



NetApp Verified Architecture

FlexPod Datacenter and Red Hat Enterprise Linux with Security Enhanced Linux

NVA Deployment

Jessica Sterling, NetApp

April 2015 | NVA-0014-DEPLOY | Version 1.0



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

TABLE OF CONTENTS

1	Program Summary	4
2	Target Audience	4
3	Solution Overview	4
4	Primary Use Case	5
5	Common Criteria Certification	5
6	Security Enhanced Linux	6
7	Technology Requirements	6
7.1	Hardware Requirements	6
7.2	Software Requirements	14
8	Configuration Guidelines	14
9	Deployment Procedures	20
9.1	FAS80XX Series Controller.....	20
9.2	Configure Cisco NX5596 Cluster Network Switch.....	21
9.3	Configure Clustered Data ONTAP 8.2.1	24
9.4	Configure Server: FlexPod Cisco UCS Base	42
9.5	Configure FlexPod Cisco UCS FCoE on Clustered Data ONTAP	43
9.6	Configure Storage Networking: FlexPod Cisco Nexus Base.....	100
9.7	Configure FlexPod Cisco Nexus FCoE Storage on Clustered Data ONTAP.....	102
9.8	Set Up Storage: Clustered Data ONTAP SAN Boot.....	109
9.9	Install Linux.....	110
10	Optional Security Enhanced Linux Procedures	111
10.1	Configure Booleans	111
10.2	Configure File Contexts	111
10.3	SELinux Policy	112
11	Solution Verification	113
12	Conclusion	114

Authors and Contributors	114
References	114
Version History	115

LIST OF TABLES

Table 1) Cisco Nexus 5548UP A cabling information.	9
Table 2) Cisco Nexus 5548UP B cabling information.	9
Table 3) Cisco Nexus 5596UP A cluster interconnect cabling information.	10
Table 4) Cisco Nexus 5596UP B cluster interconnect cabling information.	10
Table 5) NetApp controller A cabling information.	11
Table 6) NetApp controller B cabling information.	11
Table 7) Cisco UCS Fabric Interconnect A cabling information.	12
Table 8) Cisco UCS Fabric Interconnect B cabling information.	12
Table 9) Cisco UCS C-Series 1.	13
Table 10) Cisco UCS C-Series 2.	13
Table 11) Cisco UCS C-Series 3.	13
Table 12) Cisco UCS C-Series 4.	13
Table 13) FAS8040 card layout.	13
Table 14) Cisco C220M3 card layout for single-wire management.	13
Table 15) Software revisions.	14
Table 16) Necessary VLANs.	15
Table 17) Necessary VSANs.	15
Table 18) Configuration variables.	16
Table 19) FAS80XX series controller prerequisites.	20
Table 20) Cisco Nexus 5596 cluster network switch configuration prerequisites.	21
Table 21) Clustered Data ONTAP software installation prerequisites.	24
Table 22) Cluster create in clustered Data ONTAP prerequisites.	27
Table 23) Cluster join in clustered Data ONTAP prerequisites.	30
Table 24) FCP LIFs for FC WWPNs.	100
Table 25) vHBA WWPNs for Fabric A and Fabric B.	100
Table 26) FlexPod Cisco Nexus base prerequisites.	100

LIST OF FIGURES

Figure 1) FlexPod components.	5
Figure 2) FlexPod cabling diagram in clustered Data ONTAP.	8

1 Program Summary

The current industry trend in data center design is toward shared infrastructures. By using prevalidated, secure IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be deployed quickly, agilely, and at reduced costs. Cisco and NetApp have partnered to deliver FlexPod[®], which uses best-of-class storage, server, and network components that serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

2 Target Audience

This document describes the architecture and deployment procedures of an infrastructure composed of NetApp[®], Cisco, and Red Hat components that uses Fibre Channel over Ethernet (FCoE) based storage serving NAS and SAN protocols. The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy the core FlexPod architecture with NetApp clustered Data ONTAP[®] and Red Hat Enterprise Linux.

3 Solution Overview

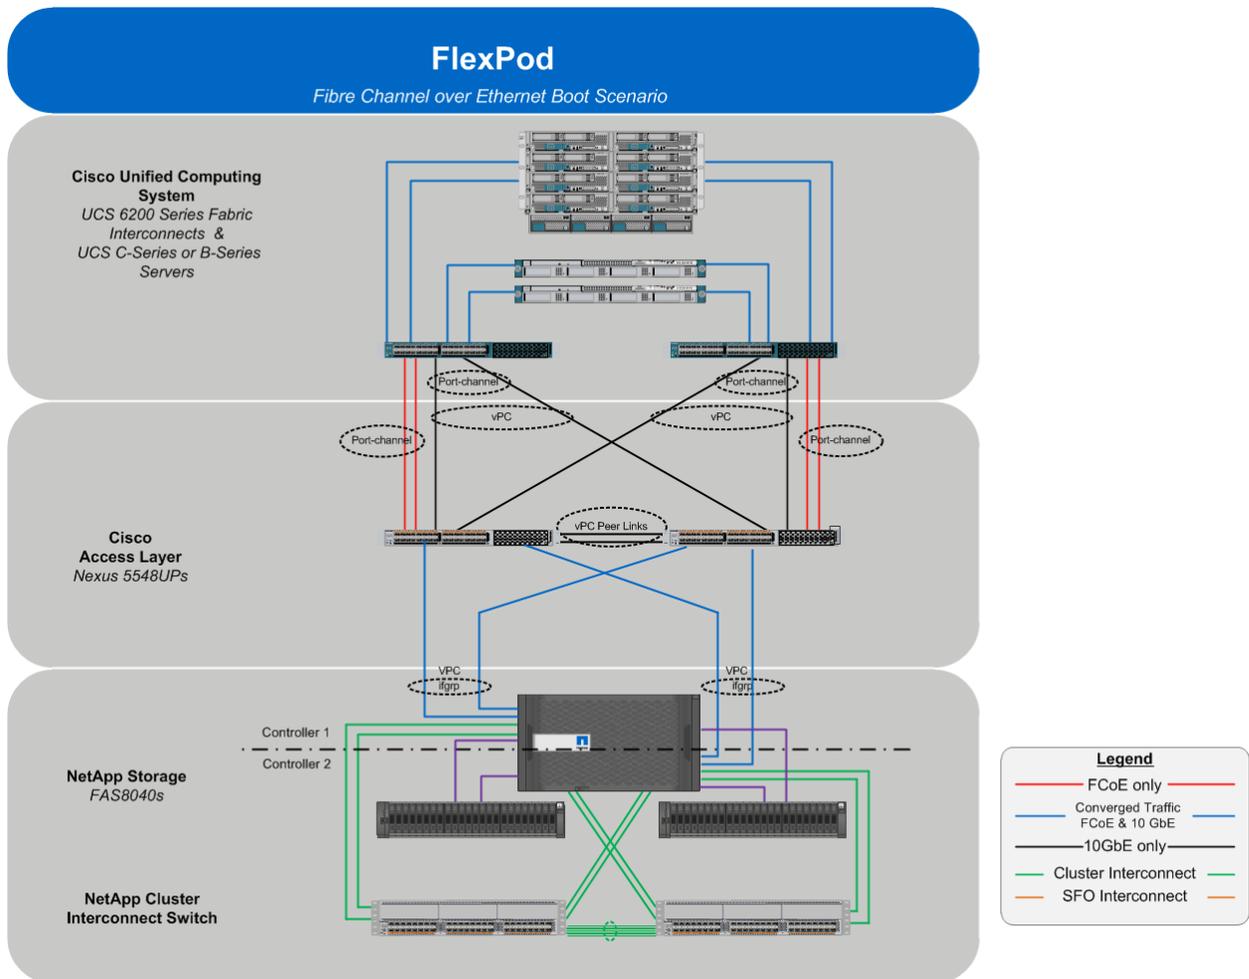
The FlexPod architecture is highly modular or “podlike.” Although each customer’s FlexPod unit varies in its exact configuration, after a FlexPod unit is built, it can easily be scaled as requirements and demands change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units).

Specifically, FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and nonvirtualized solutions. FlexPod includes NetApp storage, NetApp Data ONTAP, Cisco networking, and the Cisco Unified Computing System (Cisco UCS) in a single package. The design is flexible enough that networking, computing, and storage can fit in one data center rack or be deployed according to a customer’s data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or “flex” the environment to suit a customer’s requirements. This is why the reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an FCoE-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it is a wire-once architecture.

Figure 1 illustrates the FlexPod components and the network connections for a configuration with FCoE-based storage. This design uses the Cisco Nexus 5548UP, and Cisco UCS C-Series and B-Series with the Cisco UCS virtual interface card (VIC) and the NetApp FAS family of storage controllers connected in a highly available design that uses Cisco Virtual PortChannels (vPCs). This infrastructure is deployed to provide FCoE-booted hosts with block-level access to shared storage datastores. The reference architecture reinforces the wire-once strategy, because when additional storage is added to the architecture—be it Fibre Channel (FC), FCoE, or 10GbE—no recabling is required from the hosts to the Cisco UCS fabric interconnect.

Figure 1) FlexPod components.



This document guides you through the low-level steps for deploying the base architecture, as shown in Figure 1. These procedures cover everything from physical cabling to compute and storage configuration to installing a secure Red Hat Enterprise Linux deployment.

4 Primary Use Case

This solution is for customers who intend to leverage the core FlexPod architecture with Red Hat Enterprise Linux without a hypervisor. The end result is an infrastructure deployment with FCoE-booted bare-metal hosts with block-level access to shared storage databases.

5 Common Criteria Certification

NetApp understands the importance of security. "Trust but verify" is the foundation for NetApp's position as the [#1 provider of data storage and management to the U.S. federal government](#). Corporations and agencies in the energy, financial, healthcare, and government sectors trust NetApp because of our longstanding commitment to security certifications and verified security capabilities.

In [2005](#), NetApp became the first storage provider to achieve Common Criteria (ISO/IEC 15408) certification for Data ONTAP, an industry-leading storage operating system. The recent certifications of [Data ONTAP 8.2.1 \(7-Mode\)](#) and [clustered Data ONTAP 8.2.1](#) reflect NetApp's continued commitment to

the security principles established by the internationally recognized Common Criteria standard (ISO/IEC 15408).

FlexPod Datacenter with Red Hat Enterprise Linux 6.6 embodies the architectural best practice that certified components are the foundation for secure solutions. The following components of this FlexPod architecture are Common Criteria certified:

- Compute:
 - [Cisco UCS 5100 B/C-Series server](#) EAL 4+
 - [Cisco UCS 6248 fabric interconnect](#) EAL 4+
- Network:
 - [Cisco Nexus 5548UP switch](#) EAL 4+
- Storage:
 - NetApp FAS8040 running [clustered Data ONTAP 8.2.1](#) EAL 2+

6 Security Enhanced Linux

Security Enhanced Linux (SELinux) is a kernel-level mandatory access control mechanism developed by the National Security Agency that allows administrators to enforce rule-based controls on files and processes on a Linux system. Strict policies define which processes have access to which actions and files. These include everything from a user running a command in a terminal window to a running process trying to open or write to a file.

One benefit of SELinux is that, because the privileges associated with executing processes are limited, the scope of potential damage from software vulnerabilities is greatly reduced. This is an added layer of security for the Linux system on top of existing firewalls and discretionary access control policies, such as file permissions. It is not designed to take the place of other security measures and best practices.

SELinux is controlled by loadable, configurable policies that contain information about processes and files. Each process and file is labeled with an SELinux context that contains additional information, such as the SELinux user, role, type, and level. The default rule for most operations is to deny the action unless it has been explicitly allowed within the current SELinux policy. When an action is taking place, the operation is intercepted in the Linux kernel by SELinux. If a policy rule allows the operation, then the process continues. However, if there is no rule for the operation or if it is denied, then the operation is blocked, and an error is sent back to the original process. This action is also logged.

By default in Red Hat Enterprise Linux 6.6, SELinux is enabled and set to enforcing. SELinux can also be enabled and set to permissive, which causes SELinux not to deny access, but it continues to log actions that would have been denied had SELinux been running in enforcing mode. The permissive state is useful for testing applications in a secure staging environment before the applications are put into production. If SELinux reports that it is preventing a certain action from running, this information is logged in `/var/log/audit/audit.log` or `/var/log/messages`, if the audit service is disabled.

When making a decision, SELinux caches this information in the access vector cache (AVC). This allows SELinux to check its policy rules less as more decisions are made, decreasing its overall performance effect on the system.

7 Technology Requirements

7.1 Hardware Requirements

The reference configuration includes:

- Two Cisco Nexus 5548UP switches

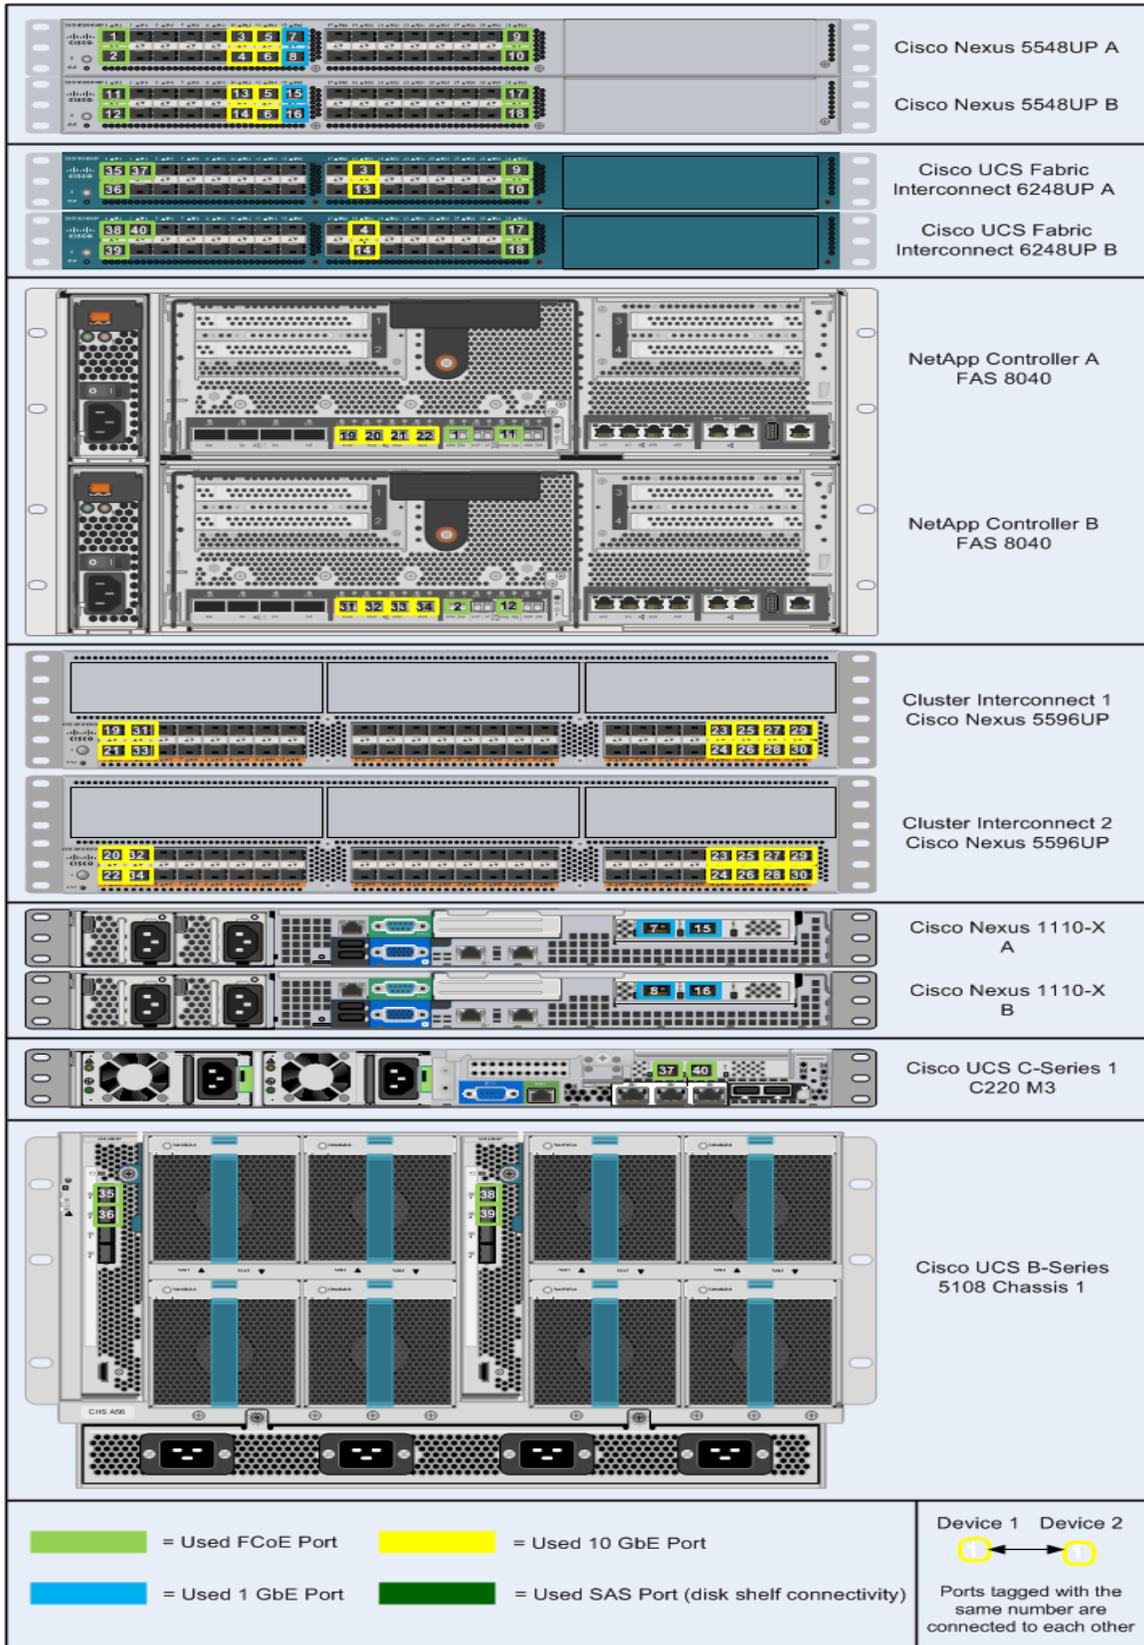
- Two Cisco UCS 6248UP fabric interconnects
- Support for 16 Cisco UCS C-Series servers without any additional networking components
- Support for 8 Cisco UCS B-Series servers without any additional blade server chassis
- Support for 160 Cisco UCS C-Series and B-Series servers
- One NetApp FAS8040 (HA pair) running clustered Data ONTAP

Storage is provided by a NetApp FAS8040 (HA configuration in one chassis) unified storage system operating in clustered Data ONTAP. All system and network links feature redundancy, providing end-to-end high availability (HA). Although this is the base design, each of the components can be scaled flexibly to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capacity and throughput, and special hardware or software features can be added to introduce new capabilities.

FlexPod Cabling on Clustered Data ONTAP

Figure 2 depicts the cabling diagram for a FlexPod configuration using clustered Data ONTAP.

Figure 2) FlexPod cabling diagram in clustered Data ONTAP.



The information provided in Table 1 through Table 9 corresponds to each connection shown in Figure 2.

Table 1) Cisco Nexus 5548UP A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5548UP A	Eth1/1	10GbE	NetApp controller A	e0e	1
	Eth1/2	10GbE	NetApp controller B	e0e	2
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/19	3
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/19	4
	Eth1/13	10GbE	Cisco Nexus 5548UP B	Eth1/13	5
	Eth1/14	10GbE	Cisco Nexus 5548UP B	Eth1/14	6
	Eth1/15	10GbE	Cisco Nexus 1110-X A	Port 1	7
	Eth1/16	10GbE	Cisco Nexus 1110-X B	Port 1	8
	Eth1/31	10GbE	Cisco UCS fabric interconnect A	Eth1/31	9
	Eth1/32	10GbE	Cisco UCS fabric interconnect A	Eth1/32	10
	MGMT0	GbE	GbE management switch	Any	

Note: For devices requiring GbE connectivity, use the GbE copper SFP+s (GLC-T=).

Table 2) Cisco Nexus 5548UP B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5548UP B	Eth1/1	10GbE	NetApp controller A	e0g	11
	Eth1/2	10GbE	NetApp controller B	e0g	12
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/20	13
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/20	14
	Eth1/13	10GbE	Cisco Nexus 5548UP A	Eth1/13	5
	Eth1/14	10GbE	Cisco Nexus 5548UP A	Eth1/14	6
	Eth1/15	10GbE	Cisco Nexus 1110-X A	Port 2	15
	Eth1/16	10GbE	Cisco Nexus 1110-X B	Port 2	16
	Eth1/31	10GbE	Cisco UCS fabric interconnect B	Eth1/31	17
	Eth1/32	10GbE	Cisco UCS fabric interconnect B	Eth1/32	18

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	MGMT0	GbE	GbE management switch	Any	

Note: For devices requiring GbE connectivity, use the GbE copper SFP+s (GLC-T=).

Table 3) Cisco Nexus 5596UP A cluster interconnect cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5596UP A	Eth1/1	10GbE	NetApp controller A	e0a	19
	Eth1/2	10GbE	NetApp controller A	e0c	21
	Eth 1/3	10GbE	NetApp controller B	e0a	31
	Eth 1/4	10GbE	NetApp controller B	e0c	33
	Eth1/41	10GbE	Cisco Nexus 5596UP B	Eth1/41	23
	Eth1/42	10GbE	Cisco Nexus 5596UP B	Eth1/42	24
	Eth1/43	10GbE	Cisco Nexus 5596UP B	Eth1/43	25
	Eth1/44	10GbE	Cisco Nexus 5596UP B	Eth1/44	26
	Eth1/45	10GbE	Cisco Nexus 5596UP B	Eth1/45	27
	Eth1/46	10GbE	Cisco Nexus 5596UP B	Eth1/46	28
	Eth1/47	10GbE	Cisco Nexus 5596UP B	Eth1/47	29
	Eth1/48	10GbE	Cisco Nexus 5596UP B	Eth1/48	30
		MGMT0	GbE	GbE management switch	Any

Table 4) Cisco Nexus 5596UP B cluster interconnect cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5596UP B	Eth1/1	10GbE	NetApp controller A	e0b	20
	Eth1/2	10GbE	NetApp controller A	e0d	22
	Eth 1/3	10GbE	NetApp controller B	e0b	32
	Eth 1/4	10GbE	NetApp controller B	e0d	34
	Eth1/41	10GbE	Cisco Nexus 5596UP A	Eth1/41	23
	Eth1/42	10GbE	Cisco Nexus 5596UP A	Eth1/42	24
	Eth1/43	10GbE	Cisco Nexus 5596UP A	Eth1/43	25
	Eth1/44	10GbE	Cisco Nexus 5596UP A	Eth1/44	26
	Eth1/45	10GbE	Cisco Nexus 5596UP A	Eth1/45	27
	Eth1/46	10GbE	Cisco Nexus 5596UP A	Eth1/46	28

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	Eth1/47	10GbE	Cisco Nexus 5596UP A	Eth1/47	29
	Eth1/48	10GbE	Cisco Nexus 5596UP A	Eth1/48	30
	MGMT0	GbE	GbE management switch	Any	

Table 5) NetApp controller A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp controller A	e0M	1000MbE	1000MbE management switch	Any	
	e0i	GbE	GbE management switch	Any	
	e0P	GbE	SAS shelves	ACP port	
	e0a	10GbE	Cisco Nexus 5596UP A	Eth1/1	19
	e0b	10GbE	Cisco Nexus 5596UP B	Eth1/1	20
	e0c	10GbE	Cisco Nexus 5596UP A	Eth1/2	21
	e0d	10GbE	Cisco Nexus 5596UP B	Eth1/2	22
	e0e	10GbE	Cisco Nexus 5548UP A	Eth1/1	1
	e0g	10GbE	Cisco Nexus 5548UP B	Eth1/1	11

Note: When the term e0M is used, the physical Ethernet port to which Table 5 is referring is the port indicated by a wrench icon on the rear of the chassis.

Table 6) NetApp controller B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp controller B	e0M	1000MbE	1000MbE management switch	Any	
	e0i	GbE	GbE management switch	Any	
	e0P	GbE	SAS shelves	ACP port	
	e0a	10GbE	Cisco Nexus 5596UP A	Eth1/1	30
	e0b	10GbE	Cisco Nexus 5596UP B	Eth1/1	31
	e0c	10GbE	Cisco Nexus 5596UP A	Eth1/2	32
	e0d	10GbE	Cisco Nexus 5596UP B	Eth1/2	33
	e0e	10GbE	Cisco Nexus 5548UP A	Eth1/2	2
	e0g	10GbE	Cisco Nexus 5548UP B	Eth1/2	12

Note: When the term e0M is used, the physical Ethernet port to which Table 6 and Table 7 are referring is the port indicated by a wrench icon on the rear of the chassis.

Table 7) Cisco UCS Fabric Interconnect A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS fabric interconnect A	Eth1/19	10GbE	Cisco Nexus 5548UP A	Eth1/11	3
	Eth1/20	10GbE	Cisco Nexus 5548UP B	Eth1/11	13
	Eth1/1	10GbE	Cisco UCS chassis 1 IOM A	Port 1	35
	Eth1/2	10GbE	Cisco UCS chassis 1 IOM A	Port 2	36
	Eth 1/3	10GbE	Cisco UCS C-Series 1	Port 0	37
	Eth 1/4	10GbE	Cisco UCS C-Series 2	Port 0	
	Eth1/31	10GbE	Cisco Nexus 5548UP A	Eth1/31	9
	Eth1/32	10GbE	Cisco Nexus 5548UP A	Eth1/32	10
	MGMT0	GbE	GbE management switch	Any	
	L1	GbE	Cisco UCS fabric interconnect B	L1	
	L2	GbE	Cisco UCS fabric interconnect B	L2	

Table 8) Cisco UCS Fabric Interconnect B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS fabric interconnect B	Eth1/19	10GbE	Cisco Nexus 5548UP A	Eth1/12	4
	Eth1/20	10GbE	Cisco Nexus 5548UP B	Eth1/12	14
	Eth1/1	10GbE	Cisco UCS chassis 1 IOM B	Port 1	38
	Eth1/2	10GbE	Cisco UCS chassis 1 IOM B	Port 2	39
	Eth 1/3	10GbE	Cisco UCS C-Series 1	Port 1	40
	Eth 1/4	10GbE	Cisco UCS C-Series 2	Port 1	
	Eth1/31	10GbE	Cisco Nexus 5548UP B	Eth1/31	17
	Eth1/32	10GbE	Cisco Nexus 5548UP B	Eth1/32	18
	MGMT0	GbE	GbE management switch	Any	
	L1	GbE	Cisco UCS fabric interconnect A	L1	
	L2	GbE	Cisco UCS fabric interconnect A	L2	

Table 9) Cisco UCS C-Series 1.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series 1	Port 0	10GbE	Fabric interconnect A	Eth 1/3	37
	Port 1	10GbE	Fabric interconnect B	Eth 1/3	40

Table 10) Cisco UCS C-Series 2.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series 2	Port 0	10GbE	Fabric interconnect A	Eth 1/4	
	Port 1	10GbE	Fabric interconnect B	Eth 1/4	

Table 11) Cisco UCS C-Series 3.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series 3	Port 0	10GbE	Fabric interconnect A	Eth 1/5	
	Port 1	10GbE	Fabric interconnect B	Eth 1/5	

Table 12) Cisco UCS C-Series 4.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series 4	Port 0	10GbE	Fabric interconnect A	Eth 1/6	
	Port 1	10GbE	Fabric interconnect B	Eth 1/6	

Table 13) FAS8040 card layout.

Slot	Part Number	Description
1	X1973A-R6	NetApp Flash Cache™ 2 – 512GB

Table 14) Cisco C220M3 card layout for single-wire management.

Slot	Part Number	Description
1	Cisco UCS VIC1225	CNA 2-port 10GbE (ports 0 and 1)

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported is available at the [NetApp Support](#) site.

When using SAS disk shelves with NetApp storage controllers, refer to the [SAS Disk Shelves Universal SAS and ACP Cabling Guide](#) for proper cabling guidelines.

NetApp Hardware Universe

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

1. Confirm that the hardware and software components are supported with the version of Data ONTAP that you plan to install by checking the [NetApp Hardware Universe](#) at the [NetApp Support](#) site.
2. Access the [Hardware Universe](#) application to view the System Configuration guides. Click the Controllers tab to view the compatibility between Data ONTAP software versions and NetApp storage appliances with the desired specifications.
3. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

7.2 Software Requirements

It is important to note the software versions used in this document. Table 15 lists the software revisions used throughout this document.

Table 15) Software revisions.

Layer	Software	Version or Release	Details
Compute	Cisco UCS fabric interconnect	2.2(2c)	Embedded management
	Cisco UCS C 220 M3	2.2(2c)	Software bundle release
	Cisco UCS B 200 M3	2.2(2c)	Software bundle release
	Cisco eNIC	2.1.2.42	Ethernet driver for Cisco VIC
	Cisco fNIC	1.6.0.5	FCoE driver for Cisco VIC
Network	Cisco Nexus fabric switch	7.0(1)N1(1)	Operating system version
Storage	NetApp FAS8040	Clustered Data ONTAP 8.2.1	Operating system version
Software	Cisco UCS host	Red Hat Enterprise Linux 6.x	Operating system version
	NetApp OnCommand®	6.1	VM (1 each): OnCommand
	Cisco Nexus 1110-x	5.2(1)SP1(7.1)	Virtual services appliance
	Cisco Nexus 1000v	4.2(1)SV2(2.2) (Advanced Edition)	Virtual services blade within the 1110-x
	Cisco UCS Central	1.1(2a)	Manager of multiple Cisco UCS domains

8 Configuration Guidelines

This document provides details for configuring a FlexPod unit with clustered Data ONTAP storage. Therefore, references are made to which component is being configured with each step, either 01 or 02. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning a Cisco UCS host, named `RHEL-Server-01`. Finally, to indicate that you should include information pertinent to your environment in a given step, `<text>` appears as part of the command structure. Refer to the following example for the `network port vlan create` command:

```

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 a0a-<vlan id>

```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 16 describes the VLANs necessary for deployment as outlined in this guide. Table 17 lists the virtual storage area networks (VSANs) necessary for deployment as outlined in this guide.

Table 18 lists the configuration variables that are used throughout this document. Table 18 can be completed based on the specific site variables and used to complete the document configuration steps.

Note: The cluster management and node management interfaces will be on the out-of-band management VLAN. Confirm that there is a layer 3 route between the out-of band and in-band management VLANs.

Table 16) Necessary VLANs.

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Mgmt in band	VLAN for in-band management interfaces	3175
Mgmt out of band	VLAN for out-of-band management interfaces	3172
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for NFS traffic	3170
FCoE-A	VLAN for FCoE traffic for Fabric A	101
FCoE-B	VLAN for FCoE traffic for Fabric B	102
vMotion	VLAN designated for the movement of VMs from one physical host to another	3173
VM traffic	VLAN for VM application traffic	3174
Packet control	VLAN for packet control traffic (Cisco Nexus 1000v)	3176

Table 17) Necessary VSANs.

VSAN Name	VSAN Purpose	ID Used in Validating This Document
VSAN A	VSAN for Fabric A traffic. ID matches FCoE-A VLAN	101
VSAN B	VSAN for Fabric B traffic. ID matches FCoE-B VLAN	102

Table 18) Configuration variables.

