



Virtualization with VMware

NetApp Solutions

NetApp
May 17, 2024

This PDF was generated from <https://docs.netapp.com/us-en/netapp-solutionshttps://docs.netapp.com/us-en/ontap-apps-dbs/vmware/vmware-vsphere-overview.html> on May 17, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- NetApp Solutions for Virtualization with VMware by Broadcom 1
 - VMware Cloud Foundation 1
 - Migration of VMs 138
 - NetApp Hybrid Multicloud with VMware Solutions 187
 - VMware Hybrid Multicloud Use Cases 187
 - Introduction to automation for ONTAP and vSphere 188
 - Virtual Desktops 200
 - NetApp All-Flash SAN Array with VMware vSphere 8 239
 - Demos and Tutorials 284

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

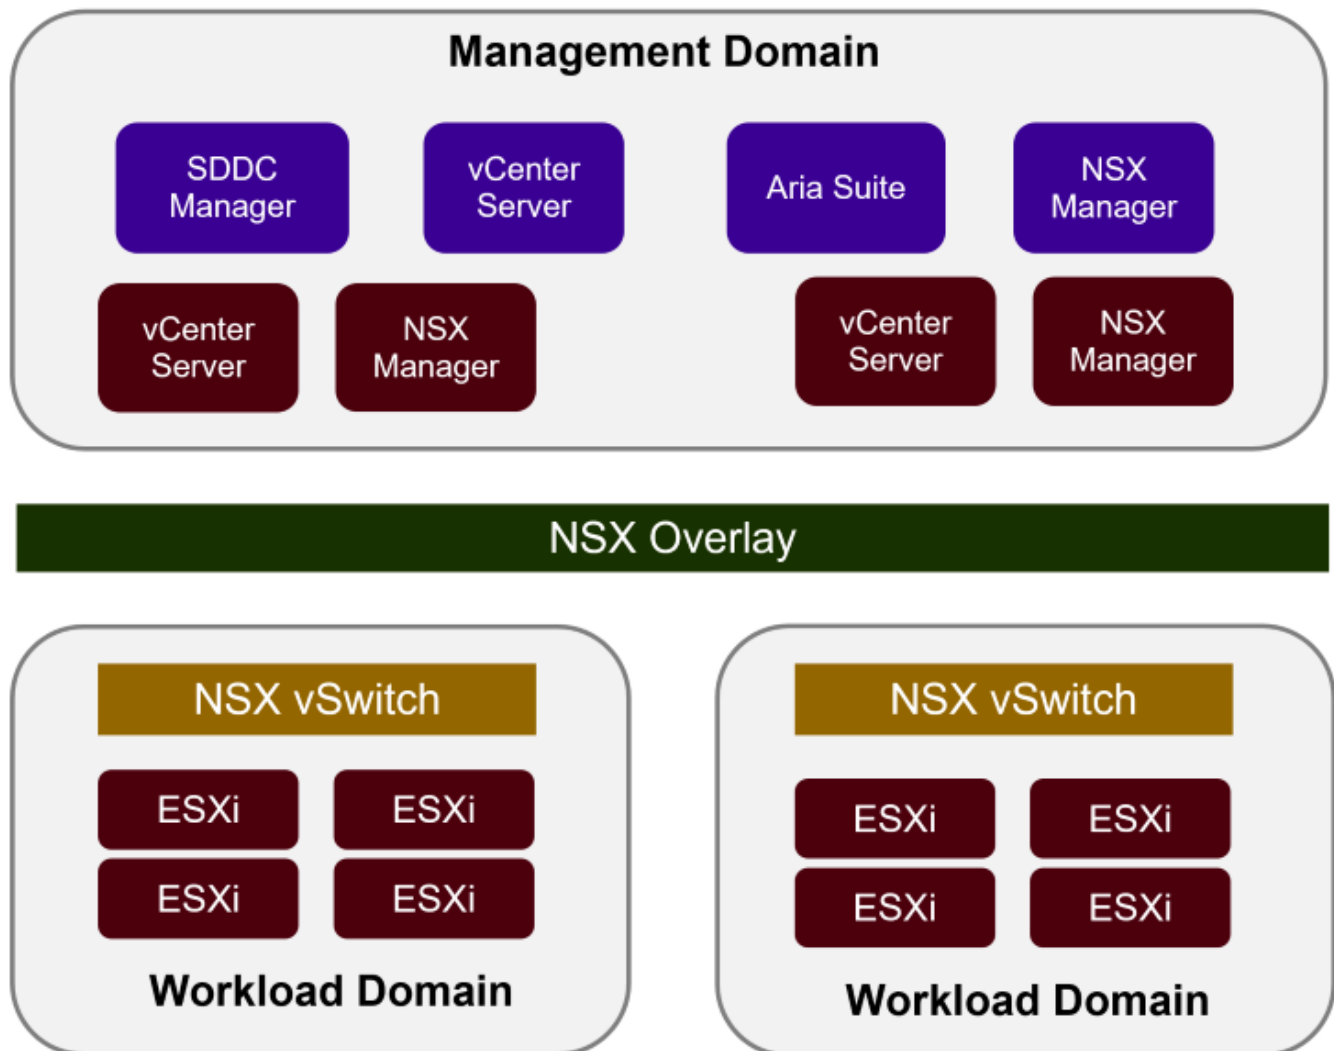
The VCF with NetApp ASA solution is comprised of the following major components:

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 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**.

Add Storage VM



STORAGE VM NAME

SVM_ISCSI

IPSPACE

Default

Access Protocol

SMB/CIFS, NFS, S3

iSCSI

FC

NVMe

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.

NETWORK INTERFACE

ntaphci-a300-01

IP ADDRESS

172.21.118.179

SUBNET MASK

24

GATEWAY

[Add optional gateway](#)

BROADCAST DOMAIN AND PORT

NFS_iSCSI

Use the same subnet mask, gateway, and broadcast domain for all of the following interfaces

IP ADDRESS

172.21.119.179

PORT

a0a-3375

ntaphci-a300-02

IP ADDRESS

172.21.118.180

PORT

a0a-3374

IP ADDRESS

172.21.119.180

PORT

a0a-3375

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


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.

New Distributed Port Group

1 Name and location

2 **Configure settings**

3 Ready to complete

Configure settings

Set general properties of the new port group.

Port binding Static binding

Port allocation Elastic ⓘ

Number of ports 8

Network resource pool (default)

VLAN

VLAN type VLAN

VLAN ID 3374

Advanced

Customize default policies configuration

CANCEL

BACK

NEXT

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....**

The screenshot shows the vSphere Client interface. At the top, there is a search bar and navigation icons. The main area displays a tree view of the environment. Under 'Management Networks', the 'vcf-m01-cl01-vds01' folder is expanded, showing several port groups. The 'vcf-m01-cl01-vds01-pg-iscsi-a' port group is selected, and a context menu is open over it. The 'Edit Settings...' option is highlighted, and a mouse cursor is pointing at it. To the right, the 'Distributed Port Group Details' panel is visible, showing the configuration for the selected port group: Port binding: Static binding, Port allocation: Elastic, VLAN ID: 3374, Distributed switch: vcf-m01-cl01-vds01, Network protocol profile: --, Network resource pool: --, and Hosts: 4.

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**.

Distributed Port Group - Edit Settings | vcf-m01-cl01-vds01-pg-iscsi-a ×

General	Load balancing	Route based on originating virtual por ▾
Advanced	Network failure detection	Link status only ▾
VLAN	Notify switches	Yes ▾
Security	Failback	Yes ▾
Traffic shaping		
Teaming and failover		
Monitoring		
Miscellaneous		

Failover order ⓘ

MOVE UP MOVE DOWN

Active uplinks

uplink1

Standby uplinks

Unused uplinks

uplink2

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 

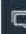
MOVE UP MOVE DOWN

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.



Select connection type

Select a connection type to create.

- VMkernel Network Adapter**
The VMkernel TCP/IP stack handles traffic for ESXi services such as vSphere vMotion, iSCSI, NFS, FCoE, Fault Tolerance, vSAN, host management and etc.
- Virtual Machine Port Group for a Standard Switch**
A port group handles the virtual machine traffic on standard switch.
- Physical Network Adapter**
A physical network adapter handles the network traffic to other hosts on the network.

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

Add Networking

- 1 Select connection type
- 2 Select target device**
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Select target device

Select a target device for the new connection.

- Select an existing network
- Select an existing standard switch
- New standard switch

Quick Filter

Enter value

	Name	NSX Port Group ID	Distributed Switch
<input type="radio"/>	SDDC-DPortGroup-VM-Mgmt	--	vcf-m01-cl01-vds01
<input checked="" type="radio"/>	vcf-m01-cl01-vds01-pg-iscsi-a	--	vcf-m01-cl01-vds01
<input type="radio"/>	vcf-m01-cl01-vds01-pg-iscsi-b	--	vcf-m01-cl01-vds01
<input type="radio"/>	vcf-m01-cl01-vds01-pg-mgmt	--	vcf-m01-cl01-vds01
<input type="radio"/>	vcf-m01-cl01-vds01-pg-vmotion	--	vcf-m01-cl01-vds01
<input type="radio"/>	vcf-m01-cl01-vds01-pg-vsan	--	vcf-m01-cl01-vds01

Manage Columns 6 items

CANCEL

BACK

NEXT

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

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties**
- 4 IPv4 settings
- 5 Ready to complete

Port properties

Specify VMkernel port settings.

Network label

MTU

TCP/IP stack

Available services

Enabled services

- vMotion
- Provisioning
- Fault Tolerance logging
- Management
- vSphere Replication
- vSphere Replication NFC
- vSAN
- vSAN Witness
- vSphere Backup NFC
- NVMe over TCP
- NVMe over RDMA

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.

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

IPv4 settings

Specify VMkernel IPv4 settings.

Obtain IPv4 settings automatically
 Use static IPv4 settings

IPv4 address

Subnet mask

Default gateway Override default gateway for this adapter

DNS server addresses

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

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Ready to complete

Review your selections before finishing the wizard

- ▼ Select target device

Distributed port group vcf-m01-cl01-vds01-pg-iscsi-a

Distributed switch vcf-m01-cl01-vds01
- ▼ Port properties

New port group vcf-m01-cl01-vds01-pg-iscsi-a (vcf-m01-cl01-vds01)

MTU 9000

vMotion Disabled

Provisioning Disabled

Fault Tolerance logging Disabled

Management Disabled

vSphere Replication Disabled

vSphere Replication NFC Disabled

vSAN Disabled

vSAN Witness Disabled

vSphere Backup NFC Disabled

NVMe over TCP Disabled

NVMe over RDMA Disabled
- ▼ IPv4 settings

IPv4 address 172.21.118.114 (static)

Subnet mask 255.255.255.0

CANCEL
BACK
FINISH

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.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

Local file

netapp-ontap-tools-for-vmware-vsphere-9.13-9554.ova

- 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.
- For the storage location of the configuration and disk files, select the vSAN datastore of the VCF management domain cluster.

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

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine [i](#)

Select virtual disk format As defined in the VM storage policy


VM Storage Policy Datastore Default

Disable Storage DRS for this virtual machine

	Name	Storage Compatibility	Capacity	Provisioned	Free	T
<input checked="" type="radio"/>	vcf-m01-c01-ds-vsan01	--	999.97 GB	7.17 TB	225.72 GB	v
<input type="radio"/>	vcf-m01-esx01-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v
<input type="radio"/>	vcf-m01-esx02-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v
<input type="radio"/>	vcf-m01-esx03-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v
<input type="radio"/>	vcf-m01-esx04-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v

Manage Columns Items per page 10 5 items

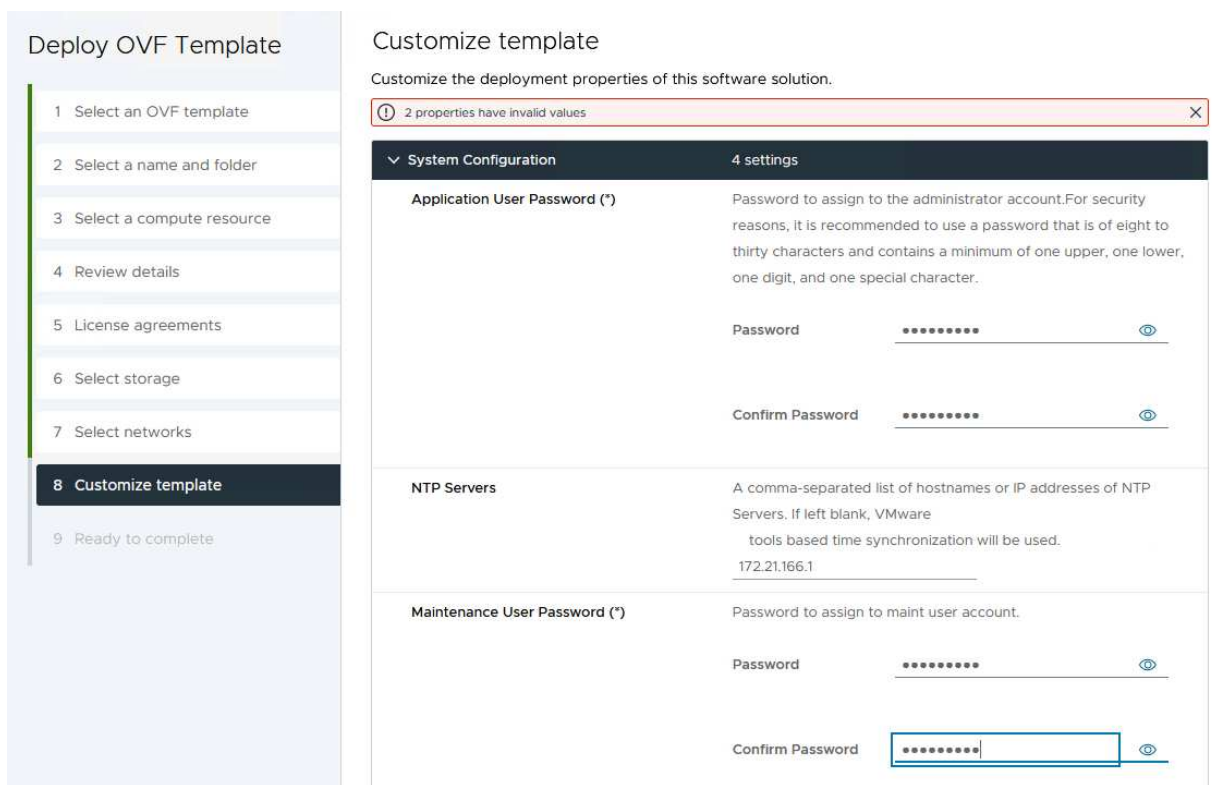
- On the Select network page select the network used for management traffic.



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.



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

Configure vCenter or Enable VCF		5 settings
Enable VMware Cloud Foundation (VCF)	vCenter server and user details are ignored when VCF is enabled. <input type="checkbox"/>	
vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to. 172.21.166.140	
Port (*)	Specify the HTTPS port of an existing vCenter to register to. 443	
Username (*)	Specify the username of an existing vCenter to register to. administrator@vsphere.local	
Password (*)	Specify the password of an existing vCenter to register to. Password: <input type="checkbox"/>	
	Confirm Password: <input type="checkbox"/>	
Network Properties		8 settings
Host Name	Specify the hostname for the appliance. (Leave blank if DHCP is desired) vcf-m01-otv9	
IP Address	Specify the IP address for the appliance. (Leave blank if DHCP is	


CANCEL BACK NEXT

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.

vSphere Client Search in all environments

NetApp ONTAP tools INSTANCE 172.21.166.139:8443

Overview

- Storage Systems
- Storage capability profile
- Storage Mapping
- Settings
- Reports
 - Datastore Report
 - Virtual Machine Report
 - vVols Datastore Report
 - vVols Virtual Machine Report
 - Log Integrity Report


ONTAP tools for VMware vSphere

Getting Started Traditional Dashboard vVols Dashboard

ONTAP tools for VMware vSphere is a vCenter Server plug-in that provides end-to-end lifecycle management for virtual machines in VMware environments using NetApp storage systems.


Add Storage System

Add storage systems to ONTAP tools for VMware vSphere.




Provision Datastore


Create traditional or vVols datastores.



Next Steps

 [View Dashboard](#)

View and monitor the datastores in ONTAP tools for VMware vSphere.

 [Settings](#)

Configure administrative settings such as credentials, alarm thresholds.

[What's new?](#)
September 4, 2023

- Qualified and supported with ONTAP 9.13.1
- Supports and interoperates with VMware vSphere 8.x releases
- Includes newer enhanced SCPs that efficiently map workloads to the newer All SAN Array platforms through policy based management



Resources

- [ONTAP tools for VMware vSphere Documentation Resources](#)
- [RBAC User Creator for Data ONTAP](#)
- [ONTAP tools for VMware vSphere REST API Documentation](#)

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

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	

CANCEL


SAVE & ADD MORE

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

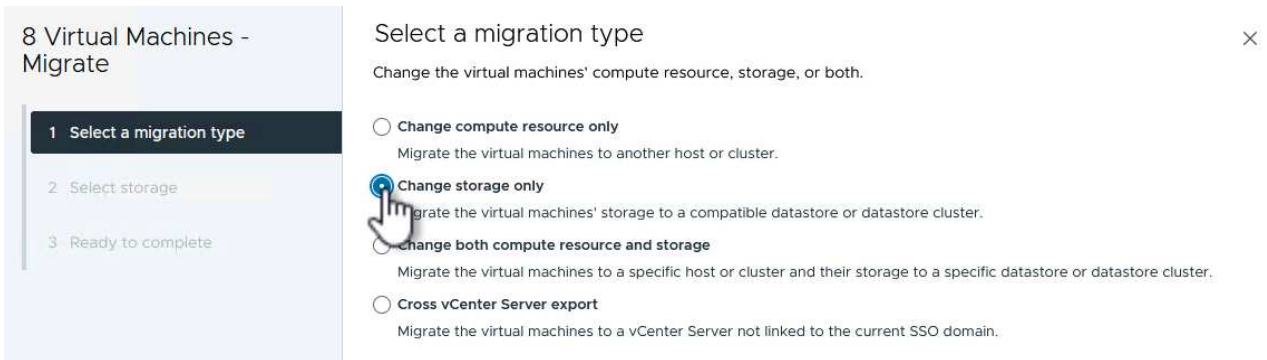
In cases where it is preferred to use ONTAP storage to protect the VCF management VM's vMotion can be used to migrate the VM's to the newly created iSCSI datastore.

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.

8 Virtual Machines - Migrate

1 Select a migration type

2 Select storage

3 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

BATCH CONFIGURE CONFIGURE PER DISK

Select virtual disk format Same format as source

VM Storage Policy Datastore Default

Disable Storage DRS for this virtual machine

Name	Storage Compatibility	Capacity	Provisioned	Free
mgmt_01_iscsi	--	3 TB	1.46 GB	3 TB
vcf-m01-cl01-ds-vsan01	--	999.97 GB	7.28 TB	52.38 GB

Manage Columns Items per page 10 2 items

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

5. Review the selections and click on **Finish** to start the migration.

6. The relocation status can be viewed from the **Recent Tasks** pane.

Task Name	Target	Status	Details
Relocate virtual machine	vcf-w01-nsx03	38%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-wkld-vc01	42%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-m01-otv9	36%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-m01-nsx01a	49%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-w01-nsx02	47%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-m01-sddcm01	39%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-w01-nsx01	42%	Migrating Virtual Machine active state
Relocate virtual machine	vcf-m01-vc01	44%	Migrating Virtual Machine active state

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**.

Add Storage VM



STORAGE VM NAME

SVM_ISCSI

IPSPACE

Default

Access Protocol

SMB/CIFS, NFS, S3

iSCSI

FC

NVMe

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.

NETWORK INTERFACE

ntaphci-a300-01

IP ADDRESS

172.21.118.179

SUBNET MASK

24

GATEWAY

Add optional gateway

BROADCAST DOMAIN AND PORT

NFS_iSCSI

Use the same subnet mask, gateway, and broadcast domain for all of the following interfaces

IP ADDRESS

172.21.119.179

PORT

a0a-3375

ntaphci-a300-02

IP ADDRESS

172.21.118.180

PORT

a0a-3374

IP ADDRESS

172.21.119.180

PORT

a0a-3375

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

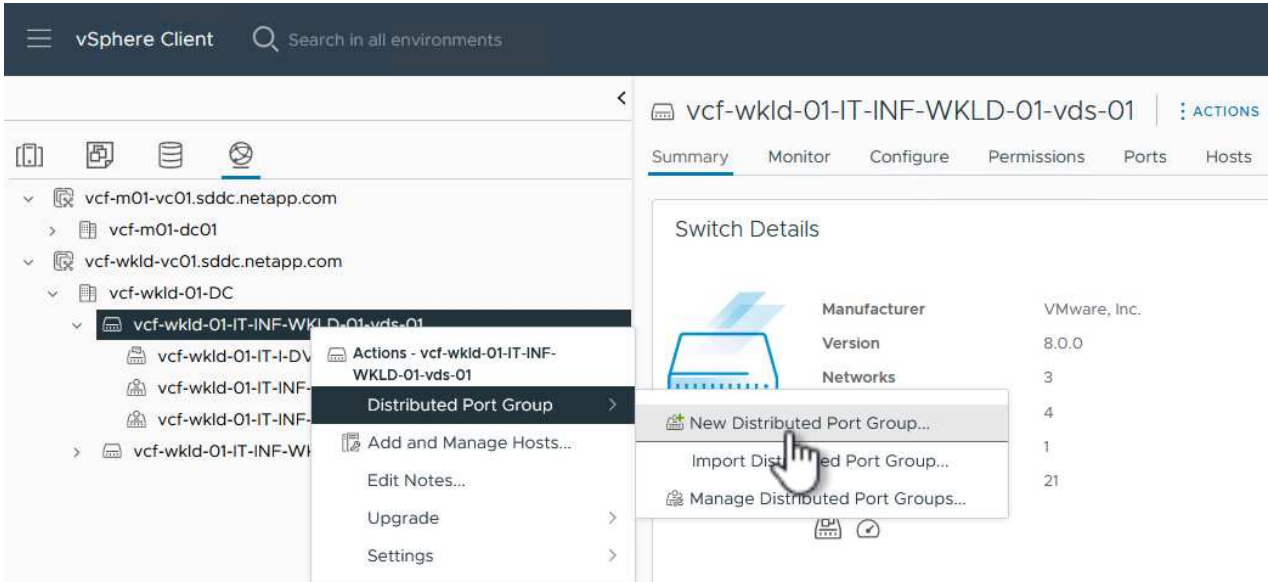
Set up networking for iSCSI on ESXi hosts

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:

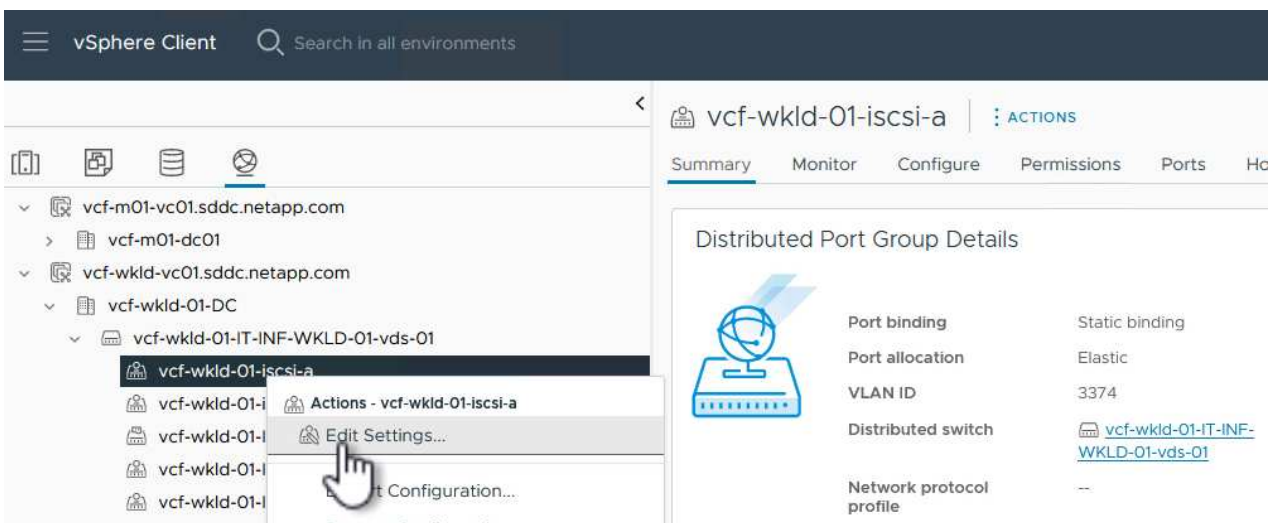
1. From the vSphere client , navigate to **Inventory > Networking** for the workload domain. Navigate to the existing Distributed Switch and choose the action to create **New Distributed Port Group....**



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**.

