NetApp Verified Architecture

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

NVA Deployment

Savita Kumari, NetApp

November 2019 | NVA-1142-DEPLOY | Version 1.0

TABLE OF CONTENTS

1 Program Summary.......... 4

2 Solution Overview 4

2.1 FlexPod Converged Infrastructure Program.......... 4

2.2 NetApp Verified Architecture Program 5

2.3 Solution Technology 5

2.4 Use Case Summary.......... 6

3 Technology Requirements 7

3.1 Hardware Requirements 7

3.2 Software Requirements 7

4 FlexPod Express Cabling Information.......... 8

5 Deployment Procedures 10

5.1 Cisco Nexus 31108PC-V Deployment Procedure.......... 11

5.2 NetApp Storage Deployment Procedure (Part 1) 18

5.3 Cisco UCS C-Series Rack Server Deployment Procedure 33

5.4 NetApp AFF Storage Deployment Procedure (Part 2) 44

5.5 VMware vSphere 6.7U2 Deployment Procedure 45

5.6 Install VMware vCenter Server 6.7U2.......... 57

5.7 Configure VMware vCenter Server 6.7U2 and vSphere Clustering 67

6 Conclusion 79

Acknowledgments 79

Where to Find Additional Information 79

Version History 80

LIST OF TABLES

Table 1) Hardware requirements for the base configuration.......... 7

Table 2) Hardware for scaling the solution by adding two hypervisor nodes. 7

Table 3) Software requirements for the base FlexPod Express implementation. 7

Table 4) Software requirements for a VMware vSphere implementation. 8

Table 5) Cabling information for Cisco Nexus switch 31108PC-V A. 8

Table 6) Cabling information for Cisco Nexus switch 31108PC-V B. 9

Table 7) Cabling information for NetApp AFF C190 storage controller A. 9

Table 8) Cabling information for NetApp AFF C190 storage controller B. 9

Table 9) Required VLANs.......... 10

Table 10) Solution vSwitches. 10
Table 11) VMware VMs created. 11
Table 12) ONTAP 9.6 installation and configuration information. 19
Table 13) Information required for NFS configuration. 29
Table 14) Information required for iSCSI configuration. 31
Table 15) Information required for NFS configuration. 32
Table 16) Information required for SVM administrator addition. 32
Table 17) Information required for CIMC configuration. 33
Table 18) Information required for iSCSI boot configuration. 36
Table 19) Information required for configuring ESXi hosts. 48

LIST OF FIGURES

Figure 1) FlexPod portfolio. 5
Figure 2) FlexPod Express with VMware vSphere 10GbE architecture. 6
Figure 3) Reference validation cabling. 8

1 Program Summary

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.

2 Solution Overview

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

2.1 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 Figure 1, 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.

Figure 1) FlexPod portfolio.

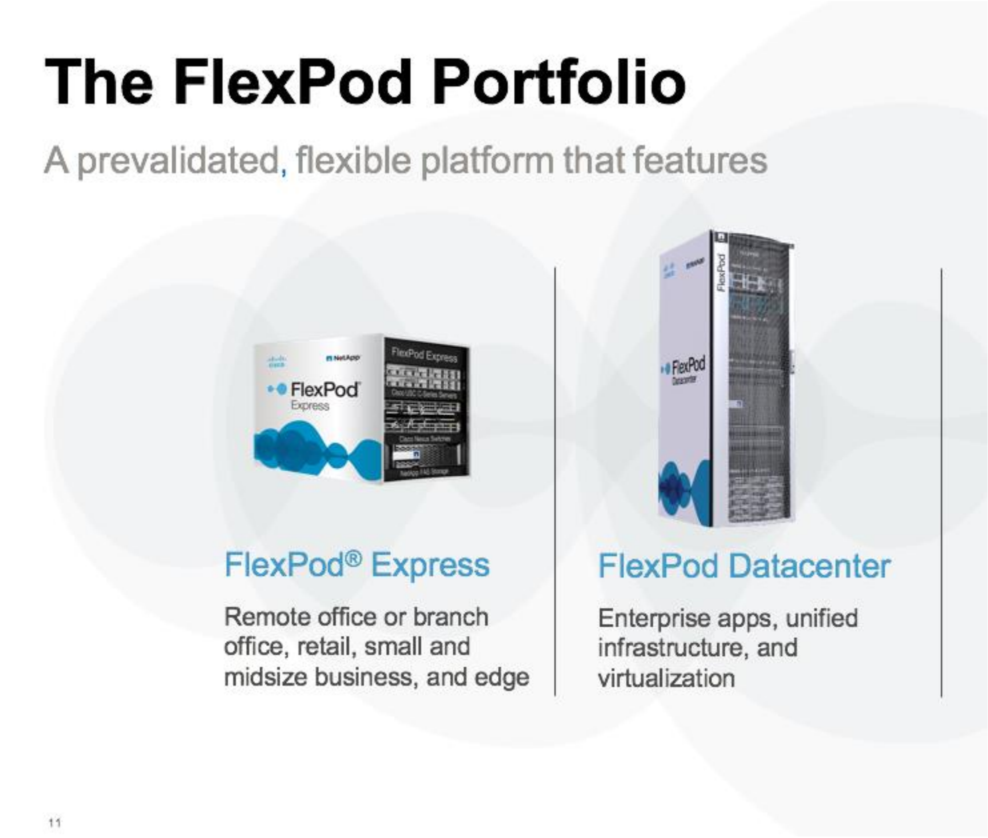

The FlexPod Portfolio

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

2.2 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.

2.3 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.

Figure 2) FlexPod Express with VMware vSphere 10GbE architecture.

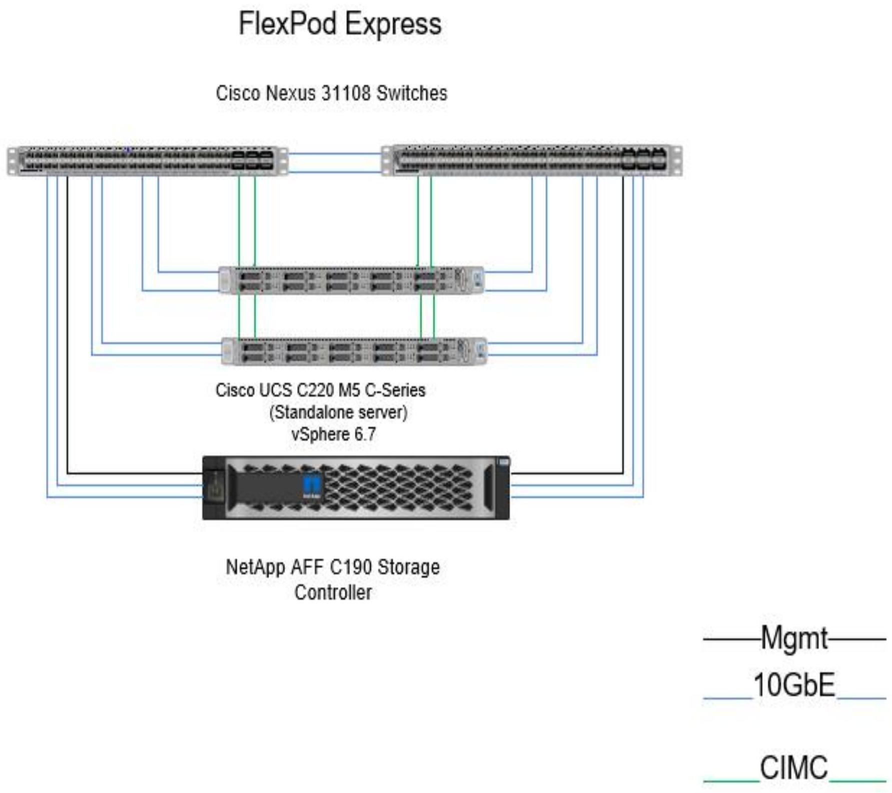

Note: To use the four physical 10GbE ports on the VIC 1457 efficiently, create two extra links from each server to the top rack switches.

2.4 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 functionally 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.

3 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.

3.1 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.

Table 1 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.

Table 1) Hardware requirements for the base configuration.

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

Table 2 lists the hardware that is required in addition to the base configuration for implementing 10GbE.

Table 2) Hardware for scaling the solution by adding two hypervisor nodes.

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

3.2 Software Requirements

Table 3 lists the software components that are required to implement the architectures of the FlexPod Express solutions.

Table 3) Software requirements for the base FlexPod Express implementation.

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

Table 4 lists the software that is required for all VMware vSphere implementations on FlexPod Express.

Table 4) Software requirements for a VMware vSphere implementation.

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

4 FlexPod Express Cabling Information

This reference validation is cabled as shown in Figure 3 and Table 5 through Table 8.

Figure 3) Reference validation cabling.

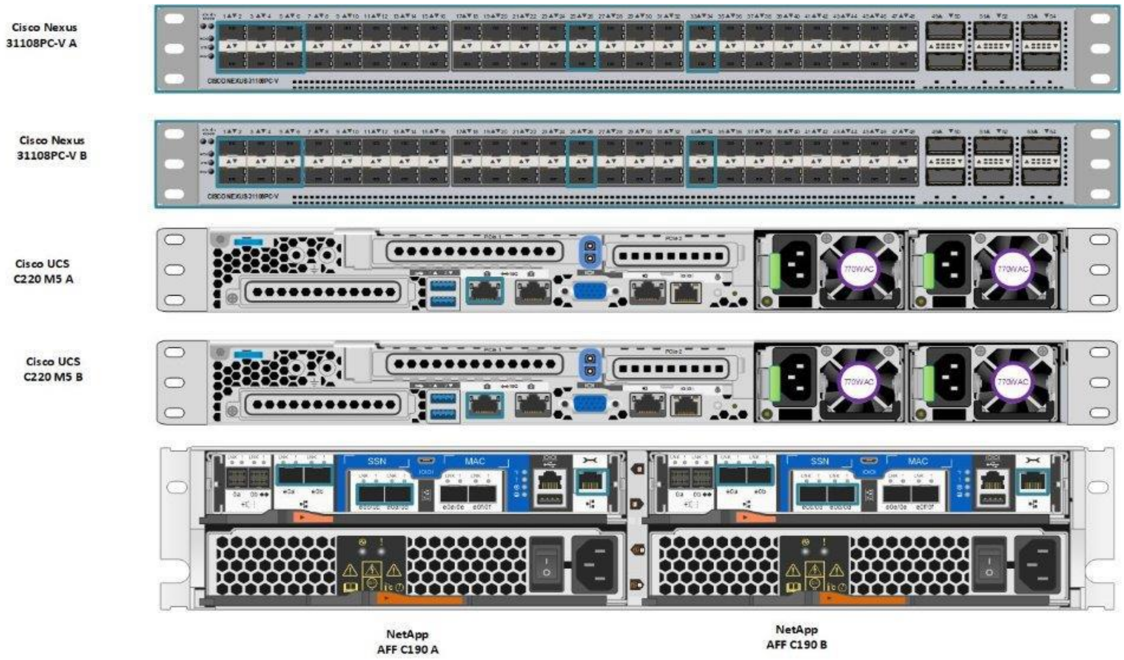

Table 5) Cabling information for Cisco Nexus switch 31108PC-V A.

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

Local Device	Local Port	Remote Device	Remote Port
	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)

Table 6) 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)

Table 7) 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

Table 8) 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

5 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> 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. Table 9 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.

Note: 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.

Table 9) Required VLANs.

VLAN Name	VLAN Purpose	VLAN ID	
Management VLAN	VLAN for management interfaces	3437	vSwitch0
NFS VLAN	VLAN for NFS traffic	3438	vSwitch0
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.

Table 10) 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.

Note: 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.

Table 11) VMware VMs created.

VM Description	Host Name
VMware vCenter Server	FlexPod-VCSA
Virtual Storage Console	FlexPod-VSC

5.1 Cisco Nexus 31108PC-V Deployment Procedure

The following 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.

Note: 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.

   Note: 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 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
```

Note: 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-l4port
```

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

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 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 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 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 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_id>>,<<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_id>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  no shut
exit
copy run start
```

Configure 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
```

Note: 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.

Note: 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 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.

5.2 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.

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

Table 12) 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.

   Note: 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.

   Note: 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.

    Note: 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.

    Note: 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.

   Note: 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.

   Note: 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.

    Note: 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.

    Note: 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 9.6 boots on the node for the first time.

Note: 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.

   Note: 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.

   Note: You can enter multiple DNS server IP addresses.

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

   Note: 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.

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

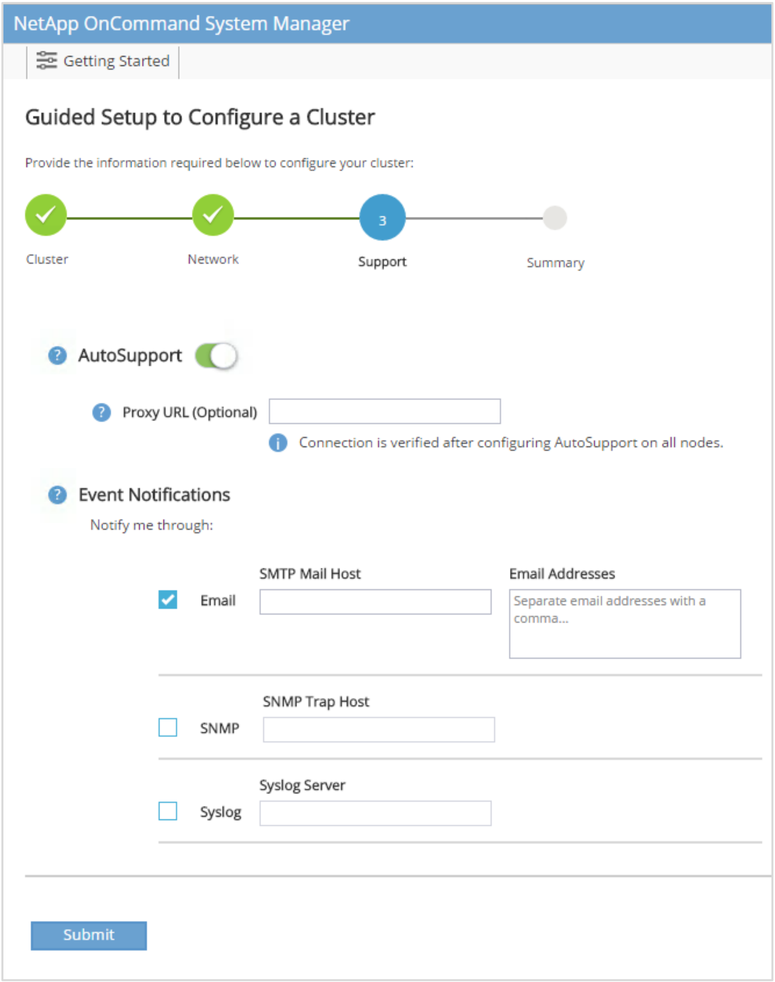

NetApp OnCommand System Manager

Getting Started

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

AutoSupport

Proxy URL (Optional)

Connection is verified after configuring AutoSupport on all nodes.

Event Notifications

Notify me through:

☑	Email	SMTP Mail Host	Email Addresses: Separate email addresses with a comma...
□	SNMP	SNMP Trap Host	
□	Syslog	Syslog Server	

When the system indicates 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 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
```

Note: 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
```

Note: 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 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
```

Setting 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>>
```

Note: 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
```

Note: 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
```

Note: Enabling failover on one node enables it for both nodes.

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

Note: 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.

Note: 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
```

Note: 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 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 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 Jumbo Frames in NetApp 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>>
```

Note: 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.

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

Note: 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>>
```

Note: 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>>
```

Note: 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
```

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

Configure NFSv3 in ONTAP

Table 13 lists the information needed to complete this configuration.

Table 13) Information required for NFS configuration.

Detail	Detail Value
ESXi host A NFS IP address	<<var_esxi_hostA_nfs_ip>>
ESXi host B NFS IP address	<<var_esxi_hostB_nfs_ip>>

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
```

Note: 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.

Creating 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
```

Creating 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.

   Note: 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_clustername>>
```

   Note: 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
```

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

Create iSCSI LIFs in ONTAP

Table 14 lists the information needed to complete this configuration.

Table 14) Information required for iSCSI 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 –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

Table 15 lists the information needed to complete this configuration.

Table 15) Information required for NFS 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 Infrastructure SVM Administrator

Table 16 lists the information needed to complete this configuration.

Table 16) Information required for SVM administrator addition.

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
```

Note: 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
```

5.3 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.

Table 17 lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.

Table 17) Information required for CIMC configuration.

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

Note: 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.

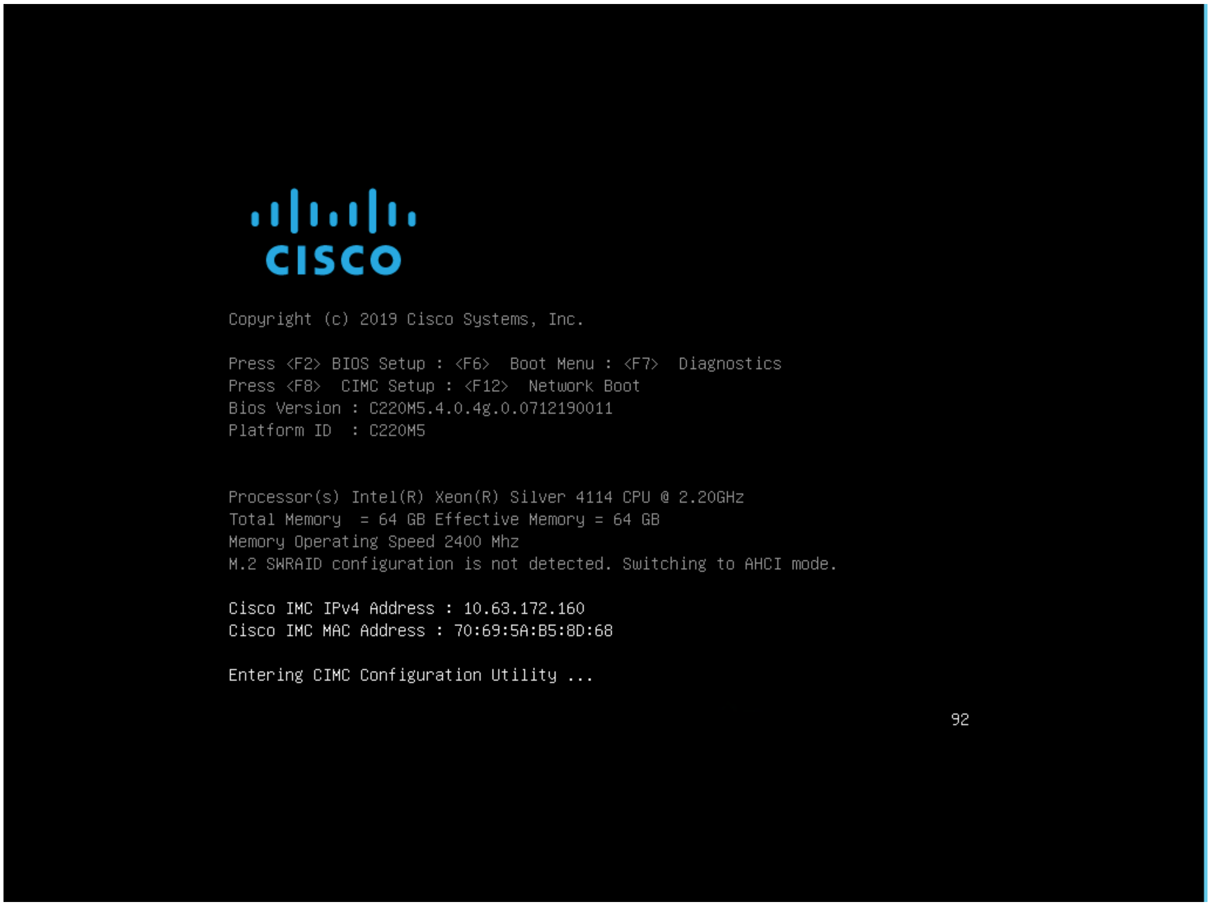

```
Copyright (c) 2019 Cisco Systems, Inc.

