

# Dell Technologies Cloud Validated Designs: Dell EMC Storage with VMware Cloud Foundation

## Abstract

This document describes how to configure the Dell Technologies™ Cloud Validated Design consisting of Dell EMC™ storage with VMware® Cloud Foundation.

October 2020

## Revisions

Date	Description
February 2019	Initial release
March 2019	SC Series added
April 2019	PowerMax NFS, VxFlex, and XtremIO added
July 2019	PowerMax and Dell EMC Unity iSCSI steps added
August 2019	Introducing Fibre Channel storage as principal storage for VMware Cloud Foundation workload domains
December 2019	VCF 3.9 updates
April 2020	PowerStore added
May 2020	Added VMware® vSphere® Virtual Volumes™ (vVols) guidance
June 2020	Clarified principal and supplemental storage; PowerFlex branding update
October 2020	Added VCF 4.1 updates for vVol support

## Acknowledgments

Authors: Nicholas Ringborg: Dell Technologies Cloud Validated Design sections

Darin Schmitz: Dell Technologies Storage configuration sections

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright ©2019-2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [10/16/2020] [Technical White Paper] [H17921]

# Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents .....	3
<b>1 Introduction.....</b>	<b>5</b>
1.1 The need for a hybrid cloud model.....	5
1.2 VMware Cloud Foundation .....	6
1.3 Dell Technologies Cloud.....	6
1.4 Dell Technologies Cloud Validated Designs .....	7
<b>2 Dell EMC storage with VMware Cloud Foundation .....</b>	<b>8</b>
2.1 Principal and supplemental storage for workload domains.....	8
2.2 Customer use cases.....	8
2.2.1 Storage-intensive applications.....	9
2.2.2 Investment protection .....	9
2.2.3 Deployment automation.....	9
2.3 Protocol deployment options .....	10
<b>3 Prerequisites.....</b>	<b>11</b>
<b>4 Workload domain configuration with NFS .....</b>	<b>12</b>
4.1 VMware Cloud Foundation network pool configuration.....	12
4.2 Dell EMC storage configuration.....	14
4.2.1 Dell EMC Unity NFS share creation .....	14
4.2.2 PowerMax NFS share creation.....	18
4.2.3 PowerStore NFS share creation.....	24
4.3 VMware Cloud Foundation workload domain deployment.....	29
<b>5 Workload domain configuration with FC, iSCSI, or vVol.....</b>	<b>31</b>
5.1 Using Fibre Channel as principal storage .....	31
5.1.1 Commission hosts for VMFS on FC .....	31
5.1.2 Create the VMFS datastore.....	32
5.1.3 Workload domain creation.....	32
5.2 Using Virtual Volumes (vVols) as principal storage.....	32
5.2.1 Create a Network Pool .....	32
5.2.2 Configure ESXi hosts for vVols .....	33
5.2.3 Create a Storage Container.....	36
5.2.4 Add Storage Settings.....	36
5.2.5 Commission hosts for vVol .....	38

## Table of contents

5.2.6 Workload domain creation .....	38
5.3 Dell EMC storage configuration .....	38
5.3.1 Dell EMC Unity .....	39
5.3.2 PowerMax .....	39
5.3.3 PowerStore .....	40
5.3.4 SC Series .....	41
5.3.5 PowerFlex System .....	41
5.3.6 XtremIO .....	42
5.4 Datastore creation .....	43
<b>A Technical support and resources .....</b>	<b>45</b>
A.1 Related resources .....	45

# 1 Introduction

Digital Transformation is no longer a buzz word, but rather a direct reflection of the massive changes that are happening across every industry. Organizations are making substantial investments in IT to succeed in this environment require leveraging data as an asset across all aspects of their business. There has also been substantial movement towards onboarding cloud services to accelerate innovation, expand to new locations, keep up with massive growth, and reduce costs of delivering IT.

As a result, IT departments are under significant pressure today. They are no longer treated as a cost center, but instead IT is playing a significant role in shaping their organization's overall digital strategy. As organizations continue to grow their application ecosystems, it represents a large shift for many IT departments. With significant challenges on the way, such as aging data centers and legacy infrastructure, it makes it difficult to support their transformation ambitions.

IT departments are now forced to reexamine what they do and how they do it—prioritizing their business objectives to deliver greater efficiency, predictability, and business agility.

## 1.1 The need for a hybrid cloud model

The explosion of data and devices in addition to the continued focus on controlling and reducing operating expenses has accelerated the adoption of both public and private cloud computing. Making data-driven decisions and acting quickly on new ideas helps to deliver outstanding customer and user experiences and achieve success. To support this new paradigm, IT is addressing the needs of both traditional three-tier applications and new cloud-native applications, each with a different set of priorities.

For existing workloads, IT departments are looking to reduce costs and improve performance and efficiency. While for cloud-native workloads, the priorities focus around access to new cloud services such as containers for AI or ML platforms. To address the specific needs of workloads, organizations are moving towards adopting public and private clouds.

However, operating in multiple clouds comes with another set of challenges including operational silos, different management and operations tools, complex workload migrations, security concerns, and inconsistent SLAs. These problems are byproducts of the lack of consistent infrastructure and operations across clouds, slowing cloud adoption and limiting its effectiveness.

To overcome this complexity, the ideal choice for organizations is to embrace a consistent hybrid cloud strategy for both public and private clouds, allowing for the optimal deployment of workloads. According to research from analyst firm ESG<sup>1</sup>, half of organizations formulating hybrid cloud strategies have cited seamless compatibility with their on-premises infrastructure as their most important consideration. In other words, they need a hybrid cloud strategy that eliminates multi-cloud complexity, while providing flexible deployment options for high-value workloads.

---

<sup>1</sup> IDC Worldwide Quarterly Cloud IT Infrastructure Tracker, Q1 2019, June 2019

## 1.2 VMware Cloud Foundation

VMware® is a leader in providing both the virtualization and management software that support a software-defined data center (SDDC). The VMware vision of a modern data center starts with a foundation of software-defined infrastructure that is simple to manage, monitor, and operate. The architecture for the SDDC empowers companies to run hybrid clouds, delivering scalability, automation, and agility. It is based on well-established products from VMware, such as vSphere®, vSAN™, and NSX®, providing compute, storage, and networking virtualization. Together with the vRealize® Suite, VCF brings additional management, self-service, automation, intelligent operations, life-cycle management, and financial transparency. With VMware Cloud Foundation, IT departments now have a powerful operational hub for their hybrid cloud environments and external storage arrays. This advantage helps deliver a strong foundation to deploy and manage both traditional and cloud-native workloads.




Figure 1 VMware software-defined data center architecture

## 1.3 Dell Technologies Cloud

The Dell Technologies™ Cloud is a set of cloud infrastructure solutions that are designed to make hybrid cloud environments simpler to deploy and manage. Combining the power of VMware Cloud Foundation and Dell Technologies, customers now have a consistent experience across public clouds, private clouds, and edge locations. With streamlined operations and lowered total cost of ownership (TCO), IT departments can now let business priorities determine where workloads reside. This flexibility ensures they are delivering the reliable infrastructure that best meets their unique business needs.

Using Dell Technologies Cloud, customers can deploy their workloads on-premises or use cloud environments—with VMware Cloud Foundation providing the consistent operational experience across all. The hybrid cloud solution is delivered through unique integration of hardware, software, services, and consumption options from Dell Technologies and VMware. Customers can consume in the way that aligns best with their workload requirements and business objectives:

- **Dell Technologies Cloud Platform** with Enterprise PKS, Tanzu, and VMware Cloud Foundation on VxRail: VxRail is the only fully integrated HCI system jointly engineered with VMware, delivered as a turnkey option. VxRail supports external storage arrays as supplemental storage only. vSAN Ready Nodes are required for Dell EMC arrays to be used as principal storage for a workload domain.

- **Dell Technologies Cloud Datacenter-as-a-Service** with VMware Cloud on Dell EMC: A fully managed service for data center and edge locations.
- **Dell Technologies Partner Clouds** with support for major cloud providers and over 4,200 additional cloud partners: Available by subscription, lease, or usage-based.
- **Dell Technologies Cloud Validated Designs**<sup>2</sup> using Dell EMC storage, compute, and networking are pretested infrastructure with deployment guidance.

## 1.4 Dell Technologies Cloud Validated Designs

Dell Technologies Cloud Validated Designs enable customers to bring cloud to a broader set of workloads that require independent scaling of storage and compute. This new option to consume Dell Technologies Cloud provides deployment guidance for pretested Dell EMC storage, compute, and networking infrastructure that has been validated with VMware Cloud Foundation. Using Validated Designs, customers can now support new and legacy workloads that have infrastructure-intensive requirements in the most efficient way possible.

Benefits of Validated Designs include the following:

- Rapid time-to-value with pretested infrastructure and deployment guidance
- Excellent performance with independent scaling of storage and compute
- Ability to leverage existing investments for hybrid cloud environments

---

<sup>2</sup> A Dell Technologies Cloud Validated Design does not equal a VMware Validated Design (VVD). A VVD is a software bill of materials (BoM) with related documentation to guide customers building their own software-defined data center (SDDC). Dell Technologies Cloud Validated Designs are focused on infrastructure, with related documentation to guide customers building out their own on-premises infrastructure.

## 2 Dell EMC storage with VMware Cloud Foundation

VMware Cloud Foundation brings immense value to simplify operations, manageability, and automation for customers looking to deploy and manage both traditional and modern applications. To address performance and scalability challenges with certain applications, VMware introduced support for external storage to address unique requirements.

Validated Designs are available for Dell EMC PowerMax™, Dell EMC PowerStore™, Dell EMC Unity™ XT, and Dell EMC PowerFlex™.

Validated Storage Product	Workload Domain Use
Dell EMC PowerMax	Principal / Supplemental
Dell EMC PowerStore	Principal / Supplemental
Dell EMC Unity XT	Principal / Supplemental
Dell EMC PowerFlex	Supplemental

Figure 2 Workload domain usage per Dell EMC product

### 2.1 Principal and supplemental storage for workload domains

Continuing a history of collaboration with VMware, Dell EMC has qualified external storage solutions for VMware Cloud Foundation workload domains. Originally, the ability to use external NFSv3 as principal with FC storage as supplemental storage became available with VMware Cloud Foundation 3.5.1. Later, the use of VMFS on FC storage as principal storage in workload domains is available with VMware Cloud Foundation version 3.9, and vVols with VCF 4.1.

VMware Cloud Foundation supports four protocols for principal storage, and six as supplemental storage for workload domains using PowerMax, PowerStore, and Dell EMC Unity XT. Also, PowerFlex system is qualified for deployment into workload domains as supplemental storage.

The principal and supplemental storage terms merely describe when the storage is configured with the workload domain. The characteristics of these storage types have no relation to features or performance profiles, as both can be equally resilient and performant. Principal storage is used as the first datastore to initially deploy a workload domain, while supplemental storage is any subsequent storage added after the workload domain is deployed. Supplemental storage including PowerFlex, VMFS on iSCSI, and NFSv4 are attached later using standard vCenter management tools.

### 2.2 Customer use cases

Customers can now use Validated Designs to build their own hybrid cloud infrastructure, combining the best of software-defined and traditional three-tier architecture. With more choices, they have deployment flexibility to meet unique external storage-intensive requirements such as the ability to scale storage capacity independent from compute capacity.

PowerMax, PowerStore, and Dell EMC Unity XT storage systems support Network File System (NFS), Fibre Channel (FC), and Virtual Volumes (vVols) as principal storage and iSCSI as supplemental storage. In addition, IP-based SAN storage products like PowerFlex are supported as supplemental storage.



	Management Domain	Workload Domains: Principal Storage	Workload Domains: Supplemental Storage
vSAN	✓	✓	✓
NFSv3	✗	✓	✓
VMFS on FC	✗	✓	✓
Virtual Volumes (vVols)	✗	✓	✓
NFSv4	✗	✗	✓
VMFS on iSCSI	✗	✗	✓

Figure 3 Protocol configuration options

### 2.2.1 Storage-intensive applications

Dell EMC storage arrays are ideal for applications with demanding throughput and capacity needs. They are powerful additions to VMware Cloud Foundation environments, delivering the flexibility to scale storage independent of compute for greater performance and application flexibility. In addition, customers can also build new levels of storage resiliency with replication for disaster recovery from on-premises data centers to multiple sites, for example.

### 2.2.2 Investment protection

For customers with existing Dell EMC storage infrastructure, they can now easily attach existing data capacity to their hybrid cloud environments. Attaching external storage enables them to take advantage of the advanced data services, performance, and capacity that PowerMax, PowerStore, Dell EMC Unity XT, and PowerFlex delivers.

