



VMware for public cloud

NetApp public and hybrid cloud solutions

NetApp
February 04, 2026

Table of Contents

VMware for public cloud	1
Overview of NetApp Hybrid Multicloud with VMware	1
VMware Cloud options in Public Cloud	1
Did you know?	1
Storage configuration	2
Benefits of NetApp cloud storage	3
Assumptions	3
Detailed architecture	4
NetApp Solutions for VMware in Hyperscalers	4
VMware in the Hyperscalers Configuration	5
NetApp Storage Options	6
NetApp / VMware Cloud Solutions	7
Supported Configurations for NetApp Hybrid Multicloud with VMware	8
VMware in the hyperscalers configuration	8
Configuring the virtualization environment in the cloud provider	8
Deploy and configure the Virtualization Environment on AWS	9
Deploy and configure the Virtualization Environment on Azure	25
Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP)	33
NetApp storage in public clouds	36
NetApp Storage options for Public Cloud Providers	36
Amazon Web Services: Options for using NetApp storage	37
TR-4938: Mount Amazon FSx ONTAP as a NFS datastore with VMware Cloud on AWS	38
Creating a Supplemental NFS Datastore in AWS	41
NetApp Guest Connected Storage Options for AWS	57
Azure Virtualization Service: Options for using NetApp storage	87
Overview of ANF Datastore Solutions	88
Creating a Supplemental NFS Datastore in Azure	90
NetApp Guest Connected Storage Options for Azure	104
Google Cloud VMware Engine: Options for using NetApp storage	120
Google Cloud VMware Engine Supplemental NFS Datastore with Google Cloud NetApp Volumes	121
NetApp Storage Options for GCP	124
Summary and Conclusion: Why NetApp Hybrid Multicloud with VMware	150
Use case #1: Optimizing storage	150
Use case #2: Cloud migration	151
Use case #3: Data center expansion	151
Use case #4: Disaster recovery to the cloud	151
Use case #5: Application modernization	151
Conclusion	151

VMware for public cloud

Overview of NetApp Hybrid Multicloud with VMware

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

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

VMware Cloud options in Public Cloud

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

Azure VMware Solution

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

VMware Cloud on AWS

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

Google Cloud VMware Engine

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



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

Did you know?

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

products:

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

Storage configuration

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

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

Notes:

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

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

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

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

For example:

- RAID 0 = effective free space = 23.4TB (usable raw capacity/1)
- RAID 1 = effective free space = 11.7TB (usable raw capacity/2)
- RAID 5 = effective free space = 17.5TB (usable raw capacity/1.33)

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



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

Points to Remember

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



This applies to guest connected storage only.

Benefits of NetApp cloud storage

NetApp cloud storage offers the following benefits:

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

Assumptions

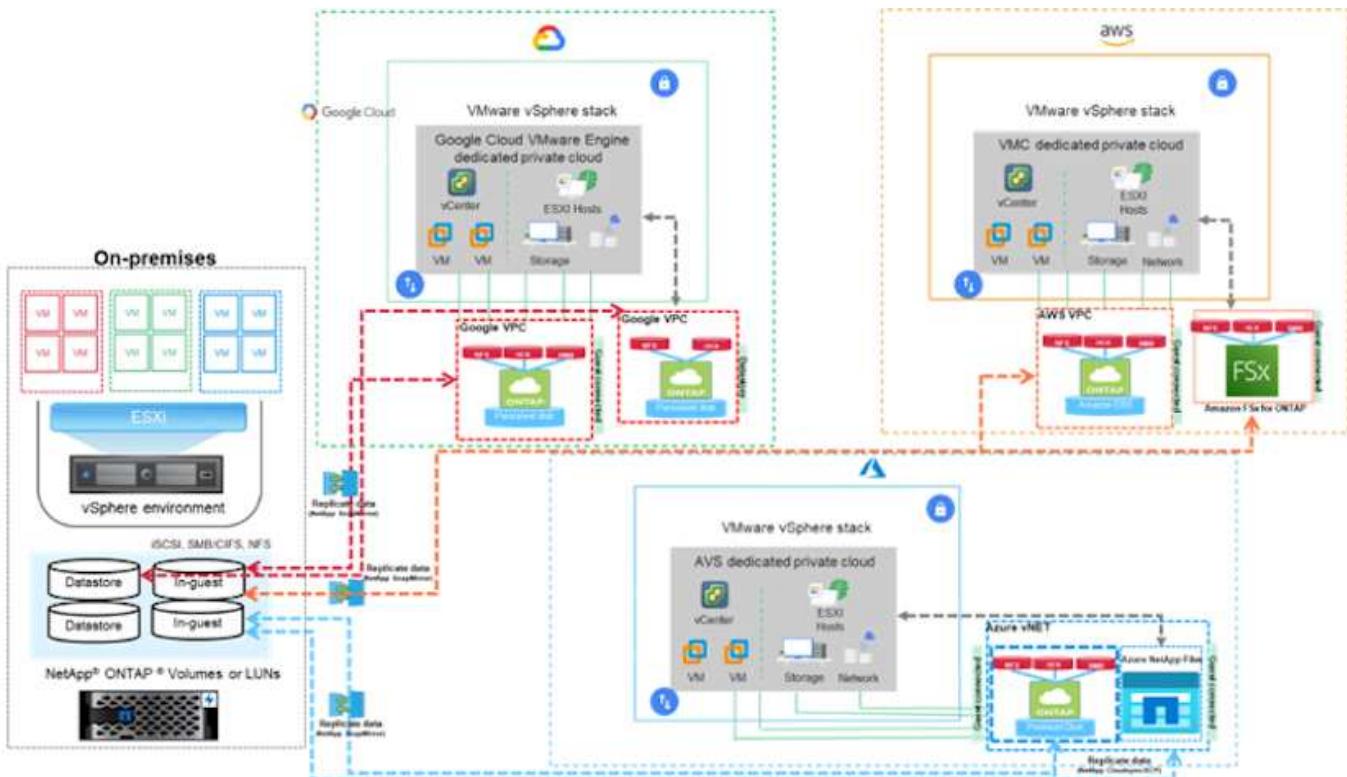
- SnapMirror technology or other relevant data migration mechanisms are enabled. There are many connectivity options, from on-premises to any hyperscaler cloud. Use the appropriate path and work with the relevant networking teams.
- In-guest storage was the only available option at the time this document was written.



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

Detailed architecture

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



NetApp Solutions for VMware in Hyperscalers

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

Pick your cloud and let NetApp do the rest!



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

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

- [VMware in the Hyperscalers Configuration](#)
- [NetApp Storage Options](#)
- [NetApp / VMware Cloud Solutions](#)

VMware in the Hyperscalers Configuration

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

AWS / VMC

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



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

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

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

View the detailed [configuration steps for VMC](#).

Azure / AVS

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



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

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

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

View the detailed [configuration steps for AVS](#).

GCP / GCVE

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



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

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

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed [configuration steps for GCVE](#).

NetApp Storage Options

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

Please visit [Supported NetApp Storage Options](#) for more information.

AWS / VMC

AWS supports NetApp storage in the following configurations:

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

View the detailed [guest connect storage options for VMC](#).

View the detailed [supplemental NFS datastore options for VMC](#).

Azure / AVS

Azure supports NetApp storage in the following configurations:

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

View the detailed [guest connect storage options for AVS](#).

View the detailed [supplemental NFS datastore options for AVS](#).

GCP / GCVE

Google Cloud supports NetApp storage in the following configurations:

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

View the detailed [guest connect storage options for GCVE](#).

View the detailed [supplemental NFS datastore options for GCVE](#).

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

NetApp / VMware Cloud Solutions

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

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

AWS / VMC

[Browse the NetApp solutions for AWS / VMC](#)

Azure / AVS

[Browse the NetApp solutions for Azure / AVS](#)

GCP / GCVE

[Browse the NetApp solutions for Google Cloud Platform \(GCP\) / GCVE](#)

Supported Configurations for NetApp Hybrid Multicloud with VMware

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

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

VMware in the hyperscalers configuration

Configuring the virtualization environment in the cloud provider

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

AWS / VMC

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



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

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

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

View the detailed [configuration steps for VMC](#).

Azure / AVS

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



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

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

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

View the detailed [configuration steps for AVS](#).

GCP / GCVE

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



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

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

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed [configuration steps for GCVE](#).

Deploy and configure the Virtualization Environment on AWS

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

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



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

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

Deploy and configure VMware Cloud for AWS

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

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



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

The setup process can be broken down into three parts:

Register for an AWS Account

Register for an [Amazon Web Services Account](#).

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

Register for a My VMware Account

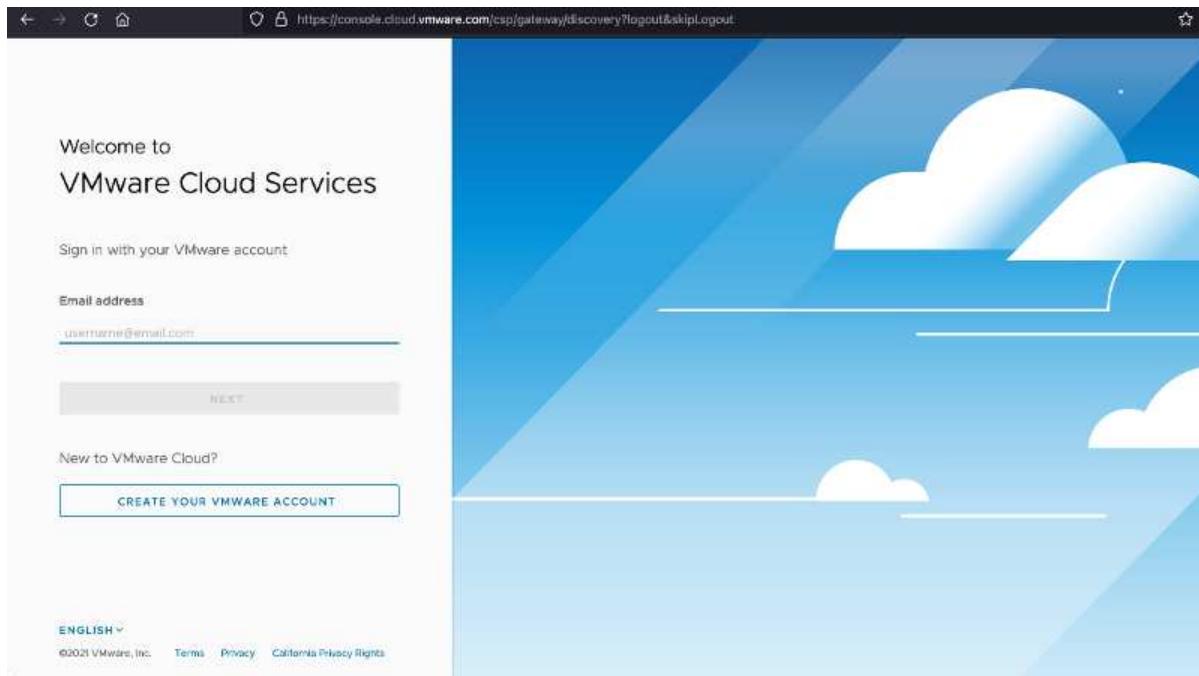
Register for a [My VMware](#) account.

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

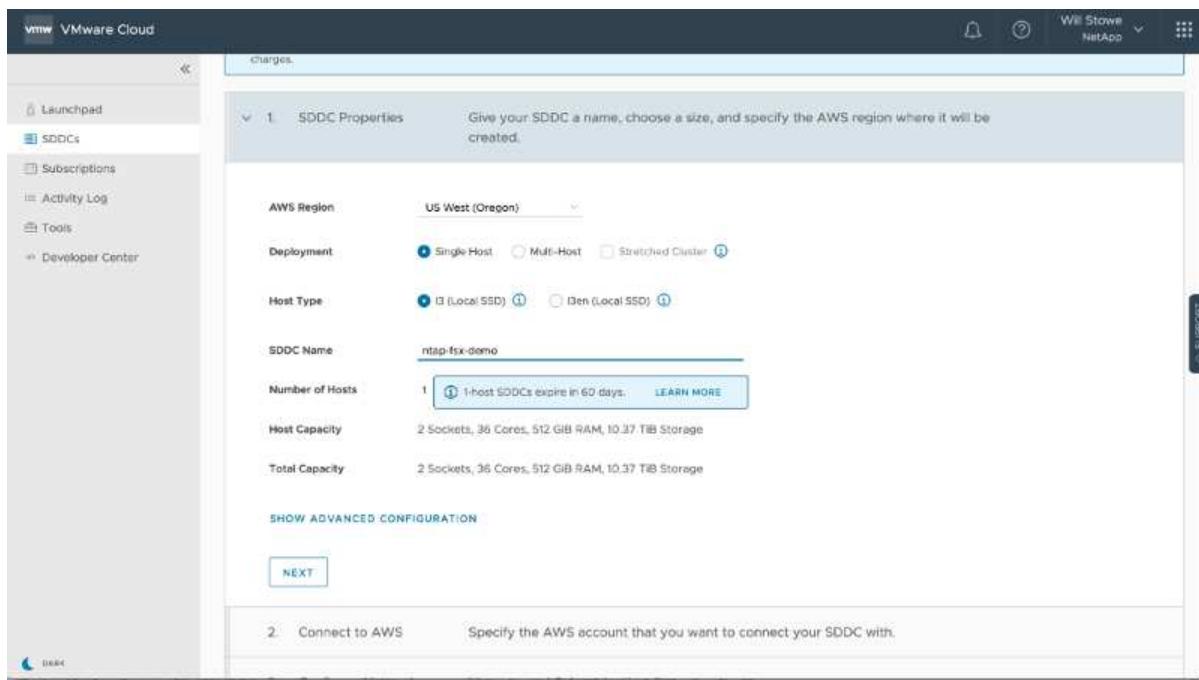
Provision SDDC in VMware Cloud

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

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



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



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

CloudFormation > Stacks > Create stack

Quick create stack

Template

Template URL:
<https://vmware-sddc.s3.us-west-2.amazonaws.com/1eb9d184-a706-4489-abb8-692aad0e25d0/mq5ijohtclieoh815b75nteq9cc4bddd7ifq07nv7v16fk36>

Stack description
 This template is created by VMware Cloud on AWS for SDDC deployment and maintenance. Please do not remove.

Stack name

Stack name:
 vmware-sddc-formation-a87f31c9-e5ac-4bb4-9d1e-9a3dabd197b7

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Feedback English (US) ▾

© 2006–2021, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use Cookie preferences

CloudFormation > Stacks > Create stack

Quick create stack

Stack name

Stack name:
 vmware-sddc-formation-a87f31c9-e5ac-4bb4-9d1e-9a3dabd197b7

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

No parameters.
 There are no parameters defined in your template

Capabilities

Important The following resource(s) require capabilities: [AWS::IAM::Role]

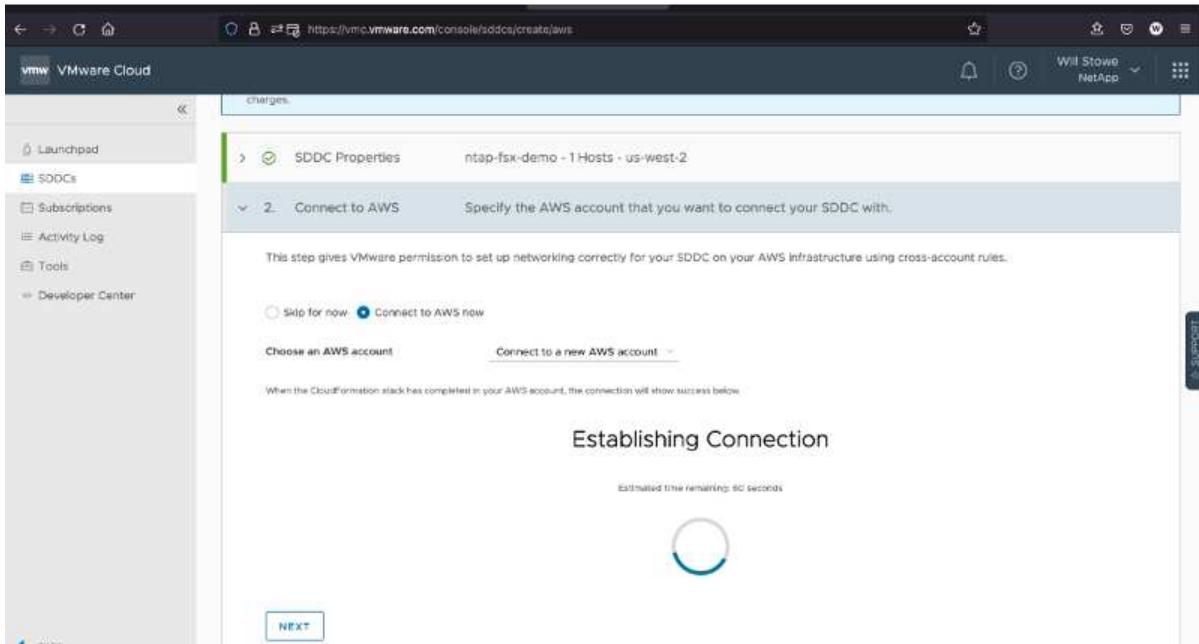
This template contains Identity and Access Management (IAM) resources that might provide entities access to make changes to your AWS account. Check that you want to create each of these resources and that they have the minimum required permissions. [Learn more](#)

I acknowledge that AWS CloudFormation might create IAM resources.

Cancel Create change set Create stack

Feedback English (US) ▾

© 2006–2021, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use Cookie preferences



VMware Cloud

SDDC Properties: ntap-fsx-demo - 1 Hosts - us-west-2

2. Connect to AWS: Specify the AWS account that you want to connect your SDDC with.

This step gives VMware permission to set up networking correctly for your SDDC on your AWS Infrastructure using cross-account rules.

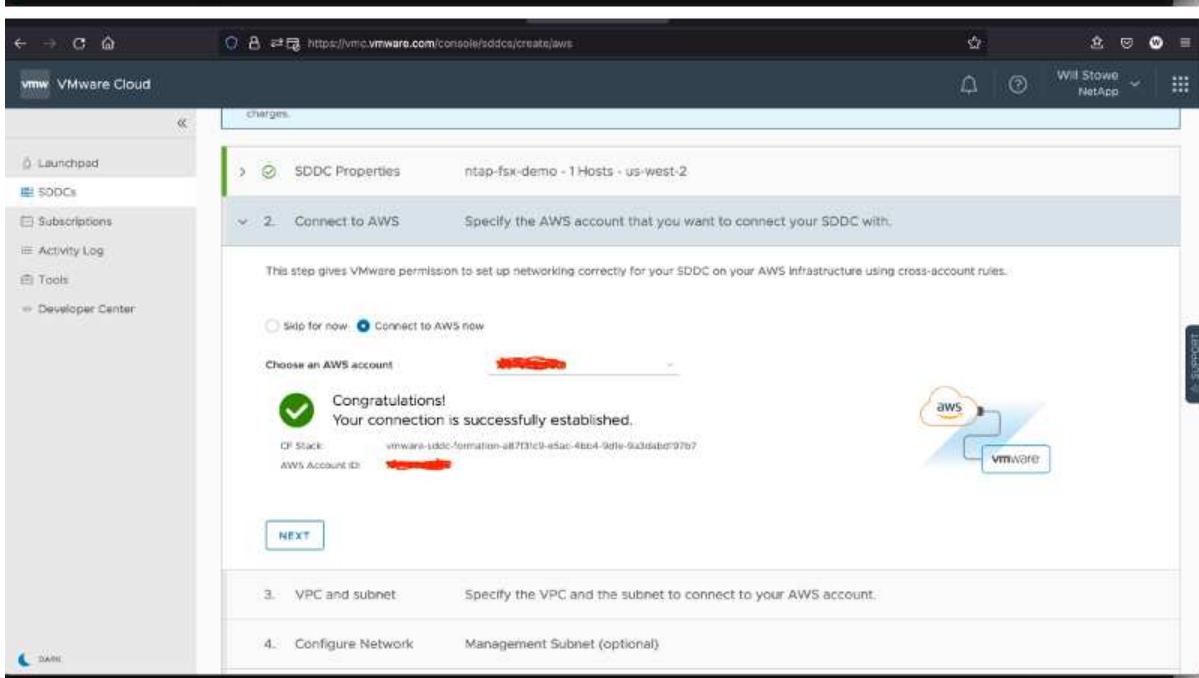
Skip for now. Connect to AWS now

Choose an AWS account: Connect to a new AWS account

Establishing Connection

Estimated time remaining: 60 seconds

NEXT



VMware Cloud

SDDC Properties: ntap-fsx-demo - 1 Hosts - us-west-2

2. Connect to AWS: Specify the AWS account that you want to connect your SDDC with.

This step gives VMware permission to set up networking correctly for your SDDC on your AWS Infrastructure using cross-account rules.

Skip for now. Connect to AWS now

Choose an AWS account: [REDACTED]

Congratulations!
Your connection is successfully established.

CF Stack: vmware-sddc-formation-487f01c9-45ac-4bb4-9d1e-9a3dab97b7b
AWS Account ID: [REDACTED]

NEXT

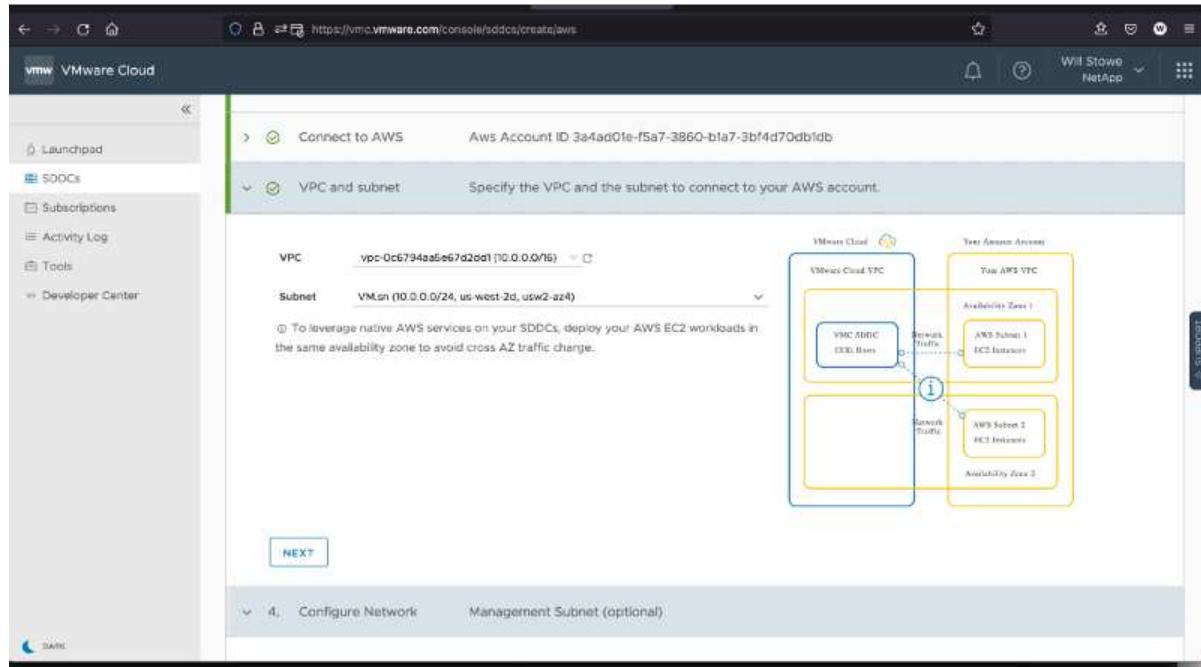
3. VPC and subnet: Specify the VPC and the subnet to connect to your AWS account.

4. Configure Network: Management Subnet (optional)

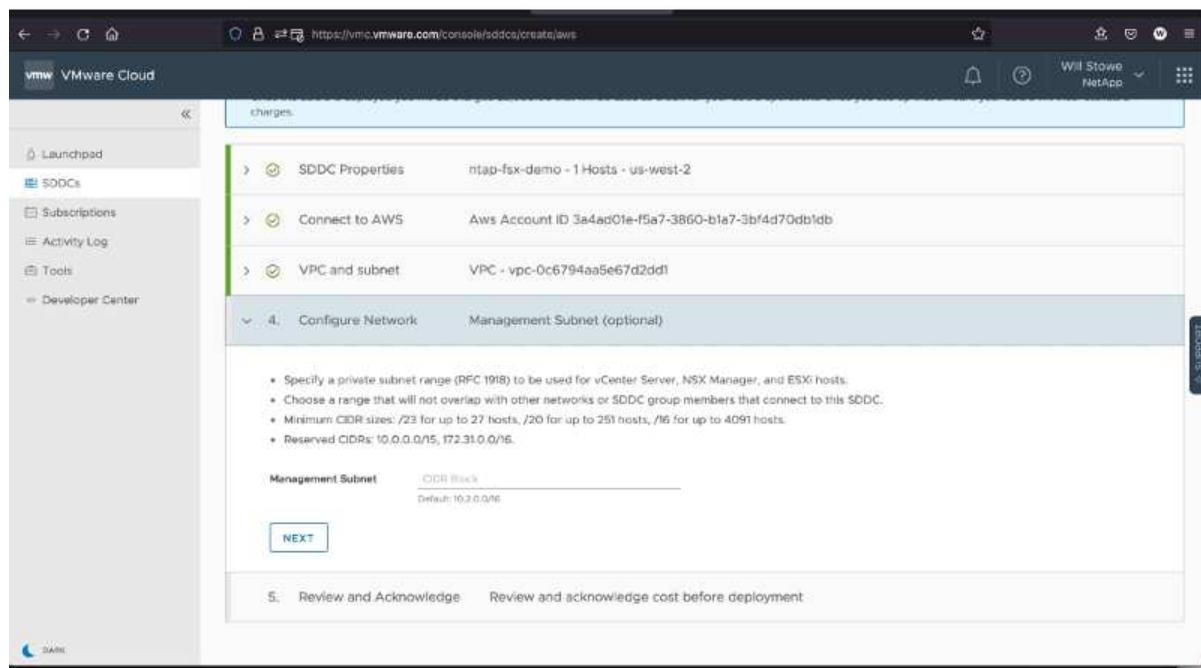


Single-host configuration is used in this validation.

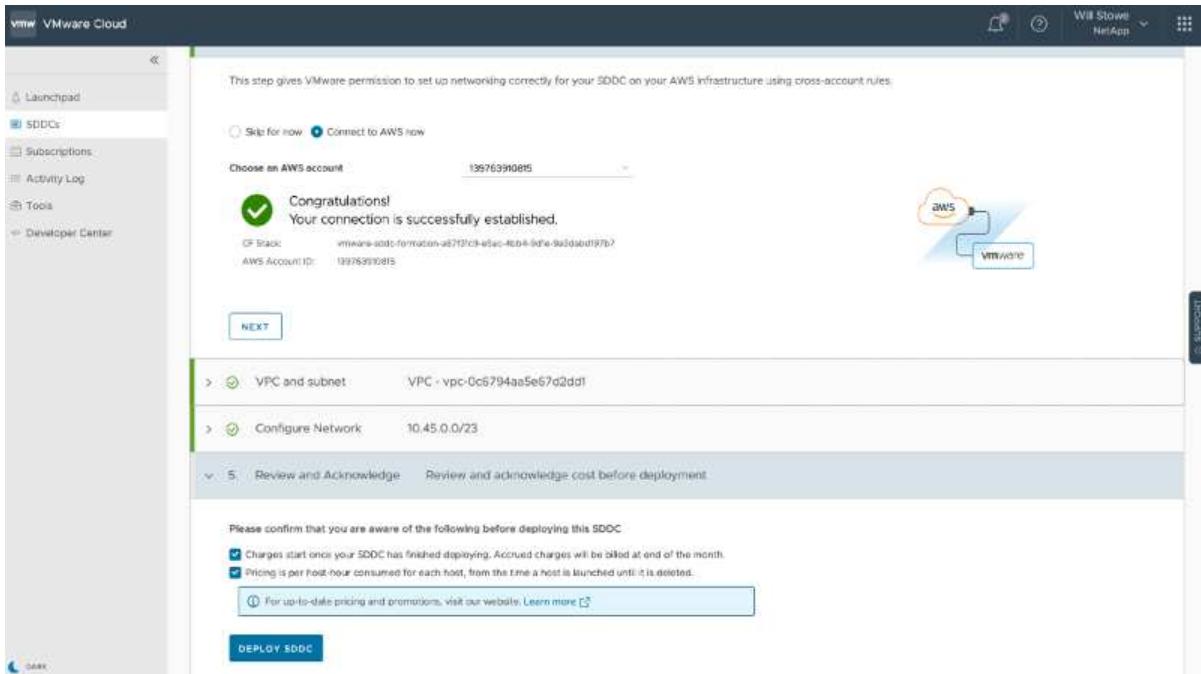
4. Select the desired AWS VPC to connect the VMC environment with.



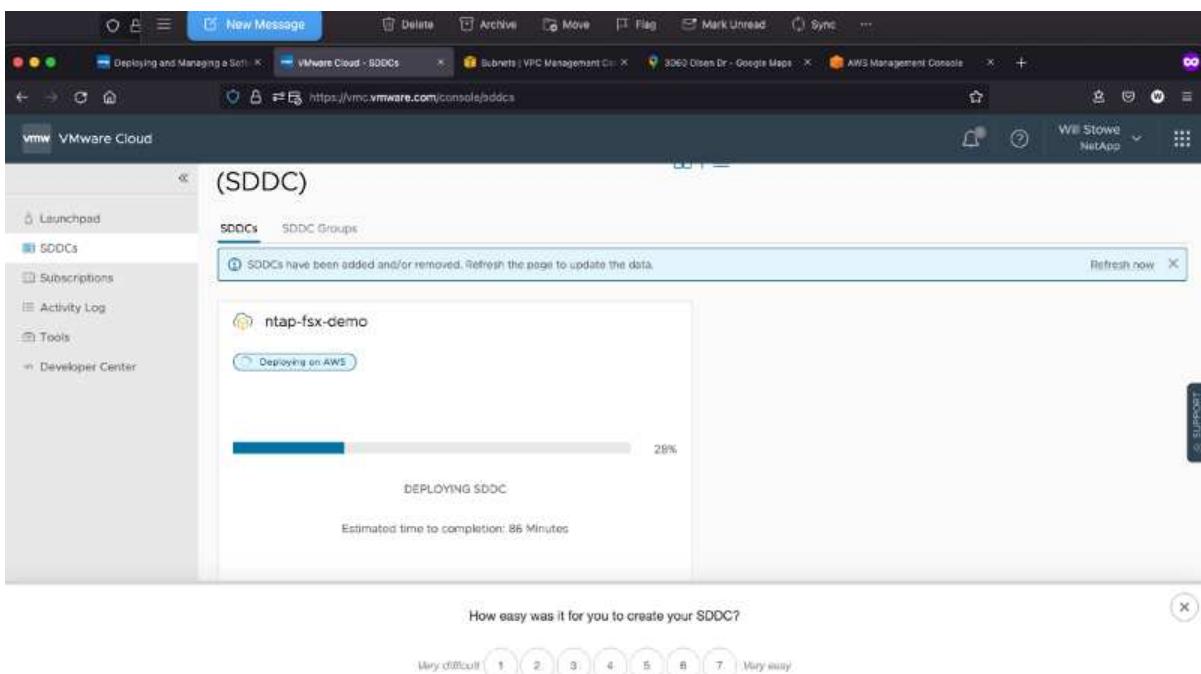
5. Configure the VMC Management Subnet; this subnet contains VMC-managed services like vCenter, NSX, and so on. Do not choose an overlapping address space with any other networks that need connectivity to the SDDC environment. Finally, follow the recommendations for CIDR size notated below.



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



The deployment process typically takes approximately two hours to complete.



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

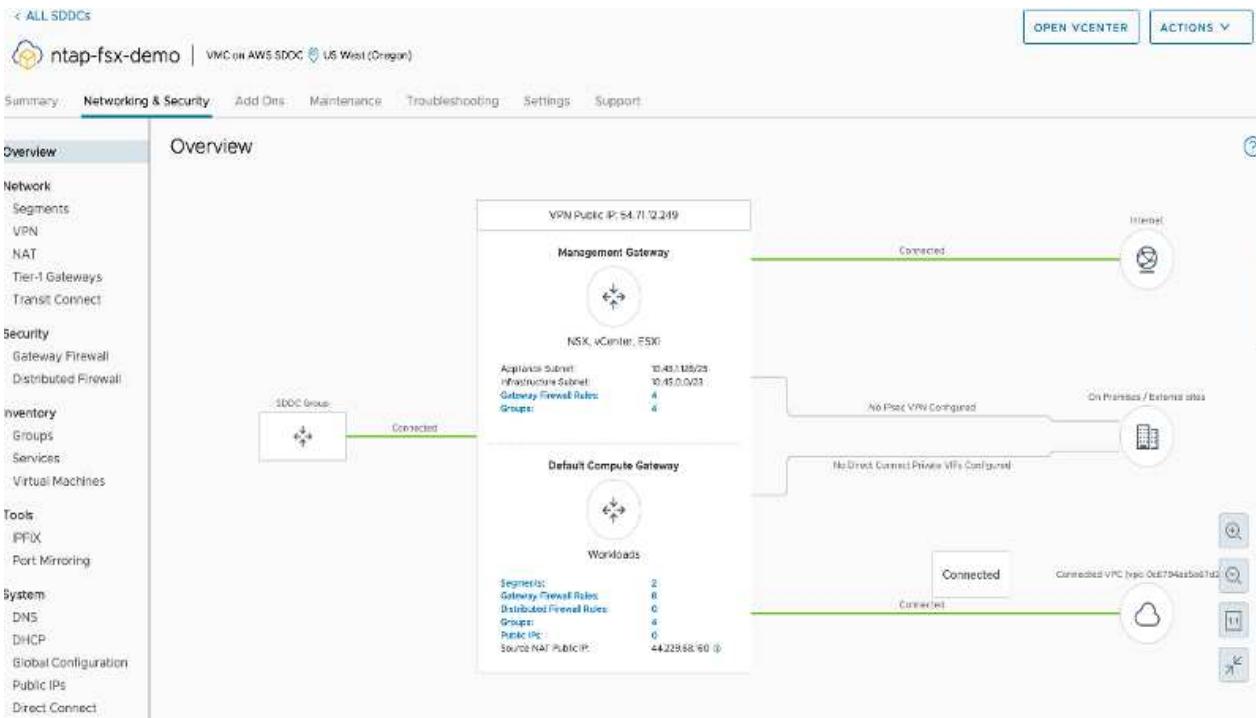
The screenshot shows the VMware Cloud Software-Defined Data Centers (SDDC) interface. The main title is "Software-Defined Data Centers (SDDC)". On the left, a sidebar lists "Launched" and "SDDCs". The "SDDCs" section is selected, showing a list with one item: "ntap-fsx-demo". The status of this SDDC is "Ready". Below the status, there are details: Region (US West (Oregon)), Type (VMC on AWS SDDC), Clusters (1), Hosts (1), Availability Zones (us-west-2a), and Capacity (VMC on AWS SDDC). Below these details, resource statistics are provided: CPU (82.8 GHz), Memory (512 GiB), and Storage (10.37 TiB). At the bottom of the card, there are links for "VIEW DETAILS", "OPEN VCENTER", and "ACTIONS".