Distributed Port Group - Edit Settings | vcf-wkld-01-iscsi-a ×

General	Load balancing	Route based on originating virtual por ▼
Advanced	Network failure detection	Link status only ▼
VLAN	Notify switches	Yes ▼
Security	Failback	Yes ▼
Traffic shaping		
Teaming and failover		
Monitoring		
Miscellaneous		

Failover order ⓘ

MOVE UP MOVE DOWN

Active uplinks

uplink1

Standby uplinks

Unused uplinks

uplink2

CANCEL OK

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-wkld-01-iscsi-b

General	Load balancing	Route based on originating virtual por ▼
Advanced	Network failure detection	Link status only ▼
VLAN	Notify switches	Yes ▼
Security	Failback	Yes ▼
Traffic shaping		
Teaming and failover		
Monitoring		
Miscellaneous		

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 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.

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Select target device

Select a target device for the new connection.

- Select an existing network
- Select an existing standard switch
- New standard switch

Quick Filter

Enter value

	Name	NSX Port Group ID	Distributed Switch
<input checked="" type="radio"/>	vcf-wkld-01-iscsi-a	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	vcf-wkld-01-iscsi-b	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-mgmt	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-nfs	--	vcf-wkld-01-IT-INF-WKLD-01-vds-02
<input type="radio"/>	vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-vmotion	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01

Manage Columns 5 items

CANCEL

BACK

NEXT

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

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Port properties

Specify VMkernel port settings.

Network label vcf-wkld-01-iscsi-a (vcf-wkld-01-IT-INF-WKLD-01-vds-01)

MTU Get MTU from switch 9000

TCP/IP stack Default

Available services

- Enabled services**
- vMotion
 - Provisioning
 - Fault Tolerance logging
 - Management
 - vSphere Replication
 - vSphere Replication NFC
 - vSAN
 - vSAN Witness
 - vSphere Backup NFC
 - NVMe over TCP
 - NVMe over RDMA

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.

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

IPv4 settings

Specify VMkernel IPv4 settings.

Obtain IPv4 settings automatically
 Use static IPv4 settings

IPv4 address

Subnet mask

Default gateway Override default gateway for this adapter

DNS server addresses

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

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Ready to complete

Review your selections before finishing the wizard

- ▼ **Select target device**

Distributed port group vcf-wkld-01-iscsi-a

Distributed switch vcf-wkld-01-IT-INF-WKLD-01-vds-01
- ▼ **Port properties**

New port group vcf-wkld-01-iscsi-a (vcf-wkld-01-IT-INF-WKLD-01-vds-01)

MTU 9000

vMotion Disabled

Provisioning Disabled

Fault Tolerance logging Disabled

Management Disabled

vSphere Replication Disabled

vSphere Replication NFC Disabled

vSAN Disabled

vSAN Witness Disabled

vSphere Backup NFC Disabled

NVMe over TCP Disabled

NVMe over RDMA Disabled
- ▼ **IPv4 settings**

IPv4 address 172.21.118.127 (static)

Subnet mask 255.255.255.0

CANCEL
BACK
FINISH

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.

Deploy OVF Template

1 Select an OVF template

- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

Local file

netapp-ontap-tools-for-vmware-vsphere-9.13-9554.ova

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.

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

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine [?](#)

Select virtual disk format As defined in the VM storage policy

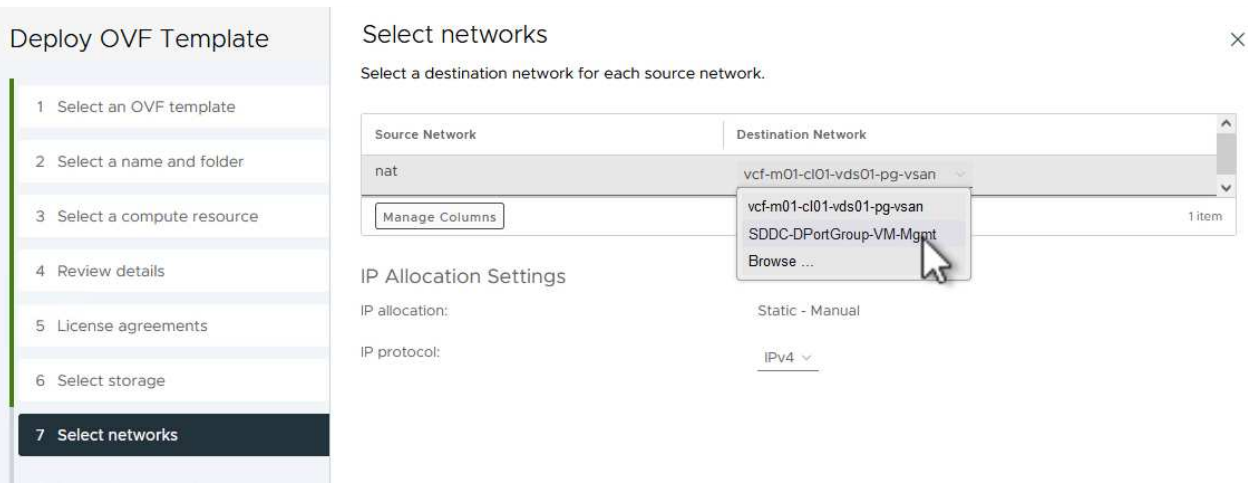
VM Storage Policy Datastore Default

Disable Storage DRS for this virtual machine

	Name	Storage Compatibility	Capacity	Provisioned	Free	
<input checked="" type="radio"/>	vcf-m01-c01-ds-vsan01	--	999.97 GB	7.17 TB	225.72 GB	v
<input type="radio"/>	vcf-m01-esx01-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v
<input type="radio"/>	vcf-m01-esx02-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v
<input type="radio"/>	vcf-m01-esx03-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v
<input type="radio"/>	vcf-m01-esx04-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	v

Manage Columns Items per page 10 5 items

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.

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.

❗ 2 properties have invalid values ✕

System Configuration		4 settings
Application User Password (*)	Password to assign to the administrator account. For security reasons, it is recommended to use a password that is of eight to thirty characters and contains a minimum of one upper, one lower, one digit, and one special character.	
	Password 👁
	Confirm Password 👁
NTP Servers	A comma-separated list of hostnames or IP addresses of NTP Servers. If left blank, VMware tools based time synchronization will be used. 172.21.166.1	
Maintenance User Password (*)	Password to assign to maint user account.	
	Password 👁
	Confirm Password 👁

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

❗ 2 properties have invalid values ✕


Configure vCenter or Enable VCF		3 settings
Enable VMware Cloud Foundation (VCF)	vCenter server and user details are ignored when VCF is enabled. <input type="checkbox"/>	
vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to. cf-wkld-vc01.sddc.netapp.com	
Port (*)	Specify the HTTPS port of an existing vCenter to register to. 443	
Username (*)	Specify the username of an existing vCenter to register to. administrator@vsphere.local	
Password (*)	Specify the password of an existing vCenter to register to.	
	Password 👁
	Confirm Password 👁
Network Properties		8 settings
Host Name	Specify the hostname for the appliance. (Leave blank if DHCP is desired) vcf-w01-otv9	
IP Address	Specify the IP address for the appliance. (Leave blank if DHCP is desired)	

CANCEL BACK NEXT

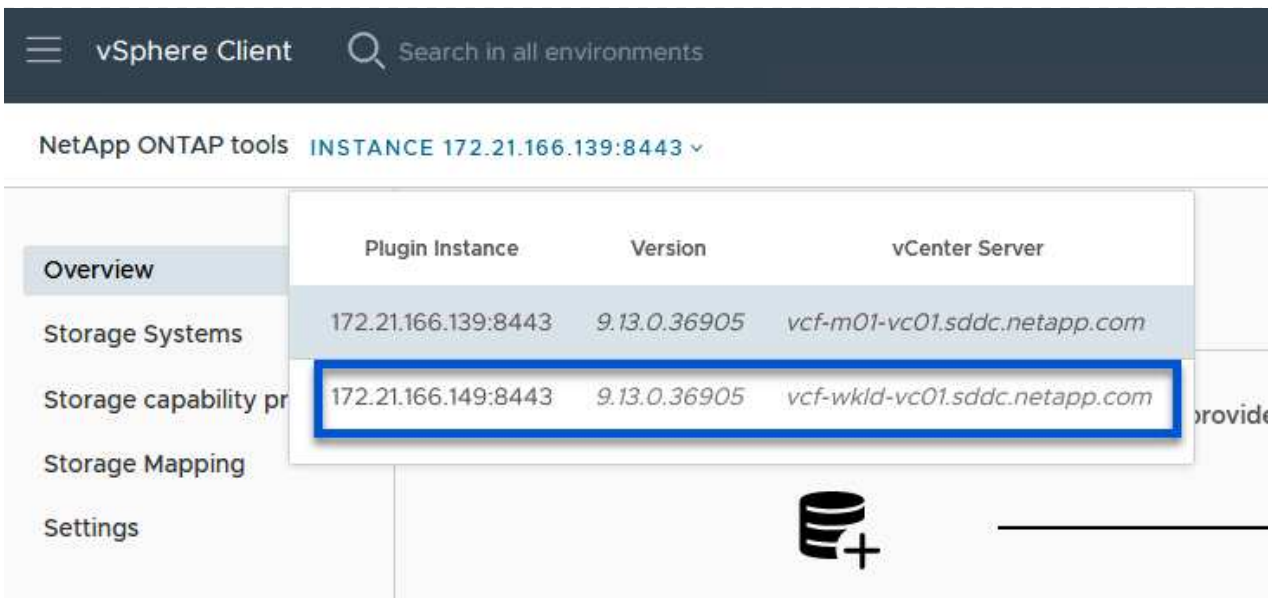
9. Review all information on the Ready to complete page and the click Finish to begin deploying the OTV appliance.

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.

Create Storage Capability Profile

1 General

2 Platform

3 Protocol

4 Performance

5 Storage attributes

6 Summary

Platform

Platform: Performance

Asymmetric:

CANCEL

BACK

NEXT

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

Create Storage Capability Profile

1 General

2 Platform

3 Protocol

4 Performance

5 Storage attributes

6 Summary

Protocol

Protocol:

Any

Any

FCP

iSCSI

NVMe/FC

CANCEL

BACK

NEXT

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

Create Storage Capability Profile

1 General

2 Platform

3 Protocol

4 Performance

5 Storage attributes

6 Summary

Performance

None ⓘ

QoS policy group ⓘ

Min IOPS:

Max IOPS:

6000

Unlimited

CANCEL

BACK

NEXT

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

Create Storage Capability Profile

1 General

2 Platform

3 Protocol

4 Performance

5 Storage attributes

6 Summary

Storage attributes

Deduplication:

Yes

Compression:

Yes

Space reserve:

Thin

Encryption:

No

Tiering policy (FabricPool):

None

CANCEL

BACK

NEXT

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.



The screenshot shows the vSphere Client interface with the ONTAP Tools dashboard. The left sidebar contains a navigation menu with 'Overview' selected. The main content area is titled 'ONTAP tools for VMware vSphere' and has three tabs: 'Getting Started', 'Traditional Dashboard', and 'vVols Dashboard'. Below the tabs, there is a diagram showing the workflow from 'Add Storage System' to 'Provision Datastore'. The 'Provision Datastore' step is highlighted with a blue box, and a green 'PROVISION' button is visible below it.

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.



The screenshot shows the 'New Datastore' wizard in the 'General' step. The left sidebar shows the wizard steps: '1 General', '2 Storage system', '3 Storage attributes', and '4 Summary'. The main content area is titled 'General' and contains the following fields:

- Provisioning destination:** IT-INF-WKLD-01 (with a 'BROWSE' button)
- Type:** NFS VMFS vVols
- Name:** VCF_WKLD_02_VVOLS
- Description:** (empty text box)
- Protocol:** NFS iSCSI FC / FCoE NVMe/FC

At the bottom right, there are 'CANCEL' and 'NEXT' buttons.

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

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage system

Specify the storage capability profiles and the storage system you want to use.

Storage capability profiles:

AFF_Encrypted_Min50_ASA_A
FAS_Default
FAS_Max20
Custom profiles
ASA_Gold_iSCSI

Storage system:

ntaphci-a300e9u25 (172.16.9.25)

Storage VM:

VCF_iSCSI

CANCEL

BACK

NEXT

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.

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage attributes

Specify the storage details for provisioning the datastore.

Volumes: Create new volumes Select volumes

Create new volumes

Name	Size	Storage Capability Profile	Aggregate
 FlexVol volumes are not added.			

Name	Size(GB) ⓘ	Storage capability profile	Aggregates	Space reserve
f_wkld_02_vvols	3000	ASA_Gold_iSCSI	EHCaggr02 - (27053.3 GE)	Thin

ADD

CANCEL

BACK

NEXT

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

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Summary

Datastore type: vVols
Protocol: iSCSI
Storage capability profile: ASA_Gold_iSCSI

Storage system details

Storage system: ntaphcl-a300e9u25
SVM: VCF_iSCSI

Storage attributes

New FlexVol Name	New FlexVol Size	Aggregate	Storage Capability Profile
vcf_wkld_02_vvols	3000 GB	EHCAGgr02	ASA_Gold_iSCSI

Click 'Finish' to provision this datastore.

CANCEL
BACK
FINISH

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**.

Add Storage VM



STORAGE VM NAME

VCF_NVMe

IPSPACE

Default

Access Protocol

SMB/CIFS, NFS, S3

iSCSI

FC

NVMe

Enable NVMe/FC

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.

NETWORK INTERFACE

ntaphci-a300-01

IP ADDRESS

172.21.118.189

SUBNET MASK

24

GATEWAY

[Add optional gateway](#)

BROADCAST DOMAIN AND PORT

NFS_iSCSI

Use the same subnet mask, gateway, and broadcast domain for all of the following interfaces

