



FlexPod Express

FlexPod

NetApp

October 30, 2025

Table of Contents

FlexPod Express	1
FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Design Guide	1
NVA-1139-DESIGN: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series	1
Program summary	1
Technology requirements	4
Design choices	5
Conclusion	11
Where to find additional information	11
FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Deployment Guide	12
NVA-1142-DEPLOY: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series - NVA	
Deployment	12
Solution overview	12
Technology requirements	15
FlexPod Express cabling information	16
Deployment procedures	19
Conclusion	105
Acknowledgments	105
Where to find additional information	106
Version history	106
FlexPod Express with Cisco UCS C-Series and AFF A220 Series Design Guide	106
NVA-1125-DESIGN: FlexPod Express with Cisco UCS C-Series and AFF A220 Series	106
Program summary	107
Solution overview	108
Technology requirements	109
Design choices	110
Solution verification	115
Conclusion	116
Where to find additional information	116
FlexPod Express with Cisco UCS C-Series and AFF A220 Series Deployment Guide	116
NVA-1123-DEPLOY: FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 deployment	
guide	116
Solution overview	117
Technology requirements	120
FlexPod Express cabling information	121
Deployment procedures	123
Conclusion	196
Where to find additional information	196
FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based	
Storage	196
NVA-1131-DEPLOY: FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with	
Direct-Attached IP-Based Storage	196
Solution Overview	197
Technology requirements	199

FlexPod Express Cabling Information	200
Deployment Procedures	202
Conclusion	304
Additional Information	304
FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS - NVA - Deployment	305

FlexPod Express

FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Design Guide

NVA-1139-DESIGN: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series

Savita Kumari, NetApp

In partnership with:



Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices that uses the technology that they are familiar with in their data center.

FlexPod Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp AFF systems. The components of FlexPod Express are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

[Next: Program summary.](#)

Program summary

FlexPod Converged Infrastructure Portfolio

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or as NetApp Verified Architectures (NVAs). Deviations that are based on customer requirements from a given CVD or NVA are permitted if those variations do not result in the deployment of unsupported configurations.

As illustrated in the following figure, the FlexPod portfolio includes the following solutions: FlexPod Express and FlexPod Datacenter.

- **FlexPod Express** is an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter** delivers an optimal multipurpose foundation for various workloads and applications.

Expanded portfolio of platforms

FlexPod® Express	FlexPod Datacenter
<p>Departmental deployments and VAR velocity</p> <p>Target: Primarily MSB, remote, and departmental deployments</p> 	<p>Massively scalable, mission-critical workloads</p> <p>Target: Enterprise/service provider</p> 
<p>Entry level: Cisco UCS, Cisco Nexus, and NetApp AFF and FAS systems</p>	<p>Cisco UCS, Cisco Nexus, and NetApp AFF and FAS systems</p>
<p>Distinct Architectures</p>	<p>Distinct Architectures</p>

NetApp Verified Architecture program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. An NVA solution has the following qualities:

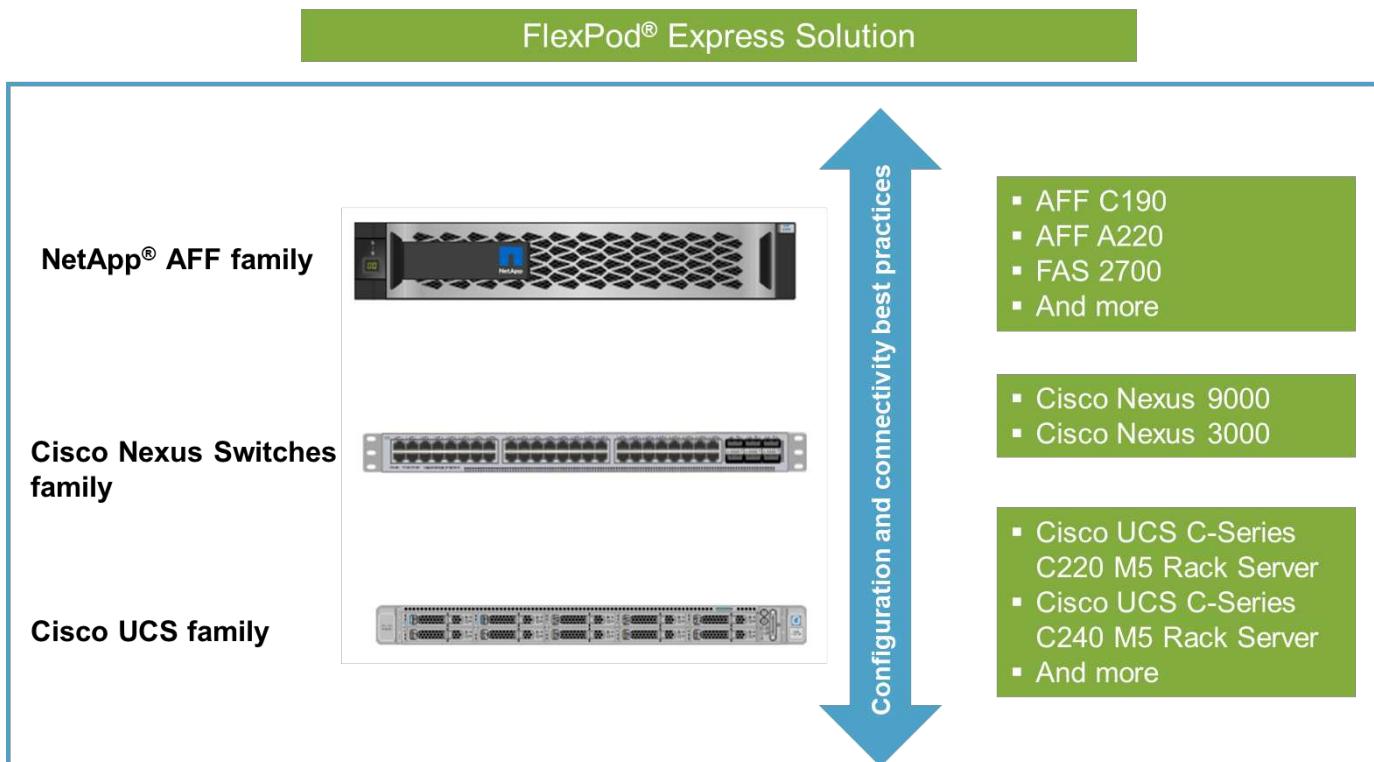
- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

In addition, this design leverages the all-new AFF C190 system, which runs NetApp ONTAP 9.6 software, Cisco Nexus 31108 switches, and Cisco UCS C220 M5 servers as hypervisor nodes.

Solution overview

FlexPod Express is designed to run mixed virtualization workloads. It is targeted for remote and branch offices and for small to midsize businesses. It is also optimal for larger businesses that want to implement a dedicated solution for a specific purpose. This new solution for FlexPod Express adds new technologies such as NetApp ONTAP 9.6, NetApp AFF C190 system, and VMware vSphere 6.7U2.

The following figure shows the hardware components that are included in the FlexPod Express solution.

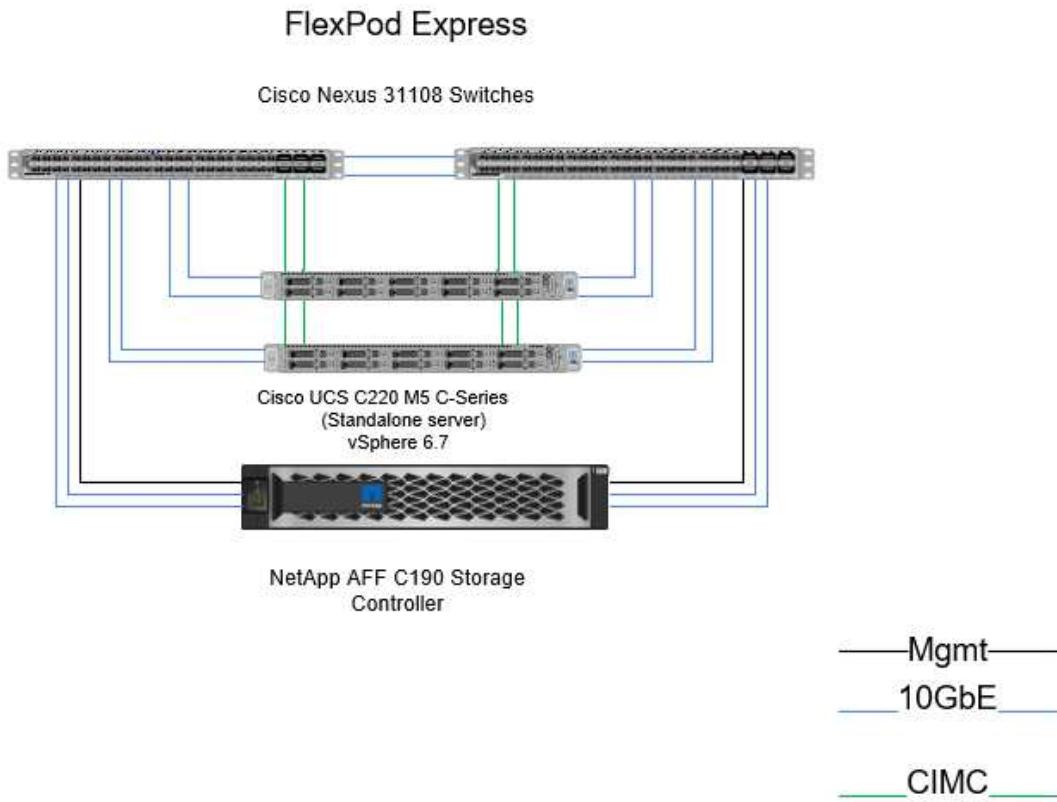


Target audience

This document is intended for people who want to take advantage of an infrastructure that is built to deliver IT efficiency and to enable IT innovation. The audience for this document includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, partner engineers, and customers.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. It features the new NetApp AFF C190 system, which runs ONTAP 9.6 software, dual Cisco Nexus 31108 switches, and Cisco UCS C220 M5 rack servers that run VMware vSphere 6.7U2. This validated solution, illustrated in the following figure, uses 10 Gigabit Ethernet (10GbE) technology. Guidance is also provided on how to scale by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.



[Next: Technology requirements.](#)

Technology requirements

FlexPod Express requires a combination of hardware and software components that depends on the selected hypervisor and network speed. In addition, FlexPod Express lays out the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, you can use a different hypervisor on the same FlexPod Express hardware.

The following table lists the hardware components that are required for this FlexPod Express configuration and to implement this solution. The hardware components that are used in any implementation of the solution can vary based on customer requirements.

Hardware	Quantity
AFF C190 2-node cluster	1
Cisco UCS C220 M5 Server	2
Cisco Nexus 31108 Switch	2

Hardware	Quantity
Cisco UCS Virtual Interface Card (VIC) 1457 for Cisco UCS C220 M5 rack server	2

Software requirements

The following table lists the software components that are required to implement the architectures of the FlexPod Express solution.

Software	Version	Details
Cisco Integrated Management Controller (CIMC)	4.0.4	For C220 M5 rack servers
Cisco NX-OS	7.0(3)I7(6)	For Cisco Nexus 31108 switches
NetApp ONTAP	9.6	For NetApp AFF C190 controllers

The following table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

Software	Version
VMware vCenter Server Appliance	6.7U2
VMware vSphere ESXi	6.7U2
NetApp VAAI Plug-In for ESXi	1.1.2
NetApp Virtual Storage Console	9.6

[Next: Design choices.](#)

Design choices

The technologies listed in this section were chosen during the architectural design phase. Each technology serves a specific purpose in the FlexPod Express infrastructure solution.

NetApp AFF C190 Series with ONTAP 9.6

This solution leverages two of the newest NetApp products: NetApp AFF C190 system and ONTAP 9.6 software.

AFF C190 system

The target group is customers who want to modernize their IT infrastructure with all- flash technology at an affordable price. The AFF C190 system comes with the new ONTAP 9.6 and flash bundle licensing, which means that the following functions are on board:

- CIFS, NFS, iSCSI, and FCP
- NetApp SnapMirror data replication software, NetApp SnapVault backup software, NetApp SnapRestore data recovery software, NetApp SnapManager storage management software product suite, and NetApp SnapCenter software

- FlexVol technology
- Deduplication, compression, and compaction
- Thin provisioning
- Storage QoS
- NetApp RAID DP technology
- NetApp Snapshot technology
- FabricPool

The following figures show the two options for host connectivity.

The following figure illustrates UTA 2 ports where SFP+ module can be inserted.



The following figure illustrates 10GBASE-T ports for connection through conventional RJ-45 Ethernet cables.



For the 10GBASE-T port option, you must have a 10GBASE-T based uplink switch.

The AFF C190 system is offered exclusively with 960GB SSDs. There are four stages of expansions from which you can choose:

- 8x 960GB
- 12x 960GB
- 18x 960GB
- 24x 960GB

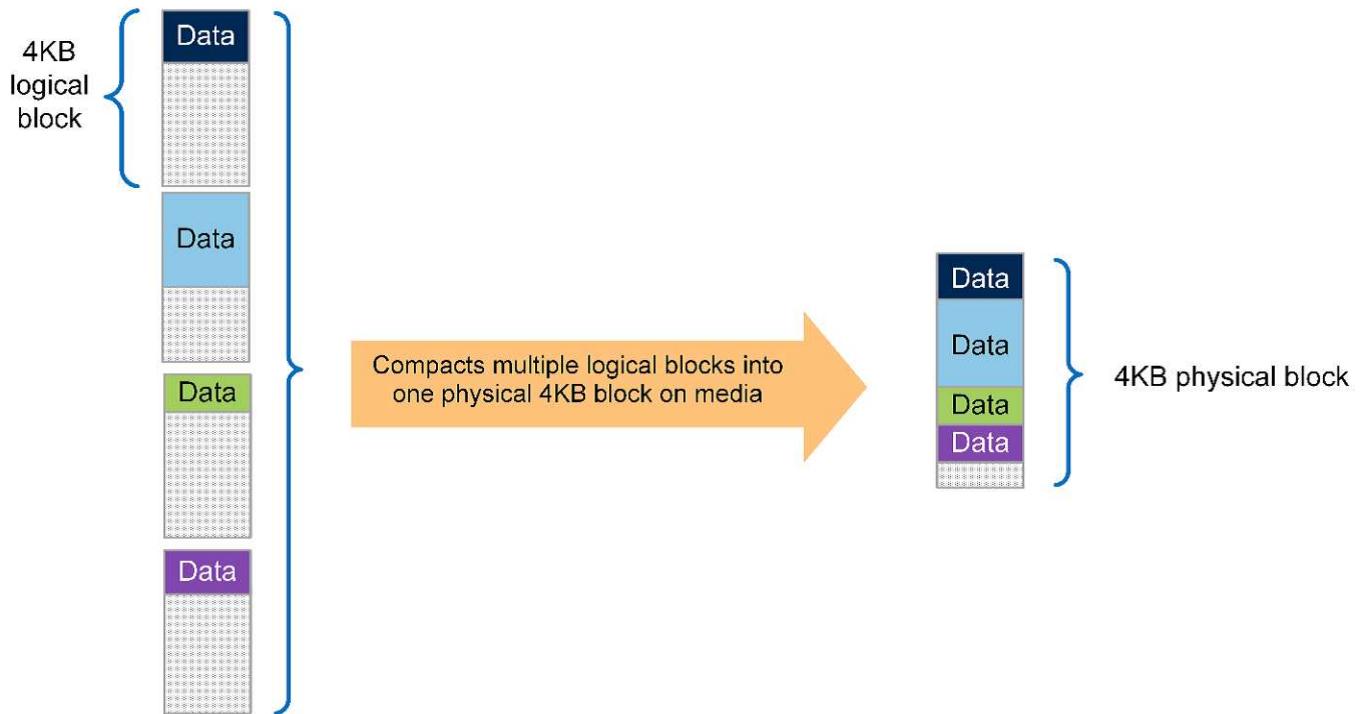
For full information about the AFF C190 hardware system, see the [NetApp AFF C190 All-Flash Array page](#).

ONTAP 9.6 software

NetApp AFF C190 systems use the new ONTAP 9.6 data management software. ONTAP 9.6 is the industry's leading enterprise data management software. It combines new levels of simplicity and flexibility with powerful

data management capabilities, storage efficiencies, and leading cloud integration.

ONTAP 9.6 has several features that are well suited for the FlexPod Express solution. Foremost is NetApp's commitment to storage efficiencies, which can be one of the most important features for small deployments. The hallmark NetApp storage efficiency features such as deduplication, compression, compaction, and thin provisioning are available in ONTAP 9.6. The NetApp WAFL system always writes 4KB blocks; therefore, compaction combines multiple logical blocks into a 4KB block when the blocks are not using their allocated space of 4KB. The following figure illustrates this process.



ONTAP 9.6 now supports an optional 512- byte block size for NVMe volumes. This capability works well with the VMware Virtual Machine File System (VMFS), which natively uses a 512-byte block. You can stay with the default 4K size or optionally set the 512-byte block size.

Other feature enhancements in ONTAP 9.6 include:

- **NetApp Aggregate Encryption (NAE).** NAE assigns keys at the aggregate level, thereby encrypting all volumes in the aggregate. This feature allows volumes to be encrypted and deduplicated at the aggregate level.
- **NetApp ONTAP FlexGroup volume enhancement.** In ONTAP 9.6, you can easily rename a FlexGroup volume. There's no need to create a new volume to migrate the data to. The volume size can also be reduced by using ONTAP System Manager or CLI.
- **FabricPool enhancement.** ONTAP 9.6 added additional support for object stores as cloud tiers. Support for Google Cloud and Alibaba Cloud Object Storage Service (OSS) was also added to the list. FabricPool supports multiple object stores, including AWS S3, Azure Blob, IBM Cloud object storage, and NetApp StorageGRID object-based storage software.
- **SnapMirror enhancement.** In ONTAP 9.6, a new volume replication relationship is encrypted by default before leaving the source array and is decrypted at the SnapMirror destination.

Cisco Nexus 3000 Series

The Cisco Nexus 31108PC-V is a 10Gbps SFP+ based top-of-rack (ToR) switch with 48 SFP+ ports and 6 QSFP28 ports. Each SFP+ port can operate in 100Mbps, 10Gbps, and each QSFP28 port can operate in native 100Gbps or 40Gbps mode or 4x 10Gbps mode, offering flexible migration options. This switch is a true PHY-less switch that is optimized for low latency and low power consumption.

The Cisco Nexus 31108PC-V specification includes the following components:

- 2.16Tbps switching capacity and forwarding rate of up to 1.2Tbps for 31108PC-V
- 48 SFP ports support 1 and 10 Gigabit Ethernet (10GbE); 6x QSFP28 ports support 4x 10GbE or 40GbE each or 100GbE

The following figure illustrates the Cisco Nexus 31108PC-V switch.



For more information about Cisco Nexus 31108PC-V switches, see [Cisco Nexus 3172PQ, 3172TQ, 3172TQ-32T, 3172PQ-XL, and 3172TQ-XL Switches Data Sheet](#).

Cisco UCS C-Series

The Cisco UCS C-Series rack server was chosen for FlexPod Express because its many configuration options allow it to be tailored for specific requirements in a FlexPod Express deployment.

Cisco UCS C-Series rack servers deliver unified computing in an industry-standard form factor to reduce TCO and to increase agility.

Cisco UCS C-Series rack servers offer the following benefits:

- A form-factor-agnostic entry point into Cisco UCS
- Simplified and fast deployment of applications
- Extension of unified computing innovations and benefits to rack servers
- Increased customer choice with unique benefits in a familiar rack package

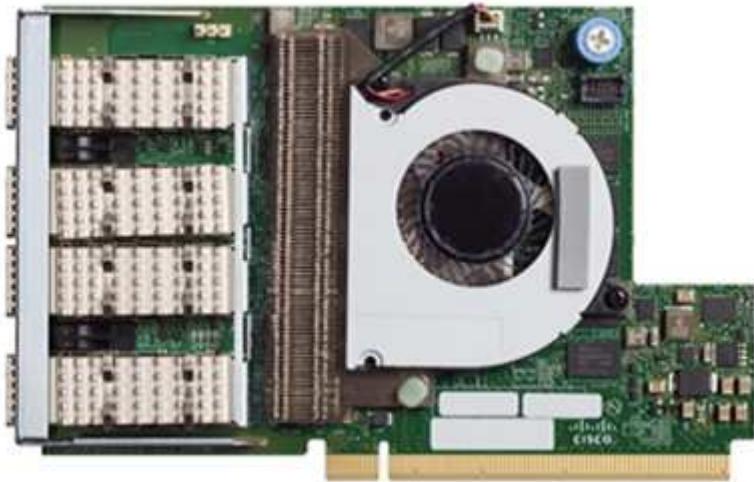


The Cisco UCS C220 M5 rack server, shown in the above figure, is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. It is a high-density two-socket rack server that delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. Cisco UCS C-Series rack servers can be deployed as standalone servers or as part of Cisco UCS to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

For more information about C220 M5 servers, see [Cisco UCS C220 M5 Rack Server Data Sheet](#).

Cisco UCS VIC 1457 connectivity for C220 M5 rack servers

The Cisco UCS VIC 1457 adapter shown in the following figure is a quad-port small form-factor pluggable (SFP28) modular LAN on motherboard (mLOM) card designed for the M5 generation of Cisco UCS C-Series Servers. The card supports 10/25Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.



For full information about the Cisco UCS VIC 1457 adapter, see [Cisco UCS Virtual Interface Card 1400 Series Data Sheet](#).

VMware vSphere 6.7U2

VMware vSphere 6.7U2 is one of the hypervisor options for use with FlexPod Express. VMware vSphere allows organizations to reduce their power and cooling footprint while confirming that the purchased compute capacity is used to its fullest. In addition, VMware vSphere allows hardware failure protection (VMware High Availability, or VMware HA) and compute resource load balancing across a cluster of vSphere hosts (VMware Distributed Resource Scheduler in maintenance mode, or VMware DRS-MM).

Because it restarts only the kernel, VMware vSphere 6.7U2 allows customers to quick boot, loading vSphere ESXi without restarting the hardware. The vSphere 6.7U2 vSphere client (HTML5-based client) has some new enhancements like Developer Center with Code Capture and API Explore. With Code Capture, you can record your actions in the vSphere client to deliver simple, usable code output. vSphere 6.7U2 also contains new features like DRS in maintenance mode (DRS-MM).

VMware vSphere 6.7U2 offers the following features:

- VMware is deprecating the external VMware Platform Services Controller (PSC) deployment model.
 - Starting with the next major vSphere release, external PSC will not be an available option.
- New protocol support for backing up and restoring a vCenter server appliance. Introducing NFS and SMB as supported protocol choices, up to 7 total (HTTP, HTTPS, FTP, FTPS, SCP, NFS, and SMB) when configuring a vCenter Server for file-based backup or restore operations.
- New functionality when using the content library. Syncing a native VM template between content libraries is now available when the vCenter Server is configured for enhanced linked mode.

- Update to the [Client Plug-Ins page](#).
- VMware vSphere Update Manager also adds enhancements to the vSphere client. You can perform attach-check compliance and remediate actions all from one screen.

For more information about VMware vSphere 6.7 U2, see the [VMware vSphere Blog page](#).

For more information about the VMware vCenter Server 6.7 U2 updates, see the [Release Notes](#).



Although this solution was validated with vSphere 6.7U2, it supports any vSphere version qualified with the other components by the [NetApp Interoperability Matrix Tool \(IMT\)](#). NetApp recommends that you deploy the next released version of vSphere for its fixes and enhanced features.

Boot architecture

The supported options for the FlexPod Express boot architecture include:

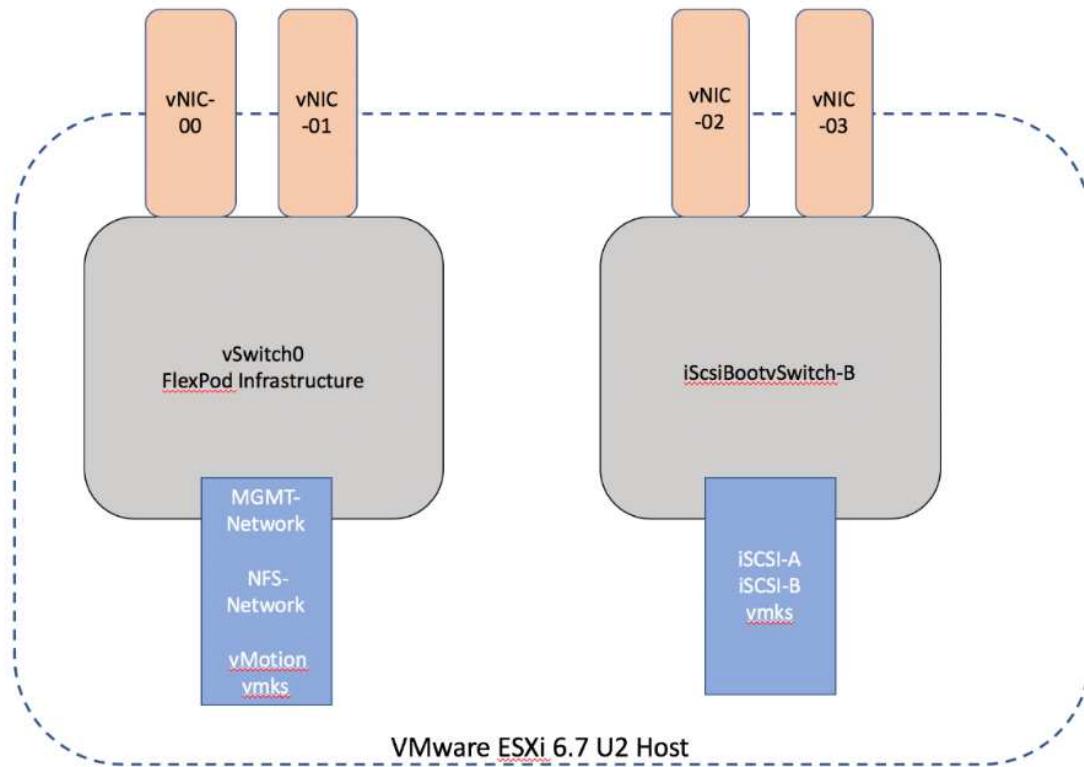
- iSCSI SAN LUN
- Cisco FlexFlash SD card
- Local disk

FlexPod Datacenter is booted from iSCSI LUNs; therefore, solution manageability is enhanced by using iSCSI boot for FlexPod Express as well.

ESXi Host Virtual Network Interface Card layout

Cisco UCS VIC 1457 has four physical ports. This solution validation includes these four physical ports in using the ESXi host. If you have a smaller or larger number of NICs, you might have different VMNIC numbers.

In an iSCSI boot implementation, iSCSI boot requires separate virtual network interface cards (vNICs) for iSCSI boot. These vNICs use the appropriate fabric's iSCSI VLAN as the native VLAN and are attached to the iSCSI boot vSwitches, as shown in the following figure.



Next: Conclusion.

Conclusion

The FlexPod Express validated design is a simple and effective solution that uses industry-leading components. By scaling and providing options for the hypervisor platform, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed for small to midsize businesses, remote and branch offices, and other businesses that require dedicated solutions.

Next: Where to find additional information.

Where to find additional information

To learn more about the information described in this document, see the following documents and websites:

- AFF and FAS System Documentation Center
<https://docs.netapp.com/platstor/index.jsp>
- AFF Documentation Resources page
<https://www.netapp.com/us/documentation/all-flash-fas.aspx>
- FlexPod Express with VMware vSphere 6.7 and NetApp AFF C190 Deployment Guide (in progress)

- NetApp documentation

<https://docs.netapp.com>

FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Deployment Guide

NVA-1142-DEPLOY: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series - NVA Deployment

Savita Kumari, NetApp

Industry trends indicate that a vast data center transformation is occurring toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices that uses technology that they are familiar with in their data center.

FlexPod® Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp® storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools to which they are accustomed. New FlexPod Express customers can easily transition to managing FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution overview

This FlexPod Express solution is part of the FlexPod Converged Infrastructure Program.

FlexPod converged infrastructure program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

The FlexPod program includes two solutions: FlexPod Express and FlexPod Datacenter.

- **FlexPod Express.** Offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter.** Delivers an optimal multipurpose foundation for various workloads and applications.

The FlexPod Portfolio

A prevalidated, flexible platform that features



FlexPod® Express

Remote office or branch office, retail, small and midsize business, and edge



FlexPod Datacenter

Enterprise apps, unified infrastructure, and virtualization

11

NetApp Verified Architecture program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. A NetApp Verified Architecture provides a NetApp solution architecture with the following qualities:

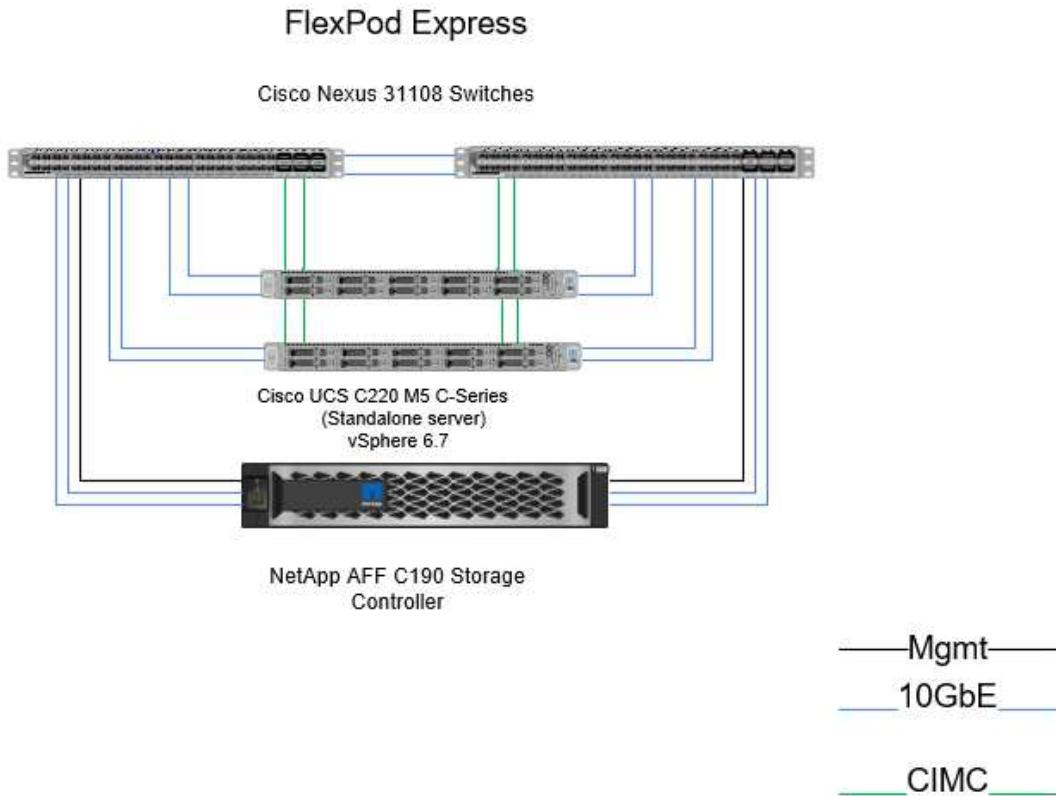
- Thoroughly tested
- Prescriptive in nature
- Minimized deployment risks
- Accelerated time to market

This guide details the design of FlexPod Express with VMware vSphere. In addition, this design uses the all-new AFF C190 system (running NetApp ONTAP® 9.6), the Cisco Nexus 31108, and Cisco UCS C-Series C220 M5 servers as hypervisor nodes.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. This solution features the new NetApp AFF C190 running ONTAP 9.6, dual Cisco Nexus 31108 switches, and Cisco UCS C220 M5 rack

servers running VMware vSphere 6.7U2. This validated solution uses 10GbE technology. Guidance is also provided on how to scale compute capacity by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.



To use the four physical 10GbE ports on the VIC 1457 efficiently, create two extra links from each server to the top rack switches.

Use case summary

The FlexPod Express solution can be applied to several use cases, including the following:

- Remote or branch offices
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is best suited for virtualized and mixed workloads. Although this solution was validated with vSphere 6.7U2, it supports any vSphere version qualified with the other components by the NetApp Interoperability Matrix Tool. NetApp recommends deploying vSphere 6.7U2 because of its fixes and enhanced features, such as the following:

- New protocol support for backing up and restoring a vCenter server appliance, including HTTP, HTTPS, FTP, FTPS, SCP, NFS and SMB.
- New functionality when utilizing the content library. Syncing of native VM templates between content libraries is now available when vCenter Server is configured for enhanced linked mode.

- An updated Client Plug-In page.
- Added enhancements in the vSphere Update Manager (VUM) and the vSphere client. You can now perform the attach, check- compliance, and remediate actions, all from one screen.

For more information on this subject, see the [vSphere 6.7U2 page](#) and the [vCenter Server 6.7U2 Release Notes](#).

Technology requirements

A FlexPod Express system requires a combination of hardware and software components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, you can use a different hypervisor on the same FlexPod Express hardware.

The following table lists the hardware components that are required for FlexPod Express configuration and implementation. The hardware components that are used in any implementation of the solution might vary based on customer requirements.

Hardware	Quantity
AFF C190 two-node cluster	1
Cisco C220 M5 server	2
Cisco Nexus 31108PC-V switch	2
Cisco UCS virtual interface card (VIC) 1457 for Cisco UCS C220 M5 rack server	2

This table lists the hardware that is required in addition to the base configuration for implementing 10GbE.

Hardware	Quantity
Cisco UCS C220 M5 server	2
Cisco VIC 1457	2

Software requirements

The following table lists the software components that are required to implement the architectures of the FlexPod Express solutions.

Software	Version	Details
Cisco Integrated Management Controller (CIMC)	4.0.4	For Cisco UCS C220 M5 rack servers
Cisco nenic driver	1.0.0.29	For VIC 1457 interface cards

Software	Version	Details
Cisco NX-OS	7.0(3)I7(6)	For Cisco Nexus 31108PC-V switches
NetApp ONTAP	9.6	For AFF C190 controllers

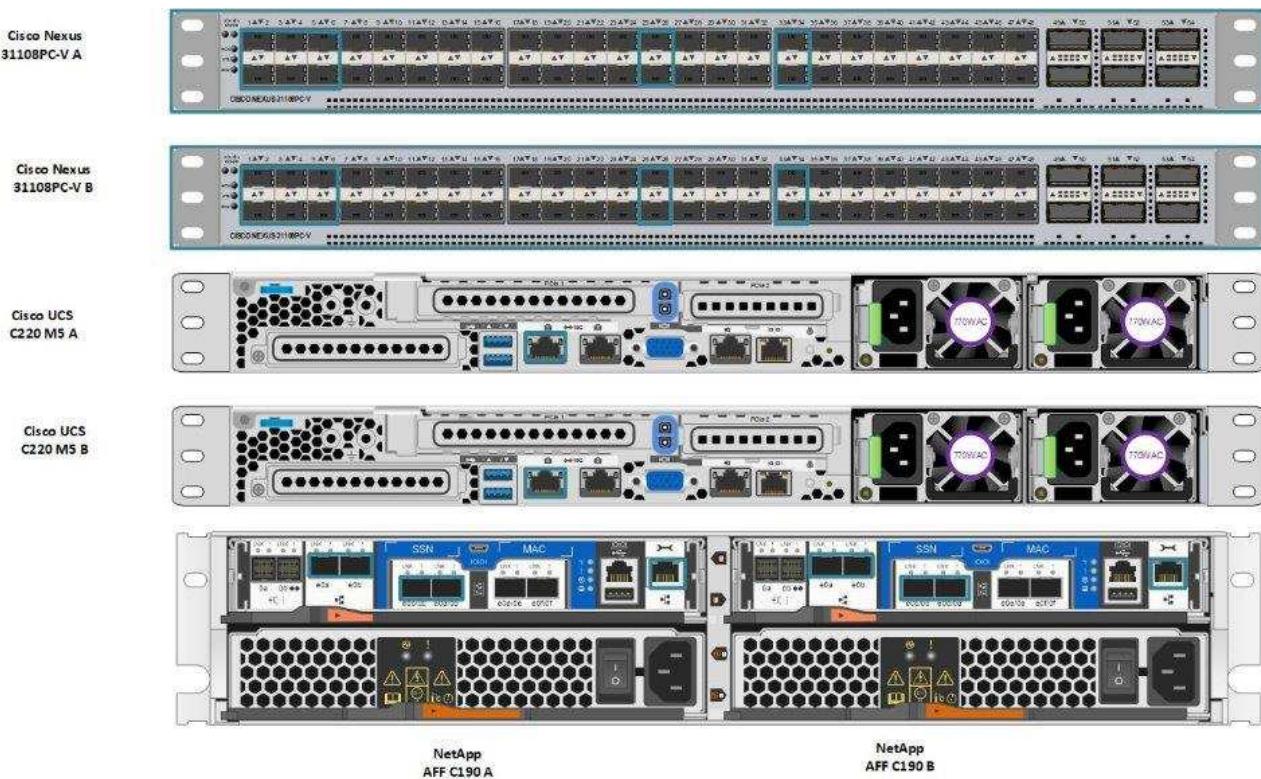
This table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

Software	Version
VMware vCenter server appliance	6.7U2
VMware vSphere ESXi hypervisor	6.7U2
NetApp VAAI Plug-In for ESXi	1.1.2
NetApp VSC	9.6

FlexPod Express cabling information

This reference validation is cabled as shown in the following figures and tables.

This figure shows the reference validation cabling.



The following table lists the cabling information for Cisco Nexus switch 31108PC-V-A.

Local device	Local port	Remote device	Remote port
Cisco Nexus switch 31108PC-V A	Eth1/1	NetApp AFF C190 storage controller A	e0c
	Eth1/2	NetApp AFF C190 storage controller B	e0c
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM0
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM0
	Eth1/5	Cisco UCS C220 C-Series standalone server A	MLOM1
	Eth1/6	Cisco UCS C220 C-Series standalone server B	MLOM1
	Eth1/25	Cisco Nexus switch 31108PC-V B	Eth1/25
	Eth1/26	Cisco Nexus switch 31108PC-V B	Eth1/26
	Eth1/33	NetApp AFF C190 storage controller A	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server A	CIMC (FEX135/1/25)

This table lists the cabling information for Cisco Nexus switch 31108PC-V- B.

Local device	Local port	Remote device	Remote port
Cisco Nexus switch 31108PC-V B	Eth1/1	NetApp AFF C190 storage controller A	e0d
	Eth1/2	NetApp AFF C190 storage controller B	e0d
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM2
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM2
	Eth1/5	Cisco UCS C220 C-Series standalone server A	MLOM3
	Eth1/6	Cisco UCS C220 C-Series standalone server B	MLOM3
	Eth1/25	Cisco Nexus switch 31108 A	Eth1/25
	Eth1/26	Cisco Nexus switch 31108 A	Eth1/26
	Eth1/33	NetApp AFF C190 storage controller B	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server B	CIMC (FEX135/1/26)

This table lists the cabling information for NetApp AFF C190 storage controller A.

Local device	Local Port	Remote device	Remote port
NetApp AFF C190 storage controller A	e0a	NetApp AFF C190 storage controller B	e0a
	e0b	NetApp AFF C190 storage controller B	e0b
	e0c	Cisco Nexus switch 31108PC-V A	Eth1/1
	e0d	Cisco Nexus switch 31108PC-V B	Eth1/1
	e0M	Cisco Nexus switch 31108PC-V A	Eth1/33

This table lists the cabling information for NetApp AFF C190 storage controller B.

Local device	Local port	Remote device	Remote port
NetApp AFF C190 storage controller B	e0a	NetApp AFF C190 storage controller A	e0a
	e0b	NetApp AFF C190 storage controller A	e0b
	e0c	Cisco Nexus switch 31108PC-V A	Eth1/2
	e0d	Cisco Nexus switch 31108PC-V B	Eth1/2
	e0M	Cisco Nexus switch 31108PC-V B	Eth1/33

Deployment procedures

Overview

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches.

In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, <<text>> appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01> network port vlan create -node <<var_nodeA>> -vlan-name
<<var_vlan-name>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. The following table describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.



If you use separate in-band and out-of-band management VLANs, you must create a layer-3 route between them. For this validation, a common management VLAN was used.

VLAN name	VLAN purpose	VLAN ID	
Management VLAN	VLAN for management interfaces	3437	vSwitch0
NFS VLAN	VLAN for NFS traffic	3438	vSwitch0

VLAN name	VLAN purpose	VLAN ID	
VMware vMotion VLAN	VLAN designated for the movement of virtual machines (VMs) from one physical host to another	3441	vSwitch0
VM traffic VLAN	VLAN for VM application traffic	3442	vSwitch0
iSCSI-A-VLAN	VLAN for iSCSI traffic on fabric A	3439	iScsiBootvSwitch
iSCSI-B-VLAN	VLAN for iSCSI traffic on fabric B	3440	iScsiBootvSwitch
Native VLAN	VLAN to which untagged frames are assigned	2	

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <<var_xxxx_vlan>>, where xxxx is the purpose of the VLAN (such as iSCSI-A).

There are two vSwitches created in this validation.

The following table lists the solution vSwitches.

vSwitch name	Active adapters	Ports	MTU	Load balancing
vSwitch0	Vmnic2, vmnic4	default (120)	9000	Route based on IP hash
iScsiBootvSwitch	Vmnic3, vmnic5	default (120)	9000	Route based on the originating virtual port ID.

 The IP hash method of load balancing requires proper configuration for the underlying physical switch using SRC-DST-IP EtherChannel with a static (mode on) port-channel. In the event of intermittent connectivity due to possible switch misconfiguration, temporarily shut down one of the two associated uplink ports on the Cisco switch to restore communication to the ESXi management vmkernel port while troubleshooting the port-channel settings.

The following table lists the VMware VMs that are created.

VM description	Host name
VMware vCenter Server	FlexPod-VCSA
Virtual Storage Console	FlexPod-VSC

Deploy Cisco Nexus 31108PC-V

This section details the Cisco Nexus 331108PC-V switch configuration used in a FlexPod Express environment.

Initial Setup of Cisco Nexus 31108PC-V Switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.



This procedure assumes that you are using a Cisco Nexus 31108PC-V running NX-OS software release 7.0(3)I7(6).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 31108PC-V switches can be connected to an existing management network, or the mgmt0 interfaces of the 31108PC-V switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.



In this deployment guide, the FlexPod Express Cisco Nexus 31108PC-V switches are connected to an existing management network.

3. To configure the Cisco Nexus 31108PC-V switches, power on the switch and follow the on-screen prompts, as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

*Note: setup is mainly used for configuring the system initially, when no configuration is present. So setup always assumes system defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): y
Do you want to enforce secure password standard (yes/no) [y]: y

Create another login account (yes/no) [n]: n

Configure read-only SNMP community string (yes/no) [n]: n

Configure read-write SNMP community string (yes/no) [n]: n

Enter the switch name : 31108PC-V-B

Continue with Out-of-band (mgmt0) management configuration? (yes/no)
[y]: y

Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>

Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>

Configure the default gateway? (yes/no) [y]: y

IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>

Configure advanced IP options? (yes/no) [n]: n

Enable the telnet service? (yes/no) [n]: n

Enable the ssh service? (yes/no) [y]: y

Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa

Number of rsa key bits <1024-2048> [1024]: <enter>

Configure the ntp server? (yes/no) [n]: y

NTP server IPv4 address : <<var_ntp_ip>>

Configure default interface layer (L3/L2) [L2]: <enter>

Configure default switchport interface state (shut/noshut) [noshut]:
<enter>

Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: <enter>

4. You then see a summary of your configuration, and you are asked if you would like to edit it. If your configuration is correct, enter n.

Would you like to edit the configuration? (y
es/no) [n]: n

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

Use this configuration and save it? (yes/no) [y]: Enter

6. Repeat this procedure for Cisco Nexus switch B.

Enable the advanced features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode using the command (config t) and run the following commands:

```
feature interface-vlan  
feature lacp  
feature vpc
```

 The default port channel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the port channel. You can achieve better distribution across the members of the port channel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

From configuration mode (config t), enter the following commands to set the global port channel load-balancing configuration on Cisco Nexus switch A and switch B:

```
port-channel load-balance src-dst ip-14port
```

Configure global spanning tree

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (config t), run the following commands to configure the default spanning tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
ntp server <<var_ntp_ip>> use-vrf management
ntp master 3
```

Define the VLANs

Before individual ports with different VLANs are configured, the layer- 2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (config t), run the following commands to define and describe the layer- 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure access and management port descriptions

As is the case with assigning names to the layer- 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (config t) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus Switch A

```

int eth1/1
  description AFF C190-A e0c
int eth1/2
  description AFF C190-B e0c
int eth1/3
  description UCS-Server-A: MLOM port 0 vSwitch0
int eth1/4
  description UCS-Server-B: MLOM port 0 vSwitch0
int eth1/5
  description UCS-Server-A: MLOM port 1 iScsiBootvSwitch
int eth1/6
  description UCS-Server-B: MLOM port 1 iScsiBootvSwitch
int eth1/25
  description vPC peer-link 31108PC-V-B 1/25
int eth1/26
  description vPC peer-link 31108PC-V-B 1/26
int eth1/33
  description AFF C190-A e0M
int eth1/34
  description UCS Server A: CIMC

```

Cisco Nexus Switch B

```

int eth1/1
  description AFF C190-A e0d
int eth1/2
  description AFF C190-B e0d
int eth1/3
  description UCS-Server-A: MLOM port 2 vSwitch0
int eth1/4
  description UCS-Server-B: MLOM port 2 vSwitch0
int eth1/5
  description UCS-Server-A: MLOM port 3 iScsiBootvSwitch
int eth1/6
  description UCS-Server-B: MLOM port 3 iScsiBootvSwitch
int eth1/25
  description vPC peer-link 31108PC-V-A 1/25
int eth1/26
  description vPC peer-link 31108PC-V-A 1/26
int eth1/33
  description AFF C190-B e0M
int eth1/34
  description UCS Server B: CIMC

```

Configure server and storage management interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (config t), enter the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus Switch A

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Cisco Nexus Switch B

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Perform the virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer- 2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree- protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back mgmt0 configuration, use the addresses defined on the interfaces and verify that they

can communicate by using the `ping <<switch_A/B_mgmt0_ip_addr>>vrf` management command.

From configuration mode (config t), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus Switch A

```
vpc domain 1
  role priority 10
  peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source
<<switch_A_mgmt0_ip_addr>> vrf
management
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>, <<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
exit
copy run start
```

Cisco Nexus Switch B

```

vpc domain 1
  peer-switch
    role priority 20
    peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source
    <<switch_B_mgmt0_ip_addr>> vrf management
    peer-gateway
    auto-recovery
    delay-restore 150
    ip arp synchronize
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<nfs_vlan_id>>, <<vmotion_vlan_id>>,
    <<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
    <<iSCSI_B_vlan_id>>
    spanning-tree port type network
    vpc peer-link
no shut
exit
copy run start

```

Configure the storage port channels

The NetApp storage controllers allow an active-active connection to the network using the Link Aggregation Control Protocol (LACP). The use of LACP is preferred because it adds both negotiation and logging between the switches. Because the network is set up for vPC, this approach enables you to have active-active connections from the storage to separate physical switches. Each controller has two links to each of the switches. However, all four links are part of the same vPC and interface group (ifgrp).

From configuration mode (config t), run the following commands on each of the switches to configure the individual interfaces and the resulting port channel configuration for the ports connected to the NetApp AFF controller.

1. Run the following commands on switch A and switch B to configure the port channels for storage controller A:

```

int eth1/1
  channel-group 11 mode active
int Po11
  description vPC to Controller-A
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<nfs_vlan_id>>,<<mgmt_vlan_id>>,<<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11
  no shut

```

2. Run the following commands on switch A and switch B to configure the port channels for storage controller B:

```

int eth1/2
  channel-group 12 mode active
int Po12
  description vPC to Controller-B
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<mgmt_vlan_id>>,
<<iSCSI_A_vlan_id>>,<<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12
  no shut
exit
copy run start

```

Configure the server connections

The Cisco UCS servers have a four-port virtual interface card, VIC1457, that is used for data traffic and booting of the ESXi operating system using iSCSI. These interfaces are configured to fail over to one another, providing additional redundancy beyond a single link. Spreading these links across multiple switches enables the server to survive even a complete switch failure.

From configuration mode (config t), run the following commands to configure the port settings for the interfaces connected to each server.

Cisco Nexus Switch A: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/5
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<iSCSI_A_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_i
d>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  no shut
exit
copy run start
```

Cisco Nexus Switch B: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/6
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<iSCSI_B_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_i
d>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  no shut
exit
copy run start
```

Configure the server port channels

Run the following commands on switch A and switch B to configure the port channels for Server-A:

```

int eth1/3
  channel-group 13 mode active
int Po13
  description vPC to Server-A
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 13
  no shut

```

Run the following commands on switch A and switch B to configure the port channels for Server-B:

```

int eth1/4
  channel-group 14 mode active
int Po14
  description vPC to Server-B
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 14
  no shut

```

 An MTU of 9000 was used in this solution validation. However, you can configure an different value for the MTU appropriate for your application requirements. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components result in packets being dropped and these packets will need to be transmitted again, affecting the overall performance of the solution.

 To scale the solution by adding additional Cisco UCS servers, run the previous commands with the switch ports that the newly added servers have been plugged into on switches A and B.

Uplink into an existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 31108 switches included in the FlexPod environment into the infrastructure. The uplinks can be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if

required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy start to save the configuration on each switch after the configuration is completed.

Next: [NetApp storage deployment procedure \(part 1\)](#).

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF C190 Series installation

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

Access the [HWU](#) application to view the system configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFFC190 Series prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections:

- Electrical Requirements
- Supported Power Cords
- Onboard Ports and Cables

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF C190 Documentation](#).

NetApp ONTAP 9.6

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the ONTAP 9.6 Software Setup Guide.



This system is set up in a two-node switchless cluster configuration.

The following table provides the ONTAP 9.6 installation and configuration information.

Cluster detail	Cluster detail value
Cluster node A IP address	<<var_nodeA_mgmt_ip>>
Cluster node A netmask	<<var_nodeA_mgmt_mask>>
Cluster node A gateway	<<var_nodeA_mgmt_gateway>>
Cluster node A name	<<var_nodeA>>
Cluster node B IP address	<<var_nodeB_mgmt_ip>>
Cluster node B netmask	<<var_nodeB_mgmt_mask>>
Cluster node B gateway	<<var_nodeB_mgmt_gateway>>
Cluster node B name	<<var_nodeB>>
ONTAP 9.6 URL	<<var_url_boot_software>>
Name for cluster	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster B gateway	<<var_clustermgmt_gateway>>
Cluster B netmask	<<var_clustermgmt_mask>>
Domain name	<<var_domain_name>>
DNS server IP (you can enter more than one)	<var_dns_server_ip>
NTP server IP (you can enter more than one)	<<var_ntp_server_ip>>

Configure Node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

Allow the system to boot.

autoboot

2. Press Ctrl-C to enter the Boot menu.



If ONTAP 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

3. To install new software, select option 7.
4. Enter y to perform an upgrade.

5. Select e0M for the network port you want to use for the download.
6. Enter y to reboot now.
7. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>
```

8. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

9. Press Enter for the user name, indicating no user name.
10. Enter y to set the newly installed software as the default to be used for subsequent reboots.
11. Enter y to reboot the node.



When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

12. Press Ctrl-C to enter the Boot menu.
13. Select option 4 for Clean Configuration and Initialize All Disks.
14. Enter y to zero disks, reset config, and install a new file system.
15. Enter y to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

While node A is initializing, begin configuring node B.

Configure Node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press Ctrl-C to enter the Boot menu.

autoboot

3. Press Ctrl-C when prompted.



If ONTAP 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.A.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.



When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation of the node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.6 boots on the node for the first time.

 The node and cluster setup procedure has changed slightly in ONTAP 9.6. The cluster setup wizard is now used to configure the first node in a cluster, and NetApp ONTAP System Manager (formerly OnCommand® System Manager) is used to configure the cluster.

1. Follow the prompts to set up node A.

```
Welcome to the cluster setup wizard.  
You can enter the following commands at any time:  
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.  
You can return to cluster setup at any time by typing "cluster setup".  
To accept a default or omit a question, do not enter a value.  
This system will send event messages and periodic reports to NetApp  
Technical  
Support. To disable this feature, enter  
autosupport modify -support disable  
within 24 hours.  
Enabling AutoSupport can significantly speed problem determination and  
resolution should a problem occur on your system.  
For further information on AutoSupport, see:  
http://support.netapp.com/autosupport  
Type yes to confirm and continue {yes}: yes  
Enter the node management interface port [e0M]:  
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>  
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>  
Enter the node management interface default gateway:  
<<var_nodeA_mgmt_gateway>>  
A node management interface on port e0M with IP address  
<<var_nodeA_mgmt_ip>> has been created.  
Use your web browser to complete cluster setup by accessing  
https://<<var\_nodeA\_mgmt\_ip>>  
Otherwise, press Enter to complete cluster setup using the command line  
interface:
```

2. Navigate to the IP address of the node's management interface.



Cluster setup can also be performed by using the CLI. This document describes cluster setup using System Manager guided setup.

3. Click Guided Setup to configure the cluster.

4. Enter <<var_clustername>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.
5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter <<var_clustermgmt_ip>> in the Cluster Management IP Address field, <<var_clustermgmt_mask>> in the Netmask field, and <<var_clustermgmt_gateway>> in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter <<var_nodeA_mgmt_ip>> for node B.
 - d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.



You can enter multiple DNS server IP addresses.

- e. Enter 10.63.172.162 in the Primary NTP Server field.



You can also enter an alternate NTP server. The IP address 10.63.172.162 from <<var_ntp_server_ip>> is the Nexus Mgmt IP.

8. Configure the support information.

- a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
- b. Enter the SMTP mail host and email address for event notifications.



You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



② AutoSupport

③ Proxy URL (Optional)

ⓘ Connection is verified after configuring AutoSupport on all nodes.

④ Event Notifications

Notify me through:

<input checked="" type="checkbox"/> Email	SMTP Mail Host	Email Addresses
	<input type="text"/>	Separate email addresses with a comma...
<hr/>		
<input type="checkbox"/> SNMP	SNMP Trap Host	
	<input type="text"/>	<hr/>
<hr/>		
<input type="checkbox"/> Syslog	Syslog Server	
	<input type="text"/>	<hr/>
<hr/>		

Submit

When the system indicates that the cluster configuration has completed, click **Manage Your Cluster** to configure the storage.

Continuation of the storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set the on-board UTA2 ports personality

1. Verify the current mode and the current type for the ports by running the `ucadmin show` command.

```
AFF C190::> ucadmin show
      Current  Current  Pending  Pending  Admin
Node     Adapter  Mode    Type     Mode    Type  Status
-----  -----  -----  -----  -----  -----
-----  -----
AFF C190_A  0c     cna    target   -      -      online
AFF C190_A  0d     cna    target   -      -      online
AFF C190_A  0e     cna    target   -      -      online
AFF C190_A  0f     cna    target   -      -      online
AFF C190_B  0c     cna    target   -      -      online
AFF C190_B  0d     cna    target   -      -      online
AFF C190_B  0e     cna    target   -      -      online
AFF C190_B  0f     cna    target   -      -      online
8 entries were displayed.
```

2. Verify that the current mode of the ports that are in use is cna and that the current type is set to target. If not, change the port personality by using the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode
cna -type target
```



The ports must be offline to run the previous command. To take a port offline, run the following command:

```
network fcp adapter modify -node <home node of the port> -adapter <port
name> -state down
```



If you changed the port personality, you must reboot each node for the change to take effect.

Rename the management logical interfaces

To rename the management logical interfaces (LIFs), complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clustername>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clustername>> -lif
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif
cluster_setup_node_mgmt_lif_AFF_C190_B_1 -newname AFF_C190-02_mgmt1
```

Set auto-revert on cluster management

Set the auto-revert parameter on the cluster management interface.

```
network interface modify -vserver <<clustername>> -lif cluster_mgmt -auto-
revert true
```

Set up the service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_nodeA>> -address
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeA_sp_ip>>
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>
system service-processor network modify -node <<var_nodeB>> -address
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeB_sp_ip>>
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```



Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```



Enabling failover on one node enables it for both nodes.

3. Verify the HA status of the two-node cluster.



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.



Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```



The message **Keep Alive Status: Error:** indicates that one of the controllers did not receive hwassist keep alive alerts from its partner, indicating that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node <<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node <<var_nodeB>>
```

Create a jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove the data ports from the default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following command:

```
net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
```

Configure the interface group LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. make sure it's configured based on the steps in this guide in section 5.1.

From the cluster prompt, complete the following steps:

```
ifgrp create -node <<var_nodeA>> -ifgrp a0a -distr-func port -mode  
multimode_lacp  
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0c  
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0d  
ifgrp create -node <<var_nodeB>> -ifgrp a0a -distr-func port -mode  
multimode_lacp  
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0c  
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0d
```

Configure the jumbo frames in ONTAP

To configure an ONTAP network port to use jumbo frames (usually with an MTU of 9,000 bytes), run the following commands from the cluster shell:

```
AFF C190::> network port modify -node node_A -port a0a -mtu 9000  
Warning: This command will cause a several second interruption of service  
on  
      this network port.  
Do you want to continue? {y|n}: y  
AFF C190::> network port modify -node node_B -port a0a -mtu 9000  
Warning: This command will cause a several second interruption of service  
on  
      this network port.  
Do you want to continue? {y|n}: y
```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```
network port vlan create -node <<var_nodeA>> -vlan-name a0a-  
<<var_nfs_vlan_id>>  
network port vlan create -node <<var_nodeB>> -vlan-name a0a-  
<<var_nfs_vlan_id>>  
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports  
<<var_nodeA>>:a0a-<<var_nfs_vlan_id>>, <<var_nodeB>>:a0a-  
<<var_nfs_vlan_id>>
```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_A_id>>,<<var_nodeB>>:a0a-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_B_id>>,<<var_nodeB>>:a0a-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<mgmt_vlan_id>>

```

Create data aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```



Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.



Start with five disks; you can add disks to an aggregate when additional storage is required.



The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1_nodeA` is online.

Configure Time Zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is America/New_York. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select md5 as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select des as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts
<<var_mailhost>> -transport https -support enable -noteto
<<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the vserver create command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate
aggr1_nodeA -rootvolume-security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fcp
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the SVM vstorage parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
vserver nfs modify -vserver Infra-SVM -vstorage enabled  
vserver nfs show
```



Commands are prefaced by `vserver` in the command line because SVMs were previously called Vservers.

Configure NFSv3 in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
ESXi host A NFS IP address	<code><<var_esxi_hostA_nfs_ip>></code>
ESXi host B NFS IP address	<code><<var_esxi_hostB_nfs_ip>></code>

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 2 -protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS C-Series servers are added.

Create the iSCSI service in ONTAP

To create the iSCSI service on the SVM, run the following command. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM  
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

To create a load-sharing mirror of the SVM root volume in ONTAP, complete the following steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate  
aggr1_nodeA -size 1GB -type DP  
volume create -vserver Infra_Vserver -volume rootvol_m02 -aggregate  
aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path  
Infra-SVM:rootvol_m01 -type LS -schedule 15min  
snapmirror create -source-path Infra-SVM:rootvol -destination-path  
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol  
snapmirror show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag  
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.



Deleting expired certificates before creating certificates is a best practice. Run the security certificate delete command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...
```

Example: security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial 552429A6

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...
```

Example: security certificate create -common-name infra-svm.netapp.com -type server -size 2048 -country US -state "North Carolina" -locality "RTP" -organization "NetApp" -unit "FlexPod" -email-addr "abc@netapp.com" -expire-days 3650 -protocol SSL -hash-function SHA256 -vserver Infra-SVM

5. To obtain the values for the parameters required in the following step, run the security certificate show command.
6. Enable each certificate that was just created using the -server-enabled true and -client-enabled false parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...
```

Example: security ssl modify -vserver Infra-SVM -server-enabled true -client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common-name infra-svm.netapp.com

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http
-vserver <<var_clusternamespace>>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow the SVM to be available by the web.

```
set -privilege admin
vserver services web modify -name spi -vserver * -enabled true
```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol® volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore -aggregate
aggr1_nodeB -size 500GB -state online -policy default -junction-path
/infra_datastore -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap
-space-guarantee none -percent-snapshot-space 0 -snapshot-policy none
-efficiency-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0
```

Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware -space-reserve disabled
```



When adding an extra Cisco UCS C-Series server, you must create an extra boot LUN.

Create iSCSI LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
Storage node A iSCSI LIF01A	<<var_nodeA_iscsi_lif01a_ip>>
Storage node A iSCSI LIF01A network mask	<<var_nodeA_iscsi_lif01a_mask>>
Storage node A iSCSI LIF01B	<<var_nodeA_iscsi_lif01b_ip>>
Storage node A iSCSI LIF01B network mask	<<var_nodeA_iscsi_lif01b_mask>>
Storage node B iSCSI LIF01A	<<var_nodeB_iscsi_lif01a_ip>>
Storage node B iSCSI LIF01A network mask	<<var_nodeB_iscsi_lif01a_mask>>
Storage node B iSCSI LIF01B	<<var_nodeB_iscsi_lif01b_ip>>
Storage node B iSCSI LIF01B network mask	<<var_nodeB_iscsi_lif01b_mask>>

Create four iSCSI LIFs, two on each node.

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up -failover-policy disabled
-firmware-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up -failover-policy disabled
-firmware-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up -failover-policy disabled
-firmware-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up -failover-policy disabled
-firmware-policy data -auto-revert false
network interface show
```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
Storage node A NFS LIF 01 IP	<<var_nodeA_nfs_lif_01_ip>>
Storage node A NFS LIF 01 network mask	<<var_nodeA_nfs_lif_01_mask>>
Storage node B NFS LIF 02 IP	<<var_nodeB_nfs_lif_02_ip>>
Storage node B NFS LIF 02 network mask	<<var_nodeB_nfs_lif_02_mask>>

Create an NFS LIF.

```
network interface create -vserver Infra-SVM -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_ip>> -netmask <<
var_nodeA_nfs_lif_01_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_ip>> -netmask <<
var_nodeB_nfs_lif_02_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface show
```

Add an infrastructure SVM administrator

The following table lists the information needed to add an SVM administrator.

Detail	Detail value
Vsmgmt IP	<<var_svm_mgmt_ip>>
Vsmgmt network mask	<<var_svm_mgmt_mask>>
Vsmgmt default gateway	<<var_svm_mgmt_gateway>>

To add the infrastructure SVM administrator and SVM administration logical interface to the management network, complete the following steps:

1. Run the following command:

```
network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true
```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway
<<var_svm_mgmt_gateway>>
network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
security login unlock -username vsadmin -vserver Infra-SVM
```

Next: [Deploy Cisco UCS C-Series rack server](#).

Deploy Cisco UCS C-Series rack server

This section provides a detailed procedure for configuring a Cisco UCS C-Series standalone rack server for use in the FlexPod Express configuration.

Perform the initial Cisco UCS C-Series standalone server setup for CIMC

Complete these steps for the initial setup of the CIMC interface for Cisco UCS C-Series standalone servers.

The following table lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.

Detail	Detail value
CIMC IP address	<<cimc_ip>>
CIMC subnet mask	\<<cimc_netmask>>
CIMC default gateway	<<cimc_gateway>>



The CIMC version used in this validation is CIMC 4.0.(4).

All servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.

Power on the server and press F8 when prompted to enter the CIMC configuration.



Copyright (c) 2019 Cisco Systems, Inc.

Press <F2> BIOS Setup : <F6> Boot Menu : <F7> Diagnostics
Press <F8> CIMC Setup : <F12> Network Boot
Bios Version : C220M5.4.0.4g.0.0712190011
Platform ID : C220M5

Processor(s) Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz
Total Memory = 64 GB Effective Memory = 64 GB
Memory Operating Speed 2400 Mhz
M.2 SWRAID configuration is not detected. Switching to AHCI mode.

Cisco IMC IPv4 Address : 10.63.172.160
Cisco IMC MAC Address : 70:69:5A:B5:8D:68

Entering CIMC Configuration Utility ...

92

2. In the CIMC configuration utility, set the following options:

a. Network interface card (NIC) mode:

Dedicated

b. IP (Basic):

IPV4:

DHCP enabled:

CIMC IP: <<cimc_ip>>

Prefix/Subnet: <<cimc_netmask>>

Gateway: <<cimc_gateway>>

c. VLAN (Advanced): Leave cleared to disable VLAN tagging.

NIC redundancy

None:

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated: [X] NIC redundancy
Shared LOM: [ ] None: [X]
Cisco Card: Active-standby: [ ]
Riser1: [ ] Active-active: [ ]
Riser2: [ ] VLAN (Advanced)
MLom: [ ] VLAN enabled: [ ]
Shared LOM Ext: [ ] VLAN ID: 1
Priority: 0
IP (Basic)
IPV4: [X] IPV6: [ ]
DHCP enabled: [ ]
CIMC IP: 10.63.172.160
Prefix/Subnet: 255.255.255.0
Gateway: 10.63.172.1
Pref DNS Server: 0.0.0.0
Smart Access USB
Enabled: [ ]
*****
<Up/Down>Selection <F10>Save <Space>Enable/Disable <F5>Refresh <ESC>Exit
<F1>Additional settings

```

3. Press F1 to see the additional settings:

a. Common properties:

Host name: <<esxi_host_name>>

Dynamic DNS: []

Factory defaults: Leave cleared.

b. Default user (basic):

Default password: <<admin_password>>

Reenter password: <<admin_password>>

Port properties: Use default values.

Port profiles: Leave cleared.

4. Press F10 to save the CIMC interface configuration.

5. After the configuration is saved, press Esc to exit.

Configure Cisco UCS C-Series Servers iSCSI boot

In this FlexPod Express configuration, the VIC1457 is used for iSCSI boot.

The following table lists the information needed to configure iSCSI boot.



An italicized font indicates variables that are unique for each ESXi host.

Detail	Detail value
ESXi host initiator A name	<<var_ucs_initiator_name_A>>
ESXi host iSCSI-A IP	<<var_esxi_host_iscsiA_ip>>
ESXi host iSCSI-A network mask	<<var_esxi_host_iscsiA_mask>>
ESXi host iSCSI A default gateway	<<var_esxi_host_iscsiA_gateway>>
ESXi host initiator B name	<<var_ucs_initiator_name_B>>
ESXi host iSCSI-B IP	<<var_esxi_host_iscsiB_ip>>
ESXi host iSCSI-B network mask	<<var_esxi_host_iscsiB_mask>>
ESXi host iSCSI-B gateway	<<var_esxi_host_iscsiB_gateway>>
IP address iscsi_lif01a	<<var_iscsi_lif01a>>
IP address iscsi_lif02a	<<var_iscsi_lif02a>>
IP address iscsi_lif01b	<<var_iscsi_lif01b>>
IP address iscsi_lif02b	<<var_iscsi_lif02b>>
Infra_SVM IQN	<<var_SVM_IQN>>

Boot order configuration

To set the boot order configuration, complete the following steps:

1. From the CIMC interface browser window, click the Compute tab and select BIOS.
2. Click Configure Boot Order and then click OK.

Cisco Integrated Management Controller

/ Compute / BIOS

BIOS Remote Management Troubleshooting Power Policies PID Catalog

Enter BIOS Setup | Clear BIOS CMOS | Restore Manufacturing Custom Settings | Restore Defaults

Configure BIOS Configure Boot Order Configure BIOS Profile

BIOS Properties

Running Version: C220M5.4.0.4g.0.0712190011
 UEFI Secure Boot
Actual Boot Mode: Uefi
Configured Boot Mode:
Last Configured Boot Order Source: BIOS
Configured One time boot device:

Save Changes

Configured Boot Devices

- Basic
- Advanced

Actual Boot Devices

- UEFI: Built-in EFI Shell (NonPolicyTarget)
- UEFI: PXE IP4 Intel(R) Ethernet Controller X550 (NonPolicyTarget)
- UEFI: PXE IP4 Intel(R) Ethernet Controller X550 (NonPolicyTarget)

Configure Boot Order

3. Configure the following devices by clicking the device under Add Boot Device and going to the Advanced tab:

a. Add Virtual Media:

Name: KVM-CD-DVD

Subtype: KVM MAPPED DVD

State: Enabled

Order: 1

b. Add iSCSI Boot:

Name: iSCSI-A

State: Enabled

Order: 2

Slot: MLOM

Port: 1

c. Click Add iSCSI Boot:

Name: iSCSI-B

State: Enabled

Order: 3

Slot: MLOM

Port: 3

4. Click Add Device.

5. Click Save Changes and then click Close.

	Name	Type	Order	State
<input checked="" type="checkbox"/>	KVM-MAPPED-DVD	VMEDIA	1	Enabled
<input type="checkbox"/>	iSCSI-A	ISCSI	2	Enabled
<input type="checkbox"/>	iSCSI-B	ISCSI	3	Enabled

6. Reboot the server to boot with your new boot order.

Disable RAID controller (if present)

Complete the following steps if your C-Series server contains a RAID controller. A RAID controller is not needed in the boot from SAN configuration. Optionally, you can also physically remove the RAID controller from the server.