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 `[X]`
   b. IP (Basic):

      IPV4: `[X]`

      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: `[X]`

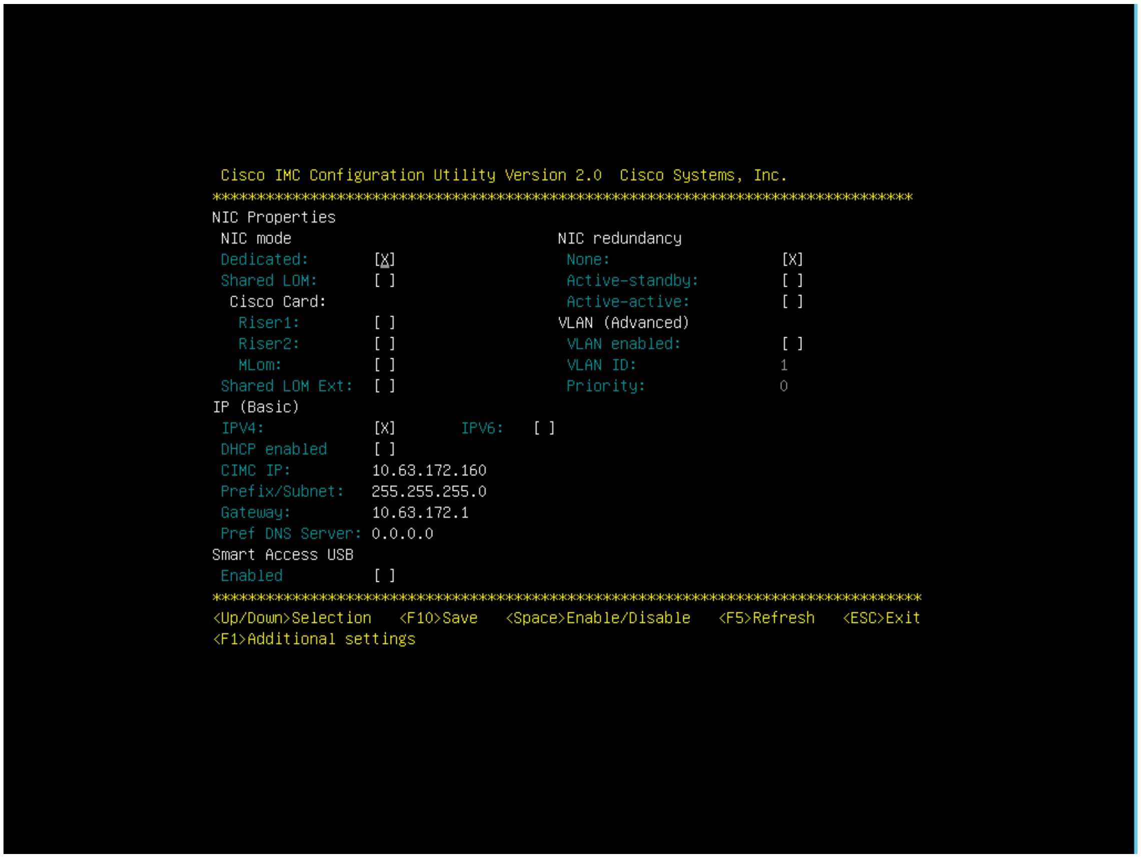

```
 Cisco IMC Configuration Utility Version 2.0  Cisco Systems, Inc.
********************************************************************************
NIC Properties
 NIC mode                              NIC redundancy
 Dedicated:       [X]                   None:                   [X]
 Shared LOM:      [ ]                   Active-standby:         [ ]
  Cisco Card:                           Active-active:          [ ]
   Riser1:        [ ]                  VLAN (Advanced)
   Riser2:        [ ]                   VLAN enabled:           [ ]
   MLom:          [ ]                   VLAN ID:                1
 Shared LOM Ext:  [ ]                   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 additional settings.
   d. Common properties:

      Host name: `<<esxi_host_name>>`

      Dynamic DNS: `[ ]`

      Factory defaults: Leave cleared.
   e. 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.

Table 18 lists the information needed to configure iSCSI boot.

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

Table 18) Information required for iSCSI boot configuration.

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.

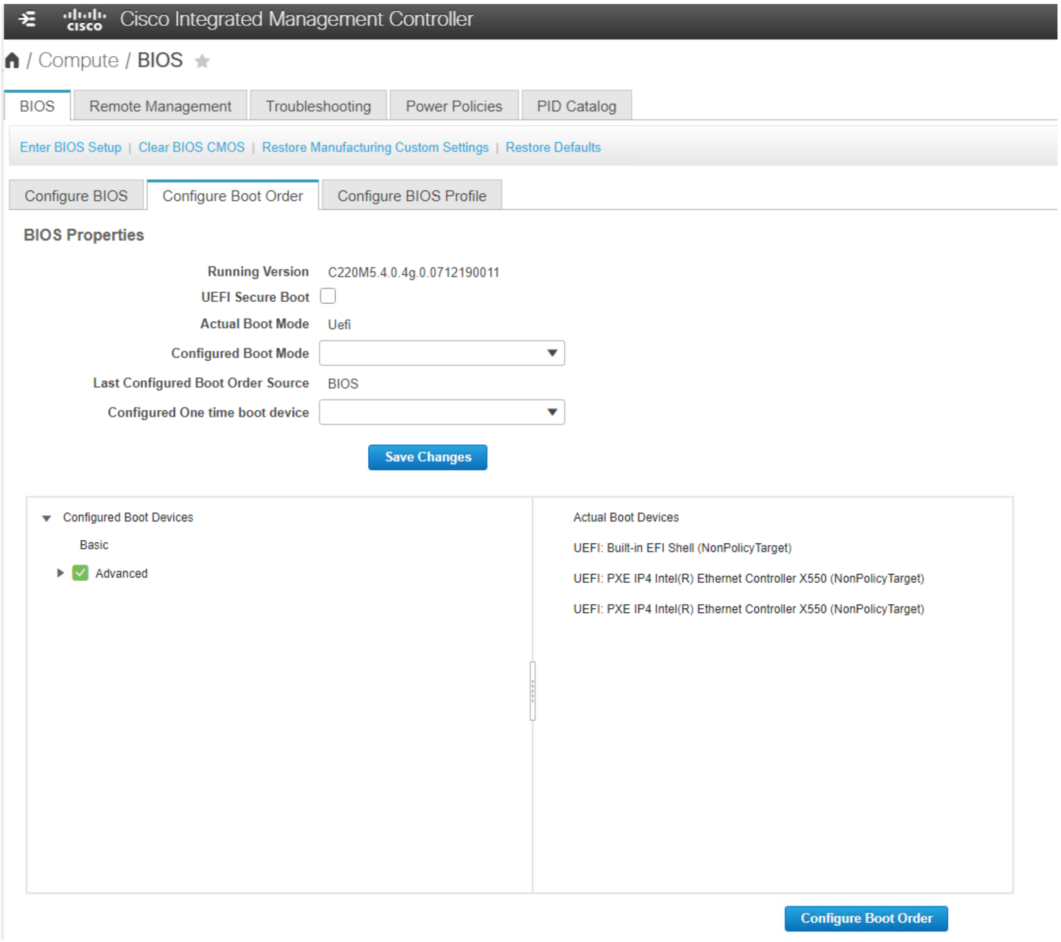

BIOS Properties

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

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)

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

    f. Add Virtual Media

       Name: KVM-CD-DVD

       Subtype: KVM MAPPED DVD

       State: Enabled

       Order: 1

    g. Add iSCSI Boot.

       Name: iSCSI-A

       State: Enabled

       Order: 2

       Slot: MLOM

       Port: 1

    h. 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.

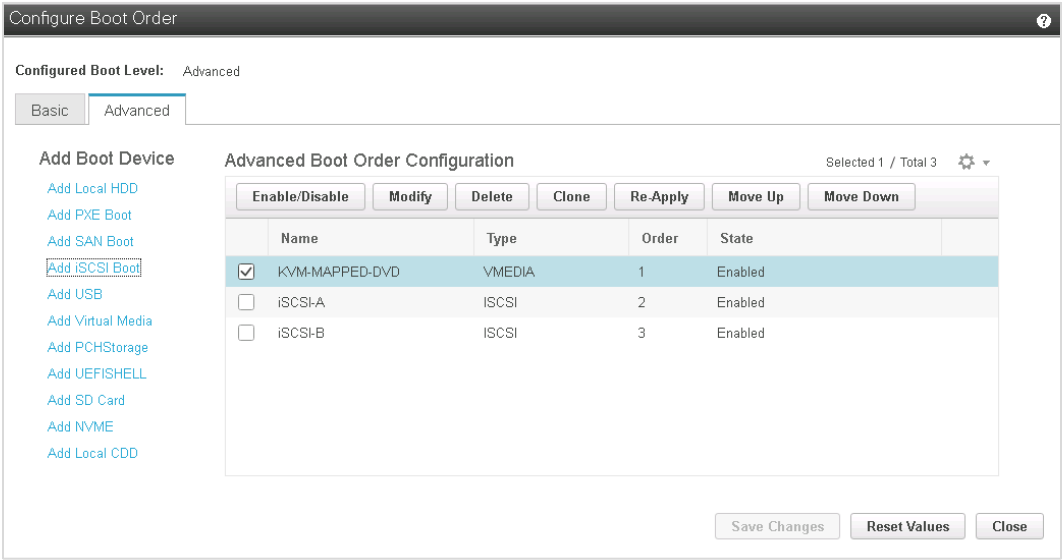

Configured Boot Level: Advanced

	Name	Type	Order	State
☑	KVM-MAPPED-DVD	VMEDIA	1	Enabled
□	iSCSI-A	ISCSI	2	Enabled
□	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.

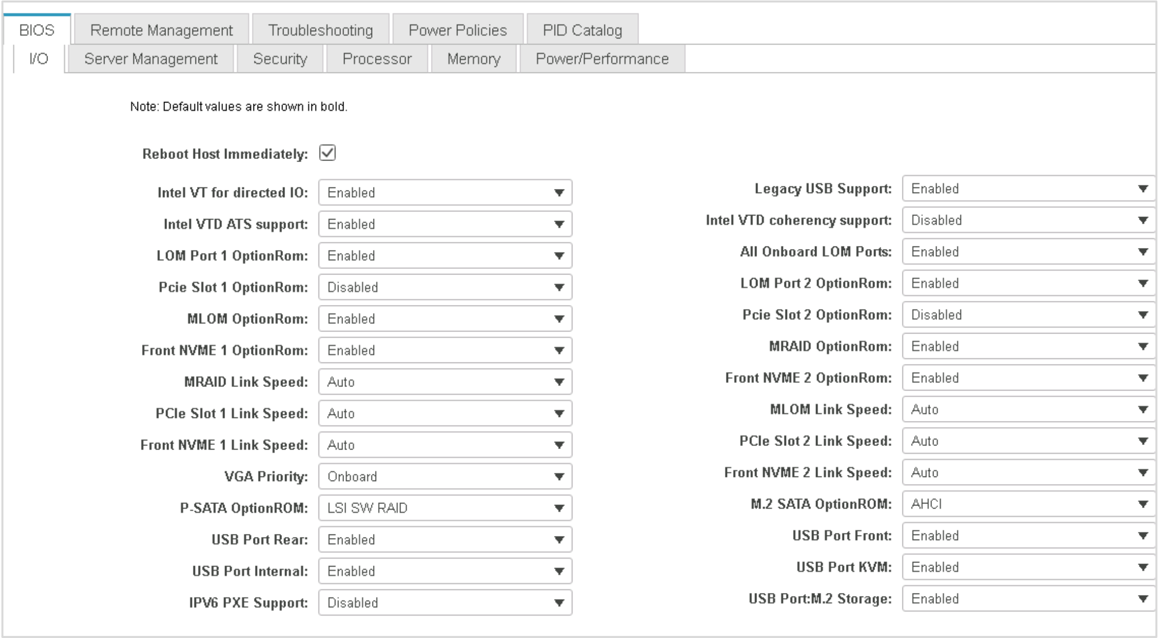

Note: Default values are shown in bold.

Reboot Host Immediately: ☑

Intel VT for directed IO:	Enabled	Legacy USB Support:	Enabled
Intel VTD ATS support:	Enabled	Intel VTD coherency support:	Disabled
LOM Port 1 OptionRom:	Enabled	All Onboard LOM Ports:	Enabled
Pcie Slot 1 OptionRom:	Disabled	LOM Port 2 OptionRom:	Enabled
MLOM OptionRom:	Enabled	Pcie Slot 2 OptionRom:	Disabled
Front NVME 1 OptionRom:	Enabled	MRAID OptionRom:	Enabled
MRAID Link Speed:	Auto	Front NVME 2 OptionRom:	Enabled
PCIe Slot 1 Link Speed:	Auto	MLOM Link Speed:	Auto
Front NVME 1 Link Speed:	Auto	PCIe Slot 2 Link Speed:	Auto
VGA Priority:	Onboard	Front NVME 2 Link Speed:	Auto
P-SATA OptionROM:	LSI SW RAID	M.2 SATA OptionROM:	AHCI
USB Port Rear:	Enabled	USB Port Front:	Enabled
USB Port Internal:	Enabled	USB Port KVM:	Enabled
IPV6 PXE Support:	Disabled	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.

Note: 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.

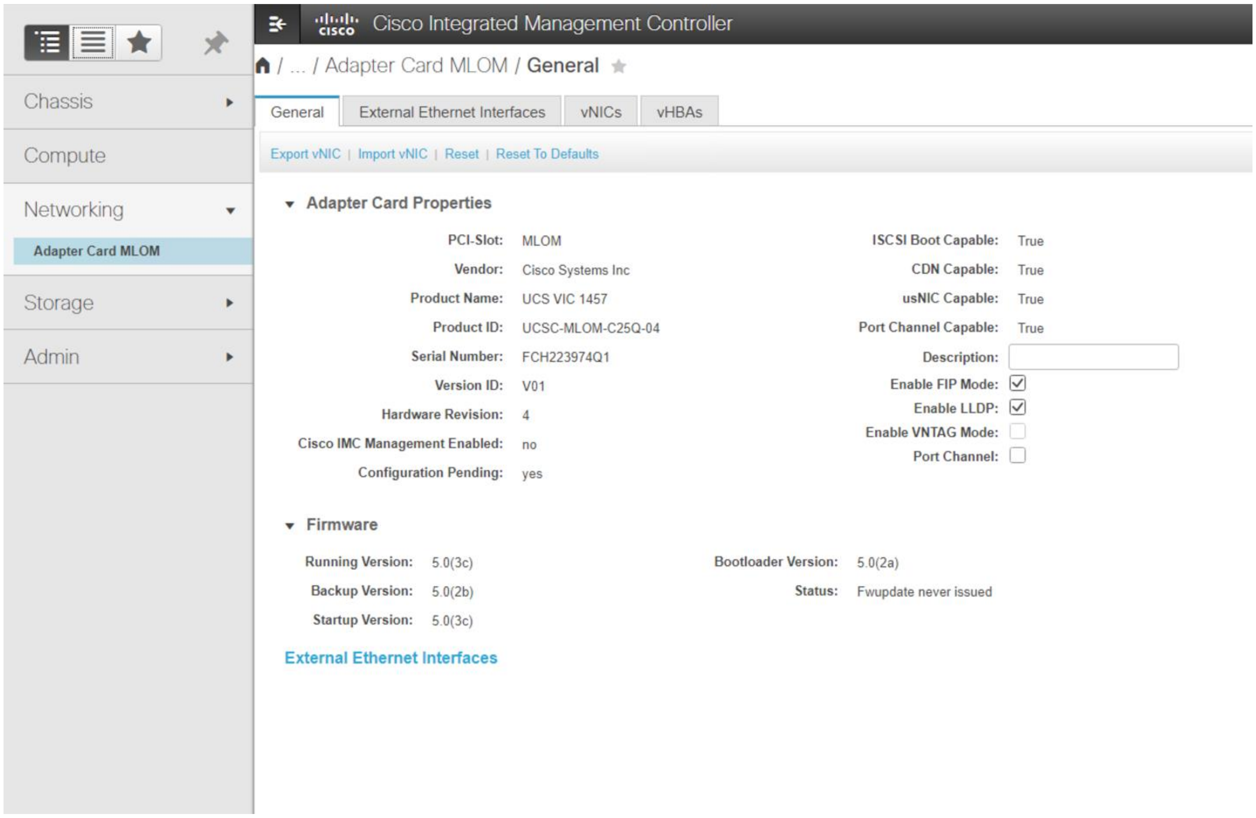
Cisco Integrated Management Controller
/ ... / Adapter Card MLOM / General
Chassis
Compute
Networking
Adapter Card MLOM
Storage
Admin
General
External Ethernet Interfaces
vNICs
vHBAs
Export vNIC | Import vNIC | Reset | Reset To Defaults
External Ethernet Interfaces

Adapter Card Properties

PCI-Slot: MLOM
Vendor: Cisco Systems Inc
Product Name: UCS VIC 1457
Product ID: UCSC-MLOM-C25Q-04
Serial Number: FCH223974Q1
Version ID: V01
Hardware Revision: 4
Cisco IMC Management Enabled: no
Configuration Pending: yes

ISCSI Boot Capable: True
CDN Capable: True
usNIC Capable: True
Port Channel Capable: True
Description:
Enable FIP Mode: ☑
Enable LLDP: ☑
Enable VNTAG Mode: □
Port Channel: □

Firmware

Running Version: 5.0(3c)
Backup Version: 5.0(2b)
Startup Version: 5.0(3c)
Bootloader Version: 5.0(2a)
Status: Fwupdate never issued

Create iSCSI vNICs

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.

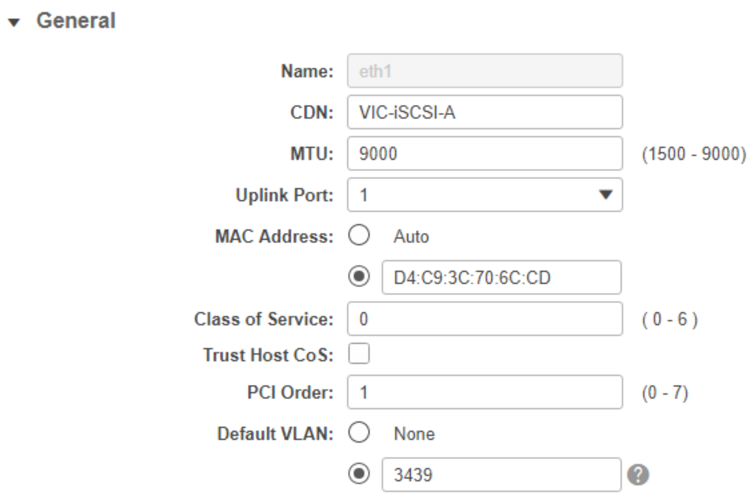
General

Name: eth1

CDN: VIC-iSCSI-A

MTU: 9000 (1500 - 9000)

Uplink Port: 1

MAC Address: ○ Auto

● D4:C9:3C:70:6C:CD

Class of Service: 0 (0 - 6)

Trust Host CoS: □

PCI Order: 1 (0 - 7)

Default VLAN: ○ None

● 3439

6. Select the vNIC eth1 on the left.

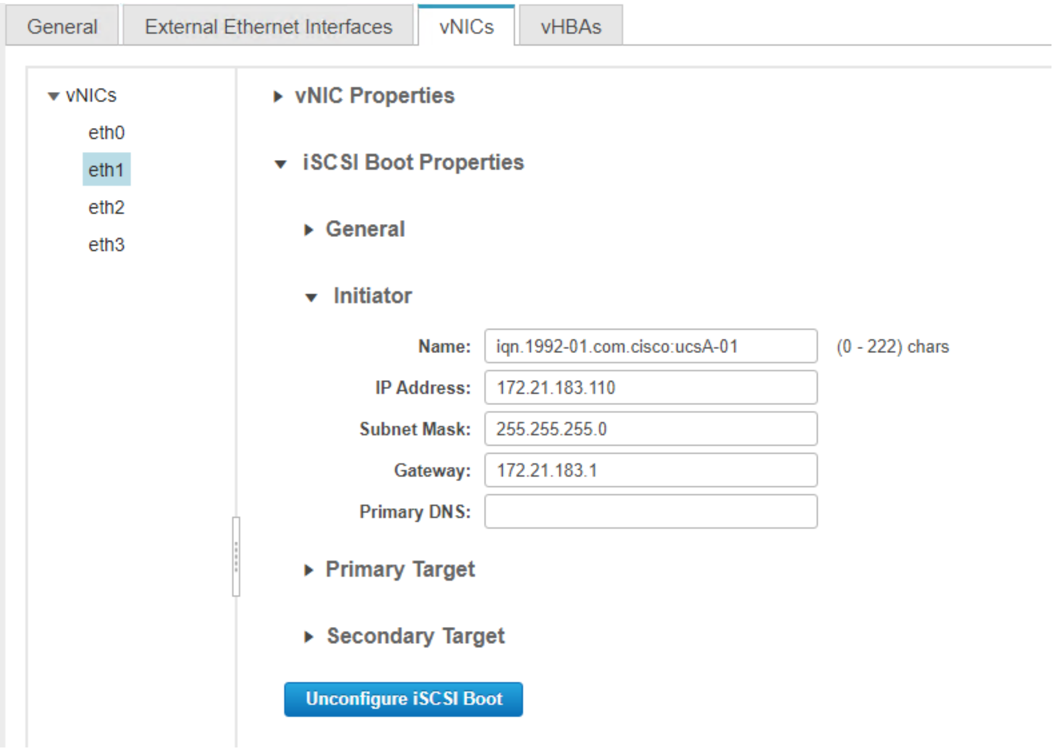

Name: iqn.1992-01.com.cisco:ucsA-01 (0 - 222) chars

IP Address: 172.21.183.110

Subnet Mask: 255.255.255.0

Gateway: 172.21.183.1

Primary DNS:

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>>`

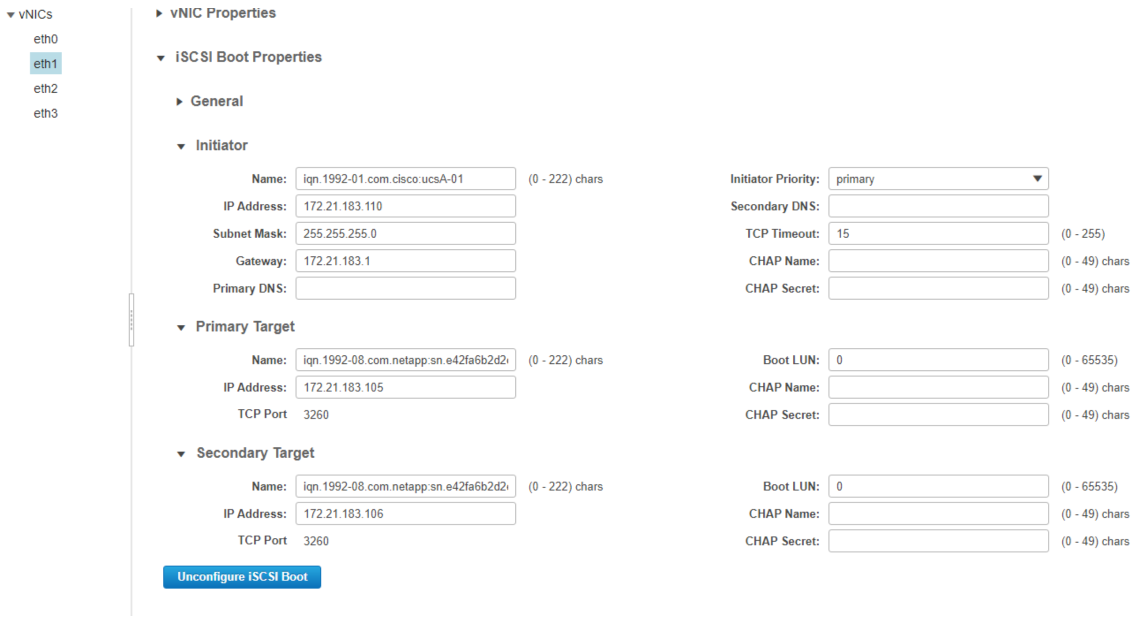

Initiator

Name:	iqn.1992-01.com.cisco:ucsA-01	(0 - 222) chars	Initiator Priority:	primary	
IP Address:	172.21.183.110		Secondary DNS:		
Subnet Mask:	255.255.255.0		TCP Timeout:	15	(0 - 255)
Gateway:	172.21.183.1		CHAP Name:		(0 - 49) chars
Primary DNS:			CHAP Secret:		(0 - 49) chars

Primary Target

Name:	iqn.1992-08.com.netapp:sn.e42fa6b2d2	(0 - 222) chars	Boot LUN:	0	(0 - 65535)
IP Address:	172.21.183.105		CHAP Name:		(0 - 49) chars
TCP Port	3260		CHAP Secret:		(0 - 49) chars

Secondary Target

Name:	iqn.1992-08.com.netapp:sn.e42fa6b2d2	(0 - 222) chars	Boot LUN:	0	(0 - 65535)
IP Address:	172.21.183.106		CHAP Name:		(0 - 49) chars
TCP Port	3260		CHAP Secret:		(0 - 49) chars

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

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

   Note: 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_ucsa_initiator_name_b>>`
    - IP address: `<<var_esxi_hostb_iscsib_ip>>`
    - Subnet mask: `<<var_esxi_hostb_iscsib_mask>>`
    - Gateway: `<<var_esxi_hostb_iscsib_gateway>>`

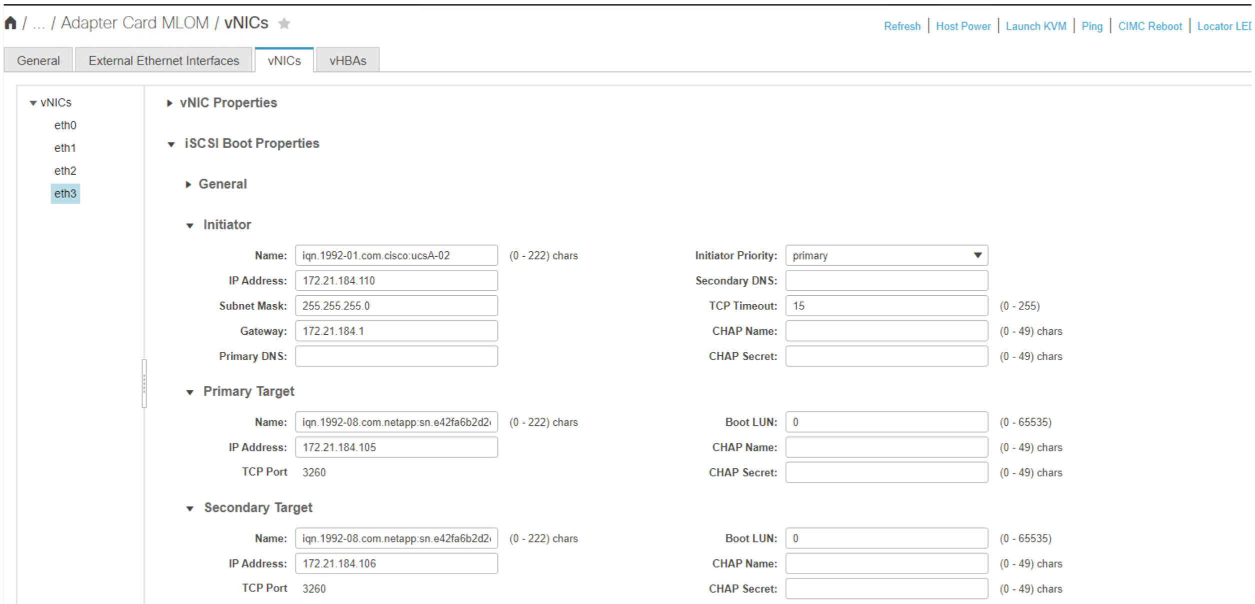
/ ... / Adapter Card MLOM / vNICs
Refresh | Host Power | Launch KVM | Ping | CIMC Reboot | Locator LEI
General
External Ethernet Interfaces
vNICs
vHBAs
vNICs
eth0
eth1
eth2
eth3
vNIC Properties
iSCSI Boot Properties
General

Initiator

Name: iqn.1992-01.com.cisco:ucsA-02 (0 - 222) chars
IP Address: 172.21.184.110
Subnet Mask: 255.255.255.0
Gateway: 172.21.184.1
Primary DNS:
Initiator Priority: primary
Secondary DNS:
TCP Timeout: 15 (0 - 255)
CHAP Name: (0 - 49) chars
CHAP Secret: (0 - 49) chars

Primary Target

Name: iqn.1992-08.com.netapp:sn.e42fa6b2d2 (0 - 222) chars
IP Address: 172.21.184.105
TCP Port 3260
Boot LUN: 0 (0 - 65535)
CHAP Name: (0 - 49) chars
CHAP Secret: (0 - 49) chars

Secondary Target

Name: iqn.1992-08.com.netapp:sn.e42fa6b2d2 (0 - 222) chars
IP Address: 172.21.184.106
TCP Port 3260
Boot LUN: 0 (0 - 65535)
CHAP Name: (0 - 49) chars
CHAP Secret: (0 - 49) chars

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

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

    Note: 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

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.

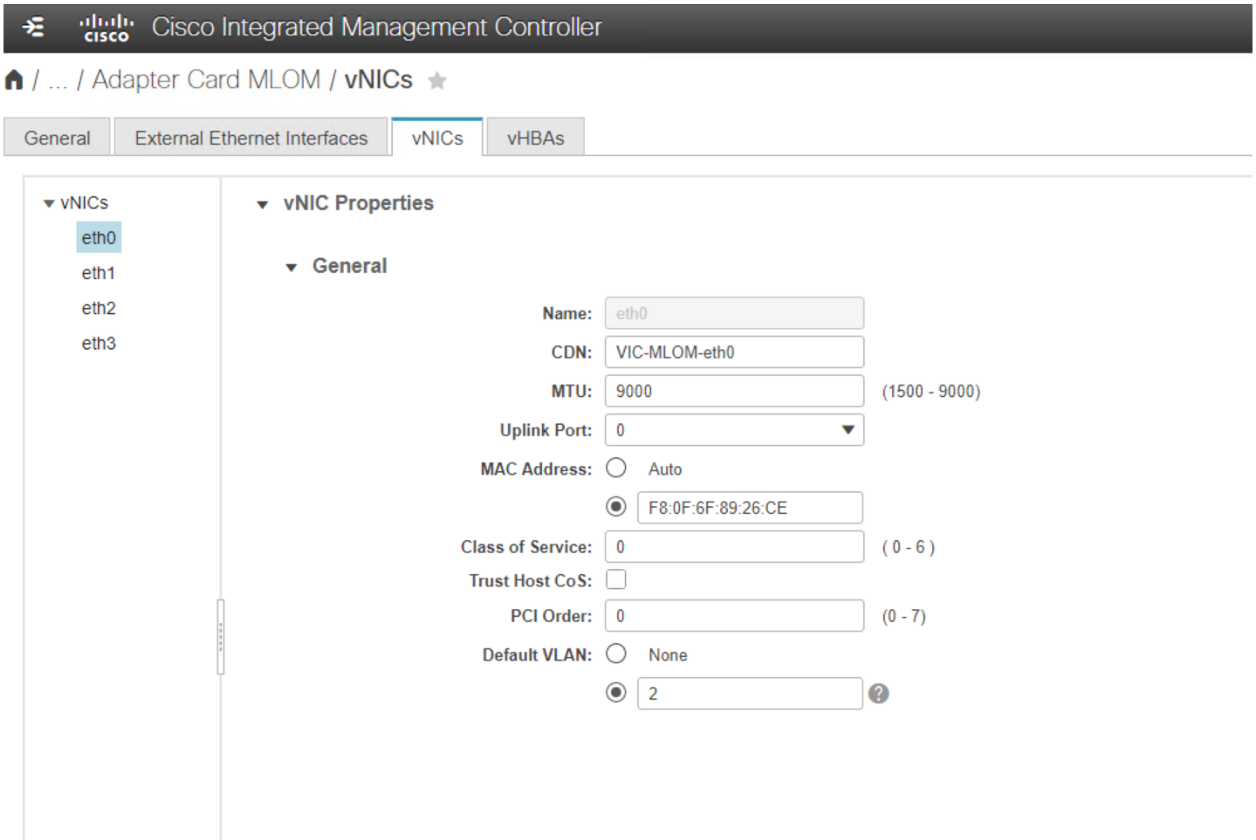
Cisco Integrated Management Controller
/ ... / Adapter Card MLOM / vNICs
General
External Ethernet Interfaces
vNICs
vHBAs
vNICs
eth0
eth1
eth2
eth3
vNIC Properties
General

Name: eth0

CDN: VIC-MLOM-eth0

MTU: 9000 (1500 - 9000)

Uplink Port: 0

MAC Address: ○ Auto

● F8:0F:6F:89:26:CE

Class of Service: 0 (0 - 6)

Trust Host CoS: □

PCI Order: 0 (0 - 7)

Default VLAN: ○ None

● 2

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

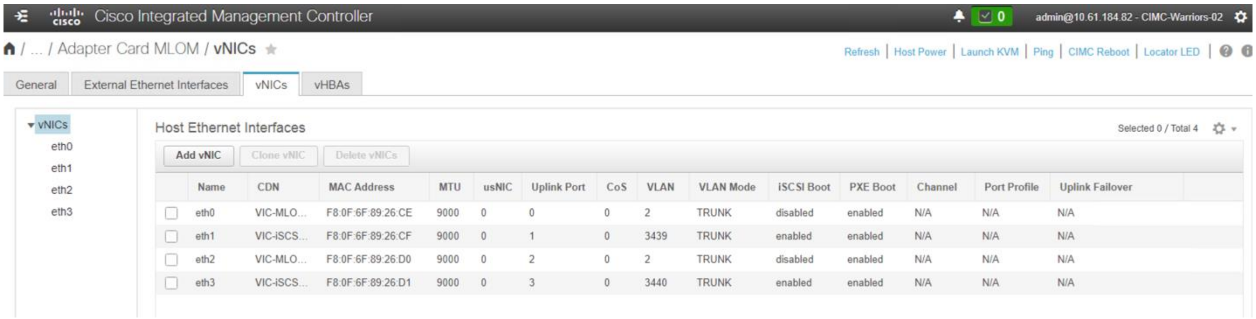
Cisco Integrated Management Controller
0
admin@10.61.184.82 - CIMC-Warriors-02
/ ... / Adapter Card MLOM / vNICs
Refresh | Host Power | Launch KVM | Ping | CIMC Reboot | Locator LED
General
External Ethernet Interfaces
vNICs
vHBAs
vNICs
eth0
eth1
eth2
eth3
Host Ethernet Interfaces
Selected 0 / Total 4
Add vNIC
Clone vNIC
Delete vNICs

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
□	eth0	VIC-MLO...	F8:0F:6F:89:26:CE	9000	0	0	0	2	TRUNK	disabled	enabled	N/A	N/A	N/A
□	eth1	VIC-iSCS...	F8:0F:6F:89:26:CF	9000	0	1	0	3439	TRUNK	enabled	enabled	N/A	N/A	N/A
□	eth2	VIC-MLO...	F8:0F:6F:89:26:D0	9000	0	2	0	2	TRUNK	disabled	enabled	N/A	N/A	N/A
□	eth3	VIC-iSCS...	F8:0F:6F:89:26:D1	9000	0	3	0	3440	TRUNK	enabled	enabled	N/A	N/A	N/A

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

5.4 NetApp AFF Storage Deployment Procedure (Part 2)

ONTAP SAN Boot Storage Setup

Create iSCSI Igroups

Note: 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>>
```

