



NetApp Verified Architecture

## FlexPod Express with VMware vSphere 6.5 and FAS2600

NVA Deployment

Melissa Palmer, Lindsey Street, NetApp  
April 2017 | NVA-1112-DEPLOY | Version 1.0

Reviewed by Cisco  
Systems, Inc.



## TABLE OF CONTENTS

<b>1 Program Summary</b>	<b>4</b>
<b>2 Solution Overview</b>	<b>4</b>
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
<b>3 Technology Requirements</b>	<b>6</b>
3.1 Hardware Requirements	7
3.2 Software Requirements	7
<b>4 FlexPod Express Cabling Information</b>	<b>8</b>
<b>5 Deployment Procedures</b>	<b>10</b>
5.1 Cisco Nexus 31108 Deployment Procedure	11
5.2 NetApp Storage Deployment Procedure (Part 1)	18
5.3 Continuation Node A Configuration and Cluster Configuration	22
5.4 Cisco UCS C-Series Rack Server Deployment Procedure	39
5.5 NetApp FAS Storage Deployment Procedure (Part 2)	50
5.6 VMware vSphere 6.5 Deployment Procedure	51
5.7 Install VMware vCenter Server 6.5	68
5.8 Configure VMware vCenter Server 6.5 and vSphere Clustering	77
5.9 NetApp Virtual Storage Console 6.2.1P1 Deployment Procedure	85
<b>Conclusion</b>	<b>89</b>
<b>About the Authors</b>	<b>90</b>
<b>Acknowledgements</b>	<b>90</b>
<b>Version History</b>	<b>90</b>

## LIST OF TABLES

Table 1) Hardware requirements for the base configuration	7
Table 2) Hardware for scaling the solution by using two hypervisor nodes	7
Table 3) Software requirements for the base FlexPod Express implementation	7
Table 4) Software requirements for a VMware vSphere implementation	7
Table 5) Cabling information for Cisco Nexus switch 31108 A	8
Table 6) Cabling information for Cisco Nexus switch 31108 B	9

Table 7) Cabling information for NetApp FAS2650 storage controller A. ....	9
Table 8) Cabling information for NetApp FAS2650 storage controller B. ....	9
Table 9) Required VLANs.....	10
Table 10) VMware virtual machines created. ....	10

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

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

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

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

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

## 2 Solution Overview

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

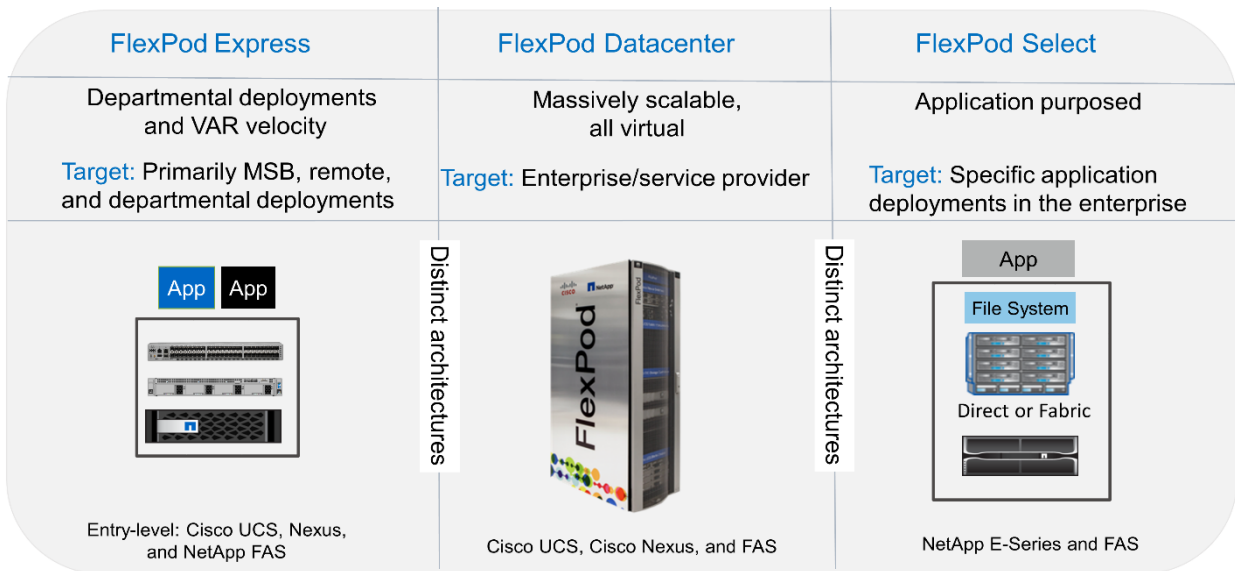
### 2.1 FlexPod Converged Infrastructure Program

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 three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

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

Figure 1) FlexPod portfolio.



## 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:


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

This guide details the design of FlexPod Express with VMware vSphere. In addition, this design uses the all-new FAS2650 system, which runs NetApp ONTAP® 9.1; the Cisco Nexus 31108; and Cisco UCS C-Series C220 M4 servers as hypervisor nodes.

## 2.3 Solution Technology

This solution leverages the latest technologies from NetApp, Cisco, VMware, and Microsoft. This solution features the new NetApp FAS2650 running ONTAP 9.1, dual Cisco Nexus 31108 switches, and Cisco UCS C220 M4 rack servers that run VMware vSphere 6.5. This validated solution uses 10 Gigabit Ethernet (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.



## 2.4 Use Case Summary

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

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

FlexPod Express is best suited for virtualization and mixed workloads.

## 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, either hypervisor can run on the same FlexPod Express hardware.

Table 1 lists the hardware components that are required for all FlexPod Express configurations.

Table 1) Hardware requirements for the base configuration.

Hardware	Quantity
FAS2650 two-node cluster	1
Cisco C220 M4 server	2
Cisco Nexus 31108 switch	2
Cisco UCS virtual interface card (VIC) 1227 for the C220 M4 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 using two hypervisor nodes.

Hardware	Quantity
Cisco UCS C220 M4 server	2
Cisco VIC 1227	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)	2.0(13h)	For Cisco UCS C220 M4 rack servers
Cisco enic driver	1.1.1.0	For VIC 1227 interface cards
Cisco NX-OS	7.0(3)I5(2)	For Cisco Nexus 31108 switches
NetApp ONTAP	9.1	For FAS2650 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	6.5

Software	Version
VMware vSphere ESXi hypervisor	6.5
NetApp Virtual Storage Console (VSC)	6.2.1P1
NetApp VAAI Plug-In for ESXi	1.1.2

## 4 FlexPod Express Cabling Information

The reference validation documented in this document 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 31108 A.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus switch 31108 A	Eth1/1	NetApp FAS2650 storage controller A	e0c
	Eth1/2	NetApp FAS2650 storage controller B	e0c
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM1
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM1



Local Device	Local Port	Remote Device	Remote Port
	Eth1/25	Cisco Nexus switch 31108 B	Eth1/25
	Eth1/26	Cisco Nexus switch 31108 B	Eth1/26
	Eth1/33	NetApp FAS2650 storage controller A	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server A	CIMC

**Table 6) Cabling information for Cisco Nexus switch 31108 B.**

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus switch 31108 B	Eth1/1	NetApp FAS2650 storage controller A	e0d
	Eth1/2	NetApp FAS2650 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
	Eth1/25	Cisco Nexus switch 31108 A	Eth1/25
	Eth1/26	Cisco Nexus switch 31108 A	Eth1/26
	Eth1/33	NetApp FAS2650 storage controller B	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server B	CIMC

**Table 7) Cabling information for NetApp FAS2650 storage controller A.**

Local Device	Local Port	Remote Device	Remote Port
NetApp FAS2650 storage controller A	e0a	NetApp FAS2650 storage controller B	e0a
	e0b	NetApp FAS2650 storage controller B	e0b
	e0c	Cisco Nexus switch 31108 A	Eth1/1
	e0d	Cisco Nexus switch 31108 B	Eth1/1
	e0M	Cisco Nexus switch 31108 A	Eth1/33

**Table 8) Cabling information for NetApp FAS2650 storage controller B.**

Local Device	Local Port	Remote Device	Remote Port
NetApp FAS2650 storage controller B	e0a	NetApp FAS2650 storage controller A	e0a
	e0b	NetApp FAS2650 storage controller A	e0b
	e0c	Cisco Nexus switch 31108 A	Eth1/2
	e0d	Cisco Nexus switch 31108 B	Eth1/2
	e0M	Cisco Nexus switch 31108 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>vlan create vif0 <<mgmt_vlan_id>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. 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.

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

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

Table 10 lists the VMware virtual machines created.

Table 10) VMware virtual machines created.

Virtual Machine Description	Host Name
VMware vCenter Server	
NetApp Virtual Storage Console	

## 5.1 Cisco Nexus 31108 Deployment Procedure

The following section details the Cisco Nexus 31108 switch configuration used in a FlexPod Express environment.

### Initial Setup of Cisco Nexus 31108 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 31108 running NX-OS software release 7.0(3)I5(2).

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 31108 switches can be connected to an existing management network, or the mgmt0 interfaces of the 31108 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 31108 switches are connected to an existing management network.

3. To configure the Cisco Nexus 31108 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 : 31108-B
```

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

```
Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>
```

```
Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>
```

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

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

```

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

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

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

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

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

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

    NTP server IPv4 address : <<var_ntp_ip>>

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

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

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

```

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

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

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

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

6. Repeat this procedure for Cisco Nexus switch B.

## Enable Advanced Features

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

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

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

```

feature interface-vlan
feature lacp
feature vpc

```

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

2. From configuration mode (`config t`), enter the following commands to set the global PortChannel 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`), type the following commands to configure the default spanning-tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

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

## Define VLANs

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

From configuration mode (`config t`), type the following commands to define and give descriptions to 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 FAS2650-A e0c
int eth1/2
  description FAS2650-B e0c
int eth1/3
  description UCS-Server-A: VIC Port 1
int eth1/4
  description UCS-Server-B: VIC Port 1
int eth1/25
  description vPC peer-link 31108-B 1/25
int eth1/26
  description vPC perr-link 31108-B 1/26
int eth1/33
  description FAS2650-A e0M
int eth1/34
  description UCS Server A: CIMC
```

### Cisco Nexus Switch B

```
int eth1/1
  description FAS2650-A e0d
int eth1/2
  description FAS2650-B e0d
int eth1/3
  description UCS-Server-A: VIC Port 2
int eth1/4
  description UCS-Server-B: VIC Port 2
int eth1/25
  description vPC peer-link 31108-A 1/25
int eth1/26
  description vPC perr-link 31108-A 1/26
int eth1/33
  description FAS2650-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 PortChannel Global Configuration

A virtual PortChannel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single PortChannel 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 PortChannel 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), type the following commands to configure the vPC global configuration for both switches:

### Cisco Nexus Switch A

```
vpc domain 1
 1-switch
  role priority 10
  peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source <<switch_A_mgmt0_ip_addr>> vrf
  management
  peer-gateway
  auto-recovery
  ip arp synchronize

int eth1/25-26
 channel-group 10 mode active
int Po10
 description vPC peer-link
 switchport
 switchport mode trunk
 switchport trunk native vlan <<native_vlan_id>>
 switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan>>,<<iSCSI_A_vlan_id>>,<<iSCSI_B_vlan_id>>
 spanning-tree port type network
 vpc peer-link
 no shut
exit
copy run start
```

## Cisco Nexus Switch B

```
vpc domain 1
  peer-switch
  role priority 20
  peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source <<switch_B_mgmt0_ip_addr>> vrf
management
  peer-gateway
  auto-recovery
  ip arp synchronize

int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>>, <<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
no shut
exit
copy run start
```

## Configure Storage PortChannels

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 PortChannel configuration for the ports connected to the NetApp FAS controller.

Run the following commands on switch A and switch B to configure the PortChannels 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
```

Run the following commands on switch A and switch B to configure the PortChannels 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

```

**Note:** Jumbo frames should be configured throughout the network to enable any applications and operating systems to transmit these larger frames without fragmentation. Both the endpoints and all the interfaces between the endpoints (layer 2 and layer 3) must support and be configured for jumbo frames to prevent performance problems caused by fragmenting frames.

## Configure Server Connections

The Cisco UCS servers have a two-port virtual interface card, VIC1227, 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`), type the following commands to configure the port settings for the interfaces connected to each server.

### Cisco Nexus Switch A: Cisco UCS Server-A and Cisco UCS Server-B Configuration

```

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

```

### Cisco Nexus Switch B: Cisco UCS Server-A and Cisco UCS Server-B Configuration

```

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

```

**Note:** Jumbo frames should be configured throughout the network to enable any applications and operating systems to transmit these larger frames without fragmentation. Both the endpoints and all the interfaces between the endpoints (layer 2 and layer 3) must support and be configured for jumbo frames to prevent performance problems by fragmenting frames.

**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 may be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

## 5.2 NetApp Storage Deployment Procedure (Part 1)

This section describes the NetApp FAS storage deployment procedure.

### NetApp Storage Controller FAS26xx 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 by using the [HWU application](#) at the [NetApp Support](#) site:

- 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 FAS26XX 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 [FAS26xx documentation](#) available at the [NetApp Support](#) site.

### NetApp ONTAP 9.1

#### Configuration Worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.1 Software Setup Guide](#) at the [NetApp Support](#) site.

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

Table 12) ONTAP 9.1 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 9.1 URL	<<var_url_boot_software>>
Name for cluster	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster B gateway	<<var_clustermgmt_gateway>>
Cluster B netmask	<<var_clustermgmt_mask>>
Domain name	<<var_domain_name>>
DNS server IP (you can enter more than one)	<var_dns_server_ip
NTP server IP (you can enter more than one)	<<var_ntp_server_ip>>

## Configure Node A

To configure node A, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

**Note:** If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 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.

```
7
```

5. Enter y to perform an upgrade.

y

6. Select e0M for the network port you want to use for the download.

e0M

7. Enter y to reboot now.

y

8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

<<var\_nodeA\_mgmt\_ip>> <<var\_nodeA\_mgmt\_mask>> <<var\_nodeA\_mgmt\_gateway>>

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

**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.

y

12. Enter y to reboot the node.

y

**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 when you see this message:

Press Ctrl-C for Boot Menu

14. Select option 4 for Clean Configuration and Initialize All Disks.

4

15. Enter y to zero disks, reset config, and install a new file system.

y

16. Enter y to erase all the data on the disks.

y

**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.

17. While node A is initializing, begin configuring node B.

## Configure Node B

To configure node B, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

**Note:** If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 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.

