VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP Deployment Guide

© VMware LLC.

Table of contents

VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP Deployment Guide

Overview

VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP is a jointly engineered, AWS-managed external NFS datastore built on NetApp's ONTAP file system that can be attached to VMware Cloud on AWS vSphere cluster. It provides customers with a flexible, high-performance virtualized storage infrastructure that scales independently of compute resources.

Purpose of This Guide

This deployment guide takes you through the steps to provision and attach an Amazon FSx for NetApp ONTAP (FSx for ONTAP) volume as an NFS datastore for VMware Cloud on AWS. In this document, you will also find best practices, supportability requirements, sizing considerations, and other information helping you plan, design, and implement the integration.

Before you provision and attach FSx for ONTAP as an NFS datastore, you must first set up an environment and deploy an SDDC from the VMC Console (vmc.vmware.com). For more information, see the Getting Started With VMware Cloud on AWS.

Audience

This tutorial is intended for Cloud Architects and Cloud Administrators familiar with VMware Cloud on AWS, VMware vSphere, AWS Console, and Amazon FSx for ONTAP.

Design, Architecture, and Supportability

Customer challenges

Inability to scale storage and compute independently increases cost: Migrating and running storage-heavy workloads on the cloud is not easy today because attaching a storage option that scales independently of compute resources can be complex and expensive.

- VMware Cloud on AWS customers have the following options for additional storage, however, these options come with their own challenges:
  - Purchase additional SDDC hosts to scale storage capacity: Customers who do not need the additional compute and memory that comes with the increase in storage capacity will under-utilize these resources, leading to a higher than desired total cost of ownership (TCO).
  - Convert i3.metal to i3en.metal hosts to get more storage capacity: Customers who have purchased i3.metal instances under a 1 or 3-year subscription agreement that is not yet near the end of its commitment term may not be able to do so without incurring unnecessary additional costs.
  - Purchase external storage from Managed Service Provider: Customers will have to engage 3rd party providers in case of support requirements and increase their vendor footprint.
- Use native AWS storage services (EFS, S3, FSx) to extend storage capacity of individual virtual machines (VMs) that are hosted in VMware Cloud on AWS SDDCs: Customers will need to manage the storage configurations at individual VM level, reducing the efficiency with which storage can be managed and scaled.
- Need a data store with agile data management capabilities: Customers also need comprehensive and consistent data management capabilities for external storage, such as snapshots, clones, replication, etc. to simplify their data storage and agile data management requirements.
- Need a datastore with elastic and flexible capabilities. Customers need the ability to modify the datastore's size, throughput and IOPs dynamically.
- Need to learn additional technology and acquiring new skills increase costs, time, and risk: Engaging new vendors, learning innovative solutions and technologies can be a drain on the productivity of IT and operations management personnel a company has. Customers need a storage solution that uses familiar technology with minimal learning curve and less complexity while migrating on-premises storage intensive workloads to the cloud.

VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP Features

- For customers requiring high storage capacity for their workloads (example, Big Data, data warehousing, and VDI), this integration provides an AWS managed, high-performance NFS datastore built on NetApp's ONTAP file system that can be attached to the VMware Cloud on AWS vSphere cluster.
- Customers can grow the storage capacity as needed without the need to purchase additional host instances.
- This service provides the same features, performance, and administrative capabilities that hundreds of thousands of NetApp customers use on-premises, with the simplicity, agility, security, and scalability of the cloud.
- This integration provides familiar NetApp ONTAP's data management capabilities in the cloud, such as space efficient snapshots, cloning, compression, deduplication, compaction, and replication, giving enhanced protection against ransomware.
- Scheduled maintenance and backup orchestration are included in the service.
- Storage capacity can be added by percentage or absolute value. The minimum storage capacity per file system (i.e., data not tiered to capacity pool storage) is 1024 GB. The maximum storage capacity per file system is 192 TiB uncompressed, however, customers can store more data by using deduplication, compression, and other efficiency mechanisms to reduce the required storage capacity.

- Ability to change data throughput, IOPS, and the size of the file system in addition to ability to increase or decrease the size of volumes using the AWS Console or CLI.

Supportability

- You must not deploy the FSx for ONTAP file system in the Connected VPC used during the SDDC deployment. Instead, you must deploy it in a new Amazon VPC that you own to enable VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP. You have the following options for connecting the FSx for ONTAP file system to your SDDC:
    - VPC Peering (Preferred option for single-AZ deployments)
    - VMware Managed Transit Gateway (vTGW) via SDDC groups
- You must deploy FSx for ONTAP within the same AWS Region as your SDDC
- It is recommended to deploy the primary FSx for ONTAP file system in the same AWS Availability Zone (AZ) as the SDDC
- For single-AZ FSx for ONTAP deployments, deploying the SDDC into the same AZ as the file system reduces latency and prevents cross-AZ network charges
- VMware Cloud on AWS supports multi-AZ and single-AZ deployments of FSx for ONTAP
- Your choice of storage network connectivity will depend on the FSx for ONTAP deployment type:
    - FSx for ONTAP Single AZ deployment: VPC Peering and VMware Transit Connect are supported
    - FSx for ONTAP Multi-AZ deployment: VMware Managed Transit Gateway (vTGW) is supported
- For the NFS Datastore integration, the VMware Cloud on AWS SDDC must be deployed with standard single-AZ clusters. A stretched clusters SDDC currently is not supported
- Access to external datastores is available on SDDC versions 1.20 and greater
- You can attach up to four (4) FSx for ONTAP volumes to a single vSphere cluster on VMware Cloud on AWS
- VMware Cloud Disaster Recovery (VCDR) is supported
- VMware Site Recovery (VSR) is currently not supported (replication sources/targets)
- VMware Cloud on AWS utilizes NFS version 3 and the TCP protocol to access NFS datastores
- Amazon VPC Peering only supports NFS datastore connectivity
- The peered Amazon VPC must not contain the default VPC CIDR Prefix 172.31.0.0/16
- A maximum of 1x Amazon VPC Peering connection can be enabled on a single SDDC

Note: All the items above are subject to change. Updated February 2024

High Level Design

The NFS Datastore provisioning is integrated into the VMware Cloud on AWS Console and API to simplify SDDC administration. As a customer, you declaratively add a datastore to one or more clusters. From that point on, the service will manage and monitor that datastore as part of the cluster. When using FSx for ONTAP as a datastore, AWS provides lifecycle management of the FSx for ONTAP file system (security updates, upgrades, and patches), VMware is responsible for the SDDC lifecycle management and the customer is responsible for establishing the network connectivity, creating, and attaching the external NFS datastore to the SDDC. In this case, the customer owns and manages the storage configuration.

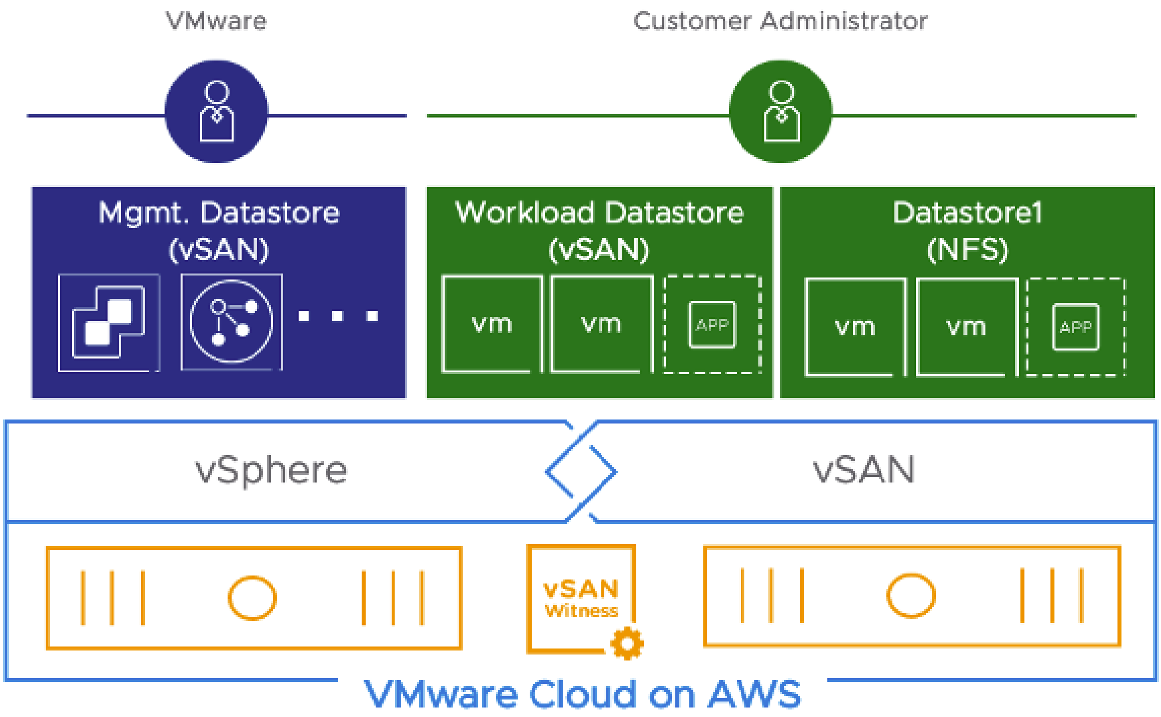

FSx for ONTAP Concepts

Take a closer look at the following concepts, as they will enhance your comprehension of the information covered later in this deployment guide

File systems

A file system is the primary FSx for ONTAP resource. When creating the file system you specify the solid-state drive (SSD) storage capacity and throughput capacity as well as an Amazon VPC from which your file system then becomes accessible. When creating the file system, you can choose from two deployment types, multi-AZ and single-AZ, that offer different levels of availability and durability. Multi-AZ deployments are designed for high availability and built-in replication across two AZs. In contrast, single-AZ deployments offer cost optimization by building the file system across separate fault domains within a single AZ. The NFS datastore integration with VMware Cloud on AWS supports both file system deployment types.

Storage Virtual Machines

The secure logical storage partition through which data is accessed in the file system is known as a storage virtual machine (SVM). An SVM is an isolated file server with its own administrative and data access points. When accessing an NFS datastore on your FSx for ONTAP file system, the VMware Cloud on AWS SDDC interfaces directly with the SVM using the SVM's NFS endpoint IP address. Although it is possible to have a single SVM per cluster, there are also many scenarios where multiple SVMs would be required or offer advantages. Review the business needs, taking into consideration any requirements for workload separation when deciding on how many SVMs to configure

Volumes

Within an SVM, volumes serve as virtual containers for organizing and grouping your data. Each volume will be treated as a separate NFS datastore and data stored in them will consume physical capacity in the file system. Volumes are thin-provisioned, meaning that they only consume storage capacity for the data stored in them.

Storage Network Connectivity

This section provides design patterns for connecting the VMware Cloud on AWS SDDC with your FSx for ONTAP file system. This will essentially represent a VPC-to-VPC connectivity model, with SDDC resources hosted in a VMware Managed VPC and the FSx for ONTAP file system accessible from a customer-managed VPC. As a rule of thumb, VPC connectivity between VPCs is best achieved when using non-overlapping IP ranges for each VPC being connected. For example, if you’d like to connect the SDDC to FSx for

ONTAP file system, make sure each VPC is configured with unique Classless Inter-Domain Routing (CIDR) ranges. Therefore, we advise you to allocate a single, contiguous, non-overlapping CIDR block to be used by each environment.

Amazon VPC Peering

VMware has introduced VPC Peering for single-AZ customers, providing them with a cost-effective advantage over the existing VMware Transit Connect approach. VPC Peering is now the preferred model for connecting single-AZ SDDC clusters to FSx for ONTAP single-AZ deployments.

Review the following list for high-level design considerations:

- The minimum SDDC version must be 1.20
- VPC Peering only supports NFS Datastore connectivity. This connectivity method does not support VM guest OS storage access
- The FSx for ONTAP file system must be deployed using the single-AZ architecture
- No network charges for creating the connection and no metering for data transfers that remain within the same AZ
- Deploying your SDDC in a different AZ to the Amazon FSx for NetApp ONTAP filesystem will result in cross-AZ data transfer charges. Normal cross-AZ data transfer costs still apply, see Data Transfer in Amazon EC2 Pricing
- The SDDC and FSx for ONTAP file system must be deployed in the same region
- Deploying both the SDDC and FSx for ONTAP file system in the same AZ is strongly recommended to improve latency and avoid cross-AZ charges

To learn more about VPC Peering, read the Feature Brief: VPC Peering for External Storage

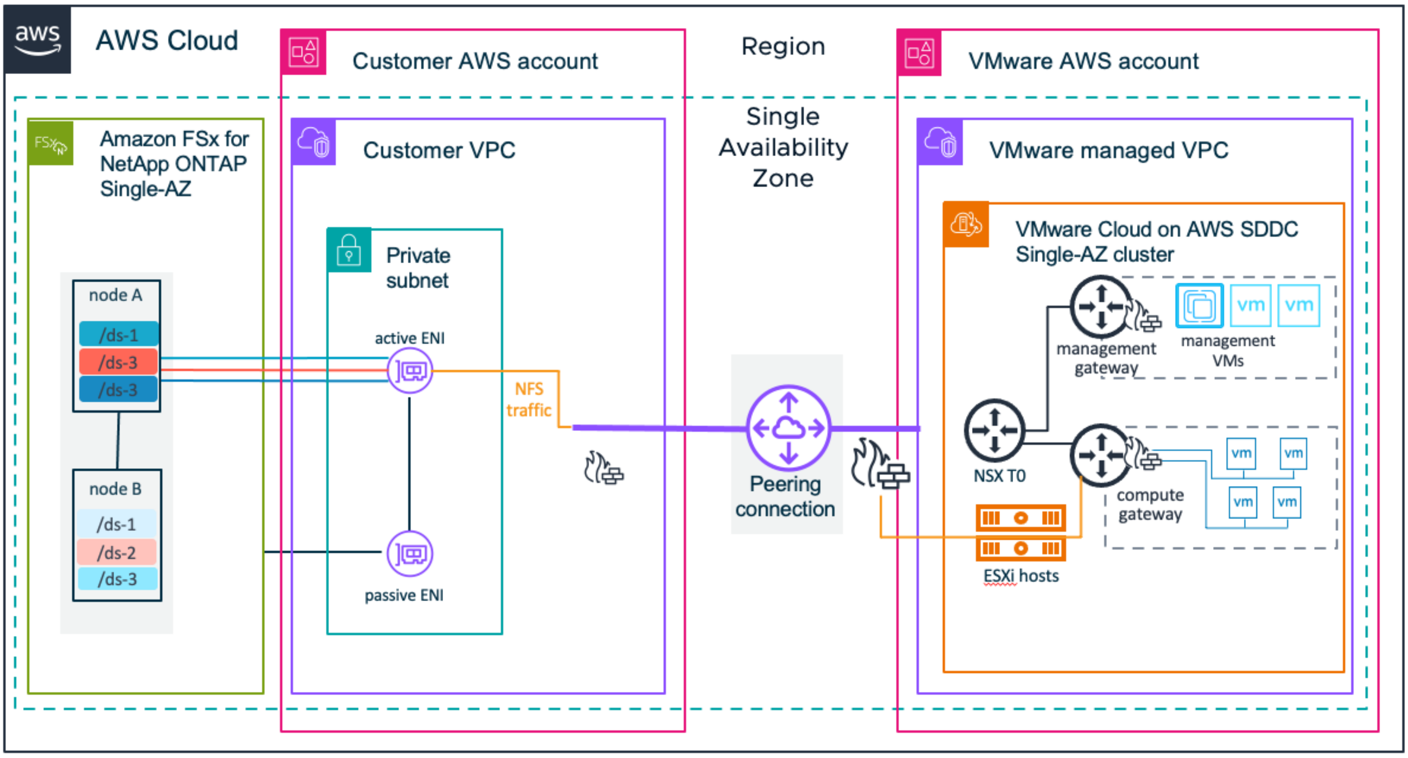

VMware Transit Connect

VMware Transit Connect enables customers to build high-speed, resilient connections between their VMware Cloud on AWS SDDCs and other resources, including SDDCs, native Amazon VPCs, and on-premises. This capability is enabled by a feature called SDDC Groups that helps customers to logically organize SDDCs together to simplify management. Behind the simplification that SDDC Groups provide is the instantiation of a VMware Managed AWS Transit Gateway, a vTGW. It is automatically deployed when an SDDC group is created. In this context, VMware Transit Connect provides the connectivity between the Amazon FSx for ONTAP NFS volumes and ESXi hosts running in the SDDC.

Data flows from the datastore mounted on the ESXi host to the VMware Transit Connect. From here the vTGW routes the storage traffic to the AWS Transit Connect attachment in the customer-managed VPC where it is finally routed to the SVM floating IP associated with the Active FSx for ONTAP ENI of the filesystem.

Review the following list for high-level design considerations:

- The minimum SDDC version must be 1.20
- VMware Transit Connect supports the flexibility of NFS Datastore connectivity and VM guest OS storage
- The FSx for ONTAP file system can be deployed using the single-AZ or multi-AZ architectures
- The overall VMware Transit Connect cost will be based on the number of connections that you make to the Transit Gateway per hour and the amount of traffic (per GB) that flows through the AWS Transit Gateway. See the Storage Network Connectivity Charges section for more details.

Review the VMware Transit Connect reference architecture for more details.