### 2.2.3 Deployment automation

The vRealize Orchestrator (vRO) plug-in for PowerMax and PowerStore allows administrators to take advantage of VMware deployment tools of choice to rapidly stand up new environments with attached external storage. Customers can now run storage operations such as provisioning storage or scheduling snapshots directly from vRO. In addition, they can further automate storage-management activities by establishing workflows through a self-service portal using vRealize Automation (vRA). Workflows allow administrators to simplify the user experience and deliver a predefined catalog of items that users can deploy without prior knowledge of specific storage platforms. Examples of self-service items include the following:

- High-performance storage: Predefined storage for workloads that require low latency such as credit card authorizations for retail transactions. Users select preconfigured high-performance NVMe storage capacity that offers high read/write performance that is designed to support Online Transaction Processing (OLTP) transactional workloads.
- Data warehouse storage: Predefined storage for workloads that require large sequential read performance such as business analytics. Users select preconfigured high-capacity storage (such as NL-SAS or SATA drives) that offer great sequential read performance supporting large database queries.

## 2.3 Protocol deployment options

Principal and supplemental storage choices depend the deployment platform selected for the VCF environment. Figure 4 and Figure 5 illustrate the various options that can be used per workload domain.

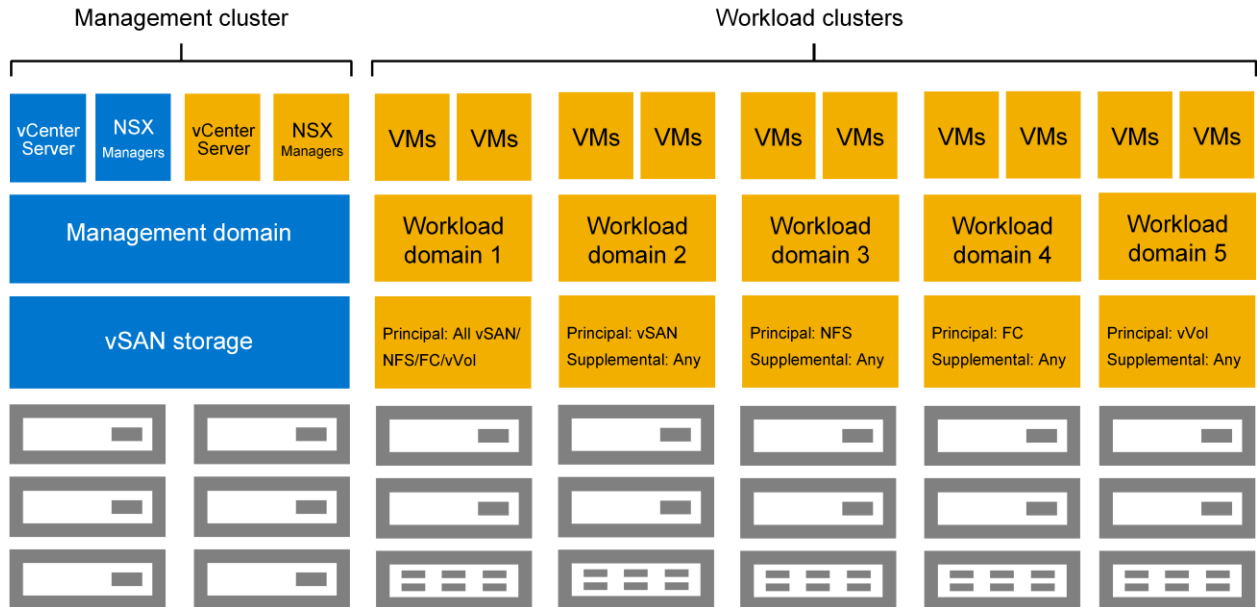


Figure 4 Dell EMC storage deployment options for vSAN ReadyNodes with VMware Cloud Foundation

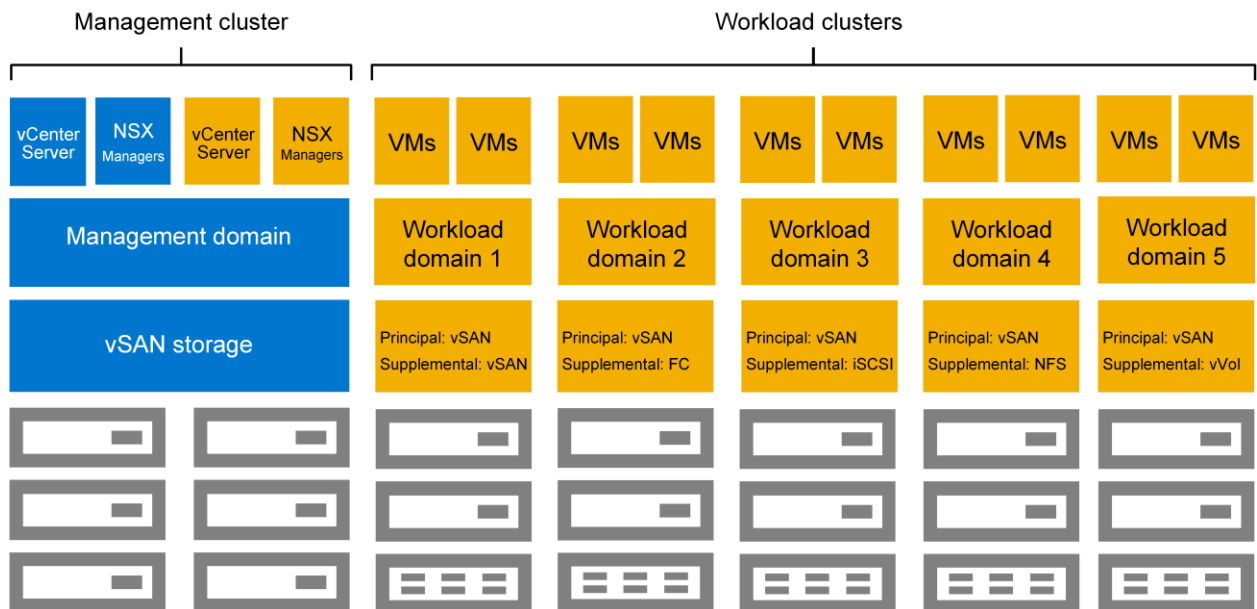


Figure 5 Dell EMC storage deployment options for VxRail with VMware Cloud Foundation

**Note:** VMFS on iSCSI and NFSv4 can only be used as supplemental storage for workload domains with supported Dell EMC arrays.

## 3 Prerequisites

The Dell Technologies Cloud Validated Designs focus on Dell EMC storage arrays using principal or supplemental storage for workload domains.

Before completing the main steps in this guide, complete the following prerequisites:

### **VMware Cloud Foundation:**

- VMware Cloud Builder has deployed the management domain.
- Hosts are commissioned and are ready to be added to a new workload domain.

### **Dell EMC storage:**

- Storage networking is configured to allow connectivity between the Dell EMC arrays and the VMware Cloud Foundation workload hosts.
  - For NFS and iSCSI configurations, it is recommended to dedicate a workload domain VLAN for an IP-based storage network pool.
  - For Fibre Channel configuration, the workload host initiator ports must be zoned with the target ports for the array within the FC switch networks.

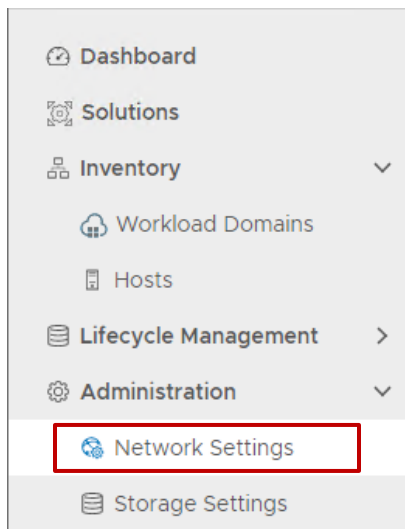
## 4 Workload domain configuration with NFS

When provisioning an NFS-based workload domain, the principal storage assignment for that workload domain must use either NFSv3 or vSAN storage for the initial cluster creation. Once the workload domain has been created, supplemental storage from other arrays or protocols can be added later. This section details how to configure NFSv3 storage to be used for workload domain creation.

### 4.1 VMware Cloud Foundation network pool configuration

Before creating the NFS share, a dedicated storage IP networking pool must be created for NFS traffic.

1. From the SDDC Manager dashboard, in the left pane, expand **Administration** and click **Network Settings**.



2. In the upper-right side of the screen, click **Create Network Pool**.



3. Create the network pool.
  - a. Specify a Network Pool Name.
  - b. For Network Type, select NFS.
  - c. Enter the IP storage network information.
  - d. When finished adding IP address ranges, click **Save**.

The screenshot shows the 'Create Network Pool' configuration page in VMware Cloud Foundation. The 'Network Type' is set to 'NFS'. The 'NFS Network Information' section is highlighted with a red box and contains the following fields:

Field	Value
VLAN ID	10
MTU	1600
Network	10.10.132.0
Subnet Mask	255.255.0.0
Default Gateway	10.10.132.1

Below the network information, the 'Included Ip Address Ranges' section shows a table with one entry:

Start IP	To	End IP
10.10.132.50	To	10.10.132.100

The 'vMotion Network Information' section is also visible, with fields for VLAN ID (21), MTU (1600), Network (192.168.21.0), Subnet Mask (255.255.255.0), and Default Gateway (192.168.21.1). It also has an 'Included Ip Address Ranges' section with a table showing a range from 192.168.21.50 to 192.168.21.100.

4. Make note of the **Included IP Address Ranges** for later use when assigning host access permissions on the NFS share.

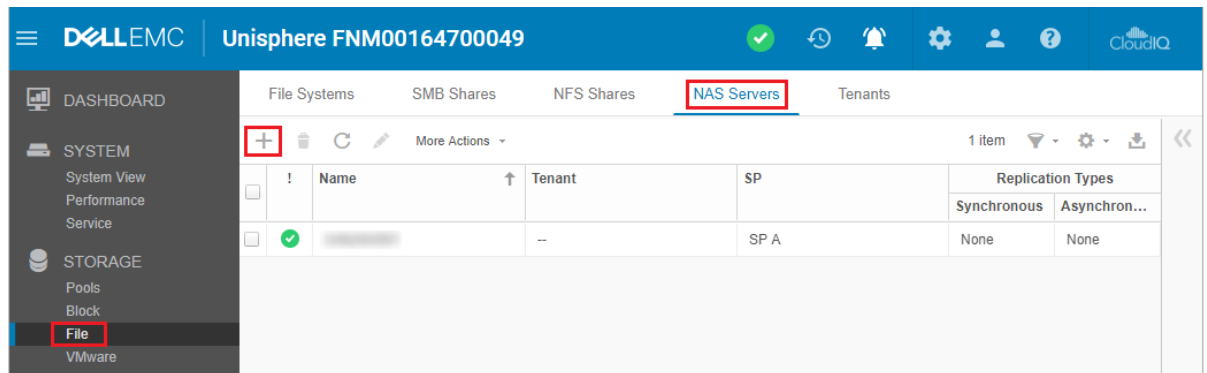
## 4.2 Dell EMC storage configuration

This section details the steps for each of the various arrays to prepare the NFS mount points for VMware Cloud Foundation.

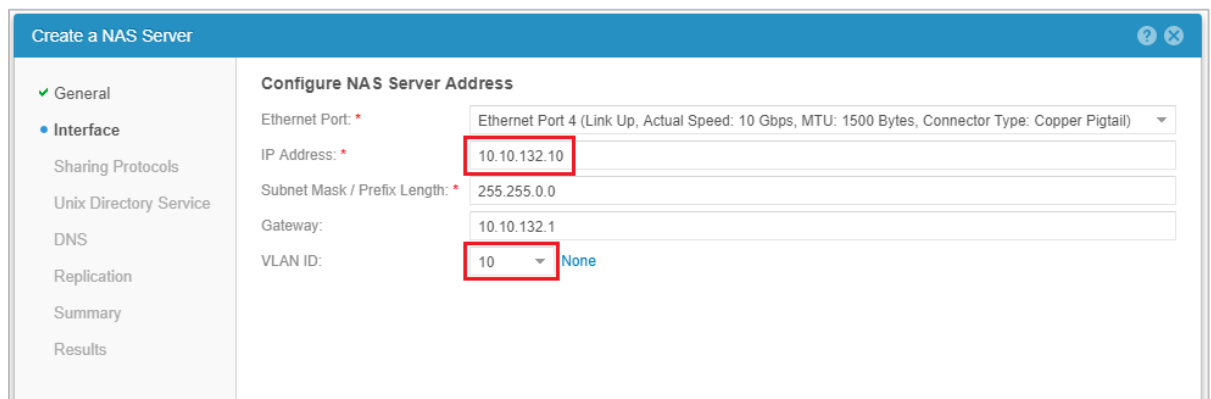
### 4.2.1 Dell EMC Unity NFS share creation

To create the NFSv3 share on the Dell EMC Unity array, follow these steps:

1. From within the Dell EMC Unity Unisphere™ interface, under **Storage** click **File**, click the **NAS Servers** section, and click the **plus (+)** to add a NAS server.



