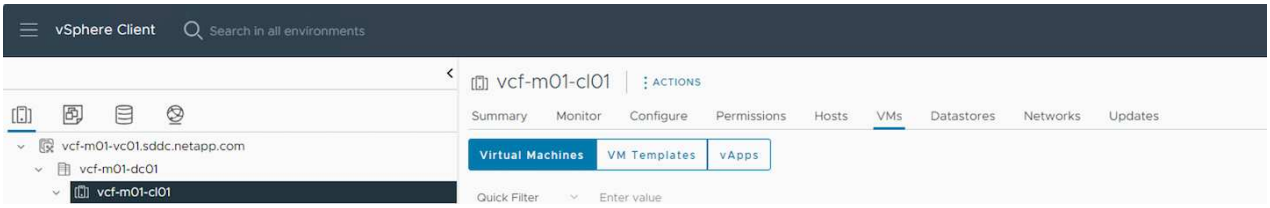

2. Select the VMs that needs to be migrated and right click to select Migrate option.

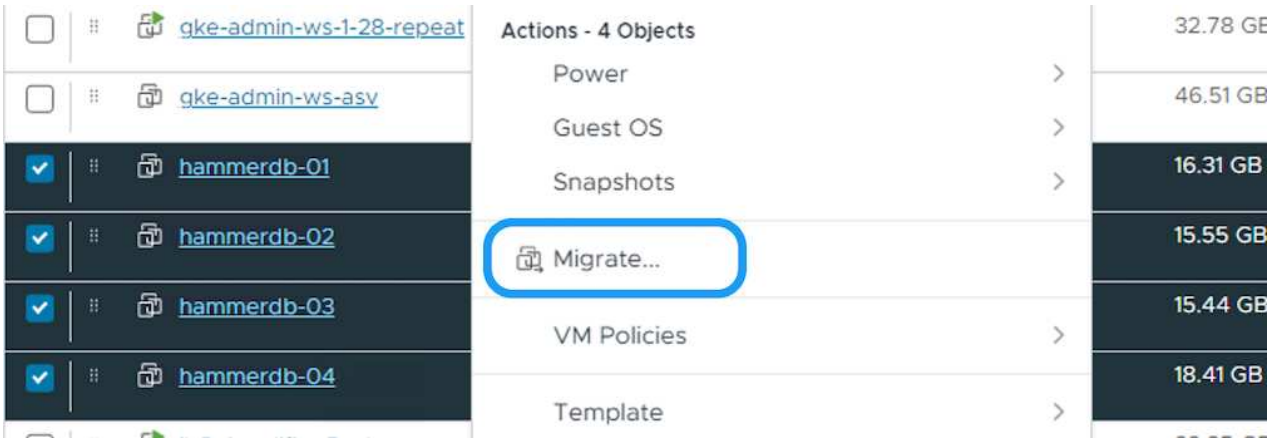

3. Choose option to change compute resource and storage, Click Next

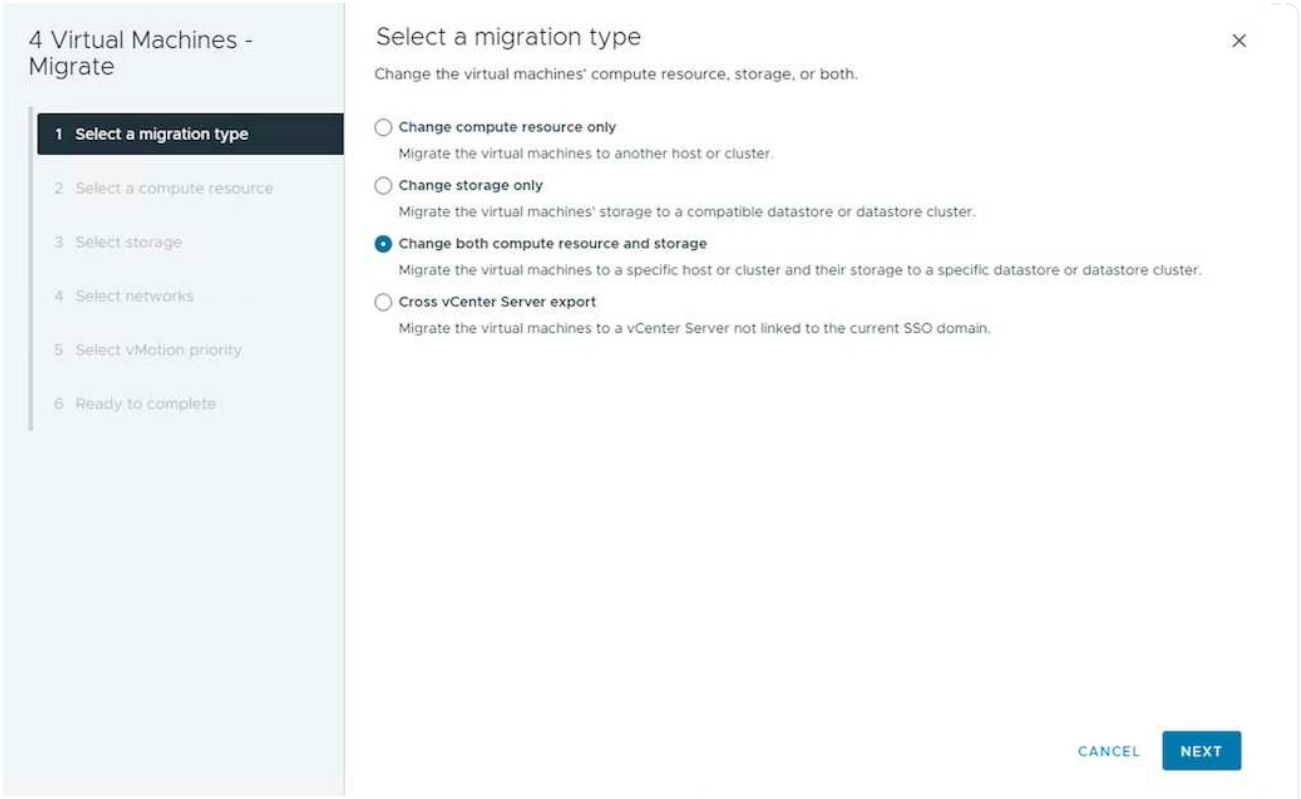

4. Navigate and pick the right cluster to migrate.

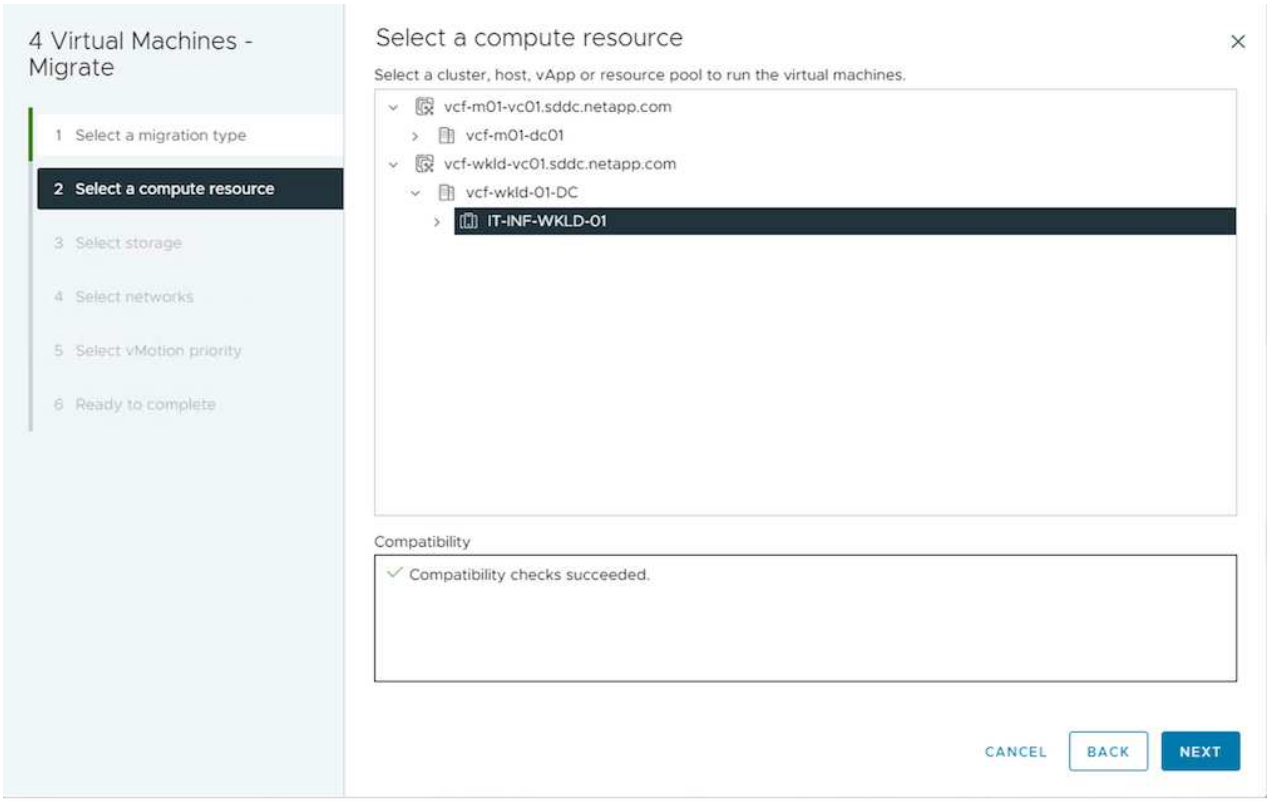

5. Select the desired VM Storage Policy and pick the datastore that is compatible. Click Next.

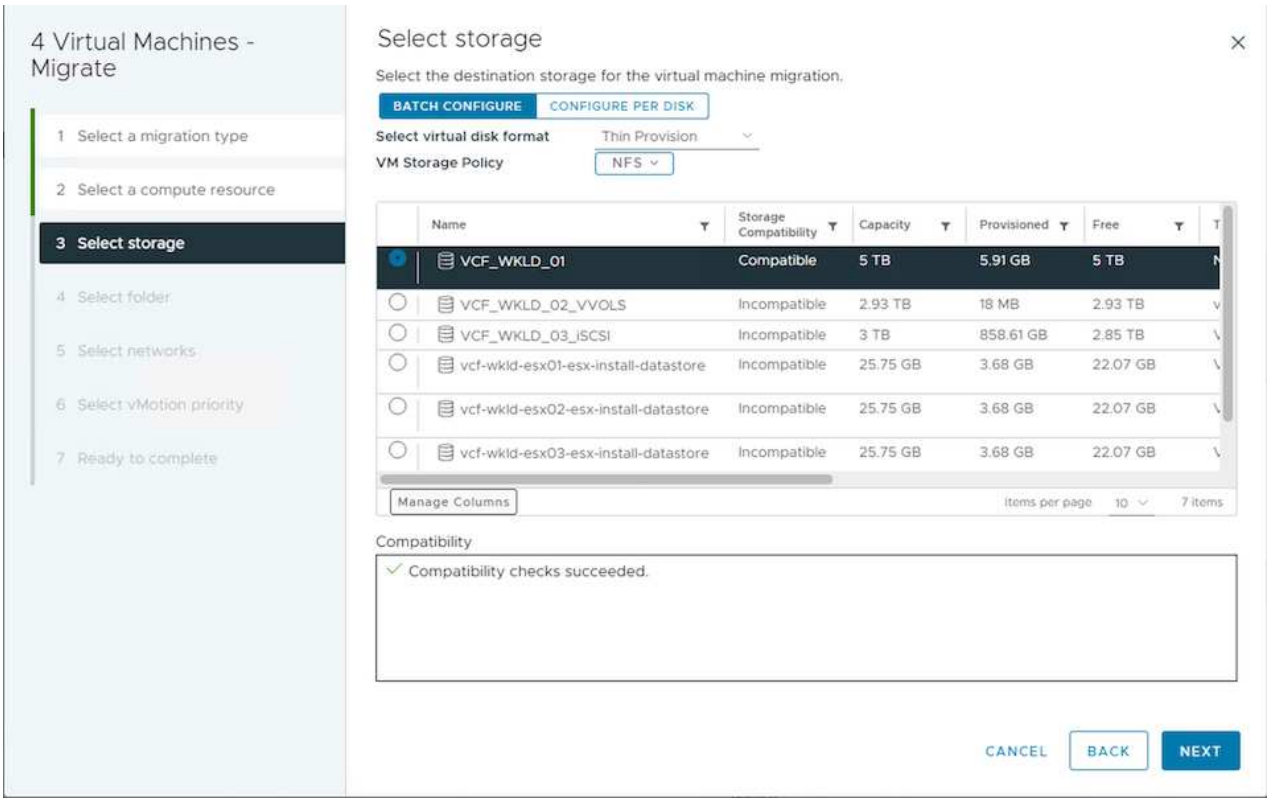

6. Pick the VM folder to place the target VMs.

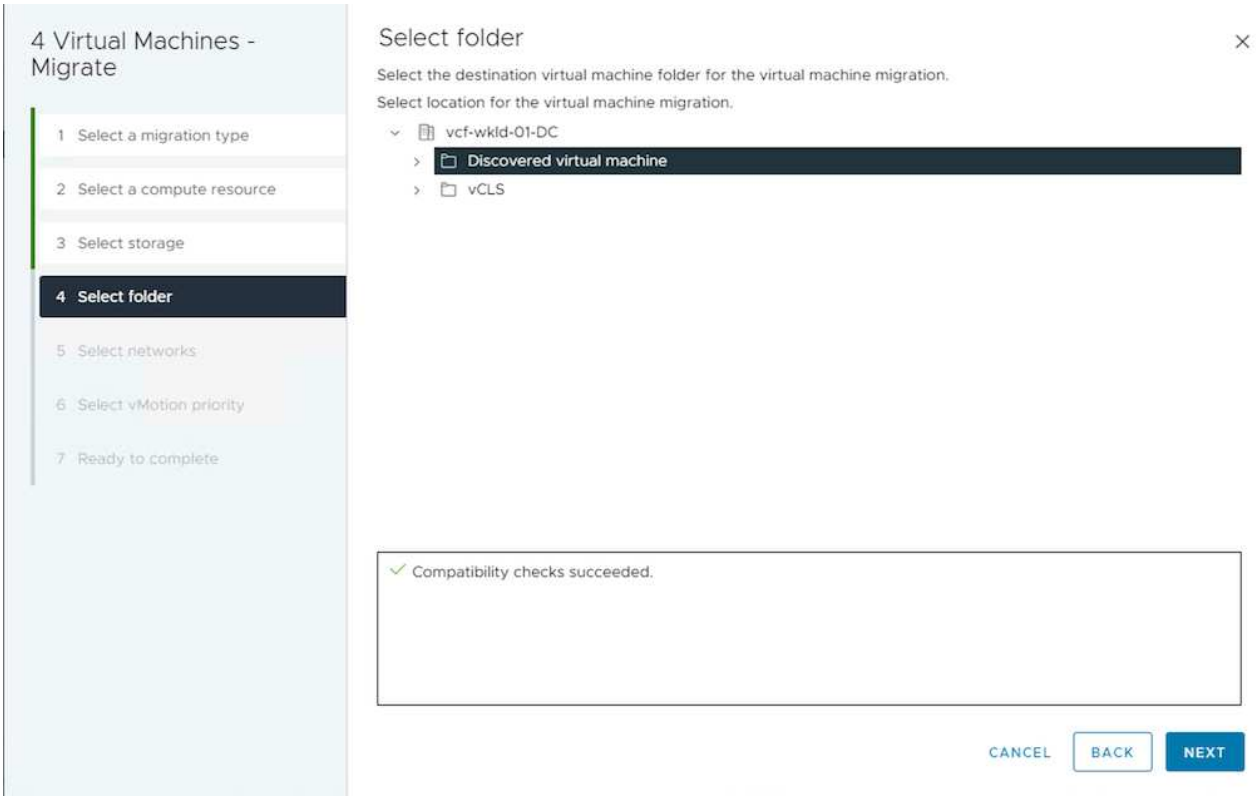

7. Select the target port group.

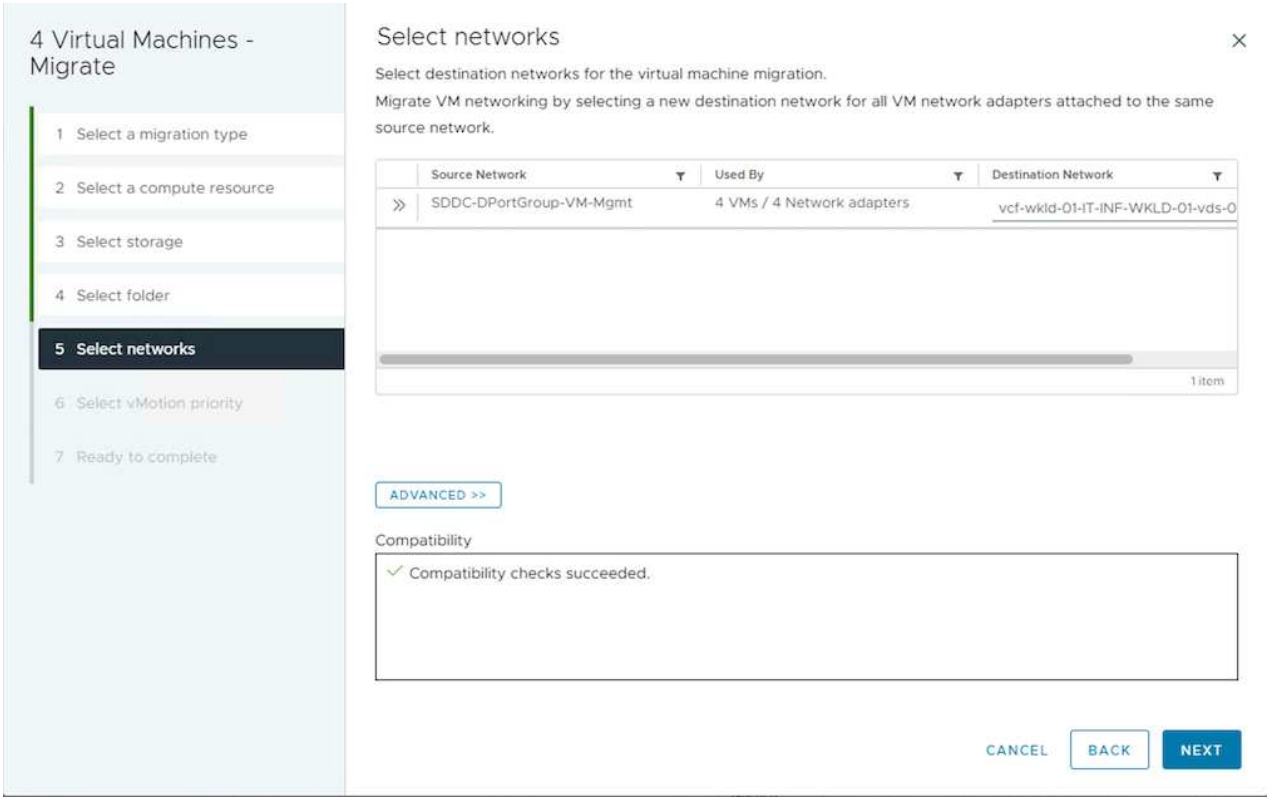

8. Review and click on Finish.

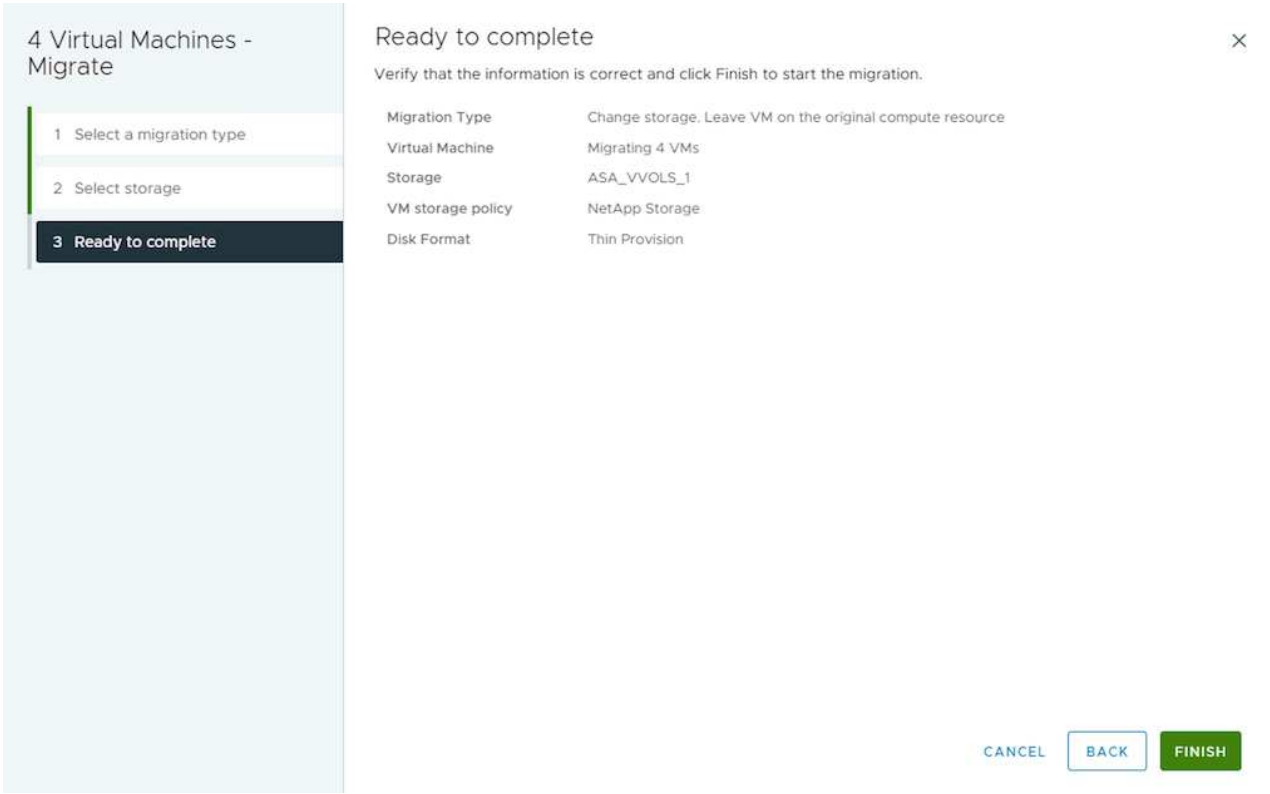

To migrate VMs using PowerCLI, here is the sample script.

```
#Authenticate to vCenter
Connect-VIServer -server vcsa.sddc.netapp.local -force

# Get all VMs with filter applied for a specific cluster
$vm = Get-Cluster 'vcf-m01-cl01' | Get-VM Aria*

#Gather VM Disk info
$vmdisk = $vm | Get-HardDisk

#Gather the desired Storage Policy to set for the VMs. Policy should be
available with valid datastores.
$storagepolicy = Get-SPBMStoragePolicy 'NetApp Storage'

#set VM Storage Policy for VM config and its data disks.
$vm, $vmdisk | Get-SPBMEntityConfiguration | Set-
SPBMEntityConfiguration -StoragePolicy $storagepolicy

#Migrate VMs to another cluster and Datastore specified by Policy
$vm | Move-VM -Destination (Get-Cluster 'Target Cluster') -Datastore
(Get-SPBMCompatibleStorage -StoragePolicy $storagepolicy)

#When Portgroup is specific to each cluster, replace the above command
with
$vm | Move-VM -Destination (Get-Cluster 'Target Cluster') -Datastore
(Get-SPBMCompatibleStorage -StoragePolicy $storagepolicy) -PortGroup
(Get-VirtualPortGroup 'VLAN 101')

#Ensure VM Storage Policy remains compliant.
$vm, $vmdisk | Get-SPBMEntityConfiguration
```

Migration of VMs across vCenter servers in same SSO domain

Follow the procedure below to migrate VMs to new vCenter server which is listed on same vSphere Client UI.

> For additional requirements like source and target vCenter versions,etc., check vSphere documentation on requirements for vMotion between vCenter server instances

1. With vSphere Web Client, select the Cluster from the Host and Cluster inventory and click on VMs tab.

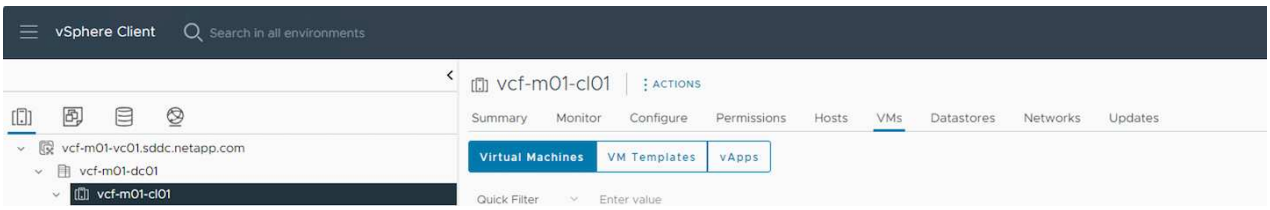

2. Select the VMs that needs to be migrated and right click to select Migrate option.

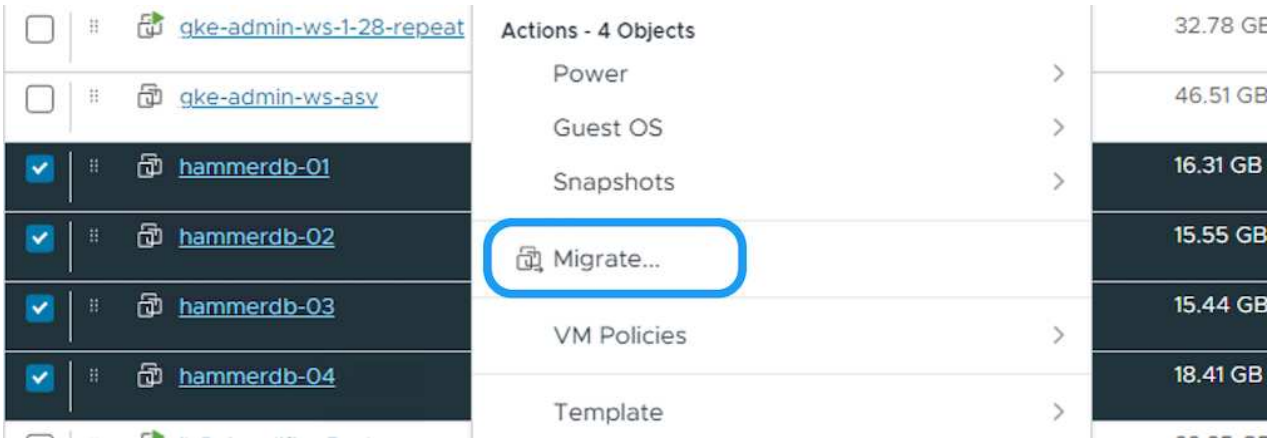

3. Choose option to change compute resource and storage, Click Next

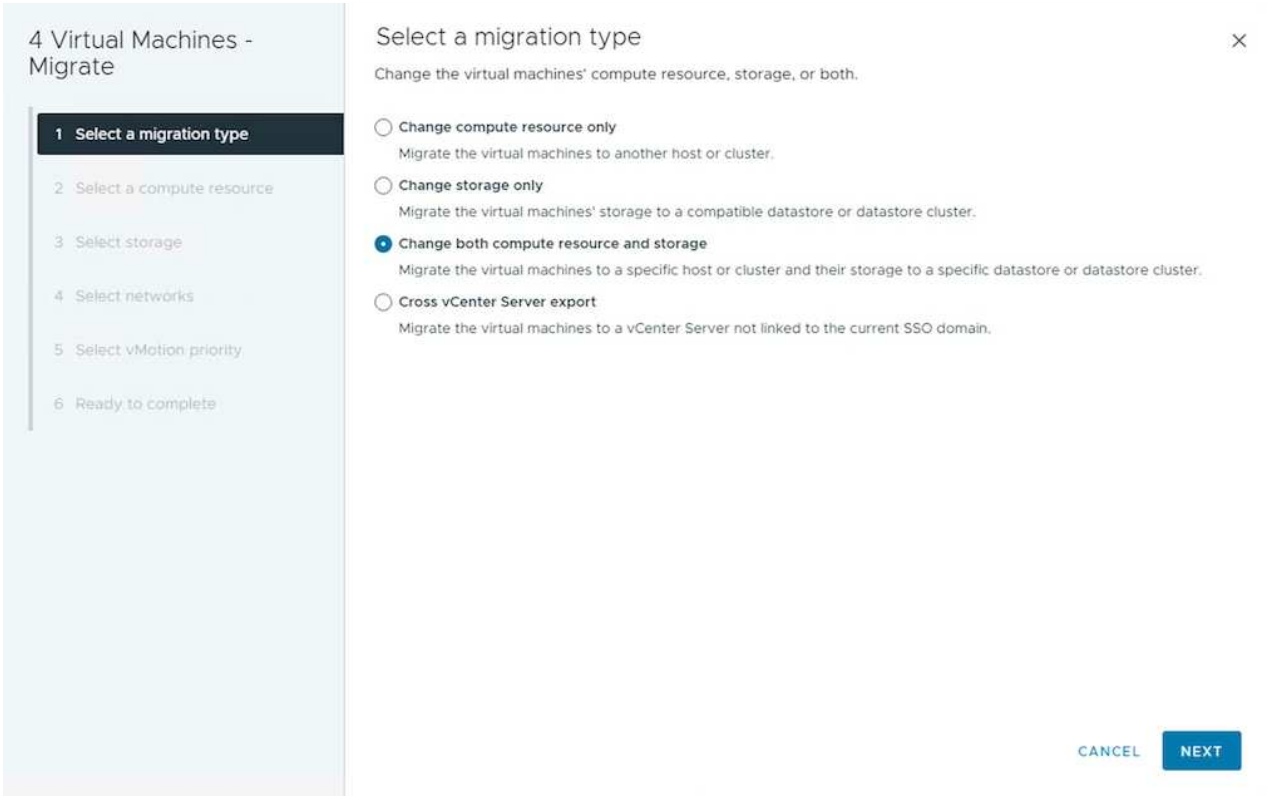

4. Select the target cluster in target vCenter server.

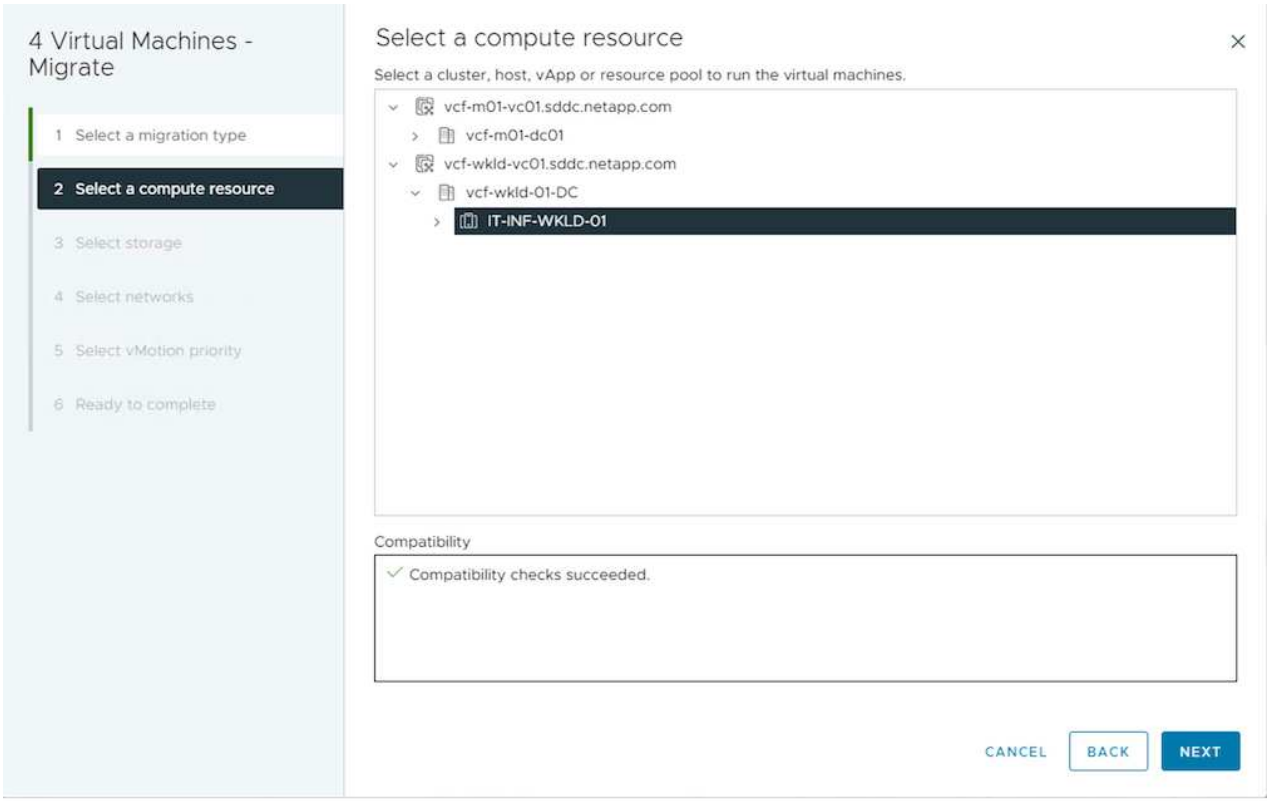

5. Select the desired VM Storage Policy and pick the datastore that is compatible. Click Next.

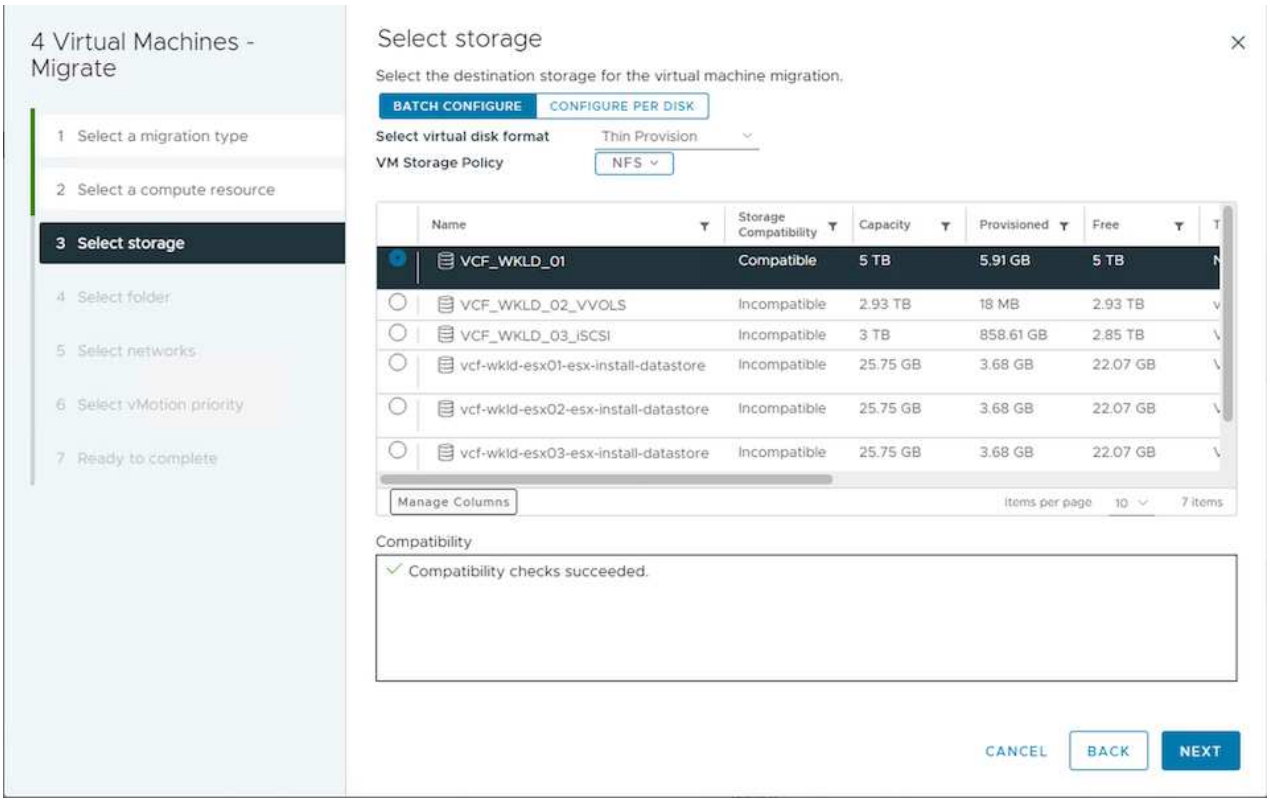

6. Pick the VM folder to place the target VMs.

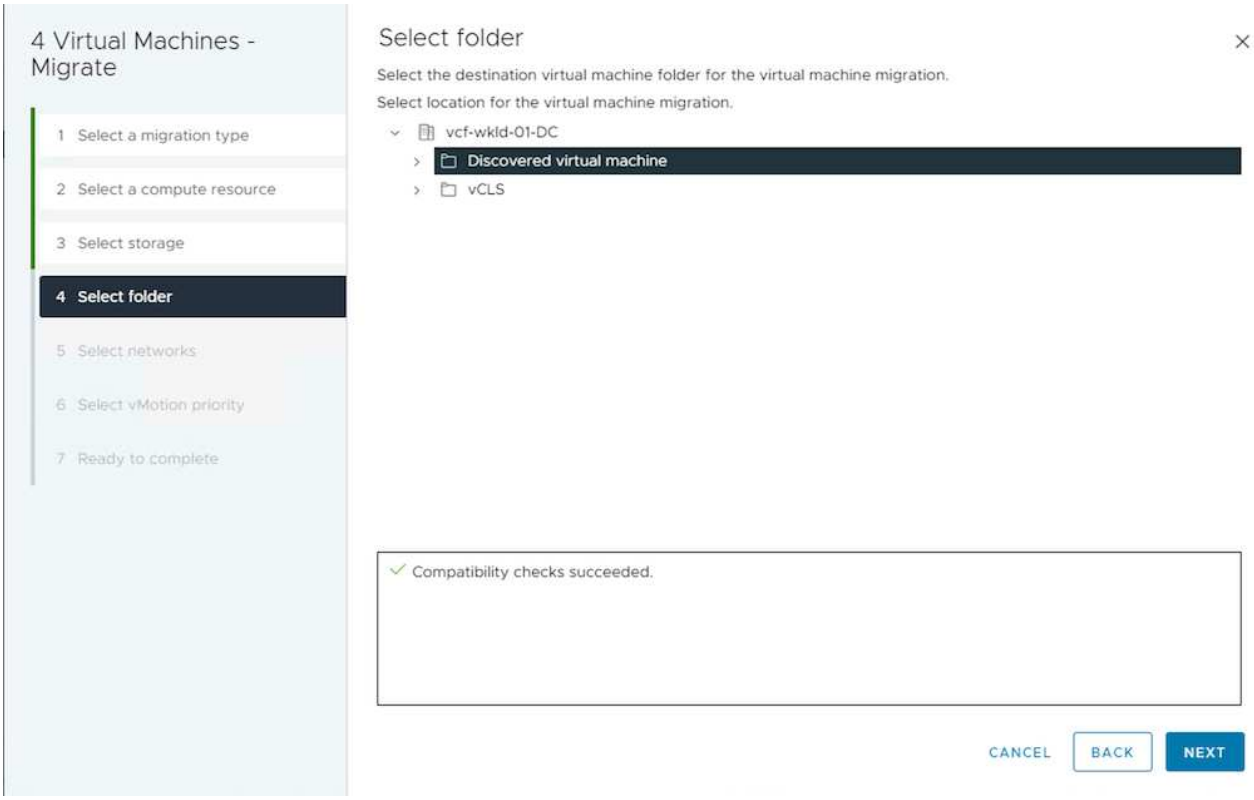

7. Select the target port group.

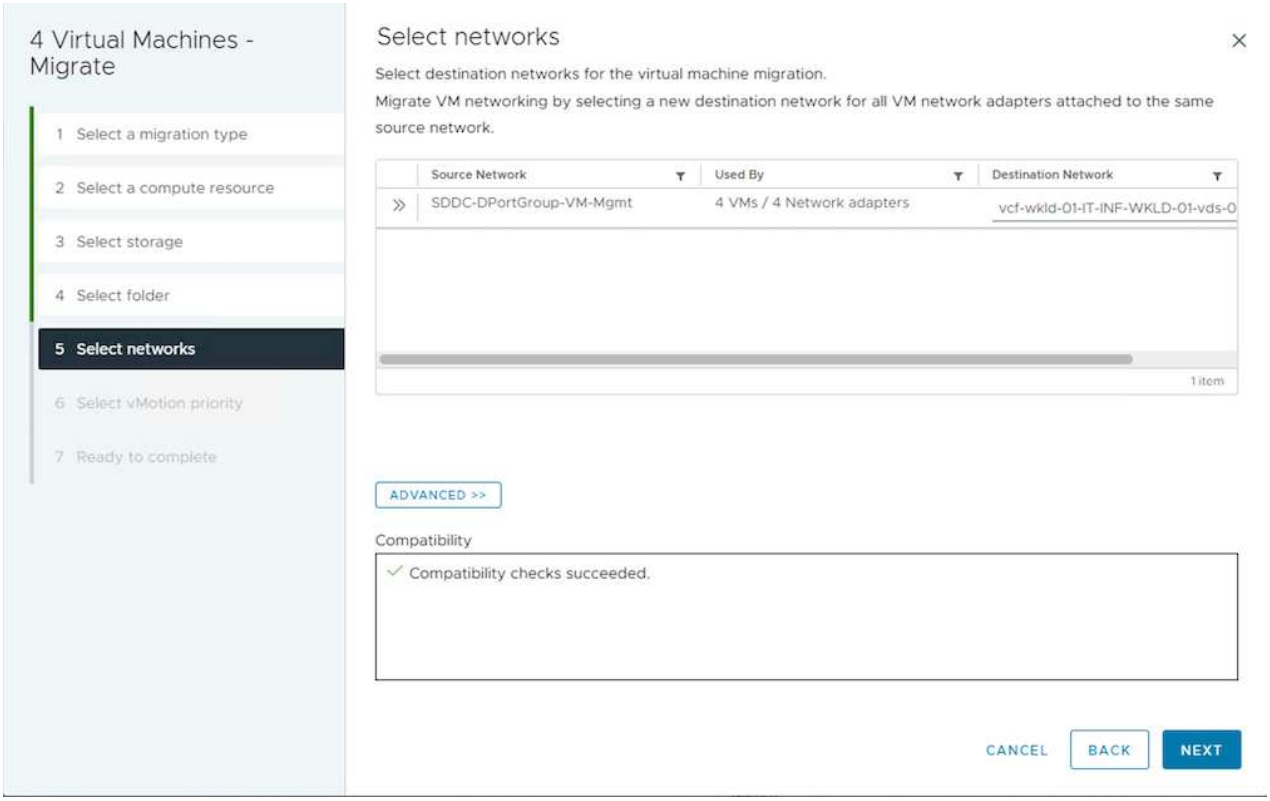

8. Review the migration options and click Finish.

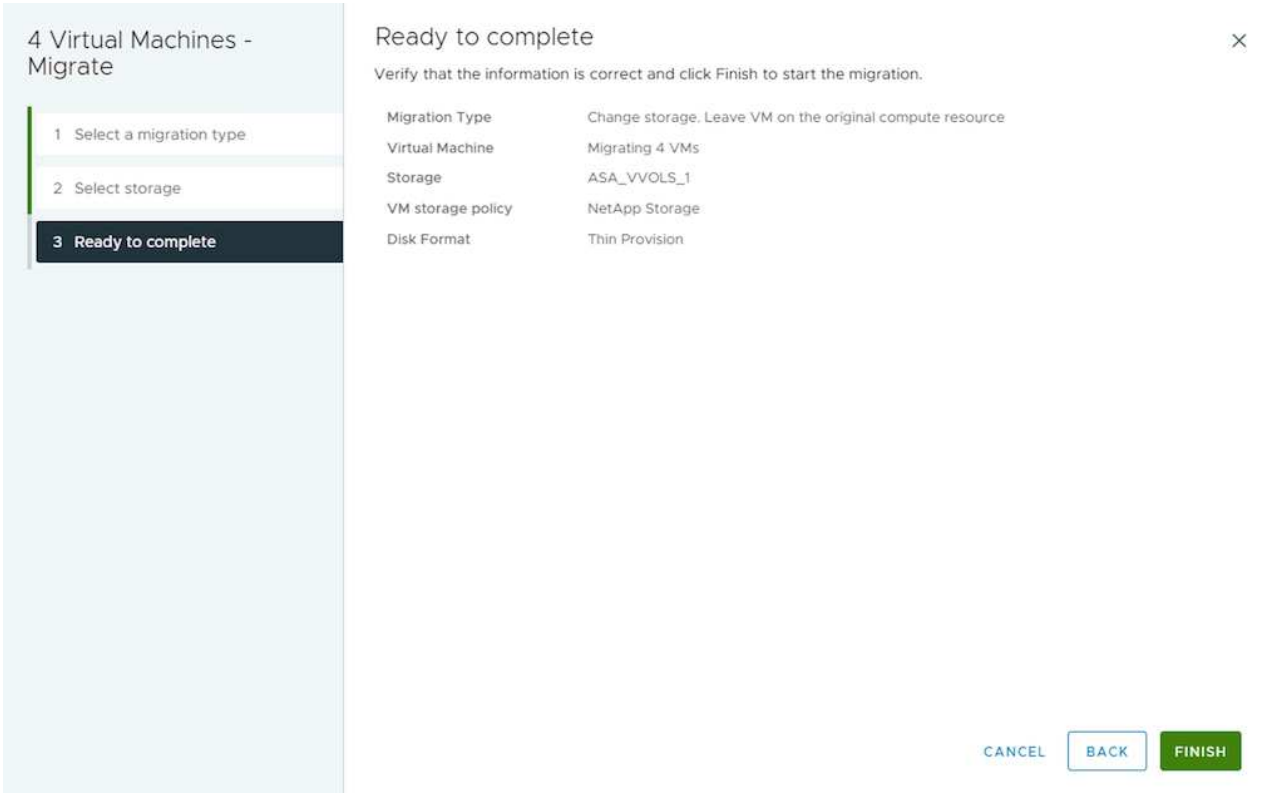

To migrate VMs using PowerCLI, here is the sample script.

```
#Authenticate to Source vCenter
$sourcevc = Connect-VIServer -server vcsa01.sddc.netapp.local -force
$targetvc = Connect-VIServer -server vcsa02.sddc.netapp.local -force

# Get all VMs with filter applied for a specific cluster
$vm = Get-Cluster 'vcf-m01-cl01'  -server $sourcevc| Get-VM Win*

#Gather the desired Storage Policy to set for the VMs. Policy should be
available with valid datastores.
$storagepolicy = Get-SPBMStoragePolicy 'iSCSI' -server $targetvc

#Migrate VMs to target vCenter
$vm | Move-VM -Destination (Get-Cluster 'Target Cluster' -server
$targetvc) -Datastore (Get-SPBMCompatibleStorage -StoragePolicy
$storagepolicy -server $targetvc) -PortGroup (Get-VirtualPortGroup
'VLAN 101' -server $targetvc)

$targetvm = Get-Cluster 'Target Cluster' -server $targetvc | Get-VM
Win*

#Gather VM Disk info
$targetvmdisk = $targetvm | Get-HardDisk

#set VM Storage Policy for VM config and its data disks.
$targetvm, $targetvmdisk | Get-SPBMEntityConfiguration | Set-
SPBMEntityConfiguration -StoragePolicy $storagepolicy

#Ensure VM Storage Policy remains compliant.
$targetvm, $targetvmdisk | Get-SPBMEntityConfiguration
```

Migration of VMs across vCenter servers in different SSO domain

This scenario assumes the communication exists between the vCenter servers. Otherwise check the across datacenter location scenario listed below. For prerequisites, check vSphere documentation on Advanced Cross vCenter vMotion

Follow the procedure below to migrate VMs to differnt vCenter server using UI.

1. With vSphere Web Client, select the source vCenter server and click on VMs tab.

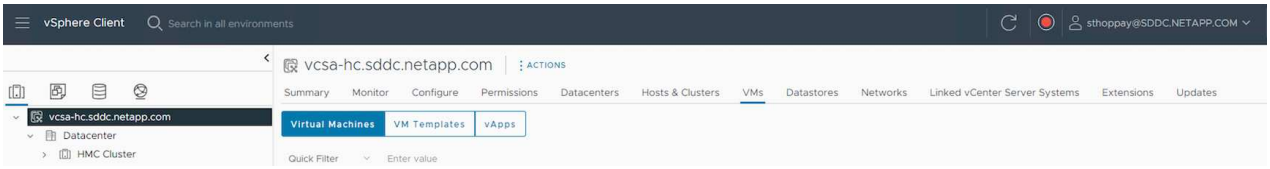

2. Select the VMs that needs to be migrated and right click to select Migrate option.

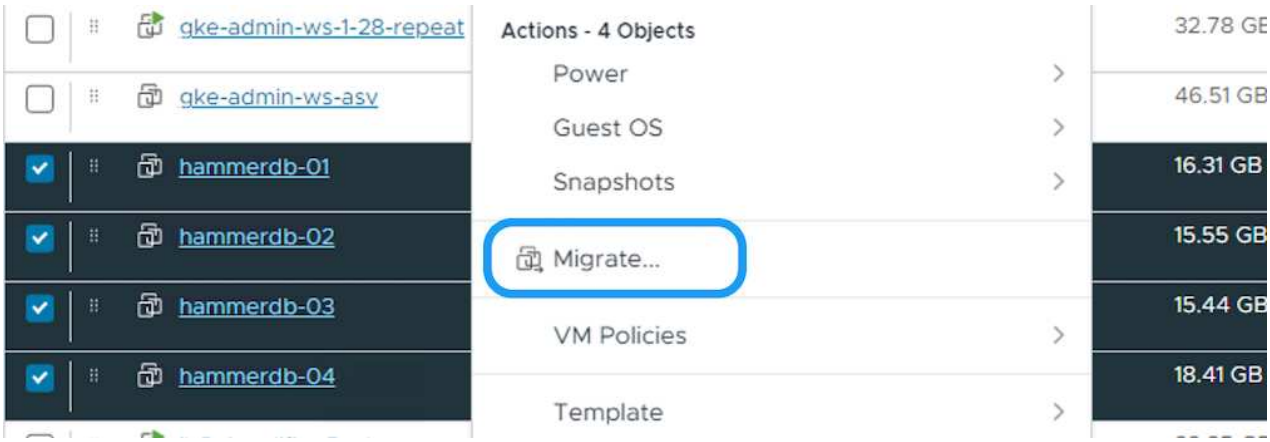

3. Choose option Cross vCenter Server export, Click Next

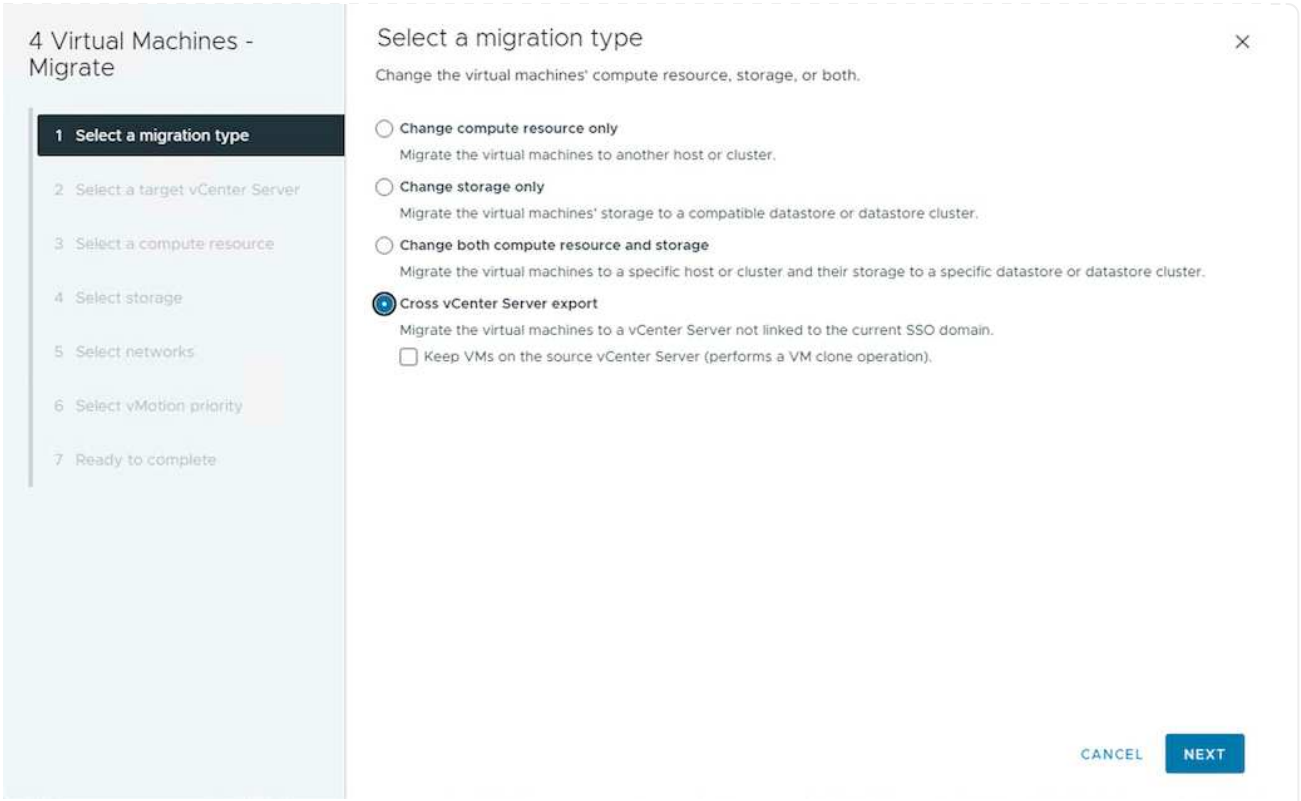

VM can also be imported from the target vCenter server. For that procedure, check Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion

4. Provide vCenter credential details and click Login.

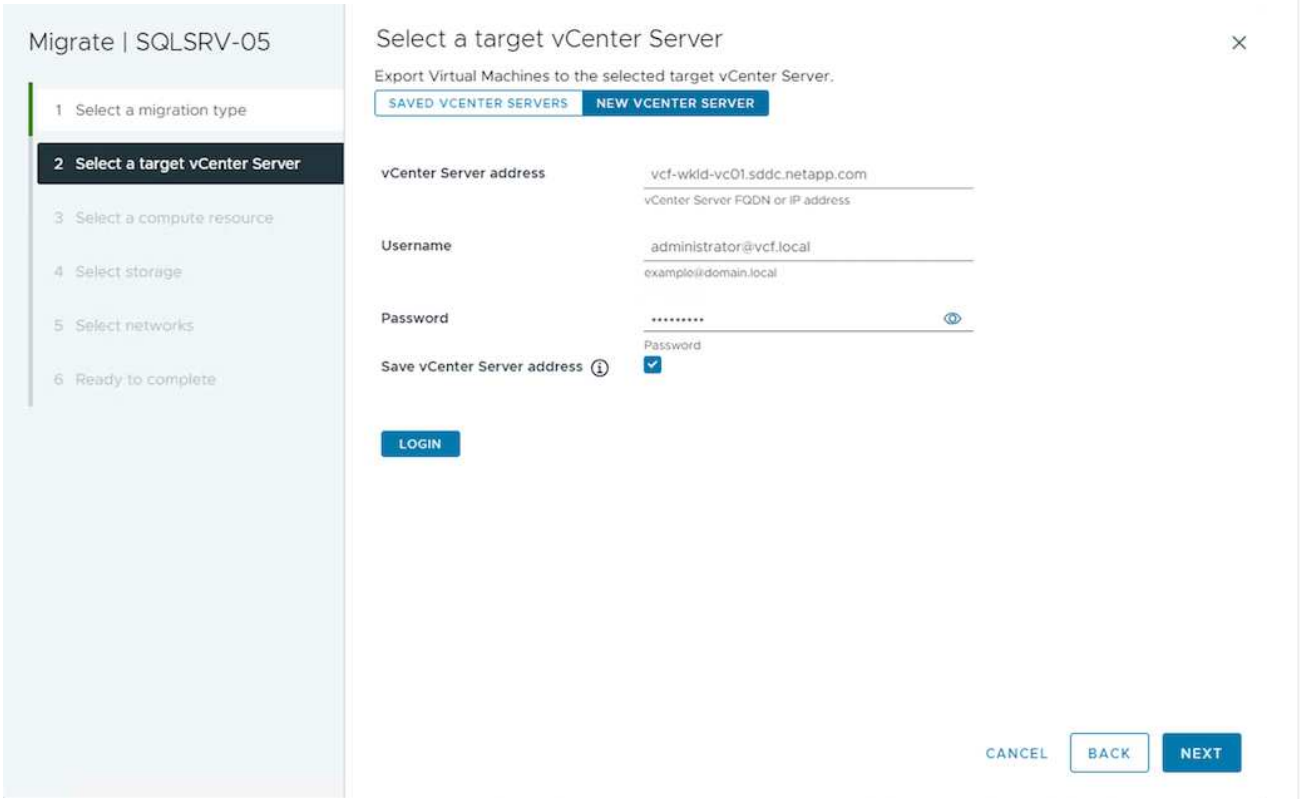

5. Confirm and Accept the SSL certificate thumbprint of vCenter server

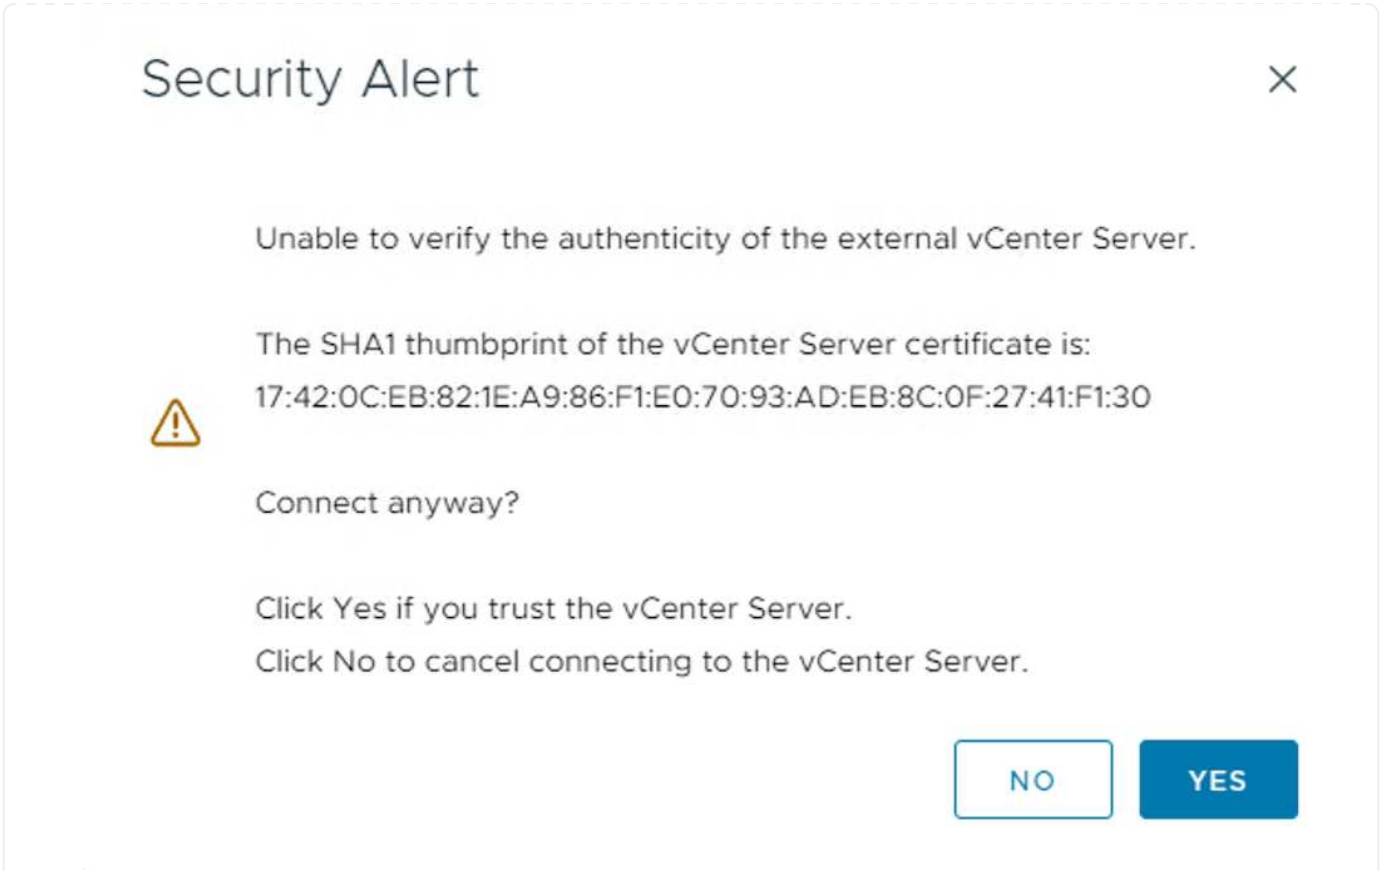

6. Expand target vCenter and select the target compute cluster.

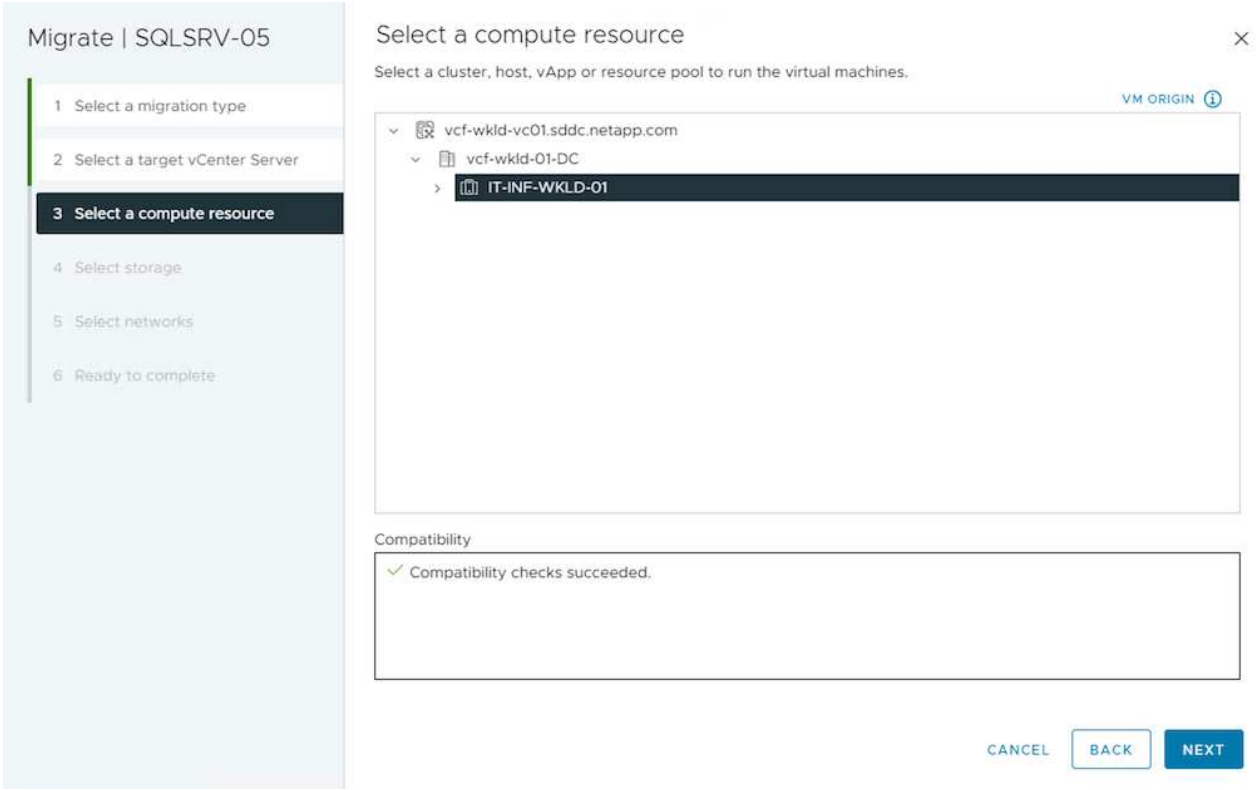

7. Select the target datastore based on the VM Storage Policy.

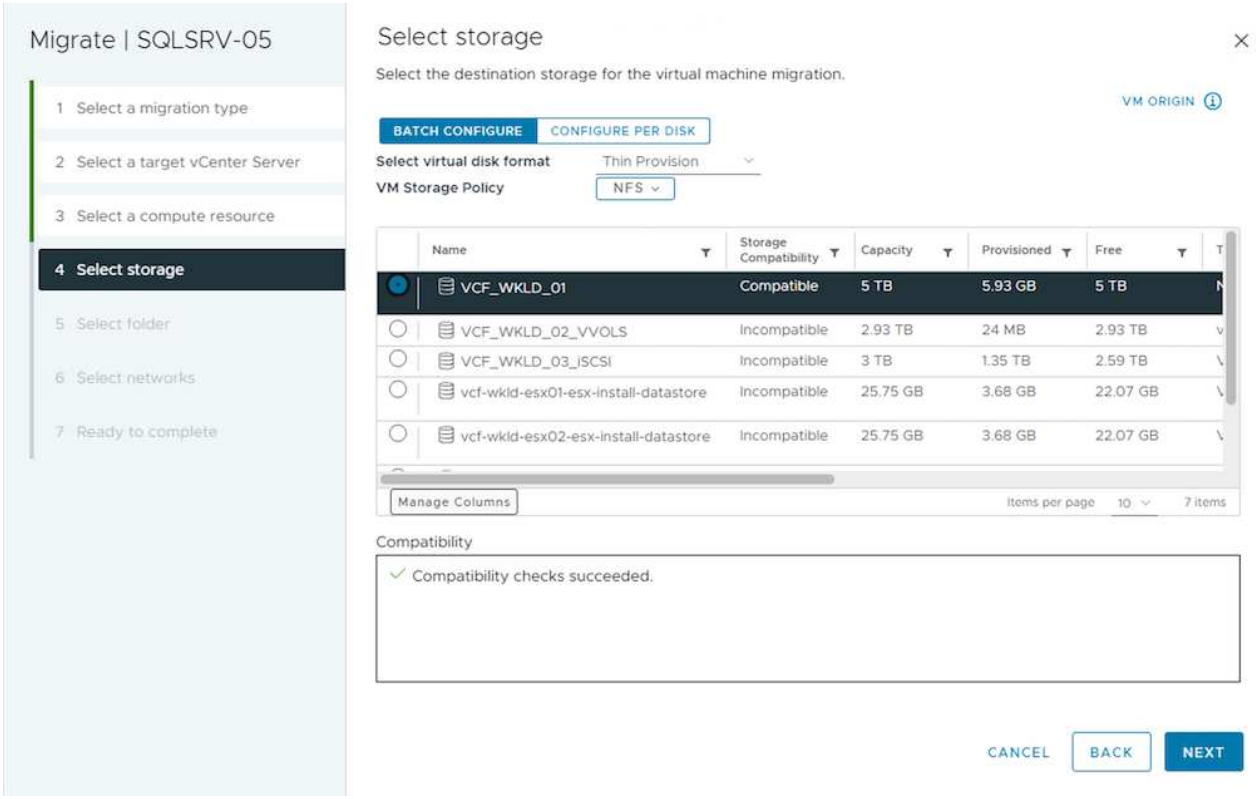

8. Select the target VM folder.

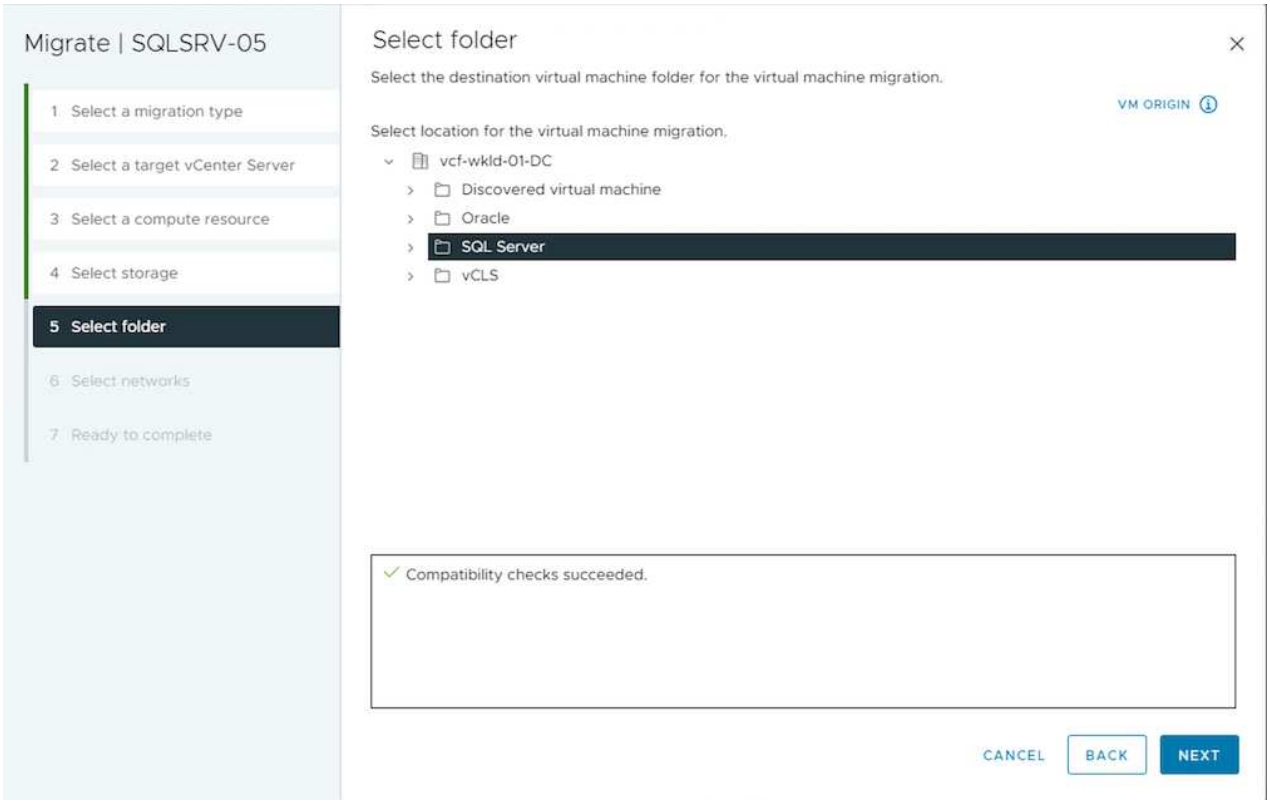

9. Pick the VM portgroup for each network interface card mapping.

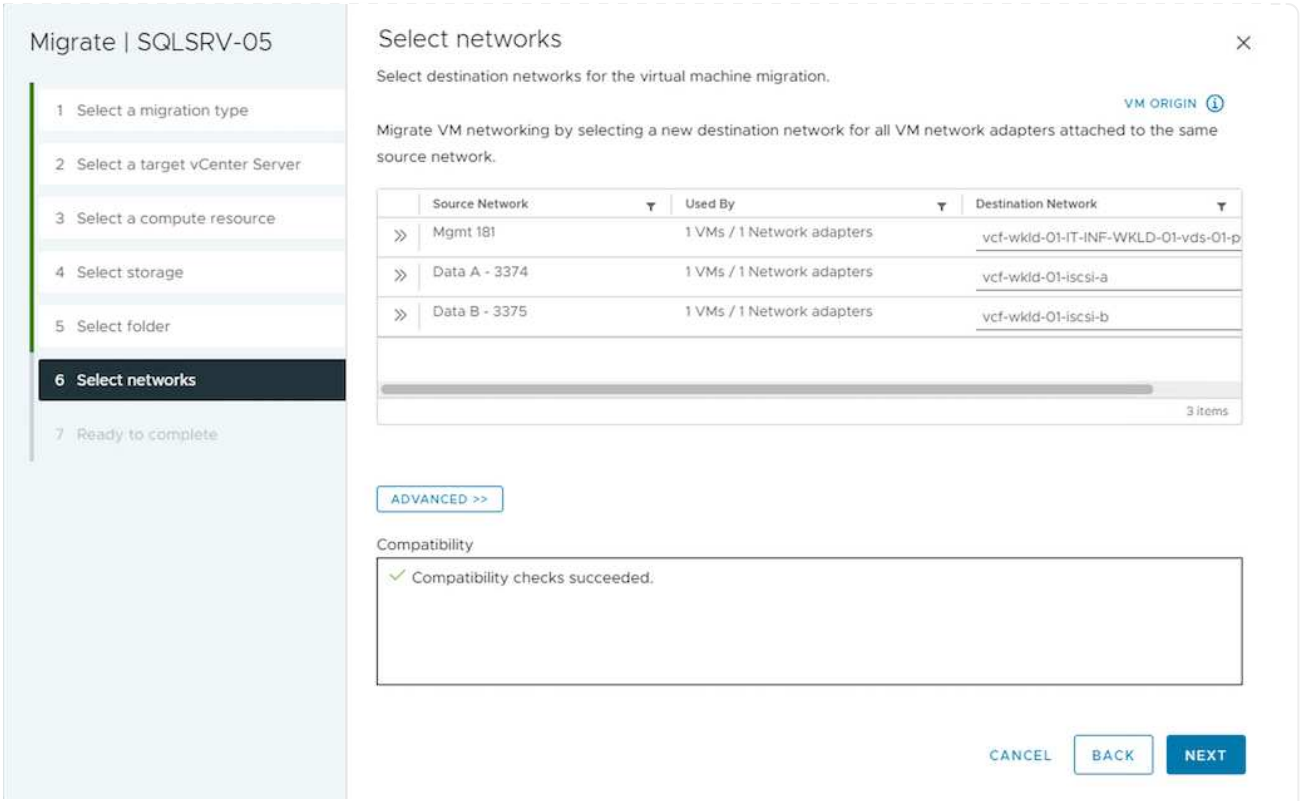

10. Review and click Finish to start the vMotion across the vCenter servers.

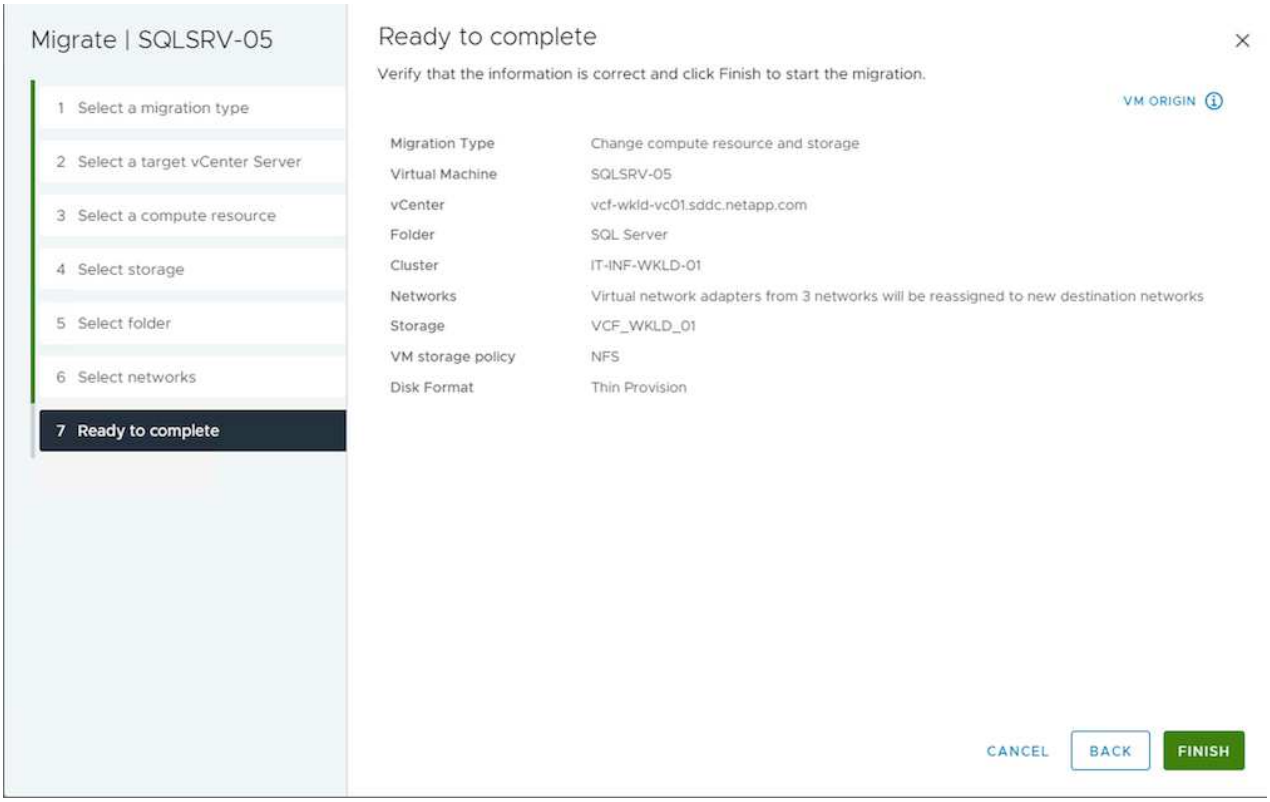

To migrate VMs using PowerCLI, here is the sample script.

```
#Authenticate to Source vCenter
$sourcevc = Connect-VIServer -server vcsa01.sddc.netapp.local -force
$targetvc = Connect-VIServer -server vcsa02.sddc.netapp.local -force

# Get all VMs with filter applied for a specific cluster
$vm = Get-Cluster 'Source Cluster'  -server $sourcevc| Get-VM Win*

#Gather the desired Storage Policy to set for the VMs. Policy should be
available with valid datastores.
$storagepolicy = Get-SPBMStoragePolicy 'iSCSI' -server $targetvc

#Migrate VMs to target vCenter
$vm | Move-VM -Destination (Get-Cluster 'Target Cluster' -server
$targetvc) -Datastore (Get-SPBMCompatibleStorage -StoragePolicy
$storagepolicy -server $targetvc) -PortGroup (Get-VirtualPortGroup
'VLAN 101' -server $targetvc)

$targetvm = Get-Cluster 'Target Cluster' -server $targetvc | Get-VM
Win*

#Gather VM Disk info
$targetvmdisk = $targetvm | Get-HardDisk

#set VM Storage Policy for VM config and its data disks.
$targetvm, $targetvmdisk | Get-SPBMEntityConfiguration | Set-
SPBMEntityConfiguration -StoragePolicy $storagepolicy

#Ensure VM Storage Policy remains compliant.
$targetvm, $targetvmdisk | Get-SPBMEntityConfiguration
```

Migration of VMs across datacenter locations

- When Layer 2 traffic is stretched across datacenters either by using NSX Federation or other options, follow the procedure for migrating VMs across vCenter servers.
- HCX provides various migration types including Replication Assisted vMotion across the datacenters to move VM without any downtime.
- Site Recovery Manager (SRM) is typically meant for Disaster Recovery purposes and also often used for planned migration utilizing storage array based replication.
- Continous Data Protection (CDP) products use vSphere API for IO (VAIO) to intercept the data and send a copy to remote location for near zero RPO solution.
- Backup and Recovery products can also be utilized. But often results in longer RTO.
- BlueXP Disaster Recovery as a Service (DRaaS) utilizes storage array based replication and automates certain tasks to recover the VMs at target site.

Migration of VMs in hybrid cloud environment

- Configure Hybrid Linked Mode and follow the procedure of Migration of VMs across vCenter servers in same SSO domain
- HCX provides various migration types including Replication Assisted vMotion across the datacenters to move VM while it is powered on.
  - TR 4942: Migrate Workloads to FSx ONTAP datastore using VMware HCX
  - TR-4940: Migrate workloads to Azure NetApp Files datastore using VMware HCX - Quickstart guide
  - Migrate workloads to NetApp Cloud Volume Service datastore on Google Cloud VMware Engine using VMware HCX - Quickstart guide
- BlueXP Disaster Recovery as a Service (DRaaS) utilizes storage array based replication and automates certain tasks to recover the VMs at target site.
- With supported Continous Data Protection (CDP) products that use vSphere API for IO (VAIO) to intercept the data and send a copy to remote location for near zero RPO solution.

When the source VM resides on block vVol datastore, it can be replicated with SnapMirror to Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP (CVO) at other supported cloud providers and consume as iSCSI volume with cloud native VMs.

VM Template Migration Scenarios

VM Templates can be managed by vCenter Server or by a content library. Distribution of VM templates, OVF and OVA templates, other types of files are handled by publishing it in local content library and remote content libraries can subscribe to it.

- VM templates stored on vCenter inventory can be converted to VM and use the VM migration options.
- OVF and OVA templates, other types of files stored on content library can be cloned to other content libraries.
- Content library VM Templates can be hosted on any datastore and needs to be added into new content library.

Migration of VM templates hosted on datastore

1. In vSphere Web Client, right click on the VM template under VM and Templates folder view and select option to convert to VM.

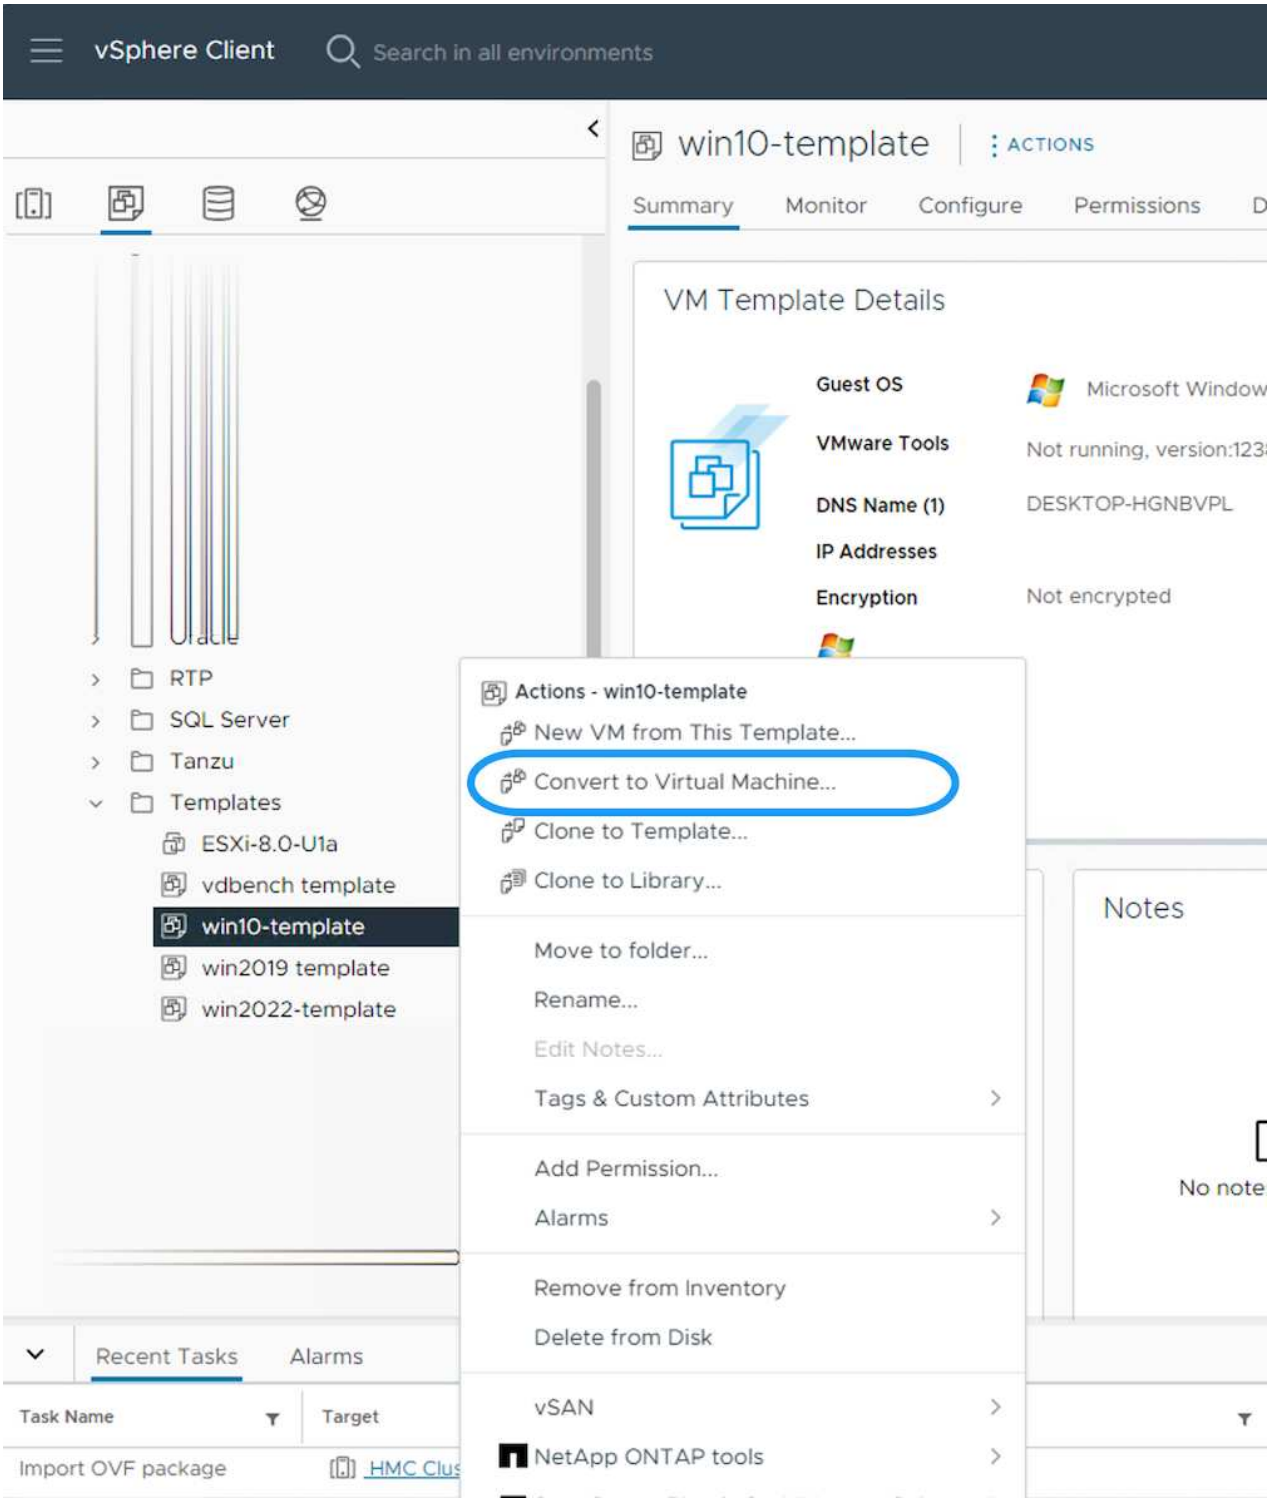

2. Once it is converted as VM, follow the VM migration options.

Clone of Content Library items

1. In vSphere Web Client, select Content Libraries

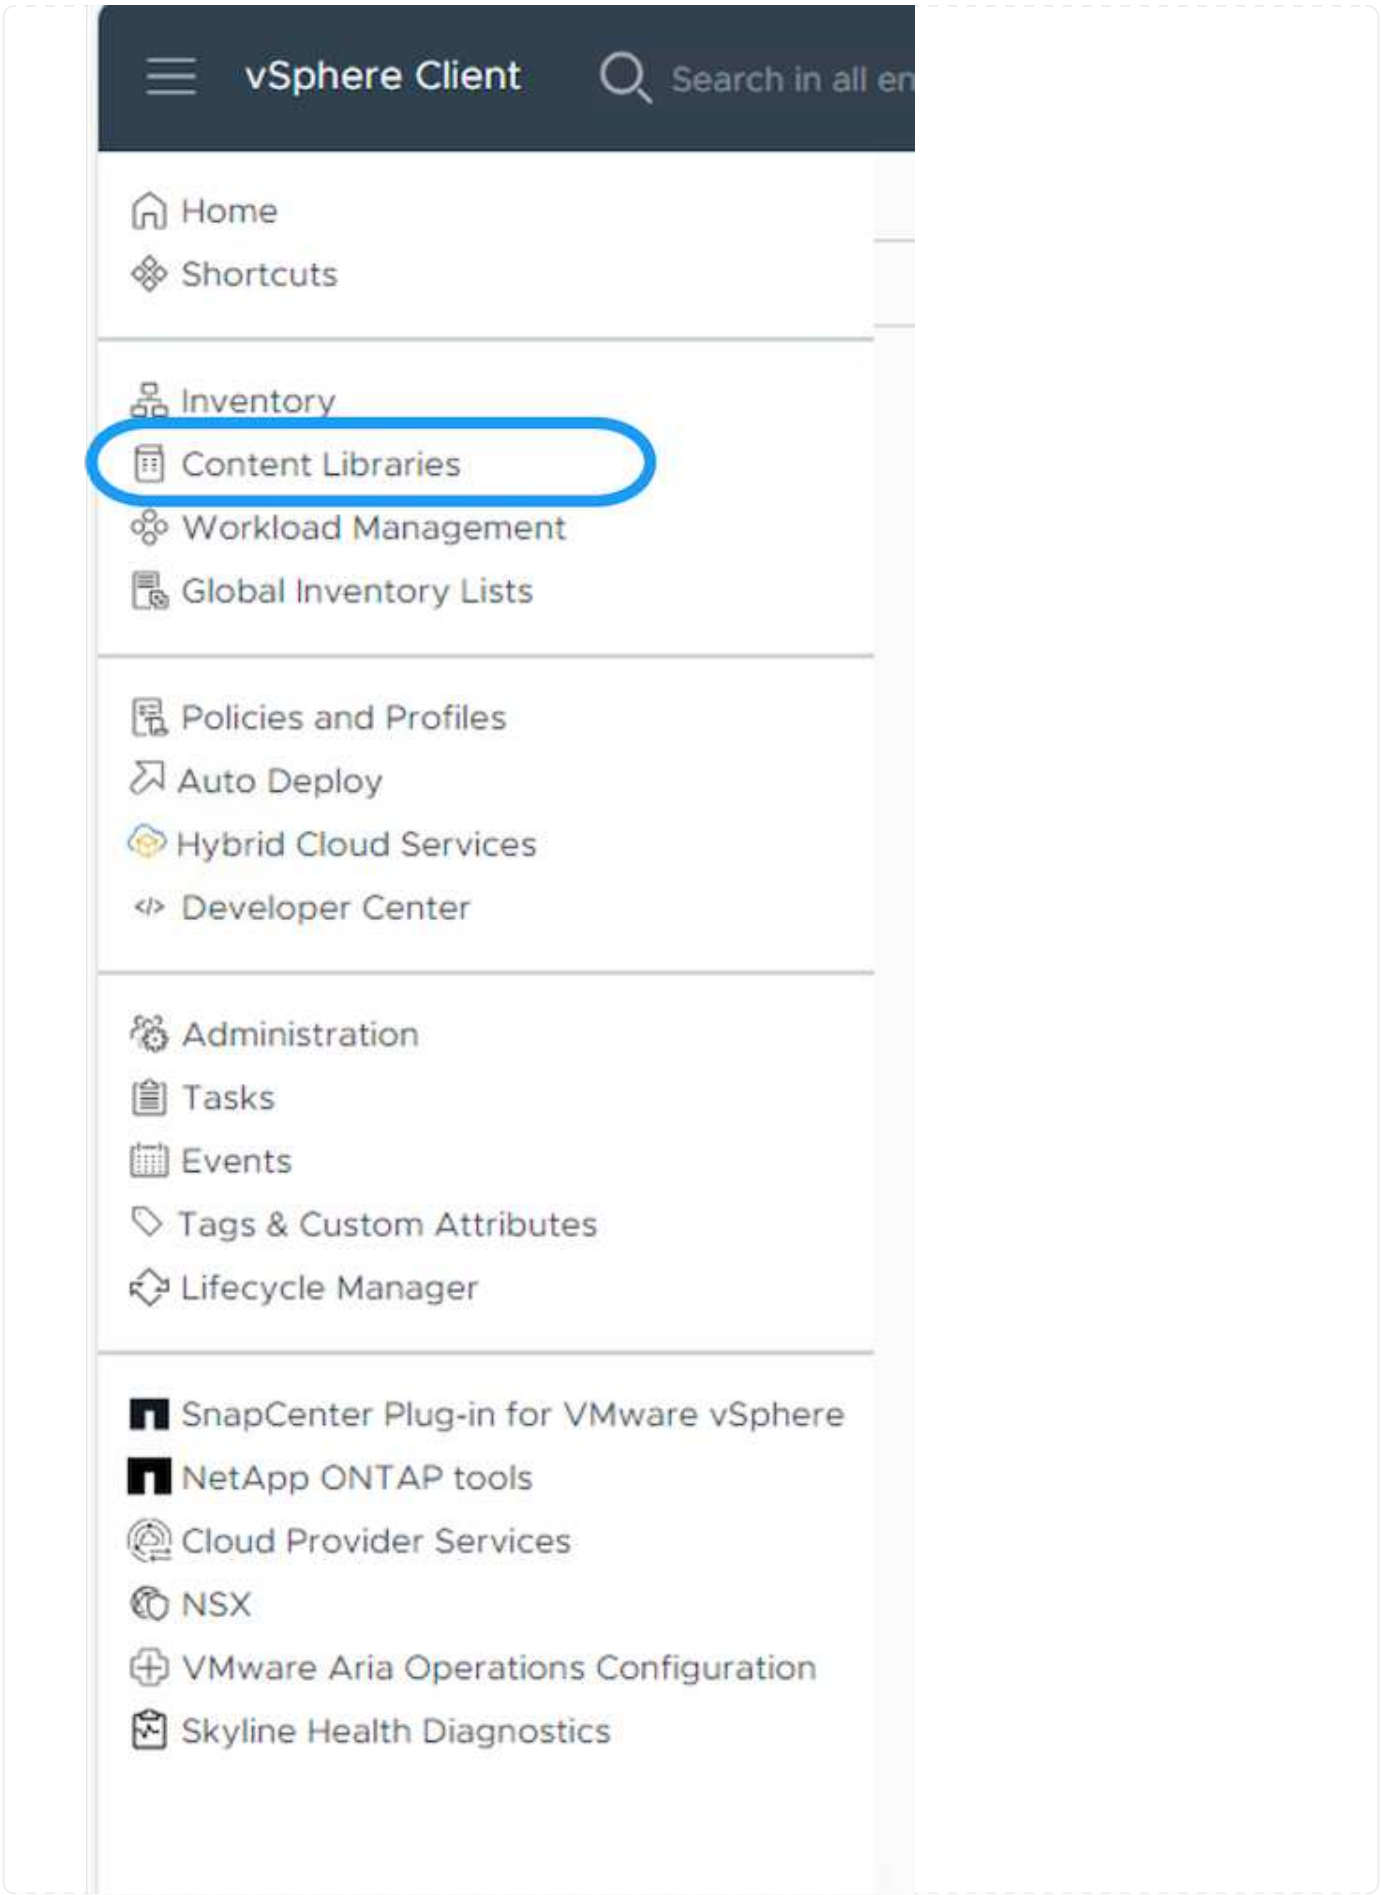
vSphere Client
Search in all en
Home
Shortcuts
Inventory
Content Libraries
Workload Management
Global Inventory Lists
Policies and Profiles
Auto Deploy
Hybrid Cloud Services
Developer Center
Administration
Tasks
Events
Tags & Custom Attributes
Lifecycle Manager
SnapCenter Plug-in for VMware vSphere
NetApp ONTAP tools
Cloud Provider Services
NSX
VMware Aria Operations Configuration
Skyline Health Diagnostics

2. Select the content library in which the item you like to clone
3. Right click on the item and click on Clone Item ..

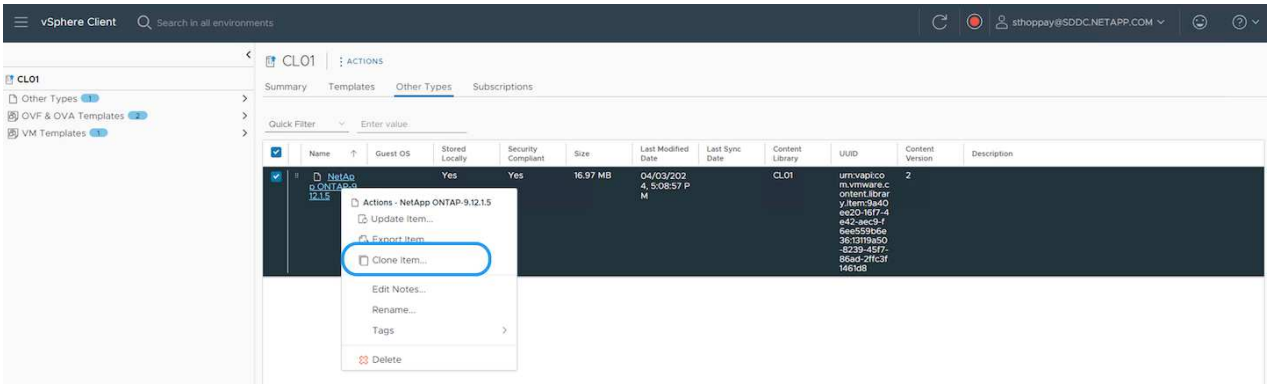

If using action menu, make sure correct target object is listed to perform action.

4. Select the target content library and click on OK.

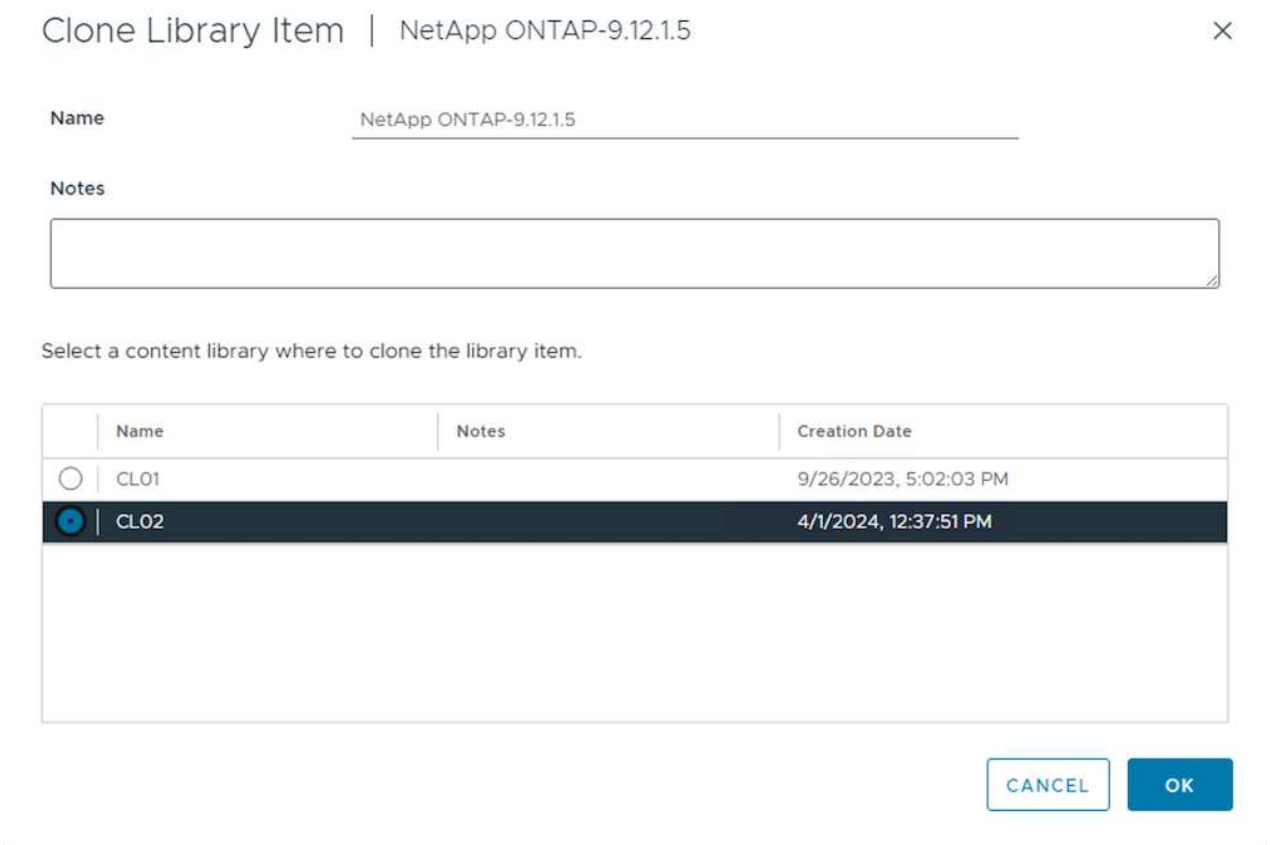

5. Validate the item is available on target content library.

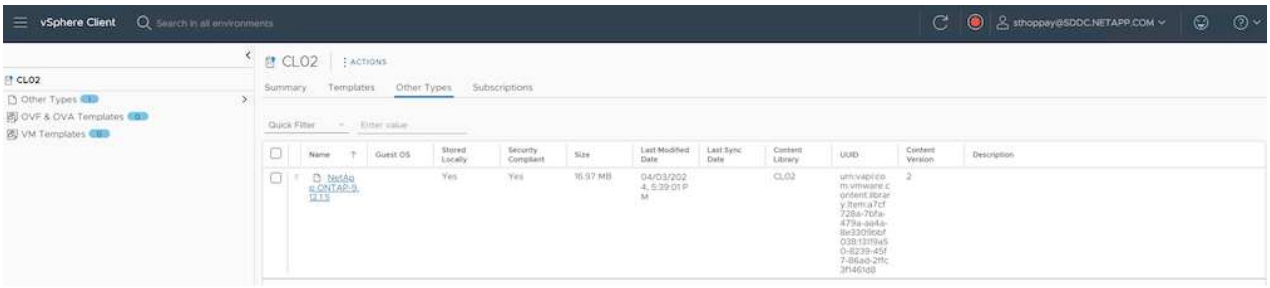

Here is the sample PowerCLI script to copy the content libary items from content library CL01 to CL02.

```
#Authenticate to vCenter Server(s)
$sourcevc = Connect-VIServer -server 'vcenter01.domain' -force
$targetvc = Connect-VIServer -server 'vcenter02.domain' -force

#Copy content library items from source vCenter content library CL01 to
target vCenter content library CL02.
Get-ContentLibaryItem -ContentLibary (Get-ContentLibary 'CL01' -Server
$sourcevc) | Where-Object { $_.ItemType -ne 'vm-template' } | Copy-
ContentLibaryItem -ContentLibrary (Get-ContentLibary 'CL02' -Server
$targetvc)
```

Adding VM as Templates in Content Library

1. In vSphere Web Client, select the VM and right click to choose Clone as Template in Library

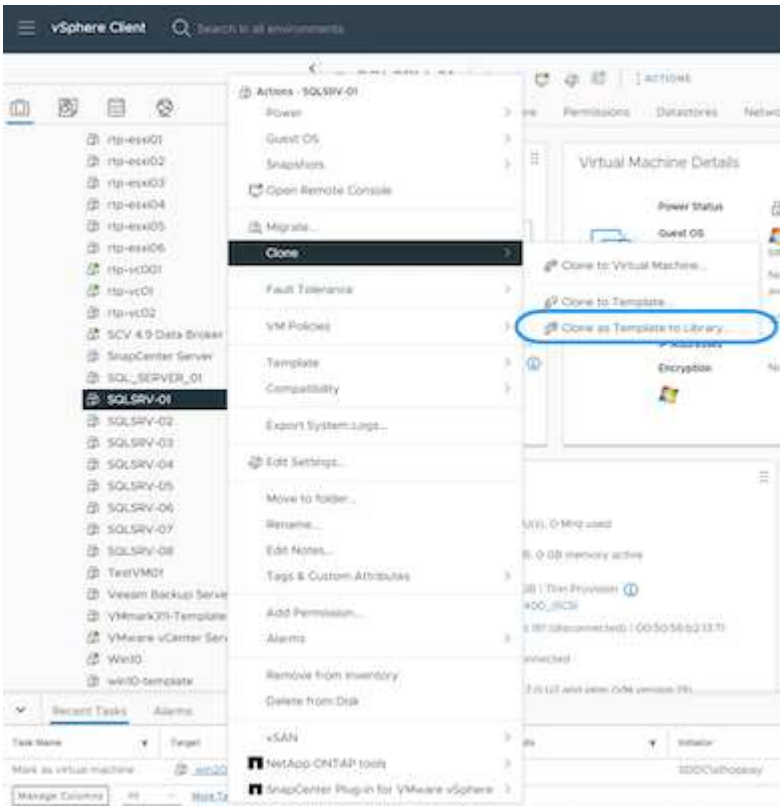

When VM template is selected to clone in libary, it can only store it as OVF & OVA template and not as VM template.