Note: 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
```

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

5.5 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.

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

Note: 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.

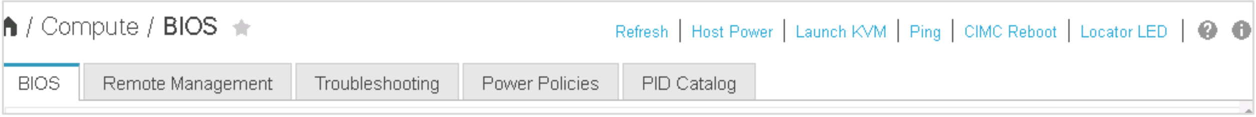

5. From the virtual KVM console, select the Virtual Media tab.
6. Select Map CD/DVD.

   Note: 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.

Installing 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.

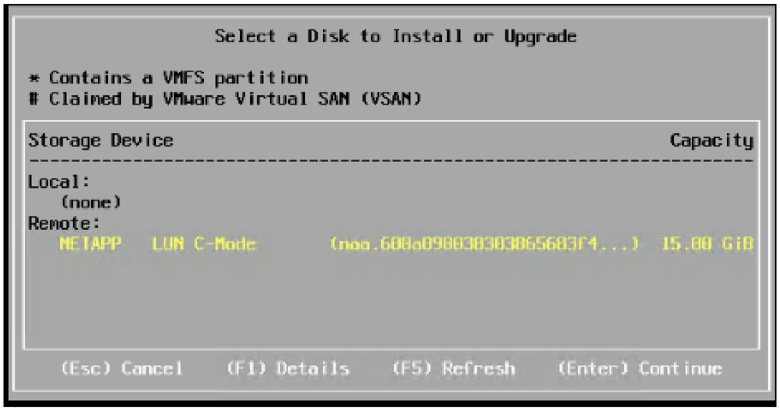

```
                 Select a Disk to Install or Upgrade

* Contains a VMFS partition
# Claimed by VMware Virtual SAN (VSAN)

Storage Device                                                  Capacity
-------------------------------------------------------------------------
Local:
   (none)
Remote:
   NETAPP   LUN C-Mode        (naa.600a098038303865683f4...)  15.00 GiB

   (Esc) Cancel     (F1) Details     (F5) Refresh     (Enter) Continue
```

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.

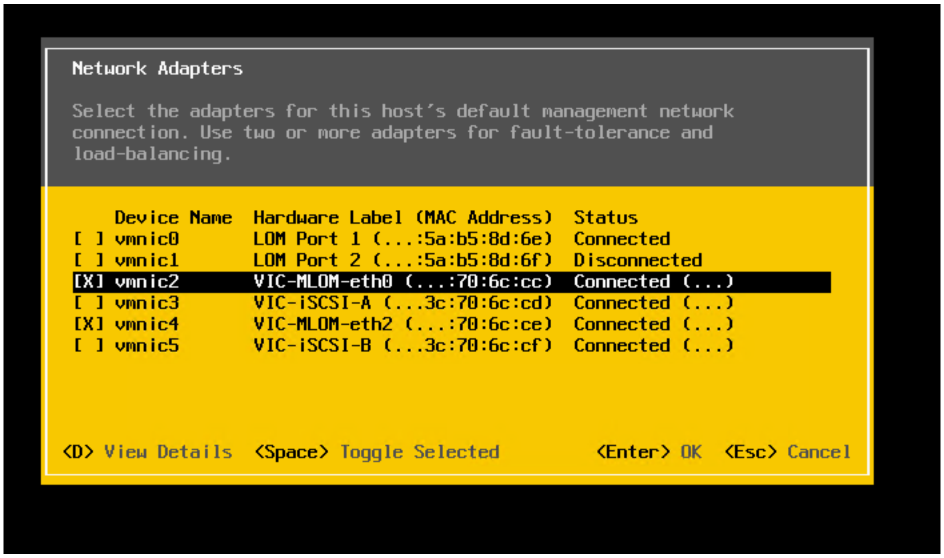

```
Network Adapters

Select the adapters for this host's default management network
connection. Use two or more adapters for fault-tolerance and
load-balancing.

    Device Name  Hardware Label (MAC Address)  Status
[ ] vmnic0       LOM Port 1 (...:5a:b5:8d:6e)  Connected
[ ] vmnic1       LOM Port 2 (...:5a:b5:8d:6f)  Disconnected
[X] vmnic2       VIC-MLOM-eth0 (...:70:6c:cc)  Connected (...)
[ ] vmnic3       VIC-iSCSI-A (...3c:70:6c:cd)  Connected (...)
[X] vmnic4       VIC-MLOM-eth2 (...:70:6c:ce)  Connected (...)
[ ] vmnic5       VIC-iSCSI-B (...3c:70:6c:cf)  Connected (...)