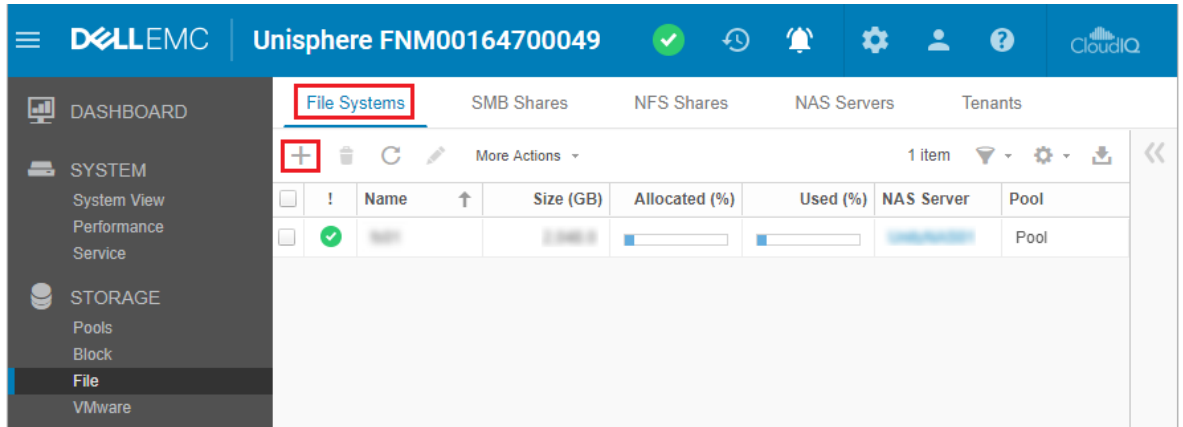
2. Follow the wizard and specify the settings unique to your environment. Any configuration pages that require special attention are covered in the following steps.
3. On the **Interface** page, specify an IP address within the VLAN of the NFS network pool that was created earlier from the SDDC Manager, and click **Next**.



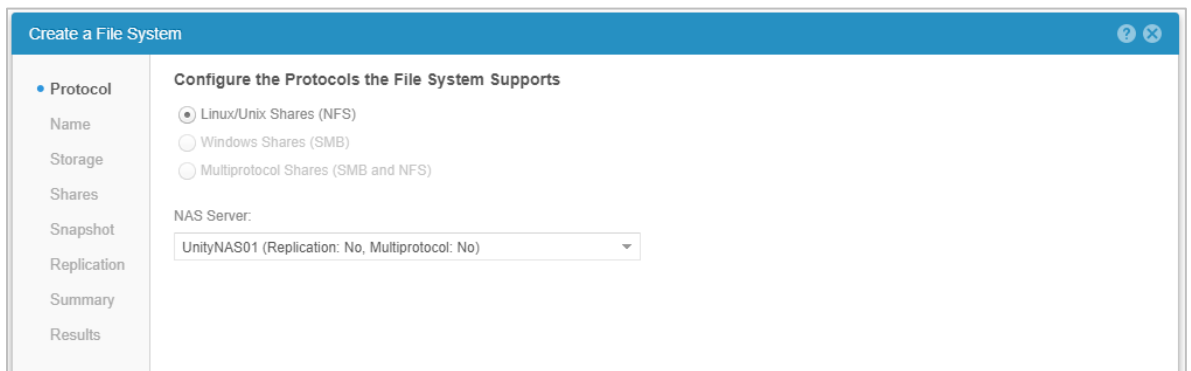
4. On the Sharing Protocols page, select Linux/Unix shares (NFS) and click Next.
5. On the **Summary** page, click **Finish** to create the NAS server.

Next, create the file system for the NFS share.

1. Select the **File Systems** section and click the **plus (+)** to add a file system.

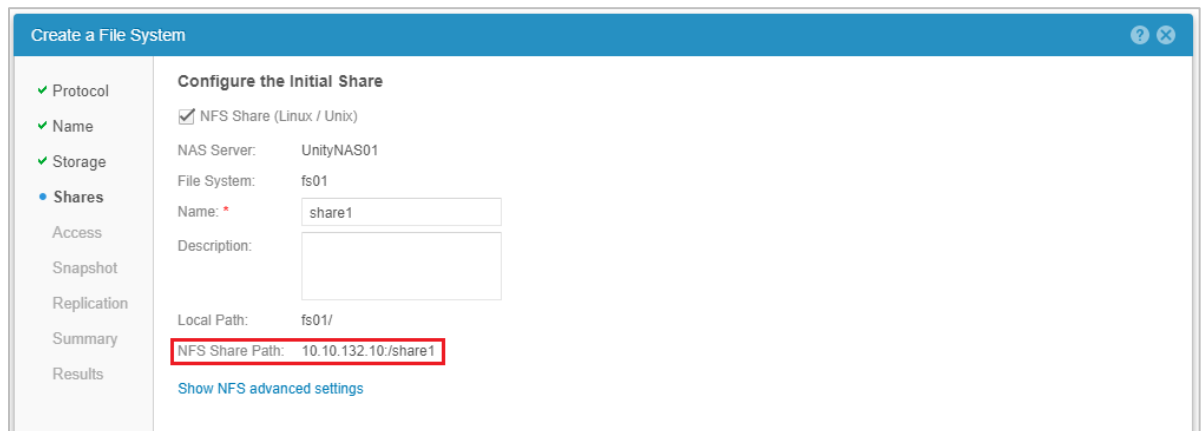


2. On the **Protocol** page, select **Linux/Unix Shares (NFS)**, select the **NAS Server** that was created previously, and click **Next**.

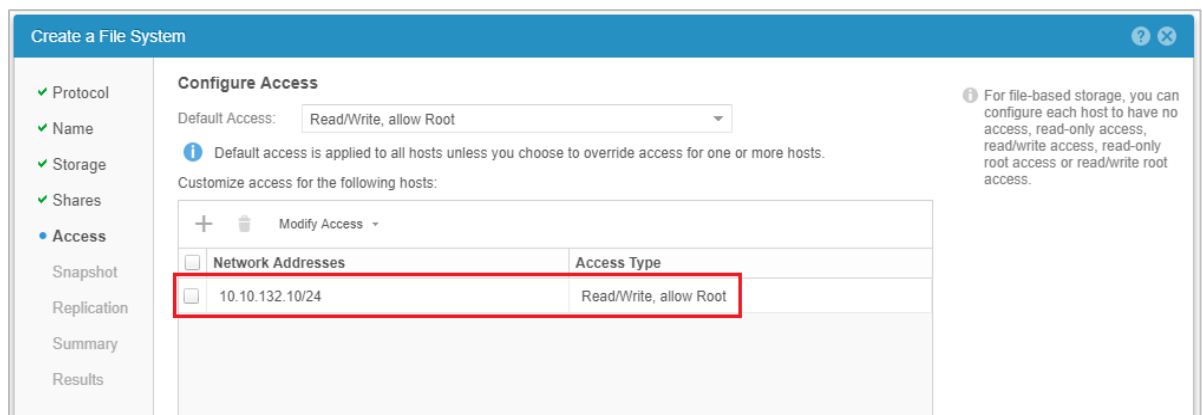


3. Specify the options for the **Name** and click **Next**.
4. On the **Shares** page, check the **NFS Share (Linux/Unix)** box, specify a share name, and click **Next**.

5. Make note of the **NFS Share Path** for use later in configuring the NFS share from within the SDDC Manager.



6. On the **Access** page, set the **Default Access** to **Read/Write, allow root**, and click the **plus sign (+)** to customize the host access. In the section **Customize access for the following hosts**, add the **Included IP Address Ranges** assigned to the NFS network pool created in section 4.1. Click **Next**.



7. Specify the options for the **Snapshot** and **Replication** pages. At the **Summary** page, click **Finish**.

Next, create the NFS share.

---

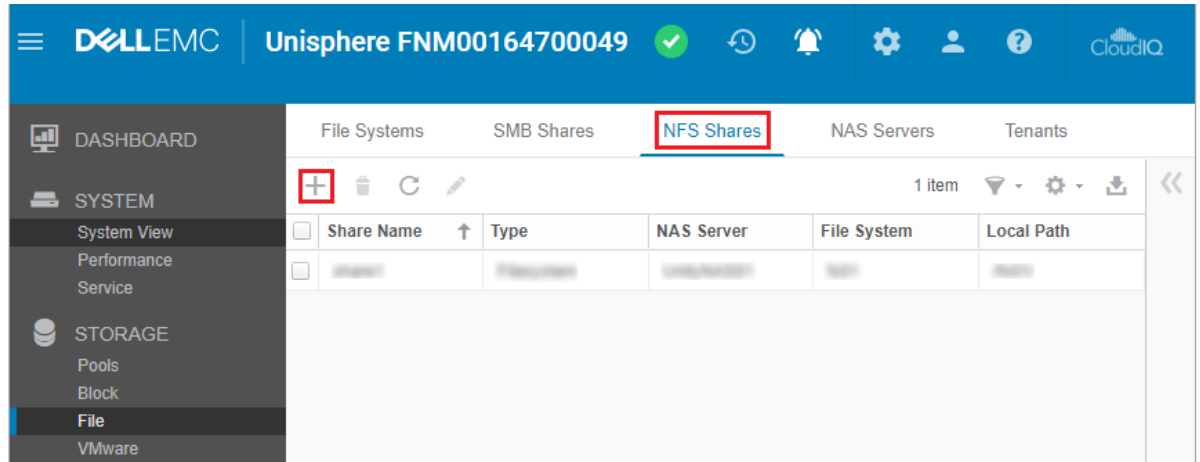
**Note:** If the NFS share was previously created in section 4.2.1, skip this section and perform the steps in section 4.2.2.

---

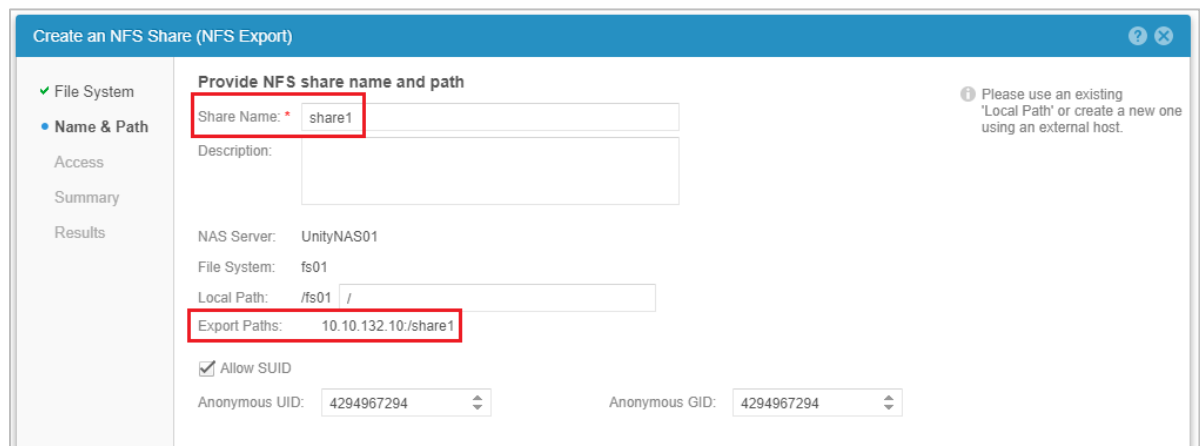


If the NFS share has not yet been created, perform the following steps.

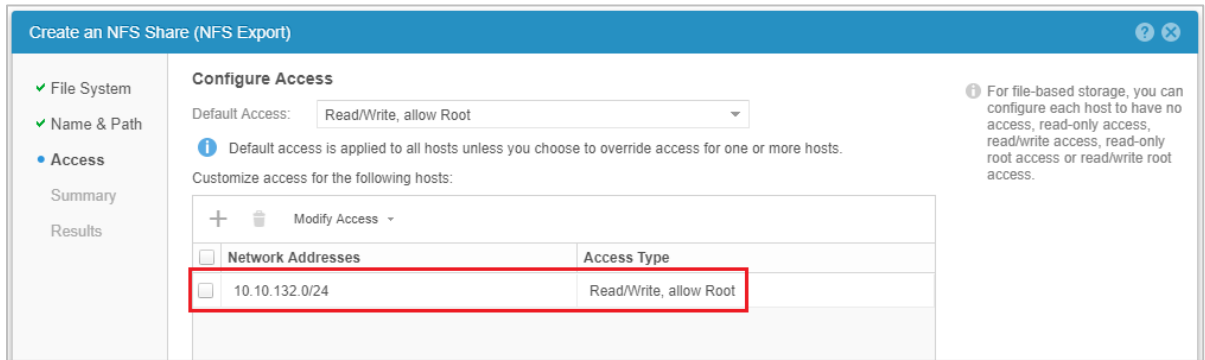
1. Select the **NFS Shares** section and click the **plus (+)** to add an NFS share.



