



NetApp Verified Architecture

FlexPod Datacenter with NetApp MetroCluster NVA Deployment

Aaron Kirk and Arvind Ramakrishnan, NetApp
August 2016 | NVA-0030-DEPLOY | Version 1.0

Reviewed by



TABLE OF CONTENTS

1	Program Summary	5
2	Solution Overview	5
2.1	Solution Technology	5
2.2	Use Case Summary	7
3	Solution Architecture	8
4	Technology Requirements	8
4.1	Hardware Requirements	8
4.2	Software Requirements	9
4.3	Infrastructure VLANs	10
4.4	Configuration Variables	10
5	Deployment Procedures	10
5.1	Configuration Guidelines	10
5.2	Configure Cisco Nexus 9372PX Switches	11
5.3	Configure OTV to Span VLANs Between Two Sites	17
5.4	Set Shelf IDs	26
5.5	Set Up ATTO FibreBridge 7500N	26
5.6	Set Up Cisco MDS 9396s	26
5.7	Controllers Setup	35
5.8	Server Configuration	47
5.9	Storage Configuration—Boot LUNs and Igroups	91
5.10	VMware vSphere 6.0 Setup	92
6	Solution Verification	123
6.1	Single Blade Failure	123
6.2	Single Host Isolation	124
6.3	Fabric Interconnect Reboot	126
6.4	Storage Isolation	126
6.5	Full Data Center Isolation	127
6.6	Storage Degradation in a Single Data Center	129
6.7	Planned Switchover and Switchback	130
6.8	Full Storage Failure in a Single Data Center	132
6.9	Volume Failure in a Single Data Center	134
6.10	Loss of a Data Center	134
6.11	Full Compute Failure in a Data Center	136

7 Conclusion	137
Acknowledgements	137
References.....	138
Appendixes.....	138
Configuration Variables	138
Cabling Details	140

LIST OF TABLES

Table 1) Hardware requirements.....	9
Table 2) Software requirements.....	9
Table 3) Software requirements.....	10
Table 4) Site A—Cisco Nexus 7004 OTV VDC connectivity.....	18
Table 5) Site A—Cisco Nexus 7004 LAN VDC connectivity.....	18
Table 6) Site B—Cisco Nexus 7004 OTV VDC connectivity.....	18
Table 7) Site B—Cisco Nexus 7004 LAN VDC connectivity.....	18
Table 8) iSCSI LIFs for iSCSI IQN.....	91
Table 9) vNIC iSCSI IQNs for fabric A and fabric B.....	91
Table 10) Single host failure in site A data center.....	123
Table 11) Single host isolation in site A data center.....	124
Table 12) Fabric interconnect test details.....	126
Table 13) Storage partition.....	126
Table 14) Data center partition.....	127
Table 15) Disk shelf failure in site A data center.....	129
Table 16) Planned switchover of storage.....	130
Table 17) Full storage failure in site A data center.....	132
Table 18) Permanent device loss.....	134
Table 19) Full compute failure in site A data center.....	135
Table 20) Loss of site A data center.....	136
Table 21) Configuration variables.....	138
Table 22) Cisco UCS 5108 cabling.....	140
Table 23) Cisco UCS 6248UP cabling.....	141
Table 24) Cisco Nexus 9372 cabling.....	141
Table 25) NetApp AFF8040 cabling.....	142
Table 26) Cisco MDS 9396s cabling.....	142
Table 27) ATTO FibreBridge 7500N cabling.....	143

LIST OF FIGURES

Figure 1) FlexPod component families.....	6
---	---

Figure 2) Cisco OTV.....	7
Figure 3) FlexPod Datacenter with NetApp MetroCluster single site.....	8
Figure 4) Cisco Nexus 9K and 7K connectivity for OTV.....	17
Figure 5) MetroCluster cabling diagram for a single site.....	31
Figure 6) Ports for FCVI zone.....	32
Figure 7) Ports for storage zone #1.....	32
Figure 8) Ports for storage zone #2.....	33
Figure 9) Ports for storage zone #3.....	33
Figure 10) Ports for storage zone #4.....	33
Figure 11) Site A, node 1 disk assignment.....	35
Figure 12) Single host failure in site A data center.....	123
Figure 13) Single host isolation in site A data center.....	125
Figure 14) Storage partition.....	127
Figure 15) Data center partition.....	128
Figure 16) Disk shelf failure in site A data center.....	130
Figure 17) Planned switchover of storage.....	131
Figure 18) Full storage failure in site A data center.....	133
Figure 19) Permanent device loss.....	134
Figure 20) Full compute failure in site A data center.....	135
Figure 21) Loss of site A data center.....	137

1 Program Summary

The NetApp Verified Architecture (NVA) program offers customers a verified architecture for NetApp® solutions. An NVA provides you with a NetApp solution architecture that:

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

This NVA deployment guide describes the deployment steps required for FlexPod® Datacenter with NetApp MetroCluster™ technology. NetApp MetroCluster provides continuous availability for business-critical applications at lower cost and complexity.

2 Solution Overview

In today's data center, unplanned downtime of business-critical infrastructure causes loss of revenues and user productivity. These losses, in addition to associated costs such as detection, recovery, equipment, and IT productivity loss, cause unplanned downtime to cost up to \$7,800 a minute. This fact drives the need to ensure that mission-critical applications are protected from data loss and downtime by infrastructure that continues to be available through large power outages.

NetApp MetroCluster provides a zero recovery point objective along with a near-zero recovery time objective to allow continuous data availability for mission-critical applications. FlexPod Datacenter with NetApp MetroCluster combines the continuous availability of MetroCluster with FlexPod Datacenter. This solution contains Cisco Unified Computing System (Cisco UCS), Cisco Nexus networking, and VMware vSphere in a validated design that reduces the risk and complexity of deploying business-critical infrastructure.

2.1 Solution Technology

FlexPod Data Center

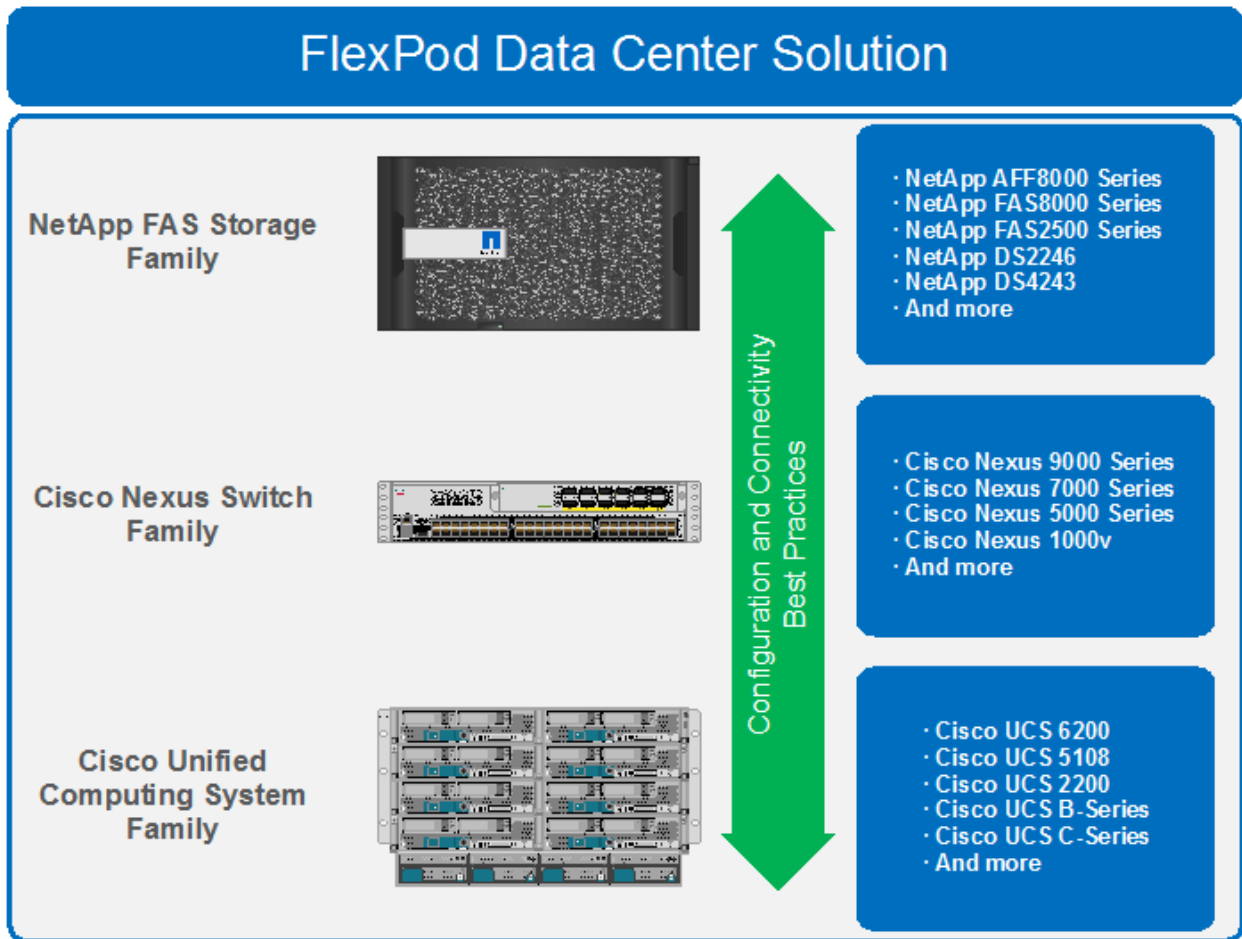
FlexPod is a best practice architecture that includes three components:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus switches
- NetApp FAS storage controllers

These components are connected and configured according to the best practices of both Cisco and NetApp and they provide the ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed). It also can scale out for environments that need multiple consistent deployments (rolling out additional FlexPod stacks). FlexPod delivers a baseline configuration and also has the flexibility to be sized and optimized to accommodate many different use cases.

Typically, the more scalable and flexible a solution is, the more difficult it becomes to maintain a single unified architecture capable of offering the same features and functionality across implementations. Doing so is one of the key benefits of FlexPod. Each of the component families shown in Figure 1 offers platform and resource options to scale the infrastructure up or down. These components do so while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod.

Figure 1) FlexPod component families.



NetApp MetroCluster

NetApp MetroCluster is a solution that combines array-based clustering with synchronous replication to deliver continuous availability and zero data loss at the lowest cost. Administration of the array-based cluster is simpler because the dependencies and complexity normally associated with host-based clustering are eliminated. MetroCluster immediately duplicates all your mission-critical data on a transaction-by-transaction basis, providing uninterrupted access to your applications and data. Unlike standard data-replication solutions, MetroCluster works seamlessly with your host environment to provide continuous data availability while eliminating the need to create and maintain failover scripts.

For more information on MetroCluster, see the [MetroCluster](#) page on netapp.com.

VMware vSphere Metro Storage Cluster (vMSC)

VMware vMSC is a specific type of configuration for a VMware vSphere cluster in a stretched cluster environment. NetApp MetroCluster is a certified vMSC configuration designed to maintain data availability beyond a single location in many different failure scenarios. This deployment guide describes the best practices to follow when configuring a vSphere cluster on FlexPod Datacenter with NetApp MetroCluster. The primary advantage of the vMSC best practices on FlexPod Datacenter with NetApp MetroCluster is that they take advantage of the active-active data availability of NetApp MetroCluster. They do so while

maintaining the capability to use vMotion with VMs between the two sites to prevent application downtime.

The best practices to follow include using these VMware technologies:

- VMware vSphere HA
- VMware vSphere DRS
- VMware VM-to-host affinity rules

Cisco Overlay Transport Virtualization (OTV)

Cisco OTV on the Cisco Nexus 7000 significantly simplifies extending layer 2 applications across distributed data centers. With OTV you can deploy virtual computing resources and clusters across geographically distributed data centers, delivering transparent workload mobility without requiring reconfiguration of the network or IP addressing. In the multisite FlexPod deployment, OTV enables seamless workload migration between sites.

Figure 2) Cisco OTV.



For more information on OTV, refer to the [Cisco OTV Overview](#).

2.2 Use Case Summary

FlexPod Datacenter with NetApp MetroCluster is a flexible architecture that can be sized and optimized to accommodate a wide variety of use cases. This deployment focuses on a fabric MetroCluster environment stretched to 300km to show the solution's capabilities to effectively span distances. This solution can be scaled down for a single campus or a single data center.


This solution applies to the following use cases:

- Maintaining data access during outages and disasters
- Simplifying disaster recovery management
- Simplifying management of infrastructure located in multiple sites

3 Solution Architecture

Figure 3 illustrates the architecture of FlexPod Datacenter with MetroCluster. You can find the cabling details for this solution in the appendix “[Cabling Details](#).”

Figure 3) FlexPod Datacenter with NetApp MetroCluster single site.



4 Technology Requirements

This section covers the technology requirements for the FlexPod Datacenter with MetroCluster solution.

4.1 Hardware Requirements

Table 1 lists the hardware components required to implement the FlexPod Datacenter with MetroCluster solution. The hardware components used in any particular implementation of the solution might vary based on customer requirements.

Table 1) Hardware requirements.

Layer	Hardware	Quantity
Compute	Cisco UCS 5108 blade server chassis	2
	Cisco UCS B-200 M4 blade server	4
	Cisco UCS 6248UP fabric interconnects	4
Network	Cisco Nexus 9372PX	4
	Cisco Nexus 7004	2
	Cisco Nexus 7000 M2-Series 24-Port 10 Gigabit Ethernet module with XL Option Data	2
	Cisco Nexus 7000 F2-Series Enhanced 48 Port Fiber 1 and 10 Gigabit Ethernet module	2
Storage	NetApp All Flash FAS8040 HA pair	2
	Cisco MDS 9396S	4
	ATTO FibreBridge 7500N	4
	NetApp DS4246 with 400G SSD all-flash disk shelf	4

4.2 Software Requirements

Table 2 lists the software components required to implement the FlexPod with MetroCluster solution. The software components used in any particular implementation of the solution might vary based on customer requirements.

Table 2) Software requirements.

Layer	Software	Version or Release
Compute	Cisco UCS fabric interconnects 6200 Series, UCS B-200 M4, UCS C-220 M4	2.2(5d)
	Cisco eNIC	2.1.2.71
	Cisco fNIC	1.6.0.17a
	Cisco UCS Central	1.4(1b)
Network	Cisco Nexus 9372PX	7.0(3)I2(1)
	Cisco Nexus 7004	7.3(0)D1(1)
Storage	NetApp All Flash FAS8040 HA pair	Data ONTAP® 8.3.2
	Cisco MDS 9396S	6.2(15)

Layer	Software	Version or Release
	ATTO FibreBridge 7500N	2.41
Software	VMware vSphere ESXi	6.0
	VMware vCenter	6.0
	NetApp Virtual Storage Console	6.2
	NetApp OnCommand® Unified Manager	6.4
	NetApp OnCommand Performance Manager	2.1

4.3 Infrastructure VLANs

Table 3 lists the VLANs that were used to implement this solution.

Table 3) Software requirements.

VLAN ID	Description
3336	Packet Control VLAN
3337	VM traffic VLAN
3338	vMotion VLAN
2	Native VLAN
3340	NFS VLAN
3341	iSCSI-A VLAN
3342	iSCSI-B VLAN
3343	IB MGMT VLAN
2000	OTV VLAN

4.4 Configuration Variables

Variables related to this configuration are located in the appendix “[Configuration Variables.](#)”

5 Deployment Procedures

5.1 Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made in each step to run the command in a single site or on both sites. For example, site A and site B are used to distinguish between the sites and node 1 and node 2 are used to distinguish between the two NetApp storage controllers in a given site. siteA_node1 refers to the first NetApp storage controller in site A. Cisco Nexus switches and

fabric interconnects are configured similarly. Finally, to indicate that you should include information pertinent to your environment in a given step, <<variables>> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?
[-node] <nodename>           Node
{ [-vlan-name] {<netport>|<ifgrp>} VLAN Name
| -port {<netport>|<ifgrp>}   Associated Network Port
[-vlan-id] <integer> }       Network Switch VLAN Identifier
```

Example:

```
network port vlan -node <<node01>> -vlan-name i0a-<<vlan id>>
```

When the variable includes A/B, the appropriate environment information should be entered for the site that is being configured.

5.2 Configure Cisco Nexus 9372PX Switches

The following section provides the detailed procedure for configuring the Cisco Nexus 9372PX switches for use in the FlexPod Datacenter with MetroCluster solution. Two switches per site were configured in this deployment.

Configure Initial Switch Settings

1. Run the Basic System Configuration Dialog utility on each of the Cisco Nexus 9372 switches. Use the default values for all inputs except the variables:

```

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): yes
Do you want to enforce secure password standard (yes/no) [y]:
  Create another login account (yes/no) [n]:
  Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
  Enter the switch name : <<var_site(A/B)_9K(1/2)>>
  Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
    Mgmt0 IPv4 address : <<var_site(A/B)_9K(1/2)>>_ip_address>>
    Mgmt0 IPv4 netmask : <<var_oob-mgmt_mask>>
  Configure the default gateway? (yes/no) [y]:
    IPv4 address of the default gateway : <<var_oob-mgmt_gateway>>
  Configure advanced IP options? (yes/no) [n]:
  Enable the telnet service? (yes/no) [n]:
  Enable the ssh service? (yes/no) [y]:
    Type of ssh key you would like to generate (dsa/rsa) [rsa]:
    Number of rsa key bits <1024-2048> [1024]:
  Configure the ntp server? (yes/no) [n]:
  Configure default interface layer (L3/L2) [L2]:
  Configure default switchport interface state (shut/noshut) [noshut]: shut
  Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:

The following configuration will be applied:
  password strength-check
  switchname <<var_site(A/B)_9K(1/2)>>
  vrf context management
  ip route 0.0.0.0/0 <<var_oob-mgmt_gateway>>
```

```

exit
  no feature telnet
  ssh key rsa 1024 force
  feature ssh
  system default switchport
  system default switchport shutdown
  copp profile strict
interface mgmt0
ip address <<var_site(A/B)_9K(1/2)_ip_address>> <<var_oob-mgmt_mask>>
no shutdown
Would you like to edit the configuration? (yes/no) [n]:

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

[#####] 100%
Copy complete.

```

Enable Licenses

1. To license the Cisco Nexus switches, run the following commands on all Cisco Nexus 9372 switches:

```

configure terminal
feature interface-vlan
feature lacp
feature vpc
feature lldp

```

Set Global Configurations

1. To set these global configurations on all Cisco Nexus 9372 switches, run the following commands:

```

configure terminal
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <<var_global_ntp_server_ip>> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <<var_ib-mgmt_vlan_gateway>>
copy run start

```

Create VLANs

1. To create the required VLANs, run the following commands on all Cisco Nexus 9372 switches:

```

vlan 2
name Native-VLAN
exit
vlan 3336
name Packet-Ctrl-VLAN
exit
vlan 3337
name VM-Traffic-VLAN
exit
vlan 3338
name vMotion-VLAN
exit
vlan 3340
name NFS-VLAN
exit
vlan 3341
name iSCSI-A-VLAN
exit
vlan 3342
name iSCSI-B-VLAN
exit
vlan 3343
name IB-MGMT-VLAN
exit

```

Add NTP Distribution Interface

To create the required NTP distribution interface, run the following commands.

Site A—Switch 1 <pre>ntp source <<var_siteA_9K1_ntp_ip_address>> interface Vlan3343 ip address <<var_site(A/B)_nexus9K(1/2)_ntp_ip_address>>/ <<var_ib-mgmt_mask>> no shutdown</pre>	Site B—Switch 1 <pre>ntp source <<var_siteB_9K1_ntp_ip_address>> interface Vlan3343 ip address <<var_site(A/B)_nexus9K(1/2)_ntp_ip_address>>/ <<var_ib-mgmt_mask>> no shutdown</pre>
Site A—Switch 2 <pre>ntp source <<var_siteB_9K2_ntp_ip_address>> interface Vlan3343 ip address <<var_site(A/B)_nexus9K(1/2)_ntp_ip_address>>/ <<var_ib-mgmt_mask>> no shutdown</pre>	Site B—Switch 2 <pre>ntp source <<var_siteB_9K2_ntp_ip_address>> interface Vlan3343 ip address <<var_site(A/B)_nexus9K(1/2)_ntp_ip_address>>/ <<var_ib-mgmt_mask>> no shutdown</pre>

Add Individual Port Description for Troubleshooting

To add individual port descriptions, run the following commands:

Site A—Switch 1 <pre>configure terminal interface Ethernet 1/1 description <<var_siteA_node1>>:e0g exit interface Ethernet 1/2 description <<var_siteA_node2>>:e0g exit interface Ethernet 1/11 description <<var_siteA_FI1>>:1/19 exit interface Ethernet 1/12 description <<var_siteA_FI2>>:1/19 exit interface Ethernet 1/47 description <<var_siteA_7K>>:eth3/1 exit interface Ethernet 1/48 description <<var_siteA_7K>>:eth3/3 exit interface Ethernet 1/49 description <<var_siteA_9K2>>:eth1/49 exit interface Ethernet 1/50 description <<var_siteA_9K2>>:eth1/50 exit</pre>	Site B—Switch 1 <pre>configure terminal interface Ethernet 1/1 description <<var_siteB_node1>>:e0g exit interface Ethernet 1/2 description <<var_siteB_node2>>:e0g exit interface Ethernet 1/11 description <<var_siteB_FI1>>:1/19 exit interface Ethernet 1/12 description <<var_siteB_FI2>>:1/19 exit interface Ethernet 1/47 description <<var_siteB_7K>>:eth3/1 exit interface Ethernet 1/48 description <<var_siteB_7K>>:eth3/3 exit interface Ethernet 1/49 description <<var_siteB_9K2>>:eth1/49 exit interface Ethernet 1/50 description <<var_siteB_9K2>>:eth1/50 exit</pre>
---	---

<p>Site A—Switch 2</p> <pre> configure terminal interface Ethernet 1/1 description <<var_siteA_node1>>:e0h exit interface Ethernet 1/2 description <<var_siteA_node2>>:e0h exit interface Ethernet 1/11 description <<var_siteA_FI1>>:1/20 exit interface Ethernet 1/12 description <<var_siteA_FI2>>:1/20 exit interface Ethernet 1/47 description <<var_siteA_7K>>:eth3/2 exit interface Ethernet 1/48 description <<var_siteA_7K>>:eth3/4 exit interface Ethernet 1/49 description <<var_siteA_9K1>>:eth1/49 exit interface Ethernet 1/50 description <<var_siteA_9K1>>:eth1/50 exit </pre>	<p>Site B—Switch 2</p> <pre> configure terminal interface Ethernet 1/1 description <<var_siteB_node1>>:e0h exit interface Ethernet 1/2 description <<var_siteB_node2>>:e0h exit interface Ethernet 1/11 description <<var_siteB_FI1>>:1/20 exit interface Ethernet 1/12 description <<var_siteB_FI2>>:1/20 exit interface Ethernet 1/47 description <<var_siteB_7K>>:eth3/2 exit interface Ethernet 1/48 description <<var_siteB_7K>>:eth3/4 exit interface Ethernet 1/49 description <<var_siteA_9K1>>:eth1/49 exit interface Ethernet 1/50 description <<var_siteA_9K1>>:eth1/50 exit </pre>
---	---

Create Port Channel

To create the required port channels, run the following commands on each switch in its respective site:

<p>Site A</p> <pre> interface Po10 description vPC peer-link exit interface Eth1/49-50 channel-group 10 mode active no shutdown exit interface Po9 description <<var_siteA_7K>>_OTV exit interface Eth1/47-48 channel-group 9 mode active no shutdown exit interface Po11 description <<var_siteA_node1>> exit interface Eth1/1 channel-group 11 mode active no shut exit interface Po12 description <<var_siteA_node2>> exit </pre>	<p>Site B</p> <pre> interface Po10 description vPC peer-link exit interface Eth1/49-50 channel-group 10 mode active no shutdown exit interface Po9 description <<var_siteB_7K>>_OTV exit interface Eth1/47-48 channel-group 9 mode active no shutdown exit interface Po11 description <<var_siteB_node1>> exit interface Eth1/1 channel-group 11 mode active no shut exit interface Po12 description <<var_siteB_node2>> exit </pre>
--	--

<pre> interface Eth1/2 channel-group 12 mode active no shut exit interface Po11 description <<var_siteA_FI1>> exit interface Eth1/11 channel-group 111 mode active no shut exit interface Po112 description <<var_siteA_FI2>> exit interface Eth1/12 channel-group 112 mode active no shut exit </pre>	<pre> interface Eth1/2 channel-group 12 mode active no shut exit interface Po11 description <<var_siteB_FI1>> exit interface Eth1/11 channel-group 111 mode active no shut exit interface Po112 description <<var_siteB_FI2>> exit interface Eth1/12 channel-group 112 mode active no shut exit </pre>
--	--

Configure Port Channel Parameters

To configure port channel parameters, run the following commands on all switches:

```

interface Po9
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3336,3337,3338,3340,3341,3342,3343
spanning-tree port type network
mtu 9216
exit
interface Po10
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3336,3337,3338,3340,3341,3342,3343
spanning-tree port type network
exit
interface Po11
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3343,3340,3341,3342
spanning-tree port type edge trunk
mtu 9216
exit
interface Po12
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3343,3340,3341,3342
spanning-tree port type edge trunk
mtu 9216
exit
interface Po111
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3336,3337,3338,3340,3341,3342,3343
spanning-tree port type edge trunk
mtu 9216
exit
interface Po112
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3336,3337,3338,3340,3341,3342,3343
spanning-tree port type edge trunk
mtu 9216
exit

```

Configure Virtual Port Channels

1. To configure a virtual port channel (vPC) relationship for site A, run the following commands on both Cisco Nexus 9372PX site A switches:

Site A—Switch 1	Site A—Switch 2
<pre>vpc domain 10 role priority 10 peer-keepalive destination <<var_siteA_9K2_ip_address>> source <<var_siteA_9K1_ip_address>> peer-switch peer-gateway auto-recovery delay restore 150 exit</pre>	<pre>vpc domain 20 role priority 20 peer-keepalive destination <<var_siteA_9K1_ip_address>> source <<var_siteA_9K2_ip_address>> peer-switch peer-gateway auto-recovery delay restore 150 exit</pre>

2. To configure a vPC relationship for site B, run the following commands on both Cisco Nexus 9372PX site B switches:

Site B—Switch 1	Site B—Switch 2
<pre>vpc domain 10 role priority 10 peer-keepalive destination <<var_siteB_9K2_ip_address>> source <<var_siteB_9K1_ip_address>> peer-switch peer-gateway auto-recovery delay restore 150 exit</pre>	<pre>vpc domain 20 role priority 20 peer-keepalive destination <<var_siteB_9K1_ip_address>> source <<var_siteB_9K2_ip_address>> peer-switch peer-gateway auto-recovery delay restore 150 exit</pre>

3. To configure vPCs, run the following commands on all Cisco Nexus 9372PX switches:

```
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po111
vpc 111
exit
interface Po112
vpc 112
exit
copy run start
```


5.3 Configure OTV to Span VLANs Between Two Sites

Because the OTV functionality was the sole focus for this deployment, a single Cisco Nexus 7004 was used on each site. In a production environment, Cisco recommends configuring two Cisco Nexus 7004 switches in each site.

Note: For more information, see the [FlexPod Datacenter with NetApp MetroCluster Design Guide](#).

Figure 4) Cisco Nexus 9K and 7K connectivity for OTV.




Table 4) Site A—Cisco Nexus 7004 OTV VDC connectivity.

Local Device	Local Port	Remote Device	Remote Port	Layer
Cisco Nexus 7004 OTV VDC	Eth 3/1	Cisco Nexus 9372 #1	Eth 1/47	2
Cisco Nexus 7004 OTV VDC	Eth 3/2	Cisco Nexus 9372 #1	Eth 1/48	2
Cisco Nexus 7004 OTV VDC	Eth 3/3	Cisco Nexus 9372 #2	Eth 1/47	2
Cisco Nexus 7004 OTV VDC	Eth 3/4	Cisco Nexus 9372 #2	Eth 1/48	2
Cisco Nexus 7004 OTV VDC	Eth 3/9	Cisco Nexus 7004 LAN VDC	Eth 4/9	2
Cisco Nexus 7004 OTV VDC	Eth 3/10	Cisco Nexus 7004 LAN VDC	Eth 4/10	3

Table 5) Site A—Cisco Nexus 7004 LAN VDC connectivity.

Local Device	Local Port	Remote Device	Remote Port	Layer
Cisco Nexus 7004 LAN VDC	Eth 4/9	Cisco Nexus 7004 OTV VDC	Eth 3/9	2
Cisco Nexus 7004 LAN VDC	Eth 4/10	Cisco Nexus 7004 OTV VDC	Eth 3/10	3
Cisco Nexus 7004 LAN VDC	Eth 4/11	Cisco Nexus 7004 LAN VDC - Site B	Eth 3/11	3
Cisco Nexus 7004 LAN VDC	Eth 4/12	Cisco Nexus 7004 LAN VDC - Site B	Eth 3/12	3

Table 6) Site B—Cisco Nexus 7004 OTV VDC connectivity.

Local Device	Local Port	Remote Device	Remote Port	Layer
Cisco Nexus 7004 OTV VDC	Eth 3/1	Cisco Nexus 9372 #3	Eth 1/47	2
Cisco Nexus 7004 OTV VDC	Eth 3/2	Cisco Nexus 9372 #3	Eth 1/48	2
Cisco Nexus 7004 OTV VDC	Eth 3/3	Cisco Nexus 9372 #4	Eth 1/47	2
Cisco Nexus 7004 OTV VDC	Eth 3/4	Cisco Nexus 9372 #4	Eth 1/48	2
Cisco Nexus 7004 OTV VDC	Eth 3/9	Cisco Nexus 7004 LAN VDC	Eth 4/9	2
Cisco Nexus 7004 OTV VDC	Eth 3/10	Cisco Nexus 7004 LAN VDC	Eth 4/10	3

Table 7) Site B—Cisco Nexus 7004 LAN VDC connectivity.

Local Device	Local Port	Remote Device	Remote Port	Layer
Cisco Nexus 7004 LAN VDC	Eth 4/9	Cisco Nexus 7004 OTV VDC	Eth 3/9	2
Cisco Nexus 7004 LAN VDC	Eth 4/10	Cisco Nexus 7004 OTV VDC	Eth 3/10	3
Cisco Nexus 7004 LAN VDC	Eth 4/11	Cisco Nexus 7004 LAN VDC - Site A	Eth 3/11	3

Local Device	Local Port	Remote Device	Remote Port	Layer
Cisco Nexus 7004 LAN VDC	Eth 4/12	Cisco Nexus 7004 LAN VDC—Site A	Eth 3/12	3

Configure the Initial Switch Settings for Switch A

Run the Basic System Configuration Dialog utility to configure switch A:

```

Abort Auto Provisioning and continue with normal setup ?(yes/no) [n]: yes

    ---- System Admin Account Setup ----

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

Enter the password for "admin":
Confirm the password for "admin":

Do you want to enable admin vdc (yes/no) [n]: y

    ---- Basic System Configuration Dialog VDC: 1 ----

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

Please register Cisco Nexus7000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus7000 devices must be registered to receive
entitled support services.

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): yes
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : <<var_siteA_7K>>
Enable license grace period? (yes/no) [n]: y
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
  Mgmt0 IPv4 address : <<var_siteA_7K_ip_address>>
  Mgmt0 IPv4 netmask : <<var_oob-mgmt_mask>>
Configure the default gateway? (yes/no) [y]:
  IPv4 address of the default gateway : <<var_oob-mgmt_gateway>>
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
  Type of ssh key you would like to generate (dsa/rsa) [rsa]:
  Number of rsa key bits <1024-2048> [1024]:
Configure the ntp server? (yes/no) [n]: y
  NTP server IPv4 address : <<var_oob-mgmt_ntp>>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [shut]:
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:

The following configuration will be applied:
password strength-check
switchname <<var_siteA_7K>>
license grace-period
vrf context management
ip route 0.0.0.0/0 <<var_oob-mgmt_gateway>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
ntp server <<var_oob-mgmt_ntp>>

```

```
system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address <<var_siteA_7K_ip_address>> <<var_oob-mgmt_mask>>
no shutdown
```

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

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

Configure the Initial Switch Settings for Switch B

Run the Basic System Configuration Dialog utility to configure switch B:

```
Abort Auto Provisioning and continue with normal setup ?(yes/no) [n]: yes
```

```
---- System Admin Account Setup ----
```

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

```
Enter the password for "admin":
Confirm the password for "admin":
```

Do you want to enable admin vdc (yes/no) [n]: y

```
---- Basic System Configuration Dialog VDC: 1 ----
```

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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): yes
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : <<var_siteB_7K>>
Enable license grace period? (yes/no) [n]: y
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Mgmt0 IPv4 address : <<var_siteB_7K_ip_address>>
Mgmt0 IPv4 netmask : <<var_oob-mgmt_mask>>
Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway : <<var_oob-mgmt_gateway>>
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
Type of ssh key you would like to generate (dsa/rsa) [rsa]:
Number of rsa key bits <1024-2048> [1024]:
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address : <<var_oob-mgmt_ntp>>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [shut]:
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
```

The following configuration will be applied:

```
password strength-check
switchname <<var_siteB_7K>>
license grace-period
vrf context management
ip route 0.0.0.0/0 <<var_oob-mgmt_gateway>>
exit
no feature telnet
```

```

ssh key rsa 1024 force
feature ssh
ntp server <<var_oob-mgmt_ntp>>
system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address <<var_siteB_7K_ip_address>> <<var_siteB_7K_mask>>
no shutdown

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

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

```

Create Virtual Device Contexts (VDCs)

Run the following commands on each switch to set up the required VDCs:

Cisco Nexus 7004—Site A	Cisco Nexus 7004—Site B
<pre> vdc LAN limit-resource module-type f2 allocate interface Ethernet 4/9-12 exit vdc OTV limit-resource module-type m2x1 allocate interface Ethernet 3/1-23 </pre>	<pre> vdc LAN limit-resource module-type f2 allocate interface Ethernet 4/9-12 exit vdc OTV limit-resource module-type m2x1 allocate interface Ethernet 3/1-23 </pre>

Configure Initial Switch Settings for Each VDC

Run the Basic System Configuration Dialog utility to configure each VDC.

Run `switchto vdc <<vdc_name>>` to begin setup on the LAN and OTV VDCs on each switch:

```

stlnexus7004-1# switchto vdc LAN

----- System Admin Account Setup -----

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

Enter the password for "admin":
Confirm the password for "admin":

----- Basic System Configuration Dialog VDC: 2 -----

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

Please register Cisco Nexus7000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus7000 devices must be registered to receive
entitled support services.

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): yes
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : <<var_site(A/B)_7K_(OTV/LAN)>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Mgmt0 IPv4 address : <<var_site(A/B)_7K_(OTV/LAN)_ip_address>>
Mgmt0 IPv4 netmask : <<var_oob-mgmt_mask>>
Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway : <<var_oob-mgmt_gateway>>

```

```

Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
  Type of ssh key you would like to generate (dsa/rsa) [rsa]:
  Number of rsa key bits <1024-2048> [1024]:
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [shut]:

```

The following configuration will be applied:

```

password strength-check
switchname A
vrf context management
ip route 0.0.0.0/0 <<var_oob-mgmt_gateway>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
system default switchport shutdown
interface mgmt0
ip address <<var_site(A/B)_7K_(OTV/LAN)_ip_address>> <<var_oob-mgmt_mask>>
no shutdown

```

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

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

Configure LAN VDC

To configure the LAN VDCs, run the following commands:

Cisco Nexus 7004—Site A—LAN VDC

- From the admin VDC:


```

switchto vdc LAN

```
- From the LAN VDC:


```

configure terminal
feature lacp
feature udld
feature ospf

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

vlan 2000
name OTV-VLAN
vlan 3336
name Packet-Ctrl-VLAN
vlan 3337
name VM-Traffic-VLAN
vlan 3338
name vMotion-VLAN
vlan 3340
name NFS-VLAN
vlan 3341
name iSCSI-A-VLAN
vlan 3342
name iSCSI-B-VLAN
vlan 3343
name IB-MGMT-VLAN

interface Po10
description <<var_siteB_7K>>
exit

```

Cisco Nexus 7004—Site B—LAN VDC

- From the admin VDC:


```

switchto vdc LAN

```
- From the LAN VDC:


```

configure terminal
feature lacp
feature udld
feature ospf

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

vlan 2000
name OTV-VLAN
vlan 3336
name Packet-Ctrl-VLAN
vlan 3337
name VM-Traffic-VLAN
vlan 3338
name vMotion-VLAN
vlan 3340
name NFS-VLAN
vlan 3341
name iSCSI-A-VLAN
vlan 3342
name iSCSI-B-VLAN
vlan 3343
name IB-MGMT-VLAN

interface Po10
description <<var_siteA_7K>>
exit

```

<pre> interface Eth4/11 description <<var_siteB_7K:eth4/11>> exit interface Eth4/12 description <<var_siteB_7K:eth4/12>> exit interface Eth4/11-12 channel-group 10 mode active exit interface Po10 no switchport mtu 9216 ip address 198.18.2.1/30 ip router ospf 10 area 0.0.0.0 no shutdown exit interface Ethernet4/9 switchport mode trunk switchport trunk native vlan 2 switchport trunk allowed vlan <<var_otv_vlan_id>> spanning-tree port type network mtu 9216 no shutdown exit interface Ethernet4/10 no switchport mtu 9216 ip address 198.18.1.2/30 ip router ospf 10 area 0.0.0.0 no shutdown exit router ospf 1 </pre>	<pre> interface Eth4/11 description <<var_siteA_7K:eth4/11>> exit interface Eth4/12 description <<var_siteA_7K:eth4/12>> exit interface Eth4/11-12 channel-group 10 mode active exit interface Po10 no switchport mtu 9216 ip address 198.18.2.2/30 ip router ospf 10 area 0.0.0.0 no shutdown exit interface Ethernet4/9 switchport mode trunk switchport trunk native vlan 2 switchport trunk allowed vlan <<var_otv_vlan_id>> spanning-tree port type network mtu 9216 no shutdown exit interface Ethernet4/10 no switchport mtu 9216 ip address 198.18.3.1/30 ip router ospf 10 area 0.0.0.0 no shutdown exit router ospf 10 </pre>
--	---

Configure the OTV VDC

To configure the OTV VDCs, run the following commands:

<p>Cisco Nexus 7004—Site A—OTV VDC</p> <ul style="list-style-type: none"> From the admin VDC, run: <pre>switchto vdc OTV</pre> From the OTV VDC, run: <pre>configure terminal feature lacp feature otv feature ospf feature interface-vlan spanning-tree port type edge bpduguard default spanning-tree port type edge bpdufilter default spanning-tree port type network default vlan 2000 name OTV-VLAN vlan 3336 name Packet-Ctrl-VLAN vlan 3337 name VM-Traffic-VLAN vlan 3338 name vMotion-VLAN</pre> 	<p>Cisco Nexus 7004—Site B—OTV VDC</p> <ul style="list-style-type: none"> From the admin VDC, run: <pre>switchto vdc OTV</pre> From the OTV VDC, run: <pre>configure terminal feature lacp feature otv feature ospf feature interface-vlan spanning-tree port type edge bpduguard default spanning-tree port type edge bpdufilter default spanning-tree port type network default vlan 2000 name OTV-VLAN vlan 3336 name Packet-Ctrl-VLAN vlan 3337 name VM-Traffic-VLAN vlan 3338 name vMotion-VLAN</pre>
---	---

<pre> vlan 3340 name NFS-VLAN vlan 3341 name iSCSI-A-VLAN vlan 3342 name iSCSI-B-VLAN vlan 3343 name IB-MGMT-VLAN interface port-channel9 description Nexus_9K exit interface Ethernet3/1 description <<var_siteA_9K1:eth1/47>> exit interface Ethernet3/2 description <<var_siteA_9K2:eth1/47>> exit interface Ethernet3/3 description <<var_siteA_9K1:eth1/48>> exit interface Ethernet3/4 description <<var_siteA_9K2:eth1/48>> exit interface Ethernet3/1-4 channel-group 9 mode active no shutdown exit interface Po9 switchport switchport mode trunk switchport trunk native vlan 2 switchport trunk allowed vlan 3336,3337,3338,3340,3341,3342,3343 spanning-tree port type network mtu 9216 no shutdown interface Ethernet3/9 description OTV L2 Interface switchport switchport mode trunk switchport trunk native vlan 2 switchport trunk allowed vlan <<var_otv_vlan_id>> spanning-tree port type network mtu 9216 no shutdown interface Ethernet3/10 description OTV L3 Interface mtu 9216 ip address 198.18.1.1/30 ip router ospf 10 area 0.0.0.0 ip igmp version 3 no shutdown interface Overlay10 otv join-interface Ethernet3/10 otv extend-vlan 3336,3337,3338,3340,3341,3342,3343otv adjacency-server unicast-only no otv suppress-arp-nd no shutdown router ospf 10 otv site-identifier 0x1 otv site-vlan 2000 </pre>	<pre> vlan 3340 name NFS-VLAN vlan 3341 name iSCSI-A-VLAN vlan 3342 name iSCSI-B-VLAN vlan 3343 name IB-MGMT-VLAN interface port-channel9 description Nexus_9K exit interface Ethernet3/1 description <<var_siteB_9K1:eth1/47>> exit interface Ethernet3/2 description <<var_siteB_9K2:eth1/47>> exit interface Ethernet3/3 description <<var_siteB_9K1:eth1/48>> exit interface Ethernet3/4 description <<var_siteB_9K2:eth1/48>> exit interface Ethernet3/1-4 channel-group 9 mode active no shutdown exit interface Po9 switchport switchport mode trunk switchport trunk native vlan 2 switchport trunk allowed vlan 3336,3337,3338,3340,3341,3342,3343 spanning-tree port type network mtu 9216 no shutdown interface Ethernet3/9 description OTV L2 Interface switchport switchport mode trunk switchport trunk native vlan 2 switchport trunk allowed vlan <<var_otv_vlan_id>> spanning-tree port type network mtu 9216 no shutdown interface Ethernet3/10 description OTV L3 Interface mtu 9216 ip address 198.18.3.2/30 ip router ospf 10 area 0.0.0.0 ip igmp version 3 no shutdown interface Overlay10 otv join-interface Ethernet3/10 otv extend-vlan 3336,3337,3338,3340,3341,3342,3343otv use- adjacency-server 198.18.1.1 unicast-only no otv suppress-arp-nd no shutdown router ospf 10 otv site-identifier 0x2 otv site-vlan 2000 </pre>
---	---

Verify the OTV Setup

To make sure that the OTV setup passes traffic between the sites, run the following commands:

Cisco Nexus 7004—Site A—OTV VDC	Cisco Nexus 7004—Site B—OTV VDC
<pre> stlnexus7004-1-OTV# ping 198.18.3.1 PING 198.18.3.1 (198.18.3.1): 56 data bytes 64 bytes from 198.18.3.1: icmp_seq=0 ttl=252 time=0.958 ms 64 bytes from 198.18.3.1: icmp_seq=1 ttl=252 time=0.769 ms 64 bytes from 198.18.3.1: icmp_seq=2 ttl=252 time=0.899 ms 64 bytes from 198.18.3.1: icmp_seq=3 ttl=252 time=1.258 ms 64 bytes from 198.18.3.1: icmp_seq=4 ttl=252 time=1.128 ms --- 198.18.3.1 ping statistics --- 5 packets transmitted, 5 packets received, 0.00% packet loss round-trip min/avg/max = 0.769/1.002/1.258 ms </pre>	<pre> stlnexus7004-2-OTV# ping 198.18.1.1 PING 198.18.1.1 (198.18.1.1): 56 data bytes 64 bytes from 198.18.1.1: icmp_seq=0 ttl=252 time=1.009 ms 64 bytes from 198.18.1.1: icmp_seq=1 ttl=252 time=0.686 ms 64 bytes from 198.18.1.1: icmp_seq=2 ttl=252 time=0.71 ms 64 bytes from 198.18.1.1: icmp_seq=3 ttl=252 time=0.699 ms 64 bytes from 198.18.1.1: icmp_seq=4 ttl=252 time=0.72 ms --- 198.18.1.1 ping statistics --- 5 packets transmitted, 5 packets received, 0.00% packet loss round-trip min/avg/max = 0.686/0.764/1.009 ms </pre>
<pre> stlnexus7004-1-OTV# show otv OTV Overlay Information Site Identifier 0000.0000.0001 Overlay interface Overlay10 VPN name : Overlay10 VPN state : UP Extended vlans : 3336-3338 3340-3343 (Total:7) Join interface(s) : Eth3/10 (198.18.1.1) Site vlan : 2000 (up) AED-Capable : Yes Capability : Unicast-Only Is Adjacency Server : Yes Adjacency Server(s) : [None] / [None] </pre>	<pre> stlnexus7004-2-OTV# show otv OTV Overlay Information Site Identifier 0000.0000.0002 Overlay interface Overlay10 VPN name : Overlay10 VPN state : UP Extended vlans : 3336-3338 3340-3343 (Total:7) Join interface(s) : Eth3/10 (198.18.3.1) Site vlan : 2000 (up) AED-Capable : Yes Capability : Unicast-Only Is Adjacency Server : No Adjacency Server(s) : 198.18.1.1 / [None] </pre>
<pre> stlnexus7004-1-OTV# show otv vlan OTV Extended VLANs and Edge Device State Information (* - AED) Legend: (NA) - Non AED, (VD) - Vlan Disabled, (OD) - Overlay Down (DH) - Delete Holddown, (HW) - HW: State Down VLAN Auth. Edge Device Vlan State Overlay ----- 3336* stlnexus7004-1-OTV active Overlay10 3337* stlnexus7004-1-OTV active Overlay10 3338* stlnexus7004-1-OTV active Overlay10 3340* stlnexus7004-1-OTV active Overlay10 3341* stlnexus7004-1-OTV active Overlay10 3342* stlnexus7004-1-OTV active Overlay10 </pre>	<pre> stlnexus7004-2-OTV# show otv vlan OTV Extended VLANs and Edge Device State Information (* - AED) Legend: (NA) - Non AED, (VD) - Vlan Disabled, (OD) - Overlay Down (DH) - Delete Holddown, (HW) - HW: State Down VLAN Auth. Edge Device Vlan State Overlay ----- 3336* stlnexus7004-2-OTV active Overlay10 3337* stlnexus7004-2-OTV active Overlay10 3338* stlnexus7004-2-OTV active Overlay10 3340* stlnexus7004-2-OTV active Overlay10 3341* stlnexus7004-2-OTV active Overlay10 3342* stlnexus7004-2-OTV active Overlay10 </pre>

3343* stlnexus7004-1-OTV active Overlay10	3343* stlnexus7004-2-OTV active Overlay10
--	--

5.4 Set Shelf IDs

Each shelf in the MetroCluster configuration requires a shelf ID unique to the FlexPod system. Power cycle each disk shelf after setting the shelf ID.