2. Confirm Template type is selected as VM Template and follow answering the wizard to complete the operation.

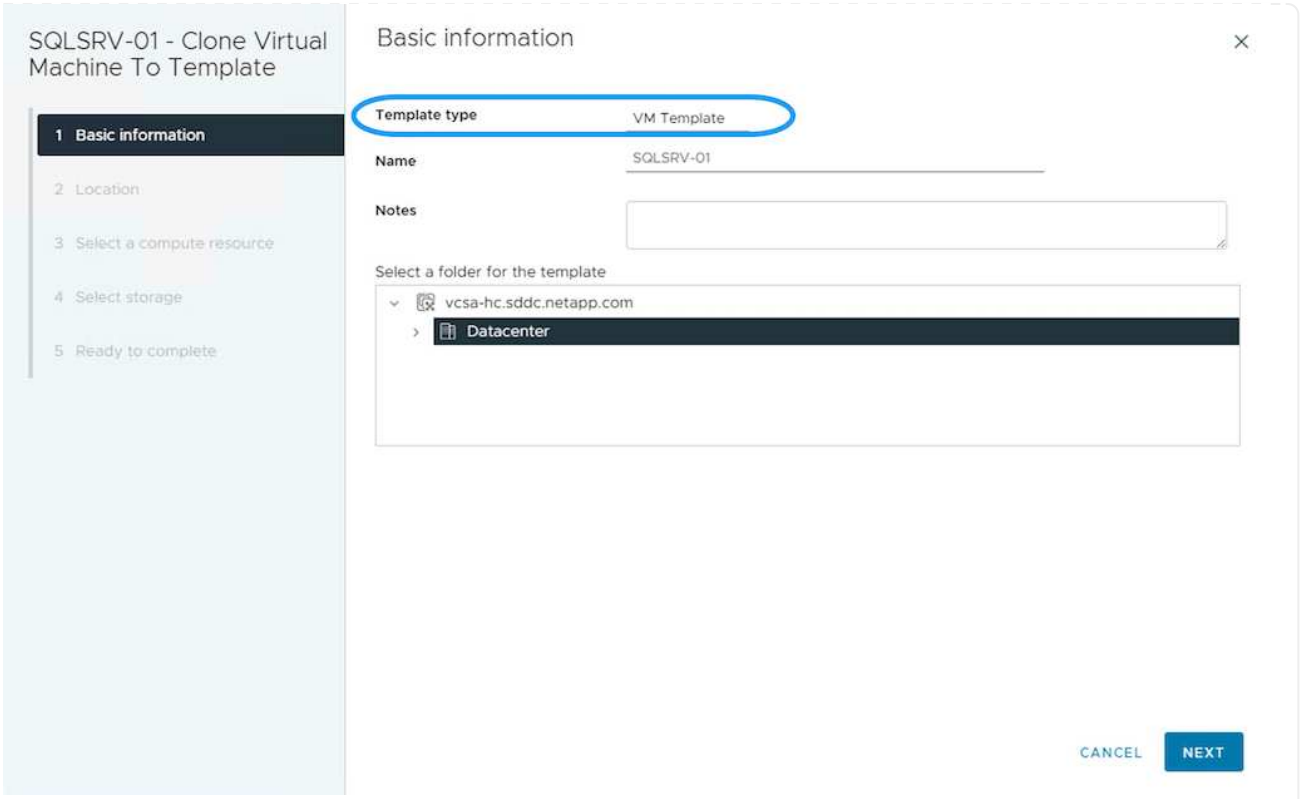

For additional details on VM templates on content library, check vSphere VM administration guide

Use Cases

Migration from third party storage systems (including vSAN) to ONTAP datastores.

- Based on where the ONTAP datastore is provisioned, pick the VM migration options from above.

Migration from previous version to latest version of vSphere.

- If in-place upgrade is not possible, can bring up new environment and use the migration options above.

In Cross vCenter migration option, import from target if export option is not available on source. For that procedure, check Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion

Migration to VCF Workload Domain.

- Migrate VMs from each vSphere Cluster to target workload domain.

> To allow network communication with existing VMs on other clusters on source vCenter, either extend NSX segment by adding the source vcenter vSphere hosts to transport zone or use L2 bridge on edge to allow L2 communication in VLAN. Check NSX documentation of Configure an Edge VM for Bridging

Additional Resources

- vSphere Virtual Machine Migration
- What's New in vSphere 8 for vMotion
- vSphere vMotion Resources
- Tier-0 Gateway Configurations in NSX Federation
- HCX 4.8 User Guide
- VMware Site Recovery Manager Documentation
- BlueXP disaster recovery for VMware

Migrate VMs to Amazon EC2 using FSxN

Migrate VMs to Amazon EC2 using FSxN: Overview

Organizations are accelerating their migrations to cloud computing solutions on AWS, taking advantage of services such as Amazon Elastic Compute Cloud (Amazon EC2) instances and Amazon FSx for NetApp ONTAP (FSx for ONTAP) to modernize their IT infrastructures, achieve cost savings, and improve operational efficiency. These AWS offerings enable migrations that optimize total cost of ownership (TCO) through consumption-based pricing models, enterprise storage features, providing the flexibility and scalability to meet evolving global business demands.

Overview

For enterprises deeply invested in VMware vSphere, migrating to AWS is a cost-effective option given the current market conditions, one that presents a unique opportunity.

As these organizations transition to AWS, they seek to capitalize on the cloud's agility and cost benefits while preserving familiar feature sets, particularly when it comes to storage. Maintaining seamless operations with familiar storage protocols—especially iSCSI—processes, tools, and skillsets is crucial when migrating workloads or setting up disaster recovery solutions.

Using the AWS managed storage service FSx for ONTAP for retaining the enterprise storage capabilities, that too coming from any third-party vendor storage from on-premises, enterprises can unlock the power of AWS while minimizing disruption and maximizing their future investments.

This technical report covers how to migrate on-premises VMware vSphere VMs to an Amazon EC2 instance with data disks placed on FSx for ONTAP iSCSI LUNs using the MigrateOps "data-mobility-as-code" functionality of Cirrus Migrate Cloud (CMC).

Solution requirements

There are a number of challenges that VMware customers are currently looking to solve. These organizations want to:

1. Leverage enterprise storage capabilities, such as thin provisioning, storage efficiency technologies, zero footprint clones, integrated backups, block-level replication, and tiering. This helps optimize migration efforts and future proof deployment on AWS from Day 1.
2. Optimize storage deployments currently on AWS that use Amazon EC2 instances by incorporating FSx for ONTAP and the cost-optimizing features it provides.
3. Reduce the total cost of ownership (TCO) of using Amazon EC2 instances with block storage solutions by rightsizing Amazon EC2 instances to meet the required IOPS and throughput parameters. With block storage, Amazon EC2 disk operations have a cap on bandwidth and I/O rates. File storage with FSx for ONTAP uses network bandwidth. In other words, FSx for ONTAP has no VM-level I/O limits.

Technical components overview

FSx for ONTAP concepts

Amazon FSx for NetApp ONTAP is a fully managed AWS storage service that provides NetApp® ONTAP® file systems with all the familiar ONTAP data management features, performance, and APIs on AWS. Its high-performance storage supports multiple protocols (NFS, SMB, iSCSI), providing a single service for workloads using Windows, Linux, and macOS EC2 instances.

Since FSx for ONTAP is an ONTAP file system, it brings a host of familiar NetApp features and services with it, including SnapMirror® data replication technology, thin clones, and NetApp Snapshot™ copies. By leveraging a low-cost capacity tier via data tiering, FSx for ONTAP is elastic and can reach a virtually unlimited scale. Plus, with signature NetApp storage efficiency technology, it reduces storage costs on AWS even further. For more, see Getting started with Amazon FSx for ONTAP.

File System

The central resource of FSx for ONTAP is its file system based on solid-state drive (SSD) storage. When provisioning an FSx for ONTAP file system, the user inputs a desired throughput and storage capacity, and selects an Amazon VPC where the file system will reside.

Users also have a choice between two built-in high-availability deployment models for the file system: Multi-Availability Zone (AZ) or single-AZ deployment. Each of these options offers its own level of durability and availability, which customers can select depending on their use case's business continuity requirements. Multi-AZ deployments consist of dual nodes that replicate seamlessly across two AZs. The more cost-optimized single-AZ deployment option structures the file system in two nodes split between two separate fault domains that both reside within a single AZ.
Storage Virtual Machines
Data in the FSx for ONTAP file system is accessed through a logical storage partition which is called a storage virtual machine (SVM). An SVM is actually its own file server equipped with its own data and admin access points. When accessing iSCSI LUNs on an FSx for ONTAP file system, the Amazon EC2 instance interfaces directly with the SVM using the SVM's iSCSI endpoint IP address.

While maintaining a single SVM in a cluster is possible, the option of running multiple SVMs in a cluster has a wide range of uses and benefits. Customers can determine the optimal number of SVMs to configure by considering their business needs, including their requirements for workload isolation.

Volumes

Data within an FSx for ONTAP SVM is stored and organized in structures known as volumes, which act as virtual containers. An individual volume can be configured with a single or multiple LUNs. The data stored in each volume consumes storage capacity in the file system. However, since FSx for ONTAP thinly provisions the volume, the volume only takes up storage capacity for the amount of data being stored.

The Cirrus Migrate Cloud MigrateOps concept

CMC is a transactable software-as-a-service (SaaS) offering from Cirrus Data Solutions, Inc. which is available via the AWS Marketplace. MigrateOps is a Data-Mobility-as-Code automation feature of CMC that allows you to declaratively manage your data mobility operations at scale using simple operation configurations in YAML. A MigrateOps configuration determines how you want your data mobility tasks to be executed. To learn more about MigrateOps, see About MigrateOps.

MigrateOps takes an automation-first approach, which is purpose-built to streamline the entire process, ensuring cloud-scale enterprise data mobility without operational disruptions. In addition to the already feature-rich functionalities that CMC offers for automation, MigrateOps further adds other automations that are often managed externally, such as:

- OS remediation
- Application cutover and approval scheduling
- Zero-downtime cluster migration
- Public/Private cloud platform integration
- Virtualization platform integration
- Enterprise storage management integration
- SAN (iSCSI) configuration

With the above tasks fully automated, all the tedious steps in preparing the on-prem source VM (such as adding AWS agents and tools), creation of destination FSx LUNs, setting up iSCSI and Multipath/MPIO at the AWS destination instance, and all the tasks of stopping/starting application services are eliminated by simply specifying parameters in a YAML file.

FSx for ONTAP is used to provide the data LUNs and rightsize the Amazon EC2 instance type, while providing all the features that organizations previously had in their on-premises environments. The MigrateOps feature of CMC will be used to automate all the steps involved, including provisioning mapped iSCSI LUNs, turning this into a predictable, declarative operation.

Note: CMC requires a very thin agent to be installed on the source and destination virtual machine instances to ensure secure data transfer from the storage source storage to FSx for ONTAP.

Benefits of using Amazon FSx for NetApp ONTAP with EC2 instances

FSx for ONTAP storage for Amazon EC2 instances provides several benefits:

- High throughput and low latency storage that provide consistent high performance for the most demanding workloads
- Intelligent NVMe caching improves performance
- Adjustable capacity, throughput, and IOPs can be changed on the fly and quickly adapt to changing storage demands
- Block-based data replication from on-premises ONTAP storage to AWS

- Multi-protocol accessibility, including for iSCSI, which is widely used in on-premises VMware deployments
- NetApp Snapshot™ technology and DR orchestrated by SnapMirror prevent data loss and speed up recovery
- Storage efficiency features that reduce storage footprint and costs, including thin provisioning, data deduplication, compression, and compaction
- Efficient replication reduces the time it takes to create backups from hours to just minutes, optimizing RTO
- Granular options for file back up and restores using NetApp SnapCenter®

Deploying Amazon EC2 instances with FSx ONTAP as the iSCSI-based storage layer delivers high performance, mission-critical data management features, and cost-reducing storage efficiency features that can transform your deployment on AWS.

Running a Flash Cache, multiple iSCSI sessions, and leveraging a working set size of 5%, it's possible for FSx for ONTAP to deliver IOPS of ~350K, providing performance levels to meet even the most intensive workloads.

Since only network bandwidth limits are applied against FSx for ONTAP, not block storage bandwidth limits, users can leverage small Amazon EC2 instance types while achieving the same performance rates as much larger instance types. Using such small instance types also keeps compute costs low, optimizing TCO.

The ability of FSx for ONTAP to serve multiple protocols is another advantage, one that helps standardize a single AWS storage service for a wide range of existing data and file services requirements.
For enterprises deeply invested in VMware vSphere, migrating to AWS is a cost-effective option given the current market conditions, one that presents a unique opportunity.

Migrate VMs to Amazon EC2 using FSxN: Architecture and Pre-Requisites

This article shows the high-level architecture and deployment pre-requisites for completing the migration.

High level architecture

The diagram below illustrates the high-level architecture of migrating Virtual Machine Disk (VMDK) data on VMware to AWS using CMC MigrateOps:

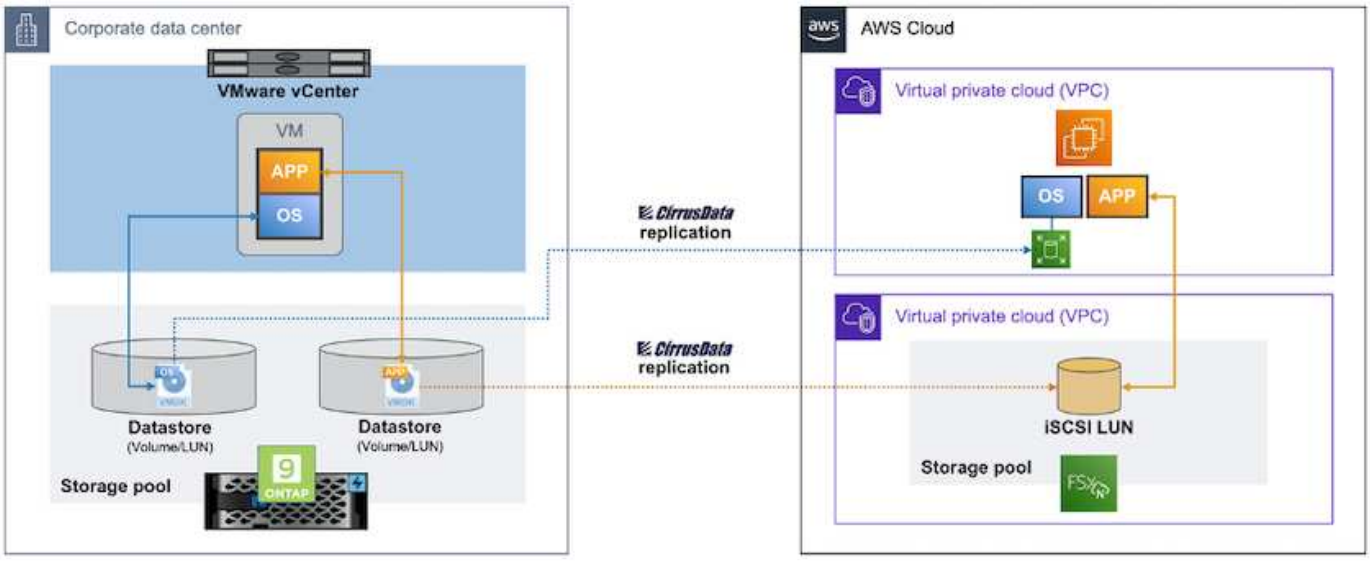

How to migrate your VMware VMs to AWS using Amazon EC2 and FSx for ONTAP iSCSI

Prerequisites

Before starting the walkthrough steps, make sure the following prerequisites are met:

On AWS

- An AWS account. This includes permissions for subnets, VPC setup, routing tables, security rule migration, security groups, and other requirements for networking such as load balancing. As with any migration, the most effort and consideration should go into networking.
- Appropriate IAM roles that allow you to provision both FSx for ONTAP and Amazon EC2 instances.
- Route tables and security groups are allowed to communicate with FSx for ONTAP.
- Add an inbound rule to the appropriate security group (see below for more details) to allow for secure data transfer from your on-premises data center to AWS.
- A valid DNS that can resolve public internet domain names.
- Check that your DNS resolution is functional and allows you to resolve host names.
- For optimal performance and rightsizing, use performance data from your source environment to rightsize your FSx for ONTAP storage.
- Each MigrateOps session uses one EIP, hence the quota for EIP should be increased for more parallelism. Keep in mind, the default EIP quota is 5.
- (If Active Directory-based workloads are being migrated) A Windows Active Directory domain on Amazon EC2.

For Cirrus Migrate Cloud

- A Cirrus Data Cloud account at cloud.cirrusdata.com must be created before using CMC. Outbound communication with the CDN, Cirrus Data endpoints, and software repository via HTTPS must be allowed.
- Allow communication (outbound) with Cirrus Data Cloud services via HTTPS protocol (Port 443).
- For a host to be managed by the CMC project, the deployed CMC software must initiate a one-way outbound TCP connection to Cirrus Data Cloud.
- Allow TCP protocol, Port 443 access to portal-gateway.cloud.cirrusdata.com which is currently at 208.67.222.222.
- Allow HTTP POST requests (via HTTPS connection) with binary data payload (application/octet-stream). This is similar to a file upload.
- Ensure that portal-gateway.cloud.cirrusdata.com is resolvable by your DNS (or via OS host file).
- If you have strict rules for prohibiting product instances to make outbound connections, the “Management Relay” feature of CMC can be used where the outbound 443 connection is from a single, secured non-production host.

Note: No storage data is ever sent to the Cirrus Data Cloud endpoint. Only management metadata is sent, and this can be optionally masked so that no real host name, volume name, network IP are included.

For migrating data from on-premises storage repositories to AWS, MigrateOps automates the management of a Host-to-Host (H2H) connection. These are optimized, one-way, TCP-based network connections that CMC uses to facilitate remote migration. This process features always-on compression and encryption that can reduce the amount of traffic by up to eight times, depending on the nature of the data.

Note: CMC is designed so that no production data / I/O leaves the production network during the entire

migration phase. As a result, direct connectivity between the source and destination host is required.

Migrate VMs to Amazon EC2 using FSxN: Deployment Guide

This article describes the deployment procedure for this migration solutions.

Configure FSx for ONTAP and Cirrus Data for migration operations

This step-by-step deployment guide shows how to add FSx for ONTAP volume to a VPC. Since these steps are sequential in nature, make sure they are covered in order.

For the purposes of this demonstration, “DRaaSDemo” is the name of the file system created.

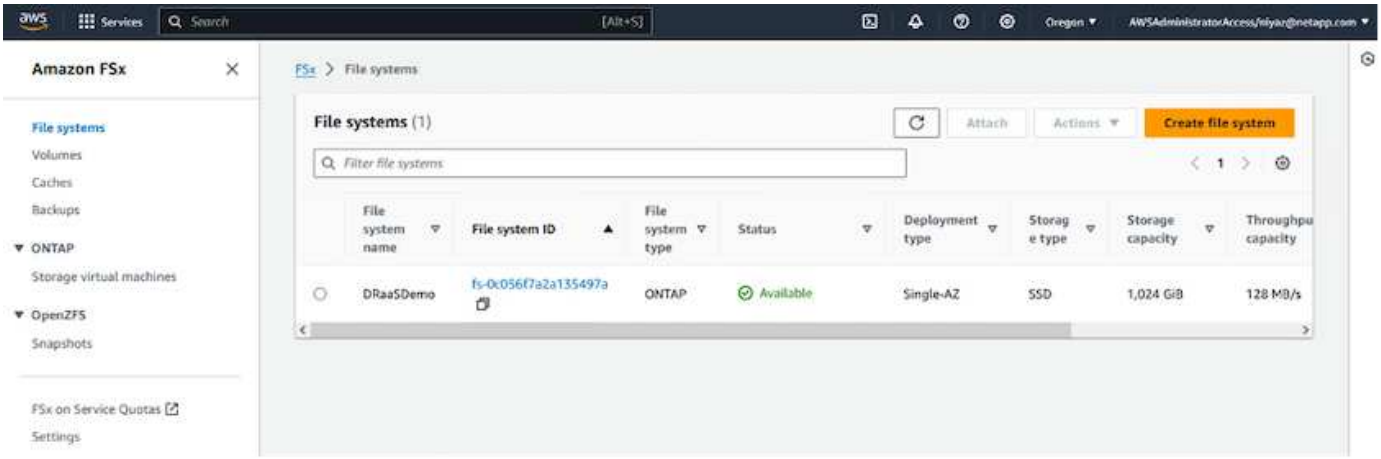

Once your AWS VPC is configured and FSx for ONTAP is provisioned based on your performance requirements, log in to cloud.cirrusdata.com and create a new project or access an existing project.

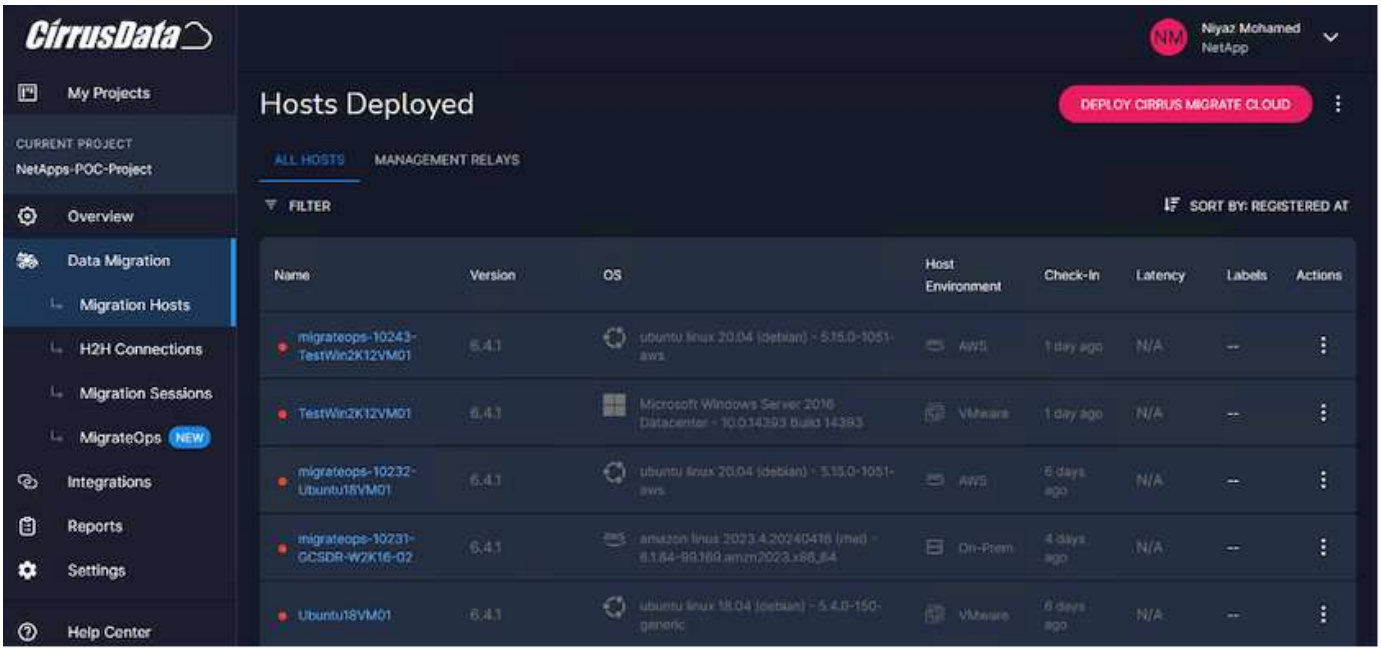

Before creating the recipe for MigrationOps, AWS Cloud should be added as an integration. CMC provides built-in integration with FSx for ONTAP and AWS. The integration for FSx for ONTAP provides the following automated functionalities:

Prepare your FSx for ONTAP file system:

- Create new volumes and LUNs that match the source volumes

Note: A destination disk in the FSx for ONTAP FS model is a “LUN” that is created on a “Volume” that has enough capacity to contain the LUN plus a reasonable amount of overhead for facilitating snapshots and meta-data. The CMC automation takes care of all these details to create the appropriate Volume and the LUN with optional user-defined parameters.

- Create Host entity (called iGroups in FSx) with the Host Initiator IQN
- Map newly created volumes to appropriate host entities using mappings
- Create all other necessary configurations

Prepare Production Host for iSCSI connection:

- If necessary, install and configure iSCSI feature and set up Initiator.
- If necessary, install and configure multipath (MPIO for Windows) with proper vendor identifiers.
- Adjust system settings, if necessary, according to vendor best practices, e.g. with udev settings on Linux.
- Create and manage iSCSI connections such as persistent/favorite iSCSI targets on Windows.

To configure CMC Integration for FSx for ONTAP and AWS, perform the following steps:

1. Log in to the Cirrus Data Cloud portal.
2. Go to the Project for which you want to enable the integration.
3. Navigate to Integrations → Goodies.
4. Scroll to find FSx for NetApp ONTAP and click ADD INTEGRATION.

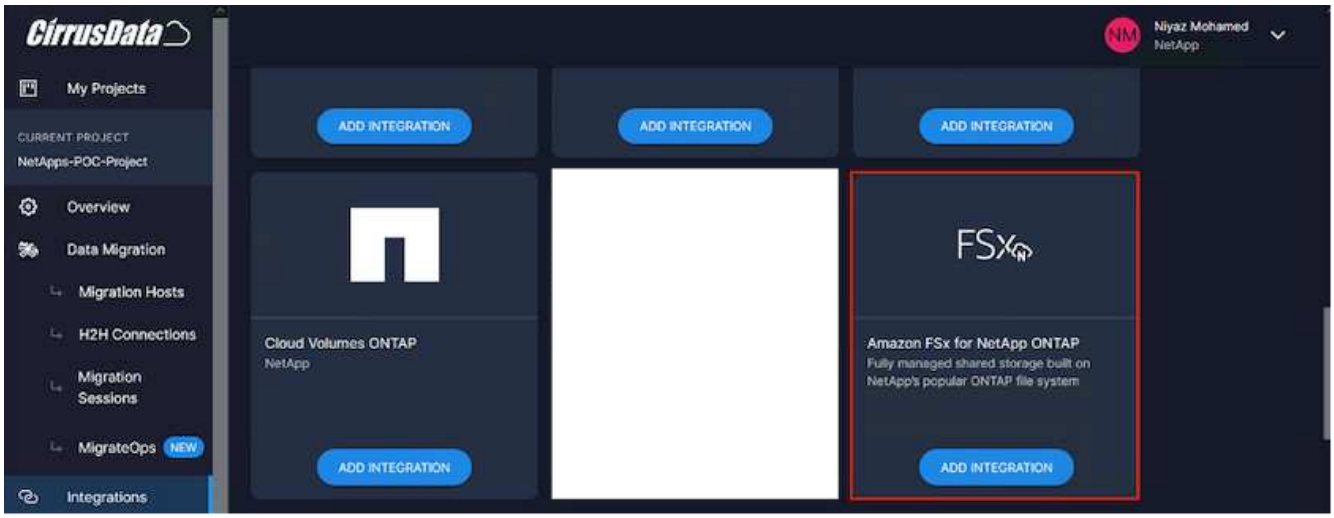

5. Provide a descriptive name (strictly for display purposes) and add the appropriate credentials.

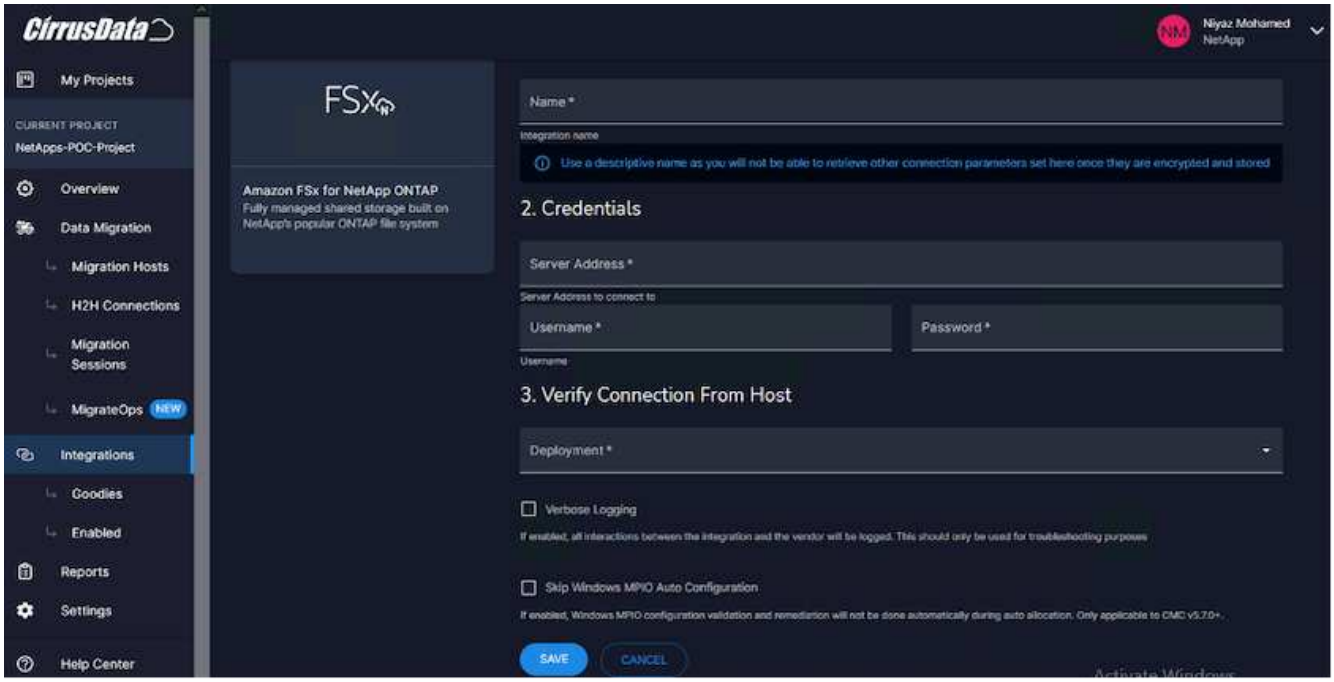

6. Once the integration is created, during the creation of a new migration session, select Auto Allocate Destination Volumes to automatically allocate new volumes on FSx for ONTAP.

   Note: New LUNs will be created with the same size as the source volume's size, unless "Migrate to Smaller Volumes" is enabled for the migration.

   Note: If a host entity (iGroup) doesn't already exist, a new one will be created. All host iSCSI Initiator IQNs will be added to that new host entity.

   Note: If an existing host entity with any of the iSCSI initiators already exists, it will be reused.

7. Once done, add the integration for AWS, following the steps on the screen.

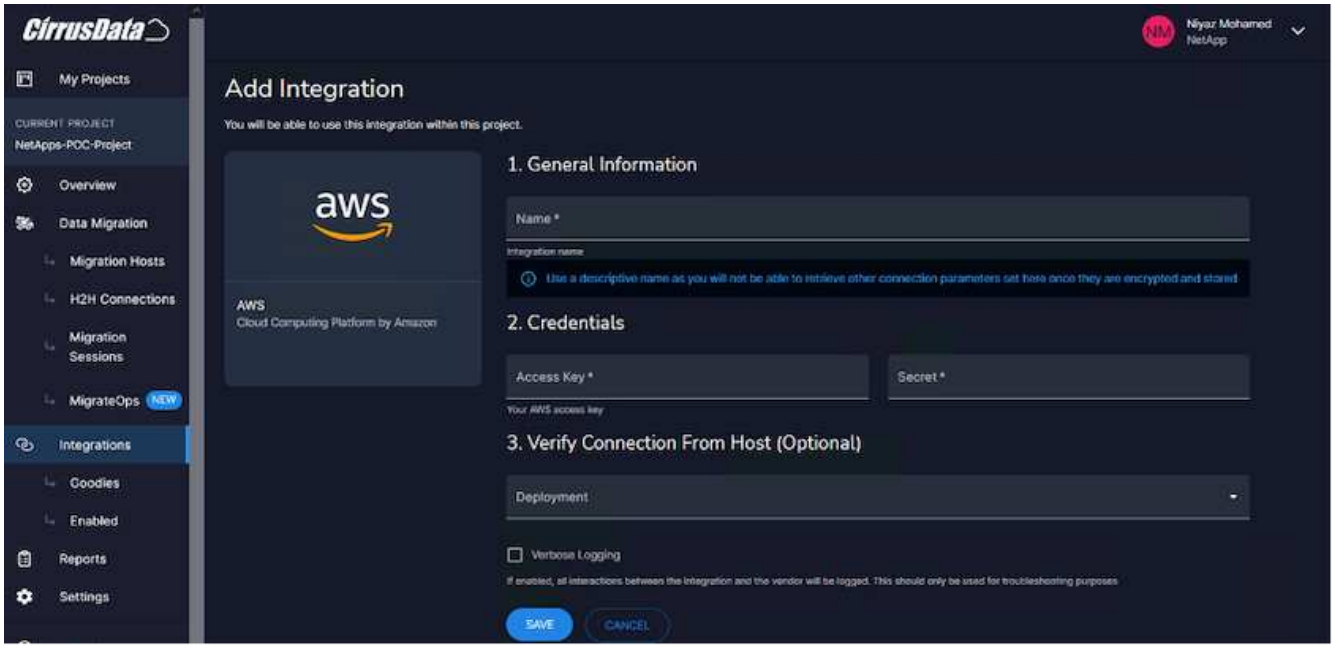

   Note: This integration is used while migrating virtual machines from on-premises storage to AWS along with FSx for ONTAP integration.

Note: Use management relays to communicate with Cirrus Data Cloud if there is no direct outbound connection for production instances to be migrated.

With Integrations added, it's time to register hosts with the Project. Let's cover this with an example scenario.

Host registration scenario

Guest VMware VMs residing on vCenter in on-premises data center:

- Windows 2016 running with SQL Server with three VMDKs including OS and data disks. It is running an active database. The database is located on a data volume backed by two VMDKs.

Note: Since the source is a VMware environment and VMDKs are used, the Windows iSCSI Initiator software is not currently configured on this guest VM. To connect to our destination storage via iSCSI, both iSCSI and MPIO will have to be installed and configured. Cirrus Data Cloud integration will perform this installation automatically during the process.

Note: The Integration configured in the previous section automates the configuration of the new destination storage in creating the new disks, setting up the host entities and their IQNs, and even remediation of the application VM (host) for iSCSI and multipath configurations.

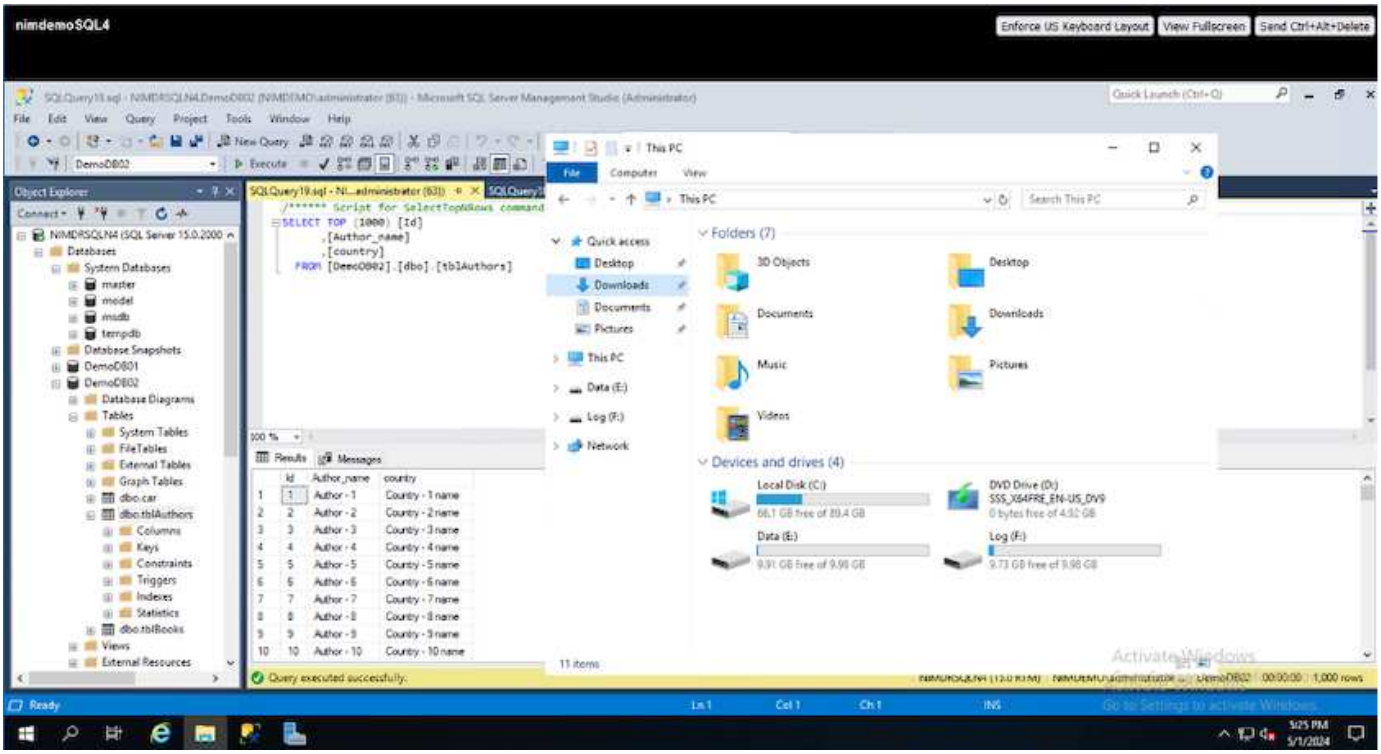

This demonstration will migrate the application VMDKs from each VM to an automatically provisioned and mapped iSCSI volume from FSx for ONTAP. The OS VMDK in this case will be migrated to an Amazon EBS volume as Amazon EC2 instances support this Amazon EBS only as the boot disk.

Note: The scale factor with this migration approach is the network bandwidth and the pipe connecting on-premises to AWS VPC. Since each VM has 1:1 host session configured, the overall migration performance depends on two factors:

- Network bandwidth
- Target instance type and ENI bandwidth

The migration steps are as follows:

1. Install CMC agent on each host (Windows and Linux) designated for the migration wave. This can be performed by executing a one-line installation command.

   To do this, access Data Migration > Migration Hosts > Click on “Deploy Cirrus Migrate Cloud” and click to select “Windows”.

   Then, copy the `iex` command to the host and run it using PowerShell. Once the deployment of the agent is successful, the host will be added to the Project under “Migration hosts”.

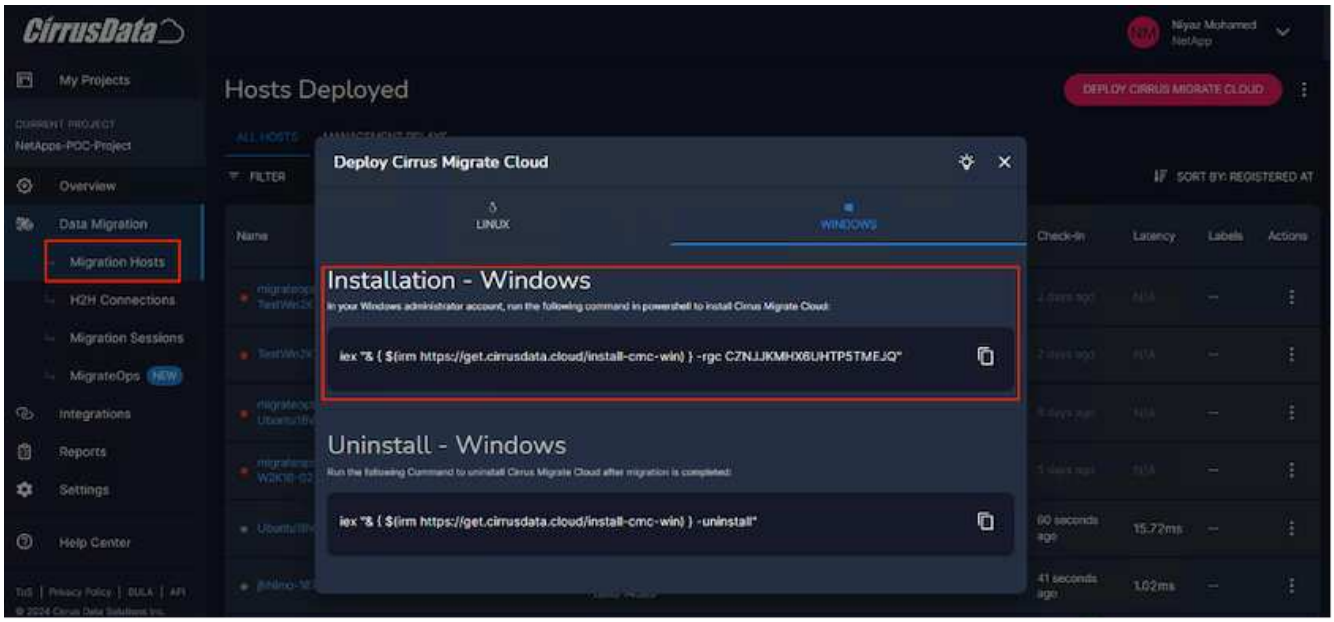

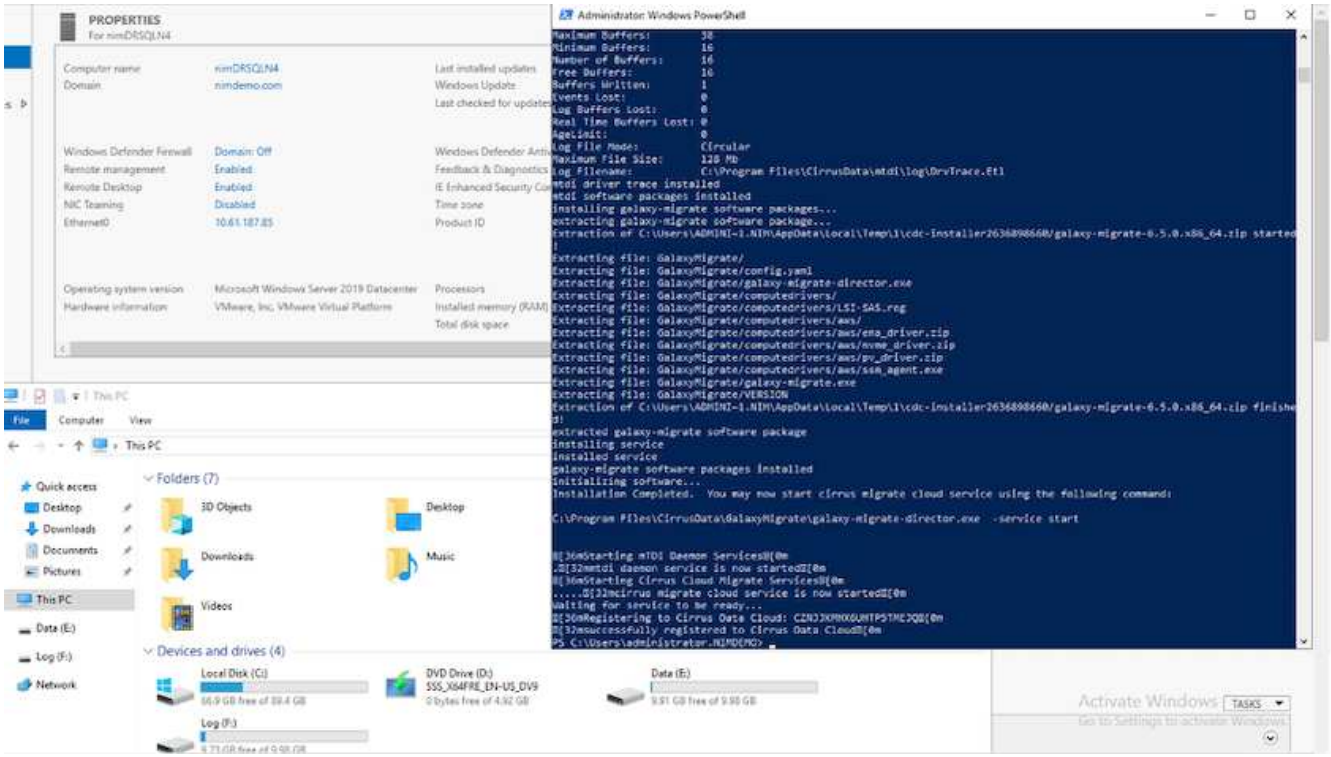

2. Prepare the YAML for each virtual machine.

Note: It is a vital step to have a YAML for each VM that specifies the necessary recipe or blueprint for the migration task.

The YAML provides the operation name, notes (description) along with the recipe name as `MIGRATEOPS_AWS_COMPUTE`, the host name (`system_name`) and integration name (`integration_name`) and the source and destination configuration. Custom scripts can be specified as a before and after cutover action.

```
operations:
    -   name: Win2016 SQL server to AWS
        notes: Migrate OS to AWS with EBS and Data to FSx for ONTAP
        recipe: MIGRATEOPS_AWS_COMPUTE
        config:
            system_name: Win2016-123
            integration_name: NimAWShybrid
            migrateops_aws_compute:
                region: us-west-2
                compute:
                    instance_type: t3.medium
                    availability_zone: us-west-2b
                network:
                    vpc_id: vpc-05596abe79cb653b7
                    subnet_id: subnet-070aeb9d6b1b804dd
                    security_group_names:
                        - default
                destination:
                    default_volume_params:
                        volume_type: GP2
                    iscsi_data_storage:
                        integration_name: DemoDRaaS
                        default_volume_params:
                            netapp:
                                qos_policy_name: ""
                migration:
                    session_description: Migrate OS to AWS with EBS and
Data to FSx for ONTAP
                    qos_level: MODERATE
                cutover:
                    stop_applications:
                        - os_shell:
                            script:
                                - stop-service -name 'MSSQLSERVER'
-Force
                                - Start-Sleep -Seconds 5
                                - Set-Service -Name 'MSSQLSERVER'
-StartupType Disabled
                                - write-output "SQL service stopped
```

```
and disabled"

                        - storage_unmount:
                              mountpoint: e
                        - storage_unmount:
                              mountpoint: f
                    after_cutover:
                        - os_shell:
                              script:
                                  - stop-service -name 'MSSQLSERVER'
-Force
                                  - write-output "Waiting 90 seconds to
mount disks..." > log.txt
                                  - Start-Sleep -Seconds 90
                                  - write-output "Now re-mounting disks
E and F for SQL..." >>log.txt
                        - storage_unmount:
                              mountpoint: e
                        - storage_unmount:
                              mountpoint: f
                        - storage_mount_all: {}
                        - os_shell:
                              script:
                                  - write-output "Waiting 60 seconds to
restart SQL Services..." >>log.txt
                                  - Start-Sleep -Seconds 60
                                  - stop-service -name 'MSSQLSERVER'
-Force
                                  - Start-Sleep -Seconds 3
                                  - write-output "Start SQL Services..."
>>log.txt
                                  - Set-Service -Name 'MSSQLSERVER'
-StartupType Automatic
                                  - start-service -name 'MSSQLSERVER'
                                  - write-output "SQL started" >>log.txt
```

3. Once the YAMLs are in place, create MigrateOps configuration. To do this, go to Data Migration > MigrateOps, click on “Start New Operation” and enter the configuration in valid YAML format.
4. Click “Create operation”.

   Note: To achieve parallelism, each host needs to have a YAML file specified and configured.

5. Unless the `scheduled_start_time` field is specified in the configuration, the operation will start immediately.
6. The operation will now execute and proceed. From the Cirrus Data Cloud UI, you can monitor the progress with detailed messages. These steps automatically include tasks that are normally done manually, such as performing auto allocation and creating migration sessions.

Note: During the host-to-host migration, an additional security group with a rule allowing Inbound 4996 port will be created, which will allow the required port for communication and it will be automatically deleted once the synchronization is complete.

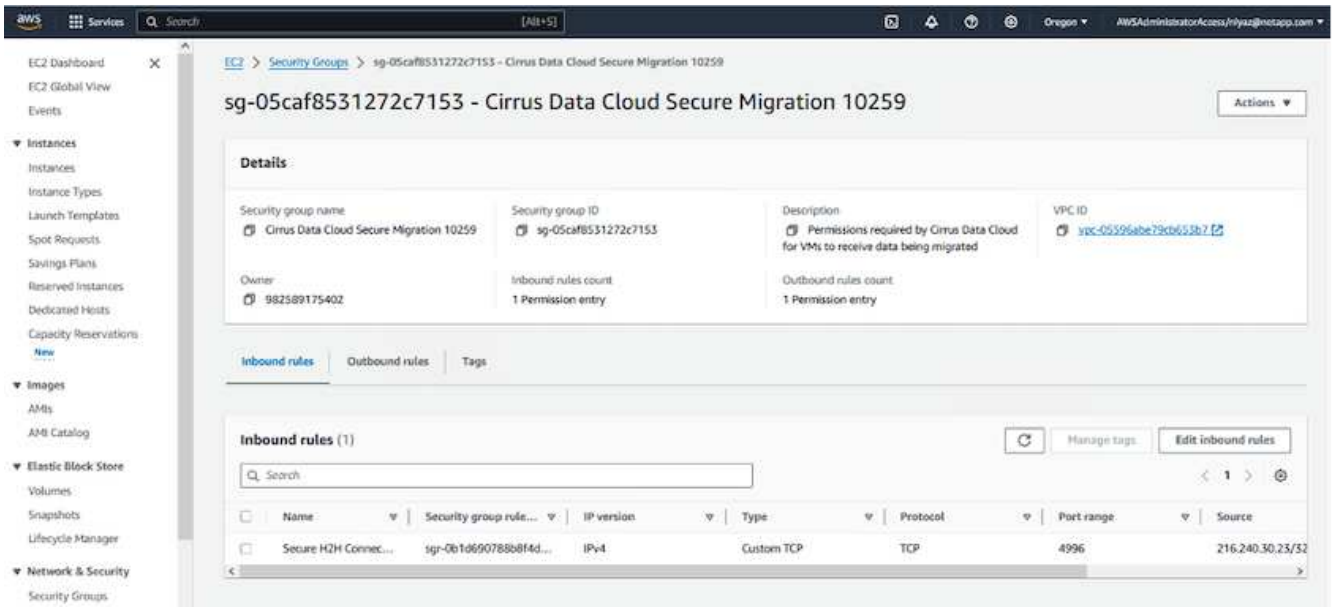

7. While this migration session is synchronizing, there is a future step in phase 3 (cutover) with the label “Approval Required.” In a MigrateOps recipe, critical tasks (such as migration cutovers) require user approval before they can be executed. Project Operators or Administrators can approve these tasks from the UI. A future approval window can also be created.

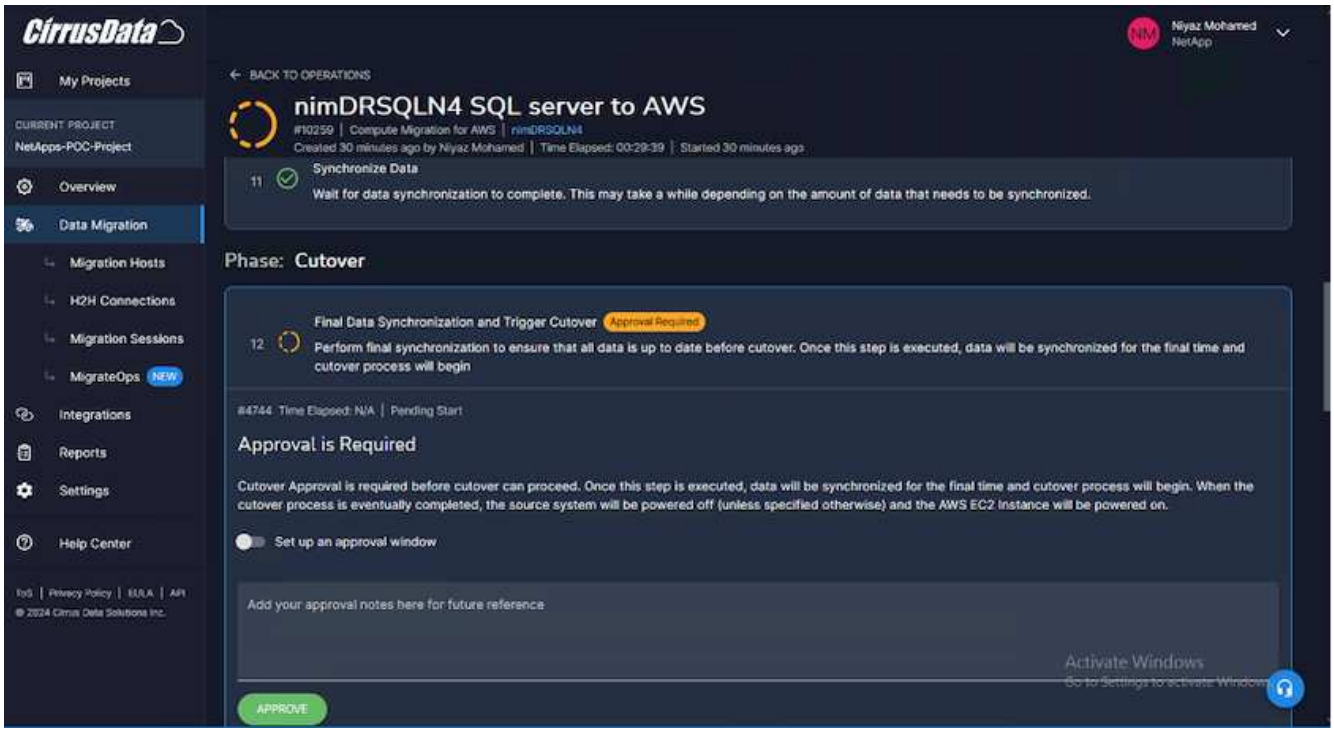

8. Once approved, the MigrateOps operation continues with the cutover.
9. After a brief moment, the operation will be completed.

Note: With the help of Cirrus Data cMotion™ technology, the destination storage has been kept up-to-date with all the latest changes. Therefore, after approval is given, this entire final cutover process will take a very short time—less than a minute—to complete.

Post-migration verification

Let's look at the migrated Amazon EC2 instance running the Windows Server OS and the following steps that have completed:

1. Windows SQL Services are now started.

2. The database is back online and is using storage from the iSCSI Multipath device.
3. All new database records added during migration can be found in the newly migrated database.
4. The old storage is now offline.

Note: With just one click to submit the data mobility operation as code, and a click to approve the cutover, the VM has successfully migrated from on-premises VMware to an Amazon EC2 instance using FSx for ONTAP and its iSCSI capabilities.

Note: Due to AWS API limitation, the converted VMs would be shown as “Ubuntu.” This is strictly a display issue and does not affect functionality of the migrated instance. An upcoming release will address this issue.

Note: The migrated Amazon EC2 instances can be accessed using the credentials that were used on the on-premises side.

Migrate VMs to Amazon EC2 using FSxN: Other Possibilities and Conclusion

This article highlight other possibilities for this migration solution as well as concluding the topic.

Other possibilities

The same approach can be extended to migrate VMs using in-guest storage on on-premises VMs. The OS VMDK can be migrated using CMC and the in-guest iSCSI LUNs can be replicated using SnapMirror. The process requires breaking the mirror and attaching the LUN to the newly migrated Amazon EC2 instance, as depicted in the diagram below.

Conclusion

This document has provided a complete walkthrough of using the MigrateOps feature of CMC to migrate data stored in on-premises VMware repositories to AWS using Amazon EC2 instances and FSx for ONTAP.

The following video demonstrates the migration process from start to finish:

Migrate VMware VMs to Amazon EC2

To check out the GUI and basic Amazon EBS to FSx for ONTAP local migration, please watch this five-minute demo video:

Migrating to any storage in scale with Cirrus Migrate Cloud

NetApp Hybrid Multicloud with VMware Solutions

VMware Hybrid Multicloud 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.

Introduction to automation for ONTAP and vSphere

This page describes the benefits of automating base ONTAP functionality in a VMware vSphere environment.

VMware automation

Automation has been an integral part of managing VMware environments since the first days of VMware ESX. The ability to deploy infrastructure as code and extend practices to private cloud operations helps to alleviate concerns surrounding scale, flexibility, self-provisioning, and efficiency.

Automation can be organized into the following categories:

- **Virtual infrastructure deployment**

- **Guest machine operations**
- **Cloud operations**

There are many options available to administrators with respect to automating their infrastructure. Whether through using native vSphere features such as Host Profiles or Customization Specifications for virtual machines to available APIs on the VMware software components, operating systems, and NetApp storage systems; there is significant documentation and guidance available.

Data ONTAP 8.0.1 and later supports certain VMware vSphere APIs for Array Integration (VAAI) features when the ESX host is running ESX 4.1 or later. VAAI is a set of APIs that enable communication between VMware vSphere ESXi hosts and storage devices. These features help offload operations from the ESX host to the storage system and increase network throughput. The ESX host enables the features automatically in the correct environment. You can determine the extent to which your system is using VAAI features by checking the statistics contained in the VAAI counters.

The most common starting point for automating the deployment of a VMware environment is provisioning block or file-based datastores. It is important to map out the requirements of the actual tasks prior to developing the corresponding automation.

For more information concerning the automation of VMware environments, see the following resources:

- The NetApp Pub. NetApp configuration management and automation.
- The Ansible Galaxy Community for VMware. A collection of Ansible resources for VMware.
- VMware {code} Resources. Resources needed to design solutions for the software-defined data center, including forums, design standards, sample code, and developer tools.

vSphere traditional block storage provisioning with ONTAP

VMware vSphere supports the following VMFS datastore options with ONTAP SAN protocol support indicated.

VMFS datastore options	ONTAP SAN protocol support
Fibre Channel (FC)	yes
Fibre Channel over Ethernet (FCoE)	yes
iSCSI	yes
iSCSI Extensions for RDMA (iSER)	no
NVMe over Fabric with FC (NVMe/FC)	yes
NVMe over Fabric with RDMA over Converged Ethernet (NVMe/RoCE)	no

If iSER or NVMe/RoCE VMFS is required, check SANtricity-based storage systems.

vSphere VMFS datastore - Fibre Channel storage backend with ONTAP

This section covers the creation of a VMFS datastore with ONTAP Fibre Channel (FC) storage.

What you need

- The basic skills necessary to manage a vSphere environment and ONTAP
- An ONTAP storage system (FAS/AFF/CVO/ONTAP Select/ASA) running ONTAP 9.8 or later
- ONTAP credentials (SVM name, userID, and password)
- ONTAP WWPN of host, target, and SVM and LUN information
- The completed FC configuration worksheet
- vCenter Server credentials
- vSphere host(s) information
  - vSphere 7.0 or later
- Fabric switch(es)
  - With connected ONTAP FC data ports and vSphere hosts
  - With the N_port ID virtualization (NPIV) feature enabled
  - Create a single initiator single target zone.
    - Create one zone for each initiator (single initiator zone).
    - For each zone, include a target that is the ONTAP FC logical interface (WWPN) for the SVMs. There should be at least two logical interfaces per node per SVM. Do not use the WWPN of the physical ports.
- An ONTAP Tool for VMware vSphere deployed, configured, and ready to consume.

Provisioning a VMFS datastore

To provision a VMFS datastore, complete the following steps:

1. Check compatability with the Interoperability Matrix Tool (IMT)
2. Verify that the FCP Configuration is supported.

ONTAP tasks

1. Verify that you have an ONTAP license for FCP.
   a. Use the `system license show` command to check that FCP is listed.
   b. Use `licen se add -license-code <license code>` to add the license.
2. Make sure that the FCP protocol is enabled on the SVM.
   a. Verify the FCP on an existing SVM.
   b. Configure the FCP on an existing SVM.
   c. Create s new SVM with the FCP.
3. Make sure that FCP logical interfaces are available on an SVM.
   a. Use `Network Interface show` to verify the FCP adapter.
   b. When an SVM is created with the GUI, logical interfaces are a part of that process.
   c. To rename network interfaces, use `Network Interface modify`.
4. Create and Map a LUN. Skip this step if you are using ONTAP tools for VMware vSphere.

VMware vSphere tasks

1. Verfiy that HBA drivers are installed. VMware supported HBAs have drivers deployed out of the box and should be visible in the Storage Adapter Information.
2. Provision a VMFS datastore with ONTAP Tools.

vSphere VMFS Datastore - Fibre Channel over Ethernet storage protocol with ONTAP

This section covers the creation of a VMFS datastore with the Fibre Channel over Ethernet (FCoE) transport protocol to ONTAP storage.

What you need

- The basic skills necessary to manage a vSphere environment and ONTAP
- An ONTAP storage system (FAS/AFF/CVO/ONTAP Select) running ONTAP 9.8 or later
- ONTAP credentials (SVM name, userID, and password)
- A supported FCoE combination
- A completed configuration worksheet
- vCenter Server credentials
- vSphere host(s) information
  - vSphere 7.0 or later
- Fabric switch(es)
  - With either ONTAP FC data ports or vSphere hosts connected
  - With the N_port ID virtualization (NPIV) feature enabled
  - Create a single initiator single target zone.
  - FC/FCoE zoning configured
- Network switch(es)
  - FCoE support
  - DCB support
  - Jumbo frames for FCoE
- ONTAP Tool for VMware vSphere deployed, configured, and ready to consume

Provision a VMFS datastore

- Check compatibility with the Interoperability Matrix Tool (IMT).
- Verify that the FCoE configuration is supported.

ONTAP tasks

1. Verify the ONTAP license for FCP.
   a. Use the `system license show` command to verify that the FCP is listed.
   b. Use `license add -license-code <license code>` to add a license.
2. Verify that the FCP protocol is enabled on the SVM.

a. Verify the FCP on an existing SVM.
b. Configure the FCP on an existing SVM.
c. Create a new SVM with the FCP.

3. Verify that FCP logical interfaces are available on the SVM.
   a. Use `Network Interface show` to verify the FCP adapter.
   b. When the SVM is created with the GUI, logical interfaces are a part of that process.
   c. To rename the network interface, use `Network Interface modify`.
4. Create and map a LUN; skip this step if you are using ONTAP tools for VMware vSphere.

VMware vSphere tasks

1. Verify that HBA drivers are installed. VMware-supported HBAs have drivers deployed out of the box and should be visible in the storage adapter information.
2. Provision a VMFS datastore with ONTAP Tools.

vSphere VMFS Datastore - iSCSI Storage backend with ONTAP

This section covers the creation of a VMFS datastore with ONTAP iSCSI storage.

For automated provisioning, use the following script: Ansible Playbook.

What you need

- The basic skills necessary to manage a vSphere environment and ONTAP.
- An ONTAP storage system (FAS/AFF/CVO/ONTAP Select/ASA) running ONTAP 9.8 or later
- ONTAP credentials (SVM name, userID, and password)
- ONTAP network port, SVM, and LUN information for iSCSI
- A completed iSCSI configuration worksheet
- vCenter Server credentials
- vSphere host(s) information
  - vSphere 7.0 or later
- iSCSI VMKernel adapter IP information
- Network switch(es)
  - With ONTAP system network data ports and connected vSphere hosts
  - VLAN(s) configured for iSCSI
  - (Optional) link aggregation configured for ONTAP network data ports
- ONTAP Tool for VMware vSphere deployed, configured, and ready to consume

Steps

1. Check compatibility with the Interoperability Matrix Tool (IMT).
2. Verify that the iSCSI configuration is supported.
3. Complete the following ONTAP and vSphere tasks.

ONTAP tasks

1. Verify the ONTAP license for iSCSI.
   a. Use the `system license show` command to check if iSCSI is listed.
   b. Use `license add -license-code <license code>` to add the license.
2. Verify that the iSCSI protocol is enabled on the SVM.
3. Verify that iSCSI network logical interfaces are available on the SVM.

   When an SVM is created using the GUI, iSCSI network interfaces are also created.

4. Use the `Network interface` command to view or make changes to the network interface.

   Two iSCSI network interfaces per node are recommended.

5. Create an iSCSI network interface. You can use the default-data-blocks service policy.
6. Verify that the data-iscsi service is included in the service policy. You can use `network interface service-policy show` to verify.
7. Verify that jumbo frames are enabled.
8. Create and map the LUN. Skip this step if you are using ONTAP tools for VMware vSphere. Repeat this step for each LUN.

VMware vSphere tasks

1. Verify that at least one NIC is available for the iSCSI VLAN. Two NICs are preferred for better performance and fault tolerance.
2. Identify the number of physical NICs available on the vSphere host.
3. Configure the iSCSI initiator. A typical use case is a software iSCSI initiator.
4. Verify that the TCPIP stack for iSCSI is available.
5. Verify that iSCSI portgroups are available.
   - We typically use a single virtual switch with multiple uplink ports.
   - Use 1:1 adapter mapping.
6. Verify that iSCSI VMKernel adapters are enabled to match the number of NICs and that IPs are assigned.
7. Bind the iSCSI software adapter to the iSCSI VMKernel adapter(s).
8. Provision the VMFS datastore with ONTAP Tools. Repeat this step for all datastores.
9. Verify hardware acceleration support.

What's next?

After these the tasks are completed, the VMFS datastore is ready to consume for provisioning virtual machines.

Ansible Playbook

```
## Disclaimer: Sample script for reference purpose only.
```

```
- hosts: '{{ vsphere_host }}'
  name: Play for vSphere iSCSI Configuration
  connection: local
  gather_facts: false
  tasks:
    # Generate Session ID for vCenter
    - name: Generate a Session ID for vCenter
      uri:
        url: "https://{{ vcenter_hostname }}/rest/com/vmware/cis/session"
        validate_certs: false
        method: POST
        user: "{{ vcenter_username }}"
       password: "{{ vcenter_password }}"
        force_basic_auth: yes
        return_content: yes
      register: vclogin

    # Generate Session ID for ONTAP tools with vCenter
    - name: Generate a Session ID for ONTAP tools with vCenter
      uri:
        url: "https://{{ ontap_tools_ip
}}:8143/api/rest/2.0/security/user/login"
        validate_certs: false
        method: POST
        return_content: yes
        body_format: json
        body:
          vcenterUserName: "{{ vcenter_username }}"
          vcenterPassword: "{{ vcenter_password }}"
      register: login

    # Get existing registered ONTAP Cluster info with ONTAP tools
    - name: Get ONTAP Cluster info from ONTAP tools
      uri:
        url: "https://{{ ontap_tools_ip
}}:8143/api/rest/2.0/storage/clusters"
        validate_certs: false
        method: Get
        return_content: yes
        headers:
          vmware-api-session-id: "{{ login.json.vmwareApiSessionId }}"
      register: clusterinfo

    - name: Get ONTAP Cluster ID
      set_fact:
        ontap_cluster_id: "{{ clusterinfo.json |
```

```
json_query(clusteridquery) }}"
      vars:
        clusteridquery: "records[?ipAddress == '{{ netapp_hostname }}' &&
type=='Cluster'].id | [0]"

    - name: Get ONTAP SVM ID
      set_fact:
        ontap_svm_id: "{{ clusterinfo.json | json_query(svmidquery) }}"
      vars:
        svmidquery: "records[?ipAddress == '{{ netapp_hostname }}' &&
type=='SVM' && name == '{{ svm_name }}'].id | [0]"

    - name: Get Aggregate detail
      uri:
        url: "https://{{ ontap_tools_ip
}}:8143/api/rest/2.0/storage/clusters/{{ ontap_svm_id }}/aggregates"
        validate_certs: false
        method: GET
        return_content: yes
        headers:
          vmware-api-session-id: "{{ login.json.vmwareApiSessionId }}"
          cluster-id: "{{ ontap_svm_id }}"
      when: ontap_svm_id != ''
      register: aggrinfo

    - name: Select Aggregate with max free capacity
      set_fact:
        aggr_name: "{{ aggrinfo.json | json_query(aggrquery) }}"
      vars:
        aggrquery: "max_by(records, &freeCapacity).name"

    - name: Convert datastore size in MB
      set_fact:
        datastoreSizeInMB: "{{ iscsi_datastore_size |
human_to_bytes/1024/1024 | int }}"

    - name: Get vSphere Cluster Info
      uri:
        url: "https://{{ vcenter_hostname }}/api/vcenter/cluster?names={{
vsphere_cluster }}"
        validate_certs: false
        method: GET
        return_content: yes
        body_format: json
        headers:
          vmware-api-session-id: "{{ vclogin.json.value }}"
```

```
      when: vsphere_cluster != ''
      register: vcenterclusterid

    - name: Create iSCSI VMFS-6 Datastore with ONTAP tools
      uri:
        url: "https://{{ ontap_tools_ip
}}:8143/api/rest/3.0/admin/datastore"
        validate_certs: false
        method: POST
        return_content: yes
        status_code: [200]
        body_format: json
        body:
          traditionalDatastoreRequest:
            name: "{{ iscsi_datastore_name }}"
            datastoreType: VMFS
            protocol: ISCSI
            spaceReserve: Thin
            clusterID:  "{{ ontap_cluster_id }}"
            svmID: "{{ ontap_svm_id }}"
            targetMoref: ClusterComputeResource:{{
vcenterclusterid.json[0].cluster }}
            datastoreSizeInMB: "{{ datastoreSizeInMB | int }}"
            vmfsFileSystem: VMFS6
            aggrName: "{{ aggr_name }}"
            existingFlexVolName: ""
            volumeStyle: FLEXVOL
            datastoreClusterMoref: ""
        headers:
          vmware-api-session-id: "{{ login.json.vmwareApiSessionId }}"
      when: ontap_cluster_id != '' and ontap_svm_id != '' and aggr_name !=
''
      register: result
      changed_when: result.status == 200
```

vSphere VMFS Datastore - NVMe/FC with ONTAP

This section covers the creation of a VMFS datastore with ONTAP storage using NVMe/FC.

What you need

- Basic skills needed to manage a vSphere environment and ONTAP.
- Basic understanding of NVMe/FC.
- An ONTAP Storage System (FAS/AFF/CVO/ONTAP Select/ASA) running ONTAP 9.8 or later
- ONTAP credentials (SVM name, userID, and password)

- ONTAP WWPN for host, target, and SVMs and LUN information
- A completed FC configuration worksheet
- vCenter Server
- vSphere host(s) information (vSphere 7.0 or later)
- Fabric switch(es)
  - With ONTAP FC data ports and vSphere hosts connected.
  - With the N_port ID virtualization (NPIV) feature enabled.
  - Create a single initiator target zone.
  - Create one zone for each initiator (single initiator zone).
  - For each zone, include a target that is the ONTAP FC logical interface (WWPN) for the SVMs. There should be at least two logical interfaces per node per SVM. DO not use the WWPN of physical ports.

Provision VMFS datastore

1. Check compatibility with the Interoperability Matrix Tool (IMT).
2. Verify that the NVMe/FC configuration is supported.

ONTAP tasks

1. Verify the ONTAP license for FCP.
   Use the `system license show` command and check if NVMe_oF is listed.
   Use `license add -license-code <license code>` to add a license.
2. Verify that NVMe protocol is enabled on the SVM.
   a. Configure SVMs for NVMe.
3. Verify that NVMe/FC Logical Interfaces are available on the SVMs.
   a. Use `Network Interface show` to verify the FCP adapter.
   b. When an SVM is created with the GUI, logical interfaces are as part of that process.
   c. To rename the network interface, use the command `Network Interface modify`.
4. Create NVMe namespace and subsystem

VMware vSphere Tasks

1. Verify that HBA drivers are installed. VMware supported HBAs have the drivers deployed out of the box and should be visible at Storage Adapter Information
2. Perform vSphere Host NVMe driver installatioln and validation tasks
3. Create VMFS Datastore

vSphere traditional file storage provisioning with ONTAP

VMware vSphere supports following NFS protocols, both of which support ONTAP.

- NFS Version 3
- NFS Version 4.1

If you need help selecting the correct NFS version for vSphere, check this comparison of NFS client versions.

Reference

vSphere datastore and protocol features: NFS

vSphere NFS datastore - Version 3 with ONTAP

Creation of NFS version 3 datastore with ONTAP NAS storage.

What you need

- The basic skill necessary to manage a vSphere environment and ONTAP.
- An ONTAP storage system (FAS/AFF/CVO/ONTAP Select/Cloud Volume Service/Azure NetApp Files) running ONTAP 9.8 or later
- ONTAP credentials (SVM name, userID, password)
- ONTAP network port, SVM, and LUN information for NFS
  - A completed NFS configuration worksheet
- vCenter Server credentials
- vSphere host(s) information for vSphere 7.0 or later
- NFS VMKernel adapter IP information
- Network switch(es)
  - with ONTAP system network data ports and connected vSphere hosts
  - VLAN(s) configured for NFS
  - (Optional) link aggregation configured for ONTAP network data ports
- ONTAP Tool for VMware vSphere deployed, configured, and ready to consume

Steps

- Check compatibility with the Interoperability Matrix Tool (IMT)
  - Verify that the NFS configuration is supported.
- Complete the following ONTAP and vSphere tasks.

ONTAP tasks

1. Verify the ONTAP license for NFS.
   a. Use the `system license show` command and check that NFS is listed.
   b. Use `license add -license-code <license code>` to add a license.
2. Follow the NFS configuration workflow.

VMware vSphere Tasks

Follow the workflow for NFS client configuration for vSphere.

Reference

vSphere datastore and protocol features: NFS

What's next?

After these tasks are completed, the NFS datastore is ready to consume for provisioning virtual machines.

vSphere NFS Datastore - Version 4.1 with ONTAP

This section describes the creation of an NFS version 4.1 datastore with ONTAP NAS storage.

What you need

- The basic skills necessary to manage a vSphere environment and ONTAP
- ONTAP Storage System (FAS/AFF/CVO/ONTAP Select/Cloud Volume Service/Azure NetApp Files) running ONTAP 9.8 or later
- ONTAP credentials (SVM name, userID, password)
- ONTAP network port, SVM, and LUN information for NFS
- A completed NFS configuration worksheet
- vCenter Server credentials
- vSphere host(s) information vSphere 7.0 or later
- NFS VMKernel adapter IP information
- Network switch(es)
  - with ONTAP system network data ports, vSphere hosts, and connected
  - VLAN(s) configured for NFS
  - (Optional) link aggregation configured for ONTAP network data ports
- ONTAP Tools for VMware vSphere deployed, configured, and ready to consume

Steps

- Check compatability with the Interoperability Matrix Tool (IMT).
  - Verify that the NFS configuration is supported.
- Complete the ONTAP and vSphere Tasks provided below.

ONTAP tasks

1. Verify ONTAP license for NFS
   a. Usethe `system license show` command to check whether NFS is listed.
   b. Use `license add -license-code <license code>` to add a license.
2. Follow the NFS configuration workflow

VMware vSphere tasks

Follow the NFS Client Configuration for vSphere workflow.

What's next?

After these tasks are completed, the NFS datastore is ready to consume for provisioning virtual machines.

Virtual Desktops

Virtual Desktop Services (VDS)

TR-4861: Hybrid Cloud VDI with Virtual Desktop Service

Suresh Thoppay, NetApp

The NetApp Virtual Desktop Service (VDS) orchestrates Remote Desktop Services (RDS) in major public clouds as well as on private clouds. VDS supports Windows Virtual Desktop (WVD) on Microsoft Azure. VDS automates many tasks that must be performed after deployment of WVD or RDS, including setting up SMB file shares (for user profiles, shared data, and the user home drive), enabling Windows features, application and agent installation, firewall, and policies, and so on.

Users consume VDS for dedicated desktops, shared desktops, and remote applications. VDS provides scripted events for automating application management for desktops and reduces the number of images to manage.

VDS provides a single management portal for handling deployments across public and private cloud environments.

Customer Value

The remote workforce explosion of 2020 has changed requirements for business continuity. IT departments are faced with new challenges to rapidly provision virtual desktops and thus require provisioning agility, remote management, and the TCO advantages of a hybrid cloud that makes it easy to provision on-premises and cloud resources. They need a hybrid-cloud solution that:

- Addresses the post-COVID workspace reality to enable flexible work models with global dynamics
- Enables shift work by simplifying and accelerating the deployment of work environments for all employees, from task workers to power users
- Mobilizes your workforce by providing rich, secure VDI resources regardless of the physical location
- Simplifies hybrid-cloud deployment
- Automates and simplifies risk reduction management

Use Cases

Hybrid VDI with NetApp VDS allows service providers and enterprise virtual desktop administrators to easily expand resources to other cloud environment without affecting their users. Having on-premises resources provides better control of resources and offers wide selection of choices (compute, GPU, storage, and network) to meet demand.

This solution applies to the following use cases:

- Bursting into the cloud for surges in demand for remote desktops and applications
- Reducing TCO for long running remote desktops and applications by hosting them on-premises with flash storage and GPU resources
- Ease of management of remote desktops and applications across cloud environments
- Experience remote desktops and applications by using a software-as-a- service model with on-premises resources

Target Audience

The target audience for the solution includes the following groups:

- EUC/VDI architects who wants to understand the requirements for a hybrid VDS
- NetApp partners who would like to assist customers with their remote desktop and application needs
- Existing NetApp HCI customers who want to address remote desktop and application demands

NetApp Virtual Desktop Service Overview

NetApp offers many cloud services, including the rapid provisioning of virtual desktop with WVD or remote applications and rapid integration with Azure NetApp Files.

Traditionally, it takes weeks to provision and deliver remote desktop services to customers. Apart from provisioning, it can be difficult to manage applications, user profiles, shared data, and group policy objects to enforce policies. Firewall rules can increase complexity and require a separate skillset and tools.

With Microsoft Azure Windows Virtual Desktop service, Microsoft takes care of maintenance for Remote Desktop Services components, allowing customers to focus on provisioning workspaces in the cloud. Customers must provision and manage the complete stack which requires special skills to manage VDI environments.

With NetApp VDS, customers can rapidly deploy virtual desktops without worrying about where to install the architecture components like brokers, gateways, agents, and so on. Customers who require complete control of their environment can work with a professional services team to achieve their goals. Customers consume VDS as a service and thus can focus on their key business challenges.

NetApp VDS is a software-as-a-service offering for centrally managing multiple deployments across AWS, Azure, GCP, or private cloud environments. Microsoft Windows Virtual Desktop is available only on Microsoft Azure. NetApp VDS orchestrates Microsoft Remote Desktop Services in other environments.

Microsoft offers multisession on Windows 10 exclusively for Windows Virtual Desktop environments on Azure. Authentication and identity are handled by the virtual desktop technology; WVD requires Azure Active Directory synced (with AD Connect) to Active Directory and session VMs joined to Active Directory. RDS requires Active Directory for user identity and authentication and VM domain join and management.

A sample deployment topology is shown in the following figure.

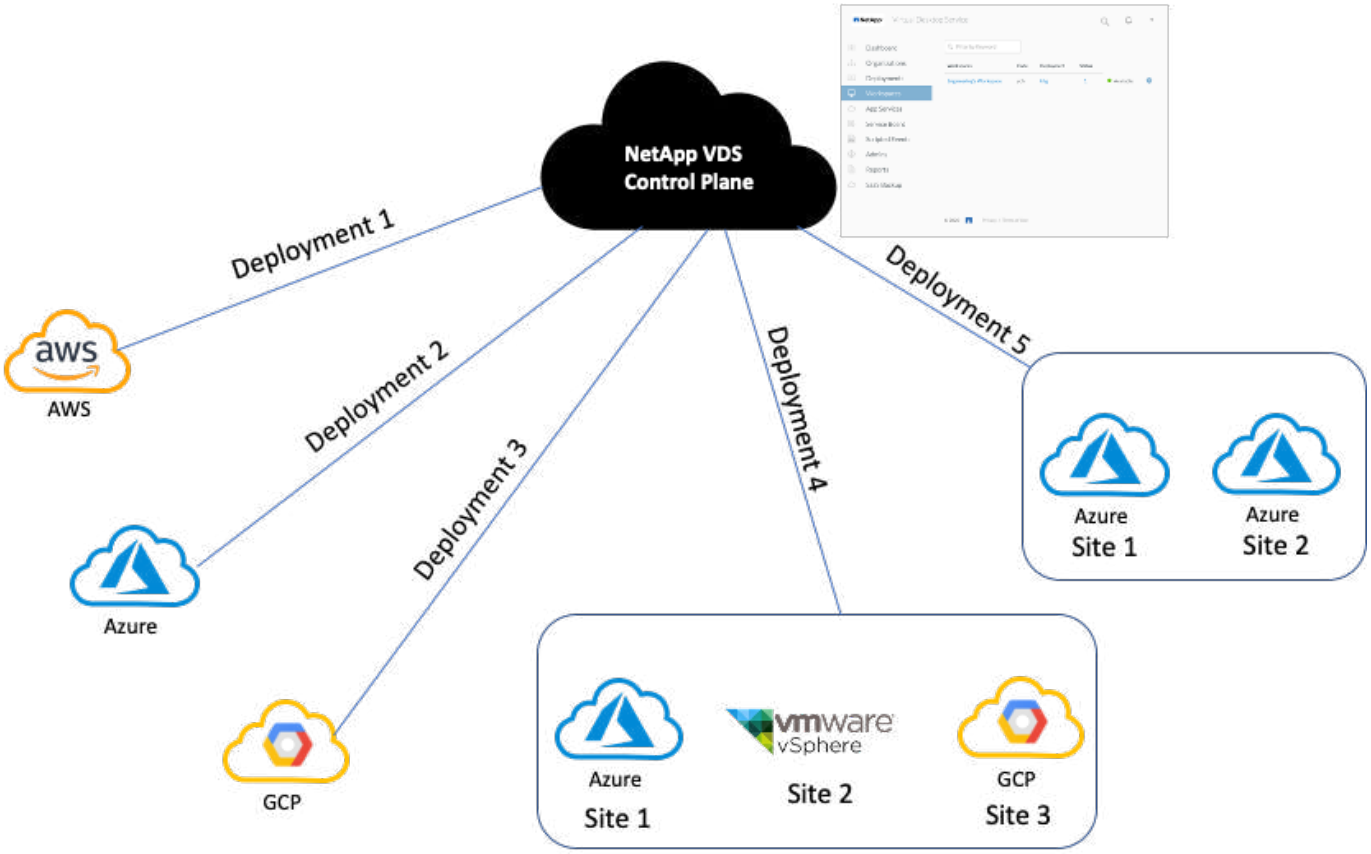

Each deployment is associated with an active directory domain and provides clients with an access entry point for workspaces and applications. A service provider or enterprise that has multiple active directory domains typically has more deployments. A single Active Directory domain that spans multiple regions typically has a single deployment with multiple sites.

For WVD in Azure, Microsoft provides a platform-as-a-service that is consumed by NetApp VDS. For other environments, NetApp VDS orchestrates the deployment and configuration of Microsoft Remote Desktop Services. NetApp VDS supports both WVD Classic and WVD ARM and can also be used to upgrade existing versions.

Each deployment has its own platform services, which consists of Cloud Workspace Manager (REST API endpoint), an HTML 5 Gateway (connect to VMs from a VDS management portal), RDS Gateways (Access point for clients), and a Domain Controller. The following figure depicts the VDS Control Plane architecture for RDS implementation.

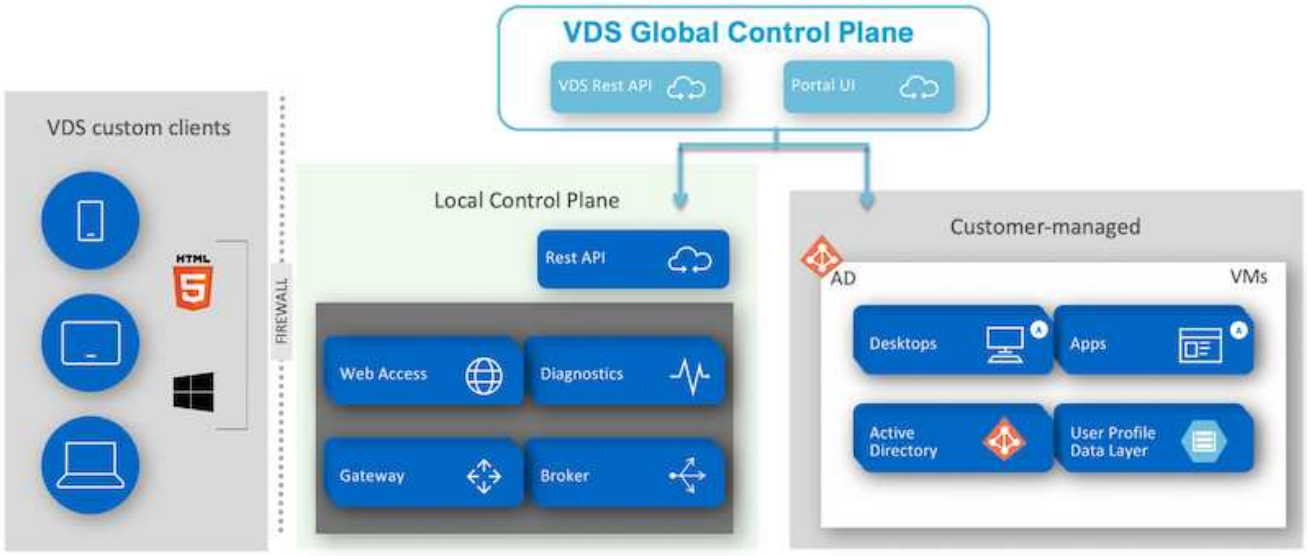

For RDS implementations, NetApp VDS can be readily accessed from Windows and browsers using client software that can be customized to include customer logo and images. Based on user credentials, it provides user access to approved workspaces and applications. There is no need to configure the gateway details.

The following figure shows the NetApp VDS client.

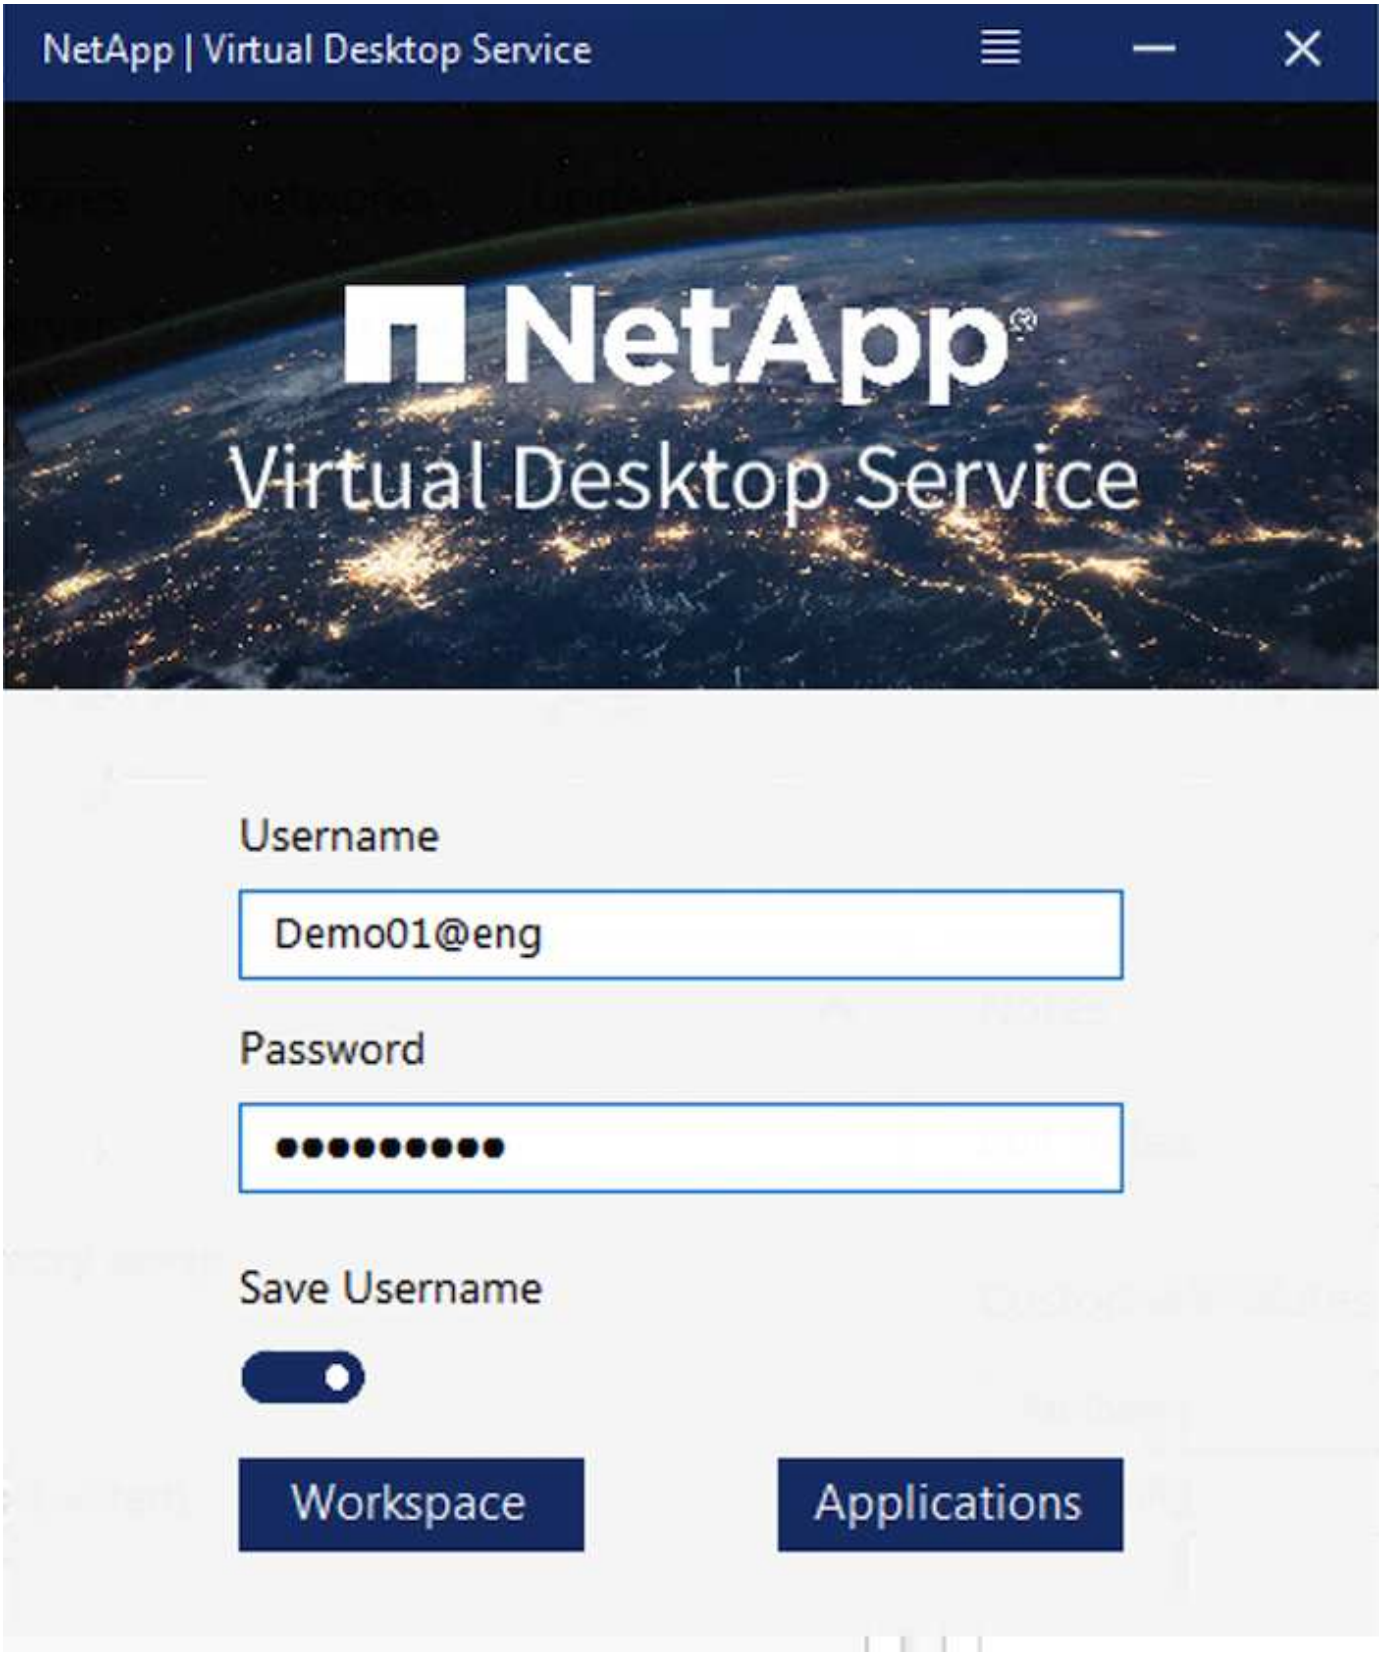

In the Azure WVD implementation, Microsoft handles the access entry point for the clients and can be consumed by a Microsoft WVD client available natively for various OSs. It can also be accessed from a web-based portal. The configuration of client software must be handled by the Group Policy Object (GPO) or in other ways preferred by customers.

The following figure depicts the VDS Control Plane architecture for Azure WVD implementations.

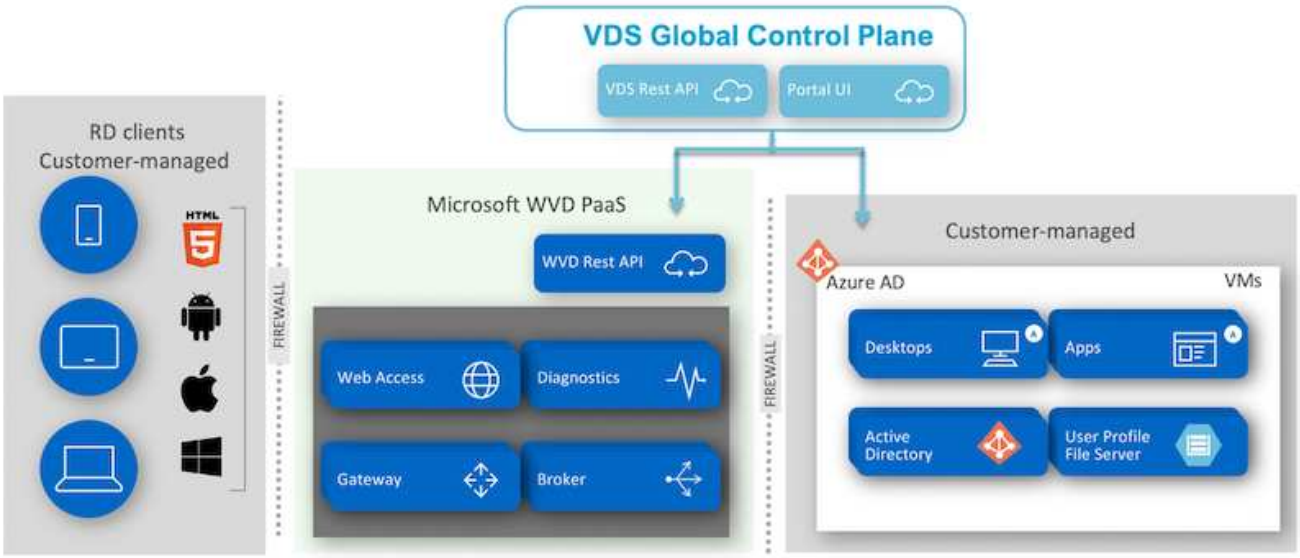

In addition to the deployment and configuration of required components, NetApp VDS also handles user management, application management, resource scaling, and optimization.

NetApp VDS can create users or grant existing user accounts access to cloud workspace or application services. The portal can also be used for password resets and the delegation of administrating a subset of components. Helpdesk administrators or Level-3 technicians can shadow user sessions for troubleshooting or connect to servers from within the portal.

NetApp VDS can use image templates that you create, or it can use existing ones from the marketplace for cloud-based provisioning. To reduce the number of images to manage, you can use a base image, and any additional applications that you require can be provisioned using the provided framework to include any command-line tools like Chocolatey, MSIX app attach, PowerShell, and so on. Even custom scripts can be used as part of machine lifecycle events.

NetApp HCI Overview

NetApp HCI is a hybrid cloud infrastructure that consists of a mix of storage nodes and compute nodes. It is available as either a two-rack unit or single-rack unit, depending on the model. The installation and configuration required to deploy VMs are automated with the NetApp Deployment Engine (NDE). Compute clusters are managed with VMware vCenter, and storage clusters are managed with the vCenter Plug-in deployed with NDE. A management VM called the mNode is deployed as part of the NDE.

NetApp HCI handles the following functions:

- Version upgrades
- Pushing events to vCenter
- vCenter Plug-In management
- A VPN tunnel for support
- The NetApp Active IQ Digital Advisor (also known as Digital Advisor) collector
- The extension of NetApp Cloud Services to on the premises, enabling a hybrid cloud infrastructure. The following figure depicts HCI components.

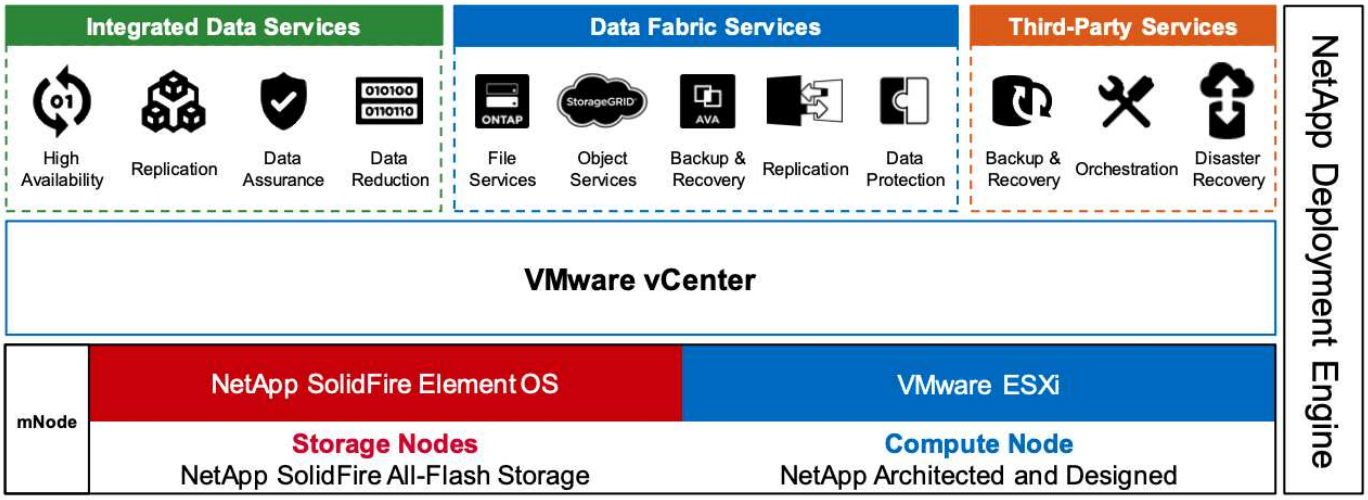

Storage Nodes

Storage nodes are available as either a half-width or full-width rack unit. A minimum of four storage nodes is required at first, and a cluster can expand to up to 40 nodes. A storage cluster can be shared across multiple compute clusters. All the storage nodes contain a cache controller to improve write performance. A single node provides either 50K or 100K IOPS at a 4K block size.

NetApp HCI storage nodes run NetApp Element software, which provides minimum, maximum, and burst QoS limits. The storage cluster supports a mix of storage nodes, although one storage node cannot exceed one-third of total capacity.

Compute Nodes

NetApp supports its storage connected to any compute servers listed in the VMware Compatability Guide.

Compute nodes are available in half-width, full-width, and two rack-unit sizes. The NetApp HCI H410C and H610C are based on scalable Intel Skylake processors. The H615C is based on second-generation scalable Intel Cascade Lake processors. There are two compute models that contain GPUs: the H610C contains two NVIDIA M10 cards and the H615C contains three NVIDIA T4 cards.

The NVIDIA T4 has 40 RT cores that provide the computation power needed to deliver real-time ray tracing. The same server model used by designers and engineers can now also be used by artists to create photorealistic imagery that features light bouncing off surfaces just as it would in real life. This RTX-capable GPU produces real-time ray tracing performance of up to five Giga Rays per second. The NVIDIA T4, when combined with Quadro Virtual Data Center Workstation (Quadro vDWS) software, enables artists to create photorealistic designs with accurate shadows, reflections, and refractions on any device from any location.

Tensor cores enable you to run deep learning inferencing workloads. When running these workloads, an NVIDIA T4 powered with Quadro vDWS can perform up to 25 times faster than a VM driven by a CPU-only server. A NetApp H615C with three NVIDIA T4 cards in one rack unit is an ideal solution for graphics and compute-intensive workloads.