2. On the **File System** page, select the file system that was created in the previous step and click **Next**.
3. On the **Name & Path** page, specify the **Share Name**. Keep note of the **Export Paths** address for later use later when configuring the NFS share from within the SDDC Manager.



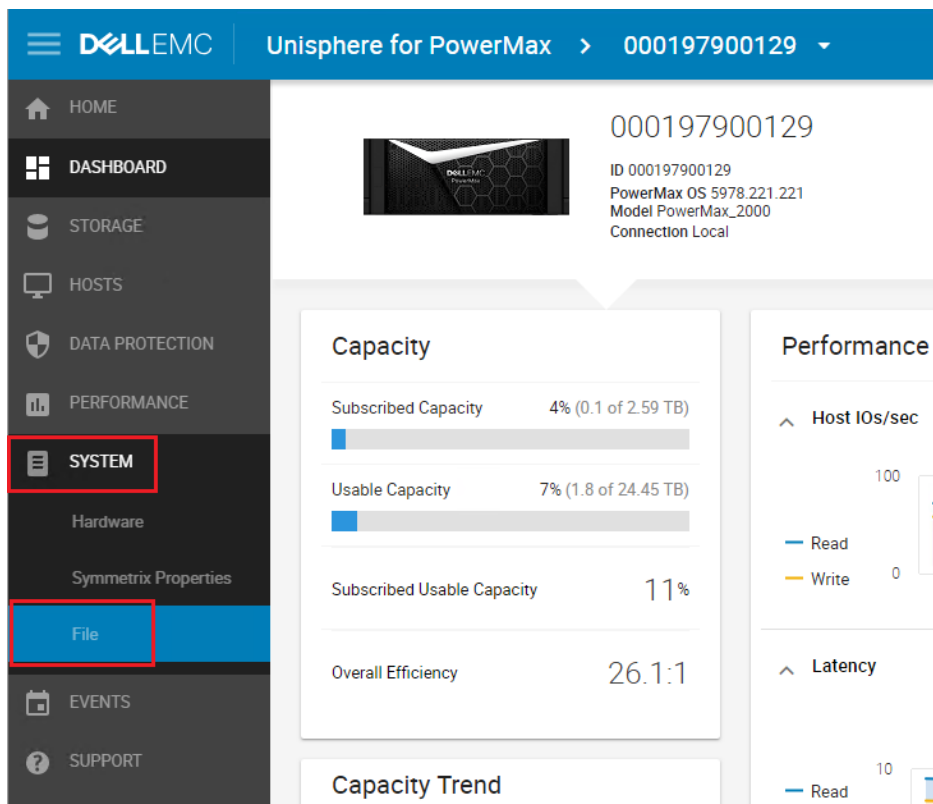
- On the **Access** page, set the **Default Access** to **Read/Write, allow root** and click the **plus (+)** to customize the host access. In the **Customize access for the following** hosts section, add the **Included IP Address Ranges** assigned to the NFS network pool created in section 4.1.



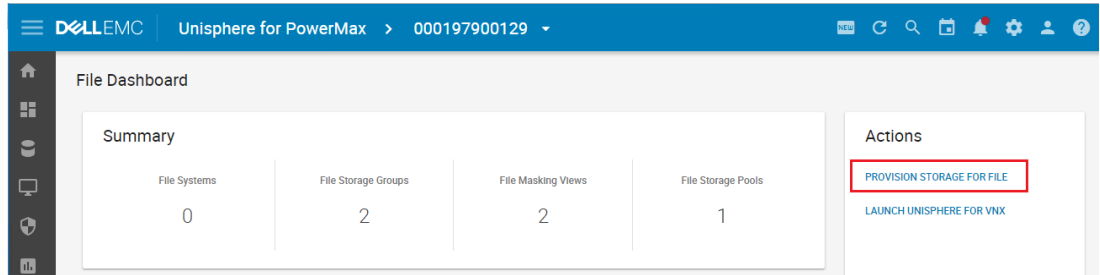
## 4.2.2 PowerMax NFS share creation

To create the NFS share on the PowerMax array, follow these steps:

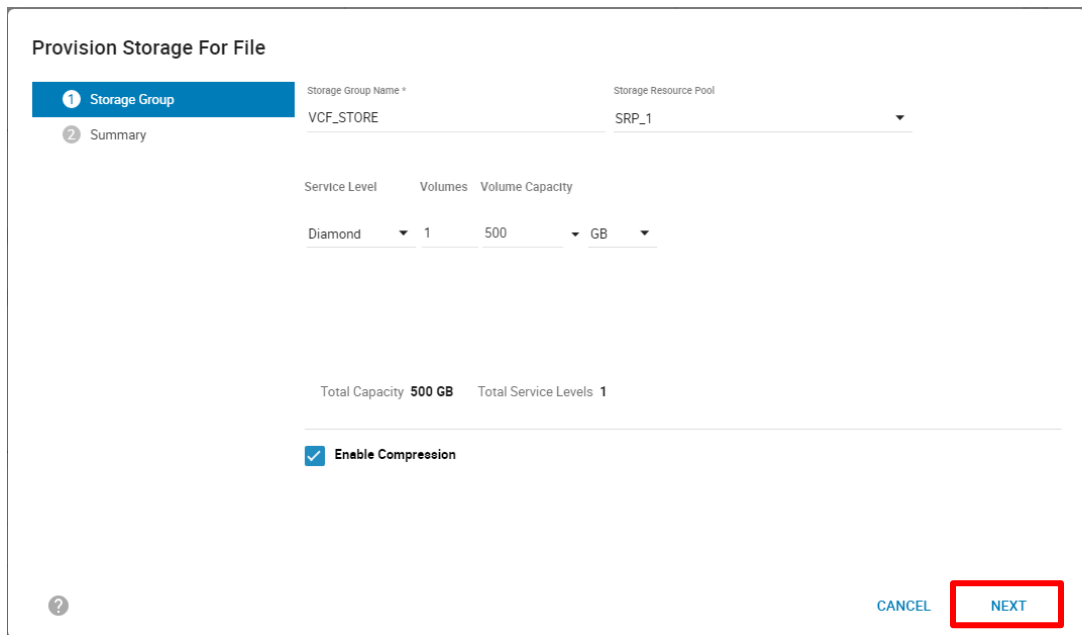
- In the Unisphere interface, click **System > File**.



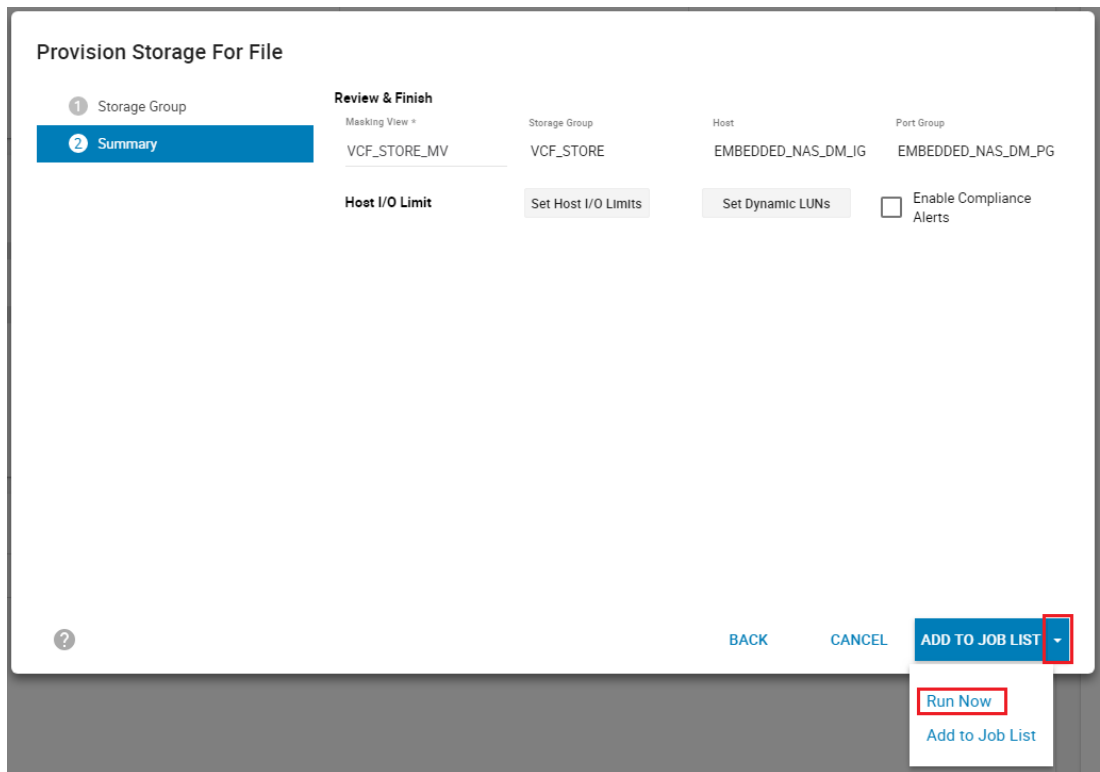
2. In the Actions pane, click Provision Storage for File.



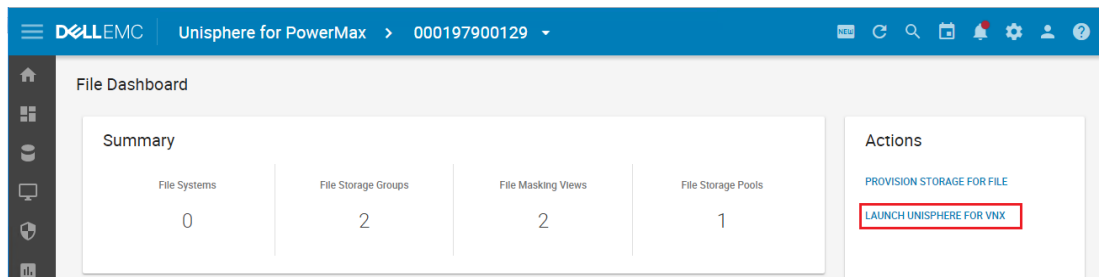
3. In the **Storage Group** section, specify the **Storage Group Name**, select the **Storage Resource Pool** and the **Service Level**. Next, set **Volumes** to **1**, set the **Volume Capacity**, and click **Next**.



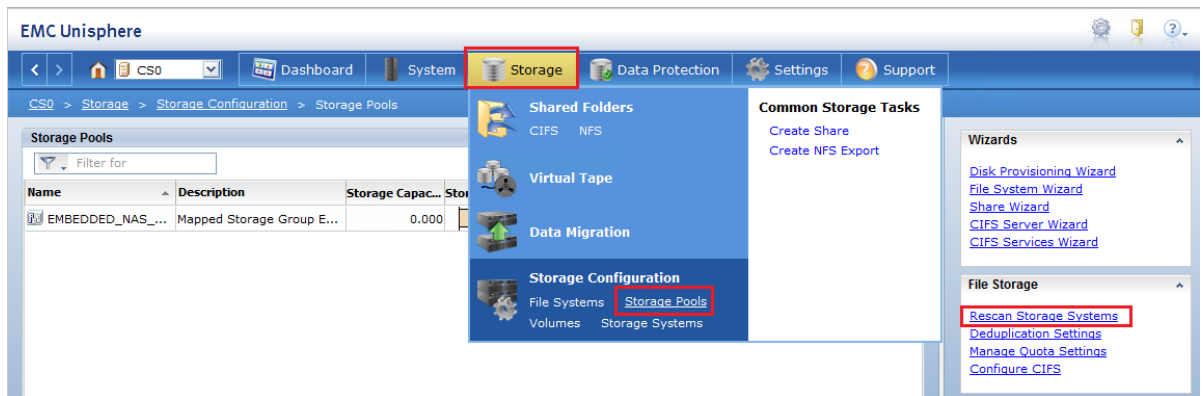
4. On the **Summary** screen, review the settings. Click the **drop-down arrow** next to **Add to Job List** and select **Run Now**.



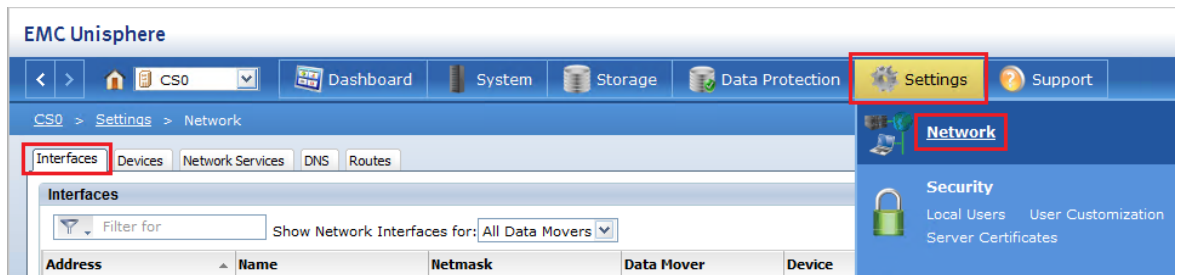
5. Launch Unisphere for VNX.



