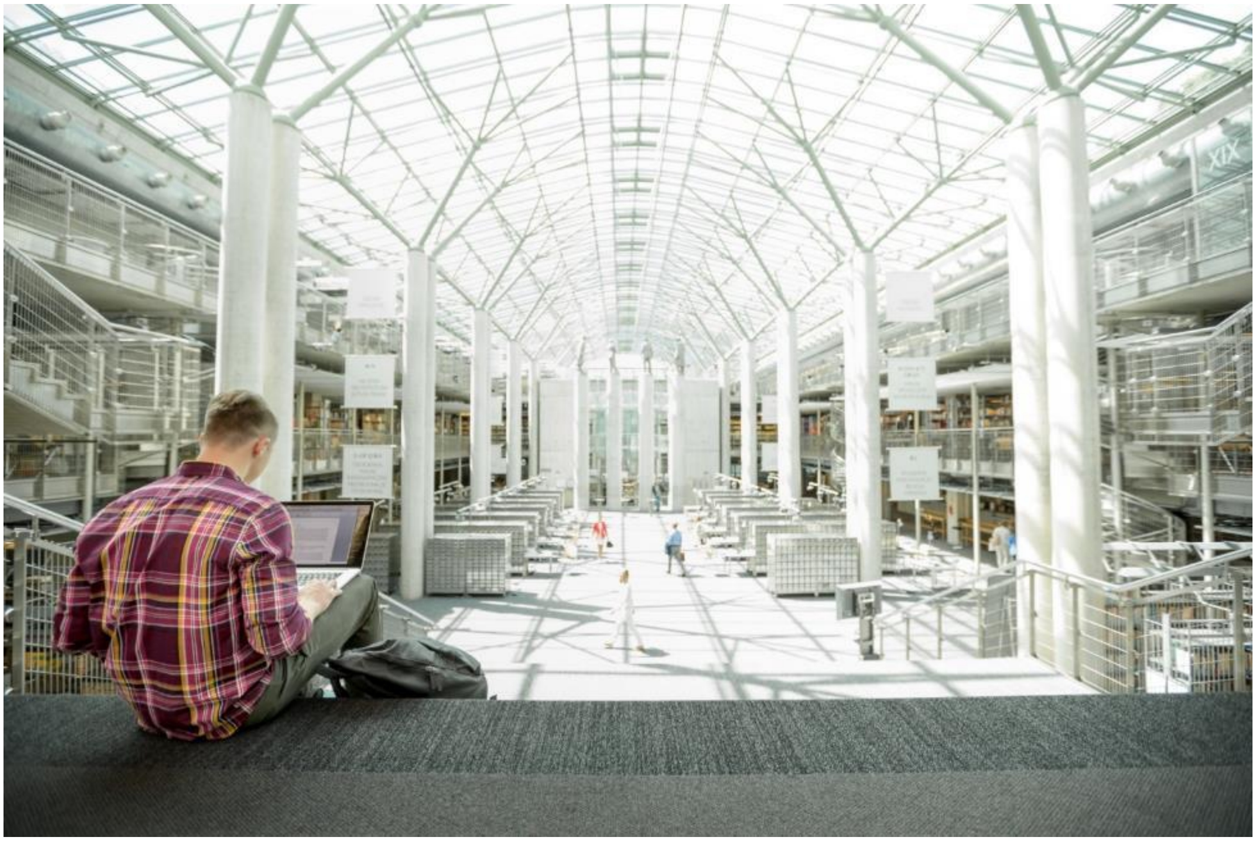

FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS 4th Generation Fabric, and NetApp AFF A-Series

Deployment Guide for FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS Manager 4.0(2), and ONTAP 9.5

Published: November 2019

In partnership with:

About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

http://www.cisco.com/go/designzone.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, " DESIGNS") IN THIS MANUAL ARE PRESENTED " AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

Table of Contents

Executive Summary

Cisco Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

This document describes the Cisco and NetApp® FlexPod Datacenter with Cisco UCS Manager unified software release 4.0(2) and VMware vSphere 6.7 U1. Cisco UCS Manager (UCSM) 4.0(2) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 6454,2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with Cisco UCS unified software release 4.0(2), and VMware vSphere 6.7 U1 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays running ONTAP® 9 storage OS.

Solution Overview

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated 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 quickly deployed, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod, which uses best of breed storage, server and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides a step by step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS Fabric Interconnects, NetApp AFF storage, Cisco MDS, and Cisco Nexus 9000 solution.

What’s New in this Release?

The following design elements distinguish this version of FlexPod from previous FlexPod models:

- Support for the Cisco UCS 4.0(2) unified software release, Cisco UCS B200-M5 servers, and Cisco UCS C220-M5 servers, Cisco 1400 Series Virtual Interface Cards (VICs)
- Support for the latest Cisco UCS 6454 Fabric Interconnect
- Support for the NetApp AFF A800 storage controller
- Support for the latest release of NetApp ONTAP® 9.5
- Support for NetApp Virtual Storage Console (VSC) 7.2.1
- Support for NetApp SnapCenter 4.1.1
- Fibre channel, NFS, iSCSI (appendix) storage design
- Validation of VMware vSphere 6.7 U1
- Unified Extensible Firmware Interface (UEFI) Secure Boot of VMware ESXi 6.7 U1
- Trusted Platform Module (TPM) 2.0 Attestation of UEFI Secure Boot of VMware ESXi 6.7 U1
- 100 Gigabit per second Ethernet Connectivity
- 32 Gigabit per second Fibre Channel Connectivity

Deployment Hardware and Software

Architecture

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere® built on FlexPod includes NetApp All Flash FAS storage, Cisco Nexus® networking, Cisco MDS storage networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the 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. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a Fibre Channel and IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

Figure 1 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with the Cisco UCS 6454 Fabric Interconnects. This design has port-channeled 10 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects, port-channeled 25 Gb Ethernet connections between the C-Series rackmounts and the Cisco UCS Fabric Interconnects, and 100 Gb Ethernet connections between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A800 storage array. This infrastructure option expanded with Cisco MDS switches sitting between the Cisco UCS Fabric Interconnect and the NetApp AFF A800 to provide FC-booted hosts with 32 Gb FC block-level access to shared storage. The reference architecture reinforces the " wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Topology

Figure 1   FlexPod with Cisco UCS 6454 Fabric Interconnects and NetApp AFF A800

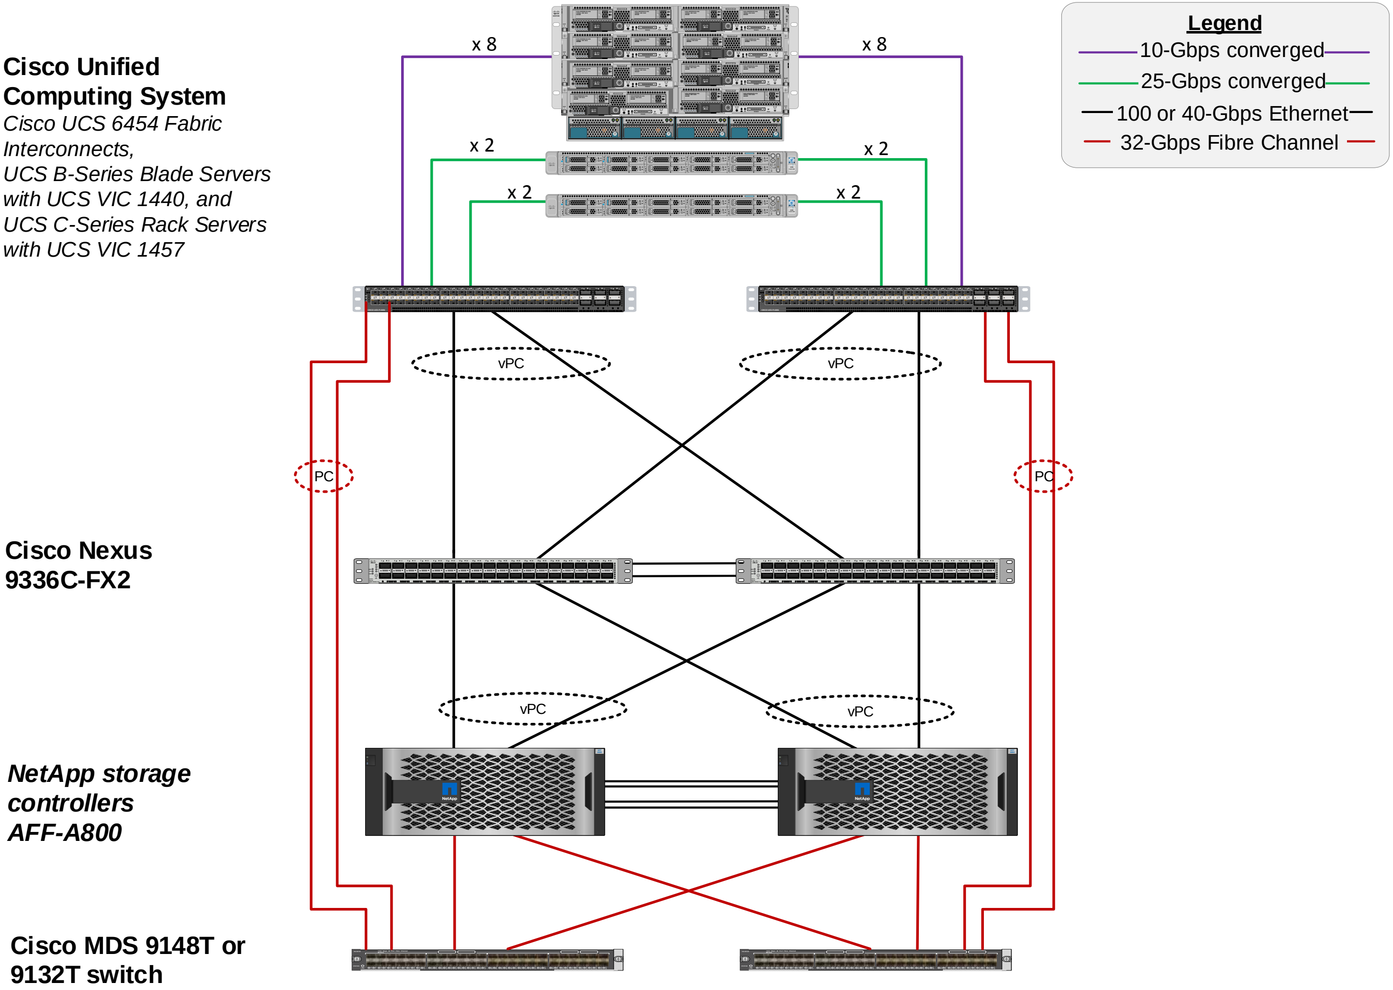
Cisco Unified Computing System
Cisco UCS 6454 Fabric Interconnects,
UCS B-Series Blade Servers with UCS VIC 1440, and
UCS C-Series Rack Servers with UCS VIC 1457
x 8
x 8
x 2
x 2
x 2
x 2
Legend
10-Gbps converged
25-Gbps converged
100 or 40-Gbps Ethernet
32-Gbps Fibre Channel
vPC
vPC
PC
PC
Cisco Nexus 9336C-FX2
vPC
vPC
NetApp storage controllers AFF-A800
Cisco MDS 9148T or 9132T switch

The reference 100Gb based hardware configuration includes:

- Two Cisco Nexus 9336c-FX2 switches
- Two Cisco UCS 6454 fabric interconnects
- Two Cisco MDS 9132T multilayer fabric switches
- One NetApp AFF A800 (HA pair) running ONTAP 9.5 with internal NVMe SSD disks

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1   Software Revisions

Layer	Device	Image	Comments

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6454, UCS B-200 M5, UCS C-220 M5	4.0(2d)	Includes the Cisco UCS-IOM 2208, Cisco UCS Manager, Cisco UCS VIC 1440 and Cisco UCS VIC 1457
Network	Cisco Nexus 9336C-FX2 NX-OS	7.0(3)I7(6)	
	Cisco MDS 9132T	8.3(2)	
Storage	NetApp AFF A800	ONTAP 9.5	
Software	Cisco UCS Manager	4.0(2d)	
	Cisco UCS Manager Plugin for VMware vSphere Web Client	2.0.4	
	VMware vSphere	6.7U1	
	VMware ESXi nfnic FC Driver	4.0.0.24	
	VMware ESXi nenic Ethernet Driver	1.0.27.0	
	NetApp Virtual Storage Console (VSC) / VASA Provider Appliance	7.2.1P1	
	NetApp NFS Plug-in for VMware VAAI	1.1.2	
	NetApp SnapCenter Backup Management	4.1.1	Includes the vSphere plug-in for SnapCenter

Configuration Guidelines

This document explains how to configure a fully redundant, highly available configuration for a FlexPod unit with ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies 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 multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-01, VM-Host-Infra-02 to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?

 [-node] <nodename>                   Node

 { [-vlan-name] {<netport>|<ifgrp>}  VLAN Name
```

```
    |  -port {<netport>|<ifgrp>}        Associated Network Port
    [-vlan-id] <integer> }              Network Switch VLAN Identifier
```

Example:

```
network port vlan create -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 2 describes the VLANs necessary for deployment as outlined in this guide.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out of Band Mgmt	VLAN for out-of-band management interfaces	13
In-Band Mgmt	VLAN for in-band management interfaces	113
Native	VLAN to which untagged frames are assigned	2
Infra-NFS	VLAN for Infrastructure NFS traffic	3050
FCoE-A	VLAN for FCoE encapsulation of VSAN-A	101
FCoE-B	VLAN for FCoE encapsulation of VSAN-B	102
vMotion	VLAN for VMware vMotion	3000
VM-Traffic	VLAN for Production VM Interfaces	900

Table 3 lists the VMs necessary for deployment as outlined in this document.

Table 3 Virtual Machines

Virtual Machine Description	Host Name	IP Address
vCenter Server		
NetApp VSC		
NetApp SnapCenter		

Physical Infrastructure

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, a cabling diagram was used.

The cabling diagram in this section contains details for the prescribed and supported configuration of the NetApp AFF A800 running NetApp ONTAP® 9.5.

For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

Be sure to use the cabling directions in this section as a guide.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to NetApp Support.

Figure 2 details the cable connections used in the validation lab for the FlexPod topology based on the Cisco UCS 6454 fabric interconnect. Two 32Gb uplinks connect as port-channels to each Cisco UCS Fabric Interconnect from the MDS switches, and a total of four 32Gb links connect the MDS switches to the NetApp AFF controllers. Also, 100Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the NetApp AFF controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has two connections to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

Figure 2 FlexPod Cabling with Cisco UCS 6454 Fabric Interconnect

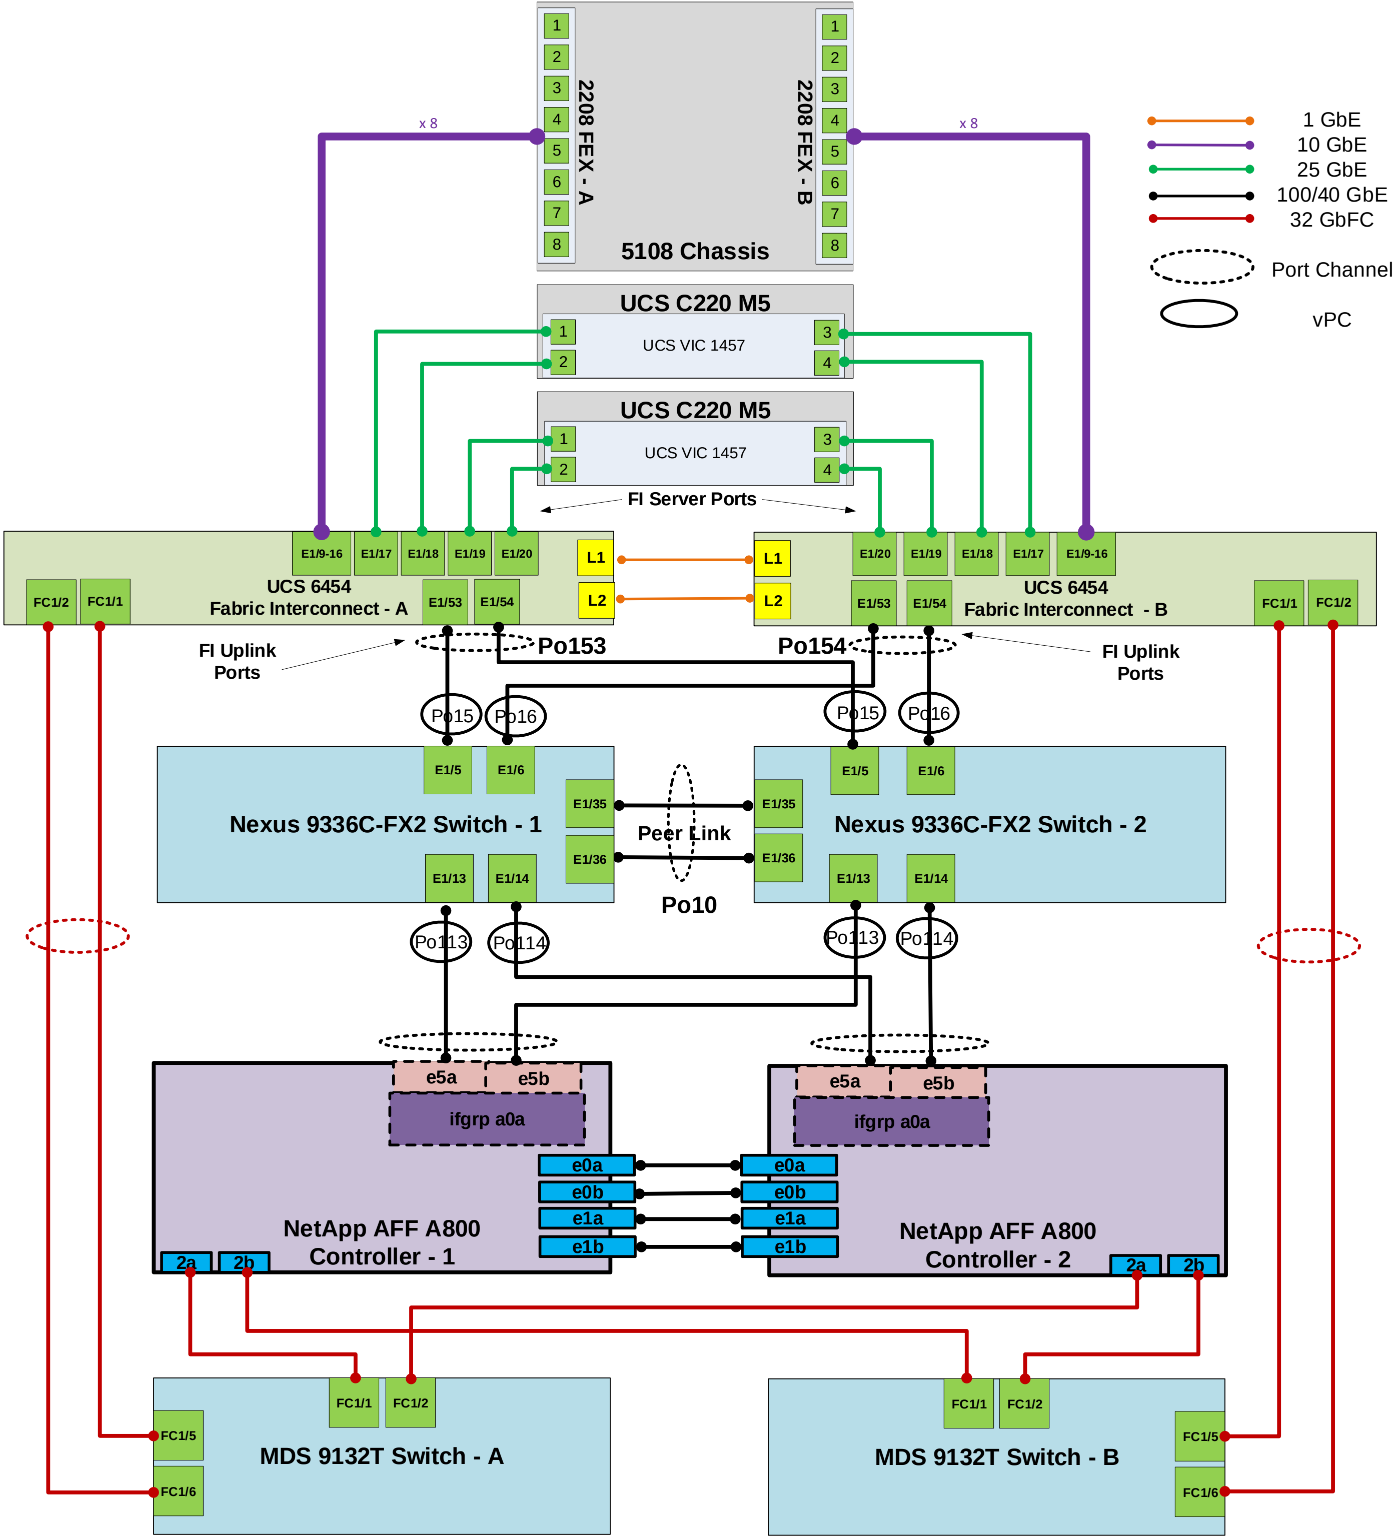

Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in the section FlexPod Cabling.

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 7.0(3)I7(6), the Cisco suggested Nexus switch release at the time of this validation.

The following procedure includes the setup of NTP distribution on both the mgmt0 port and the in-band management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes that the default VRF is used to route the in-band management VLAN.

With Cisco Nexus 9000 release 7.0(3)I7(6), autonegotiation (40G/100G) is not supported on ports 1-6 and 33-36 on the Cisco Nexus 9336C-FX2 switch. Since those ports are in use in this validation, port speed and duplex are hard set at both ends of the connection.

Set Up Initial Configuration

Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>, follow these steps:

1. Configure the switch.

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 [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no)[no]: yes
Disabling POAP.......Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

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

Do you want to enforce secure password standard (yes/no) [y]: Enter
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: <nexus-A-hostname>
```

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
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]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: 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 B

To set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>, follow these steps:

1. Configure the switch.

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 [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no)[no]: yes
Disabling POAP.......Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

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

Do you want to enforce secure password standard (yes/no) [y]: Enter
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: <nexus-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
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]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: 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
   ```

FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus A and Cisco Nexus B

To license the Cisco Nexus switches, follow these steps:

1. Log in as admin.

2. Run the following commands:

```
config t
feature udld
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

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

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

Create VLANs

Cisco Nexus A and Cisco Nexus B

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

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

```
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
vlan <infra-nfs-vlan-id>
```

```
name Infra-NFS-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus A

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

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-b-ntp-ip> use-vrf default
```

Cisco Nexus B

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

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-a-ntp-ip> use-vrf default
```

Add Individual Port Descriptions for Troubleshooting and Enable UDLD for UCS Interfaces

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A and to enable aggressive unidirectional link detection (UDLD) on copper interfaces connected to Cisco UCS systems, complete the following step:

In this step and in the later sections, configure the AFF nodename <st-node> and Cisco UCS 6454 fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

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

```
interface Eth1/5
description <ucs-clustername>-a:1/53
udld aggressive
interface Eth1/6
description <ucs-clustername>-b:1/53
udld aggressive
interface Eth1/13
description <st-node>-1:e5a
interface Eth1/14
description <st-node>-2:e5a
interface Eth1/35
description <nexus-b-hostname>:1/35
interface Eth1/36
description <nexus-b-hostname>:1/36
exit
```

Set udld to `aggressive` for copper cable or twinnax connections between the Nexus switches and the Cisco UCS Fabric Interconnects and to `enable` for fibre optic connections.

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B and to enable aggressive UDLD on copper interfaces connected to Cisco UCS systems, complete the following step:

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

```
interface Eth1/5
description <ucs-clustername>-a:1/54
udld aggressive
interface Eth1/6
description <ucs-clustername>-b:1/54
udld aggressive
interface Eth1/13
description <st-node>-1:e5b
interface Eth1/14
description <st-node>-2:e5b
interface Eth1/35
description <nexus-a-hostname>:1/35
interface Eth1/36
description <nexus-a-hostname>:1/36
exit
```

Set udld to `aggressive` for copper cable or twinnax connections between the Nexus switches and the UCS Fabric Interconnects and to `enable` for fibre optic connections.

Create Port Channels

Cisco Nexus A and Cisco Nexus B

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

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

```
interface Po10
description vPC peer-link
interface Eth1/35-36
channel-group 10 mode active
no shutdown
interface Po113
description <st-node>-1
interface Eth1/13
channel-group 13 mode active
no shutdown
interface Po114
description <st-node>-2
interface Eth1/14
channel-group 14 mode active
no shutdown
interface Po15
description <ucs-clustername>-a
interface Eth1/5
channel-group 15 mode active
no shutdown
interface Po16
description <ucs-clustername>-b
interface Eth1/6
channel-group 16 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, complete the following step on both switches:

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

```
interface Po10
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>
spanning-tree port type network
speed 100000
duplex full
no negotiate auto

interface Po113
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto

interface Po114
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto

interface Po15
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto

interface Po16
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto
exit
copy run start
```

Lab testing confirmed that the speed and duplex needed to be hard set on both ends for the NetApp storage interfaces even though these interfaces were not on ports 1-6 or 33-36.

Configure Virtual Port Channels

Cisco Nexus 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 <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po113
vpc 113
interface Po114
vpc 114
interface Po15
vpc 15
interface Po16
vpc 16
exit
copy run start
```

Cisco Nexus B

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

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

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po113
vpc 113
interface Po114
vpc 114
interface Po15
vpc 15
interface Po16
vpc 16
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, we recommend using vPCs to uplink the Cisco Nexus 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. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Switch Testing Commands

The following commands can be used to check for correct switch configuration:

Some of these commands need to run after further configuration of the FlexPod components are complete to see complete results.

```
show run
show vpc
show port-channel summary
show ntp peer-status
show cdp neighbors
show lldp neighbors
show udld neighbors
show run int
show int
```

Storage Configuration

NetApp All Flash FAS A800 Controllers

See the following section (NetApp Hardware Universe) for planning the physical location of the storage systems:

- Site Preparation
- System Connectivity Requirements
- Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements
- AFF Series Systems

NetApp Hardware Universe

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

To confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install, follow these steps at the NetApp Support site.

1. Access the HWU application to view the System Configuration guides. Click the Platforms menu to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers found in the AFF A800 Series product documentation at the NetApp Support site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported by the AFF A800 is available at the NetApp Support site.

When using SAS disk shelves with NetApp storage controllers, refer to section SAS shelves with IOM12 modules in the *AFF and FAS System Documentation Center* for proper cabling guidelines.

NetApp ONTAP 9.5

Complete Configuration Worksheet

Before running the setup script, complete the Cluster setup worksheet *in the ONTAP 9 Documentation Center*. You must have access to the NetApp Support site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the Software setup section of the *ONTAP 9 Documentation Center* to learn about configuring ONTAP. Table 4 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 4 ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
ONTAP 9.5 URL	<url-boot-software>

Configure Node 01

To configure node 01, follow these 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 the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.5 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.

5. Enter `y` to perform an upgrade.

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

7. Enter `y` to reboot now.

8. Enter the IP address, netmask, and default gateway for `e0M`.

```
<node01-mgmt-ip> <node01-mgmt-mask> <node01-mgmt-gateway>
```

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

This web server must be pingable.

```
<url-boot-software>
```

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

11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

12. Enter `y` to reboot the node.

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

13. Press **Ctrl-C** when the following message displays:

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

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

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

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

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, follow these 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 the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.5 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

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

```
<node02-mgmt-ip> <node02-mgmt-mask> <node02-mgmt-gateway>
```

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

This web server must be pingable.

```
<url-boot-software>
```

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

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

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

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

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

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

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.5 boots on the node for the first time.

1. Follow the prompts to set up node 01:

```
Welcome to node setup.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the setup wizard.
     Any changes you made before quitting will be saved.

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

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem
occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created

Use your web browser to complete cluster setup by accesing https://<node01-mgmt-ip>

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

2. To complete the cluster setup, open a web browser and navigate to `https://<node01-mgmt-ip>`.

Table 5 Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>

Cluster Detail	Cluster Detail Value
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-sp-ip>
Node 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>
NTP server B IP address	<switch-b-ntp-ip>

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

3. Click Guided Setup on the Welcome screen.

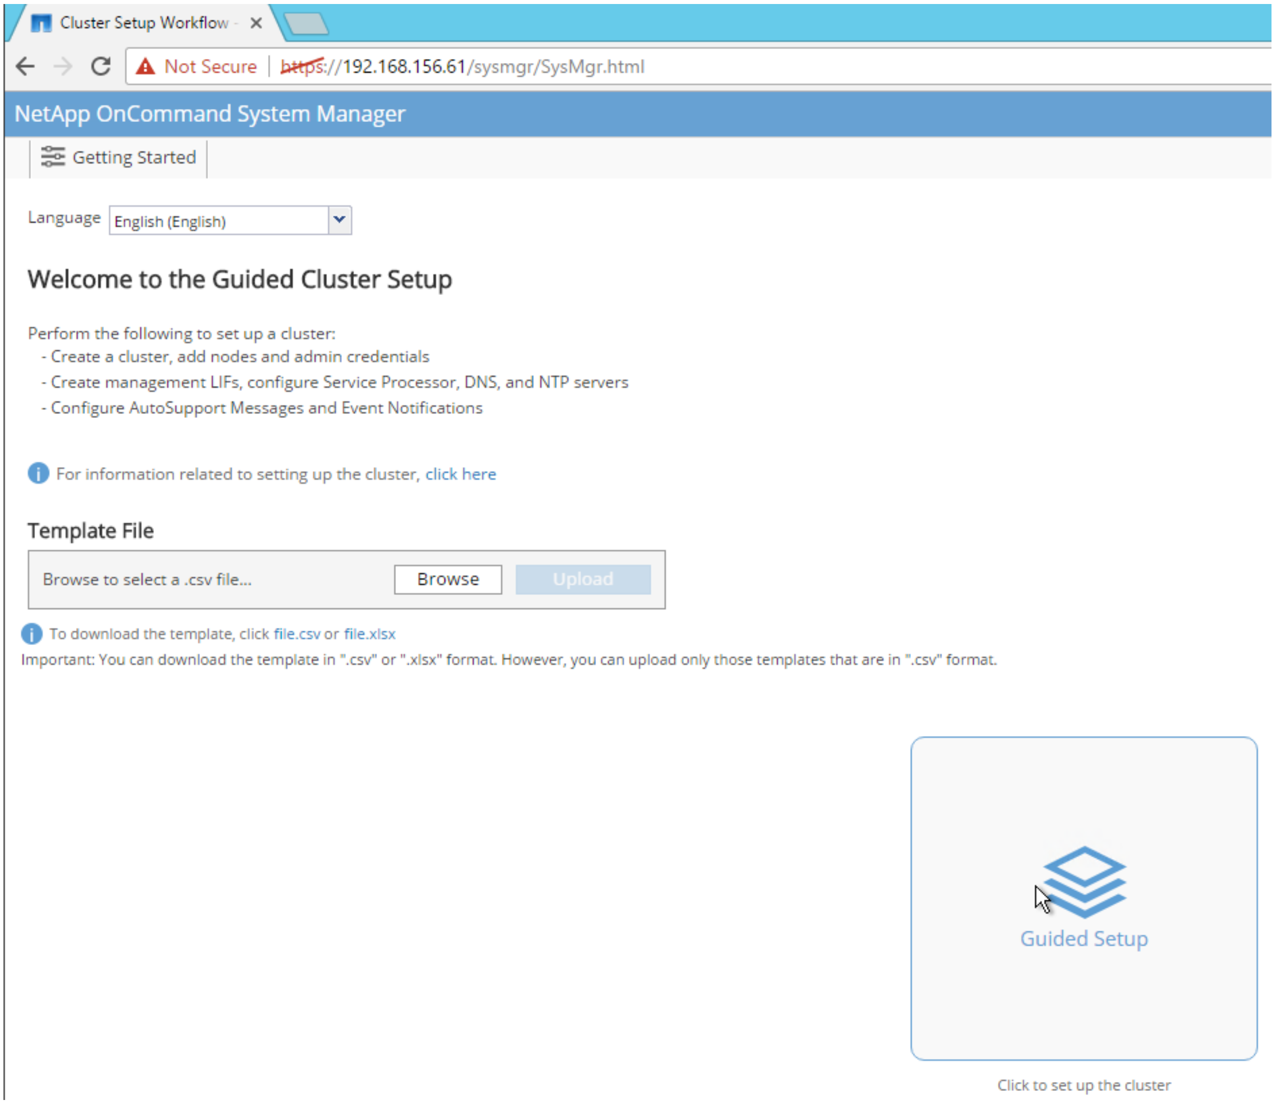

Cluster Setup Workflow
Not Secure | https://192.168.156.61/sysmgr/SysMgr.html
NetApp OnCommand System Manager
Getting Started
Language English (English)
Browse
Upload
Guided Setup

Welcome to the Guided Cluster Setup

Perform the following to set up a cluster:
- Create a cluster, add nodes and admin credentials
- Create management LIFs, configure Service Processor, DNS, and NTP servers
- Configure AutoSupport Messages and Event Notifications

For information related to setting up the cluster, click here

Template File

Browse to select a .csv file...

To download the template, click file.csv or file.xlsx

Important: You can download the template in ".csv" or ".xlsx" format. However, you can upload only those templates that are in ".csv" format.

Click to set up the cluster

4. In the Cluster screen, follow these steps:

   a. Enter the cluster and node names.

   b. Select the cluster configuration.

   c. Enter and confirm the password.

   d. (Optional) Enter the cluster base and feature licenses.

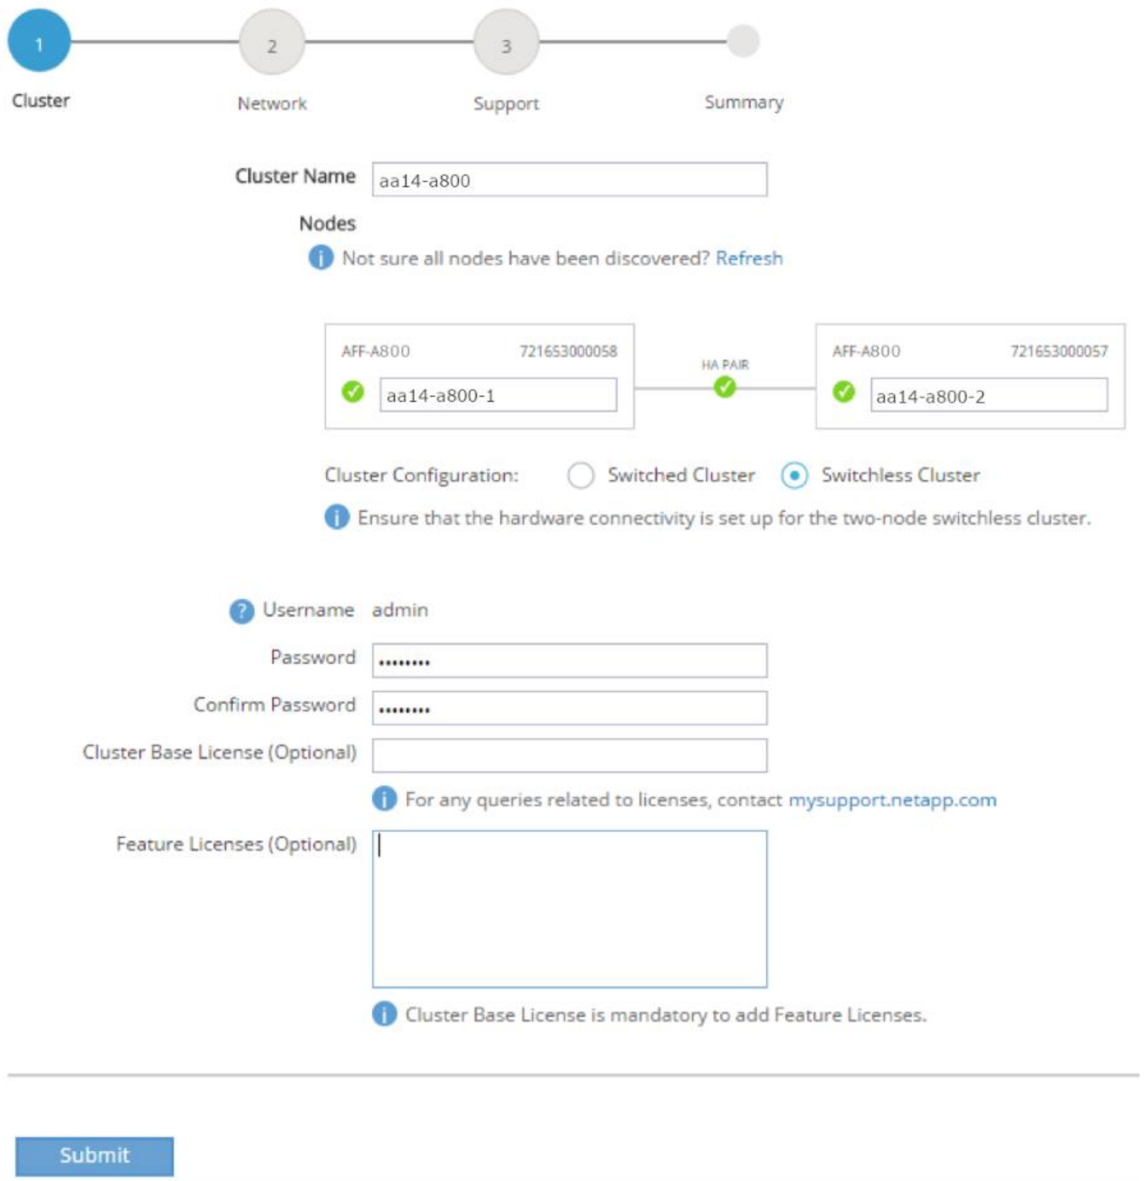

Cluster Name: aa14-a800

Nodes

Not sure all nodes have been discovered? Refresh

AFF-A800	721653000058	HA PAIR	AFF-A800	721653000057
✓ aa14-a800-1		✓	✓ aa14-a800-2	

Cluster Configuration: ○ Switched Cluster ● Switchless Cluster

Ensure that the hardware connectivity is set up for the two-node switchless cluster.

Username: admin

Password: ••••••••

Confirm Password: ••••••••

Cluster Base License (Optional):

For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional):

Cluster Base License is mandatory to add Feature Licenses.

The nodes should be discovered automatically; if they are not, click the Refresh link. By default, the cluster interfaces are created on all the new factory shipping storage controllers.

If all the nodes are not discovered, then configure the cluster using the command line.

Cluster license and feature licenses can also be installed after completing the cluster creation.

5. Click Submit.

6. In the network page, complete the following sections:

   a. Cluster Management

- Enter the IP address, netmask, gateway and port details.

b. Node Management

- Enter the node management IP addresses and port details for all the nodes.

c. Service Processor Management

- Enter the IP addresses for all the nodes.

d. DNS Details

- Enter the DNS domain names and server address.

e. NTP Details

- Enter the primary and alternate NTP server.

7. Click Submit.

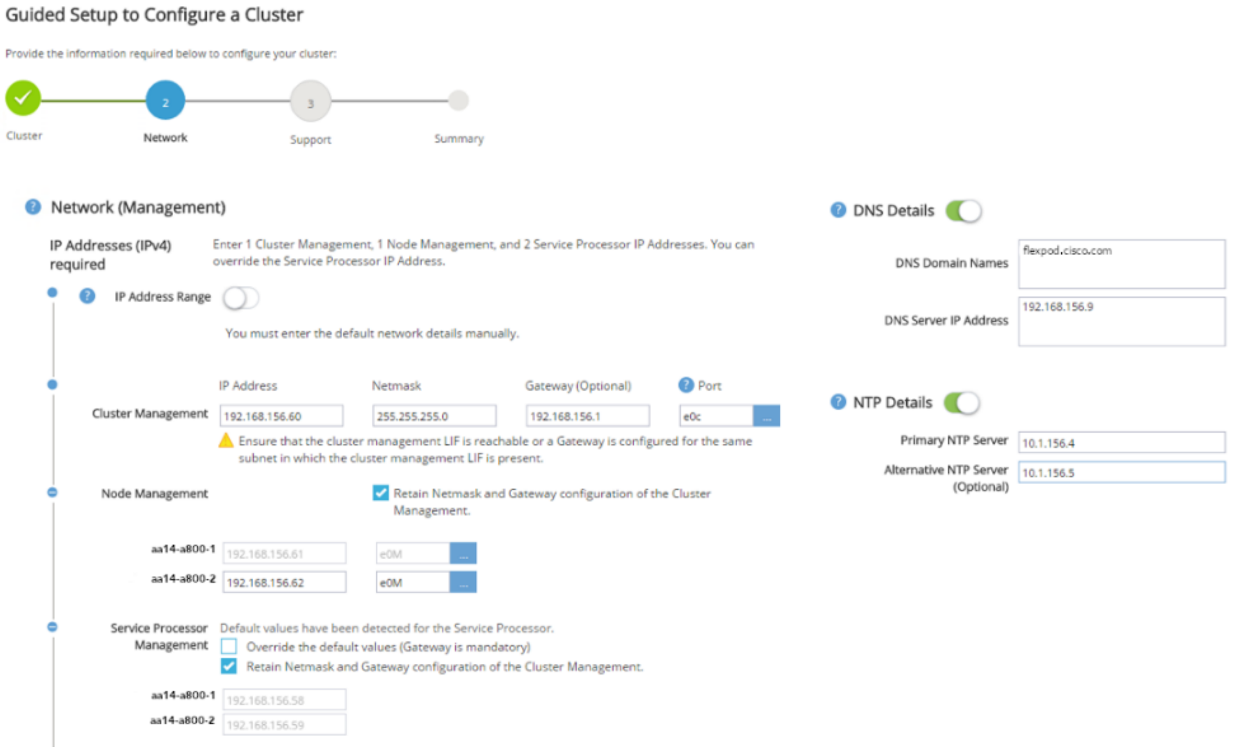

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Network (Management)

IP Addresses (IPv4) required: Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

IP Address Range: Off

You must enter the default network details manually.

	IP Address	Netmask	Gateway (Optional)	Port
Cluster Management	192.168.156.60	255.255.255.0	192.168.156.1	e0c

⚠ Ensure that the cluster management LIF is reachable or a Gateway is configured for the same subnet in which the cluster management LIF is present.

Node Management: ☑ Retain Netmask and Gateway configuration of the Cluster Management.

	IP Address	Port
aa14-a800-1	192.168.156.61	e0M
aa14-a800-2	192.168.156.62	e0M

Service Processor Management: Default values have been detected for the Service Processor.

□ Override the default values (Gateway is mandatory)

☑ Retain Netmask and Gateway configuration of the Cluster Management.

aa14-a800-1	192.168.156.58
aa14-a800-2	192.168.156.59

DNS Details: On

DNS Domain Names: flexpod.cisco.com

DNS Server IP Address: 192.168.156.9

NTP Details: On

Primary NTP Server: 10.1.156.4

Alternative NTP Server (Optional): 10.1.156.5

8. In the Support page, configure the AutoSupport and Event Notifications sections.

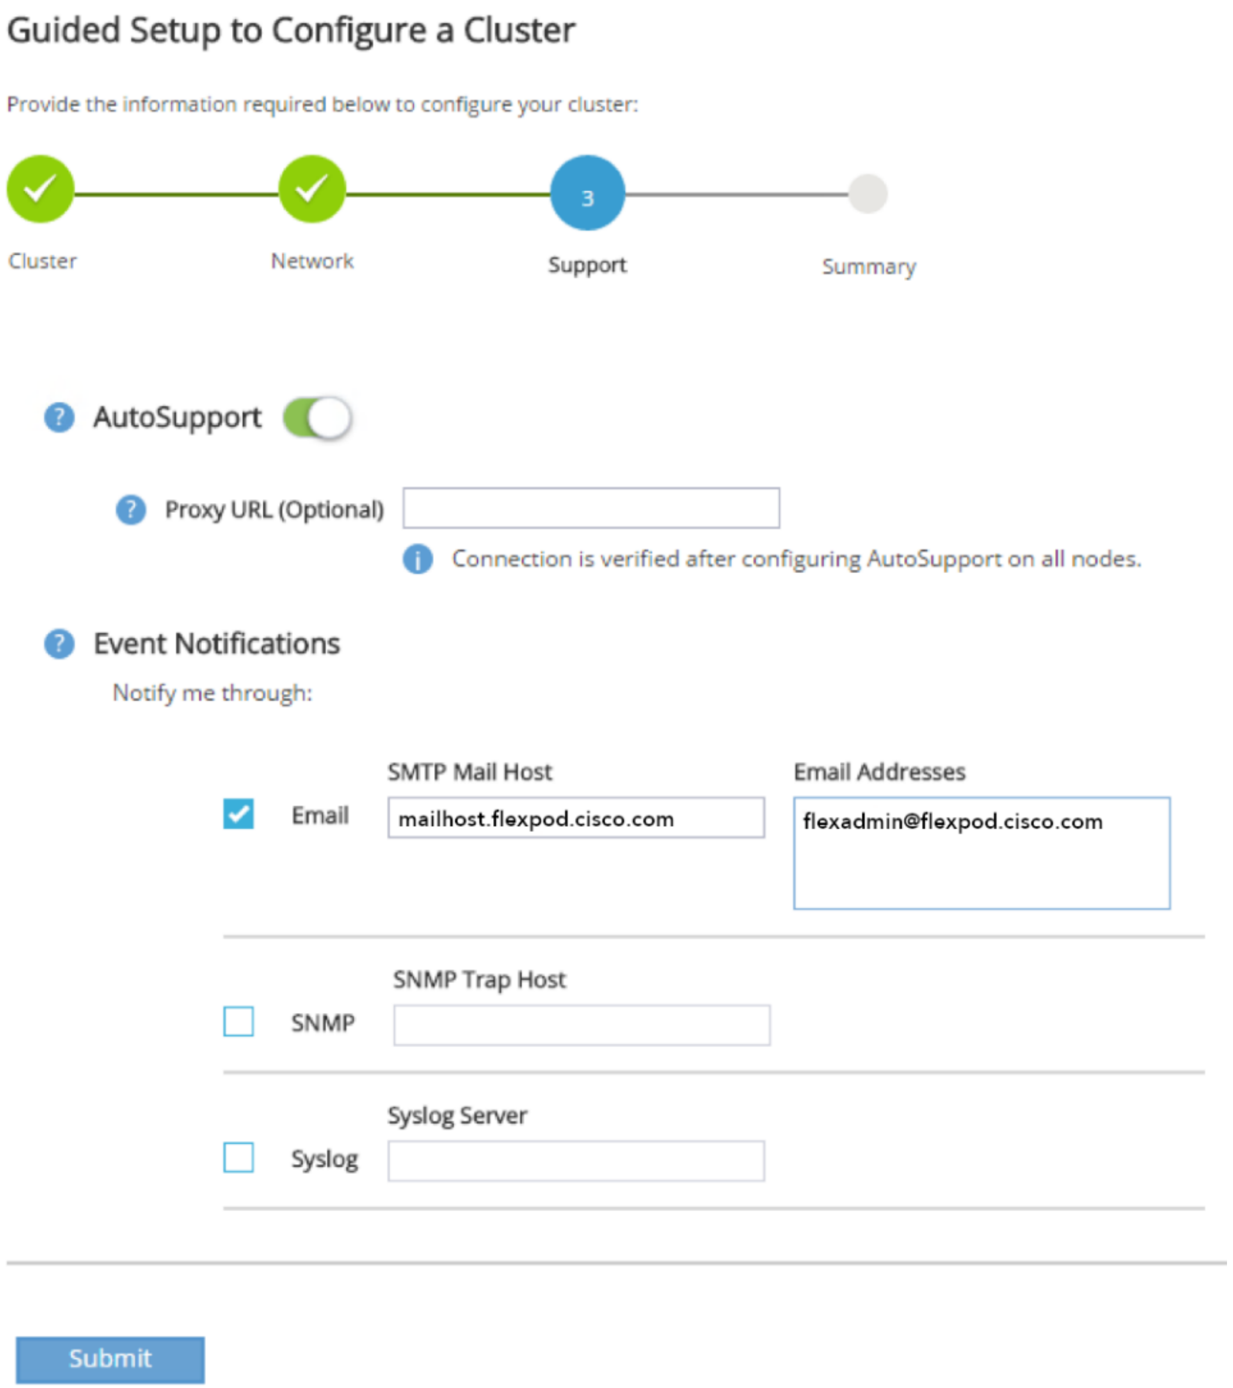

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

AutoSupport

Proxy URL (Optional)

Connection is verified after configuring AutoSupport on all nodes.

Event Notifications

Notify me through:

☑ Email

SMTP Mail Host: mailhost.flexpod.cisco.com

Email Addresses: flexadmin@flexpod.cisco.com

□ SNMP

SNMP Trap Host:

□ Syslog

Syslog Server:

9. Click Submit.

10. In the Summary page, review the configuration details if needed.

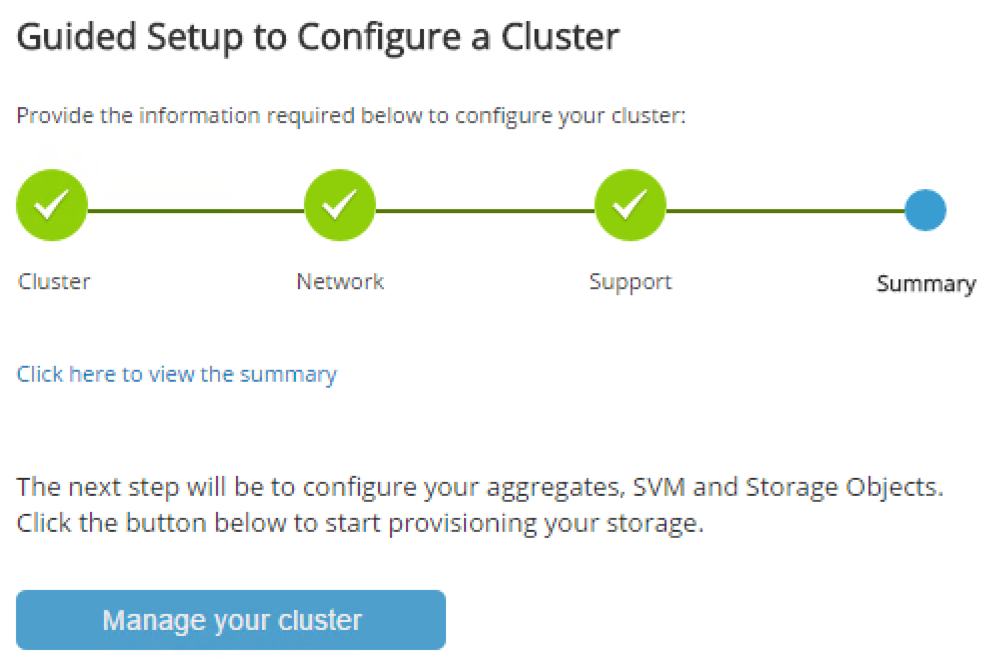

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

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

The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

Log into the Cluster

To log into the cluster, follow these steps:

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

Verify Storage Failover

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

1. Verify the status of the storage failover.