Variable	Description	Customer Implementation Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01	
<<var_node01_mgmt_mask>>	Out-of-band management network netmask	
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_url_boot_software>>	Data ONTAP 8.2.1 URL; format: http://	
<<var_#_of_disks>>	Number of disks to assign to each storage controller	
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02	
<<var_node02_mgmt_mask>>	Out-of-band management network netmask	
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_clustername>>	Storage cluster host name	
<<var_cluster_base_license_key>>	Cluster base license key	
<<var_nfs_license>>	License key for NFS protocol	
<<var_fcp_license>>	License key for FCP	
<<var_password>>	Global default administrative password	
<<var_clustermgmt_ip>>	Out-of-band management IP for the storage cluster	
<<var_clustermgmt_mask>>	Out-of-band management network netmask	
<<var_clustermgmt_gateway>>	Out-of-band management network default gateway	
<<var_dns_domain_name>>	DNS domain name	
<<var_nameserver_ip>>	DNS server IP(s)	
<<var_node_location>>	Node location string for each node	
<<var_node01>>	Cluster node 01 host name	
<<var_node02>>	Cluster node 02 host name	
<<var_num_disks>>	Number of disks to assign to each storage data aggregate	
<<var_node01_sp_ip>>	Out-of-band cluster node 01 service processor management IP	
<<var_node01_sp_mask>>	Out-of-band management network netmask	
<<var_node01_sp_gateway>>	Out-of-band management network default gateway	
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP	
<<var_node02_sp_mask>>	Out-of-band management network netmask	

Variable	Description	Customer Implementation Value
<<var_node02_sp_gateway>	Out-of-band management network default gateway	
<<var_timezone>>	FlexPod time zone (for example, America/New_York)	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_snmp_contact>>	Administrator e-mail address	
<<var_snmp_location>>	Cluster location string	
<<var_oncommand_server_fqdn>>	OnCommand virtual machine fully qualified domain name (FQDN)	
<<var_oncommand_server_ip>>	OnCommand virtual machine management IP Address	
<<var_oncommand_server_netmask>>	Out-of-band management network netmask	
<<var_oncommand_server_gateway>>	Out-of-band management network default gateway	
<<var_ucs_central_ip>>	Cisco UCS central management IP	
<<var_ucs_central_netmask>>	Out-of-band management network netmask	
<<var_ucs_central_gateway>>	Out-of-band management network default gateway	
<<var_ucs_central_hostname>>	Cisco UCS central fully qualified domain name (FQDN)	
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name	
<<var_mailhost>>	Mail server host name	
<<var_storage_admin_email>>	Administrator e-mail address	
<<var_security_cert_vserver_common_name>>	Infrastructure storage virtual machine (SVM) FQDN	
<<var_security_cert_vserver_authority>>	Infrastructure SVM Security Certificate Authority	
<<var_security_cert_vserver_serial_no>>	Infrastructure SVM security certificate serial number	
<<var_country_code>>	Two-letter country code	
<<var_state>>	State or province name	
<<var_city>>	City name	
<<var_org>>	Organization or company name	
<<var_unit>>	Organizational unit name	
<<var_security_cert_cluster_common_name>>	Storage cluster FQDN	

Variable	Description	Customer Implementation Value
<<var_security_cert_cluster_authority>>	Storage cluster security certificate authority	
<<var_security_cert_cluster_serial_no>>	Storage cluster security certificate serial number	
<<var_security_cert_node01_common_name>>	Cluster node 01 FQDN	
<<var_security_cert_node01_authority>>	Cluster node 01 security certificate authority	
<<var_security_cert_node01_serial_no>>	Cluster node 01 security certificate serial number	
<<var_security_cert_node02_common_name>>	Cluster node 02 FQDN	
<<var_security_cert_node02_authority>>	Cluster node 02 security certificate authority	
<<var_security_cert_node02_serial_no>>	Cluster node 02 security certificate serial number	
<<var_node01_nfs_lif_infra_swap_ip>>	Cluster node 01 NFS VLAN for infra_swap IP address	
<<var_node01_nfs_lif_infra_swap_mask>>	NFS VLAN for infra_swap netmask	
<<var_node02_nfs_lif_infra_datastore_1_ip>>	Cluster node 02 NFS VLAN for infra_datastore_1 IP address	
<<var_node02_nfs_lif_infra_datastore_1_mask>>	NFS VLAN for infra_datastore_1 netmask	
<<var_nexus_A_hostname>>	Cisco Nexus A host name	
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_B_hostname>>	Cisco Nexus B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_ib-mgmt_vlan_id>>	In-band management network VLAN ID	
<<var_native_vlan_id>>	Native VLAN ID	
<<var_oob-mgmt_vlan_id>>	Out-of-band management network VLAN ID	
<<var_nfs_vlan_id>>	NFS VLAN ID	

Variable	Description	Customer Implementation Value
<<var_pkt-ctrl_vlan_id>>	Cisco Nexus 1000v packet control VLAN ID	
<<var_vm-traffic_vlan_id>>	VM traffic VLAN ID	
<<var_nexus_vpc_domain_id>>	Unique Cisco Nexus switch VPC domain ID	
<<var_nexus_1110x-1>>	Cisco Nexus 1110X-1 host name	
<<var_nexus_1110x-2>>	Cisco Nexus 1110X-2 host name	
<<var_fabric_a_fcoe_vlan_id>>	Fabric A FCoE VLAN ID	
<<var_vsan_a_id>>	Fabric A VSAN ID	
<<var_fabric_b_fcoe_vlan_id>>	Fabric B FCoE VLAN ID	
<<var_vsan_b_id>>	Fabric B VSAN ID	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address	
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address	
<<var_cimc_ip>>	Out-of-band management IP for each Cisco Nexus 1110-X CIMC	
<<var_cimc_mask>>	Out-of-band management network netmask	
<<var_cimc_gateway>>	Out-of-band management network default gateway	
<<var_1110x_domain_id>>	Unique Cisco Nexus 110-X domain ID	
<<var_1110x_vsa>>	Virtual storage appliance (VSA) host name	
<<var_1110x_vsa_ip>>	In-band VSA management IP address	
<<var_1110x_vsa_mask>>	In-band management network netmask	
<<var_1110x_vsa_gateway>>	In-band management network default gateway	
<<var_vsm_domain_id>>	Unique Cisco Nexus 1000v virtual supervisor module (VSM) domain ID	
<<var_vsm_mgmt_ip>>	Cisco Nexus 1000v VSM management IP address	
<<var_vsm_mgmt_mask>>	In-band management network netmask	
<<var_vsm_mgmt_gateway>>	In-band management network default gateway	
<<var_vsm_hostname>>	Cisco Nexus 1000v VSM host name	

Variable	Description	Customer Implementation Value
<<var_nodename>>	Name of node	
<<var_node01_rootaggrname>>	Root aggregate name of node 01	
<<var_clustermgmt_port>>	Port for cluster management	
<<var_global_domain_name>>	Domain name	
<<var_dns_ip>>	IP address of the DNS server	
<<var_vsadmin_password>>	Password for VS admin account	
<<var_vserver_mgmt_ip>>	Management IP address for SVM	
<<var_vserver_mgmt_mask>>	Subnet mask for SVM	
<<var_rule_index>>	Rule index number	
<<var_ftp_server>>	IP address for FTP server	
<<var_vm_host_infra_01_A_wwpn>>	WWPN of VM-RHEL-Server-01 vHBA-A	
<<var_fcp_lif01a_wwpn>>	WWPN of FCP_LIF01a	
<<var_fcp_lif02a_wwpn>>	WWPN of FCP_LIF02a	
<<var_vm_host_infra_01_B_wwpn>>	WWPN of VM-RHEL-Server-01 vHBA-B	
<<var_fcp_lif01b_wwpn>>	WWPN of FCP_LIF01b	
<<var_fcp_lif02b_wwpn>>	WWPN of FCP_LIF02b	

9 Deployment Procedures

9.1 FAS80XX Series Controller

Table 19) FAS80XX series controller prerequisites.

Prerequisites
<p>Refer to the Site Requirements Guide for planning the physical location of the storage systems. From the downloaded guide, refer the following sections:</p> <ul style="list-style-type: none"> • Site Preparation • System Connectivity Requirements • Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements • 80xx Series Systems

Install Controllers

Follow the physical installation procedures for the controllers in the [FAS80xx series documentation](#) on the [NetApp Support](#) site.

9.2 Configure Cisco NX5596 Cluster Network Switch

Table 20) Cisco Nexus 5596 cluster network switch configuration prerequisites.

Prerequisites

- Rack and connect power to the new Cisco Nexus 5596 switches
- Provide a terminal session that connects to the switch's serial console port (9600, 8, n, 1)
- Connect the `mgmt0` port to the management network and be prepared to provide IP address information
- Obtain password for admin
- Determine switch name
- Identify SSH key type (dsa, rsa, or rsa1)
- Set up an e-mail server for Cisco Smart Call Home and IP connectivity between the switch and the e-mail server
- Provide SNMP contact information for Cisco Smart Call Home (name, phone, street address)
- Identify a CCO ID associated with an appropriate Cisco SMARTnet service contract for Cisco Smart Call Home
- Enable Cisco SMARTnet Service for the device to be registered for Cisco Smart Call home

Initial Setup of Cisco Nexus 5596 Cluster Interconnect

The first time a Cisco Nexus 5596 cluster interconnect is accessed, it runs a setup program that prompts the user to enter an IP address and other configuration information needed for the switch to communicate over the management Ethernet interface. This information is required to configure and manage the switch. If the configuration must be changed later, the setup wizard can be accessed again by running the `setup` command in EXEC mode.

To set up the Cisco Nexus 5596 cluster interconnect, complete the following steps. These steps must be completed on both the cluster interconnects.

1. Provide applicable responses to the setup prompts displayed on the Cisco Nexus 5596 cluster interconnect.

```
Abort Power On Auto Provisioning and continue with normal setup ?(yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no):yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name : <switchname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address : <ic_mgmt0_ip>
Mgmt0 IPv4 netmask : <ic_mgmt0_netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <ic_mgmt0_gw>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) : rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <ntp_server_ip>
Enter basic FC configurations (yes/no) [n]: Enter
```

2. At the end of the setup, the configuration choices are displayed. Verify the information and save the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: <n>
Use this configuration and save it? (yes/no) [y]: <y>
```

Download and Install NetApp Cluster Switch Software

When the Cisco Nexus 5596 is being used as a cluster network switch with Data ONTAP 8.2.1, it should be running NX-OS version 5.2(1)N1(1). The `show version` command from the switch command line interface shows the switch version currently running on the switch. If the currently running version is not 5.2(1)N1(1), go to the [NetApp Support](#) site and download and install NX-OS 5.2(1)N1(1) for the Cisco Nexus 5596 switch. Make sure both cluster interconnects are running NX-OS version 5.2(1)N1(1).

Download and Merge NetApp Cluster Switch Reference Configuration File

Cluster network and management network switches are shipped without the configuration files installed. These files must be downloaded to the switches during deployment. Configuration files must be downloaded when the cluster network and management network switches are first installed or after the Cisco switch software is updated or reinstalled.

After the initial setup is complete, the NetApp cluster network switch reference configuration must be transferred to the switch and merged with existing configuration. Instructions for this task and the reference configuration files for the appropriate switches are available on the [NetApp Support](#) site.

To download configuration files to a host and install them on a Cisco Nexus 5596 switch, complete the following steps on both cluster interconnects:

1. Obtain a console connection to the switch. Verify existing configuration on the switch by running the `show run` command.
2. Log in to the switch. Verify that the host recognizes the switch on the network (for example, use the ping utility).
3. Run the following command:

```
copy <transfer protocol>: bootflash: vrf management
```

4. Verify that the configuration file is downloaded.

```
***** Transfer of file Completed Successfully *****  
Copy complete, now saving to disk (please wait)...
```

5. Run the following command to view the saved configuration file.

```
dir bootflash:
```

6. Merge the configuration file into the existing `running-config`. Run the following command, where `<config file name>` is the file name for the switch type. A series of warnings regarding PortFast is displayed as each port is configured.

```
copy <config file name> running-config
```

7. Verify the success of the configuration merge by running the `show run` command and comparing its output to the contents of the configuration file (a `.txt` file) that was downloaded.
 - a. The output for both installed-base switches and new switches should be identical to the contents of the configuration file for the following items:
 - `banner` (should match the expected version)
 - Switch port descriptions such as `description Cluster Node x`
 - The new ISL algorithm `port-channel load-balance Ethernet source-dest-port`
 - b. The output for new switches should be identical to the contents of the configuration file for the following items:
 - Port channel
 - Policy map
 - System QoS

- Interface
 - Boot
 - c. The output for installed base switches should have the flow control receive and send values set to on for the following items:
 - Interface port channels 1 and 2
 - Ethernet interface 1/41 through Ethernet interface 1/48.
8. Copy the `running-config` to the `startup-config`.

```
copy running-config startup-config
```

Cisco Smart Call Home Setup

To configure Smart Call Home on a Cisco Nexus 5596 switch, complete the following steps:

1. Enter the mandatory system contact using the `snmp-server contact` command in global configuration mode. Then run the `callhome` command to enter callhome configuration mode.

```
NX-5596#config t
NX-5596(config)#snmp-server contact <sys-contact>
NX-5596(config)#callhome
```

2. Configure the mandatory contact information (phone number, e-mail address, and street address).

```
NX-5596(config-callhome)#email-contact <email-address>
NX-5596(config-callhome)#phone-contact <+1-000-000-0000>
NX-5596(config-callhome)#streetaddress <a-street-address>
```

3. Configure the mandatory e-mail server information. The server address is an IPv4 address, IPv6 address, or the domain name of a SMTP server to which Call Home will send e-mail messages. Optional port number (default=25) and VRF may be configured.

```
NX-5596(config-callhome)#transport email smtp-server <ip-address> port 25 use-vrf <vrf-name>
```

4. Set the destination profile CiscoTAC-1 e-mail address to callhome@cisco.com.

```
NX-5596(config-callhome)#destination-profile CiscoTAC-1 email-addr callhome@cisco.com
```

5. Enable periodic inventory and set the interval.

```
NX-5596(config-callhome)#periodic-inventory notification
NX-5596(config-callhome)#periodic-inventory notification interval 30
```

6. Enable `callhome`, exit, and save the configuration.

```
NX-5596(config-callhome)#enable
NX-5596(config-callhome)#end
NX-5596#copy running-config startup-config
```

7. Send a callhome inventory message to start the registration process.

```
NX-5596#callhome test inventory
trying to send test callhome inventory message
successfully sent test callhome inventory message
```

8. Watch for an e-mail from Cisco regarding the registration of the switch. Follow the instructions in the e-mail to complete the registration for Smart Call Home.

SNMP Monitoring Setup

To set up SNMP monitoring, complete the following step:

1. Configure SNMP by using the following example as a guideline. This example configures a host receiver for SNMPv1 traps and enables all link up/down traps.

```
NX-5596#config t
```

```
NX-5596(config)# snmp-server host <ip-address> traps { version 1 } <community> [udp_port
<number>]
NX-5596(config)# snmp-server enable traps link
```

9.3 Configure Clustered Data ONTAP 8.2.1

Configure Clustered Data ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [Clustered Data ONTAP 8.2 Software Setup Guide](#) to learn about the information required to configure clustered Data ONTAP. Table 21 lists the information that you will need to configure two clustered Data ONTAP nodes. You should customize the cluster detail values with the information that is applicable to your deployment.

Before running the setup script, complete the configuration worksheet from the [Clustered Data ONTAP Software Setup Guide](#) on the [NetApp Support](#) site.

Table 21) Clustered Data ONTAP software installation prerequisites.

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<<var_node01_mgmt_ip>>
Cluster node 01 netmask	<<var_node01_mgmt_mask>>
Cluster node 01 gateway	<<var_node01_mgmt_gateway>>
Cluster node 02 IP address	<<var_node02_mgmt_ip>>
Cluster node 02 netmask	<<var_node02_mgmt_mask>>
Cluster node 02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.2.1 URL	<<var_url_boot_software>>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Enable Autoboot.

```
setenv AUTOBOOT true
```

3. Allow the system to boot up.

```
boot_ontap
```

4. Press Ctrl-C when you see this message:

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

Note: If Data ONTAP 8.2.1 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.2.1 is the version being booted, then select option 8 and yes to reboot the node. Then continue with step 14.

5. To install new software, select option 7.

```
7
```

6. Answer yes to perform a nondisruptive upgrade.

y

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

e0M

8. Select yes to reboot now.

y

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

<<var_node01_mgmt_ip>> <<var_node01_mgmt_mask>> <<var_node01_mgmt_gateway>>

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

Note: This web server must be pingable.

<<var_url_boot_software>>

11. Press Enter for the user name, indicating no user name.

Enter

12. Enter yes to set the newly installed software as the default to be used for subsequent reboots.

y

13. Enter yes to reboot the node.

y

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

14. Press Ctrl-C to exit autoboot when you see this message:

Starting AUTOBOOT press Ctrl-C to abort...

15. From the loader-A prompt, enter:

printenv

Note: If `bootarg.init.boot_clustered true` is not listed, the system is not set to boot in clustered Data ONTAP.

16. If the system is not set to boot in clustered Data ONTAP, at the loader prompt, enter the following command to make sure the system boots in clustered Data ONTAP:

```
setenv bootarg.init.boot_clustered true
setenv bootarg.bsdportname e0M
```

17. At the loader-A prompt, enter:

autoboot

18. Press Ctrl- C when you see this message:

Press Ctrl-C for Boot Menu

19. Select option 4 for a clean configuration and to initialize all disks.

4

20. Answer yes to Zero disks, reset config and install a new file system.

y

21. Enter yes to erase all the data on the disks.

y

Note: The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue to node 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Enable Autoboot.

```
setenv AUTOBOOT true
```

3. Allow the system to boot up.

```
boot_ontap
```

4. Press Ctrl-C when you see this message:

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

Note: If Data ONTAP 8.2.1 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.2.1 is the version being booted, then select option 8 and *yes* to reboot the node. Then continue with step 14.

5. To install new software first, select option 7.

```
7
```

6. Answer yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

8. Select yes to reboot now.

```
y
```

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

```
<<var_node02_mgmt_ip>> <<var_node02_mgmt_mask>> <<var_node02_mgmt_gateway>>
```

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

Note: This web server must be pingable.

```
<<var_url_boot_software>>
```

11. Press Enter for the user name, indicating no user name.

```
Enter
```

12. Select yes to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

13. Select yes to reboot the node.

```
y
```

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

14. Press Ctrl-C to exit autoboot when you see this message:

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

15. From the loader-A prompt, enter:

```
printenv
```

Note: If `bootarg.init.boot_clustered true` is not listed, the system is not set to boot in clustered Data ONTAP.

16. If the system is not set to boot in clustered Data ONTAP, at the loader prompt, enter the following command to make sure the system boots in clustered Data ONTAP:

```
setenv bootarg.init.boot_clustered true
setenv bootarg.bsdportname e0M
```

17. At the loader-A prompt, enter:

```
autoboot
```

18. Press Ctrl-C when you see this message:

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

19. Select option 4 for a clean configuration and to initialize all disks.

```
4
```

20. Answer yes to Zero disks, reset config and install a new file system.

```
y
```

21. Enter yes to erase all the data on the disks.

```
y
```

Note: The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Create Cluster on Node 1

In clustered Data ONTAP, the first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01.

Table 22) Cluster create in clustered Data ONTAP prerequisites.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Clustered Data ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster node 01 IP address	<<var_node01_mgmt_ip>>
Cluster node 01 netmask	<<var_node01_mgmt_mask>>

Cluster Detail	Cluster Detail Value
Cluster node 01 gateway	<<var_node01_mgmt_gateway>>

To create a cluster on node 01, complete the following steps.

1. At the login prompt, enter `admin`.

```
admin
```

2. The Cluster Setup wizard starts running on the console.

```
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}:
```

Note: If a login prompt appears instead of the Cluster Setup wizard, start the wizard by logging in with the factory default settings and then enter the `cluster setup` command.

3. Enter the following command to create a new cluster:

```
create
```

4. Type `no` for single node cluster option

```
Do you intend for this node to be used as a single node cluster? {yes, no} [no]: no
```

5. Type `yes` for cluster network using network switches.

```
Will the cluster network be configured to use network switches? [yes]:yes
```

6. To activate HA and set storage failover, do the following:

```
Non-HA mode, Reboot node to activate HA

Do you want to reboot now to set storage failover (SFO) to HA mode? {yes, no}
[yes]: Enter
```

7. After the reboot, type `admin` at the login prompt.

```
admin
```

8. Continue creating the cluster. Enter `create` on the cluster setup wizard.

```
create
```

9. Repeat steps 3 and 4, if the cluster setup wizard prompts again.

10. The system defaults are displayed. Type `no` for using the system defaults. Follow the prompts to configure the cluster ports.

```
Existing cluster interface configuration found:

Port    MTU    IP                Netmask
e0a     9000   169.254.166.221  255.255.0.0
e0c     9000   169.254.20.239   255.255.0.0

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

System Defaults:
Private cluster network ports [e0a,e0c].
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]: no

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

List the private cluster network ports [e0a,e0c]: e0a, e0b, e0c, e0d
Enter the cluster ports' MTU size [9000]: Enter
Enter the cluster network netmask [255.255.0.0]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0a [169.254.102.165]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0b [169.254.49.135]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0c [169.254.132.54]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0d [169.254.11.125]: Enter
Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter an additional license key []:<<var_nfs_license>>
Enter an additional license key []:<<var_fcp_license>>

```

Note: The cluster is created. This can take a minute or two.

Note: For this validated architecture, NetApp recommends installing license keys for NetApp SnapRestore®, NetApp FlexClone®, and NetApp SnapManager® Suite. After you finish entering the license keys, press Enter.

```

Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
Enter the cluster management interface port [e0e]: e0i
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway: <<var_clustermgmt_gateway>>

```

11. Enter the DNS domain name.

```

Enter the DNS domain names:<<var_dns_domain_name>>
Enter the name server IP addresses:<<var_nameserver_ip>>

```

Note: If you have more than one name server IP address, separate them with a comma.

12. Set up the node.

```

Where is the controller located []:<<var_node_location>>
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
Enter the node management interface netmask:<<var_node01_mgmt_mask>>
Enter the node management interface default gateway:<<var_node01_mgmt_gateway>>

```

Note: The node management interface should be in a different subnet than the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

13. Type no for IPV4 DHCP on the service processor.

```

Enable IPv4 DHCP on the service processor interface [yes]: no

```

14. Set up the service processor (SP).

```

Enter the service processor interface IP address: <<var_node01_sp_ip>>
Enter the service processor interface netmask: <<var_node01_sp_netmask>>

```

```
Enter the service processor interface default gateway: <<var_node01_sp_gateway>>
```

15. Press Enter to accept the AutoSupport message.

16. Log in to the cluster interface with the admin user id and <<var_password>>.

17. Reboot the SP.

```
system node service-processor reboot-sp -node <<var_node01>>
Note: If your console connection is through the SP, it will be disconnected.
Do you want to reboot the SP ? {y|n}: y
```

Join Node 02 to Cluster

The first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01, and the node joining the cluster in this example is node 02. Table 23 lists the cluster network information required for joining node 02 to the existing cluster.

Table 23) Cluster join in clustered Data ONTAP prerequisites.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster node 02 IP address	<<var_node02_mgmt_ip>>
Cluster node 02 netmask	<<var_node02_mgmt_mask>>
Cluster node 02 gateway	<<var_node02_mgmt_gateway>>

To join node 02 to the existing cluster, complete the following steps:

1. At the login prompt, enter `admin`.

```
admin
```

2. The Cluster Setup wizard starts running on the console.

```
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}:
```

Note: If a login prompt is displayed instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings, and then enter the `cluster setup` command.

3. Enter the following command to join a cluster:

```
join
```

4. To activate HA and set storage failover, do the following:

```
Non-HA mode, Reboot node to activate HA

Do you want to reboot now to set storage failover (SFO) to HA mode? {yes, no}
[yes]: Enter
```

5. After the reboot, continue the cluster join process.
6. Data ONTAP detects existing cluster and agrees to join the same cluster. Follow the prompts to join the cluster.

```
Existing cluster interface configuration found:

Port      MTU      IP              Netmask
e0a       9000     169.254.226.177 255.255.0.0
e0c       9000     169.254.36.57   255.255.0.0

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

System Defaults:
Private cluster network ports [e0a,e0c].
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]: no

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

List the private cluster network ports [e0a,e0c]: e0a, e0b, e0c, e0d
Enter the cluster ports' MTU size [9000]: Enter
Enter the cluster network netmask [255.255.0.0]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0a [169.254.251.103]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0b [169.254.171.178]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0c [169.254.215.185]: Enter

Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0d [169.254.199.125]: Enter
```

Note: The cluster creation process can take a minute or two.

7. The steps to create a cluster are displayed.

```
Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter
```

Note: The node should find the cluster name.

8. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address: <<var_node02_mgmt_ip>>
Enter the node management interface netmask: <<var_node02_netmask>>Enter
Enter the node management interface default gateway: <<var_node02_gw>>Enter
```

Note: The node management interface should be in a subnet different from the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

9. Type no for IPV4 DHCP on the service processor.

```
Enable IPv4 DHCP on the service processor interface [yes]: no
```

10. Set up the SP.

```
Enter the service processor interface IP address: <<var_node01_sp_ip>>
Enter the service processor interface netmask: <<var_node01_sp_netmask>>
```

```
Enter the service processor interface default gateway: <<var_node01_sp_gateway>>
```

11. Press Enter to accept the AutoSupport message.
12. Log in to the cluster interface with the admin user id and <<var_password>>.
13. Reboot the SP.

```
system node service-processor reboot-sp -node <<var_node02>>  
Note: If your console connection is through the SP, it will be disconnected.  
Do you want to reboot the SP ? {y|n}: y
```

Log In to Cluster

To log in to the cluster, complete the following steps:

1. Open an SSH connection to either the cluster IP or host name.
2. Log in as the `admin` user with the password you provided earlier.

Zero All Spare Disks

To zero all spare disks in the cluster, complete the following step:

1. Run the following command:

```
disk zerospares
```

Set Onboard UTA2 Ports Personality

To set the personality of the onboard Unified Target Adapter 2 (UTA2) ports, complete the following steps:

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

```
icefl-stc1:> ucadmin show  
Node          Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Status  
-----  
icefl-stc1-01 0e       cna    target -        -        online  
icefl-stc1-01 0f       cna    target -        -        online  
icefl-stc1-01 0g       cna    target -        -        online  
icefl-stc1-01 0h       cna    target -        -        online  
icefl-stc1-02 0e       cna    target -        -        online  
icefl-stc1-02 0f       cna    target -        -        online  
icefl-stc1-02 0g       cna    target -        -        online  
icefl-stc1-02 0h       cna    target -        -        online  
8 entries were displayed.
```

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

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

Note: The ports must be offline to run this command.

Set Auto-Revert on Cluster Management Interface

To set the auto-revert parameter on the cluster management interface, complete the following step:

1. Run the following command:

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

Create Failover Group for Cluster Management

To create a failover group for the cluster management port, complete the following step:

1. Run the following commands:

```
network interface failover-groups create -failover-group fg-cluster-mgmt -node <<var_node01>> -port e0i
network interface failover-groups create -failover-group fg-cluster-mgmt -node <<var_node02>> -port e0i
```

Assign Management Failover Group to Cluster Management LIF

To assign the cluster management failover group to the cluster management logical interface (LIF), complete the following step:

1. Run the following command:

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt -failover-group fg-cluster-mgmt
```

Create Failover Groups Node Management Ports

To create failover groups for the node management ports, complete the following step:

1. Run the following commands:

```
network interface failover-groups create -failover-group fg-node-mgmt01 -node <<var_node01>> -port e0i
network interface failover-groups create -failover-group fg-node-mgmt01 -node <<var_node01>> -port e0M
network interface failover-groups create -failover-group fg-node-mgmt02 -node <<var_node02>> -port e0i
network interface failover-groups create -failover-group fg-node-mgmt02 -node <<var_node02>> -port e0M
```

Assign Node Management Failover Groups to Node Management LIFs

To assign the node management failover groups to the node management LIFs, complete the following step:

1. Run the following commands:

```
network interface modify -vserver <<var_node01>> -lif mgmt1 -auto-revert true -failover-group fg-node-mgmt01
network interface modify -vserver <<var_node02>> -lif mgmt1 -auto-revert true -failover-group fg-node-mgmt02
```

Enable Flash Cache

To enable Flash Cache on each node, complete the following step:

1. From the cluster management interface, run the following commands:

```
system node run -node <<var_node01>> options flexscale.enable on
system node run -node <<var_node01>> options flexscale.lopri_blocks off
system node run -node <<var_node01>> options flexscale.normal_data_blocks on
system node run -node <<var_node02>> options flexscale.enable on
system node run -node <<var_node02>> options flexscale.lopri_blocks off
system node run -node <<var_node02>> options flexscale.normal_data_blocks on
```

Note: Data ONTAP 8.2.1 and later does not require a separate license for Flash Cache.

Note: For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, refer to [TR-3832: Flash Cache Best Practice Guide](#). Before customizing the settings, determine whether the custom settings are required or if the default settings are sufficient.

Create Aggregates

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

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node1 -nodes <<var_node01>> -diskcount <<var_num_disks>>
aggr create -aggregate aggr1_node2 -nodes <<var_node02>> -diskcount <<var_num_disks>>
```

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

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

Note: The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until both `aggr1_node1` and `aggr1_node2` are online.

2. Disable NetApp Snapshot[®] copies for the two data aggregates recently created.

```
node run <<var_node01>> aggr options aggr1_node1 nosnap on
node run <<var_node02>> aggr options aggr1_node2 nosnap on
```

3. Delete any existing Snapshot copies for the two data aggregates.

```
node run <<var_node01>> snap delete -A -a -f aggr1_node1
node run <<var_node02>> snap delete -A -a -f aggr1_node2
```

4. Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02.

```
aggr show
aggr rename -aggregate aggr0 -newname <<var_node01_rootaggrname>>
```

Verify Storage Failover

To confirm that storage failover is enabled, complete the following steps for a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

2. Both the nodes `<<var_node01>>` and `<<var_node02>>` must be capable of performing a takeover.
3. Continue with step 5, if the nodes are capable of performing a takeover.
4. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_node01>> -enabled true
```

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

5. Verify the HA status for a two-node cluster.

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

```
cluster ha show
```

- Continue with step 8 if high availability is configured.
- Enable HA mode only for the two-node cluster.

Note: Do not run this command for clusters with more than two nodes because it will cause problems with failover.

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

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

```
storage failover hwassist show
storage failover modify -hwassist-partner-ip <<var_node02_mgmt_ip>> -node <<var_node01>>
storage failover modify -hwassist-partner-ip <<var_node01_mgmt_ip>> -node <<var_node02>>
```

Disable Flow Control on 10GbE and UTA2 Ports

A NetApp best practice is to disable flow control on all the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following step:

- Run the following commands:

```
net port modify -node <<var_node01>> -port e0a -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node01>> -port e0b -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node01>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node01>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node01>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node01>> -port e0g -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node02>> -port e0a -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node02>> -port e0b -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node02>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node02>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node02>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

net port modify -node <<var_node02>> -port e0g -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
```

Create LACP Interface Group

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

To create interface groups, complete the following step:

1. Run the following commands:

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0e
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0g
ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0e
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0g
```

Note: All interfaces must be in the `down` status before being added to an interface group.

Note: The interface group name must follow the standard naming convention of `a<number><letter>`, where `<number>` is an integer in the range [0-999] without leading zeros and `<letter>` is a lowercase letter.

Create VLANs

To create VLANs, complete the following step:

1. Create VLANs.

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_IB-MGMT_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_IB-MGMT_vlan_id>>
```

Configure Jumbo Frames

To configure a clustered Data ONTAP network port to use jumbo frames (which usually have an MTU of 9,000 bytes), complete the following step:

1. From the cluster shell, run the following command:

```
network port modify -node <<var_node01>> -port a0a -mtu 9000

WARNING: Changing the network port settings will cause a serveral second interruption in carrier.
Do you want to continue? {y|n}: y

network port modify -node <<var_node02>> -port a0a -mtu 9000

WARNING: Changing the network port settings will cause a serveral second interruption in carrier.
Do you want to continue? {y|n}: y
```

Note: After the MTU for ifgrp is set to 9000, all the VLAN interfaces on that ifgrp will have MTU 9,000.

Configure NTP

To configure time synchronization on the cluster, complete the following steps:

1. Set the time zone for the cluster.

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

Note: For example, in the Eastern United States, the time zone is `America/New_York`.

2. Set the date for the cluster.

```
date <<ccyymmddhhmm.ss>>
```

Note: The format for the date is

<[Century] [Year] [Month] [Day] [Hour] [Minute] . [Second]>; for example,
201309081735.17

3. Configure the Network Time Protocol (NTP) for each node in the cluster.

```
system services ntp server create -node <<var_node01>> -server <<var_global_ntp_server_ip>>
system services ntp server create -node <<var_node02>> -server <<var_global_ntp_server_ip>>
```

Configure SNMP

To configure SNMP, complete the following steps:

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

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

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <<var_oncommand_server_fqdn>>
```

Configure SNMPv1 Access

To configure SNMPv1 access, complete the following step:

1. Set the shared secret plain-text password, which is called a community.

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

Note: Use the `delete all` command with caution. If community strings are used for other monitoring products, the `delete all` command will remove them.

Create SNMPv3 User

SNMPv3 requires that a user be defined and configured for authentication. To create and configure a user for SNMPv3, complete the following steps:

1. Create a user called `snmpv3user`.

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

2. Enter the authoritative entity's EngineID and select `md5` as the authentication protocol.
3. Run the `security snmpusers` command to view the EngineID.
4. When prompted, enter an eight-character minimum-length password for the authentication protocol.
5. Select `des` as the privacy protocol.
6. When prompted, enter an eight-character minimum-length password for the privacy protocol.

Configure AutoSupport

AutoSupport sends support summary information to NetApp through HTTPS. To configure AutoSupport, complete the following step:

1. Run the following command:

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

Enable Cisco Discovery Protocol

To enable Cisco Discovery Protocol (CDP) on the NetApp storage controllers, complete the following step:

Note: To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.

1. Run the following command:

```
node run -node <<var_node01>> options cdpd.enable on
node run -node <<var_node02>> options cdpd.enable on
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

Note: Storage virtual machine (SVM) is referred to as Vserver in the GUI and CLI.

1. Start the SVM (Vserver) setup wizard.

```
vserver setup

Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.

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 your answers to previous questions, and
"exit" if you want to quit the Vserver Setup Wizard. Any changes
you made before typing "exit" will be applied.

You can restart the Vserver Setup Wizard by typing "vserver setup". To accept a default
or omit a question, do not enter a value.

Vserver Setup wizard creates and configures only data Vservers.
If you want to create a Vserver with Infinite Volume use the vserver create command.

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.
```

2. Enter the SVM name.

```
Enter the Vserver name: RHEL_Server
```

3. Select the SVM data protocols to configure.

```
Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi, ndmp}: fcp
```

4. Select the SVM client services to configure.

```
Choose the Vserver client services to configure {ldap, nis, dns}:Enter
```

5. Enter the SVM's root volume aggregate:

```
Enter the Vserver's root volume aggregate {aggr1_n1, aggr1_n2} [aggr1_n1]:aggr1_n1
```

6. Enter the SVM language setting. English is the default.

7. Enter the SVM language setting, or enter help to see all languages.

8. Enter the SVM's security style:

```
Enter the Vserver root volume's security style {mixed, ntfs, unix} [unix]: Enter
```

9. Answer no to Do you want to create a data volume?

```
Do you want to create a data volume? {yes, no} [Yes]: no
```

10. Answer no to Do you want to create a logical interface?

```
Do you want to create a logical interface? {yes, no} [Yes]: no
```

11. Answer no to Do you want to Configure FCP? {yes, no} [yes]: no.

```
Do you want to Configure FCP? {yes, no} [yes]: no
```

12. Add the two data aggregates to the RHEL_Server aggregate list for NetApp Virtual Console.

```
vserver modify -vserver RHEL_Server -aggr-list aggr1_n1, aggr1_n2
```

Create Load-Sharing Mirror of SVM Root Volume in Clustered Data ONTAP

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

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

```
volume create -vserver RHEL_Server -volume rootvol_m01 -aggregate aggr1_n1 -size 1GB -type DP
volume create -vserver RHEL_Server -volume rootvol_m02 -aggregate aggr1_n2 -size 1GB -type DP
```

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

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

3. Create the mirroring relationships.

```
snapmirror create -source-path //RHEL_Server/rootvol -destination-path //RHEL_Server/rootvol_m01
-type LS -schedule 15min
snapmirror create -source-path //RHEL_Server/rootvol -destination-path //RHEL_Server/rootvol_m02
-type LS -schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path //RHEL_Server/rootvol
```

Create FC Service

To create the FC service on each SVM, complete the following step:

1. Create the FC service on each SVM. This command also starts the FC service and sets the FC alias to the name of the SVM.

```
fcg create -vserver RHEL_Server
```

Configure HTTPS Access

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

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

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

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

```
security certificate show
```

3. To generate and install self-signed certificates, run the following commands as one-time commands:

Note: To delete existing/expired certificates before creating the certificates, run the `security certificate delete` command.

```
security certificate create -vserver RHEL_Server -common-name
<<var_security_cert_vserver_common_name>> -size 2048 -country <<var_country_code>> -state
<<var_state>> -locality <<var_city>> -organization <<var_org>> -unit <<var_unit>> -email-addr
<<var_storage_admin_email>> -type server
security certificate create -vserver <<var_clustername>> -common-name
<<var_security_cert_cluster_common_name>> -size 2048 -country <<var_country_code>> -state
<<var_state>> -locality <<var_city>> -organization <<var_org>> -unit <<var_unit>> -email-addr
<<var_storage_admin_email>> -type server
```

```

security certificate create -vserver <<var_node01>> -common-name
<<var_security_cert_node01_common_name>> -size 2048 -country <<var_country_code>> -state
<<var_state>> -locality <<var_city>> -organization <<var_org>> -unit <<var_unit>> -email-addr
<<var_storage_admin_email>> -type server
security certificate create -vserver <<var_node02>> -common-name
<<var_security_cert_node02_common_name>> -size 2048 -country <<var_country_code>> -state
<<var_state>> -locality <<var_city>> -organization <<var_org>> -unit <<var_unit>> -email-addr
<<var_storage_admin_email>> -type server

```

4. Configure and enable SSL and HTTPS access and disable Telnet access.