<D> View Details  <Space> Toggle Selected         <Enter> OK  <Esc> Cancel
```

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.

    Note: 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

You need the following information to configure each ESXi host.

Table 19) Information required for configuring ESXi hosts.

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 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.

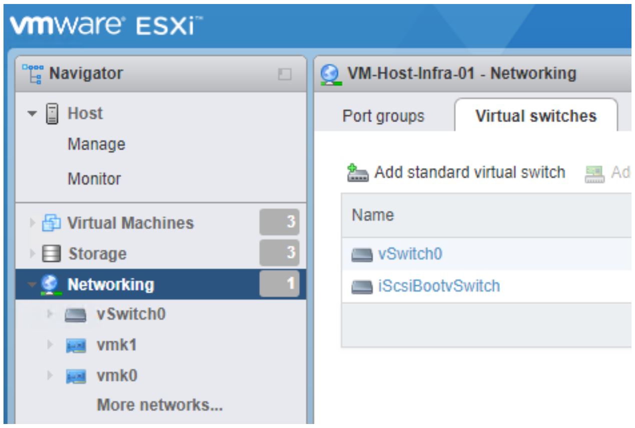

3. Click iScsiBootvSwitch.
4. Select Edit settings.
5. Change the MTU to 9000 and click Save.
6. Rename the iSCSIBootPG port to iSCSIBootPG-A.

   Note: 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.

   Note: Set the MTU to 9000 on iScsiBootPG-A.

9. To set the failover, complete the following steps:
   a. Click Edit Settings on iSCSIBootPG-A > Tiering and Failover > Failover Order > Vmnic3. Vmnic3 should be active and vmnic5 should be unused.
   b. Click Edit Settings on iSCSIBootPG-B > Teaming and Failover > Failover order > Vmnic5. Vmnic5 should be active and vmnic3 should be unused.

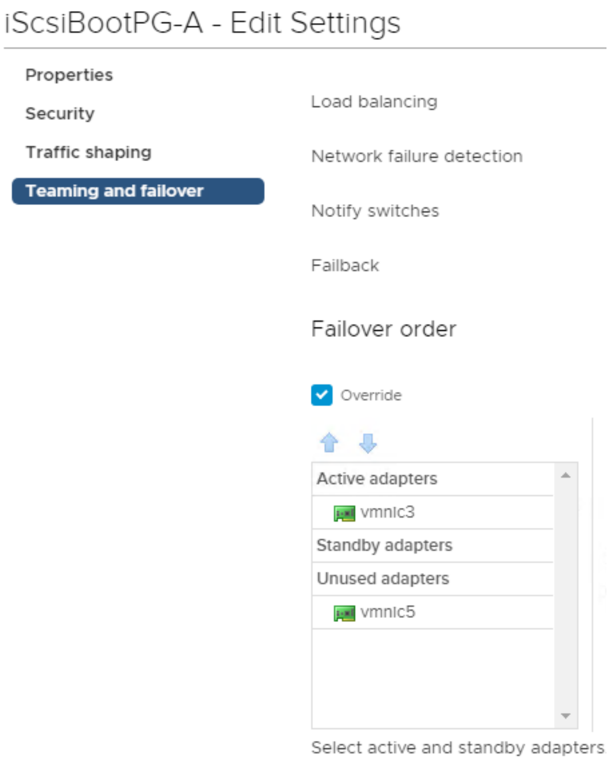

iScsiBootPG-A - Edit Settings

- Properties
- Security
- Traffic shaping
- **Teaming and failover**

Load balancing

Network failure detection

Notify switches

Failback

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:

1. Select Storage in the left navigation pane. Click Adapters.
2. Select the iSCSI software adapter and click Configure iSCSI.

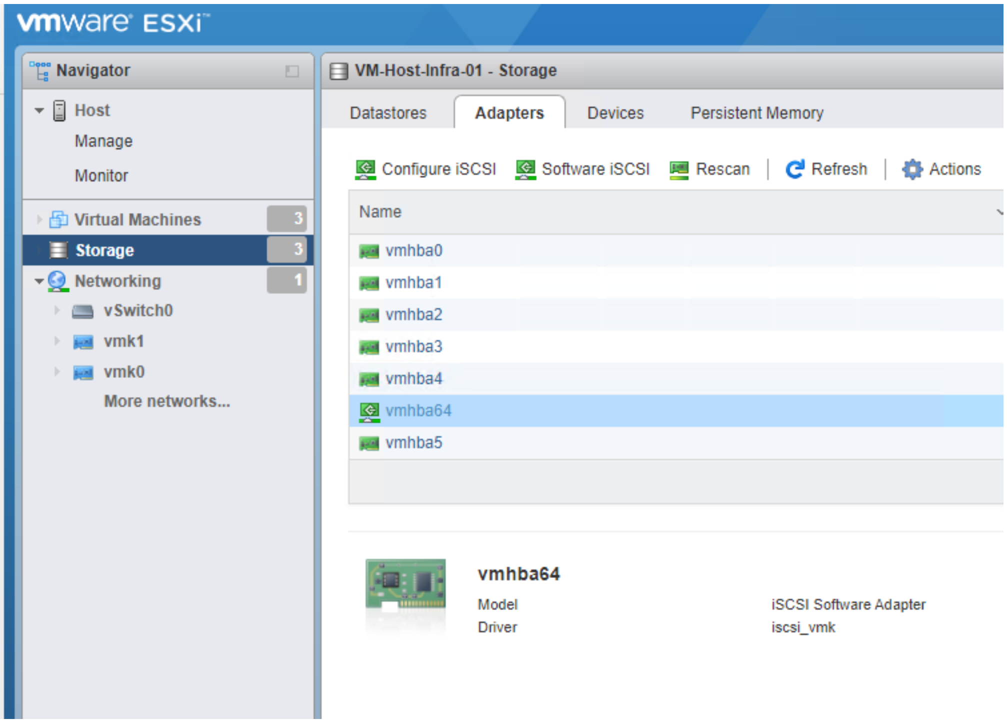

Name
vmhba0
vmhba1
vmhba2
vmhba3
vmhba4
vmhba64
vmhba5

vmhba64

Model	iSCSI Software Adapter
Driver	iscsi_vmk

3. Under Dynamic Targets, click Add Dynamic Target.

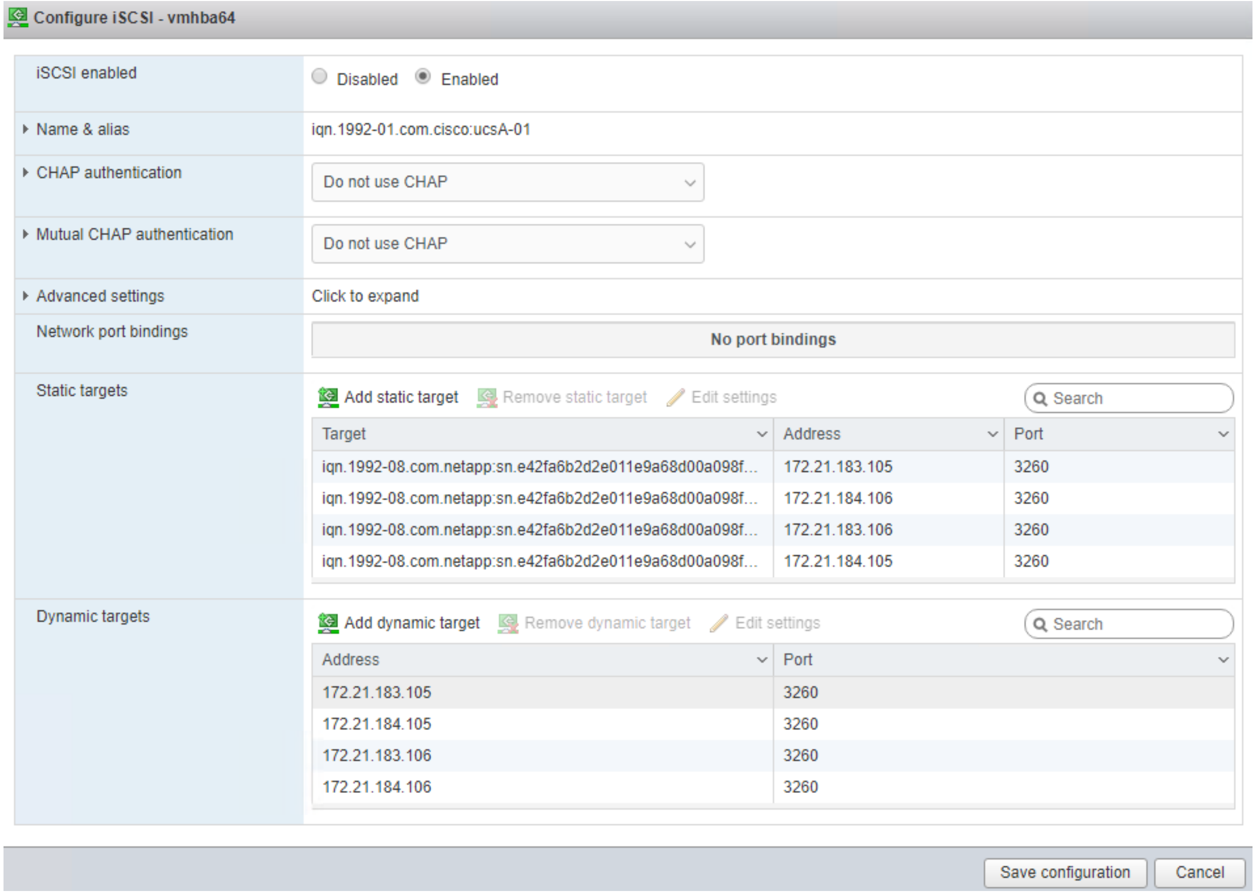

Configure iSCSI - vmhba64

iSCSI enabled: ○ Disabled ● 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

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

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.
   a. Repeat with the IP addresses iscsi_lif01b, iscsi_lif02a, and iscsi_lif02b.
   b. Click Save Configuration.

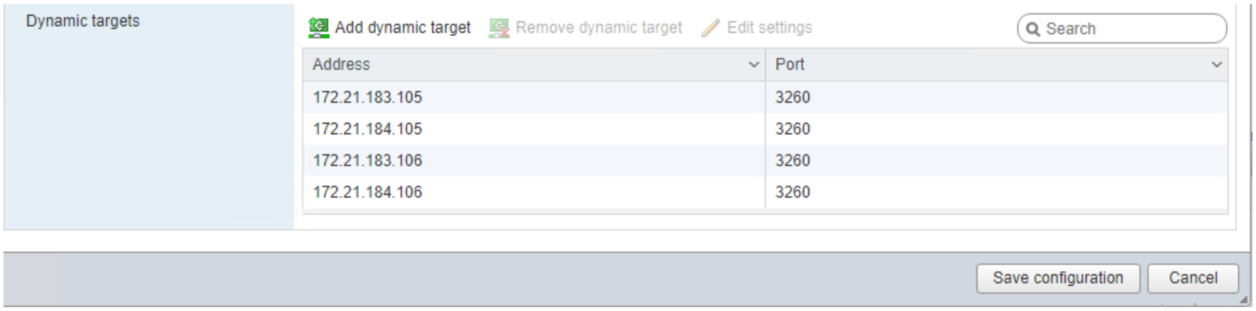

Dynamic targets

Address	Port
172.21.183.105	3260
172.21.184.105	3260
172.21.183.106	3260
172.21.184.106	3260

Note: 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 ESXi Host

1. In the left navigation pane, select Networking.
2. Select vSwitch0.

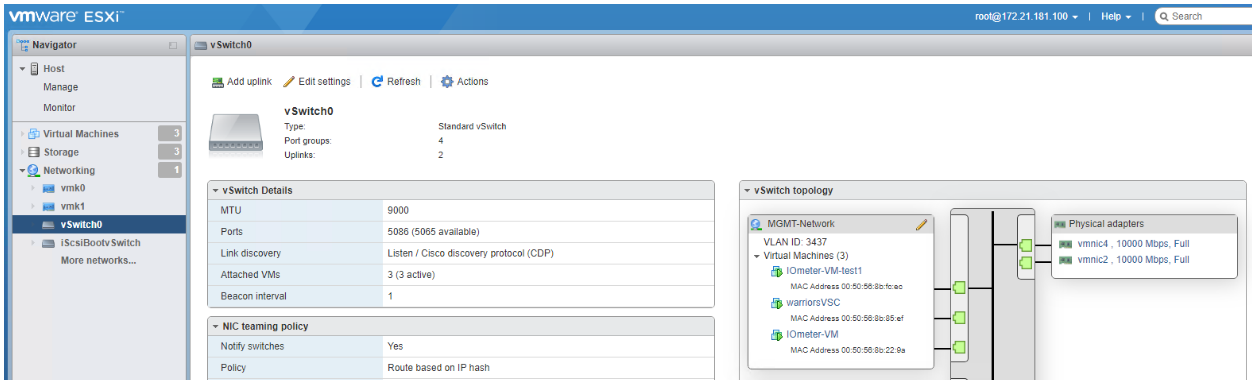

vSwitch0

Type:	Standard vSwitch
Port groups:	4
Uplinks:	2

vSwitch Details

MTU	9000
Ports	5086 (5065 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	3 (3 active)
Beacon interval	1

NIC teaming policy

Notify switches	Yes
Policy	Route based on IP hash

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.

Note: 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 Port Groups and VMkernel NICs

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

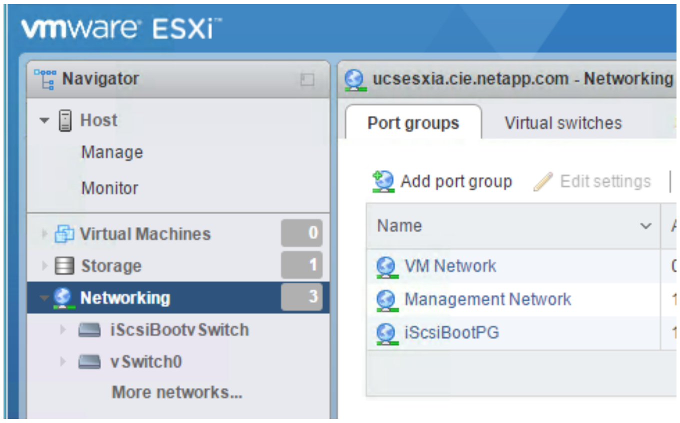

3. Right-click VM Network and select Edit. Change the VLAN ID to `<<var_vm_traffic_vlan>>`.
4. Click Add Port Group.
    a. Name the port group MGMT-Network.
    b. Enter `<<mgmt_vlan>>` for the VLAN ID.
    c. Make sure that vSwitch0 is selected.
    d. Click save.
5. Click the VMkernel NICs tab.

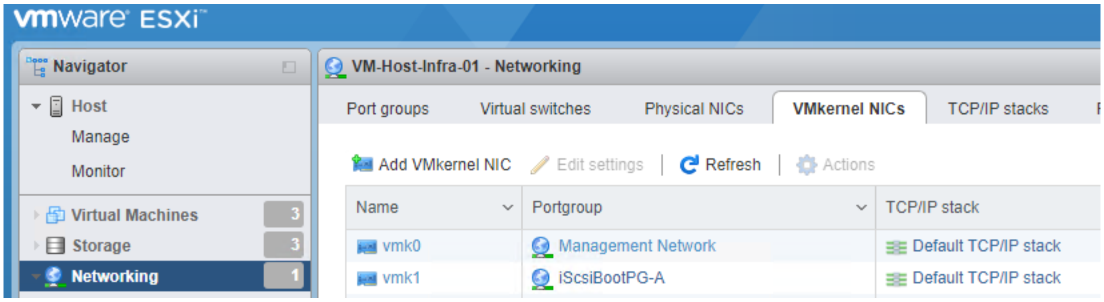

VM-Host-Infra-01 - Networking

Name	Portgroup	TCP/IP stack
vmk0	Management Network	Default TCP/IP stack
vmk1	iScsiBootPG-A	Default TCP/IP stack

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.

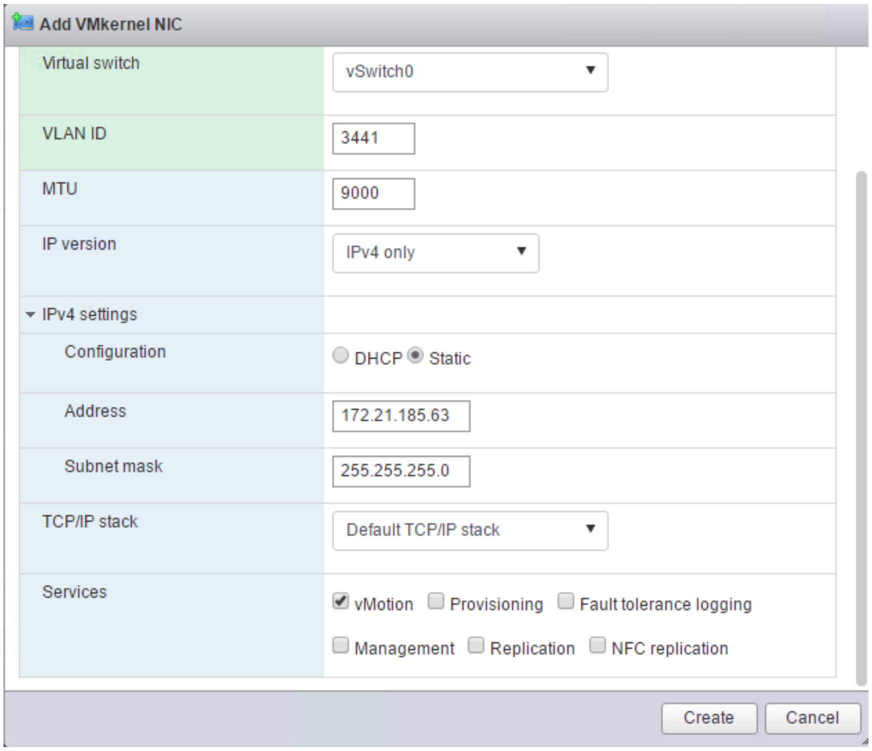

Add VMkernel NIC

Virtual switch	vSwitch0 ▼
VLAN ID	3441
MTU	9000
IP version	IPv4 only ▼
IPv4 settings	
Configuration	○ DHCP ● Static
Address	172.21.185.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack ▼
Services	☑ vMotion ☐ Provisioning ☐ Fault tolerance logging ☐ Management ☐ Replication ☐ NFC replication

Note: 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.

Mount 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.

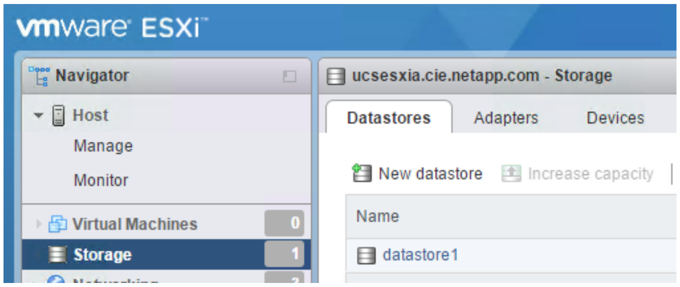

2. Select Mount NFS Datastore.

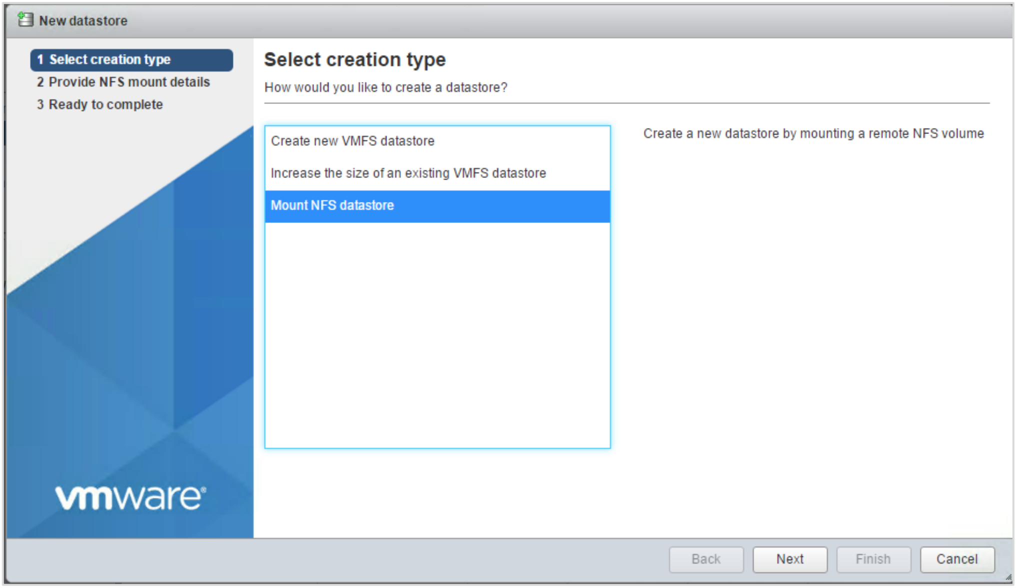

New datastore

1 Select creation type
2 Provide NFS mount details
3 Ready to complete

Select creation type

How would you like to create a datastore?