```
7
```

5. Enter y to perform an upgrade.

```
y
```

6. Select e0M for the network port you want to use for the download.

```
e0M
```

7. Enter y to reboot now.

```
y
```

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.

```
y
```

12. Enter y to reboot the node.

```
y
```

**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 when you see this message:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.

```
4
```

15. Enter `y` to zero disks, reset config, and install a new file system.

```
y
```

16. Enter `y` to erase all the data on the disks.

```
y
```

**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.

### 5.3 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 9.1 boots on the node for the first time.

**Note:** The node and cluster setup procedure has changed slightly in ONTAP 9.1. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up node A.

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".  
To accept a default or omit a question, do not enter a value.
```

```
This system will send event messages and periodic reports to NetApp Technical  
Support. To disable this feature, enter  
autosupport modify -support disable  
within 24 hours.
```

```
Enabling AutoSupport can significantly speed problem determination and  
resolution should a problem occur on your system.  
For further information on AutoSupport, see:  
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}: yes
```

```
Enter the node management interface port [e0M]:  
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>  
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>  
Enter the node management interface default gateway: <<var_nodeA_mgmt_gateway>>  
A node management interface on port e0M with IP address <<var_nodeA_mgmt_ip>> has been created.
```

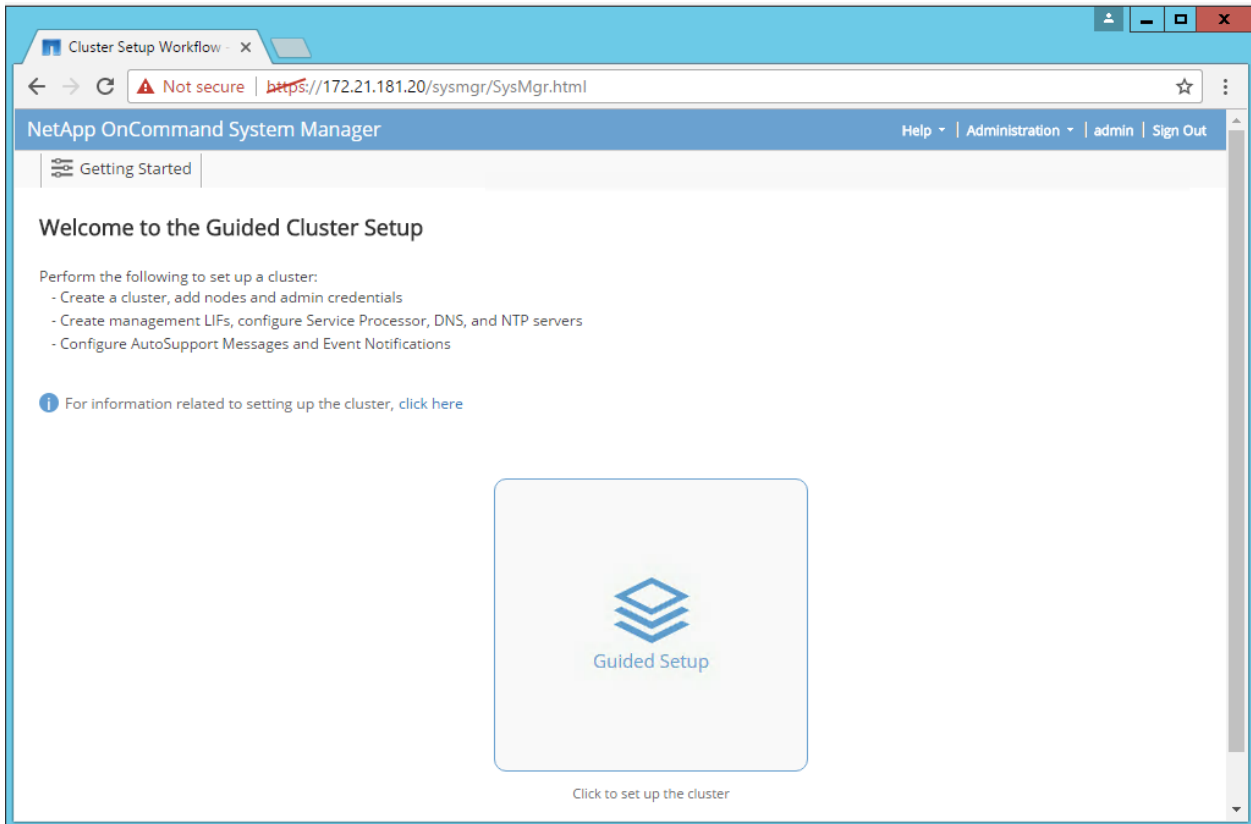
```
Use your web browser to complete cluster setup by accessing  
https://<<var\_nodeA\_mgmt\_ip>>
```

```
Otherwise, press Enter to complete cluster setup using the command line  
interface:
```

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

**Note:** Cluster setup can also be performed using the command line interface. This document describes cluster setup using NetApp System Manager guided setup.


3. Click Guided Setup to configure the cluster.




4. Enter `<<var_clustername>>` for the cluster name and `<<var_nodeA>>` and `<<var_nodeB>>` for each of the nodes you are configuring. Enter the password you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.

## Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Cluster Name 

## Nodes

 Not sure all nodes have been discovered? [Refresh](#)
Cluster Configuration: ☐ Switched Cluster ☐ Switchless Cluster


Username admin

Password Confirm Password Cluster Base License (Optional) 
 For any queries related to licenses, contact [mysupport.netapp.com](https://mysupport.netapp.com)
Feature Licenses (Optional) 

Cluster Base License is mandatory to add Feature Licenses.

Submit

5. You can also enter feature licenses.
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. Next, you are guided through the process of configuring the network. The network screen looks like this.



## Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



## Network (Management)

IP Addresses (IPv4)  
required

Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

## IP Address Range

It is recommended not to manually modify the Cluster Management, Node Management, and Service Processor Management IP addresses. If you are enabling the IP Address Range, entering the Gateway address is mandatory. If you do not want to enter the Gateway address, ensure that the IP Address Range is disabled.

Starting address to Ending address Netmask Gateway

+ Add Range      ✓ Apply Sequentially

IP Address Port

Cluster Management

Node Management

FAS2650_A	172.21.181.20	e0M
FAS2650_B		e0M

Service Processor Management  
Default values have been detected for the Service Processor.  
☐ Override the default values (Gateway is mandatory)

FAS2650_A	172.21.181.11
FAS2650_B	172.21.181.12

## DNS Details

DNS Domain Names

Enter comma separated names...

DNS Server IP Address

Enter comma separated IP Addresses...

## NTP Details

Primary NTP Server

Alternative NTP Server  
(Optional)

Submit

## 8. Begin with the Network Management section.

## ? Network (Management)

### IP Addresses (IPv4) required

Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

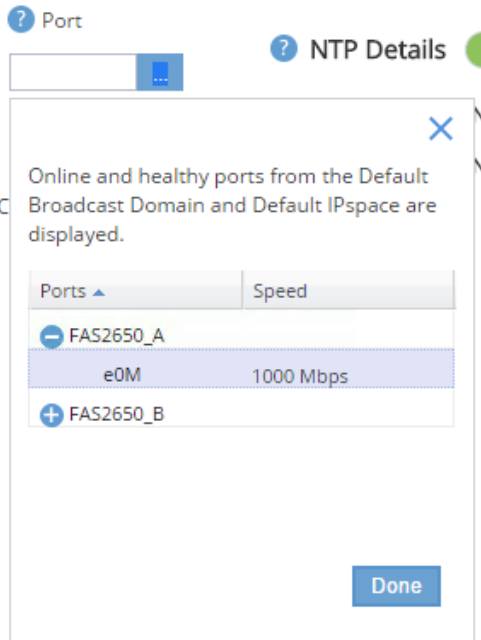
#### ? IP Address Range



You must enter the default network details manually.

	IP Address	Netmask	Gateway (Optional)	? Port
Cluster Management	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="..."/>
Node Management	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
FAS2650_A	<input type="text" value="172.21.181.20"/>	<input type="text" value="e0M"/>	<input type="text" value="..."/>	
FAS2650_B	<input type="text"/>	<input type="text" value="e0M"/>	<input type="text" value="..."/>	
Service Processor Management	Default values have been detected for the Service Processor.			
	<input type="checkbox"/> Override the default values (Gateway is mandatory)			
	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
FAS2650_A	<input type="text" value="172.21.181.11"/>			
FAS2650_B	<input type="text" value="172.21.181.12"/>			

- Deselect IP Address Range.
- 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.

9. Configure the DNS details and NTP details.

**DNS Details** ☒

DNS Domain Names

DNS Server IP Address

**NTP Details** ☒

Primary NTP Server

Alternative NTP Server (Optional)

- a. 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.

- d. Enter <<var\_ntp\_server\_ip>> in the Primary NTP Server field.

**Note:** You can also enter an alternate NTP server.

10. Click Submit.

## 11. Now configure the support information.

NetApp OnCommand System Manager

Getting Started

### Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

✓

✓

3

ClusterNetworkSupportSummary

?

 AutoSupport

?

 Proxy URL (Optional)

i

 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

Submit

- a. If your environment requires a proxy to access AutoSupport®, enter it next to the Proxy URL field.
- e. 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.

## 12. The following message indicates that the cluster configuration has completed. Click Manage Your Cluster to begin to configure the storage.

## Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



[Click here to view the summary](#)

The next step will be to configure your aggregates, SVM and Storage Objects.  
Click the button below to start provisioning your storage.

Manage your cluster

## 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

13. Verify the current mode and the current type of the ports by using the `ucadmin show` command.

```
FAS2650::> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
FAS2650_A	0c	fc	target	-	-	online
FAS2650_A	0d	fc	target	-	-	online
FAS2650_A	0e	fc	target	-	-	online
FAS2650_A	0f	fc	target	-	-	online
FAS2650_B	0c	fc	target	-	-	online
FAS2650_B	0d	fc	target	-	-	online
FAS2650_B	0e	fc	target	-	-	online
FAS2650_B	0f	fc	target	-	-	online

8 entries were displayed.

14. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If this is not the case, 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 FAS2650
```

2. Rename the cluster management LIF.

```
network interface rename -vserver FAS2650 -lif cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver FAS2650 -lif cluster_setup_node_mgmt_lif_FAS2650_B_1 -newname FAS2650-02_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: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

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

## Create Jumbo Frame MTU Broadcast Domain in ONTAP

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

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

## Remove Data Ports from Default Broadcast Domain

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

Issue the following command to remove the ports from the broadcast domain.

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

## Disable Flow Control on UTA2 Ports

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

```
net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

```

net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

```

## Configure IFGRP LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Make sure the switch is properly configured.

From the cluster prompt, complete the following steps.

```

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

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

```

## Configure Jumbo Frames in NetApp ONTAP

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

```

FAS2650::> network port modify -node FAS2650_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

FAS2650::> network port modify -node FAS2650_B -port a0a -mtu 9000

Warning: This command will cause a several second interruption of service on
this network port.
Do you want to continue? {y|n}: y

```

## Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

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

```

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

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <<var_nodeA>>:a0a-
<<var_nfs_vlan_id>>, <<var_nodeB>>:a0a-<<var_nfs_vlan_id>>

```

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



```

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

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_nodeA>>:a0a-
<<var_iscsi_vlan_A_id>>,<<var_nodeB>>:a0a-<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_nodeA>>:a0a-
<<var_iscsi_vlan_B_id>>,<<var_nodeB>>:a0a-<<var_iscsi_vlan_B_id>>

```

### 3. Create MGMT-VLAN ports.

```

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

```

## Create Aggregates in ONTAP

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

To create new 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 Simple Network Management Protocol in ONTAP

To configure the Simple Network Management Protocol (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,fc
```

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 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 storage virtual machines were previously called vservers.

## Configure NFSv3 in ONTAP

You need the following information to complete this configuration step,

**Table 11) 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 new 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, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

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

## 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. To delete the default certificates, run the following commands:

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

```
security certificate delete [TAB] ...  
Example: security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial 552429A6
```

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

```
security certificate create [TAB] ...  
Example: security certificate create -common-name infra-svm.netapp.com -type server -size 2048 -country US -state "North Carolina" -locality "RTP" -organization "NetApp" -unit "FlexPod" -email-addr "abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

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

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

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

```
system services web modify -external true -ssl3-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 SVM to be available by the web.

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

## Create a NetApp FlexVol Volume in ONTAP

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

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_nodeA -size 500GB -  
state online -policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-  
snapshot-space 0  
  
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA -size 100GB -state  
online -policy default -junction-path /infra_swap -space-guarantee none -percent-snapshot-space 0  
-snapshot-policy none  
  
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA -size 100GB -state  
online -policy default -space-guarantee none -percent-snapshot-space 0
```

## Enable Deduplication in ONTAP

To enable deduplication on appropriate volumes, run the following commands:

```
volume efficiency on -vserver Infra-SVM -volume infra_datastore_1  
volume efficiency on -vserver Infra-SVM -volume esxi_boot
```

## Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

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

**Note:** When adding an additional Cisco UCS C-Series server, an additional boot LUN must be created.

## Create iSCSI LIFs in ONTAP

You need the following information to complete this configuration step.

**Table 12) 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>>

### 1. Create four iSCSI LIFs, two on each node.

```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -
home-node <<var_nodeA>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address
<<var_nodeA_iscsi_lif01a_ip>> -netmask <<var_nodeA_iscsi_lif01a_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -
home-node <<var_nodeA>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address
<<var_nodeA_iscsi_lif01b_ip>> -netmask <<var_nodeA_iscsi_lif01b_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -
home-node <<var_nodeB>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address
<<var_nodeB_iscsi_lif01a_ip>> -netmask <<var_nodeB_iscsi_lif01a_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -
home-node <<var_nodeB>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address
<<var_nodeB_iscsi_lif01b_ip>> -netmask <<var_nodeB_iscsi_lif01b_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface show

```

## Create NFS LIFs in ONTAP

You need the following information to complete this configuration step.

**Table 13) 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>>

### 1. Create an NFS LIF.

```

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

```

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

network interface show
```

## Add Infrastructure SVM Administrator

You need the following information to complete this configuration step.

**Table 14) 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
```

## 5.4 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.

You need the following information to configure CIMC for each Cisco UCS C-Series standalone server.