5.5 Set Up ATTO FibreBridge 7500N

1. To set up the FibreBridge controller, run the following commands on each FibreBridge controller:

```
set DHCP mp1 disabled
set ipaddress mp1 <<var_site(A/B)_fibrebridge(1/2)_ip_address>>
set ipsubnetmask mp1 <<var_oob-mgmt_mask>>
set ipgateway mp1 <<var_oob-mgmt_gateway>>
set fcconnmode 1 ptp
set fcdatastrate 1 16Gb
set fcconnmode 2 ptp
set fcdatastrate 2 16Gb
set SNMP enabled
set bridgename <<var_site(A/B)_fibrebridge(1/2)>>
sasportenable a
sasportenable b
fcportenable 1
fcportenable 2
saveconfiguration
```

2. Save the configuration and restart the bridge.

```
Restart is necessary...
Do you wish to restart (y/n) ?y
```

5.6 Set Up Cisco MDS 9396s

Configure Cisco MDS 9396s Switches

Set up each of the Cisco Nexus 9372 switches by running the Basic System Configuration Dialog. Use the default settings for all values except for the variables.

```
---- Basic System Configuration Dialog ----

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): yes
Do you want to enforce secure password standard (yes/no) [y]:
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]: y
SNMP community string : public
Enter the switch name : <<var_site(A/B)_mds(1/2)>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Mgmt0 IPv4 address : <<var_site(A/B)_mds(1/2)_ip_address>>
Mgmt0 IPv4 netmask : <<var_oob-mgmt_mask>>
Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway : <<var_oob-mgmt_gateway>>
Configure advanced IP options? (yes/no) [n]:
```

```
Enable the ssh service? (yes/no) [y]:
  Type of ssh key you would like to generate (dsa/rsa) [rsa]:
  Number of rsa key bits <1024-2048> [1024]:
Enable the telnet service? (yes/no) [n]:
Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]:
Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]:
Enter milliseconds in multiples of 10 for congestion-drop for port mode F
in range (<100-500>/default), where default is 500. [d]:
```

Congestion-drop for port mode E must be greater than or equal to
Congestion-drop for port mode F. Hence, Congestion drop for port
mode E will be set as default.

```
Enable the http-server? (yes/no) [y]:
Configure clock? (yes/no) [n]:
Configure timezone? (yes/no) [n]:
Configure summertime? (yes/no) [n]:
Configure the ntp server? (yes/no) [n]:
Configure default switchport interface state (shut/noshut) [shut]:
Configure default switchport trunk mode (on/off/auto) [on]:
Configure default switchport port mode F (yes/no) [n]:
Configure default zone policy (permit/deny) [deny]:
Enable full zoneset distribution? (yes/no) [n]: yes
Configure default zone mode (basic/enhanced) [basic]:
```

The following configuration will be applied:

```
password strength-check
snmp-server community public ro
switchname <<var_site(A/B)_mds(1/2)>>
interface mgmt0
  ip address <<var_site(A/B)_mds(1/2)_ip_address>> <<var_oob-mgmt_mask>>
  no shutdown
ip default-gateway <<var_oob-mgmt_gateway>>
ssh key rsa 1024 force
feature ssh
no feature telnet
system timeout congestion-drop default mode F
system timeout congestion-drop default mode E
feature http-server
system default switchport shutdown
system default switchport trunk mode on
no system default zone default-zone permit
system default zone distribute full
no system default zone mode enhanced
```

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

```
Use this configuration and save it? (yes/no) [y]:
[#####] 100%
Copy complete.
```

Cisco License Requirements

To use the Cisco MDS 9396s in a fabric MetroCluster system, the following licenses must be installed on the switches:

- ENTERPRISE_PKG
- PORT_ACTIVATION_PKG
- FM_SERVER_PKG

Run the `show license usage` command to verify that these licenses are installed.

Feature	Ins	Lic	Status	Expiry	Date	Comments
		Count				
FM_SERVER_PKG	Yes	-	Unused	never		-
ENTERPRISE_PKG	Yes	-	In use	never		-
PORT_ACTIV_9396S_PKG	Yes	96	In use	never		-

Configure Point-to-Point Ports

Configure the following settings on each switch port connected to the controller or FibreBridge controller. In this deployment, ports fc1/1-6 were connected to NetApp controllers and ports fc1/7 and fc1/8 were connected to the ATTO FibreBridge controllers.

```
configure terminal
interface fc1/1-8
shut
port-license acquire
switchport mode F
switchport speed 16000
switchport rate-mode dedicated
no shut
end
```

Configure Inter-Switch Link (ISL) Ports

Each switch must connect to the other switch in the fabric through an ISL link. The ISL link should be a port channel consisting of one to four TE ports. Buffer-to-buffer credits should be allocated to the ISL for the proper distance. In this deployment, the ISL link is running at 16Gbps at a distance of 300km.

Note: For more information about ISL configuration, refer to the [MetroCluster Installation and Configuration Guide](#) and the [FlexPod Datacenter with NetApp MetroCluster Design Guide](#).

To configure the ISL link, complete the following steps.

1. Configure the other ports in the port group to use the minimum required number of buffer-to-buffer credits (BBCs), freeing up additional BBCs for the E-Port. For this deployment, ISL ports fc1/44 and fc1/96 were used. Other ports in the port groups were fc1/41-43 and fc1/93-96, respectively.

```
configure terminal
interface fc1/41
switchport fcrxbbcredit 2
exit
interface fc1/42
switchport fcrxbbcredit 2
exit
interface fc1/43
switchport fcrxbbcredit 2
exit
interface fc1/93
switchport fcrxbbcredit 2
exit
interface fc1/94
switchport fcrxbbcredit 2
exit
interface fc1/95
```

```
switchport fcrxbbcredit 2
exit
```

2. Run the following commands to configure the E-Ports:

```
configure terminal
interface fcl/44,fcl/92
switchport mode E
switchport trunk allowed vsan 10
switchport trunk allowed vsan add 20
switchport speed 16000
switchport rate-mode dedicated
switchport fec
switchport fcrxbbcredit extended 3600
channel-group 1
no shutdown
exit
interface port-channel 1
switchport trunk allowed vsan add 1
exit
copy run start
```

Note: The MetroCluster best practice is to configure 1.5 times the minimum required number of BBCs. In this configuration, the ISLs were configured with 300km of distance at 16Gb. Therefore, we configured $300 * 8 * 1.5$, or 3,600 BBCs.

Configure VSANs

Each fabric requires a virtual storage area network (VSAN) for the controller's FCVI ports and a separate VSAN for the storage initiator and target ports. To configure the VSANs, complete the following steps.

1. Run the following commands on both switches in fabric 1 to configure FCVI VSAN:

```
configure terminal
vsan database
vsan 10
vsan 10 name FCVI_1_10
vsan 10 interface fcl/1
vsan 10 interface fcl/4
exit
in-order-guarantee vsan 10
vsan database
vsan 10 loadbalancing src-dst-id
exit
qos enable
qos class-map FCVI_1_10_Class match-any
qos policy-map FCVI_1_10_Policy
class FCVI_1_10_Class
priority high
exit
exit
qos service policy FCVI_1_10_Policy vsan 10
end
copy run start
```

Note: Quality of service (QoS) zones are used to prioritize the FCVI traffic, which keeps the nodes' NVRAM in sync.

2. Run the following commands on both switches in fabric 2 to configure FCVI VSAN:

```
configure terminal
vsan database
vsan 30
vsan 30 name FCVI_1_30
vsan 30 interface fcl/1
vsan 30 interface fcl/4
exit
in-order-guarantee vsan 30
vsan database
vsan 30 loadbalancing src-dst-id
```

```
exit
qos enable
qos class-map FCVI_1_30_Class match-any
qos policy-map FCVI_1_30_Policy
class FCVI_1_30_Class
priority high
exit
exit
qos service policy FCVI_1_30_Policy vsan 30
end
copy run start
```

3. Run the following commands on both switches in fabric 1 to configure storage VSAN:

```
configure terminal
vsan database
vsan 20
vsan 20 name STOR_1_20
vsan 20 interface fc1/2
vsan 20 interface fc1/3
vsan 20 interface fc1/5
vsan 20 interface fc1/6
vsan 20 interface fc1/7
vsan 20 interface fc1/8
exit
```

4. Run the following commands on both switches in fabric 2 to configure storage VSAN:

```
configure terminal
vsan database
vsan 40
vsan 40 name STOR_1_40
vsan 40 interface fc1/2
vsan 40 interface fc1/3
vsan 40 interface fc1/5
vsan 40 interface fc1/6
vsan 40 interface fc1/7
vsan 40 interface fc1/8
exit
```


Zone the Switches

This deployment contains two ATTO 7500N FibreBridge controllers per site. Each FibreBridge controller has two FC target ports and is connected to both fabrics in the site. A zone must be created for each FC target port.

Additionally, an FCVI zone must be created to allow the controllers to send FCVI traffic to each other.

Figure 5 shows the connection details of one site of the MetroCluster system.

Figure 5) MetroCluster cabling diagram for a single site.



Complete the following steps to zone the switches. For more information about the zoning options, see the [FlexPod Datacenter with NetApp MetroCluster Design Guide](#).

1. Get the WWN of each switch by running the `show wwn switch` command.
2. Run the following commands on both switches in fabric 1 to set the zone distribution parameters:

```
configure terminal
no system default zone default-zone permit
system default zone distribute full
no zone default-zone permit vsan 10
no zone default-zone permit vsan 20
zoneset distribute full vsan 10
zoneset distribute full vsan 20
end
copy running-config startup-config
```

3. Run the following commands on both switches in fabric 2 to set the zone distribution parameters:

```
configure terminal
no system default zone default-zone permit
system default zone distribute full
no zone default-zone permit vsan 30
```

```

no zone default-zone permit vsan 40
zoneset distribute full vsan 30
zoneset distribute full vsan 40
end
copy running-config startup-config

```


4. Create a zone for FCVI ports shown in Figure 6 by running the following commands on one switch in fabric 1:

```

configure terminal
zone name FCVI_1 vsan 10
member interface fc1/1 swnn <<var_siteA_mds01_switch_wnn>>
member interface fc1/4 swnn <<var_siteA_mds01_switch_wnn>>
member interface fc1/1 swnn <<var_siteB_mds01_switch_wnn>>
member interface fc1/4 swnn <<var_siteB_mds01_switch_wnn>>
exit

```

Figure 6) Ports for FCVI zone.



5. Create zones for the four storage target ports:


- a. Run the following commands to create a zone for the ports shown in Figure 7.

```

configure terminal
zone name STOR_cil_fc1 vsan 20
member interface fc1/2 swnn <<var_siteA_mds01_switch_wnn>>
member interface fc1/5 swnn <<var_siteA_mds01_switch_wnn>>
member interface fc1/2 swnn <<var_siteB_mds01_switch_wnn>>
member interface fc1/5 swnn <<var_siteB_mds01_switch_wnn>>
member interface fc1/7 swnn <<var_siteA_mds01_switch_wnn>>
exit

```

Figure 7) Ports for storage zone #1.




- b. Run the following commands to create a zone for the ports shown in Figure 8.

```

configure terminal
zone name STOR_ci3_fc1 vsan 20
member interface fc1/2 swnn <<var_siteA_mds01_switch_wnn>>
member interface fc1/5 swnn <<var_siteA_mds01_switch_wnn>>
member interface fc1/2 swnn <<var_siteB_mds01_switch_wnn>>
member interface fc1/5 swnn <<var_siteB_mds01_switch_wnn>>
member interface fc1/7 swnn <<var_siteB_mds01_switch_wnn>>
exit

```



Figure 8) Ports for storage zone #2.



- c. Run the following commands to create a zone for the ports shown in Figure 9.

```
configure terminal
zone name STOR_ci2_fc1 vsan 20
member interface fc1/3 swwn <<var_siteA_mds01_switch_wwn>>
member interface fc1/6 swwn <<var_siteA_mds01_switch_wwn>>
member interface fc1/3 swwn <<var_siteB_mds01_switch_wwn>>
member interface fc1/6 swwn <<var_siteB_mds01_switch_wwn>>
member interface fc1/8 swwn <<var_siteA_mds01_switch_wwn>>
exit
```


Figure 9) Ports for storage zone #3.



- d. Run the following commands to create a zone for the ports shown in Figure 10.

```
configure terminal
zone name STOR_ci4_fc1 vsan 20
member interface fc1/3 swwn <<var_siteA_mds01_switch_wwn>>
member interface fc1/6 swwn <<var_siteA_mds01_switch_wwn>>
member interface fc1/3 swwn <<var_siteB_mds01_switch_wwn>>
member interface fc1/6 swwn <<var_siteB_mds01_switch_wwn>>
member interface fc1/8 swwn <<var_siteB_mds01_switch_wwn>>
exit
```

Figure 10) Ports for storage zone #4.



6. Run the following commands to add zones to zonesets and activate them to complete the zoning for fabric 1.

```
zoneset name FCVI_zoneset vsan 10
member FCVI_1
exit
zoneset name STOR_zoneset_20 vsan 20
member STOR_ci1_fc1
member STOR_ci2_fc1
member STOR_ci3_fc1
member STOR_ci4_fc1
zoneset activate name FCVI_zoneset vsan 10
zoneset activate name STOR_zoneset_20 vsan 20
exit
```

7. The zones to be created for fabric 2 should have identical ports as the zones from fabric 1. Run the following commands to create the zones, add them to the zonesets, and activate them.

```

configure terminal
zone name FCVI_1 vsan 30
member interface fc1/1 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/4 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/1 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/4 swnn <<var_siteB_mds02_switch_wnn>>
exit
zone name STOR_cil_fc2 vsan 40
member interface fc1/2 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/5 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/2 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/5 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/7 swnn <<var_siteA_mds02_switch_wnn>>
exit
zone name STOR_ci2_fc2 vsan 40
member interface fc1/3 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/6 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/3 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/6 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/8 swnn <<var_siteA_mds02_switch_wnn>>
exit
zone name STOR_ci3_fc2 vsan 40
member interface fc1/2 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/5 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/2 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/5 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/7 swnn <<var_siteB_mds02_switch_wnn>>
exit
zone name STOR_ci4_fc2 vsan 40
member interface fc1/3 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/6 swnn <<var_siteA_mds02_switch_wnn>>
member interface fc1/3 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/6 swnn <<var_siteB_mds02_switch_wnn>>
member interface fc1/8 swnn <<var_siteB_mds02_switch_wnn>>
exit
zoneset name FCVI_zoneset vsan 30
member FCVI_1
exit
zoneset activate name FCVI_zoneset vsan 30
zoneset name STOR_zoneset_40 vsan 40
member STOR_cil_fc2
member STOR_ci2_fc2
member STOR_ci3_fc2
member STOR_ci4_fc2
exit
zoneset activate name FCVI_zoneset vsan 30
zoneset activate name STOR_zoneset_40 vsan 40

```

Confirm Correct Zoning

To verify that all devices are logged into fabrics and zoned correctly, run `show zoneset active` and verify that each zone device's FCID shows up in the zone.

```

cisco9396s-fcs40# show zoneset active vsan 10
zoneset name FCVI_zoneset vsan 10
  zone name FCVI_1 vsan 10
    * fcid 0x120000 [interface fc1/1 swnn 20:00:8c:60:4f:bd:b2:e0]
    * fcid 0x120020 [interface fc1/4 swnn 20:00:8c:60:4f:bd:b2:e0]
    * fcid 0x860000 [interface fc1/1 swnn 20:00:8c:60:4f:bd:b0:50]
    * fcid 0x860020 [interface fc1/4 swnn 20:00:8c:60:4f:bd:b0:50]
cisco9396s-fcs40# show zoneset active vsan 20
zoneset name STOR_zoneset_1_20 vsan 20
  zone name STOR_cil_fc1 vsan 20
    * fcid 0xe40000 [interface fc1/2 swnn 20:00:8c:60:4f:bd:b2:e0]
    * fcid 0xe40020 [interface fc1/5 swnn 20:00:8c:60:4f:bd:b2:e0]
    * fcid 0x3c0000 [interface fc1/2 swnn 20:00:8c:60:4f:bd:b0:50]

```

```

* fcid 0x3c0040 [interface fc1/5 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0xe40040 [interface fc1/7 swwn 20:00:8c:60:4f:bd:b2:e0]

zone name STOR_ci3_fc1_vsan 20
* fcid 0xe40000 [interface fc1/2 swwn 20:00:8c:60:4f:bd:b2:e0]
* fcid 0xe40020 [interface fc1/5 swwn 20:00:8c:60:4f:bd:b2:e0]
* fcid 0x3c0000 [interface fc1/2 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0x3c0040 [interface fc1/5 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0x3c0080 [interface fc1/7 swwn 20:00:8c:60:4f:bd:b0:50]

zone name STOR_ci2_fc1_vsan 20
* fcid 0xe40080 [interface fc1/3 swwn 20:00:8c:60:4f:bd:b2:e0]
* fcid 0xe400a0 [interface fc1/6 swwn 20:00:8c:60:4f:bd:b2:e0]
* fcid 0x3c0020 [interface fc1/3 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0x3c0060 [interface fc1/6 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0xe40060 [interface fc1/8 swwn 20:00:8c:60:4f:bd:b2:e0]

zone name STOR_ci4_fc1_vsan 20
* fcid 0xe40080 [interface fc1/3 swwn 20:00:8c:60:4f:bd:b2:e0]
* fcid 0xe400a0 [interface fc1/6 swwn 20:00:8c:60:4f:bd:b2:e0]
* fcid 0x3c0020 [interface fc1/3 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0x3c0060 [interface fc1/6 swwn 20:00:8c:60:4f:bd:b0:50]
* fcid 0x3c00a0 [interface fc1/8 swwn 20:00:8c:60:4f:bd:b0:50]

```

5.7 Controllers Setup

The procedure described in this section should be executed on both site A and site B storage controllers.

Disk Assignment

The disks in the MetroCluster configuration should be assigned such that each node owns an equal number of local disks and remote disks.

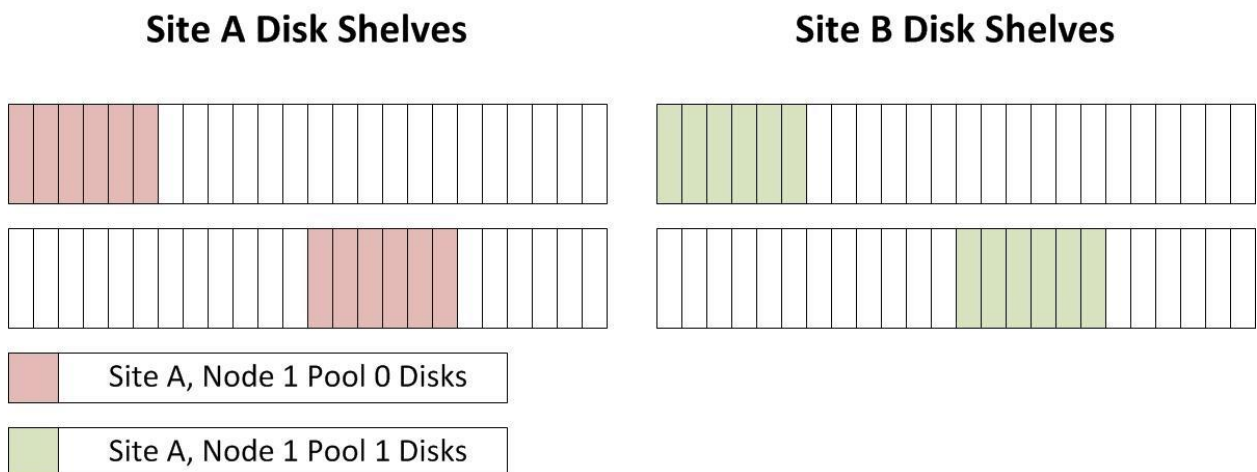
Note: Pool 0 always contains the disks that are found on the same site as the storage system that owns them.

Pool 1 always contains the disks that are found on the remote site.

All disks are assigned before shipping on new MetroCluster systems. To verify that the disks are assigned correctly, boot the nodes into maintenance mode and run the `storage show disk` command.

In this deployment, each storage system is assigned 12 pool 0 disks and 12 pool 1 disks, 6 from each shelf, as depicted in Figure 11.

Figure 11) Site A, node 1 disk assignment.



Set Up the Node

Enter the node management details on each node during the first boot.

```
Welcome to node setup.

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 setup wizard.
  Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.

Enter the node management interface port [e0M]: <<var_site(A/B)_node(1/2)_port>>
Enter the node management interface IP address [10.228.57.139]:
<<var_site(A/B)_node(1/2)_ip_address>>

Enter the node management interface netmask [255.255.252.0]: <<var_oob-mgmt_mask>>
Enter the node management interface default gateway [10.228.56.1]: <<var_oob-mgmt_gateway>>

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility
from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_siteA_node1_ip_address>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.

Wed Apr 27 17:57:27 UTC 2016
login: admin
*****
* This is a serial console session. Output from this *
* session is mirrored on the SP console session.    *
*****
::>
```

Create the Cluster

Run the cluster setup wizard on one node on each site to create two clusters.

```
::> cluster setup

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.

Do you want to create a new cluster or join an existing cluster? {create, join}: create

Do you intend for this node to be used as a single node cluster? {yes, no} [no]: no
Will the cluster network be configured to use network switches? [no]: no

System Defaults:
Private cluster network ports [e0a,e0b,e0c,e0d].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.

Do you want to use these defaults? {yes, no} [yes]: yes
e0a    9000    169.254.155.153 255.255.0.0
e0b    9000    169.254.129.10  255.255.0.0
e0c    9000    169.254.159.42  255.255.0.0
```

```

e0d      9000      169.254.205.233 255.255.0.0

Do you want to use this configuration? {yes, no} [yes]: yes

Step 1 of 5: Create a Cluster
You can type "back", "exit", or "help" at any question.

Enter the cluster name: <<var_site(A/B)_cluster_name>>
Enter the cluster base license key: <<var_cluster_base_license>>

Creating cluster <<var_site(A/B)_cluster>>

Starting cluster support services

Cluster <<var_site(A/B)_cluster_name>> has been created.

Step 2 of 5: Add Feature License Keys
You can type "back", "exit", or "help" at any question.

Enter an additional license key []:<<var_iSCSI_license_key>>
Enter an additional license key []:<<var_NFS_license_key>>
Enter an additional license key []:<<var_FlexClone_license_key>>
Enter an additional license key []:<<var_SnapMirror_license_key>>

Step 3 of 5: Set Up a Vserver for Cluster Administration
You can type "back", "exit", or "help" at any question.

Enter the cluster management interface port []: <<var_site(A/B)_cluster_port>>
Enter the cluster management interface IP address: <<var_site(A/B)_cluster_ip_address>>
Enter the cluster management interface netmask: <<var_oob-mgmt_mask>>
Enter the cluster management interface default gateway: <<var_oob-mgmt_gateway>>

A cluster management interface on port <<var_cluster_mgmt_port>> with IP address
<<var_cluster_mgmt_ip_address>> has been created. You can use this address to connect to and
manage the cluster.

Enter the DNS domain names: <<var_oob-mgmt_domain>>
Enter the name server IP addresses: <<var_oob-mgmt_dns>>
DNS lookup for the admin Vserver will use the <<var_oob-mgmt_dns>> domain.

Step 4 of 5: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.

SFO will be enabled when the partner joins the cluster.

Step 5 of 5: Set Up the Node
You can type "back", "exit", or "help" at any question.

Where is the controller located []:
Enter the node management interface port []: <<var_site(A/B)_node(1/2)_port>>
Enter the node management interface IP address []: <<var_site(A/B)_node(1/2)_ip_address>>

Enter the node management interface netmask []: <<var_oob-mgmt_mask>>
Enter the node management interface default gateway []: <<var_oob-mgmt_gateway>>

```

Join the Cluster

Run the cluster setup wizard on the other node on each site to join it to the cluster created previously.

```

::> cluster setup

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.

Do you want to create a new cluster or join an existing cluster? {create, join}: join

System Defaults:
Private cluster network ports [e0a,e0b,e0c,e0d].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.

Do you want to use these defaults? {yes, no} [yes]: <return>

It can take several minutes to create cluster interfaces...

Step 1 of 3: Join an Existing Cluster
You can type "back", "exit", or "help" at any question.

Enter the name of the cluster you would like to join [<<var_site(A/B)_cluster>>]: <return>

Joining cluster flexpod_mcc_1

Starting cluster support services ..

This node has joined the cluster flexpod_mcc_1.

Step 2 of 3: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.

SFO will be enabled when the partner joins the cluster.

Step 3 of 3: Set Up the Node
You can type "back", "exit", or "help" at any question.

Enter the node management interface port [e0M]: <<var_site(A/B)_node(1/2)_port>>
Enter the node management interface IP address []:<<var_site(A/B)_node(1/2)_ip_address>>

Enter the node management interface netmask [255.255.252.0]:<<var_oob-mgmt_mask>>
Enter the node management interface default gateway [10.228.56.1]:<<var_oob-mgmt_gateway>>

This node has been joined to cluster "flexpod_mcc_1".
To complete cluster setup, you must join each additional node to the cluster
by running "cluster setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP
Software Setup Guide for information about additional system configuration
tasks. You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager
or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster
management IP address (10.228.57.143).
To access the command-line interface, connect to the cluster management
IP address (for example, ssh admin@10.228.57.143).

```

Set Up Storage Failover Mode / ha-config

Note: A new MetroCluster configuration shipped from the factory has the correct HA state set on each controller and chassis, so this step is not required.

In a MetroCluster configuration, the HA state of each controller and chassis should be `mcc` or `mcc-2n`. Because this deployment is a four-node MetroCluster deployment, the HA state is set to `mcc`.

1. Halt each node by running the following command on each cluster:

```
halt -node * -inhibit-takeover true -skip-lif-migration true
```

2. Run `boot_ontap` menu on each node to open the boot menu.

3. Select option 5 from the boot menu and enter `y` when prompted to boot into maintenance mode.

```
Please choose one of the following:
```

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
Selection (1-8)? 5
```

4. Set the `ha-config` values.

```
ha-config modify chassis mcc
ha-config modify controller mcc
```

5. Verify that the new settings are effective.

```
ha-config show
```

6. Halt each node by running the `halt` command.
7. When each node is at the `LOADER>` prompt, boot them by running the `boot_ontap` command.

Disable Flow Control

NetApp recommends disabling flow control on all of the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps.

Note: NetApp recommends running the following commands through a console connection because these operations might result in a temporary drop of the management connection.

1. Run the following command:

```
network port modify -node * -port e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0h,e0i,e0j -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
```

2. Verify that the flow controller has been disabled:

```
flexpod_mcc_1::> network port show -fields flowcontrol-admin
node          port flowcontrol-admin
-----
flexpod_mcc_1-01 e0M full
flexpod_mcc_1-01 e0a none
flexpod_mcc_1-01 e0b none
flexpod_mcc_1-01 e0c none
flexpod_mcc_1-01 e0d none
flexpod_mcc_1-01 e0g none
flexpod_mcc_1-01 e0h none
flexpod_mcc_1-01 e0i none
flexpod_mcc_1-01 e0j none
flexpod_mcc_1-02 e0M full
flexpod_mcc_1-02 e0a none
flexpod_mcc_1-02 e0b none
flexpod_mcc_1-02 e0c none
flexpod_mcc_1-02 e0d none
flexpod_mcc_1-02 e0g none
flexpod_mcc_1-02 e0h none
flexpod_mcc_1-02 e0i none
flexpod_mcc_1-02 e0j none
32 entries were displayed.
```

Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

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

```
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
broadcast-domain create -broadcast-domain IB_MGMT -mtu 9000
```

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, run the following commands on the cluster in each site.

Note: You might need to remove ports from the default broadcast domain before adding it to an interface group.

```
ifgrp create -node <<var_site_(A/B)_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_site_(A/B)_node01>> -ifgrp a0a -port e0g
ifgrp add-port -node <<var_site_(A/B)_node01>> -ifgrp a0a -port e0h

ifgrp create -node <<var_site_(A/B)_node02>> -f ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_site_(A/B)_node02>> -ifgrp a0a -port e0g
ifgrp add-port -node <<var_site_(A/B)_node02>> -ifgrp a0a -port e0h
```

Create VLANs

Note: NetApp recommends running the following commands through a console connection because these operations might result in a temporary drop of the management connection.

Run the following commands on each site to create the VLANs:

```
network port modify -node <<var_site_(A/B)_node01>> -port a0a -mtu 9000
network port modify -node <<var_site_(A/B)_node02>> -port a0a -mtu 9000

network port vlan create -node <<var_site_(A/B)_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create -node <<var_site_(A/B)_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <<var_site_(A/B)_node01>>:a0a-
<<var_nfs_vlan_id>>, <<var_site_(A/B)_node02>>:a0a-<<var_nfs_vlan_id>>

network port vlan create -node <<var_site_(A/B)_node01>> -vlan-name a0a-<<var_iscsi_a_vlan_id>>
network port vlan create -node <<var_site_(A/B)_node02>> -vlan-name a0a-<<var_iscsi_a_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_site_(A/B)_node01>>:a0a-
<<var_iscsi_a_vlan_id>>, <<var_site_(A/B)_node02>>:a0a-<<var_iscsi_a_vlan_id>>

network port vlan create -node <<var_site_(A/B)_node01>> -vlan-name a0a-<<var_iscsi_b_vlan_id>>
network port vlan create -node <<var_site_(A/B)_node02>> -vlan-name a0a-<<var_iscsi_b_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_site_(A/B)_node01>>:a0a-
<<var_iscsi_b_vlan_id>>, <<var_site_(A/B)_node02>>:a0a-<<var_iscsi_b_vlan_id>>

network port vlan create -node <<var_site_(A/B)_node01>> -vlan-name a0a-<<var_ib_mgmt_vlan_id>>
network port vlan create -node <<var_site_(A/B)_node02>> -vlan-name a0a-<<var_ib_mgmt_vlan_id>>
broadcast-domain add-ports -broadcast-domain IB-MGMT -ports <<var_site_(A/B)_node01>>:a0a-
<<var_ib_mgmt_vlan_id>>, <<var_site_(A/B)_node02>>:a0a-<<var_ib_mgmt_vlan_id>>
```

Peer the Clusters

Before the clusters can be peered, the intercluster network interfaces need to be created on both sites. For more information about cluster peering, see the [Cluster Peering Express Guide](#).

1. Run the following commands on site A to create the intercluster network interfaces:

```
network interface create -vserver <<var_siteA_cluster_vserver>> -lif <<var_siteA_icl>> -role
intercluster -home-node <<var_siteA_node1>> -home-port a0a:<<var_ib_mgmt_vlan_id>> -address
<<var_siteA_icl_ip_address>> -netmask <<var_siteA_icl_mask>> -status-admin up
```



```
network interface create -vserver <<var_siteA_cluster_vserver>> -lif <<var_siteA_ic2>> -role
intercluster -home-node <<var_siteA_node2>> -home-port a0a:<<var_ib-mgmt_vlan_id>> -address
<<var_siteA_ic2_ip_address>> -netmask <<var_siteA_ic1_mask>> -status-admin up
```

2. Run the following commands on site B to create the intercluster network interfaces:

```
network interface create -vserver <<var_siteB_cluster_vserver>> -lif <<var_siteB_ic1>> -role
intercluster -home-node <<var_siteB_node1>> -home-port a0a:<<var_ib-mgmt_vlan_id>> -address
<<var_siteB_ic1_ip_address>> -netmask <<var_siteB_ic1_mask>> -status-admin up
```

```
network interface create -vserver <<var_siteB_cluster_vserver>> -lif <<var_siteB_ic2>> -role
intercluster -home-node <<var_siteB_node2>> -home-port a0a:<<var_ib-mgmt_vlan_id>> -address
<<var_siteB_ic2_ip_address>> -netmask <<var_siteB_ic1_mask>> -status-admin up
```

3. Run the following command on site A to create the peering relationship:

```
cluster peer create -peer-addr <<var_siteB_ic1_ip_address>>,<<var_siteB_ic2_ip_address>>
```

4. When prompted, enter a passphrase that will be entered on site B to complete the peering relationship.

5. Run the following command on site B to complete the peering relationship:

```
cluster peer create -peer-addr <<var_siteA_ic1_ip_address>>,<<var_siteA_ic2_ip_address>>
```

6. Reenter the passphrase when prompted to complete the peering between the clusters.
7. Run the `cluster peer show` command to verify that the cluster peering relationship is properly configured.

```
flexpod_mcc_1::*> cluster peer show -instance

Peer Cluster Name: flexpod_mcc_2
Cluster UUID: 2c082459-0c90-11e6-8edb-00a098543c92
Remote Intercluster Addresses: 10.228.58.86, 10.228.58.88
Availability of the Remote Cluster: Available
Remote Cluster Name: flexpod_mcc_2
Active IP Addresses: 10.228.58.86, 10.228.58.88
Cluster Serial Number: 1-80-000013
Remote Cluster Nodes: flexpod_mcc_2-01, flexpod_mcc_2-02
Remote Cluster Health: true
Unreachable Local Nodes: -
Operation Timeout (seconds): 60
Address Family of Relationship: ipv4
Authentication Status Administrative: use-authentication
Authentication Status Operational: ok
Timeout for RPC Connect: 10
Timeout for Update Pings: 5
Last Update Time: 5/31/2016 14:43:21
IPspace for the Relationship: Default
```

Mirror the Aggregates

To configure the MetroCluster relationship, all aggregates must be mirrored and one data aggregate created on each node.

Note: There is an option to create the MetroCluster system with a single data aggregate by passing the `-allow-with-one-aggregate true` option in the `configure` command. We took advantage of this option in this deployment.

1. Run the following command to mirror each aggregate:

```
storage aggregate mirror <<aggregate_id>>
```

2. Run the following command on site A to create a mirrored data aggregate:

```
storage aggregate create siteA_data -diskcount <<number_disks>> -mirror true -node
<<var_siteA_node1>>
```

3. Run the following command on site B to create a mirrored data aggregate:

```
storage aggregate create siteB_data -diskcount <<number_disks>> -mirror true -node
<<var_siteB_node1>>
```

Note: A minimum of 10 disks is needed to create a mirrored NetApp RAID DP® aggregate.

Configure MetroCluster Relationship

To configure the MetroCluster relationship, run the following command on any one site.

During testing, site A was used to set up the MetroCluster system.

```
set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only when directed
to do so by NetApp personnel.
Do you want to continue? {y|n}:y

MetroCluster configure -node-name <<var_siteA_node1>> -allow-with-one-aggregate true
Warning: Using this option will result in creation of both CRS metadata volumes on one
aggregate.

This can result in failure of configuration replication across the two DR
sites if the lone aggregate fails.

When a second data aggregate is created, one of the metadata volumes must be
moved to the new aggregate for increased resiliency using the "volume move"
command (privilege: advanced).
Do you want to continue? {y|n}:y

set -privilege admin
```

Note: The optional parameter `-allow-with-one-aggregate true` was used to save disk space in this deployment.

Create Storage Virtual Machine

To create an SVM, run the following command on both sites:

```
flexpod_mcc_1::> vserver create -vserver Infra_SVM_site_(A/B) -rootvolume rootvol -aggregate
site(A/B)_data -rootvolume-security-style unix
Vserver creation completed

flexpod_mcc_1::> vserver remove-protocols -vserver Infra_SVM_site_(A/B) -protocols fcp,cifs,ndmp
```

Create NFS Service

To create an NFS service, run the following command on both sites:

```
flexpod_mcc_1::> nfs create -vserver Infra_SVM_site_(A/B) -udp disabled

flexpod_mcc_1::> vserver nfs modify -vserver Infra_SVM_site_(A/B) -vstorage enabled

flexpod_mcc_1::> vserver nfs show

Vserver: Infra_SVM_site_(A/B)

    General Access: true
                   v3: enabled
                   v4.0: disabled
                   4.1: disabled
                   UDP: disabled
                   TCP: enabled
  Default Windows User: -
  Default Windows Group: -
```

Create iSCSI Service

To create an iSCSI Service, run the following command on both sites:

```
flexpod_mcc_1::> iscsi create -vserver Infra_SVM_site_(A/B)

flexpod_mcc_1::> iscsi show
      Target
Vserver Name          Target          Status
----- Name          Alias          Admin
-----
Infra_SVM_site_(A/B) iqn.1992-08.com.netapp:sn.8cd828a5106711e68f5400a098540c3c:vs.3
Infra_SVM_site_(A/B) up
```

Configure HTTP Access

To configure HTTP access, run the following command on both sites:

```
flexpod_mcc_1::> set -privilege diag

flexpod_mcc_1::*> security certificate delete -vserver flexpod_mcc_1 -common-name flexpod_mcc_1 -
ca flexpod_mcc_1 -type server -serial 5317BDA8D6236

Warning: Deleting a server certificate will also delete the corresponding
server-chain certificate, if one exists.
Do you want to continue? {y|n}: y

flexpod_mcc_1::*> security certificate delete -vserver Infra_SVM_site_(A/B) -common-name
Infra_SVM_site_(A/B) -ca Infra_SVM_site_(A/B) -type server -serial 531DBBFC1EF15

Warning: Deleting a server certificate will also delete the corresponding
server-chain certificate, if one exists.
Do you want to continue? {y|n}: y

flexpod_mcc_1::> security certificate create -common-name
Infra_SVM_site_(A/B).gdl.englab.netapp.com -type server -size 2048 -country US -state "North
Carolina" -locality "Morrisville" -organization "NetApp" -unit "FlexPod MCC" -email-addr
"Aaron.Kirk@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver
Infra_SVM_site_(A/B)

flexpod_mcc_1::> security certificate create -common-name flexpod-mcc-1.gdl.englab.netapp.com -
type server -size 2048 -country US -state "North Carolina" -locality "Morrisville" -organization
"NetApp" -unit "FlexPod MCC" -email-addr "Aaron.Kirk@netapp.com" -expire-days 365 -protocol SSL -
hash-function SHA256 -vserver flexpod_mcc_1

flexpod_mcc_1::*> security certificate show
Vserver      Serial Number      Common Name          Type
-----
Infra_SVM_site_(A/B) 531DBEAC7335F      Infra_SVM_site_(A/B).netapp.com      server
Certificate Authority: Infra_SVM_site_(A/B).netapp.com
Expiration Date: Tue May 02 13:24:50 2017

flexpod_mcc_1 531DBEEF97011 flexpod-mcc-1.netapp.com      server
Certificate Authority: flexpod-mcc-1.netapp.com
Expiration Date: Tue May 02 13:26:00 2017

2 entries were displayed.

flexpod_mcc_1::*>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

flexpod_mcc_1::*> system services firewall policy delete -policy mgmt -server http -vserver
flexpod_mcc_1

flexpod_mcc_1::*> set -privilege admin

flexpod_mcc_1::> vserver services web modify -name spi_|ontapi|compat -vserver * -enabled true
```

```
1 entry was modified.
```

Configure NFSv3

To configure NFSv3, run the following command on both sites:

```
vserver export-policy rule create -vserver Infra_SVM_site_(A/B) -policyname default -ruleindex 1
-protocol nfs -clientmatch 10.228.10.10 -rorule sys -rwrule sys -superuser sys -allow-suid false

vserver export-policy rule create -vserver Infra_SVM_site_(A/B) -policyname default -ruleindex 2
-protocol nfs -clientmatch 10.228.10.11 -rorule sys -rwrule sys -superuser sys -allow-suid false

vserver export-policy rule create -vserver Infra_SVM_site_(A/B) -policyname default -ruleindex 3
-protocol nfs -clientmatch 10.228.10.12 -rorule sys -rwrule sys -superuser sys -allow-suid false

vserver export-policy rule create -vserver Infra_SVM_site_(A/B) -policyname default -ruleindex 4
-protocol nfs -clientmatch 10.228.10.13 -rorule sys -rwrule sys -superuser sys -allow-suid false

volume modify -vserver Infra_SVM_site_(A/B) -volume rootvol -policy default
```

Create Volumes

To create volumes, run the following command on both sites:

```
volume create -vserver Infra_SVM_site_(A/B) -volume infra_site_(A/B)_datastore_1 -aggregate
site(A/B)_data -size 500g -state online -policy default -junction-path
/infra_site_(A/B)_datastore_1 -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra_SVM_site_(A/B) -volume infra_site_(A/B)_swap -aggregate
site(A/B)_data -size 100g -state online -policy default -junction-path /infra_site_(A/B)_swap -
space-guarantee none -percent-snapshot-space 0 -snapshot-policy none

volume create -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B) -aggregate
site(A/B)_data -size 100g -state online -policy default -space-guarantee none -percent-snapshot-
space 0
```

Create NFS Volumes for Datastore Heartbeats

To create NFS volumes, run the following command on both sites:

Site A

```
volume create -vserver Infra_SVM_site_A -volume site_A_heartbeat -aggregate siteA_data -size 20g
-state online -policy default -junction-path /site_A_heartbeat -space-guarantee none -percent-
snapshot-space 0 -snapshot-policy none
```

Site B

```
volume create -vserver Infra_SVM_site_B -volume site_B_heartbeat -aggregate siteB_data -size 20g
-state online -policy default -junction-path /site_B_heartbeat -space-guarantee none -percent-
snapshot-space 0 -snapshot-policy none
```

Create ESXi Boot LUNs—Site A and Site B

To create ESXi boot LUNs, run the following command on both sites.