1. Under the Compute tab, click BIOS in the left navigation pane in CIMC.
2. Select Configure BIOS.
3. Scroll down to PCIe Slot:HBA Option ROM.
4. If the value is not already disabled, set it to disabled.

Note: Default values are shown in bold.

Reboot Host Immediately:

Intel VT for directed IO:	Enabled
Intel VTD ATS support:	Enabled
LOM Port 1 OptionRom:	Enabled
Pcie Slot 1 OptionRom:	Disabled
MLOM OptionRom:	Enabled
Front NVME 1 OptionRom:	Enabled
MRAID Link Speed:	Auto
PCIe Slot 1 Link Speed:	Auto
Front NVME 1 Link Speed:	Auto
VGA Priority:	Onboard
P-SATA OptionROM:	LSI SW RAID
USB Port Rear:	Enabled
USB Port Internal:	Enabled
IPv6 PXE Support:	Disabled

Legacy USB Support:	Enabled
Intel VTD coherency support:	Disabled
All Onboard LOM Ports:	Enabled
LOM Port 2 OptionRom:	Enabled
Pcie Slot 2 OptionRom:	Disabled
MRAID OptionRom:	Enabled
Front NVME 2 OptionRom:	Enabled
MLOM Link Speed:	Auto
PCIe Slot 2 Link Speed:	Auto
Front NVME 2 Link Speed:	Auto
M.2 SATA OptionROM:	AHCI
USB Port Front:	Enabled
USB Port KVM:	Enabled
USB Port:M.2 Storage:	Enabled

Configure Cisco VIC1457 for iSCSI boot

The following configuration steps are for the Cisco VIC 1457 for iSCSI boot.



The default port-channeling between ports 0, 1, 2, and 3 must be turned off before the four individual ports can be configured. If port channeling is not turned off, only two ports appear for the VIC 1457. Complete the following steps to enable the port channel on the CIMC:

1. Under the networking tab, click the Adapter Card MLOM.
2. Under the General tab, uncheck the port channel.
3. Save the changes and reboot the CIMC.

Create iSCSI vNICs

To create iSCSI vNICs, complete the following steps:

1. Under the networking tab, click Adapter Card MLOM.
2. Click Add vNIC to create a vNIC.
3. In the Add vNIC section, enter the following settings:
 - Name: eth1
 - CDN Name: iSCSI-vNIC-A
 - MTU: 9000
 - Default VLAN: <<var_iscsi_vlan_a>>
 - VLAN Mode: TRUNK
 - Enable PXE boot: Check
4. Click Add vNIC and then click OK.
5. Repeat the process to add a second vNIC:
 - Name the vNIC eth3.
 - CDN Name: iSCSI-vNIC-B
 - Enter <<var_iscsi_vlan_b>> as the VLAN.
 - Set the uplink port to 3.

▼ General

Name:

CDN:

MTU: (1500 - 9000)

Uplink Port:

MAC Address: Auto

Class of Service: (0 - 6)

Trust Host CoS:

PCI Order: (0 - 7)

Default VLAN: None
 

6. Select the vNIC eth1 on the left.

General External Ethernet Interfaces vNICs vHBAs

▼ vNICs

- eth0
- eth1**
- eth2
- eth3

► vNIC Properties

▼ iSCSI Boot Properties

► General

▼ Initiator

Name: (0 - 222) chars

IP Address:

Subnet Mask:

Gateway:

Primary DNS:

► Primary Target

► Secondary Target

Unconfigure iSCSI Boot

7. Under iSCSI Boot Properties, enter the initiator details:

- Name: <>var_ucsA_initiator_name_a>>
- IP address: <>var_esxi_hostA_iscsiA_ip>>
- Subnet mask: <>var_esxi_hostA_iscsiA_mask>>
- Gateway: <>var_esxi_hostA_iscsiA_gateway>>

vNIC Properties

iSCSI Boot Properties

Initiator

Name:	iqn.1992-01.com.cisco.ucsA-01	(0 - 222) chars
IP Address:	172.21.183.110	
Subnet Mask:	255.255.255.0	
Gateway:	172.21.183.1	
Primary DNS:		

Initiator Priority: primary

Secondary DNS:

TCP Timeout: 15

CHAP Name:

CHAP Secret:

Primary Target

Name:	iqn.1992-08.com.netapp.sn.e42fa6b2d2i	(0 - 222) chars
IP Address:	172.21.183.105	
TCP Port:	3260	

Boot LUN: 0

CHAP Name:

CHAP Secret:

Secondary Target

Name:	iqn.1992-08.com.netapp.sn.e42fa6b2d2i	(0 - 222) chars
IP Address:	172.21.183.106	
TCP Port:	3260	

Boot LUN: 0

CHAP Name:

CHAP Secret:

Unconfigure iSCSI Boot

8. Enter the primary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif01a
- Boot LUN: 0

9. Enter the secondary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif02a
- Boot LUN: 0



You can obtain the storage IQN number by running the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step. In addition, the IQN names for initiators must be unique for each server and for the iSCSI vNIC.

10. Click Save Changes.

11. Select the vNIC eth3 and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.

12. Repeat the process to configure eth3.

13. Enter the initiator details:

- Name: <<var_ucs_a_initiator_name_b>>
- IP address: <<var_esxi_hostb_iscsib_ip>>
- Subnet mask: <<var_esxi_hostb_iscsib_mask>>
- Gateway: <<var_esxi_hostb_iscsib_gateway>>

14. Enter the primary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif01b
- Boot LUN: 0

15. Enter the secondary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif02b
- Boot LUN: 0



You can obtain the storage IQN number by using the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step.

16. Click Save Changes.

17. Repeat this process to configure iSCSI boot for Cisco UCS server B.

Configure vNICs for ESXi

To configure vNICs for ESXi, complete the following steps:

1. From the CIMC interface browser window, click Inventory and then click Cisco VIC adapters on the right pane.

2. Under Networking > Adapter Card MLOM, select vNICs tab and then select the vNICs underneath.
3. Select eth0 and click Properties.
4. Set the MTU to 9000. Click Save Changes.
5. Set the VLAN to native VLAN 2.

vNIC Properties - General

- Name: eth0
- CDN: VIC-MLOM-eth0
- MTU: 9000 (1500 - 9000)
- Uplink Port: 0
- MAC Address: F8:0F:6F:89:26:CE
- Class of Service: 0 (0 - 6)
- Trust Host CoS:
- PCI Order: 0 (0 - 7)
- Default VLAN: 2

6. Repeat steps 3 and 4 for eth1, verifying that the uplink port is set to 1 for eth1.

Host Ethernet Interfaces															
		Add vNIC	Clone vNIC	Delete vNICs	Selected 0 / Total 4										
		Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
vNICs		eth0	VIC-MLOM-eth0	F8:0F:6F:89:26:CE	9000	0	0	0	2	TRUNK	disabled	enabled	N/A	N/A	N/A
		eth1	VIC-ISCS-eth1	F8:0F:6F:89:26:CF	9000	0	1	0	3439	TRUNK	enabled	enabled	N/A	N/A	N/A
		eth2	VIC-MLOM-eth2	F8:0F:6F:89:26:D0	9000	0	2	0	2	TRUNK	disabled	enabled	N/A	N/A	N/A
		eth3	VIC-ISCS-eth3	F8:0F:6F:89:26:D1	9000	0	3	0	3440	TRUNK	enabled	enabled	N/A	N/A	N/A



This procedure must be repeated for each initial Cisco UCS server node and each additional Cisco UCS server node added to the environment.

Next: NetApp AFF storage deployment procedure (part 2).

NetApp AFF storage deployment procedure (part 2)

Set up ONTAP SAN boot storage

Create iSCSI igroups



You need the iSCSI initiator IQNs from the server configuration for this step.

To create igroups, run the following commands from the cluster management node SSH connection. To view the three igroups created in this step, run the `igroup show` command.

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-A -protocol iscsi  
-ostype vmware -initiator <<var_vm_host_infra_a_iSCSI-  
A_vNIC_IQN>>, <<var_vm_host_infra_a_iSCSI-B_vNIC_IQN>>  
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-B -protocol iscsi  
-ostype vmware -initiator <<var_vm_host_infra_b_iSCSI-  
A_vNIC_IQN>>, <<var_vm_host_infra_b_iSCSI-B_vNIC_IQN>>
```



This step must be completed when adding additional Cisco UCS C-Series servers.

Map boot LUNs to igroups

To map boot LUNs to igroups, run the following commands from the cluster management SSH connection:

```
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -igroup  
VM-Host-Infra-A -lun-id 0  
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -igroup  
VM-Host-Infra-B -lun-id 0
```



This step must be completed when adding additional Cisco UCS C-Series servers.

[Next: VMware vSphere 6.7U2 deployment procedure.](#)

VMware vSphere 6.7U2 deployment procedure

This section provides detailed procedures for installing VMware ESXi 6.7U2 in a FlexPod Express configuration. The deployment procedures that follow are customized to include the environment variables described in previous sections.

Multiple methods exist for installing VMware ESXi in such an environment. This procedure uses the virtual KVM console and virtual media features of the CIMC interface for Cisco UCS C-Series servers to map remote installation media to each individual server.



This procedure must be completed for Cisco UCS server A and Cisco UCS server B.



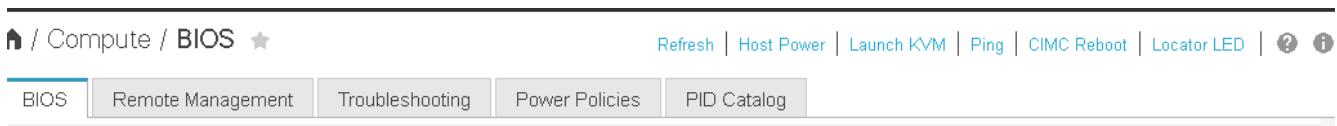
This procedure must be completed for any additional nodes added to the cluster.

Log in to CIMC interface for Cisco UCS C-Series standalone servers

The following steps detail the method for logging in to the CIMC interface for Cisco UCS C-Series standalone servers. You must log in to the CIMC interface to run the virtual KVM, which enables the administrator to begin installation of the operating system through remote media.

All hosts

1. Navigate to a web browser and enter the IP address for the CIMC interface for the Cisco UCS C-Series. This step launches the CIMC GUI application.
2. Log in to the CIMC UI using the admin user name and credentials.
3. In the main menu, select the Server tab.
4. Click Launch KVM Console.



5. From the virtual KVM console, select the Virtual Media tab.
6. Select Map CD/DVD.



You might first need to click Activate Virtual Devices. Select Accept This Session if prompted.

7. Browse to the VMware ESXi 6.7U2 installer ISO image file and click Open. Click Map Device.
8. Select the Power menu and choose Power Cycle System (Cold Boot). Click Yes.

Install VMware ESXi

The following steps describe how to install VMware ESXi on each host.

Download ESXI 6.7U2 Cisco custom image

1. Navigate to the [VMware vSphere download page](#) for custom ISOs.
2. Click Go to Downloads next to the Cisco Custom Image for the ESXi 6.7U2 Install CD.
3. Download the Cisco Custom Image for the ESXi 6.7U2 Install CD (ISO).
4. When the system boots, the machine detects the presence of the VMware ESXi installation media.
5. Select the VMware ESXi installer from the menu that appears. The installer loads, which can take several minutes.
6. After the installer has finished loading, press Enter to continue with the installation.
7. After reading the end-user license agreement, accept it and continue with the installation by pressing F11.
8. Select the NetApp LUN that was previously set up as the installation disk for ESXi, and press Enter to continue with the installation.



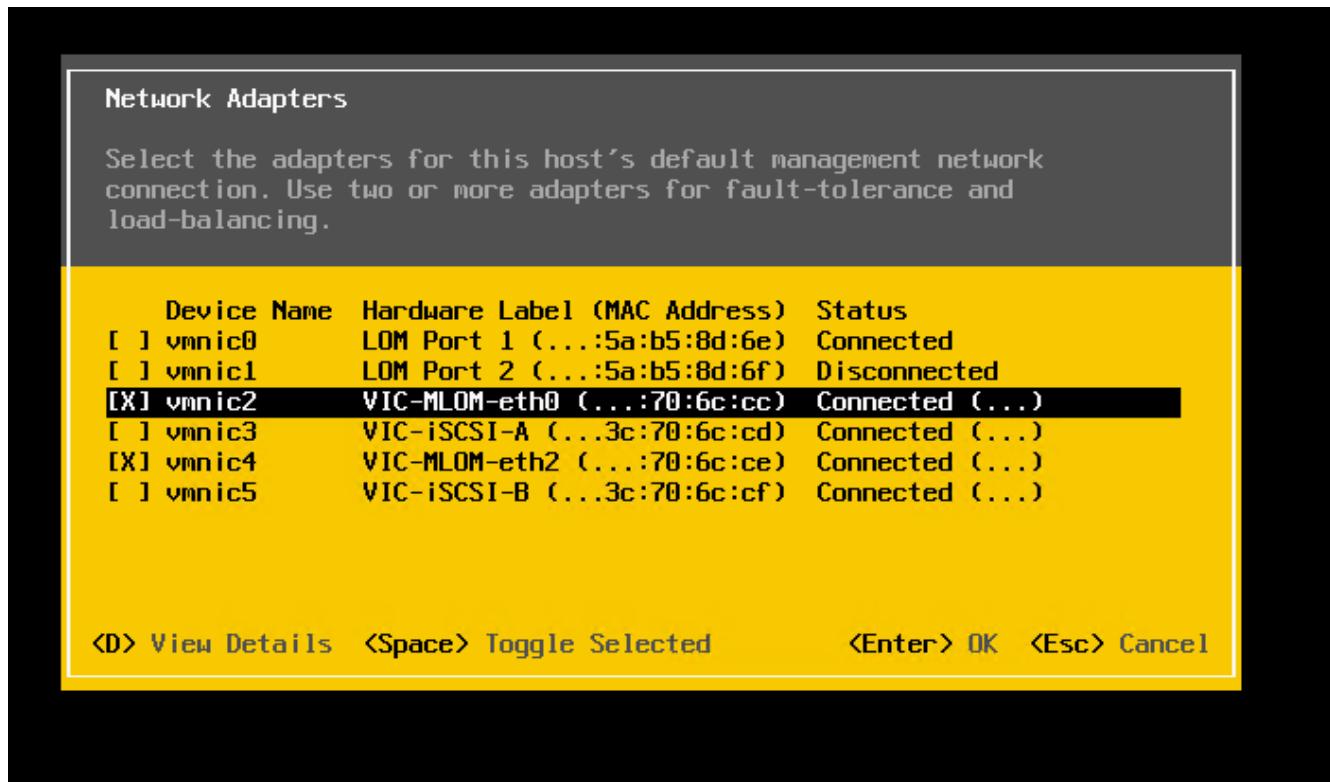
9. Select the appropriate keyboard layout and press Enter.
10. Enter and confirm the root password and press Enter.
11. The installer warns you that existing partitions are removed on the volume. Continue with the installation by pressing F11. The server reboots after the installation of ESXi.

Set up VMware ESXi host management networking

The following steps describe how to add the management network for each VMware ESXi host.

All hosts

1. After the server has finished rebooting, enter the option to customize the system by pressing F2.
2. Log in with root as the login name and the root password previously entered during the installation process.
3. Select the Configure Management Network option.
4. Select Network Adapters and press Enter.
5. Select the desired ports for vSwitch0. Press Enter.
6. Select the ports that correspond to eth0 and eth1 in CIMC.



7. Select VLAN (optional) and press Enter.
8. Enter the VLAN ID <<mgmt_vlan_id>>. Press Enter.
9. From the Configure Management Network menu, select IPv4 Configuration to configure the IP address of the management interface. Press Enter.
10. Use the arrow keys to highlight Set Static IPv4 Address and use the space bar to select this option.
11. Enter the IP address for managing the VMware ESXi host <<esxi_host_mgmt_ip>>.
12. Enter the subnet mask for the VMware ESXi host <<esxi_host_mgmt_netmask>>.
13. Enter the default gateway for the VMware ESXi host <<esxi_host_mgmt_gateway>>.
14. Press Enter to accept the changes to the IP configuration.
15. Enter the IPv6 configuration menu.
16. Use the space bar to disable IPv6 by unselecting the Enable IPv6 (restart required) option. Press Enter.
17. Enter the menu to configure the DNS settings.
18. Because the IP address is assigned manually, the DNS information must also be entered manually.
19. Enter the primary DNS server's IP address <<nameserver_ip>>.
20. (Optional) Enter the secondary DNS server's IP address.
21. Enter the FQDN for the VMware ESXi host name: <<esxi_host_fqdn>>.
22. Press Enter to accept the changes to the DNS configuration.
23. Exit the Configure Management Network submenu by pressing Esc.
24. Press Y to confirm the changes and reboot the server.
25. Select Troubleshooting Options, and then Enable ESXi Shell and SSH.



These troubleshooting options can be disabled after the validation pursuant to the customer's security policy.

26. Press Esc twice to return to the main console screen.
27. Click Alt-F1 from the CIMC Macros > Static Macros > Alt-F drop-down menu at the top of the screen.
28. Log in with the proper credentials for the ESXi host.
29. At the prompt, enter the following list of esxcli commands sequentially to enable network connectivity.

```
esxcli network vswitch standard policy failover set -v vSwitch0 -a
vmnic2,vmnic4 -l iphash
```

Configure ESXi host

Use the information in the following table to configure each ESXi host.

Detail	Detail value
ESXi host name	<<esxi_host_fqdn>>
ESXi host management IP	<<esxi_host_mgmt_ip>>
ESXi host management mask	<<esxi_host_mgmt_netmask>>
ESXi host management gateway	<<esxi_host_mgmt_gateway>>
ESXi host NFS IP	<<esxi_host_NFS_ip>>
ESXi host NFS mask	<<esxi_host_NFS_netmask>>
ESXi host NFS gateway	<<esxi_host_NFS_gateway>>
ESXi host vMotion IP	<<esxi_host_vMotion_ip>>
ESXi host vMotion mask	<<esxi_host_vMotion_netmask>>
ESXi host vMotion gateway	<<esxi_host_vMotion_gateway>>
ESXi host iSCSI-A IP	<<esxi_host_iSCSI-A_ip>>
ESXi host iSCSI-A mask	<<esxi_host_iSCSI-A_netmask>>
ESXi host iSCSI-A gateway	<<esxi_host_iSCSI-A_gateway>>
ESXi host iSCSI-B IP	<<esxi_host_iSCSI-B_ip>>
ESXi host iSCSI-B mask	<<esxi_host_iSCSI-B_netmask>>
ESXi host iSCSI-B gateway	<<esxi_host_SCSI-B_gateway>>

Log in to the ESXi host

To log in to the ESXi host, complete the following steps:

1. Open the host's management IP address in a web browser.
2. Log in to the ESXi host using the root account and the password you specified during the install process.

3. Read the statement about the VMware Customer Experience Improvement Program. After selecting the proper response, click OK.

Configure iSCSI boot

To configure iSCSI boot, complete the following steps:

1. Select Networking on the left.
2. On the right, select the Virtual Switches tab.



3. Click iScsiBootvSwitch.
4. Select Edit settings.
5. Change the MTU to 9000 and click Save.
6. Rename the iSCSIBootPG port to iSCSIBootPG-A.



Vmnic3 and vmnic5 are used for iSCSI boot in this configuration. If you have additional NICs in your ESXi host, you might have different vmnic numbers. To confirm which NICs are used for iSCSI boot, match the MAC addresses on the iSCSI vNICs in CIMC to the vmnics in ESXi.

7. In the center pane, select the VMkernel NICs tab.
8. Select Add VMkernel NIC.
 - a. Specify a new port group name of iScsiBootPG-B.
 - b. Select iScsiBootvSwitch for the virtual switch.
 - c. Enter <<iscsib_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_iscsib_ip>> for Address.

h. Enter <<var_hosta_iscsib_mask>> for Subnet Mask.

i. Click Create.



Set the MTU to 9000 on iScsiBootPG-A.

9. To set the failover, complete the following steps:

- Click Edit Settings on iSCSIBootPG-A > Teaming and Failover > Failover Order > Vmnic3. Vmnic3 should be active and vmnic5 should be unused.
- Click Edit Settings on iSCSIBootPG-B > Teaming and Failover > Failover order > Vmnic5. Vmnic5 should be active and vmnic3 should be unused.

iScsiBootPG-A - Edit Settings

Properties

Security

Load balancing

Traffic shaping

Network failure detection

Teaming and failover

Notify switches

Fallback

Failover order

Override



Active adapters

vmnic3

Standby adapters

Unused adapters

vmnic5

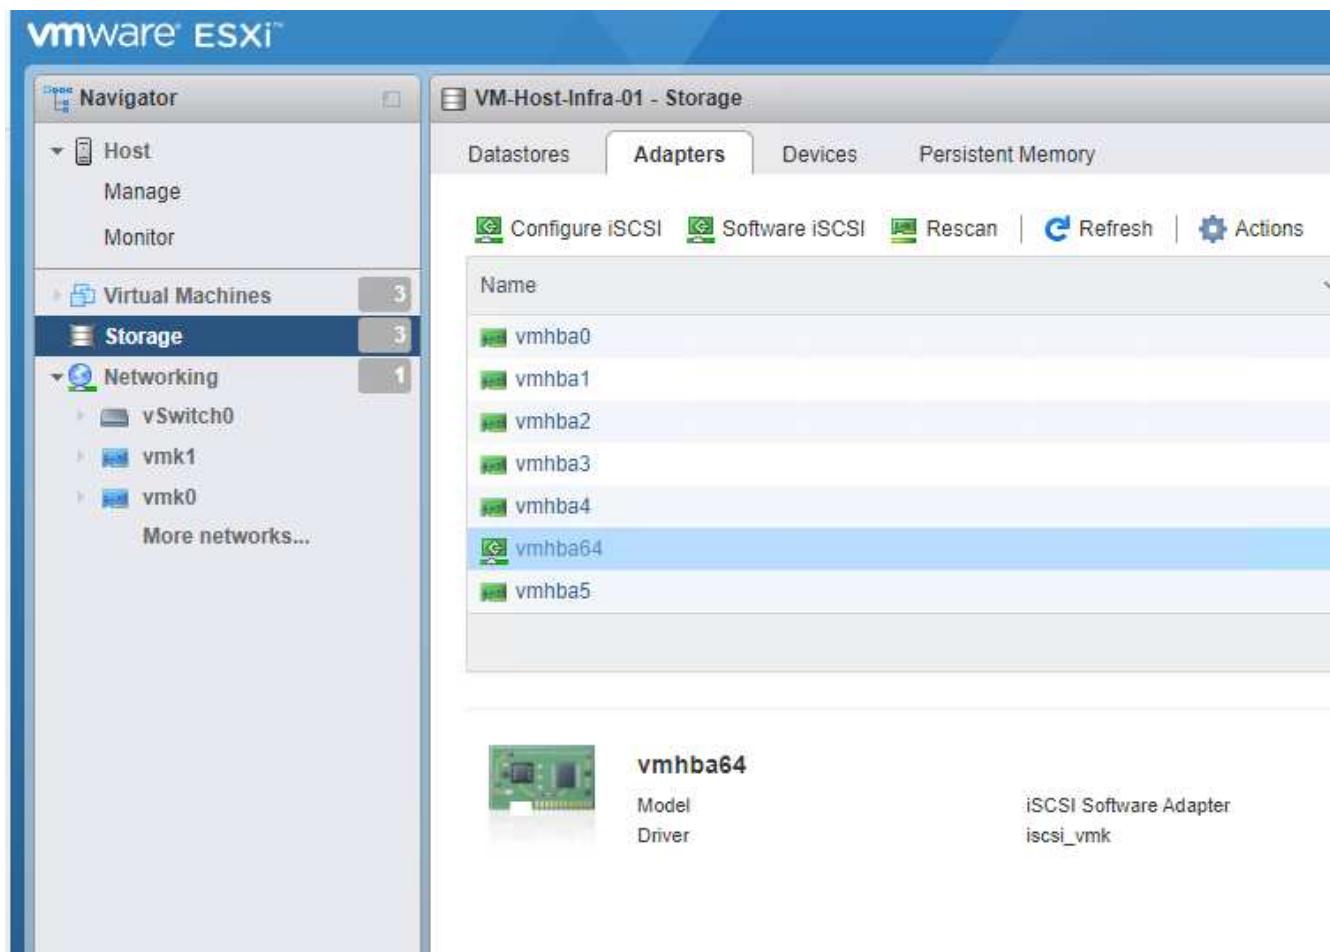
Select active and standby adapters

Configure iSCSI multipathing

To set up iSCSI multipathing on the ESXi hosts, complete the following steps:

- Select Storage in the left navigation pane. Click Adapters.

2. Select the iSCSI software adapter and click Configure iSCSI.



3. Under Dynamic Targets, click Add Dynamic Target.

Configure iSCSI - vmhba64

iSCSI enabled	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled															
▶ Name & alias	iqn.1992-01.com.cisco:ucsA-01															
▶ CHAP authentication	Do not use CHAP															
▶ Mutual CHAP authentication	Do not use CHAP															
▶ Advanced settings	Click to expand															
Network port bindings	No port bindings															
Static targets	<div style="display: flex; justify-content: space-between;"> Add static target Remove static target Edit settings <input style="border: 1px solid #ccc; padding: 2px 10px; border-radius: 4px; margin-left: 10px;" type="text"/> </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; padding: 2px;">Target</th> <th style="text-align: left; padding: 2px;">Address</th> <th style="text-align: left; padding: 2px;">Port</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px;">iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...</td> <td style="padding: 2px;">172.21.183.105</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...</td> <td style="padding: 2px;">172.21.184.106</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...</td> <td style="padding: 2px;">172.21.183.106</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...</td> <td style="padding: 2px;">172.21.184.105</td> <td style="padding: 2px;">3260</td> </tr> </tbody> </table>	Target	Address	Port	iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.183.105	3260	iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.184.106	3260	iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.183.106	3260	iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.184.105	3260
Target	Address	Port														
iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.183.105	3260														
iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.184.106	3260														
iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.183.106	3260														
iqn.1992-08.com.netapp:sn.e42fa6b2d2e011e9a68d00a098f...	172.21.184.105	3260														
Dynamic targets	<div style="display: flex; justify-content: space-between;"> Add dynamic target Remove dynamic target Edit settings <input style="border: 1px solid #ccc; padding: 2px 10px; border-radius: 4px; margin-left: 10px;" type="text"/> </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; padding: 2px;">Address</th> <th style="text-align: left; padding: 2px;">Port</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px;">172.21.183.105</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">172.21.184.105</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">172.21.183.106</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">172.21.184.106</td> <td style="padding: 2px;">3260</td> </tr> </tbody> </table>	Address	Port	172.21.183.105	3260	172.21.184.105	3260	172.21.183.106	3260	172.21.184.106	3260					
Address	Port															
172.21.183.105	3260															
172.21.184.105	3260															
172.21.183.106	3260															
172.21.184.106	3260															

4. Enter the IP address `iscsi_lif01a`.

- Repeat with the IP addresses `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
- Click Save Configuration.

Dynamic targets	<div style="display: flex; justify-content: space-between;"> Add dynamic target Remove dynamic target Edit settings <input style="border: 1px solid #ccc; padding: 2px 10px; border-radius: 4px; margin-left: 10px;" type="text"/> </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; padding: 2px;">Address</th> <th style="text-align: left; padding: 2px;">Port</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px;">172.21.183.105</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">172.21.184.105</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">172.21.183.106</td> <td style="padding: 2px;">3260</td> </tr> <tr> <td style="padding: 2px;">172.21.184.106</td> <td style="padding: 2px;">3260</td> </tr> </tbody> </table>	Address	Port	172.21.183.105	3260	172.21.184.105	3260	172.21.183.106	3260	172.21.184.106	3260
Address	Port										
172.21.183.105	3260										
172.21.184.105	3260										
172.21.183.106	3260										
172.21.184.106	3260										



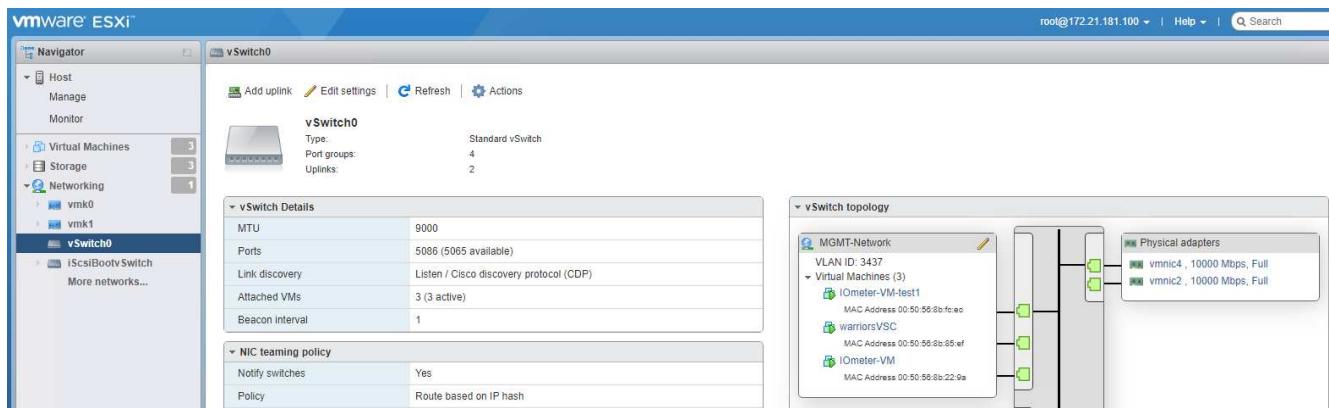
You can find the iSCSI LIF IP addresses by running the `network interface show` command on the NetApp cluster or by looking at the Network Interfaces tab in System Manager.

Configure the ESXi host

To configure ESXi boot, complete the following steps:

- In the left navigation pane, select Networking.

2. Select vSwitch0.



3. Select Edit Settings.

4. Change the MTU to 9000.

5. Expand NIC Teaming and verify that both vmnic2 and vmnic4 are set to active and NIC Teaming and Failover is set to Route Based on IP Hash.



The IP hash method of load balancing requires the underlying physical switch to be properly configured using SRC-DST-IP EtherChannel with a static (mode- on) port channel. You might experience intermittent connectivity due to possible switch misconfiguration. If so, then temporarily shut down one of the two associated uplink ports on the Cisco switch to restore communication to the ESXi management vmkernel port while troubleshooting the port-channel settings.

Configure the port groups and VMkernel NICs

To configure the port groups and VMkernel NICs, complete the following steps:

1. In the left navigation pane, select Networking.
2. Right-click the Port Groups tab.



3. Right-click VM Network and select Edit. Change the VLAN ID to <<var_vm_traffic_vlan>>.
4. Click Add Port Group.
 - a. Name the port group MGMT-Network.
 - b. Enter <<mgmt_vlan>> for the VLAN ID.
 - c. Make sure that vSwitch0 is selected.
 - d. Click save.
5. Click the VMkernel NICs tab.



6. Select Add VMkernel NIC.
 - a. Select New Port Group.
 - b. Name the port group NFS-Network.
 - c. Enter <<nfs_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_nfs_ip>> for Address.
 - h. Enter <<var_hosta_nfs_mask>> for Subnet Mask.
 - i. Click Create.
7. Repeat this process to create the vMotion VMkernel port.
8. Select Add VMkernel NIC.
 - a. Select New Port Group.
 - b. Name the port group vMotion.
 - c. Enter <<vmotion_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_vmotion_ip>> for Address.
 - h. Enter <<var_hosta_vmotion_mask>> for Subnet Mask.
 - i. Make sure that the vMotion checkbox is selected after IPv4 Settings.

 There are many ways to configure ESXi networking, including by using the VMware vSphere distributed switch if your licensing allows it. Alternative network configurations are supported in FlexPod Express if they are required to meet business requirements.

Add VMkernel NIC

Virtual switch	vSwitch0
VLAN ID	3441
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.185.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input checked="" type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication

Create **Cancel**

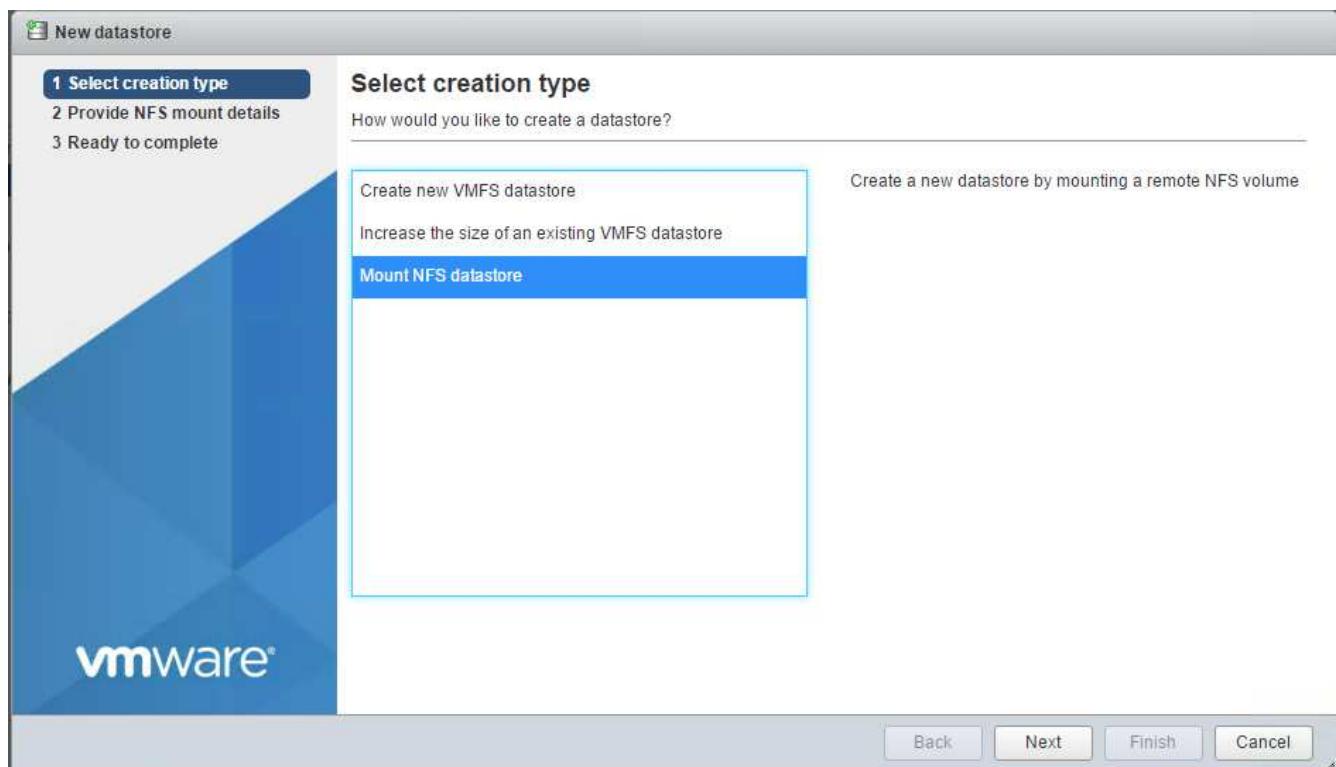
Mount the first datastores

The first datastores to be mounted are the `infra_datastore` datastore for VMs and the `infra_swap` datastore for VM swap files.

1. Click Storage in the left navigation pane, and then click New Datastore.



2. Select Mount NFS Datastore.



3. Enter the following information in the Provide NFS Mount Details page:

- Name: `infra_datastore`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_datastore`
- Make sure that NFS 3 is selected.

4. Click Finish. You can see the task completing in the Recent Tasks pane.

5. Repeat this process to mount the `infra_swap` datastore:

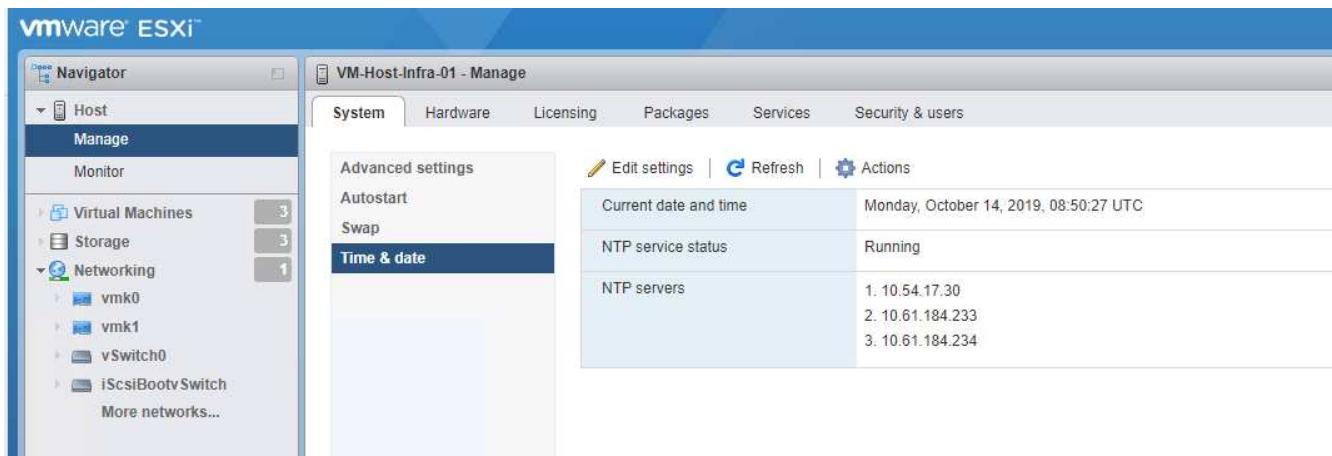
- Name: `infra_swap`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_swap`

- Make sure that NFS 3 is selected.

Configure NTP

To configure NTP for an ESXi host, complete the following steps:

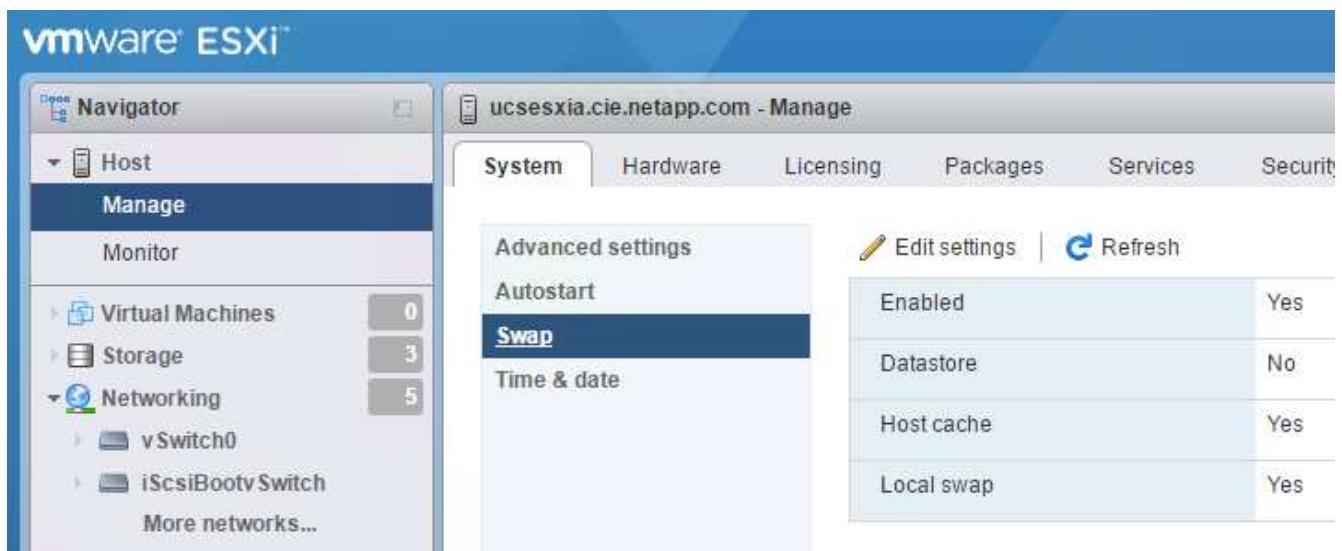
1. Click Manage in the left navigation pane. Select System in the right pane and then click Time & Date.
2. Select Use Network Time Protocol (Enable NTP Client).
3. Select Start and Stop with Host as the NTP service startup policy.
4. Enter <<var_ntp>> as the NTP server. You can set multiple NTP servers.
5. Click Save.



Move the VM swap file location

These steps provide details for moving the VM swap file location.

1. Click Manage in the left navigation pane. Select system in the right pane, then click Swap.



2. Click Edit Settings. Select `infra_swap` from the Datastore options.



3. Click Save.

[Next: VMware vCenter Server 6.7U2 installation procedure.](#)

VMware vCenter Server 6.7U2 installation procedure

This section provides detailed procedures for installing VMware vCenter Server 6.7 in a FlexPod Express configuration.

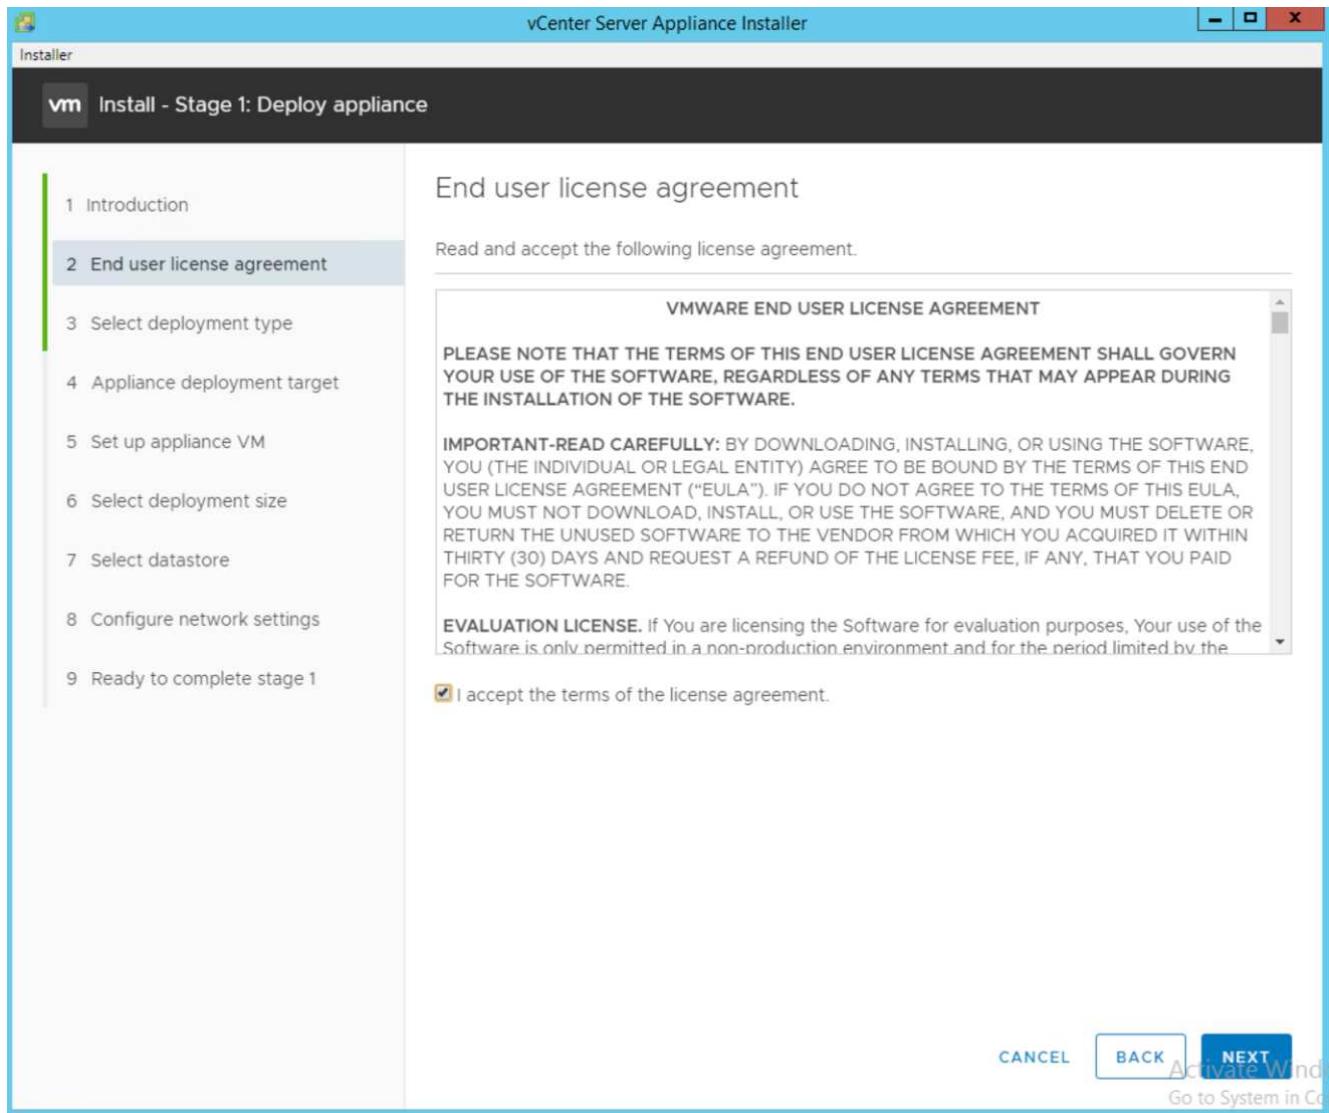


FlexPod Express uses the VMware vCenter Server Appliance (VCSA).

Download the VMware vCenter Server Appliance

To download the VMware vCenter Server Appliance (VCSA), complete the following steps:

1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.
2. Download the VCSA from the VMware site.
3. Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.
4. Mount the ISO image.
5. Navigate to the vcsa- ui-installer > win32 directory. Double-click `installer.exe`.
6. Click Install.
7. Click Next on the Introduction page.



8. Select Embedded Platform Services Controller as the deployment type.

vCenter Server Appliance Installer

Installer

vm Install - Stage 1: Deploy appliance

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

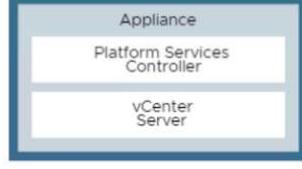
Select deployment type

Select the deployment type you want to configure on the appliance.

For more information on deployment types, refer to the vSphere 6.7 documentation.

Embedded Platform Services Controller

vCenter Server with an Embedded Platform Services Controller

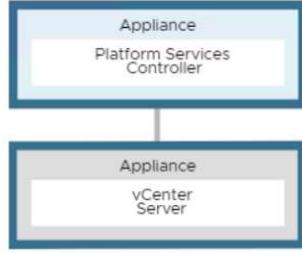


External Platform Services Controller

Deprecating soon

Platform Services Controller

vCenter Server (Requires External Platform Services Controller)

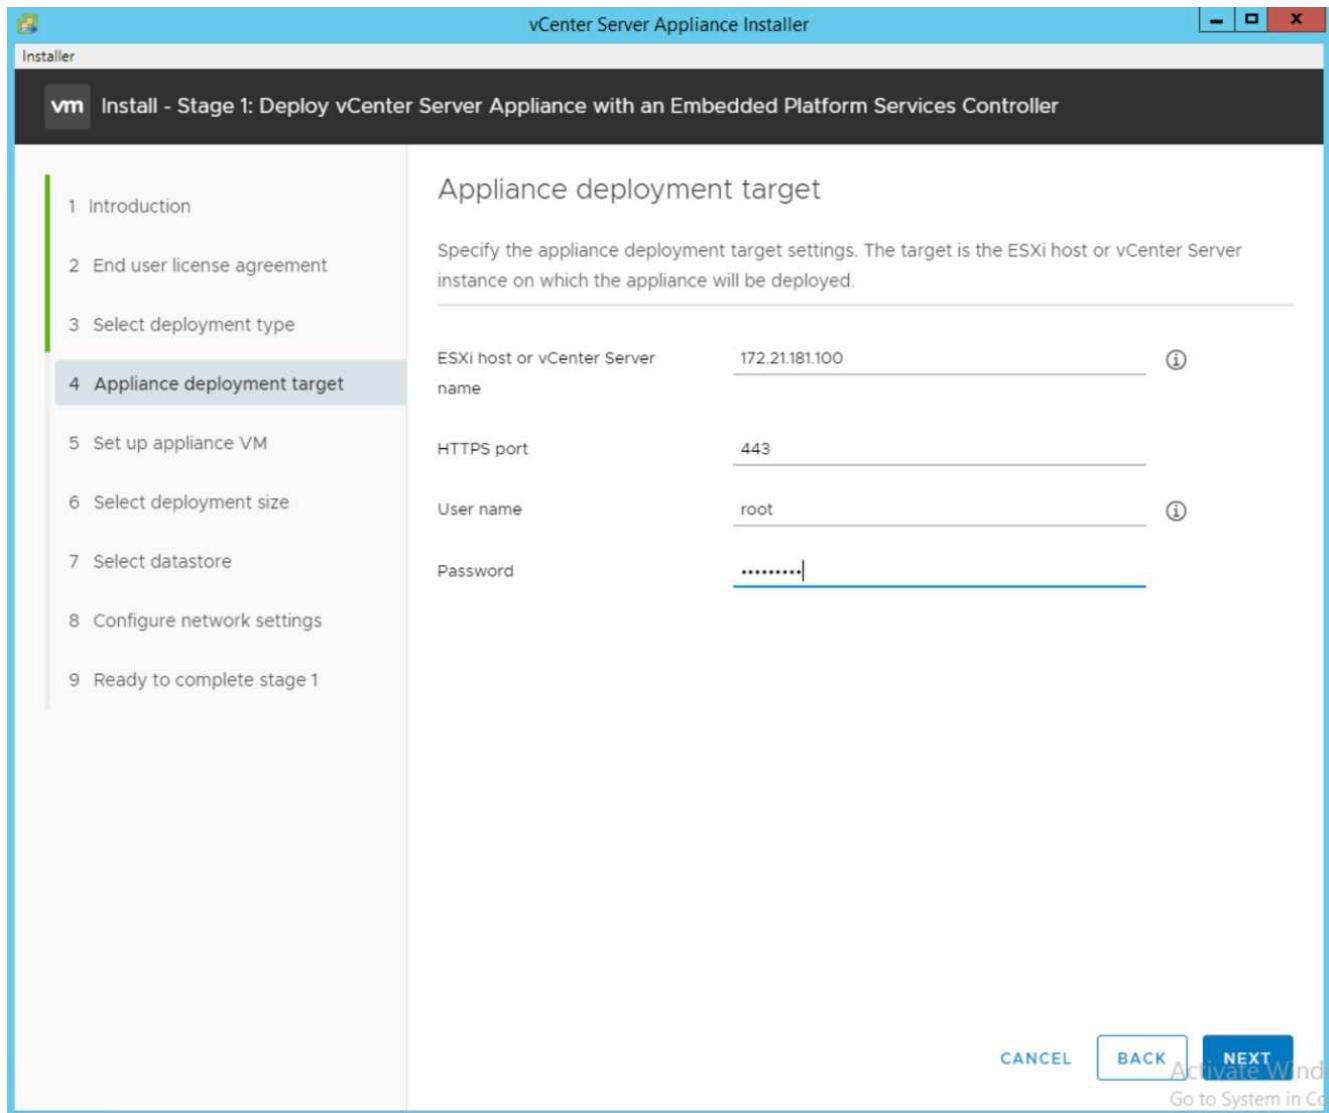


CANCEL **BACK** **NEXT** **Activate Wind**
Go to System in Co

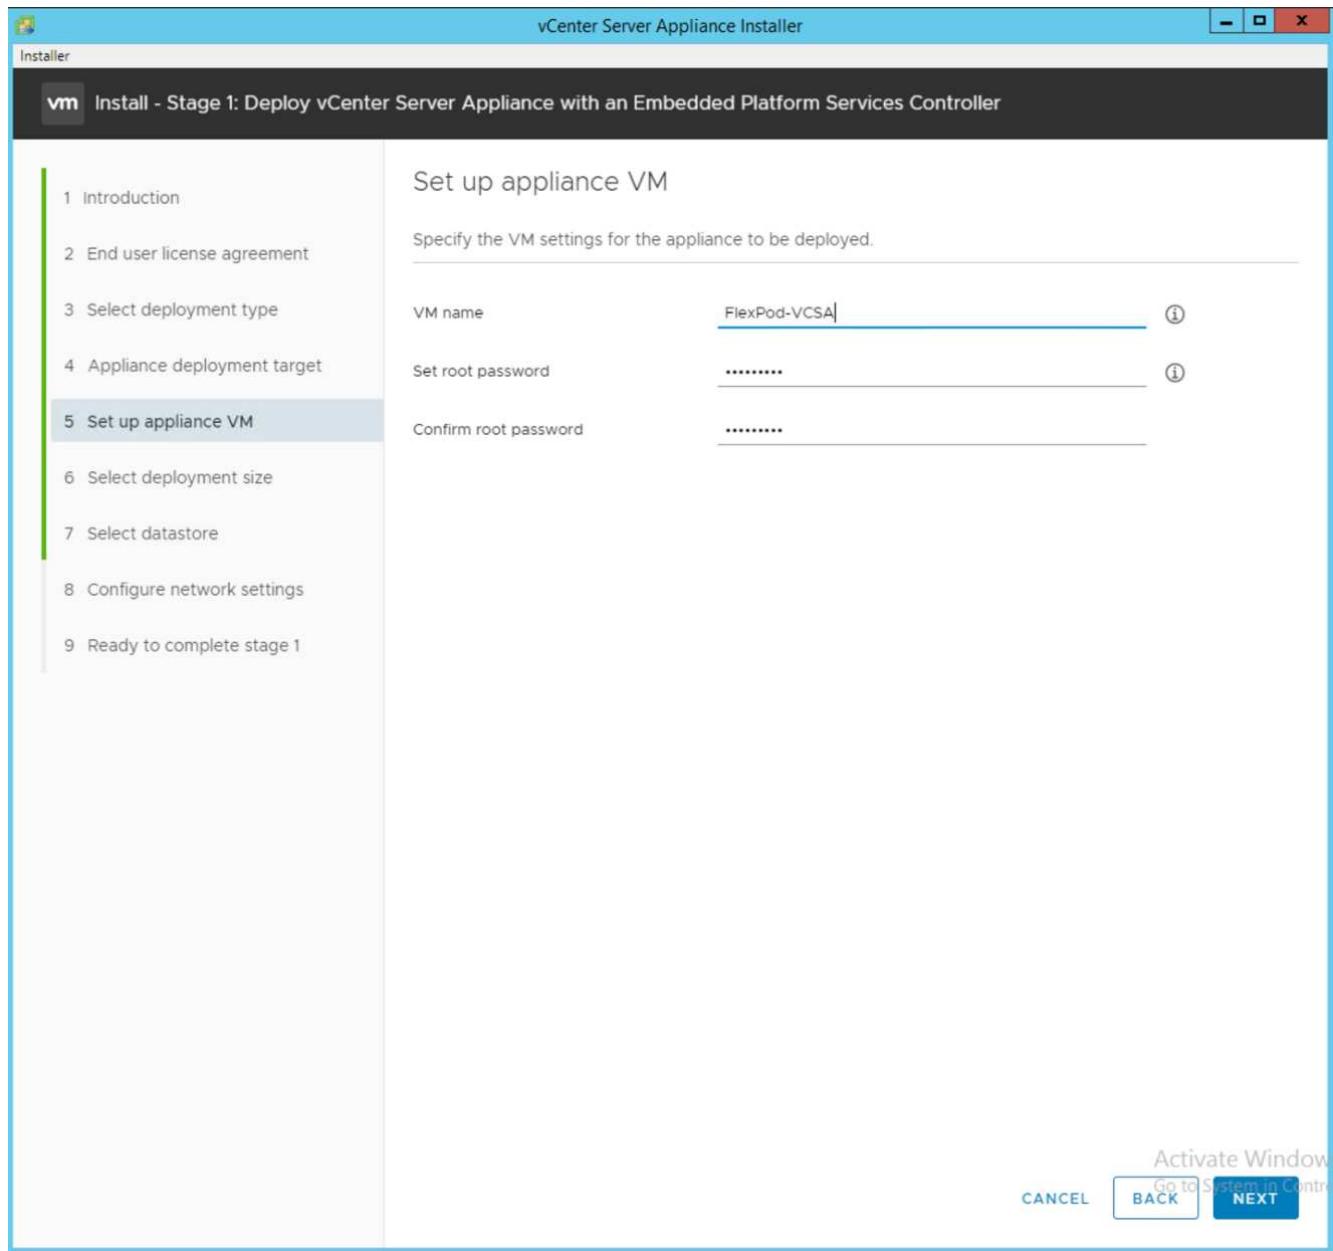


If required, the External Platform Services Controller deployment is also supported as part of the FlexPod Express solution.

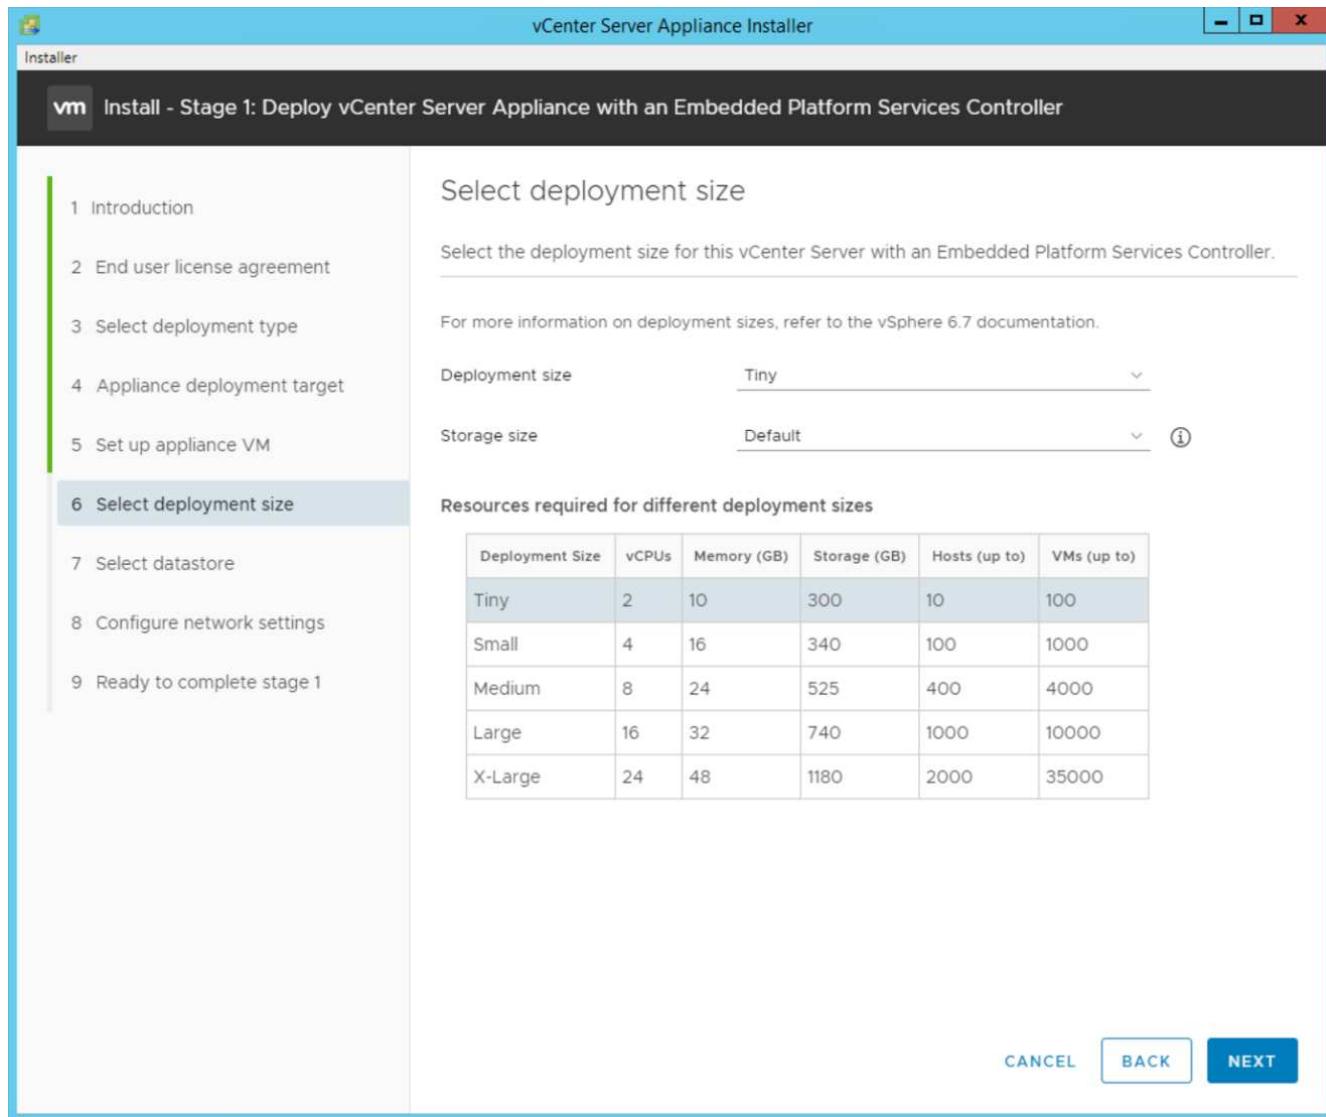
9. In the Appliance Deployment Target, enter the IP address of an ESXi host that you have deployed, the root user name, and the root password.



10. Set the appliance VM by entering VCSA as the VM name and the root password that you would like to use for the VCSA.



11. Select the deployment size that best fits your environment. Click Next.



12. Select the `infra_datastore` datastore. Click Next.
13. Enter the following information in the Configure network settings page and click Next.
 - a. Select MGMT-Network for Network.
 - b. Enter the FQDN or IP to be used for the VCSA.
 - c. Enter the IP address to be used.
 - d. Enter the subnet mask to be used.
 - e. Enter the default gateway.
 - f. Enter the DNS server.
14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.

vCenter Server Appliance Installer

Installer

VM Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Configure network settings

Configure network settings for this appliance

Network MGMT-Network ⓘ

IP version IPv4 ⓘ

IP assignment static ⓘ

FQDN FlexPod-VCSA.cle.netapp.com ⓘ

IP address 172.21.181.105

Subnet mask or prefix length 255.255.255.0 ⓘ

Default gateway 172.21.181.1

DNS servers 10.61.184.251,10.61.184.252

Common Ports

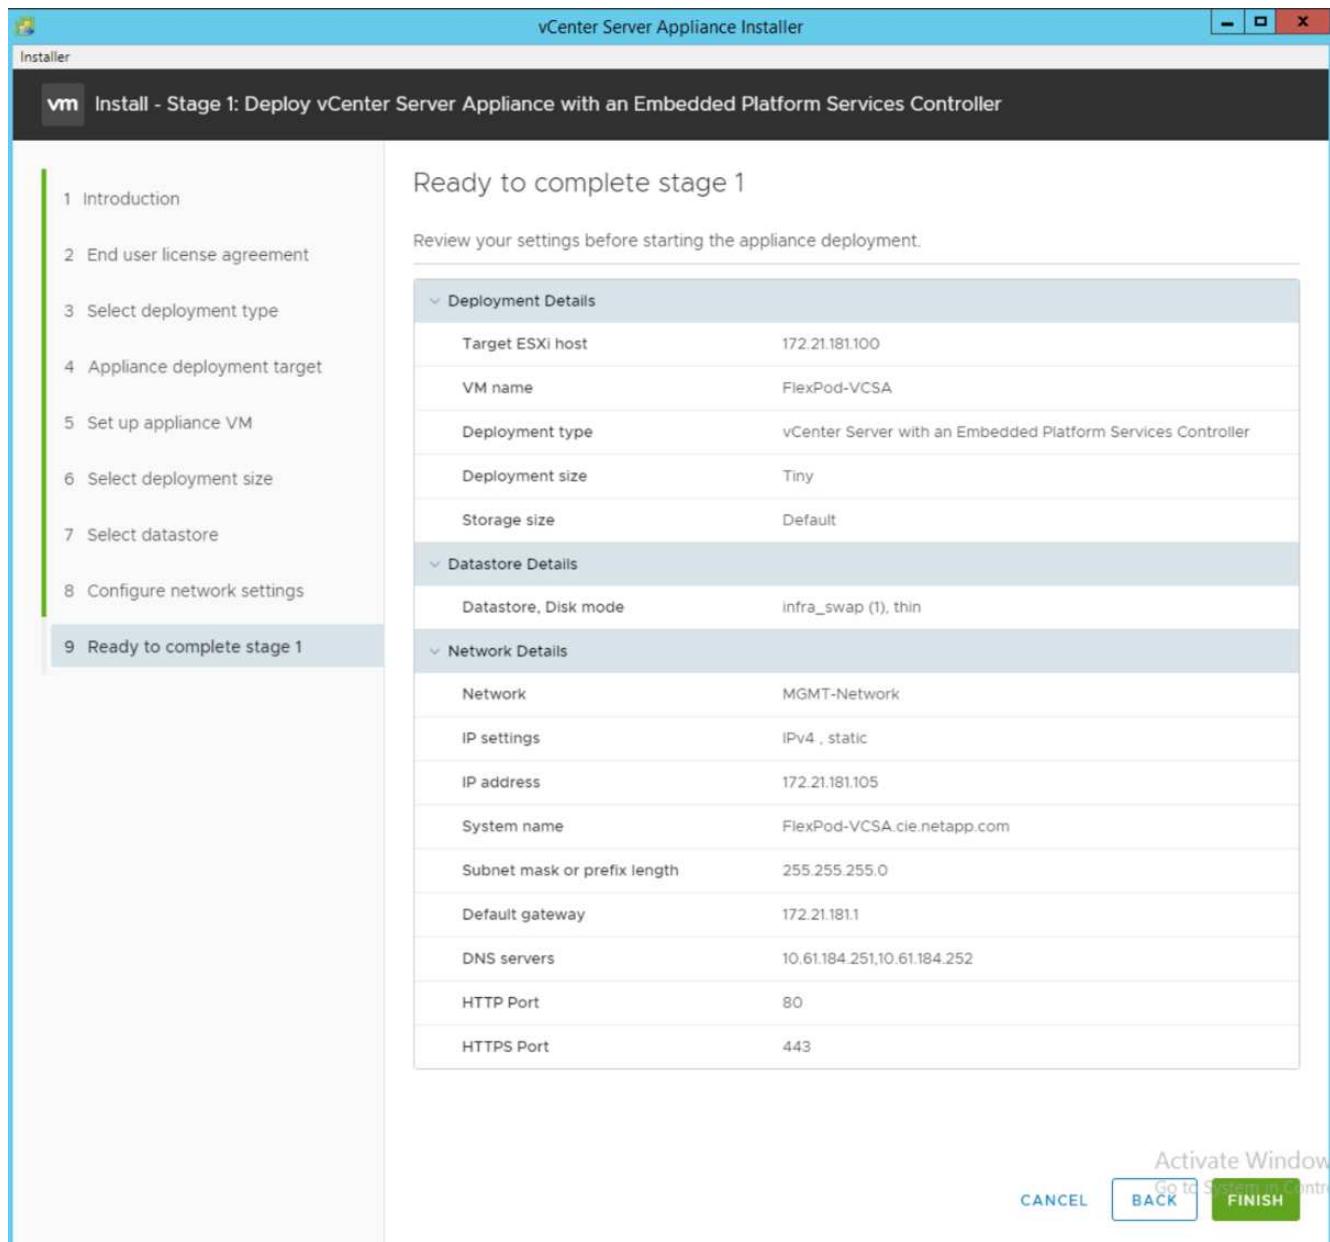
HTTP 80

HTTPS 443

Activate Window

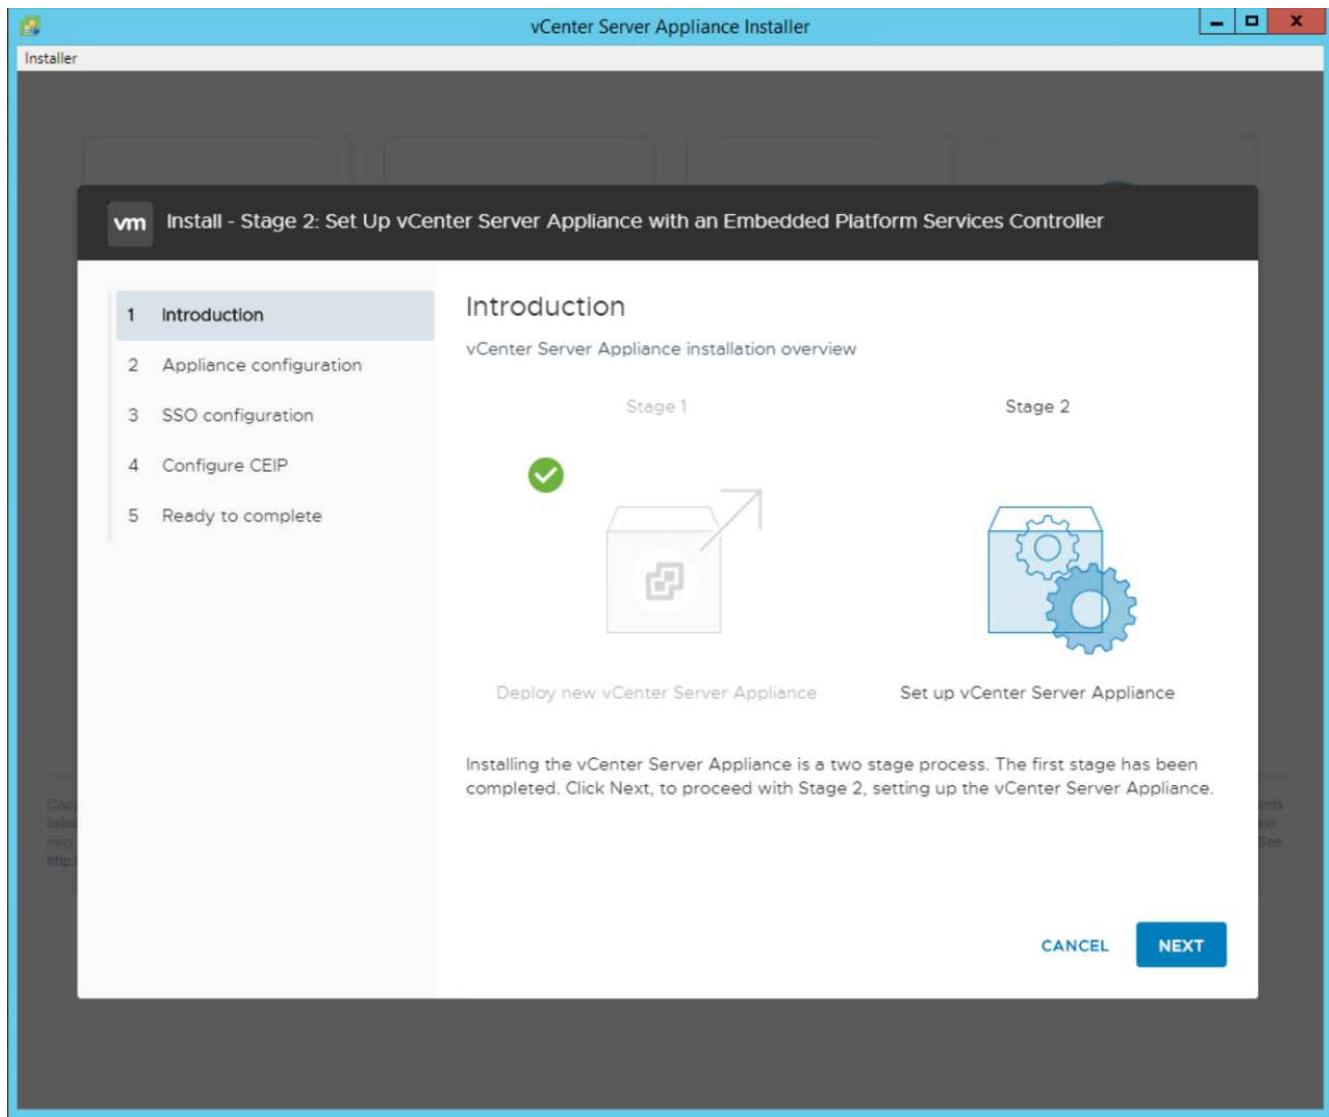
CANCEL Go to System in Control BACK NEXT

15. Review your settings on stage 1 before starting the appliance deployment.

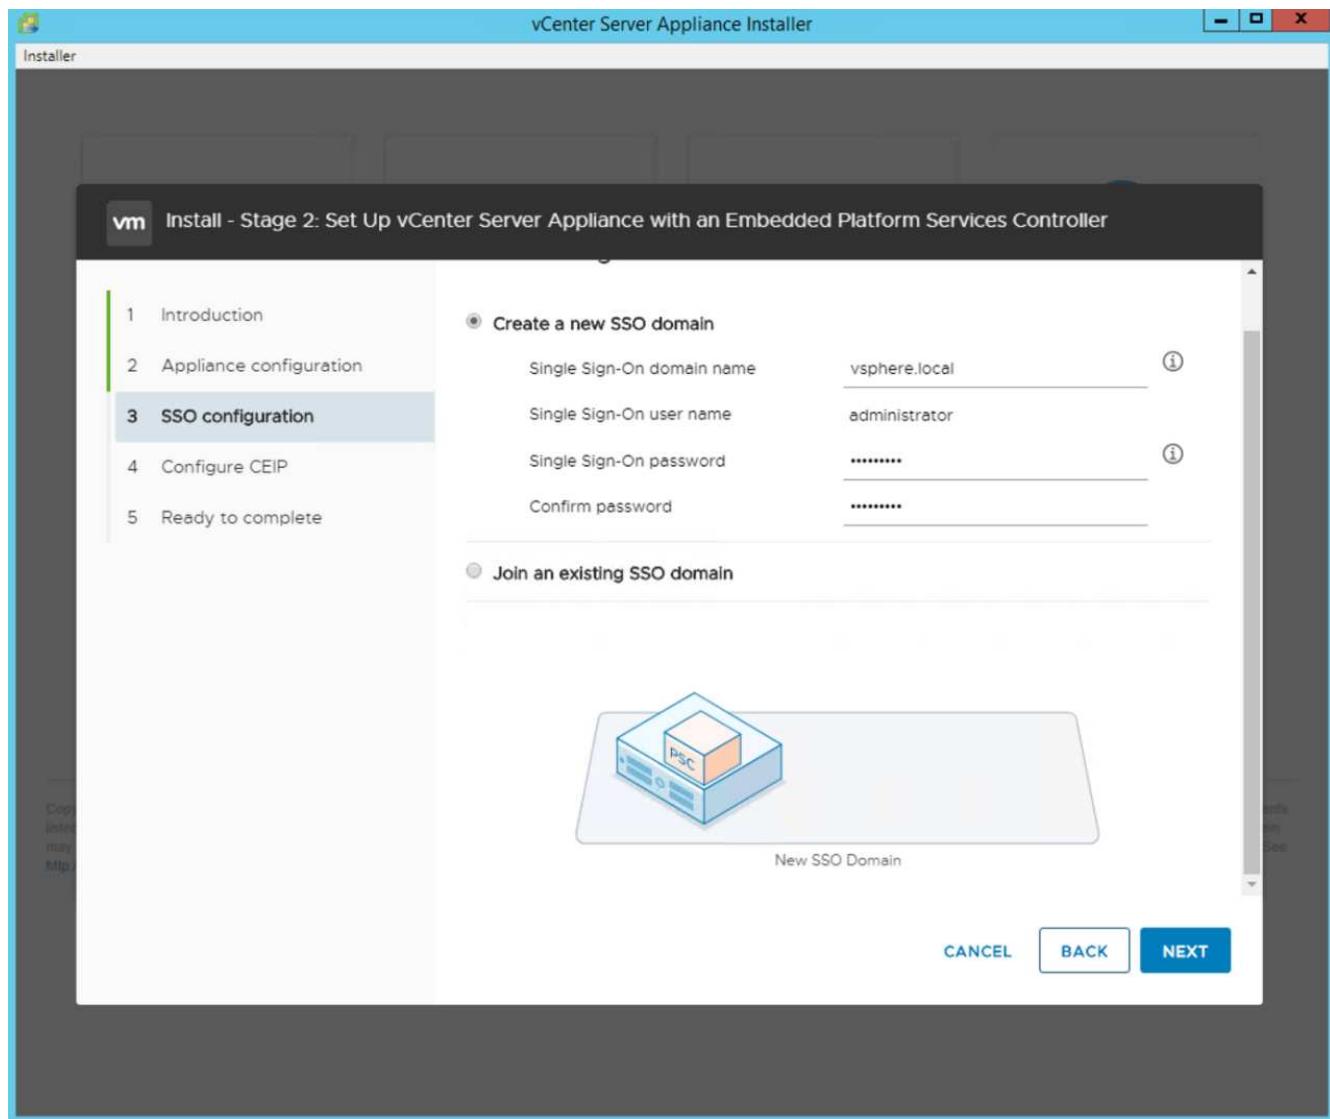


The VCSA installs now. This process takes several minutes.

16. After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.
17. On the Stage 2 Introduction page, click Next.

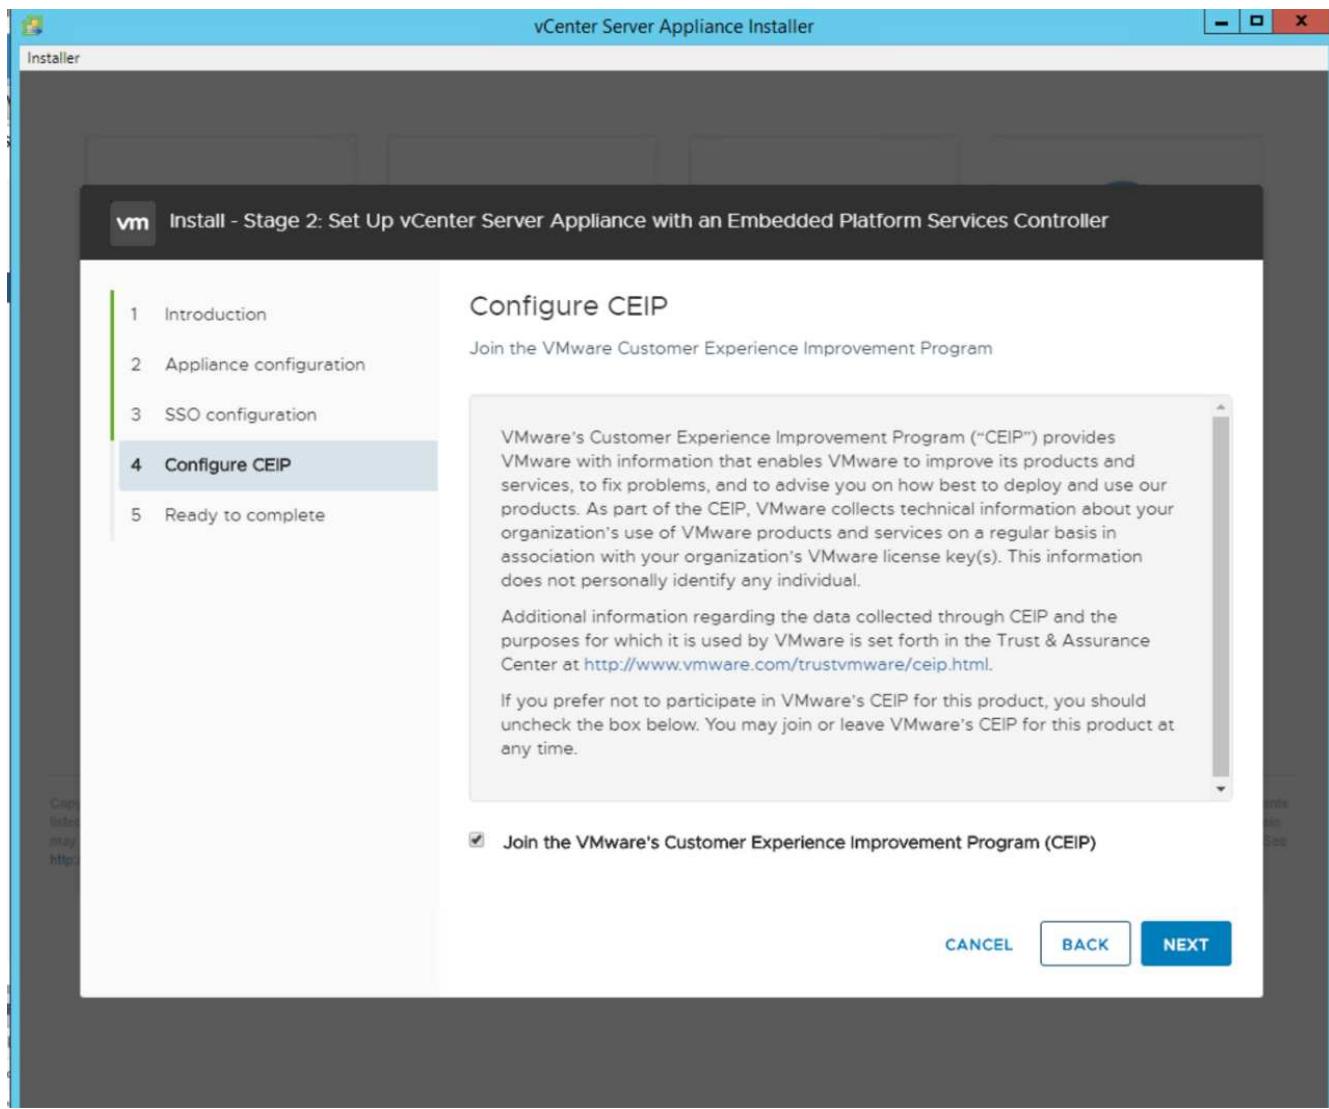


18. Enter <<var_ntp_id>> for the NTP server address. You can enter multiple NTP IP addresses.
19. If you plan to use vCenter Server high availability (HA), make sure that SSH access is enabled.
20. Configure the SSO domain name, password, and site name. Click Next.

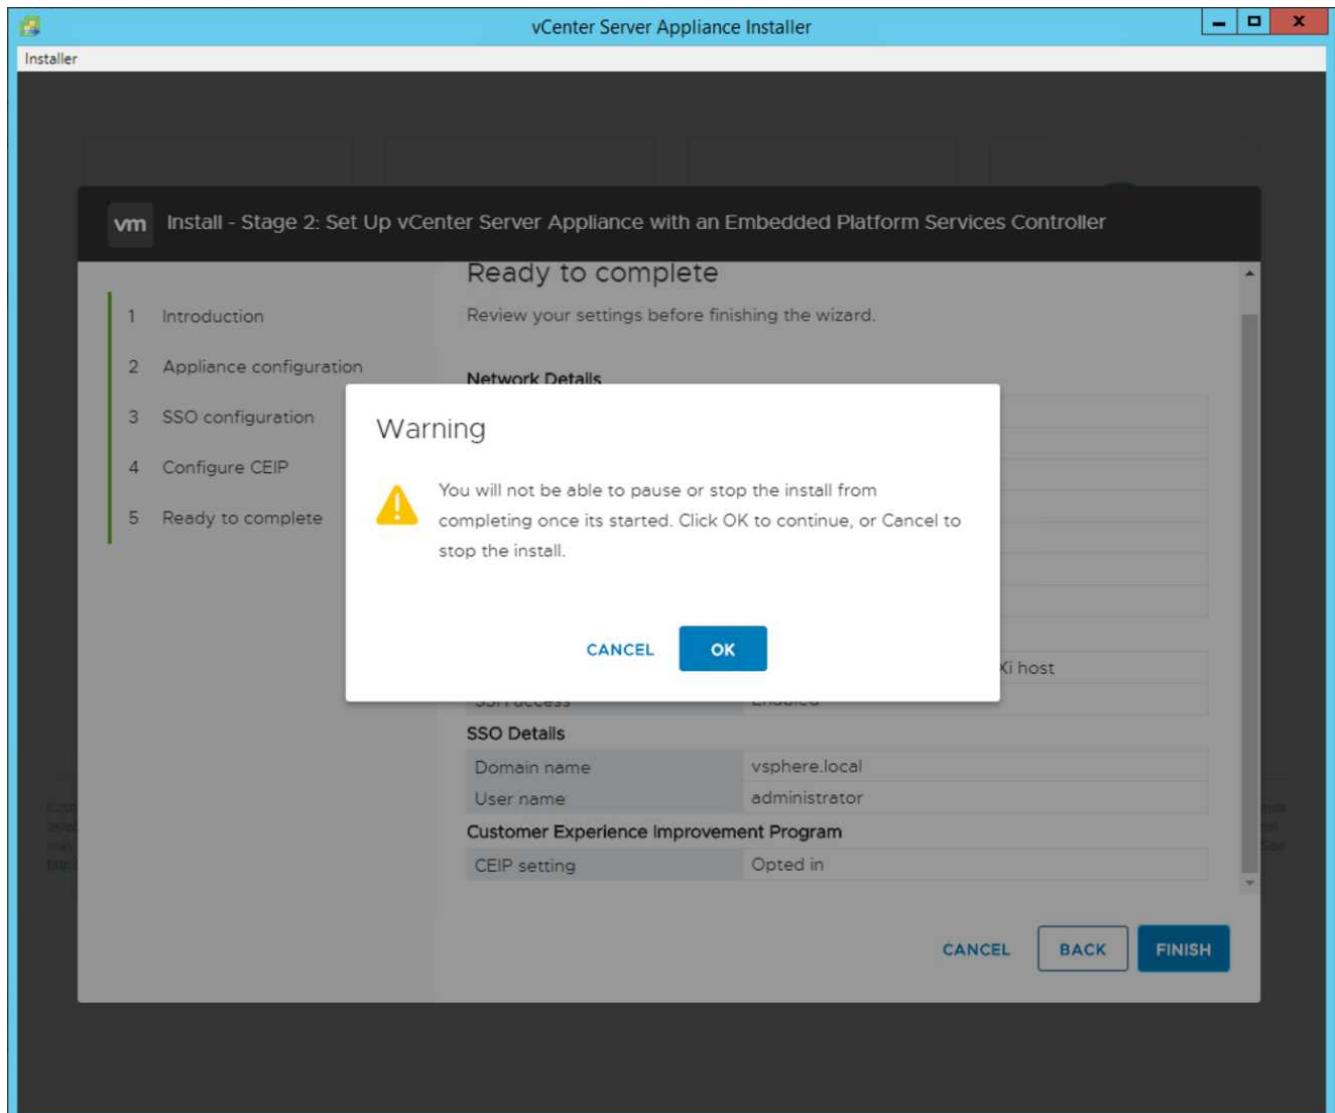


Record these values for your reference, especially if you deviate from the vsphere.local domain name.

21. Join the VMware Customer Experience Program if desired. Click Next.



22. View the summary of your settings. Click Finish or use the back button to edit settings.
23. A message appears stating that you will not be able to pause or stop the installation from completing after it has started. Click OK to continue.



The appliance setup continues. This takes several minutes.

A message appears indicating that the setup was successful.

24. The links that the installer provides to access vCenter Server are clickable.

[Next: VMware vCenter Server 6.7U2 and vSphere clustering configuration.](#)

VMware vCenter Server 6.7U2 and vSphere clustering configuration

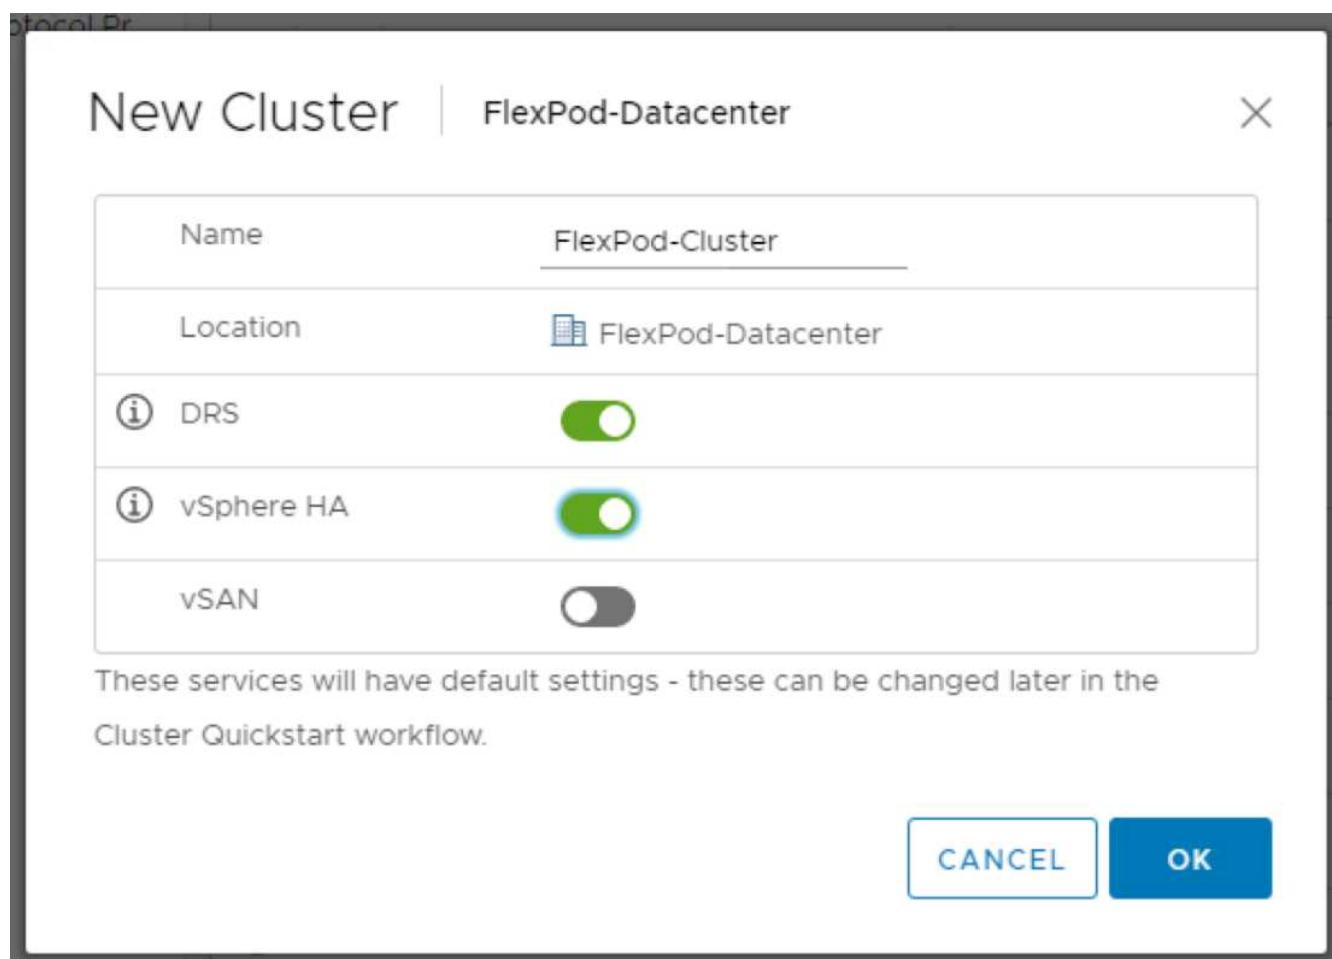
To configure VMware vCenter Server 6.7 and vSphere clustering, complete the following steps:

1. Navigate to <https://<<FQDN or IP of vCenter>>/vsphere-client/>.
2. Click Launch vSphere Client.
3. Log in with the user name administrator@vsphere.local and the SSO password you entered during the VCSA setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center and click OK.

Create a vSphere cluster

To create a vSphere cluster, complete the following steps:

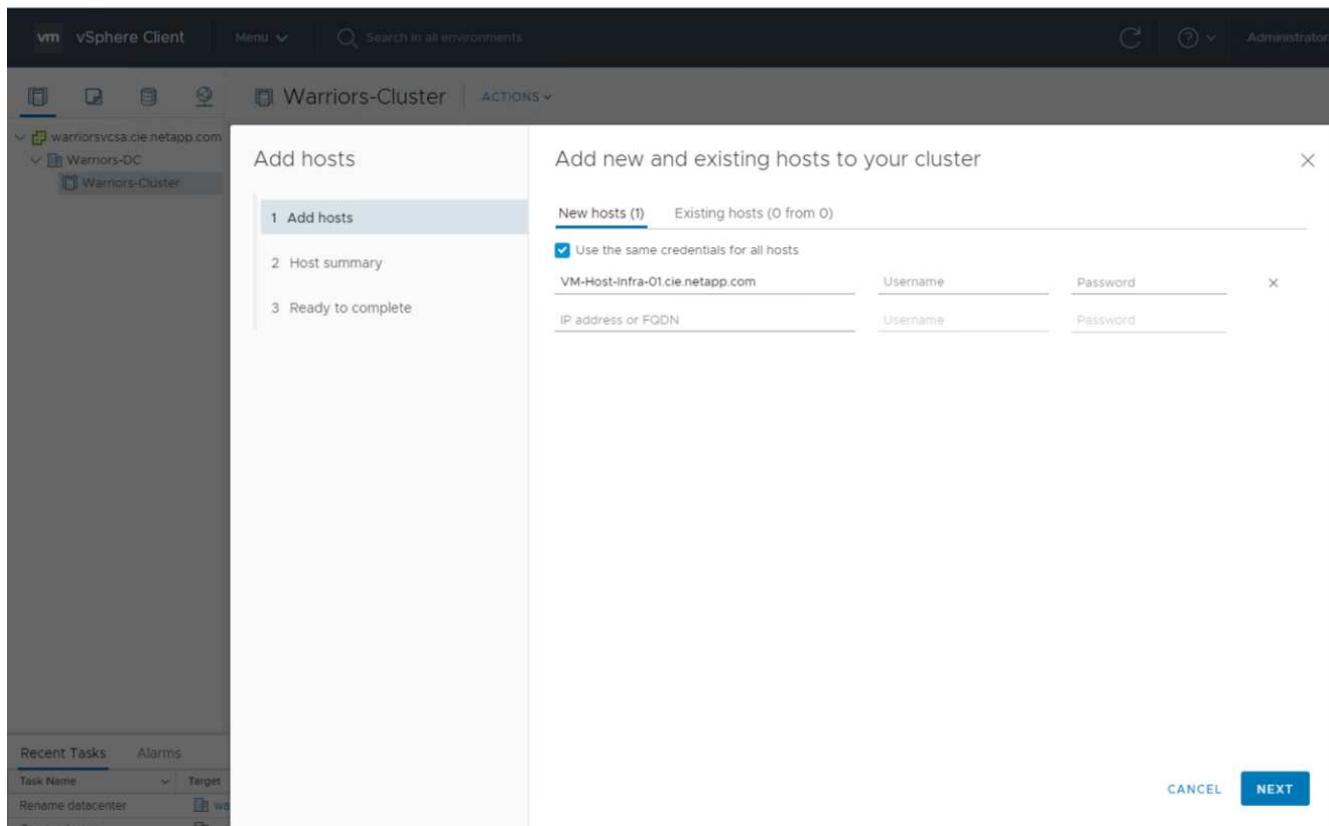
1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Enable DR and vSphere HA by selecting the checkboxes.
4. Click OK.



Add the ESXi hosts to the cluster

To add the ESXi hosts to the cluster, complete the following steps:

1. Right-click the cluster and select Add Host.



2. To add an ESXi host to the cluster, complete the following steps:
 - a. Enter the IP or FQDN of the host. Click Next.
 - b. Enter the root user name and password. Click Next.
 - c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
 - d. Click Next on the Host Summary page.
 - e. Click the green + icon to add a license to the vSphere host.
3. This step can be completed later if desired.
 - a. Click Next to leave lockdown mode disabled.
 - b. Click Next at the VM location page.
 - c. Review the Ready to Complete page. Use the back button to make any changes or select Finish.
4. Repeat steps 1 and 2 for Cisco UCS host B.



This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on the ESXi hosts

To configure coredump on the ESXi hosts, complete the following steps:

1. Log into <https://vCenter> IP:5480/, enter root for the user name, and enter the root password.
2. Click on services and select VMware vSphere ESXI Dump collector.
3. Start the VMware vSphere ESXI Dump collector service.

← → ⌂ Not secure | 172.21.181.105:5480/ui/services

Appliance Management

Mon 10-28-2019 06:51 AM UTC

Summary
RESTART
START
STOP

	Name	▼
<input type="radio"/>	vSAN health Service	
<input type="radio"/>	VMware vSphere Web Client	
<input type="radio"/>	VMware vSphere Update Manager	
<input type="radio"/>	VMware vSphere Profile-Driven Storage Service	
<input checked="" type="radio"/>	VMware vSphere ESXi Dump Collector	
<input type="radio"/>	VMware vSphere Client	
<input type="radio"/>	VMware vSphere Authentication Proxy	
<input type="radio"/>	VMware vService Manager	
<input type="radio"/>	VMware vSAN Data Protection Service	
<input type="radio"/>	VMware vCenter-Services	
<input type="radio"/>	VMware vCenter Server	
<input type="radio"/>	VMware vCenter High Availability	
<input type="radio"/>	VMware Topology Service	

4. Using SSH, connect to the management IP ESXi host, enter root for the user name, and enter the root password.
5. Run the following commands:

```
esxcli system coredump network set -i ip_address_of_core_dump_collector
-v vmk0 -o 6500
esxcli system coredump network set --enable=true
esxcli system coredump network check
```

6. The message Verified the configured netdump server is running appears after you enter the final command.

```
root@VM-Host-Infra-01:~] esxcli system coredump network set -i 172.21.181.105 -
`vmk0 -o 6500
root@VM-Host-Infra-01:~]
root@VM-Host-Infra-01:~] esxcli system coredump network set --enable=true
root@VM-Host-Infra-01:~] esxcli system coredump network check
[verified the configured netdump server is running
```



This process must be completed for any additional hosts added to FlexPod Express.



ip_address_of_core_dump_collector in this validation is the vCenter IP.

Next: NetApp Virtual Storage Console 9.6 deployment procedures.

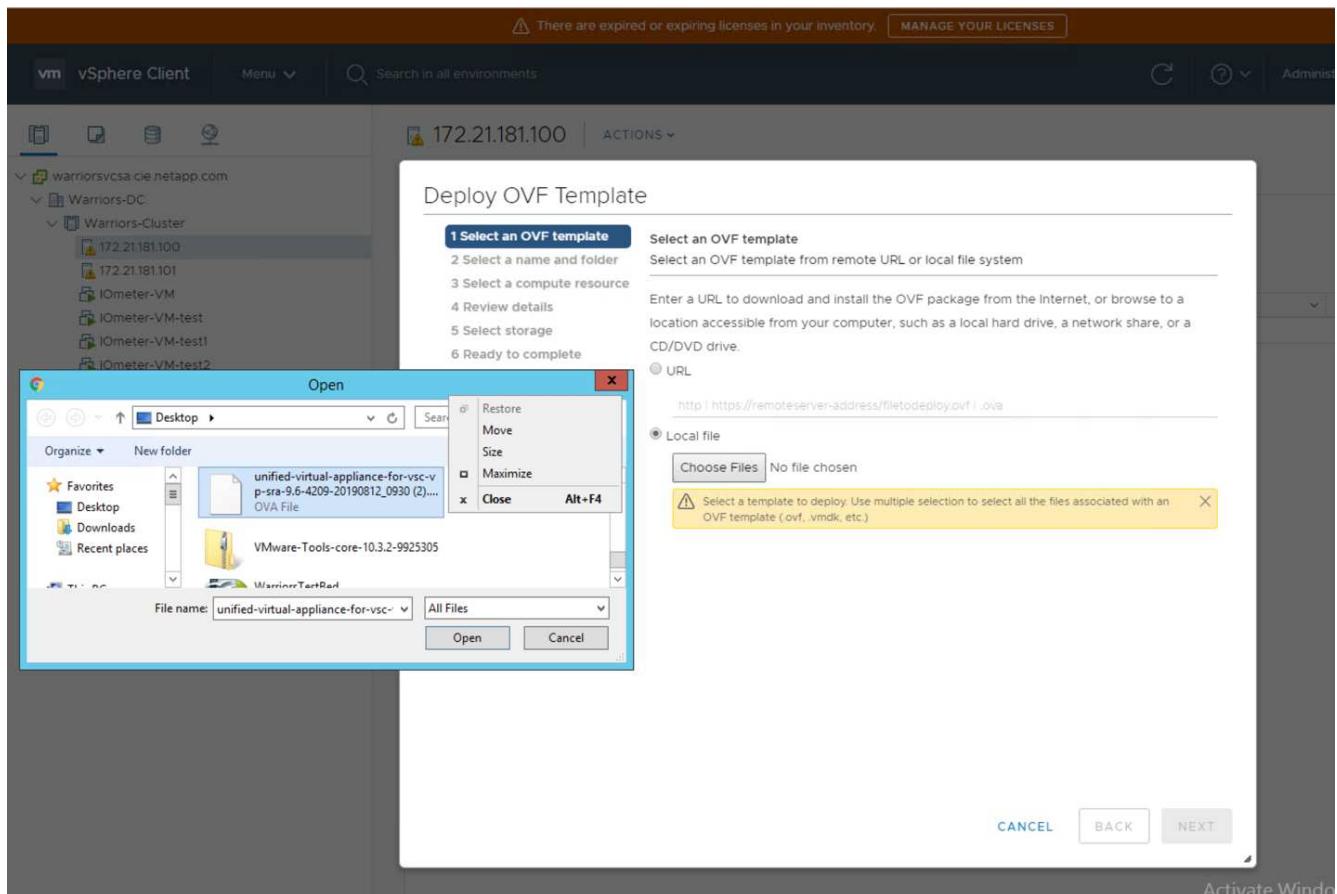
NetApp Virtual Storage Console 9.6 deployment procedures

This section describes the deployment procedures for the NetApp Virtual Storage Console (VSC).

Install Virtual Storage Console 9.6

To install the VSC 9.6 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

1. Go to vSphere Web Client > Host Cluster > Deploy OVF Template.
2. Browse to the VSC OVF file downloaded from the NetApp Support site.



3. Enter the VM name and select a datacenter or folder in which to deploy. Click Next.

Deploy OVF Template

✓ 1 Select an OVF template

✓ 2 Select a name and folder

✓ 3 Select a compute resource

✓ 4 Review details

5 License agreements

✓ 6 Select storage

7 Select networks

8 Customize template

Select a name and folder

Specify a unique name and target location

Virtual machine name: FlexPod-VSC

Select a location for the virtual machine.

✓  warriorsvcsa.cie.netapp.com
 >  FlexPod-Datacenter

4. Select the FlexPod-Cluster ESXi cluster and click Next.

5. Review the details and click Next.

Deploy OVF Template

✓ 1 Select an OVF template

✓ 2 Select a name and folder

✓ 3 Select a compute resource

4 Review details

5 License agreements

6 Select storage

7 Select networks

8 Customize template

9 Ready to complete

Review details

Verify the template details.

Publisher	No certificate present
Product	Virtual Appliance - NetApp VSC, VASA Provider and SRA for ONTAP
Version	See appliance for version
Vendor	NetApp Inc.
Description	Virtual Appliance - NetApp VSC, VASA Provider, and SRA virtual appliance for NetApp storage systems. For more information or support please visit http://www.netapp.com/
Download size	1.0 GB
Size on disk	2.1 GB (thin provisioned)
	53.0 GB (thick provisioned)

CANCEL

BACK

NEXT

6. Click Accept to accept the license and click Next.

7. Select the Thin Provision virtual disk format and one of the NFS datastores. Click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements

6 Select storage

7 Select networks

8 Customize template

9 Ready to complete

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format:

Thin Provision

VM Storage Policy:

Datastore Default

Name	Capacity	Provisioned	Free	Type
Infra_datastore	75 GB	360 KB	75 GB	NF
Infra_datastore1	475 GB	639.9 GB	276.86 GB	NF
Infra_swap (1)	100 GB	4.98 GB	95.02 GB	NF

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

8. From Select Networks, choose a destination network and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage

7 Select networks

- 8 Customize template
- 9 Ready to complete

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	MGMT-Network

1 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL

BACK

NEXT

9. From Customize Template, enter the VSC administrator password, vCenter name or IP address, and other configuration details and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

vCenter Server Address (*)

Specify the IP address/hostname of an existing vCenter to register to.

172.21.181.105

Port (*)

Specify the HTTPS port of an existing vCenter to register to.

443

Username (*)

Specify the username of an existing vCenter to register to.

administrator@vsphere.local

Password (*)

Specify the password of an existing vCenter to register to.

Password
Confirm Password

Network Properties 8 settings

Host Name

Specify the hostname for the appliance. (Leave blank if DHCP is desired)

CANCEL **BACK** **NEXT**

10. Review the configuration details entered and click Finish to complete the deployment of NetApp-VSC VM.
11. Power on the NetApp-VSC VM and open the VM console.
12. During the NetApp-VSC VM boot process, you see a prompt to install VMware Tools. From vCenter, select NetApp-VSC VM > Guest OS > Install VMware Tools.

Booting VSC, VASA Provider, and SRA virtual appliance...Please wait...

VMware Tools OVF vCenter configuration not found.

VMware Tools OVF vCenter configuration not found.

VMware Tools OVF vCenter configuration not found.

VMware Tools installation

Before you can continue the VSC, VASA Provider, and SRA virtual appliance installation, you must install the VMware Tools:

1. Select VM > Guest OS > Install VMware Tools.

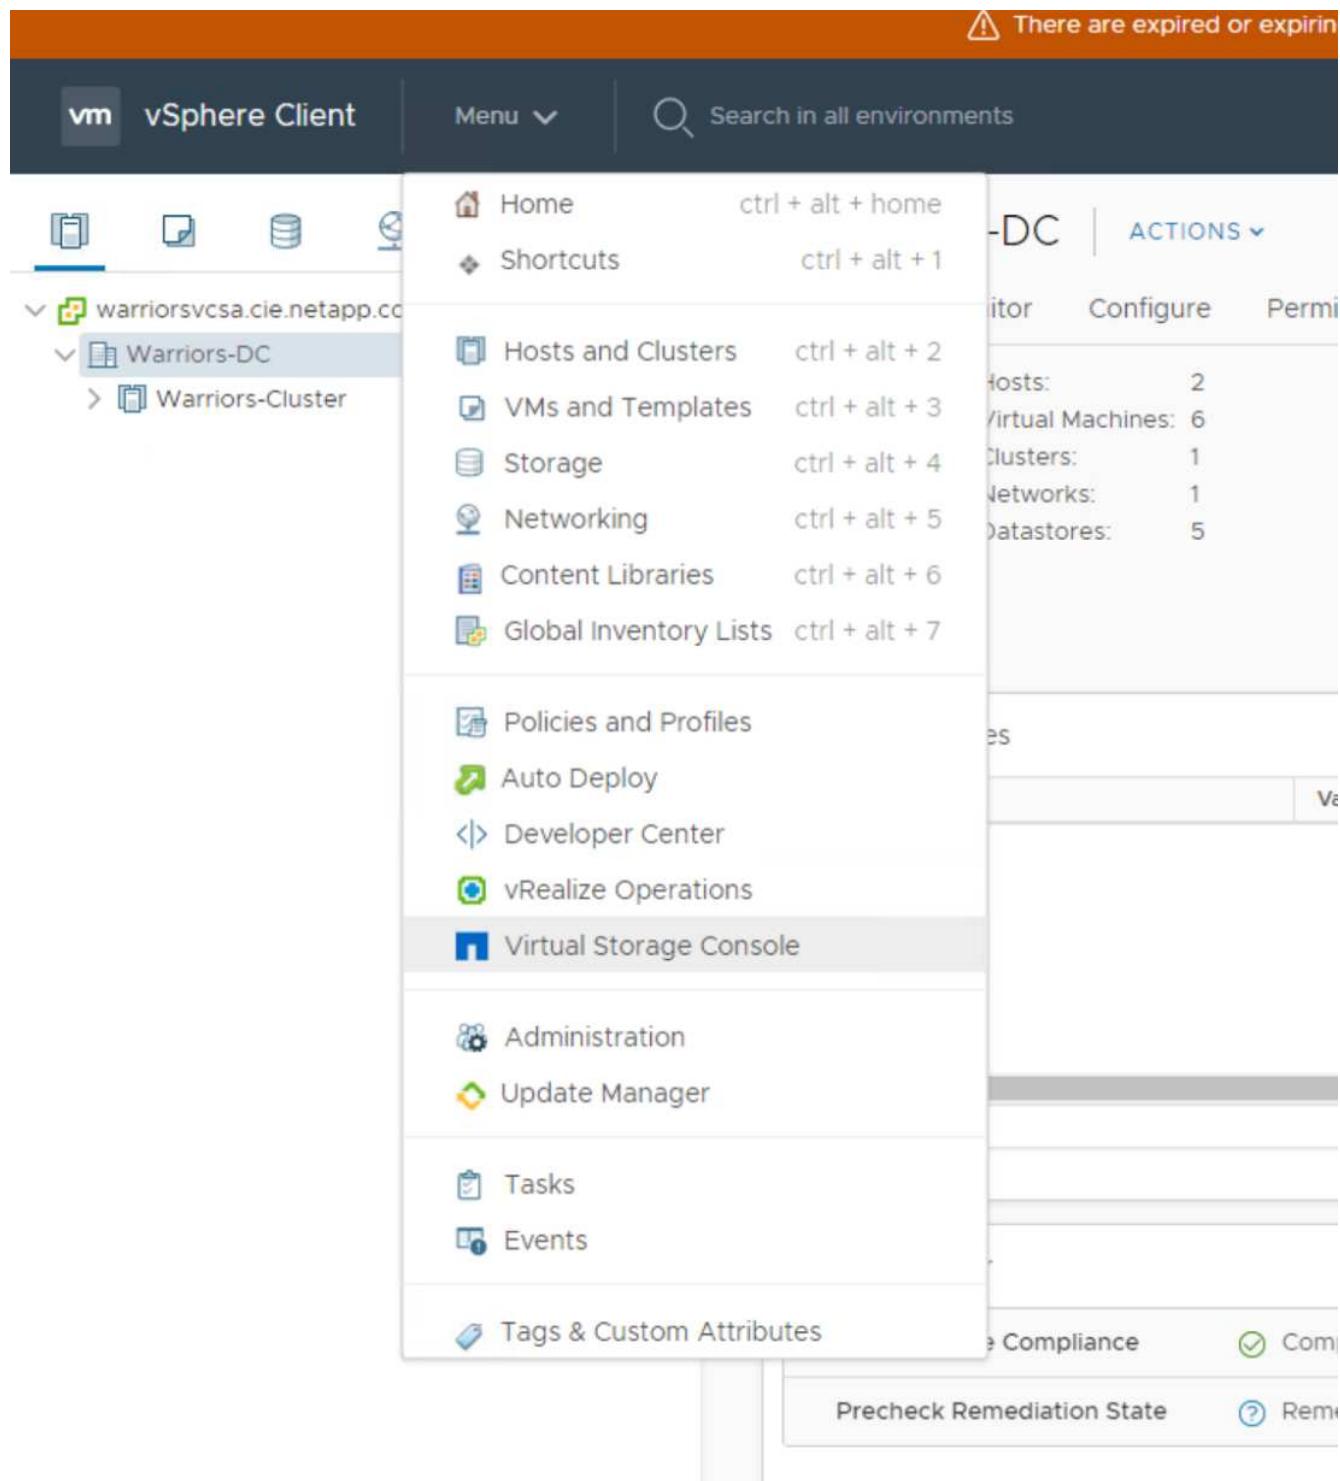
OR

Click on "Install VMware Tools" pop-up box on the vSphere Web Client.

2. Follow the prompts provided by the VMware Tools wizard.

Once you click on mount, the installation process will automatically continue.

13. Networking configuration and vCenter registration information was provided during OVF template customization. Therefore, after the NetApp-VSC VM is running, VSC, vSphere API for Storage Awareness (VASA), and VMware Storage Replication Adapter (SRA) are registered with vCenter.
14. Log out of the vCenter Client and log in again. From the Home menu, confirm that the NetApp VSC is installed.



Download and install the NetApp NFS VAAI Plug-In

To download and install the NetApp NFS VAAI Plug-In, complete the following steps:

1. Download the NetApp NFS Plug-In 1.1.2 for VMware .vib file from the NFS Plugin Download page and save it to your local machine or admin host.
2. Download the NetApp NFS Plug-in for VMware VAAI:
 - a. Go to the [software download page](#).

- b. Scroll down and click NetApp NFS Plug-in for VMware VAAI.
- c. From the Home screen in the vSphere web client, select Virtual Storage Console.
- d. Under Virtual Storage Console > Settings > NFS VAAI Tools, upload the NFS Plug-in by choosing Select File and browsing to the location where the downloaded plug-in is stored.

The screenshot shows the vSphere Client interface with the 'Virtual Storage Console' selected in the sidebar. Under 'Settings', the 'NFS VAAI Tools' tab is active. A note explains the NFS plug-in for VMware VAAI. Below it, the 'Existing version: 1.1.2-3' is shown with a 'CHANGE' link. A 'Upload NFS plug-in for VMware VAAI' button is present, with a 'BROWSE' button and a 'UPLOAD' button. A file browser window is overlaid on the client, showing the file 'NetAppNasPlugin.vib' selected for upload. The file is described as a 'VIB File' (360 KB) located in the 'Desktop > OpenOffice 4.1.1 (en-US) Installation Files' folder.

3. Click Upload to transfer the plug-in to vCenter.
4. Select the host and then select NetApp VSC > Install NFS Plug-in for VMware VAAI.

warriorsvcsa.cie.netapp.com

Summary Monitor **Configure**

Virtual

Standby

Actions - 172.21.181.100

New Virtual Machine...

Deploy OVF Template...

New Resource Pool...

New vApp...

Maintenance Mode

Connection

Power

Certificates

Storage

Add Networking...

Host Profiles

Export System Logs...

Reconfigure for vSphere HA

Assign License...

Settings

Move To...

Tags & Custom Attributes

Remove from Inventory

Add Permission...

Alarms

Update Manager

NetApp VSC

Storage

Storage Adapters

Storage Devices

Host Cache Configuration

Protocol Endpoints

I/O Filters

Networking

Virtual switches

VMkernel adapters

Physical adapters

TCP/IP configuration

Virtual Machines

VM Startup/Shutdown

Agent VM Settings

Default VM Compatibility

Swap File Location

System

Licensing

Host Profile

Time Configuration

Authentication Services

Certificate

Power Management

Advanced System Settings

Host Monitoring

Install NFS Plug-in for VMware VAAI

Update Host and Storage Data

Set Recommended Values

Mount Datastores

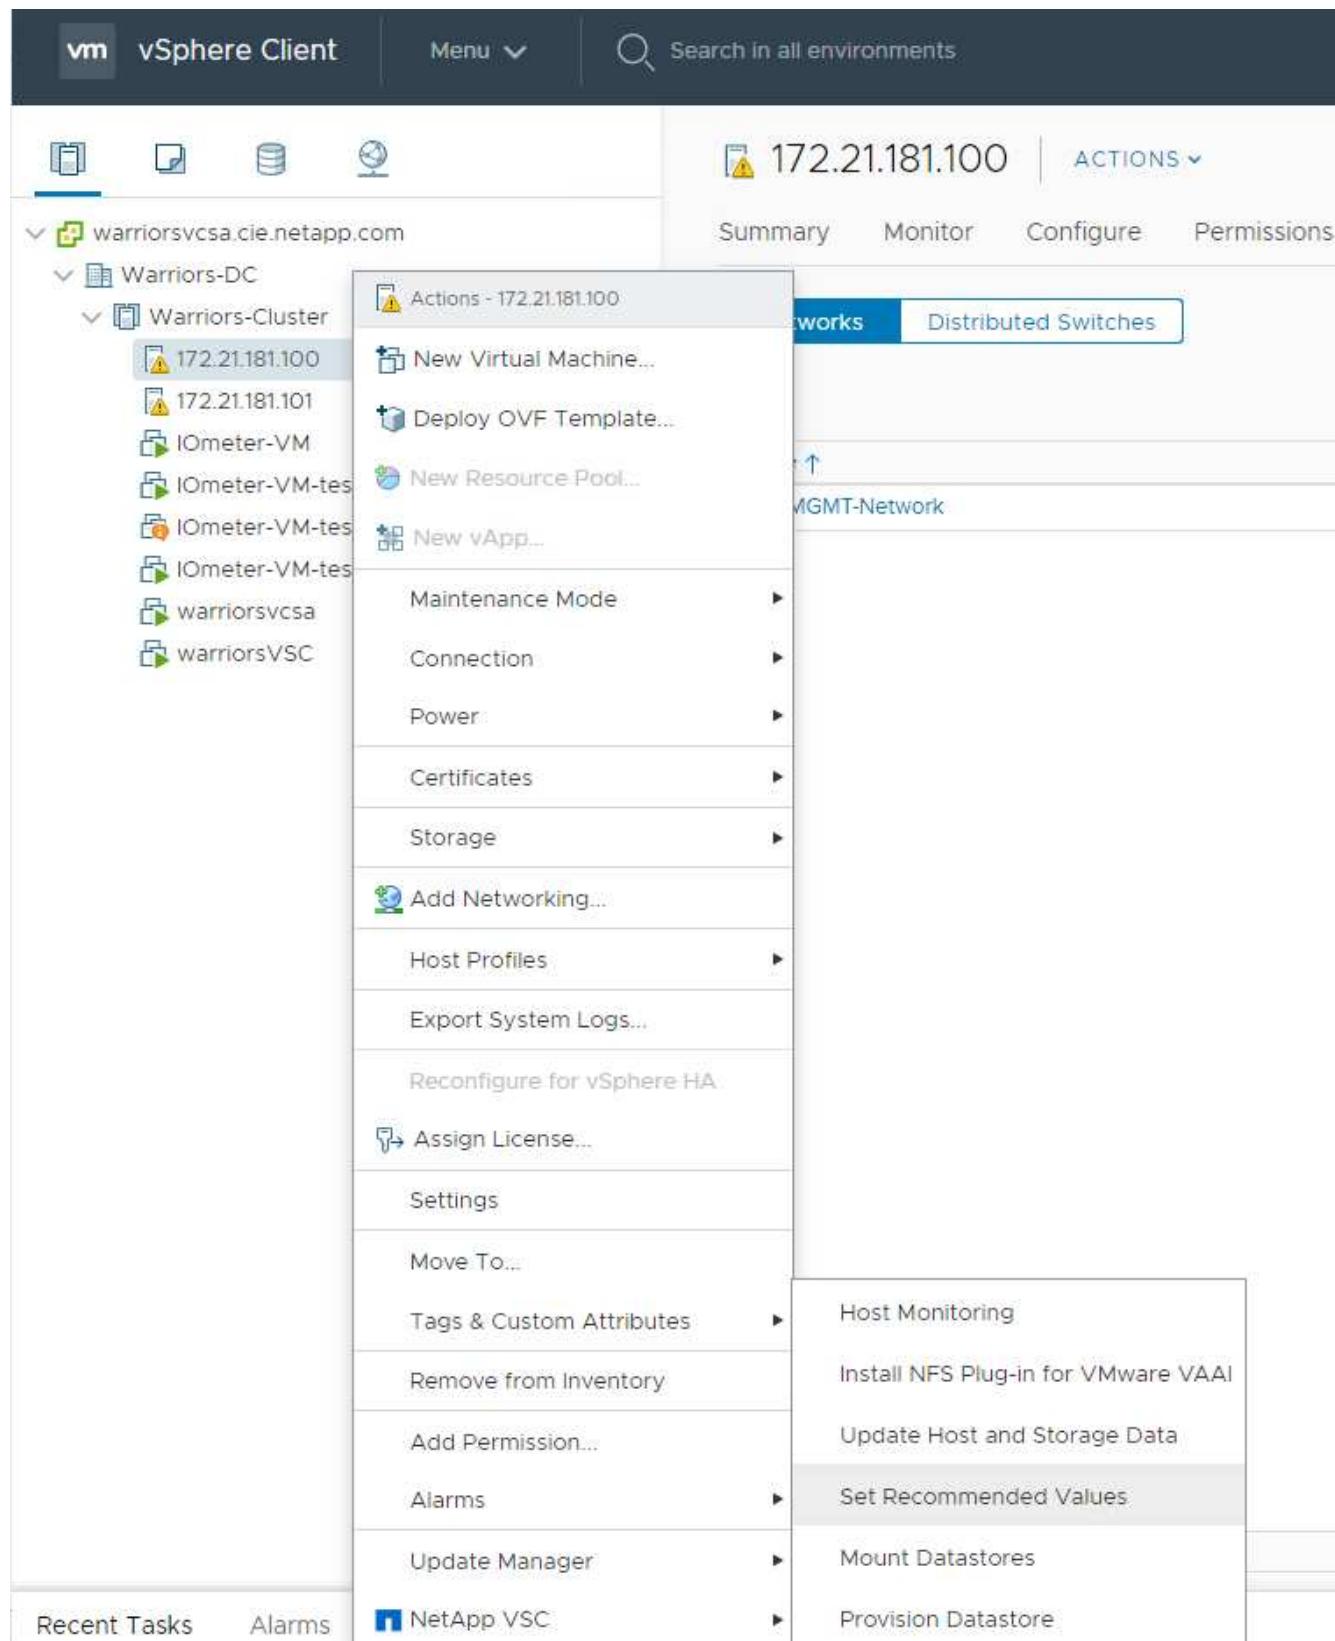
Provision Datastore

Recent Tasks Alarms

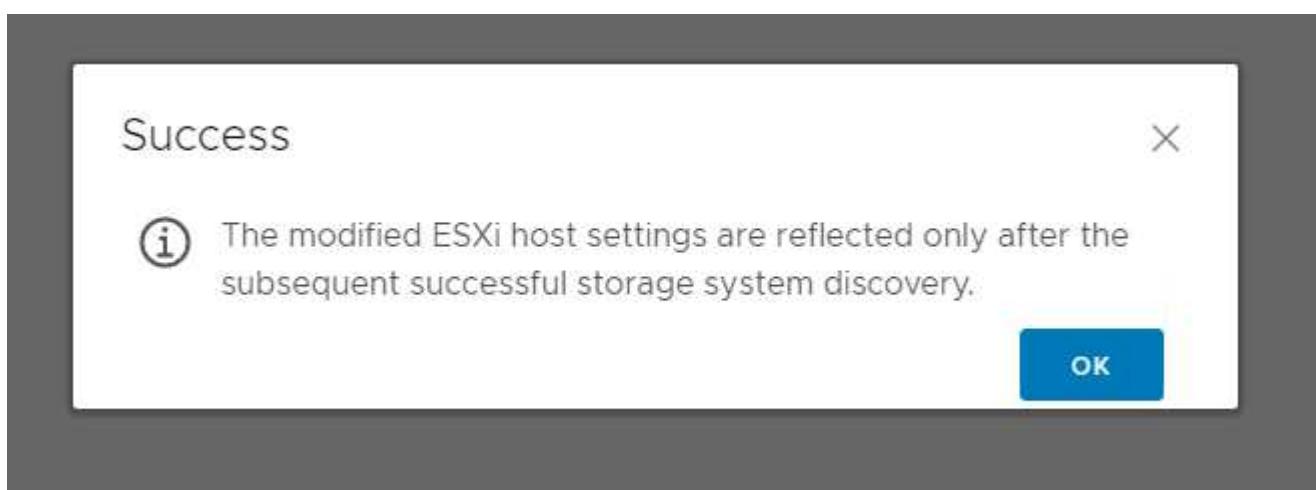
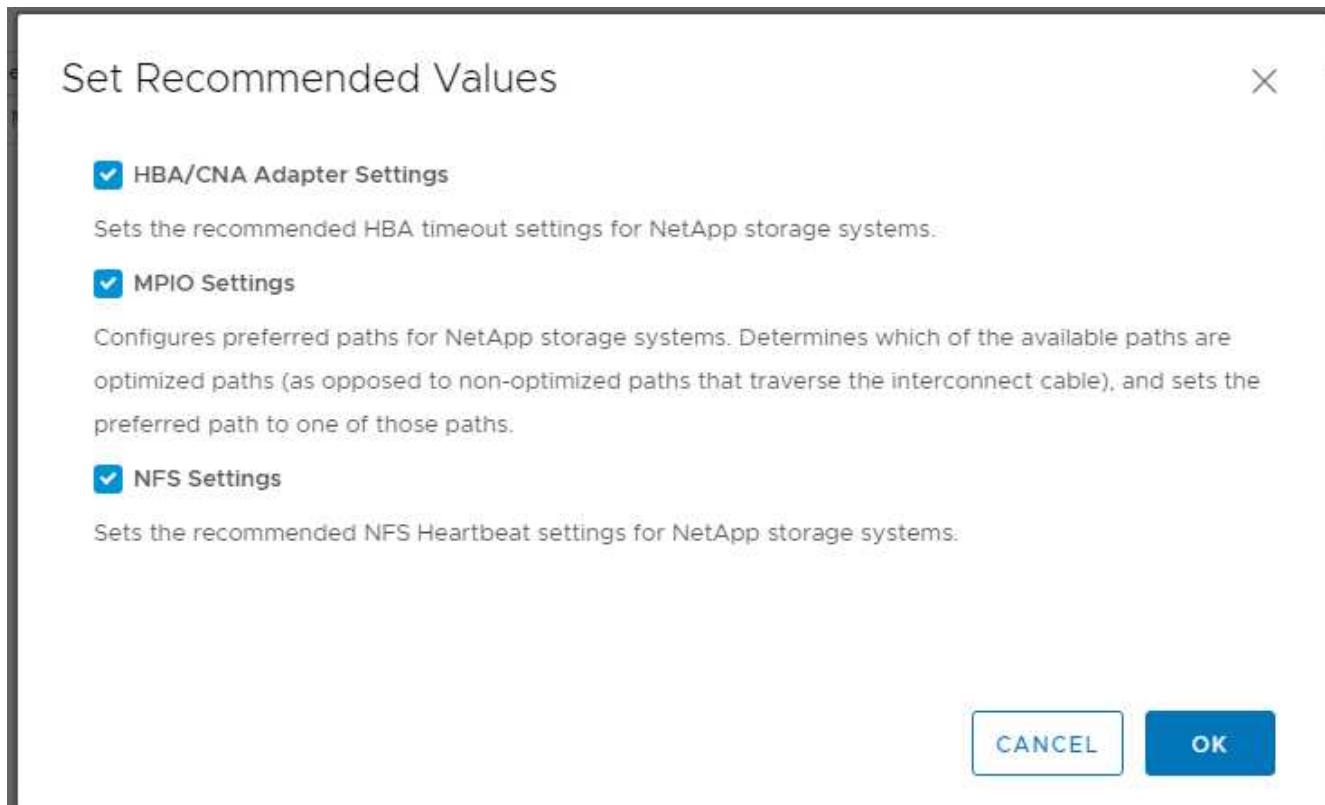
Use the optimal storage settings for the ESXi hosts

VSC enables the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the Home screen, select vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values.



2. Check the settings that you would like to apply to the selected vSphere hosts. Click OK to apply the settings.



3. Reboot the ESXi host after these settings are applied.

Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed for small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

Acknowledgments

The authors would like to acknowledge John George for his support and contribution to

this design.

Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

NetApp Product Documentation

<http://docs.netapp.com>

FlexPod Express with Guide

NVA-1139-DESIGN: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series

<https://www.netapp.com/us/media/nva-1139-design.pdf>

Version history

Version	Date	Document version history
Version 1.0	November 2019	Initial release.

FlexPod Express with Cisco UCS C-Series and AFF A220 Series Design Guide

NVA-1125-DESIGN: FlexPod Express with Cisco UCS C-Series and AFF A220 Series

*

:hardbreaks:
:icons: font
:linkattrs:
:relative_path: ./express/
:imagesdir: /tmp/d20251030-2164368-hqtqlk/source./express/./media/

Savita Kumari, NetApp

In partnership with:



Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology that they are familiar with in their data center.

FlexPod Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp AFF. The components in FlexPod Express are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are

optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

Next: Program summary.

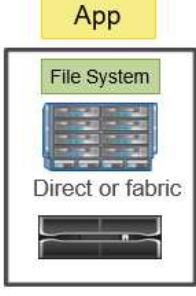
Program summary

FlexPod converged infrastructure portfolio

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or as NetApp Verified Architectures (NVAs). Deviations that are based on customer requirements from a given CVD or NVA are permitted if variations do not result in the deployment of unsupported configurations.

As depicted in the following figure, the FlexPod portfolio includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express.** Offers an entry-level solution that consists of technologies from Cisco and NetApp.
- **FlexPod Datacenter.** Delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select.** Incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.

<h2>Expanded portfolio of platforms</h2>		
FlexPod® Express	FlexPod Datacenter	FlexPod Select
Departmental deployments and VAR velocity Target: Primarily MSB, remote, and departmental deployments	Massively scalable, all virtual Target: Enterprise/service provider	Application purposed Target: Specific application deployments in the enterprise
		
Distinct Architectures Entry-level: Cisco UCS, Cisco Nexus, FAS and AFF	Distinct Architectures Cisco UCS, Cisco Nexus, FAS and AFF	Distinct Architectures Cisco UCS, Cisco Nexus, FAS and AFF

NetApp Verified Architecture program

The NVA program offers customers a verified architecture for NetApp solutions. An NVA means that the NetApp solution has the following qualities:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

This guide details the design of FlexPod Express with VMware vSphere. In addition, this design leverages the all-new AFF A220 system, which runs NetApp ONTAP 9.4 software, Cisco Nexus 3172P switches, and Cisco UCS C220 M5 servers as hypervisor nodes.

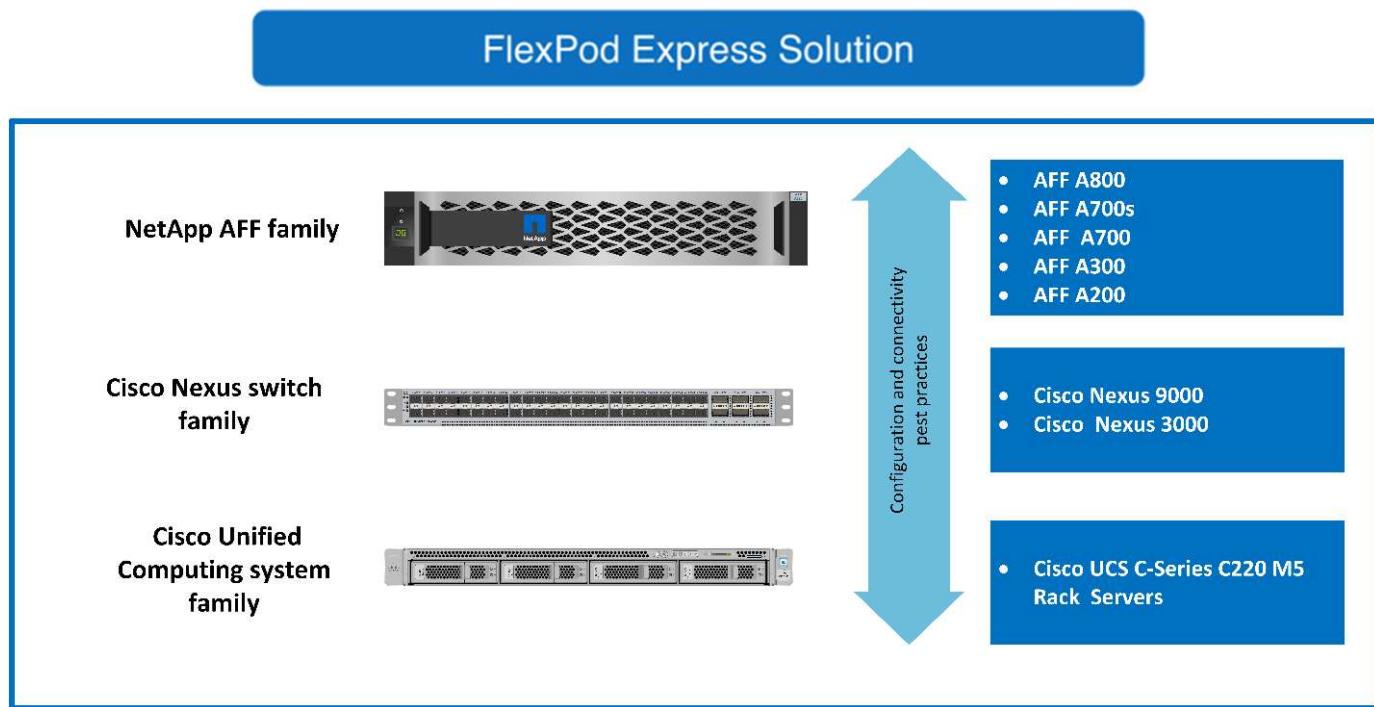
Although this document is validated for AFF A220, this solution also supports FAS2700.

[Next: Solution overview.](#)

Solution overview

FlexPod Express is designed to run mixed virtualization workloads. It is targeted for remote and branch offices and for small to midsize businesses. It is also optimal for larger businesses that want to implement a dedicated solution for a purpose. This new solution for FlexPod Express adds new technologies such as NetApp ONTAP 9.4, NetApp AFF A220, and VMware vSphere 6.7.

The following figure shows the hardware components that are included in the FlexPod Express solution.



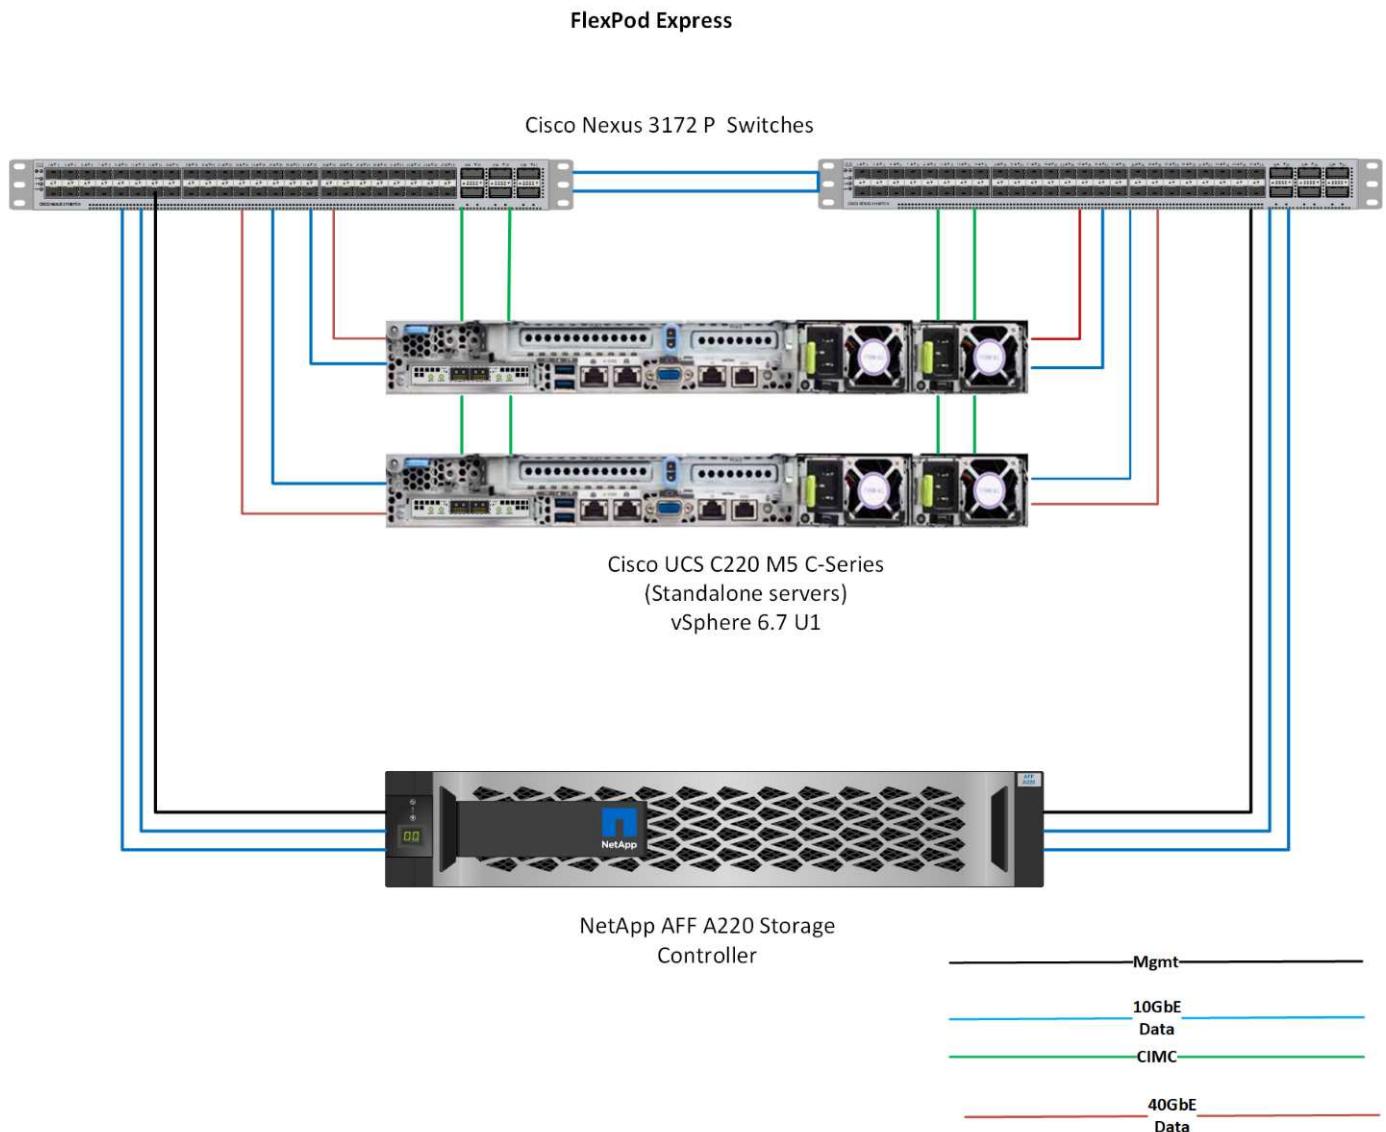
Target audience

This document is intended for those who want to take advantage of an infrastructure that is built to deliver IT efficiency and enable IT innovation. The audience for this document includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, partner engineers, and customers.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. This solution features the new NetApp AFF A220 system, which runs ONTAP 9.4 software, dual Cisco Nexus 3172P switches, and Cisco UCS C220 M5 Rack Servers that run VMware vSphere 6.7. This validated solution uses 10-Gigabit Ethernet (10GbE) technology. The following figure presents an overview. Guidance is also provided on how to scale by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's

evolving business needs.



40GbE is not validated, but it is a supported infrastructure.

Next: [Technology requirements](#).

Technology requirements

FlexPod Express requires a combination of hardware and software components that depends on the selected hypervisor and network speed. In addition, FlexPod Express lays out the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

The following table lists the hardware components that are required for all FlexPod Express configurations and to implement the solution. The hardware components that are used in any particular implementation of the solution might vary based on customer requirements.

Hardware	Quantity
AFF A220 two-node cluster	1
Cisco UCS C220 M5 server	2
Cisco Nexus 3172P switch	2
Cisco UCS Virtual Interface Card (VIC) 1387 for Cisco UCS C220 M5 Rack Server	2
Cisco CVR-QSFP-SFP10G adapter	4

Software requirements

The following tables list the software components that are required to implement the architectures of the FlexPod Express solution.

The following table lists software requirements for the base FlexPod Express implementation.

Software	Version	Details
Cisco Integrated Management Controller (CIMC)	3.1.3	For C220 M5 Rack Servers
Cisco NX-OS	nxos.7.0.3.17.5.bin	For Cisco Nexus 3172P switches
NetApp ONTAP	9.4	For AFF A220 controllers

The following table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

Software	Version
VMware vCenter Server Appliance	6.7
VMware vSphere ESXi	6.7
NetApp VAAI Plug-In for ESXi	1.1.2

[Next: Design choices.](#)

Design choices

The following technologies were chosen during the process of architecting this design. Each technology serves a specific purpose in the FlexPod Express infrastructure solution.

NetApp AFF A220 Series with ONTAP 9.4

This solution leverages two of the newest NetApp products: NetApp AFF A220 and ONTAP 9.4 software.

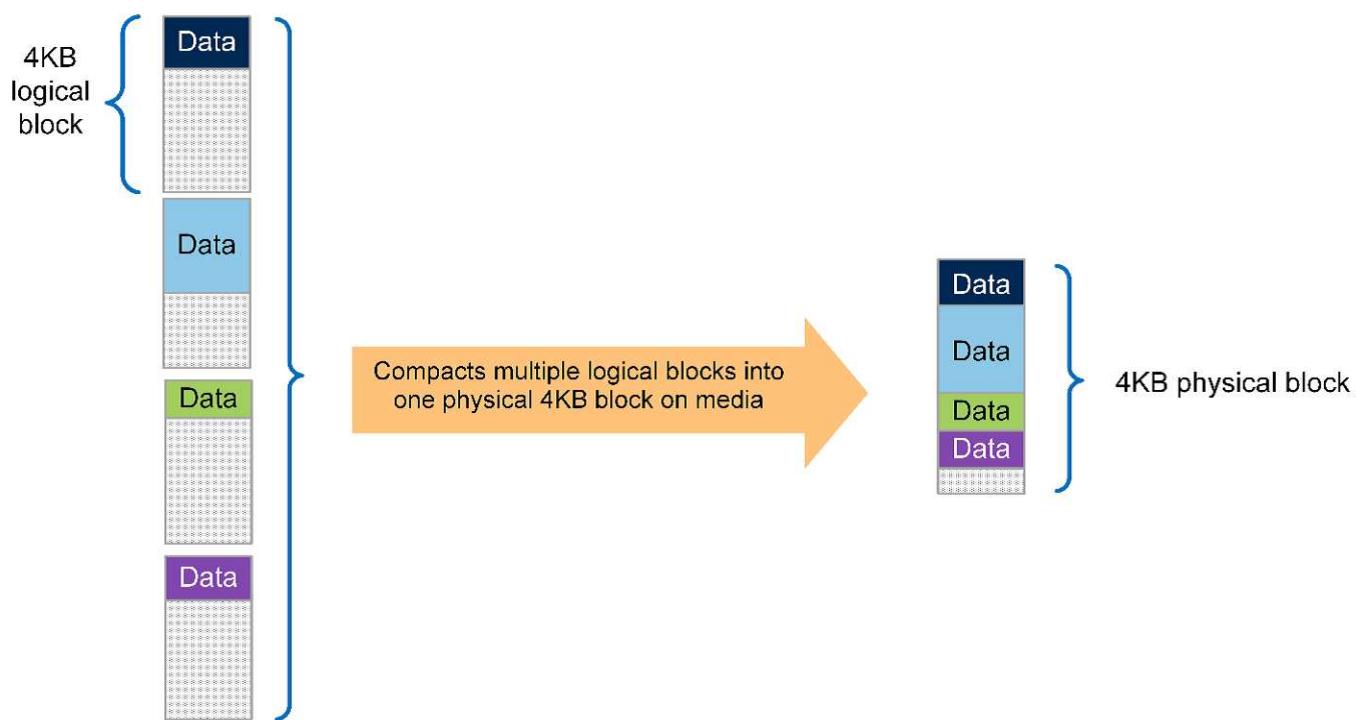
AFF A220 system

For more information about the AFF A220 hardware system, see the [AFF A-Series homepage](#).

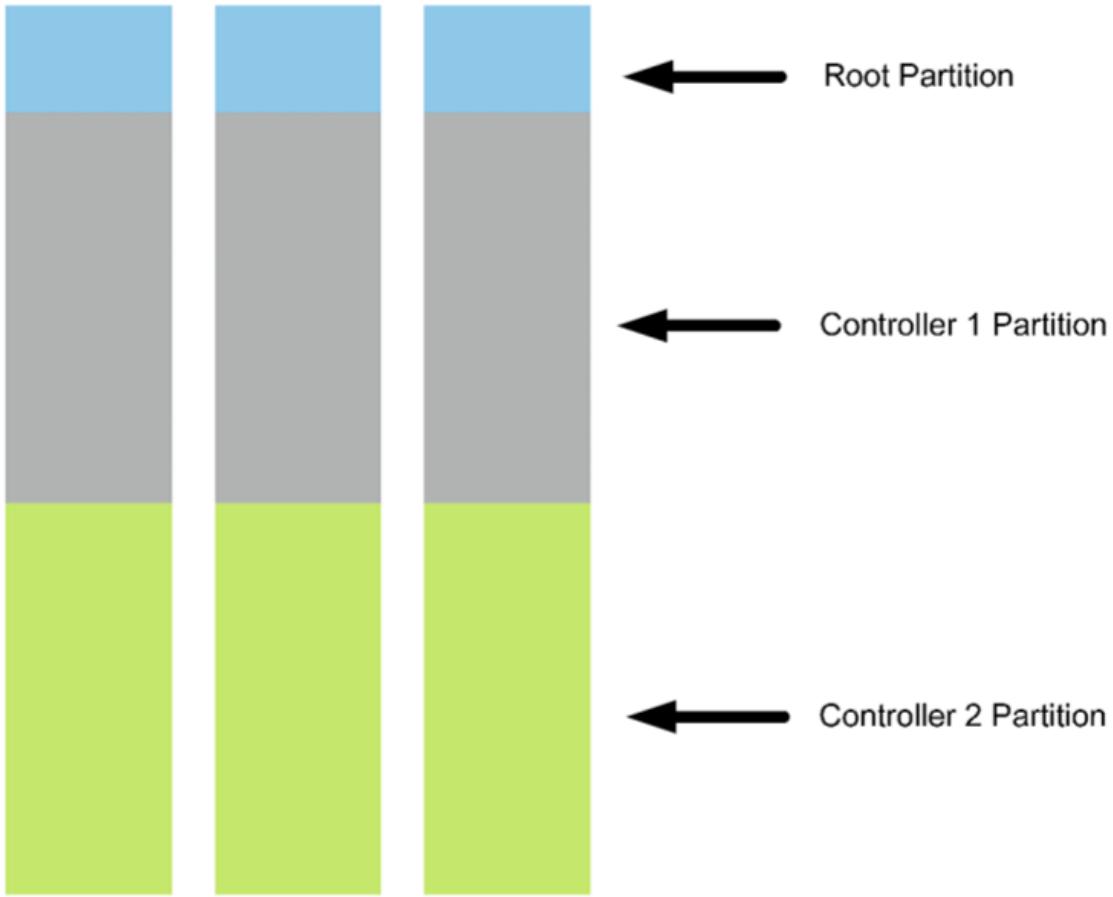
ONTAP 9.4 software

NetApp AFF A220 systems use the new ONTAP 9.4 software. ONTAP 9.4 is the industry's leading enterprise data management software. It combines new levels of simplicity and flexibility with powerful data management capabilities, storage efficiencies, and leading cloud integration.

ONTAP 9.4 has several features that are well suited for the FlexPod Express solution. Foremost is NetApp's commitment to storage efficiencies, which can be one of the most important features for small deployments. The hallmark NetApp storage efficiency features such as deduplication, compression, and thin provisioning are available in ONTAP 9.4 with a new addition, compaction. Because the NetApp WAFL system always writes 4KB blocks, compaction combines multiple blocks into a 4KB block when the blocks are not using their allocated space of 4KB. The following figure illustrates this process.



Also, root-data partitioning can be leveraged on the AFF A220 system. This partitioning allows the root aggregate and two data aggregates to be striped across the disks in the system. Therefore, both controllers in a two-node AFF A220 cluster can leverage the performance of all the disks in the aggregate. See the following figure.



These are just a few key features that complement the FlexPod Express solution. For details about the additional features and functionality of ONTAP 9.4, see the [ONTAP 9 Data Management Software datasheet](#). Also, see the NetApp [ONTAP 9 Documentation Center](#), which has been updated to include ONTAP 9.4.

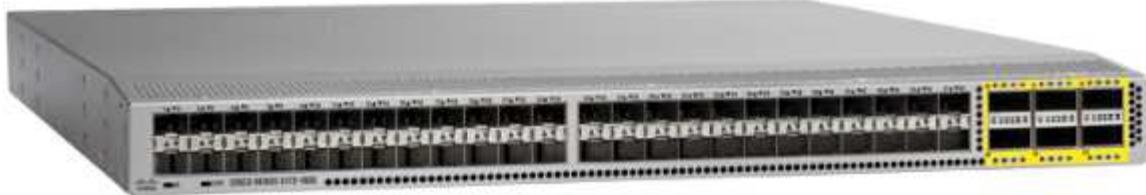
Cisco Nexus 3000 Series

The Cisco Nexus 3172P is a robust, cost- effective switch that offers 1/10/40/100Gbps switching. The Cisco Nexus 3172PQ switch, part of the Unified Fabric family, is a compact, 1-rack-unit (1RU) switch for top-of-rack data center deployments. (See the following figure.) It offers up to seventy-two 1/10GbE ports in 1RU or forty-eight 1/10GbE plus six 40GbE ports in 1RU. And for maximum physical layer flexibility, it also supports 1/10/40Gbps.

Because all the various Cisco Nexus series models run the same underlying operating system, NX-OS, multiple Cisco Nexus models are supported in the FlexPod Express and FlexPod Datacenter solutions.

Performance specifications include:

- Line-rate traffic throughput (both layers 2 and 3) on all ports
- Configurable maximum transmission units (MTUs) of up to 9216 bytes (jumbo frames)



For more information about Cisco Nexus 3172 switches, see the [Cisco Nexus 3172PQ, 3172TQ, 3172TQ-32T, 3172PQ-XL, and 3172TQ-XL switches data sheet](#).

Cisco UCS C-Series

The Cisco UCS C-Series rack server was chosen for FlexPod Express because its many configuration options allow it to be tailored for specific requirements in a FlexPod Express deployment.

Cisco UCS C-Series rack servers deliver unified computing in an industry-standard form factor to reduce TCO and to increase agility.

Cisco UCS C-Series rack servers provide the following benefits:

- A form-factor-agnostic entry point into Cisco UCS
- Simplified and fast deployment of applications
- Extension of unified computing innovations and benefits to rack servers
- Increased customer choice with unique benefits in a familiar rack package



The Cisco UCS C220 M5 rack server (in the previous figure) is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. It is a high-density two-socket rack server that delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. Cisco UCS C-Series Rack Servers can be deployed as standalone servers or as part of Cisco UCS to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

For more information about C220 M5 servers, see the [Cisco UCS C220 M5 Rack Server Data Sheet](#).

Connectivity options for C220 M5 rack servers

The connectivity options for the C220 M5 rack servers are as follows:

- **Cisco UCS VIC 1387**

The Cisco UCS VIC 1387 (in the following figure) offers dual-port enhanced QSFP+ 40GbE and FC over Ethernet (FCoE) in a modular-LAN-on-motherboard (mLOM) form factor. The mLOM slot can be used to install a Cisco VIC without consuming a Peripheral Component Interconnect Express (PCIe) slot, providing greater I/O expandability.



For more information about the Cisco UCS VIC 1387 adapter, see the [Cisco UCS Virtual Interface Card 1387](#) data sheet.

- **CVR-QSFP-SFP10G adapter**

The Cisco QSA Module converts a QSFP port into an SFP or SFP+ port. With this adapter, customers have the flexibility to use any SFP+ or SFP module or cable to connect to a lower-speed port on the other end of the network. This flexibility enables a cost-effective transition to 40GbE by maximizing the use of high-density 40GbE QSFP platforms. This adapter supports all SFP+ optics and cable reaches, and it supports several 1GbE SFP modules. Because this project has been validated by using 10GbE connectivity and because the VIC 1387 used is 40GbE, the CVR-QSFP-SFP10G adapter (in the following figure) is used for conversion.



VMware vSphere 6.7

VMware vSphere 6.7 is one hypervisor option for use with FlexPod Express. VMware vSphere allows organizations to reduce their power and cooling footprint while confirming that the purchased compute capacity is used to its fullest. In addition, VMware vSphere allows hardware failure protection (VMware High Availability, or VMware HA) and compute resource load balancing across a cluster of vSphere hosts (VMware Distributed Resource Scheduler, or VMware DRS).

Because it restarts only the kernel, VMware vSphere 6.7 allows customers to “quick boot” where it loads vSphere ESXi without restarting the hardware. This feature is available only with platforms and drivers that are

on the Quick Boot Whitelist. vSphere 6.7 extends the capabilities of the vSphere Client, which can do about 90% of what the vSphere Web Client can do.

In vSphere 6.7, VMware has extended this capability to enable customers to set Enhanced vMotion Compatibility (EVC) per virtual machine (VM) rather than per host basis. In vSphere 6.7, VMware has also exposed the APIs that can be used to create instant clones.

The following are some of the features of vSphere 6.7 U1:

- Fully featured HTML5 web-based vSphere Client
- vMotion for NVIDIA GRID vGPU VMs. Support for Intel FPGA.
- vCenter Server Converge Tool to move from external PSC to internal PCS.
- Enhancements for vSAN (HCI updates).
- Enhanced content library.

For details about vSphere 6.7 U1, see [What's New in vCenter Server 6.7 Update 1](#). Although this solution was validated with vSphere 6.7, it supports any vSphere version qualified with the other components by the NetApp Interoperability Matrix Tool. NetApp recommends deploying vSphere 6.7U1 for its fixes and enhanced features.

Boot architecture

Following are the supported options for the FlexPod Express boot architecture:

- iSCSI SAN LUN
- Cisco FlexFlash SD Card
- Local disk

Because FlexPod Datacenter is booted from iSCSI LUNs, solution manageability is enhanced by also using iSCSI boot for FlexPod Express.

[Next: Solution verification.](#)

Solution verification

Cisco and NetApp designed and built FlexPod Express to serve as a premier infrastructure platform for their customers. Because it was designed with industry-leading components, customers can trust FlexPod Express as their infrastructure foundation. In keeping with the fundamental principles of the FlexPod portfolio, the FlexPod Express architecture was thoroughly tested by Cisco and NetApp data center architects and engineers. From redundancy and availability to each individual feature, the entire FlexPod Express architecture is validated to instill confidence in our customers and to build trust in the design process.

VMware vSphere 6.7 was verified on the FlexPod Express infrastructure components. This validation included 10GbE uplink connectivity options for the hypervisor.

[Next: Conclusion.](#)

Conclusion

FlexPod Express offers a simple and effective solution by providing a validated design that uses industry-leading components. By scaling and by providing options for the hypervisor platform, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed keeping in mind small to midsize businesses, remote and branch offices, and other businesses that require dedicated solutions.

Next: [Where to find additional information](#).

Where to find additional information

To learn more about the information that is described in this document, see the following documents and websites:

- NetApp documentation

<https://docs.netapp.com>

- FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 Deployment Guide

<https://www.netapp.com/us/media/nva-1123-deploy.pdf>

FlexPod Express with Cisco UCS C-Series and AFF A220 Series Deployment Guide

NVA-1123-DEPLOY: FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 deployment guide

Savita Kumari, NetApp

In partnership with:



Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology with which they are familiar in their data center.

FlexPod Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized

and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools to which they are accustomed. New FlexPod Express customers can easily adapt to managing FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution overview

This FlexPod Express solution is part of the FlexPod Converged Infrastructure Program.

FlexPod Converged Infrastructure Program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

As depicted in the figure below, the FlexPod program includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express.** Offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter.** Delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select.** Incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.

FlexPod Express	FlexPod Datacenter	FlexPod Select
Departmental deployments and VAR velocity Target: Primarily MSB, remote, and departmental deployments	Massively scalable, all virtual Target: Enterprise/service provider	Application purposed Target: Specific application deployments in the enterprise
 Entry-level: Cisco UCS, Nexus, and NetApp FAS	 Cisco UCS, Cisco Nexus, and FAS	 NetApp E-Series and FAS

NetApp Verified Architecture Program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. A NetApp Verified Architecture provides a NetApp solution architecture with the following qualities:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

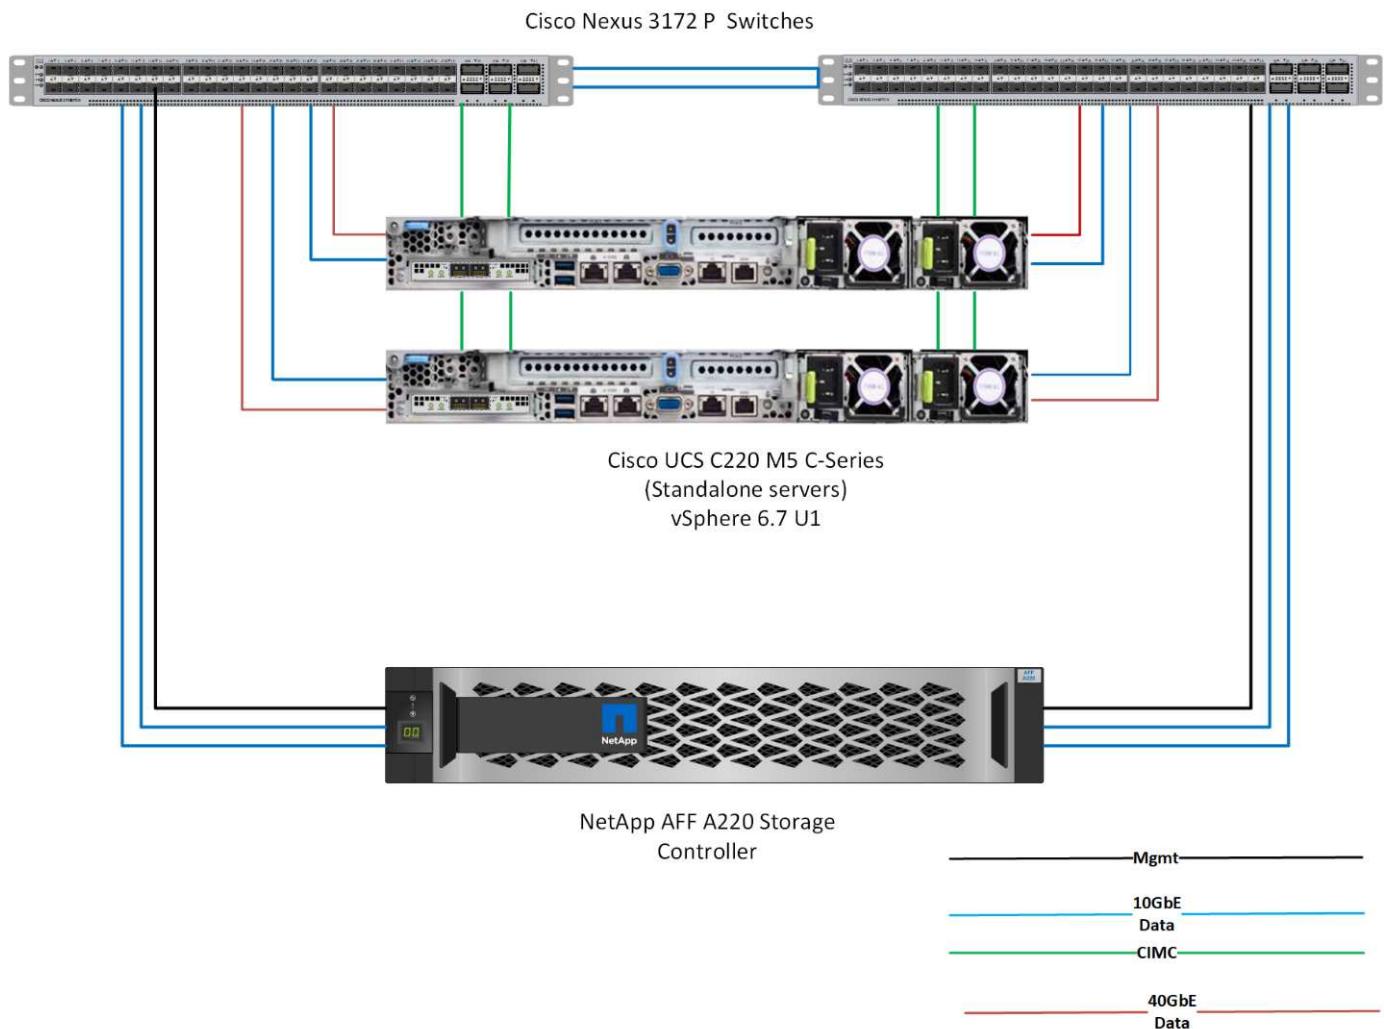
This guide details the design of FlexPod Express with VMware vSphere. In addition, this design uses the all-new AFF A220 system, which runs NetApp ONTAP 9.4; the Cisco Nexus 3172P; and Cisco UCS C-Series C220 M5 servers as hypervisor nodes.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. This solution features the new NetApp AFF A220 running ONTAP 9.4, dual Cisco Nexus 3172P switches, and Cisco UCS C220 M5 rack servers that run VMware vSphere 6.7. This validated solution uses 10GbE technology. Guidance is also provided on how to scale compute capacity by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.

The following figure shows FlexPod Express with VMware vSphere 10GbE architecture.

FlexPod Express



This validation uses 10GbE connectivity and a Cisco UCS VIC 1387, which is 40GbE. To achieve 10GbE connectivity, the CVR-QSFP-SFP10G adapter is used.

Use case summary

The FlexPod Express solution can be applied to several use cases, including the following:

- Remote offices or branch offices
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is best suited for virtualized and mixed workloads.



Although this solution was validated with vSphere 6.7, it supports any vSphere version qualified with the other components by the NetApp Interoperability Matrix Tool. NetApp recommends deploying vSphere 6.7U1 for its fixes and enhanced features.

Following are some features of vSphere 6.7 U1:

- Fully featured HTML5 web-based vSphere client
- vMotion for NVIDIA GRID vGPU VMs. Support for Intel FPGA
- vCenter Server Converge Tool to move from external PSC to internal PCS
- Enhancements for vSAN (HCI updates)
- Enhanced content library

For details about vSphere 6.7 U1, see [What's New in vCenter Server 6.7 Update 1](#).

Technology requirements

A FlexPod Express system requires a combination of hardware and software components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

The following table lists the hardware components required for all FlexPod Express configurations.

Hardware	Quantity
AFF A220 HA Pair	1
Cisco C220 M5 server	2
Cisco Nexus 3172P switch	2
Cisco UCS virtual interface card (VIC) 1387 for the C220 M5 server	2
CVR-QSFP-SFP10G adapter	4

The following table lists the hardware required in addition to the base configuration for implementing 10GbE.

Hardware	Quantity
Cisco UCS C220 M5 server	2
Cisco VIC 1387	2
CVR-QSFP-SFP10G adapter	4

Software requirements

The following table lists the software components required to implement the architectures of the FlexPod Express solutions.

Software	Version	Details
Cisco Integrated Management Controller (CIMC)	3.1(3g)	For Cisco UCS C220 M5 rack servers

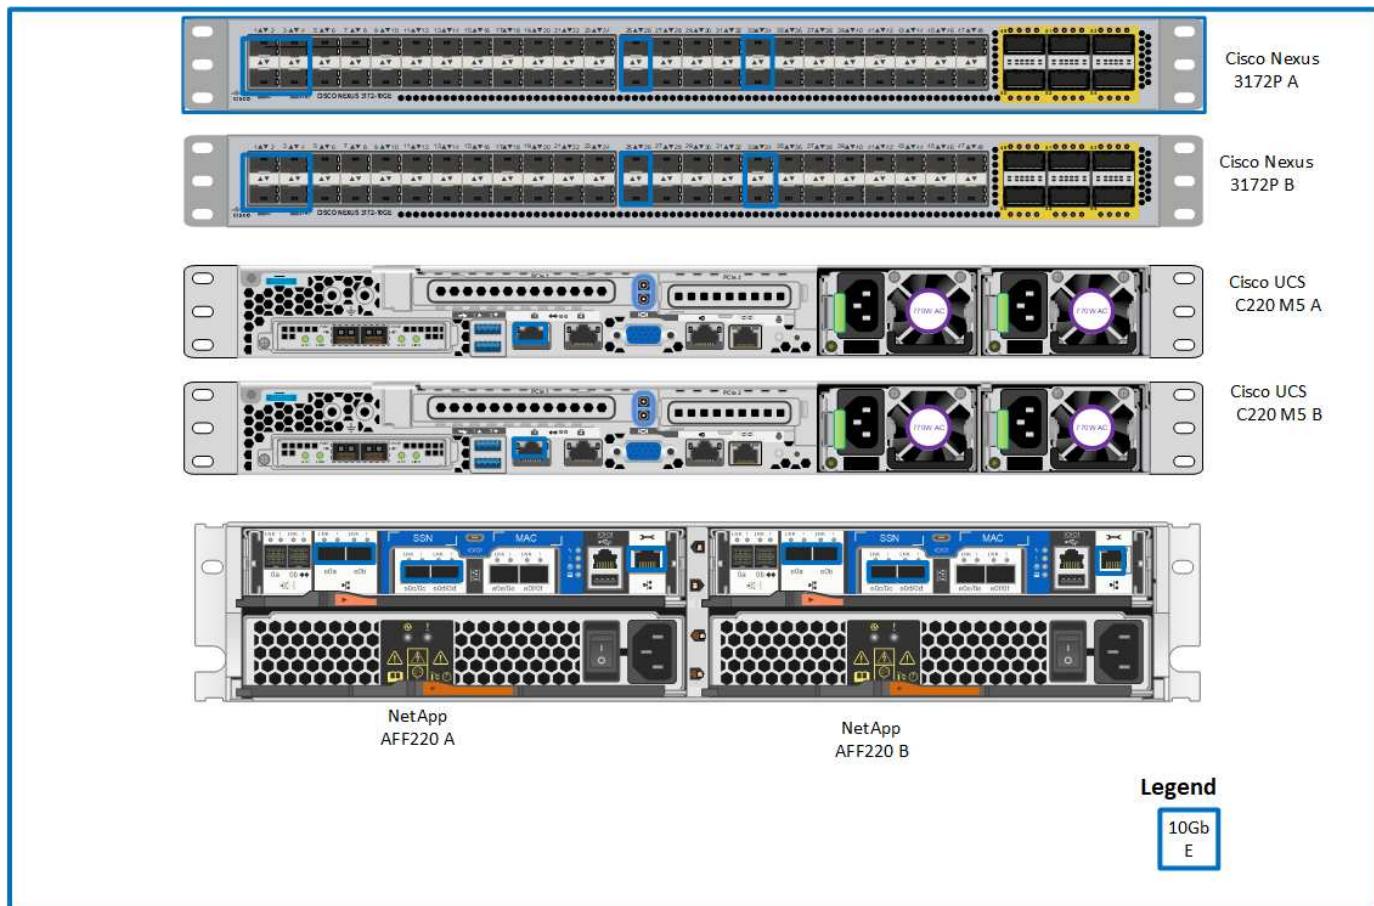
Software	Version	Details
Cisco nenic driver	1.0.25.0	For VIC 1387 interface cards
Cisco NX-OS	nxos.7.0.3.17.5.bin	For Cisco Nexus 3172P switches
NetApp ONTAP	9.4	For AFF A220 controllers

The following table lists the software required for all VMware vSphere implementations on FlexPod Express.

Software	Version
VMware vCenter server appliance	6.7
VMware vSphere ESXi hypervisor	6.7
NetApp VAAI Plug-In for ESXi	1.1.2

FlexPod Express cabling information

The following figure shows the reference validation cabling.



The following table shows cabling information for the Cisco Nexus switch 3172P A.

Local device	Local port	Remote device	Remote port
Cisco Nexus switch 3172P A	Eth1/1	NetApp AFF A220 storage controller A	e0c

Local device	Local port	Remote device	Remote port
	Eth1/2	NetApp AFF A220 storage controller B	e0c
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM1 with CVR-QSFP-SFP10G adapter
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM1 with CVR-QSFP-SFP10G adapter
	Eth1/25	Cisco Nexus switch 3172P B	Eth1/25
	Eth1/26	Cisco Nexus switch 3172P B	Eth1/26
	Eth1/33	NetApp AFF A220 storage controller A	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server A	CIMC

The following table shows cabling information for Cisco Nexus switch 3172P B.

Local device	Local port	Remote device	Remote port
Cisco Nexus switch 3172P B	Eth1/1	NetApp AFF A220 storage controller A	e0d
	Eth1/2	NetApp AFF A220 storage controller B	e0d
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM2 with CVR-QSFP-SFP10G adapter
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM2 with CVR-QSFP-SFP10G adapter
	Eth1/25	Cisco Nexus switch 3172PA	Eth1/25
	Eth1/26	Cisco Nexus switch 3172PA	Eth1/26
	Eth1/33	NetApp AFF A220 storage controller B	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server B	CIMC

The following table shows the cabling information for NetApp AFF A220 storage controller A.

Local device	Local port	Remote device	Remote port
NetApp AFF A220 storage controller A	e0a	NetApp AFF A220 storage controller B	e0a
	e0b	NetApp AFF A220 storage controller B	e0b

Local device	Local port	Remote device	Remote port
	e0c	Cisco Nexus switch 3172P A	Eth1/1
	e0d	Cisco Nexus switch 3172P B	Eth1/1
	e0M	Cisco Nexus switch 3172P A	Eth1/33

The following table shows cabling information for NetApp AFF A220 storage controller B.

Local device	Local port	Remote device	Remote port
NetApp AFF A220 storage controller B	e0a	NetApp AFF A220 storage controller A	e0a
	e0b	NetApp AFF A220 storage controller A	e0b
	e0c	Cisco Nexus switch 3172P A	Eth1/2
	e0d	Cisco Nexus switch 3172P B	Eth1/2
	e0M	Cisco Nexus switch 3172P B	Eth1/33

Deployment procedures

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches.

In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, <<text>> appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01>vlan create vif0 <<mgmt_vlan_id>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. The table below describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.



If you use separate in-band and out-of-band management VLANs, you must create a layer-3 route between them. For this validation, a common management VLAN was used.

AN Name	VLAN Purpose	ID Used in Validating This Document
Management VLAN	VLAN for management interfaces	3437
Native VLAN	VLAN to which untagged frames are assigned	2
NFS VLAN	VLAN for NFS traffic	3438
VMware vMotion VLAN	VLAN designated for the movement of virtual machines from one physical host to another	3441
Virtual machine traffic VLAN	VLAN for virtual machine application traffic	3442
iSCSI-A-VLAN	VLAN for iSCSI traffic on fabric A	3439
iSCSI-B-VLAN	VLAN for iSCSI traffic on fabric B	3440

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <<var_xxxx_vlan>>, where xxxx is the purpose of the VLAN (such as iSCSI-A).

The table below lists the VMware virtual machines created.

Virtual machine description	Host name
VMware vCenter Server	

Cisco Nexus 3172P deployment procedure

The following section details the Cisco Nexus 3172P switch configuration used in a FlexPod Express environment.

Initial setup of Cisco Nexus 3172P switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.



This procedure assumes that you are using a Cisco Nexus 3172P running NX-OS software release 7.0(3)I7(5).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 3172P switches can be connected to an existing management network, or the mgmt0 interfaces of the 3172P switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.

In this deployment guide, the FlexPod Express Cisco Nexus 3172P switches are connected to an existing

management network.

3. To configure the Cisco Nexus 3172P switches, power on the switch and follow the on- screen prompts, as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

*Note: setup is mainly used for configuring the system initially, when no configuration is present. So setup always assumes system defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): y

Do you want to enforce secure password standard (yes/no) [y]: y

Create another login account (yes/no) [n]: n

Configure read-only SNMP community string (yes/no) [n]: n

Configure read-write SNMP community string (yes/no) [n]: n

Enter the switch name : 3172P-B

Continue with Out-of-band (mgmt0) management configuration? (yes/no)

[y]: y

Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>

Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>

Configure the default gateway? (yes/no) [y]: y

IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>

Configure advanced IP options? (yes/no) [n]: n

Enable the telnet service? (yes/no) [n]: n

Enable the ssh service? (yes/no) [y]: y

Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa

Number of rsa key bits <1024-2048> [1024]: <enter>

Configure the ntp server? (yes/no) [n]: y

NTP server IPv4 address : <<var_ntp_ip>>

Configure default interface layer (L3/L2) [L2]: <enter>

Configure default switchport interface state (shut/noshut) [noshut]: <enter>

Configure CoPP system profile (strict/moderate/lenient/dense)

[strict]: <enter>

4. You then see a summary of your configuration, and you are asked if you would like to edit it. If your configuration is correct, enter n.

Would you like to edit the configuration? (yes/no) [n]: n

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

Use this configuration and save it? (yes/no) [y]: Enter

6. Repeat this procedure for Cisco Nexus switch B.

Enable advanced features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options.

 The `interface-vlan` feature is required only if you use the `back-to-back mgmt0` option described throughout this document. This feature allows you to assign an IP address to the interface VLAN (switch virtual interface), which enables in-band management communication to the switch (such as through SSH).

1. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode using the command `(config t)` and run the following commands:

```
feature interface-vlan
feature lacp
feature vpc
```

The default port channel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the port channel. You can achieve better distribution across the members of the port channel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

2. From configuration mode `(config t)`, enter the following commands to set the global port channel load-balancing configuration on Cisco Nexus switch A and switch B:

```
port-channel load-balance src-dst ip-14port
```

Perform global spanning-tree configuration

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (config t), run the following commands to configure the default spanning- tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
```

Define VLANs

Before individual ports with different VLANs are configured, the layer 2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (config t), run the following commands to define and describe the layer 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure access and management port descriptions

As is the case with assigning names to the layer 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (config t) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus Switch A

```

int eth1/1
  description AFF A220-A e0c
int eth1/2
  description AFF A220-B e0c
int eth1/3
  description UCS-Server-A: MLOM port 0
int eth1/4
  description UCS-Server-B: MLOM port 0
int eth1/25
  description vPC peer-link 3172P-B 1/25
int eth1/26
  description vPC peer-link 3172P-B 1/26
int eth1/33
  description AFF A220-A e0M
int eth1/34
  description UCS Server A: CIMC

```

Cisco Nexus Switch B

```

int eth1/1
  description AFF A220-A e0d
int eth1/2
  description AFF A220-B e0d
int eth1/3
  description UCS-Server-A: MLOM port 1
int eth1/4
  description UCS-Server-B: MLOM port 1
int eth1/25
  description vPC peer-link 3172P-A 1/25
int eth1/26
  description vPC peer-link 3172P-A 1/26
int eth1/33
  description AFF A220-B e0M
int eth1/34
  description UCS Server B: CIMC

```

Configure server and storage management interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (`config t`), enter the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus Switch A

```
int eth1/33-34
switchport mode access
switchport access vlan <<mgmt_vlan>>
spanning-tree port type edge
speed 1000
exit
```

Cisco Nexus Switch B

```
int eth1/33-34
switchport mode access
switchport access vlan <<mgmt_vlan>>
spanning-tree port type edge
speed 1000
exit
```

Perform virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer-2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back mgmt0 configuration, use the addresses defined on the interfaces and verify that they can communicate by using the ping [\[switch_A/B_mgmt0_ip_addr\] vrf](#) management command.

From configuration mode (config t), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus Switch A

```

vpc domain 1
  role priority 10
    peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source
    <<switch_A_mgmt0_ip_addr>> vrf management
    peer-gateway
    auto-recovery
    ip arp synchronize
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<nfs_vlan_id>>, <<vmotion_vlan_id>>,
    <<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
    <<iSCSI_B_vlan_id>>
    spanning-tree port type network
    vpc peer-link
    no shut
exit
copy run start

```

Cisco Nexus Switch B

```

vpc domain 1
  peer-switch
    role priority 20
    peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source
    <<switch_B_mgmt0_ip_addr>> vrf management
    peer-gateway
    auto-recovery
    ip arp synchronize
int eth1/25- 26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>, <<vmotion_vlan_id>>,
  <<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
  <<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
no shut
exit
copy run start

```

Configure storage port channels

The NetApp storage controllers allow an active-active connection to the network using the Link Aggregation Control Protocol (LACP). The use of LACP is preferred because it adds both negotiation and logging between the switches. Because the network is set up for vPC, this approach enables you to have active-active connections from the storage to separate physical switches. Each controller has two links to each of the switches. However, all four links are part of the same vPC and interface group (IFGRP).

From configuration mode (config t), run the following commands on each of the switches to configure the individual interfaces and the resulting port channel configuration for the ports connected to the NetApp AFF controller.

1. Run the following commands on switch A and switch B to configure the port channels for storage controller A:

```

int eth1/1
  channel-group 11 mode active
int Po11
  description vPC to Controller-A
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<nfs_vlan_id>>,<<mgmt_vlan_id>>,<<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11
  no shut

```

2. Run the following commands on switch A and switch B to configure the port channels for storage controller B.

```

int eth1/2
  channel-group 12 mode active
int Po12
  description vPC to Controller-B
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<mgmt_vlan_id>>,
<<iSCSI_A_vlan_id>>,<<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12
  no shut
exit
copy run start

```



In this solution validation, an MTU of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components will result in packets being dropped and these packets.

Configure server connections

The Cisco UCS servers have a two-port virtual interface card, VIC1387, that is used for data traffic and booting of the ESXi operating system using iSCSI. These interfaces are configured to fail over to one another, providing additional redundancy beyond a single link. Spreading these links across multiple switches enables the server to survive even a complete switch failure.

From configuration mode (config t), run the following commands to configure the port settings for the interfaces connected to each server.

Cisco Nexus Switch A: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/3-4
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan
<<iSCSI_A_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_id>>
spanning-tree port type edge trunk
mtu9216
no shut
exit
copy run start
```

Cisco Nexus Switch B: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/3-4
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan
<<iSCSI_B_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
no shut
exit
copy run start
```

In this solution validation, an MTU of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components will result in packets being dropped and these packets will need to be transmitted again. This will affect the overall performance of the solution.

To scale the solution by adding additional Cisco UCS servers, run the previous commands with the switch ports that the newly added servers have been plugged into on switches A and B.

Uplink into existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 3172P switches included in the FlexPod environment into the infrastructure. The uplinks may be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is

completed.

Next: [NetApp Storage Deployment Procedure \(Part 1\)](#)

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF2xx series installation

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

1. Access the [HWU](#) application to view the system configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFF2XX Series prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections: Electrical Requirements, Supported Power Cords, and Onboard Ports and Cables.

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF A220 Documentation](#).

NetApp ONTAP 9.4

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.4 Software Setup Guide](#).



This system is set up in a two-node switchless cluster configuration.

The following table shows ONTAP 9.4 installation and configuration information.

Cluster detail	Cluster detail value
Cluster node A IP address	<<var_nodeA_mgmt_ip>>
Cluster node A netmask	<<var_nodeA_mgmt_mask>>
Cluster node A gateway	<<var_nodeA_mgmt_gateway>>
Cluster node A name	<<var_nodeA>>
Cluster node B IP address	<<var_nodeB_mgmt_ip>>

Cluster detail	Cluster detail value
Cluster node B netmask	<<var_nodeB_mgmt_mask>>
Cluster node B gateway	<<var_nodeB_mgmt_gateway>>
Cluster node B name	<<var_nodeB>>
ONTAP 9.4 URL	<<var_url_boot_software>>
Name for cluster	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster B gateway	<<var_clustermgmt_gateway>>
Cluster B netmask	<<var_clustermgmt_mask>>
Domain name	<<var_domain_name>>
DNS server IP (you can enter more than one)	<<var_dns_server_ip>>
NTP server IP (you can enter more than one)	<<var_ntp_server_ip>>

Configure Node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

2. Allow the system to boot.

autoboot

3. Press Ctrl-C to enter the Boot menu.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter **y** to set the newly installed software as the default to be used for subsequent reboots.
12. Enter **y** to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press **Ctrl-C** to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter **y** to zero disks, reset config, and install a new file system.
16. Enter **y** to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

17. While node A is initializing, begin configuring node B.

Configure Node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press **Ctrl-C** to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press **Ctrl-C** to enter the Boot menu.

```
autoboot
```

3. Press **Ctrl-C** when prompted.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and **y** to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.

5. Enter **y** to perform an upgrade.
6. Select **e0M** for the network port you want to use for the download.
7. Enter **y** to reboot now.
8. Enter the IP address, netmask, and default gateway for **e0M** in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press **Enter** for the user name, indicating no user name.
11. Enter **y** to set the newly installed software as the default to be used for subsequent reboots.
12. Enter **y** to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press **Ctrl-C** to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter **y** to zero disks, reset config, and install a new file system.
16. Enter **y** to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation of Node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.4 boots on the node for the first time.



The node and cluster setup procedure has changed slightly in ONTAP 9.4. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up Node A.

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.

Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".

To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp Technical

Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.

For further information on AutoSupport, see:

<http://support.netapp.com/autosupport/>

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]:

Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>

Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>

Enter the node management interface default gateway:

<<var_nodeA_mgmt_gateway>>

A node management interface on port e0M with IP address

<<var_nodeA_mgmt_ip>> has been created.

Use your web browser to complete cluster setup by accessing

https://<<var_nodeA_mgmt_ip>>

Otherwise, press Enter to complete cluster setup using the command line interface:

2. Navigate to the IP address of the node's management interface.

Cluster setup can also be performed by using the CLI. This document describes cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup to configure the cluster.

4. Enter <<var_clustername>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster.



Cluster Name:

Nodes

i Not sure all nodes have been discovered? Refresh



Cluster Configuration:

Switched Cluster

Switchless Cluster

i Username:

Password:

Confirm Password:

Cluster Base License (Optional):

i For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional):

Enter comma separated license keys...

i Cluster Base License is mandatory to add Feature Licenses.

Submit

5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter <<var_clustermgmt_ip>> in the Cluster Management IP Address field, <<var_clustermgmt_mask>> in the Netmask field, and <<var_clustermgmt_gateway>> in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter <<var_nodeA_mgmt_ip>> for node B.

- d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.

You can enter multiple DNS server IP addresses.

- e. Enter <<var_ntp_server_ip>> in the Primary NTP Server field.

You can also enter an alternate NTP server.

8. Configure the support information.

- a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
- b. Enter the SMTP mail host and email address for event notifications.

You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



② AutoSupport

③ Proxy URL (Optional)

ⓘ Connection is verified after configuring AutoSupport on all nodes.

④ Event Notifications

Notify me through:

<input checked="" type="checkbox"/> Email	SMTP Mail Host	Email Addresses
	<input type="text"/>	Separate email addresses with a comma...
<hr/>		
<input type="checkbox"/> SNMP	SNMP Trap Host	
	<input type="text"/>	<hr/>
<hr/>		
<input type="checkbox"/> Syslog	Syslog Server	
	<input type="text"/>	<hr/>
<hr/>		

Submit

9. When indicated that the cluster configuration has completed, click **Manage Your Cluster** to configure the storage.

Continuation of storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the

storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set on-board UTA2 ports personality

1. Verify the current mode and the current type of the ports by running the `ucadmin show` command.

```
AFF A220::> ucadmin show
      Current  Current  Pending  Pending  Admin
Node      Adapter  Mode    Type     Mode    Type   Status
-----  -----
-----  -----
AFF A220_A  0c      fc      target   -      -      online
AFF A220_A  0d      fc      target   -      -      online
AFF A220_A  0e      fc      target   -      -      online
AFF A220_A  0f      fc      target   -      -      online
AFF A220_B  0c      fc      target   -      -      online
AFF A220_B  0d      fc      target   -      -      online
AFF A220_B  0e      fc      target   -      -      online
AFF A220_B  0f      fc      target   -      -      online
8 entries were displayed.
```

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If not, change the port personality by using the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode
cna -type target
```

The ports must be offline to run the previous command. To take a port offline, run the following command:

```
`network fcp adapter modify -node <home node of the port> -adapter <port
name> -state down`
```



If you changed the port personality, you must reboot each node for the change to take effect.

Rename management logical interfaces (LIFs)

To rename the management LIFs, complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clustername>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clustername>> -lif
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif
cluster_setup_node_mgmt_lif_AFF_A220_B_1 -newname AFF_A220-02_mgmt1
```

Set auto-revert on cluster management

Set the auto-revert parameter on the cluster management interface.

```
network interface modify -vserver <<clustername>> -lif cluster_mgmt -auto-
revert true
```

Set up service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_nodeA>> -address
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeA_sp_ip>>
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>
system service-processor network modify -node <<var_nodeB>> -address
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeB_sp_ip>>
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```

Enabling failover on one node enables it for both nodes.

3. Verify the HA status of the two-node cluster.

This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.



Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true  
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```

The message **Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner** indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node <<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node <<var_nodeB>>
```

Create jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove data ports from default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following command:

```
net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.

Do you want to continue? {y|n}: y
```

Configure IFGRP LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Make sure the switch is properly configured.

From the cluster prompt, complete the following steps.

```

ifgrp create -node <<var_nodeA>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0d
ifgrp create -node <<var_nodeB>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0d

```

Configure jumbo frames in NetApp ONTAP

To configure an ONTAP network port to use jumbo frames (that usually have an MTU of 9,000 bytes), run the following commands from the cluster shell:

```

AFF A220::> network port modify -node node_A -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
      this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
      this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports
<<var_nodeA>>:a0a-<<var_nfs_vlan_id>>, <<var_nodeB>>:a0a-
<<var_nfs_vlan_id>>

```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_A_id>>, <<var_nodeB>>:a0a-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_B_id>>, <<var_nodeB>>:a0a-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<mgmt_vlan_id>>

```

Create aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Start with five disks; you can add disks to an aggregate when additional storage is required.

The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1`_nodeA` is online.

Configure time zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is America/New_York. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select md5 as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select des as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>> -transport https -support enable -noteto <<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the vserver create command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_nodeA -rootvolume-security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fcp
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
`nfs create -vserver Infra-SVM -udp disabled`
```

5. Turn on the SVM vstorage parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
`vserver nfs modify -vserver Infra-SVM -vstorage enabled`  
`vserver nfs show`
```



Commands are prefaced by `vserver` in the command line because storage virtual machines were previously called servers.

Configure NFSv3 in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
ESXi host A NFS IP address	<code><<var_esxi_hostA_nfs_ip>></code>
ESXi host B NFS IP address	<code><<var_esxi_hostB_nfs_ip>></code>

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>  
-rорule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 2 -protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>>  
-rорule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS C-Series servers are added.

Create iSCSI service in ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM  
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

1. Create a volume to be the load- sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate  
aggr1_nodeA -size 1GB -type DP  
volume create -vserver Infra_Vserver -volume rootvol_m02 -aggregate  
aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path  
Infra-SVM:rootvol_m01 -type LS -schedule 15min  
snapmirror create -source-path Infra-SVM:rootvol -destination-path  
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol  
snapmirror show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag  
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.

Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...
```

Example: `security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial 552429A6`

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...
```

Example: `security certificate create -common-name infra-svm. netapp.com -type server -size 2048 -country US -state "North Carolina" -locality "RTP" -organization "NetApp" -unit "FlexPod" -email-addr "abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver Infra-SVM`

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...
```

Example: `security ssl modify -vserver Infra-SVM -server-enabled true -client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common-name infra-svm.netapp.com`

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be
    interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http
-vserver <<var_clusternamespace>>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow SVM to be available by the web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled
true
```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate
aggr1_nodeA -size 500GB -state online -policy default -junction-path
/infra_datastore_1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap
-space-guarantee none -percent-snapshot-space 0 -snapshot-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0
```

Enable deduplication in ONTAP

To enable deduplication on appropriate volumes, run the following commands:

```
volume efficiency on -vserver Infra-SVM -volume infra_datastore_1
volume efficiency on -vserver Infra-SVM -volume esxi_boot
```

Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

```

lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware -space-reserve disabled

```



When adding an extra Cisco UCS C-Series server, an extra boot LUN must be created.

Create iSCSI LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
Storage node A iSCSI LIF01A	<<var_nodeA_iscsi_lif01a_ip>>
Storage node A iSCSI LIF01A network mask	<<var_nodeA_iscsi_lif01a_mask>>
Storage node A iSCSI LIF01B	<<var_nodeA_iscsi_lif01b_ip>>
Storage node A iSCSI LIF01B network mask	<<var_nodeA_iscsi_lif01b_mask>>
Storage node B iSCSI LIF01A	<<var_nodeB_iscsi_lif01a_ip>>
Storage node B iSCSI LIF01A network mask	<<var_nodeB_iscsi_lif01a_mask>>
Storage node B iSCSI LIF01B	<<var_nodeB_iscsi_lif01b_ip>>
Storage node B iSCSI LIF01B network mask	<<var_nodeB_iscsi_lif01b_mask>>

1. Create four iSCSI LIFs, two on each node.

```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface show

```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail Value
Storage node A NFS LIF 01 IP	<<var_nodeA_nfs_lif_01_ip>>
Storage node A NFS LIF 01 network mask	<<var_nodeA_nfs_lif_01_mask>>
Storage node B NFS LIF 02 IP	<<var_nodeB_nfs_lif_02_ip>>
Storage node B NFS LIF 02 network mask	<<var_nodeB_nfs_lif_02_mask>>

1. Create an NFS LIF.

```

network interface create -vserver Infra-SVM -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_ip>> -netmask <<
var_nodeA_nfs_lif_01_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_ip>> -netmask <<
var_nodeB_nfs_lif_02_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface show

```

Add infrastructure SVM administrator

The following table lists the information needed to complete this configuration.

Detail	Detail Value
Vsmgmt IP	<<var_svm_mgmt_ip>>
Vsmgmt network mask	<<var_svm_mgmt_mask>>
Vsmgmt default gateway	<<var_svm_mgmt_gateway>>

To add the infrastructure SVM administrator and SVM administration logical interface to the management network, complete the following steps:

1. Run the following command:

```

network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true

```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```

network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway
<<var_svm_mgmt_gateway>>
network route show

```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
security login unlock -username vsadmin -vserver Infra-SVM
```

[Next: Cisco UCS C-Series Rack Server Deployment Procedure](#)

Cisco UCS C-Series rack server deployment procedure

The following section provides a detailed procedure for configuring a Cisco UCS C-Series standalone rack server for use in the FlexPod Express configuration.

Perform initial Cisco UCS C-Series standalone server setup for Cisco Integrated Management Server

Complete these steps for the initial setup of the CIMC interface for Cisco UCS C-Series standalone servers.

The following table lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.

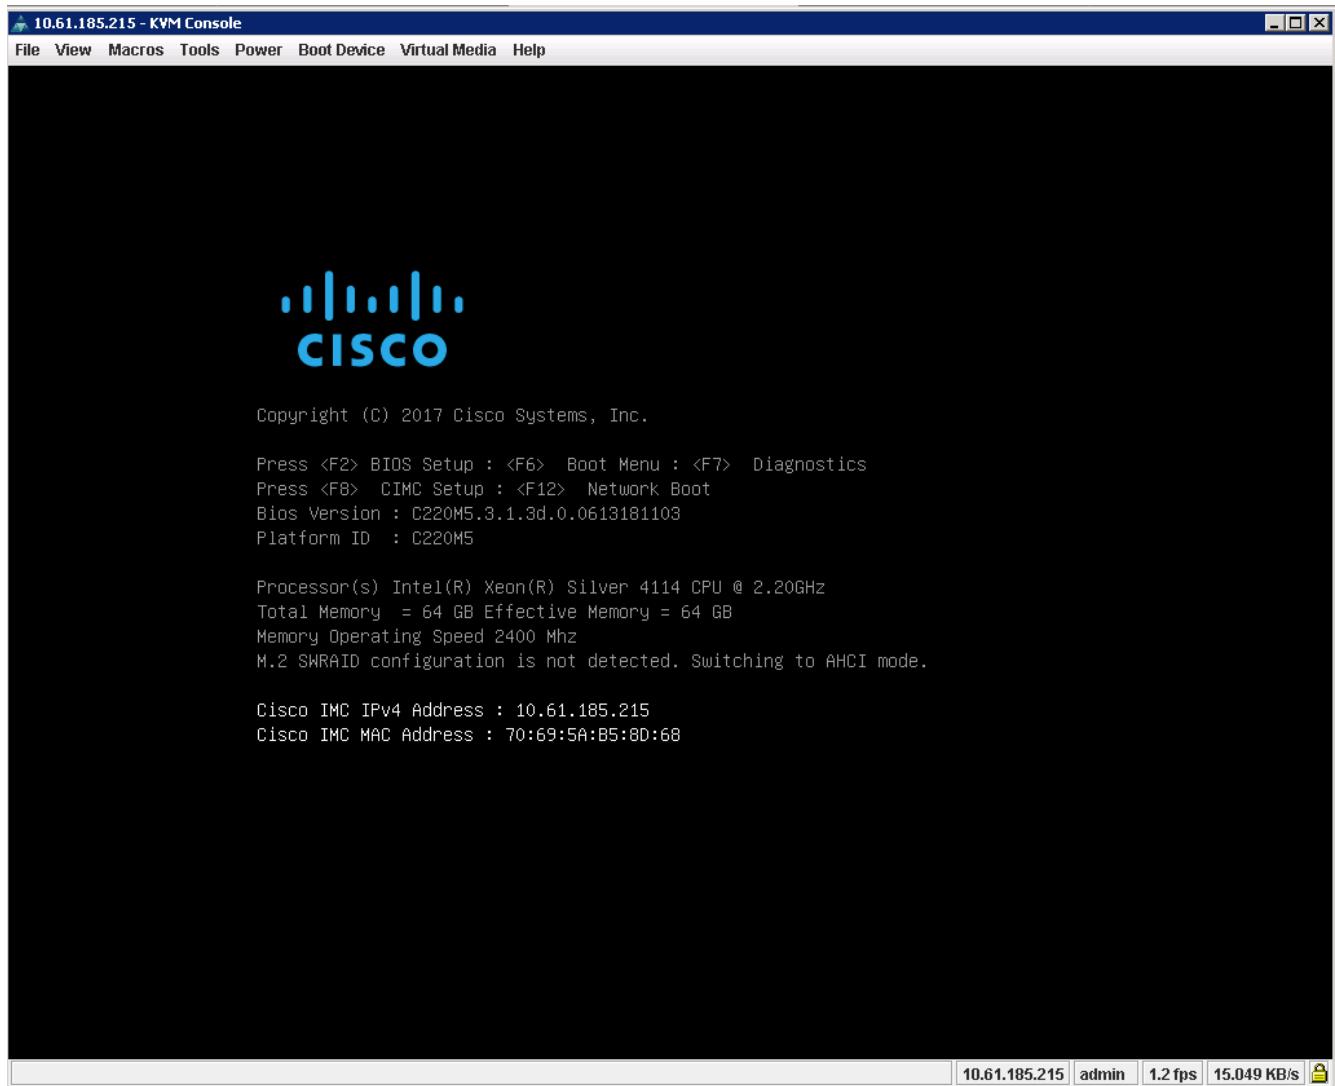
Detail	Detail value
CIMC IP address	<<cimc_ip>>
CIMC subnet mask	<<cimc_netmask>>
CIMC default gateway	<<cimc_gateway>>



The CIMC version used in this validation is CIMC 3.1.3(g).

All servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.
2. Power on the server and press F8 when prompted to enter the CIMC configuration.



10.61.185.215 admin 1.2 fps 15.049 KB/s

3. In the CIMC configuration utility, set the following options:

- Network interface card (NIC) mode:
 - Dedicated
- IP (Basic):
 - IPV4:
 - DHCP enabled:
 - CIMC IP: <<cimc_ip>>
 - Prefix/Subnet: <<cimc_netmask>>
 - Gateway: <<cimc_gateway>>
- VLAN (Advanced): Leave cleared to disable VLAN tagging.
 - NIC redundancy
 - None:

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated: [X] NIC redundancy
Shared LOM: [ ] None: [X]
Cisco Card:
  Riser1: [ ] Active-standby: [ ]
  Riser2: [ ] Active-active: [ ]
  MLom: [ ] VLAN (Advanced)
  Shared LOM Ext: [ ] VLAN enabled: [ ]
                                         VLAN ID: 1
                                         Priority: 0
IP (Basic)
  IPV4: [X]     IPV6: [ ]
  DHCP enabled: [ ]
  CIMC IP: 10.61.185.215
  Prefix/Subnet: 255.255.255.0
  Gateway: 10.61.185.1
  Pref DNS Server: 0.0.0.0
Smart Access USB
  Enabled: [ ]
*****
<Up/Down>Selection <F10>Save <Space>Enable/Disable <F5>Refresh <ESC>Exit
<F1>Additional settings

```

4. Press F1 to see additional settings.

- Common properties:
 - Host name: <<esxi_host_name>>
 - Dynamic DNS: []
 - Factory defaults: Leave cleared.
- Default user (basic):
 - Default password: <<admin_password>>
 - Reenter password: <<admin_password>>
 - Port properties: Use default values.
 - Port profiles: Leave cleared.

