VMware for Public Cloud

NetApp Solutions

NetApp
April 26, 2024

Table of Contents

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 for NetApp ONTAP, Cloud Volumes ONTAP (available in all three major hyperscalers), and Cloud Volumes Service 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 for NetApp ONTAP, Cloud Volumes ONTAP (available in all three major hyperscalers), and Cloud Volumes Service 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 for NetApp ONTAP, Cloud Volume ONTAP, Azure NetApp Files and Cloud Volumes Service 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. As supplemental NFS datastore support becomes available, additional documentation will be available here.

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 for NetApp ONTAP to ease migration using block-level mechanisms. This does not apply to Azure NetApp Files and Cloud Volumes Services. For migrating data to Azure NetApp Files or Cloud Volumes Services, 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. As supplemental NFS datastore support becomes available, additional documentation will be available here.

> 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, Cloud Volumes Service for Google Cloud and Azure NetApp Files as an additional in-guest storage option.

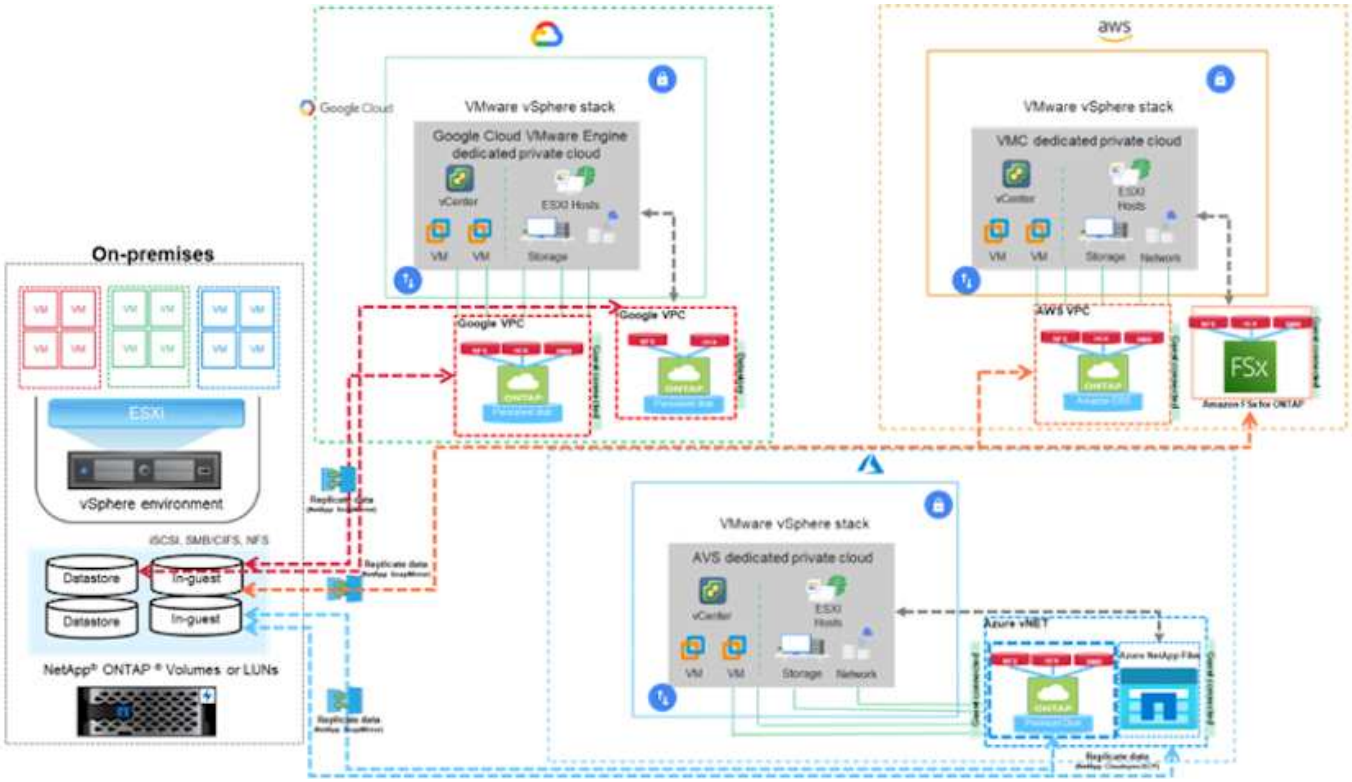

NetApp Solutions for VMware in Hyperscalers

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 Cloud Volumes Services 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
- Cloud Volumes Service (CVS) as guest connected storage
- Cloud Volumes Service (CVS) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.

Read more about NetApp Cloud Volumes Service datastore support for Google Cloud VMware Engine (NetApp blog) or How to use NetApp CVS 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 CVS Details	CVS Details

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 Cloud Volumes Services 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 for NetApp 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.

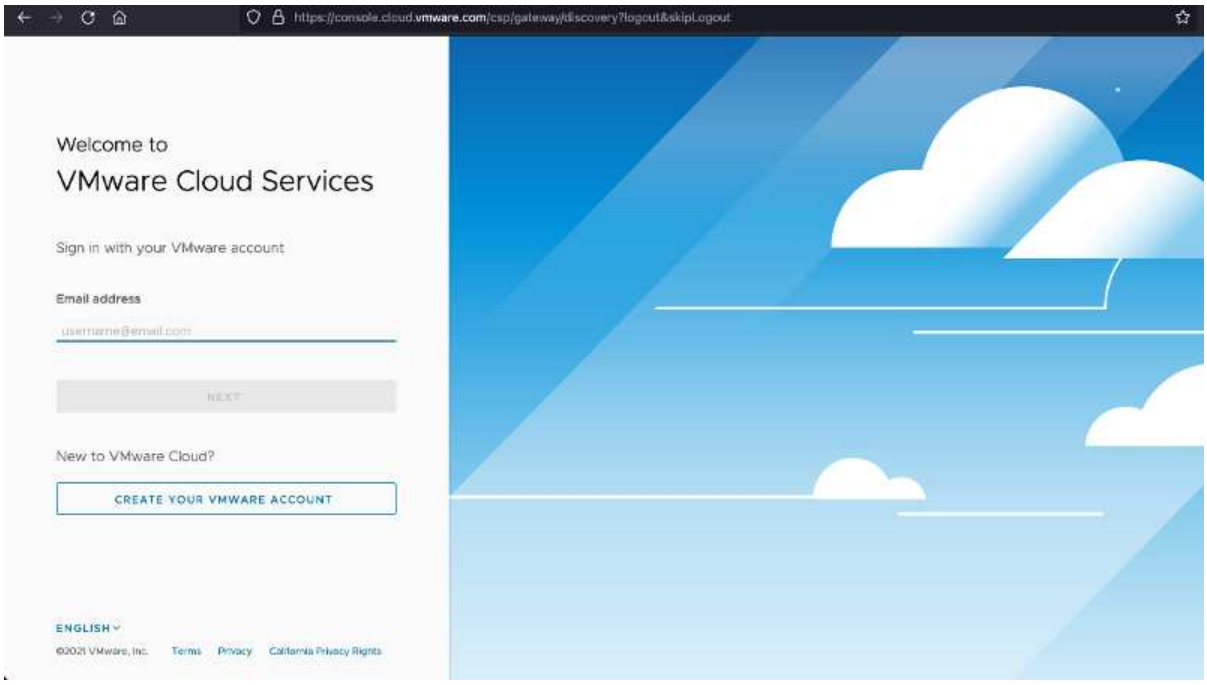

2. Configure the AWS region, deployment, and host type and the SDDC name:

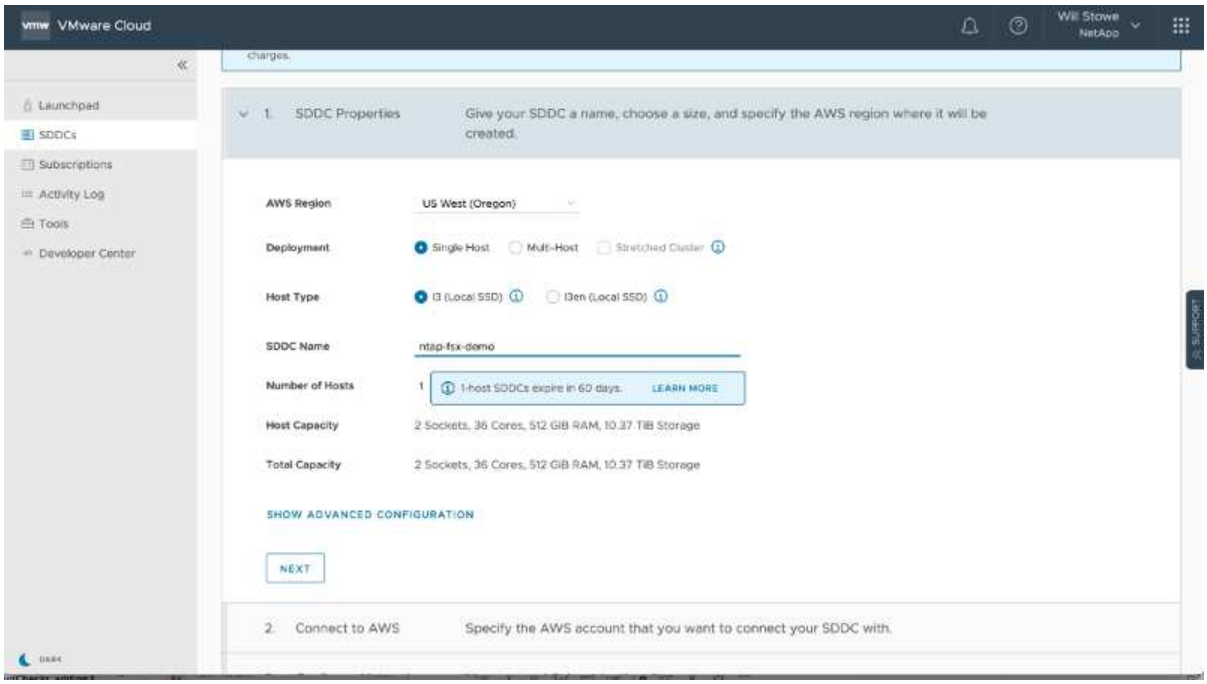

3. Connect to the desired AWS account and execute the AWS Cloud Formation stack.

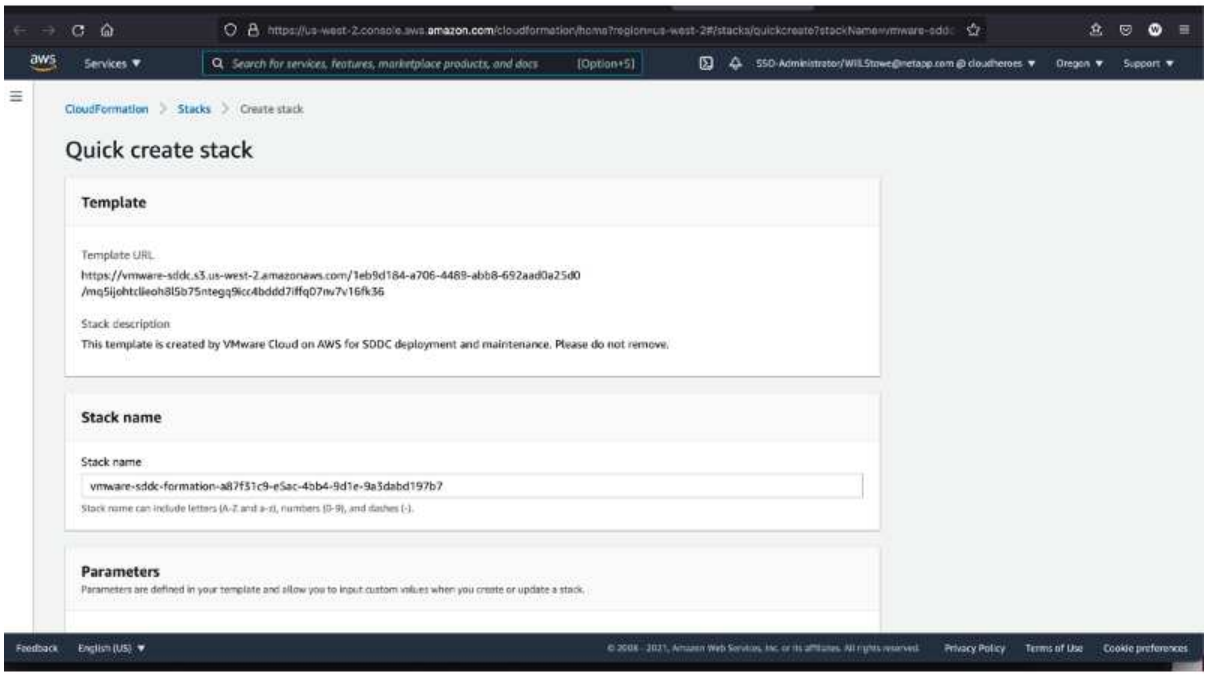
CloudFormation > Stacks > Create stack
Quick create stack
Template
Template URL
https://vmware-sddc.s3.us-west-2.amazonaws.com/1eb9d184-a706-4489-abb8-692aad0a25d0/mq5ijohtclieoh8l5b75ntegq9icc4bddd7iffq07nv7v16fk36
Stack description
This template is created by VMware Cloud on AWS for SDDC deployment and maintenance. Please do not remove.
Stack name
Stack name
vmware-sddc-formation-a87f51c9-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.

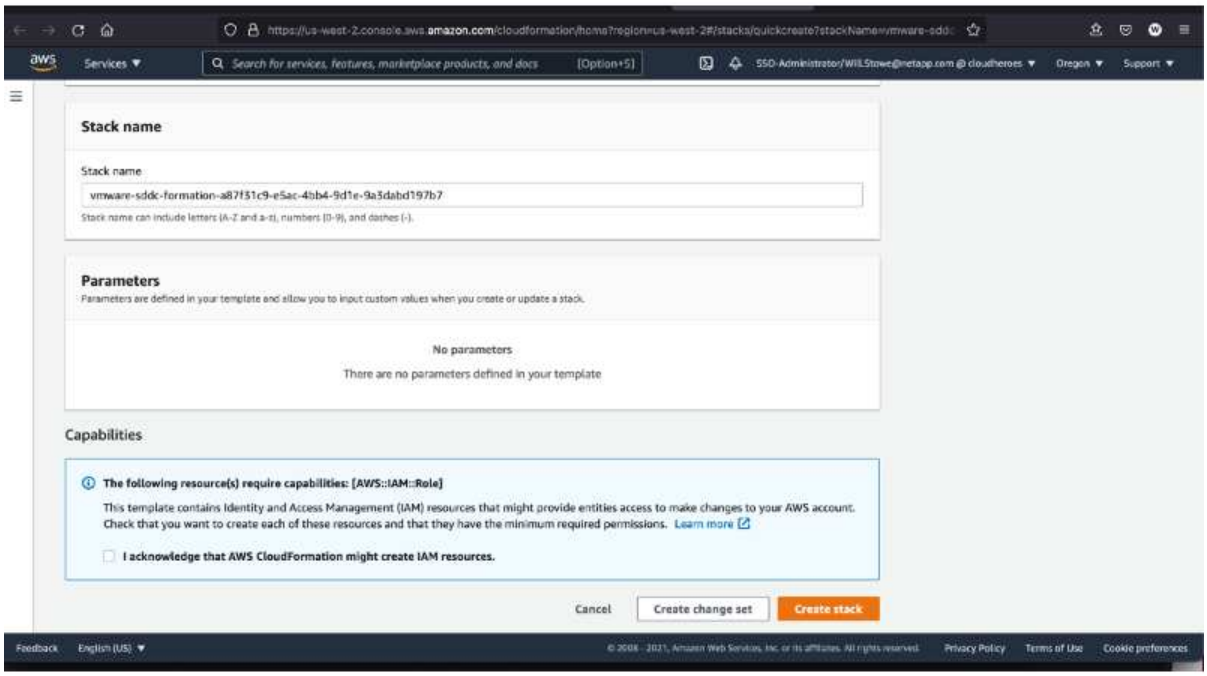
Stack name
Stack name
vmware-sddc-formation-a87f51c9-e5ac-4bb4-9d1e-9a3dabd197b7
Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).
Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.
No parameters
There are no parameters defined in your template
Capabilities
The following resource(s) require capabilities: [AWS::IAM::Role]
This template contains Identity and Access Management (IAM) resources that might provide entities access to make changes to your AWS account. Check that you want to create each of these resources and that they have the minimum required permissions. Learn more
I acknowledge that AWS CloudFormation might create IAM resources.
Cancel
Create change set
Create stack

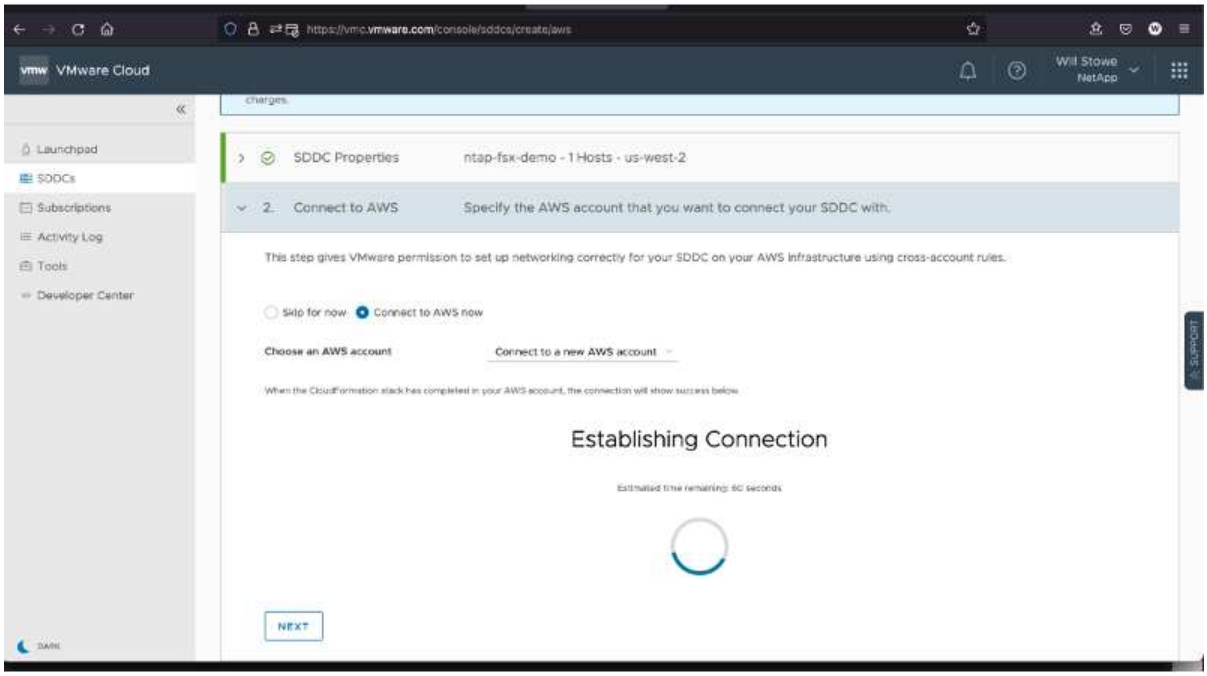

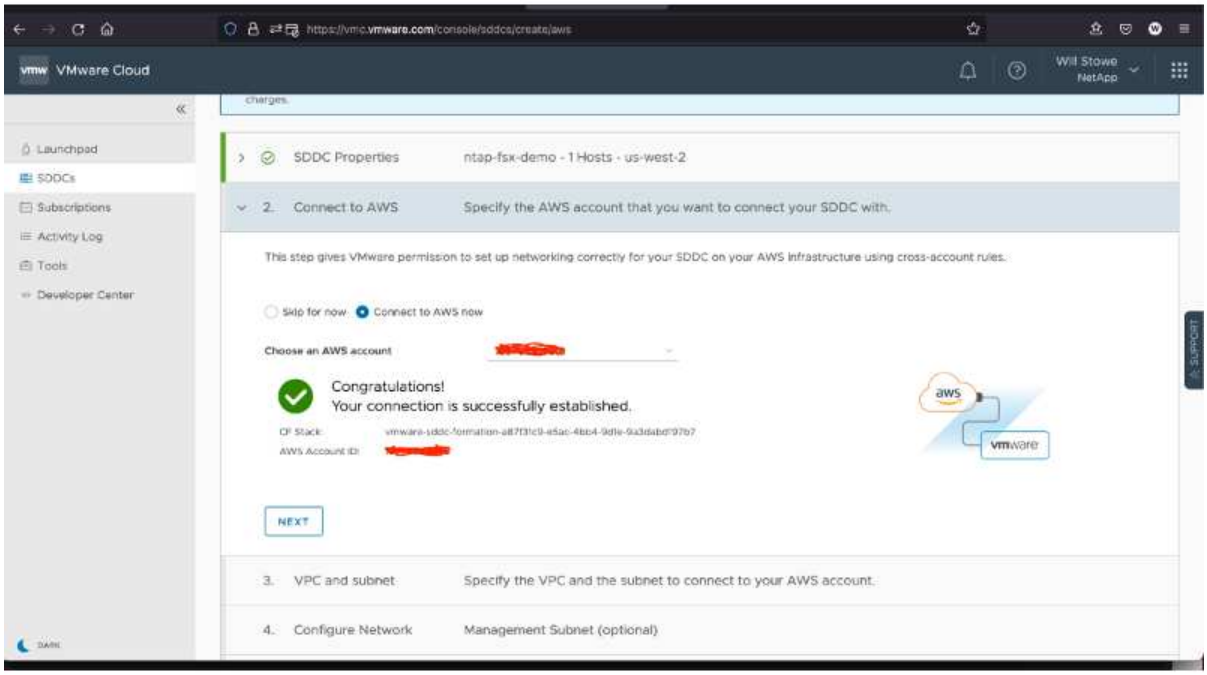

Single-host configuration is used in this validation.

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

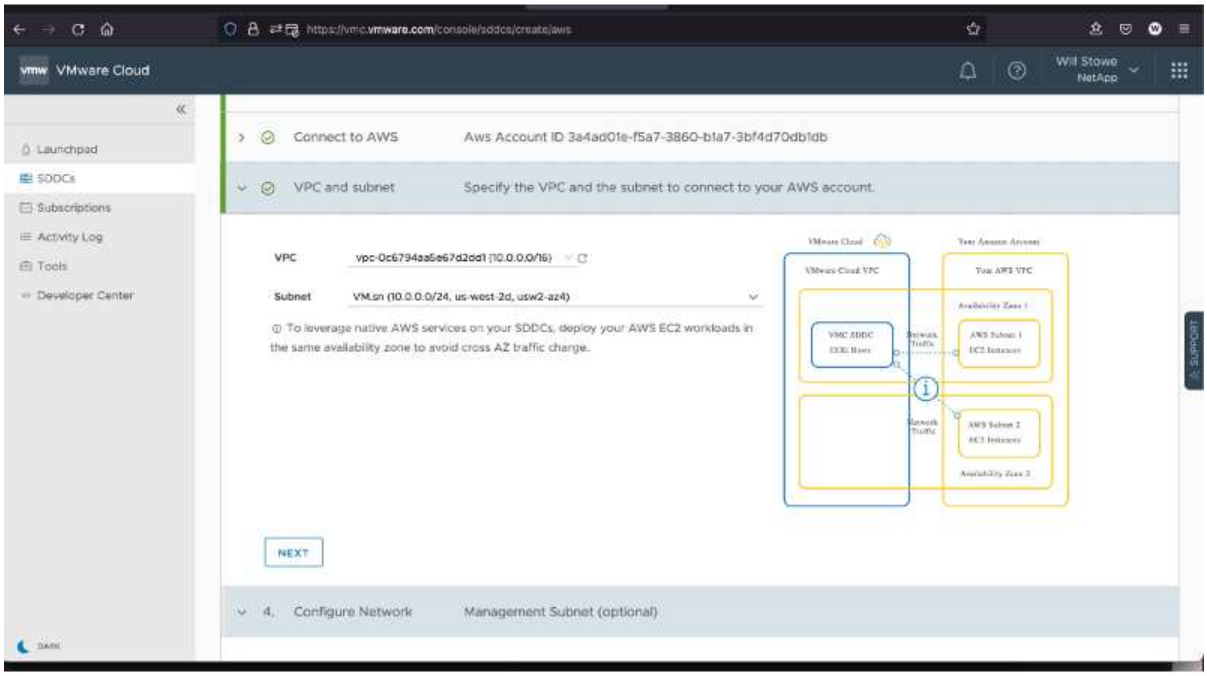

5. Configure the VMC Management Subnet; this subnet contains VMC-managed services like vCenter, NSX, and so on. Do not choose an overlapping address space with any other networks that need connectivity to the SDDC environment. Finally, follow the recommendations for CIDR size notated below.

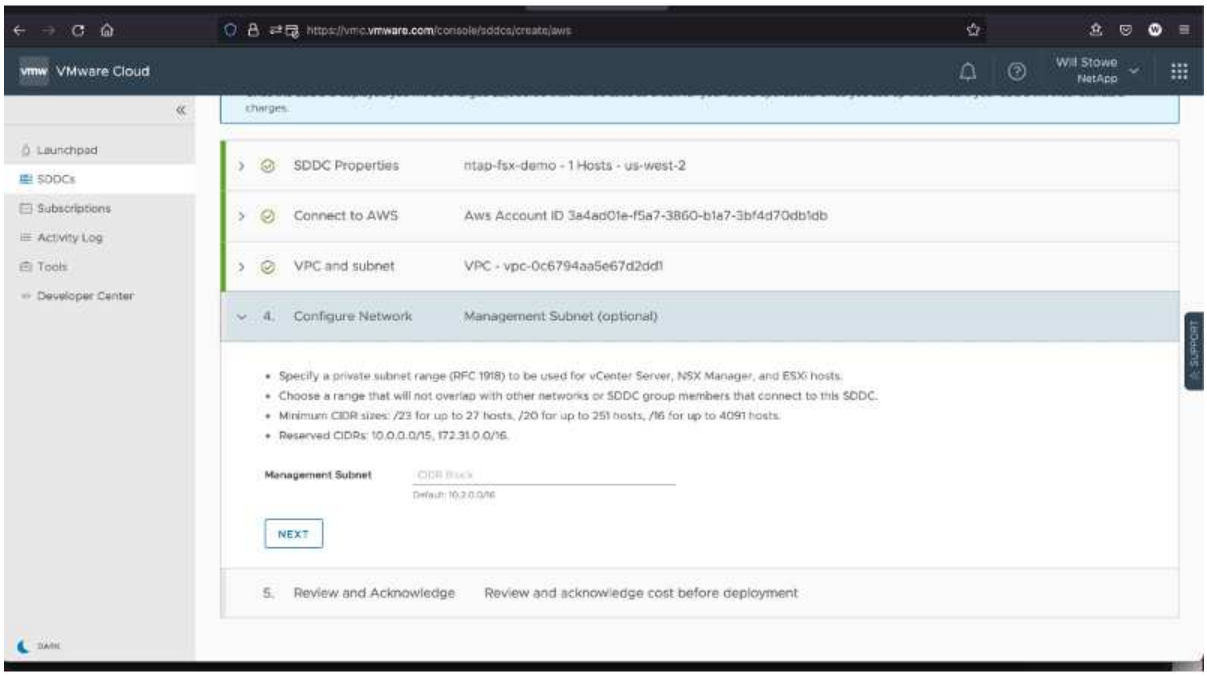

6. Review and acknowledge the SDDC configuration, and then click deploy the SDDC.

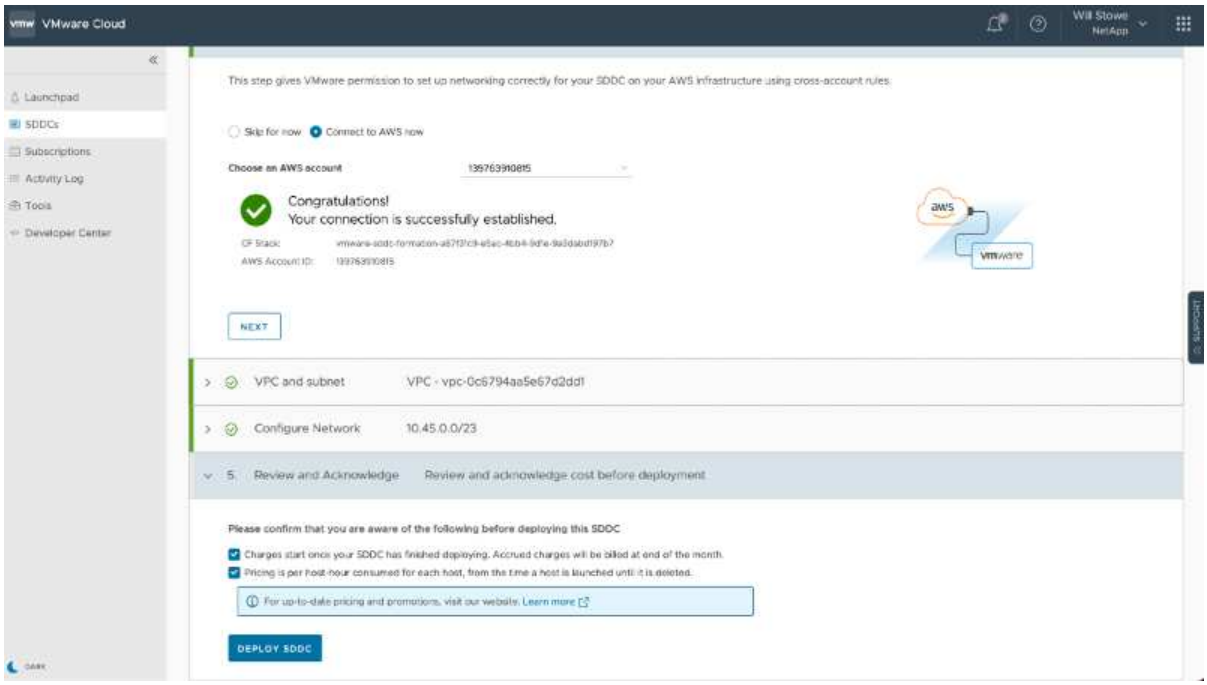

The deployment process typically takes approximately two hours to complete.

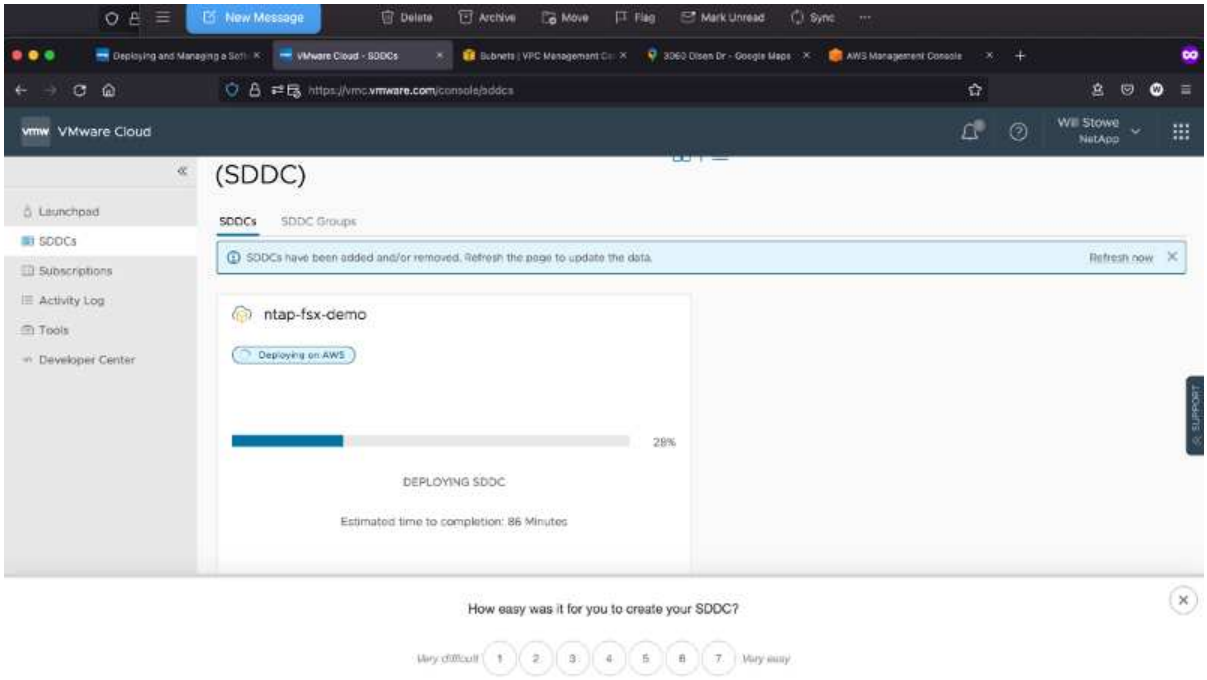

7. After completion, the SDDC is ready for use.

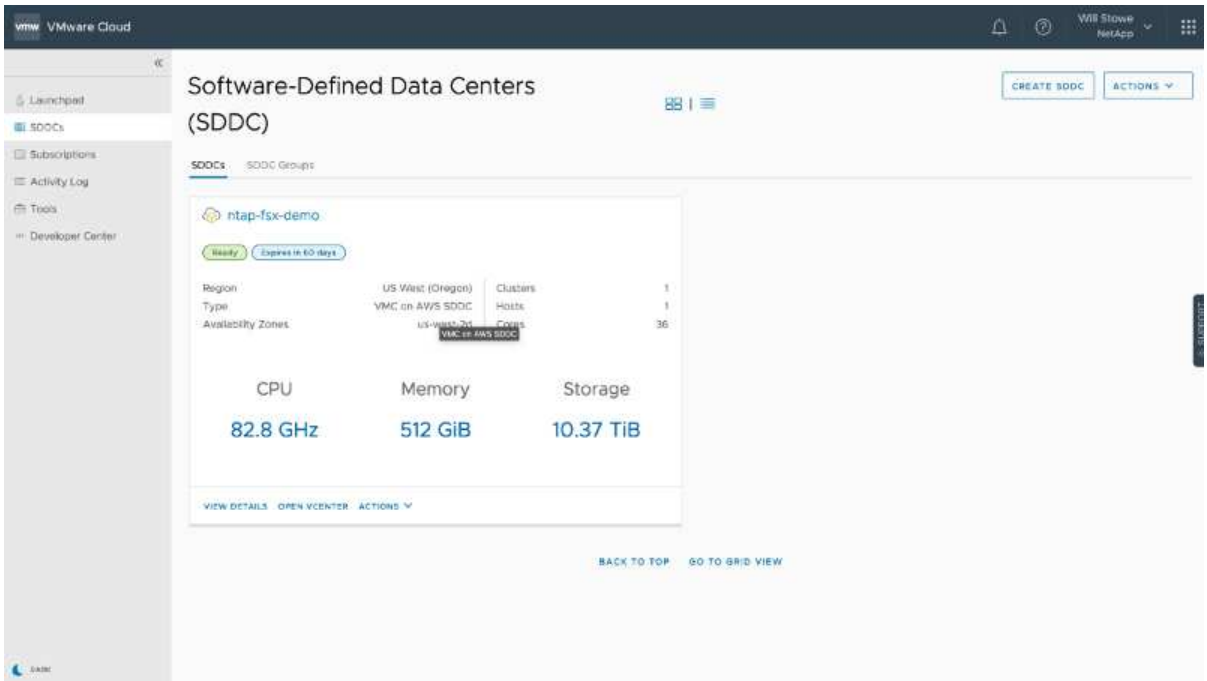

For a step-by-step guide on SDDC deployment, see Deploy an SDDC from the VMC Console.

Connect VMware Cloud to FSx ONTAP

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

1. With VMware Cloud deployment completed and connected to AWS VPC, you must deploy Amazon FSx for NetApp 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.

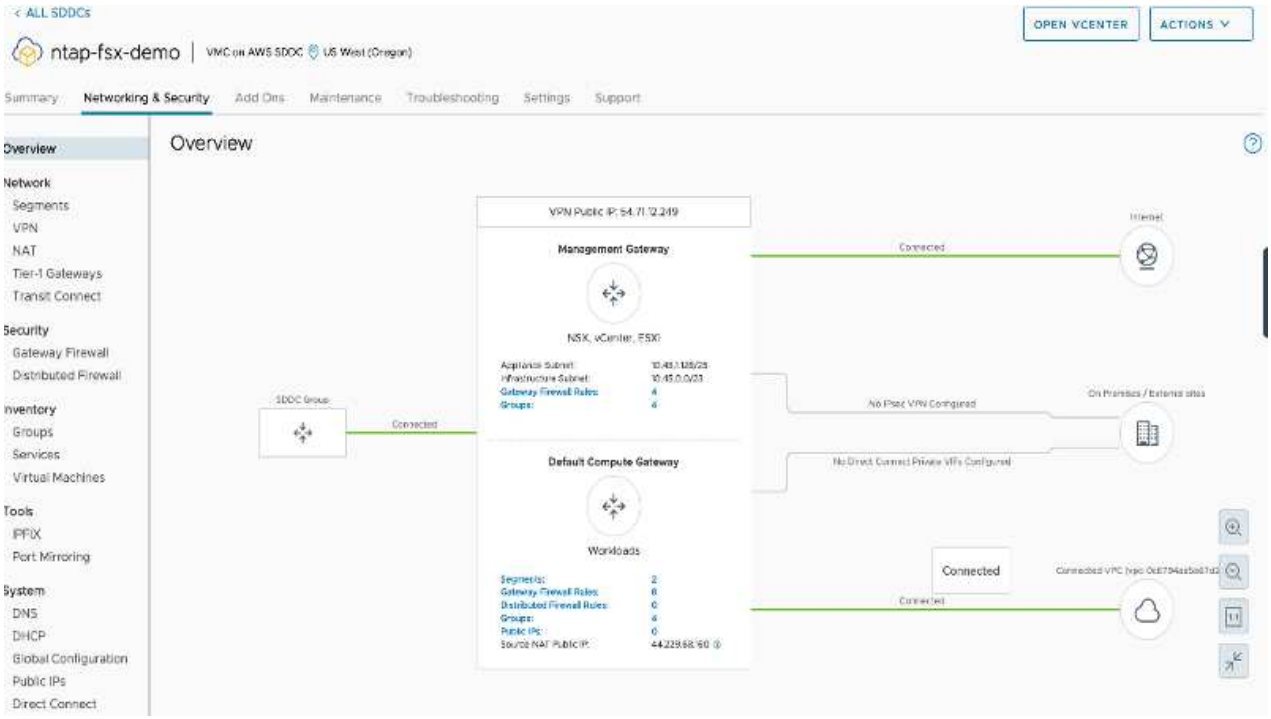

2. Deploy an additional VPC in the same region, and then deploy Amazon FSx for NetApp ONTAP into the new VPC.

   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.

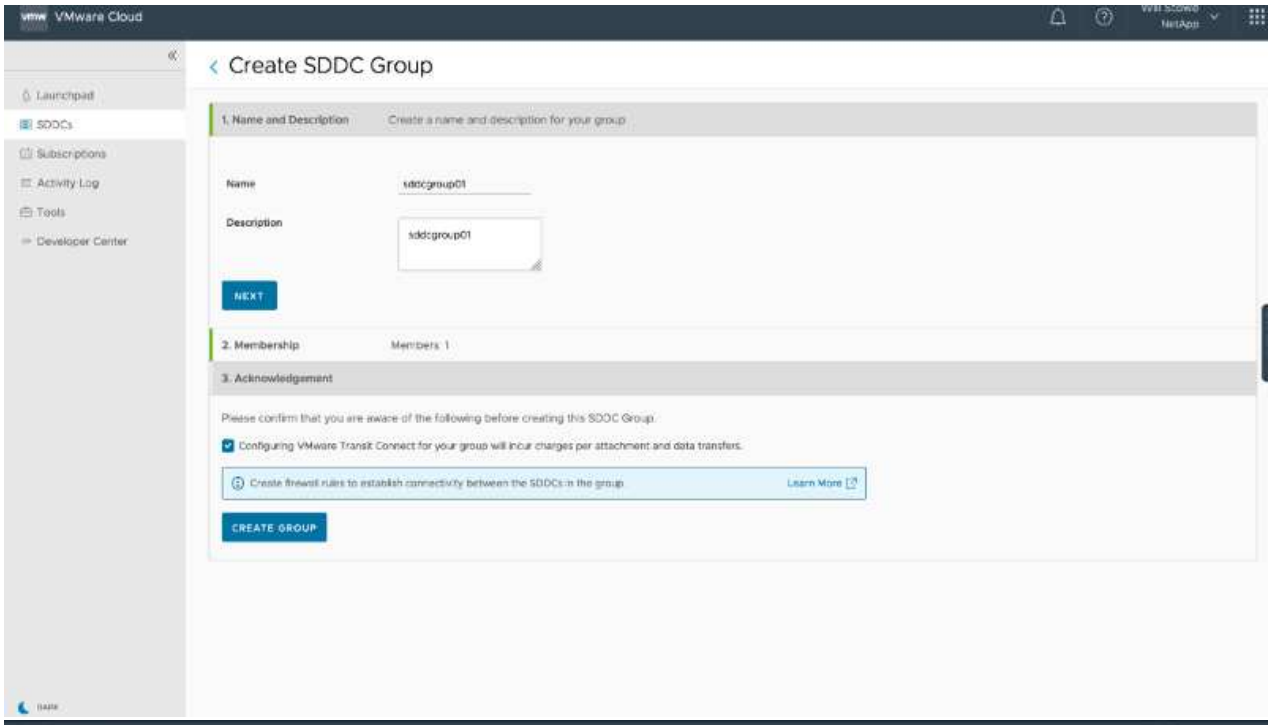
VMware Cloud
Will Stowe
NetApp
Launchpad
SDDCs
Subscriptions
Activity Log
Tools
Developer Center
Create SDDC Group
1. Name and Description
Create a name and description for your group
Name
sddcgroup01
Description
sddcgroup01
NEXT
2. Membership
Members: 1
3. Acknowledgement
Please confirm that you are 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

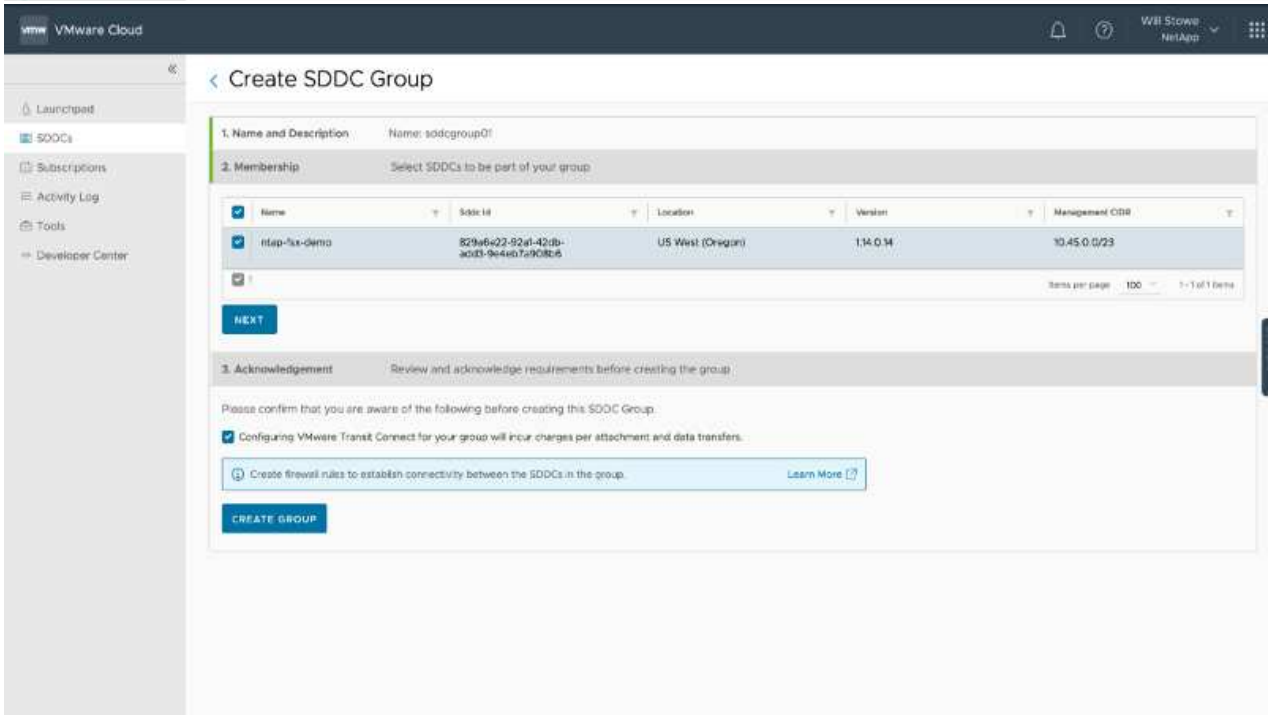
VMware Cloud
Will Stowe
NetApp
Launchpad
SDDCs
Subscriptions
Activity Log
Tools
Developer Center
Create SDDC Group
1. Name and Description
Name: sddcgroup01
2. Membership
Select SDDCs to be part of your group
Name
Sddc Id
Location
Version
Management CIDR
ntap-fsx-demo
829a6a22-92a1-42db-acd3-9e4eb7a908b6
US West (Oregon)
1.14.0.14
10.45.0.0/23
Items per page
100
1 - 1 of 1 items
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

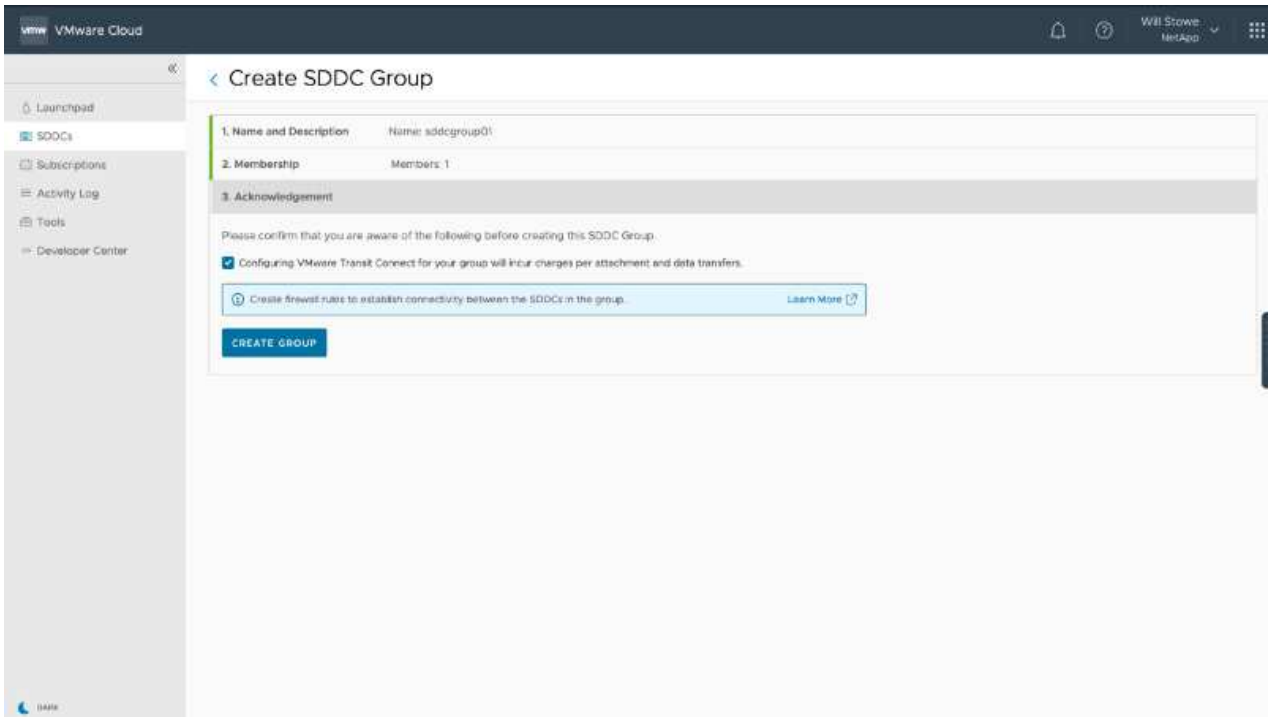

3. Attach the newly created VPC to the just created SDDC group. Select the External VPC tab and follow the instructions for attaching an External VPC to the group. This process can take 10 to 15 minutes to complete.