For a step-by-step guide on SDDC deployment, see [Deploy an SDDC from the VMC Console](#).

Connect VMware Cloud to FSx ONTAP

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

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



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

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

VMware Cloud

Create SDDC Group

1. Name and Description

Name: sddcgroup01

Description: sddcgroup01

NEXT

2. Membership

Members: 1

3. Acknowledgement

Please confirm that you are aware of the following before creating this SDDC Group.

Configuring VMware Transit Connect for your group will incur charges per attachment and data transfers.

Create firewall rules to establish connectivity between the SDDCs in the group. [Learn More](#)

CREATE GROUP

VMware Cloud

Create SDDC Group

1. Name and Description

Name: sddcgroup01

2. Membership

Select SDDCs to be part of your group.

Name	UUID	Location	Version	Management CIDR
intap-px-demo	829a6e22-92a1-42db-ad03-9e4eb7a908b6	US West (Oregon)	1.14.0.14	10.45.0.0/23
1				

NEXT

3. Acknowledgement

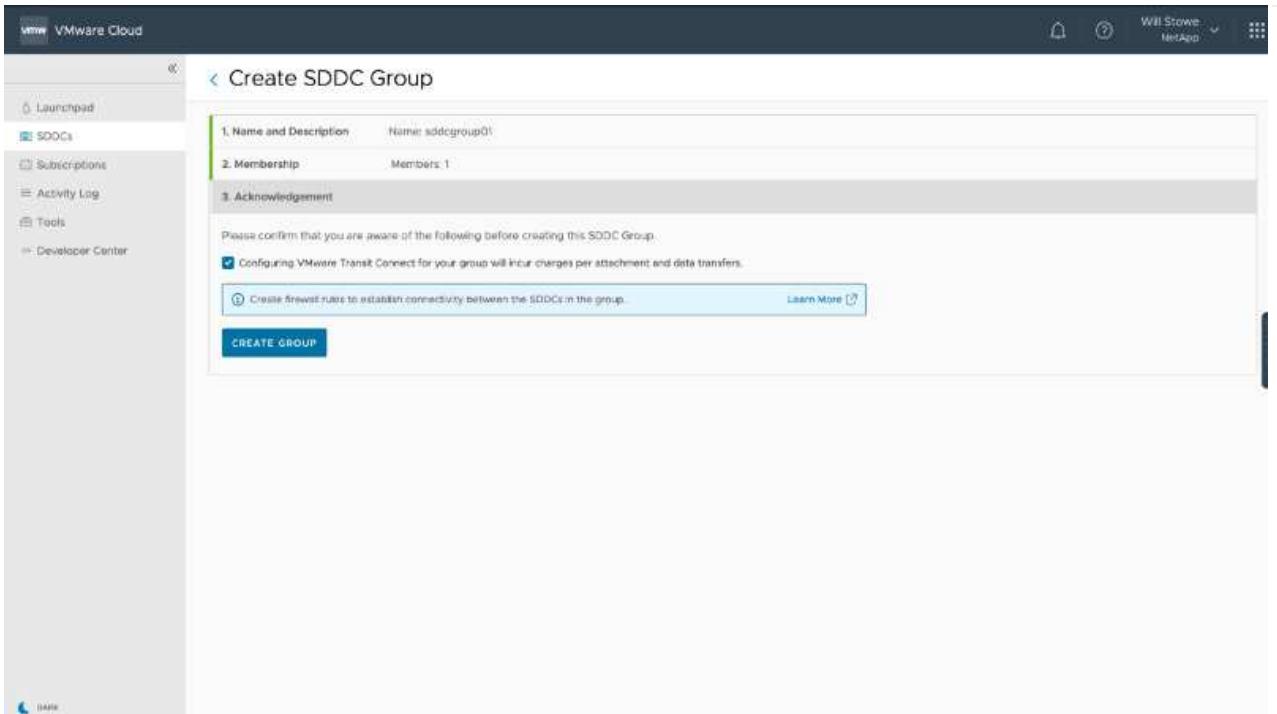
Review and acknowledge requirements before creating the group.

Please confirm that you are aware of the following before creating this SDDC Group.

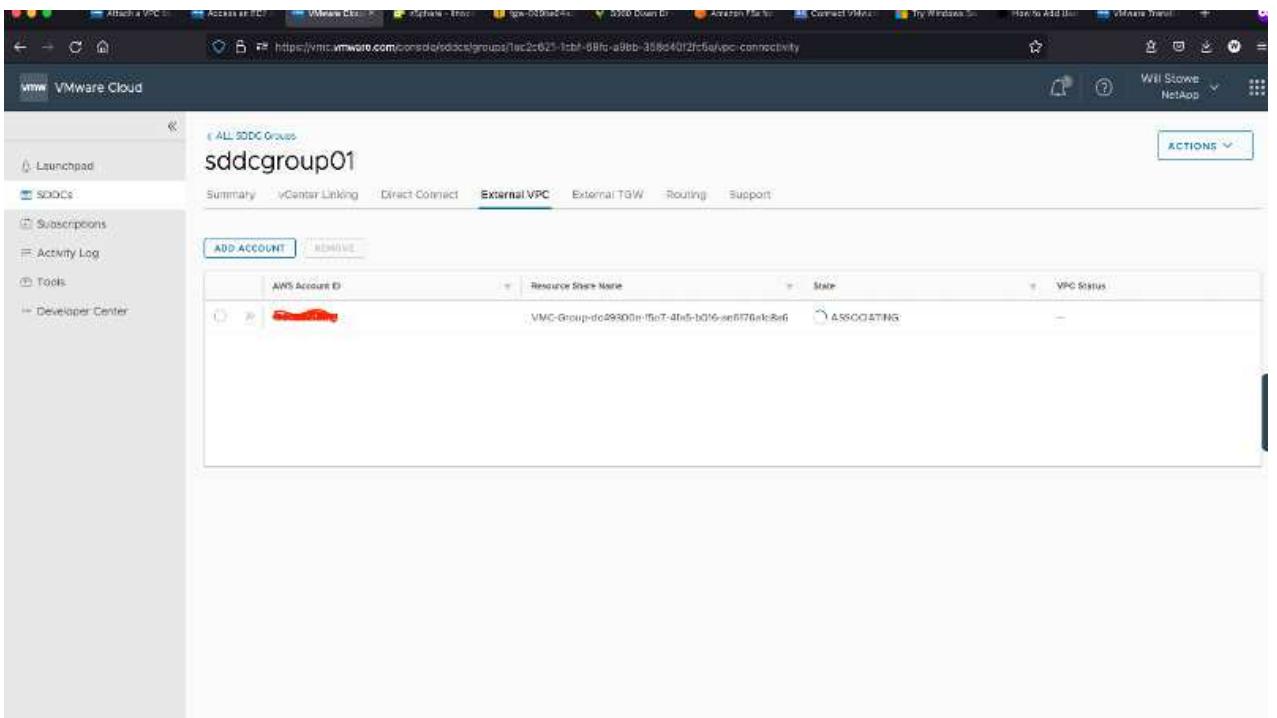
Configuring VMware Transit Connect for your group will incur charges per attachment and data transfers.

Create firewall rules to establish connectivity between the SDDCs in the group. [Learn More](#)

CREATE GROUP



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



VMware Cloud

ALL SDDC GROUPS

sddcgroup01

External VPC

AWS Account ID	Resource Share Name	Status	VPC Status
111122223333	VMC-Group-dc09300e15e74f55-b016-ee01768e86	ASSOCIATED	...

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

Resource Access Manager

Shared by me

Resource shares

Shared resources

Principals

Shared with me

Resource shares 1 in use

Shared resources

Principals

Permissions library New

Settings

AWS Resource Access Manager

Share AWS resources with other AWS accounts.

How it works

AWS Resource Access Manager allows you to share AWS resources across multiple AWS accounts or AWS Organizations by creating a Resource Share.

Select Resources

Specify Principals

Share Resources

Use cases

Manage resources centrally in a multi-account environment

Increase efficiency, decrease costs

AZ IDs provide a consistent way of identifying the location of a resource across all your accounts. This makes it easier for you to provision resources centrally in a single account and share them across multiple accounts.

5. Create the Transit Gateway Attachment.

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

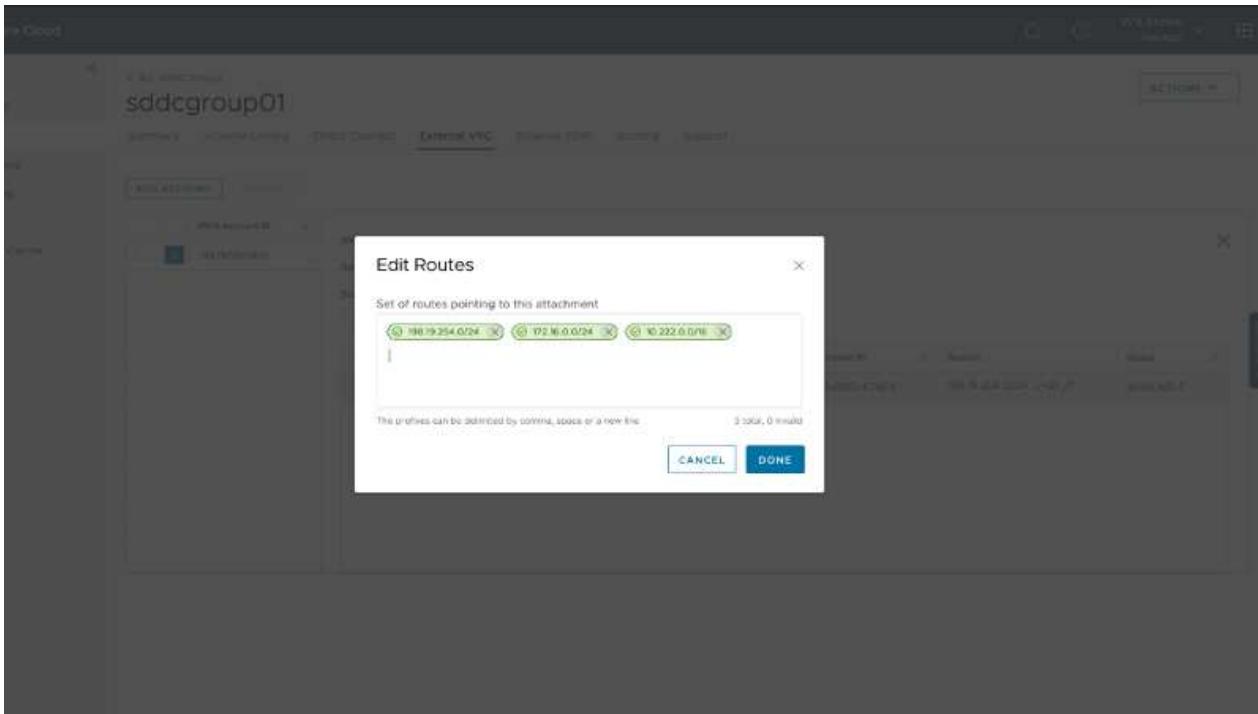
VPC ID	VNC on AWS Region	Transit Gateway Attachment ID	Routes	Status
vpc-0d1c764bcc495e805	US West (Oregon)	tgw-attach-0a4883d6f92c67d64	192.168.0.0/24	PENDING

7. While in the External VPC tab, click the edit icon in the Routes column and add in the following required routes:

- A route for the floating IP range for Amazon FSx ONTAP [floating IPs](#).
- A route for the floating IP range for Cloud Volumes ONTAP (if applicable).
- A route for the newly created external VPC address space.

VPC ID	VNC on AWS Region	Transit Gateway Attachment ID	Routes	Status
vpc-0d1c764bcc495e805	US West (Oregon)	tgw-attach-0a4883d6f92c67d64	192.168.0.0/24	AVAILABLE

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



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

Name	ID	Sources	Destinations	Services	Applied To	Action
allow Internet fro...	1019	vmc-addc vmc-addc-2	Any	Any	All Uplinks	Allow
allow VMC to VPC	1017	vmc-addc vmc-addc-2	vmc-addc vmc-addc-2	Any	All Uplinks	Allow
allow VPC to VMC	1016	Connecti...	vmc-addc	Any	All Uplinks	Allow
allow to vmcfsx2...	1022	vmc-addc vmc-addc-2	vmcfsx2.v...	Any	All Uplinks	Allow
all from vmcfsx2...	1023	vmcfsx2.v...	vmc-addc-2 vmc-addc	Any	All Uplinks	Allow
Default VTI Rule	1012	Any	Any	Any	VPN Tunnel In...	Allow
Default Uplink Ru...		Any	Any	Any	All Uplinks	Drop

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

Deploy and configure the Virtualization Environment on Azure

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

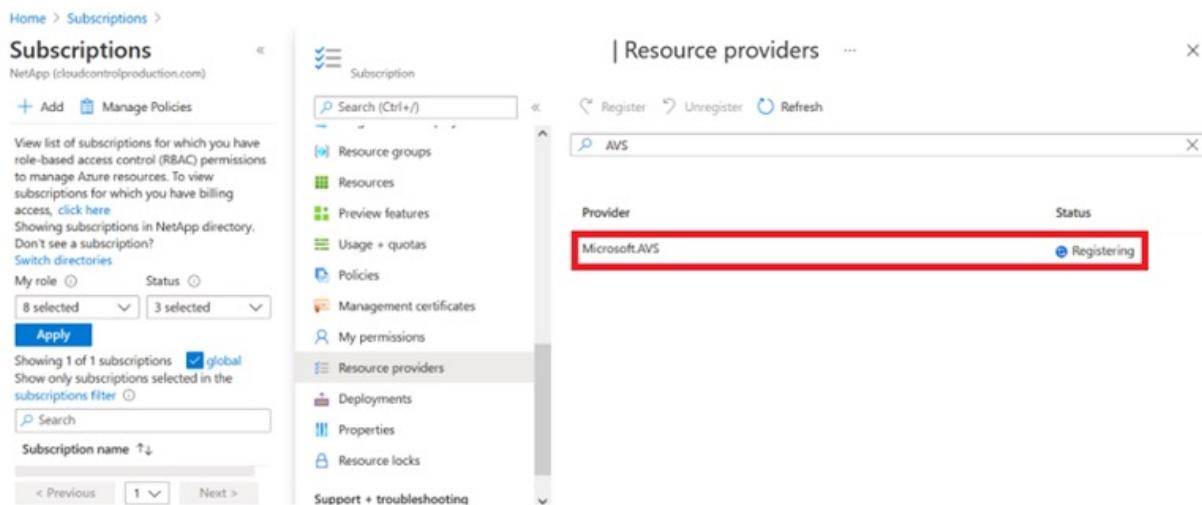
This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.

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

Register the resource provider and create a private cloud

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

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



Provider	Status
Microsoft.OperationsManagement	Registered
Microsoft.Compute	Registered
Microsoft.ContainerService	Registered
Microsoft.ManagedIdentity	Registered
Microsoft.AVS	Registered
Microsoft.OperationalInsights	Registered
Microsoft.GuestConfiguration	Registered

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

Notes:

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

Create a private cloud

Prerequisites **Basics** Tags Review and Create

Project details

Subscription * **SaaS Backup Production**
Resource group * **(New) NimoAVSDemo**
Create new

Private cloud details

Resource name * **nimoavsppriv**
Location * **(US) East US 2**
Size of host * **AV36 Trial**
Number of hosts * **3**
Find out how many hosts you need

There is no metering for the selected subscription, region, and SKU. No cost data to display.

CIDR address block

Provide IP address for private cloud for cluster management. Make sure these are unique and do not overlap with any other Azure vnets or on-premise networks.

Address block for private cloud * **10.21.0.0/22**

Review and Create Previous Next : Tags >

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

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

Home >

nimoavsppriv AVS Private cloud

Search (Ctrl+ /) Delete

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Settings

Locks

Manage

Connectivity

Identity

Clusters

Essentials

Resource group (change) NimoAVSDemo	Address block for private cloud 10.21.0.0/22
Status Succeeded	Primary peering subnet 10.21.0.232/30
Location East US 2	Secondary peering subnet 10.21.0.236/30
Subscription (change) SaaS Backup Production	Private Cloud Management network 10.21.0.0/26
Subscription ID b58a041a-e464-4497-8be9-9048369ee8e1	vMotion network 10.21.1.128/25
Tags (change) Click here to add tags	Number of hosts 3

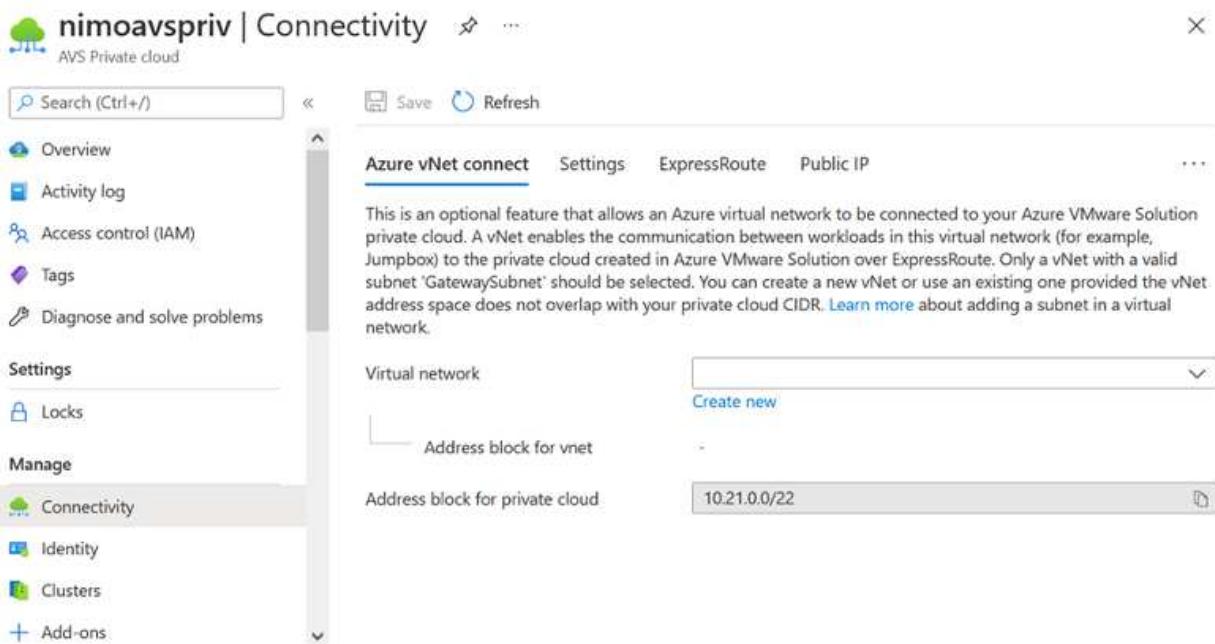
Connect to a new or existing ExpressRoute virtual network gateway

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

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

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

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



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

Create virtual network

X

This virtual network enables the communication between workloads in this virtual network (e.g. a Jumphost) to the private cloud created in Azure VMware Solution over an Express route. A default address range and a subnet is selected for this virtual network. For changing the default address range and subnet of this virtual network, follow these steps: Step 1: Change the "Address Range" to desired range (e.g. 172.16.0.0/16). Step 2: Add a subnet under "Subnets" with the name as "GatewaySubnet" and provide subnet's address range in CIDR notation (e.g. 172.16.1.0/24). [Learn more about virtual networks](#)

Name *

nimoavspiv-vnet

Address space

The virtual network's address space specified as one or more address prefixes in CIDR notation (e.g. 10.0.0.0/16).

<input type="checkbox"/> Address range	Addresses	Overlap	
<input type="checkbox"/> 172.24.0.0/16	172.24.0.4 - 172.24.255.254 (65531 addresses)	None	
	(0 Addresses)	None	

Subnets

The subnet's address range in CIDR notation (e.g. 10.0.0.0/24). It must be contained by the address space of the virtual network.

<input type="checkbox"/> Subnet name	Address range	Addresses	
<input type="checkbox"/> GatewaySubnet	172.24.0.0/24	172.24.0.4 - 172.24.0.254 (251 addresses)	
		(0 Addresses)	

OK

Discard

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



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

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

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

Create a virtual machine

Basics Disks Networking Management Advanced Tags Review + create

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization. [Learn more](#)

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * SaaS Backup Production

Resource group * NimoAVSDemo [Create new](#)

Instance details

Virtual machine name * nimAVSRH

Region * (US) East US 2

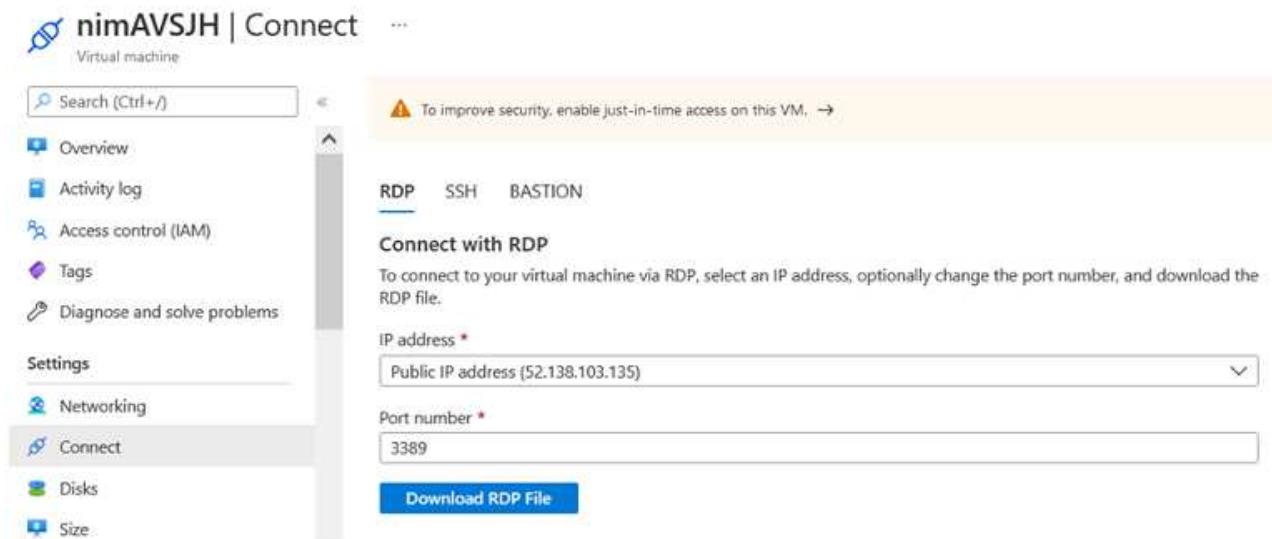
Availability options No infrastructure redundancy required

Image * Windows Server 2012 R2 Datacenter - Gen2 [See all images](#)

Azure Spot instance

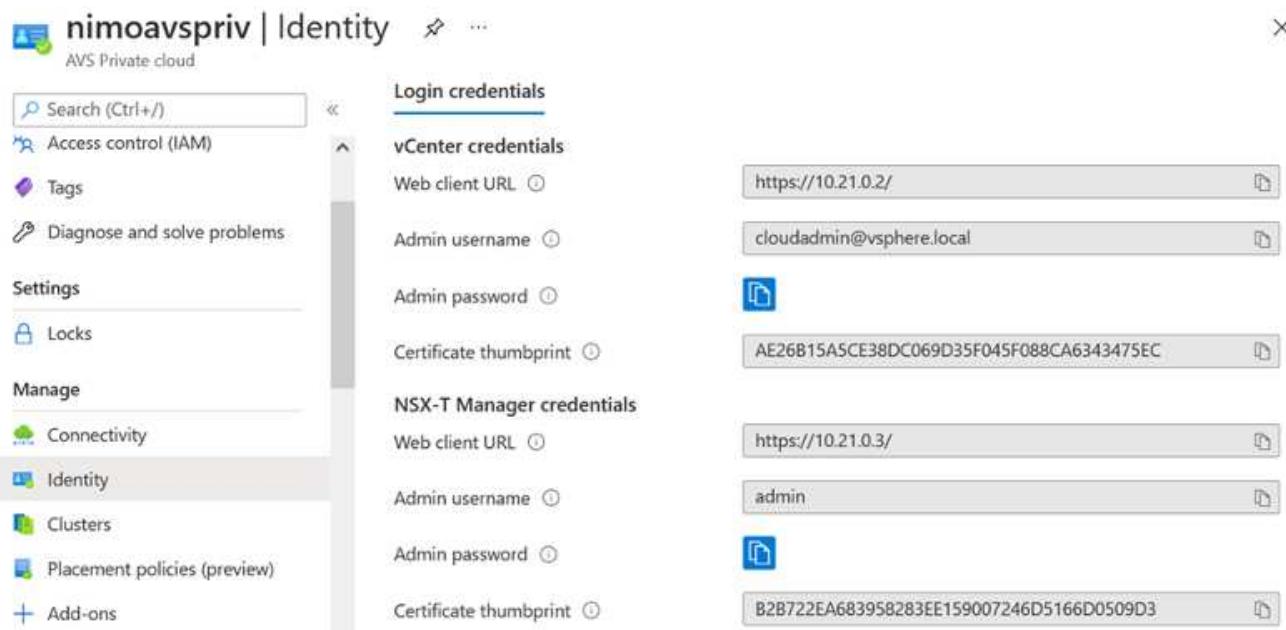
Size * Standard_D2s_v3 - 2 vcpus, 8 GiB memory (\$130.67/month) [See all sizes](#)

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



The screenshot shows the Azure portal interface for a virtual machine named 'nimAVSJH'. The left sidebar includes options like Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Settings, Networking, Connect, Disks, and Size. The 'Connect' option is currently selected. The main content area is titled 'nimAVSJH | Connect' and shows the 'Virtual machine' section. A search bar at the top left contains 'Search (Ctrl+ /)'. A warning message at the top right says 'To improve security, enable just-in-time access on this VM, →'. Below this, there are tabs for RDP, SSH, and BASTION, with 'RDP' selected. The 'Connect with RDP' section contains fields for 'IP address' (set to 'Public IP address (52.138.103.135)') and 'Port number' (set to '3389'), along with a 'Download RDP File' button.

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



The screenshot shows the Azure portal interface for a private cloud named 'nimoavsppriv'. The left sidebar includes options like Access control (IAM), Tags, Diagnose and solve problems, Settings, Locks, Manage, Connectivity, Identity (which is selected), Clusters, Placement policies (preview), and Add-ons. The main content area is titled 'nimoavsppriv | Identity' and shows the 'AVS Private cloud' section. A search bar at the top left contains 'Search (Ctrl+ /)'. The 'Identity' section contains two main sections: 'vCenter credentials' and 'NSX-T Manager credentials'. The 'vCenter credentials' section includes fields for 'Web client URL' (set to 'https://10.21.0.2/'), 'Admin username' (set to 'cloudadmin@vsphere.local'), and 'Admin password' (with a copy icon). The 'NSX-T Manager credentials' section includes fields for 'Web client URL' (set to 'https://10.21.0.3/'), 'Admin username' (set to 'admin'), and 'Admin password' (with a copy icon). Both sections also have 'Certificate thumbprint' fields.

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



The web client URLs are different for each SDDC provisioned.

VMware® vSphere

cloudadmin@vsphere.local

Use Windows session authentication

LOGIN

vSphere Client | **Menu** | **Search in all environments**

vc.beeb9fd29eab4cbea81e62.eastus2.avs.azure.com | **ACTIONS**

Sum... **Mon...** **Config...** **Permiss...** **Datacenter...** **Hosts & Cl...** **V...** **Dataset...** **Netw...** **Linked vCenter Server...** **Extens...**

SDDC-Datacenter

Virtual Machines: 0
Hosts: 3

Custom Attributes

Attribute	Value

Tags

Assigned Tag	Category	Description

Recent Tasks **Alarms**

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Undeploy plug-in	vc.beeb9fd29...	Completed	VMware vRops Client Plugin	VSPPHERE.LOCAL...	8 ms	08/12/2021, 11:38:11 AM	08/12/2021, 11:38:11 AM	vc.beeb9fd29eab...

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

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

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

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

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

Deploy and configure GCVE

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

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

Pre-requisites:

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

Google Cloud VMware Engine

← Create Private Cloud

Private Cloud name *

NIMoGCVE

Location *

us-east4 > v-zone-a > VE Placement Group 2

Node type *

ve1-standard-72
2x2.6 GHz, 36 Cores (72 HT), 768 GB RAM
19.2 TB Raw, 3.2 TB Cache (All-Flash)

Node count *

3
(3 to 3)

vSphere/vSAN subnets CIDR range *

192.168.100.0 / 22

IP Range: 192.168.100.0 - 192.168.103.255

HCX Deployment Network CIDR range

192.168.104.0 / 26

IP Range: 192.168.104.0 - 192.168.104.63

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

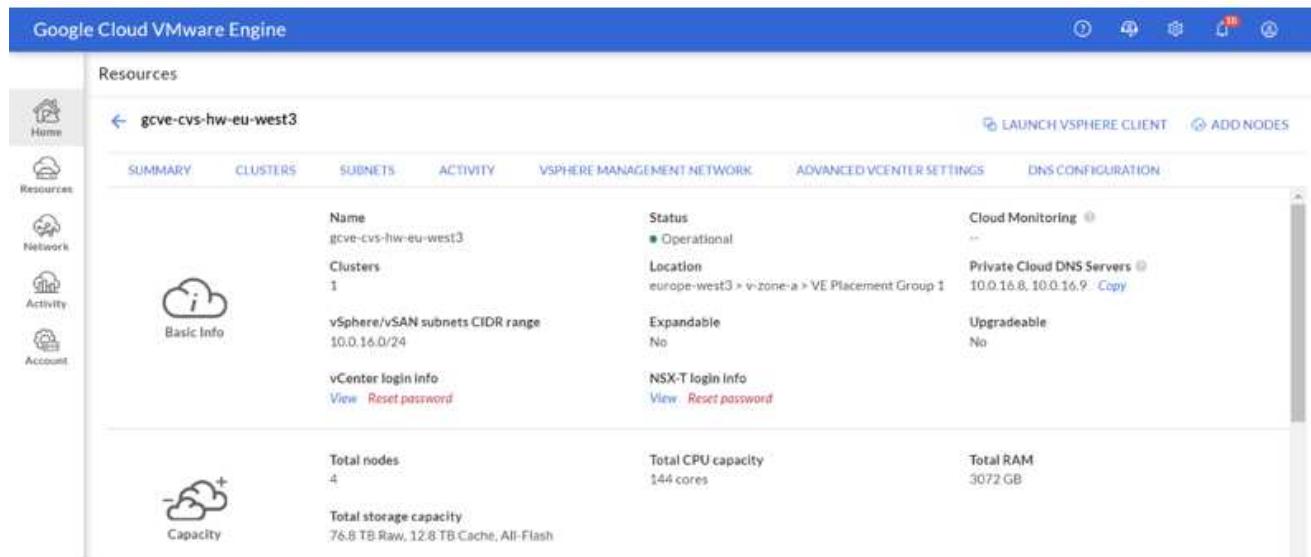
Enable Private Access to GCVE

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

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

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

Sign in to vcenter using the [CloudOwner@gve.local](#) user. To access the credentials, go to the VMware Engine portal, Go to Resources, and select the appropriate private cloud. In the Basic info section, click the View link for either vCenter login info (vCenter Server, HCX Manager) or NSX-T login info (NSX Manager).



The screenshot shows the Google Cloud VMware Engine vCenter web interface. The left sidebar has links for Home, Resources, Network, Activity, and Account. The main content area is titled 'gcve-cvs-hw-eu-west3'. It has tabs for SUMMARY, CLUSTERS, SUBNETS, ACTIVITY, VSPHERE MANAGEMENT NETWORK, ADVANCED VCENTER SETTINGS, and DNS CONFIGURATION. The SUMMARY tab is active. Under 'Basic Info', it shows the name 'gcve-cvs-hw-eu-west3', status '● Operational', location 'europe-west3 > v-zone-a > VE Placement Group 1', and a 'Private Cloud DNS Servers' section with IP addresses 10.0.16.8, 10.0.16.9. It also shows an 'Expandable' status 'No' and an 'Upgradeable' status 'No'. Under 'Capacity', it shows 'Total nodes' as 4, 'Total CPU capacity' as 144 cores, 'Total RAM' as 3072 GB, and 'Total storage capacity' as 76.8 TB Raw, 12.8 TB Cache, All-Flash.

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

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

The image shows two screenshots of the VMware vSphere interface. The top screenshot is the 'Login' screen, which includes fields for 'solution-user-01@gve.local', a password, and a checkbox for 'Use Windows session authentication'. The bottom screenshot is the 'vSphere Client' interface for the datacenter 'vcsa-57901.f7458c8f.europe-west3.gve.goog'. It displays a summary of the datacenter's resources, including its version (7.0.1), build (18392253), and last updated (Sep 22, 2021, 6:49 AM). It also shows 1 cluster, 4 hosts, and 29 virtual machines. Resource usage is shown for CPU, Memory, and Storage. The 'Datastores' section indicates 372.54 GB free CPU, 2.35 TB free Memory, and 780.88 TB free Storage.