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
Common Properties
  Hostname: CIMC-Tiger-02
  Dynamic DNS: [X]
  DDNS Domain:
  FactoryDefaults
    Factory Default: [ ]
  Default User(Basic)
    Default password: -
    Reenter password:
  Port Properties
    Auto Negotiation: [X]
      Admin Mode          Operation Mode
      Speed[1000/100/10Mbps]: Auto          1000
      Duplex mode[half/full]: Auto          full
  Port Profiles
    Reset: [ ]
    Name:

*****  

<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit  

<F2>PreviousPagegettings

```

5. Press F10 to save the CIMC interface configuration.
6. After the configuration is saved, press Esc to exit.

Configure Cisco UCS C-Series servers iSCSI boot

In this FlexPod Express configuration, the VIC1387 is used for iSCSI boot.

The following table lists the information needed to configure iSCSI boot.



Italicized font indicates variables that are unique for each ESXi host.

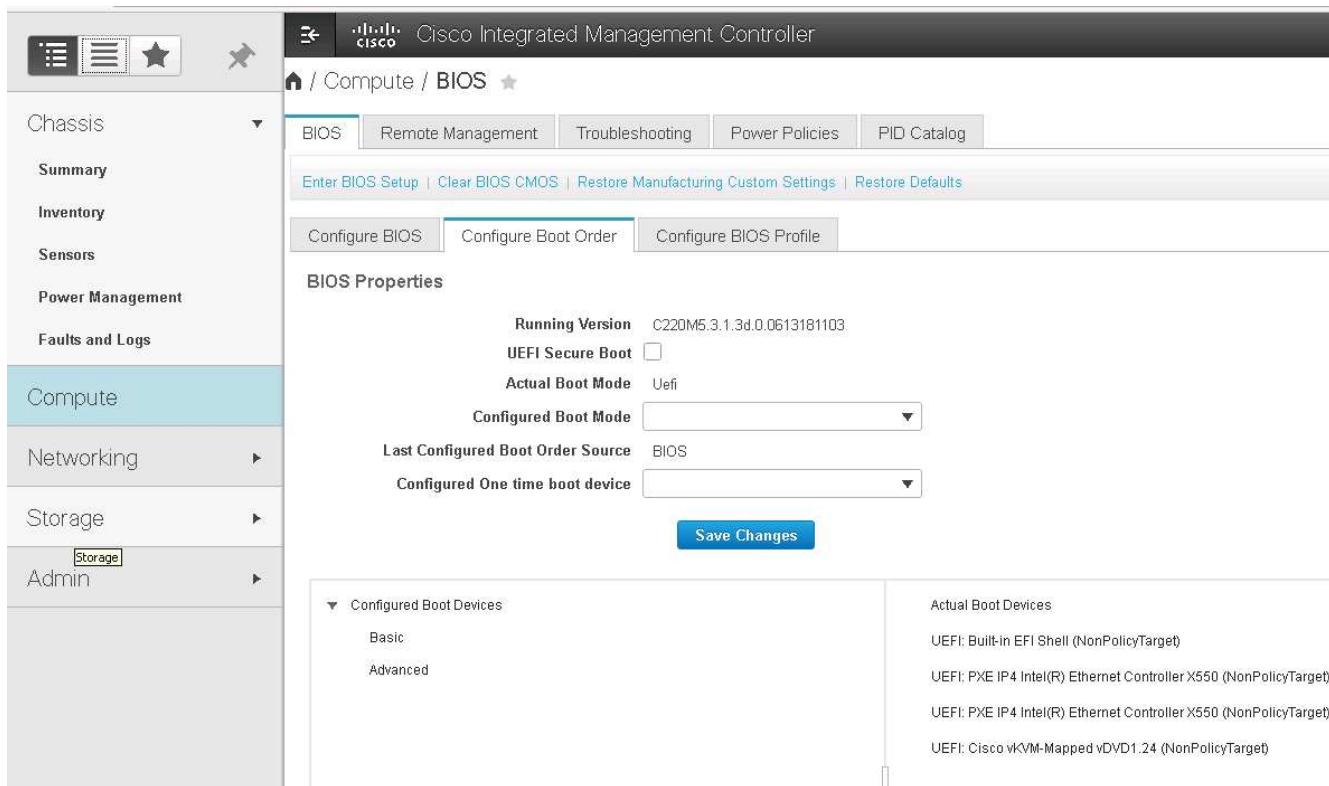
Detail	Detail value
ESXi host initiator A name	<<var_ucs_initiator_name_A>>
ESXi host iSCSI-A IP	<<var_esxi_host_iscsiA_ip>>
ESXi host iSCSI-A network mask	<<var_esxi_host_iscsiA_mask>>
ESXi host iSCSI A default gateway	<<var_esxi_host_iscsiA_gateway>>
ESXi host initiator B name	<<var_ucs_initiator_name_B>>
ESXi host iSCSI-B IP	<<var_esxi_host_iscsiB_ip>>
ESXi host iSCSI-B network mask	<<var_esxi_host_iscsiB_mask>>
ESXi host iSCSI-B gateway	<<var_esxi_host_iscsiB_gateway>>

Detail	Detail value
IP address iscsi_lif01a	
IP address iscsi_lif02a	
IP address iscsi_lif01b	
IP address iscsi_lif02b	
Infra_SVM IQN	

Boot order configuration

To set the boot order configuration, complete the following steps:

1. From the CIMC interface browser window, click the Server tab and select BIOS.
2. Click Configure Boot Order and then click OK.



3. Configure the following devices by clicking the device under Add Boot Device, and going to the Advanced tab.
 - Add Virtual Media
 - Name: KVM-CD-DVD
 - Subtype: KVM MAPPED DVD
 - State: Enabled
 - Order: 1
 - Add iSCSI Boot.
 - Name: iSCSI-A

- State: Enabled
- Order: 2
- Slot: MLOM
- Port: 0
- Click Add iSCSI Boot.
 - Name: iSCSI-B
 - State: Enabled
 - Order: 3
 - Slot: MLOM
 - Port: 1

4. Click Add Device.

5. Click Save Changes and then click Close.

Advanced Boot Order Configuration					Selected 1 / Total 3	⚙
		Name	Type	Order	State	
<input checked="" type="checkbox"/>	KVM-MAPPED-DVD	VMEDIA	1	Enabled		
<input type="checkbox"/>	iSCSI-A	ISCSI	2	Enabled		
<input type="checkbox"/>	iSCSI-B	ISCSI	3	Enabled		

6. Reboot the server to boot with your new boot order.

Disable RAID controller (if present)

Complete the following steps if your C-Series server contains a RAID controller. A RAID controller is not needed in the boot from SAN configuration. Optionally, you can also physically remove the RAID controller from the server.

1. Click BIOS on the left navigation pane in CIMC.
2. Select Configure BIOS.
3. Scroll down to PCIe Slot:HBA Option ROM.
4. If the value is not already disabled, set it to disabled.

BIOS	Remote Management	Troubleshooting	Power Policies	PID Catalog
I/O	Server Management	Security	Processor	Memory

Note: Default values are shown in bold.

Reboot Host Immediately:

Intel VT for directed IO:	Enabled
Intel VTD ATS support:	Enabled
LOM Port 1 OptionRom:	Enabled
Pcie Slot 1 OptionRom:	Disabled
MLOM OptionRom:	Enabled
Front NVME 1 OptionRom:	Enabled
MRAID Link Speed:	Auto
PCIe Slot 1 Link Speed:	Auto
Front NVME 1 Link Speed:	Auto
VGA Priority:	Onboard
P-SATA OptionROM:	LSI SW RAID
USB Port Rear:	Enabled
USB Port Internal:	Enabled
IPv6 PXE Support:	Disabled

Legacy USB Support:	Enabled
Intel VTD coherency support:	Disabled
All Onboard LOM Ports:	Enabled
LOM Port 2 OptionRom:	Enabled
Pcie Slot 2 OptionRom:	Disabled
MRAID OptionRom:	Enabled
Front NVME 2 OptionRom:	Enabled
MLOM Link Speed:	Auto
PCIe Slot 2 Link Speed:	Auto
Front NVME 2 Link Speed:	Auto
M.2 SATA OptionROM:	AHCI
USB Port Front:	Enabled
USB Port KVM:	Enabled
USB Port:M.2 Storage:	Enabled

Configure Cisco VIC1387 for iSCSI boot

The following configuration steps are for the Cisco VIC 1387 for iSCSI boot.

Create iSCSI vNICs

1. Click Add to create a vNIC.
2. In the Add vNIC section, enter the following settings:
 - Name: iSCSI-vNIC-A
 - MTU: 9000
 - Default VLAN: <>
 - VLAN Mode: TRUNK
 - Enable PXE boot: Check

▼ vNIC Properties

▼ General

Name:	iSCSI-vNIC-A	VLAN Mode:	Trunk
CDN:	VIC-MLOM-iSCSI-vNIC-A	Rate Limit:	<input checked="" type="radio"/> OFF
MTU:	9000	Channel Number:	N/A (1 - 1000)
Uplink Port:	0	PCI Link:	0 (0 - 1)
MAC Address:	<input type="radio"/> Auto	Enable NVGRE:	<input type="checkbox"/>
	<input checked="" type="radio"/> 70:69:5A:C0:98:ED	Enable VXLAN:	<input type="checkbox"/>
Class of Service:	0 (0 - 6)	Advanced Filter:	<input type="checkbox"/>
Trust Host CoS:	<input checked="" type="checkbox"/>	Port Profile:	N/A
PCI Order:	4 (0 - 5)	Enable PXE Boot:	<input checked="" type="checkbox"/>
Default VLAN:	<input type="radio"/> None	Enable VMQ:	<input type="checkbox"/>
	<input checked="" type="radio"/> 3439	Enable aRFS:	<input type="checkbox"/>
		Enable Uplink Failover:	<input type="checkbox"/>
		Fallback Timeout:	N/A (0 - 600)

3. Click Add vNIC and then click OK.
4. Repeat the process to add a second vNIC.
 - a. Name the vNIC iSCSI-vNIC-B.
 - b. Enter <<var_iscsi_vlan_b>> as the VLAN.
 - c. Set the uplink port to 1.
5. Select the vNIC iSCSI-vNIC-A on the left.

The screenshot shows the vNICs configuration interface for Adapter Card MLOM. The 'vNICs' tab is selected. On the left, a list shows vNICs: eth0, eth1, iSCSI-vNIC-A (selected and highlighted in blue), and iSCSI-vNIC-B. On the right, the 'iSCSI Boot Properties' section is expanded, showing 'General', 'Initiator', 'Primary Target', 'Secondary Target', and a 'Unconfigure iSCSI Boot' button. The 'usNIC' section is also visible.

6. Under iSCSI Boot Properties, enter the initiator details:
 - Name: <<var_ucsA_initiator_name_a>>
 - IP address: <<var_esxi_hostA_iscsiA_ip>>
 - Subnet mask: <<var_esxi_hostA_iscsiA_mask>>
 - Gateway: <<var_esxi_hostA_iscsiA_gateway>>

7. Enter the primary target details.

- Name: IQN number of infra-SVM
- IP address: IP address of `iscsi_lif01a`
- Boot LUN: 0

8. Enter the secondary target details.

- Name: IQN number of infra-SVM
- IP address: IP address of `iscsi_lif02a`
- Boot LUN: 0

You can obtain the storage IQN number by running the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step.

The screenshot shows the 'vNICs' tab selected in the navigation bar. On the left, a list of vNICs includes 'eth0', 'eth1', and 'iSCSI-v'. The 'iSCSI-v' entry is highlighted with a blue selection bar. The main panel displays 'Initiator' settings for two targets. For the 'Primary Target', the 'Name' is 'iqn.1992-08.com.netapp:sn.7e560f73a51' and the 'IP Address' is '172.21.246.16'. For the 'Secondary Target', the 'Name' is 'iqn.1992-08.com.netapp:sn.7e560f73a51' and the 'IP Address' is '172.21.246.18'. Both targets have a 'TCP Port' of '3260'. To the right, there are fields for 'Boot LUN' (set to '0'), 'CHAP Name', and 'CHAP Secret' for both targets. A blue 'Unconfigure iSCSI Boot' button is located at the bottom of the panel.

9. Click Configure iSCSI.
10. Select the vNIC **iSCSI-vNIC-B** and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.
11. Repeat the process to configure **iSCSI-vNIC-B**.
12. Enter the initiator details.
 - Name: <<var_ucsa_initiator_name_b>>
 - IP address: <<var_esxi_hostb_iscsib_ip>>
 - Subnet mask: <<var_esxi_hostb_iscsib_mask>>
 - Gateway: <<var_esxi_hostb_iscsib_gateway>>
13. Enter the primary target details.
 - Name: IQN number of infra-SVM
 - IP address: IP address of **iscsi_lif01b**
 - Boot LUN: 0
14. Enter the secondary target details.
 - Name: IQN number of infra-SVM
 - IP address: IP address of **iscsi_lif02b**
 - Boot LUN: 0

You can obtain the storage IQN number by using the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step.

15. Click Configure ISCSI.
16. Repeat this process to configure iSCSI boot for Cisco UCS server B.

Configure vNICs for ESXi

1. From the CIMC interface browser window, click Inventory and then click Cisco VIC adapters on the right pane.
2. Under Adapter Cards, select Cisco UCS VIC 1387 and then select the vNICs underneath.

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode
<input type="checkbox"/>	eth0	VIC-MLO...	70:69:5A:C0:98:49	1500	0	0	0	NONE	TRUNK
<input type="checkbox"/>	eth1	VIC-MLO...	70:69:5A:C0:98:4A	1500	0	1	0	NONE	TRUNK
<input type="checkbox"/>	iSCSI-v...	VIC-MLO...	70:69:5A:C0:98:4D	9000	0	0	0	3439	TRUNK
<input type="checkbox"/>	iSCSI-v...	VIC-MLO...	70:69:5A:C0:98:4E	9000	0	1	0	3440	TRUNK

3. Select eth0 and click Properties.
4. Set the MTU to 9000. Click Save Changes.

General External Ethernet Interfaces **vNICs** vHBAs

vNICs

eth0	Name: eth0
eth1	CDN: VIC-MLOM-eth0
iSCSI-v	MTU: 9000 (1500 - 9000)
iSCSI-v	Uplink Port: 0
	MAC Address: <input checked="" type="radio"/> Auto
	<input checked="" type="radio"/> 70:69:5A:C0:98:49
	Class of Service: 0 (0 - 6)
	Trust Host CoS: <input type="checkbox"/>
	PCI Order: 0 (0 - 5)
	Default VLAN: <input checked="" type="radio"/> None
	<input type="radio"/> ?

5. Repeat steps 3 and 4 for eth1, verifying that the uplink port is set to 1 for eth1.

1 / ... / Adapter Card MLOM / vNICs ★

General External Ethernet Interfaces **vNICs** vHBAs

vNICs

Host Ethernet Interfaces

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port
<input type="checkbox"/>	eth0	VIC-MLO...	70:69:5A:C0:98:49	9000	0	0
<input type="checkbox"/>	eth1	VIC-MLO...	70:69:5A:C0:98:4A	9000	0	1
<input type="checkbox"/>	iSCSI-v...	VIC-MLO...	70:69:5A:C0:98:4D	9000	0	0
<input type="checkbox"/>	iSCSI-v...	VIC-MLO...	70:69:5A:C0:98:4E	9000	0	1



This procedure must be repeated for each initial Cisco UCS Server node and each additional Cisco UCS Server node added to the environment.

Next: [NetApp AFF Storage Deployment Procedure \(Part 2\)](#)

NetApp AFF Storage Deployment Procedure (Part 2)

ONTAP SAN boot storage setup

Create iSCSI igroups

To create igroups, complete the following step:

You need the iSCSI initiator IQNs from the server configuration for this step.

1. From the cluster management node SSH connection, run the following commands. To view the three igroups created in this step, run the igrup show command.

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-A -protocol iscsi  
-ostype vmware -initiator <<var_vm_host_infra_a_iSCSI-A_vNIC_IQN>>,  
<<var_vm_host_infra_a_iSCSI-B_vNIC_IQN>>  
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-B -protocol iscsi  
-ostype vmware -initiator <<var_vm_host_infra_b_iSCSI-A_vNIC_IQN>>,  
<<var_vm_host_infra_b_iSCSI-B_vNIC_IQN>>
```



This step must be completed when adding additional Cisco UCS C- Series servers.

Map boot LUNs to igroups

To map boot LUNs to igroups, run the following commands from the cluster management SSH connection:

```
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- A -igroup  
VM-Host-Infra- A -lun-id 0  
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- B -igroup  
VM-Host-Infra- B -lun-id 0
```



This step must be completed when adding additional Cisco UCS C-Series servers.

Next: [VMware vSphere 6.7 Deployment Procedure.](#)

VMware vSphere 6.7 deployment procedure

This section provides detailed procedures for installing VMware ESXi 6.7 in a FlexPod Express configuration. The deployment procedures that follow are customized to include the environment variables described in previous sections.

Multiple methods exist for installing VMware ESXi in such an environment. This procedure uses the virtual KVM console and virtual media features of the CIMC interface for Cisco UCS C-Series servers to map remote installation media to each individual server.



This procedure must be completed for Cisco UCS server A and Cisco UCS server B.

This procedure must be completed for any additional nodes added to the cluster.

Log in to CIMC interface for Cisco UCS C-Series standalone servers

The following steps detail the method for logging in to the CIMC interface for Cisco UCS C-Series standalone servers. You must log in to the CIMC interface to run the virtual KVM, which enables the administrator to begin installation of the operating system through remote media.

All hosts

1. Navigate to a web browser and enter the IP address for the CIMC interface for the Cisco UCS C-Series. This step launches the CIMC GUI application.
2. Log in to the CIMC UI using the admin user name and credentials.
3. In the main menu, select the Server tab.
4. Click Launch KVM Console.



5. From the virtual KVM console, select the Virtual Media tab.
6. Select Map CD/DVD.



You might first need to click Activate Virtual Devices. Select Accept This Session if prompted.

7. Browse to the VMware ESXi 6.7 installer ISO image file and click Open. Click Map Device.
8. Select the Power menu and choose Power Cycle System (Cold Boot). Click Yes.

Install VMware ESXi

The following steps describe how to install VMware ESXi on each host.

Download ESXI 6.7 Cisco custom image

1. Navigate to the [VMware vSphere download page](#) for custom ISOs.
2. Click Go to Downloads next to the Cisco Custom Image for ESXi 6.7 GA Install CD.
3. Download the Cisco Custom Image for ESXi 6.7 GA Install CD (ISO).

All hosts

1. When the system boots, the machine detects the presence of the VMware ESXi installation media.
2. Select the VMware ESXi installer from the menu that appears.

The installer loads. This takes several minutes.

3. After the installer has finished loading, press Enter to continue with the installation.

4. After reading the end-user license agreement, accept it and continue with the installation by pressing F11.
5. Select the NetApp LUN that was previously set up as the installation disk for ESXi, and press Enter to continue with the installation.



6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer warns you that existing partitions are removed on the volume. Continue with the installation by pressing F11. The server reboots after the installation of ESXi.

Set up VMware ESXi host management networking

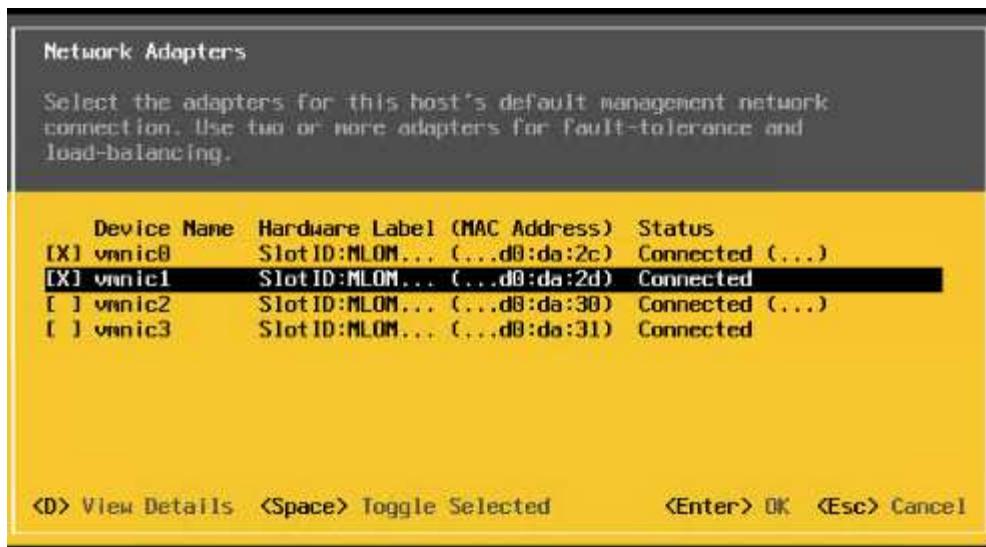
The following steps describe how to add the management network for each VMware ESXi host.

All hosts

1. After the server has finished rebooting, enter the option to customize the system by pressing F2.
2. Log in with root as the login name and the root password previously entered during the installation process.
3. Select the Configure Management Network option.
4. Select Network Adapters and press Enter.
5. Select the desired ports for vSwitch0. Press Enter.



Select the ports that correspond to eth0 and eth1 in CIMC.



6. Select VLAN (optional) and press Enter.
7. Enter the VLAN ID <<mgmt_vlan_id>>. Press Enter.
8. From the Configure Management Network menu, select IPv4 Configuration to configure the IP address of the management interface. Press Enter.
9. Use the arrow keys to highlight Set Static IPv4 address and use the space bar to select this option.
10. Enter the IP address for managing the VMware ESXi host <<esxi_host_mgmt_ip>>.
11. Enter the subnet mask for the VMware ESXi host <<esxi_host_mgmt_netmask>>.
12. Enter the default gateway for the VMware ESXi host <<esxi_host_mgmt_gateway>>.
13. Press Enter to accept the changes to the IP configuration.
14. Enter the IPv6 configuration menu.
15. Use the space bar to disable IPv6 by unselecting the Enable IPv6 (restart required) option. Press Enter.
16. Enter the menu to configure the DNS settings.
17. Because the IP address is assigned manually, the DNS information must also be entered manually.
18. Enter the primary DNS server's IP address [\[nameserver_ip\]](#).
19. (Optional) Enter the secondary DNS server's IP address.
20. Enter the FQDN for the VMware ESXi host name: [\[esxi_host_fqdn\]](#).
21. Press Enter to accept the changes to the DNS configuration.
22. Exit the Configure Management Network submenu by pressing Esc.
23. Press Y to confirm the changes and reboot the server.
24. Log out of the VMware Console by pressing Esc.

Configure ESXi host

You need the information in the following table to configure each ESXi host.

Detail	Value
ESXi host name	

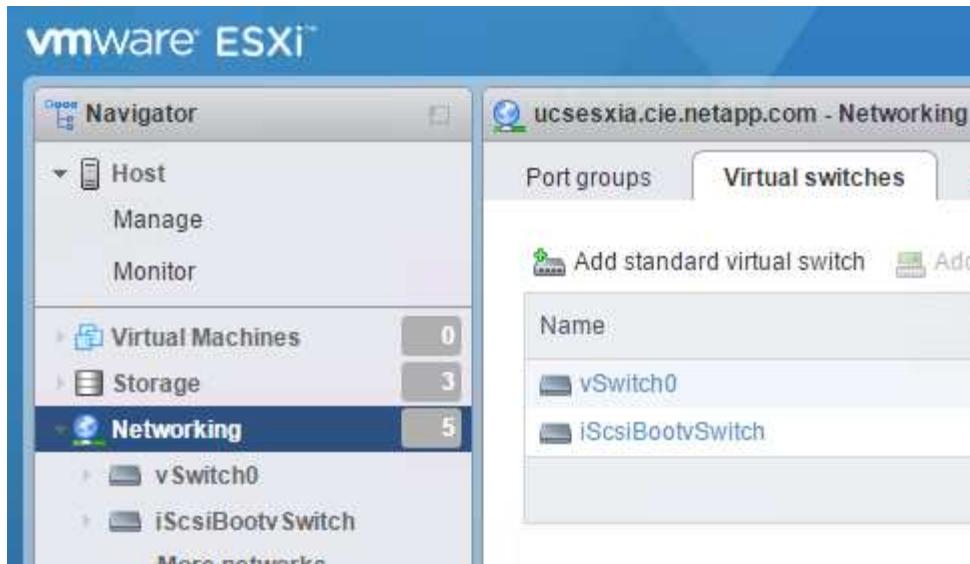
Detail	Value
ESXi host management IP	
ESXi host management mask	
ESXi host management gateway	
ESXi host NFS IP	
ESXi host NFS mask	
ESXi host NFS gateway	
ESXi host vMotion IP	
ESXi host vMotion mask	
ESXi host vMotion gateway	
ESXi host iSCSI-A IP	
ESXi host iSCSI-A mask	
ESXi host iSCSI-A gateway	
ESXi host iSCSI-B IP	
ESXi host iSCSI-B mask	
ESXi host iSCSI-B gateway	

Log in to ESXi host

1. Open the host's management IP address in a web browser.
2. Log in to the ESXi host using the root account and the password you specified during the install process.
3. Read the statement about the VMware Customer Experience Improvement Program. After selecting the proper response, click OK.

Configure iSCSI boot

1. Select Networking on the left.
2. On the right, select the Virtual Switches tab.



3. Click iScsiBootvSwitch.
4. Select Edit settings.
5. Change the MTU to 9000 and click Save.
6. Click Networking in the left navigation pane to return to the Virtual Switches tab.
7. Click Add Standard Virtual Switch.
8. Provide the name `iScsiBootvSwitch-B` for the vSwitch name.
 - Set the MTU to 9000.
 - Select `vmnic3` from the Uplink 1 options.
 - Click Add.



Vmnic2 and vmnic3 are used for iSCSI boot in this configuration. If you have additional NICs in your ESXi host, you might have different vmnic numbers. To confirm which NICs are used for iSCSI boot, match the MAC addresses on the iSCSI vNICs in CIMC to the vmnics in ESXi.

9. In the center pane, select the VMkernel NICs tab.
10. Select Add VMkernel NIC.
 - Specify a new port group name of `iScsiBootPG-B`.
 - Select `iScsiBootvSwitch-B` for the virtual switch.
 - Enter `<<iscsib_vlan_id>>` for the VLAN ID.
 - Change the MTU to 9000.
 - Expand IPv4 Settings.
 - Select Static Configuration.
 - Enter `<<var_hosta_iscsib_ip>>` for Address.
 - Enter `<<var_hosta_iscsib_mask>>` for Subnet Mask.
 - Click Create.

 Add VMkernel NIC

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	iScsiBootvSwitch-B
VLAN ID	3440
MTU	9000
IP version	IPv4 only
▼ IPv4 settings	
Configuration	<input checked="" type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.184.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input checked="" type="checkbox"/> vMotion <input checked="" type="checkbox"/> Provisioning <input checked="" type="checkbox"/> Fault tolerance logging <input checked="" type="checkbox"/> Management <input checked="" type="checkbox"/> Replication <input checked="" type="checkbox"/> NFC replication

Create Cancel

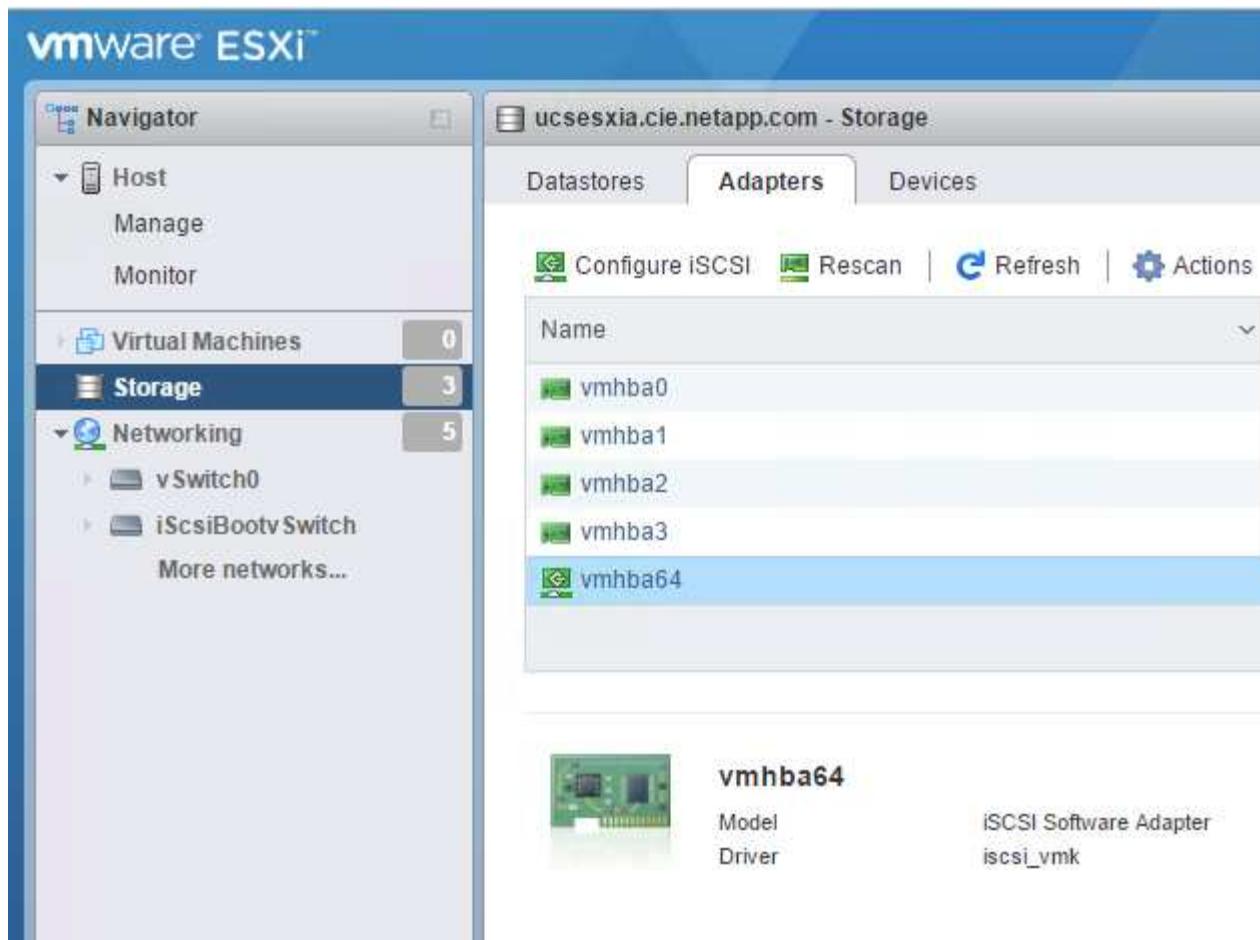


Set the MTU to 9000 on iScsiBootPG- A.

Configure iSCSI multipathing

To set up iSCSI multipathing on the ESXi hosts, complete the following steps:

1. Select Storage in the left navigation pane. Click Adapters.
2. Select the iSCSI software adapter and click Configure iSCSI.



3. Under Dynamic Targets, click Add Dynamic Target.

4. Enter the IP address `iscsi_lif01a`.

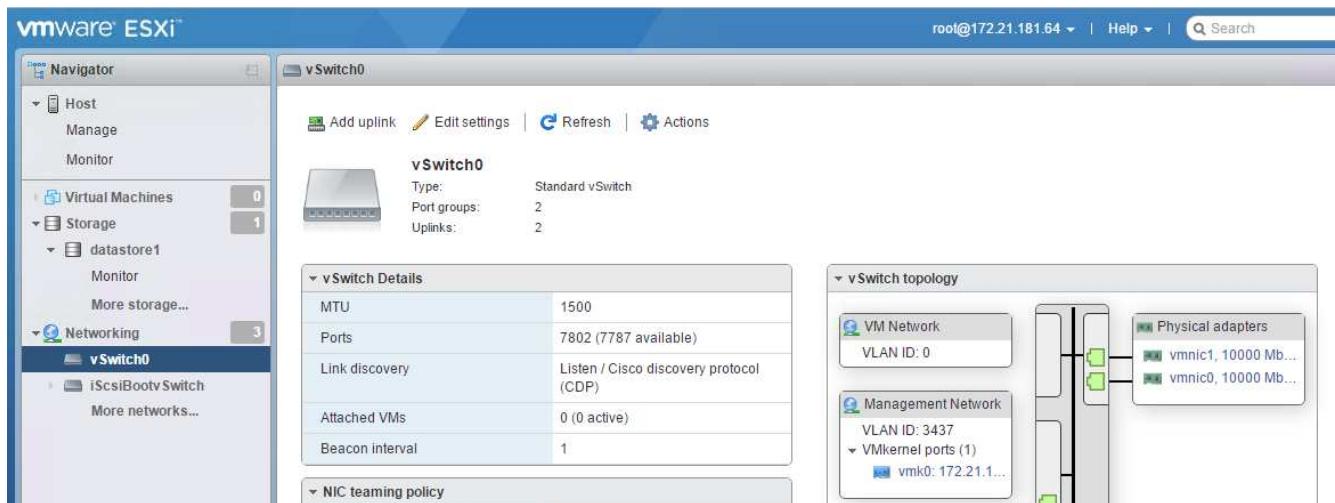
- Repeat with the IP addresses `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
- Click Save Configuration.



You can find the iSCSI LIF IP addresses by running the `network interface show` command on the NetApp cluster or by looking at the Network Interfaces tab in OnCommand System Manager.

Configure ESXi host

- In the left navigation pane, select Networking.
- Select vSwitch0.



3. Select Edit Settings.
4. Change the MTU to 9000.
5. Expand NIC Teaming and verify that both vmnic0 and vmnic1 are set to active.

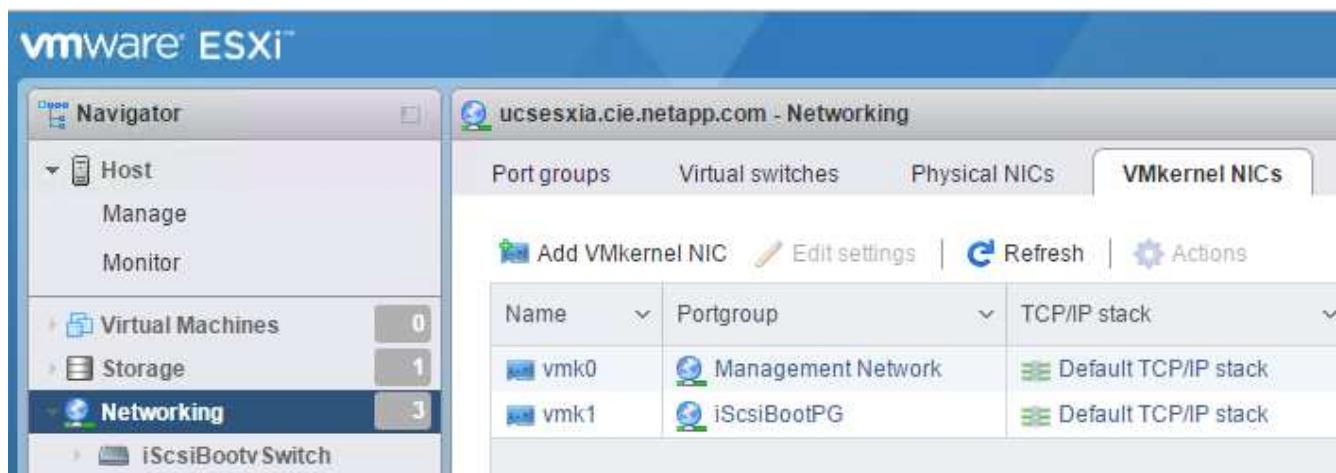
Configure port groups and VMkernel NICs

1. In the left navigation pane, select Networking.
2. Right-click the Port Groups tab.



3. Right-click VM Network and select Edit. Change the VLAN ID to <<var_vm_traffic_vlan>>.
4. Click Add Port Group.
 - Name the port group MGMT-Network.
 - Enter <<mgmt_vlan>> for the VLAN ID.
 - Make sure that vSwitch0 is selected.
 - Click Add.

5. Click the VMkernel NICs tab.



The screenshot shows the VMware ESXi interface with the 'Networking' section selected in the Navigator. The 'VMkernel NICs' tab is active, displaying a table of existing VMkernel interfaces. The table has columns for Name, Portgroup, and TCP/IP stack. Two entries are listed: 'vmk0' associated with 'Management Network' and 'Default TCP/IP stack', and 'vmk1' associated with 'iScsiBootPG' and 'Default TCP/IP stack'. Action buttons for 'Add VMkernel NIC', 'Edit settings', 'Refresh', and 'Actions' are visible at the top of the table.

Name	Portgroup	TCP/IP stack
vmk0	Management Network	Default TCP/IP stack
vmk1	iScsiBootPG	Default TCP/IP stack

6. Select Add VMkernel NIC.

- Select New Port Group.
- Name the port group NFS-Network.
- Enter <<nfs_vlan_id>> for the VLAN ID.
- Change the MTU to 9000.
- Expand IPv4 Settings.
- Select Static Configuration.
- Enter <<var_hosta_nfs_ip>> for Address.
- Enter <<var_hosta_nfs_mask>> for Subnet Mask.
- Click Create.

Add VMkernel NIC

Port group	New port group
New port group	NFS-Network
Virtual switch	vSwitch0
VLAN ID	3438
MTU	9000
IP version	IPv4 only
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.182.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack

7. Repeat this process to create the vMotion VMkernel port.
8. Select Add VMkernel NIC.
 - a. Select New Port Group.
 - b. Name the port group vMotion.
 - c. Enter <<vmotion_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_vmotion_ip>> for Address.
 - h. Enter <<var_hosta_vmotion_mask>> for Subnet Mask.
 - i. Make sure that the vMotion checkbox is selected after IPv4 Settings.

 There are many ways to configure ESXi networking, including by using the VMware vSphere distributed switch if your licensing allows it. Alternative network configurations are supported in FlexPod Express if they are required to meet business requirements.

Add VMkernel NIC

Virtual switch	vSwitch0
VLAN ID	3441
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.185.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input checked="" type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication

Create **Cancel**

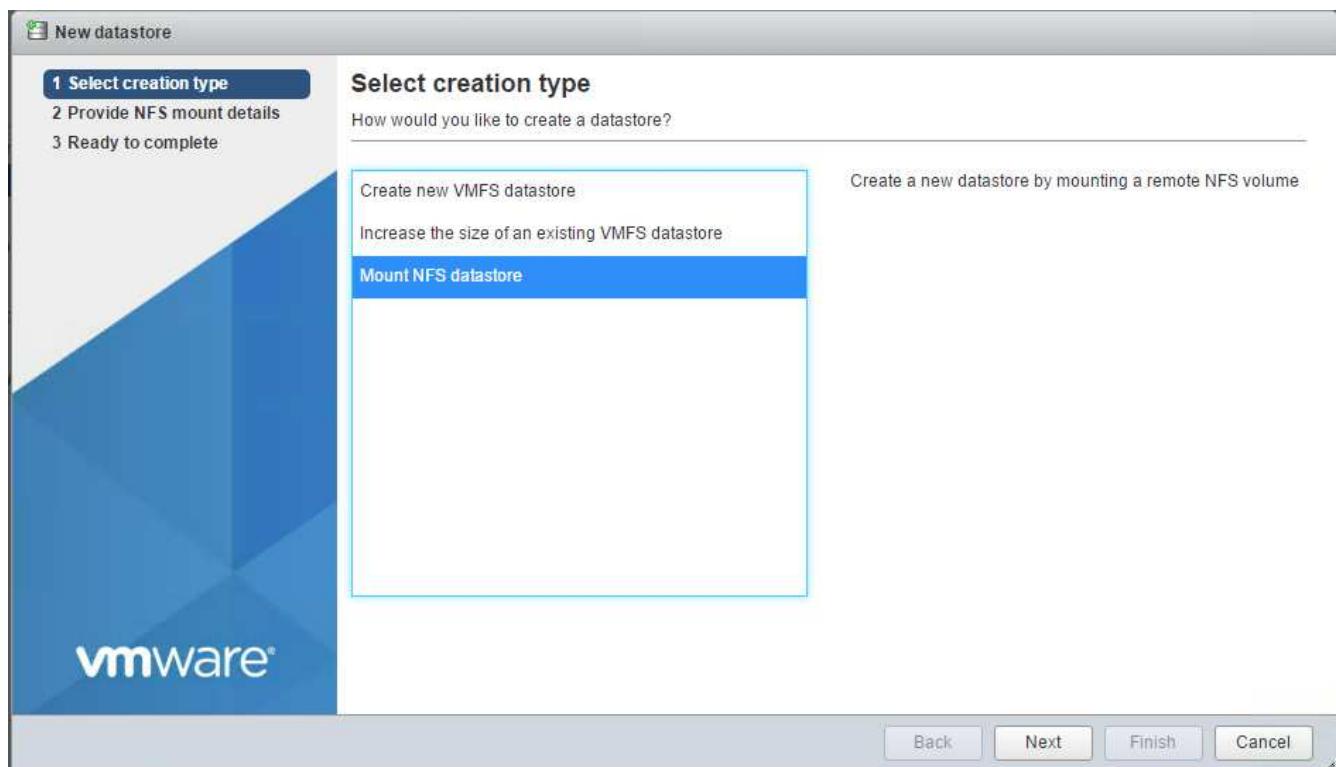
Mount first datastores

The first datastores to be mounted are the `infra_datastore_1` datastore for virtual machines and the `infra_swap` datastore for virtual machine swap files.

1. Click Storage in the left navigation pane, and then click New Datastore.



2. Select Mount NFS Datastore.



3. Next, enter the following information in the Provide NFS Mount Details page:

- Name: `infra_datastore_1`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_datastore_1`
- Make sure that NFS 3 is selected.

4. Click Finish. You can see the task completing in the Recent Tasks pane.

5. Repeat this process to mount the `infra_swap` datastore:

- Name: `infra_swap`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_swap`

- Make sure that NFS 3 is selected.

Configure NTP

To configure NTP for an ESXi host, complete the following steps:

1. Click Manage in the left navigation pane. Select System in the right pane and then click Time & Date.



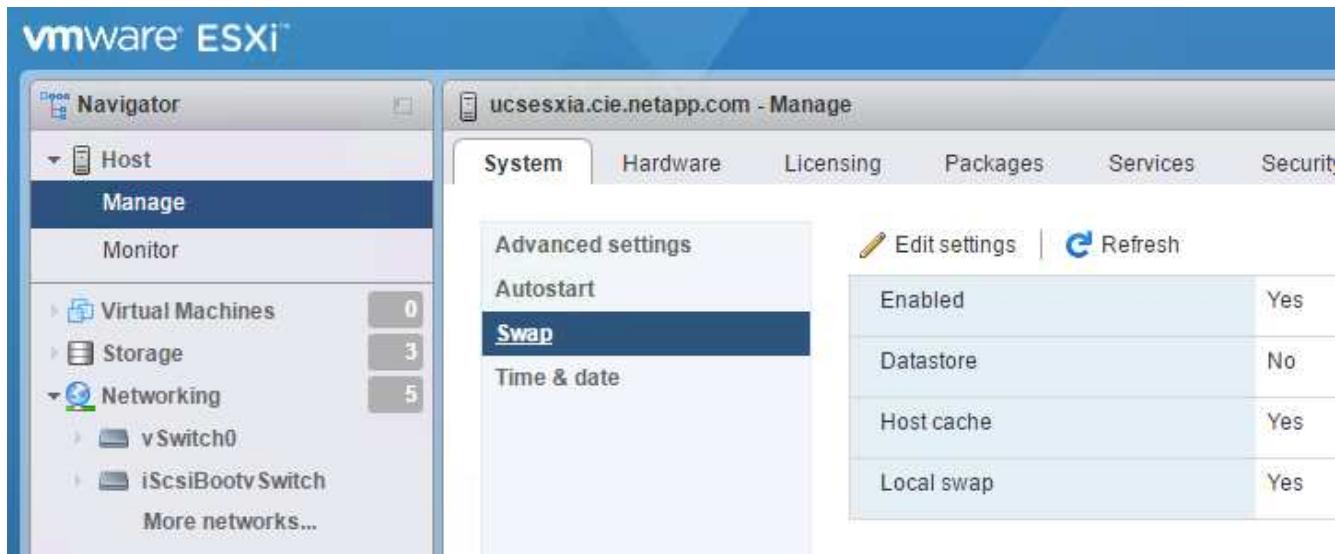
2. Select Use Network Time Protocol (Enable NTP Client).
3. Select Start and Stop with Host as the NTP service startup policy.
4. Enter <<var_ntp>> as the NTP server. You can set multiple NTP servers.
5. Click Save.



Move the virtual machine swap-file location

These steps provide details for moving the virtual machine swap-file location.

1. Click Manage in the left navigation pane. Select system in the right pane, then click Swap.



2. Click Edit Settings. Select infra_swap from the Datastore options.



3. Click Save.

Install the NetApp NFS Plug-in 1.0.20 for VMware VAAI

To install the NetApp NFS Plug-in 1.0.20 for VMware VAAI, complete the following steps.

1. Enter the following commands to verify that VAAI is enabled:

```
esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove  
esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
```

If VAAI is enabled, these commands produce the following output:

```
~ # esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
Value of HardwareAcceleratedMove is 1
~ # esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
Value of HardwareAcceleratedInit is 1
```

2. If VAAI is not enabled, enter the following commands to enable VAAI:

```
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedInit
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
```

These commands produce the following output:

```
~ # esxcfg-advcfg -s 1 /Data Mover/HardwareAcceleratedInit
Value of HardwareAcceleratedInit is 1
~ # esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
Value of HardwareAcceleratedMove is 1
```

3. Download the NetApp NFS Plug-in for VMware VAAI:

- Go to the [software download page](#).
- Scroll down and click NetApp NFS Plug-in for VMware VAAI.
- Select the ESXi platform.
- Download either the offline bundle (.zip) or online bundle (.vib) of the most recent plug-in.

4. Install the plug-in on the ESXi host by using the ESX CLI.
5. Reboot the ESXi host.

```
[root@vm-host-infra-04:~] ls /vmfs/volumes/datastore1/NetAppNasPlugin.vib
/vmfs/volumes/datastore1/NetAppNasPlugin.vib
[root@vm-host-infra-04:~] esxcli software vib install -v /vmfs/volumes/datastore1/NetAppNasPlugin.vib
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: NetApp_bootbank_NetAppNasPlugin_1.1.2-3
  VIBs Removed:
  VIBs Skipped:
[root@vm-host-infra-04:~] █
```

[Next: Install VMware vCenter Server 6.7](#)

Install VMware vCenter Server 6.7

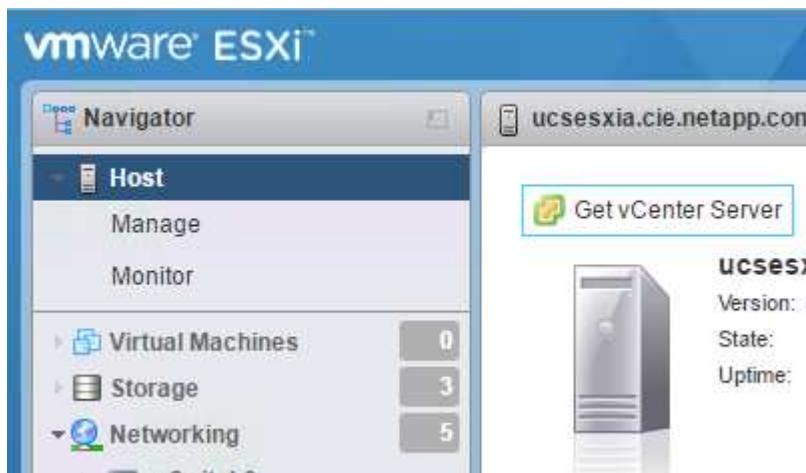
This section provides detailed procedures for installing VMware vCenter Server 6.7 in a FlexPod Express configuration.



FlexPod Express uses the VMware vCenter Server Appliance (VCSA).

Download the VMware vCenter server appliance

1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.



2. Download the VCSA from the VMware site.



Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.

3. Mount the ISO image.
4. Navigate to the vcsa-ui-installer> win32 directory. Double- click installer.exe.
5. Click Install.
6. Click Next on the Introduction page.
7. Accept the end- user license agreement.
8. Select Embedded Platform Services Controller as the deployment type.

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

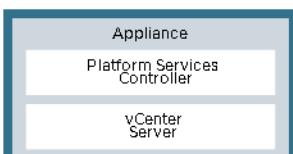
Select deployment type

Select the deployment type you want to configure on the appliance.

For more information on deployment types, refer to the vSphere 6.7 documentation.

Embedded Platform Services Controller

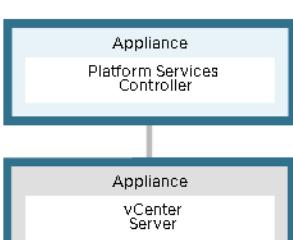
vCenter Server with an Embedded Platform Services Controller



External Platform Services Controller

Platform Services Controller

vCenter Server (Requires External Platform Services Controller)



CANCEL **BACK** **NEXT**



If required, the External Platform Services Controller deployment is also supported as part of the FlexPod Express solution.

9. In the Appliance Deployment Target, enter the IP address of an ESXi host you have deployed, and the root user name and root password.



- 1 Introduction
- 2 End user license agreement
- 3 Select deployment type
- 4 Appliance deployment target**
- 5 Set up appliance VM
- 6 Select deployment size
- 7 Select datastore
- 8 Configure network settings
- 9 Ready to complete stage 1

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

ESXi host or vCenter Server name	<input type="text" value="172.21.246.25"/> <small>(i)</small>
HTTPS port	<input type="text" value="443"/>
User name	<input type="text" value="root"/> <small>(i)</small>
Password	<input type="password" value="*****"/>

[CANCEL](#)[BACK](#)[NEXT](#)

10. Set the appliance VM by entering VCSA as the VM name and the root password you would like to use for the VCSA.

VMware vCenter Server Appliance 6.7 Installation - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name: tigervcsa

Set root password: (Info icon)

Confirm root password: (Info icon)

Cancel Back Next

1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

11. Select the deployment size that best fits your environment. Click Next.

VMware vCenter Server Appliance 6.7 Installation - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Select deployment size

Select the deployment size for this vCenter Server with an Embedded Platform Services Controller.

For more information on deployment sizes, refer to the vSphere 6.7 documentation.

Deployment size: Tiny

Storage size: Default

Resources required for different deployment sizes

Deployment Size	vCPUs	Memory (GB)	Storage (GB)	Hosts (up to)	VMs (up to)
Tiny	2	10	300	10	100
Small	4	16	340	100	1000
Medium	8	24	525	400	4000
Large	16	32	740	1000	10000
X-Large	24	48	1180	2000	35000

Cancel Back Next

1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

12. Select the `infra_datastore_1` datastore. Click Next.

vm Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Select datastore

Select the storage location for this appliance

Install on an existing datastore accessible from the target host

Name	Type	Capacity	Free	Provisioned	Thin Provisioning
infra_datastore_1	NFS	500 GB	499.98 GB	18.38 MB	Supported
infra_swap	NFS	100 GB	99.99 GB	10.95 MB	Supported

2 items

Enable Thin Disk Mode ⓘ

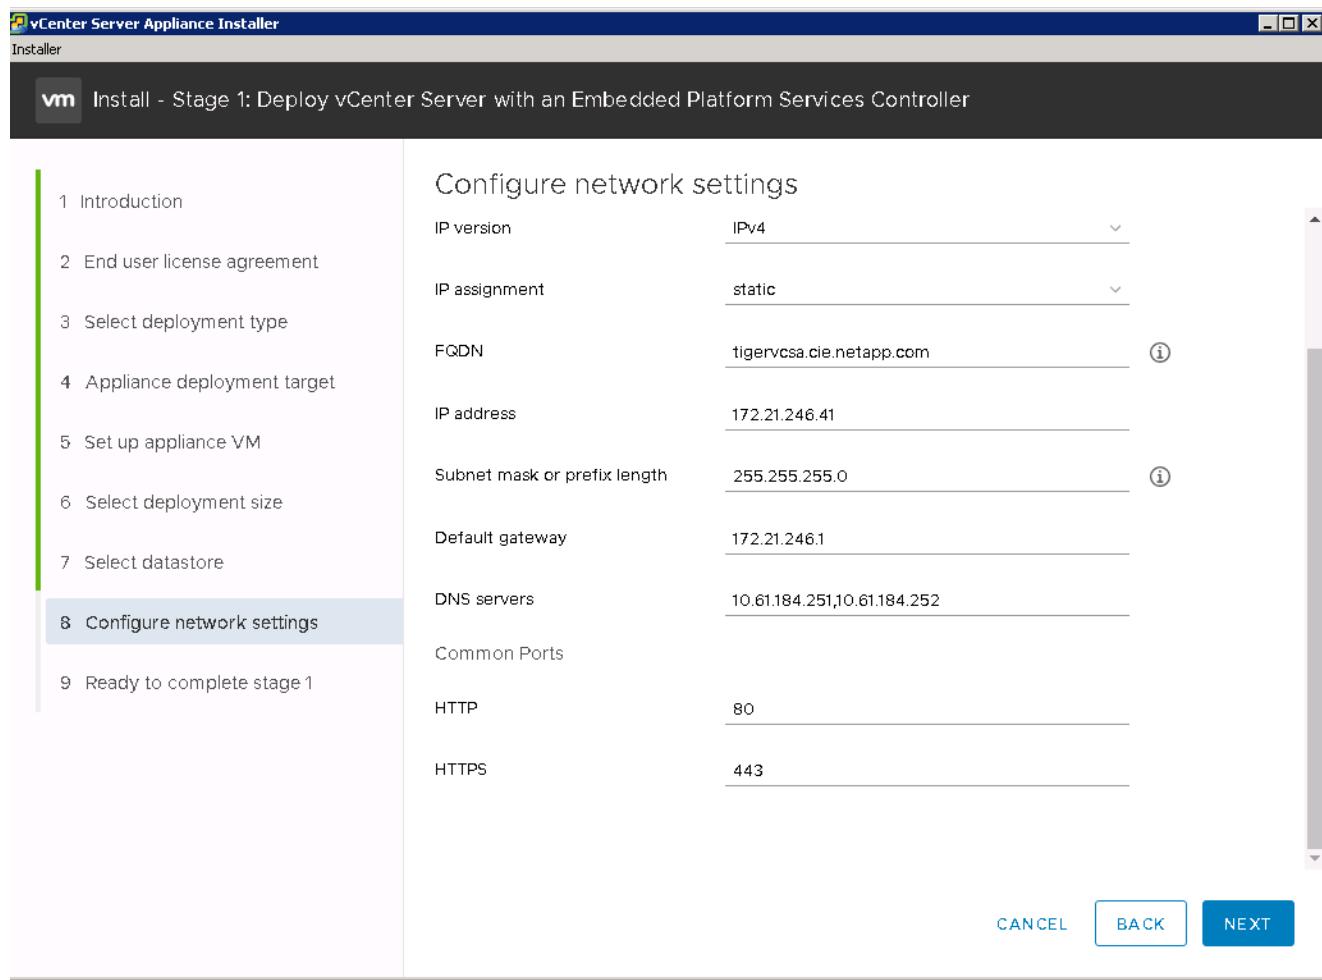
Install on a new vSAN cluster containing the target host ⓘ

CANCEL BACK NEXT

13. Enter the following information in the Configure network settings page and click Next.

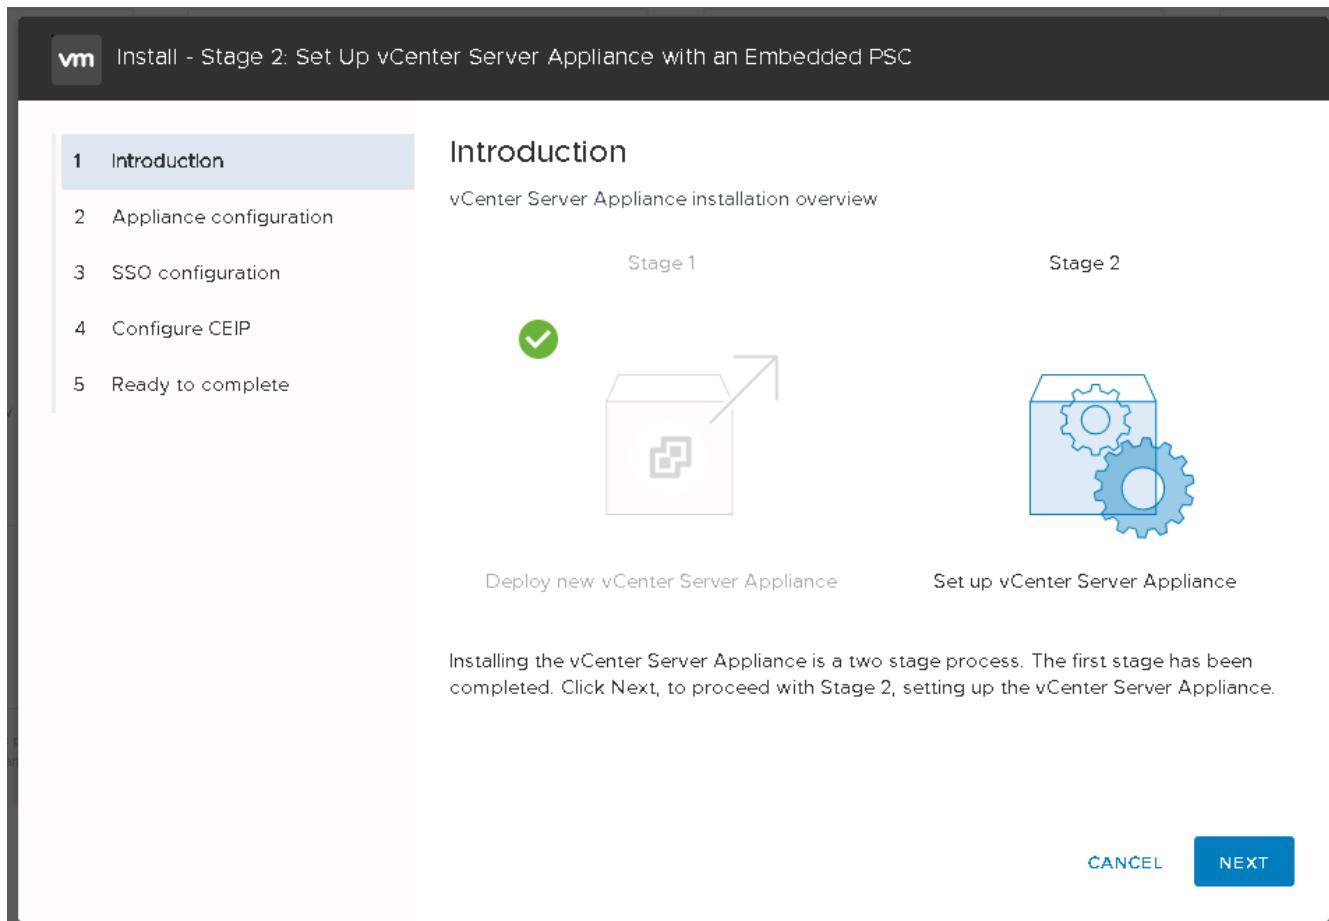
- Select MGMT-Network for Network.
- Enter the FQDN or IP to be used for the VCSA.
- Enter the IP address to be used.
- Enter the subnet mask to be used.
- Enter the default gateway.
- Enter the DNS server.

14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.



The VCSA installs now. This process takes several minutes.

15. After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.
16. On the Stage 2 Introduction page, click Next.



17. Enter <<var_ntp_id>> for the NTP server address. You can enter multiple NTP IP addresses.
If you plan to use vCenter Server high availability (HA), make sure that SSH access is enabled.
18. Configure the SSO domain name, password, and site name. Click Next.
Record these values for your reference, especially if you deviate from the vsphere.local domain name.
19. Join the VMware Customer Experience Program if desired. Click Next.
20. View the summary of your settings. Click Finish or use the back button to edit settings.
21. A message appears stating that you will not be able to pause or stop the installation from completing after it has started. Click OK to continue.
The appliance setup continues. This takes several minutes.
A message appears indicating that the setup was successful.
The links that the installer provides to access vCenter Server are clickable.

[Next: Configure VMware vCenter Server 6.7 and vSphere clustering.](#)

Configure VMware vCenter Server 6.7 and vSphere clustering

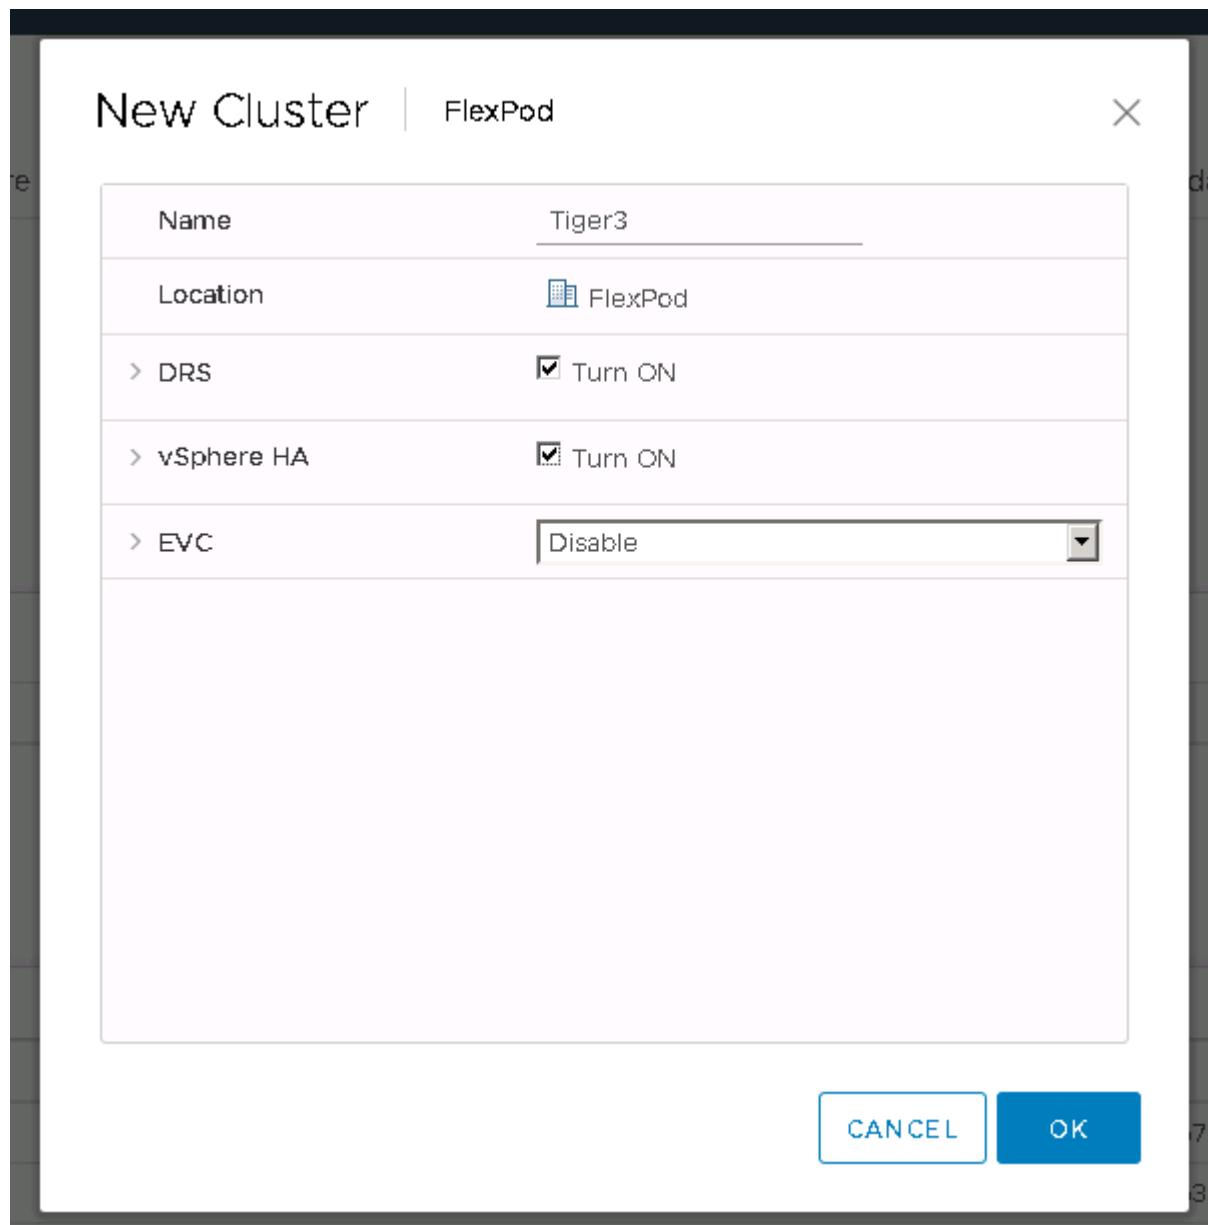
To configure VMware vCenter Server 6.7 and vSphere clustering, complete the following steps:

1. Navigate to <https://<<FQDN or IP of vCenter>>/vsphere-client/>.
2. Click Launch vSphere Client.
3. Log in with the user name administrator@vsphere.local and the SSO password you entered during the VCSA setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center and click OK.

Create vSphere cluster

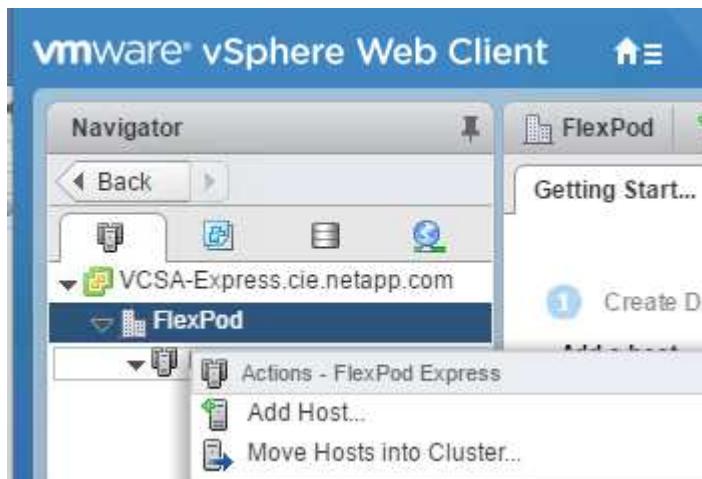
Complete the following steps to create a vSphere cluster:

1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Enable DR and vSphere HA by selecting the checkboxes.
4. Click OK.



Add ESXi hosts to cluster

1. Right-click the cluster and select Add Host.



2. To add an ESXi host to the cluster, complete the following steps:

- a. Enter the IP or FQDN of the host. Click Next.
- b. Enter the root user name and password. Click Next.
- c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
- d. Click Next on the Host Summary page.
- e. Click the green + icon to add a license to the vSphere host.



This step can be completed later if desired.

- f. Click Next to leave lockdown mode disabled.
- g. Click Next at the VM location page.
- h. Review the Ready to Complete page. Use the back button to make any changes or select Finish.

3. Repeat steps 1 and 2 for Cisco UCS host B. This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on ESXi hosts

1. Using SSH, connect to the management IP ESXi host, enter root for the user name, and enter the root password.
2. Run the following commands:

```
esxcli system coredump network set -i ip_address_of_core_dump_collector
-v vmk0 -o 6500
esxcli system coredump network set --enable=true
esxcli system coredump network check
```

3. The message **Verified the configured netdump server is running** appears after you enter the final command.

This process must be completed for any additional hosts added to FlexPod Express.

Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of additional components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed by keeping in mind small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp product documentation

<http://docs.netapp.com>

- FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 Design Guide

<https://www.netapp.com/us/media/nva-1125-design.pdf>

FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based Storage

NVA-1131-DEPLOY: FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based Storage

Sree Lakshmi Lanka, NetApp

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology with which they are familiar in their data center.

FlexPod Express is a predesigned, best practice architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal OSs and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the versatility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools to which they are accustomed. New FlexPod Express customers can easily adapt to managing FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote offices and branch offices (ROBOs) and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution Overview

This FlexPod Express solution is part of the FlexPod converged infrastructure program.

FlexPod Converged Infrastructure Program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

As depicted in the figure below, the FlexPod program includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express** offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter** delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select** incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.

The following figure shows the technical components of the solution.

FlexPod Express	FlexPod Datacenter	FlexPod Select
Departmental deployments and VAR velocity Target: Primarily MSB, remote, and departmental deployments	Massively scalable, all virtual Target: Enterprise/service provider	Application purposed Target: Specific application deployments in the enterprise
 Entry-level: Cisco UCS, Nexus, and NetApp FAS	 Cisco UCS, Cisco Nexus, and FAS	 NetApp E-Series and FAS

NetApp Verified Architecture Program

The NVA program offers customers a verified architecture for NetApp solutions. An NVA provides a NetApp solution architecture with the following qualities:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

This guide details the design of FlexPod Express with direct- attached NetApp storage. The following sections list the components used for the design of this solution.

Hardware components

- NetApp AFF A220
- Cisco UCS Mini
- Cisco UCS B200 M5
- Cisco UCS VIC 1440/1480.
- Cisco Nexus 3000 Series Switches

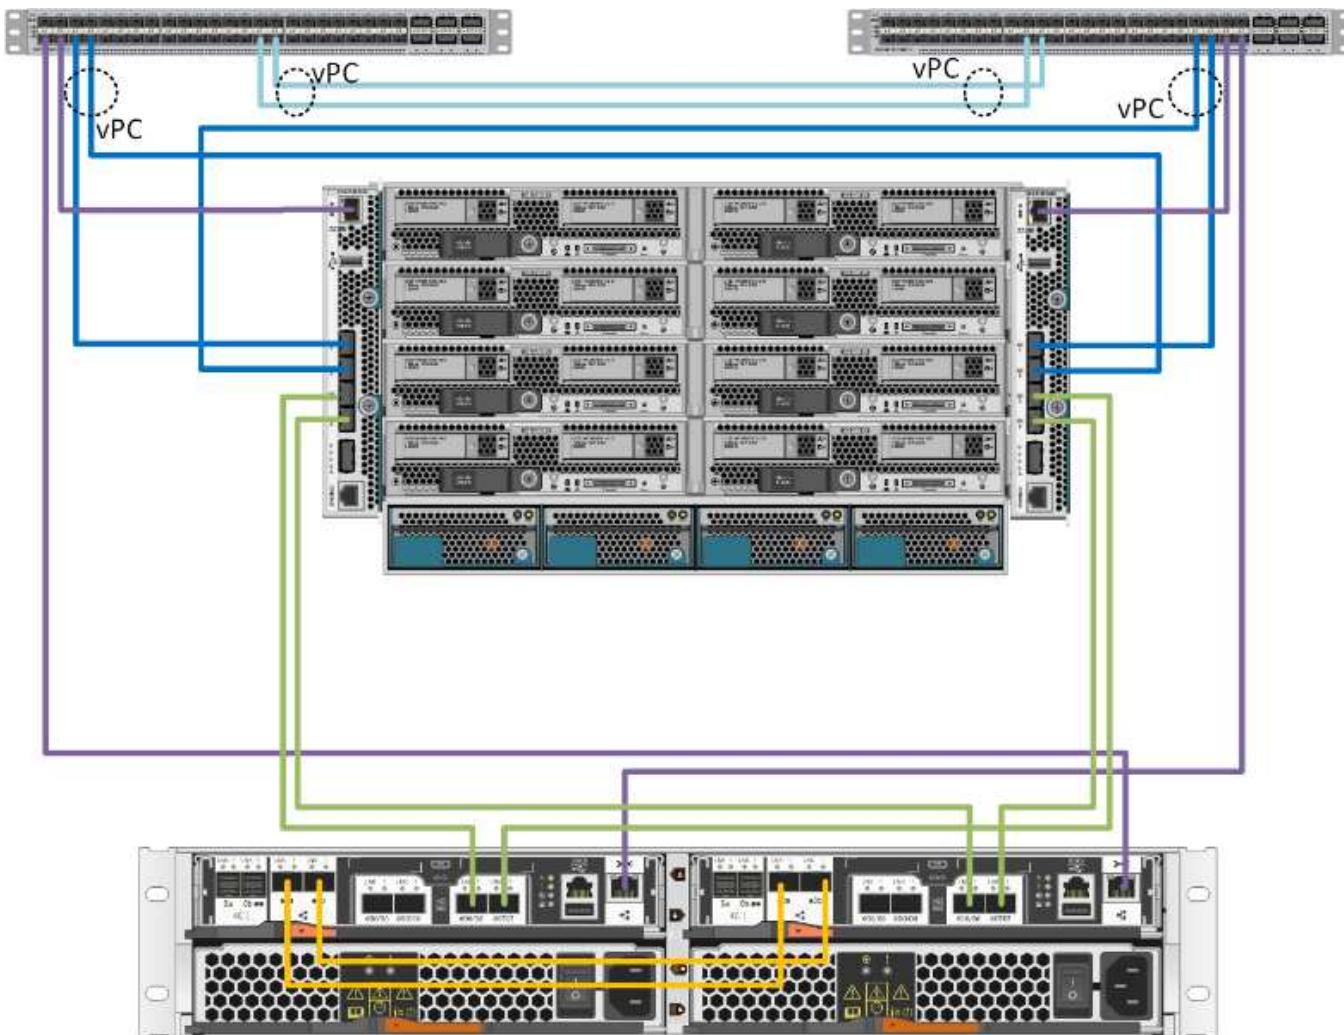
Software components

- NetApp ONTAP 9. 5
- VMWare vSphere 6.7U1
- Cisco UCS Manager 4.0(1b)
- Cisco NXOS Firmware 7.0(3)I6(1)

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. It features the new NetApp AFF A220 running ONTAP 9.5, dual Cisco Nexus 31108PCV switches, and Cisco UCS B200 M5 servers that run VMware vSphere 6.7U1. This validated solution uses Direct Connect IP storage over 10GbE technology.

The following figure illustrates FlexPod Express with VMware vSphere 6.7U1 IP-Based Direct Connect architecture.



Legend

- mgmt
- data
- IP (iSCSI/NFS)
- Cluster Interconnect
- vPC Peer

Use case summary

The FlexPod Express solution can be applied to several use cases, including the following:

- ROBOs
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is best suited for virtualized and mixed workloads.

Technology requirements

A FlexPod Express system requires a combination of hardware and software

components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

The following table lists the hardware components that are required for all FlexPod Express configurations.

Hardware	Quantity
AFF A220 HA Pair	1
Cisco UCS B200 M5 server	2
Cisco Nexus 31108PCV switch	2
Cisco UCS Virtual Interface Card (VIC) 1440 for the Cisco UCS B200 M5 server	2
Cisco UCS Mini with two Integrated UCS-FI-M-6324 fabric interconnects	1

Software requirements

The following table lists the software components that are required to implement the architectures of the FlexPod Express solutions.

Software	Version	Details
Cisco UCS Manager	4.0(1b)	For Cisco UCS Fabric Interconnect FI-6324UP
Cisco Blade software	4.0(1b)	For Cisco UCS B200 M5 servers
Cisco nenic driver	1.0.25.0	For Cisco VIC 1440 interface cards
Cisco NX-OS	7.0(3)I6(1)	For Cisco Nexus 31108PCV switches
NetApp ONTAP	9.5	For AFF A220 controllers

The following table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

Software	Version
VMware vCenter Server Appliance	6.7U1
VMware vSphere ESXi hypervisor	6.7U1

FlexPod Express Cabling Information

The reference validation cabling is documented in the following tables.

The following table lists cabling information for Cisco Nexus switch 31108PCV A.