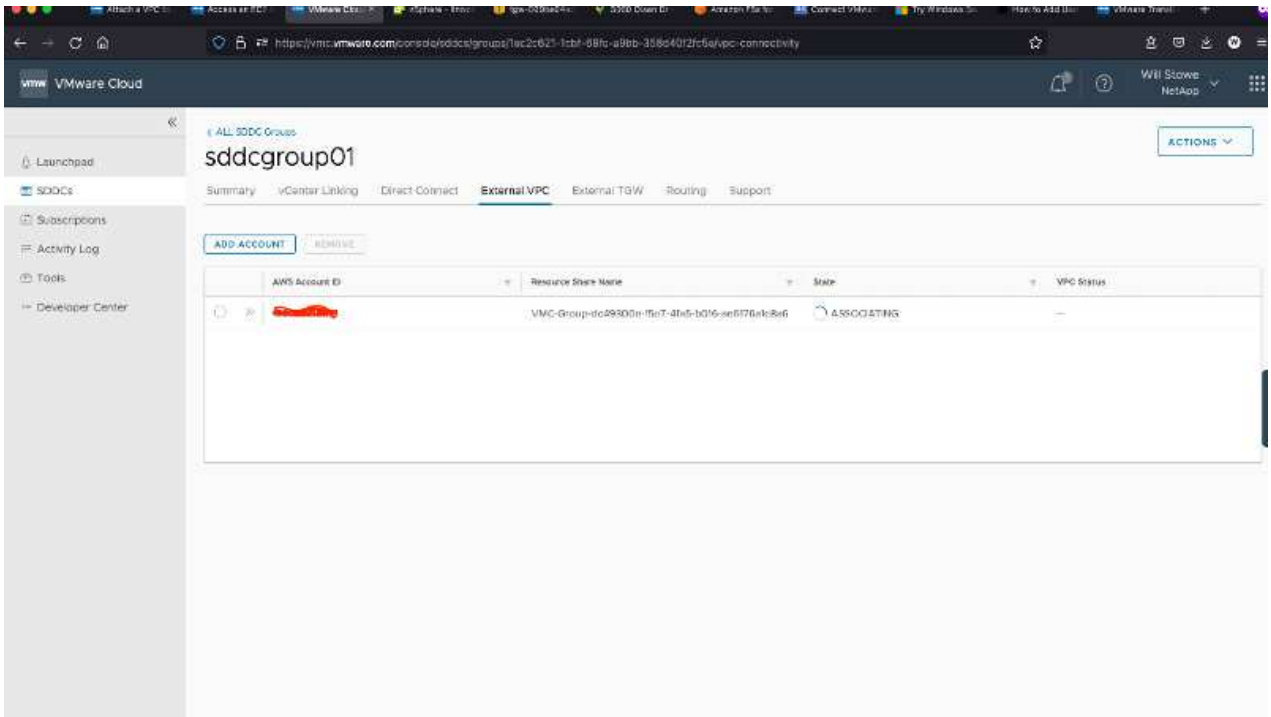

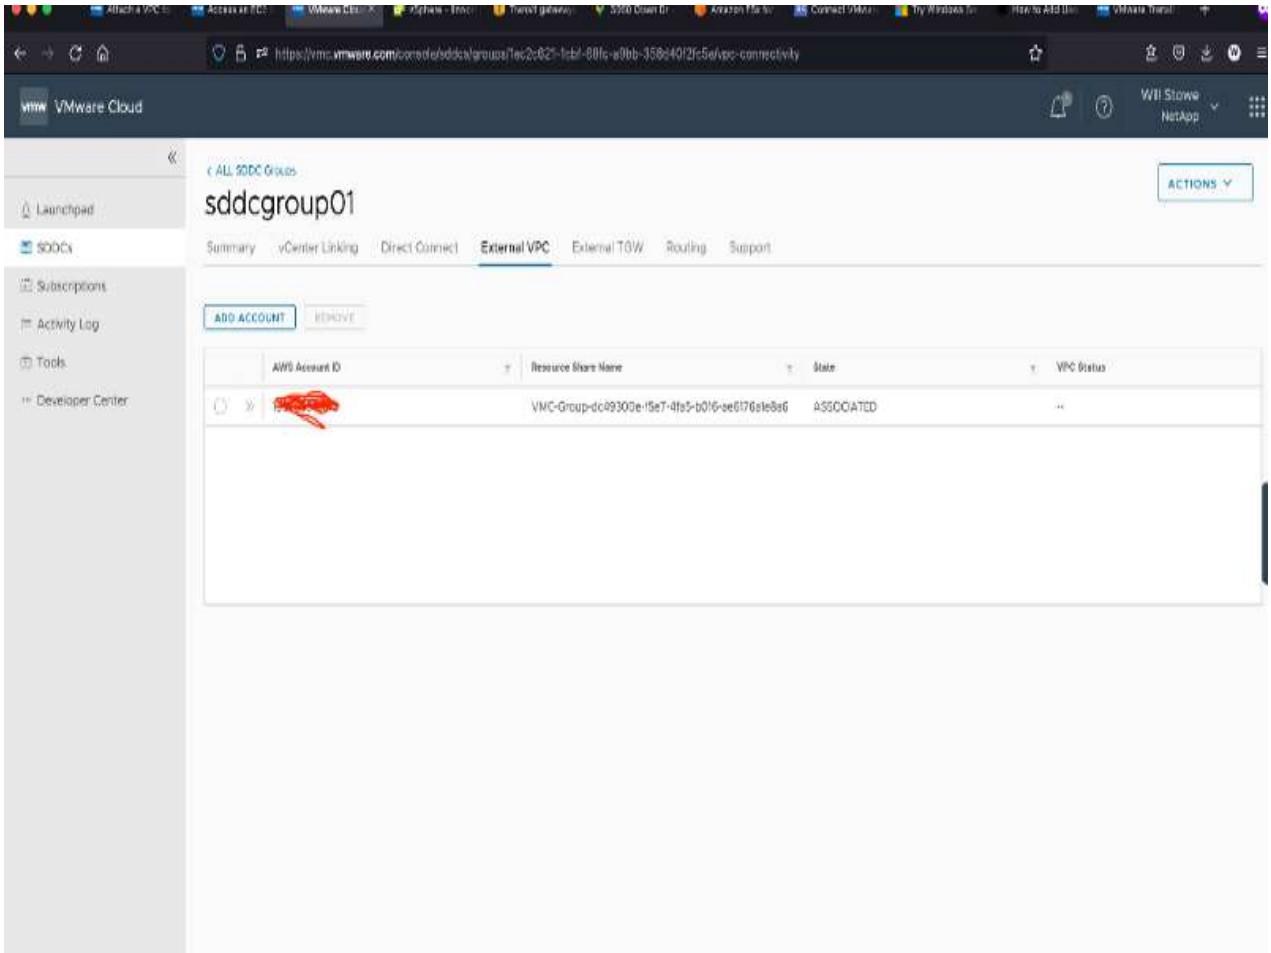

4. As part of the external VPC process, you are prompted through the AWS console to a new shared resource via the Resource Access Manager. The shared resource is the AWS Transit Gateway managed by VMware Transit Connect.

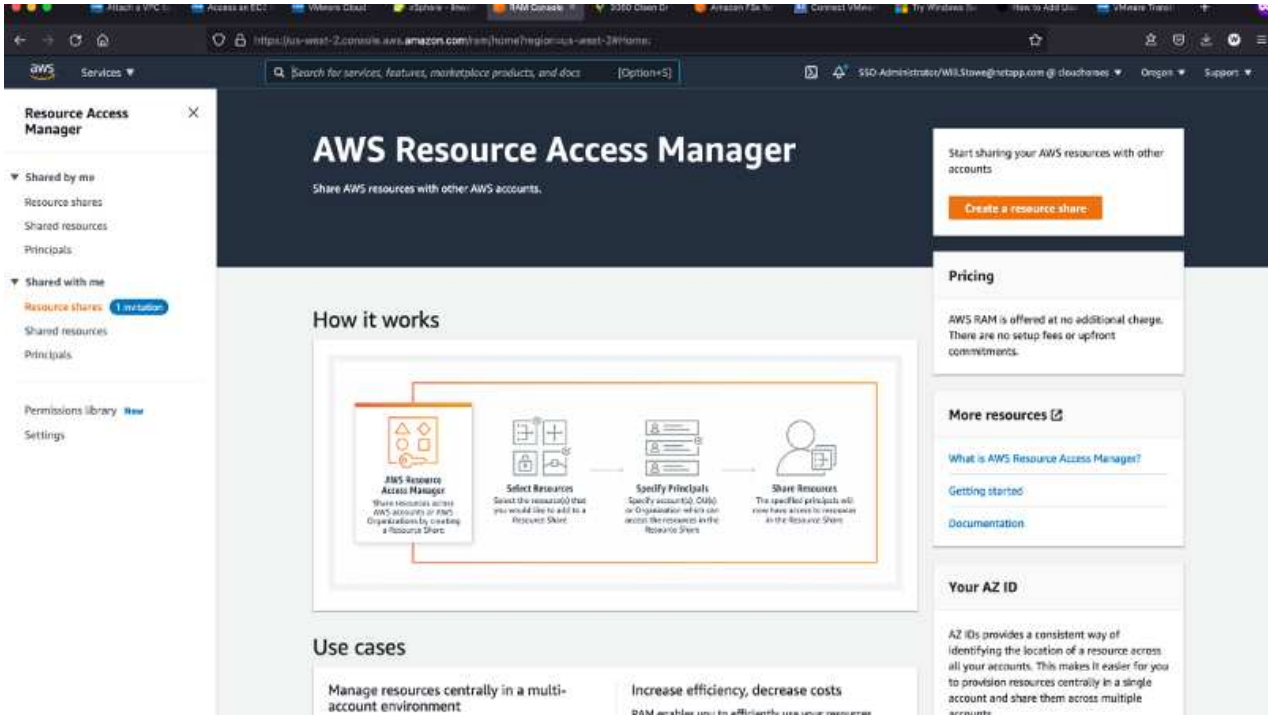

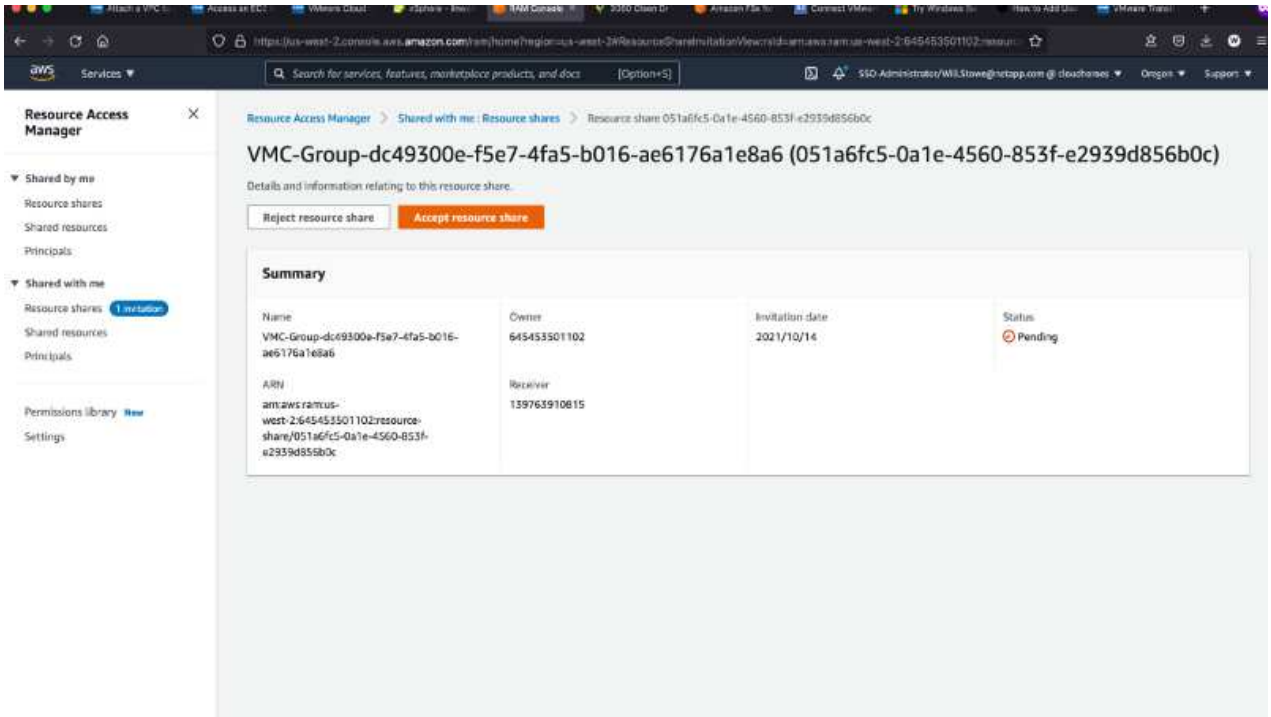

5. Create the Transit Gateway Attachment.

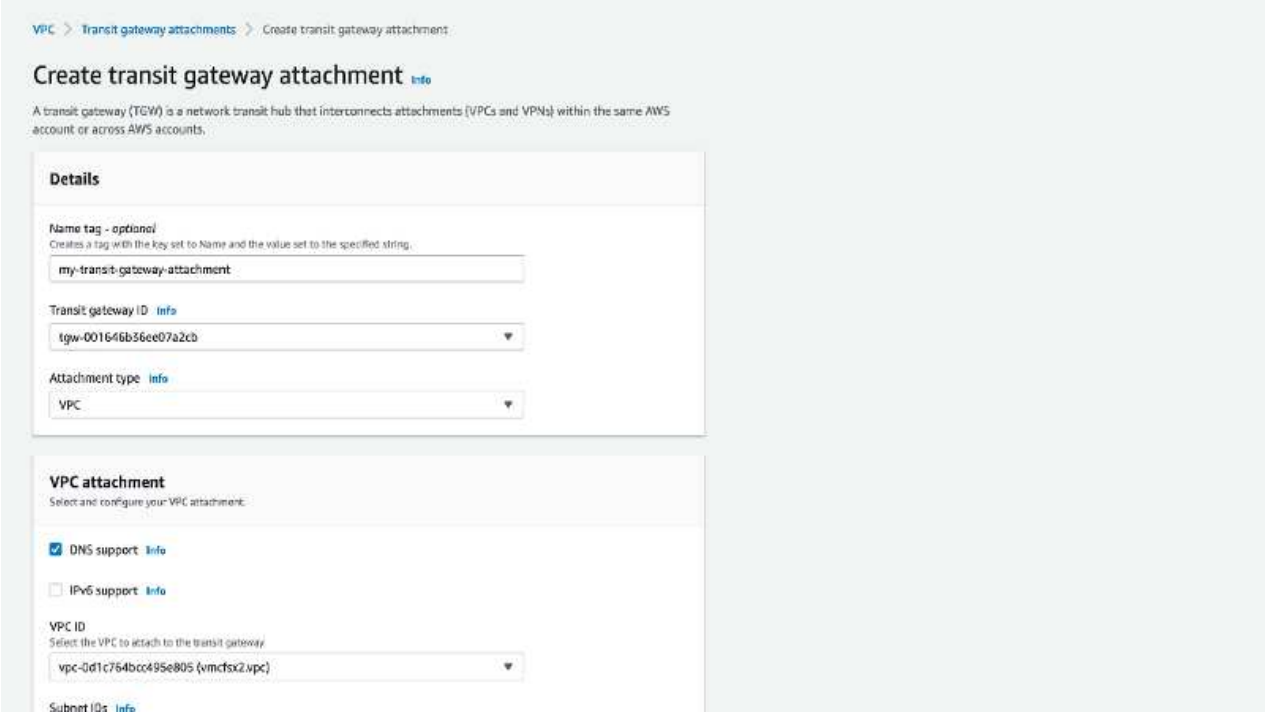

6. Back on the VMC Console, Accept the VPC attachment. This process can take approximately 10 minutes to complete.

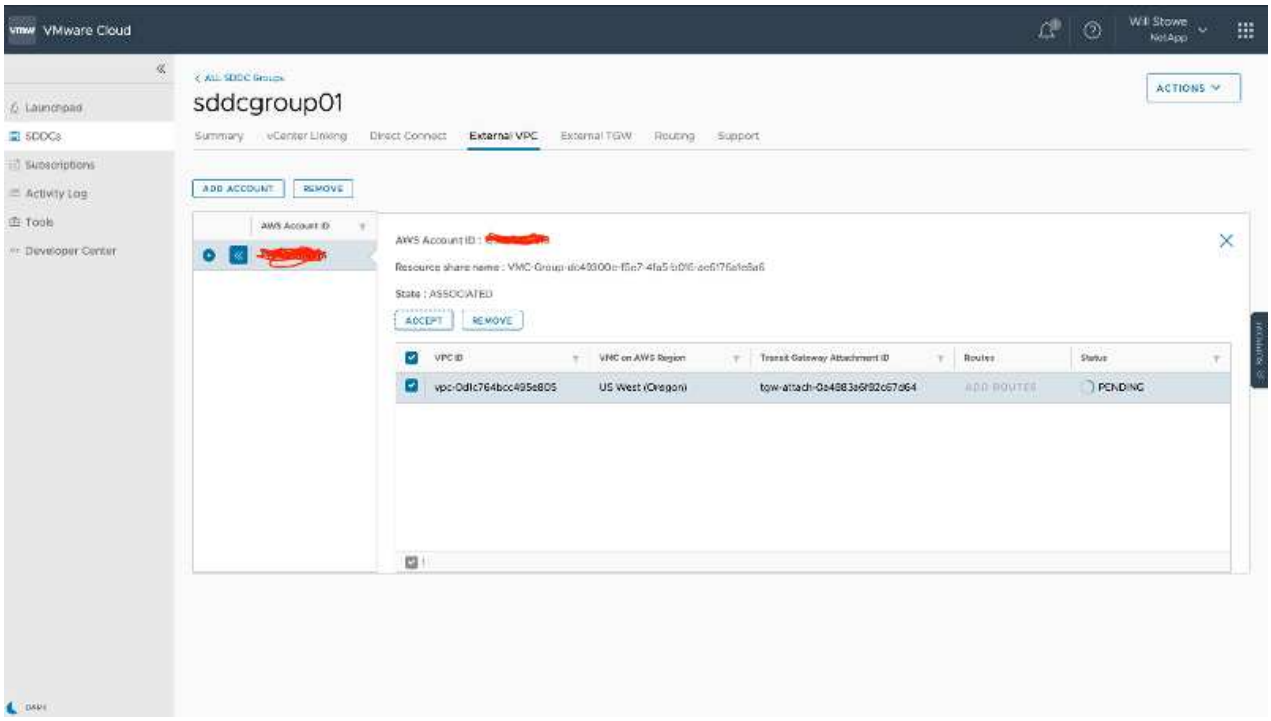

7. While in the External VPC tab, click the edit icon in the Routes column and add in the following required routes:
   - A route for the floating IP range for Amazon FSx for NetApp ONTAP floating IPs.
   - A route for the floating IP range for Cloud Volumes ONTAP (if applicable).
   - A route for the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

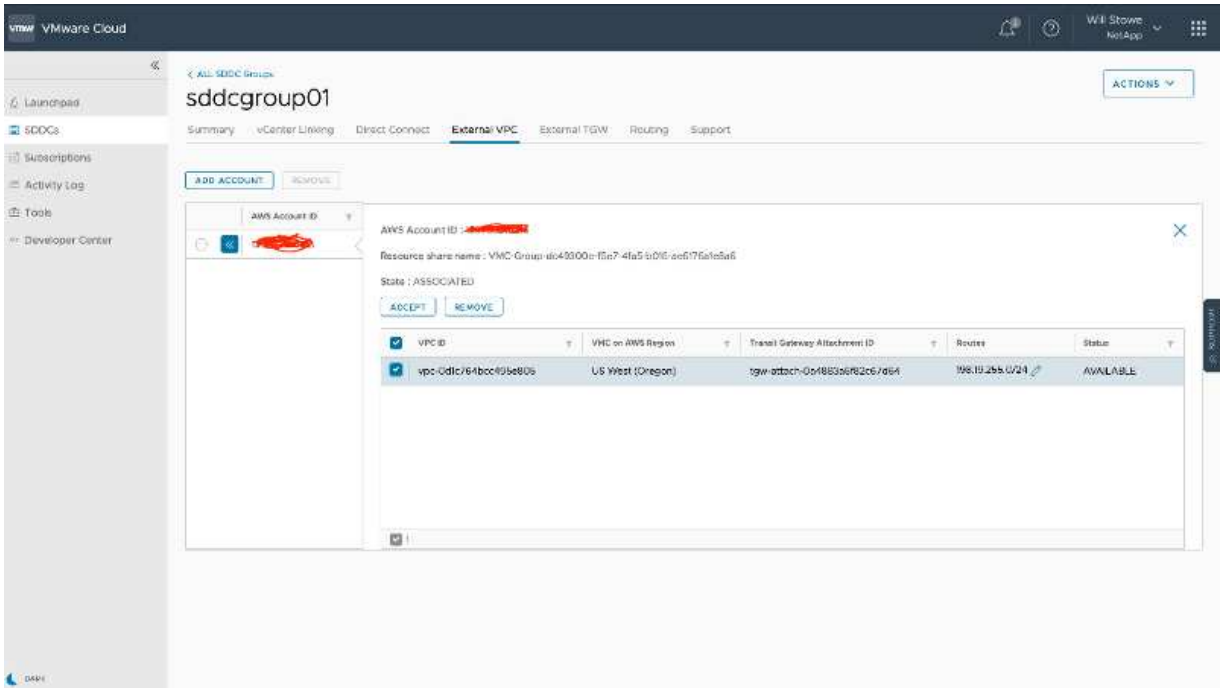

8. Finally, allow bidirectional traffic firewall rules for access to FSx/CVO. Follow these detailed steps for compute gateway firewall rules for SDDC workload connectivity.

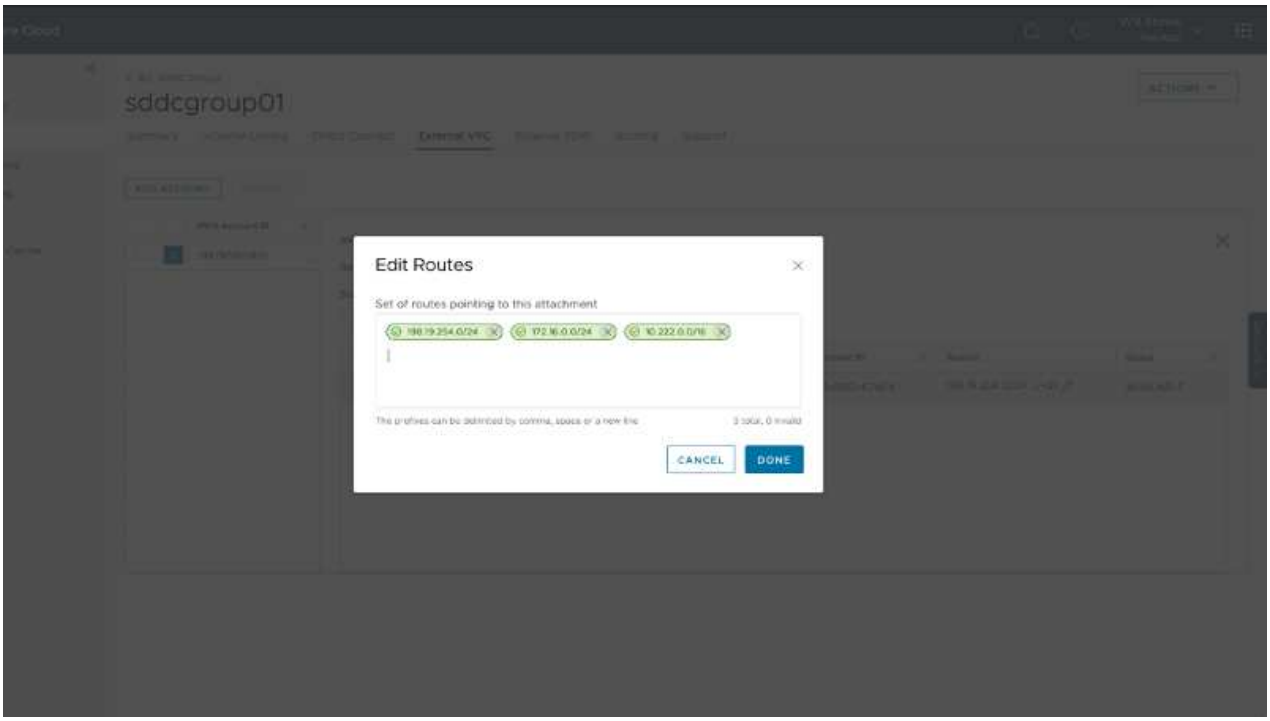

9. After the firewall groups are configured for both the Management and Compute gateway, the vCenter can be accessed as follows:

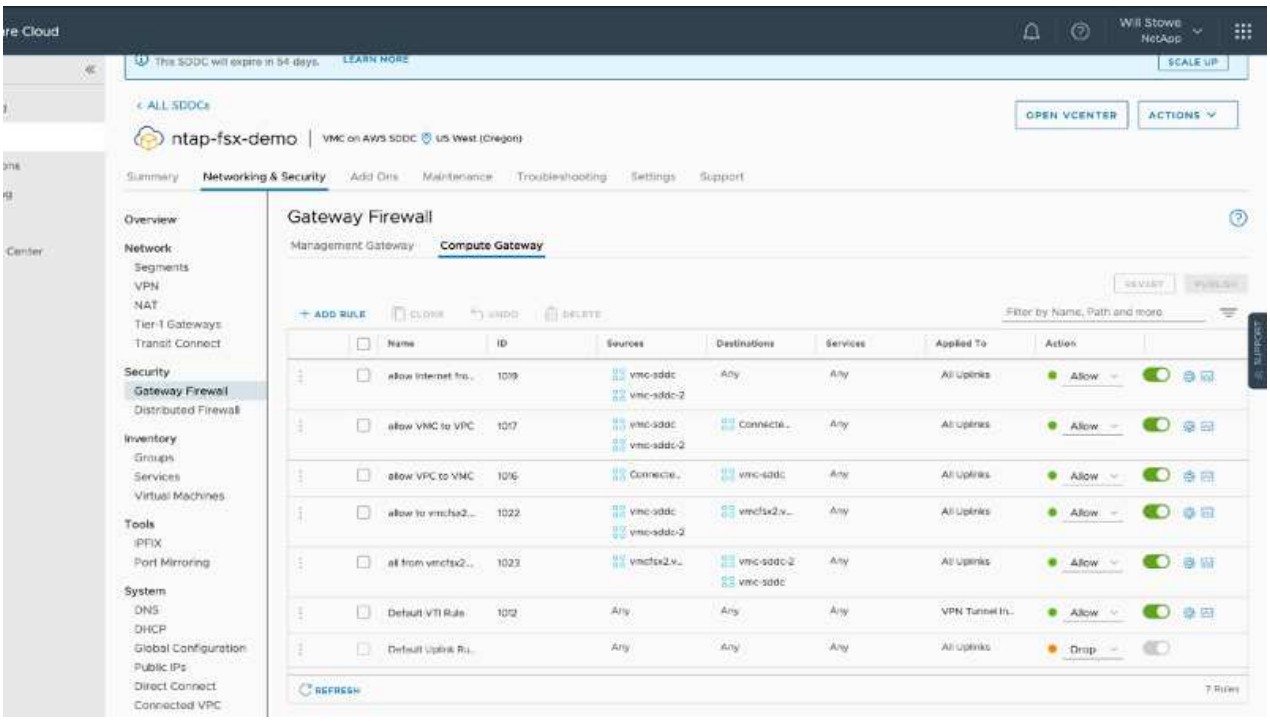

The next step is to verify that Amazon FSx ONTAP or Cloud Volumes ONTAP is configured depending on your requirements and that the volumes are provisioned to offload storage components from vSAN to optimize the deployment.

Deploy and configure the Virtualization Environment on Azure

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

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

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

Register the resource provider and create a private cloud

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

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

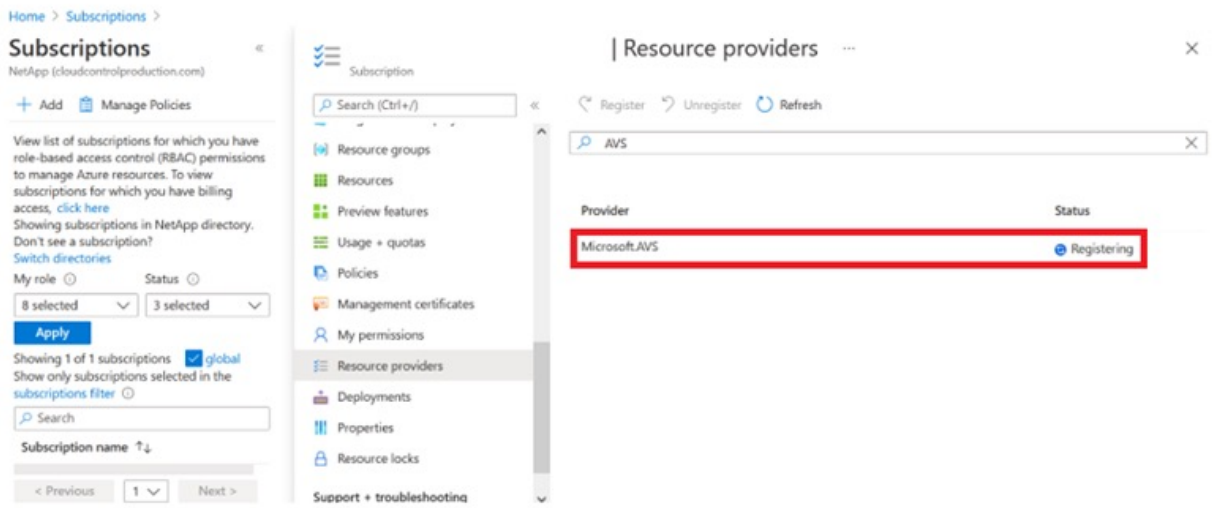

Provider	Status
Microsoft.OperationsManagement	Registered
Microsoft.Compute	Registered
Microsoft.ContainerService	Registered
Microsoft.ManagedIdentity	Registered
Microsoft.AVS	Registered
Microsoft.OperationalInsights	Registered
Microsoft.GuestConfiguration	Registered

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.

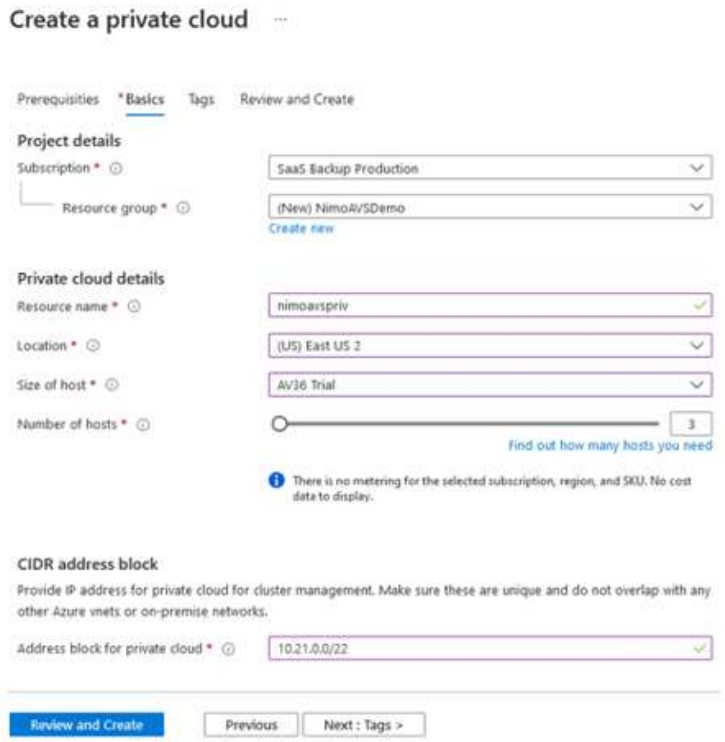

The provisioning process takes approximately 4–5 hours. After the process is complete, verify that the deployment was successful by accessing the private cloud from the Azure portal. A status of Succeeded is displayed when the deployment is complete.

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

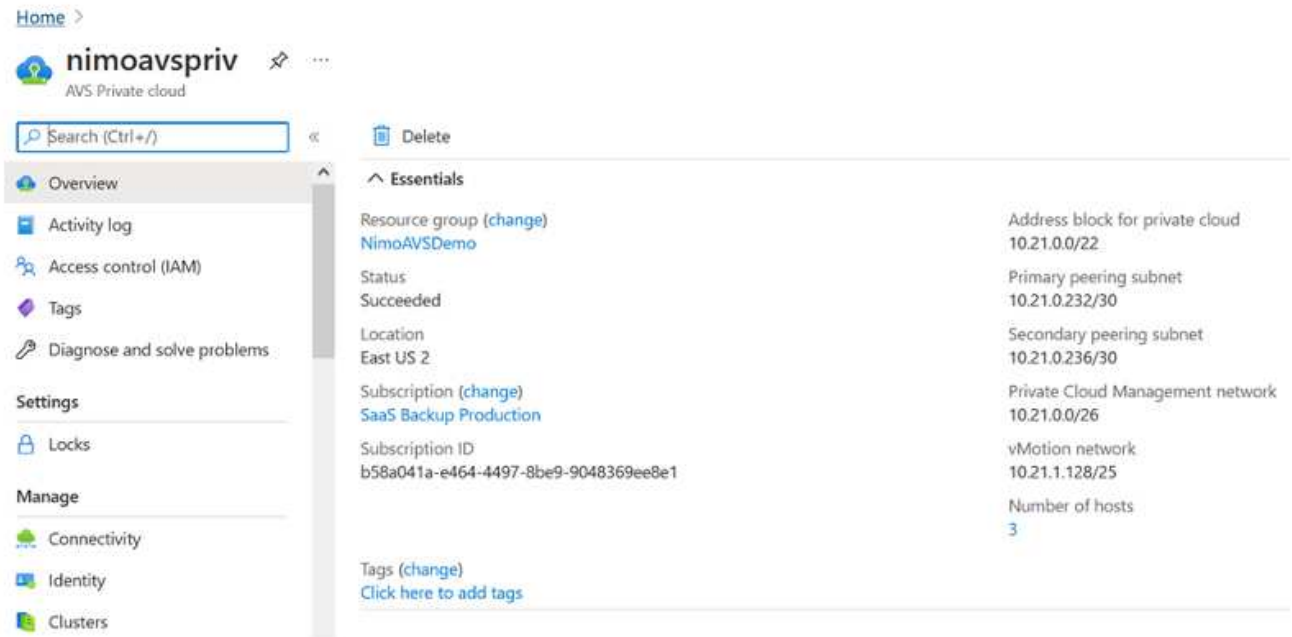

Connect to a new or existing ExpressRoute virtual network gateway

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

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

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

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

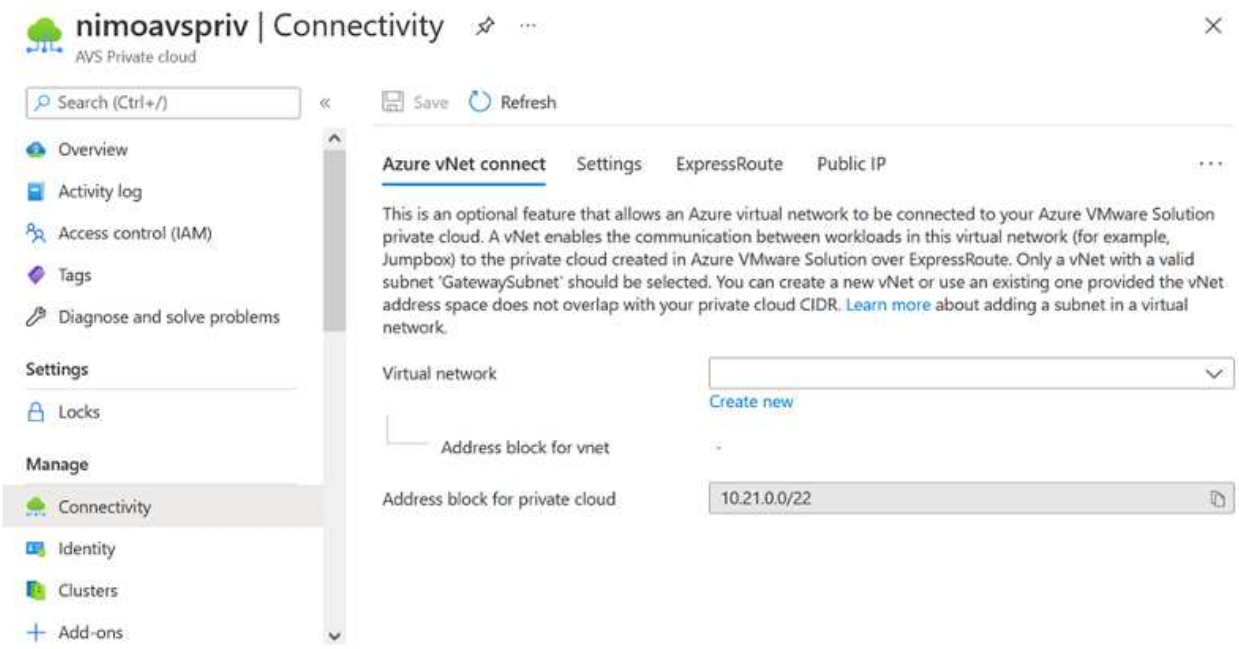

4. Provide or update the information for the new VNet and select OK.

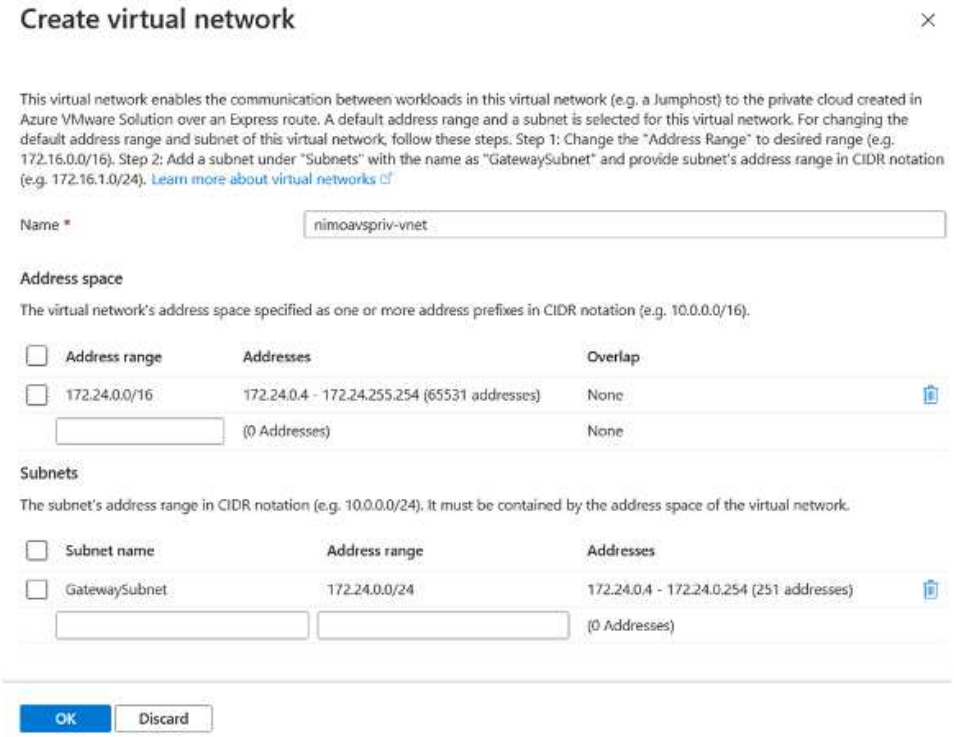

The VNet with the provided address range and gateway subnet is created in the designated subscription and resource group.

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

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

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

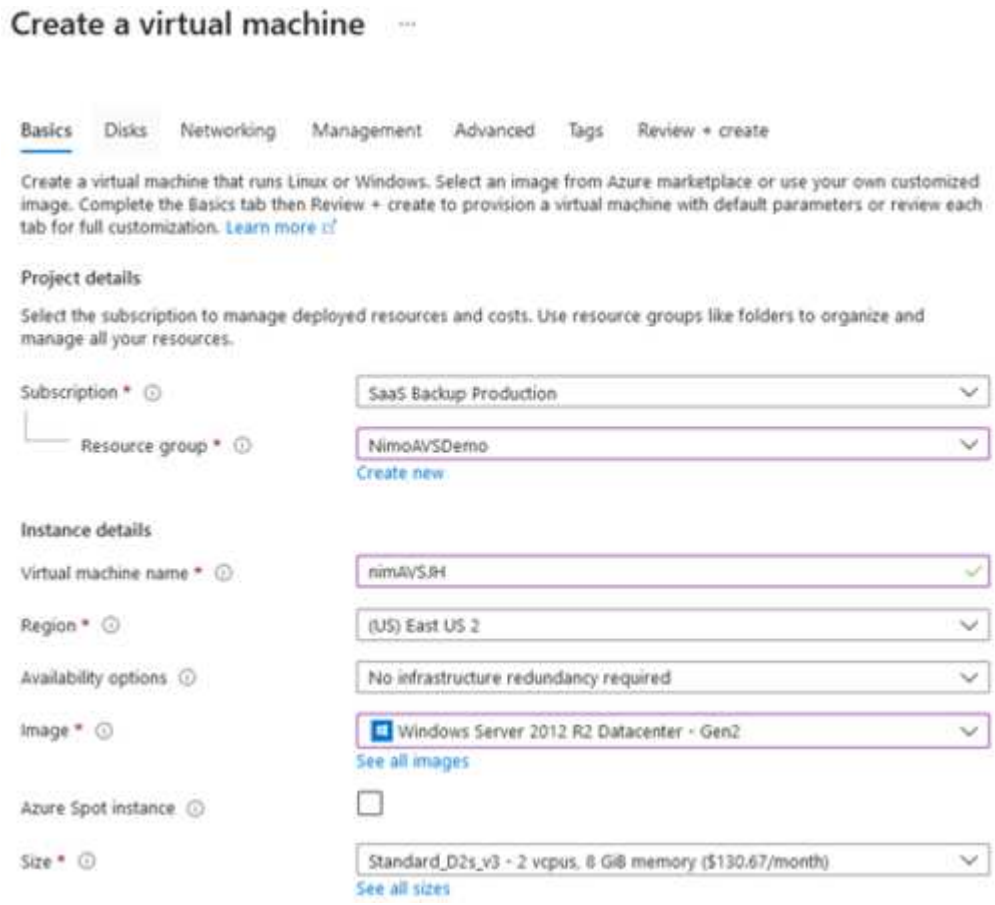

After the virtual machine is provisioned, use the Connect option to access RDP.

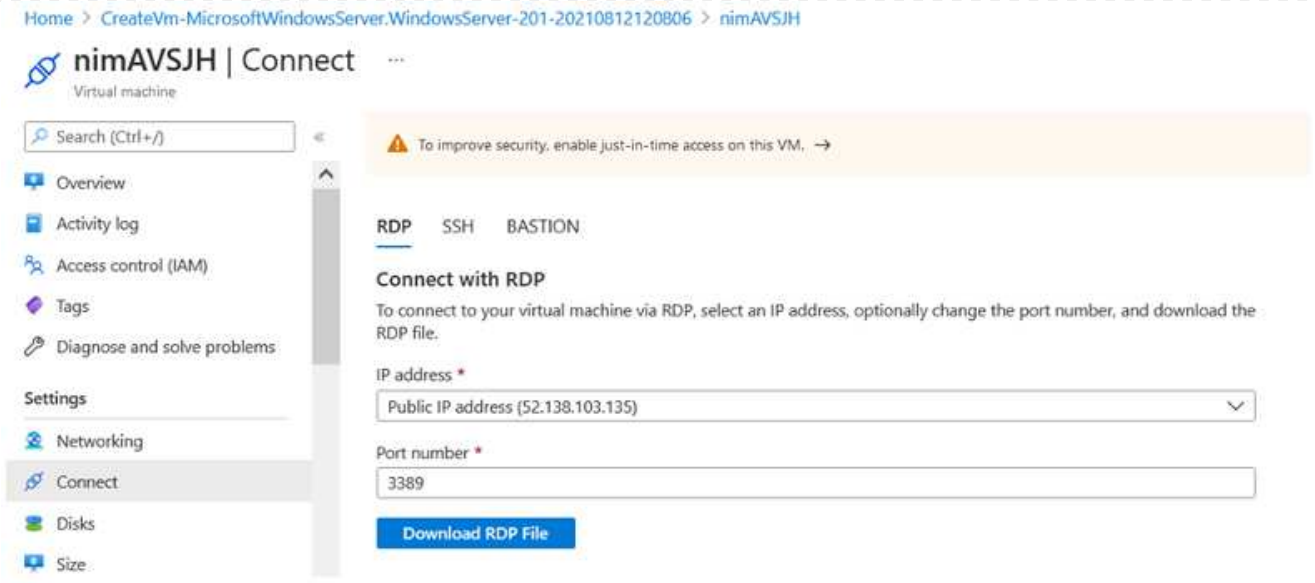

Sign in to vCenter from this newly created jump host virtual machine by using the cloud admin user . To access the credentials, go to the Azure portal and navigate to Identity (under the Manage option within the private cloud). The URLs and user credentials for the private cloud vCenter and NSX-T Manager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

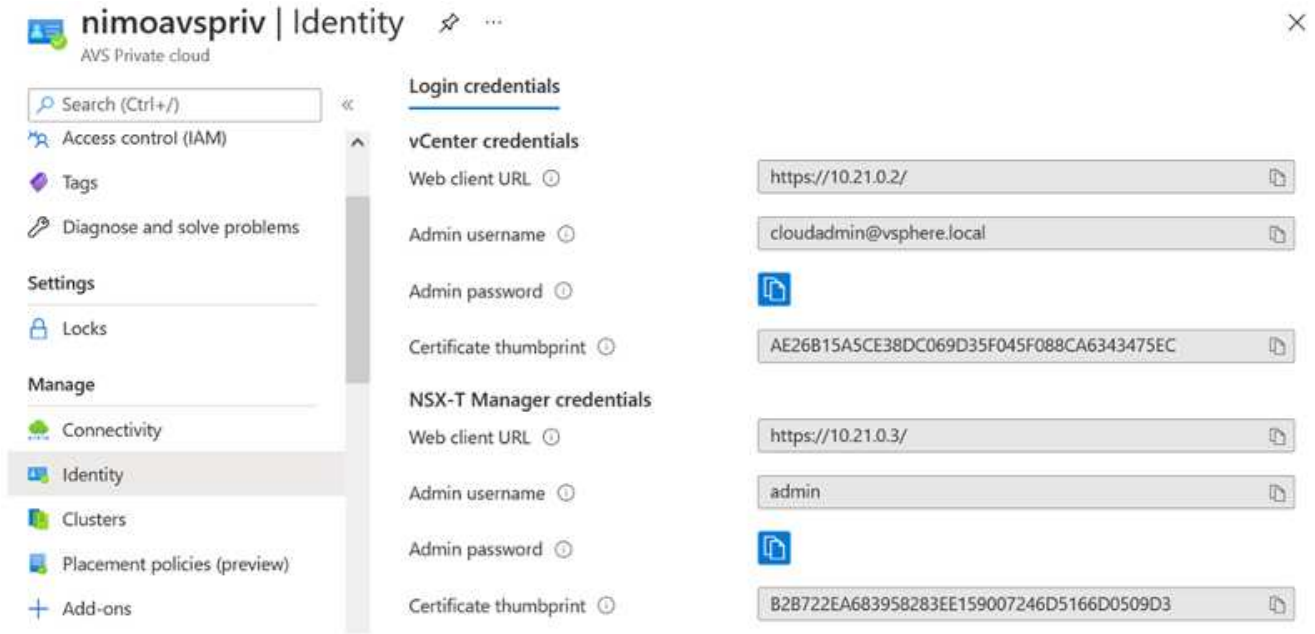

In the Windows virtual machine, open a browser and navigate to the vCenter web client URL (https://10.21.0.2/) and use the admin user name as **cloudadmin@vsphere.local** and paste the copied password. Similarly, NSX-T manager can also be accessed using the web client URL (https://10.21.0.3/) and use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

The web client URLs are different for each SDDC provisioned.