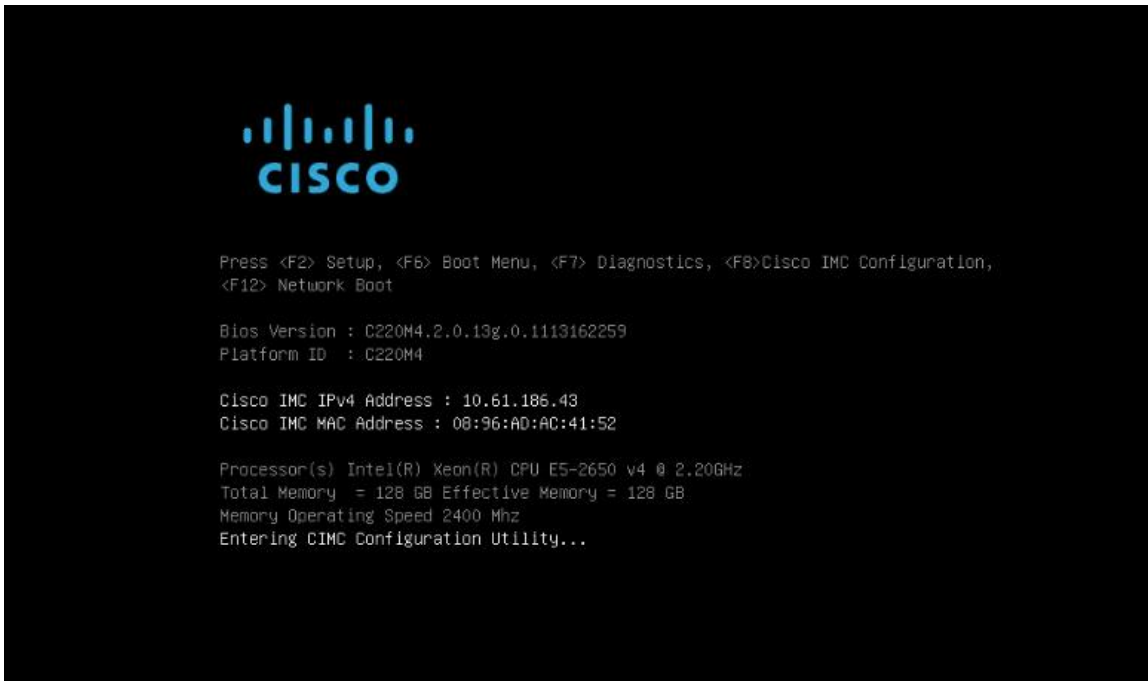
**Table 15) Information required for CIMC configuration.**

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 2.0(13h).

### All Servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.
2. Power on the server and press F8 when prompted to enter the CIMC configuration.



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

- Network interface card (NIC) mode:
  - Dedicated [X]
- IP (Basic):
  - IPV4: [X]
  - DHCP enabled: [ ]
  - CIMC IP: <<cimc\_ip>>
  - Prefix/Subnet: <<cimc\_netmask>>
  - Gateway: <<cimc\_gateway>>



- 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:    10.61.186.43
Prefix/Subnet: 255.255.255.0
Gateway:     10.61.186.1
Pref DNS Server: 0.0.0.0

*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F1>Additional settings

```

4. Press F1 to see additional settings.
  - Common properties:
    - Host name: <<esxi\_host\_name>>
    - Dynamic DNS: [ ]
    - Factory defaults: Leave cleared.
  - Default user (basic):
    - Default password: <<admin\_password>>
    - Reenter password: <<admin\_password>>
    - Port properties: Use default values.
    - Port profiles: Leave cleared.

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
Common Properties
  Hostname:      cle-c220m4-f0246
  Dynamic DNS:   [ ]
  DDNS Domain:
FactoryDefaults
  Factory Default: [ ]
Default User(Basic)
  Default password:
  Reenter password:
Port Properties
  Auto Negotiation: [X]
                        Admin Mode      Operation Mode
  Speed[1000/100/10Mbps]:      Auto          1000
  Duplex mode[half/full]:      Auto          full
Port Profiles
  Reset: [ ]
  Name:
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F2>PreviousPageettings

```

5. Press F10 to save the CIMC interface configuration.
6. After the configuration is saved, press Esc to exit.

## Configure Cisco UCS C-Series Servers iSCSI Boot

In this FlexPod Express configuration, the VIC1227 is used for iSCSI boot. You need the following information to configure each ESXi host.

**Table 16) 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	
IP address iscsi_lif02a	

Detail	Detail Value
IP address iscsi_lif01b	
IP address iscsi_lif02b	
Infra_SVM IQN	

**Note:** Italicized font indicates variables that are unique for each ESXi host.

### Boot Order Configuration

1. From the CIMC interface browser window, click the Server tab and choose BIOS.
2. Choose Configure Boot Order and click OK.

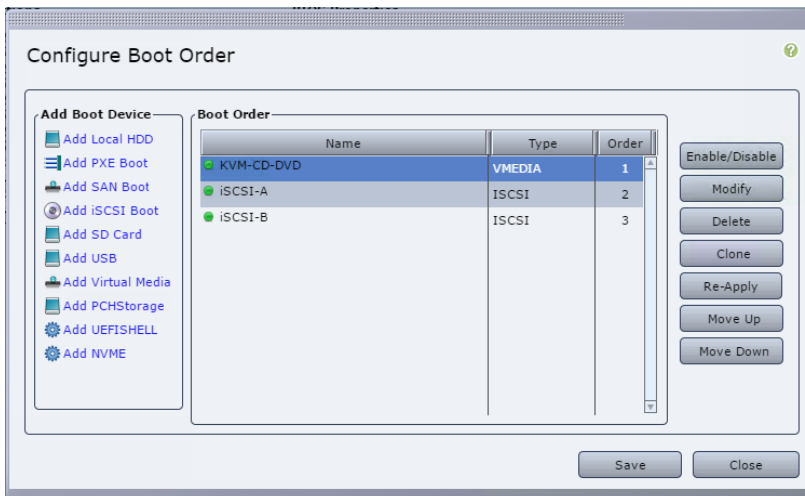


3. Configure the following devices by clicking the device under Add Boot Device.



- Add Virtual Media.

- Name: KVM-CD-DVD
- Subtype: KVM MAPPED DVD
- State: Enabled
- Add iSCSI Boot.
  - Name: iSCSI-A
  - State: Enabled
  - Order: 2
  - Slot: MLOM
  - Port: 0
- 4. Click Add Device.
  - Name: iSCSI-B
  - State: Enabled
  - Order: 3
  - Slot: MLOM
  - Port: 1
- 5. Click Add Device.
- 6. Click Save and then click Close.



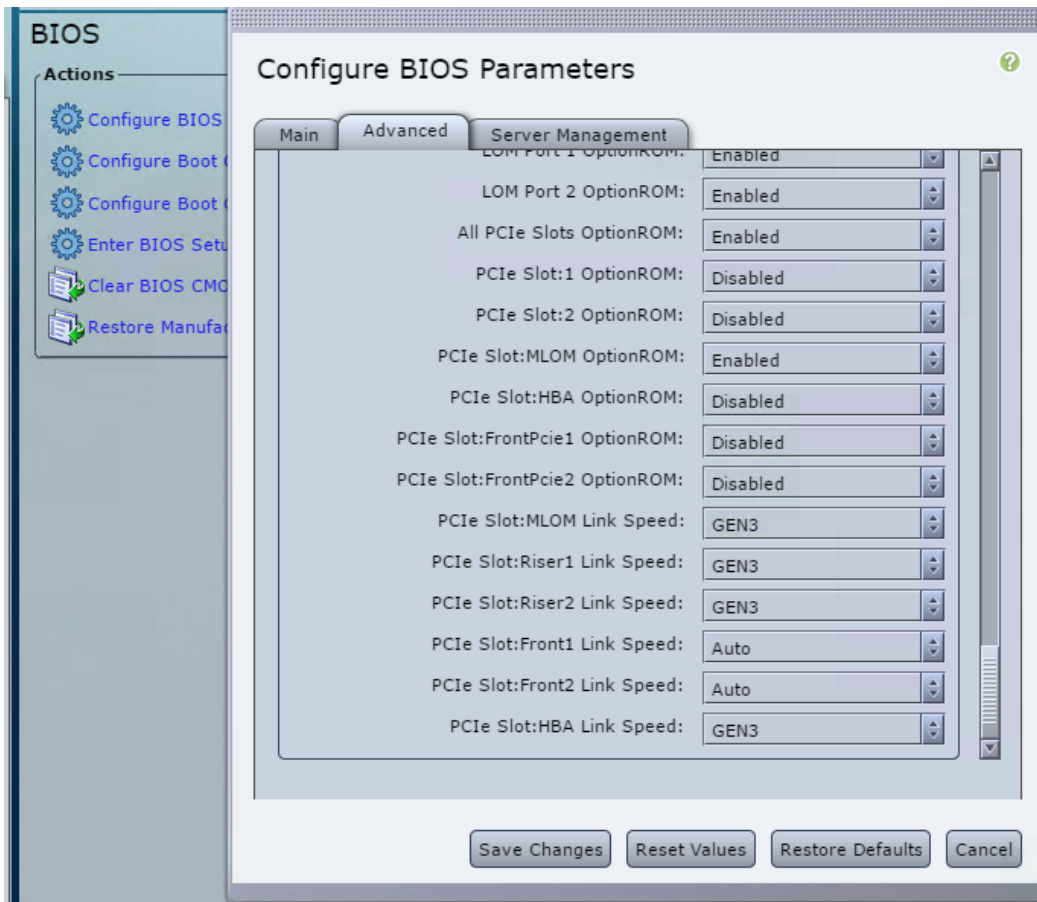
7. Reboot the server to boot with your new boot order.

### Disable RAID Controller (If Present)

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

1. Click BIOS on the left navigation pane in CIMC.
2. Select Configure BIOS.
3. Scroll down to PCIe Slot:HBA Option ROM.

4. If the value is not already disabled, set it to disabled.



## Configure Cisco VIC1227 for iSCSI Boot

The following configuration steps are for the Cisco VIC 1227 for iSCSI boot.

### Create iSCSI vNICs

1. Click Add to create a new vNIC.
2. In the Add vNIC window, complete the following settings:
  - Name: iSCSI-vNIC-A
  - MTU: 9000
  - Default VLAN: <<var\_iscsi\_vlan\_a>>
  - VLAN Mode: TRUNK
  - Enable PXE boot: Check

**Add vNIC**

**General**

Name:

MTU:  (1500 - 9000)

Uplink Port:

MAC Address: ☒ AUTO ☐

Class of Service:

Trust Host CoS: ☐

PCI Order: ☒ ANY ☐  (0 - 17)

Default VLAN: ☒ NONE ☐  (1 - 4094)

VLAN Mode:

Rate Limit: ☒ OFF ☐  (1 - 10000 Mbps)

Enable PXE Boot: ☒

Channel Number:  (1 - 1000)

3. Click Add vNIC and then click OK.
4. Repeat the process to add a second vNIC.
  - a. Name the vNIC iSCSI-vNIC-B.
  - b. Enter <<var\_iscsi\_vlan\_b>> as the VLAN.
  - c. Set the uplink port to 1.
5. Select the vNIC iSCSI-vNIC-A and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.

**Adapter Card MLOM**

General vNICs VM FEXs vHBAs

**Host Ethernet Interfaces**

Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot
eth0	VIC-3-eth0	00:81:C4:D0:DC:0C	9000	0	0	0	NONE	TRUNK	disabled
eth1	VIC-3-eth1	00:81:C4:D0:DC:0D	9000	0	1	0	NONE	TRUNK	disabled
iSCSI-vNIC-A	VIC-3-iSCS	00:81:C4:D0:DC:10	9000	0	0	0	3439	TRUNK	enabled
iSCSI-vNIC-B	VIC-3-iSCS	00:81:C4:D0:DC:11	9000	0	1	0	3440	TRUNK	enabled

6. From the iSCSI Boot Configuration window, enter the initiator details, as shown in the following screen shot.

The image shows a 'iSCSI Boot Configuration' window. At the top, it says 'IP Version: IPv4'. Below this is a section titled 'Initiator' which contains several input fields: 'Name' (with value 'iqn.1995-05.com.cisco:u' and a '(0 - 223) chars' limit), 'IP Address' (192.168.55.17), 'Subnet Mask' (255.255.255.0), 'Gateway' (192.168.55.1), 'Primary DNS' (empty), 'Secondary DNS' (empty), 'TCP Timeout' (15, with a '(0 - 255)' limit), 'CHAP Name' (empty, with a '(0 - 50) chars' limit), and 'CHAP Secret' (empty, with a '(0 - 50) chars' limit). At the bottom of the window are four buttons: 'Configure iSCSI', 'Unconfigure iSCSI', 'Reset Values', and 'Cancel'.

7. 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>>

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 using the `vserver iscsi show` command.

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

**iSCSI Boot Configuration**

**Primary Target**

Name:  (1 - 223) chars

IP Address:

TCP Port:

Boot LUN:  (0 - 65535)

CHAP Name:  (0 - 50) chars

CHAP Secret:  (0 - 50) chars

**Secondary Target**

Name:  (1 - 223) chars

IP Address:

TCP Port:

Boot LUN:  (0 - 65535)

10. Click Configure iSCSI.
11. Select the vNIC `iSCSI-vNIC-B` and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.
12. Repeat the process to configure `iSCSI-vNIC-B`.
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>>`
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.

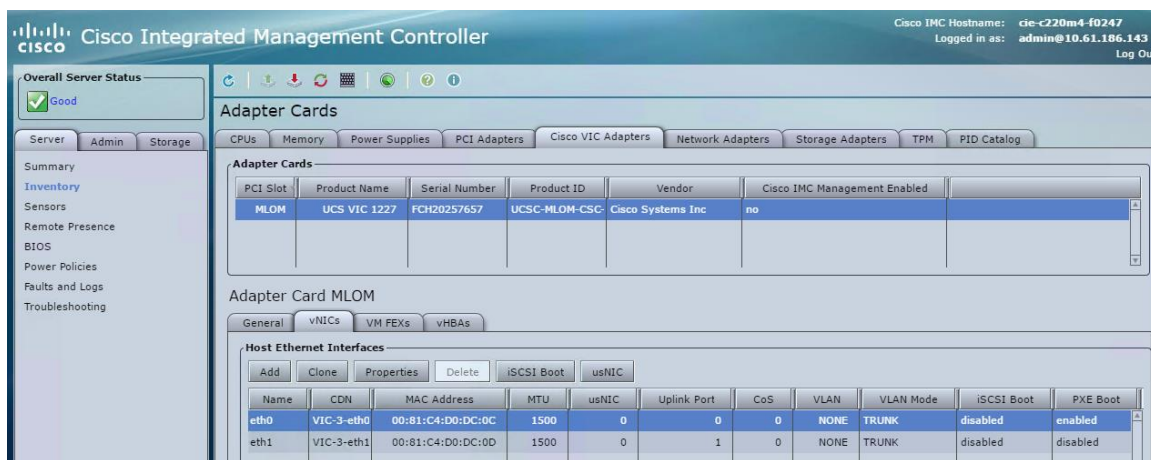
**Note:** Be sure to record the IQN names for each vNIC. You need them for a later step.

