



Deployment procedures

FlexPod

NetApp

January 21, 2025

This PDF was generated from https://docs.netapp.com/us-en/flexpod/express/express-c-series-aff220-deploy_cisco_nexus_3172p_deployment_procedure.html on January 21, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- Deployment procedures 1
 - Cisco Nexus 3172P deployment procedure 2
 - NetApp storage deployment procedure (part 1) 12
 - Cisco UCS C-Series rack server deployment procedure 36
 - NetApp AFF Storage Deployment Procedure (Part 2) 48
 - VMware vSphere 6.7 deployment procedure 48
 - Install VMware vCenter Server 6.7 64
 - Configure VMware vCenter Server 6.7 and vSphere clustering 72

Deployment procedures

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

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

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

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

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



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

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

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

The table below lists the VMware virtual machines created.

| Virtual machine description | Host name |
|-----------------------------|-----------|
| VMware vCenter Server | |

Cisco Nexus 3172P deployment procedure

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

Initial setup of Cisco Nexus 3172P switch

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



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

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

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

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

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

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

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

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

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

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

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

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

Enter the switch name : 3172P-B

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

Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>

Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>

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

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

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

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

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

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

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

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

NTP server IPv4 address : <<var_ntp_ip>>

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

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

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

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

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

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

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

6. Repeat this procedure for Cisco Nexus switch B.

Enable advanced features

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



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

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

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

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

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

```
port-channel load-balance src-dst ip-l4port
```

Perform global spanning-tree configuration

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

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

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

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

From configuration mode (`config t`), run the following commands to configure the default spanning-tree

options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

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

Define VLANs

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

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

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

Configure access and management port descriptions

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

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

Cisco Nexus Switch A

```

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

```

Cisco Nexus Switch B

```

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

```

Configure server and storage management interfaces

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

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

Cisco Nexus Switch A

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

Cisco Nexus Switch B

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

Perform virtual port channel global configuration

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

A vPC provides the following benefits:

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

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

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

Cisco Nexus Switch A

```

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

```

Cisco Nexus Switch B

```

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

```

Configure storage port channels

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

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

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

```

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

```

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

```

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

```



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

Configure server connections

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

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

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

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

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

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

Uplink into existing network infrastructure

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

completed.

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

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF2xx series installation

NetApp Hardware Universe

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

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

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

Controller AFF2XX Series prerequisites

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


Storage controllers

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

NetApp ONTAP 9.4

Configuration worksheet

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

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

The following table shows ONTAP 9.4 installation and configuration information.

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

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

Configure Node A

To configure node A, complete the following steps:

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

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

2. Allow the system to boot.

```
autoboot
```

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

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

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

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

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



This web server must be pingable.

```
<<var_url_boot_software>>
```

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

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

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

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

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

Configure Node B

To configure node B, complete the following steps:

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

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

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

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new

software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

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

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

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



This web server must be pingable.

```
<<var_url_boot_software>>
```

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

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

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

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

Continuation of Node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.4 boots on the node for the first time.



The node and cluster setup procedure has changed slightly in ONTAP 9.4. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up Node A.

```

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

```

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

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

3. Click Guided Setup to configure the cluster.

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

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Cluster Name

Nodes

Not sure all nodes have been discovered? [Refresh](#)

FAS2650

HA-PAR

FAS2650

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

Feature Licenses (Optional)

Cluster Base License is mandatory to add Feature Licenses.

5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter `<<var_clustermgmt_ip>>` in the Cluster Management IP Address field, `<<var_clustermgmt_mask>>` in the Netmask field, and `<<var_clustermgmt_gateway>>` in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter `<<var_nodeA_mgmt_ip>>` for node B.

- d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.

You can enter multiple DNS server IP addresses.

- e. Enter <<var_ntp_server_ip>> in the Primary NTP Server field.

You can also enter an alternate NTP server.

8. Configure the support information.

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

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

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? 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

- When indicated that the cluster configuration has completed, click Manage Your Cluster to configure the storage.

Continuation of storage cluster configuration

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

storage cluster.

Zero all spare disks

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

```
disk zerospares
```

Set on-board UTA2 ports personality

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

```
AFF A220::> ucadmin show
```

| Node | Adapter | Current Mode | Current Type | Pending Mode | Pending Type | Admin Status |
|------------|---------|--------------|--------------|--------------|--------------|--------------|
| AFF A220_A | 0c | fc | target | - | - | online |
| AFF A220_A | 0d | fc | target | - | - | online |
| AFF A220_A | 0e | fc | target | - | - | online |
| AFF A220_A | 0f | fc | target | - | - | online |
| AFF A220_B | 0c | fc | target | - | - | online |
| AFF A220_B | 0d | fc | target | - | - | online |
| AFF A220_B | 0e | fc | target | - | - | online |
| AFF A220_B | 0f | fc | target | - | - | online |

8 entries were displayed.

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

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

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

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



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

Rename management logical interfaces (LIFs)

To rename the management LIFs, complete the following steps:

1. Show the current management LIF names.

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

2. Rename the cluster management LIF.

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

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_node_mgmt_lif_AFF A220_B_1 -newname AFF A220-02_mgmt1
```

Set auto-revert on cluster management

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

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

Set up service processor network interface

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

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



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

Enable storage failover in ONTAP

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

1. Verify the status of storage failover.

```
storage failover show
```

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

2. Enable failover on one of the two nodes.

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

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

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

```
High Availability Configured: true
```

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



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

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

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

```
storage failover hwassist show
```

The message `Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.


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

Create jumbo frame MTU broadcast domain in ONTAP

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

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

Remove data ports from default broadcast domain

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

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

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

Disable flow control on UTA2 ports

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

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

Configure IFGRP LACP in ONTAP

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

From the cluster prompt, complete the following steps.

```

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

```

Configure jumbo frames in NetApp ONTAP

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

```

AFF A220::> network port modify -node node_A -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

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

```

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

```

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

```

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

```

3. Create MGMT-VLAN ports.

```

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

```

Create aggregates in ONTAP

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

To create aggregates, run the following commands:

```

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

```

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

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

The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1`_`nodeA` is online.

Configure time zone in ONTAP

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

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



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

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

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

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

2. Configure SNMP traps to send to remote hosts.

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

Configure SNMPv1 in ONTAP

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

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



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

Configure SNMPv3 in ONTAP

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

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

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

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

Configure AutoSupport HTTPS in ONTAP

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

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

Create a storage virtual machine

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

1. Run the `vserver create` command.

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

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

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

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

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,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 that NFS has been configured.

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



Commands are prefaced by `vserver` in the command line because storage virtual machines were previously called servers.