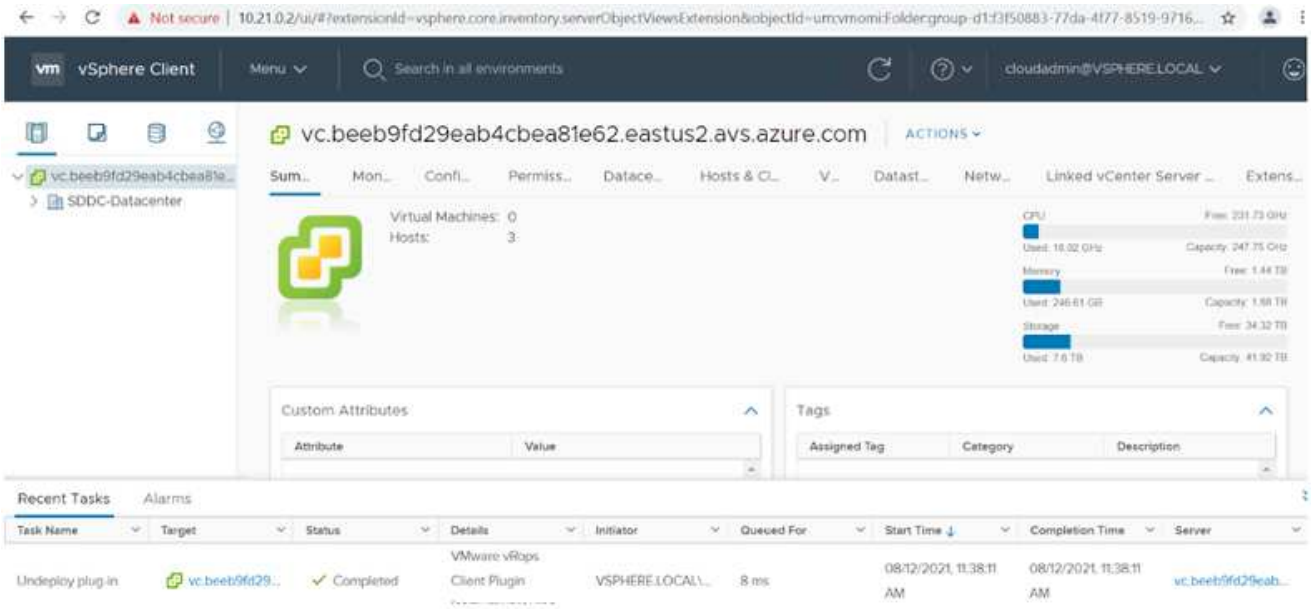

The Azure VMware Solution SDDC is now deployed and configured. Leverage ExpressRoute Global Reach to connect the on-premises environment to Azure VMware Solution private cloud. For more information, see Peer on-premises environments to Azure VMware Solution.

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

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

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

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

Deploy and configure GCVE

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

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

Pre-requisites:

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

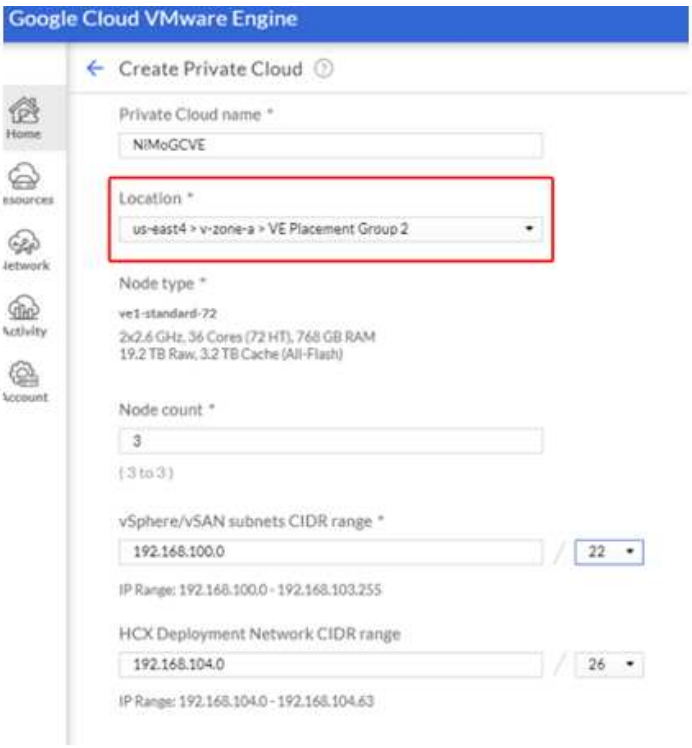

Note: Private cloud creation can take between 30 minutes to 2 hours.

Enable Private Access to GCVE

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

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

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

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

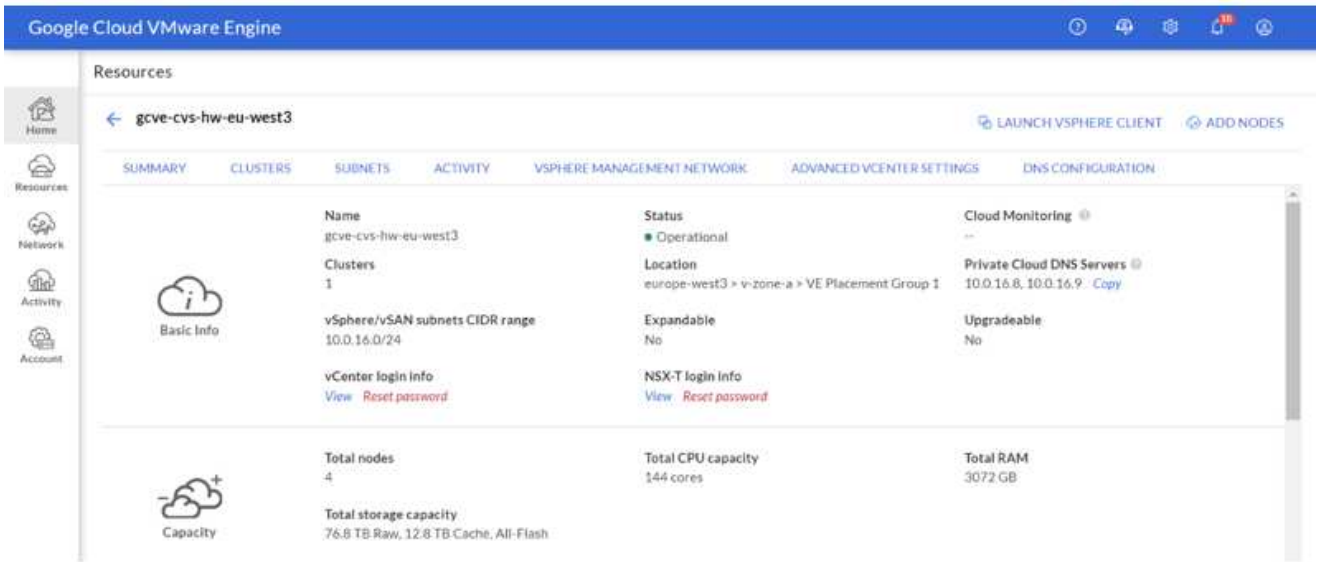

In a Windows virtual machine, open a browser and navigate to the vCenter web client URL (https://10.0.16.6/) and use the admin user name as CloudOwner@gve.local and paste the copied password. Similarly, NSX-T manager can also be accessed using the web client URL (https://10.0.16.11/) and use the admin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

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

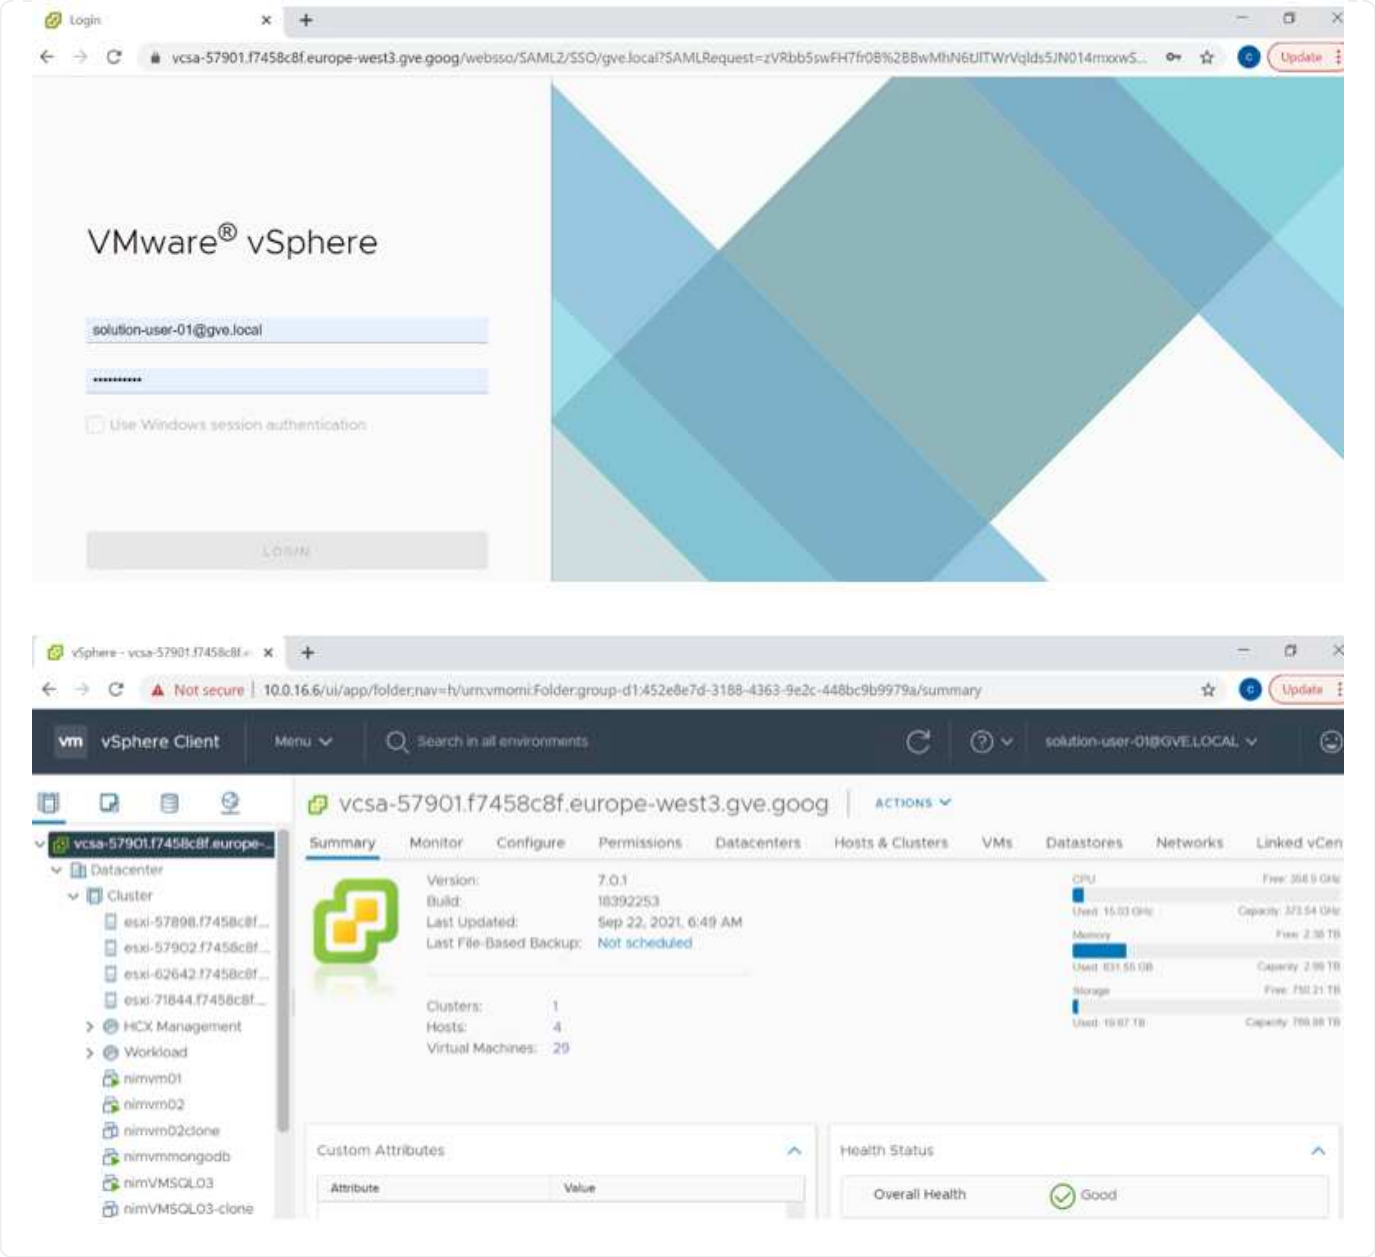

Deploy NetApp Cloud Volume Service supplemental datastore to GCVE

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

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
- Cloud Volumes Service (CVS) as guest connected storage
- Cloud Volumes Service (CVS) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.

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

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

Niyaz Mohamed, NetApp

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 for 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 for ONTAP as a NFS datastore for VMware Cloud on AWS.

> This solution is also available from VMware. Please visit the VMware Cloud Tech Zone for more information.

Connectivity options

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

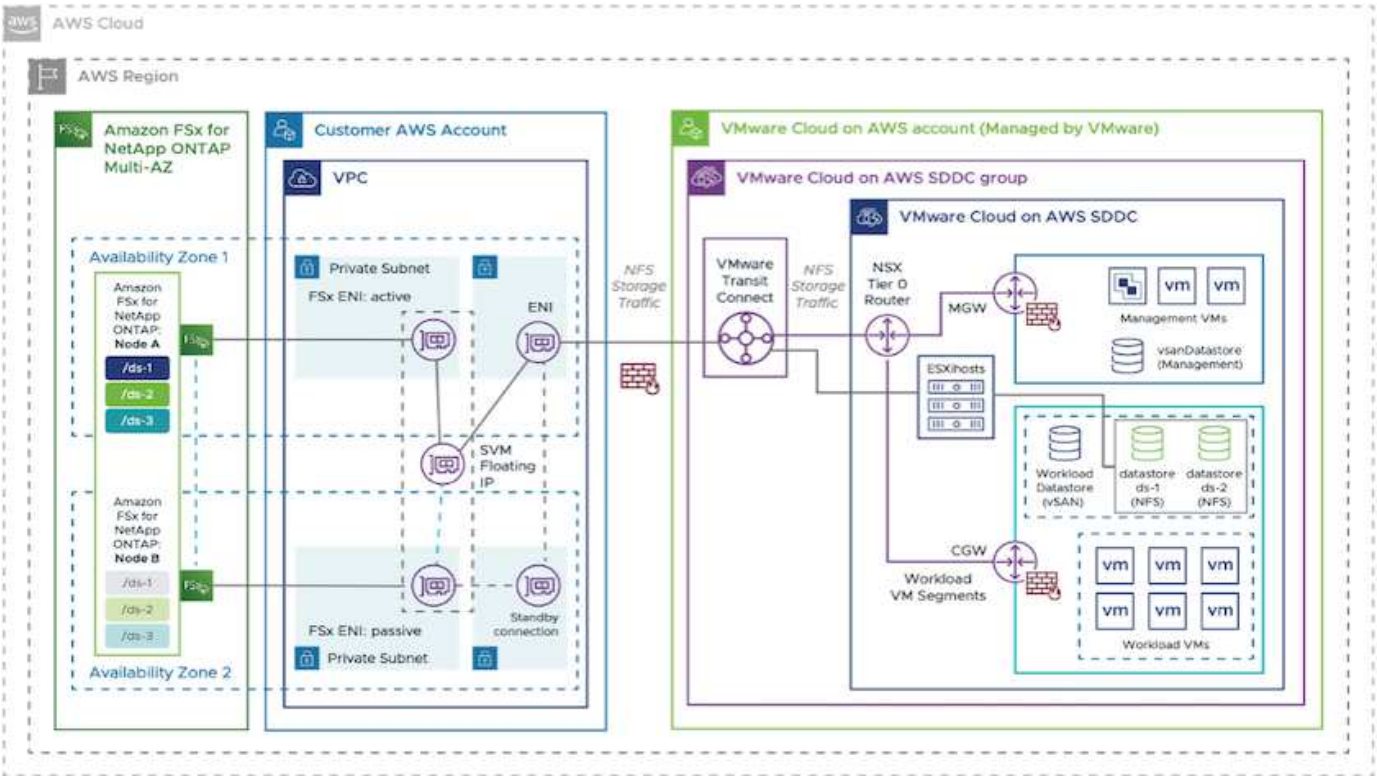

The high-level deployment steps are as follows:

1. Create Amazon FSx for 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 for 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 for ONTAP is not currently supported with stretched clusters.

Conclusion

This document covers the steps necessary to configure Amazon FSx for ONTAP with VMware cloud on AWS. Amazon FSx for 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 for 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 for ONTAP helps to optimize the overall deployment.

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

Amazon FSX for Ontap VMware Cloud

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 for NetApp 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 for 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 for 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 for ONTAP offers high-performance solid state drive (SSD) storage with submillisecond latencies. With FSx for 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 for ONTAP is easier because you can snapshot, clone, and replicate your files with the click of a button. In addition, FSx for ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

FSx for 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 for ONTAP supports popular data security and antivirus applications.

Configure Amazon FSx for NetApp ONTAP with VMware Cloud on AWS

Amazon FSx for NetApp 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 for NetApp 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 for ONTAP volumes

To create and mount Amazon FSx for NetApp 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 for NetApp ONTAP, and then choose Next. The Create File System page appears.

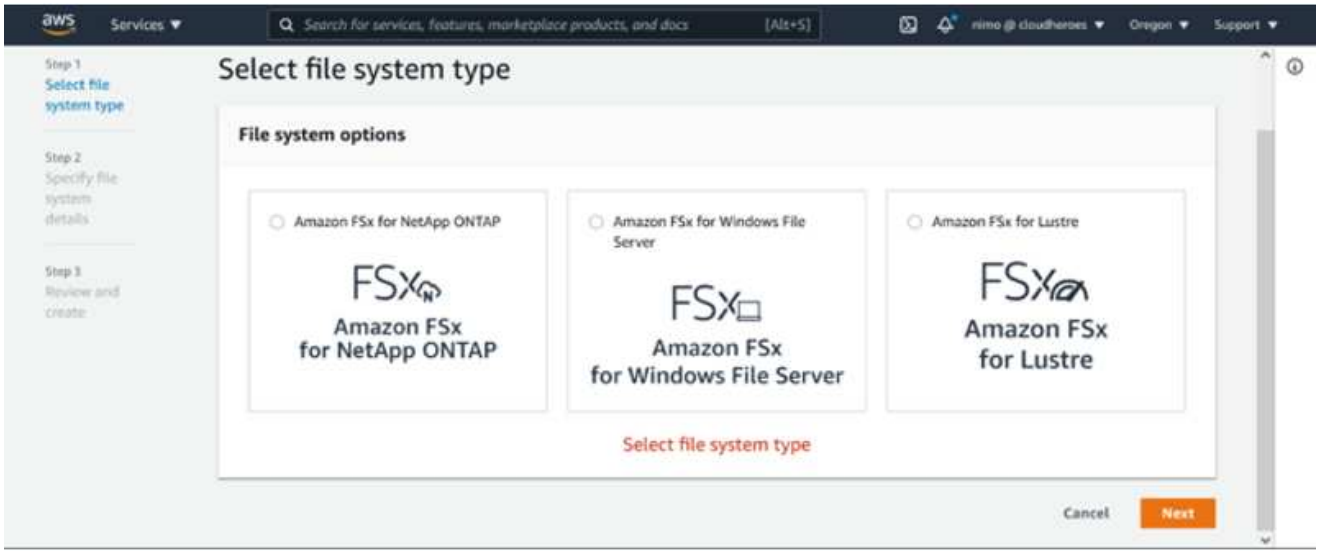

1. In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

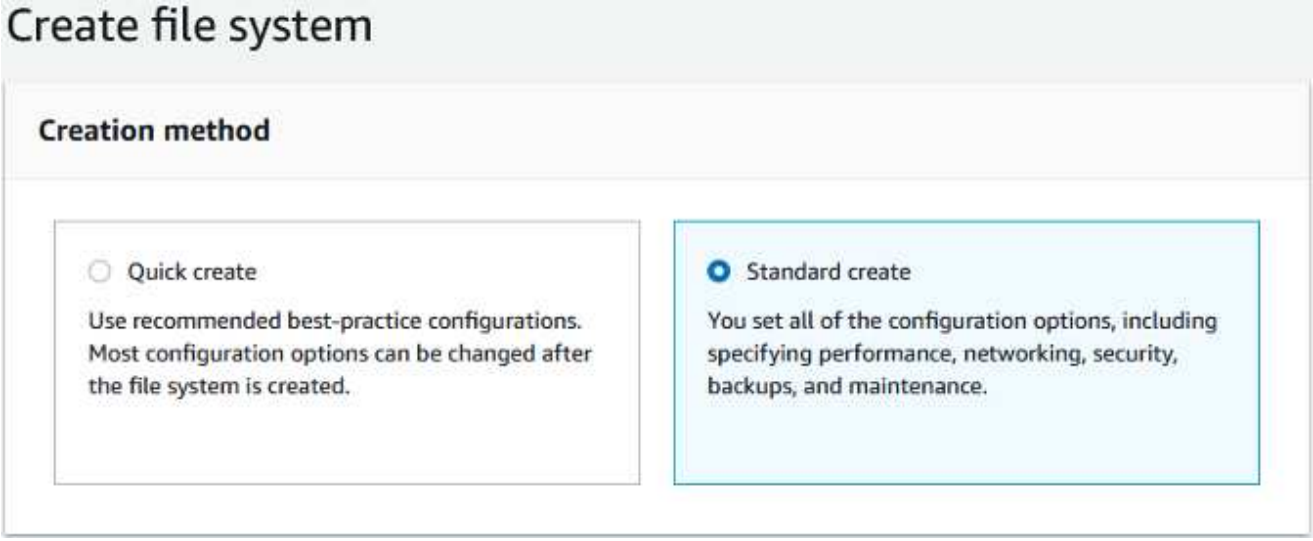

1. For the creation method, choose Standard Create. You can also choose Quick Create, but this document uses the Standard create option.

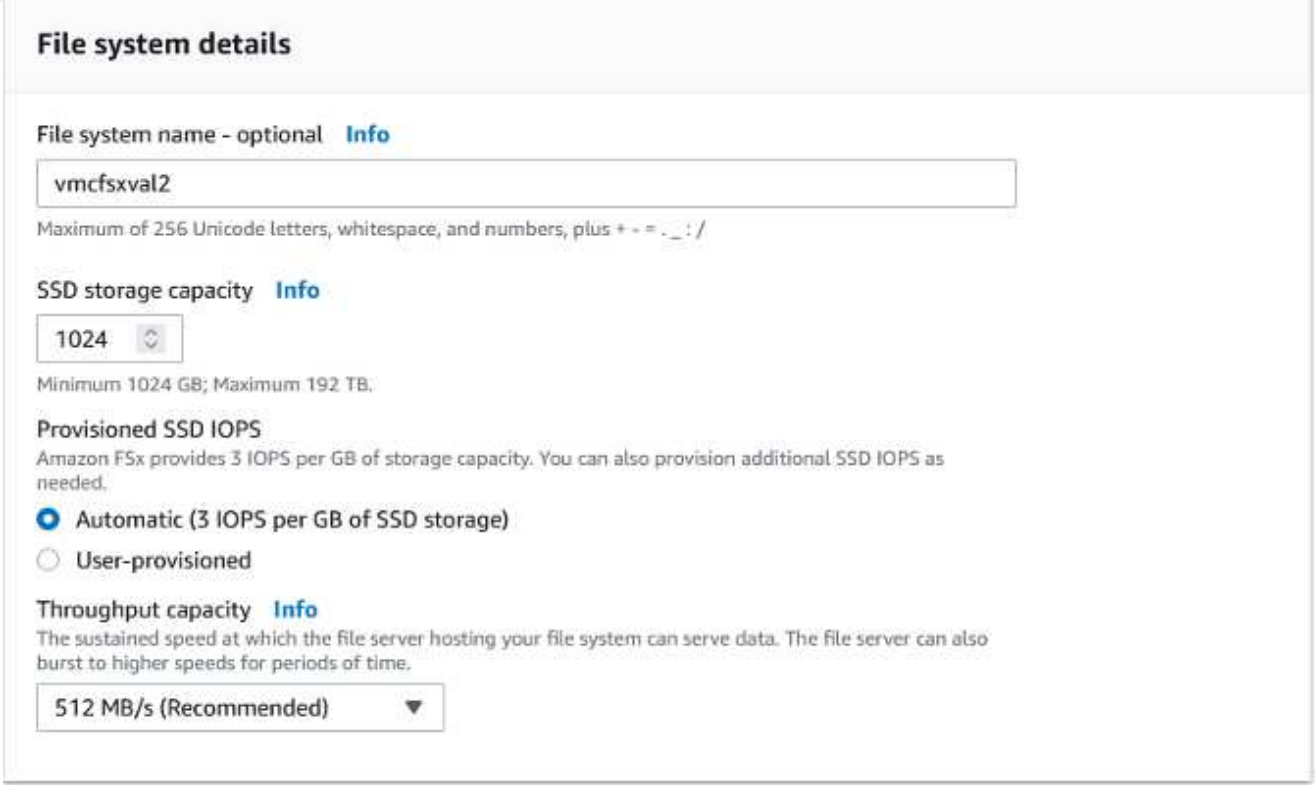

1. In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

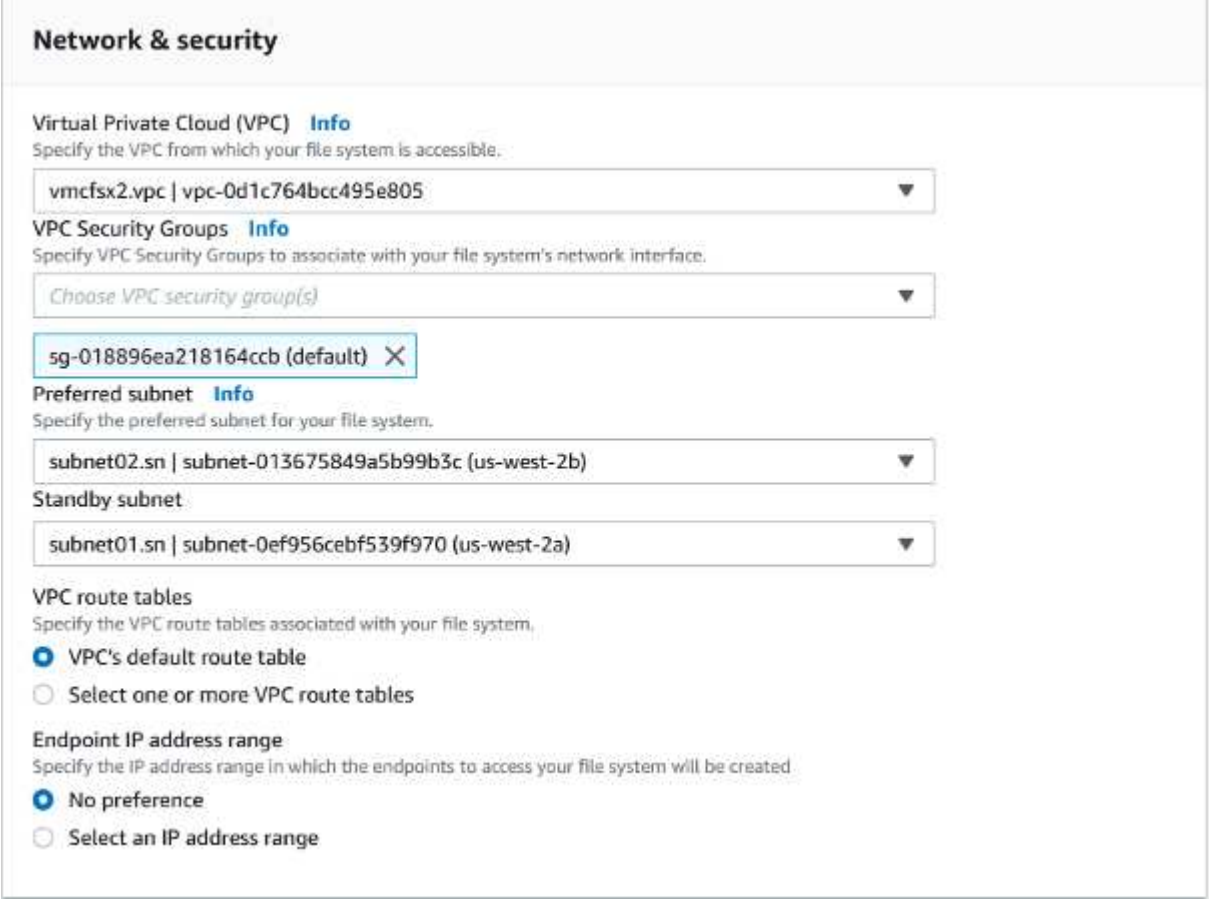

In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

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

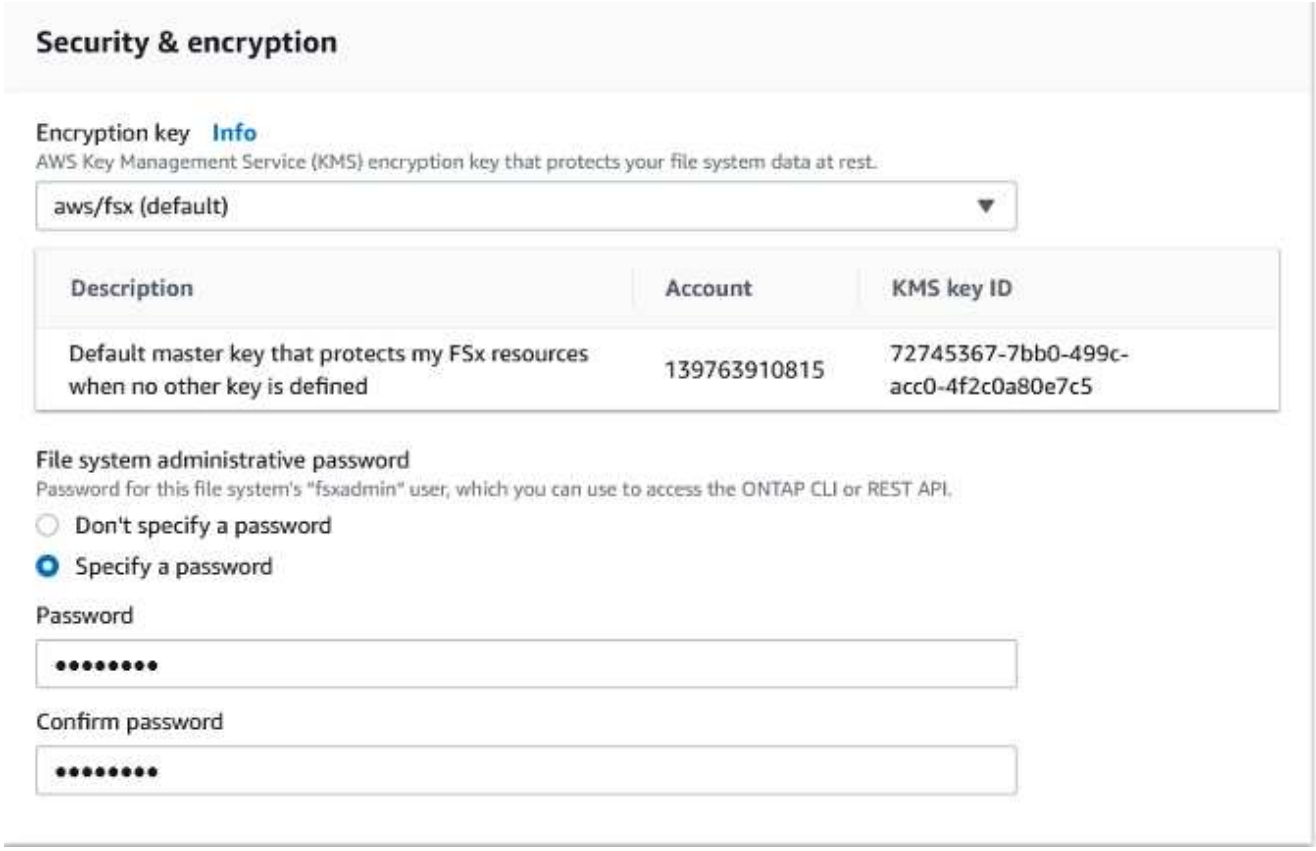

1. In virtual machine and specify the password to use with vsadmin for administering ONTAP using REST APIs or the CLI. If no password is specified, a fsxadmin user can be used for administering the SVM. In the Active Directory section, make sure to join Active Directory to the SVM for provisioning SMB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

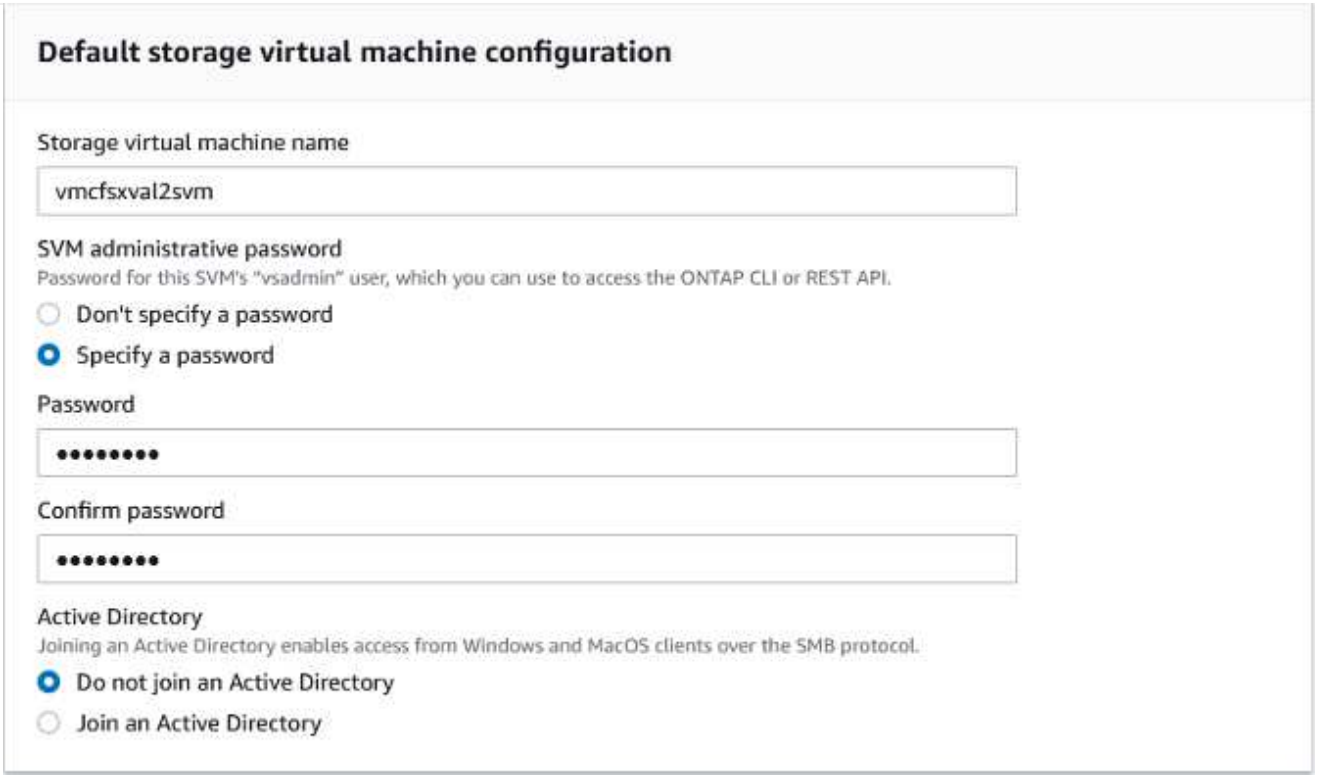

1. In the Default Volume Configuration section, specify the volume name and size. This is an NFS volume. For Storage Efficiency, choose Enabled to turn on the ONTAP storage efficiency features (compression, deduplication, and compaction) or Disabled to turn them off.

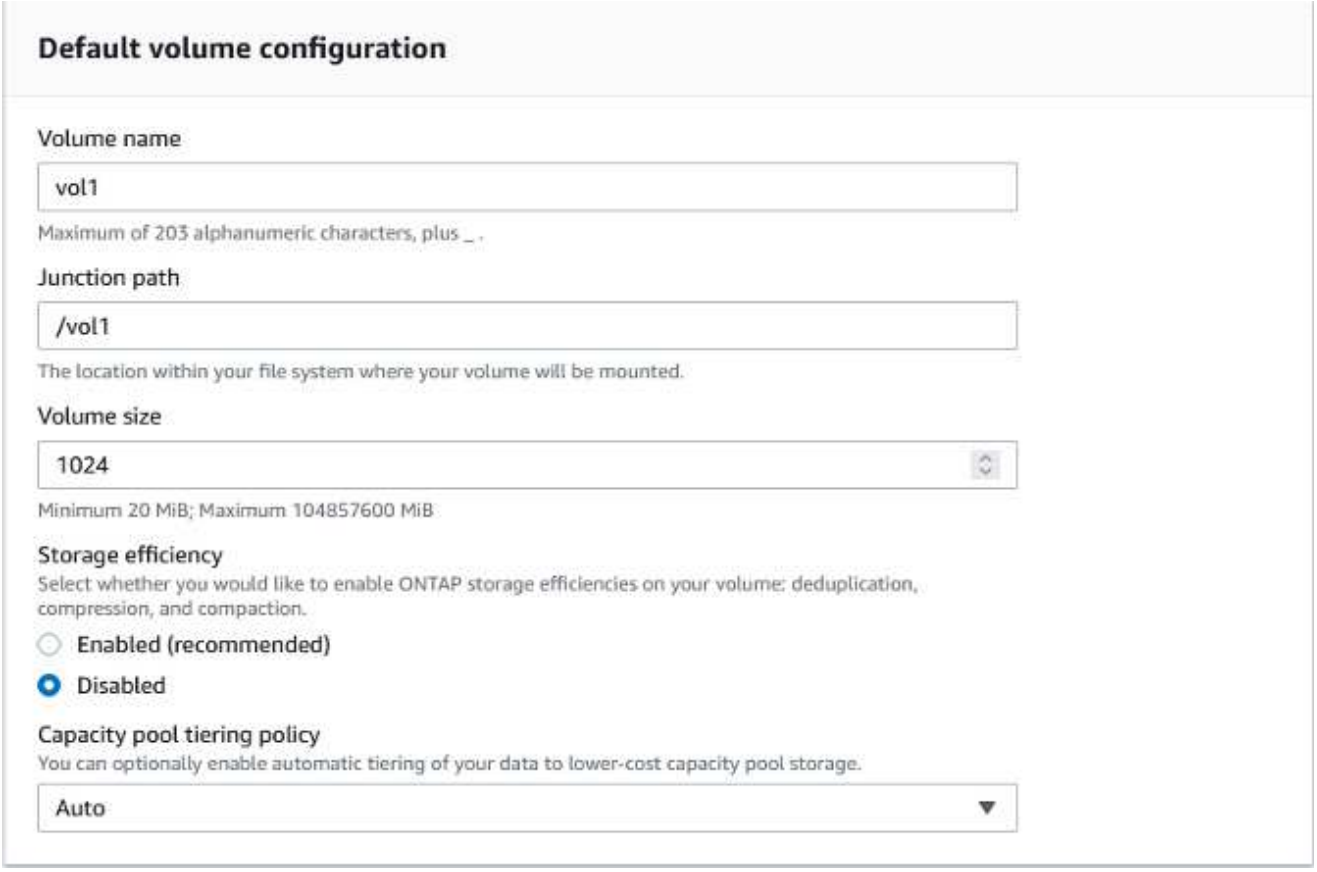

1. Review the file system configuration shown on the Create File System page.
2. Click Create File System.

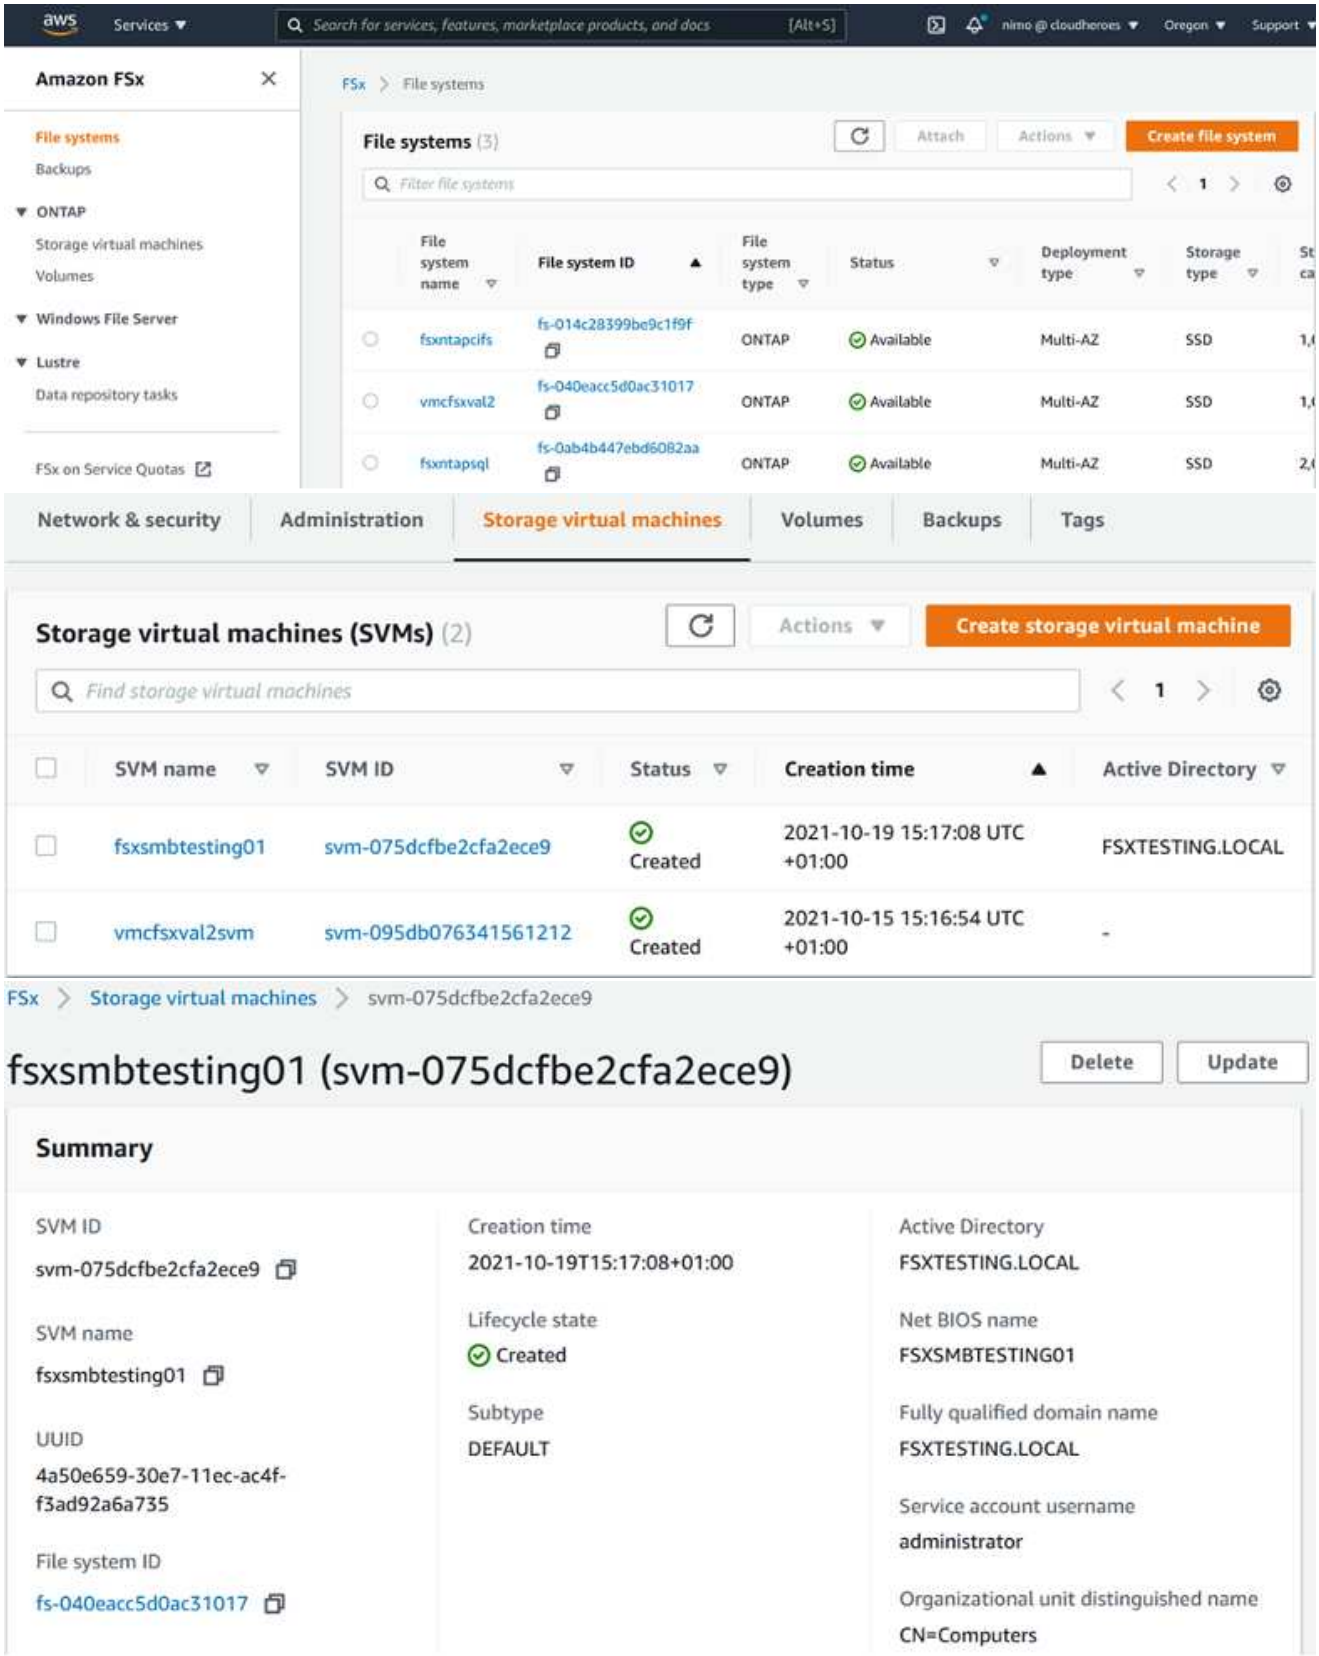

For more detailed information, see Getting started with Amazon FSx for NetApp ONTAP.