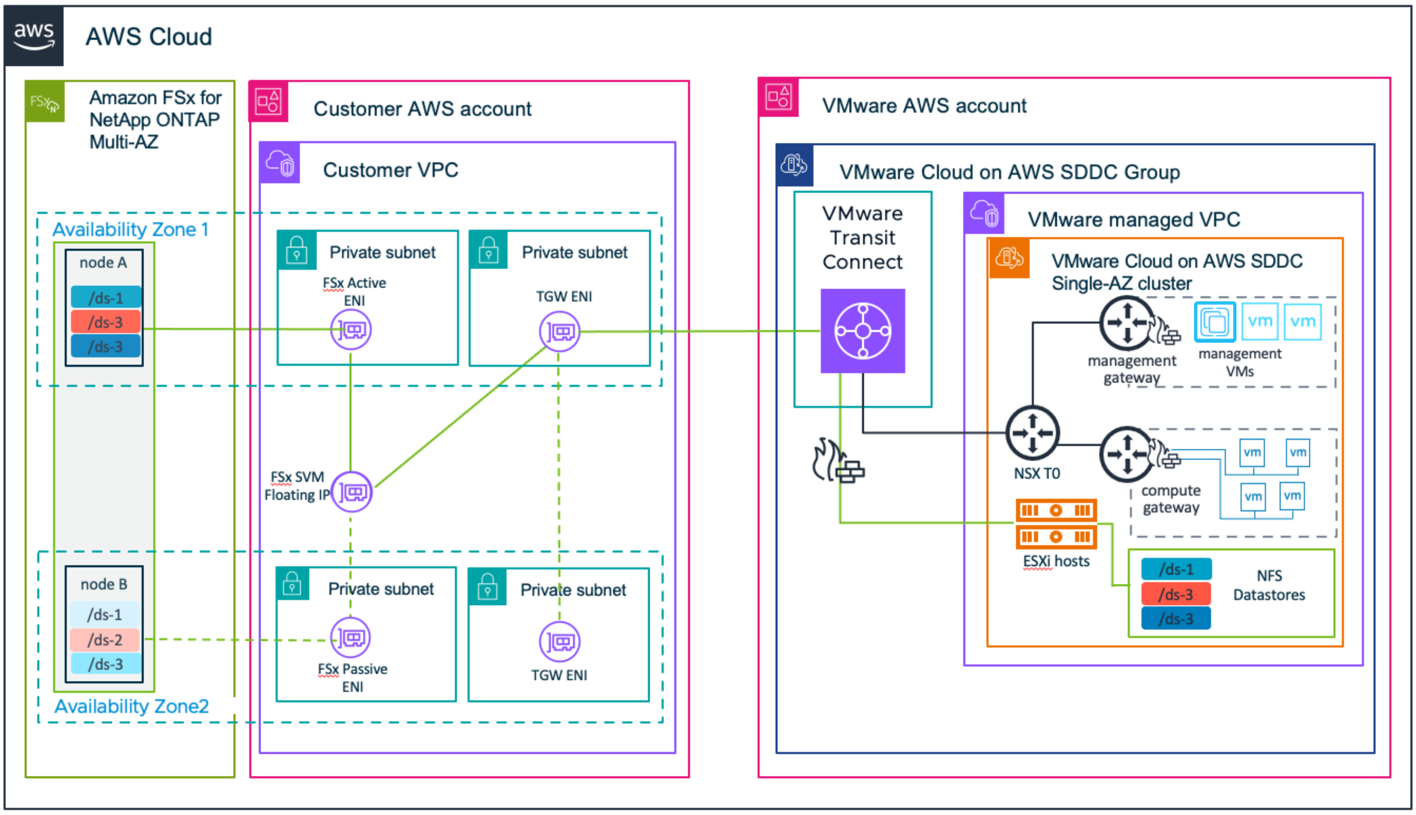

Step-by-Step Deployment Procedure

This step-by-step deployment guide depicts the procedure of adding an Amazon FSx for ONTAP volume to a vSphere cluster on VMware Cloud on AWS. The steps are sequential and build upon one another, so make sure that you complete each step before going to the next step. You can also augment this guide with the step-by-step demo.

Prerequisites

To start, deploy a VMware Cloud on AWS SDDC or use an existing SDDC (SDDC version must be 1.20 or higher). Create a new Amazon VPC in the same region and availability zone where the SDDC resides. (These steps are not covered in this guide)

High-level deployment steps

1. Create an Amazon FSx for ONTAP file system in a newly designated VPC
2. Configure storage network connectivity:
    1. Amazon VPC Peering (Preferred for single-AZ FSx for ONTAP deployments)
    2. VMware Transit Connect
3. Attach NFS volume as external storage

Create Amazon FSx for ONTAP in a designated VPC

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

1. Open the Amazon FSx console at https://console.aws.amazon.com/fsx/ and choose to **Create file system** to start the **File System Creation** wizard.

2. On the Select File System Type page, select **Amazon FSx for NetApp ONTAP** and then click **Next**. The **Create File System** page appears.

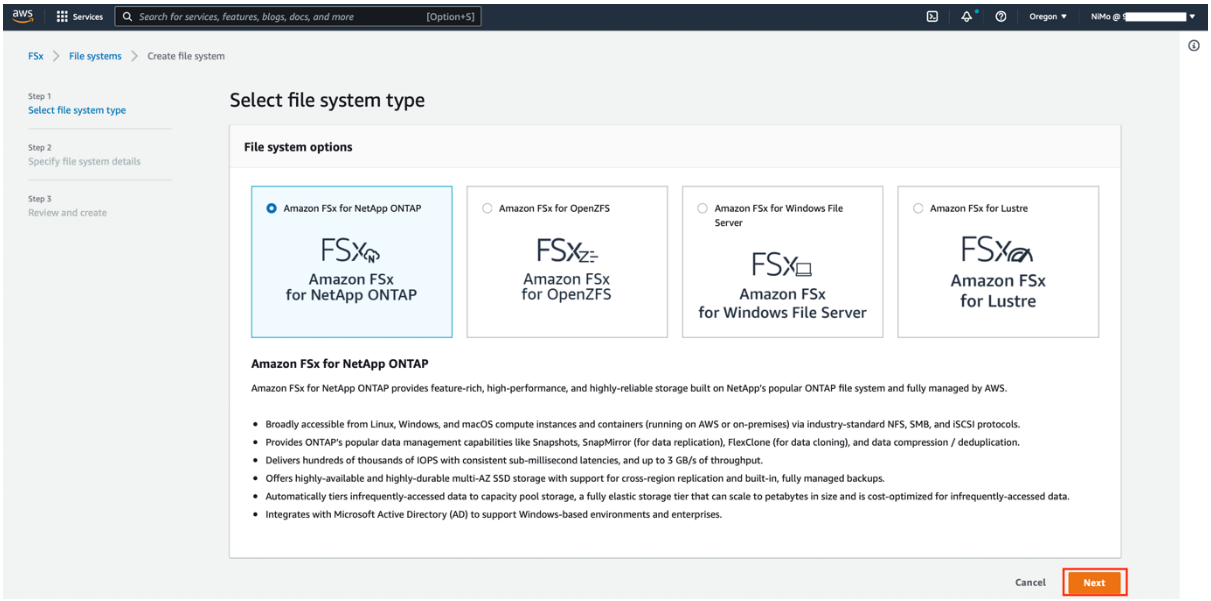
aws
Services
Search for services, features, blogs, docs, and more [Option+S]
Oregon ▾
NiMo @ [redacted] ▾
FSx > File systems > Create file system
Step 1
Select file system type
Step 2
Specify file system details
Step 3
Review and create
Amazon FSx for NetApp ONTAP
Amazon FSx for OpenZFS
Amazon FSx for Windows File Server
Amazon FSx for Lustre
Cancel
Next

Select file system type

File system options

● Amazon FSx for NetApp ONTAP ○ Amazon FSx for OpenZFS ○ Amazon FSx for Windows File Server ○ Amazon FSx for Lustre

Amazon FSx for NetApp ONTAP

Amazon FSx for NetApp ONTAP provides feature-rich, high-performance, and highly-reliable storage built on NetApp's popular ONTAP file system and fully managed by AWS.

- Broadly accessible from Linux, Windows, and macOS compute instances and containers (running on AWS or on-premises) via industry-standard NFS, SMB, and iSCSI protocols.
- Provides ONTAP's popular data management capabilities like Snapshots, SnapMirror (for data replication), FlexClone (for data cloning), and data compression / deduplication.
- Delivers hundreds of thousands of IOPS with consistent sub-millisecond latencies, and up to 3 GB/s of throughput.
- Offers highly-available and highly-durable multi-AZ SSD storage with support for cross-region replication and built-in, fully managed backups.
- Automatically tiers infrequently-accessed data to capacity pool storage, a fully elastic storage tier that can scale to petabytes in size and is cost-optimized for infrequently-accessed data.
- Integrates with Microsoft Active Directory (AD) to support Windows-based environments and enterprises.

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

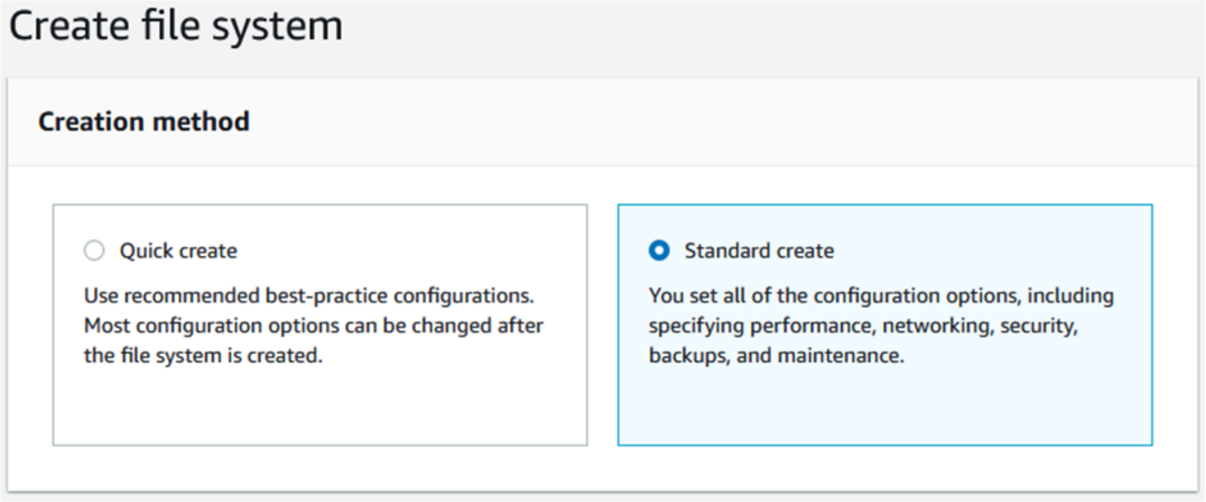

Create file system

Creation method

○ Quick create

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

● Standard create

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

4. Begin your configuration with the **File system details** section:

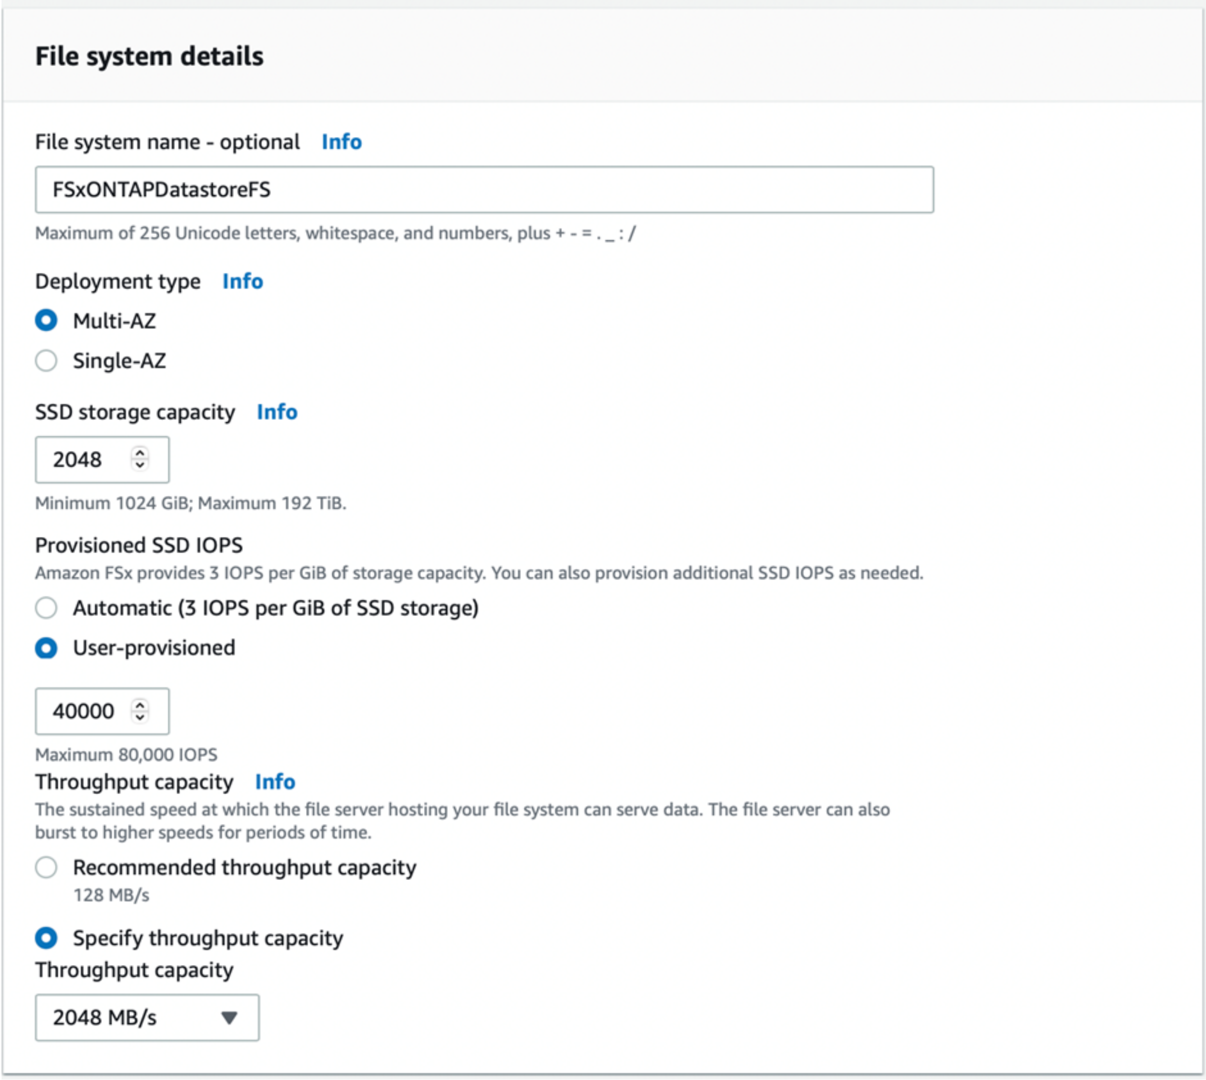

File system details

File system name - optional Info

FSxONTAPDatastoreFS

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

Deployment type Info

● Multi-AZ

○ Single-AZ

SSD storage capacity Info

2048

Minimum 1024 GiB; Maximum 192 TiB.

Provisioned SSD IOPS

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

○ Automatic (3 IOPS per GiB of SSD storage)

● User-provisioned

40000

Maximum 80,000 IOPS

Throughput capacity Info

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

○ Recommended throughput capacity
128 MB/s

● Specify throughput capacity

Throughput capacity

2048 MB/s ▼

5. For **File system name - optional**, enter a name for your NetApp ONTAP file system.

6. For **Deployment type** you can choose between **Multi-AZ** and **Single-AZ**:

   - **Multi-AZ** file systems replicate your data and support automatic failover across multiple Availability Zones in the same AWS Region.
   - **Single-AZ** file systems replicate your data and offer automatic failover within a single Availability Zone.

7. For **SSD storage capacity**, Specify the SSD storage capacity of your file system: 1024 GiB – 192 TiB. The storage capacity can be increased after you create the file system. It is recommended to provision an SDDC capacity that is right-sized to hold your production workloads. Storing your workload VMDKs within the SSD storage capacity is advised for optimal performance.

8. For **Provisioned SSD IOPS**, you have two options to provision the number of IOPS for your file system:

- Choose **Automatic** (the default) if you want Amazon FSx to automatically provision 3 IOPS per GiB of SSD storage.
- Choose User-provisioned if you want to specify the number of IOPS. You can increase your provisioned SSD IOPs after the file system is created.

9. For **Throughput capacity**, you can choose between **Recommended throughput capacity** and **Specify throughput capacity**:

- **Recommended throughput capacity:** This will be based on the SSD storage capacity you entered in step 7.
- **Specify throughput capacity:** Choose a value that is your desired throughput capacity in MB per second (MBps).

Note: For more information on the filesystem SSD IOPs and Throughput capacity, see Amazon FSx for NetApp ONTAP performance

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

Note: Leverage Live Optics or other APM-based tools to size the datastores throughput and provisioned SSD IOPS. However, if starting with assumption, begin with smaller capacity, IOPs, and throughput and scale up as required as workloads are migrated or deployed.

10. In the **Network & Security** section for Virtual Private Cloud (VPC), choose the newly created VPC and preferred subnets along with the route table. In this case, Demo-FSxforONTAP-VPC is selected from the dropdown menu.

Note: Make sure this is a newly designated VPC, not the VMware Cloud on AWS SDDC's Connected VPC.

11. Specify a subnet for your file server. If you are creating a Multi-AZ filesystem, choose a **Preferred subnet** and a **Standby subnet**. These subnets must be in different AZs. If you are creating a Single-AZ filesystem, you only select a single subnet for the primary fileserver.

Note: (Single-AZ only) It is recommended to choose a subnet hosted within the same AZ as your SDDC to reduce cross AZ network charges.