16. Click Configure iSCSI.

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 Adapter Cards, select Cisco UCS VIC 1227 and then select the vNICs underneath.



3. Select eth0 and click Properties.

4. Set the MTU to 9000. Click Save Changes.

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

Adapter Card MLOM

General VNICS VM FEXs vHBAs

Host Ethernet Interfaces

Add Clone Properties Delete iSCSI Boot usNIC

Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot
eth0	VIC-3-eth0	00:81:C4:D0:DC:0C	9000	0	0	0	NONE	TRUNK	disabled
eth1	VIC-3-eth1	00:81:C4:D0:DC:0D	9000	0	1	0	NONE	TRUNK	disabled
iSCSI-vNIC-A	VIC-3-iSCS	00:81:C4:D0:DC:10	9000	0	0	0	3439	TRUNK	enabled
iSCSI-vNIC-B	VIC-3-iSCS	00:81:C4:D0:DC:11	9000	0	1	0	3440	TRUNK	enabled

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

## 5.5 NetApp FAS Storage Deployment Procedure (Part 2)

### ONTAP SAN Boot Storage Setup

#### Create iSCSI Igroups

To create igroups, complete the following step:

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

1. From the cluster management node SSH connection, run the following commands.  
To view the three igroups created in this step, run the `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.

## 5.6 VMware vSphere 6.5 Deployment Procedure

This section provides detailed procedures for installing VMware ESXi 6.5 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 GUI 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 6.5 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

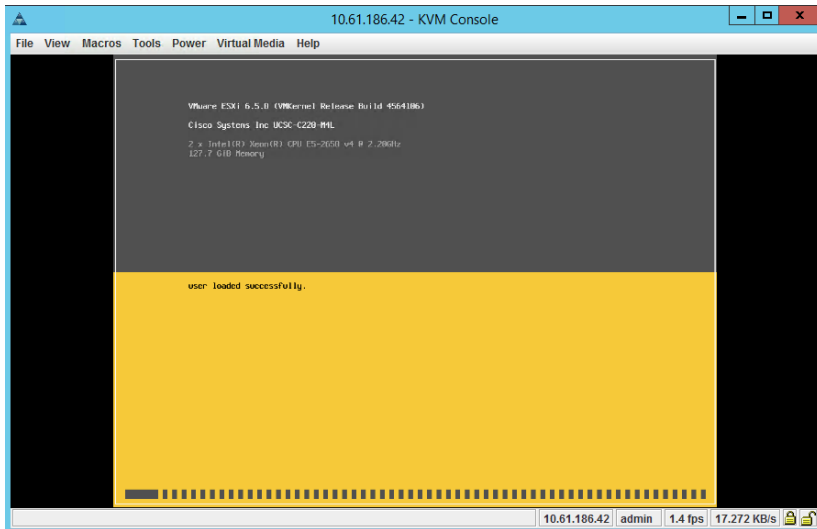
The following steps describe how to install VMware ESXi on each host.

#### Download ESXi 6.5 Cisco Custom Image

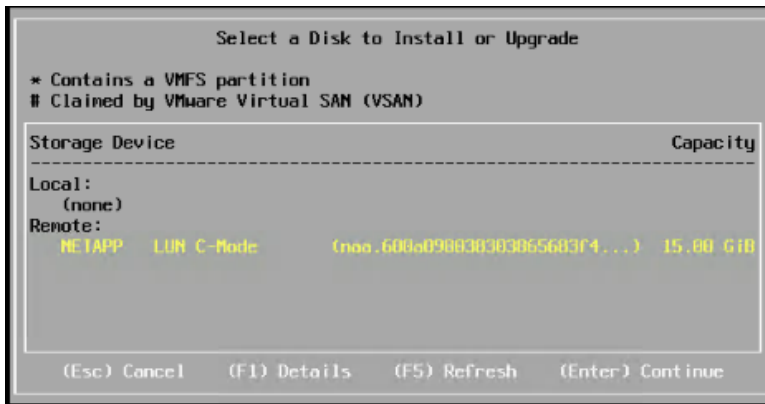
1. Navigate to the [VMware vSphere download page](#) for custom ISOs.
2. Click Go to Downloads next to the CISCO Custom Image for ESXi 6.5 GA Install CD.
3. Download the Cisco Custom Image for ESXi 6.5 GA Install CD (ISO).

#### All Hosts

1. When the system boots, the machine detects the presence of the VMware ESXi installation media.
2. Select the VMware ESXi installer from the menu that appears.
3. The installer loads. This takes several minutes.



4. After the installer has finished loading, press Enter to continue with the installation.
5. After reading the end-user license agreement, accept it and continue with the installation by pressing F11.
6. Select the NetApp LUN that was previously set up as the installation disk for ESXi, and press Enter to continue with the installation.



7. Select the appropriate keyboard layout and press Enter to continue.
8. Enter and confirm the root password and press Enter to continue.
9. The installer warns you that existing partitions are removed on the volume. Continue with the installation by pressing F11.
10. After the installation is complete, unmap the VMware ESXi installation image on the Virtual Media tab of the KVM console so that the server reboots into VMware ESXi and not the installer.
11. The Virtual Media window might warn you that it is preferable to eject the media from the guest. Because you cannot do this in this example (and the media is read-only), unmap the image anyway by selecting Yes.


## 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.

**Note:** Select the ports that correspond to eth0 and eth1 in CIMC.



6. Select VLAN (optional) and press Enter.
7. Enter the VLAN ID <<mgmt\_vlan\_id>>. Press Enter.
8. From the Configure Management Network menu, select IPv4 Configuration to configure the IP address of the management interface. Press Enter.
9. Use the arrow keys to highlight Set Static IPv4 address and use the space bar to select this option.
10. Enter the IP address for managing the VMware ESXi host <<esxi\_host\_mgmt\_ip>>.
11. Enter the subnet mask for the VMware ESXi host <<esxi\_host\_mgmt\_netmask>>.
12. Enter the default gateway for the VMware ESXi host <<esxi\_host\_mgmt\_gateway>>.
13. Press Enter to accept the changes to the IP configuration.
14. Enter the IPv6 configuration menu.

15. Use the space bar to disable IPv6 by unselecting the Enable IPv6 (restart required) option. Press Enter.
16. Enter the menu to configure the DNS settings.
17. Because the IP address is assigned manually, the DNS information must also be entered manually.
18. Enter the primary DNS server's IP address <<nameserver\_ip>>.
19. (Optional) Enter the secondary DNS server's IP address.
20. Enter the FQDN for the VMware ESXi host name: <<esxi\_host\_fqdn>>.
21. Press Enter to accept the changes to the DNS configuration.
22. Exit the Configure Management Network submenu by pressing Esc.
23. Press Y to confirm the changes and reboot the server.
24. Log out of the VMware Console by pressing Esc.

## Configure ESXi Host

You need the following information to configure each ESXi host.

Table 17) Information required for configuring ESXi hosts.

Detail	Value
ESXi host name	
ESXi host management IP	
ESXi host management mask	
ESXi host management gateway	
ESXi host NFS IP	
ESXi host NFS mask	
ESXi host NFS gateway	
ESXi host vMotion IP	
ESXi host vMotion mask	
ESXi host vMotion gateway	
ESXi host iSCSI-A IP	
ESXi host iSCSI-A mask	
ESXi host iSCSI-A gateway	
ESXi host iSCSI-B IP	
ESXi host iSCSI-B mask	
ESXi host iSCSI-B gateway	

## Log In to ESXi Host

1. Navigate to the host's management IP address in a web browser.
2. Log into 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.

## Update Cisco VIC nenic If Required

### Verify nenic Version

1. Click Manage under Host on the left navigation pane. Click Packages on the right. Scroll through the package list until you find the nenic driver. Note the version.



2. Continue with the following process to update the nenic driver if required.

**Note:** The nenic version used in this configuration is 1.1.1.0. Be sure to check the [Cisco UCS hardware and software compatibility tool](#) for information about the latest supported drivers.



## Download Required Software

The following steps provide details for downloading the Cisco virtual interface card (VIC) enic driver.

1. In a web browser, navigate to <http://www.cisco.com>.
2. Under Support, click Download Software.
3. Click Servers – Unified Computing.
4. Click UCS Virtual Interface Card.
5. Click UCS C-Series Rack-Mount Standalone Server Software in the select product screen.
6. Select your server model.
7. Select Unified Computing System (UCS) Drivers.
8. Download the VMware drivers.

**Note:** You must sign in to Cisco.com to download the drivers.

**Note:** The drivers for the Cisco UCS C220 M4 are currently located [here](#).

## Installing the Updated Cisco VIC eNIC Driver

### All Hosts

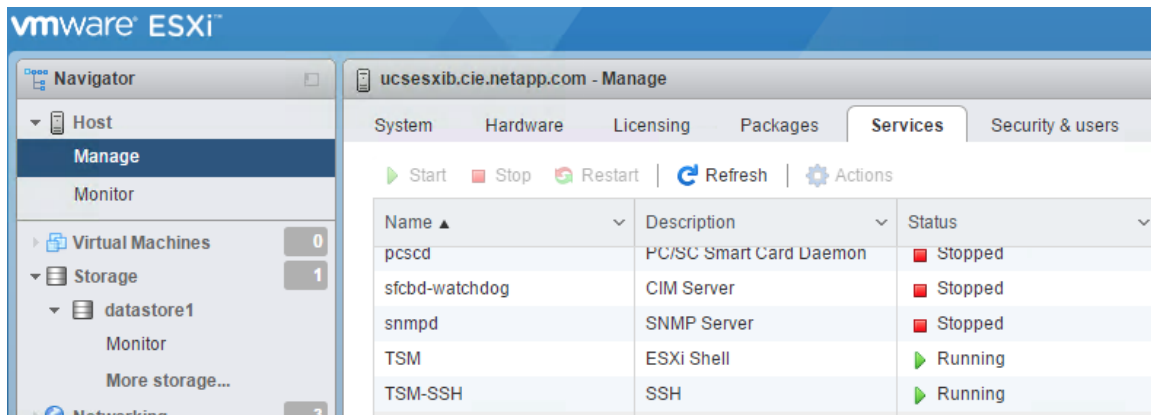
To load the updated versions of the eNIC driver for the Cisco VIC, follow these steps for all the hosts from the vSphere client:

1. Click datastore1 in the left navigation pane.
2. Click Datastore browser in the right pane.



3. Select Upload in the datastore browser and upload the vSphere Installation Bundle (VIB) file.
4. SSH to the management IP ESXi host, enter root for the user name, and enter the root password.

**Note:** You must start the SSH service by navigating to Manage in the left navigation pane and then clicking services. Right-click the Tech Support Mode (TSM) and TSM-SSH services and select Start to start the services. Be sure to stop them when you are finished.



5. Run the following commands on each host:

```
esxcli software vib update -d /vmfs/volumes/datastore1/<<driver_name>>.zip
```

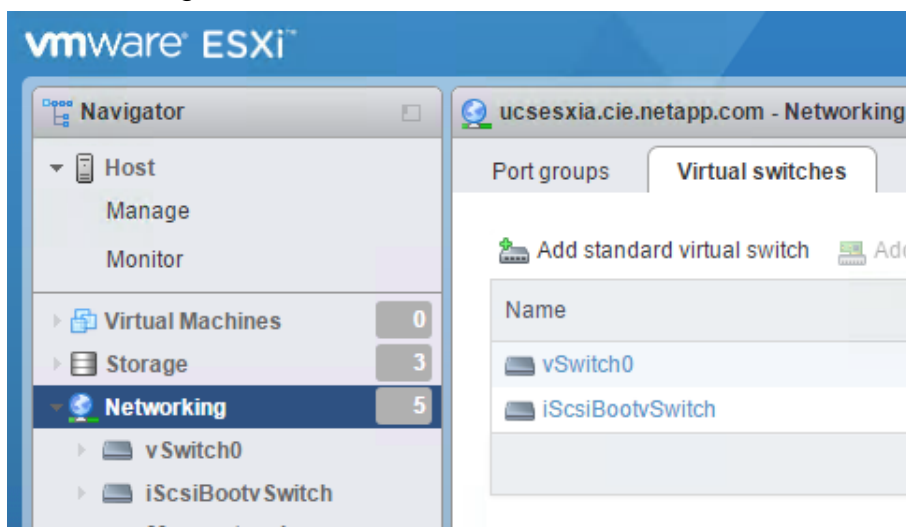
6. A message that the update has completed successfully and a reboot is required appears. You can also see the name of the VIB removed and the VIB installed.

7. Type the following command to reboot the host:

```
reboot
```

## Configure iSCSI Boot

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



3. Click iScsiBootvSwitch.
4. Select Edit settings.

5. Change the MTU to 9000 and click Save.
6. Click Networking in the left navigation pane to return to the Virtual Switches tab.
7. Click Add Standard Virtual Switch.
8. Provide the name `iScsiBootvSwitch-B` for the vSwitch name.
  - Set the MTU to 9000.
  - Select vmnic3 from the Uplink 1 pull-down options.
  - Click Add.

**Note:** Vmnic2 and vmnic3 are used for iSCSI boot in this configuration. If you have additional NICs in your ESXi host, you might have different vmnic numbers. To confirm which NICs are used for iSCSI boot, match the MAC addresses on the iSCSI vNICs in CIMC to the vmnics in ESXi.
9. In the center pane, select the VMkernel NICs tab.
10. Select Add VMkernel NIC.
  - a. Specify a new port group name of `iScsiBootPG-B`.
  - b. Select `iScsiBootvSwitch-B` 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.


Port group	New port group ▼
New port group	iScsiBootPG-B
Virtual switch	iScsilBootvSwitch-B ▼
VLAN ID	3440
MTU	9000
IP version	IPv4 only ▼
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.184.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack ▼
Services	<input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication

Create   Cancel

## 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 in the right pane.
2. Select the iSCSI software adapter and click Configure iSCSI.