```
storage failover show
```

Both `<st-node01>` **and** `<st-node02>` **must be capable of performing a takeover. Continue with step 2 if the nodes are capable of performing a takeover.**

2. Enable failover on one of the two nodes if it was not completed during the installation.

```
storage failover modify -node <st-node01> -enabled true
```

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

4. Continue with step 5 if high availability is not configured.

5. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

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

```
storage failover hwassist show
storage failover modify –hwassist-partner-ip <node02-mgmt-ip> -node <st-node01>
storage failover modify –hwassist-partner-ip <node01-mgmt-ip> -node <st-node02>
```

Set Auto-Revert on Cluster Management

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

A storage virtual machine (SVM) is referred to as a Vserver or `vserver` in the GUI and CLI.

Run the following command:

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

Zero All Spare Disks

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

```
disk zerospares
```

Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an All Flash FAS configuration. Disk autoassign should have assigned one data partition to each node in an HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Up Service Processor Network Interface

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

```
system service-processor network modify –node <st-node01> -address-family IPv4 –enable true –dhcp
none –ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>

system service-processor network modify –node <st-node02> -address-family IPv4 –enable true –dhcp
none –ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```

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

Create Aggregates

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

To create new aggregates, follow these steps:

1. Run the following commands:

```
storage aggregate create -aggregate aggr1_node01 -node <st-node01> -diskcount <num-disks>
storage aggregate create -aggregate aggr1_node02 -node <st-node02> -diskcount <num-disks>
```

You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate.

For all flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.

Start with five disks initially; you can add disks to an aggregate when additional storage is required. In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all but one remaining disk (spare) assigned to the controller.

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

Remove Ports from Default Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, `e5a`, `e5b,` and so on) should be removed from the default broadcast domain, leaving just the management network port (`e0M`). To perform this task, run the following commands:

```
network broadcast-domain remove-ports  -broadcast-domain Default -ports aa14-a800-1:e3a,aa14-a800-
1:e3b,aa14-a800-1:e3c,aa14-a800-1:e3d,aa14-a800-1:e5a,aa14-a800-1:e5b,aa14-a800-2:e3a,aa14-a800-
2:e3b,aa14-a800-2:e3c,aa14-a800-2:e3d,aa14-a800-2:e5a,aa14-a800-2:e5b

network port broadcast-domain show
```

Disable Flow Control on 10GbE and 100GbE ports

NetApp recommends disabling flow control on all the 10/40/100GbE and UTA2 ports that are connected to external devices. To disable flow control, follow these steps:

1. Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port e0a,e0b,e1a,e1b,e3a,e3b,e3c,e3d,e5a,e5b -flowcontrol-
admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <st-node02> -port e0a,e0b,e1a,e1b,e3a,e3b,e3c,e3d,e5a,e5b -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

network port show –fields flowcontrol-admin
```

Disable Auto-negotiate on 100GbE Ports

To disable auto-negotiate on the 100GbE ports, follow these steps:

1. Run the following command to configure the ports on node 01:

```
network port modify -node <st-node01> -port e5a,e5b -autonegotiate-admin false -speed-admin 100000 -
duplex-admin full -flowcontrol-admin none
```

2. Run the following command to configure the ports on node 02:

```
network port modify -node <st-node02> -port e5a,e5b -autonegotiate-admin false -speed-admin 100000 -
duplex-admin full -flowcontrol-admin none

network port show -node * -port e5a,e5b -fields speed-admin,duplex-admin,flowcontrol-admin
  (network port show)
node        port duplex-admin speed-admin flowcontrol-admin
----------- ---- ------------ ----------- -----------------
aa14-a800-1 e5a  full         100000      none
aa14-a800-1 e5b  full         100000      none
aa14-a800-2 e5a  full         100000      none
aa14-a800-2 e5b  full         100000      none
4 entries were displayed.
```

Disable Auto-negotiate on Fibre Channel Ports

In accordance with best practices for Fibre Channel host ports, disable auto-negotiate on each FCP adapter in each controller node.

1. Disable each fibre channel adapter in the controllers with the `fcp adapter modify` command.

```
fcp adapter modify -node <st-node01> -adapter 2a –status-admin down
fcp adapter modify -node <st-node01> -adapter 2b –status-admin down
fcp adapter modify -node <st-node02> -adapter 2a –status-admin down
fcp adapter modify -node <st-node02> -adapter 2b –status-admin down
```

2. Set the desired speed on the adapter and return it to the online state.

```
fcp adapter modify -node <st-node01> -adapter 2a -speed 32 -status-admin up
fcp adapter modify -node <st-node01> -adapter 2b -speed 32 -status-admin up
fcp adapter modify -node <st-node02> -adapter 2a -speed 32 -status-admin up
fcp adapter modify -node <st-node02> -adapter 2b -speed 32 -status-admin up
```

Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command to enable CDP on ONTAP:

```
node run -node * options cdpd.enable on
```

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

Enable Link-layer Discovery Protocol on all Ethernet Ports

Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches with the following step:

1. Enable LLDP on all ports of all nodes in the cluster.

```
node run * options lldp.enable on
```

Create Management Broadcast Domain

If the management interfaces are required to be on a separate VLAN, create a new broadcast domain for those interfaces.

```
network port broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500
network port broadcast-domain show
```

Create NFS Broadcast Domain

To create an NFS data broadcast domain with an MTU of 9000, run the following commands to create a broadcast domain for NFS in ONTAP:

```
network port broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
network port broadcast-domain show
```

Create Interface Groups

To create the LACP interface groups for the 100GbE data interfaces, run the following commands:

```
network port ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e5a
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e5b

network port ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e5a
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e5b

network port ifgrp show
```

Change MTU on Interface Groups

Change the MTU size on the base Interface Group ports before creating the VLAN ports.

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

Create VLANs

To create VLANs, follow these steps:

1. Create the management VLAN ports and adding them to the Management broadcast domain.

```
network port vlan create -node <st-node01> -vlan-name a0a-<ib-mgmt-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<ib-mgmt-vlan-id>

network port broadcast-domain add-ports -broadcast-domain IB-MGMGT -ports <st-node01>:a0a-<ib-mgmt-
vlan-id>,<st-node02>:a0a-<ib-mgmt-vlan-id>

network port vlan show
```

2. Create the NFS VLAN ports and add them to the `Infra_NFS` broadcast domain.

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <st-node01>:a0a-<infra-
nfs-vlan-id>,<st-node02>:a0a-<infra-nfs-vlan-id>
```

Configure Network Time Protocol

To configure time synchronization on the cluster, follow these steps:

1. Set the time zone for the cluster.

```
timezone <timezone>
```

For example, in the eastern United States, the time zone is `America/New_York`.

2. Set the date for the cluster.

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

The format for the date is <[Century][Year][Month][Day][Hour][Minute].[Second]> (for example, 201903271549.30).

3. Configure the Network Time Protocol (NTP) servers for the cluster.

```
cluster time-service ntp server create -server <nexus-A-mgmt0-ip>
cluster time-service ntp server create -server <nexus-B-mgmt0-ip>
```

Configure Simple Network Management Protocol

To configure the Simple Network Management Protocol (SNMP), follow these steps:

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

```
snmp contact <snmp-contact>
snmp location “<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 <oncommand-um-server-fqdn>
```

Configure SNMPv1 Access

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

```
snmp community add ro <snmp-community>
```

To enable SNMPv3, refer to the SNMP Configuration Express Guide to configure SNMPv3 and setup one or more security users.

Create SVM

To create an infrastructure SVM, follow these steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume infra_svm_root -aggregate aggr1_node01 -rootvolume-
security-style unix
```

2. Remove the unused data protocols from the SVM: CIFS, iSCSI and NVMe.

```
vserver remove-protocols -vserver Infra-SVM -protocols iscsi,cifs
```

3. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the NFS protocol in the Infra-SVM.

```
vserver nfs create -vserver Infra-SVM -udp disabled
```

If the NFS license was not installed during the cluster configuration, make sure to install the license before starting the NFS service.

5. Turn on the SVM vstorage parameter for the NetApp NFS VAAI plug-in.

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

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, follow these steps:

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

```
volume create -vserver Infra-SVM -volume infra_svm_root_m01 -aggregate aggr1_node01 -size 1GB -type
DP

volume create -vserver Infra-SVM -volume infra_svm_root_m02 -aggregate aggr1_node02 -size 1GB -type
DP
```

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

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

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-
SVM:infra_svm_root_m01 -type LS -schedule 15min

snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-
SVM:infra_svm_root_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship.

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

Create Block Protocol (FC) Service

Run the following command to create the FCP service on each SVM. This command also starts the FCP service and sets the WWN for the SVM.

```
vserver fcp create -vserver Infra-SVM
vserver fcp show
```

If the FC license was not installed during the cluster configuration, make sure to install the license before creating the FC service.

Configure HTTPS Access

To configure secure access to the storage controller, follow these steps:

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

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

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, the <serial-number>) by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -
serial <serial-number>
```

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

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

```
security certificate create -common-name <cert-common-name> -type  server -size 2048 -country <cert-
country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -
email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver Infra-
SVM
```

5. To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the security certificate show command.

6. Enable each certificate that was just created by using the –server-enabled true and –client-enabled false parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -
serial <cert-serial> -common-name <cert-common-name>
```

7. Disable HTTP cluster management access.

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```

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

8. Change back to the normal admin privilege level and verify the system logs are available in a web browser.

```
set -privilege admin
```

```
https://<node01-mgmt-ip>/spi

https://<node02-mgmt-ip>/spi
```

Configure NFSv3

To configure NFSv3 on the SVM, follow these steps:

1. Create a new rule for the infrastructure NFS subnet in the default export policy.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -
clientmatch <infra-nfs-subnet-cidr> -rorule sys -rwrule sys -superuser sys -allow-suid false
```

2. Assign the FlexPod export policy to the infrastructure SVM root volume.

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

Create FlexVol Volumes

The following information is required to create a NetApp FlexVol® volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

To create a FlexVol volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node01 -size 500GB -state
online -policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-snapshot-
space 0

volume create -vserver Infra-SVM -volume infra_datastore_2 -aggregate aggr1_node02 -size 500GB -state
online -policy default -junction-path /infra_datastore_2 -space-guarantee none -percent-snapshot-
space 0

volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_node01 -size 100GB -state online
-policy default -junction-path /infra_swap -space-guarantee none -percent-snapshot-space 0 -
efficiency-policy none

volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_node01 -size 100GB -state online
-policy default -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path Infra-SVM:rootvol
```

Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -size 15GB -ostype vmware -space-
reserve disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-02 -size 15GB -ostype vmware -space-
reserve disabled
```

Schedule Deduplication

On NetApp All Flash FAS systems, deduplication is enabled by default. To schedule deduplication, complete the following step:

1. After the volumes are created, assign a once-a-day deduplication schedule to `esxi_boot, infra_datastore_1 and infra_datastore_2`:

```
volume efficiency modify –vserver Infra-SVM –volume esxi_boot –schedule sun-sat@0
volume efficiency modify –vserver Infra-SVM –volume infra_datastore_1 –schedule sun-sat@0
volume efficiency modify –vserver Infra-SVM –volume infra_datastore_2 –schedule sun-sat@0
```

Create FC LIFs

Run the following commands to create four FC LIFs (two on each node):

```
network interface create –vserver Infra-SVM -lif fcp-lif01a -role data -data-protocol fcp -home-node
<st-node01> -home-port 2a –status-admin up

network interface create –vserver Infra-SVM -lif fcp-lif01b -role data -data-protocol fcp -home-node
<st-node01> -home-port 2b –status-admin up
```

```
network interface create -vserver Infra-SVM -lif fcp-lif02a -role data -data-protocol fcp -home-node
<st-node02> -home-port 2a -status-admin up

network interface create -vserver Infra-SVM -lif fcp-lif02b -role data -data-protocol fcp -home-node
<st-node02> -home-port 2b -status-admin up

network interface show
```

Create NFS LIFs

To create NFS LIFs, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs-lif01 -role data -data-protocol nfs -home-node
<st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs_lif01-ip> -netmask <node01-
nfs_lif01-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface create -vserver Infra-SVM -lif nfs-lif02 -role data -data-protocol nfs -home-node
<st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs_lif02-ip> -netmask <node02-
nfs_lif02-mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface show
```

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, follow these steps:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif svm-mgmt -role data -data-protocol none -home-node
<st-node02> -home-port  a0a-<ib-mgmt-vlan-id> -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```

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

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-gateway>

network route show
```

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

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

security login unlock -username vsadmin -vserver Infra-SVM
```

A cluster serves data through at least one and possibly several SVMs. We have just gone through creating a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create additional SVMs.

Configure and Test AutoSupport

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

```
system node autosupport modify -node * -state enable –mail-hosts <mailhost> -transport https -support
enable -noteto <storage-admin-email>
```

Test the AutoSupport configuration by sending a message from all nodes of the cluster:

```
autosupport invoke -node * -type all -message “FlexPod storage configuration completed”
```

Cisco UCS Configuration

Cisco UCS Base Configuration

This FlexPod deployment explains the configuration steps for the Cisco UCS 6454 Fabric Interconnects (FI) in a design that will support Fibre Channel SAN boot. If setting up a system with iSCSI boot, the sections with (FCP) in the heading can be skipped and then complete the Cisco UCS iSCSI Configuration section in the Appendix.

Perform Initial Setup of Cisco UCS 6454 Fabric Interconnects for FlexPod Environments

This section provides the 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 B-Series and C-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlexPod environment, follow these steps:

1. Connect to the console port on the first Cisco UCS 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]: Enter

Enter the password for "admin": <password>
Confirm the password for "admin": <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:  <ucs-cluster-name>

Physical Switch Mgmt0 IP address : <ucsa-mgmt-ip>

Physical Switch Mgmt0 IPv4 netmask : <ucsa-mgmt-mask>

IPv4 address of the default gateway : <ucsa-mgmt-gateway>

Cluster IPv4 address : <ucs-cluster-ip>

Configure the DNS Server IP address? (yes/no) [n]: y

  DNS IP address : <dns-server-1-ip>

Configure the default domain name? (yes/no) [n]: y

  Default domain name : <ad-dns-domain-name>

Join centralized management environment (UCS Central)? (yes/no) [n]: Enter

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

2. Wait for the login prompt for UCS Fabric Interconnect A before proceeding to the next section.

Cisco UCS Fabric Interconnect B

To configure the Cisco UCS for use in a FlexPod environment, follow these steps:

1. Connect to the console port on the second Cisco UCS 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: <password>
    Connecting to peer Fabric interconnect... done
    Retrieving config from peer Fabric interconnect... done
    Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
    Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucsa-mgmt-mask>
    Cluster IPv4 address          : <ucs-cluster-ip>

    Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address

  Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>

  Local fabric interconnect model(UCS-FI-6454)
  Peer fabric interconnect is compatible with the local fabric interconnect. Continuing with the
installer...

  Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

2. Wait for the login prompt for UCS Fabric Interconnect B before proceeding to the next section.

Cisco UCS Setup

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, follow these steps:

1. Open a web browser and navigate to the Cisco UCS fabric interconnect cluster address.

You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to come up.

2. Click the Launch UCS Manager link to launch Cisco 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. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 4.0(2d)

This document assumes the use of Cisco UCS 4.0(2d). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.0(2d), refer to Cisco UCS Manager Install and Upgrade Guides.

Anonymous Reporting

To enable anonymous reporting, follow this step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products. If you select Yes, enter the IP address of your SMTP Server. Click OK.

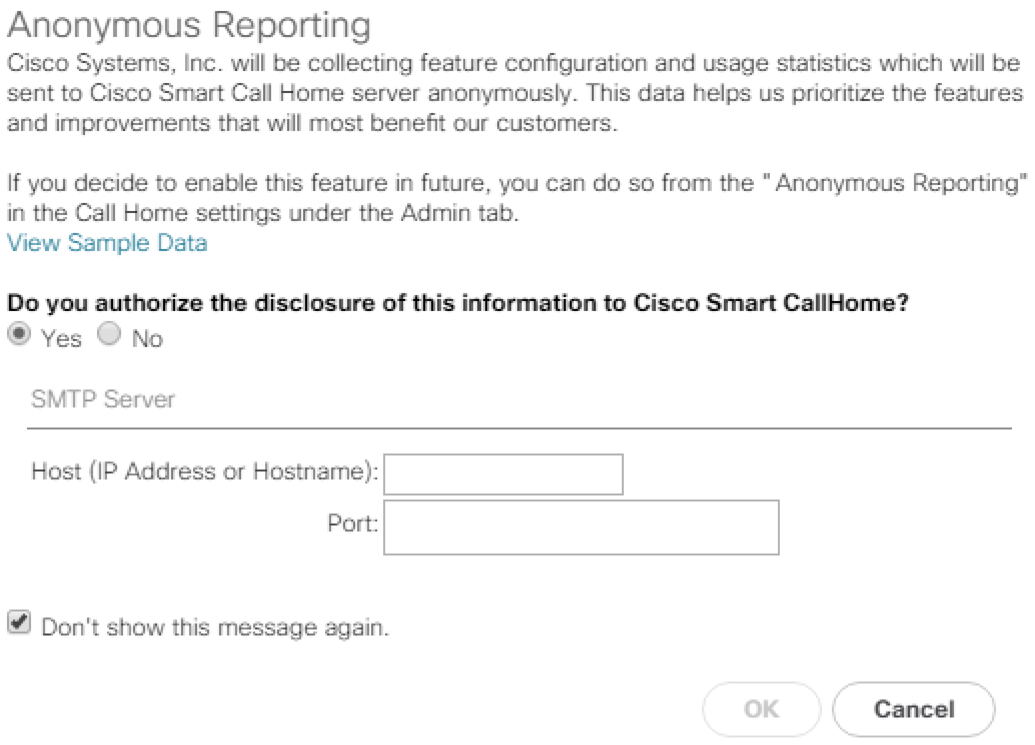

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.

If you decide to enable this feature in future, you can do so from the " Anonymous Reporting" in the Call Home settings under the Admin tab.

View Sample Data

Do you authorize the disclosure of this information to Cisco Smart CallHome?

● Yes ○ No

SMTP Server

Host (IP Address or Hostname):

Port:

☑ Don't show this message again.

Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in Cisco UCS Manager. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

Configure Unified Ports (FCP)

Fibre Channel port configurations differ between the 6454, 6332-16UP and the 6248UP Fabric Interconnects. All Fabric Interconnects have a slider mechanism within the Cisco UCS Manager GUI interface, but the fibre channel port selection options for the 6454 are from the first 8 ports starting from the first port and configured in increments of 4 ports from the left. For the 6332-16UP the port selection options are from the first 16 ports starting from the first port , and configured in increments of the first 6, 12, or all 16 of the unified ports. With the 6248UP, the port selection options will start from the right of the 32 fixed ports, or the right of the 16 ports of the expansion module, going down in contiguous increments of 2. The remainder of this section shows configuration of the 6454. Modify as necessary for the 6332-16UP or 6248UP.

To enable the fibre channel ports, follow these steps for the 6454:

1. In Cisco UCS Manager, click Equipment .

2. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate).

3. Select Configure Unified Ports.

4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.

5. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 4 or 8 ports to be set as FC Uplinks.

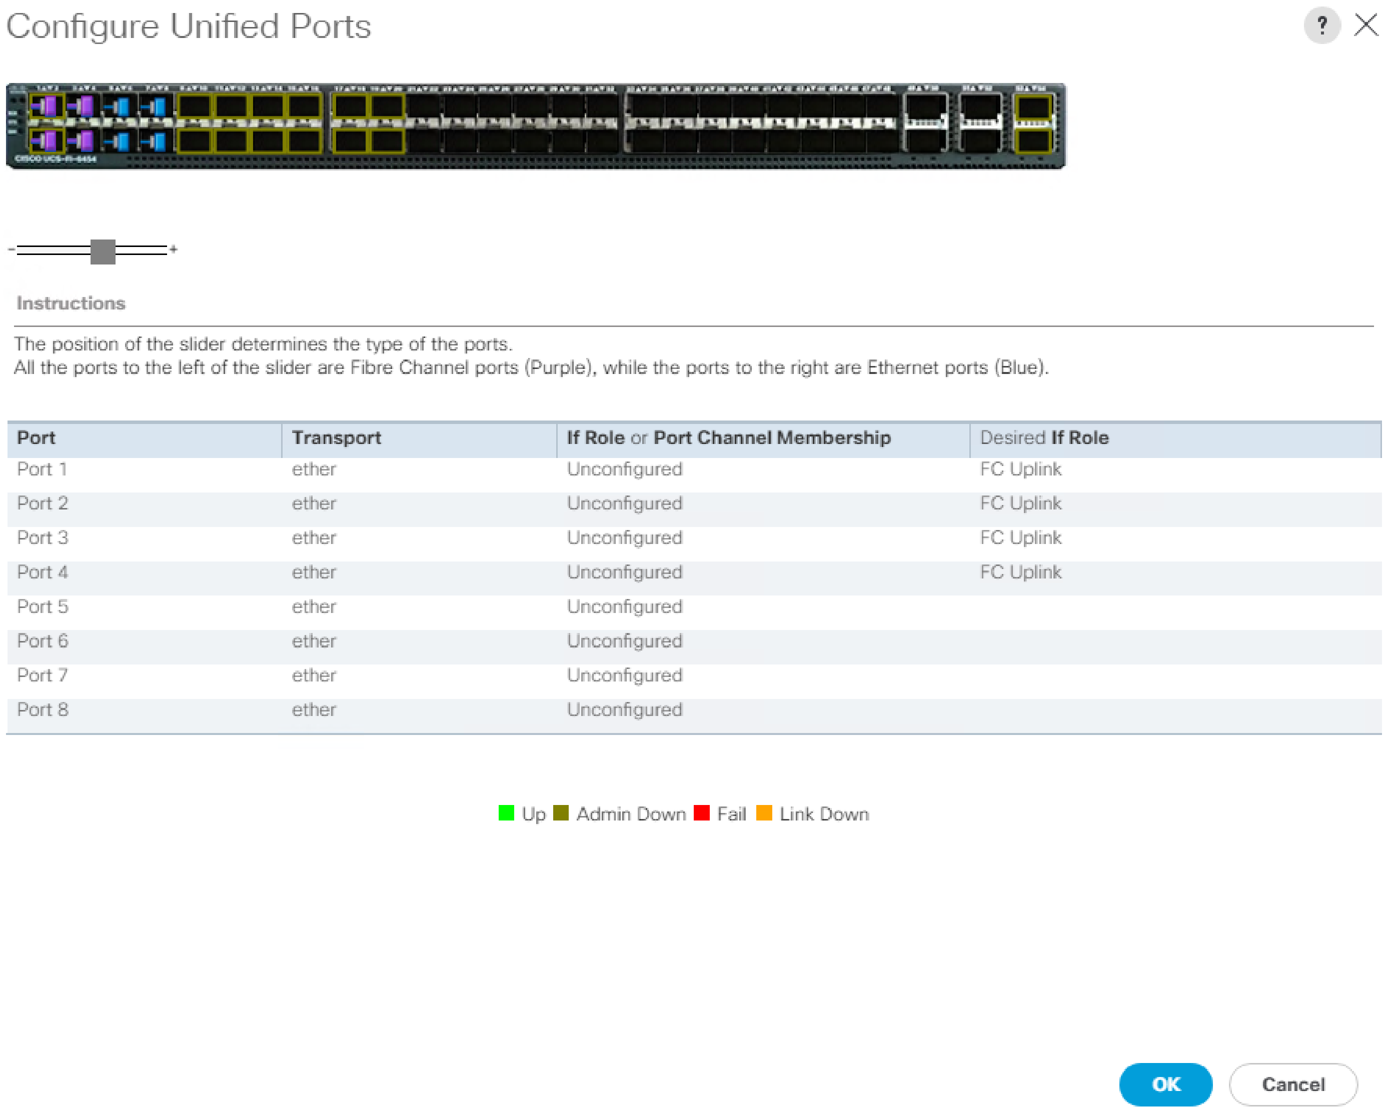

Configure Unified Ports

Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	

🟩 Up 🟫 Admin Down 🟥 Fail 🟧 Link Down

6. Click OK, then click Yes, then click OK to continue.

7. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).

8. Select Configure Unified Ports.

9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.

10. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 4 or 8 ports to be set as FC Uplinks.

11. Click OK, then click Yes, then OK to continue.

12. Wait for both Fabric Interconnects to reboot.

13. Log back into Cisco UCS Manager.

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for in band server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN .
2. Expand Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information. Optionally, enter the Primary and Secondary DNS server addresses.

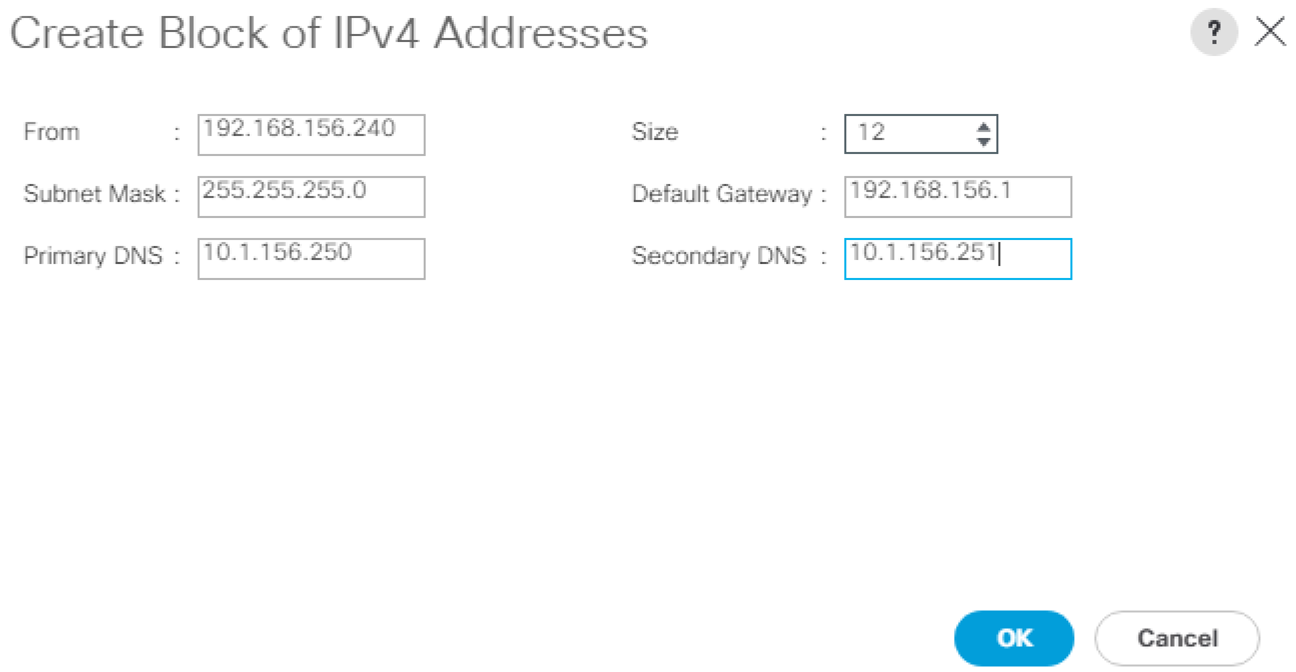

Create Block of IPv4 Addresses

From :	192.168.156.240	Size :	12
Subnet Mask :	255.255.255.0	Default Gateway :	192.168.156.1
Primary DNS :	10.1.156.250	Secondary DNS :	10.1.156.251

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

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand All > Time Zone Management.
3. Select Timezone.

4. In the Properties pane, select the appropriate time zone in the Timezone menu.

5. Click Save Changes, and then click OK.

6. Click Add NTP Server.

7. Enter <nexus-A-mgmt0-ip> and click OK. Click OK on the confirmation.

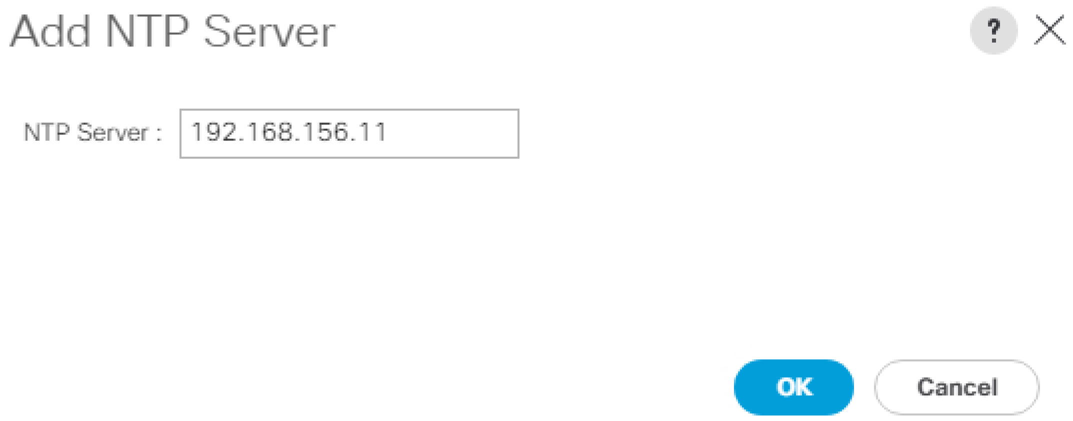

Add NTP Server

NTP Server : 192.168.156.11

We used the Nexus switch mgmt0 interface IP here because it is in the same L2 domain as the UCS mgmt0 IPs. We could also use the Nexus NTP IPs, but that traffic would then have to pass through an L3 router.

8. Click Add NTP Server.

9. Enter <nexus-B-mgmt0-ip> and click OK. Click OK.

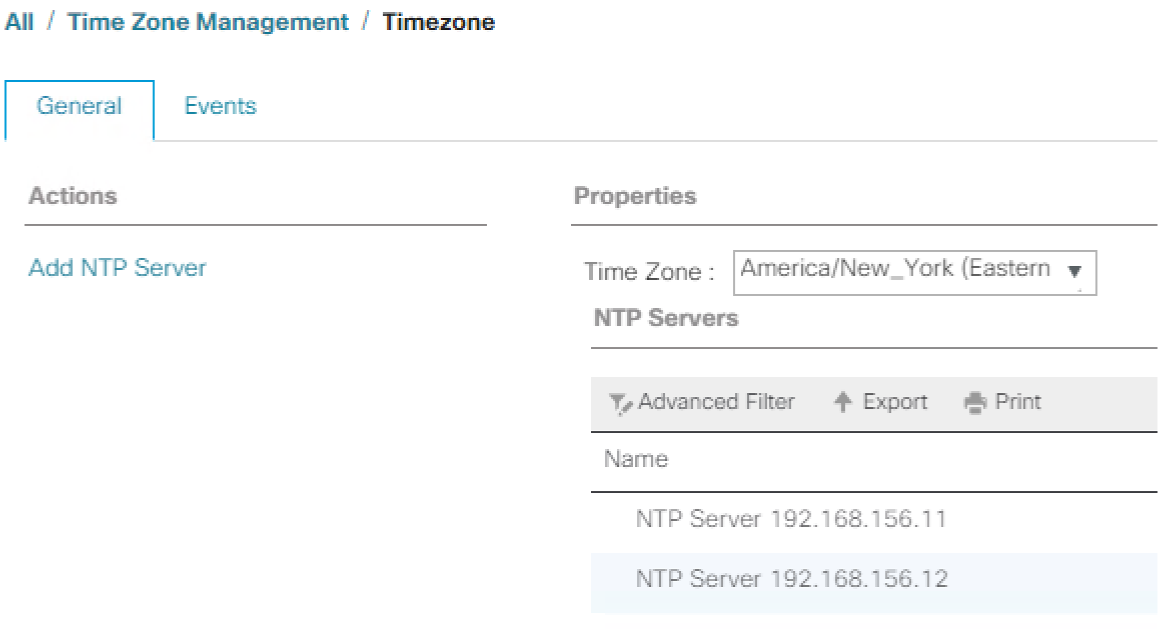

Time Zone : America/New_York (Eastern

NTP Servers

NTP Server 192.168.156.11

NTP Server 192.168.156.12

Add Additional DNS Server(s)

To add one or more additional DNS servers to the UCS environment, follow these steps:

1. In Cisco UCS Manager, click Admin .
2. Expand All > Communications Management.
3. Select DNS Management.
4. In the Properties pane, select Specify DNS Server.
5. Enter the IP address of the additional DNS server.

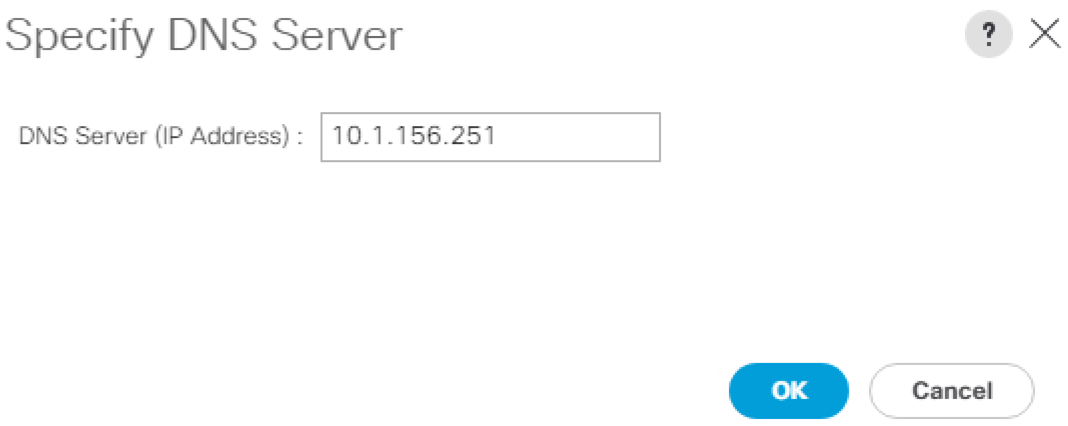

DNS Server (IP Address) : 10.1.156.251

6. Click OK and then click OK again. Repeat this process for any additional DNS servers.

Add an Additional Administrative User

To add an additional locally authenticated Administrative user (flexadmin) to the Cisco UCS environment in case issues arise with the admin user, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand User Management > User Services > Locally Authenticated Users.
3. Right-click Locally Authenticated Users and select Create User.
4. In the Create User fields it is only necessary to fill in the Login ID, Password, and Confirm Password fields. Fill in the Create User fields according to your local security policy.
5. Leave the Account Status field set to Active.
6. Set Account Expires according to your local security policy.
7. Under Roles, select admin.
8. Leave Password Required selected for the SSH Type field.

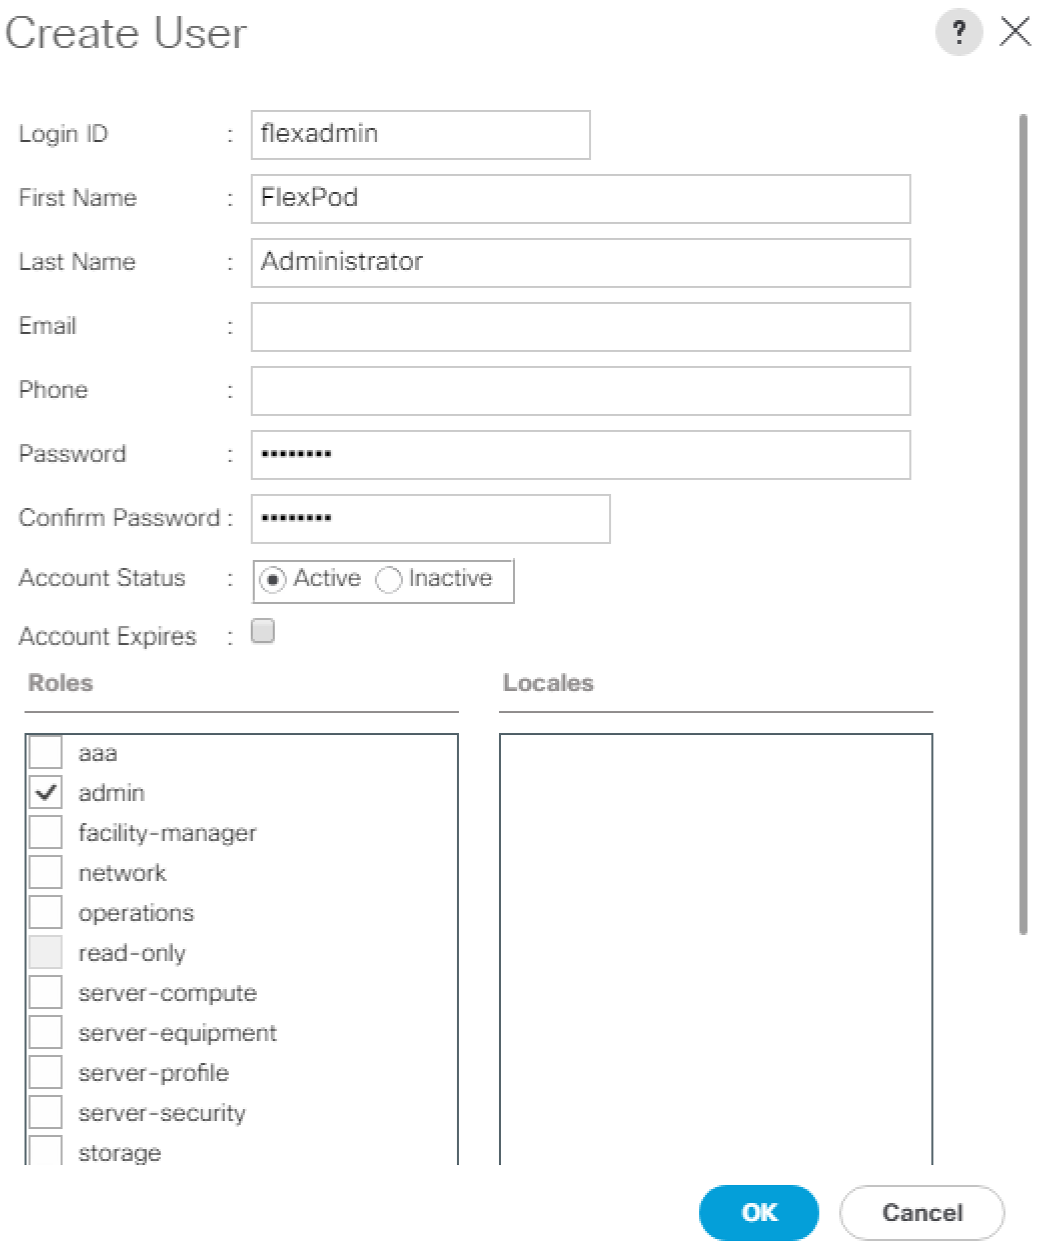

Create User

Login ID : flexadmin

First Name : FlexPod

Last Name : Administrator

Email :

Phone :

Password : •••••••••

Confirm Password : •••••••••

Account Status : ● Active ○ Inactive

Account Expires : □

Roles

- □ aaa
- ☑ admin
- □ facility-manager
- □ network
- □ operations
- □ read-only
- □ server-compute
- □ server-equipment
- □ server-profile
- □ server-security
- □ storage

Locales

9. Click OK and then click OK again to complete adding the user.

Enable Port Auto-Discovery Policy

Setting the port auto-discovery policy enables automatic discovery of Cisco UCS B-Series chassis server ports. To modify the port auto-discovery policy, follow these steps:

1. In Cisco UCS Manager, click Equipment , select All > Equipment in the Navigation Pane, and select the Policies tab on the right.

2. Under Port Auto-Discovery Policy, set Auto Configure Server Port to Enabled.

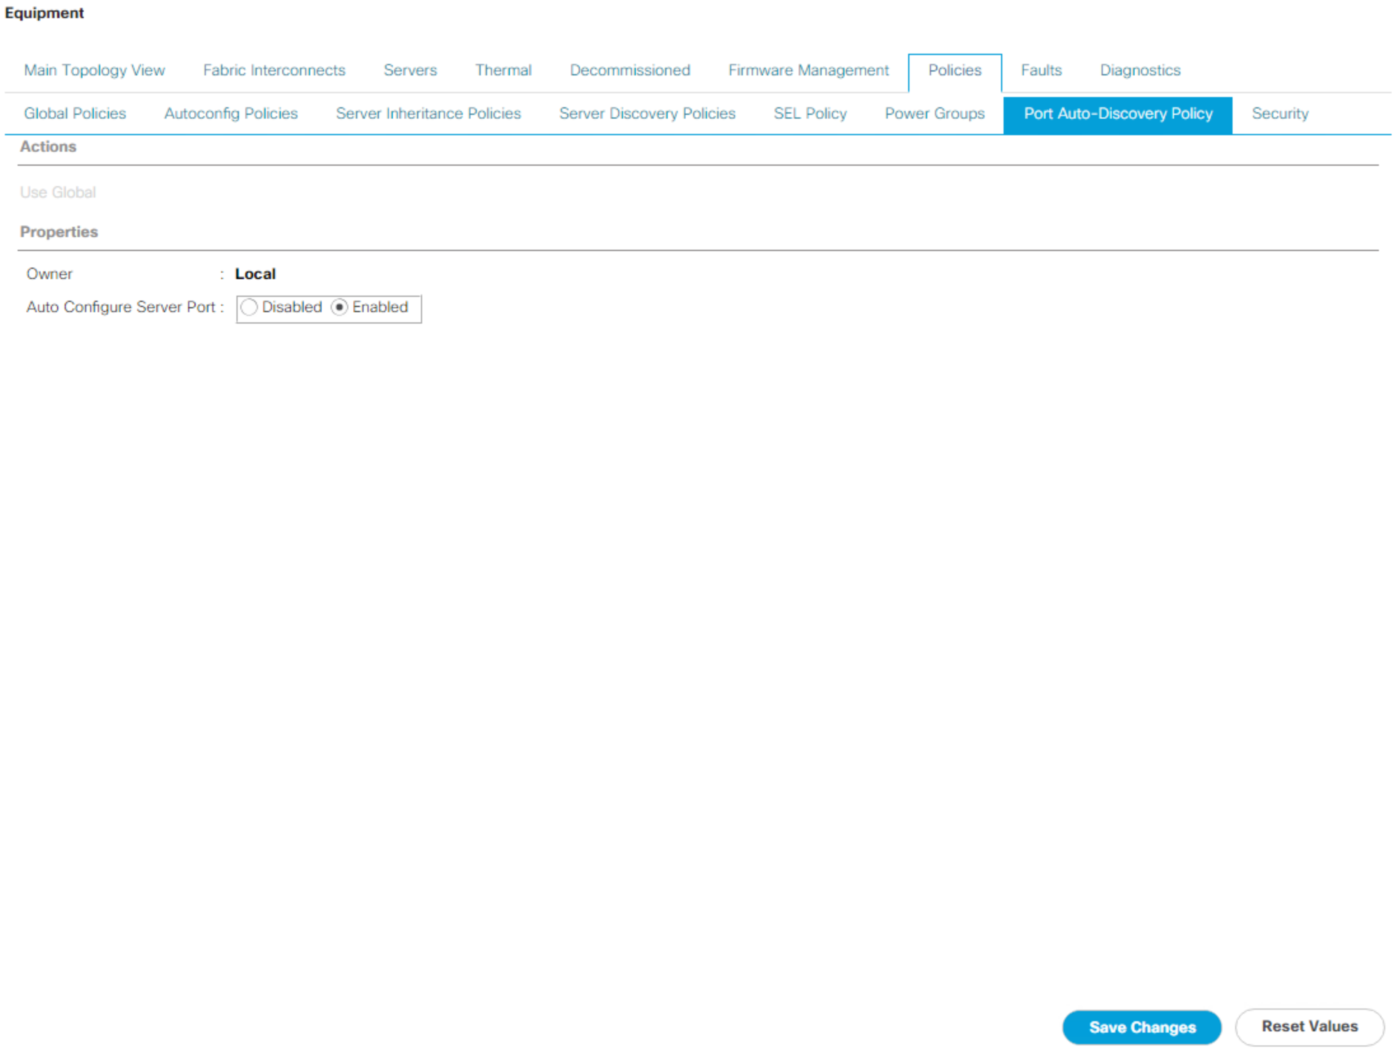

Owner : Local

Auto Configure Server Port : ○ Disabled ● Enabled

3. Click Save Changes and then OK.

Enable Info Policy for Neighbor Discovery

Enabling the info policy enables Fabric Interconnect neighbor information to be displayed. To modify the info policy, follow these steps:

1. In Cisco UCS Manager, click Equipment , select All > Equipment in the Navigation Pane, and select the Policies tab on the right.

2. Under Global Policies, scroll down to Info Policy and select Enabled for Action.

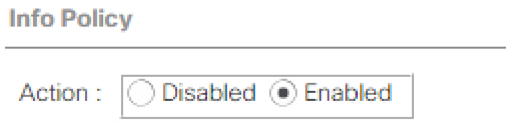

Action : ○ Disabled ● Enabled

3. Click Save Changes and then OK.

4. Under Equipment, select Fabric Interconnect A (primary). On the right, select the Neighbors tab. CDP information is shown under the LAN tab and LLDP information is shown under the LLDP tab.

Edit Chassis Discovery Policy

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

1. In Cisco UCS Manager, click Equipment and select the Policies tab on the right.

2. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.

If varying numbers of links between chassis and the Fabric Interconnects will be used, leave Action set to 1 Link.

3. On the 6454 Fabric Interconnects, the Link Grouping Preference is automatically set to Port Channel and is greyed out. On a 6300 Series or 6200 Series Fabric Interconnect, set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G.

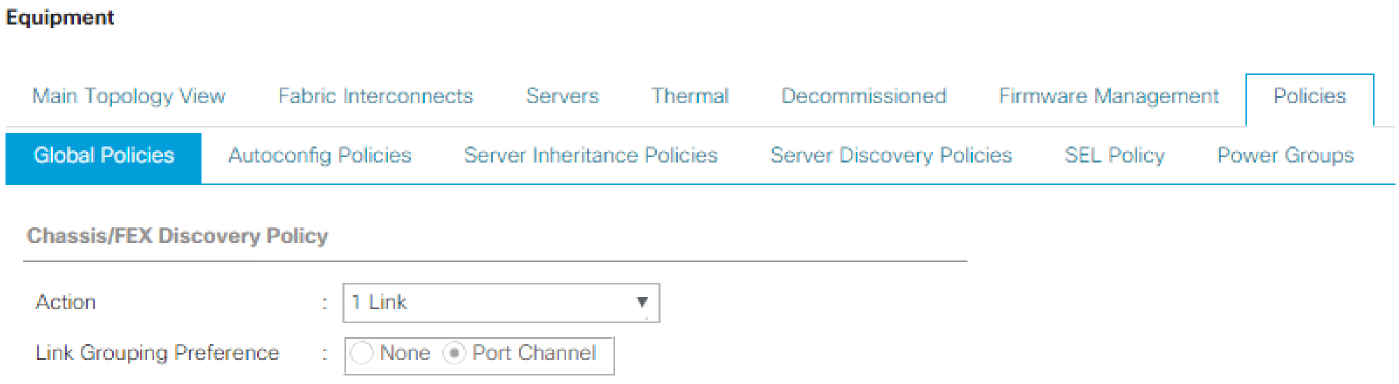

Chassis/FEX Discovery Policy

Action : 1 Link ▼

Link Grouping Preference : ○ None ● Port Channel

4. If any changes have been made, click Save Changes and then click OK.

Enable Server and Uplink Ports

To enable and verify server and uplink ports, follow these steps:

1. In Cisco UCS Manager, click Equipment.

2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.

3. Expand and select Ethernet Ports.

4. Verify that all ports connected to UCS chassis are configured as Server ports and have a status of Up.

5. Select the ports that are connected to Cisco FEXes and direct connect Cisco UCS C-Series servers, right-click them, and select Configure as Server Port.

6. Click Yes to confirm server ports and click OK.
7. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
8. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
9. Click Yes to confirm uplink ports and click OK.
10. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
11. Expand and select Ethernet Ports.
12. Verify that all ports connected to UCS chassis are configured as Server ports and have a status of Up.
13. Select the ports that are connected to Cisco FEXes and direct connect C-series servers, right-click them, and select Configure as Server Port.
14. Click Yes to confirm server ports and click OK.
15. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
16. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
17. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external FEX modules, follow these steps:

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

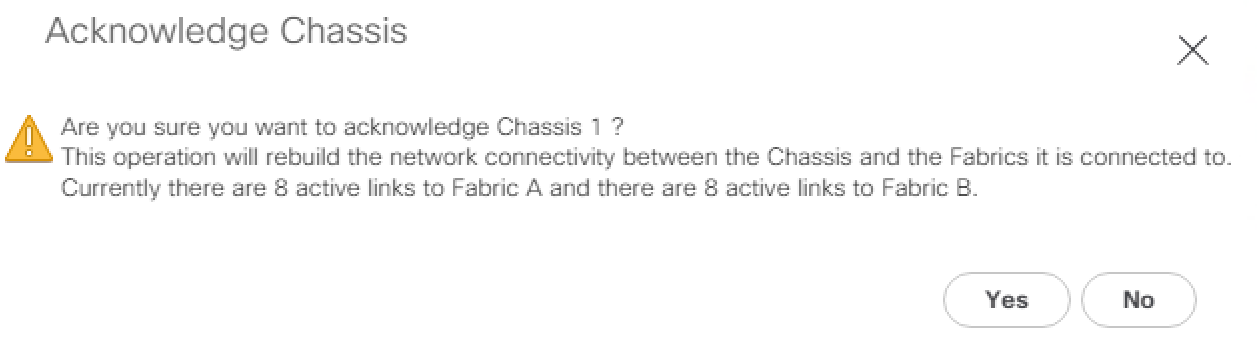

Acknowledge Chassis

Are you sure you want to acknowledge Chassis 1 ?
This operation will rebuild the network connectivity between the Chassis and the Fabrics it is connected to.
Currently there are 8 active links to Fabric A and there are 8 active links to Fabric B.

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

Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.

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

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels under Fabric A.
4. Select Create Port Channel.
5. Enter `153` as the unique ID of the port channel.
6. Enter `Po153-Nexus` as the name of the port channel.
7. Click Next.
8. Select the uplink ports connected to the Nexus switches to be added to the port channel.
9. Click >> to add the ports to the port channel.

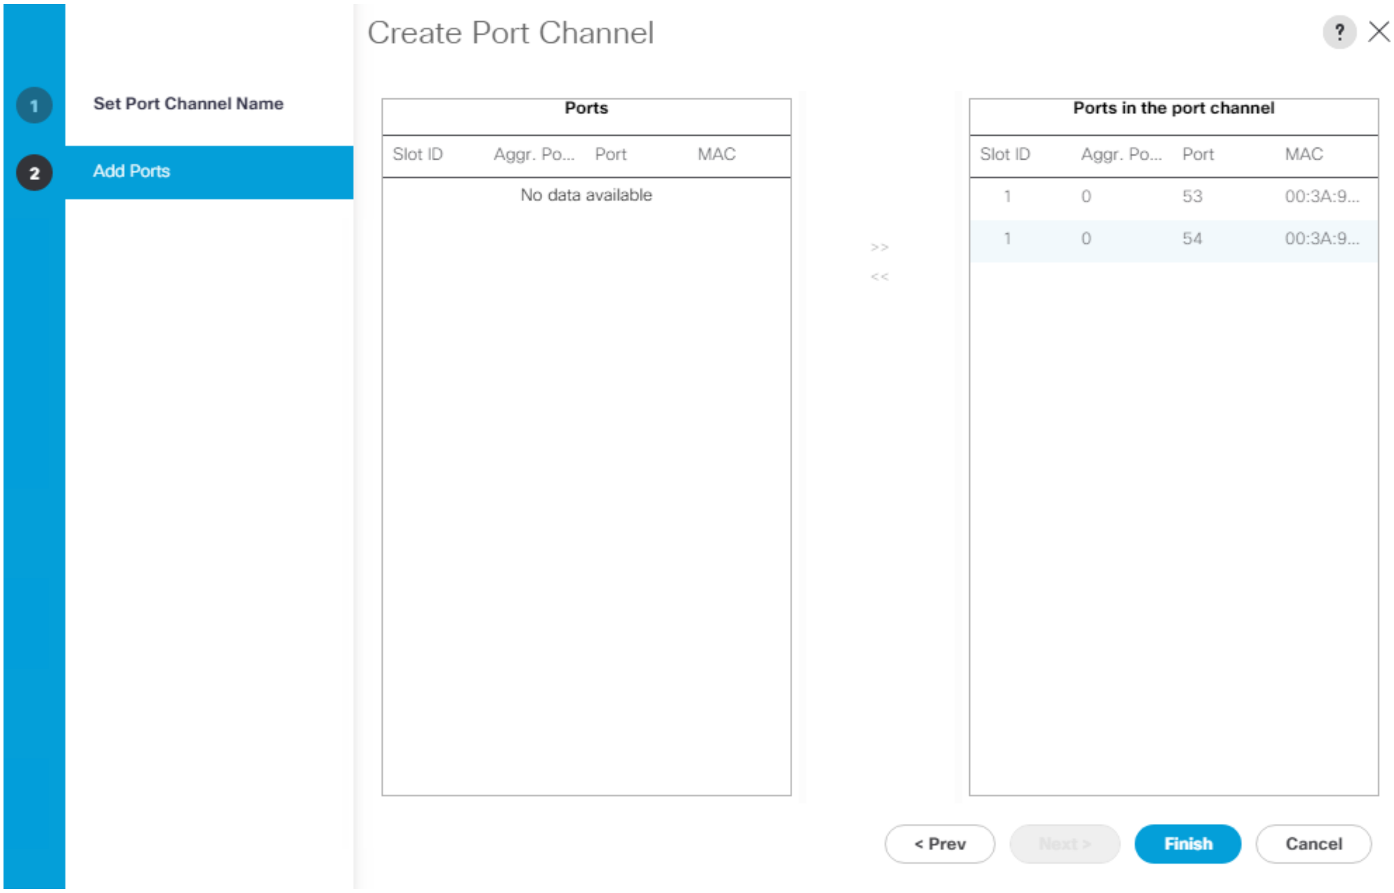

Create Port Channel

1 Set Port Channel Name

2 Add Ports

Ports

Slot ID	Aggr. Po...	Port	MAC
No data available			

Ports in the port channel

Slot ID	Aggr. Po...	Port	MAC
1	0	53	00:3A:9...
1	0	54	00:3A:9...

10. Click Finish to create the port channel.

11. Click OK.

12. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, select Port-Channel 153. Select 100 Gbps for the Admin Speed.

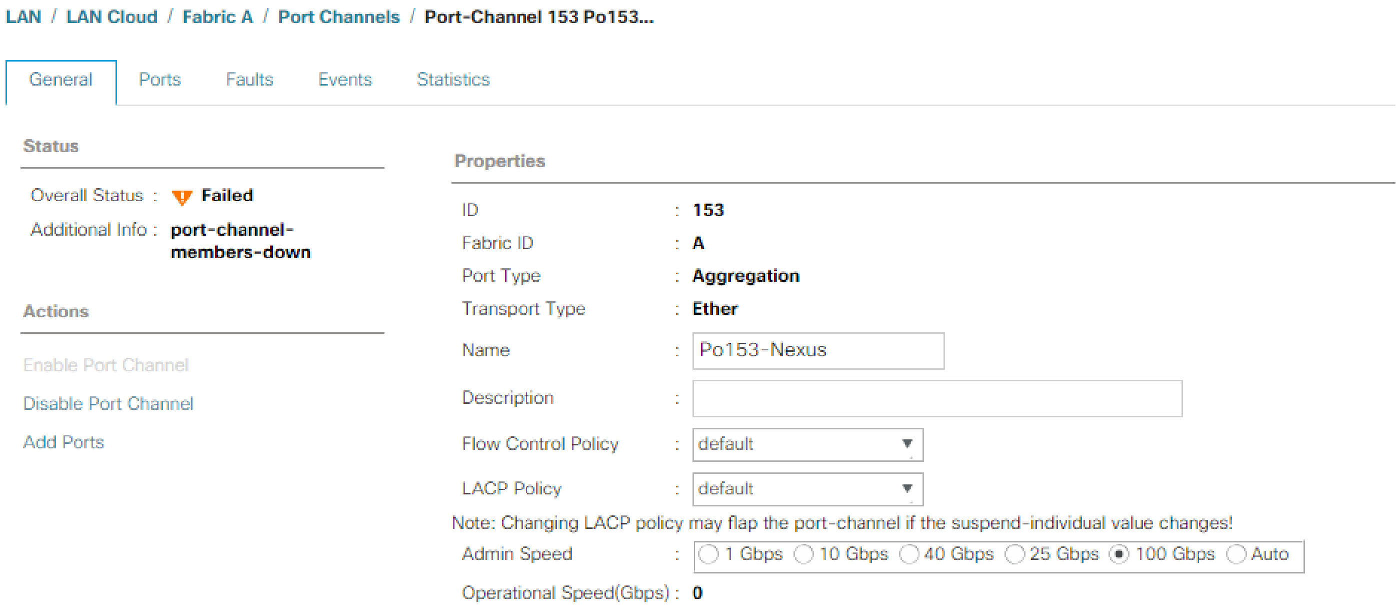

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 153 Po153...

Status

Overall Status : Failed

Additional Info : port-channel-members-down

Actions

Properties

ID : 153

Fabric ID : A

Port Type : Aggregation

Transport Type : Ether

Name : Po153-Nexus

Description :

Flow Control Policy : default

LACP Policy : default

Note: Changing LACP policy may flap the port-channel if the suspend-individual value changes!

Admin Speed : ○ 1 Gbps ○ 10 Gbps ○ 40 Gbps ○ 25 Gbps ● 100 Gbps ○ Auto

Operational Speed(Gbps) : 0

13. Click Save Changes and OK. After a few minutes, verify that the Overall Status is Up and the Operational Speed is correct.

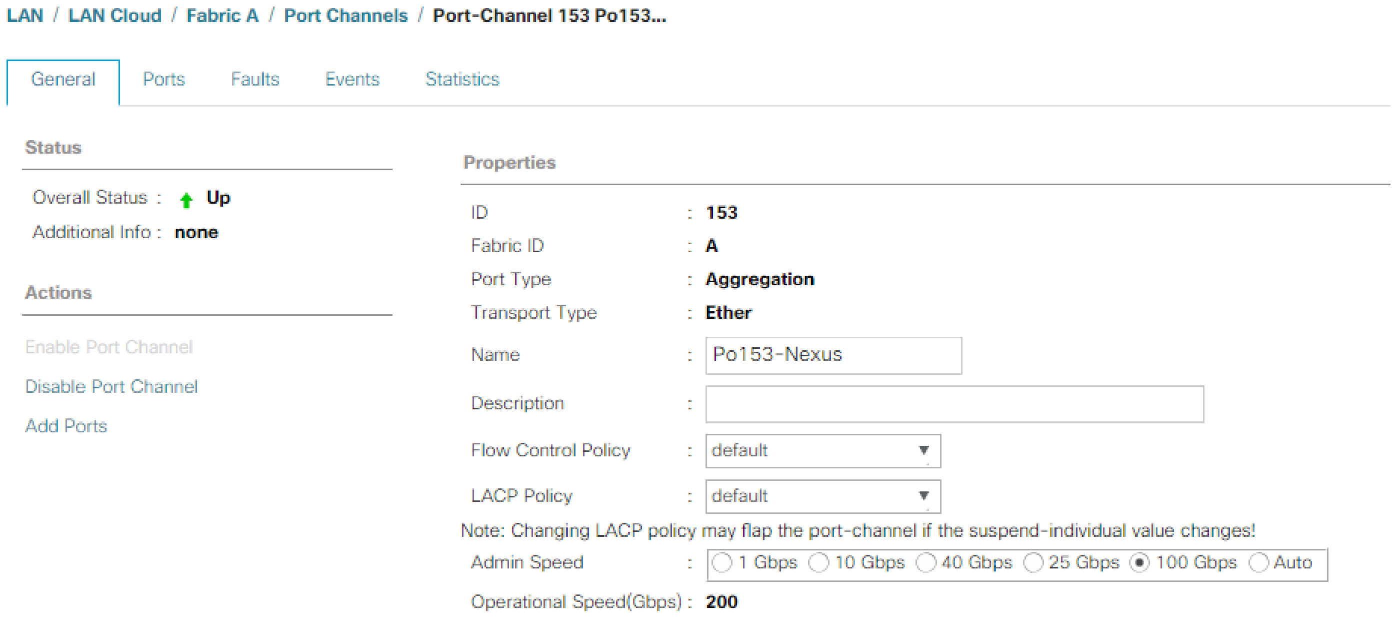

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 153 Po153...

Status

Overall Status : Up

Additional Info : none

Actions

Properties

ID : 153

Fabric ID : A

Port Type : Aggregation

Transport Type : Ether

Name : Po153-Nexus

Description :

Flow Control Policy : default

LACP Policy : default

Note: Changing LACP policy may flap the port-channel if the suspend-individual value changes!

Admin Speed : ○ 1 Gbps ○ 10 Gbps ○ 40 Gbps ○ 25 Gbps ● 100 Gbps ○ Auto

Operational Speed(Gbps) : 200

14. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.

15. Right-click Port Channels under Fabric B.

16. Select Create Port Channel.

17. Enter `154` as the unique ID of the port channel.

18. Enter `Po154-Nexus` as the name of the port channel.

19. Click Next.

20. Select the ports connected to the Nexus switches to be added to the port channel:

21. Click >> to add the ports to the port channel.

22. Click Finish to create the port channel.

23. Click OK.

24. In the navigation pane, under LAN > LAN Cloud > Fabric B > Port Channels, select Port-Channel 154. Select 100 Gbps for the Admin Speed.

25. Click Save Changes and OK. After a few minutes, verify that the Overall Status is Up and the Operational Speed is correct.

Add UDLD to Uplink Port Channels

To configure the unidirectional link detection (UDLD) on the Uplink Port Channels to the Nexus switches, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand Policies > LAN Cloud > UDLD Link Policy.

3. Right-click UDLD Link Policy and select Create UDLD Link Policy.

4. If the uplink cables to the Nexus switches are copper cables, name the Policy UDLD-Aggressive and select Enabled for the Admin State and Aggressive for the Mode. If the uplink cables to the Nexus switches are fibre optic cables, name the Policy UDLD-Normal and select Enabled for the Admin State and Normal for the Mode. In the validation lab configuration, UDLD-Aggressive was created.

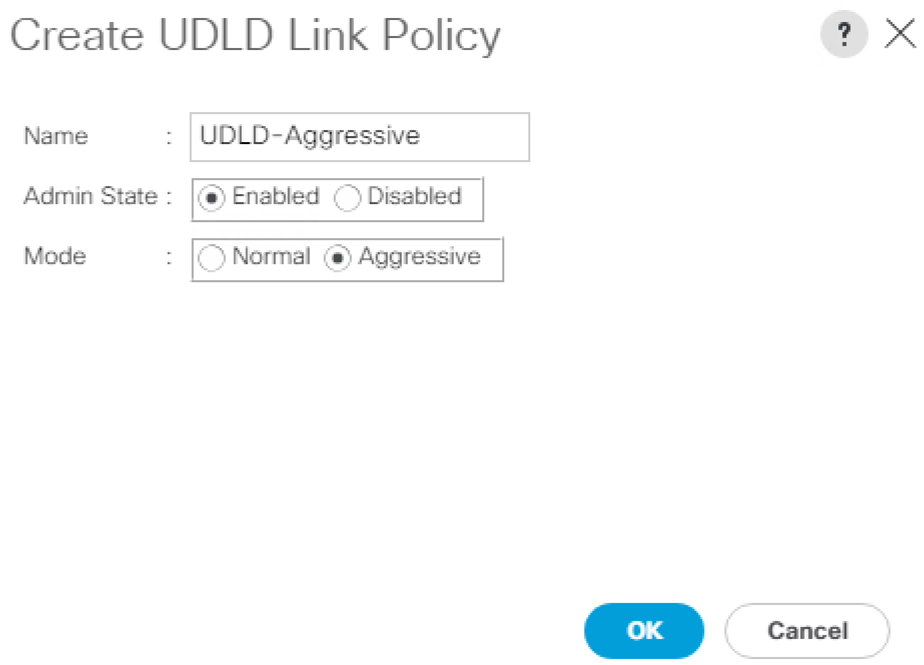

Create UDLD Link Policy

Name : UDLD-Aggressive

Admin State : ● Enabled ○ Disabled

Mode : ○ Normal ● Aggressive

5. Click OK, then OK again to complete creating the policy.

It is important that the Nexus switch port UDLD configurations (Aggressive or Normal) match the UCS port UDLD configurations.

6. Expand Policies > LAN Cloud > Link Profile.

7. Right-click Link Profile and select Create Link Profile.

8. Give the Link Profile the same name as the UDLD Link Policy above and select the UDLD Link Policy created above.

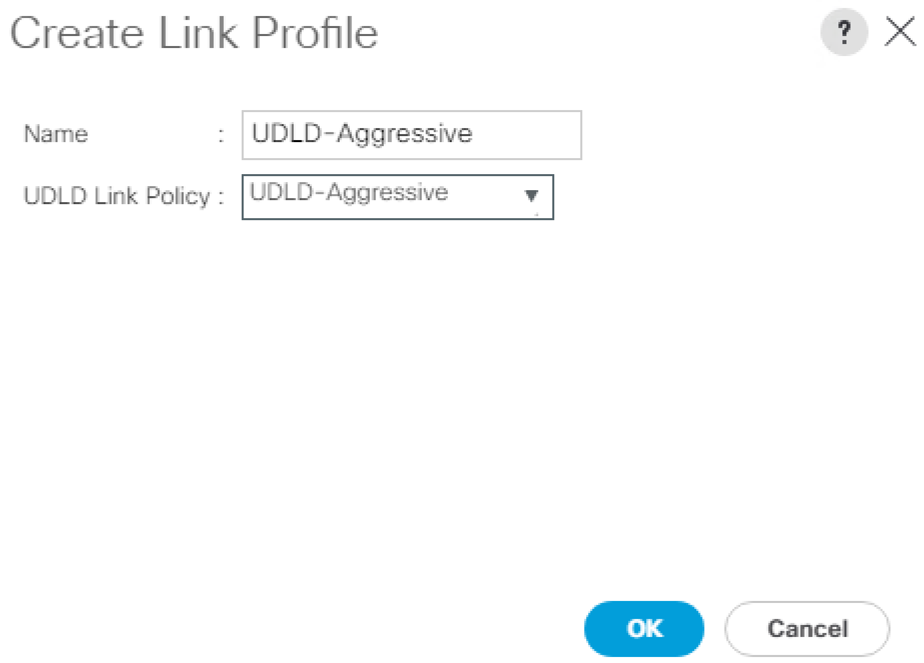

Create Link Profile

Name : UDLD-Aggressive

UDLD Link Policy : UDLD-Aggressive

9. Click OK, then click OK again to complete creating the profile.

10. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 153.

11. Select the first Eth Interface under Port-Channel 153. From the drop-down list, select the Link Profile created above, click Save Changes and OK. Repeat this process for each Eth Interface under Port-Channel 153 and for each Eth Interface under Port-Channel 154 on Fabric B.

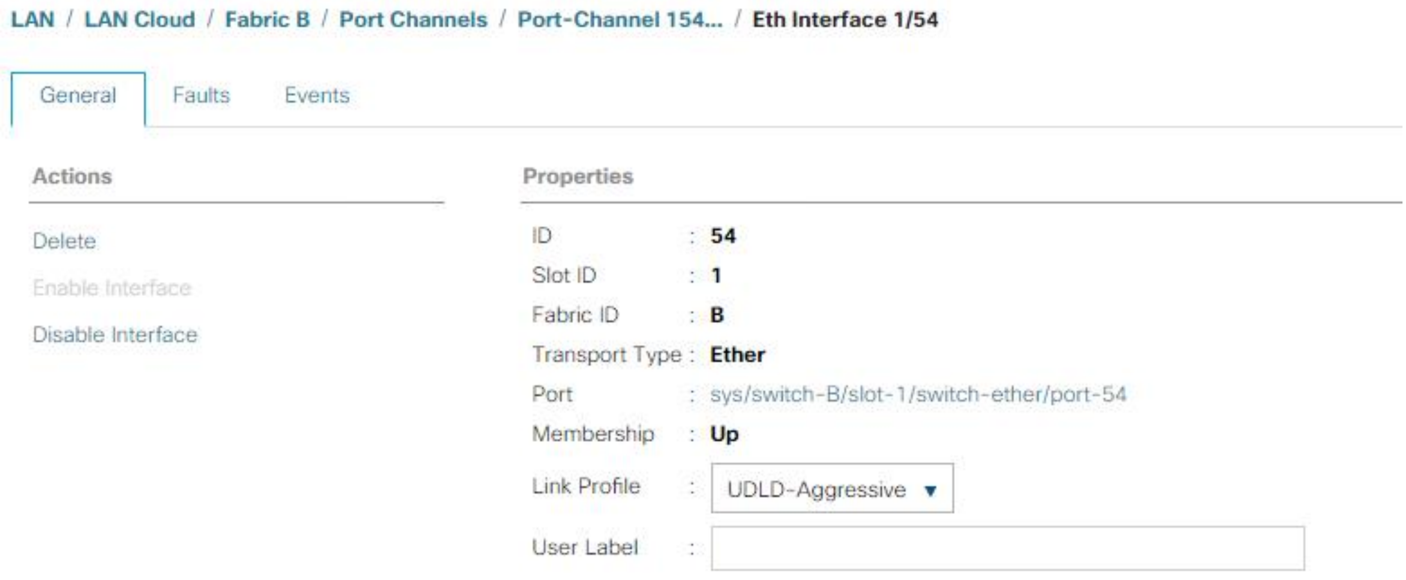