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

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

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

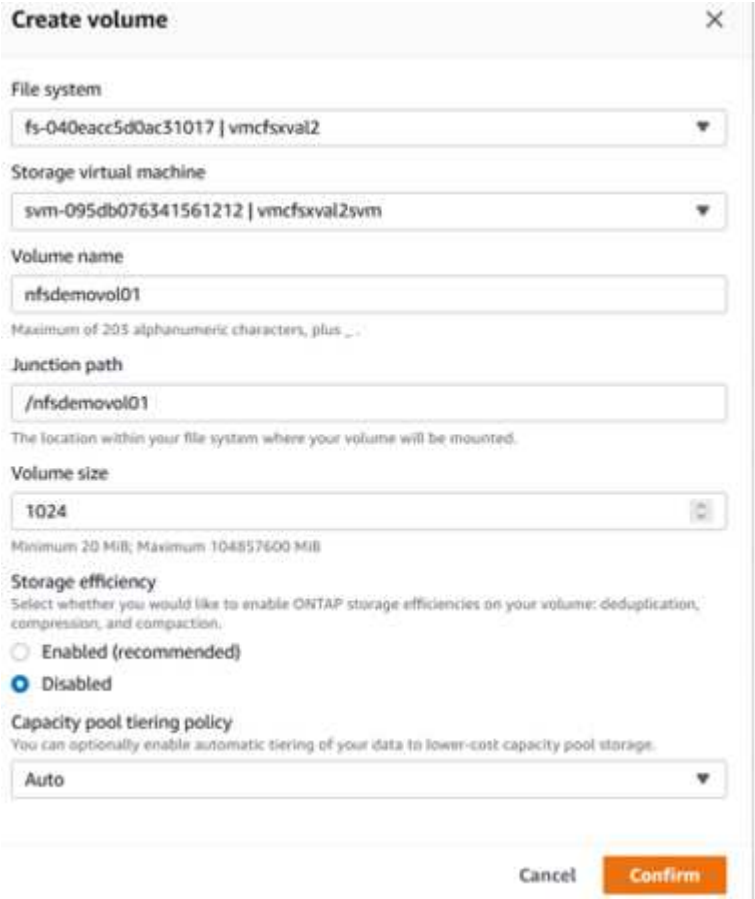

Mount FSx ONTAP volume on Linux client

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

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

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

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

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

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

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

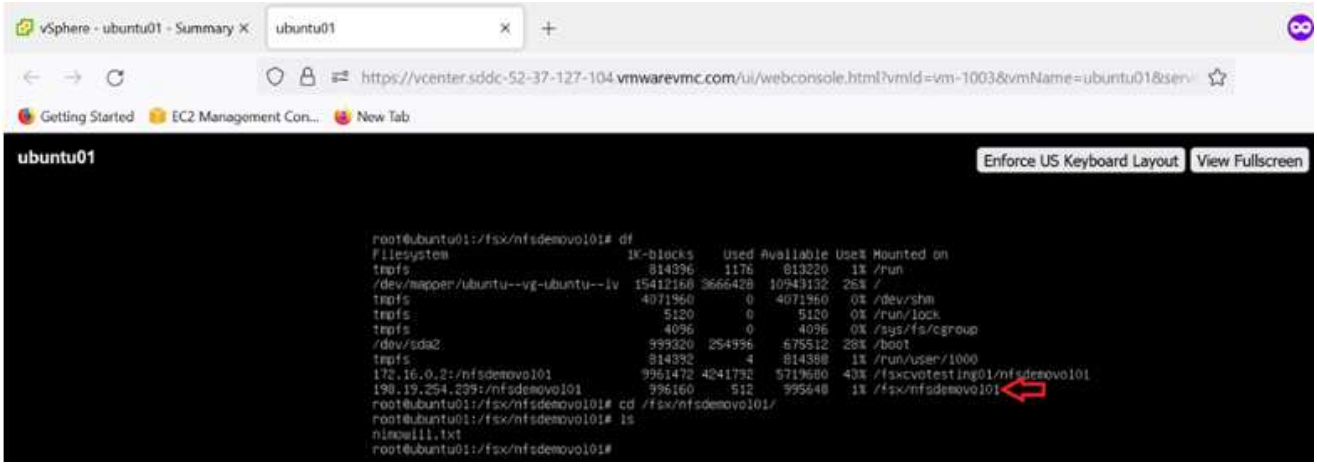

Mount FSx ONTAP volume on Linux client

Attach FSx ONTAP volumes to Microsoft Windows clients

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

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

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

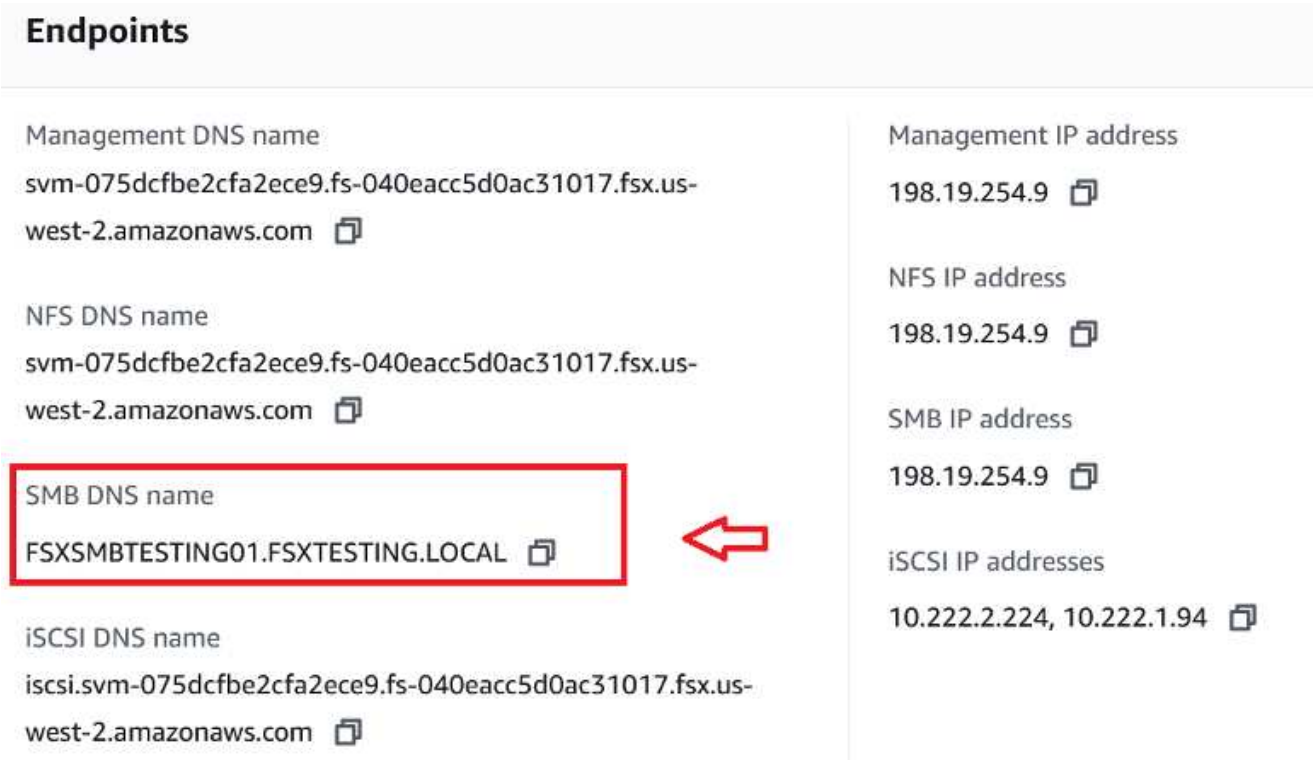

1. In the Shared Folders tool, choose Shares in the left pane to see the active shares for the Amazon FSx file system.

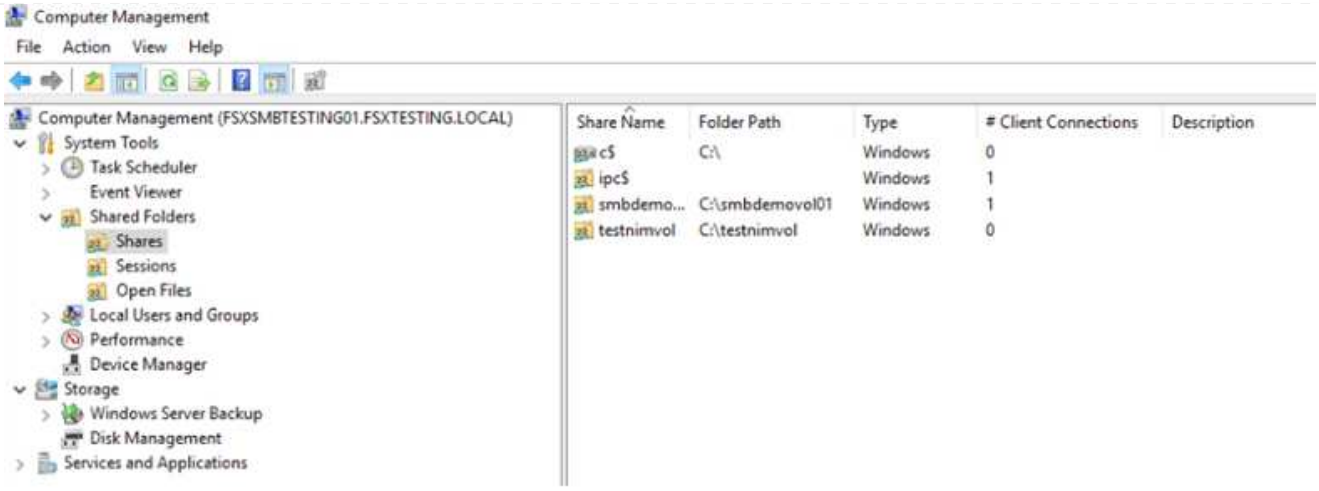

1. Now choose a new share and complete the Create a Shared Folder wizard.

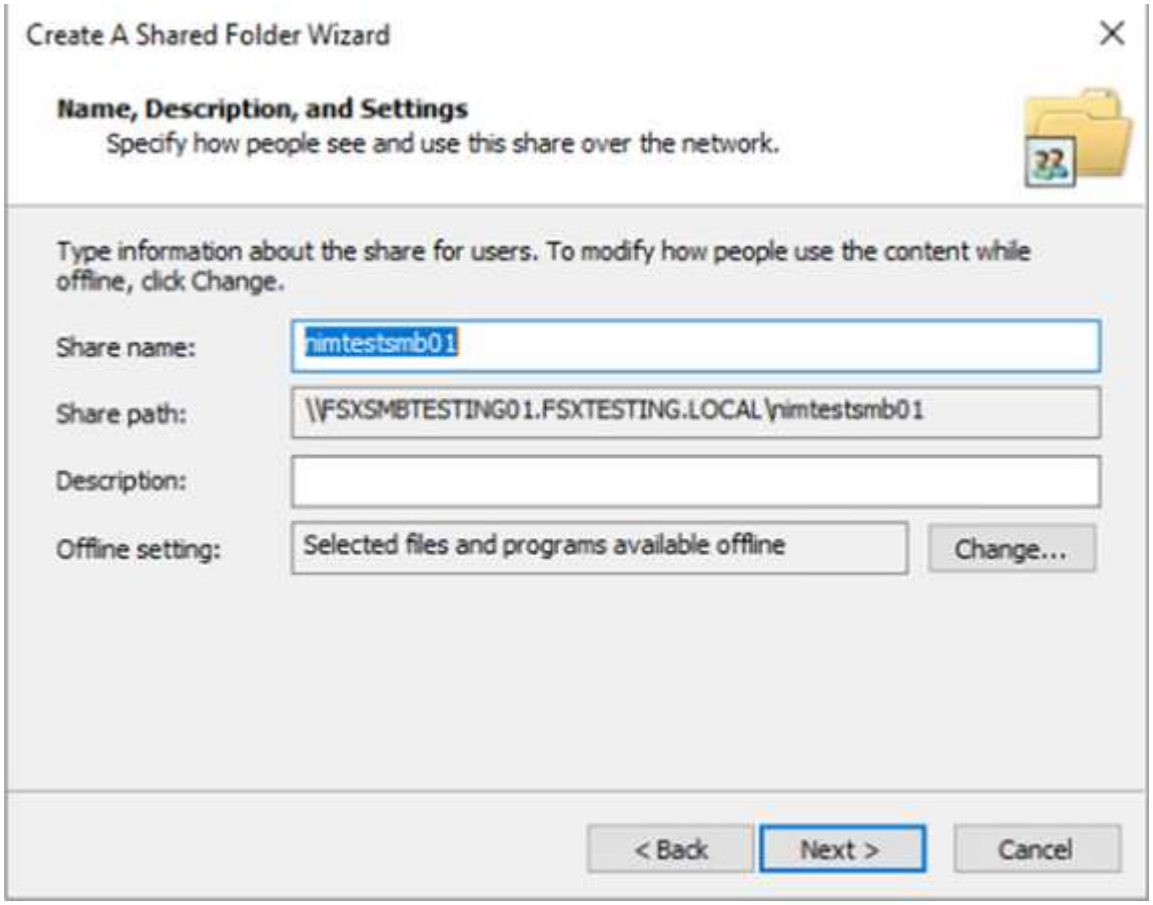

To learn more about creating and managing SMB shares on an Amazon FSx file system, see Creating SMB Shares.

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

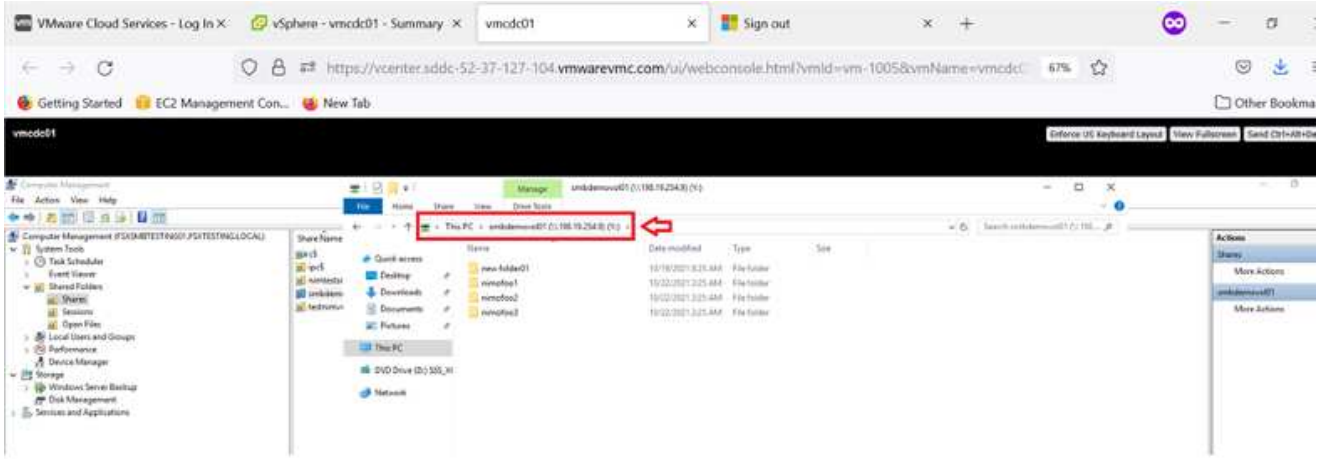

Connect a FSx for NetApp ONTAP LUN to a host using iSCSI

Connect a FSx for NetApp 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 for NetApp 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 for NetApp ONTAP:

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

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

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

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

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

FsxId040eacc5d0ac31017::> igroup show

Vserver   Igroup       Protocol OS Type  Initiators

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

vmcfsxval2svm

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

vmcfsxval2svm

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

Two entries were displayed.

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

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

FsxId040eacc5d0ac31017::> lun show

Vserver   Path                            State   Mapped   Type
Size

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

vmcfsxval2svm

          /vol/blocktest01/lun01          online  mapped   linux
5GB

vmcfsxval2svm

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

Two entries were displayed.

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

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

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

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

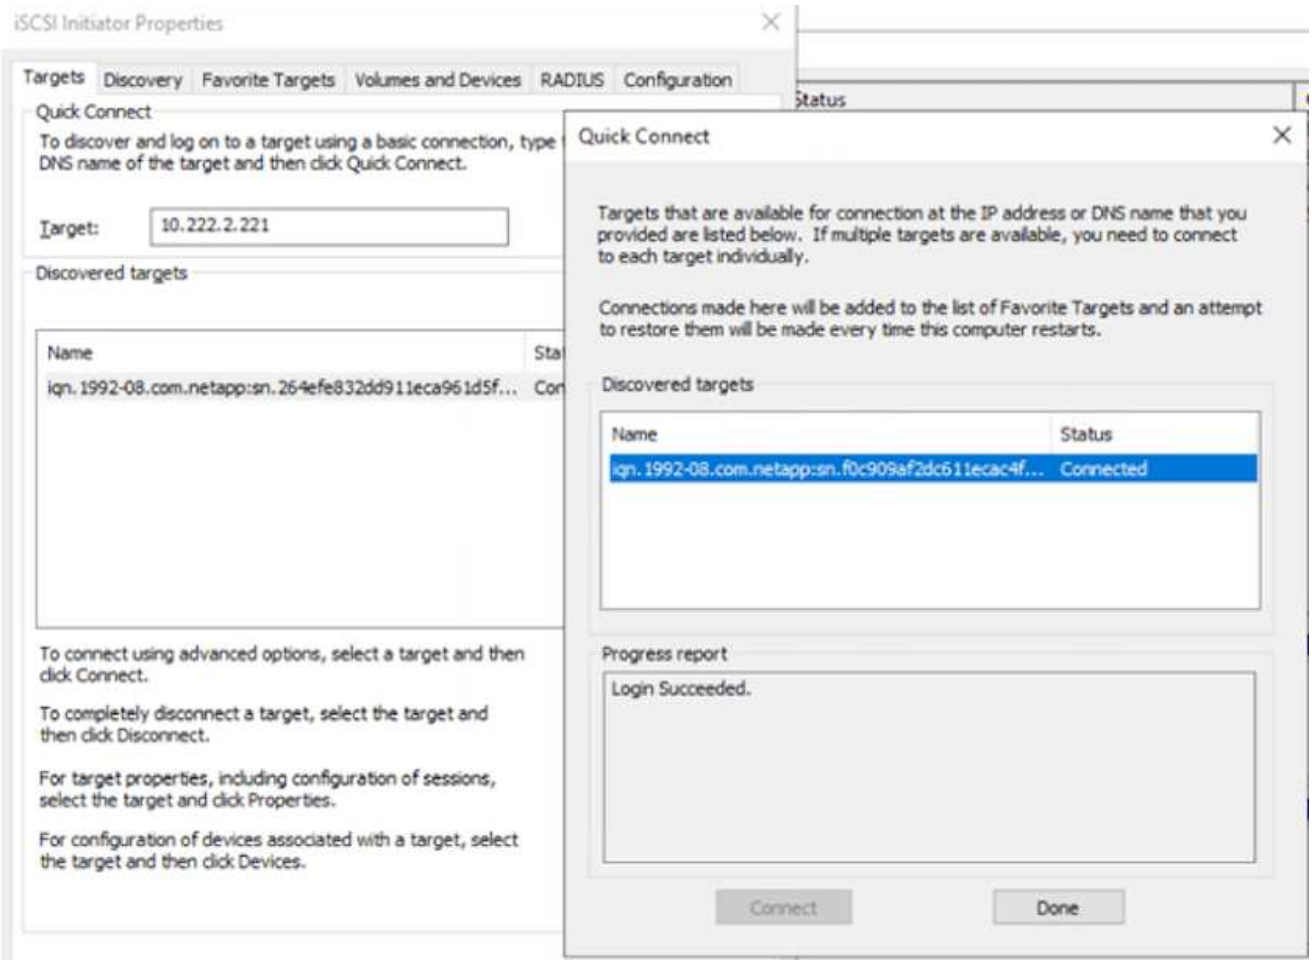

LUNs on the storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

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

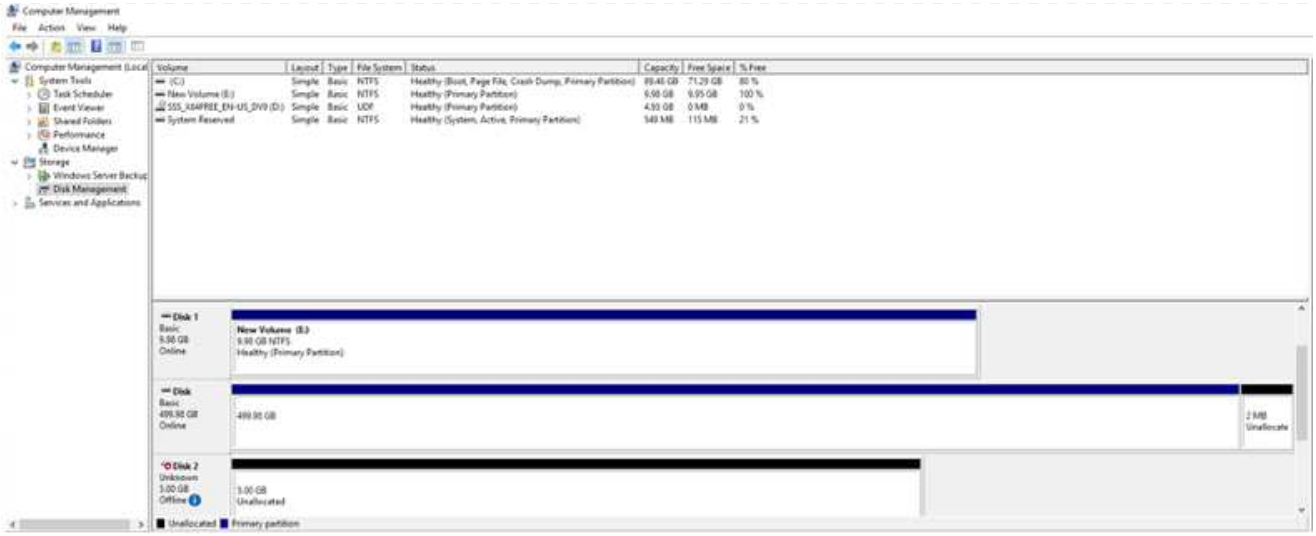

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN and, optionally, format the LUN with a file system by completing the following steps:

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

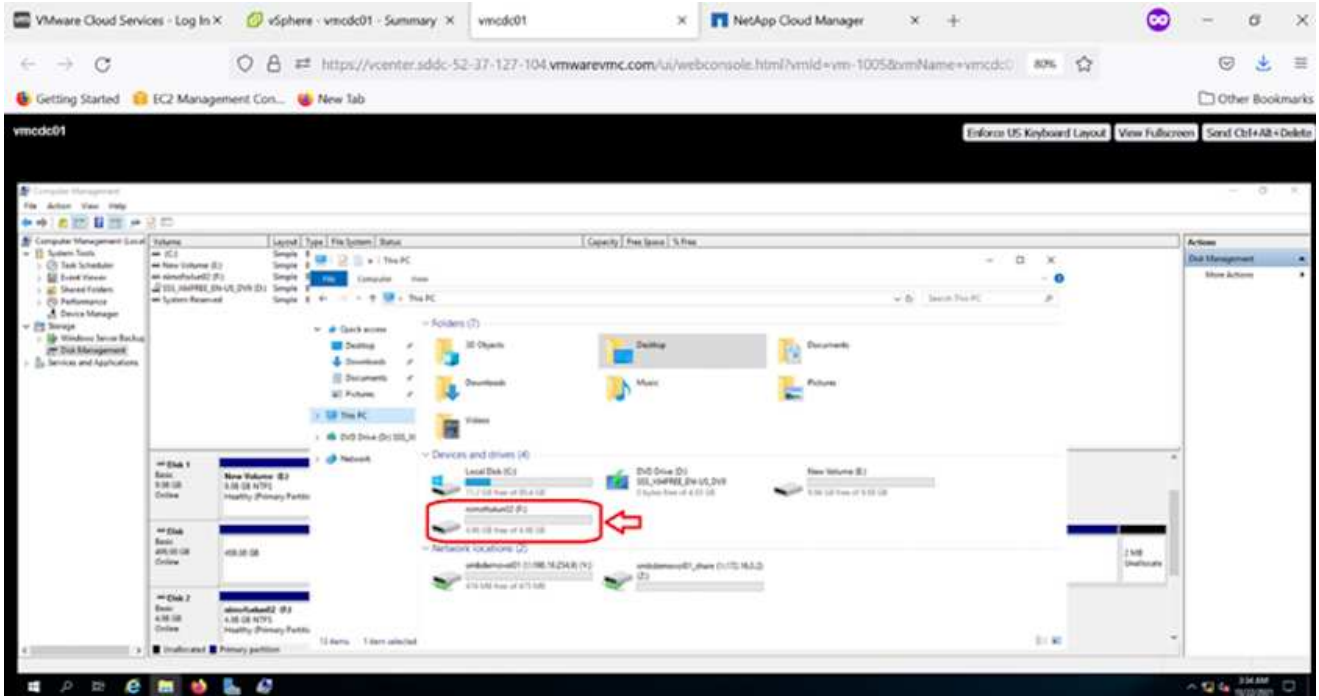

Cloud Volumes ONTAP (CVO)

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

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

manage your data.

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

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

Cloud Volumes ONTAP (CVO) as guest connected storage

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

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

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

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

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

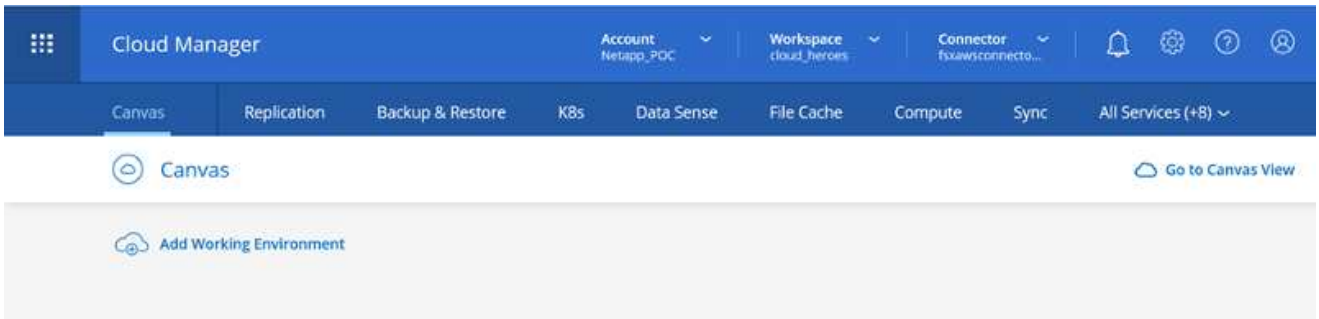

1. On the Cloud Manager home page, click Add a Working Environment and then select AWS as the cloud and the type of the system configuration.

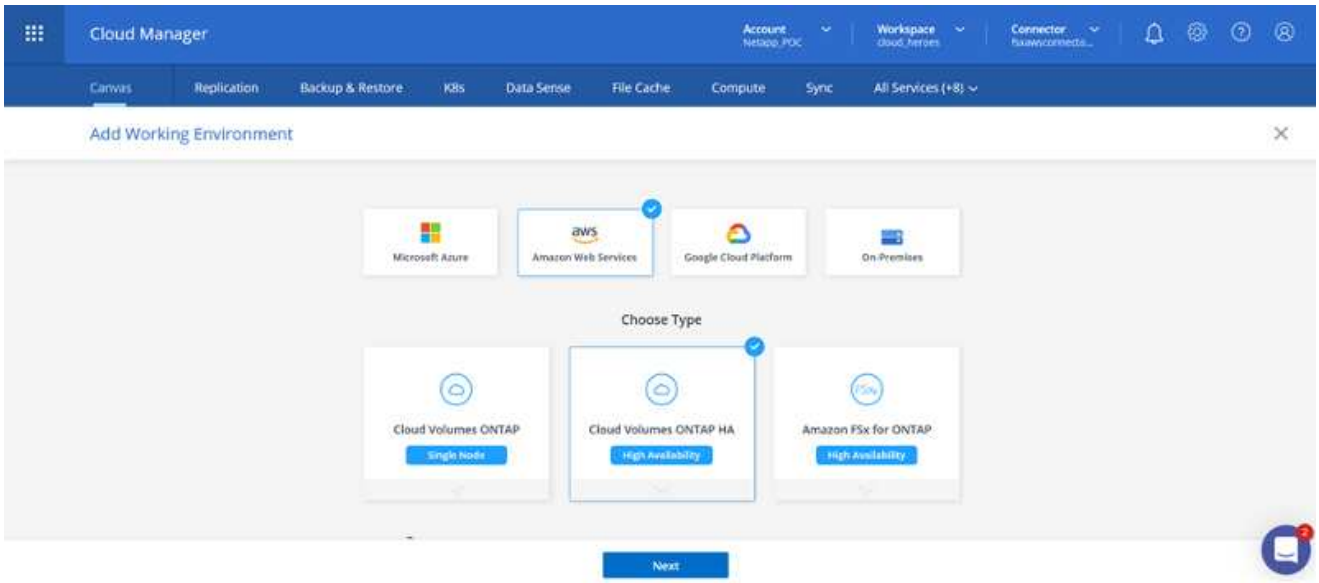

1. Provide the details of the environment to be created including the environment name and admin credentials. Click Continue.

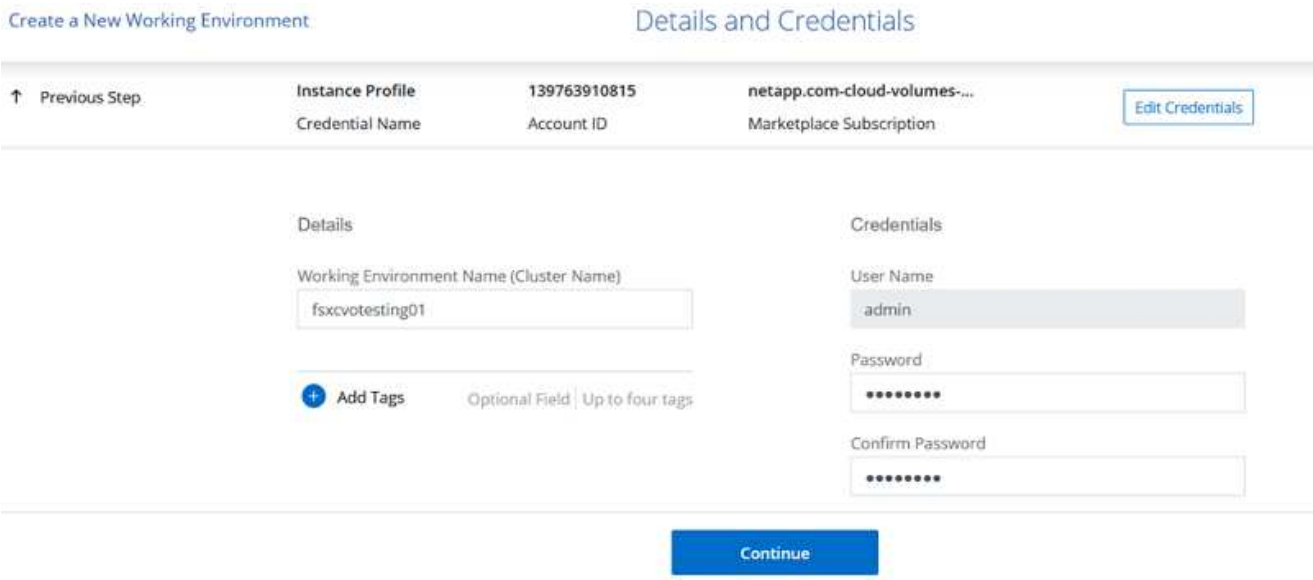

1. Select the add-on services for Cloud Volumes ONTAP deployment, including BlueXP Classification, BlueXP backup and recovery, and Cloud Insights. Click Continue.

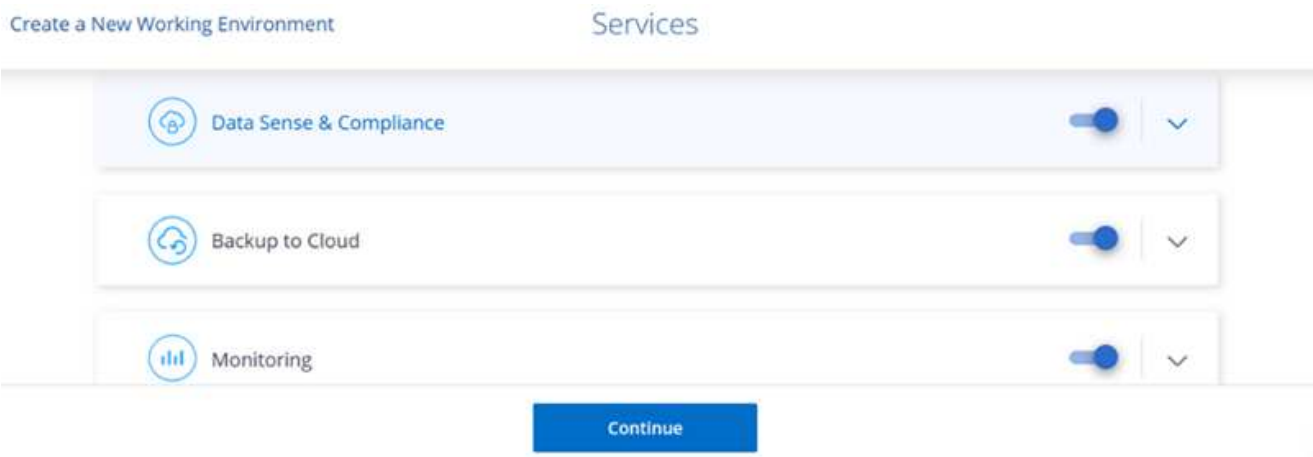

1. On the HA Deployment Models page, choose the Multiple Availability Zones configuration.

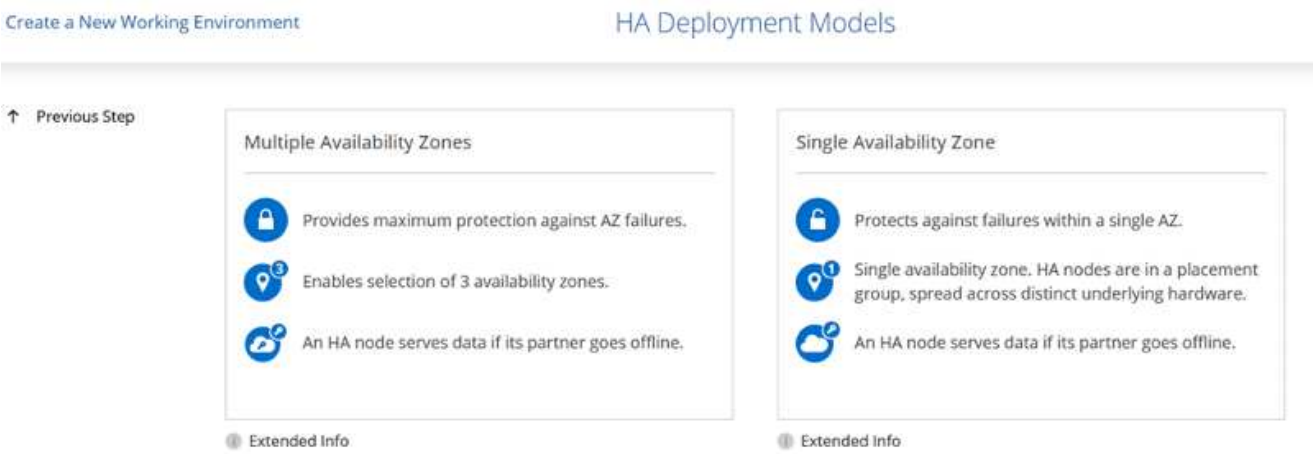

1. On the Region & VPC page, enter the network information and then click Continue.

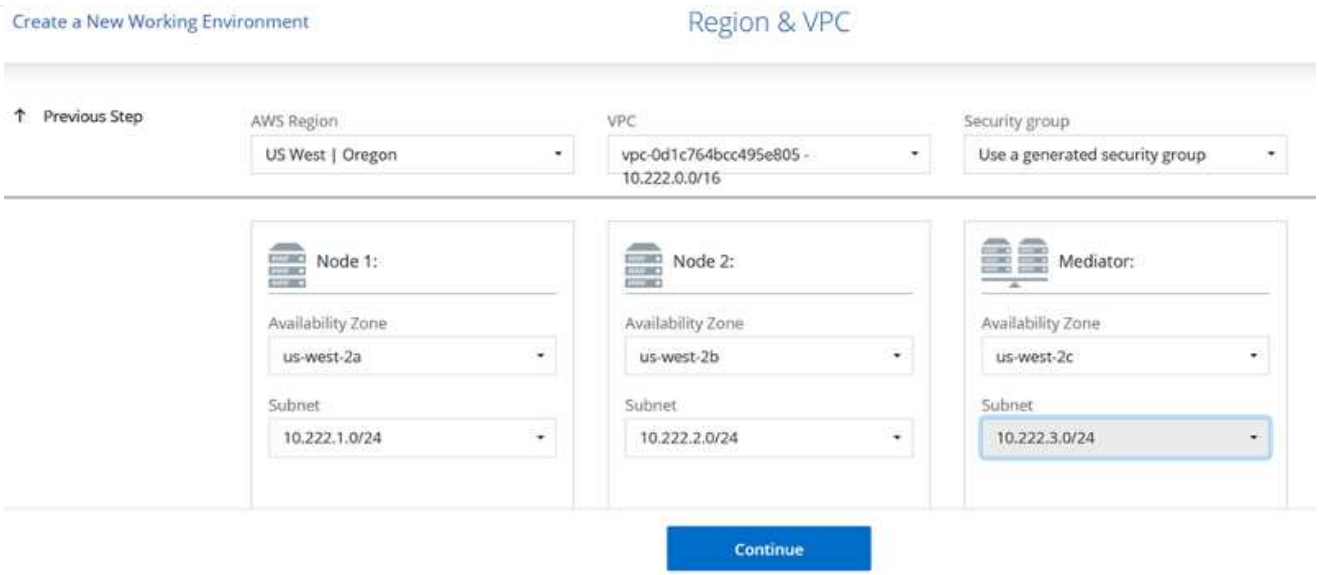

1. On the Connectivity and SSH Authentication page, choose connection methods for the HA pair and the mediator.

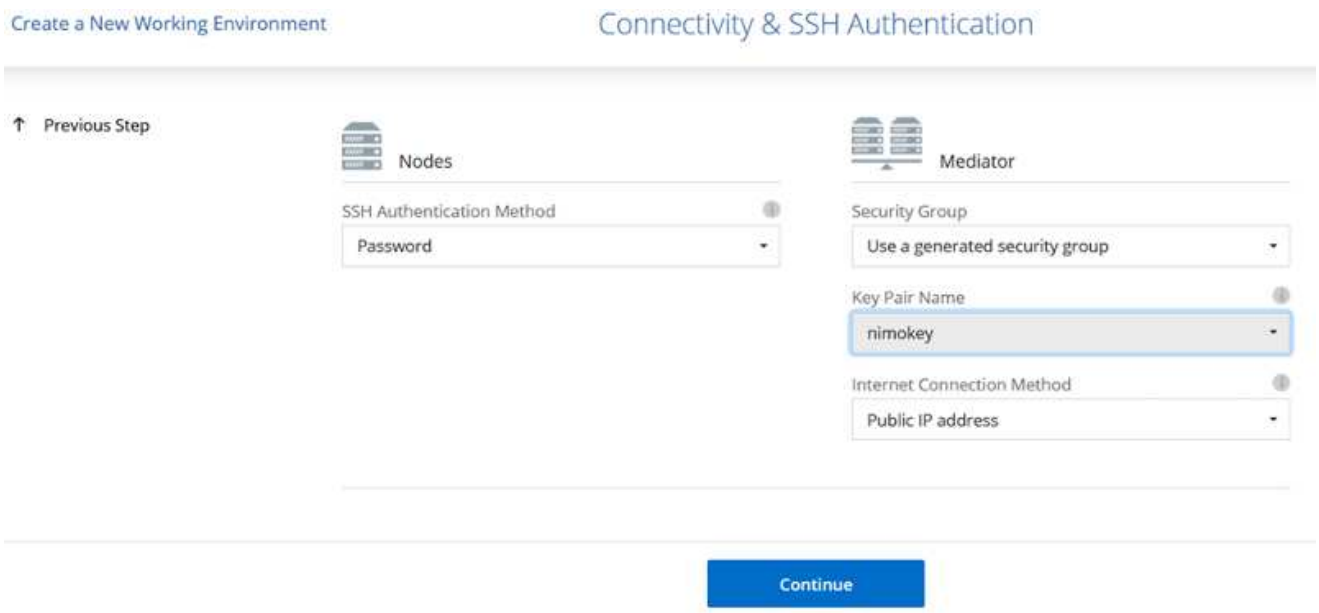

1. Specify the floating IP addresses and then click Continue.

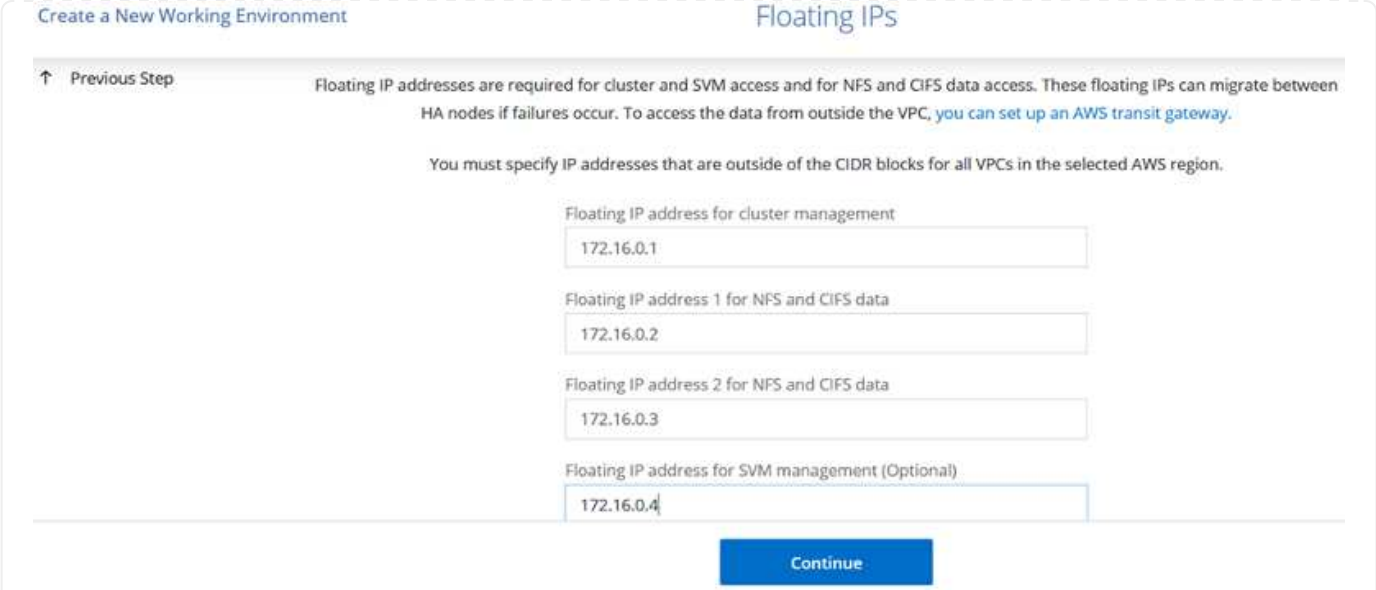

1. Select the appropriate route tables to include routes to the floating IP addresses and then click Continue.

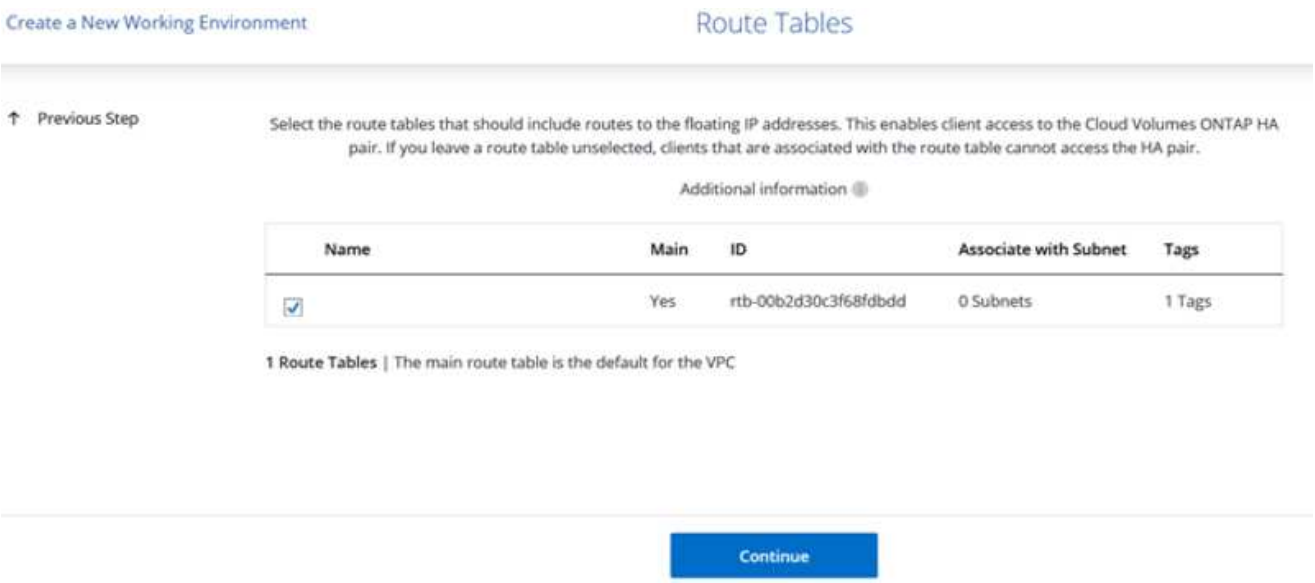

1. On the Data Encryption page, choose AWS-managed encryption.

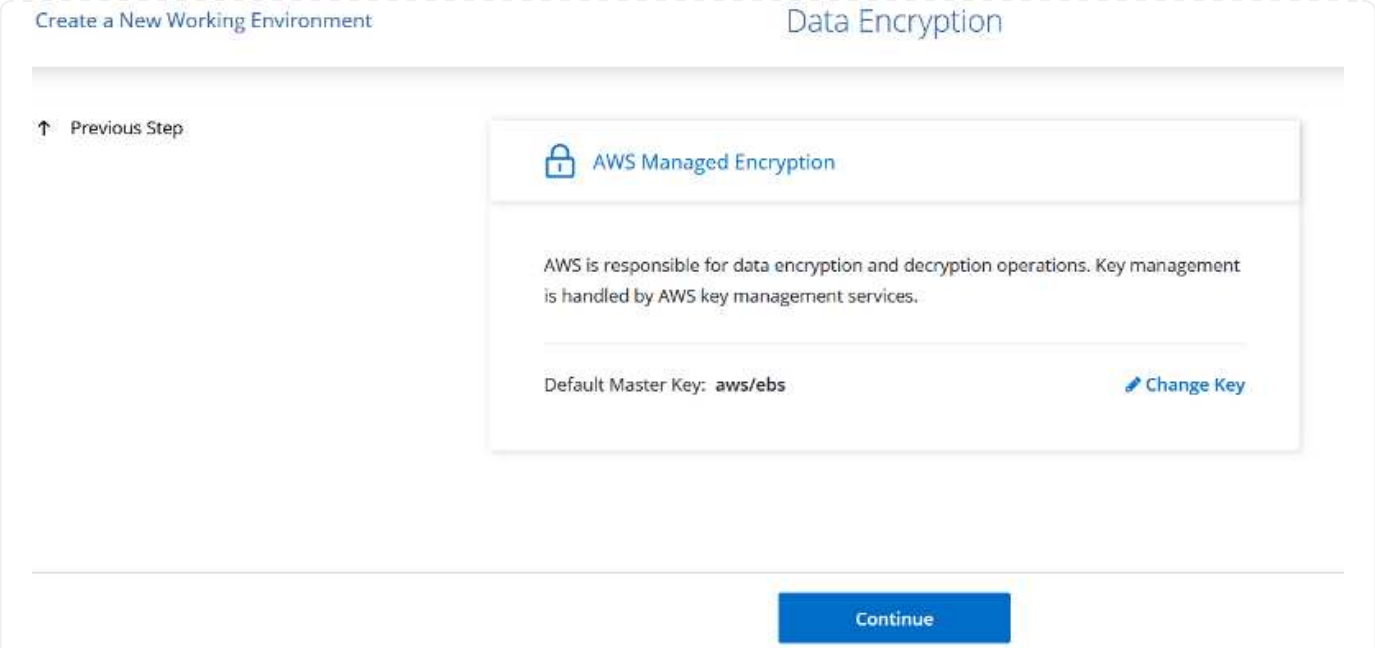

1. Select the license option: Pay-As-You-Go or BYOL for using an existing license. In this example, the Pay-As-You-Go option is used.

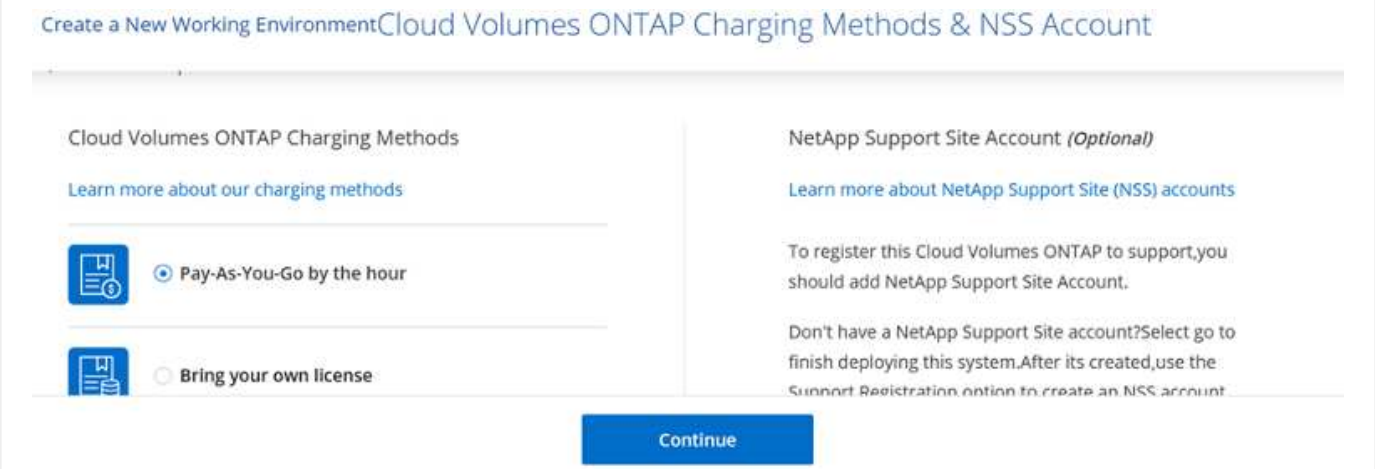

1. Select between several preconfigured packages available based on the type of workload to be deployed on the VMs running on the VMware cloud on AWS SDDC.