```
lun create -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B) -lun VM-Host-Infra-1-01 -
size 15g -ostype vmware -space-reserve disabled

Created a LUN of size 15g (16106127360)

lun create -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B) -lun VM-Host-Infra-1-02 -
size 15g -ostype vmware -space-reserve disabled
```

```
Created a LUN of size 15g (16106127360)
```

Create LUNs for Datastore Heartbeats

To create LUNS for datastore heartbeats, run the following command on both sites:

Site A

```
lun create -vserver Infra_SVM_site_A -volume site_A_heartbeat -lun site_A_heartbeat -size 10g -
ostype vmware -space-reserve disabled
```

Site B

```
lun create -vserver Infra_SVM_site_B -volume site_B_heartbeat -lun site_B_heartbeat -size 10g -
ostype vmware -space-reserve disabled
```

Schedule Deduplication—Site A and Site B

To schedule deduplication volumes, run the following command on both sites:

```
efficiency on -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B)
(volume efficiency on)
Efficiency for volume "esxi_boot_site_(A/B)" of Vserver "Infra_SVM_site_(A/B)" is enabled.

efficiency modify -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B) -schedule sun-sat@0
(volume efficiency modify)

cron create -name lmin -minute
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31,32,33,34,35
,36,37,38,39,40,41,42,43,44,45,46,47,48,48,50,51,52,53,54,55,56,57,58,59
(job schedule cron create)

Warning: Because this is a MetroCluster configuration, an additional step is
required. To complete applying the changes for schedule "lmin", execute the
same command on the remote cluster.

efficiency policy create -vserver Infra_SVM_site_(A/B) -policy Always_On_Deduplication -type
scheduled -schedule lmin -qos-policy background -enabled true
(volume efficiency policy create)

efficiency on -vserver Infra_SVM_site_(A/B) -volume infra_site_(A/B)_datastore_1
(volume efficiency on)
Efficiency for volume "infra_site_(A/B)_datastore_1" of Vserver "Infra_SVM_site_(A/B)" is
enabled.

efficiency modify -vserver Infra_SVM_site_(A/B) -volume infra_site_(A/B)_datastore_1 -policy
Always_On_Deduplication
(volume efficiency modify)
```

Create iSCSI LIFs—Site A and Site B

To create iSCSI LIFs, run the following command on both sites:

```
flexpod_mcc_1::> network interface create -vserver Infra_SVM_site_(A/B) -lif
iscsi_site(A/B)_lif01a -role data -data-protocol iscsi -home-node <<var_site_(A/B)_node01>> -
home-port a0a-<<var_iscsi_a_vlan_id>> -address 10.228.11.20 -netmask 255.255.255.0 -status-admin
up -failover-policy disabled -firewall-policy data -auto-revert false

flexpod_mcc_1::> network interface create -vserver Infra_SVM_site_(A/B) -lif
iscsi_site(A/B)_lif01b -role data -data-protocol iscsi -home-node <<var_site_(A/B)_node01>> -
home-port a0a-<<var_iscsi_b_vlan_id>> -address 10.228.12.20 -netmask 255.255.255.0 -status-admin
up -failover-policy disabled -firewall-policy data -auto-revert false

flexpod_mcc_1::> network interface create -vserver Infra_SVM_site_(A/B) -lif
iscsi_site(A/B)_lif02a -role data -data-protocol iscsi -home-node <<var_site_(A/B)_node02>> -
```

```

home-port a0a-<<var_iscsi_a_vlan_id>> -address 10.228.11.21 -netmask 255.255.255.0 -status-admin
up -failover-policy disabled -firewall-policy data -auto-revert false

flexpod_mcc_1::> network interface create -vserver Infra_SVM_site_(A/B) -lif
iscsi_site(A/B)_lif02b -role data -data-protocol iscsi -home-node <<var_site_(A/B)_node02>> -
home-port a0a-<<var_iscsi_b_vlan_id>> -address 10.228.12.21 -netmask 255.255.255.0 -status-admin
up -failover-policy disabled -firewall-policy data -auto-revert false

flexpod_mcc_1::> network interface show

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster						
	flexpod_mcc_1-01_clus1	up/up	169.254.155.153/16	flexpod_mcc_1-01	e0a	true
	flexpod_mcc_1-01_clus2	up/up	169.254.129.10/16	flexpod_mcc_1-01	e0b	true
	flexpod_mcc_1-01_clus3	up/up	169.254.159.42/16	flexpod_mcc_1-01	e0c	true
	flexpod_mcc_1-01_clus4	up/up	169.254.205.233/16	flexpod_mcc_1-01	e0d	true
	flexpod_mcc_1-02_clus1	up/up	169.254.145.213/16	flexpod_mcc_1-02	e0a	true
	flexpod_mcc_1-02_clus2	up/up	169.254.213.6/16	flexpod_mcc_1-02	e0b	true
	flexpod_mcc_1-02_clus3	up/up	169.254.46.193/16	flexpod_mcc_1-02	e0c	true
	flexpod_mcc_1-02_clus4	up/up	169.254.131.141/16	flexpod_mcc_1-02	e0d	true
Infra_SVM_site_(A/B)						
	iscsi_lif01a	up/down	10.228.11.20/24	flexpod_mcc_1-01	a0a-3341	true
	iscsi_lif01b	up/down	10.228.12.20/24	flexpod_mcc_1-01	a0a-3342	true
	iscsi_lif02a	up/down	10.228.11.21/24	flexpod_mcc_1-02	a0a-3341	true
	iscsi_lif02b	up/down	10.228.12.21/24	flexpod_mcc_1-02	a0a-3342	true
flexpod_mcc_1						
	cluster_mgmt	up/up	10.228.57.143/22	flexpod_mcc_1-01	e0M	true
	flexpod_mcc_1-01_mgmt1	up/up	10.228.57.139/22	flexpod_mcc_1-01	e0M	true
	flexpod_mcc_1-02_mgmt1	up/up	10.228.57.142/22	flexpod_mcc_1-02	e0M	true
	intercluster_1	up/up	10.228.58.82/22	flexpod_mcc_1-01	e0i	true
	intercluster_2	up/up	10.228.58.83/22	flexpod_mcc_1-01	e0i	false
	intercluster_3	up/up	10.228.58.84/22	flexpod_mcc_1-02	e0i	true
	intercluster_4	up/up	10.228.58.85/22	flexpod_mcc_1-02	e0i	false

19 entries were displayed.

Create the NFS LIF—Site A and Site B

To create NFS LIFs, run the following command on both sites:

```

flexpod_mcc_1::> network interface create -vserver Infra_SVM_site_(A/B) -lif
nfs_infra_site_(A/B)_swap -role data -data-protocol nfs -home-node <<var_site_(A/B)_node01>> -
home-port a0a-<<var_nfs_vlan_id>> -address 10.228.10.20 -netmask 255.255.255.0 -status-admin up -
failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

flexpod_mcc_1::> network interface create -vserver Infra_SVM_site_(A/B) -lif
nfs_infra_site_(A/B)_datastore_1 -role data -data-protocol nfs -home-node
<<var_site_(A/B)_node02>> -home-port a0a-<<var_nfs_vlan_id>> -address 10.228.10.21 -netmask
255.255.255.0 -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -
auto-revert true

flexpod_mcc_1::> network interface show -role data

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Infra_SVM_site_(A/B)						
	iscsi_lif01a	up/down	10.228.11.20/24	flexpod_mcc_1-01	a0a-3341	true
	iscsi_lif01b	up/down	10.228.12.20/24	flexpod_mcc_1-01	a0a-3342	true
	iscsi_lif02a	up/down	10.228.11.21/24	flexpod_mcc_1-02	a0a-3341	true
	iscsi_lif02b	up/down	10.228.12.21/24	flexpod_mcc_1-02	a0a-3342	true
	nfs_infra_site_(A/B)_datastore_1	up/down	10.228.10.21/24	flexpod_mcc_1-02	a0a-3340	true
	nfs_infra_site_(A/B)_swap	up/down	10.228.10.20/24	flexpod_mcc_1-01	a0a-3340	true

6 entries were displayed.

5.8 Server Configuration

There are two Cisco UCS domains that need to be configured for this solution. All steps need to be executed in the Cisco UCS domain on site A and the Cisco UCS domain on site B.

Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlexPod environment.

Note: This procedure provisions the Cisco UCS C-Series and B-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS 6248 #1

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect:

```
Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup. (setup/restore)? setup
You have chosen to setup a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y
Enter the password for "admin": <<var_password>>
Enter the same password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)? (yes/no) [n]: y
Which switch fabric (A|B): A
Enter the system name: <<var_ucs_site_(A/B)_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_site(A/B)_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_site(A/B)_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_site(A/B)_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_site(A/B)_cluster_ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

2. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS 6248 #2

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the second Cisco UCS 6248 fabric interconnect:

```
Enter the configuration method: console
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will
be added to the cluster. Continue (y|n)? y
Enter the admin password for the peer fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_site(A/B)_mgmt_ip>>
Apply and save the configuration (select 'no' if you want to re-enter)?
(yes/no): y
```

2. Wait for the login prompt to make sure that the configuration has been saved.

Log in to Cisco UCS Manager

To log in to the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6332-16UP fabric interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. When prompted to accept the security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.

5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 2.2(5d)

This document assumes the use of Cisco UCS 2.2(5d). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248UP fabric interconnect software to version 2.2(5d), refer to the [Cisco UCS Manager Install and Upgrade Guides](#).

Set Anonymous Reporting

To create anonymous reporting, complete the following steps:

1. In the Anonymous Reporting page, select whether or not to send anonymous data to Cisco for improving future products.

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.
If you decide to enable this feature in future, you can do so from the "Anonymous Reporting" in the Call Home settings under the Admin tab.
[View Sample Data](#)

Do you authorize the disclosure of this information to Cisco Smart CallHome?
 Yes No

Don't show this message again.

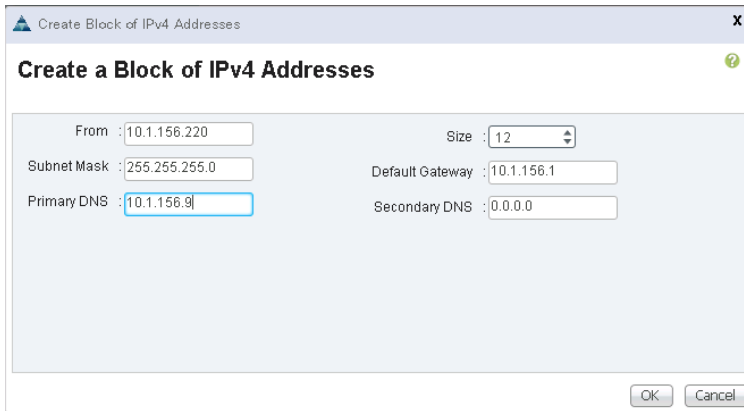
OK Cancel

Add Block of IP Addresses for In-Band KVM Access

Note: In this procedure, two blocks of IP addresses will be created, one on site A and the other on site B. Make sure that the IP addresses created do not overlap.

To create a block of IP addresses for in-band server keyboard, video, and mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.
2. Select Pools > root > IP Pools.
3. Right-click IP Pools and select Create IP Pool.
4. Name the pool `in-band-mgmt`.
5. Click Next.
6. Click Add.
7. Enter the starting IP address of the block, the number of IP addresses required, and the subnet and gateway information.
8. Click Next.
9. Click Finish to create the IP block.

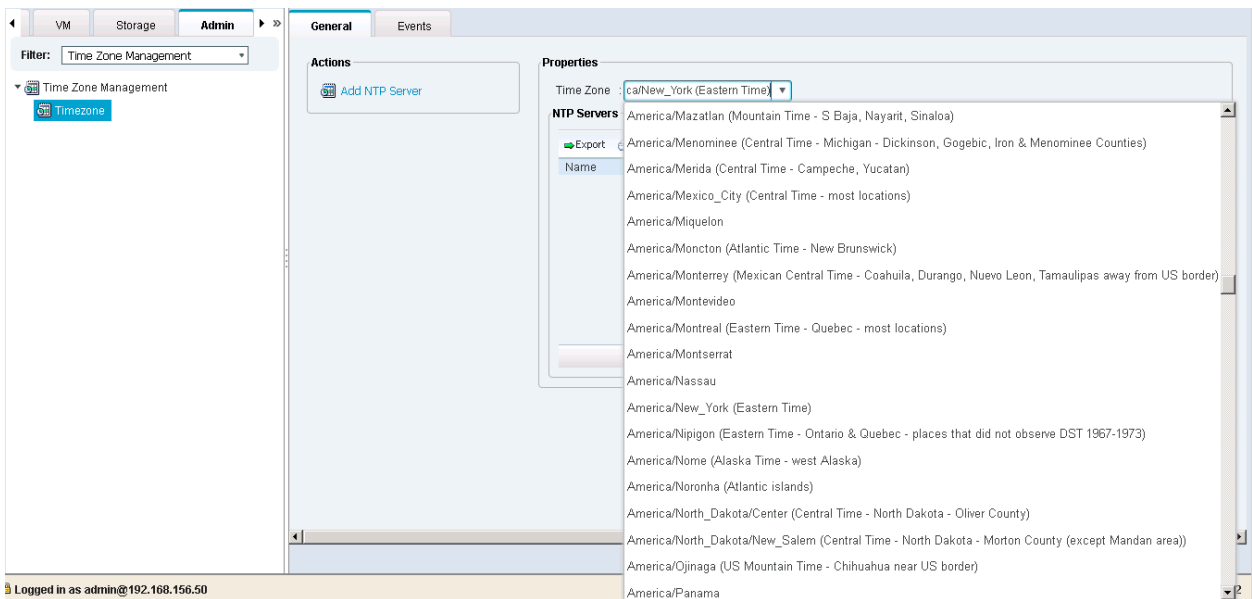


10. Click OK when prompted for confirmation.

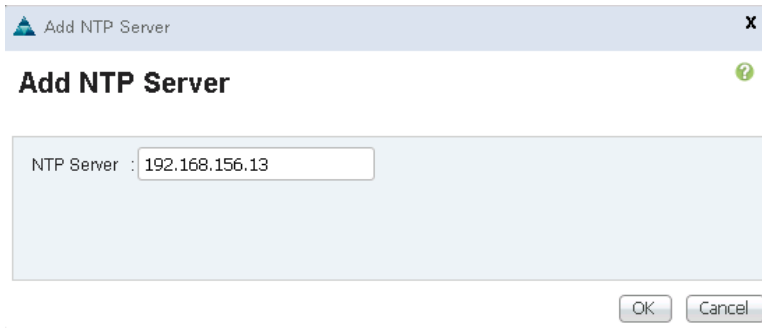
Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:


1. In Cisco UCS Manager, click the Admin tab from the navigation pane.
2. Select All > Timezone Management.



3. From the Properties pane, select the appropriate time zone from the Timezone menu.
4. Click Save Changes and then click OK.
5. Click Add NTP Server.
6. Enter <<var_switch_a_ntp_ip>> and click OK.



7. Click Add NTP Server.
8. Enter <<var_switch_b_ntp_ip>> and click OK.

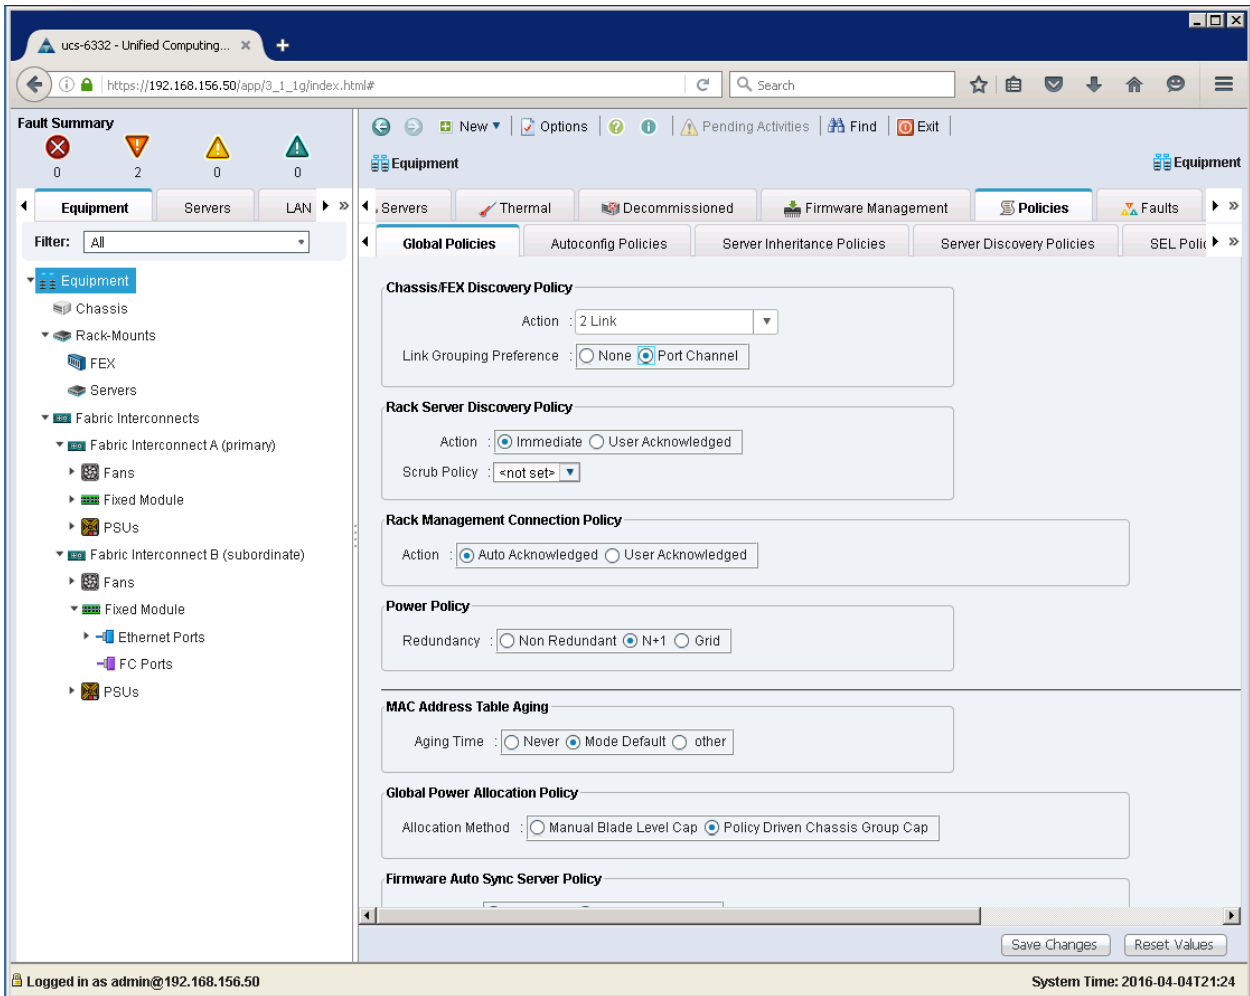


9. Click OK.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and of additional fabric extenders for further C-Series connectivity. To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab from the navigation pane and select Equipment from the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXs) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.



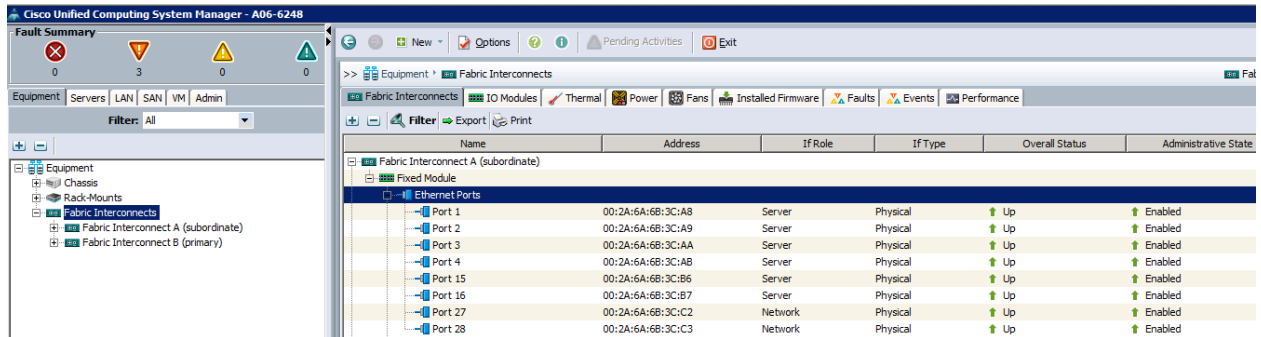
5. Click Save Changes.
6. Click OK.

Enable Server and Uplink Ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab from the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Right-click the ports connected to the chassis and select Configure as Server Port.
5. If using Cisco FEX and UCS C-Series servers, right-click the ports connected to them and select Configure as Server Port.
6. Click Yes to confirm server ports and click OK.
7. Verify that the ports connected to the chassis, C-Series servers, and Cisco FEX are now configured as server ports.
8. Select ports 19 and 20 that are connected to the Cisco Nexus switches. Right-click them and select Configure as Uplink Port.

Note: The last six ports of the Cisco UCS 6332 and 6332-16UP FIs will work only with the optical-based QSFP transceivers and AOC cables. Therefore, they can be better utilized as uplinks to upstream resources that might be optical only.

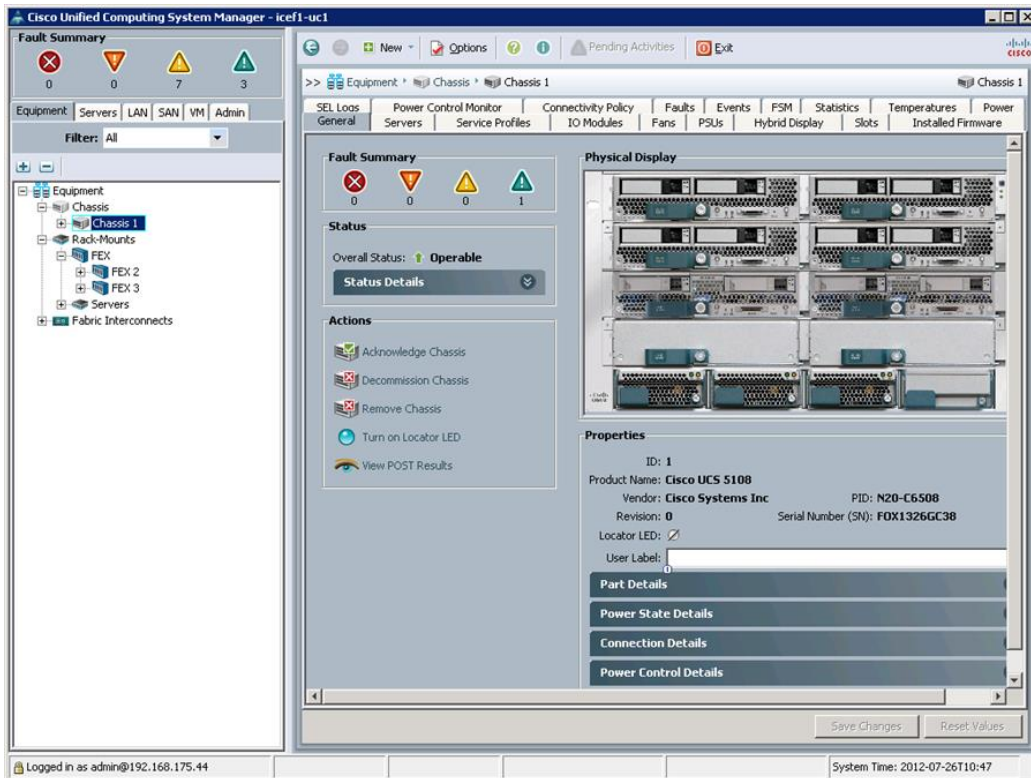


9. Click Yes to confirm uplink ports and click OK.
10. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
11. Expand Ethernet Ports.
12. Right-click the ports connected to the chassis and select Configure as Server Port.
13. If using Cisco FEX and UCS C-Series servers, right-click the ports connected to them and select Configure as Server Port.
14. Click Yes to confirm server ports and click OK.
15. Select ports 19 and 20 that are connected to the Cisco Nexus switches. Right-click them and select Configure as Uplink Port.
16. Click Yes to confirm the uplink ports and click OK.

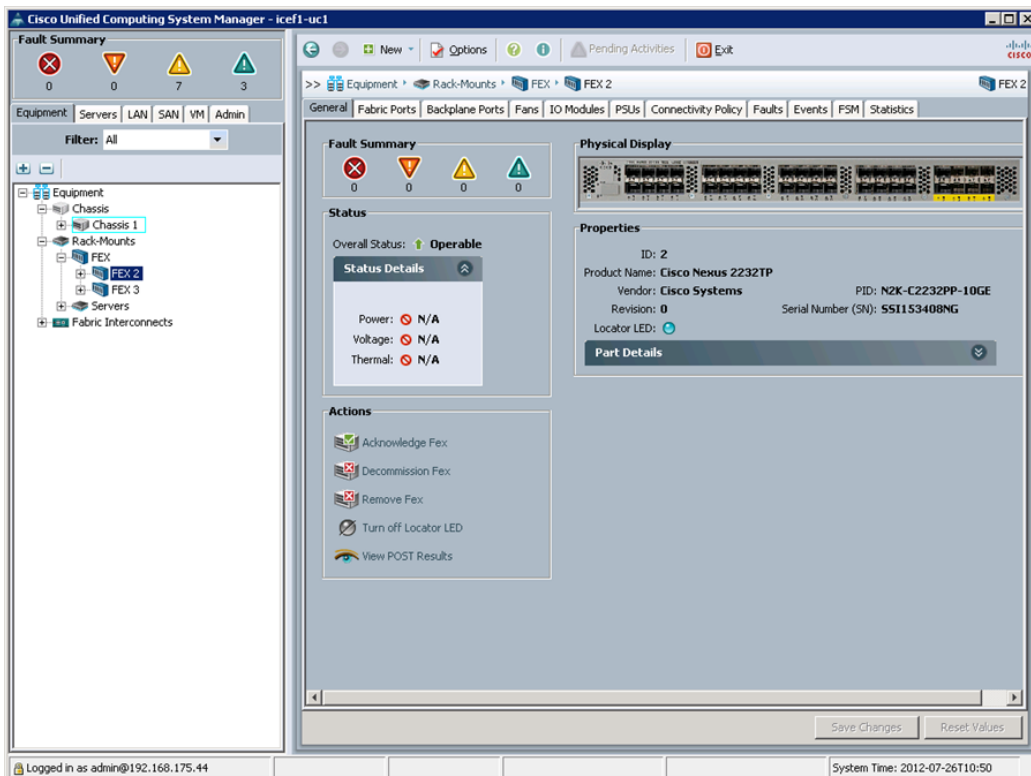
Acknowledge Cisco UCS Chassis and FEX

If using Cisco FEX, acknowledge all Cisco UCS chassis and external 2232 FEX modules by completing the following steps:

1. In Cisco UCS Manager, click the Equipment tab from the navigation pane.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.



4. Click Yes and then click OK to complete acknowledging the chassis.
5. If the Cisco Nexus 2232 FEX is part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and select Acknowledge FEX.



7. Click Yes and then click OK to complete acknowledging the FEX.

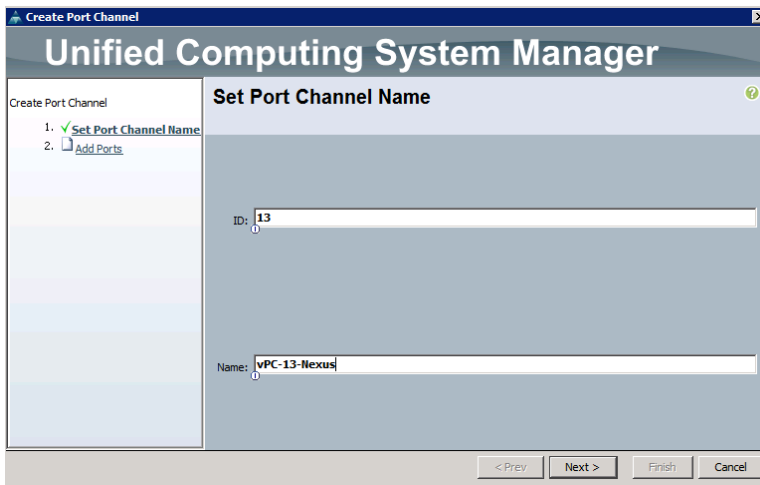
Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.

Note: In this procedure, two port channels are created: one from fabric A to both Cisco Nexus switches and one from fabric B to both Cisco Nexus switches.

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 13 as the unique ID of the port channel.
6. Enter vPC-13-Nexus as the name of the port channel.
7. Click Next.



8. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 19
 - Slot ID 1 and port 20
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter vPC-14-Nexus as the name of the port channel.
17. Click Next.
18. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 19
 - Slot ID 1 and port 20

19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps.

Note: In this procedure, a block of MAC addresses will need to be created for MAC_Pool_A and MAC_Pool_B on both sites. Make sure that the defined MAC addresses in each pool do not overlap.

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.
2. Select Pools > root.
3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC_Pool_A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Click Next.
8. Click Add.
9. Specify a starting MAC address.

Note: For the FlexPod solution, NetApp recommends placing 0A in the next-to-last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric A addresses.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.

Create a Block of MAC Addresses

First MAC Address: Size:

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

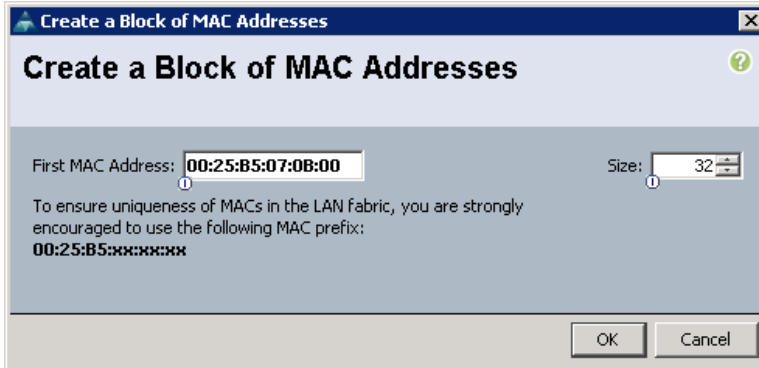
OK Cancel

11. Click OK.
12. Click Finish.
13. Click OK when prompted for confirmation.
14. Right-click MAC Pools under the root organization.
15. Select Create MAC Pool to create the MAC address pool.
16. Enter MAC_Pool_B as the name of the MAC pool.
17. Optional: Enter a description for the MAC pool.
18. Click Next.
19. Click Add.

20. Specify a starting MAC address.

Note: For the FlexPod solution, NetApp recommends placing 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses.

21. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.



22. Click OK.

23. Click Finish.

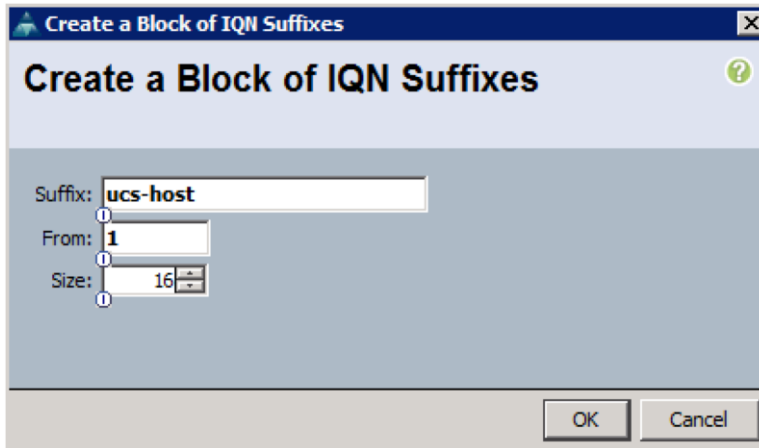
24. Click OK when prompted for confirmation.

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps.

Note: In this procedure, you need to create a block of IQN suffixes on both sites. Make sure that the defined IQN suffixes do not overlap.

1. In the UCS Manager, select the SAN tab on the left.
2. Select Pools > root.
3. Right-click IQN Pools under the root organization.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN_Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter iqn.1992-08.com.cisco as the prefix.
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter ucs-host as the suffix.
12. Enter 1 in the From field.
13. Specify a size of the IQN block sufficient to support the available server resources.
14. Click OK.



15. Click Finish.
16. Click OK when prompted for confirmation.

Create IP Pools for iSCSI Boot

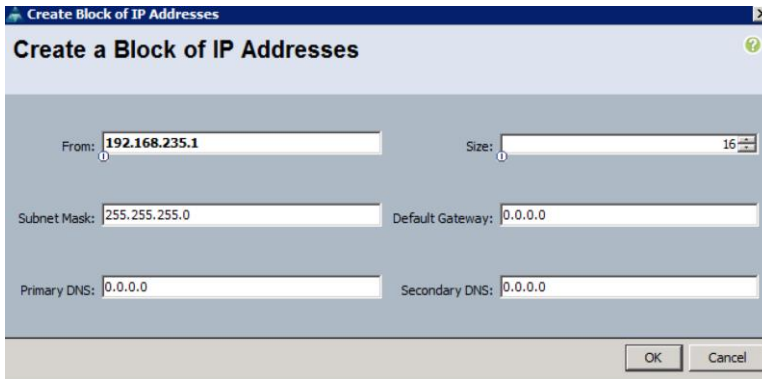
To configure the necessary IP pools for the Cisco UCS environment, complete the following steps.

1. In Cisco UCS Manager, select the LAN tab on the left.
2. Select Pools > root.

Note: Two IP pools are created per site, one for each switching fabric. Make sure that the IPs in the same switching fabric on the two sites do not overlap.

3. Right-click IP Pools under the root organization.
4. Select Create IP Pool to create the IP pool.
5. Enter iSCSI_IP_Pool_A for the name of the IP pool.
6. Optional: Enter a description of the IP pool.
7. Select Sequential for Assignment Order.
8. Click Next.
9. Click Add.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Set the size to enough addresses to accommodate the servers.
12. Click OK.
13. Click Finish.
14. Right-click IP Pools under the root organization.
15. Select Create IP Pool to create the IP pool.
16. Enter iSCSI_IP_Pool_B for the name of the IP pool.
17. Optional: Enter a description of the IP pool.
18. Select Sequential for Assignment Order.
19. Click Next.
20. Click Add.
21. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
22. Set the size to enough addresses to accommodate the servers.
23. Click OK.

24. Click Finish.



Create a Block of IP Addresses

From: Size:

Subnet Mask: Default Gateway:

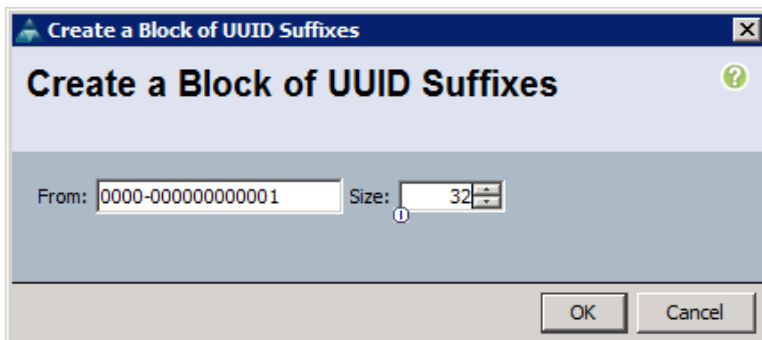
Primary DNS: Secondary DNS:

OK Cancel

Create UUID Suffix Pool

To configure the universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID_Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Click Next.
9. Click Add to add a block of UUIDs.
10. Keep the From field at the default setting.
11. Specify a size for the UUID block that is sufficient to support the available blade or server resources.



Create a Block of UUID Suffixes

From: Size:

OK Cancel

12. Click OK.
13. Click Finish.
14. Click OK.

Create Server Pool

To configure the server pools for the Cisco UCS environment, complete the following steps.

Note: Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter Infra_Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra_Pool server pool.
9. Click Finish.
10. Click OK.

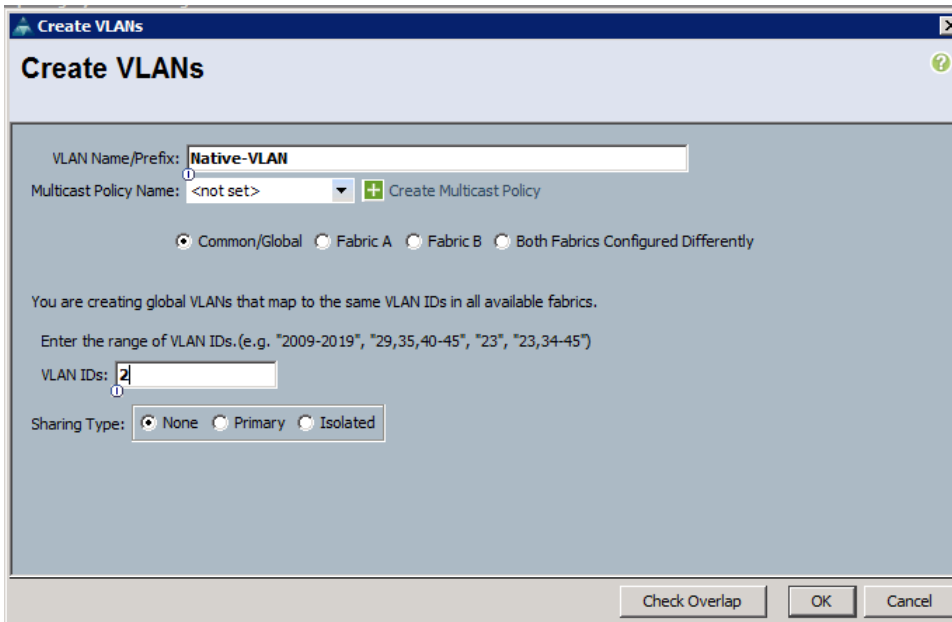
Create VLANs

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

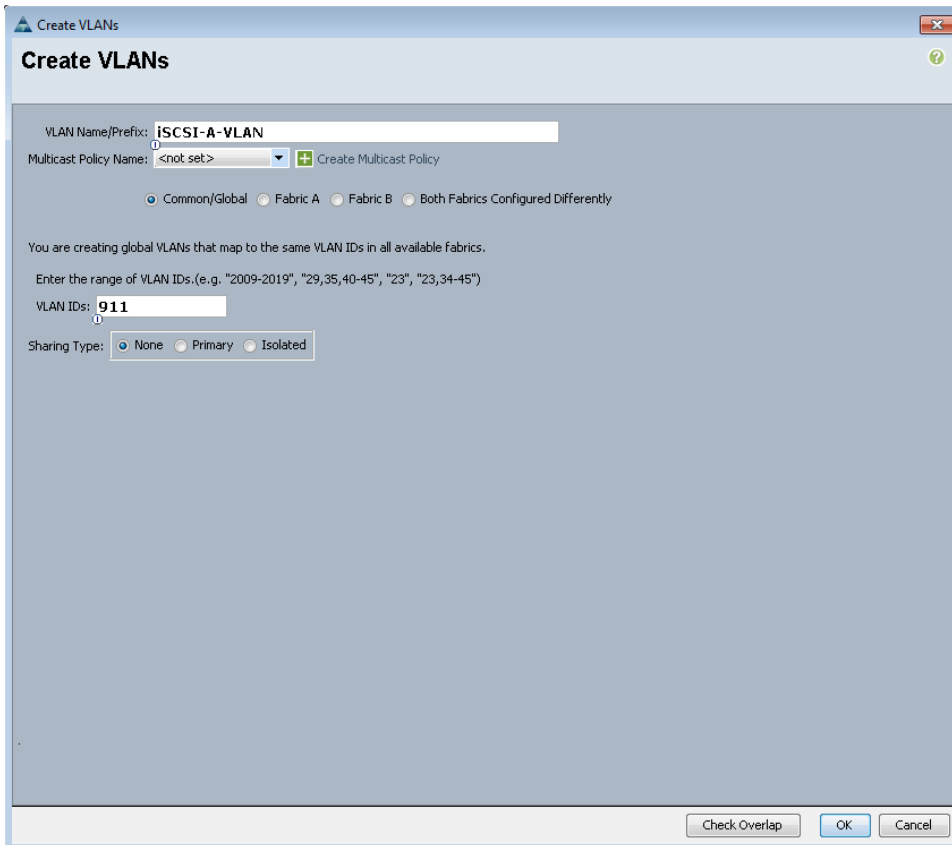
1. In Cisco UCS Manager, click the LAN tab from the navigation pane.