LAN / LAN Cloud / Fabric B / Port Channels / Port-Channel 154... / Eth Interface 1/54

ID : 54

Slot ID : 1

Fabric ID : B

Transport Type : Ether

Port : sys/switch-B/slot-1/switch-ether/port-54

Membership : Up

Link Profile : UDLD-Aggressive

User Label :

To see that UDLD is set up correctly, log into each Nexus switch and type `show udld neighbors`.

Create a WWNN Pool for FC Boot (FCP)

To configure the necessary WWNN pool for the Cisco UCS environment, follow these steps on Cisco UCS Manager.

1. Select SAN.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Select Create WWNN Pool to create the WWNN pool.
5. Enter WWNN-Pool for the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.
7. Select Sequential for Assignment Order.

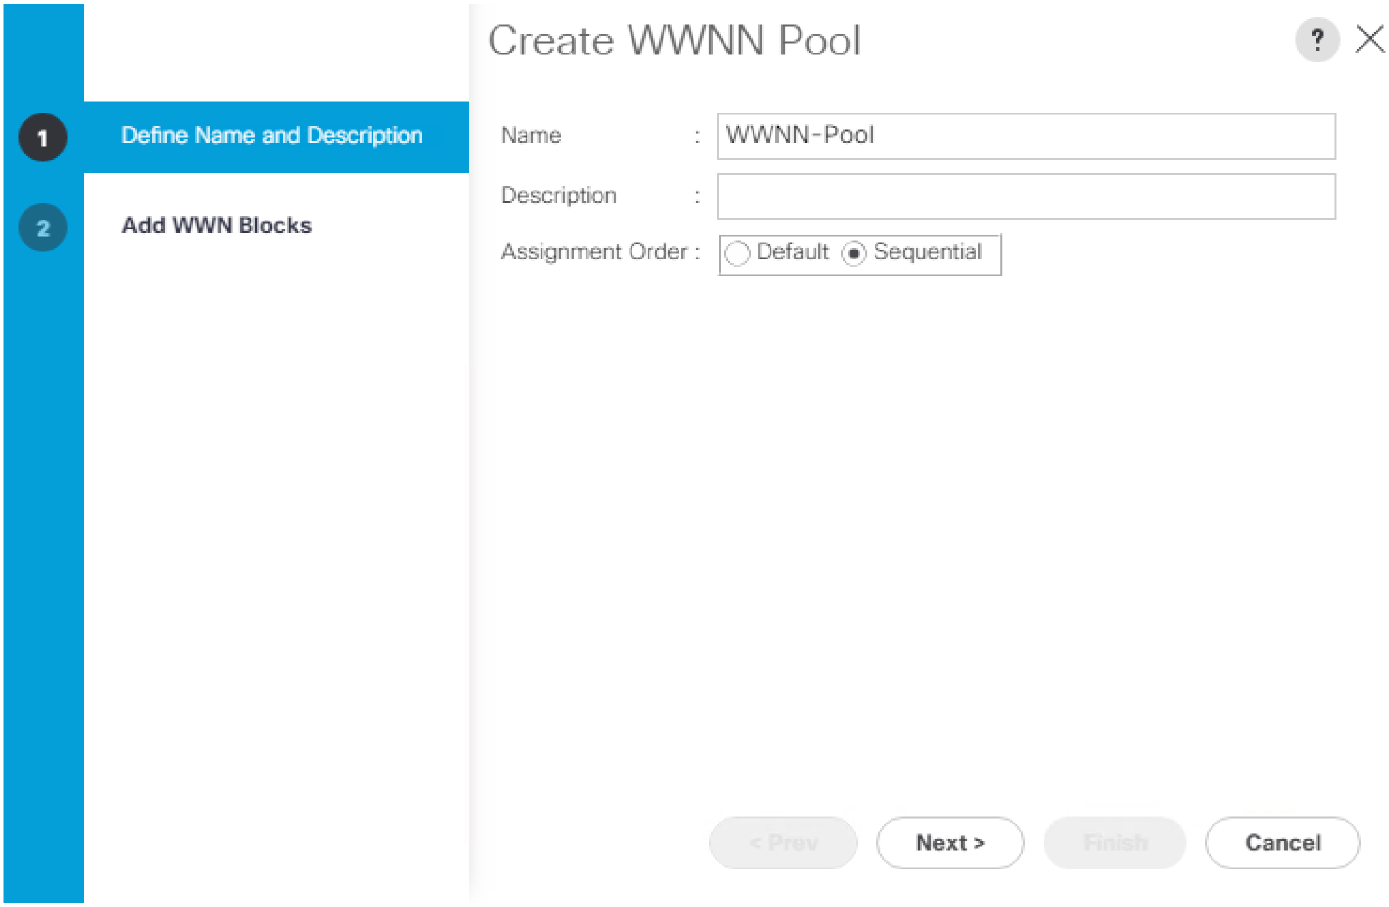

Create WWNN Pool

Name : WWNN-Pool

Description :

Assignment Order : ○ Default ● Sequential

8. Click Next.

9. Click Add.

10. Modify the From field as necessary for the UCS Environment

Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the Cisco UCS domain. Within the From field in our example, the sixth and seventh octets were changed from 00:00 to A1:30 to represent these WWNNs being in the A13 cabinet.

When there are multiple UCS domains sitting in adjacency, it is important that these blocks; the WWNN, WWPN, and MAC, hold differing values between each set.

11. Specify a size of the WWNN block sufficient to support the available server resources.

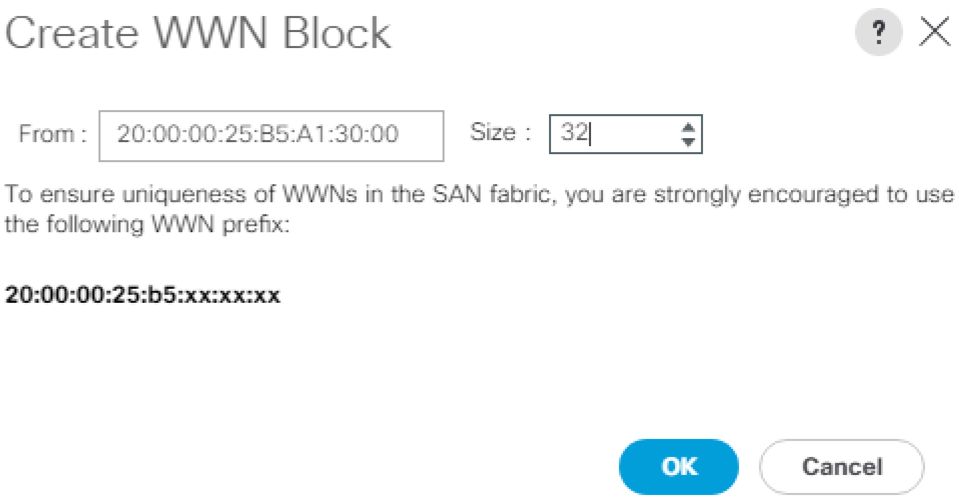

Create WWN Block

From : 20:00:00:25:B5:A1:30:00 Size : 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

12. Click OK.

13. Click Finish and click OK to complete creating the WWNN pool.

Create WWPN Pools (FCP)

To configure the necessary WWPN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.

2. Select Pools > root.

In this procedure, two WWPN pools are created, one for each switching fabric.

3. Right-click WWPN Pools under the root organization.

4. Select Create WWPN Pool to create the WWPN pool.

5. Enter WWPN-Pool-A as the name of the WWPN pool.

6. Optional: Enter a description for the WWPN pool.

7. Select Sequential for Assignment Order.

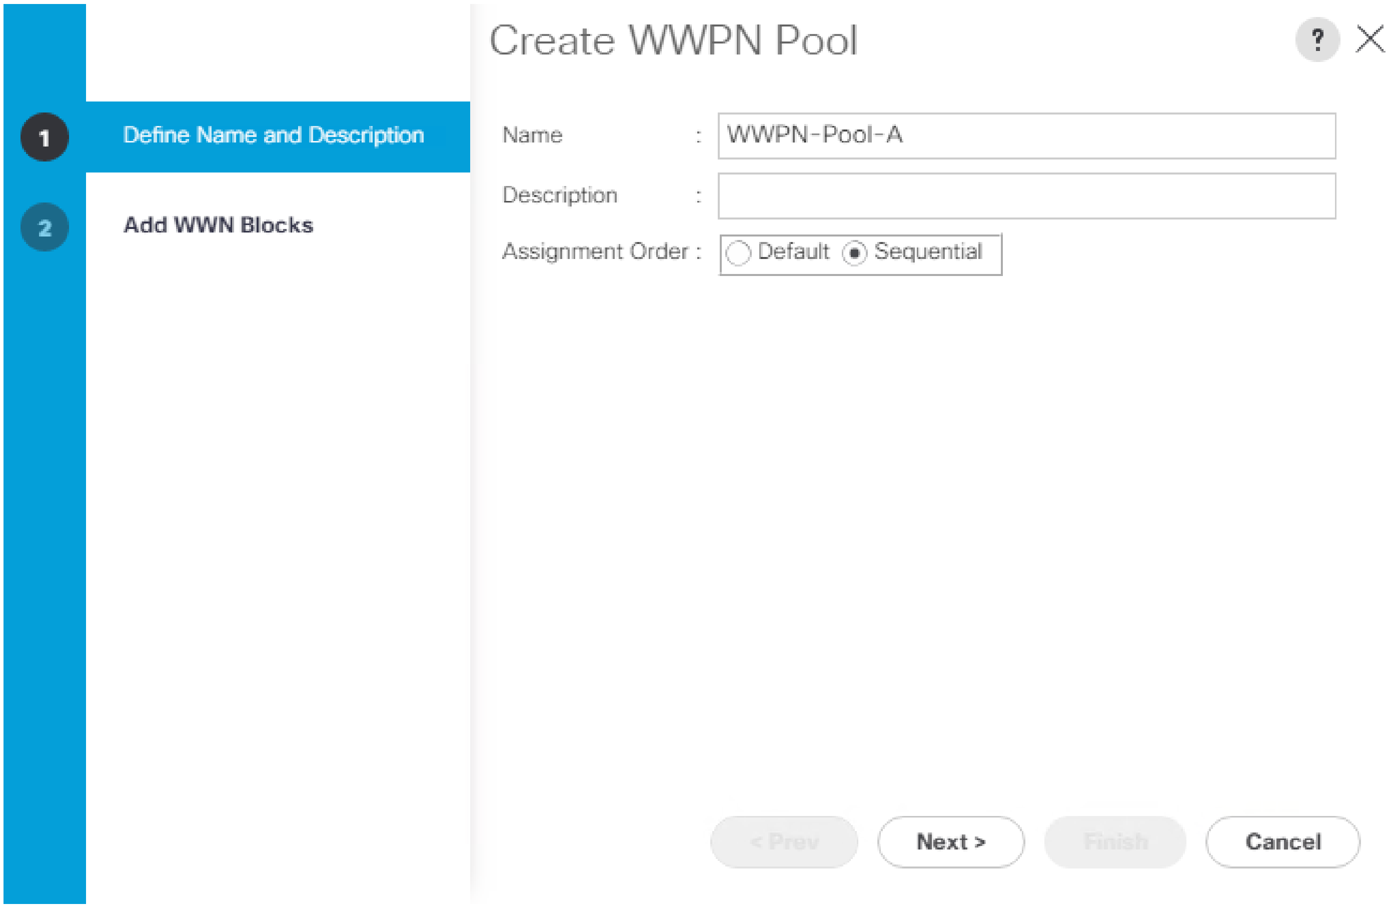

Create WWPN Pool

Name : WWPN-Pool-A

Description :

Assignment Order : ○ Default ● Sequential

8. Click Next.

9. Click Add.

10. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `A` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:A1:3A:00`

11. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources remembering that servers could have multiple vHBAs and unassociated service profiles could be created.

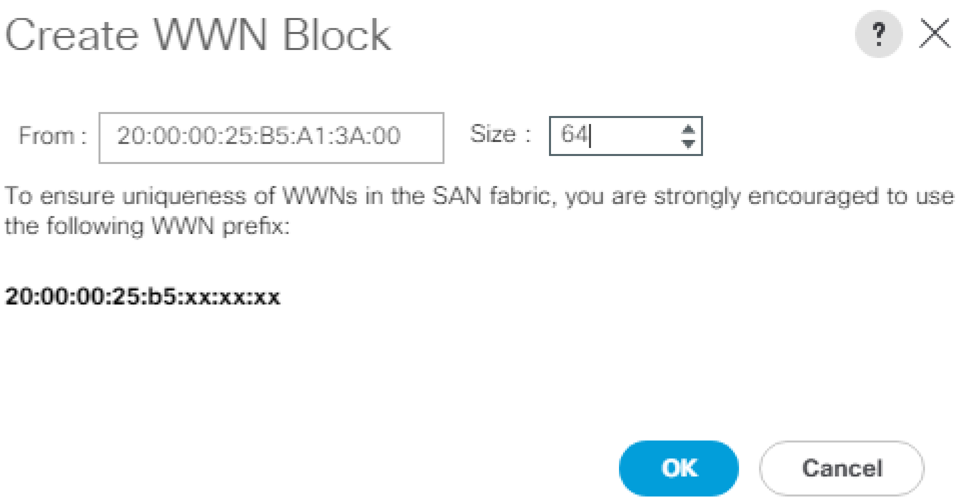

Create WWN Block

From : 20:00:00:25:B5:A1:3A:00 Size : 64

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

12. Click OK.

13. Click Finish.

14. In the confirmation message, click OK.

15. Right-click WWPN Pools under the root organization.

16. Select Create WWPN Pool to create the WWPN pool.

17. Enter WWPN-Pool-B as the name of the WWPN pool.

18. Optional: Enter a description for the WWPN pool.

19. Select Sequential for Assignment Order.

20. Click Next.

21. Click Add.

22. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `B` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric B addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:A1:3B:00`.

23. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources remembering that servers could have multiple vHBAs and unassociated service profiles could be created.

24. Click OK.

25. Click Finish.

26. In the confirmation message, click OK.

Create VSANs (FCP)

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the SAN.

In this procedure, two VSANs are created, one for each SAN switching fabric.

2. Select SAN > SAN Cloud.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter VSAN-A as the name of the VSAN to be used for Fabric A.
6. Leave FC Zoning set at Disabled.
7. Select Fabric A.
8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

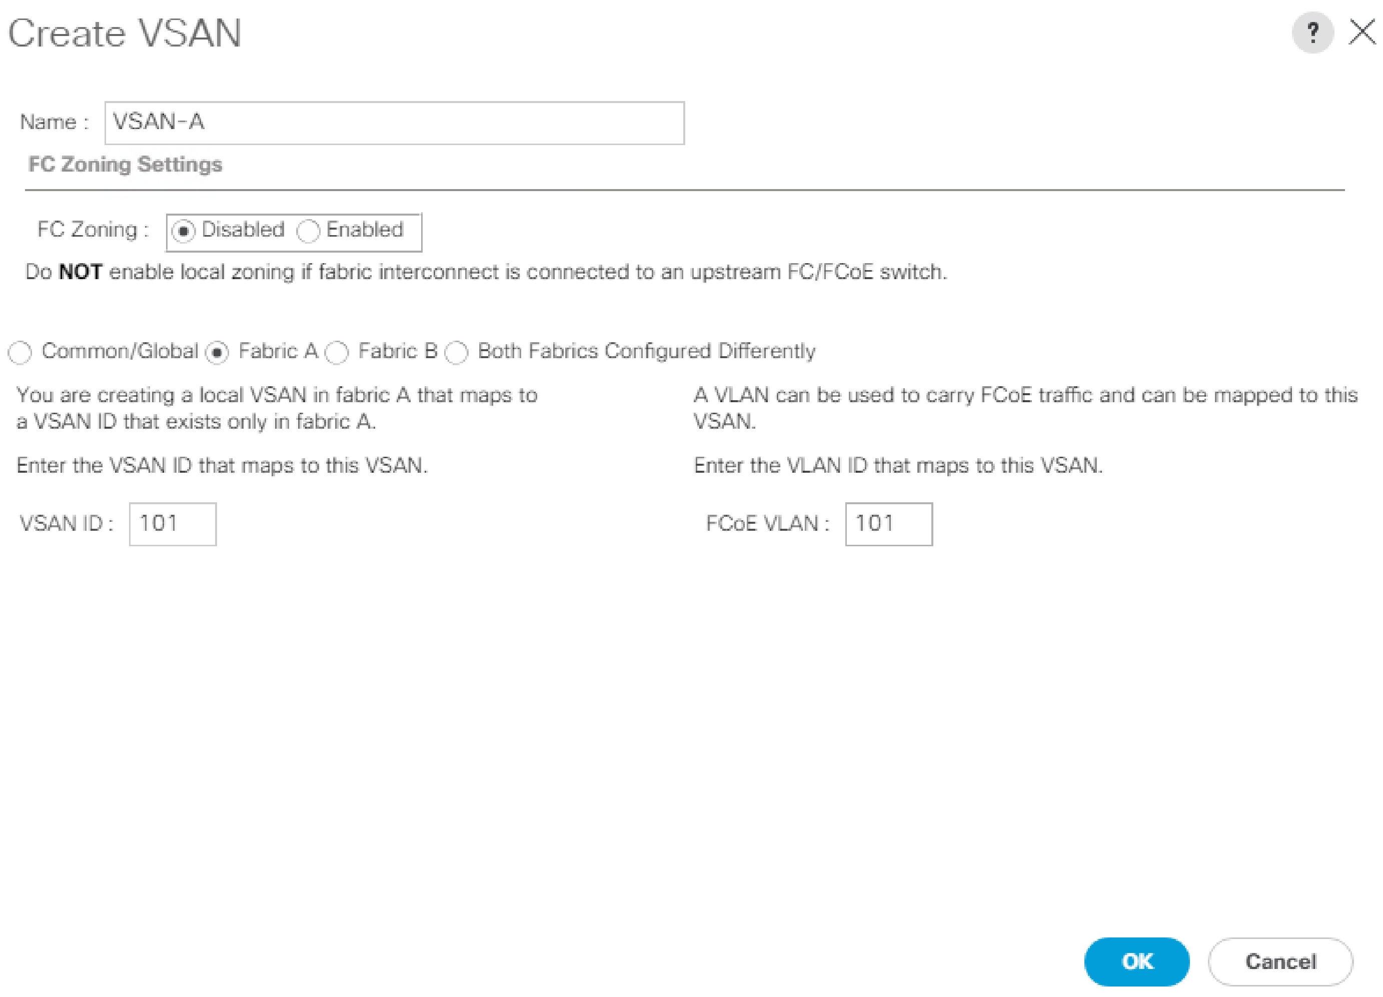

Create VSAN

Name : VSAN-A

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.

VSAN ID : 101

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.

FCoE VLAN : 101

9. Click OK and then click OK again.

10. Under SAN Cloud, right-click VSANs.

11. Select Create VSAN.

12. Enter VSAN-B as the name of the VSAN to be used for Fabric B.

13. Leave FC Zoning set at Disabled.

14. Select Fabric B.

15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric B. It is recommended use the same ID for both parameters and to use something other than 1.

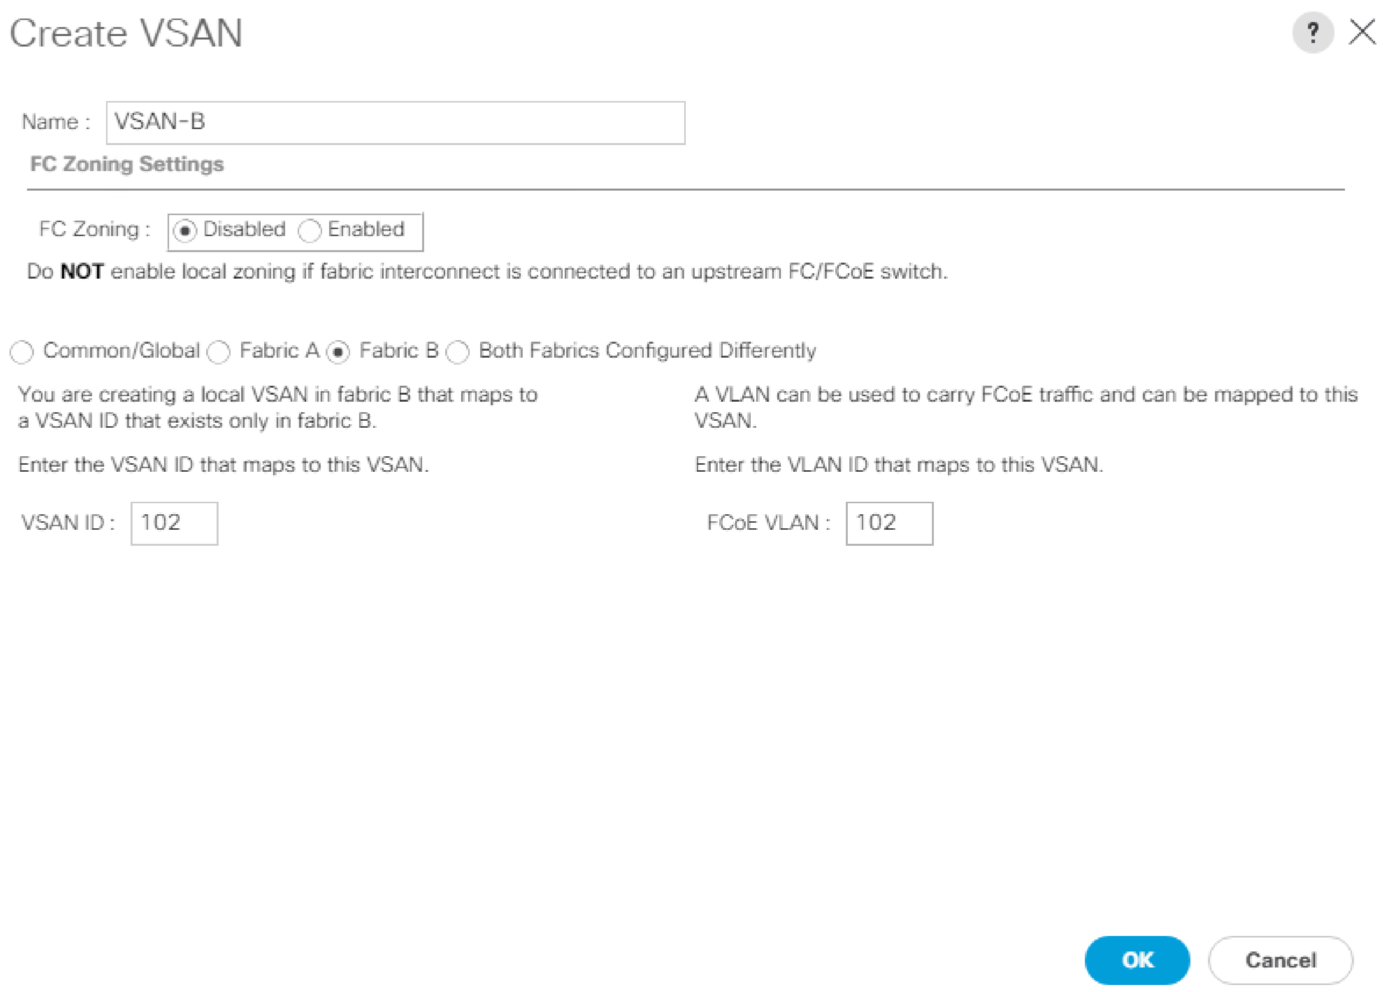

Create VSAN

Name : VSAN-B

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 B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 102

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.

FCoE VLAN : 102

16. Click OK and then click OK again.

Enable FC Uplink VSAN Trunking (FCP)

To enable VSAN trunking on the FC Uplinks in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Expand SAN > SAN Cloud.
3. Select Fabric A and in the Actions pane select Enable FC Uplink Trunking.
4. Click Yes on the Confirmation and Warning.
5. Click OK.
6. Select Fabric B and in the Actions pane select Enable FC Uplink Trunking.
7. Click Yes on the Confirmation and Warning.

8. Click OK.

Create FC Uplink Port Channels (FCP)

To create the FC Uplink Port Channels and assign the appropriate VSANs to them for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.

2. Select SAN > SAN Cloud.

3. Expand Fabric A and select FC Port Channels.

4. Right-click FC Port Channels and select Create FC Port Channel.

5. Set a unique ID for the port channel and provide a unique name for the port channel.

6. Click Next.

7. Select the appropriate Port Channel Admin Speed.

8. Select the ports connected to Cisco MDS A and use >> to add them to the port channel.

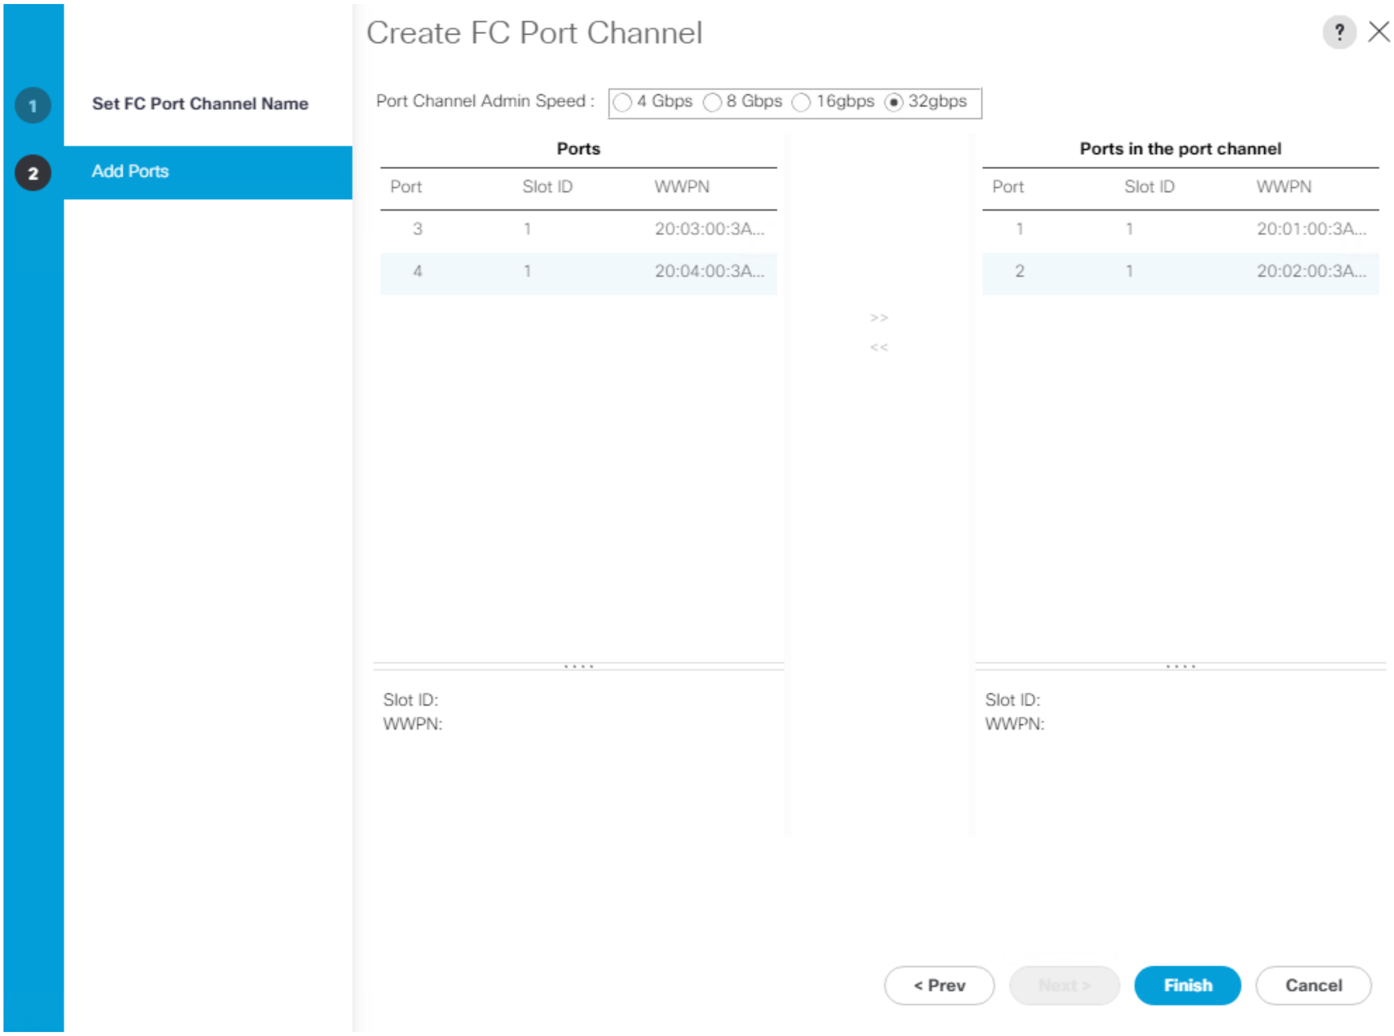

Create FC Port Channel

1 Set FC Port Channel Name

2 Add Ports

Port Channel Admin Speed : ○ 4 Gbps ○ 8 Gbps ○ 16gbps ● 32gbps

Ports

Port	Slot ID	WWPN
3	1	20:03:00:3A...
4	1	20:04:00:3A...

Slot ID:
WWPN:

Ports in the port channel

Port	Slot ID	WWPN
1	1	20:01:00:3A...
2	1	20:02:00:3A...

Slot ID:
WWPN:

9. Click Finish to complete creating the port channel.

10. Click OK the confirmation.

11. Under FC Port-Channels, select the newly created port channel.

12. In the right pane, view the drop-down list to select VSAN-A.

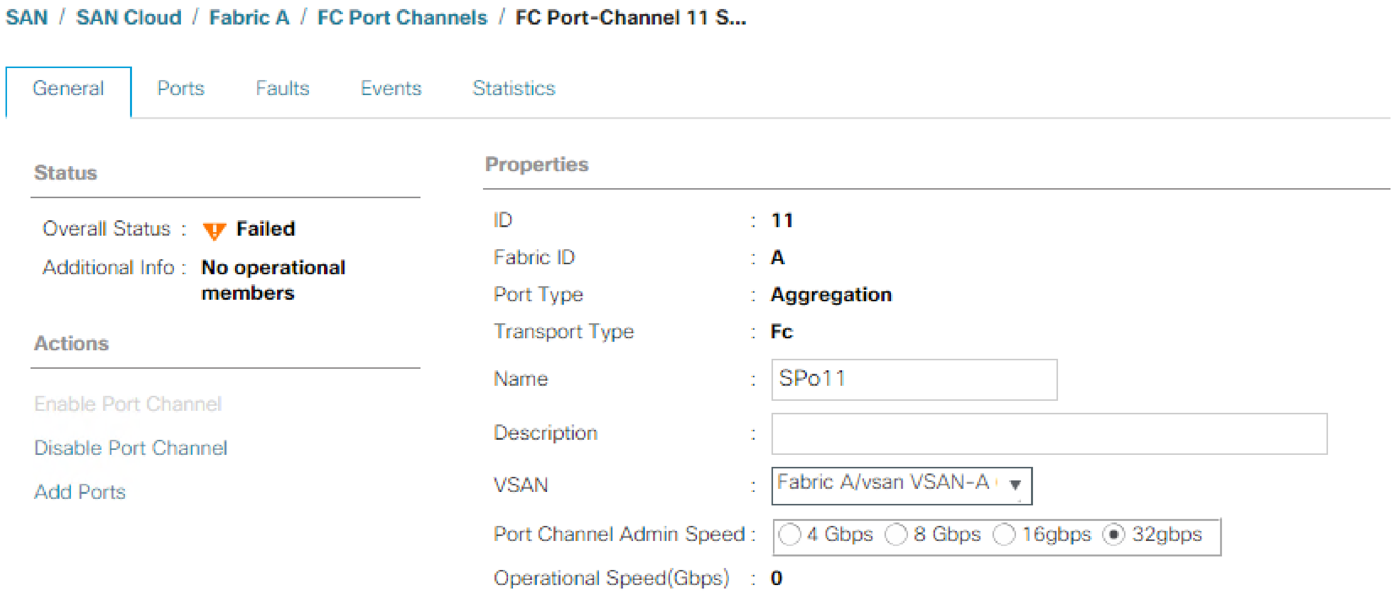

Status

Overall Status : Failed

Additional Info : No operational members

Properties

ID : 11

Fabric ID : A

Port Type : Aggregation

Transport Type : Fc

Name : SPo11

Description :

VSAN : Fabric A/vsan VSAN-A

Port Channel Admin Speed : ○ 4 Gbps ○ 8 Gbps ○ 16gbps ● 32gbps

Operational Speed(Gbps) : 0

13. Click Save Changes to assign the VSAN.

14. Click OK.

At this point in the deployment, since the MDS has not yet been configured, the SAN port-channel will not come up.

15. Expand Fabric B and select FC Port Channels.

16. Right-click FC Port Channels and select Create FC Port Channel.

17. Set a unique ID for the port channel and provide a unique name for the port channel.

18. Click Next.

19. Select the ports connected to Cisco MDS B and use >> to add them to the port channel.

20. Click Finish to complete creating the port channel.

21. Click OK on the confirmation.

22. Under FC Port-Channels, select the newly created port channel.

23. In the right pane, use the drop-down to select VSAN-B.

24. Click Save Changes to assign the VSAN.

25. Click OK.

Disable Unused FC Uplink Ports (FCP)

When Unified Ports were configured earlier in this procedure, on the Cisco UCS 6454 FI and the Cisco UCS 6332-16UP FI, FC ports were configured in groups. Because of this group configuration, some FC ports are unused and need to be disabled to prevent alerts. To disable the unused FC ports 3 and 4 on the Cisco UCS 6454 FIs, follow these steps:

1. In Cisco UCS Manager, click Equipment.
2. In the Navigation Pane, expand Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > FC Ports.
3. Select FC Port 3 and FC Port 4. Right-click and select Disable.
4. Click Yes and OK to complete disabling the unused FC ports.
5. In the Navigation Pane, expand Equipment > Fabric Interconnects > Fabric Interconnect B > Fixed Module > FC Ports.
6. Select FC Port 3 and FC Port 4. Right-click and select Disable.
7. Click Yes and OK to complete disabling the unused FC ports.

Create an Organization

To this point in the UCS deployment, all items have been deployed at the root level in Cisco UCS Manager. To allow this UCS to be shared among different projects, UCS Organizations can be created. To create an organization for this FlexPod deployment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. In the Navigation Pane, expand Servers > Service Profiles.
3. Right-click root under Service Profiles and select Create Organization.
4. Provide a name for the Organization to indicate this FlexPod deployment and optionally provide a Description.

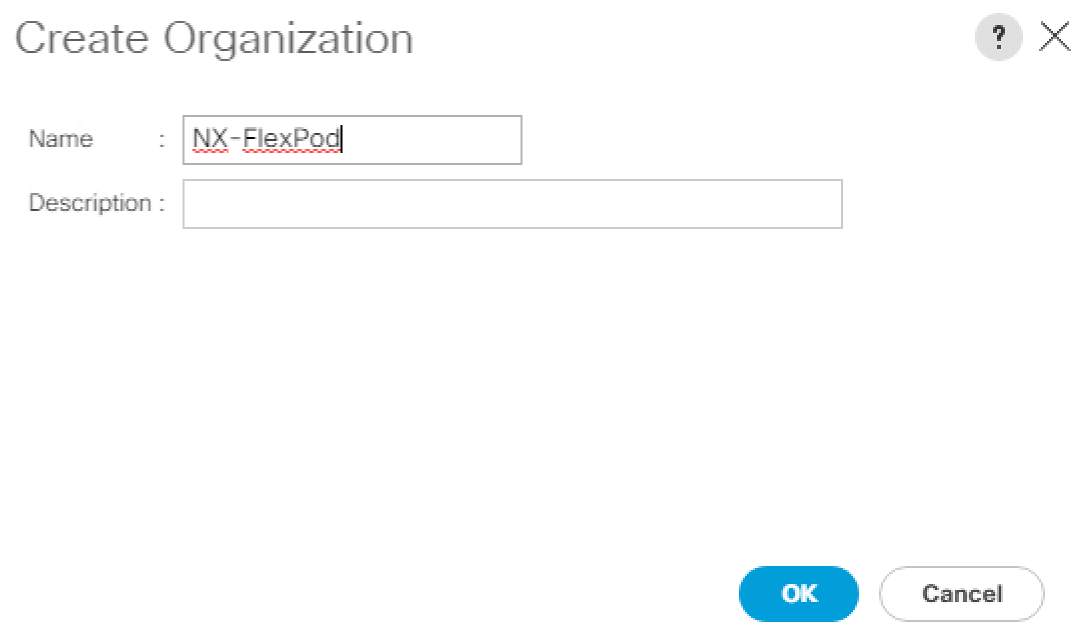

Create Organization

Name : NX-FlexPod

Description :

5. Click OK, then click OK again to complete creating the organization.

Create vHBA Templates (FCP)

To create the necessary virtual host bus adapter (vHBA) templates for the Cisco UCS environment within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click vHBA Templates under the FlexPod Organization.
4. Select Create vHBA Template.
5. Enter FCP-vHBA-A as the vHBA template name.
6. Keep Fabric A selected.
7. Leave Redundancy Type set to No Redundancy.
8. Select VSAN-A.
9. Leave Initial Template as the Template Type.
10. Select WWPN-Pool-A as the WWPN Pool.

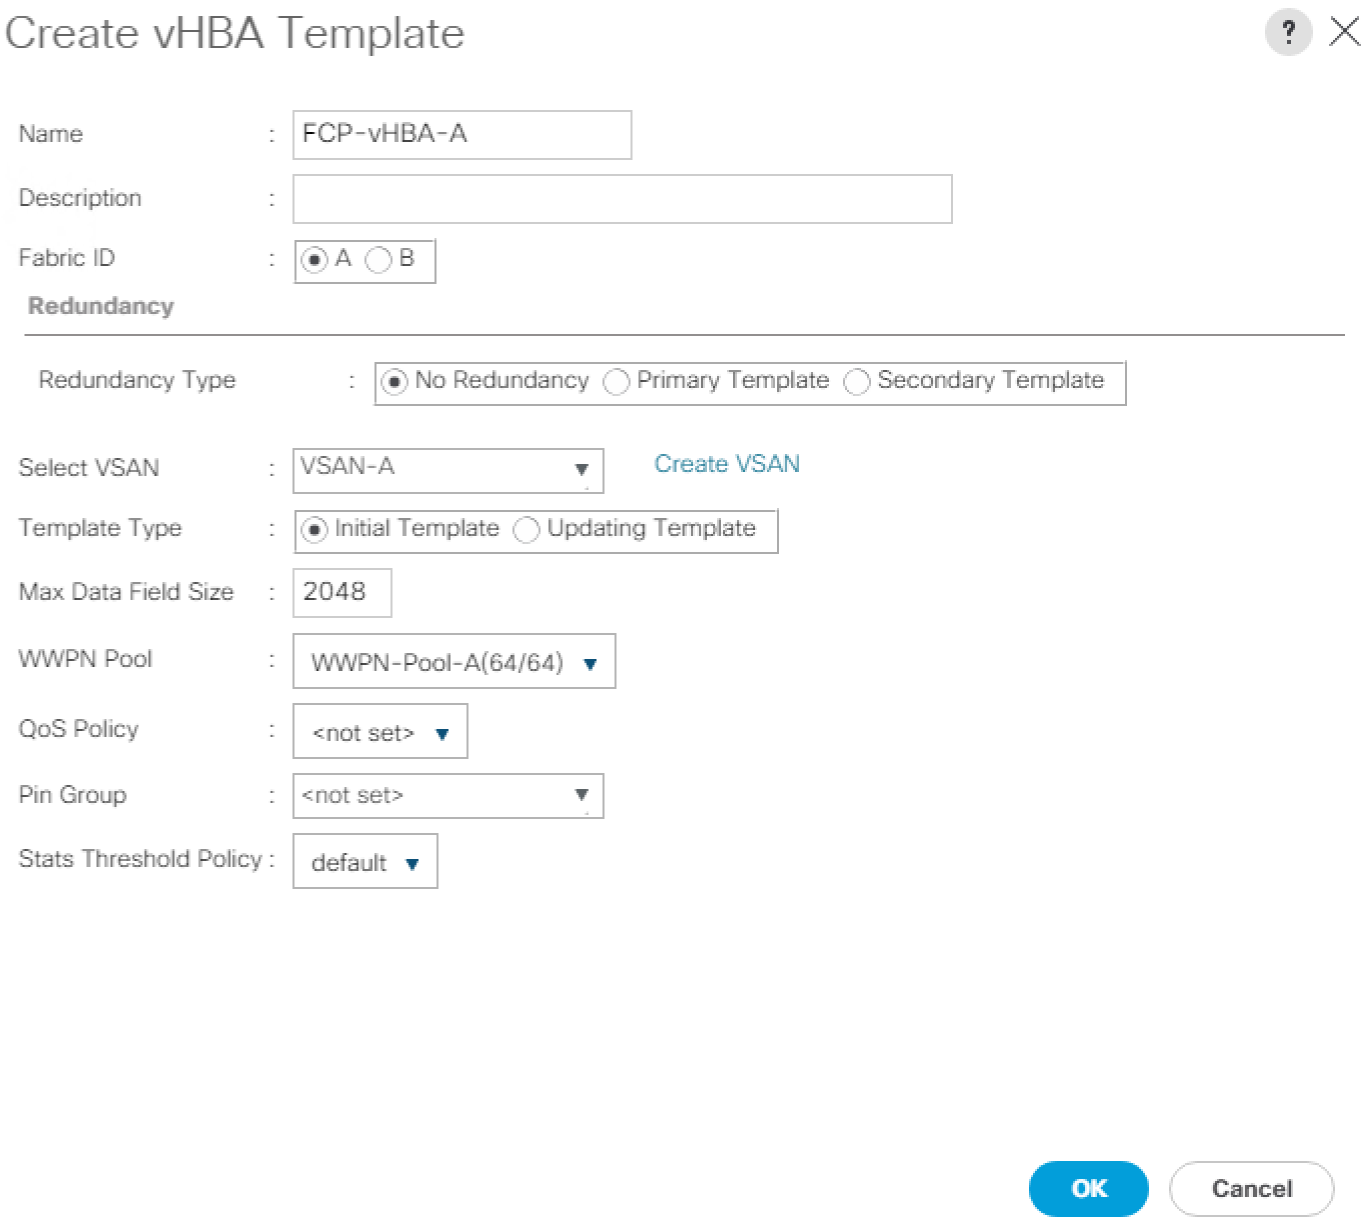

Create vHBA Template

Name : FCP-vHBA-A

Description :

Fabric ID : ● A ○ B

Redundancy

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Select VSAN : VSAN-A

Template Type : ● Initial Template ○ Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-Pool-A(64/64)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

11. Click OK to create the vHBA template.

12. Click OK.

13. Right-click vHBA Templates under the FlexPod Organization.

14. Select Create vHBA Template.

15. Enter FCP-vHBA-B as the vHBA template name.

16. Select Fabric B as the Fabric ID.

17. Leave Redundancy Type set to No Redundancy.

18. Select VSAN-B.

19. Leave Initial Template as the Template Type.

20. Select WWPN-Pool-B as the WWPN Pool.

21. Click OK to create the vHBA template.

22. Click OK.

Create SAN Connectivity Policy (FCP)

To configure the necessary Infrastructure SAN Connectivity Policy within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click SAN.

2. Select SAN > Policies > root > Sub-Organizations > FlexPod Organization.

3. Right-click SAN Connectivity Policies under the FlexPod Organization.

4. Select Create SAN Connectivity Policy.

5. Enter FC-Boot as the name of the policy.

6. Select the previously created WWNN-Pool for the WWNN Assignment.

7. Click the Add button at the bottom to add a vHBA.

8. In the Create vHBA dialog box, enter FCP-Fabric-A as the name of the vHBA.

9. Select the Use vHBA Template checkbox.

10. In the vHBA Template list, select FCP-vHBA-A.

11. In the Adapter Policy list, select VMWare.

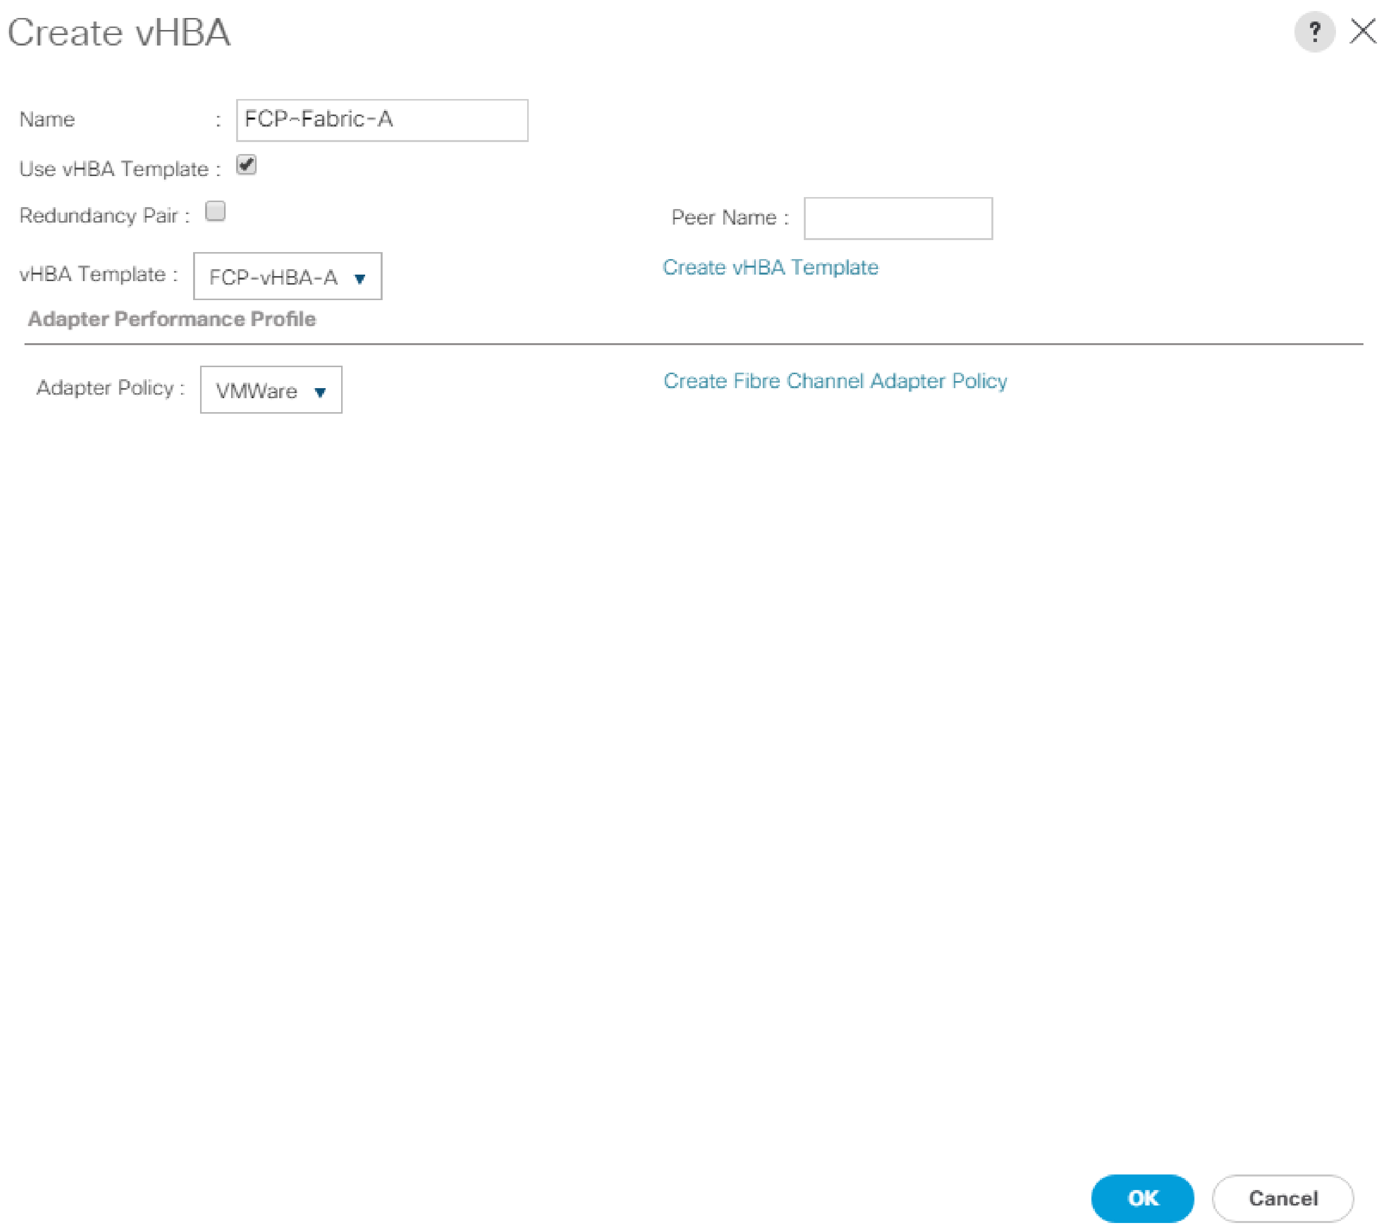

Create vHBA

Name : FCP-Fabric-A

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : FCP-vHBA-A

Adapter Performance Profile

Adapter Policy : VMWare

12. Click OK.

13. Click the Add button at the bottom to add a second vHBA.

14. In the Create vHBA dialog box, enter FCP-Fabric-B as the name of the vHBA.

15. Select the Use vHBA Template checkbox.

16. In the vHBA Template list, select FCP-vHBA-B.

17. In the Adapter Policy list, select VMWare.

18. Click OK.

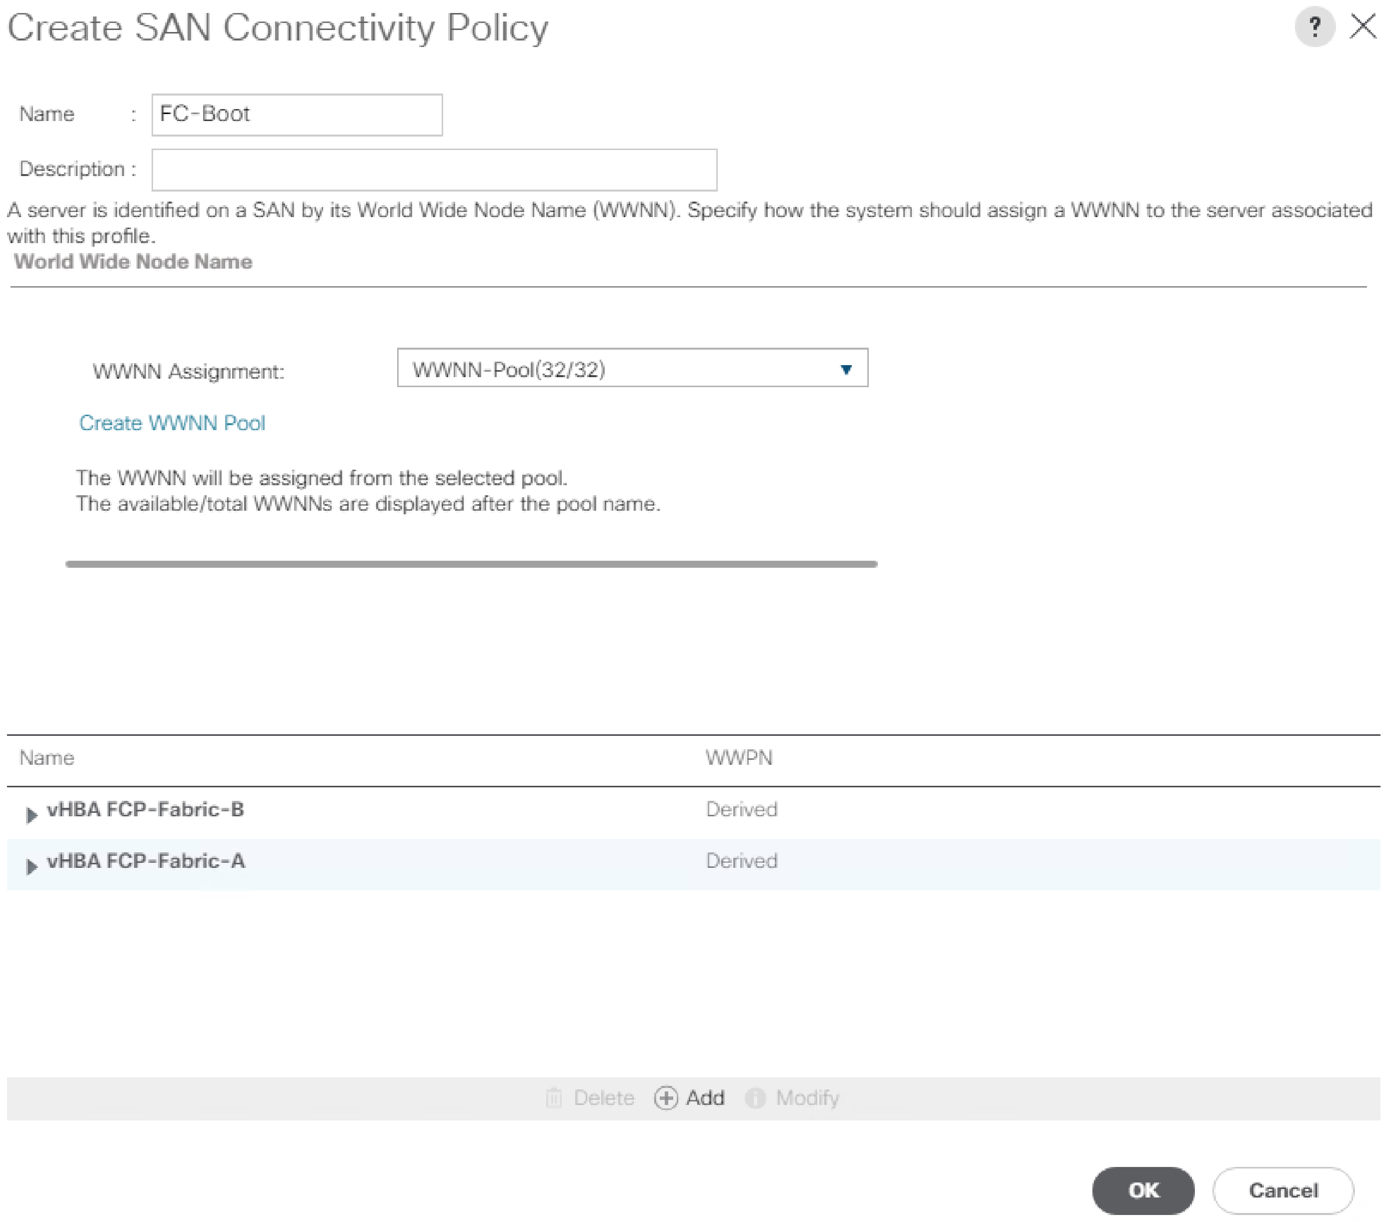

Create SAN Connectivity Policy

Name : FC-Boot

Description :

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.

World Wide Node Name

WWNN Assignment: WWNN-Pool(32/32)

The WWNN will be assigned from the selected pool.
The available/total WWNNs are displayed after the pool name.

Name	WWPN
▸ **vHBA FCP-Fabric-B**	Derived
▸ **vHBA FCP-Fabric-A**	Derived

19. Click OK to create the SAN Connectivity Policy.

20. Click OK to confirm creation.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Select Pools > root.

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 of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select Sequential as the option for Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting MAC address.

For the FlexPod solution, the recommendation is to place A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the example of also embedding the cabinet number information giving us `00:25:B5:A1:3A:00` as our first MAC address.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created.

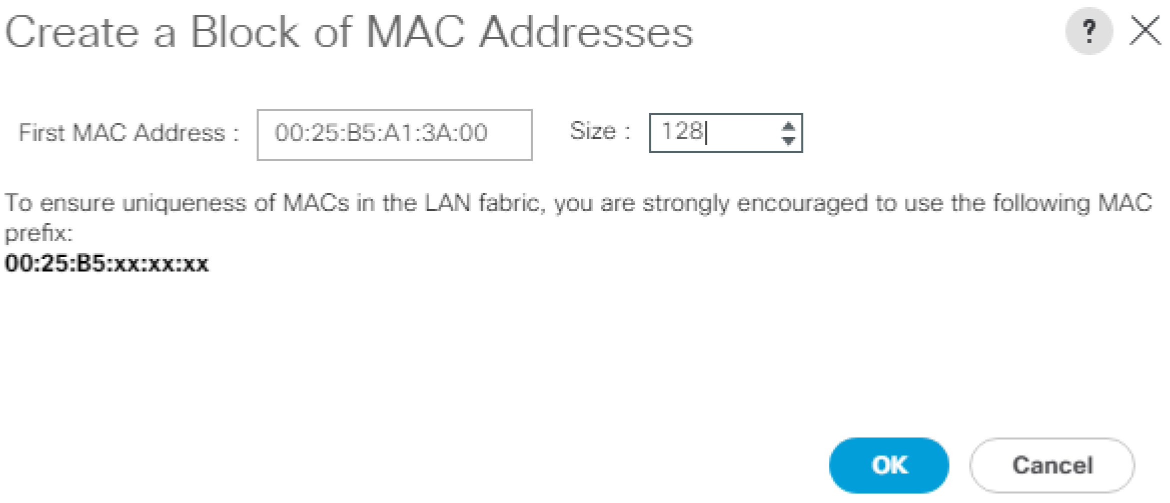

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:A1:3A:00

Size : 128

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

12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.

15. Right-click MAC Pools under the root organization.
16. Select Create MAC Pool to create the MAC address pool.
17. Enter MAC-Pool-B as the name of the MAC pool.
18. Optional: Enter a description for the MAC pool.
19. Select Sequential as the option for Assignment Order.
20. Click Next.
21. Click Add.
22. Specify a starting MAC address.

For the FlexPod solution, it is recommended to place B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward our example of also embedding the cabinet number information giving us `00:25:B5:A1:3B:00` as our first MAC address.

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created.
24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click Servers.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID-Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.

9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at 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 in the FlexPod Organization, follow these steps:

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

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

Create VLANs

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

1. In Cisco UCS Manager, click LAN.

In this procedure, five unique VLANs are created. See Table 2

2. Expand LAN > LAN Cloud.

3. Right-click VLANs.

4. Select Create VLANs.

5. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.

6. Keep the Common/Global option selected for the scope of the VLAN.

7. Enter the native VLAN ID.

8. Keep the Sharing Type as None.

9. Click OK and then click OK again.

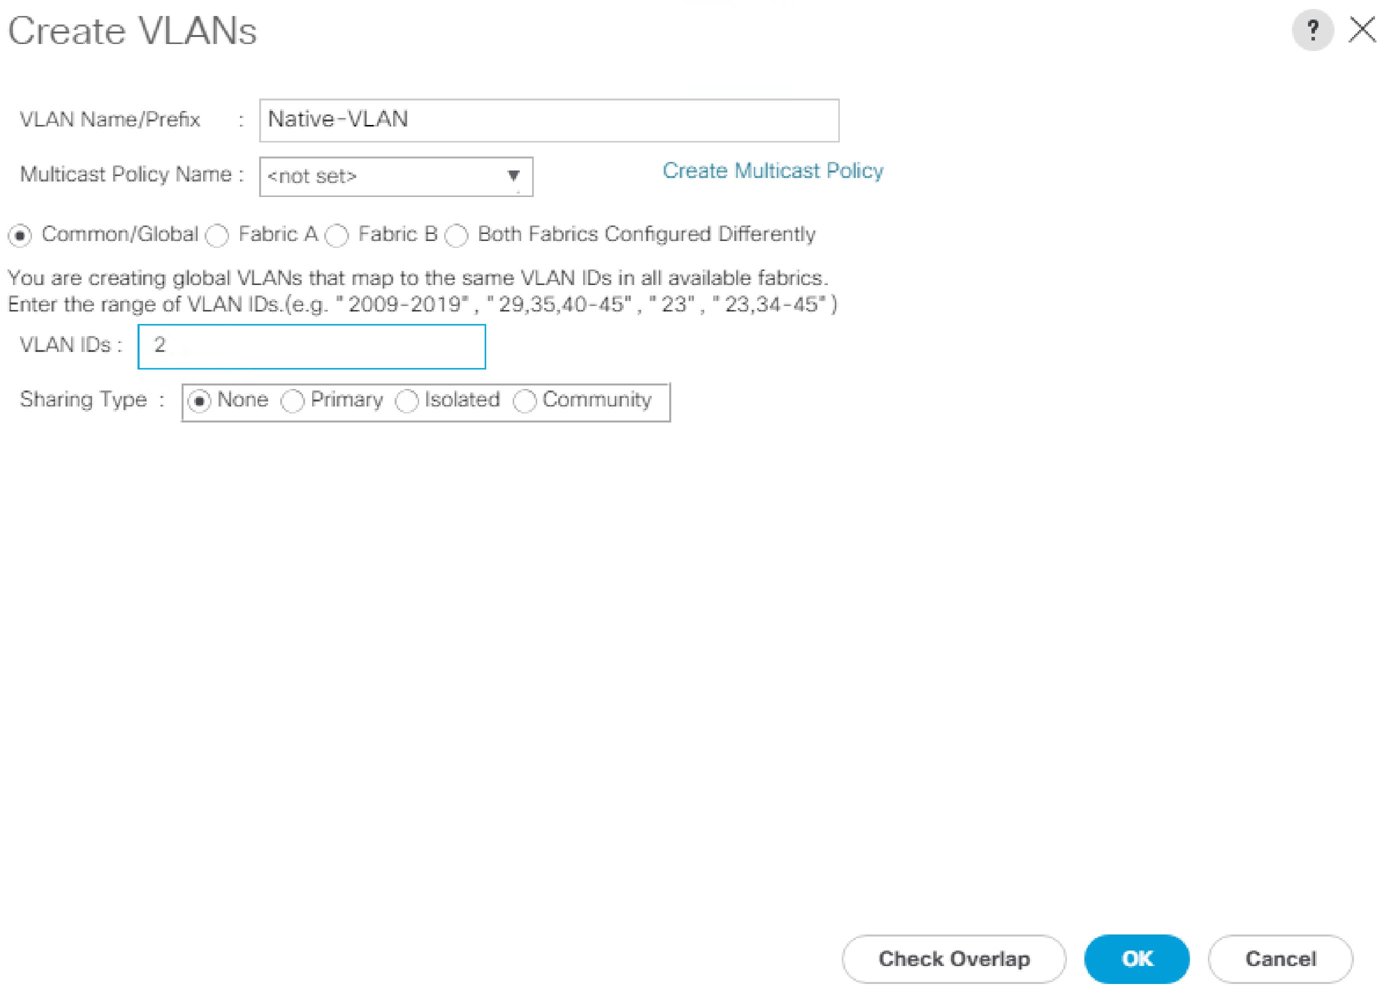

Create VLANs

VLAN Name/Prefix : Native-VLAN

Multicast Policy Name : <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. " 2009-2019" , " 29,35,40-45" , " 23" , " 23,34-45")

VLAN IDs : 2

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

10. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and select Set as Native VLAN.

11. Click Yes and then click OK.

12. Right-click VLANs.
13. Select Create VLANs
14. Enter IB-MGMT as the name of the VLAN to be used for management traffic.

Modify these VLAN names as necessary for your environment.

15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the In-Band management VLAN ID.
17. Keep the Sharing Type as None.
18. Click OK, and then click OK again.
19. Right-click VLANs.
20. Select Create VLANs.
21. Enter Infra-NFS as the name of the VLAN to be used for NFS.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the Infrastructure NFS VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK, and then click OK again.
26. Right-click VLANs.
27. Select Create VLANs.
28. Enter vMotion as the name of the VLAN to be used for vMotion.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the vMotion VLAN ID.
31. Keep the Sharing Type as None.
32. Click OK and then click OK again.
33. Select Create VLANs.
34. Enter VM-Traffic as the name of the VLAN to be used for VM Traffic.
35. Keep the Common/Global option selected for the scope of the VLAN.
36. Enter the VM-Traffic VLAN ID.

37. Keep the Sharing Type as None.

38. Click OK and then click OK again.

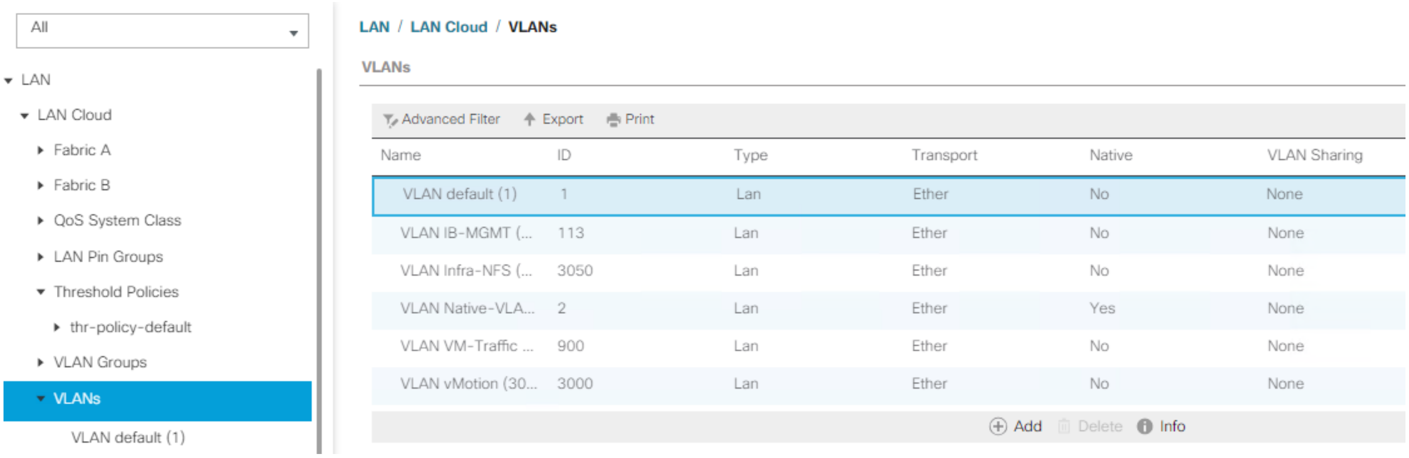

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	No	None
VLAN IB-MGMT (...	113	Lan	Ether	No	None
VLAN Infra-NFS (...	3050	Lan	Ether	No	None
VLAN Native-VLA...	2	Lan	Ether	Yes	None
VLAN VM-Traffic ...	900	Lan	Ether	No	None
VLAN vMotion (30...	3000	Lan	Ether	No	None

Modify Default Host Firmware Package

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

To modify the default firmware management policy in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers .

2. Expand Policies > root.

3. Expand Host Firmware Packages.

4. Select default.

5. In the Actions pane, select Modify Package Versions.

6. Select version 4.0(2d) for both the Blade and Rack Packages.

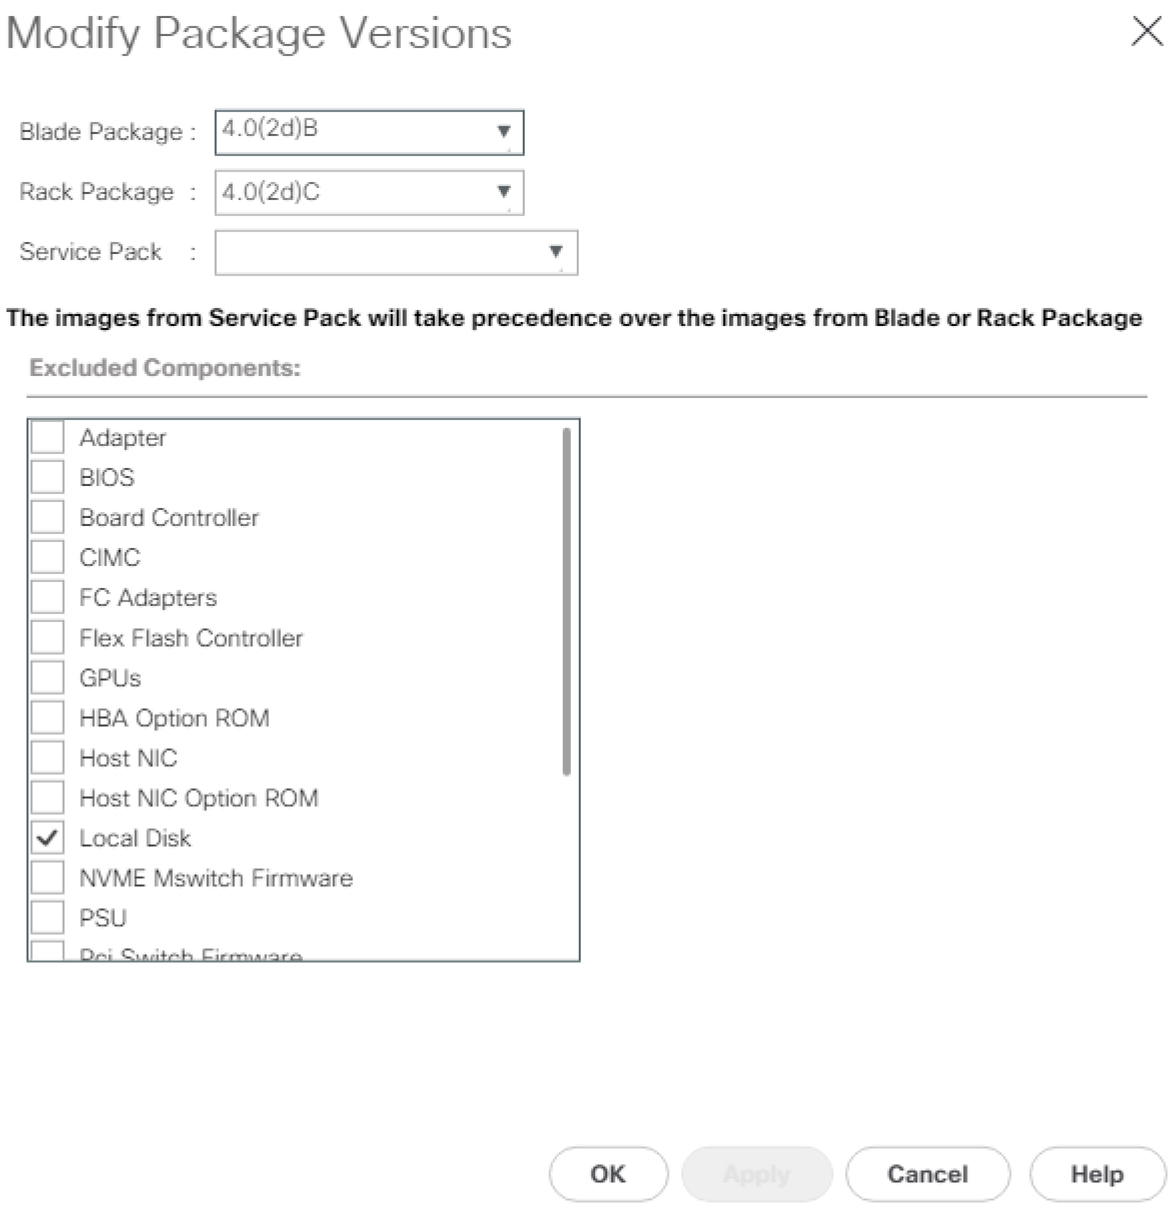

Modify Package Versions

Blade Package : 4.0(2d)B

Rack Package : 4.0(2d)C

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

- ☐ Adapter
- ☐ BIOS
- ☐ Board Controller
- ☐ CIMC
- ☐ FC Adapters
- ☐ Flex Flash Controller
- ☐ GPUs
- ☐ HBA Option ROM
- ☐ Host NIC
- ☐ Host NIC Option ROM
- ☑ Local Disk
- ☐ NVME Mswitch Firmware
- ☐ PSU
- ☐ Pci Switch Firmware

7. Click OK, then click OK again to modify the host firmware package.

Set Jumbo Frames in Cisco UCS Fabric