Local device	Local port	Remote device	Remote port
Cisco Nexus switch 31108PCV A	Eth1/1	NetApp AFF A220 storage controller A	e0M
	Eth1/2	Cisco UCS-mini FI-A	mgmt0
	Eth1/3	Cisco UCS-mini FI-A	Eth1/1
	Eth 1/4	Cisco UCS-mini FI-B	Eth1/1
	Eth 1/13	Cisco NX 31108PCV B	Eth 1/13
	Eth 1/14	Cisco NX 31108PCV B	Eth 1/14

The following table lists the cabling information for Cisco Nexus switch 31108PCV B.

Local device	Local port	Remote device	Remote port
Cisco Nexus switch 31108PCV B	Eth1/1	NetApp AFF A220 storage controller B	e0M
	Eth1/2	Cisco UCS-mini FI-B	mgmt0
	Eth1/3	Cisco UCS-mini FI-A	Eth1/2
	Eth 1/4	Cisco UCS-mini FI-B	Eth1/2
	Eth 1/13	Cisco NX 31108PCV A	Eth 1/13
	Eth 1/14	Cisco NX 31108PCV A	Eth 1/14

The following table lists cabling information for NetApp AFF A220 storage controller A.

Local device	Local port	Remote device	Remote port
NetApp AFF A220 storage controller A	e0a	NetApp AFF A220 storage controller B	e0a
	e0b	NetApp AFF A220 storage controller B	e0b
	e0e	Cisco UCS-mini FI-A	Eth1/3
	e0f	Cisco UCS-mini FI-B	Eth1/3
	e0M	Cisco NX 31108PCV A	Eth1/1

The following table lists cabling information for NetApp AFF A220 storage controller B.

Local device	Local port	Remote device	Remote port
NetApp AFF A220 storage controller B	e0a	NetApp AFF A220 storage controller B	e0a
	e0b	NetApp AFF A220 storage controller B	e0b
	e0e	Cisco UCS-mini FI-A	Eth1/4
	e0f	Cisco UCS-mini FI-B	Eth1/4
	e0M	Cisco NX 31108PCV B	Eth1/1

The following table lists cabling information for Cisco UCS Fabric Interconnect A.

Local device	Local port	Remote device	Remote port
Cisco UCS Fabric Interconnect A	Eth1/1	Cisco NX 31108PCV A	Eth1/3
	Eth1/2	Cisco NX 31108PCV B	Eth1/3
	Eth1/3	NetApp AFF A220 storage controller A	e0e
	Eth1/4	NetApp AFF A220 storage controller B	e0e
	mgmt0	Cisco NX 31108PCV A	Eth1/2

The following table lists cabling information for Cisco UCS Fabric Interconnect B.

Local device	Local port	Remote device	Remote port
Cisco UCS Fabric Interconnect B	Eth1/1	Cisco NX 31108PCV A	Eth1/4
	Eth1/2	Cisco NX 31108PCV B	Eth1/4
	Eth1/3	NetApp AFF A220 storage controller A	e0f
	Eth1/4	NetApp AFF A220 storage controller B	e0f
	mgmt0	Cisco NX 31108PCV B	Eth1/2

Deployment Procedures

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches. Fabric Interconnect A and Fabric Interconnect B are the two Integrated Nexus Fabric Interconnects.

In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified

sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, <<text>> appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01>vlan create vif0 <<mgmt_vlan_id>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. The table below describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.



If you use separate in-band and out-of-band management VLANs, you must create a layer 3 route between them. For this validation, a common management VLAN was used.

VLAN name	VLAN purpose	ID used in validating this document
Management VLAN	VLAN for management interfaces	18
Native VLAN	VLAN to which untagged frames are assigned	2
NFS VLAN	VLAN for NFS traffic	104
VMware vMotion VLAN	VLAN designated for the movement of virtual machines (VMs) from one physical host to another	103
VM traffic VLAN	VLAN for VM application traffic	102
iSCSI-A-VLAN	VLAN for iSCSI traffic on fabric A	124
iSCSI-B-VLAN	VLAN for iSCSI traffic on fabric B	125

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <<var_xxxx_vlan>>, where xxxx is the purpose of the VLAN (such as iSCSI-A).

The following table lists the VMware VMs created.

VM Description	Host Name
VMware vCenter Server	Seahawks-vcsa.cie.netapp.com

Cisco Nexus 31108PCV deployment procedure

This section details the Cisco Nexus 31308PCV switch configuration used in a FlexPod Express environment.

Initial setup of Cisco Nexus 31108PCV switch

This procedures describes how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.



This procedure assumes that you are using a Cisco Nexus 31108PCV running NX-OS software release 7.0(3)I6(1).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 31108PCV switches can be connected to an existing management network, or the mgmt0 interfaces of the 31108PCV switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.

In this deployment guide, the FlexPod Express Cisco Nexus 31108PCV switches are connected to an existing management network.

3. To configure the Cisco Nexus 31108PCV switches, power on the switch and follow the on-screen prompts, as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

*Note: setup is mainly used for configuring the system initially, when no configuration is present. So setup always assumes system defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

```

Would you like to enter the basic configuration dialog (yes/no): y
Do you want to enforce secure password standard (yes/no) [y]: y
Create another login account (yes/no) [n]: n
Configure read-only SNMP community string (yes/no) [n]: n
Configure read-write SNMP community string (yes/no) [n]: n
Enter the switch name : 31108PCV-A
Continue with Out-of-band (mgmt0) management configuration? (yes/no)
[y]: y
Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>
Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>
Configure the default gateway? (yes/no) [y]: y
IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>
Configure advanced IP options? (yes/no) [n]: n
Enable the telnet service? (yes/no) [n]: n
Enable the ssh service? (yes/no) [y]: y
Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
Number of rsa key bits <1024-2048> [1024]: <enter>
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address : <<var_ntp_ip>>
Configure default interface layer (L3/L2) [L2]: <enter>
Configure default switchport interface state (shut/noshut) [noshut]:
<enter>
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]:
<enter>

```

4. A summary of your configuration is displayed and you are asked if you would like to edit the configuration. If your configuration is correct, enter n.

Would you like to edit the configuration? (yes/no) [n]: no

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

Use this configuration and save it? (yes/no) [y]: Enter

6. Repeat steps 1 through 5 for Cisco Nexus switch B.

Enable advanced features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options.

1. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode by using the command (config t) and run the following commands:

```
feature interface-vlan
feature lacp
feature vpc
```

 The default port channel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the port channel. You can achieve better distribution across the members of the port channel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

2. From configuration mode (config t), run the following commands to set the global port channel load-balancing configuration on Cisco Nexus switch A and switch B:

```
port-channel load-balance src-dst ip-14port
```

Perform global spanning-tree configuration

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (config t), run the following commands to configure the default spanning-tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
```

Define VLANs

Before individual ports with different VLANs are configured, the layer-2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (config t), run the following commands to define and describe the layer 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure access and management port descriptions

As is the case with assigning names to the layer-2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (config t) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus switch A

```
int eth1/1
  description AFF A220-A e0M
int eth1/2
  description Cisco UCS FI-A mgmt0
int eth1/3
  description Cisco UCS FI-A eth1/1
int eth1/4
  description Cisco UCS FI-B eth1/1
int eth1/13
  description vPC peer-link 31108PVC-B 1/13
int eth1/14
  description vPC peer-link 31108PVC-B 1/14
```

Cisco Nexus switch B

```
int eth1/1
  description AFF A220-B e0M
int eth1/2
  description Cisco UCS FI-B mgmt0
int eth1/3
  description Cisco UCS FI-A eth1/2
int eth1/4
  description Cisco UCS FI-B eth1/2
int eth1/13
  description vPC peer-link 31108PVC-B 1/13
int eth1/14
  description vPC peer-link 31108PVC-B 1/14
```

Configure server and storage management interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (config t), run the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus switch A

```
int eth1/1-2
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Cisco Nexus switch B

```
int eth1/1-2
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Add NTP distribution interface

Cisco Nexus switch A

From the global configuration mode, execute the following commands.

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exitntp peer <switch-b-ntp-ip> use-vrf default
```

Cisco Nexus switch B

From the global configuration mode, execute the following commands.

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch- b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exitntp peer <switch-a-ntp-ip> use-vrf default
```

Perform virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer-2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back mgmt0 configuration, use the addresses defined on the interfaces and verify that they can communicate by using the ping <<switch_A/B_mgmt0_ip_addr>>vrf management command.

From configuration mode (config t), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus switch A

```

vpc domain 1
  role priority 10
  peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source
  <<switch_A_mgmt0_ip_addr>> vrf management
    peer-gateway
    auto-recovery
    ip arp synchronize
    int eth1/13-14
      channel-group 10 mode active
  int Po10description vPC peer-link
  switchport
    switchport mode trunkswitchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<nfs_vlan_id>>, <<vmotion_vlan_id>>,
    <<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
    <<iSCSI_B_vlan_id>> spanning-tree port type network
  vpc peer-link
  no shut
  exit
  int Po13
  description vPC ucs-FI-A
  switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
    <<mgmt_vlan>> spanning-tree port type network
  mtu 9216
  vpc 13
  no shut
  exit
  int eth1/3
    channel-group 13 mode active
  int Po14
  description vPC ucs-FI-B
  switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
    <<mgmt_vlan>> spanning-tree port type network
  mtu 9216
  vpc 14
  no shut
  exit
  int eth1/4
    channel-group 14 mode active
copy run start

```

Cisco Nexus switch B

```
vpc domain 1
peer-switch
role priority 20
peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source
<<switch_B_mgmt0_ip_addr>> vrf management
  peer-gateway
  auto-recovery
  ip arp synchronize
  int eth1/13-14
    channel-group 10 mode active
int Po10
description vPC peer-link
switchport
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<nfs_vlan_id>>, <<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>> spanning-tree port type network
vpc peer-link
no shut
exit
int Po13
description vPC ucs-FI-A
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>> spanning-tree port type network
mtu 9216
vpc 13
no shut
exit
int eth1/3
  channel-group 13 mode active
int Po14
description vPC ucs-FI-B
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>> spanning-tree port type network
mtu 9216
vpc 14
no shut
exit
int eth1/4
```

```
channel-group 14 mode active
copy run start
```

 In this solution validation, a maximum transmission unit (MTU) of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components result in packets being dropped.

Uplink into existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 31108PVC switches included in the FlexPod environment into the infrastructure. The uplinks can be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp Storage Controller AFF2xx Series Installation

NetApp Hardware Universe

The [NetApp Hardware Universe](#) (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

1. Access the [HWU](#) application to view the system configuration guides. Select the Compare Storage Systems tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFF2XX Series prerequisites

To plan the physical location of the storage systems, see the the following sections:

Electrical requirements

Supported power cords

Onboard ports and cables

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF A220 Documentation](#).

NetApp ONTAP 9.5

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.5 Software Setup Guide](#) (available in the [ONTAP 9 Documentation Center](#)). The table below illustrates ONTAP 9.5 installation and configuration information.



This system is set up in a two-node switchless cluster configuration.

Cluster Detail	Cluster Detail Value
Cluster node A IP address	<<var_nodeA_mgmt_ip>>
Cluster node A netmask	<<var_nodeA_mgmt_mask>>
Cluster node A gateway	<<var_nodeA_mgmt_gateway>>
Cluster node A name	<<var_nodeA>>
Cluster node B IP address	<<var_nodeB_mgmt_ip>>
Cluster node B netmask	<<var_nodeB_mgmt_mask>>
Cluster node B gateway	<<var_nodeB_mgmt_gateway>>
Cluster node B name	<<var_nodeB>>
ONTAP 9.5 URL	<<var_url_boot_software>>
Name for cluster	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster B gateway	<<var_clustermgmt_gateway>>
Cluster B netmask	<<var_clustermgmt_mask>>
Domain name	<<var_domain_name>>
DNS server IP (you can enter more than one)	<<var_dns_server_ip>>
NTP server A IP	<< switch-a-ntp-ip >>
NTP server B IP	<< switch-b-ntp-ip >>

Configure node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl- C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot.

```
autoboot
```

3. Press Ctrl- C to enter the Boot menu.

If ONTAP 9.5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.5 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl- C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

17. While node A is initializing, begin configuring node B.

Configure node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press Ctrl-C to enter the Boot menu.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name

11. Enter y to set the newly installed software as the default to be used for subsequent reboots.

12. Enter y to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.5 boots on the node for the first time.

The node and cluster setup procedure has changed slightly in ONTAP 9.5. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up node A.

```
Welcome to the cluster setup wizard.  
You can enter the following commands at any time:  
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.  
You can return to cluster setup at any time by typing "cluster setup".  
To accept a default or omit a question, do not enter a value.  
This system will send event messages and periodic reports to NetApp  
Technical Support. To disable this feature, enter  
autosupport modify -support disable  
within 24 hours.  
Enabling AutoSupport can significantly speed problem determination and  
resolution should a problem occur on your system.  
For further information on AutoSupport, see:  
http://support.netapp.com/autosupport/  
Type yes to confirm and continue {yes}: yes  
Enter the node management interface port [e0M]:  
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>  
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>  
Enter the node management interface default gateway:  
<<var_nodeA_mgmt_gateway>>  
A node management interface on port e0M with IP address  
<<var_nodeA_mgmt_ip>> has been created.  
Use your web browser to complete cluster setup by accessing  
https://<<var\_nodeA\_mgmt\_ip>>  
Otherwise, press Enter to complete cluster setup using the command line  
interface:
```

2. Navigate to the IP address of the node's management interface.



Cluster setup can also be performed by using the CLI. This document describes cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup to configure the cluster.

4. Enter <<var_clusternname>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system.

Select Switchless Cluster for the cluster type. Enter the cluster base license.

5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter `<<var_clustermgmt_ip>>` in the Cluster Management IP Address field, `<<var_clustermgmt_mask>>` in the Netmask field, and `<<var_clustermgmt_gateway>>` in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter `<<var_nodeA_mgmt_ip>>` for node B.
 - d. Enter `<<var_domain_name>>` in the DNS Domain Name field. Enter `<<var_dns_server_ip>>` in the DNS Server IP Address field.

You can enter multiple DNS server IP addresses.

- e. Enter `<<switch-a-ntp-ip>>` in the Primary NTP Server field.

You can also enter an alternate NTP server as `<<switch-b-ntp-ip>>`.

8. Configure the support information.
 - a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
 - b. Enter the SMTP mail host and email address for event notifications.

You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

9. When indicated that the cluster configuration has completed, click Manage Your Cluster to configure the storage.

Continuation of storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set on-board UTA2 ports personality

1. Verify the current mode and the current type of the ports by running the `ucadmin show` command.

```
AFFA220-Clus::> ucadmin show
      Current  Current  Pending  Pending  Admin
Node     Adapter  Mode    Type    Mode    Type    Status
-----  -----  -----  -----  -----  -----  -----
-----  -----
AFFA220-Clus-01
    0c      cna    target  -      -      offline
AFFA220-Clus-01
    0d      cna    target  -      -      offline
AFFA220-Clus-01
    0e      cna    target  -      -      offline
AFFA220-Clus-01
    0f      cna    target  -      -      offline
AFFA220-Clus-02
    0c      cna    target  -      -      offline
AFFA220-Clus-02
    0d      cna    target  -      -      offline
AFFA220-Clus-02
    0e      cna    target  -      -      offline
AFFA220-Clus-02
    0f      cna    target  -      -      offline
8 entries were displayed.
```

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If not, change the port personality by running the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode
cna -type target
```

The ports must be offline to run the previous command. To take a port offline, run the following command:

```
network fcp adapter modify -node <home node of the port> -adapter <port
name> -state down
```



If you changed the port personality, you must reboot each node for the change to take effect.

Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command:

```
node run -node * options cdpd.enable on
```

Enable Link-layer Discovery Protocol on all Ethernet ports

Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches by running the following command. This command enables LLDP on all ports of all nodes in the cluster.

```
node run * options lldp.enable on
```

Rename management logical interfaces

To rename the management logical interfaces (LIFs), complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clusternode>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clusternode>> -lif
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clusternode>> -lif
cluster_setup_node_mgmt_lif_AFF_A220_A_1 - newname AFF_A220-01_mgmt1
```

Set auto-revert on cluster management

Set the auto-revert parameter on the cluster management interface.

```
network interface modify -vserver <<clusternode>> -lif cluster_mgmt -auto-
revert true
```

Set up service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_nodeA>> -address
-family IPv4 -enable true - dhcp none -ip-address <<var_nodeA_sp_ip>>
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>
system service-processor network modify -node <<var_nodeB>> -address
-family IPv4 -enable true - dhcp none -ip-address <<var_nodeB_sp_ip>>
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```

3. Verify the HA status of the two-node cluster.



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.

Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```

The message `Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node <<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node <<var_nodeB>>
```

Create jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove data ports from default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following commands:

```
net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
```



The Cisco UCS Mini direct connection to ONTAP does not support LACP.

Configure jumbo frames in NetApp ONTAP

To configure an ONTAP network port to use jumbo frames (that usually have an MTU of 9,000 bytes), run the following commands from the cluster shell:

```

AFF A220::> network port modify -node node_A -port e0e -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port e0e -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_A -port e0f -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port e0f -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name e0e-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeA>> -vlan-name e0f-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0e-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0f-
<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports
<<var_nodeA>>: e0e- <<var_nfs_vlan_id>>, <<var_nodeB>>: e0e-
<<var_nfs_vlan_id>> , <<var_nodeA>>:e0f- <<var_nfs_vlan_id>>,
<<var_nodeB>>:e0f-<<var_nfs_vlan_id>>

```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name e0e-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name e0f-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0e-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0f-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>: e0e- <<var_iscsi_vlan_A_id>>, <<var_nodeB>>: e0e-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>: e0f- <<var_iscsi_vlan_B_id>>, <<var_nodeB>>: e0f-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name e0m-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0m-
<<mgmt_vlan_id>>

```

Create aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Start with five disks; you can add disks to an aggregate when additional storage is required.

The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1_nodeA` is online.

Configure time zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is America/New_York. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the sysLocation and sysContact variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the snmp community delete all command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the security snmpusers command to view the engine ID.
2. Create a user called snmpv3user.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select md5 as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select des as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>> -transport https -support enable -noteto <<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the vserver create command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_nodeA -rootvolume- security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fcp
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the SVM vstorage parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
vserver nfs modify -vserver Infra-SVM -vstorage enabled  
vserver nfs show
```



Commands are prefaced by `vserver` in the command line because SVMs were previously called servers

Configure NFSv3 in ONTAP

The table below lists the information needed to complete this configuration.

Detail	Detail Value
ESXi host A NFS IP address	<code><<var_esxi_hostA_nfs_ip>></code>
ESXi host B NFS IP address	<code><<var_esxi_hostB_nfs_ip>></code>

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid falsevserver export-  
policy rule create -vserver Infra-SVM -policyname default -ruleindex 2  
-protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>> -rorule sys -rwrule  
sys -superuser sys -allow-suid false  
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS B-Series servers are added.

Create iSCSI service in ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM  
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

To create a load-sharing mirror of the SVM root volume in ONTAP, complete the following steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate  
aggr1_nodeA -size 1GB -type DP  
volume create -vserver Infra_Vserver  
-volume rootvol_m02 -aggregate aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path  
Infra-SVM:rootvol_m01 -type LS -schedule 15min  
snapmirror create -source-path Infra-SVM:rootvol -destination-path  
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol snapmirror  
show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag  
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS fully qualified domain name (FQDN) of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.

Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...
```

Example: `security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial 552429A6`

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...
```

Example: `security certificate create -common-name infra-svm.netapp.com -type server -size 2048 -country US -state "North Carolina" -locality "RTP" -organization "NetApp" -unit "FlexPod" -email-addr "abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver Infra-SVM`

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...
```

Example: `security ssl modify -vserver Infra-SVM -server-enabled true -client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common-name infra-svm.netapp.com`

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
System services firewall policy delete -policy mgmt -service http
-vserver <<var_clusternamespace>>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow SVM to be available by the web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled
true
```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol® volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate
aggr1_nodeA -size 500GB - state online -policy default -junction-path
/infra_datastore_1 -space-guarantee none -percent- snapshot-space 0
volume create -vserver Infra-SVM -volume infra_datastore_2 -aggregate
aggr1_nodeB -size 500GB - state online -policy default -junction-path
/infra_datastore_2 -space-guarantee none -percent- snapshot-space 0
```

```
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap -space
-guarantee none -percent-snapshot-space 0 -snapshot-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0
```

Enable deduplication in ONTAP

To enable deduplication on appropriate volumes once a day, run the following commands:

```

volume efficiency modify -vserver Infra-SVM -volume esxi_boot -schedule
sun-sat@0
volume efficiency modify -vserver Infra-SVM -volume infra_datastore_1
-schedule sun-sat@0
volume efficiency modify -vserver Infra-SVM -volume infra_datastore_2
-schedule sun-sat@0

```

Create LUNs in ONTAP

To create two boot logical unit numbers (LUNs), run the following commands:

```

lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware - space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware - space-reserve disabled

```



When adding an extra Cisco UCS C-Series server, an extra boot LUN must be created.

Create iSCSI LIFs in ONTAP

The table below lists the information needed to complete this configuration.

Detail	Detail Value
Storage node A iSCSI LIF01A	<<var_nodeA_iscsi_lif01a_ip>>
Storage node A iSCSI LIF01A network mask	<<var_nodeA_iscsi_lif01a_mask>>
Storage node A iSCSI LIF01B	<<var_nodeA_iscsi_lif01b_ip>>
Storage node A iSCSI LIF01B network mask	<<var_nodeA_iscsi_lif01b_mask>>
Storage node B iSCSI LIF01A	<<var_nodeB_iscsi_lif01a_ip>>
Storage node B iSCSI LIF01A network mask	<<var_nodeB_iscsi_lif01a_mask>>
Storage node B iSCSI LIF01B	<<var_nodeB_iscsi_lif01b_ip>>
Storage node B iSCSI LIF01B network mask	<<var_nodeB_iscsi_lif01b_mask>>

1. Create four iSCSI LIFs, two on each node.

```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi - home-node <<var_nodeA>> -home-port e0e-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi - home-node <<var_nodeA>> -home-port e0f-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi - home-node <<var_nodeB>> -home-port e0e-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi - home-node <<var_nodeB>> -home-port e0f-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface show

```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
Storage node A NFS LIF 01 a IP	<<var_nodeA_nfs_lif_01_a_ip>>
Storage node A NFS LIF 01 a network mask	<<var_nodeA_nfs_lif_01_a_mask>>
Storage node A NFS LIF 01 b IP	<<var_nodeA_nfs_lif_01_b_ip>>
Storage node A NFS LIF 01 b network mask	<<var_nodeA_nfs_lif_01_b_mask>>
Storage node B NFS LIF 02 a IP	<<var_nodeB_nfs_lif_02_a_ip>>
Storage node B NFS LIF 02 a network mask	<<var_nodeB_nfs_lif_02_a_mask>>
Storage node B NFS LIF 02 b IP	<<var_nodeB_nfs_lif_02_b_ip>>
Storage node B NFS LIF 02 b network mask	<<var_nodeB_nfs_lif_02_b_mask>>

1. Create an NFS LIF.

```

network interface create -vserver Infra-SVM -lif nfs_lif01_a -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port e0e-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_a_ip>> -netmask <<
var_nodeA_nfs_lif_01_a_mask>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif01_b -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port e0f-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_b_ip>> -netmask <<
var_nodeA_nfs_lif_01_b_mask>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02_a -role data
-data-protocol nfs -home-node <<var_nodeB>> -home-port e0e-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_a_ip>> -netmask <<
var_nodeB_nfs_lif_02_a_mask>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02_b -role data
-data-protocol nfs -home-node <<var_nodeB>> -home-port e0f-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_b_ip>> -netmask <<
var_nodeB_nfs_lif_02_b_mask>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true
network interface show

```

Add infrastructure SVM administrator

The following table lists the information needed to complete this configuration.

Detail	Detail value
Vsmgmt IP	<<var_svm_mgmt_ip>>
Vsmgmt network mask	<<var_svm_mgmt_mask>>
Vsmgmt default gateway	<<var_svm_mgmt_gateway>>

To add the infrastructure SVM administrator and SVM administration LIF to the management network, complete the following steps:

1. Run the following command:

```

network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true

```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway
<<var_svm_mgmt_gateway>> network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
security login unlock -username vsadmin -vserver
```

Cisco UCS server configuration

FlexPod Cisco UCS base

Perform Initial Setup of Cisco UCS 6324 Fabric Interconnect for FlexPod Environments.

This section provides detailed procedures to configure Cisco UCS for use in a FlexPod ROBO environment by using Cisco UCS Manager.

Cisco UCS fabric interconnect 6324 A

Cisco UCS uses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

Cisco UCS Manager 4.0(1b) supports the 6324 Fabric Interconnect that integrates the Fabric Interconnect into the Cisco UCS Chassis and provides an integrated solution for a smaller deployment environment. Cisco UCS Mini simplifies the system management and saves cost for the low scale deployments.

The hardware and software components support Cisco's unified fabric, which runs multiple types of data center traffic over a single converged network adapter.

Initial system setup

The first time when you access a fabric interconnect in a Cisco UCS domain, a setup wizard prompts you for the following information required to configure the system:

- Installation method (GUI or CLI)
- Setup mode (restore from full system backup or initial setup)
- System configuration type (standalone or cluster configuration)
- System name
- Admin password
- Management port IPv4 address and subnet mask, or IPv6 address and prefix

- Default gateway IPv4 or IPv6 address
- DNS Server IPv4 or IPv6 address
- Default domain name

The following table lists the information needed to complete the Cisco UCS initial configuration on Fabric Interconnect A

Detail	Detail/value
System Name	<<var_ucs_clustername>>
Admin Password	<<var_password>>
Management IP Address: Fabric Interconnect A	<<var_ucs_mgmt_ip>>
Management netmask: Fabric Interconnect A	<<var_ucs_mgmt_mask>>
Default gateway: Fabric Interconnect A	<<var_ucs_mgmt_gateway>>
Cluster IP address	<<var_ucs_cluster_ip>>
DNS server IP address	<<var_nameserver_ip>>
Domain name	<<var_domain_name>>

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6324 Fabric Interconnect A.

```
Enter the configuration method. (console/gui) ? console

Enter the setup mode; setup newly or restore from backup.
(setup/restore) ? setup

You have chosen to setup a new Fabric interconnect. Continue? (y/n): y

Enforce strong password? (y/n) [y]: Enter

Enter the password for "admin":<<var_password>>
Confirm the password for "admin":<<var_password>>

Is this Fabric interconnect part of a cluster(select 'no' for
standalone)? (yes/no) [n]: yes

Enter the switch fabric (A/B) []: A

Enter the system name: <<var_ucs_clustername>>

Physical Switch Mgmt0 IP address : <<var_ucs_mgmt_ip>>

Physical Switch Mgmt0 IPv4 netmask : <<var_ucs_mgmt_mask>>

IPv4 address of the default gateway : <<var_ucs_mgmt_gateway>>

Cluster IPv4 address : <<var_ucs_cluster_ip>>

Configure the DNS Server IP address? (yes/no) [n]: y

DNS IP address : <<var_nameserver_ip>>

Configure the default domain name? (yes/no) [n]: y
Default domain name: <<var_domain_name>>

Join centralized management environment (UCS Central)? (yes/no) [n]:
no

NOTE: Cluster IP will be configured only after both Fabric
Interconnects are initialized. UCSM will be functional only after peer
FI is configured in clustering mode.

Apply and save the configuration (select 'no' if you want to re-
enter)? (yes/no): yes
Applying configuration. Please wait.

Configuration file - Ok
```

2. Review the settings displayed on the console. If they are correct, answer `yes` to apply and save the configuration.
3. Wait for the login prompt to verify that the configuration has been saved.

The following table lists the information needed to complete the Cisco UCS initial configuration on Fabric Interconnect B.

Detail	Detail/value
System Name	<code><<var_ucs_clustername>></code>
Admin Password	<code><<var_password>></code>
Management IP Address-FI B	<code><<var_ucs_b_mgmt_ip>></code>
Management Netmask-FI B	<code><<var_ucs_b_mgmt_mask>></code>
Default Gateway-FI B	<code><<var_ucs_b_mgmt_gateway>></code>
Cluster IP Address	<code><<var_ucs_cluster_ip>></code>
DNS Server IP address	<code><<var_nameserver_ip>></code>
Domain Name	<code><<var_domain_name>></code>

1. Connect to the console port on the second Cisco UCS 6324 Fabric Interconnect B.

```

Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect.
This Fabric interconnect will be added to the cluster. Continue (y/n) ?

Y

Enter the admin password of the peer Fabric
interconnect:<<var_password>>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <<var_ucs_mgmt_ip>>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <<var_ucs_mgmt_mask>>
Cluster IPv4 address: <<var_ucs_cluster_address>>

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric
Interconnect Mgmt0 IPv4 Address

Physical Switch Mgmt0 IP address : <<var_ucs_mgmt_ip>>

Apply and save the configuration (select 'no' if you want to re-
enter)? (yes/no): yes
Applying configuration. Please wait.

Configuration file - Ok

```

2. Wait for the login prompt to confirm that the configuration has been saved.

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address.
You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to come up.
2. Click the Launch UCS Manager link to launch Cisco UCS Manager.
3. Accept the necessary security certificates.
4. When prompted, enter admin as the user name and enter the administrator password.
5. Click Login to log in to Cisco UCS Manager.

Cisco UCS Manager software version 4.0(1b)

This document assumes the use of Cisco UCS Manager Software version 4.0(1b). To upgrade the Cisco UCS Manager software and the Cisco UCS 6324 Fabric Interconnect software refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Configure Cisco UCS Call Home

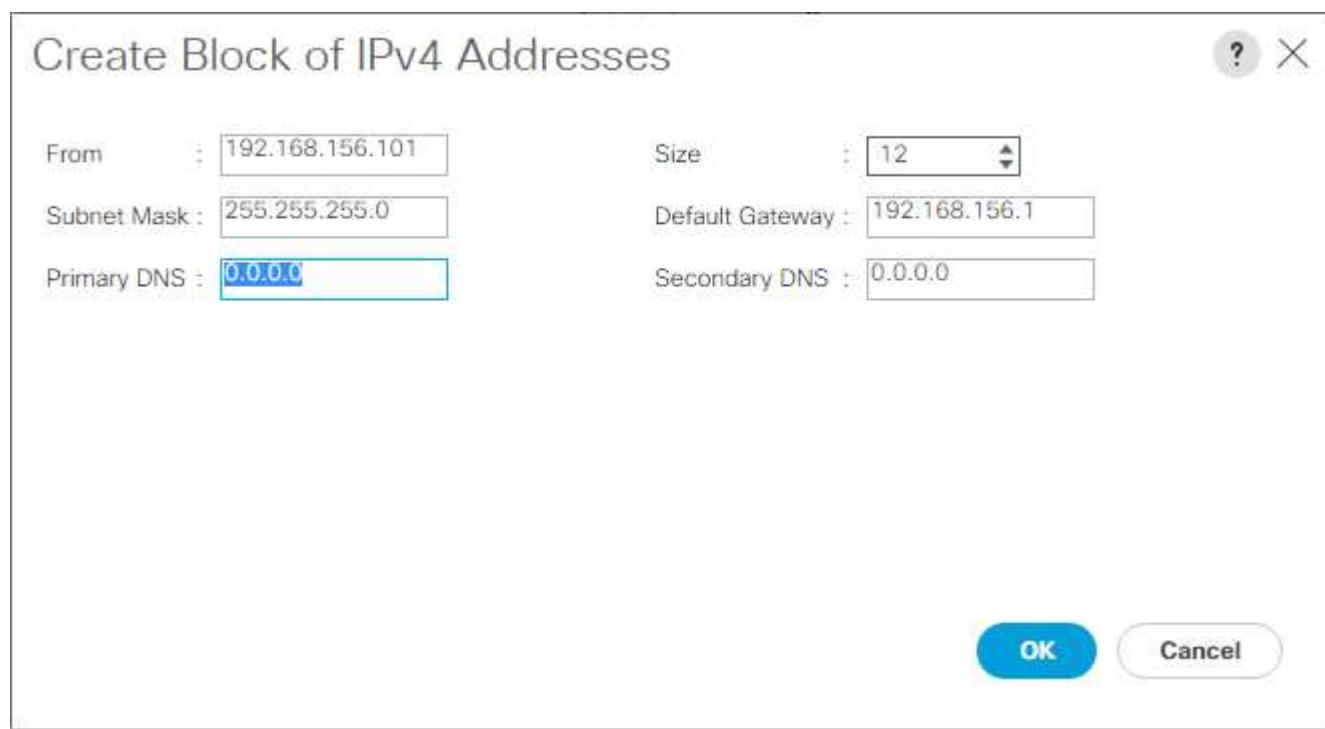
Cisco highly recommends that you configure Call Home in Cisco UCS Manager. Configuring Call Home accelerates the resolution of support cases. To configure Call Home, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

Add block of IP addresses for keyboard, video, mouse access

To create a block of IP addresses for in band server keyboard, video, mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Expand Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information.

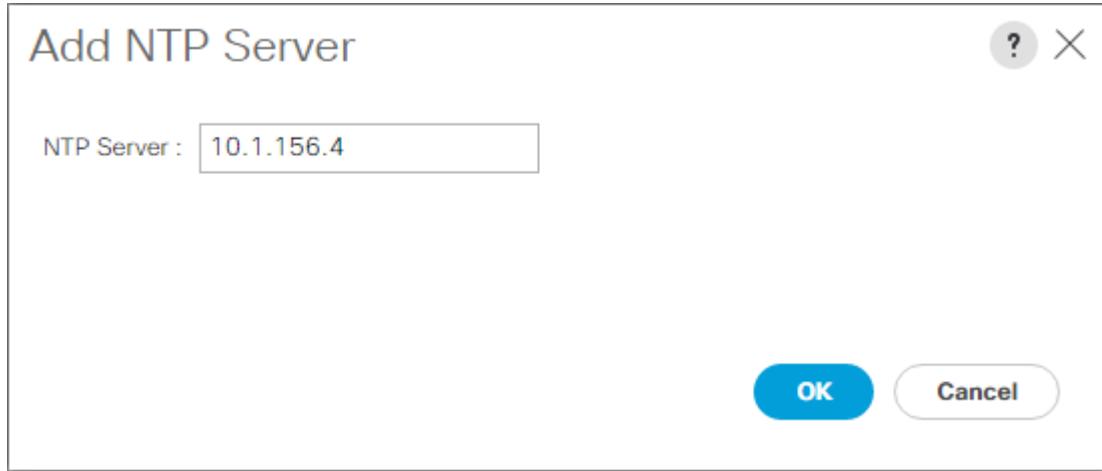


5. Click OK to create the block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Expand All > Time Zone Management.
3. Select Time Zone.
4. In the Properties pane, select the appropriate time zone in the Time Zone menu.
5. Click Save Changes and click OK.
6. Click Add NTP Server.
7. Enter <switch-a-ntp-ip> or <Nexus-A-mgmt-IP> and click OK. Click OK.



8. Click Add NTP Server.
9. Enter <switch-b-ntp-ip> or <Nexus-B-mgmt-IP> and click OK. Click OK on the confirmation.

Edit chassis discovery policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and of additional fabric extenders for further Cisco UCS C-Series connectivity. To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left and select Equipment in the second list.
2. In the right pane, select the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.
5. Click Save Changes.
6. Click OK.

Enable server, uplink, and storage ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, select the Equipment tab.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports 1 and 2 that are connected to the Cisco Nexus 31108 switches, right-click, and select Configure as Uplink Port.
5. Click Yes to confirm the uplink ports and click OK.
6. Select ports 3 and 4 that are connected to the NetApp Storage Controllers, right-click, and select Configure as Appliance Port.
7. Click Yes to confirm the appliance ports.
8. On the Configure as Appliance Port window, click OK.
9. Click OK to confirm.
10. In the left pane, select Fixed Module under Fabric Interconnect A.
11. From the Ethernet Ports tab, confirm that ports have been configured correctly in the If Role column. If any port C-Series servers were configured on the Scalability port, click on it to verify port connectivity there.

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module

General	Ethernet Ports	FC Ports	Faults	Events					
Slot	Aggr Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer	
1	0	1	00:DE:FB:30:36:88	Network	Physical	▲ Up	▲ Enabled		
1	0	2	00:DE:FB:30:36:89	Network	Physical	▲ Up	▲ Enabled		
1	0	3	00:DE:FB:30:36:8A	Appliance Storage	Physical	▲ Up	▲ Enabled		
1	0	4	00:DE:FB:30:36:8B	Appliance Storage	Physical	▲ Up	▲ Enabled		
1	5	1	00:DE:FB:30:36:8C	Unconfigured	Physical	▼ Sfp Not Present	▼ Disabled		
1	5	2	00:DE:FB:30:36:8D	Unconfigured	Physical	▼ Sfp Not Present	▼ Disabled		
1	5	3	00:DE:FB:30:36:8E	Unconfigured	Physical	▼ Sfp Not Present	▼ Disabled		
1	5	4	00:DE:FB:30:36:8F	Unconfigured	Physical	▼ Sfp Not Present	▼ Disabled		

12. Expand Equipment > Fabric Interconnects > Fabric Interconnect B > Fixed Module.
13. Expand Ethernet Ports.
14. Select Ethernet ports 1 and 2 that are connected to the Cisco Nexus 31108 switches, right-click, and select Configure as Uplink Port.
15. Click Yes to confirm the uplink ports and click OK.
16. Select ports 3 and 4 that are connected to the NetApp Storage Controllers, right-click, and select Configure as Appliance Port.
17. Click Yes to confirm the appliance ports.
18. On the Configure as Appliance Port window, click OK.
19. Click OK to confirm.
20. In the left pane, select Fixed Module under Fabric Interconnect B.
21. From the Ethernet Ports tab, confirm that ports have been configured correctly in the If Role column. If any port C-Series servers were configured on the Scalability port, click it to verify port connectivity there.

Equipment / Fabric Interconnects / Fabric Interconnect B (primary) / Fixed Module / Ethernet Ports

Ethernet Ports

Advanced Filter		Export	Print	<input checked="" type="checkbox"/> All	<input checked="" type="checkbox"/> Unconfigured	<input checked="" type="checkbox"/> Network	<input checked="" type="checkbox"/> Server	<input checked="" type="checkbox"/> FCoF Drv	<input checked="" type="checkbox"/> Unified Uplink	<input checked="" type="checkbox"/> Appliance Storage	<input checked="" type="checkbox"/> FCoF Storage	<input checked="" type="checkbox"/> Unified Storage	<input checked="" type="checkbox"/> Monitor
Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer					
1	0	1	00:DE:FB:30:34:C8	Network	Physical	Up	Enabled						
1	0	2	00:DE:FB:30:3A:C9	Network	Physical	Up	Enabled						
1	0	3	00:DE:FB:30:3A:CA	Appliance Storage	Physical	Up	Enabled						
1	0	4	00:DE:FB:30:3A:C8	Appliance Storage	Physical	Up	Enabled						
1	5	1	00:DE:FB:30:3A:CC	Unconfigured	Physical	Sfp Not Present	Disabled						
1	5	2	00:DE:FB:30:3A:CD	Unconfigured	Physical	Sfp Not Present	Disabled						
1	5	3	00:DE:FB:30:3A:CE	Unconfigured	Physical	Sfp Not Present	Disabled						
1	5	4	00:DE:FB:30:3A:CF	Unconfigured	Physical	Sfp Not Present	Disabled						

Create uplink port channels to Cisco Nexus 31108 switches

To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, select the LAN tab in the navigation pane.

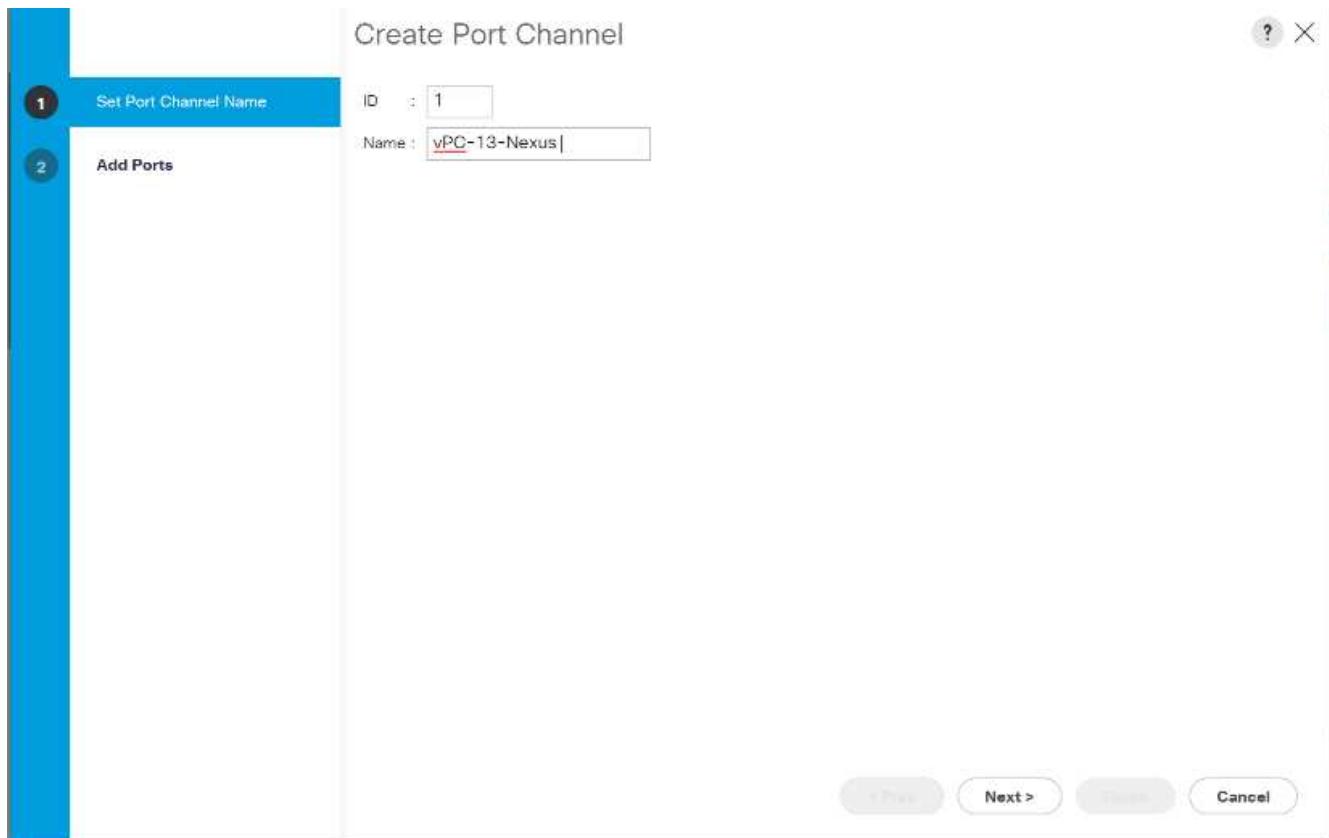


In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 31108 switches and one from Fabric B to both Cisco Nexus 31108 switches. If you are using standard switches, modify this procedure accordingly. If you are using 1 Gigabit Ethernet (1GbE) switches and GLC-T SFPs on the Fabric Interconnects, the interface speeds of Ethernet ports 1/1 and 1/2 in the Fabric Interconnects must be set to 1Gbps.

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 13 as the unique ID of the port channel.

6. Enter vPC-13-Nexus as the name of the port channel.

7. Click Next.



8. Select the following ports to be added to the port channel:

- Slot ID 1 and port 1
- Slot ID 1 and port 2

9. Click >> to add the ports to the port channel.

10. Click Finish to create the port channel. Click OK.

11. Under Port Channels, select the newly created port channel.

The port channel should have an Overall Status of Up.

12. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.

13. Right-click Port Channels.

14. Select Create Port Channel.

15. Enter 14 as the unique ID of the port channel.

16. Enter vPC-14-Nexus as the name of the port channel. Click Next.

17. Select the following ports to be added to the port channel:

- Slot ID 1 and port 1
- Slot ID 1 and port 2

18. Click >> to add the ports to the port channel.

19. Click Finish to create the port channel. Click OK.

20. Under Port Channels, select the newly created port-channel.

21. The port channel should have an Overall Status of Up.

Create an organization (optional)

Organizations are used to organizing resources and restricting access to various groups within the IT organization, thereby enabling multitenancy of the compute resources.



Although this document does not assume the use of organizations, this procedure provides instructions for creating one.

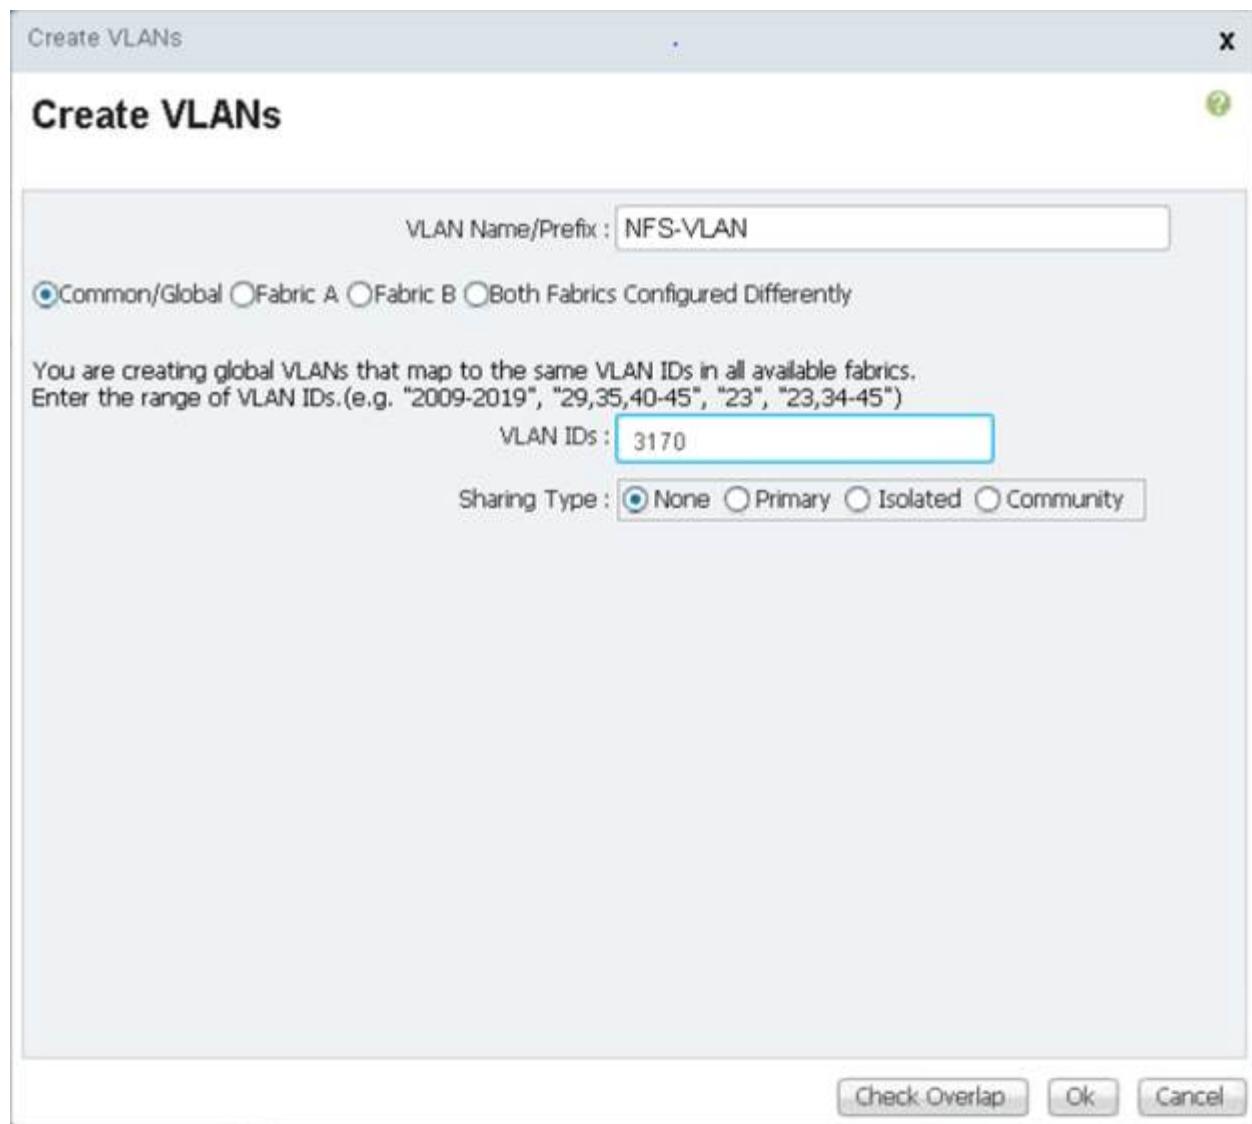
To configure an organization in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Optional: Enter a description for the organization. Click OK.
4. Click OK in the confirmation message.

Configure storage appliance ports and storage VLANs

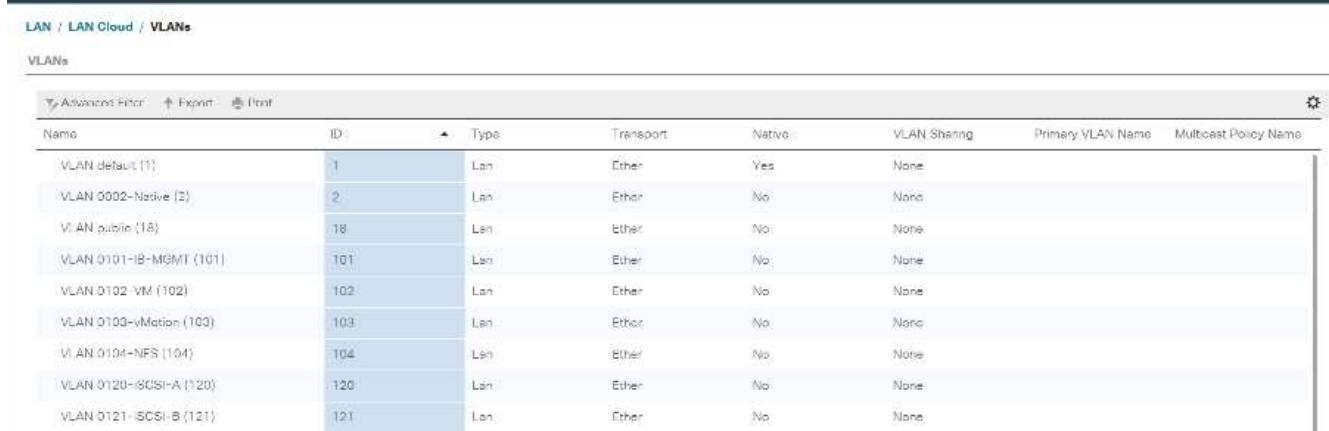
To configure the storage appliance ports and storage VLANs, complete the following steps:

1. In the Cisco UCS Manager, select the LAN tab.
2. Expand the Appliances cloud.
3. Right-click VLANs under Appliances Cloud.
4. Select Create VLANs.
5. Enter NFS-VLAN as the name for the Infrastructure NFS VLAN.
6. Leave Common/Global selected.
7. Enter <<var_nfs_vlan_id>> for the VLAN ID.
8. Leave Sharing Type set to None.



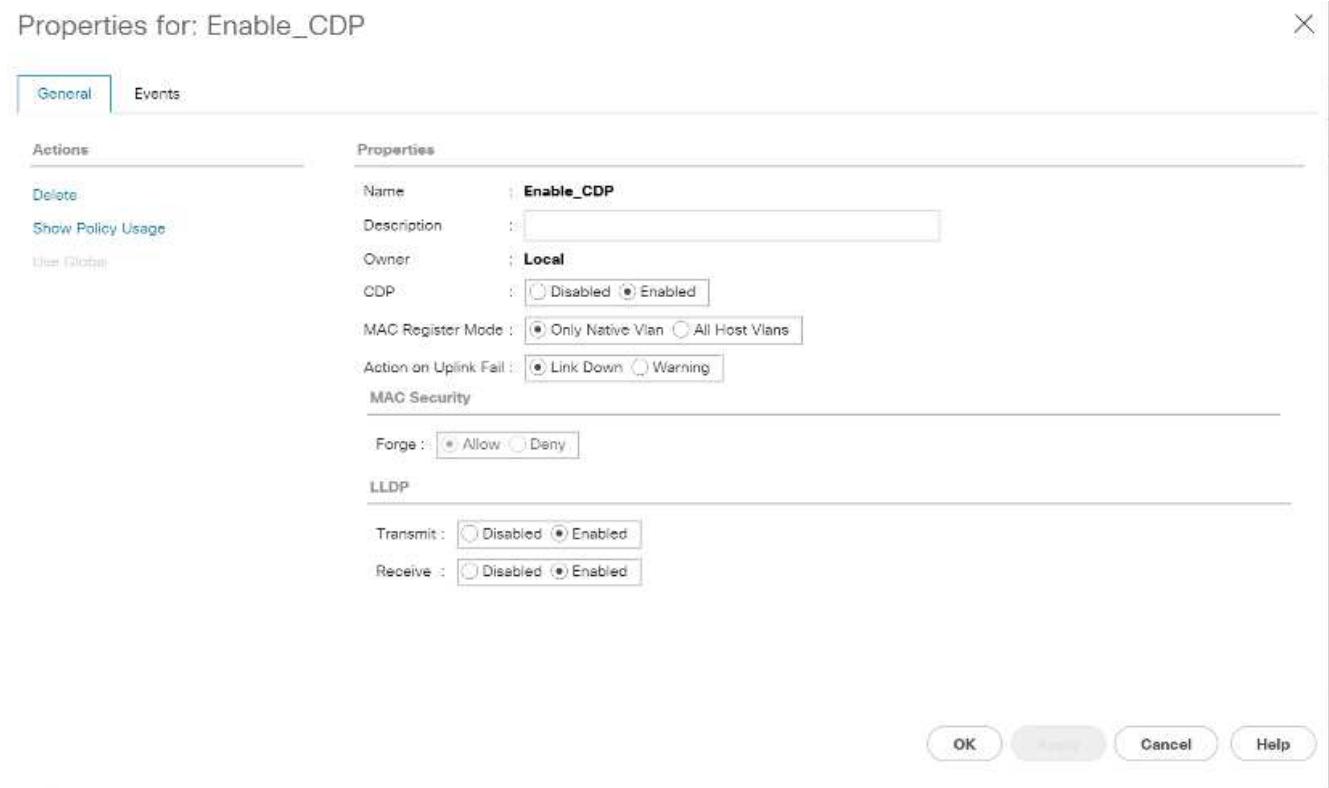
9. Click OK, and then click OK again to create the VLAN.
10. Right-click VLANs under Appliances Cloud.
11. Select Create VLANs.
12. Enter iSCSI-A-VLAN as the name for the Infrastructure iSCSI Fabric A VLAN.
13. Leave Common/Global selected.
14. Enter <<var_iscsi-a_vlan_id>> for the VLAN ID.
15. Click OK, and then click OK again to create the VLAN.
16. Right-click VLANs under Appliances Cloud.
17. Select Create VLANs.
18. Enter iSCSI-B-VLAN as the name for the Infrastructure iSCSI Fabric B VLAN.
19. Leave Common/Global selected.
20. Enter <<var_iscsi-b_vlan_id>> for the VLAN ID.
21. Click OK, and then click OK again to create the VLAN.
22. Right-click VLANs under Appliances Cloud.

23. Select Create VLANs.
24. Enter Native-VLAN as the name for the Native VLAN.
25. Leave Common/Global selected.
26. Enter <<var_native_vlan_id>> for the VLAN ID.
27. Click OK, and then click OK again to create the VLAN.



Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Name	Multicast Policy Name
VLAN default (1)	1	Lan	Ether	Yes	None		
VLAN 0002-Native (2)	2	Lan	Ether	No	None		
VLAN public (18)	18	Lan	Ether	No	None		
VLAN 0101-IB-MGMT (101)	101	Lan	Ether	No	None		
VLAN 0102-VM (102)	102	Lan	Ether	No	None		
VLAN 0103-vMotion (103)	103	Lan	Ether	No	None		
VLAN 0104-NFS (104)	104	Lan	Ether	No	None		
VLAN 0120-iSCSI-A (120)	120	Lan	Ether	No	None		
VLAN 0121-iSCSI-B (121)	121	Lan	Ether	No	None		

28. In the navigation pane, under LAN > Policies, expand Appliances and right-click Network Control Policies.
29. Select Create Network Control Policy.
30. Name the policy `Enable_CDP_LLDP` and select Enabled next to CDP.
31. Enable the Transmit and Receive features for LLDP.



Properties for: Enable_CDP

General Events

Actions

Delete Show Policy Usage Use Global

Properties

Name: **Enable_CDP**
 Description:
 Owner: **Local**
 CDP: Disabled Enabled
 MAC Register Mode: Only Native Vlan All Host Vlans
 Action on Uplink Fail: Link Down Warning
 MAC Security
 Forge: Allow Deny

LLDP

Transmit: Disabled Enabled
 Receive: Disabled Enabled

OK Apply Cancel Help

32. Click OK and then click OK again to create the policy.

33. In the navigation pane, under LAN > Appliances Cloud, expand the Fabric A tree.
34. Expand Interfaces.
35. Select Appliance Interface 1/3.
36. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_01_name>:e0e. Click Save Changes and OK.
37. Select the Enable_CDP Network Control Policy and select Save Changes and OK.
38. Under VLANs, select the iSCSI-A-VLAN, NFS VLAN, and Native VLAN. Set the Native-VLAN as the Native VLAN. Clear the default VLAN selection.
39. Click Save Changes and OK.

40. Select Appliance Interface 1/4 under Fabric A.
41. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_02_name>:e0e. Click Save Changes and OK.
42. Select the Enable_CDP Network Control Policy and select Save Changes and OK.
43. Under VLANs, select the iSCSI-A-VLAN, NFS VLAN, and Native VLAN.
44. Set the Native-VLAN as the Native VLAN.
45. Clear the default VLAN selection.
46. Click Save Changes and OK.
47. In the navigation pane, under LAN > Appliances Cloud, expand the Fabric B tree.
48. Expand Interfaces.
49. Select Appliance Interface 1/3.
50. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_01_name>:e0f. Click Save Changes and OK.
51. Select the Enable_CDP Network Control Policy and select Save Changes and OK.
52. Under VLANs, select the iSCSI-B-VLAN, NFS VLAN, and Native VLAN. Set the Native-VLAN as the Native VLAN. Unselect the default VLAN.

Actions

- Enable Interface
- Disable Interface
- Add Ethernet Target Endpoint
- Delete Ethernet Target Endpoint

Properties

ID: 3
Slot ID: 1
Fabric ID: B
Aggregated Port ID: 0
User Label: AFFA200_Clus_01:e0f
Transport Type: Ether
Port: sys/switch-B/slot-1/switch-ether/port-3
Admin Speed(gbps): 1 Gbps (selected), 10 Gbps, 40 Gbps, 25 Gbps, 100 Gbps, Auto
Priority: Best Effort
Pin Group: <not set>
Network Control Policy: Enable_CDP
Flow Control Policy: default
VLANs:

Port Mode: Trunk (selected), Access

VLAN default (1)
 VLAN iSCSI-A-VLAN (124)
 VLAN iSCSI-B-VLAN (125)
 VLAN Native-VLAN (2)
 VLAN NFS_VLAN (104)

Native VLAN: VLAN Native-VLAN (2)

Create VLAN

- Click Save Changes and OK.
- Select Appliance Interface 1/4 under Fabric B.
- In the User Label field, put in information indicating the storage controller port, such as <storage_controller_02_name>:e0f. Click Save Changes and OK.
- Select the Enable_CDP Network Control Policy and select Save Changes and OK.
- Under VLANs, select the iSCSI-B-VLAN, NFS VLAN, and Native VLAN. Set the Native-VLAN as the Native VLAN. Unselect the default VLAN.
- Click Save Changes and OK.

Set jumbo frames in Cisco UCS fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

- In Cisco UCS Manager, in the navigation pane, click the LAN tab.
- Select LAN > LAN Cloud > QoS System Class.
- In the right pane, click the General tab.
- On the Best Effort row, enter 9216 in the box under the MTU column.

5. Click Save Changes.

6. Click OK.

Acknowledge Cisco UCS chassis

To acknowledge all Cisco UCS chassis, complete the following steps:

1. In Cisco UCS Manager, select the Equipment tab, then Expand the Equipment tab on the right.
2. Expand Equipment > Chassis.
3. In the Actions for Chassis 1, select Acknowledge Chassis.
4. Click OK and then click OK to complete acknowledging the chassis.
5. Click Close to close the Properties window.

Load Cisco UCS 4.0(1b) firmware images

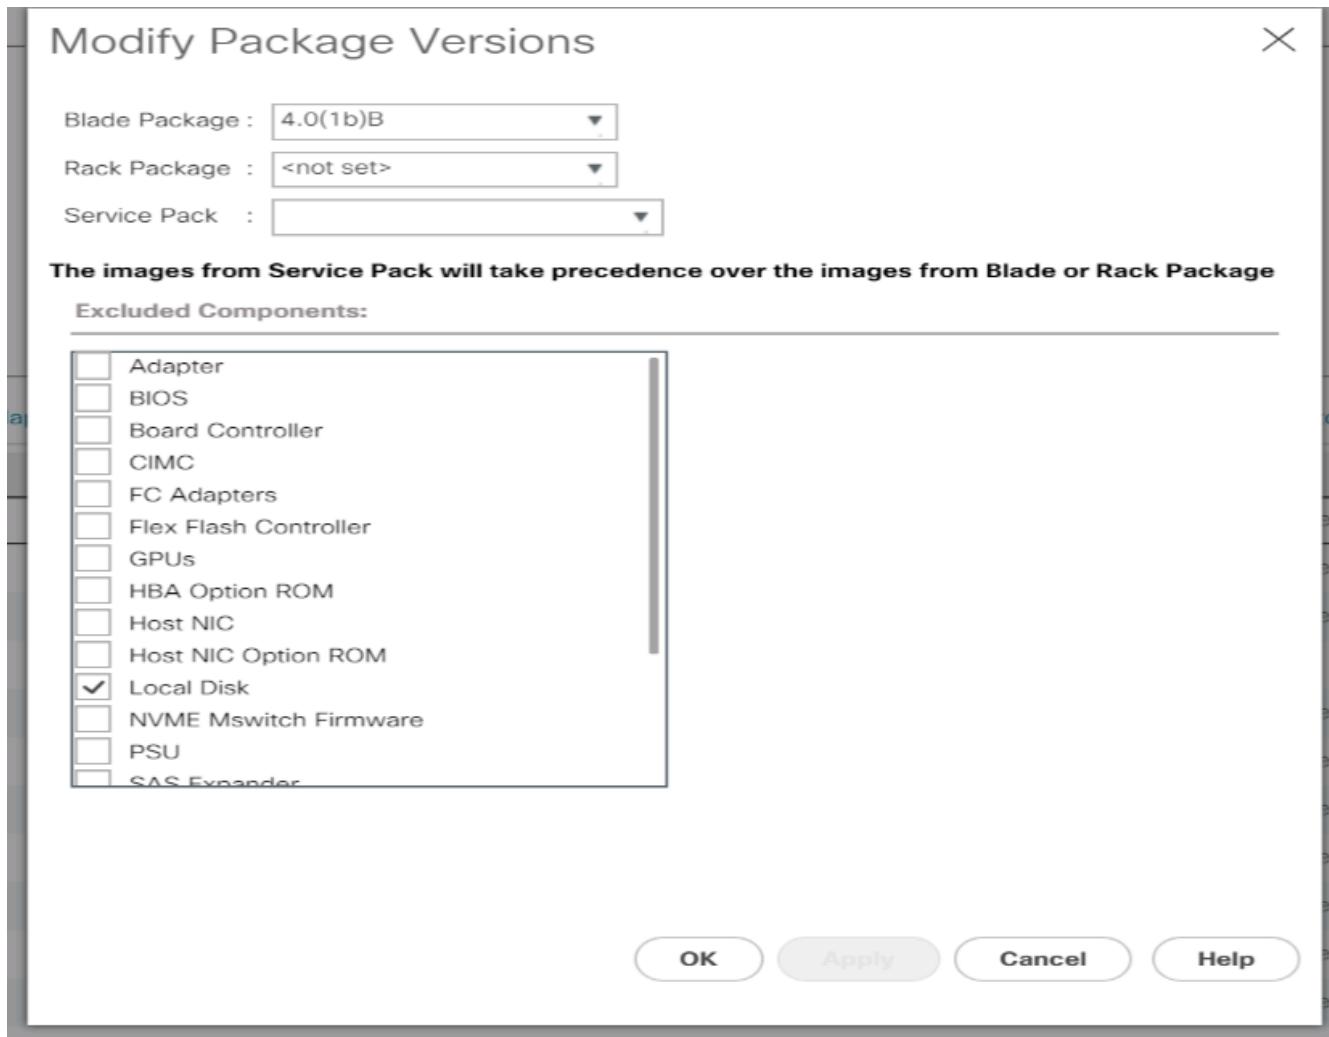
To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.0(1b) refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Create host firmware package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 4.0(1b) for both the Blade Packages.



7. Click OK then OK again to modify the host firmware package.

Create MAC address pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.

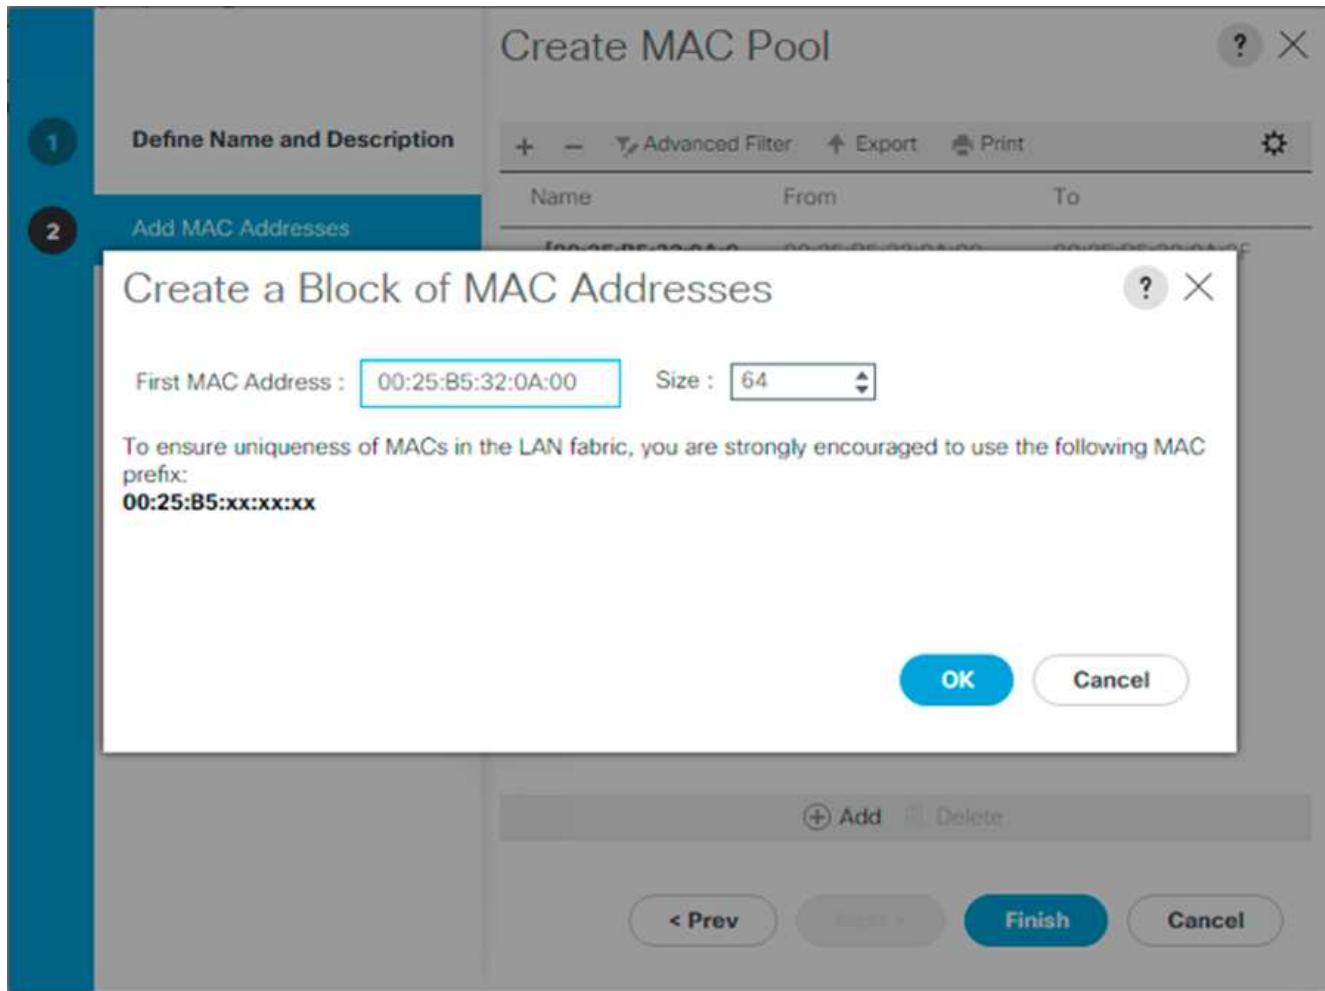
In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC-Pool-A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select Sequential as the option for Assignment Order. Click Next.
8. Click Add.
9. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0A:00 as our first MAC address.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources. Click OK.