12. Next, you specify the **VPC route tables** to be associated with the newly created filesystem:

- **VPC’s default route table**: This is the default option
- **Select one or more VPC route tables:** This option is only available for Multi-AZ

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

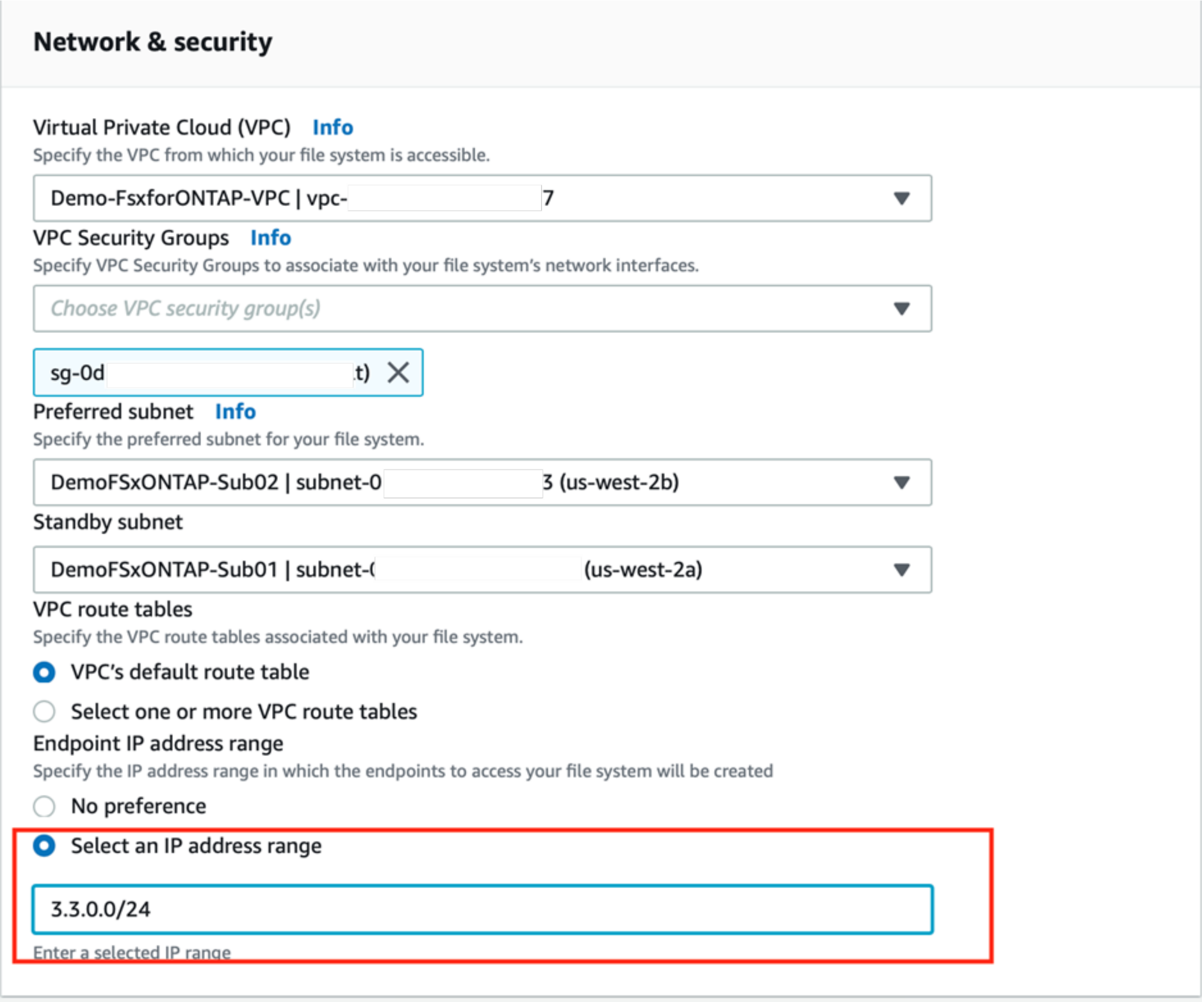

Network & security

Virtual Private Cloud (VPC) Info

Specify the VPC from which your file system is accessible.

Demo-FsxforONTAP-VPC | vpc-[redacted]7 ▼

VPC Security Groups Info

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

Choose VPC security group(s) ▼

sg-0d[redacted]t) ✕

Preferred subnet Info

Specify the preferred subnet for your file system.

DemoFSxONTAP-Sub02 | subnet-0[redacted]3 (us-west-2b) ▼

Standby subnet

DemoFSxONTAP-Sub01 | subnet-([redacted] (us-west-2a) ▼

VPC route tables

Specify the VPC route tables associated with your file system.

● VPC's default route table

○ Select one or more VPC route tables

Endpoint IP address range

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

○ No preference

● Select an IP address range

3.3.0.0/24

Enter a selected IP range

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

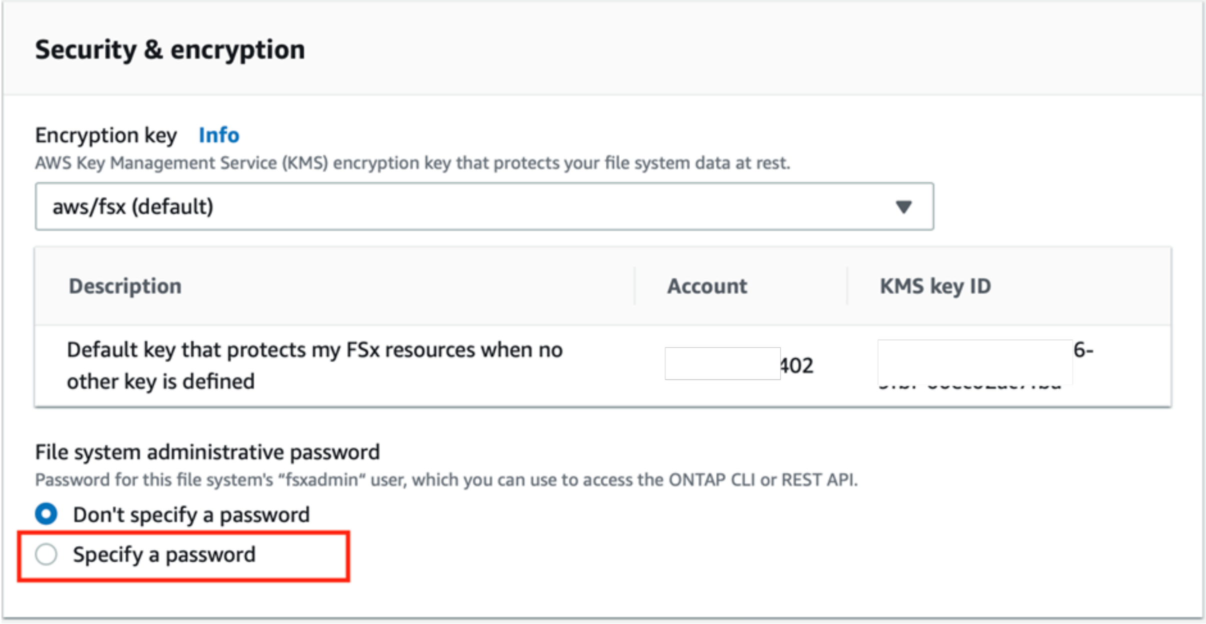

Security & encryption

Encryption key Info

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

aws/fsx (default) ▼

Description	Account	KMS key ID
Default key that protects my FSx resources when no other key is defined	[redacted]402	[redacted]6-

File system administrative password

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

● Don't specify a password

○ Specify a password

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

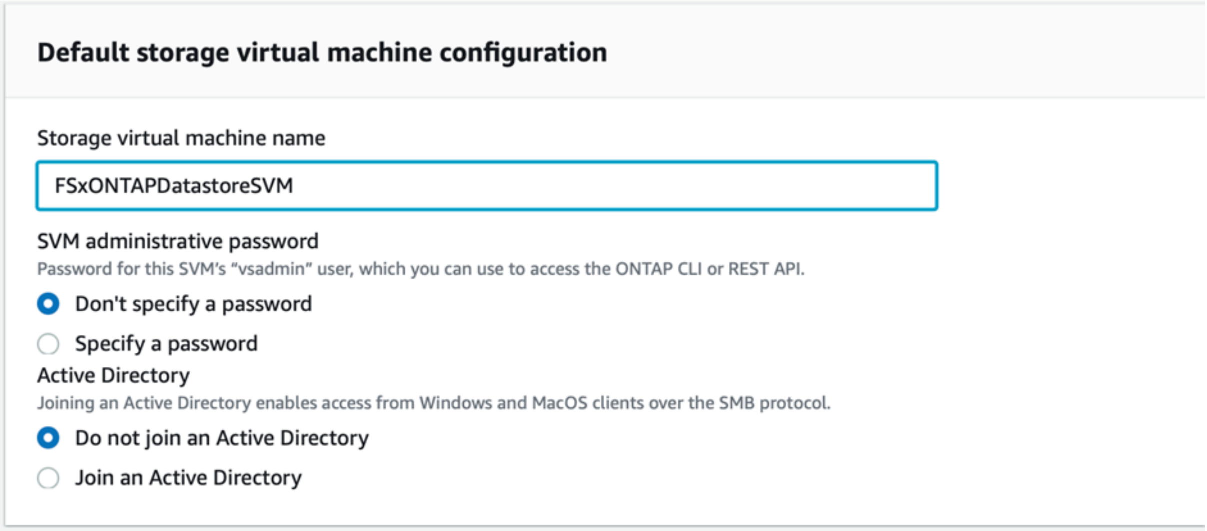

Default storage virtual machine configuration

Storage virtual machine name

FSxONTAPDatastoreSVM

SVM administrative password

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

● Don't specify a password

○ Specify a password

Active Directory

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

● Do not join an Active Directory

○ Join an Active Directory

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

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

Use the following SSH command to create and modify volumes**:**

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

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

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

You can learn more about this process in the following document.

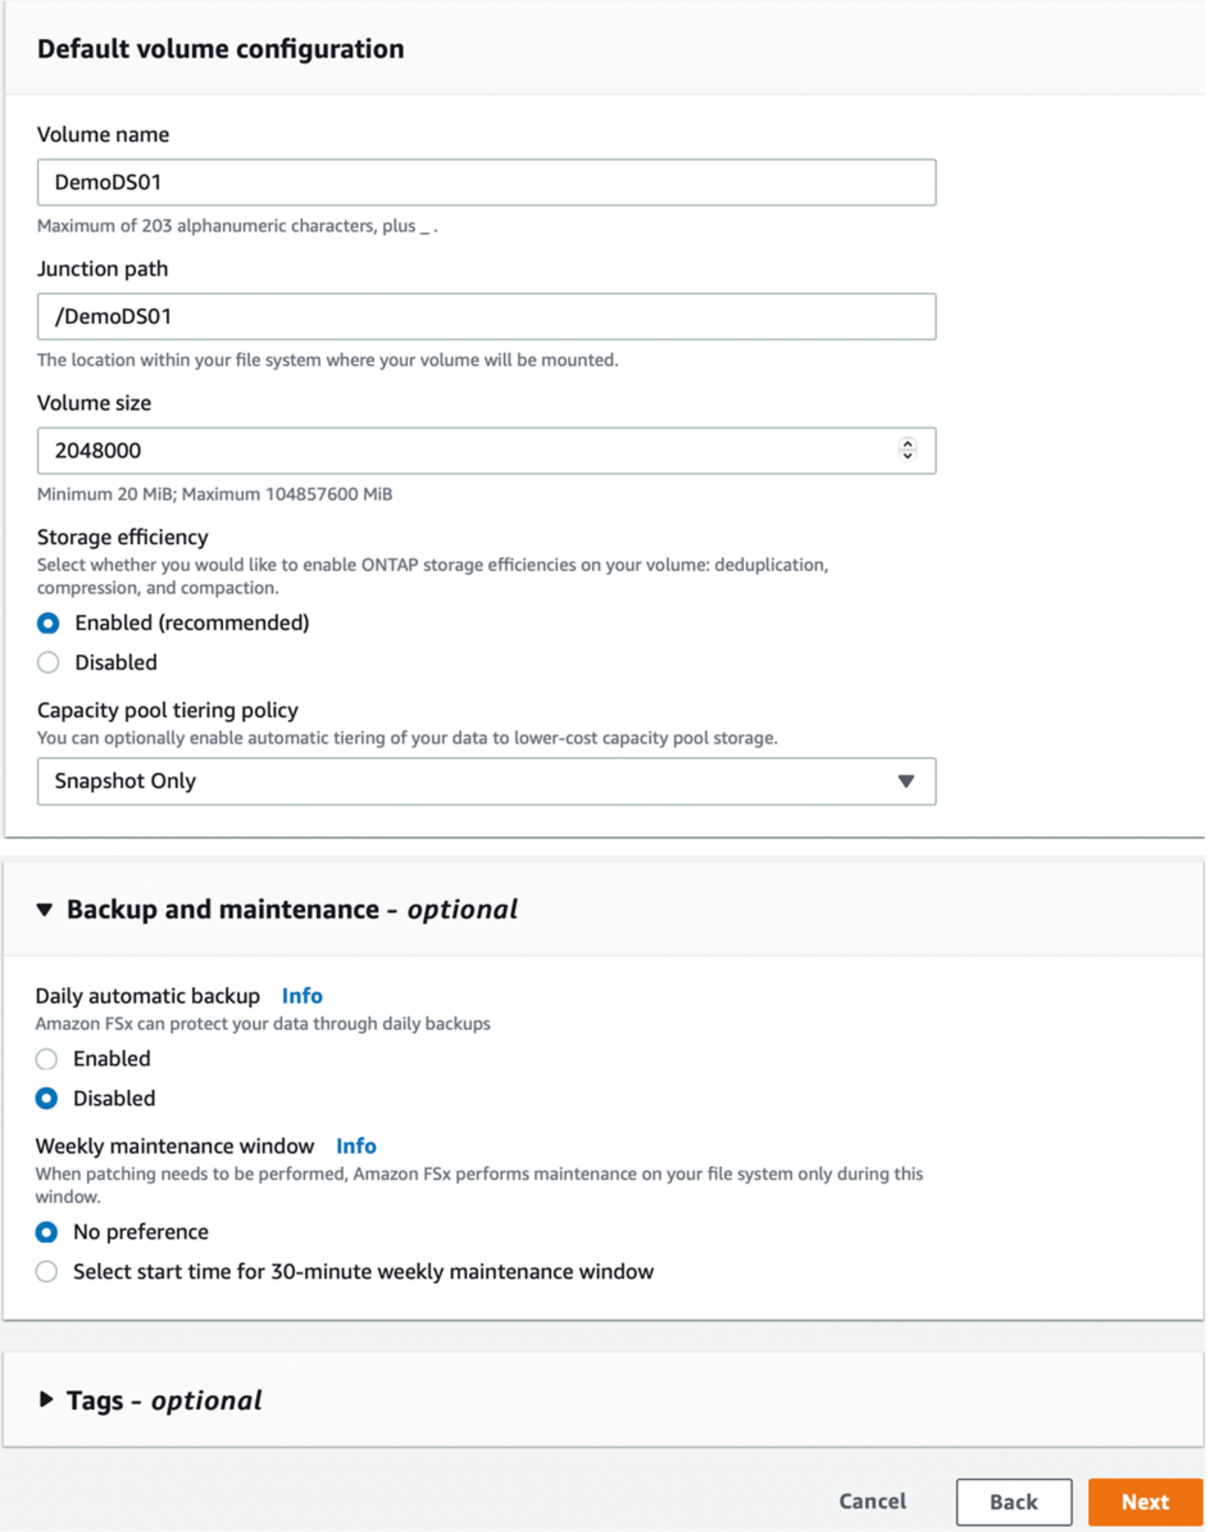

Default volume configuration

Volume name

DemoDS01

Maximum of 203 alphanumeric characters, plus _ .

Junction path

/DemoDS01

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

Volume size

2048000

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency

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

● Enabled (recommended)

○ Disabled

Capacity pool tiering policy

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

Snapshot Only

Backup and maintenance - *optional*

Daily automatic backup Info

Amazon FSx can protect your data through daily backups

○ Enabled

● Disabled

Weekly maintenance window Info

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

● No preference

○ Select start time for 30-minute weekly maintenance window

Tags - *optional*

Note: During initial migration scenario, the default snapshot policy can cause FSx for ONTAP volume to become full. To overcome it, modify the snapshot policy to suit the needs.

16. **Tags** are optional labels that you can assign to an AWS resource. Each tag consists of a key and a value. Using tags can help you manage, filter, and search your AWS resources, including your file system.

17. Click **Next**

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

19. Click **Create File System**

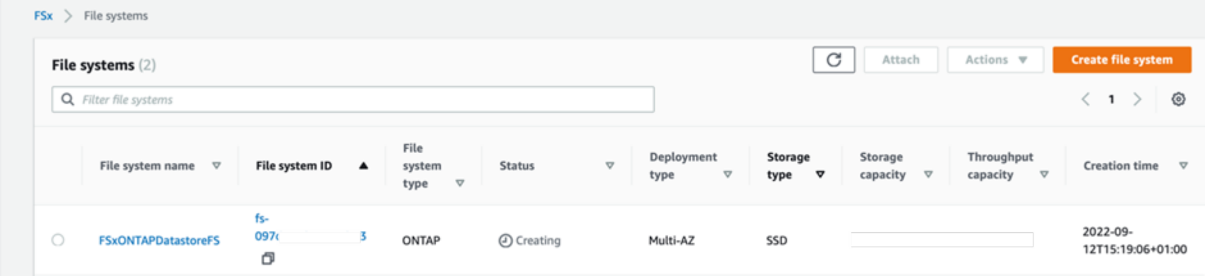

	File system name	File system ID	File system type	Status	Deployment type	Storage type	Storage capacity	Throughput capacity	Creation time
○	FSxONTAPDatastoreFS	fs-097[redacted]3	ONTAP	Creating	Multi-AZ	SSD	[redacted]		2022-09-12T15:19:06+01:00

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

Note: Currently, you can attach up to four FSx for ONTAP volumes to a single vSphere cluster on VMware Cloud on AWS.

Note: FSx for ONTAP has a default limit of 100TB for volumes and 16TB per file. For customers who intend to deploy larger volumes (up to 300TB) or VMDKs larger than 16TB (up to 62TB VMDKs can be supported), ensure to enable the large file support flag **-is-large-size-enabled true** at the volume level.

Use the following SSH command to create a new volume with large file support enabled**:**

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

Use the following SSH command to modify an existing volume to enable large file support**:**

```
volume modify -vserver FSxONTAPDatastoreSVM -volume DemoDS002 -is-large-size-enabled true
```

Note: Activating support for large files on an existing datastore will not automatically update in vCenter. As a result, vCenter will continue to show a maximum size limit of 16TB for VMDK files. Customers will need to detach and then reattach the datastore through the VMC console to refresh the details shown in vCenter. Keep in mind that detaching a datastore requires migrating all VMDKs to another datastore using Storage vMotion, or, if it is the sole external NFS datastore, the VMs must be powered off and removed from the inventory.