Jumbo Frames are used in FlexPod for the NFS and iSCSI storage protocols. The normal best practice in FlexPod has been to set the MTU of the Best Effort QoS System Class in Cisco UCS Manager to 9216 for Jumbo Frames. In the Cisco UCS 6454 Fabric Interconnect the MTU for the Best Effort QoS System Class is fixed at normal and cannot be changed. Testing has shown that even with this setting of normal in the 6454, Jumbo Frames can pass through the Cisco UCS fabric without being dropped. The screenshot below is from Cisco UCS Manager on a 6454 Fabric Interconnect, where the MTU for the Best Effort class is not settable. To configure jumbo frames in the Cisco UCS fabric in a 6300 or 6200 series Fabric Interconnect, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand 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.

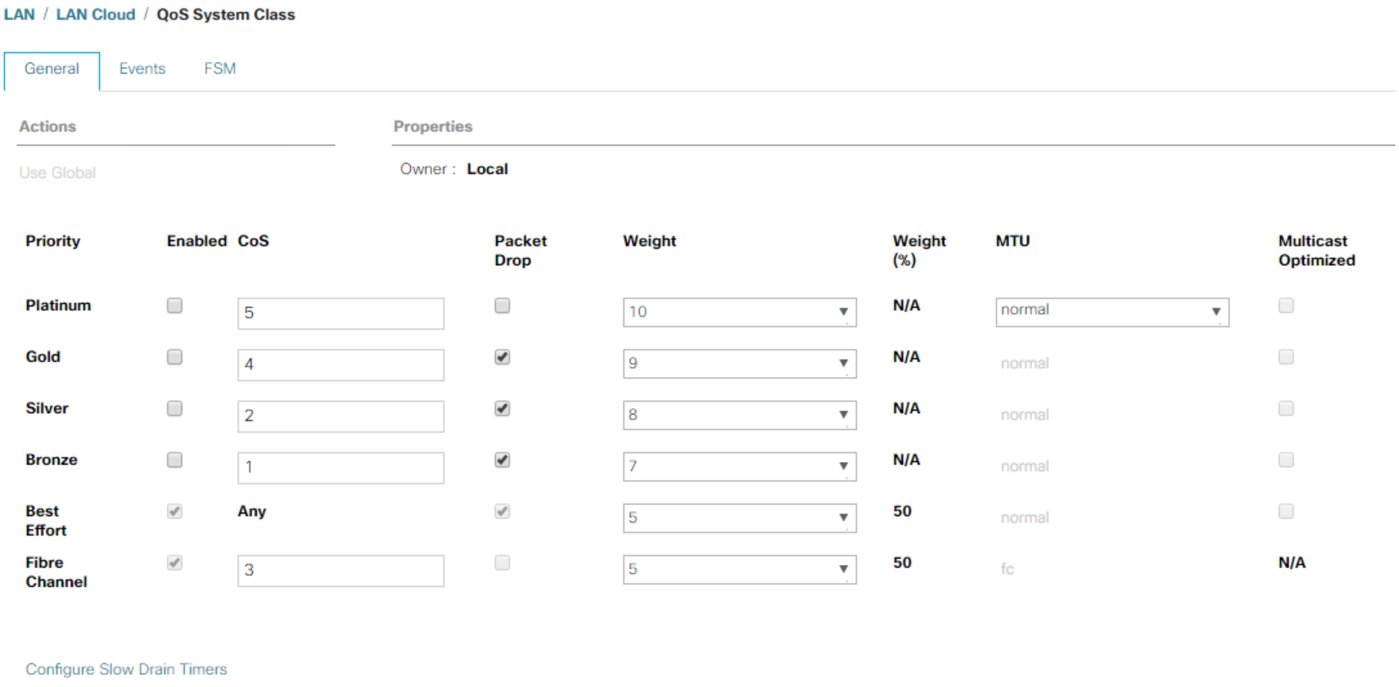

LAN / LAN Cloud / **QoS System Class**

Owner : **Local**

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	☐	5	☐	10	N/A	normal	☐
Gold	☐	4	☑	9	N/A	normal	☐
Silver	☐	2	☑	8	N/A	normal	☐
Bronze	☐	1	☑	7	N/A	normal	☐
Best Effort	☑	Any	☑	5	50	normal	☐
Fibre Channel	☑	3	☐	5	50	fc	N/A

Note that only the Fibre Channel and Best Effort QoS System Classes are enabled in this FlexPod implementation. The Cisco UCS and Nexus switches are intentionally configured this way so that all IP traffic within the FlexPod will be treated as Best Effort. Enabling the other QoS System Classes without having a comprehensive, end-to-end QoS setup in place can cause difficult to troubleshoot issues. For example, NetApp storage controllers by default mark IP-based, VLAN-tagged storage protocol packets with a CoS value of 4. With the default configuration on the Nexus switches in this implementation, storage packets will pass through the switches and into the Cisco UCS Fabric Interconnects with CoS 4 set in the packet header. If the Gold QoS System Class in the Cisco UCS is enabled and the corresponding CoS value left at 4, these storage packets will be treated according to that class and if Jumbo Frames is being used for the storage protocols, but the MTU of the Gold QoS System Class is not set to Jumbo (9216), packet drops will occur. Note also that if the Platinum class is enabled, the MTU must be set to 9216 to use Jumbo Frames in that class.

Create Local Disk Configuration Policy (Optional)

A local disk configuration specifying no local disks for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.

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

To create a local disk configuration policy, follow these steps:

1. In Cisco UCS Manager, click Servers .

2. Expand 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.

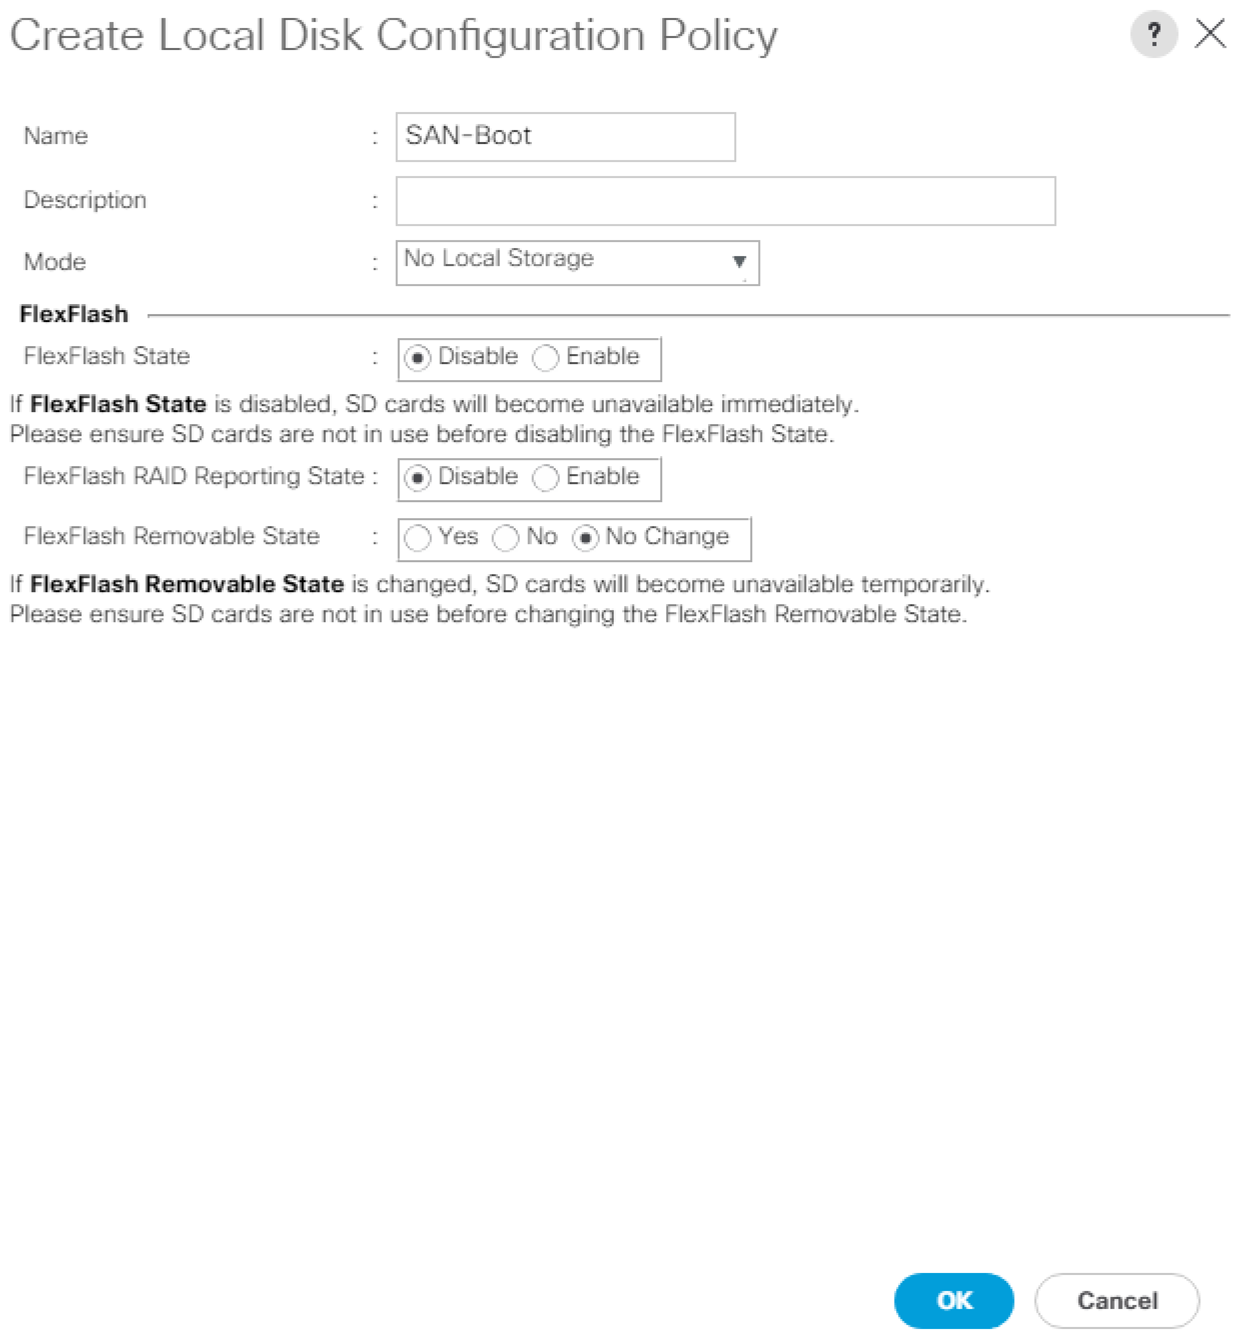

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : ● Disable ○ Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately.
Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : ● Disable ○ Enable

FlexFlash Removable State : ○ Yes ○ No ● No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily.
Please ensure SD cards are not in use before changing the FlexFlash Removable State.

7. Click OK to create the local disk configuration policy.

8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)

To create a network control policy that enables CDP and LLDP on server virtual network controller (vNIC) ports, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable-CDP-LLDP as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.

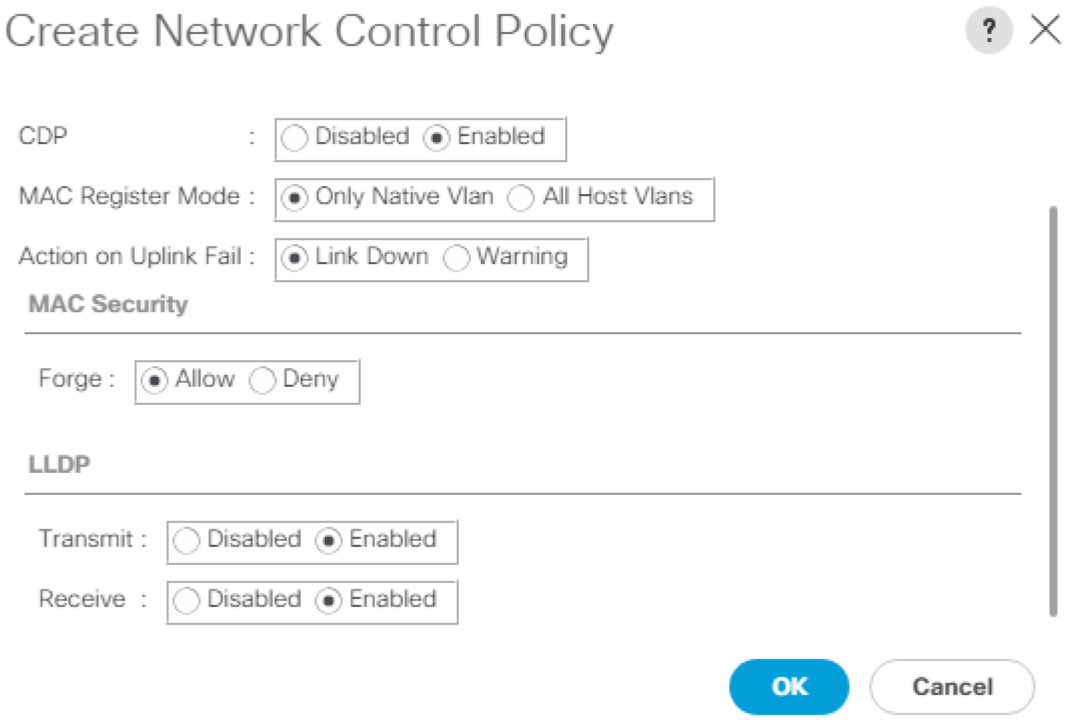

Create Network Control Policy

CDP : ○ Disabled ● Enabled

MAC Register Mode : ● Only Native Vlan ○ All Host Vlans

Action on Uplink Fail : ● Link Down ○ Warning

MAC Security

Forge : ● Allow ○ Deny

LLDP

Transmit : ○ Disabled ● Enabled

Receive : ○ Disabled ● Enabled

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

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand 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.

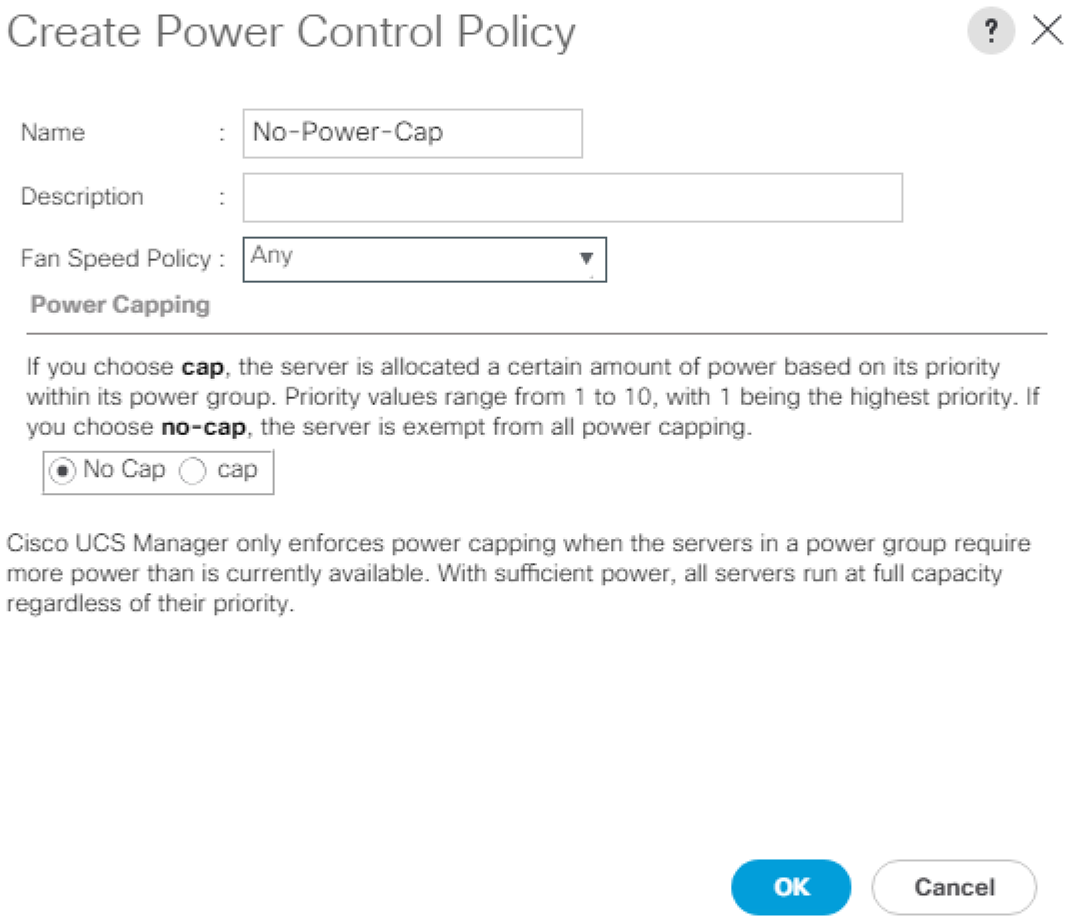

Create Power Control Policy

Name : No-Power-Cap

Description :

Fan Speed Policy : Any

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

● No Cap ○ cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

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, follow these steps:

This example creates a policy for Cisco UCS B200 M5 servers for a server pool.

1. In Cisco UCS Manager, click Servers .

2. Expand Policies > root.

3. Right-click Server Pool Policy Qualifications.

4. Select Create Server Pool Policy Qualification.

5. Name the policy UCS-B200M5.

6. Select Create Server PID Qualifications.

7. Select UCSB-B200-M5 from the PID drop-down.

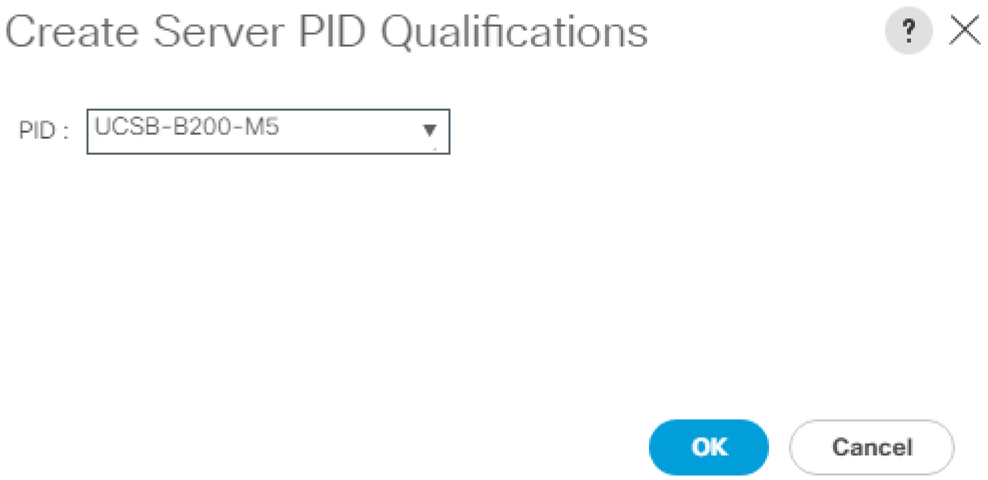

Create Server PID Qualifications

PID : UCSB-B200-M5

8. Click OK

9. Optionally select additional qualifications to refine server selection parameters for the server pool.

10. Click OK to create the policy then OK for the confirmation.

Create Server BIOS Policy

To create a server BIOS policy for VMware ESXi hosts within the FlexPod organization, follow these steps:

In this lab validation, all Cisco UCS B200 M5 and Cisco UCS C220 M5 servers had TPM2.0 modules installed. To utilize TPM2.0 functionality with VMware vSphere 6.7U1, the TPM module must be enabled and Trusted Execution Technology(TXT) disabled in BIOS. According to the Cisco UCS Server BIOS Tokens, Release 4.0 document, these settings are the default or Platform Default settings for all M5 servers. Because of this, these settings do not have to be added to this BIOS policy.

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root > Sub-Organizations > FlexPod Organization.

3. Right-click BIOS Policies under FlexPod Organization.

4. Select Create BIOS Policy.

5. Enter VM-Host as the BIOS policy name.

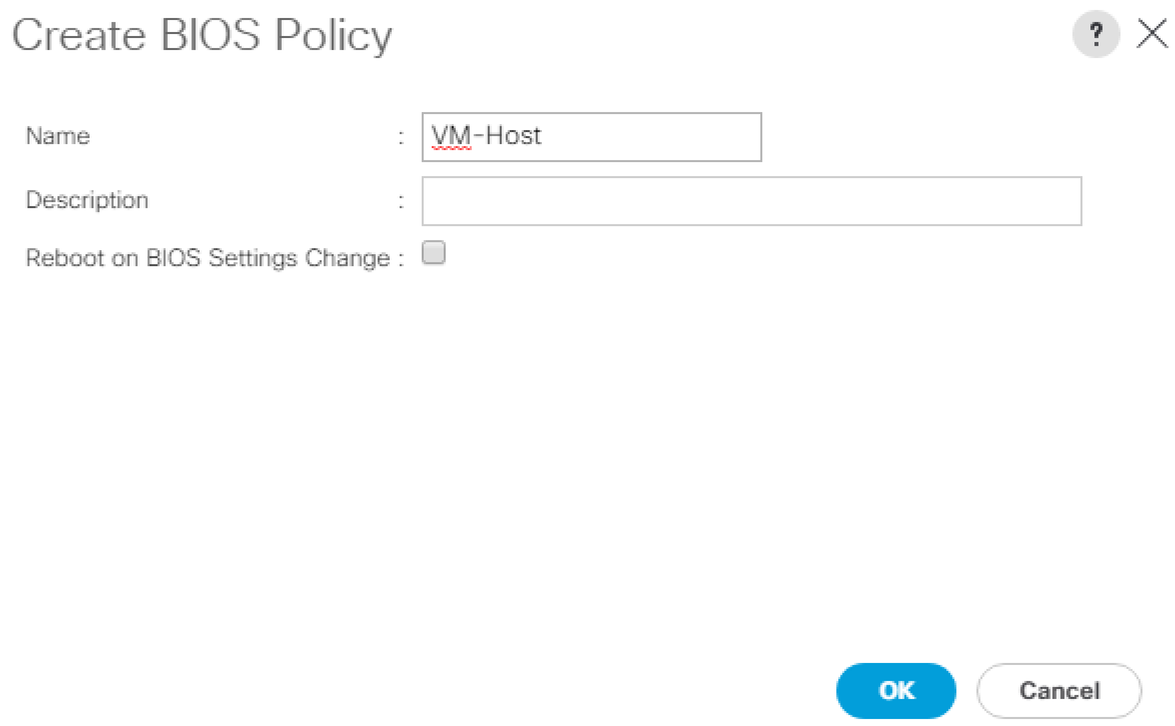

Create BIOS Policy

Name : VM-Host

Description :

Reboot on BIOS Settings Change : □

6. Click OK, then OK again to create the BIOS Policy.

7. Under the FlexPod Organization, expand BIOS Policies and select the newly created BIOS Policy. Set the following within the Main tab of the Policy:

   a. CDN Control -> Enabled

   b. Quiet Boot -> Disabled

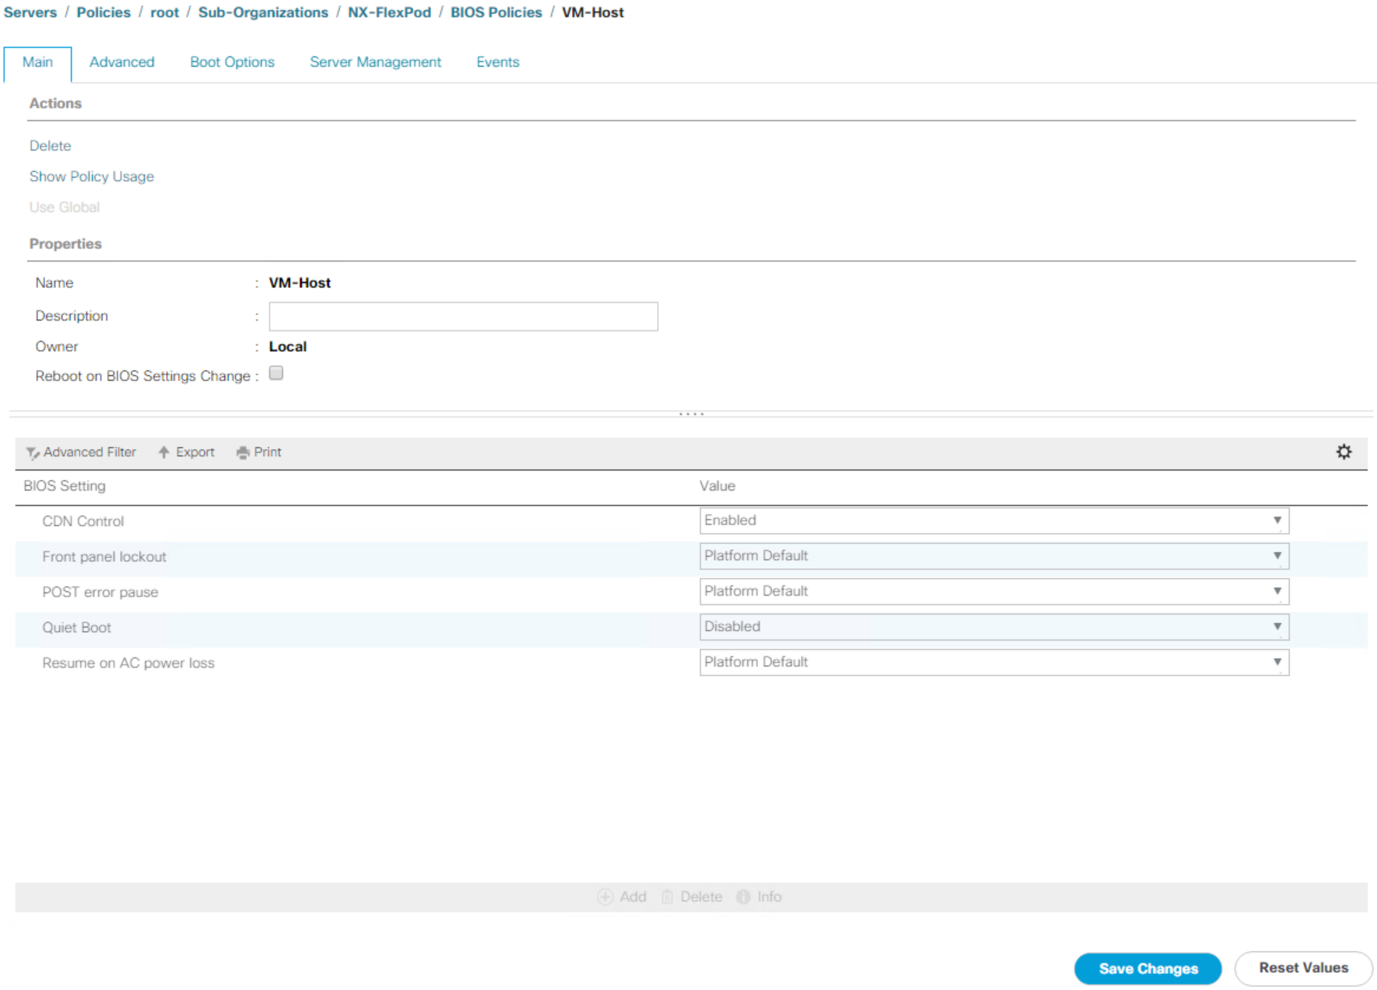

Servers / Policies / root / Sub-Organizations / NX-FlexPod / BIOS Policies / VM-Host

Name : VM-Host

Description :

Owner : Local

Reboot on BIOS Settings Change : □

BIOS Setting	Value
CDN Control	Enabled
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Disabled
Resume on AC power loss	Platform Default

8. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab. Set the following within the Processor tab:

   a. DRAM Clock Throttling -> Performance

   b. Frequency Floor Override -> Enabled

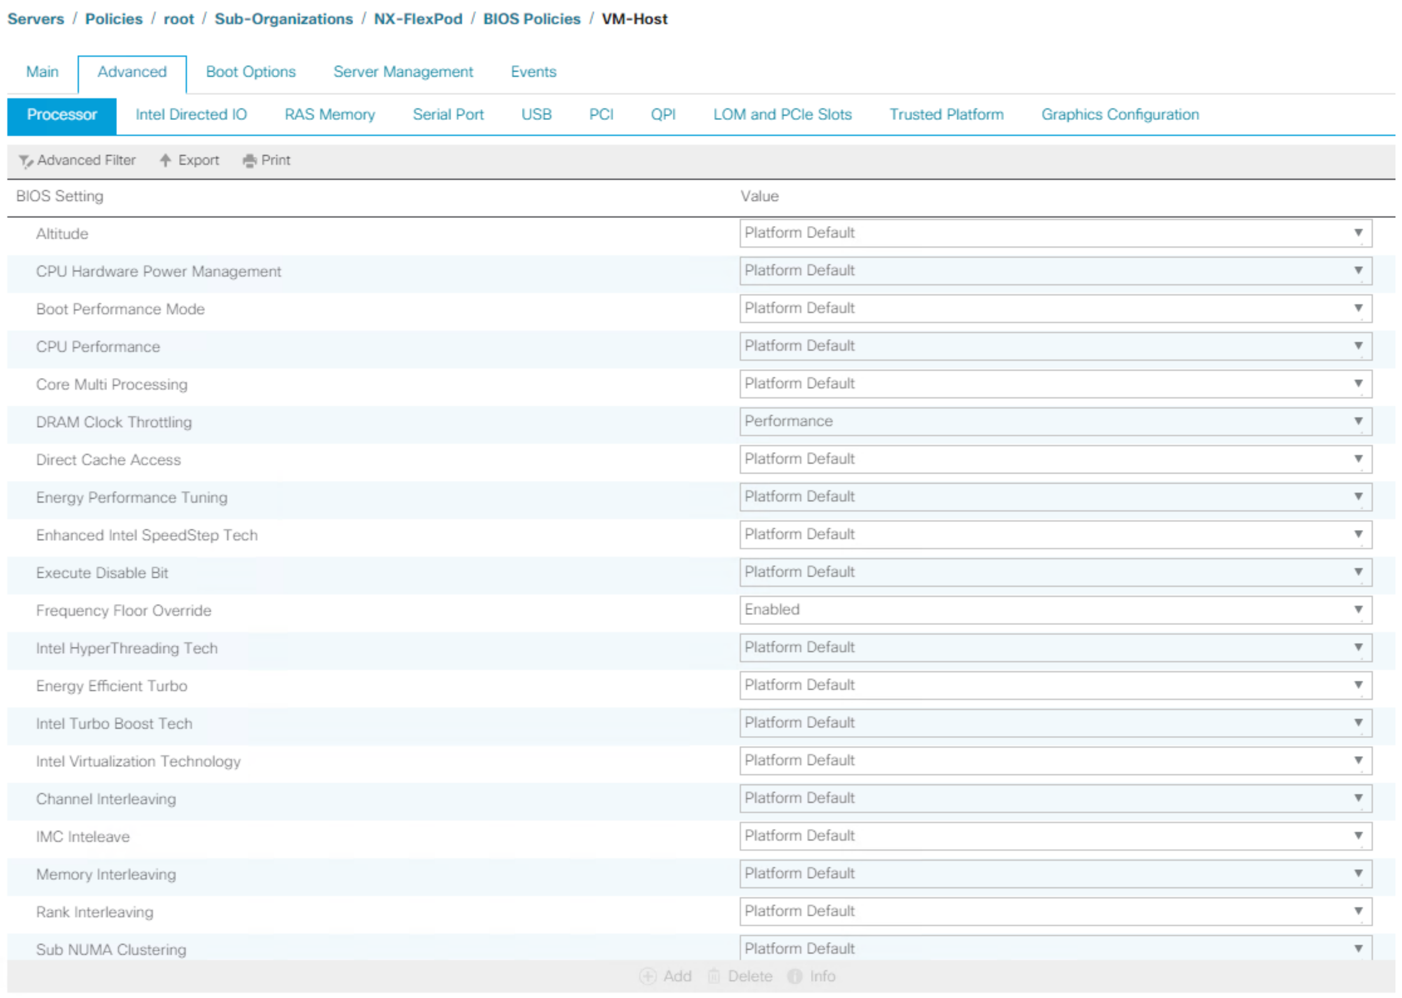

BIOS Setting	Value
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
Boot Performance Mode	Platform Default
CPU Performance	Platform Default
Core Multi Processing	Platform Default
DRAM Clock Throttling	Performance
Direct Cache Access	Platform Default
Energy Performance Tuning	Platform Default
Enhanced Intel SpeedStep Tech	Platform Default
Execute Disable Bit	Platform Default
Frequency Floor Override	Enabled
Intel HyperThreading Tech	Platform Default
Energy Efficient Turbo	Platform Default
Intel Turbo Boost Tech	Platform Default
Intel Virtualization Technology	Platform Default
Channel Interleaving	Platform Default
IMC Inteleave	Platform Default
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default

9. Scroll down to the remaining Processor options and select:

   a. Processor C State -> Disabled

   b. Processor C1E -> Disabled

   c. Processor C3 Report -> Disabled

   d. Processor C7 Report -> Disabled

   e. Energy Performance -> Performance

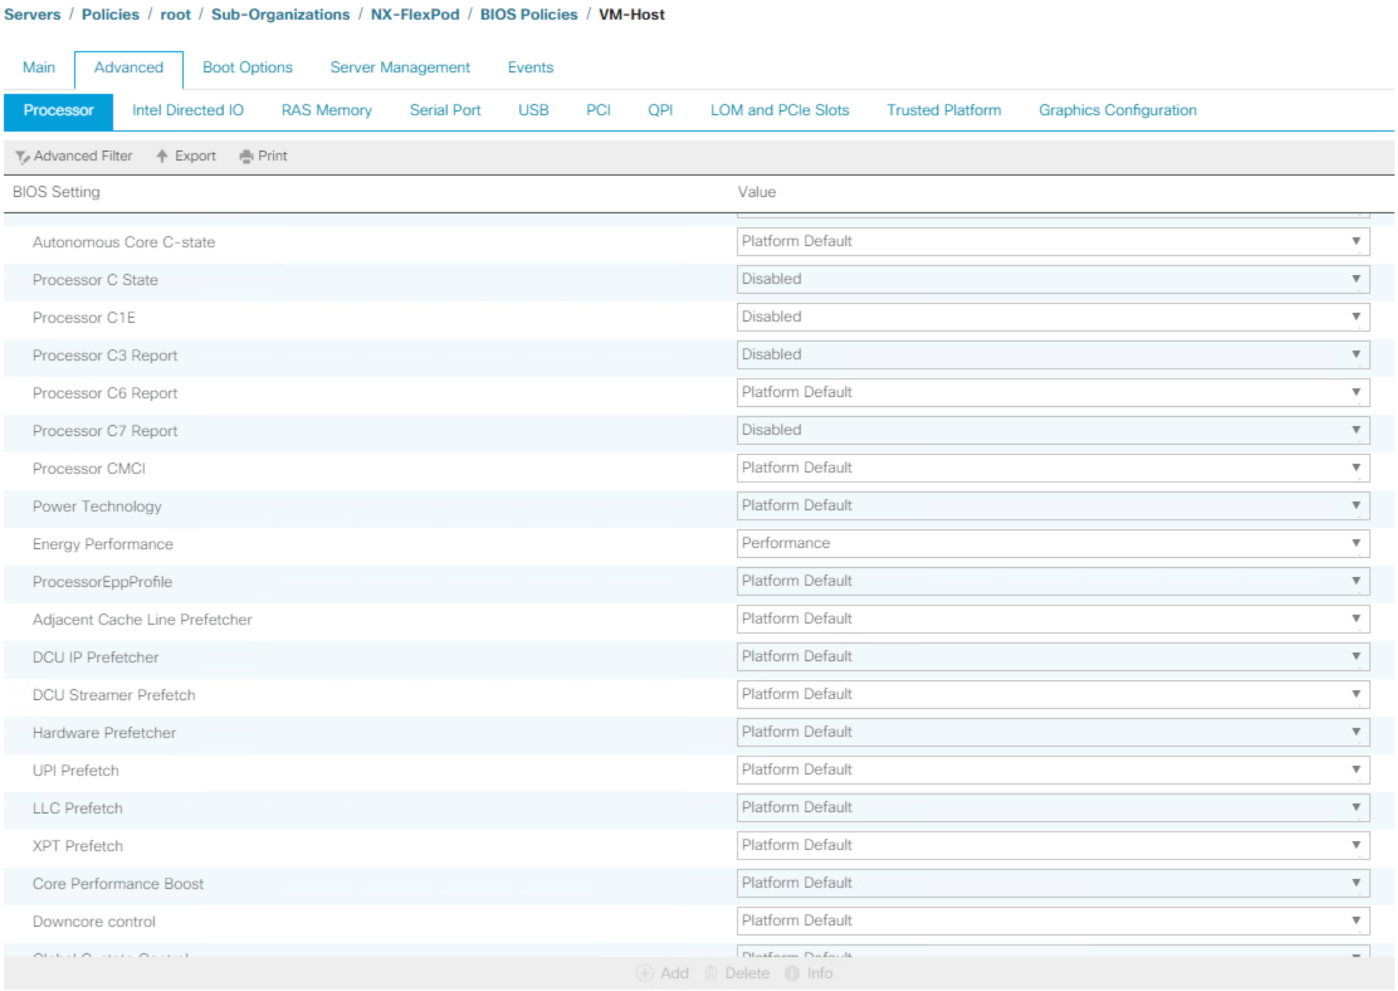

BIOS Setting	Value
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Platform Default
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Platform Default
Energy Performance	Performance
ProcessorEppProfile	Platform Default
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Core Performance Boost	Platform Default
Downcore control	Platform Default
Global C-state Control	Platform Default

10. Click the RAS Memory tab, and select:

    a. LV DDR Mode -> Performance Mode

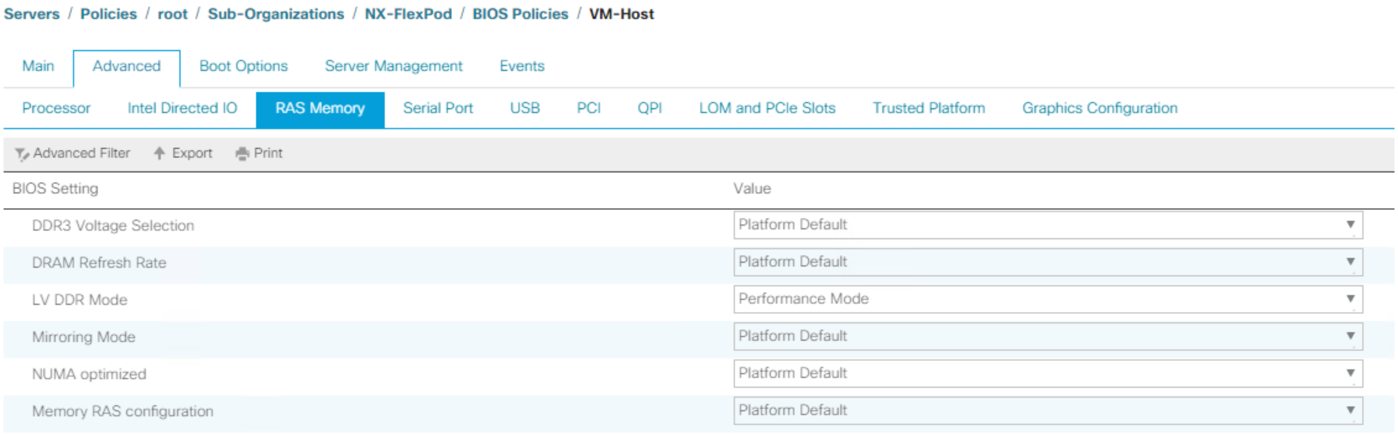

BIOS Setting	Value
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

11. Click Save Changes.

12. Click OK.

Update the Default Maintenance Policy

To update the default Maintenance Policy to either require user acknowledgement before server boot when service profiles change or to make the changes on the next server reboot, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select "On Next Boot" to delegate maintenance windows to server administrators.

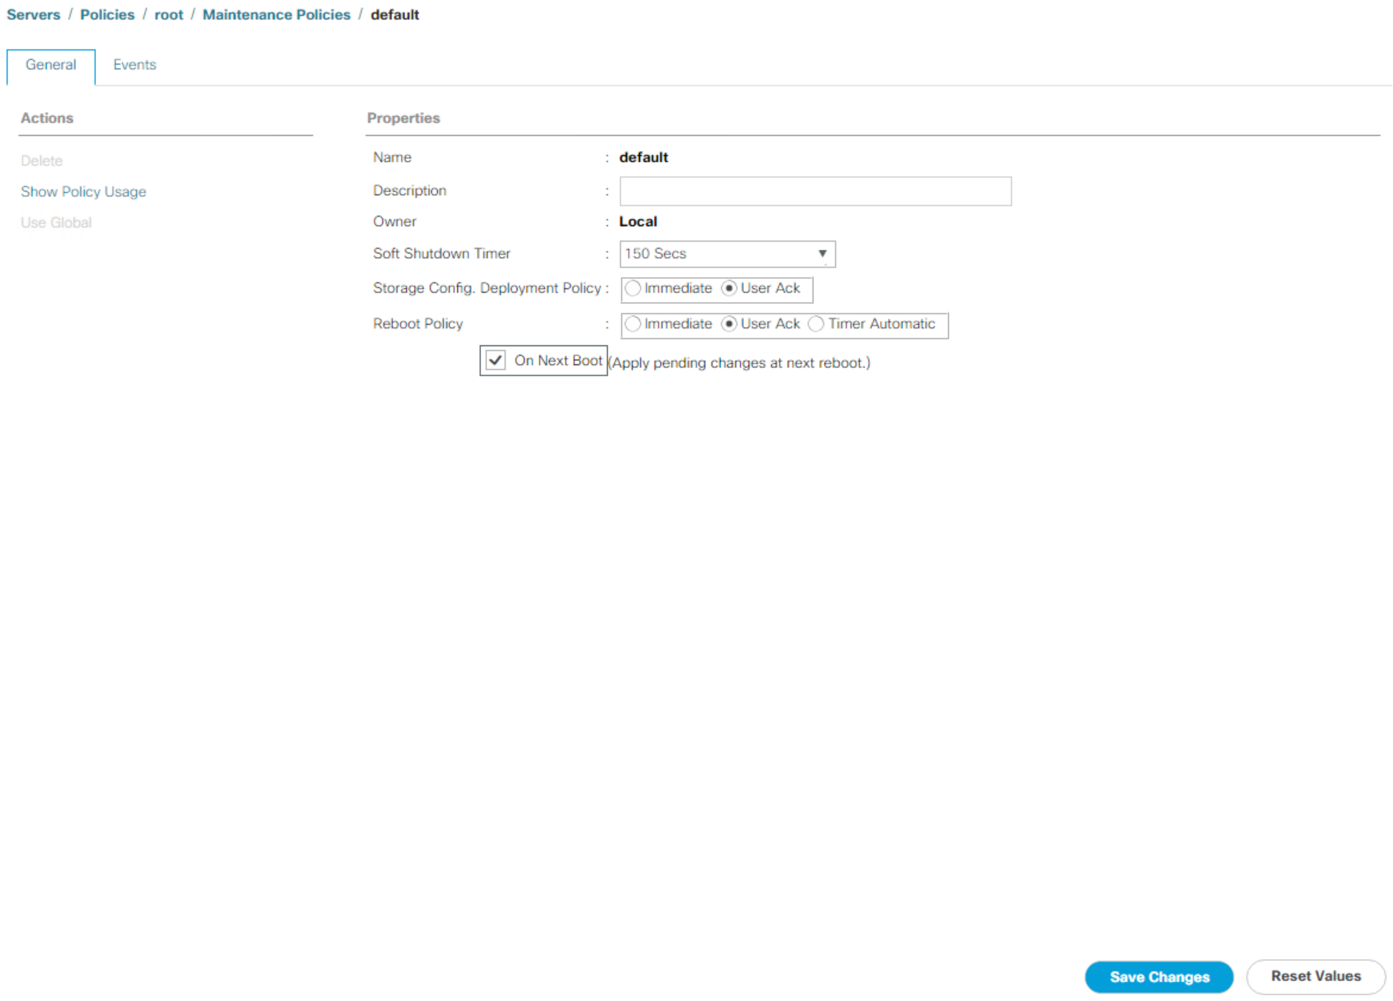

Servers / Policies / root / Maintenance Policies / default

Name : default

Description :

Owner : Local

Soft Shutdown Timer : 150 Secs ▼

Storage Config. Deployment Policy : ○ Immediate ● User Ack

Reboot Policy : ○ Immediate ● User Ack ○ Timer Automatic

☑ On Next Boot (Apply pending changes at next reboot.)

6. Click Save Changes.
7. Click OK to accept the changes.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates within the FlexPod organization, follow these steps. A total of 4 vNIC Templates will be created. Two of the vNIC templates (Infra-A and Infra-B) will be created for vNICs to connect to VMware ESXi vSwitch0. vSwitch0 will have port groups for the IB-MGMT, Infra-NFS, vMotion, and VM-Traffic VLANs. The third and fourth vNIC templates (Infra-vDS-A and Infra-vDS-B) will be created for vNICs to connect to the VMware Virtual Distributed Switch (vDS). The vDS will have port groups for the IB-MGMT, vMotion, and VM-Traffic VLANs. The vMotion VLAN is being placed on both vSwitch0 and the vDS so that the vMotion VMkernel port can initially be created on vSwitch0 then migrated to the vDS to allow QoS marking of vMotion packets to occur within the vDS if QoS policies need to be applied to vMotion in the future. The IB-MGMT VLAN is being placed on both vSwitch0 and the vDS so that management virtual machines other than vCenter can be placed on the vDS. The VM-Traffic VLAN is being placed on both vSwitch0 and the vDS so that it will be allowed on vSwitch0 in case the customer chooses not to implement the vDS on the management ESXi hosts.

Create Infrastructure vNIC Templates

To create the infrastructure vNIC templates, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Under the FlexPod Organization, right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter Infra-A as the vNIC template name.
6. Keep Fabric A selected.
7. Select the Enable Failover checkbox.

Although it is not a best practice to use vNIC Failover along with VMware NIC Teaming in the virtual switch, Failover is being used here to mitigate the effects of DDTS id: CSCvp49398 which impacts C-Series servers connected directly to Cisco UCS 6454 Fabric Interconnects using 25GbE and Cisco VIC 1455 or 1457. Due to this bug, data loss for up to 75-90 seconds can be observed when the Fabric Interconnect being actively used to forward server traffic is rebooted. This data loss issue is not observed on Cisco UCS B-Series servers and this bug does not impact the traffic forwarding under normal operation. Since a Fabric Interconnect reboot is required when upgrading the Cisco UCS system, customers should plan their upgrade during off-peak hours to minimize traffic disruption. If only B-Series servers are being used in your implementation, do not enable Failover.

8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Select Updating Template as the Template Type.

12. Under VLANs, select the checkboxes for IB-MGMT, Infra-NFS, vMotion, VM-Traffic, and Native-VLAN VLANs.

13. Set Native-VLAN as the native VLAN.

14. Select vNIC Name for the CDN Source.

15. For MTU, enter 9000.

16. In the MAC Pool list, select MAC-Pool-A.

17. In the Network Control Policy list, select Enable-CDP-LLDP.

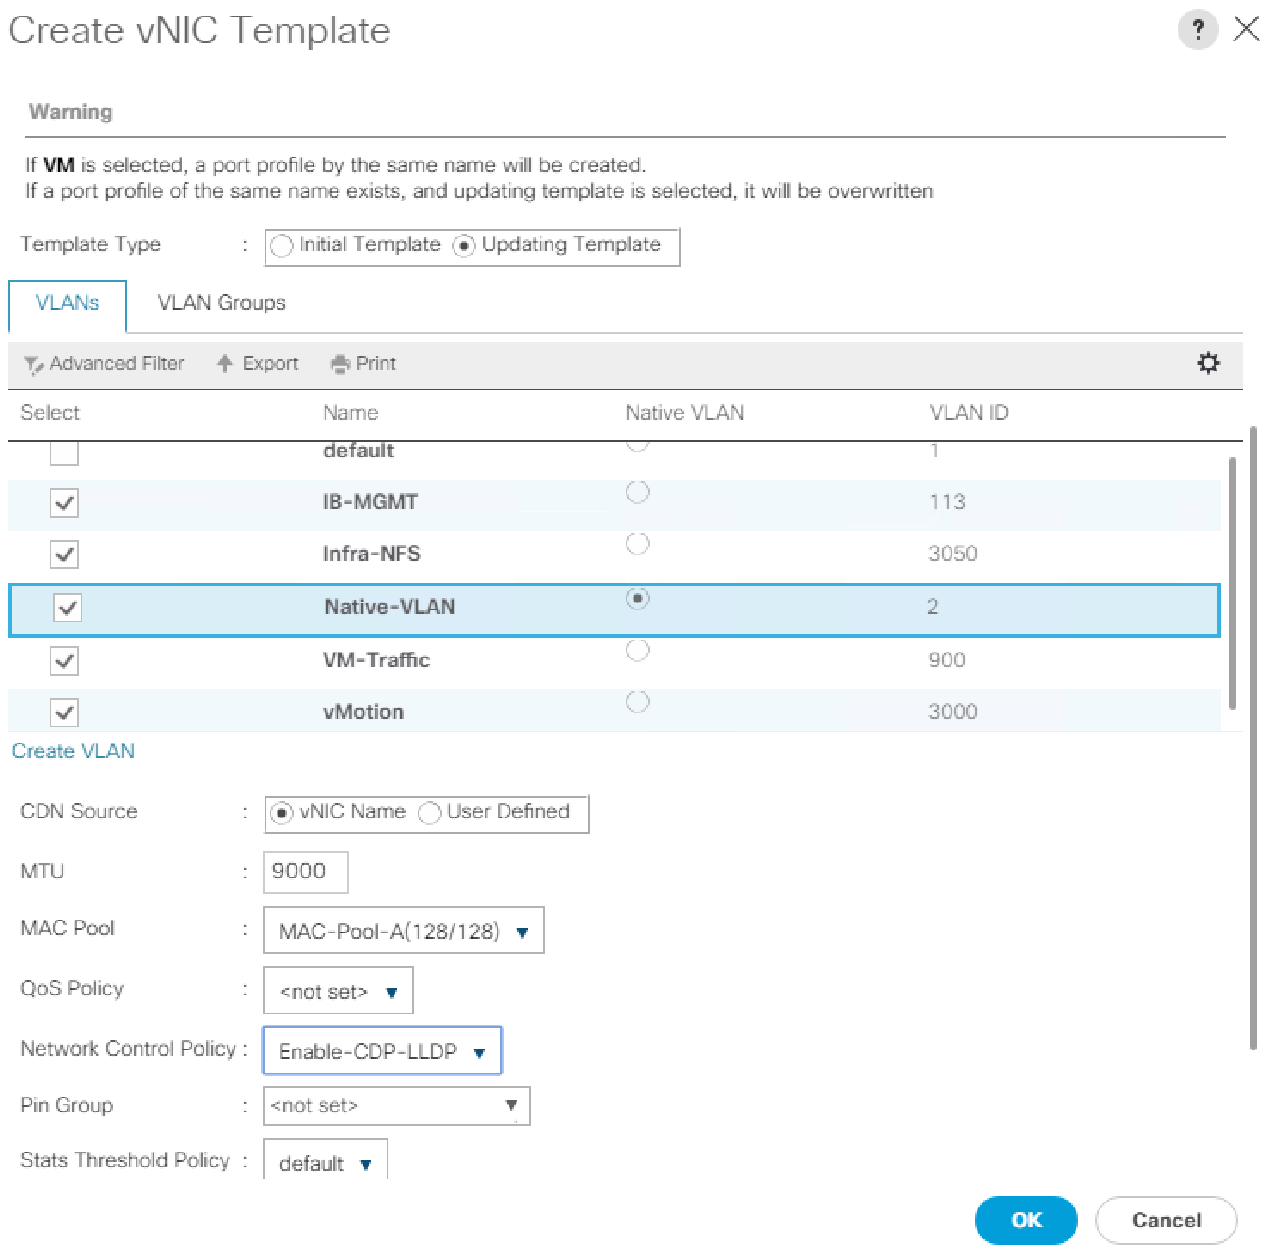

Create vNIC Template

Warning

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

Template Type : ○ Initial Template ● Updating Template

Select	Name	Native VLAN	VLAN ID
□	default	○	1
☑	IB-MGMT	○	113
☑	Infra-NFS	○	3050
☑	Native-VLAN	●	2
☑	VM-Traffic	○	900
☑	vMotion	○	3000

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(128/128)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

18. Click OK to create the vNIC template.
19. Click OK.
20. Under the FlexPod organization, right-click vNIC Templates.
21. Select Create vNIC Template.
22. Enter Infra-B as the vNIC template name.
23. Select Fabric B.
24. Select the Enable Failover checkbox.

If only Cisco UCS B-Series servers are being used in your implementation, do not enable Failover.

25. Set Redundancy Type to Secondary Template.
26. Select Infra-A for the Peer Redundancy Template.
27. In the MAC Pool list, select MAC-Pool-B.

The MAC Pool is all that needs to be selected for the Secondary Template, all other values will either be propagated from the Primary Template or set at default values.

28. Click OK to create the vNIC template.
29. Click OK.
30. Under the FlexPod Organization, right-click vNIC Templates.
31. Select Create vNIC Template.
32. Enter Infra-vDS-A as the vNIC template name.
33. Keep Fabric A selected.
34. Select the Enable Failover checkbox.

If only Cisco UCS B-Series servers are being used in your implementation, do not enable Failover.

35. Select Primary Template for Redundancy Type.
36. Leave the Peer Redundancy Template set to <not set>.
37. Under Target, make sure that only the Adapter checkbox is selected.
38. Select Updating Template as the Template Type.

39. Under VLANs, select the checkboxes for IB-MGMT, vMotion, VM-Traffic, and Native-VLAN VLANs.

40. Set Native-VLAN as the native VLAN.

41. Select vNIC Name for the CDN Source.

42. For MTU, enter 9000.

43. In the MAC Pool list, select MAC-Pool-A.

44. In the Network Control Policy list, select Enable-CDP-LLDP.

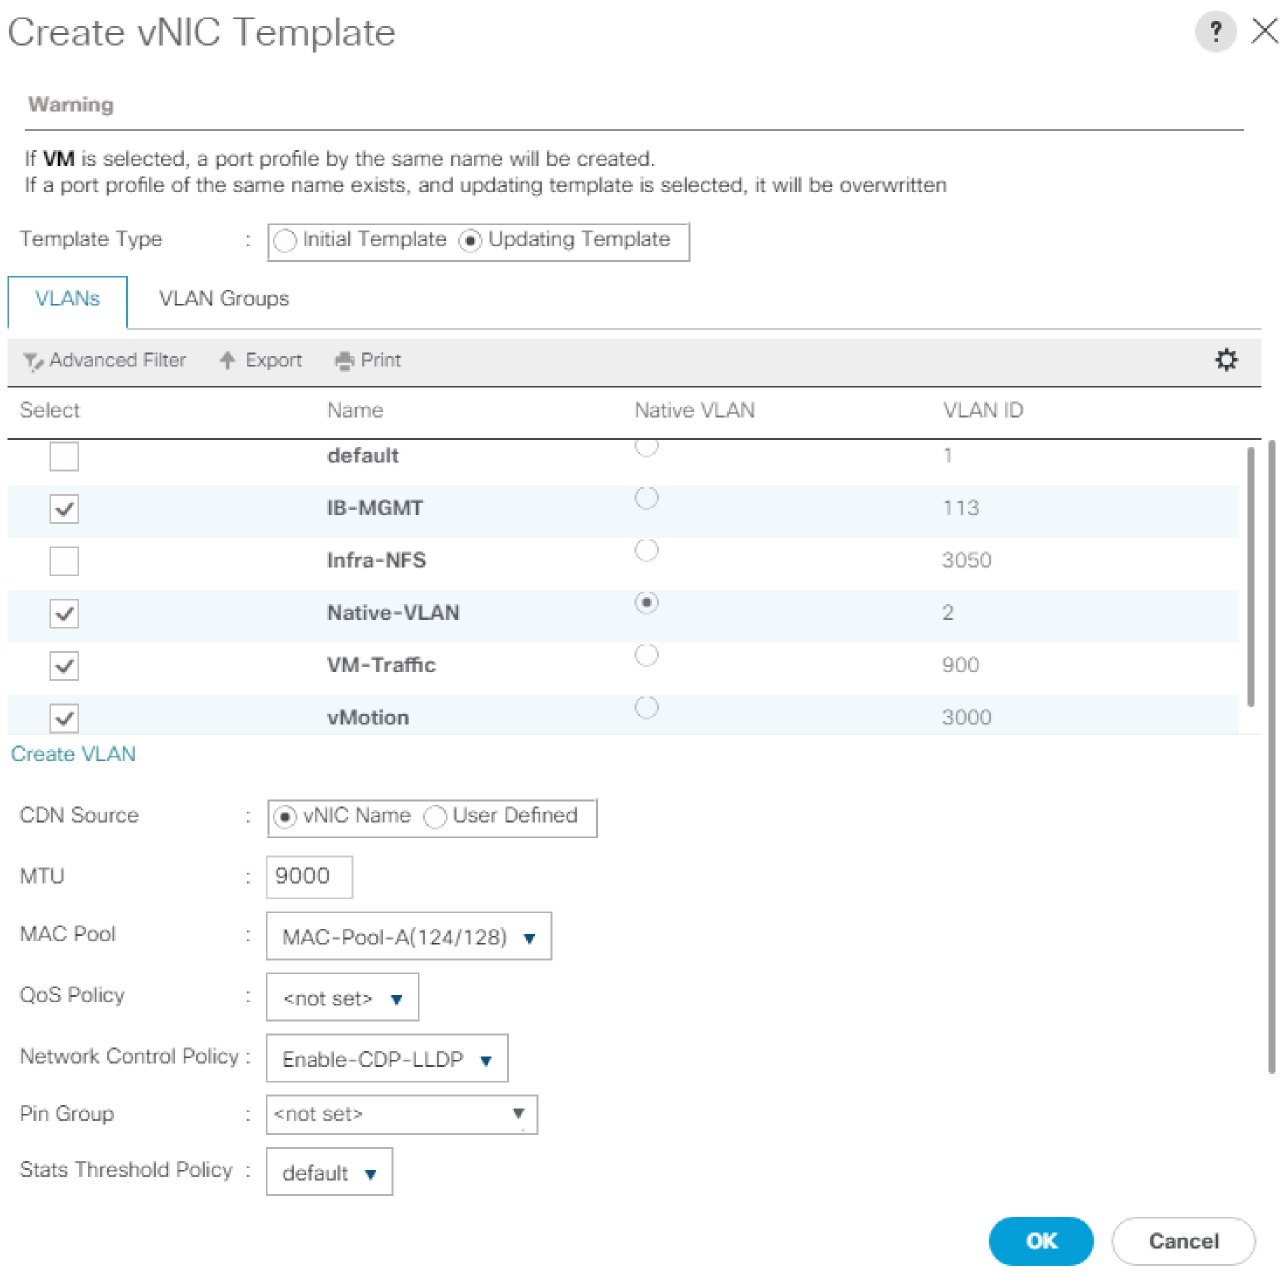

Create vNIC Template

Warning

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

Template Type : ○ Initial Template ● Updating Template

Select	Name	Native VLAN	VLAN ID
☐	**default**	○	1
☑	**IB-MGMT**	○	113
☐	**Infra-NFS**	○	3050
☑	**Native-VLAN**	●	2
☑	**VM-Traffic**	○	900
☑	**vMotion**	○	3000

CDN Source :	● vNIC Name ○ User Defined
MTU :	9000
MAC Pool :	MAC-Pool-A(124/128) ▾
QoS Policy :	<not set> ▾
Network Control Policy :	Enable-CDP-LLDP ▾
Pin Group :	<not set> ▾
Stats Threshold Policy :	default ▾

45. Click OK to create the vNIC template.

46. Click OK.

47. Under the FlexPod organization, right-click vNIC Templates.

48. Select Create vNIC Template

49. Enter Infra-vDS-B as the vNIC template name.

50. Select Fabric B.

51. Select the Enable Failover checkbox.

If only Cisco UCS B-Series servers are being used in your implementation, do not enable Failover.

52. Set Redundancy Type to Secondary Template.

53. Select Infra-vDS-A for the Peer Redundancy Template.

54. In the MAC Pool list, select MAC-Pool-B.

The MAC Pool is all that needs to be selected for the Secondary Template, all other values will either be propagated from the Primary Template or set at default values.

55. Click OK to create the vNIC template.

56. Click OK.

Create High Traffic VMware Adapter Policy

To create the optional VMware-High-Traffic Ethernet Adapter policy to provide higher vNIC performance, follow these steps:

This Ethernet Adapter policy can be attached to vNICs when creating the LAN Connectivity policy for vNICs that have large amounts of traffic on multiple flows or TCP sessions. This policy provides more hardware transmit and receive queues handled by multiple CPUs to the vNIC.

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Right-click Adapter Policies and select Create Ethernet Adapter Policy.

4. Name the policy VMware-HighTrf.

5. Expand Resources and set the values as shown below.

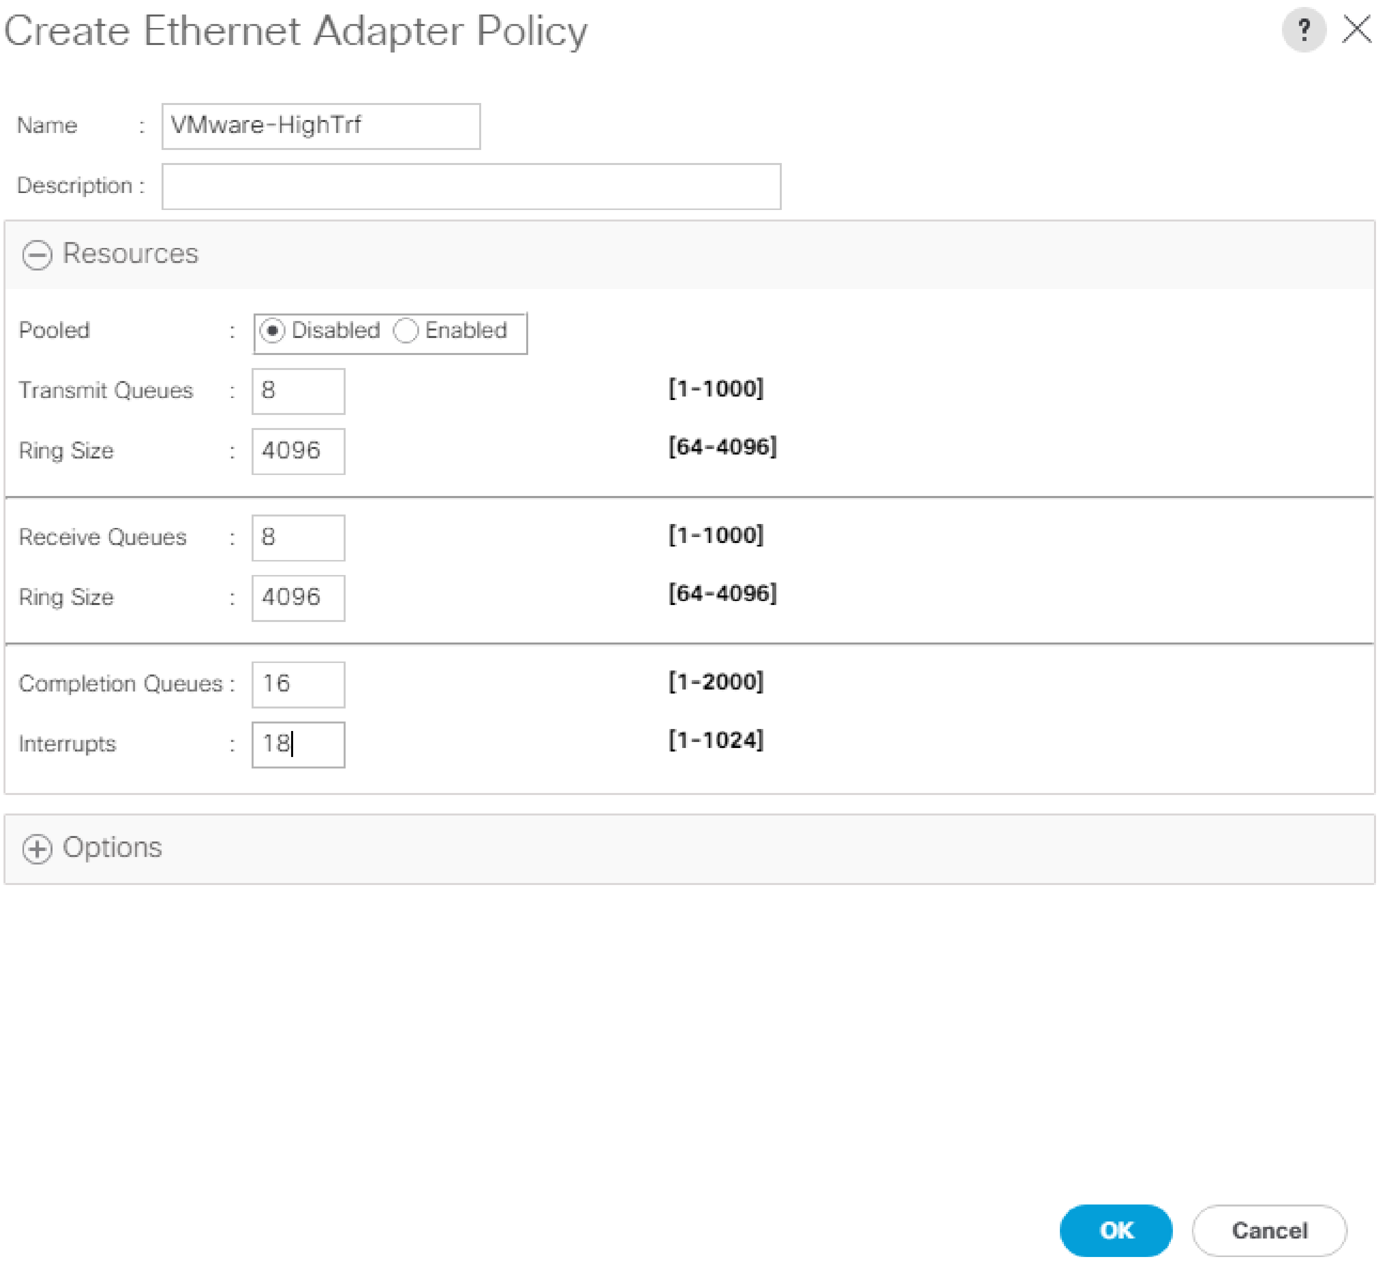

Create Ethernet Adapter Policy

Name : VMware-HighTrf

Description :

Resources

Pooled : ● Disabled ○ Enabled

Transmit Queues :	8	[1-1000]
Ring Size :	4096	[64-4096]
Receive Queues :	8	[1-1000]
Ring Size :	4096	[64-4096]
Completion Queues :	16	[1-2000]
Interrupts :	18	[1-1024]

6. Expand Options and select Enabled for Receive Side Scaling (RSS).

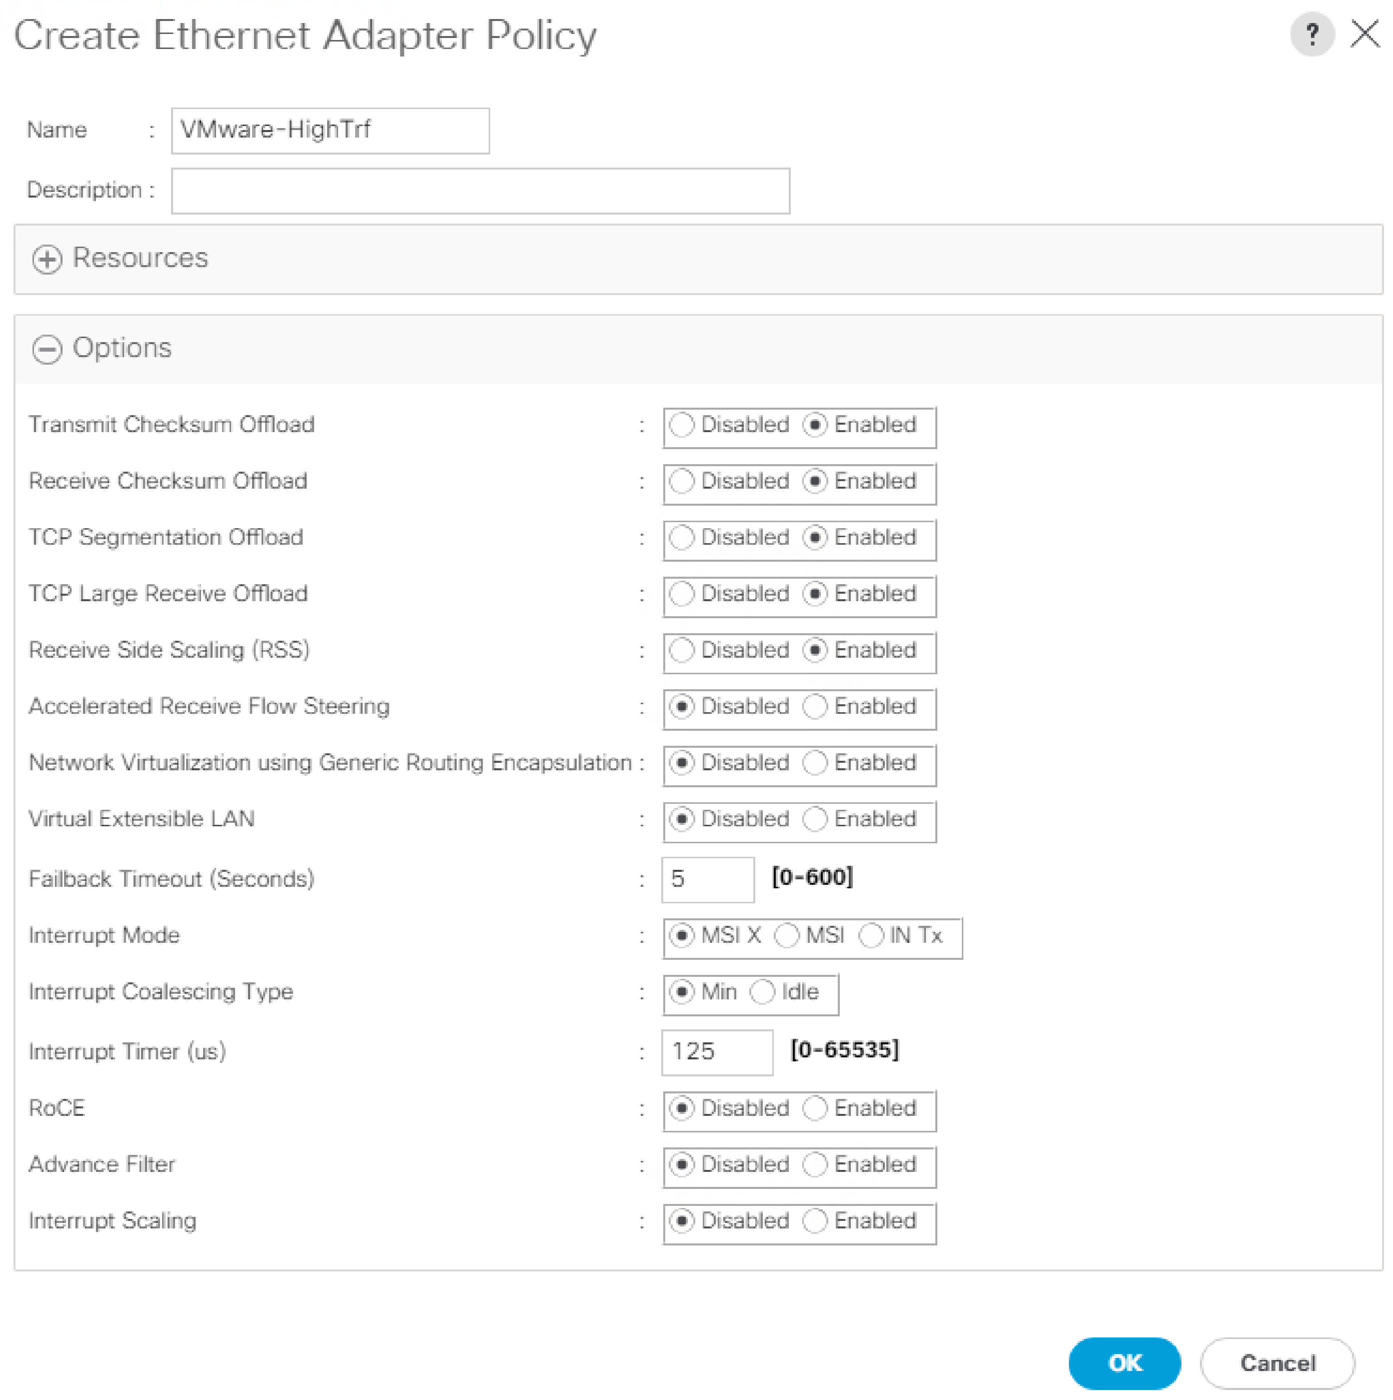

Create Ethernet Adapter Policy

Name	:	VMware-HighTrf
Description :		

⊕ Resources

⊖ Options

Transmit Checksum Offload	:	○ Disabled ● Enabled	
Receive Checksum Offload	:	○ Disabled ● Enabled	
TCP Segmentation Offload	:	○ Disabled ● Enabled	
TCP Large Receive Offload	:	○ Disabled ● Enabled	
Receive Side Scaling (RSS)	:	○ Disabled ● Enabled	
Accelerated Receive Flow Steering	:	● Disabled ○ Enabled	
Network Virtualization using Generic Routing Encapsulation	:	● Disabled ○ Enabled	
Virtual Extensible LAN	:	● Disabled ○ Enabled	
Failback Timeout (Seconds)	:	5	[0-600]
Interrupt Mode	:	● MSI X ○ MSI ○ IN Tx	
Interrupt Coalescing Type	:	● Min ○ Idle	
Interrupt Timer (us)	:	125	[0-65535]
RoCE	:	● Disabled ○ Enabled	
Advance Filter	:	● Disabled ○ Enabled	
Interrupt Scaling	:	● Disabled ○ Enabled	

7. Click OK, then click OK again to complete creating the Ethernet Adapter Policy.

Create LAN Connectivity Policy for FC Boot (FCP)

To configure the necessary Infrastructure LAN Connectivity Policy within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > Policies > root > Sub-Organizations > FlexPod Organization.