```

system services web modify -external true -ssl3-enabled true
Warning: Modifying the cluster configuration will cause pending web service requests to be
        interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -action allow
system services firewall policy create -policy mgmt -service http -action deny -ip-list 0.0.0.0/0

```

5. To obtain the values for the parameters that would be required in the following step, run the `security certificate show` command.

6. The following is an example output.

```

Cluster1:~# security certificate show
Vserver      Serial Number  Common Name                                     Type
-----
RHEL_Server
             54CFAD31      icef1-stc11.ice.rtp.netapp.com                 server
Certificate Authority: icef1-stc11.ice.rtp.netapp.com
Expiration Date: Tue Feb 02 12:00:33 2016
Cluster1:~#

```

7. The required parameters are:

- `-common-name` : `icef1-stc11.ice.rtp.netapp.com`
- `-ca` : Certificate Authority: `icef1-stc11.ice.rtp.netapp.com`
- `-serial` : `54CFAD31`

8. Use preceding values in the following command:

```

security ssl modify -vserver RHEL_Server -common-name <<var_security_cert_vserver_common_name>> -
server-enabled true -client-enabled false -ca <<var_security_certificate_vserver_authority>> -
serial <<var_security_certificate_vserver_serial_no>>

Warning: The certificate <<var_security_cert_vserver_common_name>> is a self-signed certificate,
        which offers no verification of identity by client machines. This presents the risk of
        man-in-the-middle attacks by malicious third-parties.
Do you want to continue? {y|n}:y

security ssl modify -vserver <<var_clustername>> -common-name
<<var_security_cert_cluster_common_name>> -server-enabled true -client-enabled false -ca
<<var_security_certificate_cluster_authority>> -serial
<<var_security_certificate_cluster_serial_no>>

Warning: The certificate <<var_security_cert_vserver_common_name>> is a self-signed certificate,
        which offers no verification of identity by client machines. This presents the risk of
        man-in-the-middle attacks by malicious third-parties.
Do you want to continue? {y|n}:y

security ssl modify -vserver <<var_node01>> -common-name <<var_security_cert_node01_common_name>> -
server-enabled true -client-enabled false -ca <<var_security_certificate_node01_authority>> -
serial <<var_security_certificate_node01_serial_no>>

Warning: The certificate <<var_security_cert_vserver_common_name>> is a self-signed certificate,
        which offers no verification of identity by client machines. This presents the risk of
        man-in-the-middle attacks by malicious third-parties.
Do you want to continue? {y|n}:y

```

```

security ssl modify -vserver <<var_node02>>-common-name <<var_security_cert_node02_common_name>>
-server-enabled true -client-enabled false -ca <<var_security_certificate_node02_authority>> -
serial <<var_security_certificate_node02_serial_no>>
Warning: The certificate <<var_security_cert_vserver_common_name>> is a self-signed certificate,
which offers no verification of identity by client machines. This presents the risk of
man-in-the-middle attacks by malicious third-parties.
Do you want to continue? {y|n}:y

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

```

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

Create FlexVol Volume

To create a FlexVol[®] volume, complete the following steps:

1. Collect the following information: the volume's name and size and the aggregate on which it will exist.
2. Create two datastore volumes: a server boot volume and a volume to hold the OnCommand database LUN. Also, update the SVM root volume load sharing mirrors to make the NFS mounts accessible.

```

volume create -vserver RHEL_Server -volume ucs_boot -aggregate aggr1_n1 -size 100g -state online
-policy default -space-guarantee none -percent-snapshot-space 0
[Job 114866] Job succeeded: Successful

snapmirror update-ls-set -source-path //RHEL_Server/rootvol

```

Create Boot LUNS

To create two boot LUNS, RHEL-Infra-01 and RHEL-Infra-02, complete the following step:

1. Run the following commands:

```

lun create -vserver RHEL_Server -volume ucs_boot -lun RHEL-Infra-01 -size 10g -ostype linux -
space-reserve disabled
lun create -vserver RHEL_Server -volume ucs_boot -lun RHEL-Infra-02 -size 10g -ostype linux -
space-reserve disabled

```

Enable Deduplication

To enable deduplication on the appropriate volumes, complete the following step:

1. Run the following commands.

```

volume efficiency on -vserver RHEL_Server -volume ucs_boot

```

Create FCP LIFs

To create four FCP LIFS (two on each node), complete the following step:

1. Run the following commands:

```

network interface create -vserver RHEL_Server -lif fcp_lif01a -role data -data-protocol fcp -
home-node <<var_node01>> -home-port 0e
network interface create -vserver RHEL_Server -lif fcp_lif01b -role data -data-protocol fcp -
home-node <<var_node01>> -home-port 0g
network interface create -vserver RHEL_Server -lif fcp_lif02a -role data -data-protocol fcp -
home-node <<var_node02>> -home-port 0e
network interface create -vserver RHEL_Server -lif fcp_lif02b -role data -data-protocol fcp -
home-node <<var_node02>> -home-port 0g

```

Create Failover Group for SVM Management

To create a failover group for the SVM management port, complete the following step:

1. Run the following commands:

```
network interface failover-groups create -failover-group vs_mgmt01 -node <<var_node01>> -port
a0a-<<var_IB-MGMT_vlan_id>>
network interface failover-groups create -failover-group vs_mgmt01 -node <<var_node02>> -port
a0a-<<var_IB-MGMT_vlan_id>>
```

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, complete the following step:

1. Run the following commands:

```
network interface create -vserver RHEL_Server -lif vsmgmt -role data -data-protocol none -home-
node <<var_node02>> -home-port a0a-<<var_IB-MGMT_vlan_id>> -address <<var_vserver_mgmt_ip>> -
netmask <<var_vserver_mgmt_mask>> -status-admin up -failover-policy nextavail -firewall-policy
mgmt -auto-revert true -failover-group vs_mgmt01

network routing-groups route create -vserver RHEL_Server -routing-group
<<var_vserver_mgmt_first_three_octets_0>> -destination 0.0.0.0/0 -gateway
<<var_vserver_mgmt_gateway>>

security login password -username vsadmin -vserver RHEL_Server
Enter a new password: <<var_vsadmin_password>>
Enter it again: <<var_vsadmin_password>>

security login unlock -username vsadmin -vserver RHEL_Server
```

9.4 Configure Server: FlexPod Cisco UCS Base

Perform 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. The steps are necessary to provision the Cisco UCS C-Series and B-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS Fabric Interconnect 6248 A

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/gui) ? console
Enter the setup mode; setup newly or restore from backup. (setup/restore)? setup
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
Enforce strong password? (y/n) [y]: y
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)? (yes/no) [n]: y
Enter the switch fabric (A/B) []:A
Enter the system name: <<var_ucs_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_cluster_ip>>
Configure the DNS Server IPv4 address? (yes/no) [n]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? (yes/no) [n]: y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
```

2. Review the settings displayed on the console. If they are correct, answer `yes` to apply and save the configuration.
3. Wait for the login prompt to verify that the configuration has been saved.

Cisco UCS Fabric Interconnect 6248 B

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/gui) ? console
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will
be added to the cluster. Continue (y/n) ? y
Enter the admin password of the peer Fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_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 confirm that the configuration has been saved.

9.5 Configure FlexPod Cisco UCS FCoE on Clustered Data ONTAP

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6248 fabric interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
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.

Upgrade Cisco UCS Manager Software to Version 2.2(2c)

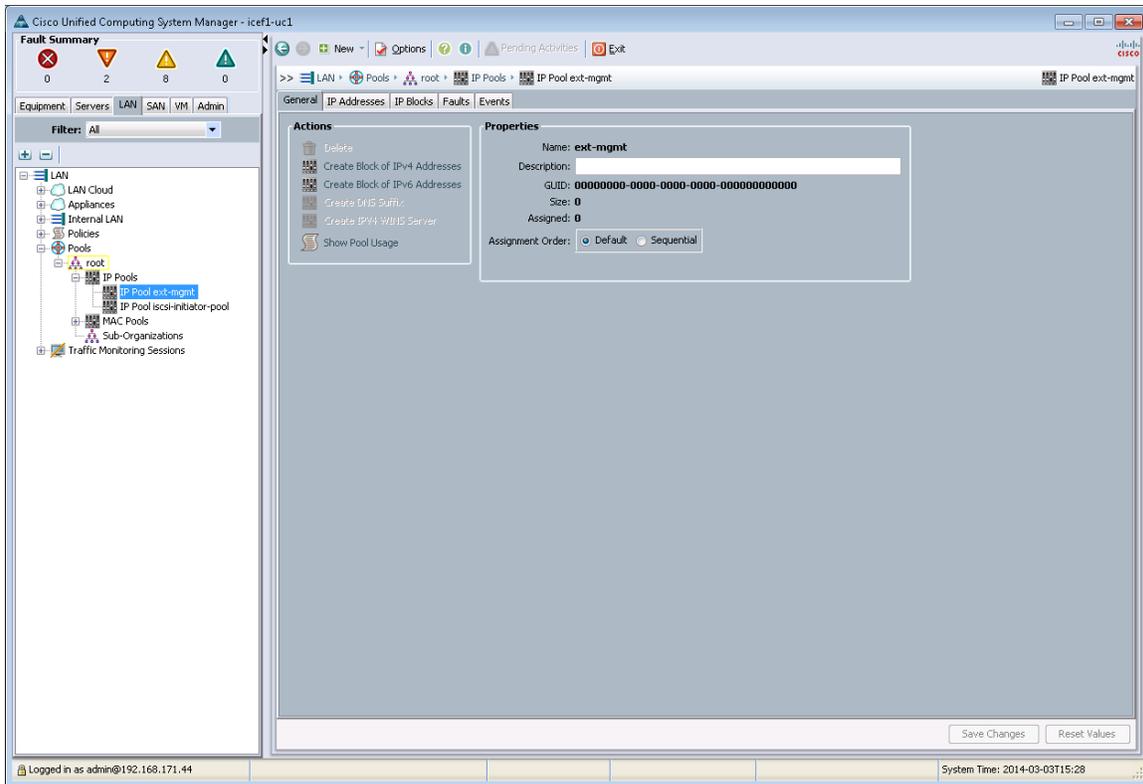
This document assumes the use of Cisco UCS Manager Software version 2.2(2c). To upgrade the Cisco UCS Manager software and the UCS 6248 fabric interconnect software to version 2.2(2c), refer to the [Cisco UCS Manager Install and Upgrade Guides](#).

Add Block of IP Addresses for Out-of-Band KVM Access

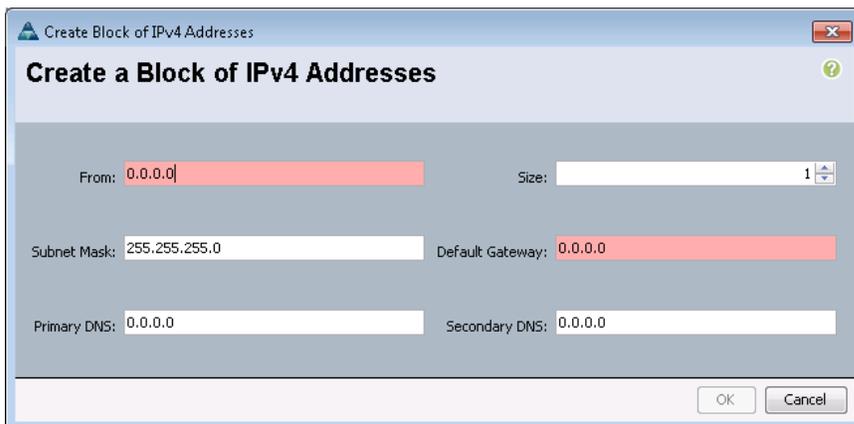
To create a block of IP addresses for server keyboard, video, and mouse (KVM) access in the Cisco UCS environment, complete the following steps:

Note: This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Pools > root > IP Pools > IP Pool ext-mgmt.



3. In the Actions pane, select Create Block of IP Addresses.
4. Enter the starting IP address of the block, the number of IP addresses required, the subnet, and the gateway information.



5. Click OK to create the IP block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

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

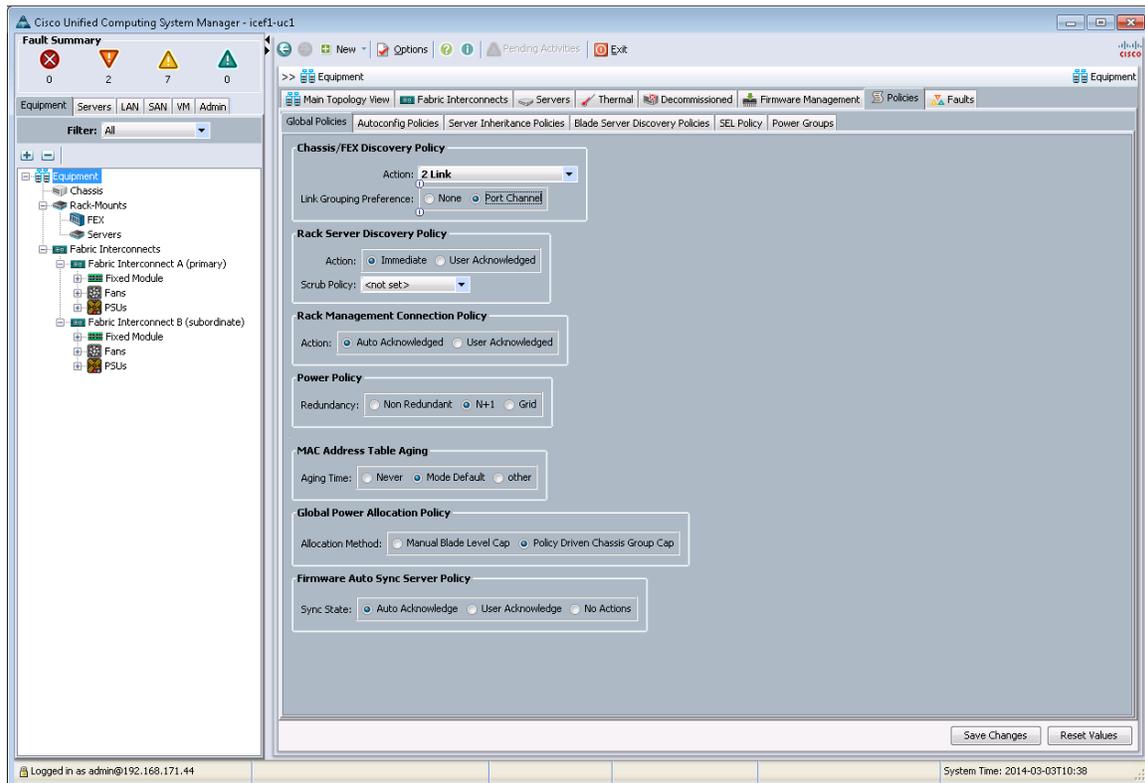
3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.
6. Enter <<var_global_ntp_server_ip>> and click OK.
7. Click OK.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and C-Series servers.

To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment node and select Equipment in 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 2-link or set it to match the number of uplink ports that are cabled between the chassis and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.



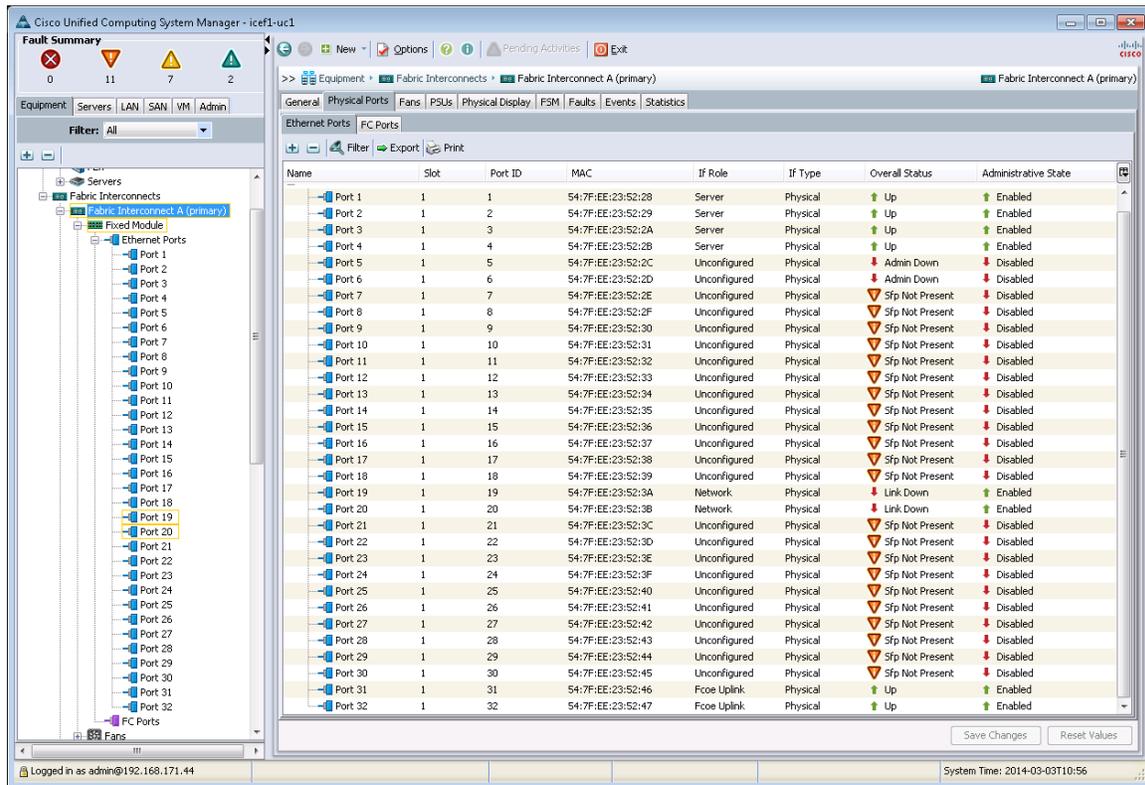
5. Click Save Changes.
6. Click OK.

Enable Server and Uplink Ports

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

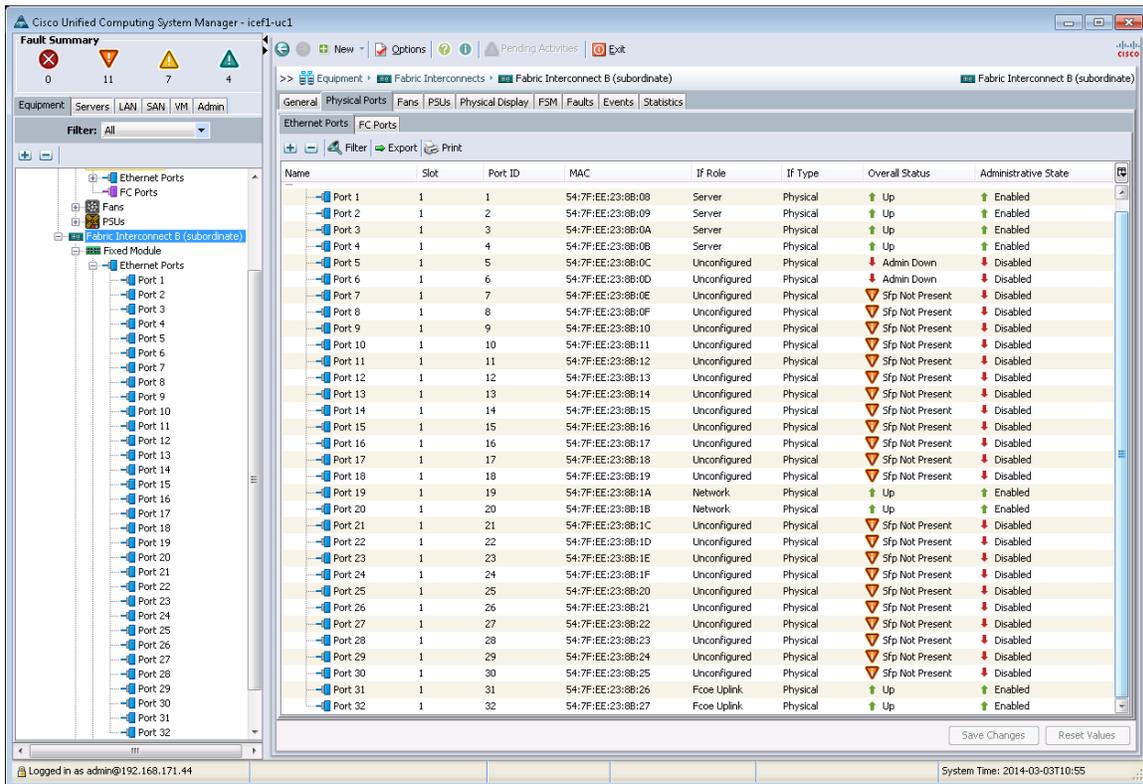
1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.

3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis or directly connected to C-Series rack servers, right-click them, and select Configure as Server Port.
5. Click Yes to confirm server ports and click OK.
6. Select ports 19 and 20 that are connected to the Cisco Nexus 5548 switches, right-click them, and select Configure as Uplink Port.
7. Click Yes to confirm uplink ports and click OK.
8. Select ports 31 and 32, which will serve as FCoE uplinks to the Cisco Nexus 5548 switches. Right-click them and select Configure as FCoE Uplink Port.
9. Click Yes to confirm FCoE uplink ports and click OK.
10. In the left pane, navigate to Fabric Interconnect A. In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that the ports have been configured correctly in the If Role column.



11. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
12. Expand Ethernet Ports.
13. Select the ports that are connected to the chassis or directly connected to C-Series rack servers, right-click 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 5548 switches, right-click them, and select Configure as Uplink Port.
16. Click Yes to confirm the uplink ports and click OK.
17. Select ports 31 and 32 that will serve as FCoE uplinks to the Cisco Nexus 5548 switches, right-click them, and select Configure as FCoE Uplink Port.
18. Click Yes to confirm FCoE uplink ports and click OK.

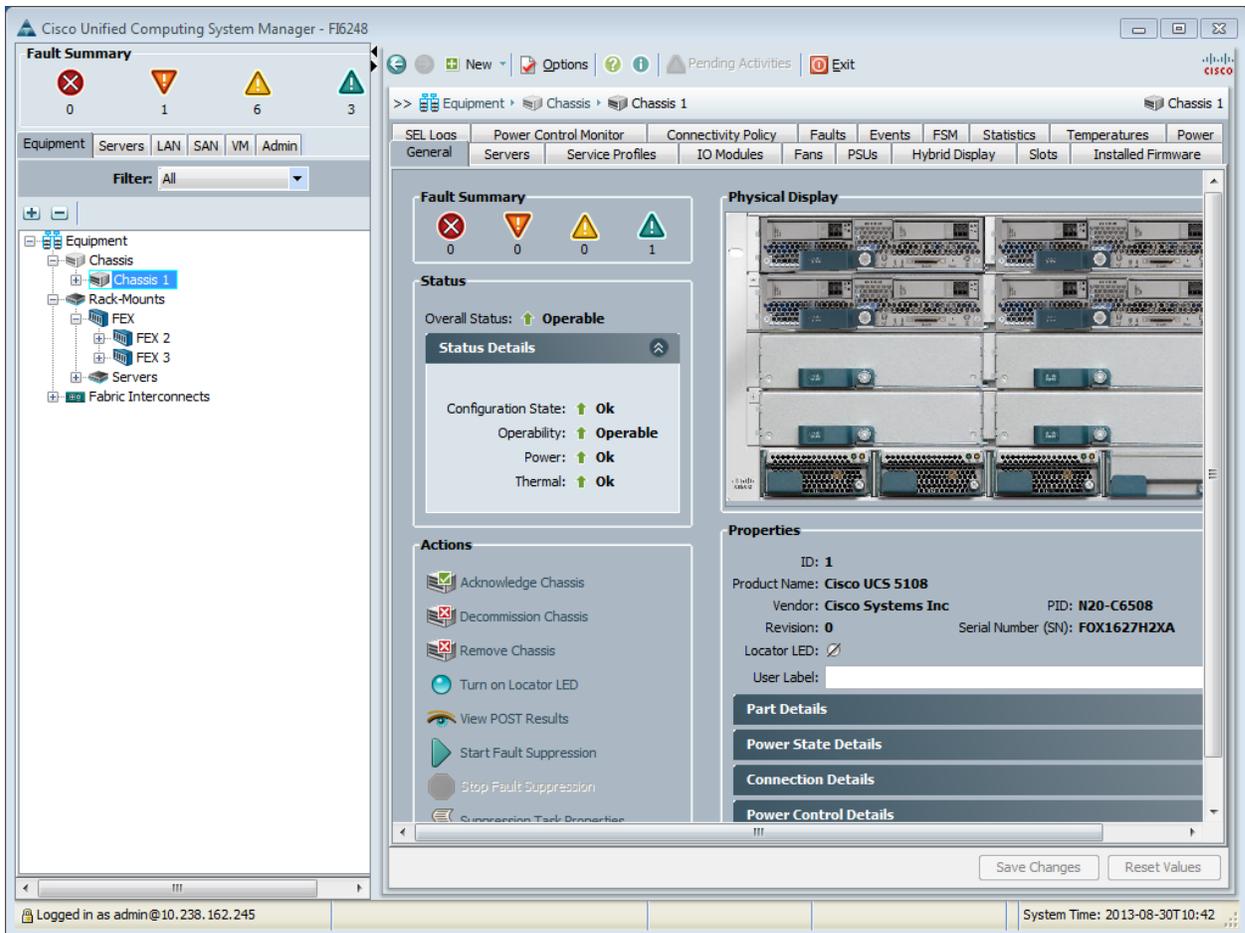
In the left pane, navigate to Fabric Interconnect B. In the right pane, navigate to the Physical Ports node > Ethernet Ports node. Confirm that the ports have been configured correctly in the If Role column.



Acknowledge Cisco UCS Chassis

To acknowledge all Cisco UCS chassis, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
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.

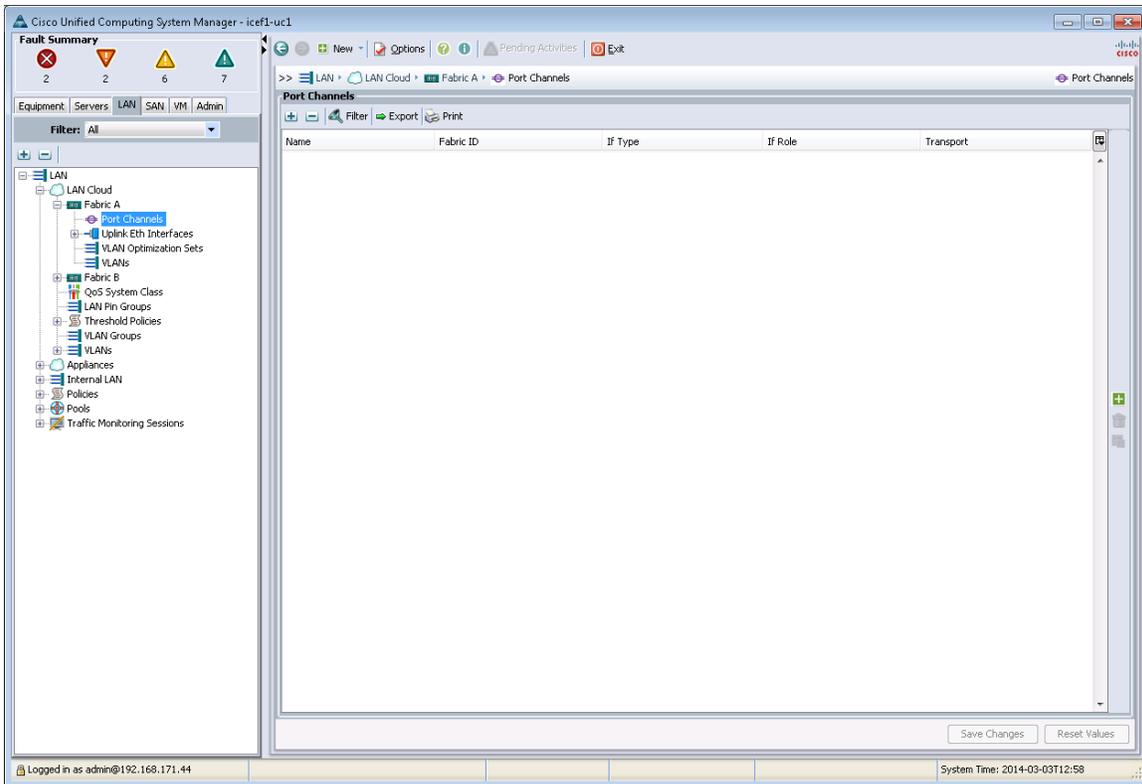
Create Uplink Port Channels to Cisco Nexus 5548UP Switches

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

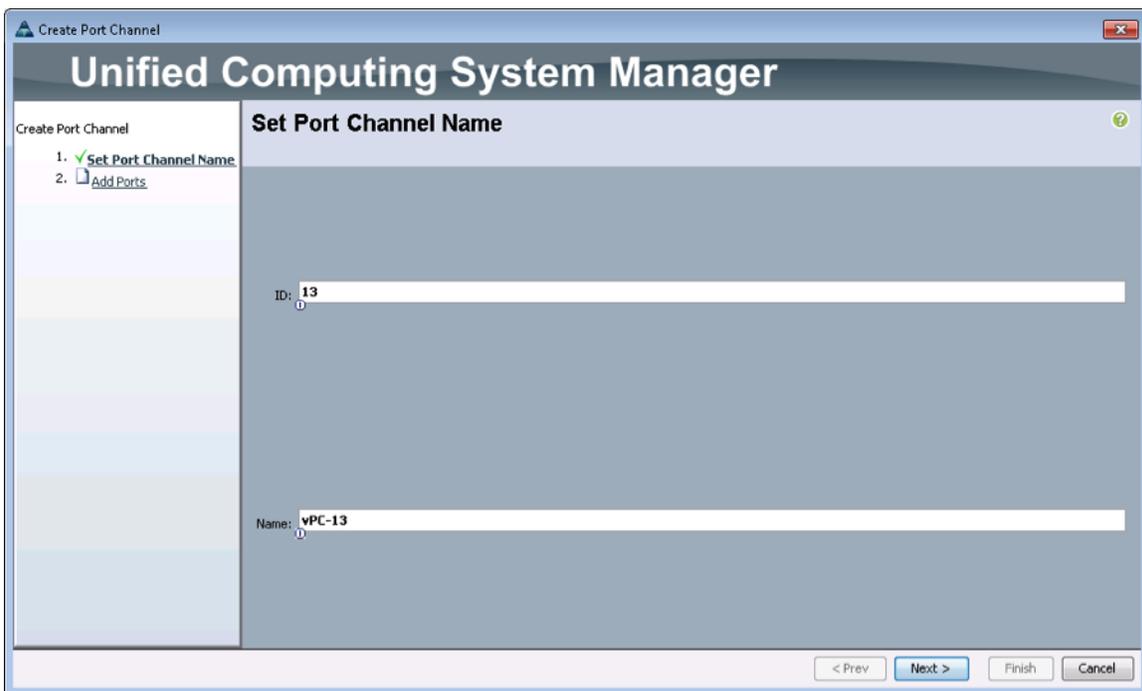
1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

Note: In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 5548UP switches and one from Fabric B to both Cisco Nexus 5548UP switches.

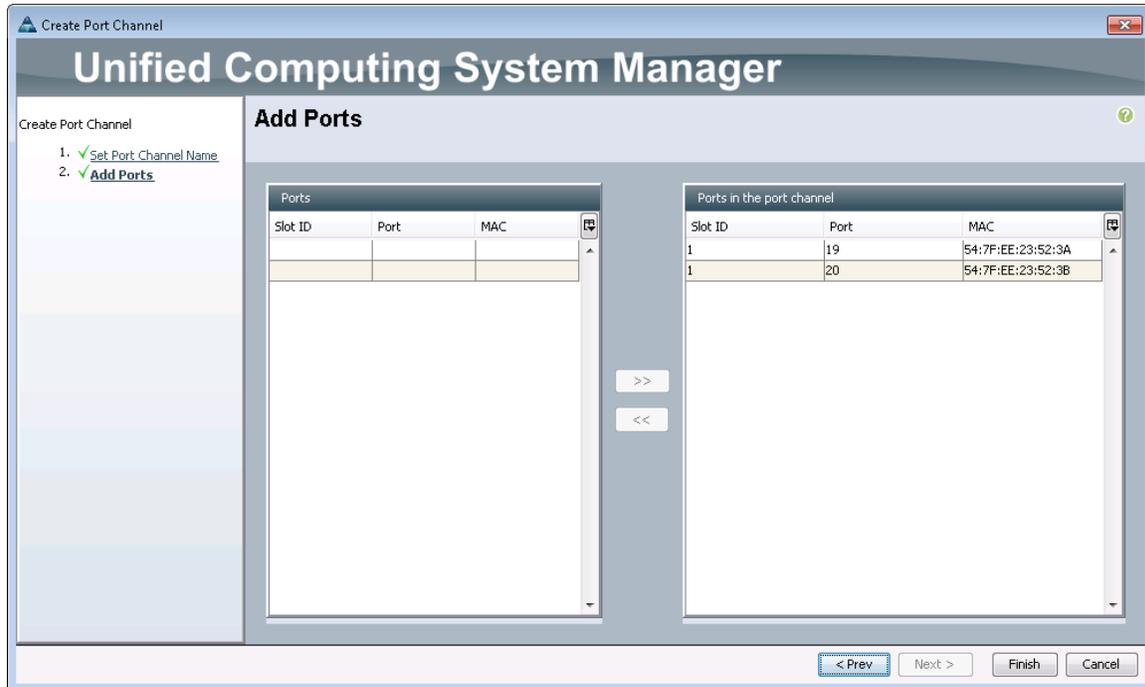
2. Under LAN > LAN Cloud, expand node Fabric A.