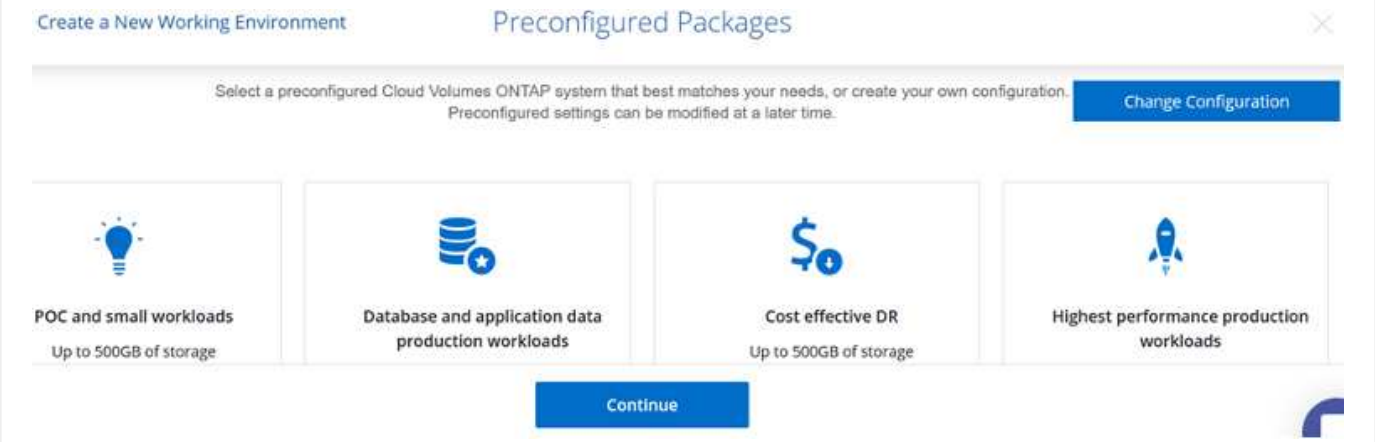

1. On the Review & Approve page, review and confirm the selections.To create the Cloud Volumes ONTAP instance, click Go.

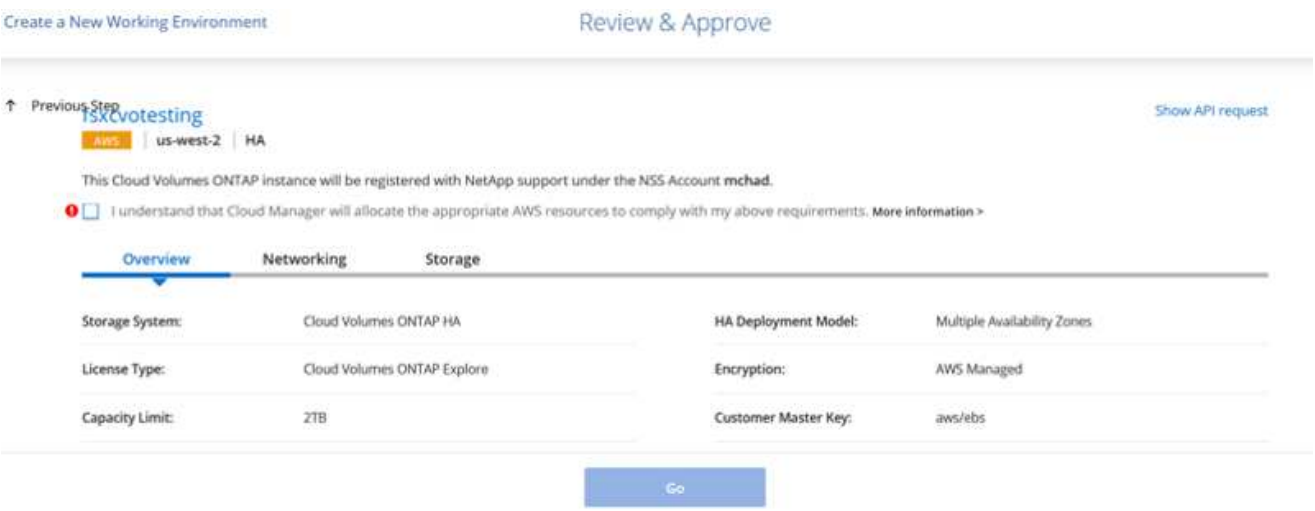

1. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.

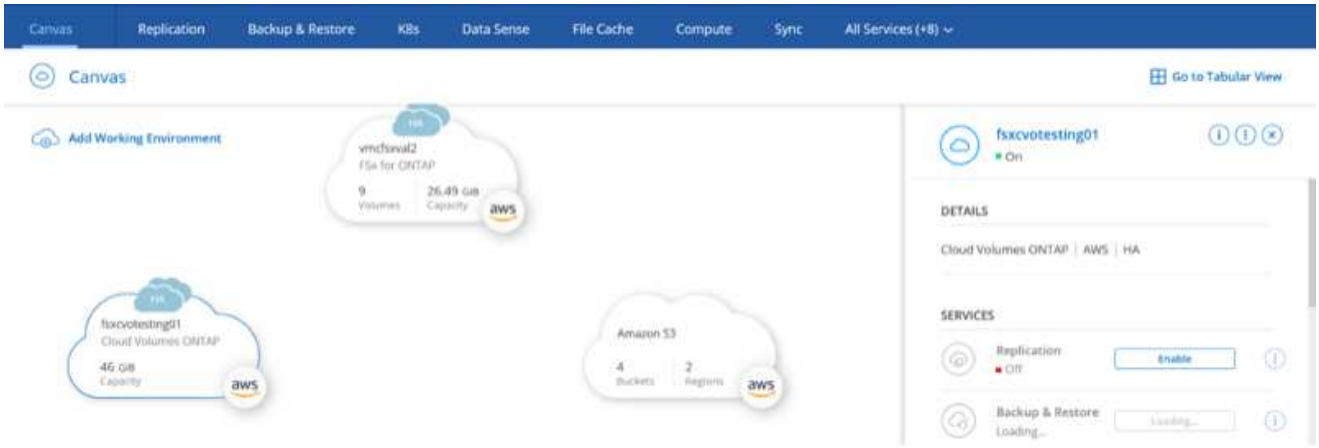

Additional configurations for SMB volumes

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

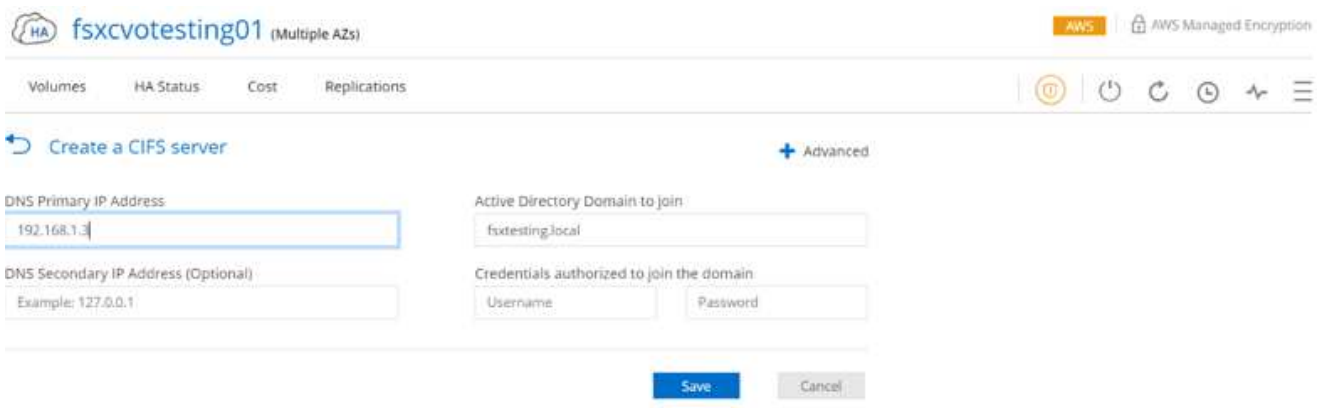

1. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected as the protocol.

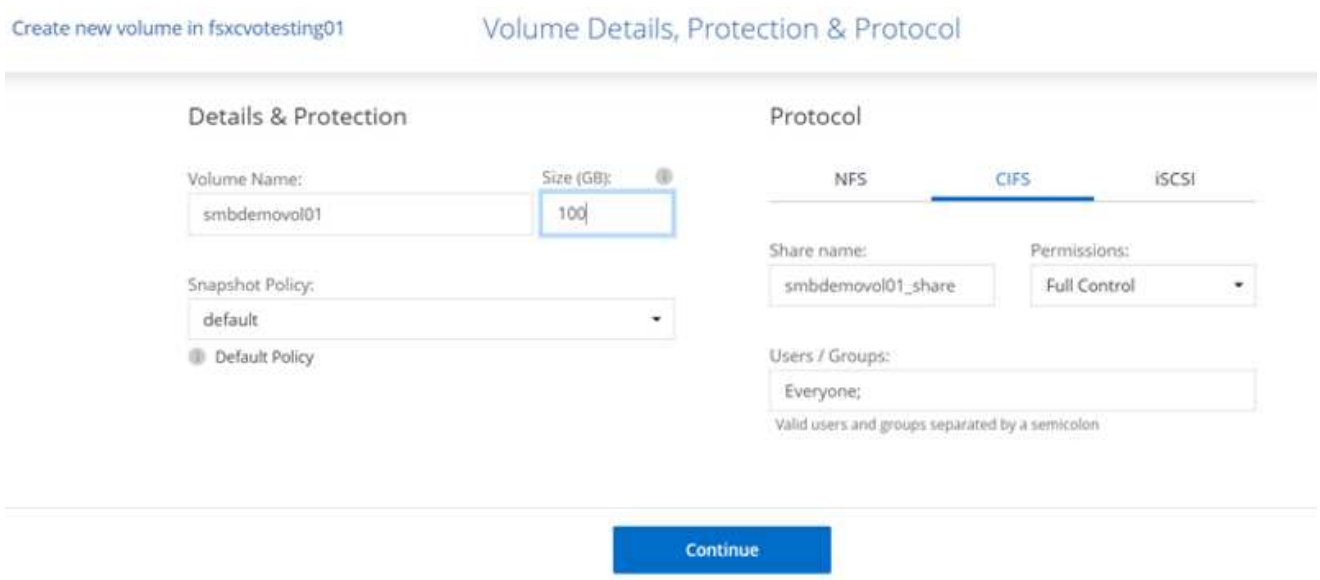

1. After the volume is provisioned, it is availabe under the Volumes pane. Because a CIFS share is provisioned, you should give your users or groups permission to the files and folders and verify that those users can access the share and create a file.

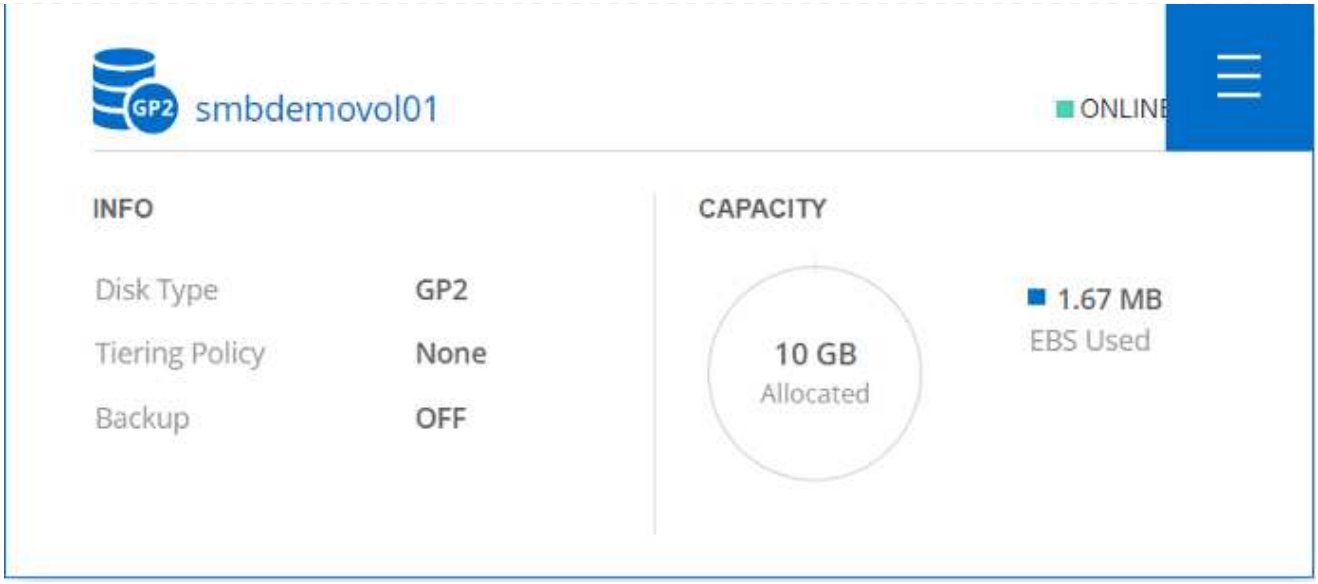

1. After the volume is created, use the mount command to connect to the share from the VM running on the VMware Cloud in AWS SDDC hosts.
2. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on the VMware Cloud in AWS SDDC.

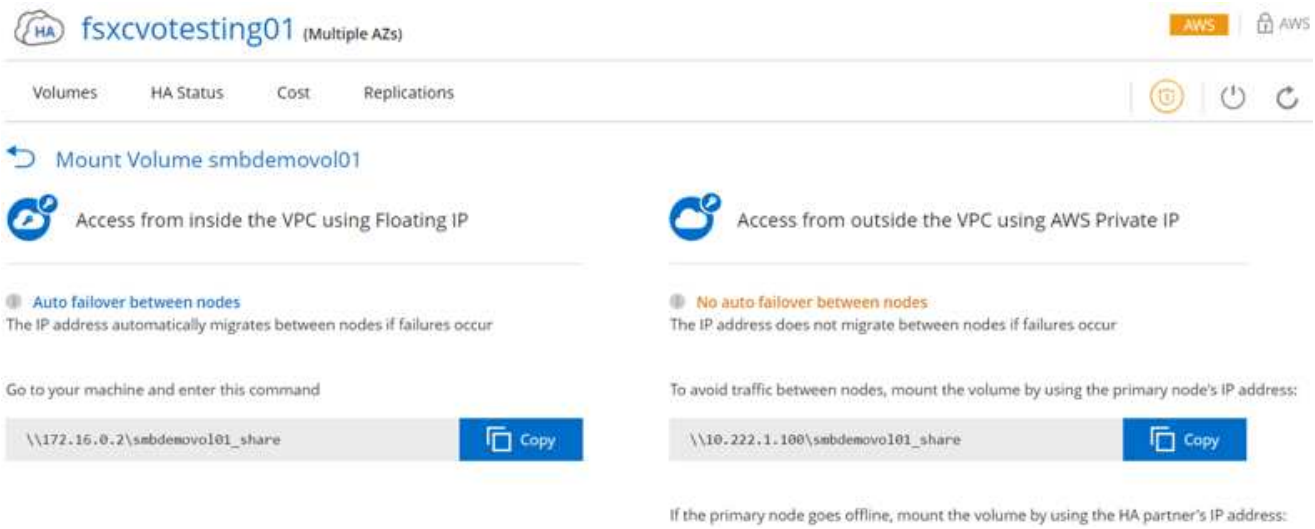

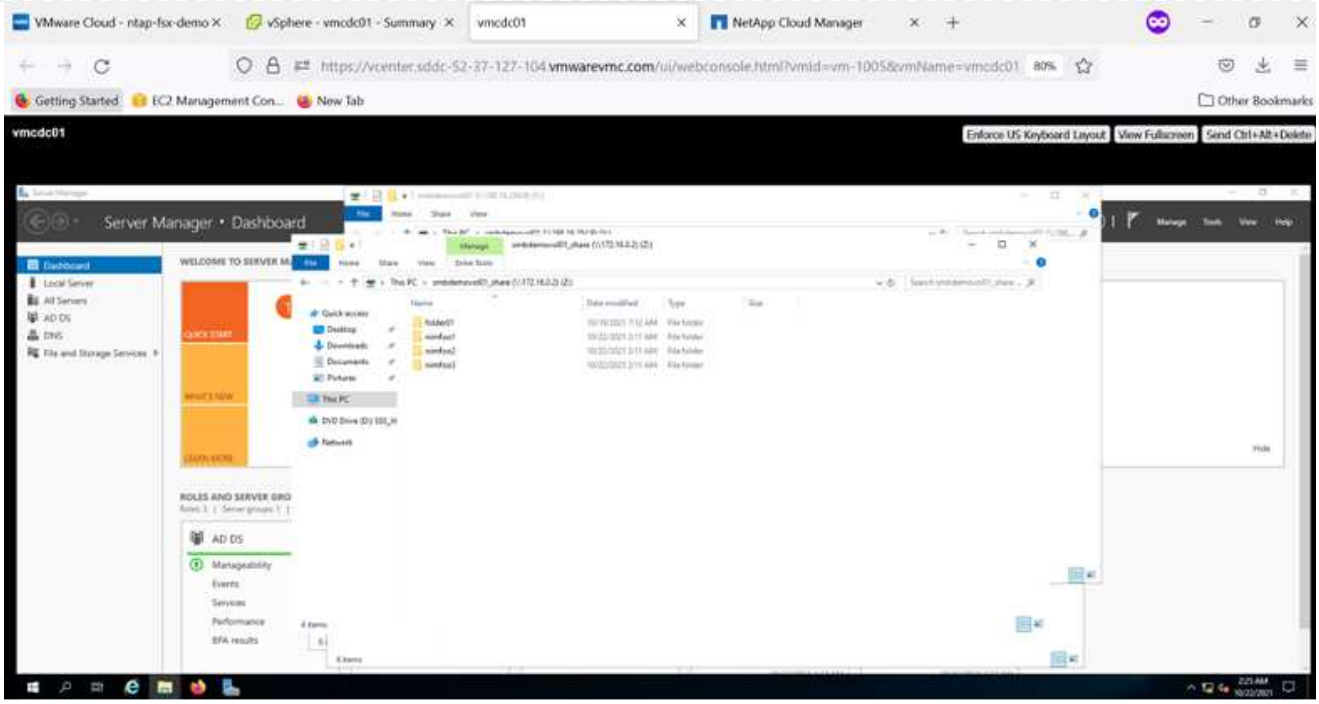
VMware Cloud - ntap-fsx-demo
vSphere - vmcdc01 - Summary
vmcdc01
NetApp Cloud Manager
https://vcenter.sddc-52-37-127-104.vmwarevmc.com/ui/webconsole.html?vmId=vm-1005&vmName=vmcdc01
Getting Started
EC2 Management Con...
New Tab
Other Bookmarks
vmcdc01
Enforce US Keyboard Layout
View Fullscreen
Send Ctrl+Alt+Delete
Server Manager • Dashboard
Dashboard
Local Server
All Servers
AD DS
DNS
File and Storage Services
AD DS
Manageability
Events
Services
Performance
BPA results

Connect the LUN to a host

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

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

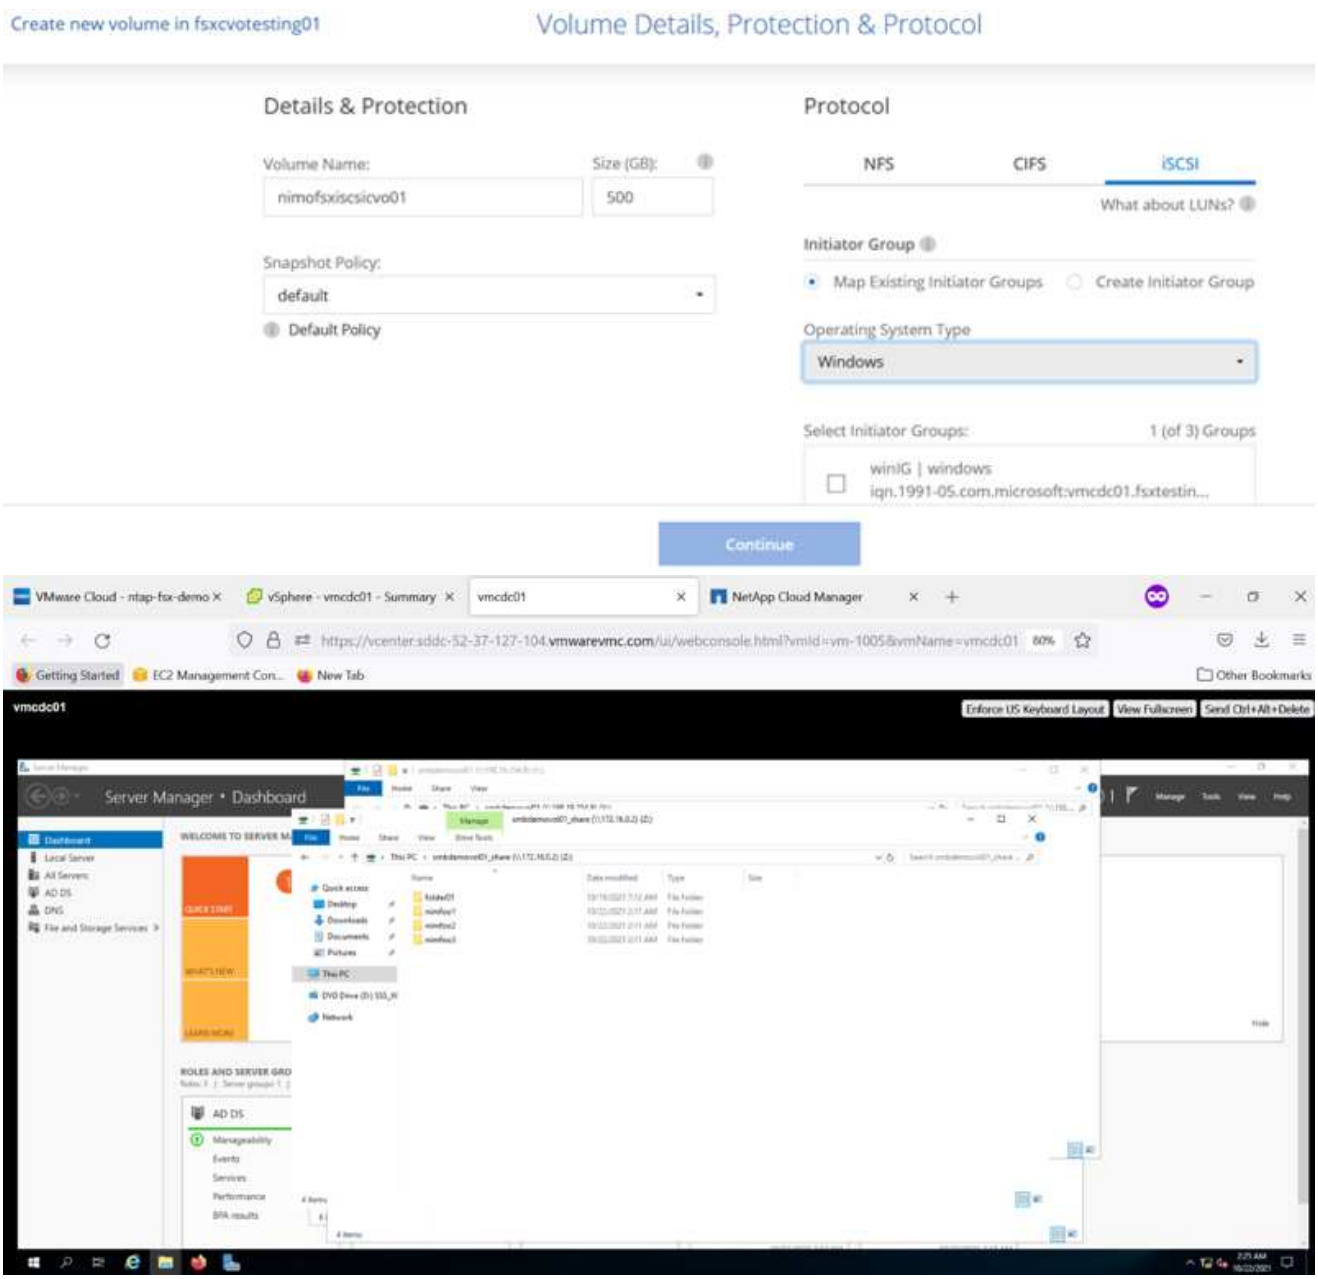

1. After the volume is provisioned, select the volume, and then click Target IQN. To copy the iSCSI Qualified Name (IQN), click Copy. Set up an iSCSI connection from the host to the LUN.

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

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

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

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

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

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

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

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

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

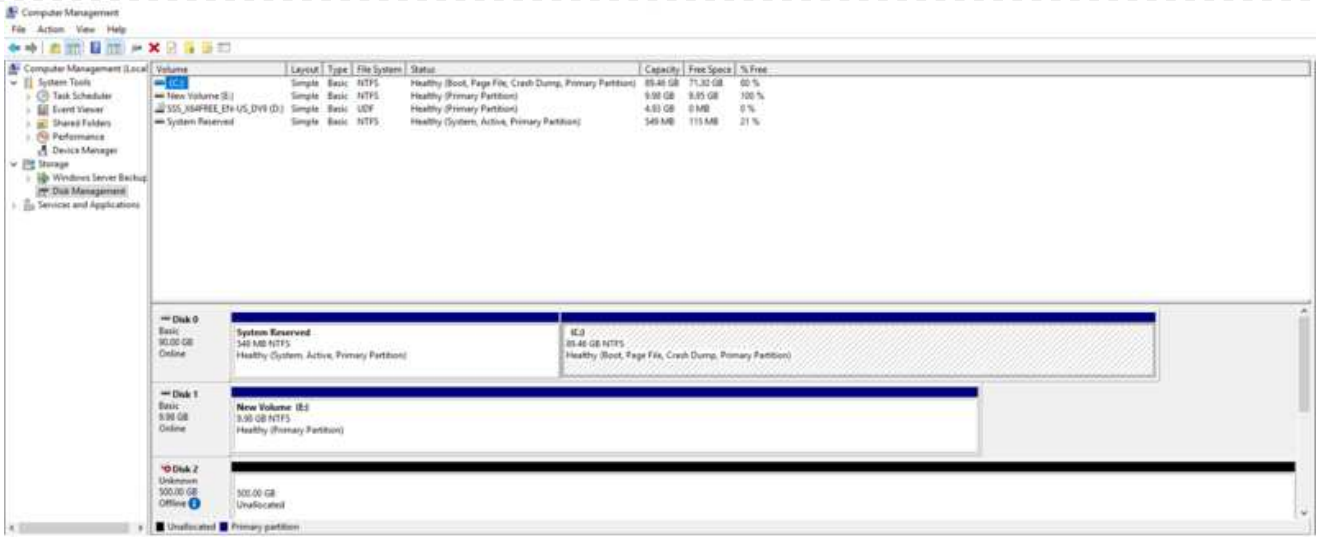

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

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

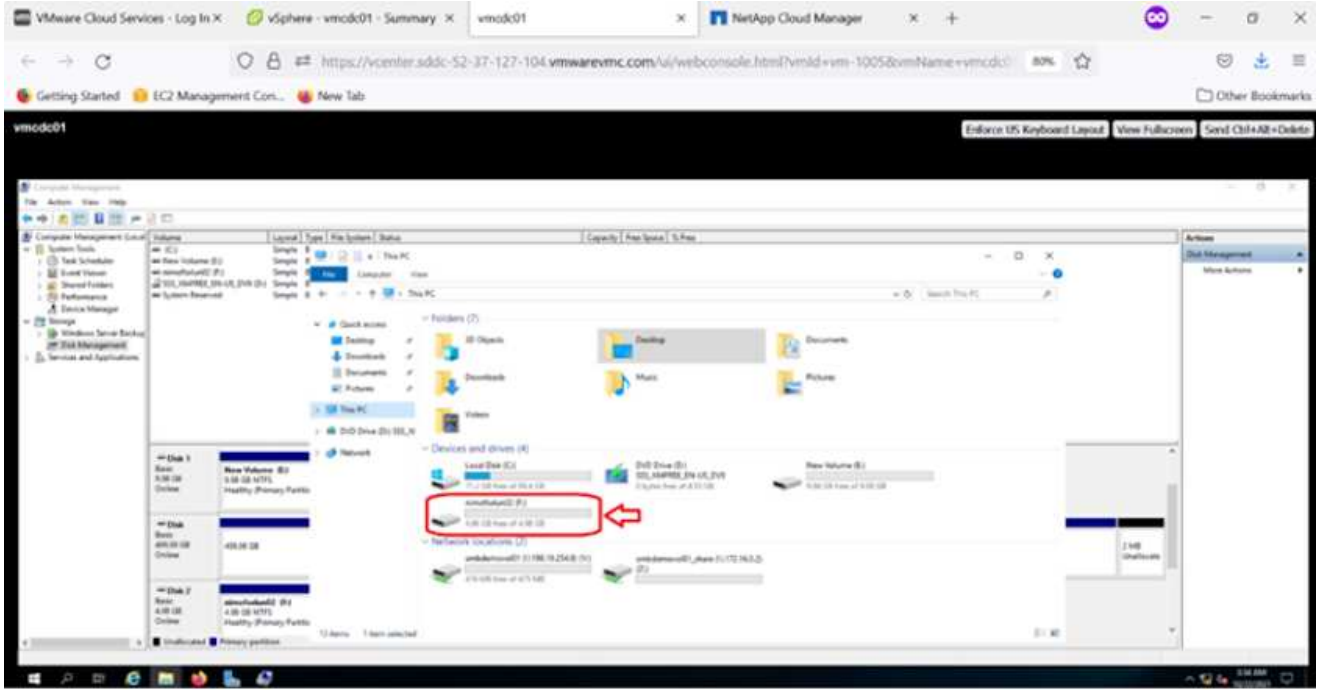

On the Linux clients, ensure the iSCSI daemon is running. After the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration for your Linux distribution. For example, Ubuntu iSCSI configuration can be found here. To verify, run lsblk cmd from the shell.

Mount Cloud Volumes ONTAP NFS volume on Linux client

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

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

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

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

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

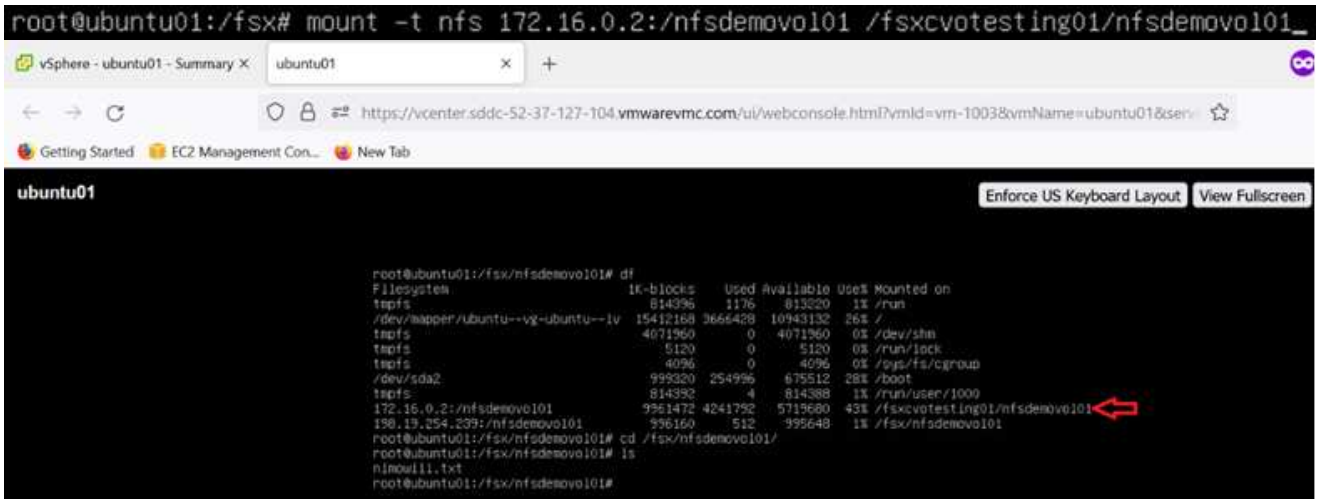

Overview of ANF Datastore Solutions

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/

NetApp Guest Connected Storage Options for Azure

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

Azure NetApp Files (ANF)

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

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

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

Azure NetApp Files (ANF) as guest connected storage

Configure Azure NetApp Files with Azure VMware Solution (AVS)

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

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

Create and mount Azure NetApp Files volumes

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

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

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

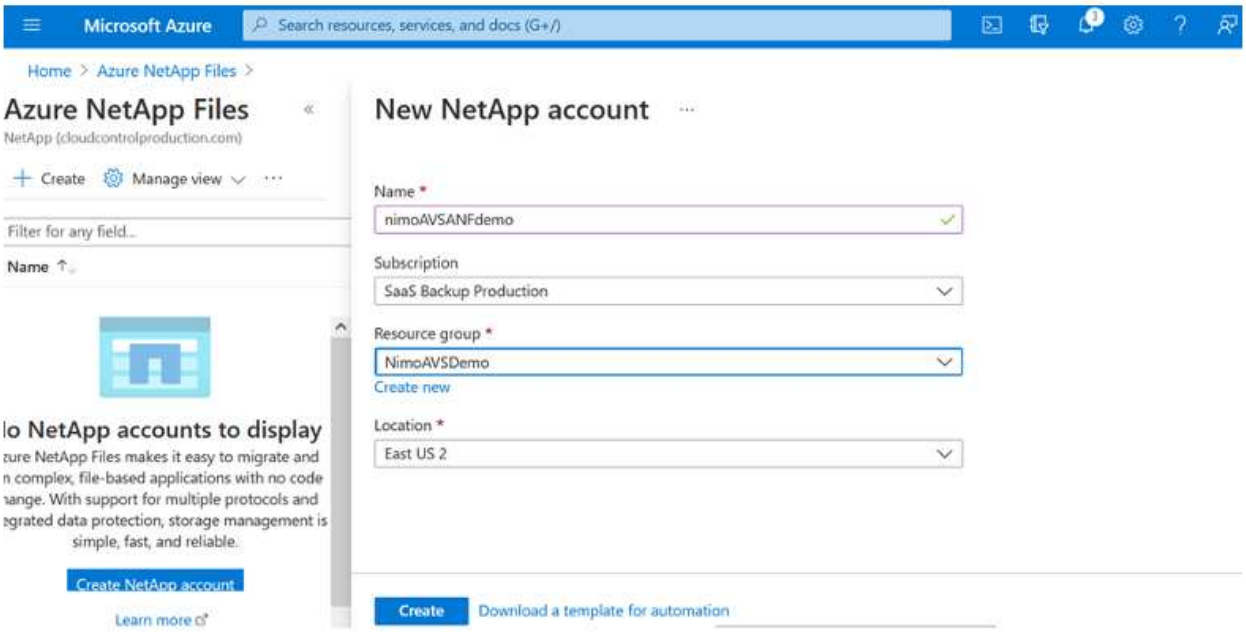

2. After the NetApp account is created, set up the capacity pools with the required service level and size.

   For more information, see Set up a capacity pool.

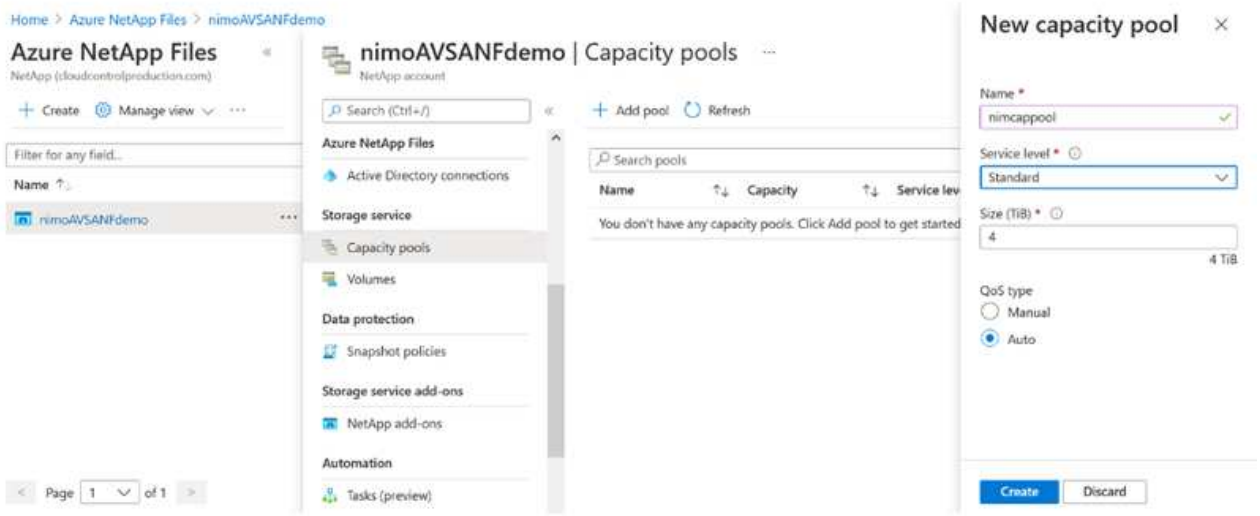

3. Configure the delegated subnet for Azure NetApp Files and specify this subnet while creating the volumes. For detailed steps to create delegated subnet, see Delegate a subnet to Azure NetApp Files.

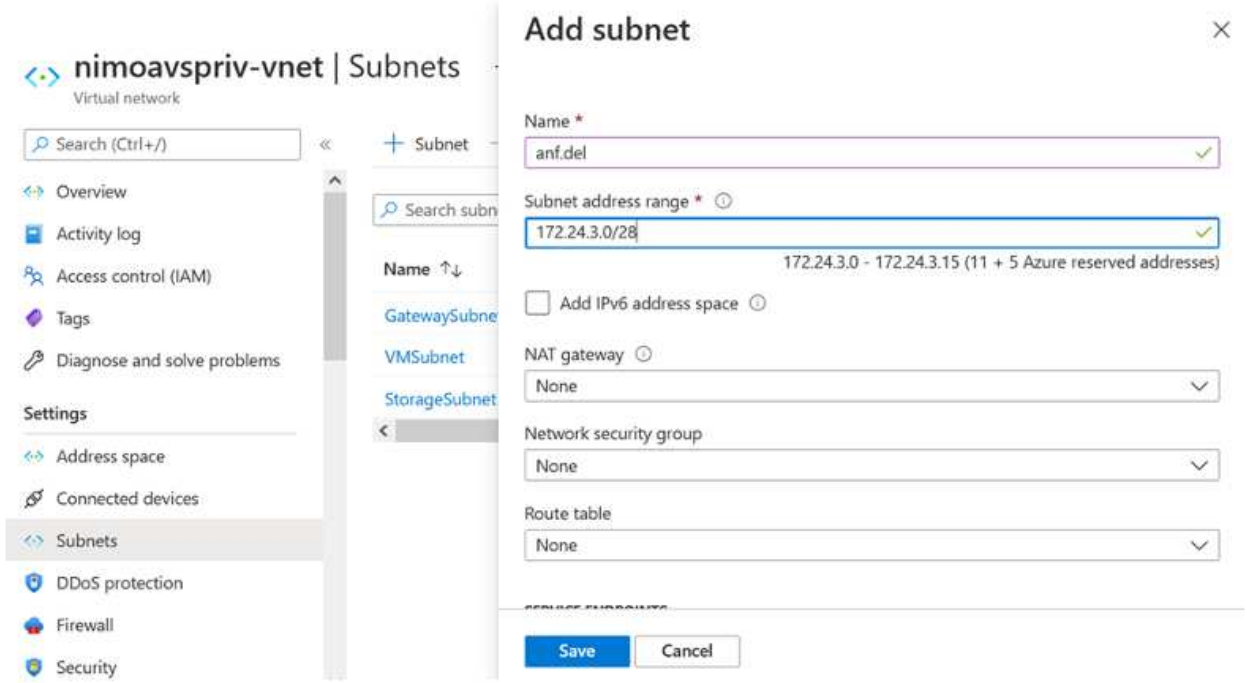

4. Add an SMB volume by using the Volumes blade under the Capacity Pools blade. Make sure the Active Directory connector is configured prior to creating the SMB volume.

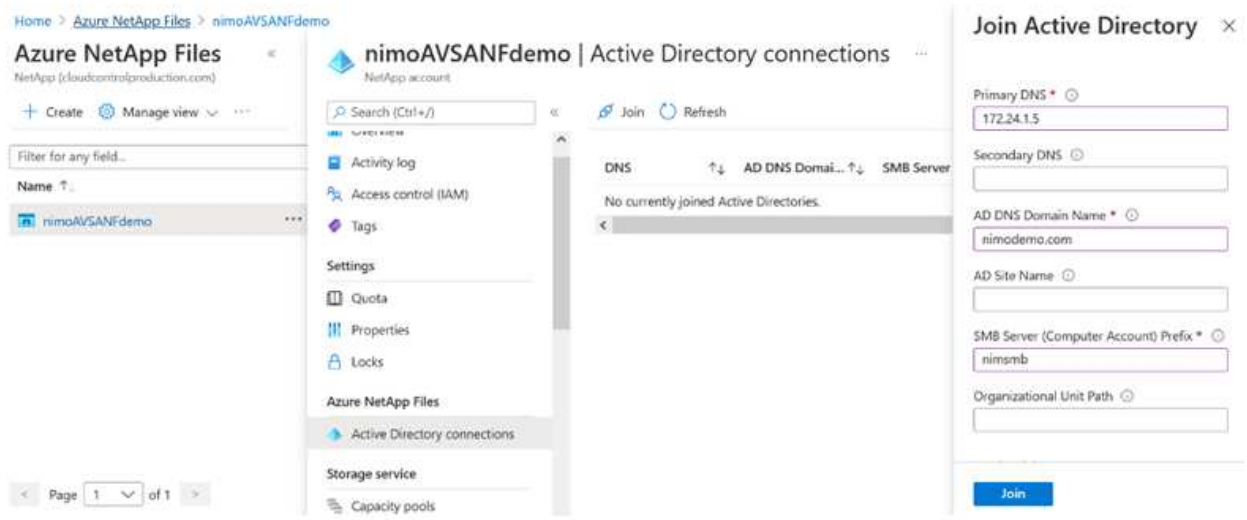

5. Click Review + Create to create the SMB volume.

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

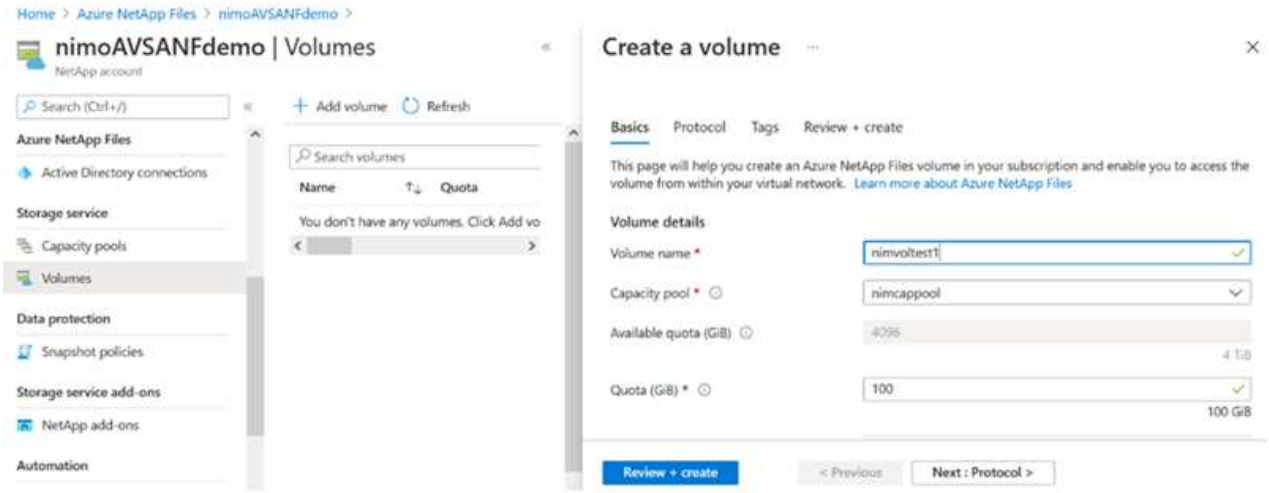

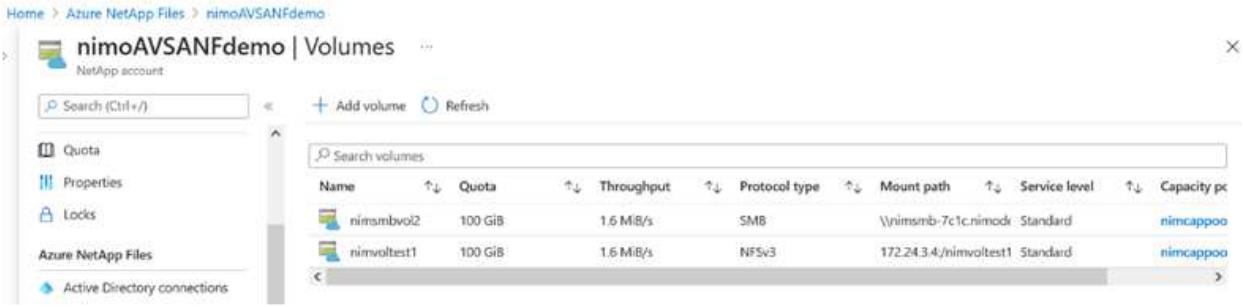

To learn more about Azure NetApp Files volume performance by size or quota, see Performance considerations for Azure NetApp Files.

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

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

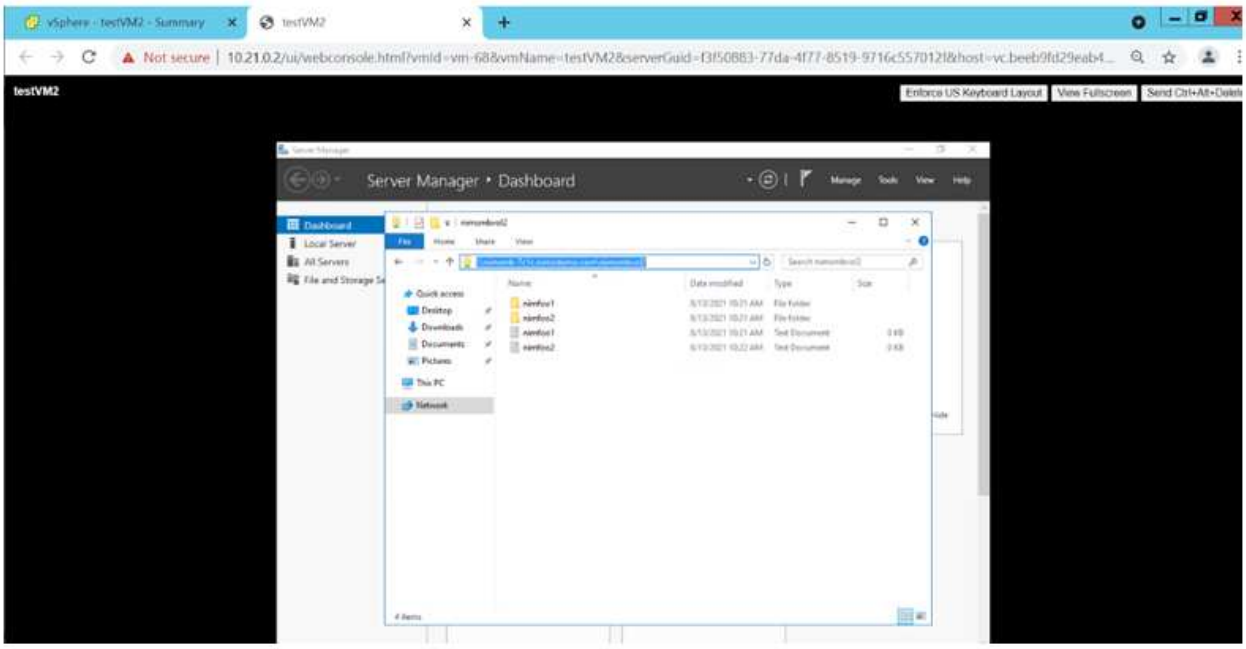

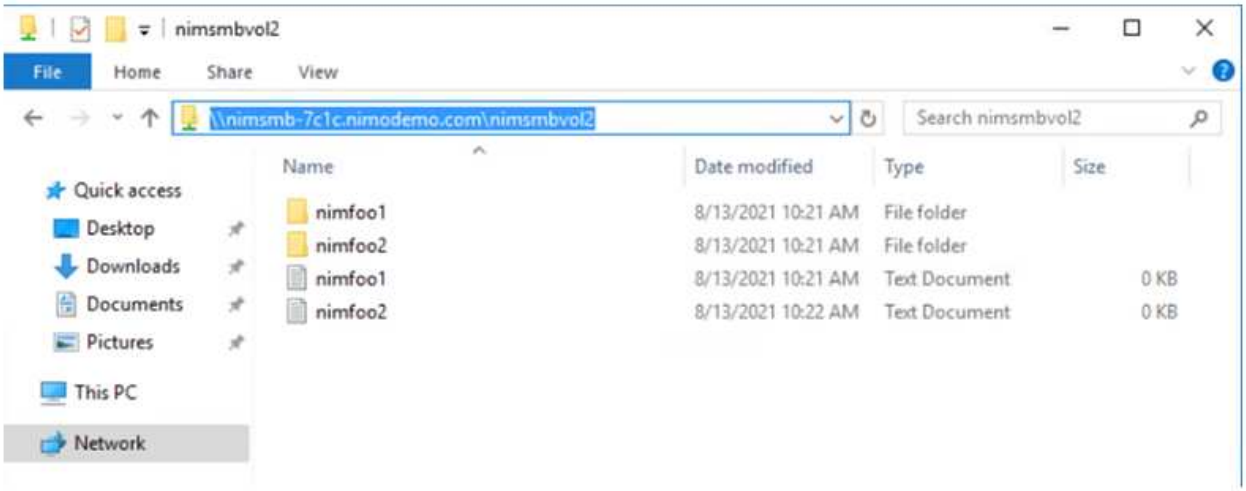

7. To mount NFS volumes on Linux VMs running on Azure VMware Solution SDDC, use this same process. Use volume reshaping or dynamic service level capability to meet the workload demands.