Deploy Google Cloud NetApp Volumes supplemental datastore to GCVE

Refer [Procedure to deploy supplemental NFS datastore with NetApp Volumes to GCVE](#)

NetApp storage in public clouds

NetApp Storage options for Public Cloud Providers

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

AWS / VMC

AWS supports NetApp storage in the following configurations:

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

View the detailed [guest connect storage options for VMC](#).

View the detailed [supplemental NFS datastore options for VMC](#).

Azure / AVS

Azure supports NetApp storage in the following configurations:

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

View the detailed [guest connect storage options for AVS](#).

View the detailed [supplemental NFS datastore options for AVS](#).

GCP / GCVE

Google Cloud supports NetApp storage in the following configurations:

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

View the detailed [guest connect storage options for GCVE](#).

View the detailed [supplemental NFS datastore options for GCVE](#).

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

Amazon Web Services: Options for using NetApp storage

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

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

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

For more information, please visit:

- [FSx ONTAP as a Supplemental NFS Datastore: Overview](#)
- [Amazon FSx for ONTAP as a Supplement datastore](#)

Amazon FSx for NetApp ONTAP as guest connected storage

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

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

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

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

For more information, please visit [FSx ONTAP as Guest Connected Storage](#)

Cloud Volumes ONTAP (CVO) as guest connected storage

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

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

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

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

For more information, please visit [CVO as Guest Connected Storage](#).

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

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

Introduction

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

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

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

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



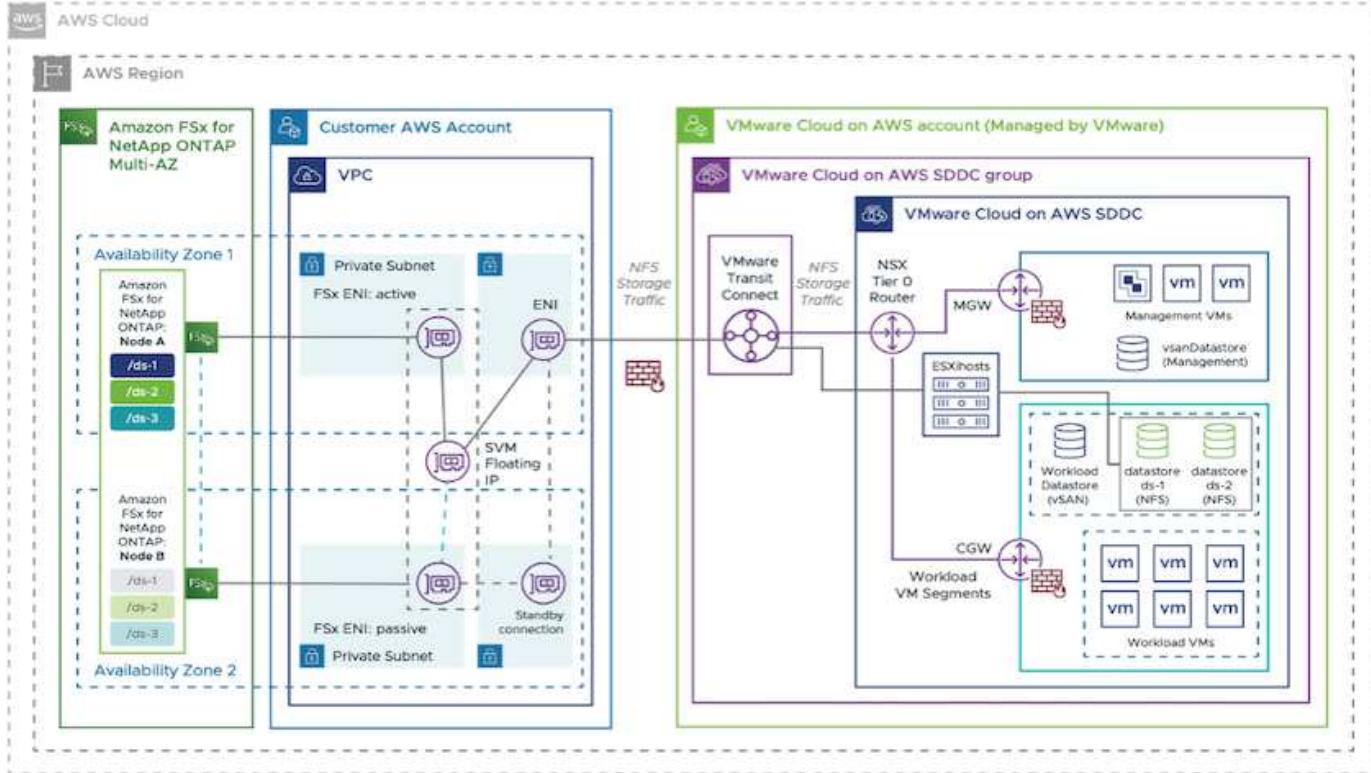
This solution is also available from VMware. Please visit the [VMware Cloud on AWS Documentation](#) for more information.

Connectivity options



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

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



The high-level deployment steps are as follows:

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

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



FSx ONTAP is not currently supported with stretched clusters.

Conclusion

This document covers the steps necessary to configure Amazon FSx ONTAP with VMware cloud on AWS. Amazon FSx ONTAP provides excellent options to deploy and manage application workloads along with file services while reducing the TCO by making data requirements seamless to the application layer. Whatever the use case, choose VMware Cloud on AWS along with Amazon FSx ONTAP for rapid realization of cloud benefits, consistent infrastructure, and operations from on-premises to AWS, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect storage. Remember, it is just the position of the data that changed along with new names; the tools and processes all remain the same, and Amazon FSx ONTAP helps to optimize the overall deployment.

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

[Amazon FSx ONTAP VMware Cloud](#)

Creating a Supplemental NFS Datastore in AWS

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

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



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

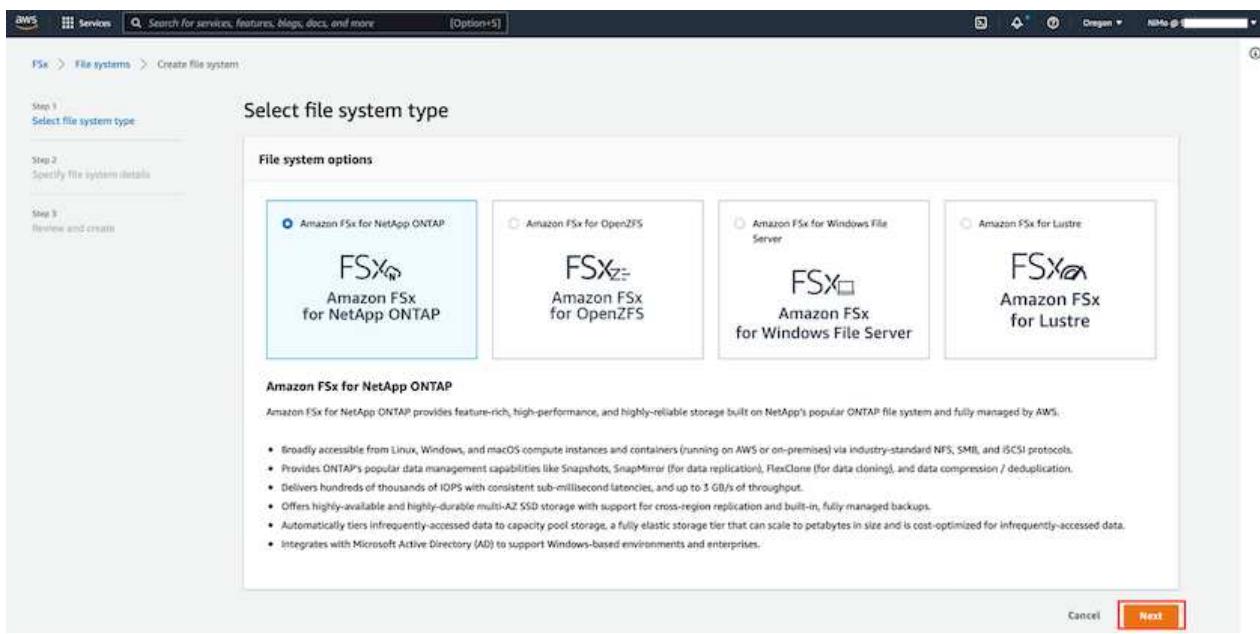


You cannot deploy FSx ONTAP in the Connected VPC. Instead, you must deploy it in a new, designated VPC and then connect the VPC to a VMware Managed Transit Gateway (vTGW) via SDDC groups.

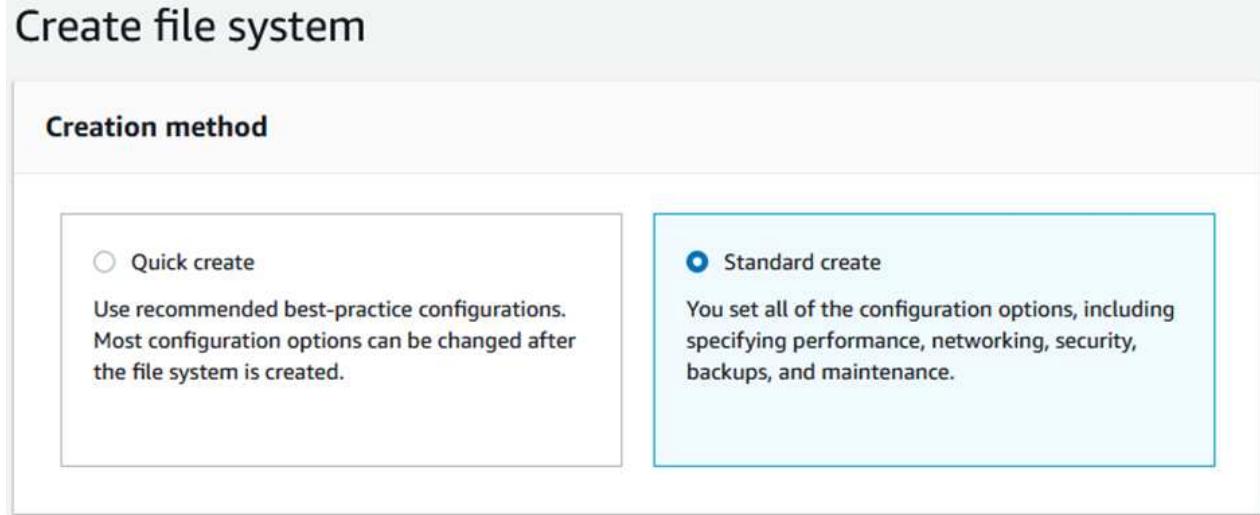
Step 1: Create Amazon FSx ONTAP in a new, designated VPC

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

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



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



File system details

File system name - optional [Info](#)

FSxONTAPDatastoreFS

Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = . _ : /

Deployment type [Info](#)

- Multi-AZ
- Single-AZ

SSD storage capacity [Info](#)

2048

Minimum 1024 GiB; Maximum 192 TiB.

Provisioned SSD IOPS

Amazon FSx provides 3 IOPS per GiB of storage capacity. You can also provision additional SSD IOPS as needed.

- Automatic (3 IOPS per GiB of SSD storage)
- User-provisioned

40000

Maximum 80,000 IOPS

Throughput capacity [Info](#)

The sustained speed at which the file server hosting your file system can serve data. The file server can also burst to higher speeds for periods of time.

- Recommended throughput capacity
128 MB/s
- Specify throughput capacity

Throughput capacity

2048 MB/s

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

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



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



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

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



As of GA, four NFS datastores are supported.

Default storage virtual machine configuration

Storage virtual machine name

FSxONTAPDatastoreSVM

SVM administrative password

Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.

Don't specify a password

Specify a password

Active Directory

Joining an Active Directory enables access from Windows and MacOS clients over the SMB protocol.

Do not join an Active Directory

Join an Active Directory

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

Setting	Configuration
Volume guarantee (Space Guarantee Style)	None (thin provisioned) – set by default
fractional_reserve (fractional-reserve)	0% – set by default
snap_reserve (percent-snapshot-space)	0%
Autosize (autosize-mode)	grow_shrink
Storage efficiency	Enabled – set by default
Autodelete	volume / oldest_first
Volume Tiering Policy	Snapshot only – set by default
try_first	Autogrow
Snapshot policy	None

Use the following SSH command to create and modify volumes:

Command to create new datastore volume from shell:

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

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

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

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

Default volume configuration

Volume name

Maximum of 203 alphanumeric characters, plus _.

Junction path

The location within your file system where your volume will be mounted.

Volume size



Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency

Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Enabled (recommended)

Disabled

Capacity pool tiering policy

You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.



▼ Backup and maintenance - *optional*

Daily automatic backup [Info](#)

Amazon FSx can protect your data through daily backups

Enabled
 Disabled

Weekly maintenance window [Info](#)

When patching needs to be performed, Amazon FSx performs maintenance on your file system only during this window.

No preference
 Select start time for 30-minute weekly maintenance window

► Tags - *optional*

Cancel

Back

Next



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

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

9. Click **Create File System**.

File system name	File system ID	File system type	Status	Deployment type	Storage type	Storage capacity	Throughput capacity	Creation time
fs-097d5cad638c5cb93	097d5cad638c5cb93	ONTAP	Creating	Multi-AZ	SSD			2022-09-12T15:19:06+01:00

File system name	File system ID	File system type	Status	Deployment type	Storage type	Storage capacity	Throughput capacity	Creation time
fs-097d5cad638c5cb93	097d5cad638c5cb93	ONTAP	Available	Multi-AZ	SSD			2022-09-12T15:49:46+01:00



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

To learn about Amazon FSx ONTAP performance, see [Amazon FSx ONTAP performance](#).

Step 2: Create SDDC group

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

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

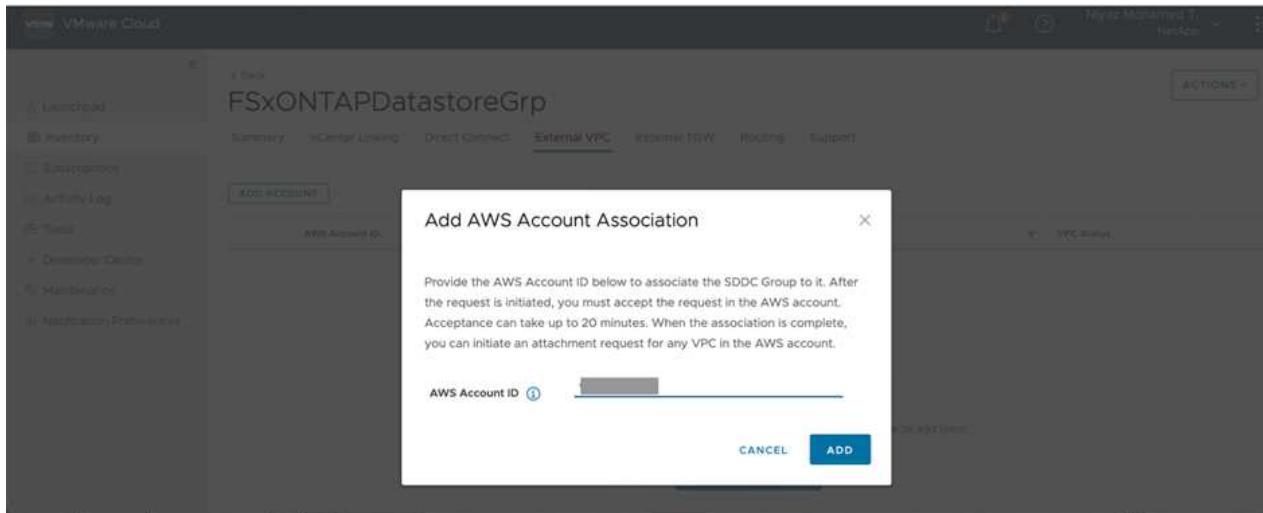
Name	Sddc Id	Location	Version	Management CIDR
FSxNDemoSDDC	c6baecd9-e01b-41d5-89e2-1f095d719a0d	US West (Oregon)	1.18.0.14	172.30.160.0/23

5. Verify that "Configuring VMware Transit Connect for your group will incur charges per attachment and data transfers" is checked, then select **Create Group**. The process can take a few minutes to complete.

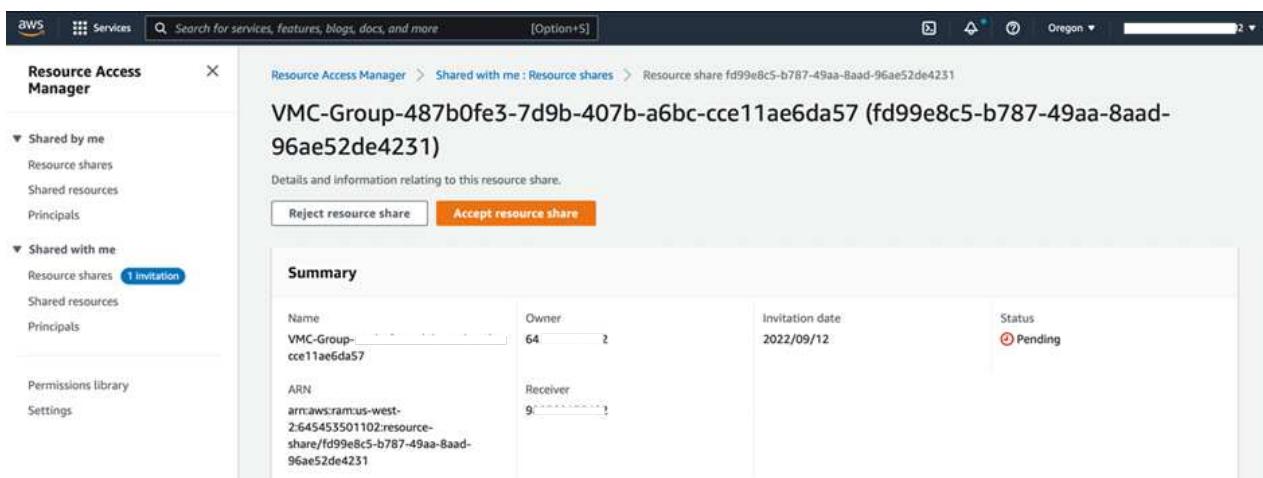
Name	SDDC ID	SDDC Version	Management CIDR	Location	Connectivity Status
FSxNDemoSDDC	c6baecd9-e01b-41d5-89e2-1f095d719a0d	1.18.0.14	172.30.160.0/23	US West (Oregon)	CONNECTED

Step 3: Configure VMware Transit connect

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

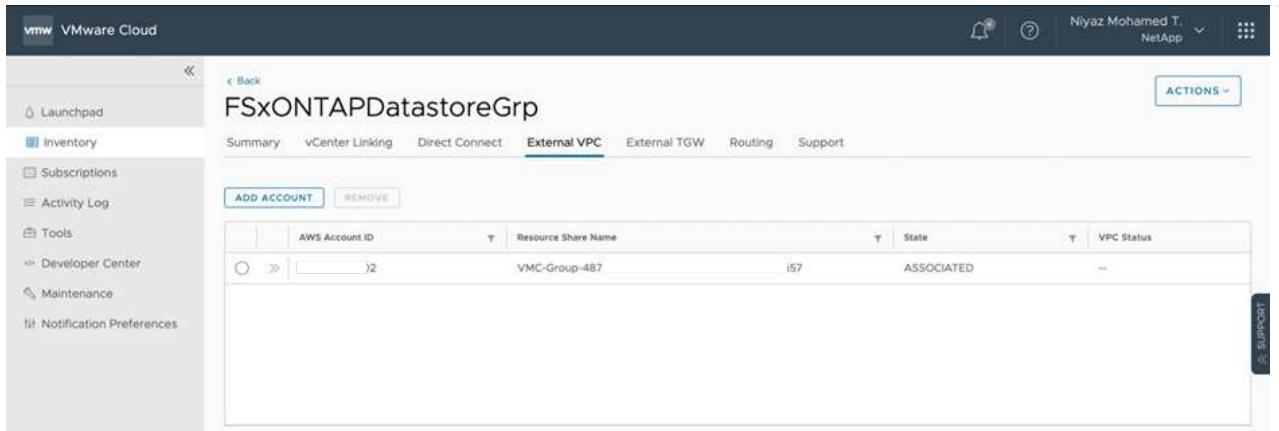


2. Click **Add Account**.
 - a. Provide the AWS account that was used to provision the FSx ONTAP file system.
 - b. Click **Add**.
3. Back in the AWS console, log into the same AWS account and navigate to the **Resource Access Manager** service page. There is a button for you to accept the resource share.



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

4. Click **Accept resource share**.



VMware Cloud

FSxONTAPDatastoreGrp

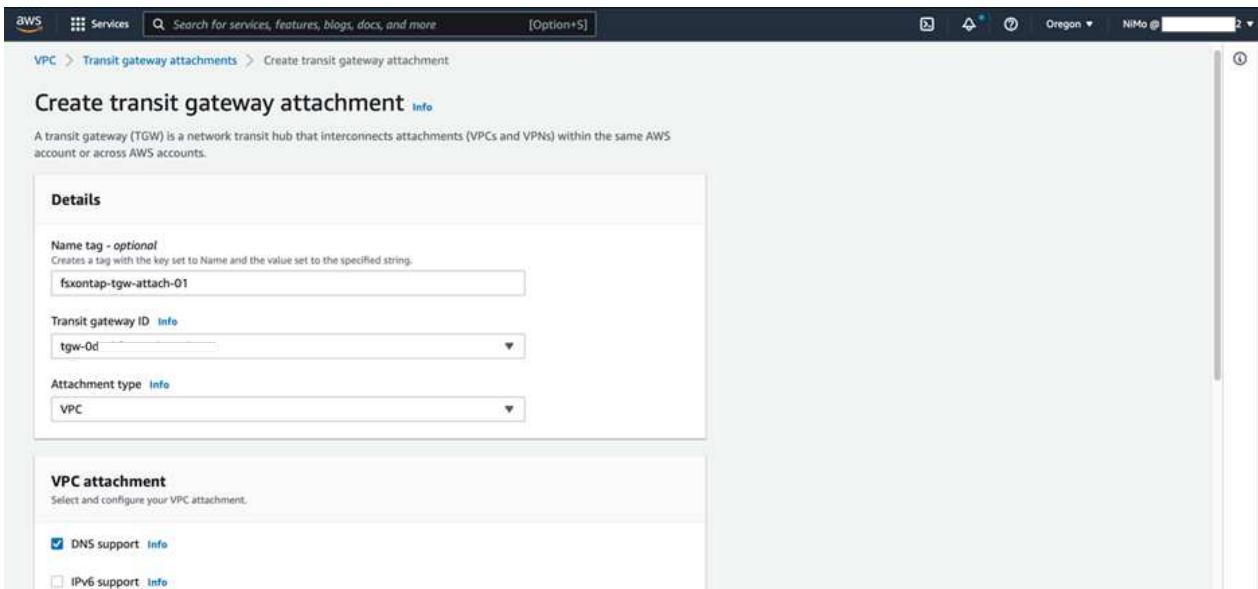
External VPC

AWS Account ID	Resource Share Name	State	VPC Status
02	VMC-Group-487	ASSOCIATED	--

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

Step 4: Create transit gateway attachment

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



Create transit gateway attachment Info

A transit gateway (TGW) is a network transit hub that interconnects attachments (VPCs and VPNs) within the same AWS account or across AWS accounts.

Details

Name tag - optional
Creates a tag with the key set to Name and the value set to the specified string.

fsxontap-tgw-attach-01

Transit gateway ID Info
tgw-0d

Attachment type Info
VPC

VPC attachment
Select and configure your VPC attachment.

DNS support Info

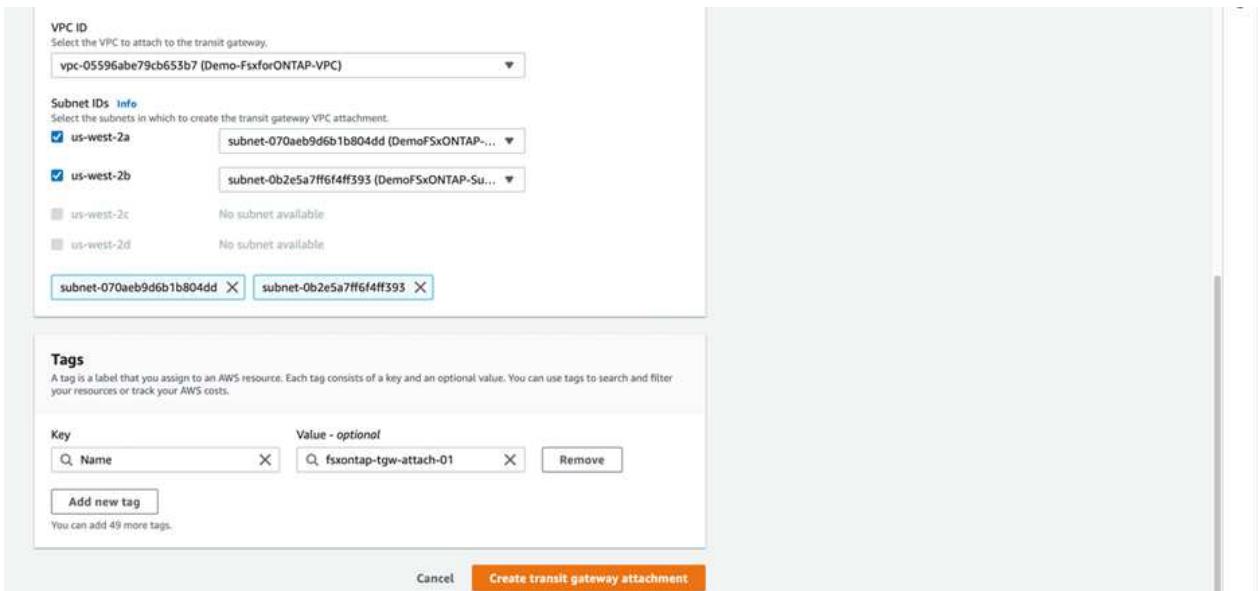
IPv6 support Info

Tags
A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

Key: Name Value - optional: fsxontap-tgw-attach-01

Add new tag

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



VPC ID
Select the VPC to attach to the transit gateway.

vp-05596abe79cb653b7 (Demo-FsxforONTAP-VPC)

Subnet IDs Info
Select the subnets in which to create the transit gateway VPC attachment.

us-west-2a subnet-070aeb9d6b1b804dd (DemoFsxONTAP-Subnet)

us-west-2b subnet-0b2e5a7ff6f4ff393 (DemoFsxONTAP-Subnet)

us-west-2c No subnet available

us-west-2d No subnet available

subnet-070aeb9d6b1b804dd subnet-0b2e5a7ff6f4ff393

Tags
A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

Key: Name Value - optional: fsxontap-tgw-attach-01

Add new tag

You can add 49 more tags.

Cancel **Create transit gateway attachment**

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

AWS Account ID	Resource Share Name	State	VPC Status
9	VMC-Group-4f	ASSOCIATED	0 Available (1 Pending Acceptance)

VPC ID	VM on AWS Region	Transit Gateway Attachment ID	Routes	Status
vpc-0	US West (Oregon)	tgw-attach-0	9	PENDING ACCEPTANCE

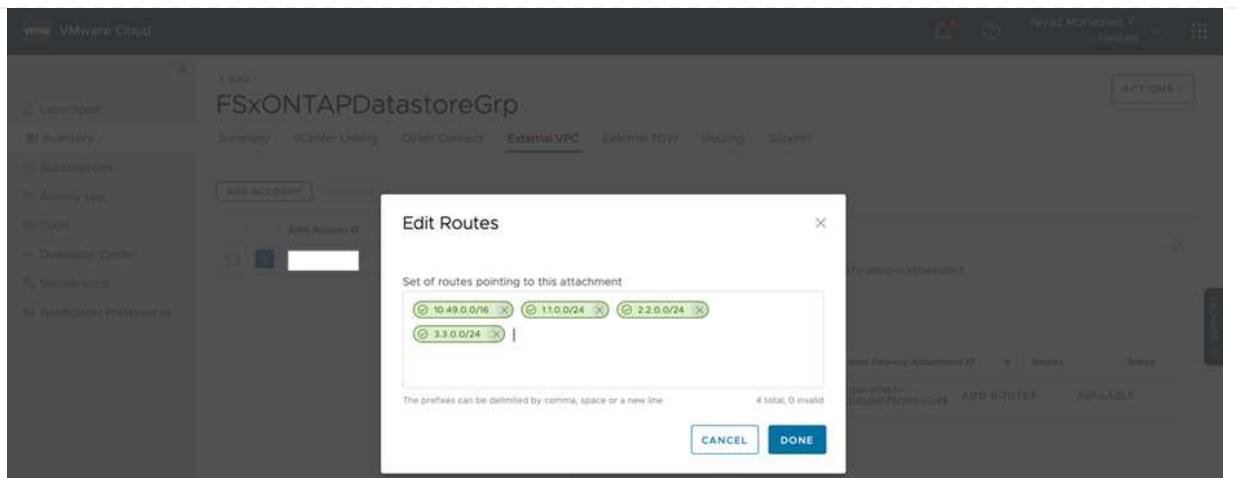


This option may take a several minutes to appear.

5. Then in the **External VPC** tab in the **Routes** column, click the **Add Routes** option and add in the required routes:

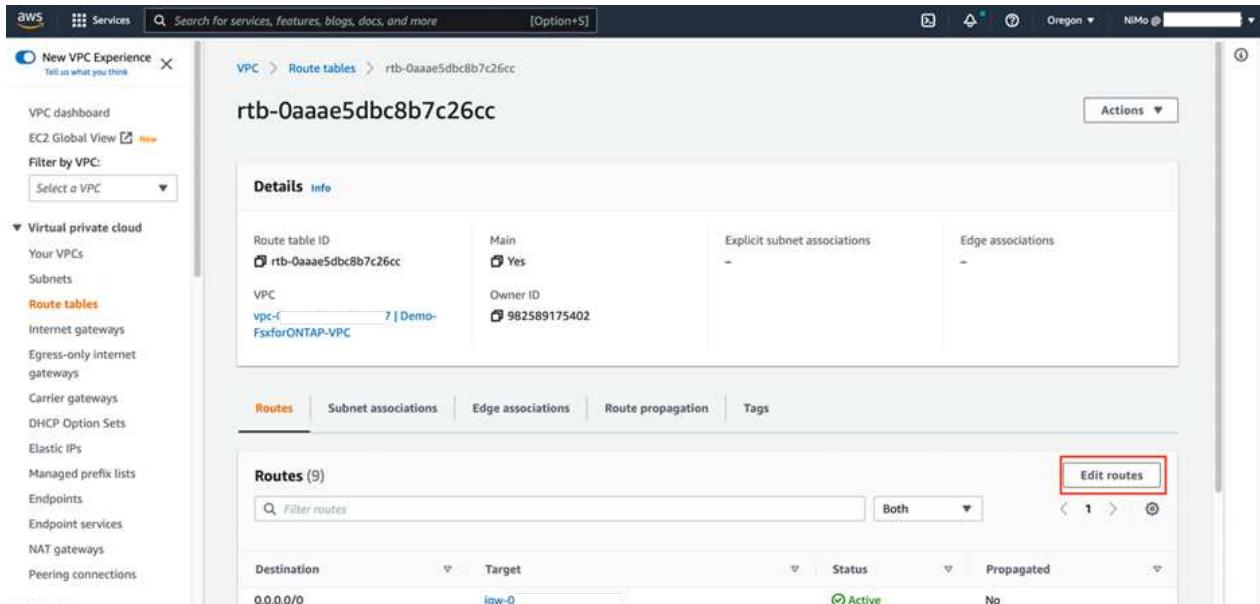
- A route for the floating IP range for Amazon FSx ONTAP floating IPs.
- A route for the newly created external VPC address space.