IP ADDRESS

172.21.119.189

PORT

a0a-3375

ntaphci-a300-02

IP ADDRESS

172.21.118.190

PORT

a0a-3374

IP ADDRESS

172.21.119.190

PORT

a0a-3375

Storage VM Administration

Manage administrator account

Save

Cancel

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
```

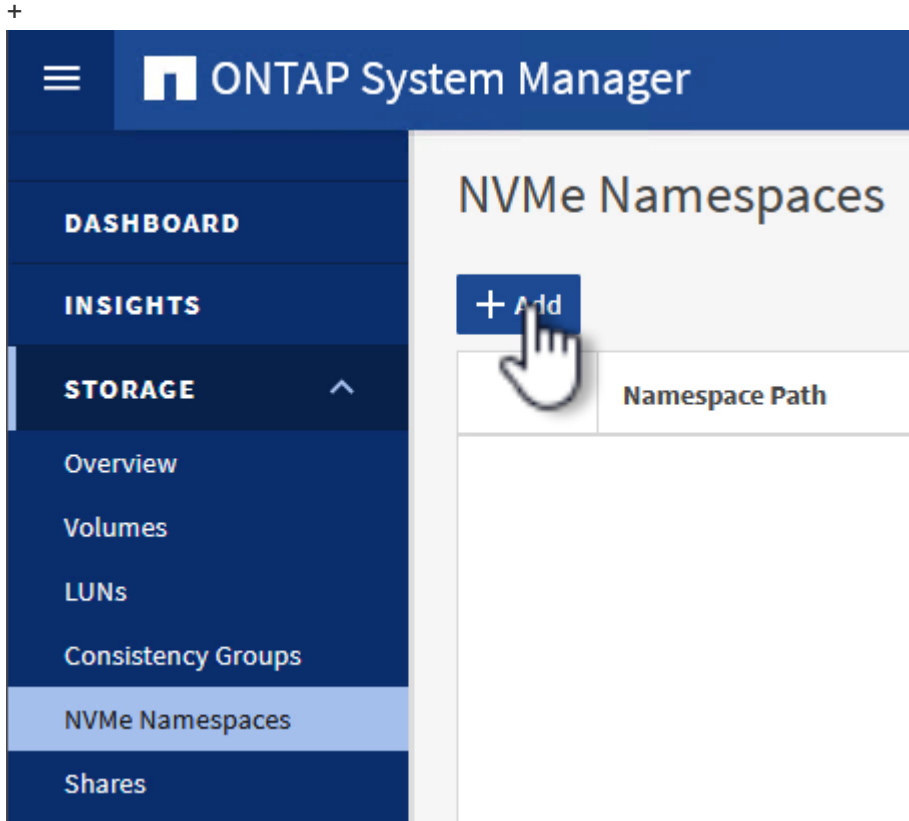
2. Record the NQN for each ESXi host in the cluster
3. From ONTAP System Manager navigate to **NVMe Namespaces** in the left-hand menu and click on **+ Add** to start.



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.




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.

New Distributed Port Group

1 Name and location

2 **Configure settings**

3 Ready to complete

Configure settings

Set general properties of the new port group.


Port binding	Static binding
Port allocation	Elastic ?
Number of ports	8
Network resource pool	(default)
VLAN	
VLAN type	VLAN
VLAN ID	3374
Advanced	
<input type="checkbox"/> Customize default policies configuration	

CANCEL

BACK

NEXT

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....**



- 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**.

Distributed Port Group - Edit Settings | vcf-wkld-01-nvme-a

General

Advanced

VLAN

Security

Traffic shaping

Teaming and failover

Monitoring

Miscellaneous

Load balancing Route based on originating virtual port

Network failure detection Link status only

Notify switches Yes

Failback Yes

Failover order ⓘ

MOVE UP MOVE DOWN

Active uplinks

uplink1

Standby uplinks

Unused uplinks

uplink2

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

Unused uplinks.

Distributed Port Group - Edit Settings | vcf-wkld-01-nvme-b

General

Advanced

VLAN

Security

Traffic shaping

Teaming and fallover

Monitoring

Miscellaneous

Load balancing

Route based on originating virtual port

Network failure detection

Link status only

Notify switches

Yes

Fallback

Yes

Failover order ⓘ

[MOVE UP](#) [MOVE DOWN](#)

Active uplinks

uplink2

Standby uplinks

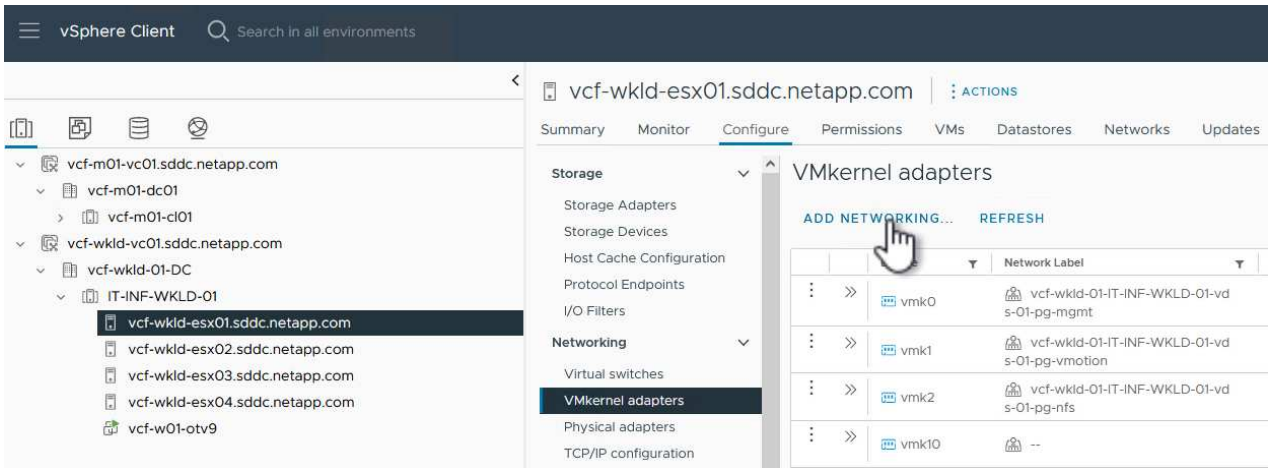
Unused uplinks

uplink1


Create VMkernel adapters on each ESXi host

Repeat this process on each ESXi host in the 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.

Add Networking

1 Select connection type

2 Select target device

3 Port properties

4 IPv4 settings

5 Ready to complete

Select target device








×

Select a target device for the new connection.

- Select an existing network
- Select an existing standard switch
- New standard switch

Quick Filter

Enter value

	Name	NSX Port Group ID	Distributed Switch
<input type="radio"/>	 vcf-wkld-01-iscsi-a	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	 vcf-wkld-01-iscsi-b	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	 vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-mgmt	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	 vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-nfs	--	vcf-wkld-01-IT-INF-WKLD-01-vds-02
<input type="radio"/>	 vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-vmotion	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input checked="" type="radio"/>	 vcf-wkld-01-nvme-a	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01
<input type="radio"/>	 vcf-wkld-01-nvme-b	--	vcf-wkld-01-IT-INF-WKLD-01-vds-01

Manage Columns 7 Items

CANCEL

BACK

NEXT



Packages

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

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Port properties ✕

Specify VMkernel port settings.

Network label vcf-wkld-01-nvme-a (vcf-wkld-01-IT-INF-WKLD-01-vds-01)

MTU Get MTU from switch ▾ 9000 ↻

TCP/IP stack Default ▾

Available services

Enabled services

<input checked="" type="checkbox"/> vMotion	<input type="checkbox"/> vSphere Replication NFC	<input type="checkbox"/> NVMe over RDMA
<input type="checkbox"/> Provisioning	<input type="checkbox"/> vSAN	
<input type="checkbox"/> Fault Tolerance logging	<input type="checkbox"/> vSAN Witness	
<input type="checkbox"/> Management	<input type="checkbox"/> vSphere Backup NFC	
<input type="checkbox"/> vSphere Replication	<input checked="" type="checkbox"/> NVMe over TCP	

CANCEL
BACK
NEXT

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.

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

IPv4 settings ✕

Specify VMkernel IPv4 settings.

Obtain IPv4 settings automatically

Use static IPv4 settings

IPv4 address 172.21.118.191

Subnet mask 255.255.255.0

Default gateway Override default gateway for this adapter

172.21.166.1

DNS server addresses 10.61.185.231

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

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Ready to complete

Review your selections before finishing the wizard

▼ Select target device

Distributed port group	vcf-wkld-01-nvme-a
Distributed switch	vcf-wkld-01-IT-INF-WKLD-01-vds-01

▼ Port properties

New port group	vcf-wkld-01-nvme-a (vcf-wkld-01-IT-INF-WKLD-01-vds-01)
MTU	9000
vMotion	Disabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled
vSAN Witness	Disabled
vSphere Backup NFC	Disabled
NVMe over TCP	Enabled
NVMe over RDMA	Disabled

▼ IPv4 settings

IPv4 address	172.21.118.191 (static)
Subnet mask	255.255.255.0

CANCEL

BACK

FINISH

Packages

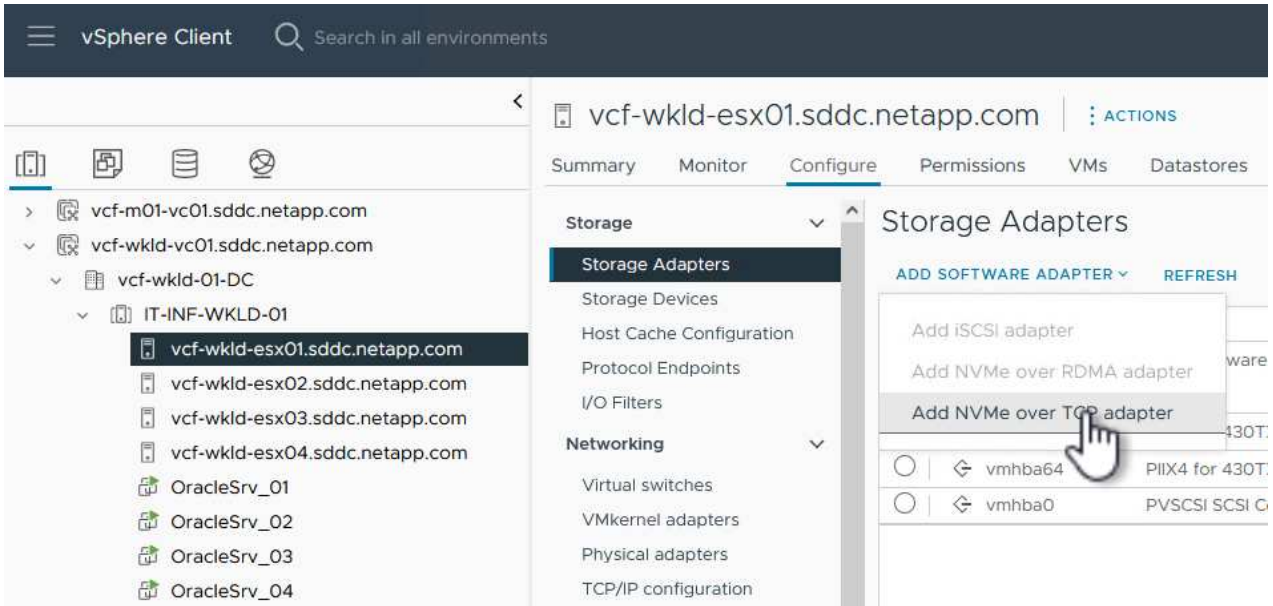
7. Repeat this process to create a VMkernel adapter for the second iSCSI network.

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.

Add Software NVMe over TCP adapter


vcf-wkld-esx01.sddc.netapp.com

Enable software NVMe adapter on the selected physical network adapter.

Physical Network 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.

Add controller | vmhba68



Automatically Manually

Host NQN

nqn.2014-08.com.netapp.sddc:nvme:vcf-wkld-...

COPY

IP

172.21.118.189

Enter IPv4 / IPv6 address

Central discovery controller

Port Number

Range more from 0

Digest parameter

Header digest

Data digest

DISCOVER CONTROLLERS

Select which controller to connect

<input type="checkbox"/>	Id	Subsystem NQN	Transport Type	IP	Port Number
<input checked="" type="checkbox"/>	65535	nqn.1992-08.com.netapp:sn.64df3069fb6411eea55100a098b46a21:subsystem.VCF_WKLD_04_NVMe_VCF_WKLD_04_NVMe	nvme	172.21.118.189	4420
<input checked="" type="checkbox"/>	65535	nqn.1992-08.com.netapp:sn.64df3069fb6411eea55100a098b46a21:subsystem.VCF	nvme	172.21.118.190	4420

4 items

3

4

OK

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

Storage Adapters

ADD SOFTWARE ADAPTER ▾ REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE

<input type="radio"/>	Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
<input type="radio"/>	vmhba65	iSCSI Software Adapter	iSCSI	Online	iscsi_vmk(ign.1998-01.com.vmware:vcf-wkld-esx01.sddc.netapp.com:794177624:65)	4	2	8
<input type="radio"/>	vmhba1	PIIX4 for 430TX/440BX/MX IDE Controller	Block SCSI	Unknown	--	1	1	1
<input type="radio"/>	vmhba64	PIIX4 for 430TX/440BX/MX IDE Controller	Block SCSI	Unknown	--	0	0	0
<input type="radio"/>	vmhba0	PVSCSI SCSI Controller	SCSI	Unknown	--	3	3	3
<input checked="" type="radio"/>	vmhba68	VMware NVMe over TCP Storage Adapter	NVMe over TCP	Online	--	1	1	1
<input type="radio"/>	vmhba69	VMware NVMe over TCP Storage Adapter	NVMe over TCP	Online	--	0	0	0

Manage Columns Export ▾ 6 items

Properties Devices Paths Namespaces Controllers

REFRESH ATTACH DETACH RENAME


<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type	Transport
<input type="checkbox"/>	NVMe TCP Disk (uuid.929a6a90457647849146e09d6e55b076)	0	disk	3.00 TB	Not Consumed	Attached	Supported	Flash	TCPRAN:RT

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.

New Datastore

1 Type

2 Name and device selection

3 VMFS version

4 Partition configuration

5 Ready to complete

Name and device selection

Specify datastore name and a disk/LUN for provisioning the datastore.

Name

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Cl V S
<input checked="" type="radio"/>	NVMe TCP Disk (uuid:929a6a90457647849146e09d6e55b076)	0	3.00 TB	Supported	Flash	512e	N
<input type="radio"/>	Local VMware Disk (naa:6000c29f83dcf1e42d230340deb66036)	0	4.00 GB	Not supported	Flash	512n	N
<input type="radio"/>	Local VMware Disk (naa:6000c291464644a835bc23d384813ac0)	0	75.00 GB	Not supported	Flash	512n	N


Manage Columns Export 3 items

CANCEL

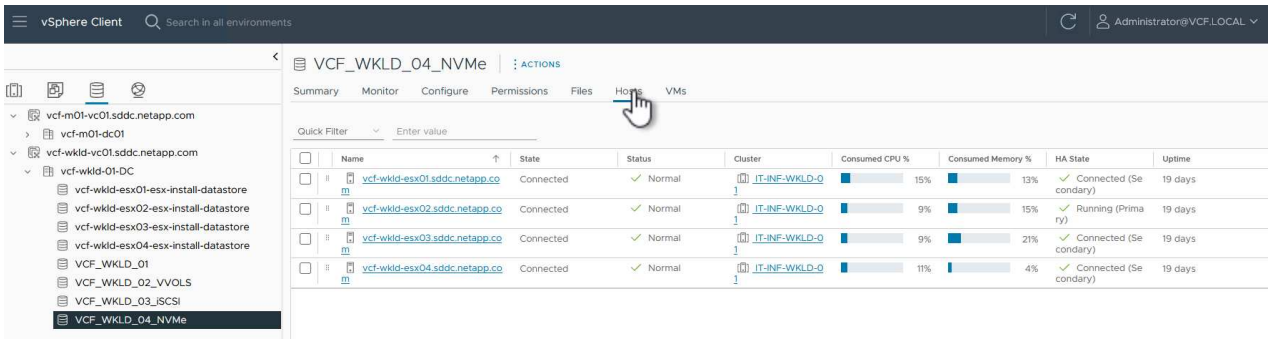
BACK

NEXT

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.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

Local file

scv-5.OP2-240310_1514.ova

- 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.
- On the **Select a compute resource** page, select the management domain cluster or specific ESXi host within the cluster to install the VM to.
- Review information pertaining to the OVF template on the **Review details** page and agree to the licensing terms on the **Licensing agreements** page.
- 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.

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

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine

Select virtual disk format

VM Storage Policy

Disable Storage DRS for this virtual machine

Name	Storage Compatibility	Capacity	Provisioned	Free	
<input checked="" type="radio"/> mgmt_01_iscsi	--	3 TB	3.71 TB	2.5 TB	
<input type="radio"/> vcf-m01-cl01-ds-vsant01	--	999.97 GB	49.16 GB	957.54 GB	
<input type="radio"/> vcf-m01-esx01-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	
<input type="radio"/> vcf-m01-esx02-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	
<input type="radio"/> vcf-m01-esx03-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	
<input type="radio"/> vcf-m01-esx04-esx-install-datastore	--	25.75 GB	4.56 GB	21.19 GB	

Manage Columns Items per page: 10 6 items

Compatibility

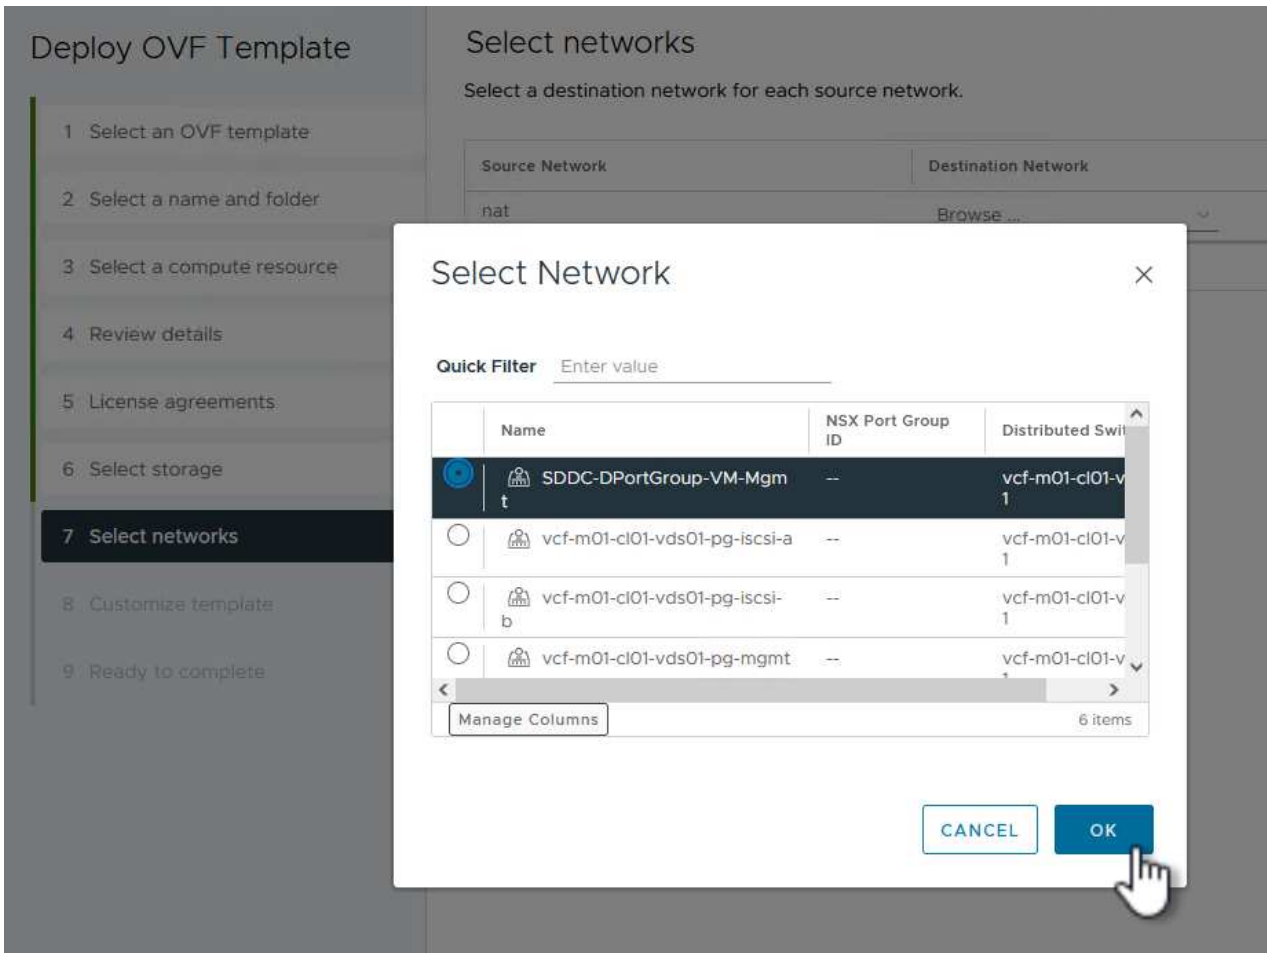
✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

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.

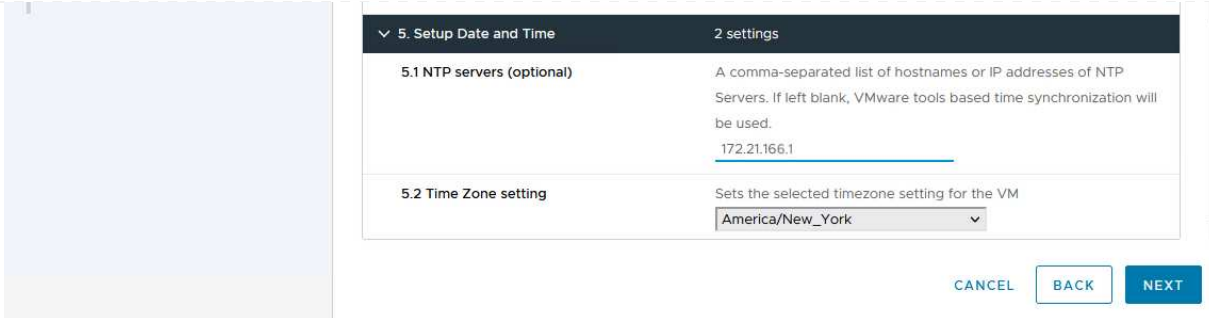
1. Register to existing vCenter		4 settings
1.1 vCenter Name(FQDN) or IP Address	<input type="text" value="cf-wkld-vc01.sddc.netapp.com"/>	
1.2 vCenter username	<input type="text" value="administrator@vcf.local"/>	
1.3 vCenter password	Password	<input type="password" value="....."/>
	Confirm Password	<input type="password" value="....."/>
1.4 vCenter port	<input type="text" value="443"/>	
2. Create SCV Credentials		2 settings
2.1 Username	<input type="text" value="admin"/>	
2.2 Password	Password	<input type="password" value="....."/>
	Confirm Password	<input type="password" value="....."/>
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) <input type="text" value="172.21.166.148"/>	
4.2.2 IPv4 Netmask	Subnet to use on the deployed network. (Leave blank if DHCP is desired) <input type="text" value="255.255.255.0"/>	
4.2.3 IPv4 Gateway	Gateway on the deployed network. (Leave blank if DHCP is desired) <input type="text" value="172.21.166.1"/>	
4.2.4 IPv4 Primary DNS	Primary DNS server's IP address. (Leave blank if DHCP is desired) <input type="text" value="10.61.185.231"/>	
4.2.5 IPv4 Secondary DNS	Secondary DNS server's IP address. (optional - Leave blank if DHCP is desired) <input type="text" value="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) <input type="text" value="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) <input type="text"/>	
4.3.2 IPv6 PrefixLen	Prefix length to use on the deployed network. (Leave blank if DHCP is desired) <input type="text"/>	




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.

Add Storage System



Storage System	<input type="text" value="172.16.9.25"/>
Authentication Method	<input checked="" type="radio"/> Credentials <input type="radio"/> Certificate
Username	<input type="text" value="admin"/>
Password	<input type="password" value="••••••••"/>
Protocol	<input type="text" value="HTTPS"/>
Port	<input type="text" value="443"/>
Timeout	<input type="text" value="60"/> Seconds
<input type="checkbox"/> Preferred IP	<input type="text" value="Preferred IP"/>

Event Management System(EMS) & AutoSupport Setting

- Log Snapcenter server events to syslog
- Send AutoSupport Notification for failed operation to storage system

CANCEL

ADD




- 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

Description

Frequency

Locking Period Enable Snapshot Locking ⓘ

Retention ⓘ

Replication Update SnapMirror after backup ⓘ
 Update SnapVault after backup ⓘ

Snapshot label

Advanced ▾ VM consistency ⓘ
 Include datastores with independent disks

Scripts ⓘ

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 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.

Create Resource Group

1. General info & notification

2. Resource

3. Spanning disks

4. Policies

5. Schedules

6. Summary

Scope: Virtual Machines

Parent entity: VCF_WKLD_03_iSCSI

Enter available entity name

Available entities

OracleSrv_01
OracleSrv_02
OracleSrv_03
OracleSrv_04

Selected entities

SQLSRV-01
SQLSRV-02
SQLSRV-03
SQLSRV-04

BACK

NEXT

FINISH

CANCEL

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

Create Resource Group

✓ 1. General info & notification

✓ 2. Resource

3. Spanning disks

4. Policies

5. Schedules

6. Summary

Always exclude all spanning datastores

This means that only the datastores directly added to the resource group and the primary datastore of VMs directly added to the resource group will be backed up

Always include all spanning datastores

All datastores spanned by all included VMs are included in this backup

Manually select the spanning datastores to be included ⓘ

You will need to modify the list every time new VMs are added

There are no spanned entities in the selected virtual entities list.

BACK

NEXT

FINISH

CANCEL

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.

Create Resource Group

- ✓ 1. General info & notification
- ✓ 2. Resource
- ✓ 3. Spanning disks
- 4. Policies**
- 5. Schedules
- 6. Summary

+ Create

<input type="checkbox"/>	Name	VM Consistent	Include independent di...	Schedule
<input checked="" type="checkbox"/>	Daily_Snapmirror	No	No	Daily
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				

BACK NEXT FINISH CANCEL

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.

Create Resource Group

✓ 1. General info & notification

✓ 2. Resource

✓ 3. Spanning disks

✓ 4. Policies

5. Schedules

6. Summary

Daily_Snapmi... ▼

Type

Daily

Every

1

Day(s)

Starting

04/04/2024



At

04



45



PM



BACK

NEXT

FINISH

CANCEL



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

Create Resource Group

- 1. General info & notification
- 2. Resource
- 3. Spanning disks
- 4. Policies
- 5. Schedules
- 6. Summary**


Name	SQL_Servers						
Description							
Send email	Never						
Latest Snapshot name	None ⓘ						
Custom snapshot format	None ⓘ						
Entities	SQLSRV-01, SQLSRV-02, SQLSRV-03, SQLSRV-04						
Spanning	False						
Policies	<table><thead><tr><th>Name</th><th>Frequency</th><th>Snapshot Locking Period</th></tr></thead><tbody><tr><td>Daily_Snapmir...</td><td>Daily</td><td>-</td></tr></tbody></table>	Name	Frequency	Snapshot Locking Period	Daily_Snapmir...	Daily	-
Name	Frequency	Snapshot Locking Period					
Daily_Snapmir...	Daily	-					

BACK NEXT **FINISH** CANCEL

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

The screenshot shows the vSphere Client interface. At the top, there is a search bar and the text "vSphere Client". Below that, it says "SnapCenter Plug-in for VMware vSphere" and "INSTANCE 172.21.166.148:8080". On the left, there is a navigation menu with options: Dashboard, Settings, **Resource Groups**, Policies, Storage Systems, and Guest File Restore. The main area shows the "Resource Groups" page. At the top of this page, there are action buttons: Create, Edit, Delete, **Run Now** (highlighted with a hand cursor), Suspend, Resume, and Export. Below the buttons is a table with columns: Name, Description, and Policy. The table contains one entry: "SQL_Servers" with a "Daily_" policy.

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.

OracleSrv_04 | Summary Monitor Configure Permissions

Guest OS Virtual Mac

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
 - Create Resource Group
 - Add to Resource Group
 - Attach Virtual Disk(s)
 - Detach Virtual Disk(s)
 - Restore
 - File Restore

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 Run



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 original 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:**

Restore ×

1. Select backup
 2. Select scope
 3. Select location
 4. Summary

Restore scope: Entire virtual machine
Restart VM:
Restore Location:

- Original Location
(This will restore the entire VM to the original Hypervisor with the original settings. Existing VM will be unregistered and replaced with this VM.)
- Alternate Location**
(This will create a new VM on selected vCenter and Hypervisor with the customized settings.)

Destination vCenter Server: 172.21.166.143
Destination ESXi host: vcf-wkld-esx04.sddc.netapp.com
Network: vcf-wkld-01-HT-INF-WKLD-01-vds-01-pg-
VM name after restore: OracleSrv_04_restored
Select Datastore: VCF_WKLD_03_ISCSI

VCF_WKLD_03_ISCSI

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.

Restore

- ✓ 1. Select backup
- ✓ 2. Select scope
- 3. Select location**
- 4. Summary

Destination datastore	Locations
VCF_WKLD_03_iSCSI	(Primary) VCF_iSCSI:VCF_WKLD_03_iSCSI
	(Primary) VCF_iSCSI:VCF_WKLD_03_iSCSI
	(Secondary) svm_iscsi:VCF_WKLD_03_iSCSI_dest
	< >

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

Restore

- ✓ 1. Select backup
- ✓ 2. Select scope
- ✓ 3. Select location
- 4. Summary**

Virtual machine to be restored	OracleSrv_04
Backup name	VCF_WKLD_iSCSI_Datastore_04-04-2024_16.50.00.0940
Restart virtual machine	No
Restore Location	Alternate Location
Destination vCenter Server	172.21.166.143
ESXi host to be used to mount the backup	vcf-wkld-esx04.sddc.netapp.com
VM Network	vcf-wkld-01-IT-INF-WKLD-01-vds-01-pg-mgmt
Destination datastore	VCF_WKLD_03_iSCSI
VM name after restore	OracleSrv_04_restored



Change IP address of the newly created VM after restore operation to avoid IP conflict.

BACK

NEXT

FINISH

CANCEL

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

- 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
VCF_WKLD_ISCI_Datastore_04-04-20...
- Backup Successful [Job ID:15] 8 min ago
VCF_WKLD_ISCI_Datastore
- Backup Successful [Job ID:12] 13 min ago
VCF_WKLD_ISCI_Datastore
- Backup Successful [Job ID:9] 13 min ago
SQL_Servers
- Backup Successful [Job ID:6] 19 min ago
SQL_Servers

[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.sdd c.netapp.com	35%	Mount operation completed successfully.	VCF.LOCAL\Administrator	6 ms	04/04/2024, 4:58:27 P M
NetApp Restore	vcf-wkld-esx04.sdd c.netapp.com	2%	Restore operation started.	VCF.LOCAL\Administrator	10 ms	04/04/2024, 4:58:27 P M


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.

Attach Virtual Disk(s)



Click here to attach to alternate VM

Backup

(This list shows primary backups. **1** modify the filter to display primary and secondary backups.)

Name	Backup Time	Mounted	Policy	VMware Snapshot
VCF_WKLD_iSCSI_Datastore_04-17-2024_09.50.01.0218	4/17/2024 9:50:01 AM	No	Hourly_Snapmirror	No
VCF_WKLD_iSCSI_Datastore_04-17-2024_08.50.01.0223	4/17/2024 8:50:01 AM	No	Hourly_Snapmirror	No
VCF_WKLD_iSCSI_Datastore_04-17-2024_07.50.01.0204	4/17/2024 7:50:00 AM	No	Hourly_Snapmirror	No
VCF_WKLD_iSCSI_Datastore_04-17-2024_06.50.01.0194	4/17/2024 6:50:00 AM	No	Hourly_Snapmirror	No
VCF_WKLD_iSCSI_Datastore_04-17-2024_05.50.01.0245	4/17/2024 5:50:01 AM	No	Hourly_Snapmirror	No
VCF_WKLD_iSCSI_Datastore_04-17-2024_04.50.01.0231	4/17/2024 4:50:01 AM	No	Hourly_Snapmirror	No

Select disks

Virtual disk	Location
<input type="checkbox"/> [VCF_WKLD_03_iSCSI] SQLSRV-01/SQLSRV-01.vmdk	Primary:VCF_iSCSI:VCF_WKLD_03_iSCSI:VCF_WKLD_iSCSI_Datastore_04-17-2024_09.50.01.0...
<input checked="" type="checkbox"/> [VCF_WKLD_03_iSCSI] SQLSRV-01/SQLSRV-01_1.v...	Primary:VCF_iSCSI:VCF_WKLD_03_iSCSI:VCF_WKLD_iSCSI_Datastore_04-17-2024_09.50.01.0...

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

Attach Virtual Disk(s)



Click here to attach to alternate VM

Backup

(This list shows primary backups)

Name
VCF_WKLD_iSCSI_Datastore_04-17-2024_09.50.01.0218
VCF_WKLD_iSCSI_Datastore_04-17-2024_08.50.01.0223
VCF_WKLD_iSCSI_Datastore_04-17-2024_07.50.01.0204
VCF_WKLD_iSCSI_Datastore_04-17-2024_06.50.01.0194
VCF_WKLD_iSCSI_Datastore_04-17-2024_05.50.01.0245
VCF_WKLD_iSCSI_Datastore_04-17-2024_04.50.01.0231

Select disks

Virtual disk	Location
<input type="checkbox"/> [VCF_WKLD_03_iSCSI] SQLSRV-01/SQLSRV-01.vmdk	Primary:VCF_iSCSI:VCF_WKLD_03_iSCSI:VCF_WKLD_iSCSI_Datastore_04-17-2024_09.50.01.0...
<input checked="" type="checkbox"/> [VCF_WKLD_03_iSCSI] SQLSRV-01/SQLSRV-01_1.v...	Primary:VCF_iSCSI:VCF_WKLD_03_iSCSI:VCF_WKLD_iSCSI_Datastore_04-17-2024_09.50.01.0...

Time range

From

Hour Minute Second

To


Hour Minute Second

VMware snapshot

Mounted

Location

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.

Run As Credentials



Run As Name	<input type="text" value="Administrator"/>	
Username	<input type="text" value="administrator"/>	
Password	<input type="password" value="••••••••"/>	
Authentication Mode	<input type="text" value="Windows"/>	
VM Name	<input type="text"/>	



CANCEL

SAVE

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**.

Proxy VM



VM Name

SQLSRV-01

Search by ESXi Host

ESXi Host

vcf-wkld-esx04.sddc.netapp.com

Virtual Machine

SQLSRV-01


Search by Virtual Machine name

CANCEL

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.

Guest File Restore



1. Restore Scope

2. Guest Details

3. Summary

Backup Name	Start Time	End Time
SQL_Servers_04-16-2024_13.52.3...	4/16/2024 1:52:34 PM	4/16/2024 1:52:40 PM
VCF_WKLD_iscsi_Datastore_04-1...	4/16/2024 1:50:01 PM	4/16/2024 1:50:08 PM

VMDK
[VCF_WKLD_03_iscsi] SQLSRV-01/SQLSRV-01.vmdk
[VCF_WKLD_03_iscsi] SQLSRV-01/SQLSRV-01_1.vmdk

Locations
Primary:VCF_iscsi:VCF_WKLD_03_iscsi:SQL_Servers_04-16-2024_13.52.34.0329
Secondary:svm_iscsi:VCF_WKLD_03_iscsi_dest:SQL_Servers_04-16-2024_13.52.34.0329

BACK

NEXT

FINISH

CANCEL

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.

Guest File Restore



1. Restore Scope

2. Guest Details

3. Summary

Use Guest VM

Guest File Restore operation will attach disk to guest VM

Run As Name	Username	Authentication Mode
Administrator	administrator	WINDOWS

Use Guest File Restore proxy VM

Send email notification

Email send from:

Email send to:

Email subject:

BACK

NEXT

FINISH

CANCEL

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.

The screenshot shows the vSphere Client interface with the SnapCenter Plug-in for VMware vSphere. The main content area displays the 'Guest File Restore' configuration page, which includes a 'Guest Configuration' section and a 'Guest Session Monitor' table. The table has the following data:

Backup Name	Source VM	Disk Path	Guest Mount Path	Time To Expire	Browse Files
SQL_Servers_04-16-2024_13:52:34.0329	SQLSRV-01	[VCF_WKLD_03]SCSI(c-202404161419...	E1	23h:58m	

Below the table, there are sections for 'Run As Credentials' and 'Proxy Credentials', both with dropdown menus.

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

	Name	Size
<input type="checkbox"/>	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:

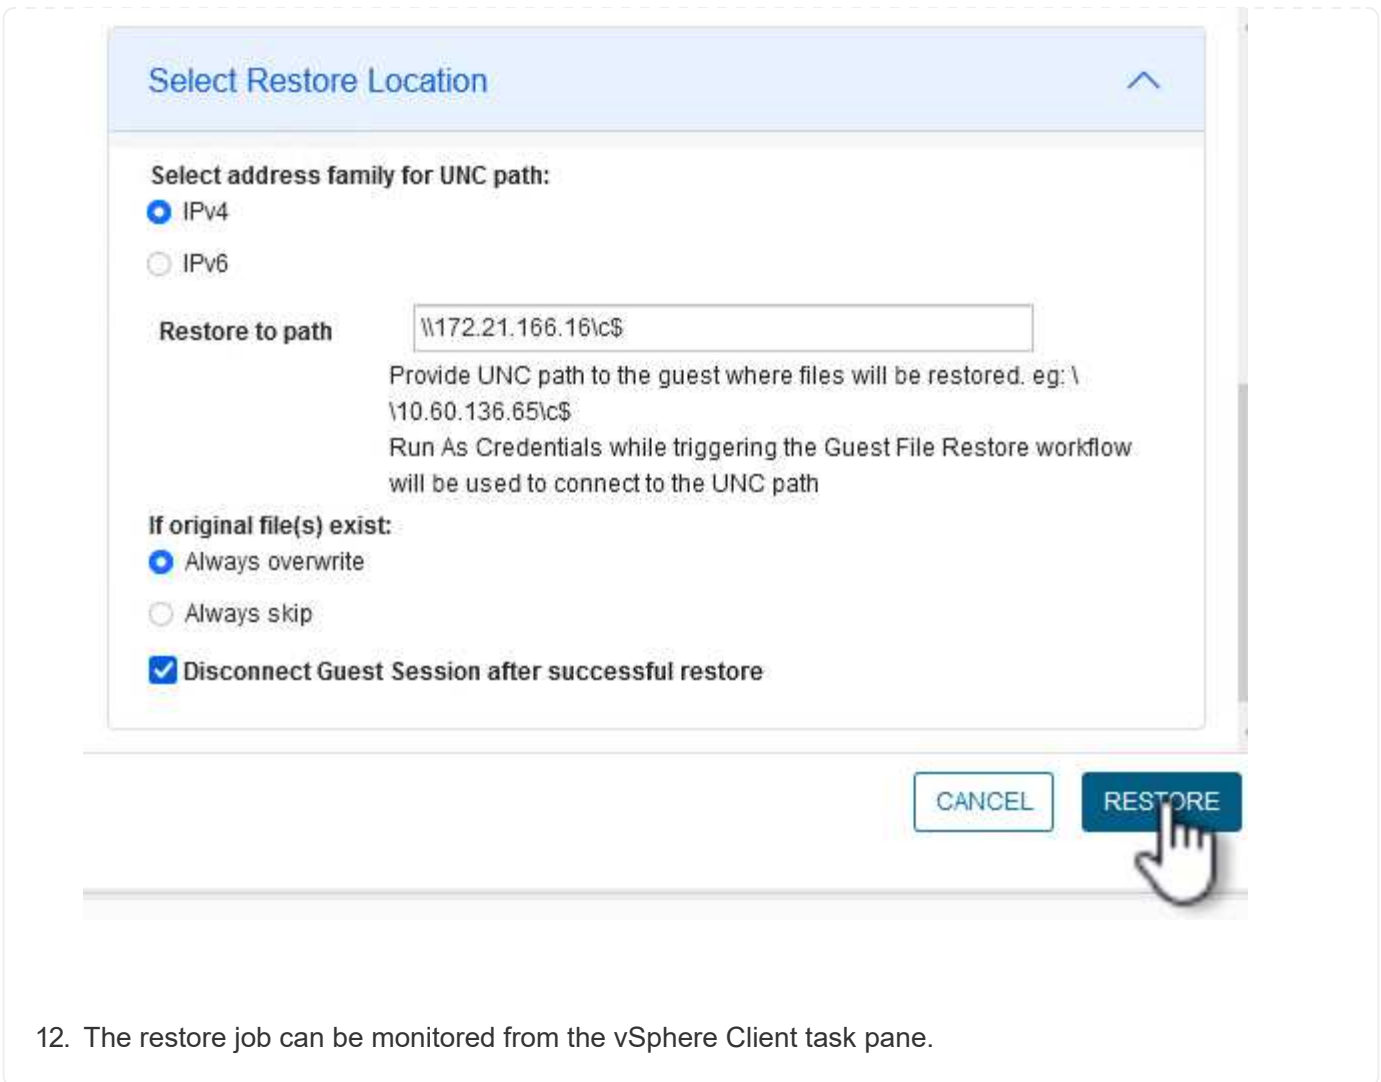
IPv4

IPv6

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.

The screenshot shows the ONTAP System Manager interface. The left-hand menu is expanded to 'STORAGE', and 'Storage VMs' is selected. The main area displays a list of Storage VMs, with 'EHC_NFS' selected. The 'Overview' tab is active, and the 'NETWORK IP INTERFACES' section is expanded, showing a list of interfaces for the 'NFS' protocol. The IP addresses 172.21.118.163 and 172.21.118.164 are highlighted with blue boxes.

Name	IP Address
EHC_NFS	172.21.118.163
EHC_NFS	172.21.118.164

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.



The screenshot displays the ONTAP System Manager interface. The left-hand navigation menu is expanded to the 'STORAGE' section, with 'Storage VMs' selected. The main content area shows a list of Storage VMs, with 'EHC_NFS' selected. The 'Settings' tab is active, and the 'Export Policies' section is visible, showing a list of policies including 'default' and 'JetStream_NFS_v02'. A hand cursor is pointing to the right arrow icon in the 'Export Policies' section, indicating the next step in the process.

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.

New export policy

NAME

WKLD_DM01

Copy rules from existing policy

STORAGE VM

svm0

EXPORT POLICY

default

RULES

No data

+ Add



Add New Rules

Save

Cancel

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.

New Rule




CLIENT SPECIFICATION

ACCESS PROTOCOLS

 SMB/CIFS FlexCache NFS NFSv3 NFSv4

ACCESS DETAILS

Type	Read-only Access	Read/Write Access	Superuser Access
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
All (As anonymous user) 	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
UNIX	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Kerberos 5	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Kerberos 5i	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Kerberos 5p	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
NTLM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Cancel

Save 




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.

- Once all rules have been entered click on the **Save** button to save the new Export Policy.
- 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.

Add Volume



NAME

VCF_WKLD_01

CAPACITY

5



TiB



STORAGE VM

EHC_NFS



Export via NFS

More Options

Cancel


Save

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.

Access Permissions

Export via NFS

GRANT ACCESS TO HOST

default 

JetStream_NFS_v04
Clients : 0.0.0.0/0 | Access protocols : Any

NFSmountTest01
3 rules

NFSmountTestReno01
Clients : 0.0.0.0/0 | Access protocols : Any

PerfTestVols
Clients : 172.21.253.0/24 | Access protocols : NFSv3, NFSv4, NFS

TestEnv_VPN
Clients : 172.21.254.0/24 | Access protocols : Any

VCF_WKLD
2 rules

WKLD_DM01
2 rules

Wkld01_NFS
Clients : 172.21.252.205, 172.21.252.206, 172.21.252.207, 172.21.2

+




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

A Network 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.

The screenshot shows the VMware Cloud Foundation interface for creating a network pool. The left sidebar contains navigation options like Dashboard, Solutions, Inventory, Lifecycle Management, Administration, Network Settings, Storage Settings, Licensing, Single Sign On, Proxy Settings, Online Depot, Composable Infrastructure, VMware Aria Suite, Backup, VMware CEP, Security, Password Management, Certificate Authority, and Developer Center. The main content area is titled 'Network Settings' and 'Create Network Pool'. It includes a breadcrumb trail: Network Pool > DNS Configuration > NTP Configuration. A note states: 'Ensure that all required networks are selected based on their usage for workload domains.' The 'Network Pool Name' is 'NFS_NPOOL'. The 'Network Type' is set to 'NFS'. Below are two configuration panels: 'NFS Network Information' and 'vMotion Network Information'. Each panel has fields for VLAN ID, MTU, Network, Subnet Mask, and Default Gateway. The NFS panel values are: VLAN ID 3374, MTU 9000, Network 172.21.118.0, Subnet Mask 255.255.255.0, and Default Gateway 172.21.118.1. The vMotion panel values are: VLAN ID 3423, MTU 9000, Network 172.21.167.0, Subnet Mask 255.255.255.0, and Default Gateway 172.21.167.1. Both panels have an 'Included IP Address Ranges' section with a warning: 'Once a network pool has been created, you are not able to edit or remove IP ranges from that pool.' The NFS panel shows a range from 172.21.118.145 to 172.21.118.148 with a 'REMOVE' button. The vMotion panel shows a range from 172.21.167.121 to 172.21.167.124 with a 'REMOVE' button. At the bottom are 'CANCEL' and 'SAVE' buttons.


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.

Checklist

Commissioning a host adds it to the VMware Cloud Foundation inventory. The host you want to commission must meet the checklist criterion below.

- Select All**
- Host for vSAN/vSAN ESA workload domain should be vSAN/vSAN ESA compliant and certified per the VMware Hardware Compatibility Guide. BIOS, HBA, SSD, HDD, etc. must match the VMware Hardware Compatibility Guide.
- Host has a standard switch with two NIC ports with a minimum 10 Gbps speed.
- Host has the drivers and firmware versions specified in the VMware Compatibility Guide.
- Host has ESXi installed on it. The host must be preinstalled with supported versions (8.0.2-22380479)
- Host is configured with DNS server for forward and reverse lookup and FQDN.
- Hostname should be same as the FQDN.
- Management IP is configured to first NIC port.
- Ensure that the host has a standard switch and the default uplinks with 10Gb speed are configured starting with traditional numbering (e.g., vmnic0) and increasing sequentially.
- Host hardware health status is healthy without any errors.
- All disk partitions on HDD / SSD are deleted.
- Ensure required network pool is created and available before host commissioning.
- Ensure hosts to be used for vSAN workload domain are associated with vSAN enabled network pool.
- Ensure hosts to be used for NFS workload domain are associated with NFS enabled network pool.
- Ensure hosts to be used for VMFS on FC workload domain are associated with NFS or VMOTION only enabled network pool.
- Ensure hosts to be used for vVol FC workload domain are associated with NFS or VMOTION only enabled network pool.
- Ensure hosts to be used for vVol NFS workload domain are associated with NFS and VMOTION only enabled network pool.
- Ensure hosts to be used for vVol iSCSI workload domain are associated with iSCSI and VMOTION only enabled network pool.
- For hosts with a DPU device, enable SR-IOV in the BIOS and in the vSphere Client (if required by your DPU vendor).

CANCEL

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.

1 Host Addition and Validation

2 Review

Add Hosts

You can either choose to add host one at a time or download [JSON](#) template and perform bulk commission.

Add new Import

Host FQDN

Storage Type vSAN NFS VMFS on FC vVol

Network Pool Name

User Name

Password

ADD

Hosts Added

Hosts added successfully. Add more or confirm fingerprint and validate host

REMOVE

Confirm all Finger Prints

VALIDATE ALL

<input checked="" type="checkbox"/>	FQDN	Network Pool	IP Address	Confirm FingerPrint	Validation Status
<input checked="" type="checkbox"/>	vcf-wkld-esx01.sddc.netapp.com	NFS_NP01	172.21.166.135	<input checked="" type="checkbox"/> SHA256:CKbsinf EOG+Hz/ lpFUoFDI2tLuY FZ47WicVdp6v EGM	<input type="checkbox"/> Not Validated

1 hosts

CANCEL

NEXT

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

Hosts Added

✔ Host Validated Successfully. ✕

REMOVE Confirm all Finger Prints i VALIDATE ALL

<input checked="" type="checkbox"/>	FGDN	Network Pool	IP Address	Confirm FingerPrint	Validation Status
<input checked="" type="checkbox"/>	vcf-wkld-esx04.sddc.netapp.com	NFS_NP01 i	172.21.166.138	✔ SHA256:9Kg+9nQaE4SQkOMsQPON/k5gZB9zyKN+6CBPmXsvLBc	✔ Valid
<input checked="" type="checkbox"/>	vcf-wkld-esx03.sddc.netapp.com	NFS_NP01 i	172.21.166.137	✔ SHA256:nPX4/mei/2zmLJHfmPwbk6zhapoUxV2IOWZDPFH+z0	✔ Valid
<input checked="" type="checkbox"/>	vcf-wkld-esx02.sddc.netapp.com	NFS_NP01 i	172.21.166.136	✔ SHA256:AMhyR60OpTQ1YYq0DJhqVbj/M/GvrQaqUy7Ce+M4IWY	✔ Valid
<input checked="" type="checkbox"/>	vcf-wkld-esx01.sddc.netapp.com	NFS_NP01 i	172.21.166.135	✔ SHA256:CKbsinfEOG+ +z/lpFUoFDI2tLuYFZ47WicVDp6vEQM	✔ Valid

CANCEL NEXT

- 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](#).

+

VI Configuration

1 General Info

2 Cluster

3 Compute

4 Networking

5 Host Selection

6 NFS Storage

7 Switch Configuration

8 License

9 Review

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.

VI Configuration

- 1 General Info
- 2 Cluster
- 3 Compute
- 4 Networking
- 5 Host Selection
- 6 NFS Storage

NFS Storage

NFS Share Details

Datastore Name ⓘ	VCF_WKLD_01
Folder ⓘ	/VCF_WKLD_01
NFS Server IP Address ⓘ	172.21.118.163

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.

VI Configuration

- 1 General Info
- 2 Cluster
- 3 Compute
- 4 Networking
- 5 Host Selection
- 6 NFS Storage
- 7 Switch Configuration
- 8 License
- 9 Review**

Review

- General**
 - Virtual Infrastructure Name: vcf-wkld-01
 - Organization Name: it-inf
 - SSO Domain Option: Joining Management SSO Domain
- Cluster**
 - Cluster Name: IT-INF-WKLD-01
- Compute**
 - vCenter IP Address: 172.21.166.143
 - vCenter DNS Name: vcf-wkld-vc01.sddc.netapp.com
 - vCenter Subnet Mask: 255.255.255.0
 - vCenter Default Gateway: 172.21.166.1
- Networking**
 - NSX Manager Instance Option: Creating new NSX instance
 - NSX Manager Cluster IP: 172.21.166.147
 - NSX Manager Cluster FQDN: vcf-w01-nsxc101.sddc.netapp.com
 - NSX Manager IP Addresses: 172.21.166.144, 172.21.166.145, 172.21.166.146

CANCEL BACK **FINISH**

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, flexibility 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.



The screenshot shows the vSphere configuration page for VMkernel adapters. The left sidebar has 'Networking' expanded, with 'VMkernel adapters' selected. The main area shows a table of VMkernel adapters with columns for Device, Network Label, Switch, IP Address, TCP/IP Stack, and Enabled Services.

Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services
vmk0	Mgmt 181	DSwitch	10.61.181.213	Default	Management
vmk1	vSAN 3376	DSwitch	172.21.120.103	Default	vSAN
vmk2	vMotion 3373	DSwitch	172.21.117.113	Default	vMotion +2
vmk3	ISCSI A - 1172	DSwitch	10.63.172.91	Default	--
vmk4	ISCSI B - 1172	DSwitch	10.63.172.92	Default	--
vmk5	Data A - 3374	DSwitch	172.21.118.123	Default	--
vmk6	VLAN 3418	DSwitch	172.21.162.103	Default	Provisioning

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



The screenshot shows the 'vmk2 - Edit Settings' dialog box. The 'Port properties' tab is active. Under 'Available services', the 'Enabled services' section has checkboxes for vMotion, Provisioning, Fault Tolerance logging, Management, vSphere Replication, vSphere Replication NFC, vSAN, vSAN Witness, vSphere Backup NFC, and NVMe over TCP. The vMotion and Fault Tolerance logging checkboxes are checked.

vmk2 - Edit Settings | esxi-hc-03.sddc.netapp.com

Port properties

TCP/IP stack: Default

MTU (Bytes): 9000

Available services

Enabled services

- vMotion
- Provisioning
- Fault Tolerance logging
- Management
- vSphere Replication
- vSphere Replication NFC
- vSAN
- vSAN Witness
- vSphere Backup NFC
- NVMe over TCP
- NVMe over RDMA

CANCEL OK



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.

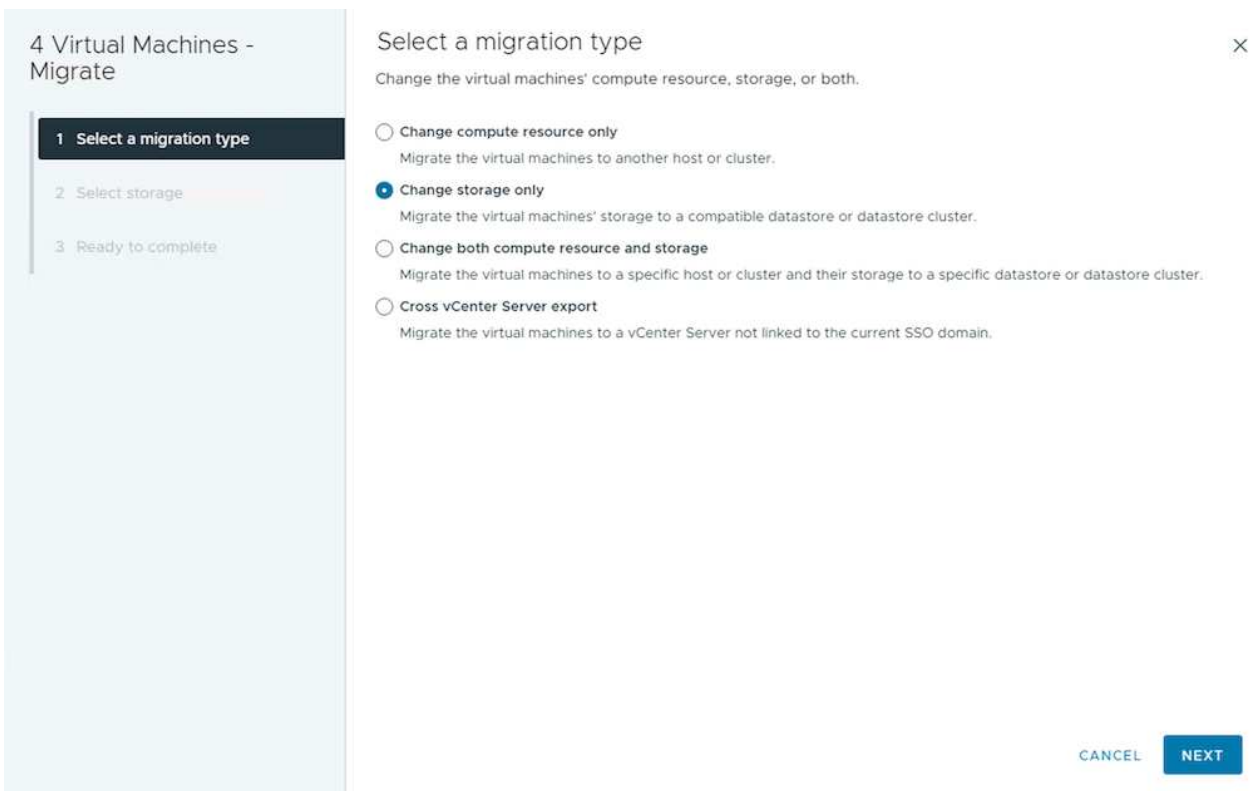
1. With vSphere Web Client, select the Datastore from the storage inventory and click on VMs tab.



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.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select storage
- 3 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

BATCH CONFIGURE CONFIGURE PER DISK

Select virtual disk format: Thin Provision

VM Storage Policy: NetApp Storage

Disable Storage DRS for this virtual machine

Name	Storage Compatibility	Capacity	Provisioned	Free
ASA_VVOLS_1	Compatible	1.95 TB	34.38 GB	1.95 TB
DemoDS	Incompatible	800 GB	7.23 GB	792.77 GB
destination	Incompatible	250 GB	31.8 MB	249.97 GB
DRaaSTest	Incompatible	1 TB	201.13 GB	880.86 GB
E13A400_JCSI	Incompatible	2 TB	858.66 GB	1.85 TB

Manage Columns | Items per page: 5 | 1 - 5 of 14 items

Compatibility

✓ Compatibility checks succeeded.

CANCEL
BACK
NEXT

5. Review and click on Finish.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select storage
- 3 Ready to complete

Ready to complete

Verify that the information is correct and click Finish to start the migration.

Migration Type	Change storage. Leave VM on the original compute resource
Virtual Machine	Migrating 4 VMs
Storage	ASA_VVOLS_1
VM storage policy	NetApp Storage
Disk Format	Thin Provision

CANCEL
BACK
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
```