The following figure lists NVIDIA GPU cards and compares their features.

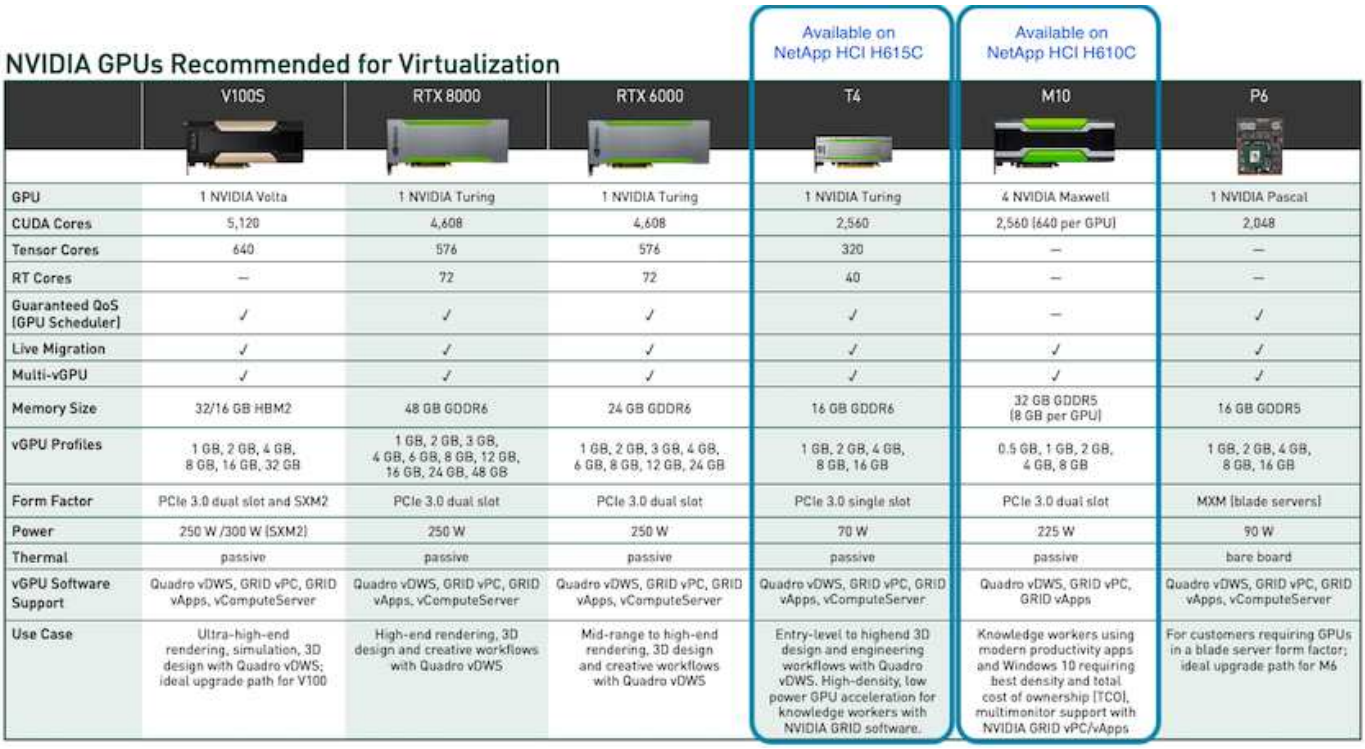

	V100S	RTX 8000	RTX 6000	T4	M10	P6
GPU	1 NVIDIA Volta	1 NVIDIA Turing	1 NVIDIA Turing	1 NVIDIA Turing	4 NVIDIA Maxwell	1 NVIDIA Pascal
CUDA Cores	5,120	4,608	4,608	2,560	2,560 (640 per GPU)	2,048
Tensor Cores	640	576	576	320	–	–
RT Cores	–	72	72	40	–	–
Guaranteed QoS (GPU Scheduler)	✓	✓	✓	✓	–	✓
Live Migration	✓	✓	✓	✓	✓	✓
Multi-vGPU	✓	✓	✓	✓	✓	✓
Memory Size	32/16 GB HBM2	48 GB GDDR6	24 GB GDDR6	16 GB GDDR6	32 GB GDDR5 (8 GB per GPU)	16 GB GDDR5
vGPU Profiles	1 GB, 2 GB, 4 GB, 8 GB, 16 GB, 32 GB	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 16 GB, 24 GB, 48 GB	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 24 GB	1 GB, 2 GB, 4 GB, 8 GB, 16 GB	0.5 GB, 1 GB, 2 GB, 4 GB, 8 GB	1 GB, 2 GB, 4 GB, 8 GB, 16 GB
Form Factor	PCIe 3.0 dual slot and SXM2	PCIe 3.0 dual slot	PCIe 3.0 dual slot	PCIe 3.0 single slot	PCIe 3.0 dual slot	MXM (blade servers)
Power	250 W /300 W (SXM2)	250 W	250 W	70 W	225 W	90 W
Thermal	passive	passive	passive	passive	passive	bare board
vGPU Software Support	Quadro vDWS, GRID vPC, GRID vApps, vComputeServer	Quadro vDWS, GRID vPC, GRID vApps, vComputeServer	Quadro vDWS, GRID vPC, GRID vApps, vComputeServer	Quadro vDWS, GRID vPC, GRID vApps, vComputeServer	Quadro vDWS, GRID vPC, GRID vApps	Quadro vDWS, GRID vPC, GRID vApps, vComputeServer
Use Case	Ultra-high-end rendering, simulation, 3D design with Quadro vDWS; ideal upgrade path for V100	High-end rendering, 3D design and creative workflows with Quadro vDWS	Mid-range to high-end rendering, 3D design and creative workflows with Quadro vDWS	Entry-level to highend 3D design and engineering workflows with Quadro vDWS. High-density, low power GPU acceleration for knowledge workers with NVIDIA GRID software.	Knowledge workers using modern productivity apps and Windows 10 requiring best density and total cost of ownership (TCO), multimonitor support with NVIDIA GRID vPC/vApps	For customers requiring GPUs in a blade server form factor; ideal upgrade path for M6

The M10 GPU remains the best TCO solution for knowledge-worker use cases. However, the T4 makes a great alternative when IT wants to standardize on a GPU that can be used across multiple use cases, such as virtual workstations, graphics performance, real-time interactive rendering, and inferencing. With the T4, IT can take advantage of the same GPU resources to run mixed workloads—for example, running VDI during the day and repurposing the resources to run compute workloads at night.

The H610C compute node is two rack units in size; the H615C is one rack unit in size and consumes less power. The H615C supports H.264 and H.265 (High Efficiency Video Coding [HEVC]) 4:4:4 encoding and decoding. It also supports the increasingly mainstrean VP9 decoder; even the WebM container package served by YouTube uses the VP9 codec for video.

The number of nodes in a compute cluster is dictated by VMware; currently, it is 96 with VMware vSphere 7.0 Update 1. Mixing different models of compute nodes in a cluster is supported when Enhanced vMotion Compatibility (EVC) is enabled.

NVIDIA Licensing

When using an H610C or H615C, the license for the GPU must be procured from NVIDIA partners that are authorized to resell the licenses. You can find NVIDIA partners with the partner locator. Search for competencies such as virtual GPU (vGPU) or Tesla.

NVIDIA vGPU software is available in four editions:

- NVIDIA GRID Virtual PC (GRID vPC)
- NVIDIA GRID Virtual Applications (GRID vApps)
- NVIDIA Quadro Virtual Data Center Workstation (Quadro vDWS)
- NVIDIA Virtual ComputeServer (vComputeServer)

GRID Virtual PC

This product is ideal for users who want a virtual desktop that provides a great user experience for Microsoft Windows applications, browsers, high-definition video, and multi-monitor support. The NVIDIA GRID Virtual PC delivers a native experience in a virtual environment, allowing you to run all your PC applications at full performance.

GRID Virtual Applications

GRID vApps are for organizations deploying a Remote Desktop Session Host (RDSH) or other app-streaming or session-based solutions. Designed to deliver Microsoft Windows applications at full performance, Windows Server-hosted RDSH desktops are also supported by GRID vApps.

Quadro Virtual Data Center Workstation

This edition is ideal for mainstream and high-end designers who use powerful 3D content creation applications like Dassault CATIA, SOLIDWORKS, 3Dexcite, Siemens NX, PTC Creo, Schlumberger Petrel, or Autodesk Maya. NVIDIA Quadro vDWS allows users to access their professional graphics applications with full features and performance anywhere on any device.

NVIDIA Virtual ComputeServer

Many organizations run compute-intensive server workloads such as artificial intelligence (AI), deep learning (DL), and data science. For these use cases, NVIDIA vComputeServer software virtualizes the NVIDIA GPU, which accelerates compute-intensive server workloads with features such as error correction code, page retirement, peer-to-peer over NVLink, and multi-vGPU.

A Quadro vDWS license enables you to use GRID vPC and NVIDIA vComputeServer.

Deployment

NetApp VDS can be deployed to Microsoft Azure using a setup app available based on the required codebase. The current release is available here and the preview release of the upcoming product is available here.

See this video for deployment instructions.

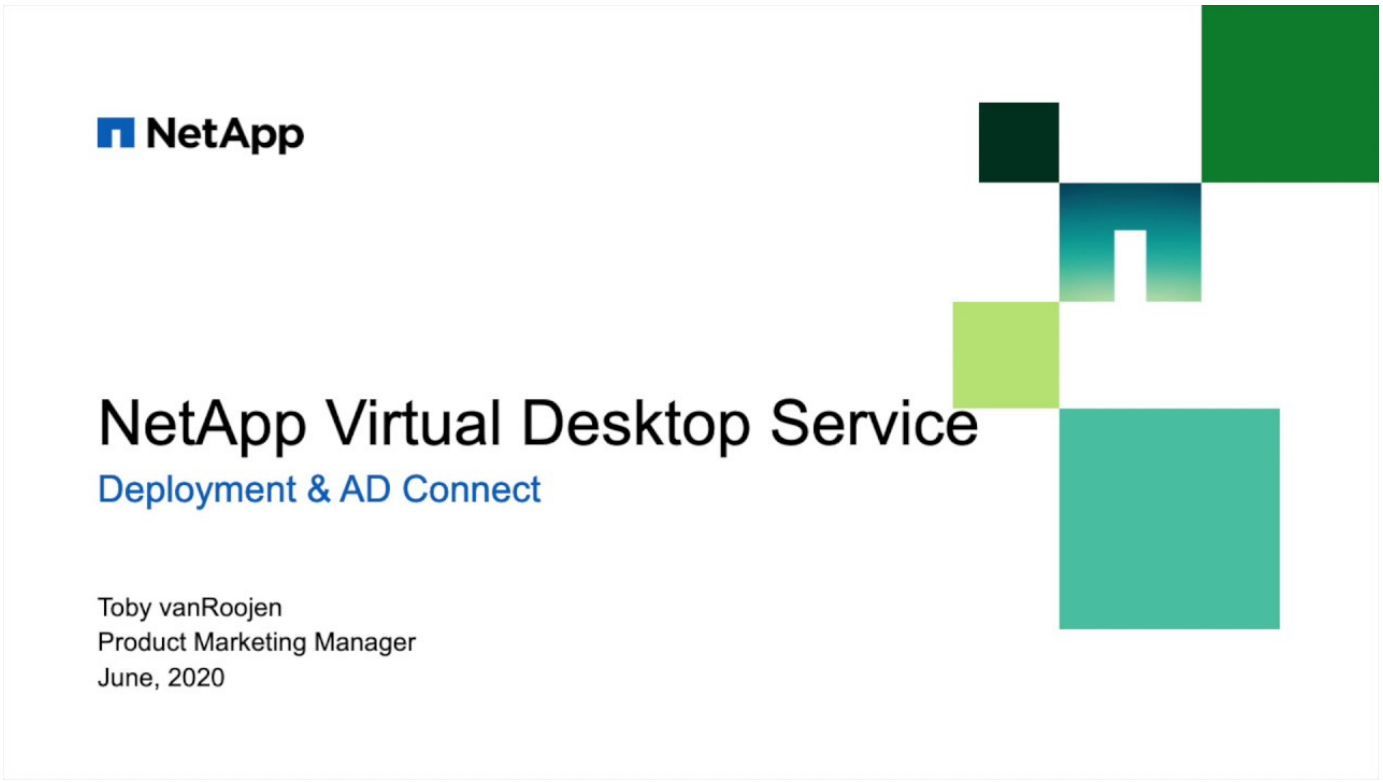

Hybrid Cloud Environment

NetApp Virtual Desktop Service can be extended to on-premises when connectivity exists between on-premises resources and cloud resources. Enterprises can establish the link to Microsoft Azure using Express Route or a site-to-site IPsec VPN connection. You can also create links to other clouds in a similar way either using a dedicated link or with an IPsec VPN tunnel.

For the solution validation, we used the environment depicted in the following figure.

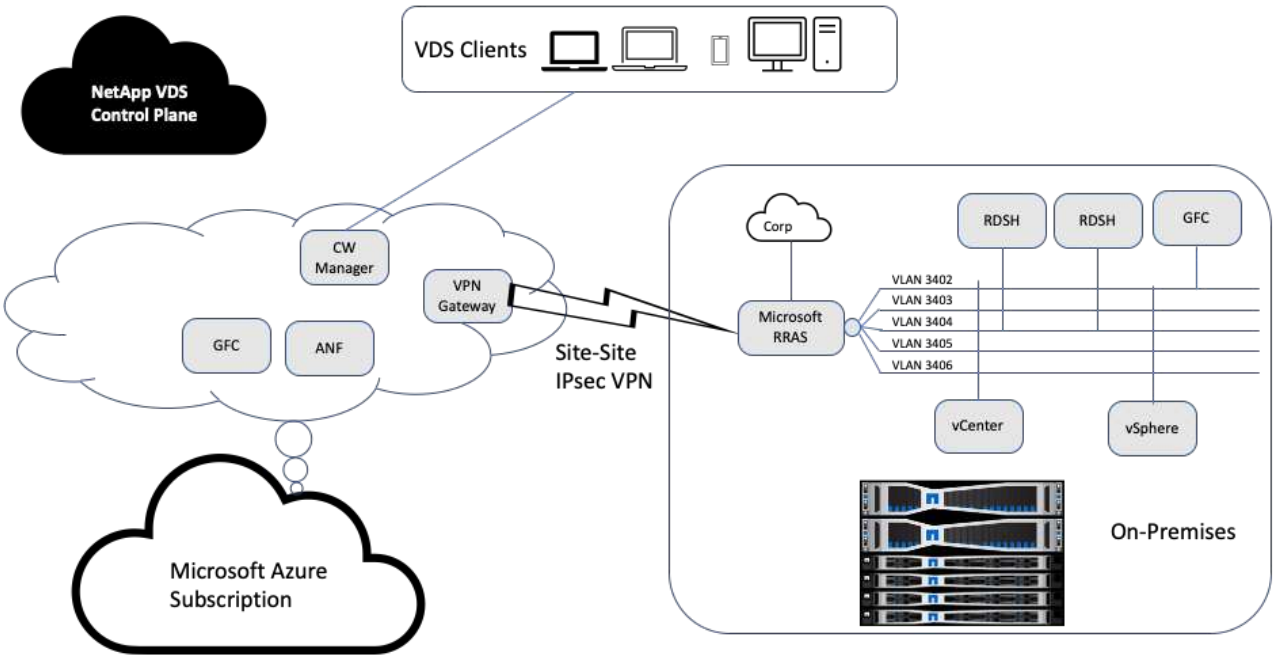

On-premises, we had multiple VLANs for management, remote-desktop-session hosts, and so on. They were on the 172.21.146-150.0/24 subnet and routed to the corporate network using the Microsoft Remote Routing Access Service. We also performed the following tasks:

1. We noted the public IP of the Microsoft Routing and Remote Access Server (RRAS; identified with IPchicken.com).
2. We created a Virtual Network Gateway resource (route-based VPN) on Azure Subscription.
3. We created the connection providing the local network gateway address for the public IP of the Microsoft RRAS server.
4. We completed VPN configuration on RRAS to create a virtual interface using pre-shared authentication that was provided while creating the VPN gateway. If configured correctly, the VPN should be in the connected state. Instead of Microsoft RRAS, you can also use pfSense or other relevant tools to create the site-to-site IPsec VPN tunnel. Since it is route-based, the tunnel redirects traffic based on the specific subnets configured.

Microsoft Azure Active Directory provides identity authentication based on oAuth. Enterprise client authentications typically require NTLM or Kerberos-based authentication. Microsoft Azure Active Directory Domain Services perform password hash sync between Azure Active Directory and on-prem domain controllers using ADConnect.

For this Hybrid VDS solution validation, we initially deployed to Microsoft Azure and added an additional site with vSphere. The advantage with this approach is that platform services were deployed to Microsoft Azure and were then readily backed up using the portal. Services can then be easily accessed from anywhere, even if the site-site VPN link is down.

To add another site, we used a tool called DCConfig. The shortcut to that application is available on the desktop of the cloud workspace manager (CWMgr) VM. After this application is launched, navigate to the DataCenter Sites tab, add the new datacenter site, and fill in the required info as shown below. The URL points to the vCenter IP. Make sure that the CWMgr VM can communicate with vCenter before adding the

configuration.

Make sure that vSphere PowerCLI 5.1 on CloudWorkspace manager is installed to enable communication with VMware vSphere environment.

The following figure depicts on- premises datacenter site configuration.

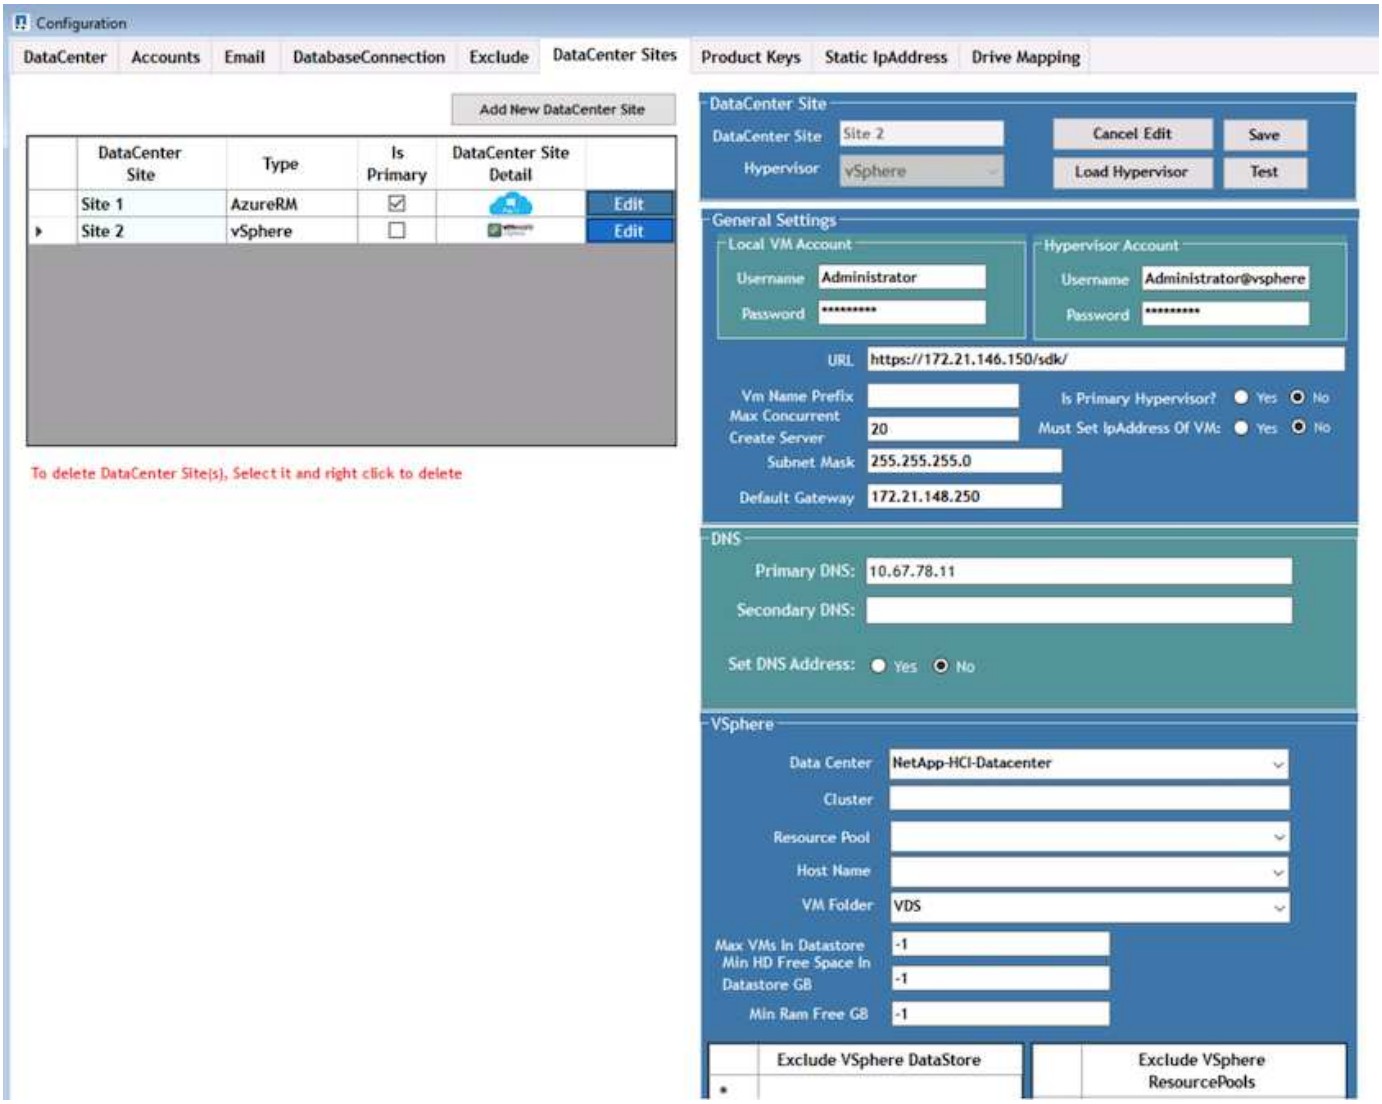

Note that there are filtering options available for compute resource based on the specific cluster, host name, or free RAM space. Filtering options for storage resource includes the minimum free space on datastores or the maximum VMs per datastore. Datastores can be excluded using regular expressions. Click Save button to save the configuration.

To validate the configuration, click the Test button or click Load Hypervisor and check any dropdown under the vSphere section. It should be populated with appropriate values. It is a best practice to keep the primary hypervisor set to yes for the default provisioning site.

The VM templates created on VMware vSphere are consumed as provisioning collections on VDS. Provisioning collections come in two forms: shared and VDI. The shared provisioning collection type is used for remote desktop services for which a single resource policy is applied to all servers. The VDI type is used for WVD instances for which the resource policy is individually assigned. The servers in a provisioning collection can be assigned one of the following three roles:

- **TSDATA.** Combination of Terminal Services and Data server role.
- **TS.** Terminal Services (Session Host).
- **DATA.** File Server or Database Server. When you define the server role, you must pick the VM template and storage (datastore). The datastore chosen can be restricted to a specific datastore or you can use the least-used option in which the datastore is chosen based on data usage.

Each deployment has VM resource defaults for the cloud resource allocation based on Active Users, Fixed, Server Load, or User Count.

Single server load test with Login VSI

The NetApp Virtual Desktop Service uses the Microsoft Remote Desktop Protocol to access virtual desktop sessions and applications, and the Login VSI tool determines the maximum number of users that can be hosted on a specific server model. Login VSI simulates user login at specific intervals and performs user operations like opening documents, reading and composing mails, working with Excel and PowerPoint, printing documents, compressing files, and taking random breaks. It then measures response times. User response time is low when server utilization is low and increases when more user sessions are added. Login VSI determines the baseline based on initial user login sessions and it reports the maximum user session when the user response exceeds 2 seconds from the baseline.

NetApp Virtual Desktop Service utilizes Microsoft Remote Desktop Protocol to access the Virtual Desktop session and Applications. To determine the maximum number of users that can be hosted on a specific server model, we used the Login VSI tool. Login VSI simulates user login at specific intervals and performs user operations like opening documents, reading and composing mails, working with Excel and PowerPoint, printing documents, compressing files, taking random breaks, and so on. It also measures response times. User response time is low when server utilization is low and increases when more user sessions are added. Login VSI determines the baseline based on the initial user login sessions and it reports maximum user sessions when the user response exceeds 2sec from the baseline.

The following table contains the hardware used for this validation.

Model	Count	Description
NetApp HCI H610C	4	Three in a cluster for launchers, AD, DHCP, and so on. One server for load testing.
NetApp HCI H615C	1	2x24C Intel Xeon Gold 6282 @2.1GHz. 1.5TB RAM.

The following table contains the software used for this validation.

Product	Description
NetApp VDS 5.4	Orchestration
VM Template Windows 2019 1809	Server OS for RDSH
Login VSI	4.1.32.1
VMware vSphere 6.7 Update 3	Hypervisor

Product	Description
VMware vCenter 6.7 Update 3f	VMware management tool

The Login VSI test results are as follows:

Model	VM configuration	Login VSI baseline	Login VSI Max
H610C	8 vCPU, 48GB RAM, 75GB disk, 8Q vGPU profile	799	178
H615C	12 vCPU, 128GB RAM, 75GB disk	763	272

Considering sub-NUMA boundaries and hyperthreading, the eight VMs chosen for VM testing and configuration depended on the cores available on the host.

We used 10 launcher VMs on the H610C, which used the RDP protocol to connect to the user session. The following figure depicts the Login VSI connection information.

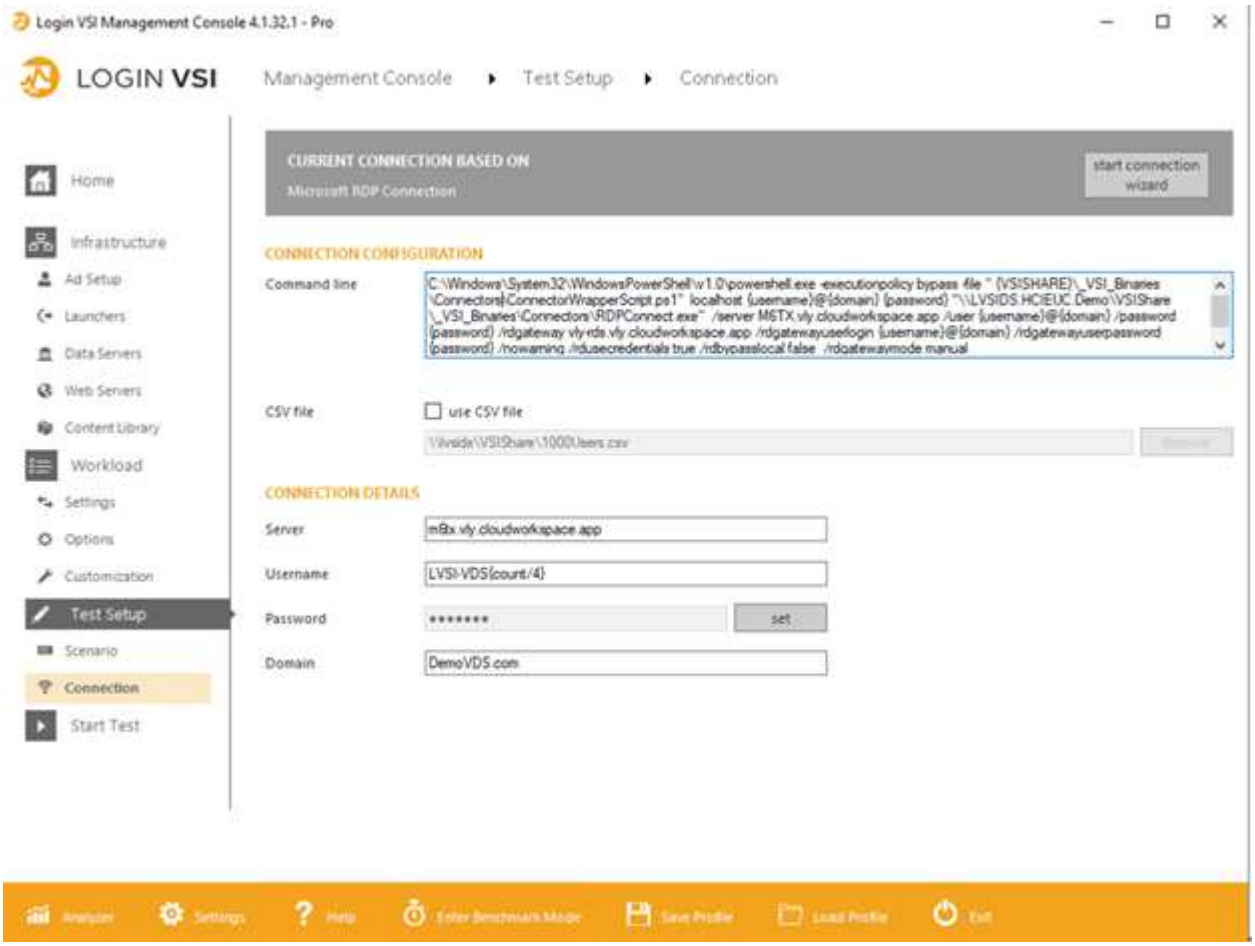

The following figure displays the Login VSI response time versus the active sessions for the H610C.

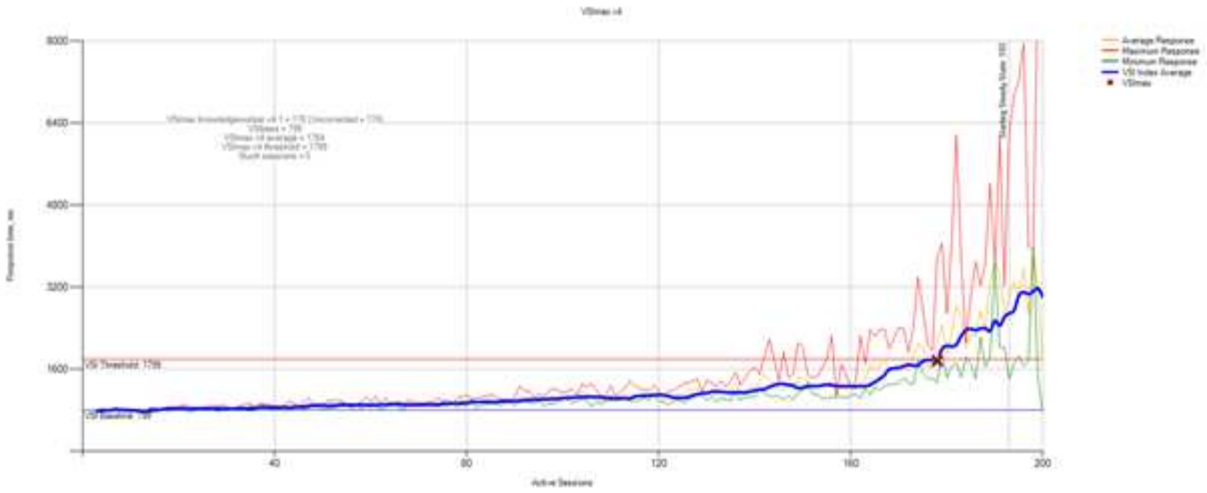

The following figure displays the Login VSI response time versus active sessions for the H615C.

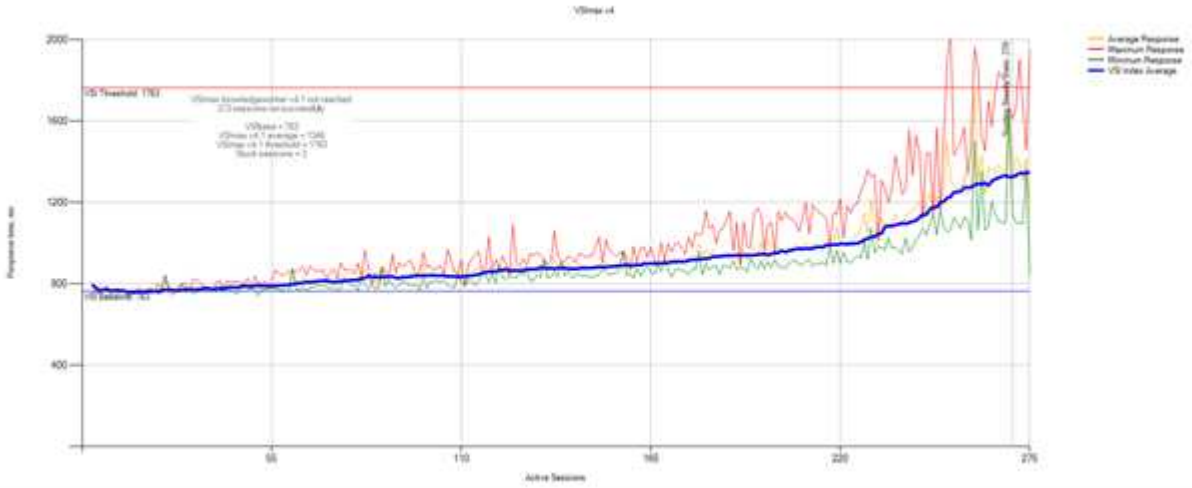

The performance metrics from Cloud Insights during H615C Login VSI testing for the vSphere host and VMs are shown in the following figure.

Management Portal

NetApp VDS Cloud Workspace Management Suite portal is available here and the upcoming version is available here.

The portal allows centralized management for various VDS deployments including one that has sites defined for on-premises, administrative users, the application catalog, and scripted events. The portal is also used by administrative users for the manual provisioning of applications if required and to connect to any machines for troubleshooting.

Service providers can use this portal to add their own channel partners and allow them to manage their own clients.

User Management

NetApp VDS uses Azure Active Directory for identity authentication and Azure Active Directory Domain Services for NTLM/Kerberos authentication. The ADConnect tool can be used to sync an on-prem Active Directory domain with Azure Active Directory.

New users can be added from the portal, or you can enable cloud workspace for existing users. Permissions for workspaces and application services can be controlled by individual users or by groups. From the management portal, administrative users can be defined to control permissions for the portal, workspaces, and so on.

The following figure depicts user management in NetApp VDS.

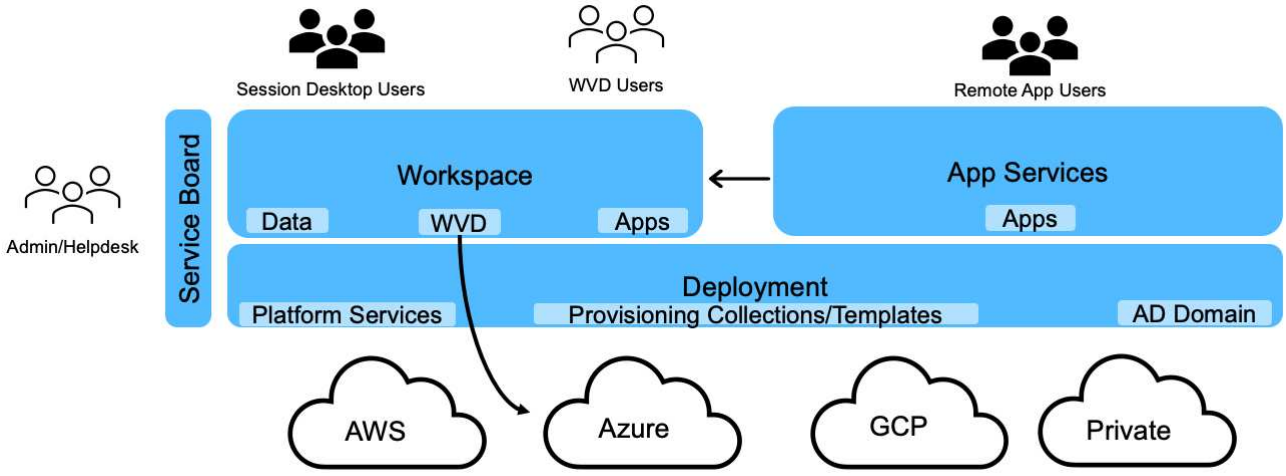

Each workspace resides in its own Active Directory organization unit (OU) under the Cloud Workspace OU as shown in the following figure.

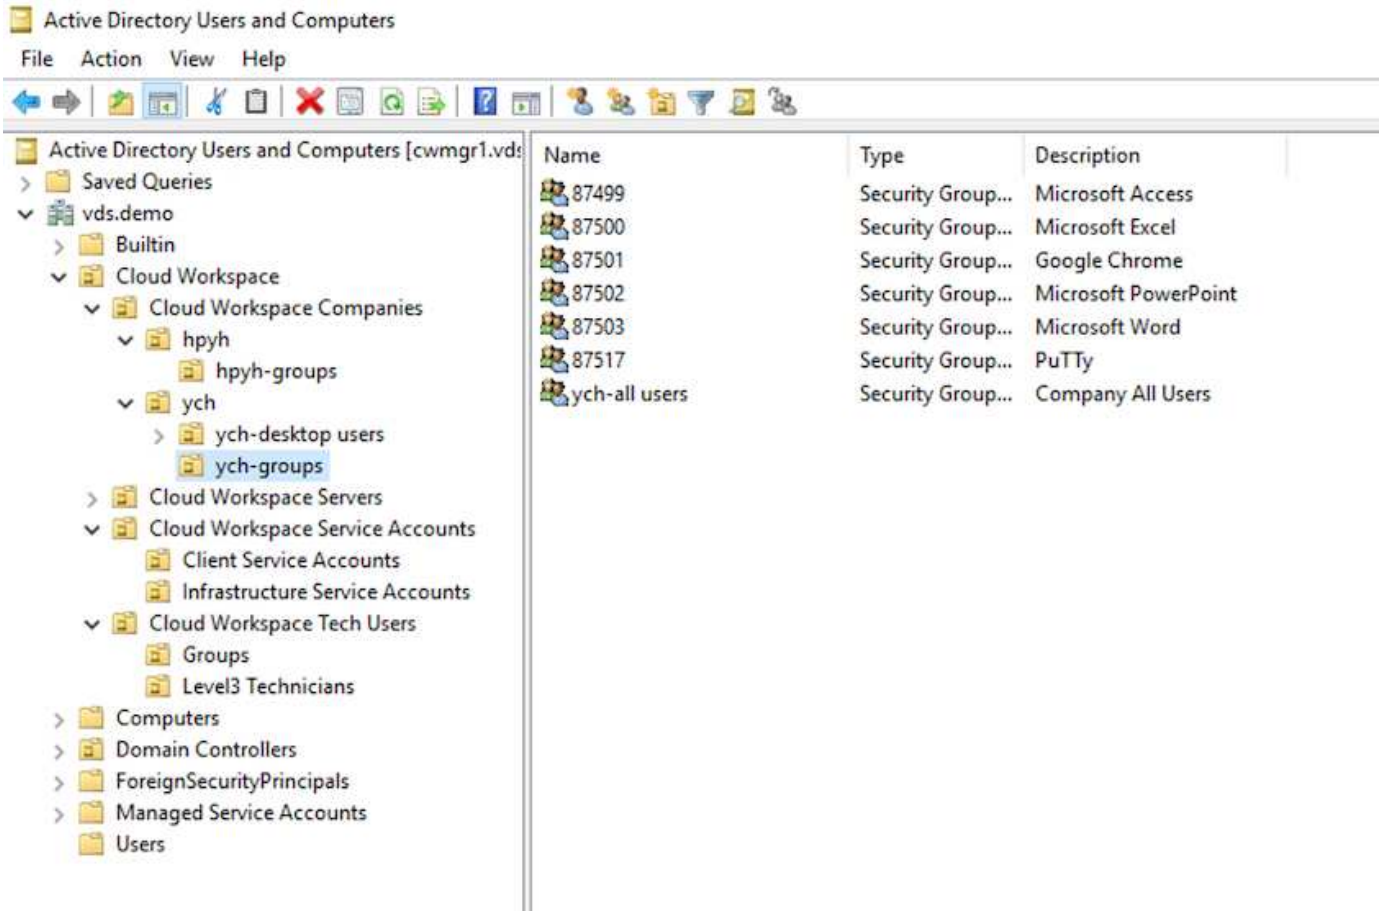

For more info, see this video on user permissions and user management in NetApp VDS.

When an Active Directory group is defined as a CRAUserGroup using an API call for the datacenter, all the users in that group are imported into the CloudWorkspace for management using the UI. As the cloud workspace is enabled for the user, VDS creates user home folders, settings permissions, user properties updates, and so on.

If VDI User Enabled is checked, VDS creates a single-session RDS machine dedicated to that user. It prompts

for the template and the datastore to provision.

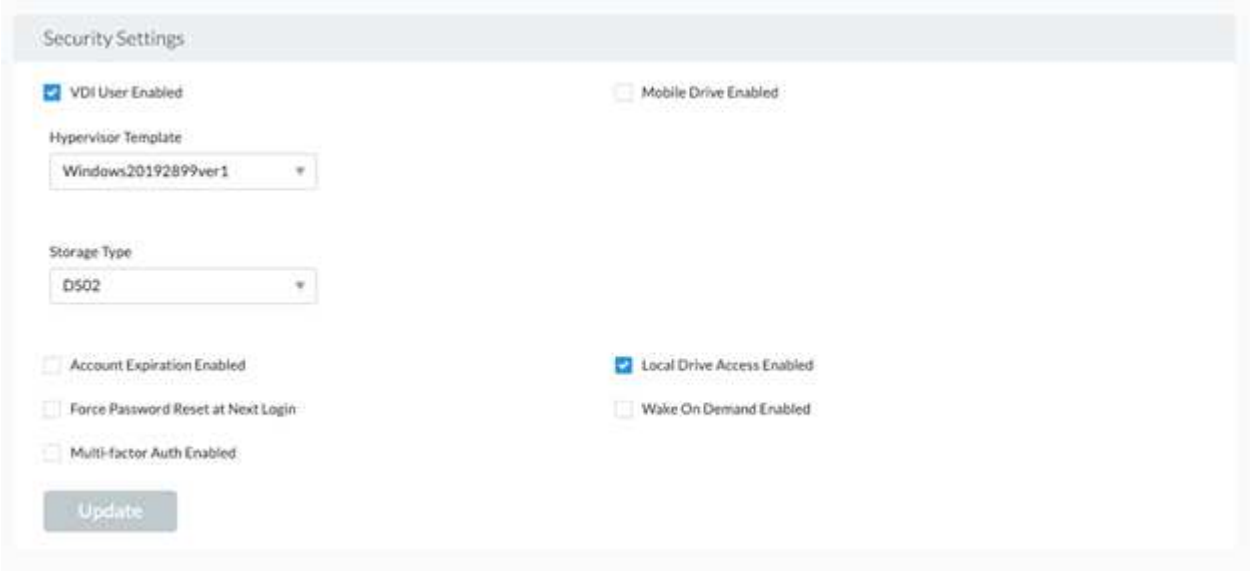

Workspace Management

A workspace consists of a desktop environment; this can be shared remote desktop sessions hosted on-premises or on any supported cloud environment. With Microsoft Azure, the desktop environment can be persistent with Windows Virtual Desktops. Each workspace is associated with a specific organization or client. Options available when creating a new workspace can be seen in the following figure.

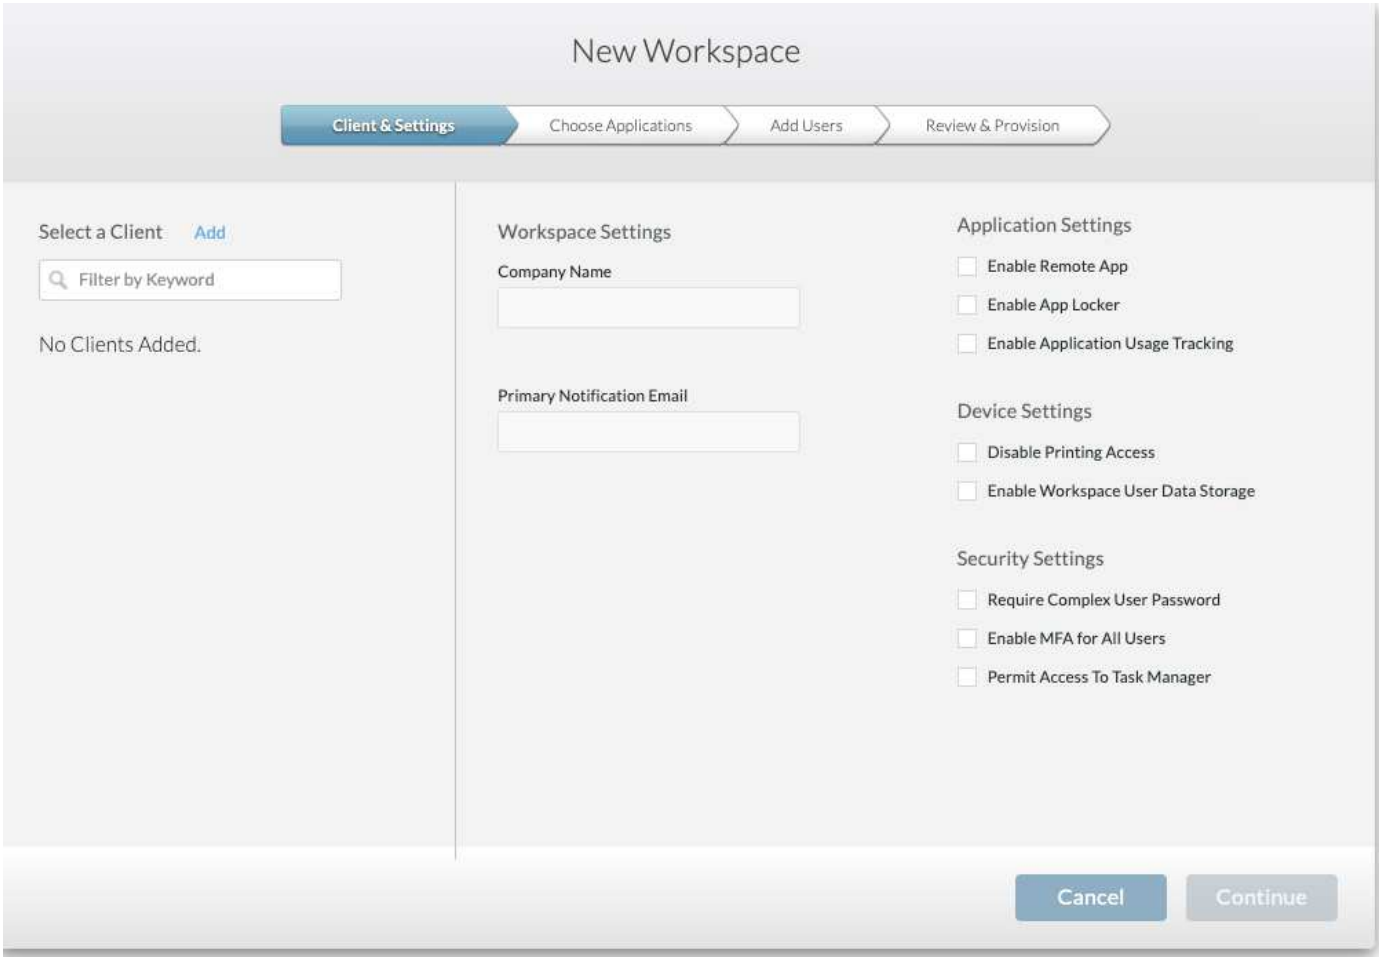

Each workspace is associated with specific deployment.

Workspaces contain associated apps and app services, shared data folders, servers, and a WVD instance. Each workspace can control security options like enforcing password complexity, multifactor authentication, file audits, and so on.

Workspaces can control the workload schedule to power on extra servers, limit the number of users per server, or set the schedule for the resources available for given period (always on/off). Resources can also be configured to wake up on demand.

The workspace can override the deployment VM resource defaults if required. For WVD, WVD host pools (which contains session hosts and app groups) and WVD workspaces can also be managed from the cloud workspace management suite portal. For more info on the WVD host pool, see this video.

Application Management

Task workers can quickly launch an application from the list of applications made available to them. App services publish applications from the Remote Desktop Services session hosts. With WVD, App Groups provide similar functionality from multi-session Windows 10 host pools.

For office workers to power users, the applications that they require can be provisioned manually using a service board, or they can be auto-provisioned using the scripted events feature in NetApp VDS.

For more information, see the NetApp Application Entitlement page.

ONTAP features for Virtual Desktop Service

The following ONTAP features make it attractive choice for use with a virtual desktop service.

- **Scale-out filesystem.** ONTAP FlexGroup volumes can grow to more than 20PB in size and can contain more than 400 billion files within a single namespace. The cluster can contain up to 24 storage nodes, each with a flexible the number of network interface cards depending on the model used.

  User's virtual desktops, home folders, user profile containers, shared data, and so on can grow based on demand with no concern for filesystem limitations.

- **File system analytics.** You can use the XCP tool to gain insights into shared data. With ONTAP 9.8+ and ActiveIQ Unified Manager, you can easily query and retrieve file metadata information and identify cold data.
- **Cloud tiering.** You can migrage cold data to an object store in the cloud or to any S3-compatible storage in your datacenter.
- **File versions.** Users can recover files protected by NetApp ONTAP Snapshot copies. ONTAP Snapshot copies are very space efficient because they only record changed blocks.
- **Global namespace.** ONTAP FlexCache technology allows remote caching of file storage making it easier to manage shared data across locations containing ONTAP storage systems.
- **Secure multi-tenancy support.** A single physical storage cluster can be presented as multiple virtual storage arrays each with its own volumes, storage protocols, logical network interfaces, identity and authentication domain, management users, and so on. Therefore, you can share the storage array across multiple business units or environments, such as test, development, and production.

  To guarantee performance, you can use adaptive QoS to set performance levels based on used or allocated space, and you can control storage capacity by using quotas.

- **VMware integration.** ONTAP tools for VMware vSphere provides a vCenter plug-in to provision datastores, implement vSphere host best practices, and monitor ONTAP resources.

  ONTAP supports vStorage APIs for Array Integration (VAAI) for offloading SCSI/file operations to the storage array. ONTAP also supports vStorage APIs for Storage Awareness (VASA) and Virtual Volumes support for both block and file protocols.

  The Snapcenter Plug-in for VMware vSphere provides an easy way to back up and restore virtual machines using the Snapshot feature on a storage array.

  ActiveIQ Unified Manager provides end-to-end storage network visibility in a vSphere environment. Administrators can easily identify any latency issues that might occur on virtual desktop environments hosted on ONTAP.

- **Security compliance.** With ActiveIQ Unified Manager, you can monitor multiple ONTAP systems with alerts for any policy violations.
- **Multi-protocol support.** ONTAP supports block (iSCSI, FC, FCoE, and NVMe/FC), file (NFSv3, NFSv4.1, SMB2.x, and SMB3.x), and object (S3) storage protocols.
- **Automation support.** ONTAP provides REST API, Ansible, and PowerShell modules to automate tasks with the VDS Management Portal.

Data Management

As a part of deployment, you can choose the file-services method to host the user profile, shared data, and the home drive folder. The available options are File Server, Azure Files, or Azure NetApp Files. However, after deployment, you can modify this choice with the Command Center tool to point to any SMB share. There are various advantages to hosting with NetApp ONTAP. To learn how to change the SMB share, see Change Data Layer.

Global File Cache

When users are spread across multiple sites within a global namespace, Global File Cache can help reduce latency for frequently accessed data. Global File Cache deployment can be automated using a provisioning collection and scripted events. Global File Cache handles the read and write caches locally and maintains file locks across locations. Global File Cache can work with any SMB file servers, including Azure NetApp Files.

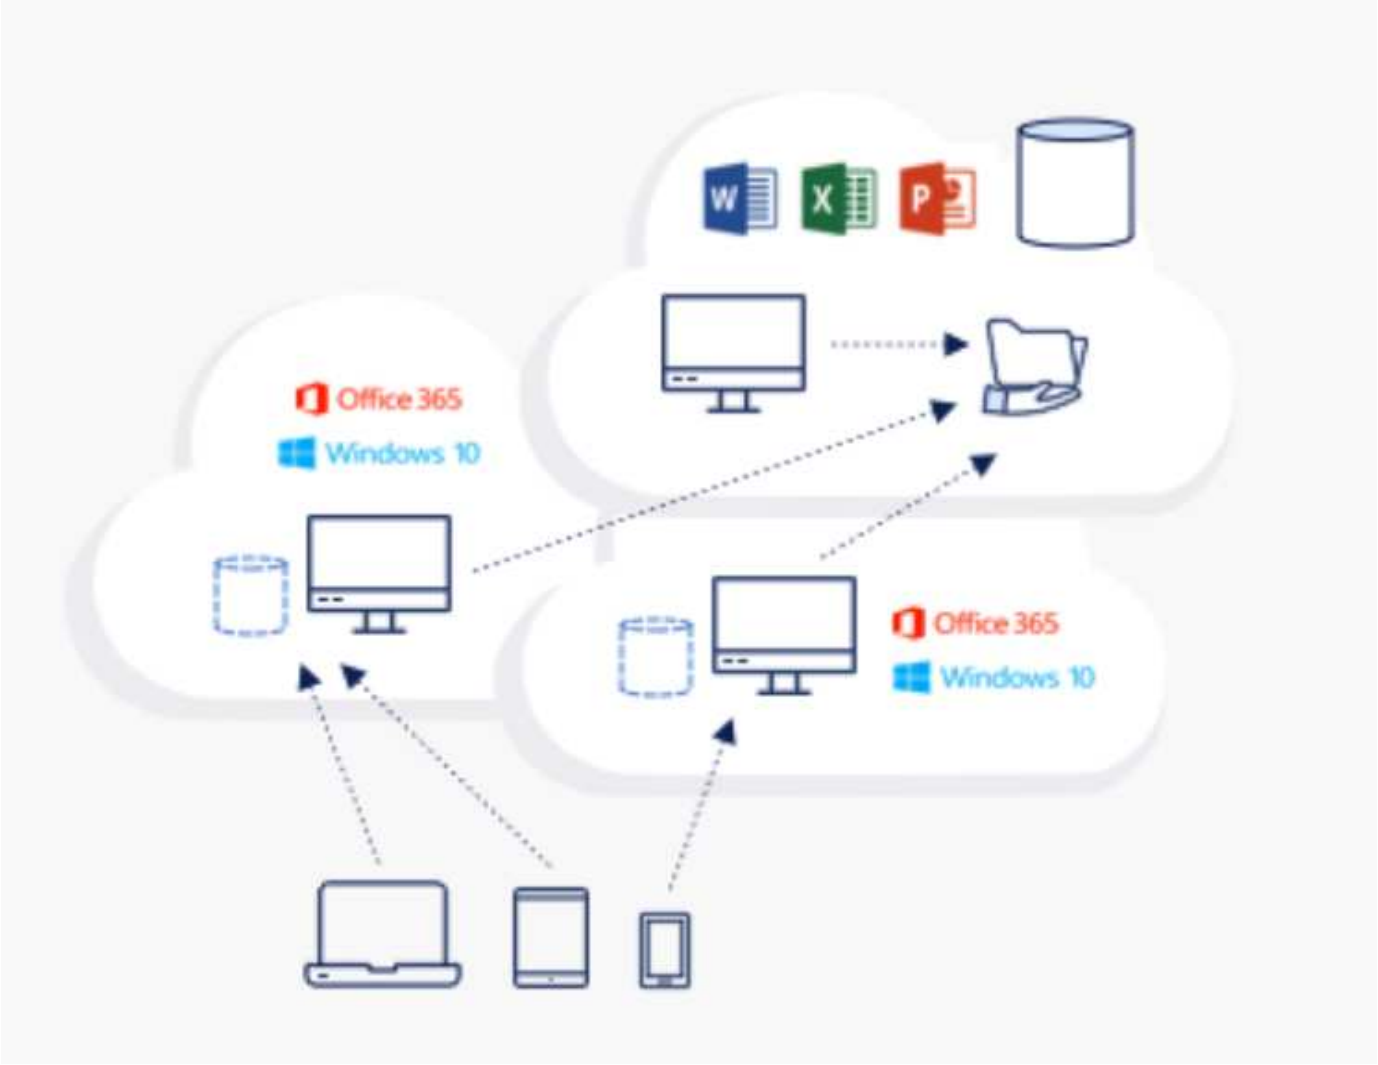

Global File Cache requires the following:

- Management server (License Management Server)
- Core

- Edge with enough disk capacity to cache the data

To download the software and to calculate the disk cache capacity for Edge, see the GFC documentation.

For our validation, we deployed the core and management resources on the same VM at Azure and edge resources on NetApp HCI. Please note that the core is where high-volume data access is required and the edge is a subset of the core. After the software is installed, you must activate the license activated before use. To do so, complete the following steps:

1. Under the License Configuration section, use the link Click Here to complete the license activation. Then register the core.

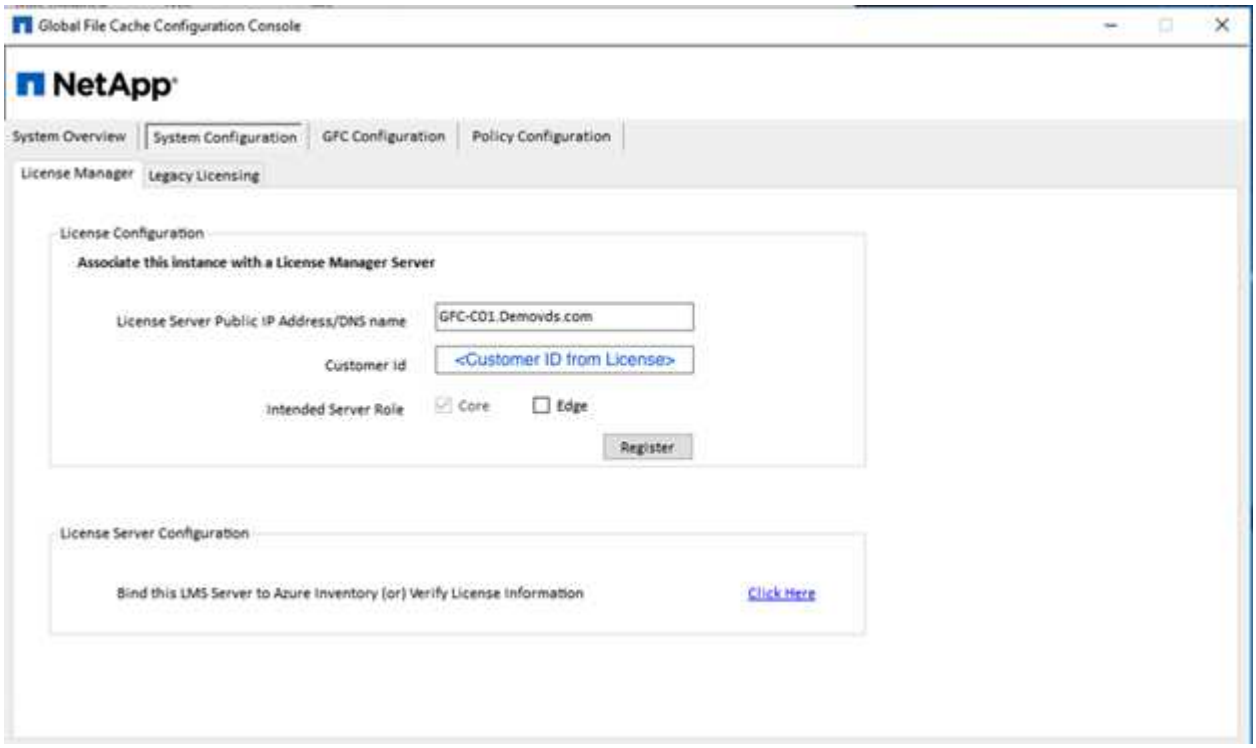

2. Provide the service account to be used for the Global File Cache. For the required permissions for this account, see the GFC documentation.

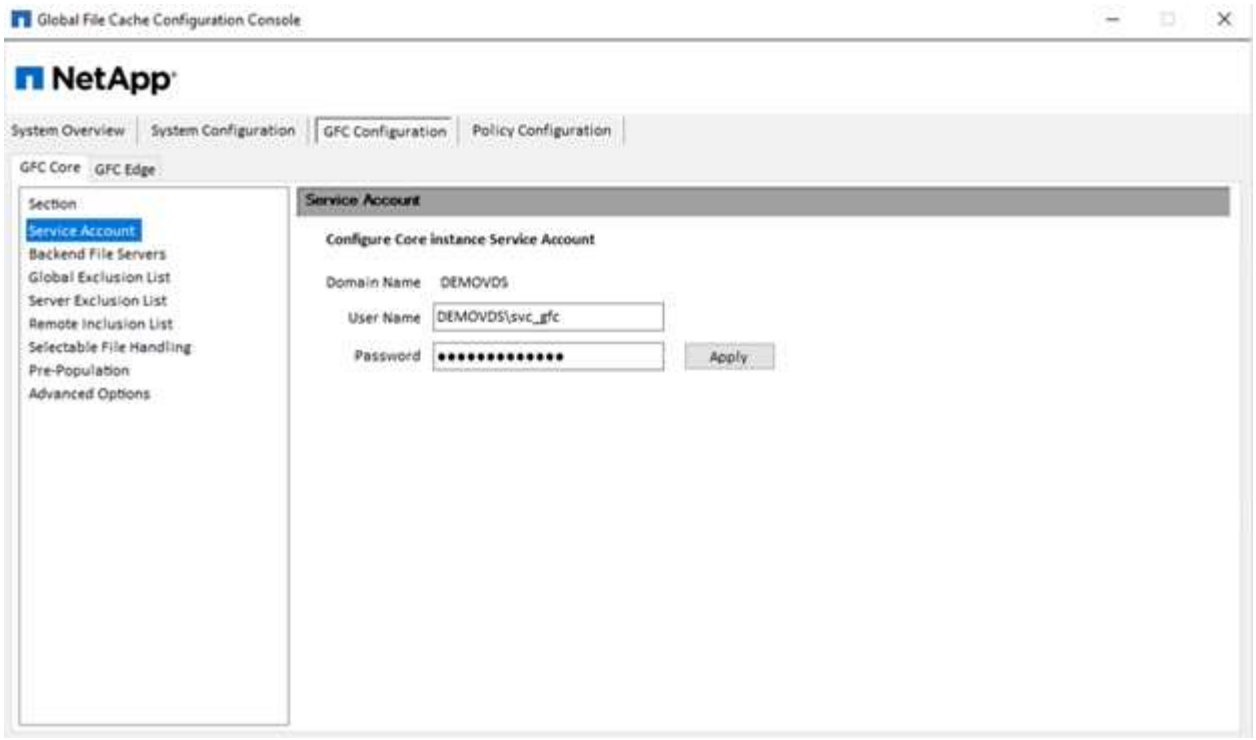

3. Add a new backend file server and provide the file server name or IP.

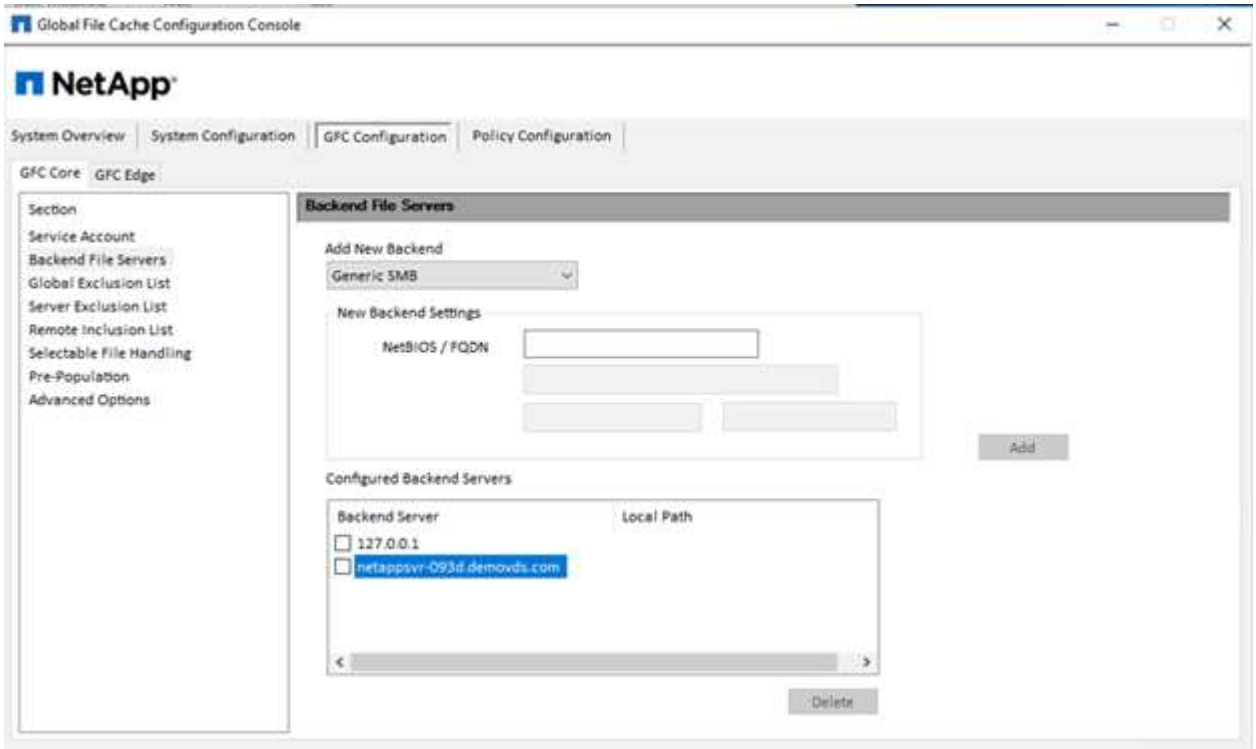

4. On the edge, the cache drive must have the drive letter D. If it does not, use diskpart.exe to select the volume and change drive letter. Register with the license server as edge.

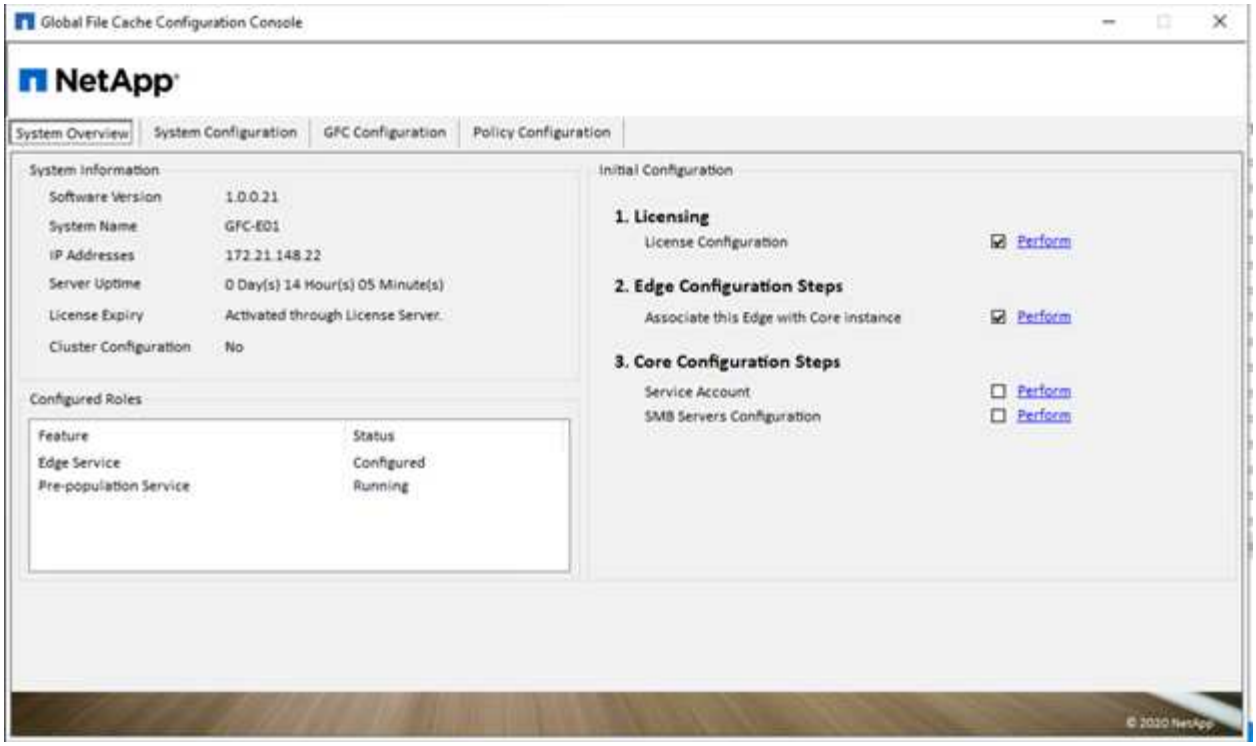

If core auto-configuration is enabled, core information is retrieved from the license management server automatically.

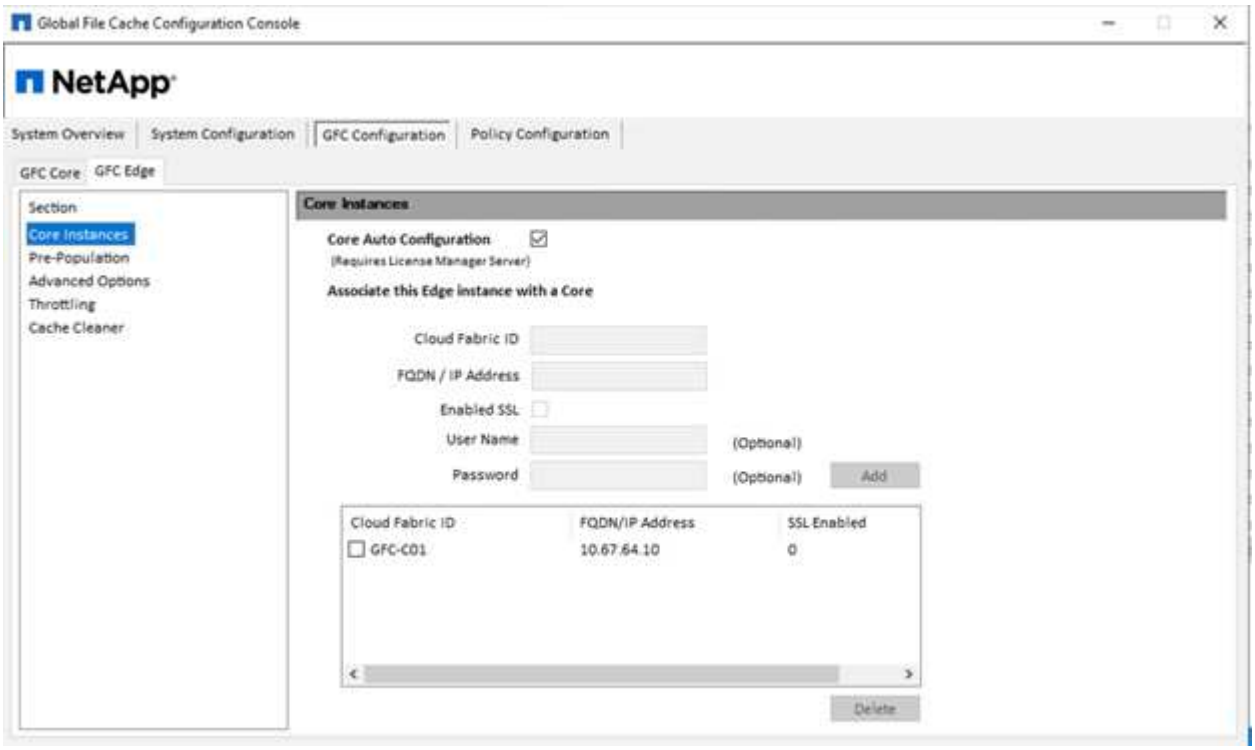

From any client machine, the administrators that used to access the share on the file server can access it with GFC edge using UNC Path `\\<edge server name>\FASTDATA\<core server name>\<backend file server name>\<share name>`. Administrators can include this path in user logonscript or GPO for users drive mapping at the edge location.

To provide transparent access for users across the globe, an administrator can setup the Microsoft Distributed

Filesystem (DFS) with links pointing to file server shares and to edge locations.

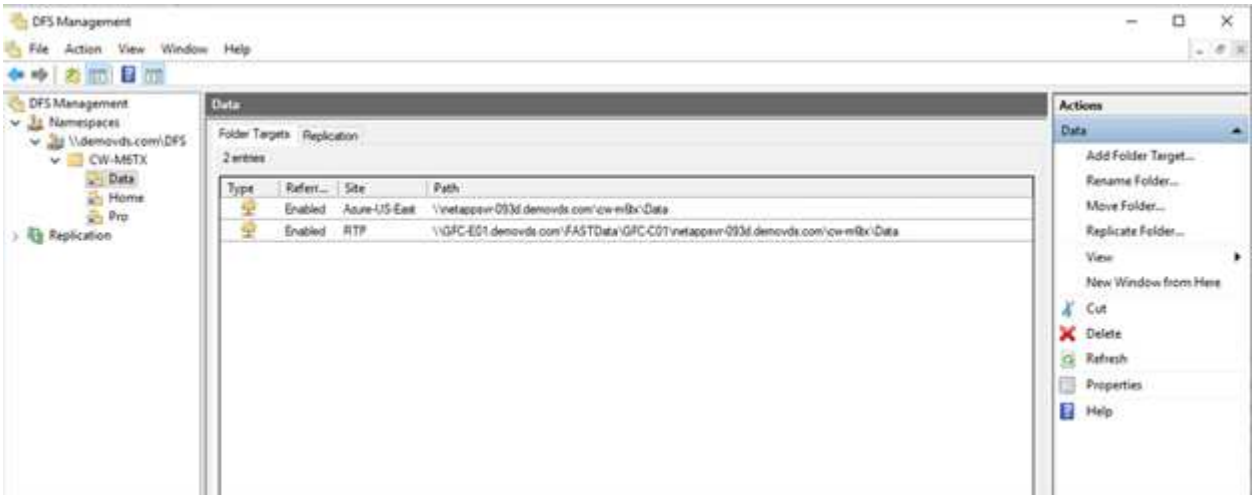

When users log in with Active Directory credentials based on the subnets associated with the site, the appropriate link is utilized by the DFS client to access the data.

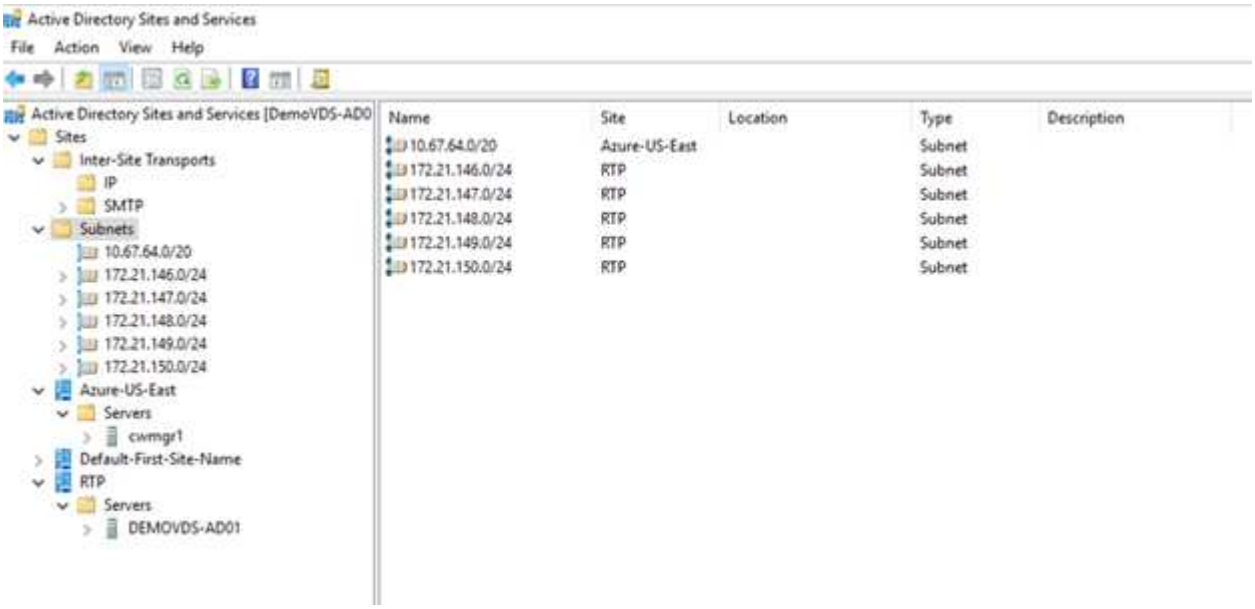

File icons change depending on whether a file is cached; files that are not cached have a grey X on the lower left corner of the icon. After a user in an edge location accesses a file, that file is cached, and the icon changes.

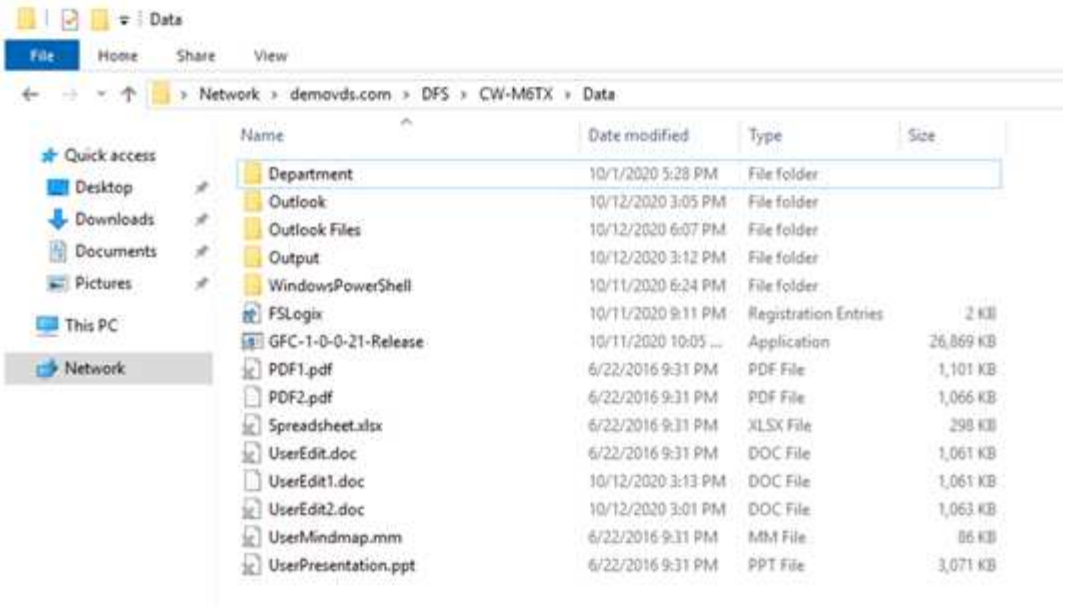

When a file is open and another user is trying to open the same file from an edge location, the user is prompted with the following selection:

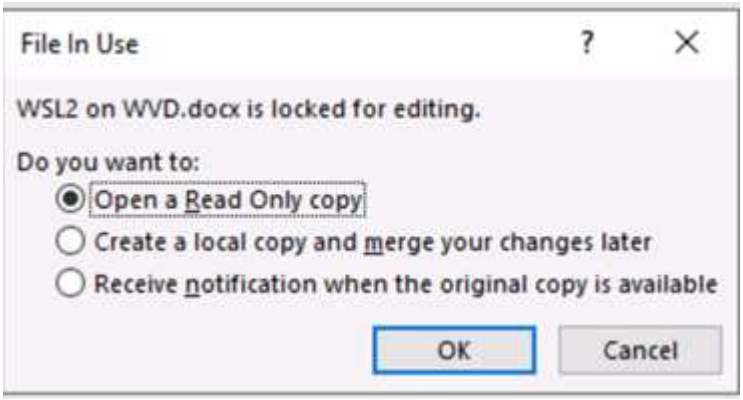

If the user selects the option to receive a notification when the original copy is available, the user is notified as follows:

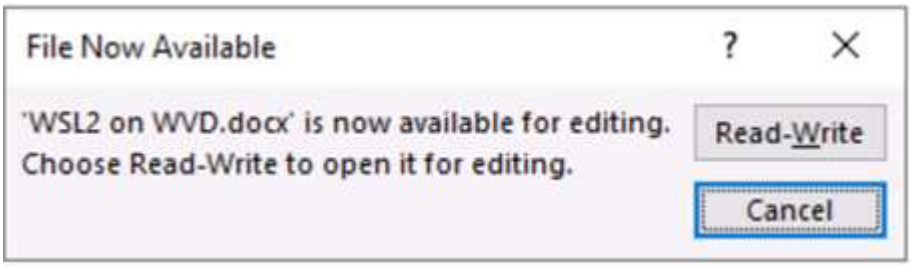

For more information, see this video on Talon and Azure NetApp Files Deployment.

SaaS Backup

NetApp VDS provides data protection for Salesforce and Microsoft Office 365, including Exchange, SharePoint, and Microsoft OneDrive. The following figure shows how NetApp VDS provides SaaS Backup for these data services.

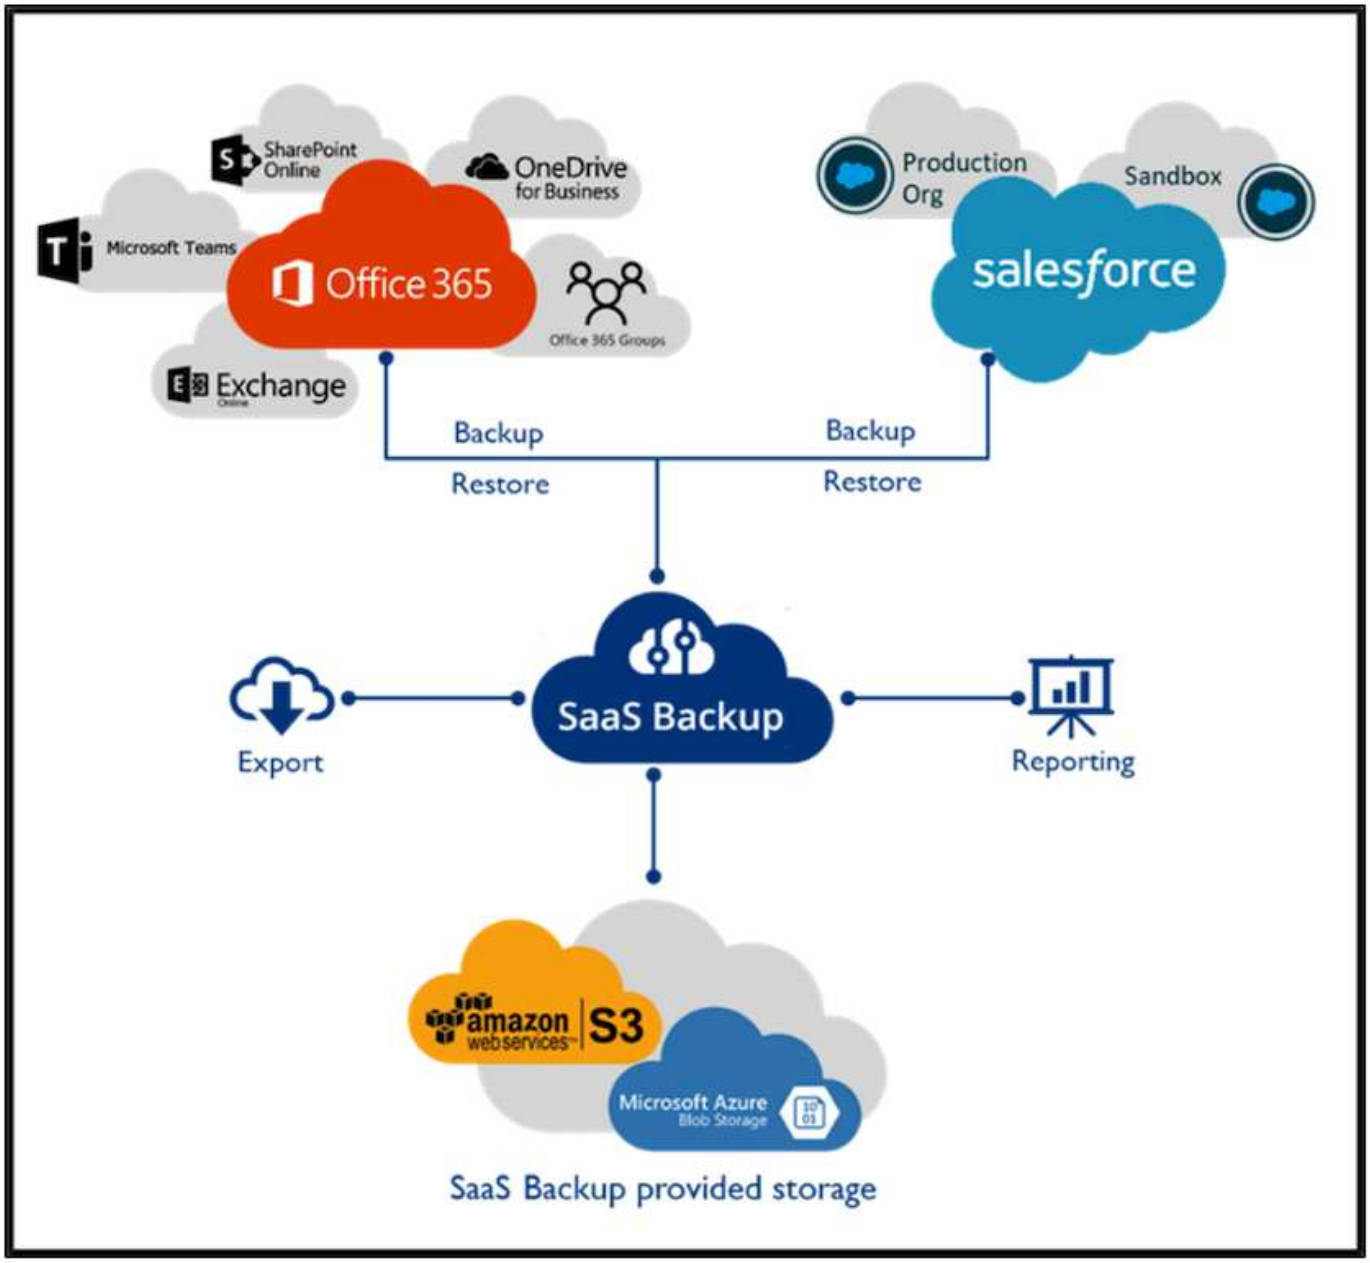

For a demonstration of Microsoft Office 365 data protection, see this video.

For a demonstration of Salesforce data protection, see this video.

Operation management

With NetApp VDS, administrators can delegate tasks to others. They can connect to deployed servers to troubleshoot, view logs, and run audit reports. While assisting customers, helpdesk or level-3 technicians can shadow user sessions, view process lists, and kill processes if required.

For information on VDS logfiles, see the Troubleshooting Failed VDA Actions page.

For more information on the required minimum permissions, see the VDA Components and Permissions page.

If you would like to manually clone a server, see the Cloning Virtual Machines page.

To automatically increase the VM disk size, see the Auto-Increase Disk Space Feature page.

To identify the gateway address to manually configure the client, see the End User Requirements page.

Cloud Insights

NetApp Cloud Insights is a web-based monitoring tool that gives you complete visibility into infrastructure and applications running on NetApp and other third-party infrastructure components. Cloud Insights supports both private cloud and public clouds for monitoring, troubleshooting, and optimizing resources.

Only the acquisition unit VM (can be Windows or Linux) must be installed on a private cloud to collect metrics from data collectors without the need for agents. Agent-based data collectors allow you to pull custom metrics from Windows Performance Monitor or any input agents that Telegraf supports.

The following figure depicts the Cloud Insights VDS dashboard.

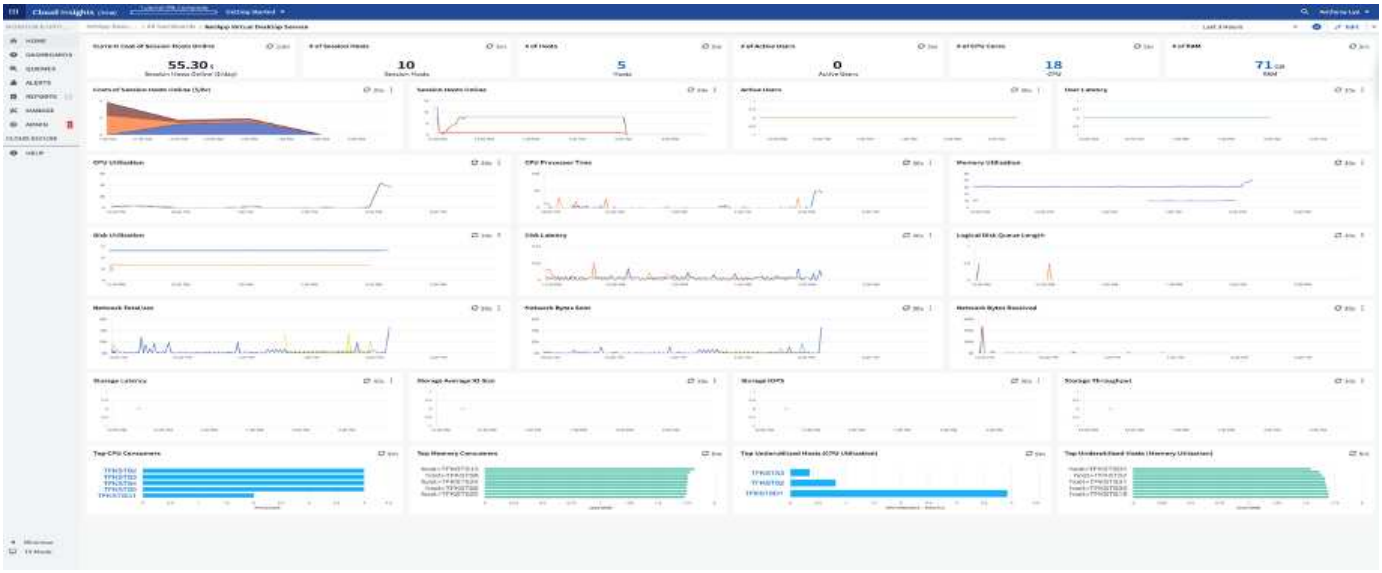

For more info on NetApp Cloud Insights, see this video.

Tools and Logs

This page discusses the DCConfig Tool, TestVdc Tools, and log files.

DCConfig Tool

The DCCconfig tool supports the following hypervisor options for adding a site:

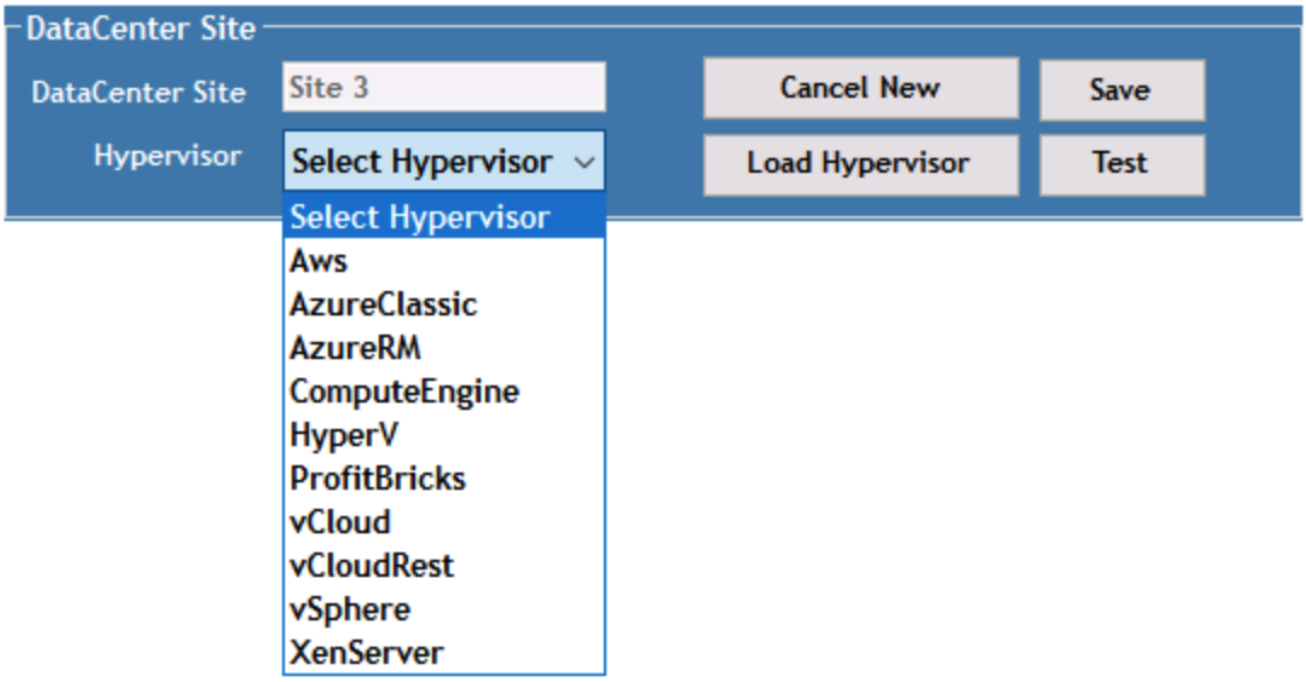

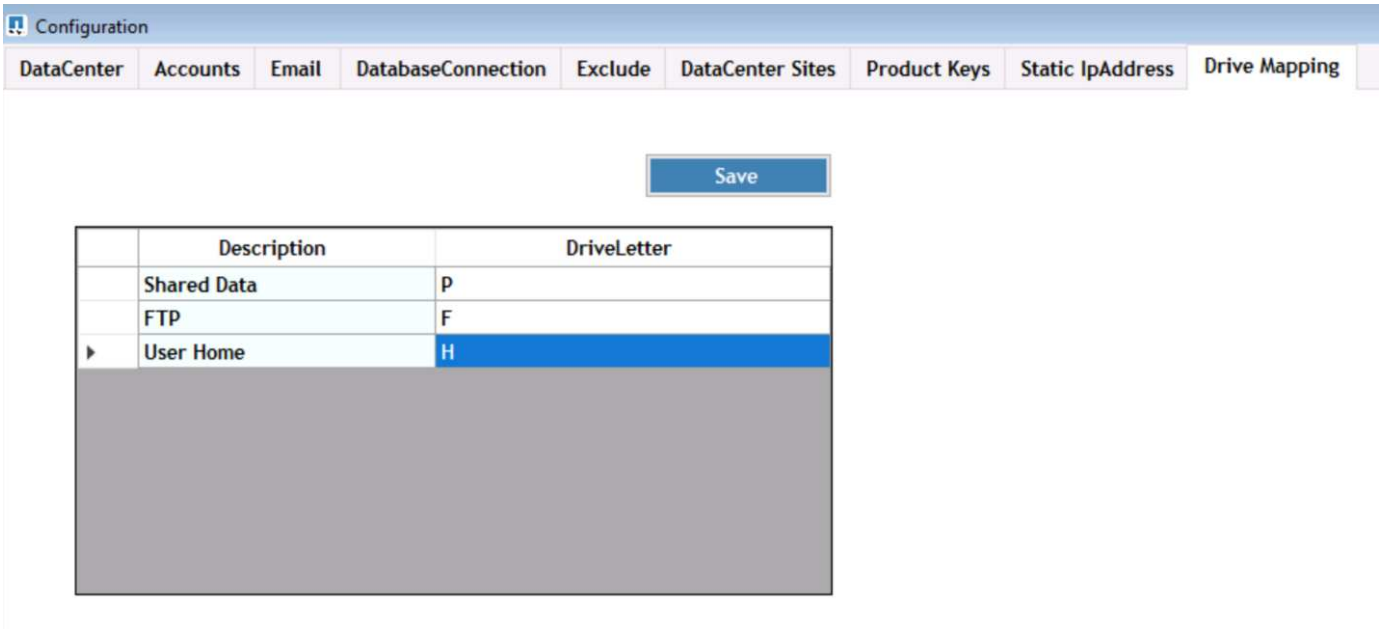

Workspace-specific drive-letter mapping for shared data can be handled using GPO. Professional Services or the support team can use the advanced tab to customize settings like Active Directory OU names, the option to enable or disable deployment of FSLogix, various timeout values, and so on.

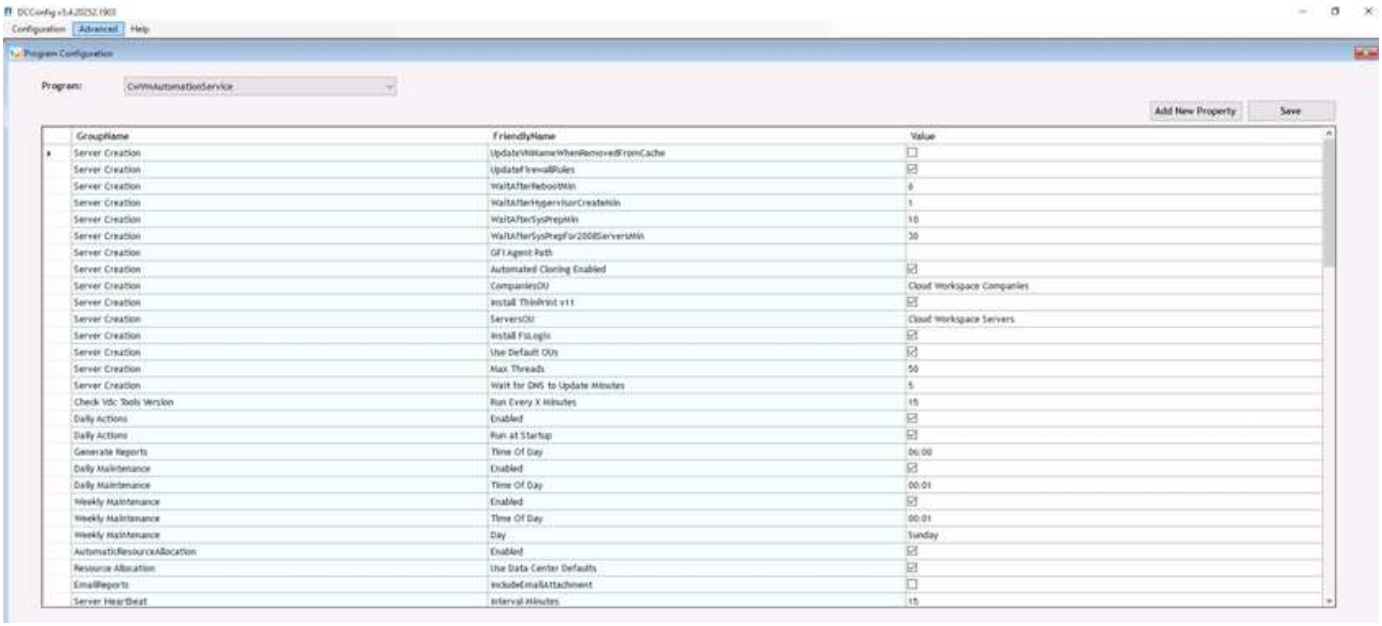

Command Center (Previously known as TestVdc Tools)

To launch Command Center and the required role, see the Command Center Overview.

You can perform the following operations:

- Change the SMB Path for a workspace.

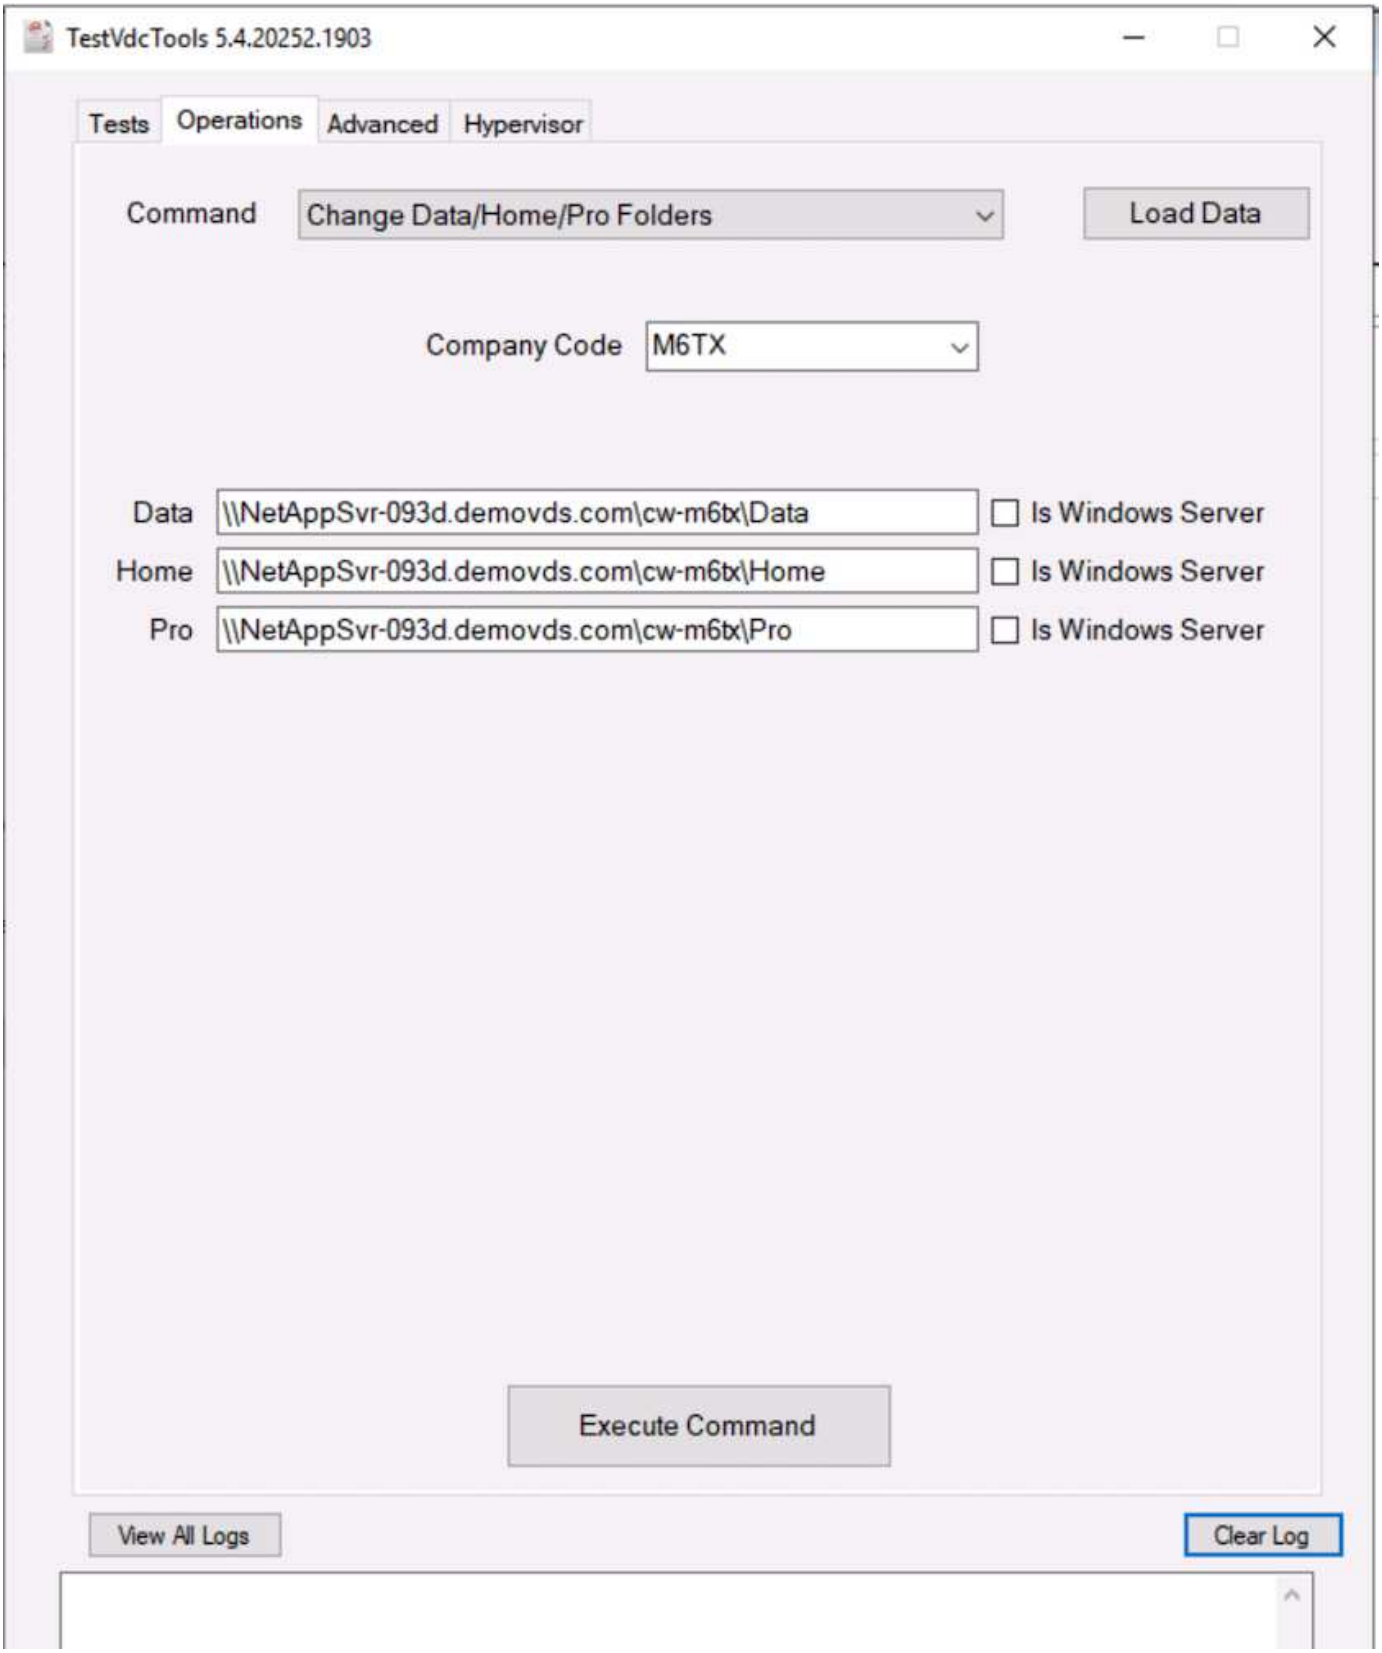

- Change the site for provisioning collection.