VPC ID	VM on AWS Region	Transit Gateway Attachment ID	Routes	Status
vpc-0	US West (Oregon)	tgw-attach-0	9	AVAILABLE



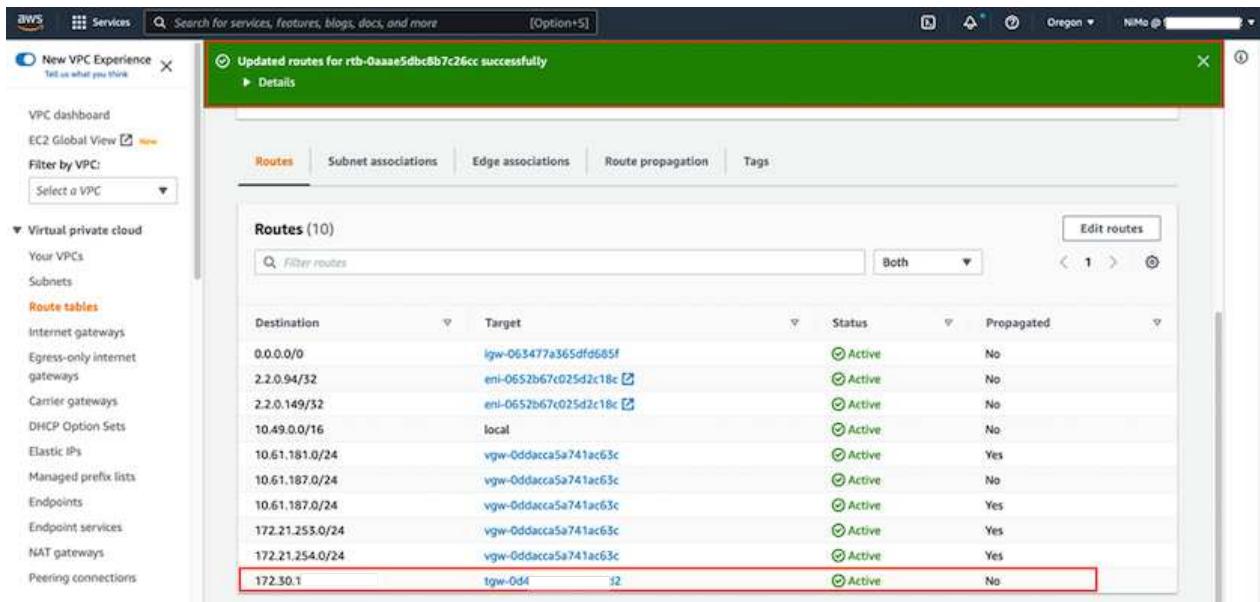
Step 5: Configure routing (AWS VPC and SDDC) and security groups

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



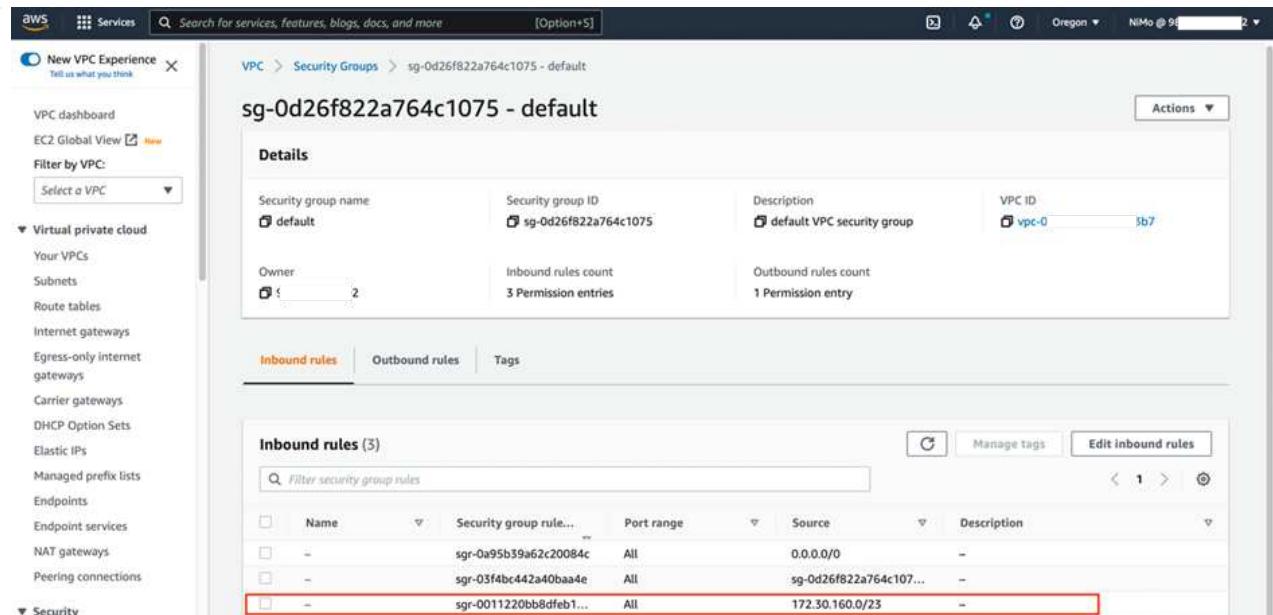
The screenshot shows the AWS VPC service page with the 'Route tables' section selected. A specific route table, 'rtb-Oaaaae5dbc8b7c26cc', is displayed. The 'Edit routes' button in the top right corner of the route table details section is highlighted with a red box.

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



The screenshot shows the same VPC Route Tables page after adding a new route. The 'Edit routes' button is highlighted with a red box. A success message at the top of the page reads: 'Updated routes for rtb-Oaaaae5dbc8b7c26cc successfully'.

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



The screenshot shows the AWS VPC Security Groups page. The left sidebar lists various VPC-related services and options. The main content area shows the details of the 'sg-0d26f822a764c1075 - default' security group. The 'Details' section includes fields for Security group name (default), Security group ID (sg-0d26f822a764c1075), Description (default VPC security group), and VPC ID (vpc-0). It also shows the Owner (AWS), Inbound rules count (3 Permission entries), and Outbound rules count (1 Permission entry). Below this, the 'Inbound rules' tab is selected, showing a table with three entries. The third entry, which is the rule allowing traffic from 172.30.160.0/23, is highlighted with a red border. The table columns are: Name, Security group rule..., Port range, Source, and Description.

Name	Security group rule...	Port range	Source	Description
-	sgr-0a95b39a62c20084c	All	0.0.0.0/0	-
-	sgr-03f4bc442a40baade	All	sg-0d26f822a764c107...	-
-	sgr-0011220bb8df...	All	172.30.160.0/23	-



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



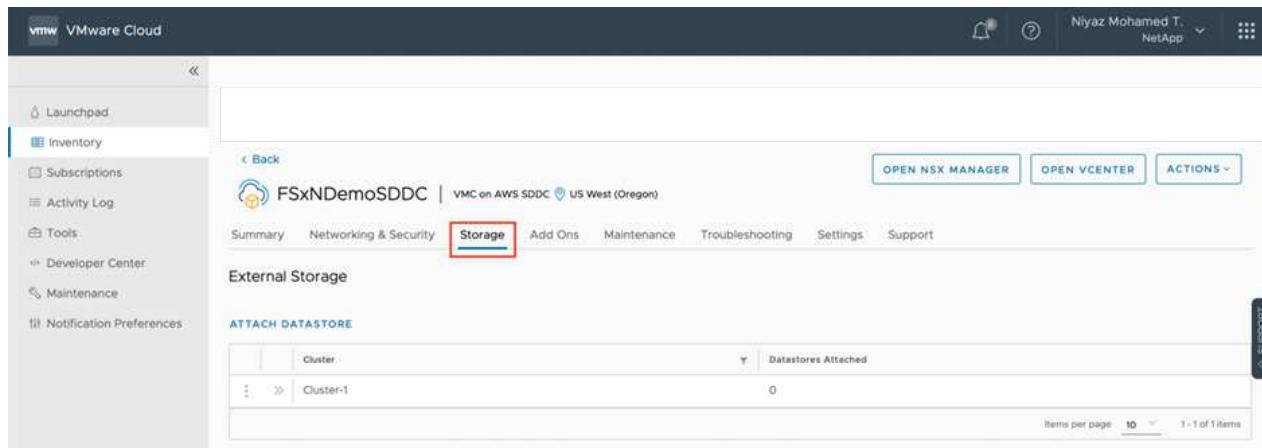
Update the security group to accept NFS traffic.

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

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

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

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

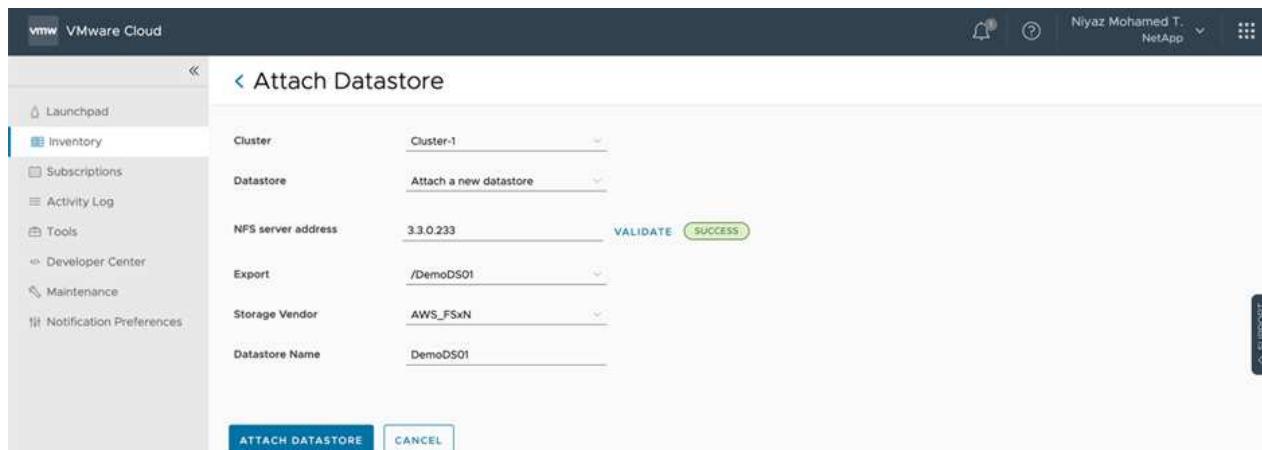


The screenshot shows the VMware Cloud Console interface. The left sidebar is titled 'Inventory' and includes 'Launchpad', 'Subscriptions', 'Activity Log', 'Tools', 'Developer Center', 'Maintenance', and 'Notification Preferences'. The main content area is titled 'FSxNDemoSDDC' and shows 'VMC on AWS SDDC' and 'US West (Oregon)'. The top navigation bar includes 'OPEN NSX MANAGER', 'OPEN VCENTER', and 'ACTIONS'. The 'Storage' tab is selected and highlighted with a red box. Below it, the 'External Storage' section is visible. A table titled 'ATTACH DATASTORE' shows two entries: 'Cluster' and 'Cluster-1', both with '0' datastores attached. The right side of the table shows 'Datastores Attached'. The bottom right corner of the table has a 'Support' link.

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



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



The screenshot shows the 'Attach Datastore' dialog box. The left sidebar is the same as the previous screenshot. The dialog box has a title 'Attach Datastore' with a back arrow. It contains the following fields:

- Cluster: Cluster-1
- Datastore: Attach a new datastore
- NFS server address: 3.3.0.233 (with a 'VALIDATE' button next to it, which has a green 'SUCCESS' message)
- Export: /DemoDS01
- Storage Vendor: AWS_FSxN
- Datastore Name: DemoDS01

At the bottom of the dialog box are 'ATTACH DATASTORE' and 'CANCEL' buttons. The right side of the dialog box has a 'Support' link.

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

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

NetApp Guest Connected Storage Options for AWS

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

FSx ONTAP

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

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

Managing your data with FSx ONTAP is easier because you can snapshot, clone, and replicate your files with

the click of a button. In addition, FSx ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

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

FSx ONTAP as guest connected storage

Configure Amazon FSx ONTAP with VMware Cloud on AWS

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

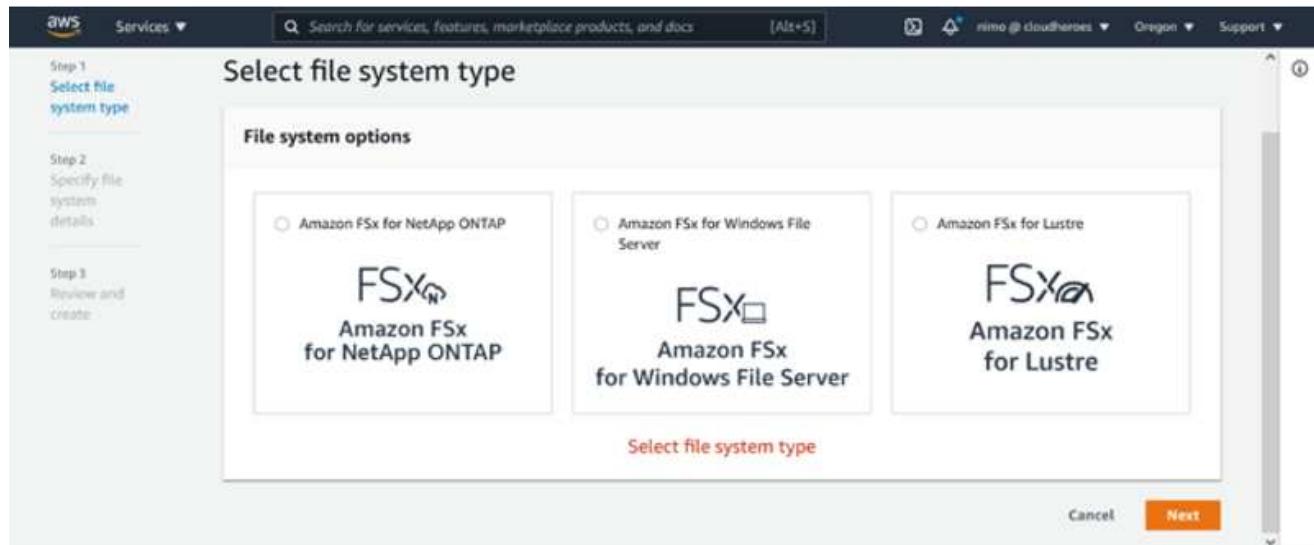


Amazon FSx ONTAP and VMware Cloud on AWS must be in the same availability zone to achieve better performance and avoid data transfer charges between availability zones.

Create and mount Amazon FSx ONTAP volumes

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

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



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

Create file system

Creation method

Quick create

Use recommended best-practice configurations. Most configuration options can be changed after the file system is created.

Standard create

You set all of the configuration options, including specifying performance, networking, security, backups, and maintenance.

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

File system details

File system name - optional [Info](#)

vmcfsxval2

Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = . _ : /

SSD storage capacity [Info](#)

1024

Minimum 1024 GB; Maximum 192 TB.

Provisioned SSD IOPS

Amazon FSx provides 3 IOPS per GB of storage capacity. You can also provision additional SSD IOPS as needed.

- Automatic (3 IOPS per GB of SSD storage)
- User-provisioned

Throughput capacity [Info](#)

The sustained speed at which the file server hosting your file system can serve data. The file server can also burst to higher speeds for periods of time.

512 MB/s (Recommended)

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

Network & security

Virtual Private Cloud (VPC) [Info](#)

Specify the VPC from which your file system is accessible.

vmcfsx2.vpc | vpc-0d1c764bcc495e805

VPC Security Groups [Info](#)

Specify VPC Security Groups to associate with your file system's network interface.

Choose VPC security group(s)

sg-018896ea218164ccb (default) 

Preferred subnet [Info](#)

Specify the preferred subnet for your file system.

subnet02.sn | subnet-013675849a5b99b3c (us-west-2b)

Standby subnet

subnet01.sn | subnet-0ef956cebf539f970 (us-west-2a)

VPC route tables

Specify the VPC route tables associated with your file system.

- VPC's default route table
- Select one or more VPC route tables

Endpoint IP address range

Specify the IP address range in which the endpoints to access your file system will be created.

- No preference
- Select an IP address range



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

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

Security & encryption

Encryption key [Info](#)

AWS Key Management Service (KMS) encryption key that protects your file system data at rest.

aws/fsx (default)

Description	Account	KMS key ID
Default master key that protects my FSx resources when no other key is defined	139763910815	72745367-7bb0-499c-acc0-4f2c0a80e7c5

File system administrative password

Password for this file system's "fsxadmin" user, which you can use to access the ONTAP CLI or REST API.

Don't specify a password
 Specify a password

Password

Confirm password

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

Default storage virtual machine configuration

Storage virtual machine name

vmcfsxval2svm

SVM administrative password

Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.

- Don't specify a password
- Specify a password

Password

Confirm password

Active Directory

Joining an Active Directory enables access from Windows and MacOS clients over the SMB protocol.

- Do not join an Active Directory
- Join an Active Directory

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

Default volume configuration

Volume name

vol1

Maximum of 203 alphanumeric characters, plus _.

Junction path

/vol1

The location within your file system where your volume will be mounted.

Volume size

1024



Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency

Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

- Enabled (recommended)
- Disabled

Capacity pool tiering policy

You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

Auto



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

Screenshot of the AWS FSx console showing the creation of a Storage Virtual Machine (SVM).

File systems (3)

File system name	File system ID	File system type	Status	Deployment type	Storage type	Size
fsxntapcifs	fs-014c28399be9c1f9f	ONTAP	Available	Multi-AZ	SSD	1.00 TB
vmcfsxval2	fs-040eacc5d0ac31017	ONTAP	Available	Multi-AZ	SSD	1.00 TB
fsxntapsql	fs-0ab4b447ebd6082aa	ONTAP	Available	Multi-AZ	SSD	2.00 TB

Storage virtual machines (SVMs) (2)

SVM name	SVM ID	Status	Creation time	Active Directory
fsxsmbtesting01	svm-075dcfbe2cfa2ece9	Created	2021-10-19 15:17:08 UTC +01:00	FSXTESTING.LOCAL
vmcfsxval2svm	svm-095db076341561212	Created	2021-10-15 15:16:54 UTC +01:00	-

fsxsmbtesting01 (svm-075dcfbe2cfa2ece9)

Summary

SVM ID svm-075dcfbe2cfa2ece9	Creation time 2021-10-19T15:17:08+01:00	Active Directory FSXTESTING.LOCAL
SVM name fsxsmbtesting01	Lifecycle state Created	Net BIOS name FSXSMBTESTING01
UUID 4a50e659-30e7-11ec-ac4f-f3ad92a6a735	Subtype DEFAULT	Fully qualified domain name FSXTESTING.LOCAL
File system ID fs-040eacc5d0ac31017		Service account username administrator
		Organizational unit distinguished name CN=Computers

For more detailed information, see [Getting started with Amazon FSx ONTAP](#).

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

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

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

Create volume

File system
fs-040eacc5d0ac31017 | vmcfsxval2

Storage virtual machine
svm-095db076341561212 | vmcfsxval2svm

Volume name
nfsdemovol01

Maximum of 205 alphanumeric characters, plus _.

Junction path
/nfsdemovol01

The location within your file system where your volume will be mounted.

Volume size
1024

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency
Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Enabled (recommended)

Disabled

Capacity pool tiering policy
You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

Auto

Cancel Confirm

Mount FSx ONTAP volume on Linux client

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

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

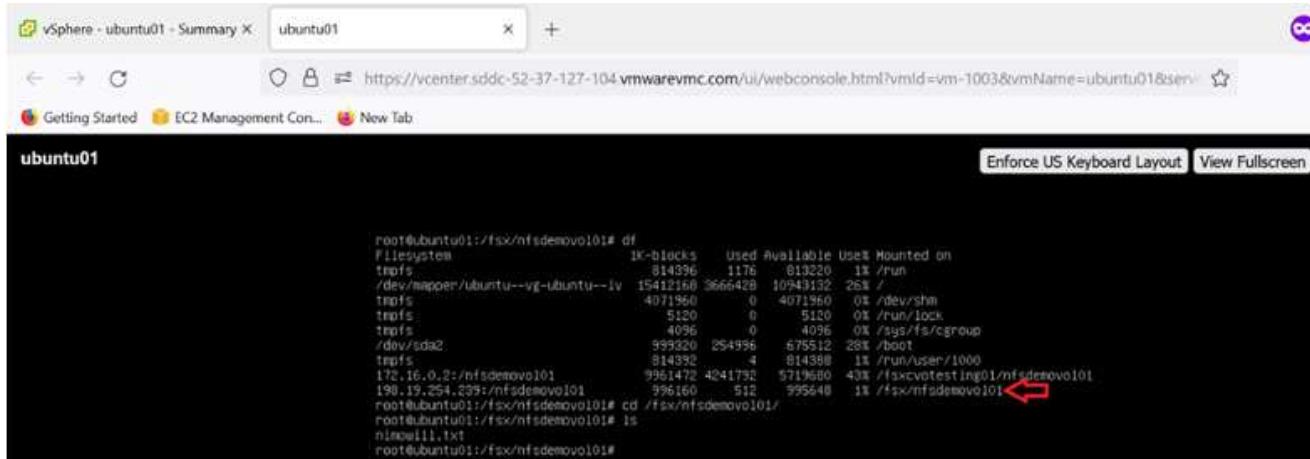
```
$ sudo mkdir /fsx/nfsdemovol01
```

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

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

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

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



```
root@ubuntu01:/fsx/nfsdemovol01# df  
Filesystem 1K-blocks Used Available Use% Mounted on  
tmpfs 814396 1176 813220 1% /run  
/dev/mapper/ubuntu--vg-ubuntu--lv 15412160 3666428 10943132 26% /  
tmpfs 4071960 0 4071960 0% /dev/shm  
tmpfs 5120 0 5120 0% /run/lock  
tmpfs 4096 0 4096 0% /sys/fs/cgroup  
/dev/sda2 99320 254936 675512 28% /boot  
tmpfs 814332 4 814388 1% /run/user/1000  
198.19.254.239:/nfsdemovol01 9961472 4041752 5719680 43% /sixxvotesting01/nfsdemovol01  
198.19.254.239:/nfsdemovol01 396160 512 395648 1% /fsx/nfsdemovol01
```

Mount FSx ONTAP volume on Linux client

Attach FSx ONTAP volumes to Microsoft Windows clients

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

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



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

Endpoints

Management DNS name	Management IP address
svm-075dcfbe2cfa2ece9.fs-040eacc5d0ac31017.fsx.us-west-2.amazonaws.com	198.19.254.9
NFS DNS name	NFS IP address
svm-075dcfbe2cfa2ece9.fs-040eacc5d0ac31017.fsx.us-west-2.amazonaws.com	198.19.254.9
SMB DNS name	SMB IP address
FSXSMBTESTING01.FSXTESTING.LOCAL	198.19.254.9
iSCSI DNS name	iSCSI IP addresses
iscsi.svm-075dcfbe2cfa2ece9.fs-040eacc5d0ac31017.fsx.us-west-2.amazonaws.com	10.222.2.224, 10.222.1.94

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

Computer Management

File Action View Help

Computer Management (FSXSMBTESTING01.FSXTESTING.LOCAL)

System Tools

- Task Scheduler
- Event Viewer
- Shared Folders
 - Shares
 - Sessions
 - Open Files
- Local Users and Groups
- Performance
- Device Manager

Storage

- Windows Server Backup
- Disk Management

Services and Applications

Share Name	Folder Path	Type	# Client Connections	Description
c\$	C:\	Windows	0	
ipc\$		Windows	1	
smbdemo...	C:\smbdemovol01	Windows	1	
testnimvol	C:\testnimvol	Windows	0	

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

Create A Shared Folder Wizard

Name, Description, and Settings

Specify how people see and use this share over the network.

Type information about the share for users. To modify how people use the content while offline, click Change.

Share name:

Share path:

Description:

Offline setting:

< Back Cancel

Create A Shared Folder Wizard

Sharing was Successful



Status:

You have successfully completed the Share a Folder Wizard.

Summary:

You have selected the following share settings on \\FSXSMBTESTING01.FSXTESTING.LOCAL:
Folder path: C:\\nimtestsmb01
Share name: nimtestsmb01
Share path: \\FSXSMBTESTING01.FSXTESTING.LOCAL\\nimtestsmb01

When I click Finish, run the wizard again to share another folder

To close this wizard, click Finish.

Finish

Cancel

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

1. After connectivity is in place, the SMB share can be attached and used for application data. To accomplish this, Copy the share path and use the Map Network Drive option to mount the volume on the VM running on VMware Cloud on the AWS SDDC.

The screenshot shows the Windows File Explorer interface within a browser window. The address bar shows the URL: https://vcenter.sddc-52-37-127-104.vmwarevmc.com/u/webconsole.html?vmId=vm-1005&vmName=vmcd01. The left sidebar shows the 'Computer Management' interface with various tools like Task Scheduler, Event Viewer, and Shared Folders. The 'Shared Folders' section is expanded, showing a list of shares. One share, 'nimtestsmb01', is highlighted with a red box and an arrow pointing to it. The share path is 'This PC > enddemovm01 (C:\192.168.10.254\0)'. The contents of this share are listed in the main pane, showing subfolders 'new_folder01', 'nimtest01', 'nimtest02', and 'nimtest03'. The 'Actions' column on the right shows options like 'Share' and 'More Actions' for each folder.

Connect a FSx ONTAP LUN to a host using iSCSI

Connect a FSx ONTAP LUN to a host using iSCSI

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

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

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

Provision a LUN in FSx ONTAP:

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

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

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

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

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

```
FsxId040eacc5d0ac31017::> igrup show
```

Vserver	Igroup	Protocol	OS	Type	Initiators
---------	--------	----------	----	------	------------

```
-----
```

```
-----
```

```
vmcfsxval2svm
```

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

```
vmcfsxval2svm
```

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

Two entries were displayed.

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

```

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

FsxId040eacc5d0ac31017::> lun show

Vserver      Path          State   Mapped   Type
Size

-----
-----
vmcfsxval2svm

/vol/blocktest01/lun01      online  mapped   linux
5GB

vmcfsxval2svm

/vol/nimfsxscsivol/nimofsxlun01  online  mapped   windows
5GB

```

Two entries were displayed.

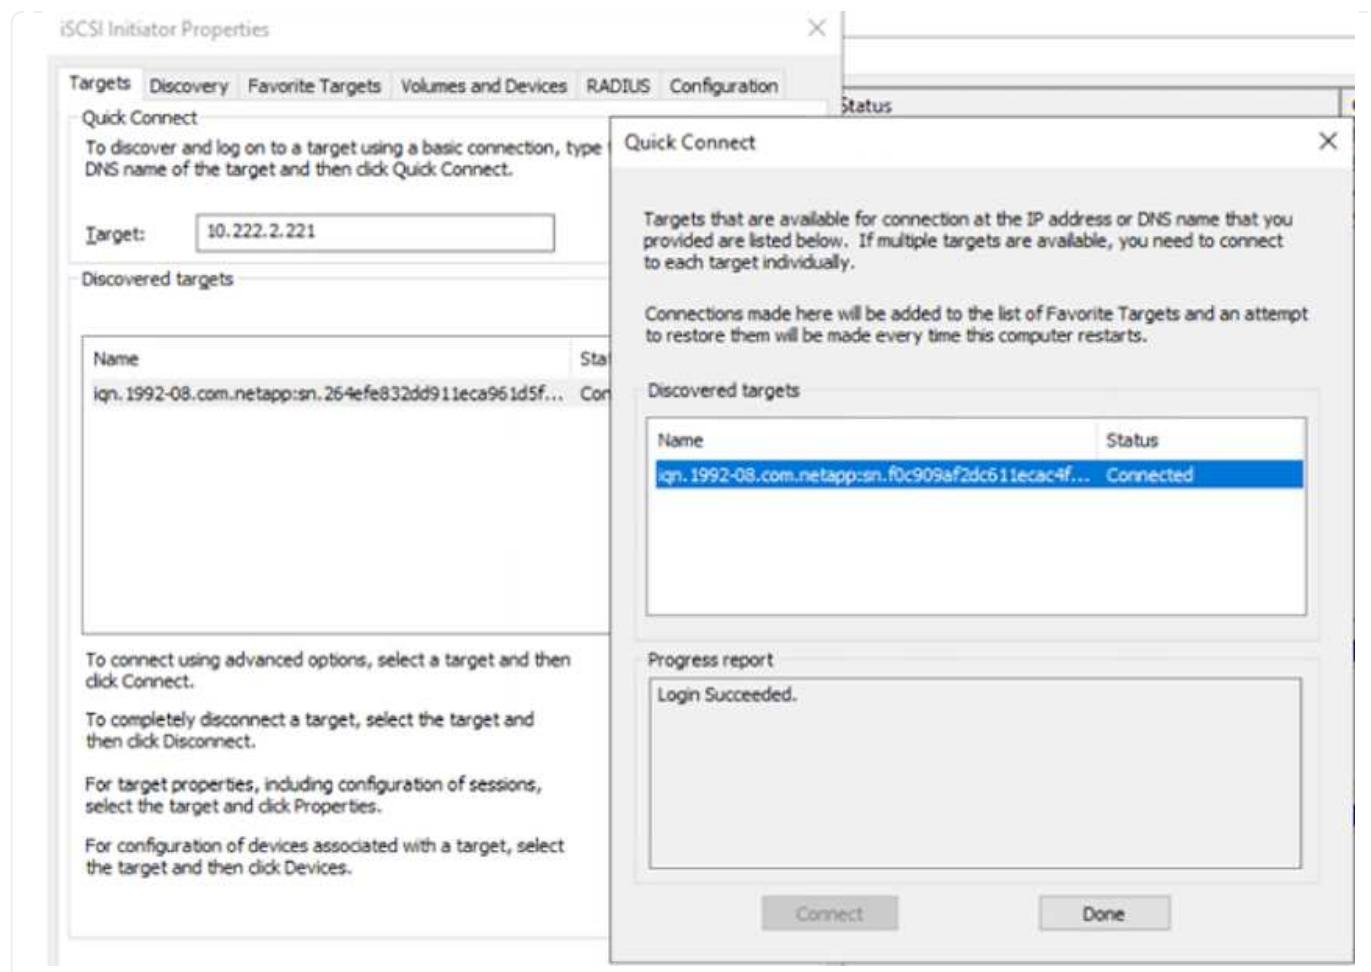
1. Connect the newly provisioned LUN to a Windows VM:

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

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

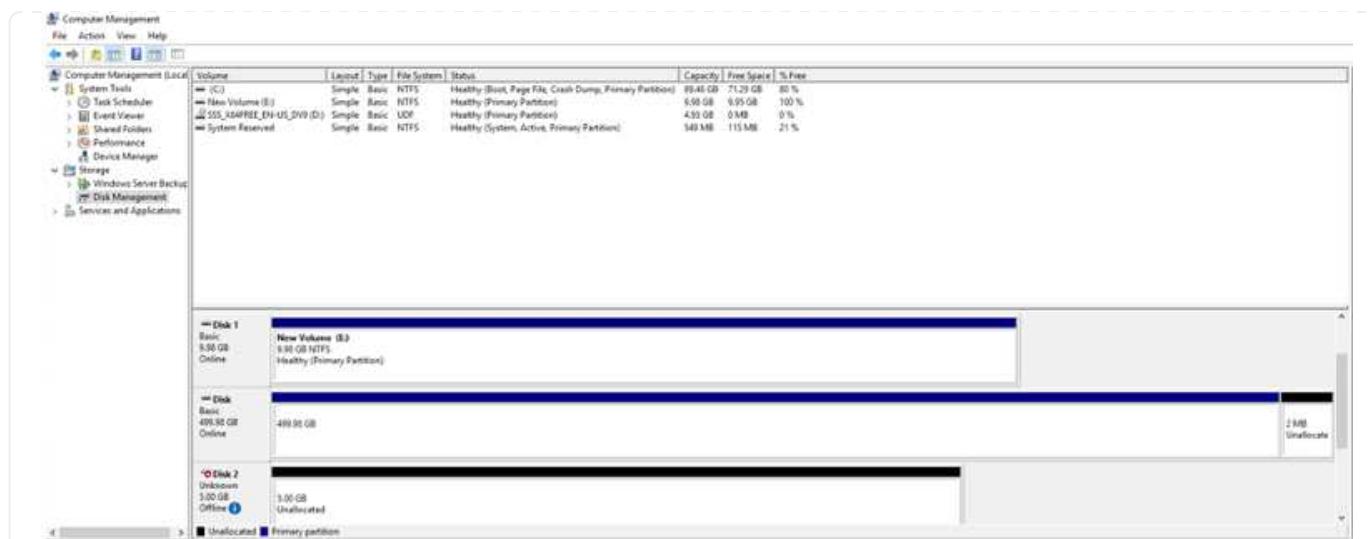


The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



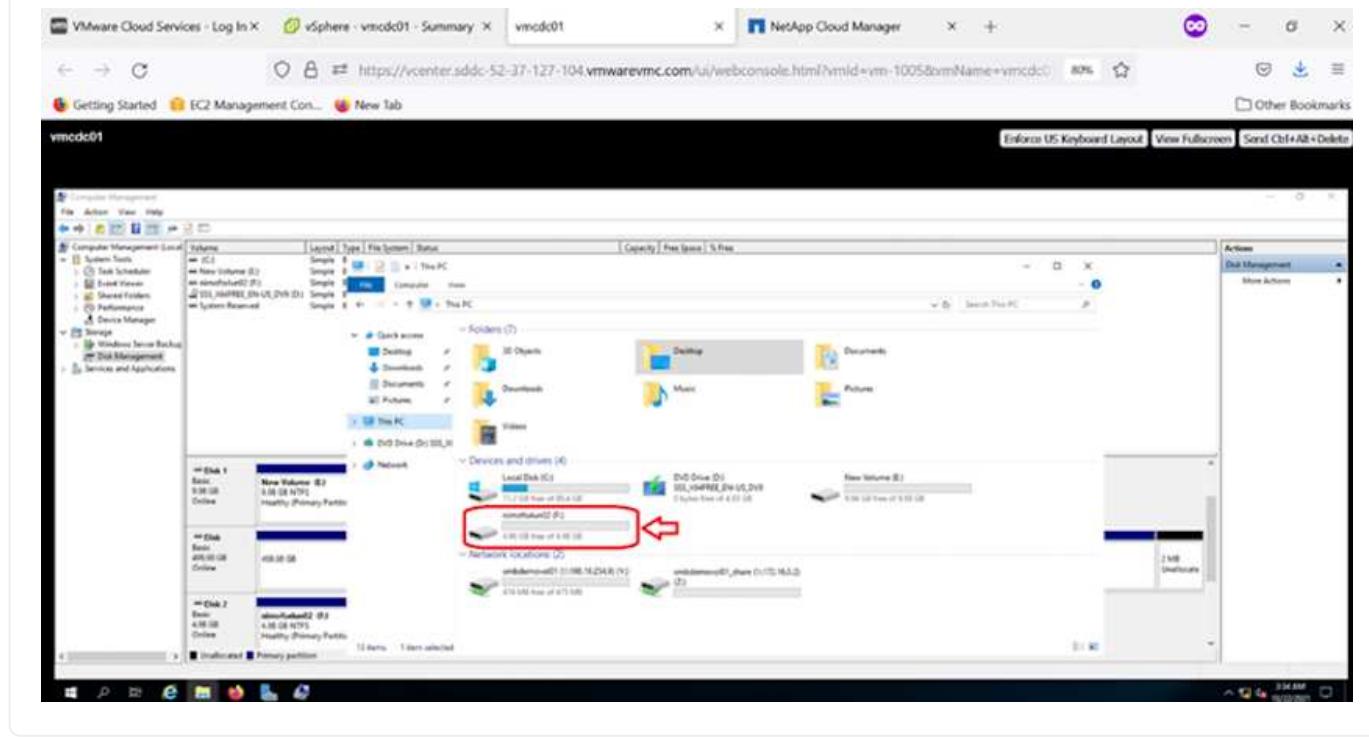
LUNs on the storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

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



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

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



Cloud Volumes ONTAP (CVO)

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

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

manage your data.