Configure NFSv3 in ONTAP

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

| Detail | Detail value |
|----------------------------|---------------------------|
| ESXi host A NFS IP address | <<var_esxi_hostA_nfs_ip>> |
| ESXi host B NFS IP address | <<var_esxi_hostB_nfs_ip>> |

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

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

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

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

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



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

Create iSCSI service in ONTAP

To create the iSCSI service, complete the following step:

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

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

Create load-sharing mirror of SVM root volume in ONTAP

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

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

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

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

3. Create the mirroring relationships.

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

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

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

Configure HTTPS access in ONTAP

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

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

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

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


```
security certificate show
```

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

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

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

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

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

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

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

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

```
system services web modify -external true -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>>
```



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

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

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

Create a NetApp FlexVol volume in ONTAP

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

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

Enable deduplication in ONTAP

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

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

Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

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



When adding an extra Cisco UCS C-Series server, an extra boot LUN must be created.

Create iSCSI LIFs in ONTAP

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

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

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

```

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

```

Create NFS LIFs in ONTAP

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

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

1. Create an NFS LIF.

```

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

```

Add infrastructure SVM administrator

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

| Detail | Detail Value |
|------------------------|--------------------------|
| Vsmgmt IP | <<var_svm_mgmt_ip>> |
| Vsmgmt network mask | <<var_svm_mgmt_mask>> |
| Vsmgmt default gateway | <<var_svm_mgmt_gateway>> |

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

1. Run the following command:

```

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

```



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

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

```

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

```

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

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

[Next: Cisco UCS C-Series Rack Server Deployment Procedure](#)

Cisco UCS C-Series rack server deployment procedure

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

Perform initial Cisco UCS C-Series standalone server setup for Cisco Integrated Management Server

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

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


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



The CIMC version used in this validation is CIMC 3.1.3(g).

All servers

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



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
Dedicated:      [X]          NIC redundancy
Shared LOM:     [ ]          None: [X]
Cisco Card:     [ ]          Active-standby: [ ]
Riser1:        [ ]          Active-active: [ ]
Riser2:        [ ]          VLAN (Advanced)
MLom:          [ ]          VLAN enabled: [ ]
Shared LOM Ext: [ ]          VLAN ID: 1
Priority: 0
IP (Basic)
IPv4: [X] IPv6: [ ]
DHCP enabled [ ]
CIMC IP: 10.61.185.215
Prefix/Subnet: 255.255.255.0
Gateway: 10.61.185.1
Pref DNS Server: 0.0.0.0
Smart Access USB
Enabled [ ]
*****
<Up/Down>Selection <F10>Save <Space>Enable/Disable <F5>Refresh <ESC>Exit
<F1>Additional settings
```

4. Press F1 to see additional settings.

- Common properties:
 - Host name: <<esxi_host_name>>
 - Dynamic DNS: []
 - Factory defaults: Leave cleared.
- Default user (basic):
 - Default password: <<admin_password>>
 - Reenter password: <<admin_password>>
 - Port properties: Use default values.
 - Port profiles: Leave cleared.


```

Cisco IMC Configuration Utility Version 2.0  Cisco Systems, Inc.
*****
Common Properties
  Hostname:      CIMC-Tiger-02
  Dynamic DNS:   [X]
  DDNS Domain:
FactoryDefaults
  Factory Default:      [ ]
Default User(Basic)
  Default password:      -
  Reenter password:
Port Properties
  Auto Negotiation:      [X]
                                Admin Mode      Operation Mode
  Speed[1000/100/10Mbps]:      Auto              1000
  Duplex mode[half/full]:      Auto              full
Port Profiles
  Reset:                  [ ]
  Name:
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F2>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 VIC1387 is used for iSCSI boot.

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



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

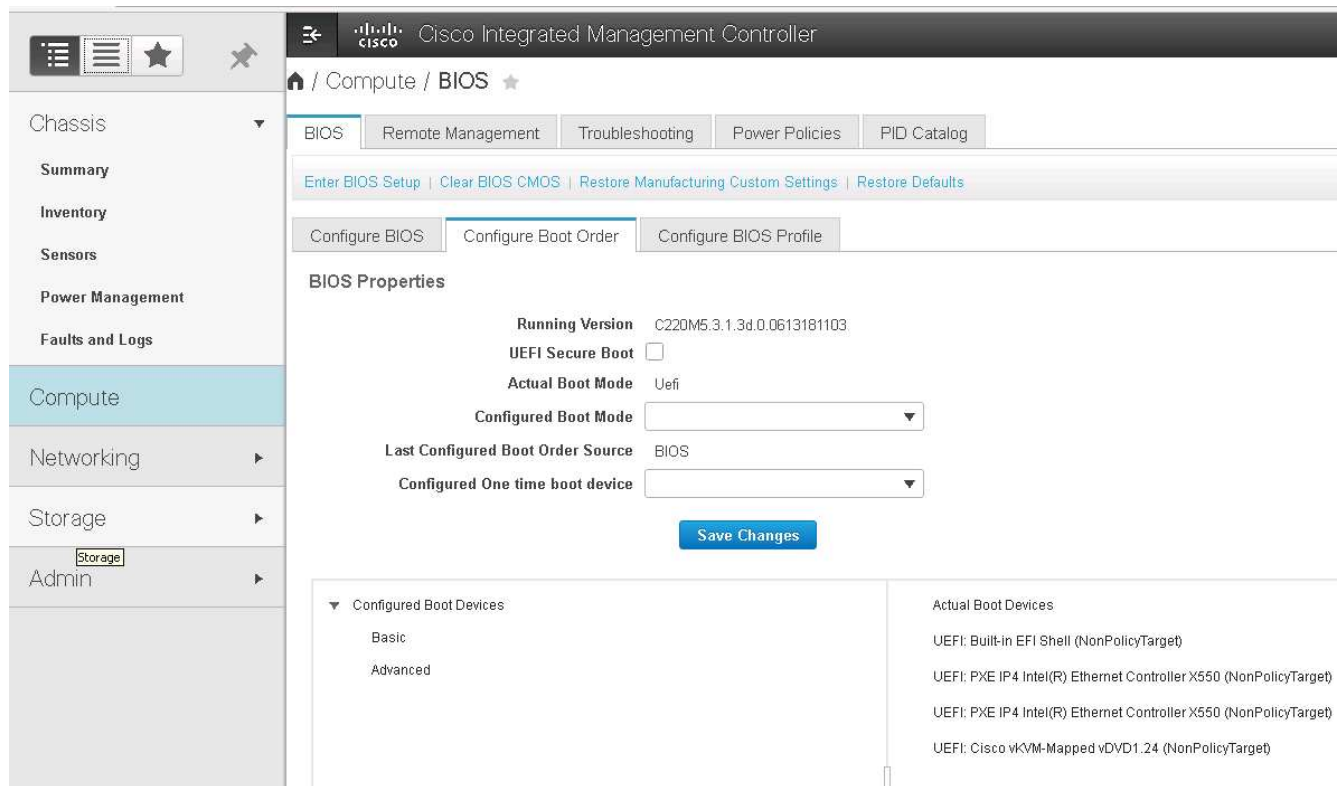
| Detail | Detail value |
|-----------------------------------|----------------------------------|
| ESXi host initiator A name | <<var_ucs_initiator_name_A>> |
| ESXi host iSCSI-A IP | <<var_esxi_host_iscsiA_ip>> |
| ESXi host iSCSI-A network mask | <<var_esxi_host_iscsiA_mask>> |
| ESXi host iSCSI A default gateway | <<var_esxi_host_iscsiA_gateway>> |
| ESXi host initiator B name | <<var_ucs_initiator_name_B>> |
| ESXi host iSCSI-B IP | <<var_esxi_host_iscsiB_ip>> |
| ESXi host iSCSI-B network mask | <<var_esxi_host_iscsiB_mask>> |
| ESXi host iSCSI-B gateway | <<var_esxi_host_iscsiB_gateway>> |