- When the administrative console opens, click **Storage, Storage Pool**. When the pane opens, click **Rescan Storage Systems**. When the dialog opens, click **OK**.

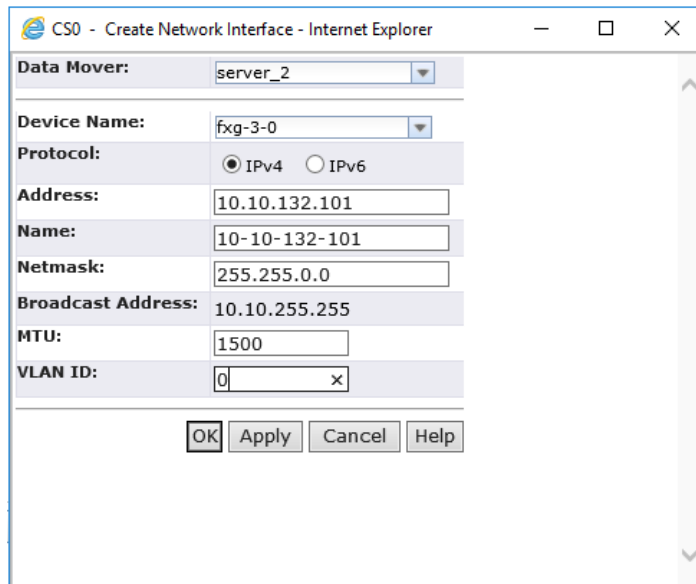


- To create the interface, click **Settings > Network** and click **Interfaces**.

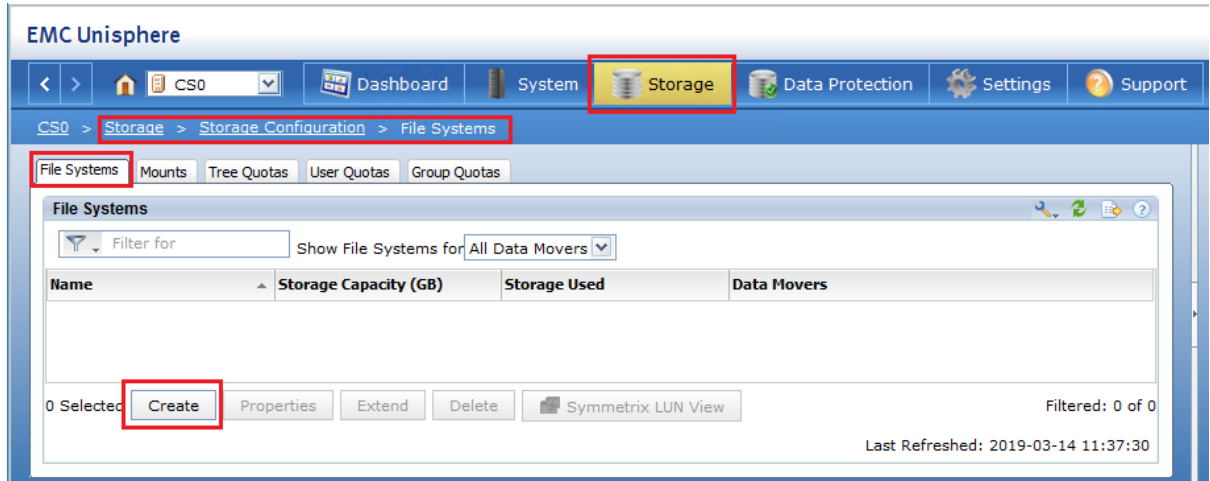


- At the bottom of the pane, click **Create**. Specify the interface settings and click **OK**.

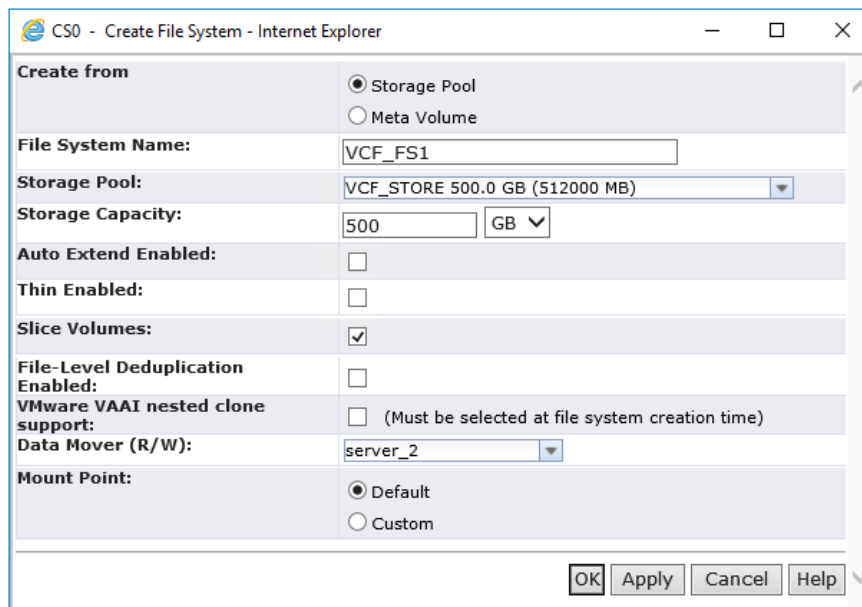
Note this IP address for use in section 4.3.



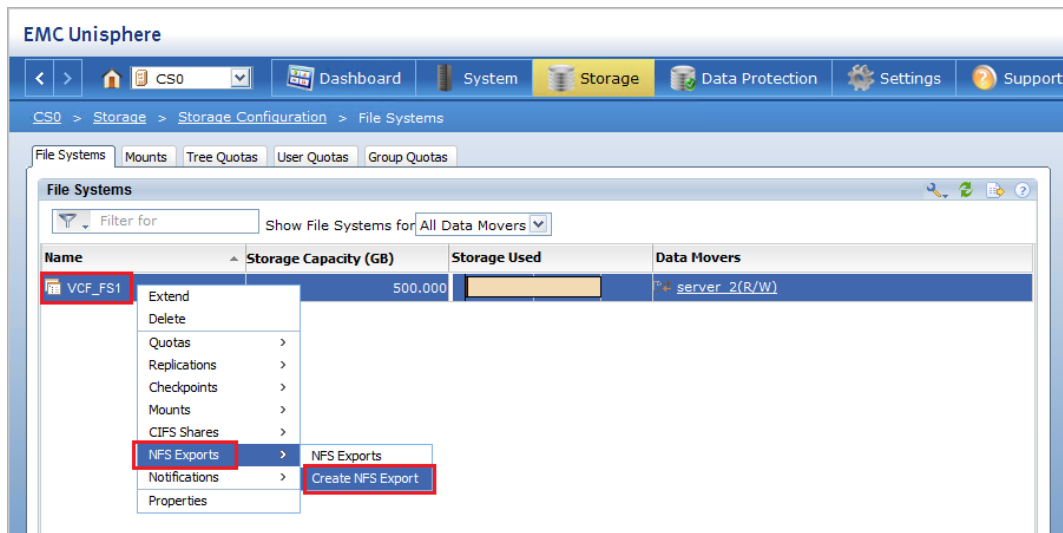
- To create the file system, click **Storage > Storage Configuration > File Systems** and click **Create**.



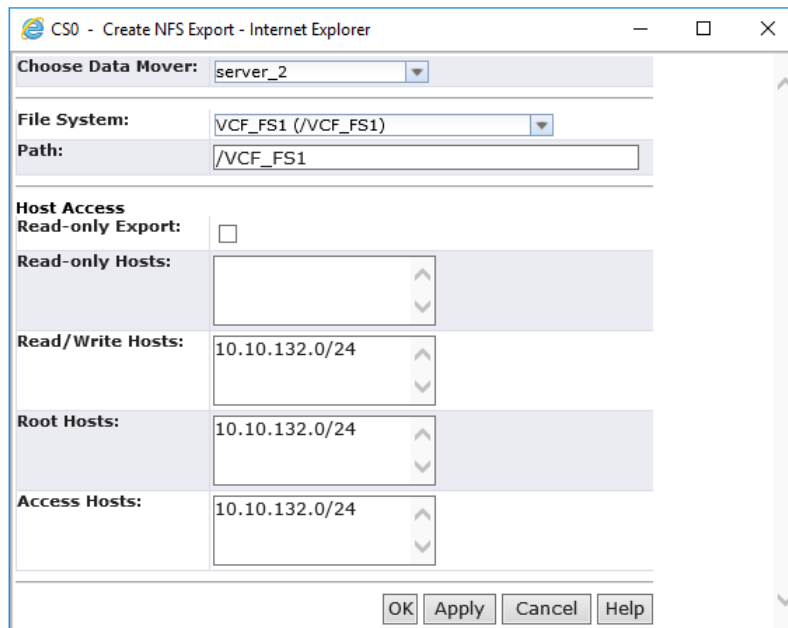
- Specify the options for the file system and click **OK**.



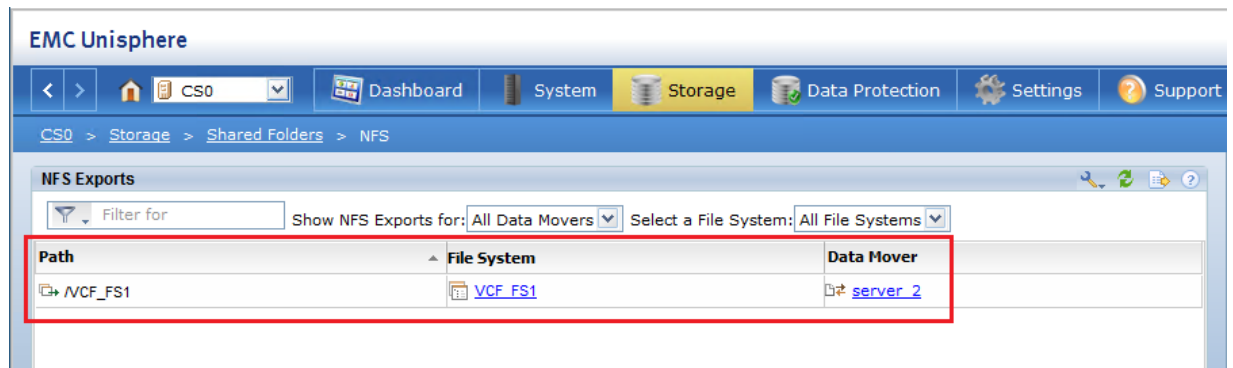
11. To create the NFS export, right-click the file system, click **NFS Exports**, and click **Create NFS Export**.



12. When the **Create NFS Export** configuration screen appears, enter the export and host information and click **OK**.



- Click **Storage** > **Shared Folders** and click **NFS** to reveal the share information needed for section 4.3.

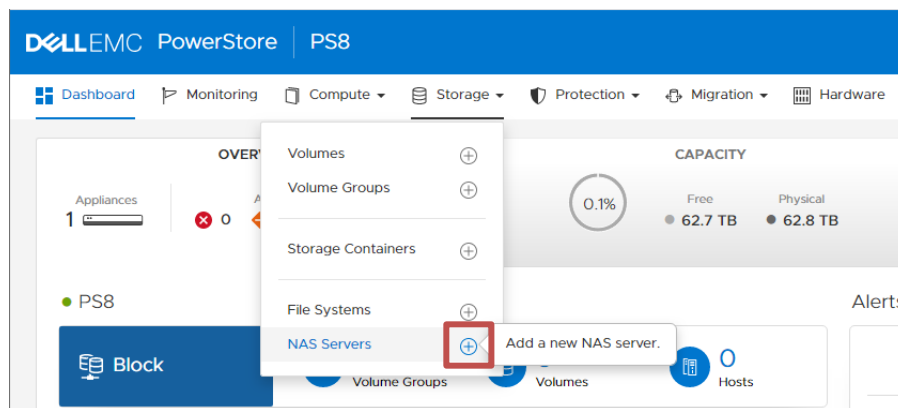


In this example, the NFS share is: **10.10.132.101:VCF\_FS1**

### 4.2.3 PowerStore NFS share creation

To create the NFSv3 share on the PowerStore appliance, follow these steps:

- From within the PowerStore Manager, click **Storage**, and click the **plus sign (+)** next to **NAS Servers**.



- Follow through the wizard to specify the settings unique to your environment. Any configuration pages that require special attention are covered in the following steps.