20. Once all the required SVMs and volumes have been created; note down the relevant IP connectivity endpoints required for vSphere datastore mounting. You can find the NFS mount endpoints in the Amazon FSx Console by selecting your **filesystem**, choosing **Storage virtual machines**, and then choosing the SVM where your volume resides.

Amazon VPC Peering

Your SDDC offers various network options for connecting to the FSx for ONTAP file system. In this section, we will explore the VPC Peering capability and the necessary steps to extend network connectivity to the external NFS storage.

Request VPC Peering

1. Login to the VMware Cloud on AWS Console and locate the intended SDDC.
2. Open the Storage tab of the SDDC and click Create Peering Connection under VPC Peering

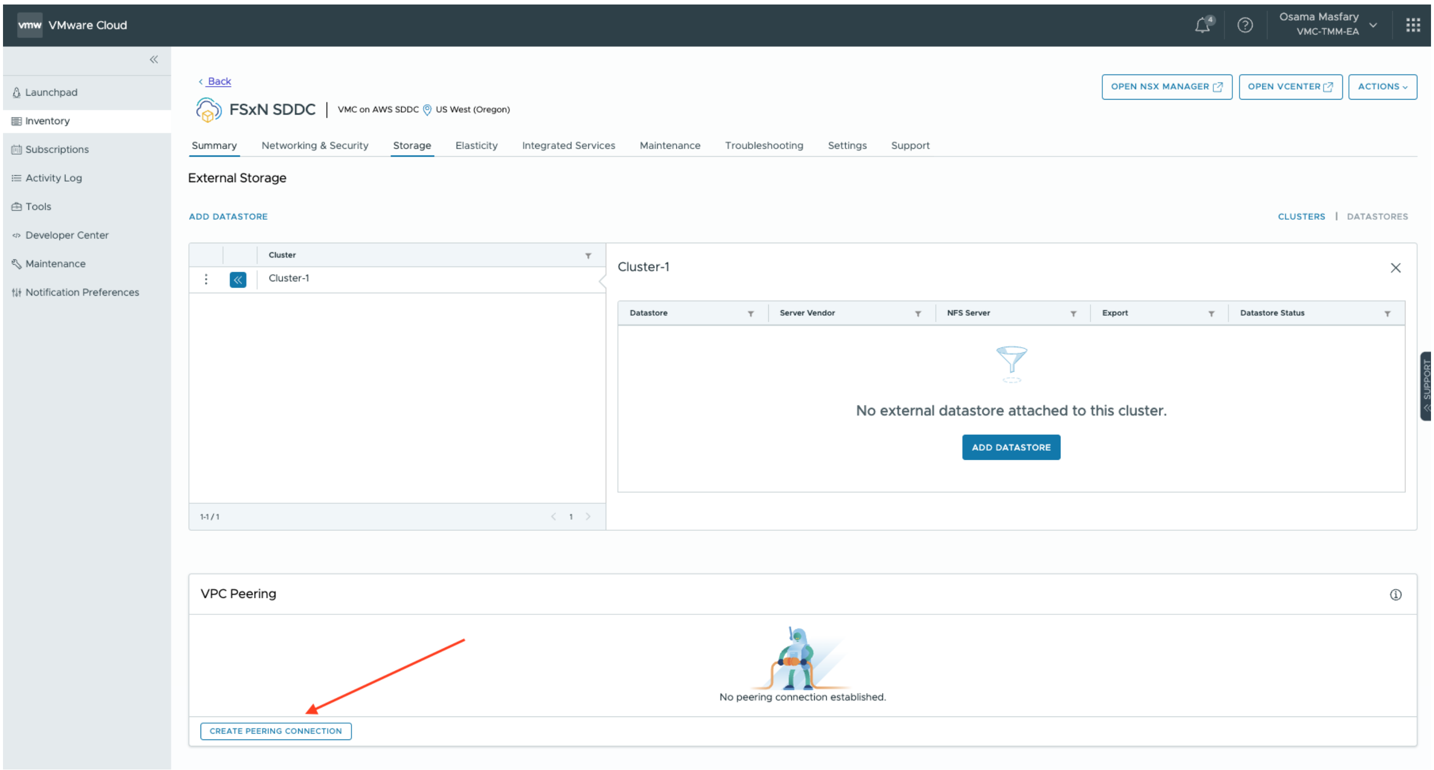

The following information will be required to complete the request:

- AWS Account ID (AWS Account hosting the FSx ONTAP file system)
- Amazon VPC ID (Amazon VPC providing access to the FSx ONTAP file system)

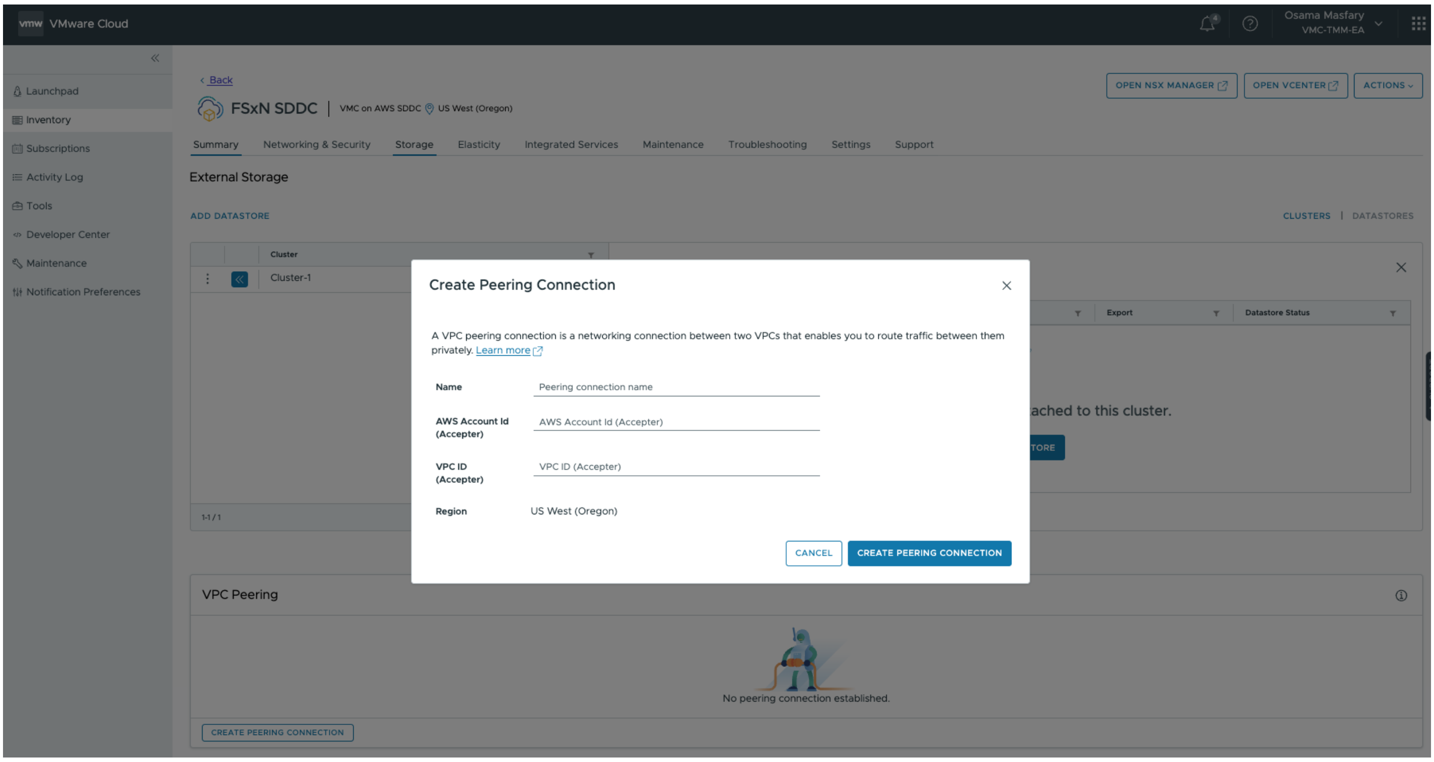

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them privately. Learn more

Accept VPC Peering (your AWS Account)

1. Log in to AWS Console: Access your AWS account at https://aws.amazon.com/console/ using your credentials.
2. Once logged in, click on the "Services" dropdown and select "VPC" under "Networking & Content Delivery."

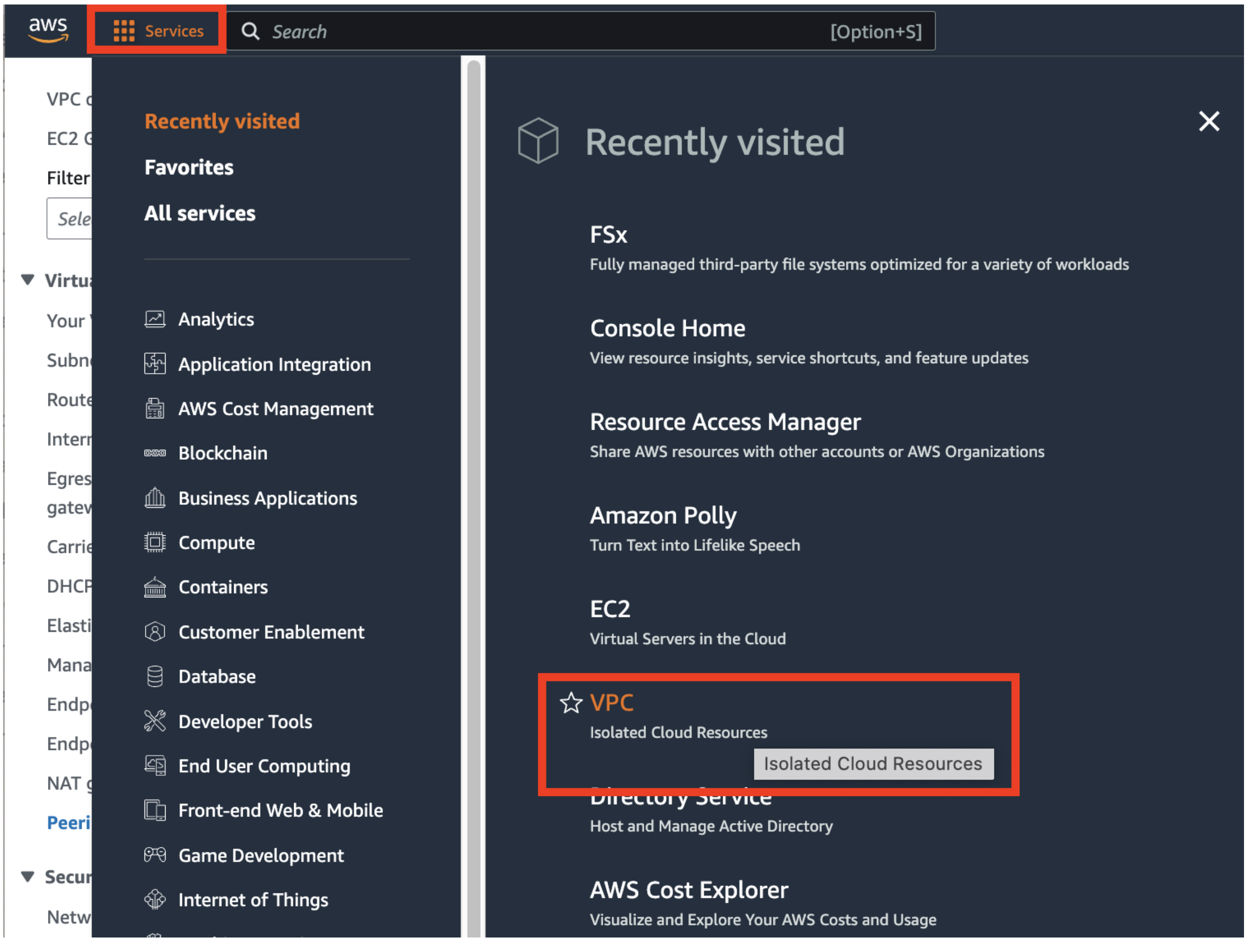

3. Access Peering Connections: In the VPC Dashboard, find and click on "Peering Connections" from the left-hand menu.

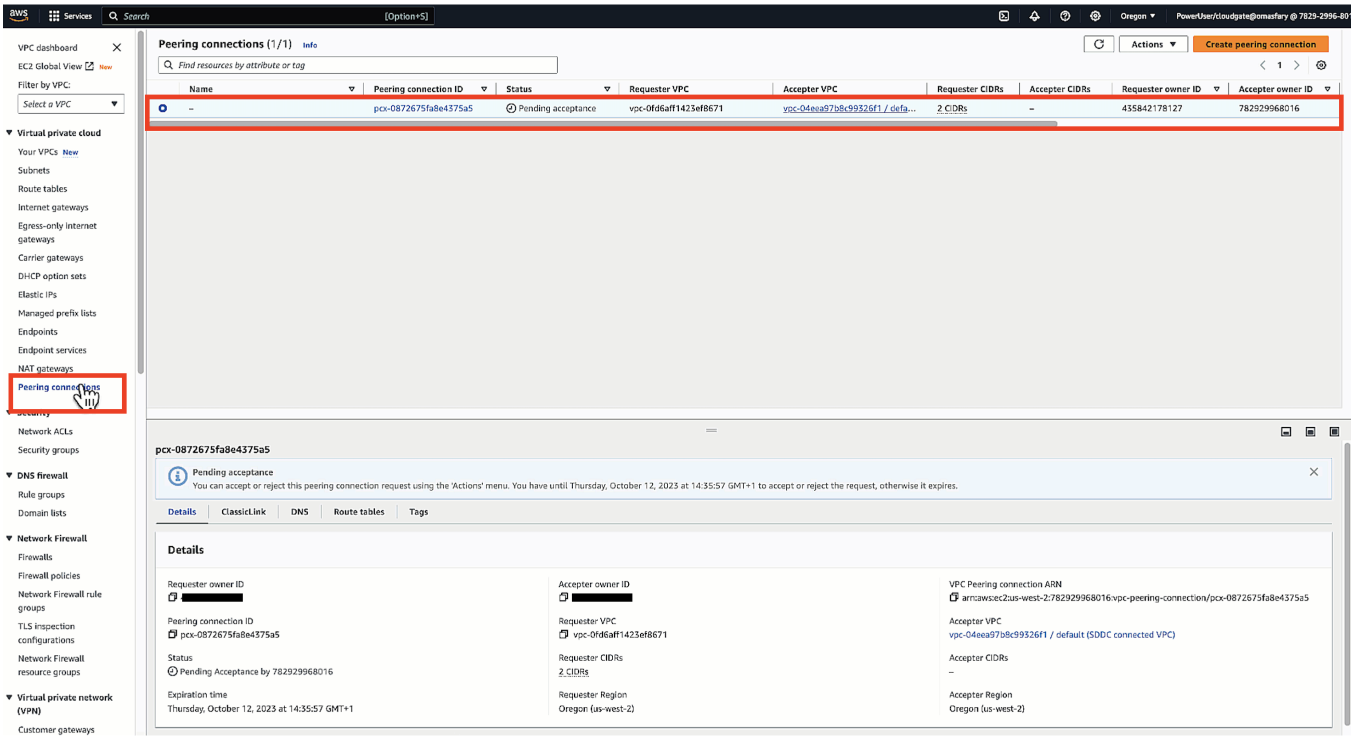

Peering connections (1/1)

	Name	Peering connection ID	Status	Requester VPC	Accepter VPC	Requester CIDRs	Accepter CIDRs	Requester owner ID	Accepter owner ID
●	–	pcx-0872675fa8e4375a5	Pending acceptance	vpc-0fd6aff1423ef8671	vpc-04eea97b8c99326f1 / defa...	2 CIDRs	–	435842178127	782929968016

pcx-0872675fa8e4375a5

Pending acceptance
You can accept or reject this peering connection request using the 'Actions' menu. You have until Thursday, October 12, 2023 at 14:35:57 GMT+1 to accept or reject the request, otherwise it expires.