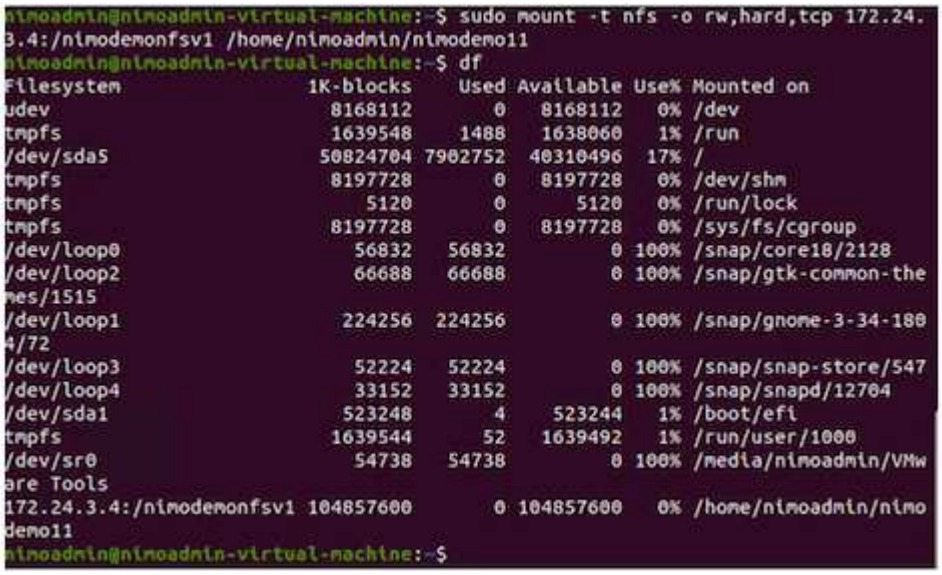

```
nimoadmin@nimoadmin-virtual-machine:~$ sudo mount -t nfs -o rw,hard,tcp 172.24.
3.4:/nimodemonfsv1 /home/nimoadmin/nimodemo11
nimoadmin@nimoadmin-virtual-machine:~$ df
Filesystem                  1K-blocks    Used Available Use% Mounted on
udev                          8168112       0   8168112   0% /dev
tmpfs                         1639548    1488   1638060   1% /run
/dev/sda5                    50824704 7902752  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-the
mes/1515
/dev/loop1                     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/VMw
are Tools
172.24.3.4:/nimodemonfsv1 104857600       0 104857600   0% /home/nimoadmin/nimo
demo11
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 you IT staff in all-new methods to manage your data.

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

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

Cloud Volumes ONTAP (CVO) as guest connected storage

Deploy new Cloud Volumes ONTAP in Azure

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

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

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

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

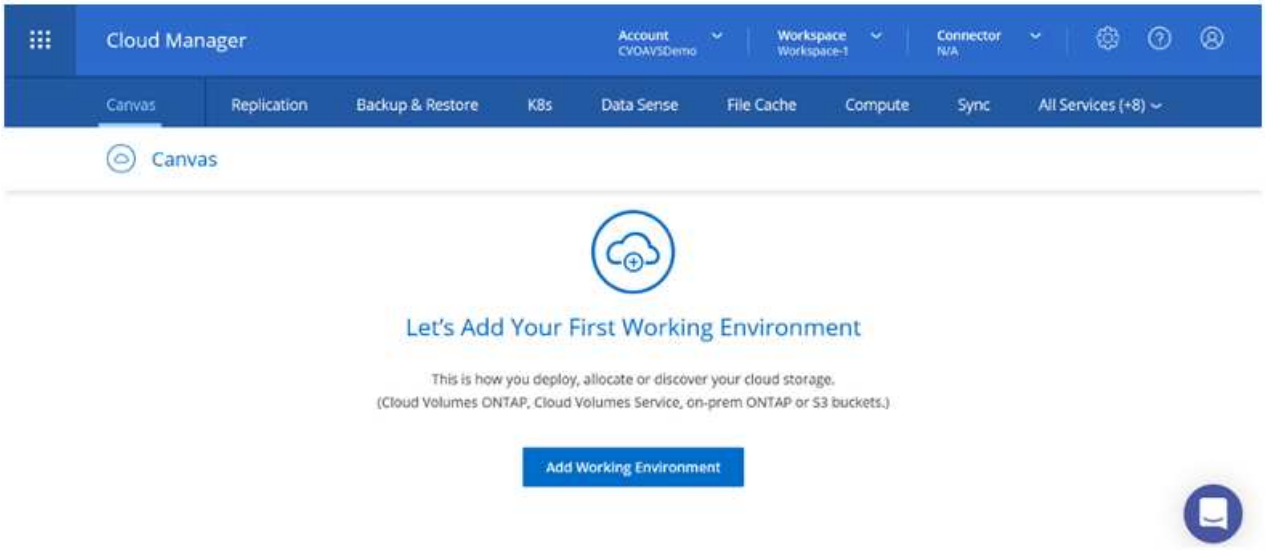

2. On the Cloud Manager home page, click Add a Working Environment and then select Microsoft Azure as the cloud and the type of the system configuration.

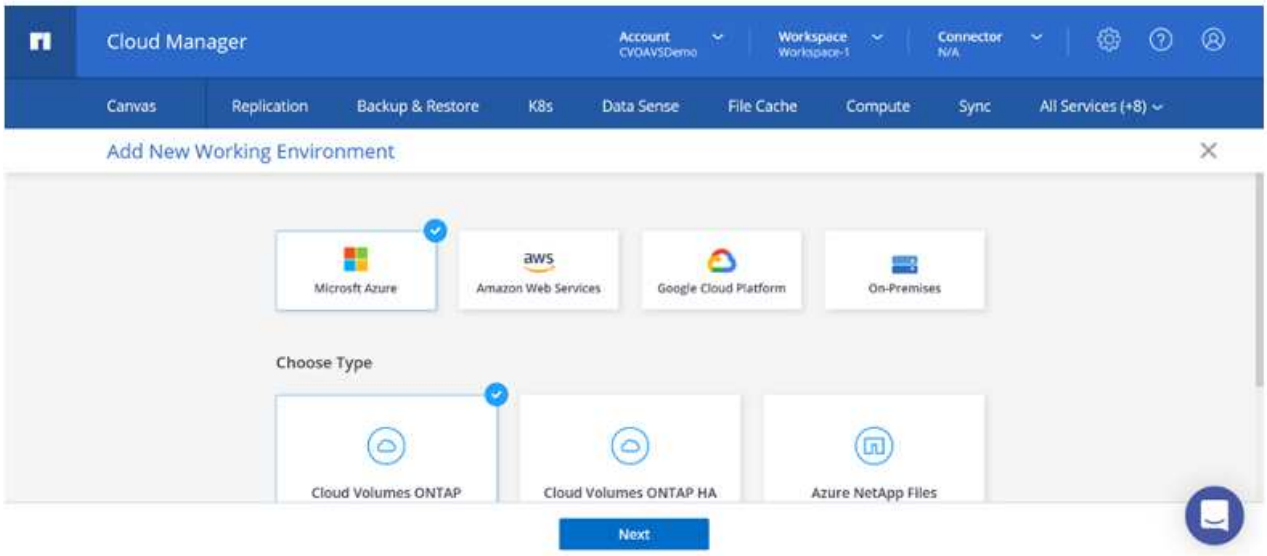

3. When creating the first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to deploy a Connector.

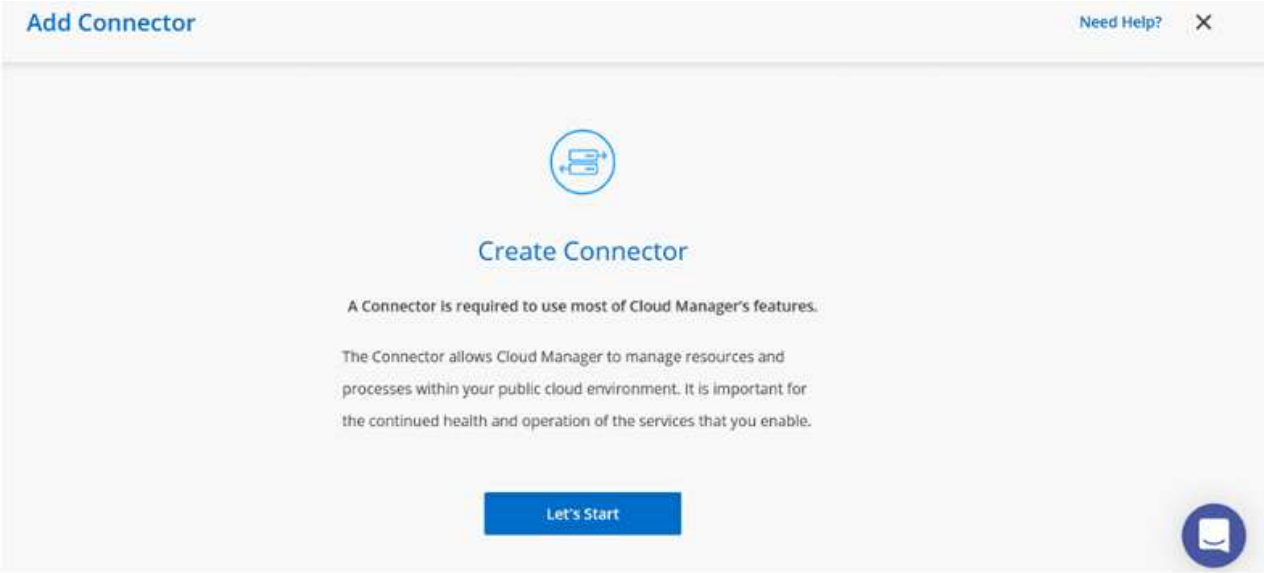

4. After the connector is created, update the Details and Credentials fields.

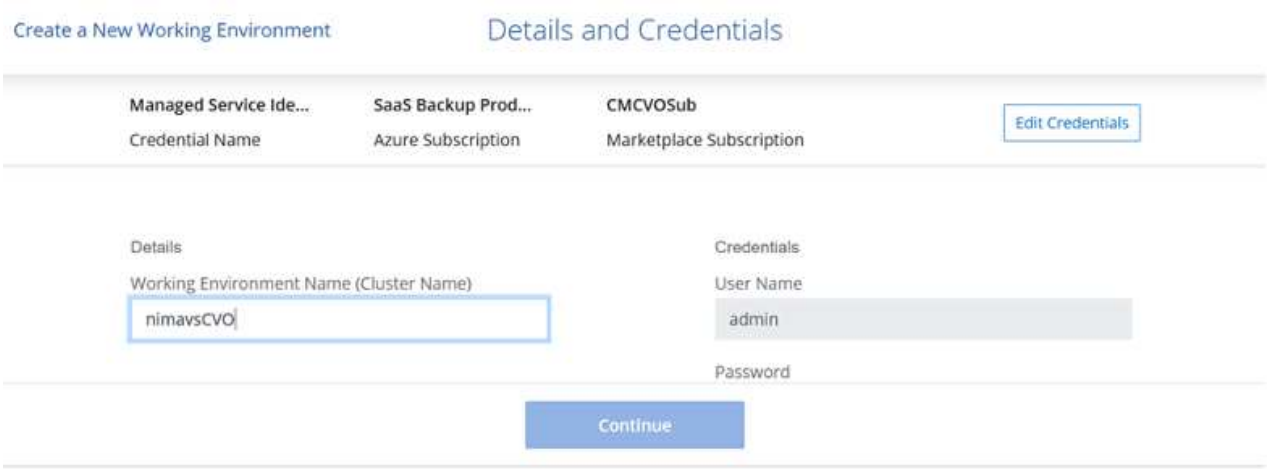

5. Provide the details of the environment to be created including the environment name and admin credentials. Add resource group tags for the Azure environment as an optional parameter. After you are done, click Continue.

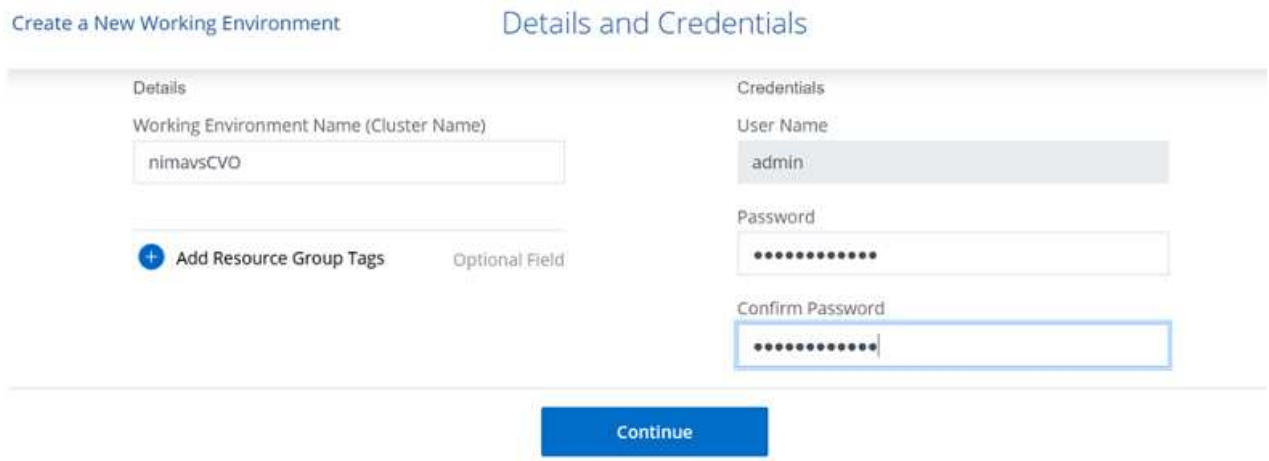

6. Select the add-on services for Cloud Volumes ONTAP deployment, including BlueXP Classification, BlueXP backup and recovery, and Cloud Insights. Select the services and then click Continue.

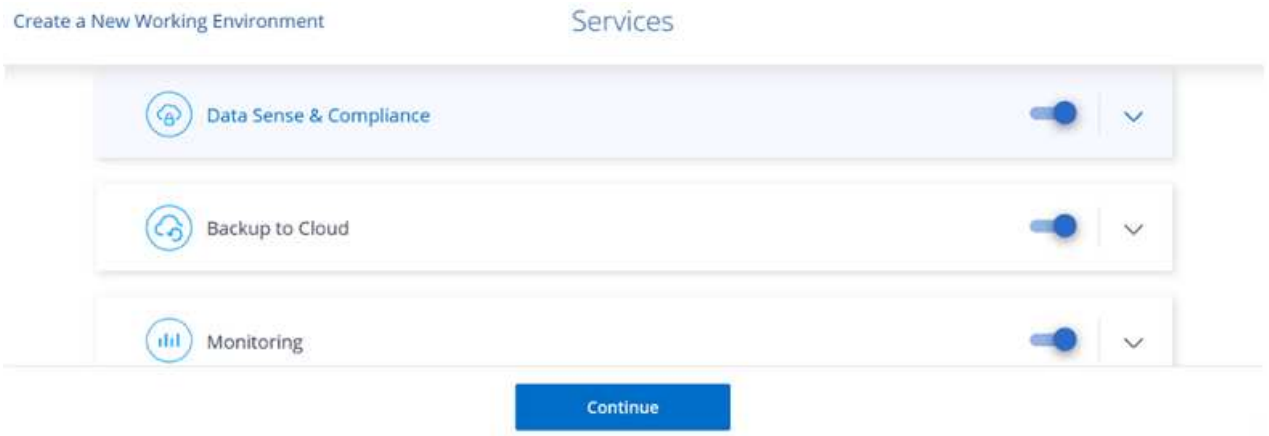

7. Configure the Azure location and connectivity. Select the Azure Region, resource group, VNet, and subnet to be used.

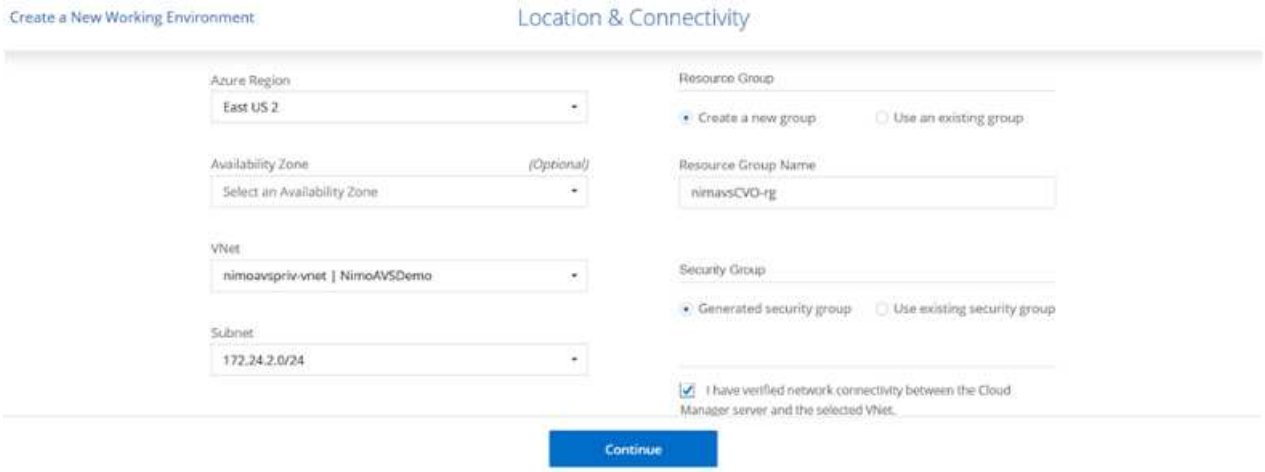

8. Select the license option: Pay-As-You-Go or BYOL for using existing license. In this example, Pay-As-You-Go option is used.

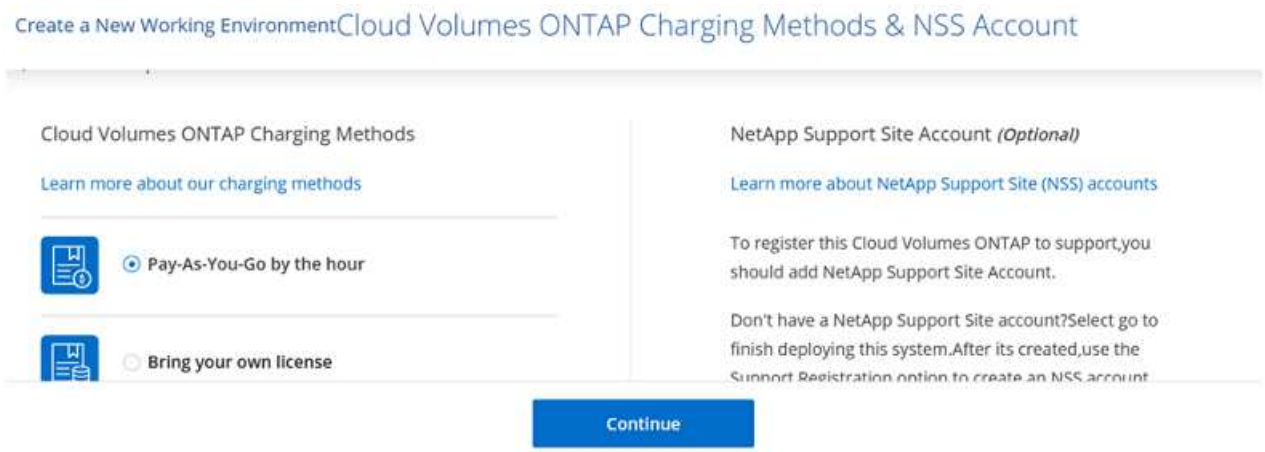

9. Select between several preconfigured packages available for the various types of workloads.

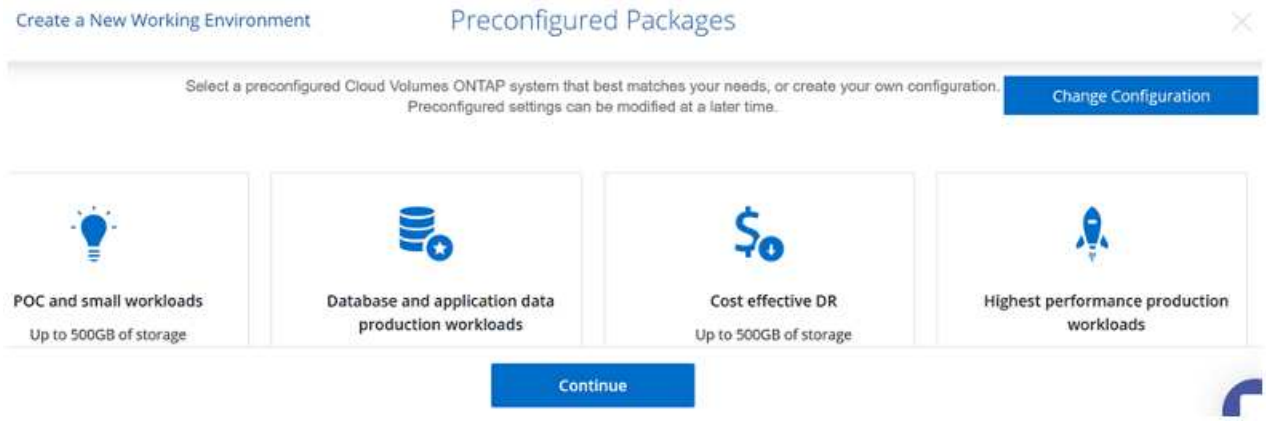

10. Accept the two agreements regarding activating support and allocation of Azure resources.To create the Cloud Volumes ONTAP instance, click Go.

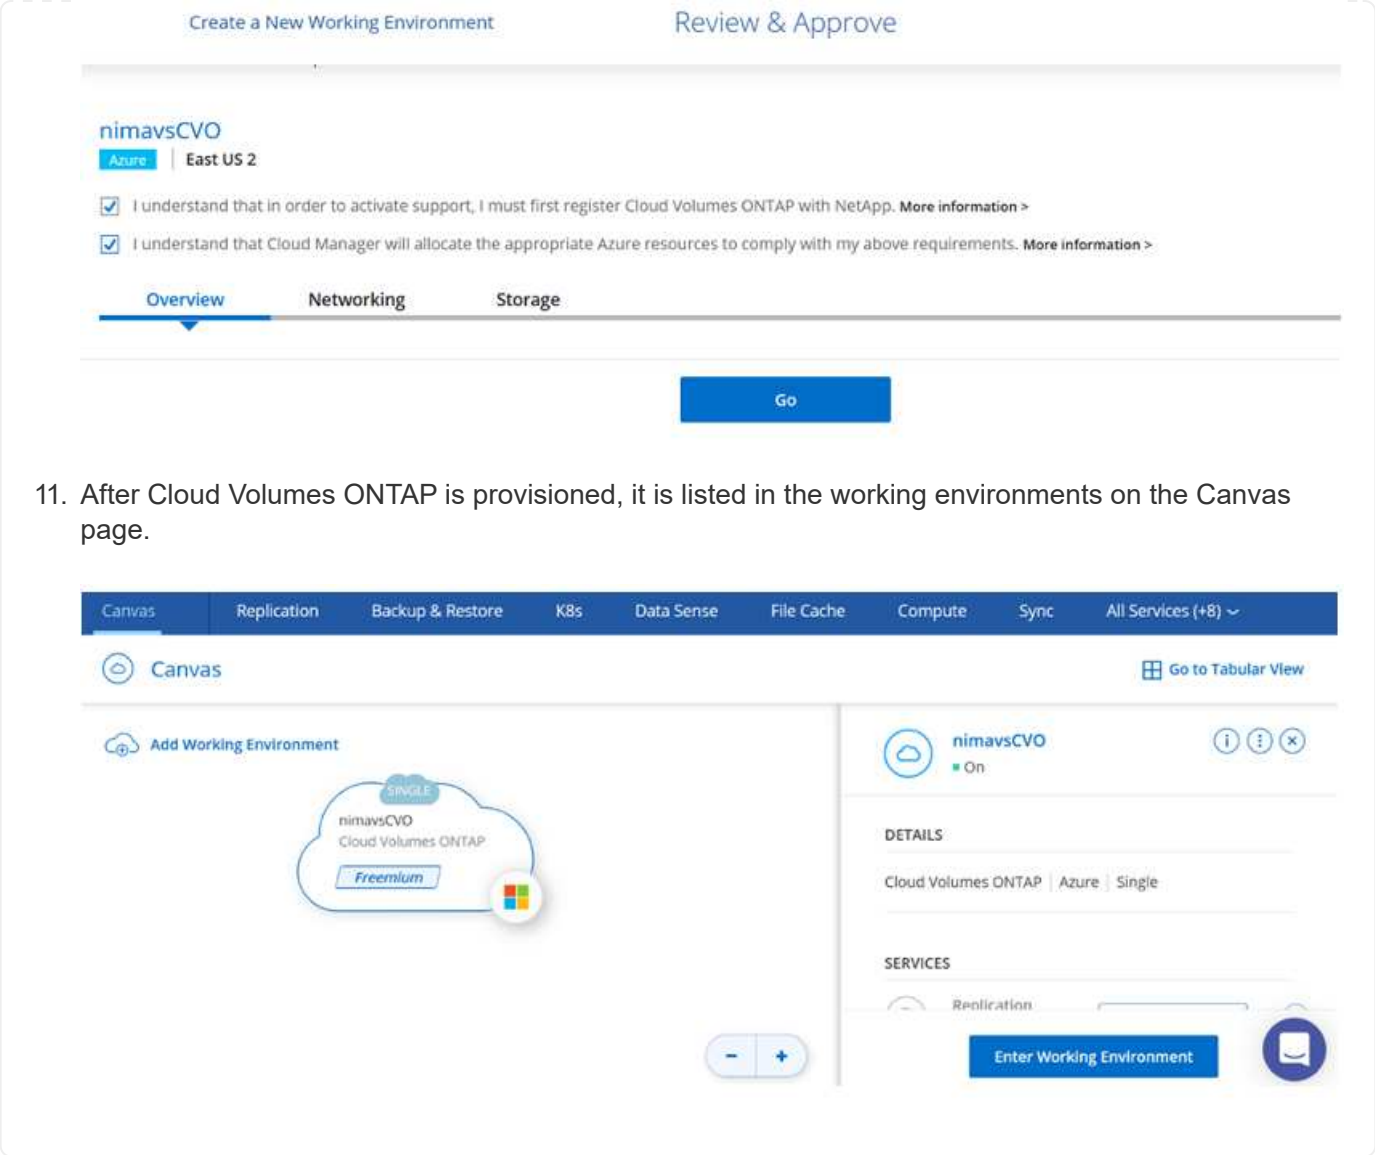

11. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.

Additional configurations for SMB volumes

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

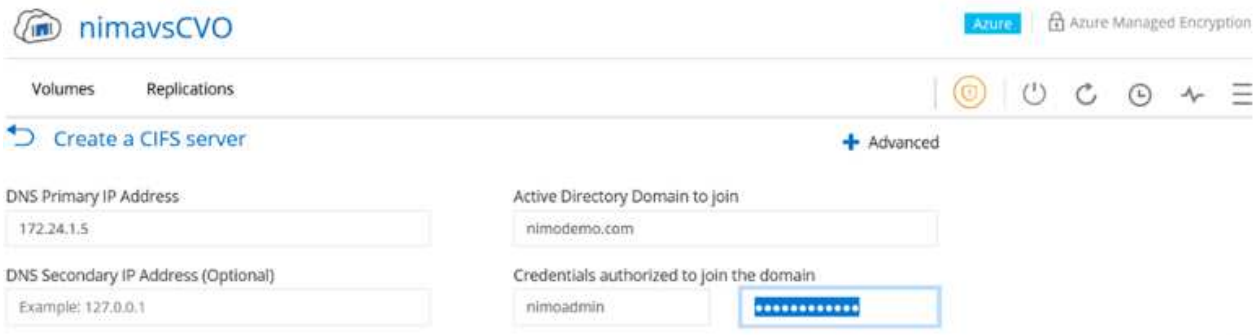

2. Creating the SMB volume is an easy process. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected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

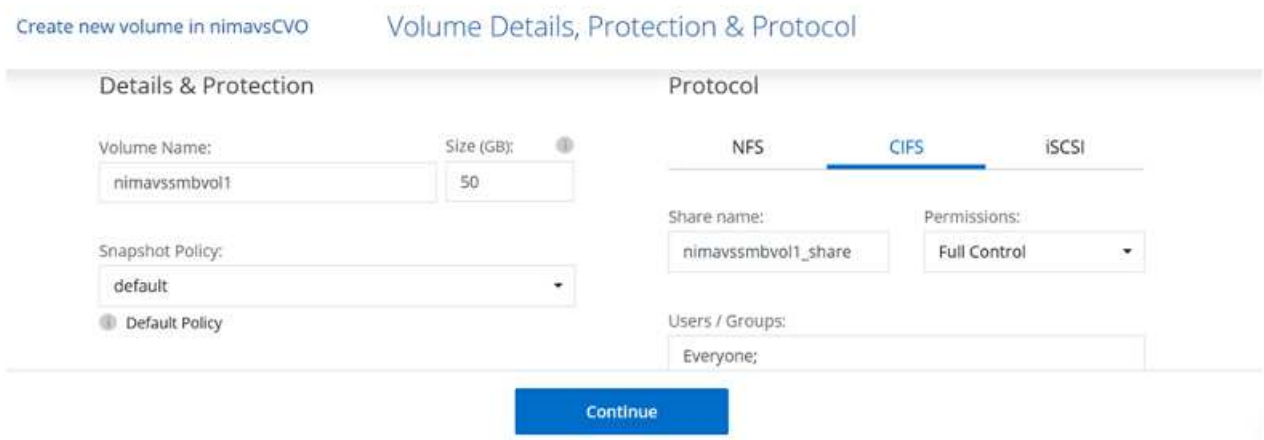

3. After the volume is provisioned, it will be availabe under the Volumes pane. Because a CIFS share is provisioned, give your users or groups permission to the files and folders and verify that those users can access the share and create a file. This step is not required if the volume is replicated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

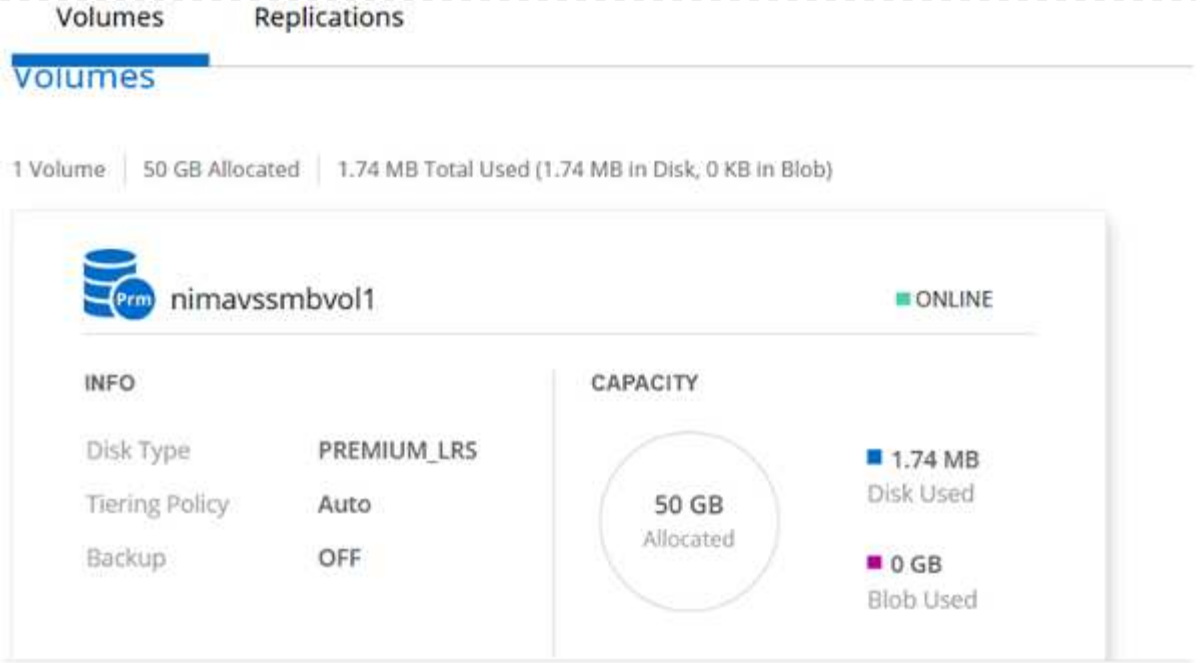

4. After the volume is created, use the mount command to connect to the share from the VM running on the Azure VMware Solution SDDC hosts.
5. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on Azure VMware Solution SDDC.

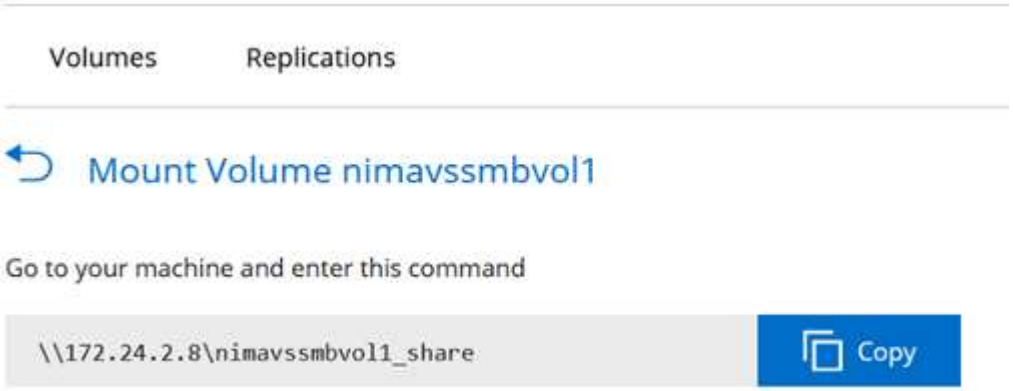

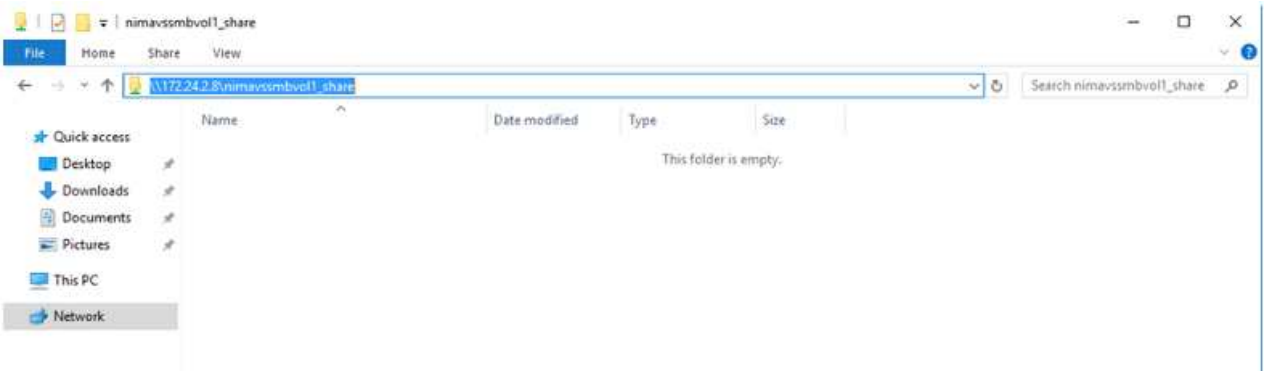

Connect the LUN to a host

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

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

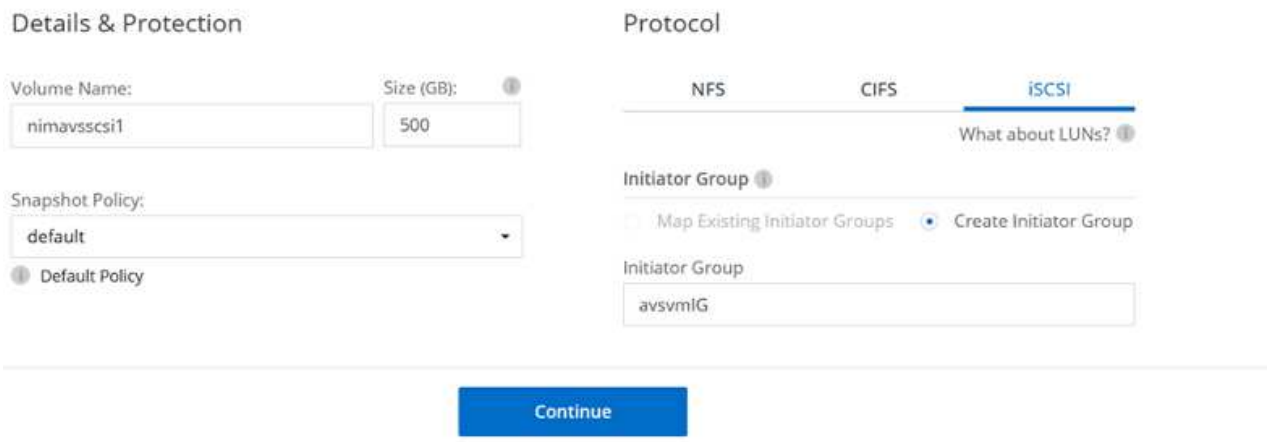

3. After the volume is provisioned, select the volume, and then click Target IQN. To copy the iSCSI Qualified Name (IQN), click Copy. Set up an iSCSI connection from the host to the LUN.

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

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

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

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

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

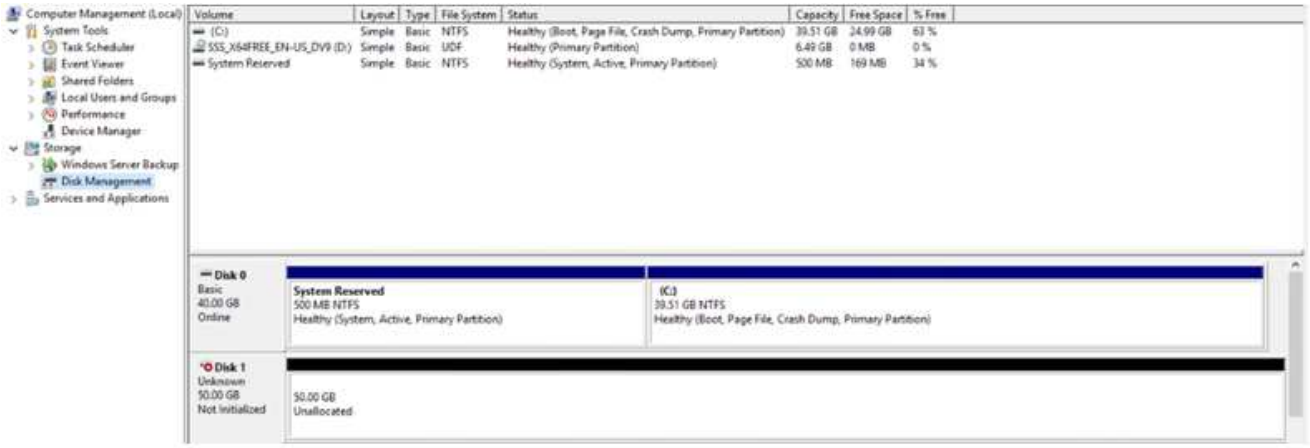

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

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

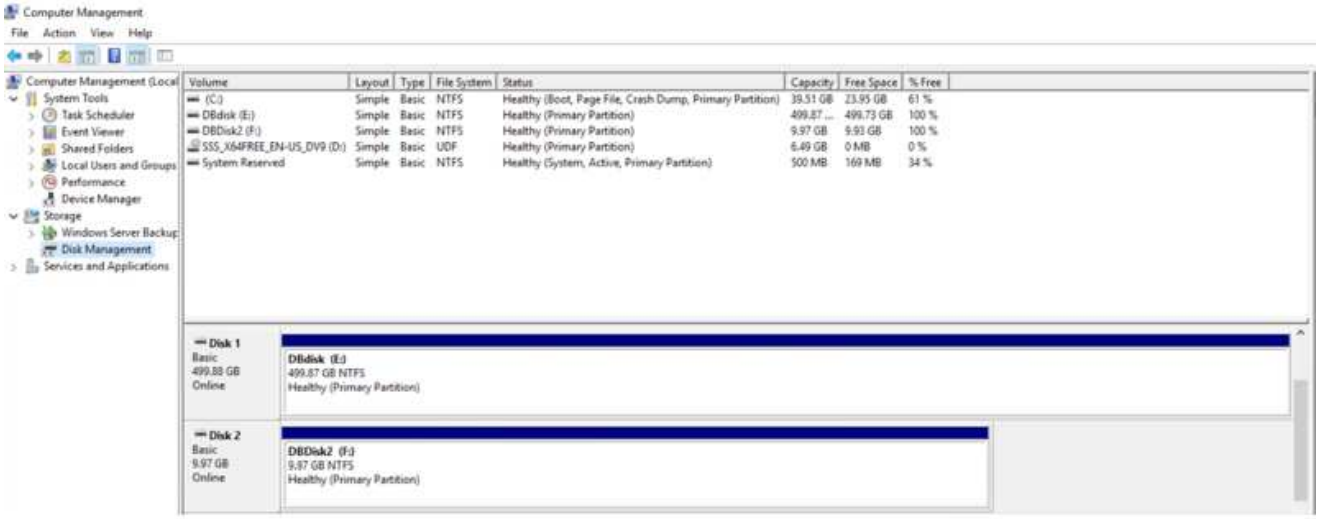

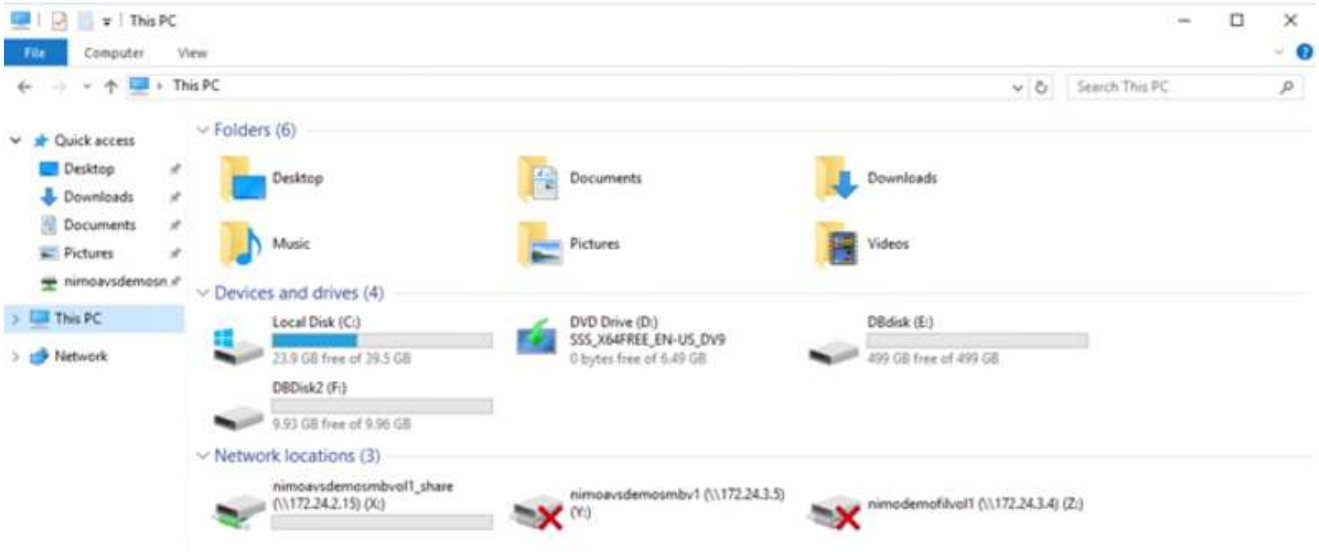

Google Cloud VMware Engine Supplemental NFS Datastore with NetApp Cloud Volume Service

Overview

Authors: Suresh Thoppay, NetApp

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

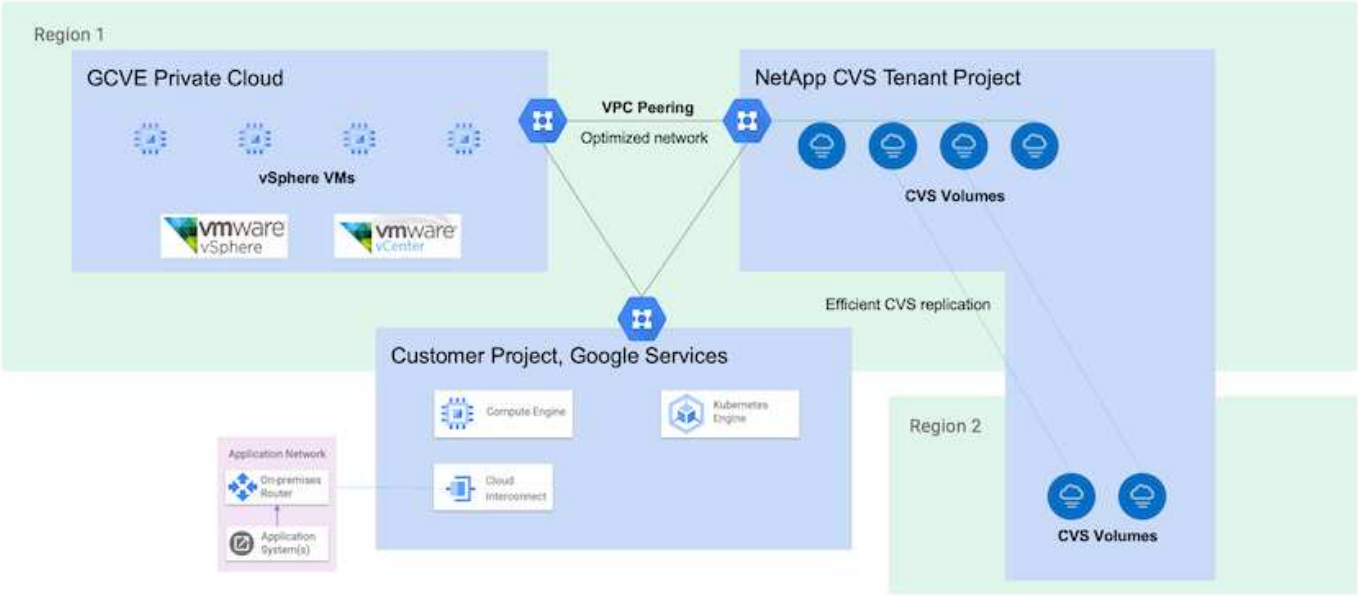

Deployment steps to mount NFS datastore from NetApp CVS on GCVE

Provision CVS-Performance Volume

The NetApp Cloud Volume Service volume can be either provisioned by
Using Google Cloud Console
Using NetApp BlueXP portal or API

Mark that CVS volume as non-deletable

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

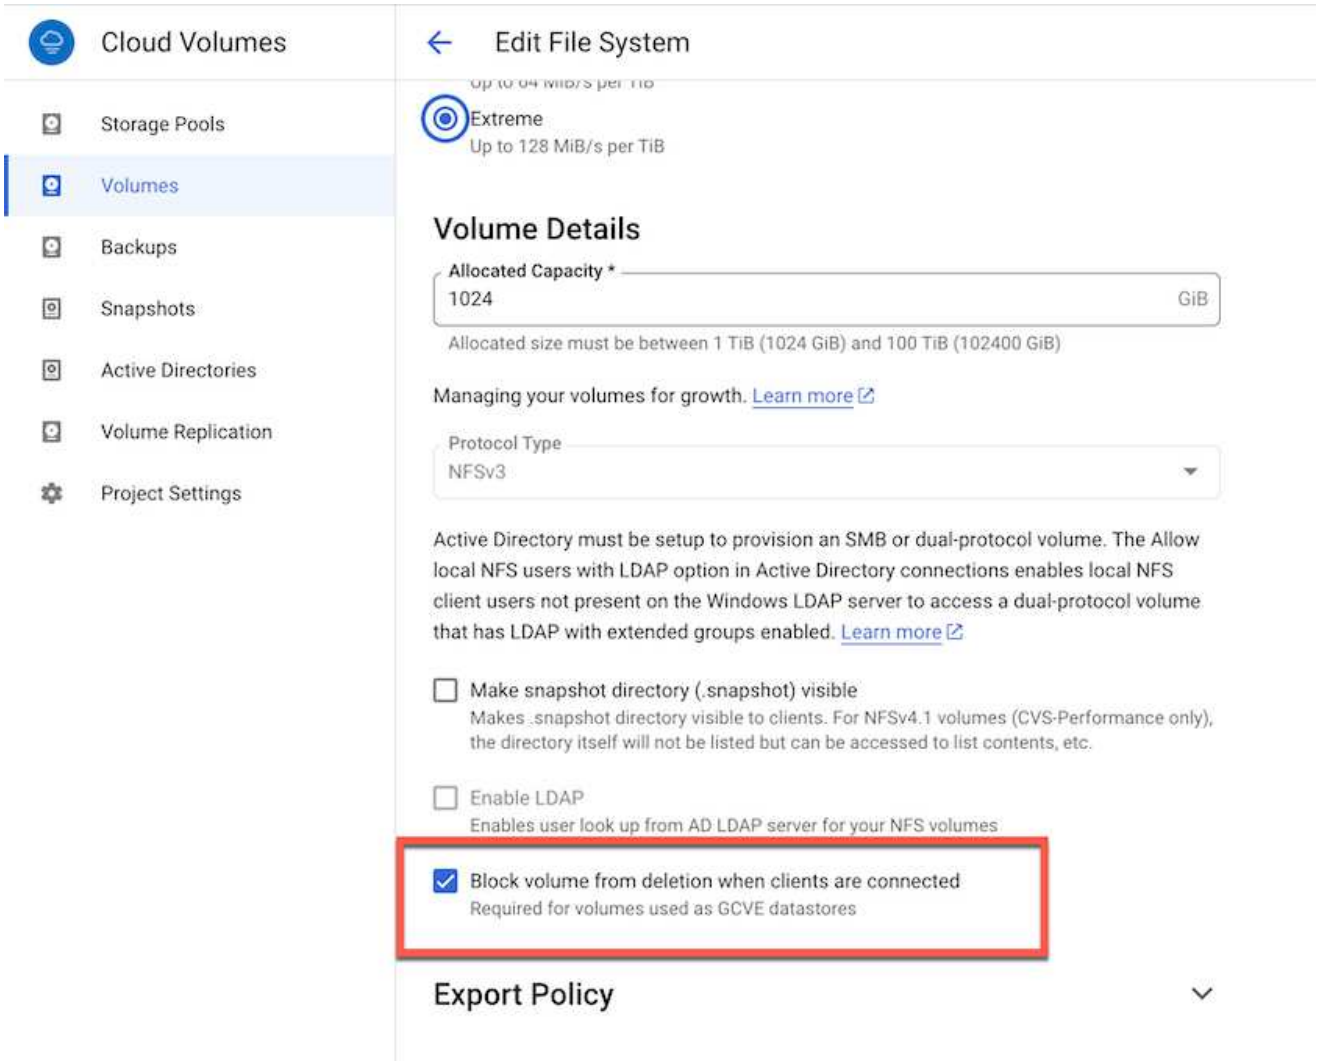

For more info, please refer Creating NFS Volume documentation.

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

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

> 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 NetApp Cloud Volume Service

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

Reference Links

- Google Blog - How to use NetApp CVS 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 Cloud Volumes Service (CVS).

Cloud Volumes ONTAP (CVO)

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

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

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

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

Cloud Volumes ONTAP (CVO) as guest connected storage

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

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

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

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

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

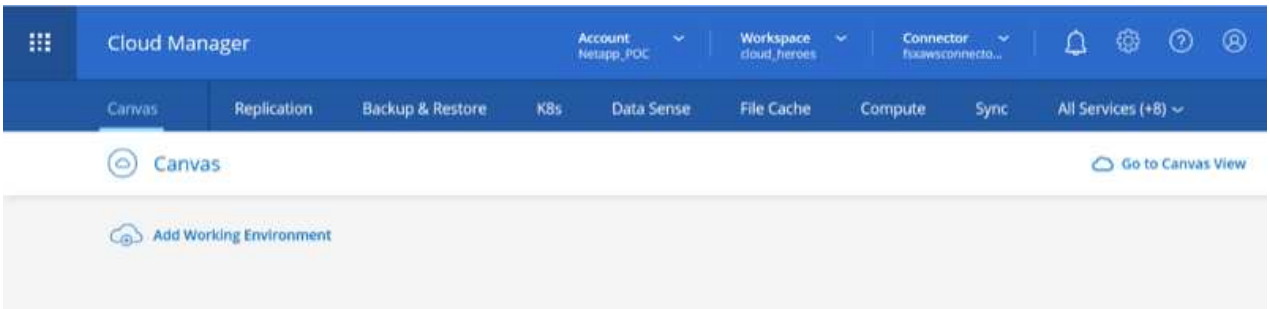

2. On the Cloud Manager Canvas tab, click Add a Working Environment and then select Google Cloud Platform as the cloud and the type of the system configuration. Then, click Next.

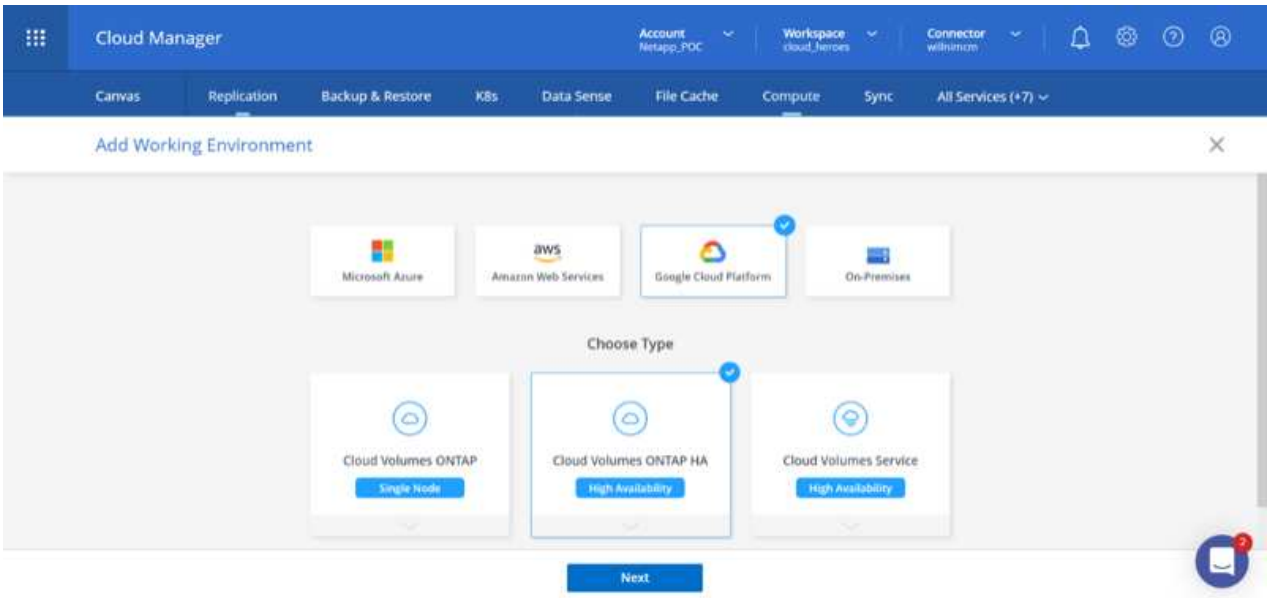

3. Provide the details of the environment to be created including the environment name and admin

credentials. After you are done, click Continue.

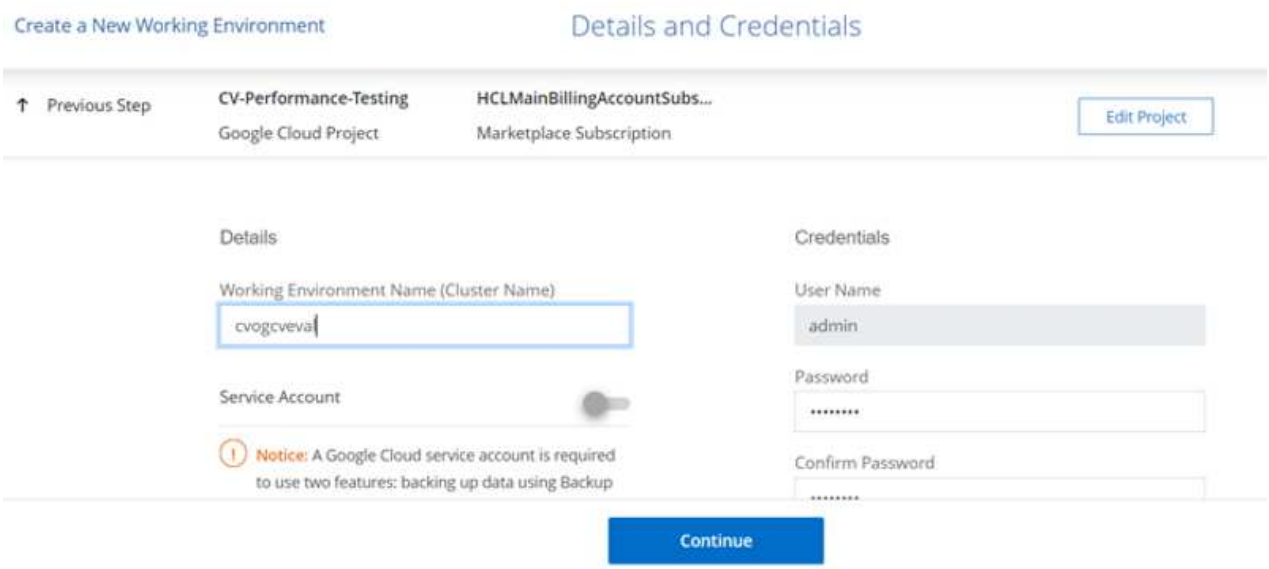

4. Select or deselect the add-on services for Cloud Volumes ONTAP deployment, including Data Sense & Compliance or Backup to Cloud. Then, click Continue.

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

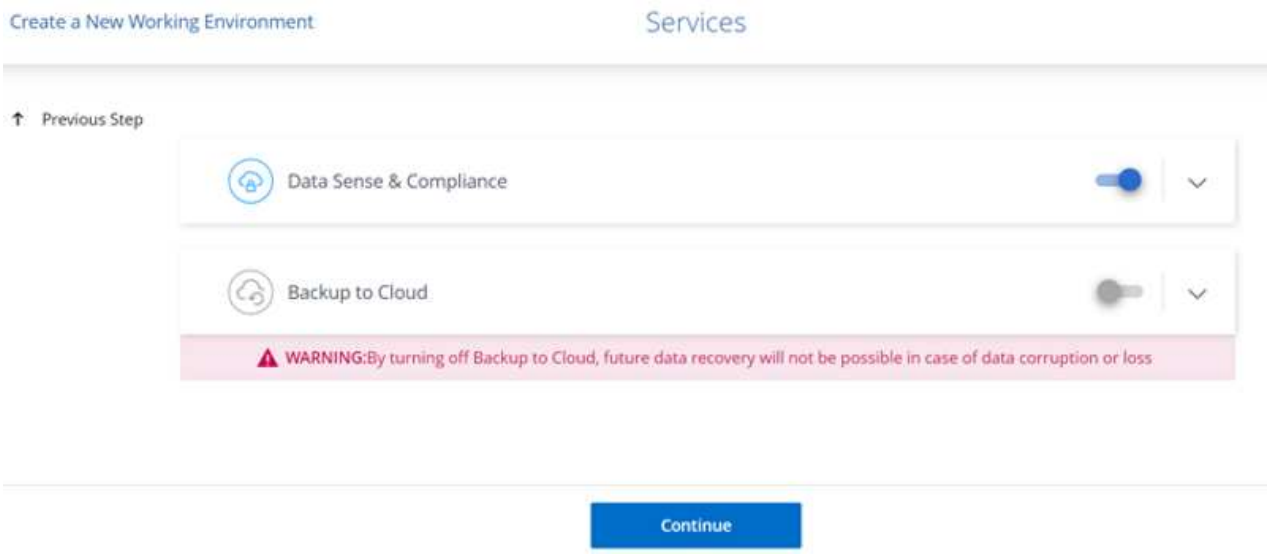

5. Select a location, choose a firewall policy, and select the checkbox to confirm network connectivity to Google Cloud storage.

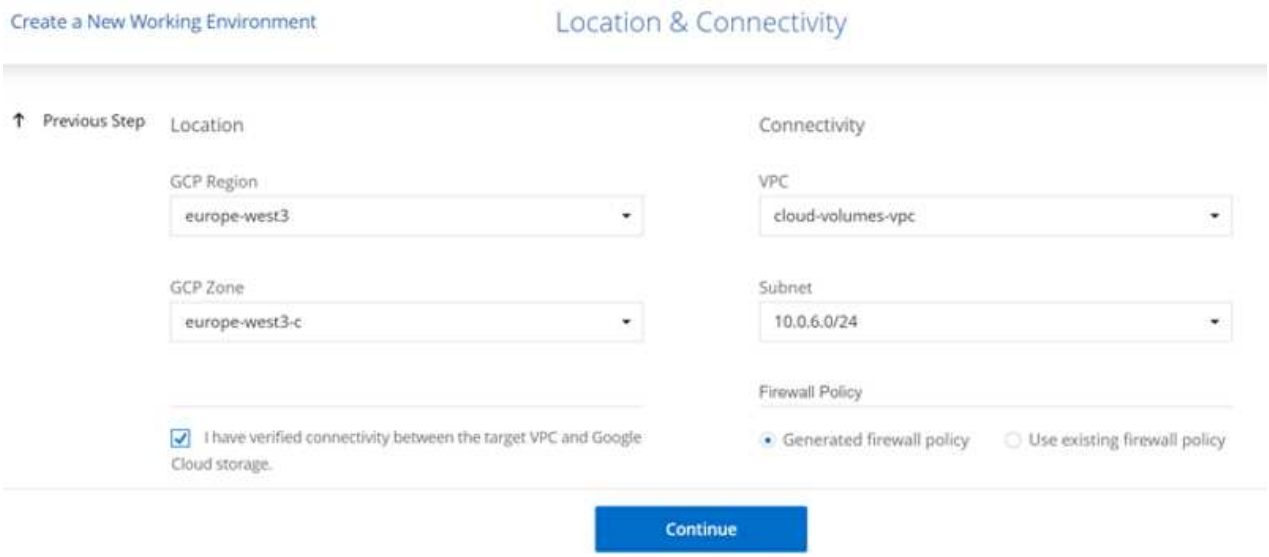

6. Select the license option: Pay-As-You-Go or BYOL for using existing license. In this example, Freemium option is used. Then, click on Continue.

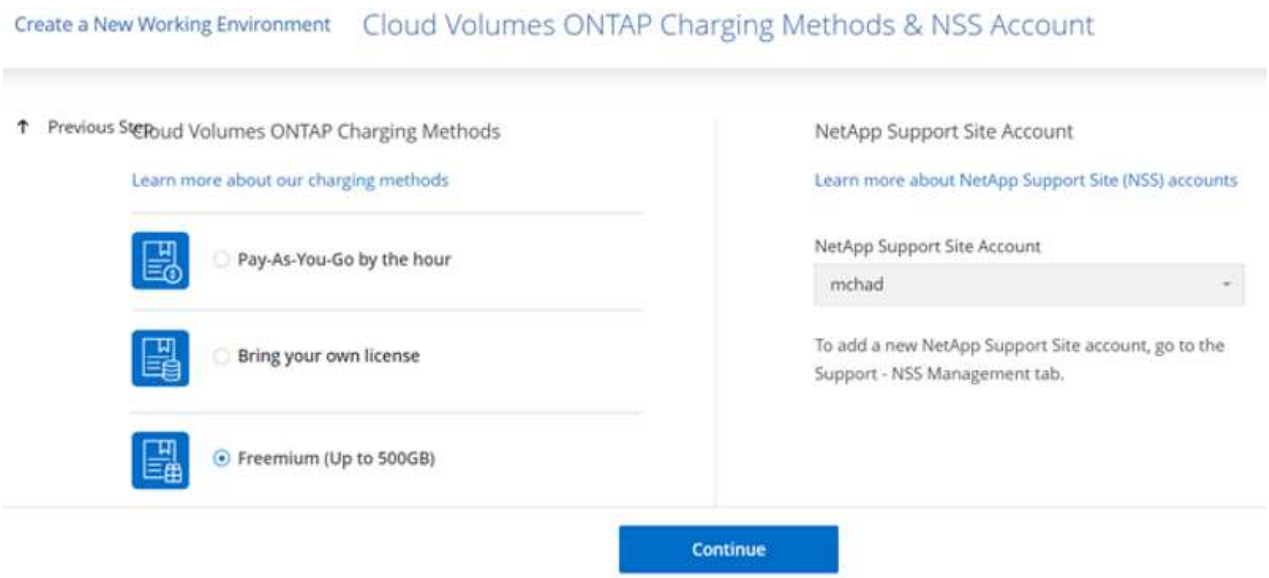

7. Select between several preconfigured packages available based on the type of workload that will be deployed on the VMs running on VMware cloud on AWS SDDC.

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

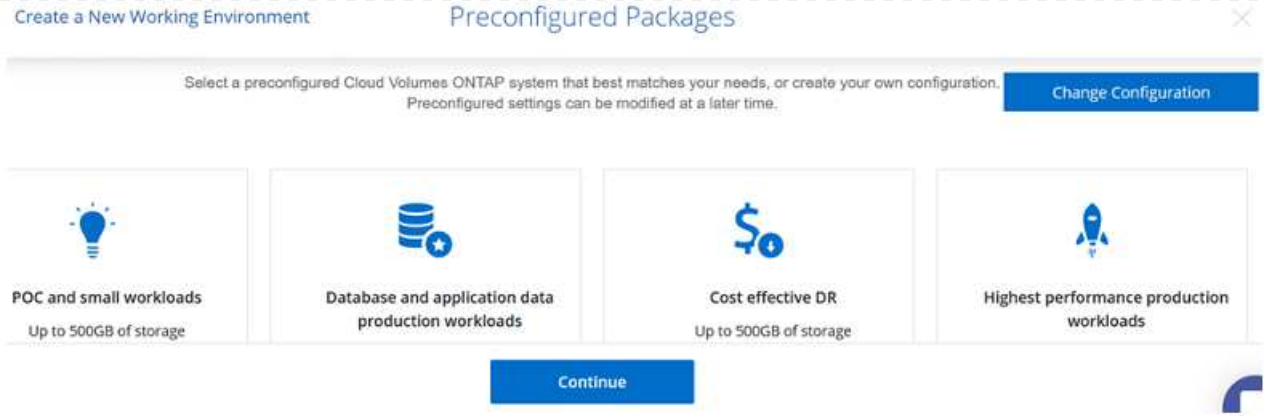

8. On the Review & Approve page, review and confirm the selections.To create the Cloud Volumes ONTAP instance, click Go.

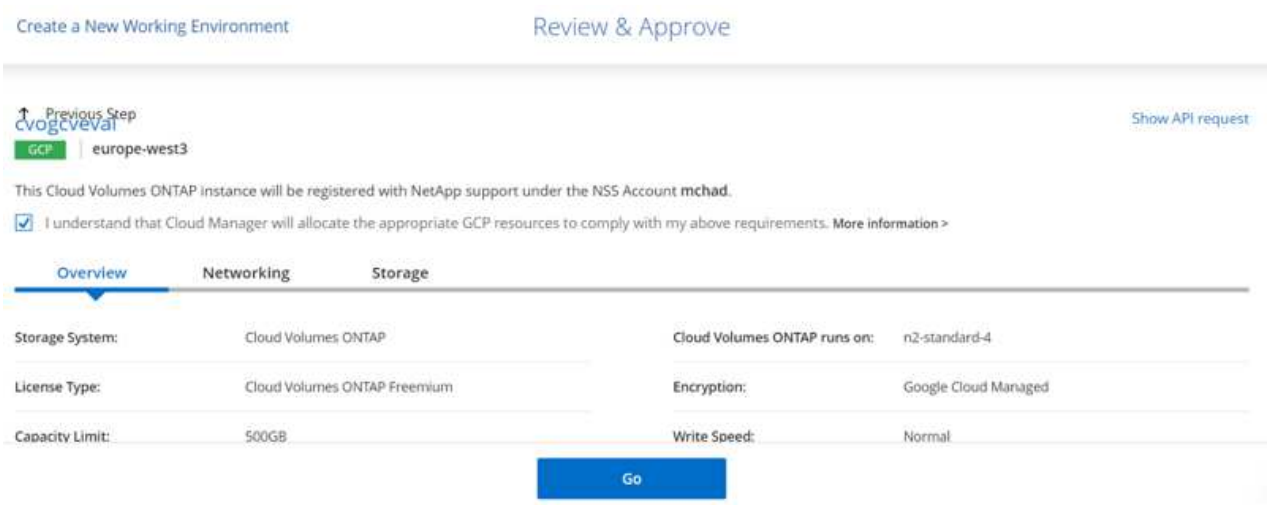

9. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.

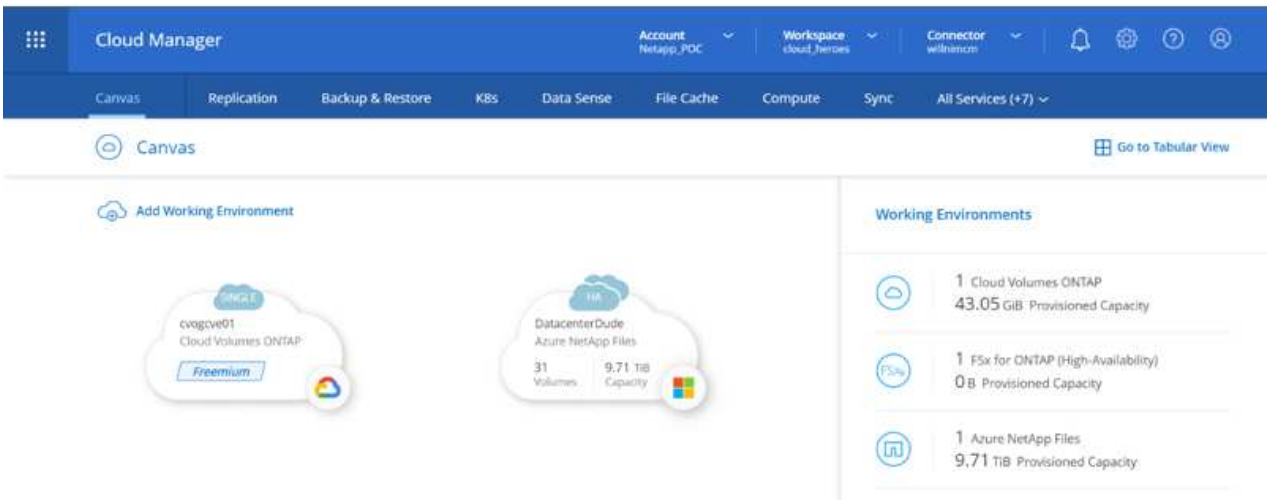

Additional configurations for SMB volumes

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

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

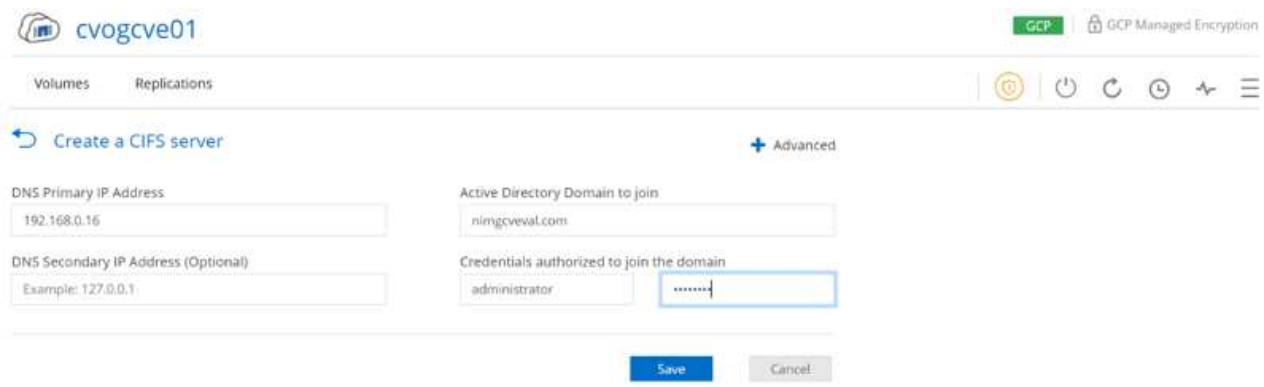

2. Creating the SMB volume is an easy process. At Canvas, double-click the Cloud Volumes ONTAP working environment to create and manage volumes and click on the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to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

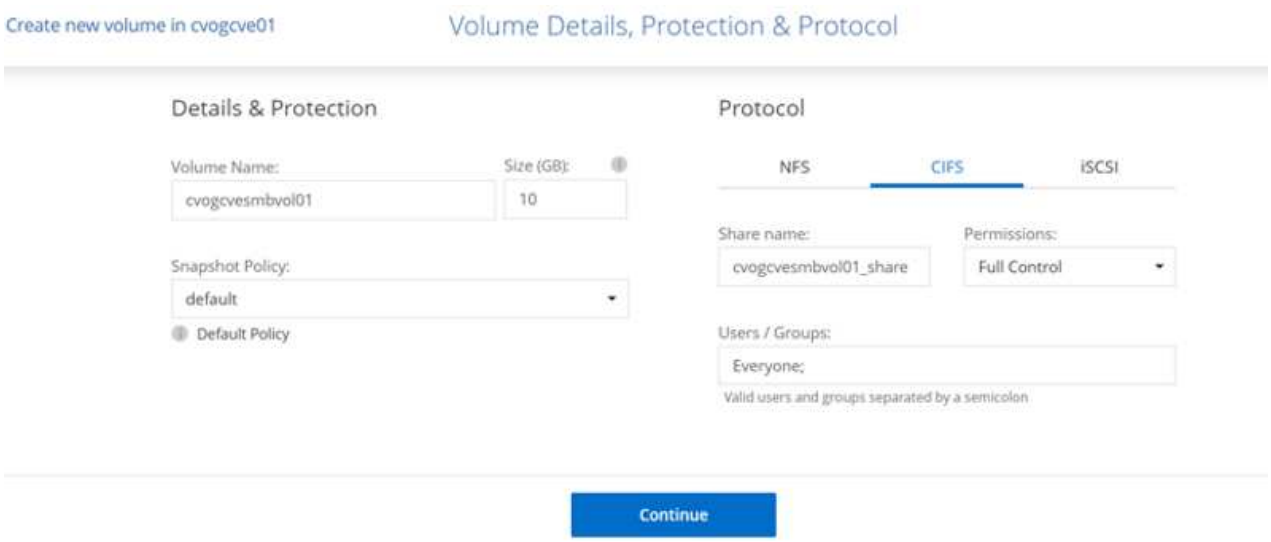

3. After the volume is provisioned, it will be availabe under the Volumes pane. Because a CIFS share is provisioned, give your users or groups permission to the files and folders and verify that those users can access the share and create a file. This step is not required if the volume is replicated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

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

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

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

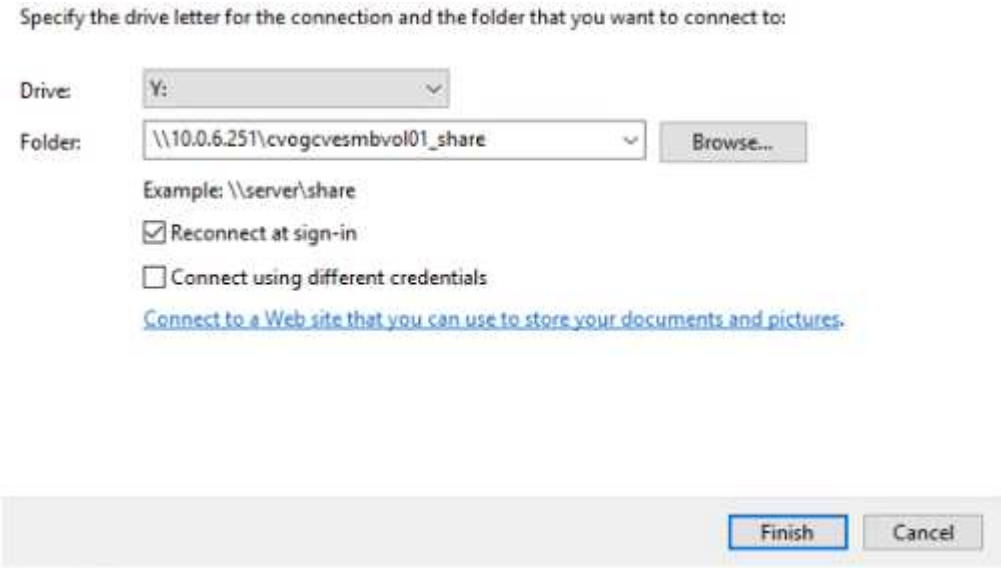

Once mapped, it can be easily accessed, and the NTFS permissions can be set accordingly.

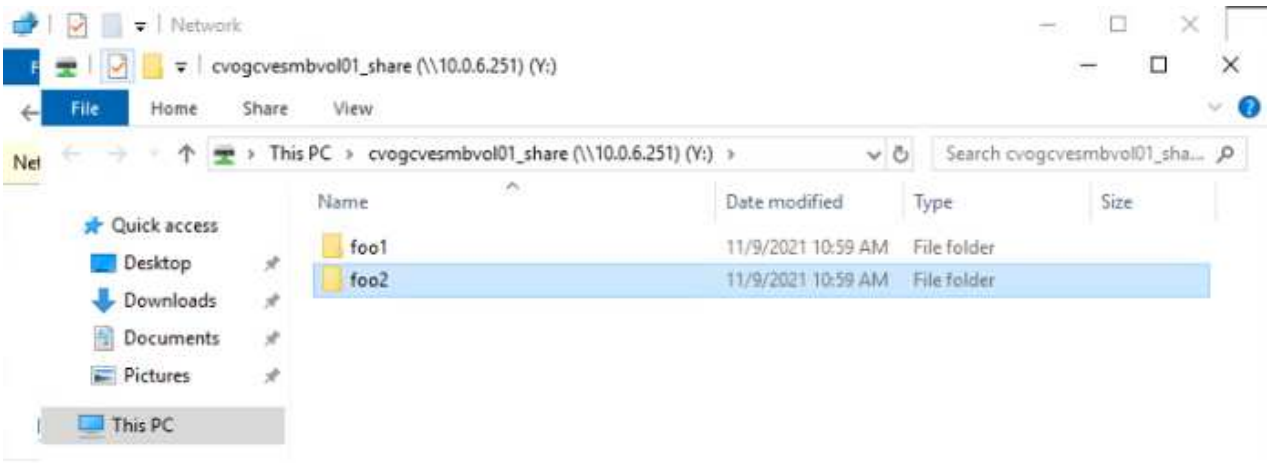

Connect the LUN on Cloud Volumes ONTAP to a host

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

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

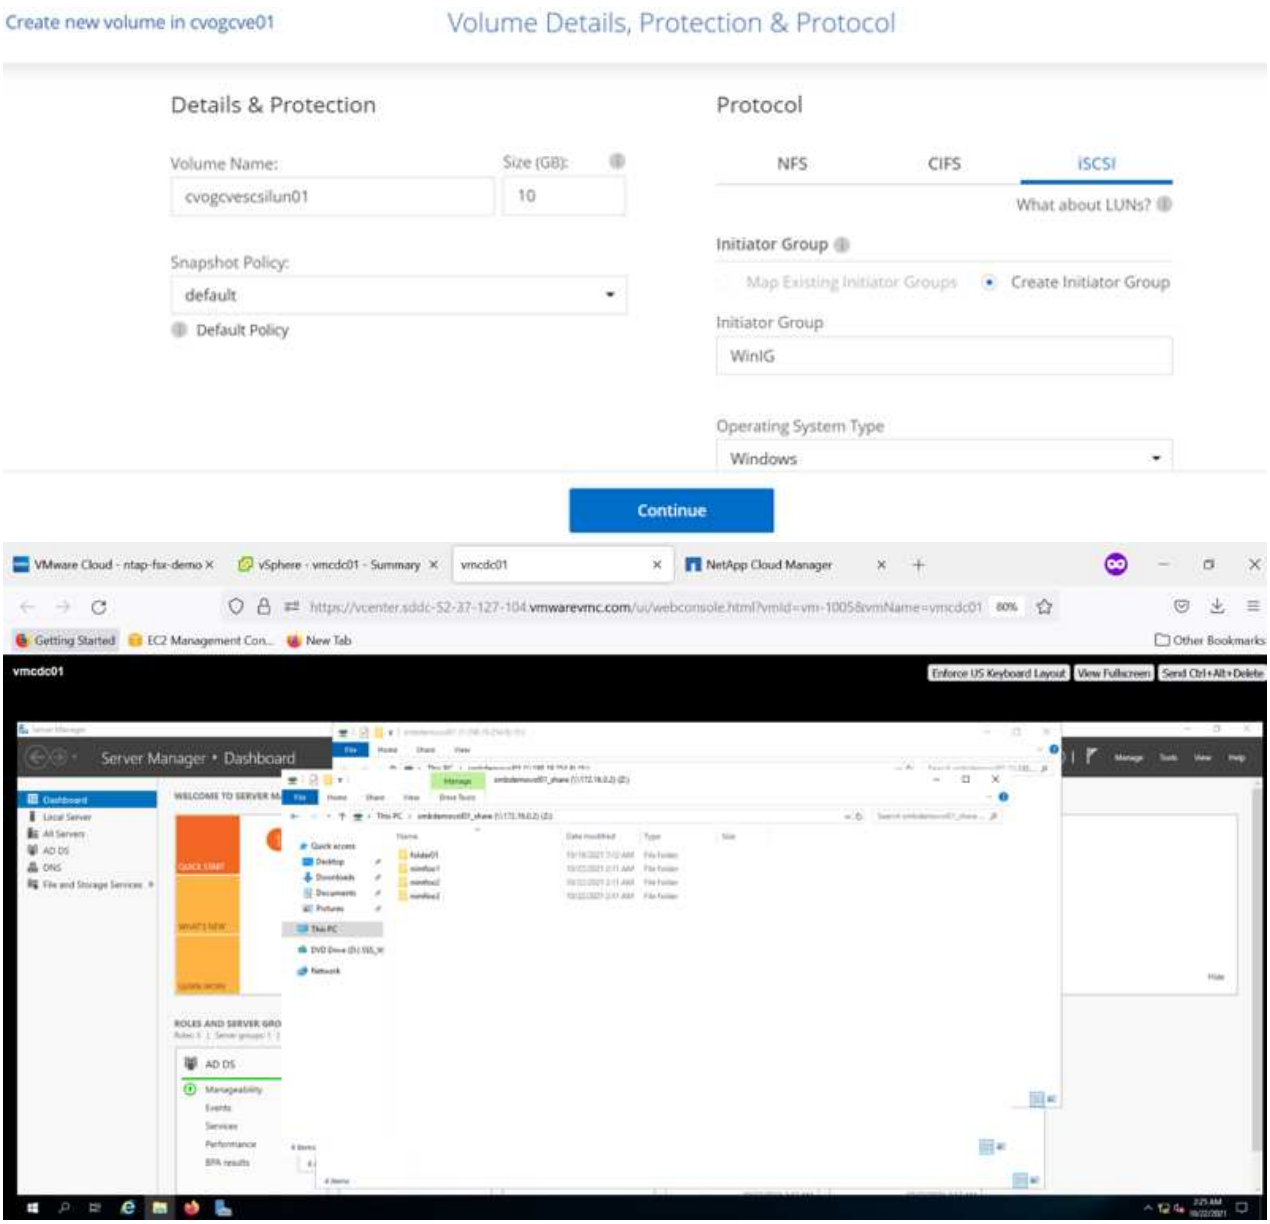

3. After the volume is provisioned, select the volume menu (º), and then click Target iQN. To copy the iSCSI Qualified Name (iQN), click Copy. Set up an iSCSI connection from the host to the LUN.

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

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

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

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

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

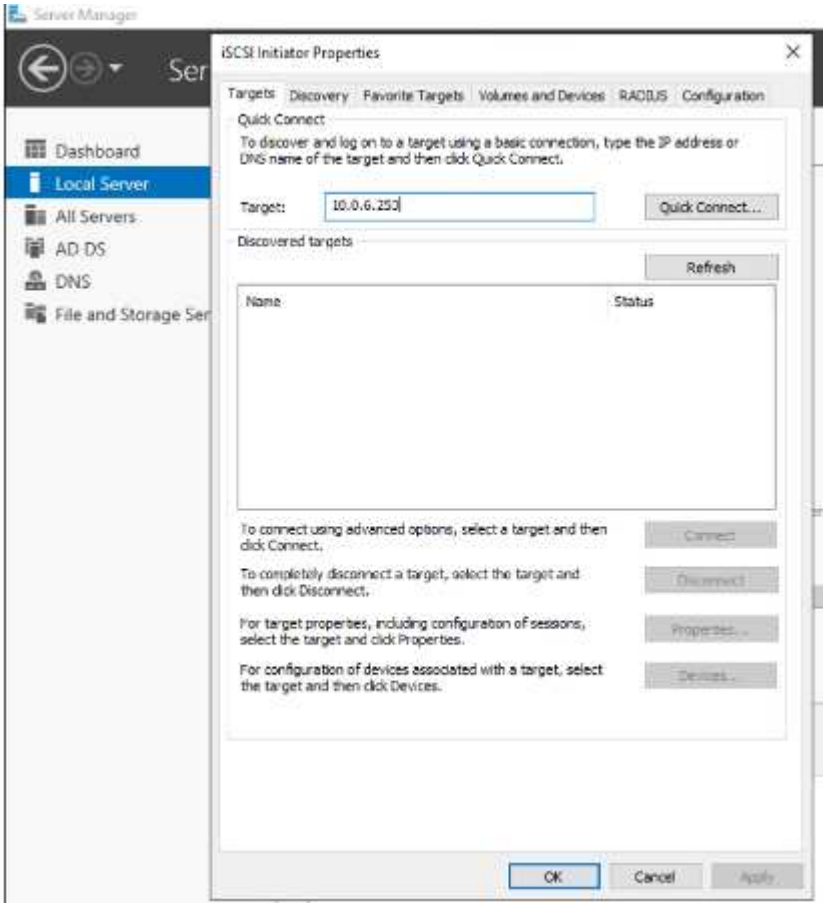

LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

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

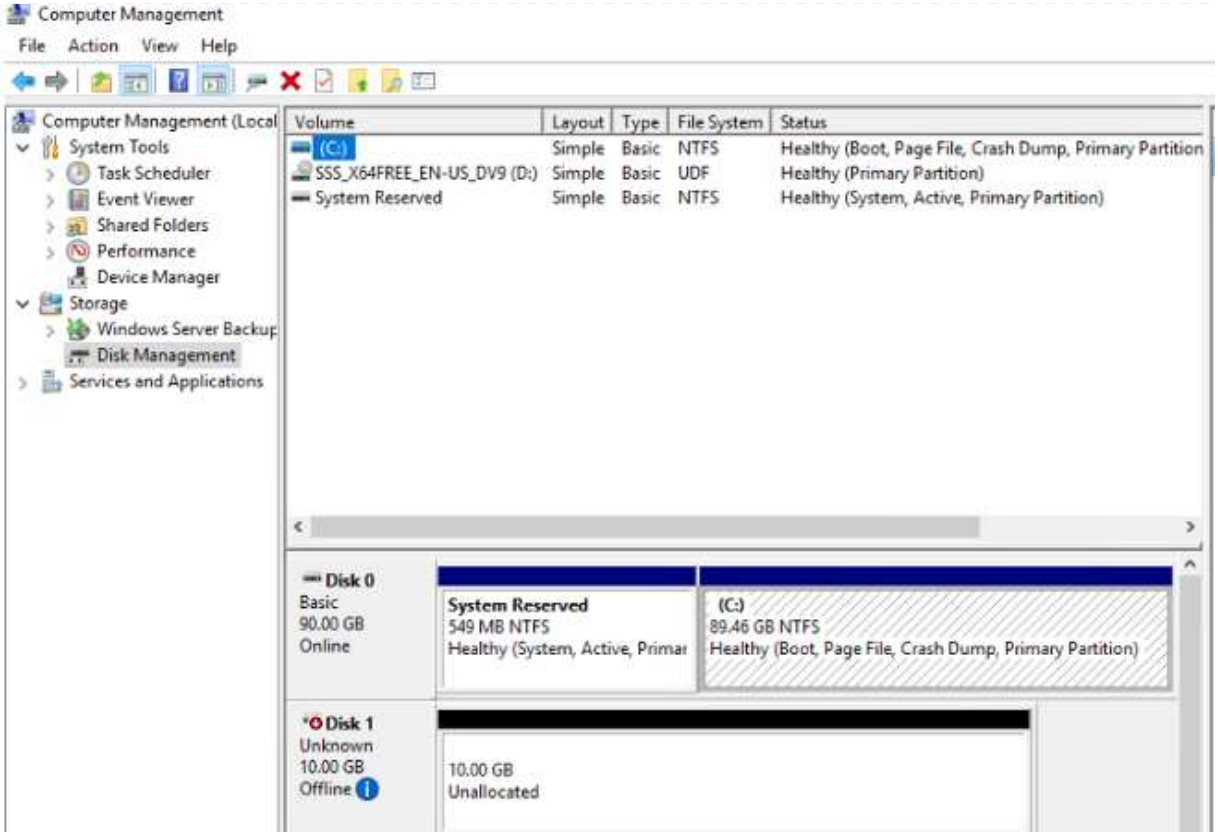

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

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

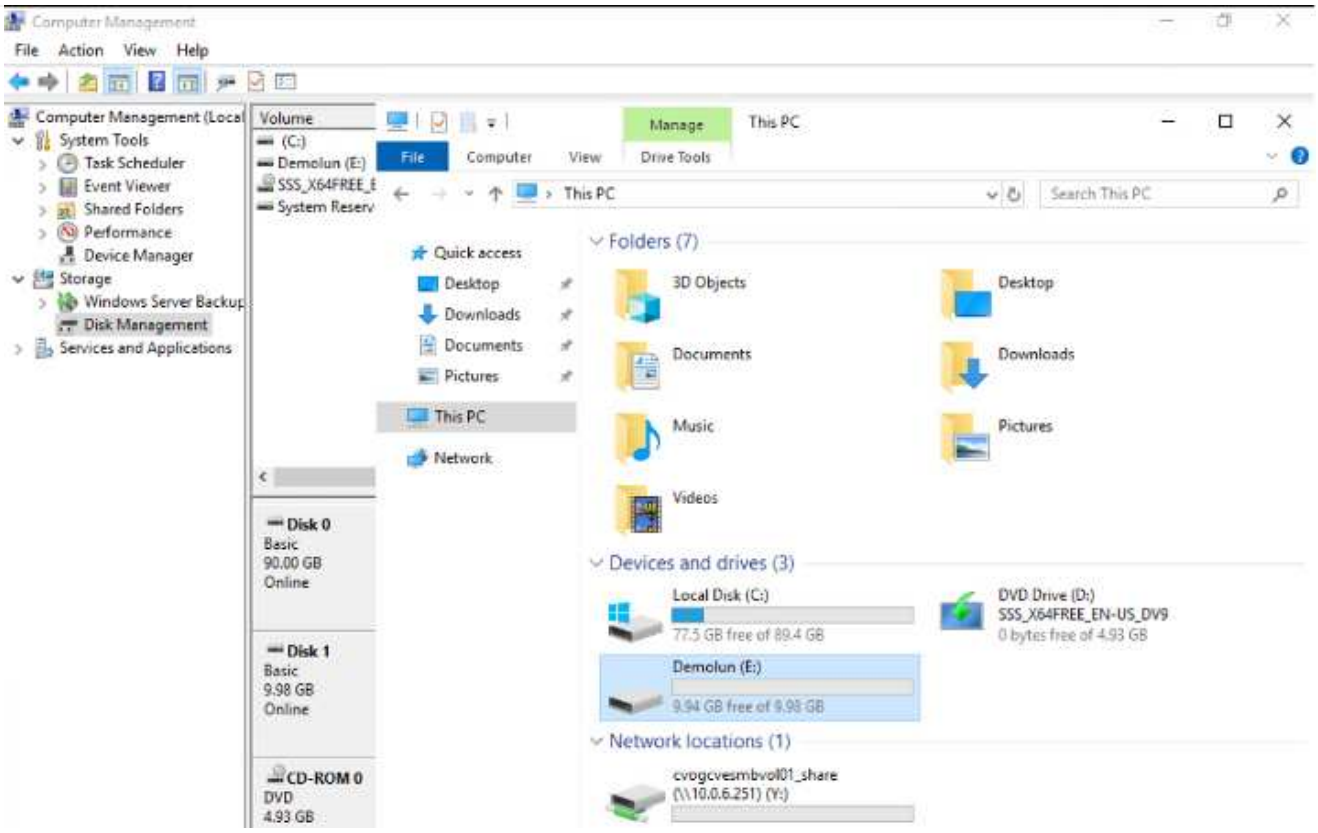

On the Linux clients, ensure the iSCSI daemon is running. Once the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration with Ubuntu as an example here. To verify, run lsblk cmd from the shell.