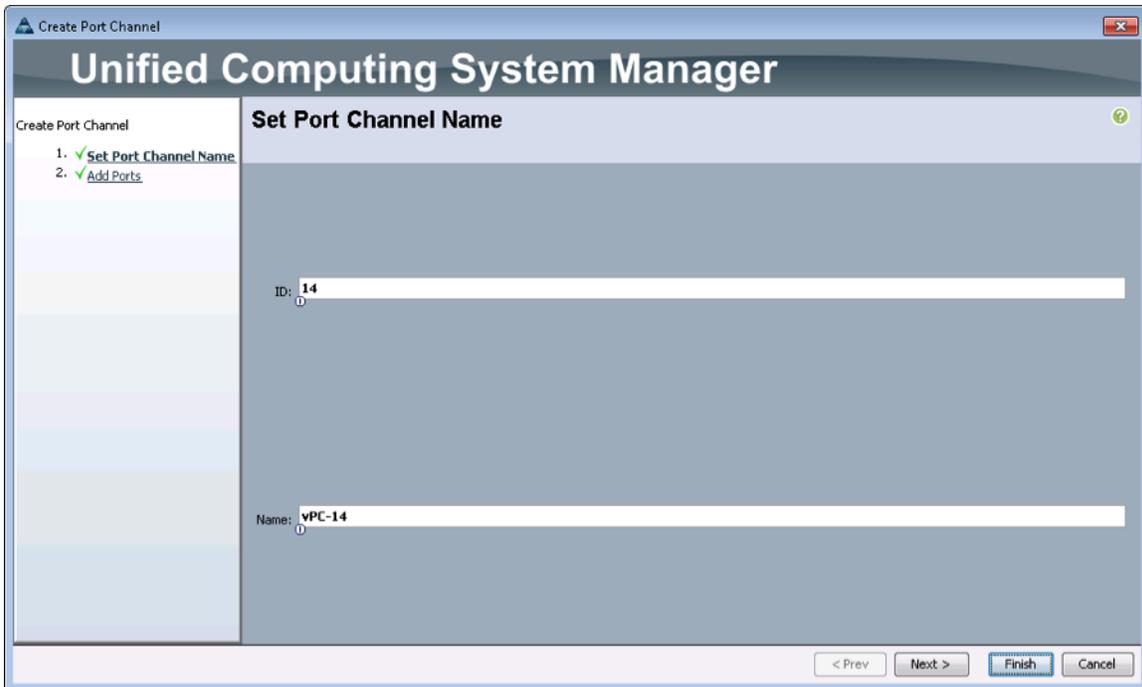
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 13 as the unique ID of the port channel.
6. Enter vPC-13 as the name for port channel.



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 node Fabric B.
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 as the name for port channel.

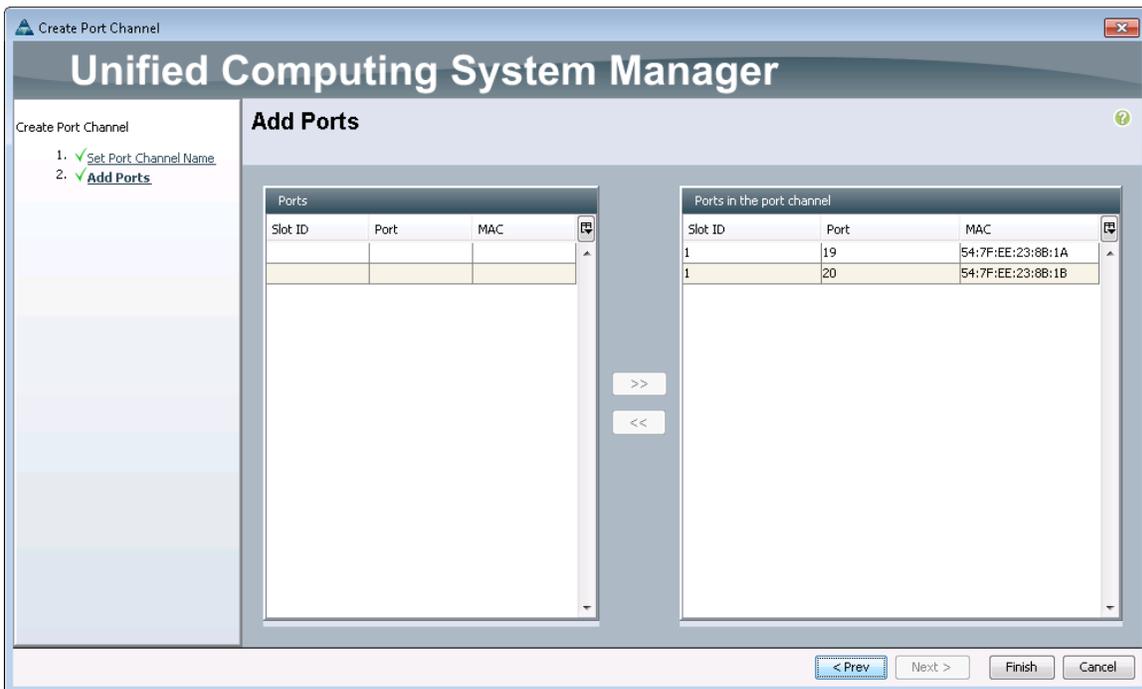


17. Click Next.

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 an Organization

Organizations are used to organize resources and restrict access to various groups within the IT infrastructure, thereby enabling multi-tenancy of the compute resources.

Note: Although this document does not assume the use of organizations, this procedure provides instructions for creating one.

To configure an organization in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Optional: Enter a description for the organization.
4. Click OK.
5. Click OK in the confirmation message.

Create MAC Address Pools

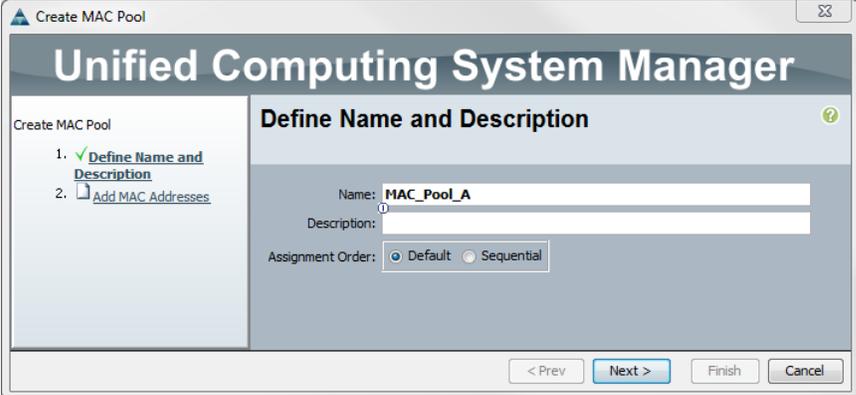
To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root.

Note: In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter `MAC_Pool_A` as the name for MAC pool.
6. Optional: Enter a description for the MAC pool.

Note: Keep the Assignment Order set to Default.

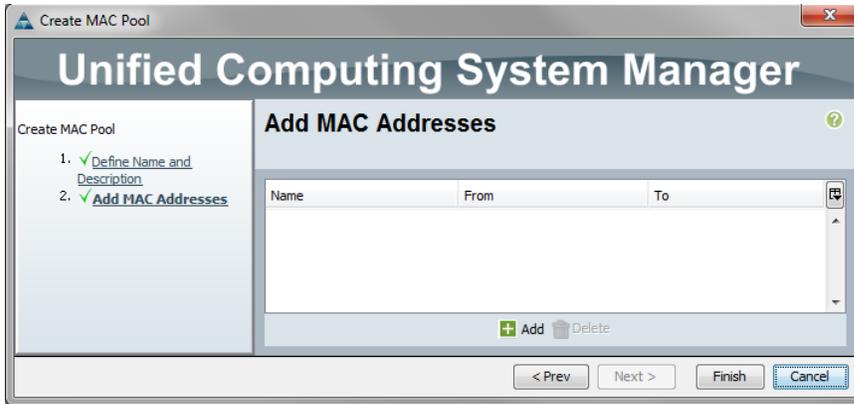


The screenshot shows the 'Create MAC Pool' dialog box in Cisco UCS Manager. The dialog is titled 'Unified Computing System Manager' and 'Create MAC Pool'. It has a progress bar on the left with two steps: '1. Define Name and Description' (checked) and '2. Add MAC Addresses'. The main area is titled 'Define Name and Description' and contains the following fields and options:

- Name: `MAC_Pool_A`
- Description: (empty field)
- Assignment Order: Default Sequential

At the bottom of the dialog are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

7. Click Next.

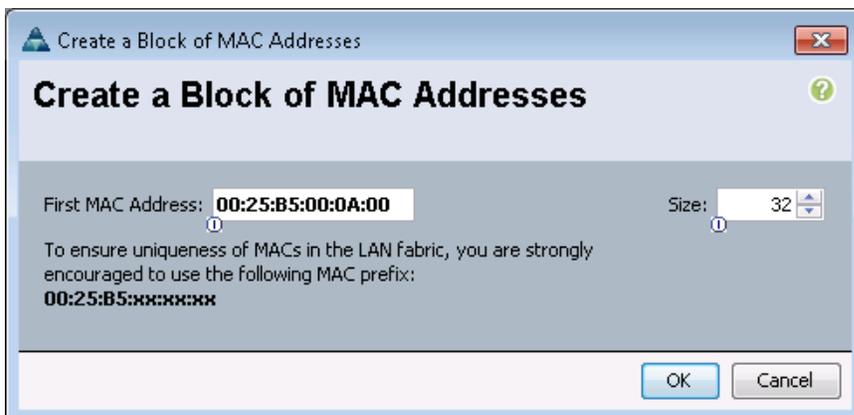


8. Click Add.

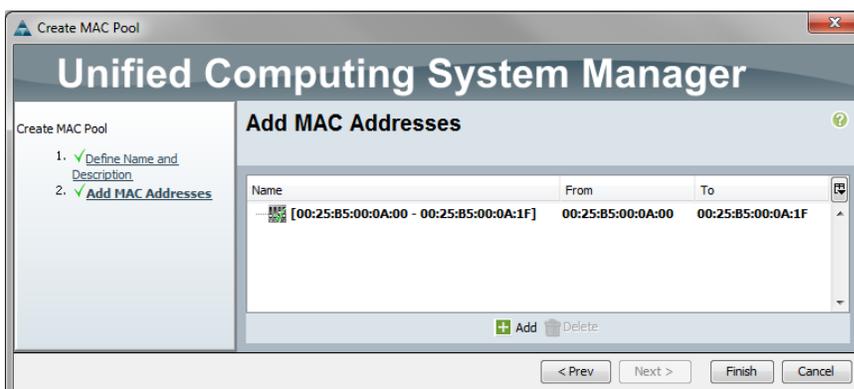
9. Specify a starting MAC address.

Note: For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as Fabric A addresses.

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



11. Click OK.

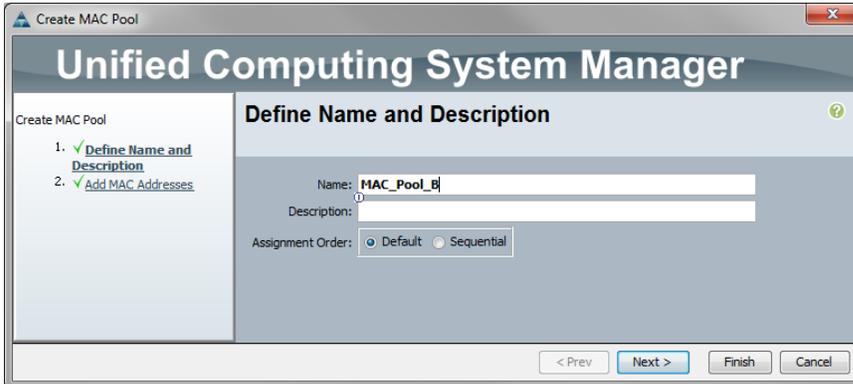


12. Click Finish.

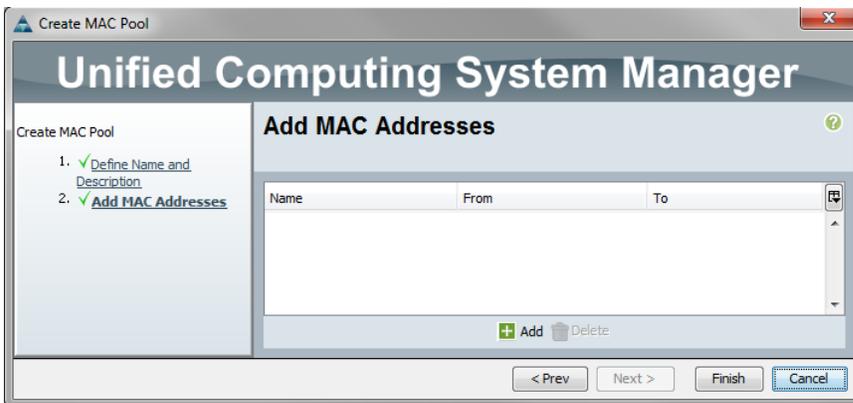
13. In the confirmation message, click OK.

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 for MAC pool.
17. Optional: Enter a description for the MAC pool.

Note: Select Default for the Assignment Order.



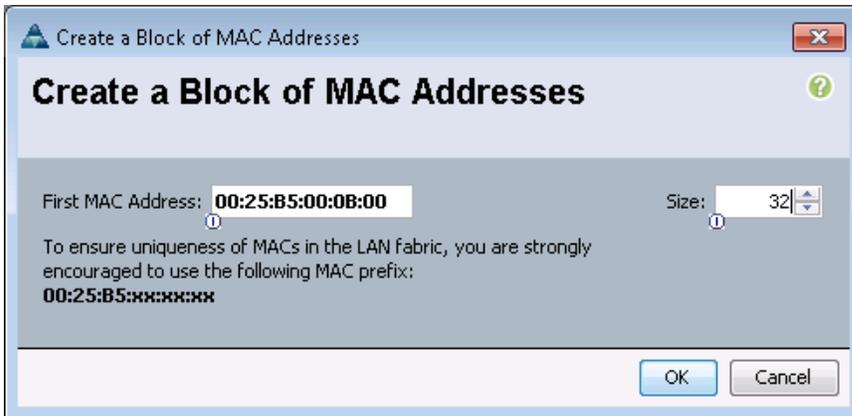
18. Click Next.



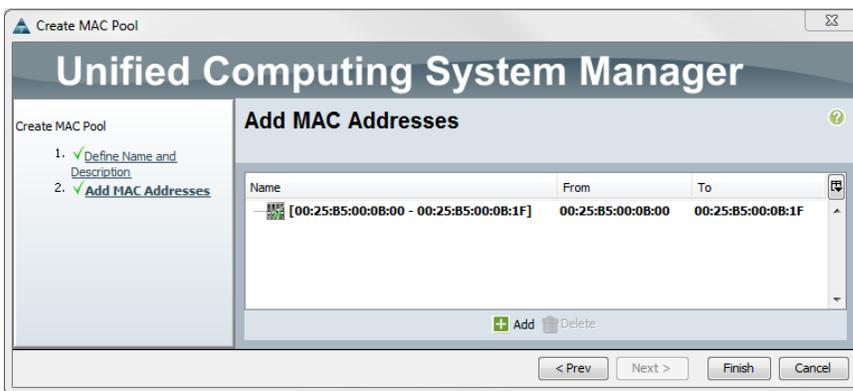
19. Click Add.
20. Specify a starting MAC address.

Note: For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting MAC address to identify all the MAC addresses in this pool as Fabric B addresses.

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



22. Click OK.



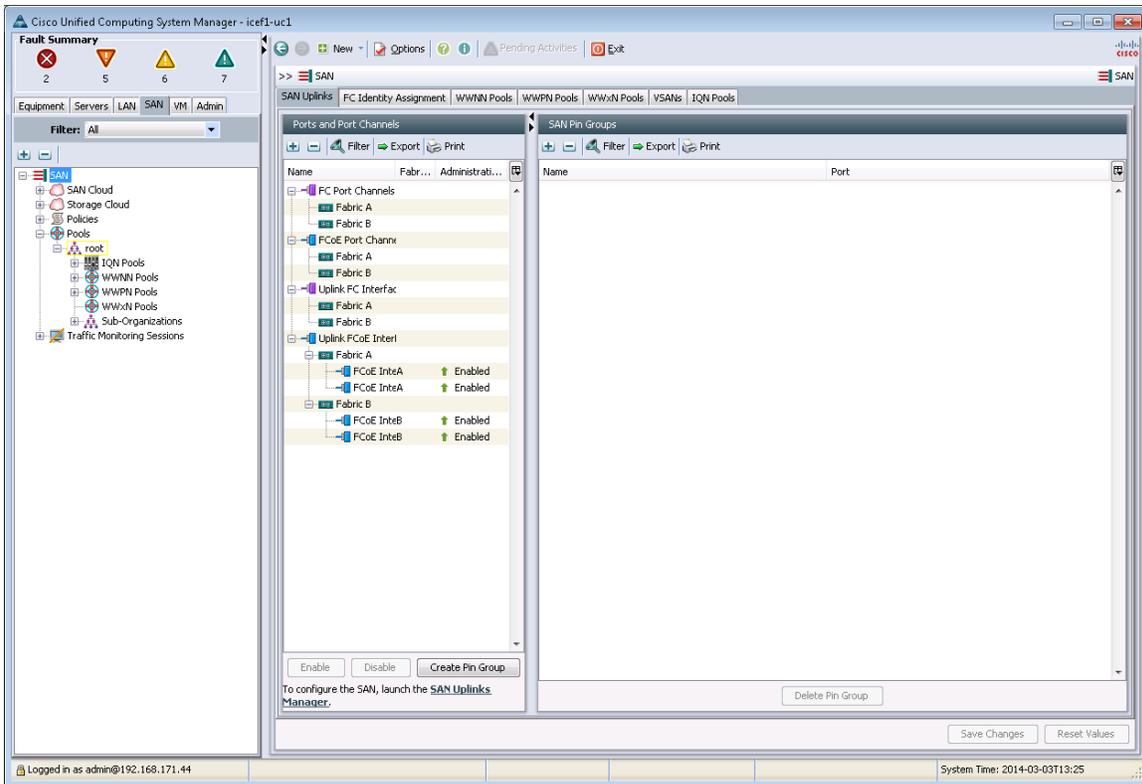
23. Click Finish.

24. In the confirmation message, click OK.

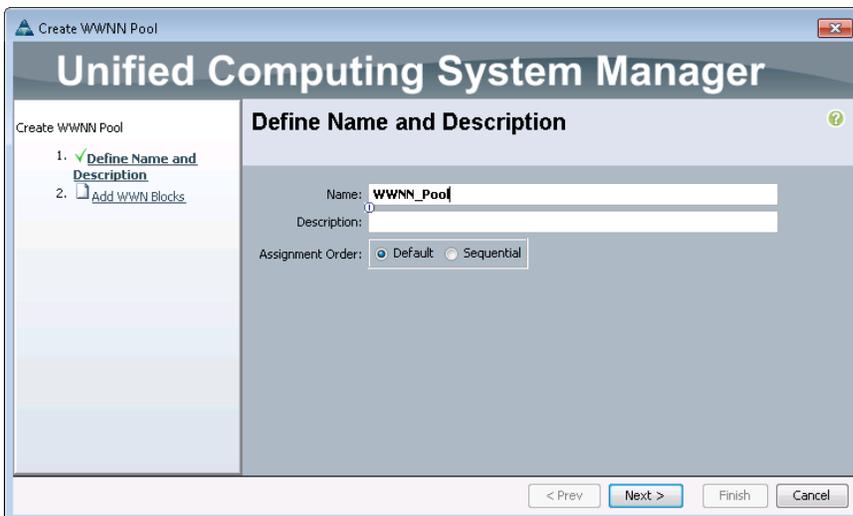
Create WWNN Pools

To configure the necessary worldwide node name (WWNN) pools for the Cisco UCS environment, complete the following steps:

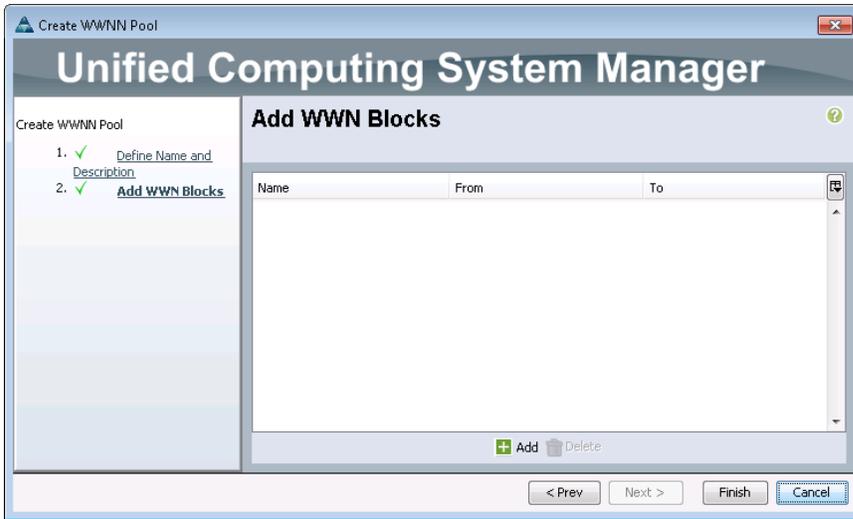
1. In Cisco UCS Manager, click the SAN node in the navigation pane.
2. Select Pools > root.



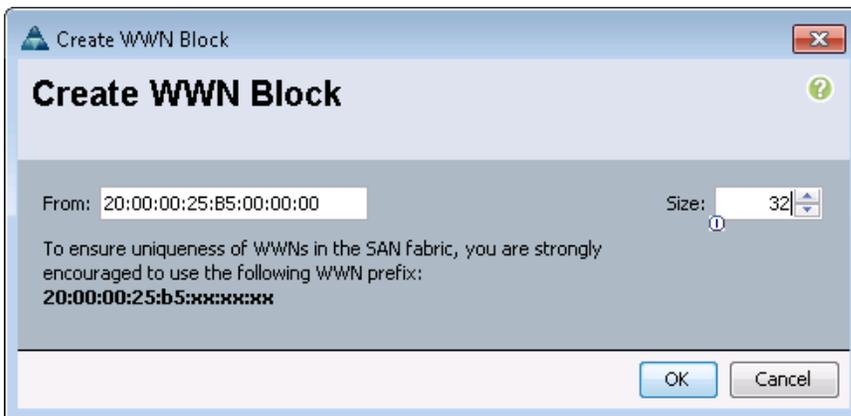
3. Right-click WWNN Pools.
4. Select Create WWNN Pool.
5. Enter `WWNN_pool` as the name for WWNN pool.
6. Optional: Add a description for the WWNN pool.
7. Select Default for the Assignment Order.



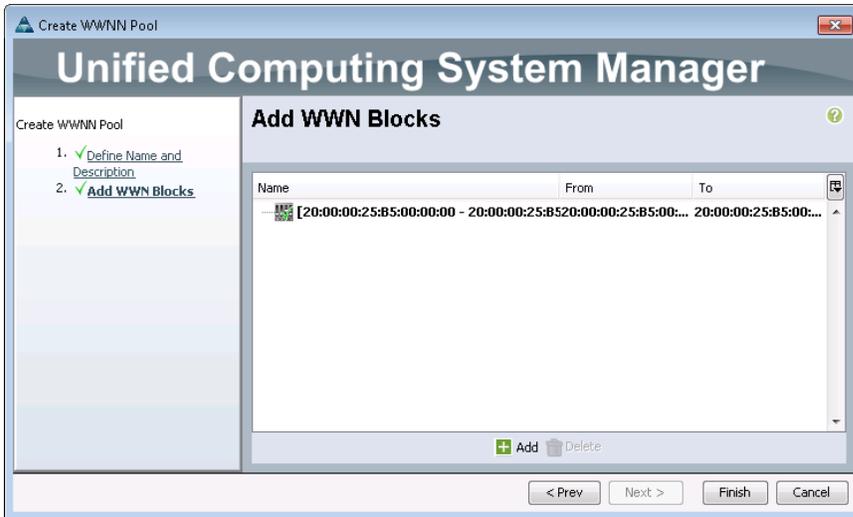
8. Click Next.



9. Click Add to add a block of WWNNs.
10. Either retain the default block of WWNNs, or specify a base WWNN.
11. Specify a size for the WWNN block that is sufficient to support the available blade or server resources.



12. Click OK.



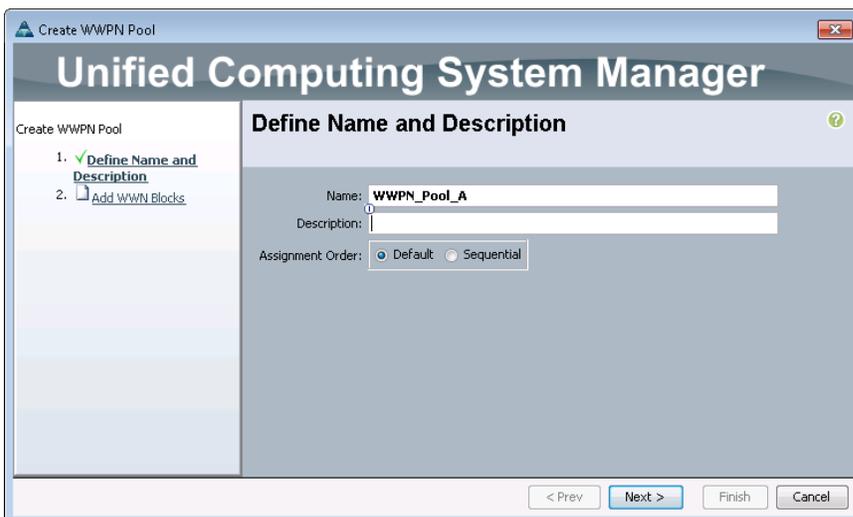
13. Click Finish.

14. Click OK.

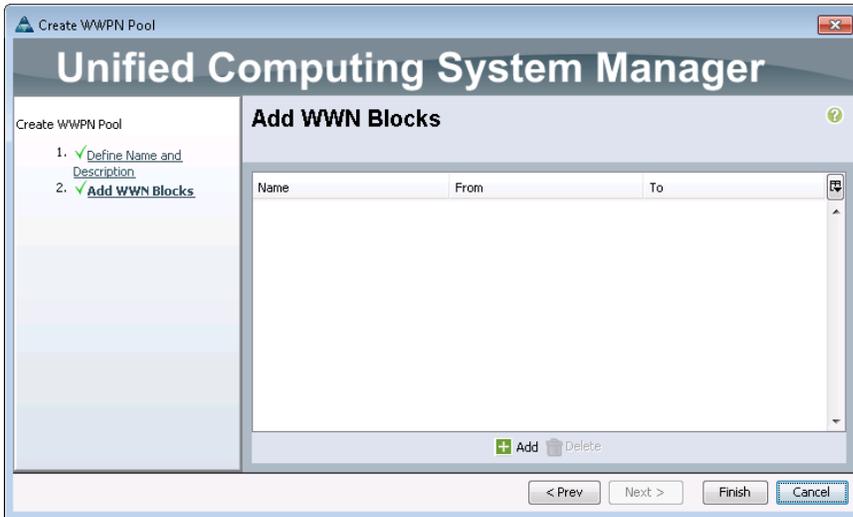
Create WWPN Pools

To configure the necessary worldwide port name (WWPN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
 - Note:** In this procedure, two WWPN pools are created: one for Fabric A and one for Fabric B.
3. Right-click WWPN Pools.
4. Select Create WWPN Pool.
5. Enter `WWPN_Pool_A` as the name for WWPN pool for Fabric A.
6. Optional: Enter a description for this WWPN pool.
7. Select Default for Assignment Order.



8. Click Next.

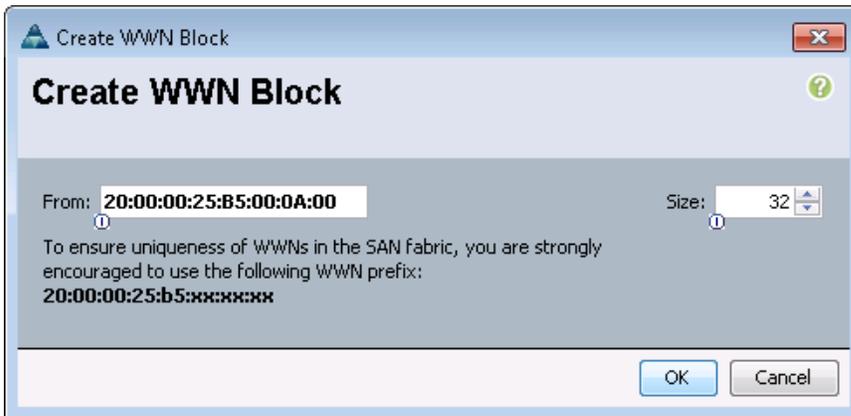


9. Click Add to add a block of WWPNs.

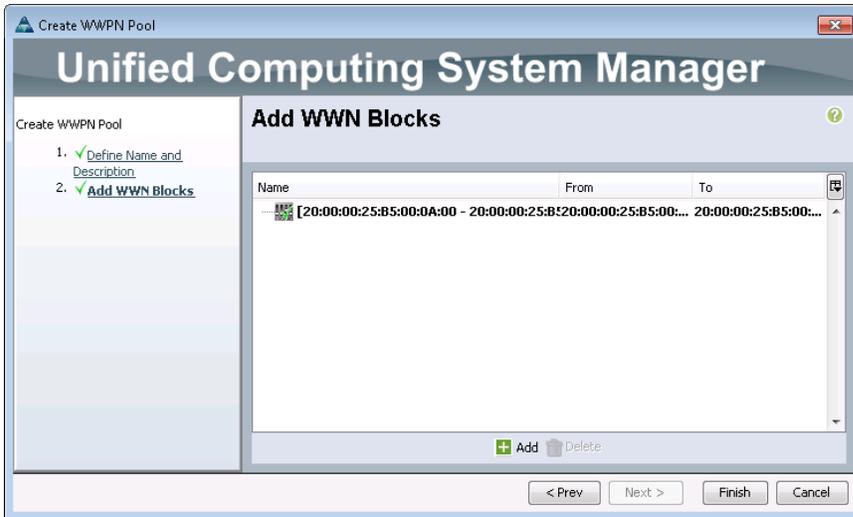
10. Specify the starting WWPN in the block for Fabric A.

Note: For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting WWPN to identify all the WWPNs in this pool as Fabric A addresses.

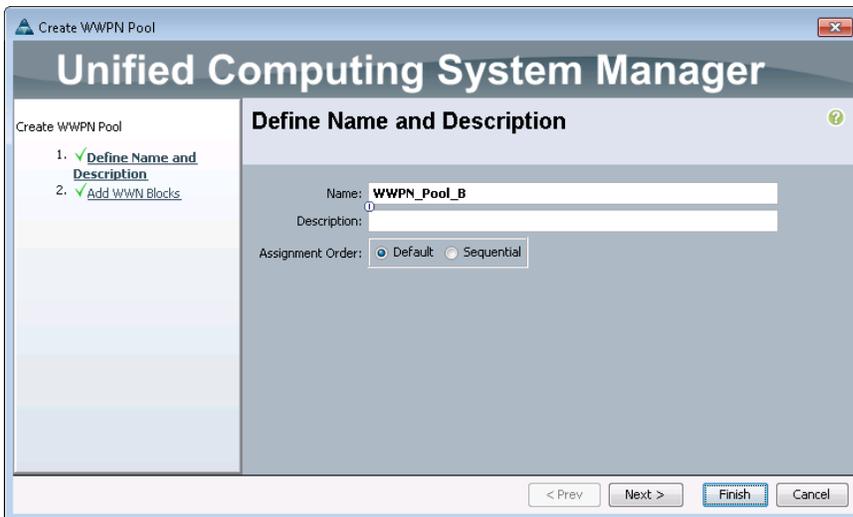
11. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.



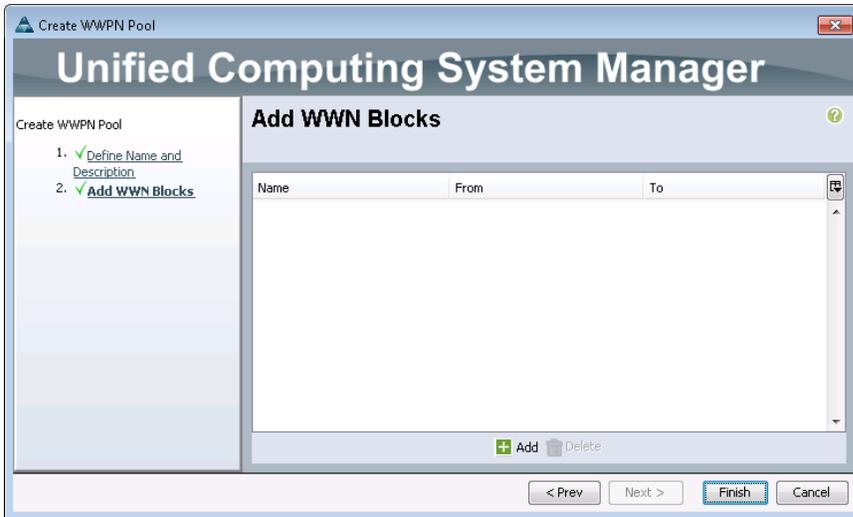
12. Click OK.



13. Click Finish to create the WWPN pool.
14. Click OK.
15. Right-click WWPN Pools.
16. Select Create WWPN Pool.
17. Enter `WWPN_Pool_B` as the name for the WWPN pool for Fabric B.
18. Optional: Enter a description for this WWPN pool.
19. Select Default for the Assignment Order.



20. Click Next.

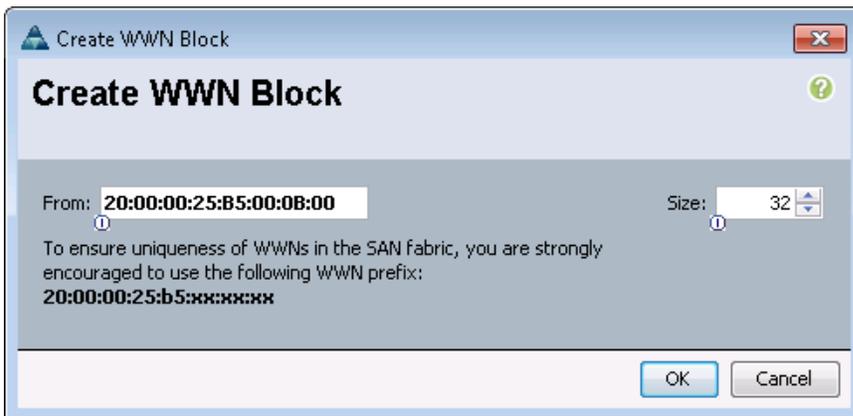


21. Click Add to add a block of WWPNs.

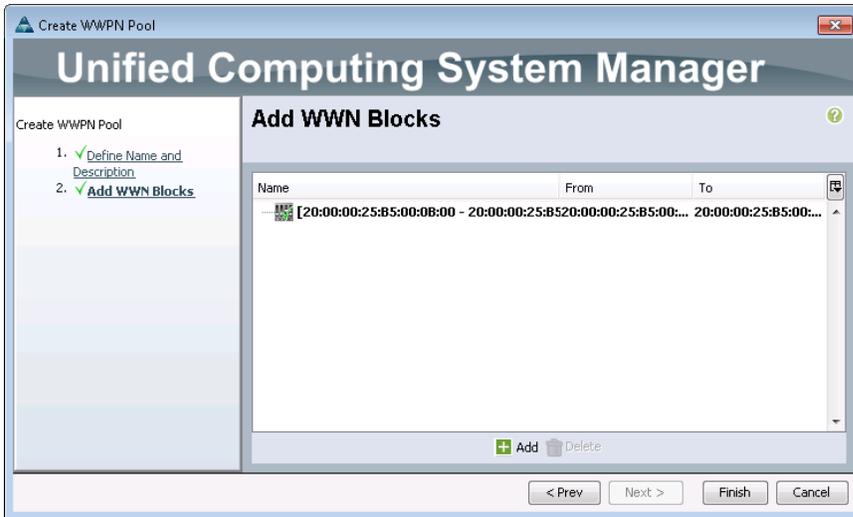
22. Enter the starting WWPN address in the block for Fabric B.