3. In the **Details** section, specify the name, and IP address for the NAS server. Ensure that the NAS server IP address is within the VLAN of the NFS network pool that was created earlier from the SDDC Manager and click **Next**.

**Create NAS Server** [Close]

Details      Sharing Protocol      Summary

Give your NAS server a friendly name, configure File Interface and Sharing Protocol

**Name**

**Description (Optional)**

**IP Address**

**Subnet Mask/Prefix Length**

**Gateway**

**VLAN ID**

[CANCEL] [NEXT]

4. In the **Sharing Protocol** section, select **NFSv3**, and click **Next**.

**Create NAS Server** [Close]

Details      Sharing Protocol      Summary

Select the type of File Sharing Protocol your NAS server will use. Selecting both SMB and NFS will enable Multiprotocol.

**Select Sharing Protocol**

- Windows Server Settings
- Unix Directory Services
- DNS

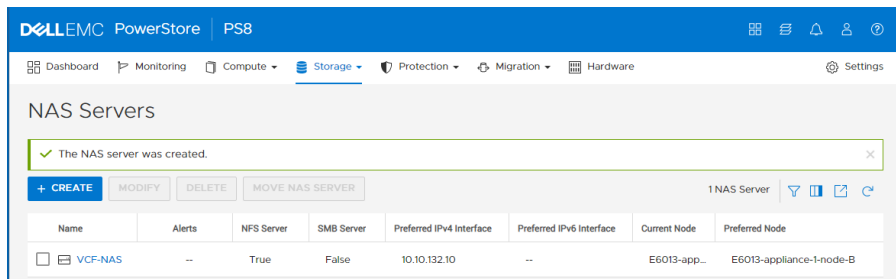
**Protocol Type**

- SMB
- NFSv3
- NFSv4

[CANCEL] [BACK] [NEXT]

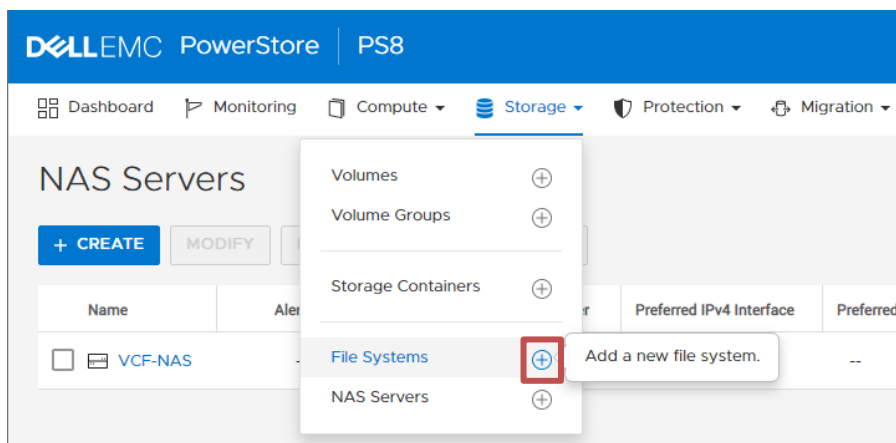
5. (Optional) If using UNIX Directory Services, configure the settings, and click **Next**.

6. (Optional) Configure the NAS Server DNS and click **Next**.
7. At the Summary page, review the settings, and click **CREATE NAS SERVER**.

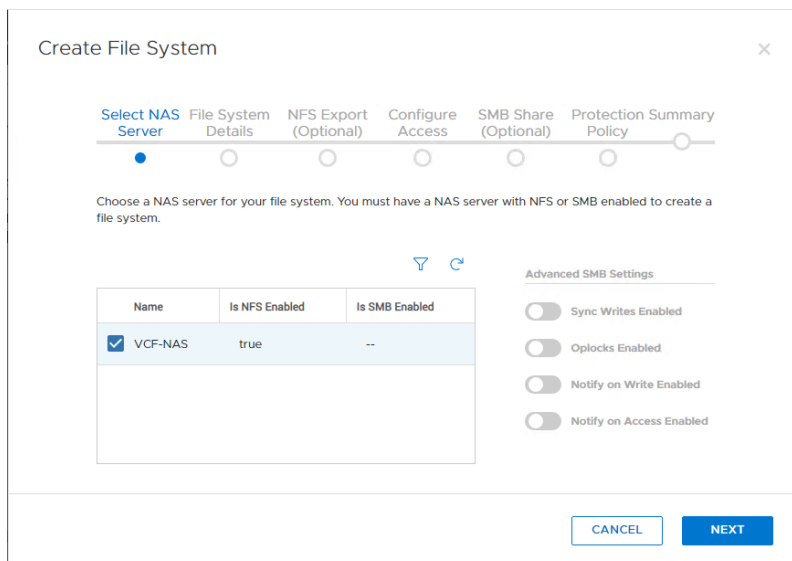


Next, create the file system for the NFS share.

1. To create the file system, click **Storage**, and click the **plus (+)** next to **File Systems**.



2. In the **Select NAS Server** section, select the NAS server that was created previously, and click **Next**.



3. In the **File System Details**, specify a name and size for the file system, and click **Next**.

The screenshot shows the 'Create File System' wizard with the 'File System Details' step selected. The progress bar indicates the following steps: Select NAS Server (completed), File System Details (current), NFS Export (Optional) (pending), Configure Access (pending), Protection Policy (pending), and Summary (pending).

Provide file system name and size

**Name**

**Description** (Optional)

**Size**

4. On the **NFS Export** page, specify the **Name** for the export. Make note of the **NFS Export Path** for use later in configuring NFS options from within the SDDC Manager.

The screenshot shows the 'Create File System' wizard with the 'NFS Export (Optional)' step selected. The progress bar indicates the following steps: Select NAS Server (completed), File System Details (completed), NFS Export (Optional) (current), Configure Access (pending), Protection Policy (pending), and Summary (pending).

Provide NFS Export name

**Name**

**Description** (Optional)

**NFS Export**

NAS Server	VCF-NAS
Local Path	/FS01
File System Name	FS01
NFS Export Path	10.10.132.10/share1

- On the **Configure Access** page, click **ADD HOST** to add the **Included IP Address Ranges** assigned to the NFS network pool created in section 4.1. Next, assign **Read/Write, allow root** permissions.

Minimum Security: Sys

Default Access: No Access

Name/Network Address	Access Type
10.10.132.0/24	Read/Write, allow Root

- On the **Protection Policy** page, select the protection policy for the file system then click **Next**.
- At the **Summary** page, review the settings, and click **CREATE FILE SYSTEM**.

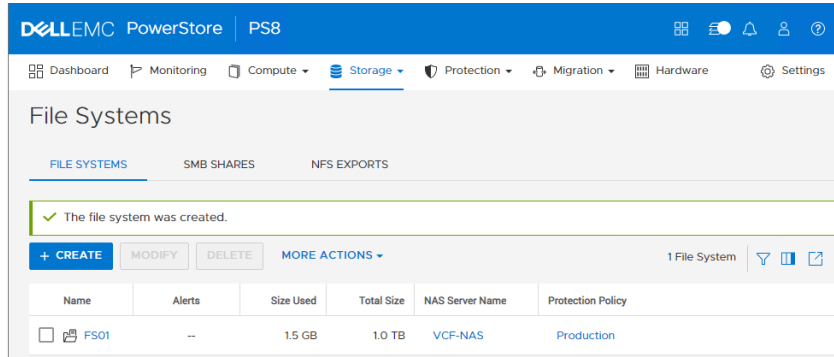
**File System Details**

NAS Server	VCF-NAS
Description	NA
Size	1.0 TB
Protection Policy	Production

**NFS Export**

Name	share1
Description	
Local Path	F501/
NFS Export Path	10.10.132.10/share1
Default Access	No Access
Minimum Security	No Access
Configured Host Access	Yes


- Once the file system is created, go to section 4.3.



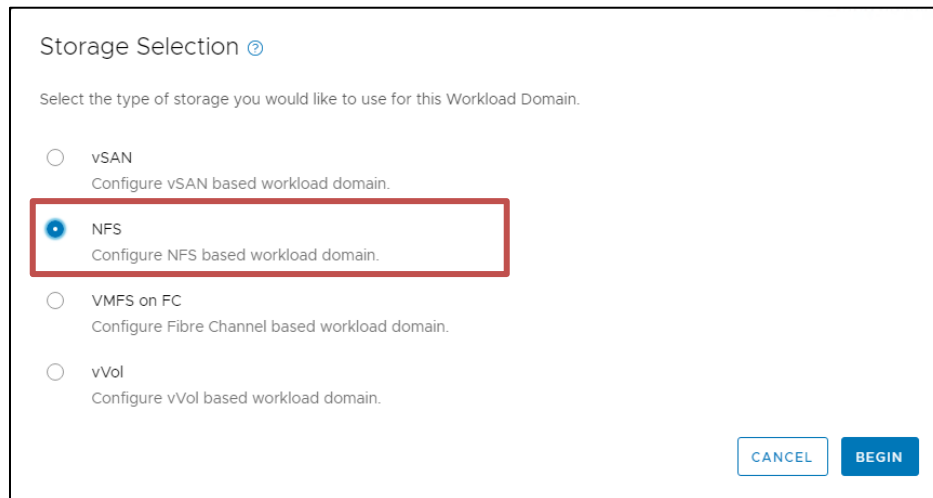
## 4.3 VMware Cloud Foundation workload domain deployment

This section describes how to deploy the NFS workload domain using the NFS share settings noted previously.

- From the SDDC Manager dashboard, click the **+Workload Domain** button and click **VI - Virtual Infrastructure**.



- For **Storage Selection**, select **NFS** and click **Next**.



- When the **VI Configuration** wizard begins, enter the environment-specific details into the **Name**, **Compute**, **Networking**, and **Host Selection** pages.
- When the **NFS Storage** page is reached, enter the NFS share **Export Paths** info previously noted in section 4.2.

The screenshot displays the 'VI Configuration' sidebar on the left, with '6 NFS Storage' highlighted. The main content area is titled 'NFS Storage' and contains a section for 'NFS Share Details' enclosed in a red box. This section includes three input fields:

Field	Value
Datastore Name	msp01-w01-nfs
Folder	/share1
NFS Server IP Address	10.10.132.10

5. On the next screen, specify the **Licenses** and **Object Names**. On the **Review** page, click **Finish** to begin the workload domain deployment.
6. When the workload domain has finished deployment, the NFS share information appears in the SDDC Manager interface.

## 5 Workload domain configuration with FC, iSCSI, or vVol

When presenting external storage to VCF vSphere hosts, remember that VMFS datastores on FC can be principal or supplemental, while iSCSI is considered only supplemental storage for a workload domain. On creation, principal storage is required for initial workload domain deployment. Afterwards, the supplemental datastores can be presented to the workload hosts.

In addition, while principal storage can be monitored from within the SDDC Manager interface, all supplemental datastores must be managed independently through the workload domain VMware vCenter® instance.

---

**Note:** Presenting FC or iSCSI datastores to the management domain is not supported. FC or iSCSI datastores should only be presented to workload domains.

---

**Caution:** When upgrading VMware ESXi™ hosts within the VMware Cloud Foundation workload cluster, hosts requiring custom Fibre Channel VIBs may be overwritten, potentially causing an outage. Consult the VMware KB article [How to upgrade ESXi hosts in VMware Cloud Foundation 3.5 with a vendor-specific ISO image \(65047\)](#) for more information.

---

### 5.1 Using Fibre Channel as principal storage

This section covers the prerequisite steps that are required to use Fibre Channel as principal storage for workload domain hosts.

#### 5.1.1 Commission hosts for VMFS on FC

Before the workload domain can be created using FC as principal storage, the hosts must be commissioned with the **VMFS on FC** storage type.

##### Host Addition and Validation [?](#)

▼ Add Hosts

---

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

Add new     Import

Host FQDN

Storage Type     VSAN     NFS     **VMFS on FC**     vVol

Network Pool Name [i](#)  ▼

User Name

Password  [👁](#)

**ADD**

Figure 6 Selecting the VMFS on FC storage type for host commissioning

## 5.1.2 Create the VMFS datastore

The wizard for workload domain creation requires a preconfigured VMFS datastore that is visible to all the ESXi hosts used in the workload domain.

To accomplish this task, follow these general steps:

1. Verify that all the workload domain ESXi hosts have been zoned to the storage array.
2. Create a datastore volume and present it to all the workload domain hosts. Section 5.3 includes instructions for each Dell EMC storage array.
3. Log in to the web client of one of the ESXi hosts and create the datastore.
  - a. Rescan the storage adapters for new storage devices.
  - b. Create a new VMFS-6 datastore and note the datastore name for workload domain creation later.
4. Log in to the remaining ESXi hosts web client to verify that the hosts can see the newly created datastore.