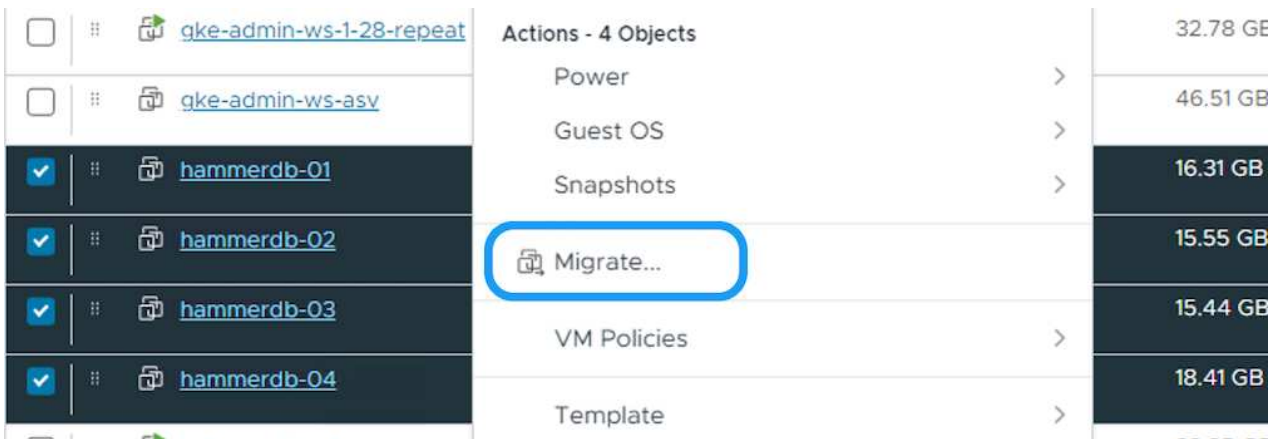
Migration of VMs in same vSphere cluster

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

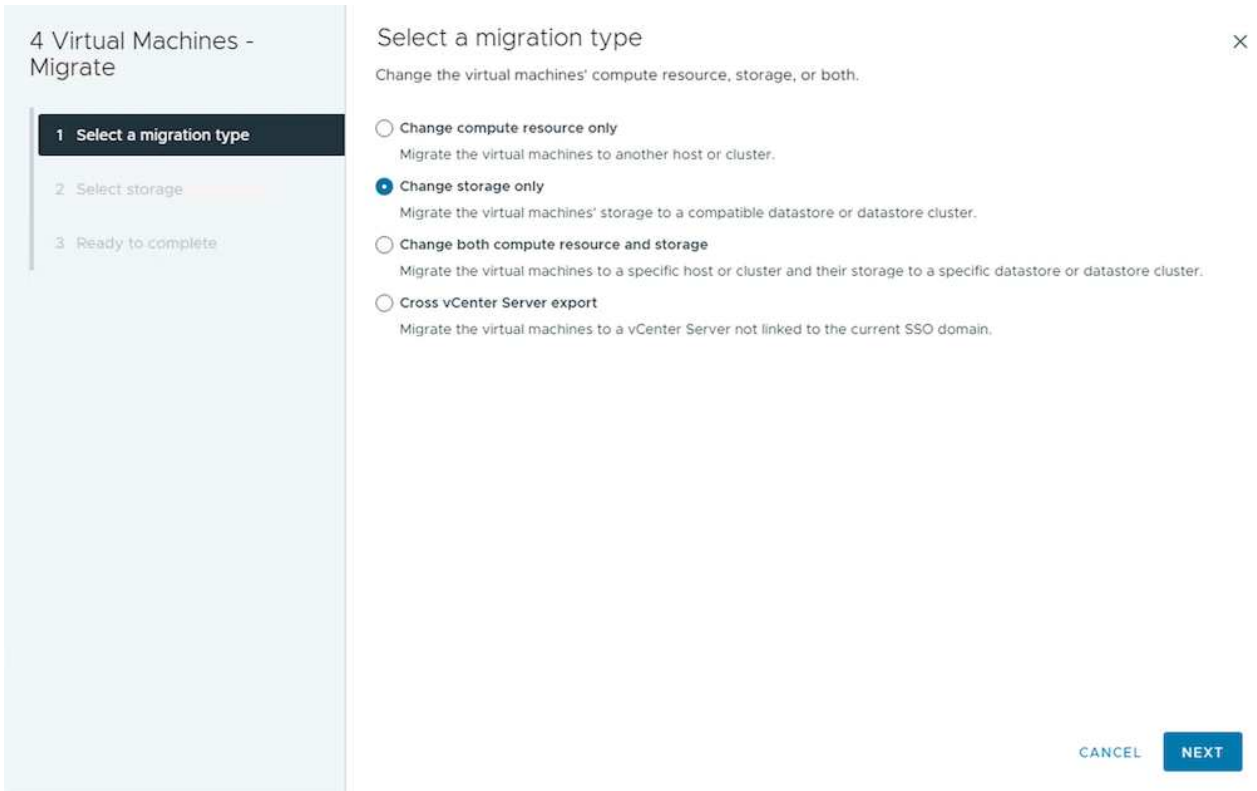
1. With vSphere Web Client, select the Cluster from the Host and Cluster inventory and click on VMs tab.