```
niyaz@nimubu01:~$ lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
loop0    7:0    0 55.4M  1 loop /snap/core18/2128
loop1    7:1    0  219M  1 loop /snap/gnome-3-34-1804/72
loop2    7:2    0 65.1M  1 loop /snap/gtk-common-themes/1515
loop3    7:3    0   51M  1 loop /snap/snap-store/547
loop4    7:4    0 32.3M  1 loop /snap/snapd/12704
loop5    7:5    0 32.5M  1 loop /snap/snapd/13640
loop6    7:6    0 55.5M  1 loop /snap/core18/2246
loop7    7:7    0    4K  1 loop /snap/bare/5
loop8    7:8    0 65.2M  1 loop /snap/gtk-common-themes/1519
sda      8:0    0   16G  0 disk
├─sda1   8:1    0  512M  0 part /boot/efi
├─sda2   8:2    0    1K  0 part
└─sda5   8:5    0 15.5G  0 part /
sdb      8:16   0    1G  0 disk
```

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

Mount Cloud Volumes ONTAP NFS volume on Linux client

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

Provision the volume following the below steps

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

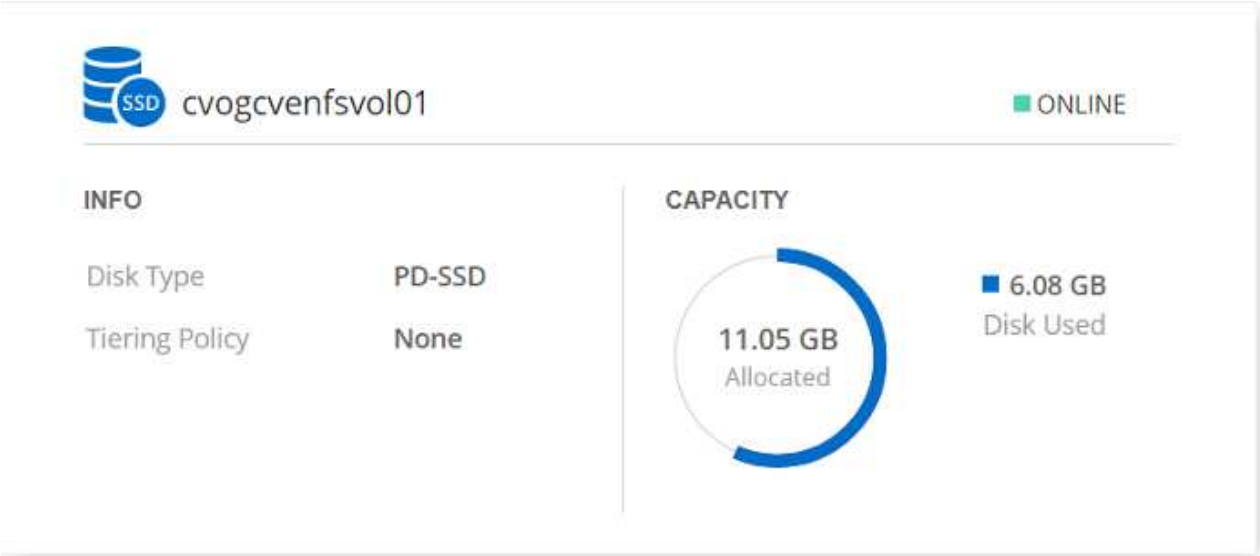

3. In the Volumes tab, place your mouse cursor over the volume, select the menu icon (º), and then click Mount Command.

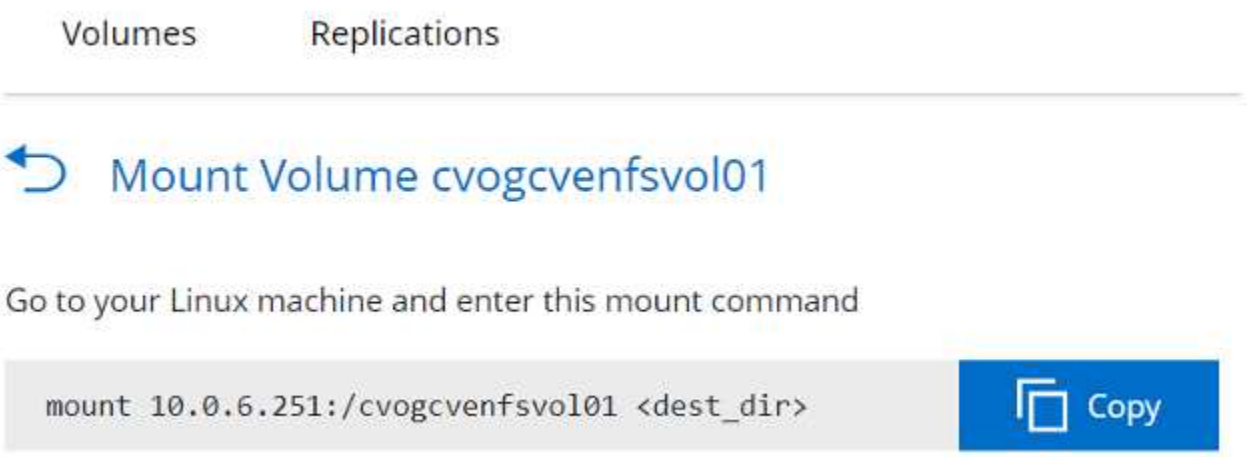

4. Click Copy.
5. Connect to the designated Linux instance.
6. Open a terminal on the instance using secure shell (SSH) and log in with the appropriate credentials.
7. Make a directory for the volume's mount point with the following command.