Note: In this procedure, four unique VLANs are created.

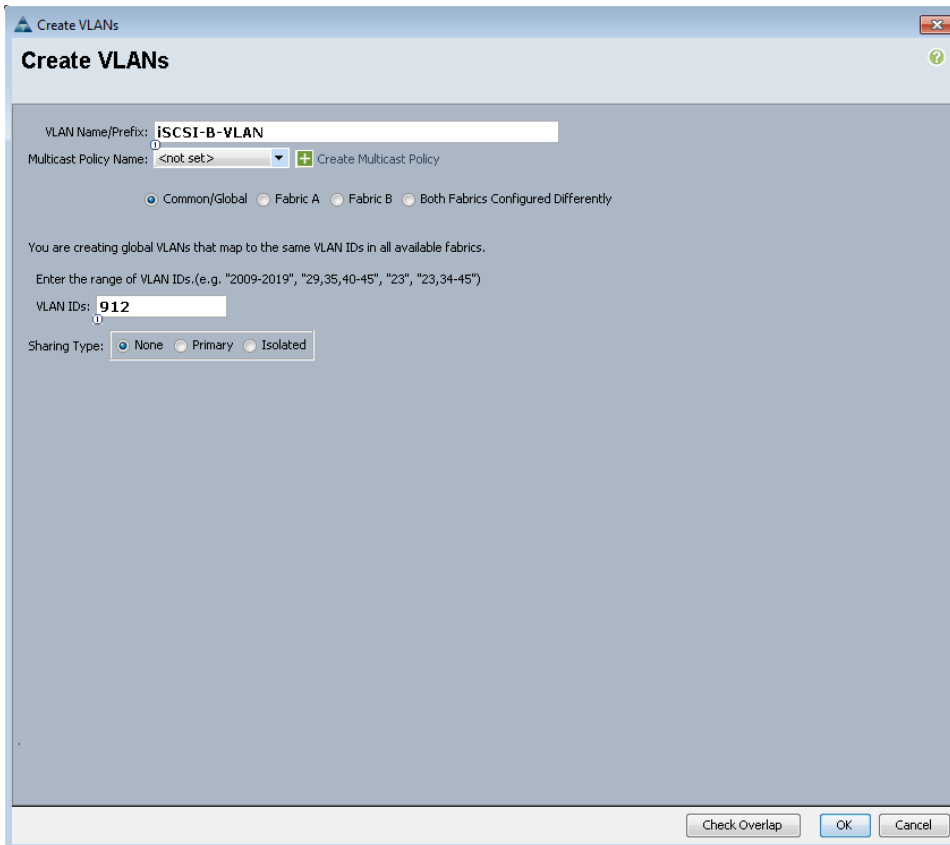
2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `Native-VLAN` as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.



10. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN, and select Set as Native VLAN.
11. Click Yes and then click OK.
12. Right-click VLANs.
13. Select Create VLANs.
14. Enter iSCSI-A-VLAN as the name of the VLAN to be used for the first iSCSI VLAN.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the VLAN ID for the first iSCSI VLAN.
17. Click OK and then click OK.



18. Right-click VLANs.
19. Select Create VLANs.
20. Enter `iSCSI-B-VLAN` as the name of the VLAN to be used for the second iSCSI VLAN.
21. Keep the Common/Global option selected for the scope of the VLAN.
22. Enter the VLAN ID for the second iSCSI VLAN.
23. Click OK and then click OK.



24. Right-click VLANs.
25. Select Create VLANs.
26. Enter `IB-Mgmt` as the name of the VLAN to be used for management traffic.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the in-band management VLAN ID.
29. Keep the Sharing Type as None.
30. Click OK and then click OK again.
31. Right-click VLANs.
32. Select Create VLANs.
33. Enter `INFRA-NFS` as the name of the VLAN to be used for NFS.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the NFS VLAN ID.
36. Keep the Sharing Type as None.
37. Click OK and then click OK again.
38. Right-click VLANs.
39. Select Create VLANs.
40. Enter `vMotion` as the name of the VLAN to be used for vMotion.
41. Keep the Common/Global option selected for the scope of the VLAN.
42. Enter the vMotion VLAN ID.
43. Keep the Sharing Type as None.

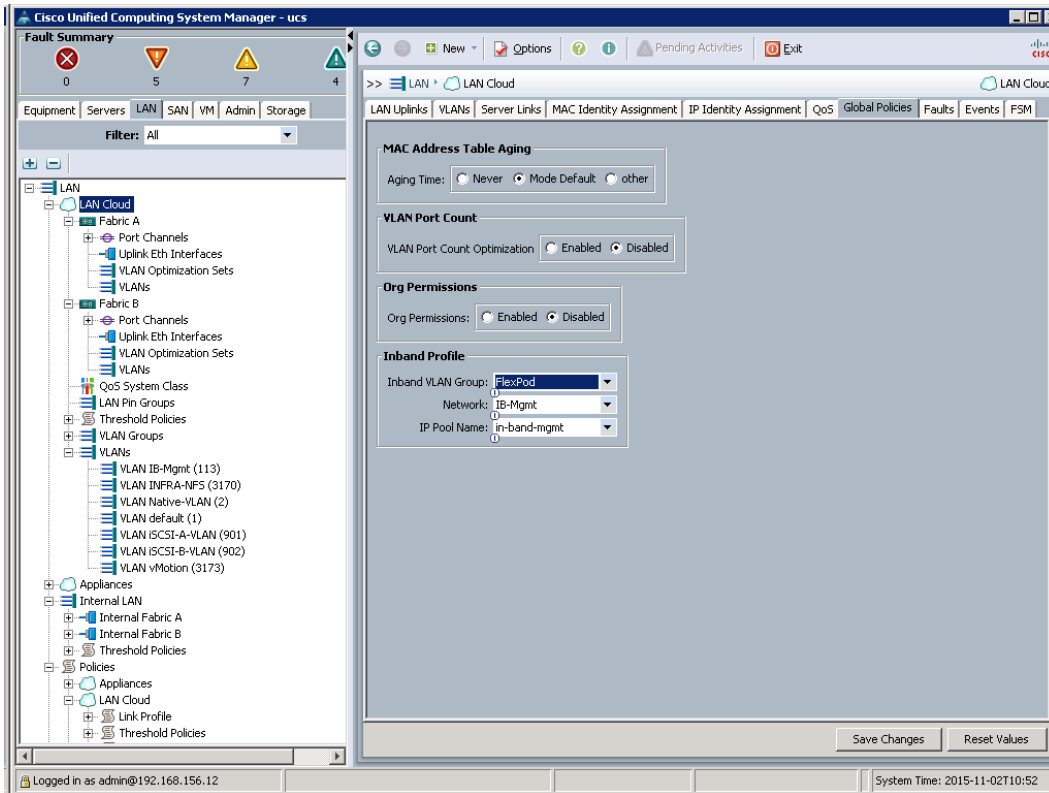
44. Click OK and then click OK again.
45. Right-click VLANs.
46. Select Create VLANs.
47. Enter `VM-Traffic` as the name of the VLAN to be used for VM Traffic.
48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VM-Traffic VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK and then click OK again.

Create VLAN Group and Assign In-Band Profile

A VLAN group is required to set up in-band KVM access.

To create VLAN groups, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLAN Groups and select Create VLAN Group.
4. Name the VLAN group `FlexPod` and select all VLANs.
5. Select the Native-VLAN button and click Next.
6. Click Next.
7. Select the two uplink port channels and use the >> button to add them to the VLAN group.
8. Click Finish.
9. Click OK.
10. Select LAN > LAN Cloud. Then, select the Global Policies tab.
11. In the Inband Profile box, select the FlexPod VLAN group, the IB-MGMT network, and the in-band-mgmt IP pool name.
12. Select Save Changes and click OK.




Create Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

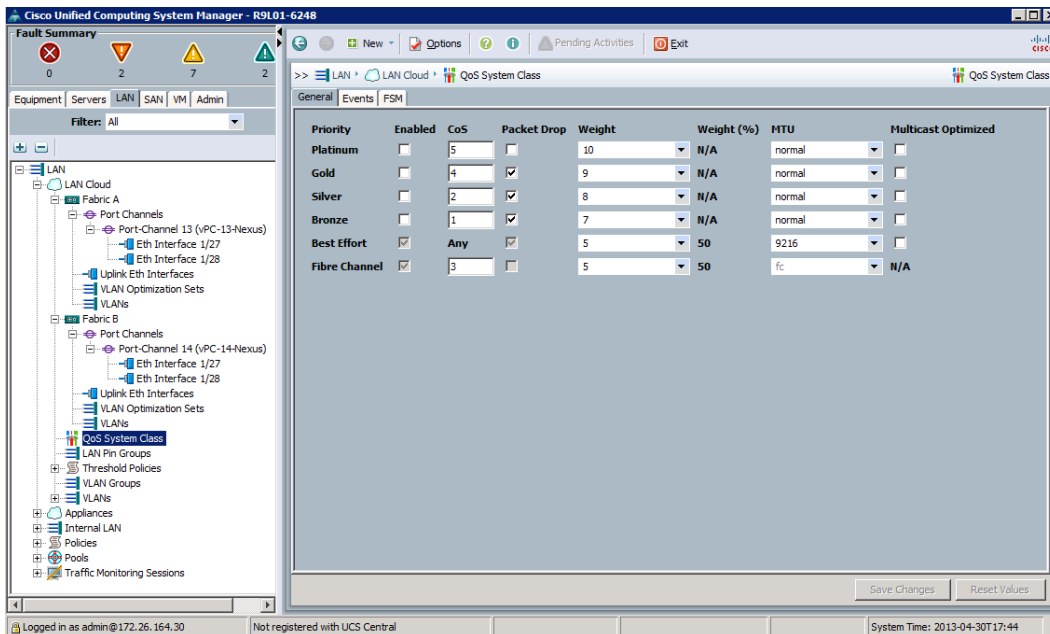
1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. From the Actions pane, select Modify Package Versions.
6. Select version 2.2(5d) for both the Blade and Rack Packages.
7. Click OK to modify the host firmware package.



Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes.
6. Click OK



Create Local Disk Configuration Policy

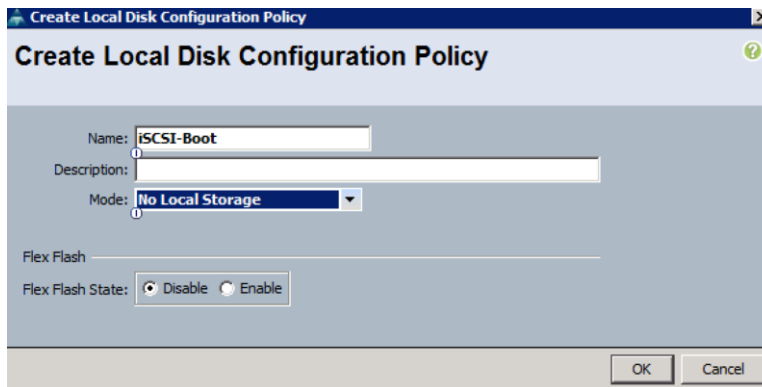
NetApp recommends the use of SAN boot to provide storage to the servers for loading the operating system.

A local disk configuration policy will be created to define No Local Storage.

Note: This policy should not be used on servers that contain local disks.

To create a local disk configuration policy, complete the following steps on each site:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter `iSCSI-Boot` as the local disk configuration policy name.
6. Change the mode to No Local Storage.
7. Click OK to create the local disk configuration policy.



8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable_CDP` as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.

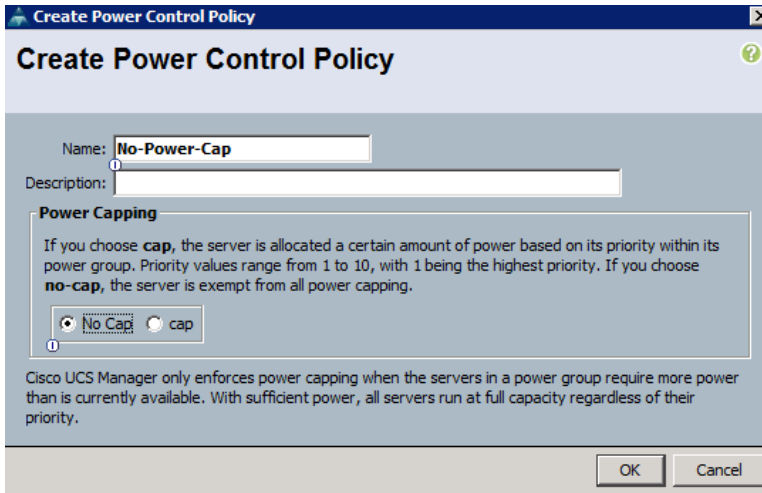


8. Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

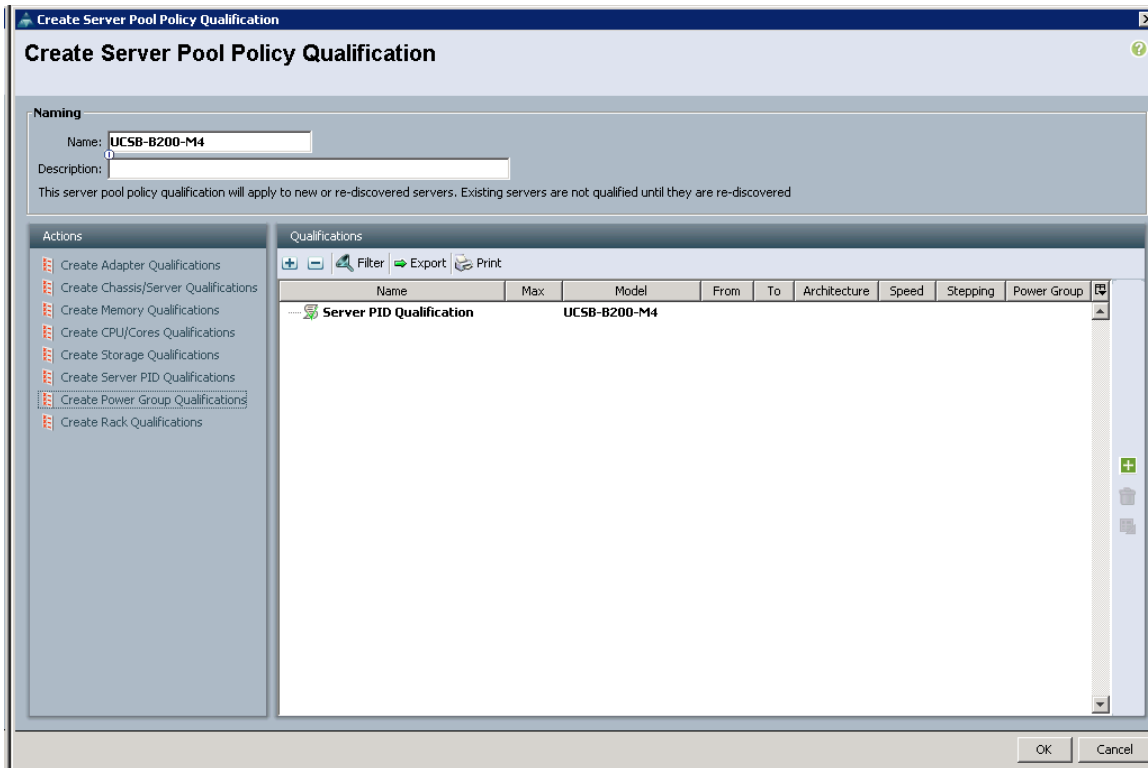


Create Server Pool Qualification Policy (Optional)

To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:

Note: This example creates a policy for a Cisco UCS B200-M4 server.

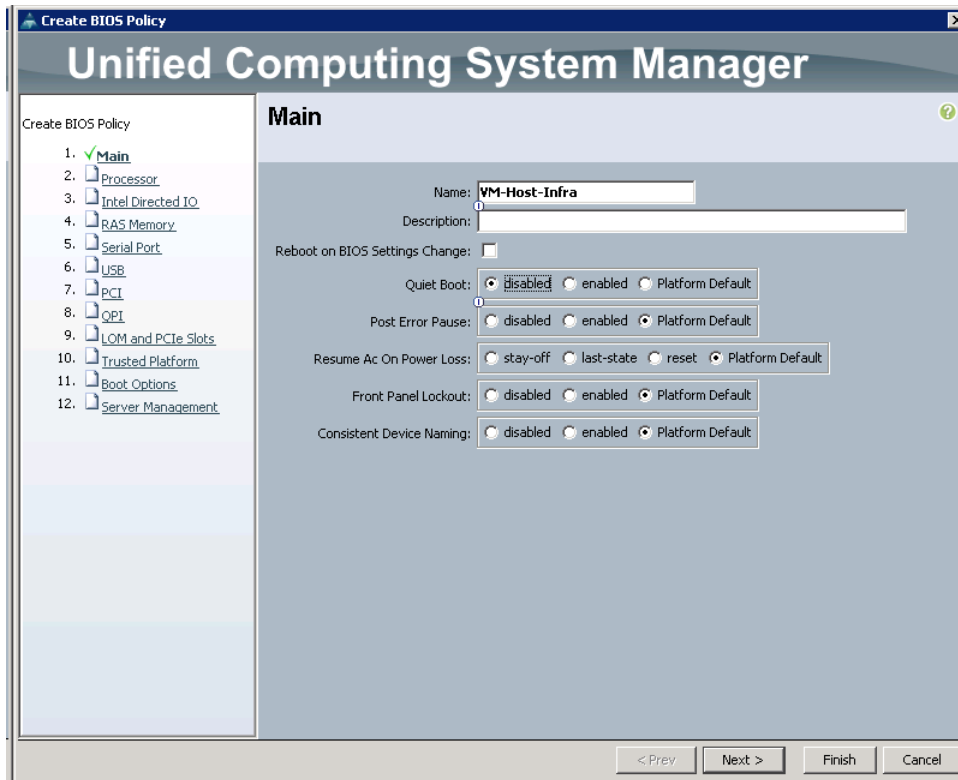
1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCSB-B200-M4.
6. Select Create Server PID Qualifications.
7. Select UCSB-B200-M4 as the name.
8. Click OK to create the server PID qualification.
9. Click OK to create the policy, then click OK when prompted for confirmation.



Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter VM-Host-Infra as the BIOS policy name.
6. Change the Quiet Boot setting to Disabled.
7. Click Finish to create the BIOS policy.

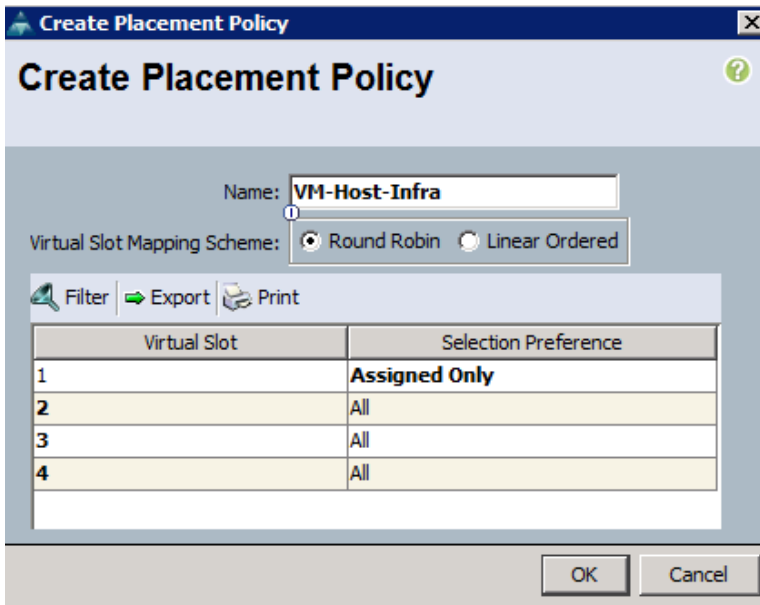


8. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

To create a vNIC/vHBA placement policy for the infrastructure hosts, complete the following steps:

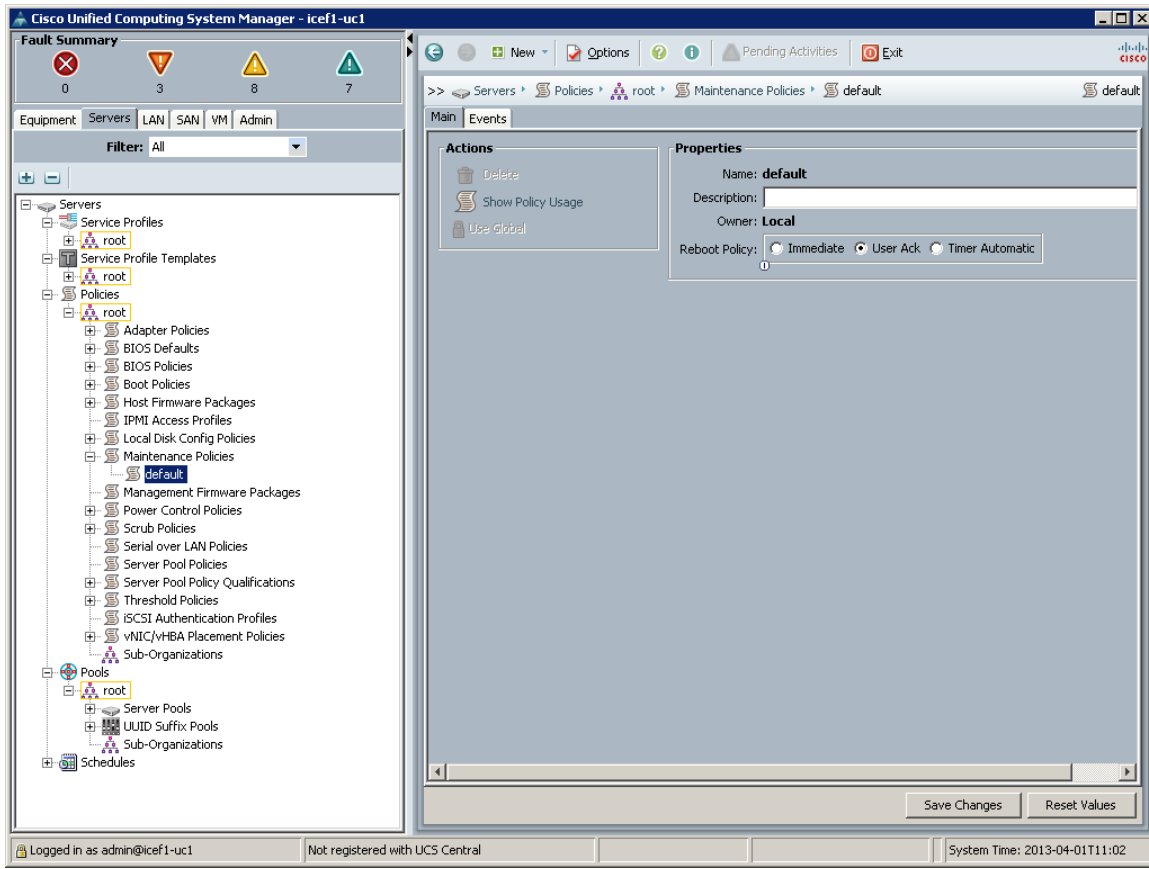
1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter VM-Host-Infra as the name of the placement policy.
6. Click 1 and select Assigned Only.
7. Click OK and then click OK again.



Update Default Maintenance Policy

To update the default maintenance policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Click Save Changes.
6. Click OK to accept the change.



Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of four vNIC templates will be created.

Data vNICs

1. In Cisco UCS Manager, click the LAN tab from the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.
10. Under VLANs, select the checkboxes for IB-MGMT, INFRA-NFS, Native-VLAN, VM-Traffic, and vMotion VLANs.
11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. From the MAC Pool list, select MAC_Pool_A.
14. From the Network Control Policy list, select Enable_CDP.

15. Click OK to create the vNIC template.
16. Click OK.

Create vNIC Template

Name:

Description:

Fabric ID: Fabric A Fabric B Enable Failover

Target

Adapter
 VM

Warning
If VM is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: Initial Template Updating Template

VLANs

Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	INFRRA-NFS	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	VM-Traffic	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A-VLAN	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

+ Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy:

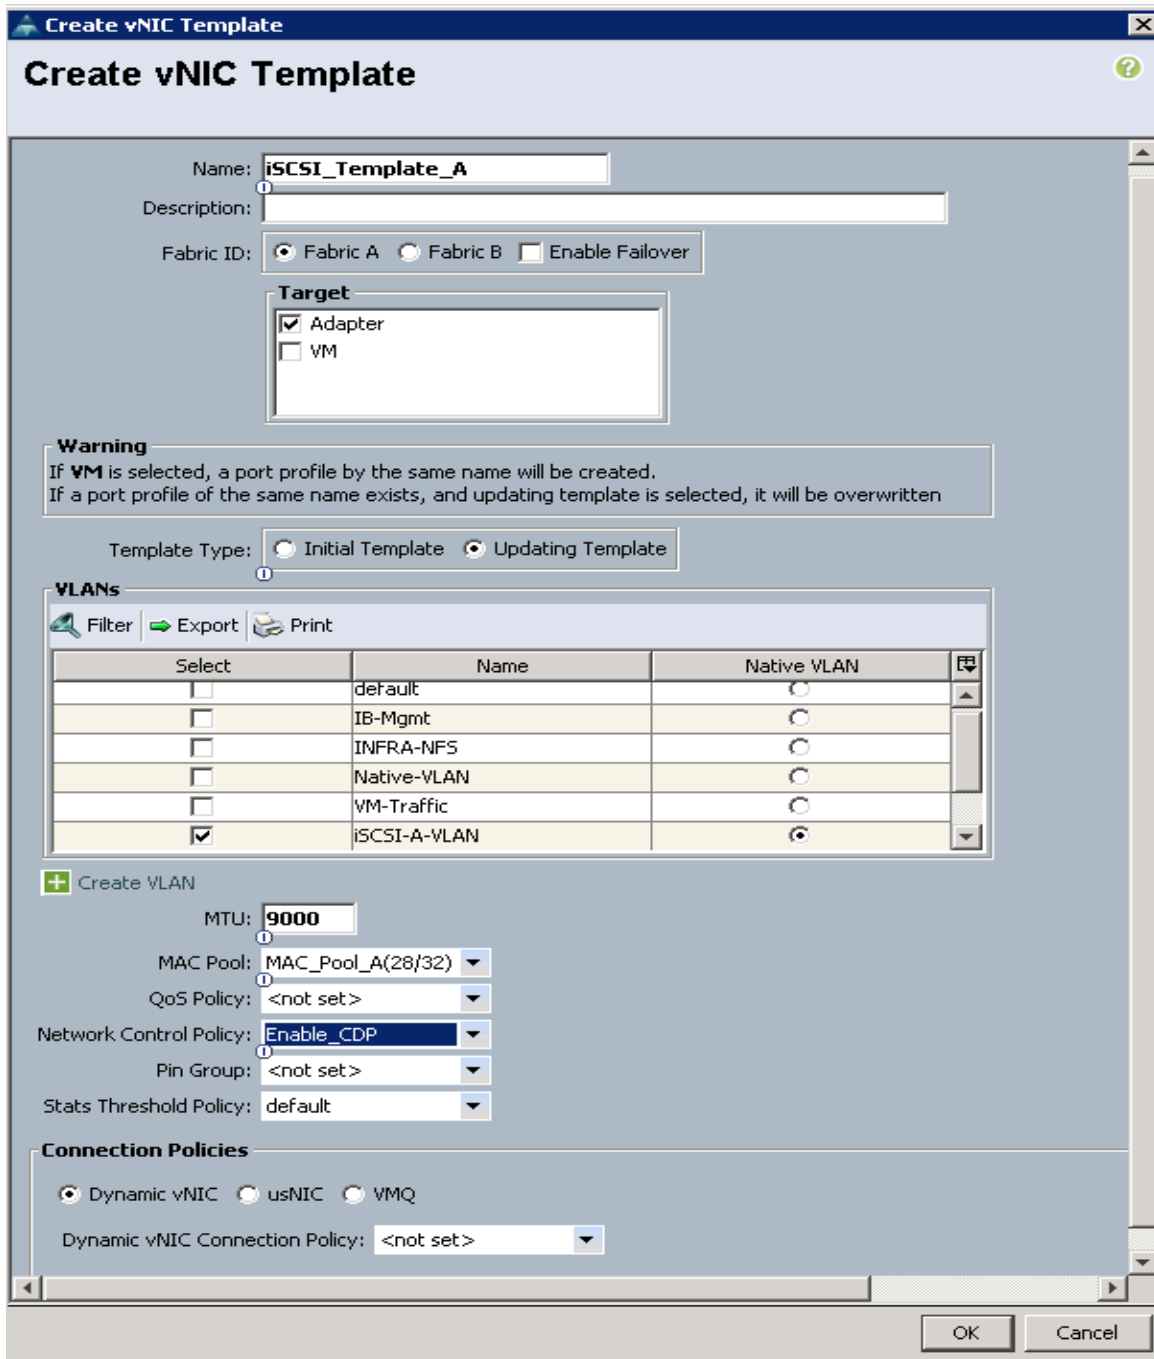
OK Cancel

17. From the navigation pane, select the LAN tab.
18. Select Policies > root.
19. Right-click vNIC Templates.

20. Select Create vNIC Template.
21. Enter `vNIC_Template_B` as the vNIC template name.
22. Select Fabric B.
23. Do not select the Enable Failover checkbox.
24. Under Target, make sure that the VM checkbox is not selected.
25. Select Updating Template as the template type.
26. Under VLANs, select the checkboxes for IB-MGMT, INFRA-NFS, Native-VLAN, and vMotion VLANs.
27. Set default as the native VLAN.
28. Enter 9000 for the MTU.
29. From the MAC Pool list, select `MAC_Pool_B`.
30. From the Network Control Policy list, select `Enable_CDP`.
31. Click OK to create the vNIC template.
32. Click OK.

Create iSCSI vNICs

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `iSCSI_Template_A` as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.
9. Under VLANs, select iSCSI-A-VLAN.
10. Set iSCSI-A-VLAN as the native VLAN.
11. Under MTU, enter 9000.
12. From the MAC Pool list, select `MAC_Pool_A`.
13. From the Network Control Policy list, select `Enable_CDP`.
14. Click OK to complete creating the vNIC template.
15. Click OK.



16. Select the LAN tab on the left.
17. Select Policies > root.
18. Right-click vNIC Templates.
19. Select Create vNIC Template.
20. Enter `iSCSI_Template_B` as the vNIC template name.
21. Select Fabric B. Do not select the Enable Failover checkbox.
22. Under Target, make sure that the VM checkbox is not selected.

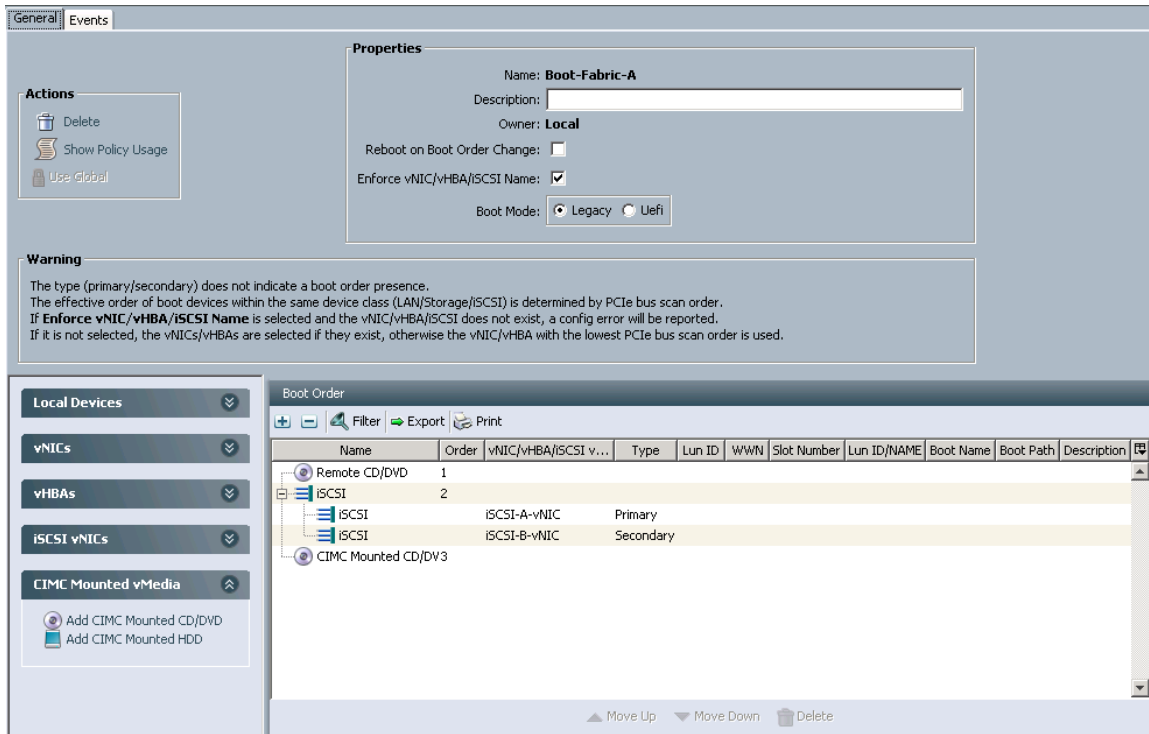
23. Select Updating Template for Template Type.
24. Under VLANs, select iSCSI-B-VLAN.
25. Set iSCSI-B-VLAN as the native VLAN.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select MAC_Pool_B.
28. From the Network Control Policy list, select Enable_CDP.
29. Click OK to complete creating the vNIC template.
30. Click OK.

Create Boot Policies

This procedure applies to a Cisco UCS environment in which two iSCSI LIFs are on cluster node 1 (iscsi lif01a and iscsi lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi lif02a and iscsi lif02b). One boot policy is configured in this procedure. This policy configures the primary target to be iscsi_lif01a.

To create boot policies for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-Fabric-A` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
9. Expand the iSCSI vNICs section and select Add iSCSI Boot.
10. In the Add iSCSI Boot dialog box, enter iSCSI-A-vNIC.
11. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
14. Click OK.
15. Expand CIMC Mounted vMedia.
16. Select Add CIMC Mounted CD/DVD.
17. Click OK.
18. Click OK to save the boot policy. Click OK to close the Boot Policy page.

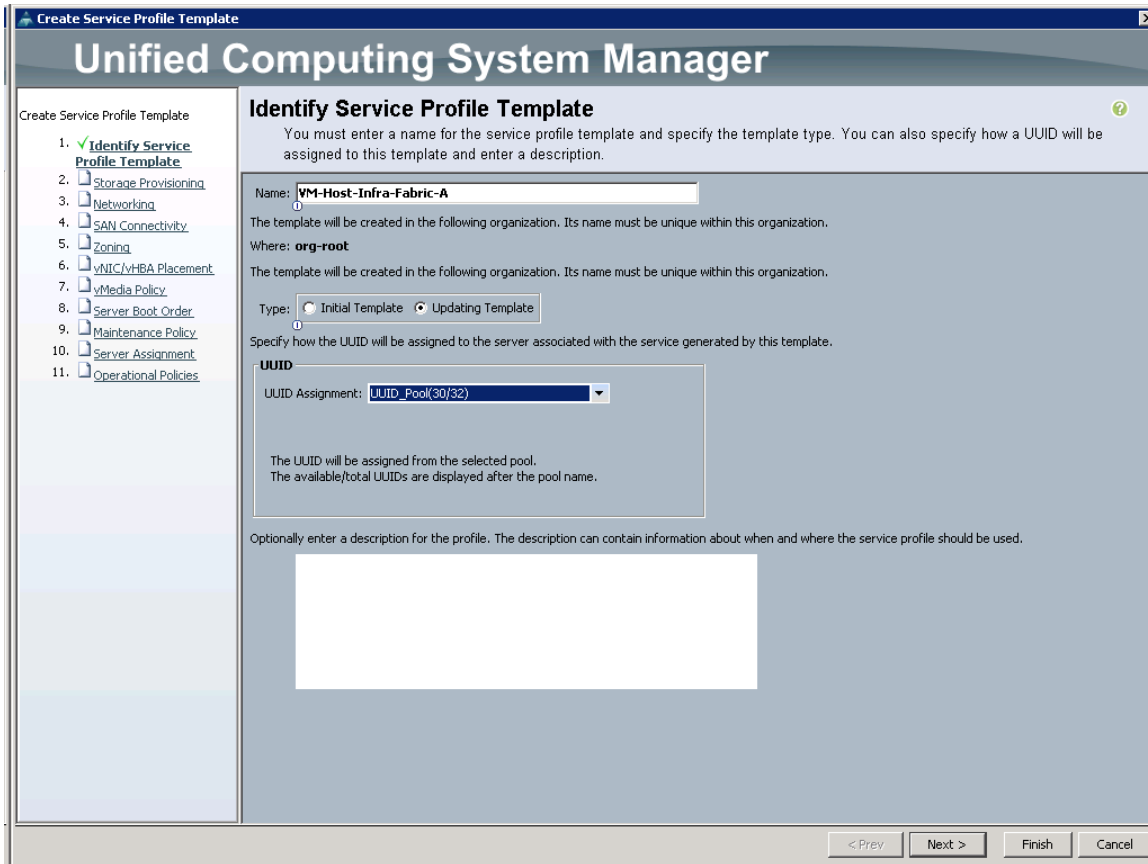


Create Service Profile Template

In this procedure, one service profile template for Infrastructure ESXi hosts is created for fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the Updating Template option.
7. Under UUID, select UUID_Pool as the UUID pool.
8. Click Next.

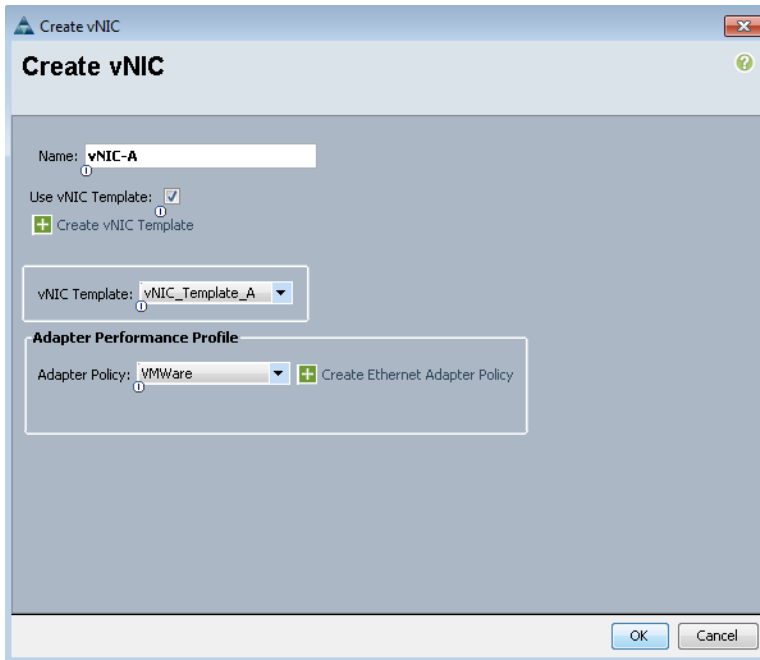


Configure Storage Provisioning

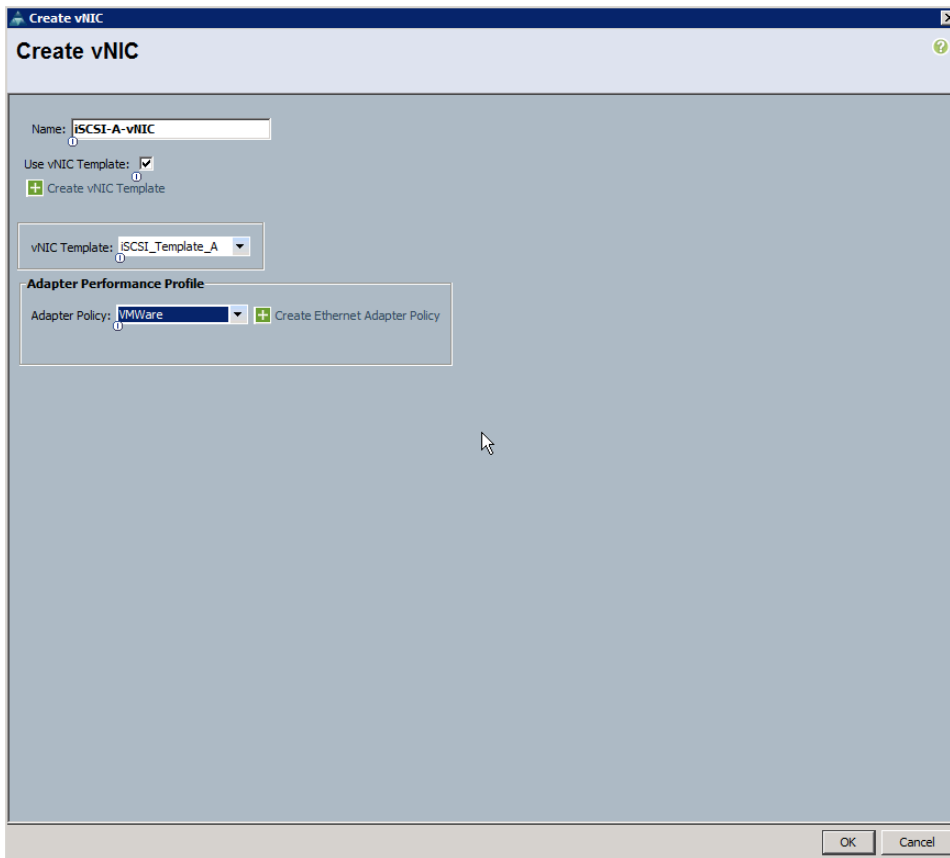
1. If you have servers with no physical disks, select iSCSI-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

Configure Networking Options

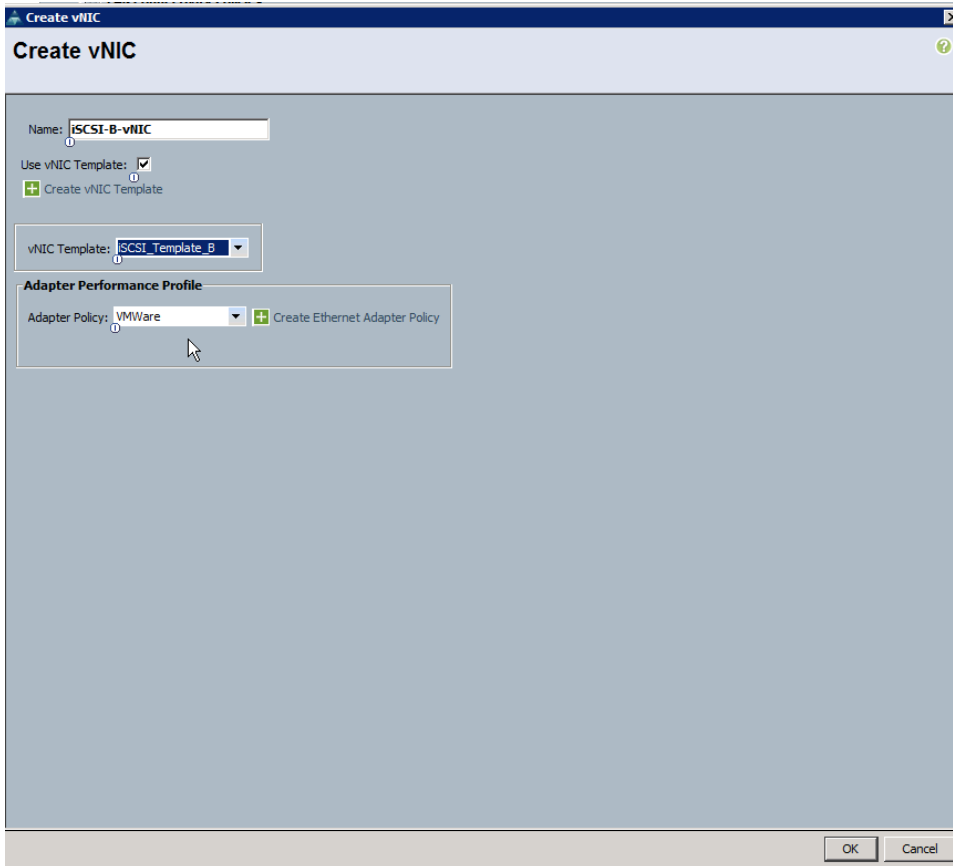
1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Expert option to configure the LAN connectivity.
3. Click the upper Add button to add a vNIC to the template.
4. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
5. Select the Use vNIC Template checkbox.
6. From the vNIC Template list, select vNIC_Template_A.
7. From the Adapter Policy list, select VMWare.
8. Click OK to add this vNIC to the template.