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.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select storage
- 3 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

BATCH CONFIGURE **CONFIGURE PER DISK**

Select virtual disk format: Thin Provision

VM Storage Policy: **NetApp Storage**

Disable Storage DRS for this virtual machine

Name	Storage Compatibility	Capacity	Provisioned	Free
ASA_VVOLS_1	Compatible	1.95 TB	34.38 GB	1.95 TB
DemoDS	Incompatible	800 GB	7.23 GB	792.77 GB
destination	Incompatible	250 GB	31.8 MB	249.97 GB
DRaaSTest	Incompatible	1 TB	201.13 GB	880.86 GB
E13A400_JCSI	Incompatible	2 TB	858.66 GB	1.85 TB

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

5. Review and click on Finish.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select storage
- 3 Ready to complete

Ready to complete

Verify that the information is correct and click Finish to start the migration.

Migration Type	Change storage. Leave VM on the original compute resource
Virtual Machine	Migrating 4 VMs
Storage	ASA_VVOLS_1
VM storage policy	NetApp Storage
Disk Format	Thin Provision

CANCEL BACK 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



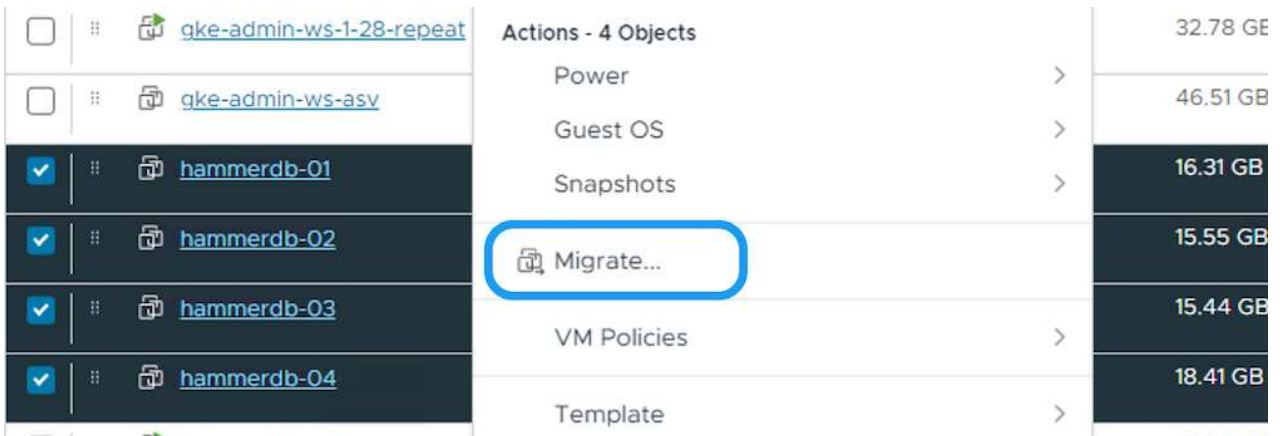
Refer [CPU Compatibility](#) and [vSphere Enhanced vMotion Compatibility](#) when source and target hosts are of different CPU family or model.

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 Virtual Machines - Migrate

1 Select a migration type

2 Select a compute resource

3 Select storage

4 Select networks

5 Select vMotion priority

6 Ready to complete

Select a migration type

Change the virtual machines' compute resource, storage, or both.

Change compute resource only

Migrate the virtual machines to another host or cluster.

Change storage only

Migrate the virtual machines' storage to a compatible datastore or datastore cluster.

Change both compute resource and storage

Migrate the virtual machines to a specific host or cluster and their storage to a specific datastore or datastore cluster.

Cross vCenter Server export

Migrate the virtual machines to a vCenter Server not linked to the current SSO domain.

CANCEL

NEXT

4. Navigate and pick the right cluster to migrate.

4 Virtual Machines - Migrate

1 Select a migration type

2 Select a compute resource

3 Select storage

4 Select networks

5 Select vMotion priority

6 Ready to complete

Select a compute resource

Select a cluster, host, vApp or resource pool to run the virtual machines.

- ▼ vcf-m01-vc01.sddc.netapp.com
 - > vcf-m01-dc01
- ▼ vcf-wkld-vc01.sddc.netapp.com
 - ▼ vcf-wkld-01-DC
 - > IT-INF-WKLD-01

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

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

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select folder
- 5 Select networks
- 6 Select vMotion priority
- 7 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

Select virtual disk format Thin Provision

VM Storage Policy NFS

Name	Storage Compatibility	Capacity	Provisioned	Free	T
<input checked="" type="radio"/> VCF_WKLD_01	Compatible	5 TB	5.91 GB	5 TB	
<input type="radio"/> VCF_WKLD_02_VVOLS	Incompatible	2.93 TB	18 MB	2.93 TB	v
<input type="radio"/> VCF_WKLD_03_ISCSI	Incompatible	3 TB	858.61 GB	2.85 TB	v
<input type="radio"/> vcf-wkld-esx01-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB	v
<input type="radio"/> vcf-wkld-esx02-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB	v
<input type="radio"/> vcf-wkld-esx03-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB	v

Items per page 10 7 items

Compatibility

✓ Compatibility checks succeeded.

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

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select folder
- 5 Select networks
- 6 Select vMotion priority
- 7 Ready to complete

Select folder

Select the destination virtual machine folder for the virtual machine migration.

Select location for the virtual machine migration.

- ✖ vcf-wkld-01-DC
 - > Discovered virtual machine
 - > vCLS

✓ Compatibility checks succeeded.

7. Select the target port group.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select folder
- 5 Select networks**
- 6 Select vMotion priority
- 7 Ready to complete

Select networks

Select destination networks for the virtual machine migration.

Migrate VM networking by selecting a new destination network for all VM network adapters attached to the same source network.

Source Network	Used By	Destination Network
SDDC-DPortGroup-VM-Mgmt	4 VMs / 4 Network adapters	vcf-wkld-01-IT-INF-WKLD-01-vds-0

ADVANCED >>

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

8. Review and click on Finish.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select storage
- 3 Ready to complete**

Ready to complete

Verify that the information is correct and click Finish to start the migration.

Migration Type	Change storage. Leave VM on the original compute resource
Virtual Machine	Migrating 4 VMs
Storage	ASA_VVOLS_1
VM storage policy	NetApp Storage
Disk Format	Thin Provision

CANCEL

BACK

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.

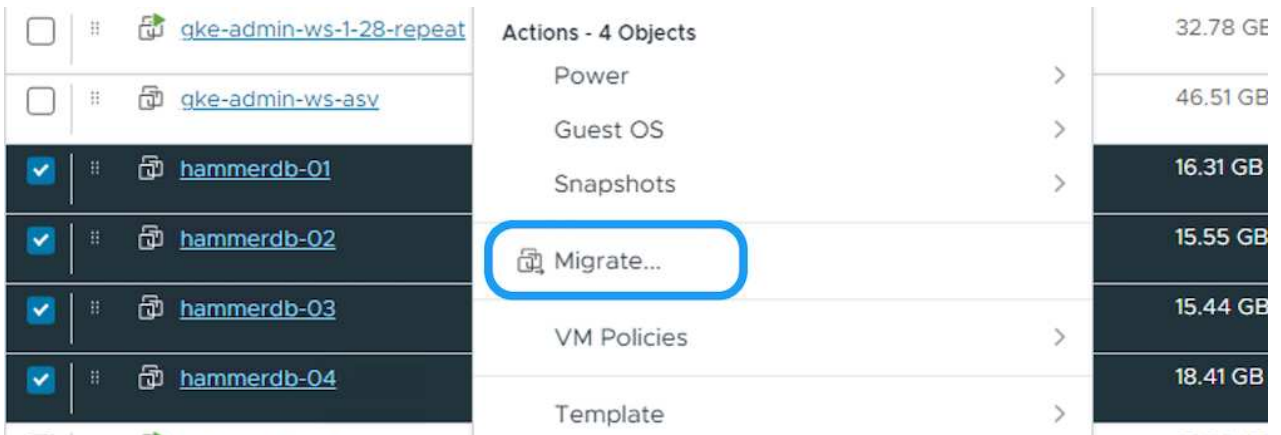


For additional requirements like source and target vCenter versions, etc., check [vSphere documentation on requirements for vMotion between vCenter server instances](#)

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 Virtual Machines - Migrate

1 Select a migration type

2 Select a compute resource

3 Select storage

4 Select networks

5 Select vMotion priority

6 Ready to complete

Select a migration type

×

Change the virtual machines' compute resource, storage, or both.

Change compute resource only

Migrate the virtual machines to another host or cluster.

Change storage only

Migrate the virtual machines' storage to a compatible datastore or datastore cluster.

Change both compute resource and storage

Migrate the virtual machines to a specific host or cluster and their storage to a specific datastore or datastore cluster.

Cross vCenter Server export

Migrate the virtual machines to a vCenter Server not linked to the current SSO domain.

CANCEL

NEXT

4. Select the target cluster in target vCenter server.

4 Virtual Machines - Migrate

1 Select a migration type

2 Select a compute resource

3 Select storage

4 Select networks

5 Select vMotion priority

6 Ready to complete

Select a compute resource

×

Select a cluster, host, vApp or resource pool to run the virtual machines.

- ▼ vcf-m01-vc01.sddc.netapp.com
 - > vcf-m01-dc01
- ▼ vcf-wkld-vc01.sddc.netapp.com
 - ▼ vcf-wkld-01-DC
 - > IT-INF-WKLD-01

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

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

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage**
- 4 Select folder
- 5 Select networks
- 6 Select vMotion priority
- 7 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

BATCH CONFIGURE **CONFIGURE PER DISK**

Select virtual disk format Thin Provision
VM Storage Policy NFS

	Name	Storage Compatibility	Capacity	Provisioned	Free
<input checked="" type="radio"/>	VCF_WKLD_01	Compatible	5 TB	5.91 GB	5 TB
<input type="radio"/>	VCF_WKLD_02_VVOLS	Incompatible	2.93 TB	18 MB	2.93 TB
<input type="radio"/>	VCF_WKLD_03_ISCSI	Incompatible	3 TB	858.61 GB	2.85 TB
<input type="radio"/>	vcf-wkld-esx01-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB
<input type="radio"/>	vcf-wkld-esx02-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB
<input type="radio"/>	vcf-wkld-esx03-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB

Manage Columns Items per page 10 7 items

Compatibility

✓ Compatibility checks succeeded.

CANCEL **BACK** **NEXT**

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

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select folder**
- 5 Select networks
- 6 Select vMotion priority
- 7 Ready to complete

Select folder

Select the destination virtual machine folder for the virtual machine migration.

Select location for the virtual machine migration.

vcf-wkld-01-DC

Discovered virtual machine

vCLS

✓ Compatibility checks succeeded.

CANCEL **BACK** **NEXT**

7. Select the target port group.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select folder
- 5 Select networks**
- 6 Select vMotion priority
- 7 Ready to complete

Select networks

Select destination networks for the virtual machine migration.

Migrate VM networking by selecting a new destination network for all VM network adapters attached to the same source network.

Source Network	Used By	Destination Network
SDDC-DPortGroup-VM-Mgmt	4 VMs / 4 Network adapters	vcf-wkld-01-IT-INF-WKLD-01-vds-0

ADVANCED >>

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

8. Review the migration options and click Finish.

4 Virtual Machines - Migrate

- 1 Select a migration type
- 2 Select storage
- 3 Ready to complete**

Ready to complete

Verify that the information is correct and click Finish to start the migration.

Migration Type	Change storage. Leave VM on the original compute resource
Virtual Machine	Migrating 4 VMs
Storage	ASA_VVOLS_1
VM storage policy	NetApp Storage
Disk Format	Thin Provision

CANCEL

BACK

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 different vCenter server using UI.

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



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



3. Choose option Cross vCenter Server export, Click Next

4 Virtual Machines - Migrate

1 Select a migration type

- 2 Select a target vCenter Server
- 3 Select a compute resource
- 4 Select storage
- 5 Select networks
- 6 Select vMotion priority
- 7 Ready to complete

Select a migration type

Change the virtual machines' compute resource, storage, or both.

- Change compute resource only**
Migrate the virtual machines to another host or cluster.
- Change storage only**
Migrate the virtual machines' storage to a compatible datastore or datastore cluster.
- Change both compute resource and storage**
Migrate the virtual machines to a specific host or cluster and their storage to a specific datastore or datastore cluster.
- Cross vCenter Server export**
Migrate the virtual machines to a vCenter Server not linked to the current SSO domain.
 - Keep VMs on the source vCenter Server (performs a VM clone operation).

CANCEL 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.

Migrate | SQLSRV-05

1 Select a migration type

- 2 **Select a target vCenter Server**
- 3 Select a compute resource
- 4 Select storage
- 5 Select networks
- 6 Ready to complete

Select a target vCenter Server

Export Virtual Machines to the selected target vCenter Server.

SAVED VCENTER SERVERS

NEW VCENTER SERVER

vCenter Server address
vCenter Server FQDN or IP address

Username
example@domain.local

Password
Password

Save vCenter Server address

LOGIN

CANCEL BACK NEXT

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

Security Alert



Unable to verify the authenticity of the external vCenter Server.

The SHA1 thumbprint of the vCenter Server certificate is:

17:42:0C:EB:82:1E:A9:86:F1:E0:70:93:AD:EB:8C:0F:27:41:F1:30



Connect anyway?

Click Yes if you trust the vCenter Server.

Click No to cancel connecting to the vCenter Server.

NO

YES

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

The screenshot shows a migration wizard window titled "Migrate | SQLSRV-05". On the left is a vertical navigation pane with six steps: "1 Select a migration type", "2 Select a target vCenter Server", "3 Select a compute resource" (which is highlighted), "4 Select storage", "5 Select networks", and "6 Ready to complete". The main area of the dialog is titled "Select a compute resource" and contains the instruction "Select a cluster, host, vApp or resource pool to run the virtual machines." Below this is a tree view showing a hierarchy: "vcf-wkld-vc01.sddc.netapp.com" (expanded) -> "vcf-wkld-01-DC" (expanded) -> "IT-INF-WKLD-01" (selected). A "VM ORIGIN" icon is visible in the top right of the tree view. Below the tree view is a "Compatibility" section with a green checkmark and the text "Compatibility checks succeeded." At the bottom right of the dialog are three buttons: "CANCEL", "BACK", and "NEXT".

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

Migrate | SQLSRV-05

- 1 Select a migration type
- 2 Select a target vCenter Server
- 3 Select a compute resource
- 4 Select storage**
- 5 Select folder
- 6 Select networks
- 7 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

VM ORIGIN ⓘ

BATCH CONFIGURE CONFIGURE PER DISK

Select virtual disk format Thin Provision ▾

VM Storage Policy NFS ▾

	Name	Storage Compatibility	Capacity	Provisioned	Free	T
<input checked="" type="radio"/>	VCF_WKLD_01	Compatible	5 TB	5.93 GB	5 TB	N
<input type="radio"/>	VCF_WKLD_02_VVOLS	Incompatible	2.93 TB	24 MB	2.93 TB	v
<input type="radio"/>	VCF_WKLD_03_JSCSI	Incompatible	3 TB	1.35 TB	2.59 TB	v
<input type="radio"/>	vcf-wkld-esx01-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB	v
<input type="radio"/>	vcf-wkld-esx02-esx-install-datastore	Incompatible	25.75 GB	3.68 GB	22.07 GB	v

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

8. Select the target VM folder.

Migrate | SQLSRV-05

- 1 Select a migration type
- 2 Select a target vCenter Server
- 3 Select a compute resource
- 4 Select storage
- 5 Select folder**
- 6 Select networks
- 7 Ready to complete

Select folder

Select the destination virtual machine folder for the virtual machine migration.

VM ORIGIN ⓘ

Select location for the virtual machine migration.

- vcf-wkld-01-DC
 - Discovered virtual machine
 - Oracle
 - SQL Server**
 - vCLS

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

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

Migrate | SQLSRV-05

- 1 Select a migration type
- 2 Select a target vCenter Server
- 3 Select a compute resource
- 4 Select storage
- 5 Select folder
- 6 Select networks
- 7 Ready to complete

Select networks

Select destination networks for the virtual machine migration. VM ORIGIN ⓘ

Migrate VM networking by selecting a new destination network for all VM network adapters attached to the same source network.

	Source Network	Used By	Destination Network
»	Mgmt 181	1 VMs / 1 Network adapters	vcf-wkld-01-IT-INF-WKLD-01-vds-01-p
»	Data A - 3374	1 VMs / 1 Network adapters	vcf-wkld-01-iscsi-a
»	Data B - 3375	1 VMs / 1 Network adapters	vcf-wkld-01-iscsi-b

3 Items

ADVANCED >>

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

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

Migrate | SQLSRV-05

- 1 Select a migration type
- 2 Select a target vCenter Server
- 3 Select a compute resource
- 4 Select storage
- 5 Select folder
- 6 Select networks
- 7 Ready to complete

Ready to complete

Verify that the information is correct and click Finish to start the migration. VM ORIGIN ⓘ

Migration Type	Change compute resource and storage
Virtual Machine	SQLSRV-05
vCenter	vcf-wkld-vc01.sddc.netapp.com
Folder	SQL Server
Cluster	IT-INF-WKLD-01
Networks	Virtual network adapters from 3 networks will be reassigned to new destination networks
Storage	VCF_WKLD_01
VM storage policy	NFS
Disk Format	Thin Provision

CANCEL BACK 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 '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.
- Continuous 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 Continuous 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.

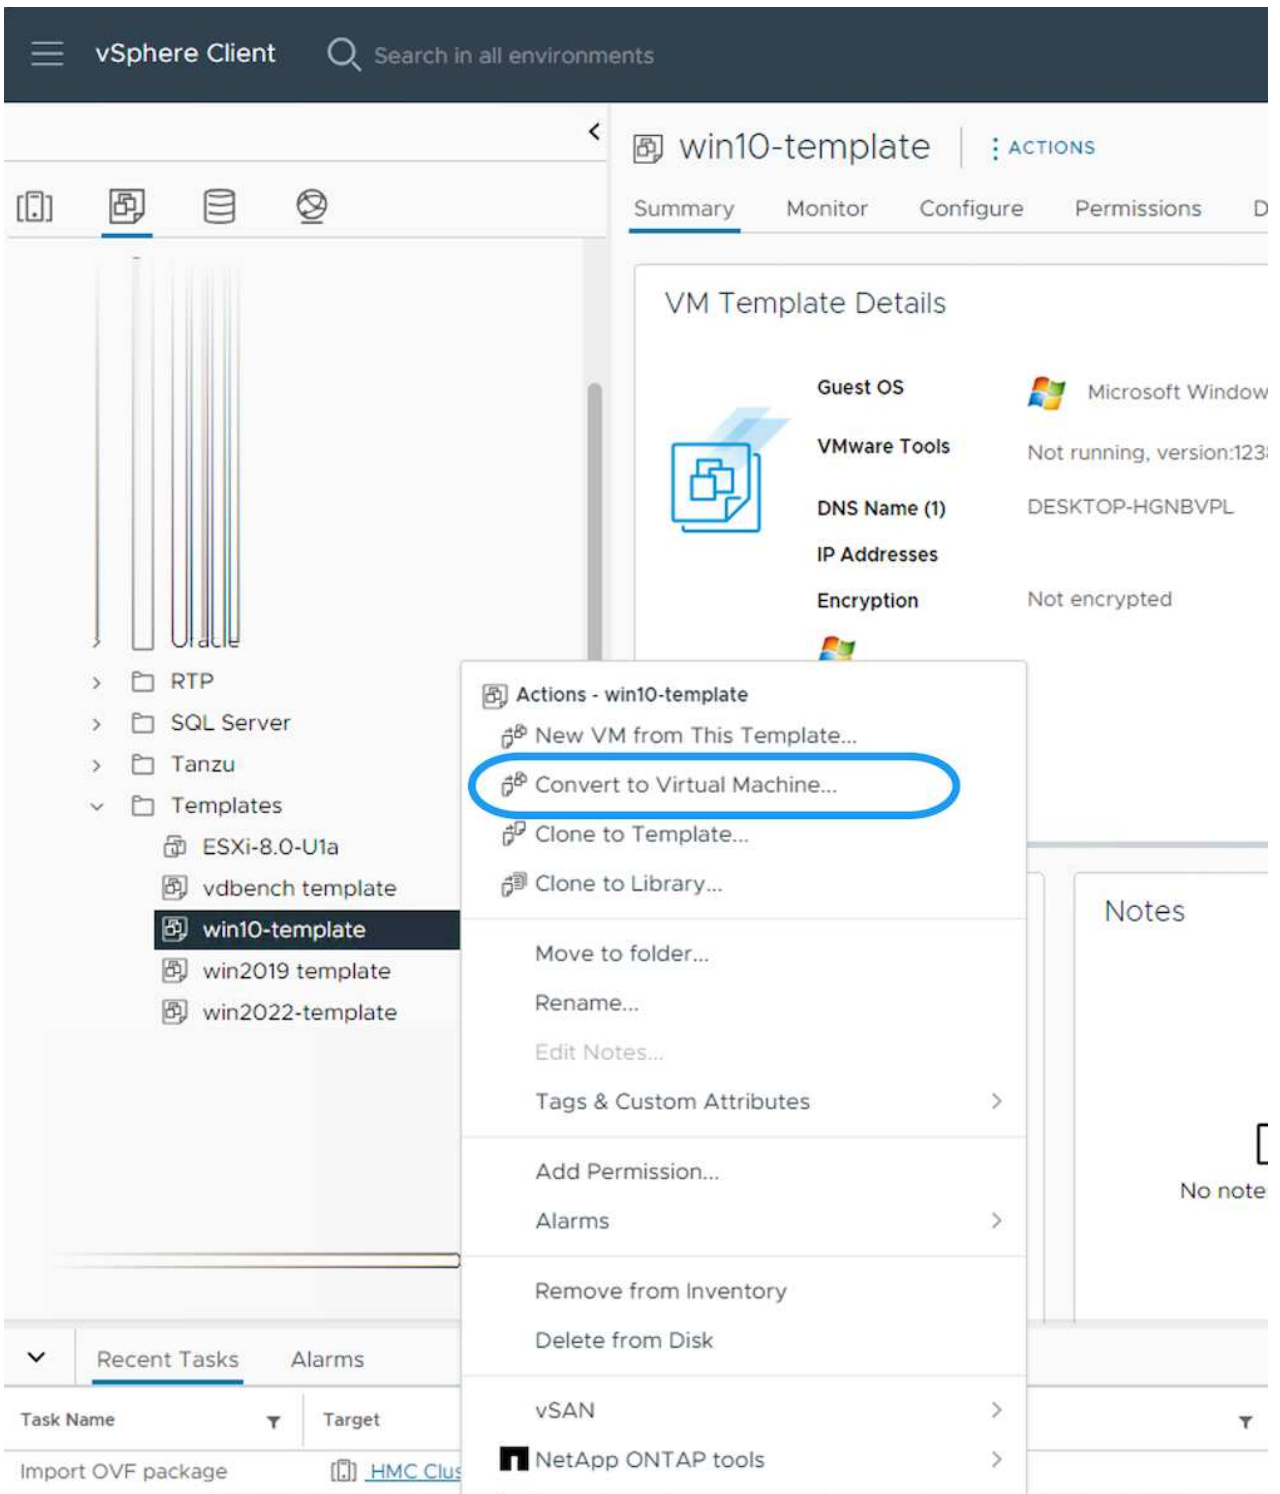
VM Template Migration Scenarios

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




 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

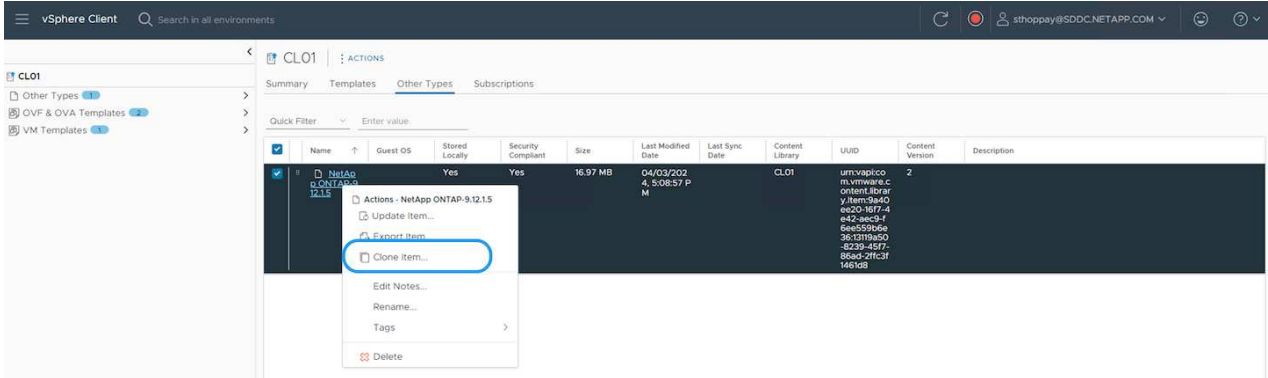
 Cloud Provider Services

 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.