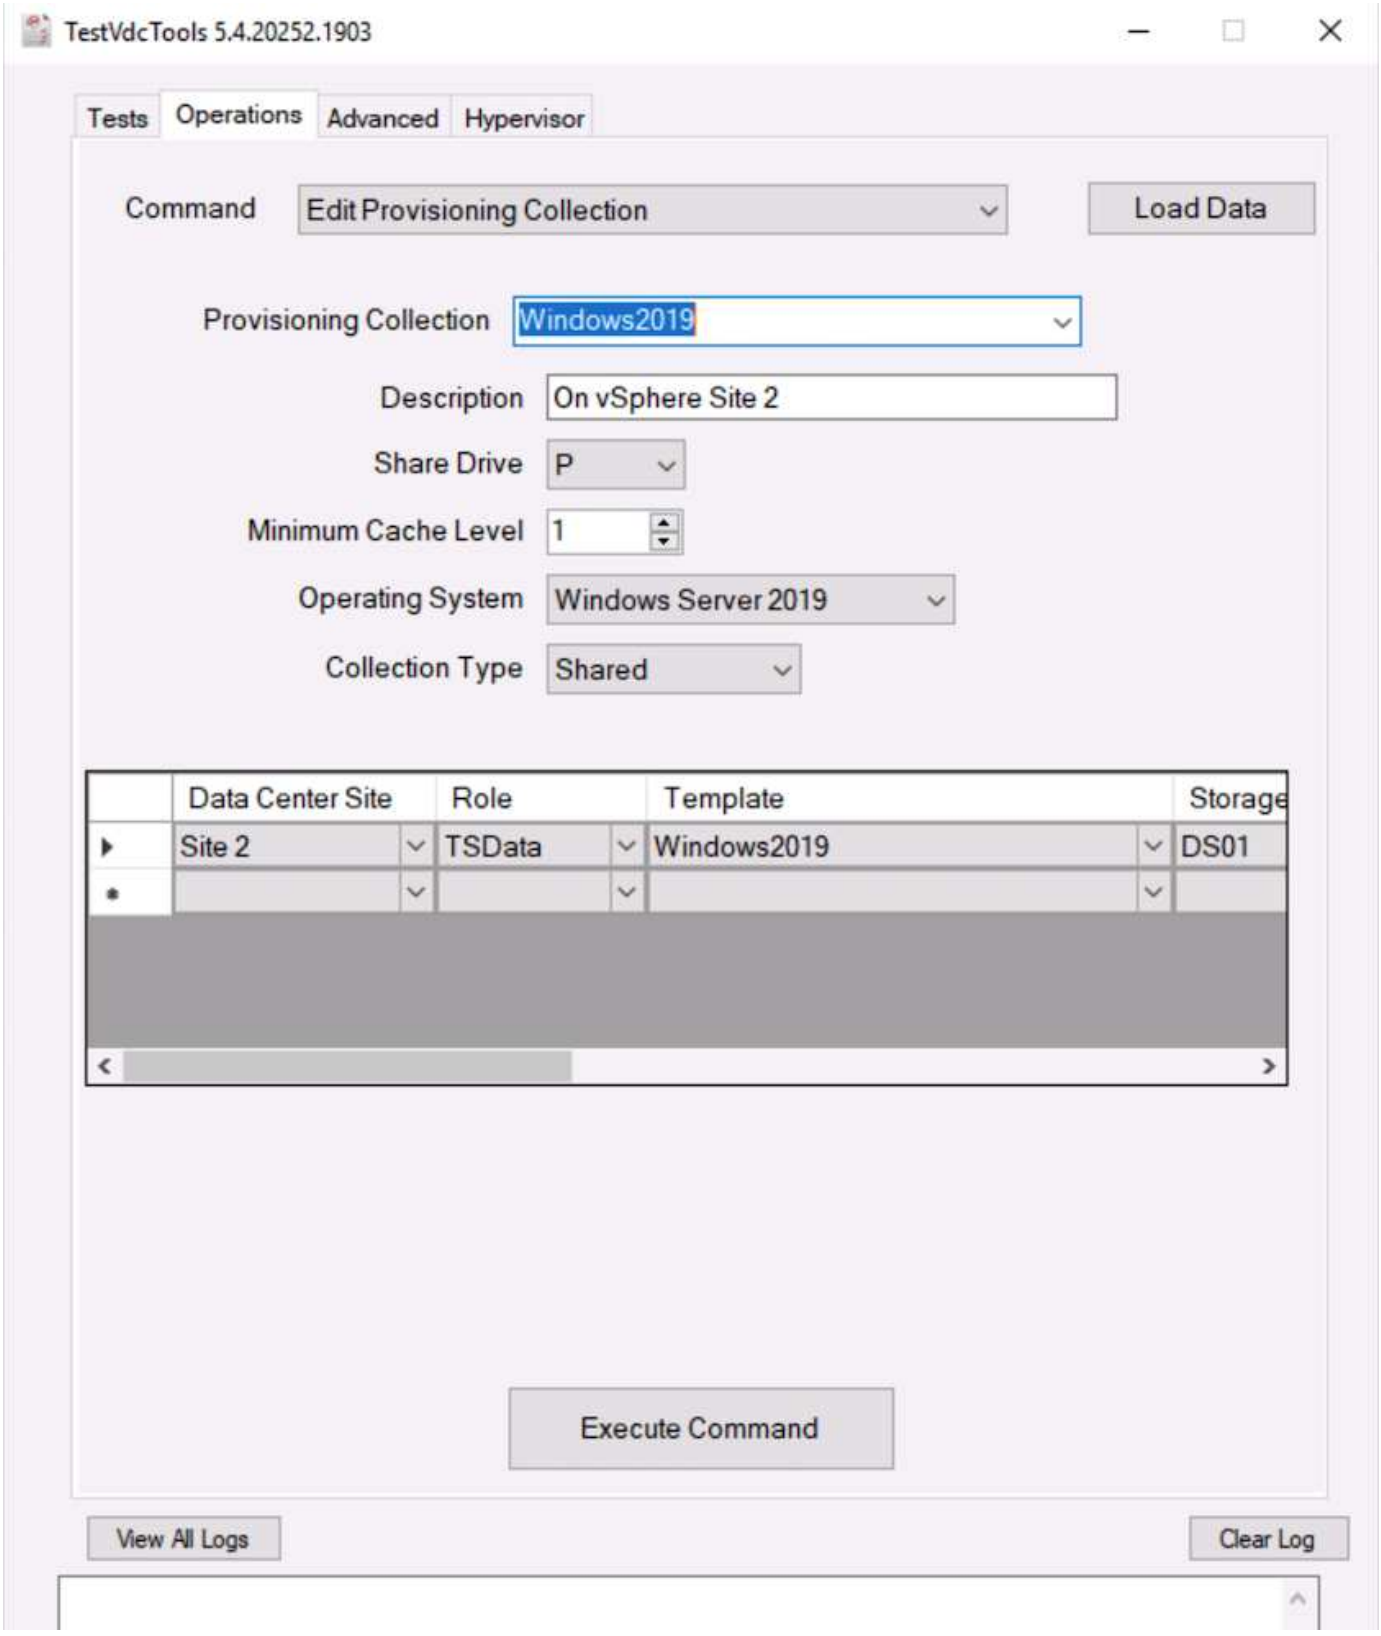
TestVdcTools 5.4.20252.1903
Tests
Operations
Advanced
Hypervisor
Command
Edit Provisioning Collection
Load Data
Provisioning Collection
Windows2019
Description
On vSphere Site 2
Share Drive
P
Minimum Cache Level
1
Operating System
Windows Server 2019
Collection Type
Shared
Data Center Site
Role
Template
Storage
Site 2
TSData
Windows2019
DS01
Execute Command
View All Logs
Clear Log

Log Files

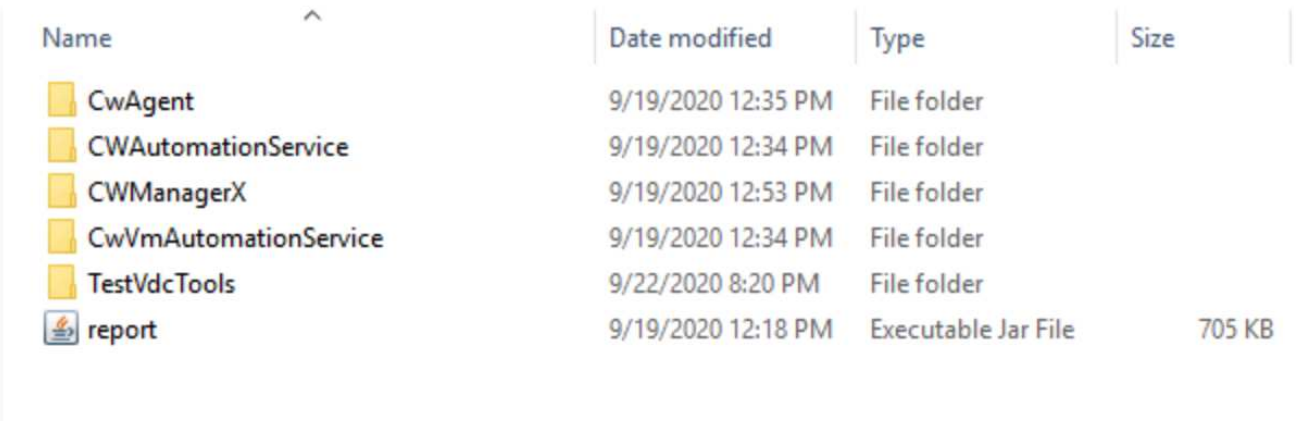

Name	Date modified	Type	Size
CwAgent	9/19/2020 12:35 PM	File folder	
CWAutomationService	9/19/2020 12:34 PM	File folder	
CWManagerX	9/19/2020 12:53 PM	File folder	
CwVmAutomationService	9/19/2020 12:34 PM	File folder	
TestVdcTools	9/22/2020 8:20 PM	File folder	
report	9/19/2020 12:18 PM	Executable Jar File	705 KB

Check automation logs for more info.

GPU considerations

GPUs are typically used for graphic visualization (rendering) by performing repetitive arithmetic calculations. This repetitive compute capability is often used for AI and deep learning use cases.

For graphic intensive applications, Microsoft Azure offers the NV series based on the NVIDIA Tesla M60 card with one to four GPUs per VM. Each NVIDIA Tesla M60 card includes two Maxwell-based GPUs, each with 8GB of GDDR5 memory for a total of 16GB.

An NVIDIA license is included with the NV series.

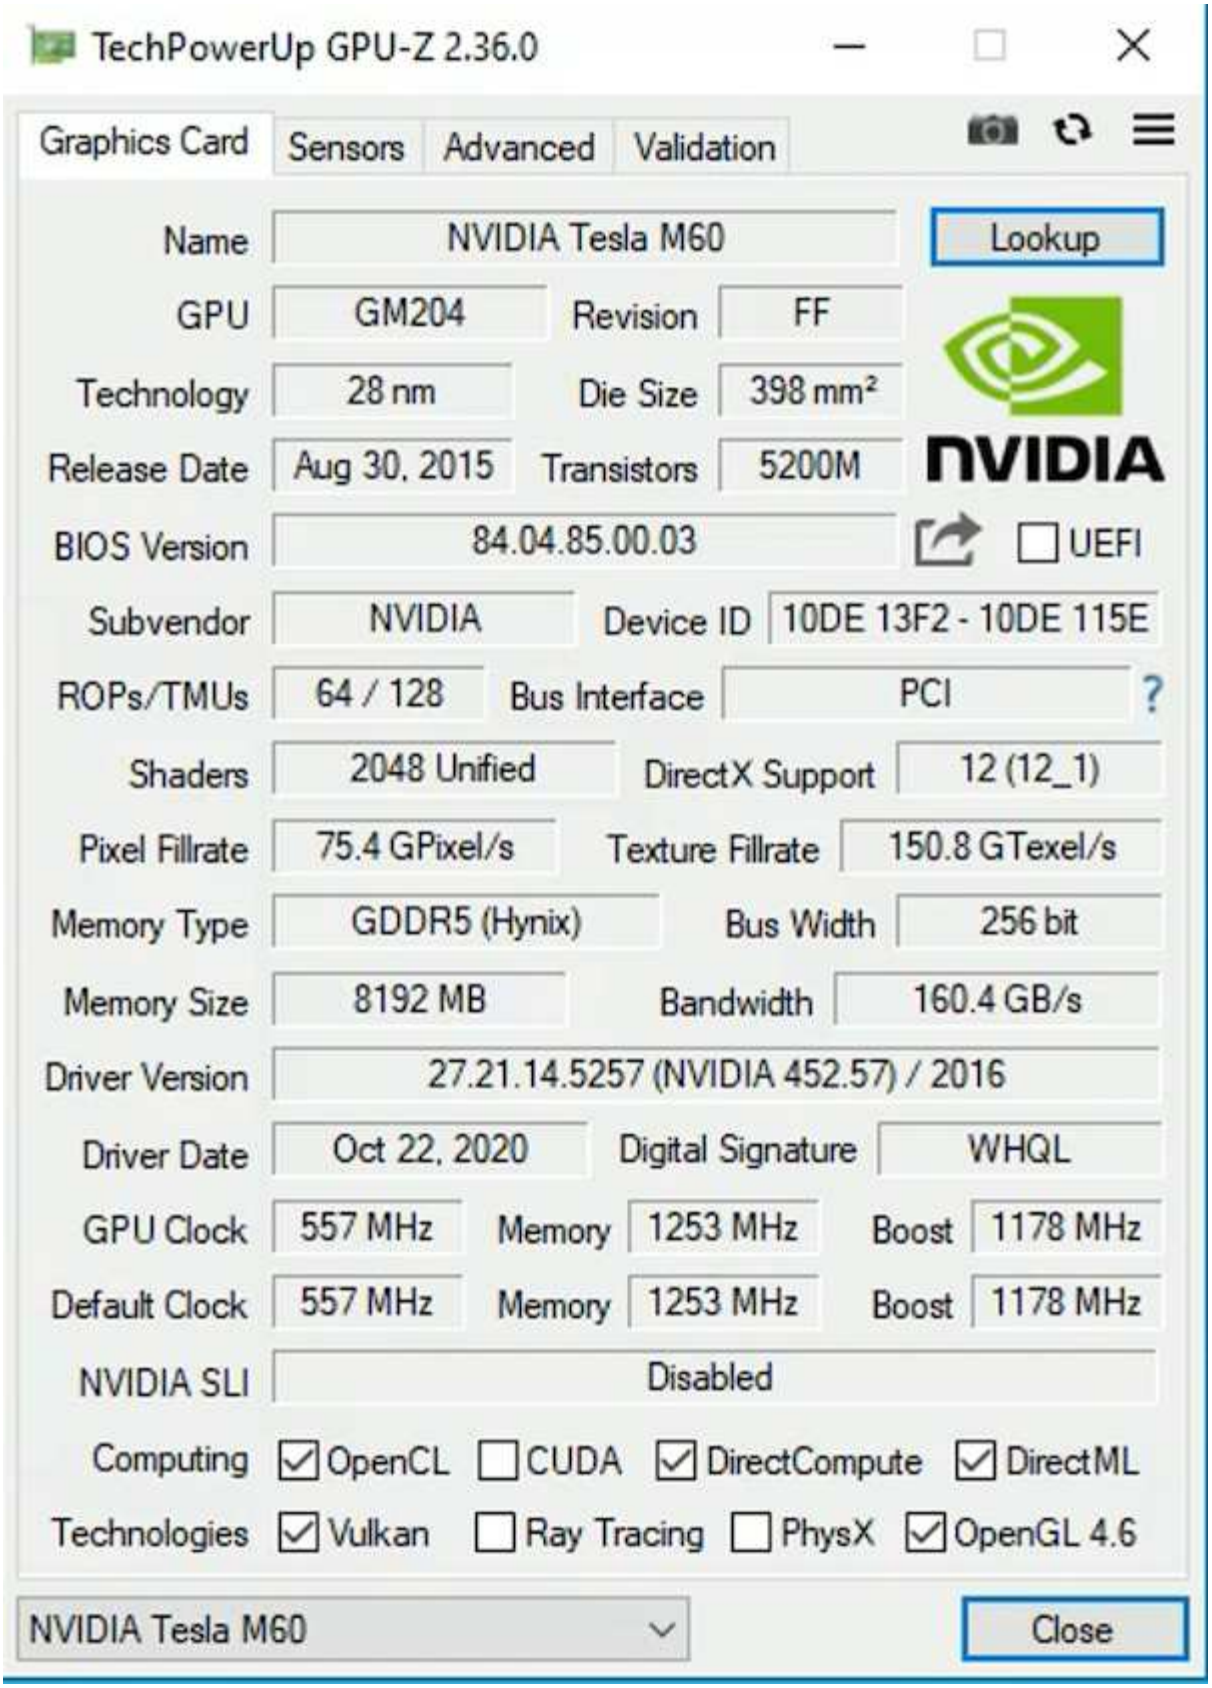

With NetApp HCI, the H615C GPU contains three NVIDIA Tesla T4 cards. Each NVIDIA Tesla T4 card has a Touring-based GPU with 16GB of GDDR6 memory. When used in a VMware vSphere environment, virtual machines are able to share the GPU, with each VM having dedicated frame buffer memory. Ray tracing is available with the GPUs on the NetApp HCI H615C to produce realistic images including light reflections. Please note that you need to have an NVIDIA license server with a license for GPU features.

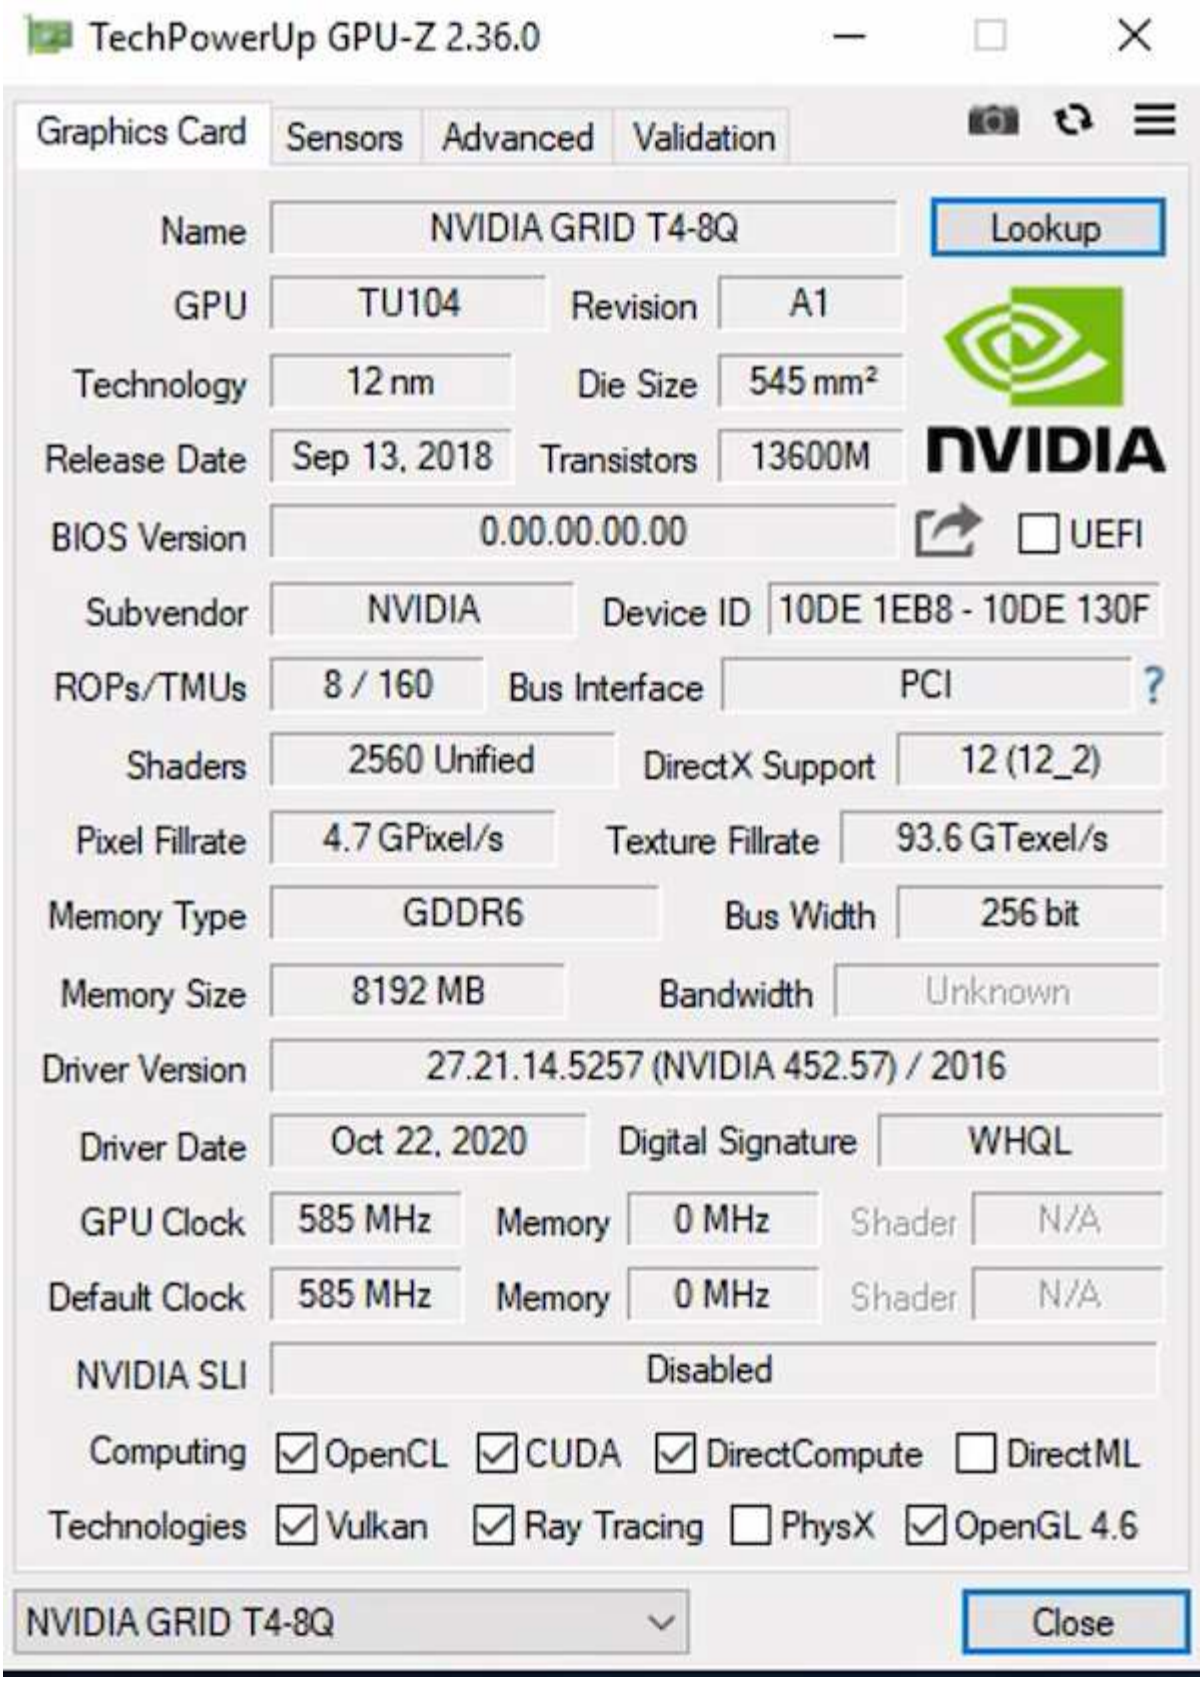

To use the GPU, you must install the appropriate driver, which can be downloaded from the NVIDIA license portal. In an Azure environment, the NVIDIA driver is available as GPU driver extension. Next, the group policies in the following screenshot must be updated to use GPU hardware for remote desktop service sessions. You should prioritize H.264 graphics mode and enable encoder functionality.

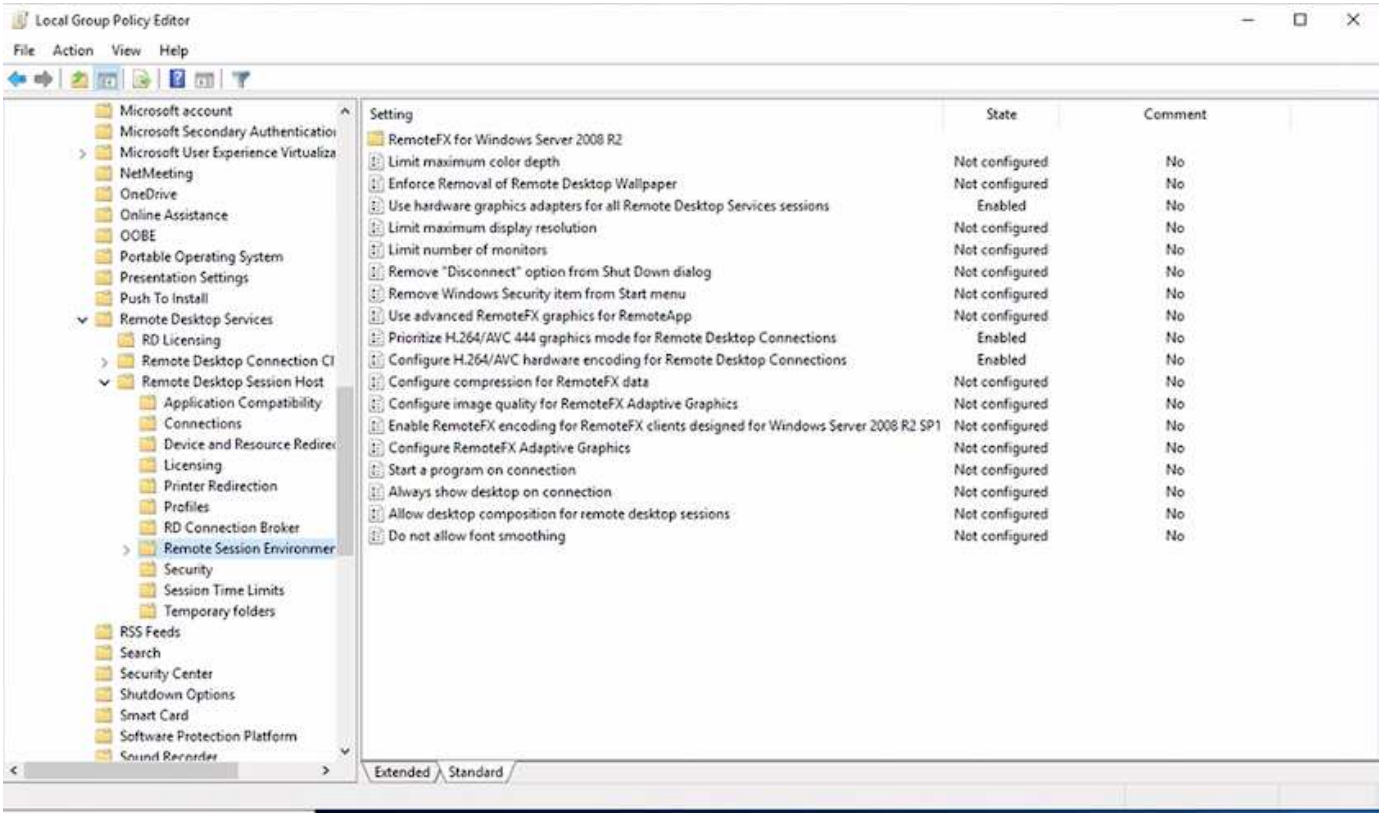

Validate GPU performance monitoring with Task Manager or by using the nvidia-smi CLI when running WebGL samples. Make sure that GPU, memory, and encoder resources are being consumed.

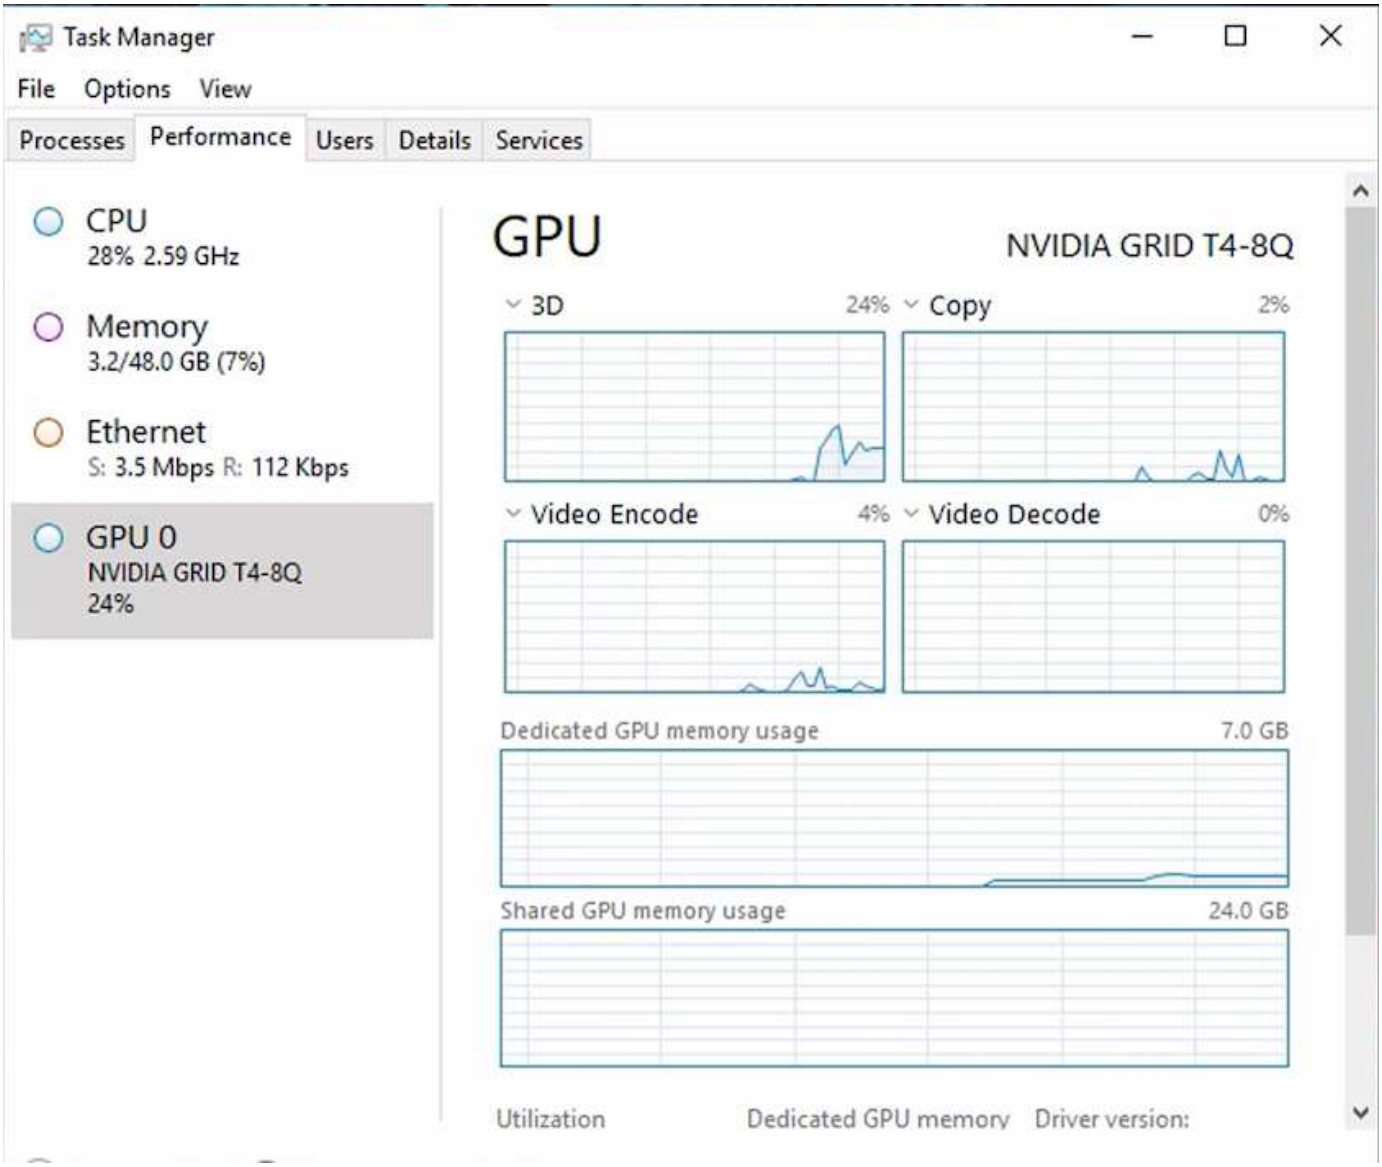

To make sure that the virtual machine is deployed to the NetApp HCI H615C with Virtual Desktop Service, define a site with the vCenter cluster resource that has H615C hosts. The VM template must have the required vGPU profile attached.

For shared multi-session environments, consider allocating multiple homogenous vGPU profiles. However, for high end professional graphics application, it is better to have each VM dedicated to a user to keep VMs isolated.

The GPU processor can be controlled by a QoS policy, and each vGPU profile can have dedicated frame buffers. However, the encoder and decoder are shared for each card. The placement of a vGPU profile on a GPU card is controlled by the vSphere host GPU assignment policy, which can emphasize performance (spread VMs) or consolidation (group VMs).

Solutions for Industry

Graphics workstations are typically used in industries such as manufacturing, healthcare, energy, media and entertainment, education, architecture, and so on. Mobility is often limited for graphics-intensive applications.

To address the issue of mobility, Virtual Desktop Services provide a desktop environment for all types of workers, from task workers to expert users, using hardware resources in the cloud or with NetApp HCI, including options for flexible GPU configurations. VDS enables users to access their work environment from anywhere with laptops, tablets, and other mobile devices.

To run manufacturing workloads with software like ANSYS Fluent, ANSYS Mechanical, Autodesk AutoCAD, Autodesk Inventor, Autodesk 3ds Max, Dassault Systèmes SOLIDWORKS, Dassault Systèmes CATIA, PTC Creo, Siemens PLM NX, and so on, the GPUs available on various clouds (as of Jan 2021) are listed in the following table.

GPU Model	Microsoft Azure	Google Compute (GCP)	Amazon Web Services (AWS)	On-Premises (NetApp HCI)
NVIDIA M60	Yes	Yes	Yes	No
NVIDIA T4	No	Yes	Yes	Yes
NVIDIA P100	No	Yes	No	No
NVIDIA P4	No	Yes	No	No

Shared desktop sessions with other users and dedicated personal desktops are also available. Virtual desktops can have one to four GPUs or can utilize partial GPUs with NetApp HCI. The NVIDIA T4 is a versatile GPU card that can address the demands of a wide spectrum of user workloads.
Each GPU card on NetApp HCI H615C has 16GB of frame buffer memory and three cards per server. The number of users that can be hosted on single H615C server depends on the user workload.

Users/Server	Light (4GB)	Medium (8GB)	Heavy (16GB)
H615C	12	6	3

To determine the user type, run the GPU profiler tool while users are working with applications performing typical tasks. The GPU profiler captures memory demands, the number of displays, and the resolution that users require. You can then pick the vGPU profile that satisfies your requirements.

Virtual desktops with GPUs can support a display resolution of up to 8K, and the utility nView can split a single monitor into regions to work with different datasets.

With ONTAP file storage, you can realize the following benefits:

- A single namespace that can grow up to 20PB of storage with 400 billion of files, without much administrative input
- A namespace that can span the globe with a Global File Cache
- Secure multitenancy with managed NetApp storage
- The migration of cold data to object stores using NetApp FabricPool
- Quick file statistics with file system analytics
- Scaling a storage cluster up to 24 nodes increasing capacity and performance
- The ability to control storage space using quotas and guaranteed performance with QoS limits
- Securing data with encryption
- Meeting broad requirements for data protection and compliance
- Delivering flexible business continuity options

Conclusion

The NetApp Virtual Desktop Service provides an easy-to-consume virtual desktop and application environment with a sharp focus on business challenges. By extending VDS with the on-premises ONTAP environment, you can use powerful NetApp features in a VDS environment, including rapid clone, in-line deduplication, compaction, thin provisioning, and compression. These features save storage costs and improve performance with all-flash storage. With VMware vSphere hypervisor, which minimizes server-provisioning time by using Virtual Volumes and vSphere API for Array integration. Using the hybrid cloud, customers can pick the right environment for their demanding workloads and save money. The desktop session running on-premises can access cloud resources based on policy.

Where to Find Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp Cloud
- NetApp VDS Product Documentation
- Connect your on-premises network to Azure with VPN Gateway
- Azure Portal
- Microsoft Windows Virtual Desktop
- Azure NetApp Files Registration

VMware Horizon

NVA-1132-DESIGN: VMware end-user computing with NetApp HCI

Suresh Thoppay, NetApp

VMware end-user computing with NetApp HCI is a prevalidated, best-practice data center architecture for deploying virtual desktop workloads at an enterprise scale. This document describes the architectural design and best practices for deploying the solution at production scale in a reliable and risk-free manner.

NVA-1132-DESIGN: VMware end-user computing with NetApp HCI

NVA-1129-DESIGN: VMware end-user computing with NetApp HCI and NVIDIA GPUs

Suresh Thoppay, NetApp

VMware end-user computing with NetApp HCI is a prevalidated, best-practice data center architecture for deploying virtual desktop workloads at an enterprise scale. This document describes the architectural design and best practices for deploying the solution at production scale in a reliable and risk-free manner.

NVA-1129-DESIGN: VMware end-user computing with NetApp HCI and NVIDIA GPUs

NVA-1129-DEPLOY: VMware end-user Computing with NetApp HCI and NVIDIA GPUs

Suresh Thoppay, NetApp

VMware end-user Computing with NetApp HCI is a prevalidated, best-practice, data center architecture for deploying virtual desktop workloads at an enterprise scale. This document describes how to deploy the solution at production scale in a reliable and risk-free manner

NVA-1129-DEPLOY: VMware end-user Computing with NetApp HCI and NVIDIA GPUs

NetApp HCI for virtual desktop infrastructure with VMware Horizon 7 - Empower your power users with 3D Graphics

Suresh Thoppay, NetApp

TR-4792 provides guidance on using the NetApp H615C compute node for 3D graphics workloads in a VMware Horizon environment powered by NVIDIA graphics processing units (GPUs) and virtualization software. It also provides the results from the preliminary testing of SPECviewperf 13 for the H615C.

NetApp HCI for virtual desktop infrastructure with VMware Horizon 7 - Empower your power users with 3D Graphics

FlexPod desktop virtualization solutions

Learn more about FlexPod virtualization solutions by reviewing the FlexPod design guides

NetApp All-Flash SAN Array with VMware vSphere 8

For nearly two decades, NetApp ONTAP software has established itself as a premier storage solution for VMware vSphere environments, continually introducing innovative features that simplify management and decrease costs. NetApp is an established leader in the development of NAS and unified storage platforms that offer a wide range of protocol and connectivity support. Alongside this market segment, there are many customers who prefer the simplicity and cost benefits of block-based SAN storage platforms that are focused on doing one job well. NetApp's All-Flash SAN Array (ASA) delivers on that promise with simplicity at scale and with consistent management and automation features for all applications and cloud providers.

Author: Josh Powell - NetApp Solutions Engineering

Solution Overview

Purpose of This Document

In this document we will cover the unique value of using NetApp ASA storage systems with VMware vSphere and provide a technology overview of the NetApp All-Flash SAN Array. In addition, we will look at additional tools for simplifying storage provisioning, data protection, and monitoring of your VMware and ONTAP datacenter.

Deployment sections of this document cover creating vVol datastores with ONTAP Tools for VMware vSphere, and observability for the modern datacenter with NetApp Cloud Insights.

Technology Overview

This solution includes innovative technologies from VMware and NetApp.

VMware vSphere 8.0

VMware vSphere is a virtualization platform that transforms physical resources into pools of compute, network and storage which can be used to satisfy customers' workload and application requirements. The main components of VMware vSphere include:

- **ESXi** - VMware's hypervisor which enables the abstraction of compute processors, memory, network and other resources and makes them available to virtual machines and container workloads.
- **vCenter** - VMware vCenter is a centralized management platform for interacting with compute resources, networking and storage as part of a virtual infrastructure. vCenter plays a crucial role in simplifying the administration of virtualized infrastructure.

New Improvements in vSphere 8.0

vSphere 8.0 introduces some new improvements including, but not limited to:

Scalability - vSphere 8.0 supports the latest Intel and AMD CPUs and has extended limits for vGPU devices, ESXi hosts, VMs per cluster, and VM DirectPath I/O devices.

Distributed Services Engine - Network offloading with NSX to Data Processing Units (DPUs).

Enhanced Device Efficiency - vSphere 8.0 boosts device management capabilities with features like device groups and Device Virtualization Extensions (DVX).

Improved Security - The inclusion of an SSH timeout and TPM Provision Policy strengthens the security framework.

Integration with Hybrid Cloud Services - This feature facilitates seamless transition between on-premises and cloud workloads.

Integrated Kubernetes Runtime - With the inclusion of Tanzu, vSphere 8.0 simplifies container orchestration.

For more information refer to the blog, What's New in vSphere 8?.

VMware Virtual Volumes (vVols)

vVols are a revolutionary new approach to storage management in vSphere clusters, providing simplified management and more granular control of storage resources. In a vVols datastore each virtual disk is a vVol and becomes a native LUN object on the storage system. The integration of the storage system and vSphere takes place through the **VMware API's for Storage Awareness (VASA)** provider and allows the storage system to be aware of the VM data and manage it accordingly. Storage policies, defined in the vCenter Client are used to allocate and manage storage resources.

vVols are a simplified approach to storage management and are preferred in some use cases.

For more information on vVols see the vVols Getting Started Guide.

NVMe over Fabrics

With the release of vSphere 8.0, NVMe is now supported end-to-end with full support for vVols with NVMe-TCP and NVMe-FC.

For detailed information on using NVMe with vSphere refer to About VMware NVMe Storage in the vSphere Storage documentation.

NetApp ONTAP

NetApp ONTAP software has been a leading storage solution for VMware vSphere environments for almost two decades and continues to add innovative capabilities to simplify management while reducing costs. Using ONTAP together with vSphere is a great combination that lets you reduce host hardware and VMware software expenses. You can also protect your data at lower cost with consistent high performance while taking advantage of native storage efficiencies.

Base ONTAP Features

NetApp Snapshot copies: Snapshot copies of a VM or datastore, ensuring no performance impact upon the creation or utilization of a Snapshot. These replicas can serve as restoration points for VMs or as a simple data safeguard. These array-based snapshots are different than VMware (consistency) snapshots. The most straightforward method to generate an ONTAP Snapshot copy is through the SnapCenter Plug-In for VMware vSphere, backing up VMs and datastores.

- **Storage Efficiency** - ONTAP provides real-time and background deduplication and compression, zero-block deduplication, and data compaction.
- **Volume and LUN move** - Allows non-disruptive movement of volumes and LUNs supporting vSphere datastores and vVols within the ONTAP cluster to balance performance and capacity or support non-disruptive maintenance and upgrades.
- **Relocation of Volume and LUN** - ONTAP allows non-disruptive movement of volumes and LUNs that host vSphere datastores and vVols within the ONTAP cluster. This aids in balancing performance and capacity, and allows for non-disruptive upgrades.
- **Quality of Service** - QoS is a feature that enables the management of performance on an individual LUN, volume, or file. It can be used to limit an aggressive VM or to ensure that a critical VM receives sufficient performance resources.
- **Encryption** - NetApp Volume Encryption and NetApp Aggregate Encryption. These options provide a straightforward software-based approach to encrypting data at rest, ensuring its protection.
- **Fabric Pool** - This feature tiers less frequently accessed data to a separate object store, freeing up valuable flash storage. By operating at the block level, it efficiently identifies and tiers colder data, helping to optimize storage resources and reduce costs.
- **Automation** - Simplifies storage and data management tasks by utilizing ONTAP REST APIs for automation, and leveraging Ansible modules for seamless configuration management of ONTAP systems. Ansible modules offer a convenient solution for efficiently managing the configurations of ONTAP systems. The combination of these powerful tools enables the streamlining of workflows and enhancement of the overall management of storage infrastructure.

ONTAP Disaster Recovery Features

NetApp ONTAP provides robust disaster recovery solutions for VMware environments. These solutions leverage SnapMirror replication technologies between primary and secondary storage systems to allow failover

and quick recovery in the case of failure.

Storage Replication Adapter:
The NetApp Storage Replication Adapter (SRA) is a software component that provides integration between NetApp storage systems and VMware Site Recovery Manager (SRM). It facilitates replication of virtual machine (VM) data across NetApp storage arrays, delivering robust data protection and disaster recovery capabilities. The SRA uses SnapMirror and SnapVault to achieve the replication of VM data across disparate storage systems or geographical locations.

The adapter provides asynchronous replication at the storage virtual machine (SVM) level using SnapMirror technology and extends support for both VMFS in SAN storage environments (iSCSI and FC) and NFS in NAS storage environments.

The NetApp SRA is installed as part of ONTAP Tools for VMware vSphere.

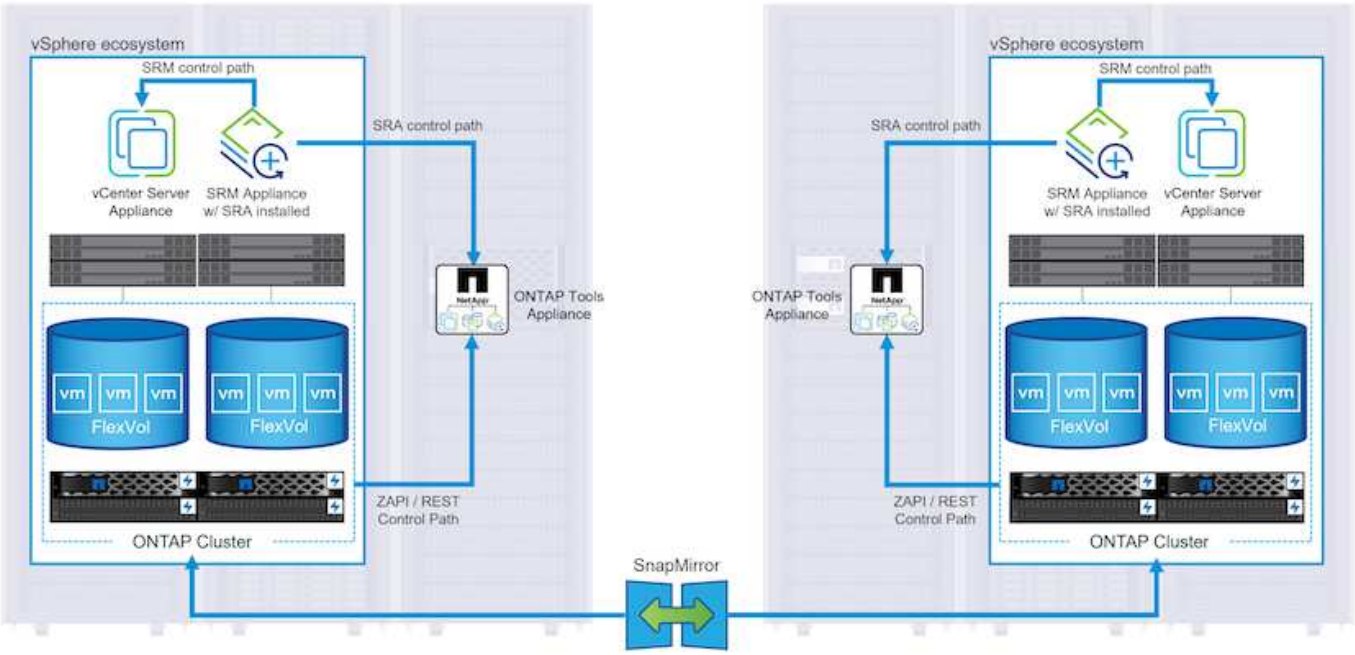

For information on the NetApp Storage Replication Adapter for SRM refer to VMware Site Recovery Manager with NetApp ONTAP.

SnapMirror Business Continuity:
SnapMirror is a NetApp data replication technology that provides synchronous replication of data between storage systems. It allows for the creation of multiple copies of data at different locations, providing the ability to recover data in case of a disaster or data loss event. SnapMirror provides flexibility in terms of replication frequency and allows for the creation of point-in-time copies of data for backup and recovery purposes. SM-BC replicates data at the Consistency Group level.

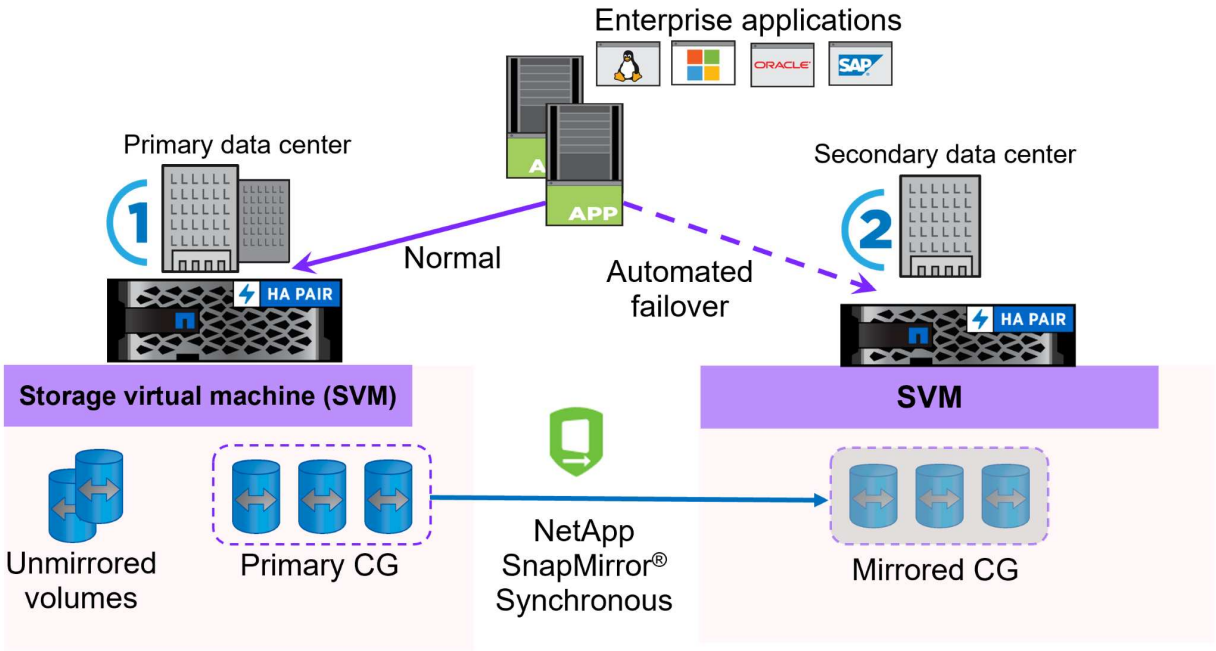

For more information refer to SnapMirror Business Continuity overview.

NetApp MetroCluster:
NetApp MetroCluster is a high-availability and disaster recovery solution that provides synchronous data replication between two geographically dispersed NetApp storage systems. It is designed to ensure continuous data availability and protection in the event of a site-wide failure.

MetroCluster uses SyncMirror to synchronously replicate data just above the RAID level. SyncMirror is designed to efficiently transition between synchronous and asynchronous modes. This allows the primary storage cluster to continue operating in a non-replicated state in situations where the secondary site becomes temporarily inaccessible. SyncMirror will also replicate back to a RPO = 0 state when connectivity is restored.

MetroCluster can operate over IP based networks or using fibre channel.

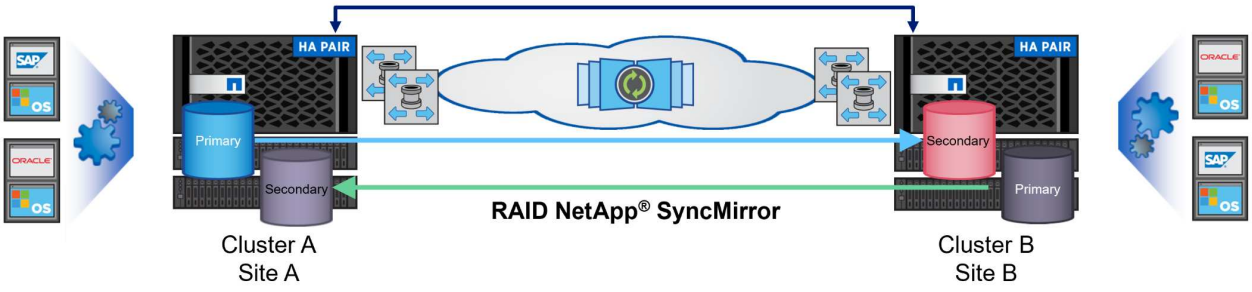

For detailed information on MetroCluster architecture and configuration refer to the MetroCluster documentation site.

ONTAP One Licensing Model

ONTAP One is a comprehensive licensing model that provides access to all features of ONTAP without requiring additional licenses. This includes data protection, disaster recovery, high availability, cloud integration, storage efficiency, performance, and security. Customers with NetApp storage systems licensed with Flash, Core plus Data Protection, or Premium are entitled to ONTAP One licensing, ensuring they can maximize the use of their storage systems.

ONTAP One licensing includes all of the following features:

NVMeoF – Enables the use of NVMe over Fabrics for front end client IO, both NVMe/FC and NVMe/TCP.

FlexClone – Enables rapid creation of space efficient cloning of data based on snapshots.

S3 – Enables the S3 protocol for front end client IO.

SnapRestore – Enables rapid recovery of data from snapshots.

Autonomous Ransomware Protection - Enables the automatic protection of NAS file shares when abnormal filesystem activity is detected.

Multi Tenant Key Manager - Enables the ability to have multiple key managers for different tenants on the system.

SnapLock – Enables the protection of data from modification, deletion or corruption on the system.

SnapMirror Cloud – Enables the replication of system volumes to object targets.

S3 SnapMirror – Enables the replication of ONTAP S3 objects to alternate S3 compatible targets.

NetApp All-Flash SAN Array

The NetApp All-Flash SAN Array (ASA) is a high-performance storage solution designed to meet the demanding requirements of modern data centers. It combines the speed and reliability of flash storage with NetApp's advanced data management features to deliver exceptional performance, scalability, and data protection.

The ASA lineup is comprised of both A-Series and C-Series models.

The NetApp A-Series all-NVMe flash arrays are designed for high-performance workloads, offering ultra-low latency and high resiliency, making them suitable for mission-critical applications.

C-Series QLC flash arrays are aimed at higher-capacity use cases, delivering the speed of flash with the economy of hybrid flash.

For detailed information see the NetApp ASA landing page.

NetApp ASA features

The NetApp All-Flash SAN Array includes the following features:

Performance - The All-Flash SAN Array leverages solid-state drives (SSDs), with an end-to-end NVMe architecture, to provide lightning-fast performance, significantly reducing latency and improving application response times. It delivers consistent high IOPS and low latency, making it suitable for latency-sensitive workloads such as databases, virtualization, and analytics.

Scalability - NetApp All-Flash SAN Arrays are built with a scale-out architecture, allowing organizations to seamlessly scale their storage infrastructure as their needs grow. With the ability to add additional storage nodes, organizations can expand capacity and performance without disruption, ensuring that their storage can keep up with increasing data demands.

Data Management - NetApp's Data ONTAP operating system powers the All-Flash SAN Array, providing a comprehensive suite of data management features. These include thin provisioning, deduplication, compression, and data compaction, which optimize storage utilization and reduce costs. Advanced data protection features like snapshots, replication, and encryption ensure the integrity and security of stored data.

Integration and Flexibility - The All-Flash SAN Array integrates with NetApp's broader ecosystem, enabling seamless integration with other NetApp storage solutions, such as hybrid cloud deployments with NetApp Cloud Volumes ONTAP. It also supports industry-standard protocols like Fibre Channel (FC) and iSCSI, enabling easy integration into existing SAN infrastructures.

Analytics and Automation - NetApp's management software, including NetApp Cloud Insights, provides comprehensive monitoring, analytics, and automation capabilities. These tools enable administrators to gain

insights into their storage environment, optimize performance, and automate routine tasks, simplifying storage management and improving operational efficiency.

Data Protection and Business Continuity - The All-Flash SAN Array offers built-in data protection features such as point-in-time snapshots, replication, and disaster recovery capabilities. These features ensure data availability and facilitate rapid recovery in the event of data loss or system failures.

Protocol Support

The ASA supports all standard SAN protocols including, iSCSI, Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and NVME over fabrics.

iSCSI - NetApp ASA provides robust support for iSCSI, allowing block-level access to storage devices over IP networks. It offers seamless integration with iSCSI initiators, enabling efficient provisioning and management of iSCSI LUNs. ONTAP's advanced features, such as multi-pathing, CHAP authentication, and ALUA support.

For design guidance on iSCSI configurations refer to .

Fibre Channel - NetApp ASA offers comprehensive support for Fibre Channel (FC), a high-speed network technology commonly used in storage area networks (SANs). ONTAP seamlessly integrates with FC infrastructure, providing reliable and efficient block-level access to storage devices. It offers features like zoning, multi-pathing, and fabric login (FLOGI) to optimize performance, enhance security, and ensure seamless connectivity in FC environments.

For design guidance on Fibre Channel configurations refer to the SAN Configuration reference documentation.

NVMe over Fabrics - NetApp ONTAP and ASA support NVMe over fabrics. NVMe/FC enables the use of NVMe storage devices over Fibre Channel infrastructure, and NVMe/TCP over storage IP networks.

For design guidance on NVMe refer to NVMe configuration, support and limitations.

Active-active technology

NetApp All-Flash SAN Arrays allows for active-active paths through both controllers, eliminating the need for the host operating system to wait for an active path to fail before activating the alternative path. This means that the host can utilize all available paths on all controllers, ensuring active paths are always present regardless of whether the system is in a steady state or undergoing a controller failover operation.

Furthermore, the NetApp ASA offers a distinctive feature that greatly enhances the speed of SAN failover. Each controller continuously replicates essential LUN metadata to its partner. As a result, each controller is prepared to take over data serving responsibilities in the event of a sudden failure of its partner. This readiness is possible because the controller already possesses the necessary information to start utilizing the drives that were previously managed by the failed controller.

With active-active pathing, both planned and unplanned takeovers have IO resumption times of 2-3 seconds.

For more information see TR-4968, NetApp All-SAS Array – Data Availability and Integrity with the NetApp ASA.

Storage guarantees

NetApp offers a unique set of storage guarantees with NetApp All-flash SAN Arrays. The unique benefits include:

Storage efficiency guarantee: Achieve high performance while minimizing storage cost with the Storage

Efficiency Guarantee. 4:1 for SAN workloads.

6 Nines (99.9999%) data availability guarantee: Guarantees remediation for unplanned downtime in excess of 31.56 seconds per year.

Ransomware recovery guarantee: Guaranteed data recovery in the event of a ransomware attack.

See the NetApp ASA product portal for more information.

NetApp Plug-ins for VMware vSphere

NetApp storage services are tightly integrated with VMware vSphere through the use of the following plug-ins:

ONTAP Tools for VMware vSphere

The ONTAP Tools for VMware allows administrators to manage NetApp storage directly from within the vSphere Client. ONTAP Tools allows you to deploy and manage datastores, as well as provision vVol datastores.
ONTAP Tools allows mapping of datastores to storage capability profiles which determine a set of storage system attributes. This allows the creation of datastores with specific attributes such as storage performance and QoS.

ONTAP Tools includes the following components:

Virtual Storage Console (VSC): The VSC includes the interface integrated with the vSphere client where you can add storage controllers, provision datastores, monitor performance of datastores, and view and update ESXi host settings.

VASA Provider: The VMware vSphere APIs for Storage Awareness (VASA) Provider for ONTAP send information about storage used by VMware vSphere to the vCenter Server, enabling provisioning of VMware Virtual Volumes (vVols) datastores, creation and use of storage capability profiles, compliance verification, and performance monitoring.

Storage Replication Adapter (SRA): When enabled and used with VMware Site Recovery Manager (SRM), SRA facilitates the recovery of vCenter Server datastores and virtual machines in the event of a failure, allowing configuration of protected sites and recovery sites for disaster recovery.

For more information on NetApp ONTAP tools for VMware see ONTAP tools for VMware vSphere Documentation.

SnapCenter Plug-in for VMware vSphere

The SnapCenter Plug-in for VMware vSphere (SCV) is a software solution from NetApp that offers comprehensive data protection for VMware vSphere environments. It is designed to simplify and streamline the process of protecting and managing virtual machines (VMs) and datastores.

The SnapCenter Plug-in for VMware vSphere provides the following capabilities in a unified interface, integrated with the vSphere client:

Policy-Based Snapshots - SnapCenter allows you to define policies for creating and managing application-consistent snapshots of virtual machines (VMs) in VMware vSphere.

Automation - Automated snapshot creation and management based on defined policies help ensure consistent and efficient data protection.

VM-Level Protection - Granular protection at the VM level allows for efficient management and recovery of individual virtual machines.

Storage Efficiency Features - Integration with NetApp storage technologies provides storage efficiency features like deduplication and compression for snapshots, minimizing storage requirements.

The SnapCenter Plug-in orchestrates the quiescing of virtual machines in conjunction with hardware-based snapshots on NetApp storage arrays. SnapMirror technology is utilized to replicate copies of backups to secondary storage systems including in the cloud.

For more information refer to the SnapCenter Plug-in for VMware vSphere documentation.

BlueXP integration enables 3-2-1 backup strategies that extend copies of data to object storage in the cloud.

For more information on 3-2-1 backup strategies with BlueXP visit 3-2-1 Data Protection for VMware with SnapCenter Plug-in and BlueXP backup and recovery for VMs.

NetApp Cloud Insights

NetApp Cloud Insights simplifies observation of on-prem and cloud infrastructure and provides analytics and troubleshooting capabilities to help solve complex problems. Cloud Insights works by collecting data from a data center environment and sending that data to the cloud. This is done with locally installed software called an Acquisition Unit and with specific collectors enabled for the assets in the data center.

The assets in Cloud Insights can be tagged with annotations that provide a method of organizing and classifying data. Dashboard can be created using a wide variety of widgets for displaying the data and Metric Queries can be created for detailed tabular views of data.

Cloud Insights comes with a large number of ready-made dashboards that help to zero in on specific types of problem areas and categories of data.

Cloud Insights is a heterogeneous tool designed to collect data from a wide range of devices. However, there is a library of templates, called ONTAP Essentials, that makes it easy for NetApp customers to get started quickly.

For detailed information on how to get started with Cloud Insights refer to the NetApp BlueXP and Cloud Insights landing page.

NetApp All-Flash SAN Array with VMware vSphere 8

The ONTAP Tools for VMware allows administrators to manage NetApp storage directly from within the vSphere Client. ONTAP Tools allows you to deploy and manage datastores, as well as provision vVol datastores.
ONTAP Tools allows mapping of datastores to storage capability profiles which determine a set of storage system attributes. This allows the creation of datastores with specific attributes such as storage performance and QoS.

Author: Josh Powell - NetApp Solutions Engineering

Managing Block Storage with ONTAP Tools for VMware vSphere

ONTAP Tools includes the following components:

Virtual Storage Console (VSC): The VSC includes the interface integrated with the vSphere client where you can add storage controllers, provision datastores, monitor performance of datastores, and view and update ESXi host settings.

VASA Provider: The VMware vSphere APIs for Storage Awareness (VASA) Provider for ONTAP send information about storage used by VMware vSphere to the vCenter Server, enabling provisioning of VMware Virtual Volumes (vVols) datastores, creation and use of storage capability profiles, compliance verification, and performance monitoring.

Storage Replication Adapter (SRA): When enabled and used with VMware Site Recovery Manager (SRM), SRA facilitates the recovery of vCenter Server datastores and virtual machines in the event of a failure, allowing configuration of protected sites and recovery sites for disaster recovery.

For more information on NetApp ONTAP tools for VMware see ONTAP tools for VMware vSphere Documentation.

Solution Deployment Overview

In this solution we will demonstrate the use of the ONTAP Tools for VMware vSphere to provision a VMware Virtual Volumes (vVol) datastores and create a virtual machine on a vVol datastore.

In a vVols datastore each virtual disk is a vVol and becomes a native LUN object on the storage system. The integration of the storage system and vSphere takes place through the VMware API's for Storage Awareness (VASA) provider (installed with ONTAP Tools) and allows the storage system to be aware of the VM data and manage it accordingly. Storage policies, defined in the vCenter Client are used to allocate and manage storage resources.

For detailed information on vVols with ONTAP refer to Virtual Volumes vVols) with ONTAP.

This solution covers the following high level steps:

1. Add a storage system in ONTAP Tools.
2. Create a storage capability profile in ONTAP Tools.
3. Create a vVols datastore in ONTAP Tools.
4. Create a VM storage policy in the vSphere client.
5. Create a new virtual machine on the vVol datastore.

Prerequisites

The following components were used in this solution:

1. NetApp All-Flash SAN Array A400 with ONTAP 9.13.
2. iSCSI SVM created on the ASA with network connectivity to the ESXi hosts.
3. ONTAP Tools for VMware vSphere 9.13 (VASA provider enabled by default).
4. vSphere 8.0 cluster (vCenter appliance, and ESXi hosts).

Solution Deployment

Create a vVols datastore in ONTAP Tools

To create a vVols datastore in ONTAP Tools complete the following steps:

Add a storage system to ONTAP Tools.

1. Access NetApp ONTAP Tools by selecting it from the main menu in the vSphere client.

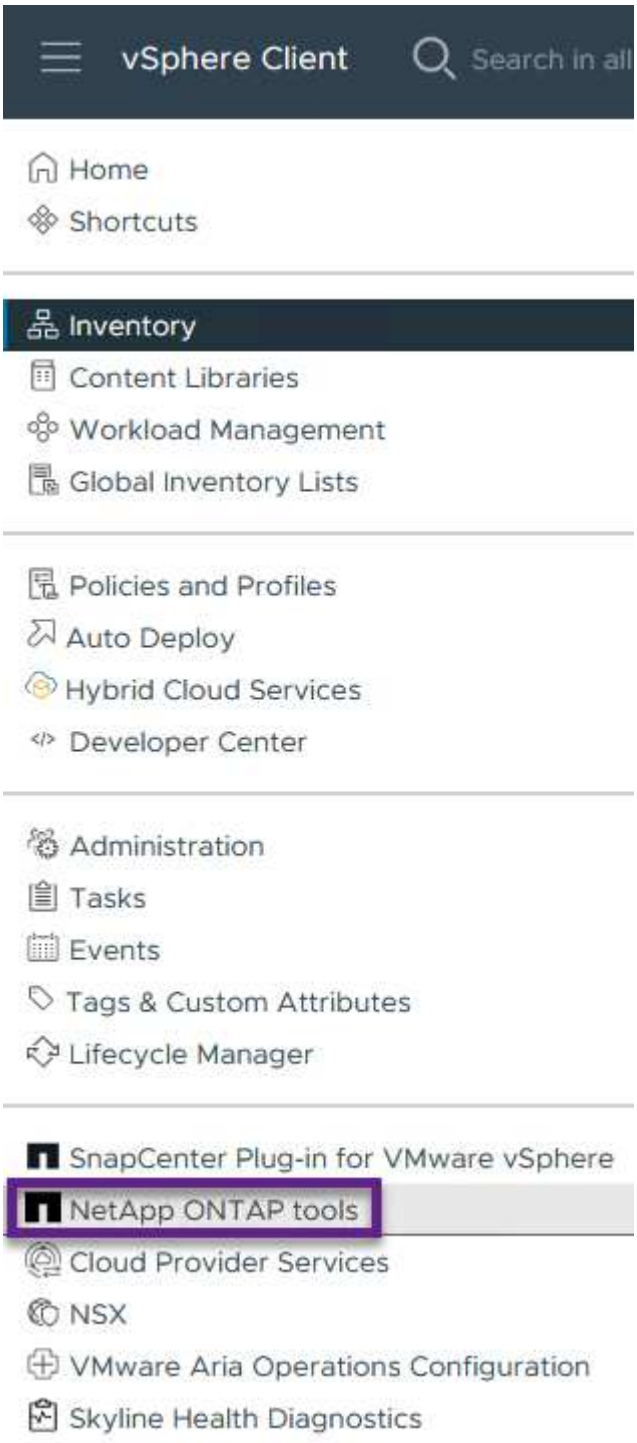

2. In ONTAP Tools select **Storage Systems** from the left hand menu and then press **Add**.

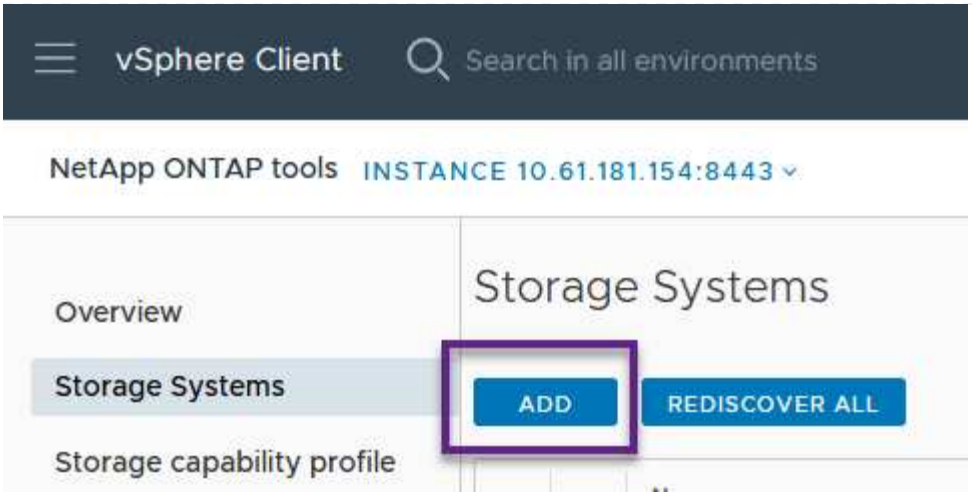

3. Fill out the IP Address, credentials of the storage system and the port number. Click on **Add** to start the discovery process.

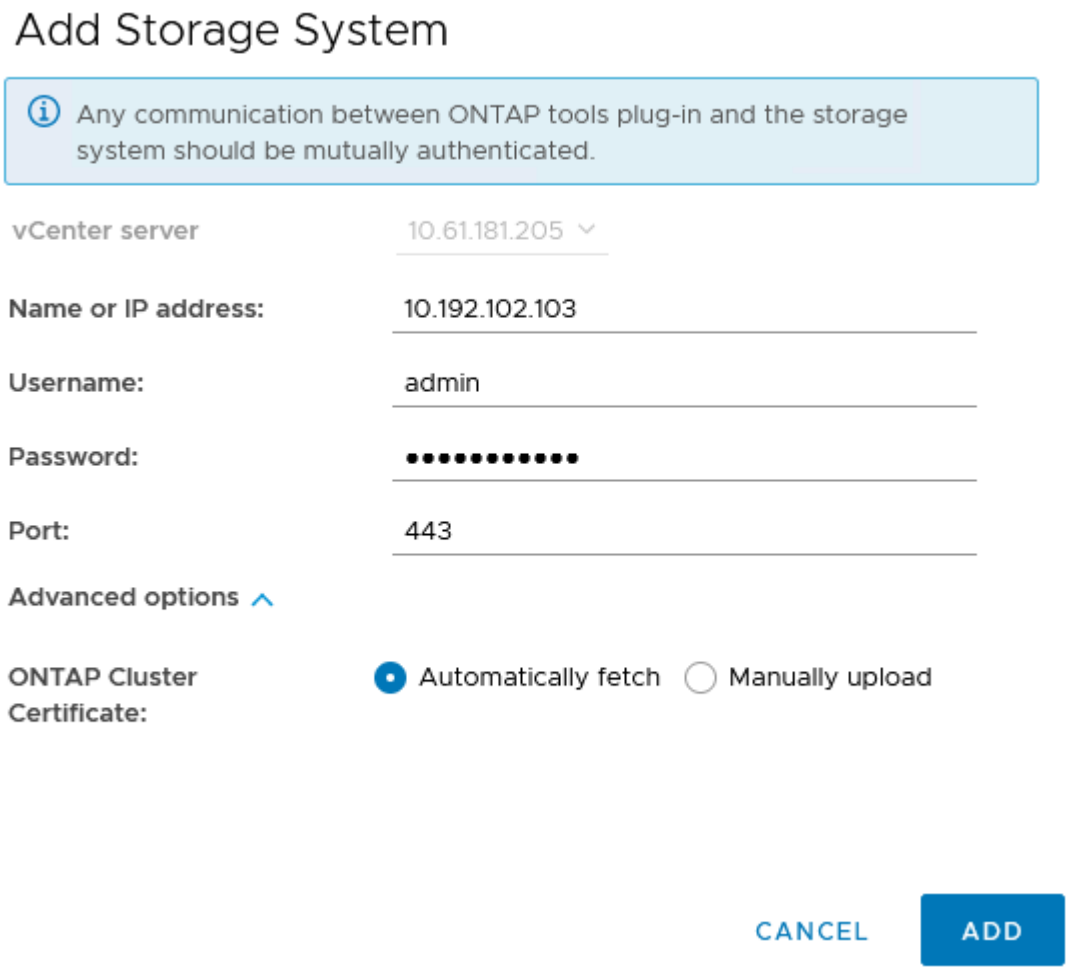

Create a storage capability profile in ONTAP Tools

Storage capability profiles describe the features provided by a storage array or storage system. They include quality of service definitions and are used to select storage systems that meet the parameters defined in the profile.

To create a storage capability profile in ONTAP Tools complete the following steps:

1. In ONTAP Tools select **Storage capability profile** from the left hand menu and then press **Create**.

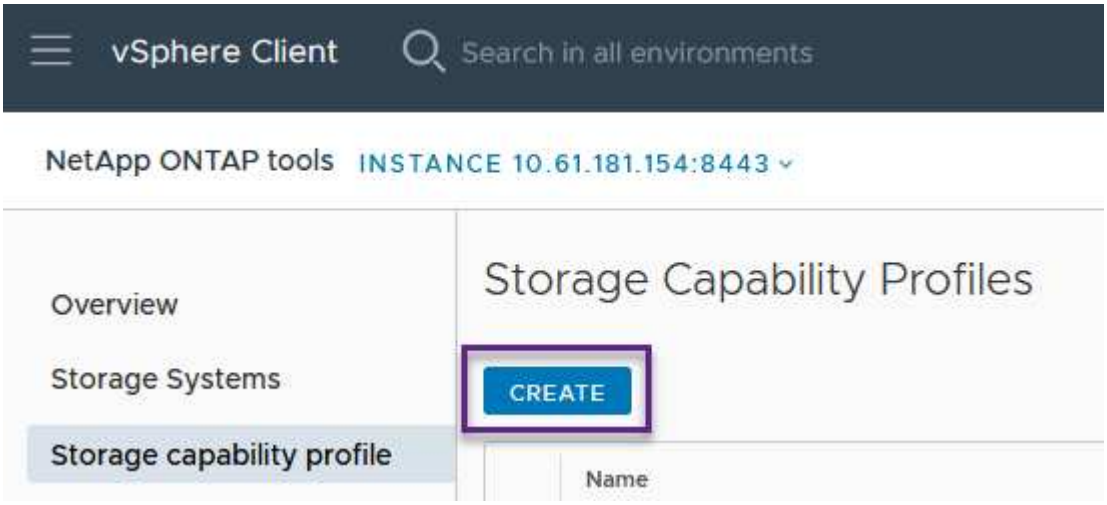

2. In the **Create Storage Capability profile** wizard provide a name and description of the profile and click on **Next**.

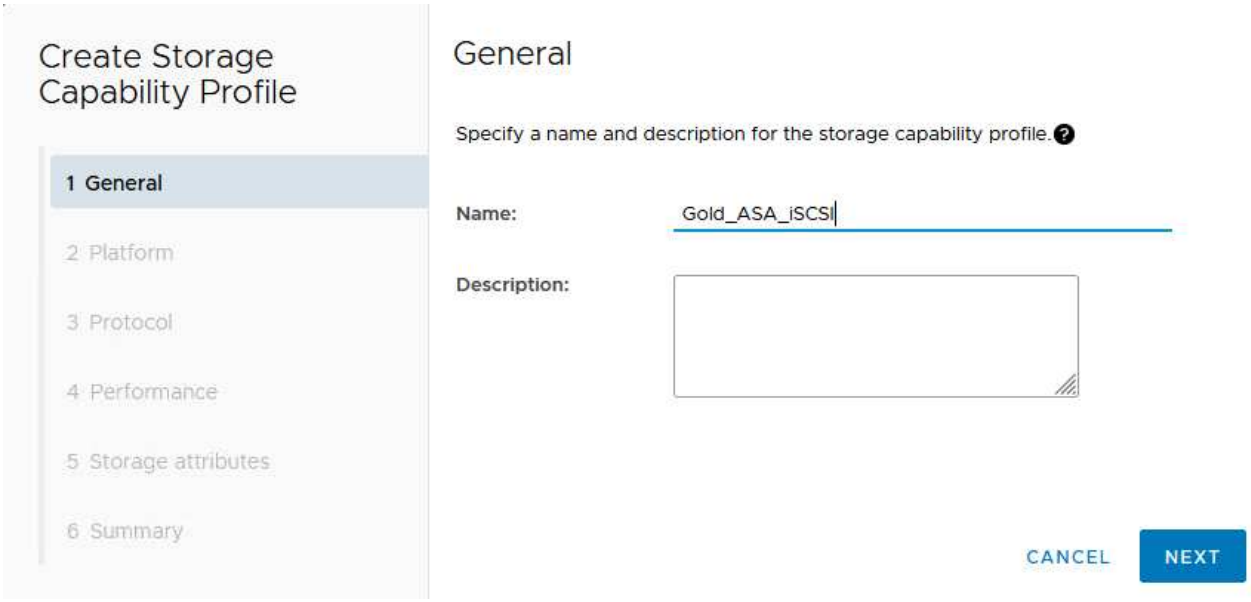

3. Select the platform type and to specify the storage system is to be an All-Flash SAN Array set **Asymmetric** to false.

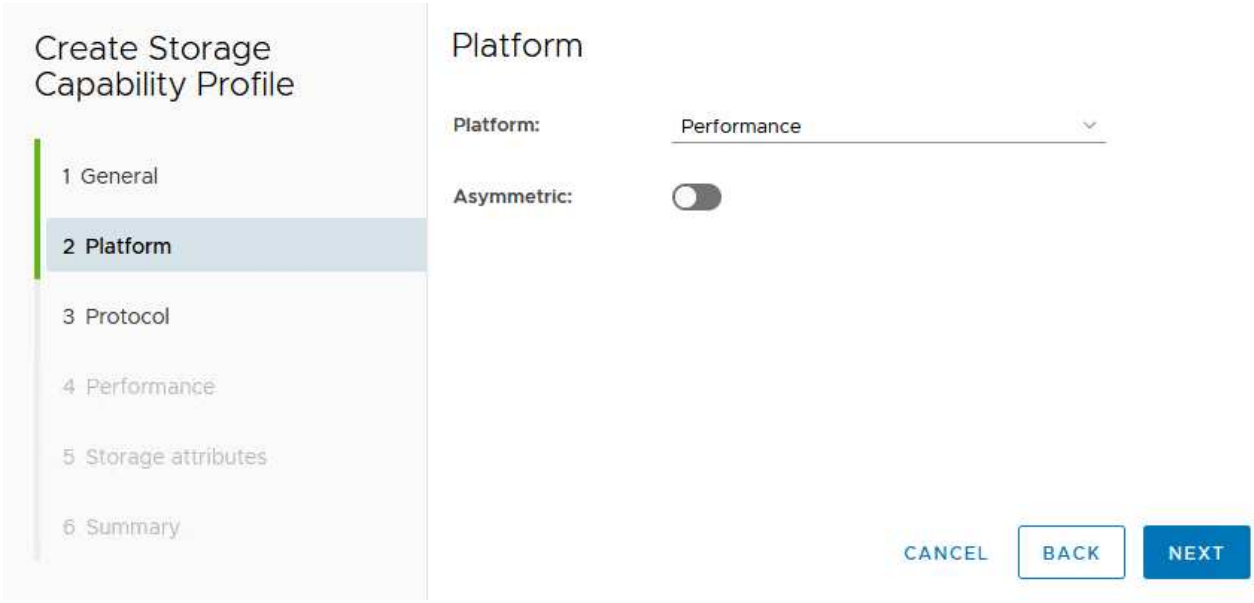

4. Next, select choice of protocol or **Any** to allow all possible protocols. Click **Next** to continue.

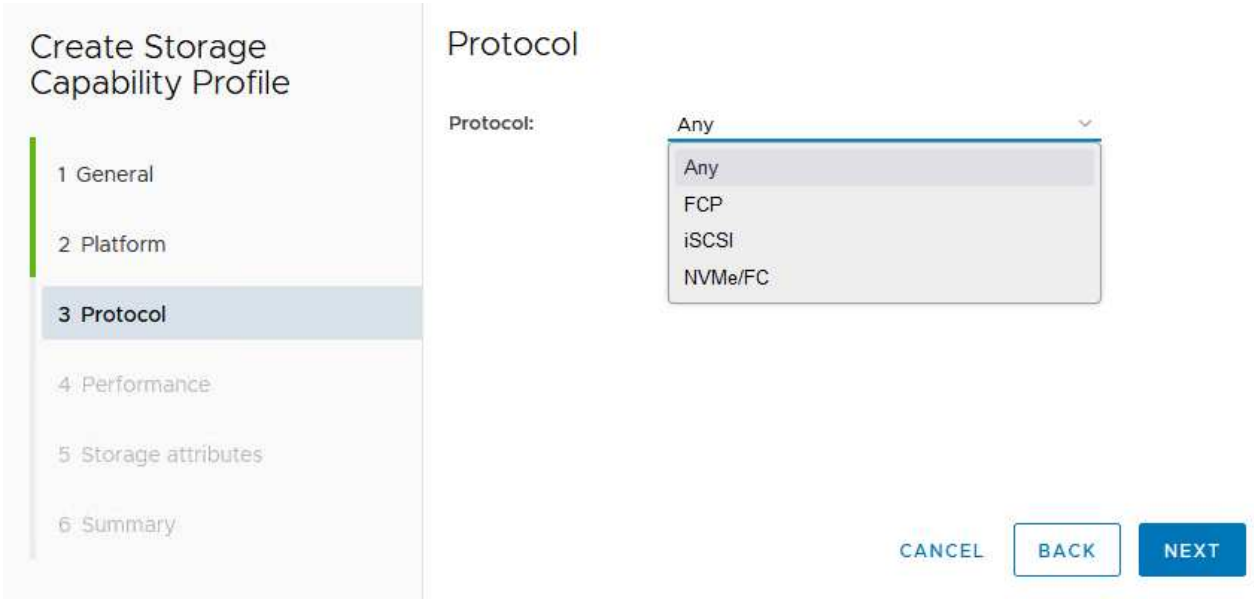

5. The **performance** page allows setting of quality of service in form of minimum and maximum IOPs allowed.

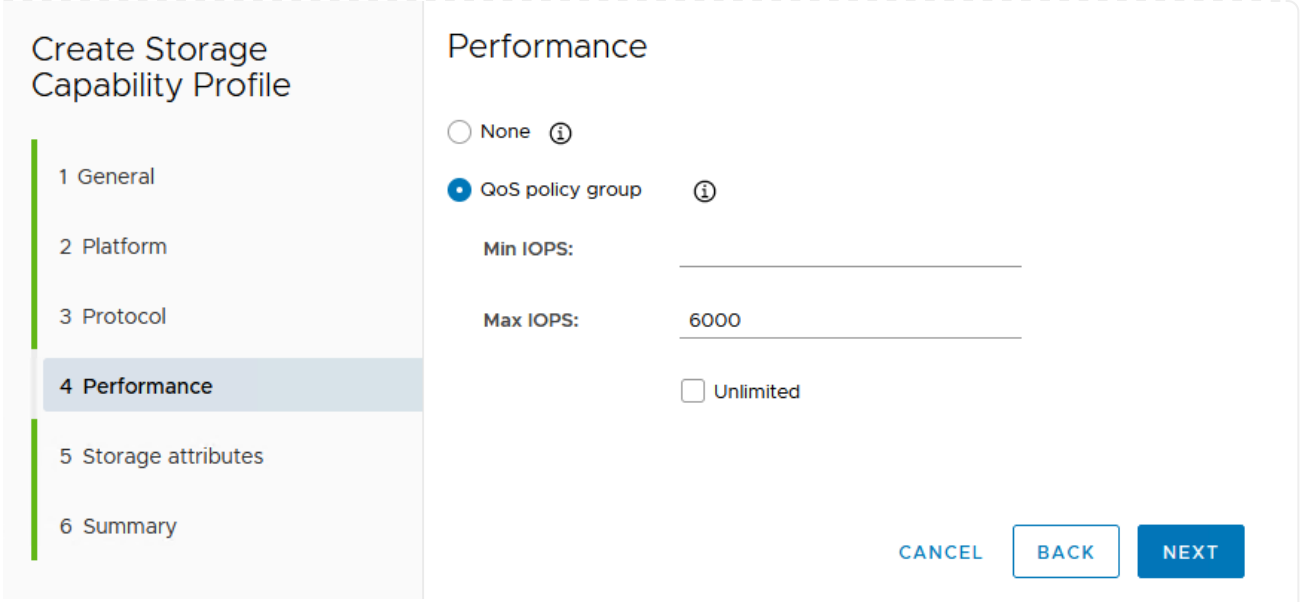

6. Complete the **storage attributes** page selecting storage efficiency, space reservation, encryption and any tiering policy as needed.

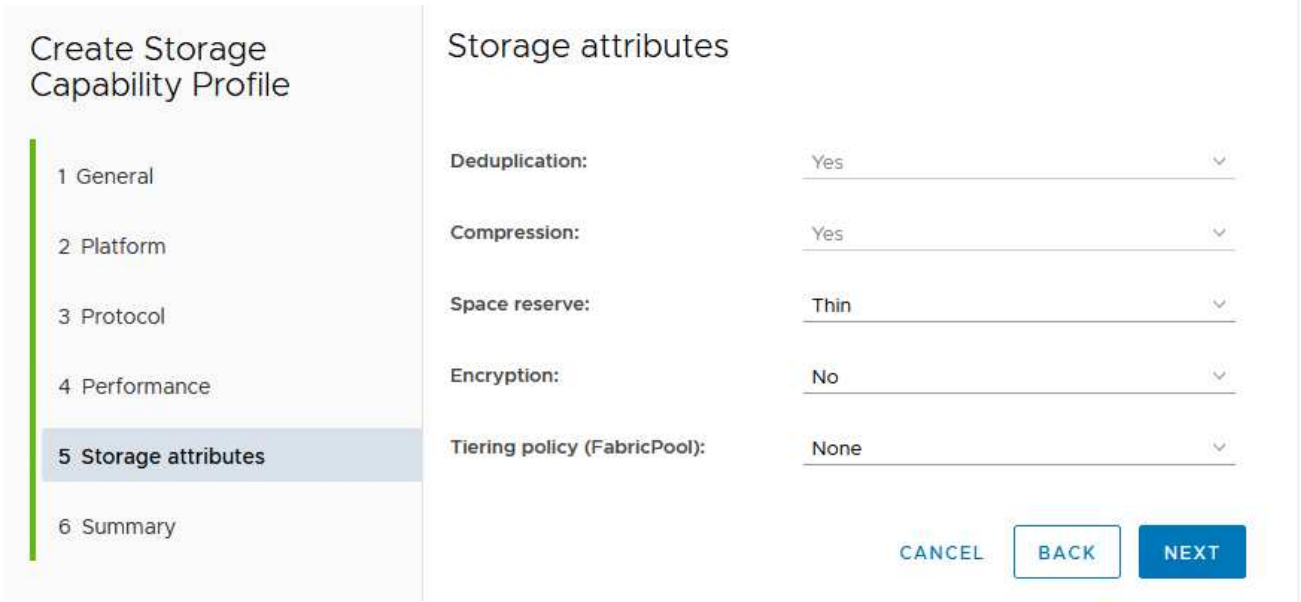

7. Finally, review the summary and click on Finish to create the profile.

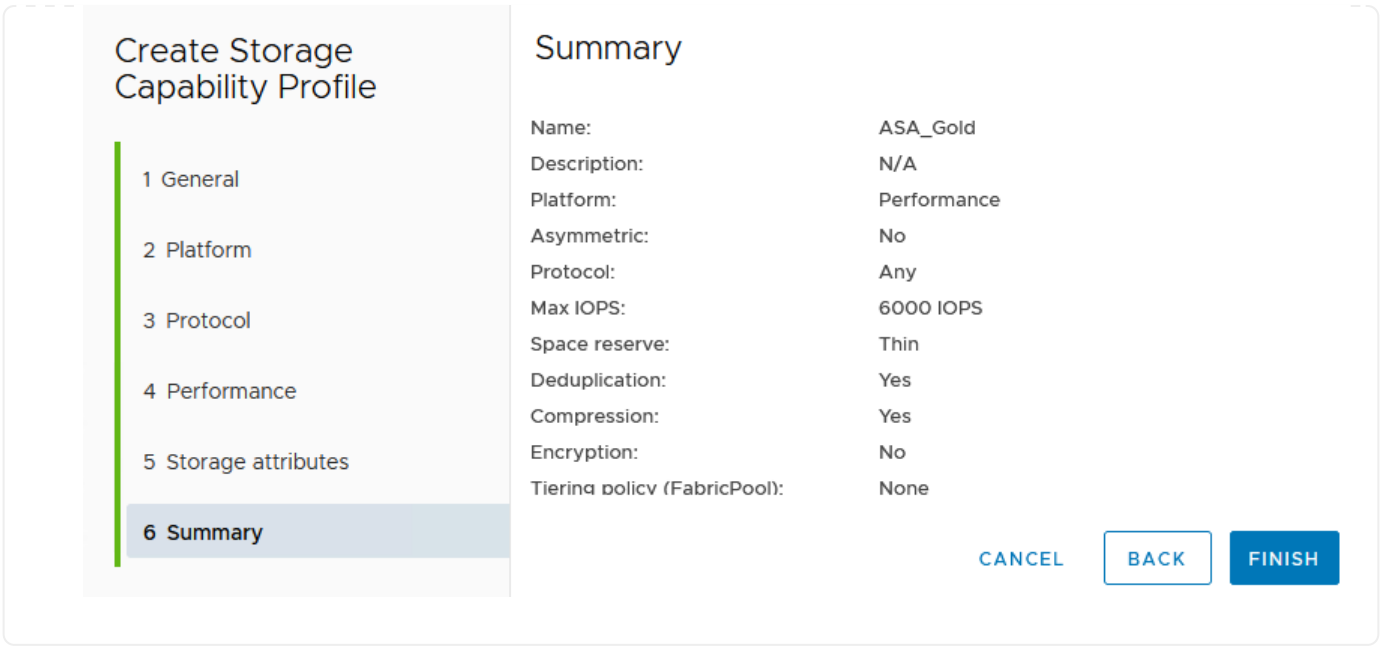
Create Storage Capability Profile
1 General
2 Platform
3 Protocol
4 Performance
5 Storage attributes
6 Summary
Summary
Name: ASA_Gold
Description: N/A
Platform: Performance
Asymmetric: No
Protocol: Any
Max IOPS: 6000 IOPS
Space reserve: Thin
Deduplication: Yes
Compression: Yes
Encryption: No
Tiering policy (FabricPool): None
CANCEL
BACK
FINISH

Create a vVols datastore in ONTAP Tools

To create a vVols datastore in ONTAP Tools complete the following steps:

1. In ONTAP Tools select **Overview** and from the **Getting Started** tab click on **Provision** to start the wizard.

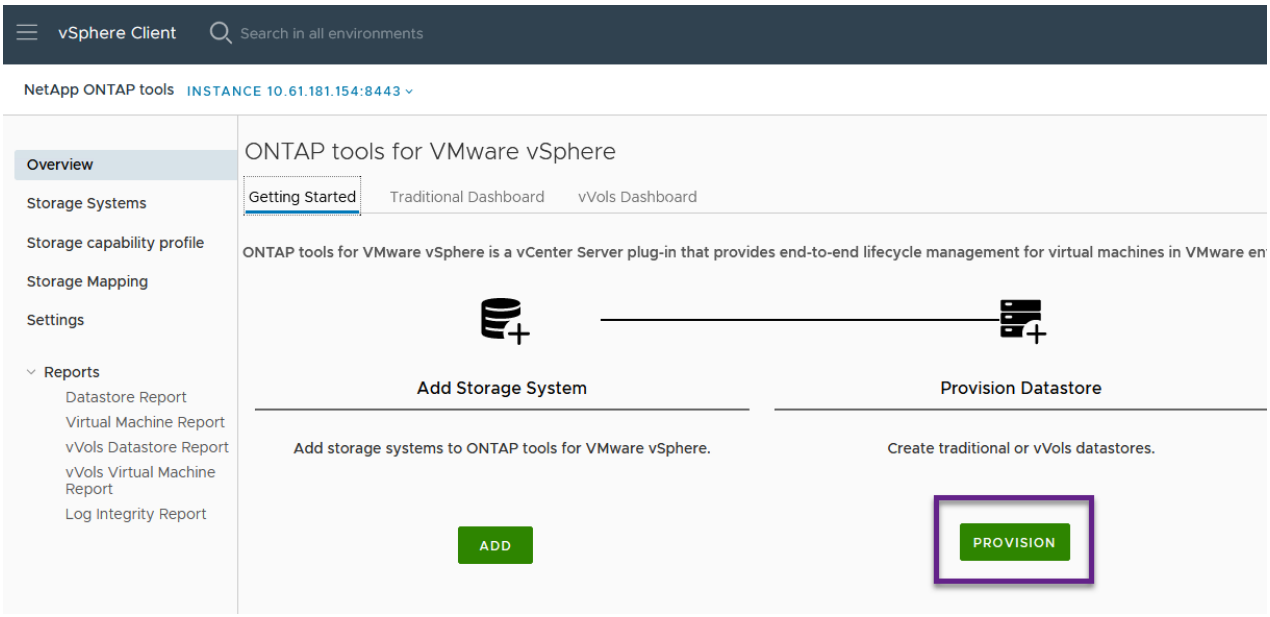

2. On the **General** page of the New Datastore wizard select the vSphere datacenter or cluster destination. Select **vVols** as the dastatore type, fill out a name for the datastore, and select the protocol.

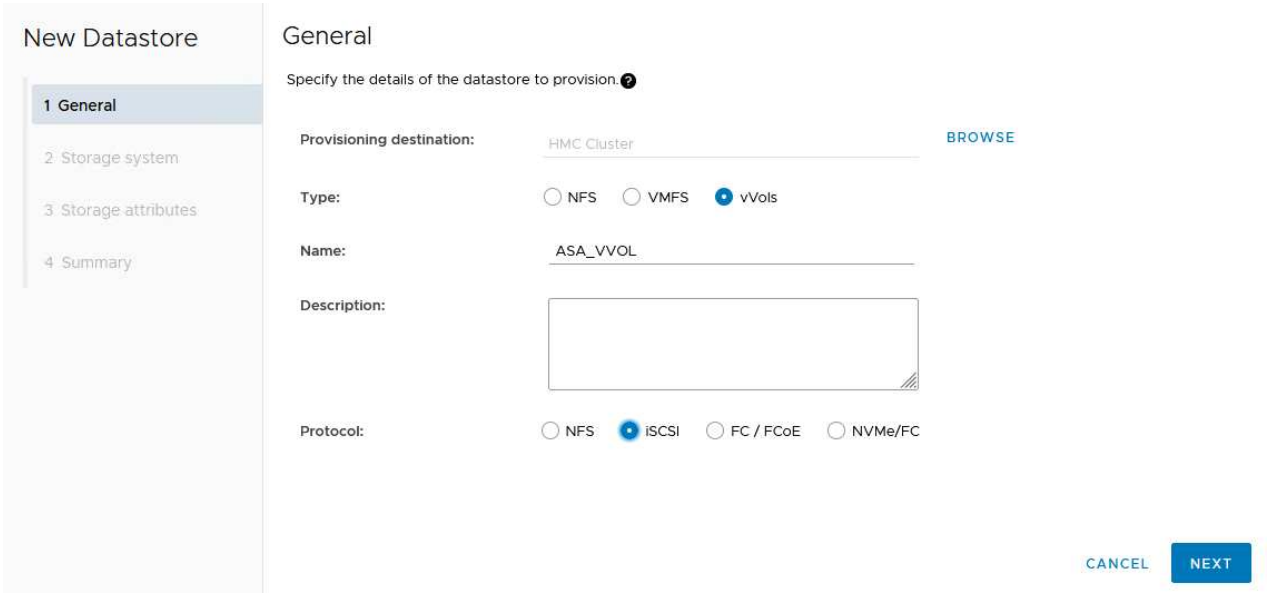

3. On the **Storage system** page select the select a storage capability profile, the storage system and SVM. Click on **Next** to continue.

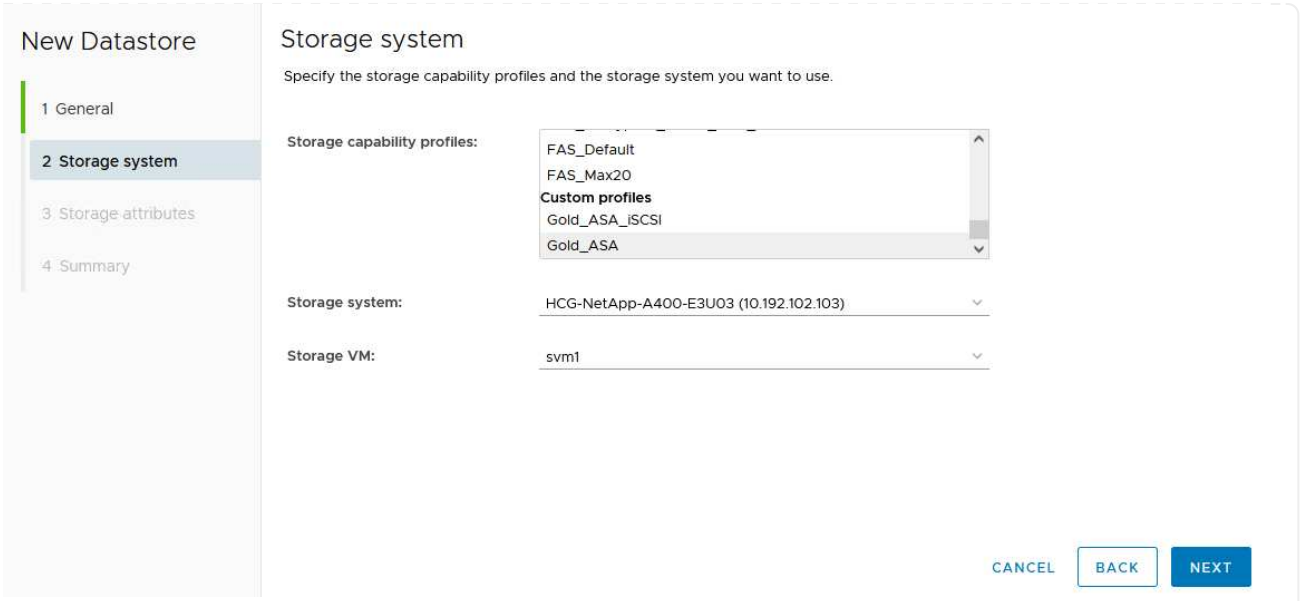

4. On the **Storage attributes** page select to create a new volume for the datastore and fill out the storage attributes of the volume to be created. Click on **Add** to create the volume and then **Next** to continue.