## 5.1.3 Workload domain creation

When creating the workload domain, the **storage** section of the wizard prompts for the VMFS-6 datastore name that was created previously. If the correct datastore name is not provided, the interface produces an error message like Figure 7.




Figure 7 Error message displayed when SDDC manager cannot detect the datastore on the workload domain hosts.

## 5.2 Using Virtual Volumes (vVols) as principal storage

For Dell EMC storage arrays that support VMware vSphere Virtual Volumes™ (vVols), they can be used as principal or supplemental storage to workload domains. The [VMware HCL](#) for vVols can provide array, protocol, and VASA-provider versions that are required to use vVols with ESXi hosts in a VCF environment. See the respective Dell EMC product administration guides listed in appendix A for specific array configuration instructions.

This section covers the prerequisite steps that are required to use vVols as principal storage for workload domain hosts.

### 5.2.1 Create a Network Pool

Create a network pool that aligns with the vVol storage protocol supported by the array.

- FC vVol: NFS or vMotion only enabled network pool
- iSCSI vVol: iSCSI and vMotion only enabled network pool
- NFS vVol: NFS and vMotion only enabled network pool



## Create Network Pool

Ensure that all required networks are selected based on their usage for workload domains.

Network Pool Name

Network Type  vSAN  NFS  iSCSI  vMotion

### vMotion Network Information

VLAN ID	<input type="text" value="21"/>
MTU	<input type="text" value="1600"/>
Network	<input type="text" value="192.168.21.0"/>
Subnet Mask	<input type="text" value="255.255.255.0"/>
Default Gateway	<input type="text" value="192.168.21.1"/>

**Included Ip Address Ranges**

Once a network pool has been created, you are not able to edit or remove IP ranges from that pool.

To

To

### iSCSI Network Information

VLAN ID	<input type="text" value="10"/>
MTU	<input type="text" value="1600"/>
Network	<input type="text" value="10.10.132.0"/>
Subnet Mask	<input type="text" value="255.255.0.0"/>
Default Gateway	<input type="text" value="10.10.132.1"/>

**Included Ip Address Ranges**

Once a network pool has been created, you are not able to edit or remove IP ranges from that pool.

To

To

Figure 8 Example of a network pool supporting iSCSI vVols

### 5.2.2 Configure ESXi hosts for vVols

When configuring ESXi hosts to be used in a vVol workload domain, there are several prerequisite configuration items that must be completed.

### 5.2.2.1 Complete all host commissioning checklist steps

Ensure that the basic configuration of each ESXi host has been finished according to the current host commissioning checklist. This checklist can be found as the first screen in the host commissioning wizard.

Checklist

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

- Select All**
- Host for vSAN workload domain should be vSAN compliant and certified per the VMware Hardware Compatibility Guide. BIOS, HBA, SSD, HDD, etc. must match the VMware Hardware Compatibility Guide.
- Host has a standard switch with two NIC ports with a minimum 10 Gbps speed.
- Host has the drivers and firmware versions specified in the VMware Compatibility Guide.
- Host has ESXi installed on it. The host must be preinstalled with supported versions (7.0.1-16841703)
- SSH and syslog are enabled on the host.
- Host is configured with DNS server for forward and reverse lookup and FQDN.
- Hostname should be same as the FQDN.
- Management IP is configured to first NIC port.
- Ensure that the host has a standard switch and the default uplinks with 10Gb speed are configured starting with traditional numbering (e.g., vmnic0) and increasing sequentially.
- Host hardware health status is healthy without any errors.
- All disk partitions on HDD / SSD are deleted.
- Ensure required network pool is created and available before host commissioning.
- Ensure hosts to be used for VSAN workload domain are associated with VSAN enabled network pool.
- Ensure hosts to be used for NFS workload domain are associated with NFS enabled network pool.
- Ensure hosts to be used for VMFS on FC workload domain are associated with NFS or VMOTION only enabled network pool.
- Ensure hosts to be used for vVol FC workload domain are associated with NFS or VMOTION only enabled network pool.
- Ensure hosts to be used for vVol NFS workload domain are associated with NFS and VMOTION only enabled network pool.
- Ensure hosts to be used for vVol iSCSI workload domain are associated with iSCSI and VMOTION

CANCEL PROCEED

Figure 9 Example of Host Commissioning checklist in VCF 4.1

### 5.2.2.2 Configure ESXi hosts for FC vVol

In addition to the items on the host commissioning checklist, the following general configuration steps must be completed on the Dell EMC storage array.

1. Zone the Fibre Channel fabric so that each initiator in the ESXi host is visible to the array
2. Create the host objects (hosts, host groups, port groups) on the array
3. If not automatically configured by the array, provision protocol endpoints to each ESXi host

---

**Note:** Before deploying the workload domain, testing to ensure that vVols are working properly with the ESXi hosts is highly recommended.

---

### 5.2.2.3 Configure ESXi hosts for iSCSI vVol

To configure iSCSI vVols, there are multiple configuration steps for both the ESXi host and the storage array.

1. Log directly into the ESXi user interface, and enable the software iSCSI initiator by selecting **Storage > Adapters > Software iSCSI > Enabled** (See Figure 10)
  - a. Copy the full IQN from the **Name & Alias** field.

- i. For example: **iqn.1998-01.com.vmware:myhostname-1a2b3c4d**
  - ii. This IQN will be used later to create the host object on the array.
- b. Click **Add dynamic target** and then add one or more target IPs for the storage array.
  - c. Click **Save configuration**.

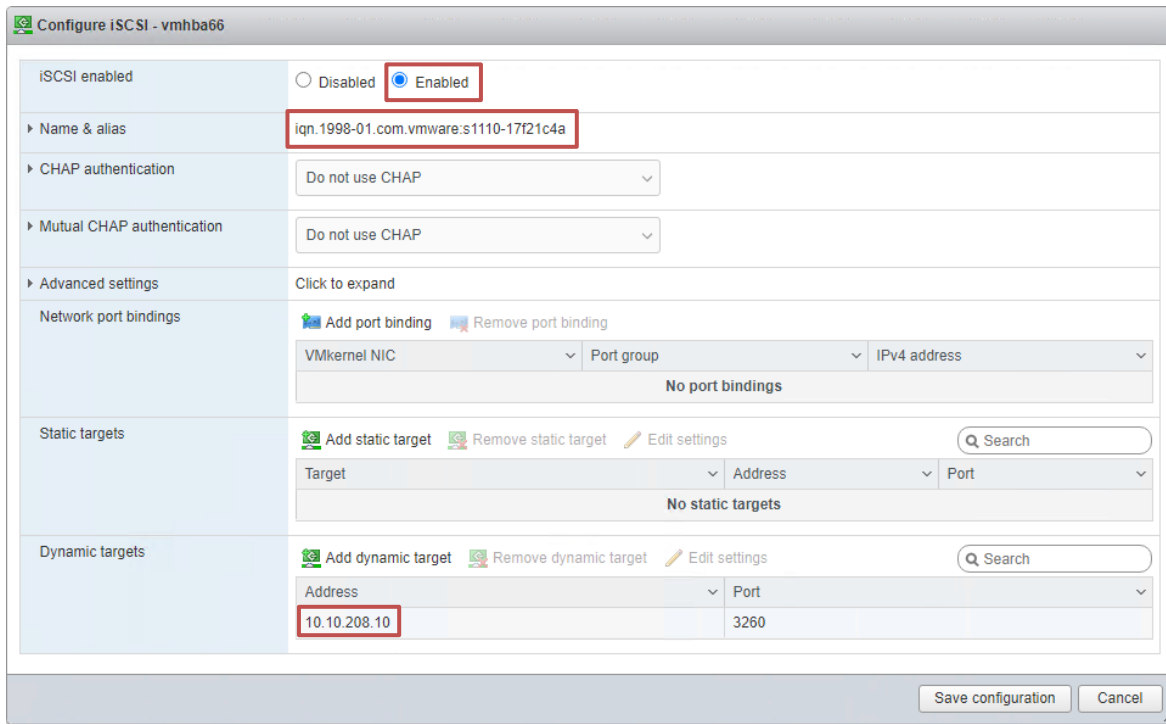


Figure 10 ESXi iSCSI vmhba configuration screen

2. Create the host objects (hosts, host groups, port groups) on the array.
  - a. The IQN address copied previously can be used to manually create the host objects.
3. If not automatically configured by the array, provision protocol endpoints to each ESXi host

**Note:** Before deploying the workload domain, testing to ensure that vVols are working properly with the ESXi hosts is highly recommended. To test iSCSI connectivity, temporarily add a VMkernel NIC with an unused static IP address within the iSCSI network pool range. Use vmkping to ensure connectivity to all the iSCSI targets. When finished testing, the temporary VMkernel IP and associated port group must be removed from the ESXi host before commissioning, or validations will fail.

#### 5.2.2.4 Configure ESXi hosts for NFS vVol

While there are no prerequisite configurations for the ESXi hosts to use NFS vVols, it is recommended to test connectivity between the ESXi host and the NAS server IP address.

**Note:** To test NFS connectivity, temporarily add a VMkernel NIC with an unused static IP address within the NFS network pool range. Use vmkping to test connectivity to the NAS protocol endpoints IP addresses. When finished testing, the temporary VMkernel IP and associated port group must be removed from the ESXi host before commissioning, or validations will fail.

### 5.2.3 Create a Storage Container

The steps to create the vVol storage container differ between Dell EMC arrays, so consult the respective array model's documentation for more detailed instructions. Once the storage container is created, make note of the exact storage container name for future use when creating the workload domain.

---

**Note:** When creating the workload domain with vVols, the storage container name specified in the wizard must exactly match the container name on the array.

---

### 5.2.4 Add Storage Settings

The **Storage Settings** screen allows administrators to add the VASA provider information for one or more arrays. To begin the wizard, click the **ADD VASA PROVIDER** button.

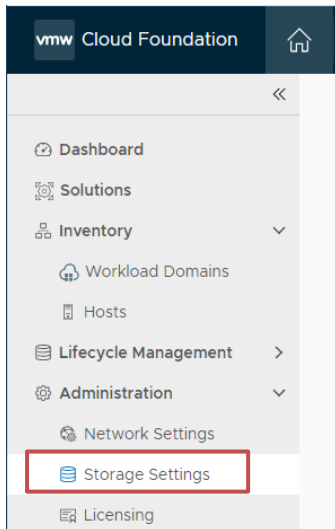


Figure 11 Storage Settings screen location

Enter the array's VASA provider information, making sure to exactly match the storage container name as specified on the array. If needed, multiple container names and credentials can be added to each provider.

## Storage Settings

VASA Providers

### Add VASA Provider

Ensure that all the required details are entered for adding a VASA provider

Name ⓘ	Dell EMC PowerStore	📄
URL ⓘ	<a href="https://100.88.130.80:8443/version.xml">https://100.88.130.80:8443/version.xml</a>	
User Credentials ⓘ		
User Name	vmadmin	
Password	.....	🔒 🔍
<a href="#">ADD ADDITIONAL USER CREDENTIALS</a>		
Storage Containers ⓘ		
Container Name	VCF	
Container Type	ISCSI ▾	
<a href="#">REMOVE</a>		
<a href="#">ADD ADDITIONAL STORAGE CONTAINERS</a>		
<input type="button" value="SAVE"/> <input type="button" value="CANCEL"/>		

Figure 12 Example of entering VASA Provider information

Dell EMC best practice is to create a dedicated security account on the array for the VASA provider to use.

---

**Note:** The VASA provider IP address must reside on the same management network as the vCenter server for the workload domain. vCenter must have network connectivity to the VASA provider in order to issue vVol instructions to the array.

---

## 5.2.5 Commission hosts for vVol

Before the workload domain can be created using vVols as principal storage, the hosts must be commissioned with the **vVol** storage type.

Host Addition and Validation [?](#)

▼ Add Hosts

---

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

Add new     Import

Host FQDN

Storage Type     VSAN     NFS     VMFS on FC     **vVol**

vVol Storage Protocol Type  ▼

Network Pool Name  ▼

User Name

Password  [?](#)

**ADD**

Selecting the vVol storage type for host commissioning

## 5.2.6 Workload domain creation

When creating the workload domain, the **vVol Storage** section of the wizard prompts for the VASA provider information that was created previously.

vVol Storage [?](#)

### vVol Storage Details

Select VASA Protocol Type <a href="#">?</a>	<input type="text" value="ISCSI"/> ▼
Select VASA Provider <a href="#">?</a>	<input type="text" value="Dell EMC PowerStore"/> ▼
Select Storage Container <a href="#">?</a>	<input type="text" value="VCF"/> ▼
Select VASA User <a href="#">?</a>	<input type="text" value="vmadmin"/> ▼
Datastore Name <a href="#">?</a>	<input type="text" value="msp-w01-vvol"/>

Figure 13 vVol Storage screen example

## 5.3 Dell EMC storage configuration

For each Dell EMC storage product, the steps to provision storage to VMware Cloud Foundation vary from model to model. This section details the high-level steps for each.