Clone Library Item | NetApp ONTAP-9.12.15 ✕


Name

Notes

Select a content library where to clone the library item.

	Name	Notes	Creation Date
<input type="radio"/>	CL01		9/26/2023, 5:02:03 PM
<input checked="" type="radio"/>	CL02		4/1/2024, 12:37:51 PM

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




Here is the sample PowerCLI script to copy the content library 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-ContentLibraryItem -ContentLibrary (Get-ContentLibrary 'CL01' -Server
$sourcevc) | Where-Object { $_.ItemType -ne 'vm-template' } | Copy-
ContentLibraryItem -ContentLibrary (Get-ContentLibrary '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 library, 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.

SQLSRV-01 - Clone Virtual Machine To Template
Basic information ×

- 1 Basic information
- 2 Location
- 3 Select a compute resource
- 4 Select storage
- 5 Ready to complete

Template type VM Template

Name SQLSRV-01

Notes

Select a folder for the template

- vcsa-hc.sddc.netapp.com
 - Datacenter

CANCEL
NEXT

i

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.

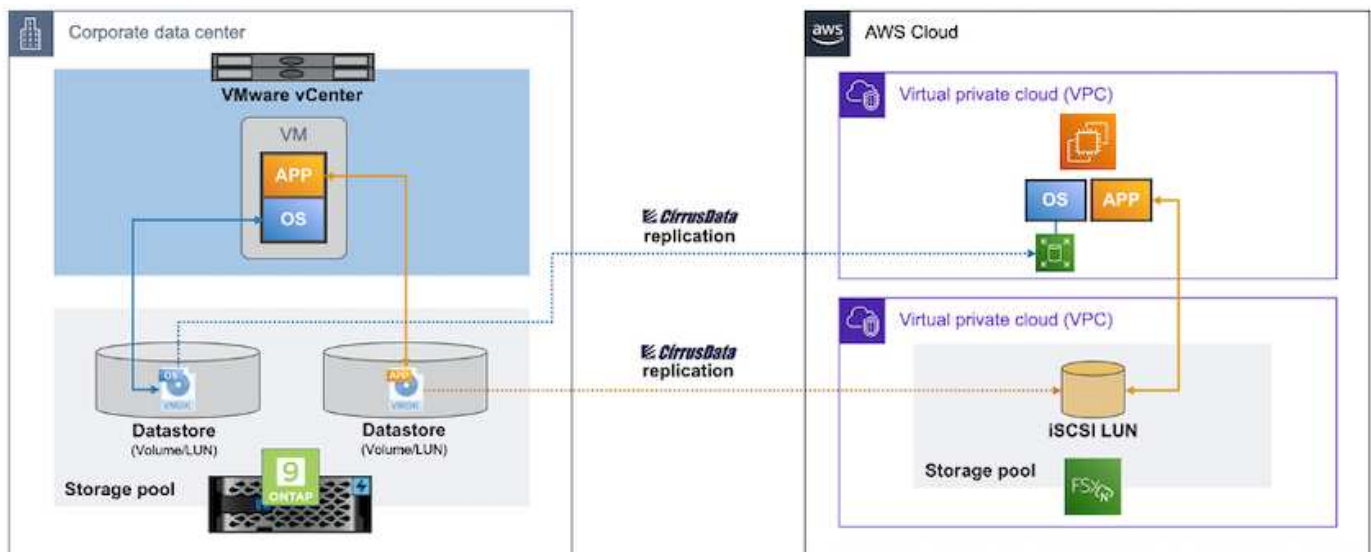
The ability of FSx for ONTAP to serve multiple protocols is another advantage, one that helps standardize a single AWS storage service for a wide range of existing data and file services requirements. For enterprises deeply invested in VMware vSphere, migrating to AWS is a cost-effective option given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

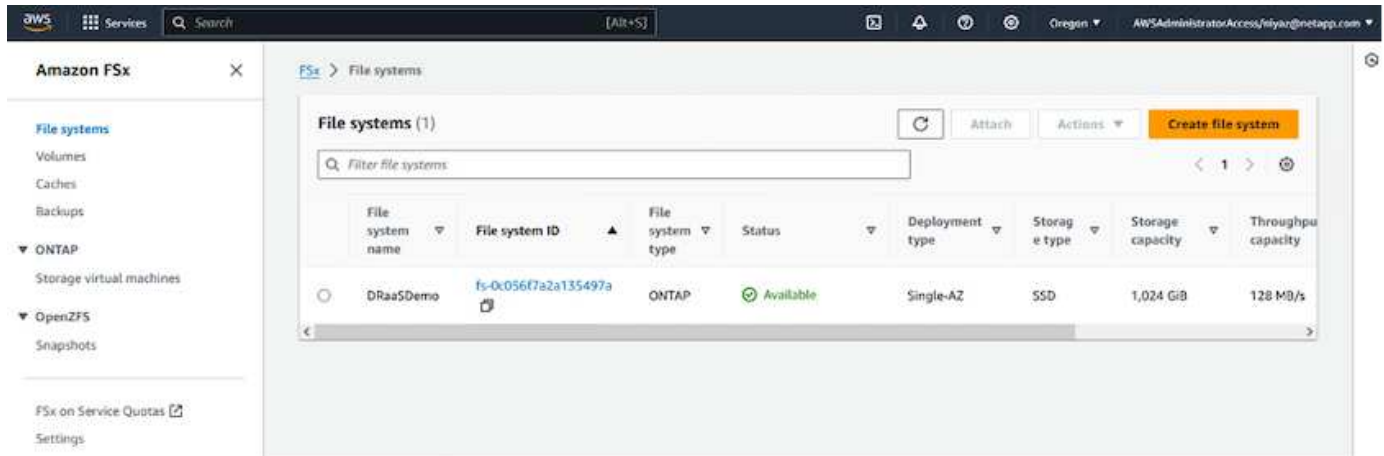
Migrate VMs to Amazon EC2 using FSxN: Deployment Guide

This article describes the deployment procedure for this migration solutions.

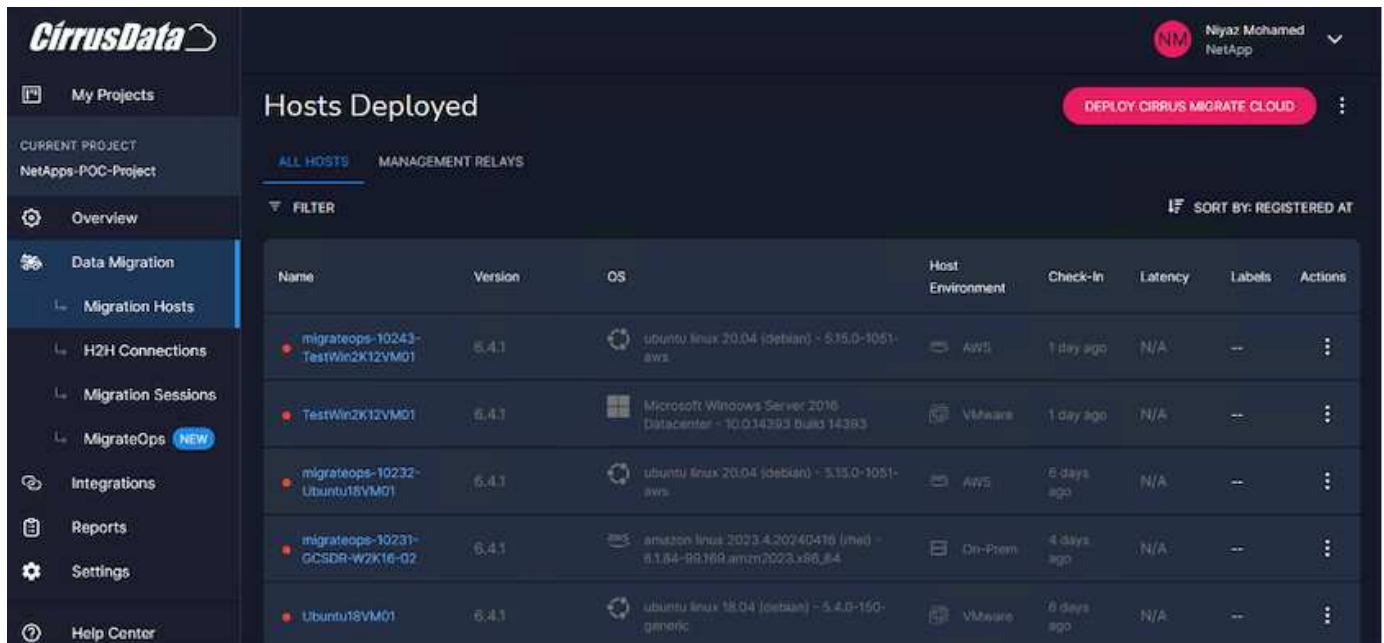
Configure FSx for ONTAP and Cirrus Data for migration operations

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, “DRaaS Demo” 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 metadata. The CMC automation takes care of all these details to create the appropriate Volume and the LUN with optional user-defined parameters.

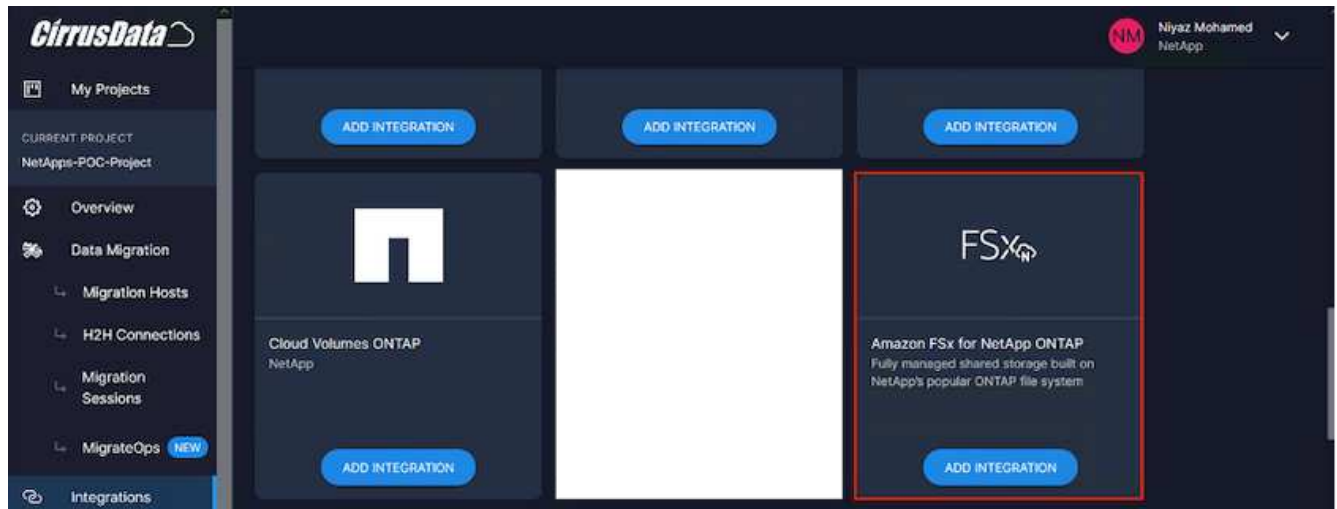
- Create Host entity (called iGroups in FSx) with the Host Initiator IQN
- Map newly created volumes to appropriate host entities using mappings
- Create all other necessary configurations

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.



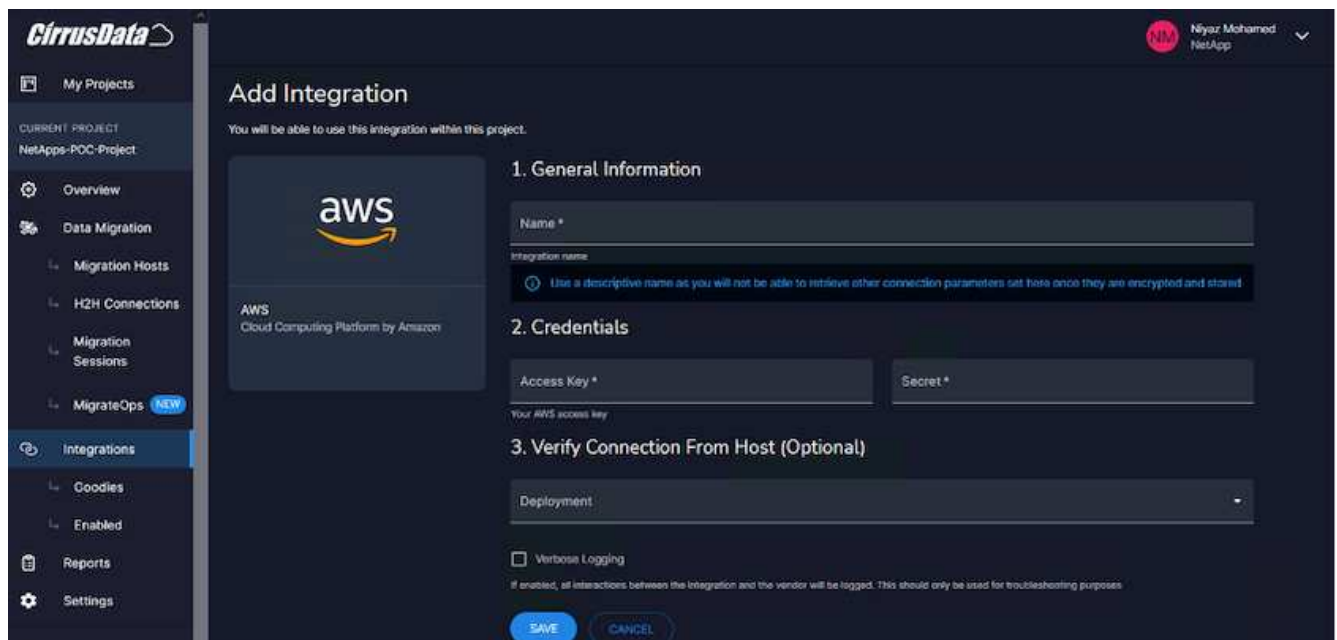
- Once the integration is created, during the creation of a new migration session, select Auto Allocate Destination Volumes to automatically allocate new volumes on FSx for ONTAP.

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.

- 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

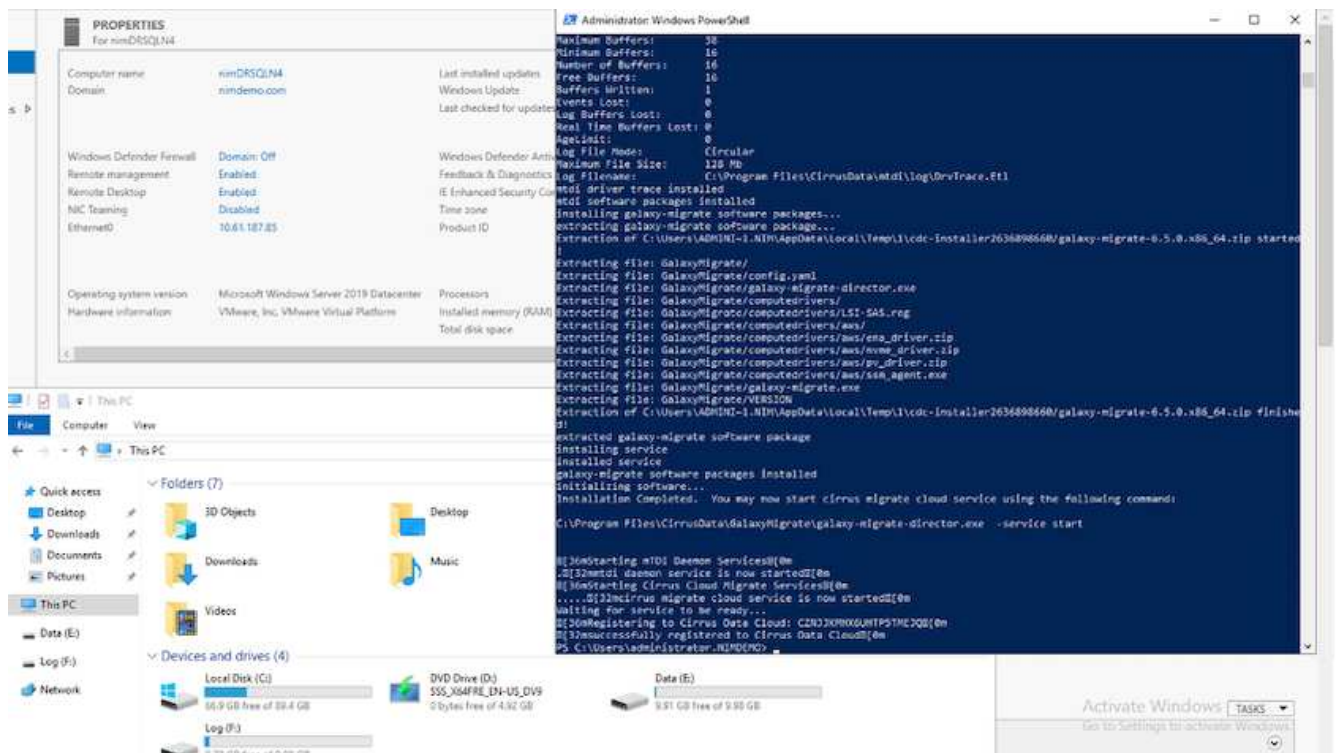
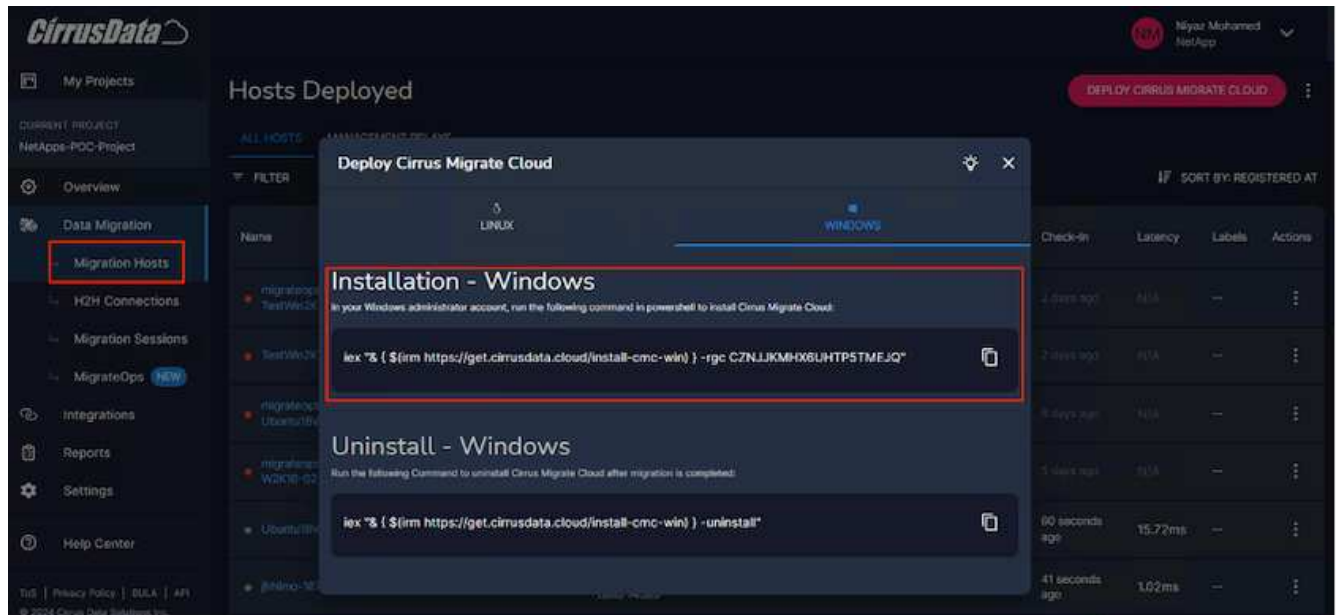
- Target instance type and ENI bandwidth

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 `lex` 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”.

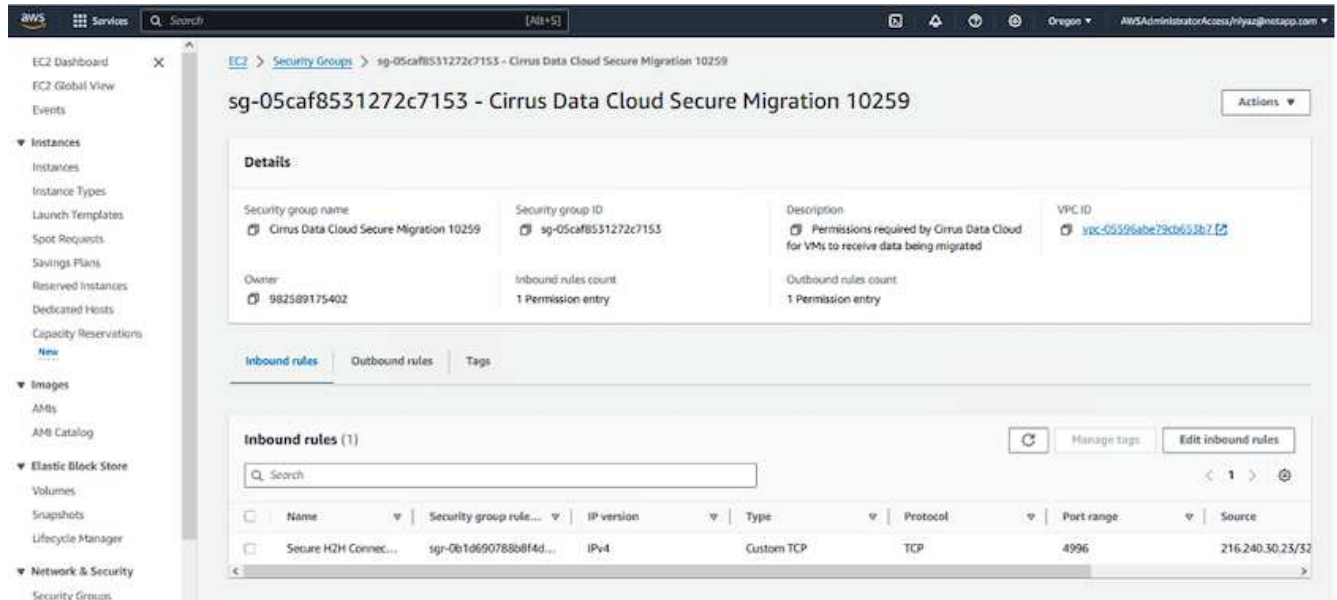
Note: To achieve parallelism, each host needs to have a YAML file specified and configured.