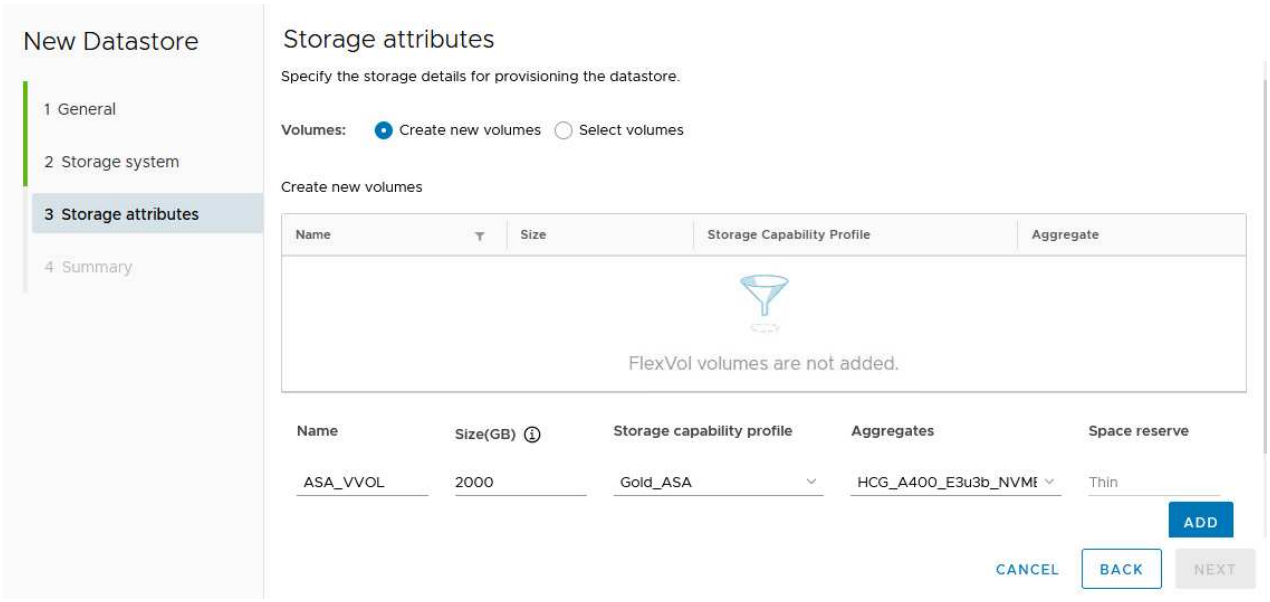

5. Finally, review the summary and click on **Finish** to start the vVol datastore creation process.

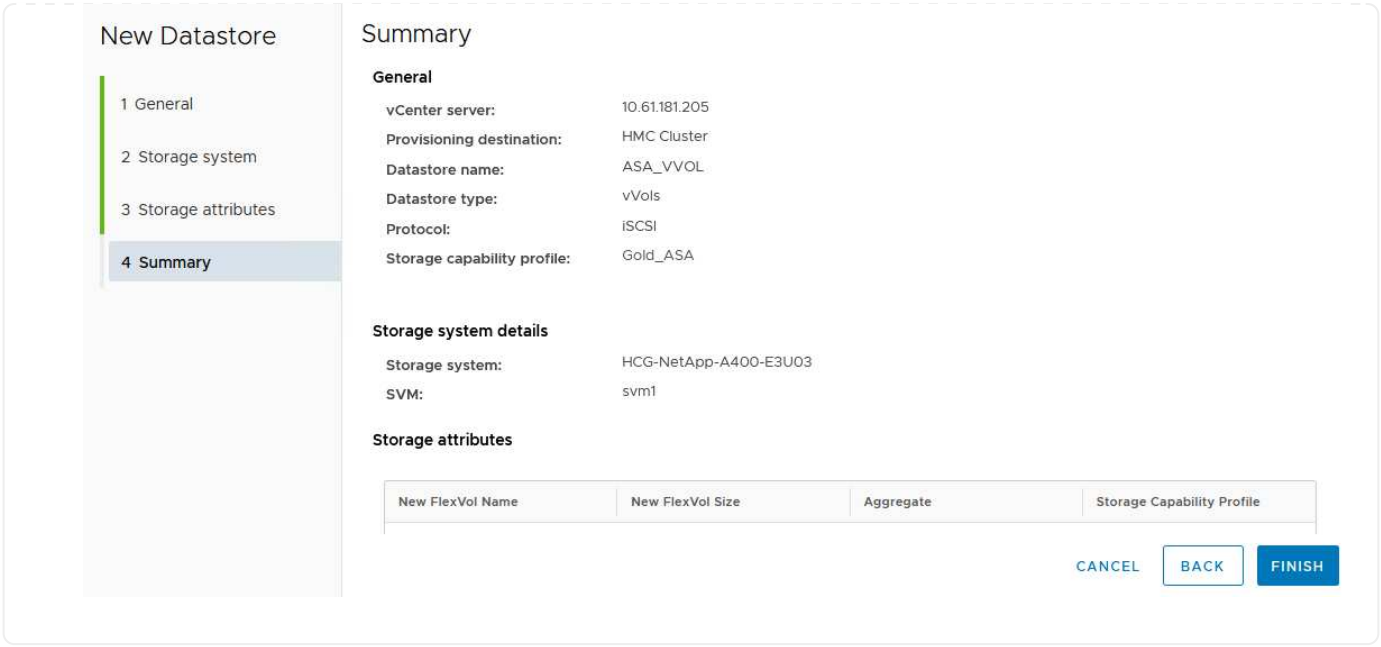

Create a VM storage policy in the vSphere client

A VM storage policy is a set of rules and requirements that define how virtual machine (VM) data should be stored and managed. It specifies the desired storage characteristics, such as performance, availability, and data services, for a particular VM.

In this case, the task involves creating a VM storage policy to specify that a virtual machine will be generated on vVol datastores and to establish a one-to-one mapping with the previously generated storage capability profile.

Create a VM storage policy

To create a VM storage policy complete the following steps:

1. From the vSphere clients main menu select **Policies and Profiles**.

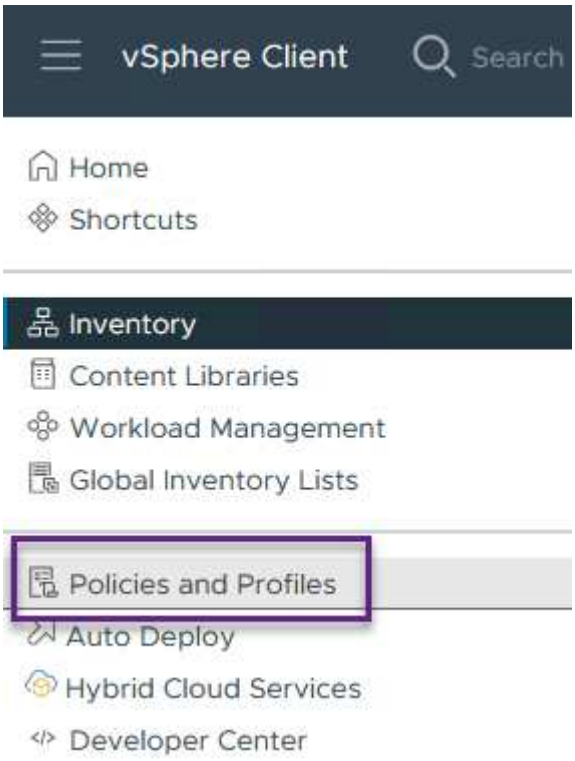

2. In the **Create VM Storage Policy** wizard, first fill out a name and description for the policy and click on **Next** to continue.

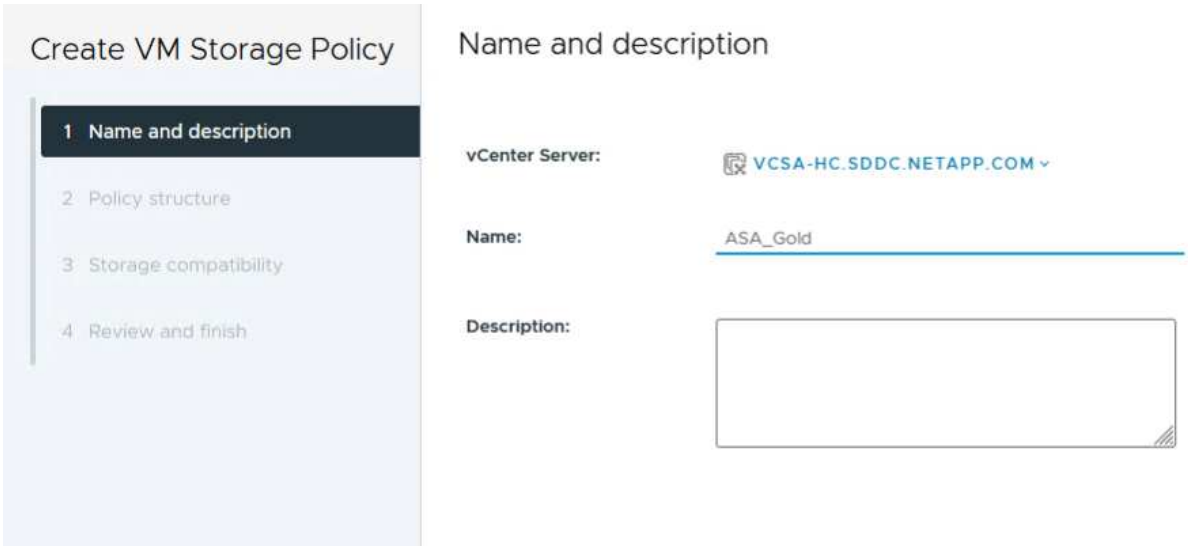

3. On the **Policy structure** page select to enable rules for NetApp clustered data ontap vVol storage and click on **Next**.

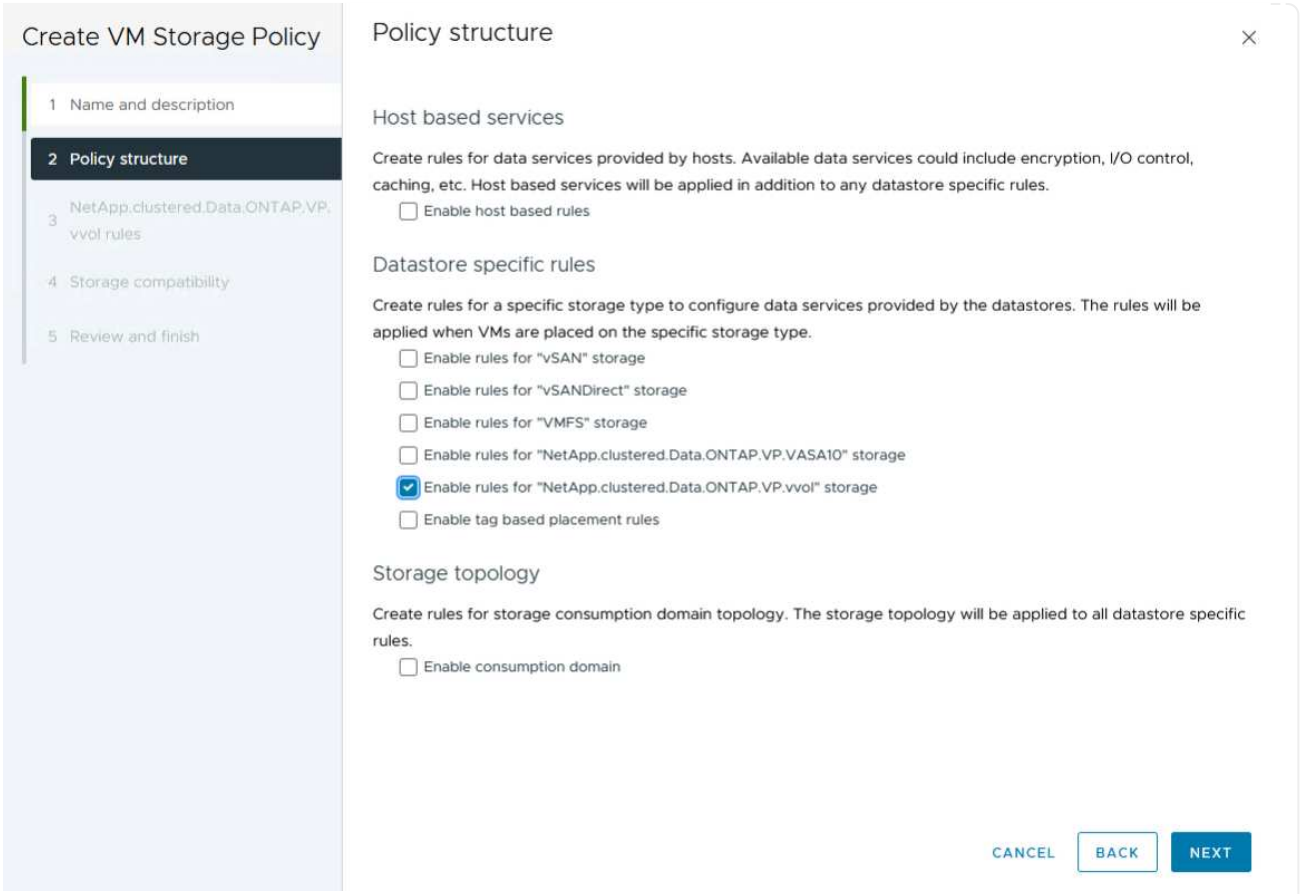

4. On the next page specific to the policy structure chosen, select the storage capability profile that describes the storage system(s) to be used in the VM storage policy. Click on **Next** to continue.

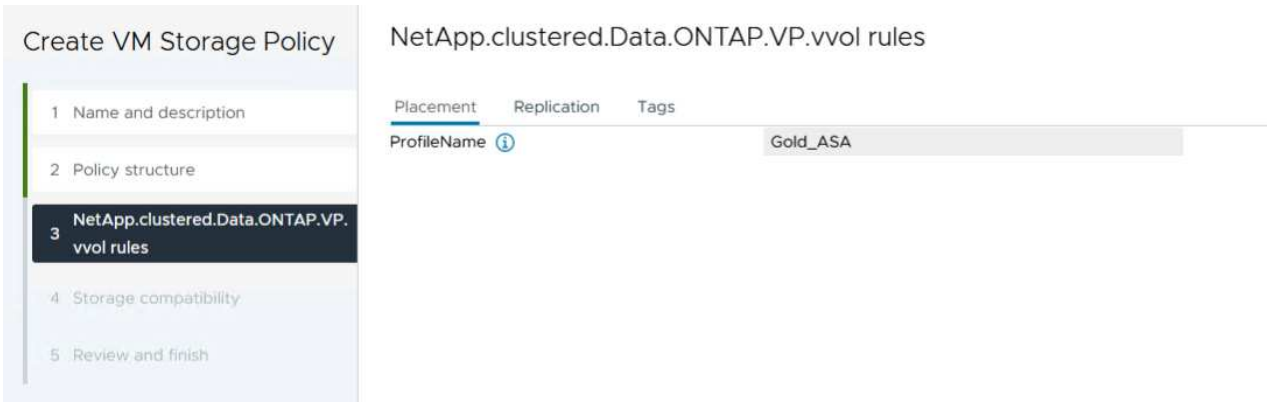

5. On the **Storage compatibility** page, review the list of vSAN datastores that match this policy and click **Next**.
6. Finally, review the policy to be implemented and click on **Finish** to create the policy.

Create a VM storage policy in the vSphere client

A VM storage policy is a set of rules and requirements that define how virtual machine (VM) data should be stored and managed. It specifies the desired storage characteristics, such as performance, availability, and data services, for a particular VM.

In this case, the task involves creating a VM storage policy to specify that a virtual machine will be generated

on vVol datastores and to establish a one-to-one mapping with the previously generated storage capability profile.

Create a virtual machine on a vVol datastore

The final step is to create a virtual machine using the VM storage policies previously created:

1. From the **New Virtual Machine** wizard select **Create a new virtual machine** and select **Next** to continue.

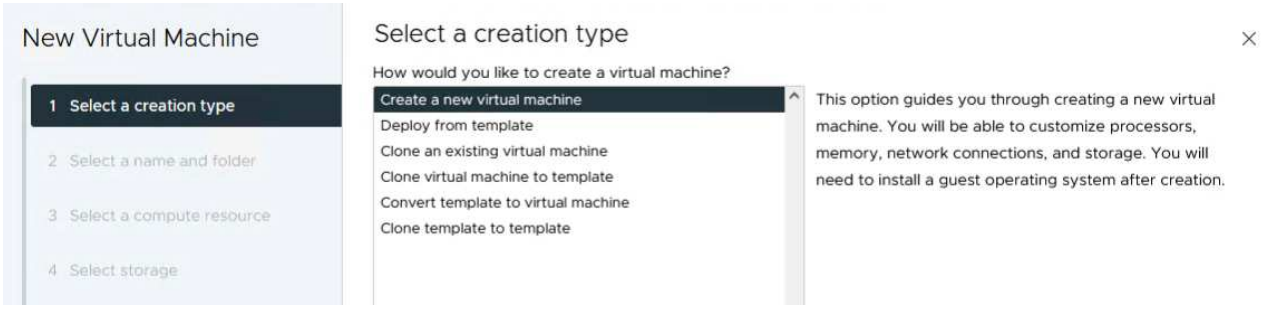

2. Fill in a name and select a location for the virtual machine and click on **Next**.
3. On the **Select a compute resource** page select a destination and click on **Next**.

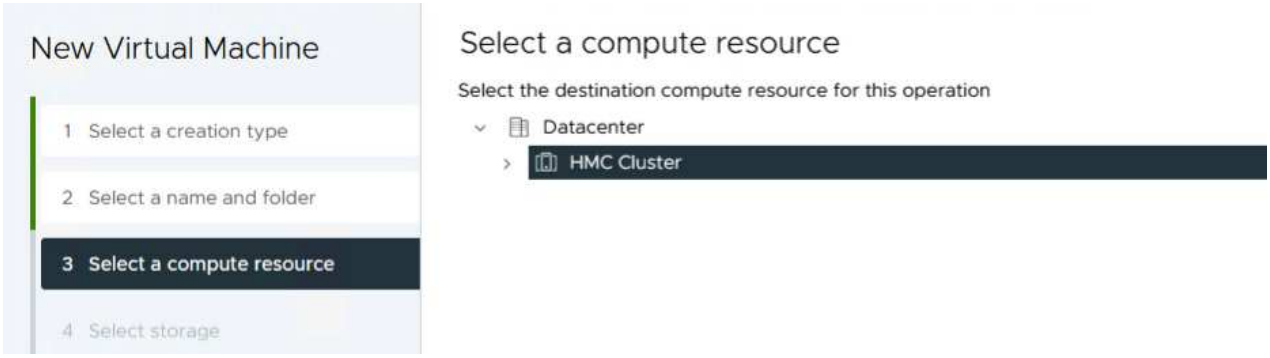

4. On the **Select storage** page select a VM Storage Policy and the vVols datastore that will be the destination for the VM. Click on **Next**.

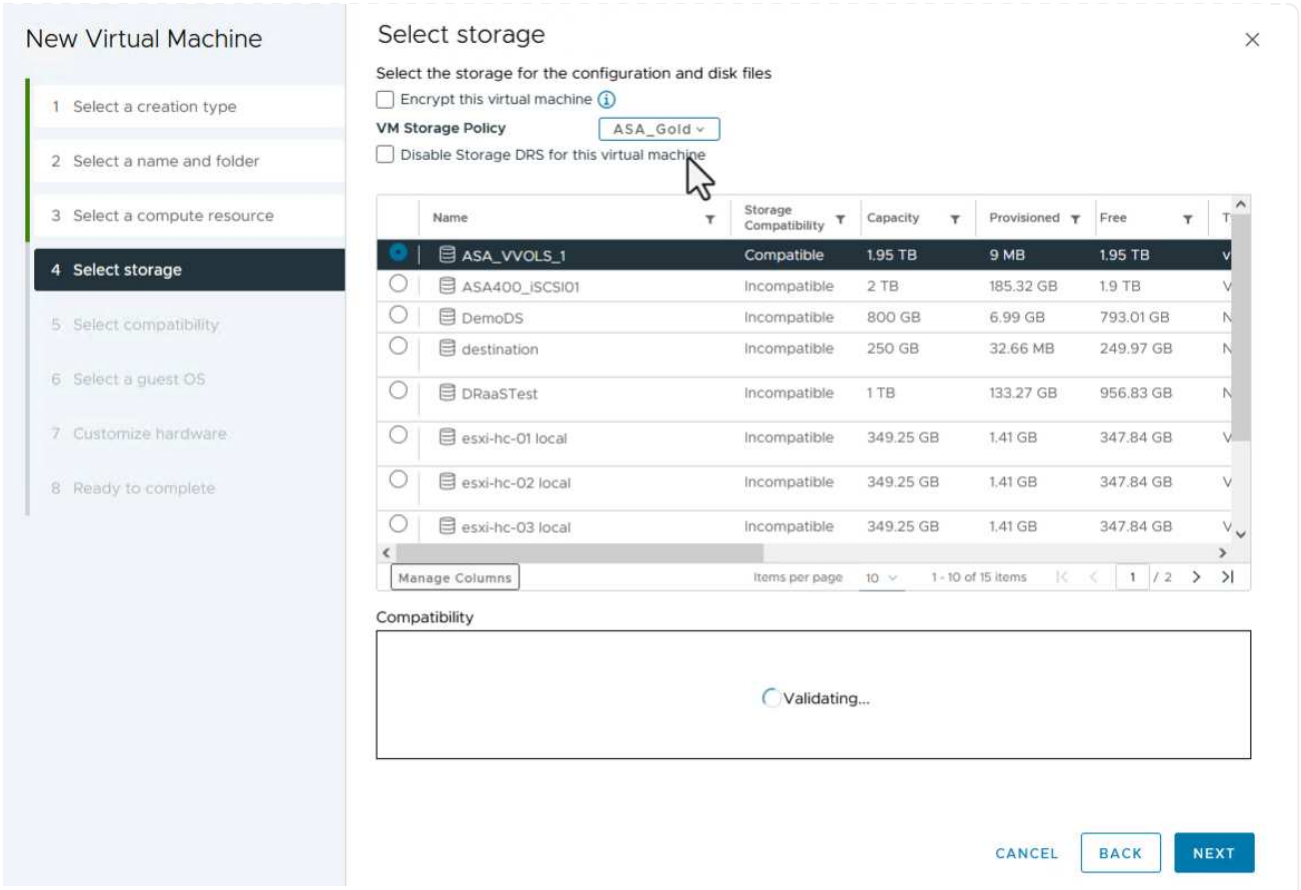

5. On the **Select compatibility** page choose the vSphere version(s) that the VM will be compatible with.
6. Select the guest OS family and version for the new VM and click on **Next**.
7. Fill out the **Customize hardware** page. Note that a separate VM storage policy can be selected for each hard disk (VMDK file).

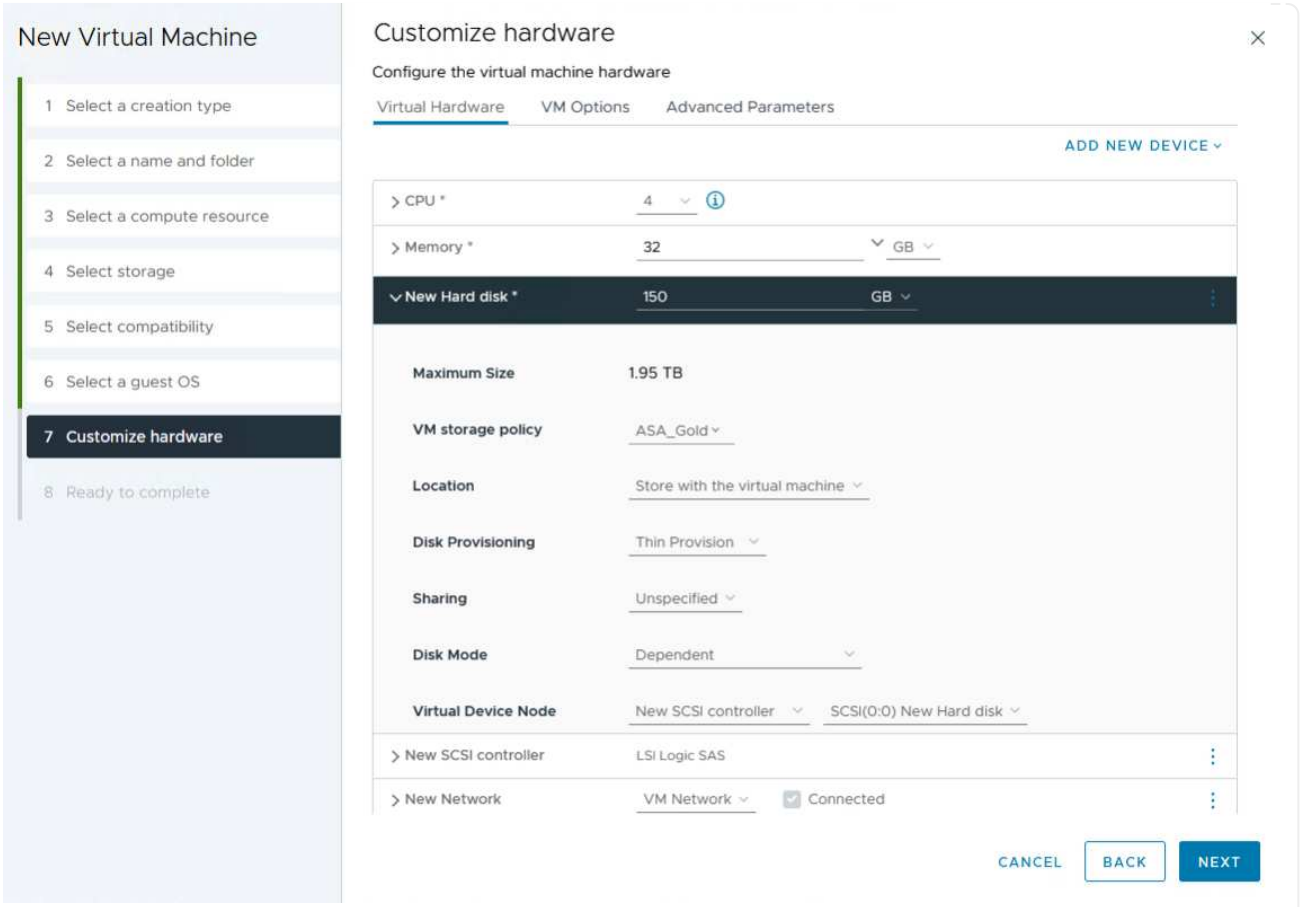

8. Finally, review the summary page and click on **Finish** to create the VM.

In summary, NetApp ONTAP Tools automates the process of creating vVol datastores on ONTAP storage systems. Storage capability profiles define not only the storage systems to be used for datastore creation but also dictate QoS policies that can be implemented on an individual VMDK basis. vVols provide a simplified storage management paradigm and tight integration between NetApp and VMware make this a practical solution for streamlined, efficient, and granular control over virtualized environments.

NetApp All-Flash SAN Array with VMware vSphere 8

NetApp Cloud Insights is a cloud-based infrastructure monitoring and analytics platform designed to provide comprehensive visibility and insights into the performance, health, and costs of IT infrastructures, both on-premises and in the cloud. Key features of NetApp Cloud Insights include real-time monitoring, customizable dashboards, predictive analytics, and cost optimization tools, allowing organizations to effectively manage and optimize their on-premises and cloud environments.

Author: Josh Powell - NetApp Solutions Engineering

Monitoring On-Premises Storage with NetApp Cloud Insights

NetApp Cloud Insights operates through Acquisition Unit software, which is set up with data collectors for assets such as VMware vSphere and NetApp ONTAP storage systems. These collectors gather data and transmit it to Cloud Insights. The platform then utilizes a variety of dashboards, widgets, and metric queries to organize the data into insightful analyses for users to interpret.

Cloud Insights architecture diagram:

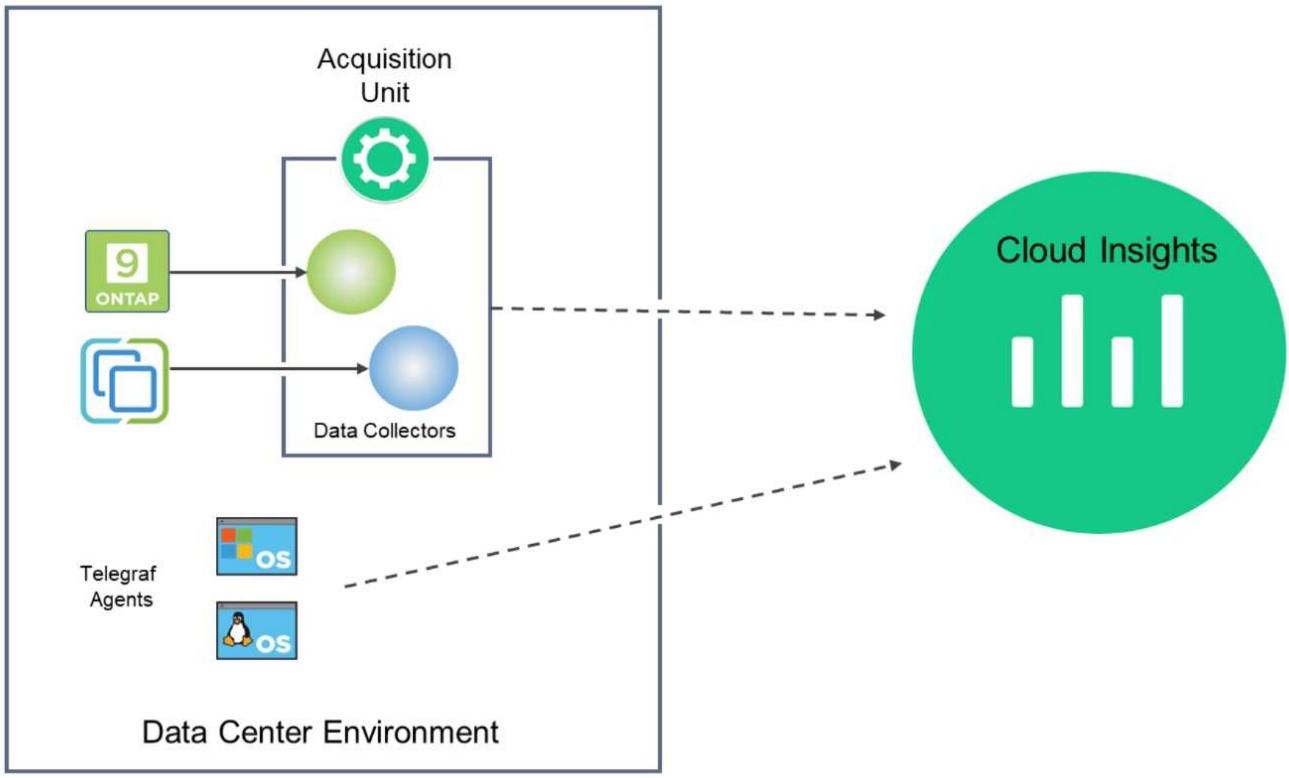

Solution Deployment Overview

This solution provides an introduction to monitoring on-premises VMware vSphere and ONTAP storage systems using NetApp Cloud Insights.

This list provides the high level steps covered in this solution:

1. Configure Data Collector for a vSphere cluster.
2. Configure Data Collector for an ONTAP storage system.
3. Use Annotation Rules to tag assets.
4. Explore and correlate assets.
5. Use a Top VM Latency dashboard to isolate noisy neighbors.
6. Identify opportunities to rightsize VMs.
7. Use queries to isolate and sort metrics.

Prerequisites

This solution uses the following components:

1. NetApp All-Flash SAN Array A400 with ONTAP 9.13.
2. VMware vSphere 8.0 cluster.
3. NetApp Cloud Insights account.
4. NetApp Cloud Insights Acqusition Unit software installed on a local VM with network connectivity to assets

for data collection.

Solution Deployment

Configure Data Collectors

To configure Data Collectors for VMware vSphere and ONTAP storage systems complete the following steps:

Add a Data Collector for an ONTAP storage systems

1. Once logged into Cloud Insights, navigate to **Observability > Collectors > Data Collectors** and press the button to install a new Data Collector.

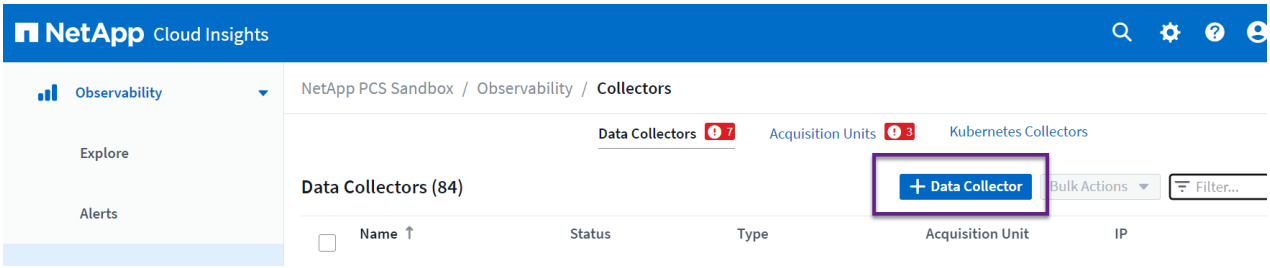

2. From here search for **ONTAP** and click on **ONTAP Data Management Software**.

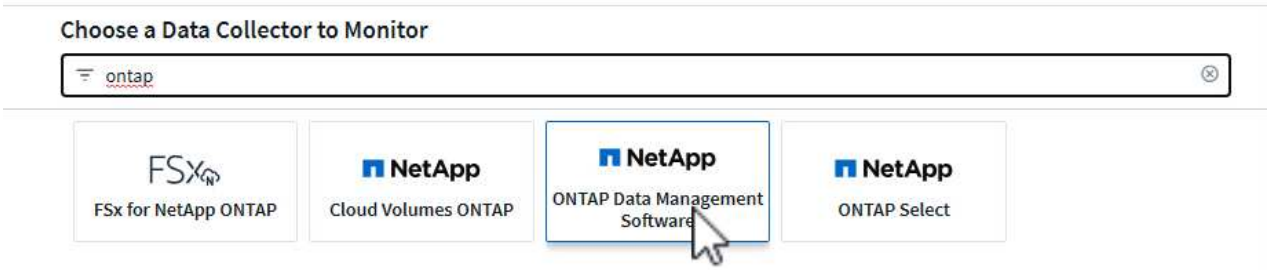

3. On the **Configure Collector** page fill out a name for the collector, specify the correct **Acquisition Unit** and provide the credentials for the ONTAP storage system. Click on **Save and Continue** and then **Complete Setup** at the bottom of the page to complete the configuration.

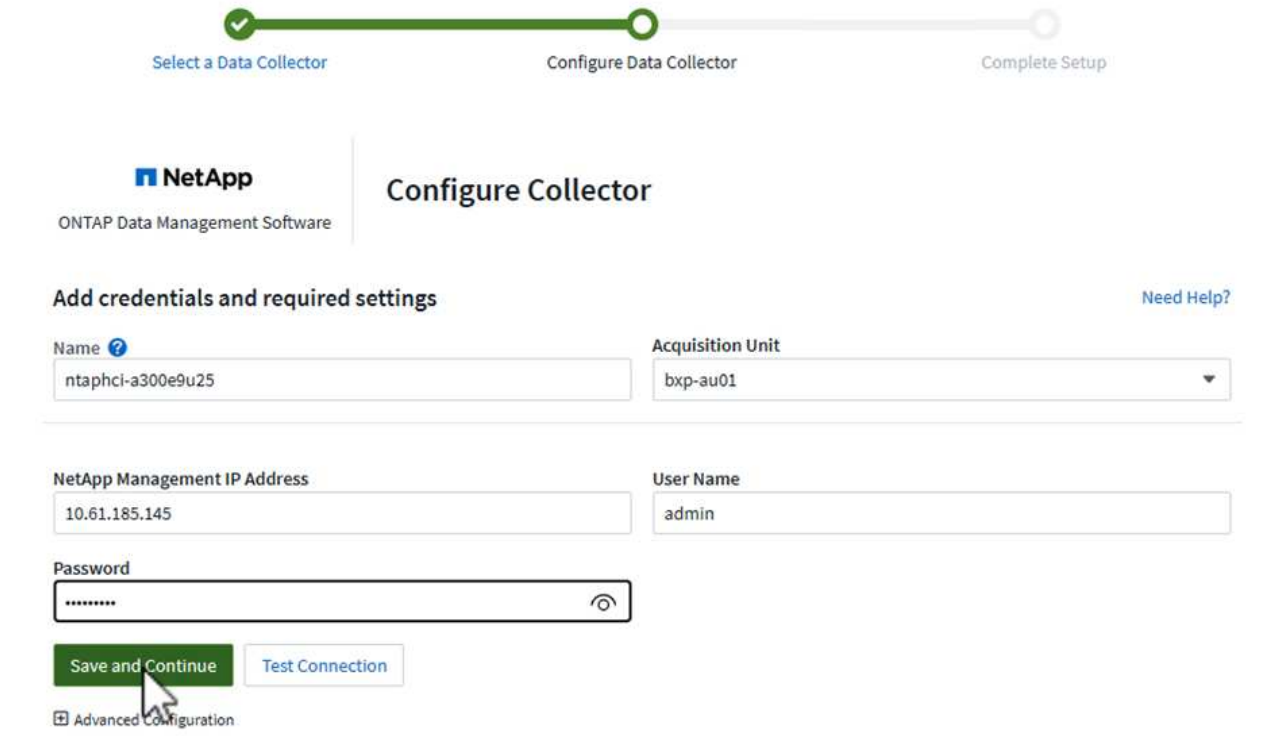

Add a Data Collector for a VMware vSphere cluster

1. Once again, navigate to **Observability > Collectors > Data Collectors** and press the button to install a new Data Collector.

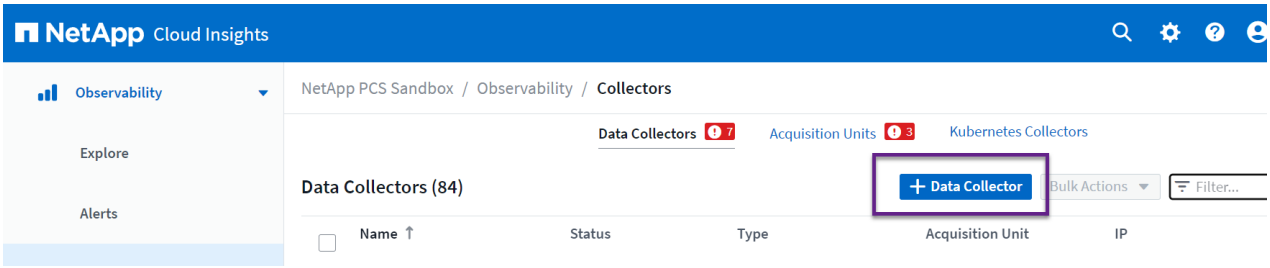

2. From here search for **vSphere** and click on **VMware vSphere**.

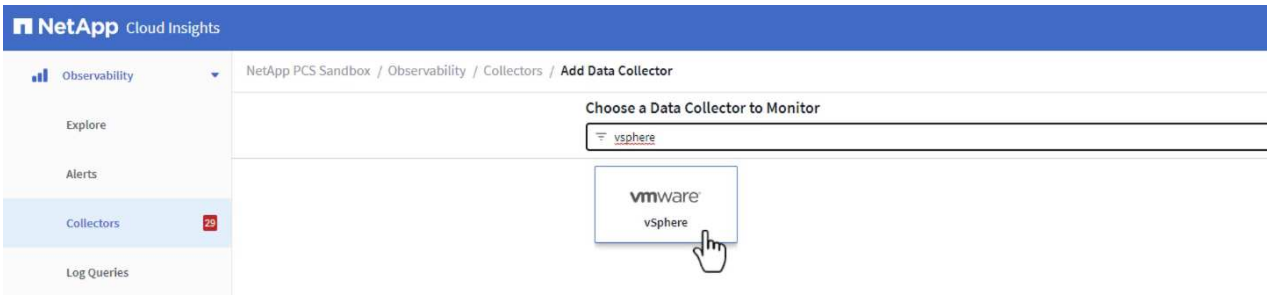

3. On the **Configure Collector** page fill out a name for the collector, specify the correct **Acquisition Unit** and provide the credentials for the vCenter server. Click on **Save and Continue** and then **Complete Setup** at the bottom of the page to complete the configuration.

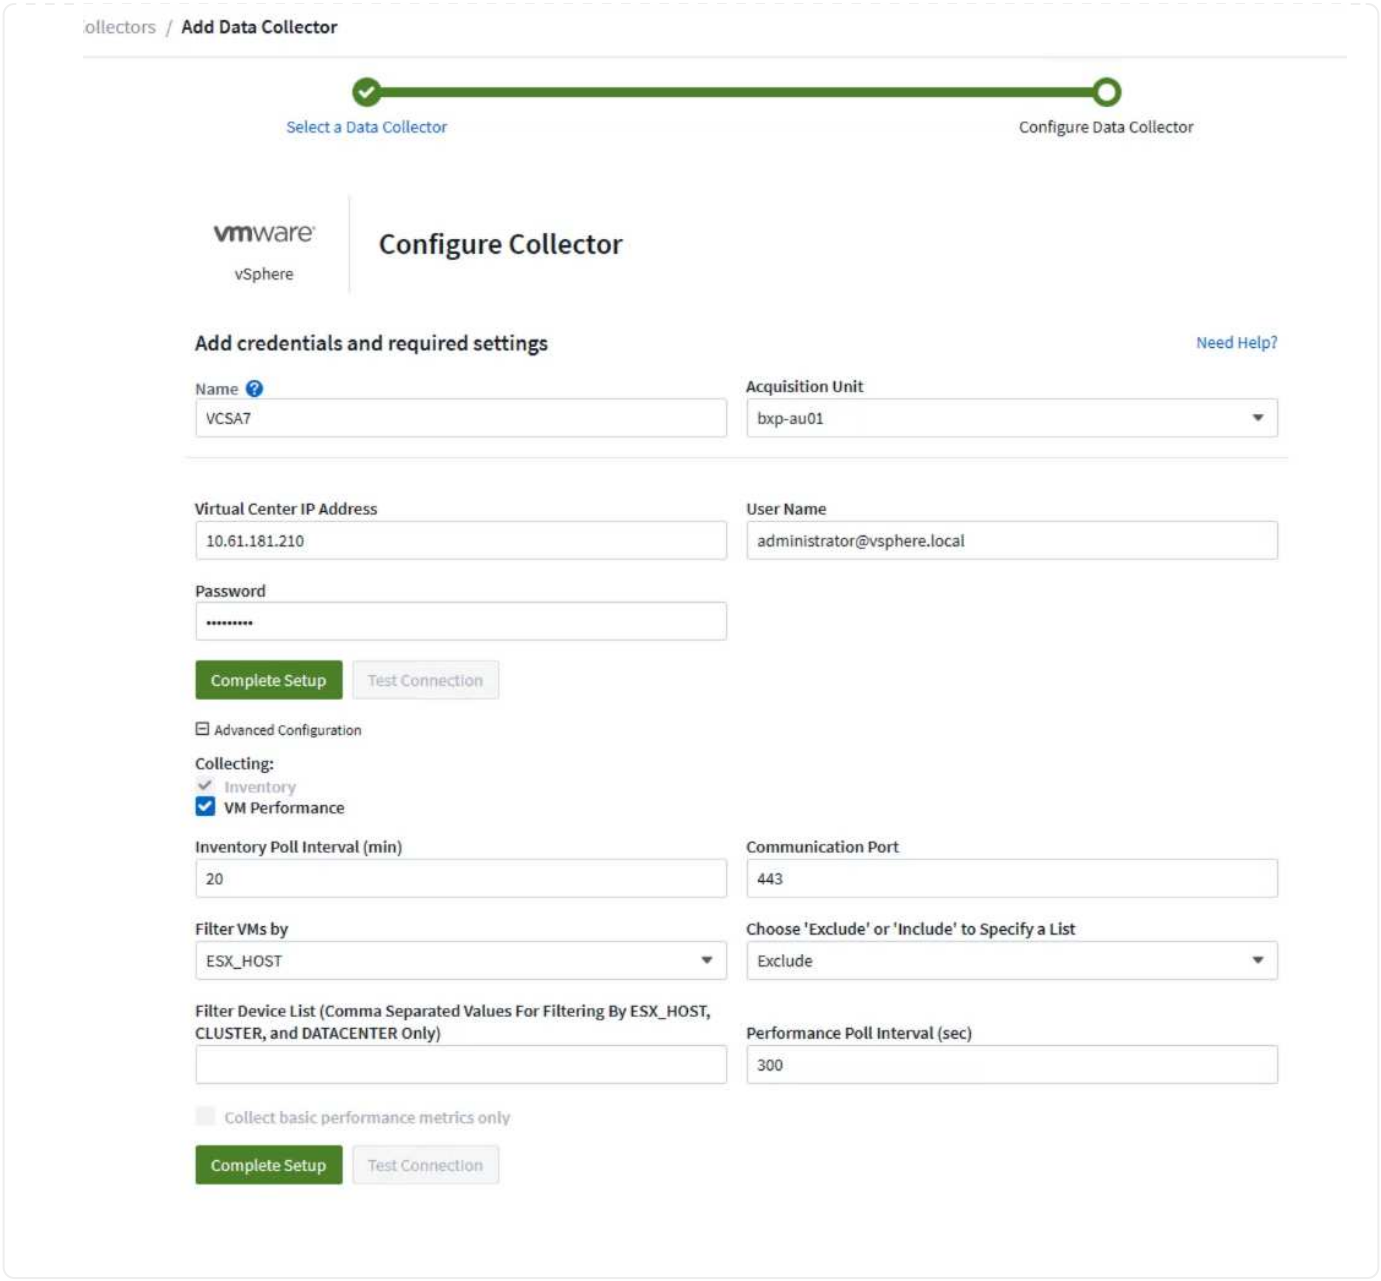

Add Annotations to assets

Annotations are a useful method of tagging assets so that they can be filtered and otherwise identified in the various views and metric queries available in Cloud Insights.

In this section, annotations will be added to virtual machine assets for filtering by **Data Center**.

Use Annotation Rules to tag assets

1. In the left-hand menu, navigate to **Observability > Enrich > Annotation Rules** and click on the **+ Rule** button in the upper right to add a new rule.

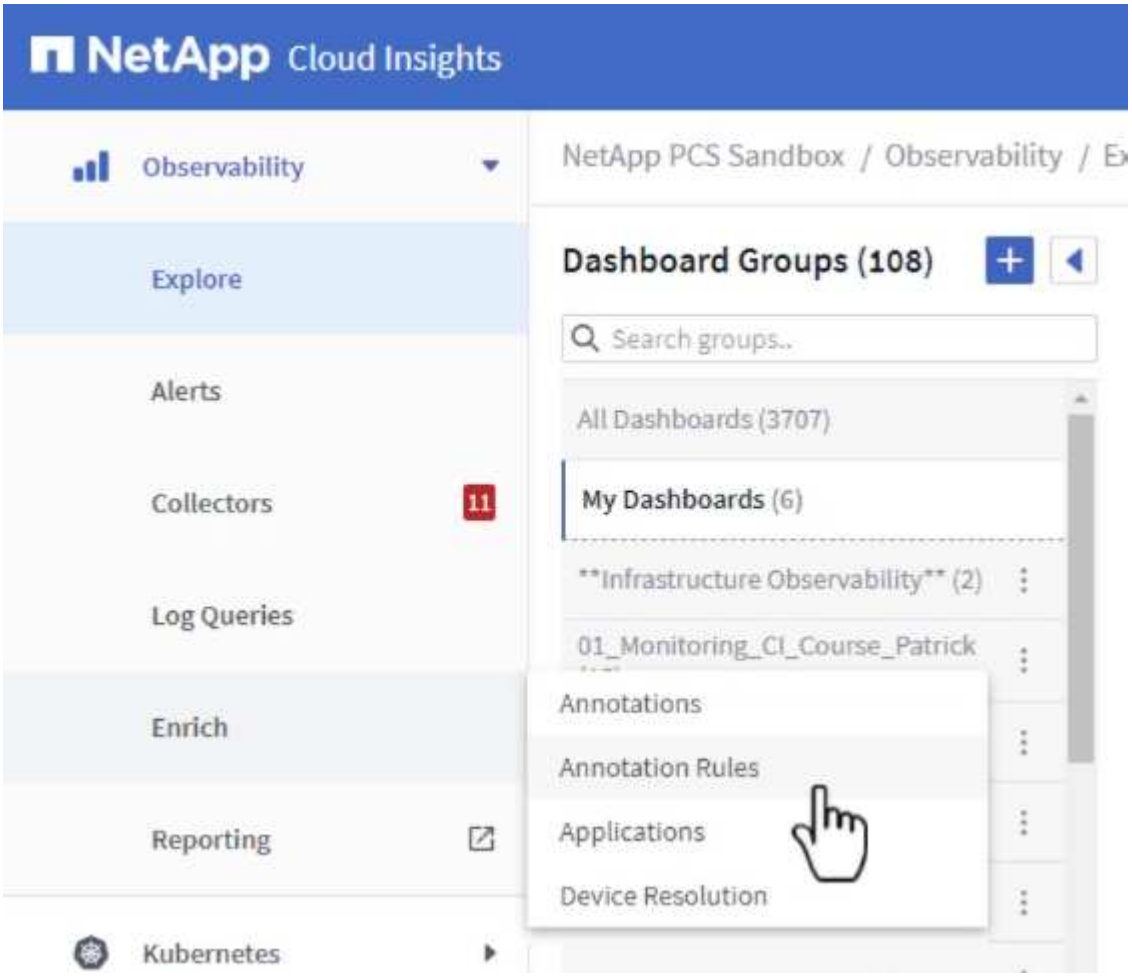

2. In the **Add Rule** dialog box fill in a name for the rule, locate a query to which the rule will be applied, the annotation field affected, and the value to be populated.

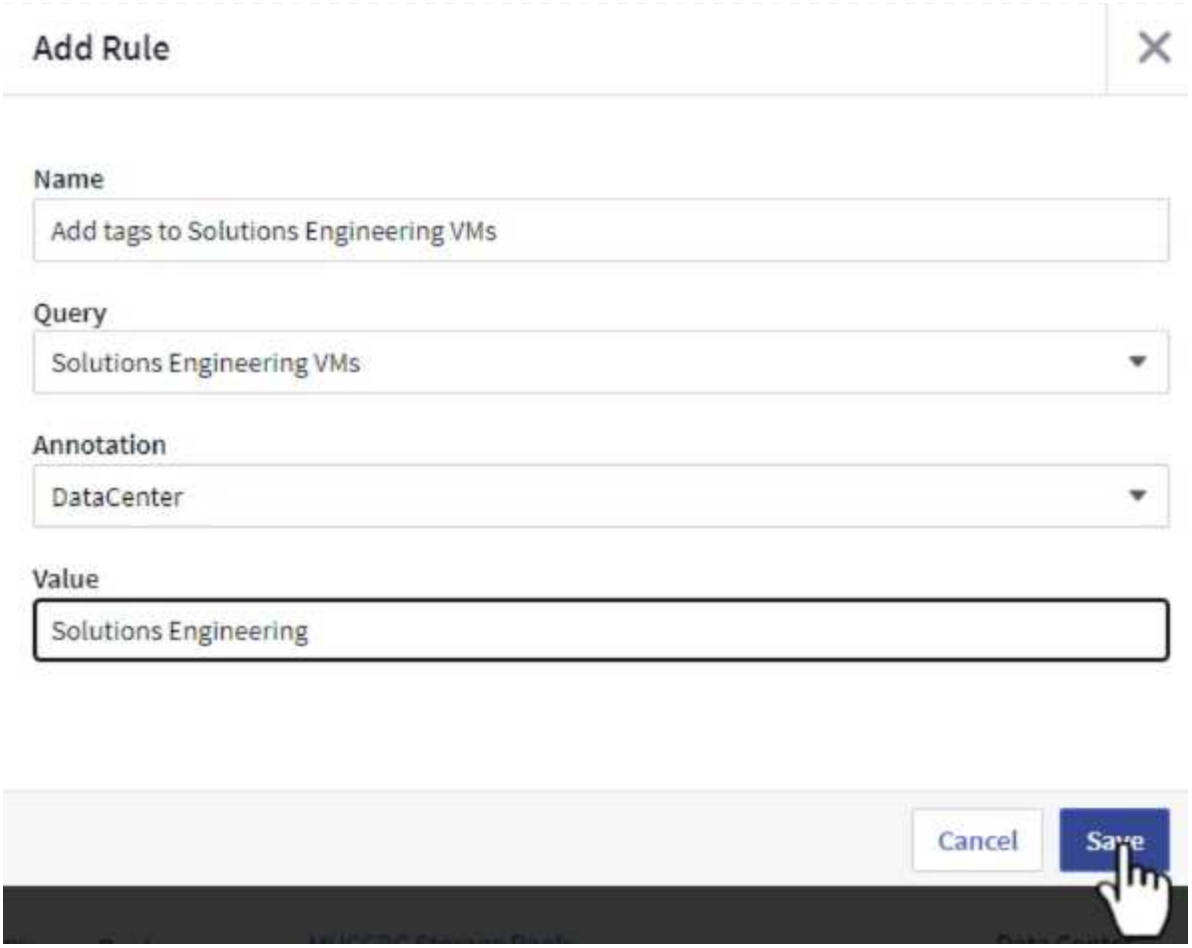

3. Finally, in the upper right hand corner of the **Annotation Rules** page click on **Run All Rules** to run the rule and apply the annotation to the assets.

Explore and correlate assets

Cloud Insights draws logical conclusions about the assets that are running together on your storage systems and vsphere clusters.

This sections illustrates how to use dashboards to correlate assets.

Correlating assets from a storage performance Dashboard

1. In the left-hand menu, navigate to **Observability > Explore > All Dashboards**.

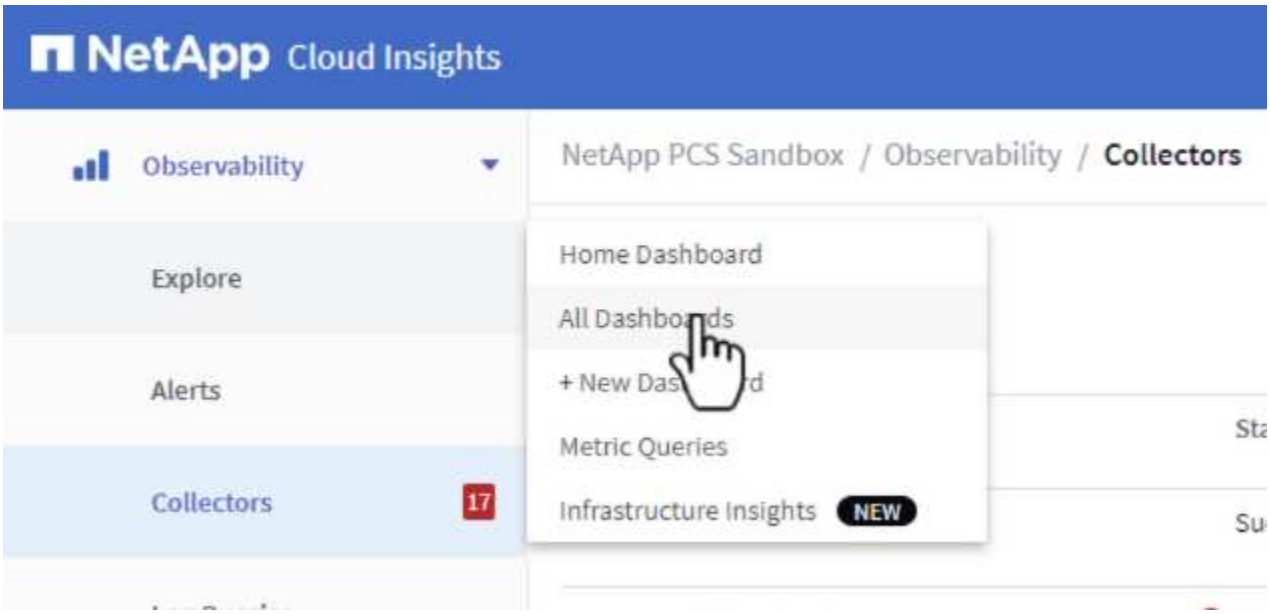

2. Click on the **+ From Gallery** button to view a list of ready-made dashboards that can be imported.

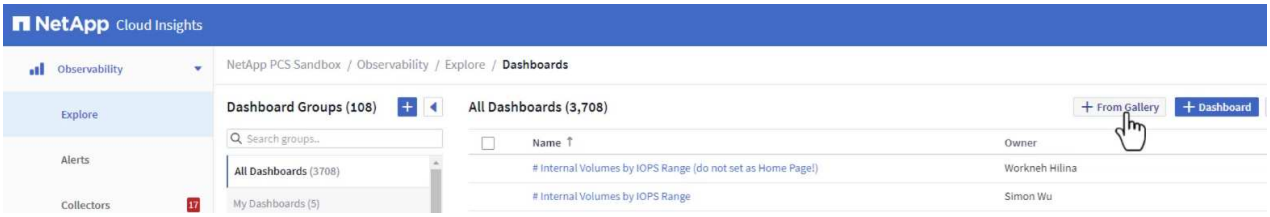

3. Choose a dashboard for FlexVol performance from the list and click on the **Add Dashboards** button at the bottom of the page.

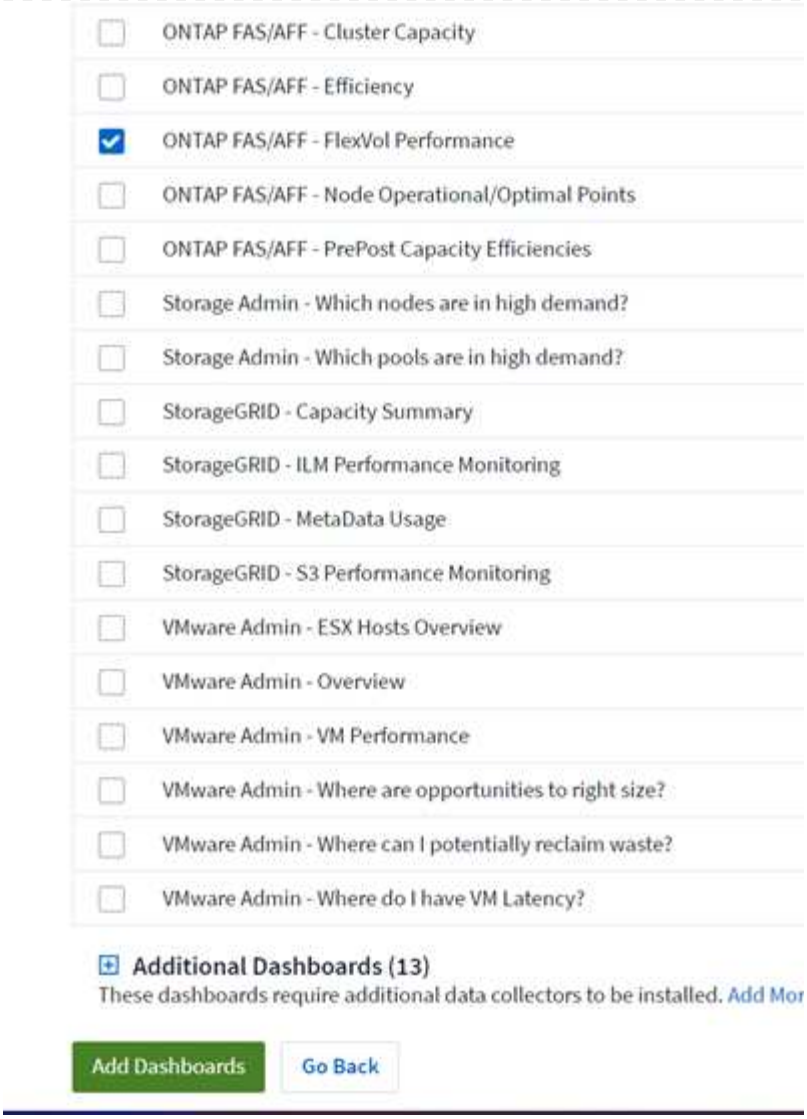

4. Once imported, open the dashboard. From here you can see various widgets with detailed performance data. Add a filter to view a single storage system and select a storage volume to drill into it's details.

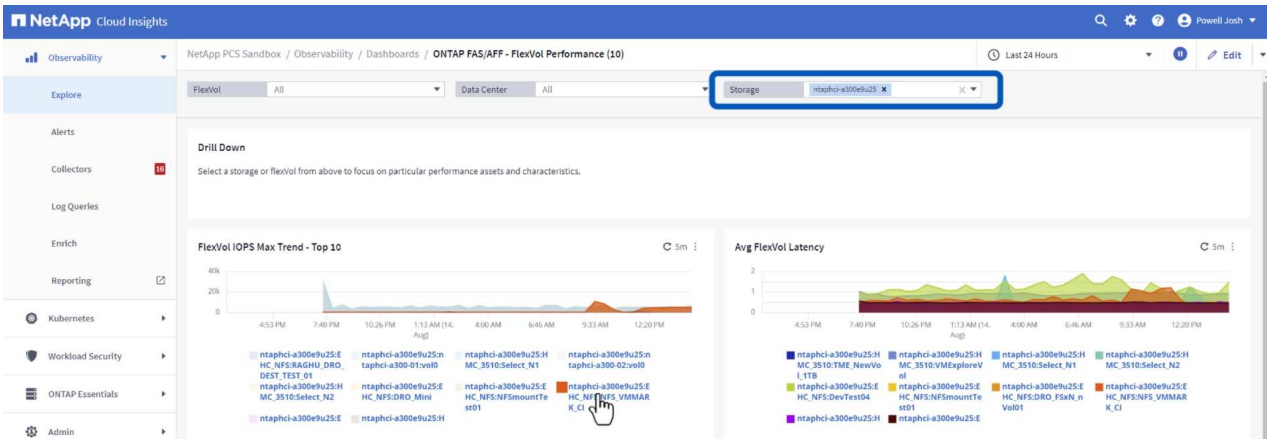

5. From this view you can see various metrics related to this storage volume and the top utilized and correlated virtual machines running on the volume.

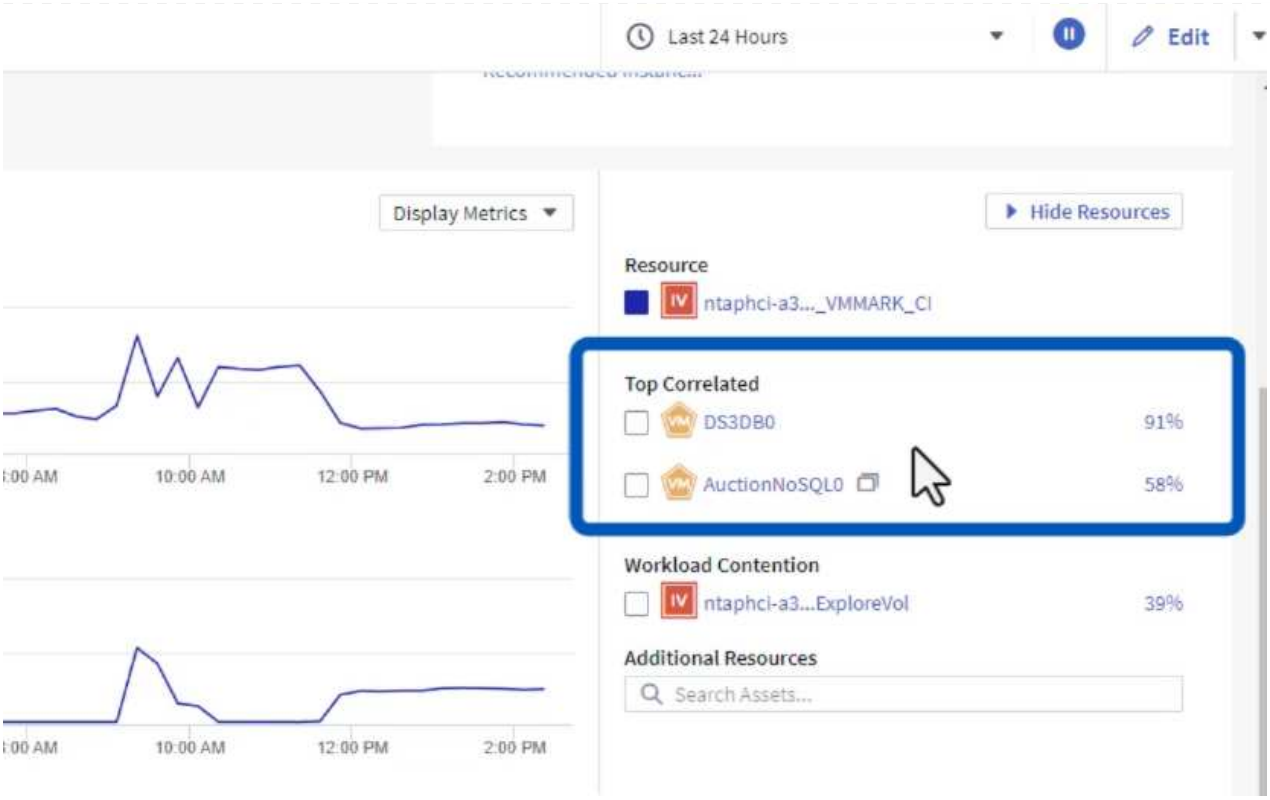

6. Clicking on the VM with the highest utilization drills into the metrics for that VM to view any potential issues.

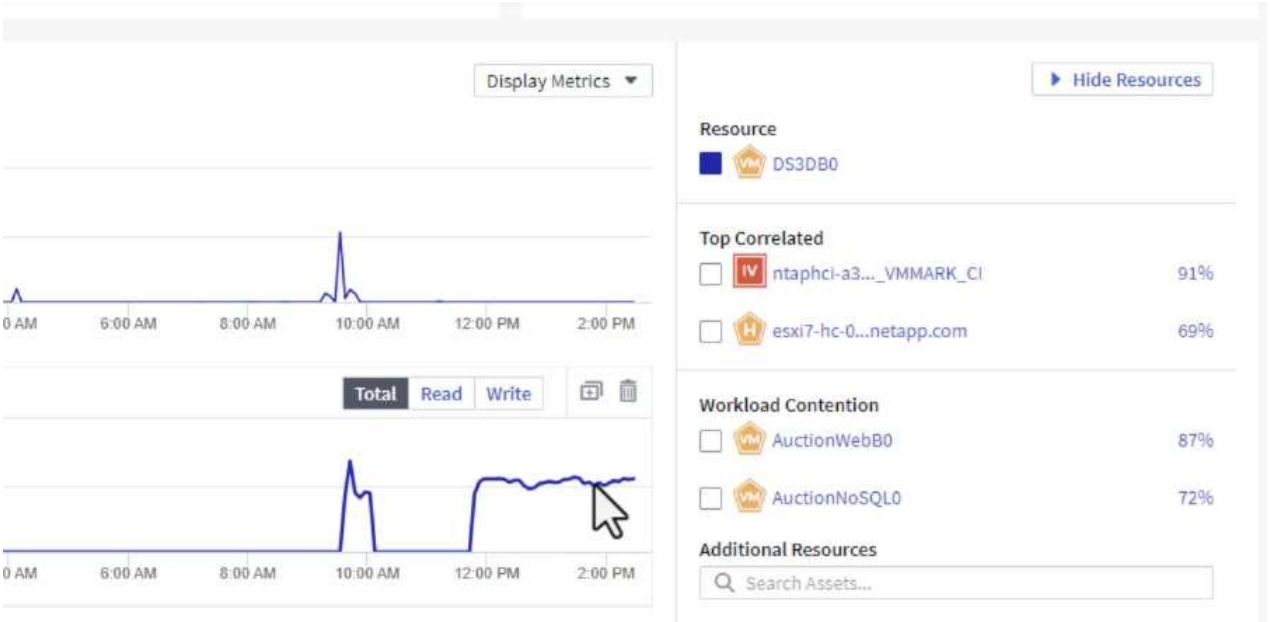

Use Cloud Insights to identify noisy neighbors

Cloud Insights features dashboards that can easily isolate peer VMs that are negatively impacting other VMs running on the same storage volume.

Use a Top VM Latency dashboard to isolate noisy neighbors

1. In this example access a dashboard available in the **Gallery** called **VMware Admin - Where do I have VM Latency?**

2. Next, filter by the **Data Center** annotation created in a previous step to view a subset of assets.

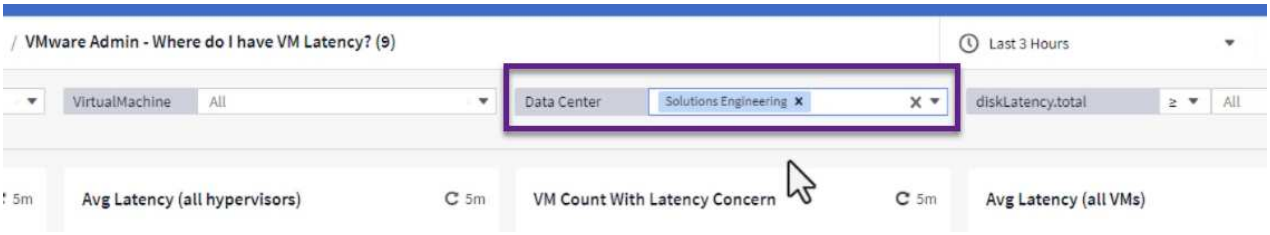

3. This dashboard shows a list of the top 10 VMs by average latency. From here click on the VM of concern to drill into its details.

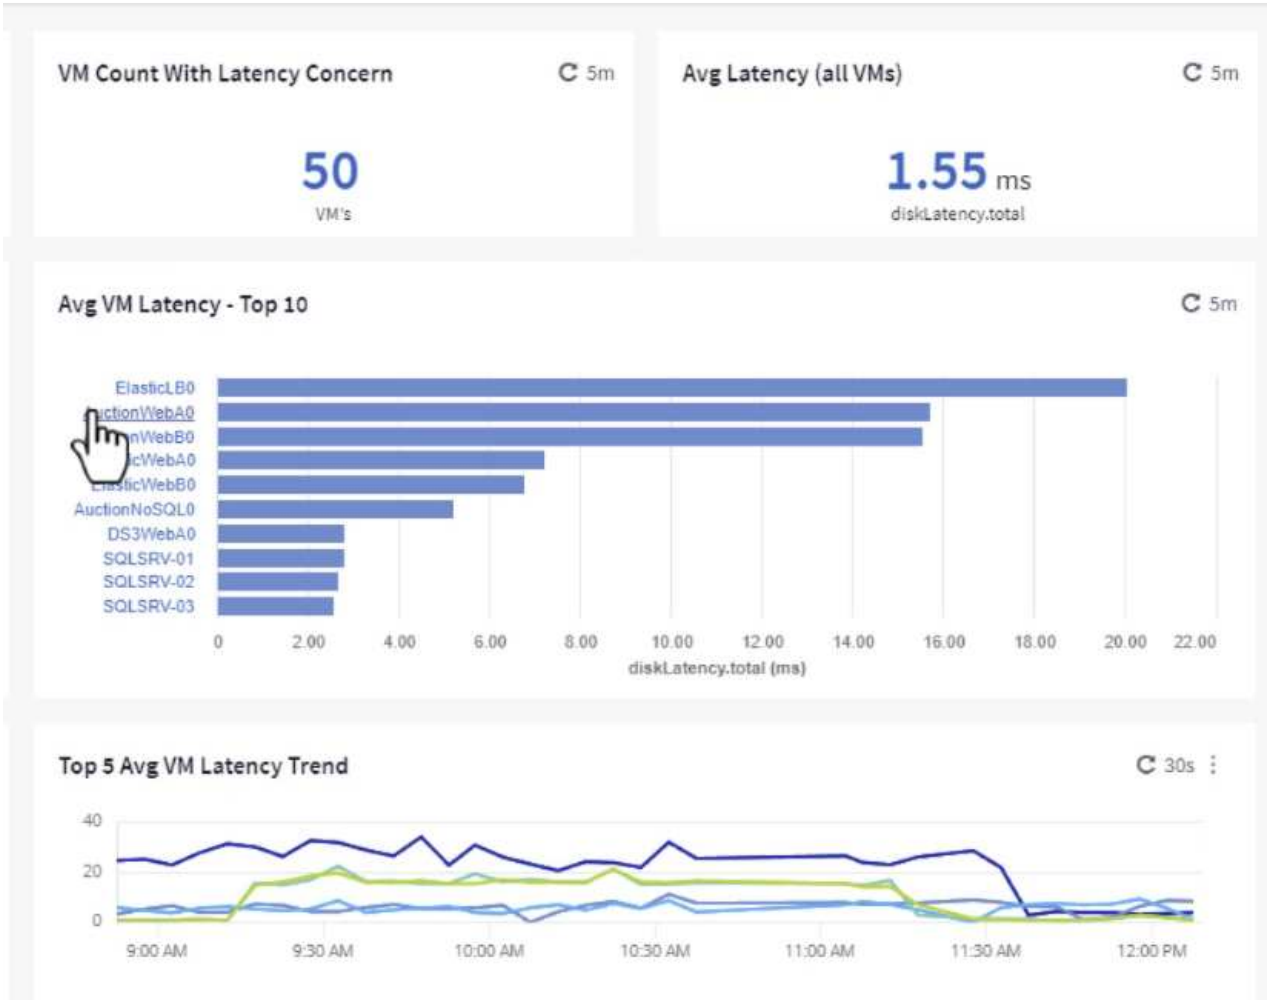

4. The VMs potentially causing workload contention are listed and available. Drill into these VMs performance metrics to investigate any potential issues.

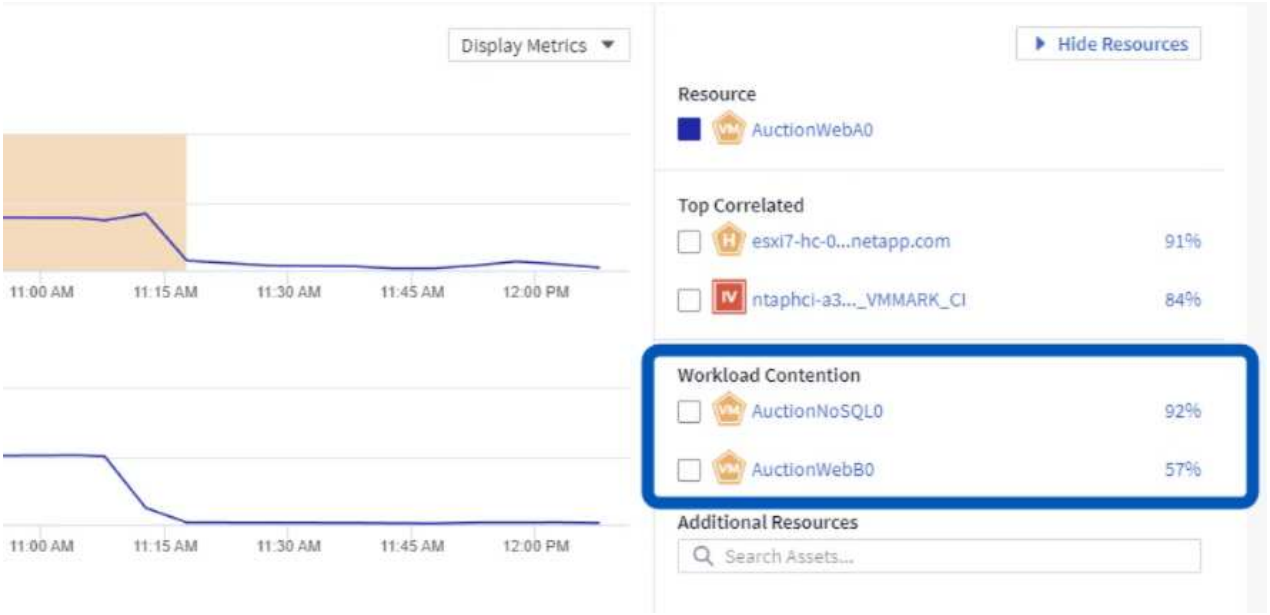

View over and under utilized resources in Cloud Insights

By matching VM resources to actual workload requirements, resource utilization can be optimized, leading to cost savings on infrastructure and cloud services. Data in Cloud Insights can be customized to easily display over or under utilized VMs.

Identify opportunities to right size VMs

1. In this example access a dashboard available in the **Gallery** called **VMware Admin - Where are opportunities to right size?**

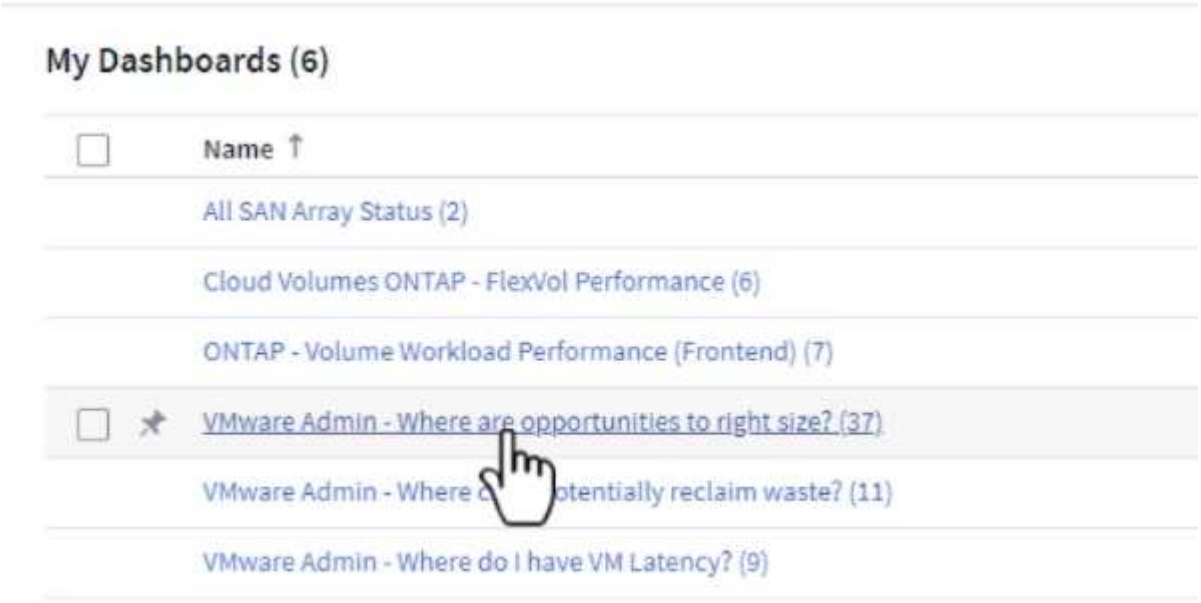

2. First filter by all of the ESXi hosts in the cluster. You can then see ranking of the top and bottom VMs by memory and CPU utilization.

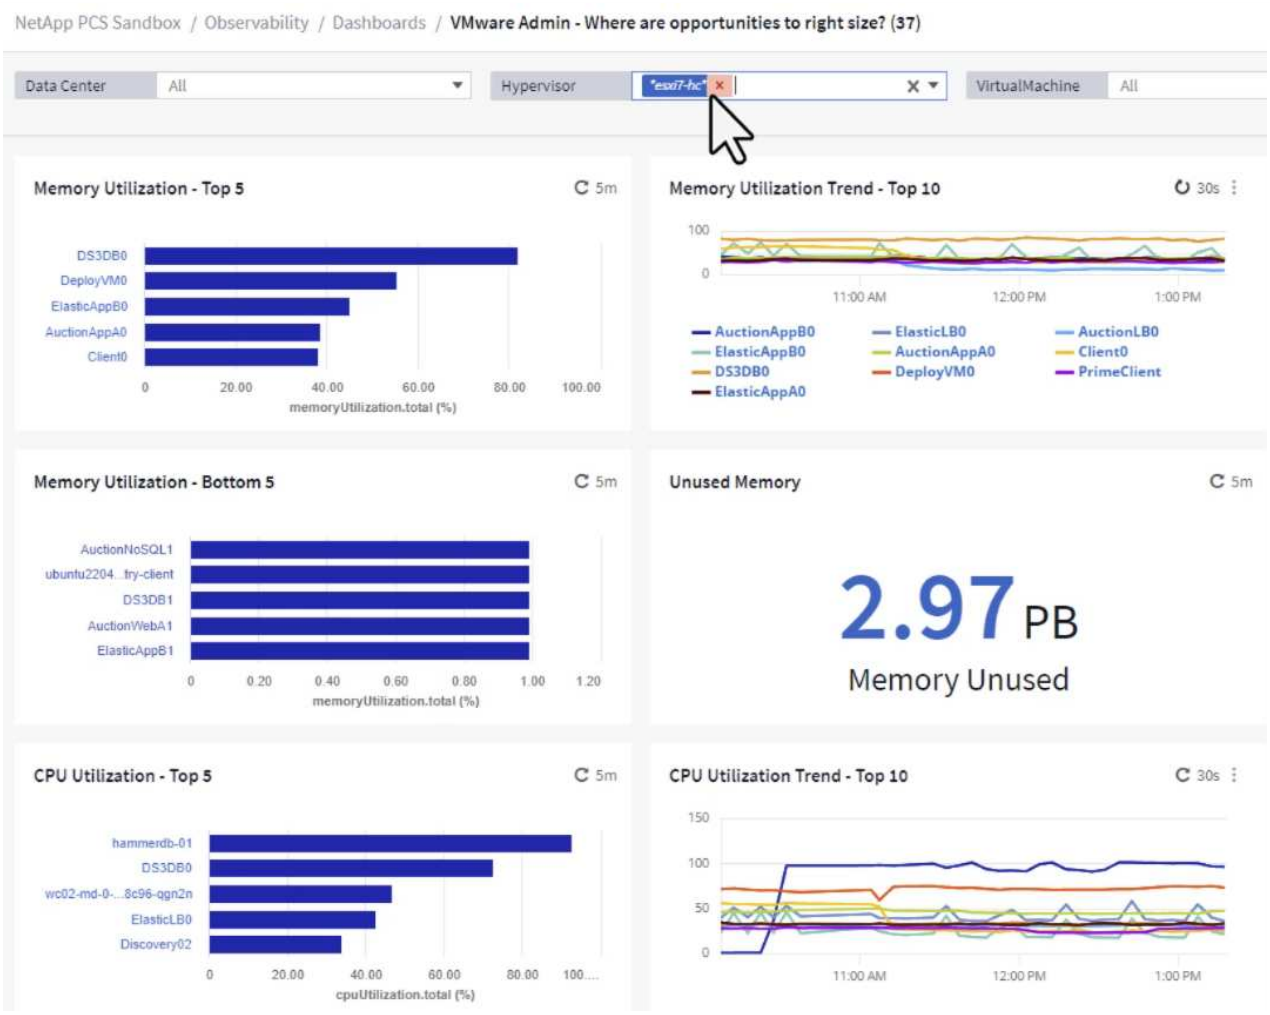

3. Tables allow sorting and provide more detail based on the columns of data chosen.

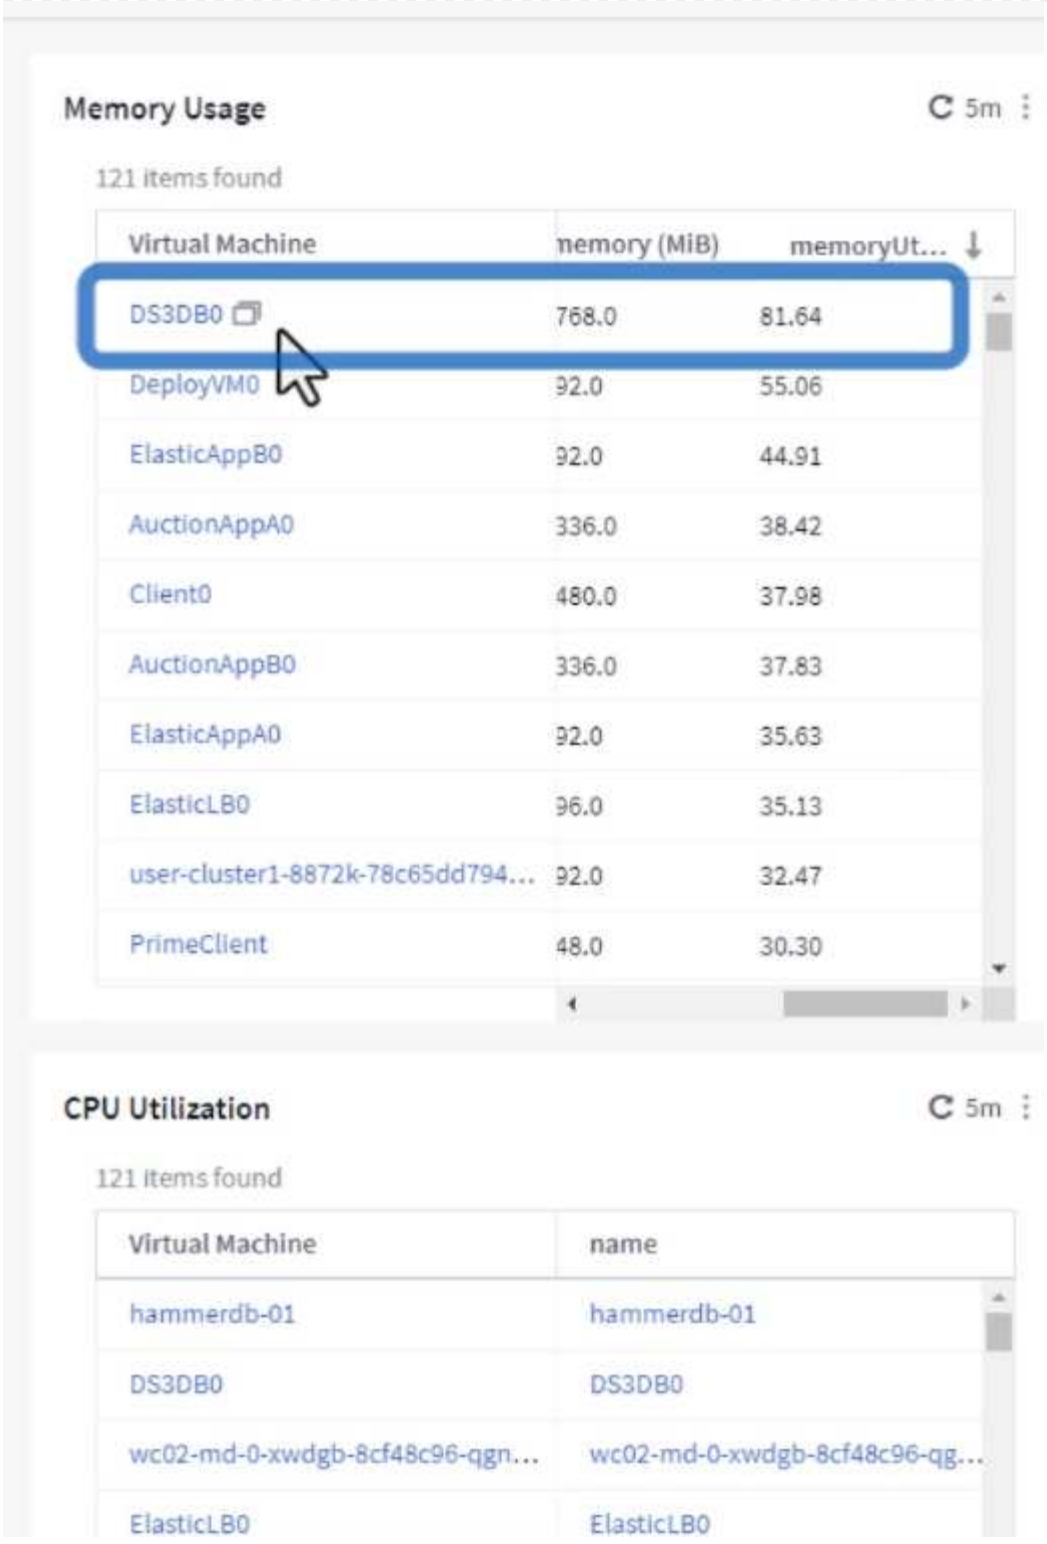

4. Another dashboard called **VMware Admin - Where can I potentially reclaim waste?** shows powered off VM's sorted by their capacity use.

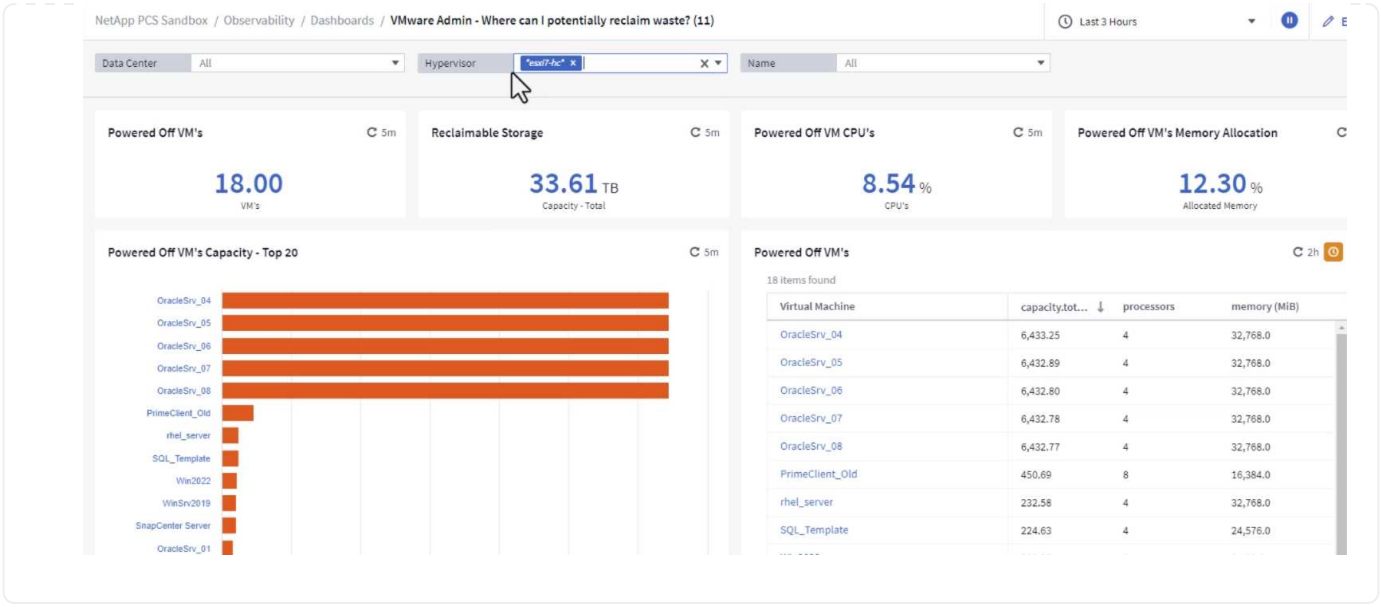

Use queries to isolate and sort metrics

The amount of data captured by Cloud Insights is quite comprehensive. Metric queries provide a powerful way to sort and organize large amounts of data in useful ways.

View a detailed VMware query under ONTAP Essentials

1. Navigate to **ONTAP Essentials > VMware** to access a comprehensive VMware metric query.

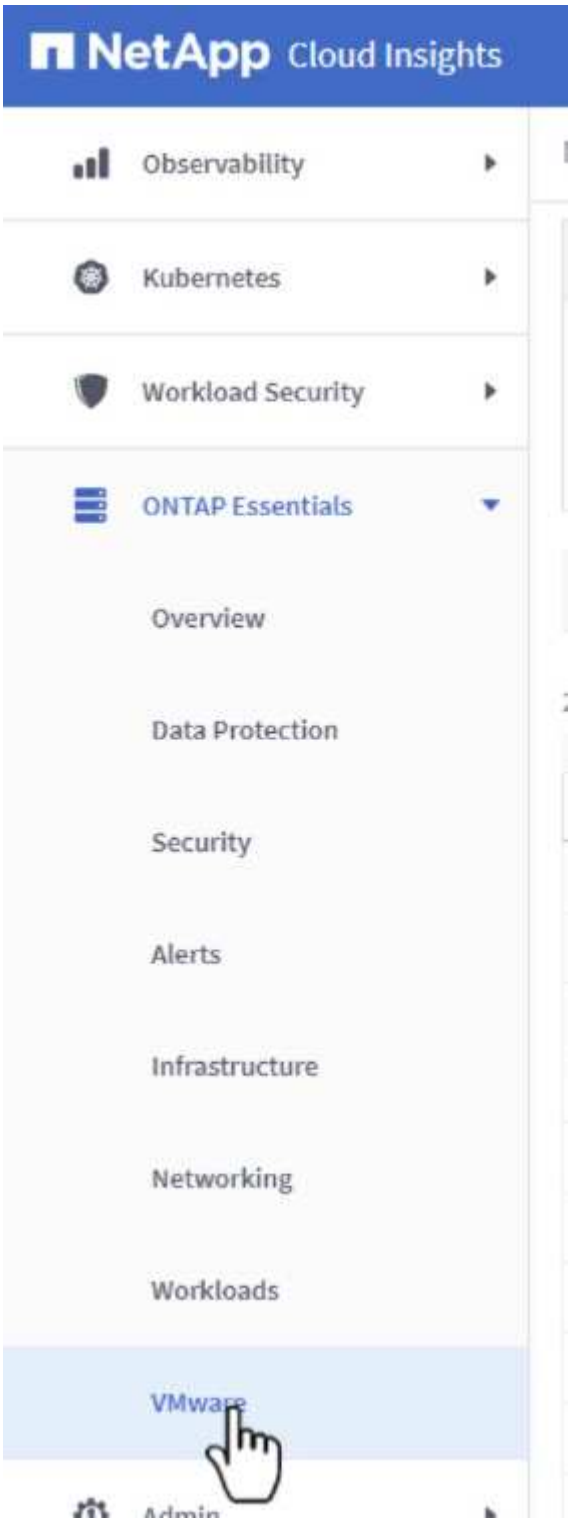

2. In this view you are presented with multiple options for filtering and grouping the data at the top. All columns of data are customizable and additional columns can be easily added.

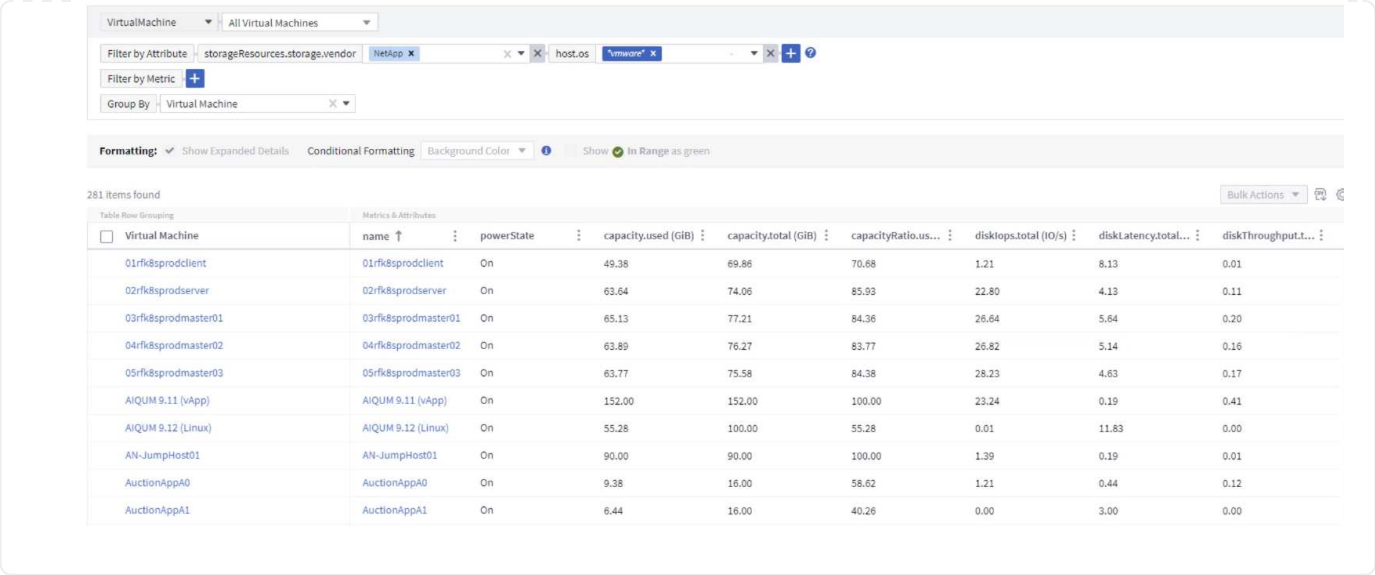

Conclusion

This solution was designed as a primer to learn how to get started with NetApp Cloud Insights and show some of the powerful capabilities that this observability solution can provide. There are hundreds of dashboards and metric queries built into the product which makes it easy to get going immediately. The full version of Cloud Insights is available as a 30-day trial and the basic version is available free to NetApp customers.

Additional Information

To learn more about the technologies presented in this solution refer to the following additional information.

- NetApp BlueXP and Cloud Insights landing page
- NetApp Cloud Insights documentation

Demos and Tutorials

Virtualization Videos and Demos

See the following videos and demos highlighting specific features of the hybrid cloud, virtualization, and container solutions.

NetApp ONTAP Tools for VMware vSphere

ONTAP Tools for VMware - Overview

VMware iSCSI Datastore Provisioning with ONTAP

VMware NFS Datastore Provisioning with ONTAP

SnapCenter Plug-in for VMware vSphere

NetApp SnapCenter software is an easy-to-use enterprise platform to securely coordinate and manage data protection across applications, databases, and file systems.

The SnapCenter Plug-in for VMware vSphere allows you to perform backup, restore, and attach operations for VMs and backup and mount operations for datastores that are registered with SnapCenter directly within VMware vCenter.

For more information about NetApp SnapCenter Plug-in for VMware vSphere, see the NetApp SnapCenter Plug-in for VMware vSphere Overview.

SnapCenter Plug-in for VMware vSphere - Solution Pre-Requisites

SnapCenter Plug-in for VMware vSphere - Deployment

SnapCenter Plug-in for VMware vSphere - Backup Workflow

SnapCenter Plug-in for VMware vSphere - Restore Workflow

SnapCenter - SQL Restore Workflow

3-2-1 Data Protection Solutions

3-2-1 data protection solutions combine on-premises primary and secondary backups, using SnapMirror technology, with replicated copies to object storage using BlueXP backup and recovery.

3-2-1 Data Protection for VMFS Datastores with SnapCenter Plug-in for VMware vSphere and BlueXP Backup and Recovery for Virtual Machines

VMware Cloud on AWS with AWS FSx for NetApp ONTAP

Windows Guest Connected Storage with FSx ONTAP using iSCSI

Linux Guest Connected Storage with FSx ONTAP using NFS

VMware Cloud on AWS TCO savings with Amazon FSx for NetApp ONTAP

VMware Cloud on AWS supplemental datastore w/ Amazon FSx for NetApp ONTAP

VMware HCX Deployment and Configuration Setup for VMC

vMotion Migration Demonstration with VMware HCX for VMC and FSxN

Cold Migration Demonstration with VMware HCX for VMC and FSxN

Azure VMware Services on Azure with Azure NetApp Files (ANF)

- Azure VMware Solution supplemental datastore overview with Azure NetApp Files
- Azure VMware Solution DR with Cloud Volumes ONTAP, SnapCenter and JetStream
- 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

VMware Cloud Foundation with NetApp ONTAP

- NFS Datastores as Principal Storage for VCF Workload Domains
- iSCSI Datastores as Supplemental Storage for VCF Management Domains

NetApp with VMware Tanzu

VMware Tanzu enables customers to deploy, administer, and manage their Kubernetes environment through vSphere or the VMware Cloud Foundation. This portfolio of products from VMware allows customer to manage all their relevant Kubernetes clusters from a single control plane by choosing the VMware Tanzu edition that best suits their needs.

For more information about VMware Tanzu, see the VMware Tanzu Overview. This review covers use cases, available additions, and more about VMware Tanzu.

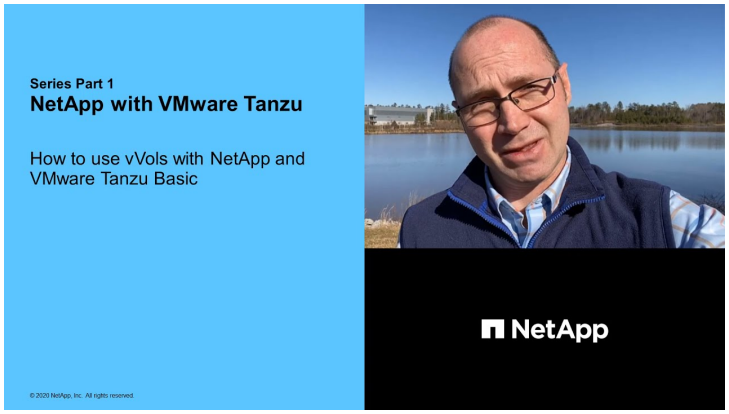

How to use vVols with NetApp and VMware Tanzu Basic, part 1

How to use vVols with NetApp and VMware Tanzu Basic, part 2

How to use vVols with NetApp and VMware Tanzu Basic, part 3

NetApp Cloud Insights

NetApp Cloud Insights is comprehensive monitoring and analytics platform designed to provide visibility and control over your on-premises and cloud infrastructure.

NetApp Cloud Insights - Observability for the Modern Datacenter

NetApp Hyper-V Virtualization Solutions

Getting Started

Deploying Microsoft Hyper-V on NetApp Storage

Deploying Microsoft Hyper-V on NetApp Storage

The Windows Server platform uses the Hyper-V role to provide virtualization technology. Hyper-V is one of many optional roles that are offered with Windows Server.

Overview

The Hyper-V role enables us to create and manage a virtualized computing environment by using virtualization technology built into Windows Server. The Hyper-V technology virtualizes hardware to provide an environment in which you can run multiple operating systems at the same time on one physical computer. Hyper-V enables you to create and manage virtual machines and their resources. Each virtual machine is an isolated, virtualized computer system that can run its own operating system. Hyper-V provides infrastructure to virtualize applications and workloads that supports a variety of business goals aimed at improving efficiency and reducing costs which is a perfect alternative to VMware® vSphere, especially when organizations are looking for co-existence of multiple hypervisors during the current market conditions.

Audience

This document describes the architecture and deployment procedures for the Hyper-V Cluster configuration with the NetApp ONTAP systems. The intended audience for this document includes sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to deploy Hyper-V as the primary or as an alternate hypervisor.

Architecture

The architecture described in this document specifically includes Microsoft® Windows Server® 2022 and Hyper-V® virtualization. NetApp strongly recommends virtualization software and infrastructure management software as part of every deployment. The configuration uses the best practices for each component to enable a reliable, enterprise-class infrastructure.