3. Under Dynamic Targets, click Add Dynamic Target.

Configure iSCSI - vmhba64

iSCSI enabled	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled								
Name & alias	iqn.1992-08.com.cisco:ucsaiscsia								
CHAP authentication	Do not use CHAP								
Mutual CHAP authentication	Do not use CHAP								
Advanced settings	Click to expand								
Network port bindings	<div> <span>Add port binding</span> <span>Remove port binding</span> </div> <table border="1"> <thead> <tr> <th>VMkernel NIC</th> <th>Port group</th> <th>IPv4 address</th> </tr> </thead> <tbody> <tr> <td colspan="3">No port bindings</td> </tr> </tbody> </table>			VMkernel NIC	Port group	IPv4 address	No port bindings		
VMkernel NIC	Port group	IPv4 address							
No port bindings									
Static targets	<div> <span>Add static target</span> <span>Remove static target</span> <span>Edit settings</span> </div> <table border="1"> <thead> <tr> <th>Target</th> <th>Address</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td>iqn.1992-08.com.netapp:sn.09591199033811e78eb...</td> <td>172.21.183.34</td> <td>3260</td> </tr> </tbody> </table>			Target	Address	Port	iqn.1992-08.com.netapp:sn.09591199033811e78eb...	172.21.183.34	3260
Target	Address	Port							
iqn.1992-08.com.netapp:sn.09591199033811e78eb...	172.21.183.34	3260							
Dynamic targets	<div> <span>Add dynamic target</span> <span>Remove dynamic target</span> <span>Edit settings</span> </div> <table border="1"> <thead> <tr> <th>Address</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td colspan="2">No dynamic targets</td> </tr> </tbody> </table>			Address	Port	No dynamic targets			
Address	Port								
No dynamic targets									

Save configuration Cancel

4. Enter the IP address `iscsi_lif01a`.

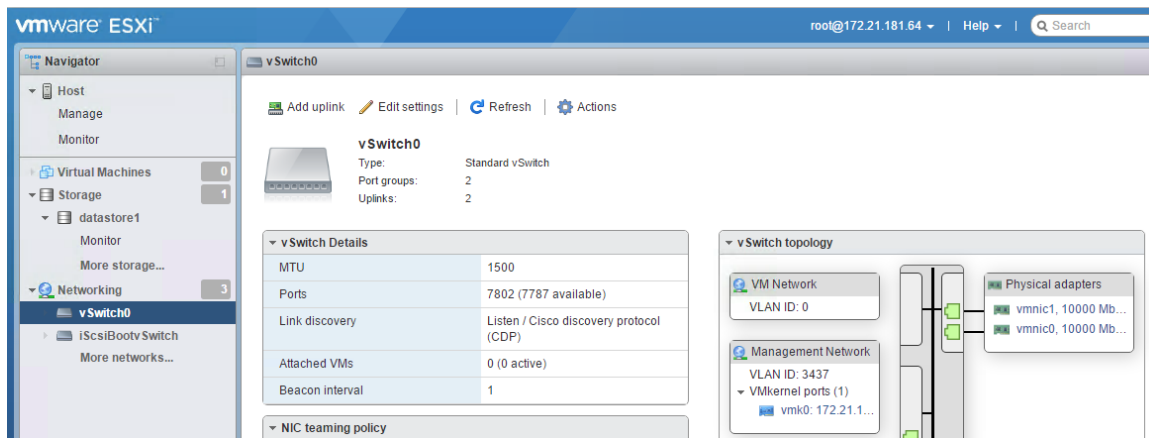
- Repeat with the IP addresses `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
- Click Save Configuration.

Dynamic targets		<a href="#">Add dynamic target</a> <a href="#">Remove dynamic target</a> <a href="#">Edit settings</a>
Address	Port	
172.21.183.33	3260	
172.21.183.34	3260	
172.21.184.33	3260	
172.21.184.34	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 OnCommand® System Manager.

## Configure ESXi Host

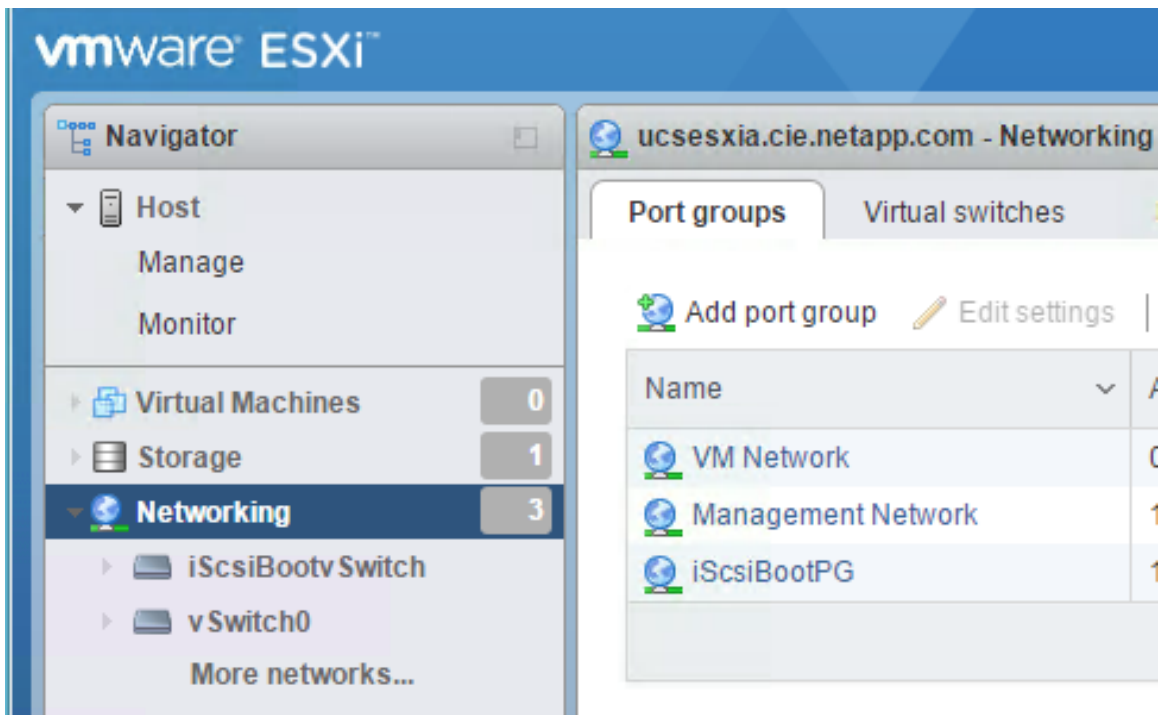
1. In the left navigation pane, select Networking.
2. Select vSwitch0.



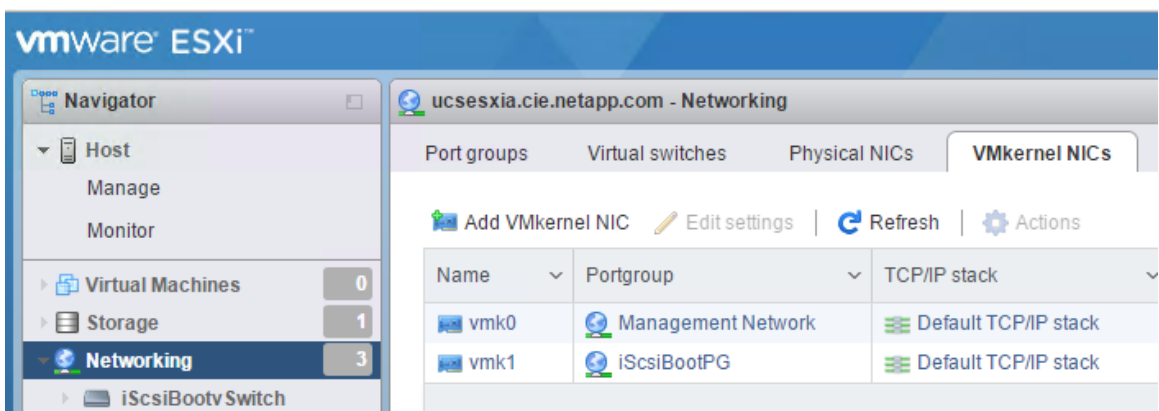
3. Select Edit Settings.
4. Change the MTU to 9000.
5. Expand NIC Teaming and verify that both vmnic0 and vmnic1 are set to active.

## Configure Port Groups and VMkernel NICs

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



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



6. Select Add VMkernel NIC.
  - a. Select New Port Group.
  - b. Name the port group `NFS-Network`.

- c. Enter <<nfs\_vlan\_id>> for the VLAN ID.
- d. Change the MTU to 9000.
- e. Expand IPv4 Settings.
- f. Select Static Configuration.
- g. Enter <<var\_hosta\_nfs\_ip>> for Address.
- h. Enter <<var\_hosta\_nfs\_mask>> for Subnet Mask.
- i. Click Create.

**Add VMkernel NIC**

Port group	New port group ▼
New port group	NFS-Network
Virtual switch	vSwitch0 ▼
VLAN ID	3438
MTU	9000
IP version	IPv4 only ▼
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.182.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack ▼

Create Cancel

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.

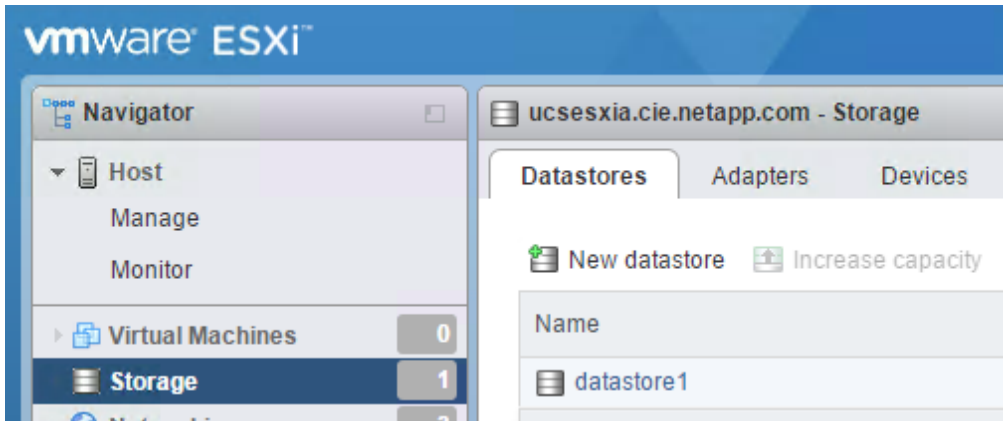
Add VMkernel NIC	
Virtual switch	vSwitch0
VLAN ID	3441
MTU	9000
IP version	IPv4 only
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.21.185.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input checked="" type="checkbox"/>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

**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.

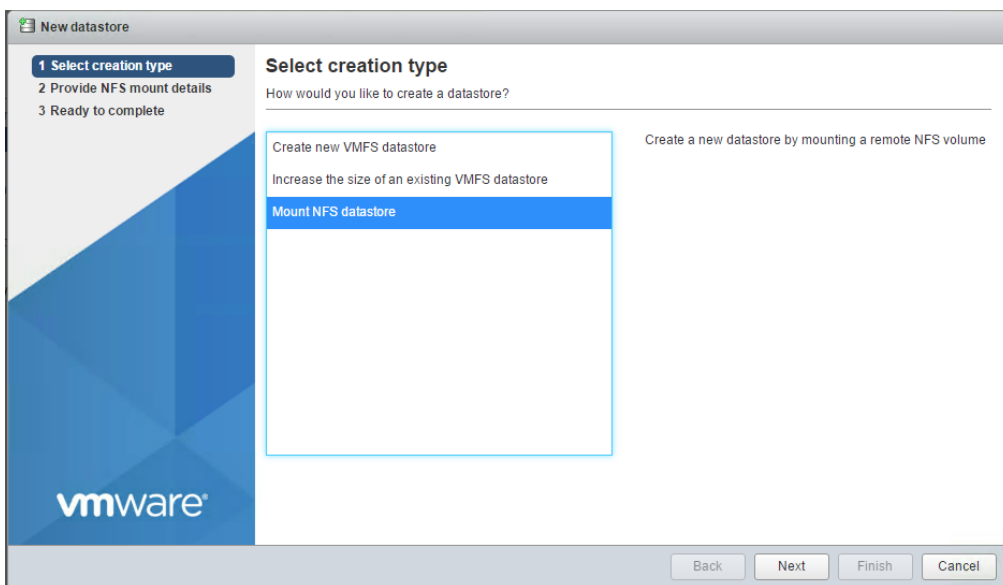
### Mount First Datastores

The first datastores to be mounted are the infra\_datastore\_1 datastore for virtual machines and the infra\_swap datastore for virtual machine swap files.

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



## 2. Select Mount NFS Datastore.



## 3. Next, enter the following information in the Provide NFS Mount Details screen:

- Name: `infra_datastore_1`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_datastore_1`
- Make sure that NFS 3 is selected.

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

## 5. Repeat this process to mount the infra\_swap datastore:

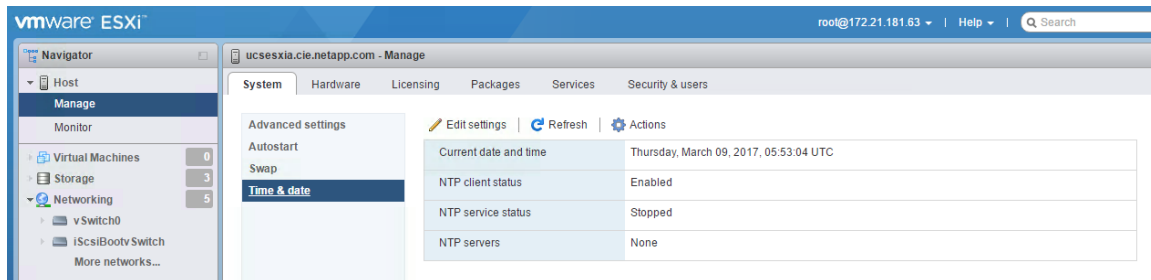
- Name: `infra_swap`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_swap`
- Make sure that NFS 3 is selected.

**Note:** Additional datastores can be created by using NetApp VSC, which is installed in a later step. VSC enhances NetApp and VMware vSphere integration and simplifies management.