Create new VMFS datastore
Increase the size of an existing VMFS datastore
Mount NFS datastore

Create a new datastore by mounting a remote NFS volume

vmware
Back
Next
Finish
Cancel

3. Next, 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.

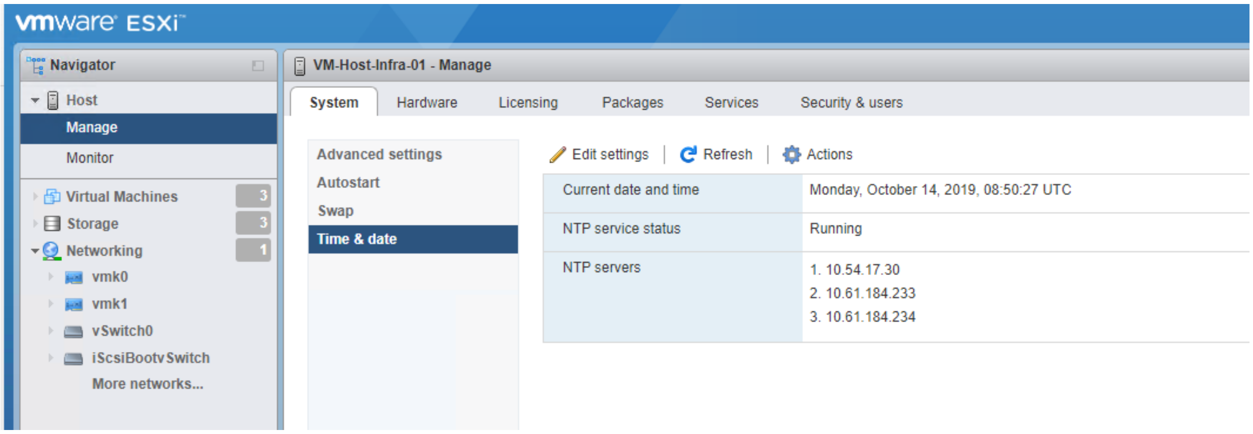

Current date and time	Monday, October 14, 2019, 08:50:27 UTC
NTP service status	Running

NTP servers	1. 10.54.17.30
	2. 10.61.184.233
	3. 10.61.184.234

Move the Virtual Machine 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.

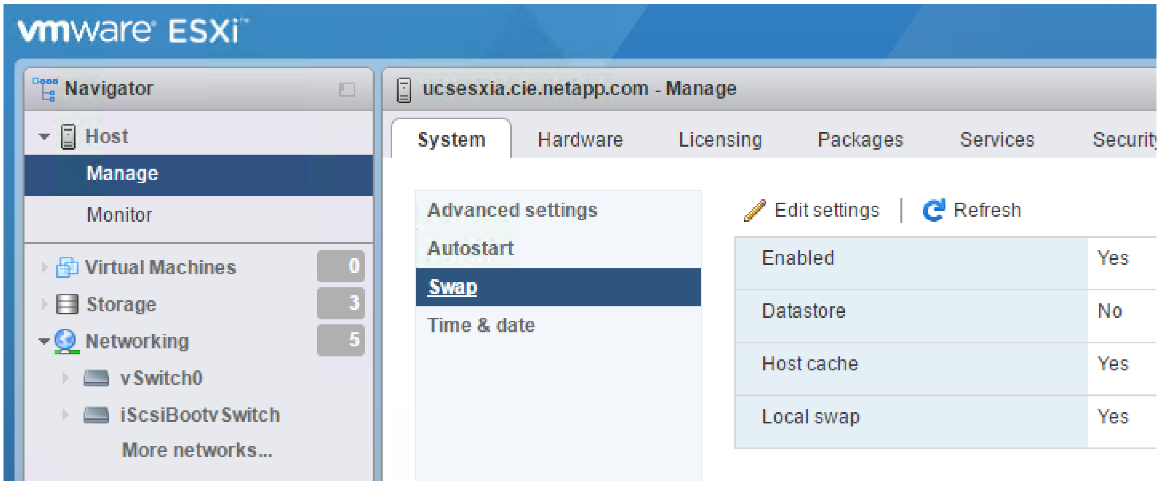

Enabled	Yes
Datastore	No
Host cache	Yes
Local swap	Yes

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

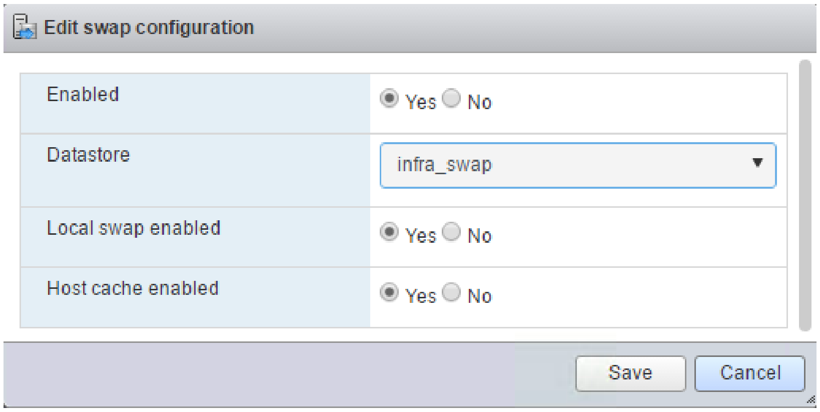

Enabled	● Yes ○ No
Datastore	infra_swap ▼
Local swap enabled	● Yes ○ No
Host cache enabled	● Yes ○ No

3. Click Save.

5.6 Install VMware vCenter Server 6.7U2

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

Note: 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.
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.

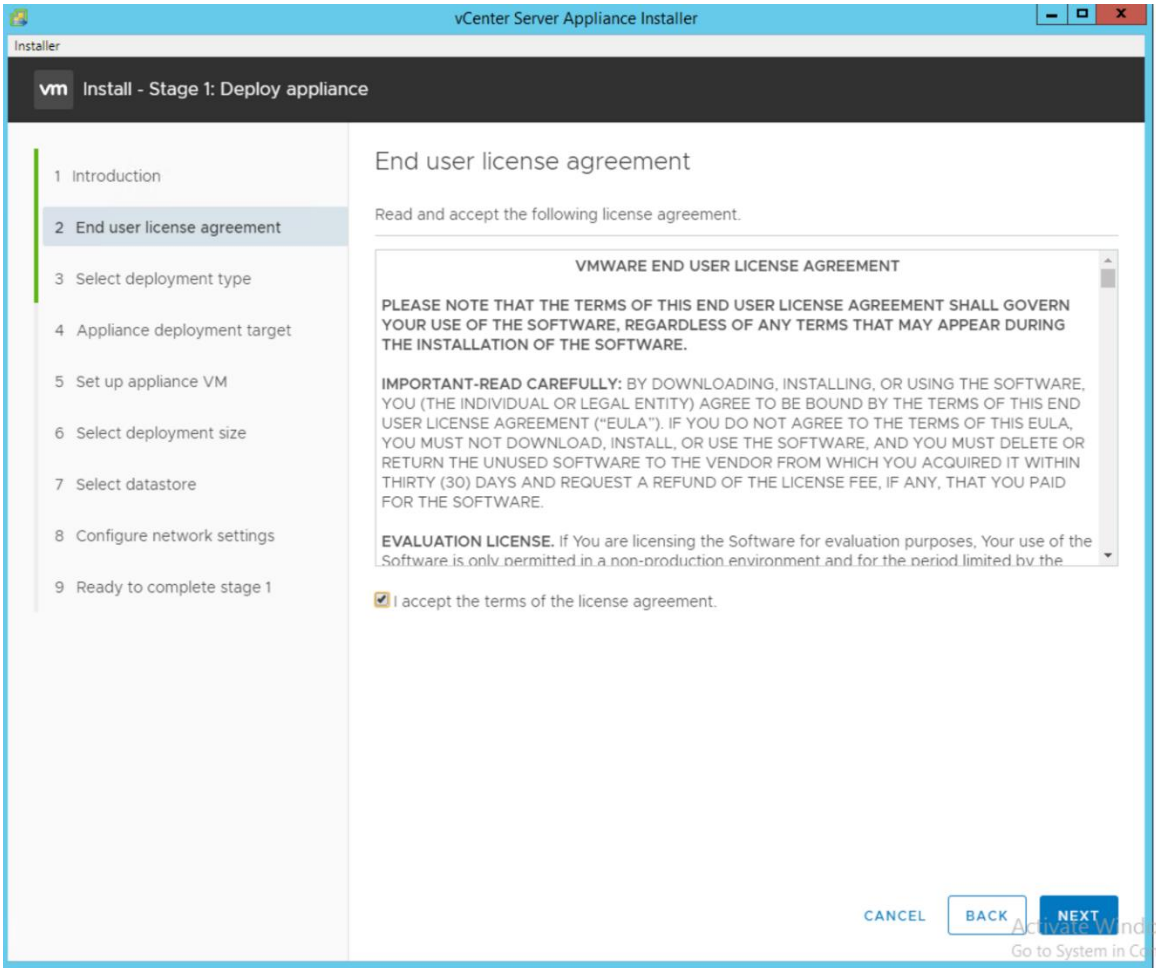
vCenter Server Appliance Installer
Installer
vm Install - Stage 1: Deploy appliance
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
CANCEL
BACK
NEXT
Activate Windo
Go to System in Co

End user license agreement

Read and accept the following license agreement.

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE VENDOR FROM WHICH YOU ACQUIRED IT WITHIN THIRTY (30) DAYS AND REQUEST A REFUND OF THE LICENSE FEE, IF ANY, THAT YOU PAID FOR THE SOFTWARE.

EVALUATION LICENSE. If You are licensing the Software for evaluation purposes, Your use of the Software is only permitted in a non-production environment and for the period limited by the

☑ I accept the terms of the license agreement.

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

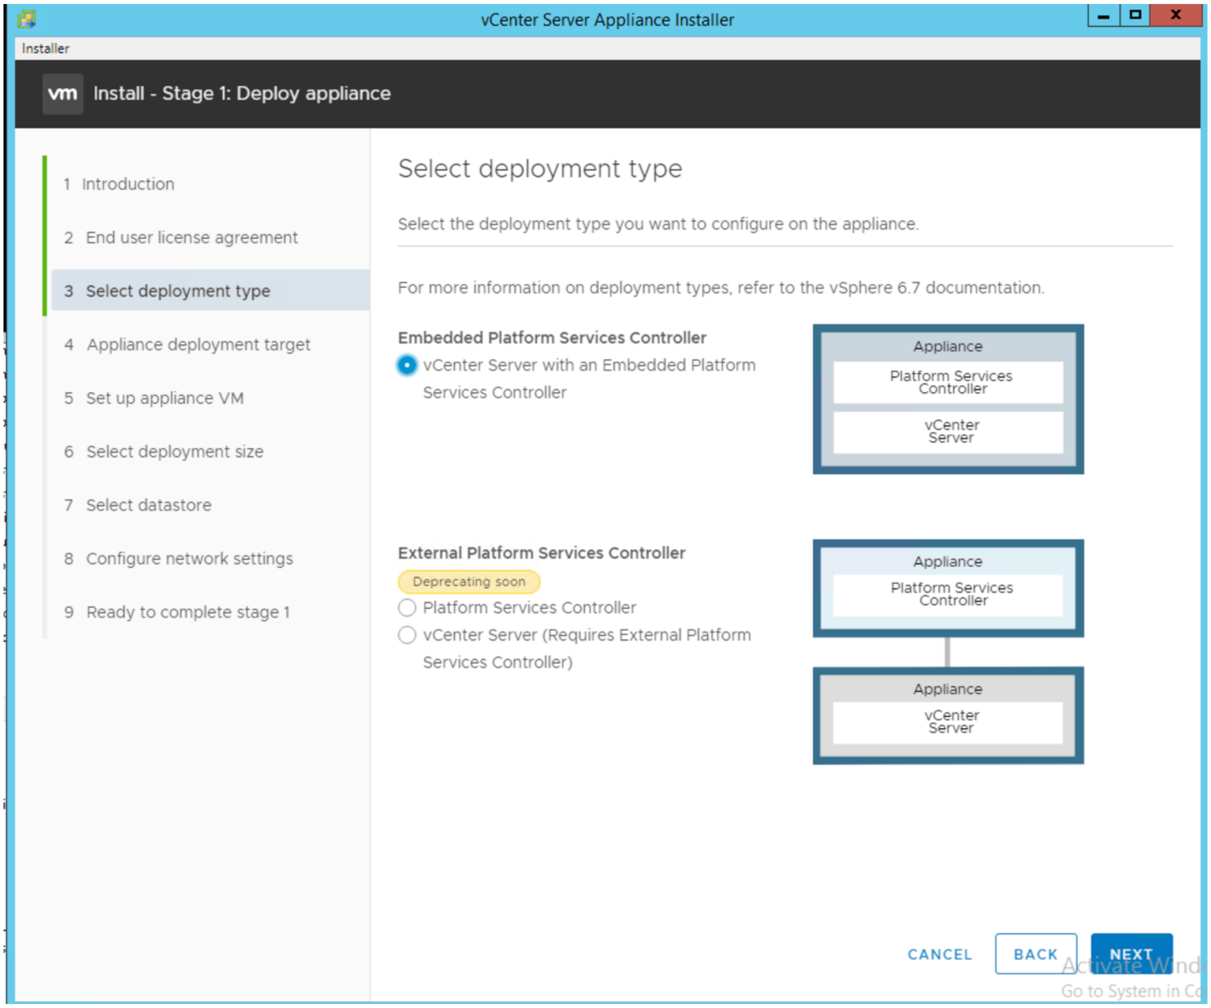

vCenter Server Appliance Installer
Installer
vm Install - Stage 1: Deploy appliance
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Appliance
Platform Services Controller
vCenter Server
Appliance
Platform Services Controller
Appliance
vCenter Server
CANCEL
BACK
NEXT

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)

Note: 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.

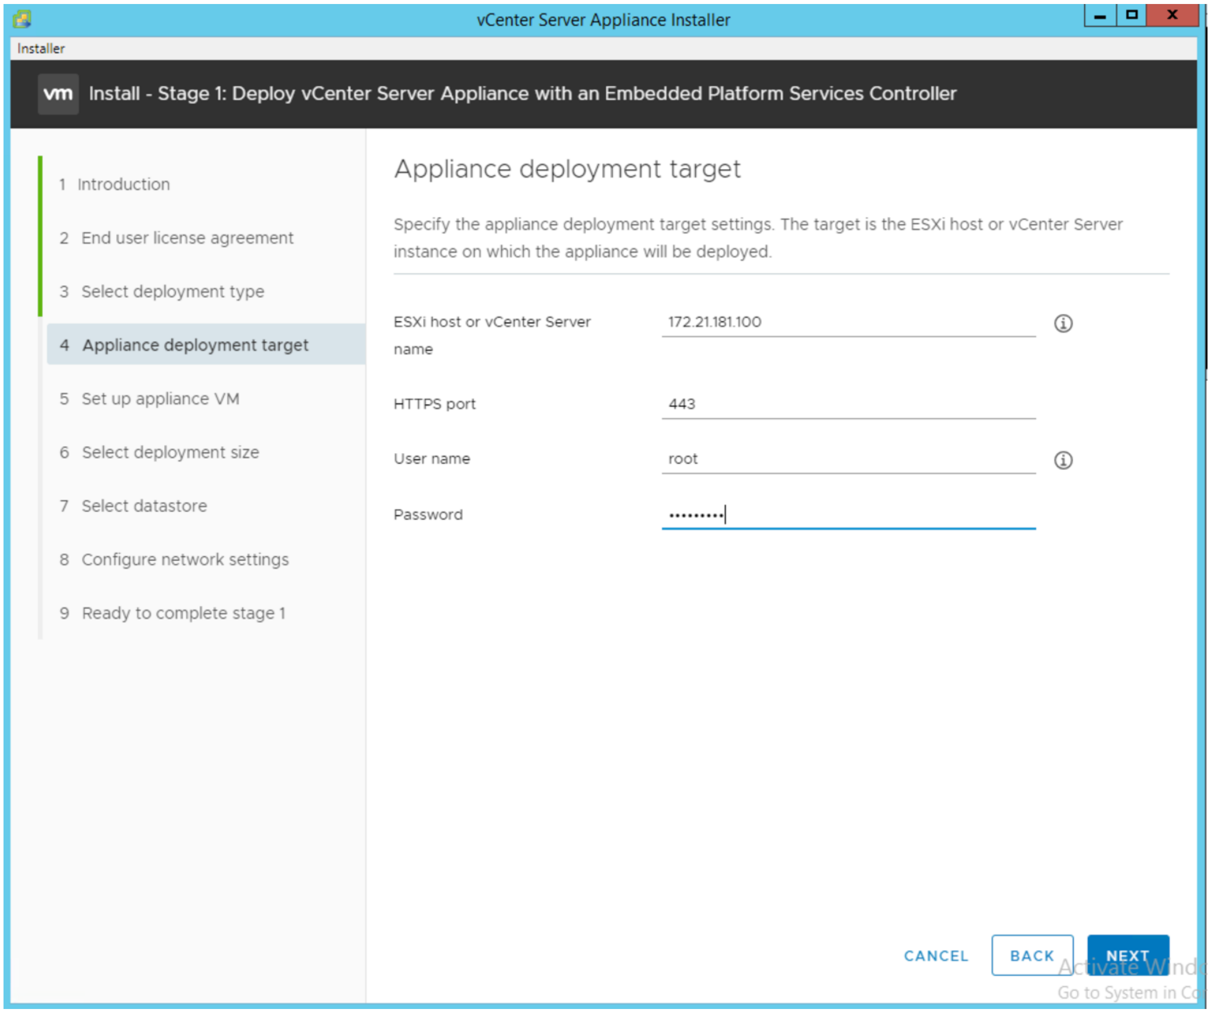
vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
CANCEL
BACK
NEXT

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: 172.21.181.100

HTTPS port: 443

User name: 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.

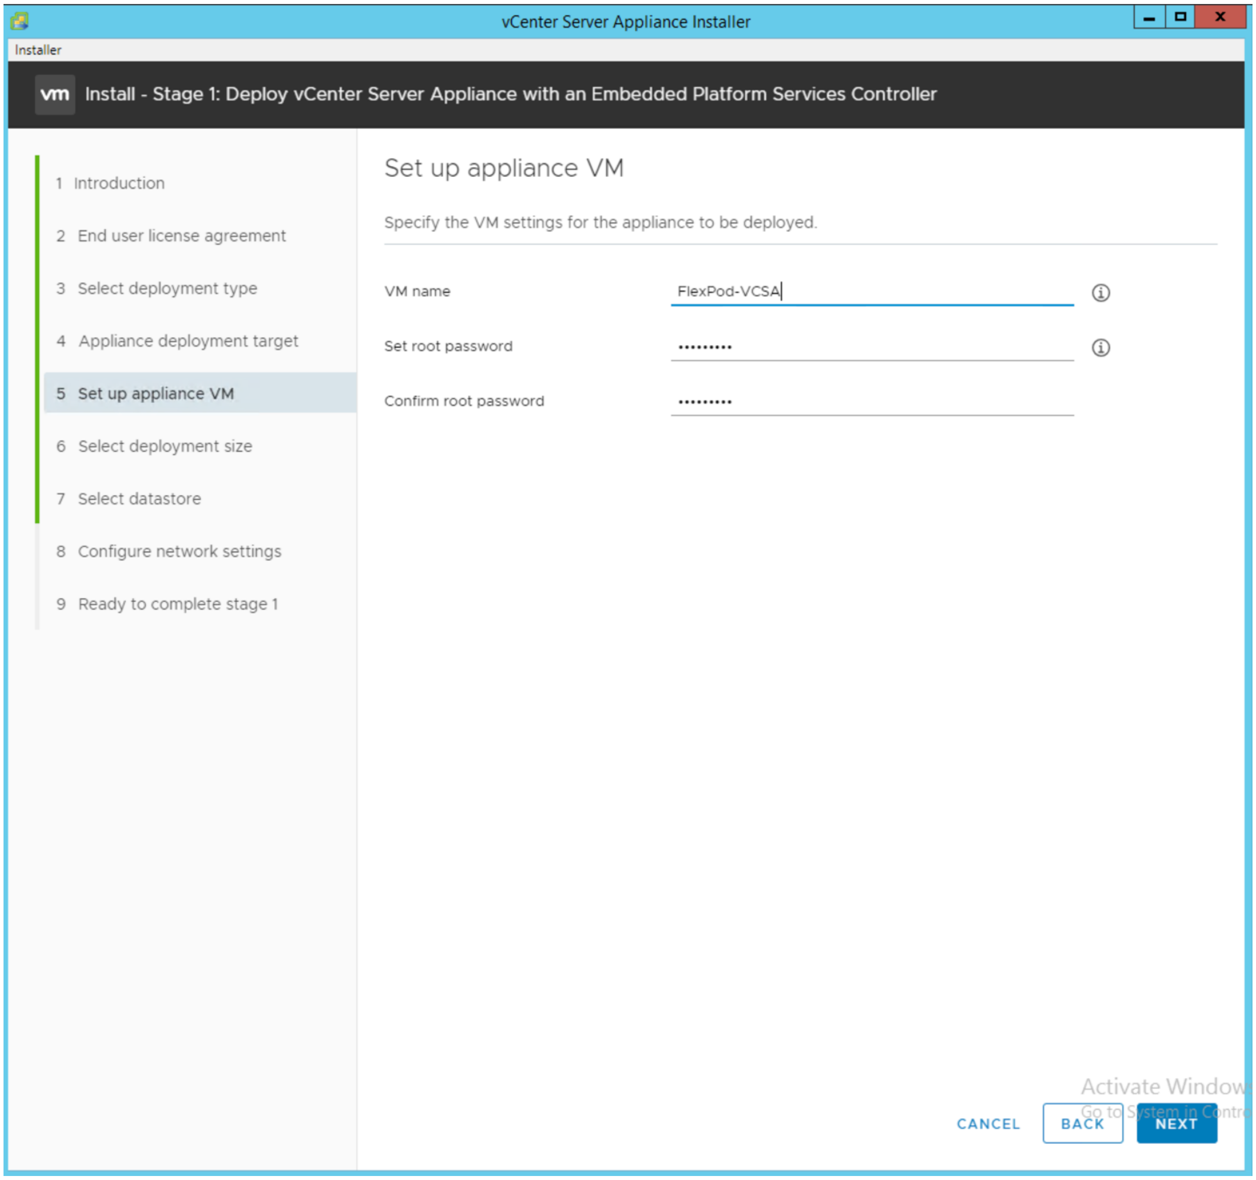
vCenter Server Appliance Installer
Installer
vm
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Activate Window
Go to System in Contr
CANCEL
BACK
NEXT

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name	FlexPod-VCSA
Set root password	•••••••••
Confirm root password	•••••••••

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

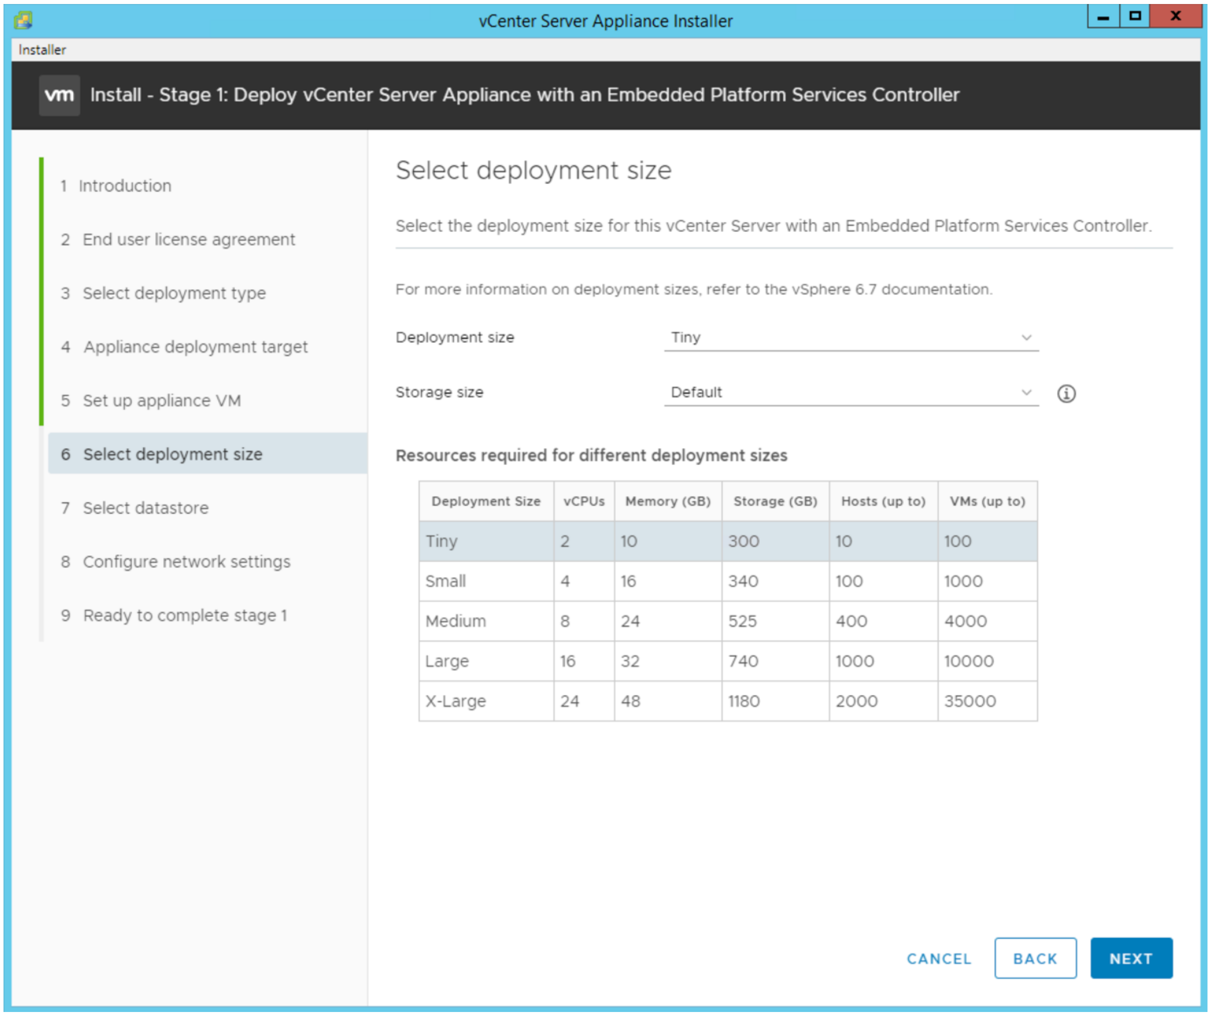
vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
CANCEL
BACK
NEXT

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

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.

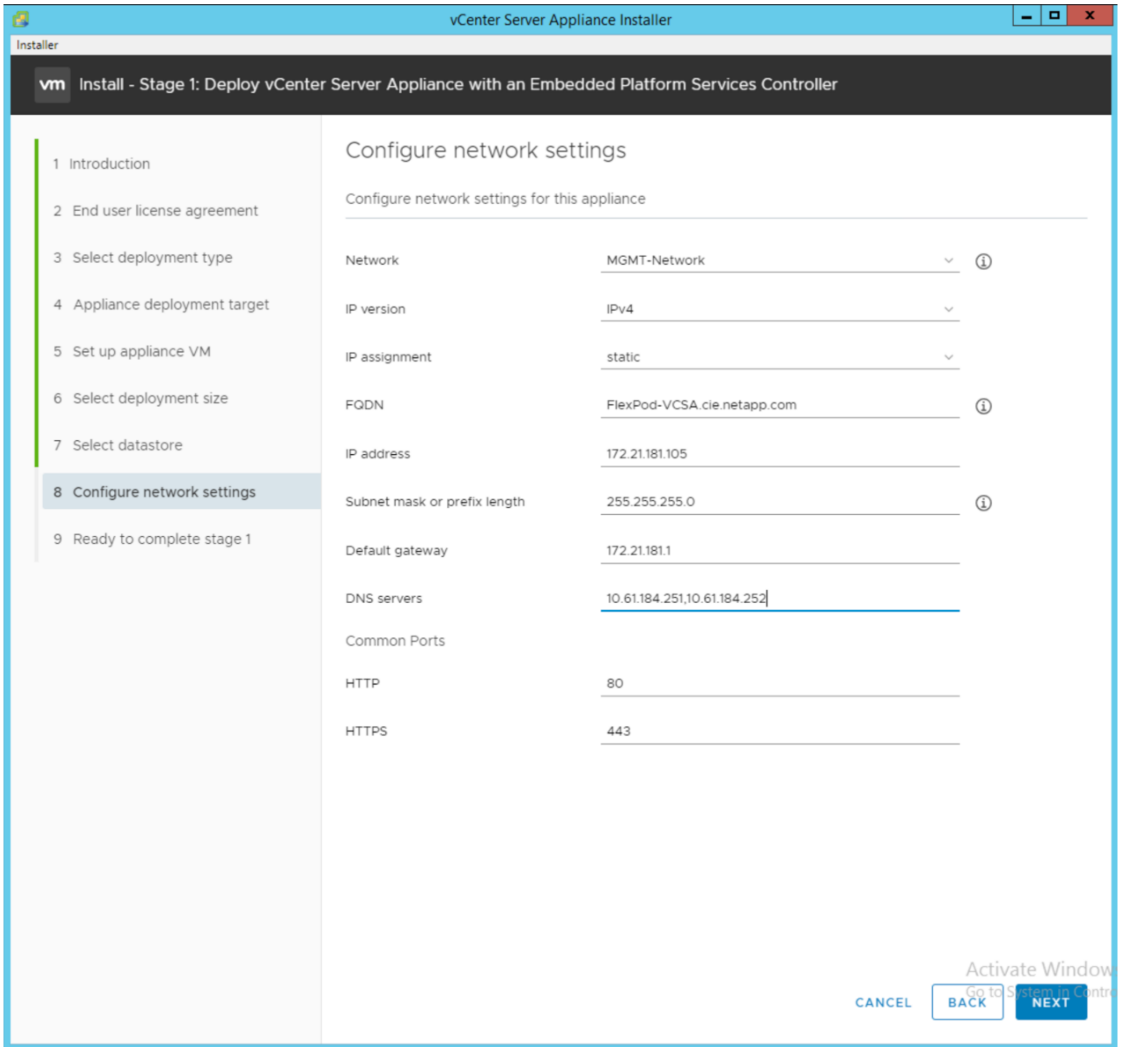
vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
CANCEL
BACK
NEXT

Configure network settings

Configure network settings for this appliance

Network	MGMT-Network
IP version	IPv4
IP assignment	static
FQDN	FlexPod-VCSA.cie.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

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

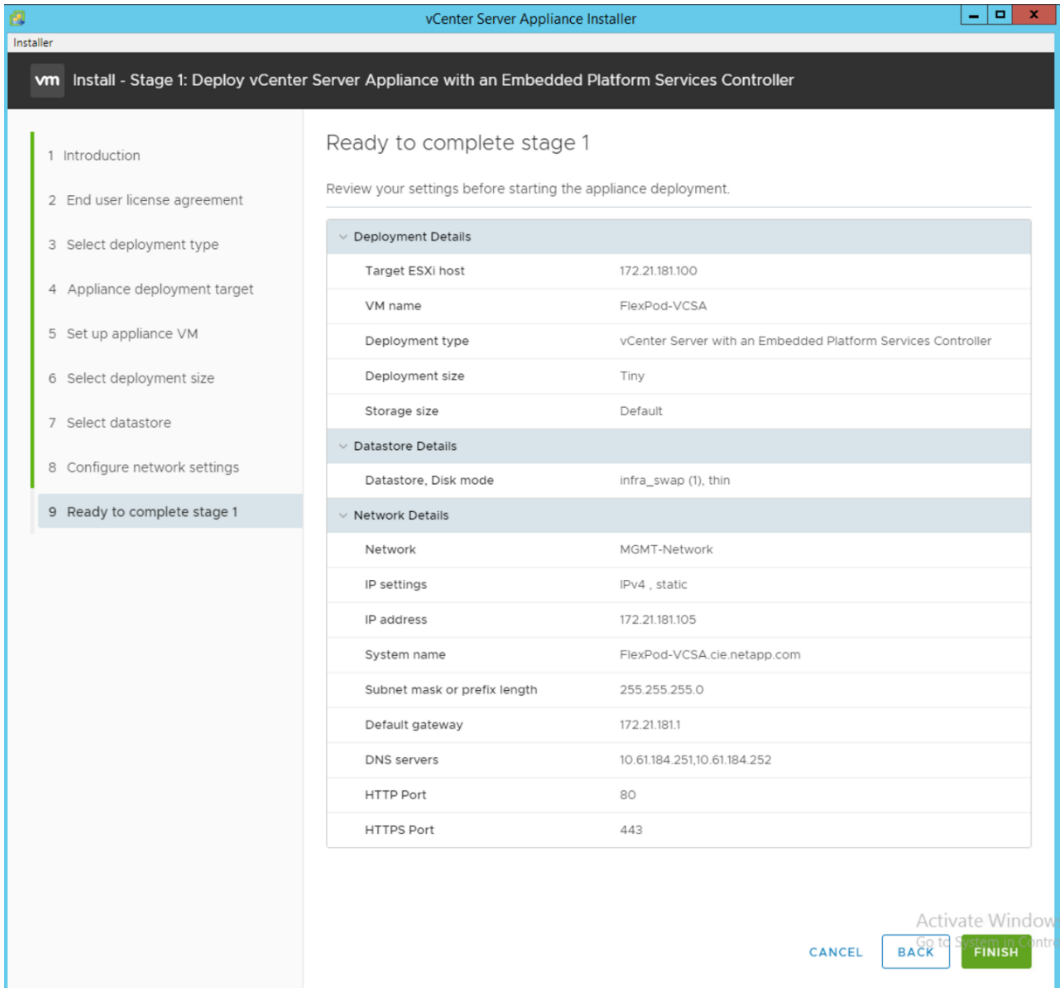

vCenter Server Appliance Installer
Installer
vm Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
CANCEL
BACK
FINISH
Activate Windows
Go to System in Control

Ready to complete stage 1

Review your settings before starting the appliance deployment.

Deployment Details	
Target ESXi host	172.21.181.100
VM name	FlexPod-VCSA
Deployment type	vCenter Server with an Embedded Platform Services Controller
Deployment size	Tiny
Storage size	Default
Datastore Details	
Datastore, Disk mode	infra_swap (1), thin
Network Details	
Network	MGMT-Network
IP settings	IPv4 , static
IP address	172.21.181.105