Details

Requester owner ID
[redacted]

Peering connection ID
pcx-0872675fa8e4375a5

Status
Pending Acceptance by 782929968016

Expiration time
Thursday, October 12, 2023 at 14:35:57 GMT+1

Accepter owner ID
[redacted]

Requester VPC
vpc-0fd6aff1423ef8671

Requester CIDRs
2 CIDRs

Requester Region
Oregon (us-west-2)

VPC Peering connection ARN
arn:aws:ec2:us-west-2:782929968016:vpc-peering-connection/pcx-0872675fa8e4375a5

Accepter VPC
vpc-04eea97b8c99326f1 / default (SDDC connected VPC)

Accepter CIDRs
–

Accepter Region
Oregon (us-west-2)

4. Select Pending Peering Connection: Identify the pending peering connection vpc-peering-fsx-*UUID*

5. Highlight the pending peering connection, then click "Actions" and choose "Accept Request." Confirm your decision.

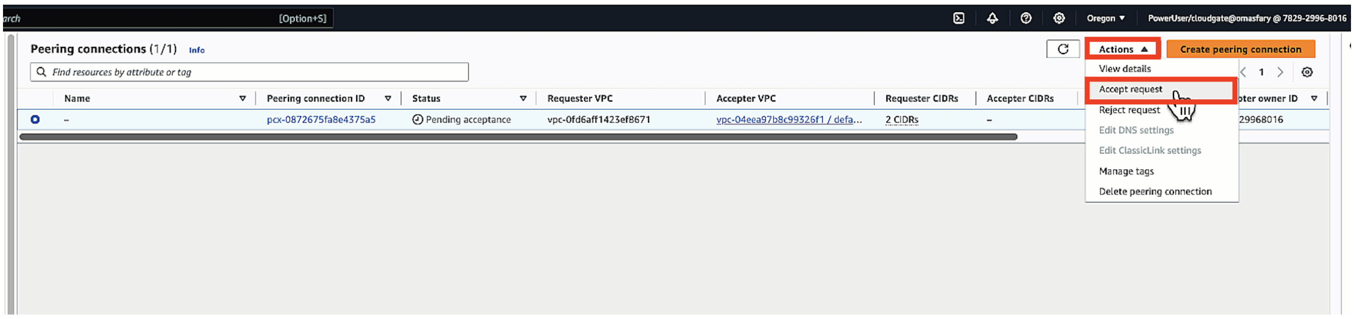

[Option+S]
Oregon ▾
PowerUser/cloudgate@omasfary @ 7829-2996-8016
Actions ▴
Create peering connection
View details
Accept request
Reject request
Edit DNS settings
Edit ClassicLink settings
Manage tags
Delete peering connection

Peering connections (1/1) Info

Find resources by attribute or tag

	Name	Peering connection ID	Status	Requester VPC	Accepter VPC	Requester CIDRs	Accepter CIDRs	…pter owner ID
●	–	pcx-0872675fa8e4375a5	Pending acceptance	vpc-0fd6aff1423ef8671	vpc-04eea97b8c99326f1 / defa...	2 CIDRs	–	…29968016

6. Status Update: The peering connection's status will change to "Active" once the acceptance process is complete.

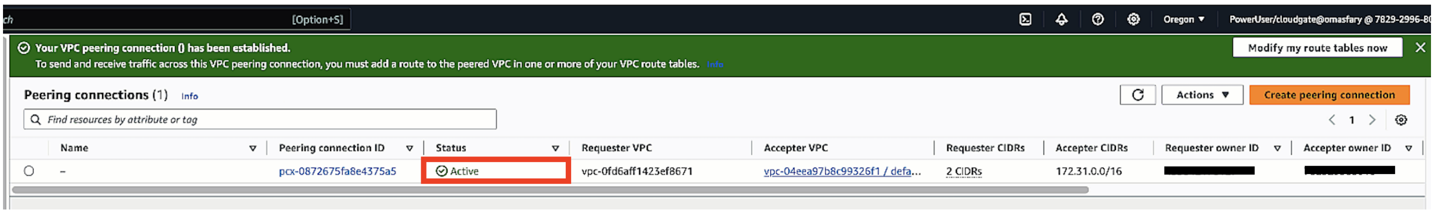

[Option+S]
Oregon ▾
PowerUser/cloudgate@omasfary @ 7829-2996-8
Modify my route tables now
Actions ▾
Create peering connection

Your VPC peering connection () has been established.
To send and receive traffic across this VPC peering connection, you must add a route to the peered VPC in one or more of your VPC route tables. Info

Peering connections (1) Info

Find resources by attribute or tag

	Name	Peering connection ID	Status	Requester VPC	Accepter VPC	Requester CIDRs	Accepter CIDRs	Requester owner ID	Accepter owner ID
○	–	pcx-0872675fa8e4375a5	Active	vpc-0fd6aff1423ef8671	vpc-04eea97b8c99326f1 / defa...	2 CIDRs	172.31.0.0/16	[redacted]	[redacted]

Note: Proceed to the next step once the state of the peering connection is Active in both the AWS and VMware Cloud on AWS Consoles.

Configure Security Group and VPC Route Table

Once the VPC Peering is in an Active state, the next step is to confirm the Security Group in the associated VPC is configured with the correct inbound rules detailing the SDDC management subnet CIDR and NFS network ports.

1. Update the Inbound rule with the CIDR block of the SDDC management. You could either allow all the traffic from the SDDC or specify the ports required for NFS traffic, depending on your security requirements.

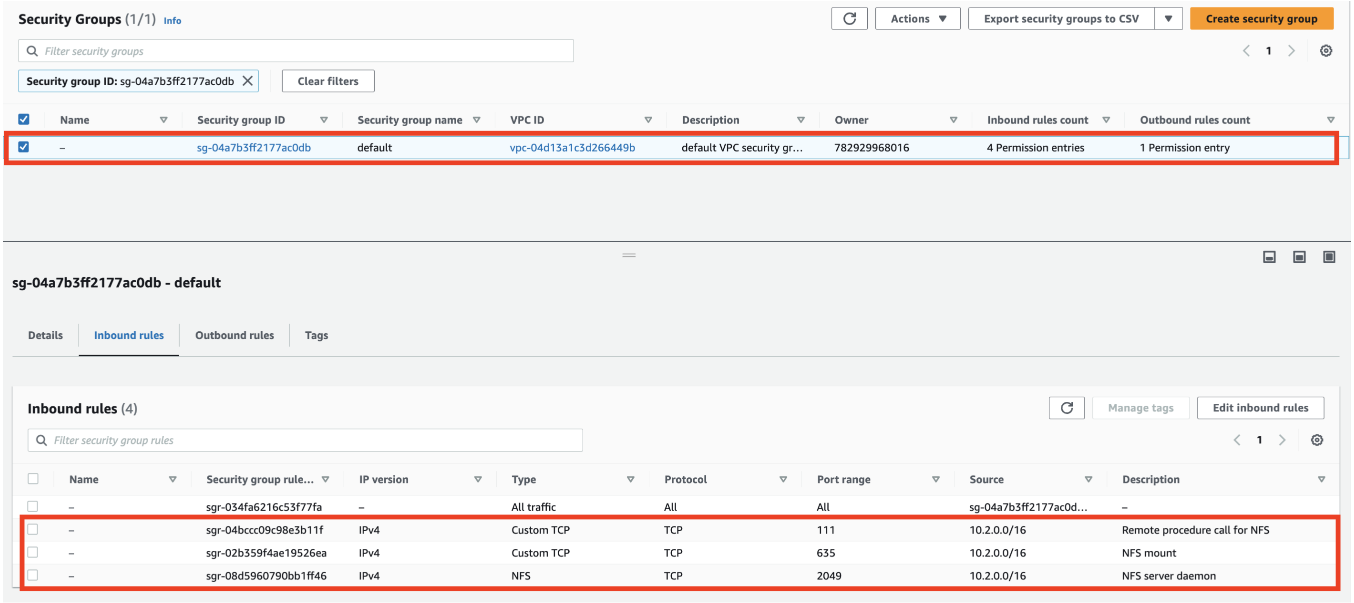

Security Groups (1/1) Info

Filter security groups

Security group ID: sg-04a7b3ff2177ac0db

☑	Name	Security group ID	Security group name	VPC ID	Description	Owner	Inbound rules count	Outbound rules count
☑	–	sg-04a7b3ff2177ac0db	default	vpc-04d13a1c3d266449b	default VPC security gr...	782929968016	4 Permission entries	1 Permission entry

sg-04a7b3ff2177ac0db - default

Inbound rules (4)

Filter security group rules

□	Name	Security group rule...	IP version	Type	Protocol	Port range	Source	Description
□	–	sgr-034fa6216c53f77fa	–	All traffic	All	All	sg-04a7b3ff2177ac0d...	–
□	–	sgr-04bccc09c98e3b11f	IPv4	Custom TCP	TCP	111	10.2.0.0/16	Remote procedure call for NFS
□	–	sgr-02b359f4ae19526ea	IPv4	Custom TCP	TCP	635	10.2.0.0/16	NFS mount
□	–	sgr-08d5960790bb1ff46	IPv4	NFS	TCP	2049	10.2.0.0/16	NFS server daemon

2. Update the routes in the VPC Route Table with Destination as SDDC's Management CIDR and Target as VPC Peering Id.

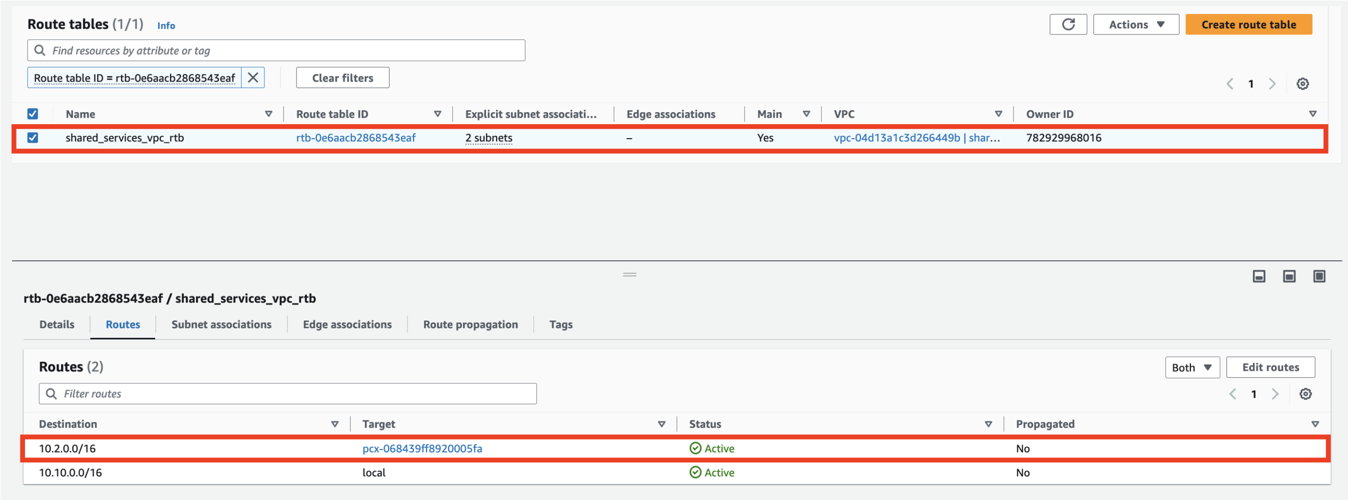

Route table ID = rtb-0e6aacb2868543eaf

☑	Name	Route table ID	Explicit subnet associati...	Edge associations	Main	VPC	Owner ID
☑	shared_services_vpc_rtb	rtb-0e6aacb2868543eaf	2 subnets	–	Yes	vpc-04d13a1c3d266449b | shar...	782929968016

Destination	Target	Status	Propagated
10.2.0.0/16	pcx-068439ff8920005fa	Active	No
10.10.0.0/16	local	Active	No

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

This is the last step in preparing the VPC Peering connectivity. With the file system configured, routes added, and security groups updated, it is time to mount the datastores.

VMware Transit Connect

In this section, we will explore the VMware Transit Connect capability and the necessary steps to extend network connectivity to the external NFS storage.

Create an SDDC Group

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

1. Log into the VMC Console at https://vmc.vmware.com.

2. On the **Inventory** page, click **SDDC Groups**.

3. On the **SDDC Groups** tab, click **ACTIONS** and select **Create SDDC Group**. For demo purposes, the SDDC group is called FSxONTAPDatastoreGrp.

4. On the Membership grid, select the SDDCs to include as group member. You can add a single or multiple SDDCs to the group. All SDDCs in the group will have access to Amazon FSx volumes.

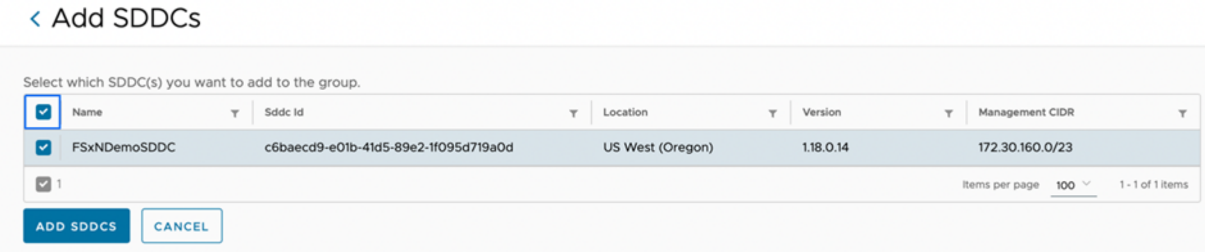

Select which SDDC(s) you want to add to the group.

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

☑ 1 Items per page 100 1 - 1 of 1 items

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

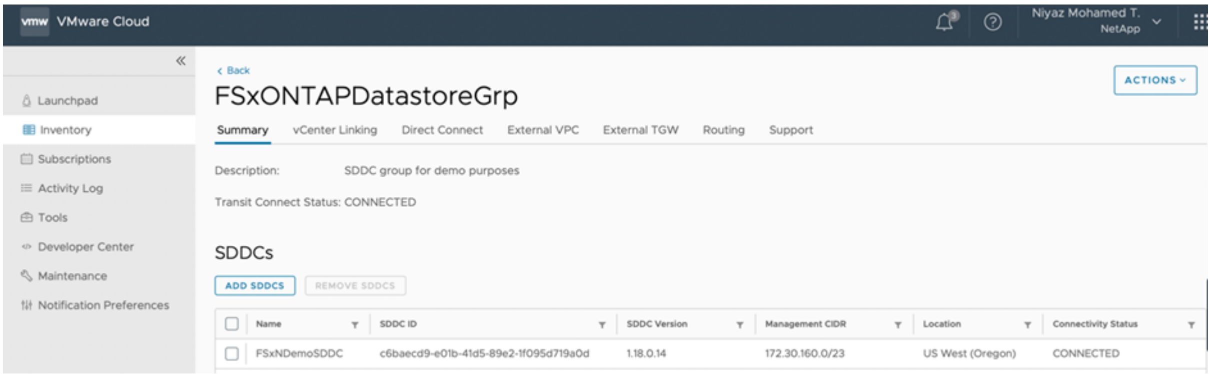

FSxONTAPDatastoreGrp

Description: SDDC group for demo purposes

Transit Connect Status: CONNECTED

SDDCs

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

Note: SDDC Groups are an organization-level object. An SDDC Group cannot contain SDDCs from more than one organization. An SDDC Group can include members from up to three AWS regions. You can have a single SDDC as a member of an SDDC Group.

Configure VMware Transit Connect

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

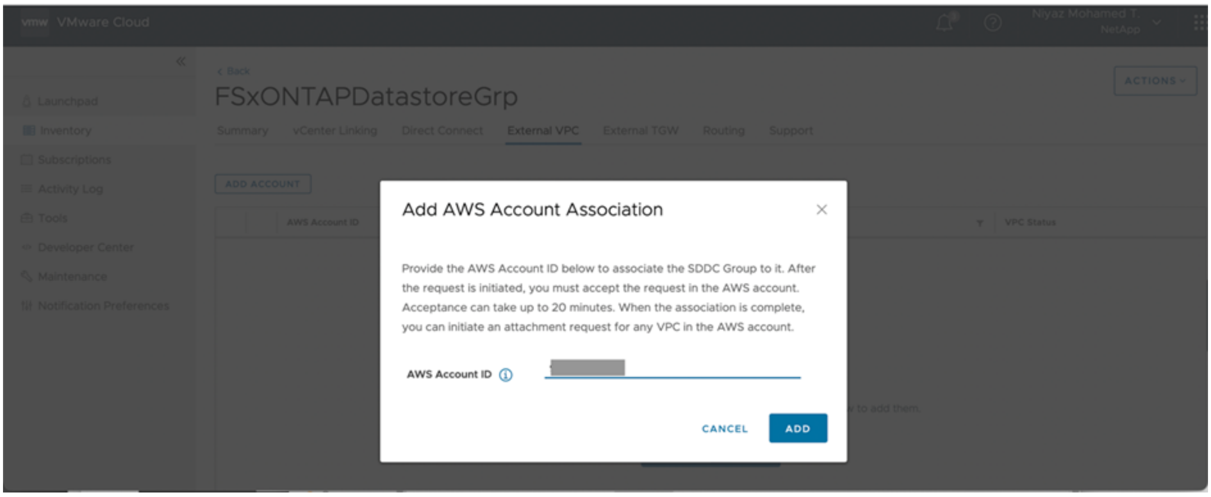

Add AWS Account Association

Provide the AWS Account ID below to associate the SDDC Group to it. After the request is initiated, you must accept the request in the AWS account. Acceptance can take up to 20 minutes. When the association is complete, you can initiate an attachment request for any VPC in the AWS account.

AWS Account ID [redacted]

2. Click **Add Account**.

- Provide the AWS account that was used to provision the FSx for ONTAP file system.
- Click **Add**.

3. Back in the AWS console, log into the same AWS account and navigate to the **Resource Access Manager** service page. There is a button for you to accept the resource share.

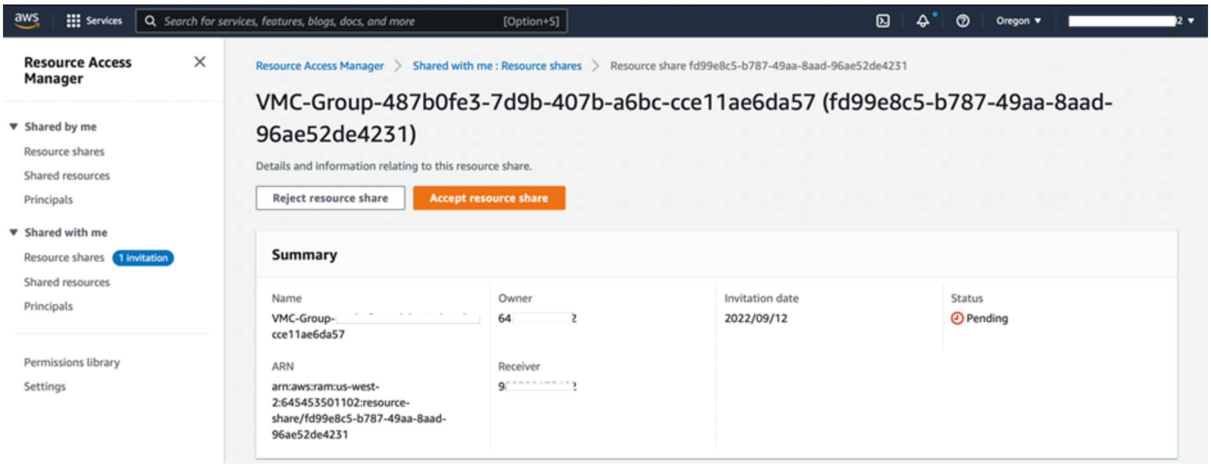
aws
Services
Search for services, features, blogs, docs, and more [Option+S]
Oregon
Resource Access Manager
Shared by me
Resource shares
Shared resources
Principals
Shared with me
Resource shares 1 invitation
Shared resources
Principals
Permissions library
Settings
Resource Access Manager > Shared with me : Resource shares > Resource share fd99e8c5-b787-49aa-8aad-96ae52de4231
Reject resource share
Accept resource share

VMC-Group-487b0fe3-7d9b-407b-a6bc-cce11ae6da57 (fd99e8c5-b787-49aa-8aad-96ae52de4231)

Details and information relating to this resource share.