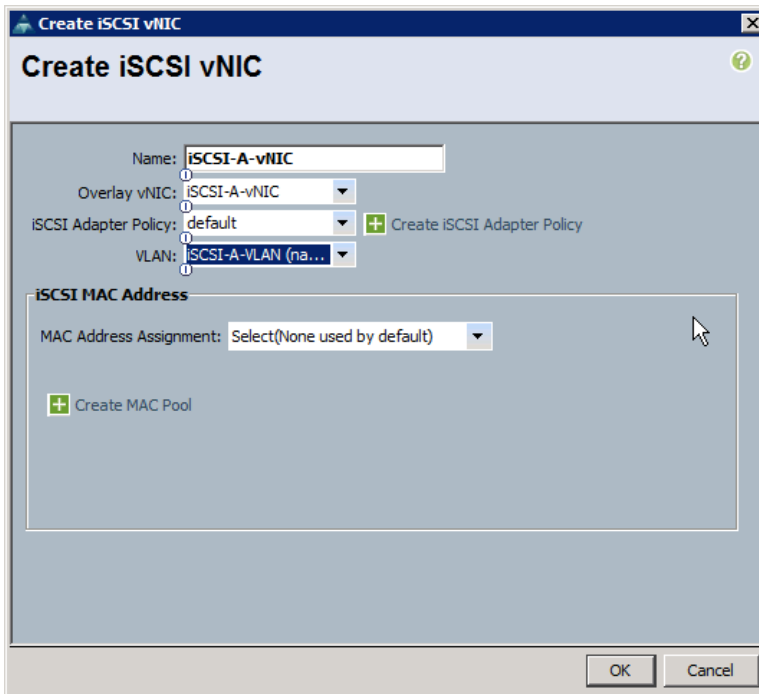
9. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.
10. In the Create vNIC box, enter `vNIC-B` as the name of the vNIC.
11. Select the Use vNIC Template checkbox.
12. From the vNIC Template list, select `vNIC_Template_B`.
13. From the Adapter Policy list, select `VMWare`.
14. Click OK to add the vNIC to the template.
15. Click the upper Add button to add a vNIC to the template.
16. In the Create vNIC dialog box, enter `iSCSI-A-vNIC` as the name of the vNIC.
17. Select the Use vNIC Template checkbox.
18. From the vNIC Template list, select `iSCSI_Template_A`.
19. From the Adapter Policy list, select `VMWare`.
20. Click OK to add this vNIC to the template.



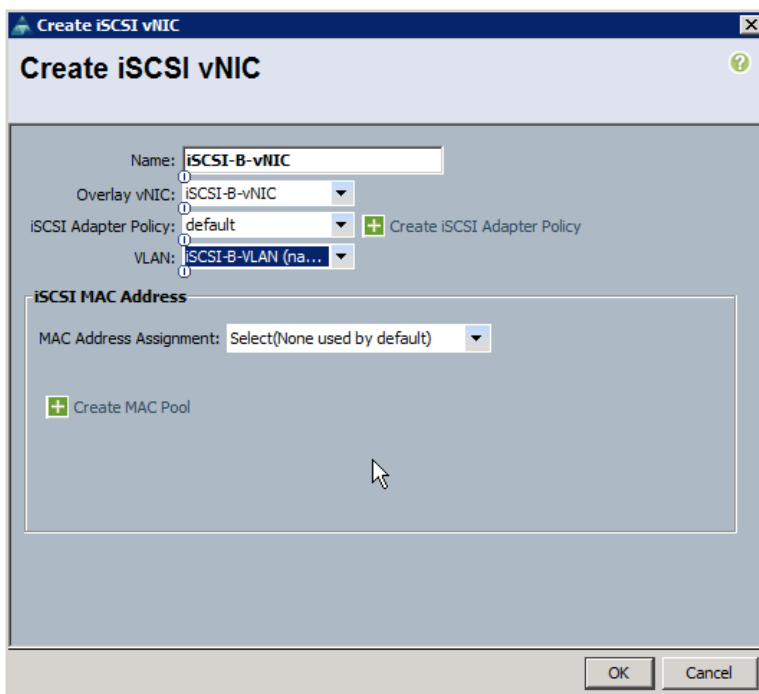
21. Click the upper Add button to add a vNIC to the template.
22. In the Create vNIC dialog box, enter `iSCSI-B-vNIC` as the name of the vNIC.
23. Select the Use vNIC Template checkbox.
24. From the vNIC Template list, select `iSCSI_Template_B`.
25. From the Adapter Policy list, select `VMWare`.
26. Click OK to add this vNIC to the template.



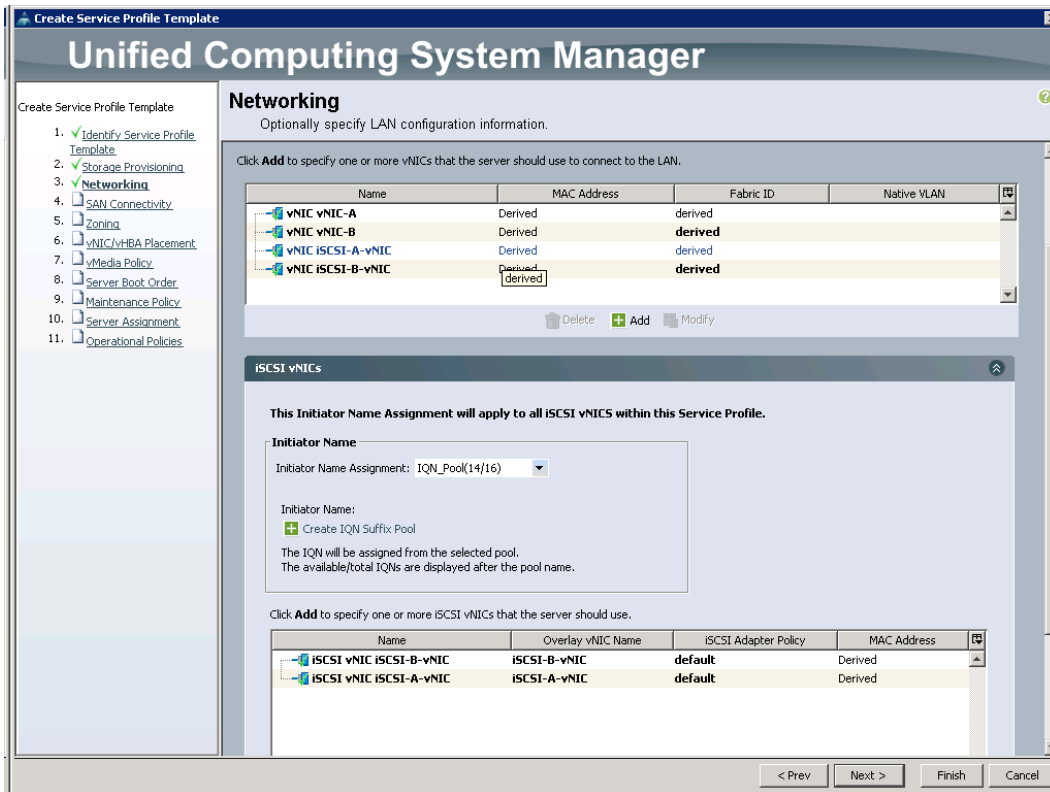
27. Expand the iSCSI vNICs section (if not already expanded).
28. Select IQN-Pool under Initiator Name Assignment.
29. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
30. Enter `iSCSI-A-vNIC` as the name of the vNIC.
31. Select `iSCSI-A-vNIC` for Overlay vNIC.
32. Set the iSCSI Adapter Policy to default.
33. Set the VLAN to `iSCSI-A-VLAN`.
34. Leave the MAC Address set to None.
35. Click OK.



36. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
37. Enter iSCSI-B-vNIC as the name of the vNIC.
38. Set the Overlay vNIC to iSCSI-B-vNIC
39. Set the iSCSI Adapter Policy to default.
40. Set the VLAN to iSCSI-B-VLAN.
41. Leave the MAC Address set to None.
42. Click OK.



43. Click OK.
44. Review the table in the Networking page to make sure that all vNICs were created.
45. Click Next.



Configure Storage Options

1. Select the No vHBAs option for the “How would you like to configure SAN connectivity?” field.
2. Click Next.

Configure Zoning Options

1. Set no Zoning options and click Next.

Configure vNIC/HBA Placement


1. From the Select Placement list, select the VM-Host-Infra placement policy.
2. Select vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - a. vNIC-A
 - b. vNIC-B
 - c. iSCSI-vNIC-A
 - d. iSCSI-vNIC-B
3. Review the table to verify that all vNICs were assigned to the policy in the appropriate order.
4. Click Next.

Configure vMedia Policy

Do not configure a vMedia Policy at this time.

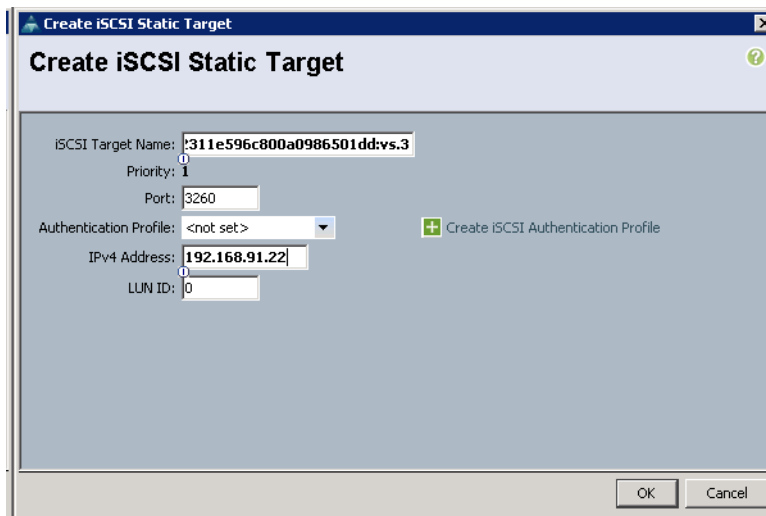
1. Click Next.


Configure Server Boot Order

1. Select Boot-Fabric-A for Boot Policy.
2. From the Boot Order pane, select iSCSI-A-vNIC.
3. Click the Set iSCSI Boot Parameters button.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you independently created one appropriate to your environment.
5. Leave the Initiator Name Assignment dialog box to <not set> to use the single Service Profile Initiator Name defined in the previous steps.
6. Set iSCSI_IP_Pool_A for the Initiator IP address policy.
7. Keep the iSCSI Static Target Interface button selected and click the  button at the bottom right.
8. Log in to the storage cluster management interface and run the following command:

```
iscsi show
-----
Vserver      Target                Target                Status
Name         Name                 Alias                 Admin
-----
Infra_SVM_site_(A/B)  iqn.1992-08.com.netapp:sn.cbc5f0dff5b911e5aaa600a0985b4a74:vs.3
                                      Infra_SVM_site_(A/B)  up
```

9. Note or copy the iSCSI target name for Infra_SVM_site_(A/B).
10. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra_SVM_site_(A/B).
11. Enter the IP address of iSCSI_lif02a in the IPv4 Address field.



12. Click OK to add the iSCSI static target.
13. Keep the iSCSI Static Target Interface option selected and click the  button.
14. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra_SVM_site_(A/B) into the iSCSI Target Name field.
15. Enter the IP address of iscsi_lif01a in the IPv4 Address field.

16. Click OK.

Set iSCSI Boot Parameters
✕

Set iSCSI Boot Parameters ?

Name: **iSCSI-A-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set>

+ Create IQN Suffix Pool

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_A(14/16)

IPv4 Address: **0.0.0.0**
 Subnet Mask: **255.255.255.0**
 Default Gateway: **0.0.0.0**
 Primary DNS: **0.0.0.0**
 Secondary DNS: **0.0.0.0**

+ Create IP Pool


The IP address will be automatically assigned from the selected pool.

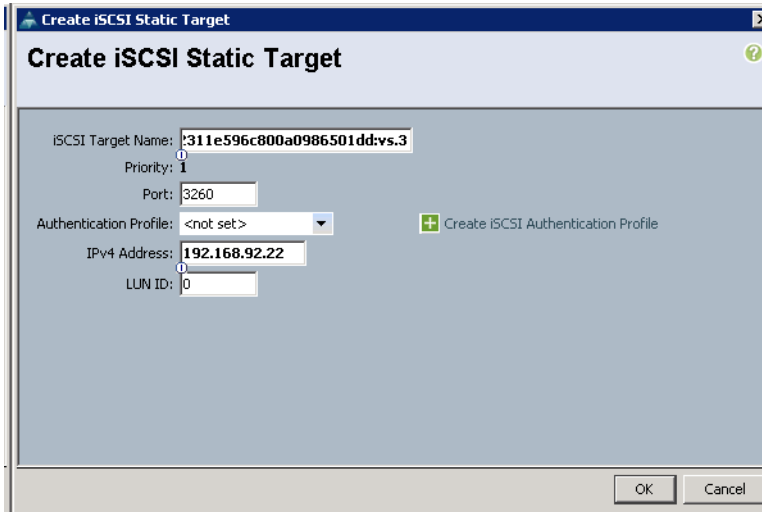
iSCSI Static Target Interface
 iSCSI Auto Target Interface


Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	IPv4 Address	LUN Id
iqn.1992-08.c...	2	3260		192.168.91.21	0
iqn.1992-08.c...	1	3260		192.168.91.22	0

OK
Cancel

17. Click OK.
18. From the Boot Order pane, select iSCSI-vNIC-B.
19. Click the Set iSCSI Boot Parameters button.
20. In the Set iSCSI Boot Parameters dialog box, set the Leave the Initiator Name Assignment to <not set>.
21. In the Set iSCSI Boot Parameters dialog box, set the initiator IP address policy to iSCSI_IP_Pool_B.
22. Keep the iSCSI Static Target Interface option selected and click the  button at the bottom right.
23. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra_SVM_site_(A/B) into the iSCSI Target Name field (same target name as above).
24. Enter the IP address of iscsi_lif02b in the IPv4 Address field.



25. Click OK to add the iSCSI static target.
26. Keep the iSCSI Static Target Interface option selected and click the  button.
27. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra_SVM_site_(A/B) into the iSCSI Target Name field.
28. Enter the IP address of iscsi_lif01b in the IPv4 Address field.
29. Click OK.

Set iSCSI Boot Parameters

Name: **iSCSI-B-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set>

+ Create IQN Suffix Pool

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_B(14/16)

IPv4 Address: **0.0.0.0**
 Subnet Mask: **255.255.255.0**
 Default Gateway: **0.0.0.0**
 Primary DNS: **0.0.0.0**
 Secondary DNS: **0.0.0.0**

+ Create IP Pool

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

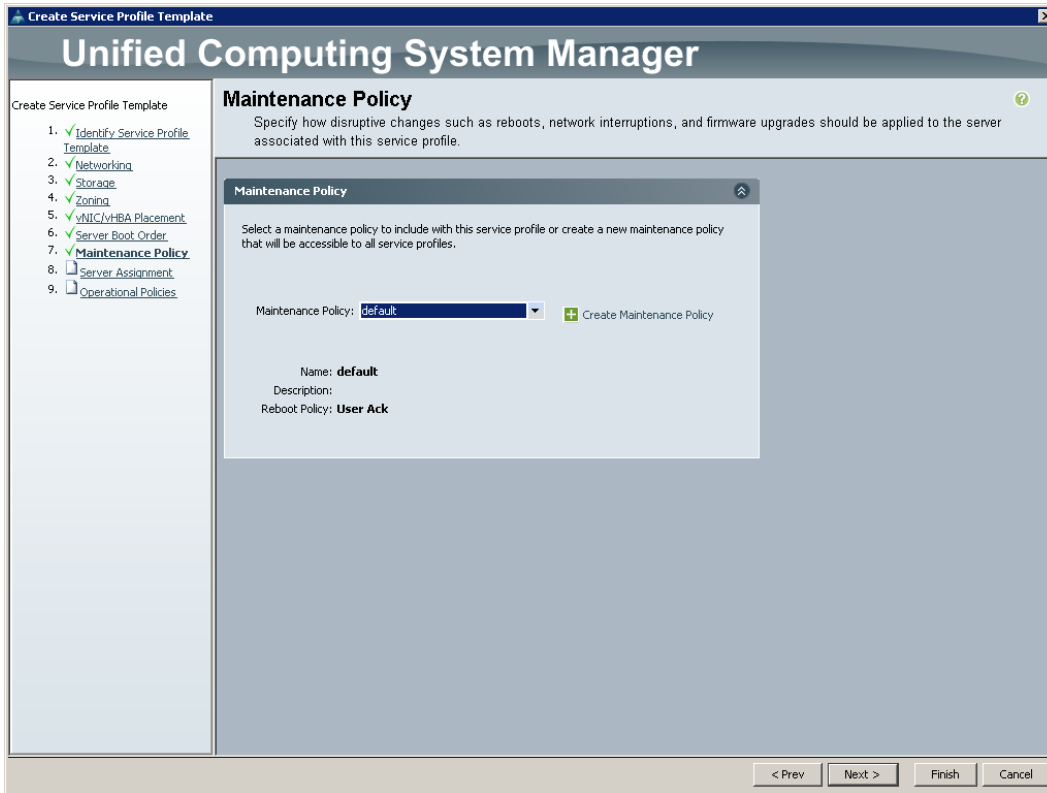
Name	Priority	Port	Authentication Profile	IPv4 Address	LUN Id
iqn.1992-08.c...	2	3260		192.168.92.21	0
iqn.1992-08.c...	1	3260		192.168.92.22	0

OK Cancel

30. Click OK.
31. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
32. Click Next to continue to the next section.

Configure Maintenance Policy

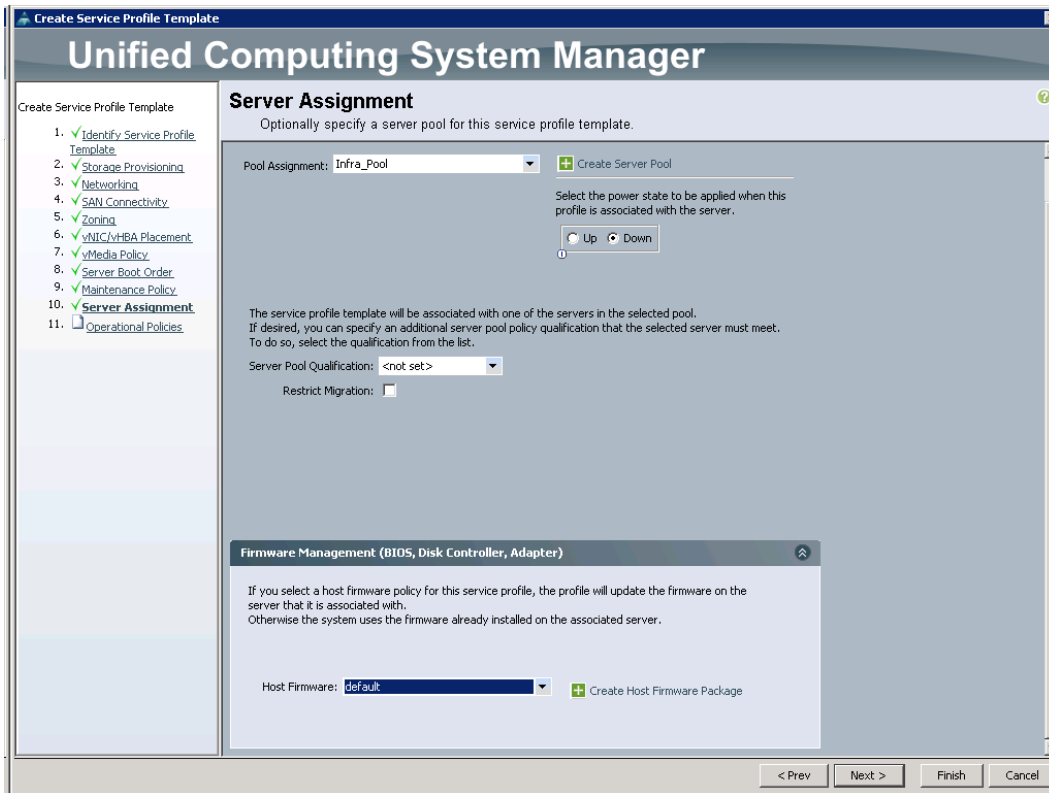
1. Leave the default Maintenance Policy selected.
2. Click Next.



Configure Server Assignment

To configure server assignment, complete the following steps:

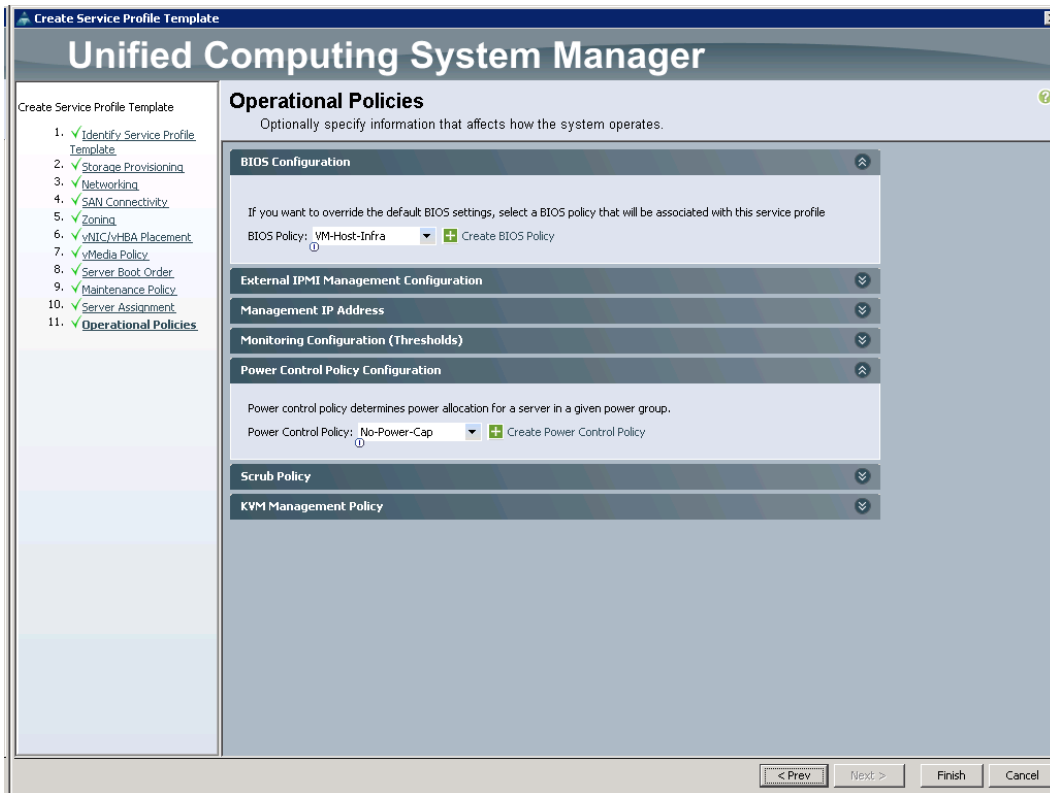
1. From the Pool Assignment list, select Infra_Pool.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management at the bottom of the page and select default from the Host Firmware list.
5. Click Next.



Configure Operational Policies

To configure the operational policies, complete the following steps:

1. From the BIOS Policy list, select VM-Host-Infra.
2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

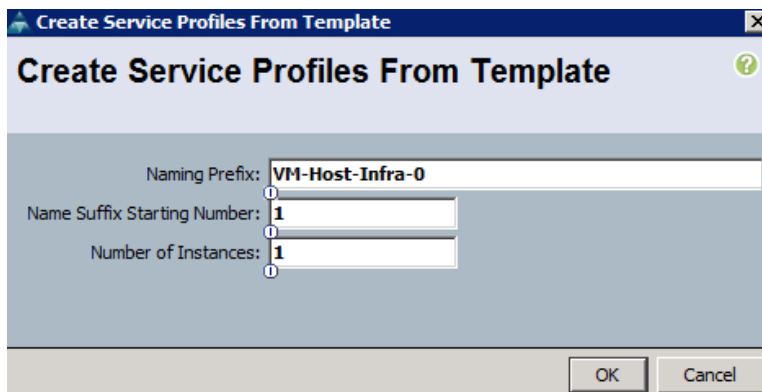


3. Click Finish to create the service profile template.
4. Click OK when prompted for confirmation.

Create Service Profiles

To create service profiles from the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab from the navigation pane.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-Fabric-A.
3. Right-click VM-Host-Infra-Fabric-A and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Enter 1 for Name Suffix Starting Number.
6. Enter 2 for Number of Instances.
7. Click OK to create the service profile.



- Click OK when prompted for confirmation.

Adding Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

Gather Necessary Information

- After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into Table 8 and Table 9.
- To obtain the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface on both sites.

Table 8) iSCSI LIFs for iSCSI IQN.

Vserver	IQN (iSCSI)
Infra_SVM_site_A	<<var_iscsi_name_site_A>>
Infra_SVM_site_B	<<var_iscsi_name_site_B>>

- To obtain the iSCSI vNIC IQN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the iSCSI vNICs tab on the right. The Initiator Name is displayed at the top of the page under the Service Profile Initiator Name. Repeat the same process for the other site.

Table 9) vNIC iSCSI IQNs for fabric A and fabric B.

Cisco UCS Service Profile Name	Initiator: IQN (iSCSI)	Variables
VM-Host-Infra-Fabric-A1		<<var_vm_host_infra_fabric_A1_iqn>>
VM-Host-Infra-Fabric-A2		<<var_vm_host_infra_fabric_A2_iqn>>
VM-Host-Infra-Fabric-B1		<<var_vm_host_infra_fabric_B1_iqn>>
VM-Host-Infra-Fabric-B2		<<var_vm_host_infra_fabric_B2_iqn>>

5.9 Storage Configuration—Boot LUNs and Igroups

Clustered Data ONTAP Boot Storage Setup

Create Igroups

On site A and site B, from the cluster management node SSH connection, run the following commands:

```
igroup create -vserver Infra_SVM_site_(A/B) -igroup VM-Host-Infra-Site-(A/B)1 -protocol iscsi -
ostype vmware -initiator <<var_vm_host_infra_site_(A/B)_1_iqn>>

igroup create -vserver Infra_SVM_site_(A/B) -igroup VM-Host-Infra-Site-(A/B)2 -protocol iscsi -
ostype vmware -initiator <<var_vm_host_infra_site_(A/B)_2_iqn>>
```

```
igroup create -vserver Infra_SVM_site_(A/B) -igroup MGMT-Hosts -protocol iscsi -ostype vmware -
initiator <<var_vm_host_infra_site_(A/B)_1_ign>>, <<var_vm_host_infra_site_(A/B)_2_ign>>
```

Map Boot LUNs to Igroups

On site A and site B, from the storage cluster management SSH connection, run the following commands:

```
lun map -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B) -lun VM-Host-Infra-Site-(A/B)-
01 -igroup VM-Host-Infra-Site-(A/B)1 -lun-id 0

lun map -vserver Infra_SVM_site_(A/B) -volume esxi_boot_site_(A/B) -lun VM-Host-Infra-Site-(A/B)-
02 -igroup VM-Host-Infra-Site-(A/B)2 -lun-id 0
```

Map Site A Heartbeat LUNs to Igroups

Run the following commands on site A to map site A heartbeat LUNs to igroups:

```
lun map -vserver Infra_SVM_site_A -volume site_A_heartbeat -lun site_A_heartbeat -igroup VM-Host-
Infra-Site-A1 -lun-id 1

lun map -vserver Infra_SVM_site_A -volume site_A_heartbeat -lun site_A_heartbeat -igroup VM-Host-
Infra-Site-A2 -lun-id 1
```

Map Site B Heartbeat LUNs to Igroups

Run the following commands on site B to map site B heartbeat LUNs to igroups:

```
lun map -vserver Infra_SVM_site_B -volume site_B_heartbeat -lun site_B_heartbeat -igroup VM-Host-
Infra-Site-B1 -lun-id 1

lun map -vserver Infra_SVM_site_B -volume site_B_heartbeat -lun site_B_heartbeat -igroup VM-Host-
Infra-Site-B2 -lun-id 1
```

5.10 VMware vSphere 6.0 Setup

This section provides detailed instructions for installing VMware ESXi 6.0 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned at each site.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the in-built KVM console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot LUNs.

Run these steps for the Cisco UCS domain in site A and site B.

Download Cisco Custom Image for ESXi 6.0

1. Go to my.vmware.com.
2. Type your e-mail or customer number and the password and then click Log in.
3. Download the [Custom Image for ESXi 6.0.0](#).
4. Save the image to your destination folder.

Note: The Cisco custom image for ESXi 6.0 includes updates for the fNIC and eNIC drivers. The versions that are part of this image are eNIC: 2.1.2.71 and fNIC: 1.6.0.17a.

Log in to Cisco UCS 6200 Fabric Interconnect

The IP KVM enables the administrator to begin the installation of the OS through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. To download the Cisco UCS Manager software, click the Launch UCS Manager link.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click the Servers tab.
7. Select Servers > Service Profiles > root > VM-Host-Infra-Site-<A/B>-01.
8. Right-click VM-Host-Infra-Site-<A/B>-01 and select KVM Console.
9. If prompted to accept an Unencrypted KVM session, accept it.
10. Select Servers > Service Profiles > root > VM-Host-Infra-Site-<A/B>-02.
11. Right-click VM-Host-Infra-Site-<A/B>-02 and select KVM Console.
12. If prompted to accept an Unencrypted KVM session, accept it.

Set Up VMware ESXi Installation

Note: Skip this step if using vMedia policies. The ISO file will already be connected to KVM.

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM page, click Virtual Media.
2. Click Activate Virtual Devices.
3. If prompted to accept an unencrypted virtual media session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.
8. Boot the server by selecting Boot Server and clicking OK. Click OK again.

Install ESXi

To install VMware ESXi to the iSCSI-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
4. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, click the Virtual Media tab and clear the ✓ mark next to the ESXi installation media. Click Yes.

Note: The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

9. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for each VMware host, complete the following steps on the host.

1. After the server has finished booting, press F2 to customize the system.
2. Log in as root, enter the corresponding password, and press Enter to log in.
3. Select the Configure Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the in-band-management VLAN ID and press Enter.
6. Select the Network Adapters option and select the VMNICs that are used for vNIC-A and vNIC-B. Press Enter.

Note: From the UCS Manager, select the service profile for the corresponding server and click the Network tab. The network adapters will be listed with their MAC address.

7. From the Configure Management Network menu, select IPv4 Configuration and press Enter.
8. Select the Set Static IPv4 Address and Network Configuration option using the space bar.
9. Enter the IP address for managing the ESXi host.
10. Enter the subnet mask for the ESXi host.
11. Enter the default gateway for the ESXi host.
12. Press Enter to accept the changes to the IP configuration.
13. Select the IPv6 Configuration option and press Enter.
14. Using the space bar, select Disable IPv6 (restart required) and press Enter.
15. Select the DNS Configuration option and press Enter.

Note: Because the IP address is assigned manually, the DNS information must also be entered manually.

16. Enter the IP address of the primary DNS server.
17. Optional: Enter the IP address of the secondary DNS server.
18. Enter the fully qualified domain name (FQDN) for the first ESXi host.
19. Press Enter to accept the changes to the DNS configuration.
20. Press Esc to exit the Configure Management Network submenu.
21. Press Y to confirm the changes and return to the main menu.
22. The ESXi host reboots. After reboot, press F2 and log back in as root.
23. Select Test Management Network to verify that the management network is set up correctly. Press Enter.
24. Press Enter to run the test.
25. Press Enter to exit the window.
26. Press Esc to log out of the VMware console.

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-Site-A-01 management IP address.
2. Download and install the vSphere Client.

Note: This application is downloaded from the VMware website and Internet access is required on the management workstation.

Download VMware vSphere CLI 6.0

Note: Install VMware vSphere CLI 6.0 on the management workstation.

1. Click the following link: [VMware vSphere CLI 6.0](#).
2. Select your OS and click Download.
3. Save it to your destination folder.
4. Run the `VMware-vSphere-CLI-6.0.0-2503617.exe`.
5. Click Next.
6. Accept the terms for the license and click Next.
7. Click Next on the Destination Folder screen.
8. Click Install.
9. Click Finish.

Log in to VMware ESXi Hosts Using VMware vSphere Client

To log in to the ESXi host using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded vSphere Client and enter the IP address ESXi to which you want to connect.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

Set Up VMkernel Ports and Virtual Switches

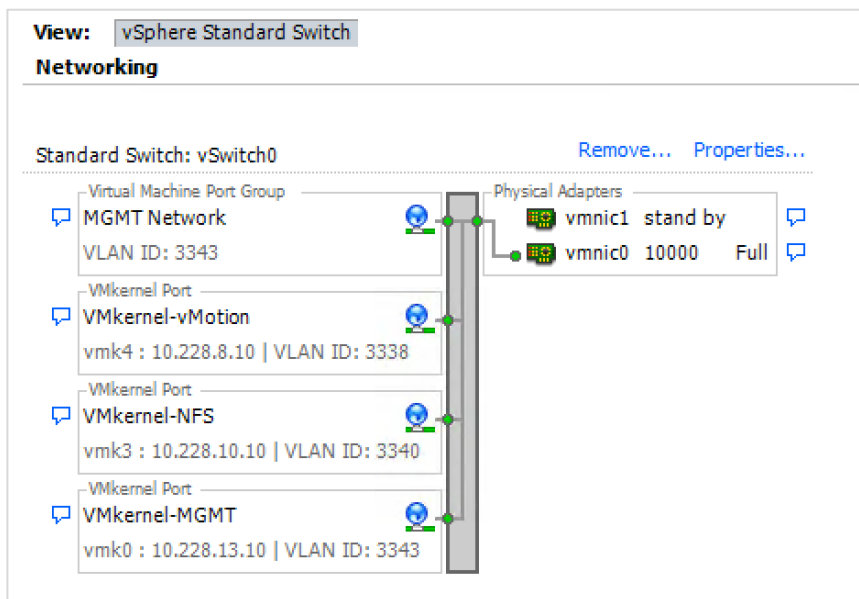
To set up the VMkernel ports and the virtual switches on each ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of vSwitch0, click Properties.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to 9000.
7. Click OK.
8. Select the Management Network configuration and click Edit.
9. Change the network label to VMkernel-MGMT and select the Management Traffic checkbox.
10. Click OK to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to MGMT Network and enter `<<var_ib-mgmt_vlan_id>>` in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.
14. Click Close.
15. On the right side of iScsiBootvSwitch, click Properties.
16. Select the vSwitch configuration and click Edit.

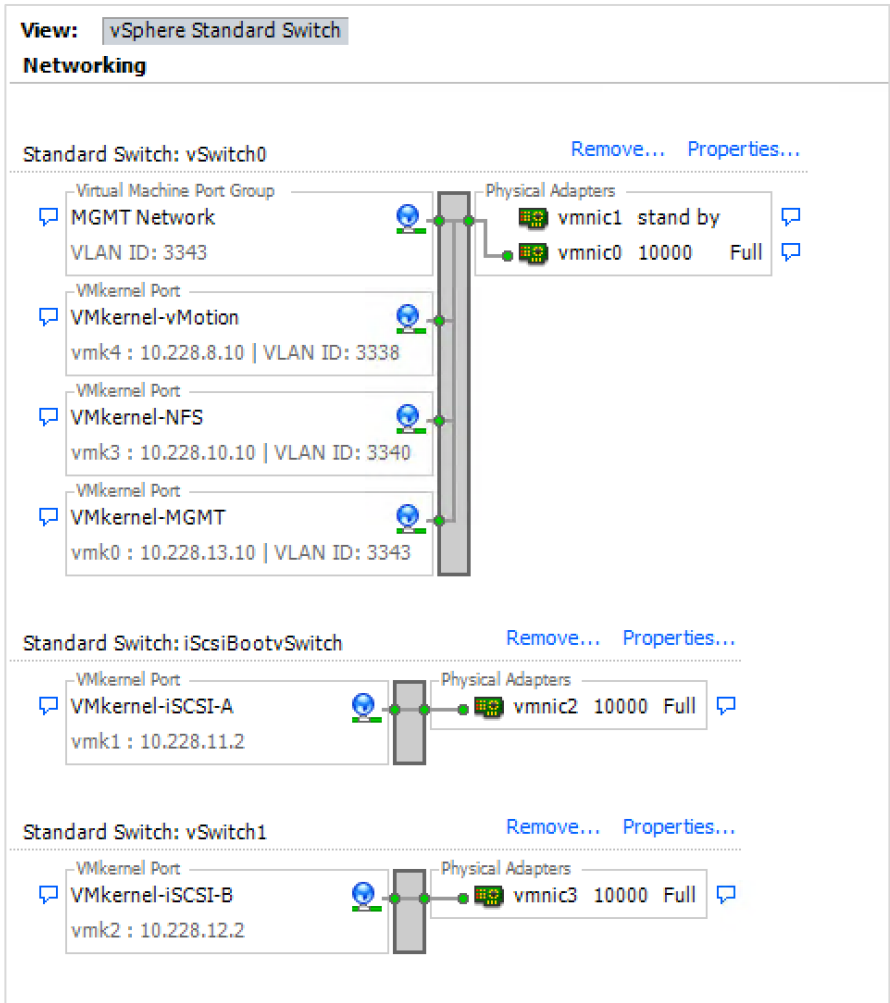
17. Change the MTU to 9000.
 18. Click OK.
 19. Select iScsiBootPG and click Edit.
 20. Change the Network Label to VMkernel-iSCSI-A.
 21. Change the MTU to 9000.
 22. Click OK.
 23. Click Close.
 24. In the vSphere Standard Switch view, click Add Networking.
 25. Select VMkernel and click Next.
 26. Select Create a vSphere standard switch to create a new vSphere standard switch.
 27. Select the checkboxes for the network adapter VMNIC3.
 28. Click Next.
 29. Change the network label to VMkernel-iSCSI-B.
 30. Click Next.
 31. Enter the IP address and the subnet mask for the iSCSI VLAN B interface for VM-Host-Infra-01.
- Note:** To obtain the iSCSI IP address information, log in to Cisco UCS Manager. In the Servers tab, select the corresponding service profiles. In the right pane, click Boot Order and select iSCSI-B-vNIC. Click Set iSCSI Boot Parameters. The IP address should appear as the initiator IP address.
32. Click Next.
 33. Click Finish.
 34. On the right side of vSwitch1, click Properties.
 35. Select the vSwitch configuration and click Edit.
 36. Change the MTU to 9000.
 37. Click OK.
 38. Select VMkernel-iSCSI-B and click Edit.
 39. Change the MTU to 9000.
 40. Click OK.
 41. Click Close.
 42. On the right side of vSwitch0, click Properties.
 43. Click Add.
 44. Select VMkernel, click Next.
 45. Change the network label to VMkernel-NFS and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.
 46. Click Next.
 47. Enter the IP address <<var_nfs_vlan_ip_host>> and the subnet mask <<var_nfs_vlan_ip_mask_host>> for the NFS VLAN interface for VM-Host-Infra-01.
 48. To continue with the NFS VMkernel creation, click Next.
 49. To finalize the creation of the NFS VMkernel interface, click Finish.
 50. Select the VMkernel-NFS configuration and click Edit.
 51. Change the MTU to 9000.
 52. Click OK to finalize the edits for the VMkernel-NFS network.
 53. Click Add.

54. Select VMkernel and click Next.
55. Change the network label to VMkernel-vMotion and enter <<var_vmotion_vlan_id>> in the VLAN ID (Optional) field.
56. Click Next.
57. Enter the IP address <<var_vmotion_vlan_ip_host>> and the subnet mask <<var_vmotion_vlan_ip_mask_host>> for the vMotion VLAN interface for VM-Host-Infra-01.
58. To continue with the vMotion VMkernel creation, click Next.
59. To finalize the creation of the vMotion VMkernel interface, click Finish.
60. Select the VMkernel-vMotion configuration and click Edit.
61. Change the MTU to 9000.
62. Click OK to finalize the edits for the VMkernel-vMotion network.

The properties for vSwitch0 should be similar to the following example:




To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:



Set Up iSCSI Multipathing

To set up four iSCSI paths between each storage cluster and the ESXi host, complete the following steps on each ESXi host:

1. From the vSphere Client, click Storage Adapters in the Hardware pane.
2. Select the iSCSI software adapter and click Properties.
3. Select the Dynamic Discovery tab and click Add.
4. Enter the IP address of the first LIF in the storage cluster in site A and click OK.
5. Repeat step 4, entering the IP addresses of the remaining LIFs in the storage cluster in site A.
6. Repeat step 4, entering the IP addresses of the LIFs in the storage cluster in site B.



7. Click Close and then click Yes to rescan the HBA.
You should now see eight connected paths in the Details pane.

Install VMware Drivers for the Cisco Virtual Interface Card (Optional)

To update the VMware eNIC and fNIC drivers, complete the following steps.

Download and extract the VMware VIC drivers to the management workstation.

For example:

- fNIC driver version 1.6.0.17a
- eNIC driver version 2.2.2.71

Note: VMware vSphere 5.5 drivers are supported to work with vSphere 6.0.