System name	FlexPod-VCSA.cie.netapp.com
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
HTTP Port	80
HTTPS Port	443

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.

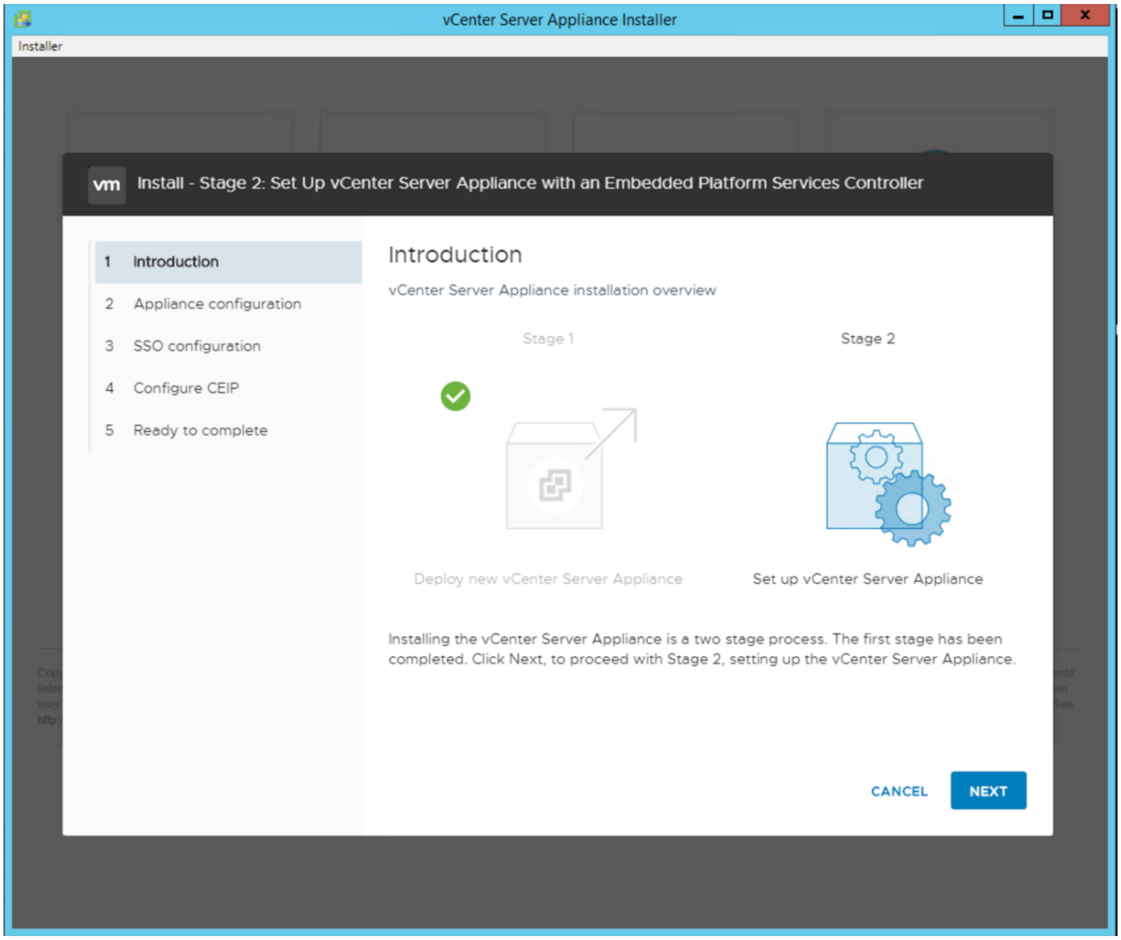

vCenter Server Appliance Installer

Installer

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded Platform Services Controller

1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete

Introduction

vCenter Server Appliance installation overview

Stage 1
Stage 2
Deploy new vCenter Server Appliance
Set up vCenter Server Appliance

Installing the vCenter Server Appliance is a two stage process. The first stage has been completed. Click Next, to proceed with Stage 2, setting up the vCenter Server Appliance.

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.

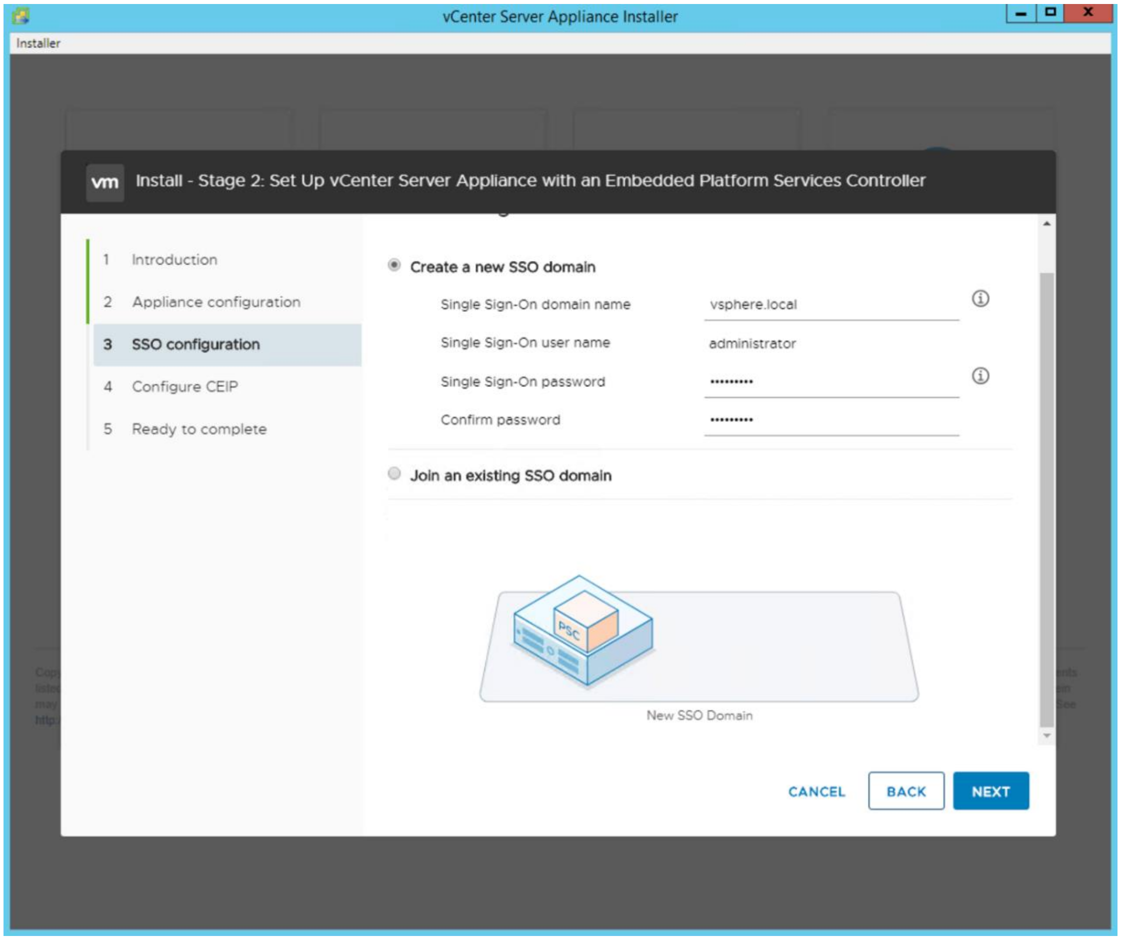

vCenter Server Appliance Installer

Installer

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded Platform Services Controller

● Create a new SSO domain

Single Sign-On domain name: vsphere.local

Single Sign-On user name: administrator

Single Sign-On password: •••••••••

Confirm password: •••••••••

○ Join an existing SSO domain

Note: 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.

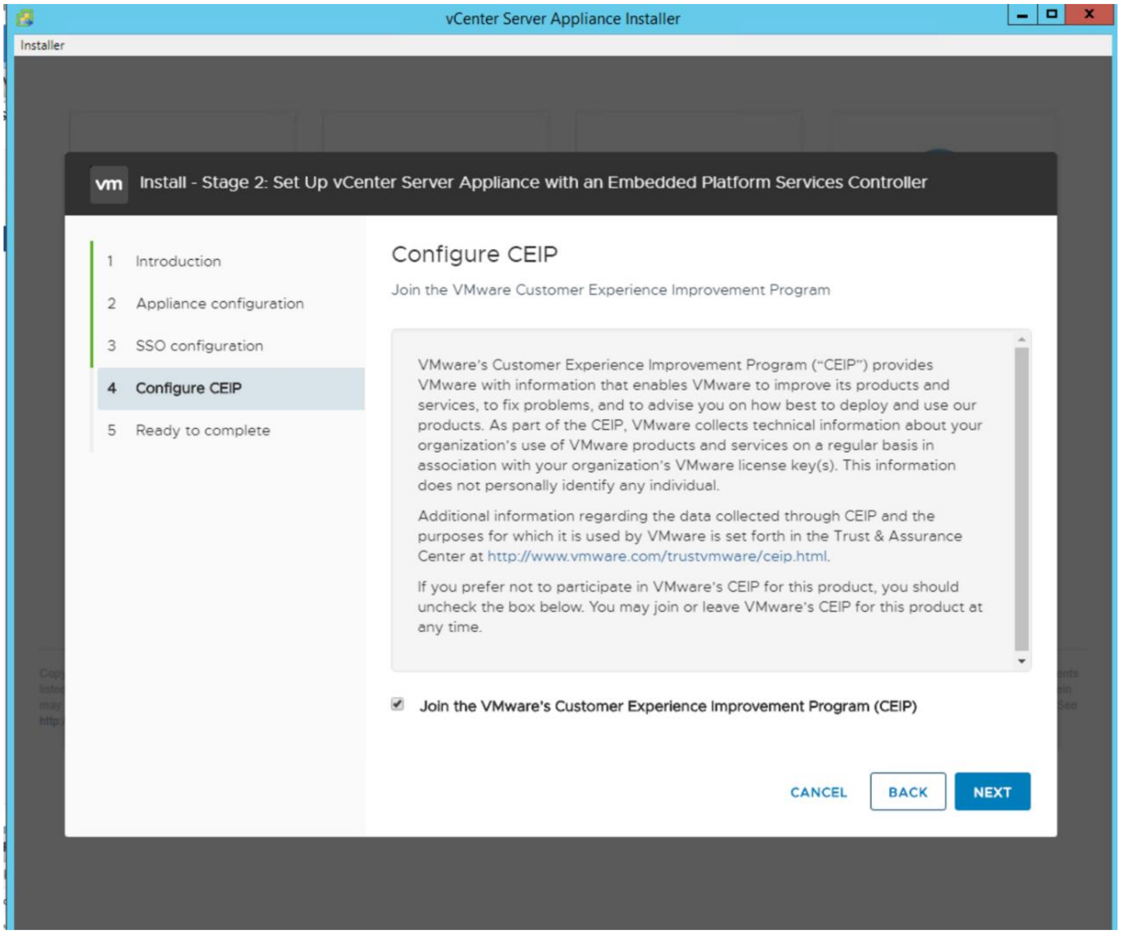

vCenter Server Appliance Installer
Installer
vm Install - Stage 2: Set Up vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
CANCEL
BACK
NEXT

Configure CEIP

Join the VMware Customer Experience Improvement Program

VMware's Customer Experience Improvement Program ("CEIP") provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects technical information about your organization's use of VMware products and services on a regular basis in association with your organization's VMware license key(s). This information does not personally identify any individual.

Additional information regarding the data collected through CEIP and the purposes for which it is used by VMware is set forth in the Trust & Assurance Center at http://www.vmware.com/trustvmware/ceip.html.

If you prefer not to participate in VMware's CEIP for this product, you should uncheck the box below. You may join or leave VMware's CEIP for this product at any time.

☑ Join the VMware's Customer Experience Improvement Program (CEIP)

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.

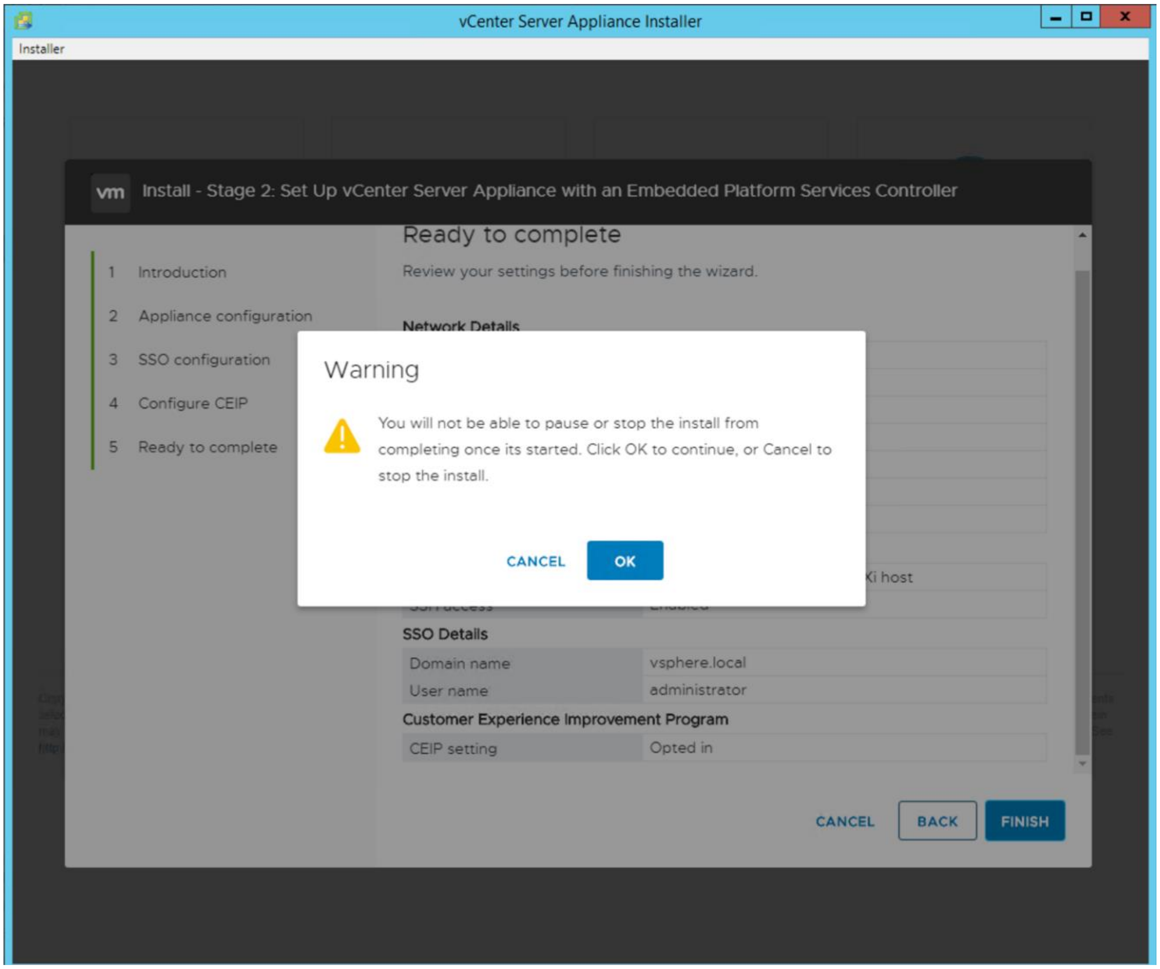
vCenter Server Appliance Installer
Installer
vm Install - Stage 2: Set Up vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
CANCEL
OK
CANCEL
BACK
FINISH

Ready to complete

Review your settings before finishing the wizard.

Network Details

Warning

You will not be able to pause or stop the install from completing once its started. Click OK to continue, or Cancel to stop the install.

SSO Details

Domain name	vsphere.local
User name	administrator

Customer Experience Improvement Program

CEIP setting	Opted in

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.

5.7 Configure VMware vCenter Server 6.7U2 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.

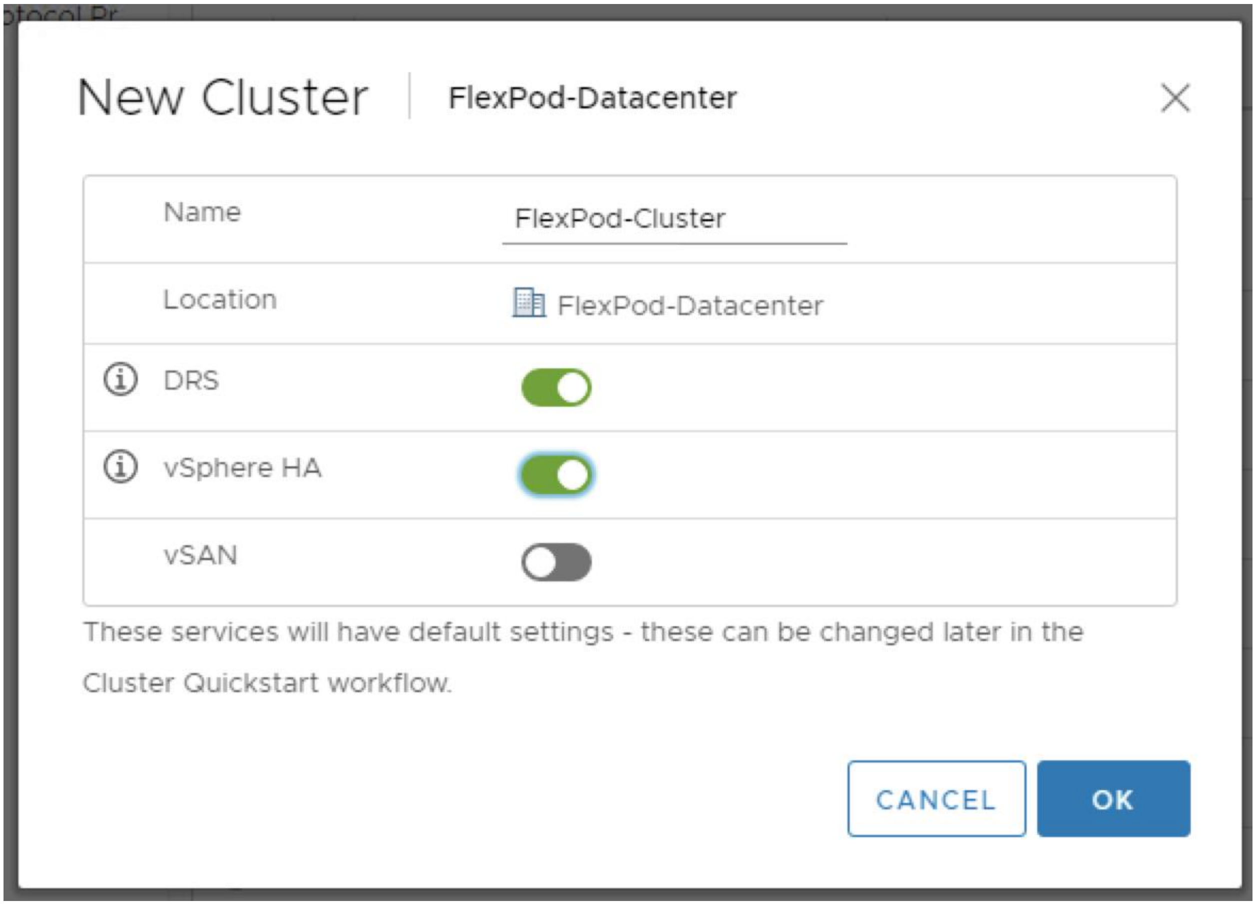

New Cluster | FlexPod-Datacenter

Name	FlexPod-Cluster
Location	FlexPod-Datacenter
DRS	● On
vSphere HA	● On
vSAN	○ Off

These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

Add ESXi Hosts to Cluster

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

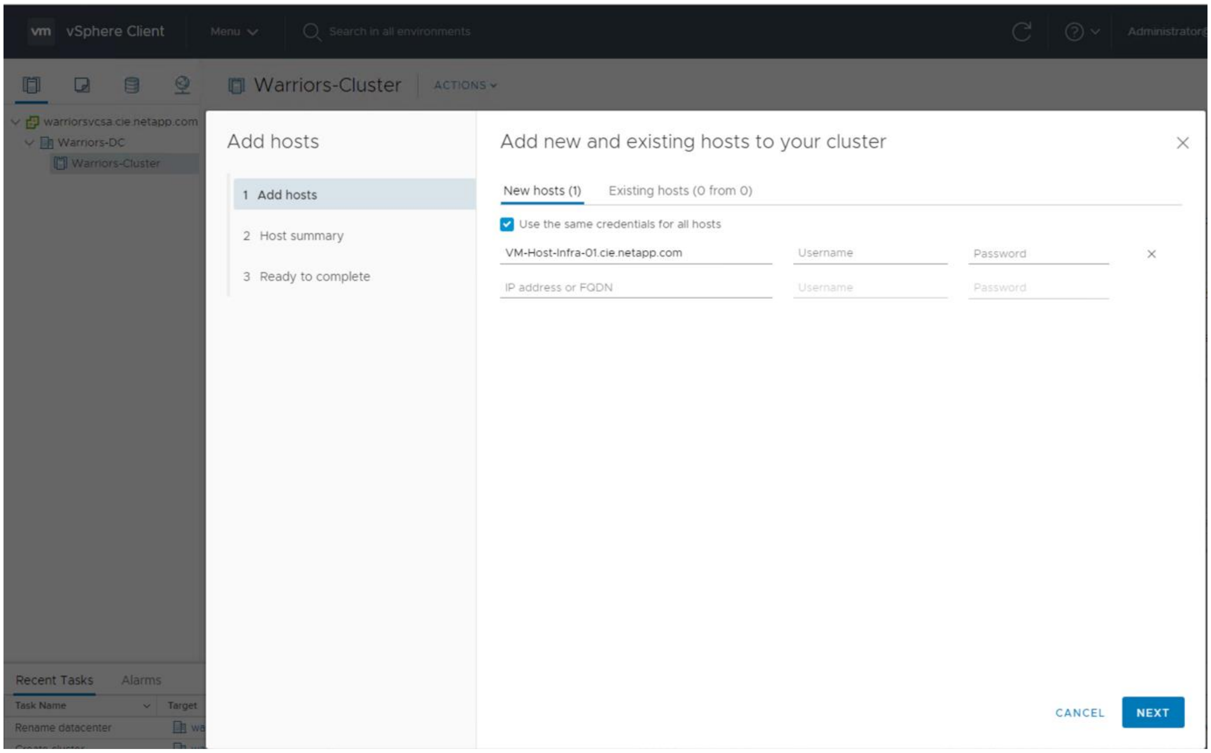

☑ Use the same credentials for all hosts

VM-Host-Infra-01.cie.netapp.com

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.

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

Configure Coredump on ESXi Hosts

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.

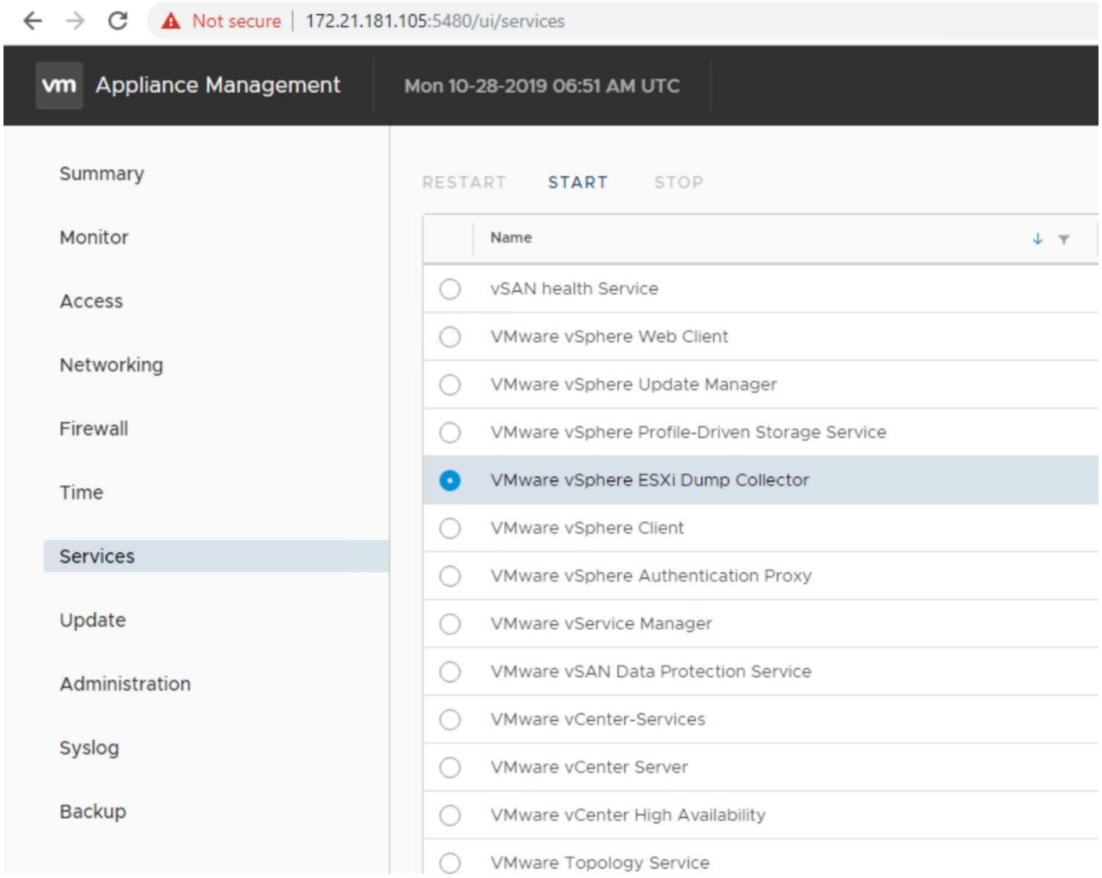

	Name
○	vSAN health Service
○	VMware vSphere Web Client
○	VMware vSphere Update Manager
○	VMware vSphere Profile-Driven Storage Service
●	VMware vSphere ESXi Dump Collector
○	VMware vSphere Client
○	VMware vSphere Authentication Proxy
○	VMware vService Manager
○	VMware vSAN Data Protection Service
○	VMware vCenter-Services
○	VMware vCenter Server
○	VMware vCenter High Availability
○	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.

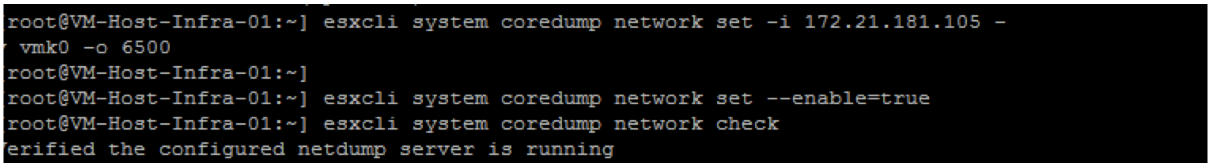