| Detail | Detail value |
|-------------------------|--------------|
| IP address iscsi_lif01a | |
| IP address iscsi_lif02a | |
| IP address iscsi_lif01b | |
| IP address iscsi_lif02b | |
| Infra_SVM IQN | |

Boot order configuration

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

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



3. Configure the following devices by clicking the device under Add Boot Device, and going to the Advanced tab.
 - Add Virtual Media
 - Name: KVM-CD-DVD
 - Subtype: KVM MAPPED DVD
 - State: Enabled
 - Order: 1
 - Add iSCSI Boot.
 - Name: iSCSI-A

- State: Enabled
- Order: 2
- Slot: MLOM
- Port: 0
- Click Add iSCSI Boot.
 - Name: iSCSI-B
 - State: Enabled
 - Order: 3
 - Slot: MLOM
 - Port: 1

4. Click Add Device.

5. Click Save Changes and then click Close.

Configure Boot Order

Configured Boot Level: Advanced

Basic Advanced

Add Boot Device

- Add Local HDD
- Add PXE Boot
- Add SAN Boot
- Add iSCSI Boot
- Add USB
- Add Virtual Media
- Add PCHStorage
- Add UEFISHELL
- Add SD Card
- Add NVME
- Add Local CDD

Advanced Boot Order Configuration

Selected 1 / Total 3

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

Save Changes Reset Values Close

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

Disable RAID controller (if present)

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

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

| | | | | | |
|------|-------------------|-----------------|----------------|-------------|-------------------|
| BIOS | Remote Management | Troubleshooting | Power Policies | PID Catalog | |
| I/O | Server Management | Security | Processor | Memory | Power/Performance |

Note: Default values are shown in bold.

Reboot Host Immediately: ☒

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

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

Configure Cisco VIC1387 for iSCSI boot

The following configuration steps are for the Cisco VIC 1387 for iSCSI boot.

Create iSCSI vNICs

1. Click Add to create a vNIC.
2. In the Add vNIC section, enter the following settings:
 - Name: iSCSI-vNIC-A
 - MTU: 9000
 - Default VLAN: <<var_iscsi_vlan_a>>
 - VLAN Mode: TRUNK
 - Enable PXE boot: Check

▼ vNIC Properties

▼ General

Name:

iSCSI-vNIC-A

CDN:

VIC-MLOM-iSCSI-vNIC-A

MTU:

9000

(1500 - 9000)

Uplink Port:

0

MAC Address:

☐ Auto
 ☒ 70:69:5A:C0:98:ED

Class of Service:

0

(0 - 6)

Trust Host CoS:

☒

PCI Order:

4

(0 - 5)

Default VLAN:

☐ None
 ☒ 3439

VLAN Mode:

Trunk

Rate Limit:

☒ OFF
 ☐

(1 - 1000)

Channel Number:

N/A

(0 - 1)

PCI Link:

0

(0 - 1)

Enable NVGRE:

☐

Enable VXLAN:

☐

Advanced Filter:

☐

Port Profile:

N/A

Enable PXE Boot:

☒

Enable VMQ:

☐

Enable aRFS:

☐

Enable Uplink Failover:

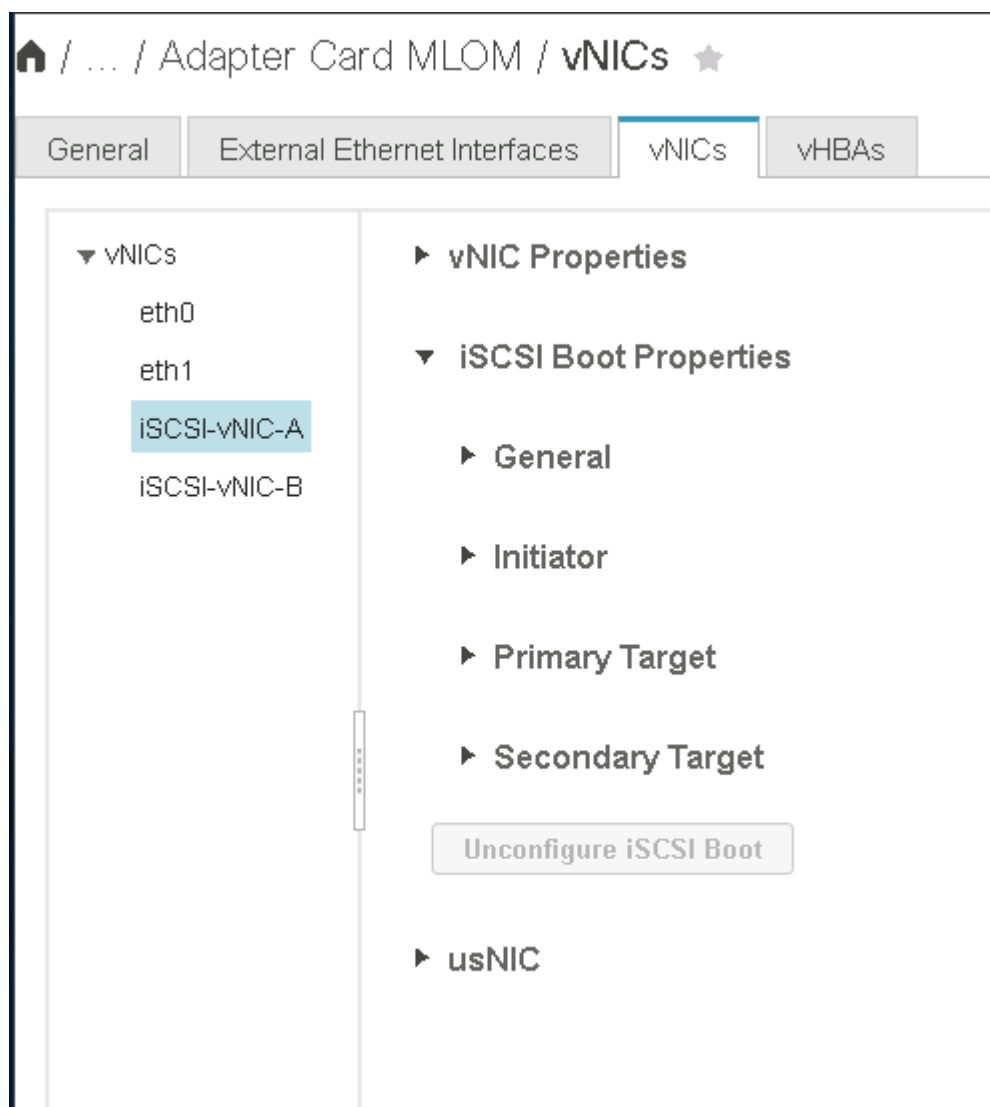
☐

Failback Timeout:

N/A

(0 - 600)

3. Click Add vNIC and then click OK.
4. Repeat the process to add a second vNIC.
 - a. Name the vNIC iSCSI-vNIC-B.
 - b. Enter <<var_iscsi_vlan_b>> as the VLAN.
 - c. Set the uplink port to 1.
5. Select the vNIC iSCSI-vNIC-A on the left.



6. Under iSCSI Boot Properties, enter the initiator details:
 - Name: <<var_ucsa_initiator_name_a>>
 - IP address: <<var_esxi_hostA_iscsiA_ip>>
 - Subnet mask: <<var_esxi_hostA_iscsiA_mask>>
 - Gateway: <<var_esxi_hostA_iscsiA_gateway>>

▼ vNICs
 eth0
 eth1
 iSCSI-v
 iSCSI-v

▼ iSCSI Boot Properties

▶ General

▼ Initiator

Name: (0 - 233) chars
 Initiator Priority:

IP Address:
 Secondary DNS:

Subnet Mask:
 TCP Timeout:

Gateway:
 CHAP Name:

Primary DNS:
 CHAP Secret:

▶ Primary Target

▶ Secondary Target

7. Enter the primary target details.

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

8. Enter the secondary target details.

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

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



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

General
External Ethernet Interfaces
vNICs
vHBAs

vNICs
eth0
eth1
iSCSI-v
iSCSI-v

Initiator

Primary Target

Name: iqn.1992-08.com.netapp:sn.7e560f73a51 (0 - 233) chars
IP Address: 172.21.246.16
TCP Port: 3260

Boot LUN: 0
CHAP Name:
CHAP Secret:

Secondary Target

Name: iqn.1992-08.com.netapp:sn.7e560f73a51 (0 - 233) chars
IP Address: 172.21.246.18
TCP Port: 3260

Boot LUN: 0
CHAP Name:
CHAP Secret:

Unconfigure iSCSI Boot

9. Click Configure iSCSI.
10. Select the vNIC `iSCSI-vNIC-B` and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.
11. Repeat the process to configure `iSCSI-vNIC-B`.
12. Enter the initiator details.
 - Name: `<<var_ucsa_initiator_name_b>>`
 - IP address: `<<var_esxi_hostb_iscsib_ip>>`
 - Subnet mask: `<<var_esxi_hostb_iscsib_mask>>`
 - Gateway: `<<var_esxi_hostb_iscsib_gateway>>`
13. Enter the primary target details.
 - Name: IQN number of infra-SVM
 - IP address: IP address of `iscsi_lif01b`
 - Boot LUN: 0
14. Enter the secondary target details.
 - Name: IQN number of infra-SVM
 - IP address: IP address of `iscsi_lif02b`
 - Boot LUN: 0

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



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

15. Click Configure iSCSI.
16. Repeat this process to configure iSCSI boot for Cisco UCS server B.

Configure vNICs for ESXi

- 1. From the CIMC interface browser window, click Inventory and then click Cisco VIC adapters on the right pane.
- 2. Under Adapter Cards, select Cisco UCS VIC 1387 and then select the vNICs underneath.

Home / ... / Adapter Card

MLOM / vNICs ★

Refresh | Host Power | Launch KVM | Ping | CIMC Reboot | Locat

General | External Ethernet Interfaces | vNICs | vHBAs

▼ vNICs

eth0

eth1

iSCSI-v

iSCSI-v

Host Ethernet Interfaces

Selected 0

Add vNIC | Clone vNIC | Delete vNICs

| | Name | CDN | MAC Address | MTU | usNIC | Uplink Port | CoS | VLAN | VLAN Mode |
|--------------------------|------------|------------|-------------------|------|-------|-------------|-----|------|-----------|
| <input type="checkbox"/> | eth0 | VIC-MLO... | 70:69:5A:C0:98:49 | 1500 | 0 | 0 | 0 | NONE | TRUNK |
| <input type="checkbox"/> | eth1 | VIC-MLO... | 70:69:5A:C0:98:4A | 1500 | 0 | 1 | 0 | NONE | TRUNK |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4D | 9000 | 0 | 0 | 0 | 3439 | TRUNK |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4E | 9000 | 0 | 1 | 0 | 3440 | TRUNK |

- 3. Select eth0 and click Properties.
- 4. Set the MTU to 9000. Click Save Changes.

General

External Ethernet Interfaces

vNICs

vHBAs

▼ vNICs

eth0

eth1

ISCSI-v

ISCSI-v

Name:

eth0

CDN:

VIC-MLOM-eth0

MTU:

9000

(1500 - 9000)

Uplink Port:

0

MAC Address:

☐ Auto
 ☒ 70:69:5A:C0:98:49

Class of Service:

0

(0 - 6)

Trust Host CoS:

☐

PCI Order:

0

(0 - 5)

Default VLAN:

☒ None
 ☐ ?

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

[/ ... / Adapter Card MLOM / vNICs](#) ★

General

External Ethernet Interfaces

vNICs

vHBAs

▼ vNICs

eth0

eth1

ISCSI-vNIC-A

ISCSI-vNIC-B

Host Ethernet Interfaces

Add vNIC

Clone vNIC

Delete vNICs

| | Name | CDN | MAC Address | MTU | usNIC | Uplink Port |
|--------------------------|------------|------------|-------------------|------|-------|-------------|
| <input type="checkbox"/> | eth0 | VIC-MLO... | 70:69:5A:C0:98:49 | 9000 | 0 | 0 |
| <input type="checkbox"/> | eth1 | VIC-MLO... | 70:69:5A:C0:98:4A | 9000 | 0 | 1 |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4D | 9000 | 0 | 0 |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4E | 9000 | 0 | 1 |



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

NetApp AFF Storage Deployment Procedure (Part 2)

ONTAP SAN boot storage setup

Create iSCSI igroups

To create igroups, complete the following step:

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

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

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



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

Map boot LUNs to igroups

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

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



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

[Next: VMware vSphere 6.7 Deployment Procedure.](#)

VMware vSphere 6.7 deployment procedure

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

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



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


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

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

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

All hosts

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



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



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

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

Install VMware ESXi

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

Download ESXi 6.7 Cisco custom image

1. Navigate to the [VMware vSphere download page](#) for custom ISOs.
2. Click Go to Downloads next to the Cisco Custom Image for ESXi 6.7 GA Install CD.
3. Download the Cisco Custom Image for ESXi 6.7 GA Install CD (ISO).