Summary

Name: VMC-Group-[redacted]cce11ae6da57

Owner: 64[redacted]2

Invitation date: 2022/09/12

Status: Pending

ARN: arn:aws:ram:us-west-2:645453501102:resource-share/fd99e8c5-b787-49aa-8aad-96ae52de4231

Receiver: 9[redacted]2

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

4. Click **Accept resource share**.

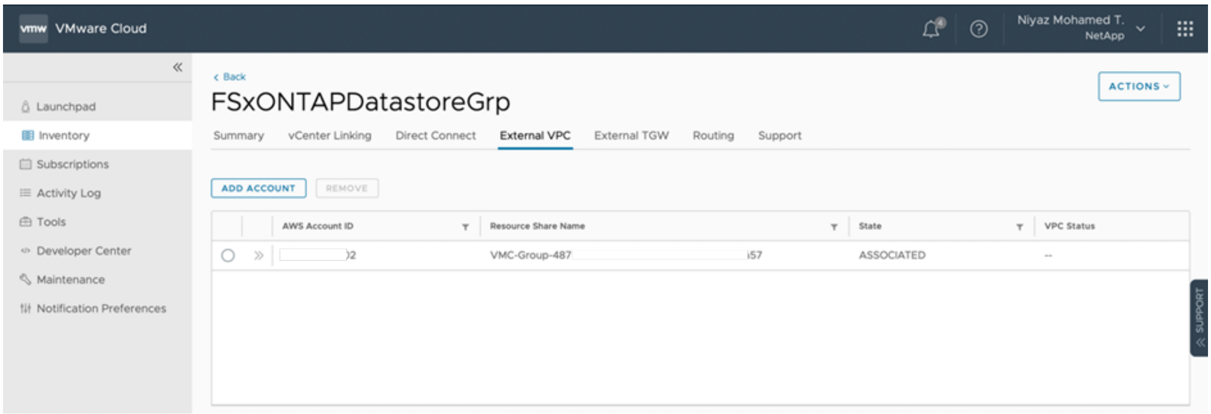

FSxONTAPDatastoreGrp

AWS Account ID	Resource Share Name	State	VPC Status
2	VMC-Group-487 57	ASSOCIATED	--

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

Create transit gateway attachment, configure routing and security groups

1. In the AWS Console, go to the VPC service page and navigate to the VPC that was used for provisioning the Amazon FSx for ONTAP file system. Here you create a transit gateway attachment by clicking **Transit Gateway Attachment** on the navigation pane on the right.

2. Under **VPC Attachment**, make sure that DNS Support is checked and select the VPC in which FSx for ONTAP was deployed.

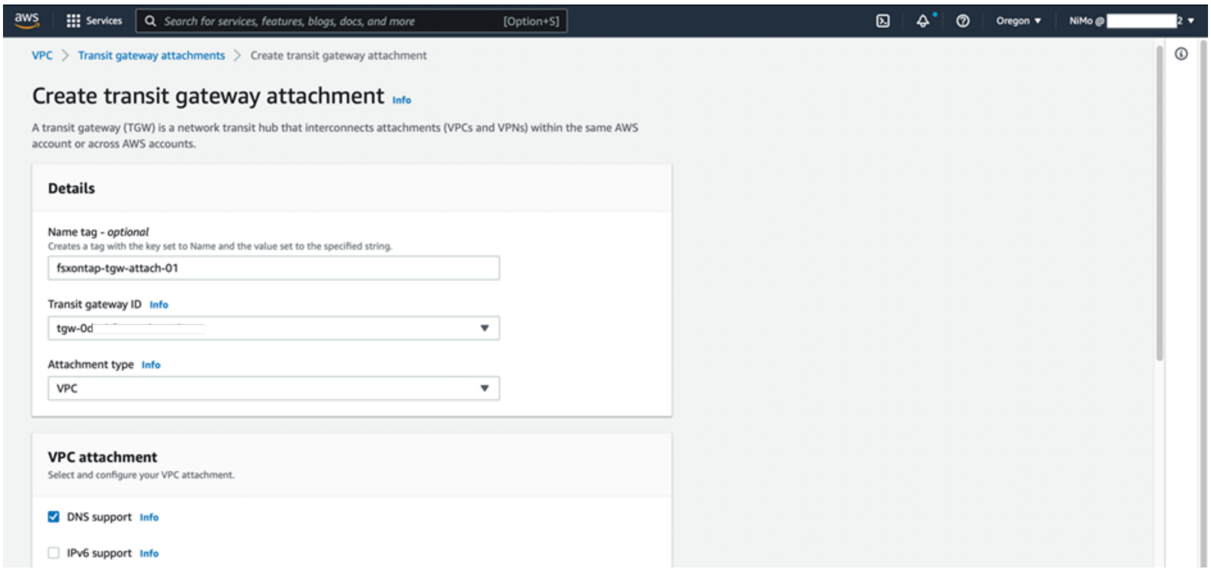
aws
Services
Search for services, features, blogs, docs, and more
[Option+S]
Oregon
NiMo @
VPC
Transit gateway attachments
Create transit gateway attachment
Info

Create transit gateway attachment

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

Details

Name tag - *optional*
Creates a tag with the key set to Name and the value set to the specified string.
fsxontap-tgw-attach-01

Transit gateway ID
tgw-0d

Attachment type
VPC

VPC attachment
Select and configure your VPC attachment.

☑ DNS support

☐ IPv6 support

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

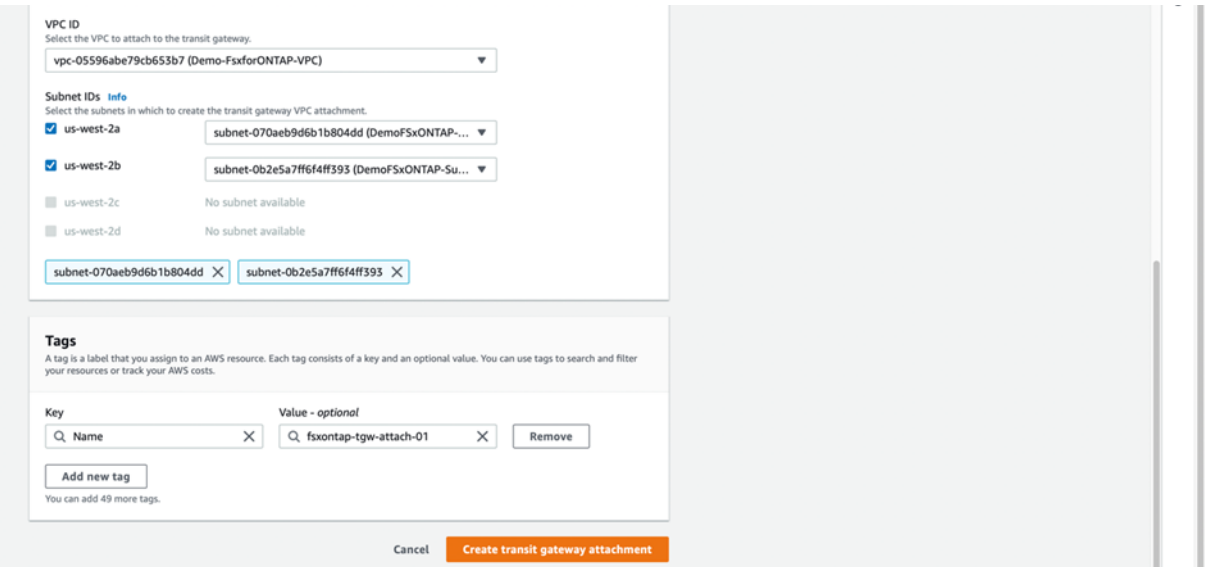

VPC ID
Select the VPC to attach to the transit gateway.
vpc-05596abe79cb653b7 (Demo-FsxforONTAP-VPC)

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

☑ us-west-2a — subnet-070aeb9d6b1b804dd (DemoFSxONTAP-...
☑ us-west-2b — subnet-0b2e5a7ff6f4ff393 (DemoFSxONTAP-Su...
□ us-west-2c — No subnet available
□ us-west-2d — No subnet available

subnet-070aeb9d6b1b804dd ✕ subnet-0b2e5a7ff6f4ff393 ✕

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

Key: Name
Value - *optional*: fsxontap-tgw-attach-01
Remove
Add new tag
You can add 49 more tags.

Cancel Create transit gateway attachment

4. Back in the VMC console, navigate to **SDDC Group** > **External VPC** tab. Select the AWS account ID used for Amazon FSx for ONTAP and click the VPC and click **Accept**.

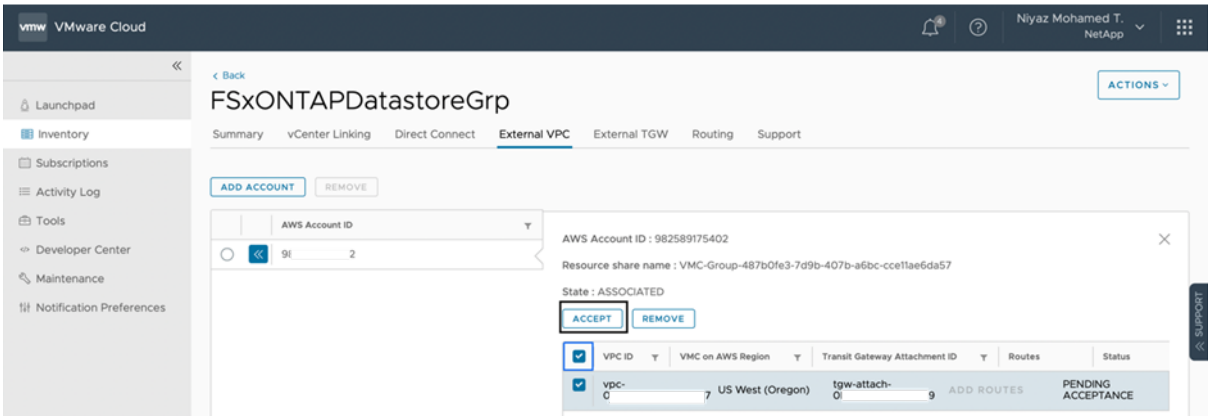

‹ Back
FSxONTAPDatastoreGrp
ACTIONS

Summary | vCenter Linking | Direct Connect | External VPC | External TGW | Routing | Support

ADD ACCOUNT REMOVE

AWS Account ID: 9[redacted]2

AWS Account ID : 982589175402
Resource share name : VMC-Group-487b0fe3-7d9b-407b-a6bc-cce11ae6da57
State : ASSOCIATED
ACCEPT REMOVE

☑	VPC ID	VMC on AWS Region	Transit Gateway Attachment ID	Routes	Status
☑	vpc-0[redacted]7	US West (Oregon)	tgw-attach-0[redacted]9	ADD ROUTES	PENDING ACCEPTANCE

Note: This option may take several minutes to appear.

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

- A route for the NFS IP range for Amazon FSx for ONTAP
- A route for the newly created AWS VPC CIDRs

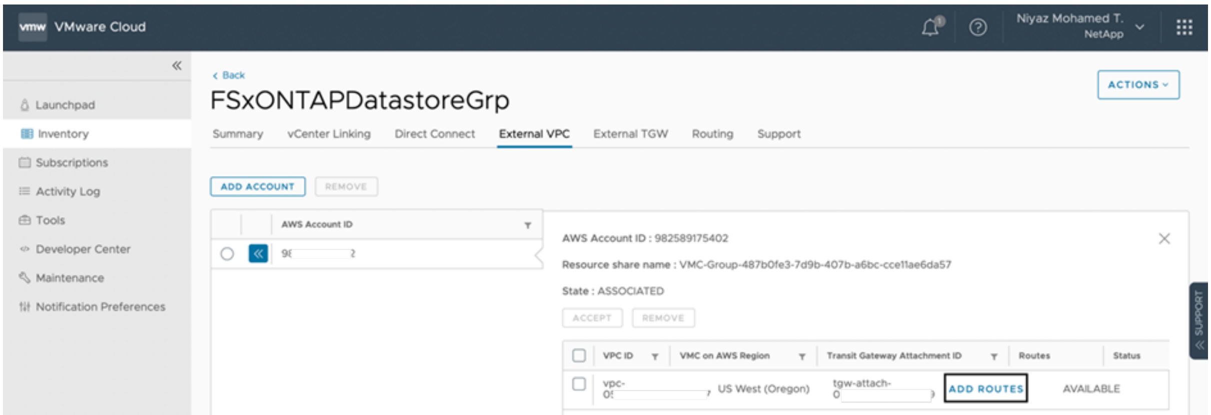

‹ Back
FSxONTAPDatastoreGrp
ACTIONS

Summary | vCenter Linking | Direct Connect | External VPC | External TGW | Routing | Support

ADD ACCOUNT REMOVE

AWS Account ID: 9[redacted]2

AWS Account ID : 982589175402
Resource share name : VMC-Group-487b0fe3-7d9b-407b-a6bc-cce11ae6da57
State : ASSOCIATED
ACCEPT REMOVE

□	VPC ID	VMC on AWS Region	Transit Gateway Attachment ID	Routes	Status
□	vpc-05[redacted]7	US West (Oregon)	tgw-attach-0[redacted])	ADD ROUTES	AVAILABLE

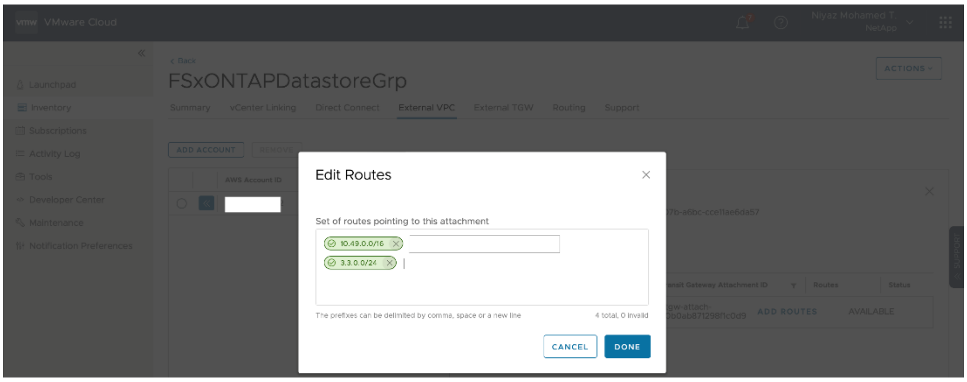

Edit Routes

Set of routes pointing to this attachment

10.49.0.0/16

3.3.0.0/24

The prefixes can be delimited by comma, space or a new line — 4 total, 0 invalid

6. In the AWS Console, create the route back to the SDDC by locating the VPC where Amazon FSx for ONTAP is provisioned in the VPC service page and select the main route table for the VPC.

7. Browse to the route table in the lower panel and click **Edit routes**.

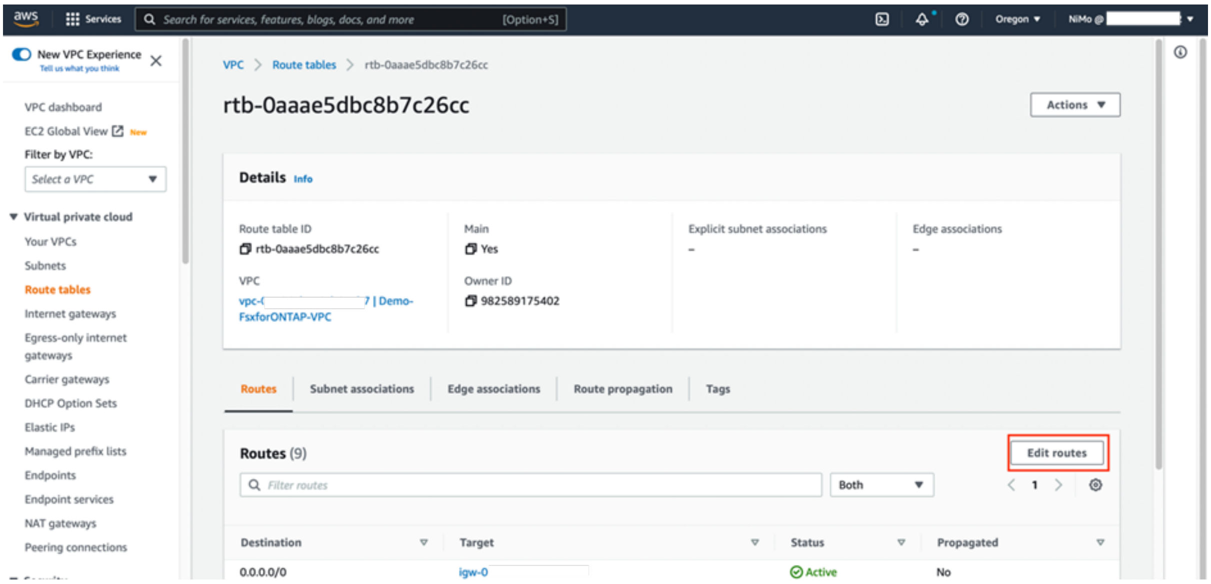

VPC > Route tables > rtb-0aaae5dbc8b7c26cc

rtb-0aaae5dbc8b7c26cc

Details

Route table ID	Main	Explicit subnet associations	Edge associations
rtb-0aaae5dbc8b7c26cc	Yes	–	–

VPC: vpc-[redacted]7 | Demo-FsxforONTAP-VPC

Owner ID: 982589175402

Routes (9)

Destination	Target	Status	Propagated
0.0.0.0/0	igw-0[redacted]	Active	No

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

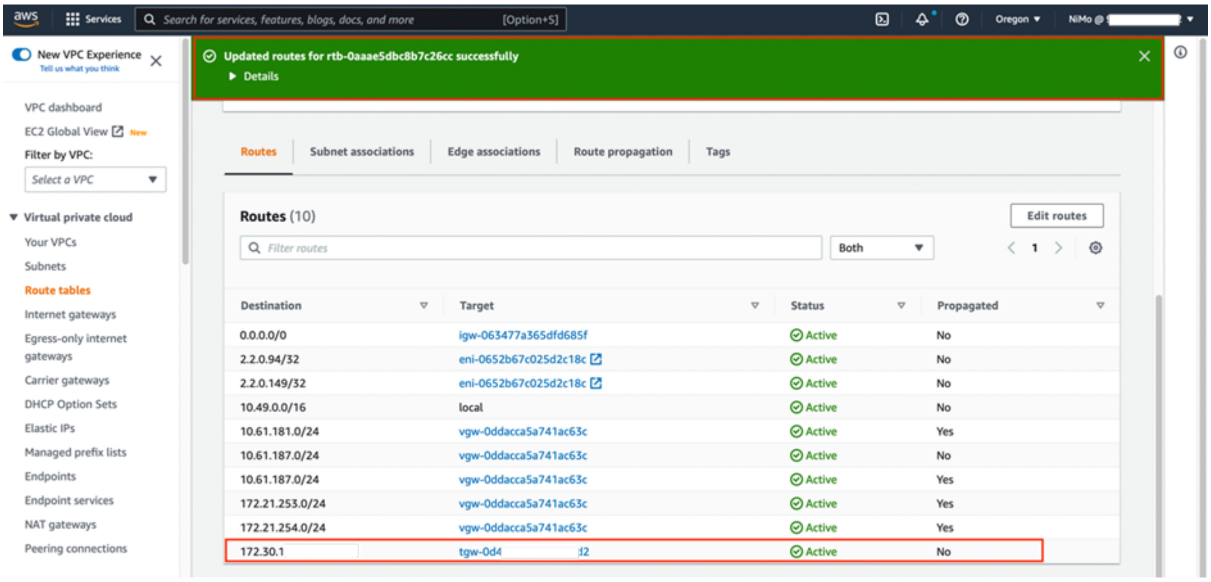

Updated routes for rtb-0aaae5dbc8b7c26cc successfully

Routes (10)

Destination	Target	Status	Propagated
0.0.0.0/0	igw-063477a365dfd685f	Active	No
2.2.0.94/32	eni-0652b67c025d2c18c	Active	No
2.2.0.149/32	eni-0652b67c025d2c18c	Active	No
10.49.0.0/16	local	Active	No
10.61.181.0/24	vgw-0ddacca5a741ac63c	Active	Yes
10.61.187.0/24	vgw-0ddacca5a741ac63c	Active	No
10.61.187.0/24	vgw-0ddacca5a741ac63c	Active	Yes
172.21.253.0/24	vgw-0ddacca5a741ac63c	Active	Yes
172.21.254.0/24	vgw-0ddacca5a741ac63c	Active	Yes
172.30.1[redacted]	tgw-0d4[redacted]2	Active	No

The next step is to verify that the security group in the associated VPC is updated with the correct inbound rules for the SDDC management subnet CIDR.