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

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

Cloud Volumes ONTAP (CVO) as guest connected storage

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

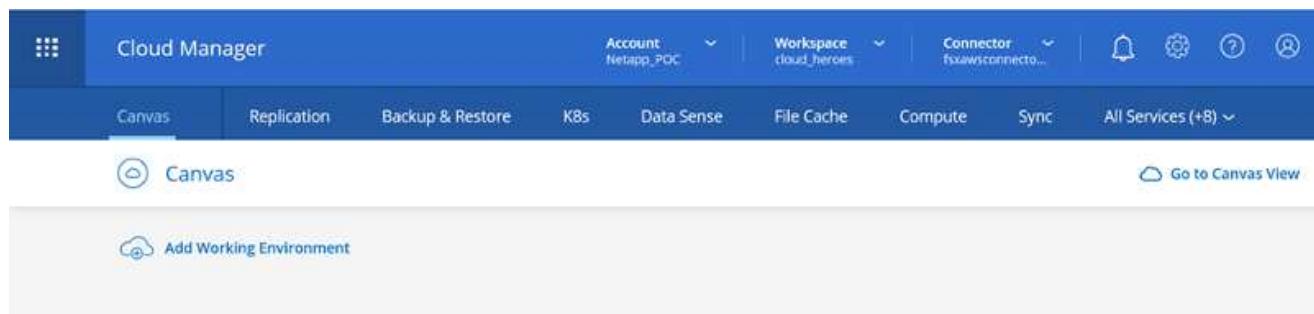
Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the VMware Cloud on AWS SDDC environment. The volumes can also be mounted on native AWS VM Linux Windows clients, and LUNs can be accessed on Linux or Windows clients as block devices when mounted over iSCSI because Cloud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

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

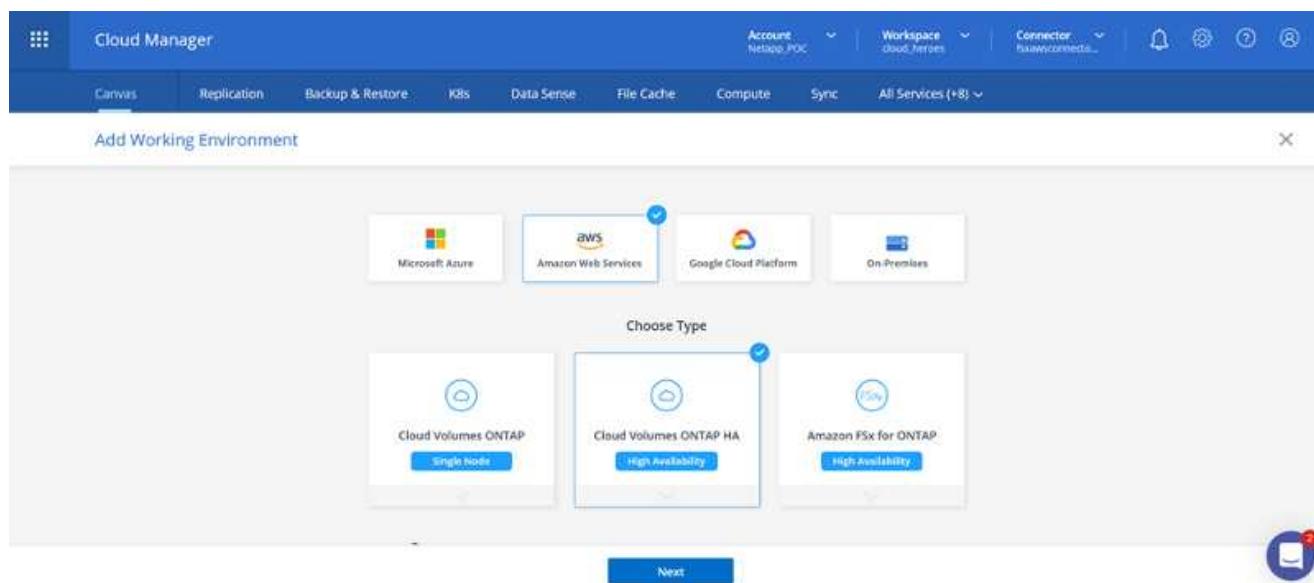


Use the [Cloud Volumes ONTAP sizer](#) to accurately size the Cloud Volumes ONTAP instances. Also, monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

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



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



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

Create a New Working Environment

Details and Credentials

↑ Previous Step	Instance Profile Credential Name	139763910815 Account ID	netapp.com-cloud-volumes-... Marketplace Subscription	Edit Credentials
-----------------	-------------------------------------	----------------------------	--	----------------------------------

Details	Credentials
Working Environment Name (Cluster Name) fsxvotesting01	User Name admin
Add Tags Optional Field Up to four tags	Password *****
	Confirm Password *****

[Continue](#)

1. Select the add-on services for Cloud Volumes ONTAP deployment, including BlueXP Classification, BlueXP backup and recovery, and Cloud Insights. Click Continue.

Create a New Working Environment

Services

 Data Sense & Compliance	<input checked="" type="checkbox"/> View
 Backup to Cloud	<input checked="" type="checkbox"/> View
 Monitoring	<input checked="" type="checkbox"/> View

[Continue](#)

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

Create a New Working Environment

HA Deployment Models

↑ Previous Step	<p>Multiple Availability Zones</p> <ul style="list-style-type: none">  Provides maximum protection against AZ failures.  Enables selection of 3 availability zones.  An HA node serves data if its partner goes offline. <p>Extended Info</p>	<p>Single Availability Zone</p> <ul style="list-style-type: none">  Protects against failures within a single AZ.  Single availability zone. HA nodes are in a placement group, spread across distinct underlying hardware.  An HA node serves data if its partner goes offline. <p>Extended Info</p>
-----------------	--	---

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

[↑ Previous Step](#)

AWS Region

US West | Oregon

VPC

vpc-0d1c764bcc495e805 -
10.222.0.0/16

Security group

Use a generated security group



Node 1:

Availability Zone

us-west-2a

Subnet

10.222.1.0/24



Node 2:

Availability Zone

us-west-2b

Subnet

10.222.2.0/24



Mediator:

Availability Zone

us-west-2c

Subnet

10.222.3.0/24

[Continue](#)

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

[↑ Previous Step](#)

Nodes



Mediator

SSH Authentication Method

Password



Nodes

Security Group

Use a generated security group

Key Pair Name

nimokey

Internet Connection Method

Public IP address

[Continue](#)

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

[↑ Previous Step](#)

Floating IP addresses are required for cluster and SVM access and for NFS and CIFS data access. These floating IPs can migrate between HA nodes if failures occur. To access the data from outside the VPC, [you can set up an AWS transit gateway](#).

You must specify IP addresses that are outside of the CIDR blocks for all VPCs in the selected AWS region.

Floating IP address for cluster management

172.16.0.1

Floating IP address 1 for NFS and CIFS data

172.16.0.2

Floating IP address 2 for NFS and CIFS data

172.16.0.3

Floating IP address for SVM management (Optional)

172.16.0.4

[Continue](#)

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

[↑ Previous Step](#)

Select the route tables that should include routes to the floating IP addresses. This enables client access to the Cloud Volumes ONTAP HA pair. If you leave a route table unselected, clients that are associated with the route table cannot access the HA pair.

[Additional information](#) 

Name	Main	ID	Associate with Subnet	Tags
<input checked="" type="checkbox"/>	Yes	rtb-00b2d30c3f68fdbdd	0 Subnets	1 Tags

[1 Route Tables](#) | The main route table is the default for the VPC

[Continue](#)

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

[↑ Previous Step](#) AWS Managed Encryption

AWS is responsible for data encryption and decryption operations. Key management is handled by AWS key management services.

Default Master Key: **aws/ebs** [Change Key](#)[Continue](#)

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

Create a New Working Environment [Cloud Volumes ONTAP Charging Methods & NSS Account](#)

Cloud Volumes ONTAP Charging Methods

[Learn more about our charging methods](#) Pay-As-You-Go by the hour Bring your own licenseNetApp Support Site Account *(Optional)*[Learn more about NetApp Support Site \(NSS\) accounts](#)

To register this Cloud Volumes ONTAP to support, you should add NetApp Support Site Account.

Don't have a NetApp Support Site account? Select go to finish deploying this system. After it's created, use the [Support Registration](#) option to create an NSS account.

[Continue](#)

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



Select a preconfigured Cloud Volumes ONTAP system that best matches your needs, or create your own configuration. Preconfigured settings can be modified at a later time.

[Change Configuration](#)

POC and small workloads
Up to 500GB of storage



Database and application data production workloads



Cost effective DR
Up to 500GB of storage



Highest performance production workloads

[Continue](#)

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

Create a New Working Environment Review & Approve

↑ Previous Step **fsxvotesting** Show API request

AWS | us-west-2 | HA

This Cloud Volumes ONTAP instance will be registered with NetApp support under the NSS Account **mchad**.

I I understand that Cloud Manager will allocate the appropriate AWS resources to comply with my above requirements. [More information >](#)

Overview **Networking** **Storage**

Storage System:	Cloud Volumes ONTAP HA	HA Deployment Model:	Multiple Availability Zones
License Type:	Cloud Volumes ONTAP Explore	Encryption:	AWS Managed
Capacity Limit:	2TB	Customer Master Key:	aws/ebs

Go

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

Canvas Replication Backup & Restore K8s Data Sense File Cache Compute Sync All Services (+8) ▾

Canvas Go to Tabular View

Add Working Environment

fsxvotesting01 On

Cloud Volumes ONTAP | AWS | HA

DETAILS

SERVICES

fsxvotesting01 Cloud Volumes ONTAP

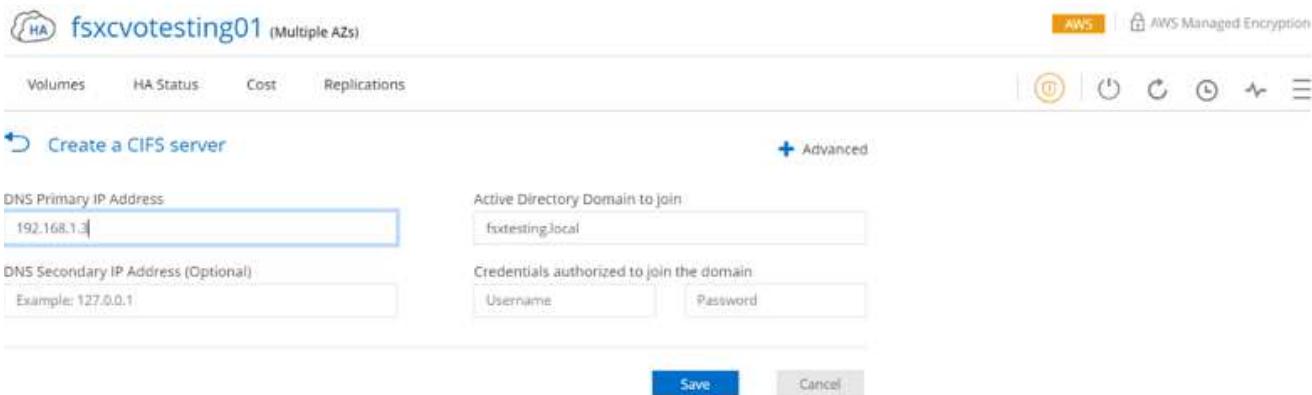
46.08 Capacity

Amazon S3

4 Buckets **2** Regions

Additional configurations for SMB volumes

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



DNS Primary IP Address: 192.168.1.3

Active Directory Domain to join: foxtesting.local

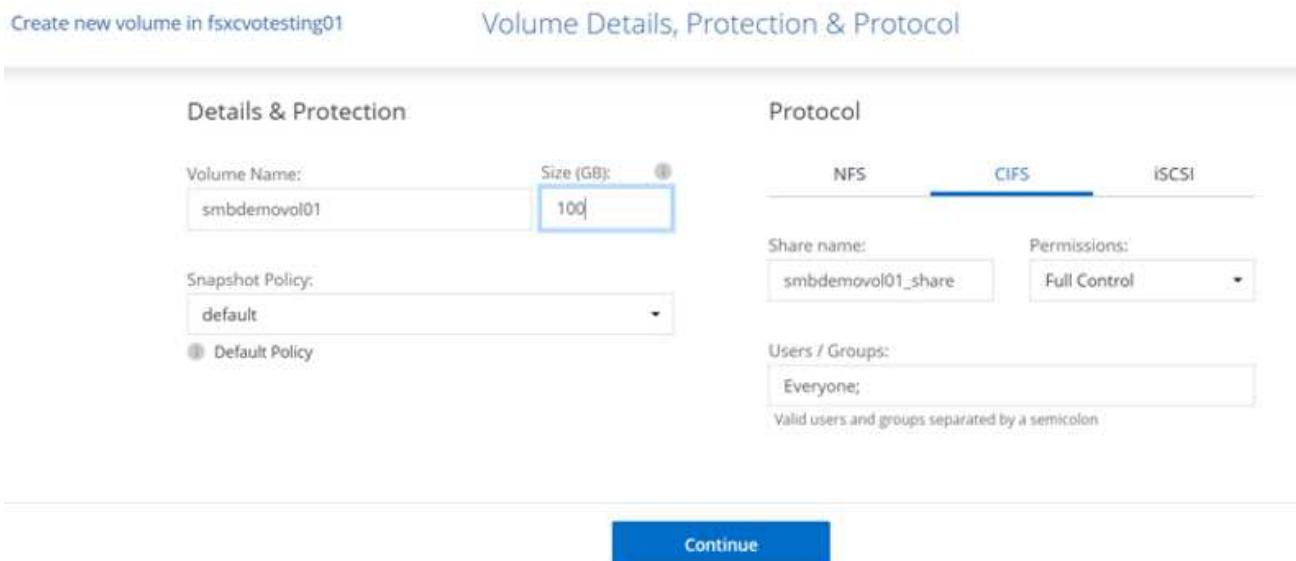
DNS Secondary IP Address (Optional): Example: 127.0.0.1

Credentials authorized to join the domain:

Username: Password:

Save Cancel

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



Create new volume in fsxvotesting01

Volume Details, Protection & Protocol

Details & Protection

Volume Name: smbdemovol01

Size (GB): 100

Snapshot Policy: default

Default Policy

Protocol

NFS CIFS iSCSI

Share name: smbdemovol01_share

Permissions: Full Control

Users / Groups: Everyone

Valid users and groups separated by a semicolon

Continue

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



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

fsxcvotesting01 (Multiple AZs)

AWS AWS

Volumes HA Status Cost Replications

Mount Volume **smbdemovol01**

Access from inside the VPC using Floating IP

Auto failover between nodes
The IP address automatically migrates between nodes if failures occur

Go to your machine and enter this command

`\\"172.16.0.2\\smbdemovol01_share`

Copy

Access from outside the VPC using AWS Private IP

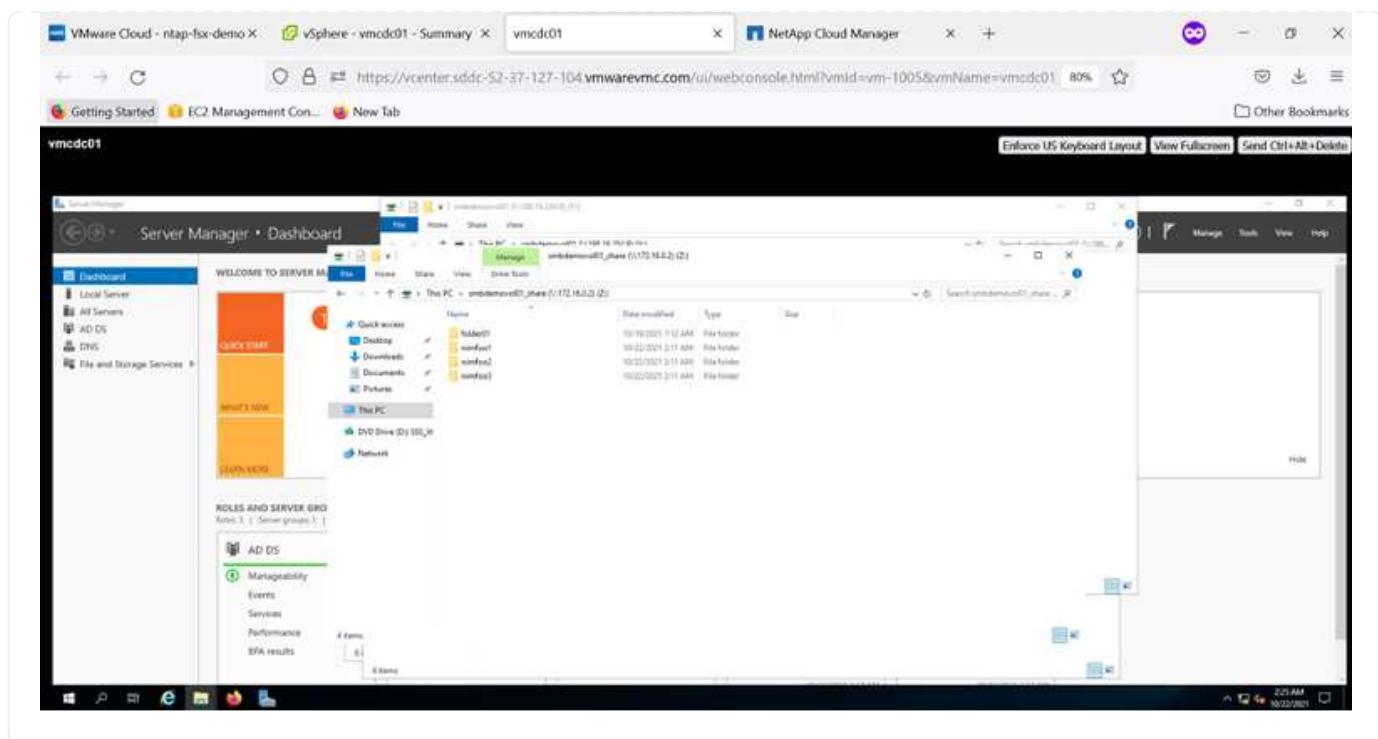
No auto failover between nodes
The IP address does not migrate between nodes if failures occur

To avoid traffic between nodes, mount the volume by using the primary node's IP address:

`\\"10.222.1.100\\smbdemovol01_share`

Copy

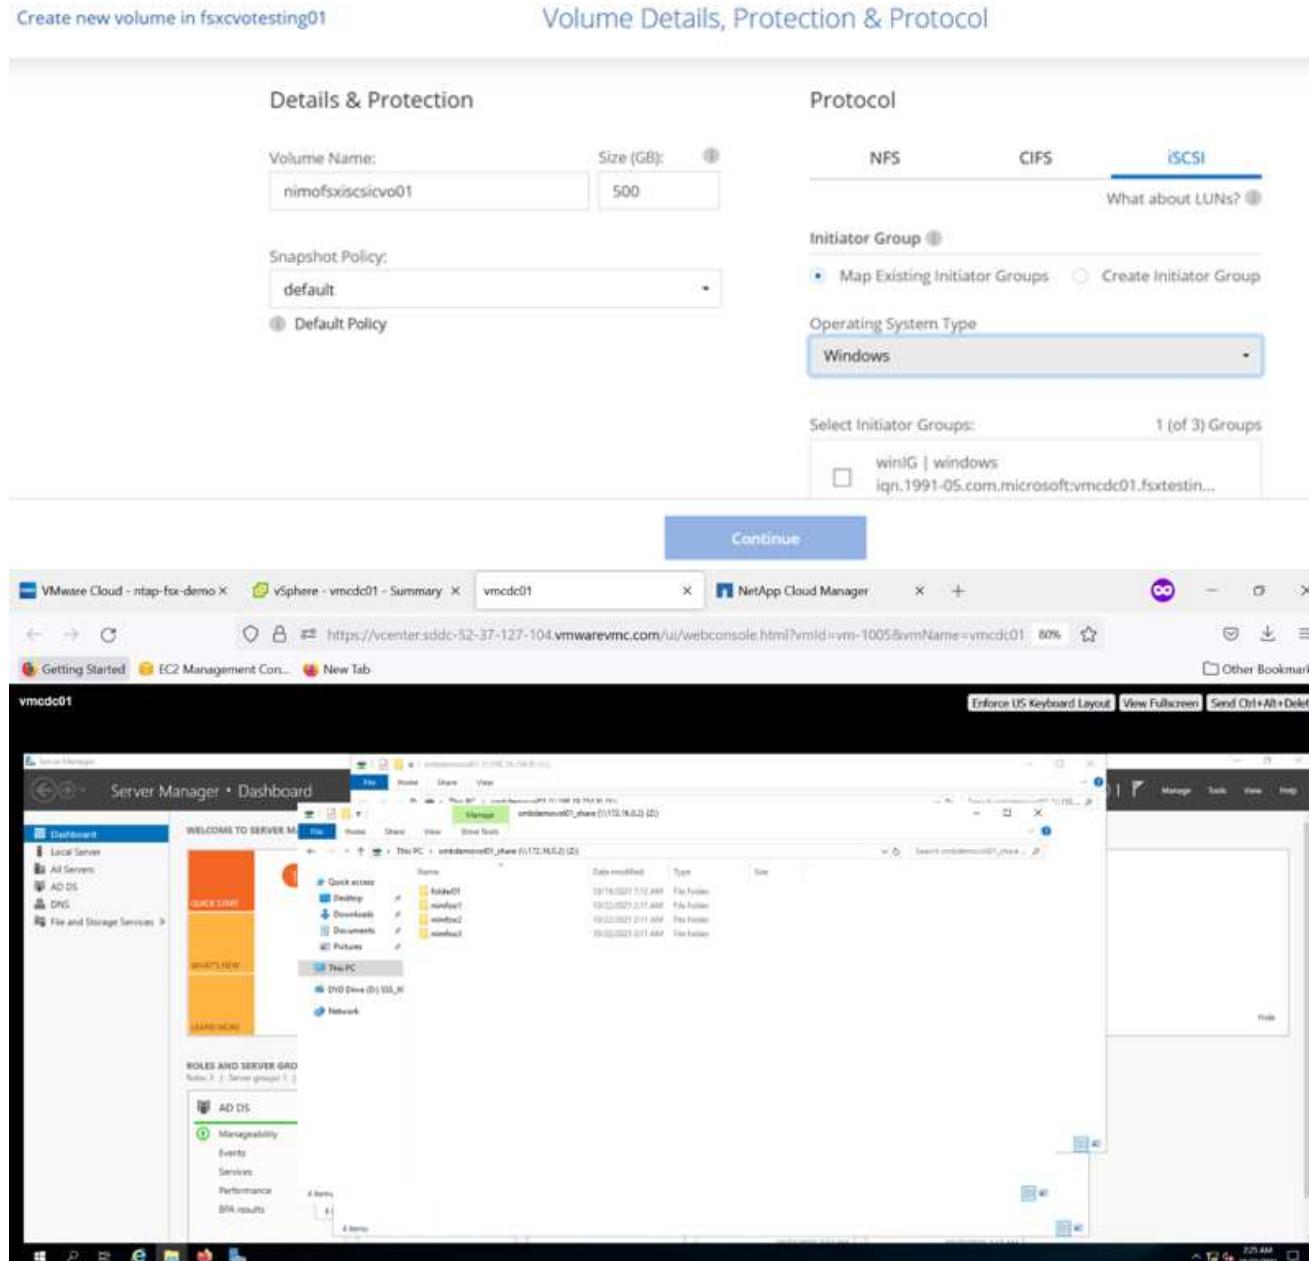
If the primary node goes offline, mount the volume by using the HA partner's IP address:



Connect the LUN to a host

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

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



The screenshot shows the 'Volume Details, Protection & Protocol' page in the NetApp Cloud Manager. The 'iSCSI' tab is selected. The 'Volume Name' is 'nimofsxiscsicvo01' and the 'Size (GB)' is '500'. The 'Snapshot Policy' is 'default'. Under 'Protocol', 'iSCSI' is selected. The 'Initiator Group' section shows 'Map Existing Initiator Groups' selected. The 'Operating System Type' is 'Windows'. The 'Select Initiator Groups' section shows '1 (of 3) Groups' with 'winIG | windows' and 'iqn.1991-05.com.microsoft:vmcd01.fsxtesting01' listed. A 'Continue' button is at the bottom. Below this, a screenshot of a Windows Server Manager Dashboard shows a mounted volume 'vmcd01' with a folder structure including 'folder01', 'monkey1', 'monkey2', and 'monkey3'.

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

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

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

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

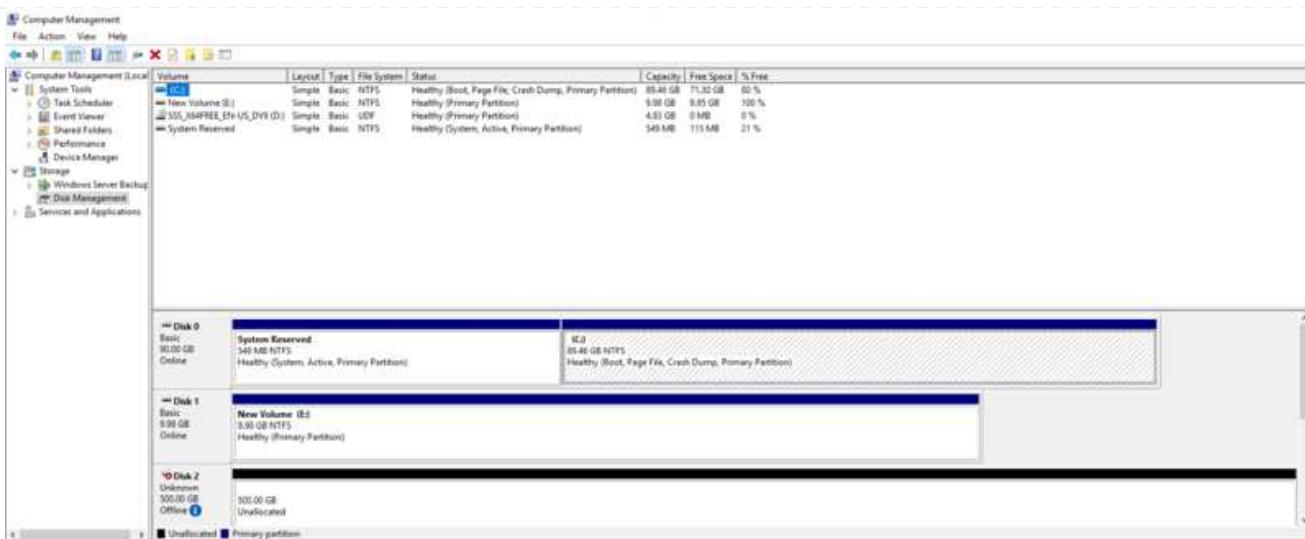


The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



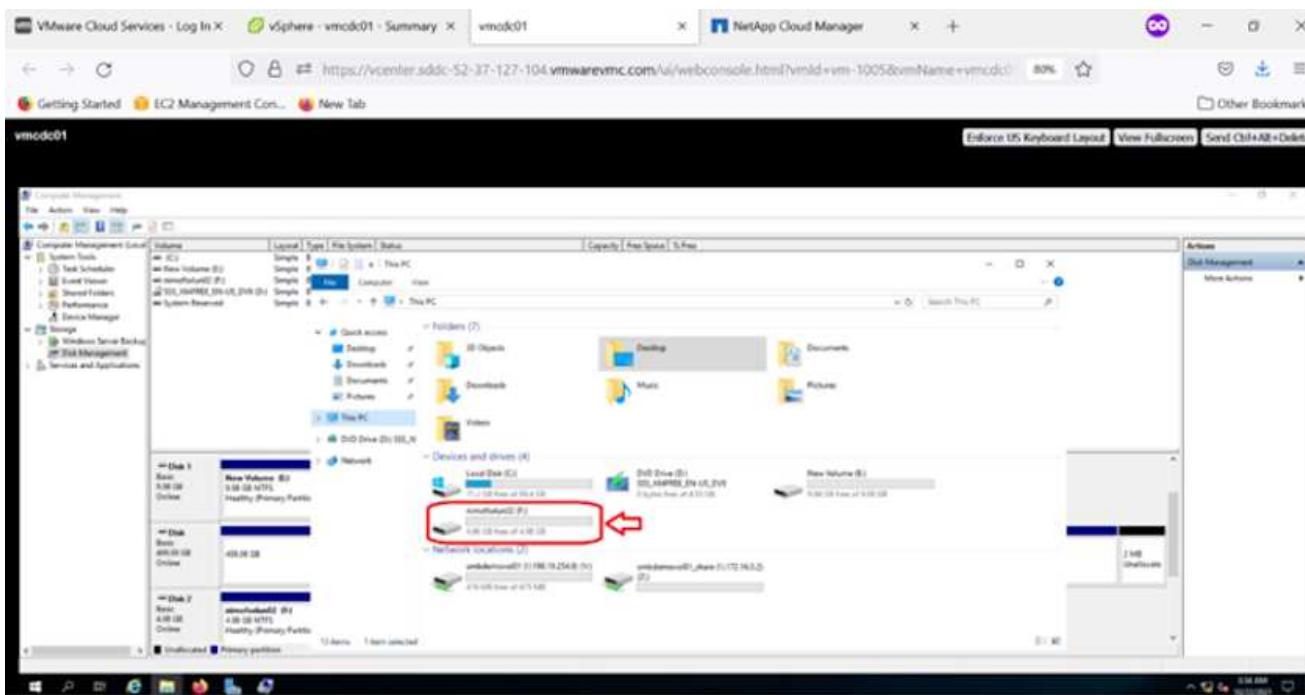
LUNs from the SVM appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

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



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

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



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

Mount Cloud Volumes ONTAP NFS volume on Linux client

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

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

```
$ sudo mkdir /fsxvcotesting01/nfsdemovol01
```

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

```
sudo mount -t nfs nfsvers=4.1,172.16.0.2:/nfsdemovol01  
/fsxvcotesting01/nfsdemovol01
```



```
root@ubuntu01:/fsx/nfsdemovol01# mount -t nfs 172.16.0.2:/nfsdemovol01 /fsxvcotesting01/nfsdemovol01
root@ubuntu01:~# df
Filesystem      1K-blocks   Used Available Use% Mounted on
tmpfs            614936    116   613820  1% /run
/dev/mapper/ubuntu--vg-ubuntu--1v 15412168 3666428 10943132 28% /
tmpfs            4071960    0  4071960  0% /dev/shm
tmpfs            5120      0   5120  0% /run/lock
tmpfs            4096      0   4096  0% /sys/fs/cgroup
/dev/sda2        999320  254995   675512 28% /boot
tmpfs            814382      4   814388  1% /run/user/1000
172.16.0.2:/nfsdemovol01 9961472 4241792 5715680 43% /fsxvcotesting01/nfsdemovol01
root@ubuntu01:~# cd /fsx/nfsdemovol01/
root@ubuntu01:~/fsx/nfsdemovol01# ls
nmonv111.txt
root@ubuntu01:~/fsx/nfsdemovol01#
```

Azure Virtualization Service: Options for using NetApp storage

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

Azure NetApp Files (ANF) as a supplemental NFS datastore

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

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

Azure NetApp Files is built on industry-leading, highly available NetApp ONTAP data management software.

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

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

- [ANF as a Supplemental NFS Datastore: Overview](#)
- [Supplemental NFS Datastore Option in Azure](#)

Azure NetApp Files (ANF) as guest connected storage

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

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

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

For more information, please visit [ANF as Guest Connected Storage](#).

Cloud Volumes ONTAP (CVO) as guest connected storage

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

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

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

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

For more information, please visit [CVO as Guest Connected Storage](#).

Overview of ANF Datastore Solutions

Every successful organization is on a path of transformation and modernization. As part of this process, companies typically use their existing VMware investments while leveraging cloud benefits and exploring how to make migration, burst, extend, and disaster recovery processes as seamless as possible. Customers migrating to the cloud must evaluate the issues of elasticity and burst, data center exit, data center

consolidation, end-of-life scenarios, mergers, acquisitions, and so on. The approach adopted by each organization can vary based on their respective business priorities. When choosing cloud-based operations, selecting a low-cost model with appropriate performance and minimal hindrance is a critical goal. Along with choosing the right platform, storage and workflow orchestration is particularly important to unleash the power of cloud deployment and elasticity.

Use Cases

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

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

Another common use case for Azure VMware Solution is disaster recovery (DR). Most organizations do not have a fool-proof DR strategy, or they might struggle to justify running a ghost datacenter just for DR. Administrators might explore zero-footprint DR options with a pilot-light cluster or an on-demand cluster. They could then scale the storage without adding additional hosts, potentially an attractive option.

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

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



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

VMware Cloud options in Azure

Azure VMware Solution

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

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

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

Conclusion

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

Takeaways

The key points of this document include:

- You can now use Azure NetApp Files as a datastore on AVS SDDC.
- Boost the application response times and deliver higher availability to provide access workload data when and where it is needed.
- Simplify the overall complexity of the vSAN storage with simple and instant resizing capabilities.
- Guaranteed performance for mission-critical workloads using dynamic reshaping capabilities.
- If Azure VMware Solution Cloud is the destination, Azure NetApp Files is the right storage solution for optimized deployment.

Where to find additional information

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

- Azure VMware Solution documentation

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

- Azure NetApp Files documentation

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

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

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

Creating a Supplemental NFS Datastore in Azure

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

Running vSphere on NFS is a widely adopted option for virtualization deployments on-premises because it offers strong performance and stability. If you have significant network-attached storage (NAS) in an on-

premises datacenter, you should consider deploying an Azure VMware solution SDDC in Azure with Azure NetApp File datastores to overcome capacity and performance challenges.