11. Click Finish.
12. In the confirmation message, click OK.
13. Right-click MAC Pools under the root organization.
14. Select Create MAC Pool to create the MAC address pool.
15. Enter MAC-Pool-B as the name of the MAC pool.
16. Optional: Enter a description for the MAC pool.
17. Select Sequential as the option for Assignment Order. Click Next.
18. Click Add.
19. Specify a starting MAC address.



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0B:00 as our first MAC address.

20. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources. Click OK.
21. Click Finish.
22. In the confirmation message, click OK.

Create iSCSI IQN pool

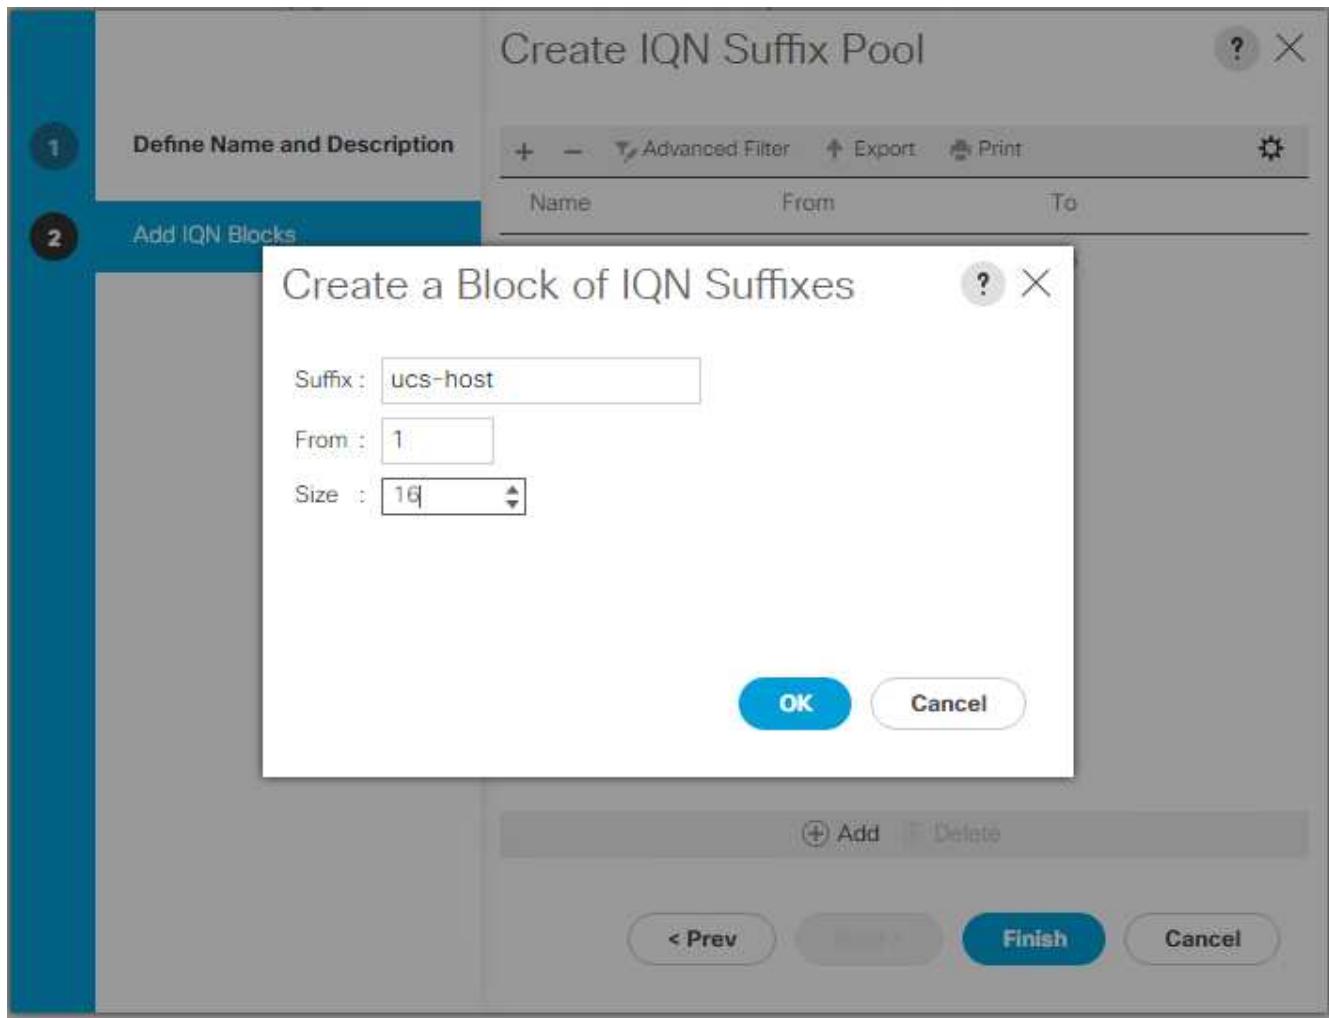
To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. Right- click IQN Pools.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter `iqn.1992-08.com.cisco` as the prefix.
8. Select Sequential for Assignment Order. Click Next.
9. Click Add.
10. Enter `ucs-host` as the suffix.



If multiple Cisco UCS domains are being used, a more specific IQN suffix might need to be used.

11. Enter 1 in the From field.
12. Specify the size of the IQN block sufficient to support the available server resources. Click OK.



13. Click Finish.

Create iSCSI initiator IP address pools

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.
3. Right-click IP Pools.
4. Select Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Select Sequential for the assignment order. Click Next.
8. Click Add to add a block of IP address.
9. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
10. Set the size to enough addresses to accommodate the servers. Click OK.
11. Click Next.
12. Click Finish.

13. Right-click IP Pools.
14. Select Create IP Pool.
15. Enter iSCSI-IP-Pool-B as the name of IP pool.
16. Optional: Enter a description for the IP pool.
17. Select Sequential for the assignment order. Click Next.
18. Click Add to add a block of IP address.
19. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
20. Set the size to enough addresses to accommodate the servers. Click OK.
21. Click Next.
22. Click Finish.

Create UUID suffix pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID-Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources. Click OK.
13. Click Finish.
14. Click OK.

Create server pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



Consider creating unique server pools to achieve the granularity that is required in your environment.

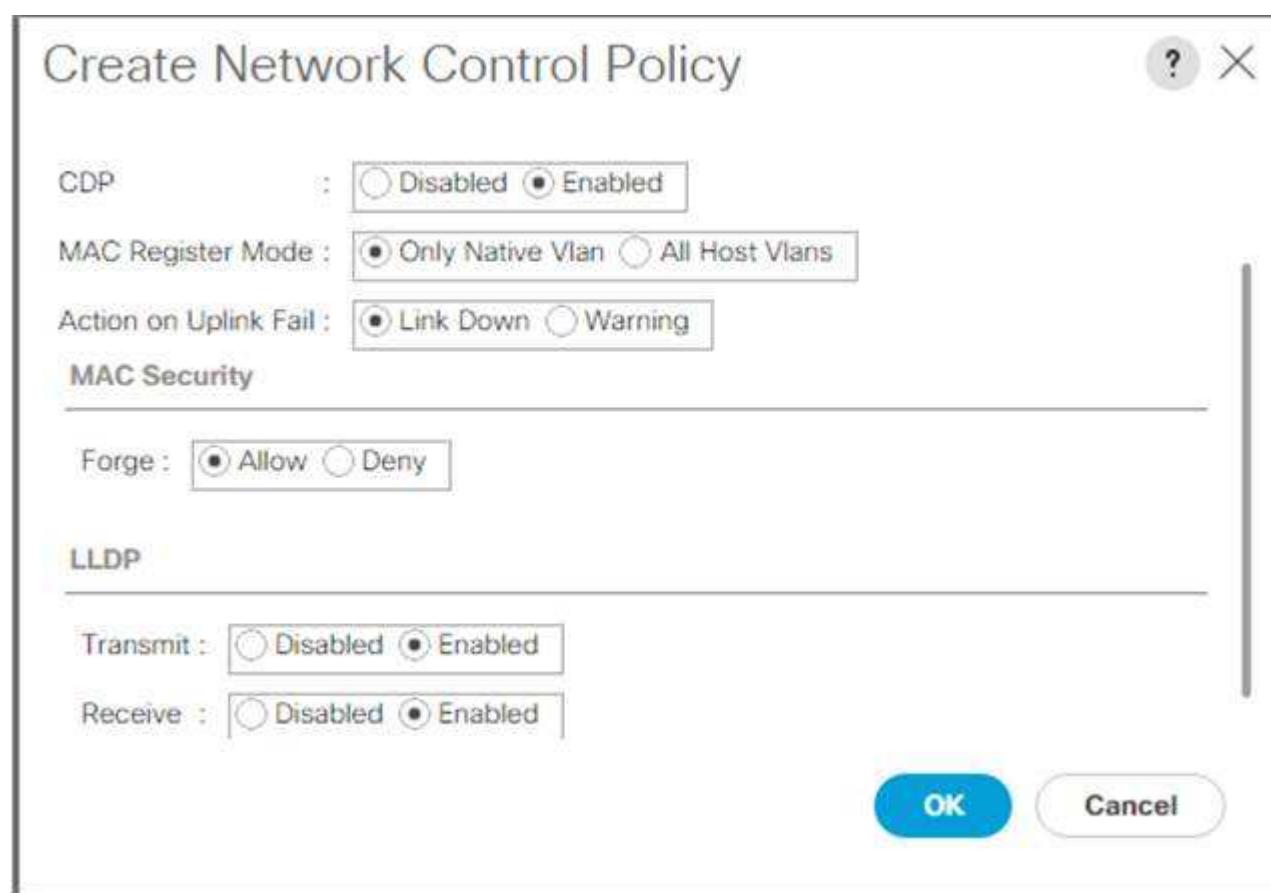
1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click Server Pools.

4. Select Create Server Pool.
5. Enter 'Infra-Pool' as the name of the server pool.
6. Optional: Enter a description for the server pool. Click Next.
7. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the 'Infra-Pool' server pool.
8. Click Finish.
9. Click OK.

Create Network Control Policy for Cisco Discovery Protocol and Link Layer Discovery Protocol

To create a Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP), complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable-CDP-LLDP policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy. Click OK.



Create power control policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab on the left.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy. Click OK.

Create Power Control Policy

Name :

Description :

Fan Speed Policy :

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

No Cap cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK **Cancel**

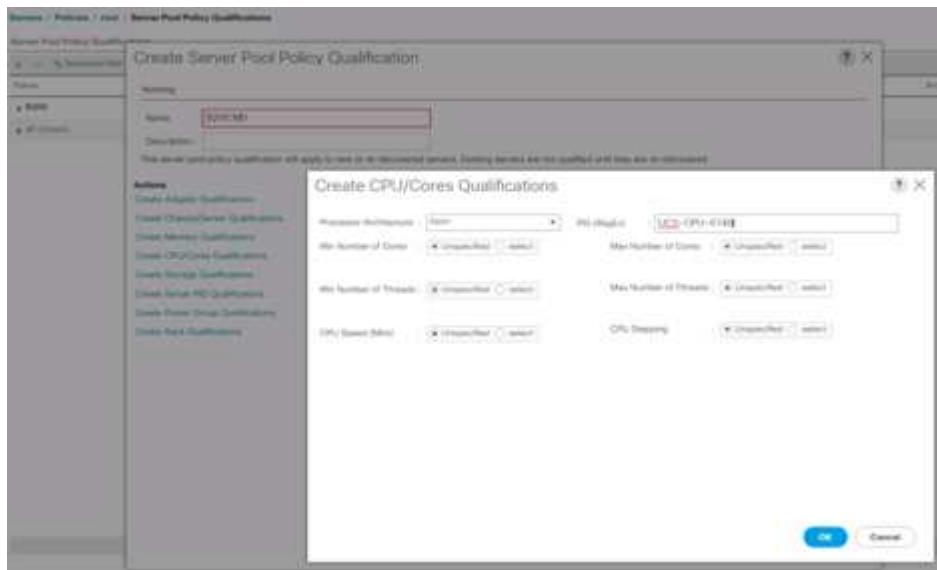
Create server pool qualification policy (Optional)

To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:



This example creates a policy for Cisco UCS B-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

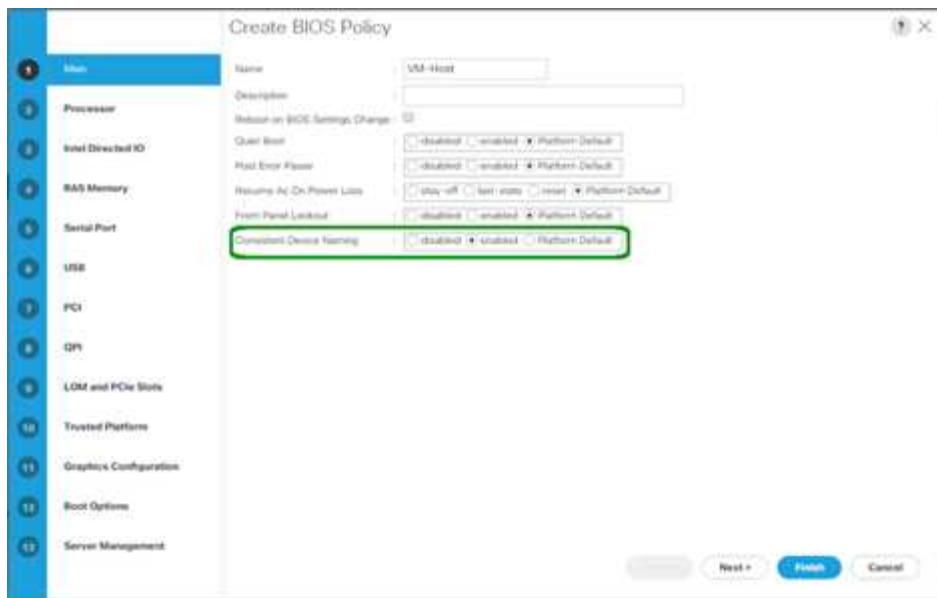
1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification or Add.
5. Name the policy Intel.
6. Select Create CPU/Cores Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Enter <UCS-CPU- PID> as the process ID (PID).
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy, and then click OK for the confirmation.



Create server BIOS policy

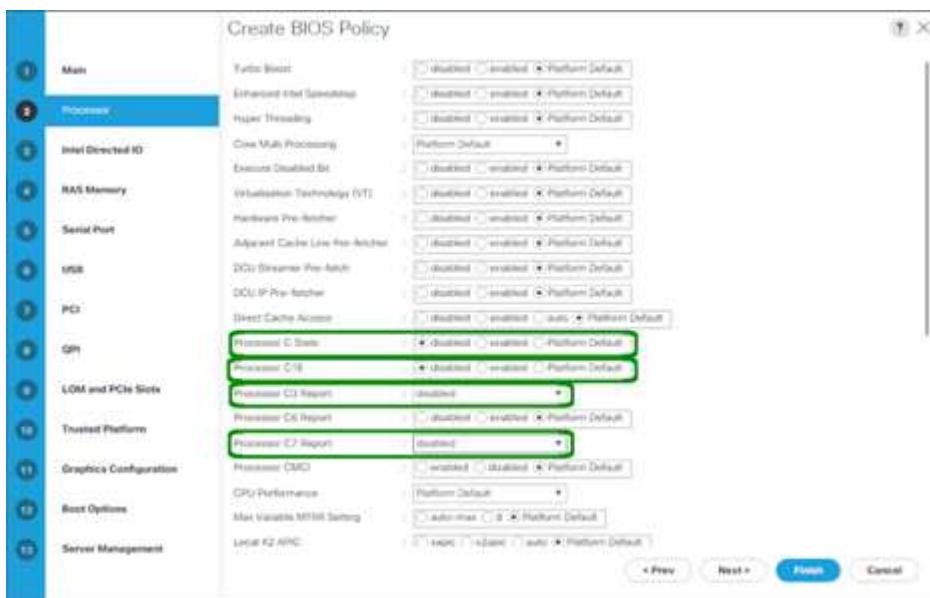
To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter VM-Host as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.



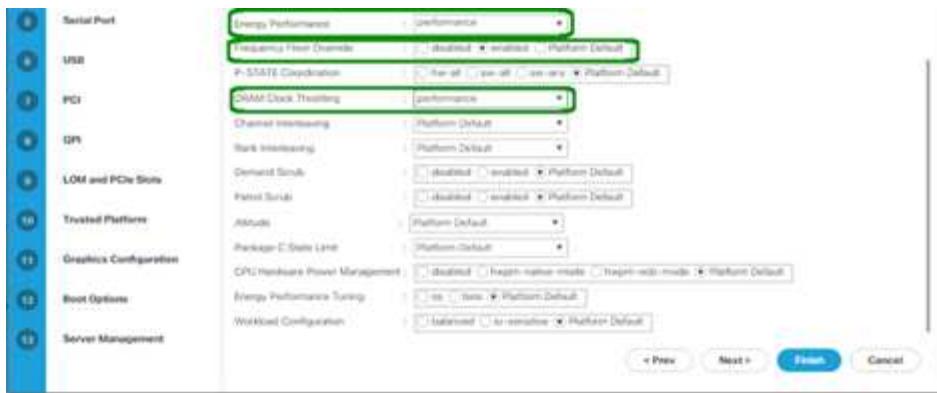
8. Select the Processor tab and set the following parameters:

- Processor C State: disabled
- Processor C1E: disabled
- Processor C3 Report: disabled
- Processor C7 Report: disabled



9. Scroll down to the remaining Processor options and set the following parameters:

- Energy Performance: performance
- Frequency Floor Override: enabled
- DRAM Clock Throttling: performance



10. Click RAS Memory and set the following parameters:

- LV DDR Mode: performance mode



11. Click Finish to create the BIOS policy.

12. Click OK.

Update the default maintenance policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select On Next Boot to delegate maintenance windows to server administrators.

6. Click Save Changes.
7. Click OK to accept the change.

Create vNIC templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the procedures described in this section.



A total of four vNIC templates are created.

Create infrastructure vNICs

To create an infrastructure vNIC, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter Site-XX-vNIC_A as the vNIC template name.
6. Select updating-template as the Template Type.
7. For Fabric ID, select Fabric A.
8. Ensure that the Enable Failover option is not selected.
9. Select Primary Template for Redundancy Type.
10. Leave the Peer Redundancy Template set to <not set>.
11. Under Target, make sure that only the Adapter option is selected.
12. Set Native-VLAN as the native VLAN.
13. Select vNIC Name for the CDN Source.
14. For MTU, enter 9000.
15. Under Permitted VLANs, select Native-VLAN, Site-XX-IB-MGMT, Site-XX-NFS, Site-XX-VM-Traffic, and Site-XX-vMotion. Use the Ctrl key to make this multiple selection.
16. Click Select. These VLANs should now appear under Selected VLANs.
17. In the MAC Pool list, select MAC_Pool_A.

18. In the Network Control Policy list, select Pool-A.
19. In the Network Control Policy list, select Enable-CDP-LLDP.
20. Click OK to create the vNIC template.
21. Click OK.

The screenshot shows the 'vNIC Templates' section in the Cisco UCS Manager. A new template, 'vNIC_Template_A', is being created. The 'Template Type' is set to 'Updating Template'. The 'vNIC Name' is 'User Defined'. The 'Port ID' is '9002'. Under 'Policies', 'MAC Pool' is 'MAC_Pool_AEBSV5', 'QoS Policy' is 'violated', 'Network Control Policy' is 'Create_CLP', and 'Fabric ID' is 'Fabric A'. The 'Enable Failover' checkbox is checked. The 'Failover Template' is set to 'vNIC_Template_B'. The 'Target' section shows 'VM-1' selected. A 'Create vNIC Template' button is visible at the bottom right.

To create the secondary redundancy template Infra-B, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter 'Site-XX-vNIC_B' as the vNIC template name.
6. Select updating-template as the Template Type.
7. For Fabric ID, select Fabric B.
8. Select the Enable Failover option.



Selecting Failover is a critical step to improve link failover time by handling it at the hardware level, and to guard against any potential for NIC failure not being detected by the virtual switch.

9. Select Primary Template for Redundancy Type.
10. Leave the Peer Redundancy Template set to vNIC_Template_A.
11. Under Target, make sure that only the Adapter option is selected.
12. Set Native-VLAN as the native VLAN.
13. Select vNIC Name for the CDN Source.
14. For MTU, enter 9000.
15. Under Permitted VLANs, select Native-VLAN, Site-XX-IB-MGMT, Site-XX-NFS, Site-XX-VM-Traffic, and Site-XX-vMotion. Use the Ctrl key to make this multiple selection.
16. Click Select. These VLANs should now appear under Selected VLANs.
17. In the MAC Pool list, select MAC_Pool_B.
18. In the Network Control Policy list, select Pool-B.
19. In the Network Control Policy list, select Enable-CDP-LLDP.
20. Click OK to create the vNIC template.
21. Click OK.

LAN / Policies / root / vNIC Templates / vNIC Template vNIC_Template_B

Actions: Modify VLAN, Modify VLAN Groups, Delete, Show Policy Usage, Use Wizard

Properties:

- Name: vNIC_Template_B
- Description:
- Owner: Local
- Fabric ID: Fabric A Fabric B Enable Fabric
- Redundancy Type: No Redundancy Primary Template Secondary Template
- Peer Redundancy Template: vNIC_Template_A

Target:

- Adapter
- VM

Template Type: Initial Template Updating Template

CDN Source: vNIC Name User Defined

MTU: 9000

Policies:

- MAC Pool: MAC Pool_B(56/54)
- QoS Policy: vSphere QoS
- Network Control Policy: Enable CDP
- Rin Group: vSphere RIN
- QoS Threshold Policy: vSphere

Connection Policies:

- Dynamic vNIC sVNIC VMQ
- Dynamic vNIC Connection Policy: vSphere vNIC

Create vNIC Template

Create iSCSI vNICs

To create iSCSI vNICs, complete the following steps:

1. Select LAN on the left.

2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter Site- 01-iSCSI_A as the vNIC template name.
6. Select Fabric A. Do not select the Enable Failover option.
7. Leave Redundancy Type set at No Redundancy.
8. Under Target, make sure that only the Adapter option is selected.
9. Select Updating Template for Template Type.
10. Under VLANs, select only Site- 01-iSCSI_A_VLAN.
11. Select Site- 01-iSCSI_A_VLAN as the native VLAN.
12. Leave vNIC Name set for the CDN Source.
13. Under MTU, enter 9000.
14. From the MAC Pool list, select MAC-Pool-A.
15. From the Network Control Policy list, select Enable-CDP-LLDP.
16. Click OK to complete creating the vNIC template.
17. Click OK.

LAN / Policies / root / vNIC Templates / vNIC Template Site_01_ISCSI-A

General VLANs VLAN Groups Faults Events

Actions

- Modify VLANs
- Modify VLAN Groups
- Delete
- Show Policy Usage
- Use Global

Properties

Name: Site_01_ISCSI-A
 Description:
 Owner: Local
 Fabric ID: Fabric A Fabric B Enable Failover
 Redundancy Type: No Redundancy Primary Template Secondary Template

Target

- Adapter
- VM

Template Type: Initial Template Updating Template
CDN Source: vNIC Name User Defined
MTU: 9000

Policies

MAC Pool: MAC_Pool_A(56/64)
 QoS Policy: <not set>
 Network Control Policy: Enable_CDP
 Pin Group: <not set>
 Stats Threshold Policy: default

Connection Policies

Dynamic vNIC usNIC VMQ
 Dynamic vNIC Connection Policy: <not set>

18. Select LAN on the left.
19. Select Policies > root.
20. Right-click vNIC Templates.
21. Select Create vNIC Template.
22. Enter **Site- 01-iSCSI_B** as the vNIC template name.
23. Select Fabric B. Do not select the Enable Failover option.
24. Leave Redundancy Type set at No Redundancy.
25. Under Target, make sure that only the Adapter option is selected.
26. Select Updating Template for Template Type.
27. Under VLANs, select only **Site- 01-iSCSI_B_VLAN**.
28. Select **Site- 01-iSCSI_B_VLAN** as the native VLAN.
29. Leave vNIC Name set for the CDN Source.
30. Under MTU, enter 9000.
31. From the MAC Pool list, select **MAC-Pool-B**.
32. From the Network Control Policy list, select **Enable-CDP-LLDP**.
33. Click OK to complete creating the vNIC template.
34. Click OK.

General VLANs VLAN Groups Fautes Events

Actions

- Modify VLANs
- Modify VLAN Groups
- Delete
- Show Policy Usage
- Link Global

Properties

Name: Site_01_ISCSI-B
 Description:
 Owner: Local
 Fabric ID: Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type: No Redundancy Primary Template Secondary Template

Target

Adapter
 VM

Template Type: Initial Template Updating Template

CDN Source: vNIC Name User Defined

MTU: 9000

Policies

MAC Pool: MAC_Pool_B[64/64]
 QoS Policy: <not set>
 Network Control Policy: Enable_CDP
 Pin Group: <not set>
 Stats Threshold Policy: default

Connection Policies

Dynamic vNIC usNIC VMQ
 Dynamic vNIC Connection Policy: <not set>

Create LAN connectivity policy for iSCSI boot

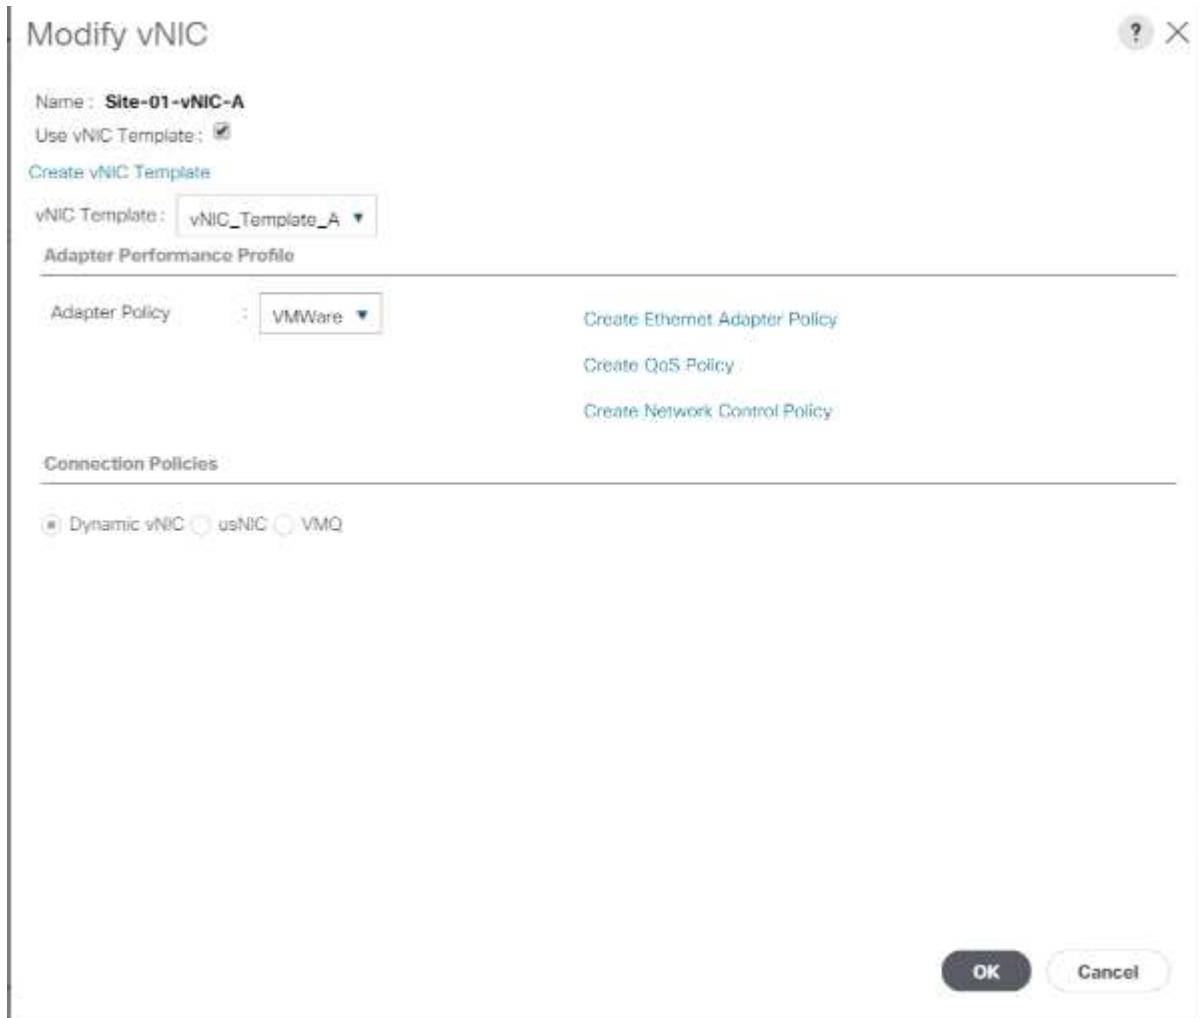
This procedure applies to a Cisco UCS environment in which two iSCSI LIFs are on cluster node 1 (iscsi_lif01a and iscsi_lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi_lif02a and iscsi_lif02b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS 6324 A) and the B LIFs are connected to Fabric B (Cisco UCS 6324 B).

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter Site-XX-Fabric-A as the name of the policy.
6. Click the upper Add option to add a vNIC.
7. In the Create vNIC dialog box, enter Site-01-vNIC-A as the name of the vNIC.
8. Select the Use vNIC Template option.
9. In the vNIC Template list, select vNIC_Template_A.

10. From the Adapter Policy drop-down list, select VMWare.

11. Click OK to add this vNIC to the policy.



12. Click the upper Add option to add a vNIC.

13. In the Create vNIC dialog box, enter Site-01-vNIC-B as the name of the vNIC.

14. Select the Use vNIC Template option.

15. In the vNIC Template list, select vNIC_Template_B.

16. From the Adapter Policy drop-down list, select VMWare.

17. Click OK to add this vNIC to the policy.

18. Click the upper Add option to add a vNIC.

19. In the Create vNIC dialog box, enter Site-01-iSCSI-A as the name of the vNIC.

20. Select the Use vNIC Template option.

21. In the vNIC Template list, select Site-01-iSCSI-A.

22. From the Adapter Policy drop-down list, select VMWare.

23. Click OK to add this vNIC to the policy.

24. Click the upper Add option to add a vNIC.

25. In the Create vNIC dialog box, enter `Site-01-iSCSI-B` as the name of the vNIC.
26. Select the Use vNIC Template option.
27. In the vNIC Template list, select `Site-01-iSCSI-B`.
28. From the Adapter Policy drop-down list, select VMWare.
29. Click OK to add this vNIC to the policy.
30. Expand the Add iSCSI vNICs option.
31. Click the Lower Add option in the Add iSCSI vNICs space to add the iSCSI vNIC.
32. In the Create iSCSI vNIC dialog box, enter `Site-01-iSCSI-A` as the name of the vNIC.
33. Select the Overlay vNIC as `Site-01-iSCSI-A`.
34. Leave the iSCSI Adapter Policy option to Not Set.
35. Select the VLAN as `Site-01-iSCSI-Site-A (native)`.
36. Select None (used by default) as the MAC address assignment.
37. Click OK to add the iSCSI vNIC to the policy.

Modify iSCSI vNIC

Name : **Site-01-ISCSI-A**

Overlay vNIC : **Site-01-ISCSI-A**

iSCSI Adapter Policy : <not set> [Create iSCSI Adapter Policy](#)

VLAN : **Site_01_ISCSI-A (native)**

iSCSI MAC Address

MAC Address Assignment: **Select(None used by default)**

[Create MAC Pool](#)

OK **Cancel**

38. Click the Lower Add option in the Add iSCSI vNICs space to add the iSCSI vNIC.
39. In the Create iSCSI vNIC dialog box, enter `Site-01-ISCSI-B` as the name of the vNIC.
40. Select the Overlay vNIC as Site-01-ISCSI-B.
41. Leave the iSCSI Adapter Policy option to Not Set.
42. Select the VLAN as Site-01-ISCSI-Site-B (native).
43. Select None(used by default) as the MAC Address Assignment.
44. Click OK to add the iSCSI vNIC to the policy.
45. Click Save Changes.

The screenshot shows the 'LAN Policies' interface in Cisco UCS Manager. A new policy named 'Site01-iSCSIBoot' is being created. The 'Description' field is empty, and the 'Owner' is set to 'Local'. The 'Native VLAN' is set to 'Native'. Under the 'VMs' tab, four vNICs are listed: 'vNIC Site-01-iSCSI-A' (Defined), 'vNIC Site-01-iSCSI-B' (Defined), 'vNIC Site-01-vNIC-A' (Defined), and 'vNIC Site-01-vNIC-B' (Defined). Below this, the 'iSCSI vNICs' tab shows two entries: 'iSCSI vNIC Site-01-iSCSI-A' (Defined) and 'iSCSI vNIC Site-01-iSCSI-B' (Defined). Both entries have the same 'Overlay vNIC Name' of 'Site-01-iSCSI-A' and 'iSCSI Adapter Policy' of 'Site-01-iSCSI-B'.

Create vMedia policy for VMware ESXi 6.7U1 install boot

In the NetApp Data ONTAP setup steps an HTTP web server is required, which is used for hosting NetApp Data ONTAP as well as VMware software. The vMedia Policy created here maps the VMware ESXi 6.7U1 ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, complete the following steps:

1. In Cisco UCS Manager, select Servers on the left.
2. Select Policies > root.
3. Select vMedia Policies.
4. Click Add to create new vMedia Policy.
5. Name the policy ESXi-6.7U1-HTTP.
6. Enter Mounts ISO for ESXi 6.7U1 in the Description field.
7. Select Yes for Retry on Mount failure.
8. Click Add.
9. Name the mount ESXi-6.7U1-HTTP.
10. Select the CDD Device Type.
11. Select the HTTP Protocol.
12. Enter the IP Address of the web server.



The DNS server IPs were not entered into the KVM IP earlier, therefore, it is necessary to enter the IP of the web server instead of the hostname.

13. Enter `VMware-VMvisor-Installer-6.7.0.update01-10302608.x86_64.iso` as the Remote File name.

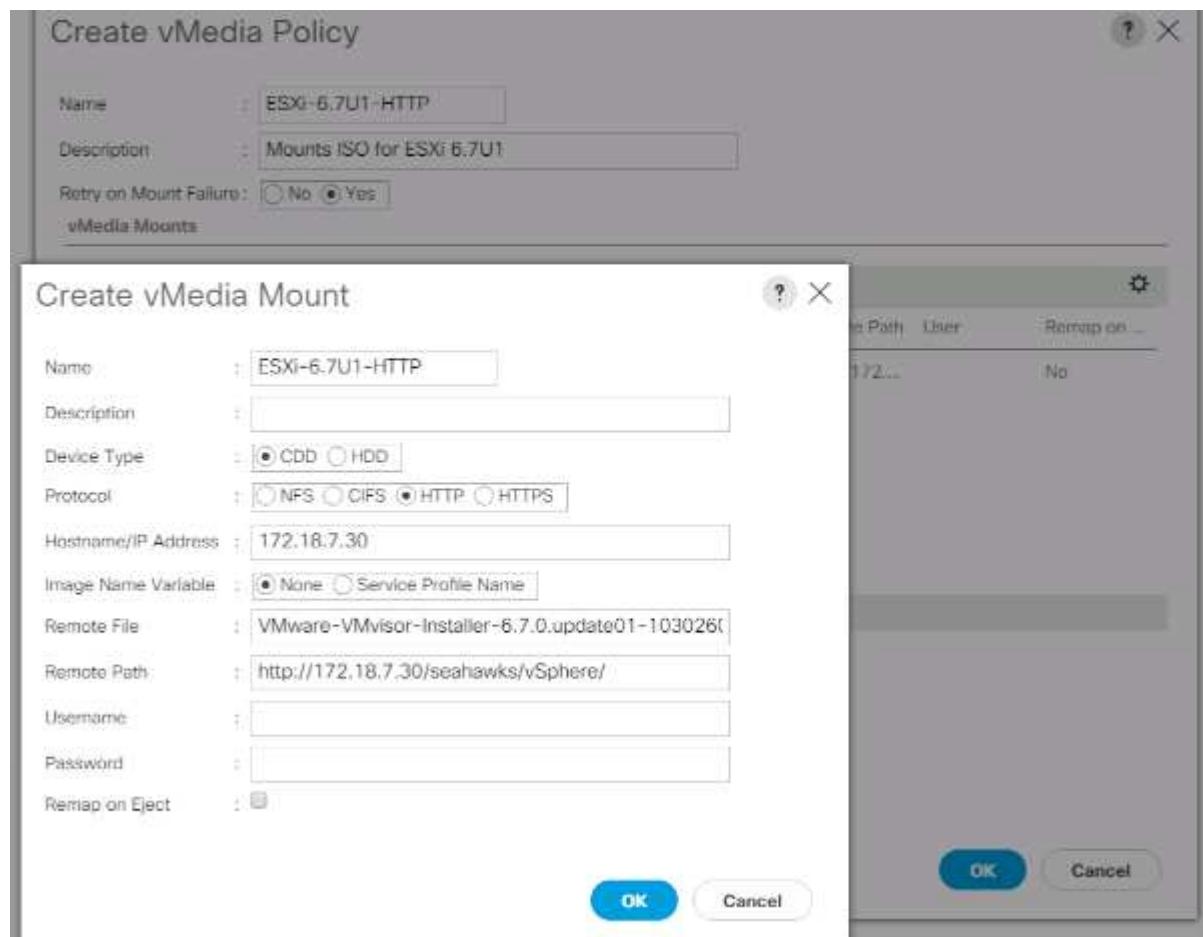
This VMware ESXi 6.7U1 ISO can be downloaded from [VMware Downloads](#).

14. Enter the web server path to the ISO file in the Remote Path field.

15. Click OK to create the vMedia Mount.

16. Click OK then OK again to complete creating the vMedia Policy.

For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi host. On first boot, the host boots into the ESXi installer since the SAN mounted disk is empty. After ESXi is installed, the vMedia is not referenced as long as the boot disk is accessible.



Create iSCSI boot policy

The procedure in this section applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (`iscsi_lif01a` and `iscsi_lif01b`) and two iSCSI LIFs are on cluster node 2 (`iscsi_lif02a` and `iscsi_lif02b`). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).

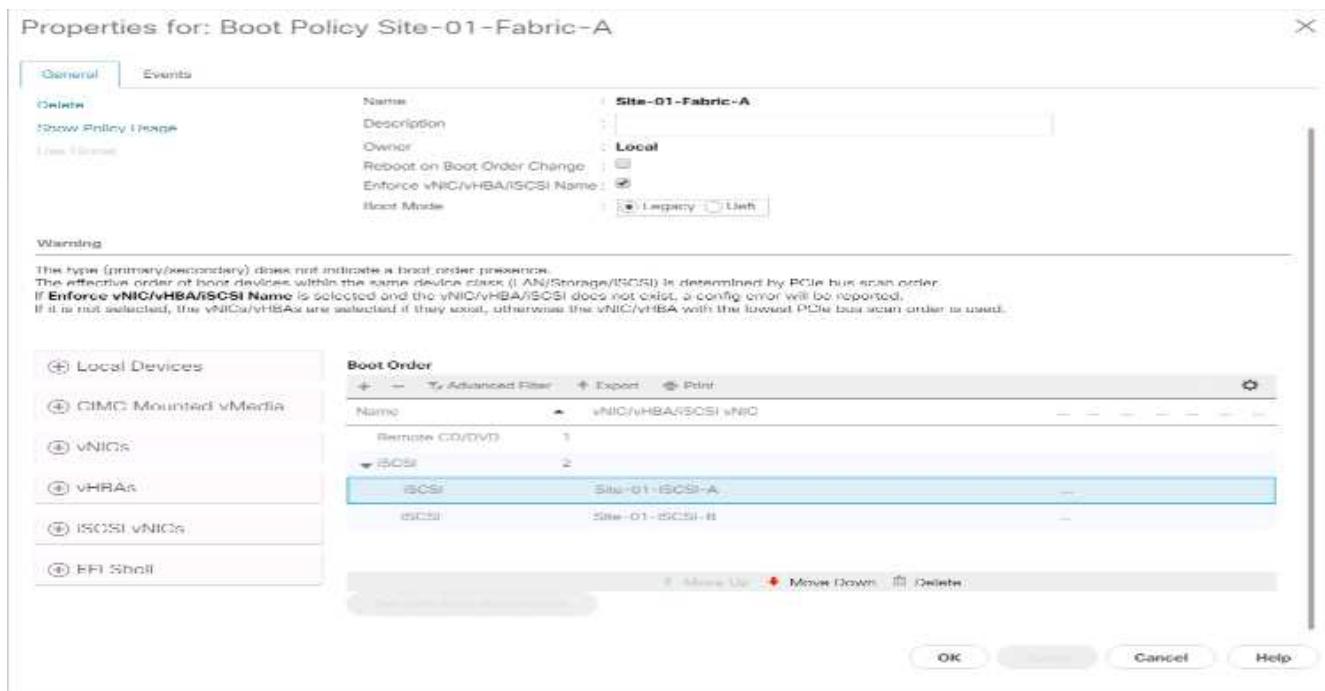


One boot policy is configured in this procedure. The policy configures the primary target to be `iscsi_lif01a`.

To create a boot policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.

5. Enter Site-01-Fabric-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Boot Mode is Legacy.
9. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
10. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
11. In the Add iSCSI Boot dialog box, enter Site-01-iSCSI-A. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter Site-01-iSCSI-B. Click OK.
14. Click OK to create the policy.

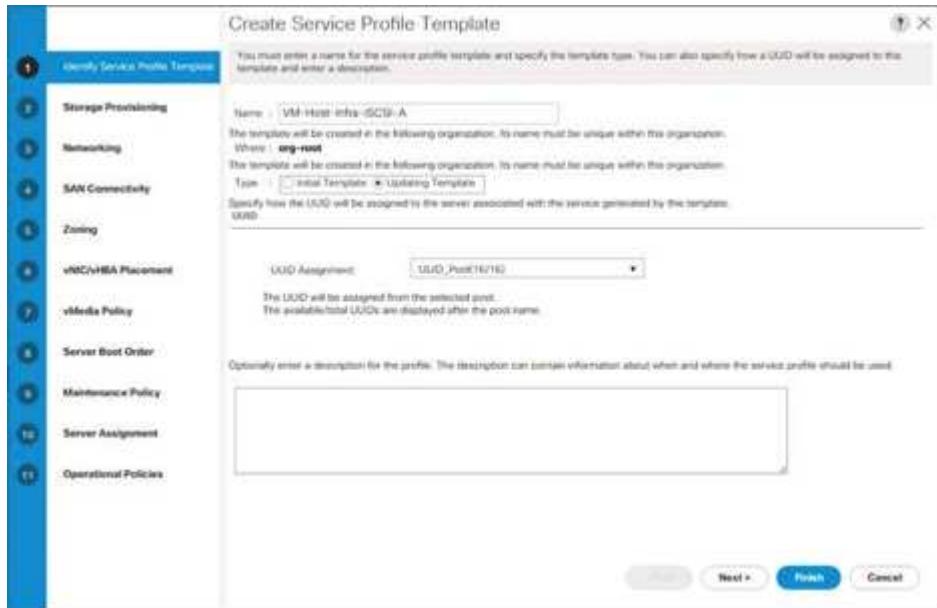


Create service profile template

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the Updating Template option.
7. Under UUID, select UUID_Pool as the UUID pool. Click Next.



Configure storage provisioning

To configure storage provisioning, complete the following steps:

1. If you have servers with no physical disks, click Local Disk Configuration Policy and select the SAN Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

Configure networking options

To configure the networking options, complete the following steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Use Connectivity Policy option to configure the LAN connectivity.
3. Select iSCSI-Boot from the LAN Connectivity Policy drop-down menu.
4. Select IQN_Pool1 in Initiator Name Assignment. Click Next.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionaly specify LAN configuration information:

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

LAN Connectivity Policy: Site01 iSCSIBoot

Initiator Name

Initiator Name Assignment: IQN_Pool(60/64)

Initiator Name: Create IQN Suffix Pool

The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name...

< Prev Next >

Configure SAN connectivity

To configure SAN connectivity, complete the following steps:

1. For the vHBAs, select No for the How Would you Like to Configure SAN Connectivity? option.
2. Click Next.

Configure zoning

To configure zoning, simply click Next.

Configure vNIC/HBA placement

To configure vNIC/HBA placement, complete the following steps:

1. From the Select Placement drop-down list, leave the placement policy as Let System Perform Placement.
2. Click Next.

Configure vMedia policy

To configure the vMedia policy, complete the following steps:

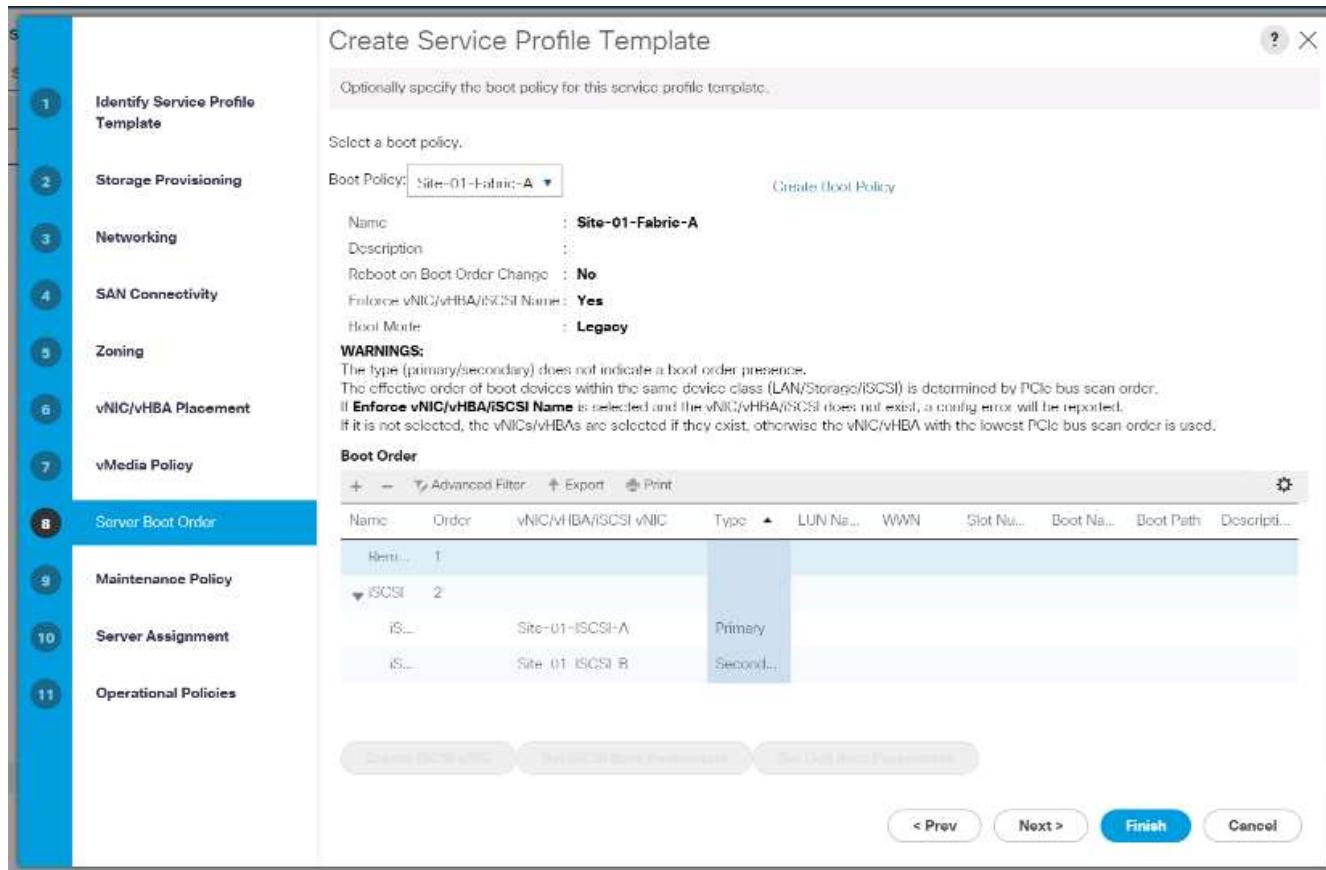
1. Do not select a vMedia Policy.
2. Click Next.



Configure server boot order

To configure the server boot order, complete the following steps:

1. Select Boot-Fabric-A for Boot Policy.



2. In the Boot order, select Site-01- iSCSI-A.
3. Click Set iSCSI Boot Parameters.
4. In the Set iSCSI Boot Parameters dialog box, leave the Authentication Profile option to Not Set unless you have independently created one appropriate for your environment.
5. Leave the Initiator Name Assignment dialog box Not Set to use the single Service Profile Initiator Name defined in the previous steps.
6. Set iSCSI_IP_Pool_A as the Initiator IP address Policy.
7. Select iSCSI Static Target Interface option.
8. Click Add.
9. Enter the iSCSI target name. To get the iSCSI target name of Infra-SVM, log in into storage cluster management interface and run the `iscsi show` command.

```
bb04-aff300::> iscsi show
      Target          Target          Status
  Vaserver  Name          Alias          Admin
-----
Infra-SVM  iqn.1992-08.com.netapp:sn.b5acab9ef1c811e68d9d00a098a9fec2:vs.3
              Infra-SVM          up
```

10. Enter the IP address of `iscsi_lif_02a` for the IPv4 Address field.

Create iSCSI Static Target

iSCSI Target Name :

Priority :

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

[OK](#) [Cancel](#)

11. Click OK to add the iSCSI static target.
12. Click Add.
13. Enter the iSCSI target name.
14. Enter the IP address of `iscsi_lif_01a` for the IPv4 Address field.

Create iSCSI Static Target

iSCSI Target Name :

Priority :

Port :

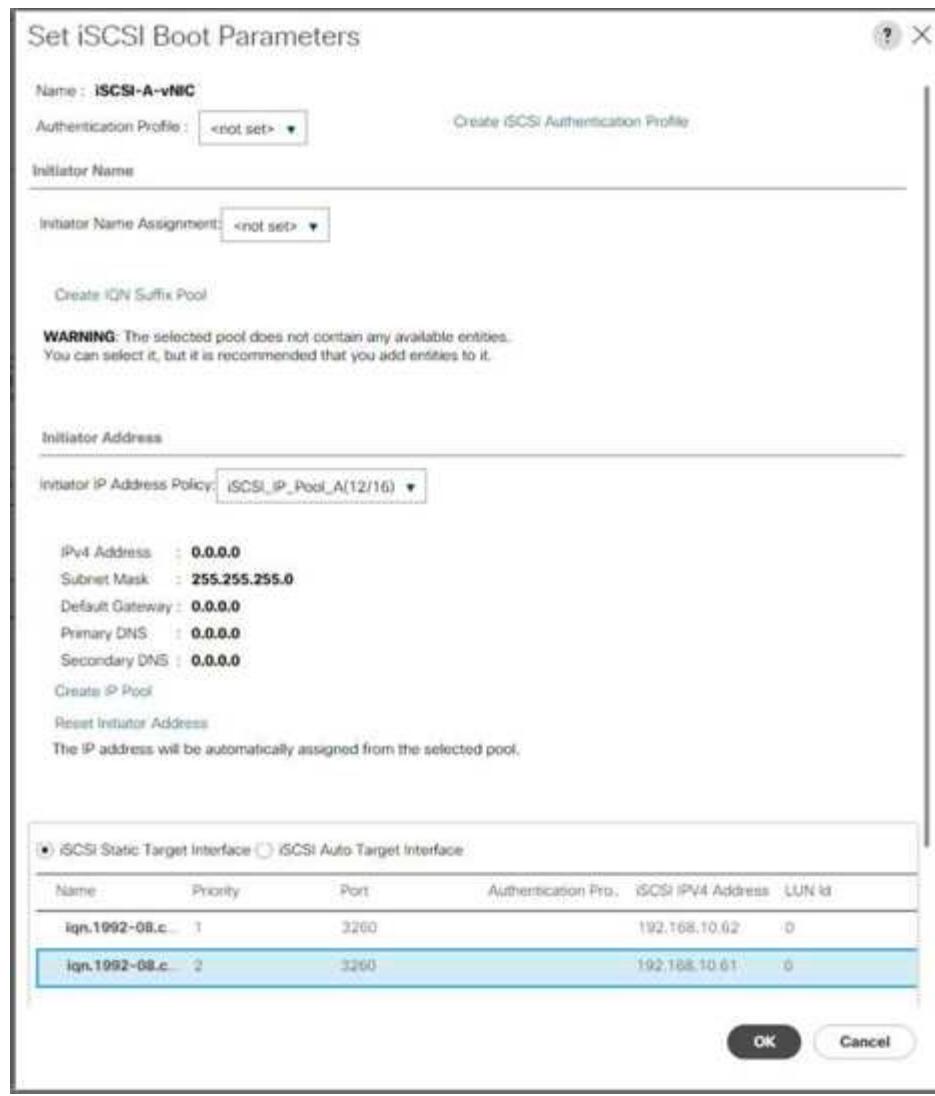
Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

[OK](#) [Cancel](#)

15. Click OK to add the iSCSI static target.



The target IPs were put in with the storage node 02 IP first and the storage node 01 IP second. This is assuming the boot LUN is on node 01. The host boots by using the path to node 01 if the order in this procedure is used.

16. In the Boot order, select iSCSI-B-vNIC.
17. Click Set iSCSI Boot Parameters.
18. In the Set iSCSI Boot Parameters dialog box, leave the Authentication Profile option as Not Set unless you have independently created one appropriate to your environment.
19. Leave the Initiator Name Assignment dialog box Not Set to use the single Service Profile Initiator Name defined in the previous steps.
20. Set iSCSI_IP_Pool_B as the initiator IP address policy.
21. Select the iSCSI Static Target Interface option.
22. Click Add.
23. Enter the iSCSI target name. To get the iSCSI target name of Infra-SVM, log in into storage cluster management interface and run the `iscsi show` command.

Vserver	Target Name	Target Alias	Status
Infra-SVM	iqn.1992-08.com.netapp:sn.b5acab9ef1c811e68d9d00a098a9fec2:vs.3	Infra-SVM	up

24. Enter the IP address of `iscsi_lif_02b` for the IPv4 Address field.

Create iSCSI Static Target

iSCSI Target Name	:	iqn.1992-08.com.netapp::
Priority	:	1
Port	:	3260
Authentication Profile	:	<not set> ▼ Create iSCSI Authentication Profile
IPv4 Address	:	192.168.20.62
LUN ID	:	0

[OK](#) [Cancel](#)

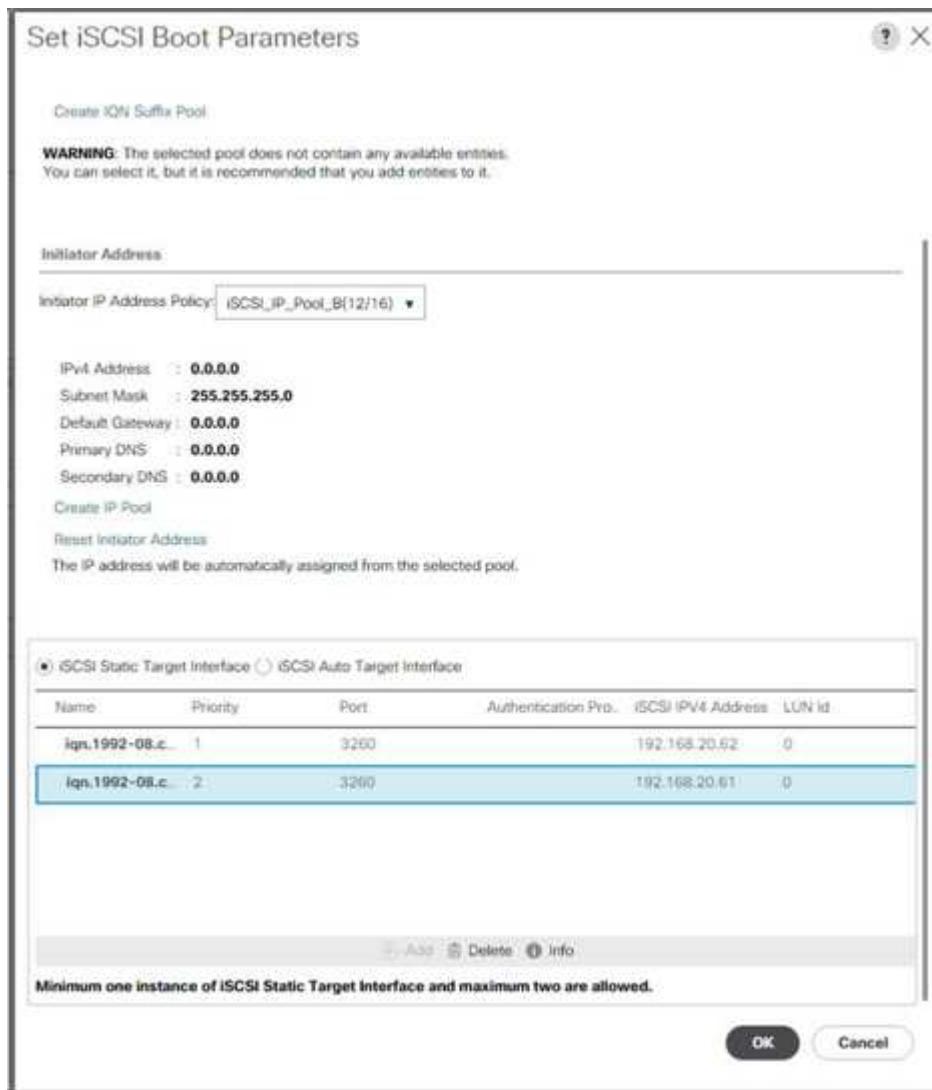
25. Click OK to add the iSCSI static target.
 26. Click Add.
 27. Enter the iSCSI target name.
 28. Enter the IP address of `iscsi_lif_01b` for the IPv4 Address field.

Create iSCSI Static Target

iSCSI Target Name	:	iqn.1992-08.com.netapp::
Priority	:	2
Port	:	3260
Authentication Profile	:	<not set> ▼ Create iSCSI Authentication Profile
IPv4 Address	:	192.168.20.61
LUN ID	:	0

[OK](#) [Cancel](#)

29. Click OK to add the iSCSI static target.

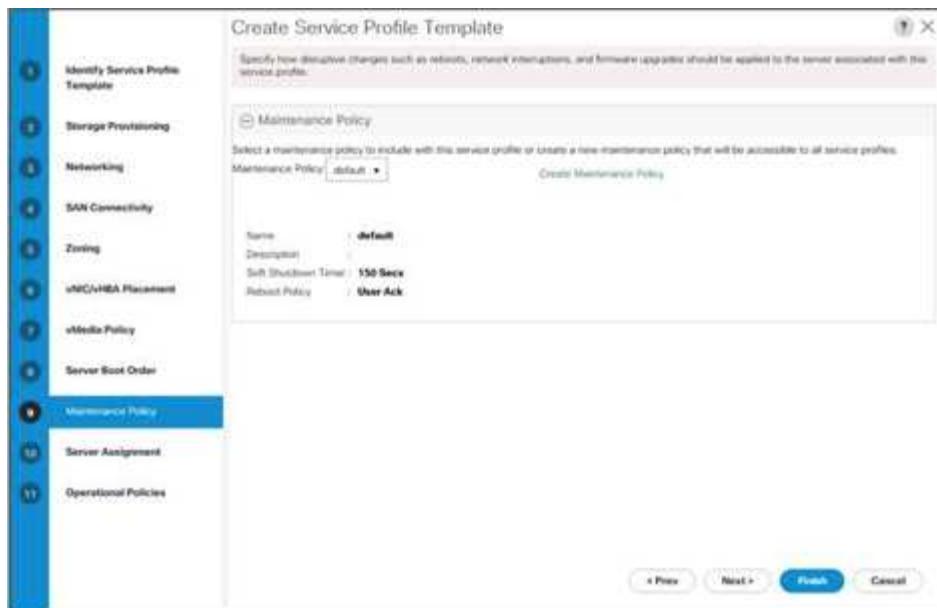


30. Click Next.

Configure maintenance policy

To configure the maintenance policy, complete the following steps:

1. Change the maintenance policy to default.

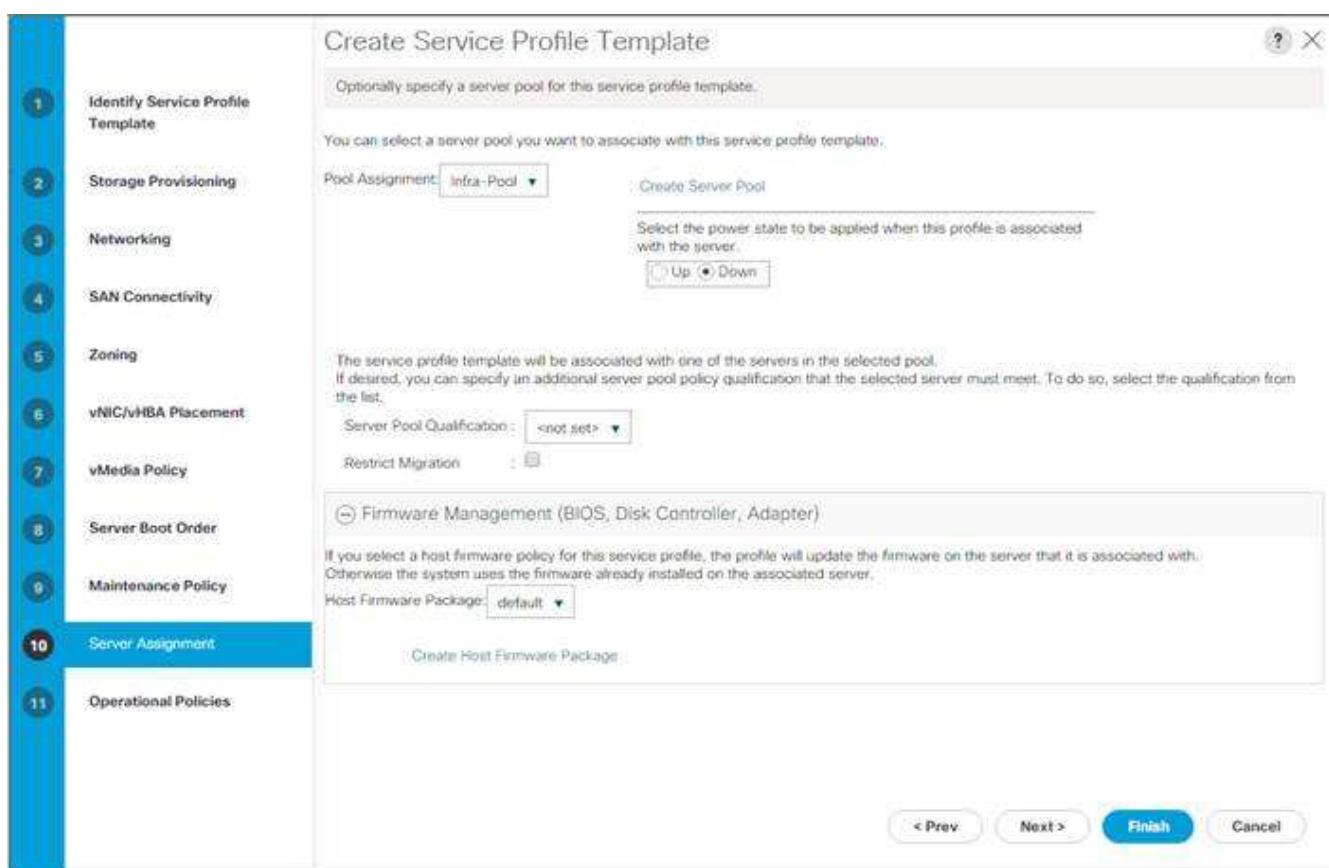


2. Click Next.

Configure server assignment

To configure the server assignment, complete the following steps:

1. In the Pool Assignment list, select Infra-Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Expand Firmware Management at the bottom of the page and select the default policy.

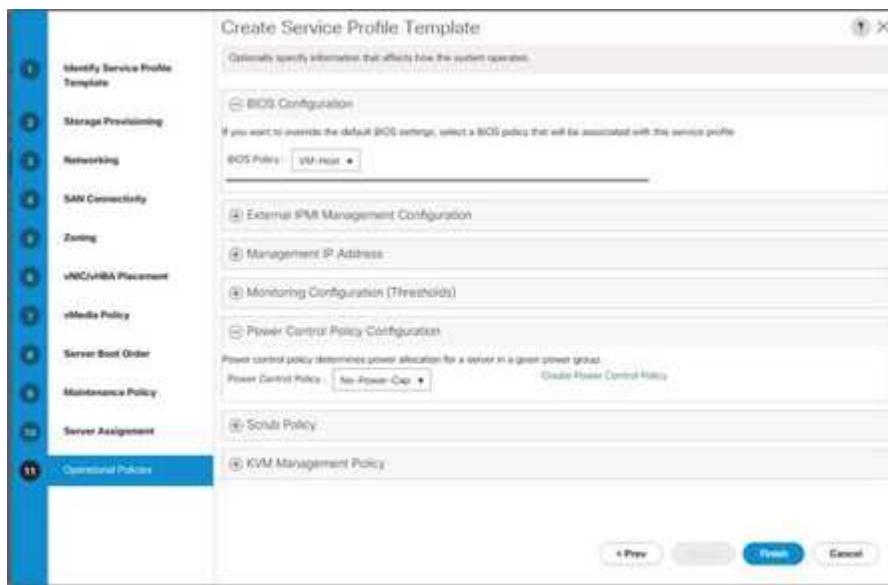


4. Click Next.

Configure operational policies

To configure the operational policies, complete the following steps:

1. From the BIOS Policy drop-down list, select VM-Host.
2. Expand Power Control Policy Configuration and select No-Power-Cap from the Power Control Policy drop-down list.



3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create vMedia-enabled service profile template

To create a service profile template with vMedia enabled, complete the following steps:

1. Connect to UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and select Create a Clone.
4. Name the clone VM-Host-Infra-iSCSI-A-vM.
5. Select the newly created VM-Host-Infra-iSCSI-A-vM and select the vMedia Policy tab on the right.
6. Click Modify vMedia Policy.
7. Select the ESXi-6. 7U1-HTTP vMedia Policy and click OK.
8. Click OK to confirm.