Note: For the FlexPod solution, the recommendation is to place 0B in the next-to-last octet of the starting WWPN to identify all the WWPNs in this pool as Fabric B addresses.

23. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.



24. Click OK.



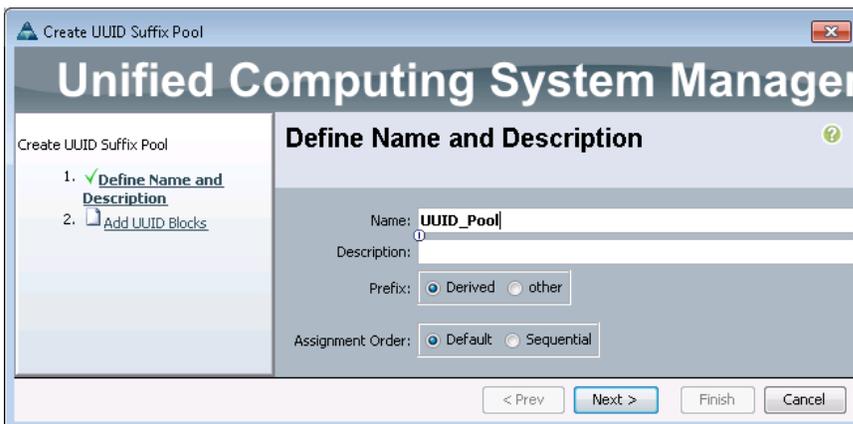
25. Click Finish.

26. Click OK.

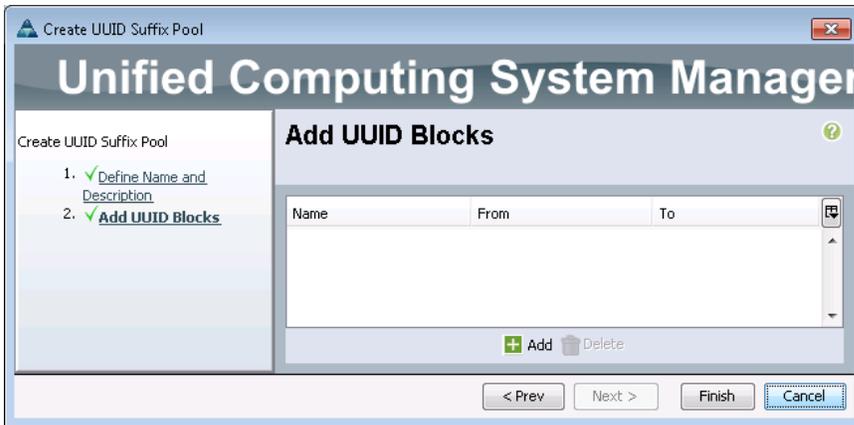
Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in 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 for UUID suffix pool.
6. Optional: Enter a description for UUID suffix pool.
7. Select the Derived option for Prefix.
8. Select Default for the Assignment Order.



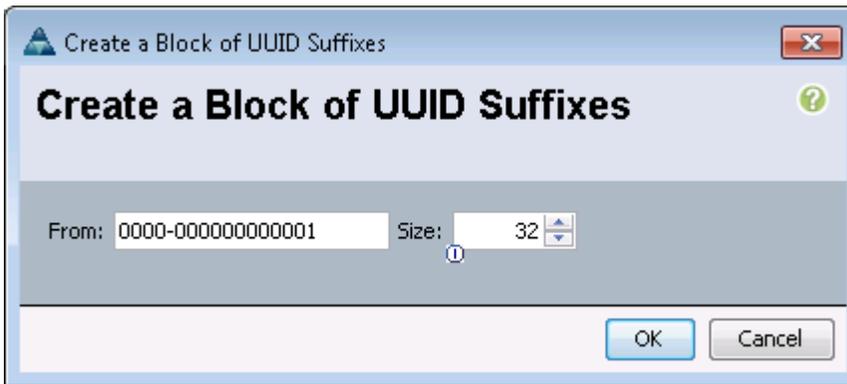
9. Click Next.



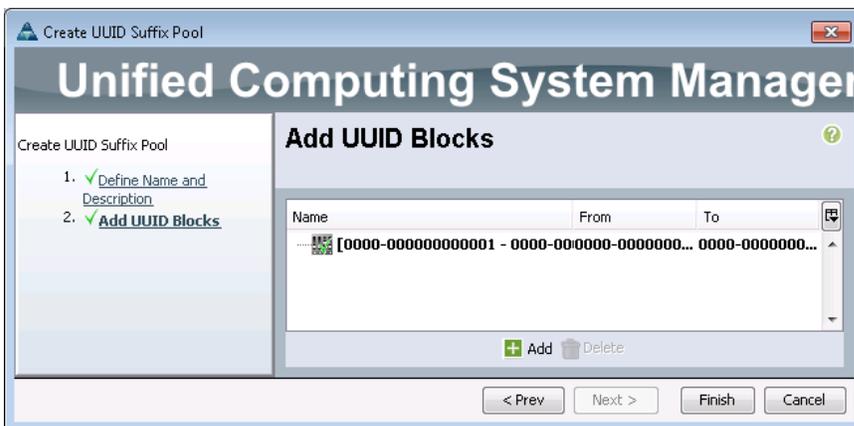
10. Click Add to add a block of UUIDs.

11. Select From as the default setting.

12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.



13. Click OK.



14. Click Finish.

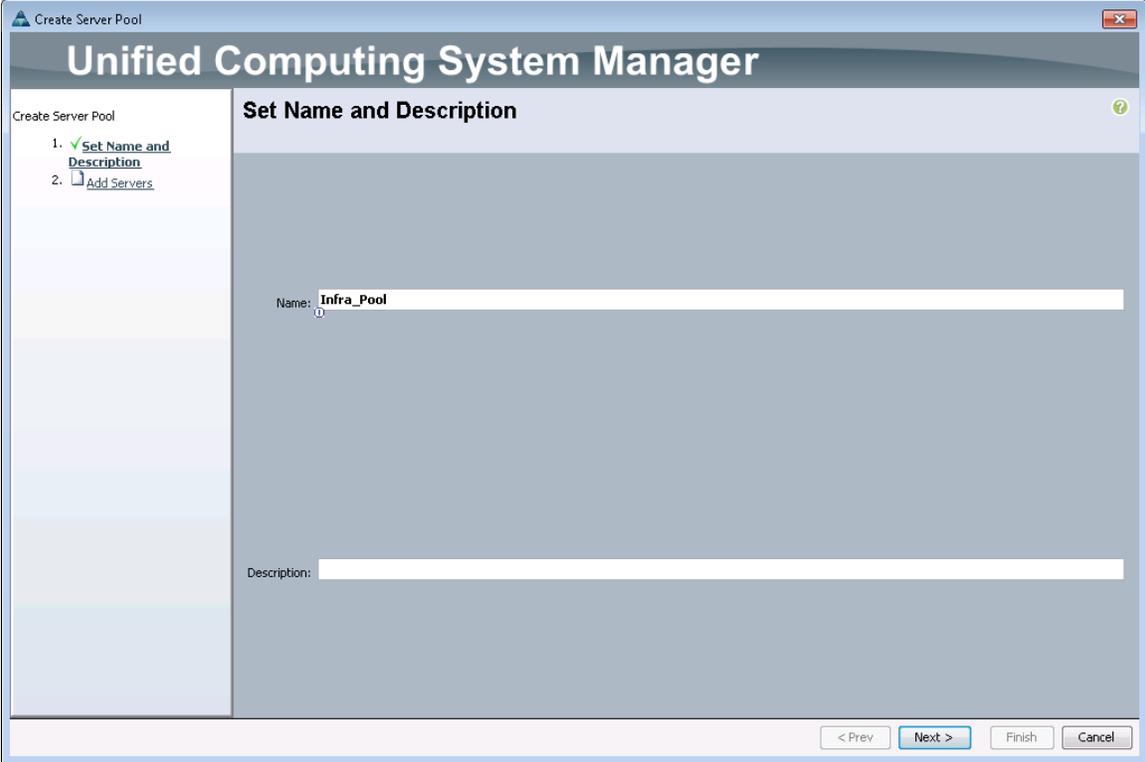
15. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:

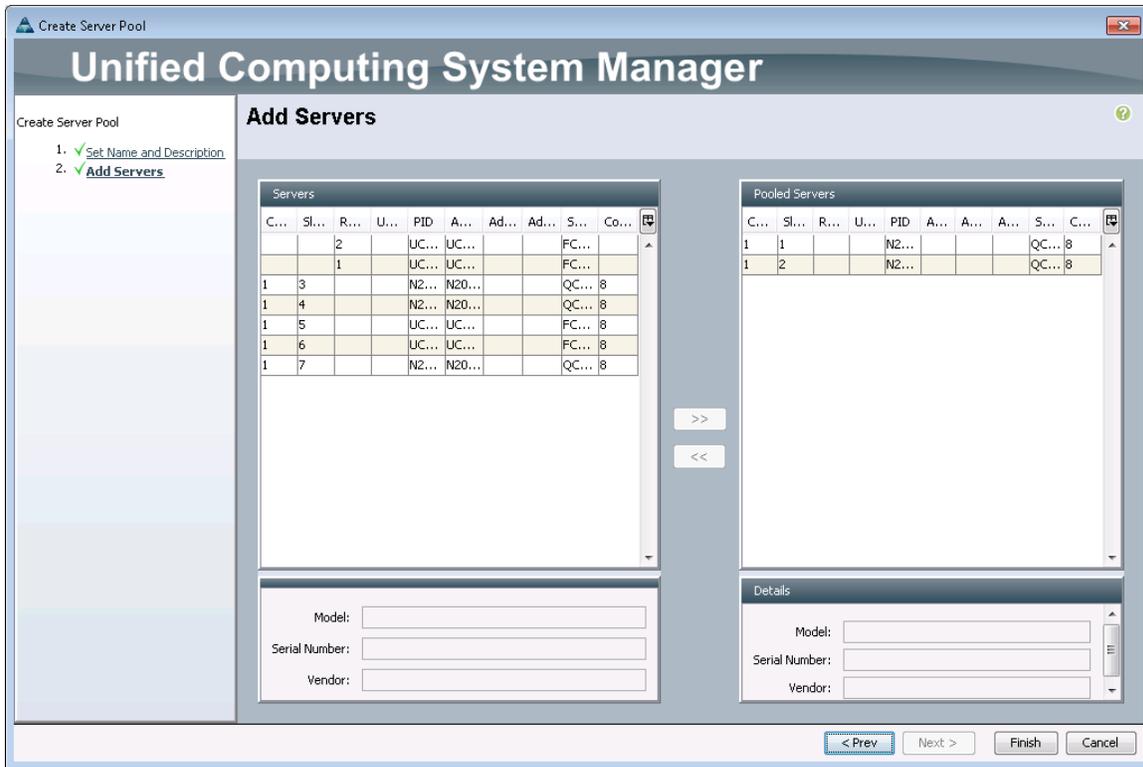
Note: Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `Infra_Pool` as the name for server pool.
6. Optional: Enter a description for the server pool.



The screenshot shows the 'Create Server Pool' wizard in the Cisco Unified Computing System Manager. The window title is 'Create Server Pool'. The main heading is 'Unified Computing System Manager'. The current step is 'Set Name and Description'. On the left, a progress pane shows two steps: '1. Set Name and Description' (completed with a checkmark) and '2. Add Servers' (pending). The main area contains a 'Name:' field with the value 'Infra_Pool' and a 'Description:' field which is currently empty. At the bottom right, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

7. Click Next.
8. Select a server and click >> to add it to the `Infra_Pool` server pool.



9. Click Finish.
10. Click OK.

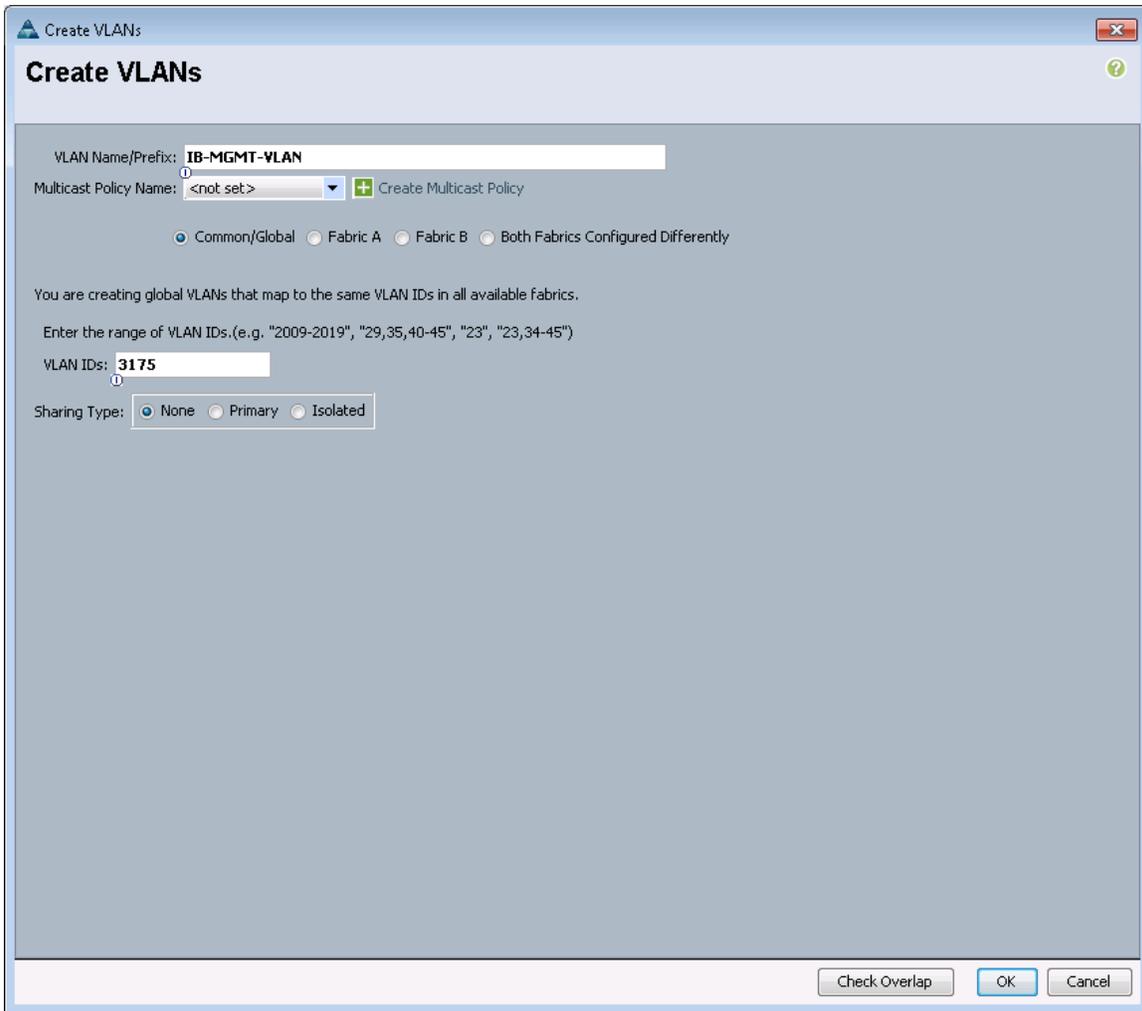
Create VLANs

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

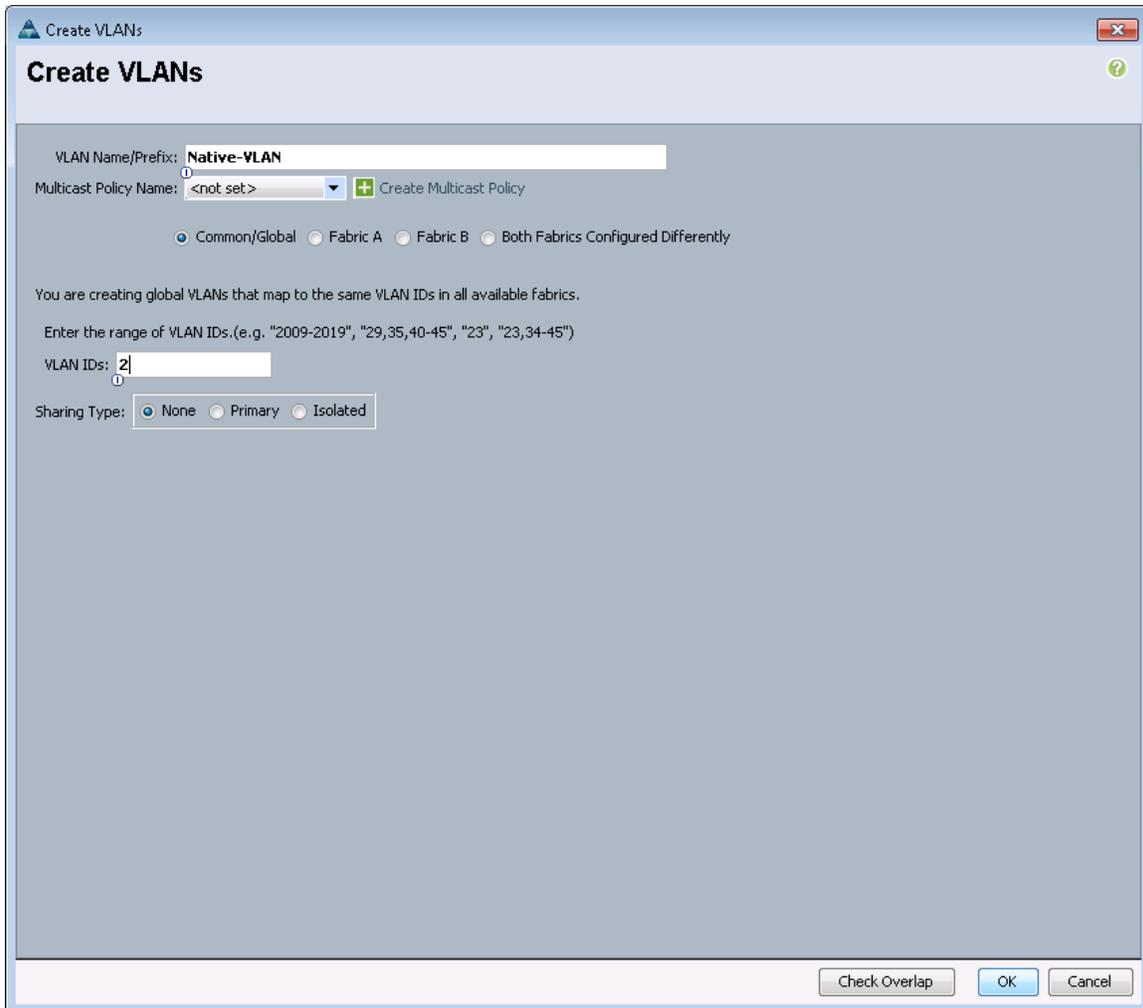
1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

Note: In this procedure, five VLANs are created.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `IB-MGMT-VLAN` as the name for VLAN to be used for management traffic.
6. Retain the Common/Global option selected for the scope of the VLAN.
7. Enter `<<var_ib-mgmt_vlan_id>>` as the ID of the management VLAN.
8. Retain the Sharing Type as None.



9. Click OK, and then click OK again.
10. Right-click VLANs.
11. Select Create VLANs.
12. Enter `Native-VLAN` as the name for the VLAN to be used as the native VLAN.
13. Keep the Common/Global option selected for the scope of the VLAN.
14. Enter the `<<var_native_vlan_id>>` as the ID for the native VLAN.
15. Keep the Sharing Type as None.



16. Click OK, and then click OK again.
17. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and select Set as Native VLAN.
18. Click Yes, and then click OK.

Create VSANs and FCoE Port Channels

To configure the necessary virtual storage area networks (VSANs) and FCoE uplink port channels for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Expand the SAN > SAN Cloud tree.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter `VSAN_A` as the name for the VSAN for Fabric A.
6. Keep the Disabled option selected for FC Zoning.
7. Select Fabric A.
8. Enter `<<var_vsan_a_id>>` as the VSAN ID for Fabric A.

9. Enter `<<var_fabric_a_fcoe_vlan_id>>` as the FCoE VLAN ID for Fabric A.

Note: For the FlexPod solution, it is recommended to use the same ID for the VSAN and the FCoE VLAN required for Fabric A.

Create VSAN

Name:

FC Zoning Settings

FC Zoning: Disabled Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating a local VSAN in Fabric A that maps to a VSAN ID that exists only in Fabric A.
Enter the VSAN ID that maps to this VSAN.

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
Enter the VLAN ID that maps to this VSAN.

VSAN ID:

FCoE VLAN:

OK Cancel

10. Click OK, and then click OK again to create the VSAN.

11. Right-click VSANs.

12. Select Create VSAN.

13. Enter `VSAN_B` as the name for the VSAN for Fabric B.

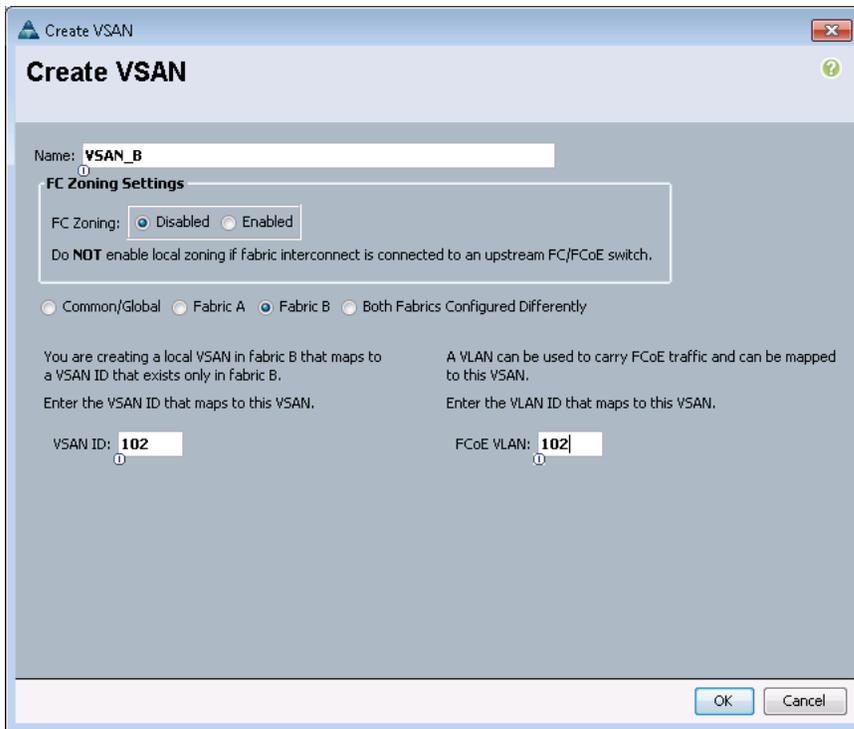
14. Retain the Disabled option selected for FC Zoning.

15. Select Fabric B.

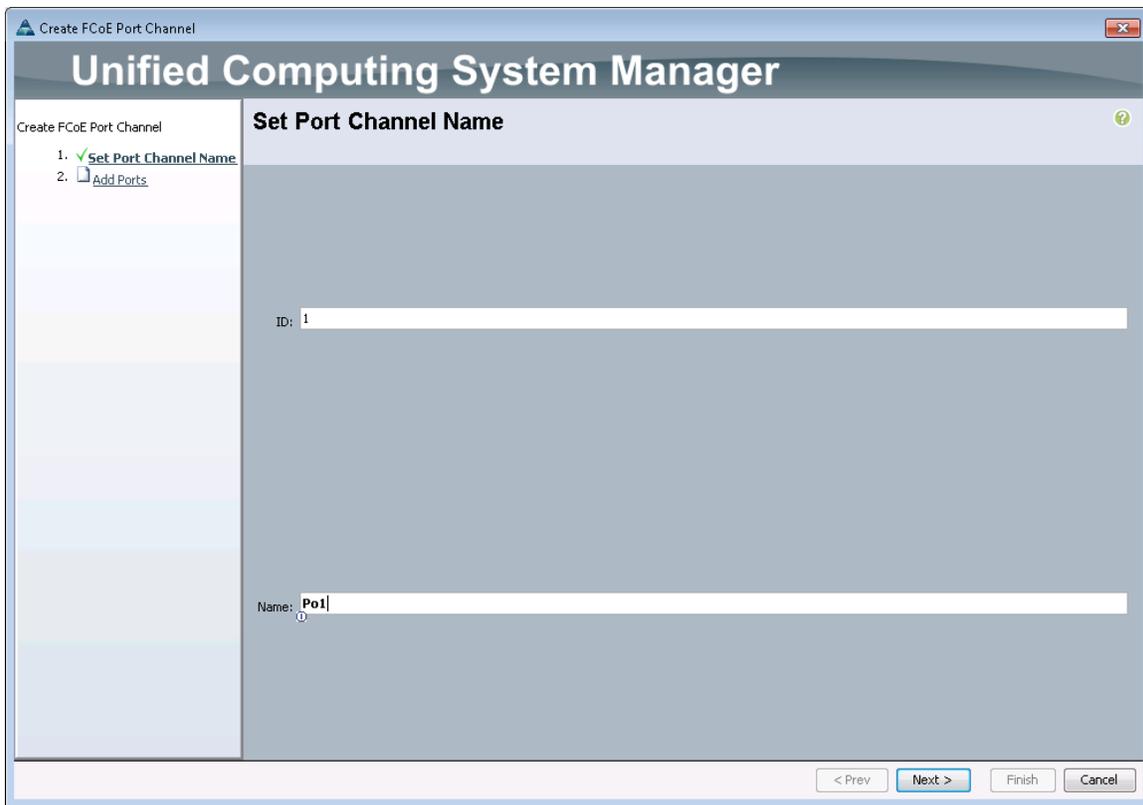
16. Enter `<<var_vsan_b_id>>` as the VSAN ID for Fabric B.

17. Enter `<<var_fabric_b_fcoe_vlan_id>>` as the FCoE VLAN ID for Fabric B.

Note: NetApp recommends using the same ID for the VSAN and the FCoE VLAN required for Fabric B.

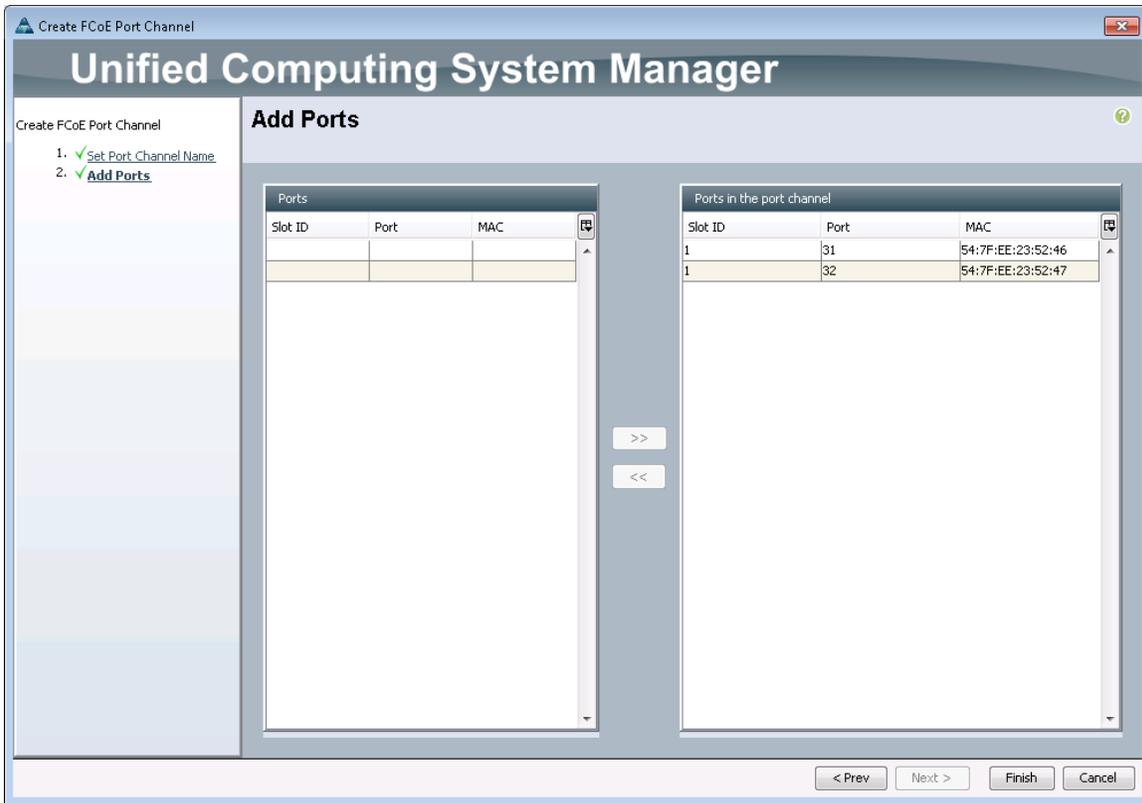


18. Click OK, and then click OK again to create the VSAN.
19. In the navigation pane, under SAN > SAN Cloud, expand the Fabric A tree.
20. Right-click FCoE Port Channels.
21. Select Create FCoE Port Channel.
22. Enter 1 for the port channel ID and P01 for the port channel name.



23. Click Next.

24. Select ports 31 and 32 and click >> to add the ports to the port channel.

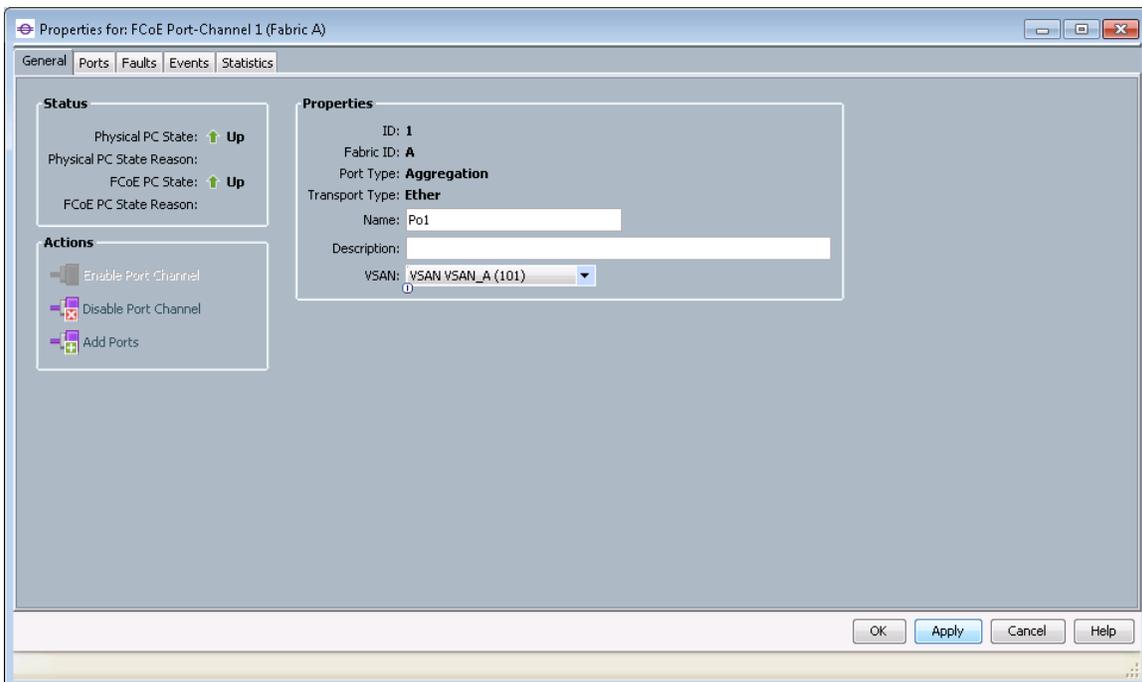


25. Click Finish.

26. Select the checkbox for Show Navigator for FCoE Port-Channel 1 (Fabric A).

27. Click OK to create the port channel.

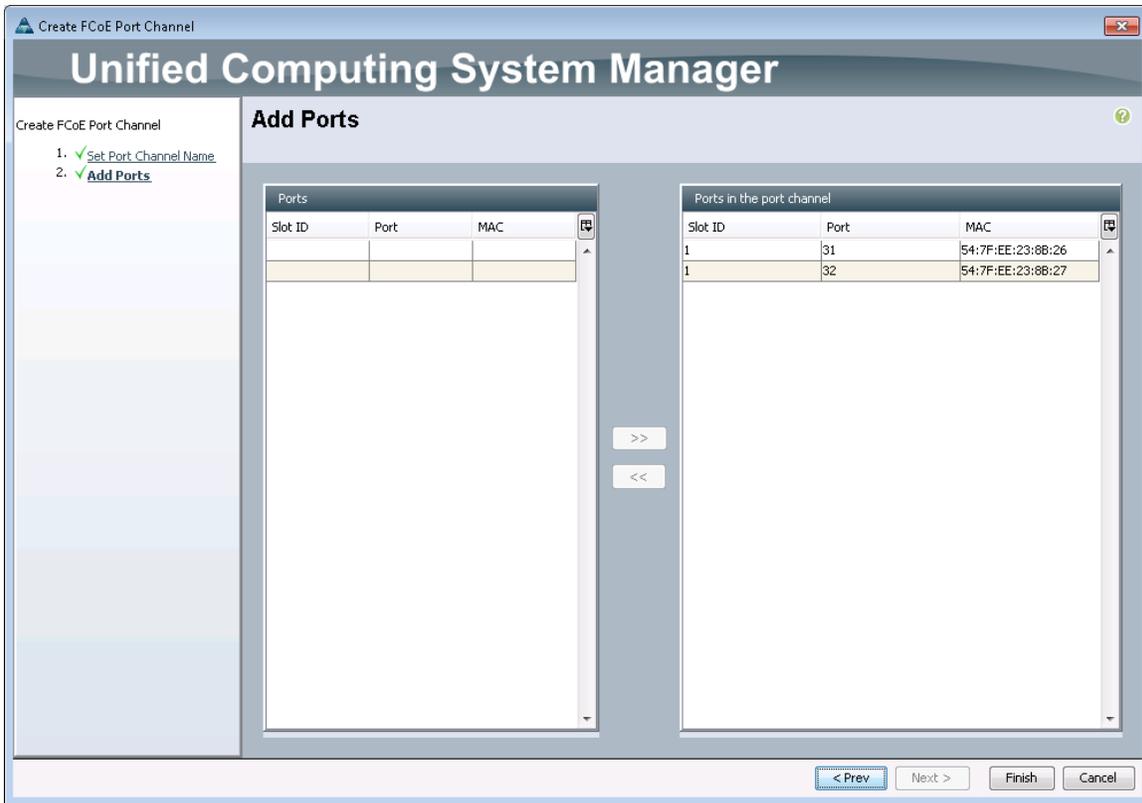
28. In the right pane, under Properties, select VSAN VSAN_A for Fabric A in the VSAN list.