All hosts

1. When the system boots, the machine detects the presence of the VMware ESXi installation media.
2. Select the VMware ESXi installer from the menu that appears.

The installer loads. This takes several minutes.

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



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

Set up VMware ESXi host management networking


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

All hosts

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



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



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

Configure ESXi host

You need the information in the following table to configure each ESXi host.

| Detail | Value |
|----------------|-------|
| ESXi host name | |

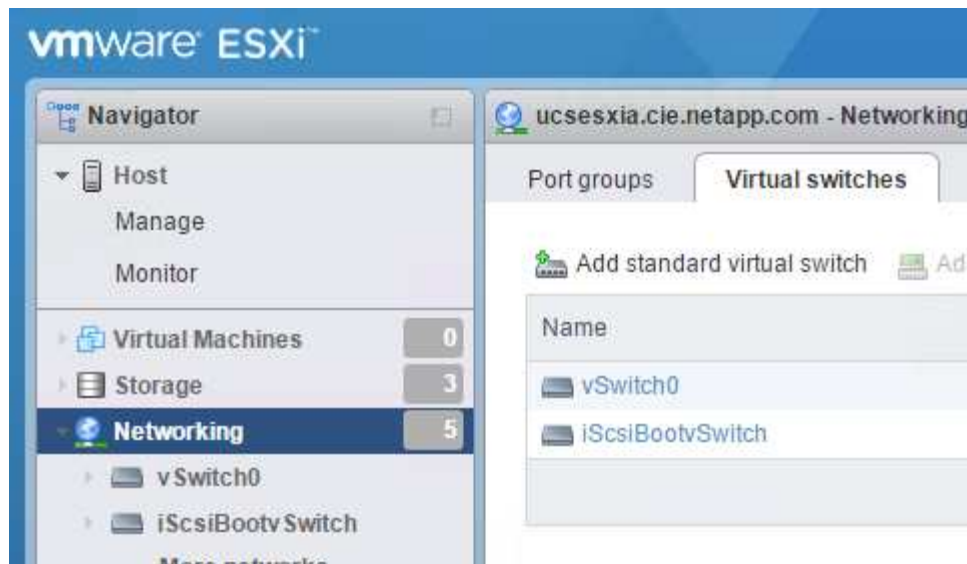
| Detail | Value |
|------------------------------|-------|
| ESXi host management IP | |
| ESXi host management mask | |
| ESXi host management gateway | |
| ESXi host NFS IP | |
| ESXi host NFS mask | |
| ESXi host NFS gateway | |
| ESXi host vMotion IP | |
| ESXi host vMotion mask | |
| ESXi host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

Log in to ESXi host

1. Open the host's management IP address in a web browser.
2. Log in to the ESXi host using the root account and the password you specified during the install process.
3. Read the statement about the VMware Customer Experience Improvement Program. After selecting the proper response, click OK.

Configure iSCSI boot

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



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



Vmnic2 and vmnic3 are used for iSCSI boot in this configuration. If you have additional NICs in your ESXi host, you might have different vmnic numbers. To confirm which NICs are used for iSCSI boot, match the MAC addresses on the iSCSI vNICs in CIMC to the vmnics in ESXi.

9. In the center pane, select the VMkernel NICs tab.
10. Select Add VMkernel NIC.
 - Specify a new port group name of iScsiBootPG-B.
 - Select iScsiBootvSwitch-B for the virtual switch.
 - Enter <<iscsib_vlan_id>> for the VLAN ID.
 - Change the MTU to 9000.
 - Expand IPv4 Settings.
 - Select Static Configuration.
 - Enter <<var_hosta_iscsib_ip>> for Address.
 - Enter <<var_hosta_iscsib_mask>> for Subnet Mask.
 - Click Create.

| | |
|-----------------|--|
| Port group | New port group ▼ |
| New port group | iScsiBootPG-B |
| Virtual switch | iScsiBootvSwitch-B ▼ |
| VLAN ID | 3440 |
| MTU | 9000 |
| IP version | IPv4 only ▼ |
| ▼ IPv4 settings | |
| Configuration | <input 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