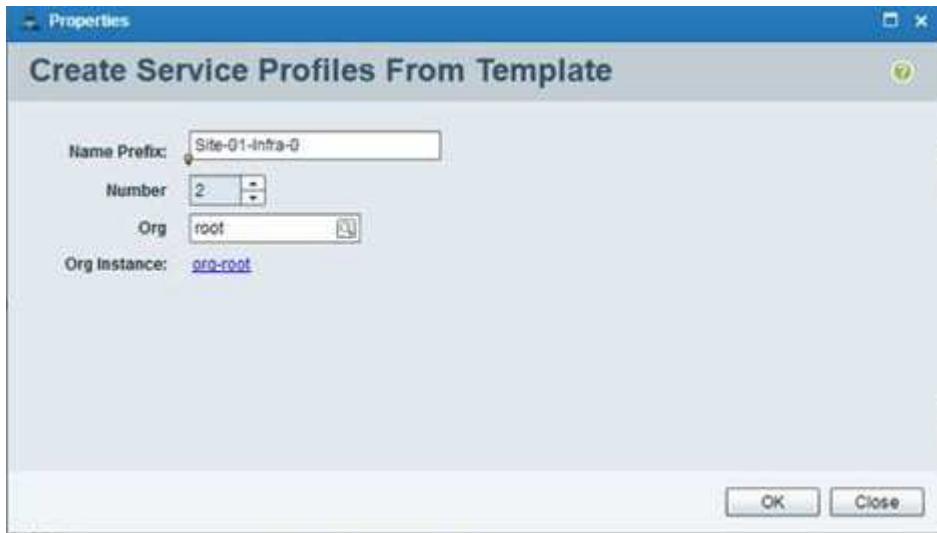
Create service profiles

To create service profiles from the service profile template, complete the following steps:

1. Connect to Cisco UCS Manager and click Servers on the left.
2. Expand Servers > Service Profile Templates > root > Service Template <name>.

3. In Actions, click Create Service Profile from Template and complete the following steps:

- Enter Site- 01-Infra-0 as the naming prefix.
- Enter 2 as the number of instances to create.
- Select root as the org.
- Click OK to create the service profiles.



4. Click OK in the confirmation message.

5. Verify that the service profiles Site-01-Infra-01 and Site-01-Infra-02 have been created.



The service profiles are automatically associated with the servers in their assigned server pools.

Storage configuration part 2: boot LUNs and initiator groups

ONTAP boot storage setup

Create initiator groups

To create initiator groups (igroups), complete the following steps:

1. Run the following commands from the cluster management node SSH connection:

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol
iscsi -ostype vmware -initiator <vm-host-infra-01-iqn>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol
iscsi -ostype vmware -initiator <vm-host-infra-02-iqn>
igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol iscsi
-ostype vmware -initiator <vm-host-infra-01-iqn>, <vm-host-infra-02-iqn>
```



Use the values listed in Table 1 and Table 2 for the IQN information.

2. To view the three igroups just created, run the `igroup show` command.

Map boot LUNs to igroups

To map boot LUNs to igroups, complete the following step:

1. From the storage cluster management SSH connection, run the following commands:

```
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- A
-igroup VM-Host-Infra-01 -lun-id 0
lun map -vserver Infra-SVM -volume
esxi_boot -lun VM-Host-Infra- B -igroup VM-Host-Infra-02 -lun-id 0
```

VMware vSphere 6.7U1 deployment procedure

This section provides detailed procedures for installing VMware ESXi 6.7U1 in a FlexPod Express configuration. After the procedures are completed, two booted ESXi hosts are provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in KVM console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot LUNs.

Download Cisco custom image for ESXi 6.7U1

If the VMware ESXi custom image has not been downloaded, complete the following steps to complete the download:

1. Click the following xref:./express/ [VMware vSphere Hypervisor \(ESXi\) 6.7U1](#).
2. You need a user ID and password on [vmware.com](#) to download this software.
3. Download the .iso file.

Cisco UCS Manager

The Cisco UCS IP KVM enables the administrator to begin the installation of the OS through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click Servers on the left.
7. Select Servers > Service Profiles > root > VM-Host-Infra-01.
8. Right-click VM-Host-Infra-01 and select KVM Console.

9. Follow the prompts to launch the Java-based KVM console.
10. Select Servers > Service Profiles > root > VM-Host-Infra-02.
11. Right-click VM-Host-Infra-02. and select KVM Console.
12. Follow the prompts to launch the Java-based KVM console.

Set up VMware ESXi installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware ESXi to the iSCSI-bootable LUN of the hosts, complete the following steps on each host:

1. Boot the server by selecting Boot Server and clicking OK. Then click OK again.
2. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
5. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. After the installation is complete, select the Virtual Media tab and clear the P mark next to the ESXi installation media. Click Yes.



The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

10. After the installation is complete, press Enter to reboot the server.
11. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

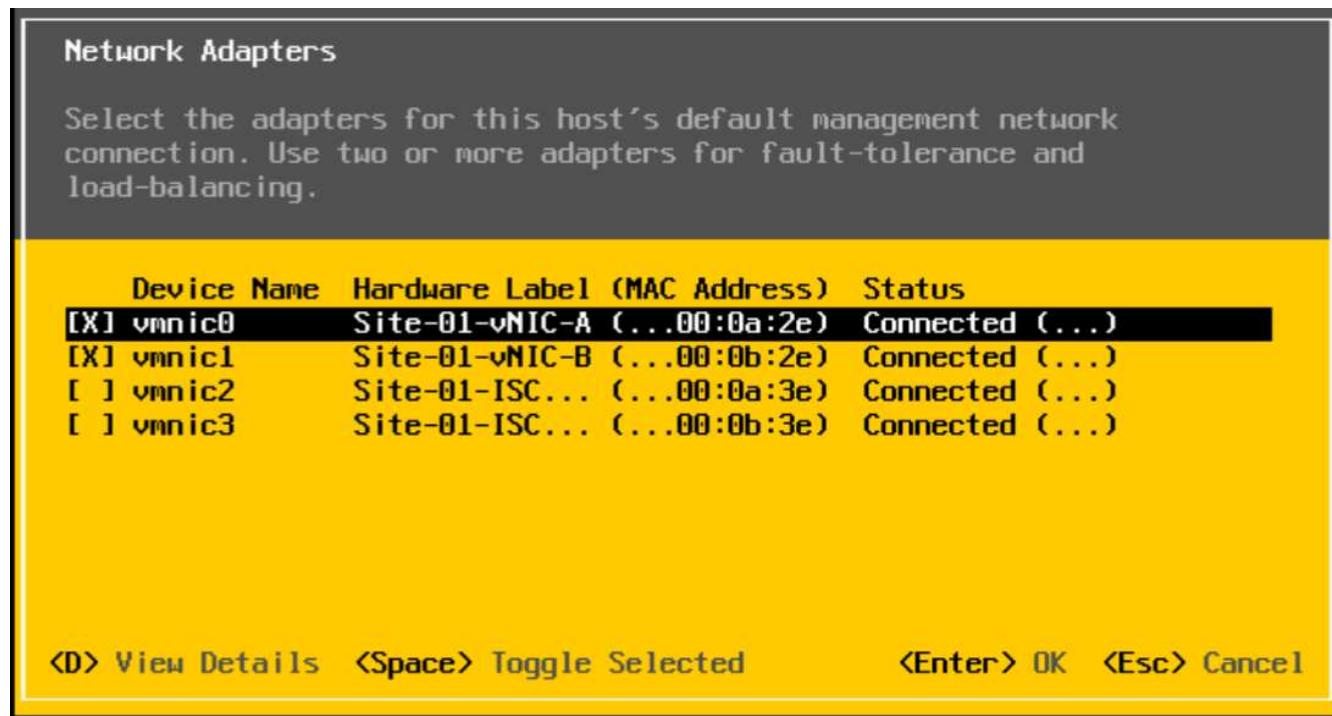
Set up management networking for ESXi hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host:

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To configure each ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select Troubleshooting Options and press Enter.
4. Select Enable ESXi Shell and press Enter.
5. Select Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Select the Configure Management Network option and press Enter.
8. Select Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field.
10. Press Enter.



11. Select the VLAN (Optional) option and press Enter.
12. Enter the `<ib-mgmt-vlan-id>` and press Enter.
13. Select IPv4 Configuration and press Enter.
14. Select the Set Static IPv4 Address and Network Configuration option by using the space bar.
15. Enter the IP address for managing the first ESXi host.
16. Enter the subnet mask for the first ESXi host.

17. Enter the default gateway for the first ESXi host.
18. Press Enter to accept the changes to the IP configuration.
19. Select the DNS Configuration option and press Enter.



Because the IP address is assigned manually, the DNS information must also be entered manually.

20. Enter the IP address of the primary DNS server.
21. Optional: Enter the IP address of the secondary DNS server.
22. Enter the FQDN for the first ESXi host.
23. Press Enter to accept the changes to the DNS configuration.
24. Press Esc to exit the Configure Management Network menu.
25. Select Test Management Network to verify that the management network is set up correctly and press Enter.
26. Press Enter to run the test, press Enter again once the test has completed, review environment if there is a failure.
27. Select the Configure Management Network again and press Enter.
28. Select the IPv6 Configuration option and press Enter.
29. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
30. Press Esc to exit the Configure Management Network submenu.
31. Press Y to confirm the changes and reboot the ESXi host.

Reset VMware ESXi host VMkernel port vmk0 MAC address (optional)

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port on which it is placed. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will occur because vmk0 retains the assigned MAC address unless the ESXi system configuration is reset. To reset the MAC address of vmk0 to a random VMware-assigned MAC address, complete the following steps:

1. From the ESXi console menu main screen, press Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of static macros.
2. Log in as root.
3. Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk0. vmk0 should be a part of the Management Network port group. Note the IP address and netmask of vmk0.
4. To remove vmk0, enter the following command:

```
esxcfg-vmknic -d "Management Network"
```

5. To add vmk0 again with a random MAC address, enter the following command:

```
esxcfg-vmknic -a -i <vmk0-ip> -n <vmk0-netmask> "Management Network"".
```

6. Verify that vmk0 has been added again with a random MAC address

```
esxcfg-vmknic -l
```

7. Type `exit` to log out of the command line interface.
8. Press `Ctrl-Alt-F2` to return to the ESXi console menu interface.

Log into VMware ESXi hosts with VMware host client

ESXi Host VM-Host-Infra-01

To log in to the VM-Host-Infra-01 ESXi host by using the VMware Host Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.
2. Click Open the VMware Host Client.
3. Enter `root` for the user name.
4. Enter the root password.
5. Click Login to connect.
6. Repeat this process to log in to VM-Host-Infra-02 in a separate browser tab or window.

Install VMware drivers for the Cisco Virtual Interface Card (VIC)

Download and extract the offline bundle for the following VMware VIC driver to the Management workstation:

- `nenic` Driver version 1.0.25.0

ESXi hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each Host Client, select Storage.
2. Right-click `datastore1` and select `Browse`.
3. In the Datastore browser, click `Upload`.
4. Navigate to the saved location for the downloaded VIC drivers and select `VMW-ESX-6.7.0-nenic-1.0.25.0-offline_bundle-11271332.zip`.
5. In the Datastore browser, click `Upload`.
6. Click `Open` to upload the file to `datastore1`.
7. Make sure the file has been uploaded to both ESXi hosts.
8. Place each host into Maintenance mode if it isn't already.
9. Connect to each ESXi host through ssh from a shell connection or putty terminal.

10. Log in as root with the root password.
11. Run the following commands on each host:

```
esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.7.0-  
nenic-1.0.25.0-offline_bundle-11271332.zip  
reboot
```

12. Log into the Host Client on each host once reboot is complete and exit Maintenance Mode.

Set up VMkernel ports and virtual switch

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To set up the VMkernel ports and the virtual switches on the ESXi hosts, complete the following steps:

1. From the Host Client, select Networking on the left.
2. In the center pane, select the Virtual switches tab.
3. Select vSwitch0.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, select vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.
9. Click Save.
10. Select Networking on the left.
11. In the center pane, select the Virtual switches tab.
12. Select iScsiBootvSwitch.
13. Select Edit settings.
14. Change the MTU to 9000
15. Click Save.
16. Select the VMkernel NICs tab.
17. Select vmk1 iScsiBootPG.
18. Select Edit settings.
19. Change the MTU to 9000.
20. Expand IPv4 settings and change the IP address to an address outside of the UCS iSCSI-IP-Pool-A.



To avoid IP address conflicts if the Cisco UCS iSCSI IP Pool addresses should get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

21. Click Save.
22. Select the Virtual switches tab.

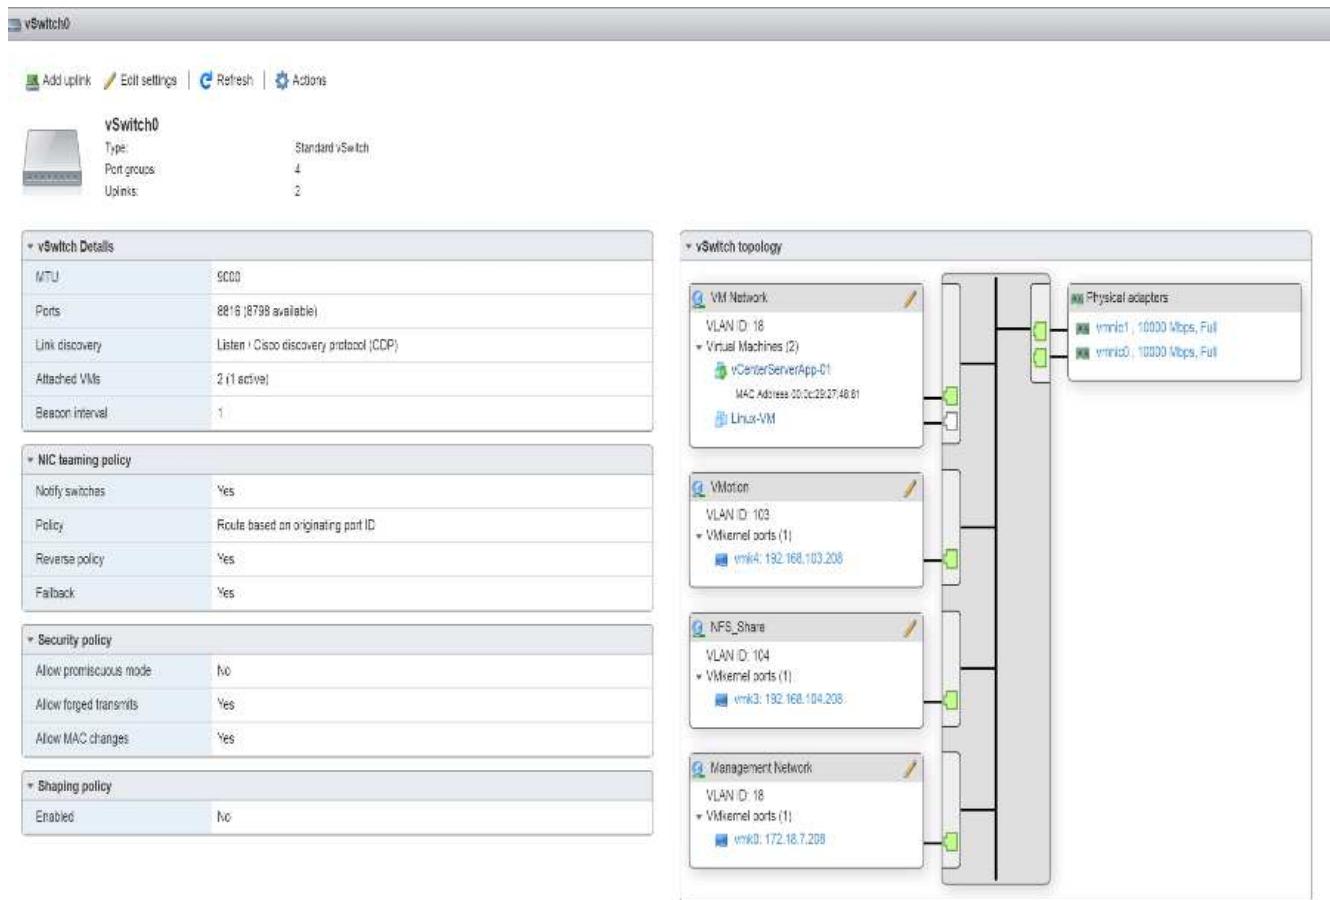
23. Select the Add standard virtual switch.
24. Provide a name of `iScscliBootvSwitch-B` for the vSwitch Name.
25. Set the MTU to 9000.
26. Select `vmnic3` from the Uplink 1 drop-down menu.
27. Click Add.
28. In the center pane, select the VMkernel NICs tab.
29. Select Add VMkernel NIC
30. Specify a New port group name of `iScsliBootPG-B`.
31. Select `iScscliBootvSwitch-B` for Virtual switch.
32. Set the MTU to 9000. Do not enter a VLAN ID.
33. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask within the Configuration.



To avoid IP address conflicts, if the Cisco UCS iSCSI IP Pool addresses should get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

34. Click Create.
35. On the left, select Networking, then select the Port groups tab.
36. In the center pane, right-click VM Network and select Remove.
37. Click Remove to complete removing the port group.
38. In the center pane, select Add port group.
39. Name the port group Management Network and enter `<ib-mgmt-vlan-id>` in the VLAN ID field, and make sure Virtual switch `vSwitch0` is selected.
40. Click Add to finalize the edits for the IB-MGMT Network.
41. At the top, select the VMkernel NICs tab.
42. Click Add VMkernel NIC.
43. For New port group, enter VMotion.
44. For Virtual switch, select `vSwitch0` selected.
45. Enter `<vmotion-vlan-id>` for the VLAN ID.
46. Change the MTU to 9000.
47. Select Static IPv4 settings and expand IPv4 settings.
48. Enter the ESXi host VMotion IP address and netmask.
49. Select the VMotion stack TCP/IP stack.
50. Select VMotion under Services.
51. Click Create.
52. Click Add VMkernel NIC.
53. For New port group, enter `NFS_Share`.
54. For Virtual switch, select `vSwitch0` selected.

55. Enter <infra-nfs-vlan-id> for the VLAN ID
56. Change the MTU to 9000.
57. Select Static IPv4 settings and expand IPv4 settings.
58. Enter the ESXi host Infrastructure NFS IP address and netmask.
59. Do not select any of the Services.
60. Click Create.
61. Select the Virtual Switches tab, then select vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:



62. Select the VMkernel NICs tab to confirm the configured virtual adapters. The adapters listed should be similar to the following example:

localhost.localdomain - Networking						
Port groups		Virtual switches		Physical NICs		VMkernel NICs
						TCP/IP stacks
						Firewall rules
Add VMkernel NIC		Edit settings		Refresh	Actions	
Name	Portgroup	TCP/IP stack	Services	IPv4 ad...	IPv6 addresses	
vmk0	Management Network	Default TCP/IP stack	Management	172.18.7....	fe80::225:b5ff:fe00:a2e/64	
vmk1	iScsiBootPG	Default TCP/IP stack		192.168....	fe80::225:b5ff:fe00:a3e/64	
vmk2	iScsiBootPG-B	Default TCP/IP stack		192.168....	fe80::250:56ff:fe64:1248...	
vmk3	NFS_Share	Default TCP/IP stack		192.168....	fe80::250:56ff:fe65:29a4...	
vmk4	VMotion	Default TCP/IP stack	vMotion	192.168....	fe80::250:56ff:fe6c:2650...	
5 items						

Setup iSCSI multipathing

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

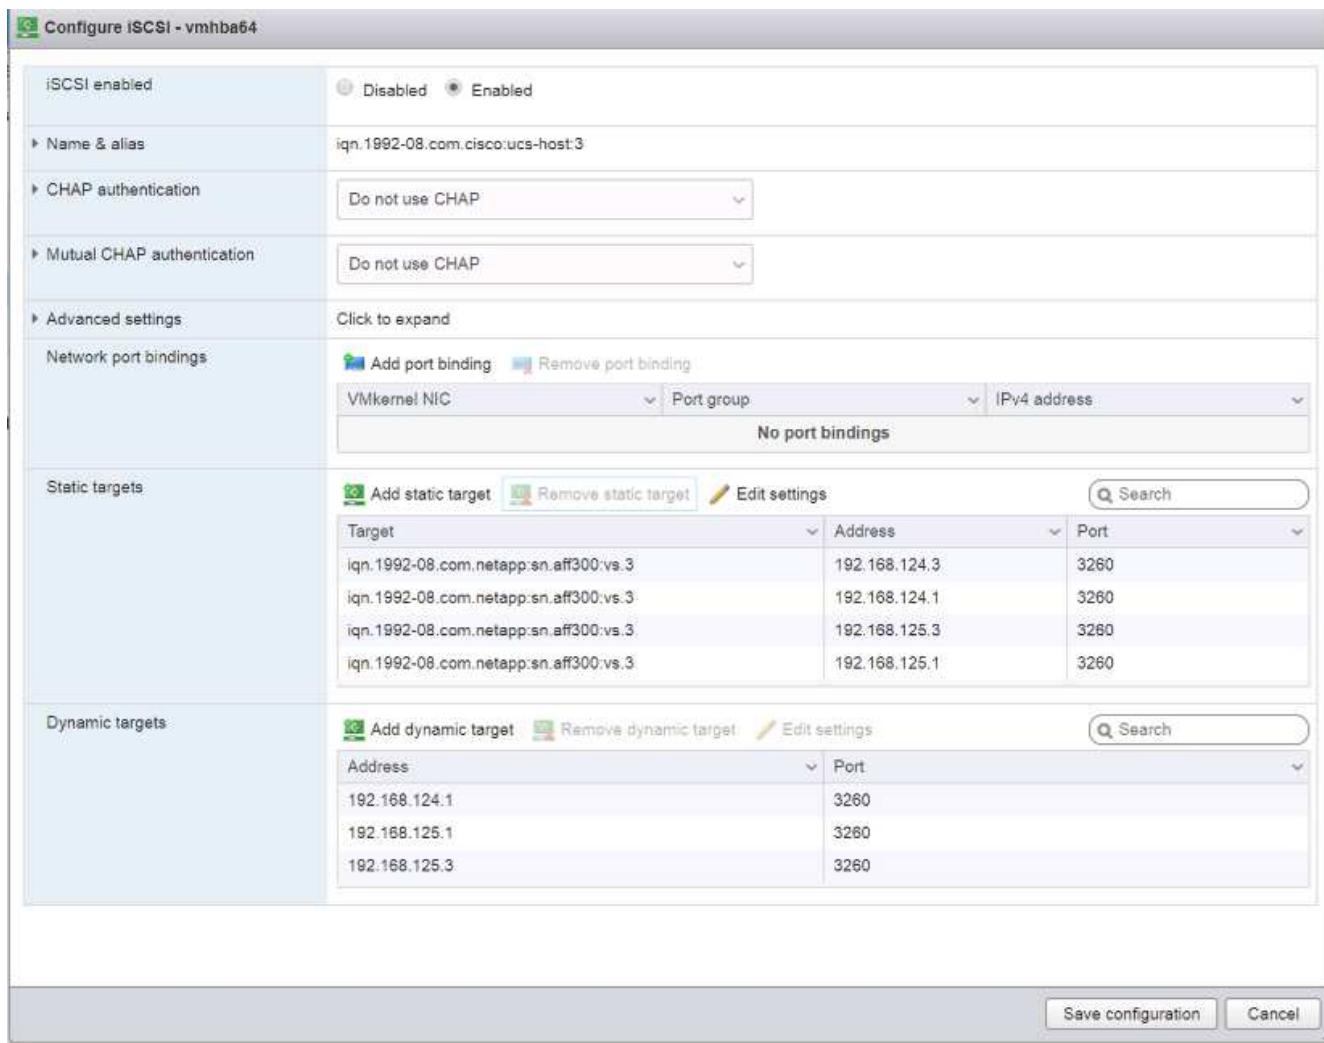
To set up the iSCSI multipathing on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each Host Client, select Storage on the left.
2. In the center pane, click Adapters.
3. Select the iSCSI software adapter and click Configure iSCSI.

localhost.localdomain - Storage						
Datastores		Adapters	Devices	Persistent Memory		
Configure iSCSI		Software iSCSI	Rescan	Refresh	Actions	
Name		Model		Status	Driver	
vmhba0		Lewisburg SATA AHCI Controller		Unknown	vmw_ahci	
vmhba64		iSCSI Software Adapter		Online	iscsi_vmk	
2 Items						



4. Under Dynamic targets, click Add dynamic target.
5. Enter the IP Address of `iscsi_lif01a`.
6. Repeat entering these IP addresses: `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
7. Click Save Configuration.



To obtain all of the `iscsi_lif` IP addresses, log in to NetApp storage cluster management interface and run the `network interface show` command.



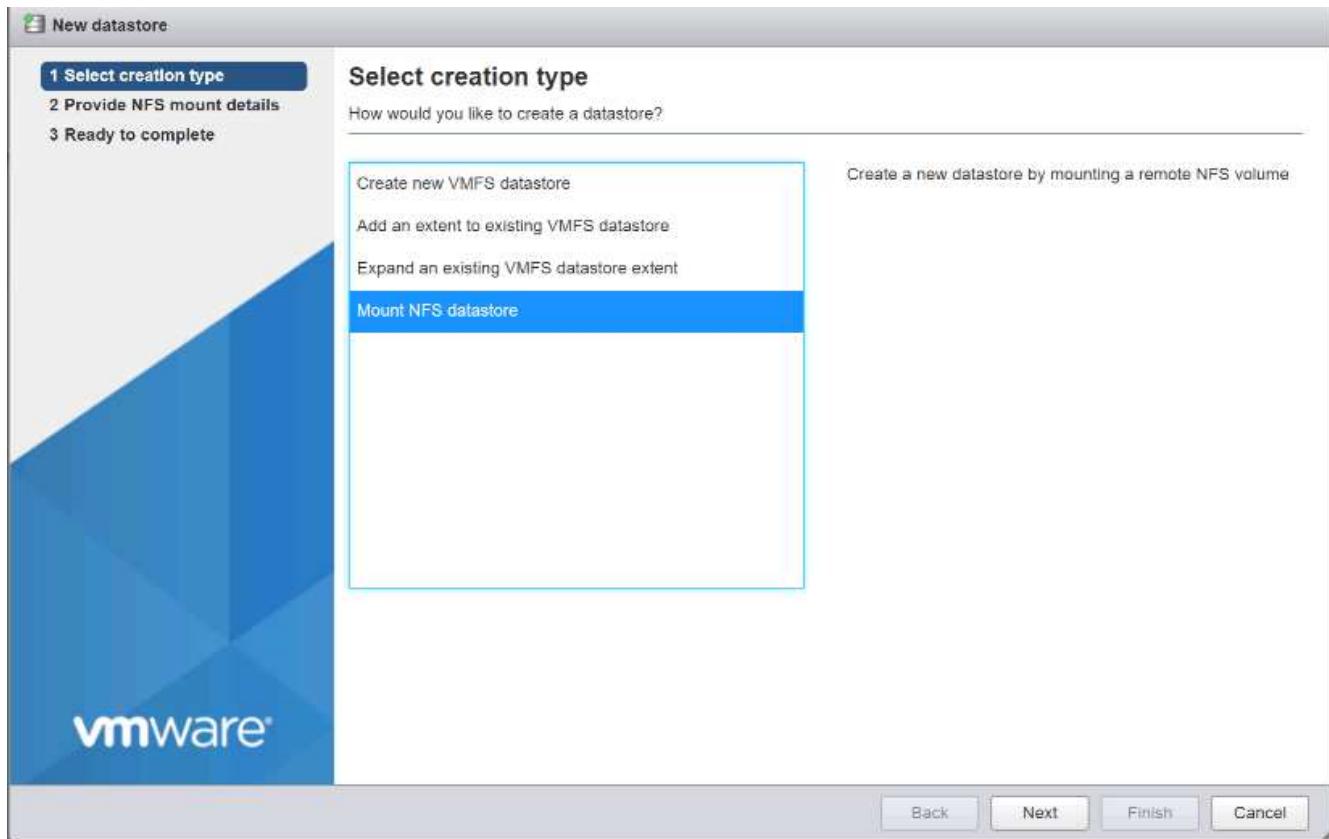
The host automatically rescans the storage adapter and the targets are added to static targets.

Mount required datastores

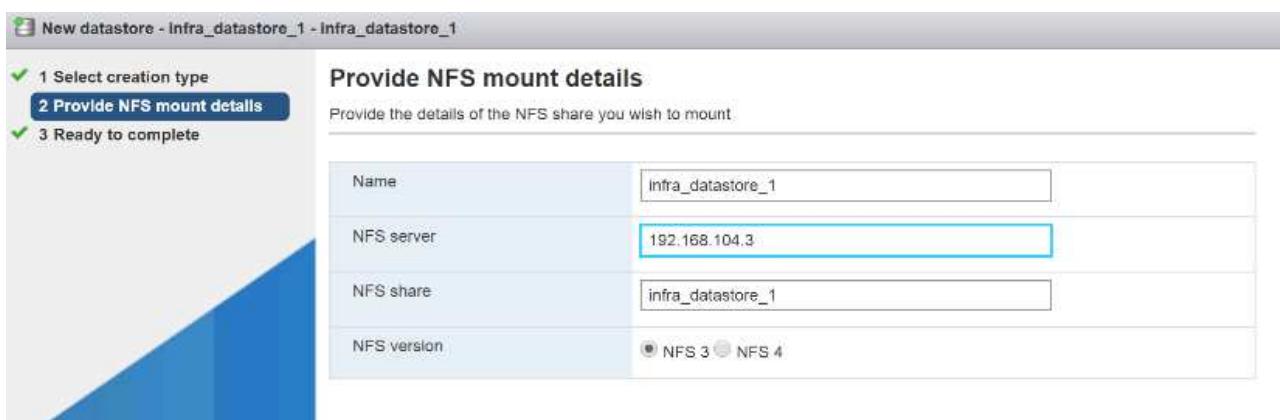
ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To mount the required datastores, complete the following steps on each ESXi host:

1. From the Host Client, select Storage on the left.
2. In the center pane, select Datastores.
3. In the center pane, select New Datastore to add a new datastore.
4. In the New datastore dialog box, select Mount NFS datastore and click Next.



5. On the provide NFS Mount Details page, complete these steps:
 - a. Enter `infra_datastore_1` for the datastore name.
 - b. Enter the IP address for the `nfs_lif01_a` LIF for the NFS server.
 - c. Enter `/infra_datastore_1` for the NFS share.
 - d. Leave the NFS version set at NFS 3.
 - e. Click Next.



6. Click Finish. The datastore should now appear in the datastore list.
7. In the center pane, select New Datastore to add a new datastore.
8. In the New Datastore dialog box, select Mount NFS Datastore and click Next.
9. On the provide NFS Mount Details page, complete these steps:

- a. Enter `infra_datastore_2` for the datastore name.
- b. Enter the IP address for the `nfs_lif02_a` LIF for the NFS server.
- c. Enter `/infra_datastore_2` for the NFS share.
- d. Leave the NFS version set at NFS 3.
- e. Click Next.

10. Click Finish. The datastore should now appear in the datastore list.

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provision	Access
datastore1	Non-SSD	7.5 GB	3.95 GB	3.55 GB	VMFS6	Supported	Single
infra_datastore_1	Unknown	500 GB	37.19 GB	462.81 GB	NFS	Supported	Single
infra_datastore_2	Unknown	500 GB	60.79 GB	439.21 GB	NFS	Supported	Single

11. Mount both datastores on both ESXi hosts.

Configure NTP on ESXi hosts

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To configure NTP on the ESXi hosts, complete the following steps on each host:

1. From the Host Client, select Manage on the left.
2. In the center pane, select the Time & Date tab.
3. Click Edit Settings.
4. Make sure Use Network Time Protocol (enable NTP client) is selected.
5. Use the drop-down menu to select Start and Stop with Host.
6. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.



7. Click Save to save the configuration changes.
8. Select Actions > NTP service > Start.
9. Verify that NTP service is now running and the clock is now set to approximately the correct time



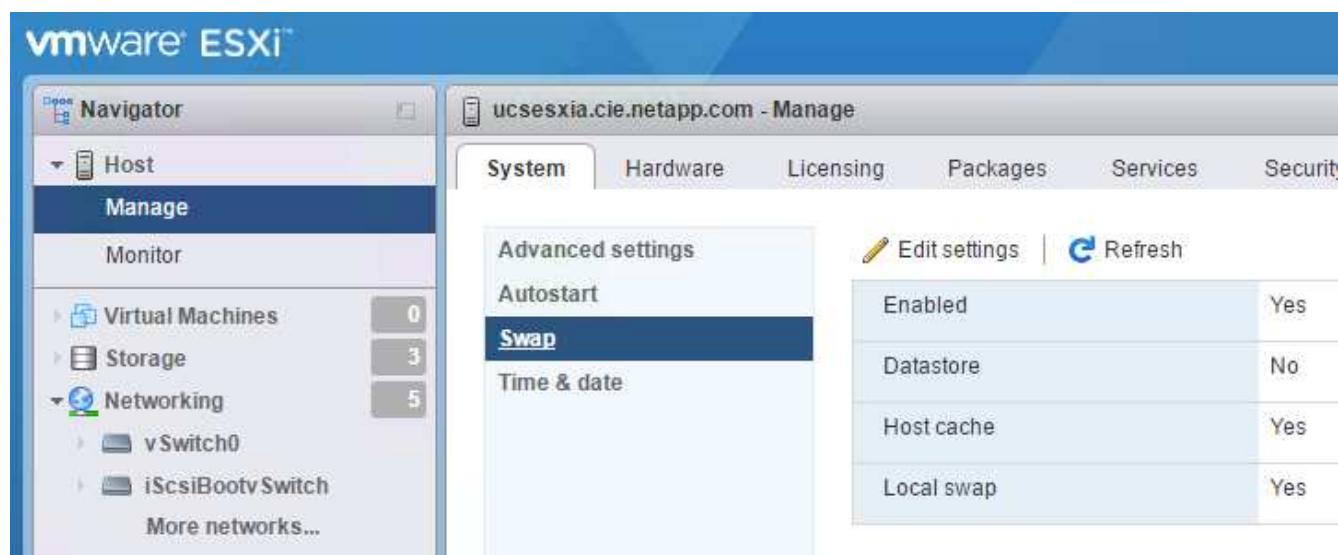
The NTP server time might vary slightly from the host time.

Configure ESXi host swap

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To configure host swap on the ESXi hosts, follow these steps on each host:

1. Click Manage in the left navigation pane. Select System in the right pane and click Swap.



2. Click Edit Settings. Select `infra_swap` from the Datastore options.



3. Click Save.

Install the NetApp NFS Plug-in 1.1.2 for VMware VAAI

To install the NetApp NFS Plug-in 1.1.2 for VMware VAAI, complete the following steps.

1. Download the NetApp NFS Plug-in for VMware VAAI:
 - a. Go to the [NetApp software download page](#).
 - b. Scroll down and click NetApp NFS Plug-in for VMware VAAI.
 - c. Select the ESXi platform.
 - d. Download either the offline bundle (.zip) or online bundle (.vib) of the most recent plug-in.
2. The NetApp NFS plug-in for VMware VAAI is pending IMT qualification with ONTAP 9.5 and interoperability details will be posted to the NetApp IMT soon.
3. Install the plug-in on the ESXi host by using the ESX CLI.
4. Reboot the ESXi host.

Install VMware vCenter Server 6.7

This section provides detailed procedures for installing VMware vCenter Server 6.7 in a FlexPod Express configuration.

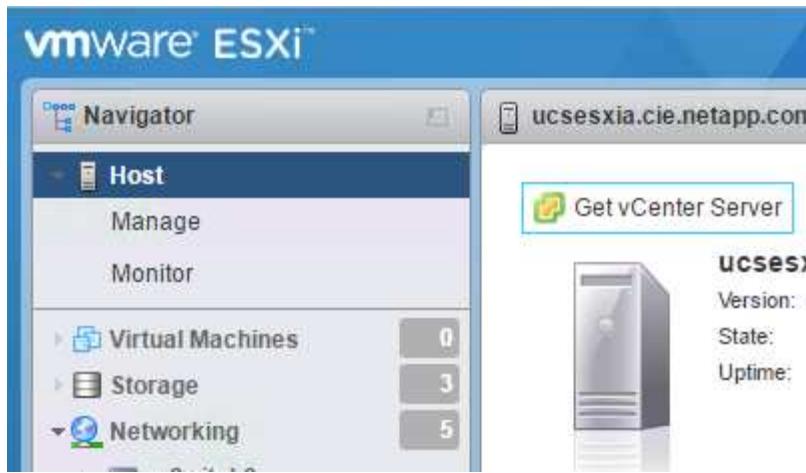


FlexPod Express uses the VMware vCenter Server Appliance (VCSA).

Install VMware vCenter server appliance

To install VCSA, complete the following steps:

1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.

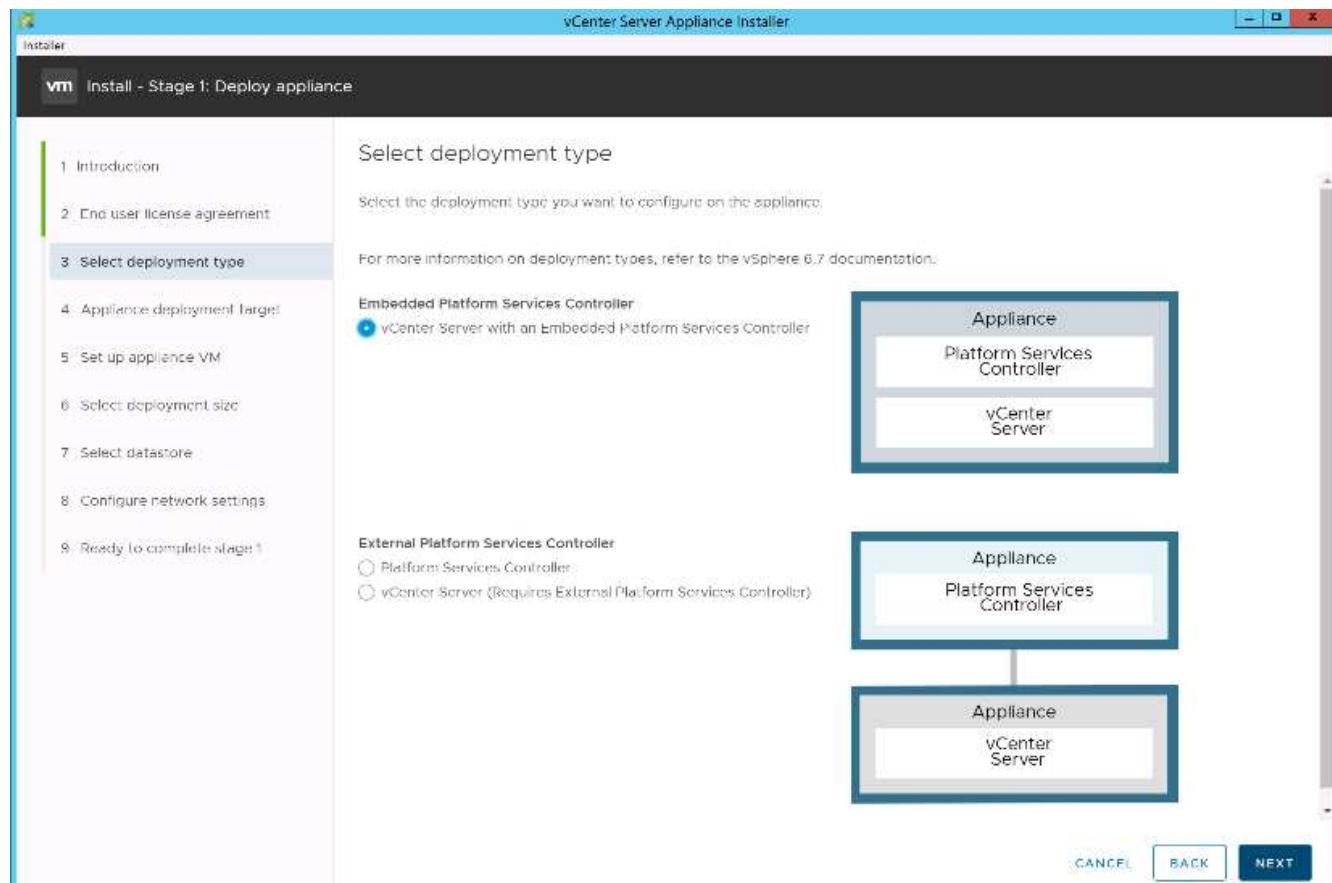


2. Download the VCSA from the VMware site.



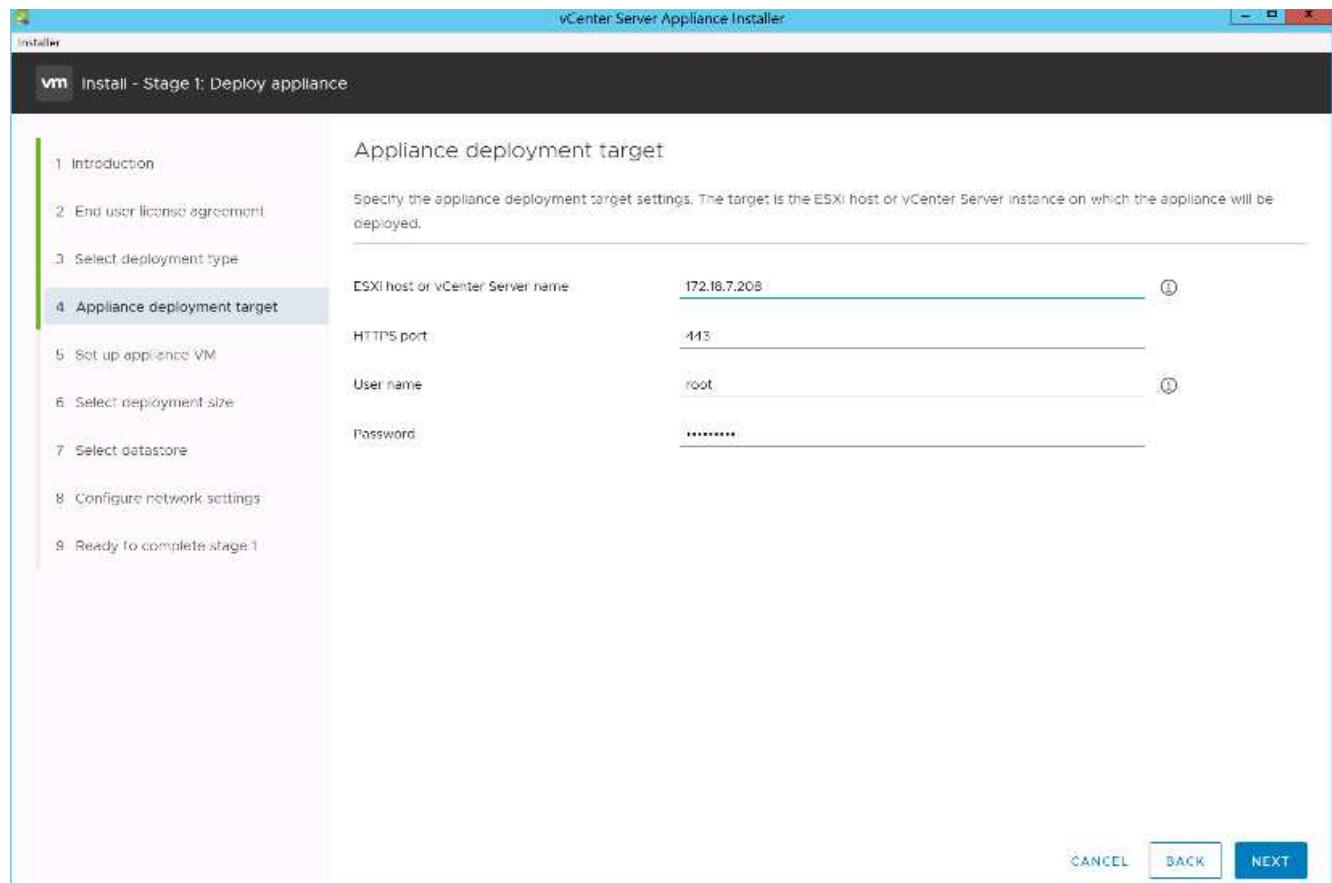
Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.

3. Mount the ISO image.
4. Navigate to the `vcsa-ui-installer > win32` directory. Double-click `installer.exe`.
5. Click Install.
6. Click Next on the Introduction page.
7. Accept the EULA.
8. Select Embedded Platform Services Controller as the deployment type.

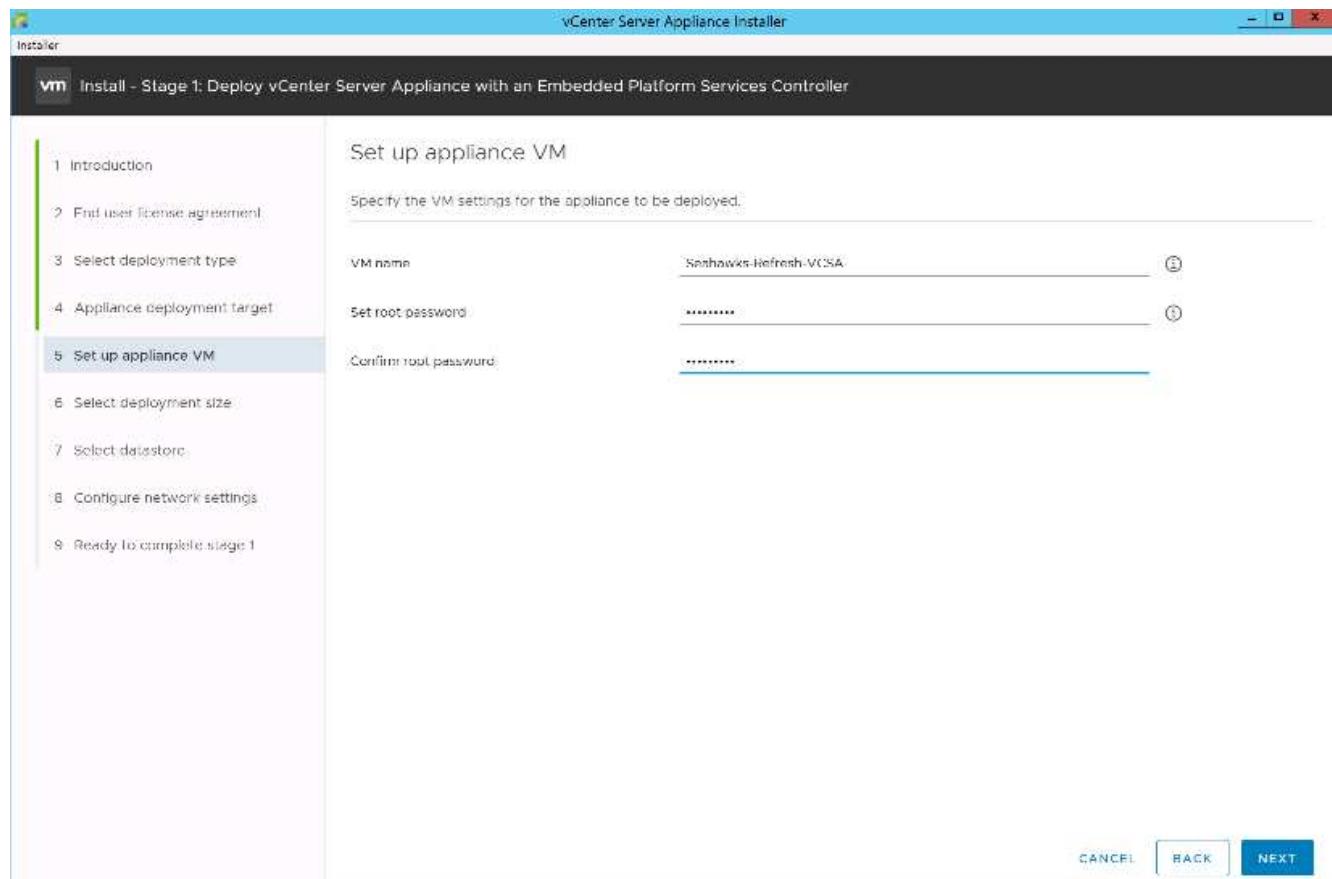


If required, the External Platform Services Controller deployment is also supported as part of the FlexPod Express solution.

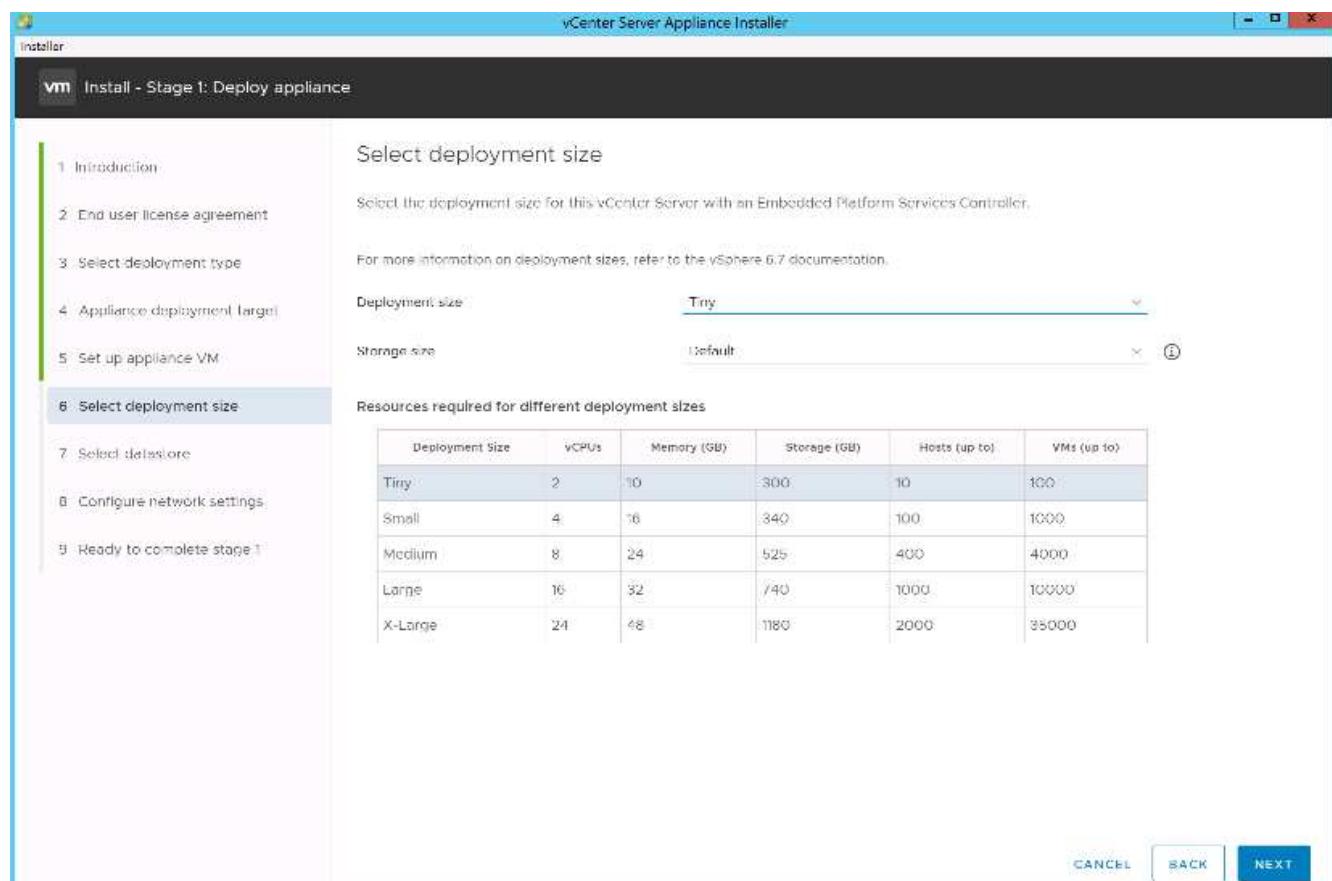
9. On the Appliance Deployment Target page, enter the IP address of an ESXi host you have deployed, the root user name, and the root password. Click Next.



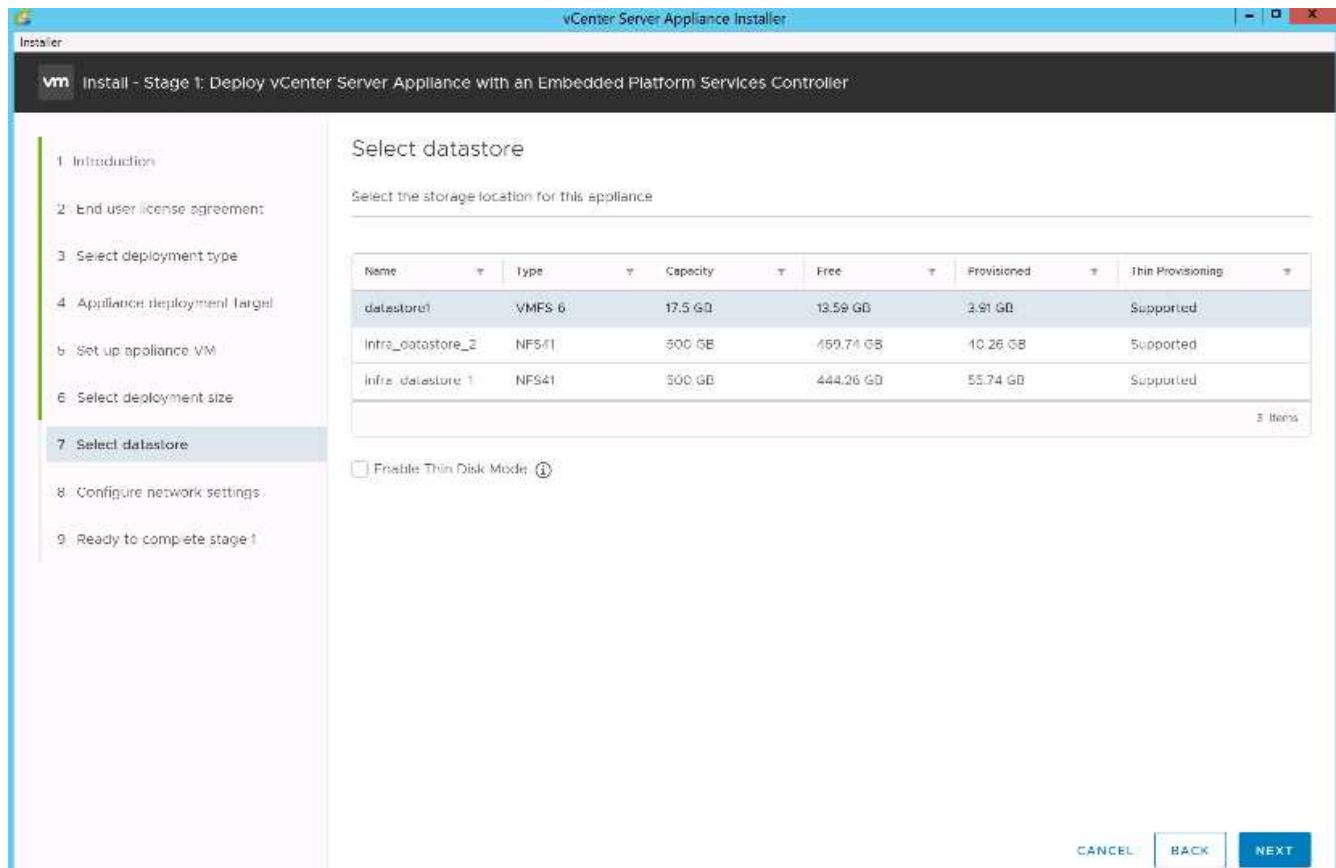
10. Set the appliance VM by entering VCSA as the VM name and the root password you would like to use for the VCSA. Click Next.



11. Select the deployment size that best fits your environment. Click Next.

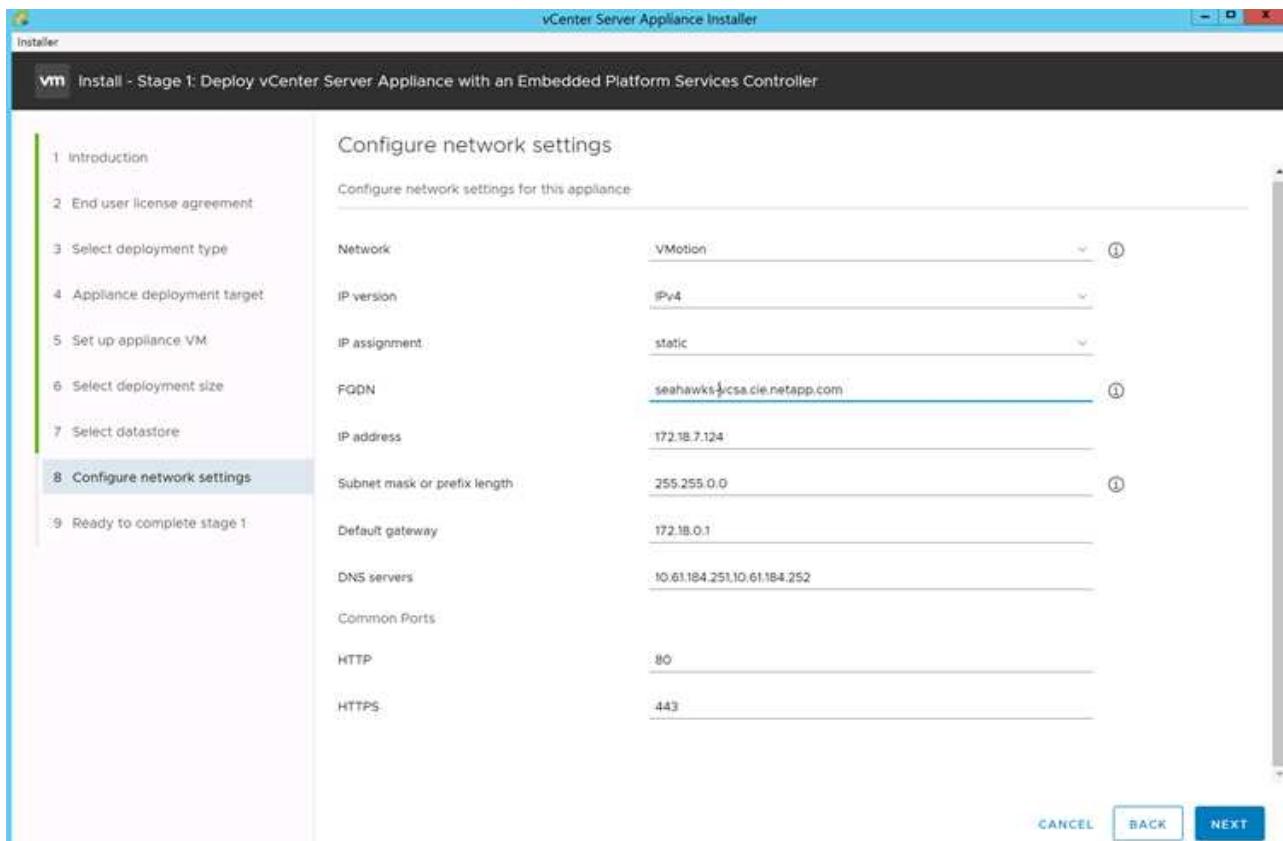


12. Select the `infra_datastore_1` datastore. Click Next.



13. Enter the following information on the Configure Network Settings page and click Next.

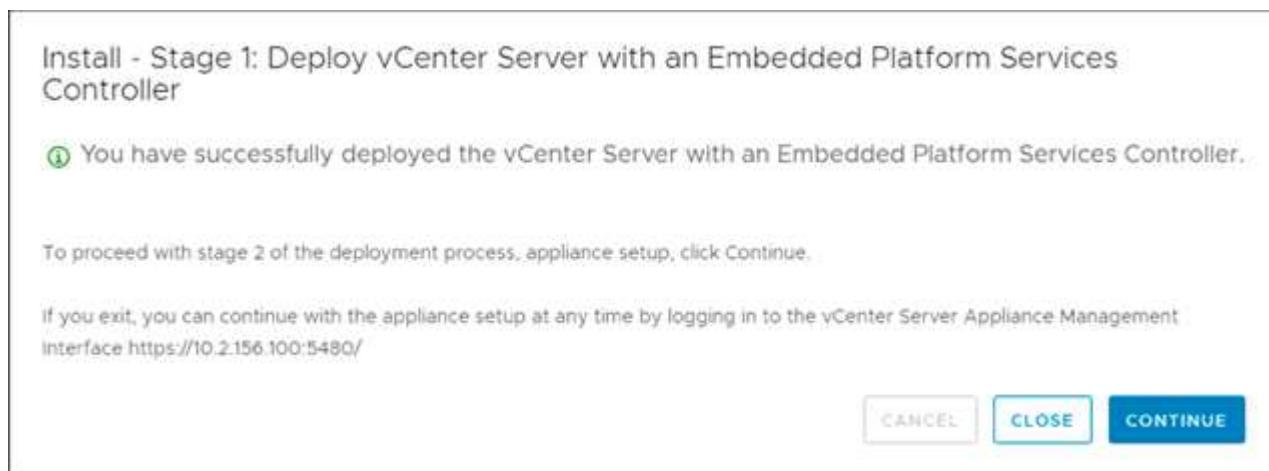
- Select MGMT-Network as your network.
- Enter the FQDN or IP to be used for the VCSA.
- Enter the IP address to be used.
- Enter the subnet mask to be used.
- Enter the default gateway.
- Enter the DNS server.



14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.

The VCSA installs now. This process takes several minutes.

15. After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.



16. On the Stage 2 Introduction page, click Next.
17. Enter <<var_ntp_id>> for the NTP server address. You can enter multiple NTP IP addresses.

If you plan to use vCenter Server high availability, make sure that SSH access is enabled.

18. Configure the SSO domain name, password, and site name. Click Next.

Record these values for your reference, especially if you deviate from the `vsphere.local` domain name.

19. Join the VMware Customer Experience Program if desired. Click Next.
20. View the summary of your settings. Click Finish or use the back button to edit settings.
21. A message appears stating that you are not able to pause or stop the installation from completing after it has started. Click OK to continue.

The appliance setup continues. This takes several minutes.

A message appears indicating that the setup was successful.



The links that the installer provides to access vCenter Server are clickable.

Configure VMware vCenter Server 6.7 and vSphere clustering

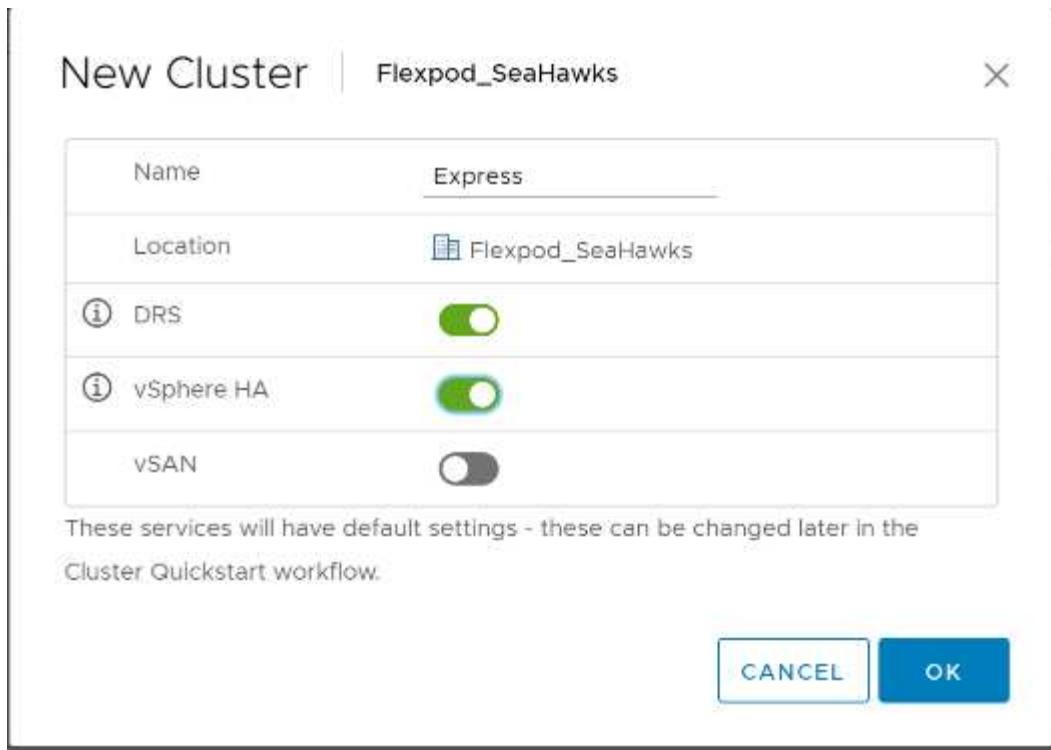
To configure VMware vCenter Server 6.7 and vSphere clustering, complete the following steps:

1. Navigate to `https://<<FQDN or IP of vCenter>>/vsphere-client/`.
2. Click Launch vSphere Client.
3. Log in with the user name `administrator@vsphere.local` and the SSO password you entered during the VCSA setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center and click OK.

Create vSphere Cluster.

To create a vSphere cluster, complete the following steps:

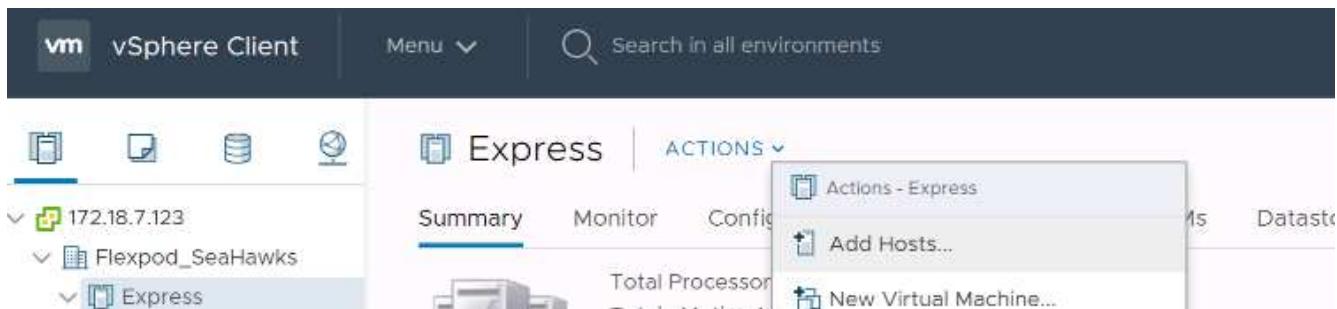
1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Select and enable DRS and vSphere HA options.
4. Click OK.



Add ESXi Hosts to Cluster

To add ESXi hosts to the cluster, complete the following steps:

1. Select Add Host in the Actions menu of the cluster.



2. To add an ESXi host to the cluster, complete the following steps:

- Enter the IP or FQDN of the host. Click Next.
- Enter the root user name and password. Click Next.
- Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
- Click Next on the Host Summary page.
- Click the green + icon to add a license to the vSphere host.



This step can be completed later if desired.

- Click Next to leave lockdown mode disabled.
- Click Next at the VM location page.
- Review the Ready to Complete page. Use the back button to make any changes or select Finish.

3. Repeat steps 1 and 2 for Cisco UCS host B.

This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on ESXi hosts

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. This procedure should be run at the end of the vCenter deployment section. To setup the ESXi Dump Collector, follow these steps:

1. Log in to the vSphere Web Client as administrator@vsphere.local and select Home.
2. In the center pane, click System Configuration.
3. In the left pane, select Services.
4. Under Services, click VMware vSphere ESXi Dump Collector.
5. In the center pane, click the green start icon to start the service.
6. In the Actions menu, click Edit Startup Type.
7. Select Automatic.
8. Click OK.
9. Connect to each ESXi host using ssh as root.
10. Run the following commands:

```
esxcli system coredump network set -v vmk0 -j <vcenter-ip>
esxcli system coredump network set -e true
esxcli system coredump network check
```

The message Verified the configured netdump server is running appears after you run the final command.



This process must be completed for any additional hosts added to FlexPod Express.

Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of additional components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed by keeping in mind small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NVA- 1130-DESIGN: FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based Storage NVA Design

<https://www.netapp.com/us/media/nva-1130-design.pdf>

- AFF and FAS Systems Documentation Center

<http://docs.netapp.com/platstor/index.jsp>

- ONTAP 9 Documentation Center

<http://docs.netapp.com/ontap-9/index.jsp>

- NetApp Product Documentation

<https://docs.netapp.com>

FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS - NVA - Deployment

Jyh-shing Chen, NetApp

The FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS solution leverages Cisco UCS Mini with B200 M5 blade servers, Cisco UCS 6324 in-chassis Fabric Interconnects, Cisco Nexus 31108PC-V switches, or other compliant switches, and NetApp AFF A220, C190, or the FAS2700 series controller HA pair, which runs NetApp ONTAP 9.7 data management software. This NetApp Verified Architecture (NVA) deployment document provides the detailed steps needed to configure the infrastructure components and to deploy VMware vSphere 7.0 and the associated tools to create a highly reliable and highly available FlexPod Express-based virtual infrastructure.

[FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS - NVA - Deployment](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.