3. Under the FlexPod Organization, right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter FC-Boot as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Infra-A.
10. In the Adapter Policy list, select VMWare.

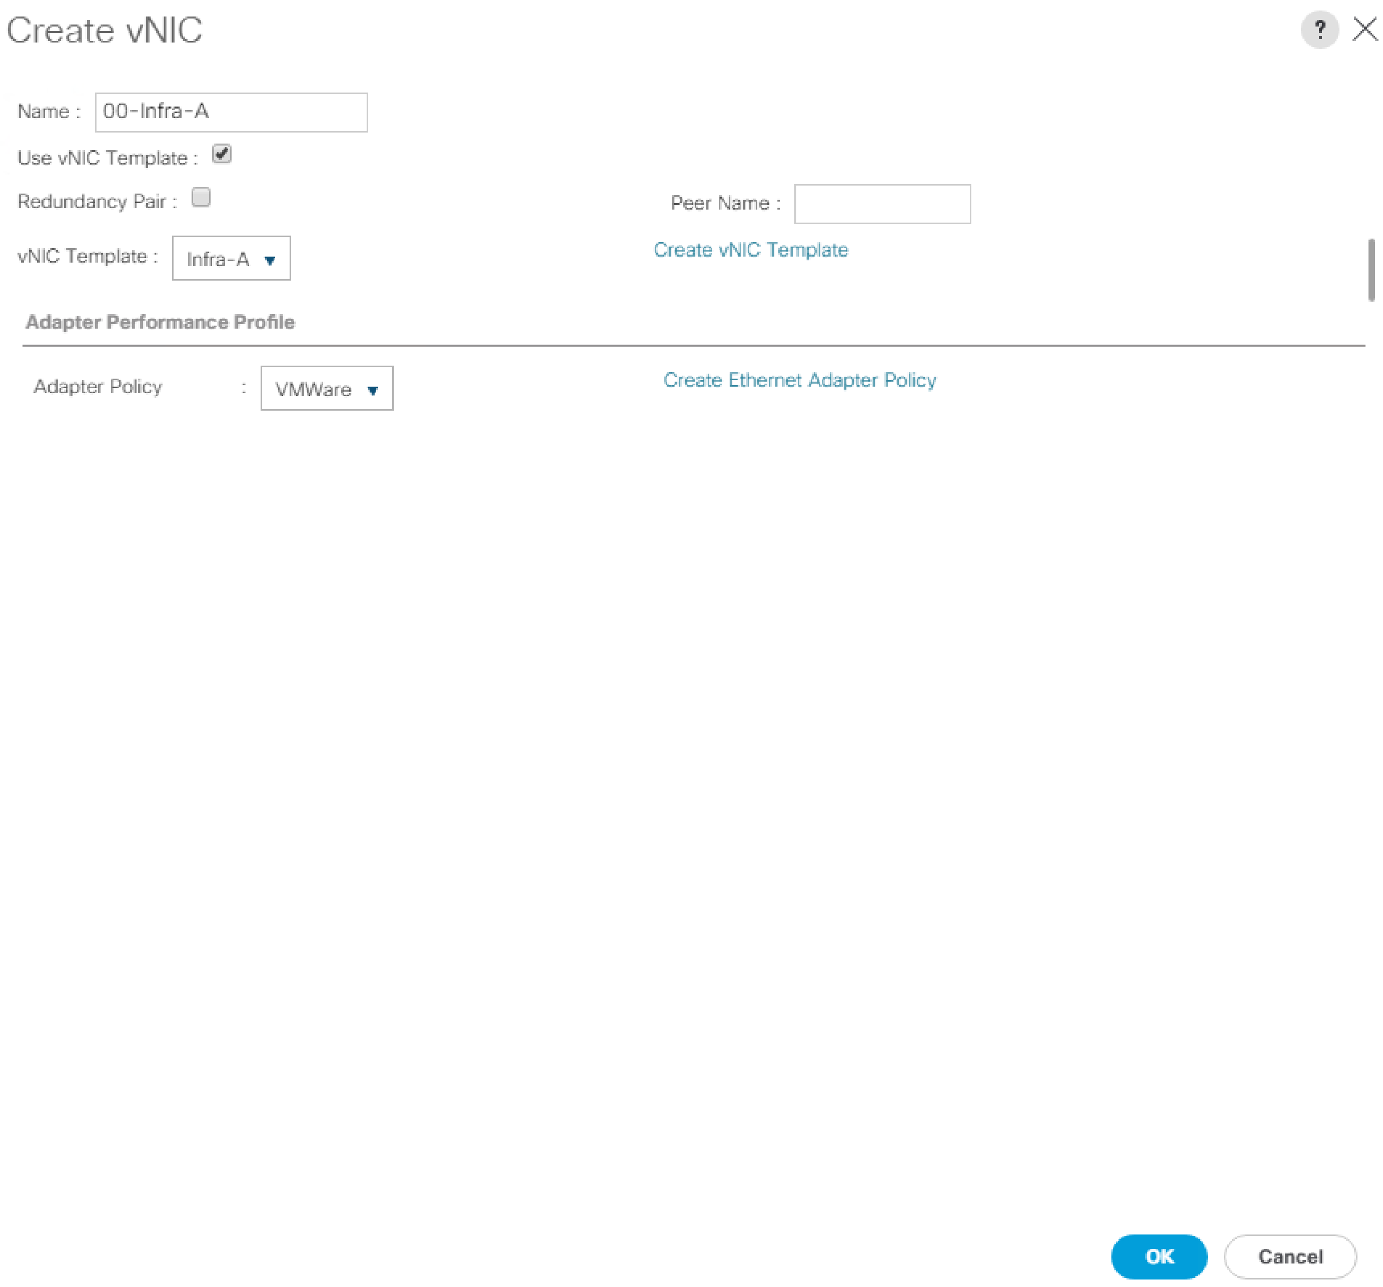

Create vNIC

Name : 00-Infra-A

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : Infra-A

Adapter Performance Profile

Adapter Policy : VMWare

11. Click OK to add this vNIC to the policy.

12. Click Add to add another vNIC to the policy.

13. In the Create vNIC box, enter 01-Infra-B as the name of the vNIC.

14. Select the Use vNIC Template checkbox.

15. In the vNIC Template list, select Infra-B.

16. In the Adapter Policy list, select VMWare.

17. Click OK to add the vNIC to the policy.
18. Click Add to add another vNIC to the policy.
19. In the Create vNIC dialog box, enter 02-Infra-vDS-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select Infra-vDS-A.
22. In the Adapter Policy list, select VMWare.

The VMware-HighTrf Adapter Policy can also be selected for this vNIC.

23. Click OK to add this vNIC to the policy.
24. Click Add to add another vNIC to the policy.
25. In the Create vNIC box, enter 03-Infra-vDS-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select Infra-vDS-B.
28. In the Adapter Policy list, select VMWare.

Select the same Adapter Policy that was selected for 02-Infra-vDS-A.

29. Click OK to add this vNIC to the policy.

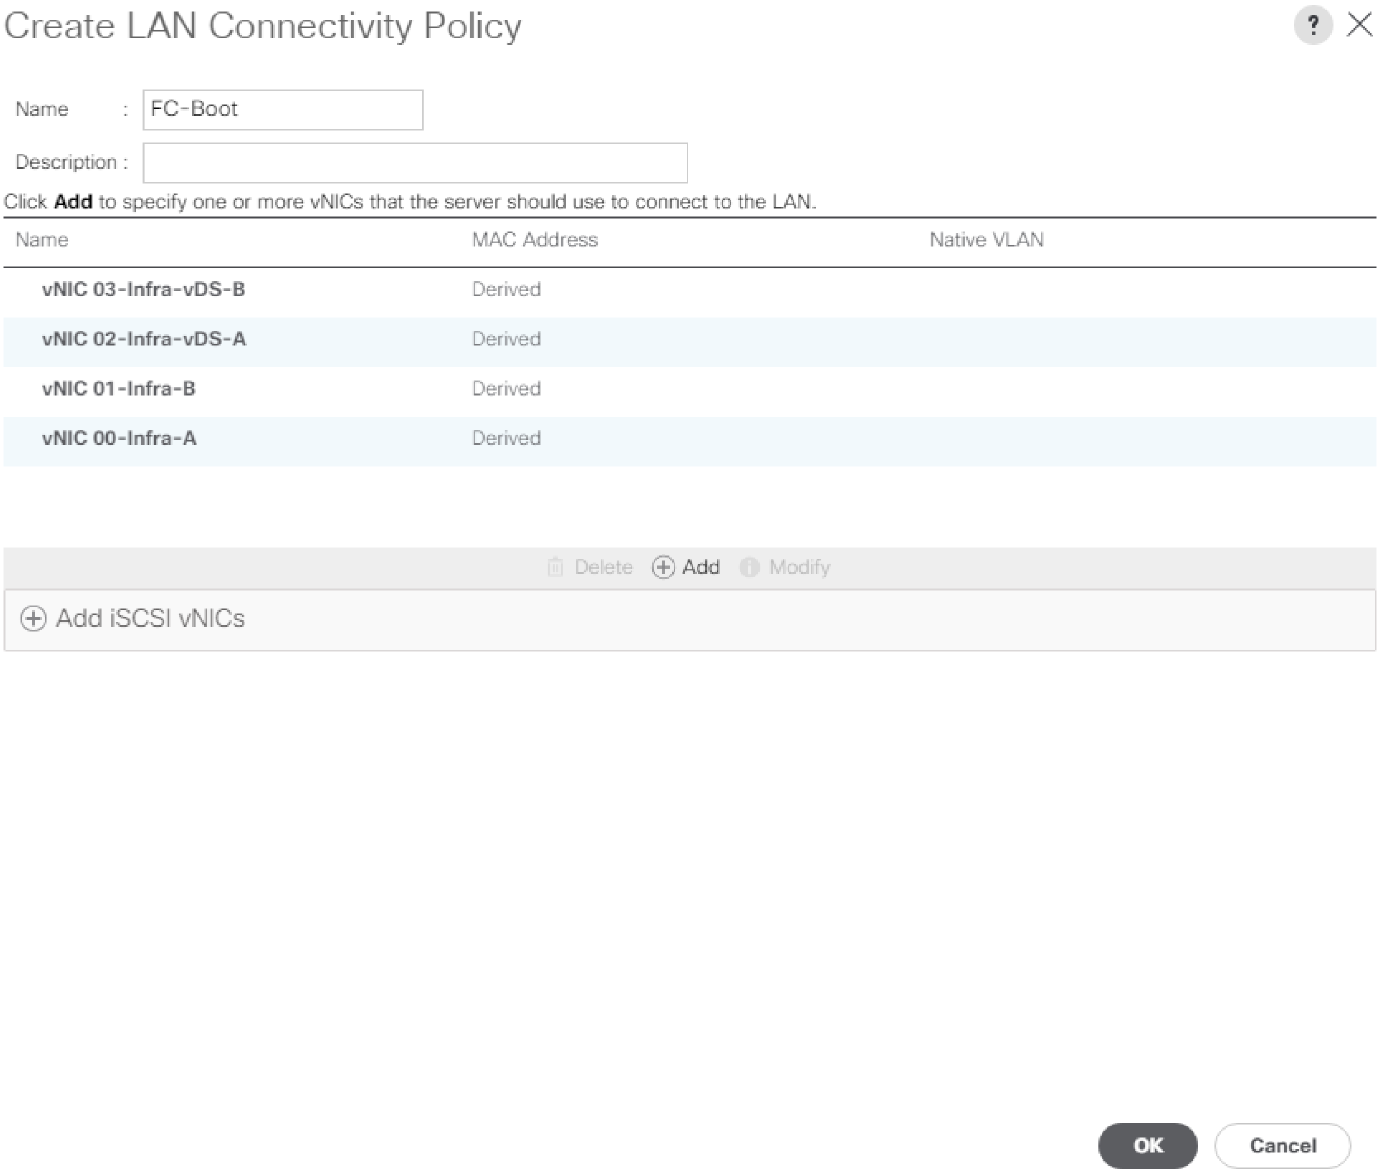
Create LAN Connectivity Policy

Name : FC-Boot

Description :

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Native VLAN
vNIC 03-Infra-vDS-B	Derived	
vNIC 02-Infra-vDS-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

30. Click OK, then click OK again to create the LAN Connectivity Policy.

Create vMedia Policy for VMware ESXi 6.7U1 ISO Install Boot

In the NetApp ONTAP setup steps, an HTTP web server is required, which is used for hosting ONTAP as well as VMware software. The vMedia Policy created will map the VMware ESXi 6.7U1 Cisco Custom ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, follow these steps:

1. In Cisco UCS Manager, select Servers .
2. Select Policies > root.
3. Right-click vMedia Policies.

4. Select Create vMedia Policy.

5. Name the policy ESXi-6.7U1-HTTP.

6. Enter “Mounts Cisco Custom ISO for ESXi 6.7U1” in the Description field.

7. Click Add.

8. Name the mount ESXi-6.7U1-HTTP.

9. Select the CDD Device Type.

10. Select the HTTP Protocol.

11. Enter the IP Address of the web server.

To avoid any DNS lookup issues, enter the IP of the web server instead of the hostname.

12. Enter VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso as the Remote File name.

This VMware ESXi 6.7U1 Cisco Custom ISO can be downloaded from VMware Downloads.

If a working vCenter 6.7U1 installation is already in your environment, a FlexPod custom ISO for installing ESXi 6.7U1 with all necessary drivers for this FlexPod deployment can be created. Please see the Appendix for a procedure for building this custom ISO.

13. Enter the web server path to the ISO file in the Remote Path field.

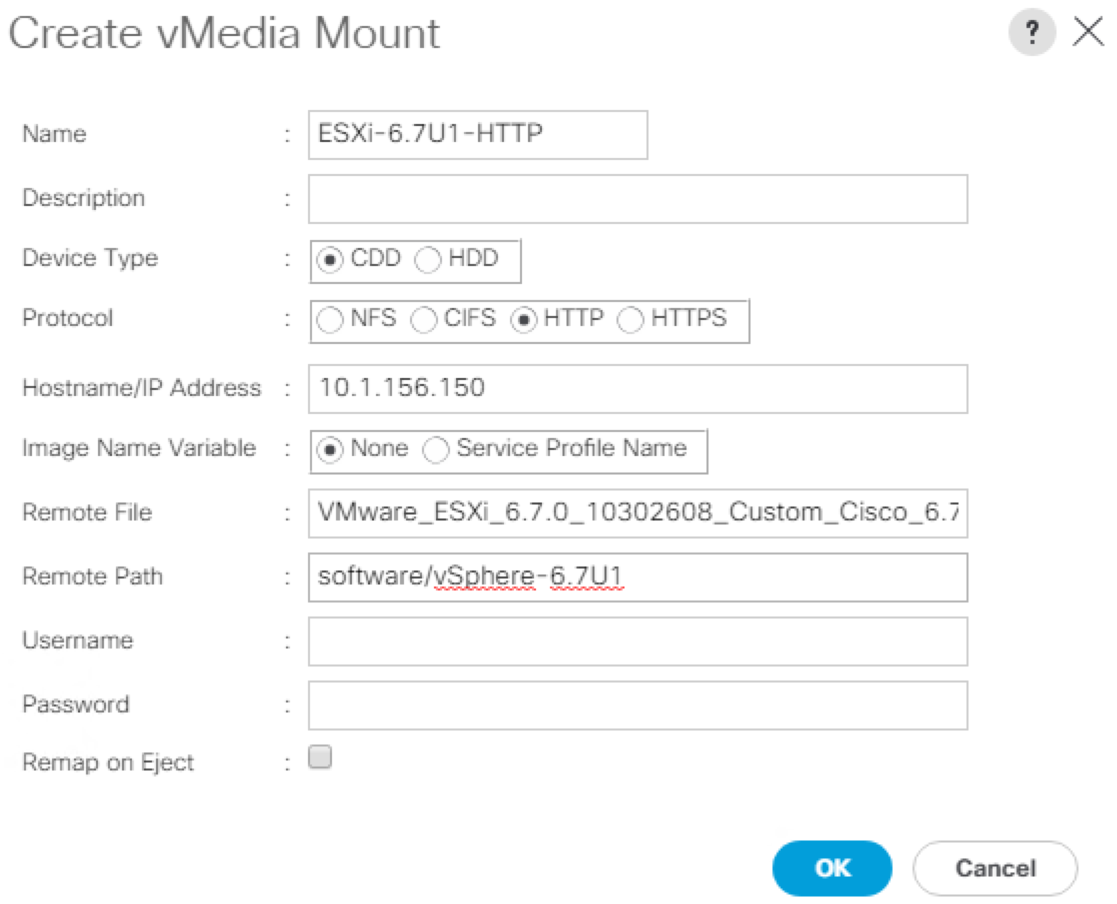
Create vMedia Mount

Name : ESXi-6.7U1-HTTP

Description :

Device Type : ● CDD ○ HDD

Protocol : ○ NFS ○ CIFS ● HTTP ○ HTTPS

Hostname/IP Address : 10.1.156.150

Image Name Variable : ● None ○ Service Profile Name

Remote File : VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7

Remote Path : software/vSphere-6.7U1

Username :

Password :

Remap on Eject : □

14. Click OK to create the vMedia Mount.

15. Click OK then click OK again to complete creating the vMedia Policy.

For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi host. On first boot the host will boot into the ESXi installer since the SAN mounted disk is empty. After ESXi is installed, the vMedia will not be referenced as long as the boot disk is accessible.

Create FC Boot Policy (FCP)

This procedure applies to a Cisco UCS environment in which two Fibre Channel logical interfaces (LIFs) are on cluster node 1 (fcp-lif01a and fcp-lif01b) and two Fibre Channel 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 UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).

One boot policy is configured in this procedure. The policy configures the primary target to be fcp-lif01a.

To create a boot policy for the within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click Servers .

2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Under the FlexPod Organization, right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter Boot-FCP-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Select the Uefi Boot Mode.
9. Select the Boot Security checkbox.

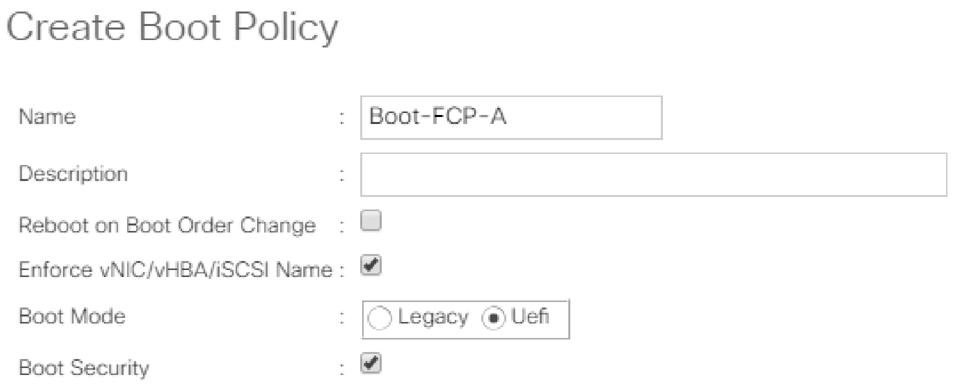

Create Boot Policy

Name	:	Boot-FCP-A
Description	:	
Reboot on Boot Order Change	:	□
Enforce vNIC/vHBA/iSCSI Name	:	☑
Boot Mode	:	○ Legacy ● Uefi
Boot Security	:	☑

UEFI Secure Boot can be used to boot VMware ESXi 6.7U1 with or without a TPM 2.0 module in the UCS server.

10. Expand Local Devices and select Add Remote CD/DVD.
11. Expand vHBAs and select Add SAN Boot.
12. Select Primary for the type field.
13. Enter FCP-Fabric-A in the vHBA field.

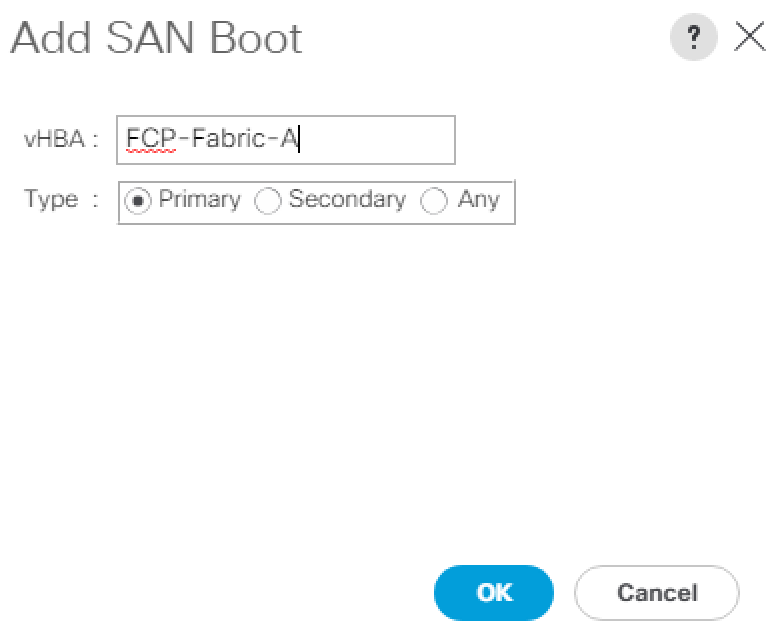

Add SAN Boot

vHBA : FCP-Fabric-A

Type : ● Primary ○ Secondary ○ Any

14. Click OK.

15. From vHBAs, select Add SAN Boot Target.

16. Keep 0 as the value for Boot Target LUN.

17. Enter the WWPN for fcp-lif01a.

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

18. Select Primary for the SAN boot target type.

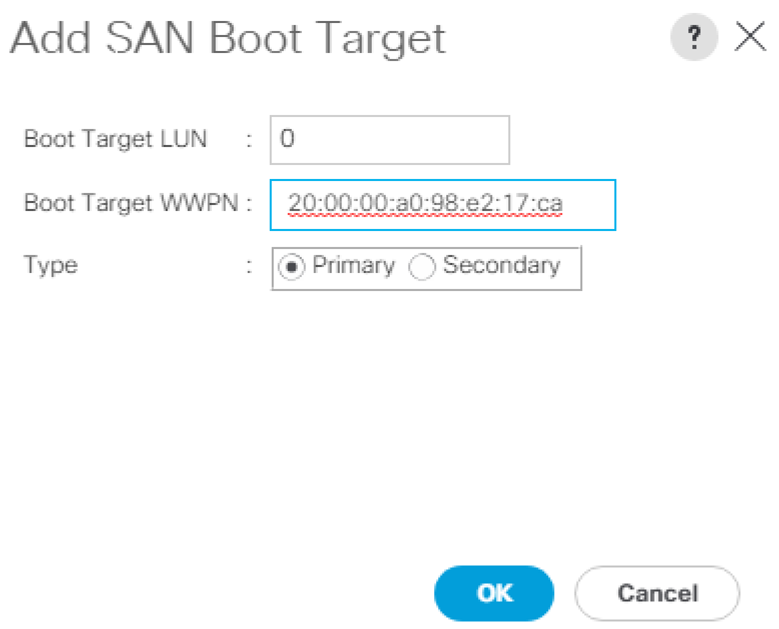

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:00:00:a0:98:e2:17:ca

Type : ● Primary ○ Secondary

19. Click OK to add the SAN boot target.

20. From vHBAs, select Add SAN Boot Target.

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

22. Enter the WWPN for fcp-lif02a.

23. Click OK to add the SAN boot target.

24. From vHBAs, select Add SAN Boot.

25. In the Add SAN Boot dialog box, enter FCP-Fabric-B in the vHBA box.

26. The SAN boot type should automatically be set to Secondary.

27. Click OK.

28. From vHBAs, select Add SAN Boot Target.

29. Keep 0 as the value for Boot Target LUN.

30. Enter the WWPN for fcp-lif01b.

31. Select Primary for the SAN boot target type.

32. Click OK to add the SAN boot target.

33. From vHBAs, select Add SAN Boot Target.

34. Keep 0 as the value for Boot Target LUN.

35. Enter the WWPN for fcp-lif02b.

36. Click OK to add the SAN boot target.

37. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

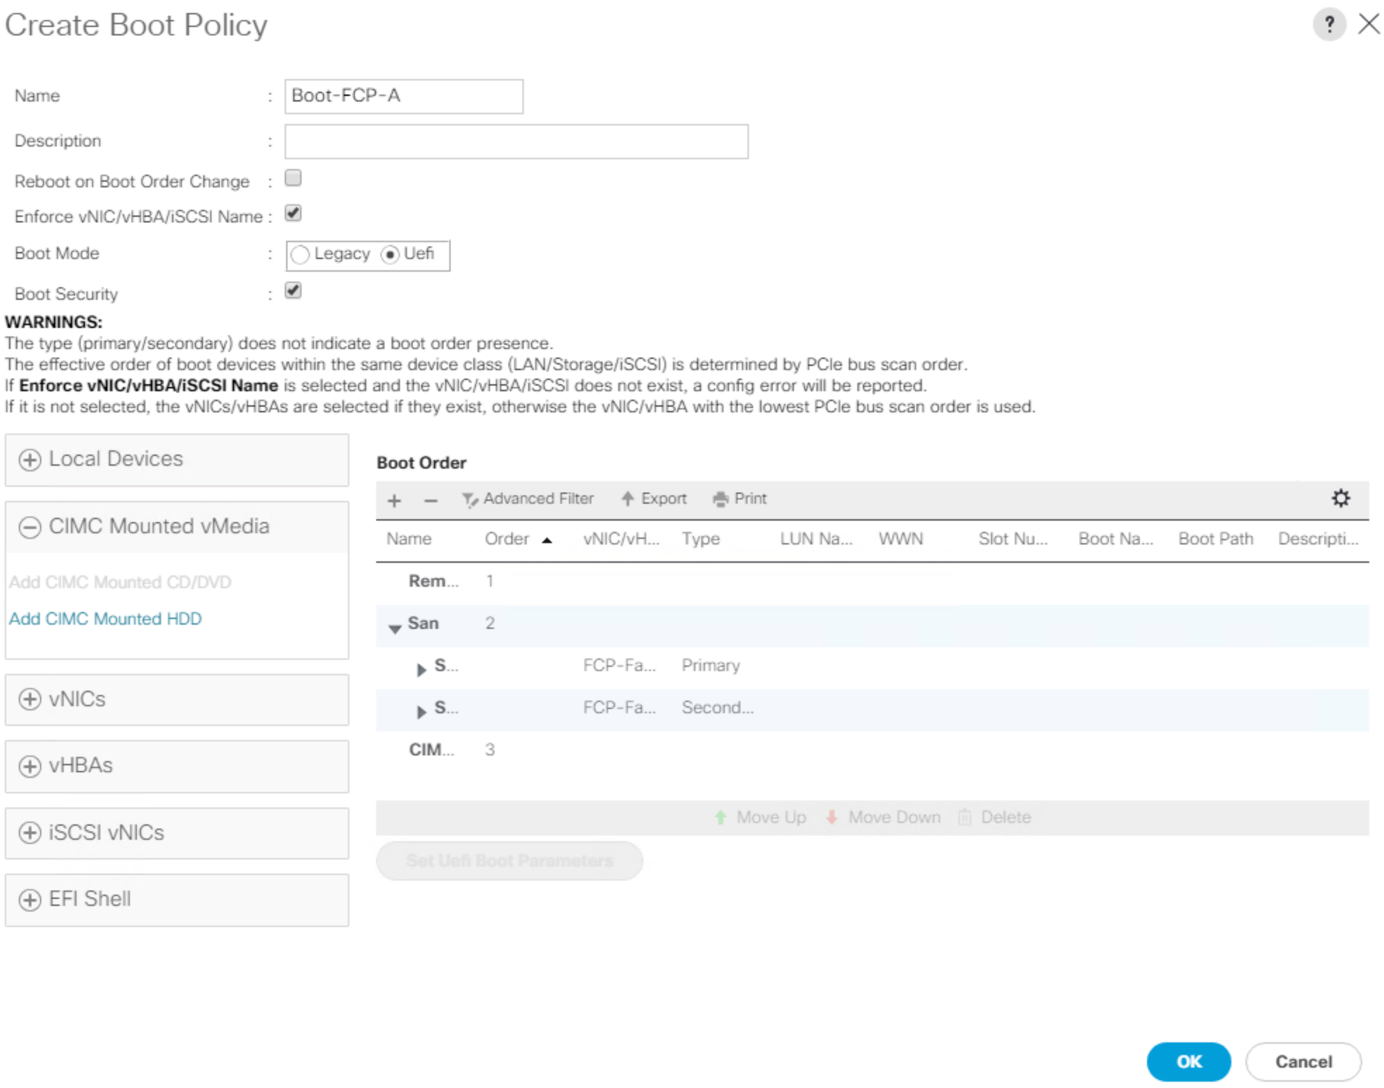

Create Boot Policy

Name : Boot-FCP-A

Description :

Reboot on Boot Order Change : □

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ○ Legacy ● Uefi

Boot Security : ☑

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
Rem...	1								
San	2								
S...		FCP-Fa...	Primary						
S...		FCP-Fa...	Second...						
CIM...	3								

38. Expand San > SAN Primary and select SAN Target Primary. Select Set Uefi Boot Parameters.

39. Fill in the Set Uefi Boot Parameters exactly as shown in the following screenshot:

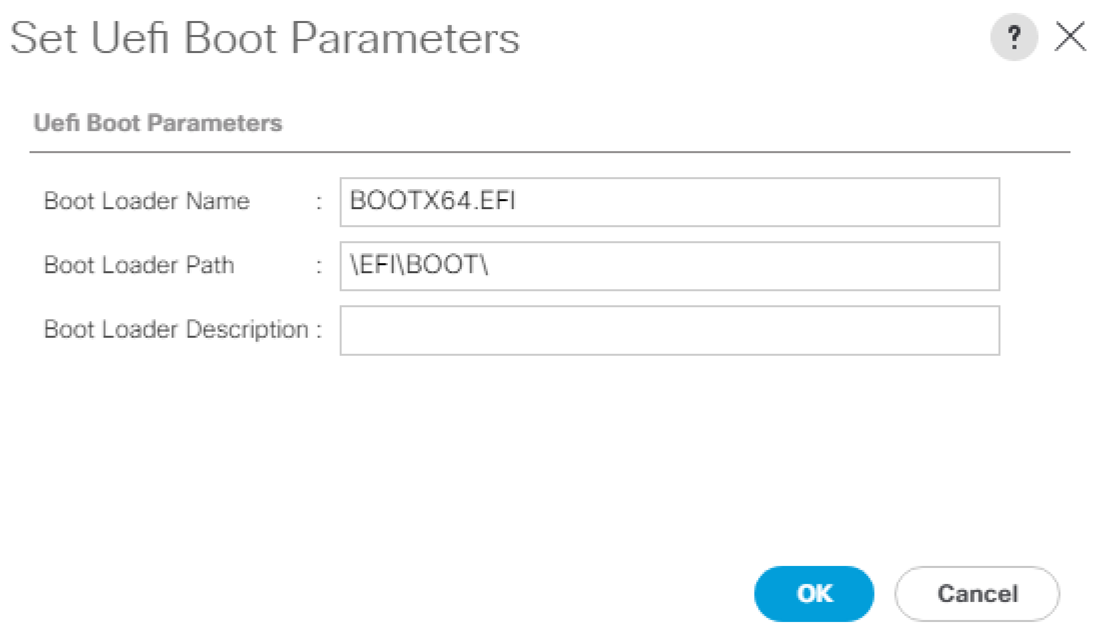

Set Uefi Boot Parameters

Uefi Boot Parameters

Boot Loader Name : BOOTX64.EFI

Boot Loader Path : \EFI\BOOT\

Boot Loader Description :

40. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target.

41. Repeat this process to set Uefi Boot Parameters for each of the 4 SAN Boot Targets.

For Cisco UCS B200 M5 and Cisco UCS C220 M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for M4 and earlier servers, VMware ESXi 6.7U1 will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

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

Create Service Profile Template (FCP)

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot within the FlexPod organization. To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Service Profile Templates > root > Sub-Organizations > FlexPod Organization.

3. Right-click the FlexPod Organization.

4. Select Create Service Profile Template to open the Create Service Profile Template wizard.

5. Enter VM-Host-Infra-FCP-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.

6. Select the “Updating Template” option.

7. Under UUID, select UUID_Pool as the UUID pool.

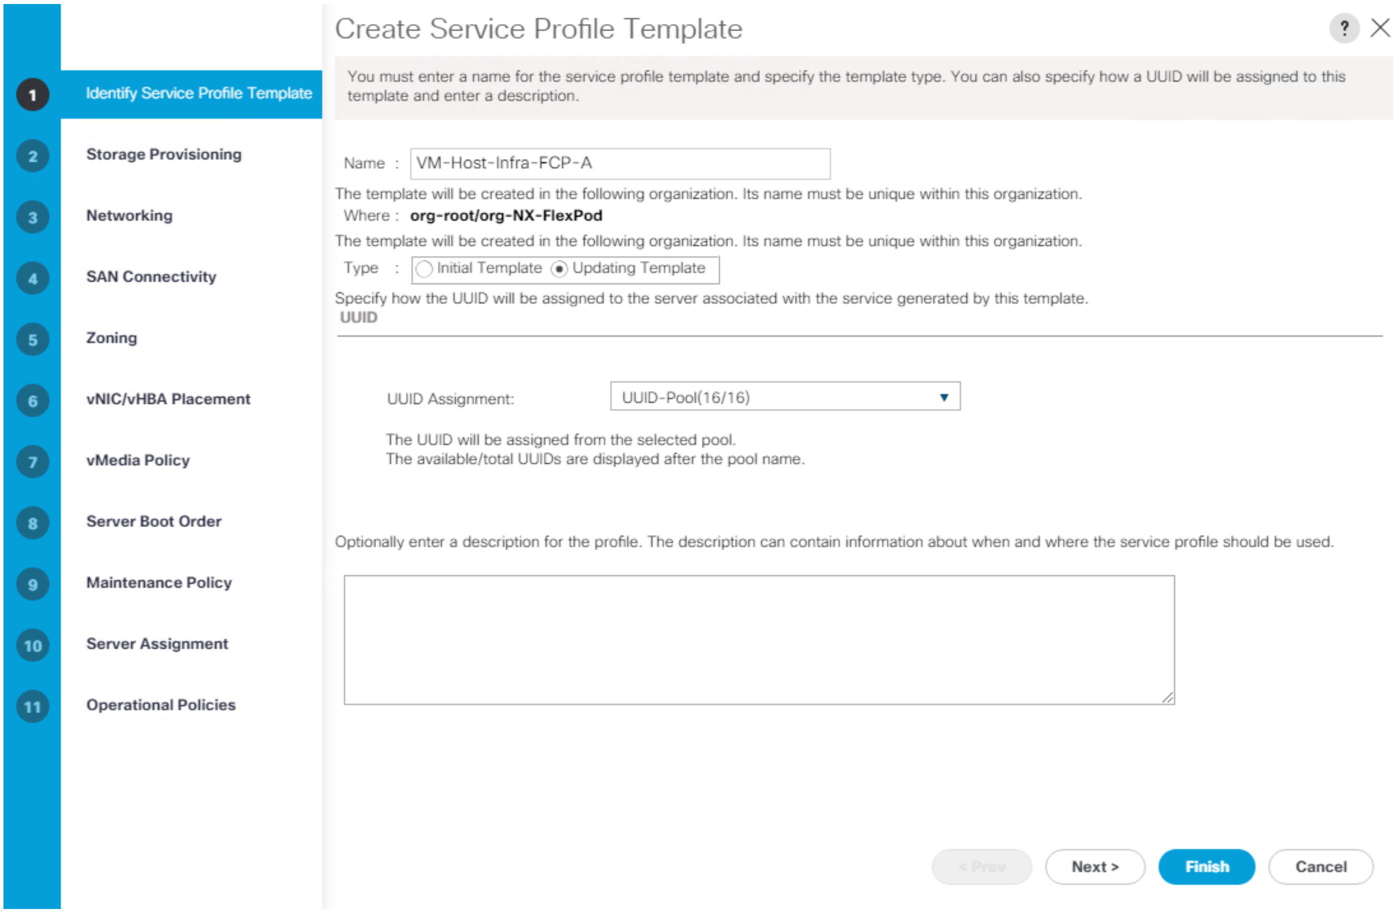

Create Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name : VM-Host-Infra-FCP-A

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-NX-FlexPod**

The template will be created in the following organization. Its name must be unique within this organization.

Type : ○ Initial Template ● Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

UUID

UUID Assignment: UUID-Pool(16/16)

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

8. Click Next.

Configure Storage Provisioning

To configure storage provisioning, follow these steps:

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy and select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.

2. Click Next.

Configure Networking

To configure networking, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.

2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.

3. Select FC-Boot from the LAN Connectivity Policy drop-down list.

4. Leave Initiator Name Assignment at <not set>.

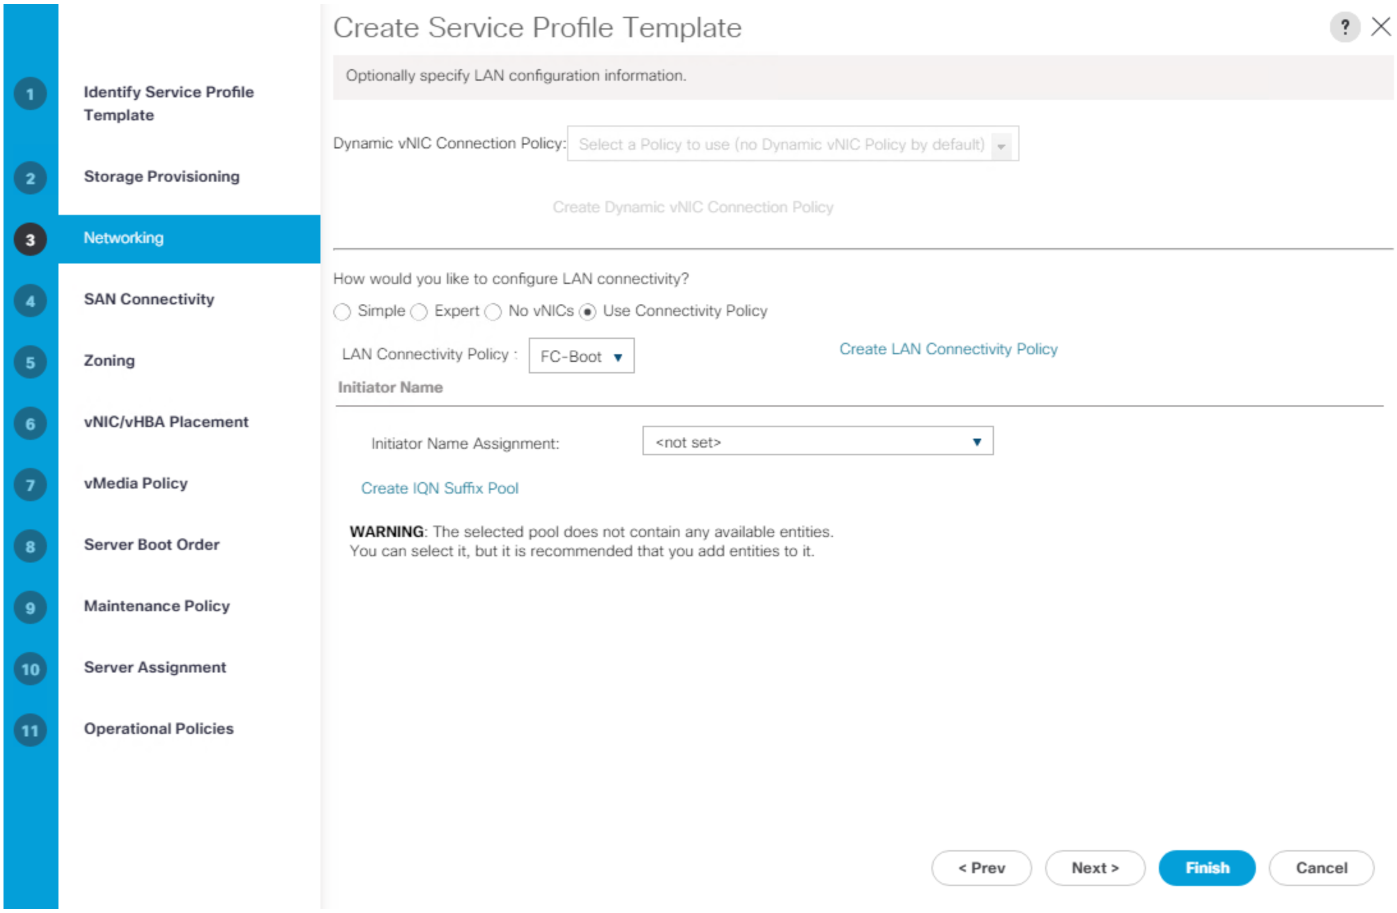

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default)

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

○ Simple ○ Expert ○ No vNICs ● Use Connectivity Policy

LAN Connectivity Policy : FC-Boot

Create LAN Connectivity Policy

Initiator Name

Initiator Name Assignment: <not set>

Create IQN Suffix Pool

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

5. Click Next.

Configure SAN Connectivity

To configure SAN connectivity, follow these steps:

1. Select the Use Connectivity Policy option for the “How would you like to configure SAN connectivity?” field.

2. Select the FC-Boot option from the SAN Connectivity Policy drop-down list.

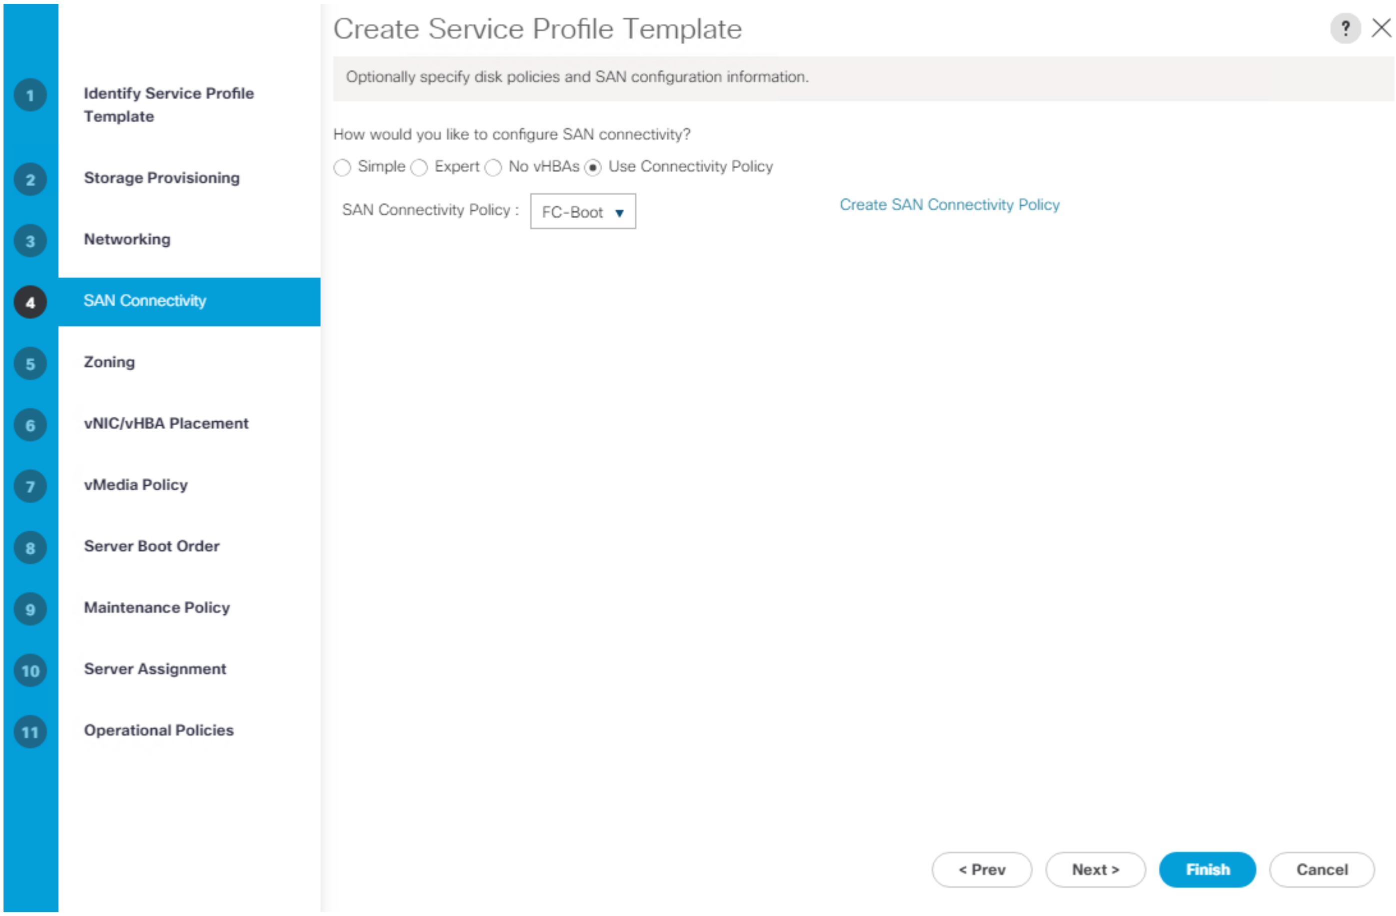

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

○ Simple ○ Expert ○ No vHBAs ● Use Connectivity Policy

SAN Connectivity Policy : FC-Boot

Create SAN Connectivity Policy

3. Click Next.

Configure Zoning

To configure zoning, follow this step:

1. Set no zoning options and click Next.

Configure vNIC/HBA Placement

To configure vNIC/HBA placement, follow these steps:

1. In the Select Placement list, retain the placement policy as Let System Perform Placement.
2. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

1. Do not select a vMedia Policy.
2. Click Next.

Configure Server Boot Order

To configure the server boot order, follow these steps:

1. Select Boot-FCP-A for Boot Policy.

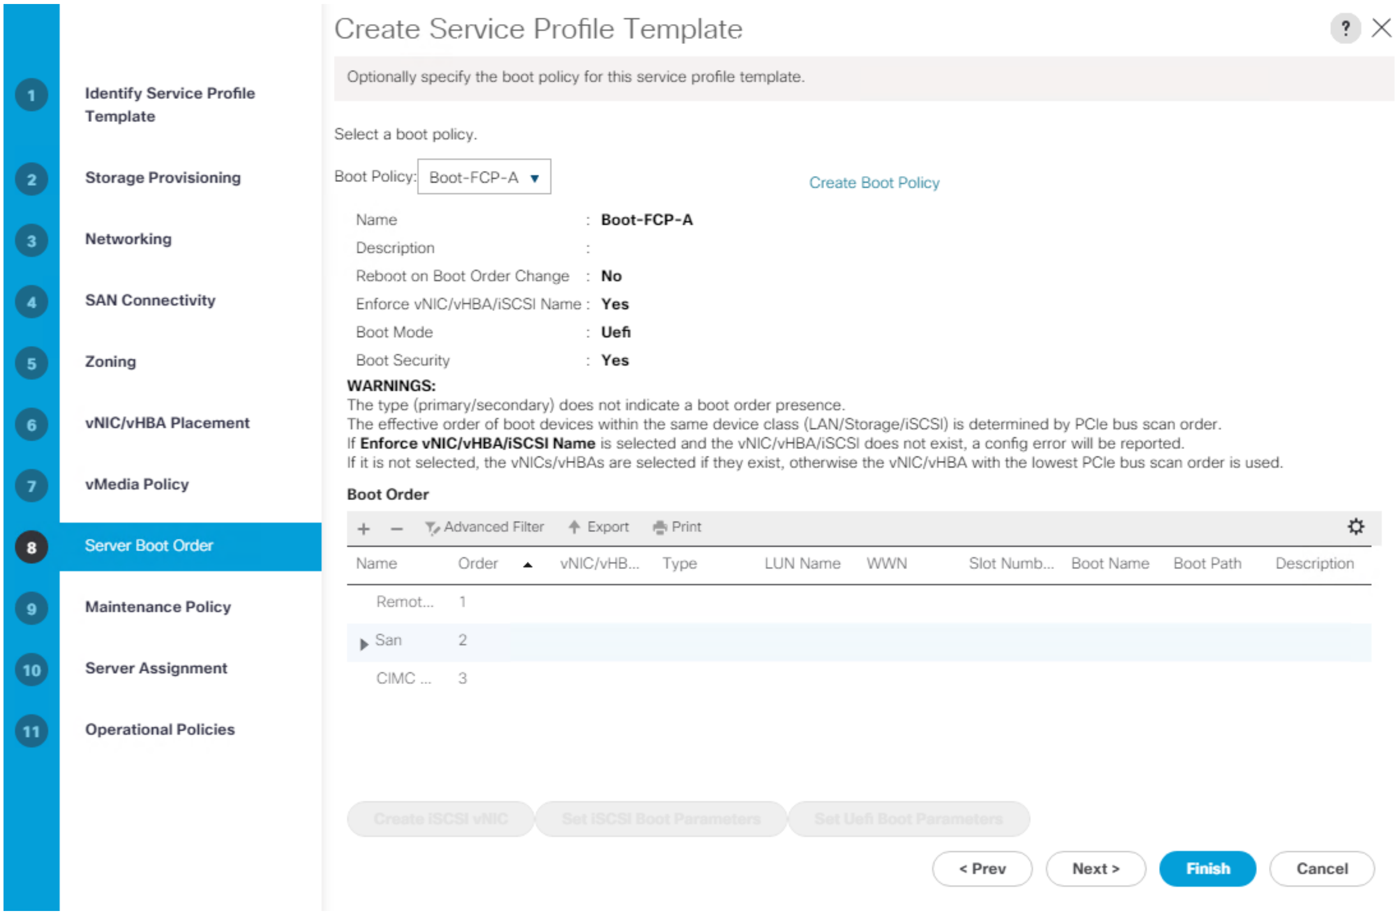

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Boot-FCP-A ▾

Name	: **Boot-FCP-A**
Description	:
Reboot on Boot Order Change	: **No**
Enforce vNIC/vHBA/iSCSI Name	: **Yes**
Boot Mode	: **Uefi**
Boot Security	: **Yes**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order ▲	vNIC/vHB...	Type	LUN Name	WWN	Slot Numb...	Boot Name	Boot Path	Description
Remot...	1								
▸ San	2								
CIMC ...	3								

2. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.

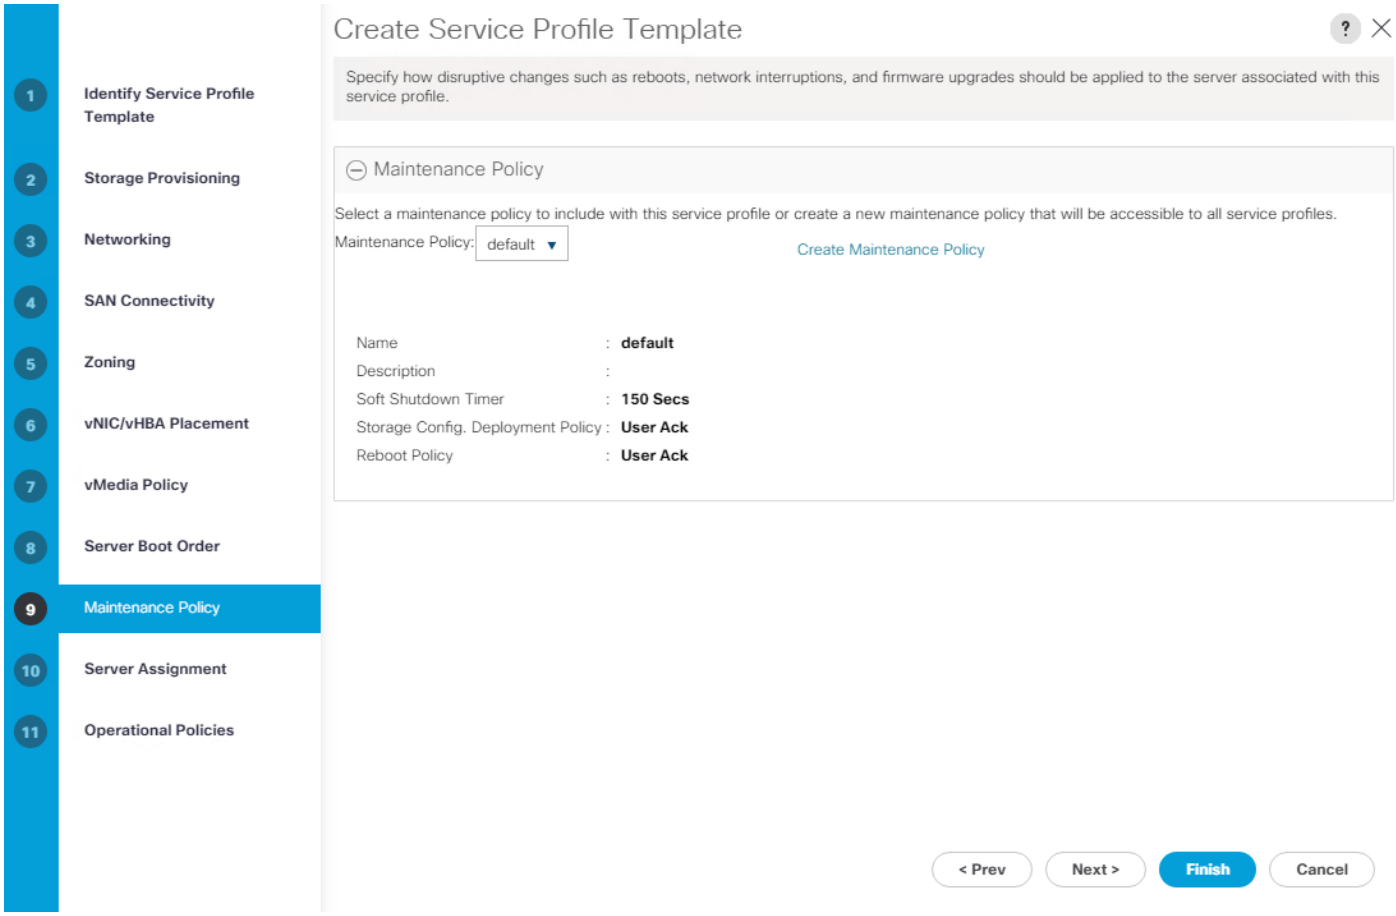

Create Service Profile Template

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: default

Name	:	default
Description	:	
Soft Shutdown Timer	:	150 Secs
Storage Config. Deployment Policy	:	User Ack
Reboot Policy	:	User Ack

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, select Infra-Pool.

2. Select Down as the power state to be applied when the profile is associated with the server.

3. Optional: select “B200-M5” for the Server Pool Qualification to select only B200 M5 servers in the pool.

4. Expand Firmware Management and select the default Host Firmware Package.

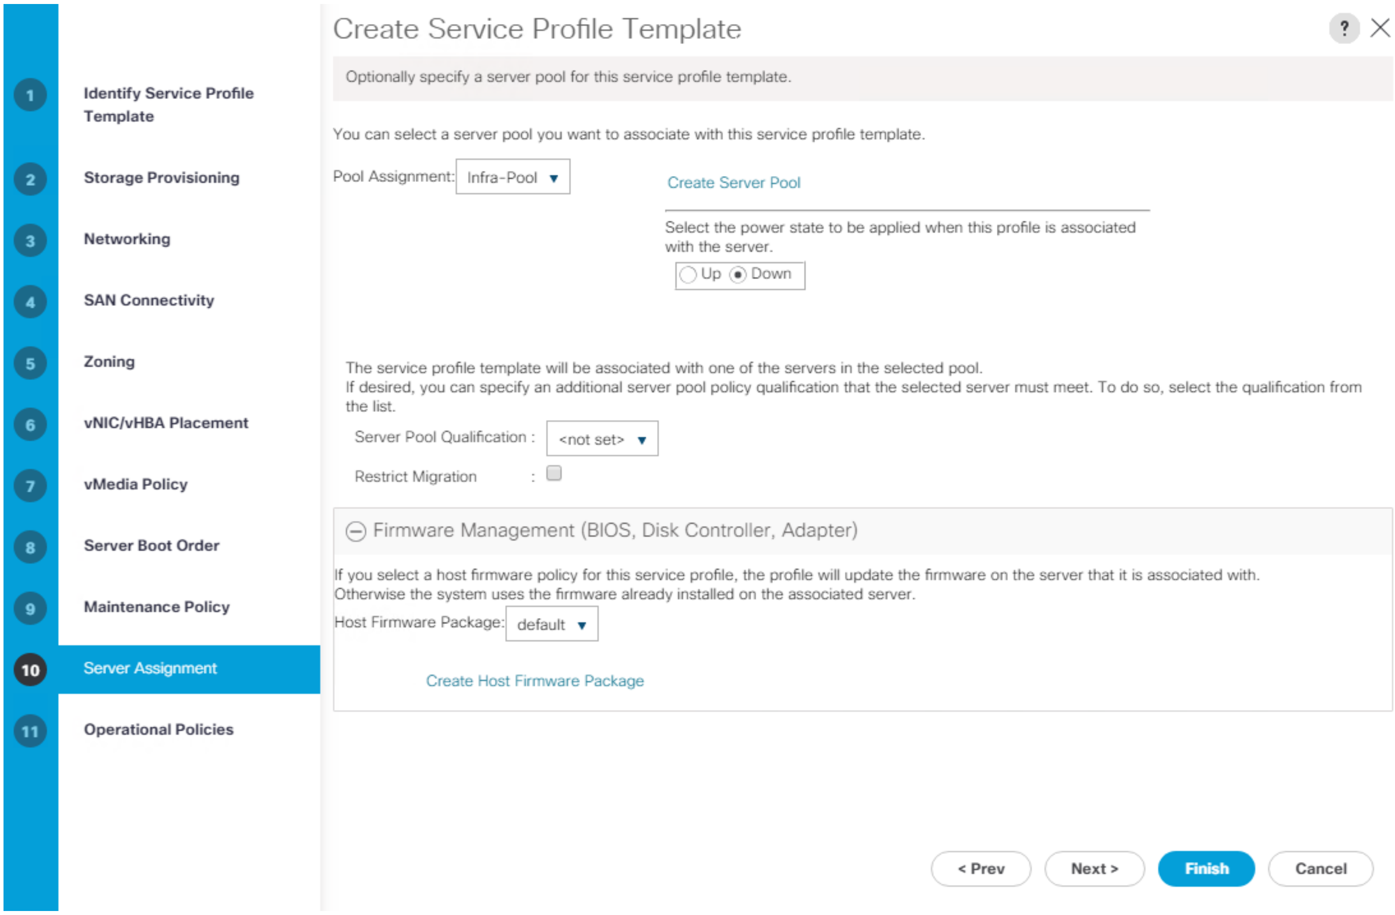

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: Infra-Pool

Select the power state to be applied when this profile is associated with the server.

○ Up ● Down

The service profile template will be associated with one of the servers in the selected pool.
If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification : <not set>

Restrict Migration : □

⊖ Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with.
Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: default

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, select VM-Host.

2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

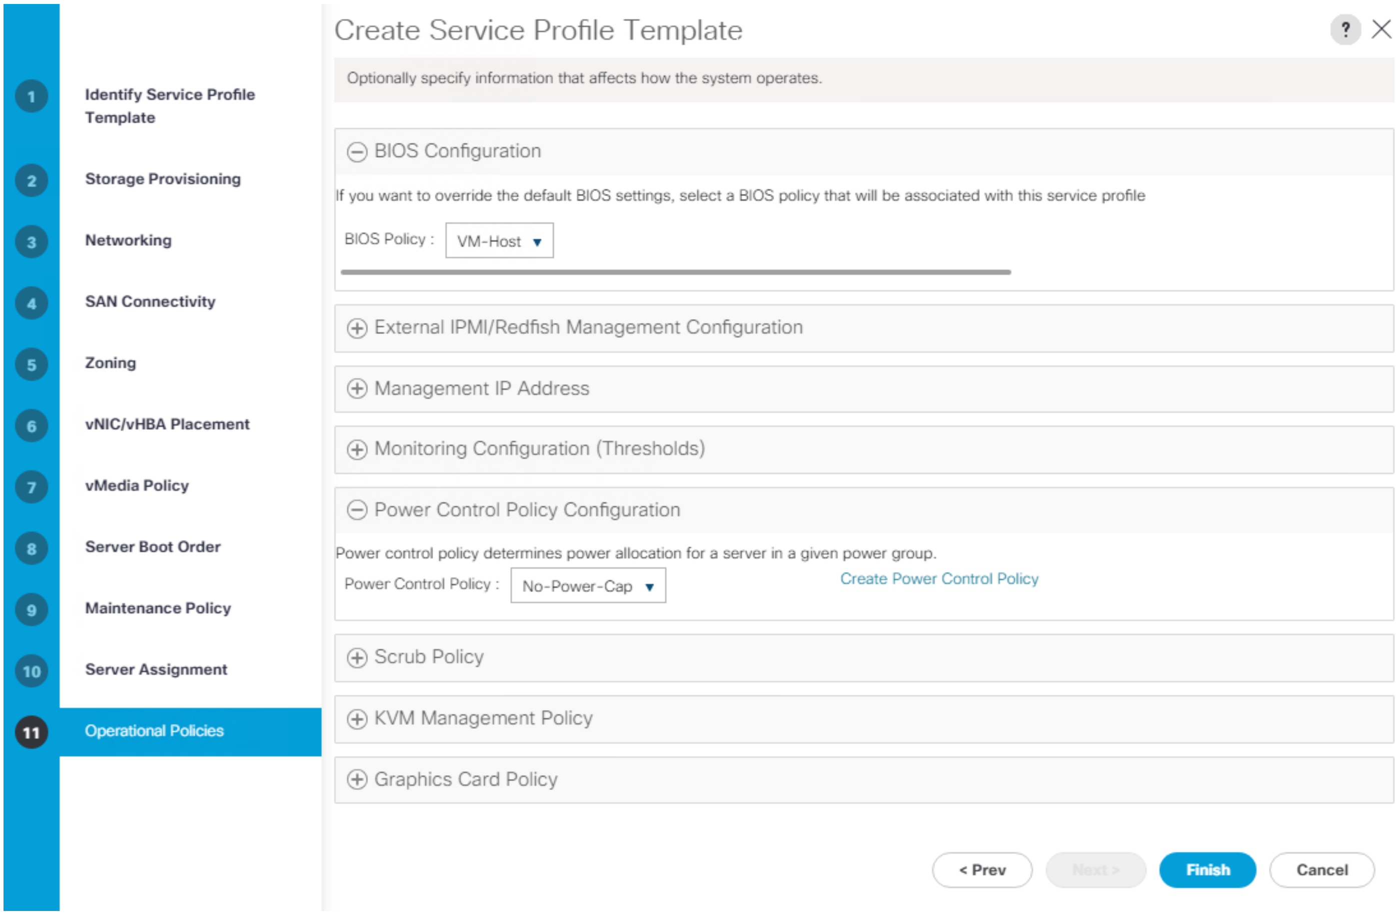

Create Service Profile Template

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy : VM-Host

Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : No-Power-Cap

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create vMedia-Enabled Service Profile Template

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers .

2. Select Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template VM-Host-Infra-FCP-A.

3. Right-click VM-Host-Infra-FCP-A and select Create a Clone.

4. Name the clone VM-Host-Infra-FC-A-vM.

5. Click OK then OK again to create the Service Profile Template clone.

6. Select the newly-created VM-Host-Infra-FCP-A-vM and select the vMedia Policy tab.

7. Click Modify vMedia Policy.

8. Select the ESXi-6.7U1-HTTP vMedia Policy and click OK.

9. Click OK to confirm.

Create Service Profiles

To create service profiles from the service profile template within the FlexPod organization, follow these steps:

1. Connect to UCS Manager and click Servers .

2. Select Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template VM-Host-Infra-FCP-A-vM.

3. Right-click VM-Host-Infra-FCP-A-vM and select Create Service Profiles from Template.

4. Enter VM-Host-Infra-0 as the service profile prefix.

5. Enter 1 as "Name Suffix Starting Number."

6. Enter 2 as the "Number of Instances."

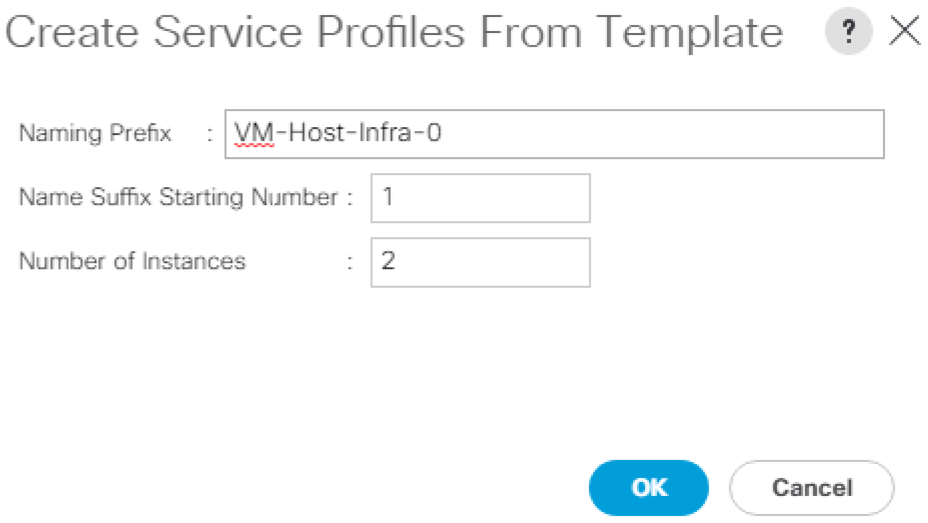

Create Service Profiles From Template

Naming Prefix : VM-Host-Infra-0

Name Suffix Starting Number : 1

Number of Instances : 2

7. Click OK to create the service profiles.

8. Click OK in the confirmation message.

9. When VMware ESXi 6.7U1 has been installed on the hosts, the host Service Profiles can be bound to the VM-Host-Infra-FCP-A Service Profile Template to remove the vMedia Mapping from the host.

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 server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 6 and Table 7 .

Table 6 WWPNs from NetApp Storage

SVM	Adapter	MDS Switch	Target: WWPN
Infra-SVM	fcp-lif01a	Fabric A	`<fcp-lif01a-wwpn>`
	fcp-lif01b	Fabric B	`<fcp-lif01b-wwpn>`
	fcp-lif02a	Fabric A	`<fcp-lif02a-wwpn>`
	fcp-lif02b	Fabric B	`<fcp-lif02b-wwpn>`

To obtain the FC WWPNs, run the `network interface show` command on the storage cluster management interface.

Table 7 WWPNs for Cisco UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
VM-Host-Infra-01	Fabric A	`<vm-host-infra-01-wwpna>`
	Fabric B	`<vm-host-infra-01-wwpnb>`
VM-Host-Infra-02	Fabric A	`<vm-host-infra-02-wwpna>`
	Fabric B	`<vm-host-infra-02-wwpnb>`

To obtain the FC vHBA WWPN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root > Sub-Organizations > Organization. Click each service profile and then click the “Storage” tab, then “vHBAs” tab on the right. The WWPNs are displayed in the table at the bottom of the page.

SAN Switch Configuration

This section explains how to configure the Cisco MDS 9000s for use in a FlexPod environment. Follow the steps precisely because failure to do so could result in an improper configuration.

If directly connecting storage to the Cisco UCS fabric interconnects, skip this section.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in the section FlexPod Cabling.

FlexPod Cisco MDS Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes you are using the Cisco MDS 9132T with NX-OS 8.3(2).

Cisco MDS 9132T A

To set up the initial configuration for the Cisco MDS A switch, <mds-A-hostname>, follow these steps:

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. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

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


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

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 : <mds-A-hostname>

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

Mgmt0 IPv4 address : <mds-A-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-A-mgmt0-netmask>

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

IPv4 address of the default gateway : <mds-A-mgmt0-gw>

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

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

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

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

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

Configure congestion/no_credit drop for fc interfaces? (yes/no)     [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500.  [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter

Configure timezone? (yes/no) [n]: Enter

Configure summertime? (yes/no) [n]: Enter

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

NTP server IPv4 address : <nexus-A-mgmt0-ip>

Configure default switchport interface state (shut/noshut) [shut]: Enter

Configure default switchport trunk mode (on/off/auto) [on]: auto

Configure default switchport port mode F (yes/no) [n]: yes

Configure default zone policy (permit/deny) [deny]: Enter

Enable full zoneset distribution? (yes/no) [n]: Enter

Configure default zone mode (basic/enhanced) [basic]: Enter
```