```
$ sudo mkdir /cvogcvetst
```

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

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

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

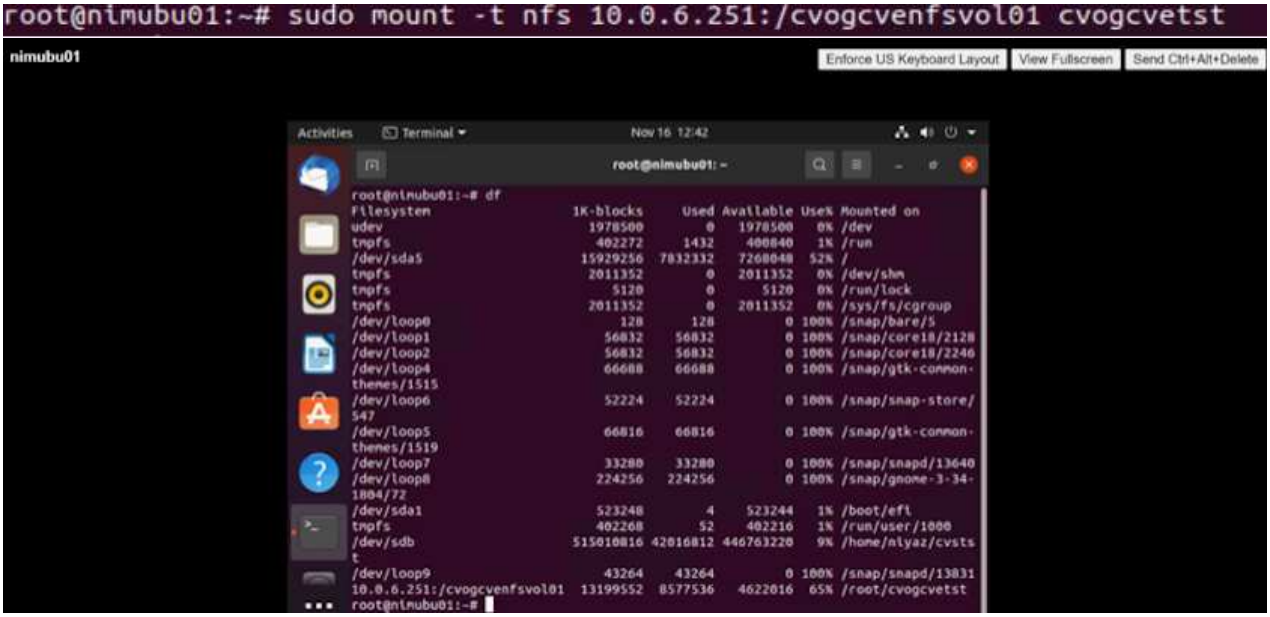

Cloud Volumes Service (CVS)

Cloud Volumes Services (CVS) is a complete portfolio of data services to deliver advanced cloud solutions. Cloud Volumes Services 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.

Cloud Volumes Service (CVS) as guest connected storage

Configure Cloud Volumes Service with VMware Engine

Cloud Volumes Service 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 Cloud Volumes Service supports SMB and NFS protocols. Cloud Volumes Service volumes can be set up in simple steps.

Cloud Volume Service and Google Cloud VMware Engine private cloud must be in the same region.

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

Create a CVS NFS volume to GCVE private cloud

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

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

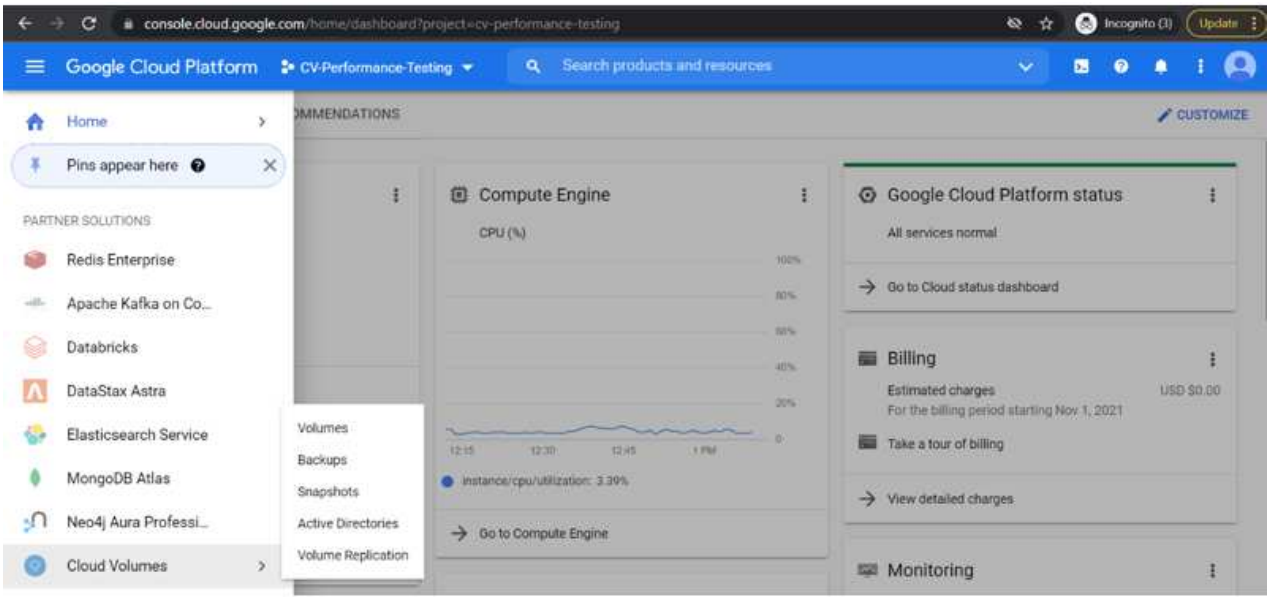

2. In the Cloud Volumes Console, go to the Volumes page and click Create.

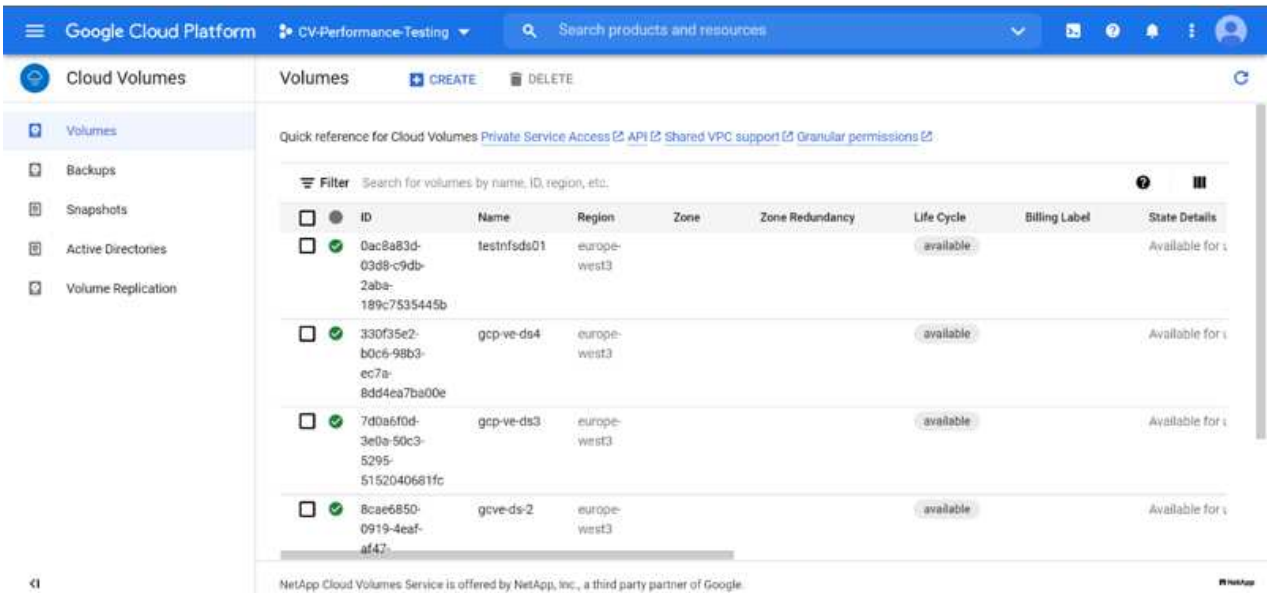

3. On the Create File System page, specify the volume name and billing labels as required for chargeback mechanisms.

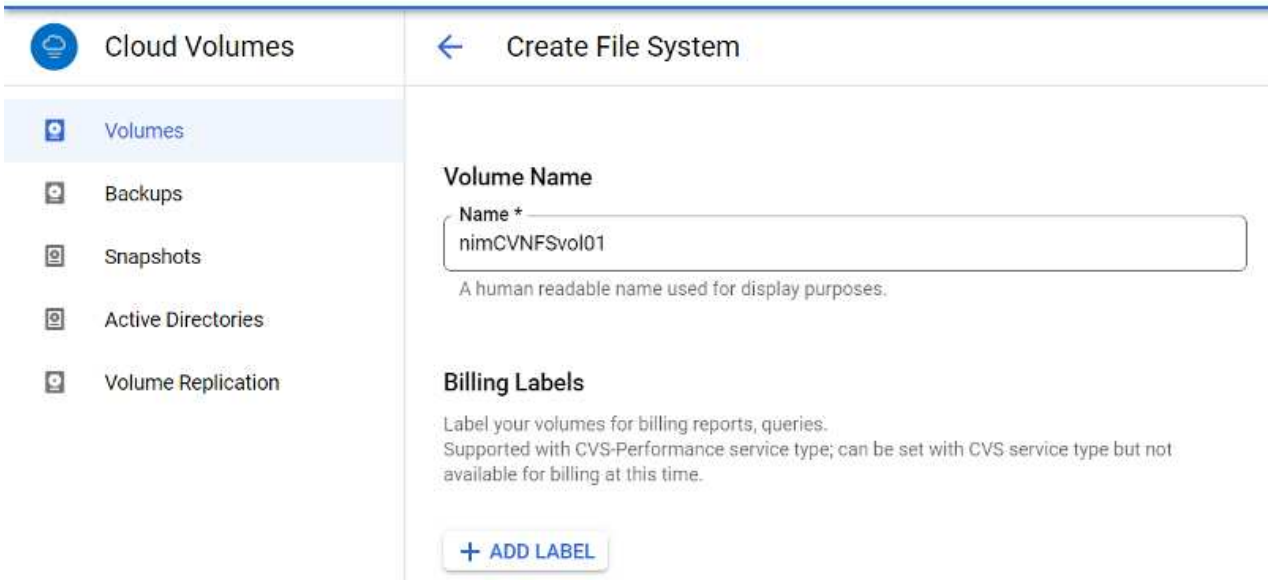

4. Select the appropriate service. For GCVE, choose CVS-Performance and desired service level for improved latency and higher performance based on the application workload requirements.

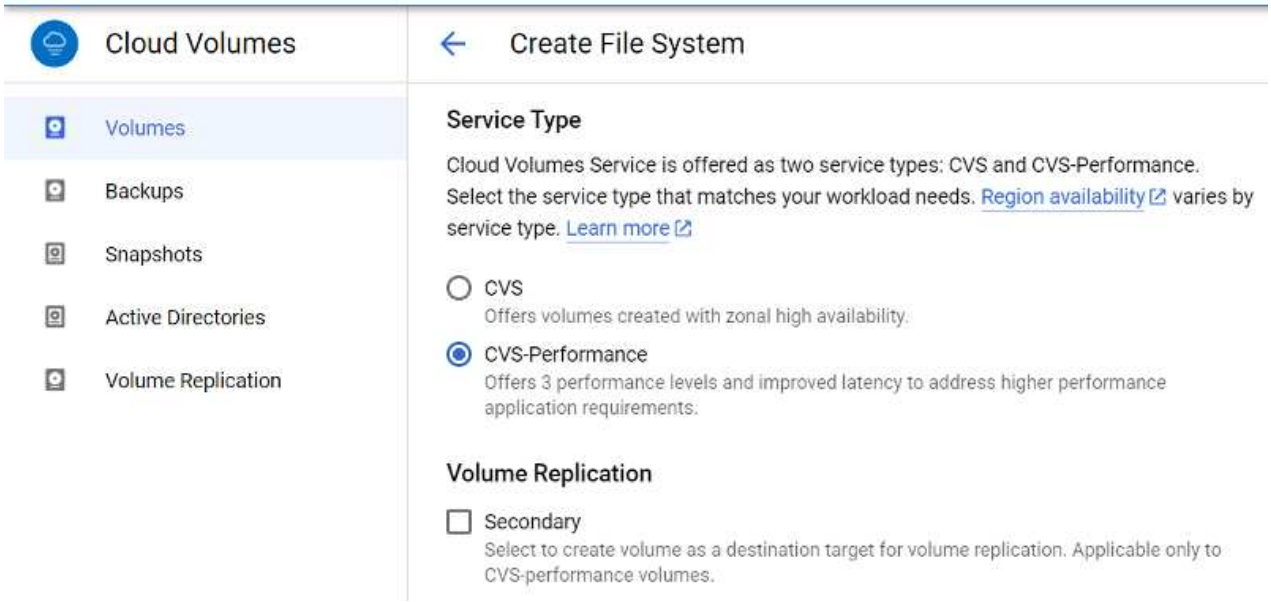

5. Specify the Google Cloud region for the volume and volume path (The volume path must be unique across all of cloud volumes in the project)

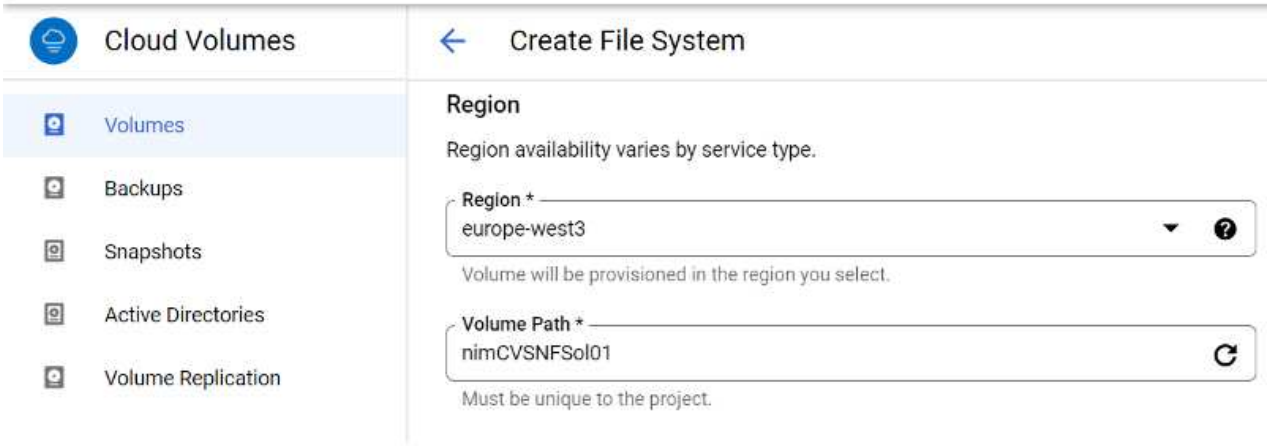

6. Select the level of performance for the volume.

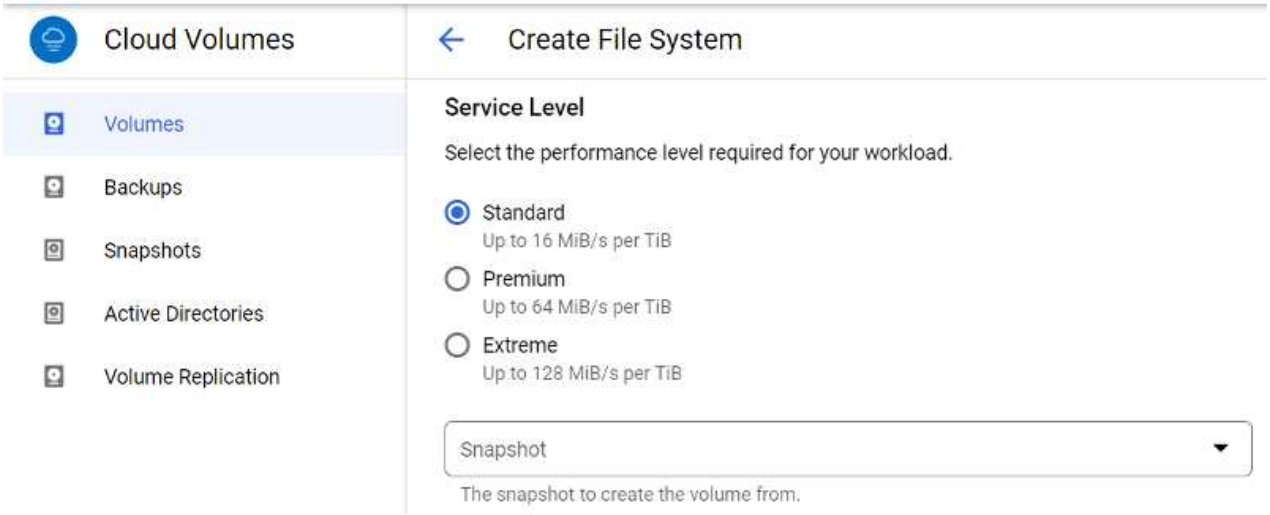

7. Specify the size of the volume and the protocol type. In this testing, NFSv3 is used.

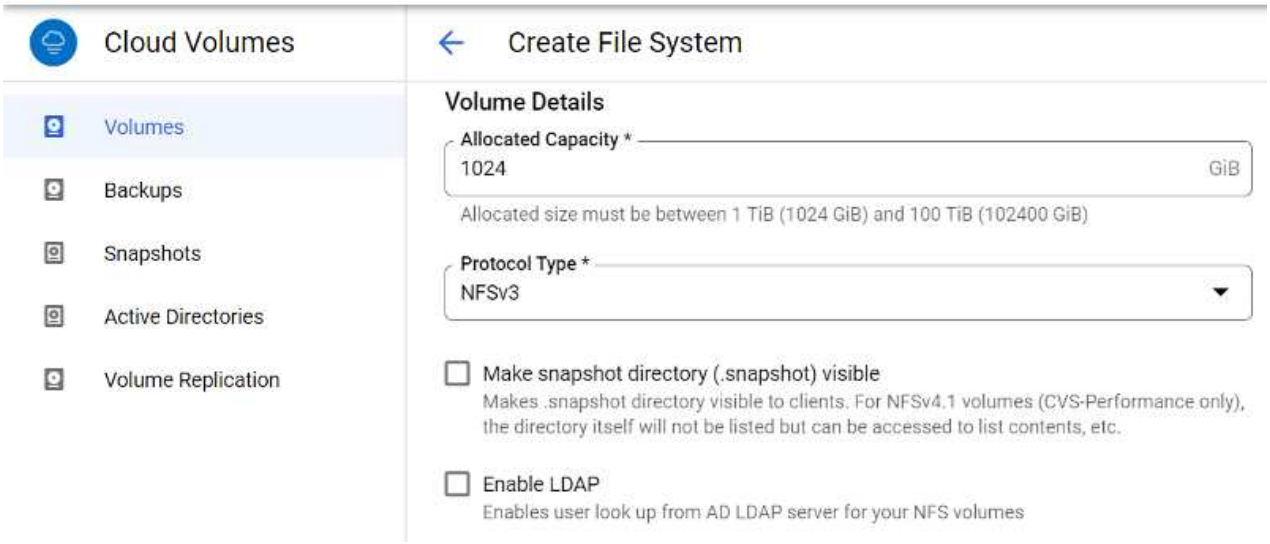

8. In this step, select the VPC Network from which the volume will be accessible. Ensure VPC peering is in place.

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

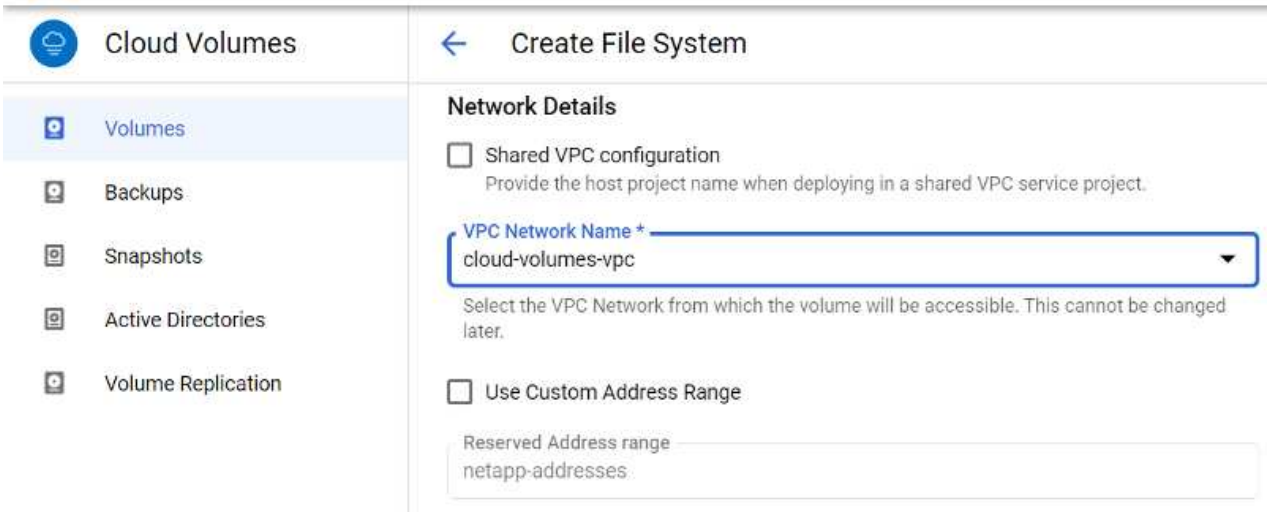

9. Manage the Export policy rules by adding the appropriate rules and Select the checkbox for the corresponding NFS version.

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

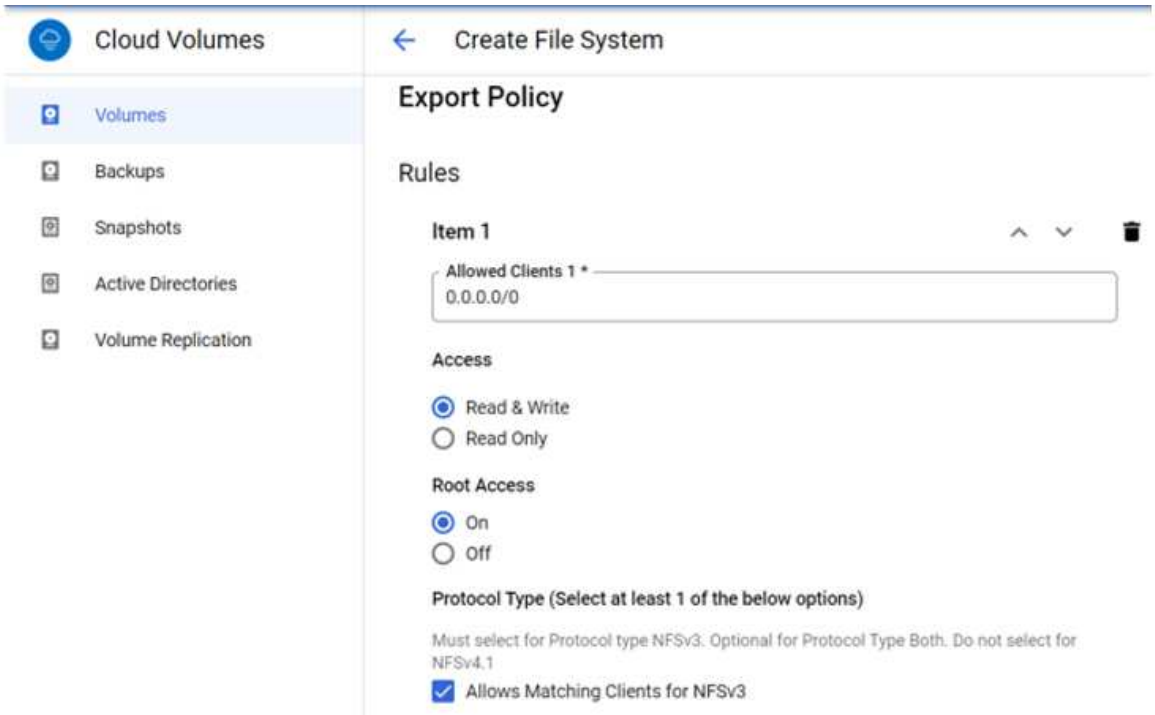

10. Click Save to create the volume.

Mounting NFS exports to VMs running on VMware Engine

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

To mount an NFS volume, do the following:

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

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

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

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

   b. On an Ubuntu or Debian instance:

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

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

   ```
   sudo mkdir /nimCVSNFSol01
   ```

   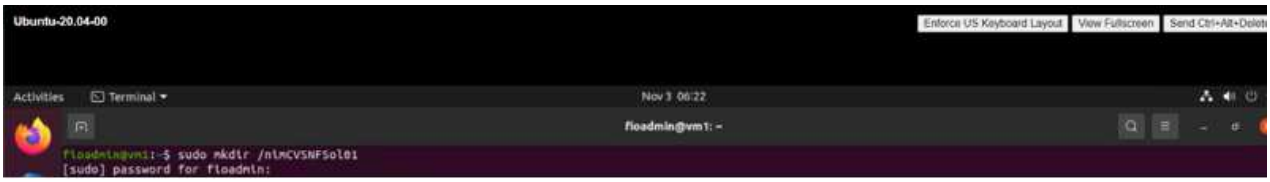

4. Mount the volume using the appropriate command. Example command from the lab is below:

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

   


```
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/mapper/nfsprdvg1-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:/nlmCVSNFSol01      107374182400        256 107374182144   1% /nlmCVSNFSol01
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.

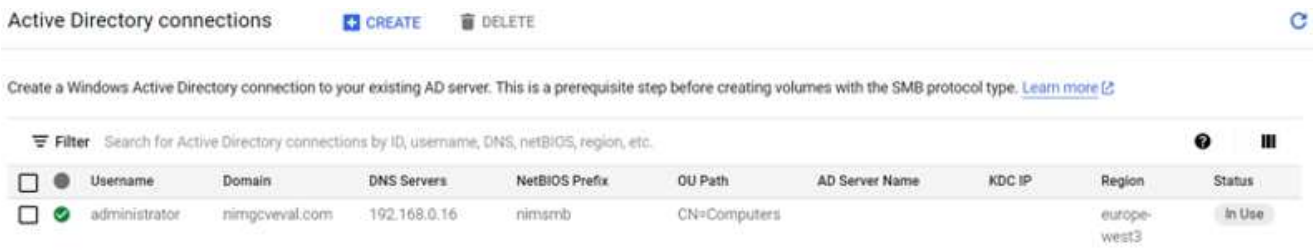

Once the AD connection is in place, create the volume with the desired service level. The steps are like creating NFS volume except selecting the appropriate protocol.

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

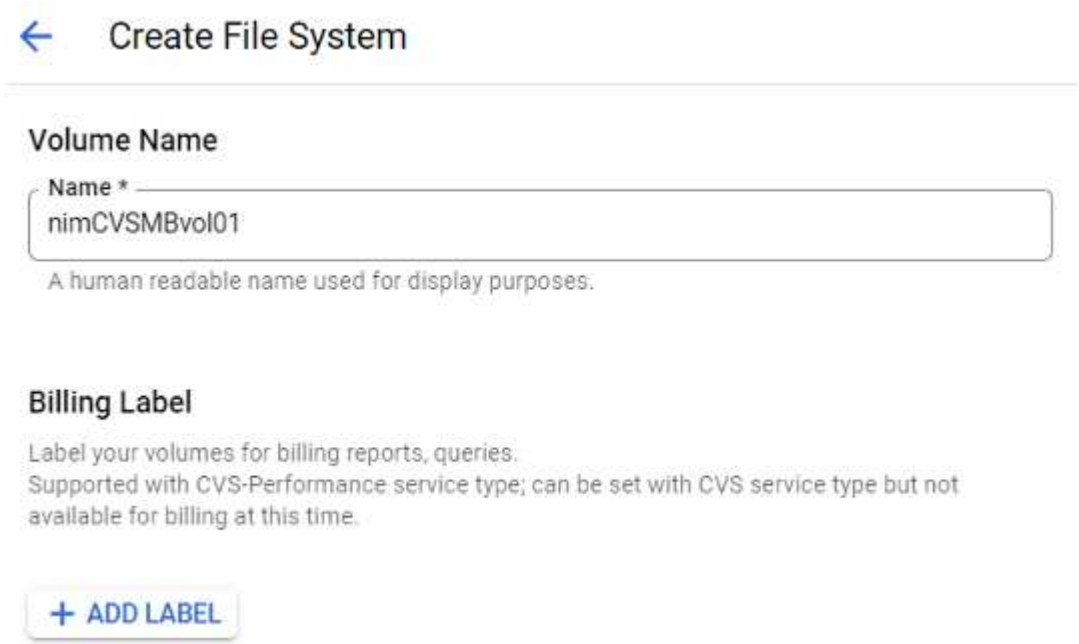

3. Select the appropriate service. For GCVE, choose CVS-Performance and desired service level for improved latency and higher performance based on the workload requirements.

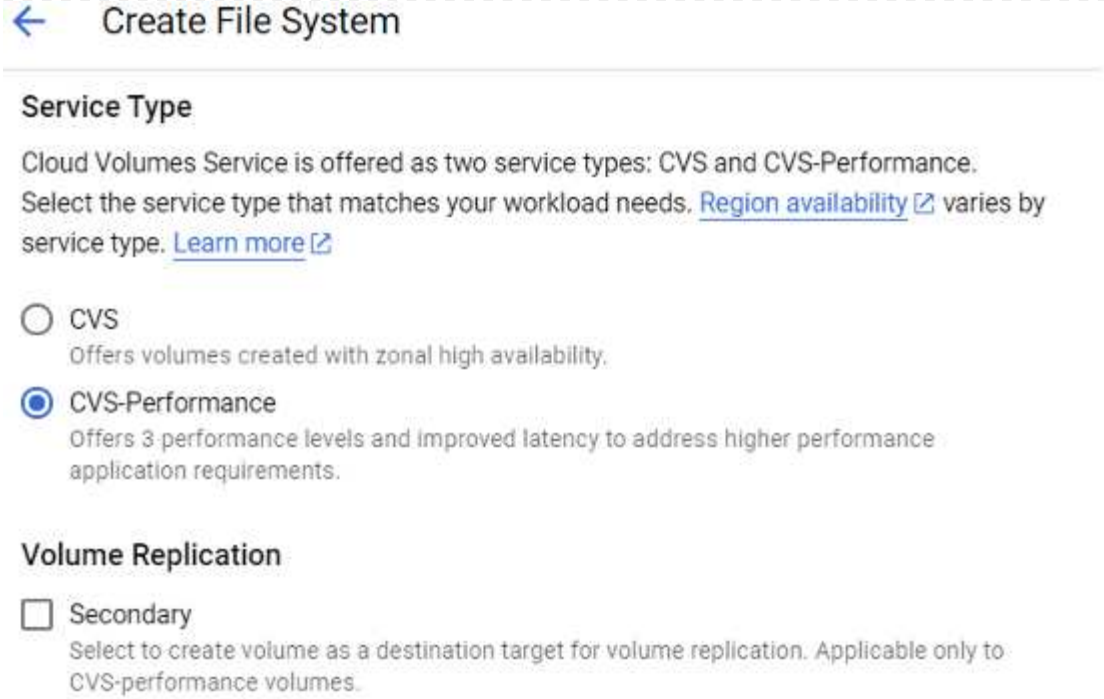

4. Specify the Google Cloud region for the volume and volume path (The volume path must be unique across all of cloud volumes in the project)

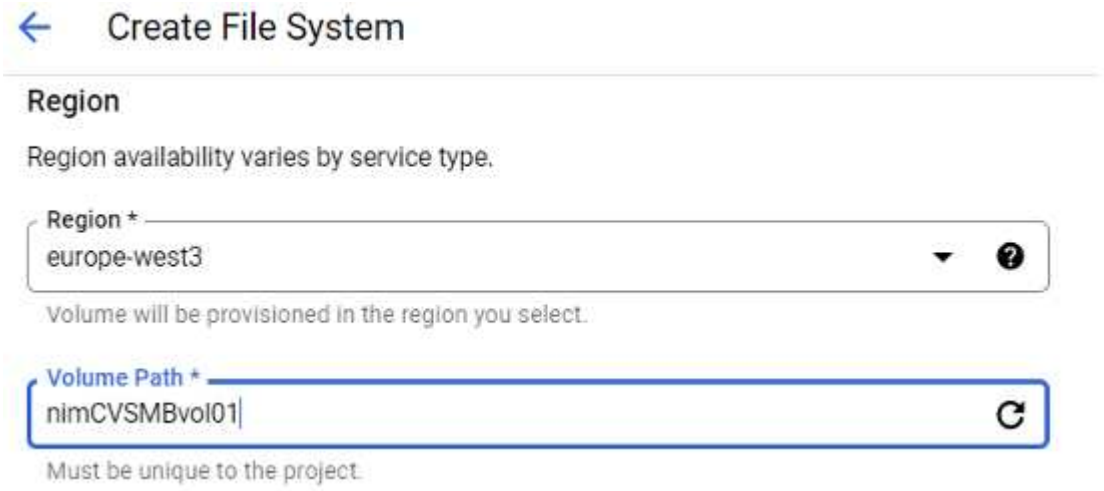

5. Select the level of performance for the volume.

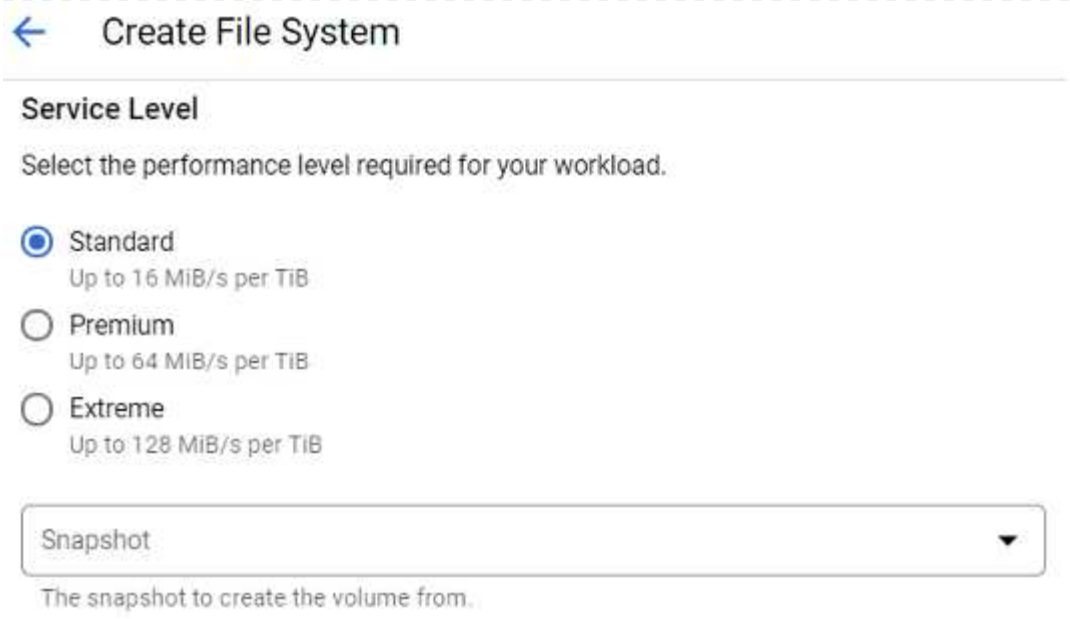

6. Specify the size of the volume and the protocol type. In this testing, SMB is used.

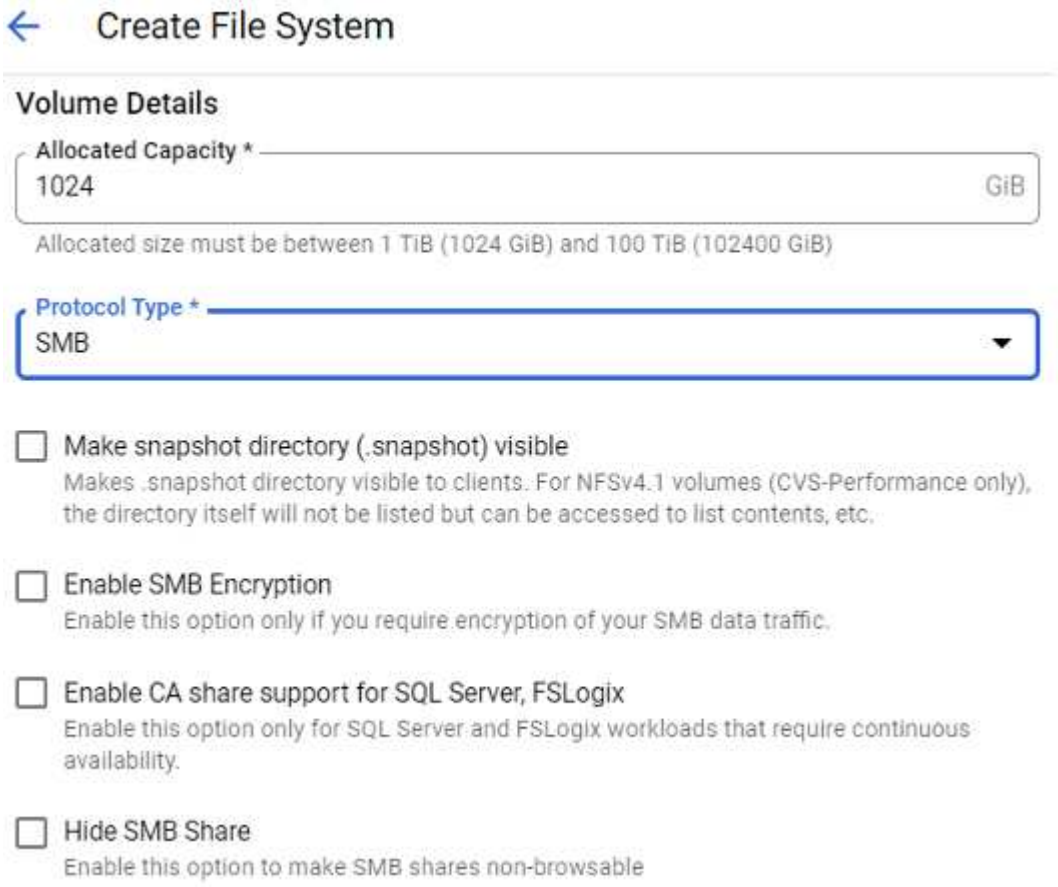

7. In this step, select the VPC Network from which the volume will be accessible. Ensure VPC peering is in place.

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

beforehand, refer to these instructions.

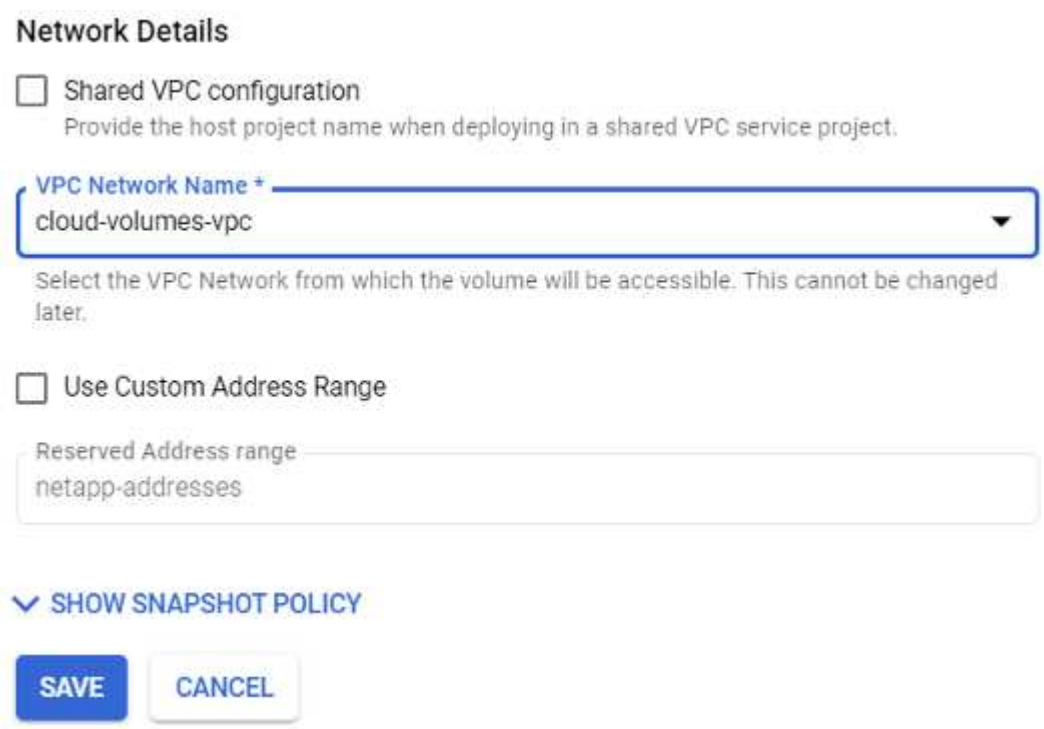

8. Click Save to create the volume.

To mount the SMB volume, do the following:

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

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

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

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

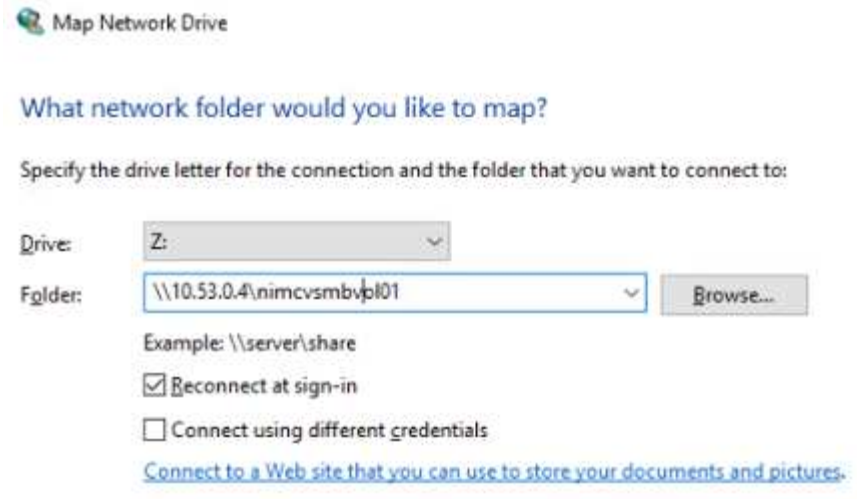

To connect every time you log on to your computer, select the Reconnect at sign-in check box.

5. Click Finish.

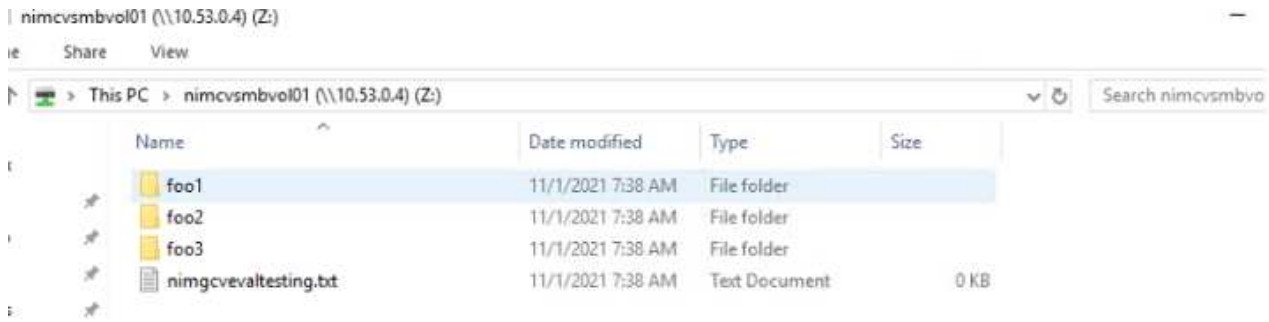

Region Availability for Supplemental NFS datastores on AWS, Azure, and GCP

Learn more about the the Global Region support for supplemental NFS datastores on AWS, Azure and Google Cloud Platform (GCP).

AWS Region Availability

The availability of supplemental NFS datastores on AWS / VMC is defined by Amazon. First, you need to determine if both VMC and FSxN are available in a specified region. Next, you need to determine if the FSxN supplemental NFS datastore is supported in that region.

- Check the availability of VMC here.
- Amazon's pricing guide offers information on where FSxN (FSx ONTAP) is available. You can find that information here.
- Availability of the FSxN supplemental NFS datastore for VMC is coming soon.

While information is still being released, the following chart identifies the current support for VMC, FSxN and FSxN as a supplemental NFS datastore.

Americas

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
US East (Northern Virginia)	Yes	Yes	Yes
US East (Ohio)	Yes	Yes	Yes
US West (Northern California)	Yes	No	No
US West (Oregon)	Yes	Yes	Yes
GovCloud (US West)	Yes	Yes	Yes
Canada (Central)	Yes	Yes	Yes
South America (Sao Paulo)	Yes	Yes	Yes

Last updated on: June 2, 2022.

EMEA

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
Europe (Ireland)	Yes	Yes	Yes
Europe (London)	Yes	Yes	Yes
Europe (Frankfurt)	Yes	Yes	Yes
Europe (Paris)	Yes	Yes	Yes
Europe (Milan)	Yes	Yes	Yes
Europe (Stockholm)	Yes	Yes	Yes

Last updated on: June 2, 2022.

Asia Pacific

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
Asia Pacific (Sydney)	Yes	Yes	Yes
Asia Pacific (Tokyo)	Yes	Yes	Yes
Asia Pacific (Osaka)	Yes	No	No
Asia Pacific (Singapore)	Yes	Yes	Yes
Asia Pacific (Seoul)	Yes	Yes	Yes
Asia Pacific (Mumbai)	Yes	Yes	Yes
Asia Pacific (Jakarta)	No	No	No
Asia Pacific (Hong Kong)	Yes	Yes	Yes

Last updated on: September 28, 2022.

Azure Region Availability

The availability of supplemental NFS datastores on Azure / AVS is defined by Microsoft. First, you need to determine if both AVS and ANF are available in a specific region. Next, you need to determine if the ANF supplemental NFS datastore is supported in that region.

- Check the availability of AVS and ANF here.
- Check the availability of the ANF supplemental NFS datastore here.

GCP Region Availability

GCP region availability will be released when GCP enters public availability.

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.