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

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

Keep in mind the following issues:

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

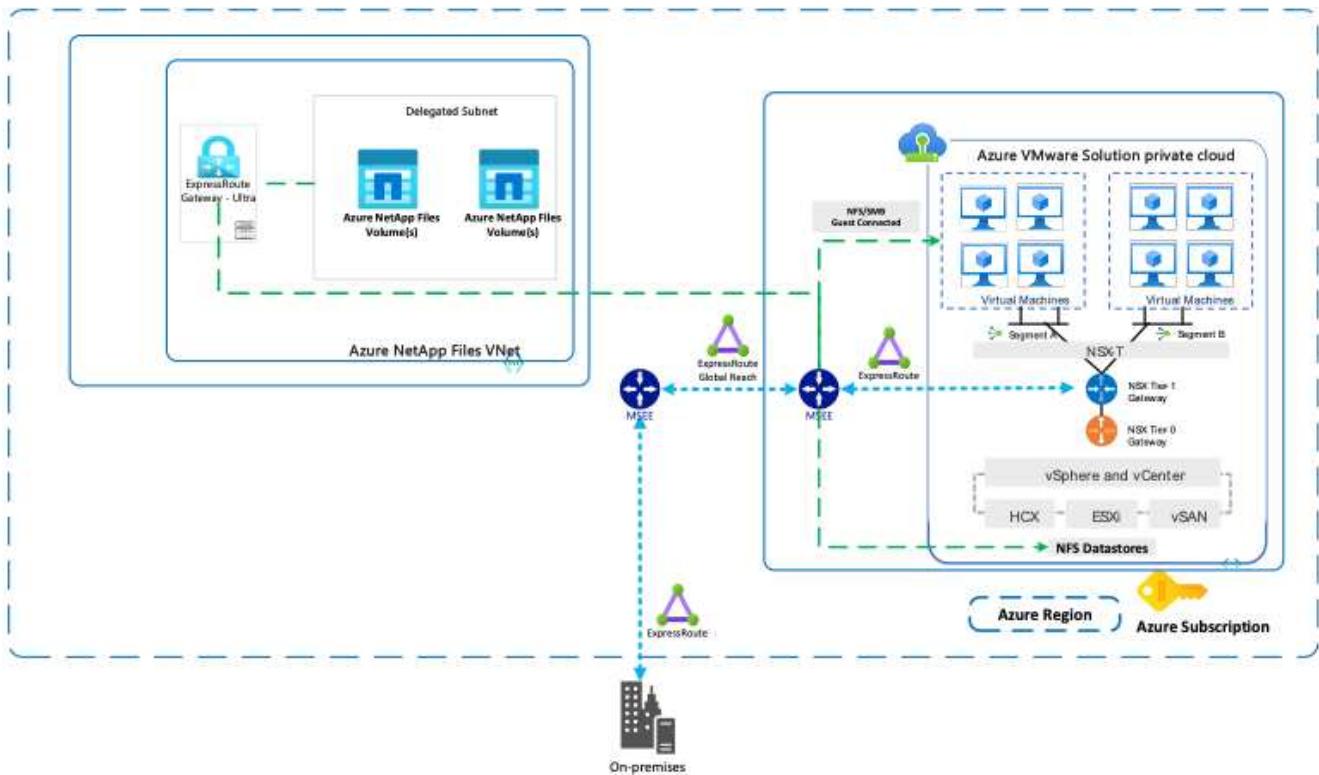
Azure NetApp Files also enables you to deploy multiple datastores, which helps to mimic an on-premises deployment model by placing virtual machines in the appropriate datastore and by assigning the required service level to meet workload performance requirements. With the unique capability of multi-protocol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



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

Detailed architecture

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



Sizing

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

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

To run RVtools, complete the following steps:

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



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

Deploying and configuring Azure VMware Solution

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

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

The setup process can be broken down into three parts:

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

Configure Azure NetApp Files with Azure VMware Solution

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



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

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

Prerequisites

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

```
az login -tenant xcvxvcxc- vxvcv- xcvx- cvxc- vxcvxvcvxcv  
az feature show --name "CloudSanExperience" --namespace "Microsoft.AVS"
```

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

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



Registration can take approximately 15 minutes to complete.

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

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

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

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

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

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

1. Verify that the vmware extension is installed.

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

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

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

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

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

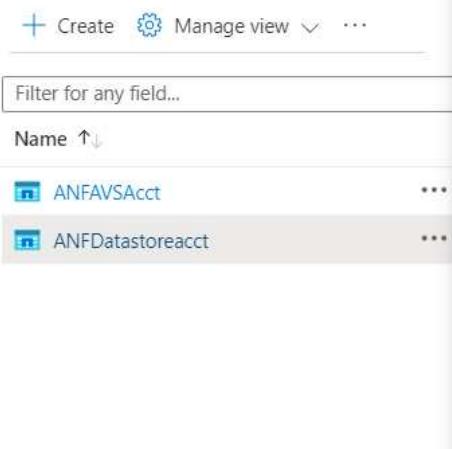
Create and mount Azure NetApp Files volumes

1. Log into the Azure Portal and access Azure NetApp Files. Verify access to the Azure NetApp Files service and register the Azure NetApp Files Resource Provider by using the `az provider register --namespace Microsoft.NetApp --wait` command. After registration, create a NetApp account. Refer to this [link](#) for detailed steps.

Home > Azure NetApp Files >

Azure NetApp Files

Hybrid Cloud TME



Filter for any field...

Name ↑↓

ANFAVSAcct

ANFDatstoreacct

New NetApp account

Name *

Enter the name

Subscription

Hybrid Cloud TME Onprem

Resource group *

ANFAVSAcct2

Create new

Location *

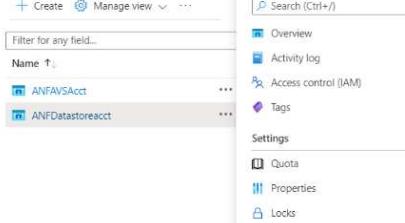
West Europe

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

Home > Azure NetApp Files > ANFDatstoreacct

Azure NetApp Files

Hybrid Cloud TME



Filter for any field...

Name ↑↓

ANFAVSAcct

ANFDatstoreacct

ANFDatstoreacct | Capacity pools

+ Add pool

Refresh

Search pools

Name Capacity Service level QoS

ANFRec0DS 4 TiB Premium Auto

ANFRec0DSU 4 TiB Ultra Auto

New capacity pool

Name *

Enter the name

Service level *

Premium

Size (TiB) *

4

4 TiB

QoS type

Manual

Auto

Points to Remember

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

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

ANFDatastoreacct | Volumes X

NetApp account

Search (Ctrl+ /) Add volume + Add data replication + Refresh

Overview Activity log + Access control (IAM) + Tags + Settings

Search volumes

Name	Quota	Throughput	Protocol type	Mount path	Service level	Network features	Capacity
anfrecods001	4 TiB	262.144 MiB/s	NFSv3	172.30.153.132:/ANFRe	Premium	Standard	anfrecods
anfrecodsu002	4 TiB	524.288 MiB/s	NFSv3	172.30.153.132:/anfrec	Ultra	Standard	anfrecodsu

To learn about Azure NetApp Files volume performance by size or quota, see [Performance considerations for Azure NetApp Files](#).

Add Azure NetApp files datastore to private cloud



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

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

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

```
C:\Users\niyaz>az vmware datastore netapp-volume create --name
ANFRecoDSU002 --resource-group anfavsva12 --cluster Cluster-1 --private
-cloud ANFDataClus --volume-id /subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.NetApp/netAp
pAccounts/anfdatastoreacct/capacityPools/anfrecodsu/volumes/anfrecodsU0
02
{
  "diskPoolVolume": null,
  "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.AVS/privateC
louds/ANFDataClus/clusters/Cluster-1/datastores/ANFRecoDSU002",
  "name": "ANFRecoDSU002",
  "netAppVolume": {
    "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.NetApp/netAp
pAccounts/anfdatastoreacct/capacityPools/anfrecodsu/volumes/anfrecodsU0
02",
    "resourceGroup": "anfavsva12"
  },
  "provisioningState": "Succeeded",
  "resourceGroup": "anfavsva12",
  "type": "Microsoft.AVS/privateClouds/clusters/datastores"
}

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

```
C:\Users\niyaz>az vmware datastore list --resource-group anfavsva12 --cluster Cluster-1 --private-cloud
ANFDataClus
[
{
  "diskPoolVolume": null,
  "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.AVS/privateClouds/ANFDataClus/clusters/
Cluster-1/datastores/ANFRecoDS001",
```

```

"name": "ANFRecoDS001",
"netAppVolume": {
  "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.NetApp/netAppAccounts/anfdastoreacct/
capacityPools/anfrecods/volumes/ANFRecoDS001",
"resourceGroup": "anfavsva12"
},
"provisioningState": "Succeeded",
"resourceGroup": "anfavsva12",
"type": "Microsoft.AVS/privateClouds/clusters/datastores"
},
{
"diskPoolVolume": null,
"id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.AVS/privateClouds/ANFDataClus/clusters/
Cluster-1/datastores/ANFRecoDSU002",
"name": "ANFRecoDSU002",
"netAppVolume": {
  "id": "/subscriptions/0efa2dfb-917c-4497-b56a-
b3f4eadb8111/resourceGroups/anfavsva12/providers/Microsoft.NetApp/netAppAccounts/anfdastoreacct/
capacityPools/anfrecodsu/volumes/anfrecodsU002",
"resourceGroup": "anfavsva12"
},
"provisioningState": "Succeeded",
"resourceGroup": "anfavsva12",
"type": "Microsoft.AVS/privateClouds/clusters/datastores"
}
]

```

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

The screenshot shows the vSphere Client interface for managing datastores. The left sidebar shows a tree structure with 'vc.3960e358ae6a4df38d9...' and 'anf-z1-stds01' selected. The main pane is titled 'anf-z1-stds01' and shows the 'Summary' tab. The 'Details' section displays the following information:

Location	ds://vmfs/volumes/d2abbb2d-2c68af49/
Type	NFS 3
Hosts	3
Virtual machines	2
VM templates	0
Server	172.30.136.70
Folder	anf-z1-stds01

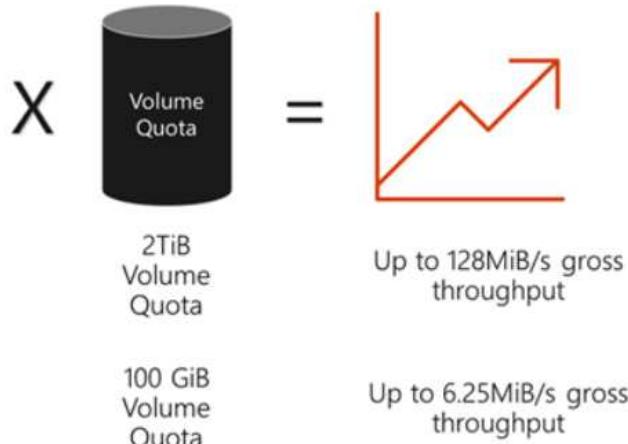
The 'Related Objects' section shows 'Custom Attributes' with an empty table.

Sizing and performance optimization

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

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

Service Level	Throughput
Ultra	128MiB/s per 1TiB quota
Premium	64MiB/s per 1TiB quota
Standard	16MiB/s per 1TiB quota



For more information, see [Service levels for Azure NetApp Files](#).

Refer to this [link from Microsoft](#) for detailed performance benchmarks that can be used during a sizing exercise.

Points to Remember

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

Performance considerations

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

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

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

Refer to this [link from Microsoft](#) for detailed performance benchmarks that can be used during a sizing exercise.

Points to Remember

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

Increasing the size of the datastore

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

Points to remember

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

Workloads

Migration

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

Data Protection

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

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



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

BlueXP backup and recovery for Virtual Machines

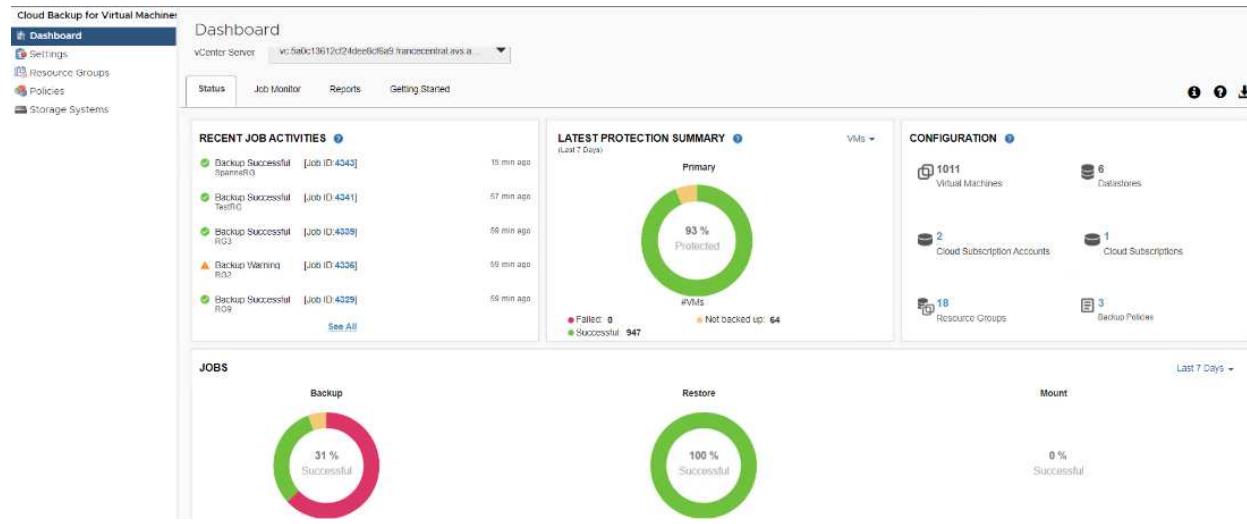
BlueXP backup and recovery for Virtual Machines provides a vSphere web client GUI on vCenter to protect Azure VMware Solution virtual machines and Azure NetApp files datastores via backup policies. These policies can define schedule, retention, and other capabilities. The BlueXP backup and recovery for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

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

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



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



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

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

[DR solution with ANF, JetStream and AVS.](#)

NetApp Guest Connected Storage Options for Azure

Azure supports guest connected NetApp storage with the native Azure NetApp Files (ANF) service or with Cloud Volumes ONTAP (CVO).

Azure NetApp Files (ANF)

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

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

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

Azure NetApp Files (ANF) as guest connected storage

Configure Azure NetApp Files with Azure VMware Solution (AVS)

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

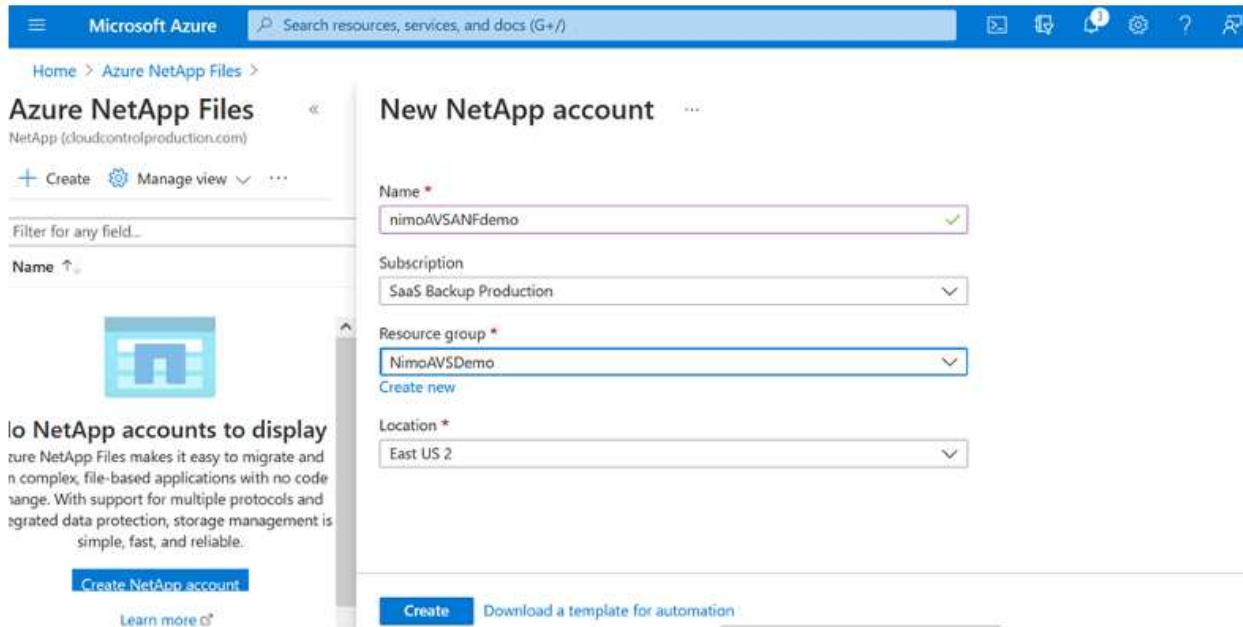
Azure NetApp Files and Azure VMware Solution must be in the same Azure region.

Create and mount Azure NetApp Files volumes

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

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

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



The screenshot shows the 'New NetApp account' creation page in the Azure Portal. The 'Name' field is filled with 'nimoAVSANFDemo'. The 'Subscription' is set to 'SaaS Backup Production'. The 'Resource group' is set to 'NimoAVSDemo'. The 'Location' is set to 'East US 2'. At the bottom, there are 'Create' and 'Download a template for automation' buttons.

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

For more information, see [Set up a capacity pool](#).

The screenshot shows the Azure NetApp Files portal with the 'Capacity pools' blade open. A 'New capacity pool' dialog is displayed on the right, containing the following fields:

- Name: nimappool
- Service level: Standard
- Size (TiB): 4
- QoS type: Auto

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

The screenshot shows the Azure portal with the 'Subnets' blade open for the 'nimoavspiv-vnet' virtual network. A 'Add subnet' dialog is displayed on the right, containing the following fields:

- Name: anf.del
- Subnet address range: 172.24.3.0/28 (172.24.3.0 - 172.24.3.15 (11 + 5 Azure reserved addresses))
- Add IPv6 address space:
- NAT gateway: None
- Network security group: None
- Route table: None

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

The screenshot shows the Azure NetApp Files interface. On the left, a sidebar lists 'Azure NetApp Files' and 'nimoAVSANFdemo'. The main area is titled 'nimoAVSANFdemo | Active Directory connections'. It includes a search bar, a 'Join' button, and a table for joining Active Directories. The table has columns for 'DNS', 'AD DNS Domain', and 'SMB Server'. A note says 'No currently joined Active Directories.' On the right, a 'Join Active Directory' form is open, showing fields for 'Primary DNS' (172.24.1.5), 'AD DNS Domain Name' (nimodemo.com), 'SMB Server (Computer Account) Prefix' (nimsmrb), and an 'Organizational Unit Path'. A 'Join' button is at the bottom.

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

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

The screenshot shows the 'Create a volume' blade. The 'Basics' tab is selected. It displays the following details:

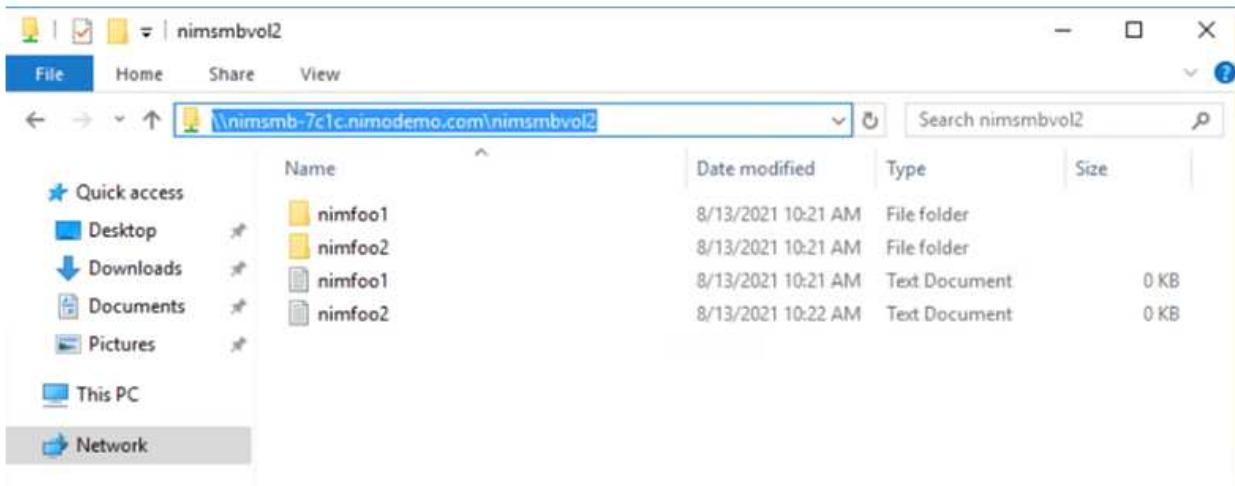
- Volume details:**
 - Volume name: nimvoltest1
 - Capacity pool: nimcappool
 - Available quota (GiB): 4096
 - Quota (GiB): 100

At the bottom, there are buttons for 'Review + create', '< Previous', and 'Next : Protocol >'. The 'Review + create' button is highlighted in blue.

To learn more about Azure NetApp Files volume performance by size or quota, see [Performance considerations for Azure NetApp Files](#).

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

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



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