Use Case Summary

This document describes the deployment procedures and best practices to set up Hyper-V cluster to optimally perform as a workload on Microsoft Windows Server 2022 using NetApp All-flash FAS and ASA arrays models. The server operating system/hypervisor is Microsoft Windows Server 2022. The guidance covers NetApp storage systems that serve data over storage area network (SAN) and network-attached storage (NAS) protocols.

Deploying Microsoft Hyper-V on NetApp Storage: Pre-Requisities

This topic provides steps to configure and deploy a two-node failover cluster and clustered Hyper-V virtual machines leveraging ONTAP storage system.

Pre-requisites for Deployment Procedure

- All hardware must be certified for the version of Windows Server that you are running, and the complete failover cluster solution must pass all tests in the Validate a Configuration Wizard
- Hyper-V nodes joined to the domain controller (recommended) and appropriate connectivity between each other.
- Every Hyper-V node should be configured identically.
- Network adapters and designated virtual switches configured on each Hyper-V server for segregated traffic for mgmt, ISCSI, SMB, live migrate.
- The failover cluster feature is enabled on each Hyper-V server.
- SMB shares or CSVs are used as shared storage to store VMs and their disks for Hyper-V clustering.
- Storage should not be shared between different clusters. Plan for one or multiple CSV/CIFS share per cluster.
- If the SMB share is used as shared storage, then permissions on the SMB share must be configured to grant access to the computer accounts of all the Hyper-V nodes in the cluster.

For more information, see:

- System Requirements for Hyper-V on Windows Server
- Validate Hardware for a Failover Cluster
- Deploy a Hyper-V Cluster

Installing Windows Features

The following steps describe how to install the required Windows Server 2022 features.

All Hosts

1. Prepare the windows OS 2022 with necessary updates and device drivers on all the designated nodes.
2. Log into each Hyper-V node using the administrator password entered during installation.
3. Launch a PowerShell prompt by right clicking the PowerShell icon in the taskbar and selecting `Run as Administrator`.
4. Add the Hyper-V, MPIO, and clustering features.

```
Add-WindowsFeature Hyper-V, Failover-Clustering, Multipath-IO `-
IncludeManagementTools –Restart
```

Configuring Networks

Proper network planning is key to achieving fault tolerant deployment. Setting up distinct physical network adapters for each type of traffic was the standard suggestion for a failover cluster. With the ability to add virtual

network adapters, switch embedded teaming (SET) and features like Hyper-V QoS introduced, condense network traffic on fewer physical adapters. Design the network configuration with quality of service, redundancy, and traffic isolation in mind. Configuring network isolation techniques like VLANs in conjunction with traffic isolation techniques provides redundancy for the traffic and quality of service which would improve and add consistency to storage traffic performance.

It is advised to separate and isolate specific workloads using multiple logical and/or physical networks. Typical network traffic examples that are typically divided into segments are as follows:

- ISCSI Storage network.
- CSV (Cluster Shared Volume) or Heartbeat network.
- Live Migration
- VM network
- Management network

Note: When iSCSI is used with dedicated NICs, then using any teaming solution is not recommended and MPIO/DSM should be used.

Note: Hyper-V networking best practices also do not recommend using NIC teaming for SMB 3.0 storage networks in Hyper-V environment.

For additional information, refer to Plan for Hyper-V networking in Windows Server

Deciding on Storage Design for Hyper-V

Hyper-V supports NAS (SMB3.0) and Block storage (iSCSI/FC) as the backing storage for virtual machines. NetApp supports SMB3.0, iSCSI and FC protocol which can be used as native storage for VMs - Cluster Shared Volumes (CSV) using iSCSI/FC and SMB3. Customers can also use SMB3 and iSCSI as guest connected storage options for workloads that require direct access to the storage. ONTAP provides flexible options with unified storage (All Flash Array) - for workload that requires mixed protocol access and SAN optimized storage (All SAN Array) for SAN only configurations.

The decision to use SMB3 vs iSCSI/FC is driven by the existing infrastructure in place today, SMB3/iSCSI allow customers to use existing network infrastructure. For customers that have existing FC infrastructure can leverage that infrastructure and present storage as FC based Clustered Shared Volumes.

Note: A NetApp storage controller running ONTAP software can support the following workloads in a Hyper-V environment:

- VMs hosted on continuously available SMB 3.0 shares
- VMs hosted on Cluster Shared Volume (CSV) LUNs running on iSCSI or FC
- In-Guest storage and pass through disks to guest virtual machines

Note: Core ONTAP features such as thin provisioning, deduplication, compression, data compaction, flex clones, snapshots, and replication work seamlessly in the background regardless of the platform or operating system and provide significant value for the Hyper-V workloads. The default settings for these features are optimal for Windows Server and Hyper-V.

Note: MPIO is supported on the guest VM using in-guest initiators if multiple paths are available to the VM, and the multipath I/O feature is installed and configured.

Note: ONTAP supports all major industry-standard client protocols: NFS, SMB, FC, FCoE, iSCSI, NVMe/FC, and S3. However, NVMe/FC and NVMe/TCP are not supported by Microsoft.

Installing NetApp Windows iSCSI Host Utilities

The following section describes how to perform an unattended installation of the NetApp Windows iSCSI Host Utilities. For detailed information regarding the installation see the Install Windows Unified Host Utilities 7.2 (or the latest supported version)

All Hosts

1. Download Windows iSCSI Host Utilities
2. Unblock the downloaded file.

```
Unblock-file ~\Downloads\netapp_windows_host_utilities_7.2_x64.msi
```

3. Install the Host Utilities.

```
~\Downloads\netapp_windows_host_utilities_7.2_x64.msi /qn
"MULTIPATHING=1"
```

Note: The system will reboot during this process.

Configuring Windows Host iSCSI initiator

The following steps describe how to configure the built in Microsoft iSCSI initiator.

All Hosts

1. Launch a PowerShell prompt by right clicking the PowerShell icon in the taskbar and selecting Run as Administrator.
2. Configure the iSCSI service to start automatically.

```
Set-Service -Name MSiSCSI -StartupType Automatic
```

3. Start the iSCSI service.

```
Start-Service -Name MSiSCSI
```

4. Configure MPIO to claim any iSCSI device.

```
Enable-MSDSMAutomaticClaim -BusType iSCSI
```

5. Set the default load balance policy of all newly claimed devices to round robin.

```
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy RR
```

6. Configure an iSCSI target for each controller.

```
New-IscsiTargetPortal -TargetPortalAddress <<iscsia_lif01_ip>>
-InitiatorPortalAddress <iscsia_ipaddress>

New-IscsiTargetPortal -TargetPortalAddress <<iscsib_lif01_ip>>
-InitiatorPortalAddress <iscsib_ipaddress

New-IscsiTargetPortal -TargetPortalAddress <<iscsia_lif02_ip>>
-InitiatorPortalAddress <iscsia_ipaddress>

New-IscsiTargetPortal -TargetPortalAddress <<iscsib_lif02_ip>>
-InitiatorPortalAddress <iscsib_ipaddress>
```

7. Connect a session for each iSCSI network to each target.

```
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true
-IsMultipathEnabled $true -InitiatorPo rtalAddress <iscsia_ipaddress>

Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true
-IsMultipathEnabled $true -InitiatorPo rtalAddress <iscsib_ipaddress>
```

Note: Add multiple sessions (min of 5-8) for increased performance and utilizing the bandwidth.

Creating a Cluster

One Server Only

1. Launch a PowerShell prompt with administrative permissions, by right clicking the PowerShell icon and selecting `Run as Administrator`.
2. Create a new cluster.

```
New-Cluster -Name <cluster_name> -Node <hostnames> -NoStorage
-StaticAddress <cluster_ip_address>
```

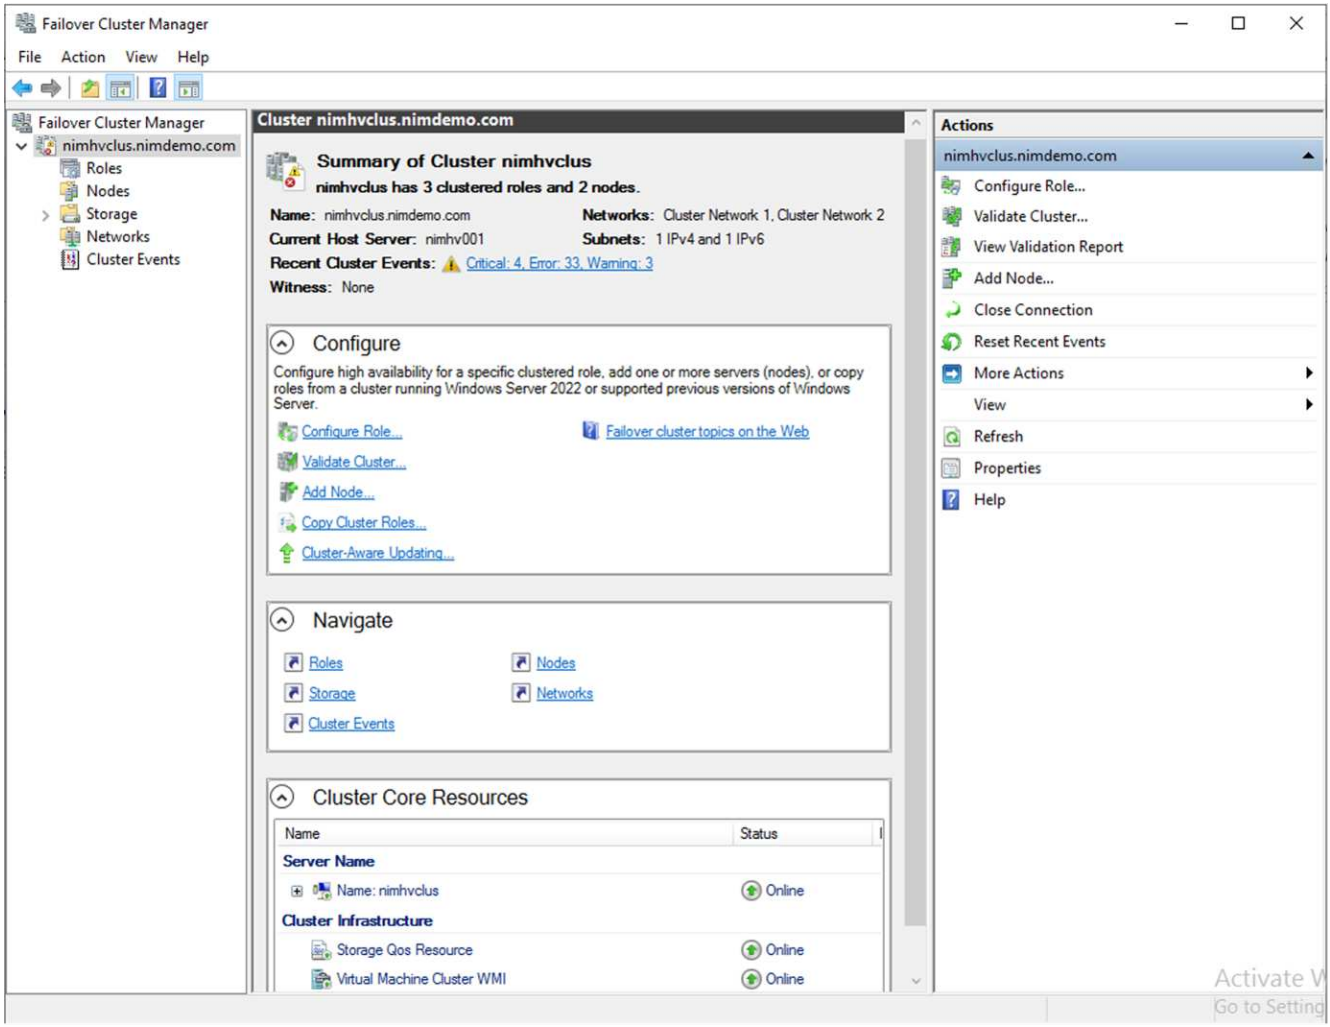

3. Select the appropriate cluster network for Live migration.
4. Designate the CSV network.

```
(Get-ClusterNetwork -Name Cluster).Metric = 900
```

5. Change the cluster to use a quorum disk.
   a. Launch a PowerShell prompt with administrative permissions by right clicking the PowerShell icon and selecting 'Run as Administrator'.

```
start-ClusterGroup "Available Storage"| Move-ClusterGroup -Node
$env:COMPUTERNAME
```

   b. In Failover Cluster Manager, select `Configure Cluster Quorum Settings`.

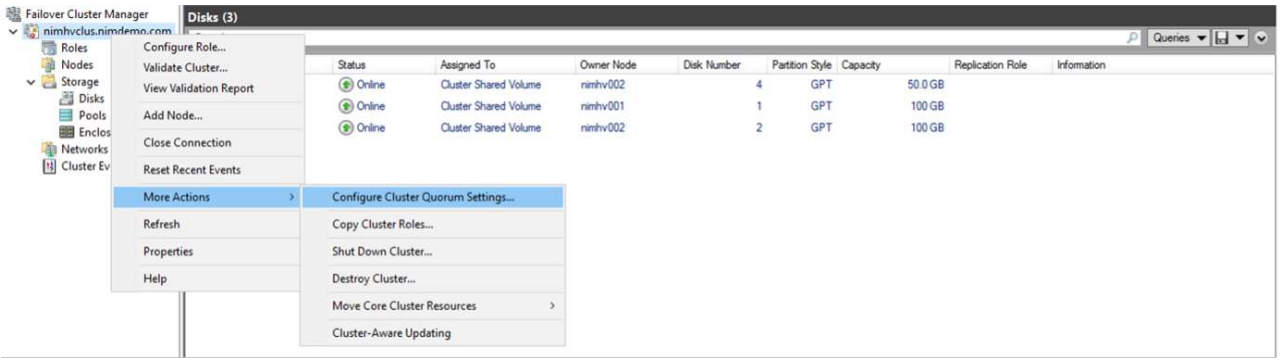

c. Click Next through the Welcome page.

d. Select the quorum witness and click Next.

e. Select Configure a disk witness` and click Next.

f. Select Disk W: from the available storage and click Next.

g. Click Next through the confirmation page and Finish on the summary page.

   For more detailed information about quorum and witness, see Configuring and manage quorum

6. Run the Cluster Validation wizard from Failover Cluster Manager to validate deployment.
7. Create CSV LUN to store virtual machine data and create highly available virtual machines via Roles within Failover Cluster Manager.

Deploying Microsoft Hyper-V on NetApp Storage: Considerations

This step is vital to ascertain that the applications, services, and workloads can operate effectively in the Hyper-V environment. Compatibility checks must encompass operating system versions, Windows server versions, application dependencies, database systems, and any specific configurations or customisations that exist in the existing environment.

Right sizing the storage

Before deploying the workload or migrating from existing hypervisor, ensure the workload is sized to meet the required performance. This can be easily done by collecting performance data for each individual VM that collects statistics for CPU (used/provisioned), Memory (used/provisioned), Storage (provisioned/utilized), Network throughput and latency along with aggregation of the Read/Write IOPs, throughput and block size. These parameters are mandatory for have a successful deployment and to correctly size the storage array and workload hosts.

Note: Plan for IOPS and capacity when sizing storage for Hyper-V and associated workloads.

Note: For higher-I/O intensive VMs or those that require lots of resources and capacity, segregate the OS and data disks. Operating system and application binaries change infrequently, and volume crash consistency is acceptable.

Note: Use Guest connected storage (aka in-guest) for high performance data disks than using VHDs. This helps with easier cloning process as well.

Enhance Virtual Machine performance

Choose the right amount of RAM and vCPUs for optimal performance along with attaching multiple disks to a

single virtual SCSI controller. Using fixed VHDx is still recommended as the primary choice for virtual disks for deployments and there are no restrictions for using any type of VHDX virtual disks.

Note: Avoid installing unnecessary roles on Windows Server that will not be utilized.

Note: Choose Gen2 as the generation for virtual machines able to load VMs from the SCSI controller and is based on the VMBUS and VSP / VSC architecture for the boot level, which significantly increases the overall VM performance.

Note: Avoid making frequent checkpoints because it has a negative impact on the performance of the VM.

SMB3.0 Design and Consideration

SMB 3.0 file shares can be used as shared storage for Hyper-V. ONTAP supports nondisruptive operations over SMB shares for Hyper-V. Hyper-V can use SMB file shares to store virtual machine files, such as configuration, snapshots, and virtual hard disk (VHD) files. Use dedicated ONTAP CIFS SVM for SMB3.0 based shares for Hyper-V. The volumes used to store virtual machine files must be created with NTFS security-style volumes. Connectivity between Hyper-V hosts and the NetApp array is recommended on a 10GB network if one is available. In case of 1GB network connectivity, NetApp recommends creating an interface group consisting of multiple 1GB ports. Connect each NIC serving SMB multichannel to its dedicated IP subnet so that each subnet provides a single path between the client and server.

Key Points

- Enable SMB multi-channel on ONTAP SVM
- ONTAP CIFS SVMs should have at least one data LIF on each node in a cluster.
- Shares used must be configured with the continuously available property set.
- ONTAP One is now included on every AFF (A-Series and C-Series), All-SAN Array (ASA), and FAS system. Hence there is no separate licenses needed.
- For Shared VHDx, use guest connected iSCSI LUN

Note: ODX is supported and works across protocols. Copying data between a file share and iSCSI or an FCP-attached LUN also utilizes ODX.

Note: Time settings on nodes in the cluster should be set up accordingly. Network Time Protocol (NTP) should be used if the NetApp CIFS server must participate in the Windows Active Directory (AD) domain.

Note: Large MTU values must be enabled through the CIFS server. Small packet sizes might result in performance degradation.

Provisioning SMB volume

1. Verify that the required CIFS server options are enabled on the storage virtual machine (SVM)
2. The following options should be set to true: smb2-enabled smb3-enabled copy-offload-enabled shadowcopy-enabled is-multichannel-enabled is-large-mtu-enabled

```
HV_NestedCluster::*> vserver cifs options show -vserver NestedHVsvm01 -fields copy-offload-enabled, is-multichannel-enabled, is-large-mtu-enabled, smb2-enabled, smb3-enabled, copy-offloa
d-enabled, shadowcopy-enabled
vserver       smb2-enabled smb3-enabled copy-offload-enabled shadowcopy-enabled is-multichannel-enabled is-large-mtu-enabled
------------- ------------ ------------ -------------------- ------------------ ----------------------- --------------------
NestedHVsvm01 true         true         true                 true               true                    true
```

Activate Windows
Go to Settings to activate Windows.

3. Create NTFS data volumes on the storage virtual machine (SVM) and then configure continuously available shares for use with Hyper-V

```
HV_NestedCluster::*> volume create -vserver NestedHVsvm01 -volume hvdemosmb -aggregate HV_NestedCluster_01_VM_DISK_1 -size 500GB -security-style ntfs -junction-path /hvdemosmb
[Job 169] Job succeeded: Successful
```

Note: Nondisruptive operations for Hyper-V over SMB do not work correctly unless the volumes used in the configuration are created as NTFS security-style volumes.

4. Enable continuously available and configure NTFS permissions on the share to include Hyper-V nodes with full control.

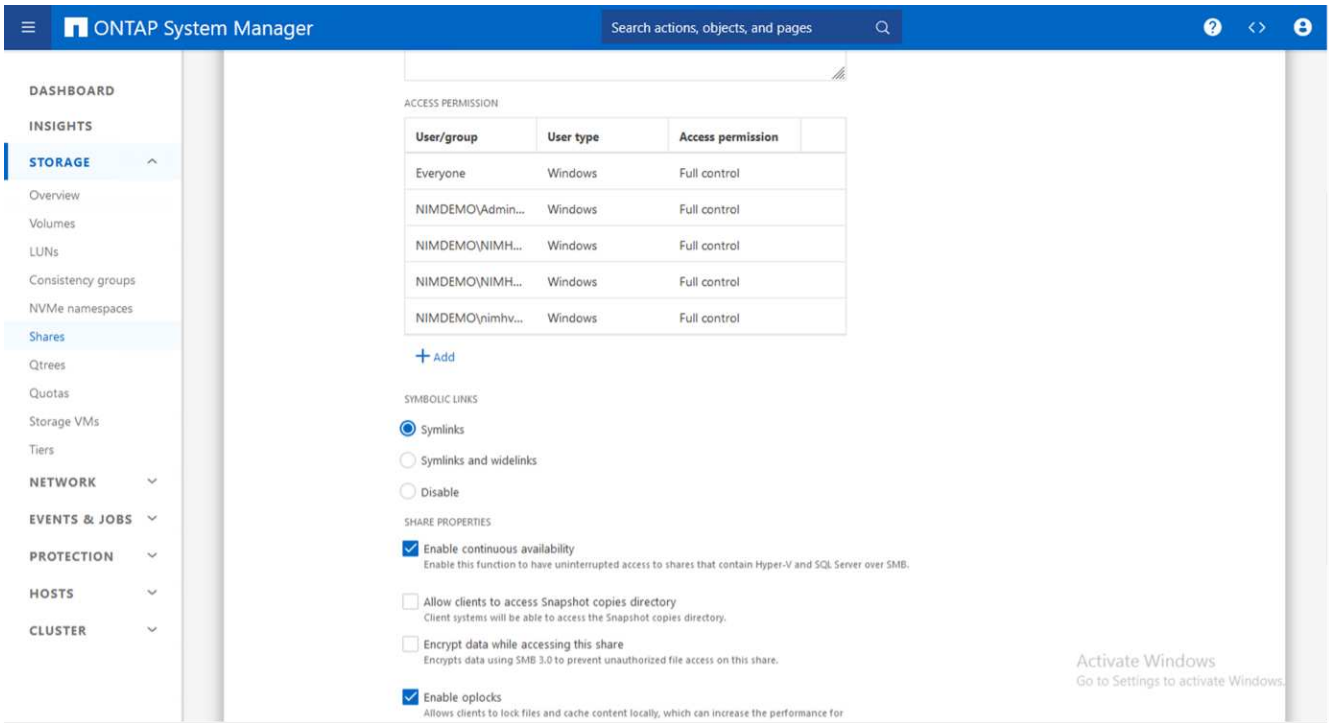

For detailed best practices guidance, see Deployment Guidelines and best practices for Hyper-V.

For additional information, refer to SMB server and volume requirements for Hyper-V over SMB
.

Block Protocol Design and Consideration

Key Points

- Use multipathing (MPIO) on hosts to manage the multiple paths. Create more paths as needed, either to facilitate data mobility operations or to leverage additional I/O resources, but do not exceed the maximum number of paths a host OS can support.
- Install the Host Utilities Kit on hosts accessing the LUNs.
- Create a minimum of 8 volumes.

Note: Use one LUN per volume, thus having 1:1 mapping for LUN to CSV ratio.

- An SVM should have one LIF per Ethernet network or Fibre Channel fabric on every storage controller that is going to serve data using iSCSI or Fibre Channel.
- SVMs serving data with FCP, or iSCSI need an SVM management interface.

Provisioning ISCSI volume

To provision ISCSI volume, ensure the following pre-requisites are met.

- The storage virtual machine (SVM) should have the iSCSI protocol enabled and the appropriate logical interfaces (LIFs) created.
- The designated aggregate must have enough free space to contain the LUN.

Note: By default, ONTAP uses Selective LUN Map (SLM) to make the LUN accessible only through paths on the node owning the LUN and its high-availability (HA) partner.

- Configure all the iSCSI LIFs on every node for LUN mobility in case the LUN is moved to another node in the cluster.

Steps

1. Use System Manager and navigate to the LUNs window (ONTAP CLI can be used for the same operation).
2. Click Create.
3. Browse and select the designated SVM in which the LUNs to be created and the Create LUN Wizard is displayed.
4. On the General Properties page, select Hyper-V for LUNs containing virtual hard disks (VHDs) for Hyper-V virtual machines.

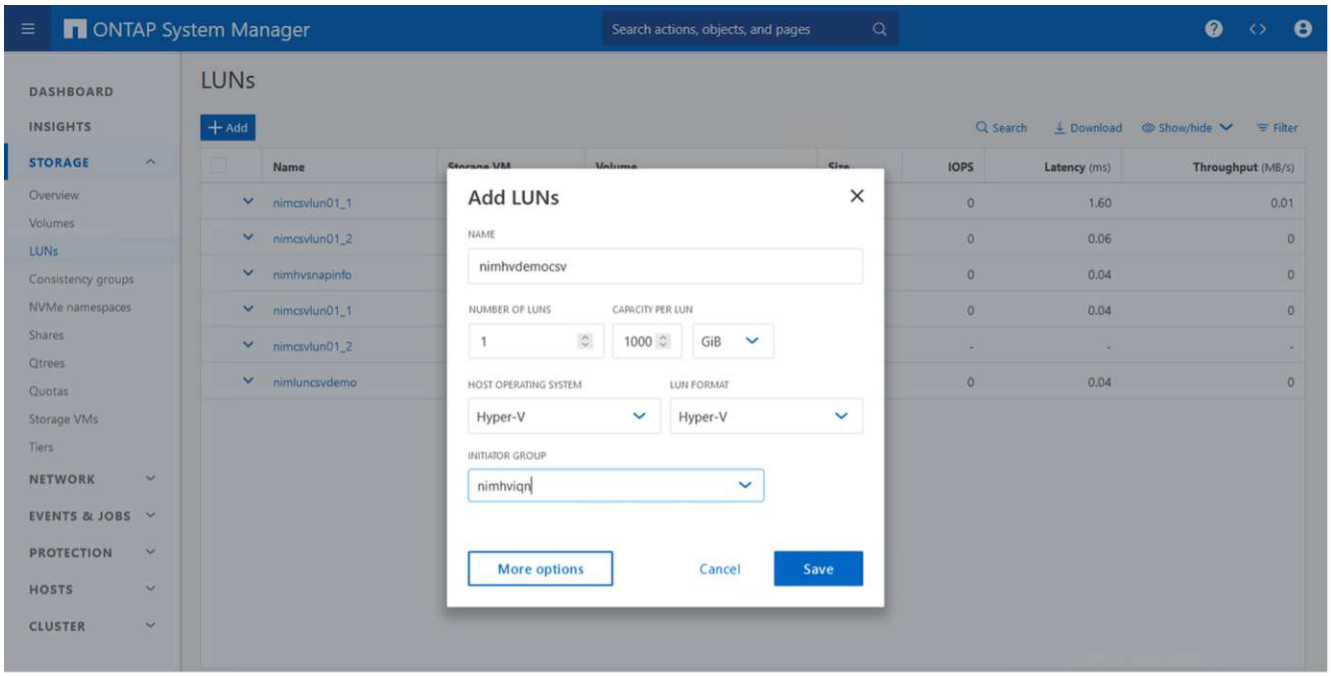

5. <click on More options> On the LUN Container page, select an existing FlexVol volume otherwise a new volume will be created.
6. <click on More options> On the Initiators Mapping page, click Add Initiator Group, enter the required information on the General tab, and then on the Initiators tab, enter the iSCSI initiator node name of the hosts.
7. Confirm the details, and then click Finish to complete the wizard.

Once the LUN is created, go to the Failover Cluster Manager. To add a disk to CSV, the disk must be added to the Available Storage group of the cluster (if it is not already added), and then add the disk to CSV on the

cluster.

Note: The CSV feature is enabled by default in Failover Clustering.

Adding a disk to Available Storage:

1. In Failover Cluster Manager, in the console tree, expand the name of the cluster, and then expand Storage.
2. Right-click Disks, and then select Add Disk. A list appears showing the disks that can be added for use in a failover cluster.
3. Select the disk or disks you want to add, and then select OK.
4. The disks are now assigned to the Available Storage group.
5. Once done, select the disk that was just assigned to Available Storage, right-click the selection, and then select Add to Cluster Shared Volumes.

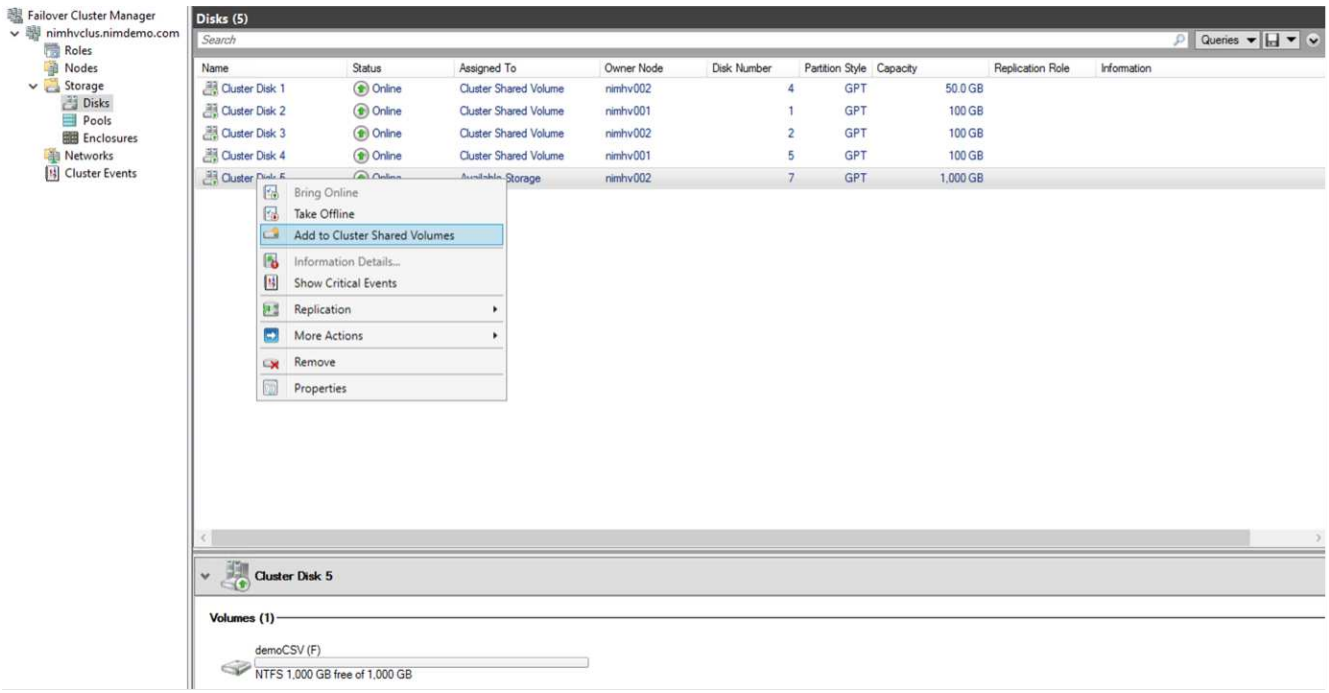

6. The disks are now assigned to the Cluster Shared Volume group in the cluster. The disks are exposed to each cluster node as numbered volumes (mount points) under the %SystemDrive%ClusterStorage folder. The volumes appear in the CSVFS file system.

For additional information, refer to Use Cluster Shared Volumes in a failover cluster.

Create highly available virtual machines:

To create a highly available virtual machine, follow the below steps:

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. Click Roles.
3. In the Actions pane, click Virtual Machines, and then click New Virtual Machine. The New Virtual Machine Wizard appears. Click Next.
4. On the Specify Name and Location page, specify a name for the virtual machine, such as nimdemo. Click Store the virtual machine in a different location, and then type the full path or click Browse and navigate to

the shared storage.

5. Assign Memory and configure network adapter to the virtual switch that is associated with the physical network adapter.
6. On the Connect Virtual Hard Disk page, click Create a virtual hard disk.
7. On the Installation Options page, click Install an operating system from a boot CD/DVD-ROM. Under Media, specify the location of the media, and then click Finish.
8. The virtual machine is created. The High Availability Wizard in Failover Cluster Manager then automatically configures the virtual machine for high availability.

Fast Provisioning of Virtual Disks Using ODX Feature

The ODX feature in ONTAP allows making copies of master VHDXs by simply copying a master VHDX file hosted by ONTAP storage system. Because an ODX-enabled copy does not put any data on the network wire, the copy process happens on the NetApp storage side and as a result can be up to six to eight times faster. General considerations for fast provisioning include master sysprepped images stored on file shares and regular copy processes initiated by the Hyper-V host machines.

Note: ONTAP supports ODX for both the SMB and SAN protocols.

Note: To take advantage of the use cases for ODX copy offload pass-through with Hyper-V, the guest operating system must support ODX, and the guest operating system's disks must be SCSI disks backed by storage (either SMB or SAN) that supports ODX. IDE disks on the guest operating system do not support ODX pass-through.

Performance optimization

Although the recommended number of VMs per CSV is subjective, numerous factors determine the optimum number of VMs that can be placed on each CSV or SMB volume. Although most administrators only consider capacity, the amount of concurrent I/O being sent to the VHDx is one of the most key factors for overall performance. The easiest way to control performance is by regulating the number of virtual machines that are placed on each CSV or share. If the concurrent virtual machine I/O patterns are sending too much traffic to the CSV or share, the disk queues fill, and higher latency are generated.

SMB Volume and CSV sizing

Ensure the solution is adequately sized end-to-end to avoid bottlenecks and when a volume is created for Hyper-V VM storage purposes, the best practice is to create a volume no larger than required. Right sizing volumes prevent accidentally placing too many virtual machines on the CSV and decreases the probability of resource contention. Each cluster shared volume (CSV) supports one VM or multiple VMs. The number of VMs to place on a CSV is determined by the workload and business preferences, and how ONTAP storage features such as snapshots and replication will be used. Placing multiple VMs on a CSV is a good starting point in most deployment scenarios. Adjust this approach for specific use cases to meet performance and data protection requirements.

Since volumes and VHDx sizes can be easily increased, if a VM needs extra capacity, it is not necessary to size CSVs larger than required. Diskpart can be used for extending the CSV size or an easier approach is to create a new CSV and migrate the required VMs to the new CSV. For optimal performance, the best practice is to increase the number of CSVs rather than increase their size as an interim measure.

Migration

One of the most common use cases in the current market condition is migration. Customers can use VMM Fabric or other third-party migration tools to migrate VMs. These tools use host level copy to move data form

the source platform to the destination platform, which can be time consuming depending on number of virtual machines that are in scope of migration.

Using ONTAP in such scenario's enable quicker migration than using host based migrationprocess. ONTAP also enables swift migration of VMs from one hypervisor to another (ESXi in this case to Hyper-V). VMDK of any size can be converted to VHDx in seconds on NetApp Storage. That is our PowerShell way - It leverages NetApp FlexClone® technology for the rapid conversion of VM hard disks. It also handles the creation and configuration of target and destination VMs.

This process helps minimize downtime and enhances business productivity. It also offers choice and flexibility by reducing licensing costs, lock-in, and commitments to a single vendor. This is also beneficial for organizations looking to optimize VM licensing costs and extend IT budgets.

The following video demonstates the process to migrate virtual machines from VMware ESX to Hyper-V.

Zero touch migration from ESX to Hyper-V

For additional information about migration using Flexclone and PowerShell, see the PowerShell script for migration.

Deploying Microsoft Hyper-V on NetApp Storage: Data Protection

Data protection is a key tenant for any production workload. This section describes how to backup and restore Hyper-V virtual machines.

Restore using NetApp Storage snapshot

Backing up VMs and quickly recovering or cloning them are among the great strengths of ONTAP volumes. Use Snapshot copies to make quick FlexClone copies of the VMs or even the whole CSV volume without affecting performance. This enables working with production data without the risk of data corruption when cloning production data volumes and mounting them on QA, staging and development environments. FlexClone volumes are useful for making test copies of production data, without having to double the amount of space required to copy the data.

Keep in mind, Hyper-V nodes assign each disk a unique ID and taking a snapshot of the volume that has respective partition (MBR or GPT) will carry the same unique identification. MBR uses disk signatures and GPT uses GUIDs (Global Unique Identifiers). In case of standalone Hyper-V host, the FlexClone volume can be easily mounted without any conflicts. This is because stand-alone Hyper-V servers can automatically detect duplicate disk IDs and change them dynamically without user intervention. This approach can be used to recover the VM(s) by copying the VHDs as the scenario demands.

While it is straightforward with standalone Hyper-V hosts, the procedure is different for Hyper-V clusters. The recovery process involves mapping the FlexClone volume to a standalone Hyper-V host or using diskpart to manually change the signature by mapping FlexClone volume to a standalone Hyper-V host (it is important because a disk ID conflict results in inability to bring the disk online) and once done, map the FlexClone volume to the cluster.

Backup and Restore using Third party solution

Note: This section uses Commvault, however this is applicable to other third-party solutions.

Leveraging ONTAP snapshots, CommVault IntelliSnap® creates hardware-based snapshots of Hyper-V. Backups can be automated based on the configuration for a Hyper-V hypervisor or VM group, or manually for a VM group or a specific VM. IntelliSnap enables fast protection of Hyper-V environments placing minimal load on the production Virtualization Farm. The integration of IntelliSnap technology with the Virtual

Server Agent (VSA) enables NetApp ONTAP Array to complete backups with a large number of virtual machines and data stores in a matter of minutes. Granular access provides individual file and folder recovery from the secondary tier of storage along with the full guest .vhd files.

Prior to configuring the virtualization environment, deploy the proper agents requiring snapshot integration with the Array. Microsoft Hyper-V virtualization environments require the following agents:

- MediaAgent
- Virtual Server Agent (VSA)
- VSS Hardware Provider (Windows Server 2012 and newer operating systems)

Configure NetApp Array using Array Management

The following steps show how to configure IntelliSnap virtual machine backups in an environment utilizing an ONTAP array and Hyper-V.

1. On the ribbon in the CommCell Console, click the Storage tab, and then click Array Management.
2. The Array Management dialog box appears.
3. Click Add.

   The Array Properties dialog box appears.

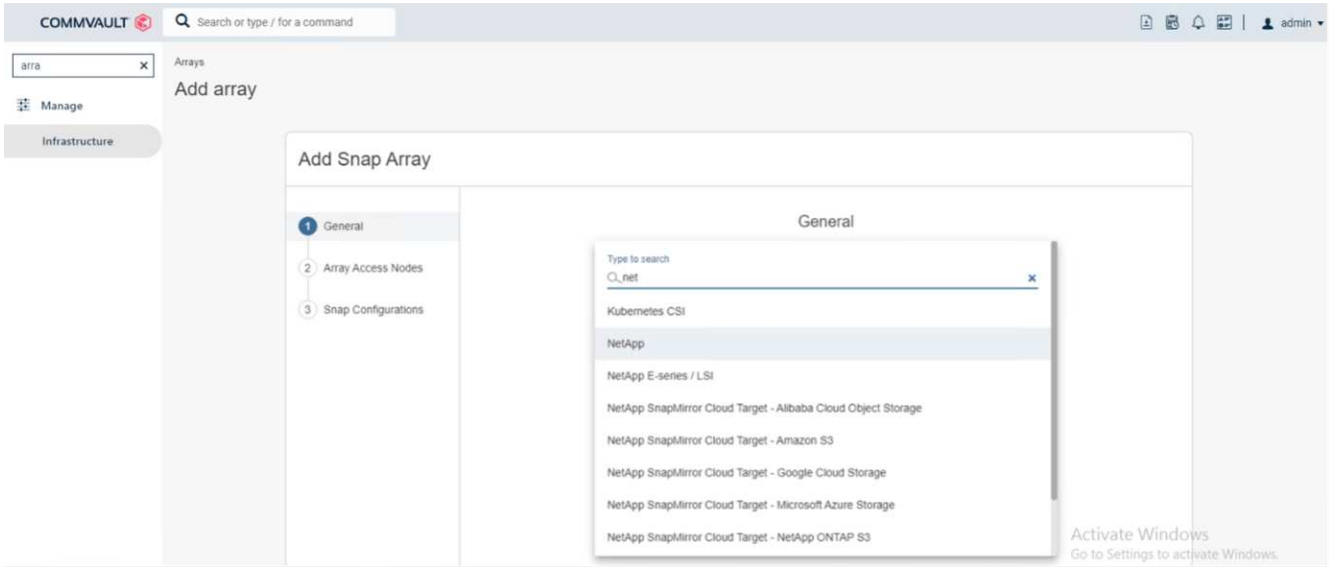

4. On the General tab, specify the following information:
5. From the Snap Vendor list, select NetApp.
6. In the Name box, enter the host name, the fully qualified domain name (FQDN), or the TCP/IP address of the primary file server.
7. On the Array Access Nodes tab, select available media agents.
8. On the Snap Configuration tab, configure Snapshot Configuration Properties according to your needs.
9. Click OK.
10. <Mandatory step> Once done, also configure SVM on the NetApp storage array by using the detect option to automatically detect storage virtual machines (SVM), then choose an SVM, and with the add option, add the SVM in the CommServe database, as an array management entry.

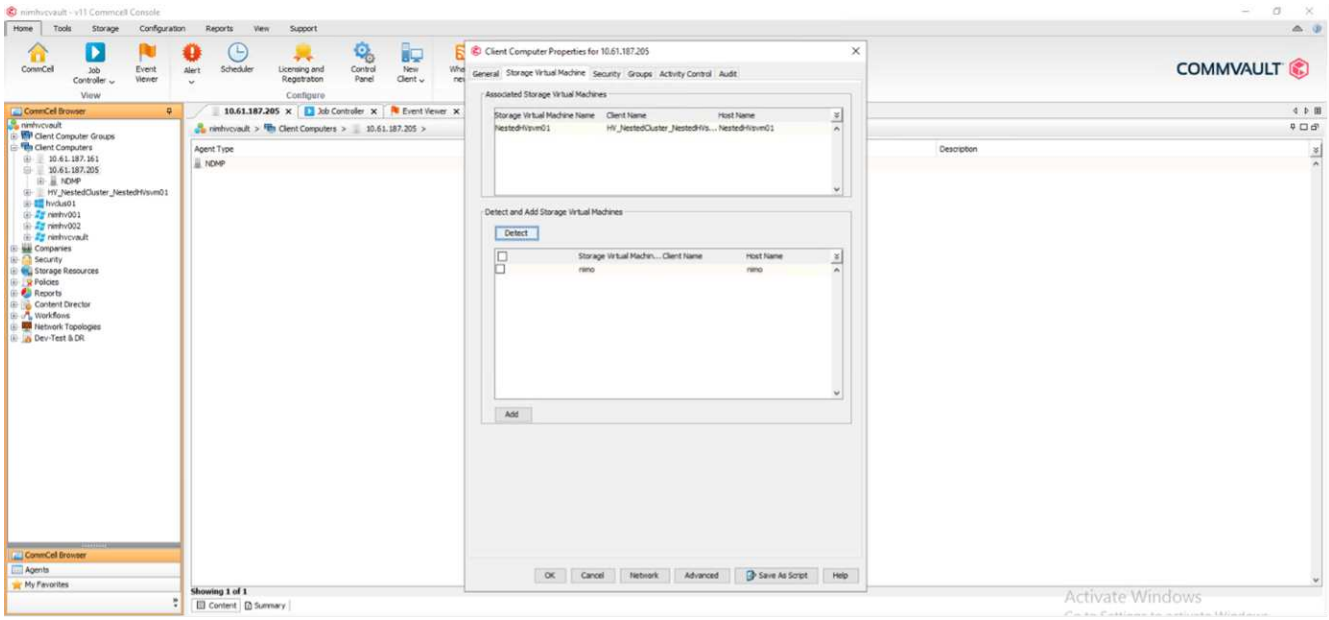

11. Click on Advanced (as shown in the below graphics) and select “Enable IntelliSnap” checkbox.

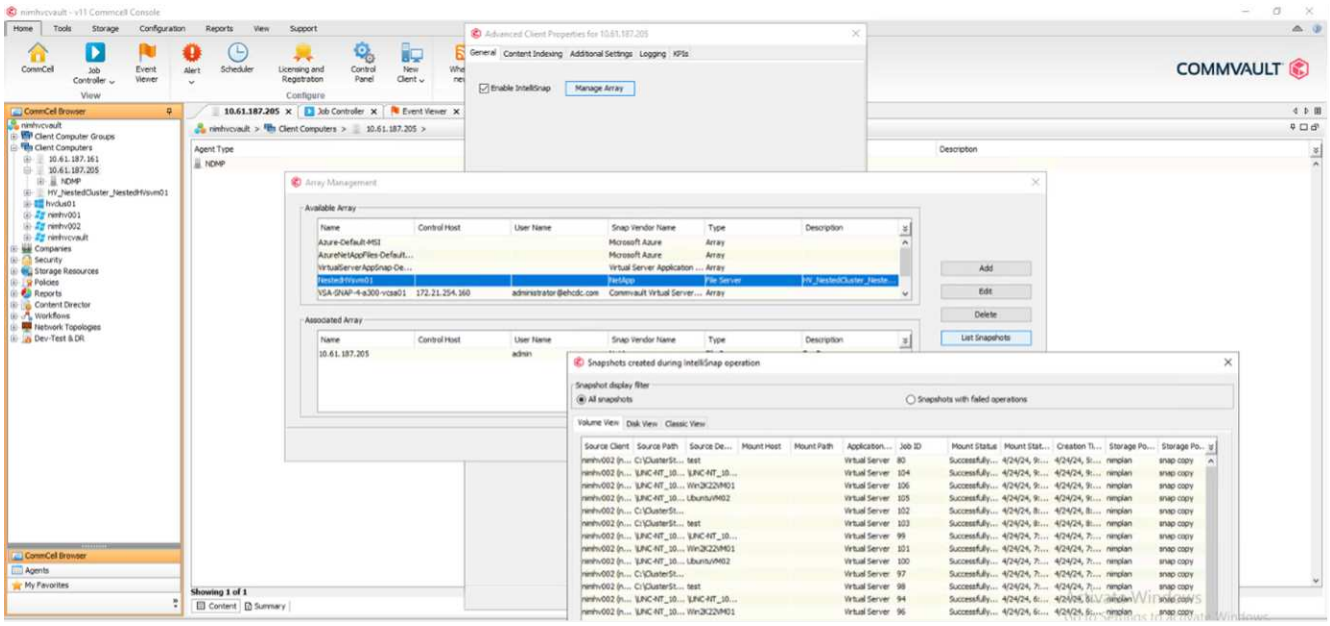

For detailed steps about configuring the array, see Configuring NetApp Array and Configuring Storage Virtual machines on NetApp Arrays

Add Hyper-V as the Hypervisor

Next step is to add Hyper-V hypervisor and adding a VM group.

Pre-requisites

- The hypervisor can be a Hyper-V cluster, a Hyper-V server in a cluster, or a standalone Hyper-V server.
- The user must belong to the Hyper-V administrators' group for Hyper-V Server 2012 and later. For a Hyper-V cluster, the user account must have full cluster permissions (Read and Full Control).
- Identify one or more nodes on which you will install the Virtual Server Agent (VSA) to create access nodes (VSA proxies) for backup and restore operations. To discover Hyper-V servers, the CommServe system

must have the VSA installed.

- To use Changed Block Tracking for Hyper-V 2012 R2, select all nodes in the Hyper-V cluster.

The following steps show how to add Hyper-V as a hypervisor.

1. After the core setup is complete, on the Protect tab, click the Virtualization tile.
2. On the Create server backup plan page, type a name for the plan, then provide information about storage, retention, and backup schedules.
3. Now the Add hypervisor page appears > Select vendor: Select Hyper-V (Enter the IP address or FQDN and user credentials)
4. For a Hyper-V server, click Discover nodes. When the Nodes field is populated, select one or more nodes on which to install the Virtual Server Agent.

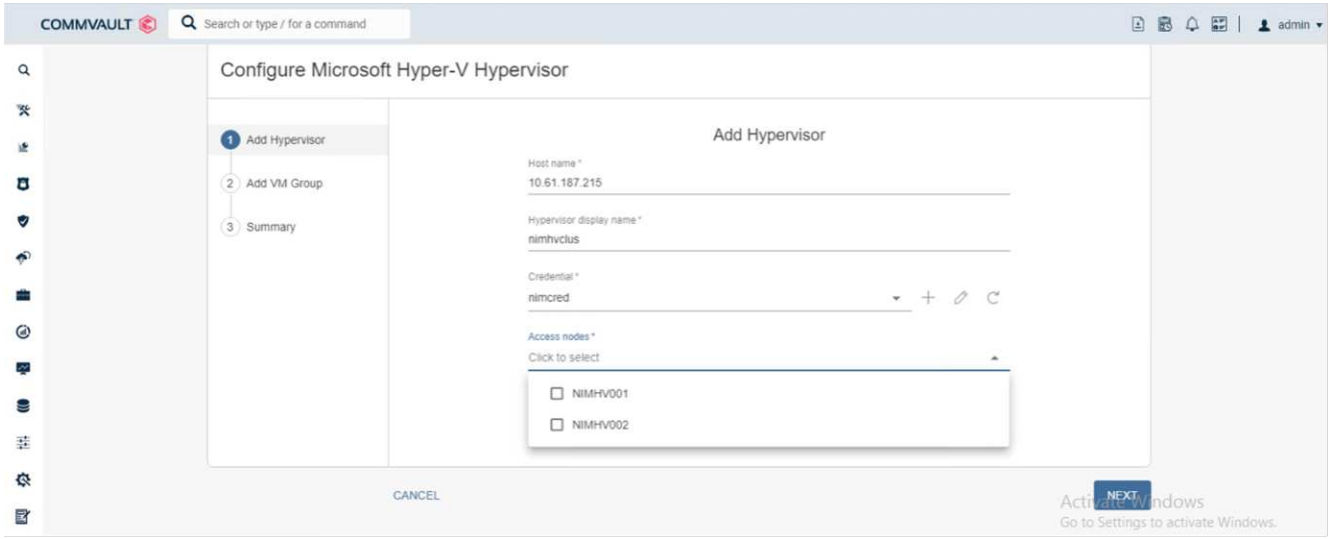

5. Click Next and the Save.

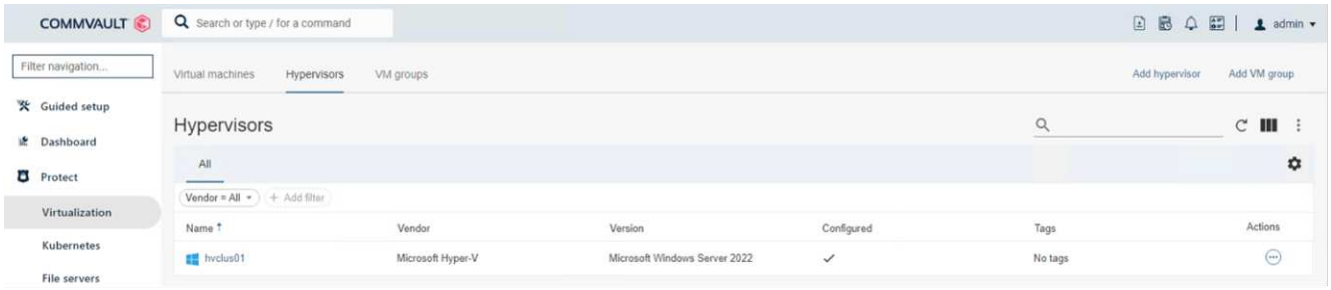

6. On the Add VM group page, select the virtual machines to be protected (Demogrp is the VM group created in this case) and enable IntelliSnap option as shown below.

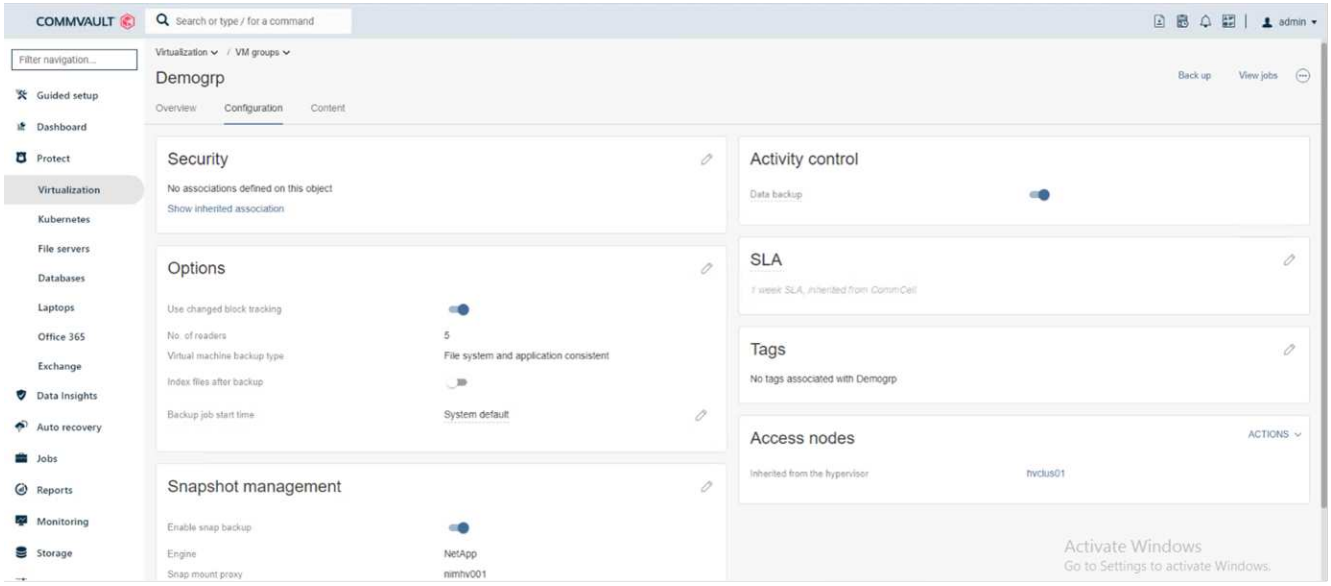

Note: When IntelliSnap is enabled on a VM group, Commvault automatically creates schedule policies for the primary (snap) and backup copies.

7. Click Save.

For detailed steps about configuring the array, see Adding a Hypervisor.

Performing a backup:

1. From the navigation pane, go to Protect > Virtualization. The Virtual machines page appears.
2. Back up the VM or the VM group. In this demo, VM group is selected. In the row for the VM group, click the action button action_button, and then select Back up. In this case, nimplan is the plan associated against Demogrp and Demogrp01.

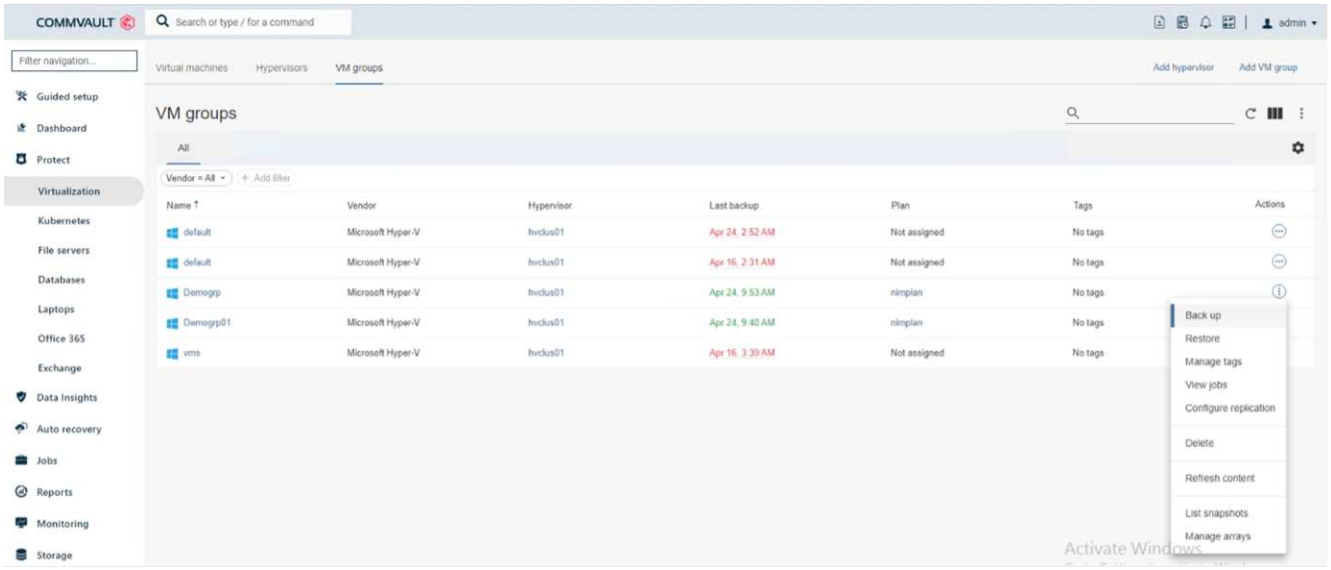

3. Once the backup is successful, restore points are available as shown in the screen capture. From the snap copy, restore of full VM and restore of guest files and folders can be performed.

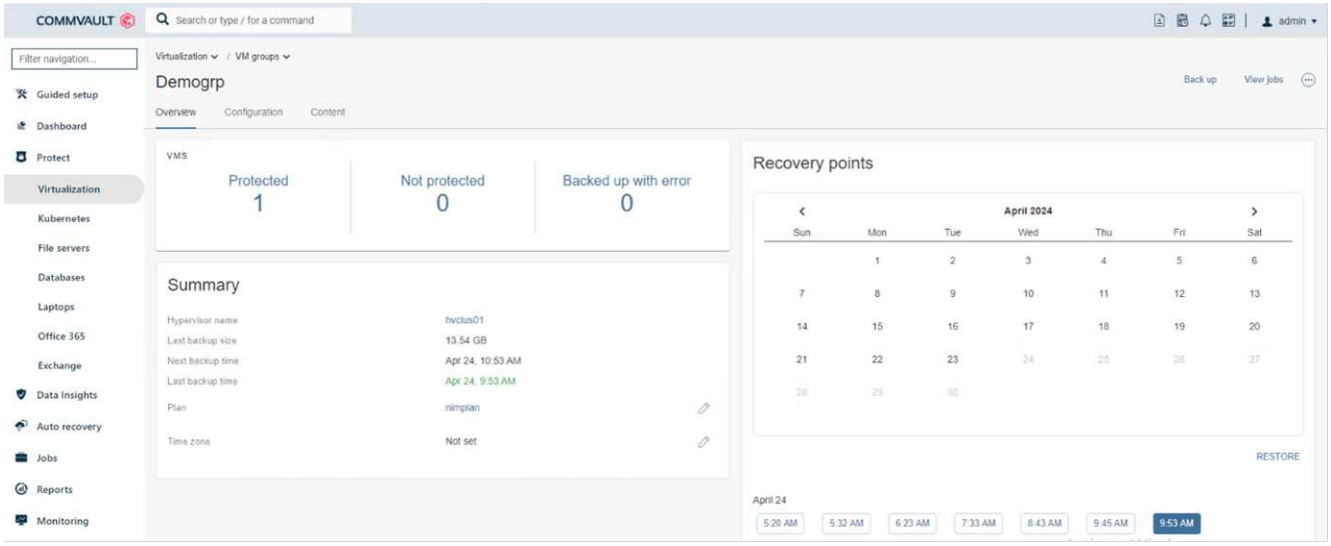

Note: For critical and heavily utilized virtual machines, keep fewer virtual machines per CSV

Performing a restore operation:

Restore full VMs, guest files and folders, or virtual disk files via the restore points.

1. From the navigation pane, go to Protect > Virtualization, the Virtual machines page appears.
2. Click the VM groups tab.
3. The VM group page appears.
4. In the VM groups area, click Restore for the VM group that contains the virtual machine.
5. The Select restore type page appears.

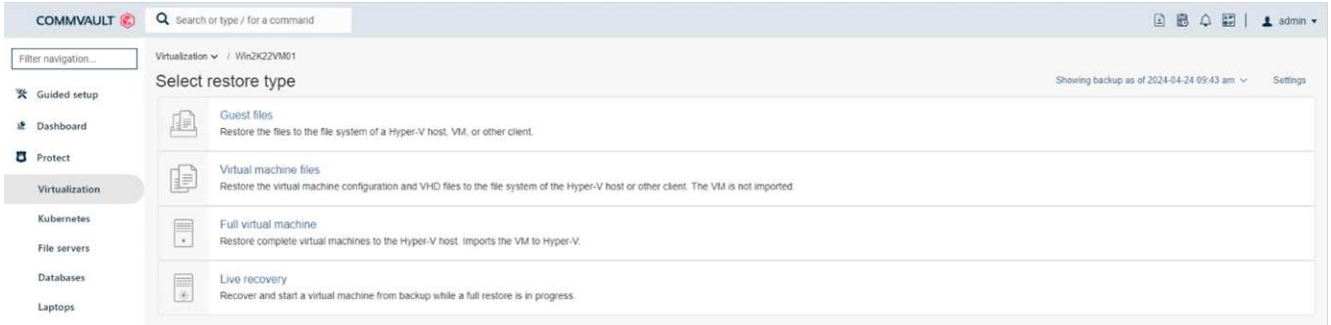

6. Select Guest files or Full virtual machine depending on the selection and trigger the restore.

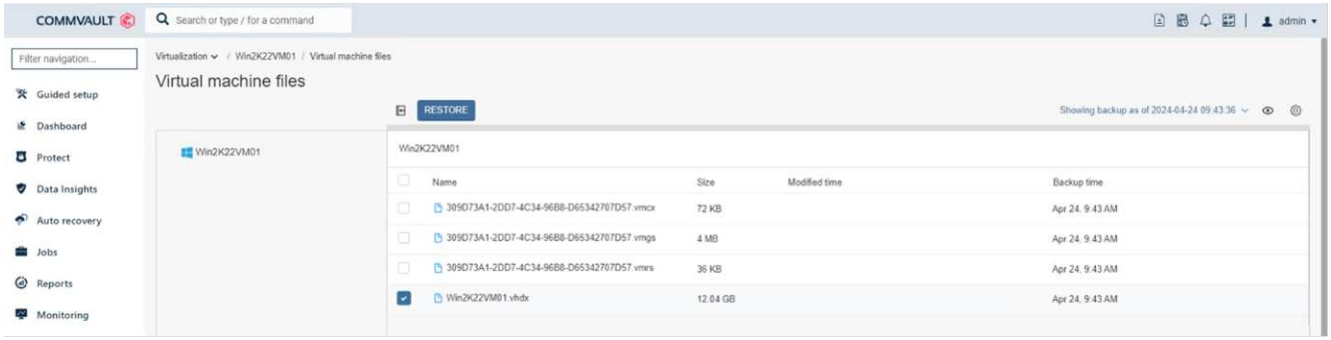

For detailed steps for all supported restore options, see Restores for Hyper-V.

Advanced NetApp ONTAP options

NetApp SnapMirror enables efficient site-to-site storage replication, making disaster recovery rapid, reliable, and manageable to suit today's global enterprises. Replicating data at high speeds over LANs and WANs, SnapMirror provides high data availability and fast recovery for mission-critical applications, as well as outstanding storage deduplication and network compression capabilities. With NetApp SnapMirror technology, disaster recovery can protect the entire data center. Volumes can back up to an off-site location incrementally. SnapMirror performs incremental, block-based replication as frequently as the required RPO. The block-level updates reduce bandwidth and time requirements, and data consistency is maintained at the DR site.

An important step is to create a one-time baseline transfer of the entire dataset. This is required before incremental updates can be performed. This operation includes the creation of a Snapshot copy at the source and the transfer of all the data blocks referenced by it to the destination file system. After the initialization is complete, scheduled or manually triggered updates can occur. Each update transfers only the new and changed blocks from the source to the destination file system. This operation includes creating a Snapshot copy at the source volume, comparing it with the baseline copy, and transferring only the changed blocks to the destination volume. The new copy becomes the baseline copy for the next update. Because the replication is periodic, SnapMirror can consolidate the changed blocks and conserve network bandwidth. The impact on write throughput and write latency is minimal.

Recovery is performed by completing the following steps:

1. Connect to the storage system on the secondary site.
2. Break the SnapMirror relationship.
3. Map the LUNs in the SnapMirror volume to the initiator group (igroup) for the Hyper-V servers on the secondary site.
4. Once the LUNs are mapped to the Hyper-V cluster, make these disks online.
5. Using the failover-cluster PowerShell cmdlets, add the disks to available storage and convert them to CSVs.
6. Import the virtual machines in the CSV to the Hyper-V manager, make them highly available, and then add them to the cluster.
7. Turn on the VMs.

Deploying Microsoft Hyper-V on NetApp Storage: Conclusion

ONTAP is the optimal shared storage foundation to deploy a variety of IT workloads. ONTAP AFF or ASA platforms are both flexible and scalable for multiple use cases and applications. Windows Server 2022 and Hyper-V enabled on it is one common use case as the virtualization solution, which is described in this document. The flexibility and scalability of ONTAP storage and associated features enable customers to start out with a right-sized storage layer that can grow with and adapt to their evolving business requirements. In current market conditions, Hyper-V offers a perfect alternate hypervisor option which provides most of the functionalities that was provided VMware.

Deploying Microsoft Hyper-V on NetApp Storage: Migration Script

This section contains a PowerShell script that can be used for migration using Flexclone.

Powershell script

```
param (
    [Parameter(Mandatory=$True, HelpMessage="VCenter DNS name or IP
Address")]
    [String]$VCENTER,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP NFS Datastore
name")]
    [String]$DATASTORE,
    [Parameter(Mandatory=$True, HelpMessage="VCenter credentials")]
    [System.Management.Automation.PSCredential]$VCENTER_CREDS,
    [Parameter(Mandatory=$True, HelpMessage="The IP Address of the ONTAP
Cluster")]
    [String]$ONTAP_CLUSTER,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP VServer/SVM
name")]
    [String]$VSERVER,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP NSF,SMB Volume
name")]
    [String]$ONTAP_VOLUME_NAME,
    [Parameter(Mandatory=$True, HelpMessage="ONTAP NFS/CIFS Volume mount
Drive on Hyper-V host")]
    [String]$ONTAP_NETWORK_SHARE_ADDRESS,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP Volume QTree
folder name")]
    [String]$VHDX_QTREE_NAME,
    [Parameter(Mandatory=$True, HelpMessage="The Credential to connect to
the ONTAP Cluster")]
    [System.Management.Automation.PSCredential]$ONTAP_CREDS,
    [Parameter(Mandatory=$True, HelpMessage="Hyper-V VM switch name")]
    [String]$HYPERV_VM_SWITCH
)

function main {

    ConnectVCenter

    ConnectONTAP

    GetVMList

    GetVMInfo

    #PowerOffVMs

    CreateOntapVolumeSnapshot
```

```
    Shift

    ConfigureVMsOnHyperV
}

function ConnectVCenter {
    Write-Host
"-------------------------------------------------------------------------
-----" -ForegroundColor Cyan
    Write-Host "Connecting to vCenter $VCENTER" -ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan

    [string]$vmwareModuleName = "VMware.VimAutomation.Core"

    Write-Host "Importing VMware $vmwareModuleName Powershell module"
    if ((Get-Module|Select-Object -ExpandProperty Name) -notcontains
$vmwareModuleName) {
        Try {
            Import-Module $vmwareModuleName -ErrorAction Stop
            Write-Host "$vmwareModuleName imported successfully"
-ForegroundColor Green
        } Catch {
            Write-Error "Error: $vmwareMdouleName PowerShell module not
found"
            break;
        }
    }
    else {
        Write-Host "$vmwareModuleName Powershell module already imported"
-ForegroundColor Green
    }

    Write-Host "`nConnecting to vCenter $VCENTER"
    Try {
        $connect = Connect-VIServer -Server $VCENTER -Protocol https
-Credential $VCENTER_CREDS -ErrorAction Stop
        Write-Host "Connected to vCenter $VCENTER" -ForegroundColor Green
    } Catch {
        Write-Error "Failed to connect to vCenter $VCENTER. Error : $($_
.Exception.Message)"
        break;
    }
}
```

```
function ConnectONTAP {
    Write-Host "`n
------------------------------------------------------------------------
----" -ForegroundColor Cyan
    Write-Host "Connecting to VSerevr $VSERVER at ONTAP Cluster
$ONTAP_CLUSTER" -ForegroundColor Magenta
    Write-Host
"------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan

    [string]$ontapModuleName = "NetApp.ONTAP"

    Write-Host "Importing NetApp ONTAP $ontapModuleName Powershell module"
    if ((Get-Module|Select-Object -ExpandProperty Name) -notcontains
$ontapModuleName) {
        Try {
            Import-Module $ontapModuleName -ErrorAction Stop
            Write-Host "$ontapModuleName imported successfully"
-ForegroundColor Green
        } Catch {
            Write-Error "Error: $vmwareMdouleName PowerShell module not
found"
            break;
        }
    }
    else {
        Write-Host "$ontapModuleName Powershell module already imported"
-ForegroundColor Green
    }

    Write-Host "`nConnecting to ONTAP Cluster $ONTAP_CLUSTER"
    Try {
        $connect = Connect-NcController -Name $ONTAP_CLUSTER -Credential
$ONTAP_CREDS -Vserver $VSERVER
        Write-Host "Connected to ONTAP Cluster $ONTAP_CLUSTER"
-ForegroundColor Green
    } Catch {
        Write-Error "Failed to connect to ONTAP Cluster $ONTAP_CLUSTER.
Error : $($_.Exception.Message)"
        break;
    }
}

function GetVMList {
    Write-Host "`n
------------------------------------------------------------------------
```

```
----" -ForegroundColor Cyan
    Write-Host "Fetching powered on VMs list with Datastore $DATASTORE"
-ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan
    try {
        $vmList = VMware.VimAutomation.Core\Get-VM -Datastore $DATASTORE
-ErrorAction Stop| Where-Object {$_.PowerState -eq "PoweredOn"} | OUT-
GridView -OutputMode Multiple
        #$vmList = Get-VM -Datastore $DATASTORE -ErrorAction Stop| Where-
Object {$_.PowerState -eq "PoweredOn"}

        if($vmList) {
            Write-Host "Selected VMs for Shift" -ForegroundColor Green
            $vmList | Format-Table -Property Name
            $Script:VMList = $vmList
        }
        else {
            Throw "No VMs selected"
        }
    }
    catch {
        Write-Error "Failed to get VM List. Error : $($_.Exception.
Message)"
        Break;
    }
}

function GetVMInfo {
    Write-Host
"-------------------------------------------------------------------------
-----" -ForegroundColor Cyan
    Write-Host "VM Information" -ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----" -ForegroundColor Cyan
    $vmObjArray = New-Object System.Collections.ArrayList

    if($VMList) {
        foreach($vm in $VMList) {
            $vmObj = New-Object -TypeName System.Object

            $vmObj | Add-Member -MemberType NoteProperty -Name ID -Value
$vm.Id
            $vmObj | Add-Member -MemberType NoteProperty -Name Name -Value
```

```
$vm.Name
            $vmObj | Add-Member -MemberType NoteProperty -Name NumCpu
-Value $vm.NumCpu
            $vmObj | Add-Member -MemberType NoteProperty -Name MemoryGB
-Value $vm.MemoryGB
            $vmObj | Add-Member -MemberType NoteProperty -Name Firmware
-Value $vm.ExtensionData.Config.Firmware

            $vmDiskInfo = $vm | VMware.VimAutomation.Core\Get-HardDisk

            $vmDiskArray = New-Object System.Collections.ArrayList
            foreach($disk in $vmDiskInfo) {
                $diskObj = New-Object -TypeName System.Object

                $diskObj | Add-Member -MemberType NoteProperty -Name Name
-Value $disk.Name

                $fileName = $disk.Filename
                if ($fileName -match '\[(.*?)\]') {
                    $dataStoreName = $Matches[1]
                }

                $parts = $fileName -split " "
                $pathParts = $parts[1] -split "/"
                $folderName = $pathParts[0]
                $fileName = $pathParts[1]

                $diskObj | Add-Member -MemberType NoteProperty -Name
DataStore -Value $dataStoreName
                $diskObj | Add-Member -MemberType NoteProperty -Name
Folder -Value $folderName
                $diskObj | Add-Member -MemberType NoteProperty -Name
Filename -Value $fileName
                $diskObj | Add-Member -MemberType NoteProperty -Name
CapacityGB -Value $disk.CapacityGB

                $null = $vmDiskArray.Add($diskObj)
            }

            $vmObj | Add-Member -MemberType NoteProperty -Name
PrimaryHardDisk -Value "[$($vmDiskArray[0].DataStore)] $($vmDiskArray[0]
.Folder)/$($vmDiskArray[0].Filename)"
            $vmObj | Add-Member -MemberType NoteProperty -Name HardDisks
-Value $vmDiskArray

            $null = $vmObjArray.Add($vmObj)
```

```
            $vmNetworkArray = New-Object System.Collections.ArrayList

            $vm |
            ForEach-Object {
              $VM = $_
              $VM | VMware.VimAutomation.Core\Get-VMGuest | Select-Object
-ExpandProperty Nics |
              ForEach-Object {
                $Nic = $_
                foreach ($IP in $Nic.IPAddress)
                {
                  if ($IP.Contains('.'))
                  {
                    $networkObj = New-Object -TypeName System.Object

                    $vlanId = VMware.VimAutomation.Core\Get-
VirtualPortGroup | Where-Object {$_.Key -eq $Nic.NetworkName}
                    $networkObj | Add-Member -MemberType NoteProperty
-Name VLanID -Value $vlanId
                    $networkObj | Add-Member -MemberType NoteProperty
-Name IPv4Address -Value $IP

                    $null = $vmNetworkArray.Add($networkObj)
                  }
                }
              }
            }

            $vmObj | Add-Member -MemberType NoteProperty -Name PrimaryIPv4
-Value $vmNetworkArray[0].IPv4Address
            $vmObj | Add-Member -MemberType NoteProperty -Name
PrimaryVLanID -Value $vmNetworkArray.VLanID
            $vmObj | Add-Member -MemberType NoteProperty -Name Networks
-Value $vmNetworkArray

            $guest = $vm.Guest
            $parts = $guest -split ":"
            $afterColon = $parts[1]

            $osFullName = $afterColon

            $vmObj | Add-Member -MemberType NoteProperty -Name OSFullName
-Value $osFullName
            $vmObj | Add-Member -MemberType NoteProperty -Name GuestID
-Value $vm.GuestId
        }
```

```
    }

    $vmObjArray | Format-Table -Property ID, Name, NumCpu, MemoryGB,
PrimaryHardDisk, PrimaryIPv4, PrimaryVLanID, GuestID, OSFullName, Firmware

    $Script:VMObjList = $vmObjArray
}

function PowerOffVMs {
    Write-Host "`n
-------------------------------------------------------------------------
----" -ForegroundColor Cyan
    Write-Host "Power Off VMs" -ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan
    foreach($vm in $VMObjList) {
        try {
            Write-Host "Powering Off VM $($vm.Name) in vCenter $($VCENTER
)"
            $null = VMware.VimAutomation.Core\Stop-VM -VM $vm.Name
-Confirm:$false -ErrorAction Stop
            Write-Host "Powered Off VM $($vm.Name)" -ForegroundColor Green
        }
        catch {
            Write-Error "Failed to Power Off VM $($vm.Name). Error :
$._Exception.Message"
            Break;
        }
        Write-Host "`n"
    }
}

function CreateOntapVolumeSnapshot {
    Write-Host "`n
-------------------------------------------------------------------------
----" -ForegroundColor Cyan
    Write-Host "Taking ONTAP Snapshot for Volume $ONTAP_VOLUME_NAME"
-ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan

    Try {
        Write-Host "Taking snapshot for Volume $ONTAP_VOLUME_NAME"
        $timestamp = Get-Date -Format "yyyy-MM-dd_HHmmss"
```

```
        $snapshot = New-NcSnapshot -VserverContext $VSERVER -Volume
$ONTAP_VOLUME_NAME -Snapshot "snap.script-$timestamp"

        if($snapshot) {
            Write-Host "Snapshot ""$($snapshot.Name)"" created for Volume
$ONTAP_VOLUME_NAME" -ForegroundColor Green
            $Script:OntapVolumeSnapshot = $snapshot
        }
    } Catch {
        Write-Error "Failed to create snapshot for Volume
$ONTAP_VOLUME_NAME. Error : $_.Exception.Message"
        Break;
    }
}

function Shift {
    Write-Host
"-------------------------------------------------------------------------
-----" -ForegroundColor Cyan
    Write-Host "VM Shift" -ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan

    $Script:HypervVMList = New-Object System.Collections.ArrayList
    foreach($vmObj in $VMObjList) {

        Write-Host "***********************************************"
        Write-Host "Performing VM conversion for $($vmObj.Name)"
-ForegroundColor Blue
        Write-Host "***********************************************"

        $hypervVMObj = New-Object -TypeName System.Object

        $directoryName = "/vol/$($ONTAP_VOLUME_NAME)/$($VHDX_QTREE_NAME)
/$($vmObj.HardDisks[0].Folder)"

        try {
            Write-Host "Creating Folder ""$directoryName"" for VM $(
$vmObj.Name)"
            $dir = New-NcDirectory -VserverContext $VSERVER -Path
$directoryName -Permission 0777 -Type directory -ErrorAction Stop
            if($dir) {
                Write-Host "Created folder ""$directoryName"" for VM
$($vmObj.Name)`n" -ForegroundColor Green
            }
```

```
        }
        catch {
            if($_.Exception.Message -eq "[500]: File exists") {
                Write-Warning "Folder ""$directoryName"" already exists!
`n"
            }
            Else {
                Write-Error "Failed to create folder ""$directoryName""
for VM $($vmObj.Name). Error : $($_.Exception.Message)"
                Break;
            }
        }

        $vmDiskArray = New-Object System.Collections.ArrayList

        foreach($disk in $vmObj.HardDisks) {
            $vmDiskObj = New-Object -TypeName System.Object
            try {
                Write-Host "`nConverting $($disk.Name)"
                Write-Host "--------------------------------"

                $vmdkPath = "/vol/$($ONTAP_VOLUME_NAME)/$($disk.Folder)/
$($disk.Filename)"
                $fileName = $disk.Filename -replace '\.vmdk$', ''
                $vhdxPath = "$($directoryName)/$($fileName).vhdx"

                Write-Host "Converting ""$($disk.Name)"" VMDK path ""
$($vmdkPath)"" to VHDX at Path ""$($vhdxPath)"" for VM $($vmObj.Name)"
                $convert = ConvertTo-NcVhdx -SourceVmdk $vmdkPath
-DestinationVhdx $vhdxPath  -SnapshotName $OntapVolumeSnapshot
-ErrorAction Stop -WarningAction SilentlyContinue
                if($convert) {
                    Write-Host "Successfully converted VM ""$($vmObj.Name
)"" VMDK path ""$($vmdkPath)"" to VHDX at Path ""$($vhdxPath)"""
-ForegroundColor Green

                    $vmDiskObj | Add-Member -MemberType NoteProperty -Name
Name -Value $disk.Name
                    $vmDiskObj | Add-Member -MemberType NoteProperty -Name
VHDXPath -Value $vhdxPath

                    $null = $vmDiskArray.Add($vmDiskObj)
                }
            }
            catch {
                Write-Error "Failed to convert ""$($disk.Name)"" VMDK to
```

```
VHDX for VM $($vmObj.Name). Error : $($_.Exception.Message)"
                Break;
            }
        }

        $hypervVMObj | Add-Member -MemberType NoteProperty -Name Name
-Value $vmObj.Name
        $hypervVMObj | Add-Member -MemberType NoteProperty -Name HardDisks
-Value $vmDiskArray
        $hypervVMObj | Add-Member -MemberType NoteProperty -Name MemoryGB
-Value $vmObj.MemoryGB
        $hypervVMObj | Add-Member -MemberType NoteProperty -Name Firmware
-Value $vmObj.Firmware
        $hypervVMObj | Add-Member -MemberType NoteProperty -Name GuestID
-Value $vmObj.GuestID



        $null = $HypervVMList.Add($hypervVMObj)
        Write-Host "`n"

    }
}

function ConfigureVMsOnHyperV {
    Write-Host
"-------------------------------------------------------------------------
-----" -ForegroundColor Cyan
    Write-Host "Configuring VMs on Hyper-V" -ForegroundColor Magenta
    Write-Host
"-------------------------------------------------------------------------
-----`n" -ForegroundColor Cyan

    foreach($vm in $HypervVMList) {
        try {

            # Define the original path
            $originalPath = $vm.HardDisks[0].VHDXPath
            # Replace forward slashes with backslashes
            $windowsPath = $originalPath -replace "/", "\"

            # Replace the initial part of the path with the Windows drive
letter
            $windowsPath = $windowsPath -replace "^\\vol\\", "\\
$($ONTAP_NETWORK_SHARE_ADDRESS)\"

            $vmGeneration = if ($vm.Firmware -eq "bios") {1} else {2};
```

```
            Write-Host "************************************************"
            Write-Host "Creating VM $($vm.Name)" -ForegroundColor Blue
            Write-Host "************************************************"
            Write-Host "Creating VM $($vm.Name) with Memory $($vm.
MemoryGB)GB, vSwitch $($HYPERV_VM_SWITCH), $($vm.HardDisks[0].Name) ""
$($windowsPath)"", Generation $($vmGeneration) on Hyper-V"

            $createVM = Hyper-V\New-VM -Name $vm.Name -VHDPath
$windowsPath -SwitchName $HYPERV_VM_SWITCH -MemoryStartupBytes (Invoke-
Expression "$($vm.MemoryGB)GB") -Generation $vmGeneration -ErrorAction
Stop
            if($createVM) {
                Write-Host "VM $($createVM.Name) created on Hyper-V host
`n" -ForegroundColor Green


                $index = 0
                foreach($vmDisk in $vm.HardDisks) {
                    $index++
                    if ($index -eq 1) {
                        continue
                    }

                    Write-Host "`nAttaching $($vmDisk.Name) for VM $($vm
.Name)"
                    Write-Host
"---------------------------------------------"

                    $originalPath = $vmDisk.VHDXPath

                    # Replace forward slashes with backslashes
                    $windowsPath = $originalPath -replace "/", "\"

                    # Replace the initial part of the path with the
Windows drive letter
                    $windowsPath = $windowsPath -replace "^\\vol\\", "\\
$($ONTAP_NETWORK_SHARE_ADDRESS)\"

                    try {
                        $attachDisk = Hyper-v\Add-VMHardDiskDrive -VMName
$vm.Name -Path $windowsPath -ErrorAction Stop
                        Write-Host "Attached $($vmDisk.Name) ""
$($windowsPath)"" to VM $($vm.Name)" -ForegroundColor Green
                    }
                    catch {
```

```
                        Write-Error "Failed to attach $($vmDisk.Name)
$($windowsPath) to VM $($vm.Name): Error : $($_.Exception.Message)"
                        Break;
                    }
                }

                if($vmGeneration -eq 2 -and $vm.GuestID -like "*rhel*") {
                    try {
                        Write-Host "`nDisabling secure boot"
                        Hyper-V\Set-VMFirmware -VMName $createVM.Name
-EnableSecureBoot Off -ErrorAction Stop
                        Write-Host "Secure boot disabled" -ForegroundColor
Green
                    }
                    catch {
                        Write-Error "Failed to disable secure boot for VM
$($createVM.Name). Error : $($_.Exception.Message)"
                    }
                }

                try {
                    Write-Host "`nStarting VM $($createVM.Name)"
                    Hyper-v\Start-VM -Name $createVM.Name -ErrorAction
Stop
                    Write-Host "Started VM $($createVM.Name)`n"
-ForegroundColor Green
                }
                catch {
                    Write-Error "Failed to start VM $($createVM.Name).
Error : $($_.Exception.Message)"
                    Break;
                }
            }
        }
        catch {
            Write-Error "Failed  to create VM $($vm.Name) on Hyper-V.
Error : $($_.Exception.Message)"
            Break;
        }
    }
}

main
```

Additional Resources

NetApp OpenShift Virtualization Solutions

Overview

Red Hat OpenShift Virtualization with NetApp ONTAP

Depending on the specific use case, both containers and virtual machines (VMs) can serve as optimal platforms for different types of applications. Therefore, many organizations run some of their workloads on containers and some on VMs. Often, this leads organizations to face additional challenges by having to manage separate platforms: a hypervisor for VMs and a container orchestrator for applications.

To address this challenge, Red Hat introduced OpenShift Virtualization (formerly known as Container Native Virtualization) starting from OpenShift version 4.6. The OpenShift Virtualization feature enables you to run and manage virtual machines alongside containers on the same OpenShift Container Platform installation, providing hybrid management capability to automate deployment and management of VMs through operators. In addition to creating VMs in OpenShift, with OpenShift Virtualization, Red Hat also supports importing VMs from VMware vSphere, Red Hat Virtualization, and Red Hat OpenStack Platform deployments.

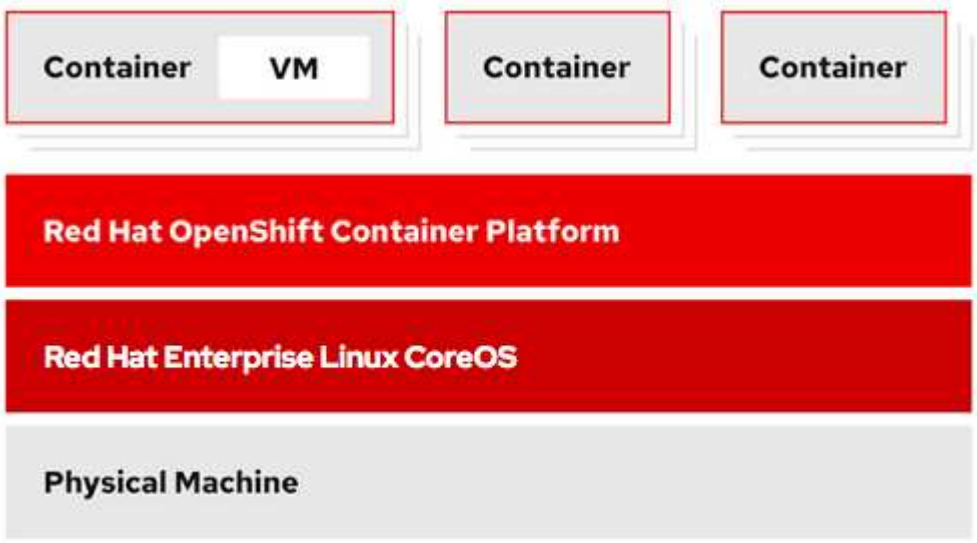

Certain features like live VM migration, VM disk cloning, VM snapshots and so on are also supported by OpenShift Virtualization with assistance from Astra Trident when backed by NetApp ONTAP. Examples of each of these workflows are discussed later in this document in their respective sections.

To learn more about Red Hat OpenShift Virtualization, see the documentation here.

NetApp Storage Overview

NetApp has several storage platforms that are qualified with our Astra Trident Storage Orchestrator to provision storage for applications deployed on Red Hat OpenShift.

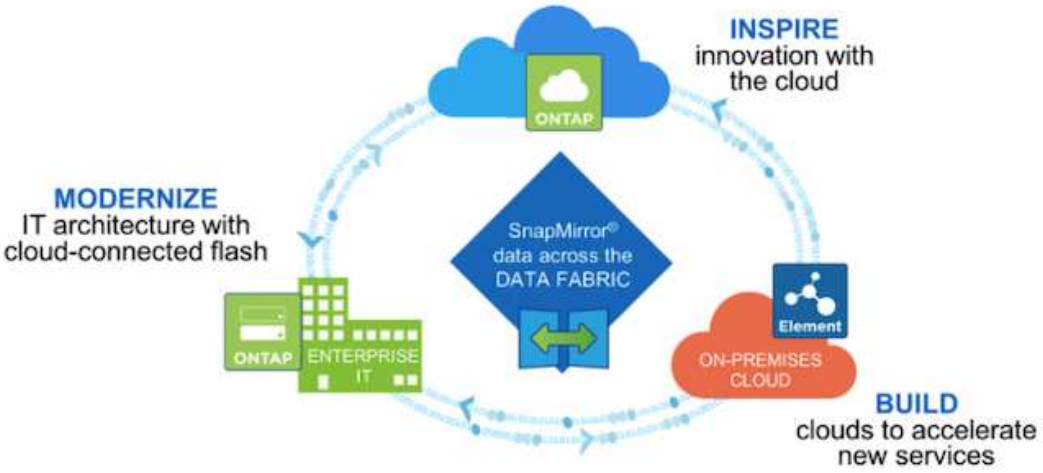

- AFF and FAS systems run NetApp ONTAP and provide storage for both file-based (NFS) and block-based (iSCSI) use cases.
- Cloud Volumes ONTAP and ONTAP Select provide the same benefits in the cloud and virtual space respectively.
- NetApp Cloud Volumes Service (AWS/GCP) and Azure NetApp Files provide file-based storage in the cloud.
- NetApp Element storage systems provide for block-based (iSCSI) use cases in a highly scalable environment.

Each storage system in the NetApp portfolio can ease both data management and movement between on-premises sites and the cloud, ensuring that your data is where your applications are.

The following pages have additional information about the NetApp storage systems validated in the Red Hat OpenShift with NetApp solution:

- NetApp ONTAP
- NetApp Element

NetApp Storage Integration Overview

NetApp provides a number of products to help you with orchestrating and managing persistent data in container based environments, such as Red Hat OpenShift.

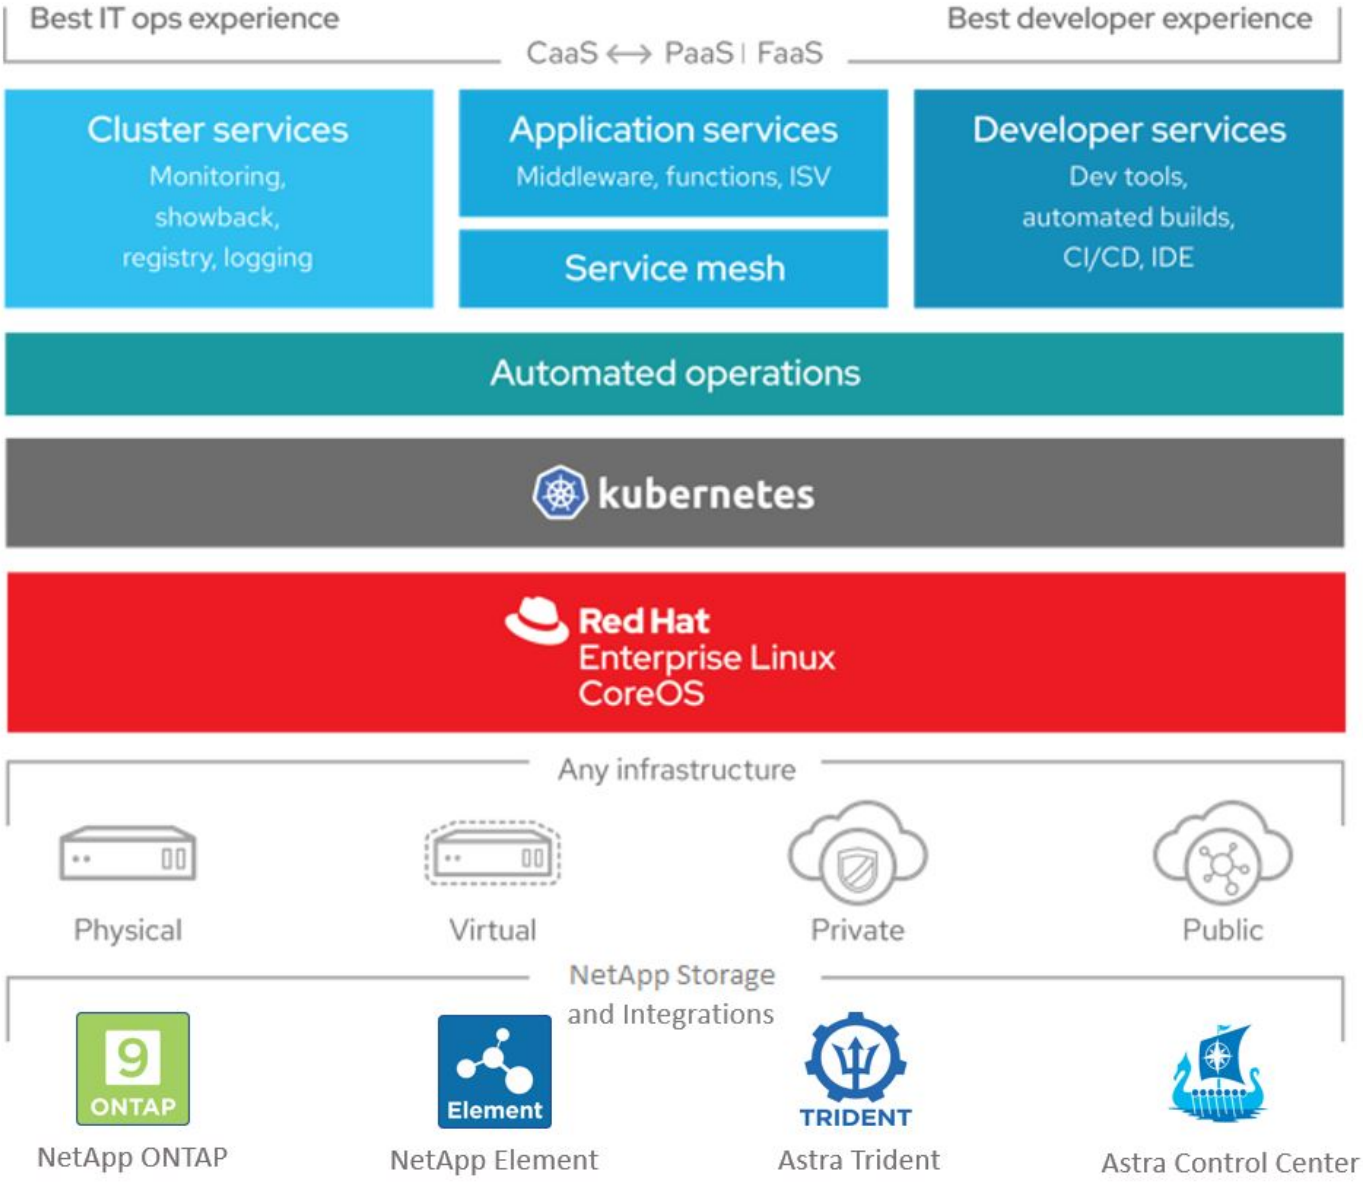

NetApp Astra Control offers a rich set of storage and application-aware data management services for stateful Kubernetes workloads, powered by NetApp data protection technology. The Astra Control Service is available to support stateful workloads in cloud-native Kubernetes deployments. The Astra Control Center is available to support stateful workloads in on-premises deployments, like Red Hat OpenShift. For more information visit the NetApp Astra Control website here.

NetApp Astra Trident is an open-source and fully-supported storage orchestrator for containers and Kubernetes distributions, including Red Hat OpenShift. For more information, visit the Astra Trident website here.

The following pages have additional information about the NetApp products that have been validated for application and persistent storage management in the Red Hat OpenShift with NetApp solution:

- NetApp Astra Control Center
- NetApp Astra Trident

Videos and Demos: Red Hat OpenShift with NetApp

The following videos demonstrate some of the capabilities documented in this document:

Cloud Insights integration with Openshift Virtualization

Using Red Hat MTV to migrate VMs to OpenShift Virtualization with NetApp ONTAP Storage

Accelerate Software Development with Astra Control and NetApp FlexClone Technology - Red Hat OpenShift with NetApp

Leverage NetApp Astra Control to Perform Post-mortem Analysis and Restore Your Application

Data Protection in CI/CD pipeline with Astra Control Center

Workload Migration using Astra Control Center - Red Hat OpenShift with NetApp

Workload Migration - Red Hat OpenShift with NetApp

Installing OpenShift Virtualization - Red Hat OpenShift with NetApp

Deploying a Virtual Machine with OpenShift Virtualization - Red Hat OpenShift with NetApp

NetApp HCI for Red Hat OpenShift on Red Hat Virtualization

Deployment

Deploy Red Hat OpenShift Virtualization with NetApp ONTAP

This section details how to deploy Red Hat OpenShift Virtualization with NetApp ONTAP.

Prerequisites

- A Red Hat OpenShift cluster (later than version 4.6) installed on bare-metal infrastructure with RHCOS worker nodes
- The OpenShift cluster must be installed via installer provisioned infrastructure (IPI)
- Deploy Machine Health Checks to maintain HA for VMs
- A NetApp ONTAP cluster
- Astra Trident installed on the OpenShift cluster
- A Trident backend configured with an SVM on ONTAP cluster
- A StorageClass configured on the OpenShift cluster with Astra Trident as the provisioner
- Cluster-admin access to Red Hat OpenShift cluster
- Admin access to NetApp ONTAP cluster
- An admin workstation with tridentctl and oc tools installed and added to $PATH

Because OpenShift Virtualization is managed by an operator installed on the OpenShift cluster, it imposes additional overhead on memory, CPU, and storage, which must be accounted for while planning the hardware requirements for the cluster. See the documentation here for more details.

Optionally, you can also specify a subset of the OpenShift cluster nodes to host the OpenShift Virtualization operators, controllers, and VMs by configuring node placement rules. To configure node placement rules for

OpenShift Virtualization, follow the documentation here.

For the storage backing OpenShift Virtualization, NetApp recommends having a dedicated StorageClass that requests storage from a particular Trident backend, which in turn is backed by a dedicated SVM. This maintains a level of multitenancy with regard to the data being served for VM-based workloads on the OpenShift cluster.

Deploy Red Hat OpenShift Virtualization with NetApp ONTAP

To install OpenShift Virtualization, complete the following steps:

1. Log into the Red Hat OpenShift bare-metal cluster with cluster-admin access.
2. Select Administrator from the Perspective drop down.
3. Navigate to Operators > OperatorHub and search for OpenShift Virtualization.