2. Review the configuration.

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

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

Cisco MDS 9132T B

To set up the initial configuration for the Cisco MDS B switch, <mds-B-hostname>, follow these steps:

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. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

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


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

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 : <mds-B-hostname>

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

Mgmt0 IPv4 address : <mds-B-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-B-mgmt0-netmask>

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

IPv4 address of the default gateway : <mds-B-mgmt0-gw>

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

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

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

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

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

Configure congestion/no_credit drop for fc interfaces? (yes/no)     [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500.  [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter

Configure timezone? (yes/no) [n]: Enter

Configure summertime? (yes/no) [n]: Enter

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

NTP server IPv4 address : <nexus-A-mgmt0-ip>

Configure default switchport interface state (shut/noshut) [shut]: Enter

Configure default switchport trunk mode (on/off/auto) [on]: auto

Configure default switchport port mode F (yes/no) [n]: yes

Configure default zone policy (permit/deny) [deny]: Enter

Enable full zoneset distribution? (yes/no) [n]: Enter

Configure default zone mode (basic/enhanced) [basic]: Enter
```

2. Review the configuration.

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

FlexPod Cisco MDS Switch Configuration

Enable Licenses

Cisco MDS 9132T A and Cisco MDS 9132T B

To enable the correct features on the Cisco MDS switches, follow these steps:

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

```
configure terminal
feature npiv
feature fport-channel-trunk
```

Add Second NTP Server

Cisco MDS 9132T A and Cisco MDS 9132T B

To configure the second NTP server, complete the following step:

From the global configuration mode, run the following command:

```
ntp server <nexus-B-mgmt0-ip>
```

Configure Individual Ports

Cisco MDS 9132T A

To configure individual ports and port-channels for switch A, complete the following step:

From the global configuration mode, run the following commands:

```
interface fc1/1
switchport description <st-node01>:2a
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description <st-node02>:2a
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description <ucs-clustername>-a:1/1
channel-group 15
no shutdown
exit

interface fc1/6
switchport description <ucs-clustername>-a:1/2
channel-group 15
no shutdown
exit
```

```
interface port-channel15
channel mode active
switchport trunk allowed vsan <vsan-a-id>
switchport description <ucs-clustername>-a
switchport speed 32000
no shutdown
exit
```

Cisco MDS 9132T B

To configure individual ports and port-channels for switch B, complete the following step:

From the global configuration mode, run the following commands:

```
interface fc1/1
switchport description <st-node01>:2b
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description <st-node02>:2b
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description <ucs-clustername>-b:1/1
channel-group 15
no shutdown
exit

interface fc1/6
switchport description <ucs-clustername>-b:1/2
channel-group 15
no shutdown
exit

interface port-channel15
channel mode active
switchport trunk allowed vsan <vsan-b-id>
switchport description <ucs-clustername>-b
switchport speed 32000
no shutdown
exit
```

Create VSANs

Cisco MDS 9132T A

To create the necessary VSANs for fabric A and add ports to them, follow these steps:

From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fc1/1
vsan <vsan-a-id> interface fc1/2
vsan <vsan-a-id> interface port-channel15
exit
```

Cisco MDS 9132T B

To create the necessary VSANs for fabric B and add ports to them, follow these steps:

From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fc1/1
vsan <vsan-b-id> interface fc1/2
vsan <vsan-b-id> interface port-channel15
exit
```

At this point, it may be necessary to go into UCS Manager and disable and enable the FC port-channel interfaces to get the port-channels to come up.

Create Device Aliases

Cisco MDS 9132T A

To create device aliases for Fabric A that will be used to create zones, follow these steps:

From the global configuration mode, run the following commands:

```
device-alias database
device-alias name Infra-SVM-fcp-lif01a pwwn <fcp-lif01a-wwpn>
device-alias name Infra-SVM-fcp-lif02a pwwn <fcp-lif02a-wwpn>
device-alias name VM-Host-Infra-01-A pwwn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwwn <vm-host-infra-02-wwpna>
device-alias commit
```

Cisco MDS 9132T B

To create device aliases for Fabric B that will be used to create zones, follow these steps:

From the global configuration mode, run the following commands:

```
device-alias database
device-alias name Infra-SVM-fcp-lif01b pwwn <fcp-lif01b-wwpn>
device-alias name Infra-SVM-fcp-lif02b pwwn <fcp-lif02b-wwpn>
device-alias name VM-Host-Infra-01-B pwwn <vm-host-infra-01-wwpnb>
device-alias name VM-Host-Infra-02-B pwwn <vm-host-infra-02-wwpnb>
device-alias commit
```

Create Zones and Zoneset

Cisco MDS 9132T A

To create the required zones and zoneset on Fabric A, run the following commands:

```
configure terminal
zone name FlexPod-Boot-Fabric-A vsan <vsan-a-id>
```

```
member device-alias VM-Host-Infra-01-A init
member device-alias VM-Host-Infra-02-A init
member device-alias Infra-SVM-fcp-lif01a target
member device-alias Infra-SVM-fcp-lif02a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member FlexPod-Boot-Fabric-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
show zoneset active
copy r s
```

Since Smart Zoning is enabled, a single zone is created with all host boot initiators and boot targets for the Infra-SVM instead of creating a separate zone for each host with the host initiator and boot targets. If a new host is added, its boot initiator can simply be added to the single zone in each MDS switch and then the zoneset reactivated. If another SVM is added to the FlexPod with FC targets, a new zone can be added for that SVM.

Cisco MDS 9132T B

To create the required zones and zoneset on Fabric B, run the following commands:

```
configure terminal
zone name FlexPod-Boot-Fabric-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-01-B init
member device-alias VM-Host-Infra-02-B init
member device-alias Infra-SVM-fcp-lif01b target
member device-alias Infra-SVM-fcp-lif02b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member FlexPod-Boot-Fabric-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active
copy r s
```

Storage Configuration – Boot LUNs

ONTAP Boot Storage Setup

Create igroups

Create igroups by entering the following commands from the storage cluster management node SSH connection:

```
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-01 -protocol fcp -ostype vmware -initiator
<vm-host-infra-01-wwpna>, <vm-host-infra-01-wwpnb>

lun igroup create -vserver Infra-SVM -igroup vm-host-infra-02 -protocol fcp -ostype vmware -initiator
<vm-host-infra-02-wwpna>, <vm-host-infra-02-wwpnb>

lun igroup create -vserver Infra-SVM -igroup mgmt-hosts -protocol fcp -ostype vmware -initiator <vm-
host-infra-01-wwpna>, <vm-host-infra-01-wwpnb>, <vm-host-infra-02-wwpna>, <vm-host-infra-02-wwpnb>
```

Use the values listed in Table 6 and Table 7 for the WWPN information.

To view the three igroups just created, use the command lun igroup show.