9. Update the Inbound rule with the CIDR block of the SDDC management. You could either allow all the traffic from the SDDC or specify the ports required for NFS traffic, depending on your security requirements.

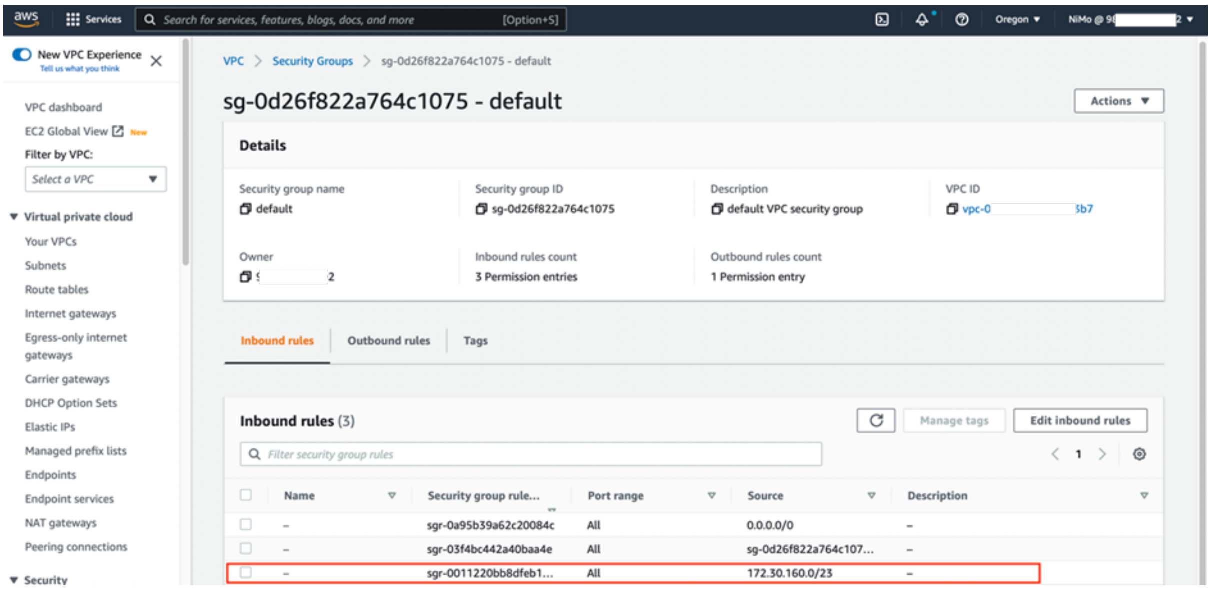

sg-0d26f822a764c1075 - default

Details

Security group name: default
Security group ID: sg-0d26f822a764c1075
Description: default VPC security group
VPC ID: vpc-0[redacted]3b7
Owner: [redacted]2
Inbound rules count: 3 Permission entries
Outbound rules count: 1 Permission entry

Inbound rules (3)

	Name	Security group rule...	Port range	Source	Description
□	–	sgr-0a95b39a62c20084c	All	0.0.0.0/0	–
□	–	sgr-03f4bc442a40baa4e	All	sg-0d26f822a764c107...	–
□	–	sgr-0011220bb8dfeb1...	All	172.30.160.0/23	–

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

This is the last step in preparing the VMware Transit Connect connectivity. With the file system configured, routes added, and security groups updated, it is time to mount the datastores.

Attach NFS volume as external storage

After the file system is provisioned and the connectivity is in place, access the VMC Console to mount NFS datastores.

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

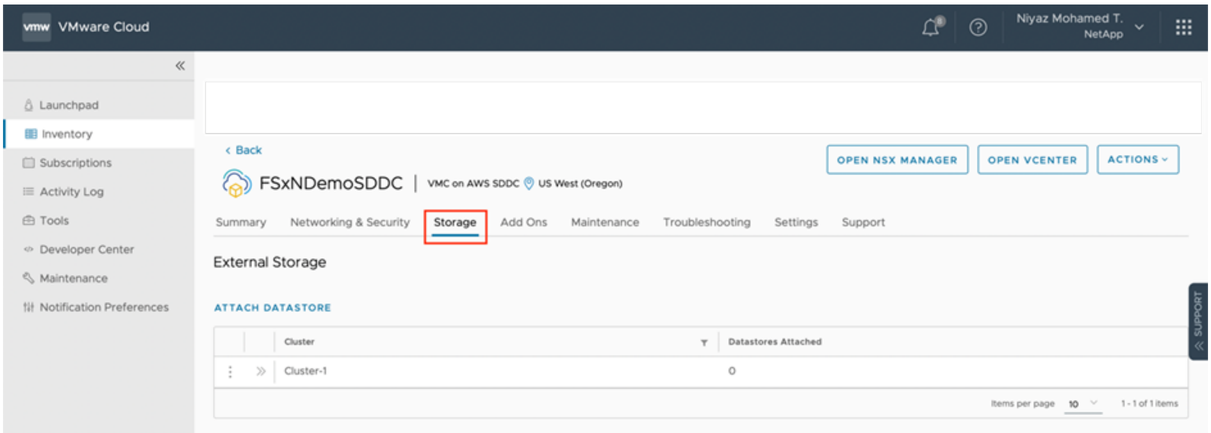

External Storage

Cluster	Datastores Attached
Cluster-1	0

Items per page 10 1 - 1 of 1 items

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

3. Click **Validate** on the right from the NFS Server address If the validation was not successful:

- Check your security group inbound rules on the AWS VPC.
- When using a multi-AZ FSx for ONTAP deployment, ensure to use the SAME route table for both FSx for NetApp ONTAP VPC subnets.
- FSx for NetApp ONTAP should have been deployed in a new AWS VPC and subnet.

© VMware LLC.

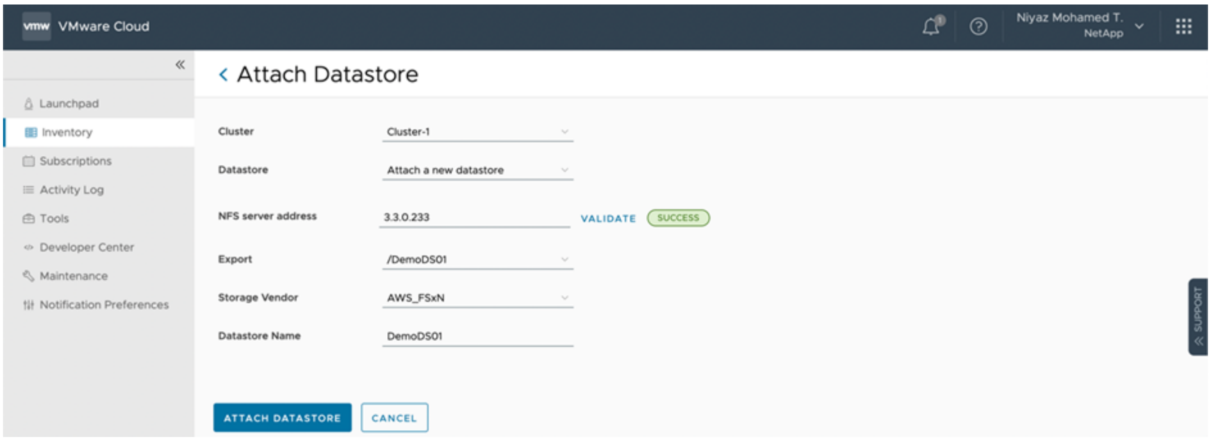

Attach Datastore

Cluster	Cluster-1
Datastore	Attach a new datastore
NFS server address	3.3.0.233 VALIDATE SUCCESS
Export	/DemoDS01
Storage Vendor	AWS_FSxN
Datastore Name	DemoDS01

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

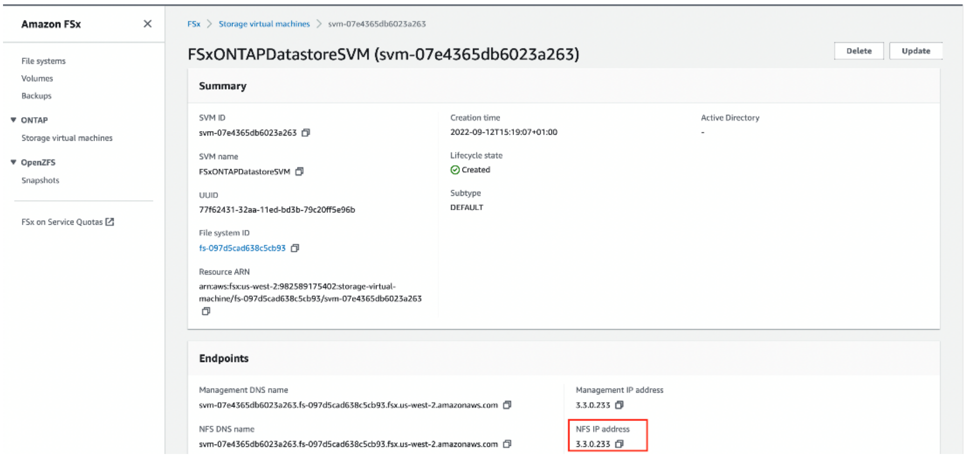

FSxONTAPDatastoreSVM (svm-07e4365db6023a263)

Summary

SVM ID
svm-07e4365db6023a263

SVM name
FSxONTAPDatastoreSVM

UUID
77f62431-32aa-11ed-bd3b-79c20ff5e96b

File system ID
fs-097d5cad638c5cb93

Resource ARN
arn:aws:fsx:us-west-2:982589175402:storage-virtual-machine/fs-097d5cad638c5cb93/svm-07e4365db6023a263

Creation time
2022-09-12T15:19:07+01:00

Lifecycle state
Created

Subtype
DEFAULT

Active Directory
-

Endpoints

Management DNS name
svm-07e4365db6023a263.fs-097d5cad638c5cb93.fsx.us-west-2.amazonaws.com

NFS DNS name
svm-07e4365db6023a263.fs-097d5cad638c5cb93.fsx.us-west-2.amazonaws.com

Management IP address
3.3.0.233

NFS IP address
3.3.0.233

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

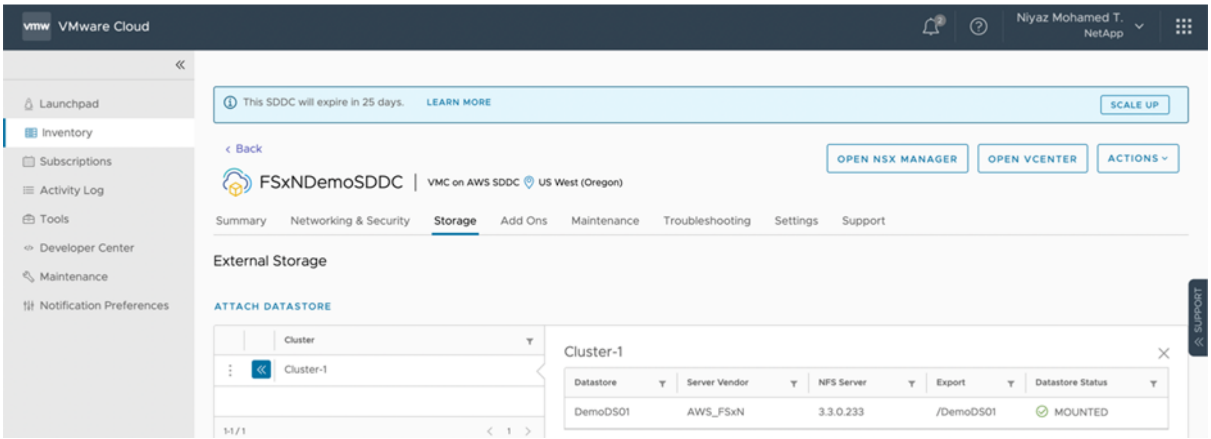

This SDDC will expire in 25 days.

FSxNDemoSDDC | VMC on AWS SDDC US West (Oregon)

External Storage

Cluster
Cluster-1

1-1 / 1

Cluster-1

Datastore	Server Vendor	NFS Server	Export	Datastore Status
DemoDS01	AWS_FSxN	3.3.0.233	/DemoDS01	MOUNTED

5. Validate the NFS datastore by accessing vCenter Web Client as shown below:

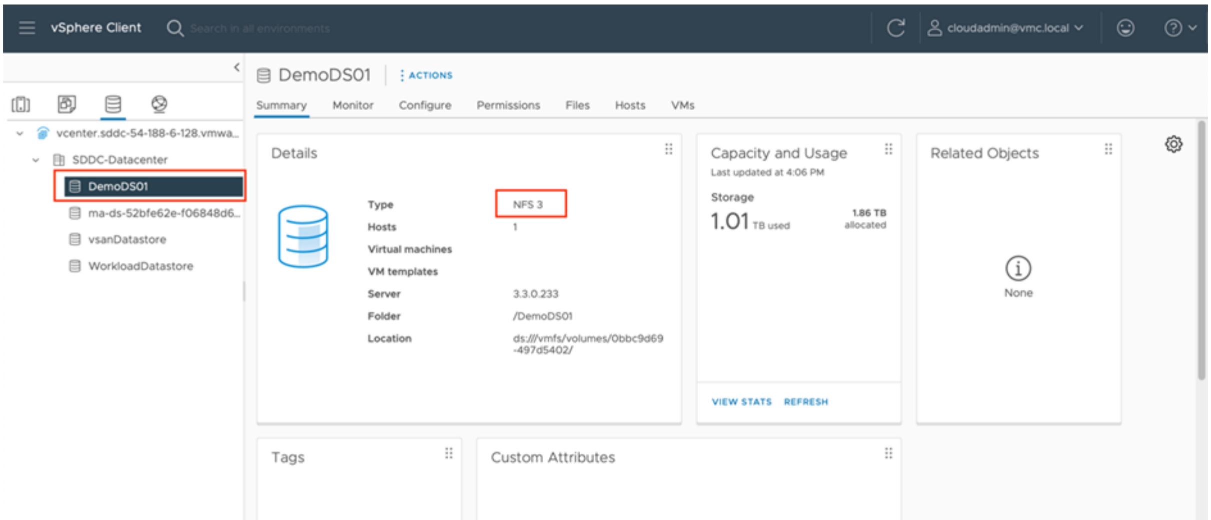

Type	NFS 3
Hosts	1
Virtual machines	
VM templates	
Server	3.3.0.233
Folder	/DemoDS01
Location	ds:///vmfs/volumes/0bbc9d69 -497d5402/

Last updated at 4:06 PM

Storage

1.01 TB used

1.86 TB allocated

Features, Considerations, and Integrations

Deciding on Storage Connectivity

Choosing between VPC Peering and VMware Transit Connect is contingent upon the storage needs of the SDDC and the intended deployment architecture of the FSxN file system. For customers focused solely on decoupling their SDDC storage using external NFS datastores, VPC Peering is advantageous for decreasing the overall cost by removing the data transfer charges linked to VMware Transit Connect. This is a network connection that is free to setup and use was long as the SDDC and the FSxN file system reside within the same AZ. However, it's essential to recognize, though, that this choice is exclusive for NFS datastores only and necessitates that the FSxN file system be implemented in a single-AZ architecture. Additionally, the SDDC can only be linked to a single designated FSxN VPC at a time. Overall, this connectivity model provides a balance of cost efficiency, performance, and simple setup.

In contrast, VMware Transit Connect is the preferred solution for customers looking to consolidate their connectivity approach, build scalability and have the flexibility of deploying either FSxN single-AZ or multi-AZ file systems. This method allows customers to connect multiple SDDCs and FSxN designated VPCs, facilitating scalability and centralized management. It enables the convergence of both workload and storage traffic over a singular network path, offering a comprehensive network solution that supports external NFS datastores, Guest-OS storage, and workload connectivity within a unified, scalable architecture. Nevertheless, this convenience comes with associated costs based on the number of TGW attachments and the quantity of data (measured in GB) transferred.