```
root@VM-Host-Infra-01:~] esxcli system coredump network set -i 172.21.181.105 -
 vmk0 -o 6500
root@VM-Host-Infra-01:~]
root@VM-Host-Infra-01:~] esxcli system coredump network set --enable=true
root@VM-Host-Infra-01:~] esxcli system coredump network check
erified the configured netdump server is running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

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

Note: `ip_address_of_core_dump_collector` in this validation is the vCenter IP.

NetApp Virtual Storage Console 9.6 Deployment Procedure

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.

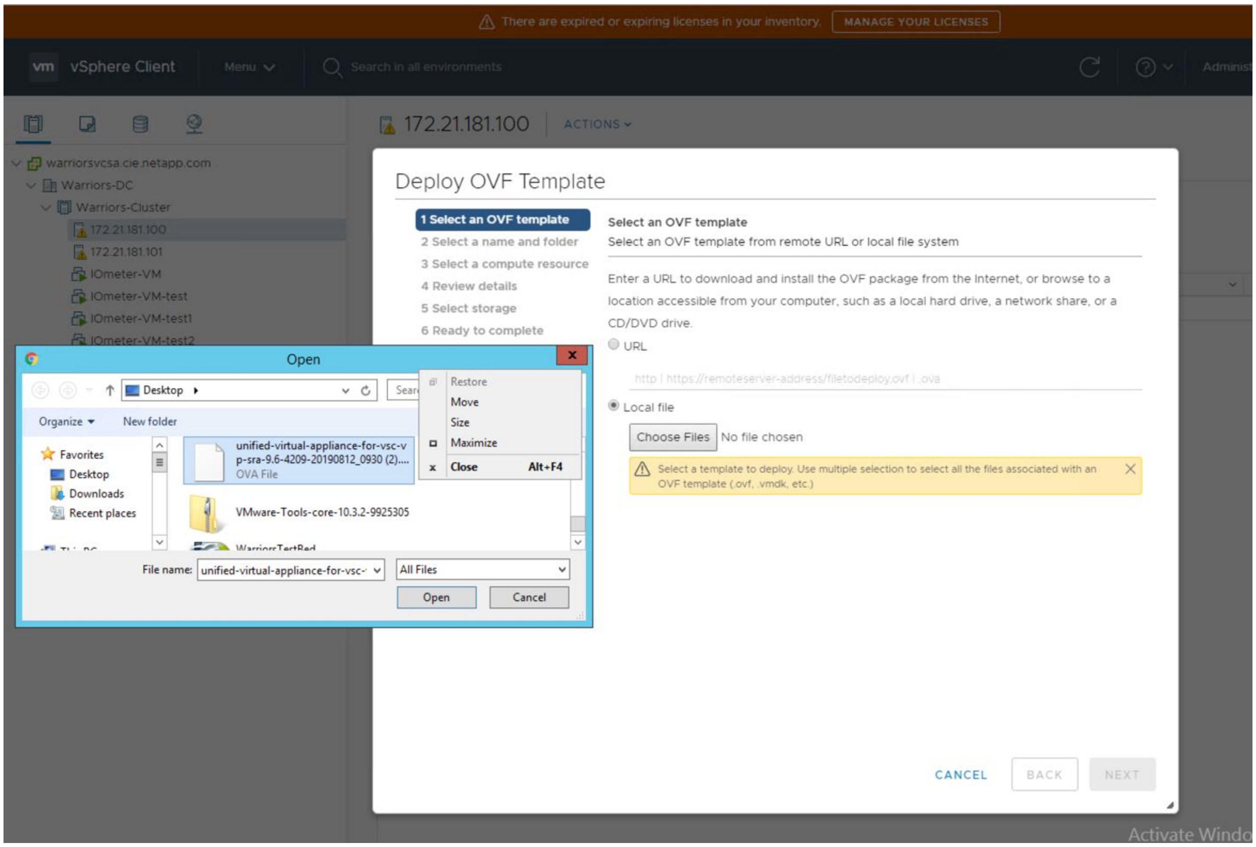

There are expired or expiring licenses in your inventory. MANAGE YOUR LICENSES

Deploy OVF Template

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

● Local file

Choose Files No file chosen

Select a template to deploy. Use multiple selection to select all the files associated with an OVF template (.ovf, .vmdk, etc.)

File name: unified-virtual-appliance-for-vsc-v

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

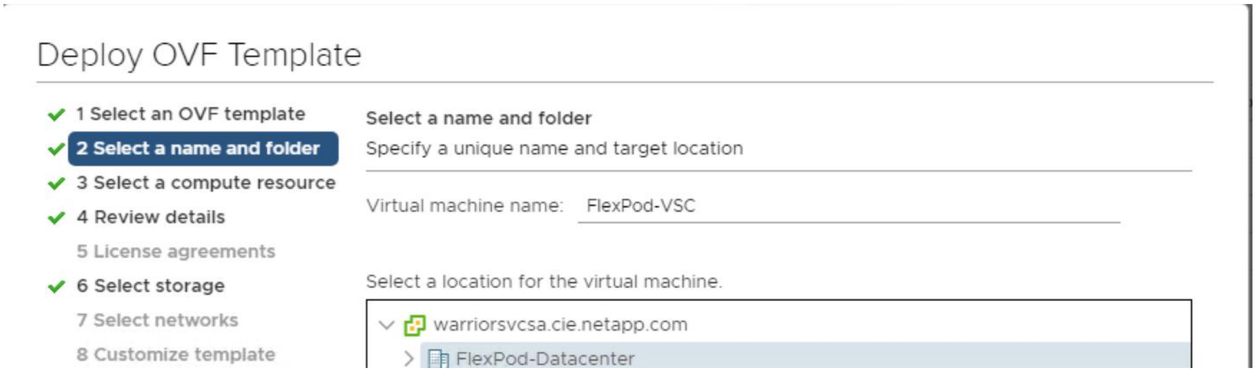

Deploy OVF Template

Select a name and folder

Specify a unique name and target location

Virtual machine name: FlexPod-VSC

Select a location for the virtual machine.

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

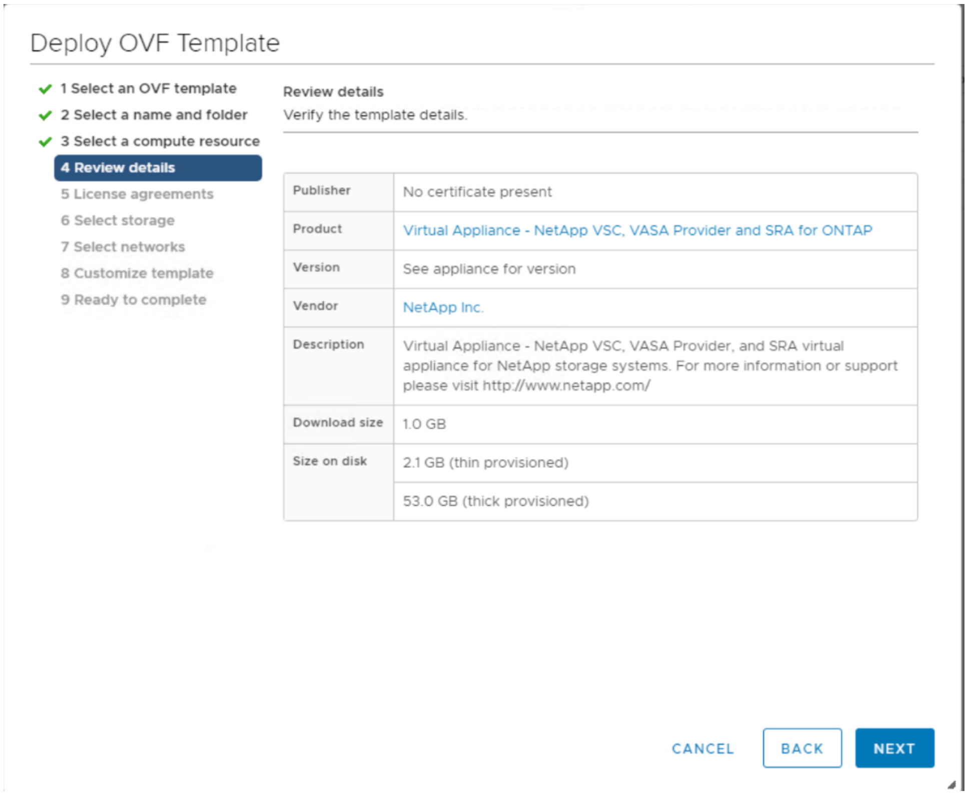

Deploy OVF Template

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)

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.

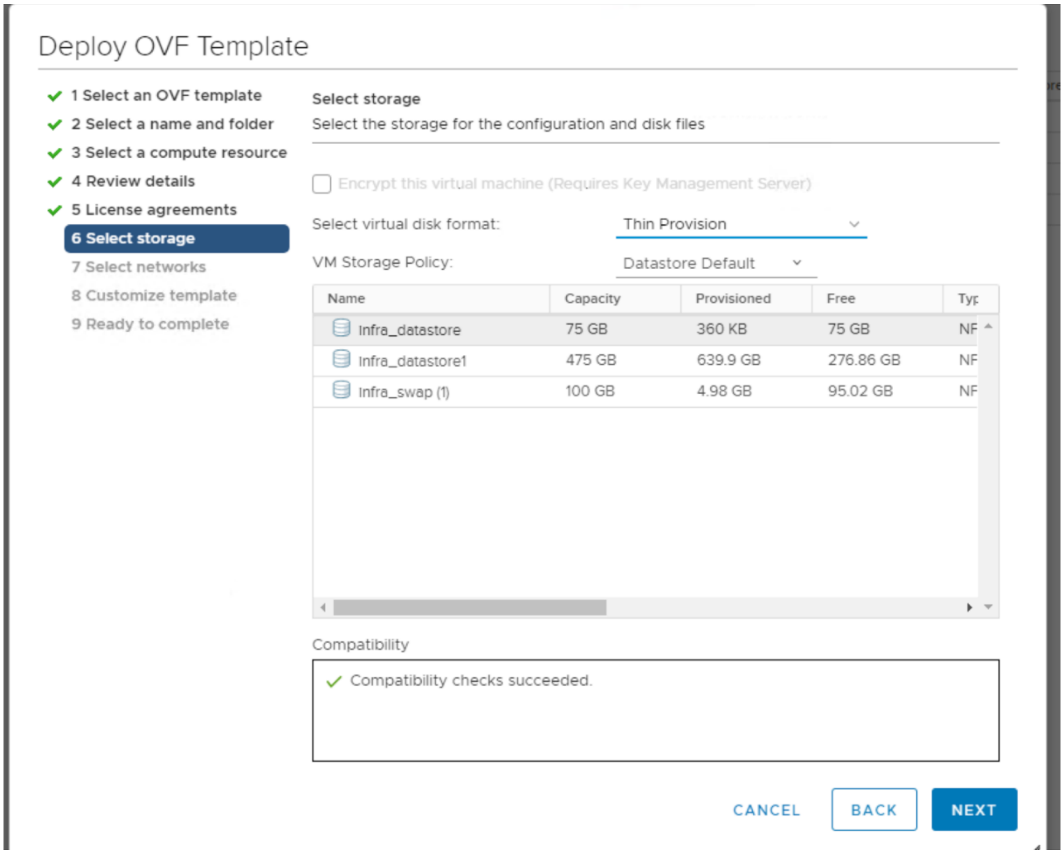

Select storage

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	Typ
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.