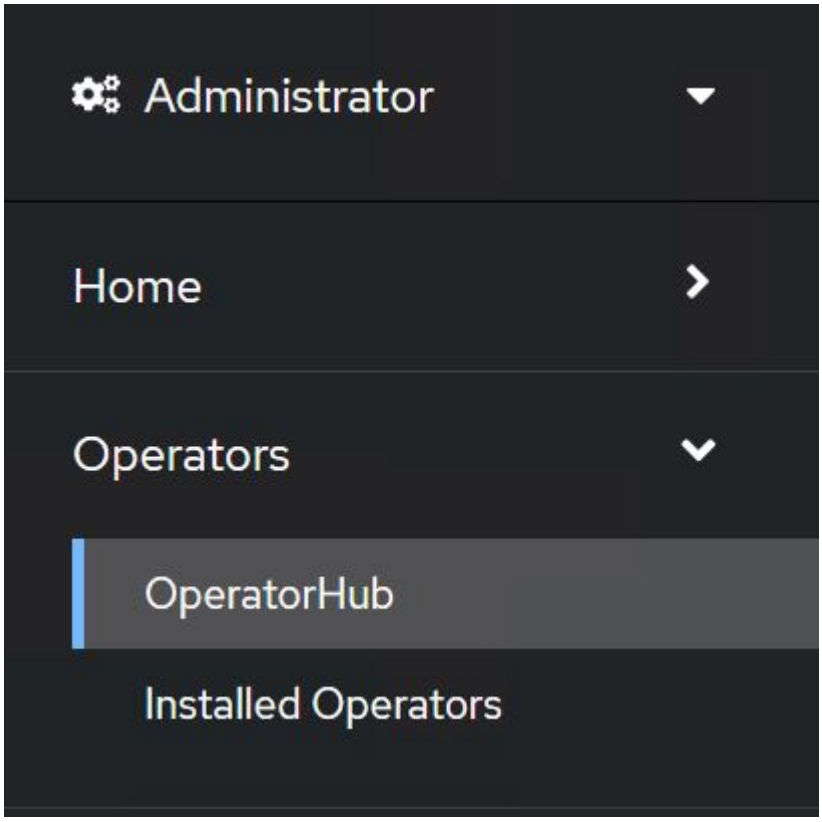

4. Select the OpenShift Virtualization tile and click Install.

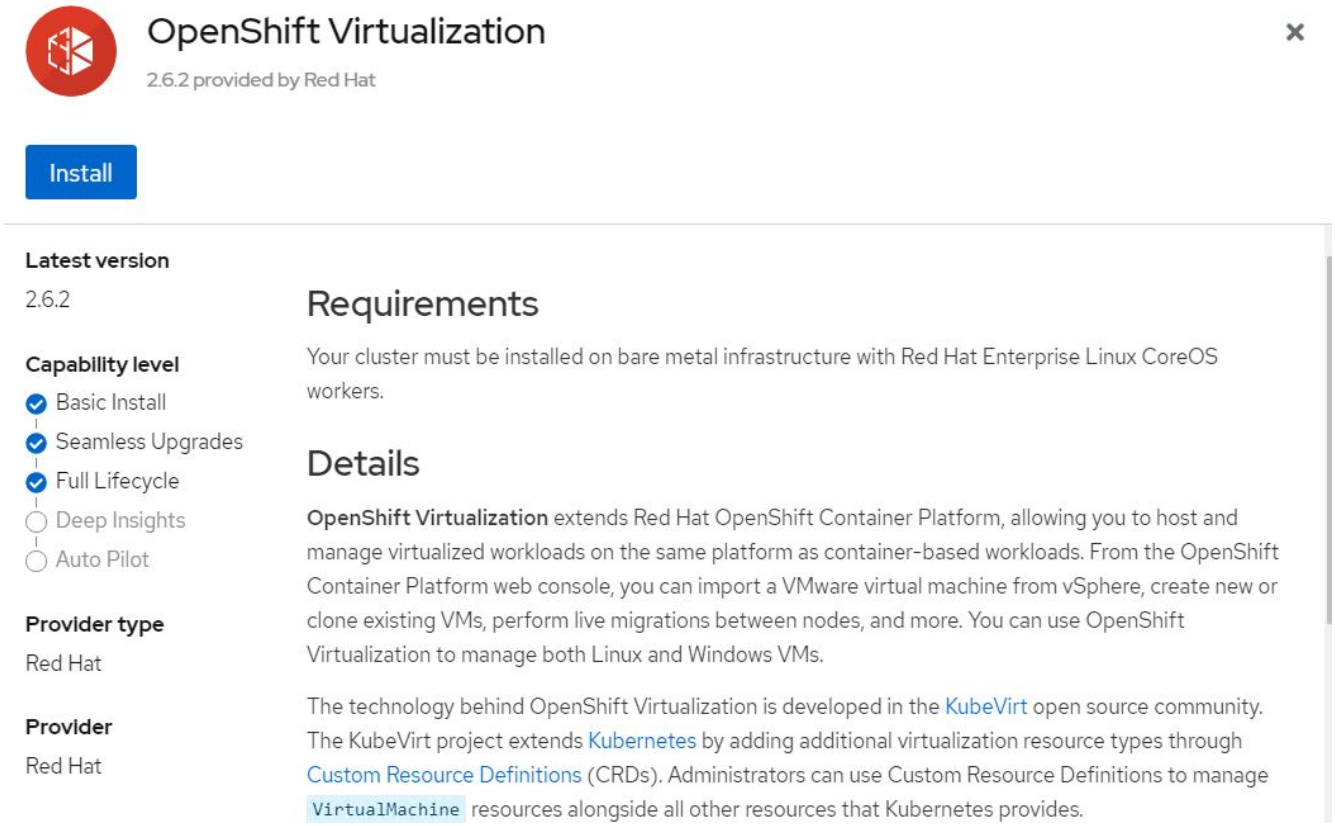

5. On the Install Operator screen, leave all default parameters and click Install.

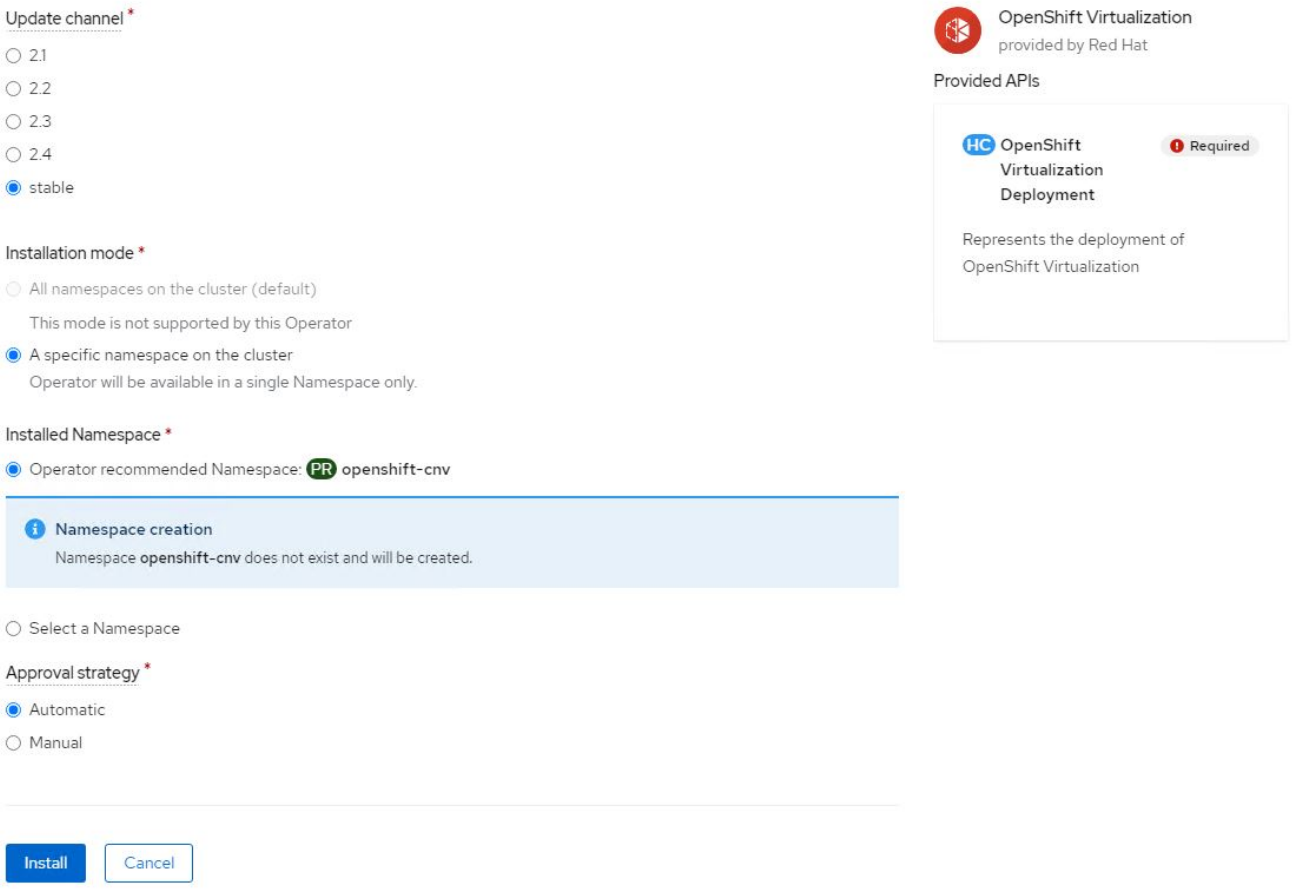

6. Wait for the operator installation to complete.

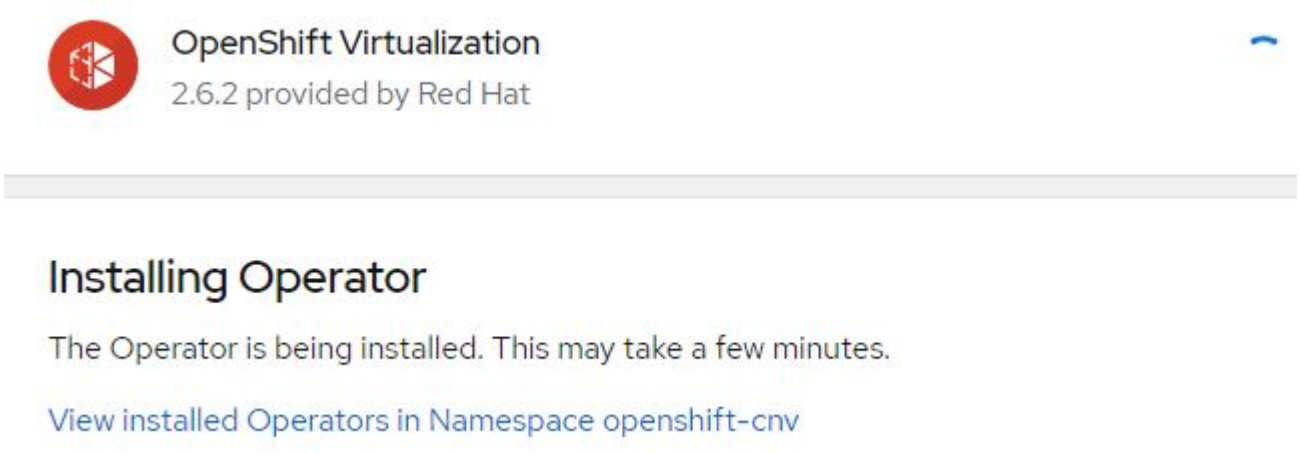

7. After the operator has installed, click Create HyperConverged.

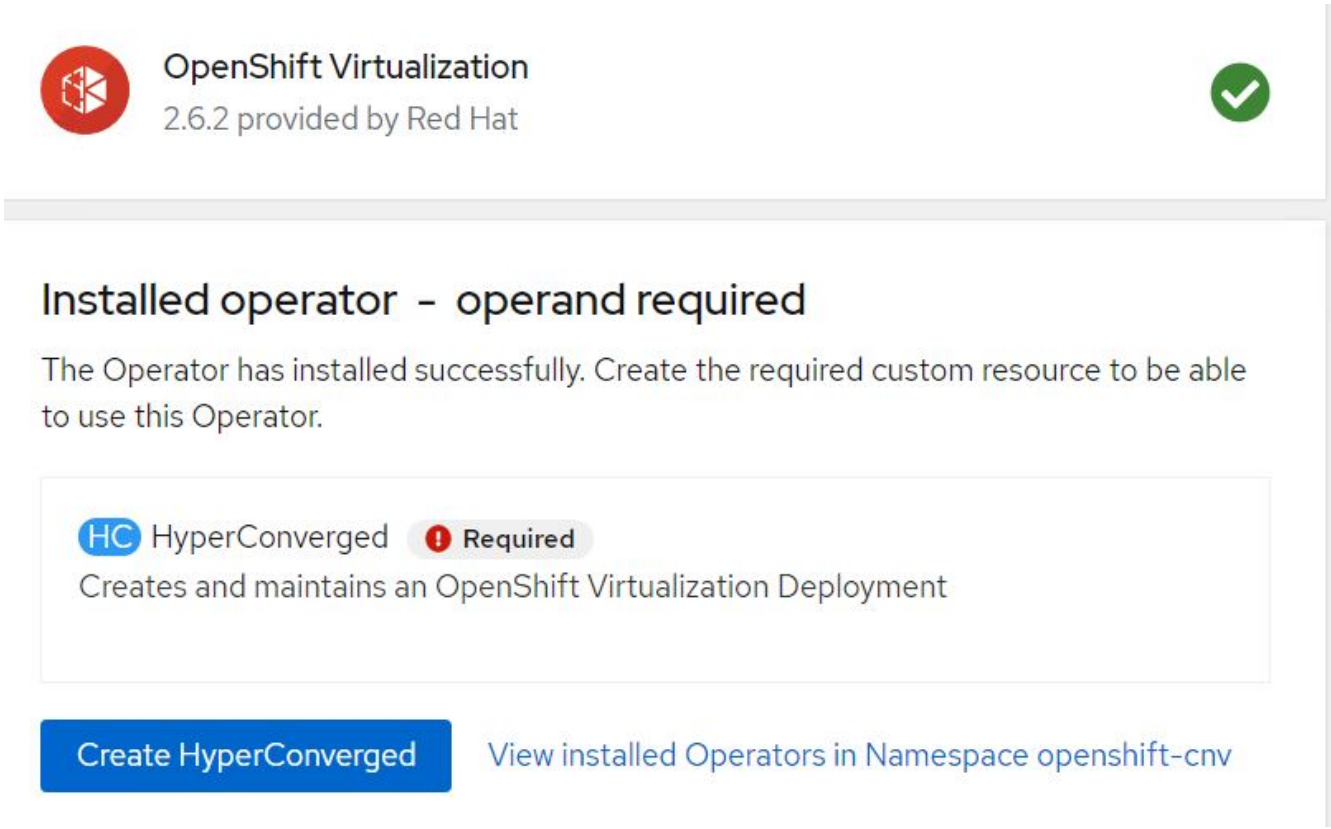

8. On the Create HyperConverged screen, click Create, accepting all default parameters. This step starts the installation of OpenShift Virtualization.

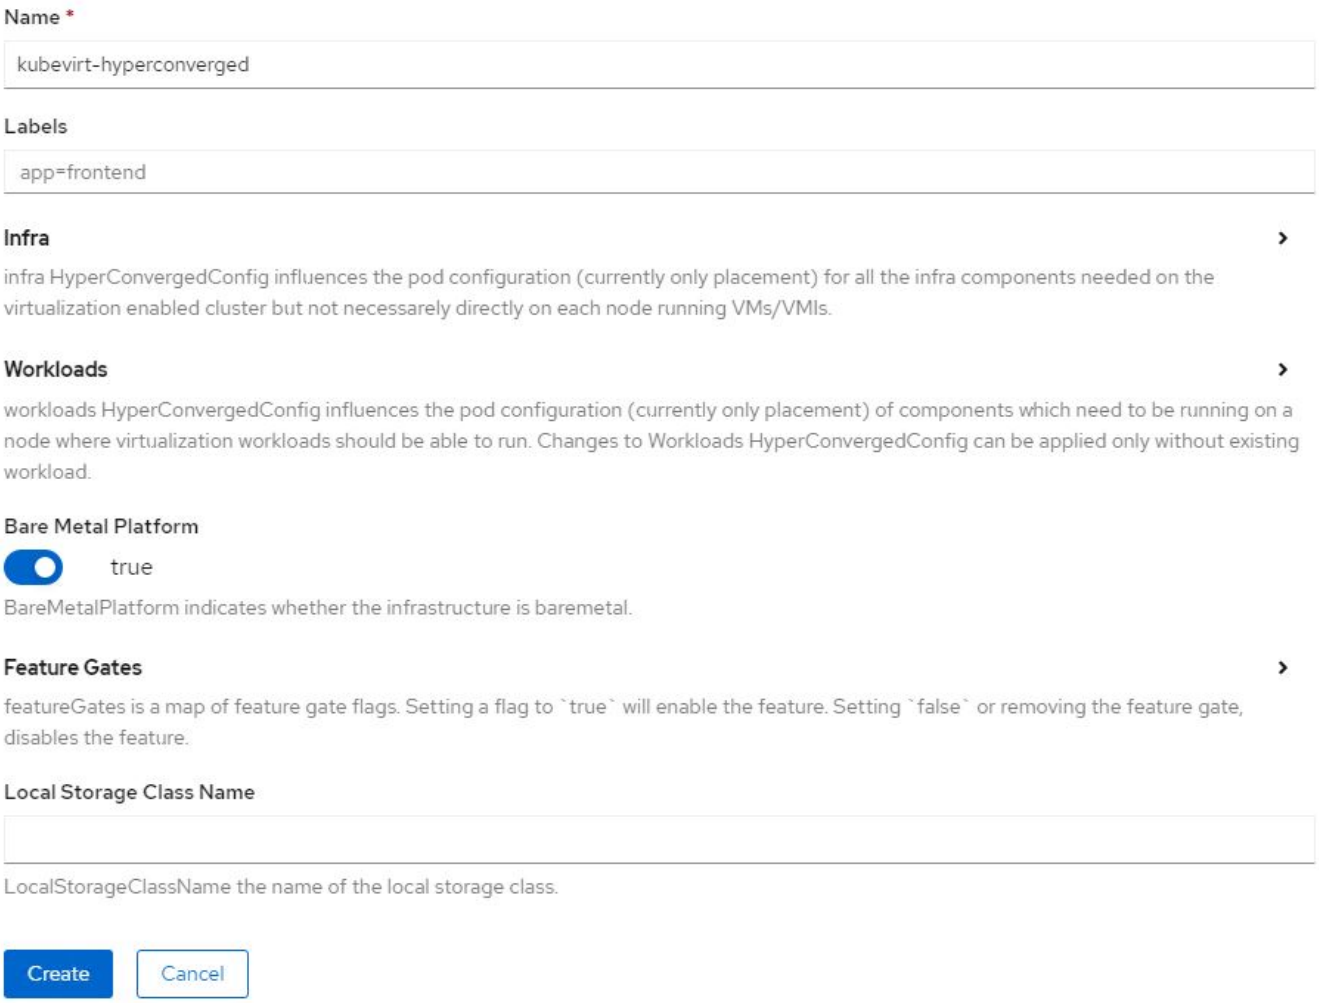

9. After all the pods move to the Running state in the openshift-cnv namespace and the OpenShift Virtualization operator is in the Succeeded state, the operator is ready to use. VMs can now be created on the OpenShift cluster.

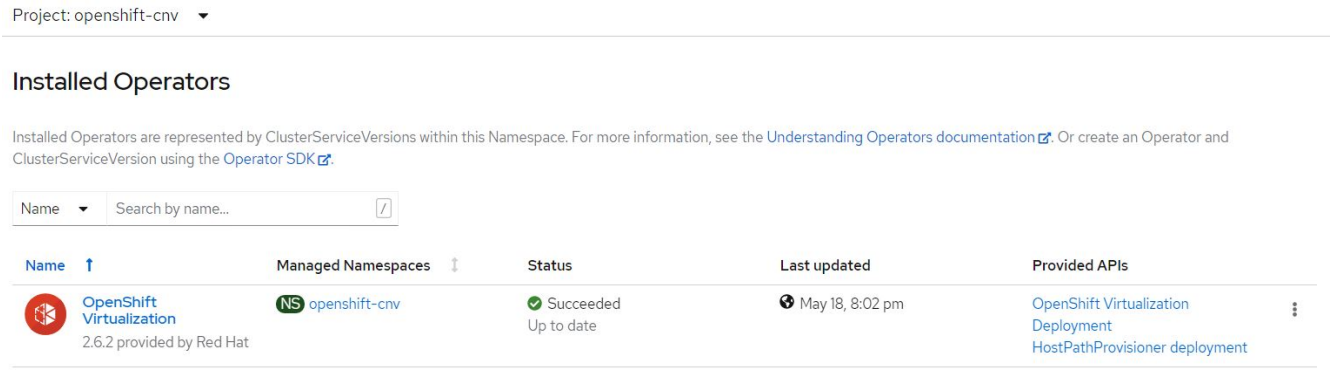

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section covers the how to create a virtual machine with Red Hat OpenShift Virtualization.

Create VM

VMs are stateful deployments that require volumes to host the operating system and data. With CNV, because the VMs are run as pods, the VMs are backed by PVs hosted on NetApp ONTAP through Trident. These volumes are attached as disks and store the entire filesystem including the boot source of the VM.

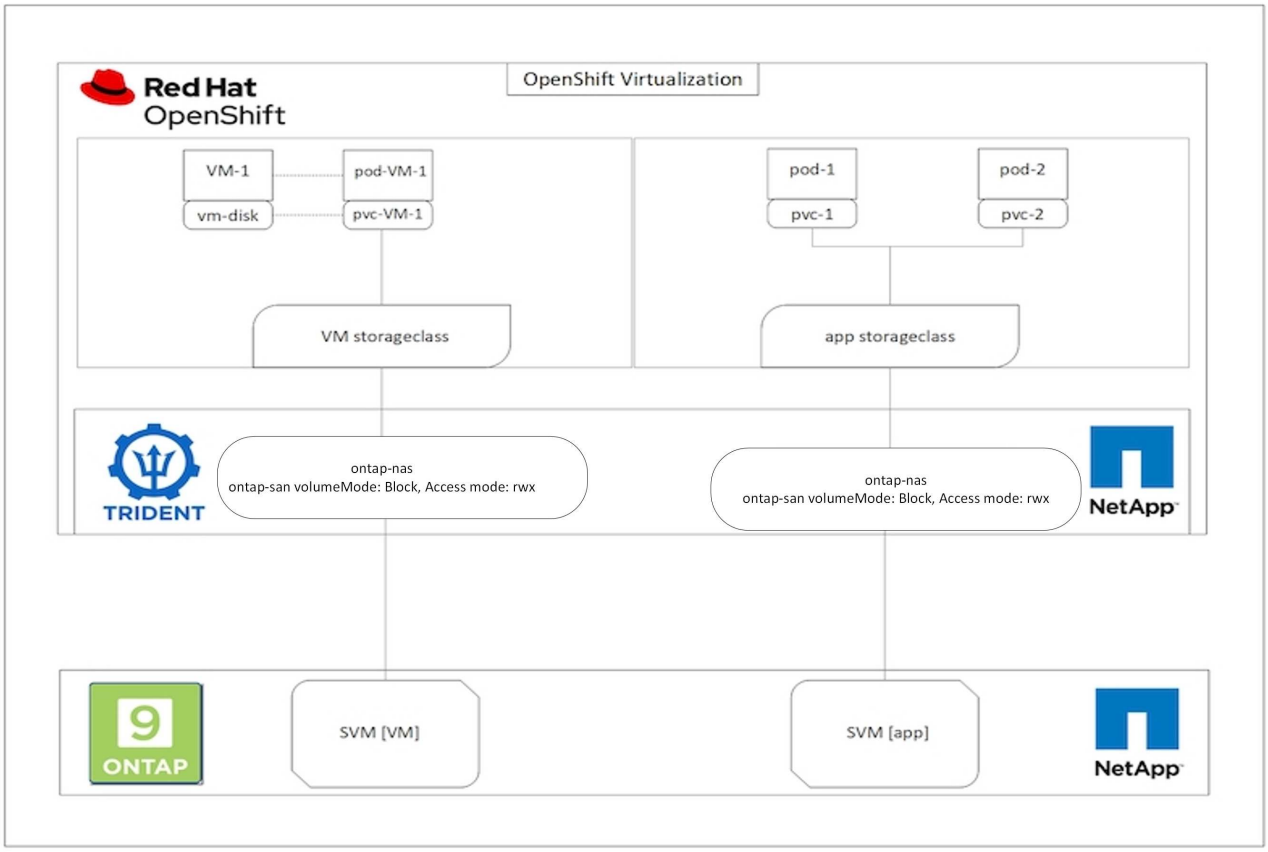

To quickly create a virtual machine on the OpenShift cluster, complete the following steps:

1. Navigate to Virtualization > Virtual Machines and click Create.
2. Select From template.
3. Select the desired operating system for which the boot source is available.
4. Check the checkbox Start the VirtualMachine after creation.
5. Click Quick create VirtualMachine.

The virtual machine is created and started and comes to the **Running** state. It automatically creates a PVC and a corresponding PV for the boot disk using the default storage class. In order to be able to live migrate the VM in the future, you must ensure that the storage class used for the disks can support RWX volumes. This is a requirement for live migration. ontap-nas and ontap-san (volumeMode block for iSCSI and NVMe/TCP protocols) can support RWX access modes for the volumes created using the respective storage classes.

To configure ontap-san storage class on the cluster see the Section for Migrating a VM from VMware to OpenShift Virtualization.

You can set up ontap NAS or iSCSI as the default storage class for the cluster. Clicking on Quick create VirtualMachine will use the default storage class to create the PVC and PV for the bootable root disk for the VM. If your default storage class is not ontap-nas or ontap-san, you can select the storage class for the disk, by selecting Customize VirtualMachine > Customize VirtualMachine parameters > Disks and then editing the disk to use the required storage class.

Typically block access mode is preferred compared to file systems while provisioning the VM disks.

To customize the virtual machine creation after you have selected the OS template, click on Customize VirtualMachine instead of Quick create.

1. If the selected operating system has boot source configured, you can click on **Customize VirtualMachine parameters**.
2. If the selected operating system has no boot source configured, you must configure it. You can see details about the procedures shown in the documentation.
3. After Configuring the boot disk, you can click on **Customize VirtualMachine parameters**.
4. You can customize the VM from the tabs on this page. For eg. click on the **Disks** tab and then click on **Add disk** to add another disk to the VM.
5. Click Create Virtual Machine to create the virtual machine; this spins up a corresponding pod in the background.

When a boot source is configured for a template or an operating system from an URL or from a registry, it creates a PVC in the `openshift-virtualization-os-images` project and downloads the KVM guest image to the PVC. You must make sure that template PVCs have enough provisioned space to accommodate the KVM guest image for the corresponding OS. These PVCs are then cloned and attached as rootdisk to virtual machines when they are created using the respective templates in any project.

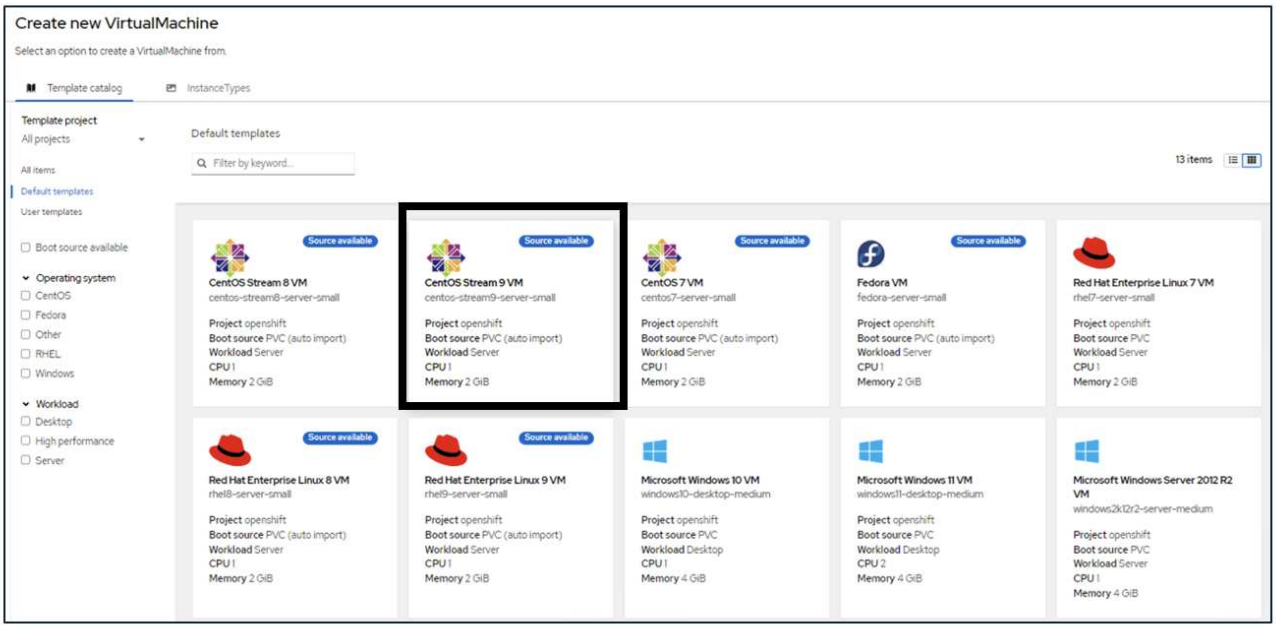
Create new VirtualMachine
Select an option to create a VirtualMachine from.
Template catalog
InstanceTypes
Template project
All projects
All items
Default templates
User templates
Boot source available
Operating system
CentOS
Fedora
Other
RHEL
Windows
Workload
Desktop
High performance
Server
Default templates
Filter by keyword...
13 items
Source available
CentOS Stream 8 VM
centos-stream8-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB
Source available
CentOS Stream 9 VM
centos-stream9-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB
Source available
CentOS 7 VM
centos7-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB
Source available
Fedora VM
fedora-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB
Red Hat Enterprise Linux 7 VM
rhel7-server-small
Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 2 GiB
Source available
Red Hat Enterprise Linux 8 VM
rhel8-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB
Source available
Red Hat Enterprise Linux 9 VM
rhel9-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB
Microsoft Windows 10 VM
windows10-desktop-medium
Project openshift
Boot source PVC
Workload Desktop
CPU 1
Memory 4 GiB
Microsoft Windows 11 VM
windows11-desktop-medium
Project openshift
Boot source PVC
Workload Desktop
CPU 2
Memory 4 GiB
Microsoft Windows Server 2012 R2 VM
windows2k12r2-server-medium
Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 4 GiB

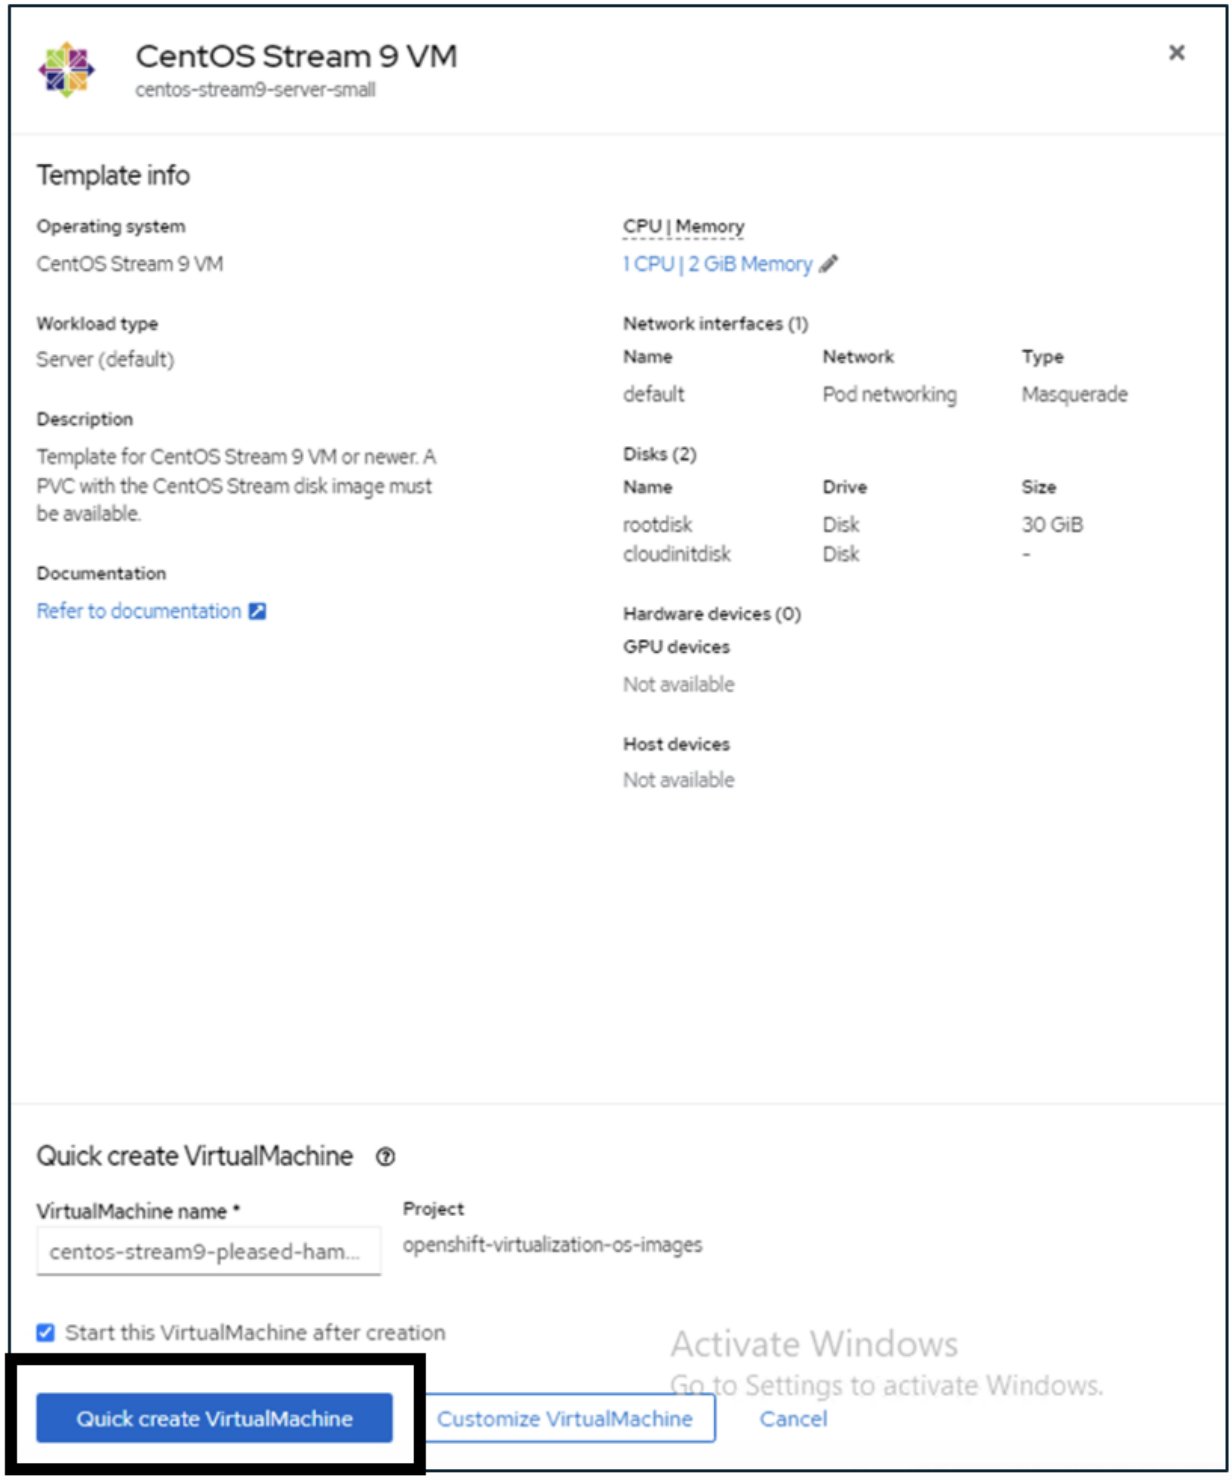
CentOS Stream 9 VM
centos-stream9-server-small
Template info
Operating system
CentOS Stream 9 VM
Workload type
Server (default)
Description
Template for CentOS Stream 9 VM or newer. A PVC with the CentOS Stream disk image must be available.
Documentation
Refer to documentation
CPU | Memory
1 CPU | 2 GiB Memory
Network interfaces (1)
Name Network Type
default Pod networking Masquerade
Disks (2)
Name Drive Size
rootdisk Disk 30 GiB
cloudinitdisk Disk -
Hardware devices (0)
GPU devices
Not available
Host devices
Not available
Quick create VirtualMachine
VirtualMachine name *
centos-stream9-pleased-ham...
Project
openshift-virtualization-os-images
Start this VirtualMachine after creation
Activate Windows
Go to Settings to activate Windows.
Quick create VirtualMachine
Customize VirtualMachine
Cancel

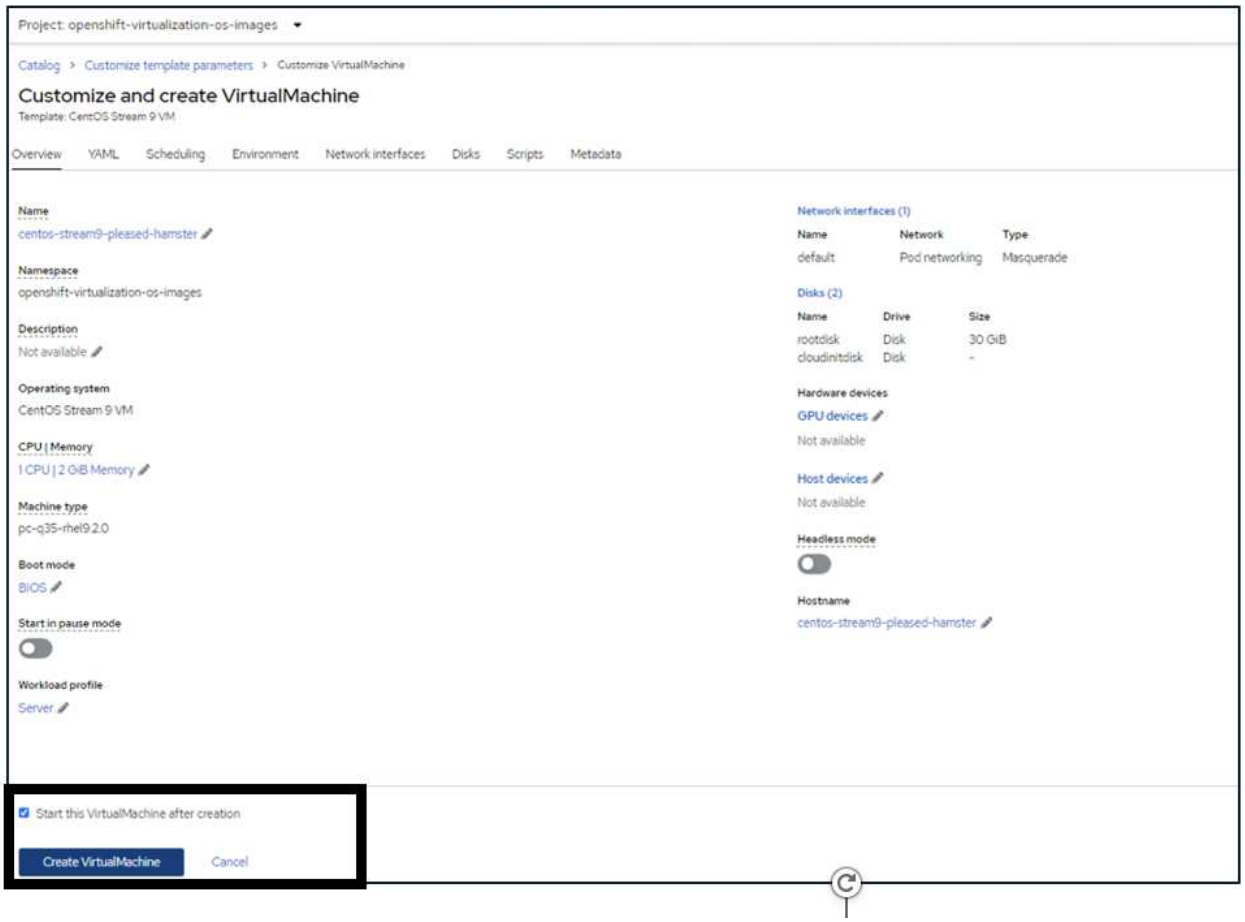

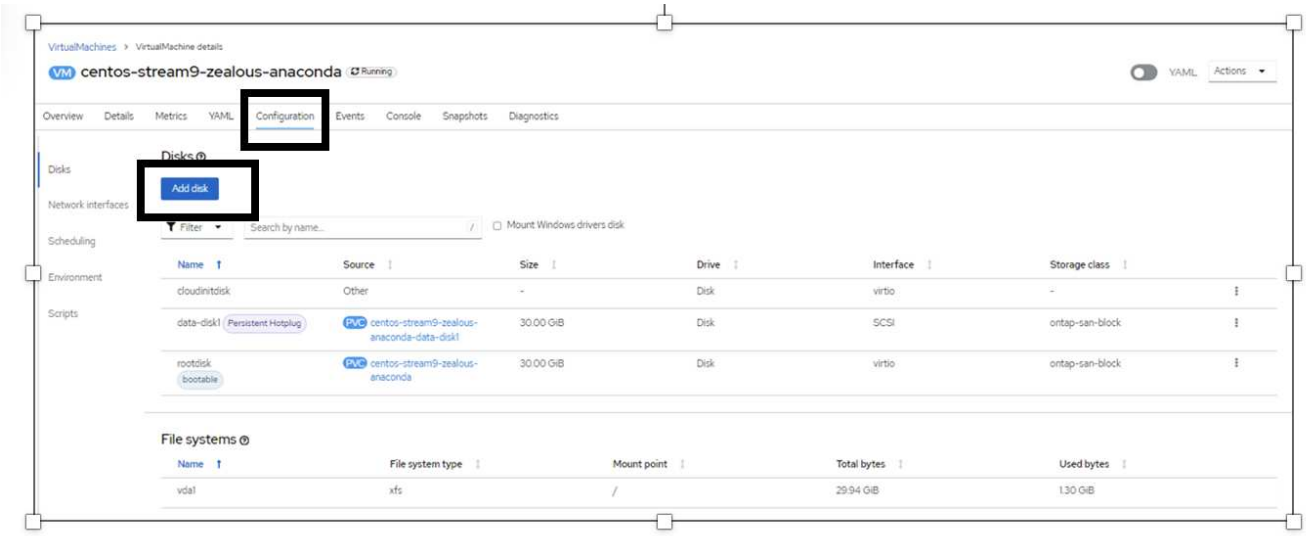

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section covers the how to migrate a virtual machine between from VMware to an OpenShift Cluster using Red Hat OpenShift Virtualization migration toolkit.

Migration of VM from VMware to OpenShift Virtualization using Migration Toolkit for Virtualization

In this section, we will see how to use the Migration Toolkit for Virtualization (MTV) to migrate virtual machines from VMware to OpenShift Virtualization running on OpenShift Container platform and integrated with NetApp

ONTAP storage using Astra Trident.

The following video shows a demonstration of the migration of a RHEL VM from VMware to OpenShift Virtualization using ontap-san storage class for persistent storage.

Using Red Hat MTV to migrate VMs to OpenShift Virtualization with NetApp ONTAP Storage

The following diagram shows a high level view of the migration of a VM from VMware to Red Hat OpenShift Virtualization.

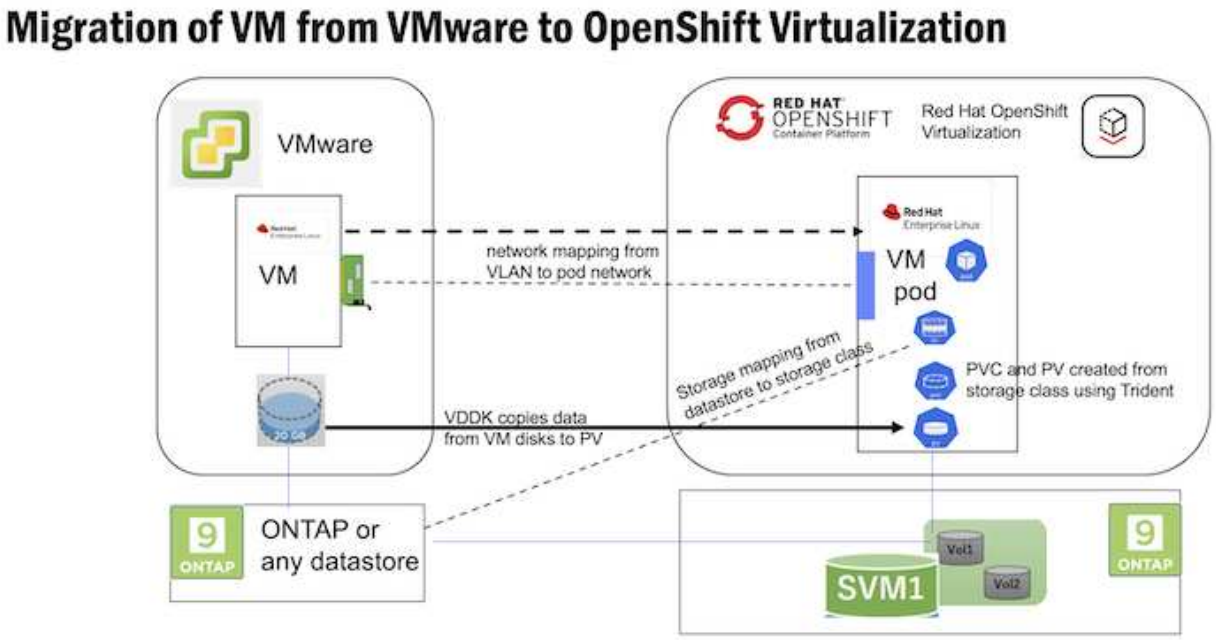

Prerequisites for the sample migration

On VMware

- A RHEL 9 VM using rhel 9.3 with the following configurations were installed:
  - CPU: 2, Memory: 20 GB, Hard disk: 20 GB
  - user credentials: root user and an admin user credentials
- After the VM was ready, postgresql server was installed.
  - postgresql server was started and enabled to start on boot

```
systemctl start postgresql.service`
systemctl enable postgresql.service
The above command ensures that the server can start in the VM in
OpenShift Virtualization after migration
```

  - Added 2 databases, 1 table and 1 row in the table were added. Refer here for the instructions for installing postgresql server on RHEL and creating database and table entries.

Ensure that you start the postgresql server and enable the service to start at boot.

On OpenShift Cluster

The following installations were completed before installing MTV:

- OpenShift Cluster 4.13.34
- Astra Trident 23.10
- Multipath on the cluster nodes enabled for iSCSI (for ontap-san storage class). See the provided yaml to create a daemon set that enables iSCSI on each node in the cluster.
- Trident backend and Storage class for ontap SAN using iSCSI. See the provided yaml files for trident backend and storage class.
- OpenShift Virtualization

To install iscsi and multipath on the OpenShift Cluster nodes use the yaml file given below
Preparing the cluster nodes for iSCSI

```
apiVersion: apps/v1
kind: DaemonSet
metadata:
  namespace: trident
  name: trident-iscsi-init
  labels:
    name: trident-iscsi-init
spec:
  selector:
    matchLabels:
      name: trident-iscsi-init
  template:
    metadata:
      labels:
        name: trident-iscsi-init
    spec:
      hostNetwork: true
      serviceAccount: trident-node-linux
      initContainers:
      - name: init-node
        command:
          - nsenter
          - --mount=/proc/1/ns/mnt
          - --
          - sh
          - -c
        args: ["$(STARTUP_SCRIPT)"]
        image: alpine:3.7
        env:
```

```
        - name: STARTUP_SCRIPT
          value: |
            #! /bin/bash
            sudo yum install -y lsscsi iscsi-initiator-utils sg3_utils
device-mapper-multipath
            rpm -q iscsi-initiator-utils
            sudo sed -i 's/^\(node.session.scan\).*/\1 = manual/'
/etc/iscsi/iscsid.conf
            cat /etc/iscsi/initiatorname.iscsi
            sudo mpathconf --enable --with_multipathd y --find_multipaths
n
            sudo systemctl enable --now iscsid multipathd
            sudo systemctl enable --now iscsi
        securityContext:
          privileged: true
      hostPID: true
      containers:
      - name: wait
        image: k8s.gcr.io/pause:3.1
      hostPID: true
      hostNetwork: true
      tolerations:
      - effect: NoSchedule
        key: node-role.kubernetes.io/master
  updateStrategy:
    type: RollingUpdate
```

Use the following yaml file to create trident backend configuration for using ontap san storage
Trident backend for iSCSI

```
apiVersion: v1
kind: Secret
metadata:
  name: backend-tbc-ontap-san-secret
type: Opaque
stringData:
  username: <username>
  password: <password>
---
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: ontap-san
spec:
  version: 1
  storageDriverName: ontap-san
  managementLIF: <management LIF>
  backendName: ontap-san
  svm: <SVM name>
  credentials:
    name: backend-tbc-ontap-san-secret
```

Use the following yaml file to create trident storage class configuration for using ontap san storage
Trident storage class for iSCSI

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ontap-san
provisioner: csi.trident.netapp.io
parameters:
  backendType: "ontap-san"
  media: "ssd"
  provisioningType: "thin"
  snapshots: "true"
allowVolumeExpansion: true
```

Install MTV

Now you can install the Migration Toolkit for virtualization (MTV). Refer to the instructions provided here for help with the installation.

The Migration Toolkit for Virtualization (MTV) user interface is integrated into the OpenShift web console. You can refer here to start using the user interface for various tasks.

Create Source Provider

In order to migrate the RHEL VM from VMware to OpenShift Virtualization, you need to first create the source provider for VMware. Refer to the instructions here to create the source provider.

You need the following to create your VMware source provider:

- VCenter url
- VCenter Credentials
- VCenter server thumbprint
- VDDK image in a repository

Sample source provider creation:

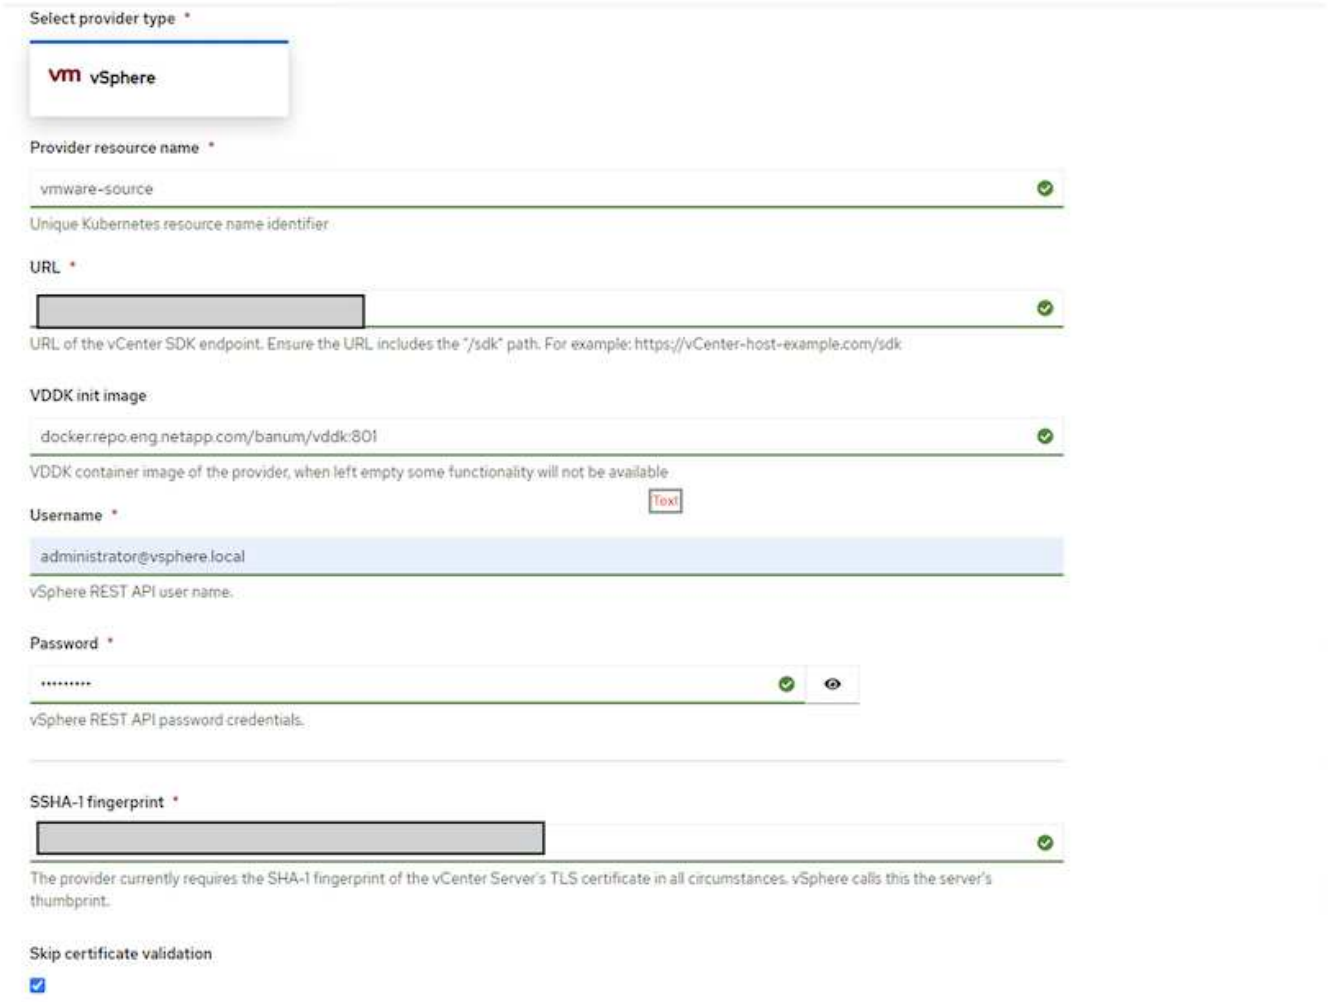

The Migration Toolkit for Virtualization (MTV) uses the VMware Virtual Disk Development Kit (VDDK) SDK to accelerate transferring virtual disks from VMware vSphere. Therefore, creating a VDDK image, although optional, is highly recommended.
To make use of this feature, you download the VMware Virtual Disk Development Kit (VDDK), build a VDDK image, and push the VDDK image to your image registry.

Follow the instructions provided here to create and push the VDDK image to a registry accessible from the OpenShift Cluster.

Create Destination provider

The host cluster is automatically added as the OpenShift virtualization provider is the source provider.

Create Migration Plan

Follow the instructions provided here to create a migration plan.

While creating a plan, you need to create the following if not already created:

- A network mapping to map the source network to the target network.
- A storage mapping to map the source datastore to the target storage class. For this you can choose ontap-san storage class.
  Once the migration plan is created, the status of the plan should show **Ready** and you should now be able to **Start** the plan.

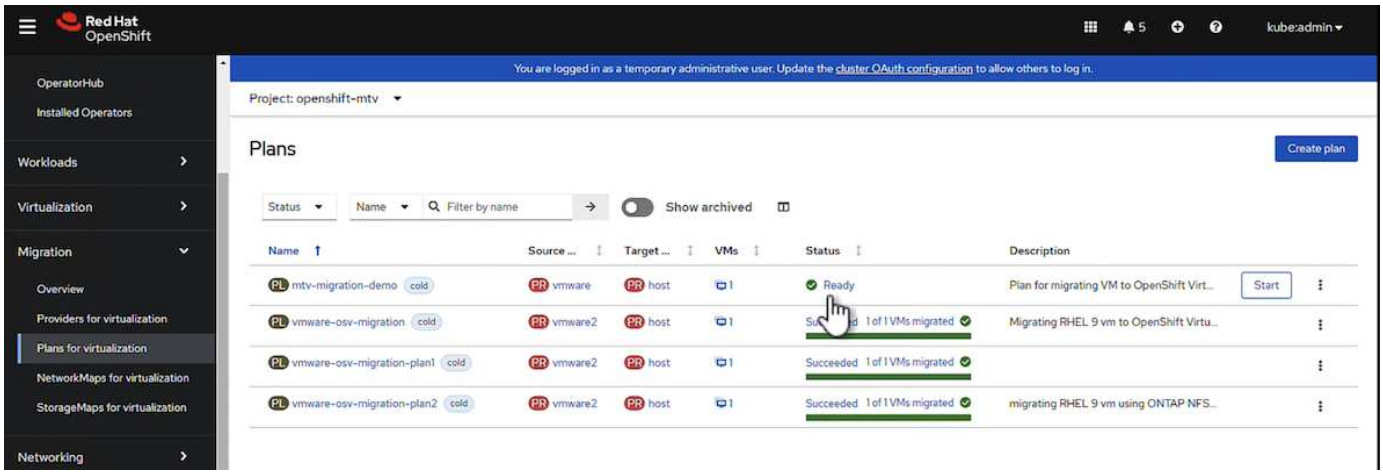

Clicking on **Start** will run through a sequence of steps to complete the migration of the VM.

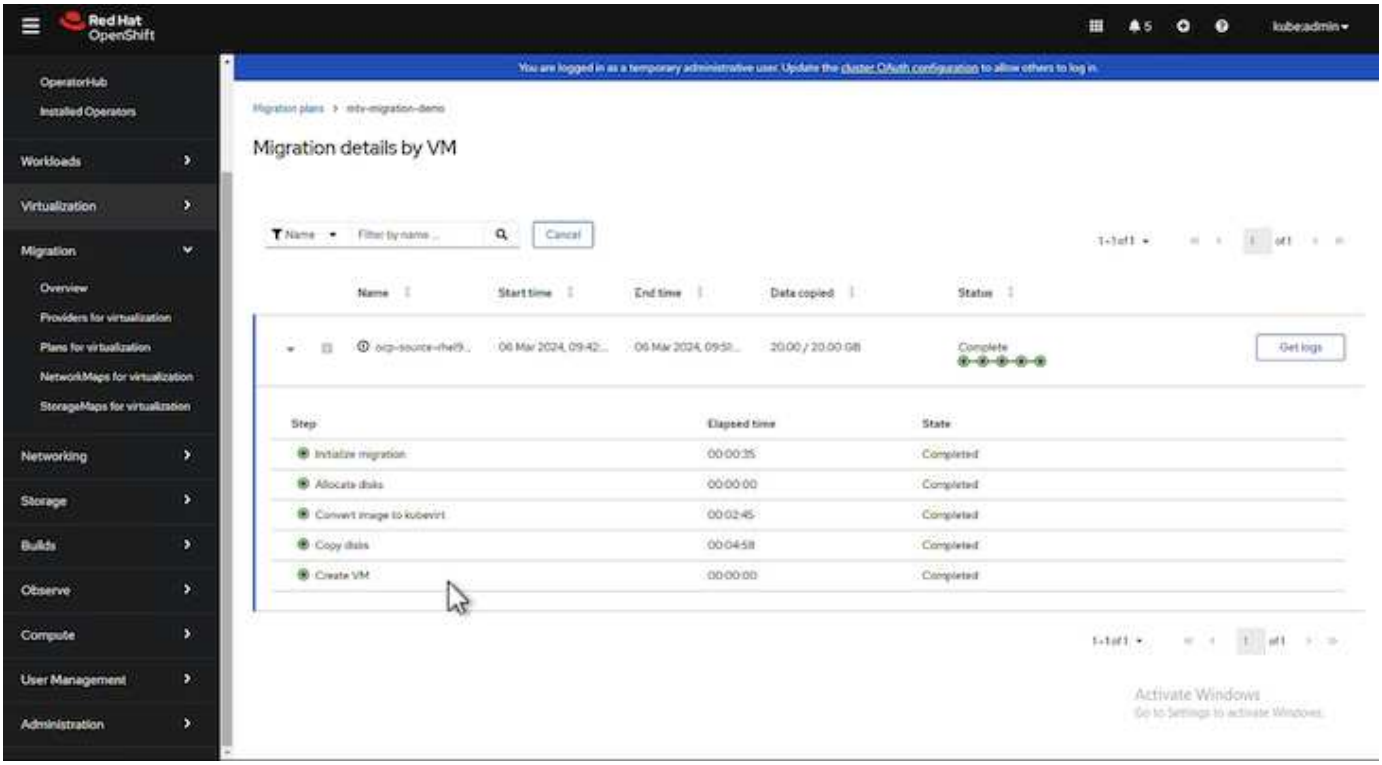

When all steps are completed, you can see the migrated VMs by clicking on the **virtual machines** under **Virtualization** in the left-side navigation menu.
Instructions to access the virtual machines are provided here.

You can log into the virtual machine and verify the contents of the posgresql databases. The databases, tables and the entries in the table should be the same as what was created on the source VM.

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section shows how to migrate a virtual machine in OpenShift Virtualization between nodes in the cluster .

VM Live Migration

Live Migration is a process of migrating a VM instance from one node to another in an OpenShift cluster with no downtime. For live migration to work in an OpenShift cluster, VMs must be bound to PVCs with shared ReadWriteMany access mode. Astra Trident backends configured using ontap-nas drivers support RWX access mode for FileSystem protocols nfs and smb. Refer to the documentation here. Astra Trident backends configured using ontap-san drivers support RWX access mode for block volumeMode for iSCSI and NVMe/TCP protocols. Refer to the documentation here.

Therefore, for live migration to succeed, the VMs must be provisioned with disks (boot disks and additional hot plug disks) with PVCs using ontap-nas or ontap-san (volumeMode: Block) storage classes. When the PVCs are created, Trident creates ONTAP volumes in an SVM which is NFS-enabled or iSCSI enabled.

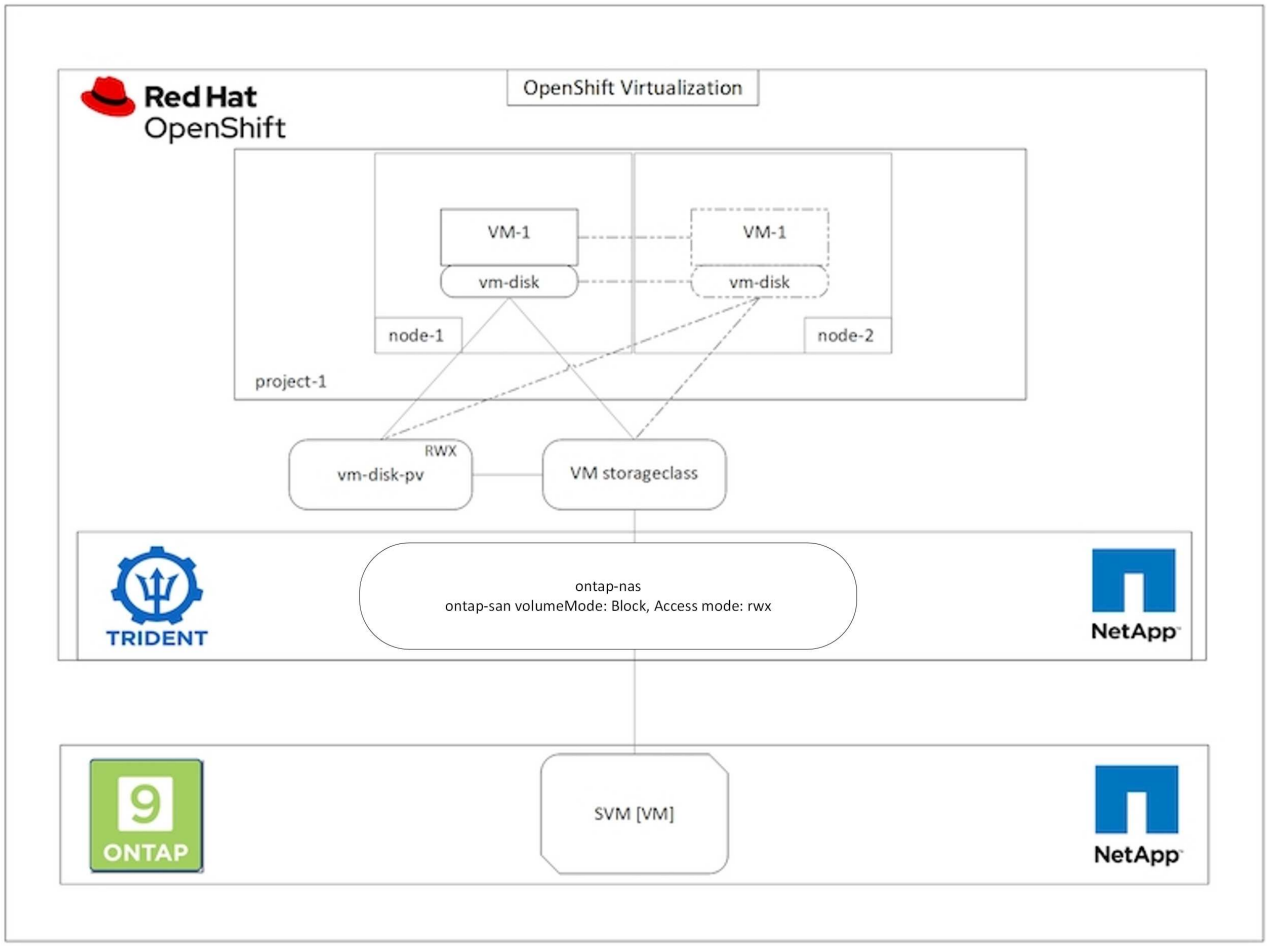

To perform a live migration of a VM that has been created previously and is in a Running state perform the following steps:

1. Select the VM that you want to live-migrate.
2. Click on **Configuration** tab.
3. Ensure that all the disks of the VM are created using Storage classes that can support RWX access mode.
4. Click on **Actions** on the right corner and then select **Migrate**.
5. To look at the progression of the Migration, go to Virtualization > Overview on the left hand side menu and then click on the **Migrations** tab.

The Migration of the VM will transition from **Pending** to **Scheduling** to **Succeeded**

A VM instance in an OpenShift cluster automatically migrates to another node when the original node is placed into maintenance mode if the evictionStrategy is set to LiveMigrate.

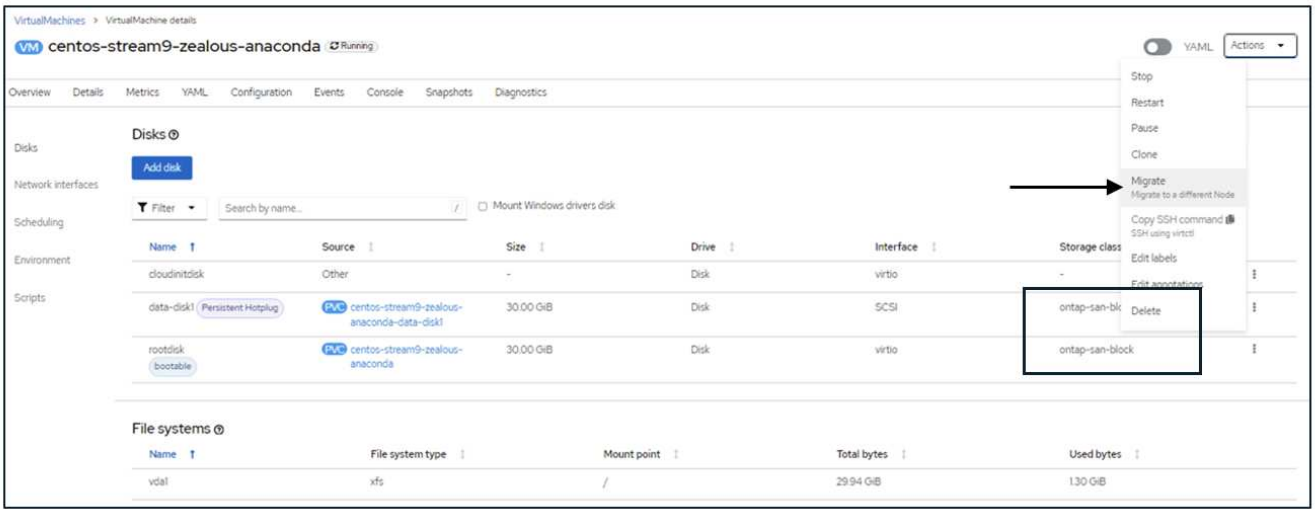

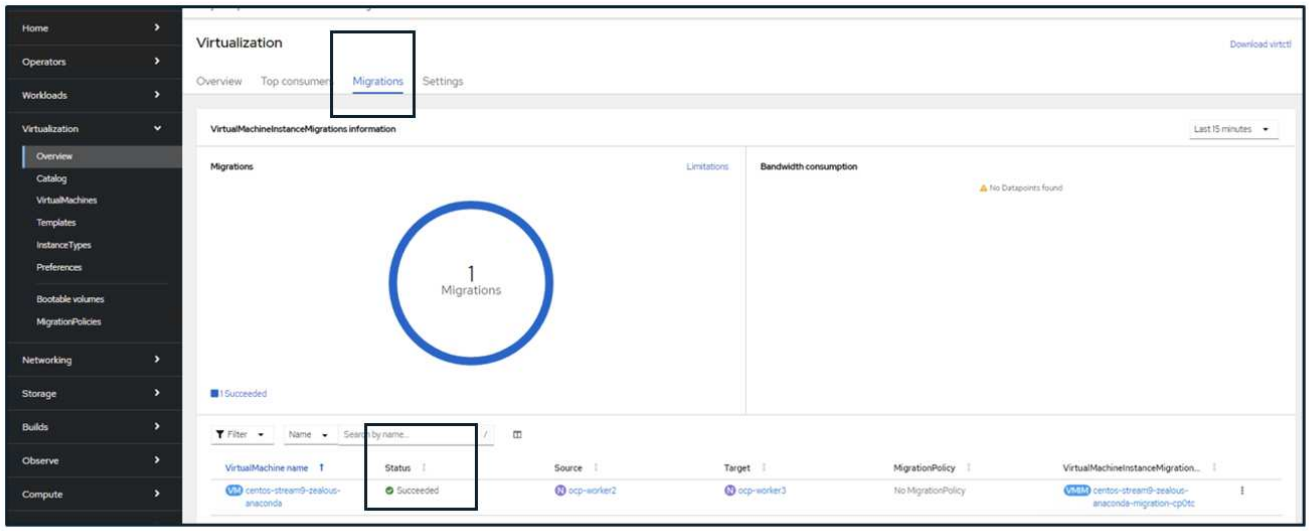

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section covers the how to clone a virtual machine with Red Hat OpenShift Virtualization.

VM cloning

Cloning an existing VM in OpenShift is achieved with the support of Astra Trident's Volume CSI cloning feature. CSI volume cloning allows for creation of a new PVC using an existing PVC as the data source by duplicating its PV. After the new PVC is created, it functions as a separate entity and without any link to or dependency on the source PVC.

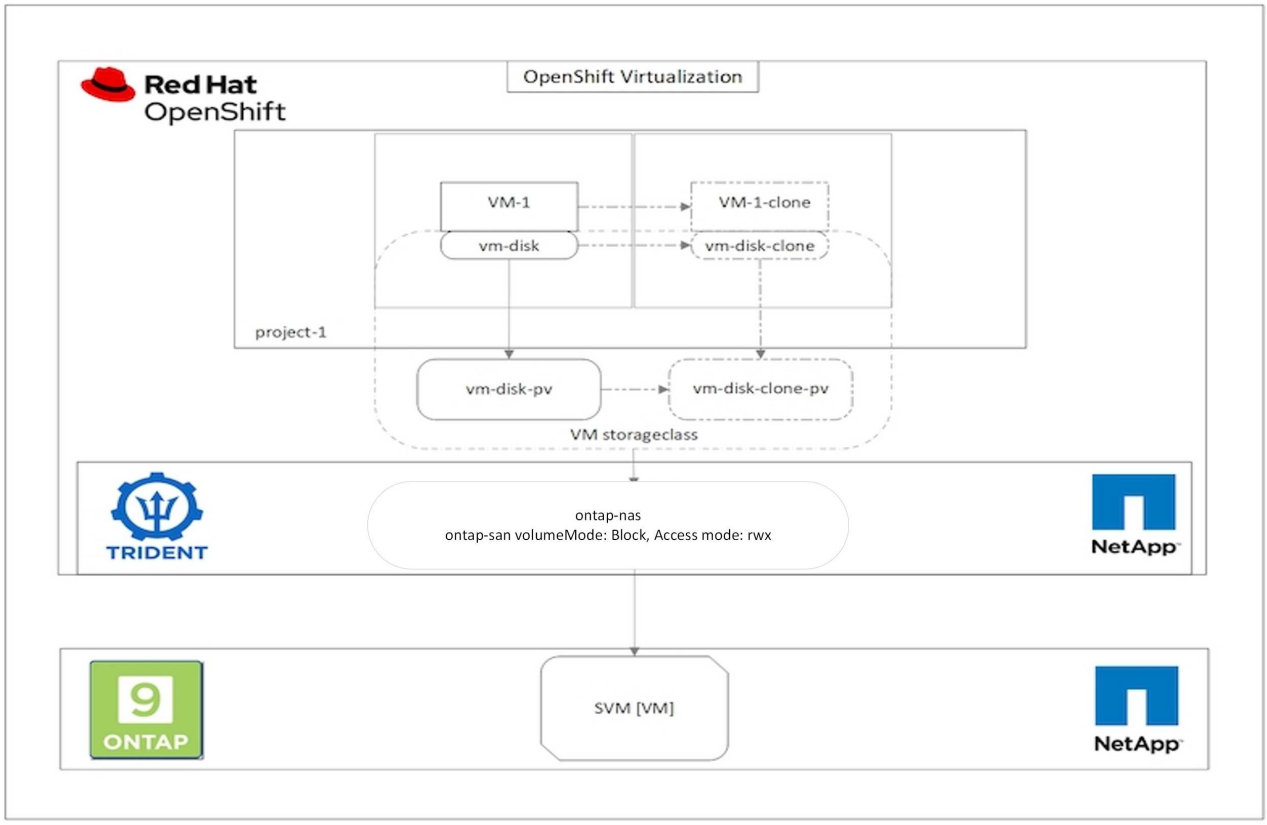

There are certain restrictions with CSI volume cloning to consider:

1. Source PVC and destination PVC must be in the same project.
2. Cloning is supported within the same storage class.
3. Cloning can be performed only when source and destination volumes use the same VolumeMode setting; for example, a block volume can only be cloned to another block volume.

VMs in an OpenShift cluster can be cloned in two ways:

1. By shutting down the source VM
2. By keeping the source VM live

By Shutting down the source VM

Cloning an existing VM by shutting down the VM is a native OpenShift feature that is implemented with support from Astra Trident. Complete the following steps to clone a VM.

1. Navigate to Workloads > Virtualization > Virtual Machines and click the ellipsis next to the virtual machine you wish to clone.
2. Click Clone Virtual Machine and provide the details for the new VM.

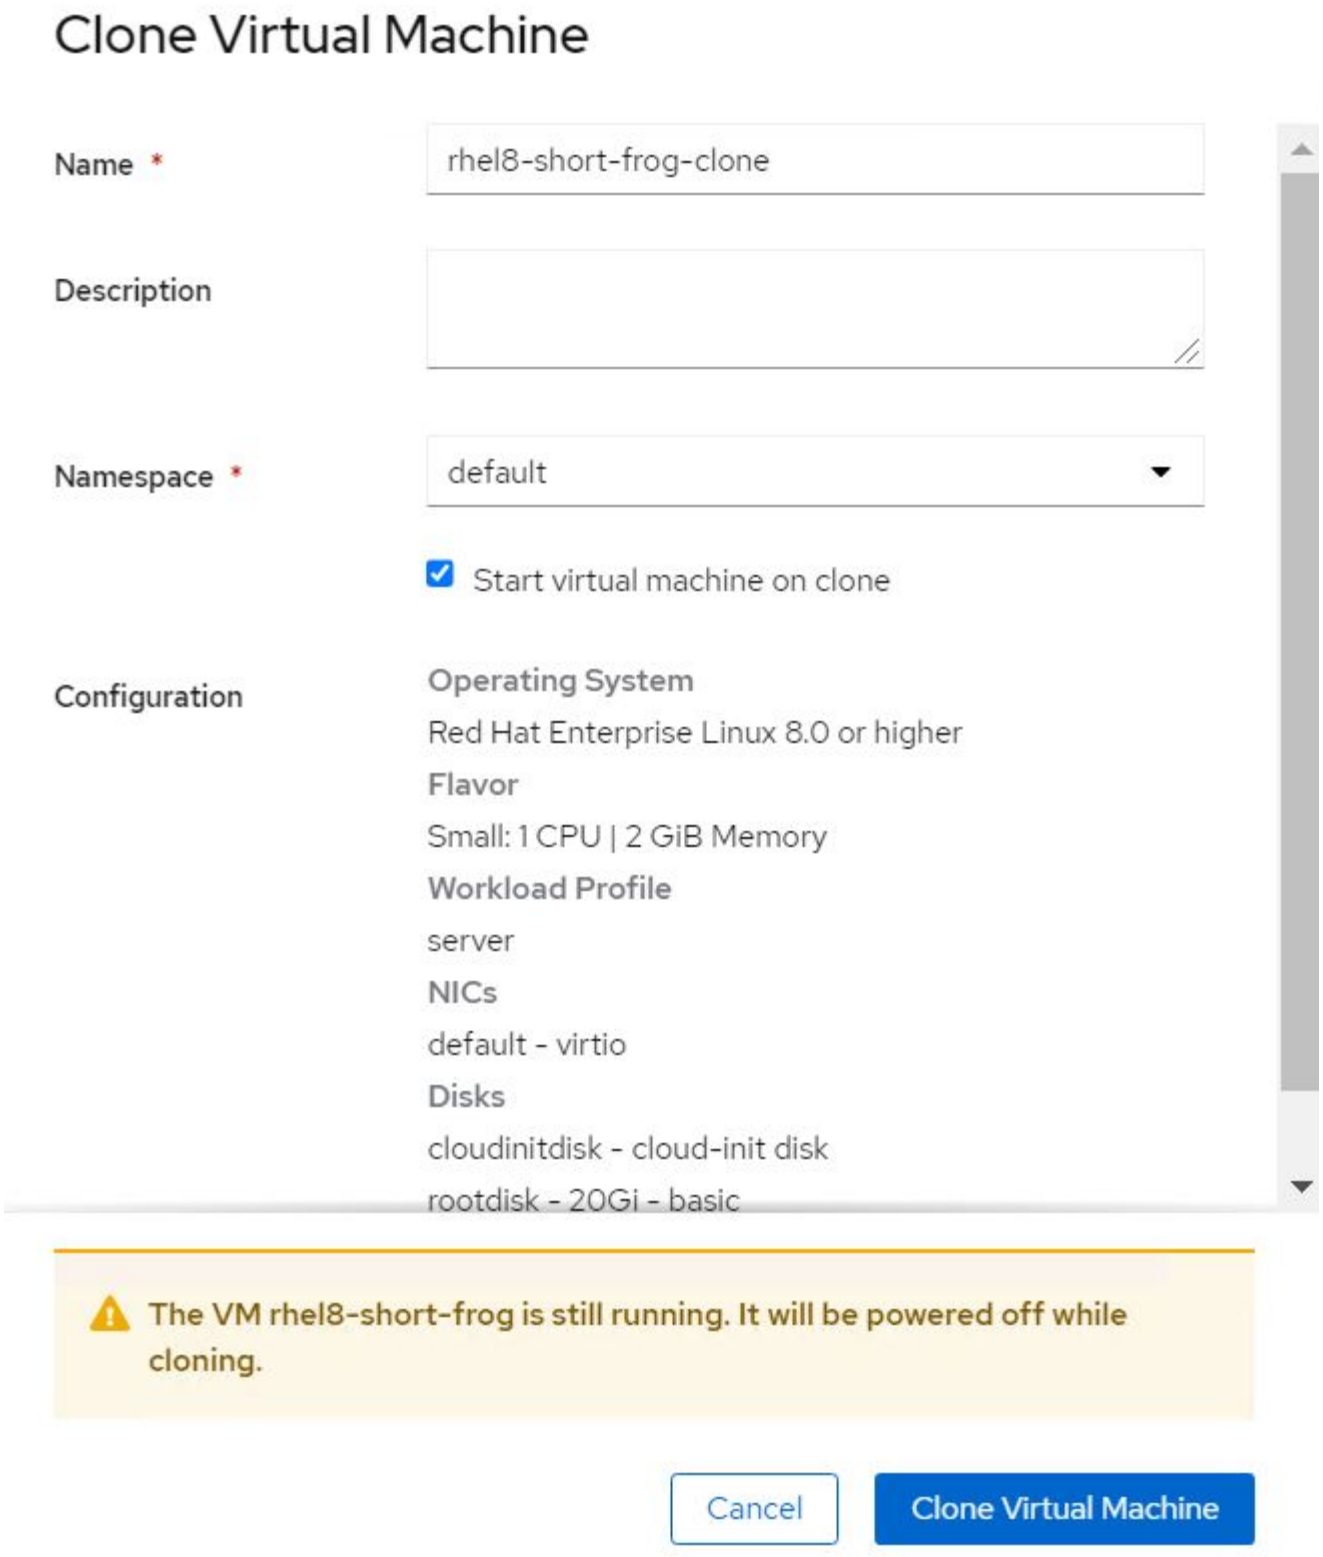

3. Click Clone Virtual Machine; this shuts down the source VM and initiates the creation of the clone VM.
4. After this step is completed, you can access and verify the content of the cloned VM.

By keeping the source VM live

An existing VM can also be cloned by cloning the existing PVC of the source VM and then creating a new VM using the cloned PVC. This method does not require you to shut down the source VM. Complete the following steps to clone a VM without shutting it down.

1. Navigate to Storage > PersistentVolumeClaims and click the ellipsis next to the PVC that is attached to the source VM.
2. Click Clone PVC and furnish the details for the new PVC.

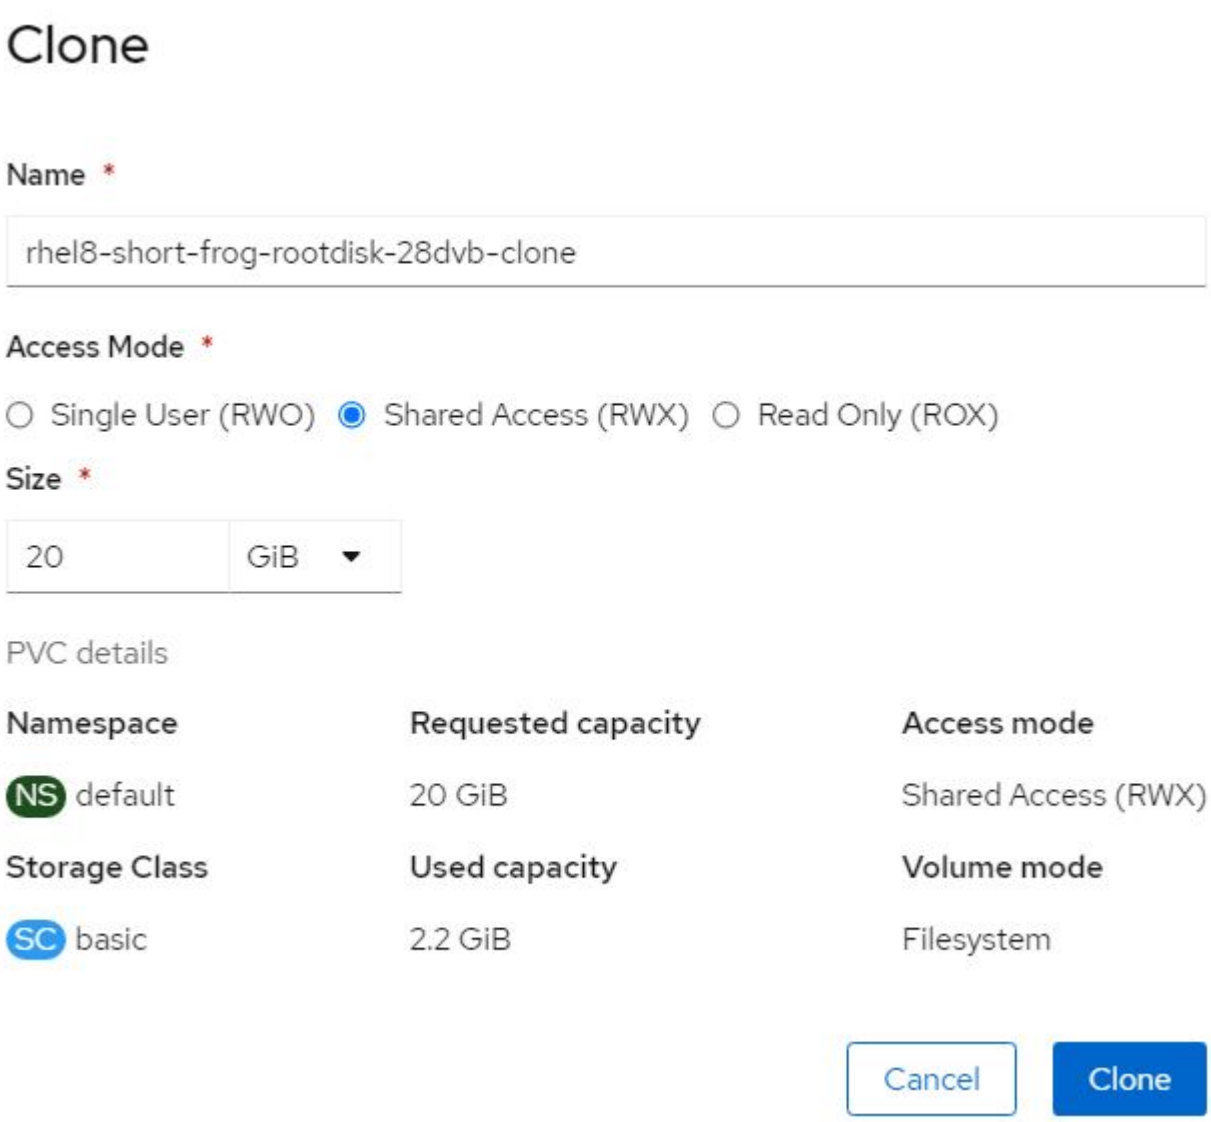

3. Then click Clone. This creates a PVC for the new VM.
4. Navigate to Workloads > Virtualization > Virtual Machines and click Create > With YAML.
5. In the spec > template > spec > volumes section, attach the cloned PVC instead of the container disk. Provide all other details for the new VM according to your requirements.

```
- name: rootdisk
  persistentVolumeClaim:
    claimName: rhel8-short-frog-rootdisk-28dvb-clone
```

6. Click Create to create the new VM.
7. After the VM is created successfully, access and verify that the new VM is a clone of the source VM.

Workflows: Red Hat OpenShift Virtualization with NetApp ONTAP

This section shows how to create a virtual machine from a Snapshot with Red Hat OpenShift Virtualization.

Create VM from a Snapshot

With Astra Trident and Red Hat OpenShift, users can take a snapshot of a persistent volume on Storage Classes provisioned by it. With this feature, users can take a point-in-time copy of a volume and use it to create a new volume or restore the same volume back to a previous state. This enables or supports a variety of use-cases, from rollback to clones to data restore.

For Snapshot operations in OpenShift, the resources VolumeSnapshotClass, VolumeSnapshot, and VolumeSnapshotContent must be defined.

- A VolumeSnapshotContent is the actual snapshot taken from a volume in the cluster. It is cluster-wide resource analogous to PersistentVolume for storage.
- A VolumeSnapshot is a request for creating the snapshot of a volume. It is analogous to a PersistentVolumeClaim.
- VolumeSnapshotClass lets the administrator specify different attributes for a VolumeSnapshot. It allows you to have different attributes for different snapshots taken from the same volume.

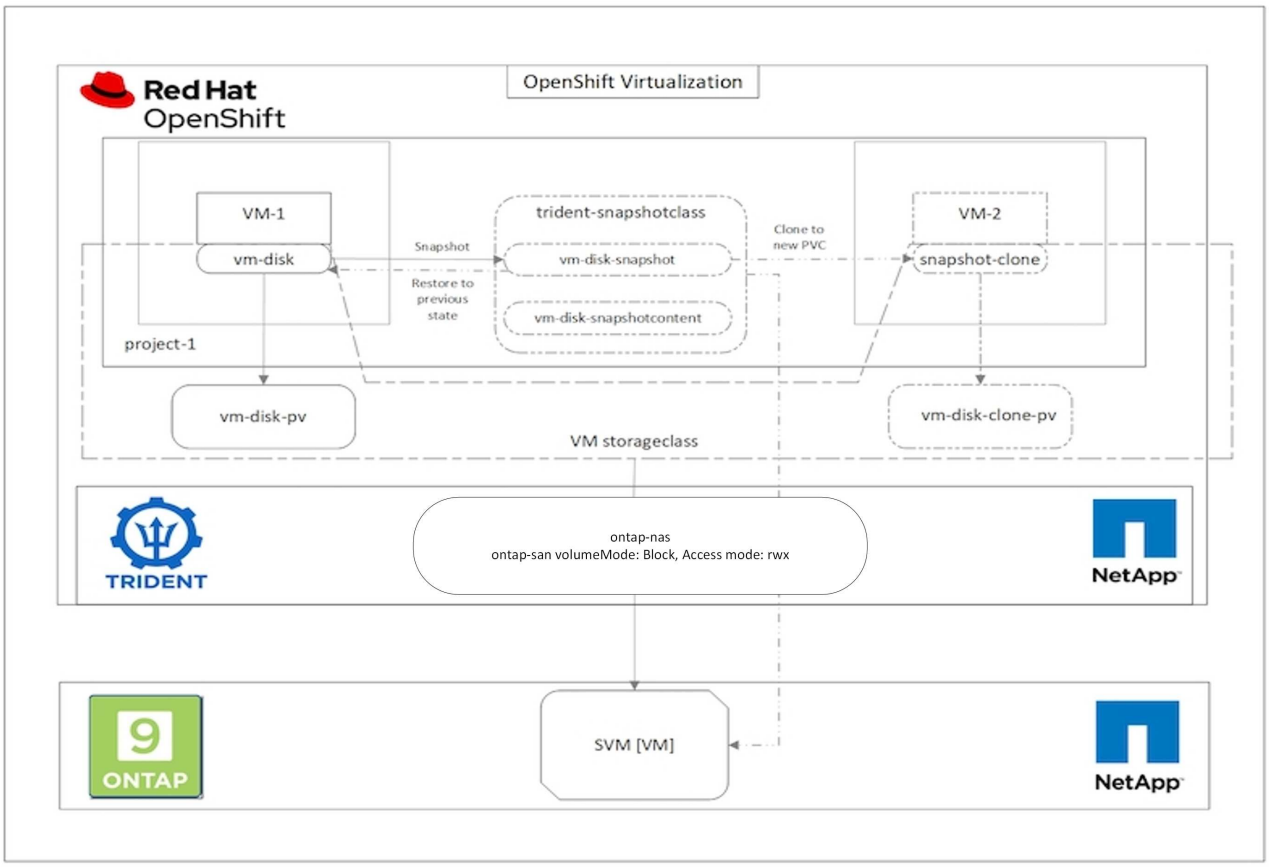

To create Snapshot of a VM, complete the following steps:

1. Create a VolumeSnapshotClass that can then be used to create a VolumeSnapshot. Navigate to Storage > VolumeSnapshotClasses and click Create VolumeSnapshotClass.
2. Enter the name of the Snapshot Class, enter csi.trident.netapp.io for the driver, and click Create.

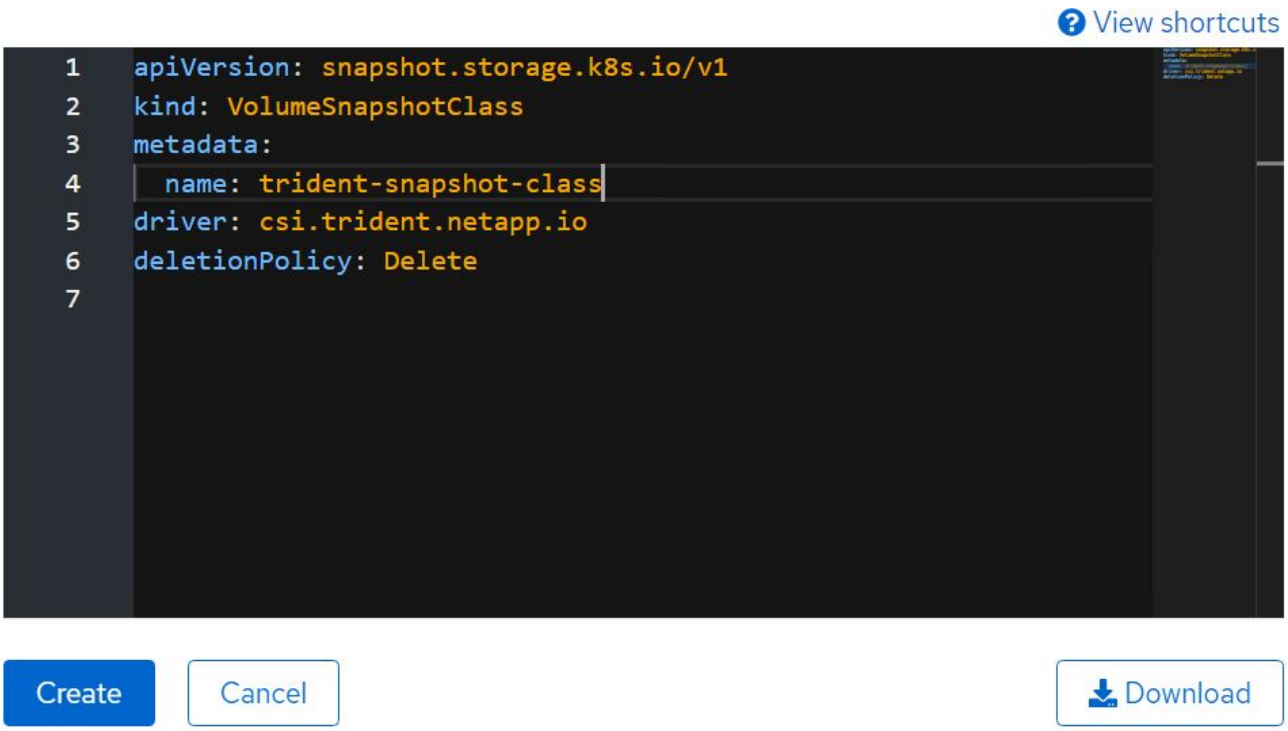

3. Identify the PVC that is attached to the source VM and then create a Snapshot of that PVC. Navigate to `Storage > VolumeSnapshots` and click Create VolumeSnapshots.
4. Select the PVC that you want to create the Snapshot for, enter the name of the Snapshot or accept the default, and select the appropriate VolumeSnapshotClass. Then click Create.

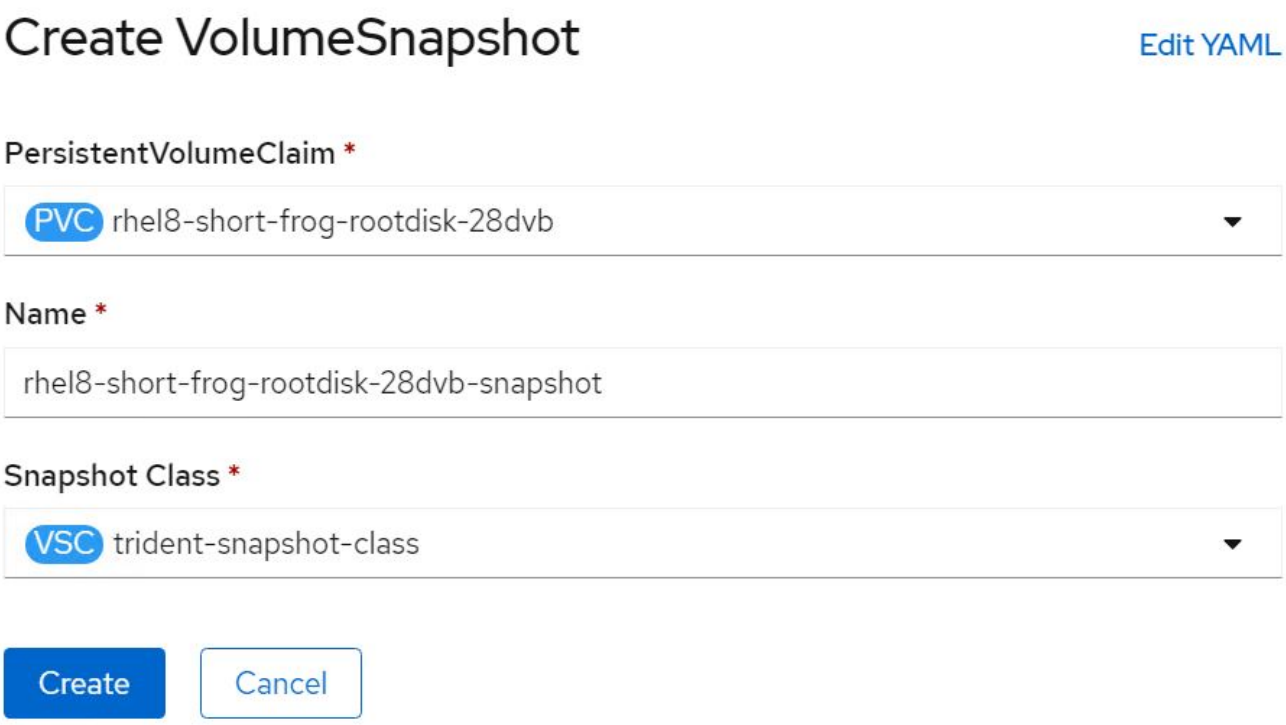

5. This creates the snapshot of the PVC at that point in time.

Create a new VM from the snapshot

1. First, restore the Snapshot into a new PVC. Navigate to Storage > VolumeSnapshots, click the ellipsis next to the Snapshot that you wish to restore, and click Restore as new PVC.
2. Enter the details of the new PVC and click Restore. This creates a new PVC.

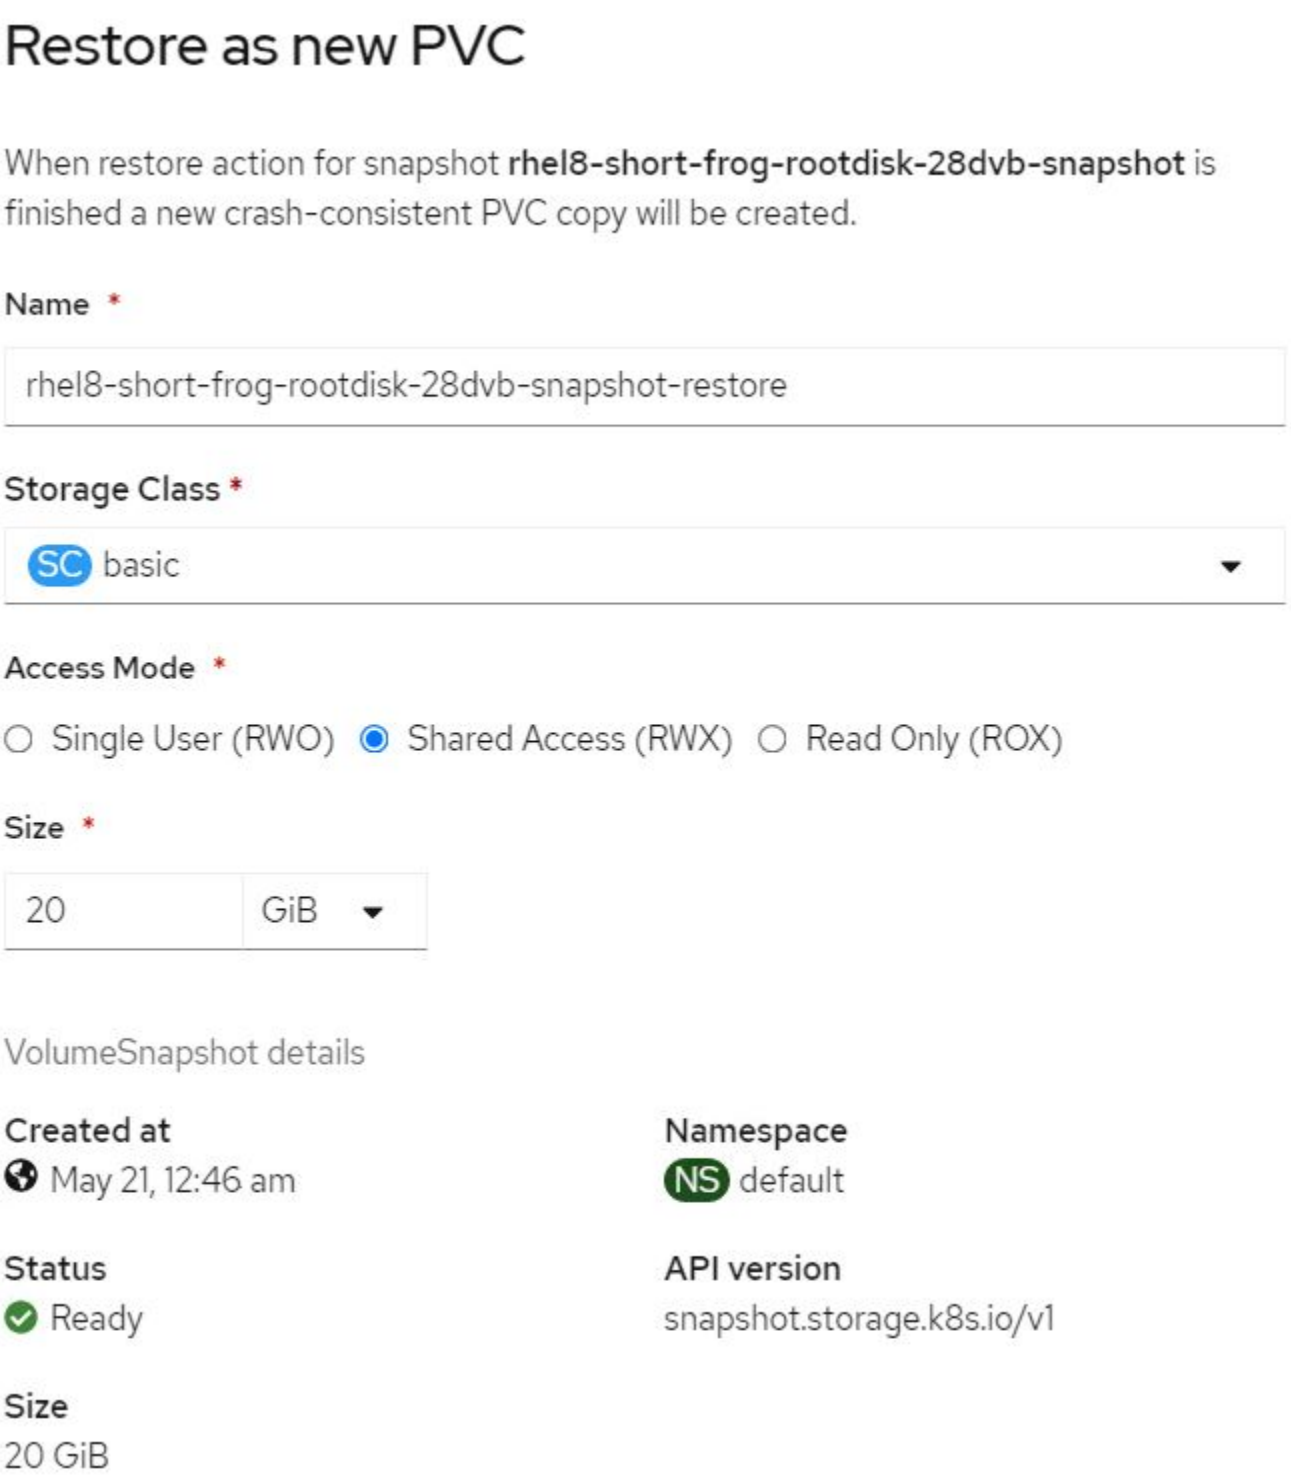

3. Next, create a new VM from this PVC. Navigate to Virtualization > Virtual Machines and click Create > With YAML.
4. In the spec > template > spec > volumes section, specify the new PVC created from Snapshot instead of

from the container disk. Provide all other details for the new VM according to your requirements.