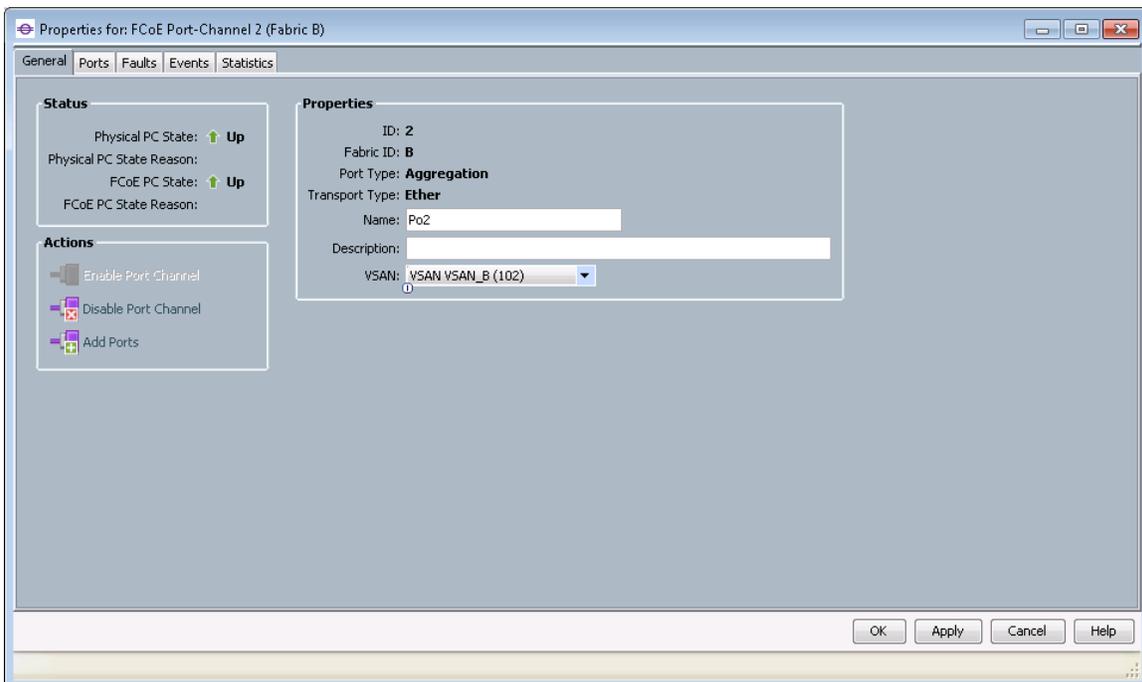
29. Click Apply, and then click OK.
30. Click OK to close the navigator.
31. In the navigation pane, under SAN > SAN Cloud, expand the Fabric B tree.
32. Right-click FCoE Port Channels.
33. Select Create FCoE Port Channel.
34. Enter 2 for the port channel ID and Po2 for the port channel name.

The screenshot shows a web-based configuration window titled "Create FCoE Port Channel" within the "Unified Computing System Manager". The interface is split into a left sidebar and a main content area. The sidebar contains a progress indicator with two steps: "1. Set Port Channel Name" (marked with a green checkmark) and "2. Add Ports". The main content area is titled "Set Port Channel Name" and features two text input fields. The first field is labeled "ID:" and contains the value "2". The second field is labeled "Name:" and contains the value "Po2". At the bottom of the window, there are four buttons: "< Prev", "Next >", "Finish", and "Cancel". The "Next >" button is highlighted in blue, indicating it is the active step.

35. Click Next.
36. Select ports 31 and 32 and click >> to add the ports to the port channel.



37. Click Finish.
38. Select the checkbox for Show Navigator for FCoE Port-Channel 2 (Fabric B).
39. Click OK to create the port channel.
40. In the right pane, under Properties, select VSAN VSAN_B for Fabric B.



41. Click Apply, and then click OK.
42. Click OK to close the navigator.

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 adapters, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

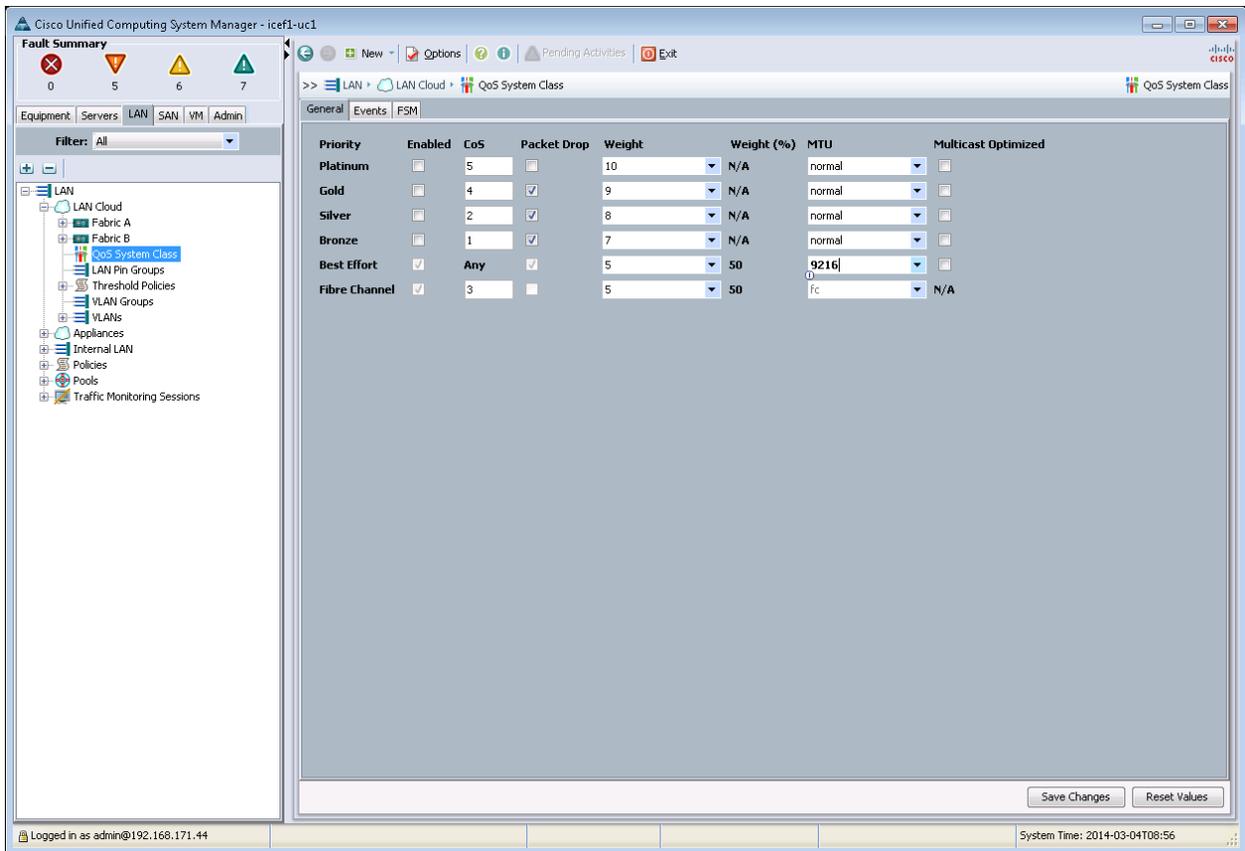
To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter `RHEL-Server` as the name for the host firmware package.
6. Keep Simple selected.
7. Select the version 2.2(2c) for both the Blade and Rack Packages.
8. Click OK to create the host firmware package.
9. Click OK.

Set Jumbo Frames in Cisco UCS Fabric

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

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

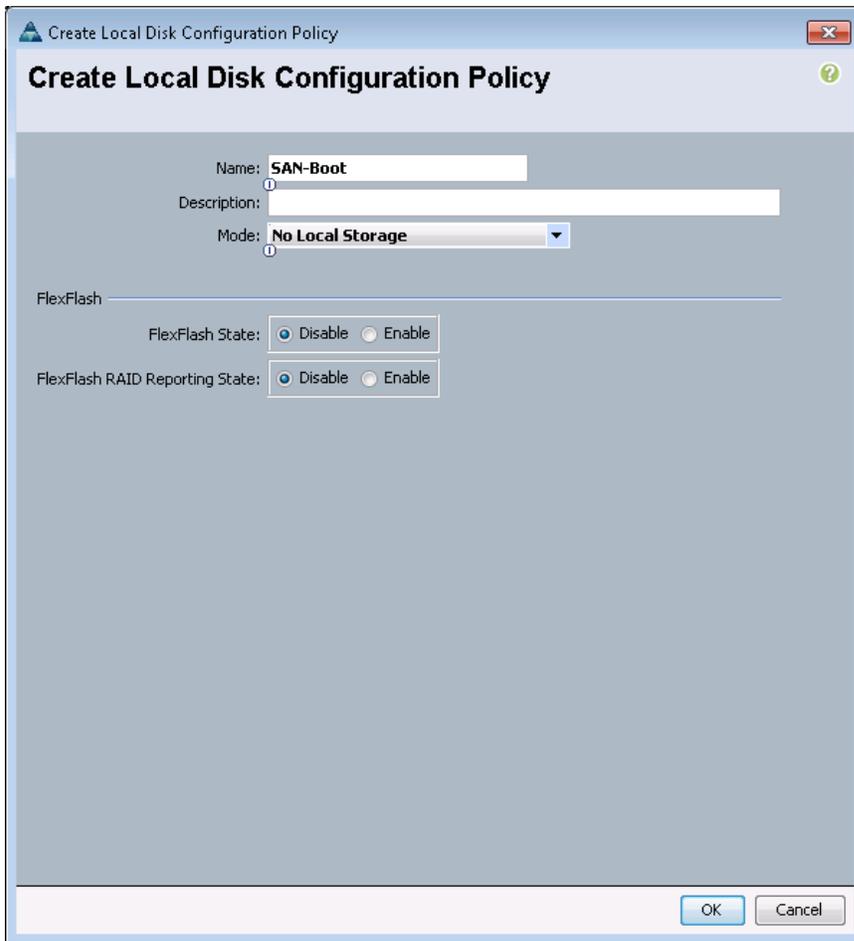
Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.

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

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

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > root.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter `SAN-Boot` as the local disk configuration policy name.
6. Change the Mode to No Local Storage.
7. Keep FlexFlash State and FlexFlash Raid Reporting State set to Disable.

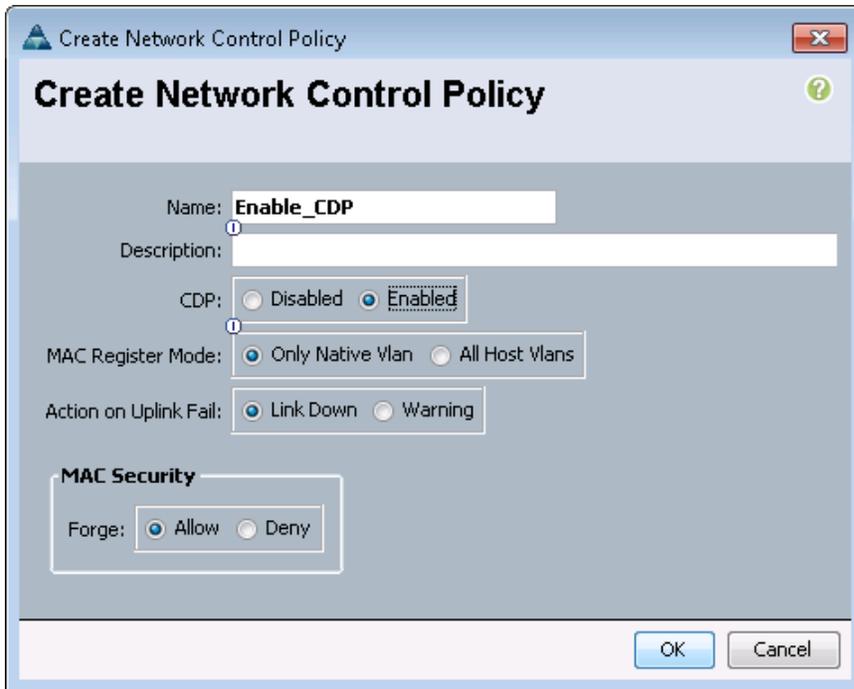


8. Click OK to create the local disk configuration policy.
9. Click OK.

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, in the navigation pane, click the LAN tab.
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.

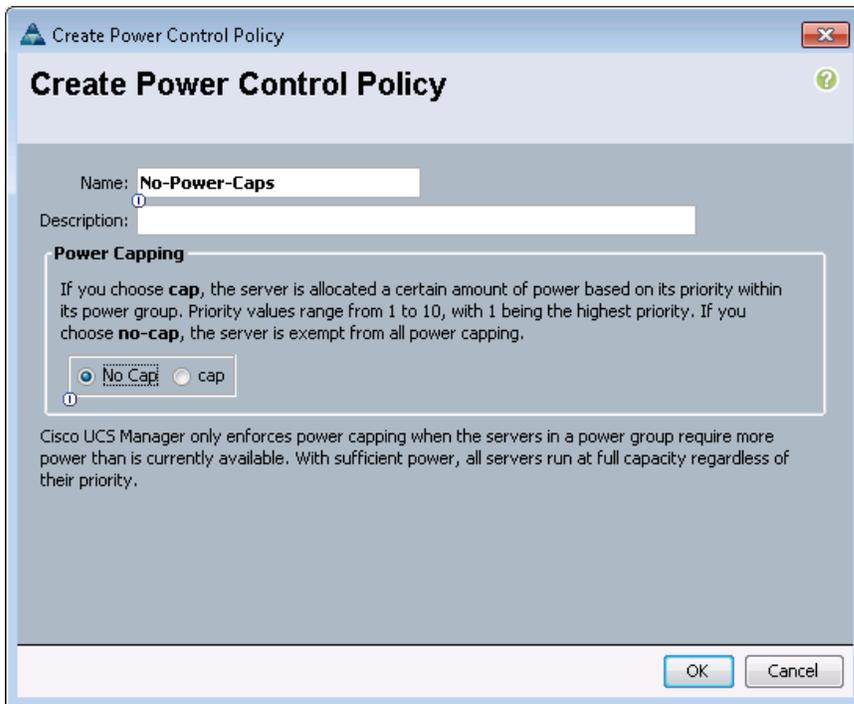


7. Click OK to create the network control policy.
8. Click OK.

Create Power Control Policy

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

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
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.



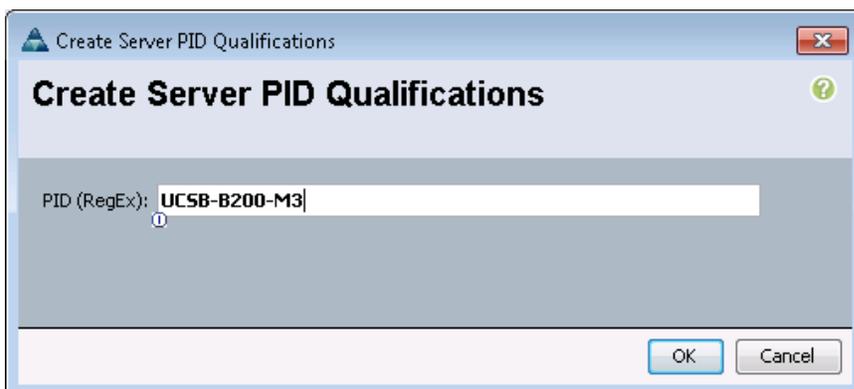
7. Click OK to create the power control policy.
8. Click OK.

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 B200-M3 server.

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Enter UCSB-B200-M3 as the name for the policy.
6. In the left pane, under Actions Select Create Server PID Qualifications.
7. Enter UCSB-B200-M3 as the PID.



8. Click OK to create the server pool qualification policy.
9. Click OK, and then click OK again.

Create Server BIOS Policy

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

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `RHEL-Server` 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

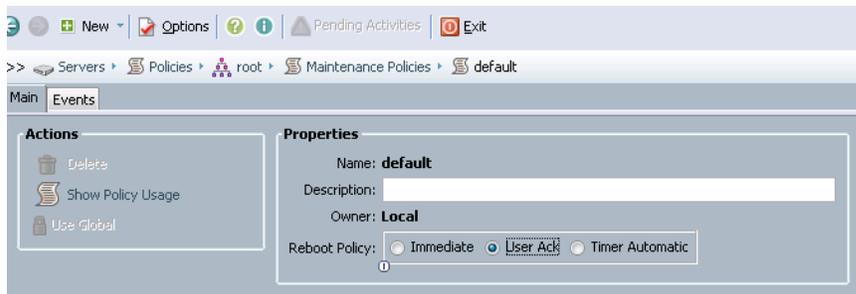
To create a vNIC/vHBA placement policy for the infrastructure hosts, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter `RHEL-Server` as the name for the placement policy.
6. Click 1 and under the Selection Preference 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, in the navigation pane, click the Servers tab.
2. Choose Policies > root.
3. Choose Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.



5. Click Save Changes.
6. Click OK to acknowledge the change.

Create vNIC Templates for Fabric A and Fabric B

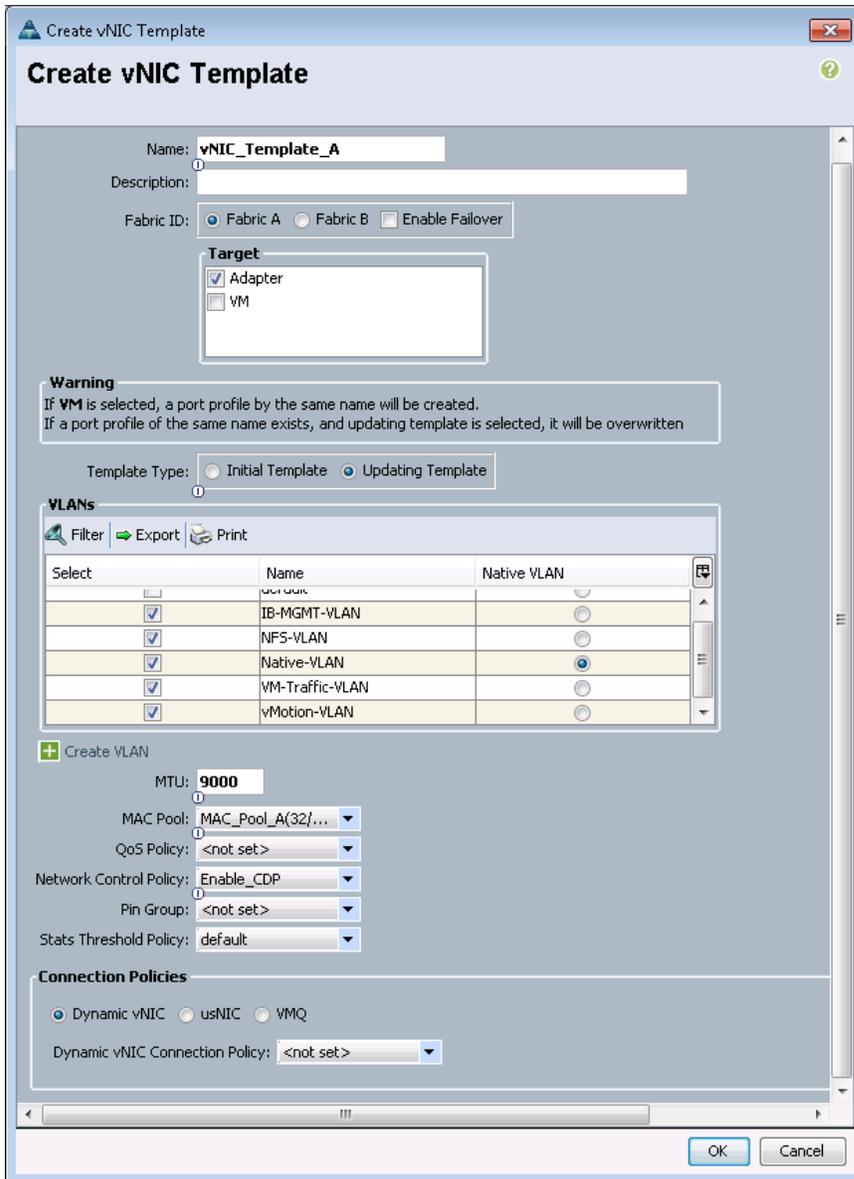
To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
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. For Fabric ID, select Fabric A.

Note: Select the Enable Failover checkbox.

Note: Under Target, do not select the VM checkbox.

7. Select Updating Template as the Template Type.
8. Under VLANs, select the checkbox for `IB-MGMT-VLAN`.
9. Set `Native-VLAN` as the native VLAN.
10. For MTU, enter `9000`.
11. In the MAC Pool list, select `MAC_Pool_A`.
12. In the Network Control Policy list, select `Enable_CDP`.



13. Click OK to create the vNIC template.
 14. Click OK.
 15. In the navigation pane, click the LAN tab.
 16. Select Policies > root.
 17. Right-click vNIC Templates.
 18. Select Create vNIC Template.
 19. Enter vNIC_Template_B as the vNIC template name.
 20. For Fabric ID, select Fabric B.
- Note:** Select the Enable Failover checkbox.
- Note:** Under Target, do not select the VM checkbox.
21. Select Updating Template as the Template Type.

22. Under VLANs, select the checkbox for `IB-MGMT-VLAN`.
23. Set `Native-VLAN` as the native VLAN.
24. For MTU, enter `9000`.
25. In the MAC Pool list, select `MAC_Pool_B`.
26. In the Network Control Policy list, select `Enable_CDP`.
27. Click OK to create the vNIC template.
28. Click OK.

Create vHBA Templates for Fabric A and Fabric B

To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the SAN tab.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter `vHBA_Template_A` as the vHBA template name.
6. Select A for Fabric ID.
7. In the Select VSAN list, select `VSAN_A`.
8. In the WWPN Pool list, select `WWPN_Pool_A`.

The screenshot shows a 'Create vHBA Template' dialog box with the following fields and values:

- Name: `vHBA_Template_A`
- Description: (empty)
- Fabric ID: A B
- Select VSAN: `VSAN_A` (with a '+ Create VSAN' button)
- Template Type: Initial Template Updating Template
- Max Data Field Size: `2048`
- WWPN Pool: `WWPN_Pool_A(32/32)`
- QoS Policy: `<not set>`
- Pin Group: `<not set>`
- Stats Threshold Policy: `default`

Buttons: OK, Cancel

9. Click OK to create the vHBA template.
10. Click OK.
11. In the navigation pane, click the SAN tab.
12. Select Policies > root.
13. Right-click vHBA Templates.

14. Select Create vHBA Template.
15. Enter `vHBA_Template_B` as the vHBA template name.
16. Select B for Fabric ID.
17. In the Select VSAN list, select `VSAN_B`.
18. In the WWPN Pool, select `WWPN_Pool_B`.

The screenshot shows a 'Create vHBA Template' dialog box with the following fields and values:

- Name: `vHBA_Template_B`
- Description: (empty)
- Fabric ID: A B
- Select VSAN: `VSAN_B` (with a '+ Create VSAN' button)
- Template Type: Initial Template Updating Template
- Max Data Field Size: `2048`
- WWPN Pool: `WWPN_Pool_B(32/32)`
- QoS Policy: `<not set>`
- Pin Group: `<not set>`
- Stats Threshold Policy: `default`

Buttons: OK, Cancel

19. Click OK to create the vHBA template.
20. Click OK.

Create Boot Policies

This procedure applies to a Cisco UCS environment in which two FCoE logical interfaces (LIFs) are on cluster node 1 (`fcp_lif01a` and `fcp_lif01b`) and two FCoE LIFs are on cluster node 2 (`fcp_lif02a` and `fcp_lif02b`). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco Nexus 5548UP A) and the B LIFs are connected to Fabric B (Cisco Nexus 5548UP B).

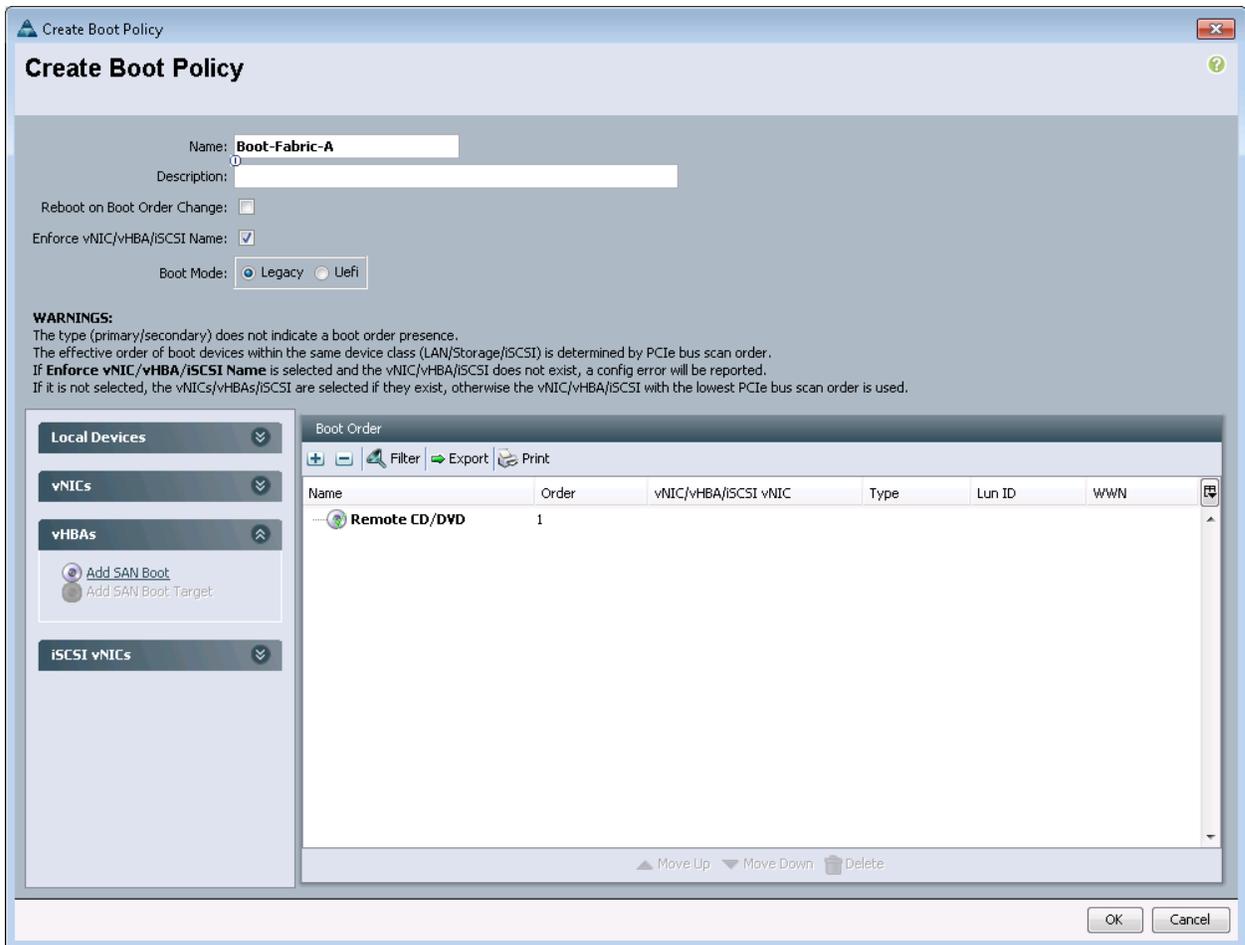
Two boot policies are configured in this procedure. The first policy configures the primary target to be `fcp_lif01a` and the second boot policy configures the primary target to be `fcp_lif01b`.

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

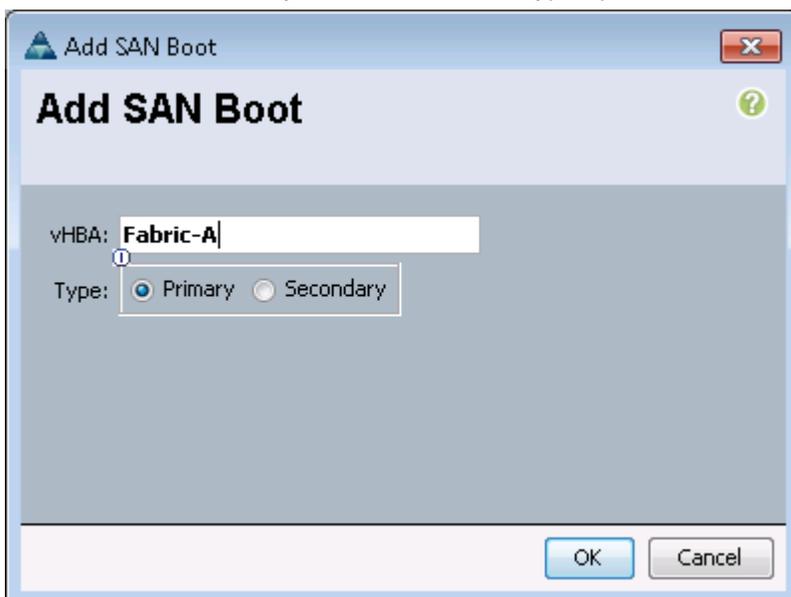
1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-Fabric-A` as the name for the boot policy.
6. Optional: Enter a description for the boot policy.

Note: Do not select the Reboot on Boot Order Change checkbox.

7. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.



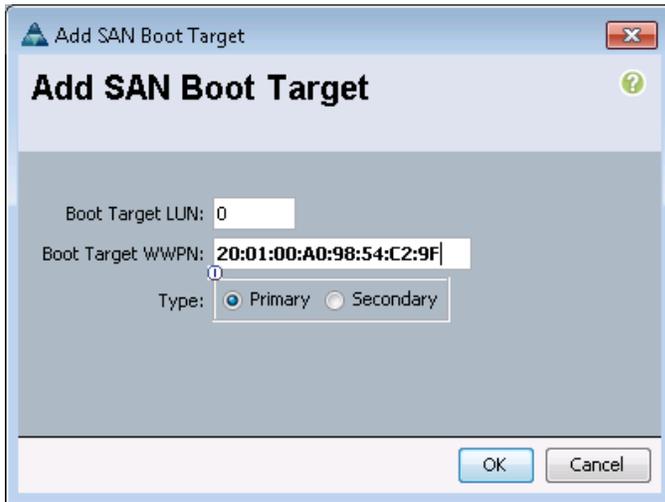
8. Expand the vHBAs drop-down menu and select Add SAN Boot.
9. In the Add SAN Boot dialog box, enter `Fabric-A` in the vHBA field.
10. Confirm that Primary is selected for the Type option.



11. Click OK to add the SAN boot initiator.
12. From the vHBA drop-down menu, select Add SAN Boot Target.
13. Keep 0 as the value for Boot Target LUN.
14. Enter the WWPN for `fcplif01a`.

Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.

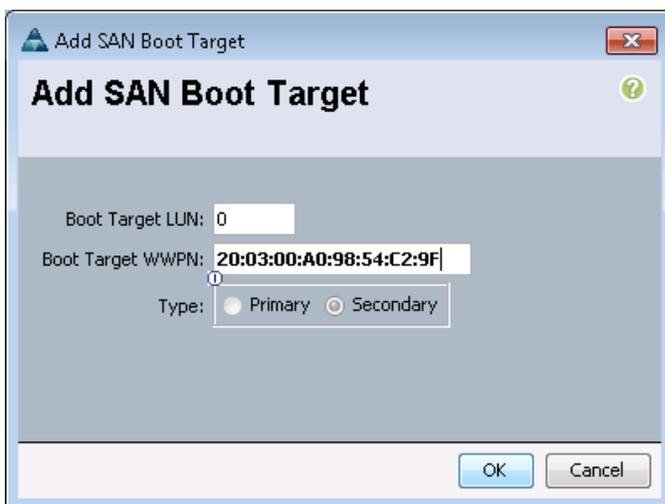
15. Select Primary for the SAN boot target type.



The screenshot shows a dialog box titled "Add SAN Boot Target". It has a title bar with a close button. The main area contains three input fields: "Boot Target LUN:" with the value "0", "Boot Target WWPN:" with the value "20:01:00:A0:98:54:C2:9F", and "Type:" with two radio buttons, "Primary" (selected) and "Secondary". At the bottom right, there are "OK" and "Cancel" buttons.

16. Click OK to add the SAN boot target.
17. From the vHBA drop-down menu, select Add SAN Boot Target.
18. Enter 0 as the value for Boot Target LUN.
19. Enter the WWPN for `fcplif02a`.

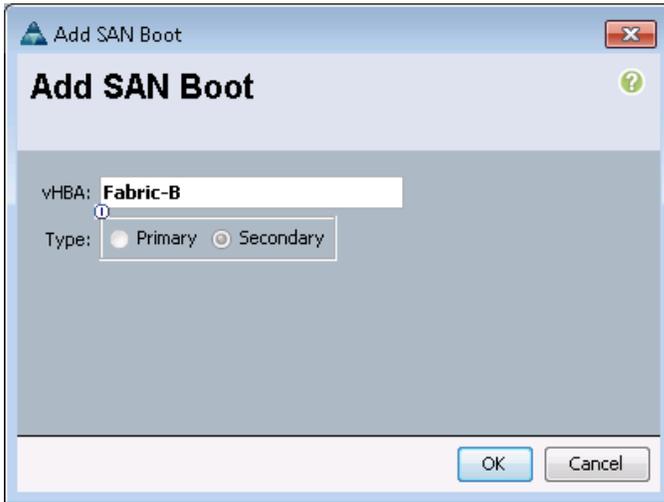
Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.



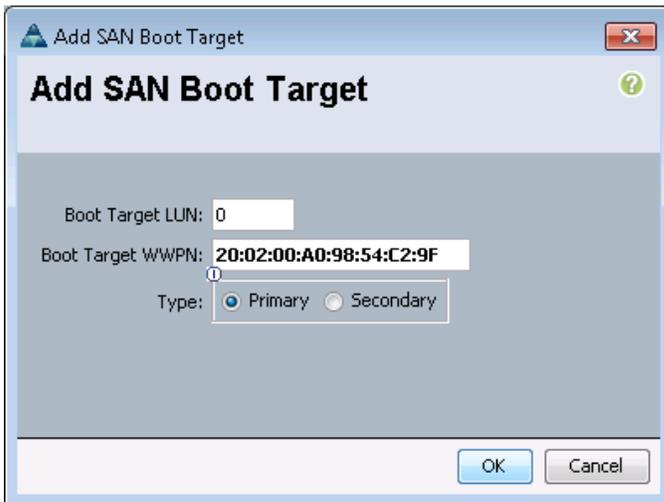
The screenshot shows a dialog box titled "Add SAN Boot Target". It has a title bar with a close button. The main area contains three input fields: "Boot Target LUN:" with the value "0", "Boot Target WWPN:" with the value "20:03:00:A0:98:54:C2:9F", and "Type:" with two radio buttons, "Primary" and "Secondary" (selected). At the bottom right, there are "OK" and "Cancel" buttons.