```
nimoadmin@nimoadmin-virtual-machine:~$ sudo mount -t nfs -o rw,hard,tcp 172.24.3.4:/nimodemonfs1 /home/nimoadmin/nimo11
nimoadmin@nimoadmin-virtual-machine:~$ df
Filesystem      1K-blocks   Used Available Use% Mounted on
udev             8168112     0  8168112  0% /dev
tmpfs            1639548  1488  1638060  1% /run
/dev/sdas        50824704 7982752 40310496 17% /
tmpfs            8197728     0  8197728  0% /dev/shm
tmpfs             5120      0   5120  0% /run/lock
tmpfs            8197728     0  8197728  0% /sys/fs/cgroup
/dev/loop0        56832    56832      0 100% /snap/core18/2128
/dev/loop2        66688    66688      0 100% /snap/gtk-common-themes/1515
tmpfs            224256  224256      0 100% /snap/gnome-3-34-180
4/72
/dev/loop3        52224    52224      0 100% /snap/snap-store/547
/dev/loop4        33152    33152      0 100% /snap/snapd/12704
/dev/sda1        523248      4  523244  1% /boot/efi
tmpfs            1639544    52  1639492  1% /run/user/1000
/dev/sr0          54738    54738      0 100% /media/nimoadmin/VMware Tools
172.24.3.4:/nimodemonfs1 104857600      0 104857600  0% /home/nimoadmin/nimo11
nimoadmin@nimoadmin-virtual-machine:~$
```

For more information, see [Dynamically change the service level of a volume](#).

Cloud Volumes ONTAP (CVO)

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

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

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

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

Cloud Volumes ONTAP (CVO) as guest connected storage

Deploy new Cloud Volumes ONTAP in Azure

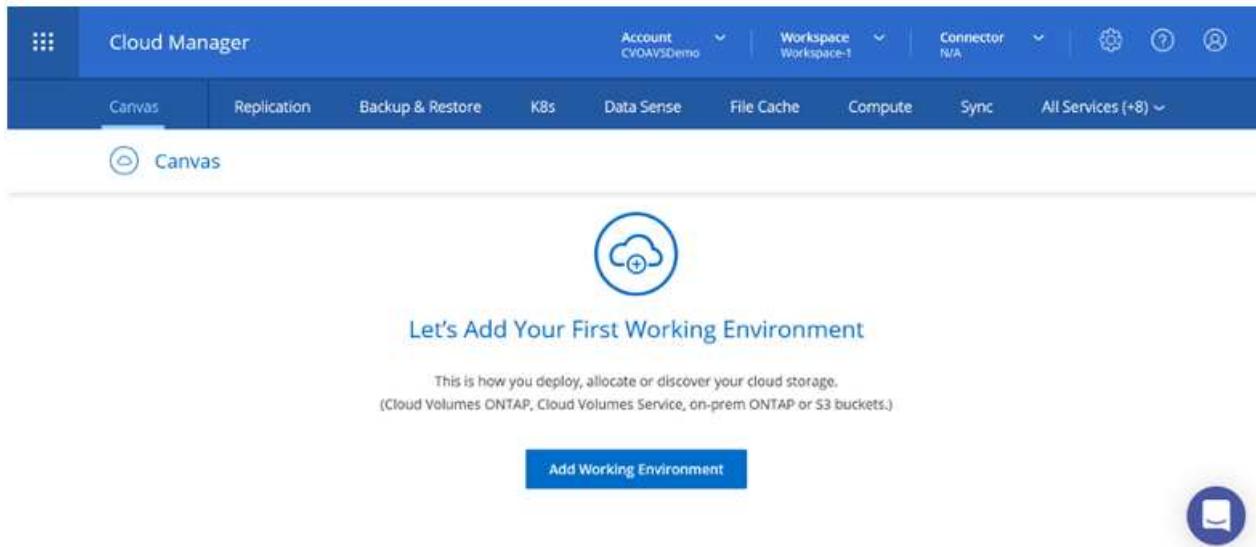
Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the Azure VMware Solution SDDC environment. The volumes can also be mounted on the Linux client and on Windows client because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be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

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

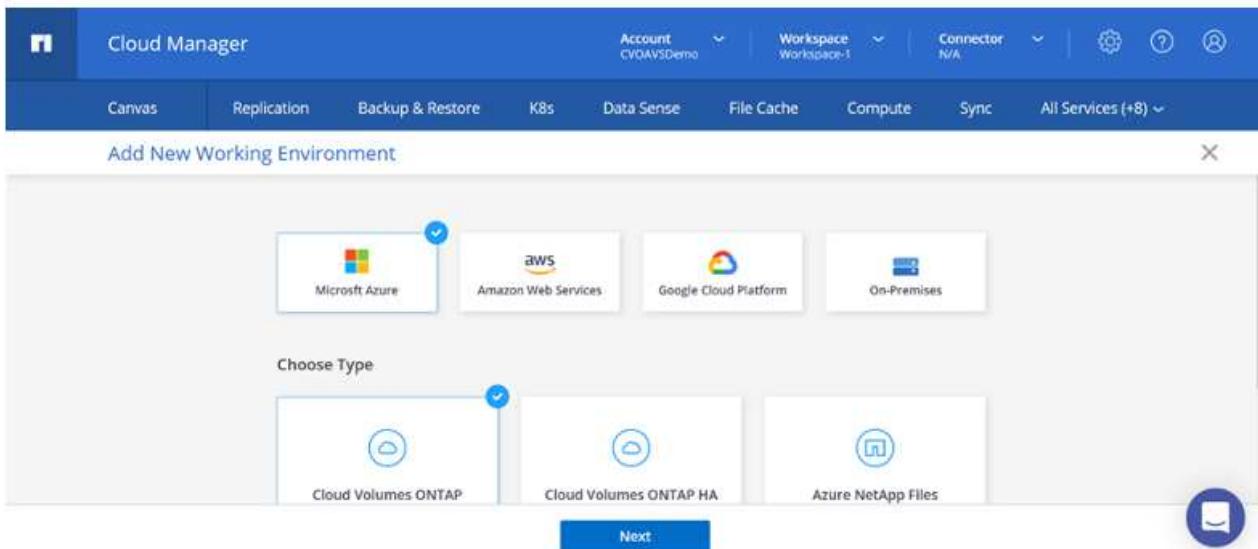


Use [Cloud Volumes ONTAP sizer](#) to accurately size the Cloud Volumes ONTAP instances. Also monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

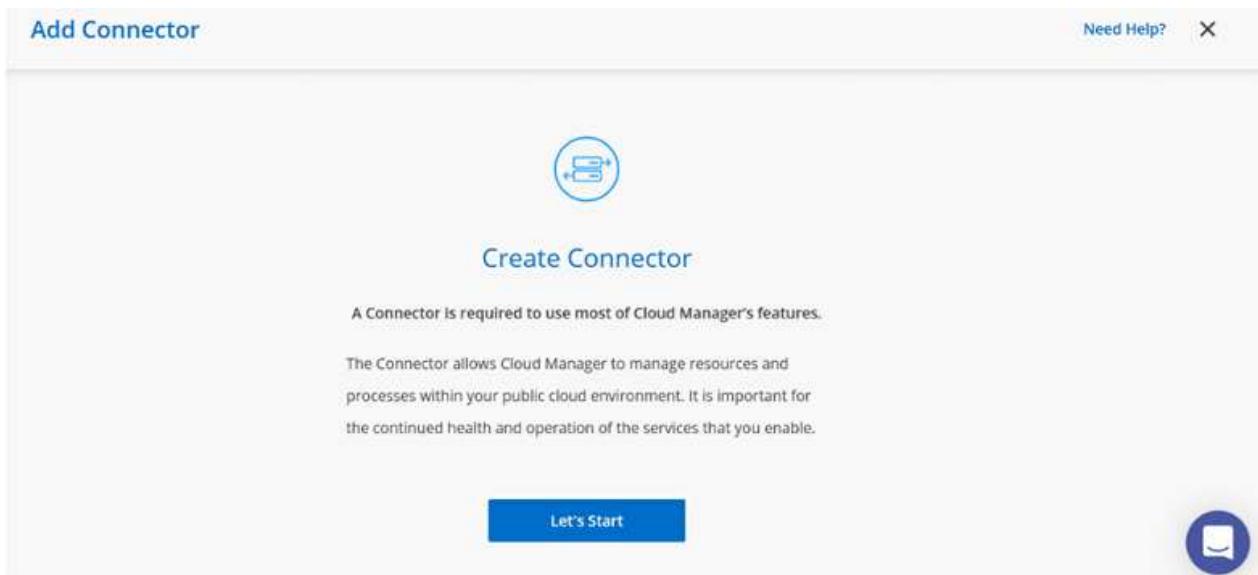
1. Log in to NetApp Cloud Central—the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.



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



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



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

[Create a New Working Environment](#)

Details and Credentials

Managed Service Ide...	SaaS Backup Prod...	CMCVOSub	Edit Credentials
Credential Name	Azure Subscription	Marketplace Subscription	

Details

Working Environment Name (Cluster Name)
nimavscVO

Credentials

User Name
admin

Password

[Continue](#)

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

[Create a New Working Environment](#)

Details and Credentials

Details	Credentials
Working Environment Name (Cluster Name) nimavscVO	User Name admin
+ Add Resource Group Tags Optional Field	Password *****
	Confirm Password *****

[Continue](#)

6. Select the add-on services for Cloud Volumes ONTAP deployment, including BlueXP Classification, BlueXP backup and recovery, and Cloud Insights. Select the services and then click Continue.

[Create a New Working Environment](#)

Services

 Data Sense & Compliance	<input checked="" type="checkbox"/>	▼
 Backup to Cloud	<input checked="" type="checkbox"/>	▼
 Monitoring	<input checked="" type="checkbox"/>	▼

[Continue](#)

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

Azure Region

Availability Zone (Optional)

VNet

Subnet

Resource Group
 Create a new group Use an existing group.

Resource Group Name

Security Group
 Generated security group Use existing security group

I have verified network connectivity between the Cloud Manager server and the selected VNet.

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

Cloud Volumes ONTAP Charging Methods

[Learn more about our charging methods](#)

Pay-As-You-Go by the hour

Bring your own license

NetApp Support Site Account (Optional)

[Learn more about NetApp Support Site \(NSS\) accounts](#)

To register this Cloud Volumes ONTAP to support, you should add NetApp Support Site Account.

Don't have a NetApp Support Site account? Select go to finish deploying this system. After it's created, use the [Support Registration](#) option to create an NSS account.

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

Select a preconfigured Cloud Volumes ONTAP system that best matches your needs, or create your own configuration. Preconfigured settings can be modified at a later time.

POC and small workloads
Up to 500GB of storage

Database and application data production workloads

Cost effective DR
Up to 500GB of storage

Highest performance production workloads

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



nimavscVO

Azure | East US 2

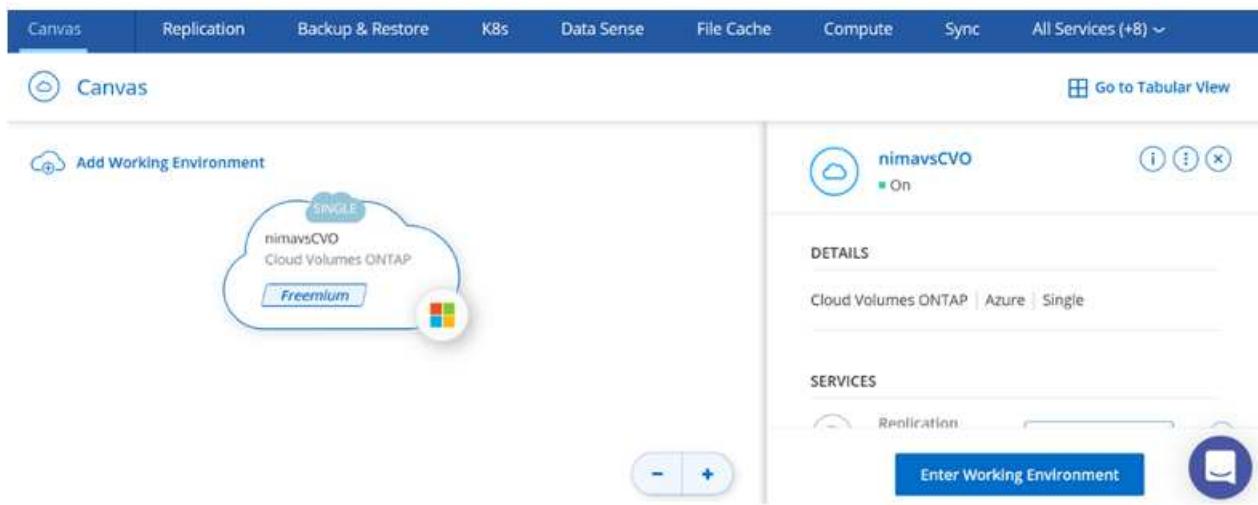
I understand that in order to activate support, I must first register Cloud Volumes ONTAP with NetApp. [More information >](#)

I understand that Cloud Manager will allocate the appropriate Azure resources to comply with my above requirements. [More information >](#)

Overview Networking Storage

Go

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



Canvas Replication Backup & Restore K8s Data Sense File Cache Compute Sync All Services (+8) ▾

Cloud Canvas Go to Tabular View

Add Working Environment

nimavscVO
Cloud Volumes ONTAP
Freemium

DETAILS

Cloud Volumes ONTAP | Azure | Single

SERVICES

Replication

Enter Working Environment

Additional configurations for SMB volumes

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

The screenshot shows the 'Create a CIFS server' configuration page in the nimavsCVO interface. The 'Volumes' tab is selected. The 'DNS Primary IP Address' field contains '172.24.1.5'. The 'Active Directory Domain to join' field contains 'nimodemo.com'. The 'Credentials authorized to join the domain' field contains 'nimoadmin'. There is an 'Advanced' button in the top right corner.

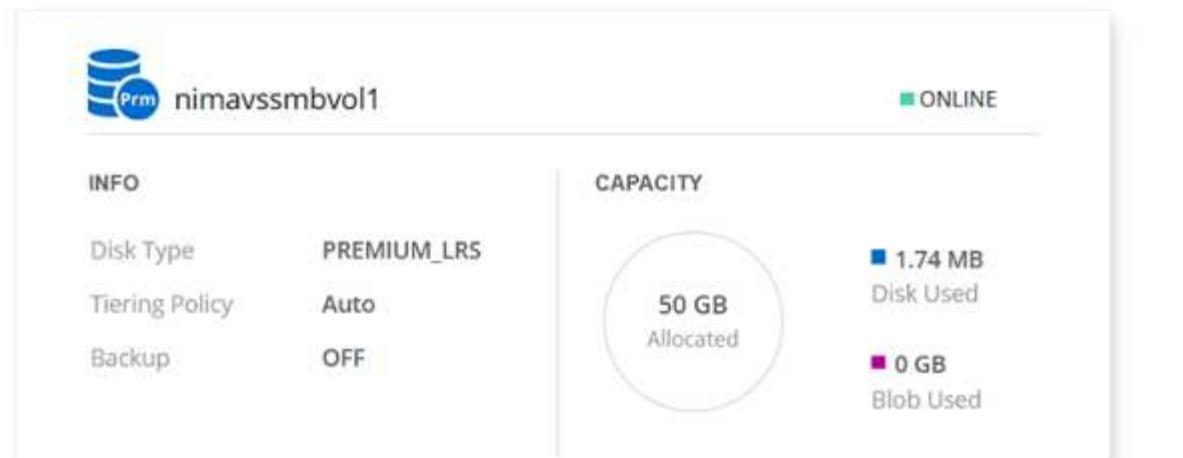
2. Creating the SMB volume is an easy process. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected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

The screenshot shows the 'Volume Details, Protection & Protocol' configuration page in the nimavsCVO interface. The 'Protocol' tab is selected, showing 'CIFS' is chosen. The 'Share name' is 'nimavssmbvol1_share' and 'Permissions' is 'Full Control'. The 'Details & Protection' tab shows 'Volume Name' as 'nimavssmbvol1' and 'Size (GB)' as '50'. The 'Snapshot Policy' dropdown shows 'default'.

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

Volumes

1 Volume | 50 GB Allocated | 1.74 MB Total Used (1.74 MB in Disk, 0 KB in Blob)



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

 Mount Volume nimavssmbvol1

Go to your machine and enter this command

\\\172.24.2.8\vimavssmbvol1_share

 Copy

Connect the LUN to a host

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

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

Details & Protection

Protocol

Volume Name: nimavsscsi1

Size (GB): 500

NFS CIFS iSCSI

What about LUNs?

Snapshot Policy: default

Initiator Group

Map Existing Initiator Groups Create Initiator Group

Initiator Group: avsvmlG

Continue

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

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

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

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



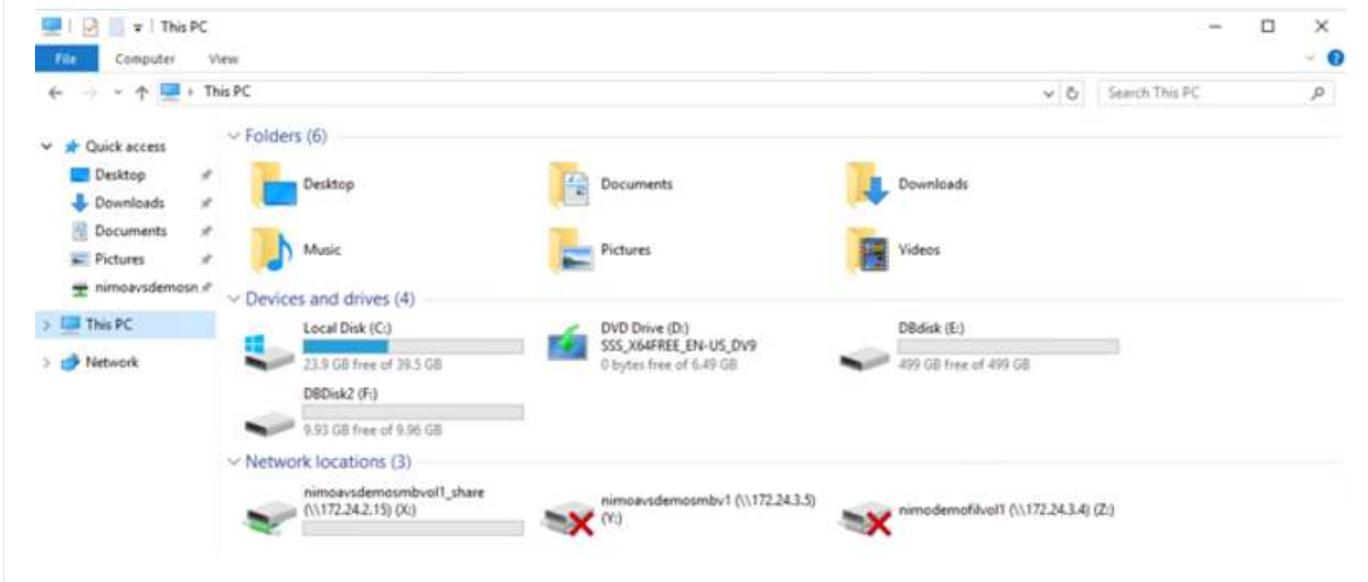
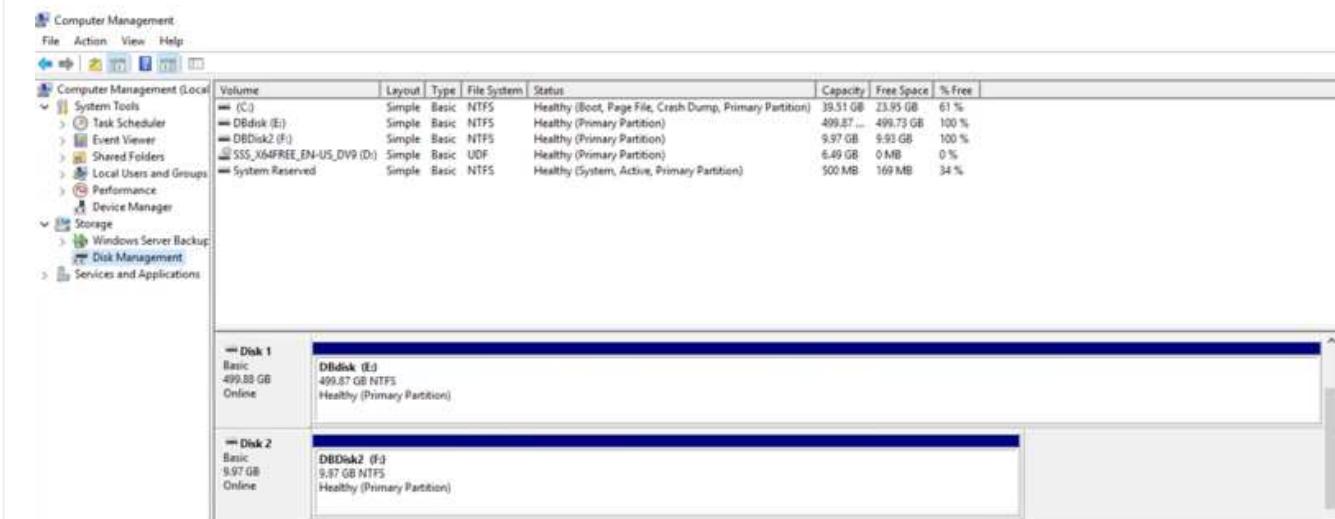
LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

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



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

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



Google Cloud VMware Engine: Options for using NetApp storage

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

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

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

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

NetApp CVO as Guest Connected Storage

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

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

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

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

For more information, please visit [NetApp CVO as Guest Connected Storage](#)

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

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

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

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

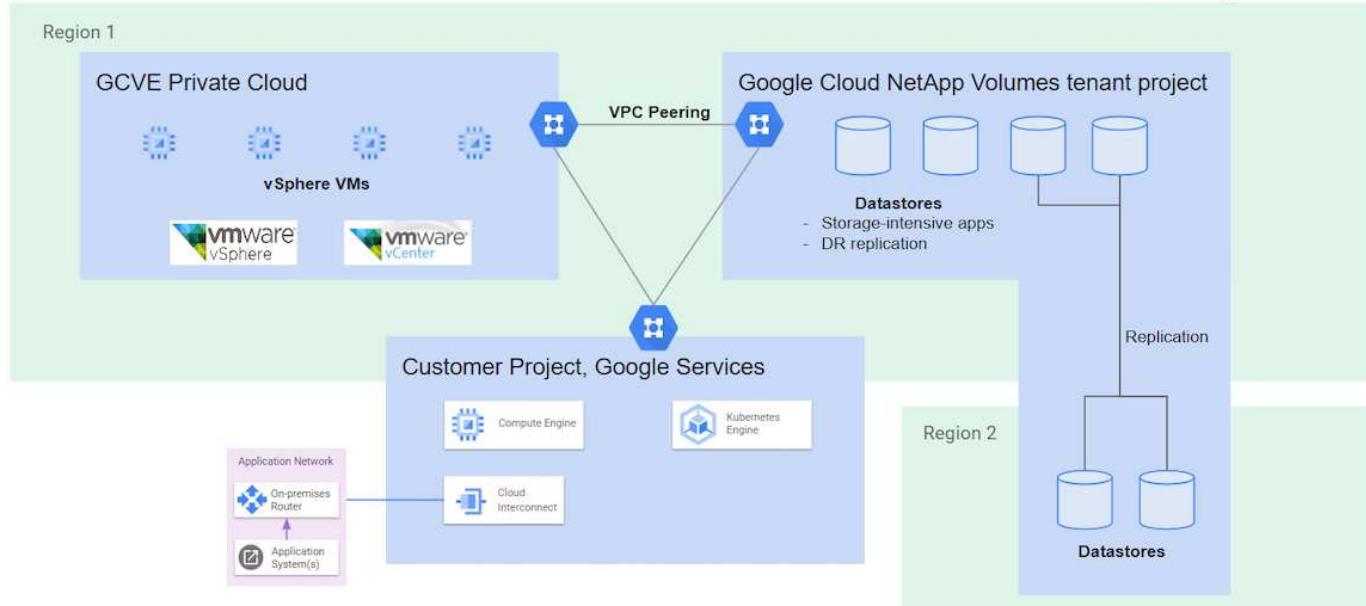
Google Cloud VMware Engine Supplemental NFS Datastore with Google Cloud NetApp Volumes

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

Overview

Customers that require additional storage capacity on their Google Cloud VMware Engine (GCVE) environment can utilize Netapp Cloud Volume Service to mount as supplemental NFS datastore.

Storing data on Google Cloud NetApp Volumes allows customers to replicate between regions to protect from disaster.



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

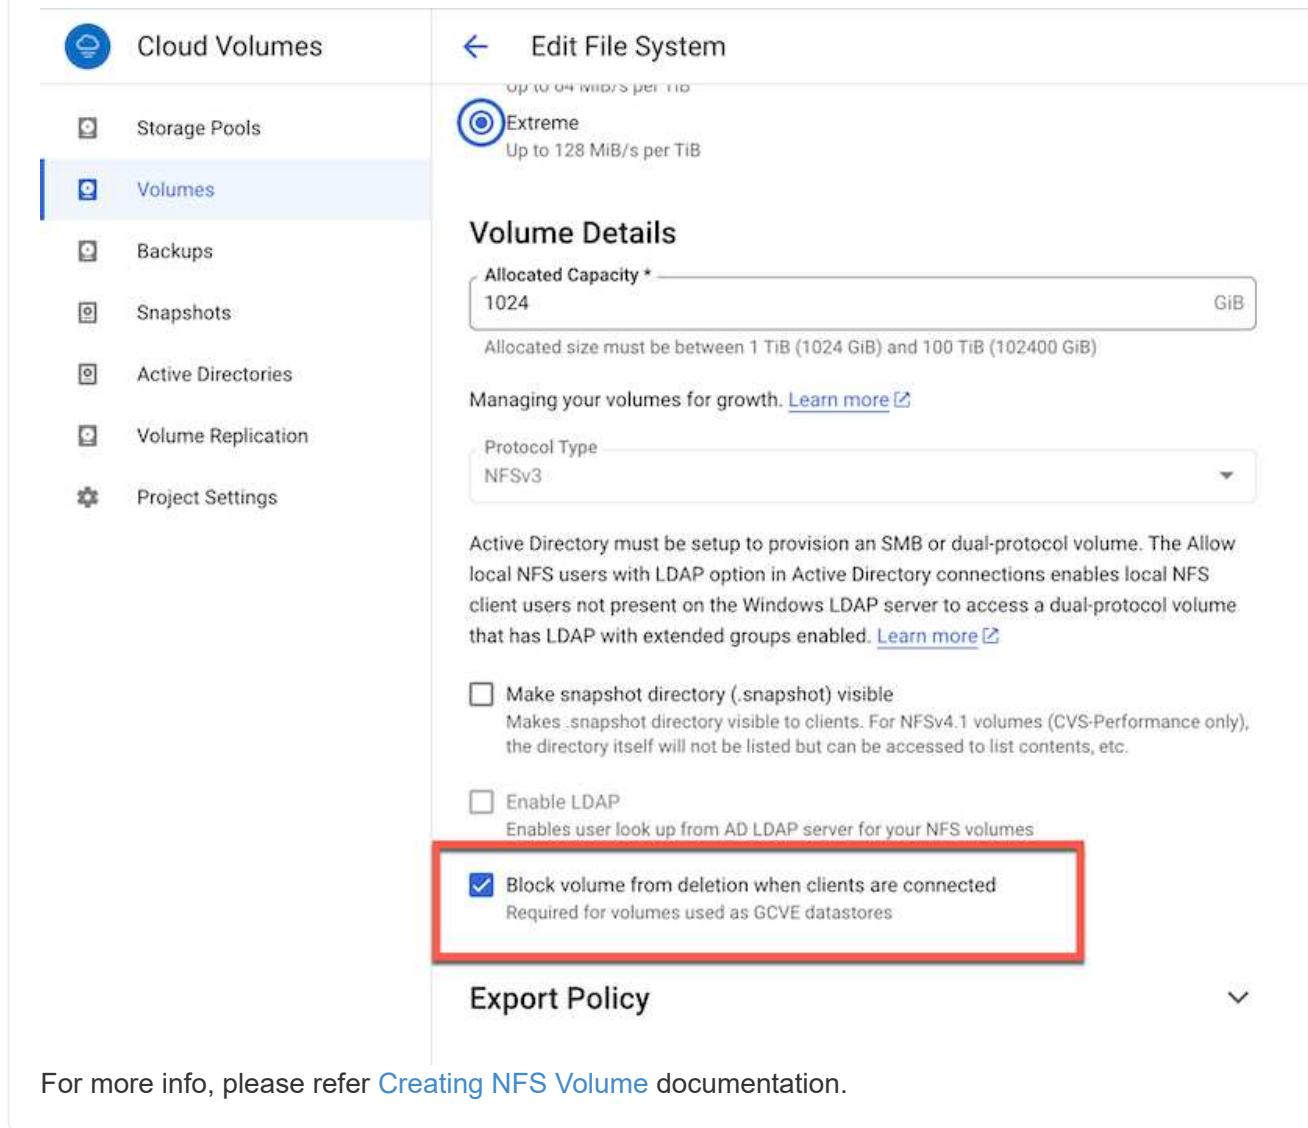
Provision NetApp Volumes-Performance Volume

The Google Cloud NetApp Volumes volume can be either provisioned by

- [Using Google Cloud Console](#)
- [Using NetApp BlueXP portal or API](#)

Mark that NetApp Volumes volume as non-deletable

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



The screenshot shows the 'Edit File System' interface for a volume. On the left, a sidebar lists 'Cloud Volumes', 'Storage Pools', 'Volumes' (which is selected and highlighted in blue), 'Backups', 'Snapshots', 'Active Directories', 'Volume Replication', and 'Project Settings'. The main area is titled 'Edit File System' with a back arrow. It shows 'Volume Details' with an 'Allocated Capacity' of 1024 GiB (with a note that it must be between 1 TiB and 100 TiB). It also shows 'Protocol Type' as NFSv3. Below this, a note states that Active Directory must be setup to provision an SMB or dual-protocol volume. The 'Export Policy' section contains three checkboxes: 'Make snapshot directory (.snapshot) visible' (unchecked), 'Enable LDAP' (unchecked), and 'Block volume from deletion when clients are connected' (checked). The last checkbox is highlighted with a red box. At the bottom, a note says 'For more info, please refer [Creating NFS Volume](#) documentation.'

For more info, please refer [Creating NFS Volume](#) documentation.

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

To mount NFS Datastore, there should be a private connection exists between GCVE and NetApp Volumes project.

For more info, please refer [How to setup Private Service Access](#)

Mount NFS datastore

For instructions on how to mount NFS datastore on GCVE, please refer [How to create NFS datastore with NetApp Volumes](#)



As vSphere hosts are managed by Google, you don't have access to install NFS vSphere API for Array Integration (VAAI) vSphere Installation Bundle (VIB).

If you need support for Virtual Volumes (vVol), please let us know.

If you like to use Jumbo Frames, please refer [Maximum supported MTU sizes on GCP](#)

Savings with Google Cloud NetApp Volumes

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

Reference Links

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

NetApp Storage Options for GCP

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

Cloud Volumes ONTAP (CVO)

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

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

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

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

Cloud Volumes ONTAP (CVO) as guest connected storage

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

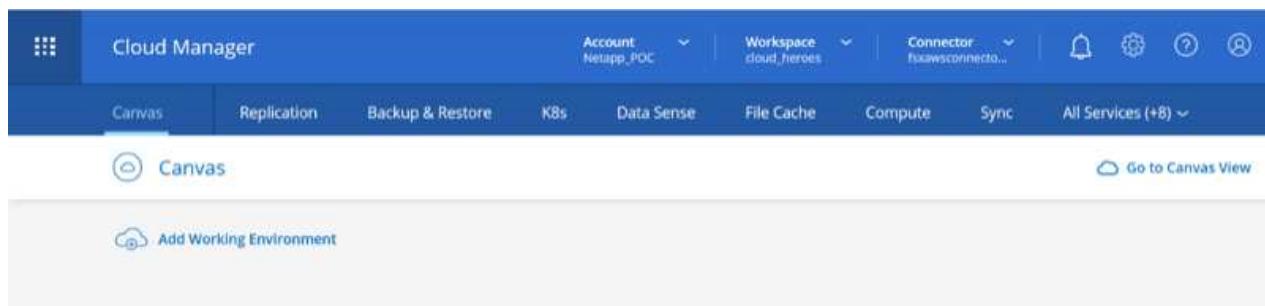
Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the GCVE private cloud environment. The volumes can also be mounted on the Linux client and on Windows client and LUNs can be accessed on Linux or Windows clients as block devices when mounted over iSCSI because Cloud Volumes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

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

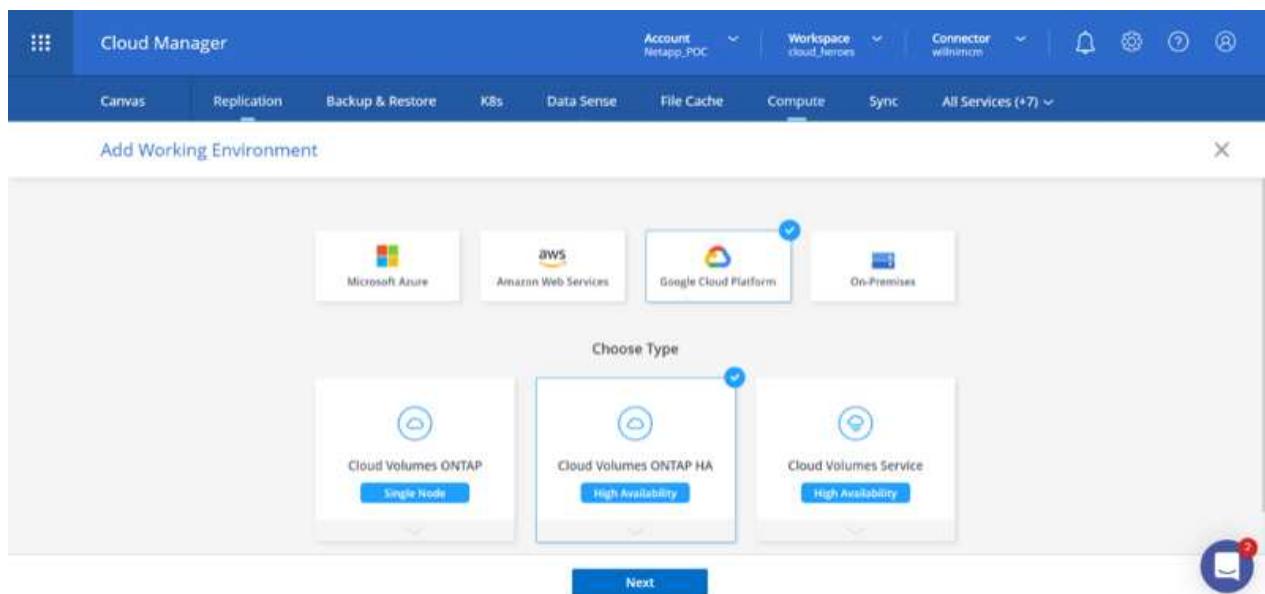


Use [Cloud Volumes ONTAP sizer](#) to accurately size the Cloud Volumes ONTAP instances. Also monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

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



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



3. Provide the details of the environment to be created including the environment name and admin

credentials. After you are done, click Continue.

Create a New Working Environment Details and Credentials

↑ Previous Step CV-Performance-Testing HCLMainBillingAccountSubs...
Google Cloud Project Marketplace Subscription Edit Project

Details Credentials

Working Environment Name (Cluster Name)
evogcveva

Service Account 

Notice: A Google Cloud service account is required to use two features: backing up data using Backup & Compliance or Backup to Cloud.

User Name
admin

Password

Confirm Password

Continue

4. Select or deselect the add-on services for Cloud Volumes ONTAP deployment, including Data Sense & Compliance or Backup to Cloud. Then, click Continue.

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

Create a New Working Environment Services

↑ Previous Step

 Data Sense & Compliance  

 Backup to Cloud  

WARNING: By turning off Backup to Cloud, future data recovery will not be possible in case of data corruption or loss

Continue

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

[↑ Previous Step](#) Location

Connectivity

GCP Region

europe-west3

VPC

cloud-volumes-vpc

GCP Zone

europe-west3-c

Subnet

10.0.6.0/24

I have verified connectivity between the target VPC and Google Cloud storage.

Firewall Policy

Generated firewall policy Use existing firewall policy

[Continue](#)

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

[↑ Previous Step](#)

Cloud Volumes ONTAP Charging Methods

[Learn more about our charging methods](#)

Pay-As-You-Go by the hour



Bring your own license



Freemium (Up to 500GB)

NetApp Support Site Account

[Learn more about NetApp Support Site \(NSS\) accounts](#)

NetApp Support Site Account

mchad

To add a new NetApp Support Site account, go to the Support - NSS Management tab.

[Continue](#)

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

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

Select a preconfigured Cloud Volumes ONTAP system that best matches your needs, or create your own configuration. Preconfigured settings can be modified at a later time.

[Change Configuration](#)



POC and small workloads
Up to 500GB of storage



Database and application data production workloads



Cost effective DR
Up to 500GB of storage



Highest performance production workloads

[Continue](#)



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

[↑ Previous Step](#)

[cvogcveval](#) | europe-west3

[Show API request](#)

This Cloud Volumes ONTAP instance will be registered with NetApp support under the NSS Account **mchad**.

I understand that Cloud Manager will allocate the appropriate GCP resources to comply with my above requirements. [More information >](#)

[Overview](#)

[Networking](#)

[Storage](#)

Storage System:

Cloud Volumes ONTAP

Cloud Volumes ONTAP runs on:

n2-standard-4

License Type:

Cloud Volumes ONTAP Freemium

Encryption:

Google Cloud Managed

Capacity Limit:

500GB

Write Speed:

Normal

[Go](#)

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

Cloud Manager

Account: NetApp_POC | Workspace: cloud_heroes | Connector: willmccorm

Canvas | Replication | Backup & Restore | K8s | Data Sense | File Cache | Compute | Sync | All Services (+7) | Go to Tabular View

Add Working Environment

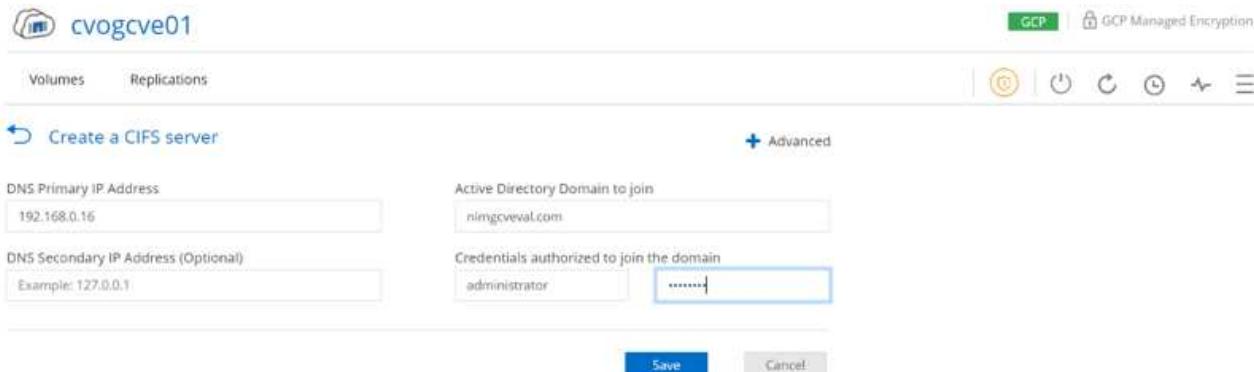
Working Environments

cvogcve01 Cloud Volumes ONTAP Freemium	DatacenterDude Azure NetApp Files 31 Volumes, 9.71 TiB Capacity	1 Cloud Volumes ONTAP 43.05 GiB Provisioned Capacity
		1 FSx for ONTAP (High-Availability) 0 B Provisioned Capacity
		1 Azure NetApp Files 9.71 TiB Provisioned Capacity

Additional configurations for SMB volumes

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

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



2. Creating the SMB volume is an easy process. At Canvas, double-click the Cloud Volumes ONTAP working environment to create and manage volumes and click on the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, CIFS/SMB is selected as the protocol.

Create new volume in cvogcve01

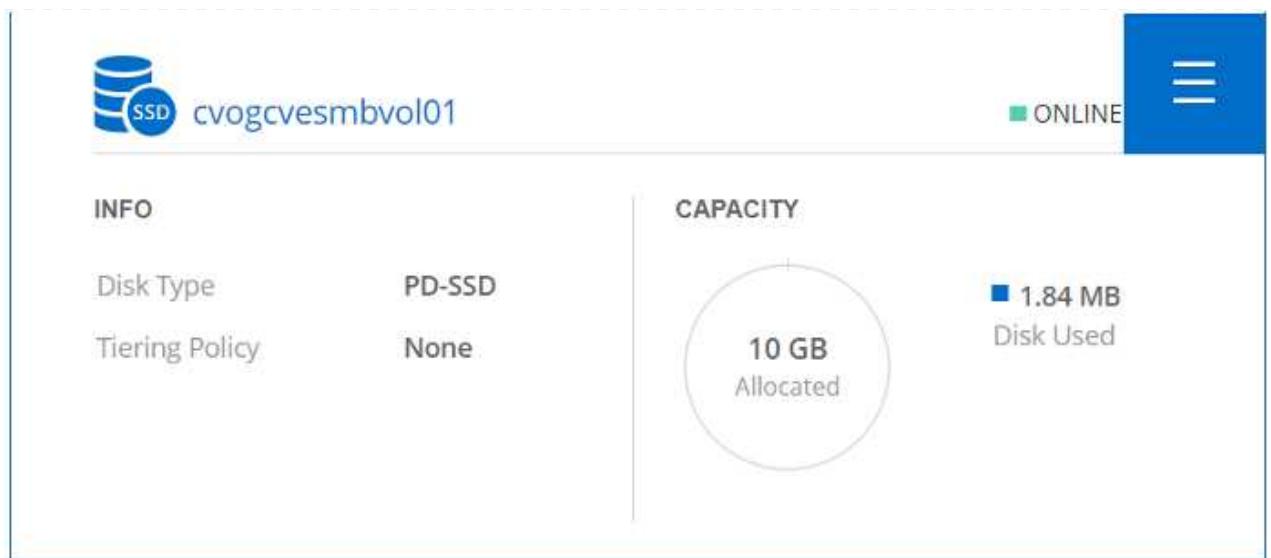
Volume Details, Protection & Protocol

Details & Protection		Protocol
Volume Name: cvogcvesmbvol01	Size (GB): 10	NFS
Snapshot Policy: default		CIFS
<input checked="" type="radio"/> Default Policy	Share name: cvogcvesmbvol01_share	Permissions: Full Control
	Users / Groups: Everyone	Valid users and groups separated by a semicolon

Continue

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

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



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

cvogcve01

Volumes Replications

Mount Volume cvogcvesmbvol01

Go to your machine and enter this command

\\10.0.6.251\cvogcvesmbvol01_share

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

Specify the drive letter for the connection and the folder that you want to connect to:

Drive:

Folder:

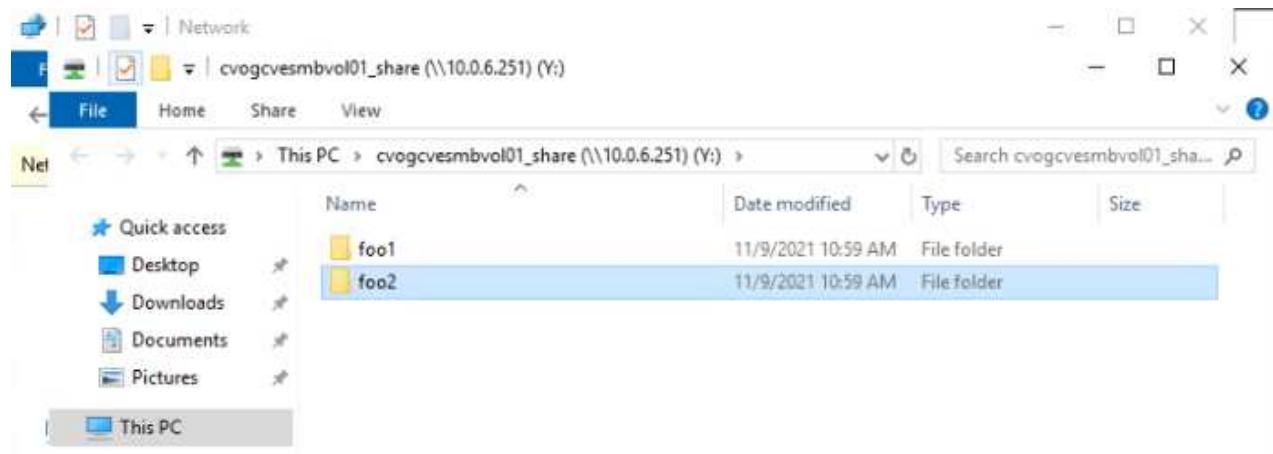
Example: \\server\share

Reconnect at sign-in

Connect using different credentials

[Connect to a Web site that you can use to store your documents and pictures.](#)

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



Connect the LUN on Cloud Volumes ONTAP to a host

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

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

The image shows two screenshots. The top screenshot is a 'Volume Details, Protection & Protocol' page from NetApp Cloud Manager. It shows a 'Details & Protection' section with 'Volume Name: cvogcvescsilun01' and 'Size (GB): 10'. The 'Protocol' section is set to 'iSCSI'. An 'Initiator Group' section shows 'WinIG' selected. The bottom screenshot is a 'Server Manager - Dashboard' window from a VMware host (vmcdc01). It shows a file share named 'embdem001_share' with contents like 'Desktop', 'Downloads', 'Documents', and 'Pictures'. The taskbar at the bottom shows the host's IP (172.16.0.2) and the date/time (10/23/2017 2:05 AM).

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

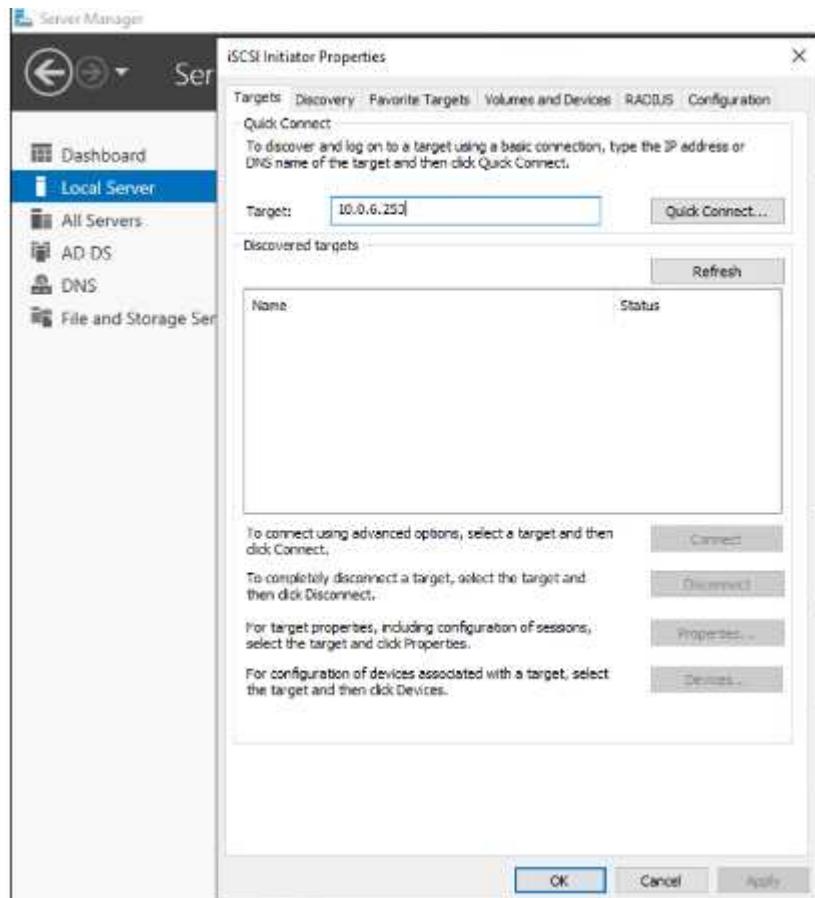
To accomplish the same for the host residing on Google Cloud VMware Engine:

- a. RDP to the VM hosted on Google Cloud VMware Engine.
- b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.
- c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.

- d. From the Targets tab, select the target discovered and then click Log on or Connect.
- e. Select Enable multipath, and then select Automatically Restore This Connection When the Computer Starts or Add This Connection to the List of Favorite Targets. Click Advanced.

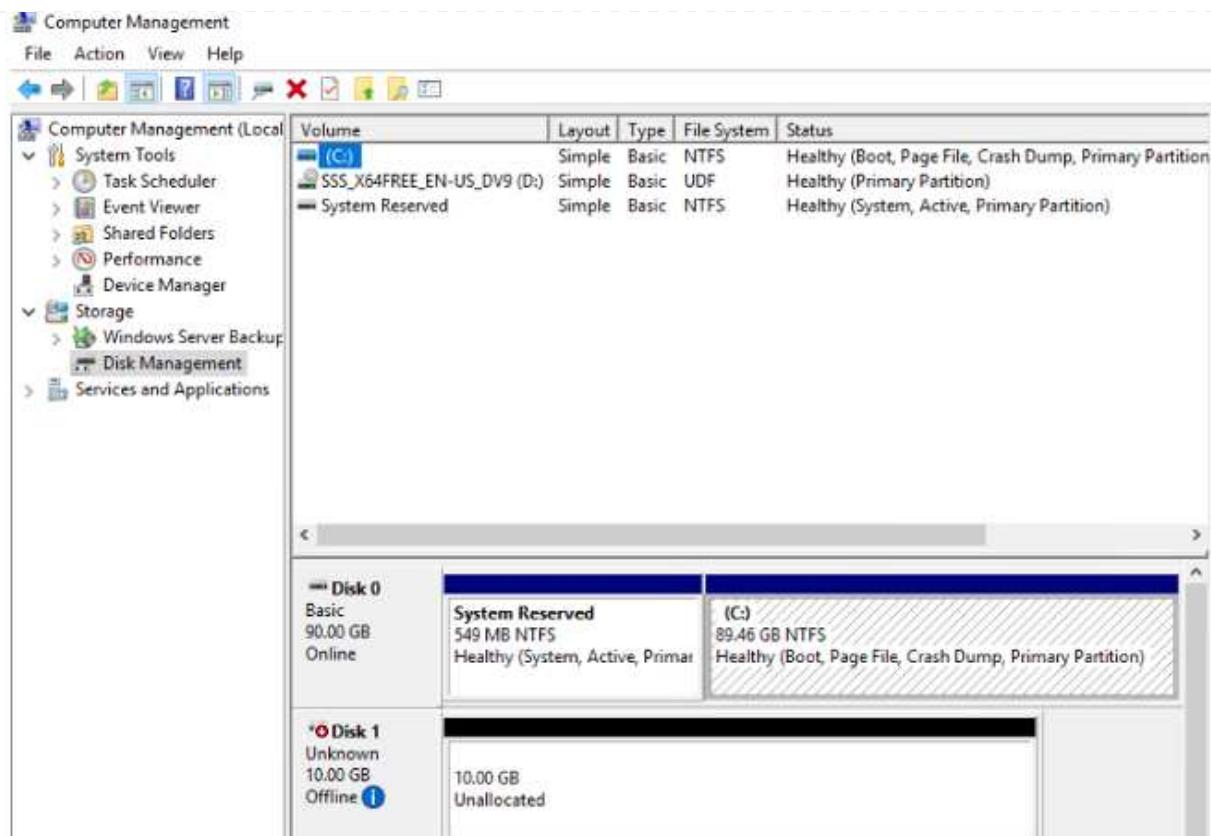


The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.



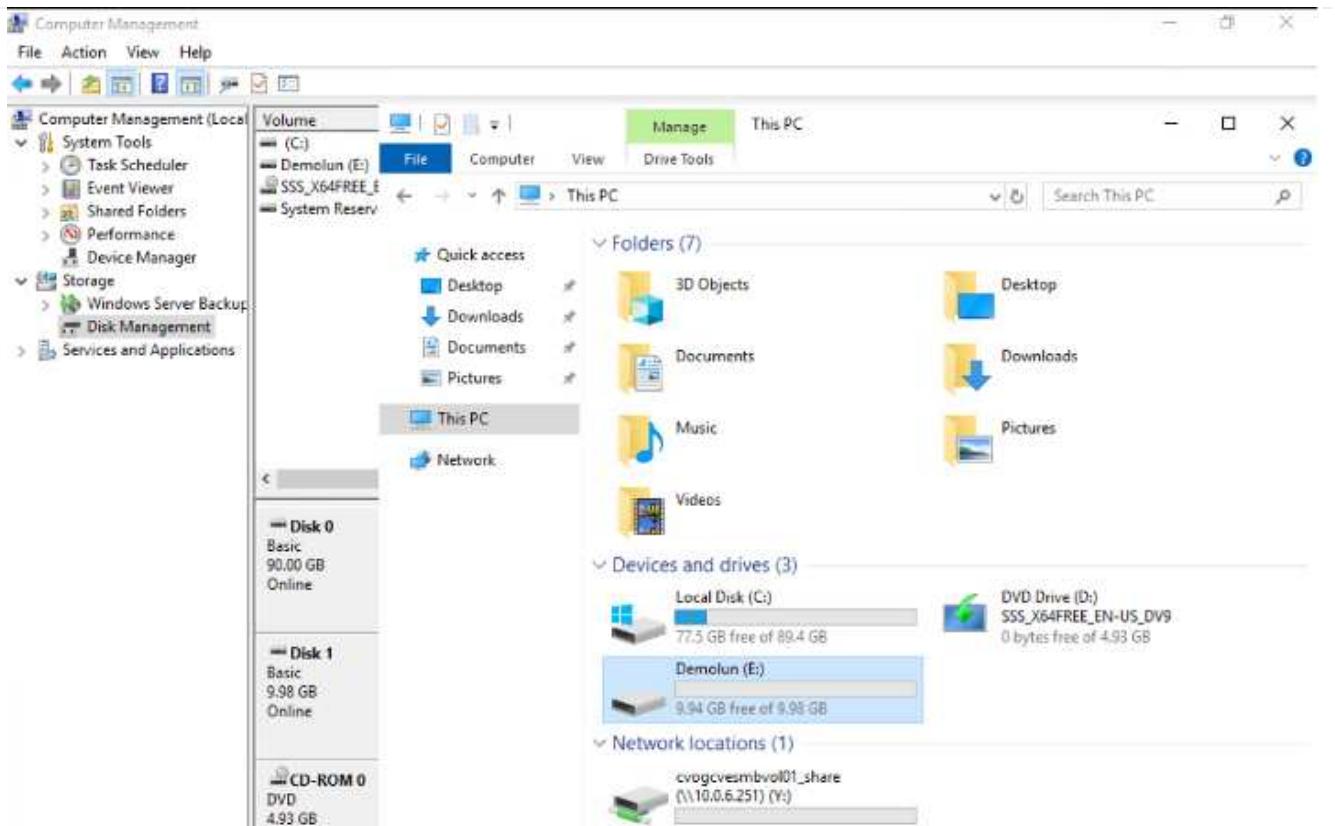
LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

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



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

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



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