## 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 Time & Date.



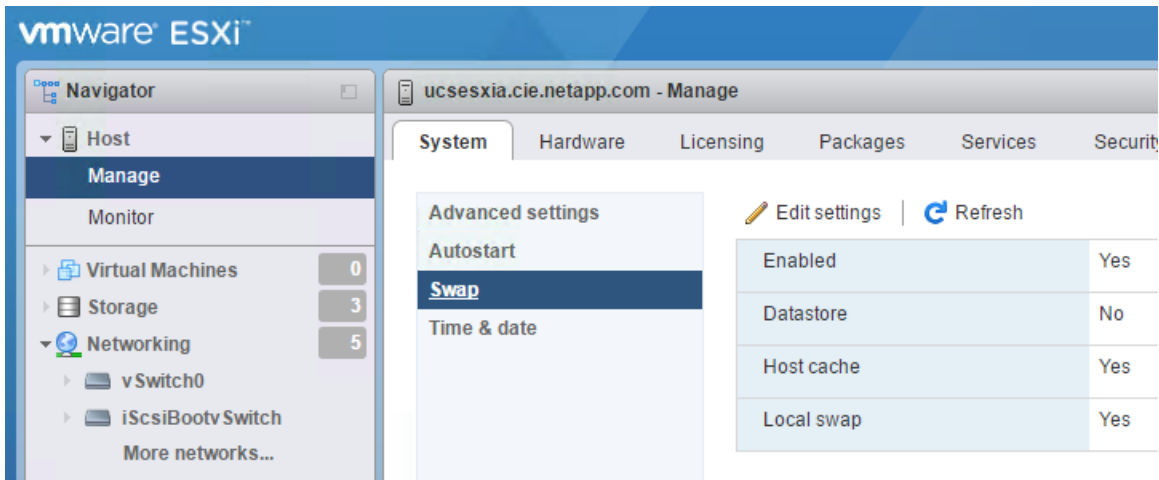
2. Select Use Network Time Protocol (Enable NTP Client).
3. Select Start and Stop with Host as the NTP service startup policy.
4. Enter <<var\_ntp>> as the NTP server. You can set multiple NTP servers.
5. Click Save.

The 'Edit time configuration' dialog box is shown. It has two radio buttons: 'Manually configure the date and time on this host' (unselected) and 'Use Network Time Protocol (enable NTP client)' (selected). Under the NTP section, 'NTP service startup policy' is set to 'Start and stop with host'. The 'NTP servers' field contains '10.61.184.251' and has a hint text below it: 'Separate servers with commas, e.g. 10.31.21.2, fe00::2800'. At the bottom right are 'Save' and 'Cancel' buttons.

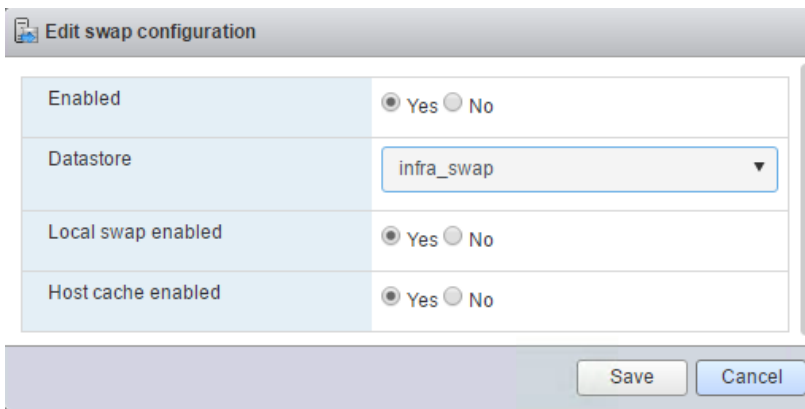
## Move the Virtual Machine Swap-File Location

These steps provide details for moving the virtual machine swap-file location.

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



2. Click Edit Settings. Select infra\_swap from the Datastore pull-down menu.



3. Click Save.

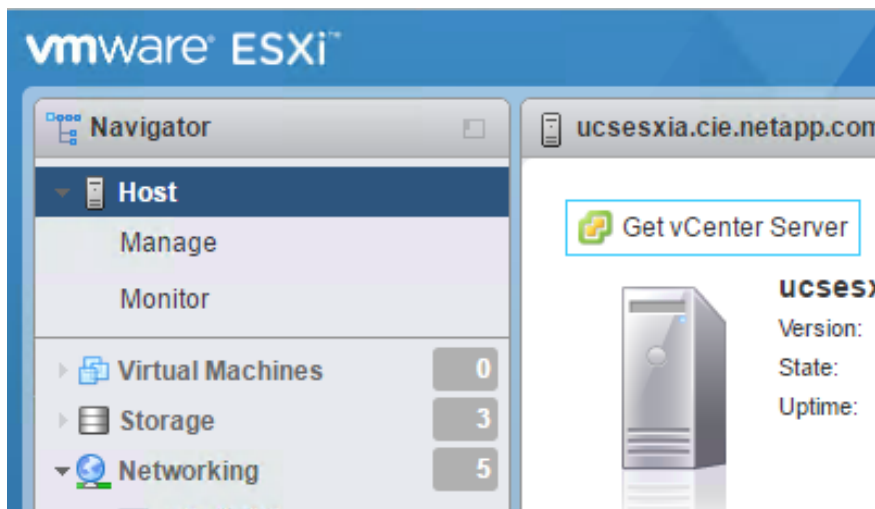
## 5.7 Install VMware vCenter Server 6.5

This section provides detailed procedures for installing VMware vCenter Server 6.5 in a FlexPod Express configuration.

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

### Download the VMware vCenter Server Appliance

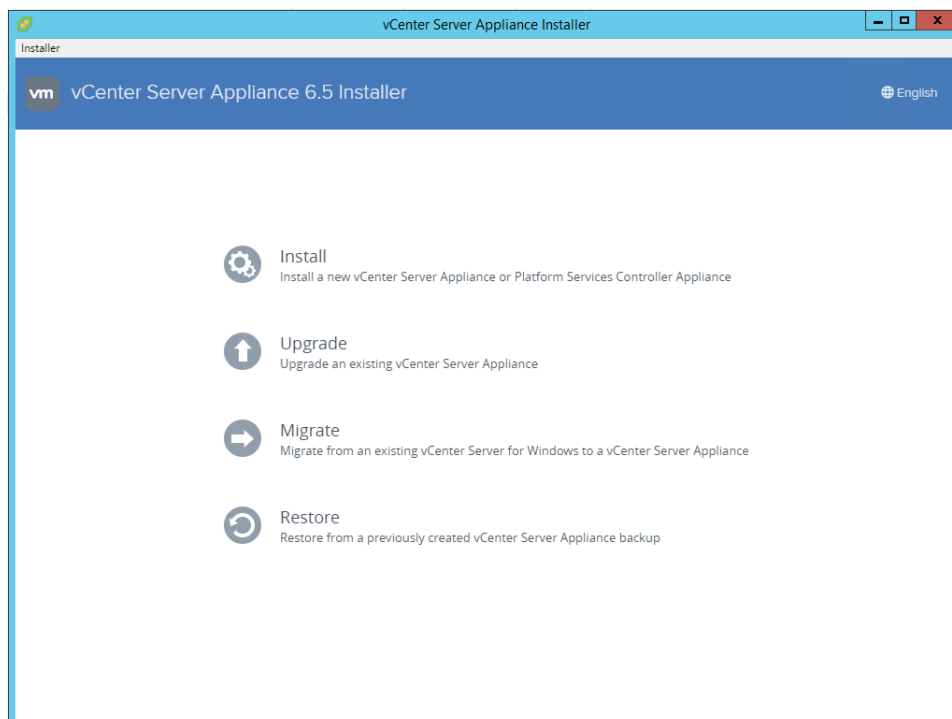
1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.



2. Download the VCSA from the VMware site.

**Note:** Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.

3. Mount the ISO image.
4. Navigate to the vcsa-ui-installer> win32 directory. Double-click installer.exe.
5. Click Install.



6. Click Next at the Introduction screen.
7. Accept the end-user license agreement on the next screen. Click Next.
8. Select Embedded Platform Services Controller as the deployment type.

Install - Stage 1: Deploy appliance

✓ 1 Introduction

✓ 2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Select deployment type

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

For more information on deployment types, refer to the vSphere 6.5 documentation.

Embedded Platform Services Controller

☒ vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

☐ Platform Services Controller
☐ vCenter Server (Requires External Platform Services Controller)

Back

Next

Finish

Cancel

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

- Enter the IP address of an ESXi host you have deployed, as well as the root user name and root password.

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

✓ 1 Introduction

✓ 2 End user license agreement

✓ 3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

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

172.21.181.63

ⓘ

HTTPS port

443

User name

root

ⓘ

Password

\*\*\*\*\*

Back

Next

Finish

Cancel

70

FlexPod Express with VMware vSphere  
6.5 and FAS2600

© 2017 NetApp, Inc. All Rights Reserved.

10. Enter `VCSA` as the VM name along with the root password you would like to use for the VCSA. Click Next.

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

✓ 1 Introduction

✓ 2 End user license agreement

✓ 3 Select deployment type

✓ 4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name

VCSA

?

Root password

.....

?

Confirm root password

.....

Back

Next

Finish

Cancel

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

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

✓ 1 Introduction

✓ 2 End user license agreement

✓ 3 Select deployment type

✓ 4 Appliance deployment target

✓ 5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Select 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 6.5 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	250	10	100
Small	4	16	290	100	1000
Medium	8	24	425	400	4000
Large	16	32	640	1000	10000
X-Large	24	48	980	2000	35000

Back

Next

Finish

Cancel

12. Select the infra\_datastore\_1 datastore. Click Next.

13. Enter the following information in the Configure network settings screen:

- Select MGMT-Network for Network.
- Enter the FQDN or IP to be used for the VCSA.
- Enter the IP address to be used.
- Enter the subnet mask to be used.
- Enter the default gateway.
- Enter the DNS server.

14. Click Next.



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

✓ 1 Introduction

✓ 2 End user license agreement

✓ 3 Select deployment type

✓ 4 Appliance deployment target

✓ 5 Set up appliance VM

✓ 6 Select deployment size

✓ 7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Configure network settings

Configure network settings for this vCenter Server with an Embedded Platform Services Controller.

Network

MGMT-Network

IP version

IPv4

IP assignment

static

System name

VCSA.cie.netapp.com

IP address

172.21.181.99

Subnet mask or prefix length

255.255.255.0

Default gateway

172.21.181.1

DNS servers

10.61.184.251

Back

Next

Finish

Cancel

15. On the Ready to Complete Stage 1 screen, verify that the settings you have entered are correct. Click Finish.
16. The VCSA installs now. This process takes several minutes. A status bar appears during this time.

73 FlexPod Express with VMware vSphere  
6.5 and FAS2600

© 2017 NetApp, Inc. All Rights Reserved.

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




Cancel

17. After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.

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

✓ You have successfully deployed the vCenter Server with an Embedded Platform Services Controller.



Deployment complete

To proceed with stage 2 of the deployment process, appliance setup, click Continue.

If you exit, you can continue with the appliance setup at any time by logging in to the vCenter Server Appliance Management Interface <https://vc.vikings.cisco.com:5480/>

Continue

Close

18. At the Stage 2 Introduction screen, click Next.

## Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

1 Introduction

2 Appliance configuration

3 SSO configuration

4 Configure CEIP

5 Ready to complete

Introduction

vCenter Server Appliance installation overview

Stage 1

  
Deploy new vCenter Server Appliance

Stage 2

  
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.

Back

Next

Finish

Cancel

19. Enter `<<var_ntp_id>>` for the NTP server address. You can enter multiple NTP IP addresses.

**Note:** If you plan to use vCenter Server high availability (HA), make sure that SSH access is enabled.

✓ 1 Introduction

2 Appliance configuration

3 SSO configuration

4 Configure CEIP

5 Ready to complete

Appliance configuration

Time synchronization mode

Synchronize time with NTP servers ▼

NTP servers (comma-separated list)

10.61.184.251

SSH access

Enabled ▼

Back

Next

Finish

Cancel

20. Configure the SSO domain name, password, and site name. Click Next.

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

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

✓ 1 Introduction

✓ 2 Appliance configuration

3 SSO configuration

4 Configure CEIP

5 Ready to complete

SSO configuration

SSO domain name  ⓘ

SSO user name

SSO password  ⓘ

Confirm password

Site name  ⓘ

ⓘ In vCenter 6.5, joining a vCenter with embedded PSC to an external PSC is not supported. For more information on recommended vCenter and PSC topologies, refer to the vCenter Server documentation.

Back

Next

Finish

Cancel

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

22. View the summary of your settings. Click Finish or use the back button to edit settings.

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

✓ 1 Introduction

✓ 2 Appliance configuration

✓ 3 SSO configuration

✓ 4 Configure CEIP

✓ 5 Ready to complete

Ready to complete  
Review your settings before finishing the wizard.

**Network Details**

Network configuration	Assign static IP address
IP version	IPv4
Host name	VCSA-Express.cie.netapp.com
IP Address	172.21.181.99
Subnet mask	255.255.255.0
Gateway	172.21.181.1
DNS servers	10.61.184.251

**Appliance Details**

Time synchronization mode	Synchronize time with NTP servers
NTP Server	10.61.184.251
SSH access	Enabled

**SSO Details**

Domain name (new)	vsphere.local
Site name (new)	FlexPod
User name	administrator


Back

Next

Finish

Cancel

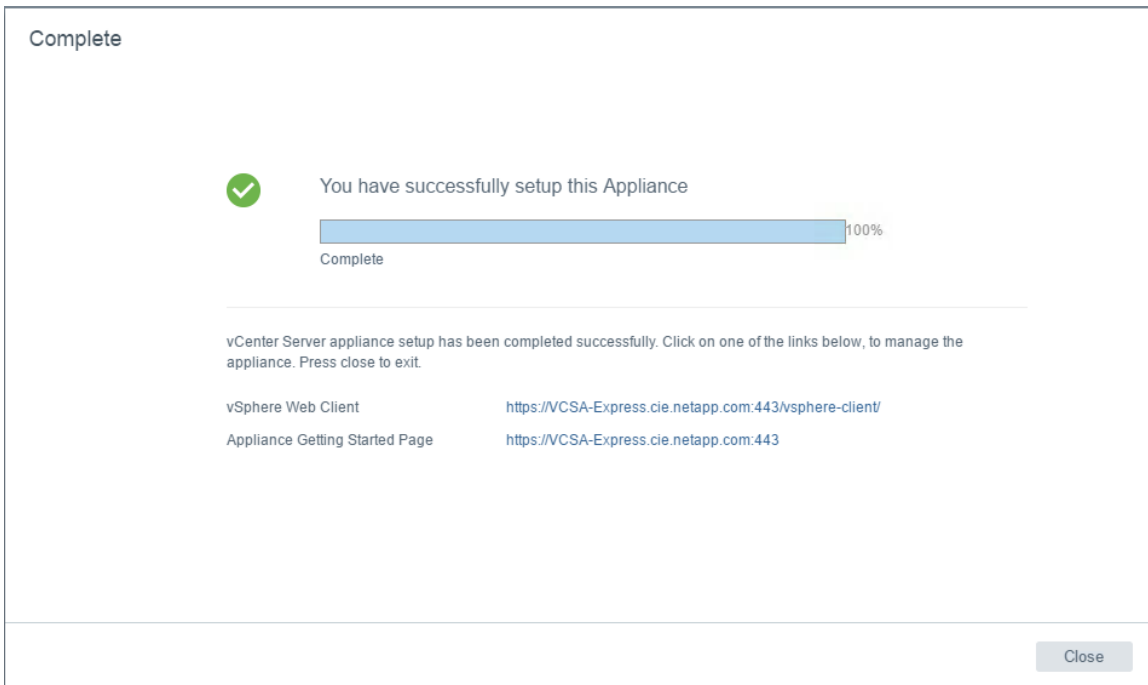
23. A warning message appears stating that you are not able to pause or stop the installation from completing after it has started. Click OK to continue.