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

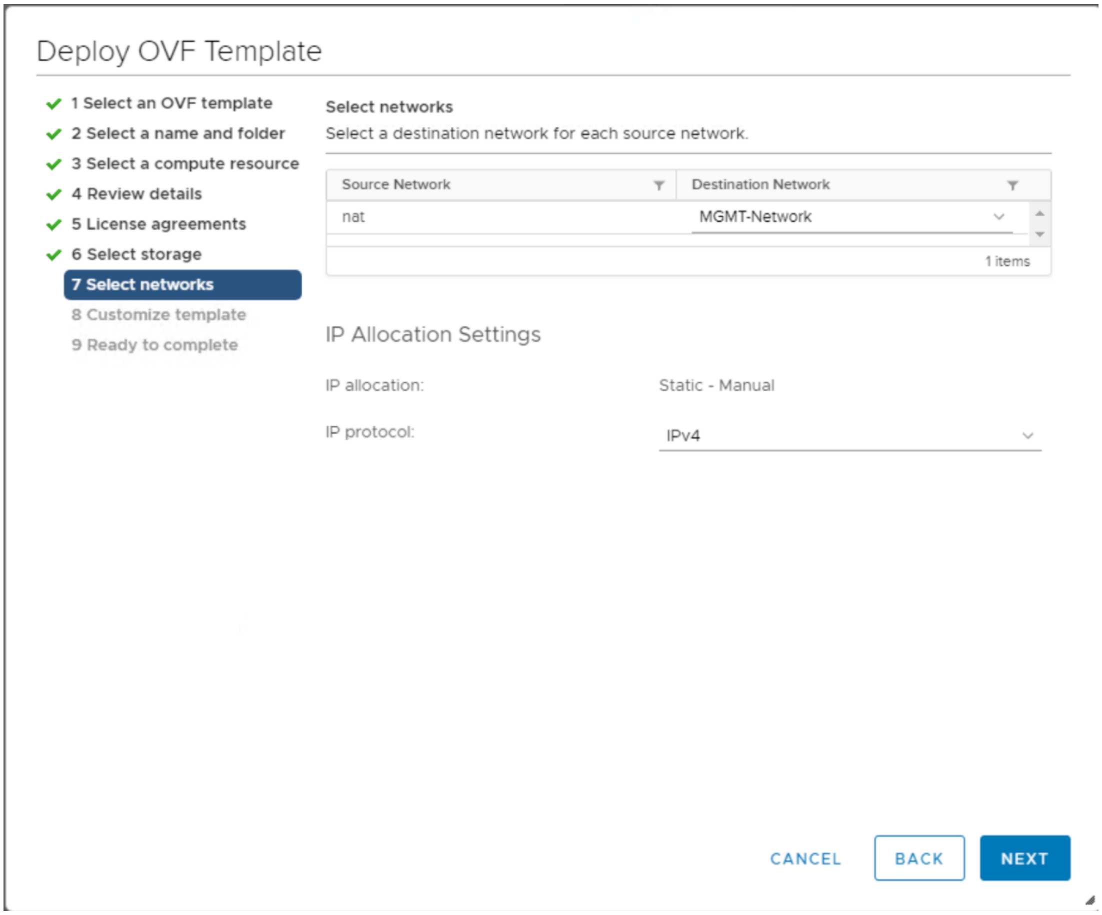

Deploy OVF Template

- ✓ 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

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

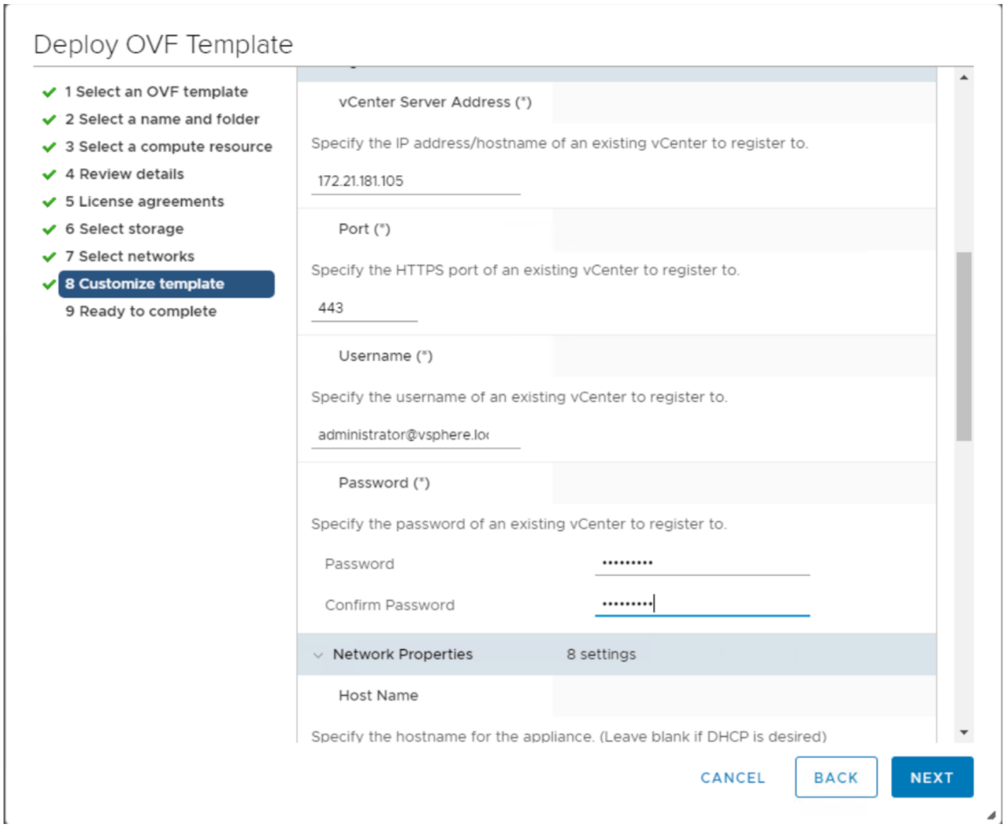

Deploy OVF Template

- ✓ 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.lo

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)

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.

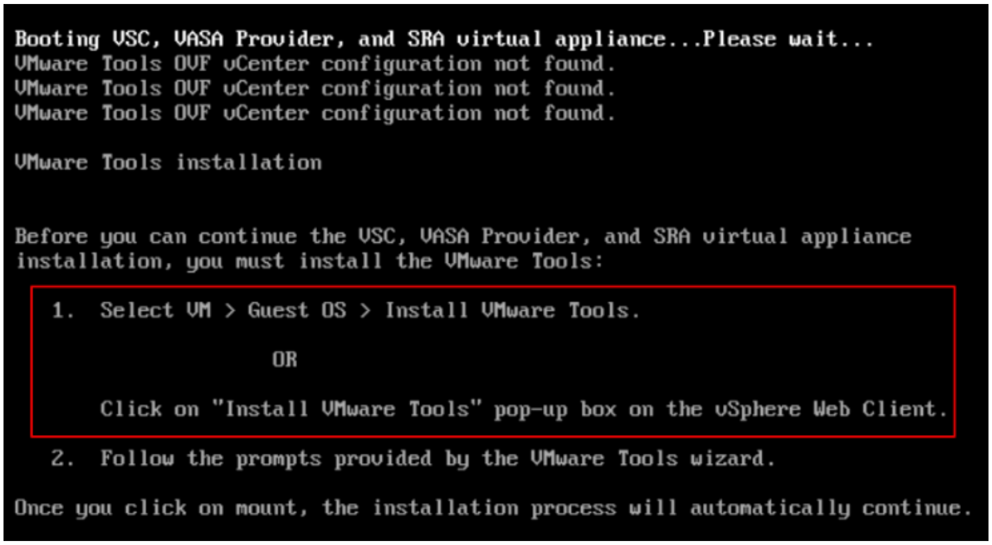

```
Booting VSC, VASA Provider, and SRA virtual appliance...Please wait...
VMware Tools OVF vCenter configuration not found.
VMware Tools OVF vCenter configuration not found.
VMware Tools OVF vCenter configuration not found.

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.

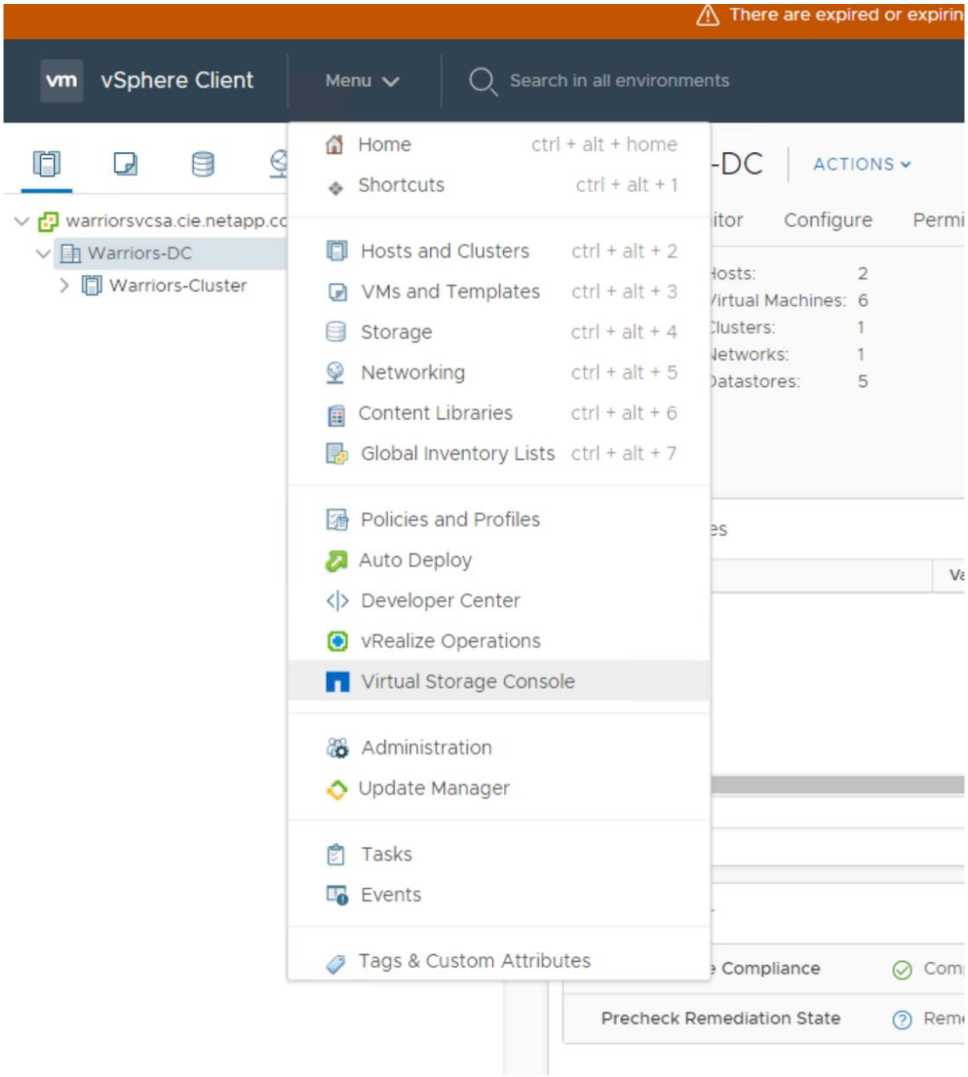

Download the NetApp NFS VAAI Plug-In

To 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.

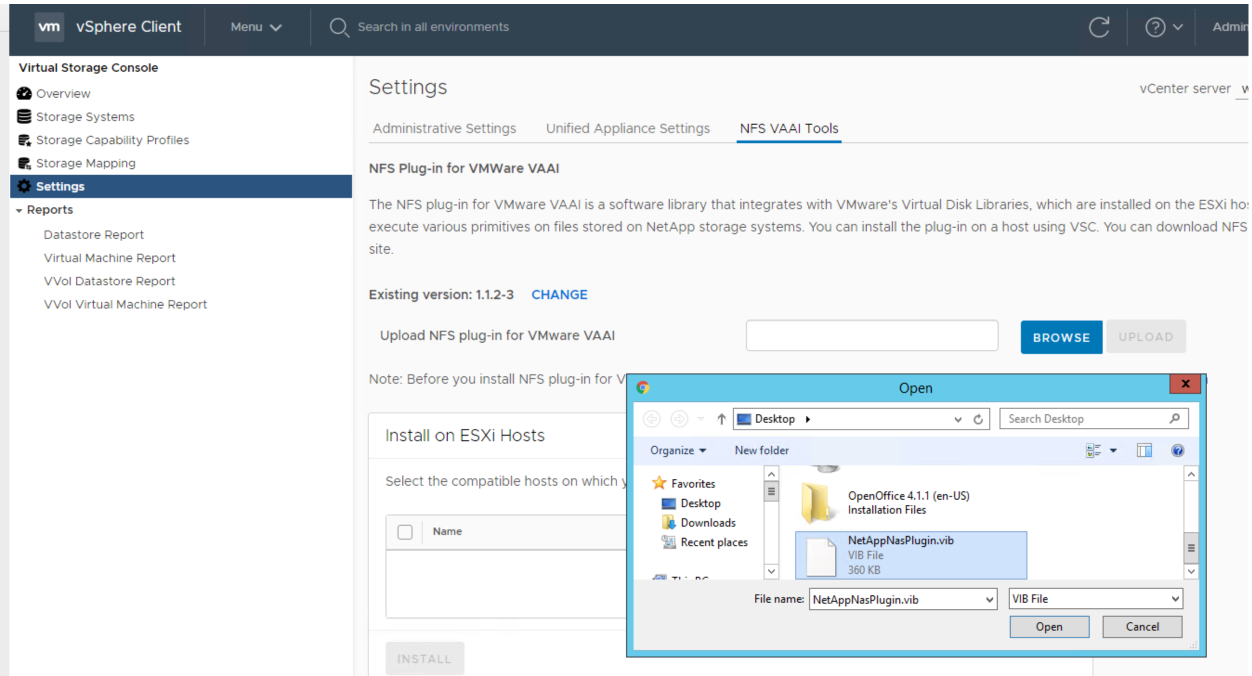

vm vSphere Client | Menu | Search in all environments | Admin

Virtual Storage Console

- Overview
- Storage Systems
- Storage Capability Profiles
- Storage Mapping
- Settings
- Reports
  - Datastore Report
  - Virtual Machine Report
  - VVol Datastore Report
  - VVol Virtual Machine Report

Settings — vCenter server

Administrative Settings | Unified Appliance Settings | NFS VAAI Tools

NFS Plug-in for VMWare VAAI

The NFS plug-in for VMware VAAI is a software library that integrates with VMware's Virtual Disk Libraries, which are installed on the ESXi ho execute various primitives on files stored on NetApp storage systems. You can install the plug-in on a host using VSC. You can download NFS site.

Existing version: 1.1.2-3 CHANGE

Upload NFS plug-in for VMware VAAI — BROWSE | UPLOAD

Note: Before you install NFS plug-in for V

Install on ESXi Hosts

Select the compatible hosts on which y

□ Name

INSTALL

Open

Desktop ▸ | Search Desktop

Organize ▾ | New folder

Favorites: Desktop, Downloads, Recent places

OpenOffice 4.1.1 (en-US) Installation Files

NetAppNasPlugin.vib
VIB File
360 KB

File name: NetAppNasPlugin.vib | VIB File

Open | Cancel

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.

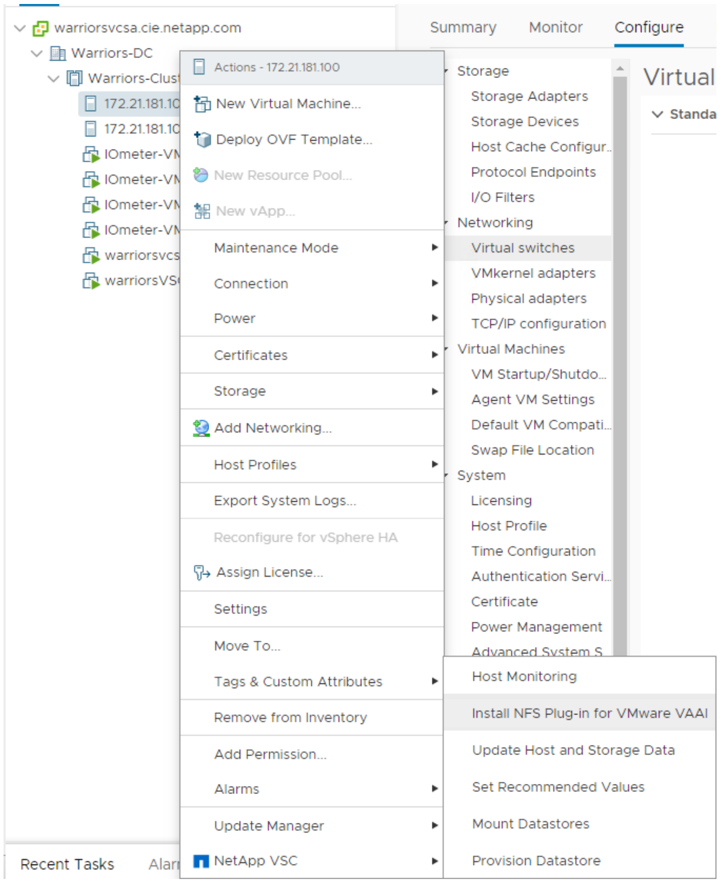

- warriorsvcsa.cie.netapp.com
  - Warriors-DC
    - Warriors-Clust
      - 172.21.181.10
      - 172.21.181.10
      - IOmeter-VI
      - IOmeter-VI
      - IOmeter-VI
      - IOmeter-VI
      - warriorsvcs
      - warriorsVS

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 ▸

Summary | Monitor | Configure

- Storage: Storage Adapters, Storage Devices, Host Cache Configur..., Protocol Endpoints, I/O Filters
- Networking: Virtual switches, VMkernel adapters, Physical adapters, TCP/IP configuration
- Virtual Machines: VM Startup/Shutdo..., Agent VM Settings, Default VM Compati..., Swap File Location
- System: Licensing, Host Profile, Time Configuration, Authentication Servi..., Certificate, Power Management, Advanced System S

Virtual — Standa

NetApp VSC submenu:

- Host Monitoring
- Install NFS Plug-in for VMware VAAI
- Update Host and Storage Data
- Set Recommended Values
- Mount Datastores
- Provision Datastore

Recent Tasks | Alar

Optimal Storage Settings for 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, follow these steps:

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

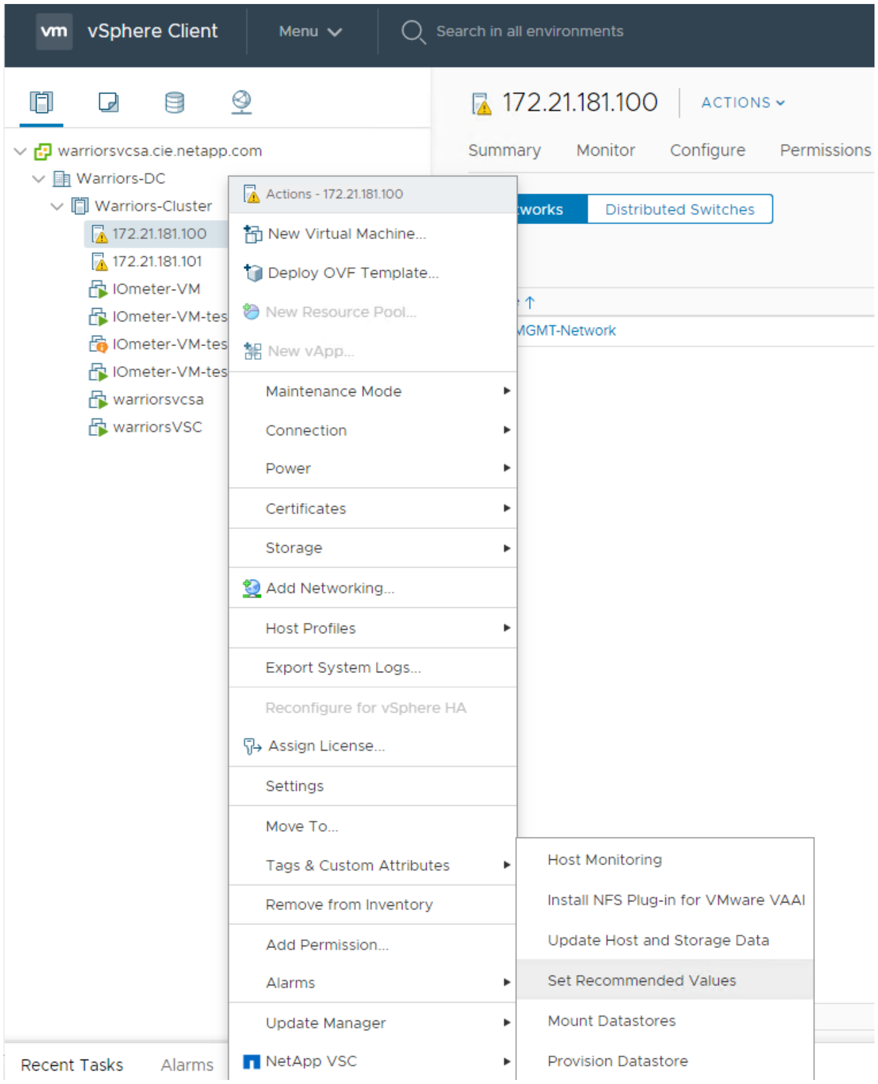

2. Check the settings that you would like to apply to the selected vSphere hosts. Click OK to apply the settings.

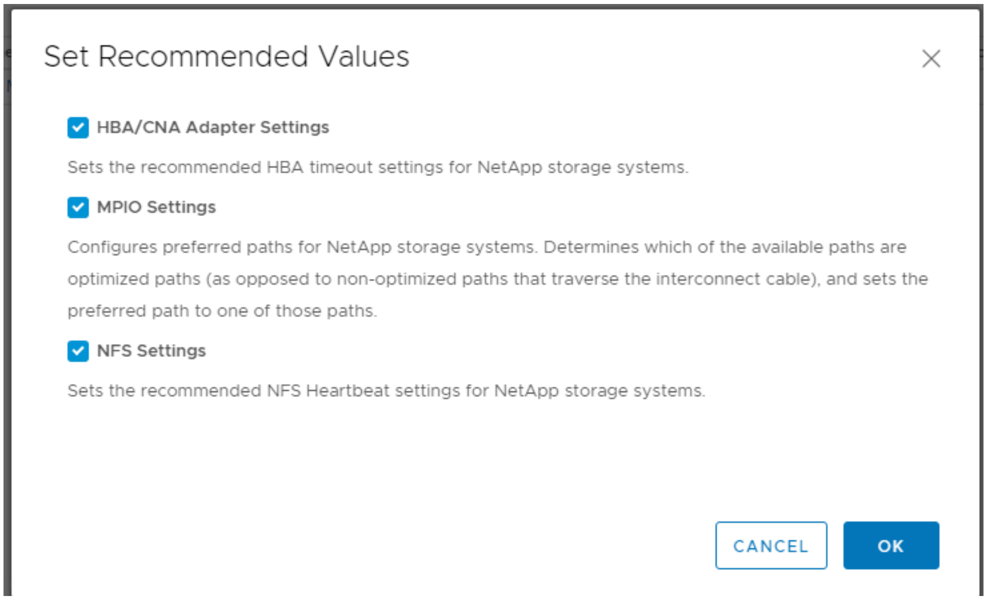

Set Recommended Values

☑ **HBA/CNA Adapter Settings**

Sets the recommended HBA timeout settings for NetApp storage systems.

☑ **MPIO Settings**

Configures preferred paths for NetApp storage systems. Determines which of the available paths are optimized paths (as opposed to non-optimized paths that traverse the interconnect cable), and sets the preferred path to one of those paths.

☑ **NFS Settings**

Sets the recommended NFS Heartbeat settings for NetApp storage systems.

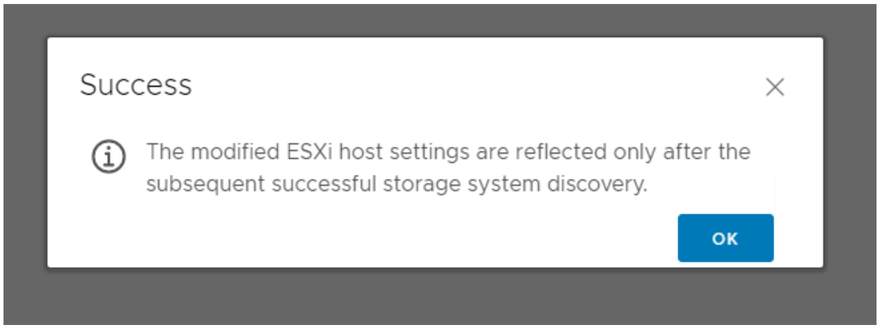

Success

The modified ESXi host settings are reflected only after the subsequent successful storage system discovery.

Note: Reboot the ESXI host after these settings are applied.

6 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.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Trademark Information