Switching from vTGW to VPC Peering

Switching from vTGW to VPC Peering is a straightforward and, most importantly, non-disruptive process for our customers. This seamless migration consists of two key phases. Initially, it begins with a request to create and configure a VPC peering connection, which extends network connectivity between your Amazon FSx for NetApp ONTAP file system and the VMware Cloud on AWS SDDC. The steps for establishing this network connectivity align with those previously outlined in this guide. Customers should initiate a request and then accept the VPC Peering connection before applying the final security and network configurations. Once your VPC peering connection is Active, you can proceed to the switchover phase, which involves a simple update to the route table used by the Amazon VPC through which the file system is accessible. In this update, set the VPC Peering connection as the new target for the SDDC Management CIDR prefix, replacing the vTGW attachment. The switchover is seamless and non-disruptive to the workloads hosted on the NFS datastore.

Please review the following steps to summarize the necessary actions:

Phase 1: VPC Peering Preparation

1. Identify the target SDDC and Amazon FSx for NetApp ONTAP file system
2. Note the Amazon VPC associated with the file system
3. Note the VPC route table currently utilizing the TGW attachment

Phase 2: VPC Peering Configuration

1. Create VPC Peering (VMC Console)
2. Accept VPC Peering (AWS Console)

Phase 3: Switchover

1. Locate the Amazon VPC associated with the file system
2. Edit the route table of this Amazon VPC. Update the existing route entry of the SDDC Management CIDR, replacing the vTGW attachment with the VPC peering connection as the target

Verify Network Connectivity

In scenarios where the FSxN VPC is connected to the SDDC using both VPC Peering and VMware Transit Connect, with the former dedicated to NFS datastore storage and the latter to workload connectivity, it's important to ascertain the path of network traffic. To verify if traffic is passing through the peering connection or the Transit Gateway, start by examining the VPC route tables tied to the FSxN Elastic Network Interface (ENI) subnets. AWS prioritizes the most specific route that matches the traffic, known as the longest prefix match, within the VPC route tables. Review the route table to determine which CIDR Prefix is being utilized, whether it's for the VPC peering or the AWS Transit Gateway.

Furthermore, to confirm the actual traffic flow, whether it's via the VPC Peering connection or the Transit Gateway, consider setting up a test EC2 instance in the same subnet as the FSxN ENI interfaces. Then, implement Transit Gateway Flow Logs to observe whether traffic from this test instance is going through the TGW. If the traffic is correctly using the peering connection, which should have the most specific route, you would not detect the Source IP of the EC2 instance in the flow logs attributed to the Transit Gateway.

Data Transfer Costs

Data transfer charges are often overlooked while architecting a solution in AWS. Considering data transfer charges while making architectural decisions can help save costs. This section will help identify potential data transfer charges you may encounter while operating your workload on VMware Cloud on AWS.

VPC Peering

This feature proves to be the most cost-efficient solution for external datastore storage. The data transfer of storage traffic between VMware Cloud on AWS and an Amazon FSx for NetApp ONTAP file system within the same Availability Zone is **free**.

VMware Transit Connect

VMware Transit Connect enables customers to build high-speed, resilient connections between their VMware Cloud on AWS SDDCs and VPC resources hosting the FSx for ONTAP file system. This capability is enabled via a feature called SDDC Groups. Behind the simplification that SDDC Groups provide is the instantiation of a VMware Managed AWS Transit Gateway, a vTGW. It is automatically deployed when an SDDC group is created. The vTGW provides connectivity between the Amazon FSx for ONTAP NFS volumes and ESXi hosts running in the SDDC.

Data flows from the datastore mounted on the ESXi host to the VMware Transit Connect. From here the vTGW routes the storage traffic to the AWS Transit Connect attachment in the customer-managed VPC where it is finally routed to the SVM floating IP associated with the Active FSx for ONTAP ENI to the FSx for ONTAP filesystem. The data flow is depicted on the reference architecture diagram.

The transit gateway cost is based on the number of attachments made to the Transit Gateway and the amount of egress traffic that flows through AWS Transit Gateway. You can find more information in the following guide.

The table below will help with estimating the data transfer cost for “low IO” workloads. The recommendation is to engage AWS, VMware, or NetApp representatives to help with a detailed projection.

Projected NFS datastore Capacity in TB	**R/W Change rate in %***	**Egress Capacity (daily)**	**No. of TGW attachments**	**TGW single attachment monthly cost in $****	**Total TGW data processing monthly cost in $ ****	**Total TGW data processing and attachments monthly cost in $ ****
50TB	35%	17.5TB	2	$36.50	$10,752.00	**$10,825**
100TB	35%	35TB	2	$36.50	$21,504.00	**$21,577**
200TB	35%	70TB	2	$36.50	$43,008.00	**$43,081**

* Change rate Includes both daily churn rate, including updates and the amount of reads from the existing data (covers both reads and writes). Your actual charge rate will vary. Use this as an estimate only and use appropriate tools to identify the daily network transfers to accurately identify the transfer amount.

** We used AWS US East (N Virginia) Region for cost estimates. For other regions, costs might be different.

Note: Prices are actual as of September 2022. Prices vary per AWS Regions and are subject to change.

Another approach to identifying the accurate total data transfer capacity is to leverage vROPs, Live Optics, NetApp Cloud Insights, or similar performance tools to capture the storage network In/Out for virtual machines or at the ESXi host level. Once the data is collected, use the AWS pricing calculator to calculate the egress cost.

Pricing example:

Consider a real-time scenario wherein data egress on VMware Cloud on AWS SDDC is 133MB/s and data egress from FSx for NetApp ONTAP is 140MB/s, then the net data transfer is 273 MB/s. This would equate to 23.5TB of data per day spread across two TGW attachments.

Total data processed per all Transit Gateway attachments: 23,5TB x 30 days = 705TB per month. Let us convert to GB: 705TB x 1024 GB = 721,920 GB per month.

Pricing calculations*

- Total Monthly Transit Gateway attachments cost: 730 hours in a month x 0.05 USD* = 36.50 USD x 2 TGW attachments = 73 USD.
- Total Monthly Transit Gateway data processing cost: 721,920 GB per month x 0.02 USD = 14,438.4 USD.
- **Total Transit Gateway monthly cost**: 73 USD attachments cost + 14,438.4 USD data processing cost = **14,511.4** USD.

Total cost:

For the scenario outlined above, estimated monthly cost is **14,511.4** USD*.

* We used AWS US East (N Virginia) Region for cost estimates. For other regions, costs might be different.

Performance considerations

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

To leverage more available bandwidth with Amazon FSx for ONTAP volumes you can configure multiple datastores, with each datastore using separate connections between the ESXi hosts and the FSx for ONTAP filesystem.

Note: VMware Cloud on AWS supports up to four NFS datastores per each vSphere cluster in the SDDC.

Performance optimization

Although the recommended number of VM per NFS datastore is subjective, numerous factors determine the optimum number of VMs that can be placed on each datastore. Although most administrators only consider capacity, the amount of concurrent I/O being sent to the VMDKs is one of the most key factors for overall performance. The ESXi host has various mechanisms to ensure fairness between VMs competing for datastore resources. However, the easiest way to control performance is by regulating the number of virtual machines that are placed on each datastore. If the concurrent virtual machine I/O patterns are sending too much traffic to the datastore, the disk queues fill, and higher latency are generated.

Volume and datastore sizing

When a volume is created on Amazon FSx for ONTAP for datastore purposes, the best practice is to create a volume no larger than required. A general recommendation is to begin with a small datastore capacity and increase it as needed. Right sizing datastores prevent accidentally placing too many virtual machines on the datastore and decreases the probability of resource contention. Since datastores and VMDK sizes can be easily increased, if a VM needs extra capacity, it is not necessary to size datastores larger than required. For optimal performance, the best practice is to increase the number of datastores rather than increase their size as an interim measure.

Capacity monitoring for overprovisioned NFS volumes

To ensure your storage system operates efficiently and you accurately monitor available space, it's critical not to overprovision your NFS volume or datastore beyond the total SSD capacity of your FSx for NetApp ONTAP file system. This approach allows you to leverage VMware vCenter for straightforward storage space management within the SDDC.

Overprovisioning the NFS volume beyond the SSD capacity of the FSx for NetApp ONTAP system can lead to reporting discrepancies. VMware vCenter might indicate a higher usage level than what FSx for NetApp ONTAP actually reports. This discrepancy could cause vCenter to falsely alert you of nearing full capacity, even though ample space remains according to the FSx for NetApp ONTAP system viewed on the AWS console. Ignoring this issue might result in unnecessary warnings and potential operational disruptions, particularly if vCenter signals that the datastore is close to 95% capacity.

If your NFS volume is overprovisioned, consider these actions:

- Increase the SSD capacity of your FSxN file system to accommodate your storage needs.
- For scenarios requiring overprovisioning, such as backup or disaster recovery, activating logical space reporting on FSxN may be beneficial. This feature, when paired with direct monitoring via AWS CloudWatch—configurable through the AWS Management Console—enables strategic storage management. The logical space reporting can also be used in scenarios where the datastore volume and SSD tier is rightsized, however customer wants to hide storage efficiency savings from vSphere. Read the following guide for more information on enabling logical space reporting.

Note: An average production size for an FSx for ONTAP NFS datastore is between 10TB to 20TB.

Contact AWS, VMware, and NetApp to plan and size storage and host requirements accurately. We recommend identifying storage performance requirements before finalizing the datastore layout for production deployments.

Increasing the size of the datastore

Resizing the datastore volume is completely transparent to the SDDC. Increase the size of NFS datastores by resizing the volume from the AWS console or by using the FSx for ONTAP CLI. After you are done, access vCenter, go to the datastore tab, right-click the appropriate datastore and select Refresh Capacity Information. This process is also completely transparent to VMs and applications consuming the datastore.

Migration

One of the most common use cases is migration based on various factor. Customers can use VMware HCX or other migrations tools to migrate VMs.

For additional information about migration options and on how to migrate workloads from on-premises to VMware Cloud on AWS, see VMware Cloud TechZone, VMware HCX User Guide, and Migrate workloads to FSx for ONTAP datastore using VMware HCX Guide.

Advanced Amazon FSX for NetApp ONTAP options

Backing up VMs and quickly recovering them are among the great strengths of AWS FSx for ONTAP datastores. Use Snapshot copies to make quick copies of your VMs or even the whole NFS datastore without affecting performance, and then send them to a secondary AWS region of your choice using NetApp SnapMirror replication for disaster recovery purposes. This approach minimizes storage space and network bandwidth by storing changed information only.

Use Amazon FSx for ONTAP snapshot copies for general protection and use application tools like NetApp SnapCenter (for guest-connected storage) or third-party tools to protect virtual machines and transactional data such as Microsoft SQL Server or Oracle residing on the guest VMs.

Note: NetApp SnapCenter Plug-in for VMware vSphere is currently not supported with AWS FSx for ONTAP

ROI Calculations

You can calculate how much you could save by using VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP. VMware, NetApp, and AWS have built a new ROI tool based on the current VMC Cloud Sizer to help estimate the savings potential. The tool can be accessed here.

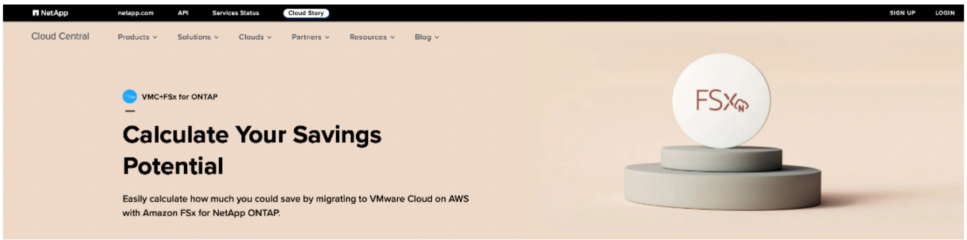

VMC+FSx for ONTAP

Calculate Your Savings Potential

Easily calculate how much you could save by migrating to VMware Cloud on AWS with Amazon FSx for NetApp ONTAP.

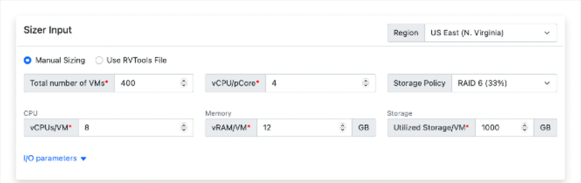

Sizer Input

Region: US East (N. Virginia)

● Manual Sizing ○ Use RVTools File

Total number of VMs*: 400
vCPU/pCore*: 4
Storage Policy: RAID 6 (33%)

CPU
vCPUs/VM*: 8

Memory
vRAM/VM*: 12 GB

Storage
Utilized Storage/VM*: 1000 GB

The sizer works with manual Inputs as well as with RVtools and considers reserved instance pricing as well.

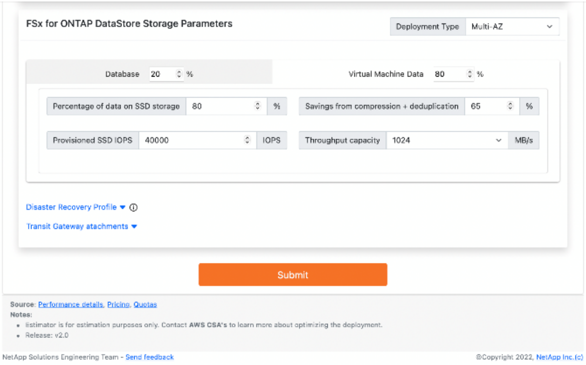

FSx for ONTAP DataStore Storage Parameters

Deployment Type: Multi-AZ

Database: 20 %

Virtual Machine Data: 80 %

Percentage of data on SSD storage: 80 %

Savings from compression + deduplication: 65 %

Provisioned SSD IOPS: 40000 IOPS

Throughput capacity: 1024 MB/s

Source: Performance details, Pricing, Quotas

Notes:

- Estimator is for estimation purposes only. Contact **AWS CSA's** to learn more about optimizing the deployment.
- Release: v2.0

NetApp Solutions Engineering Team - Send feedback

©Copyright 2022, NetApp Inc.(c)

Here is a quick sizing result which highlights the TCO optimization whilst using FSx for ONTAP as the supplemental datastore.

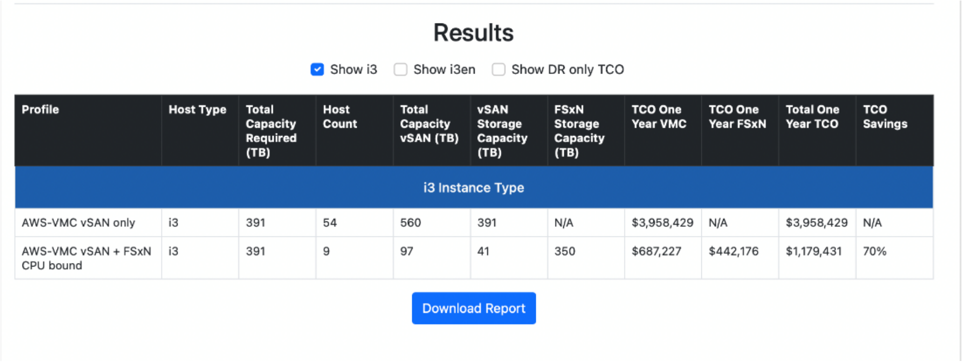

Results

☑ Show i3 □ Show i3en □ Show DR only TCO

Profile	Host Type	Total Capacity Required (TB)	Host Count	Total Capacity vSAN (TB)	vSAN Storage Capacity (TB)	FSxN Storage Capacity (TB)	TCO One Year VMC	TCO One Year FSxN	Total One Year TCO	TCO Savings
i3 Instance Type										
AWS-VMC vSAN only	i3	391	54	560	391	N/A	$3,958,429	N/A	$3,958,429	N/A
AWS-VMC vSAN + FSxN CPU bound	i3	391	9	97	41	350	$687,227	$442,176	$1,179,431	70%

Conclusion

VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP is a jointly engineered, AWS-managed external NFS datastore built on NetApp's ONTAP file system that can be attached to VMware Cloud on AWS vSphere cluster. It provides customers with a flexible, high-performance virtualized storage infrastructure that scales independently of compute resources.

Takeaways

1. Lower TCO: Optimize costs for storage heavy workloads by attaching a cost-effective external storage option that scales independently of compute resources thus avoiding any unused/wasted capacity. Customers can also leverage NetApp ONTAP's comprehensive and consistent data management capabilities such as snapshots, clones, replication, etc., to simplify and make data management more agile and to lower costs.
2. Increased flexibility: Attach storage to multiple SDDC's and have the flexibility to define performance and storage requirements based on specific workload needs.
3. Familiar technology leading to better productivity: Reduce risks and complexity associated with migration and modernization initiatives by using familiar VMware and NetApp ONTAP technology with minimal learning curve while migrating storage intensive workloads to the cloud. No need to reengineer the data layer or modify the application code, minimizing the need to re-architect the storage design.
4. Improved time to market: Increase flexibility to scale the storage infrastructure and familiar technology across hybrid cloud environment allows customers to accelerate their workload migrations and bring out new applications and features faster.

You can take the VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP interactive demo to gain additional experience with this feature.

Summary and Additional Resources

Additional Resources

For more information about VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP, you can explore the following resources:

- VMware Cloud Techzone
- How Amazon FSx for NetApp ONTAP works
- What is Amazon FSx for NetApp ONTAP?

Changelog

01 October 2022	

About the Author and Contributors

- Niyaz Mohamed (NiMo), Principal Architect, Cloud & Hybrid Cloud Solutions, NetApp
- Oleg Ulyanov, Staff Cloud Solutions Architect, Technical Marketing, VMware
- Kiran Reid, Senior Partner Solutions Architect, AWS
- Osama Masfary, Staff Cloud Solutions Architect, Technical Marketing, VMware

Feedback

Your feedback is valuable.

To comment on this paper, contact VMware Public Cloud Technical Marketing

© VMware LLC.

Copyright © 2005-2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.