24. The appliance setup continues. This takes several minutes.

25. A message appears indicating that the setup was successful.

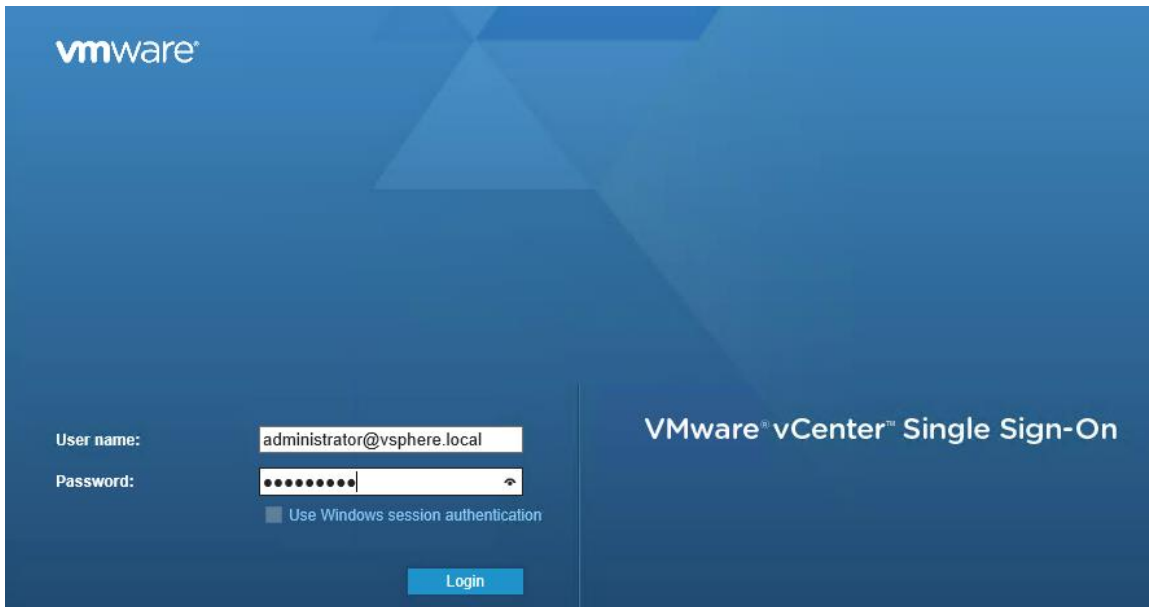
**Note:** The links that the installer provides to access vCenter Server are clickable.



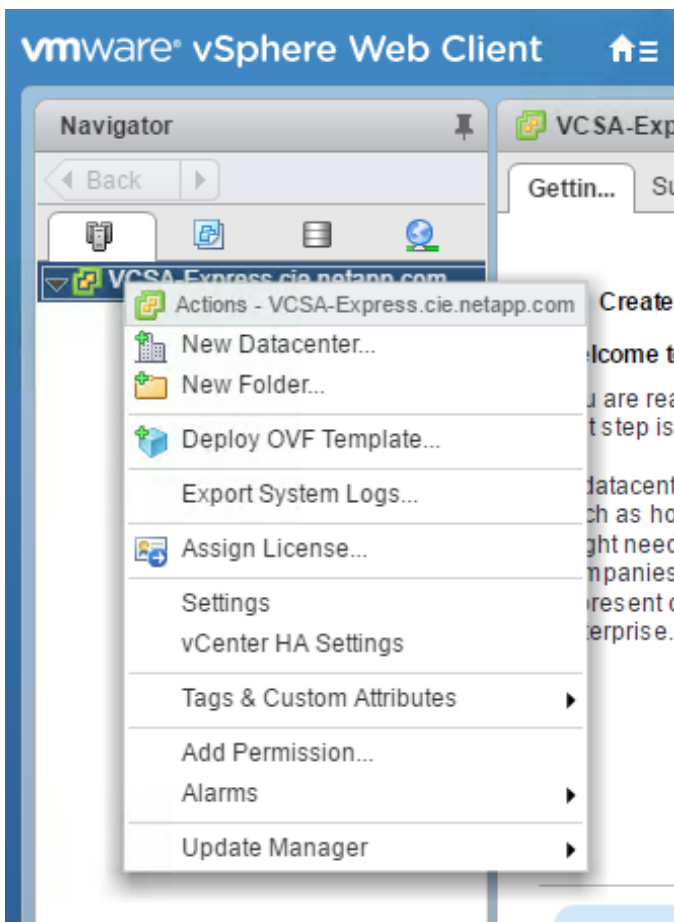
## 5.8 Configure VMware vCenter Server 6.5 and vSphere Clustering

To configure VMware vCenter Server 6.5 and vSphere clustering, complete the following steps:


1. Navigate to <https://<<FQDN or IP of vCenter>>/vsphere-client/>.
2. Log in with the user name [administrator@vsphere.local](mailto:administrator@vsphere.local) and the SSO password you entered during the VCSA setup process.



3. Right-click the vCenter name and select New Datacenter.

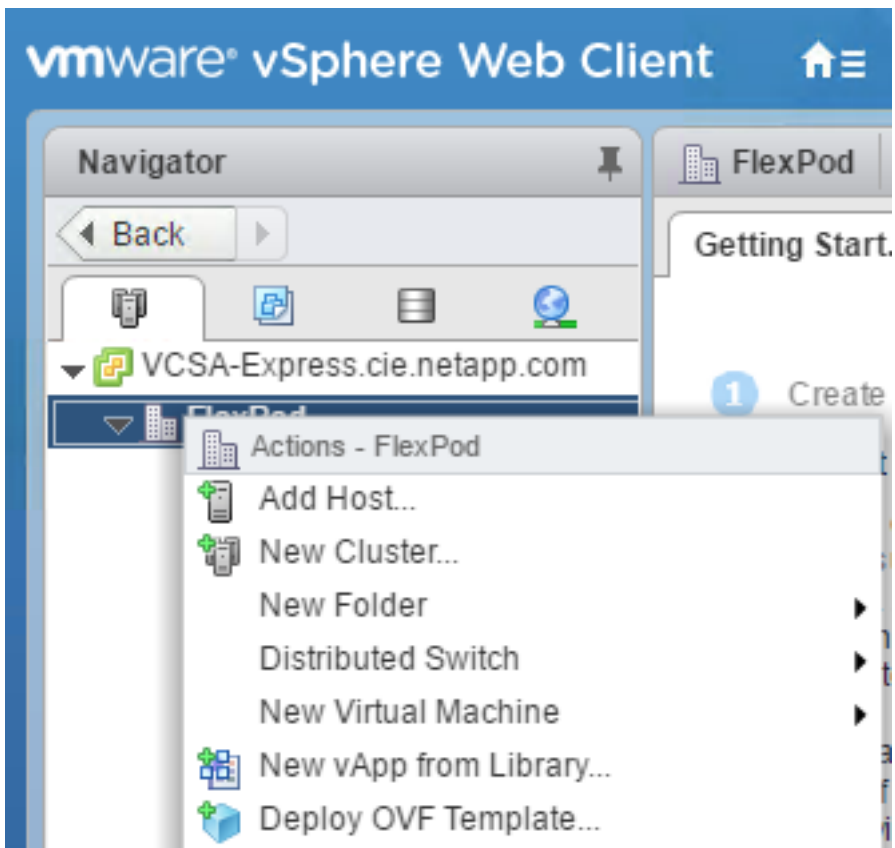


4. Enter a name for the data center.



### Create vSphere Cluster

1. Right-click the newly created data center and select New Cluster.




2. Complete the following steps to create a vSphere cluster:
  - a. Name the cluster.
  - b. Click the checkbox to Turn on DRS.
  - c. Click the checkbox to Turn on HA.
  - d. Click OK.

Name	FlexPod Express
Location	FlexPod
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	<input checked="" type="checkbox"/> Enable admission control
VM Monitoring	<input type="checkbox"/> Disabled <small>Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.</small>
Monitoring Sensitivity	Low ——— High
EVC	Disable
Virtual SAN	<input type="checkbox"/> Turn ON

OK Cancel

## Add ESXi Hosts to Cluster

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



2. To add an ESXi host to the cluster, complete the following steps:
  - a. Enter the IP or FQDN of the host. Click Next.



- b. Enter the root user name and password. Click Next.
  - c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
  - d. Click Next on the host summary screen.
  - e. Click the green + symbol to add a license to the vSphere host.
- Note:** This step can be completed later if desired.
- f. Click Next to leave lockdown mode disabled.
  - g. Click Next at the VM location screen.
  - h. Review the Ready to Complete screen. Use the back button to make any changes or select Finish.

Add Host	
1 Name and location	Name: ucsesxia.cie.netapp.com
2 Connection settings	Version: VMware ESXi 6.5.0 build-4564106
3 Host summary	License: Evaluation License
4 Assign license	Networks: MGMT-Network VM Network
5 Lockdown mode	Datastores: datastore1 test_vol infra_swap infra_datastore
6 Resource pool	Lockdown mode: Disabled
7 Ready to complete	Resources destination: FlexPod Express

Back Next Finish Cancel

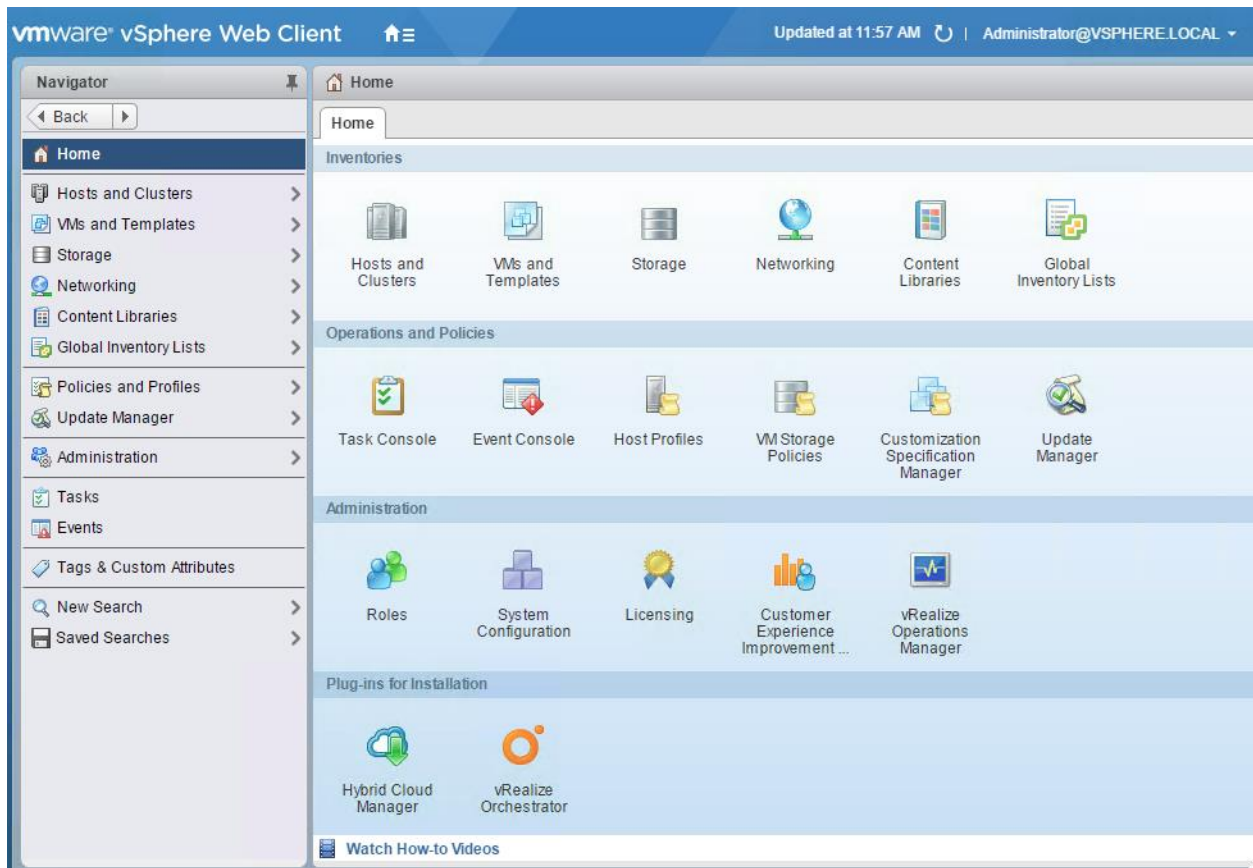
3. Repeat this process for Cisco UCS host B.