5. Unless the `scheduled_start_time` field is specified in the configuration, the operation will start immediately.


- 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.



- 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 highlights 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 compatibility 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 `license 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 a 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. 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 - 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 TCP/IP 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: "{{ vcllogin.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 installation 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 compatibility 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. Use the `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.

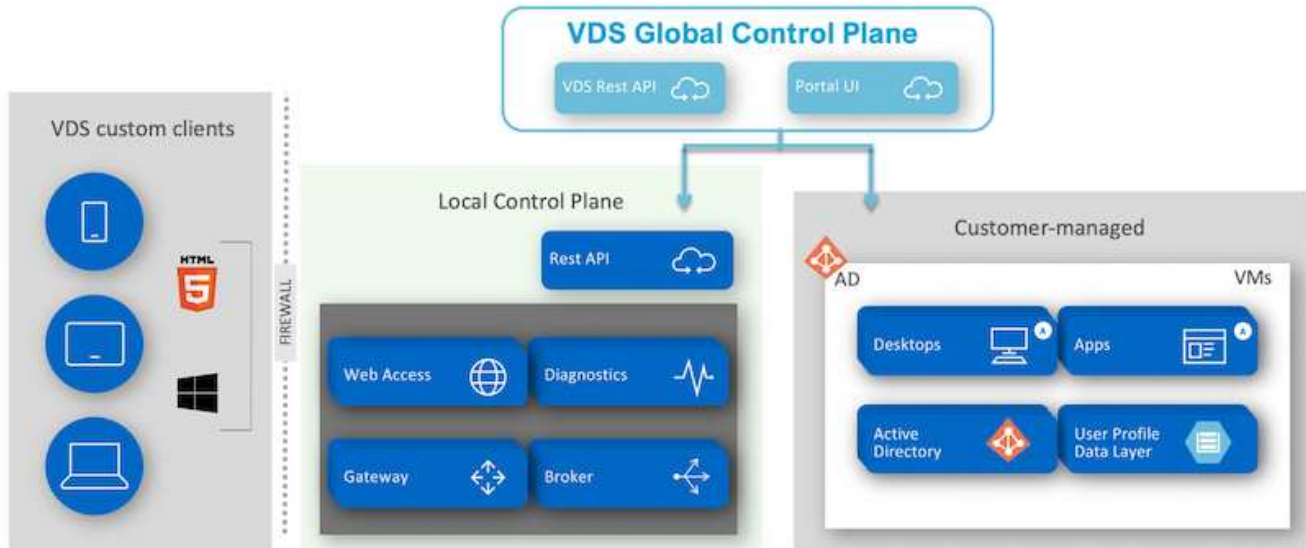
A sample deployment topology is shown in the following figure.



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.

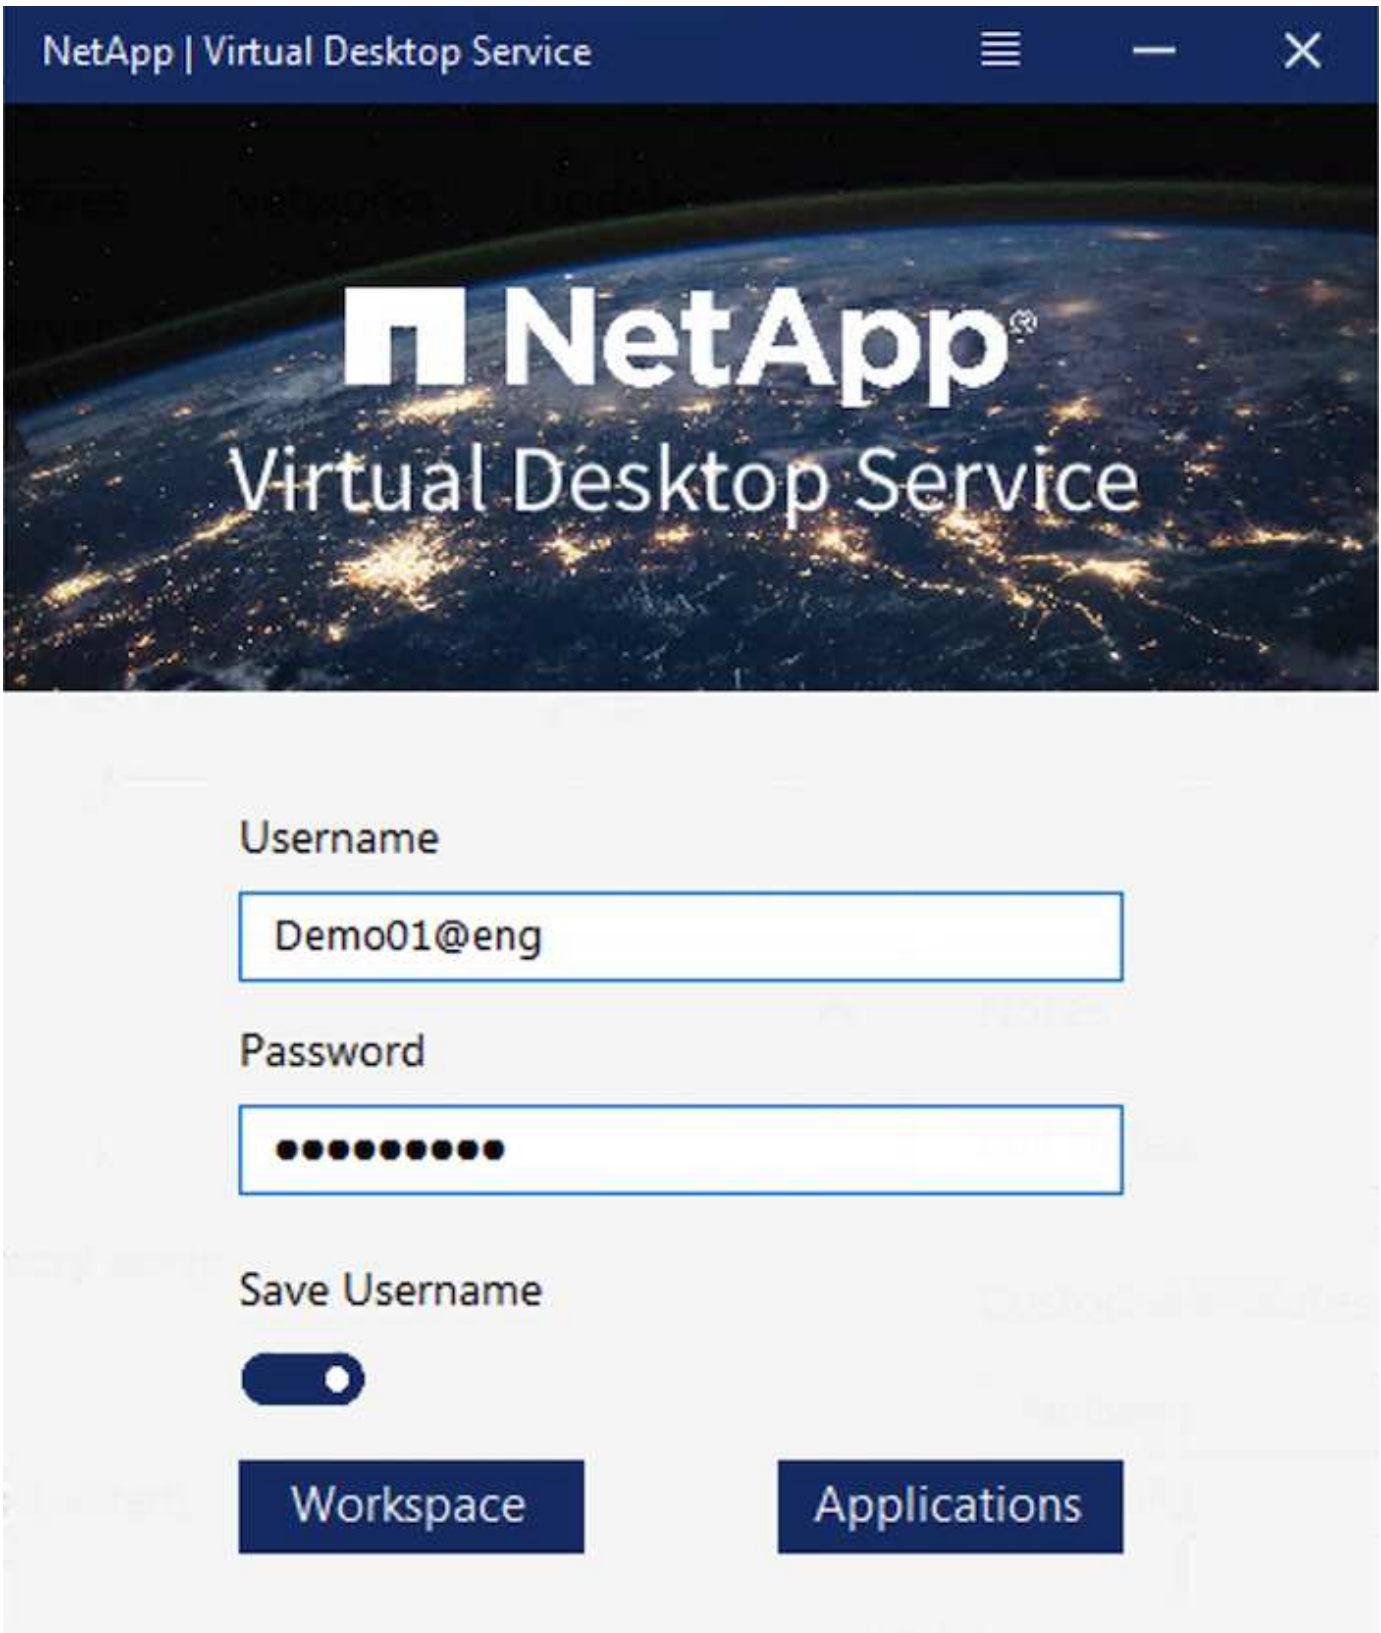
For WVD in Azure, Microsoft provides a platform-as-a-service that is consumed by NetApp VDS. For other environments, NetApp VDS orchestrates the deployment and configuration of Microsoft Remote Desktop Services. NetApp VDS supports both WVD Classic and WVD ARM and can also be used to upgrade existing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

NVIDIA GPUs Recommended for Virtualization

	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 high-end 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 mainstream 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.



NetApp Virtual Desktop Service


Deployment & AD Connect

Toby vanRoojen
Product Marketing Manager
June, 2020

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.

The screenshot shows the 'Configuration' window with the 'DataCenter Sites' tab selected. On the left, a table lists two sites: Site 1 (AzureRM, Is Primary checked) and Site 2 (vSphere, Is Primary unchecked). Below the table is a red instruction: 'To delete DataCenter Site(s), Select it and right click to delete'. The right pane shows the configuration for Site 2. The 'DataCenter Site' is 'Site 2' and the 'Hypervisor' is 'vSphere'. Under 'General Settings', the 'Local VM Account' has Username 'Administrator' and Password '*****'. The 'Hypervisor Account' has Username 'Administrator@vsphere' and Password '*****'. The 'URL' is 'https://172.21.146.150/sdk/'. Other settings include 'Vm Name Prefix', 'Max Concurrent' (20), 'Create Server', 'Subnet Mask' (255.255.255.0), and 'Default Gateway' (172.21.148.250). The 'DNS' section has 'Primary DNS' (10.67.78.11) and 'Secondary DNS' (empty). The 'Set DNS Address' is set to 'No'. The 'vSphere' section has 'Data Center' (NetApp-HCI-Datacenter), 'Cluster', 'Resource Pool', 'Host Name', and 'VM Folder' (VDS). Filtering options include 'Max VMs In Datastore' (-1), 'Min HD Free Space In Datastore GB' (-1), and 'Min Ram Free GB' (-1). At the bottom, there are checkboxes for 'Exclude vSphere DataStore' and 'Exclude vSphere ResourcePools'.

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

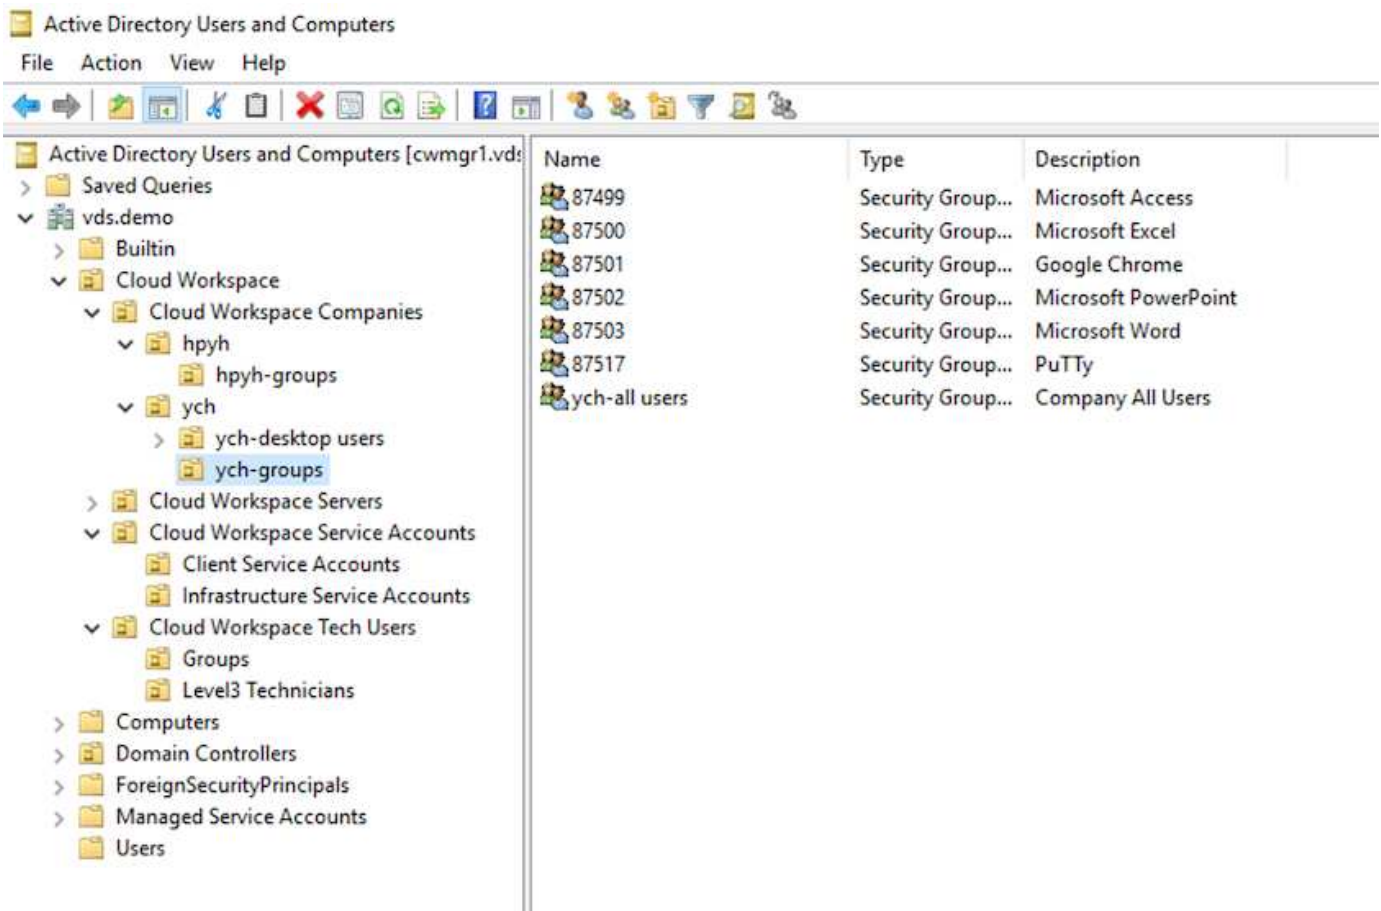
NetApp VDS uses Azure Active Directory for identity authentication and Azure Active Directory Domain Services for NTLM/Kerberos authentication. The ADConnect tool can be used to sync an on-prem Active Directory domain with Azure Active Directory.

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.



The screenshot shows a 'Security Settings' window with the following configuration:

- VDI User Enabled
- Mobile Drive Enabled
- Hypervisor Template: Windows20192899ver1
- Storage Type: DS02
- Account Expiration Enabled
- Local Drive Access Enabled
- Force Password Reset at Next Login
- Wake On Demand Enabled
- Multi-factor Auth Enabled

An 'Update' button is located at the bottom left of the settings panel.

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.

New Workspace

Client & Settings
Choose Applications
Add Users
Review & Provision

Select a Client [Add](#)

No Clients Added.

Workspace Settings

Company Name

Primary Notification Email

Application Settings

Enable Remote App

Enable App Locker

Enable Application Usage Tracking

Device Settings

Disable Printing Access

Enable Workspace User Data Storage

Security Settings

Require Complex User Password

Enable MFA for All Users

Permit Access To Task Manager

Cancel
Continue



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](#).

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 migrate 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.




The screenshot displays the NetApp Global File Cache Configuration Console. The interface includes a navigation menu with tabs for System Overview, System Configuration, GFC Configuration, and Policy Configuration. The current view is the License Manager section, specifically Legacy Licensing. The License Configuration section contains the following fields and options:

- License Configuration**
 - Associate this instance with a License Manager Server
 - License Server Public IP Address/DNS name:
 - Customer id:
 - Intended Server Role: Core Edge
 -
- License Server Configuration**
 - Bind this LMS Server to Azure Inventory (or) Verify License Information [Click here](#)


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.



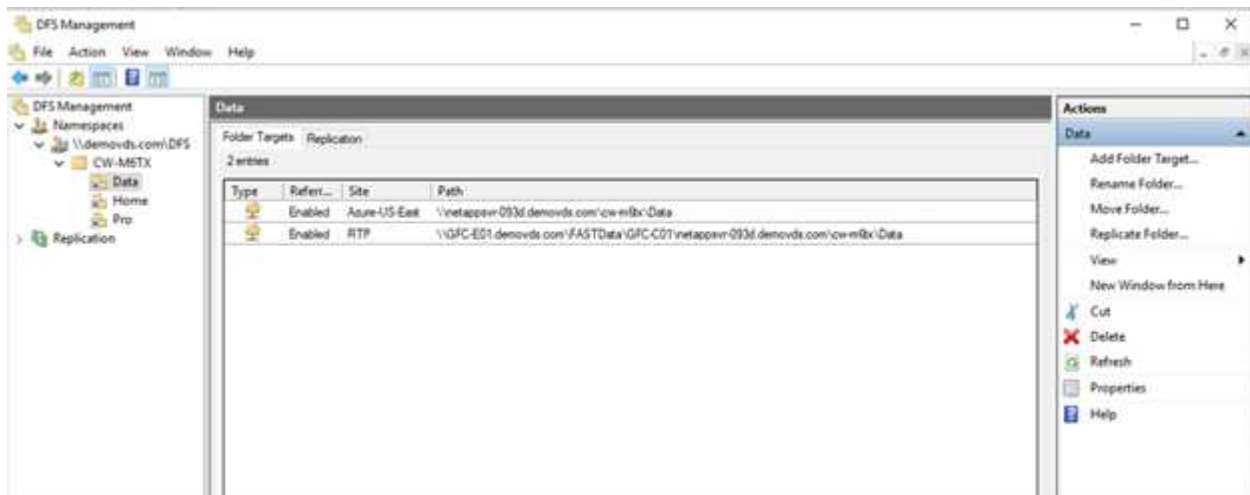
If core auto-configuration is enabled, core information is retrieved from the license management server automatically.




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.



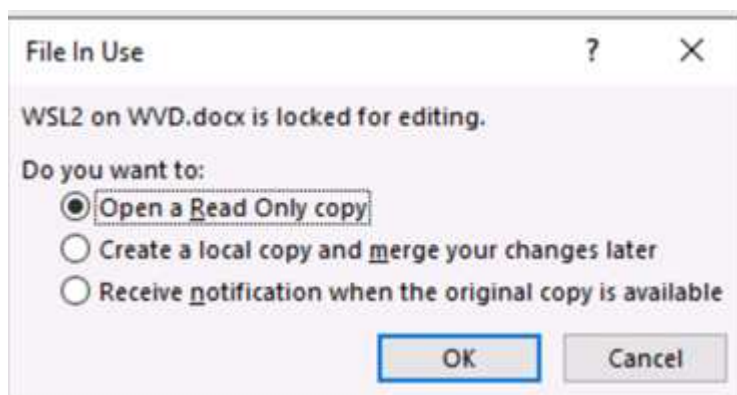
When users log in with Active Directory credentials based on the subnets associated with the site, the appropriate link is utilized by the DFS client to access the data.




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 DCCconfig Tool, TestVdc Tools, and log files.

DCCconfig Tool

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

DataCenter Site

DataCenter Site

Hypervisor

Cancel New Save

Load Hypervisor Test

Select Hypervisor

- Aws
- AzureClassic
- AzureRM
- ComputeEngine
- HyperV
- ProfitBricks
- vCloud
- vCloudRest
- vSphere
- XenServer


Configuration

DataCenter Accounts Email DatabaseConnection Exclude DataCenter Sites Product Keys Static IpAddress Drive Mapping

Save

	Description	DriveLetter
	Shared Data	P
	FTP	F
▶	User Home	H

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.

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 ▾	TSDData ▾	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.

TechPowerUp GPU-Z 2.36.0

Graphics Card | Sensors | Advanced | Validation

Name: NVIDIA Tesla M60 [Lookup](#)

GPU: GM204 Revision: FF

Technology: 28 nm Die Size: 398 mm²

Release Date: Aug 30, 2015 Transistors: 5200M

BIOS Version: 84.04.85.00.03 UEFI

Subvendor: NVIDIA Device ID: 10DE 13F2 - 10DE 115E

ROPs/TMUs: 64 / 128 Bus Interface: PCI ?

Shaders: 2048 Unified DirectX Support: 12 (12_1)

Pixel Fillrate: 75.4 GPixel/s Texture Fillrate: 150.8 GTexel/s

Memory Type: GDDR5 (Hynix) Bus Width: 256 bit

Memory Size: 8192 MB Bandwidth: 160.4 GB/s

Driver Version: 27.21.14.5257 (NVIDIA 452.57) / 2016

Driver Date: Oct 22, 2020 Digital Signature: WHQL

GPU Clock: 557 MHz Memory: 1253 MHz Boost: 1178 MHz

Default Clock: 557 MHz Memory: 1253 MHz Boost: 1178 MHz

NVIDIA SLI: Disabled

Computing OpenCL CUDA DirectCompute DirectML


Technologies Vulkan Ray Tracing PhysX OpenGL 4.6

NVIDIA Tesla M60 [Close](#)


With NetApp HCI, the H615C GPU contains three NVIDIA Tesla T4 cards. Each NVIDIA Tesla T4 card has a Turing-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.

Graphics Card		Sensors	Advanced	Validation
Name	NVIDIA GRID T4-8Q			Lookup
GPU	TU104	Revision	A1	
Technology	12 nm	Die Size	545 mm ²	
Release Date	Sep 13, 2018	Transistors	13600M	
BIOS Version	0.00.00.00.00			<input type="checkbox"/> UEFI
Subvendor	NVIDIA	Device ID	10DE 1EB8 - 10DE 130F	
ROPs/TMUs	8 / 160	Bus Interface	PCI ?	
Shaders	2560 Unified	DirectX Support	12 (12_2)	
Pixel Fillrate	4.7 GPixel/s	Texture Fillrate	93.6 GTexel/s	
Memory Type	GDDR6	Bus Width	256 bit	
Memory Size	8192 MB	Bandwidth	Unknown	
Driver Version	27.21.14.5257 (NVIDIA 452.57) / 2016			
Driver Date	Oct 22, 2020	Digital Signature	WHQL	
GPU Clock	585 MHz	Memory	0 MHz	Shader N/A
Default Clock	585 MHz	Memory	0 MHz	Shader N/A
NVIDIA SLI	Disabled			
Computing	<input checked="" type="checkbox"/> OpenCL	<input checked="" type="checkbox"/> CUDA	<input checked="" type="checkbox"/> DirectCompute	<input type="checkbox"/> DirectML
Technologies	<input checked="" type="checkbox"/> Vulkan	<input checked="" type="checkbox"/> Ray Tracing	<input type="checkbox"/> PhysX	<input checked="" type="checkbox"/> OpenGL 4.6
NVIDIA GRID T4-8Q				Close

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.