20. Click OK to add the SAN boot target.
21. From the vHBA drop-down menu, select Add SAN Boot.

22. In the Add SAN Boot dialog box, enter `Fabric-B` in the vHBA box.
23. The SAN boot type should automatically be set to `Secondary`, and the `Type` option should be unavailable.

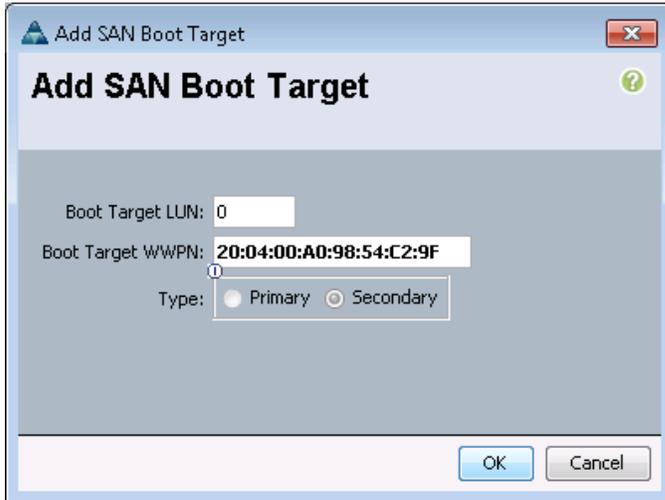


24. Click `OK` to add the SAN boot initiator.
25. From the vHBA drop-down menu, select `Add SAN Boot Target`.
26. Keep `0` as the value for `Boot Target LUN`.
27. Enter the WWPN for `fcplif01b`.
Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.
28. Select `Primary` for the SAN boot target type.

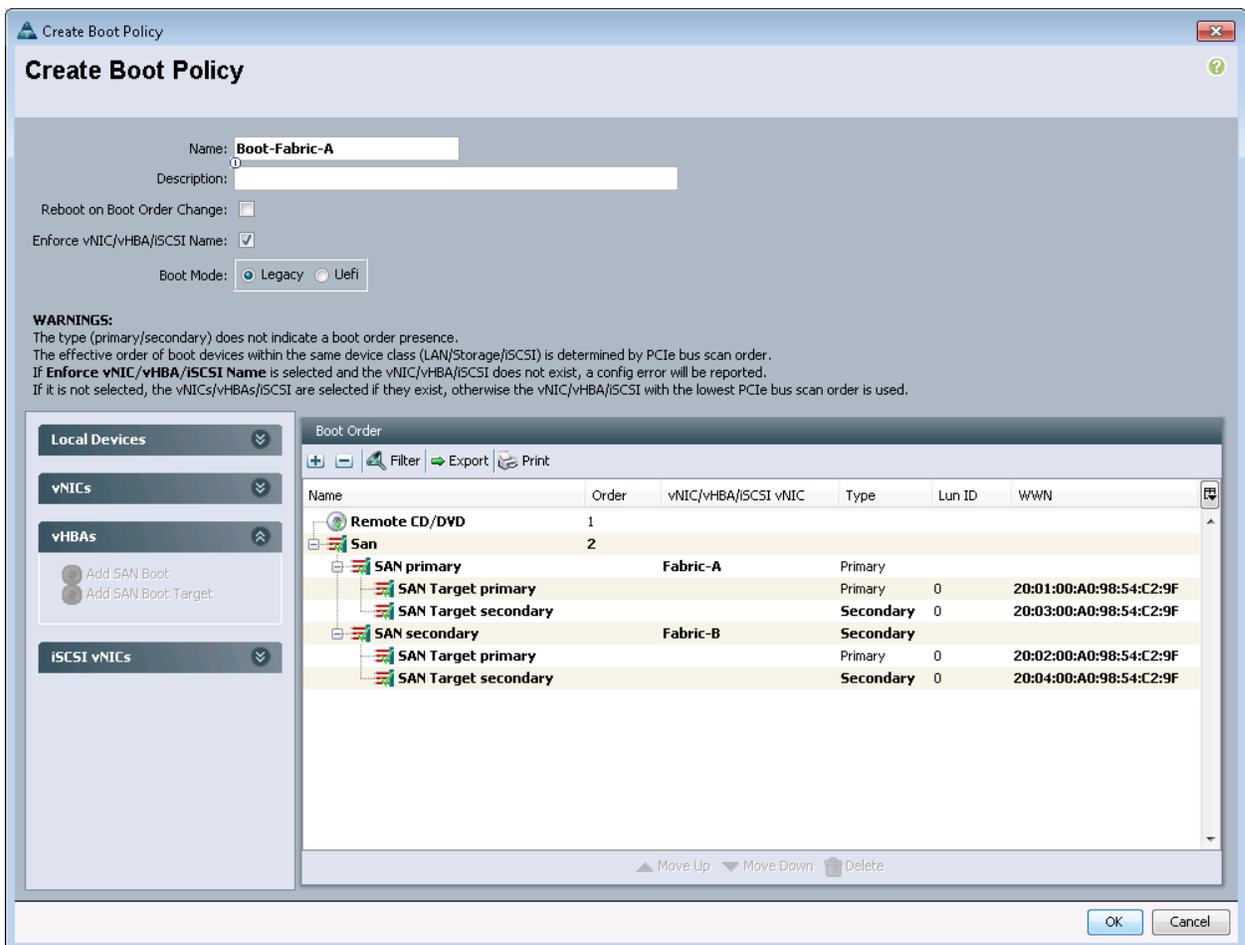


29. Click `OK` to add the SAN boot target.
30. From the vHBA drop-down menu, select `Add SAN Boot Target`.
31. Keep `0` as the value for `Boot Target LUN`.
32. Enter the WWPN for `fcplif02b`.

Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.

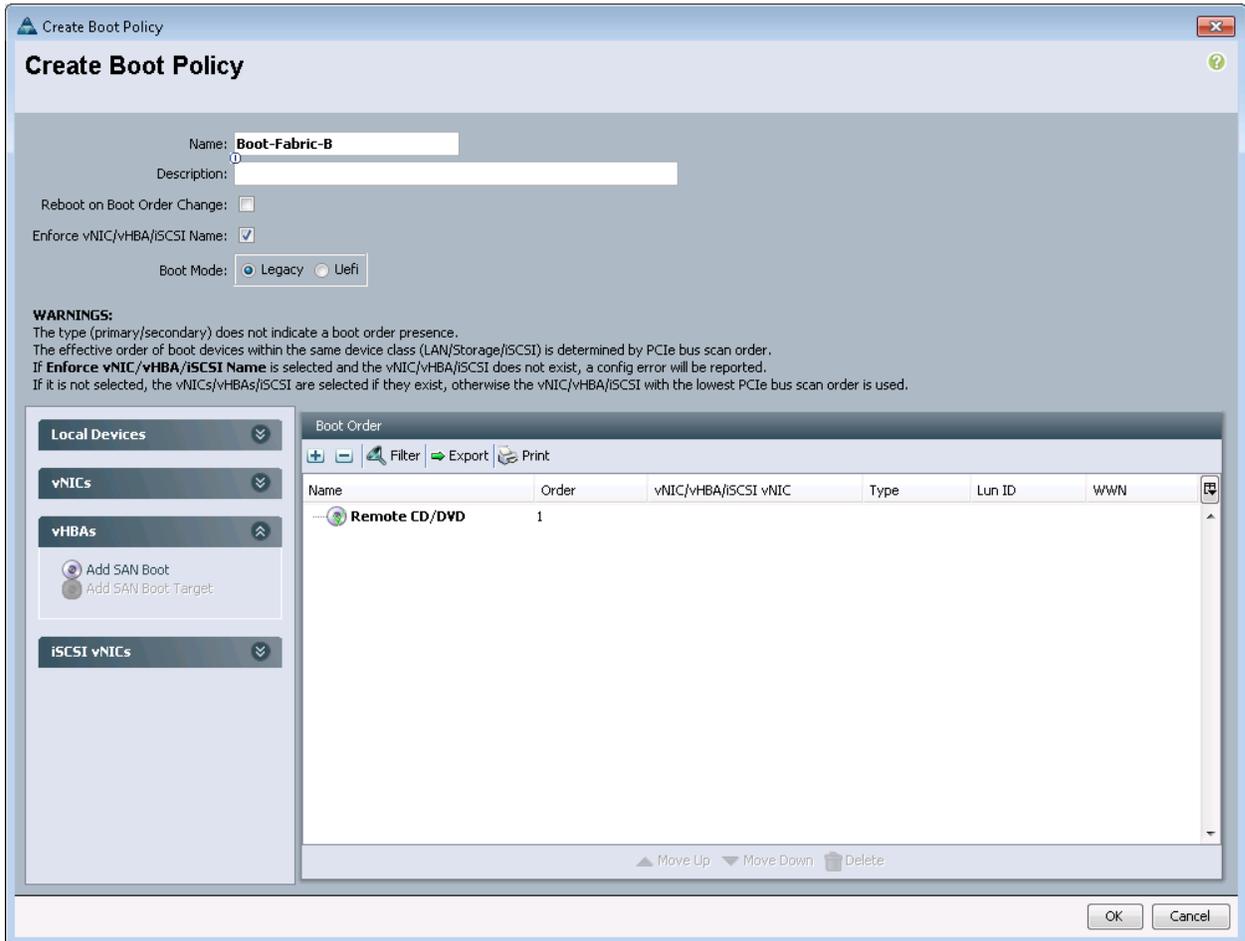


33. Click OK to add the SAN boot target.

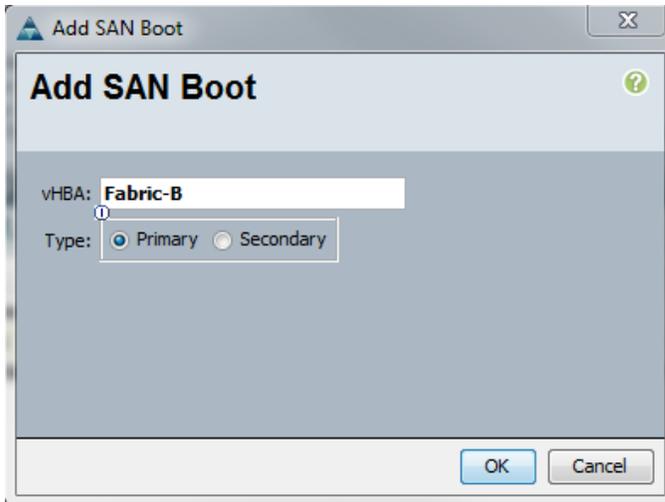


34. Click OK, then click OK again to create the boot policy.

35. Right-click Boot Policies again.
36. Select Create Boot Policy.
37. Enter `Boot-Fabric-B` as the name for the boot policy.
38. Optional: Enter a description of the boot policy.
39. Do not select the Reboot on Boot Order Change option.
40. From the Local Devices drop-down menu, select Add Remote CD/DVD.



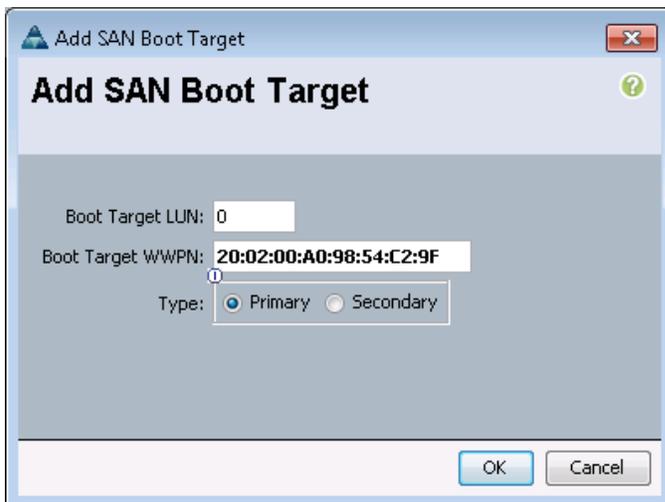
41. From the vHBA drop-down menu, select Add SAN Boot.
42. In the Add SAN Boot dialog box, enter `Fabric-B` in the vHBA box.
43. Confirm that Primary option is selected for the SAN boot type.



44. Click OK to add the SAN boot initiator.
45. From the vHBA drop-down menu, select Add SAN Boot Target.
46. Enter 0 as the value for Boot Target LUN.
47. Enter the WWPN for fcp_lif01b.

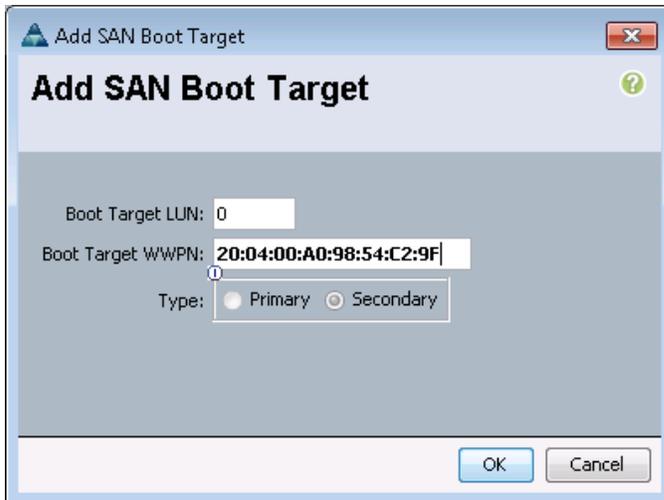
Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.

48. Select Primary as the SAN boot target type.

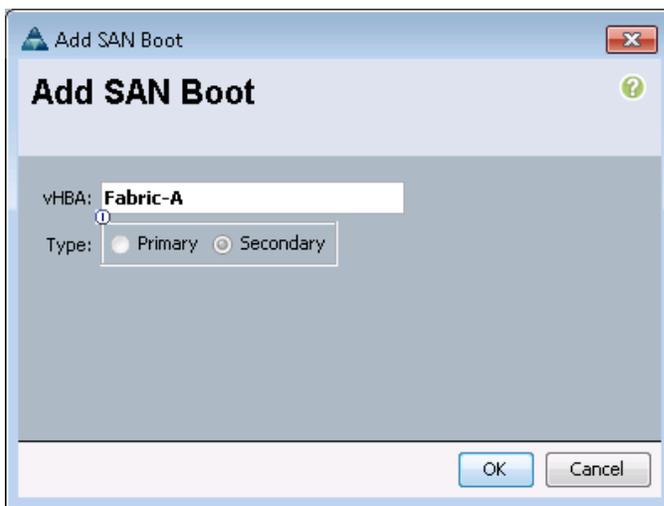


49. Click OK to add the SAN boot target.
50. From the vHBA drop-down menu, select Add SAN Boot Target.
51. Enter 0 as the value for Boot Target LUN.
52. Enter the WWPN for fcp_lif02b.

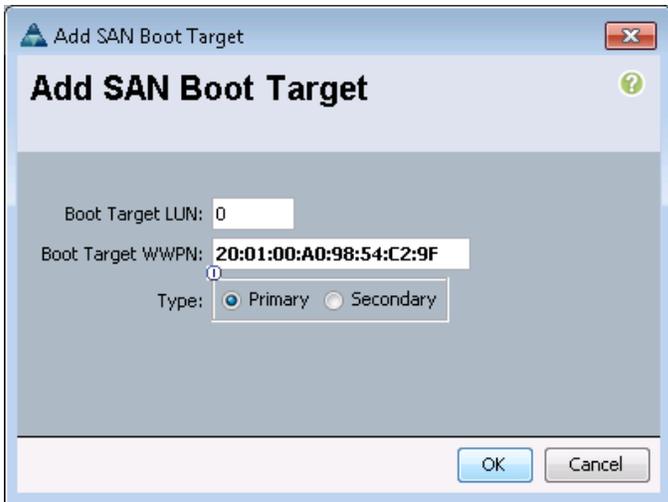
Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.



53. Click OK to add the SAN boot target.
54. From the vHBA menu, select Add SAN Boot.
55. In the Add SAN Boot dialog box, enter `Fabric-A` in the vHBA box.
56. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.



57. Click OK to add the SAN boot initiator.
58. From the vHBA menu, select Add SAN Boot Target.
59. Enter 0 as the value for Boot Target LUN.
60. Enter the WWPN for `fcp_lif01a`.
Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.
61. Select the Primary option for the SAN boot target type.



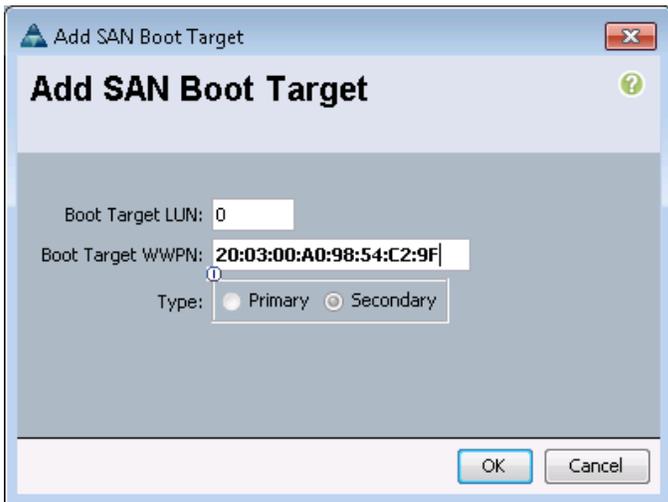
62. Click OK to add the SAN boot target.

63. From the vHBA drop-down menu, select Add SAN Boot Target.

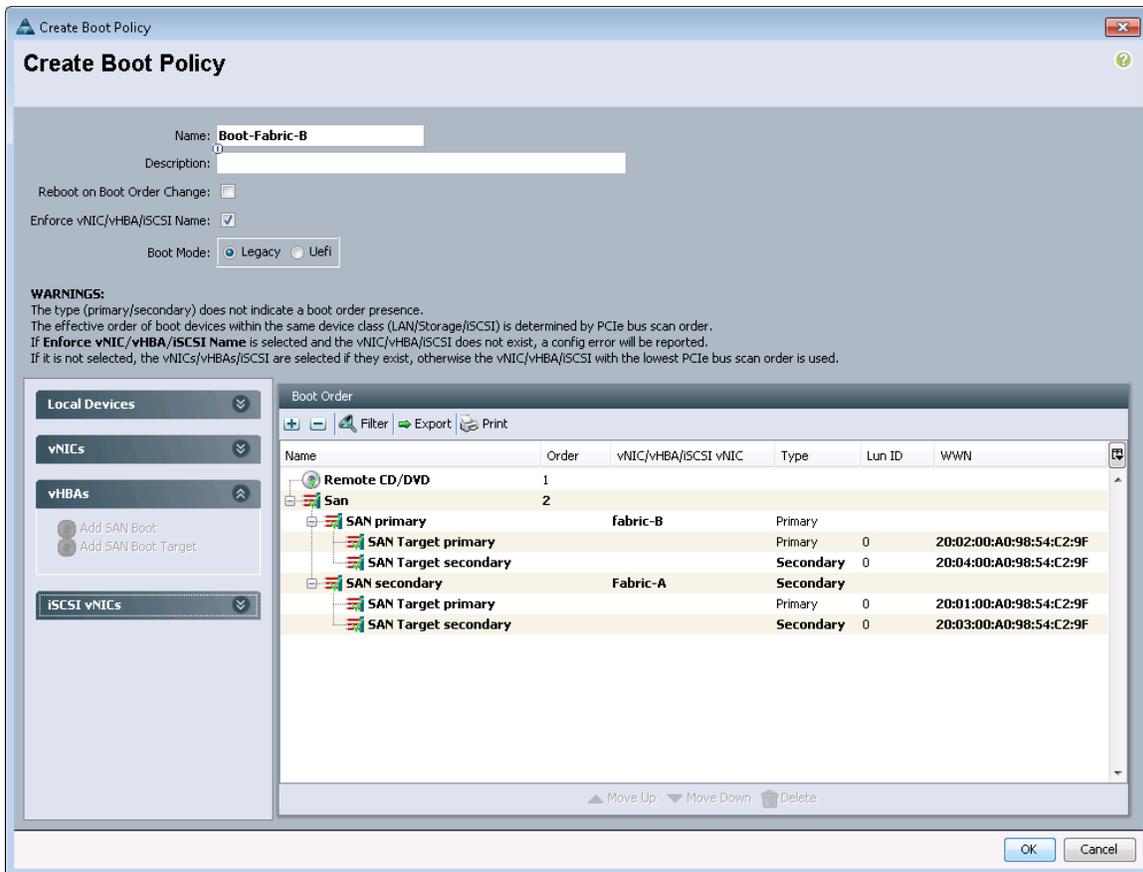
64. Enter 0 as the value for Boot Target LUN.

65. Enter the WWPN for `fcv_lif02a`.

Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.



66. Click OK to add the SAN boot target.

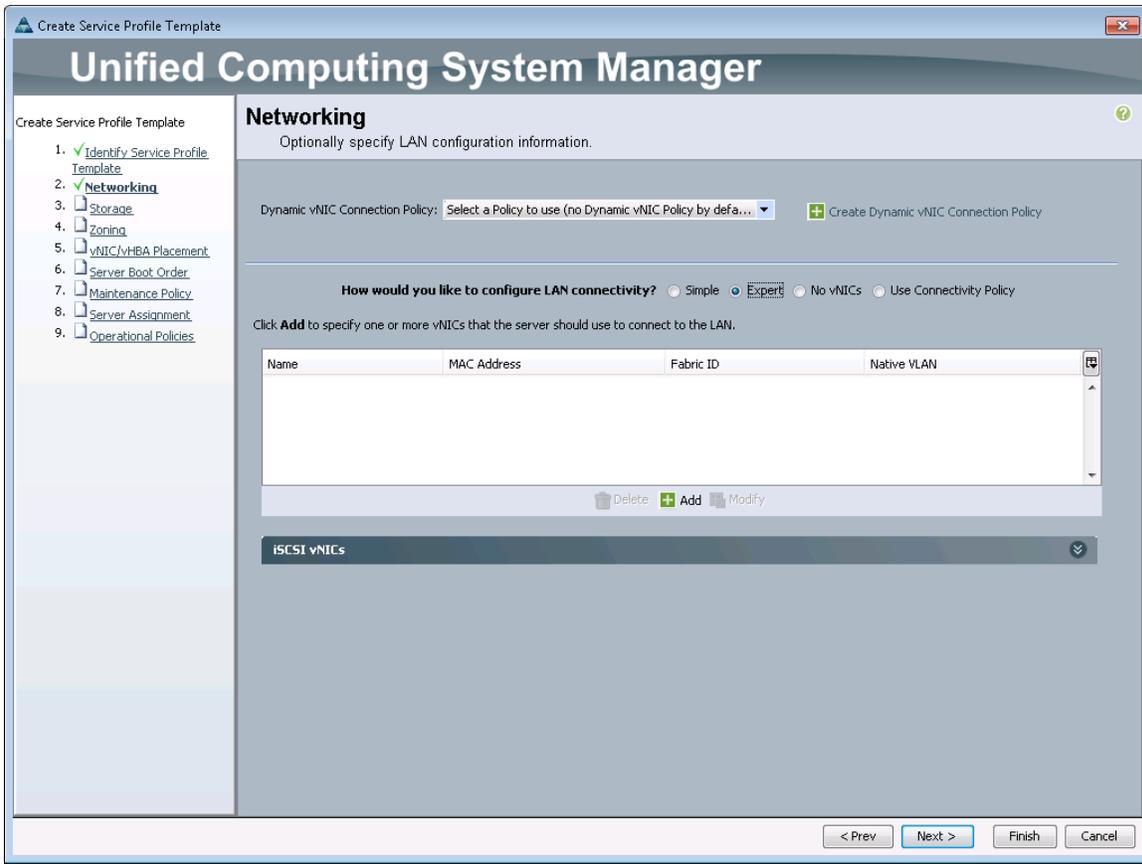


67. Click OK, and then click OK again to create the boot policy.

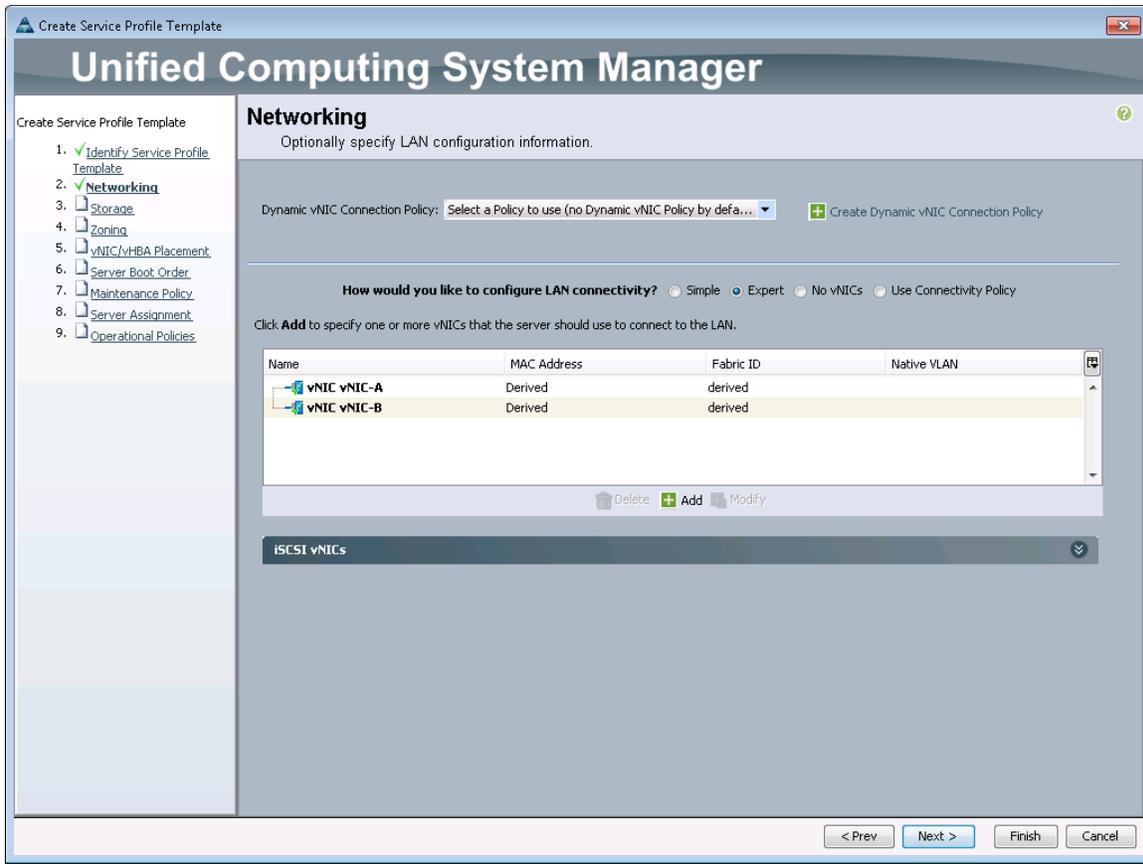
Create Service Profile Template

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

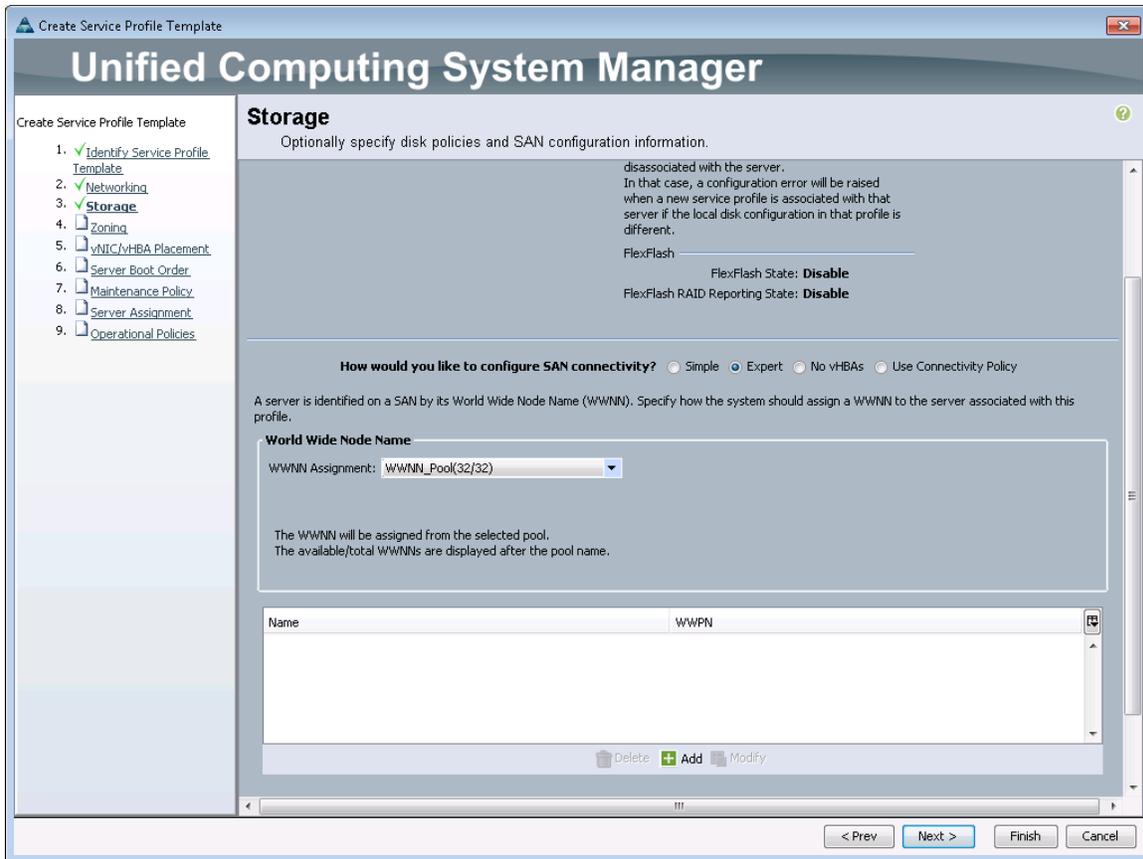
1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the Service Profile Template:
 - a. Enter `RHEL_Server-Fabric-A` as the name for the service profile template. This service profile template is configured to boot from node 1 on Fabric A.
 - b. Select the Updating Template option.
 - c. Under UUID, select `UUID_Pool` as the UUID pool.
 - d. Click Next.
6. Configure the Networking options:
 - a. Retain the default setting for Dynamic vNIC Connection Policy.
 - b. Select the Expert option to configure the LAN connectivity.



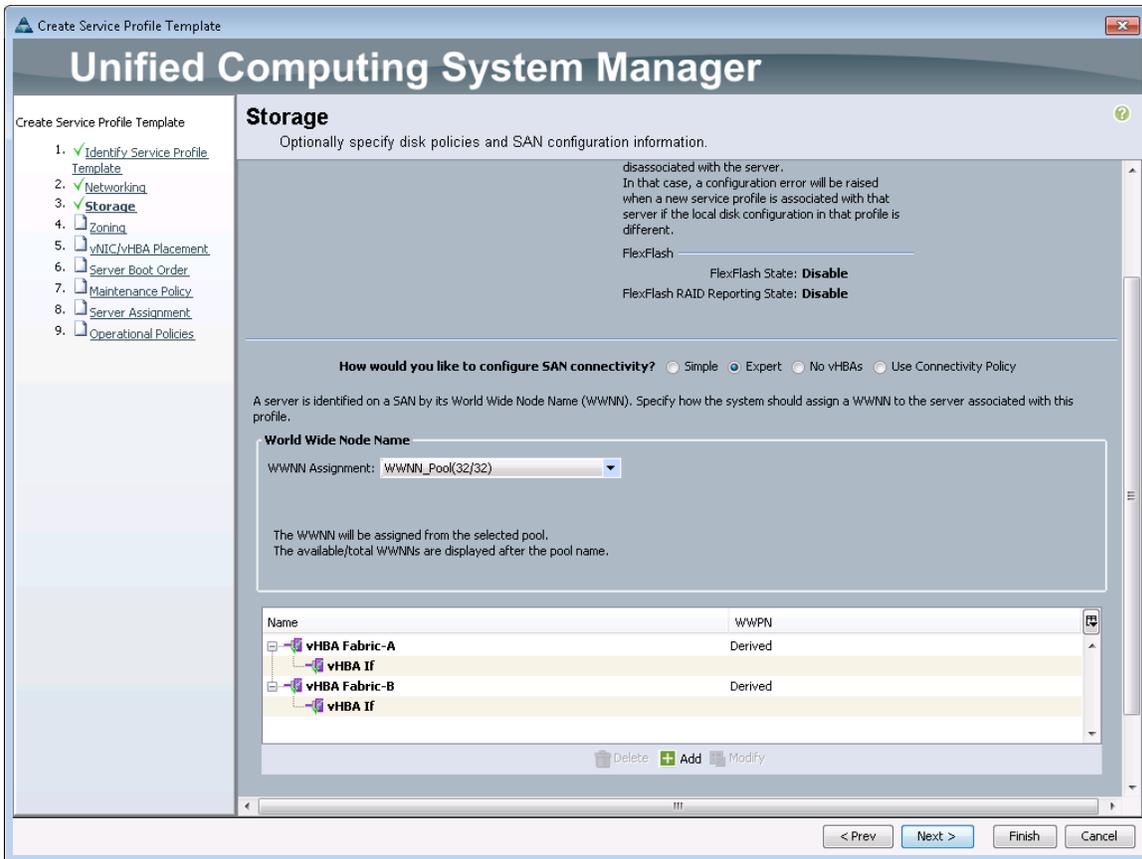
- c. Click Add in the middle of the screen to add a vNIC to the template.
- d. In the Create vNIC dialog box, enter vNIC-A as the name for vNIC.
- e. Select the Use vNIC Template checkbox.
- f. In the vNIC Template list, select vNIC_Template_A.
- g. In the Adapter Policy list, select Linux.
- h. Click OK to add this vNIC to the template.
- i. On the Networking page of the wizard, click the Add button to add another vNIC to the template.
- j. In the Create vNIC box, enter vNIC-B as the name for vNIC.
- k. Select the Use vNIC Template checkbox.
- l. In the vNIC Template list, select vNIC_Template_B.
- m. In the Adapter Policy list, select Linux.
- n. Click OK to add the vNIC to the template.
- o. Review the table on the Networking page to confirm that both vNICs were created.