```
nllyaz@nllyaz-OptiPlex-5070:~$ lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
loop0    7:0    0 55.4M  1 loop /snap/core18/2128
loop1    7:1    0 219M  1 loop /snap/gnome-3-34-1804/72
loop2    7:2    0 65.1M  1 loop /snap/gtk-common-themes/1515
loop3    7:3    0  51M  1 loop /snap/snap-store/547
loop4    7:4    0 32.3M  1 loop /snap/snapd/12704
loop5    7:5    0 32.5M  1 loop /snap/snapd/13640
loop6    7:6    0 55.5M  1 loop /snap/core18/2246
loop7    7:7    0   4K  1 loop /snap/bare/5
loop8    7:8    0 65.2M  1 loop /snap/gtk-common-themes/1519
sda      8:0    0   16G  0 disk
└─sda1   8:1    0 512M  0 part /boot/efi
└─sda2   8:2    0   1K  0 part
└─sda5   8:5    0 15.5G  0 part /
sdb      8:16   0   1G  0 disk
```

```
piyaz@nimubu01:~$ df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            1.9G   0  1.9G  0% /dev
tmpfs           394M  1.5M 392M  1% /run
/dev/sda5        16G   7.6G  6.9G  53% /
tmpfs           2.0G   0  2.0G  0% /dev/shm
tmpfs           5.0M   0  5.0M  0% /run/lock
tmpfs           2.0G   0  2.0G  0% /sys/fs/cgroup
/dev/loop1       219M  219M   0 100% /snap/gnome-3-34-1804/72
/dev/loop2       66M   66M   0 100% /snap/gtk-common-themes/1515
/dev/loop3       51M   51M   0 100% /snap/snap-store/547
/dev/loop0       56M   56M   0 100% /snap/core18/2128
/dev/loop4       33M   33M   0 100% /snap/snapd/12704
/dev/sda1       511M  4.0K  511M  1% /boot/efi
tmpfs           394M  64K  394M  1% /run/user/1000
/dev/loop5       33M   33M   0 100% /snap/snapd/13640
/dev/loop6       56M   56M   0 100% /snap/core18/2246
/dev/loop7       128K  128K   0 100% /snap/bare/5
/dev/loop8       66M   66M   0 100% /snap/gtk-common-themes/1519
/dev/sdb        976M  2.6M  967M  1% /mnt
```

Mount Cloud Volumes ONTAP NFS volume on Linux client

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

Provision the volume following the below steps

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



cvogcvenfsvol01

ONLINE

INFO

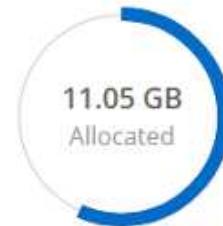
Disk Type

PD-SSD

Tiering Policy

None

CAPACITY



11.05 GB

Allocated

6.08 GB

Disk Used

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

Volumes

Replications



Mount Volume cvogcvenfsvol01

Go to your Linux machine and enter this mount command

```
mount 10.0.6.251:/cvogcvenfsvol01 <dest_dir>
```

Copy

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

```
$ sudo mkdir /cvogcvetst
```

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

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

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

```
root@nimubu01:~# sudo mount -t nfs 10.0.6.251:/cvogcvenfsvol01 cvogcvetst
```

nimubu01

Activities Terminal Nov 16 12:42

root@nimubu01:~

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
udev	1978500	0	1978500	0%	/dev
tmpfs	402272	1432	400840	1%	/run
/dev/sda5	15929256	7832332	7268048	52%	/
tmpfs	2011352	0	2011352	0%	/dev/shm
tmpfs	5120	0	5120	0%	/run/lock
tmpfs	2011352	0	2011352	0%	/sys/fs/cgroup
/dev/loop0	128	128	0	100%	/snap/bare/5
/dev/loop1	56832	56832	0	100%	/snap/core18/2128
/dev/loop2	56832	56832	0	100%	/snap/core18/2246
/dev/loop4	66688	66688	0	100%	/snap/gtk-common-themes/1515
/dev/loop5	52224	52224	0	100%	/snap/snap-store/547
/dev/loop6	66816	66816	0	100%	/snap/gtk-common-themes/1519
/dev/loop7	33280	33280	0	100%	/snap/snapd/13640
/dev/loop8	224256	224256	0	100%	/snap/gnome-3-34-1804/72
/dev/sda1	523248	4	523244	1%	/boot/cfs
tmpfs	402268	52	402216	1%	/run/user/1000
/dev/sdb	515010816	42016812	446763220	9%	/home/niyaz/cvsts
...					
10.0.6.251:/cvogcvenfsvol01	13199552	8577536	4622016	65%	/root/cvogcvetst
root@nimubu01:~					

Google Cloud NetApp Volumes (NetApp Volumes)

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

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

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

Configure NetApp Volumes with VMware Engine

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

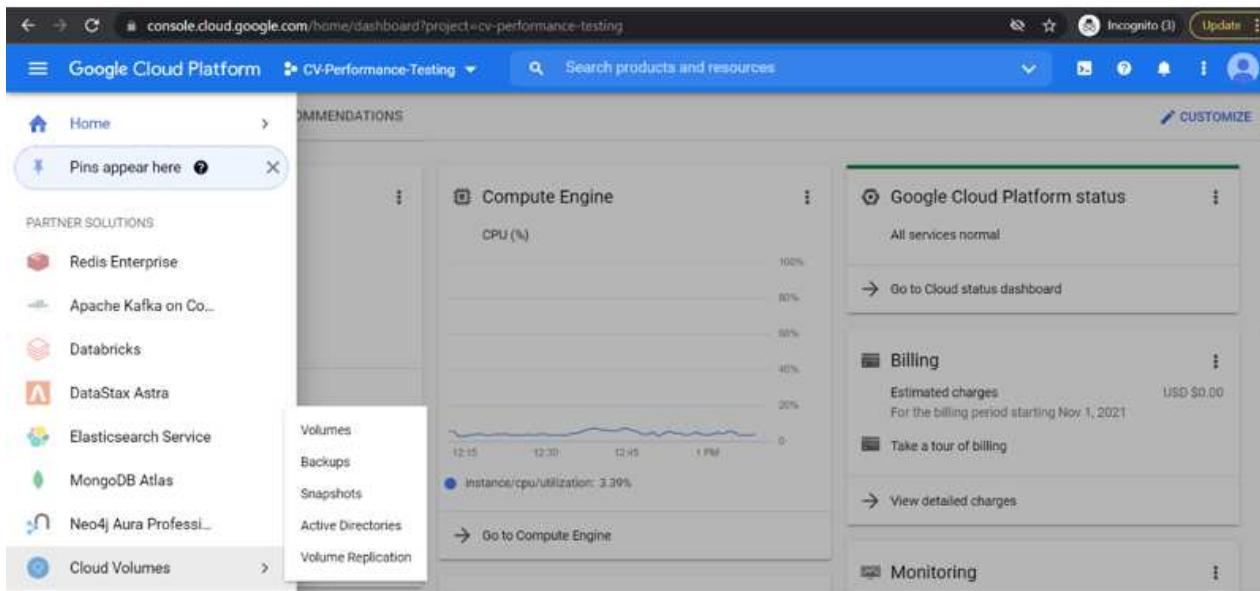
Google Cloud NetApp Volumes and Google Cloud VMware Engine private cloud must be in the same region.

To purchase, enable and configure Google Cloud NetApp Volumes for Google Cloud from the Google Cloud Marketplace, follow this detailed [guide](#).

Create a NetApp Volumes NFS volume to GCVE private cloud

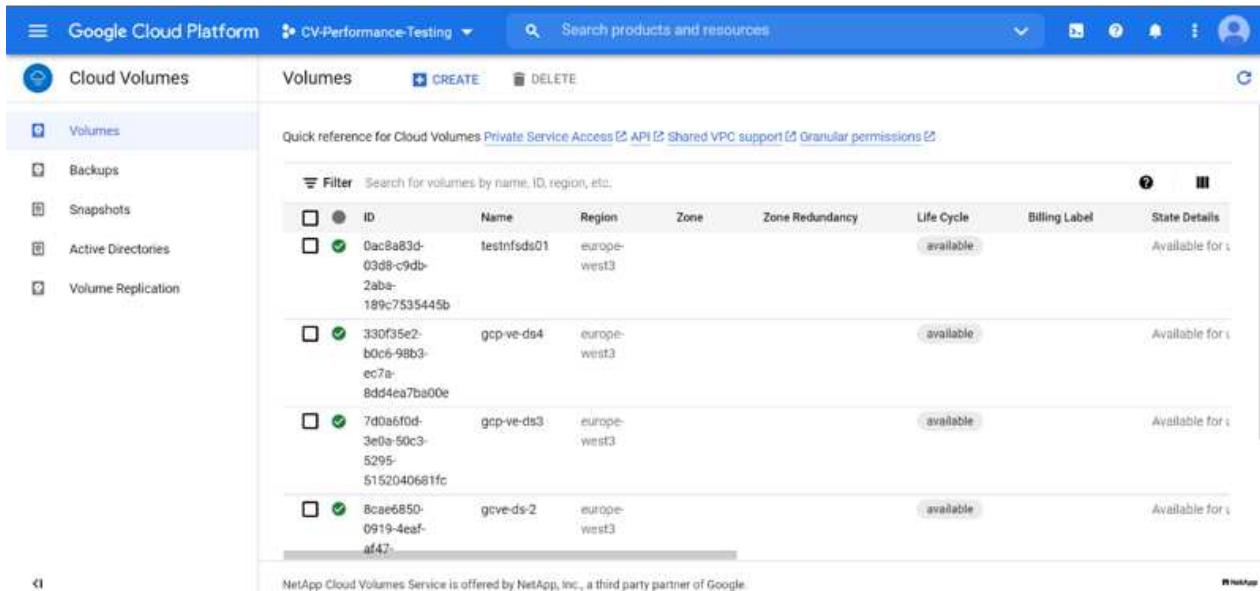
To create and mount NFS volumes, complete the following steps:

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



The screenshot shows the Google Cloud Platform dashboard for the project 'CV-Performance-Testing'. On the left, there is a sidebar with 'PARTNER SOLUTIONS' listed, including Redis Enterprise, Apache Kafka on Cloud, Databricks, DataStax Astra, Elasticsearch Service, MongoDB Atlas, Neo4j Aura Professional, and Cloud Volumes. The 'Cloud Volumes' option is selected, and a dropdown menu is open, showing 'Volumes' as the current selection. The main content area displays a 'Compute Engine' dashboard with a graph showing CPU utilization over time, with a data point 'instance/cpu/utilization: 3.30%'. To the right, there are sections for 'Google Cloud Platform status' (All services normal), 'Billing' (Estimated charges USD \$0.00 for the billing period starting Nov 1, 2021), and 'Monitoring'.

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



The screenshot shows the 'Volumes' page within the Cloud Volumes service. The left sidebar has options for 'Cloud Volumes', 'Volumes' (which is selected), 'Backups', 'Snapshots', 'Active Directories', and 'Volume Replication'. The main area has a 'CREATE' button and a table listing existing volumes. The table columns are: ID, Name, Region, Zone, Zone Redundancy, Life Cycle, Billing Label, and State Details. The table data is as follows:

ID	Name	Region	Zone	Zone Redundancy	Life Cycle	Billing Label	State Details
Dac8a83d-03db-c9db-2aba-189c7535445b	testnfsds01	europe-west3			available		Available for use
330f35e2-b0f6-98b3-ec7a-8dd4ea7ba00e	gcp-ve-ds4	europe-west3			available		Available for use
7d0a6f0d-3e0a-50c3-5295-5152040681fc	gcp-ve-ds3	europe-west3			available		Available for use
8cae6850-0919-4eaf-af47-	gcve-ds-2	europe-west3			available		Available for use

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

 Cloud Volumes <ul style="list-style-type: none">  Volumes  Backups  Snapshots  Active Directories  Volume Replication 	<p style="text-align: center;">← Create File System</p> <p> Volumes</p> <p> Backups</p> <p> Snapshots</p> <p> Active Directories</p> <p> Volume Replication</p>
--	---

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

 Cloud Volumes <ul style="list-style-type: none">  Volumes  Backups  Snapshots  Active Directories  Volume Replication 	<p style="text-align: center;">← Create File System</p> <p> Volumes</p> <p> Backups</p> <p> Snapshots</p> <p> Active Directories</p> <p> Volume Replication</p>
---	--

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

 Cloud Volumes <ul style="list-style-type: none">  Volumes  Backups  Snapshots  Active Directories  Volume Replication 	<p>← Create File System</p> <p>Region Region availability varies by service type.</p> <p>Region * <input type="text" value="europe-west3"/> ? </p> <p>Volume will be provisioned in the region you select.</p> <p>Volume Path * <input type="text" value="nimCVSNFS01"/> </p> <p>Must be unique to the project.</p>
--	--

6. Select the level of performance for the volume.

 Cloud Volumes <ul style="list-style-type: none">  Volumes  Backups  Snapshots  Active Directories  Volume Replication 	<p>← Create File System</p> <p>Service Level Select the performance level required for your workload.</p> <p><input checked="" type="radio"/> Standard Up to 16 MiB/s per TiB</p> <p><input type="radio"/> Premium Up to 64 MiB/s per TiB</p> <p><input type="radio"/> Extreme Up to 128 MiB/s per TiB</p> <p>Snapshot ▼</p> <p>The snapshot to create the volume from.</p>
---	---

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

 Cloud Volumes <ul style="list-style-type: none">  Volumes  Backups  Snapshots  Active Directories  Volume Replication 	<p>← Create File System</p> <p>Volume Details</p> <p>Allocated Capacity * <input type="text" value="1024"/> GiB</p> <p>Allocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB)</p> <p>Protocol Type * <input type="text" value="NFSv3"/> ▼</p> <p><input type="checkbox"/> Make snapshot directory (.snapshot) visible Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only), the directory itself will not be listed but can be accessed to list contents, etc.</p> <p><input type="checkbox"/> Enable LDAP Enables user look up from AD LDAP server for your NFS volumes</p>
--	---

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

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

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

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

10. Click Save to create the volume.

<input type="checkbox"/>	<input checked="" type="checkbox"/> 4b5ed909- bc6d-f3d5- 5a0f- 7da26aed3ed0	nimnfsdemods02	europe- west3	Available for use	CVS- Performance	Primary	Extreme	NFSv3 : 10.53.0.4/nimnfsdemods02
--------------------------	--	----------------	------------------	-------------------	---------------------	---------	---------	----------------------------------

Mounting NFS exports to VMs running on VMware Engine

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

To mount an NFS volume, do the following:

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

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

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

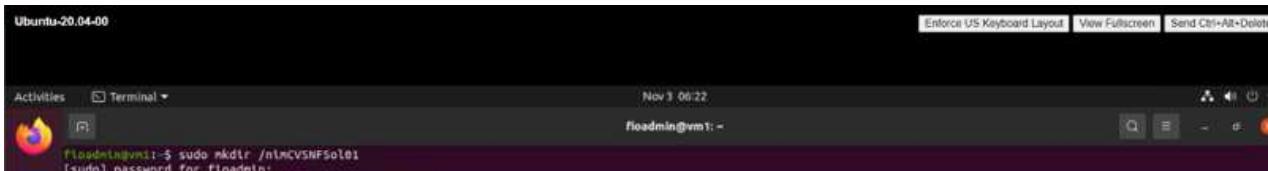
```
sudo yum install -y nfs-utils
```

- b. On an Ubuntu or Debian instance:

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

3. Create a new directory on the instance, such as "/nimCVSNFS01":

```
sudo mkdir /nimCVSNFS01
```



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

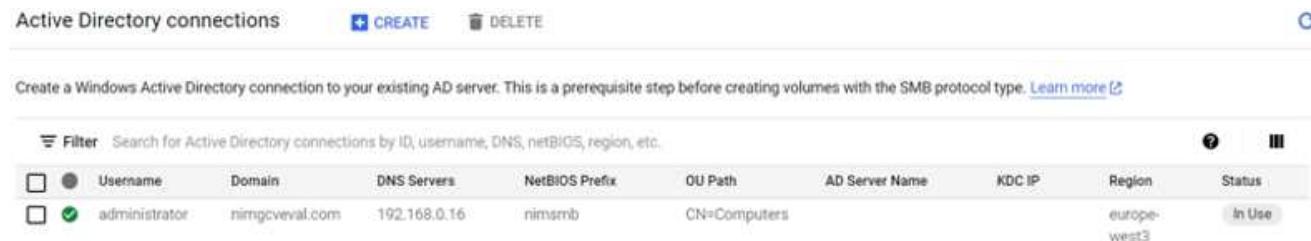
```
sudo mount -t nfs -o rw,hard,rsize=65536,wszie=65536,vers=3,tcp  
10.53.0.4:/nimCVSNFS01 /nimCVSNFS01
```

```
root@vm1:~# sudo mkdir nimCVSNFS01  
root@vm1:~# sudo mount -t nfs -o rw,hard,rsize=65536,wszie=65536,vers=3,tcp 10.53.0.4:/nimCVSNFS01 /nimCVSNFS01
```

```
root@vm1:~# df
Filesystem      1K-blocks   Used   Available Use% Mounted on
udev             16409952     0   16409952  0% /dev
tmpfs            3288328   1580   3286748  1% /run
/dev/sdb5        61145932 19231356  38778832  34% /
tmpfs            16441628     0   16441628  0% /dev/shm
tmpfs            5120       0     5120  0% /run/lock
tmpfs            16441628     0   16441628  0% /sys/fs/cgroup
/dev/loop0        128       128     0 100% /snap/bare/5
/dev/loop1        56832     56832     0 100% /snap/core18/2128
/dev/loop2        66688     66688     0 100% /snap/gtk-common-themes/1515
/dev/loop4        66816     66816     0 100% /snap/gtk-common-themes/1519
/dev/loop3        52224     52224     0 100% /snap/snap-store/547
/dev/loop5        224256    224256     0 100% /snap/gnome-3-34-1804/72
/dev/sdb1        523248     4     523244  1% /boot/efi
tmpfs            3288324    28     3288296  1% /run/user/1000
10.53.0.4:/gcve-ds-1 107374182400 1136086016 106238096384  2% /base
/dev/napperv/nfsprdvgl-prod01 419155968 55384972 363770996 14% /datastore1
/dev/loop8        33280     33280     0 100% /snap/snapd/13270
/dev/loop6        33280     33280     0 100% /snap/snapd/13640
/dev/loop7        56832     56832     0 100% /snap/core18/2246
10.53.0.4:/nimCVSNFSol01 107374182400   256 107374182144  1% /nimCVSNFSol01
root@vm1:~#
```

Creating and Mounting SMB Share to VMs running on VMware Engine

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



Username	Domain	DNS Servers	NetBIOS Prefix	OU Path	AD Server Name	KDC IP	Region	Status
administrator	nimcvceval.com	192.168.0.16	nimsmb	CN=Computers			europe-west3	In Use

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

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

[←](#) Create File System

Volume Name

Name * nimCVSMBvol01

A human readable name used for display purposes.

Billing Label

Label your volumes for billing reports, queries.

Supported with CVS-Performance service type; can be set with CVS service type but not available for billing at this time.

[+ ADD LABEL](#)

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

Create File System

Service Type

Cloud Volumes Service is offered as two service types: CVS and CVS-Performance.

Select the service type that matches your workload needs. [Region availability](#) varies by service type. [Learn more](#)

CVS

Offers volumes created with zonal high availability.

CVS-Performance

Offers 3 performance levels and improved latency to address higher performance application requirements.

Volume Replication

Secondary

Select to create volume as a destination target for volume replication. Applicable only to CVS-performance volumes.

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

Create File System

Region

Region availability varies by service type.

Region *
europe-west3



Volume will be provisioned in the region you select.

Volume Path *
nimCVSMBvol01



Must be unique to the project.

5. Select the level of performance for the volume.

Create File System

Service Level

Select the performance level required for your workload.

Standard

Up to 16 MiB/s per TiB

Premium

Up to 64 MiB/s per TiB

Extreme

Up to 128 MiB/s per TiB

Snapshot



The snapshot to create the volume from.

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

Create File System

Volume Details

Allocated Capacity *

1024

GiB

Allocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB)

Protocol Type *

SMB



Make snapshot directory (.snapshot) visible

Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only), the directory itself will not be listed but can be accessed to list contents, etc.

Enable SMB Encryption

Enable this option only if you require encryption of your SMB data traffic.

Enable CA share support for SQL Server, FSLogix

Enable this option only for SQL Server and FSLogix workloads that require continuous availability.

Hide SMB Share

Enable this option to make SMB shares non-browsable

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

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

beforehand, refer to these [instructions](#).

Network Details

Shared VPC configuration

Provide the host project name when deploying in a shared VPC service project.

VPC Network Name *

cloud-volumes-vpc

Select the VPC Network from which the volume will be accessible. This cannot be changed later.

Use Custom Address Range

Reserved Address range

netapp-addresses

SHOW SNAPSHOT POLICY

SAVE

CANCEL

8. Click Save to create the volume.

	6a4552ed-7378-7302-be28-21a16934f28	nimCVSMBvol01	europe-west3	Available for use	CVS-Performance	Primary	Standard	SMB : \\\nimsmb-3830.nimgcveval.com\\nimCVSMBvol01
<input type="checkbox"/>								

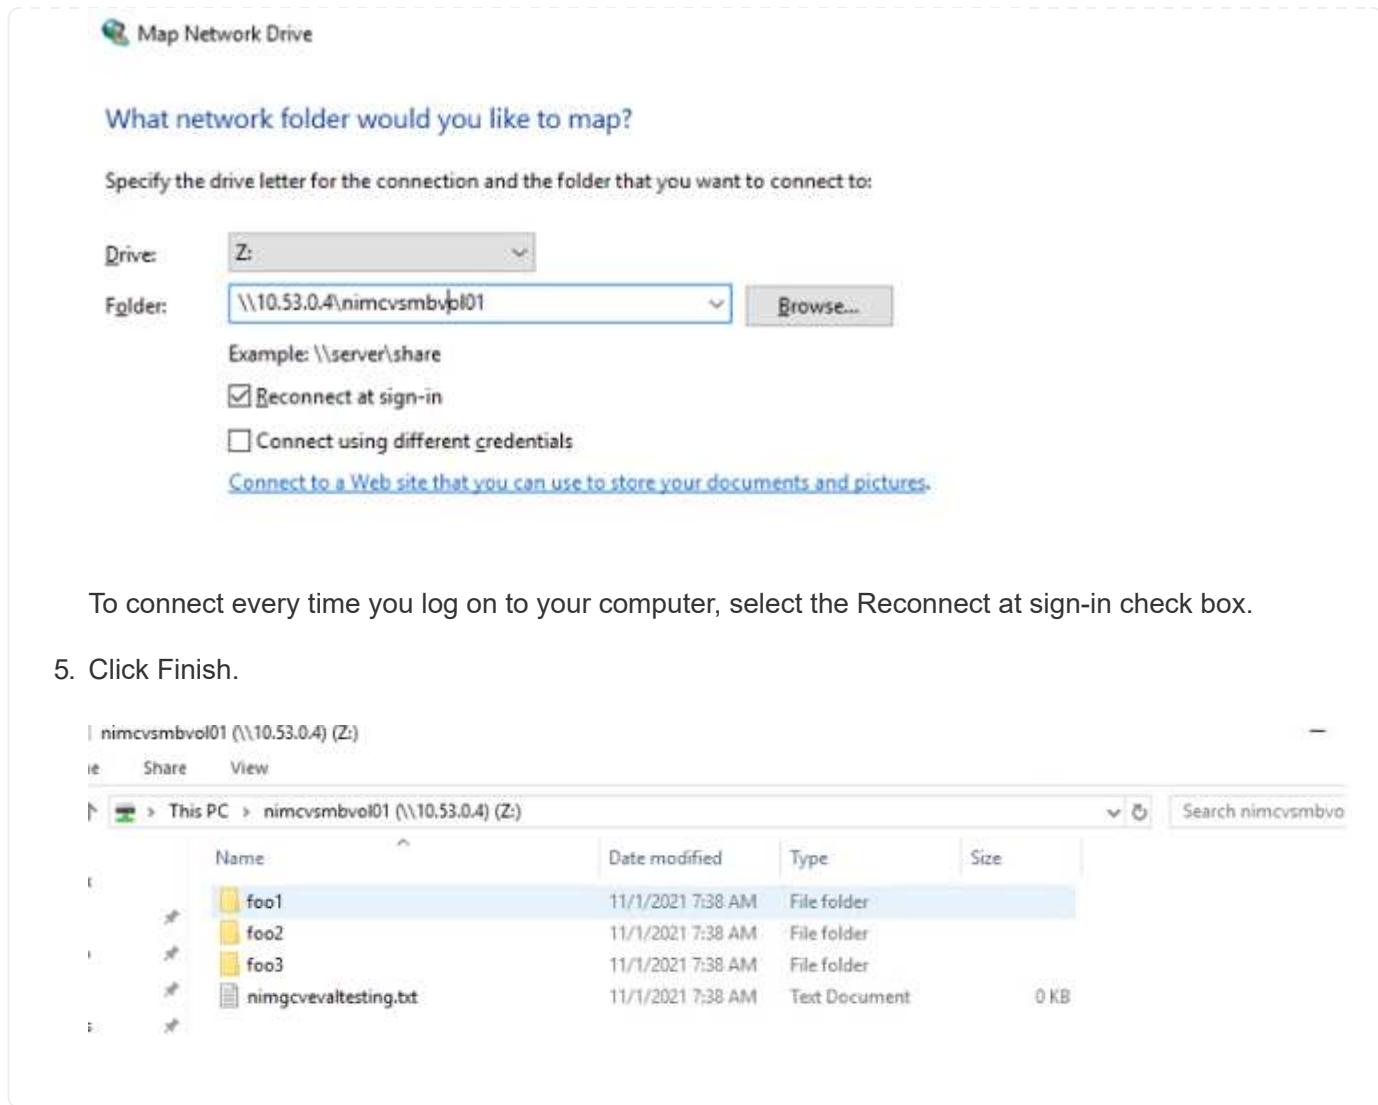
To mount the SMB volume, do the following:

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

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

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

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



Map Network Drive

What network folder would you like to map?

Specify the drive letter for the connection and the folder that you want to connect to:

Drive: Z:

Folder: \\10.53.0.4\imcvsmbv01

Example: \\server\share

Reconnect at sign-in

Connect using different credentials

[Connect to a Web site that you can use to store your documents and pictures.](#)

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

5. Click Finish.

Name	Date modified	Type	Size
foo1	11/1/2021 7:38 AM	File folder	
foo2	11/1/2021 7:38 AM	File folder	
foo3	11/1/2021 7:38 AM	File folder	
nimgcvevaltesting.txt	11/1/2021 7:38 AM	Text Document	0 KB

Summary and Conclusion: Why NetApp Hybrid Multicloud with VMware

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

Use case #1: Optimizing storage

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

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

from warm storage to production-level performance in seconds.

Use case #2: Cloud migration

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

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

Use case #3: Data center expansion

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

Use case #4: Disaster recovery to the cloud

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

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

Use case #5: Application modernization

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

Conclusion

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

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

It is the same familiar process and procedures that are used to connect the storage. Remember, it is just the

position of the data that changed with new names; the tools and processes all remain the same and NetApp Cloud Volumes helps in optimizing the overall deployment.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.