To install VMware VIC drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each vSphere Client, select the host from the inventory.
2. Click the Summary tab to view the environment summary.
3. From Resources > Storage, right-click datastore1 and select Browse Datastore.
4. Click the fourth button and select Upload File.
5. Navigate to the saved location for the downloaded VIC drivers and select `fnic_driver_1.6.0.17a_ESX55-offline_bundle-2774889.zip`.
6. Click Open and Yes to upload the file to datastore1.
7. Click the fourth button and select Upload File.

8. Navigate to the saved location for the downloaded VIC drivers and select enic-2.1.2.71_esx55-offline_bundle-2739120.zip.
9. Click Open and Yes to upload the file to datastore1.
10. Make sure the files have been uploaded to both ESXi hosts.
11. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
12. At the command prompt, run the following commands to account for each host:

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> --thumbprint <host_thumbprint>
software vib update -d /vmfs/volumes/datastore1/fnic_driver_1.6.0.17a_ESX55-offline_bundle-
2774889.zip
```

Note: To get the host thumbprint, type the following command without the `--thumbprint` option and then copy and paste the thumbprint into the command.

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> --thumbprint <host_thumbprint>
software vib update -d /vmfs/volumes/datastore1/fnic_driver_1.6.0.17a_ESX55-offline_bundle-
2774889.zip

esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> --thumbprint <host_thumbprint>
software vib update -d /vmfs/volumes/datastore1/enic-2.1.2.71_esx55-offline_bundle-2739120.zip

esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> --thumbprint <host_thumbprint>
software vib update -d /vmfs/volumes/datastore1/enic-2.1.2.71_esx55-offline_bundle-2739120.zip
```

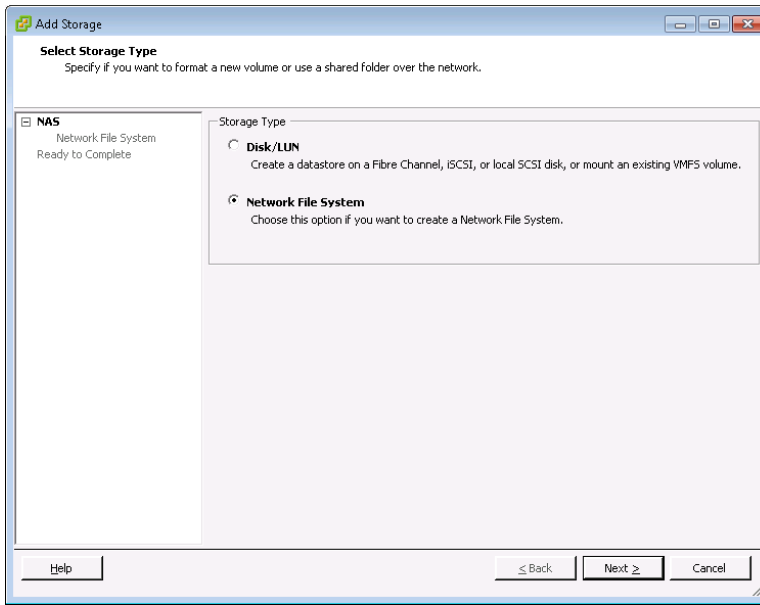
13. In the vSphere Client for each host, right-click the host and select Reboot.
14. Click Yes and OK to reboot the host.
15. Log back into each host with vSphere Client.

Mount Required Datastores

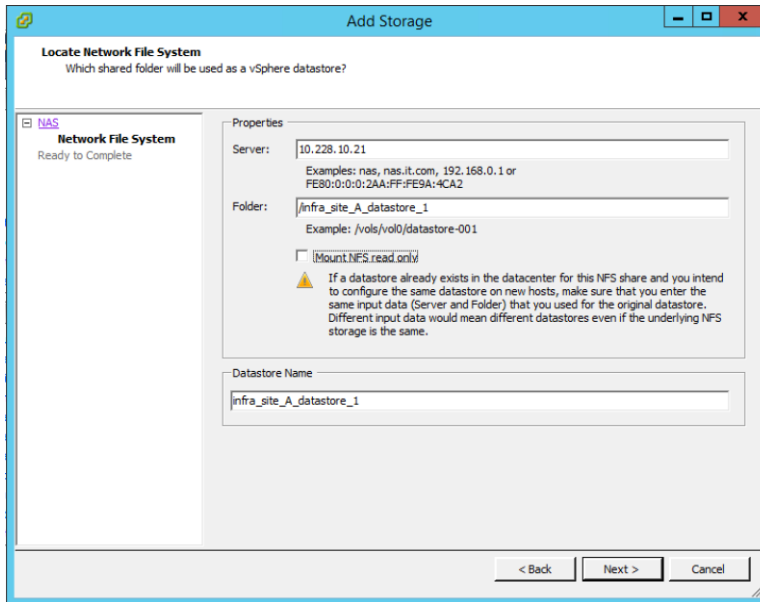
To mount the required datastores, complete the following steps on each ESXi host.

Note: Each ESXi server will be mapped to four datastores meant for VM data: one Infra_Datastore per site and one SWAP datastore per site.

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the Configuration tab.
3. Click Storage in the Hardware pane.
4. From the Datastores area, click Add Storage to open the Add Storage wizard.



5. Select Network File System and click Next.
6. The wizard prompts for the location of the NFS export. Enter the IP address for `nfs_lif_infra_datastore_1`.
7. Enter `/infra_site_A_datastore_1` as the path for the NFS export.
8. Verify that the Mount NFS read-only checkbox is not selected.
9. Enter `infra_site_A_datastore_1` as the datastore name.



10. To continue with the NFS datastore creation, click Next.
11. To finalize the creation of the NFS datastore, click Finish.
12. Repeat the previous steps for the following NFS datastores. Use the correct network interface IP address for each datastore.
 - `/infra_site_A_swap`

- /infra_site_B_datastore_1
- /infra_site_B_swap

Configure NTP on ESXi Hosts

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

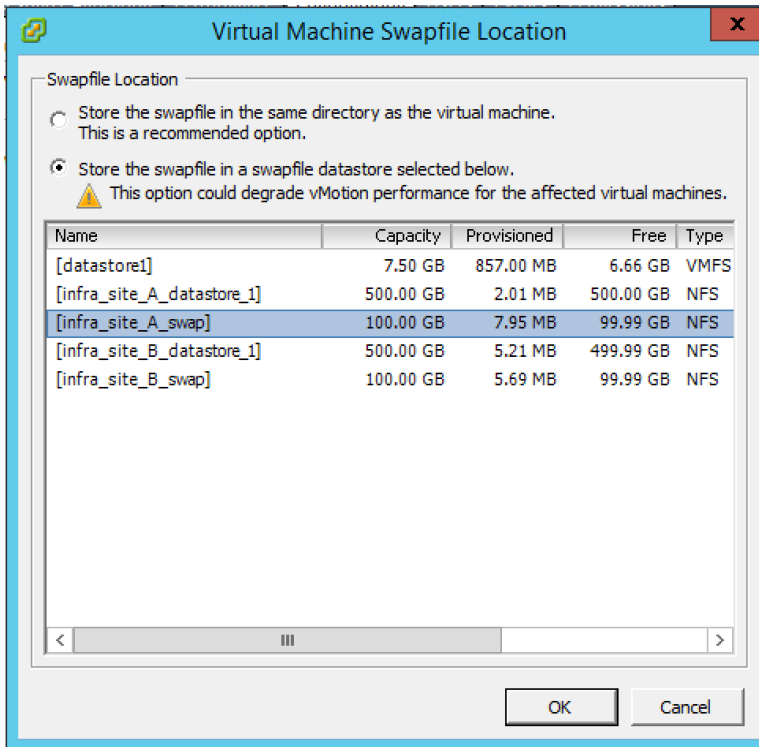
1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Time Configuration in the Software pane.
4. Click Properties at the upper-right side of the page.
5. At the bottom of the Time Configuration dialog box, click Options.
6. In the NTP Daemon (NTPD) Options dialog box, complete the following steps:
7. Click General in the left pane and select Start and stop with host.
8. Click NTP Settings in the left pane and click Add.
9. In the Add NTP Server dialog box, enter the NTP distribution IP address of the first Cisco Nexus 9K on the same site as the ESXi host.
10. Click Add.
11. In the Add NTP Server dialog box, enter the NTP distribution IP address of the second Cisco Nexus 9K on the same site as the ESXi host.
12. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes to the checkbox and click OK.
13. In the Time Configuration dialog box, complete the following steps:
 - a. Select the NTP Client Enabled checkbox and click OK.
 - b. Verify that the clock is now set to approximately the correct time.

Note: The NTP server time might vary slightly from the host time.

Move VM Swap File Location

To move the VM swap file location, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the Configuration tab.
3. Click Virtual Machine Swapfile Location in the Software pane.
4. Click Edit at the upper-right side of the page.
5. Select Store the Swapfile in a Swapfile Datastore Selected Below.
6. Select the <infra_swap> datastore of the local site in which to house the swap files.



7. Click OK to finalize moving the swap file location.

VMware vCenter Server Appliance 6.0

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter Server appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Note: The vCenter Server can be installed in either site. The ESXi hosts of both the sites will be added to the same vCenter cluster.

Install Client Integration Plug-In

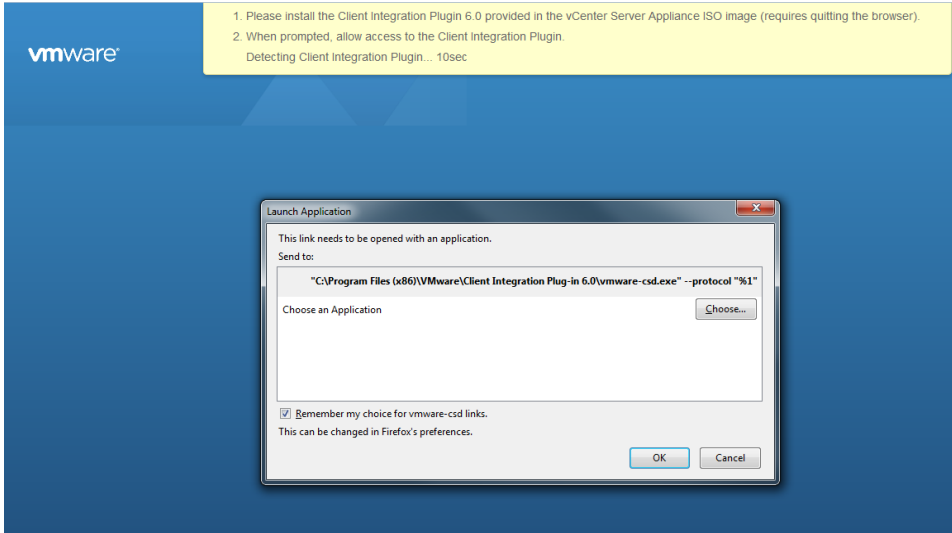
To install the Client Integration Plug-In, complete the following steps:

1. Download the .iso installer for the vCenter Server appliance and Client Integration Plug-In.
2. Mount the ISO image to the Windows virtual machine or physical server on which you want to install the Client Integration Plug-In to deploy the vCenter Server appliance.
3. In the software installer directory, navigate to the VCSA directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe. The Client Integration Plug-In installation wizard appears.
4. On the Welcome page, click Next.
5. Read and accept the terms in the End-User License Agreement and click Next.
6. Click Next.
7. Click Install.
8. Click Finish to exit the wizard.


Build VMware vCenter Server Appliance

To build the VMware vCenter virtual machine, complete the following steps:

1. In the software installer directory, double-click vcsa-setup.html.
2. Allow the plug-in to run on the browser when prompted.

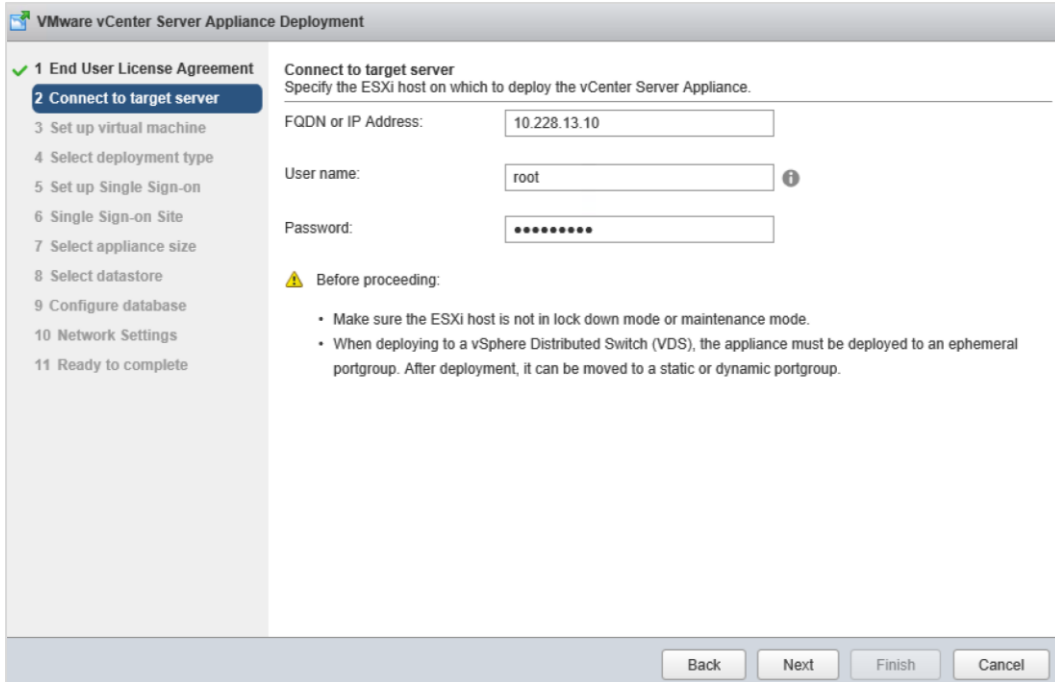


3. On the Home page, click Install to start the vCenter Server appliance deployment wizard.

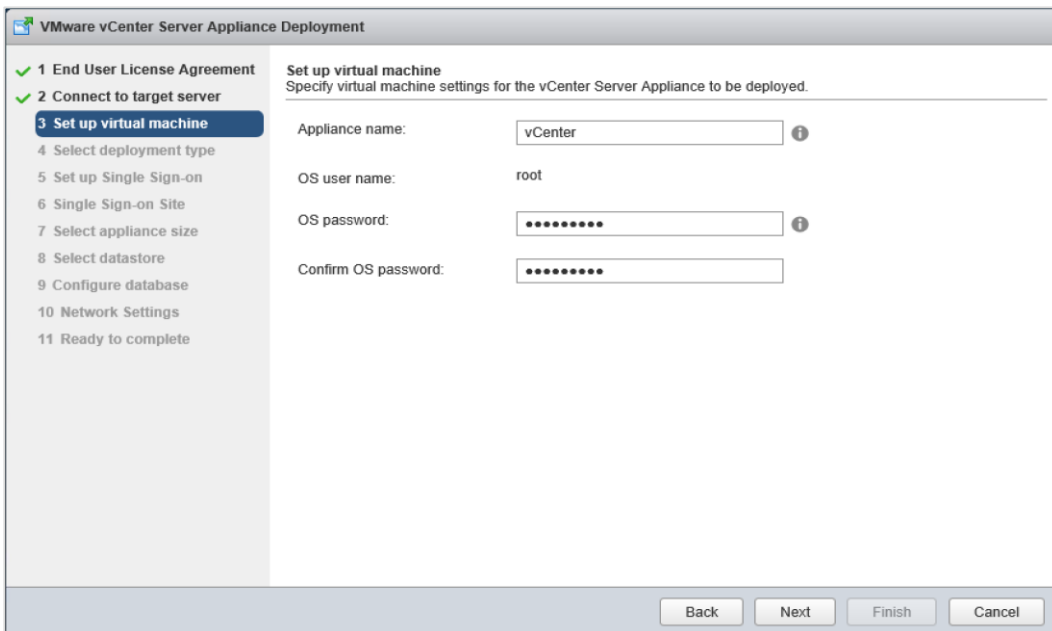


4. Read and accept the EULA and click Next.
5. In the Connect to Target Server dialog box, enter the ESXi host name, user name, and password. Click Next.

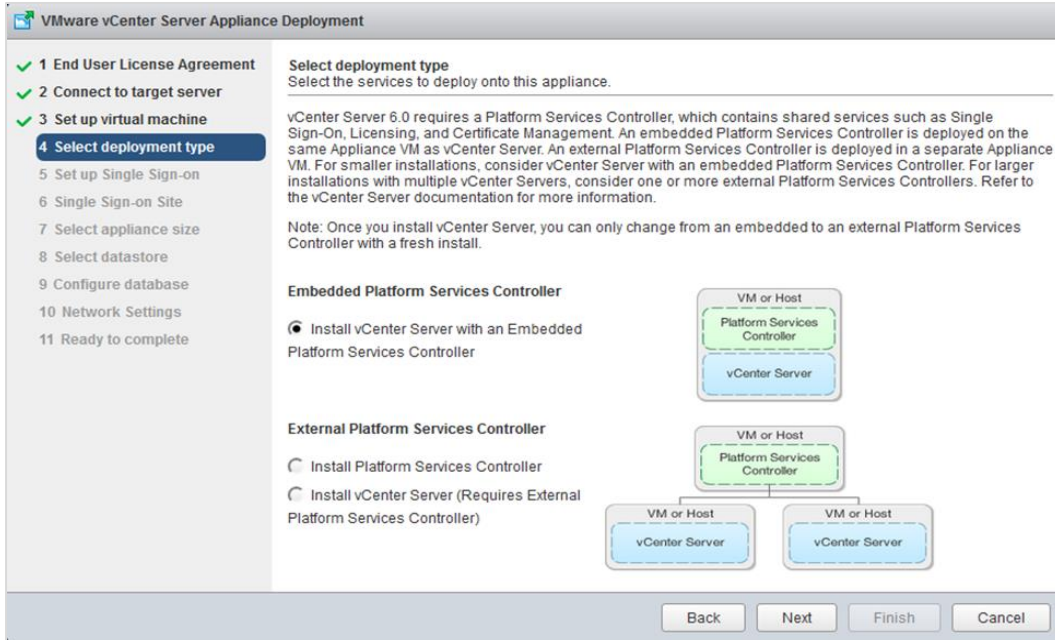
Note: The vCenter Server can reside on either site; choose any ESXi host from both sites.



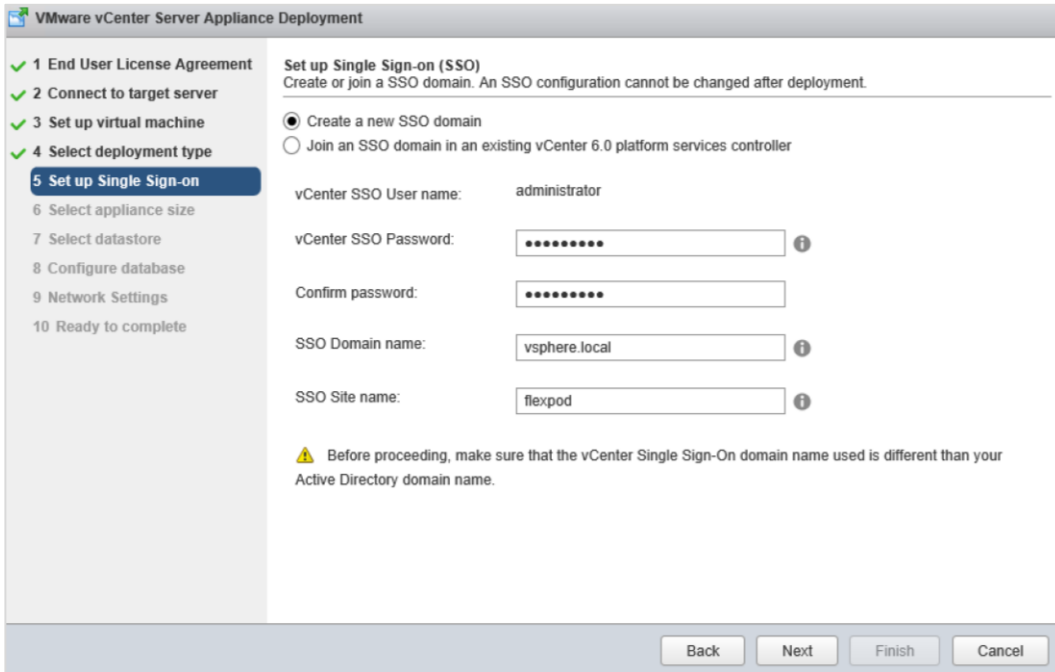
6. Accept the Certificate Warning by selecting Yes.
7. In the Set Up Virtual Machine dialog box, enter the appliance name and password. Click Next.



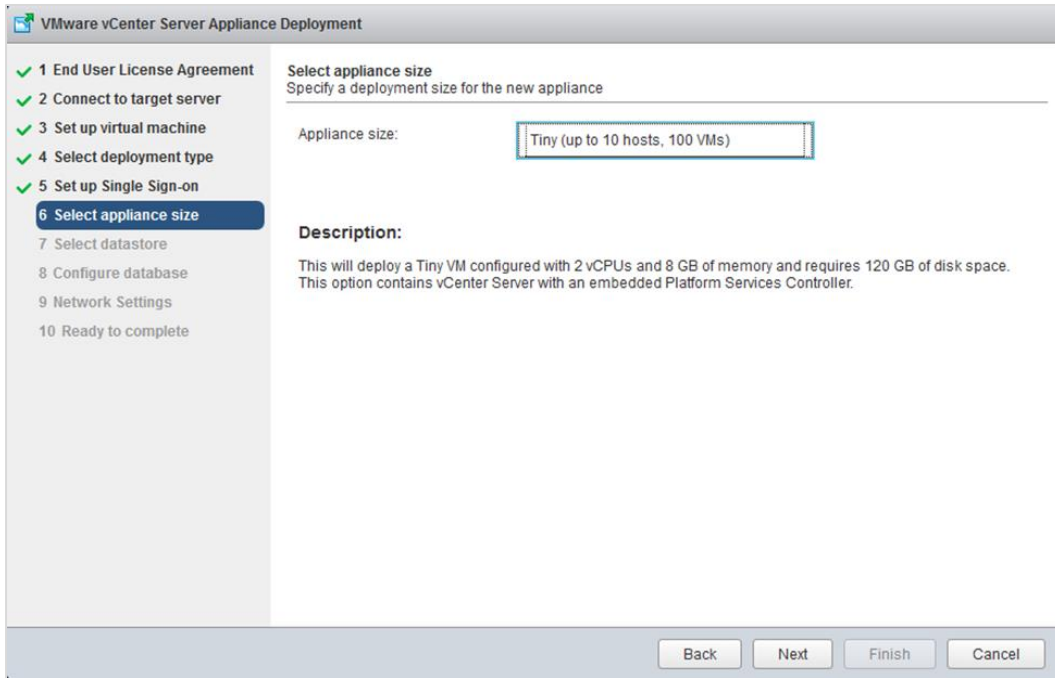
8. In the Select Deployment Type dialog box, select Install vCenter Server with an embedded Platform Services Controller. Click Next.



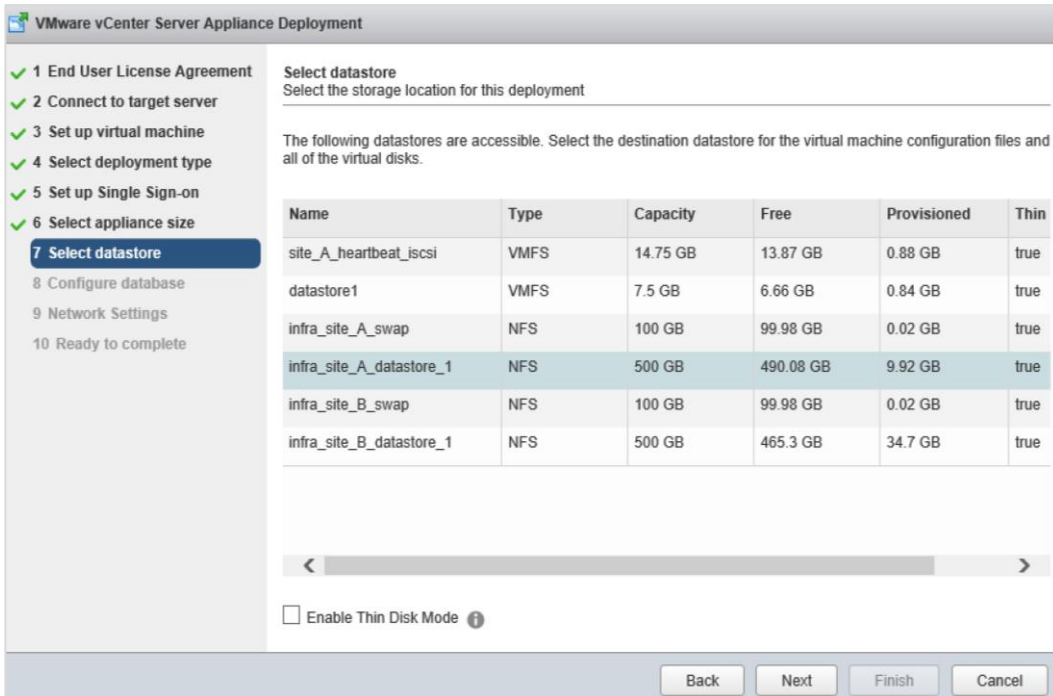
- In the Set Up Single Sign-On dialog box, select Create a New SSO Domain. Enter the SSO password, the domain name, and the site name and click Next.



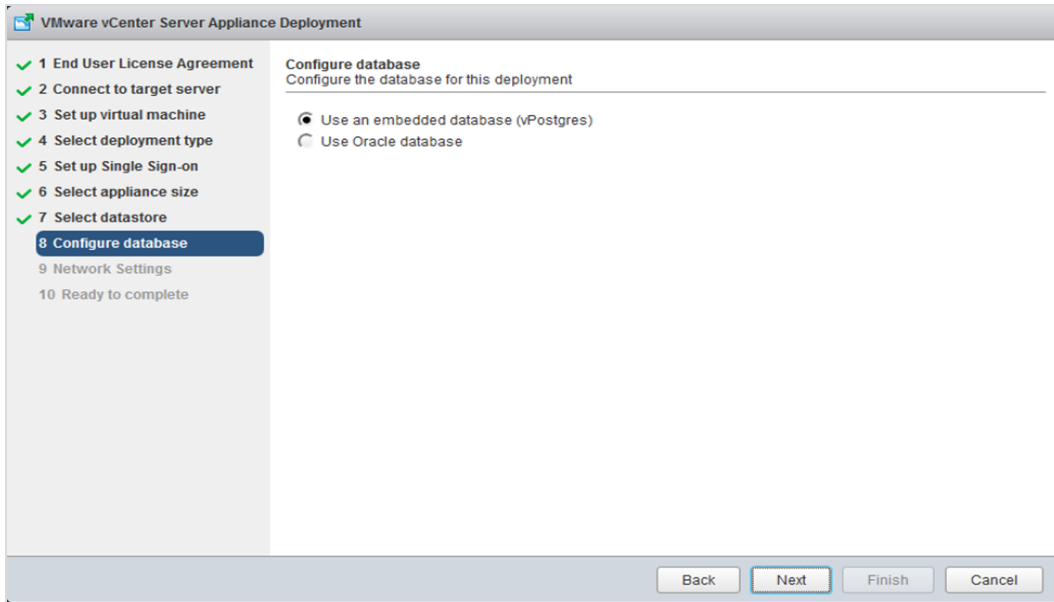
- Select the appliance size; for example, Tiny (up to 10 hosts, 100 VMs). Click Next.



11. In the Select Datastore dialog box, select `infra_site_(A/B)datastore_1`, which resides on the local site. Click Next.



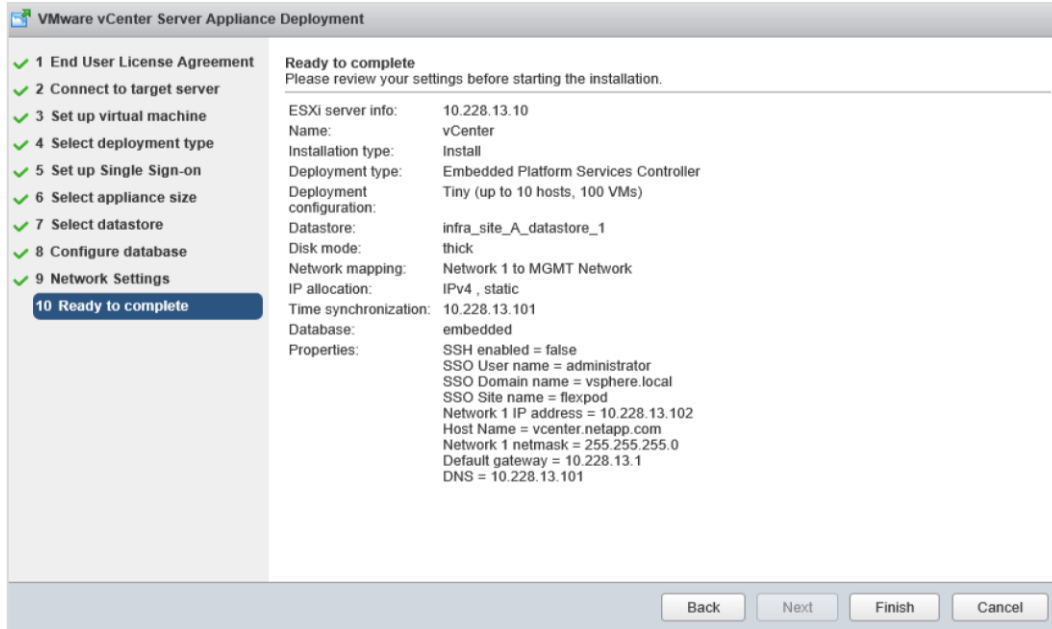
12. In the Configure Database dialog box, select Embedded Database. Click Next.



13. In the Network Settings dialog box, configure the following settings and click Next:

- Choose a network: MGMT-Network
- IP address family: IPV4
- Network type: static
- Network address: <<var_vcenter_ip>>
- System name: <<var_vcenter_fqdn>>
- Subnet mask: <<var_vcenter_subnet_mask>>
- Network gateway: <<var_vcenter_gateway>>
- Network DNS servers: <<var_dns_server>>
- Configure time sync: Use NTP servers
- (Optional): Enable SSH

14. Review the configuration and click Finish.



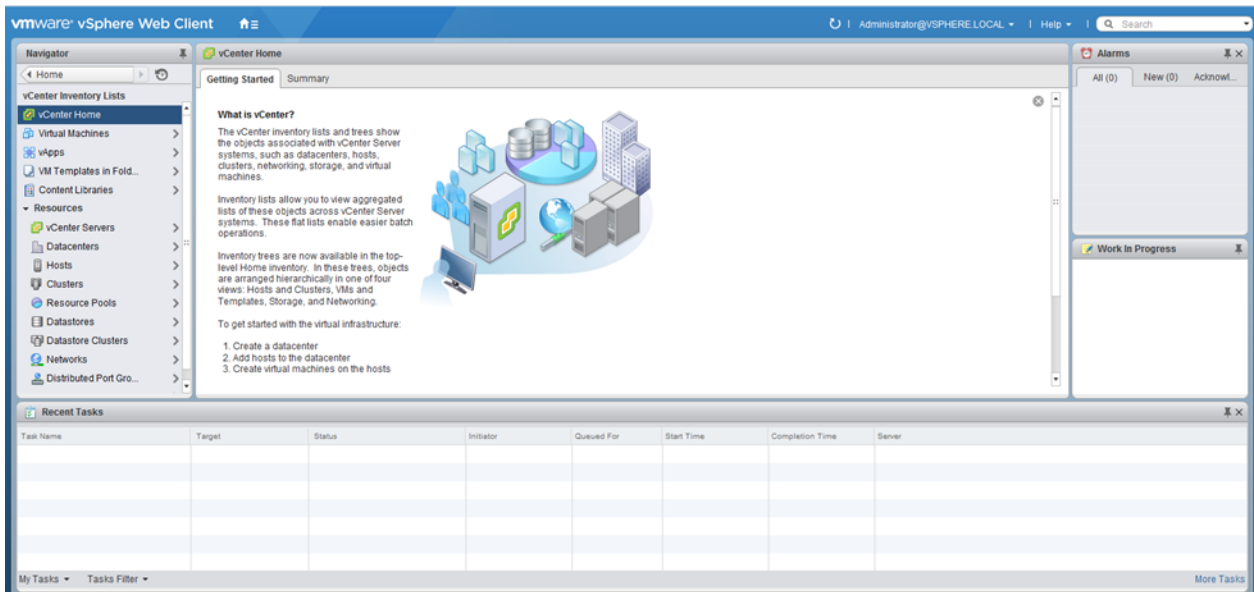
15. The vCenter appliance installation takes a few minutes to complete.

Set Up VMware vCenter Server

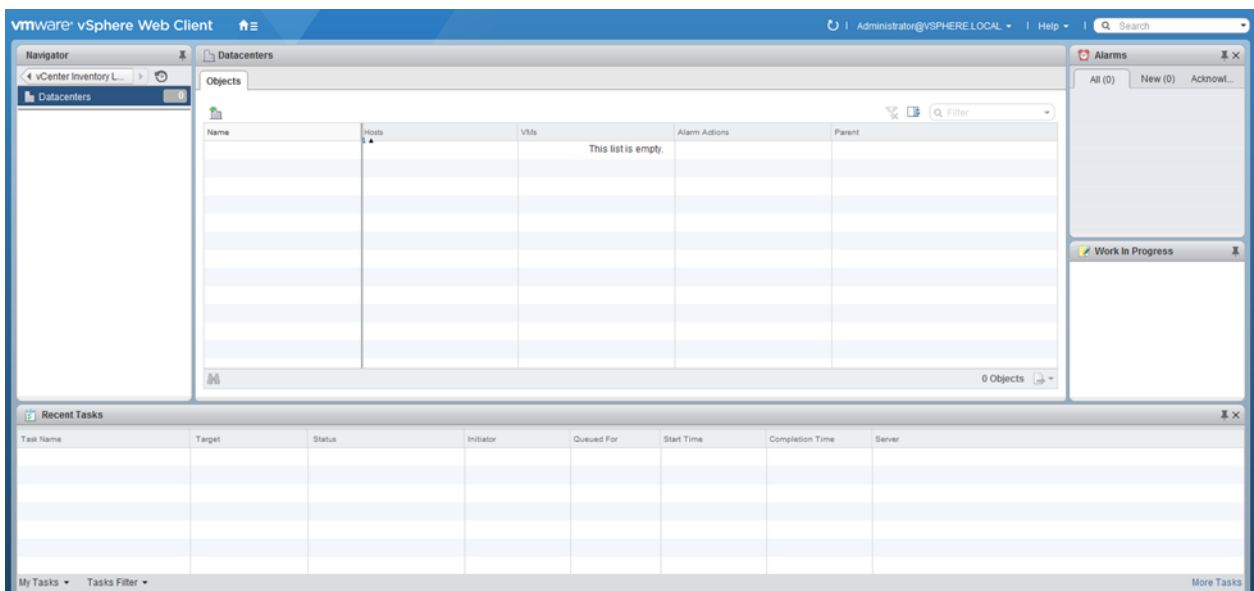
1. Using a web browser, navigate to https://<<var_vcenter_ip>.




2. Click Log in to vSphere Web Client.
3. Click OK in the Launch Application page.
4. Log in using the single-sign-on user name and password created during the vCenter installation.
5. Navigate to vCenter Inventory Lists on the left pane.

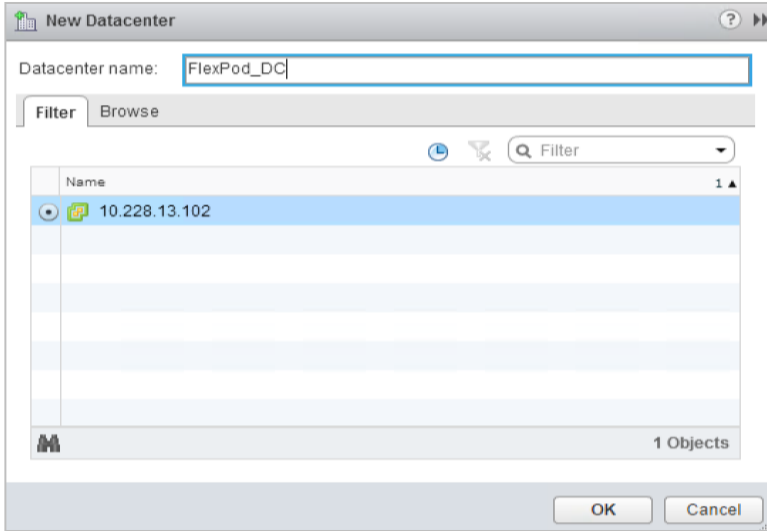


6. Under Resources, click Datacenters in the left plane.

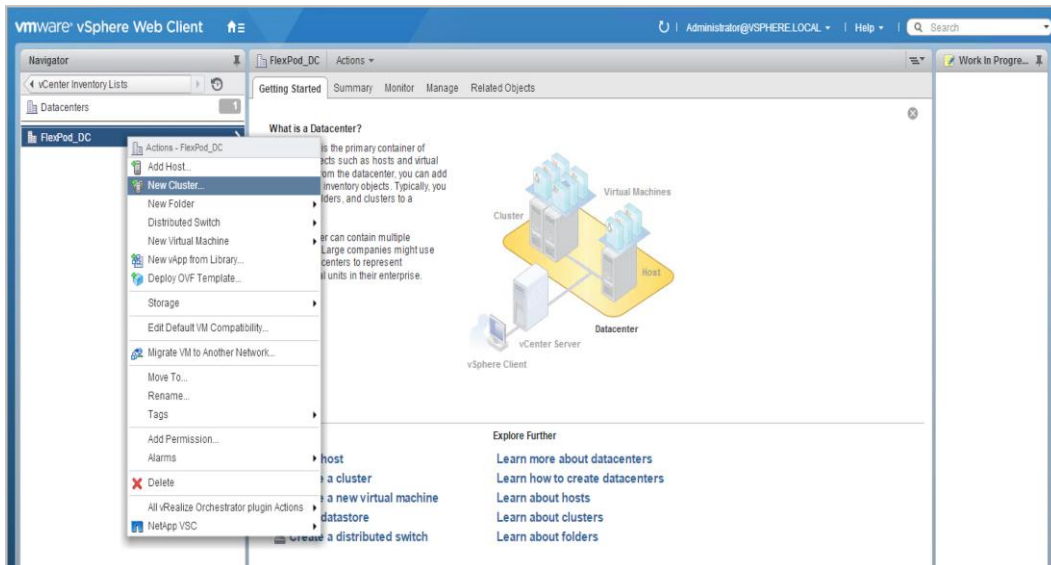


7. To create a data center, click the  icon in the center pane.

8. Type `FlexPod_DC` in the Datacenter Name field. Select the vCenter Name/IP option and click OK.



9. Right-click the data center FlexPod_DC from the list in the center pane. Click New Cluster.



10. Name the cluster FlexPod_Management.

11. Select the checkbox to enable DRS. Retain the default values for this setting.

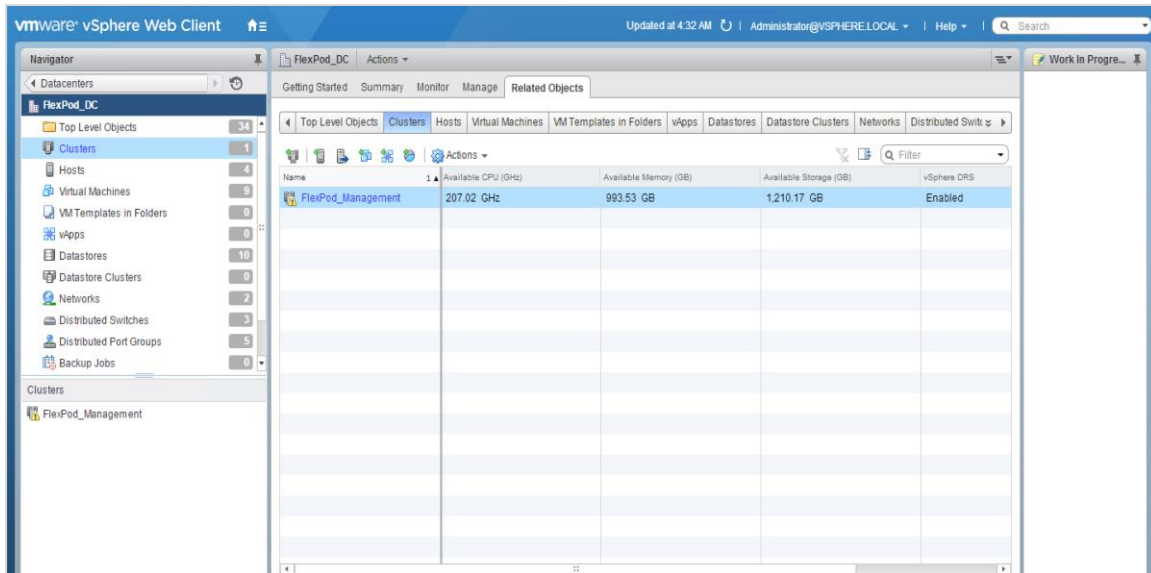
12. Select the checkbox to enable vSphere HA. Retain the default values for this setting.