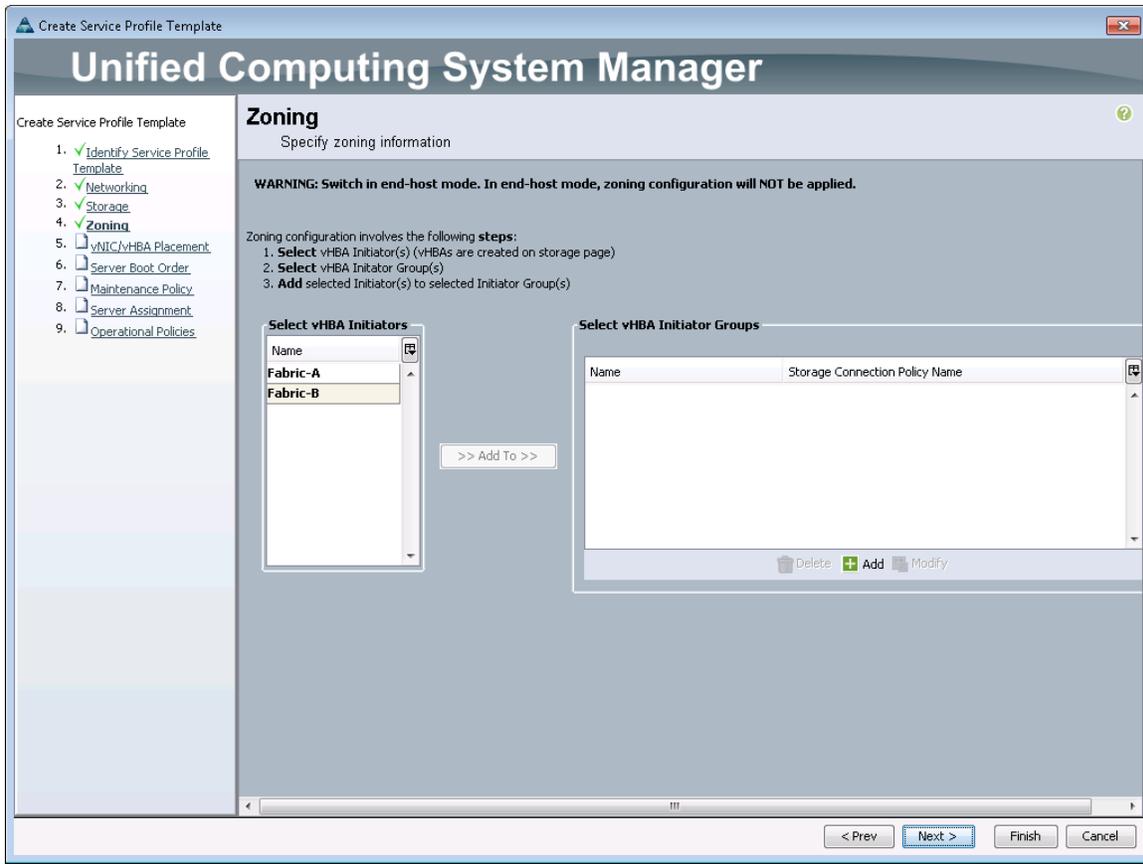
- p. Click Next.
7. Configure the Storage options:
- Select a local disk configuration policy:
 - If the server in question has local disks, select default in the Local Storage list.
 - If the server in question does not have local disks, select SAN-Boot.
 - Select the Expert option to configure the SAN connectivity.
 - In the WWNN Assignment list, select WWNN_Pool.



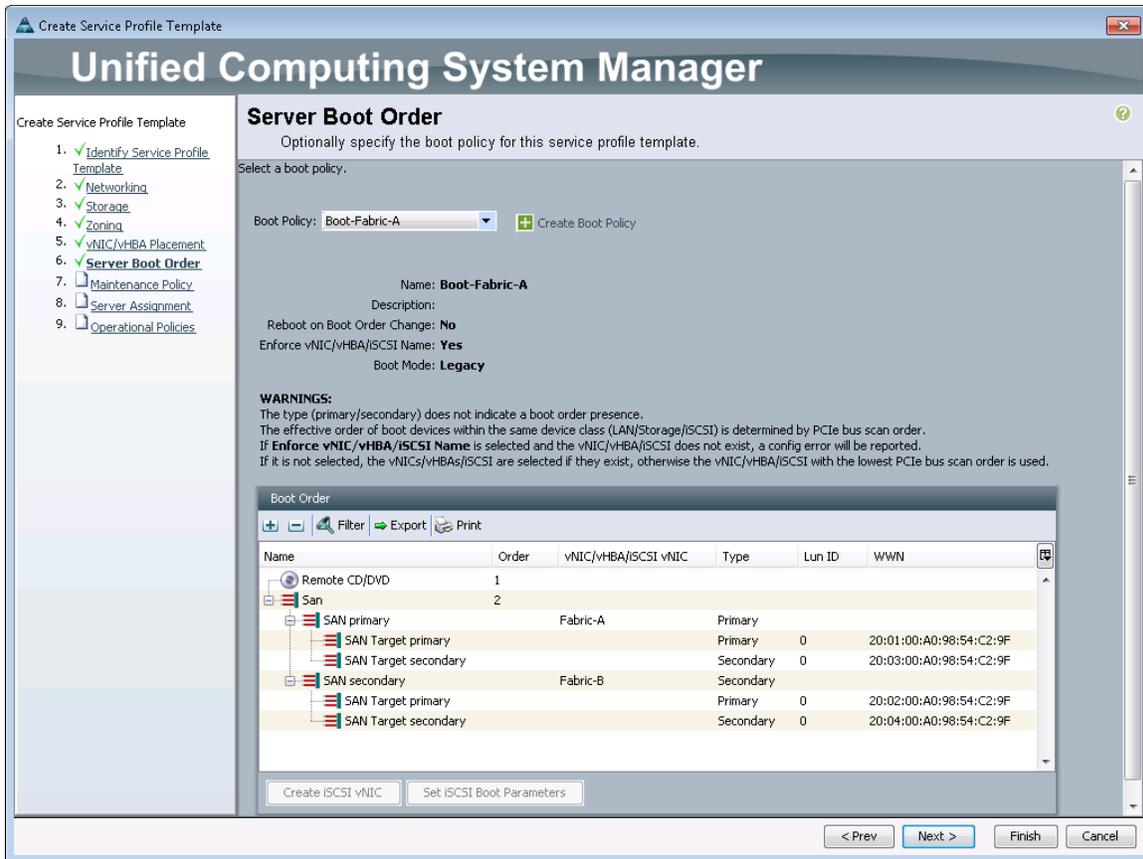
- d. Click Add to add a vHBA to the template.
- e. In the Create vHBA dialog box, enter `Fabric-A` as the name for vHBA.
- f. Select the Use vHBA Template checkbox.
- g. In the vHBA Template list, select `vHBA_Template_A`.
- h. In the Adapter Policy list, select Linux.
- i. Click OK to add this vHBA to the template.
- j. On the Storage page of the wizard, click Add to add another vHBA to the template.
- k. In the Create vHBA dialog box, enter `Fabric-B` as the name for vHBA.
- l. Select the checkbox for Use HBA Template.
- m. In the vHBA Template list, select `vHBA_Template_B`.
- n. In the Adapter Policy list, select Linux.
- o. Click OK to add the vHBA to the template.
- p. Review the table on the Storage page to verify that both vHBAs were created.



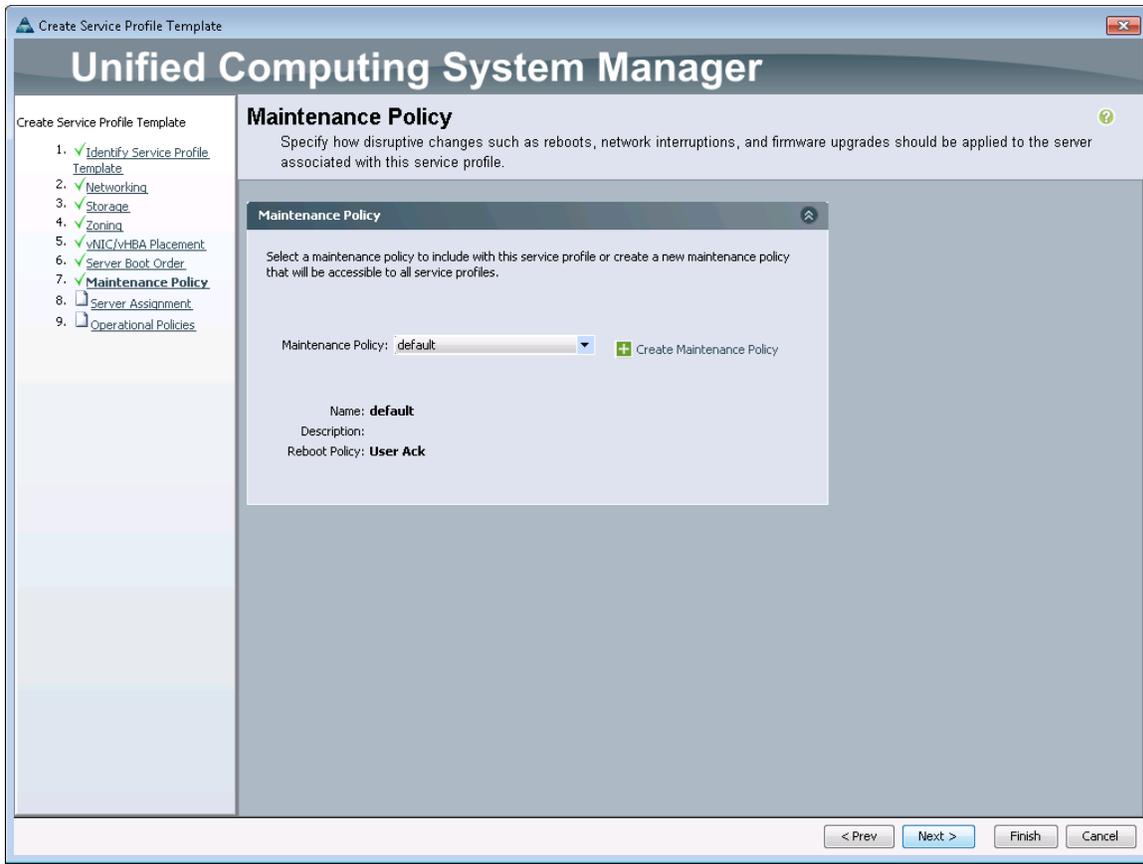
q. Click Next.



8. Select no zoning options and click Next.
9. Set the vNIC/vHBA placement options.
 - a. In the Select Placement list, select the RHEL-Server placement policy.
 - b. Select vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - vHBA Fabric-A
 - vHBA Fabric-B
 - vNIC-A
 - vNIC-B
 - c. Review the table to verify that all of the vNICs and vHBAs were assigned to the policy in the appropriate order.
 - d. Click Next.
10. Set the Server Boot Order:
 - a. In the Boot Policy list, select Boot-Fabric-A.
 - b. Review the table to verify that all of the boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.



- c. Click Next.
11. Configure the maintenance policy:
 - a. Confirm that Maintenance Policy is set to default.



- b. Click Next.
12. Specify the server assignment:
 - a. In the Pool Assignment list, select Infra_Pool.
 - b. Optional: Select a Server Pool Qualification policy.
 - c. Select Down as the power state to be applied when the profile is associated with the server.
 - d. Expand Firmware Management and select default from the Host Firmware list.
 - e. Click Next.
13. Add operational policies:
 - a. In the BIOS Policy list, select RHEL-Server.
 - b. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.
14. Click Finish to create the service profile template.
15. Click OK in the confirmation message.

Create Service Profile

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Template RHEL-Server-Fabric-A.
3. Right-click RHEL-Server-Fabric-A and select Create Service Profiles from Template.
4. Enter RHEL-Server-0 as the Naming Prefix.

5. Enter 1 as the Suffix Starting Number
6. Enter 1 as the Number of Instances to create.
7. Click OK to create the service profile.
8. Click OK in the confirmation message.
9. Verify that the service profile `RHEL-Server-01` was created. The service profiles are automatically associated with the servers in their assigned server pools.
10. Optional: Select each newly created service profile and enter the server host name or the FQDN in the User Label field on the General tab. Click Save Changes to map the server host name to the service profile name.

Add More 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 24 and Table 25.

Table 24) FCP LIFs for FC WWPNS.

FCP LIFs	FC WWPNS
fcp_lif01a	
fcp_lif01b	
fcp_lif02a	
fcp_lif02b	

Note: To gather the FC WWPNS, log in to the storage cluster and run the `network interface show` command.

Table 25) vHBA WWPNS for Fabric A and Fabric B.

Cisco UCS Service Profile Name	vHBA Fabric A WWPNS	vHBA Fabric B WWPNS
RHEL-Server-01		

Note: To gather the vHBA WWPNS information, start the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root. Click each service profile, then click the Storage tab in the right pane, and then click the vHBAs tab. In Table 25, record the WWPNS information that is displayed in the right pane for both vHBA Fabric A and vHBA Fabric B for each service profile.

9.6 Configure Storage Networking: FlexPod Cisco Nexus Base

Table 26) FlexPod Cisco Nexus base prerequisites.

Prerequisite
The Cisco Nexus switch must be running Cisco Nexus NX-OS 7.0(0)N1(1) or later.

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

Set Up Initial Configuration

Cisco Nexus 5548UP A

To set up the initial configuration for the Cisco Nexus A switch on <<var_nexus_A_hostname>>, complete the following steps:

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <<var_nexus_A_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_A_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_A_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <<var_nexus_A_mgmt0_gw>>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

  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_global_ntp_server_ip>>
Enter basic FC configurations (yes/no) [n]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus 5548UP B

To set up the initial configuration for the Cisco Nexus B switch on <<var_nexus_B_hostname>>, complete the following steps:

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <<var_nexus_B_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_B_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_B_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
```

```
IPv4 address of the default gateway: <<var_nexus_B_mgmt0_gw>>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

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_global_ntp_server_ip>>
Enter basic FC configurations (yes/no) [n]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

9.7 Configure FlexPod Cisco Nexus FCoE Storage on Clustered Data ONTAP

Enable Licenses

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To license the Cisco Nexus switches, complete the following steps:

1. Log in as admin.
2. Run the following commands:

```
config t
feature fcoe
feature npiv
feature lacp
feature vpc
```

Set Global Configurations

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To set global configurations, complete the following step on both the switches:

1. Run the following commands to set global configurations and jumbo frames in QoS:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
port-channel load-balance ethernet source-dest-port
policy-map type network-qos jumbo
class type network-qos class-default
mtu 9216
exit
class type network-qos class-fcoe
pause no-drop
mtu 2158
exit
exit
system qos
service-policy type network-qos jumbo
exit
copy run start
```

Create VLANs

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To create the necessary virtual local area networks (VLANs), complete the following step on both the switches:

1. From the global configuration mode, run the following commands:

```
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
exit
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 5548UP A

To add individual port descriptions for troubleshooting activity and verification for switch A, complete the following step:

1. From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_node01>>:e0e
exit
interface Eth1/2
description <<var_node02>>:e0e
exit
interface Eth1/11
description <<var_ucs_clustername>>-A:1/19
exit
interface Eth1/12
description <<var_ucs_clustername>>-B:1/19
exit
interface Eth1/13
description <<var_nexus_B_hostname>>:1/13
exit
interface Eth1/14
description <<var_nexus_B_hostname>>:1/14
exit
interface eth1/31
description <<var_ucs_clustername>>-A:1/31
exit
interface eth1/32
description <<var_ucs_clustername>>-A:1/32
exit
```

Cisco Nexus 5548UP B

To add individual port descriptions for troubleshooting activity and verification for switch B, complete the following step:

1. From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_node01>>:e0g
exit
interface Eth1/2
description <<var_node02>>:e0g
exit
interface Eth1/11
description <<var_ucs_clustername>>-A:1/20
exit
interface Eth1/12
description <<var_ucs_clustername>>-B:1/20
exit
interface Eth1/13
description <<var_nexus_A_hostname>>:1/13
exit
interface Eth1/14
description <<var_nexus_A_hostname>>:1/14
exit
```

```
interface eth1/31
description <<var_ucs_clustername>>-B:1/31
exit
interface eth1/32
description <<var_ucs_clustername>>-B:1/32
exit
```

Create Port Profiles

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

Port profiles are used to simplify ongoing network administration and configuration. Ports with similar configurations can be grouped within port profiles. Configuration changes can then be made to the port profile and will be applied to all port members of the port profile. NetApp recommends port profiles for the following port types:

- FAS uplink ports
- UCS Ethernet ports
- UCS FCoE ports
- Nexus VPC ports
- Nexus 1110-X ports

To create the Ethernet traffic port profiles, complete the following step on both the switches:

1. From the Global configuration mode, run the following commands

```
port-profile default max-ports 512
port-profile type port-channel FAS-data-uplink
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk
load-interval counter 3 60
state enabled
port-profile type port-channel UCS-Ethernet
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk
state enabled
port-profile type port-channel 1110X
switchport mode trunk
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type edge trunk

state enabled
port-profile type port-channel vPC-Peer-Link
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>
spanning-tree port type network
state enabled
exit
```

Create Port Channels

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To create the necessary port channels between devices, complete the following step on both the switches:

1. From the global configuration mode, run the following commands:

```

interface Po10
description vPC peer-link
exit
interface Eth1/13-14
channel-group 10 mode active
no shutdown
exit
interface Po11
description <<var_node01>>
exit
interface Eth1/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
exit
interface Eth1/2
channel-group 12 mode active
no shutdown
exit
interface Po13
description <<var_ucs_clustername>>-A
exit
interface Eth1/11
channel-group 13 mode active
no shutdown
exit
interface Po14
description <<var_ucs_clustername>>-B
exit
interface Eth1/12
channel-group 14 mode active
no shutdown
exit
copy run start

```

Add Port Profiles to Port Channels

Port channels and their member ports inherit their configuration from the previously configured port profiles.

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To assign port profiles to the appropriate port channels, complete the following step on both switches:

1. From the global configuration mode, run the following commands:

```

interface Po10
inherit port-profile vPC-Peer-Link
exit
interface Po11
inherit port-profile FAS-data-uplink
exit
interface Po12
inherit port-profile FAS-data-uplink
exit
interface Po13
inherit port-profile UCS-Ethernet
exit
interface Po14
inherit port-profile UCS-Ethernet
exit
copy run start

```

Configure Virtual Port Channels

Cisco Nexus 5548UP A

To configure virtual port channels (vPCs) for switch A, complete the following step:

1. From the global configuration mode, run the following commands:

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 10
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source <<var_nexus_A_mgmt0_ip>>
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Cisco Nexus 5548UP B

To configure vPCs for switch B, complete the following step:

1. From the global configuration mode, run the following commands.

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 20
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source <<var_nexus_B_mgmt0_ip>>
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 5548UP switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Run `copy run start` to save the configuration on each switch after the configuration is completed.

Create VSANs, Assign, and Enable Virtual Fibre Channel Ports

This procedure configures the FCoE connections between the Cisco Nexus 5548UP switches, the Cisco UCS fabric interconnects, and the NetApp storage systems.

Cisco Nexus 5548UP A

To configure virtual storage area networks (VSANs), create and update relevant port profiles, assign virtual Fibre Channel (vFC) ports, and enable vFC ports on switch A, complete the following step:

1. From the global configuration mode, run the following commands:

```
vlan <<var_fabric_a_fcoe_vlan_id>>
name FCoE_Fabric_A
fcoe vsan <<var_vsan_a_id>>
exit

port-profile type port-channel FAS-data-uplink
  switchport trunk allowed vlan add <<var_fabric_a_fcoe_vlan_id>>
  exit
port-profile type port-channel UCS-FCOE-FABRIC-A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan <<var_fabric_a_fcoe_vlan_id>>
  spanning-tree port type edge trunk
  state enabled
  exit

interface vfc11
switchport description <<var_node01>>:0e
bind interface Eth1/1
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
exit
interface vfc12
switchport description <<var_node02>>:0e
bind interface Eth1/2
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
exit
interface po15
description <<var_ucs_clustername>>-A:FCoE
exit
interface Eth1/31-32
channel-group 15 mode active
exit
interface po15
inherit port-profile UCS-FCOE-FABRIC-A
exit
interface vfc15
switchport description <<var_ucs_clustername>>-A:FCoE
bind interface po15
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
vsan database
vsan <<var_vsan_a_id>> name Fabric_A
vsan <<var_vsan_a_id>> interface vfc11
vsan <<var_vsan_a_id>> interface vfc12
vsan <<var_vsan_a_id>> interface vfc15
exit
copy run start
```

Cisco Nexus 5548UP B

To configure VSANs, create and update relevant port profiles, assign vFC ports, and enable vFC ports on switch B, complete the following step:

1. From the global configuration mode, run the following commands:

```

vlan <<var_fabric_b_fcoe_vlan_id>>
name FCoE_Fabric_B
fcoe vsan <<var_vsan_b_id>>
exit
interface po11
switchport trunk allowed vlan add <<var_fabric_b_fcoe_vlan_id>>
exit
interface vfc11
switchport description <<var_node01>>:0g
bind interface Eth1/1
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
exit
interface vfc12
switchport description <<var_node02>>:0g
bind interface Eth1/2
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
exit
interface po15
description <<var_ucs_clustername>>-B:FCoE
exit
interface Eth1/31-32
channel-group 15 mode active
exit
interface po15
inherit port-profile UCS-FCOE-FABRIC-B
exit
interface vfc15
switchport description <<var_ucs_clustername>>-B:FCoE
bind interface po15
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
vsan database
vsan <<var_vsan_b_id>> name Fabric_B
vsan <<var_vsan_b_id>> interface vfc11
vsan <<var_vsan_b_id>> interface vfc12
vsan <<var_vsan_b_id>> interface vfc15
exit
copy run start

```

Create Device Aliases

Cisco Nexus 5548UP A

To configure device aliases and zones for the primary boot paths of switch A on <<var_nexus_A_hostname>>, complete the following step:

1. From the global configuration mode, run the following commands:

```

device-alias database
device-alias name RHEL-Server-01_A pwn <<var_rhel-server_01_A_wwpn>>
device-alias name fcp_lif01a pwn <<var_fcp_lif01a_wwpn>>
device-alias name fcp_lif02a pwn <<var_fcp_lif02a_wwpn>>
exit
device-alias commit

```

Cisco Nexus 5548UP B

To configure device aliases and zones for the boot paths of switch B on <<var_nexus_B_hostname>>, complete the following step:

1. From the global configuration mode, run the following commands:

```

device-alias database
device-alias name RHEL-Server-01_B pwn <<var_rhel-server_01_B_wwpn>>
device-alias name fcp_lif01b pwn <<var_fcp_lif01b_wwpn>>
device-alias name fcp_lif02b pwn <<var_fcp_lif02b_wwpn>>

```

```
exit
device-alias commit
```

Create Zones

Cisco Nexus 5548UP A

To create zones for the service profile on switch A, complete the following steps:

1. Create a zone for the service profile.

```
zone name RHEL-Server-01_A vsan <<var_vsan_a_id>>
member device-alias RHEL-Server-01_A
member device-alias fcp_lif01a
member device-alias fcp_lif02a
exit
```

2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members.

```
zoneset name FlexPod vsan <<var_vsan_a_id>>
member RHEL-Server-01_A
exit
```

3. Activate the zone set.

```
zoneset activate name FlexPod vsan <<var_vsan_a_id>>
exit
copy run start
```

Cisco Nexus 5548UP B

To create zones for the service profile on switch B, complete the following steps:

1. Create a zone for each service profile.

```
zone name RHEL-Server-01_B vsan <<var_vsan_b_id>>
member device-alias RHEL-Server-01_B
member device-alias fcp_lif01b
member device-alias fcp_lif02b
exit
```

2. After the zones for the Cisco UCS service profile have been created, create the zone set and add the necessary members.

```
zoneset name FlexPod vsan <<var_vsan_b_id>>
member RHEL-Server-01_B
exit
```

3. Activate the zone set.

```
zoneset activate name FlexPod vsan <<var_vsan_b_id>>
exit
copy run start
```

9.8 Set Up Storage: Clustered Data ONTAP SAN Boot

Create Igroups

To create igroups, complete the following step:

1. From the cluster management node SSH connection, enter the following:

```
igroup create -vserver RHEL_Server -igroup RHEL_Server-01 -protocol fcp -ostype linux -initiator
<<var_rhel-server_01_A_wwpn>>, <<var_rhel-server_01_B_wwpn>>
```

```
igroup create -vserver RHEL_Server -igroup MGMT-Hosts -protocol fcp -ostype linux -initiator <<var_rhel-server_01_A_wwpn>>, <<var_rhel-server_01_B_wwpn>>
```

Note: To view the recently created igroups, enter `igroup show`.

Map Boot LUNs to Igroups

To map boot LUNs to igroups, complete the following step:

1. From the cluster management SSH connection, run the following command:

```
lun map -vserver RHEL_Server -volume ucs_boot -lun RHEL-Infra-01 -igroup RHEL_Server-01 -lun-id 0
```

Log in to Cisco UCS 6200 Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (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 starts the Cisco UCS Manager application.
2. To download the Cisco UCS Manager software, click Launch UCS Manager.
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 > RHEL-Server-01.
8. Right-click `RHEL-Server-01` and select KVM Console.
9. If prompted to accept an unencrypted KVM session, accept as necessary.

9.9 Install Linux

Set Up Red Hat Enterprise Linux Installation

To prepare the server for the OS installation, complete the following steps:

1. In the KVM window, click the Virtual Media node.
2. If prompted to accept an unencrypted KVM session, accept as necessary.
3. Click Add Image.
4. Browse to the Red Hat Enterprise Linux ISO image file and click Open.
5. Select the Mapped checkbox to map the newly added image.
6. Click the KVM tab to monitor the server boot.
7. Boot the server by selecting Boot Server and clicking OK. Then click OK again.

Install Red Hat Enterprise Linux

To install the OS, complete the following steps:

1. When prompted to check the installation media, select Skip by pressing the tab key and then pressing Enter.
2. Select the appropriate language for the installation and click Next.

3. Select the appropriate keyboard layout and click Next.
4. Select Specialized Storage Devices and click Next.
5. From the Multipath Devices tab, select the appropriate identifier by looking for NetApp in the vendor column and click Next.
6. Verify the WWID and size of the multipath device in the Storage Device Warning window to make sure that the correct device was selected. Then select the Apply My Choice to All Devices with Undetected Partitions or File Systems checkbox and click Yes, Discard Any Data.
7. Enter a hostname for this machine and click Next.
8. Select the appropriate city and time zone and click Next.
9. Set the root password and click Next.
10. Select Use All Space and click Next.
11. In the confirmation message, click Write Changes to Disk.
12. Select Basic Server as the installation type. Select Red Hat Enterprise Linux only in the software repositories box, and keep Customize Later selected. Click Next.
13. Click Reboot after the installation is complete.
14. After the machine reboots, log in as the root user by using the previously entered password.

10 Optional Security Enhanced Linux Procedures

If you choose to use SELinux in your deployment, the following sections explain SELinux enablement and configuration. For more information about SELinux, refer to section 6.

10.1 Configure Booleans

Parts of the SELinux policy can be changed at runtime through the use of Booleans. Using Booleans permits changes without reloading or recompiling the SELinux policy. For example, by default, Apache HTTP Server scripts and modules are prohibited from connecting to database servers.

To change this, complete the following steps:

1. Run the following command:

```
setsebool httpd_can_network_connect_db on
```

2. Run the following command to check the value of the Boolean:

```
getsebool httpd_can_network_connect_db
```

3. To make this change persistent across reboots, run the following command:

```
setsebool -P httpd_can_network_connect_db
```

For more information about Booleans, review the man pages for `getsebool`, `setsebool`, `booleans`, `selinux`, and `togglesebool`.

For a list of Booleans, their current state (on or off), and a brief explanation of each, run the `semanage boolean -l` command.

Note: This command is included in the `policycoreutils-python` package. To install this package, run the `yum -y install policycoreutils-python` command.

10.2 Configure File Contexts

File contexts can be viewed by running the `ls -Z` command, as shown:

```
[root@rhel6 ~]# ls -Z .bashrc
-rw-r--r--. root root system_u:object_r:admin_home_t:s0 .bashrc
[root@rhel6 ~]# _
```

In the previous example, the `.bashrc` file in the root user's home directory is assigned a user (`system_u`), a role (`object_r`), a type (`admin_home_t`), and a level (`s0`). For the default SELinux policy used, the main permission control used is `Type Enforcement`.

To change the SELinux context on a file, use the `chcon` command. For example, `chcon -t <type> <filename>` changes the file type (such as `httpd_sys_content_t`) for the specified file name.

A directory can also be passed as a file name. To change the type of a directory and its contents, run the `chcon -R -t <type> <directoryname>` command. To restore the original SELinux context of a file, run the `restorecon <filename>` command. For more information, review the man page for `chcon`.

10.3 SELinux Policy

By default, the targeted policy is used for SELinux. For more complex control, a policy generation tool is included in the `selinux-policy` package (installed by default on RHEL 6.6) that provides examples and allows administrators to create their own custom SELinux policy.

Enable SELinux

To enable SELinux, complete the following steps:

1. Verify the current status of SELinux by running the `sestatus` command and note the value of the `Current mode`.

Note: Even if SELinux is enabled, it will not block any operations on the system unless the current mode is set to `enforcing`.

```
[root@rhel6 ~]# sestatus
SELinux status:                enabled
SELinuxfs mount:              /selinux
Current mode:                  enforcing
Mode from config file:        enforcing
Policy version:               24
Policy from config file:      targeted
[root@rhel6 ~]# _
```

2. Take one of the following actions:
 - If SELinux is set to `disabled`, each file on the file system must be labeled with an SELinux context before it can be enabled, and the system must be rebooted. To do this, set `SELINUX=permissive` in `/etc/selinux/config` and then run the `reboot` command. Then run `setenforce 1` and continue with step 3.
 - If `Current mode` is set to `permissive`, run `setenforce 1` to set SELinux to `enforcing` mode for the current session.
3. Run `sestatus` again to make sure that `Current mode` is set to `enforcing`.
4. To make this change persistent across reboots, set `SELINUX=enforcing` in `/etc/selinux/config`.

Disable SELinux

To disable SELinux, complete the following steps:

1. To stop SELinux from blocking operations on the machine for the current session, run `setenforce 0` to set the current mode to permissive.
2. To make this change persistent across reboots, set `SELINUX=permissive` in `/etc/selinux/config`.
3. To completely disable SELinux, set `SELINUX=disabled` in `/etc/selinux/config` and then run `reboot`.

11 Solution Verification

This solution was verified by confirming that the basic FlexPod configuration was complete and included the following elements:

- Cabling
- NetApp FAS8040 controller deployment
- Cisco Nexus 5596UP cluster network switch deployment
- Clustered Data ONTAP installation and deployment
- Cisco Nexus 5548UP switch deployment

The following elements were verified by logging in to the NetApp storage cluster:

- SVM configuration
- Load sharing mirror of SVM root volume creation
- FlexVol configuration
- Boot LUN creation and mapping
- Deduplication configuration
- FCP LIF configuration
- SVM management failover group configuration
- Network interface configuration
- Network routing groups configuration

The configuration on each of the Cisco UCS 6248 fabric interconnects was verified to contain the correct fabric information, Cisco UCS cluster name, IP address, netmask, gateway, DNS server address, and domain name.

The following elements were verified by logging in to the Cisco UCS Manager web interface: The block of IP addresses for out-of-band KVM access in the Cisco UCS environment was verified to be in the same subnet as the management IP addresses for the Cisco UCS Manager by navigating to the LAN tab. The MAC address pools assigned to each switching fabric were verified by navigating to LAN > Pools > root. The WWNN and WWPN pools were verified for each fabric by navigating to SAN > Pools > root. The UUID suffix pools and server pool were verified by navigating to Servers > Pools > root. The VLANs were verified by navigating to LAN > LAN Cloud. The VSANs and FCoE uplink port channels for each fabric were verified by navigating to SAN > SAN Cloud tree. The vNIC templates for each fabric were verified by navigating to LAN > Policies > root. The vHBA templates for each fabric were verified by navigating to SAN > Policies > root. The boot policies for each cluster node and fabric were verified by navigating to Servers > Policies > root. The service profile template was verified by navigating to Service Profile Templates > root.

The device aliases and zones for the boot paths were verified on the Cisco Nexus 5548UP switches for each fabric.

SAN boot was verified by logging in to the cluster management node, running `igroup show`, and checking for the `RHEL_Server-01` igroup. Running `lun mapped show -vserver RHEL_Server` verified that the boot LUNs were mapped to the igroups successfully.

The service profile `RHEL-Server-01` was verified by right-clicking the service profile and selecting KVM Console. No errors were shown in the service profile. The KVM console was verified by navigating to the Cisco UCS cluster web address and selecting Launch KVM Manager.

To verify the previous steps during the Red Hat Enterprise Linux 6.6 installation, the correct multipath device was shown in the Specialized Storage Devices option, and the installation completed. Upon reboot, the system booted from SAN successfully. Running `sestatus` as root verified that SELinux was enabled and set to enforcing.

12 Conclusion

This solution enables customers to leverage the core FlexPod architecture by using Red Hat Enterprise Linux without the need for a hypervisor. This reduces the overall operational costs and saves time spent on hypervisor configuration steps. Using Red Hat Enterprise Linux with Security Enhanced Linux limits the scope of potential damage from software and other security vulnerabilities. This solution provides the steps to create an infrastructure deployment with FCoE-booted bare metal hosts with block-level access to shared storage datastores.

Authors and Contributors

Much of this document is based [TR-4328: FlexPod Express with VMware vSphere 5.5 Update 1: Small and Medium Configurations Implementation Guide](#) and [TR-4331: FlexPod Express with VMware vSphere 5.5 Update 1: Large Configurations Implementation Guide](#), which were created by Arvind Ramakrishnan, Karthick Radhakrishnan, Lindsey Street, and John George of NetApp. The content specific to Red Hat was added by Jessica Sterling and Mike Scanlin of NetApp.

References

This report references the following documents and resources:

- Certification Report: NetApp Clustered Data ONTAP 8.2.1
www.commoncriteriaportal.org/files/epfiles/383-4-263%20CR%20v1.0e.pdf
- Certification Report: NetApp Data ONTAP v8.2.1 7-Mode
www.commoncriteriaportal.org/files/epfiles/netapp-data-v8217-cert-eng.pdf
- Clustered Data ONTAP Software Setup Guide
https://library.netapp.com/ecm/ecm_download_file/ECMP1368696
- Cisco UCS Manager Install and Upgrade Guides
www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html
- National Information Assurance Partnership Common Criteria Evaluation and Validation Scheme
www.commoncriteriaportal.org/files/epfiles/st_vid10023-vr.pdf
- NetApp Hardware Universe
<http://mysupport.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/>
- NetApp Support site
<http://mysupport.netapp.com/>
- SAS Disk Shelves Universal SAS and ACP Cabling Guide
http://mysupport.netapp.com/NOW/knowledge/docs/hardware/filer/215-05500_A0.pdf?isLegacy=true

- TR-3832: Flash Cache Best Practice Guide
<http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-3832.pdf&cc=us>
- TR-4323: FlexPod Express with VMware vSphere 5.5 Update 1: Small and Medium Configurations Implementation Guide
<https://fieldportal.netapp.com/myFieldPortal.aspx?oparams=263592>
- TR-4331: FlexPod Express with VMware vSphere 5.5 Update 1: Large Configurations Implementation Guide
<https://fieldportal.netapp.com/?oparams=267052>
- U.S. Federal Government Customers Choose NetApp More Than Any Other Storage Vendor
www.netapp.com/us/company/news/press-releases/news-rel-20121129-909054.aspx

Version History

Version	Date	Document Version History
Version 1.0	April 2015	First release

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

Copyright Information

Copyright © 1994–2015 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, 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-0014-DEPLOY-0415