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

### ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator must be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter appliance. To set up the ESXi Dump Collector, complete the following steps:

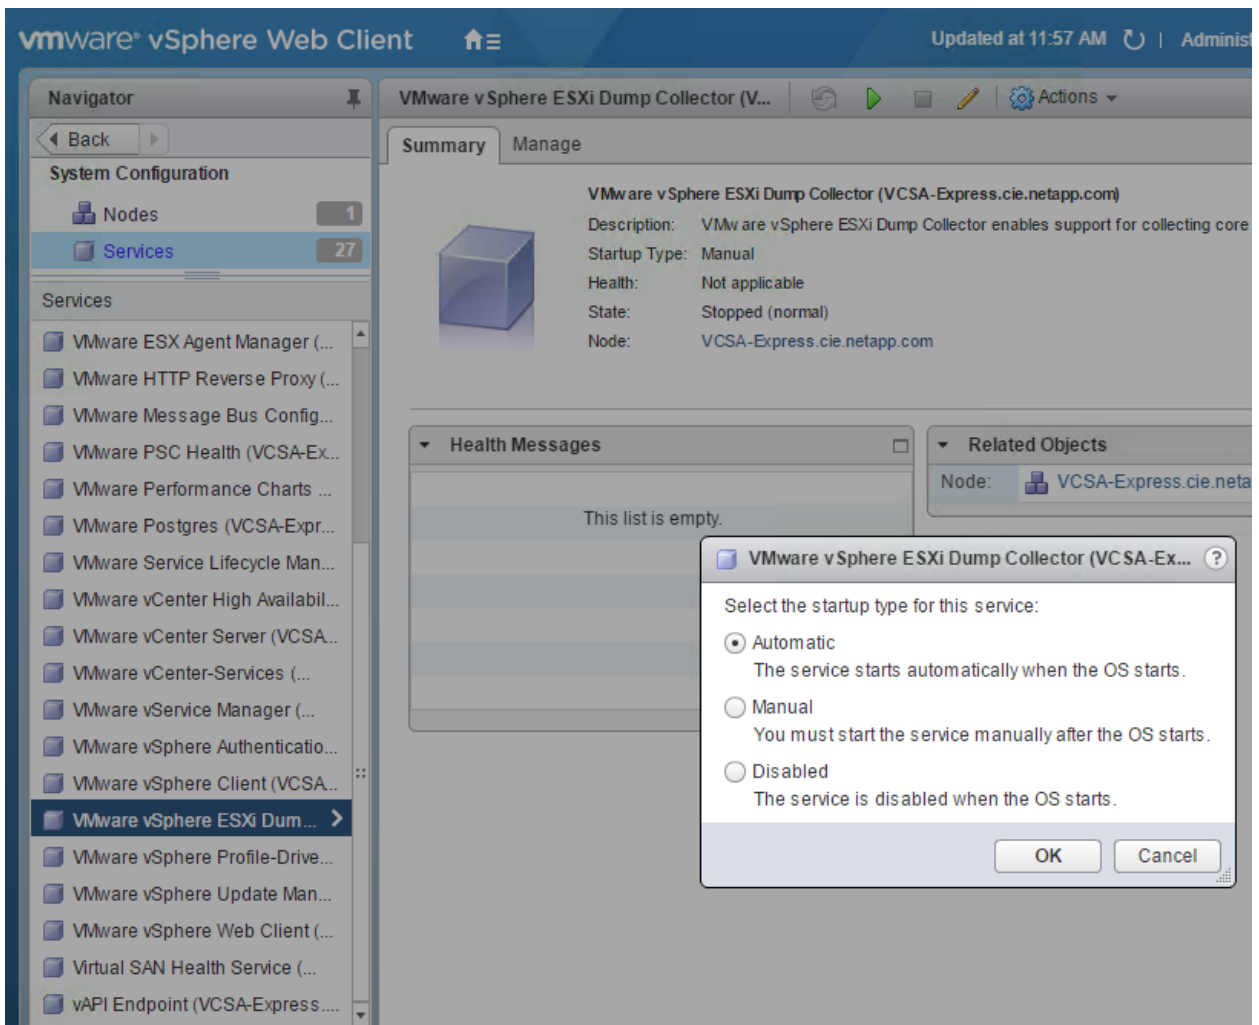
1. Select Home in the vSphere Client when logged in to vCenter. This is the house icon with three lines at the top of the client next to the phrase VMware vSphere Web Client.



2. Click Administration.
3. Click System Configuration under Deployment.



4. You are brought to the System Configuration screen. Click Services.
5. Locate the VMware vSphere ESXi Dump Collector Service.
  - a. Right-click the VMware vSphere ESXi Dump Collector Service and select Edit Startup Type.
  - b. Change the startup type to Automatic and click OK.
  - c. Click the triangular green play button to start the service.



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

### Configure Coredump on ESXi Hosts

1. SSH to the management IP ESXi host, enter root for the user name, and enter the root password.
2. Run the following commands:

```
esxcli system coredump network set --interface-name=vmk0 --server-ip4=<<vcenter-ip>> --server-port=6500
esxcli system coredump network set --enable=true
esxcli system coredump network check
```

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

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

## 5.9 NetApp Virtual Storage Console 6.2.1P1 Deployment Procedure

This section describes the deployment procedures for the NetApp VSC.

### Virtual Storage Console 6.2.1P1 Preinstallation Considerations


The following licenses are required for VSC on storage systems that run ONTAP 9.1:

- Protocol licenses (NFS and iSCSI)
- NetApp FlexClone® technology (for provisioning and cloning only)
- NetApp SnapRestore® technology (for backup and recovery)
- The NetApp SnapManager® Suite

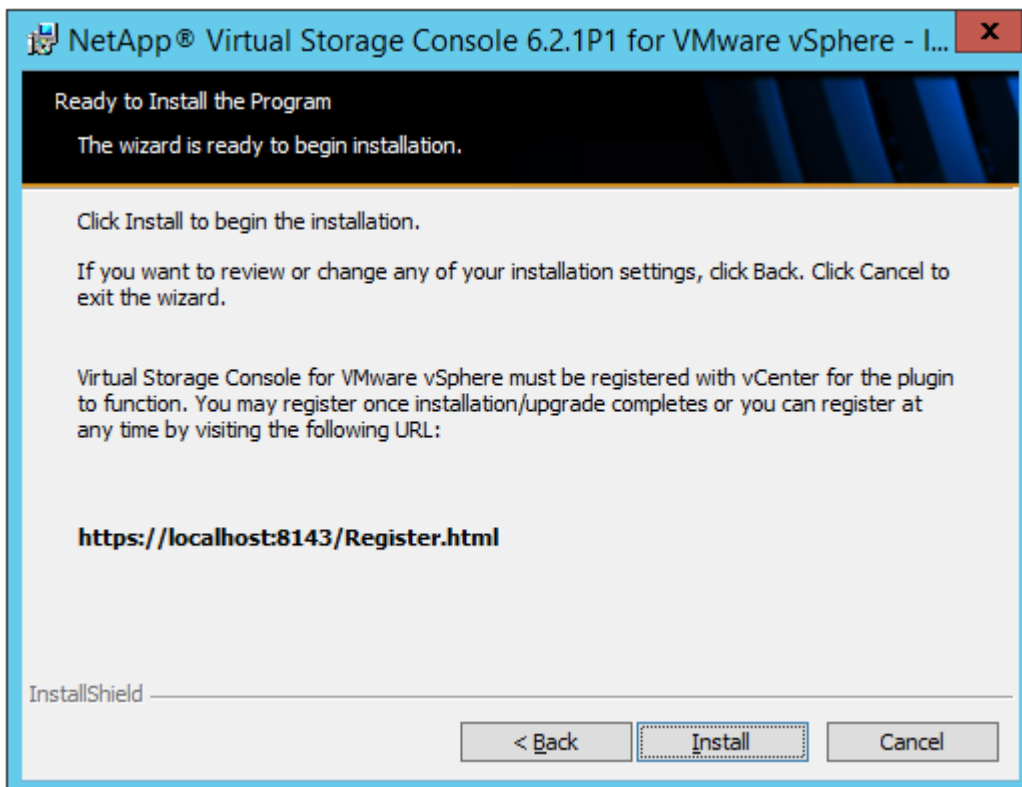
### Install Virtual Storage Console 6.2.1P1

To install the VSC 6.2.1P1 software, complete the following steps:

1. Build a VSC VM with Windows Server 2012 R2, 4GB of RAM, two CPUs, and one virtual network interface in the MGMT network port group. The virtual network interface should be a VMXNET 3 adapter.
2. Bring up the VM, install VMware Tools, assign the IP address and gateway in the MGMT subnet, and join the machine to the Active Directory domain.
3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the user interfaces and infrastructure feature on the VM.
4. Install all Windows updates on the VM.
5. Log in to the VSC VM as an administrative user using the VMware console.  
**Note:** Alternatively, you may use the Remote Desktop Protocol to access the server after you have added an administrative user to the Remote Desktop Users group.
6. From the VMware console on the VSC VM, download the x64 version of VSC 6.2.1P1 from the NetApp Support site.
7. Right-click the `VSC-6.2.1P1-win64.exe` file downloaded in step 6 and select Run as Administrator.
8. Select the appropriate language and click OK.
9. On the installation wizard welcome page, click Next.
10. Select the checkbox to accept the message and click Next.
11. Click Next to accept the default installation location.



12. Click Install.



13. Click Finish.

### Register Virtual Storage Console with vCenter Server

To register the VSC with the vCenter Server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open <https://localhost:8143/Register.html> in Internet Explorer.
2. Click Continue to This Website (Not Recommended).
3. In the Plug-in Service Information section, select the local IP address of the VSC VM.
4. In the vCenter Server Information section, enter the host name or IP address, the user name (FlexPod admin user or root), and the user password for the vCenter Server. Click Register to complete the registration.

#### vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.

Plugin service information

Host name or IP Address:

vCenter Server information

Host name or IP Address:

Port:

User name:

User password:


Register

5. Upon successful registration, storage controller discovery begins automatically.

### Install NetApp NFS VAAI Plug-In on VSC Server

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 plug-in download on the VSC VM.
2. Rename the downloaded file `NetAppNasPlugin.vib`.
3. Move the file to the `C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web` folder.



## Discover and Add Storage Resources

To discover storage resources for the monitoring and host configuration capability and the provisioning and cloning capability, complete the following steps:

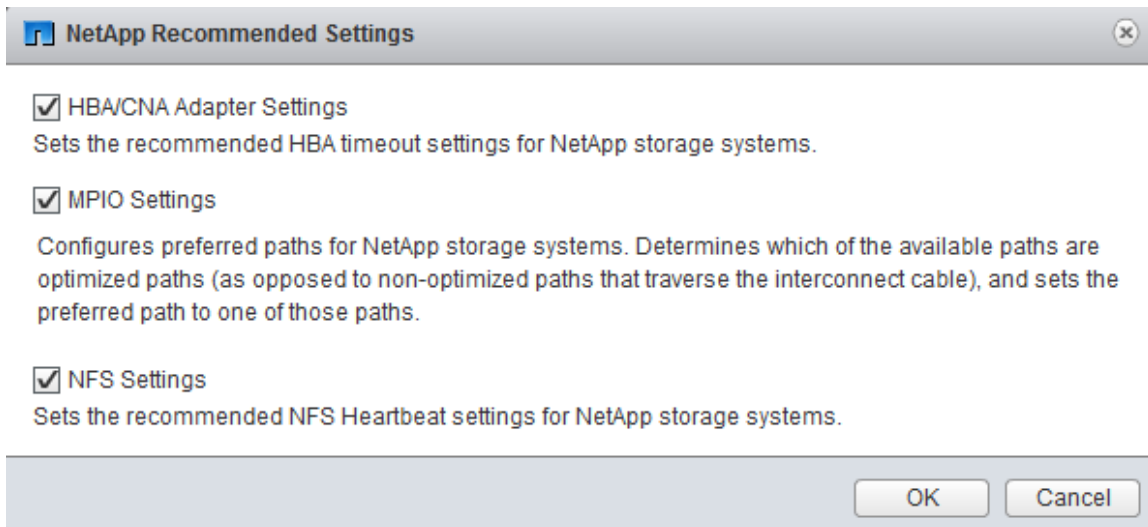
1. Using the vSphere web client, log in to the vCenter Server as the FlexPod admin user. If the vSphere web client was previously opened, close it and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.
3. Select Storage Systems. Under the Objects tab, click Actions > Modify.
4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name and the admin password for password. Confirm that Use TLS to Connect to This Storage System is selected. Click OK.
5. Click OK to accept the controller privileges.
6. Wait for the storage systems to update. You might need to click Refresh to complete this update.

## Optimal Storage Settings for ESXi Hosts

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

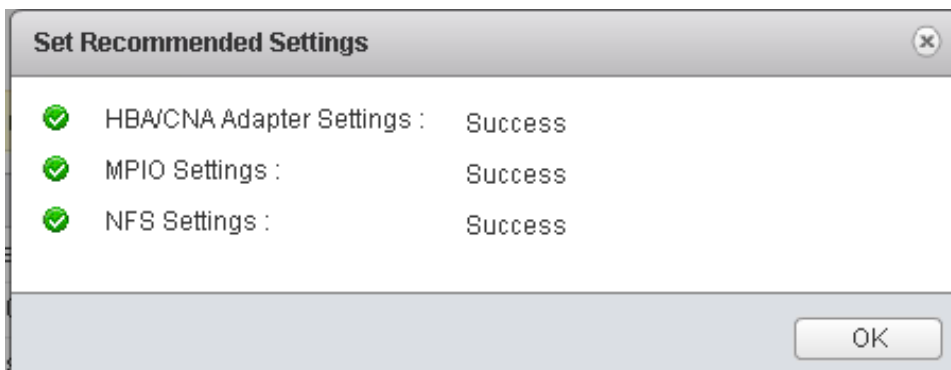
1. From the Home screen, click vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for these hosts.





2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings.

**Note:** This functionality sets values for HBAs and converged network adapters, sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).



3. From the Home screen in the vSphere Web Client, select Virtual Storage Console.
4. On the left under Virtual Storage Console, select NFS VAAI Tools.  
**Note:** Make sure that NFS Plug-in for VMware VAAI Version 1.1.2-3 is shown.
5. Click Install on Host.
6. Select both ESXi hosts and click Install.
7. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.

**Note:** This process must be completed when adding additional servers.

## Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of

additional components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed by keeping in mind small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

## About the Authors

### **Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp**

Lindsey Street is a solutions architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has a bachelor of science degree in computer networking and a master of science degree in information security from East Carolina University.

### **Melissa Palmer, Solutions Architect, Infrastructure and Cloud Engineering, NetApp**

Melissa Palmer is a solutions architect in the NetApp Infrastructure and Cloud Engineering team. She is also VMware Certified Design Expert (VCDX) #236. Prior to joining the Infrastructure and Cloud Engineering team, Melissa was a systems engineer for NetApp and a VMware engineer for a number of enterprise environments. Melissa has bachelor of engineering and master of engineering degrees from Stevens Institute of Technology.

## Acknowledgements

The authors would like to acknowledge the following people for their support and contribution to this design:

- Chris O'Brien, Cisco Systems, Inc.
- John George, Cisco Systems, Inc.
- Dave Derry, NetApp
- Karthick Radhakrishnan, NetApp

## Version History

Version	Date	Document Version History
Version 1.0	April 2017	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

Copyright © 2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS

AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

### **Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.