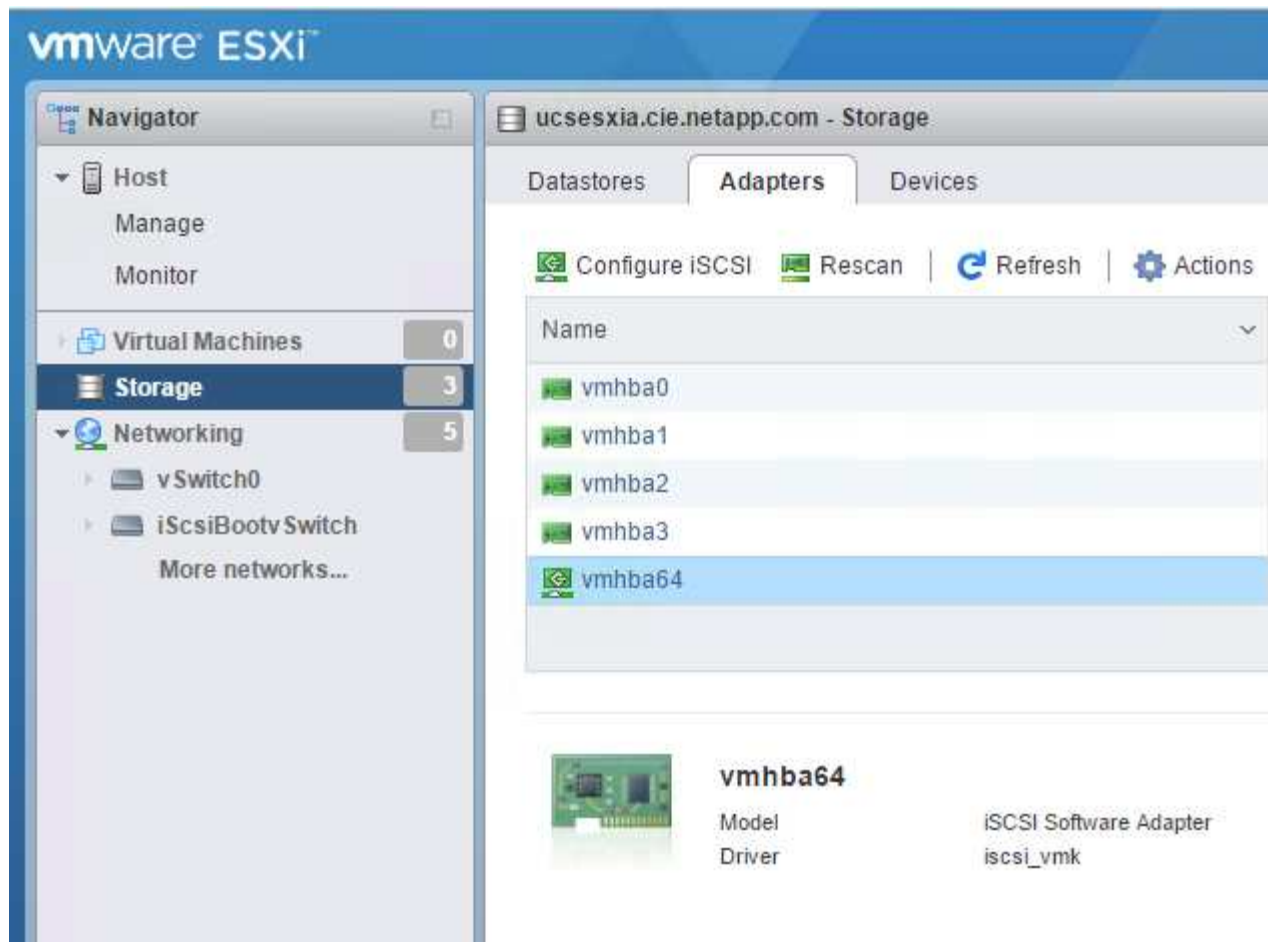


Set the MTU to 9000 on iScsiBootPG- A.

Configure iSCSI multipathing

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

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



3. Under Dynamic Targets, click Add Dynamic Target.

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> Add port binding Remove port binding </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> Add static target Remove static target Edit settings <input type="text" value="Search"/> </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> Add dynamic target Remove dynamic target Edit settings <input type="text" value="Search"/> </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 | | | | | | | | | |

4. Enter the IP address `iscsi_lif01a`.
 - Repeat with the IP addresses `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
 - Click Save Configuration.

Dynamic targets

Add dynamic target
 Remove dynamic target
 Edit settings

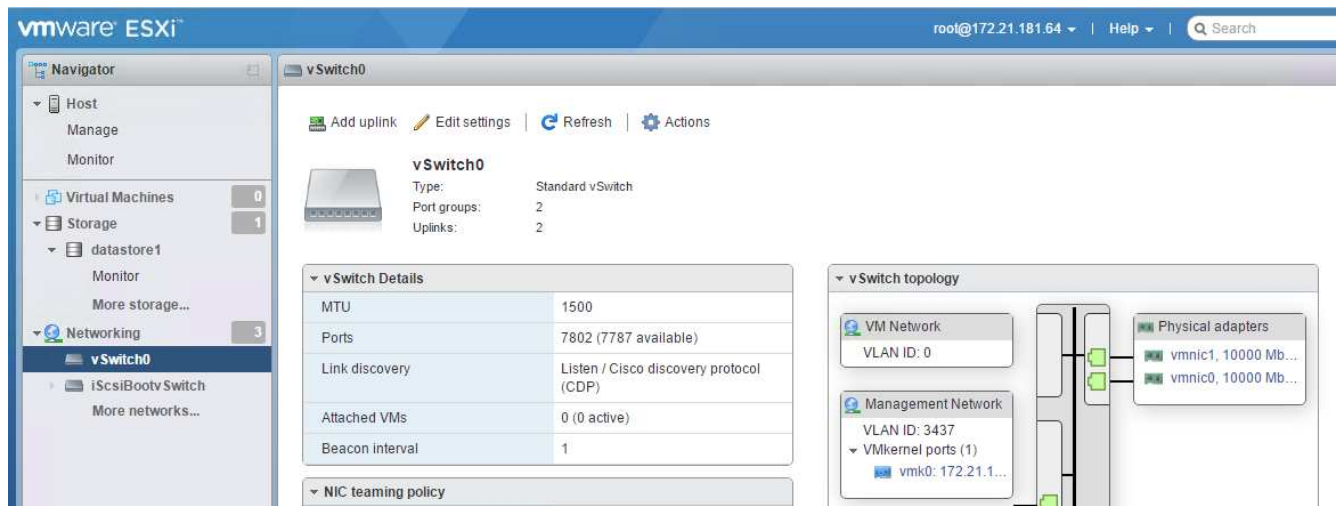
| Address | Port |
|---------------|------|
| 172.21.183.33 | 3260 |
| 172.21.183.34 | 3260 |
| 172.21.184.33 | 3260 |
| 172.21.184.34 | 3260 |



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

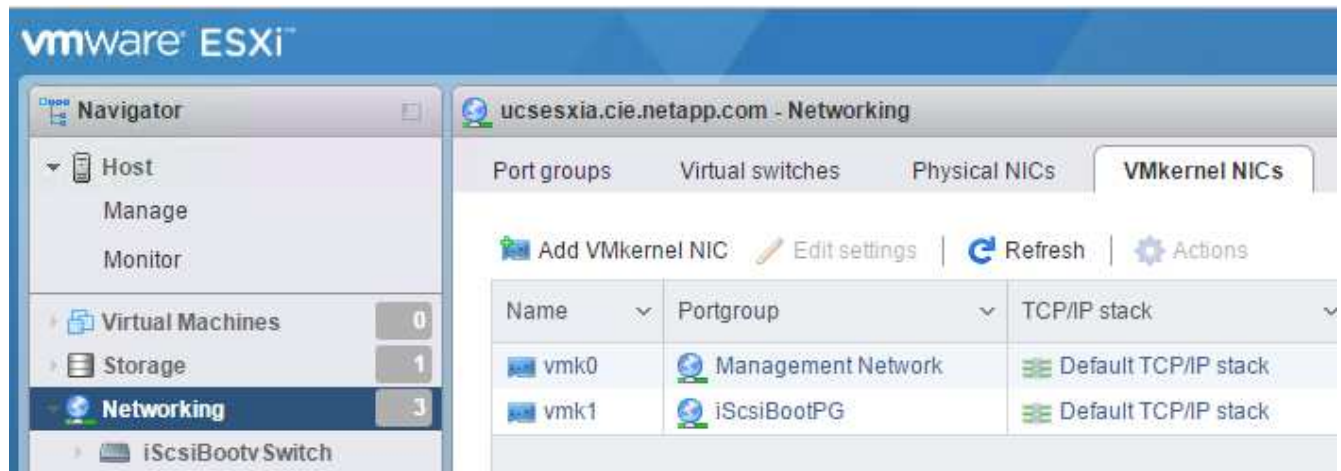
Configure port groups and VMkernel NICs

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.
 - Name the port group MGMT-Network.
 - Enter <<mgmt_vlan>> for the VLAN ID.
 - Make sure that vSwitch0 is selected.
 - Click Add.