```
- name: rootdisk
  persistentVolumeClaim:
    claimName: rhel8-short-frog-rootdisk-28dvb-snapshot-restore
```

5. Click Create to create the new VM.
6. After the VM is created successfully, access and verify that the new VM has the same state as that of the VM whose PVC was used to create the snapshot at the time when the snapshot was created.

Data Protection

Data protection for VMs in OpenShift Virtualization using OpenShift API for Data Protection (OADP)

Author: Banu Sundhar, NetApp

This section of the reference document provides details for creating backups of VMs using the OpenShift API for Data Protection (OADP) with Velero on NetApp ONTAP S3 or NetApp StorageGRID S3. The backups of Persistent Volumes(PVs) of the VM disks are created using CSI Astra Trident Snapshots.

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 Astra Trident CSI. In this section we use OpenShift API for Data Protection (OADP) to perform backup of VMs including its data volumes to

- ONTAP Object Storage
- StorageGrid

We then restore from the backup when needed.

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.

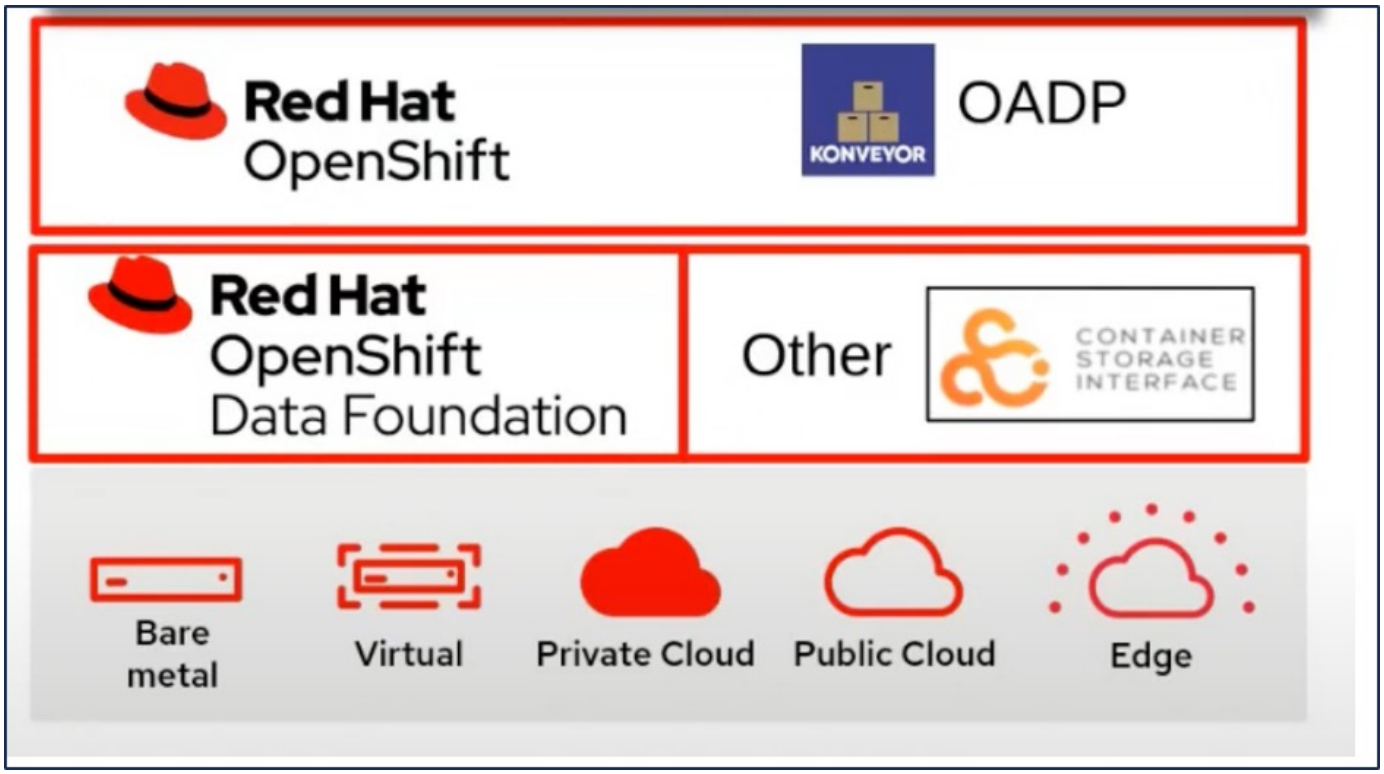

Red Hat OpenShift has leveraged the solutions developed by the OpenSource communities for data protection. Velero is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes. To use Velero easily, OpenShift has developed 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.

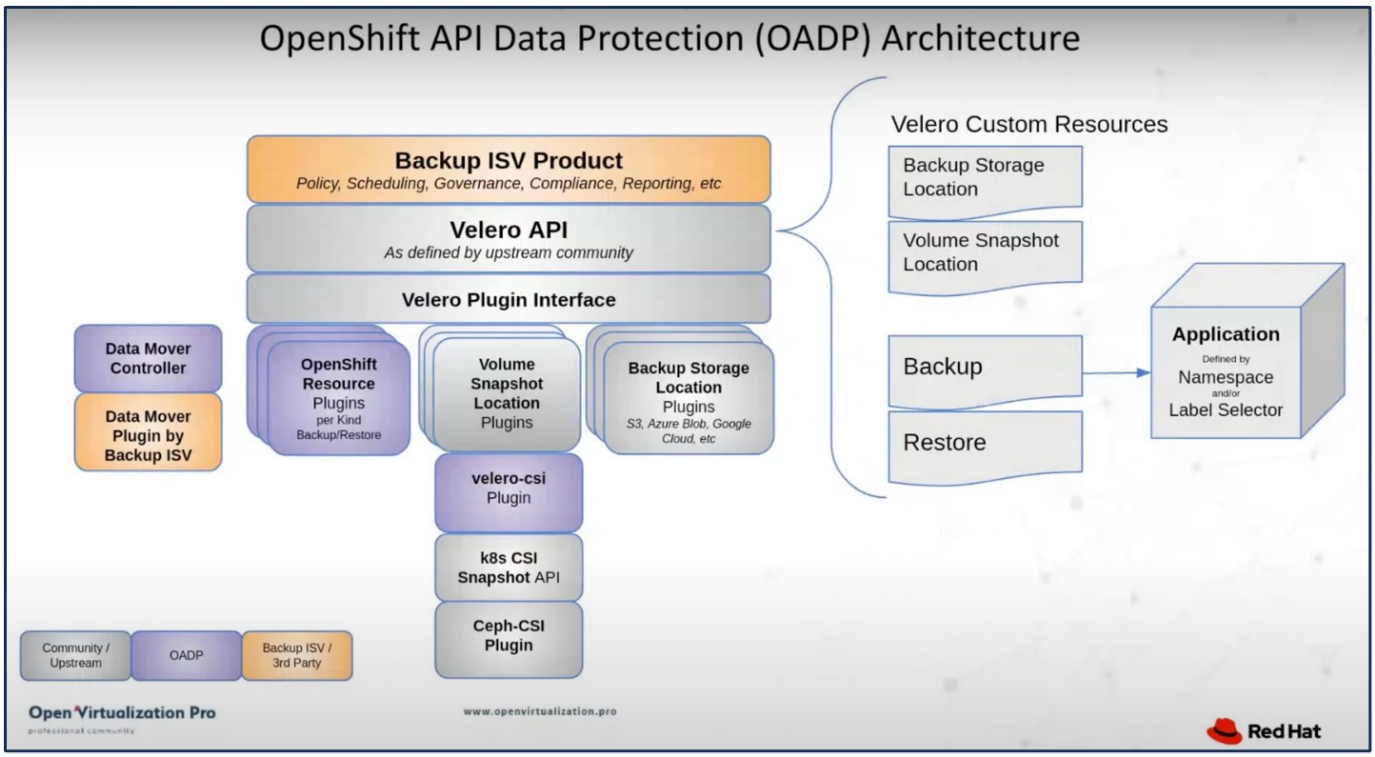

OADP 1.3 is available from the operator hub of OpenShift cluster 4.12 and later. It has a built-in Data Mover that can move CSI volume snapshots to a remote object store. This provides portability and durability by

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
- OpenShift Virtualization installed via OperatorOpenShift Virtualization Operator provided by Red Hat
- OADP Operator 1.13 provided by Red Hat
- Velero CLI 1.13 for Linux
- Astra Trident 24.02
- ONTAP 9.12

Astra 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 Astra Trident
- A Trident backend configured with an SVM on ONTAP cluster
- A StorageClass configured on the OpenShift cluster with Astra 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
- OpenShift Virtualization operator installed and configured
- VMs deployed in a Namespace on OpenShift Virtualization
- An admin workstation with tridentctl and oc tools installed and added to $PATH

> If you want to take a backup of a VM when it is in the Running state, then you must install the QEMU guest agent on that virtual machine. If you install the VM using an existing template, then QEMU agent is installed automatically. QEMU allows the guest agent to quiesce in-flight data in the guest OS during the snapshot process, and avoid possible data corruption. If you do not have QEMU installed, you can stop the virtual machine before taking a backup.

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.

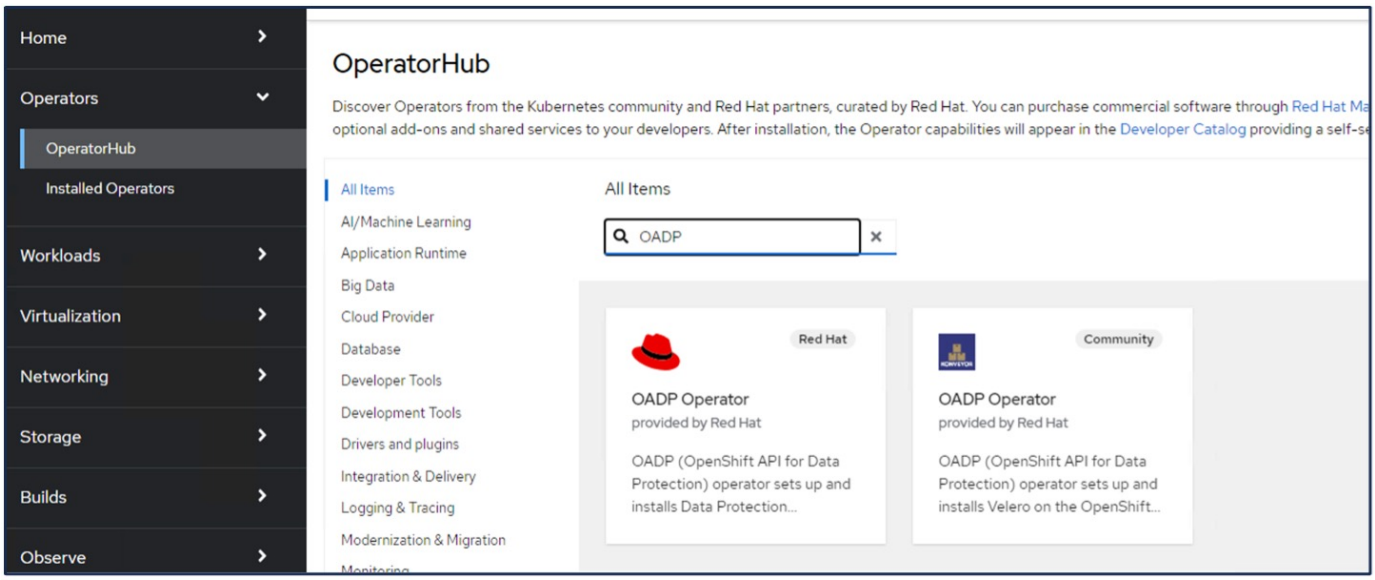
Home
Operators
OperatorHub
Installed Operators
Workloads
Virtualization
Networking
Storage
Builds
Observe
OperatorHub
Discover Operators from the Kubernetes community and Red Hat partners, curated by Red Hat. You can purchase commercial software through Red Hat Ma
optional add-ons and shared services 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...

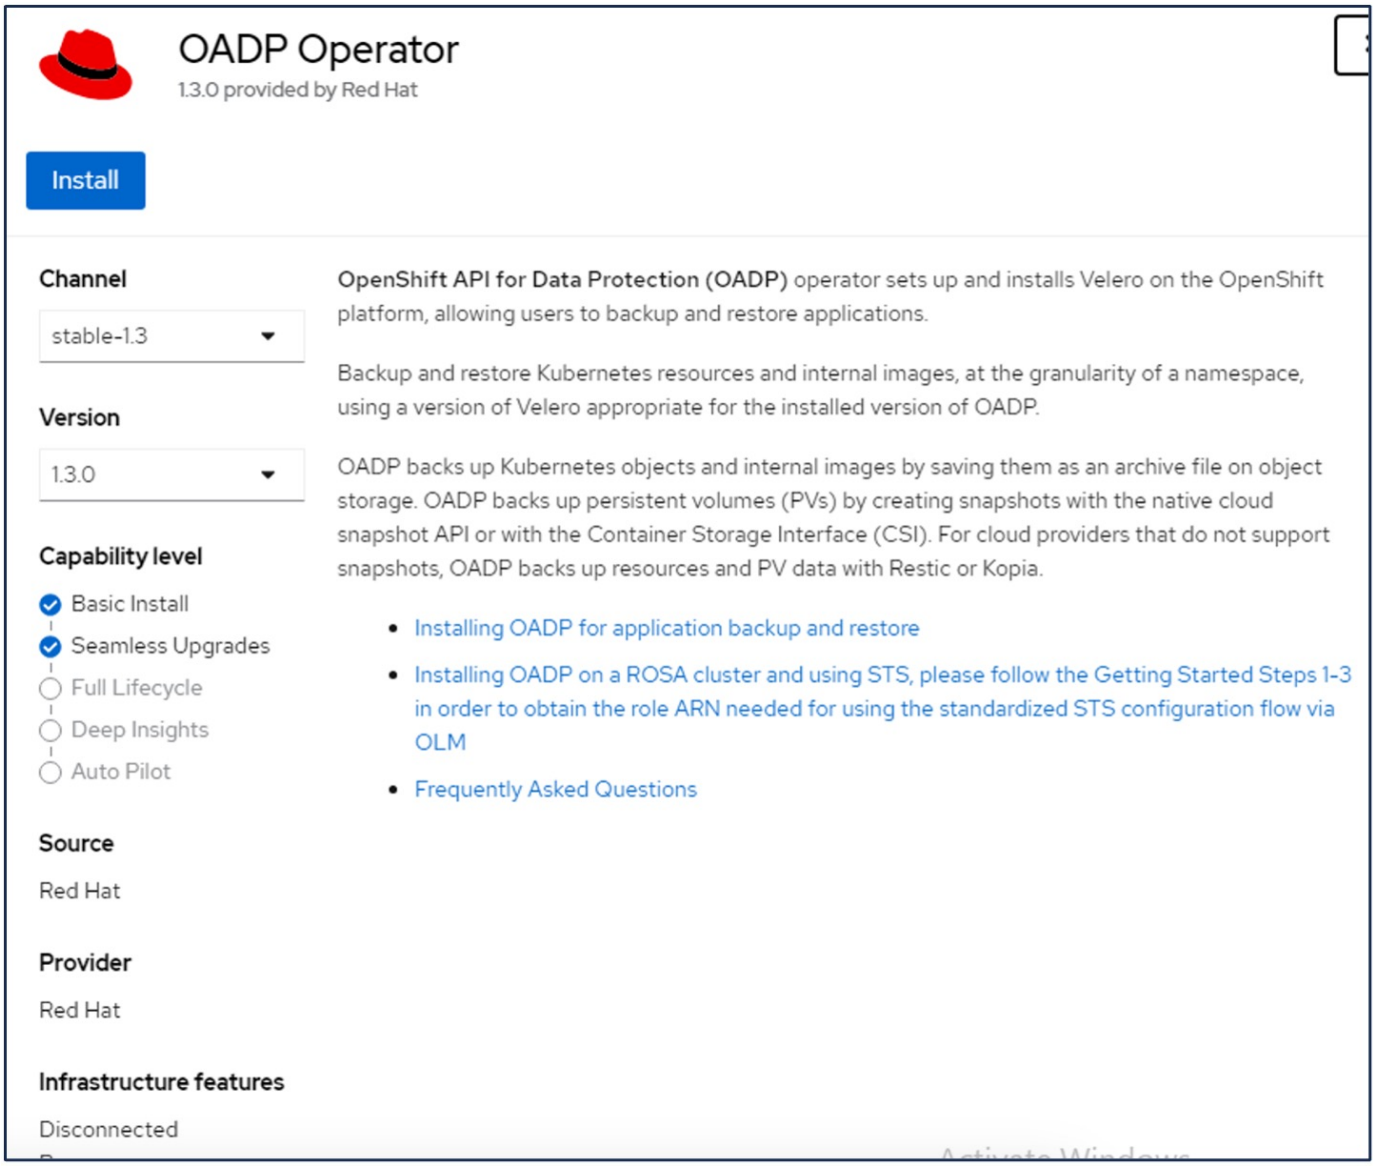
OADP Operator
1.3.0 provided by Red Hat
Install
Channel
stable-1.3
Version
1.3.0
Capability level
Basic Install
Seamless Upgrades
Full Lifecycle
Deep Insights
Auto Pilot
Source
Red Hat
Provider
Red Hat
Infrastructure features
Disconnected
OpenShift API for Data Protection (OADP) operator sets up and 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

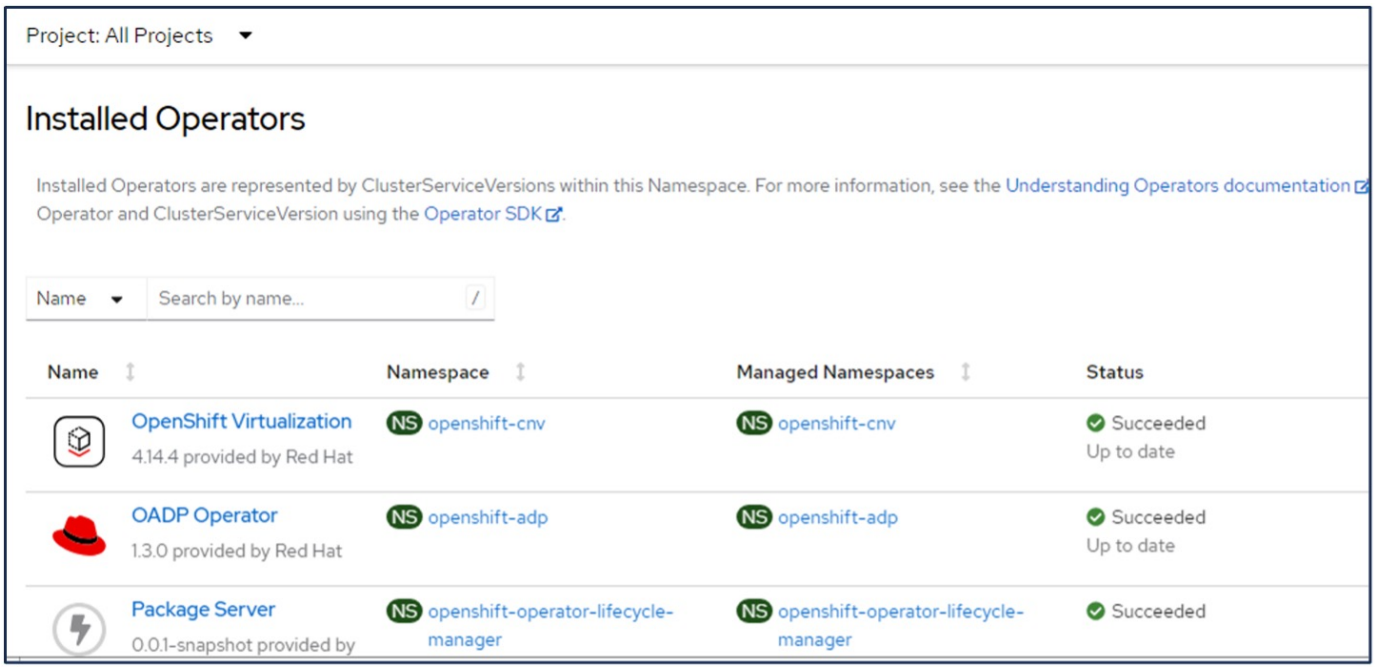

Prerequisites for Velero configuration with Ontap S3 details

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.

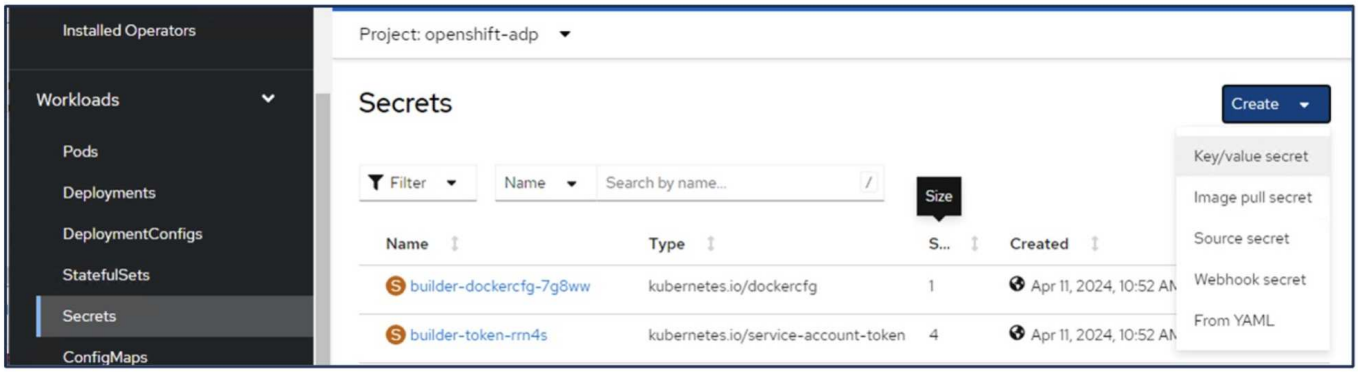

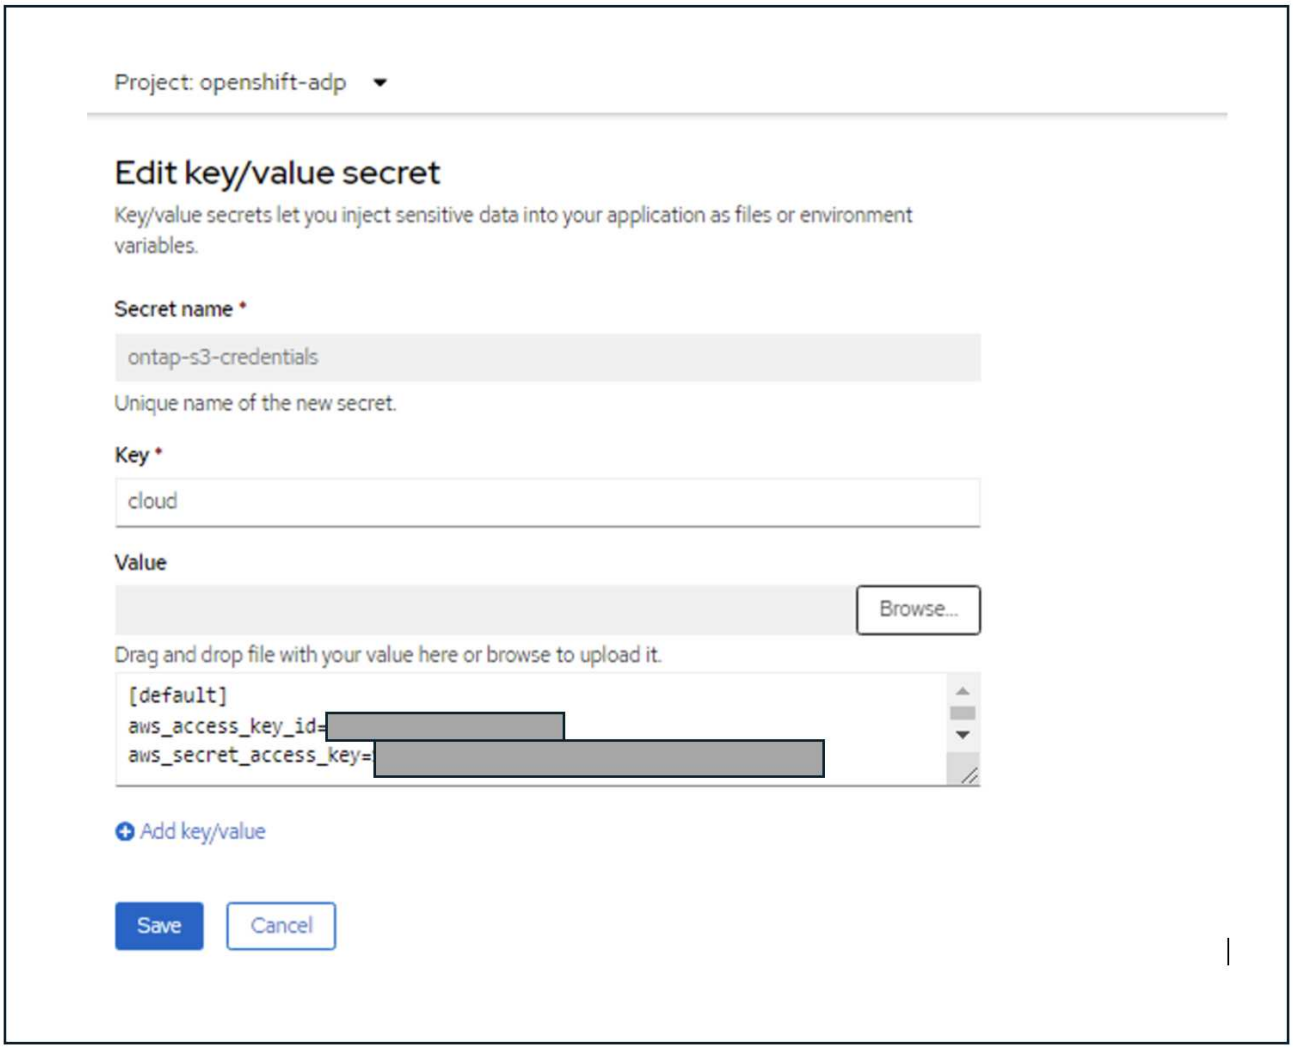

To create a secret named sg-s3-credentials from the CLI you can use the following command.

```
# 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.

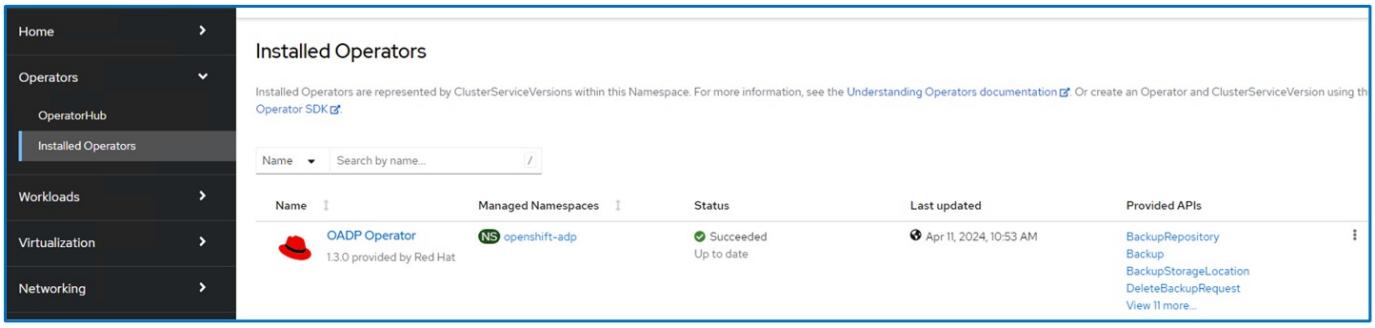

Click on Create DataProtectionApplication. In the form view, provide a name for the DataProtection Application or use the default name.

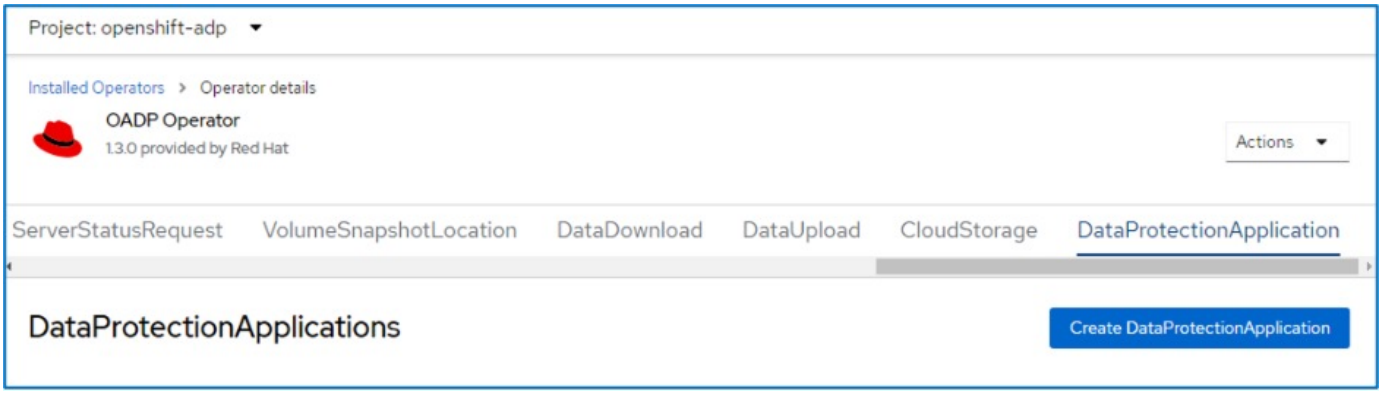

Now go to the YAML view and replace the spec information as shown in the yaml file examples below.

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.xx.xx.xx' ->LIF to access S3. 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: demobackup ->The folder that will be created in the
bucket
        provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
      #default Data Mover uses Kopia to move snapshots to Object Storage
    velero:
      defaultPlugins:
        - csi ->Add this plugin
        - openshift
        - aws
        - kubevirt ->Add this plugin
```

Sample yaml file for configuring Velero with StorageGrid S3 as the backupLocation and snapshotLocation

```
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 Astra 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.

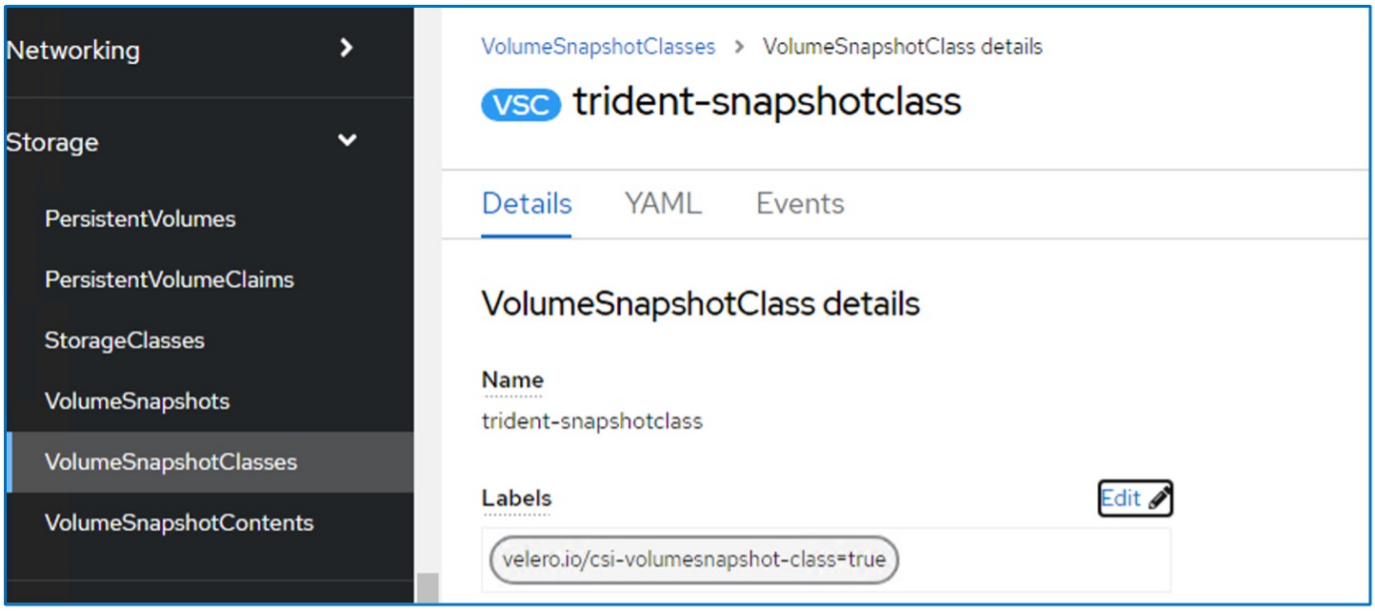

Ensure that the snapshots can persist even if the VolumeSnapshot objects are deleted. This can be done by setting the **deletionPolicy** to Retain. If not, deleting a namespace will completely lose all PVCs ever backed up in it.

```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: trident-snapshotclass
driver: csi.trident.netapp.io
deletionPolicy: Retain
```

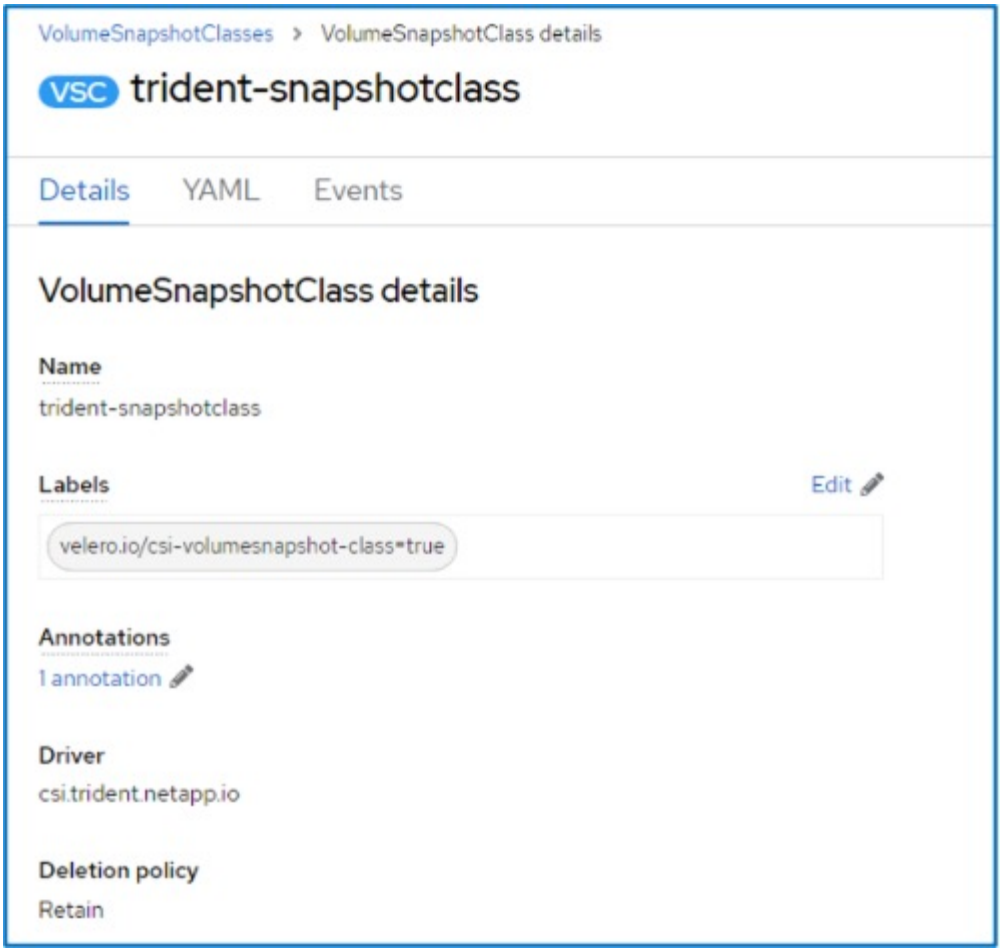

Ensure that the DataProtectionApplication is created and is in condition:Reconciled.

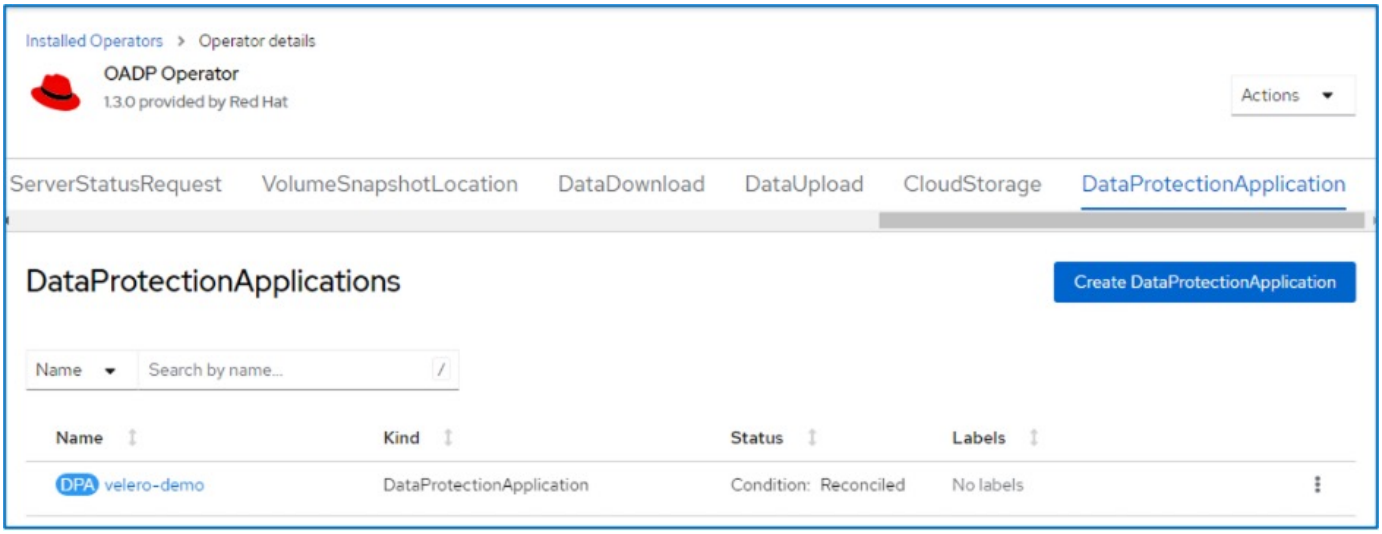

The OADP operator will create a corresponding BackupStorageLocation.This will be used when creating a backup.

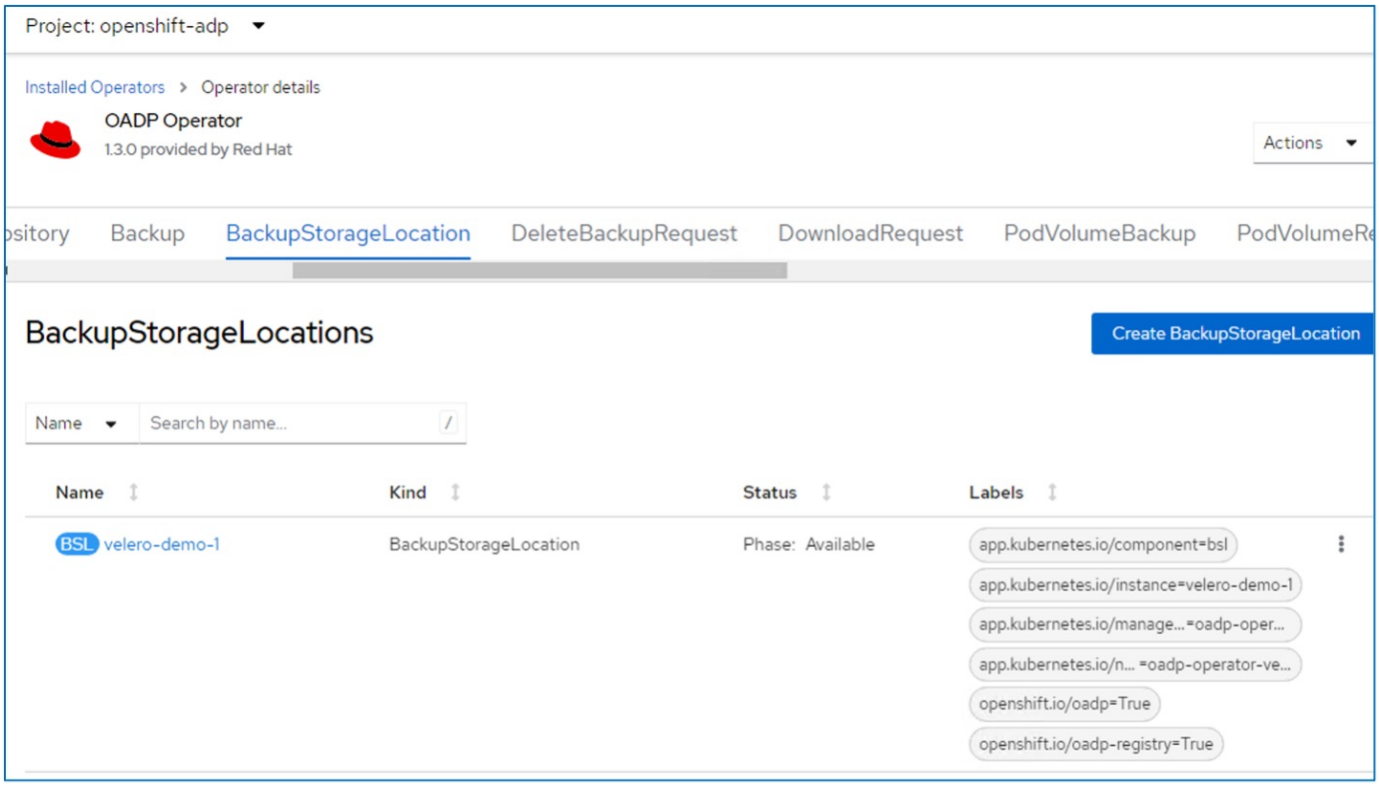

Creating on-demand backup for VMs in OpenShift Virtualization

This section outlines how to create on-demand backup for VMs in OpenShift Virtualization.

Steps to create a backup of a VM

To create an on-demand backup of the entire VM (VM metadata and VM disks), click on the **Backup** tab. This creates a Backup Custom Resource (CR). A sample yaml is provided to create the Backup CR. Using this yaml, the VM and its disks in the specified namespace will be backed up. Additional parameters can be set as shown in the documentation.

A snapshot of the persistent volumes backing the disks will be created by the CSI. A backup of the VM along with the snapshot of its disks are created and stored in the backup location specified in the yaml. The backup will remain in the system for 30 days as specified in the ttl.

```
apiVersion: velero.io/v1
kind: Backup
metadata:
  name: backup1
  namespace: openshift-adp
spec:
  includedNamespaces:
  - virtual-machines-demo
  snapshotVolumes: true
  storageLocation: velero-demo-1 -->this is the backupStorageLocation
previously created
                                    when Velero is configured.
  ttl: 720h0m0s
```

Once the backup completes, its Phase will show as completed.

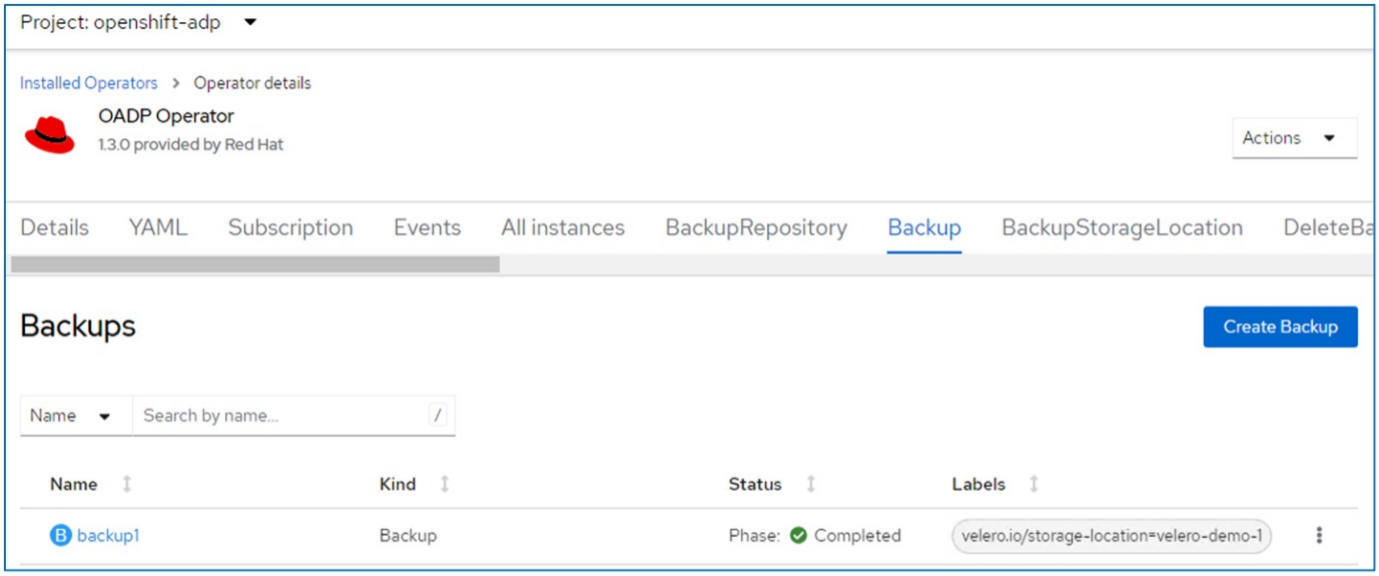

You can inspect the backup in the Object storage with the help of an S3 browser application. The path of the backup shows in the configured bucket with the prefix name (velero/demobackup). You can see the contents of the backup includes the volume snapshots, logs, and other metadata of the virtual machine.

In StorageGrid, you can also use the S3 console that is available from the Tenant Manager to view the backup objects.

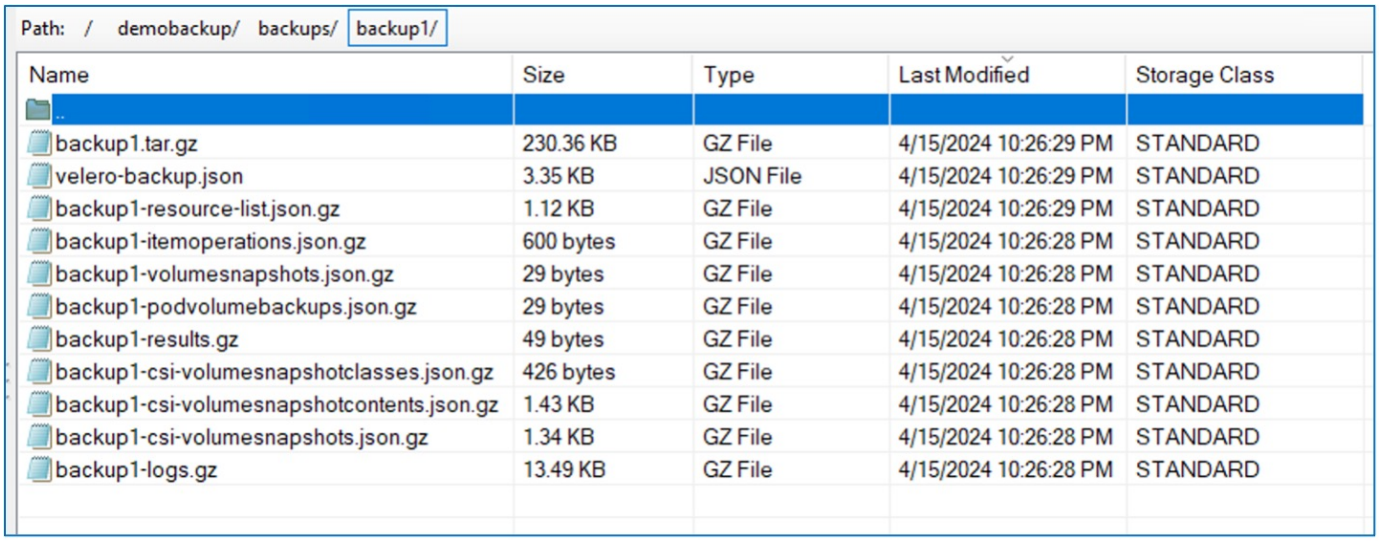

Creating scheduled backups for VMs in OpenShift Virtualization

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.

```
apiVersion: velero.io/v1
kind: Schedule
metadata:
  name: <schedule>
  namespace: openshift-adp
spec:
  schedule: 0 7 * * *
  template:
    hooks: {}
    includedNamespaces:
    - <namespace>
    storageLocation: velero-demo-1
    defaultVolumesToFsBackup: true
    ttl: 720h0m0s
```

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.

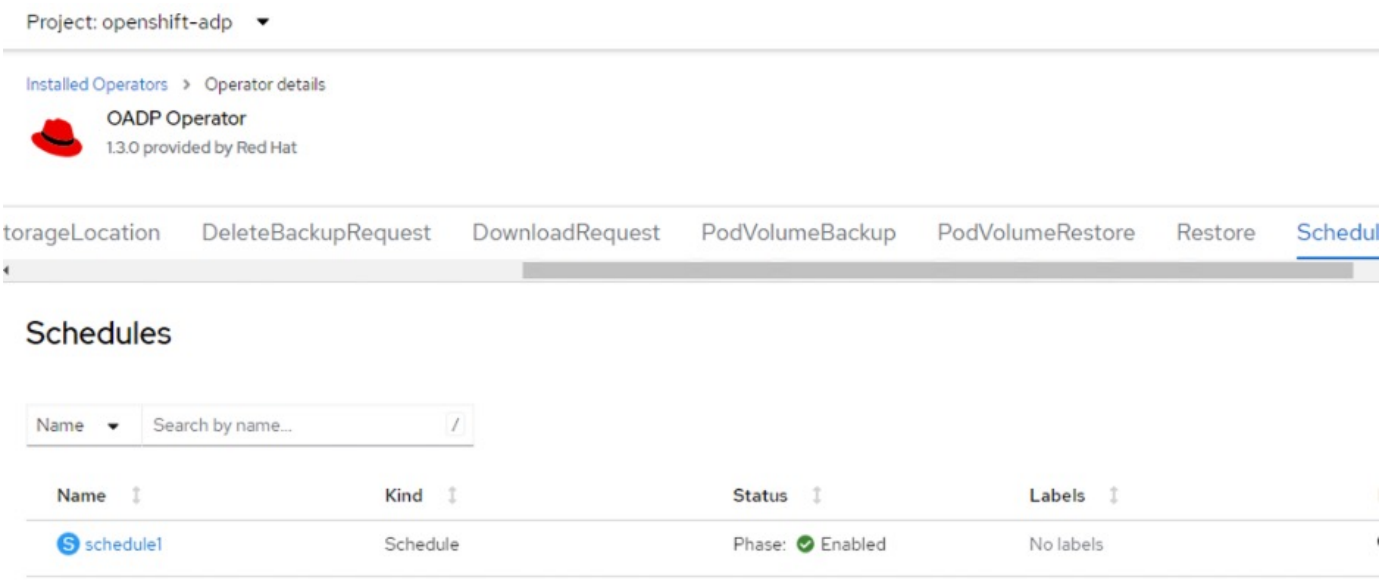

Backups will be created according to this schedule, and can be viewed from the Backup tab.

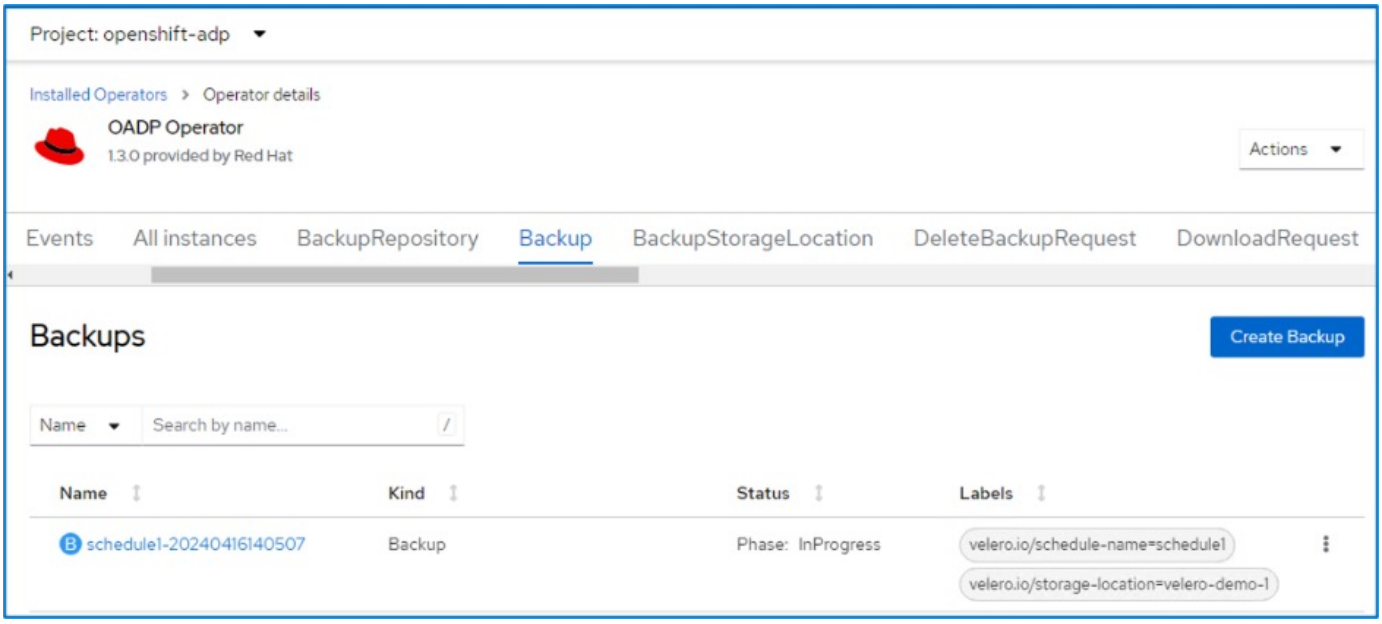

Restore a VM from a backup

This section describes how to restore virtual machine(s) from a backup.

Prerequisites

To restore from a backup, let us assume that the namespace where the virtual machine existed got accidentally deleted.

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.

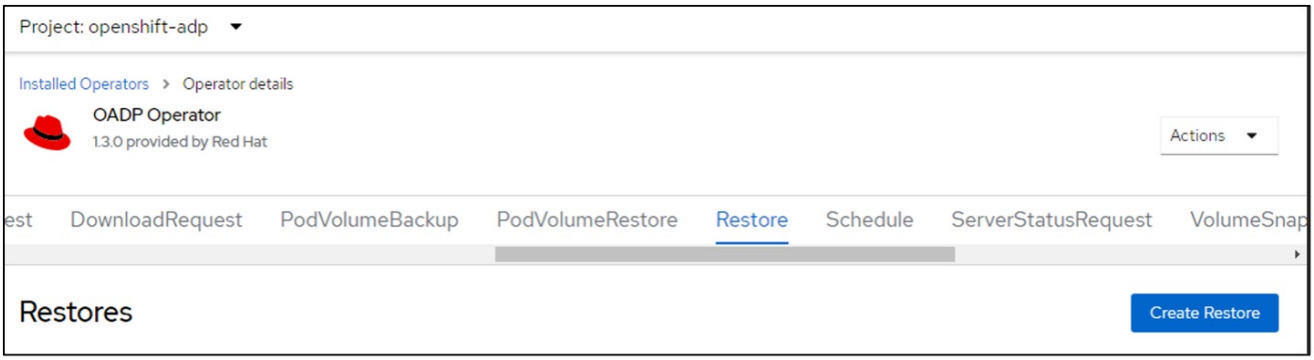


```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: restore1
  namespace: openshift-adp
spec:
  backupName: backup1
  restorePVs: true
```

When the phase shows completed, you can see that the virtual machines have been restored to the state when the snapshot was taken. (If the backup was created when the VM was running, restoring the VM from the backup will start the restored VM and bring it to a running state). The VM is restored to the same namespace.

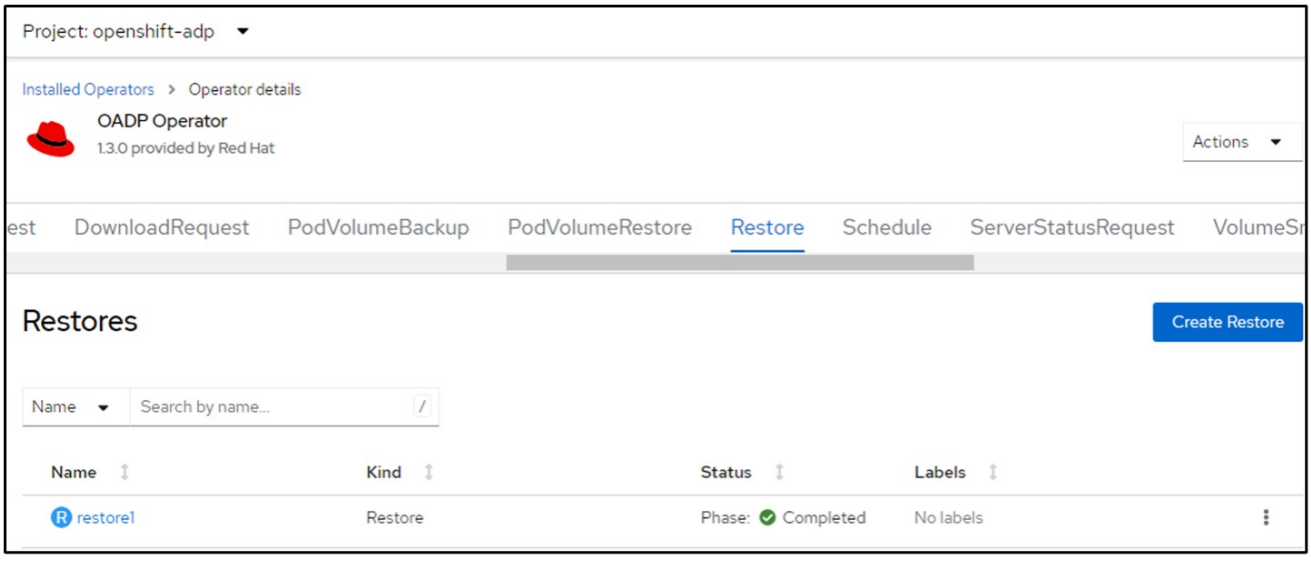

Restore to a different namespace

To restore the VM 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 a VM and its disks in the virtual-machines-demo namespace when the backup was taken to the virtual-machines namespace.

```
apiVersion: velero.io/v1
kind: Restore
metadata:
  name: restore-to-different-ns
  namespace: openshift-adp
spec:
  backupName: backup
  restorePVs: true
  includedNamespaces:
  - virtual-machines-demo
  namespaceMapping:
    virtual-machines-demo: virtual-machines
```

When the phase shows completed, you can see that the virtual machines have been restored to the state when the snapshot was taken. (If the backup was created when the VM was running, restoring the VM from the backup will start the restored VM and bring it to a running state). The VM is restored to a different namespace as specified in the yaml.

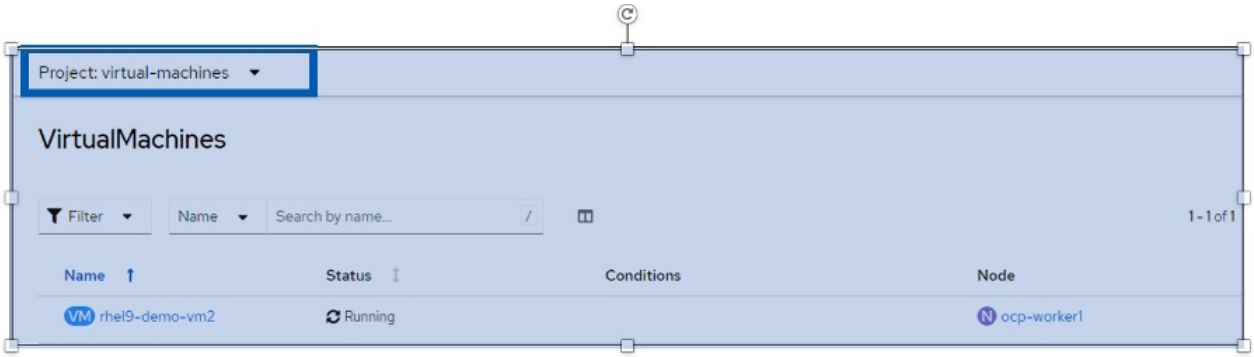

Restore to a different storage class

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 virtual machine, during creation uses ontap-nas as the storage class for its disks. A backup of the virtual machine named backup1 is created.

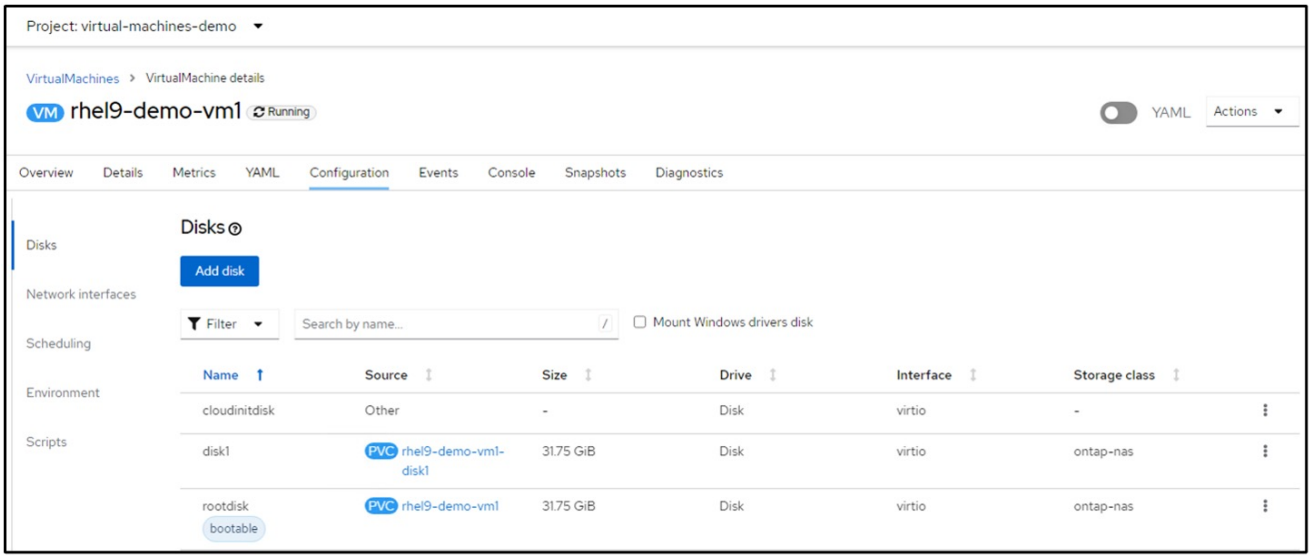

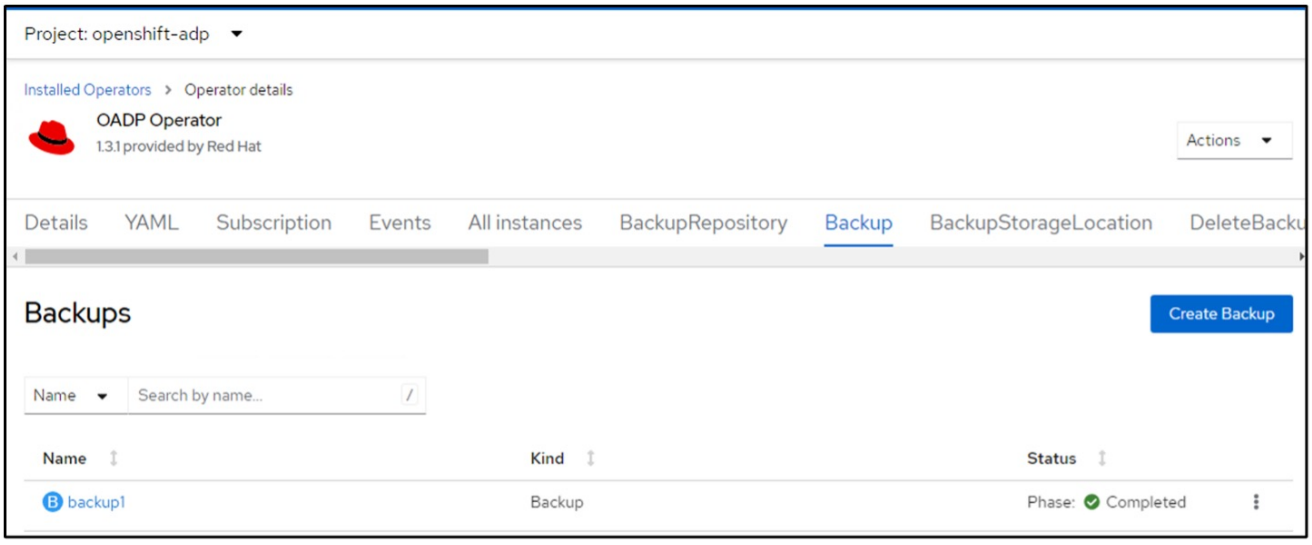

Simulate a loss of the VM by deleting the VM.

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-storage-class-config (can be any name)

Key: change-storage-class-config.yaml:
Value:

```
version: v1
    resourceModifierRules:
    - conditions:
         groupResource: persistentvolumeclaims
         resourceNameRegex: "^rhel*"
         namespaces:
         - virtual-machines-demo
      patches:
      - operation: replace
        path: "/spec/storageClassName"
        value: "ontap-nas-eco"
```

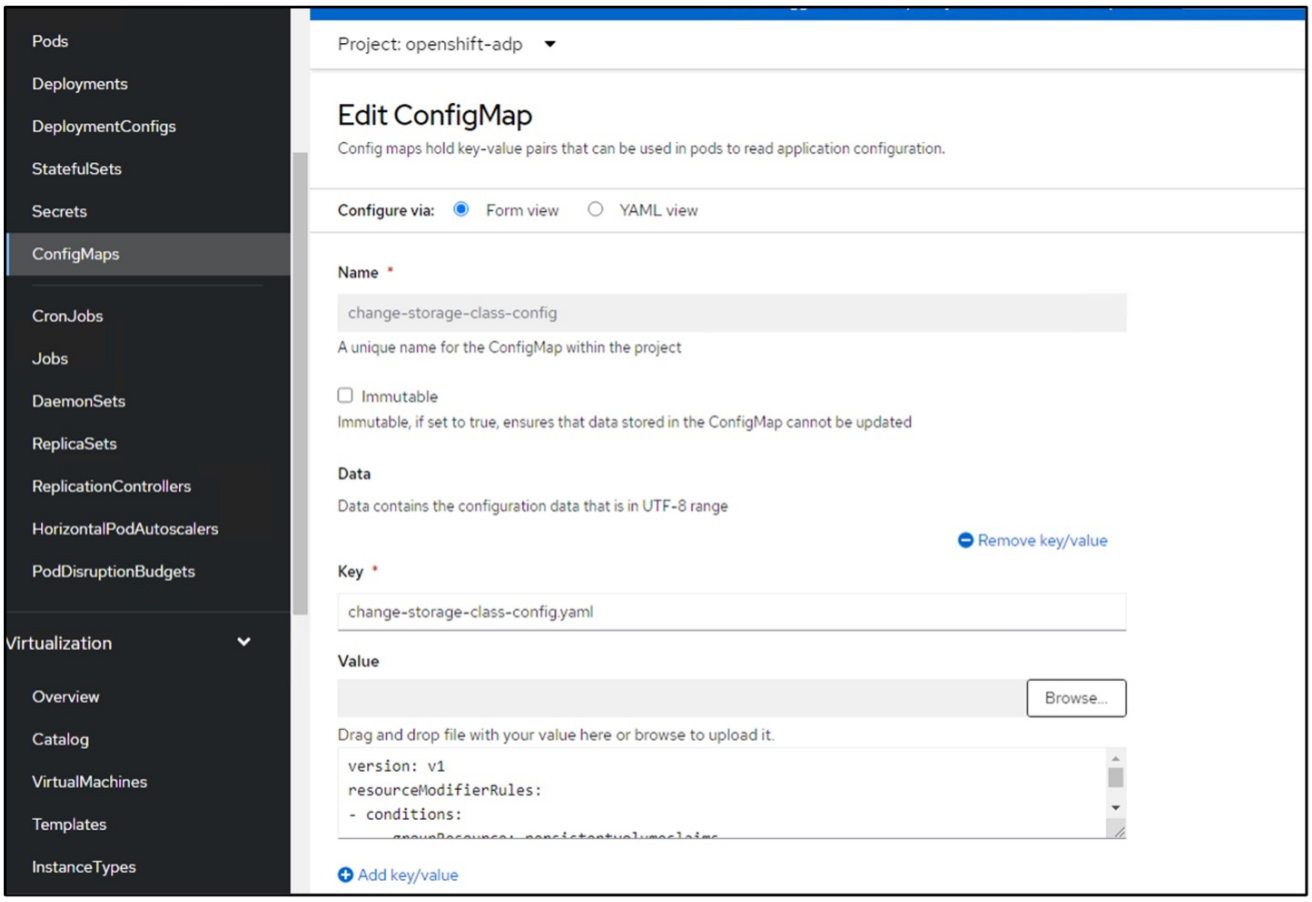

The resulting config map object should look like this (CLI):

```
# 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 VM is restored in the same namespace with the disks created using the storage class ontap-nas-eco.

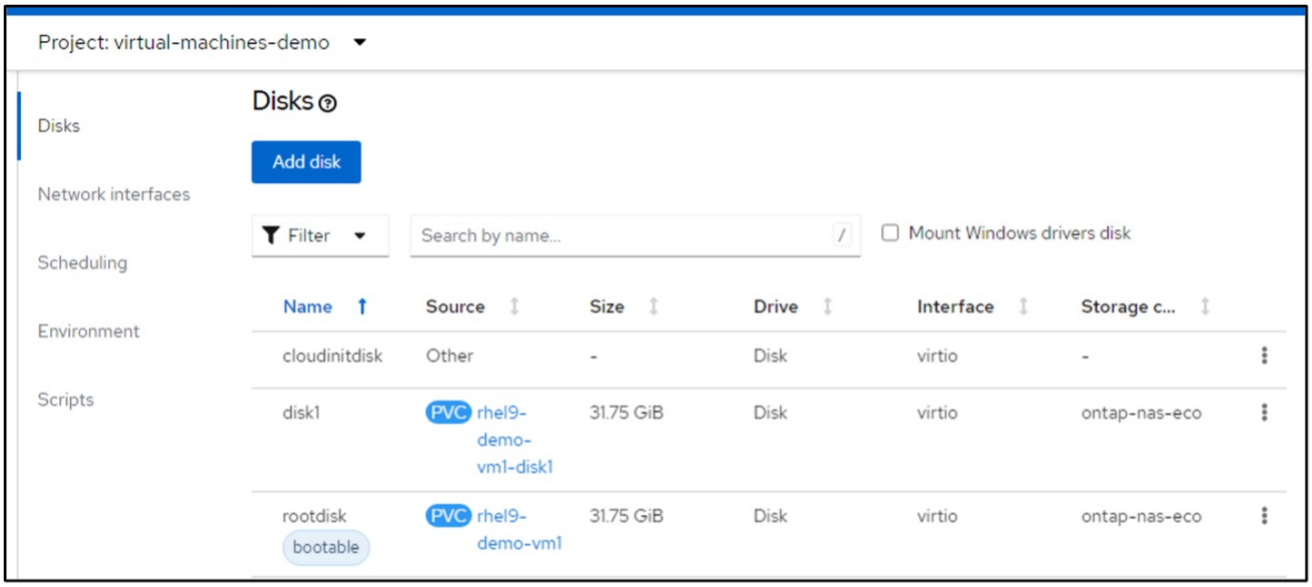

Deleting backups and restores in using Velero

This section outlines how to delete backups and restores for VMs in OpenShift Virtualization using Velero.

Deleting a backup

You can delete a Backup CR without deleting the Object Storage data by using the OC CLI tool.

```
oc delete backup <backup_CR_name> -n <velero_namespace>
```

If you want the delete the Backup CR and delete the associated object storage data, you can do so by using the Velero CLI tool.

Download the CLI as given in the instructions in the Velero documentation.

Execute the following delete command using the Velero CLI

```
velero backup delete <backup_CR_name> -n <velero_namespace>
```

Deleting a Restore

You can delete the Restore CR using the Velero CLI

```
velero restore delete restore --namespace openshift-adp
```

You can use oc command as well as the UI to delete the restore CR

```
oc delete backup <backup_CR_name> -n <velero_namespace>
```

Monitoring

Monitoring using Cloud Insights for VMs in Red Hat OpenShift Virtualization

Author: Banu Sundhar, NetApp

This section of the reference document provides details for integrating NetApp Cloud Insights with a Red Hat OpenShift Cluster to monitor OpenShift Virtualization VMs.

NetApp Cloud Insights is a cloud infrastructure monitoring tool that gives you visibility into your complete infrastructure. With Cloud Insights, you can monitor, troubleshoot, and optimize all your resources including your public clouds and your private data centers. For more information about NetApp Cloud Insights, refer to the Cloud Insights documentation.

To start using Cloud Insights, you must sign up on the NetApp BlueXP portal. For details, refer to the Cloud Insights Onboarding

Cloud Insights has several features that enable you to quickly and easily find data, troubleshoot issues, and provide insights into your environment. You can find data easily with powerful queries, you can visualize data in dashboards, and send email alerts for data thresholds you set. Refer to the video tutorials to help you understand these features.

For Cloud Insights to start collecting data you need the following

Data Collectors
There are 3 types of Data Collectors:
* Infrastructure (storage devices, network switches, compute infrastructure)
* Operating Systems (such as VMware or Windows)
* Services (such as Kafka)

Data Collectors discover information from the data sources, such as ONTAP storage device (infrastructure data collector). The information gathered is used for analysis, validation, monitoring, and troubleshooting.

Acquisition Unit
If you are using an infrastructure Data Collector, you also need an Acquisition Unit to inject data into Cloud Insights. An Acquisition Unit is a computer dedicated to hosting data collectors, typically a Virtual Machine. This computer is typically located in the same data center/VPC as the monitored items.

Telegraf Agents
Cloud Insights also supports Telegraf as its agent for collection of integration data. Telegraf is a plugin-driven server agent that can be used to collect and report metrics, events, and logs.

Cloud Insights Architecture

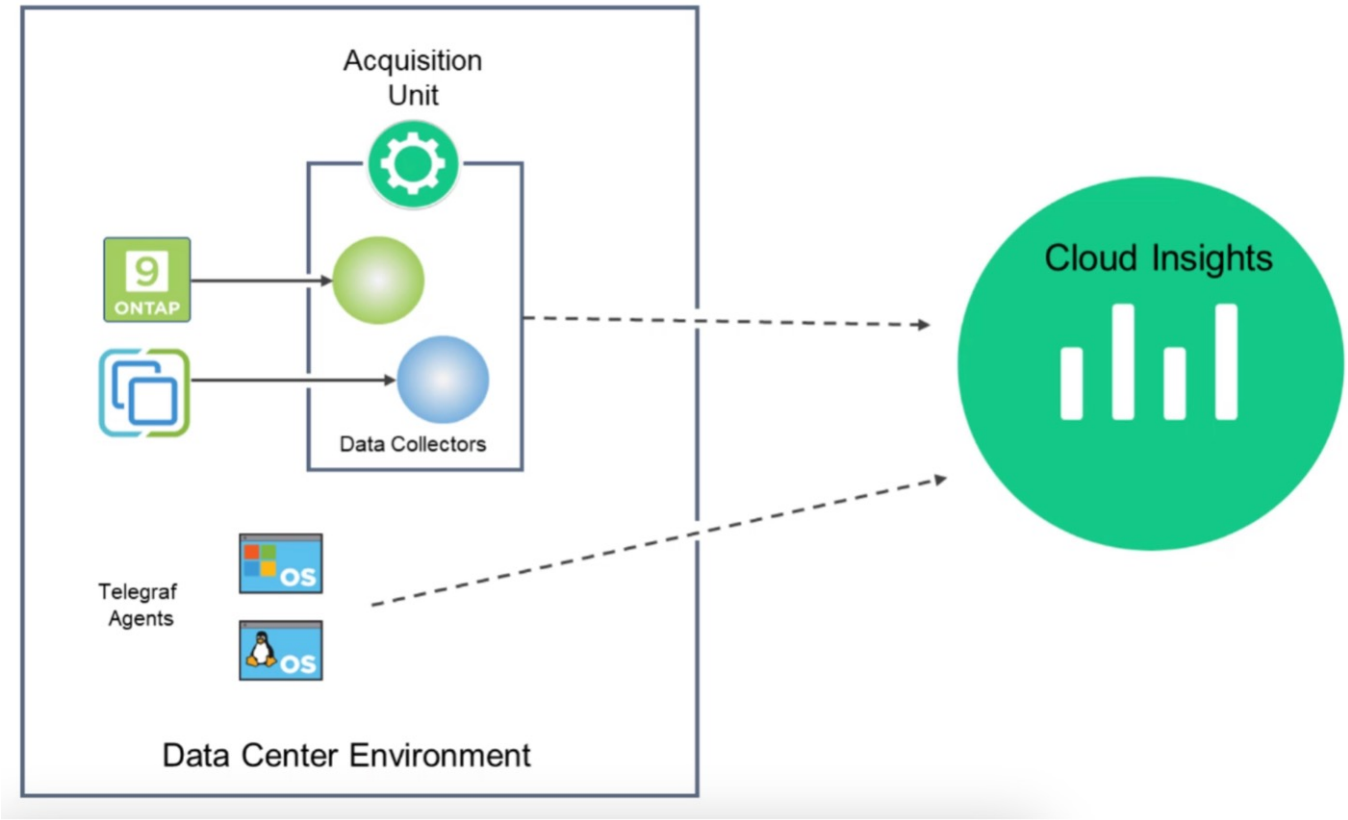

Integration with Cloud Insights for VMs in Red Hat OpenShift Virtualization

To start collecting data for VMs in OpenShift Virtualization you will need to install:

1. A Kubernetes monitoring operator and data collector to collect Kubernetes data
   For complete instructions, refer to the documentation.
2. An acquisition unit to collect data from ONTAP storage that provides persistent storage for the VM disks
   For complete instructions, refer to the documentation.
3. A data collector for ONTAP
   For complete instructions, refer to the documentation

Additionally, if you are using StorageGrid for VM backups, you need a data collector for the StorageGRID as well.

Sample Monitoring capabilities for VMs in Red Hat OpenShift Virtualization

This section discusses monitoring using Cloud Insights for VMs in Red Hat OpenShift Virtualization.

Monitoring based on events and creating Alerts

Here is a sample where the namespace that contains a VM in OpenShift Virtualization is monitored based on events. In this example, a monitor is created based on **logs.kubernetes**.event for the specified namespace in the cluster.

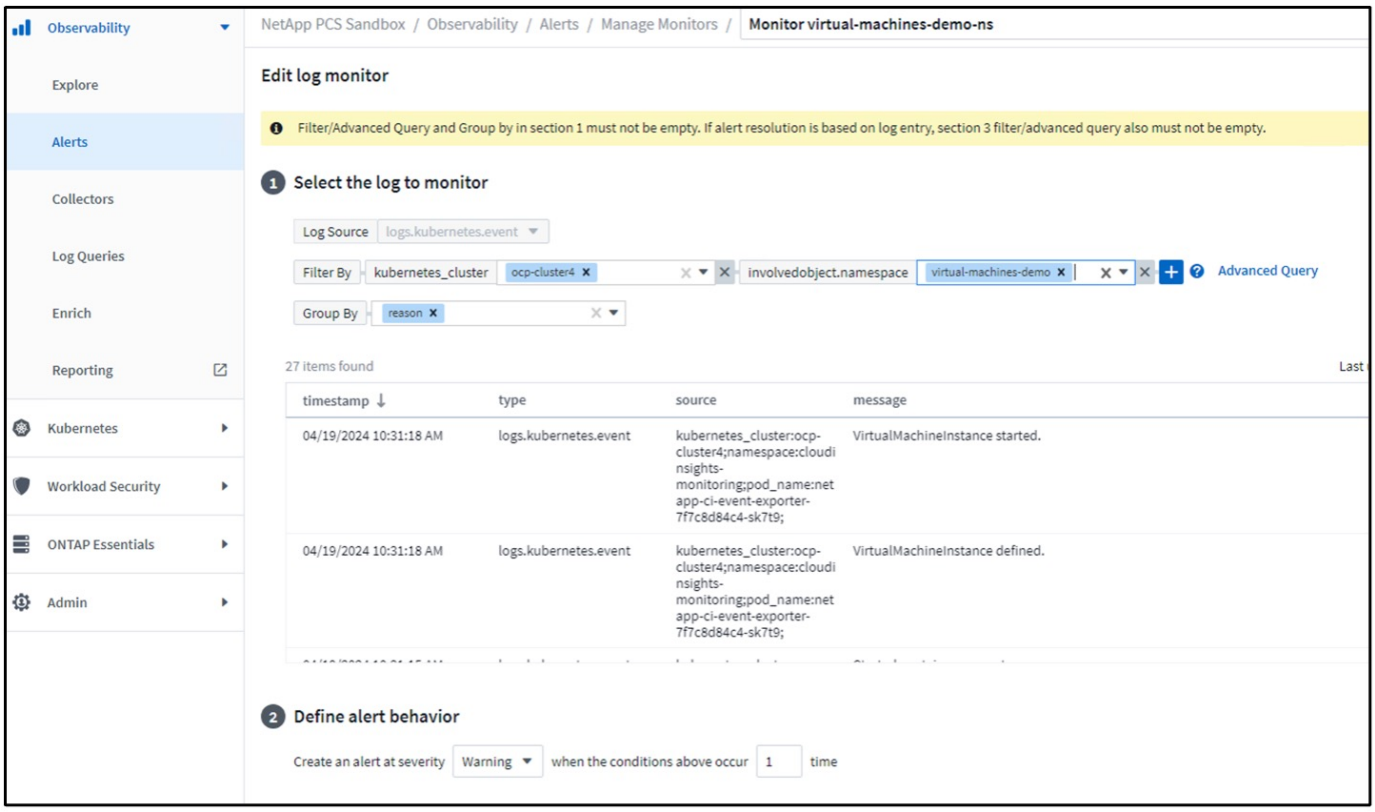

This query provides all the events for the virtual machine in the namespace. (There is only one virtual machine in the namespace). An advanced query can also be constructed to filter based on the event where the reason is “failed” or “FailedMount” These events are typically created when there is an issue in creating a PV or mounting the PV to a pod indicating issues in the dynamic provisioner for creating persistent volumes for the VM.
While creating the Alert Monitor as shown above, you can also configure notification to recipients. You can also provide corrective actions or additional information that can be useful to resolve the error. In the above example, additional information could be to look into the Trident backend configuration and storage class definitions for resolving the issue.

Change Analytics

With Change Analytics, you can get a view of what changed in the state of your cluster including who made that change which can help in troubleshooting issues.

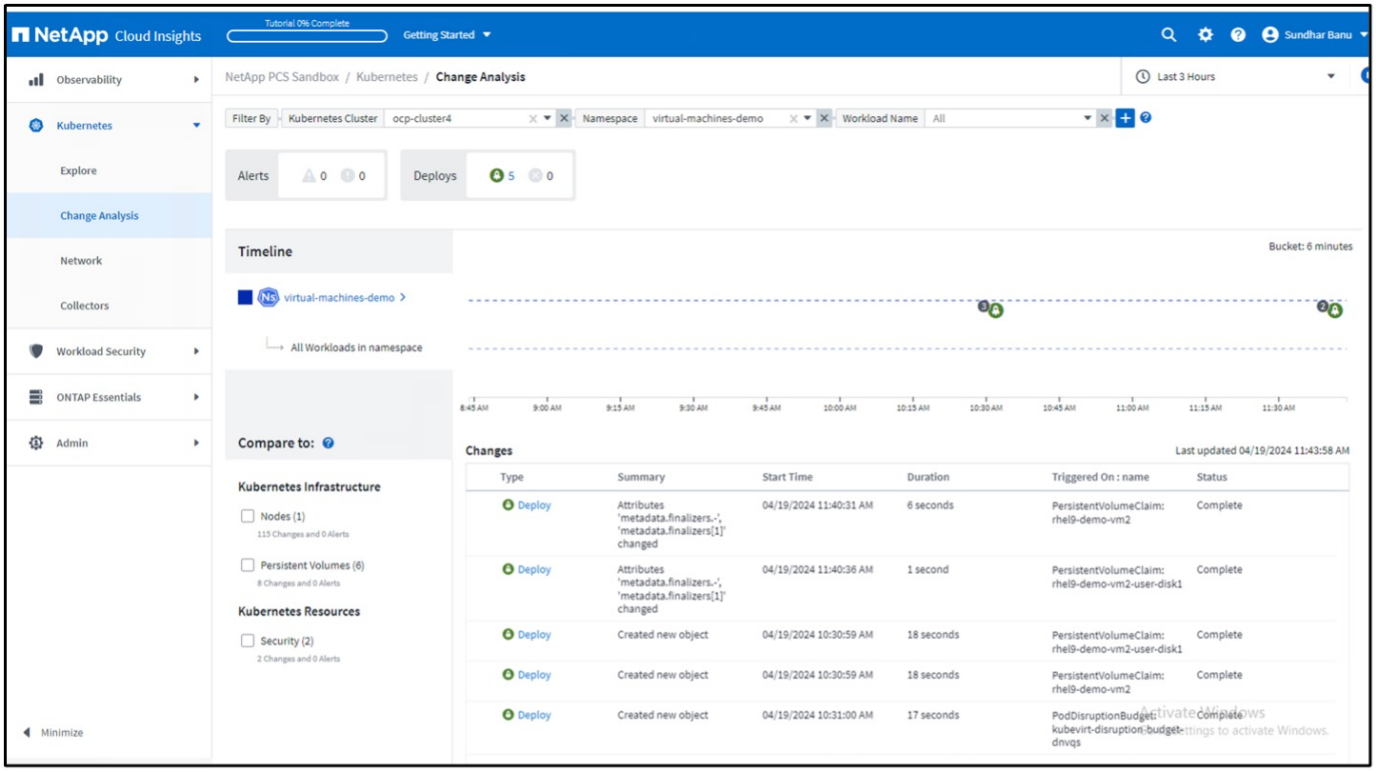

In the above example, Change Analysis is configured on the OpenShift cluster for the namespace that contains an OpenShift Virtualization VM. The dashboard shows changes against the timeline. You can drill down to see what changed and the click on All Changes Diff to see the diff of the manifests. From the manifest, you can see that a new backup of the persistent disks was created.

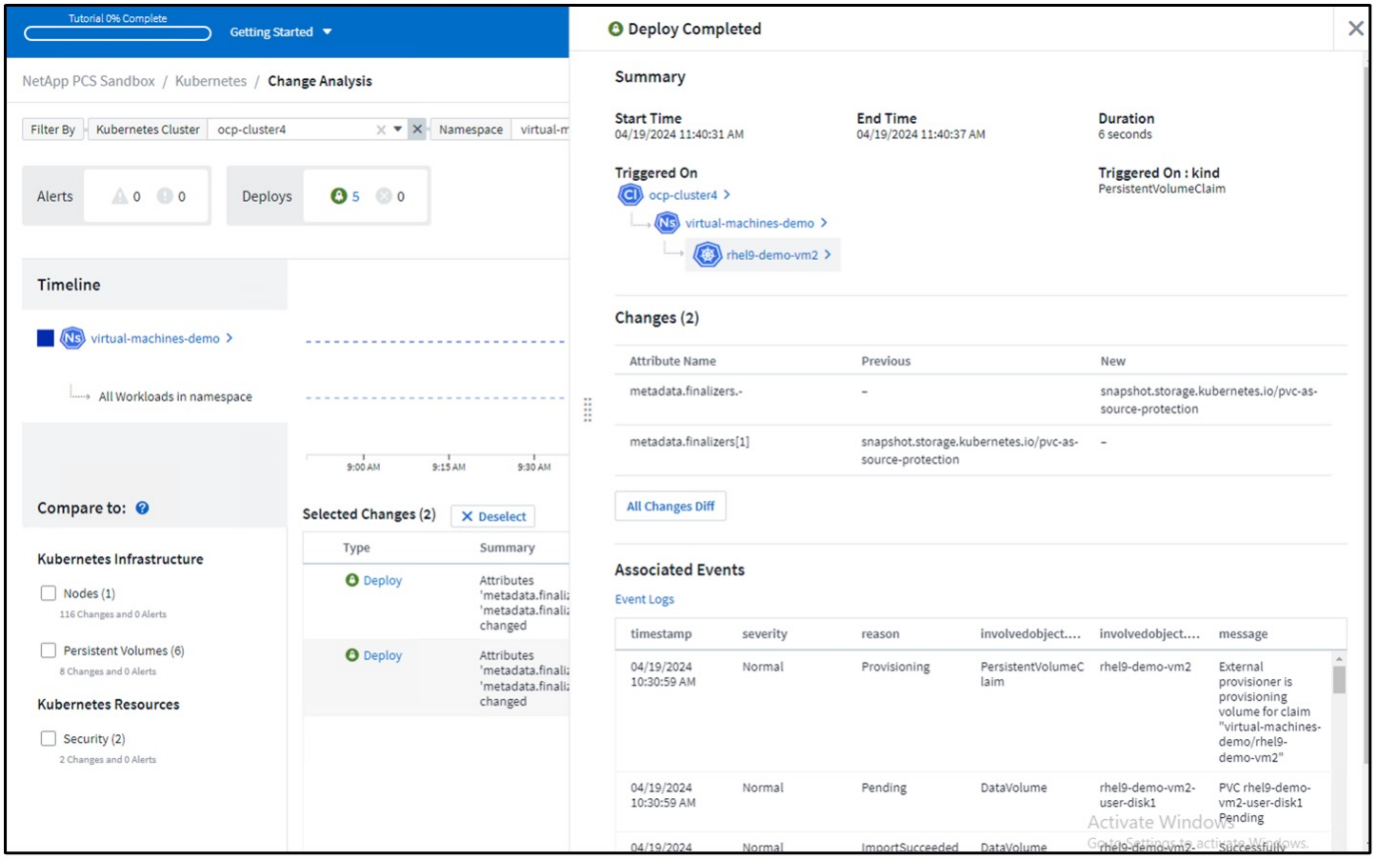

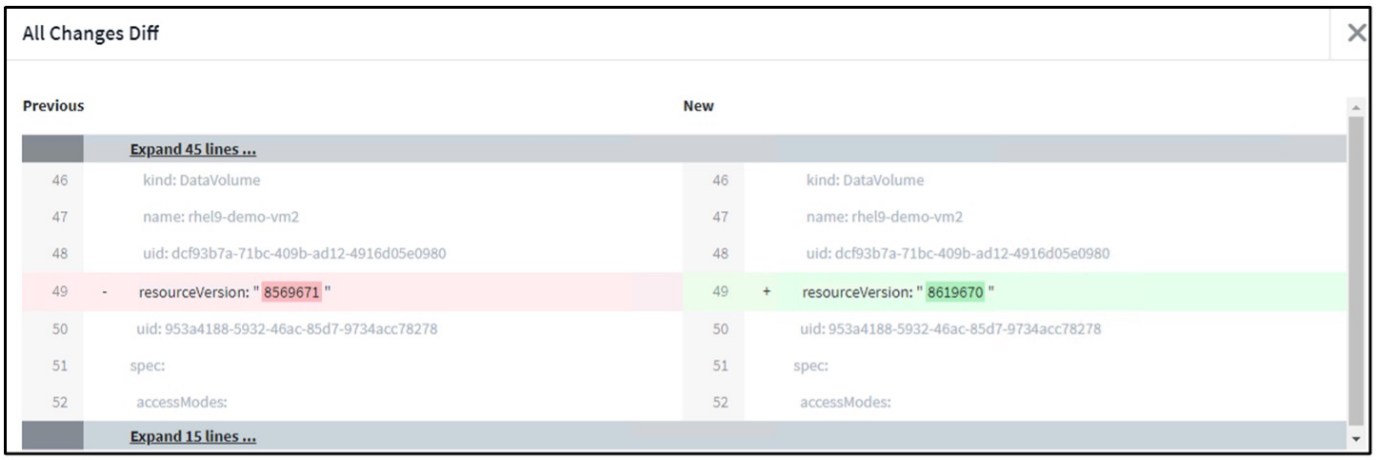

Backend Storage Mapping

With Cloud Insights, you can easily see the backend storage of the VM disks and several statistics about the PVCs.

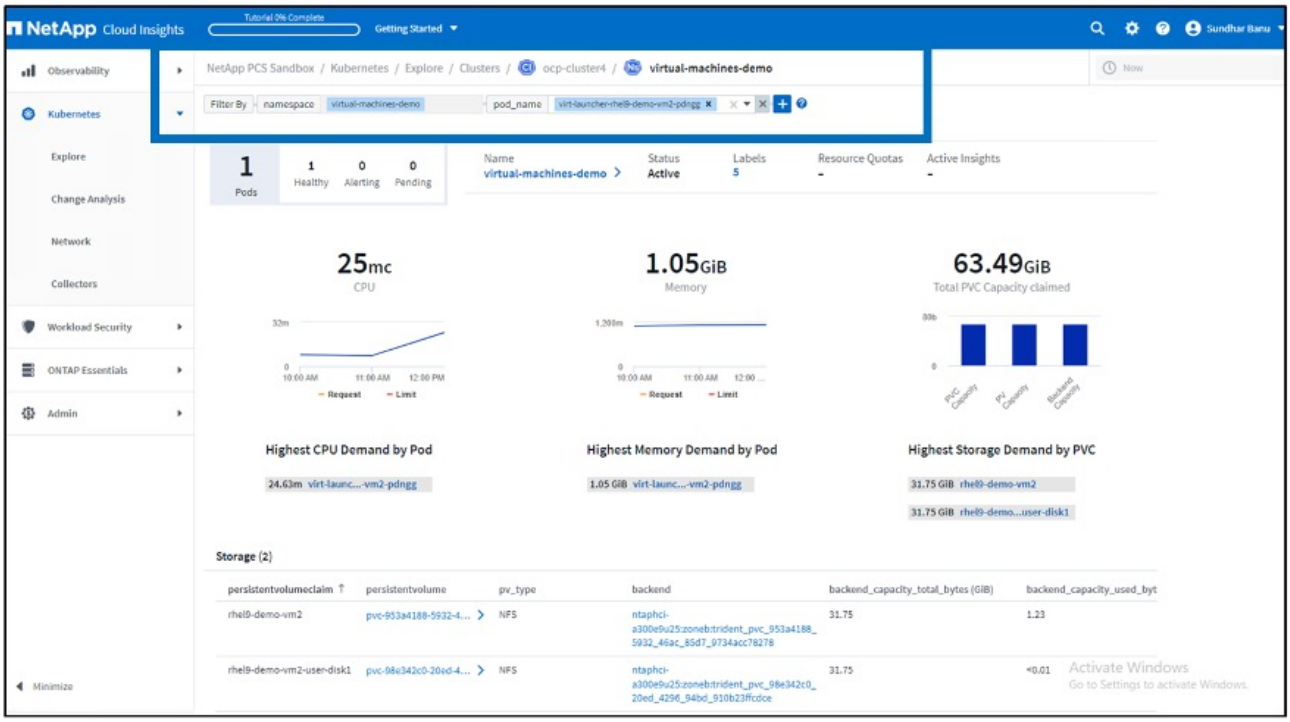

You can click on the links under the backend column, which will pull data directly from the backend ONTAP storage.

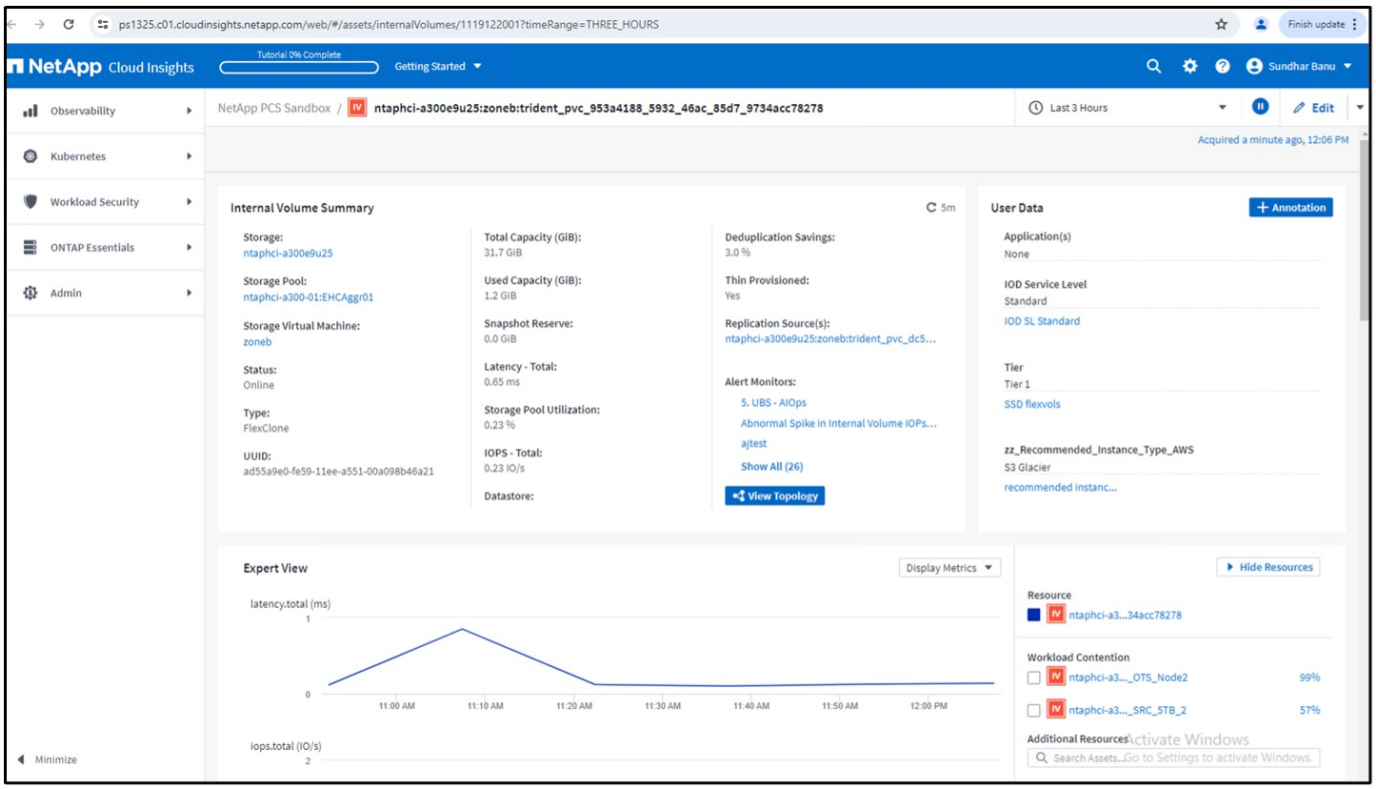

Another way to look at all the pod to storage mapping is creating an All Metrics query From Observability menu under Explore.

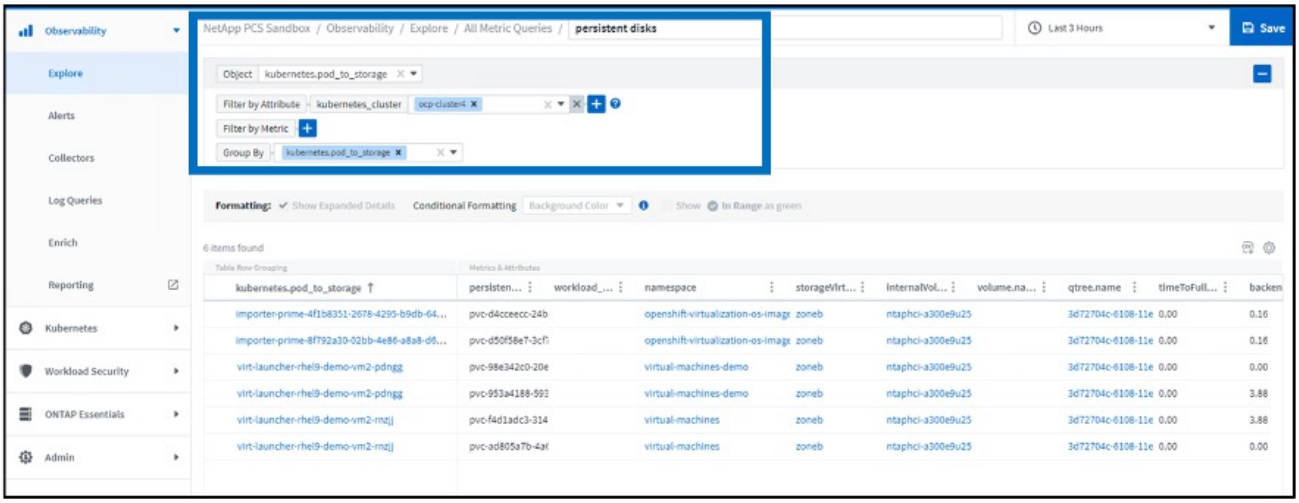

Clicking on any of the links will give you the corresponding details from ONTP storage. For example, clicking on an SVM name in the storageVirtualMachine column will pull details about the SVM from ONTAP. Clicking on an internal volume name will pull details about the volume in ONTAP.

	storageVirtualMachin...	internalVolume.name	volume.na..
zation-os-image	zoneb	ntaphci-a300e9u25:zoneb:trident_p	
zation-os-image	zoneb	ntaphci-a300e9u25:zoneb:trident_p	
demo	zoneb	ntaphci-a300e9u25:zoneb:trident_p	
demo	zoneb	ntaphci-a300e9u25:zoneb:trident_p	
	zoneb	ntaphci-a300e9u25:zoneb:trident_p	
	zoneb	ntaphci-a300e9u25:zoneb:trident_p	

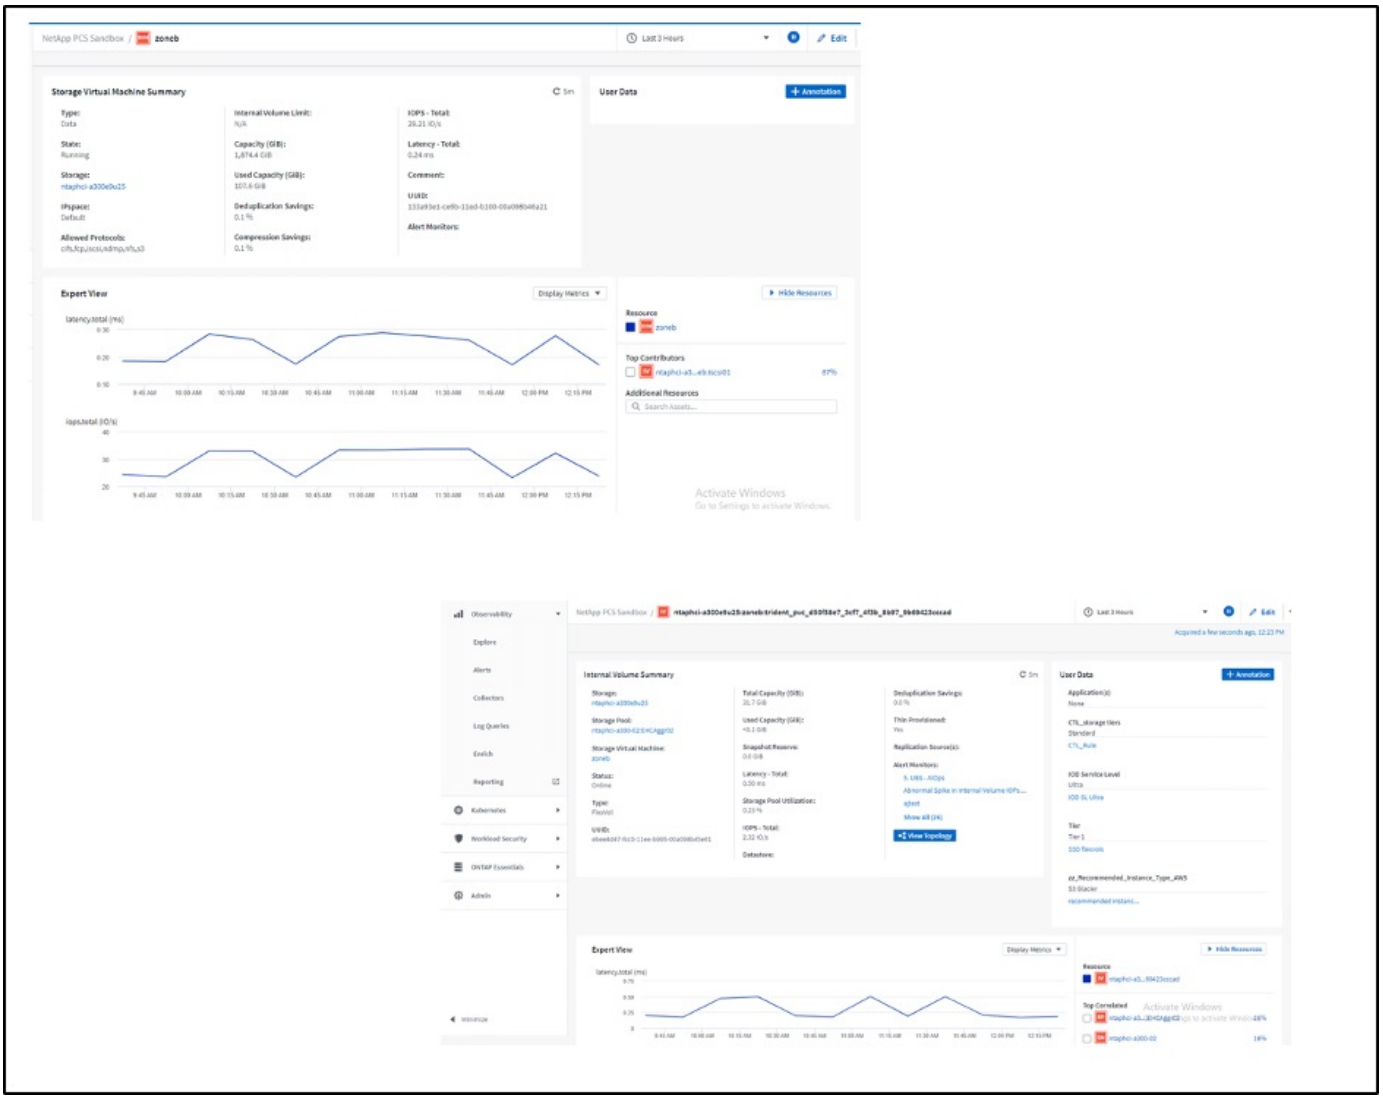

Additional Resources

Copyright information

Copyright © 2024 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.