New Cluster

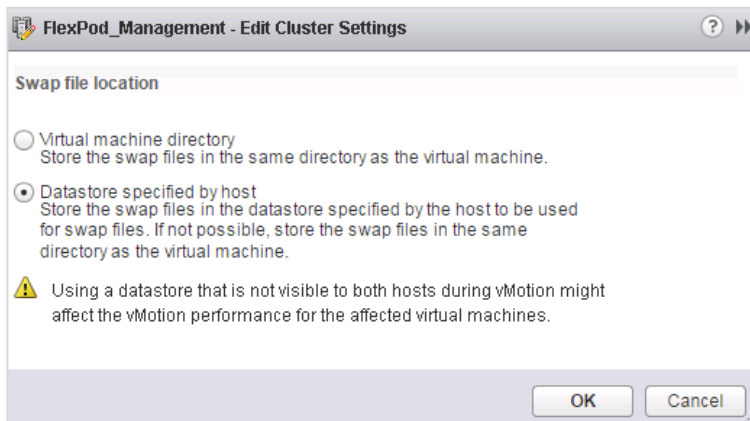
Name	FlexPod_Management
Location	FlexPod_DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	
Admission Control Status	Admission control will prevent powering on VMs that violate availability constraints <input checked="" type="checkbox"/> Enable admission control
Policy	Specify the type of the policy that admission control should enforce. <input checked="" type="radio"/> Host failures cluster tolerates: 1 <input type="radio"/> Percentage of cluster resources reserved as failover spare capacity: Reserved failover CPU capacity: 25 % CPU Reserved failover Memory capacity: 25 % Memory
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low ——— High
EVC	Disable
Virtual SAN	<input type="checkbox"/> Turn ON

OK Cancel

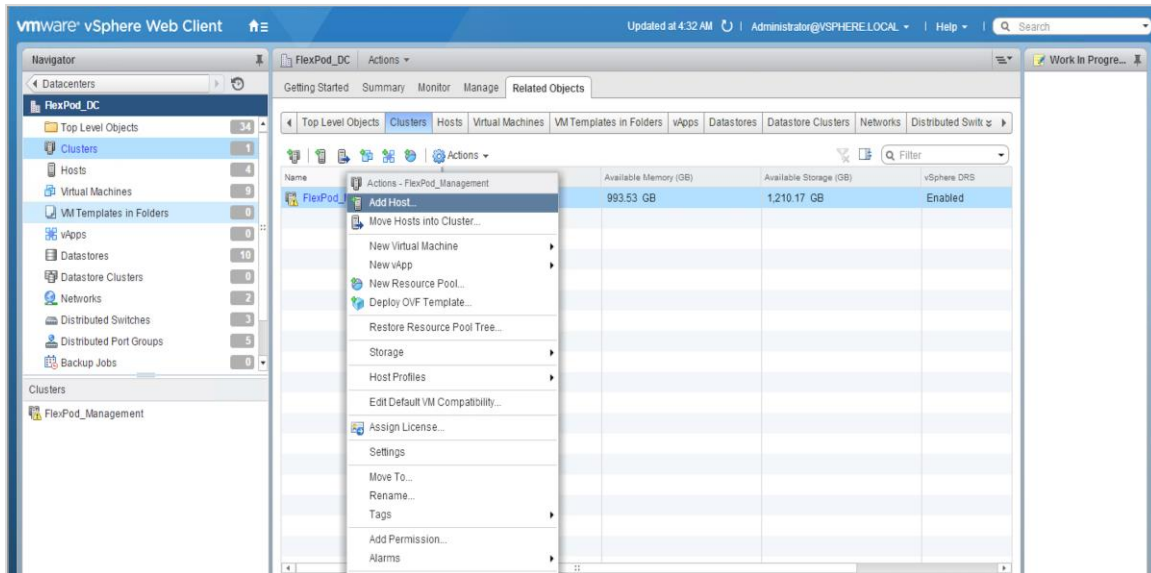
13. Click OK to create the new cluster.
14. On the left pane, double-click FlexPod_DC.
15. Click Clusters.



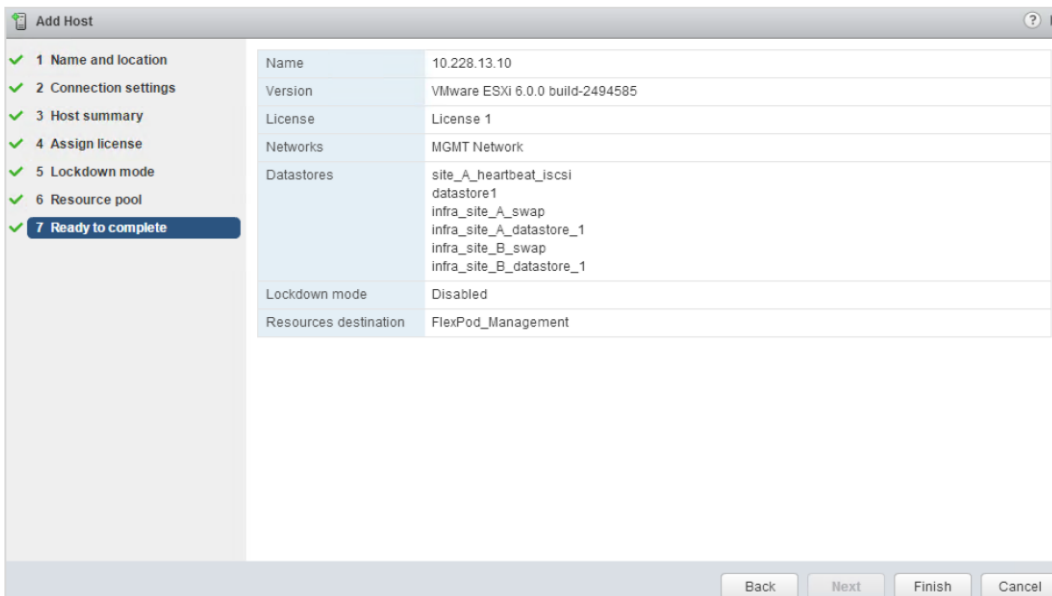
16. Under the Clusters pane, right-click FlexPod_Management and select Settings.
17. Select Configuration > General from the list on the left and click Edit to the right of General.
18. Select the Datastore Specified by Host option and click OK.



19. Under the Clusters pane, right-click FlexPod_Management and click Add Host.



20. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.
21. Type root as the user name and the root password. Click Next.
22. Click Yes to accept the certificate.
23. Review the host details and click Next.
24. Assign a license and click Next to continue.
25. Click Next and then click Next to continue.
26. Review the configuration parameters and click Finish to add the host.



27. Repeat steps 21 to 26 to add the remaining VMware ESXi hosts from both sites to the cluster.

Note: Four VMware ESXi hosts will be added to the cluster, two from each site.

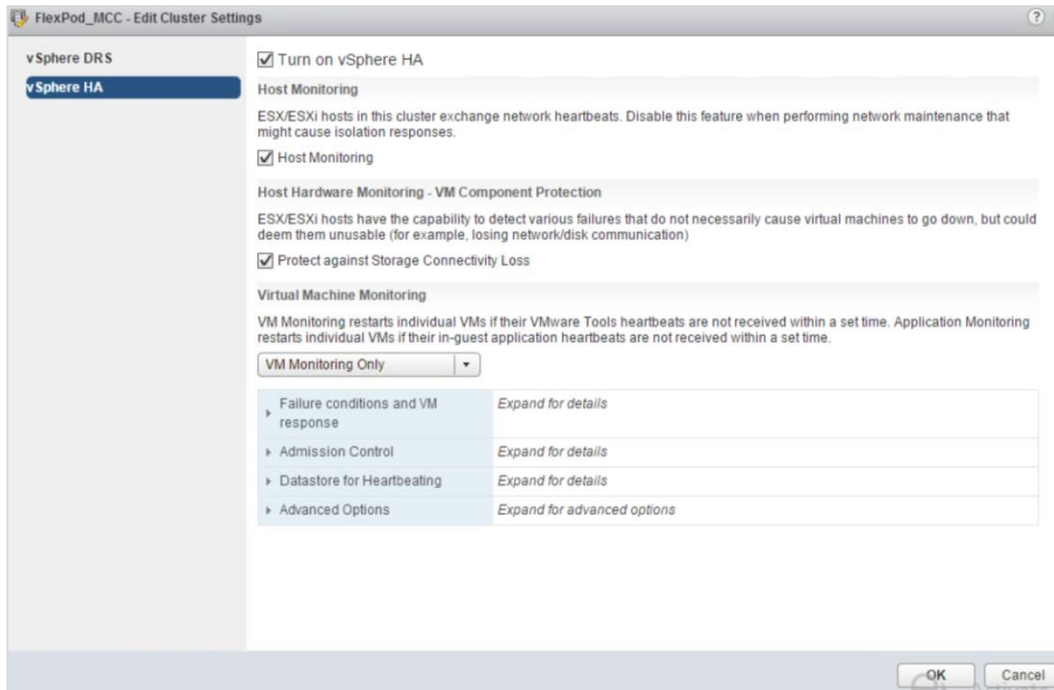
Best Practices for vSphere Metro Storage Cluster (vMSC) 6.0

This section describes best practices for configuring vSphere HA in a stretched cluster environment, such as MetroCluster.

Configure vSphere HA

Set the following best practices for vMSC for vSphere HA.

1. In vSphere HA settings, select Host Monitoring.
2. Select Protect against Storage Connectivity Loss.



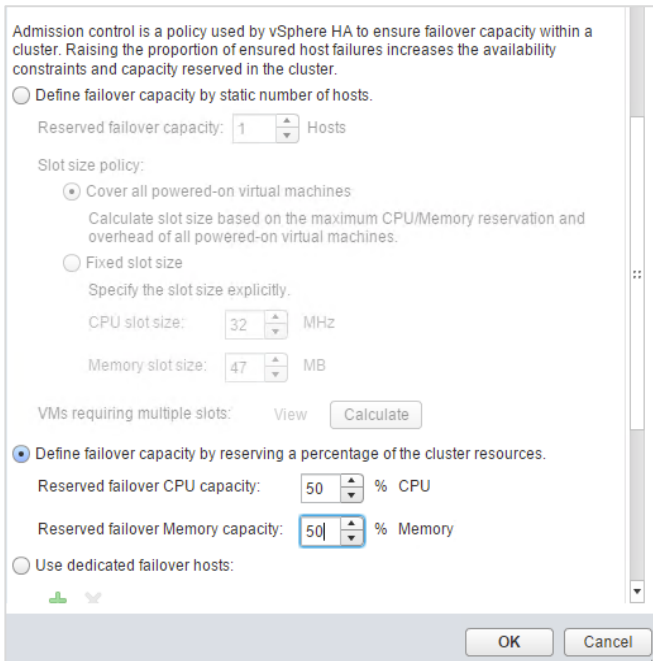
3. Expand Failure Conditions and VM Response:

- Set Response for Host Isolation to Power Off and Restart VMs.
- Set Response for Datastore with Permanent Device Loss (PDL) to Power Off and Restart VMs.
- Set Response for Datastore with All Paths Down (APD) to Power Off and Restart VMs (Conservative).
- Set Response for APD Recovery After APD Timeout to Disabled.

VM restart priority	Medium	⌵	When Disabled is selected, virtual machines are not restarted in the event of a host failure. In addition, they remain Protected when Turn on vSphere HA is enabled.
Response for Host Isolation	Power off and restart VMs	⌵	
Response for Datastore with Permanent Device Loss (PDL)	Power off and restart VMs	⌵	
Response for Datastore with All Paths Down (APD)	Power off and restart VMs (...)	⌵	
Delay for VM failover for APD	3	⬆️ ⬆️ ⬆️ ⬆️ ⬆️ ⬆️ ⬆️ ⬆️ ⬆️ ⬆️	minutes
Response for APD recovery after APD timeout	Disabled	⌵	

4. Expand Admission Control:

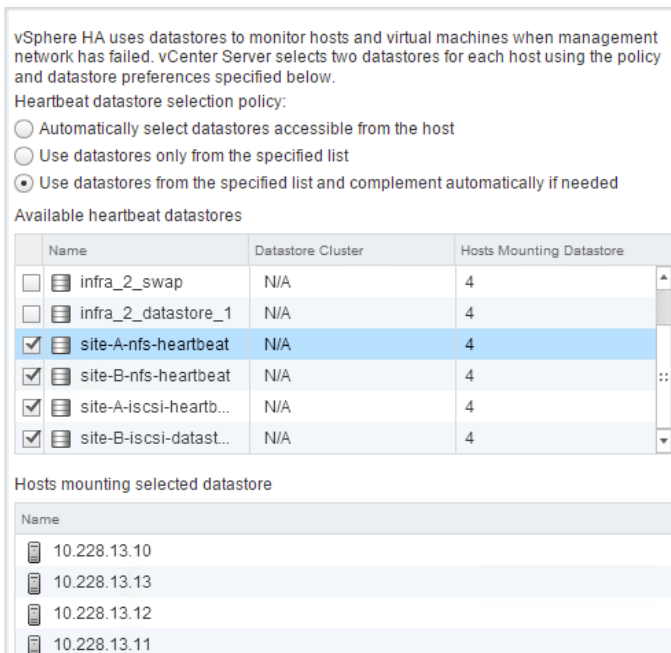
- Select Define Failover Capacity by Reserving a Percentage of the Cluster Resources.
- Set the Reserved CPU Capacity and the Reserved Failover Memory Capacity to 50%.



5. Expand the Datastore for Heartbeating option:

- Select the Use Datastores from the Specified List and Complement Automatically If Needed option.
- Select the datastores that should be used for datastore heartbeating.

Note: For the FlexPod MCC deployment, one NFS datastore and one iSCSI datastore were provisioned from each site.



6. Expand Advanced Options and set the configuration parameters as follows:

- das.heartbeatDsPerHost: 4
- das.isolationAddress0: <<vsphere isolation address 1>>
- das.isolationAddress1: <<vsphere isolation address 2>>

Advanced Options	
Configuration Parameters	
<input type="button" value="Add"/> <input type="button" value="Delete"/>	
Option	Value
das.heartbeatDsPerHost	4
das.isolationAddress0	10.228.13.14
das.isolationAddress1	10.228.13.16
das.respectVmHostSoftAffinityRules	true
das.respectVmVmAntiAffinityRules	false
das.useDefaultIsolationAddress	false

7. Click OK to finish editing the settings for vSphere HA.

Configure vSphere DRS Settings

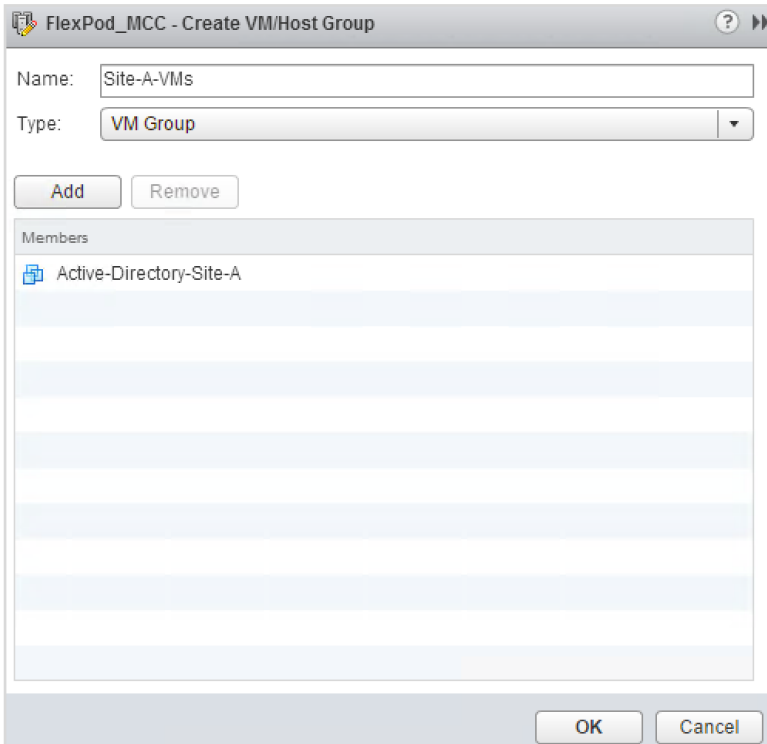
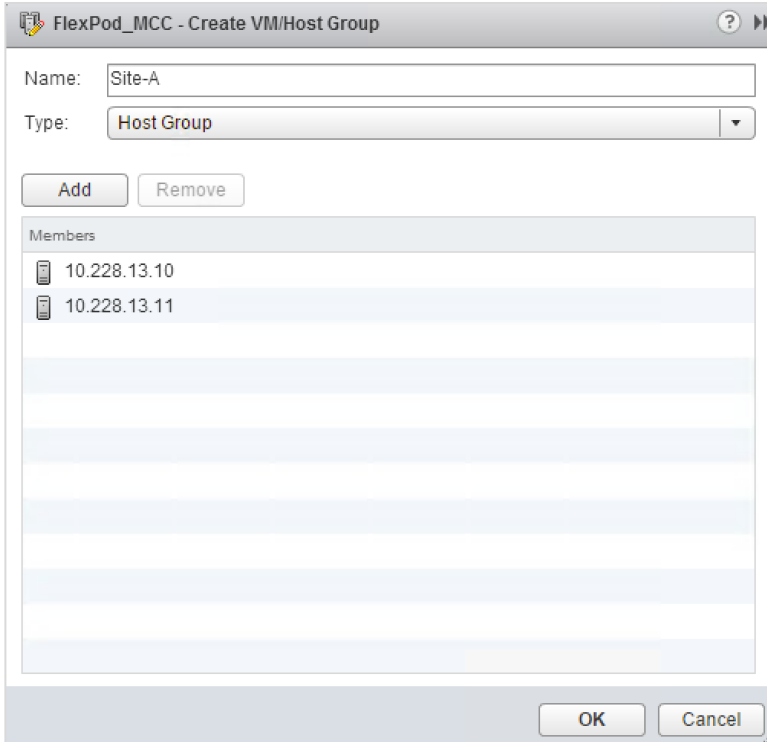
To control when VM migrations occur, VMware recommends configuring vSphere DRS in manual mode.

vSphere DRS is Turned ON		Schedule DRS...	Edit...
<p>▼ DRS Automation</p>			
Automation Level	Manual	vCenter Server will suggest migration recommendations for virtual machines.	
Migration Threshold	Apply only priority 1 recommendations. vCenter Server will only apply recommendations that must be taken to satisfy cluster constraints like affinity rules and host maintenance.		
Virtual Machine Automation	Individual virtual machine automation levels enabled.		
<p>▼ Power Management</p>			
Automation Level	Off	vCenter Server will not provide power management recommendations. Individual host overrides may be set, but will not become active until the cluster default is either Manual or Automatic.	
DPM Threshold	N/A		
▶ Advanced Options	None		

Create VM/Host Groups

To create VM and host groups, complete the following steps:

1. In the Create VM/Host Group dialog box, click the vSphere cluster, Manage, and Settings.
2. Click VM/Host Groups and click Add.
 - Create one host group for site A containing the ESXi hosts on site A.
 - Create one host group for site B containing the ESXi hosts on site B.
 - Create one VM group for site A containing the VMs that should run on site A hosts.
 - Create one VM group for site B containing the VMs that should run on site B hosts.



Create VM-to-Host Affinity Rules

1. Click VM/Host Rules and then click Add.
 - Create a VM/Host rule for site B VMs to run on site B ESXi hosts.

VM/Host Rules

Add... Edit... Delete

Name	Type	Enabled	Conflicts	Defined By
Site-B	Run VMs on Hosts	Yes	0	User
Site-A	Run VMs on Hosts	Yes	0	User

VM/Host Rule Details

Virtual Machines that are members of the VM Group should run on hosts that are members of the Host Group.

Add... Remove

Site-B-VMs Group Members	Site-B Group Members
Site-B-1	10.228.13.13
	10.228.13.12

- Create a VM/Host rule for site A VMs to run on site A ESXi hosts.

VM/Host Rules

Add... Edit... Delete

Name	Type	Enabled	Conflicts	Defined By
Site-B	Run VMs on Hosts	Yes	0	User
Site-A	Run VMs on Hosts	Yes	0	User

VM/Host Rule Details

Virtual Machines that are members of the VM Group must run on hosts that are members of the Host Group.

Add... Remove

Site-A-VMs Group Members	Site-A Group Members
Site-A-1	10.228.13.11
	10.228.13.10

- In the vSphere HA Rule Settings page, click Edit.
- Select the vSphere HA Should Respect Rules During Failover option.

vSphere HA Rule Settings Edit...

vSphere HA can enforce VM/Host rules when restarting virtual machines.

VM anti-affinity rules	Ignore rules
VM to Host affinity rules	vSphere HA should respect rules during failover

Set Up ESXi Dump Collector for iSCSI-Booted Hosts

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

1. In vSphere Web Client, click System Configuration.
2. In the left pane, click Services and then VMware vSphere ESXi Dump Collector.
3. From the Actions menu, select Start and click Edit Startup Type.
4. Select Automatic and click OK.
5. On the management workstation, open the VMware vSphere CLI command prompt.
6. Set each iSCSI-booted ESXi host to coredump to the ESXi Dump Collector.

```
esxcli -s <<var_vm_host_infra_XX_ip>> -u root -p <<var_password>> --thumbprint <host_thumbprint>
system coredump network set --interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip> --
server-port 6500
```

To get the host thumbprint, type the following command without the `--thumbprint` option, then copy and paste the thumbprint into the command.

```
esxcli -s <<var_vm_host_infra_XX_ip>> -u root -p <<var_password>> system coredump network set --
enable true
```

Cisco UCS Virtual Media (vMedia) Policy for VMware ESXi Installation

Set Up Storage Controller for vMedia Policy

To set up the NetApp storage cluster, complete the following steps:

1. Connect to the NetApp storage cluster through SSH.
2. Allow the in-band management VLAN to access the infrastructure datastores.
3. Enter the following command:

```
vserver export-policy rule create -policyname default -<<var_inband_mgmt_subnet_cidr>> -rorule
sys -rwrule sys -allow-suid false -vserver Infra_SVM_site_(A/B) -ruleindex 3 -protocol nfs -
superuser sys
```

4. Create two ports and then add those ports to the newly created broadcast domain.

```
network port vlan create -node clus-01 -vlan-name a0a-<<var_ib_mgmt_vlan_id>>
network port vlan create -node clus-01 -vlan-name a0a-<<var_ib_mgmt_vlan_id>>
broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500 -ports clus-01:a0a-
<<var_ib_mgmt_vlan_id>>, clus-02:a0a-<<var_ib_mgmt_vlan_id>>
```

5. Create an additional network interface for the infrastructure datastore to be accessed from the management VLAN network.

```
network interface create -vserver Infra_SVM_site_(A/B) -lif nfs_IB-MGMT -role data -data-protocol
nfs -home-node clus-02 -home-port a0a-<<var_ib_mgmt_vlan_id>> -address <<var_inband_nfs_ip>> -
netmask <<var_inband_mgmt_vlan_mask>> -status-admin up -failover-policy broadcast-domain-wide -
firewall-policy data -auto-revert true -failover-group IB-MGMT
```

6. From the vSphere interface, click the Home tab and select Storage.
7. Expand FlexPod_DC and Infra_datastore 1.
8. Right-click and select Browse Files.
9. Click the third icon to create a new folder; name this folder Software and click Create.
10. Select the software folder in the list on the left; click the first icon to upload a file to the datastore.
11. Browse to the VMware Cisco custom ISO and upload it to the datastore.
12. Log in to the Cisco UCS Manager.

13. Select the Servers tab. Go to Policies > root > vMedia Policies.
14. Right-click vMedia Policies and select Create vMedia Policy.
15. Enter the policy name and then click the green plus sign icon.
16. In the Create vMedia Mount dialog box, enter the vMedia mount name, IP address, and remote file name and its path. Click OK.

Create vMedia Mount

Name:

Device Type: CDD HDD

Protocol: NFS CIFS HTTP HTTPS

Hostname/IP Address:

Image Name Variable: None Service Profile Name

Remote File:

Remote Path:

OK Cancel

17. Click OK in the Create vMedia Policy page.

Create vMedia Policy

Name:

Description:


Retry on Mount Failure: No Yes

vMedia Mounts

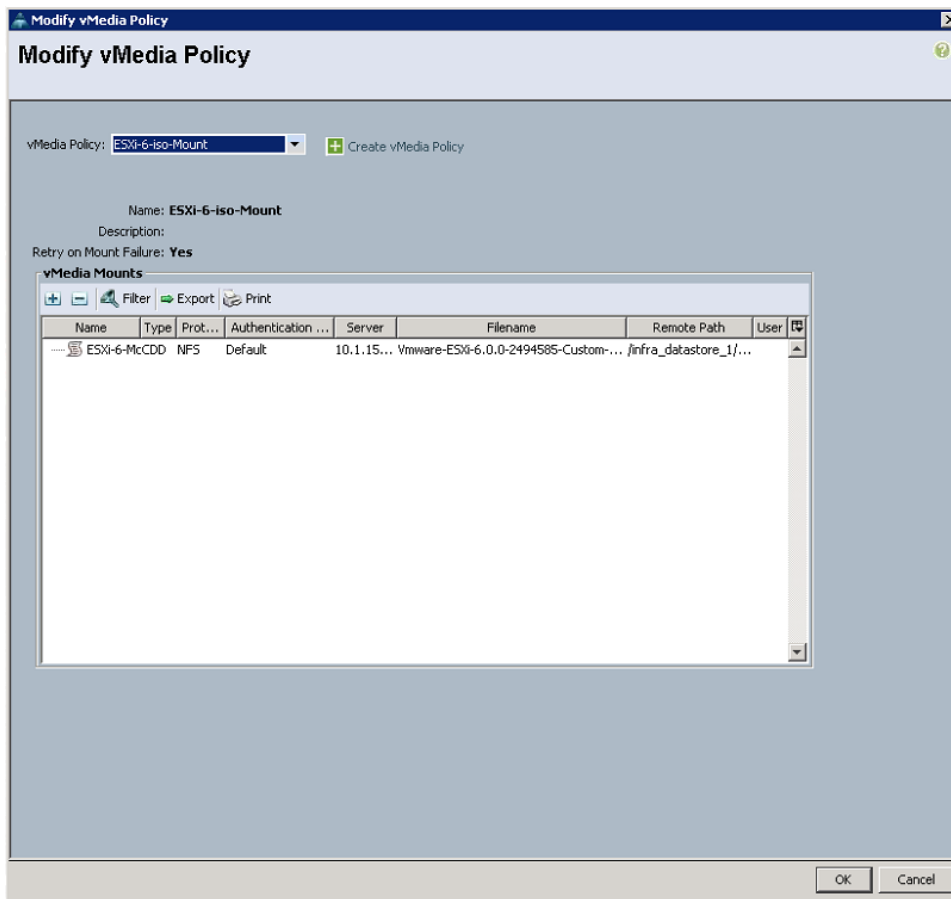
Name	Type	Protocol	Authentication Protocol	Server	Filename	Remote Path	User
ESXCDD	NFS	NFS	Default	10.1.156...	Vmware-E...	/infra_datast...	

OK Cancel

18. Select Service Profile Templates > root.
19. Right-click Service Profile Template and select Create a Clone.
20. Name the clone and select root for Org and click OK.



21. Select the template that was just created and select the vMedia Policy tab.
22. In the Actions pane, select Modify vMedia Policy.
23. Select the ESXi-6-iso-Mount policy from the drop-down menu and click OK.



24. Click OK in the pop-up.

Note: For any new servers added to the Cisco UCS environment, the vMedia service profile template can be used to create the service profile. On first boot, the host will boot into the ESXi installer. After ESXi is installed, you can unbind from the vMedia Service Profile Template and bind to the original Service Profile Template. The ESXi installer will not be automatically mounted.

6 Solution Verification

This solution validation leverages the VMware vSphere MetroCluster storage cluster configuration. This section describes the failure scenarios that were tested in the infrastructure.


6.1 Single Blade Failure

Table 10 summarizes the test scenario in which one blade in the infrastructure was failed. The VMs were checked to make sure that they recovered correctly.

Table 10) Single host failure in site A data center.

Test Case	Details
Test number	Test-001
Date	5/16/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	The loss of one blade in the Cisco UCS chassis causes the loss of one of the four ESXi hosts in the VMware vSphere cluster. When the master node in the cluster notices that there is no heartbeat or datastore heartbeat from the node, it powers off and restarts all VMs, which previously ran on that host in accordance with the VM-to-host affinity rules.
Test results	The test completed as expected.
Comments	

Figure 12) Single host failure in site A data center.



Test Procedure

1. Make sure that the FlexPod with MetroCluster system runs in a steady state.
2. Identify an ESXi server on site A that you want to fail.
3. Identify the corresponding blade number of the ESXi.
4. Manually pull out the blade server from the chassis while it runs.
5. Verify that vCenter shows the disconnected host and that the VMs are restarted on the correct host according to the VM-to-host affinity rules.

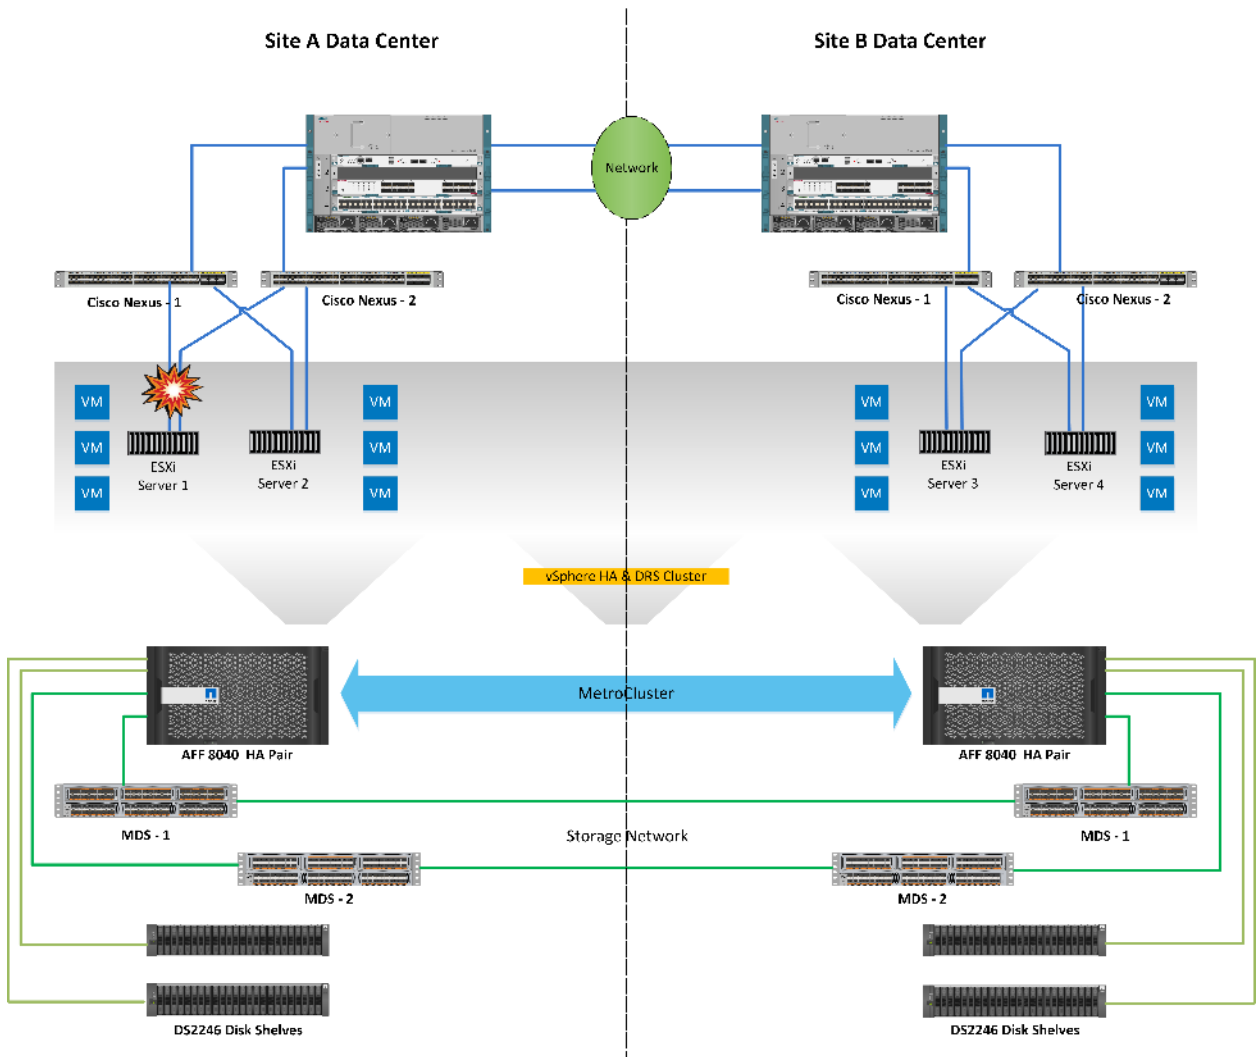
6.2 Single Host Isolation

Table 11 summarizes the test scenario in which access to the in-band management network for one ESXi host was removed. The VMs that were restarted on the other host were verified for accordance with the VM-to-host affinity rules.

Table 11) Single host isolation in site A data center.

Test Case	Details
Test number	Test-002
Date	6/21/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	vCenter should detect the loss of management access to the ESXi host. Because this is a converged solution with NFS traffic and management traffic passing over the same port channel, the VMware vSphere setting has been set to power off and restart VMs. All VMs should restart in accordance with the VM-to-host affinity rules.
Test results	The test completed as expected.
Comments	Another response to host isolation could be to make no changes to the VMs running on the host because it is just isolated, not powered down. However, the most obvious reason that the host would be isolated in this configuration is that a networking error occurred. Because the traffic is being passed on the same port channel, NetApp recommends that the VMs be restarted on a different host.

Figure 13) Single host isolation in site A data center.



Test Procedure

1. Make sure that site A and site B run in an ideal state.
2. Identify an ESXi server on site A that you want to isolate.
3. Identify the VMNICs that are used for management traffic.
4. Enable the ESXi shell for the ESXi.
5. From the ESXi shell, run the following command for each VMNIC:

```
esxcli network nic down -n vmnic<<X>>
```

6.3 Fabric Interconnect Reboot

Table 12 summarizes the test scenario in which a fabric interconnect is rebooted and then verified to make sure that the traffic fails over to the surviving fabric interconnect.

Table 12) Fabric interconnect test details.

Test Case	Details
Test number	Test-003
Date	6/22/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	Traffic should fail over to the available fabric interconnect on the same site and all VMs should continue to operate normally.
Test results	The test completed as expected.
Comments	

Test Procedure

1. Make sure that site A and site B run in an ideal state.
2. Log into one fabric interconnect and reboot it.

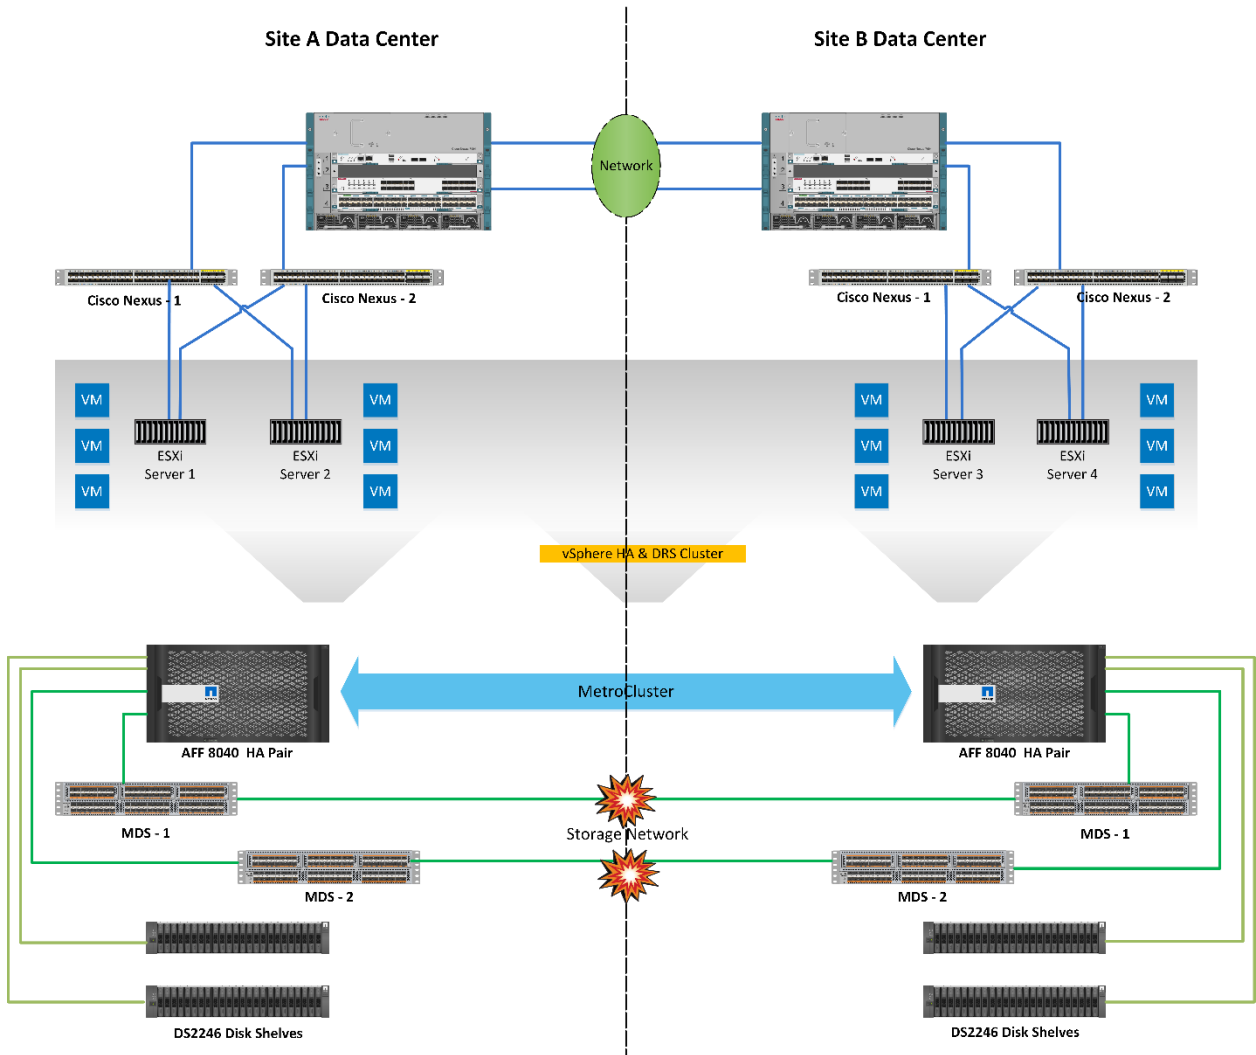
6.4 Storage Isolation

Table 13 summarizes the test scenario in which the storage network between the two data centers is brought down.

Table 13) Storage partition.

Test Case	Details
Test number	Test-004
Date	6/29/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	There is no change because of the storage isolation. No takeover will occur and both NetApp clusters will continue to serve data to their respective sites. After the ISLs between the Cisco MDS switches are back up, the new data will be resynced back to the other site.
Test results	The test completed as expected.
Comments	

Figure 14) Storage partition.



Test Procedure

3. Make sure that site A and site B run in an ideal state.
4. Shut down the ports on each MDS9396s switch that are used to interconnect the MDS switches on either site.

6.5 Full Data Center Isolation

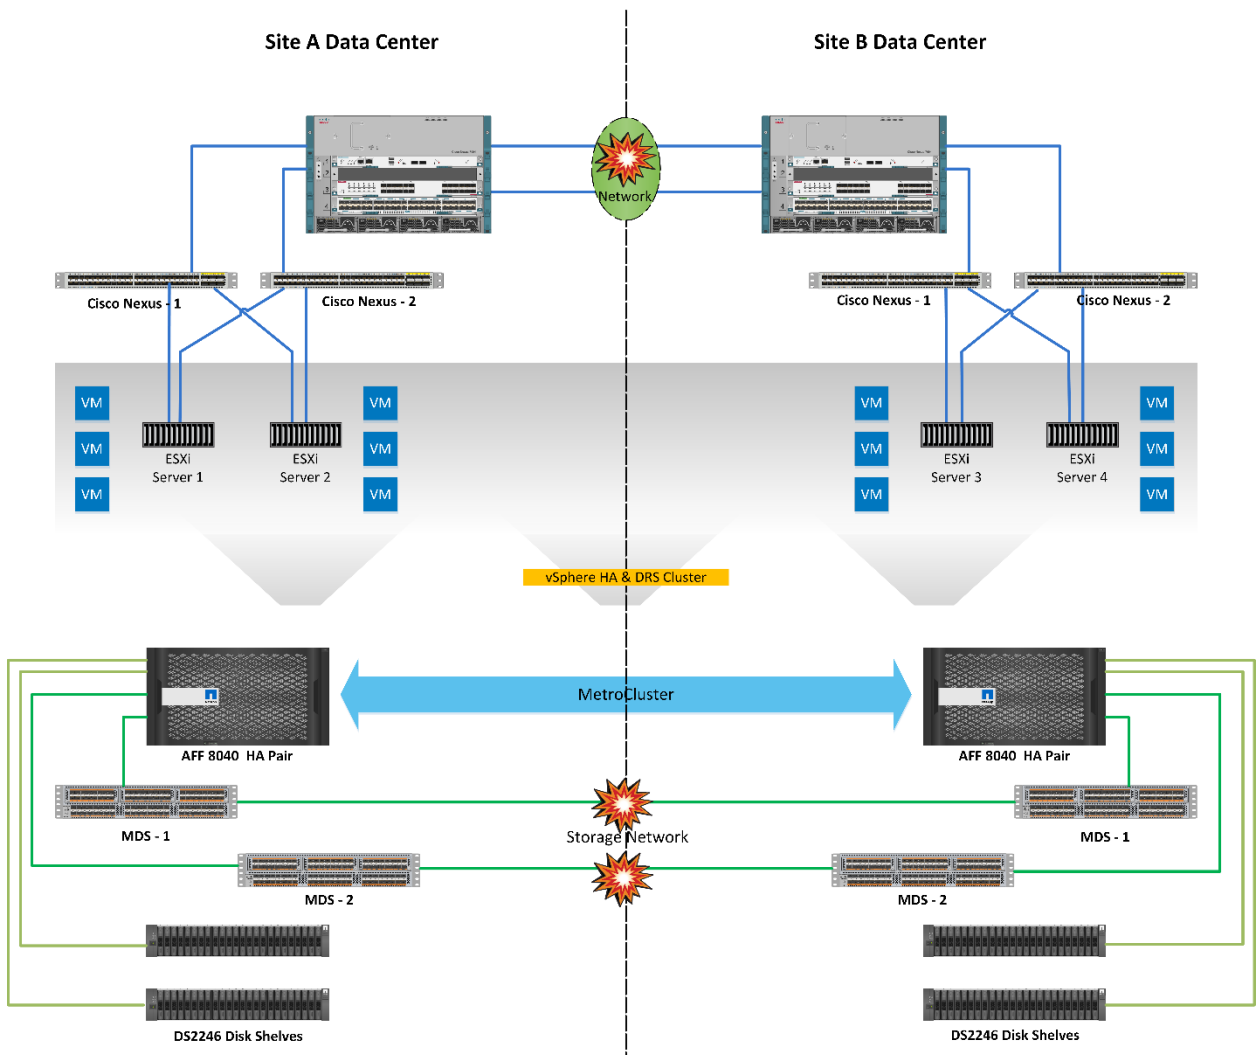
Table 14 summarizes the test scenario in which the site A data center and site B data center are isolated from each other.

Table 14) Data center partition.

Test Case	Details
Test number	Test-005
Date	6/22/16

Test Case	Details
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	NetApp MetroCluster will continue to operate normally during full data center isolation. ESXi hosts will lose access to heartbeat and storage datastores from the other site because of the loss of connection between the two sites. If a VM was running on a datastore presented from the other site, it will be restarted on the site that still has access to the datastore. All other VMs will continue to operate normally.
Test results	The test completed as expected.
Comments	Depending on the location of the vCenter server, it might lose access to the two sites. NetApp assumes that the vCenter server was able to access both sites during the test through an out-of-band management network.