5. Click the VMkernel NICs tab.



6. Select Add VMkernel NIC.
- Select New Port Group.
 - Name the port group NFS-Network.
 - Enter <<nfs_vlan_id>> for the VLAN ID.
 - Change the MTU to 9000.
 - Expand IPv4 Settings.
 - Select Static Configuration.
 - Enter <<var_hosta_nfs_ip>> for Address.
 - Enter <<var_hosta_nfs_mask>> for Subnet Mask.
 - Click Create.

| | |
|-----------------|--|
| 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.

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

Create Cancel

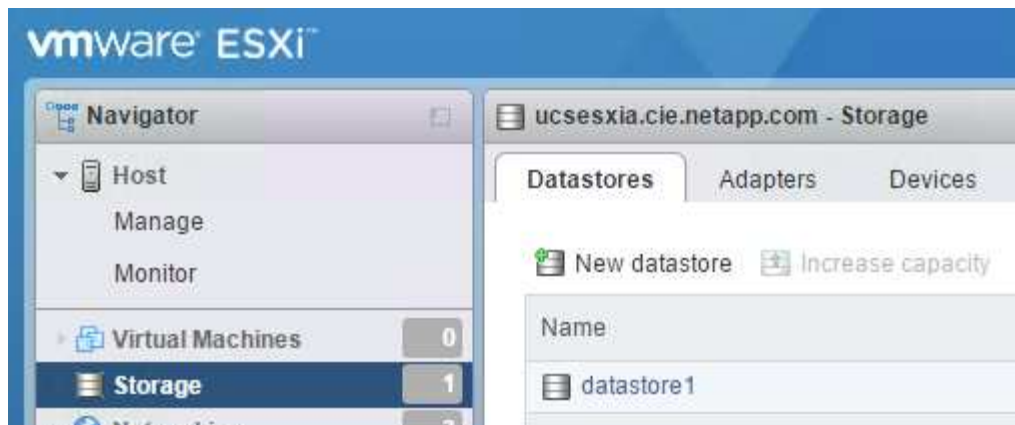


There are many ways to configure ESXi networking, including by using the VMware vSphere distributed switch if your licensing allows it. Alternative network configurations are supported in FlexPod Express if they are required to meet business requirements.


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 click New Datastore.



2. Select Mount NFS Datastore.



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

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

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

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

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

- Make sure that NFS 3 is selected.


Configure NTP

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

1. Click Manage in the left navigation pane. Select System in the right pane and then click Time & Date.




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



Move the virtual machine swap-file location

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

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



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



3. Click Save.

Install the NetApp NFS Plug-in 1.0.20 for VMware VAAI

To install the NetApp NFS Plug-in 1.0.20 for VMware VAAI, complete the following steps.

1. Enter the following commands to verify that VAAI is enabled:

```
esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
```

If VAAI is enabled, these commands produce the following output:

```
~ # esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
Value of HardwareAcceleratedMove is 1
~ # esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
Value of HardwareAcceleratedInit is 1
```

2. If VAAI is not enabled, enter the following commands to enable VAAI:

```
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedInit
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
```

These commands produce the following output:

```
~ # esxcfg-advcfg -s 1 /Data Mover/HardwareAcceleratedInit
Value of HardwareAcceleratedInit is 1
~ # esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
Value of HardwareAcceleratedMove is 1
```

3. Download the NetApp NFS Plug-in for VMware VAAI:
 - a. Go to the [software download page](#).
 - b. Scroll down and click NetApp NFS Plug-in for VMware VAAI.
 - c. Select the ESXi platform.
 - d. Download either the offline bundle (.zip) or online bundle (.vib) of the most recent plug-in.
4. Install the plug-in on the ESXi host by using the ESX CLI.
5. Reboot the ESXi host.

```
[root@vm-host-infra-04:~] ls /vmfs/volumes/datastore1/NetAppNasPlugin.vib
/vmfs/volumes/datastore1/NetAppNasPlugin.vib
[root@vm-host-infra-04:~] esxcli software vib install -v /vmfs/volumes/datastore1/NetAppNasPlugin.vib
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: NetApp_bootbank_NetAppNasPlugin_1.1.2-3
  VIBs Removed:
  VIBs Skipped:
[root@vm-host-infra-04:~] █
```

Next: [Install VMware vCenter Server 6.7](#)

Install VMware vCenter Server 6.7


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



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

Download the VMware vCenter server appliance

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



2. Download the VCSA from the VMware site.



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

3. Mount the ISO image.
4. Navigate to the vcsa-ui-installer> win32 directory. Double-click installer.exe.
5. Click Install.
6. Click Next on the Introduction page.
7. Accept the end-user license agreement.
8. Select Embedded Platform Services Controller as the deployment type.

Installer

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

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

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 | 172.21.246.25 | i |
| HTTPS port | 443 | |
| User name | root | i |
| Password | ***** | |

CANCEL

BACK

NEXT

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

VM

Install

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

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name

tigervcsa

Set root password

.....

Confirm root password

.....

CANCEL

BACK

NEXT

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

VM

Install

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

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Select deployment size

Select the deployment size for this vCenter Server with an Embedded Platform Services Controller.

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

Deployment size

Tiny

Storage size

Default

Resources required for different deployment sizes

| Deployment Size | vCPUs | Memory (GB) | Storage (GB) | Hosts (up to) | VMs (up to) |
|-----------------|-------|-------------|--------------|---------------|-------------|
| Tiny | 2 | 10 | 300 | 10 | 100 |
| Small | 4 | 16 | 340 | 100 | 1000 |
| Medium | 8 | 24 | 525 | 400 | 4000 |
| Large | 16 | 32 | 740 | 1000 | 10000 |
| X-Large | 24 | 48 | 1180 | 2000 | 35000 |

CANCEL

BACK

NEXT

12. Select the infra_datastore_1 datastore. Click Next.

vm

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

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Select datastore

Select the storage location for this appliance

☒ Install on an existing datastore accessible from the target host

| Name | Type | Capacity | Free | Provisioned | Thin Provisioning |
|-------------------|------|----------|-----------|-------------|-------------------|
| infra_datastore_1 | NFS | 500 GB | 499.98 GB | 18.38 MB | Supported |
| infra_swap | NFS | 100 GB | 99.99 GB | 10.95 MB | Supported |

2 items

☒ Enable Thin Disk Mode ⓘ

☐ Install on a new vSAN cluster containing the target host ⓘ

CANCEL

BACK

NEXT

13. Enter the following information in the Configure network settings page and click Next.

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

14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.