Add Storage System

 Any communication between ONTAP tools plug-in and the storage system should be mutually authenticated.

vCenter server	10.61.181.205 ▾
Name or IP address:	10.192.102.103
Username:	admin
Password:	●●●●●●●●
Port:	443

Advanced options ^

ONTAP Cluster Certificate: Automatically fetch Manually upload

CANCEL


ADD

Create a storage capability profile in ONTAP Tools

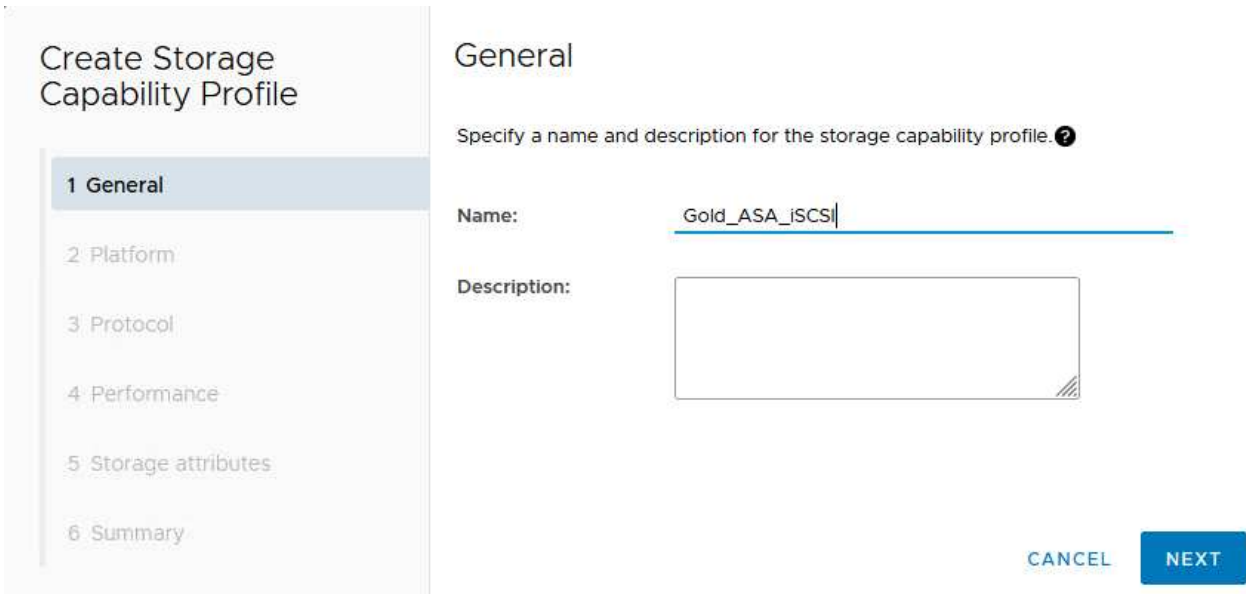
Storage capability profiles describe the features provided by a storage array or storage system. They include quality of service definitions and are used to select storage systems that meet the parameters defined in the profile.

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.

Create Storage Capability Profile

- 1 General
- 2 Platform**
- 3 Protocol
- 4 Performance
- 5 Storage attributes
- 6 Summary

Platform

Platform: Performance

Asymmetric:

CANCEL

BACK

NEXT

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

Create Storage Capability Profile

- 1 General
- 2 Platform
- 3 Protocol**
- 4 Performance
- 5 Storage attributes
- 6 Summary

Protocol

Protocol: Any

- Any
- FCP
- iSCSI
- NVMe/FC

CANCEL

BACK

NEXT

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

Create Storage Capability Profile

1 General

2 Platform

3 Protocol

4 Performance

5 Storage attributes

6 Summary

Performance

None ⓘ

QoS policy group ⓘ

Min IOPS:

Max IOPS:

6000

Unlimited

CANCEL

BACK

NEXT

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

Create Storage Capability Profile

1 General

2 Platform

3 Protocol

4 Performance

5 Storage attributes

6 Summary

Storage attributes

Deduplication:

Yes

Compression:

Yes

Space reserve:

Thin

Encryption:

No

Tiering policy (FabricPool):

None

CANCEL

BACK

NEXT

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 datastore 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.

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage system

Specify the storage capability profiles and the storage system you want to use.

Storage capability profiles:

FAS_Default
FAS_Max20
Custom profiles
Gold_ASA_JSCSI
Gold_ASA

Storage system:

HCG-NetApp-A400-E3U03 (10.192.102.103)

Storage VM:

svml

CANCEL

BACK

NEXT

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.

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage attributes

Specify the storage details for provisioning the datastore.

Volumes: Create new volumes Select volumes

Create new volumes

Name	Size	Storage Capability Profile	Aggregate
 FlexVol volumes are not added.			

Name	Size(GB) ⓘ	Storage capability profile	Aggregates	Space reserve
ASA_VVOL	2000	Gold_ASA	HCG_A400_E3u3b_NVMe	Thin

ADD

CANCEL

BACK

NEXT

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

New Datastore

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Summary

General

vCenter server: 10.61.181.205

Provisioning destination: HMC Cluster

Datastore name: ASA_VVOL

Datastore type: vVols

Protocol: iSCSI

Storage capability profile: Gold_ASA

Storage system details

Storage system: HCG-NetApp-A400-E3U03

SVM: svm1

Storage attributes

New FlexVol Name	New FlexVol Size	Aggregate	Storage Capability Profile

CANCEL
BACK
FINISH

Create a VM storage policy in the vSphere client

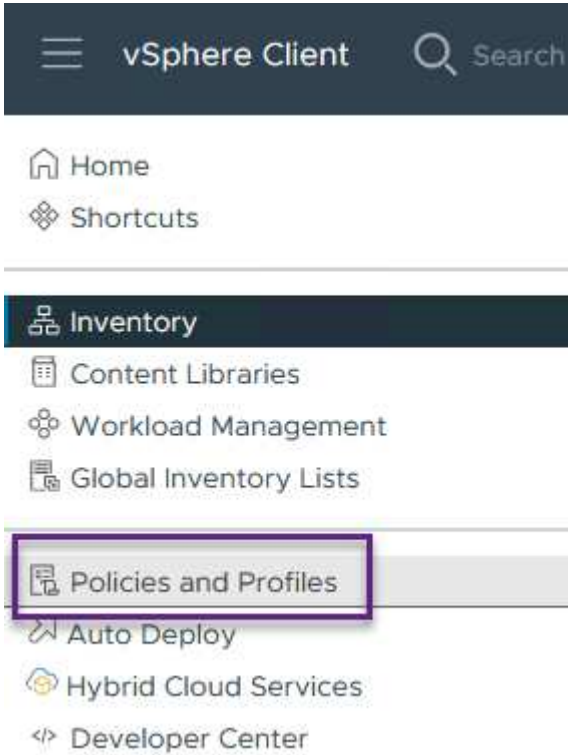
A VM storage policy is a set of rules and requirements that define how virtual machine (VM) data should be stored and managed. It specifies the desired storage characteristics, such as performance, availability, and data services, for a particular VM.

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.

A screenshot of the 'Create VM Storage Policy' wizard. The left sidebar shows a progress indicator with four steps: '1 Name and description' (active), '2 Policy structure', '3 Storage compatibility', and '4 Review and finish'. The main area is titled 'Name and description' and contains three fields: 'vCenter Server' with a dropdown menu showing 'VCSA-HC.SDDC.NETAPP.COM', 'Name' with a text input field containing 'ASA_Gold', and 'Description' with a large empty text area.

3. On the **Policy structure** page select to enable rules for NetApp clustered data ontap vVol storage and click on **Next**.

- 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.

- On the **Storage compatibility** page, review the list of vSAN datastores that match this policy and click **Next**.
- 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.

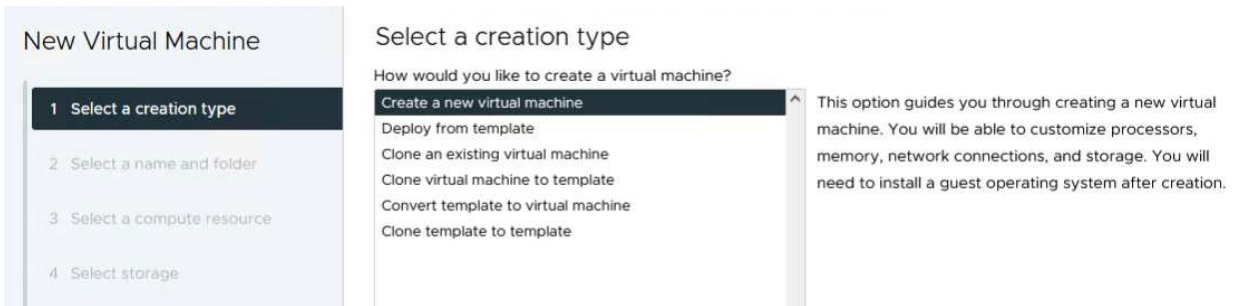
In this case, the task involves creating a VM storage policy to specify that a virtual machine will be generated

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**.

New Virtual Machine

- 1 Select a creation type
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Select storage**
- 5 Select compatibility
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine [?](#)

VM Storage Policy ASA_Gold ▾

Disable Storage DRS for this virtual machine

	Name	Storage Compatibility	Capacity	Provisioned	Free	
<input checked="" type="radio"/>	ASA_VVOLS_1	Compatible	1.95 TB	9 MB	1.95 TB	V
<input type="radio"/>	ASA400_ISCSI01	Incompatible	2 TB	185.32 GB	1.9 TB	V
<input type="radio"/>	DemoDS	Incompatible	800 GB	6.99 GB	793.01 GB	N
<input type="radio"/>	destination	Incompatible	250 GB	32.66 MB	249.97 GB	N
<input type="radio"/>	DRaaSTest	Incompatible	1 TB	133.27 GB	956.83 GB	N
<input type="radio"/>	esxi-hc-01 local	Incompatible	349.25 GB	1.41 GB	347.84 GB	V
<input type="radio"/>	esxi-hc-02 local	Incompatible	349.25 GB	1.41 GB	347.84 GB	V
<input type="radio"/>	esxi-hc-03 local	Incompatible	349.25 GB	1.41 GB	347.84 GB	V

Manage Columns Items per page: 10 1 - 10 of 15 items 1 / 2

Compatibility

Validating...

CANCEL

BACK

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 Acquisition Unit software installed on a local VM with network connectivity to assets

for data collection.

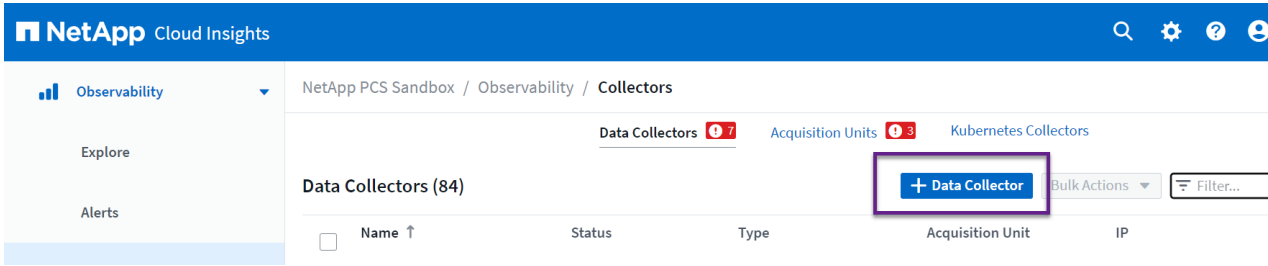
Solution Deployment

Configure Data Collectors

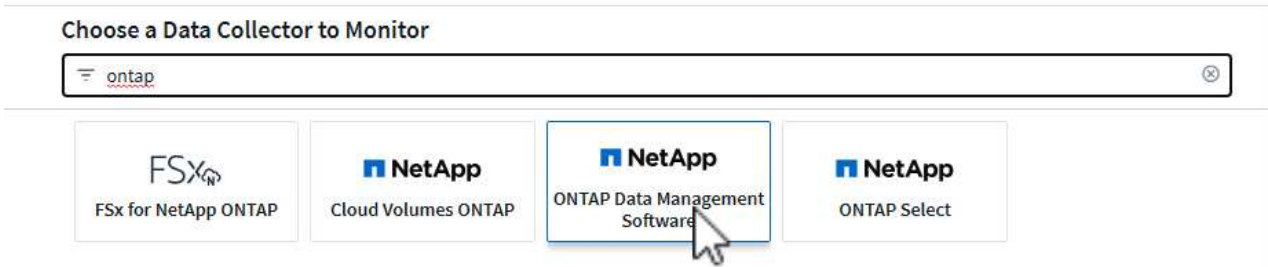
To configure Data Collectors for VMware vSphere and ONTAP storage systems complete the following steps:

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.

✓
○

Select a Data Collector
Configure Data Collector

Configure Collector

[Need Help?](#)

Add credentials and required settings

Name ?

Acquisition Unit

Virtual Center IP Address

User Name

Password

Complete Setup
Test Connection

Advanced Configuration

Collecting:

- Inventory
- VM Performance

Inventory Poll Interval (min)

Communication Port

Filter VMs by

Choose 'Exclude' or 'Include' to Specify a List

Filter Device List (Comma Separated Values For Filtering By ESX_HOST, CLUSTER, and DATACENTER Only)

Performance Poll Interval (sec)

Collect basic performance metrics only

Complete Setup
Test Connection

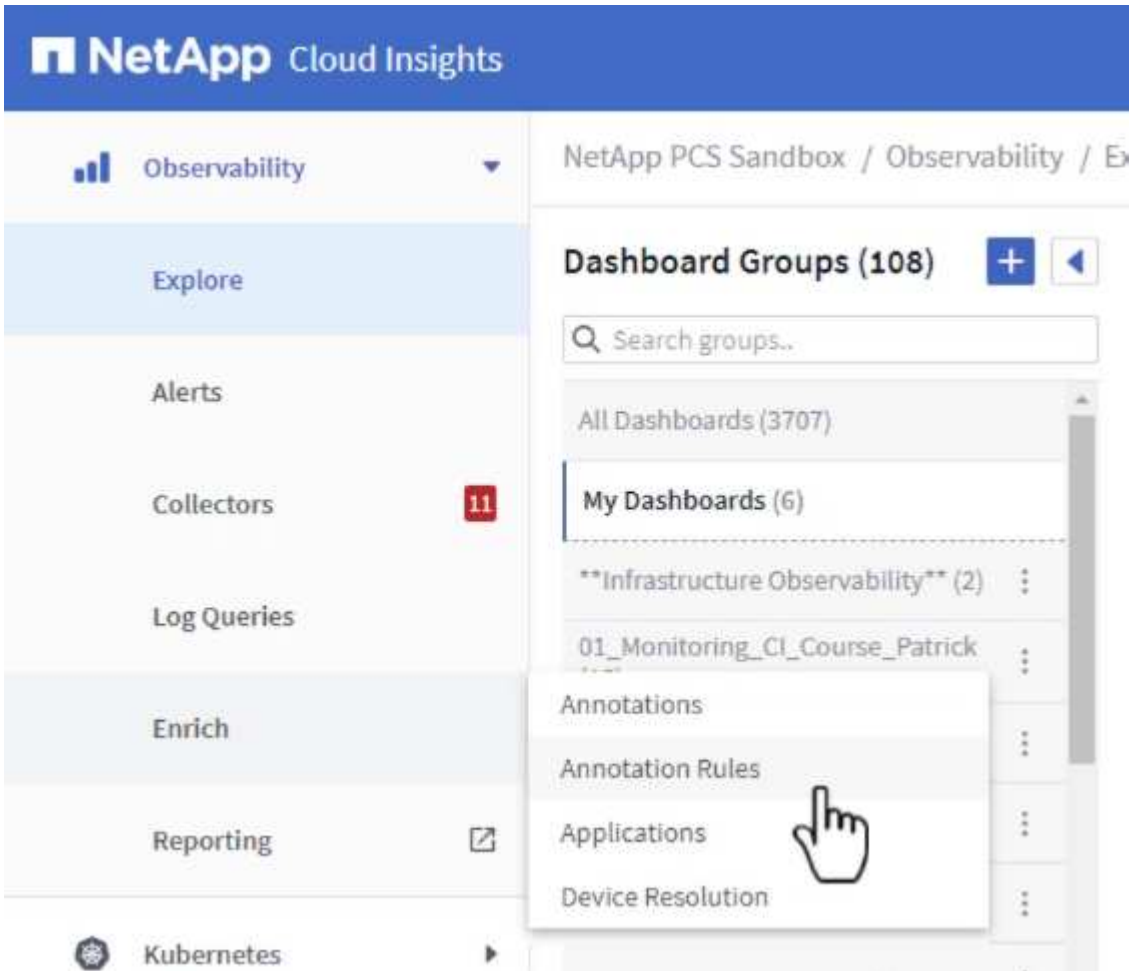
Add Annotations to assets

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.

Add Rule
✕

Name

Query

Annotation

Value

- 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.

NetApp PCS Sandbox / Observability / Enrich / **Annotation Rules**

Rules running... **Run All Rules**

Annotation rules (217) + Rule Filter...

Name	Resource Type	Query	Annotation	Value
Annotate Tier 1 Storage Pools	Storage Pool	Find Storage Pools (no aggro) for Tier...	Tier	Tier 1
Annotate Tier 2 Storage Pools	Storage Pool	Find Storage Pools (no aggro) for Tier...	Tier	Tier 2


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.

- ONTAP FAS/AFF - Cluster Capacity
- ONTAP FAS/AFF - Efficiency
- ONTAP FAS/AFF - FlexVol Performance
- ONTAP FAS/AFF - Node Operational/Optimal Points
- ONTAP FAS/AFF - PrePost Capacity Efficiencies
- Storage Admin - Which nodes are in high demand?
- Storage Admin - Which pools are in high demand?
- StorageGRID - Capacity Summary
- StorageGRID - ILM Performance Monitoring
- StorageGRID - MetaData Usage
- StorageGRID - S3 Performance Monitoring
- VMware Admin - ESX Hosts Overview
- VMware Admin - Overview
- VMware Admin - VM Performance
- VMware Admin - Where are opportunities to right size?
- VMware Admin - Where can I potentially reclaim waste?
- VMware Admin - Where do I have VM Latency?


+ Additional Dashboards (13)
These dashboards require additional data collectors to be installed. [Add More](#)

[Add Dashboards](#) [Go Back](#)


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?**

NetApp PCS Sandbox / Observability / Explore / Dashboards

Dashboard Groups (108) + My Dashboards (6) + From Gallery + Dashboard

Name ↑	Owner
All SAN Array Status (2)	Powell Josh
Cloud Volumes ONTAP - FlexVol Performance (6)	Powell Josh
ONTAP - Volume Workload Performance (Frontend) (7)	Powell Josh
VMware Admin - Where are opportunities to right size? (37)	Powell Josh
VMware Admin - Where can I potentially reclaim waste? (11)	Powell Josh
<input checked="" type="checkbox"/> VMware Admin - Where do I have VM Latency? (9)	Powell Josh

2. Next, filter by the **Data Center** annotation created in a previous step to view a subset of assets.

/ VMware Admin - Where do I have VM Latency? (9) Last 3 Hours

VirtualMachine All Data Center Solutions Engineering diskLatency.total ≥ All

! 5m Avg Latency (all hypervisors) C 5m VM Count With Latency Concern C 5m Avg Latency (all VMs)

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.

VM Count With Latency Concern

5m

50

VM's

Avg Latency (all VMs)


5m

1.55 ms

diskLatency.total


Avg VM Latency - Top 10

5m




Top 5 Avg VM Latency Trend

30s



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?**

My Dashboards (6)

<input type="checkbox"/>	Name ↑
	All SAN Array Status (2)
	Cloud Volumes ONTAP - FlexVol Performance (6)
	ONTAP - Volume Workload Performance (Frontend) (7)
<input type="checkbox"/>	VMware Admin - Where are opportunities to right size? (37)
	VMware Admin - Where do I have VMs that potentially reclaim waste? (11)
	VMware Admin - Where do I have VM Latency? (9)

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.

Memory Usage

5m

121 items found

Virtual Machine	memory (MiB)	memoryUt... ↓
DS3DB0	768.0	81.64
DeployVM0	92.0	55.06
ElasticAppB0	92.0	44.91
AuctionAppA0	336.0	38.42
Client0	480.0	37.98
AuctionAppB0	336.0	37.83
ElasticAppA0	92.0	35.63
ElasticLB0	96.0	35.13
user-cluster1-8872k-78c65dd794...	92.0	32.47
PrimeClient	48.0	30.30

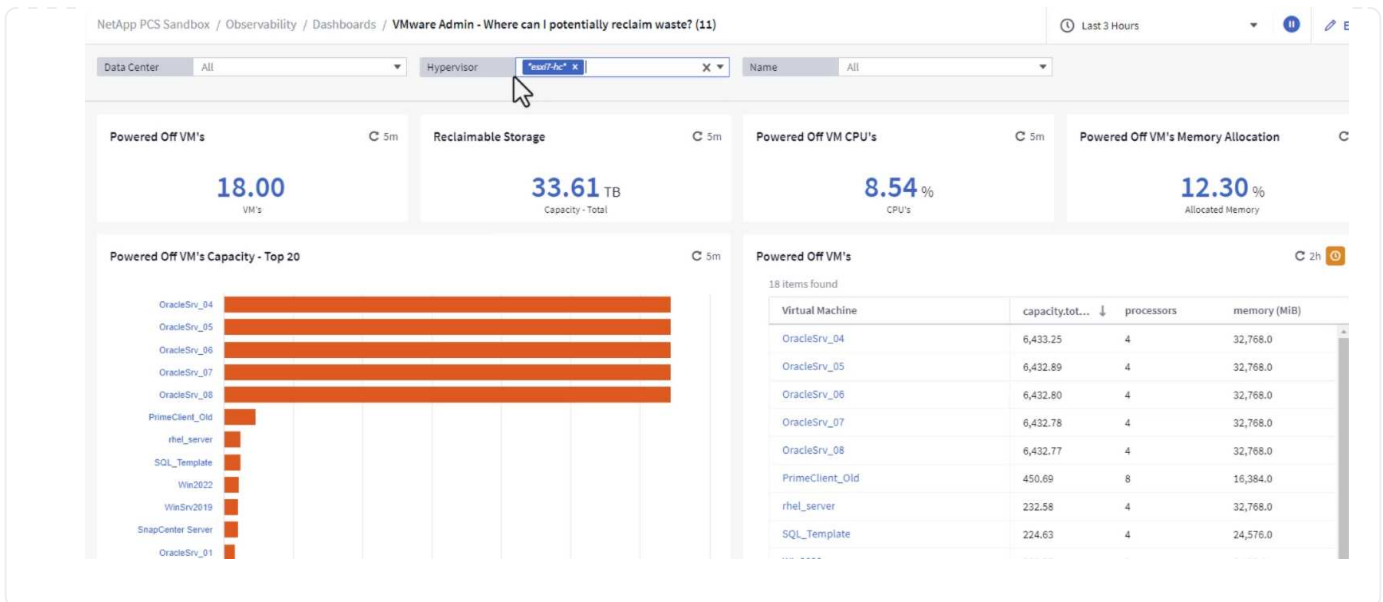
CPU Utilization

5m

121 items found

Virtual Machine	name
hammerdb-01	hammerdb-01
DS3DB0	DS3DB0
wc02-md-0-xwdgb-8cf48c96-qgn...	wc02-md-0-xwdgb-8cf48c96-qg...
ElasticLB0	ElasticLB0

- 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.

VirtualMachine | All Virtual Machines

Filter by Attribute: storageResources.storage.vendor: NetApp | host.Los: VMware

Filter by Metric: +

Group By: Virtual Machine

Formatting: Show Expanded Details | Conditional Formatting | Background Color | Show In Range as green

281 Items found

Virtual Machine	name	powerState	capacity.used (GiB)	capacity.total (GiB)	capacityRatio.us...	diskIops.total (I/O/s)	diskLatency.total...	diskThroughput...
01rfk8prodclient	01rfk8prodclient	On	49.38	69.86	70.68	1.21	8.13	0.01
02rfk8prodserver	02rfk8prodserver	On	63.64	74.06	85.93	22.80	4.13	0.11
03rfk8prodmaster01	03rfk8prodmaster01	On	65.13	77.21	84.36	26.64	5.64	0.20
04rfk8prodmaster02	04rfk8prodmaster02	On	63.89	76.27	83.77	26.82	5.14	0.16
05rfk8prodmaster03	05rfk8prodmaster03	On	63.77	75.58	84.38	28.23	4.63	0.17
AIQUM 9.11 (vApp)	AIQUM 9.11 (vApp)	On	152.00	152.00	100.00	23.24	0.19	0.41
AIQUM 9.12 (Linux)	AIQUM 9.12 (Linux)	On	55.28	100.00	55.28	0.01	11.83	0.00
AN-JumpHost01	AN-JumpHost01	On	90.00	90.00	100.00	1.39	0.19	0.01
AuctionAppA0	AuctionAppA0	On	9.38	16.00	58.62	1.21	0.44	0.12
AuctionAppA1	AuctionAppA1	On	6.44	16.00	40.26	0.00	3.00	0.00

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](#)

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.