Figure 15) Data center partition.



Test Procedure

1. Make sure that site A and site B run in an ideal state.
2. Shut down the ports on each MDS9396s switch that are used to interconnect the MDS switches on either site.
3. Shut down the ports on both the Cisco Nexus 7000 switches that are used to interconnect with each other.

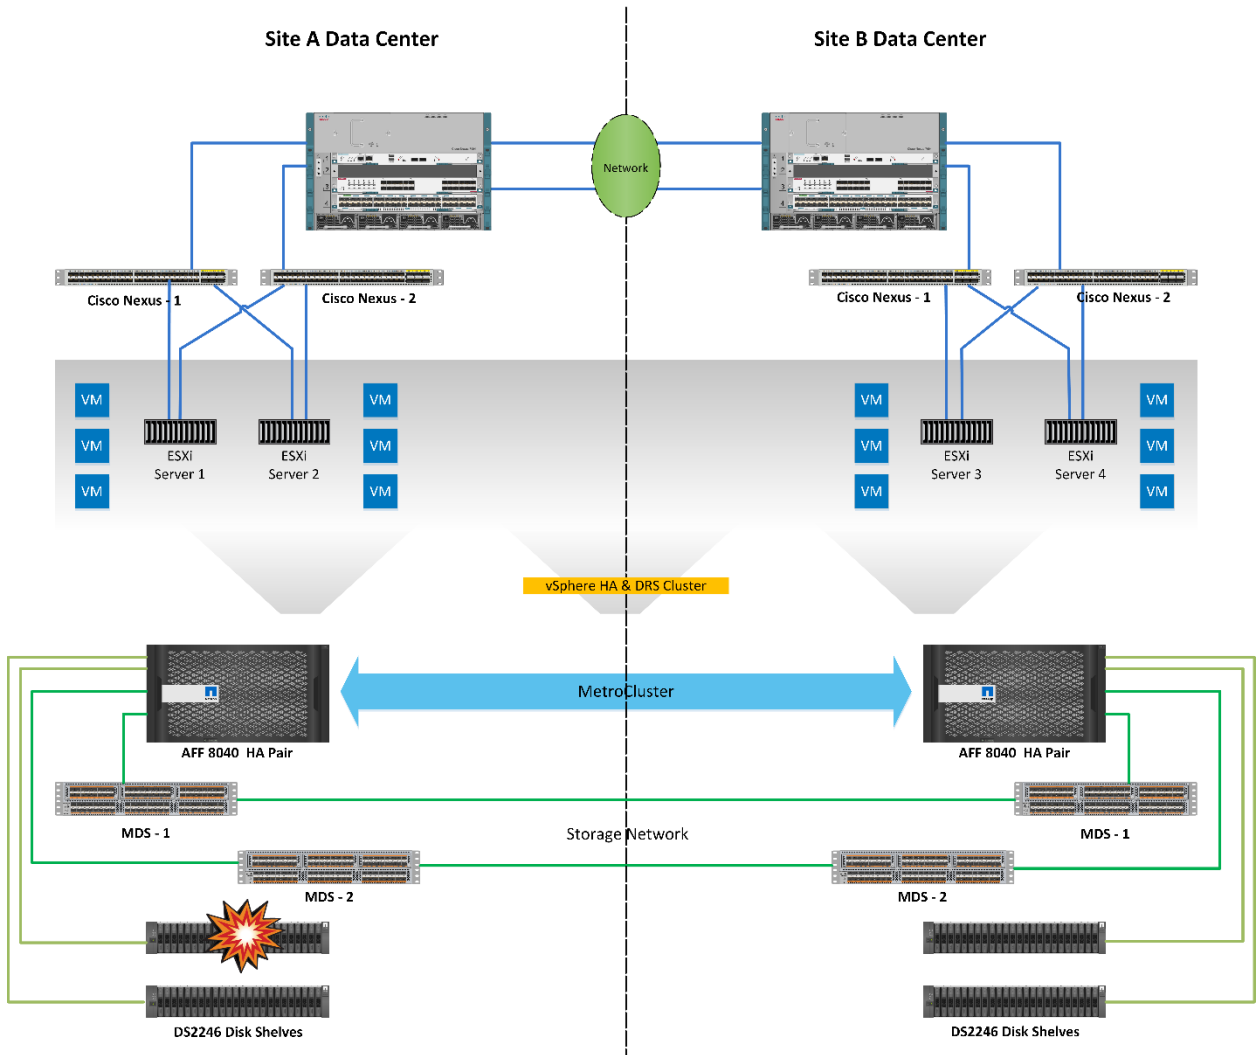
6.6 Storage Degradation in a Single Data Center

In this test, one of the disk shelves in site A was made unavailable. Table 15 summarizes the test results of this scenario.

Table 15) Disk shelf failure in site A data center.

Test Case	Details
Test number	Test-006
Date	1/27/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	The VMs will continue to run with no impact. NetApp storage will be aware of this and the data will be served from the mirror copy available in site B.
Test results	One shelf disappeared from the <code>storage shelf show</code> output. The MetroCluster remained in a steady state and continued to serve data.
Comments	The storage latency might go up for a short period of time during the switchover.

Figure 16) Disk shelf failure in site A data center.



Test Procedure

1. Make sure that site A and site B run in an ideal state.
2. Identify the disk shelf that you would like to make unavailable.
3. Power off the disk shelf.
4. Make sure that the NetApp storage remains up and serving data.

6.7 Planned Switchover and Switchback

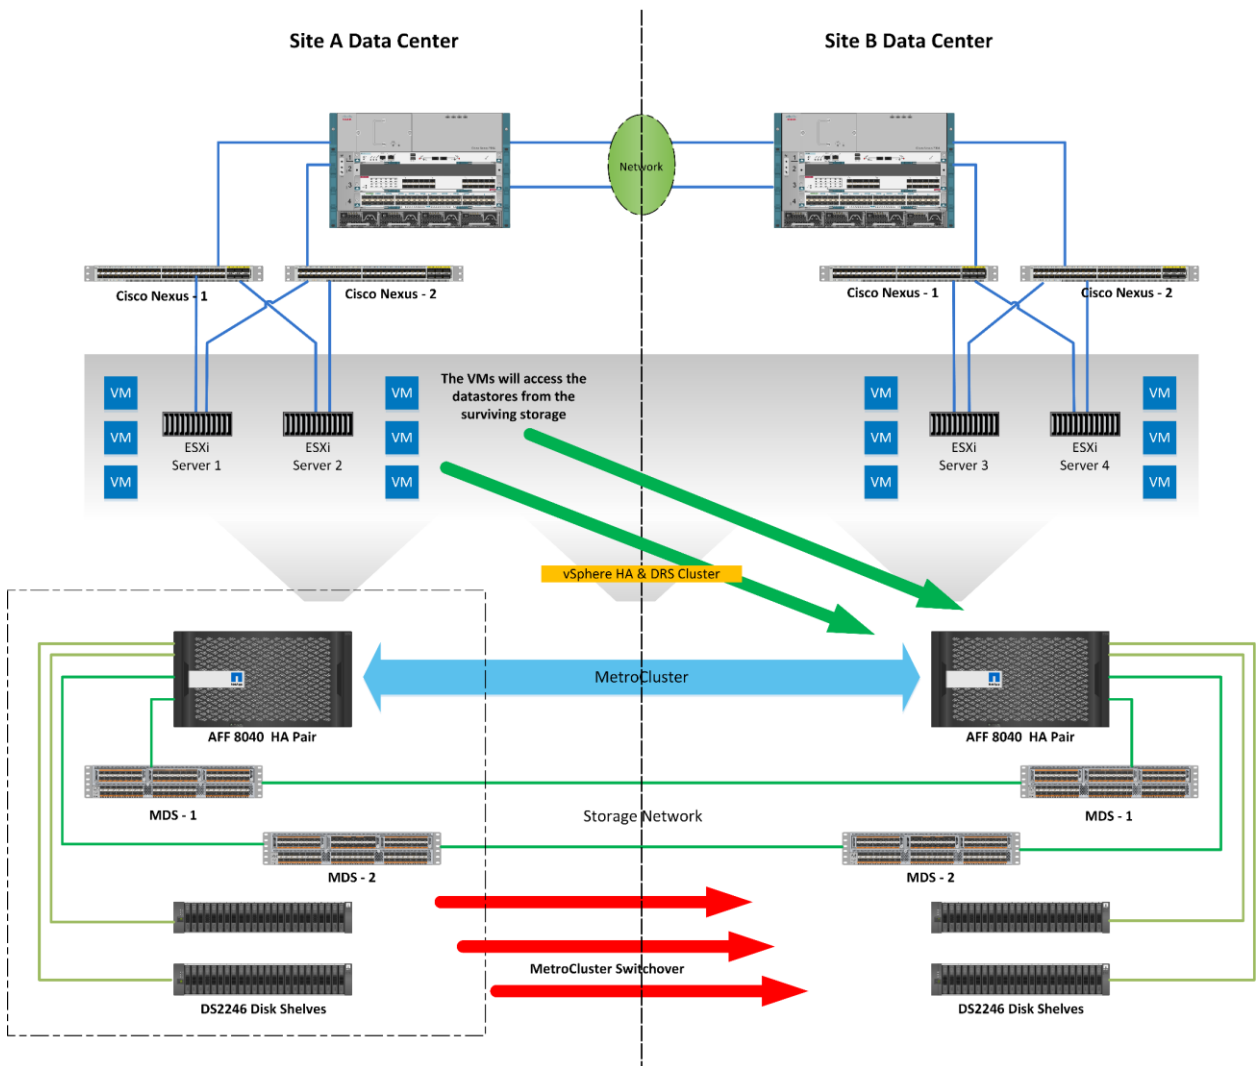
In this test, all data SVMs (formerly Vservers) in the MetroCluster system were switched over to a single site. Table 16 summarizes the test results of this scenario.

Table 16) Planned switchover of storage.

Test Case	Details
Test number	Test-007
Date	6/15/16

Test Case	Details
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	The VMs will continue to run with no impact. All paths will be presented from the surviving storage.
Test results	The vSphere cluster temporarily saw only two out of the four heartbeat datastores after performing this procedure. After about a minute, the cluster recognized that the datastores were present. No issues were observed from the ESXi cluster.
Comments	A planned switchover could be executed for maintenance purposes or testing for a full disaster scenario.

Figure 17) Planned switchover of storage.



Test Procedure

1. Make sure that the FlexPod MetroCluster runs in a steady state.
2. Execute `metrocluster switchover` from the cluster interface on one site.
3. Answer `y` when prompted.
4. Wait for the switchover to succeed.
5. Make sure that the vSphere cluster runs in the normal state.
 - All iSCSI-booted ESXi hosts should be running.
 - All VMs should continue to run on the same servers.
 - No alarms should be seen in the cluster.
6. Run `metrocluster heal -phase aggregates` on the surviving site.
7. Run `metrocluster heal -phase root-aggregates` on the surviving site.
8. Boot the nodes that were previously halted.
9. Run `metrocluster switchover` on the surviving site.

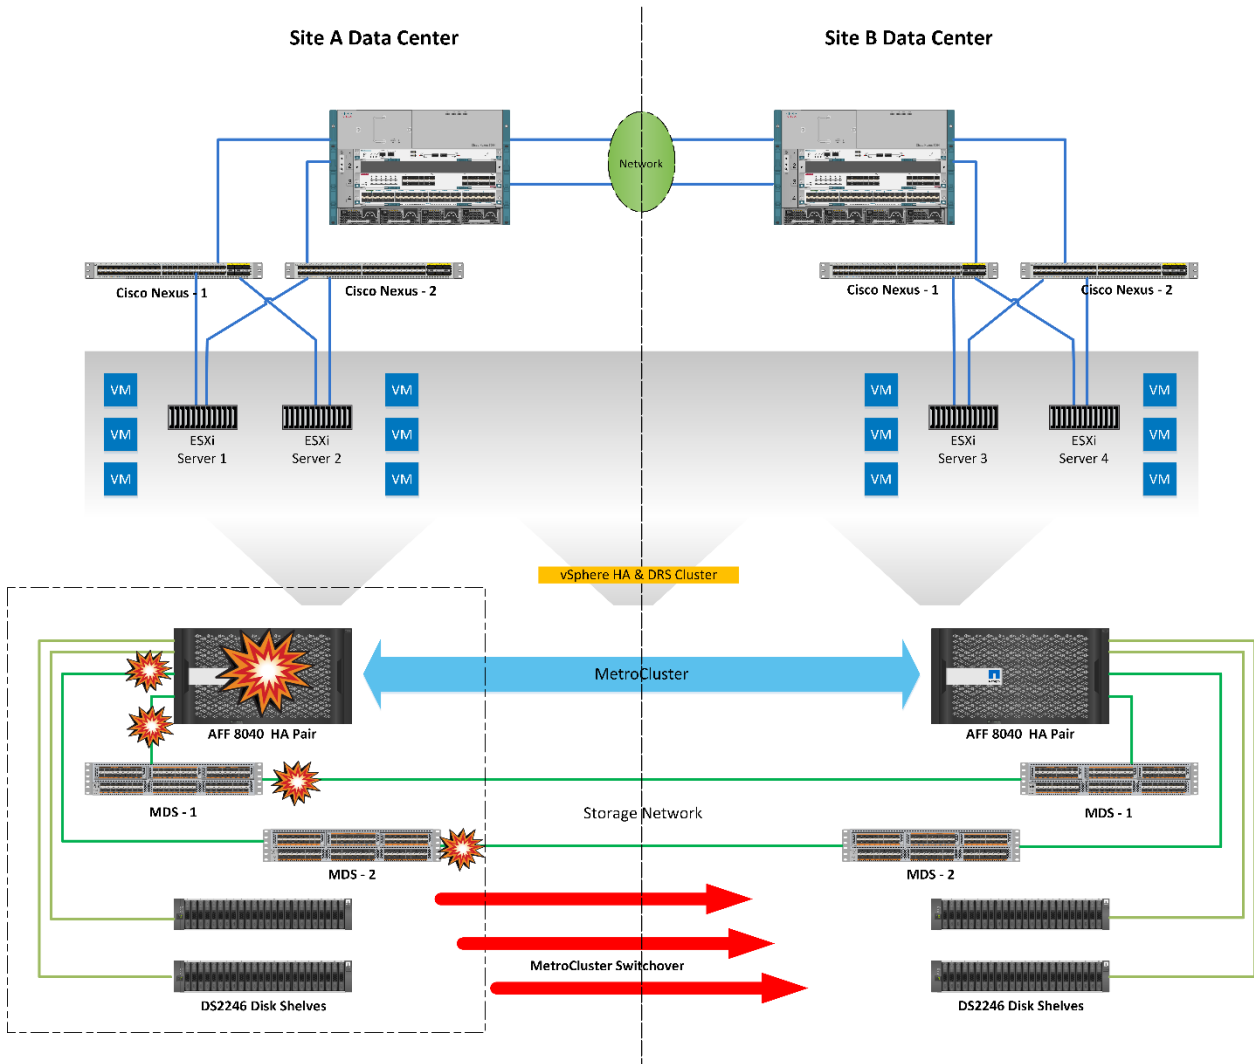
6.8 Full Storage Failure in a Single Data Center

In this test, the storage is failed in one data center and a switchover is executed on the surviving site so that there is no downtime for VMs and applications in the vSphere cluster. Table 17 summarizes the test results of this scenario.

Table 17) Full storage failure in site A data center.

Test Case	Details
Test number	Test-008
Date	6/15/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	The vSphere cluster will not be affected because of a full storage disaster. Control will be switched to the surviving NetApp MetroCluster site shortly after the disaster takes place to prevent the all-paths-down scenario from taking too long. Because of this, all target LIFs will appear on the surviving site within 120 seconds of the disaster.
Test results	Switchover completed successfully. The two ESXi nodes on the disaster site experienced a temporary loss of all paths to their iSCSI boot LUN. They temporarily ran from memory until the LUN came back online. ESXi hosts saw a temporary loss of two of the heartbeat datastores before the switchover was complete. No VMs were rebooted or moved.
Comments	The error about the loss of the boot LUN persists even after the paths are back up. Restarting the management network will fix this issue.

Figure 18) Full storage failure in site A data center.



Test Procedure

1. Make sure that the FlexPod MetroCluster runs in a steady state.
2. Simulate a disaster by executing a `shutdown` command on all MDS switch ports on the disaster site.
3. Power off both nodes in the site using the service processor.
4. Run `metrocluster switchover -forced-on-disaster` on the surviving site.
5. Wait for switchover to complete.
6. Make sure that the vSphere cluster remains in a steady state and that all ESXi hosts are running.

Note: Another option is to pull out the power cables of the storage controllers.

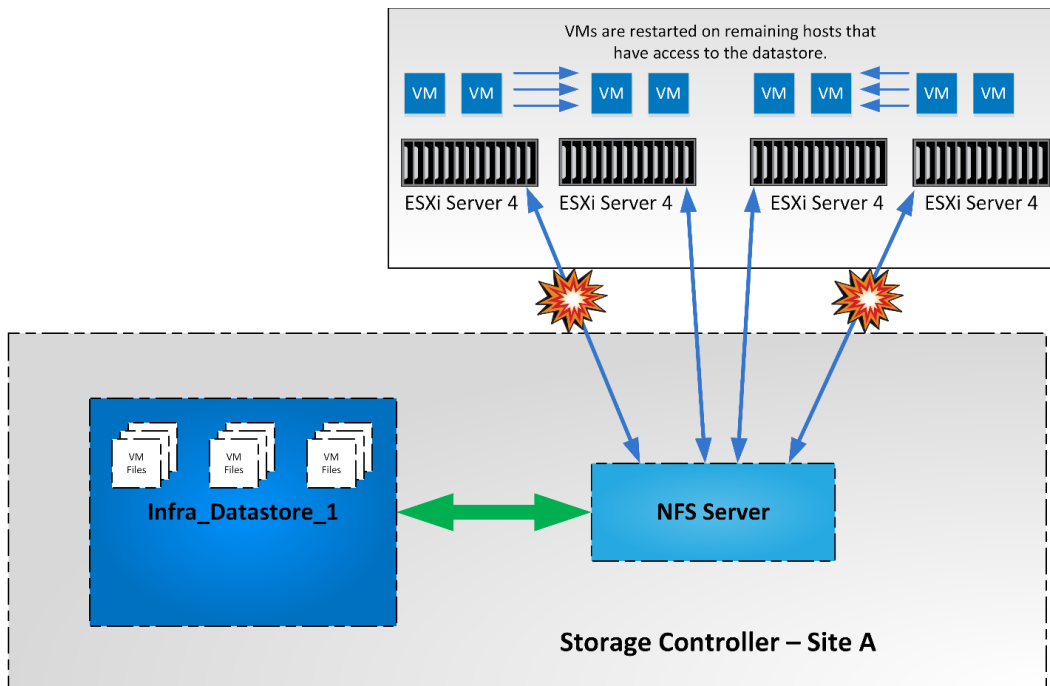
6.9 Volume Failure in a Single Data Center

In this test, datastore access to certain ESXi hosts is removed. The VMs should transition to the surviving datastore. Table 18 summarizes the test results of this scenario.

Table 18) Permanent device loss.

Test Case	Details
Test number	Test-009
Date	6/22/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	The VMs running on the host will be restarted on remaining hosts that have access to the datastore.
Test results	The test completed as expected.
Comments	

Figure 19) Permanent device loss.



Test Procedure

1. Make sure that the FlexPod MetroCluster runs in a steady state.
2. From the NetApp storage console, remove the `vserver export-policy` rules for the ESXi hosts that should not have access to the datastore.

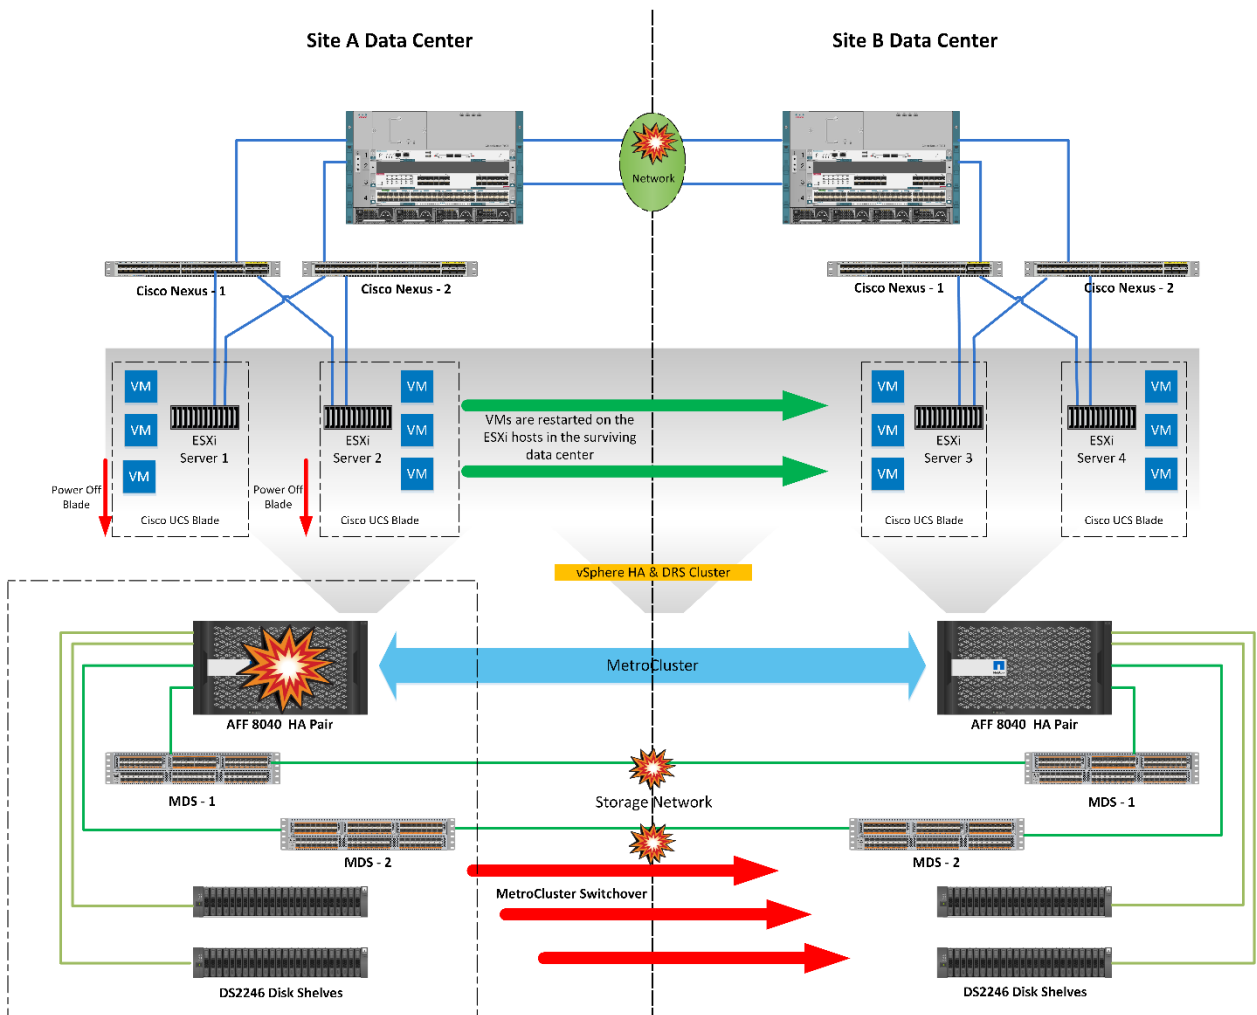
6.10 Loss of a Data Center

In this test scenario, an entire site is lost and the surviving site must take control of all operations. Table 19 summarizes the test results of this scenario.

Table 19) Full compute failure in site A data center.

Test Case	Details
Test number	Test-010
Date	6/27/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	After a loss of the compute from one site and the storage from the same site, the <code>metrocluster switchover -forced-on-disaster true</code> command needs to be run on the surviving site to make all datastores available to the surviving compute nodes. VMware vSphere HA will detect the loss of half of the nodes from the cluster and will restart VMs on the surviving nodes even though doing so violates VM-to-host affinity rules.
Test results	The test completed as expected.
Comments	

Figure 20) Full compute failure in site A data center.



Test Procedure

1. Make sure that the FlexPod MetroCluster runs in a steady state.
2. Shut the ports on both the Cisco Nexus 7000 switches that are used to interconnect them.
3. Shut the ports on each MDS9396s switch that are used to interconnect the MDS switches on either site.
4. Cut the power to the storage nodes on one site using the service processor.
5. Remove access to the compute nodes by powering off the blades on the same site.
6. Run the `metrocluster switchover -forced-on-disaster true` command from the surviving site.
7. All VMs should be restarted on the surviving nodes of the vSphere cluster.

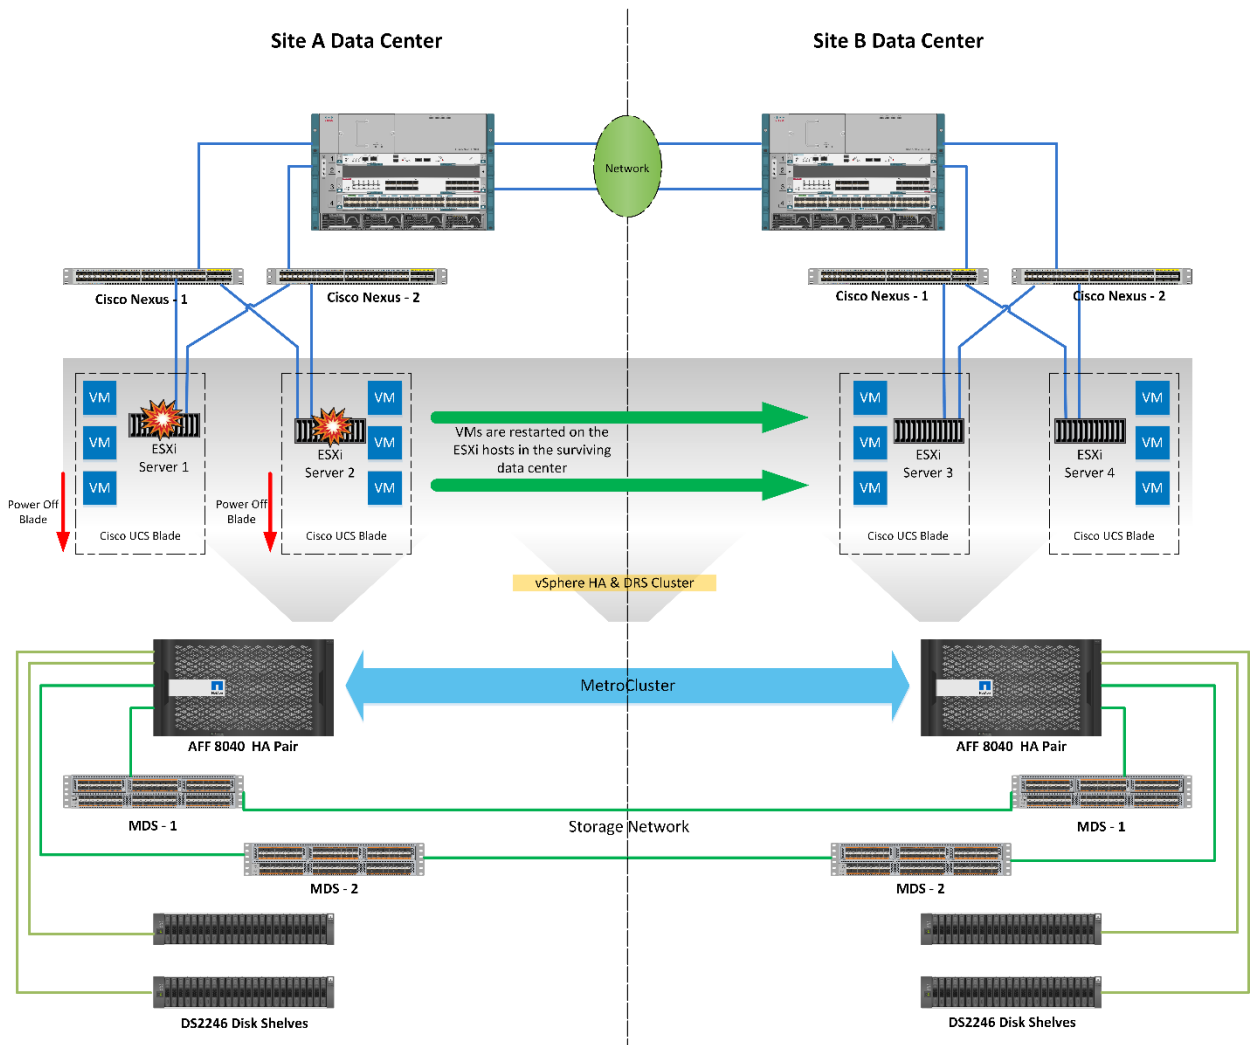
6.11 Full Compute Failure in a Data Center

In this test scenario, the loss of all compute in one site is simulated. Table 20 summarizes the test results of this scenario.

Table 20) Loss of site A data center.

Test Case	Details
Test number	Test-011
Date	6/24/16
Test prerequisites	The FlexPod MetroCluster system should be configured according to the best practices for FlexPod, MetroCluster, and VMware, as described in this document.
Expected outcome	The VMware vSphere HA will detect the loss of half of the nodes from the cluster and restart all VMs in the surviving site.
Test results	The test completed as expected.
Comments	

Figure 21) Loss of site A data center.



Test Procedure

1. Power off the blades in one site.
2. Observe that the VMs are restarted on the surviving site by vSphere HA.

7 Conclusion

Businesses often risk loss of revenue, employee productivity, and brand perception when key infrastructure and services go offline. Although many of the reasons for downtime can be mitigated, it is difficult to maintain business continuity after the loss of power to an entire rack or data center. Because of the flexibility of FlexPod Datacenter to be deployed over two sites in combination with NetApp MetroCluster, businesses can be sure that key applications remain online—even in a large power outage.

Acknowledgements

The authors of this document would like to thank David Klem of NetApp for his support and contributions to this project.

References

This report references the following documents and resources:

- [FlexPod Datacenter with NetApp MetroCluster Design Guide](#)
- [MetroCluster Installation and Configuration Guide](#)
- [2016 - Cost of Data Center Outages](#)

Appendixes

Configuration Variables

Table 21 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used during deployment.

Table 21) Configuration variables.

Variable	Value
License Configuration Variables	
<<var_cluster_base_license>>	Storage cluster base license key
<<var_iSCSI_license_key>>	Storage iSCSI license key
<<var_NFS_license_key>>	Storage NFS license key
<<var_FlexClone_license_key>>	Storage NetApp FlexClone® license key
<<var_SnapMirror_license_key>>	Storage NetApp SnapMirror® license key
Global Configuration Variables	
<<var_oob-mgmt_gateway>>	Out-of-band network gateway
<<var_oob-mgmt_mask>>	Out-of-band network mask
<<var_oob-mgmt_ntp>>	IP address of the NTP server on the out-of-band network
<<var_oob-mgmt_domain>>	Domain name of the out-of-band network
<<var_oob-mgmt_dns>>	DNS IP address of the out-of-band network
<<var_oob-mgmt_ntp>>	Domain name of the out-of-band network
<<var_ib-mgmt_mask>>	In-band network mask
Cisco Nexus 9372PX Configuration Variables	
<<var_site(A/B)_9K(1/2)>>	Cisco Nexus 9K switch name
<<var_site(A/B)_9K(1/2)_ip_address>>	Cisco Nexus 9K switch IP address
<<var_siteA_9K1_ntp_ip_address>>	Site A - Cisco Nexus 9K #1 NTP distribution IP address

Variable	Value
<<var_siteA_9K2_ntp_ip_address>>	Site A - Cisco Nexus 9K #2 NTP distribution IP address
<<var_siteB_9K1_ntp_ip_address>>	Site B - Cisco Nexus 9K #1 NTP distribution IP address
<<var_siteB_9K2_ntp_ip_address>>	Site B - Cisco Nexus 9K #2 NTP distribution IP address
NetApp AFF8040 Configuration Variables	
<<var_siteA_node1>>	Site A - NetApp AFF8040 #1 host name
<<var_siteA_node1_ip_address>>	Site A - NetApp AFF8040 #1 IP address
<<var_siteA_node1_port>>	Site A - NetApp AFF8040 #1 MGMT port
<<var_siteA_node2>>	Site A - NetApp AFF8040 #2 host name
<<var_siteA_node2_ip_address>>	Site A - NetApp AFF8040 #2 IP address
<<var_siteA_node2_port>>	Site A - NetApp AFF8040 #2 MGMT port
<<var_siteB_node1>>	Site B - NetApp AFF8040 #1 host name
<<var_siteB_node1_ip_address>>	Site B - NetApp AFF8040 #1 IP address
<<var_siteB_node1_port>>	Site B - NetApp AFF8040 #1 MGMT port
<<var_siteB_node2>>	Site B - NetApp AFF8040 #2 host name
<<var_siteB_node2_ip_address>>	Site B - NetApp AFF8040 #2 IP address
<<var_siteB_node2_port>>	Site B - NetApp AFF8040 #2 MGMT port
NetApp AFF8040 Cluster Configuration Variables	
<<var_site(A/B)_cluster>>	Host name for cluster in site A/B
<<var_site(A/B)_cluster_ip_address>>	IP address for cluster in site A/B
<<var_site(A/B)_cluster_port>>	MGMT port for cluster in site A/B
Cisco UCS 6248UP Configuration Variables	
<<var_siteA_FI1>>	Site A - fabric interconnect #1 host name
<<var_siteA_FI2>>	Site A - fabric interconnect #2 host name
<<var_siteB_FI1>>	Site B - fabric interconnect #1 host name
<<var_siteB_FI2>>	Site B - fabric interconnect #2 host name
Cisco Nexus 7004 Configuration Variables	

Variable	Value
<<var_site(A/B)_7K_(OTV/LAN)>>	Host name for LAN/OTV VDC on site A or B
<<var_site(A/B)_7K_(OTV/LAN)_ip_address>>	IP address for LAN/OTV VDC on site A or B
<<var_siteA_7K>>	Site A - Cisco Nexus 7004 host name
<<var_siteA_7K_ip_address>>	Site A - Cisco Nexus 7004 IP address
<<var_siteB_7K>>	Site B - Cisco Nexus 7004 host name
<<var_siteB_7K_ip_address>>	Site B - Cisco Nexus 7004 IP address
ATTO FibreBridge 7500N Configuration Variables	
<<var_site(A/B)_fibrebridge_(1/2)>>	Host name of the FibreBridge controllers in site A or B
<<var_site(A/B)_fibrebridge_(1/2)_ip_address>>	IP address of the FibreBridge controllers in site A or B
Cisco MDS 9396S Configuration Variables	
<<var_site(A/B)_mds(1/2)>>	Host name of the MDS switches in site A/B
<<var_site(A/B)_mds(1/2)_ip_address>>	IP address of the MDS switches in site A/B
Cisco UCS Configuration Variables	
<<var_vm_host_infra_fabric_A1_iqn>>	IQN initiator of the ESXi blade 1 in site A
<<var_vm_host_infra_fabric_A2_iqn>>	IQN initiator of the ESXi blade 2 in site A
<<var_vm_host_infra_fabric_B1_iqn>>	IQN initiator of the ESXi blade 1 in site B
<<var_vm_host_infra_fabric_B2_iqn>>	IQN initiator for ESXi blade 2 in site B

Cabling Details

This section details the connections made for a single site of FlexPod Datacenter with NetApp MetroCluster. These connections should be made on both sites to complete the solution.

Table 22) Cisco UCS 5108 cabling.

Local Device	Local Port	Remote Device	Remote Port
Cisco UCS 5108	Fabric Port 1/1	Cisco UCS 6248UP #1	Port 31
Cisco UCS 5108	Fabric Port 1/2	Cisco UCS 6248UP #1	Port 32
Cisco UCS 5108	Fabric Port 2/1	Cisco UCS 6248UP #2	Port 31
Cisco UCS 5108	Fabric Port 2/2	Cisco UCS 6248UP #2	Port 32

Table 23) Cisco UCS 6248UP cabling.

Local Device	Local Port	Remote Device	Remote Port
Cisco UCS 6248UP #1	Port 19	Cisco Nexus 9372 #1	Eth 1/11
Cisco UCS 6248UP #1	Port 20	Cisco Nexus 9372 #2	Eth 1/11
Cisco UCS 6248UP #1	Port 31	Cisco UCS 5108	Fabric port 1/1
Cisco UCS 6248UP #1	Port 32	Cisco UCS 5108	Fabric port 1/2
Cisco UCS 6248UP #2	Port 19	Cisco Nexus 9372 #1	Eth 1/12
Cisco UCS 6248UP #2	Port 20	Cisco Nexus 9372 #2	Eth 1/12
Cisco UCS 6248UP #2	Port 31	Cisco UCS 5108	Fabric port 2/1
Cisco UCS 6248UP #2	Port 32	Cisco UCS 5108	Fabric port 2/2

Table 24) Cisco Nexus 9372 cabling.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 9372 #1	Eth 1/1	NetApp AFF8040 #1	e0g
Cisco Nexus 9372 #1	Eth 1/2	NetApp AFF8040 #2	e0g
Cisco Nexus 9372 #1	Eth 1/11	Cisco UCS 6248UP #1	Port 19
Cisco Nexus 9372 #1	Eth 1/12	Cisco UCS 6248UP #2	Port 19
Cisco Nexus 9372 #1	Eth 1/47	Cisco Nexus 7004	Eth 3/1
Cisco Nexus 9372 #1	Eth 1/48	Cisco Nexus 7004	Eth 3/2
Cisco Nexus 9372 #1	Eth 1/49	Cisco Nexus 9372 #2	Eth 1/49
Cisco Nexus 9372 #1	Eth 1/50	Cisco Nexus 9372 #2	Eth 1/50
Cisco Nexus 9372 #2	Eth 1/1	NetApp AFF8040 #1	e0h
Cisco Nexus 9372 #2	Eth 1/2	NetApp AFF8040 #2	e0h
Cisco Nexus 9372 #2	Eth 1/11	Cisco UCS 6248UP #1	Port 20
Cisco Nexus 9372 #2	Eth 1/12	Cisco UCS 6248UP #2	Port 20
Cisco Nexus 9372 #2	Eth 1/47	Cisco Nexus 7004	Eth 3/3
Cisco Nexus 9372 #2	Eth 1/48	Cisco Nexus 7004	Eth 3/4
Cisco Nexus 9372 #2	Eth 1/49	Cisco Nexus 9372 #2	Eth 1/49

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 9372 #2	Eth 1/50	Cisco Nexus 9372 #2	Eth 1/50

Table 25) NetApp AFF8040 cabling.

Local Device	Local Port	Remote Device	Remote Port
NetApp AFF8040 #1	1a	Cisco MDS 9396s #1	fc 1/1
NetApp AFF8040 #1	1b	Cisco MDS 9396s #2	fc 1/1
NetApp AFF8040 #1	0e	Cisco MDS 9396s #1	fc 1/2
NetApp AFF8040 #1	0f	Cisco MDS 9396s #2	fc 1/2
NetApp AFF8040 #1	e0g	Cisco Nexus 9372 #1	Eth 1/1
NetApp AFF8040 #1	e0h	Cisco Nexus 9372 #2	Eth 1/1
NetApp AFF8040 #1	3c	Cisco MDS 9396s #1	fc 1/3
NetApp AFF8040 #1	3d	Cisco MDS 9396s #2	fc 1/3
NetApp AFF8040 #2	1a	Cisco MDS 9396s #1	fc 1/4
NetApp AFF8040 #2	1b	Cisco MDS 9396s #2	fc 1/4
NetApp AFF8040 #2	0e	Cisco MDS 9396s #1	fc 1/5
NetApp AFF8040 #2	0f	Cisco MDS 9396s #2	fc 1/5
NetApp AFF8040 #2	e0g	Cisco Nexus 9372 #1	Eth 1/2
NetApp AFF8040 #2	e0h	Cisco Nexus 9372 #2	Eth 1/2
NetApp AFF8040 #2	3c	Cisco MDS 9396s #1	fc 1/6
NetApp AFF8040 #2	3d	Cisco MDS 9396s #2	fc 1/6

Table 26) Cisco MDS 9396s cabling.

Local Device	Local Port	Remote Device	Remote Port
Cisco MDS 9396s #1	fc 1/1	NetApp AFF8040 #1	1a
Cisco MDS 9396s #1	fc 1/2	NetApp AFF8040 #1	0e
Cisco MDS 9396s #1	fc 1/3	NetApp AFF8040 #1	3c
Cisco MDS 9396s #1	fc 1/4	NetApp AFF8040 #2	1a
Cisco MDS 9396s #1	fc 1/5	NetApp AFF8040 #2	0e

Local Device	Local Port	Remote Device	Remote Port
Cisco MDS 9396s #1	fc 1/6	NetApp AFF8040 #2	3c
Cisco MDS 9396s #1	fc 1/7	ATTO FibreBridge 7500N #1	fc1
Cisco MDS 9396s #1	fc 1/8	ATTO FibreBridge 7500N #2	fc1
Cisco MDS 9396s #1	fc1/48	Site B - Cisco MDS 9396s #1	fc1/48
Cisco MDS 9396s #1	fc1/96	Site B - Cisco MDS 9396s #1	fc1/96
Cisco MDS 9396s #2	fc 1/1	NetApp AFF8040 #1	1b
Cisco MDS 9396s #2	fc 1/2	NetApp AFF8040 #1	0f
Cisco MDS 9396s #2	fc 1/3	NetApp AFF8040 #1	3d
Cisco MDS 9396s #2	fc 1/4	NetApp AFF8040 #2	1b
Cisco MDS 9396s #2	fc 1/5	NetApp AFF8040 #2	0f
Cisco MDS 9396s #2	fc 1/6	NetApp AFF8040 #2	3d
Cisco MDS 9396s #2	fc1/48	Site B - Cisco MDS 9396s #2	fc1/48
Cisco MDS 9396s #2	fc1/96	Site B - Cisco MDS 9396s #2	fc1/96
Cisco MDS 9396s #2	fc 1/7	ATTO FibreBridge 7500N #1	fc2
Cisco MDS 9396s #2	fc 1/8	ATTO FibreBridge 7500N #2	fc2

Table 27) ATTO FibreBridge 7500N cabling.

Local Device	Local Port	Remote Device	Remote Port
ATTO FibreBridge 7500N #1	SAS port A	DS2246 #1	Top
ATTO FibreBridge 7500N #1	SAS port B	DS2246 #2	Top
ATTO FibreBridge 7500N #2	SAS port A	DS2246 #1	Bottom
ATTO FibreBridge 7500N #2	SAS port B	DS2246 #2	Bottom

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

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS.

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

Copyright Information

Copyright © 1994–2016 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.

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

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. NVA-0030-DEPLOY