vCenter Server Appliance Installer

Installer

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

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

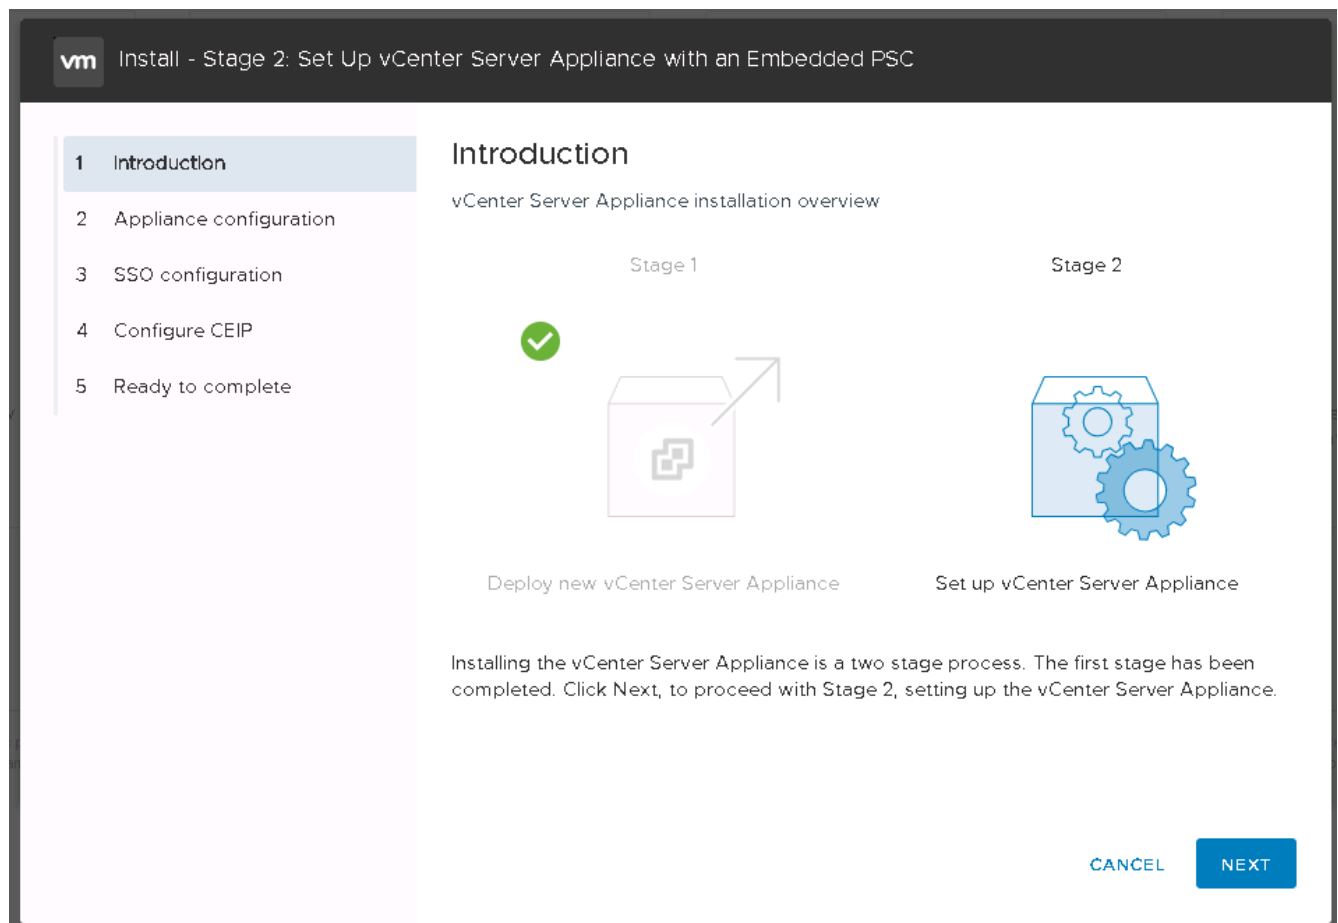
Configure network settings

| | | |
|------------------------------|-----------------------------|---|
| IP version | IPv4 | |
| IP assignment | static | |
| FQDN | tigervcsa.cle.netapp.com | i |
| IP address | 172.21.246.41 | |
| Subnet mask or prefix length | 255.255.255.0 | i |
| Default gateway | 172.21.246.1 | |
| DNS servers | 10.61.184.251,10.61.184.252 | |
| Common Ports | | |
| HTTP | 80 | |
| HTTPS | 443 | |

CANCEL BACK NEXT

The VCSA installs now. This process takes several minutes.

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



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

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

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

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

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

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

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

The appliance setup continues. This takes several minutes.

A message appears indicating that the setup was successful.

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

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

Configure VMware vCenter Server 6.7 and vSphere clustering

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

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

Create vSphere cluster

Complete the following steps to create a vSphere cluster:

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

New Cluster | FlexPod

Name

Tiger3

Location

FlexPod

> DRS

☒ Turn ON

> vSphere HA

☒ Turn ON

> EVC

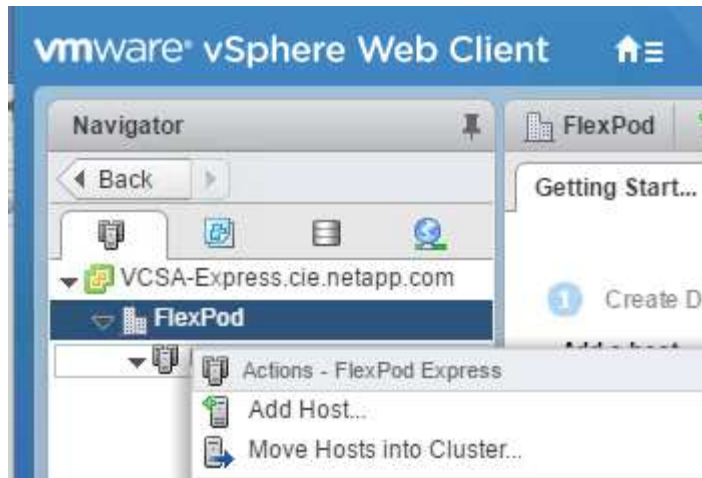
Disable

CANCEL

OK

Add ESXi hosts to cluster

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



2. To add an ESXi host to the cluster, complete the following steps:
 - a. Enter the IP or FQDN of the host. Click Next.
 - b. Enter the root user name and password. Click Next.
 - c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
 - d. Click Next on the Host Summary page.
 - e. Click the green + icon to add a license to the vSphere host.



This step can be completed later if desired.

- f. Click Next to leave lockdown mode disabled.
 - g. Click Next at the VM location page.
 - h. Review the Ready to Complete page. Use the back button to make any changes or select Finish.
3. Repeat steps 1 and 2 for Cisco UCS host B. This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on ESXi hosts

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

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

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

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

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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