### 5.3.1 Dell EMC Unity

When presenting storage from a Dell EMC Unity array, the array can present multiple protocols for the principal workload domain storage, and iSCSI datastores for supplemental storage.

Detailed steps for provisioning datastores to vSphere hosts are not covered within this document. For more information, refer to the document [Dell EMC Unity Storage with VMware vSphere](#).

The high-level steps to provision storage to the workload cluster are as follows:

1. If using VMFS on FC, perform the following:
  - a. Create Fibre Channel zones for the workload domain ESXi hosts.
  - b. Perform the following in SDDC Manager:
    - i. Commission the workload domain hosts.
    - ii. Create a workload domain from the commissioned hosts.
  - > If using a Dell EMC Unity NFS share for the storage selection, see the steps in section 4.
2. If using VMFS on iSCSI, perform the following:
  - a. Add an iSCSI Distributed Port Group and assign it the iSCSI VLAN.
  - b. Add a VMkernel adapter to each workload domain host with an IP in the storage network.
  - c. Add a software iSCSI adapter to each workload domain host.
  - d. Add the iSCSI targets for the array to the iSCSI adapter and then rescan.
3. Perform the following in the Dell EMC Unity Unisphere interface:
  - a. Create the hosts for the workload domain ESXi® hosts.
  - b. Create a block LUN specifying the name, capacity, and other options needed for the datastore. Configure the host access to specify the workload domain vSphere hosts.
  - c. Optionally, use the **Create Datastore Wizard** from the VMware storage section to automatically provision a datastore to the vSphere hosts, which bypasses the following manual datastore-creation steps.
4. Create the datastore using the steps outlined in section 5.4.

### 5.3.2 PowerMax

When presenting storage from a Dell EMC PowerMax array, the array can present multiple protocols for the principal workload domain storage, and iSCSI datastores for supplemental storage. When presenting storage from PowerMax arrays to a workload domain, the process is completed in the same manner as presenting storage to conventional vSphere clusters.

Detailed steps for provisioning datastores to vSphere hosts are not covered within this document. For more information, refer to the document [Using Dell EMC VMAX and PowerMax in VMware vSphere Environments](#).

The high-level steps to provision storage to the workload cluster are as follows:

1. If using VMFS on FC, perform the following:
  - a. Create Fibre Channel zones for the workload domain VMware ESXi hosts.
  - b. Perform the following in SDDC Manager:

- i. Commission the workload domain hosts.
  - ii. Create a workload domain from the commissioned hosts.
2. If using VMFS on iSCSI, perform the following:
  - a. Add an iSCSI distributed port group and assign it the iSCSI VLAN.
  - b. Add a VMkernel adapter to each workload domain host with an IP in the storage network.
  - c. Add a software iSCSI adapter to each workload domain host.
  - d. Add the iSCSI targets for the array to the iSCSI adapter and then rescan.
3. Perform the following in the Unisphere for PowerMax interface:
  - a. Create the port group for the PowerMax target ports zoned to the ESXi hosts.
  - b. Create the hosts and hosts group for initiators belonging to the workload domain ESXi hosts.
  - c. Create a storage group specifying the name and volume capacity for the datastore.
  - d. Create the masking view to present the new volume to the workload domain vSphere cluster.
4. Create the datastore using the steps outlined in section 5.4.

### 5.3.3 PowerStore

When presenting storage from a PowerStore T model appliance, it can present multiple protocols for the principal workload domain storage, and iSCSI datastores for supplemental storage. The PowerStore X model appliance can present storage externally to workload domain hosts as a VMFS on FC datastore for principal workload domain storage, and iSCSI datastores for supplemental storage. PowerStore X model ESXi hosts cannot be used to create or be used as part of a VCF workload domain cluster.

Detailed steps for provisioning datastores to vSphere hosts are not covered within this document. For more information, see the document [Dell EMC PowerStore Series Virtualization Guide](#).

The high-level steps to provision storage to the workload cluster are as follows:

1. If using VMFS on FC, perform the following:
  - a. Create Fibre Channel zones for the workload domain ESXi hosts.
  - b. Perform the following in SDDC Manager:
    - i. Commission the workload domain hosts.
    - ii. Create a workload domain from the commissioned hosts.
  - > The workload domain storage selection must use:
    - PowerStore T model – VMFS on FC, NFSv3, vSAN, or vVols
    - PowerStore X model – VMFS on FC, or vSAN
  - > If using a PowerStore NFSv3 share, see the steps in section 4.
2. If using VMFS on iSCSI, perform the following:
  - a. Add an iSCSI Distributed Port Group and assign it the iSCSI VLAN.
  - b. Add a VMkernel adapter to each workload domain host with an IP in the storage network.
  - c. Add a software iSCSI adapter to each workload domain host.
  - d. Add the iSCSI targets or Global Storage Discovery IP for the appliance to the iSCSI adapter and then rescan.



3. Perform the following in the PowerStore Manager:
  - a. Create the hosts and host group for the workload domain ESXi hosts.
  - b. Create a volume specifying the name, size, and other options for the datastore. Configure the host mappings to specify the workload domain vSphere host group.
  - c. Optionally, use the Virtual Storage Integrator (VSI) vSphere plug-in to automatically provision a datastore to the vSphere hosts, which bypasses the manual datastore-creation steps in section 5.4.
4. Create the datastore using the steps outlined in section 5.4.

### 5.3.4 SC Series

When presenting storage from Dell EMC SC Series arrays to a workload domain, the process is completed in the same manner as presenting storage to conventional vSphere clusters. Since FC datastores are added and managed outside of VMware Cloud Foundation, normal SC Series administrative tools can be used.

Detailed steps for provisioning datastores to vSphere hosts are not covered within this document. For more information, see the current *Dell Storage Manager Administrator's Guide* on [Dell.com/support](http://Dell.com/support).

The high-level steps to provision storage to the workload cluster are as follows:

1. If using VMFS on FC, perform the following:
  - a. Create Fibre Channel zones for the workload domain ESXi hosts.
  - b. Perform the following in SDDC Manager:
    - i. Commission the workload domain hosts.
    - ii. Create a workload domain from the commissioned hosts.
2. If using VMFS on iSCSI, perform the following:
  - a. Add an iSCSI Distributed Port Group and assign it the iSCSI VLAN.
  - b. Add a VMkernel adapter to each workload domain host with an IP in the storage network.
  - c. Add a software iSCSI adapter to each workload domain host.
  - d. Add the iSCSI targets for each controller to the iSCSI adapter and then rescan.
3. Perform the following in Dell Storage Manager or Unisphere for SC Series interface:
  - a. Create the Server objects and Server Cluster for the workload domain ESXi hosts.
  - b. Create a volume specifying the name, size, and other options for the datastore.
    - i. For the **Server** (Dell Storage Manager Client) or **Server Mapping** (Unisphere for SC Series web interface), select the server cluster previously created.
  - c. Configure the server mappings to specify the workload domain vSphere cluster.
4. Create the datastore using the steps outlined in section 5.4.

### 5.3.5 PowerFlex System

When presenting storage from a PowerFlex [two-layer system](#) (where storage and compute nodes remain separate operationally) to workload domains for supplemental storage, the process is the same as presenting

storage to conventional vSphere clusters. Since the datastores are managed outside of VMware Cloud Foundation, PowerFlex storage administrative tools should only be used for storage-specific operations.

The high-level steps to provision storage to the workload cluster are as follows:

1. Prerequisites
  - a. Set up and deploy the VCF management domain per VMware best practices
  - b. Set up and deploy the VCF workload domain per VMware best practices
  - c. Set up and deploy the PowerFlex two-layer configuration using PowerFlex Manager (VxFlex Manager) per PowerFlex Systems best practices
2. Installation
  - a. Install and configure the PowerFlex SDC driver on the VCF workload domain compute nodes.
  - b. Configure PowerFlex-specific network options on the VCF workload domain compute nodes through vCenter.
3. Post installation
  - a. Create PowerFlex LUNs through PowerFlex UI.
  - b. Expose PowerFlex LUNs to workload domain nodes.
  - c. Map PowerFlex LUNs to VCF workload domain nodes through vCenter.
4. Create the datastore using the steps outlined in section 5.4.

For more detailed information about deployment of PowerFlex system as supplemental storage with VCF, see the [Implementation guide for VMware Cloud Foundation with PowerFlex \(formerly VxFlex\) Systems](#).

### 5.3.6 XtremIO

When presenting storage from XtremIO storage controllers to a workload domain, the process is completed in the same manner as presenting storage to conventional vSphere clusters. Since the datastores are managed outside of VMware Cloud Foundation, XtremIO administrative tools should only be used for storage-specific operations.

Detailed steps for provisioning datastores to vSphere hosts are not covered within this document. For more information, refer to the latest *XtremIO Storage Array User Guide* on [Dell.com/support](http://Dell.com/support).

The high-level steps to provision storage to the workload cluster are as follows:

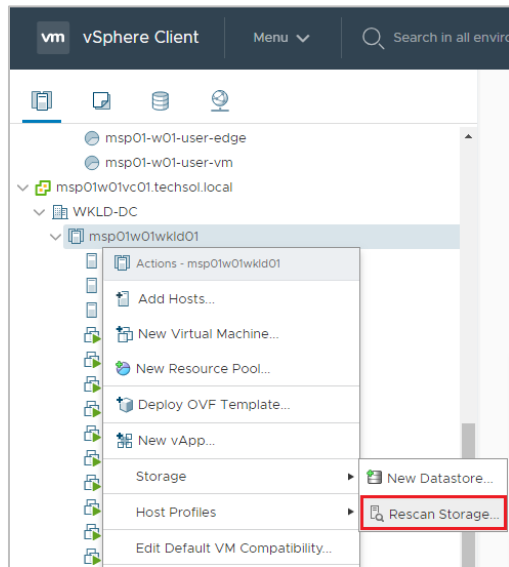
1. Perform the following in SDDC Manager:
  - a. Commission the workload domain hosts.
  - b. Create a workload domain from the commissioned hosts. The workload domain storage selection must use either VMFS on FC for Fibre Channel arrays, or an alternative principal storage option such as vSAN for iSCSI arrays.
2. Perform the following in the XtremIO Storage Management interface:
  - a. Create the initiator group for initiators belonging to the workload domain ESXi hosts.
  - b. Create a volume specifying the name, size, and other options for the datastore.
  - c. Map the volume to the initiator group for the workload domain ESXi hosts.

3. Create the datastore using the steps outlined in section 5.4.

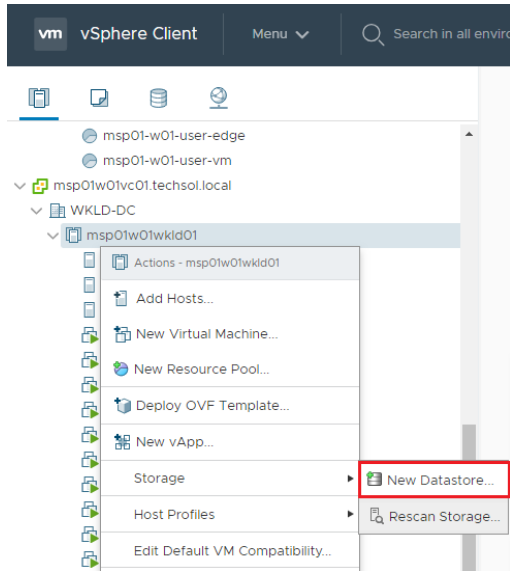
## 5.4 Datastore creation

Once the storage device is presented to the ESXi cluster, switch to the workload domain's VMware vCenter client to finish provisioning the datastore.

1. To rescan for new storage, right-click the workload vSphere Cluster, click **Storage**, and click **Rescan Storage**.



2. Right-click the vSphere Cluster, click **Storage**, and click **New Datastore**.



3. Complete the New Datastore wizard to create the datastore for the workload cluster.

## A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

### A.1 Related resources

- VMware Cloud Foundation resources
  - [VMware Cloud Foundation documentation](#)
  - [Additional VMware Cloud Foundation documents](#)
- Dell EMC Unity resources
  - [Dell EMC Unity Storage with VMware vSphere](#)
  - [Dell EMC Unity: Best Practices Guide - Best Practices for Performance and Availability](#)
  - [Dell EMC Unity: Virtualization Integration](#)
- PowerMax resources
  - [PowerMax Product Guide](#)
  - [Using Dell EMC VMAX and PowerMax in VMware vSphere Environments](#)
- PowerStore resources
  - [PowerStore Product Documentation](#)
  - [PowerStore Info Hub](#)
- PowerFlex resources
  - [PowerFlex Product Documentation](#)
- XtremIO resources
  - [Dell EMC XtremIO Storage Array Host Configuration Guide](#)