```
lun igroup show -protocol fcp
```

Map Boot LUNs to igroups

From the storage cluster management SSH connection, enter the following commands:

```
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/vm-host-infra-01 -igroup vm-host-infra-01
-lun-id 0

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/vm-host-infra-02 -igroup vm-host-infra-02
-lun-id 0
```

VMware vSphere 6.7U1 Setup

VMware ESXi 6.7U1

This section provides detailed instructions for installing VMware ESXi 6.7U1 in a FlexPod environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

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

Download ESXi 6.7U1 from VMware

If the VMware ESXi ISO has not been downloaded, follow these steps:

1. Click the following link: Cisco Custom ESXi 6.7U1 ISO.

2. You will need a user id and password on vmware.com to download this software.

3. Download the .iso file.

Log into Cisco UCS 6454 Fabric Interconnect

Cisco UCS Manager

The Cisco UCS 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 into the Cisco UCS environment, follow these steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.

2. Click the Launch UCS Manager link to launch the HTML 5 UCS Manager GUI.

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 Servers .

7. Select Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-01.

8. In the Actions pane, click >> to the right of KVM Console.

9. Select the “Launch Java KVM Console” checkbox and click OK to launch the console.

10. Follow the prompts to launch the Java-based KVM console.

11. Select Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-02.
12. In the Actions pane, click >> to the right of KVM Console.
13. Select the "Launch Java KVM Console" checkbox and click OK to launch the console.
14. Follow the prompts to launch the Java-based KVM console.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

Skip this section if using vMedia policies. ISO file will already be connected to KVM.

To prepare the server for the OS installation, follow these steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Select Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM Console tab to monitor the server boot.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware ESXi to the bootable LUN of the hosts, follow these steps on each host:

1. Boot the server by selecting Boot Server in the KVM and click OK, then click OK again.
2. On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

It may be necessary to map function keys as User Defined Macros under the Macros menu in the UCS KVM console.

5. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

6. Select the appropriate keyboard layout and press Enter.

7. Enter and confirm the root password and press Enter.

8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.

9. After the installation is complete, press Enter to reboot the server.

The ESXi installation image will be automatically unmapped in the KVM when Enter is pressed.

10. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, follow these steps on each ESXi host:

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To configure each ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, press F2 to customize VMware ESXi.

2. Log in as root, enter the corresponding password, and press Enter to log in.

3. Select Troubleshooting Options and press Enter.

4. Select Enable ESXi Shell and press Enter.

5. Select Enable SSH and press Enter.

6. Press Esc to exit the Troubleshooting Options menu.

7. Select the Configure Management Network option and press Enter.

8. Select Network Adapters and press Enter.

9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field.

10. Using the spacebar, select vmnic1.

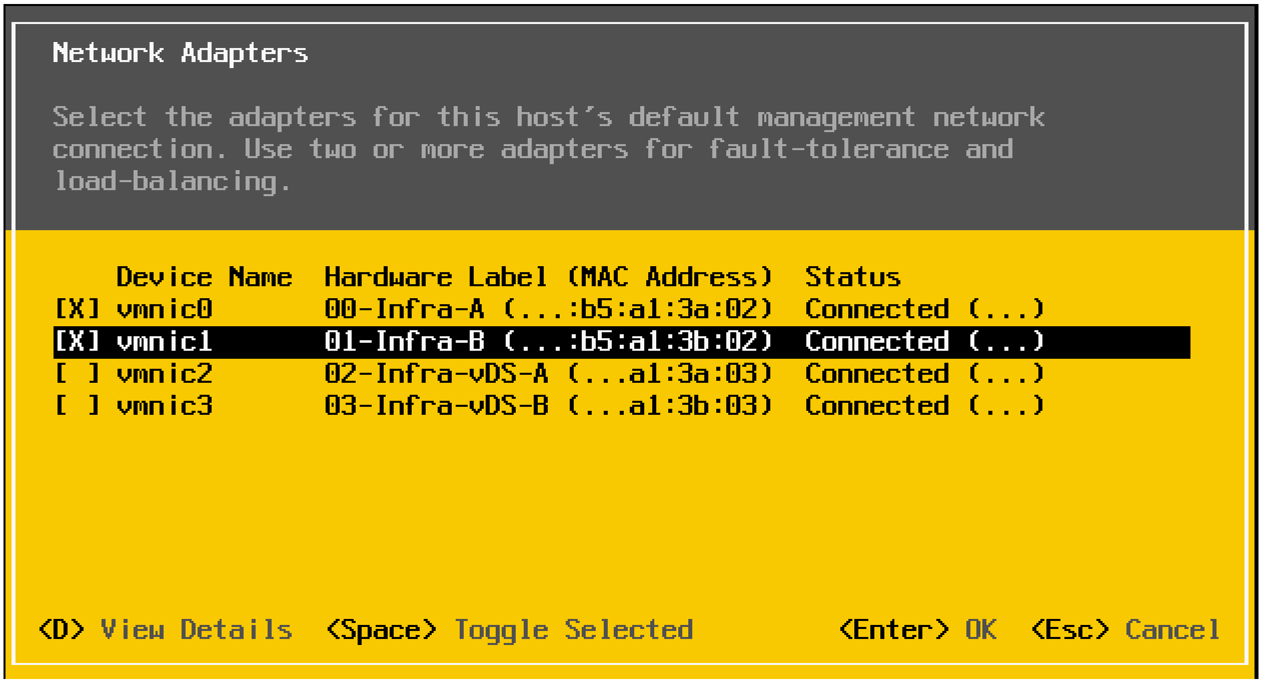

```
Network Adapters

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

    Device Name  Hardware Label (MAC Address)  Status
[X] vmnic0       00-Infra-A (...:b5:a1:3a:02)  Connected (...)
[X] vmnic1       01-Infra-B (...:b5:a1:3b:02)  Connected (...)
[ ] vmnic2       02-Infra-vDS-A (...a1:3a:03)  Connected (...)
[ ] vmnic3       03-Infra-vDS-B (...a1:3b:03)  Connected (...)
```

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

In lab testing, examples have been seen where the vmnic and device ordering do not match. If this is the case, use the Consistent Device Naming (CDN) to note which vmnics are mapped to which vNICs and adjust the upcoming procedure accordingly.

11. Press Enter.

12. Select the VLAN (Optional) option and press Enter.

13. Enter the <ib-mgmt-vlan-id> and press Enter.

14. Select IPv4 Configuration and press Enter.

15. Select the “Set static IPv4 address and network configuration” option by using the arrow keys and space bar.

16. Move to the IPv4 Address field and enter the IP address for managing the ESXi host.

17. Move to the Subnet Mask field and enter the subnet mask for the ESXi host.

18. Move to the Default Gateway field and enter the default gateway for the ESXi host.

19. Press Enter to accept the changes to the IP configuration.

20. Select the IPv6 Configuration option and press Enter.

21. Using the spacebar, select Disable IPv6 (restart required) and press Enter.

22. Select the DNS Configuration option and press Enter.

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

23. Using the spacebar, select “Use the following DNS server addresses and hostname:”

24. Move to the Primary DNS Server field and enter the IP address of the primary DNS server.

25. Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.

26. Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.

27. Press Enter to accept the changes to the DNS configuration.

28. Press Esc to exit the Configure Management Network submenu.

29. Press Y to confirm the changes and reboot the ESXi host.

Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional)

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port it is placed on. If the ESXi host’s boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will exist because vmk0 will retain the assigned MAC address unless the ESXi System Configuration is reset. To reset the MAC address of vmk0 to a random VMware-assigned MAC address, follow these steps:

1. From the ESXi console menu main screen, type Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of Static Macros.

2. Log in as root.

3. Type “esxcfg-vmknic –l” to get a detailed listing of interface vmk0. vmk0 should be a part of the “Management Network” port group. Note the IP address and netmask of vmk0.

4. To remove vmk0, type “esxcfg-vmknic –d “Management Network””.

5. To re-add vmk0 with a random MAC address, type “esxcfg-vmknic –a –i <vmk0-ip> -n <vmk0-netmask> “Management Network””.

6. Verify vmk0 has been re-added with a random MAC address by typing “esxcfg-vmknic –l”.

7. Type “exit” to log out of the command line interface.

8. Type Ctrl-Alt-F2 to return to the ESXi console menu interface.

Log into VMware ESXi Hosts by Using VMware Host Client

ESXi Host VM-Host-Infra-01

To log into the VM-Host-Infra-01 ESXi host by using the VMware Host Client, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.

2. Click Open the VMware Host Client.

3. Enter root for the user name.

4. Enter the root password.

5. Click Login to connect.

6. Decide whether to join the VMware Customer Experience Improvement Program and click OK.

7. Repeat this process to log into VM-Host-Infra-02 in a separate browser tab or window.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To set up the VMkernel ports and the virtual switches on the ESXi hosts, follow these steps:

1. From the Host Client, select Networking .

2. In the center pane, select the Virtual switches tab.

3. Highlight the vSwitch0 line.

4. Select Edit settings.

5. Change the MTU to 9000.

6. Expand NIC teaming.

7. In the Failover order section, select vmnic1 and click Mark active.

8. Verify that vmnic1 now has a status of Active.

9. Click Save.

10. Select Networking, then select the Port groups tab.

11. In the center pane, right-click VM Network and select Remove.

12. Click Remove to complete removing the port group.

13. In the center pane, select Add port group.

14. Name the port group IB-MGMT Network and enter <ib-mgmt-vlan-id> in the VLAN ID field, and make sure Virtual switch vSwitch0 is selected.

15. Click Add to finalize the edits for the IB-MGMT Network.

16. At the top, select the VMkernel NICs tab.

17. Click Add VMkernel NIC.

18. For New port group, enter VMkernel-vMotion.

19. For Virtual switch, select vSwitch0.

20. Enter <vmotion-vlan-id> for the VLAN ID.

21. Change the MTU to 9000.

22. Select Static IPv4 settings and expand IPv4 settings.

23. Enter the ESXi host vMotion IP address and netmask.

24. Select the vMotion stack for TCP/IP stack.

25. Click Create.

26. Click Add VMkernel NIC.

27. For New port group, enter VMkernel-Infra-NFS

28. For Virtual switch, select vSwitch0.

29. Enter <infra-nfs-vlan-id> for the VLAN ID

30. Change the MTU to 9000.

31. Select Static IPv4 settings and expand IPv4 settings.

32. Enter the ESXi host Infrastructure NFS IP address and netmask.

33. Do not select any of the Services.

34. Click Create.

35. Select the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:

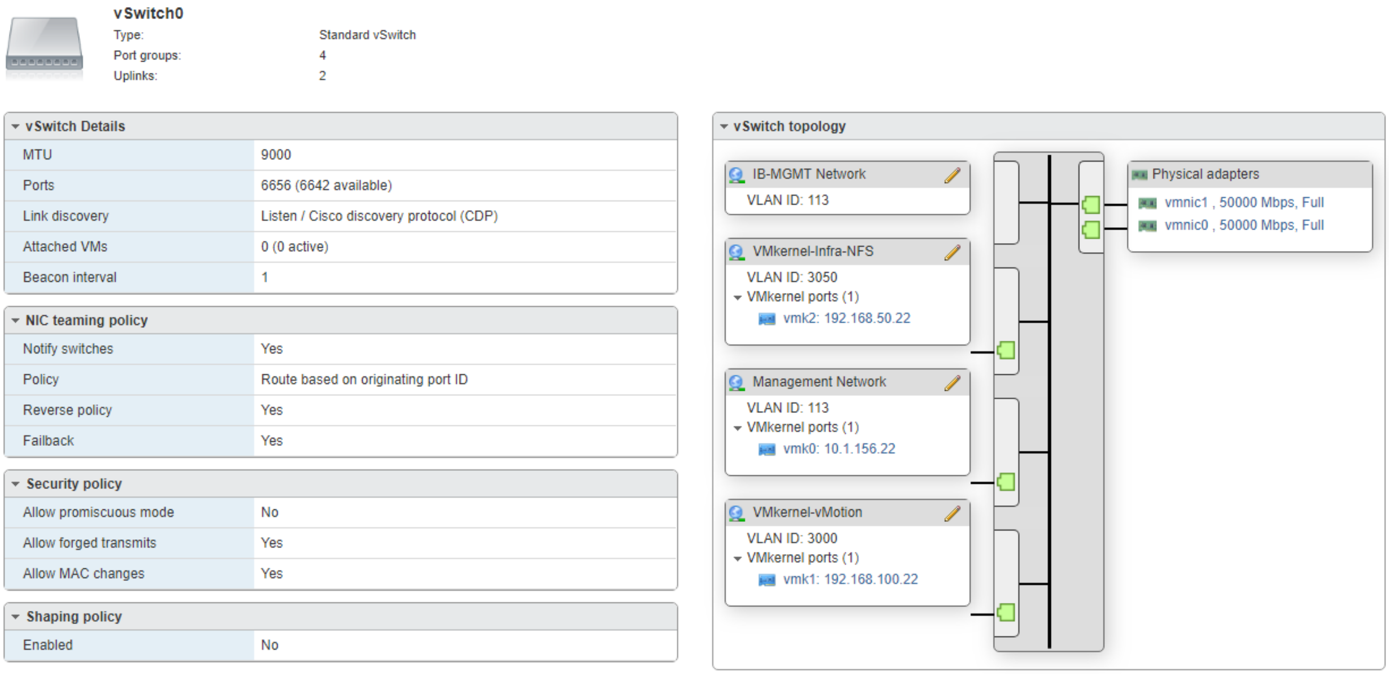

vSwitch0

Type:	Standard vSwitch
Port groups:	4
Uplinks:	2

▾ vSwitch Details

MTU	9000
Ports	6656 (6642 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

▾ NIC teaming policy

Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Failback	Yes

▾ Security policy

Allow promiscuous mode	No
Allow forged transmits	Yes
Allow MAC changes	Yes

▾ Shaping policy

Enabled	No

▾ vSwitch topology

IB-MGMT Network
VLAN ID: 113

VMkernel-Infra-NFS
VLAN ID: 3050
▾ VMkernel ports (1)
vmk2: 192.168.50.22

Management Network
VLAN ID: 113
▾ VMkernel ports (1)
vmk0: 10.1.156.22

VMkernel-vMotion
VLAN ID: 3000
▾ VMkernel ports (1)
vmk1: 192.168.100.22

Physical adapters
vmnic1 , 50000 Mbps, Full
vmnic0 , 50000 Mbps, Full

36. Select Networking and the VMkernel NICs tab to confirm configured virtual adapters. The adapters listed should be similar to the following example:

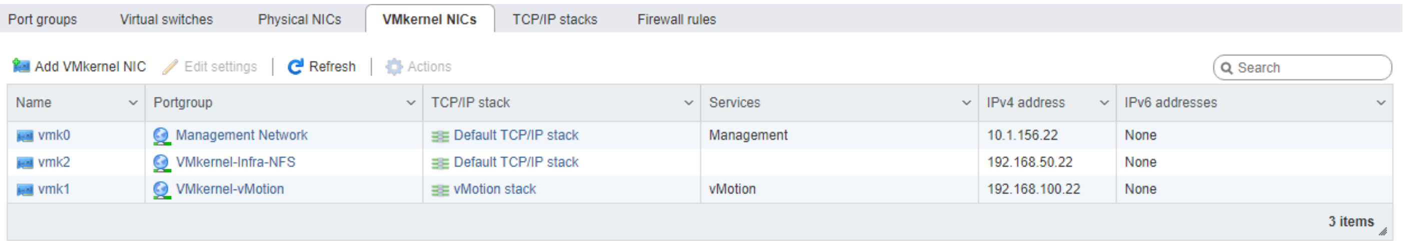

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.156.22	None
vmk2	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.22	None
vmk1	VMkernel-vMotion	vMotion stack	vMotion	192.168.100.22	None

3 items

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

Download and extract the offline bundle for the following VMware VIC Driver to the Management workstation:

nfnic Driver version 4.0.0.24

nenic Driver version 1.0.27.0

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, follow these steps:

1. From each Host Client, select Storage .

2. Right-click datastore1 and select Browse.

3. In the Datastore browser, click Upload.

4. Navigate to the saved location for the downloaded VIC drivers and select VMW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-11271332.zip.

5. In the Datastore browser, click Upload.

6. Navigate to the saved location for the downloaded VIC drivers and select VMW-ESX-6.7.0-nfnic-4.0.0.24-offline_bundle-12271979.zip.

7. Click Open to upload the file to datastore1.

8. Make sure the file has been uploaded to both ESXi hosts.

9. Place each host into Maintenance mode if it isn't already.

10. Connect to each ESXi host through ssh from a shell connection or putty terminal.

11. Login as root with the root password.

12. Run the following commands on each host:

```
esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-
11271332.zip

esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.7.0-nfnic-4.0.0.24-offline_bundle-
12271979.zip

reboot
```

13. Log into the Host Client on each host once reboot is complete and exit Maintenance Mode.

Mount Required Datastores

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To mount the required datastores, follow these steps on each ESXi host:

1. From the Host Client, select Storage .

2. In the center pane, select the Datastores tab.

3. In the center pane, select New Datastore to add a new datastore.

4. In the New datastore popup, select Mount NFS datastore and click Next.

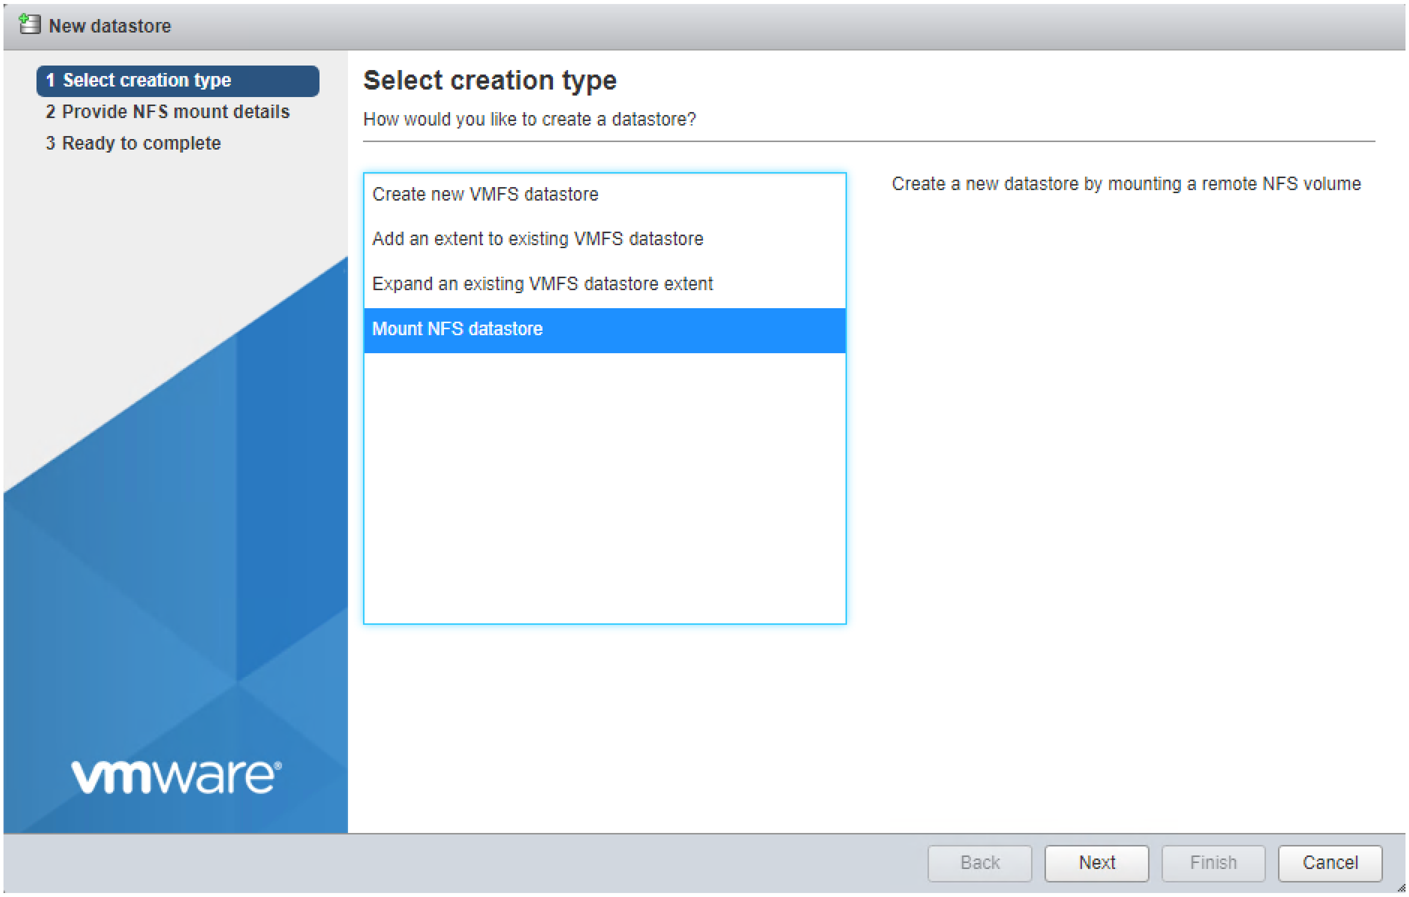

Select creation type

How would you like to create a datastore?

Create new VMFS datastore

Add an extent to existing VMFS datastore

Expand an existing VMFS datastore extent

Mount NFS datastore

Create a new datastore by mounting a remote NFS volume

5. Input infra_datastore_1 for the datastore name. Input the IP address for the nfs_lif01 LIF for the NFS server. Input /infra_datastore_1 for the NFS share. Leave the NFS version set at NFS 3. Click Next.

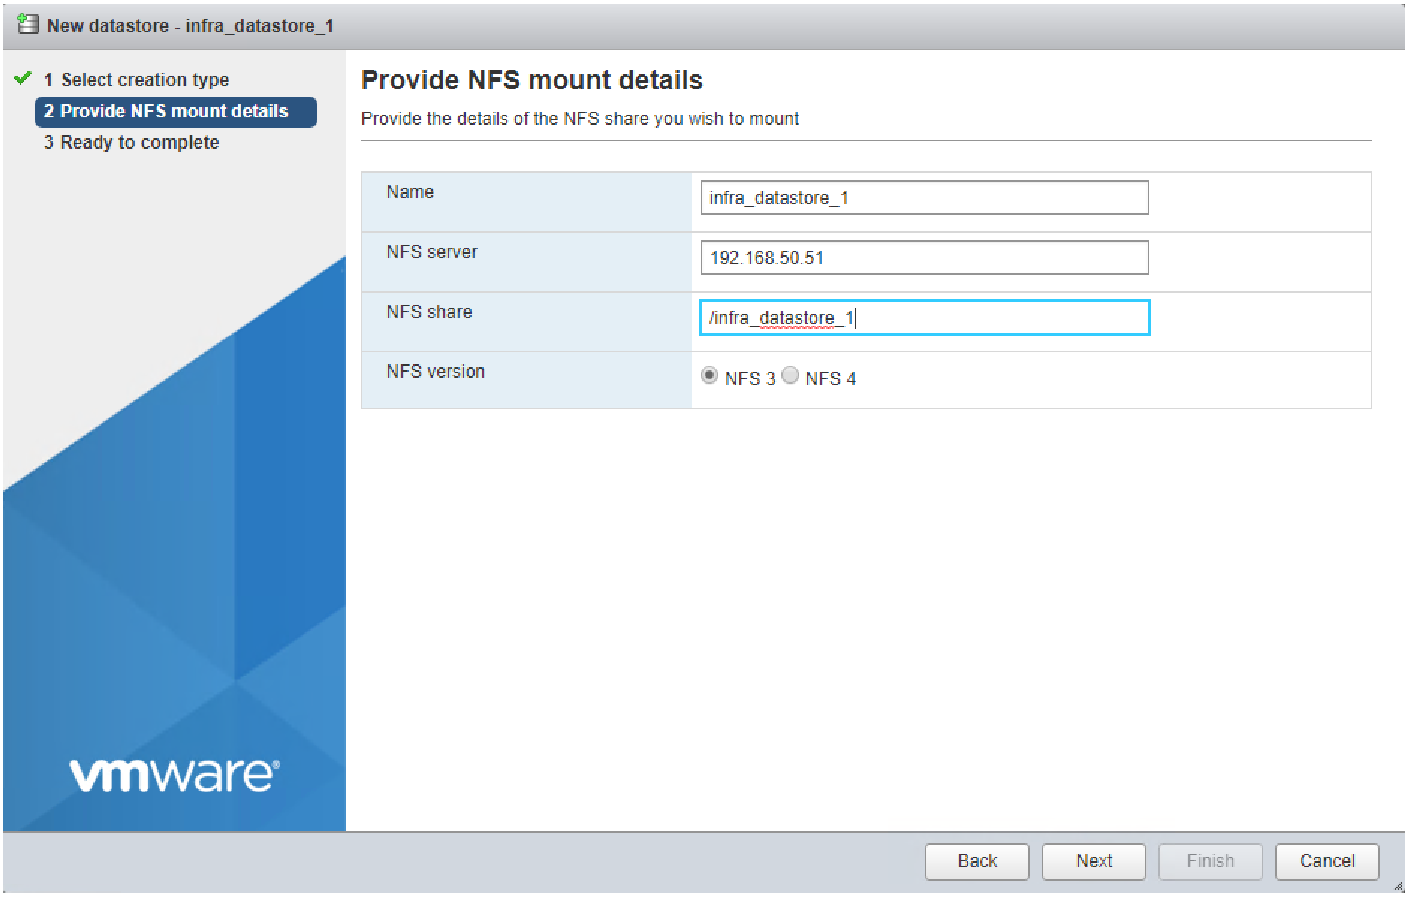

Provide NFS mount details

Provide the details of the NFS share you wish to mount

Name	infra_datastore_1
NFS server	192.168.50.51
NFS share	/infra_datastore_1
NFS version	● NFS 3 ○ NFS 4

6. Click Finish. The datastore should now appear in the datastore list.

7. In the center pane, select New Datastore to add a new datastore.

8. In the New datastore popup, select Mount NFS datastore and click Next.

9. Input infra_datastore_2 for the datastore name. Input the IP address for the nfs_lif02 LIF for the NFS server. Input /infra_datastore_2 for the NFS share. Leave the NFS version set at NFS 3. Click Next.

10. Click Finish. The datastore should now appear in the datastore list.

11. In the center pane, select New Datastore to add a new datastore.

12. In the New datastore popup, select Mount NFS datastore and click Next.

13. Input infra_swap for the datastore name. Input the IP address for the nfs_lif01 LIF for the NFS server. Input /infra_swap for the NFS share. Leave the NFS version set at NFS 3. Click Next.

14. Click Finish. The datastore should now appear in the datastore list.

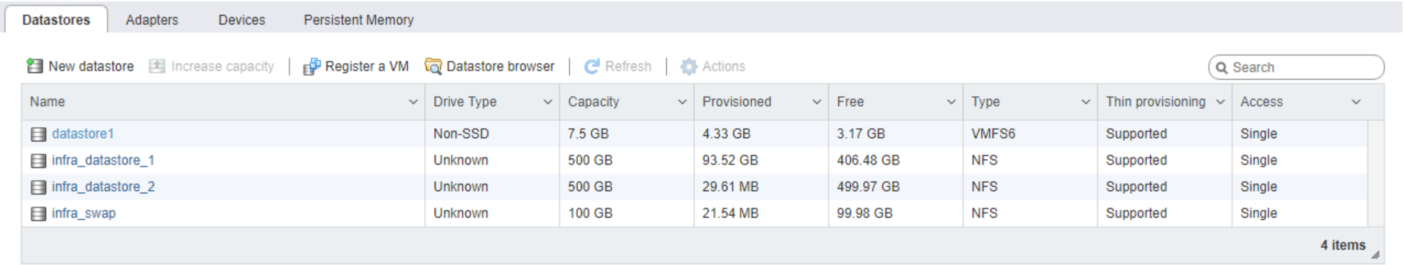

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
datastore1	Non-SSD	7.5 GB	4.33 GB	3.17 GB	VMFS6	Supported	Single
infra_datastore_1	Unknown	500 GB	93.52 GB	406.48 GB	NFS	Supported	Single
infra_datastore_2	Unknown	500 GB	29.61 MB	499.97 GB	NFS	Supported	Single
infra_swap	Unknown	100 GB	21.54 MB	99.98 GB	NFS	Supported	Single

4 items

15. Mount all three datastores on both ESXi hosts.

Configure NTP on ESXi Hosts

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

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

1. From the Host Client, select Manage.

2. In the center pane, select Time & date.

3. Click Edit settings.

4. Make sure Use Network Time Protocol (enable NTP client) is selected.

5. Use the drop-down to select Start and stop with host.

6. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.

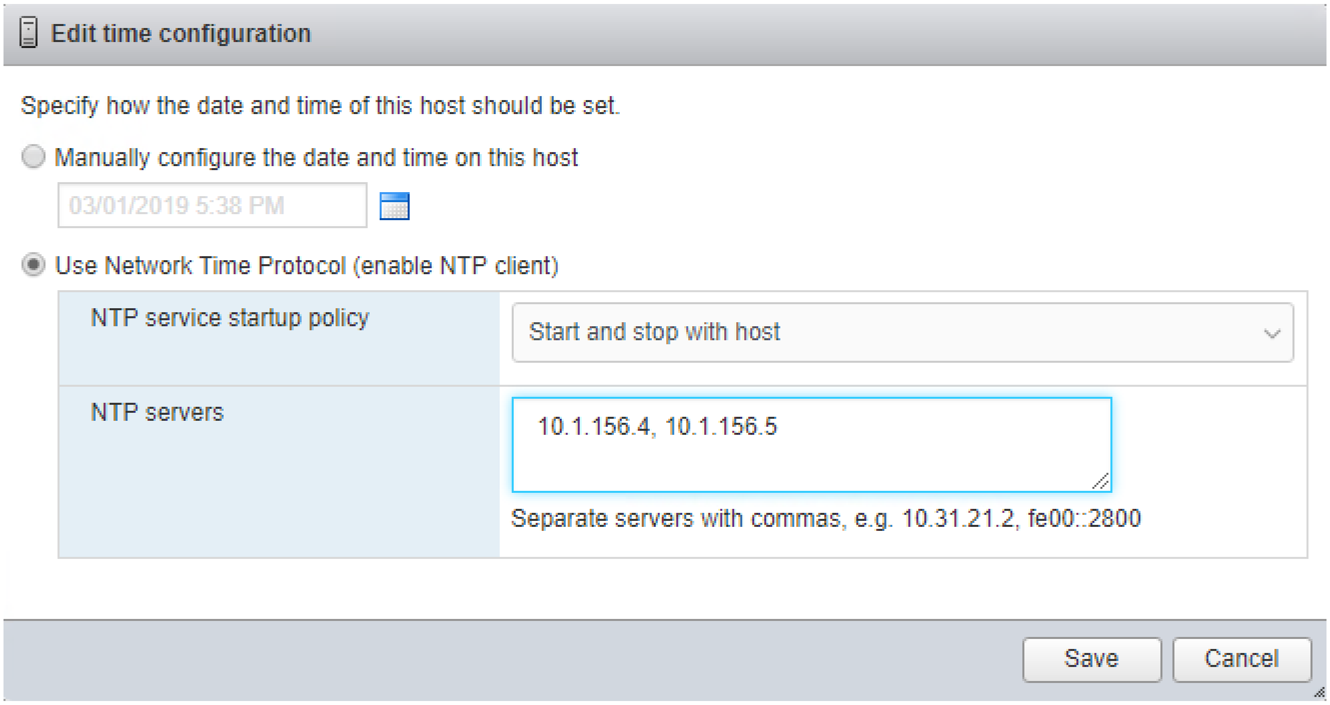

Edit time configuration

Specify how the date and time of this host should be set.

○ Manually configure the date and time on this host

03/01/2019 5:38 PM

● Use Network Time Protocol (enable NTP client)

NTP service startup policy	Start and stop with host
NTP servers	10.1.156.4, 10.1.156.5 Separate servers with commas, e.g. 10.31.21.2, fe00::2800

7. Click Save to save the configuration changes.

8. Select Actions > NTP service > Start.

9. Verify that NTP service is now running and the clock is now set to approximately the correct time.

The NTP server time may vary slightly from the host time.

Configure ESXi Host Swap

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To configure host swap on the ESXi hosts, follow these steps on each host:

1. From the Host Client, select Manage.

2. In the center pane, select Swap.

3. Click Edit settings.

4. Use the drop-down list to select infra_swap. Leave all other settings unchanged.

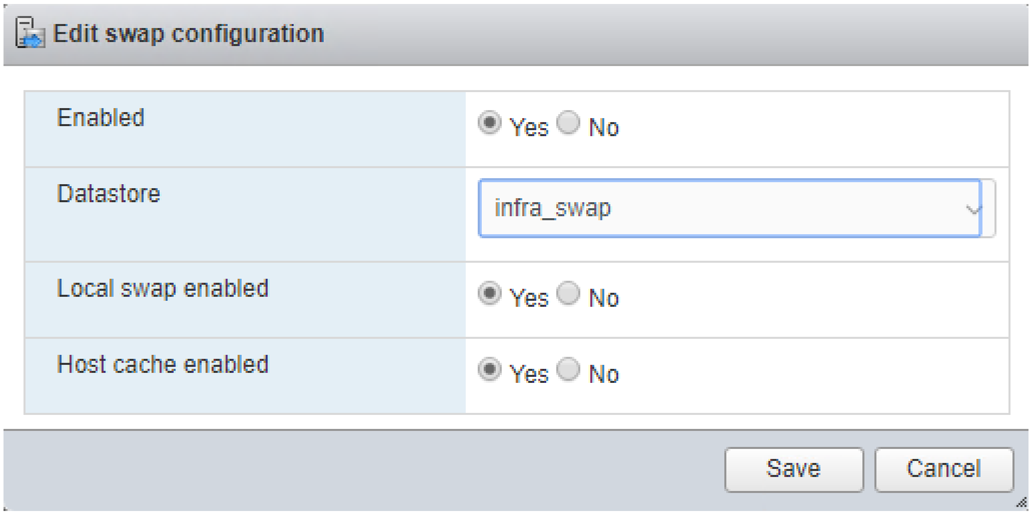

Enabled	● Yes ○ No
Datastore	infra_swap
Local swap enabled	● Yes ○ No
Host cache enabled	● Yes ○ No

5. Click Save to save the configuration changes.

6. If you are implementing iSCSI boot, execute the VMware ESXi setup scripts in the iSCSI Addition Appendix.

VMware vCenter 6.7U1

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.7U1 Server Appliance in a FlexPod environment. After the procedures are completed, a VMware vCenter Server will be configured.

Build the VMware vCenter Server Appliance

The VCSA deployment consists of 2 stages: install and configuration. To build the VMware vCenter virtual machine, follow these steps:

1. Locate and copy the VMware-VCSA-all-6.7.0-11726888.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 6.5 U1 vCenter Server Appliance.

2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012).

3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appears.

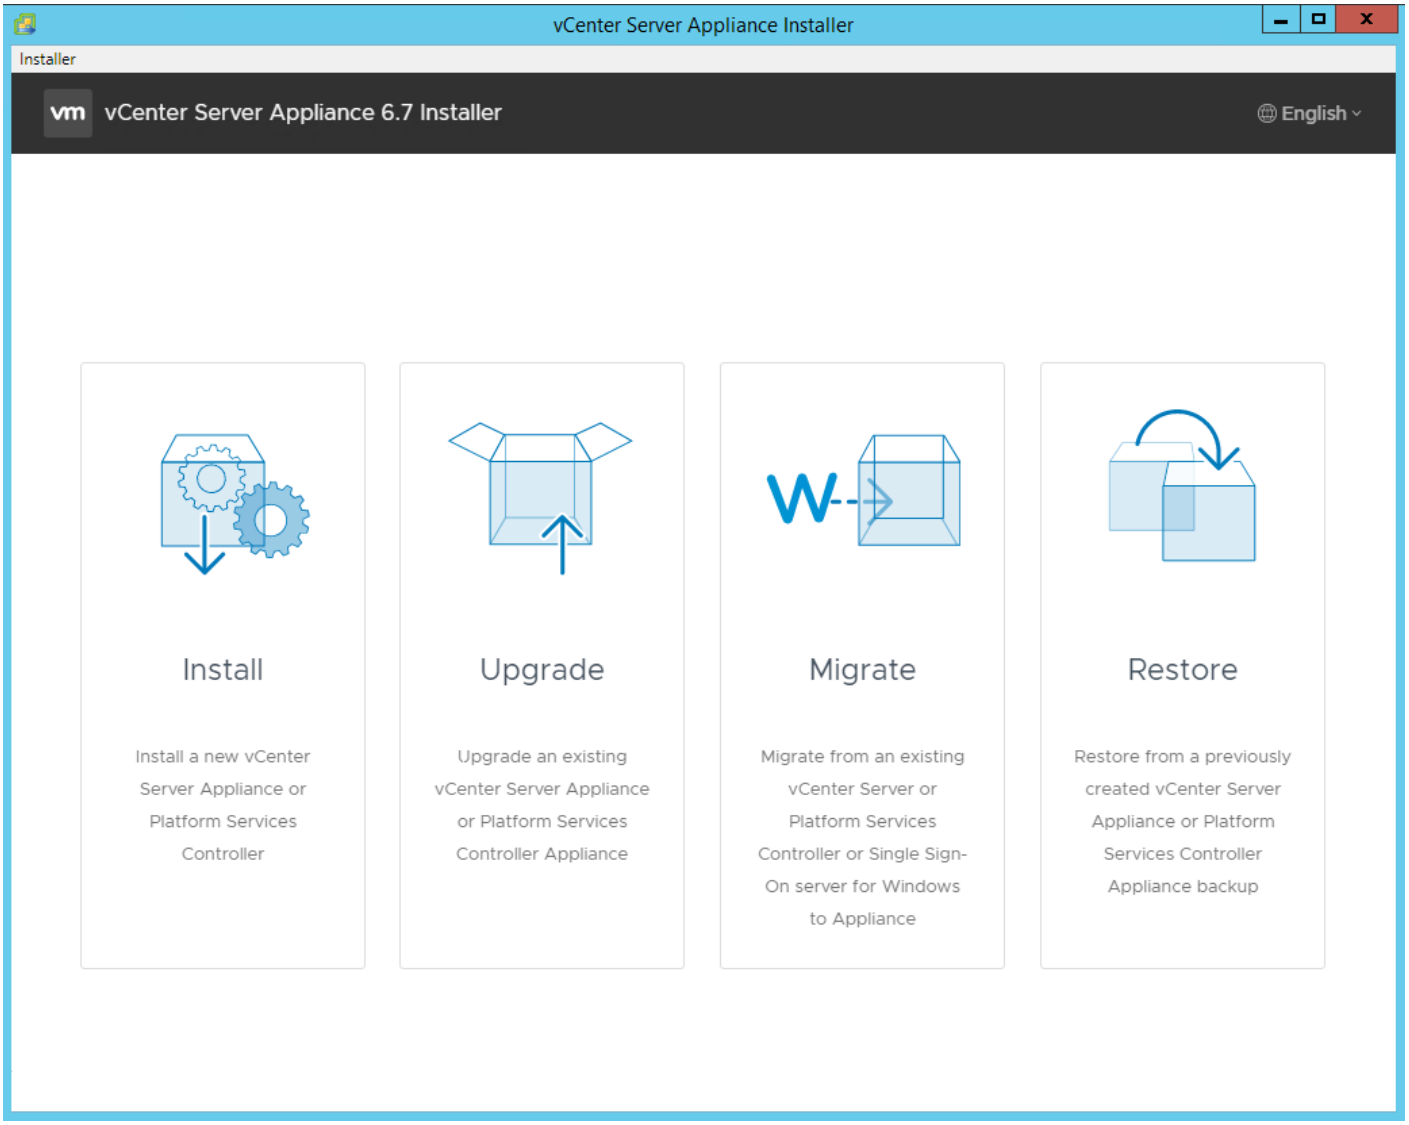

vCenter Server Appliance Installer

Installer

vCenter Server Appliance 6.7 Installer

Install

Install a new vCenter Server Appliance or Platform Services Controller

Upgrade

Upgrade an existing vCenter Server Appliance or Platform Services Controller Appliance

Migrate

Migrate from an existing vCenter Server or Platform Services Controller or Single Sign-On server for Windows to Appliance

Restore

Restore from a previously created vCenter Server Appliance or Platform Services Controller Appliance backup

4. Click Install to start the vCenter Server Appliance deployment wizard.

5. Click Next in the Introduction section.

6. Read and accept the license agreement and click NEXT.

7. In the "Select deployment type" section, select Embedded Platform Services Controller and click NEXT.

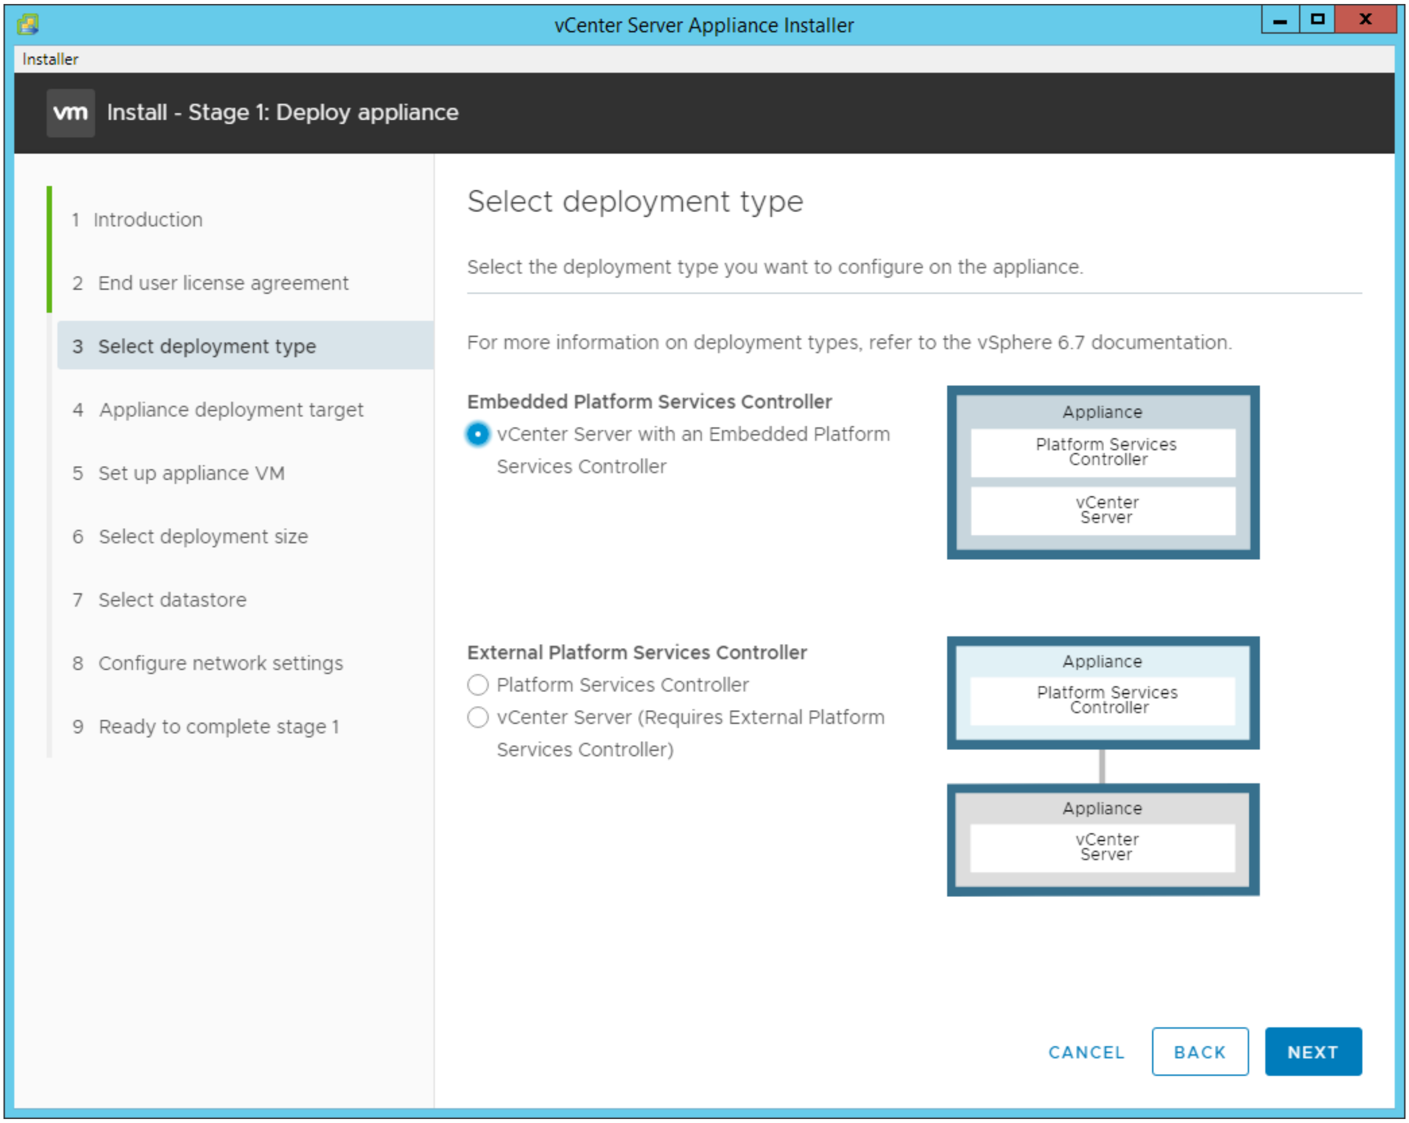

vCenter Server Appliance Installer

Installer

Install - Stage 1: Deploy appliance

1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

Select deployment type

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

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

Embedded Platform Services Controller

● vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

○ Platform Services Controller

○ vCenter Server (Requires External Platform Services Controller)

8. In the "Appliance deployment target", enter the host name or IP address of the first ESXi host, User name (root) and Password.

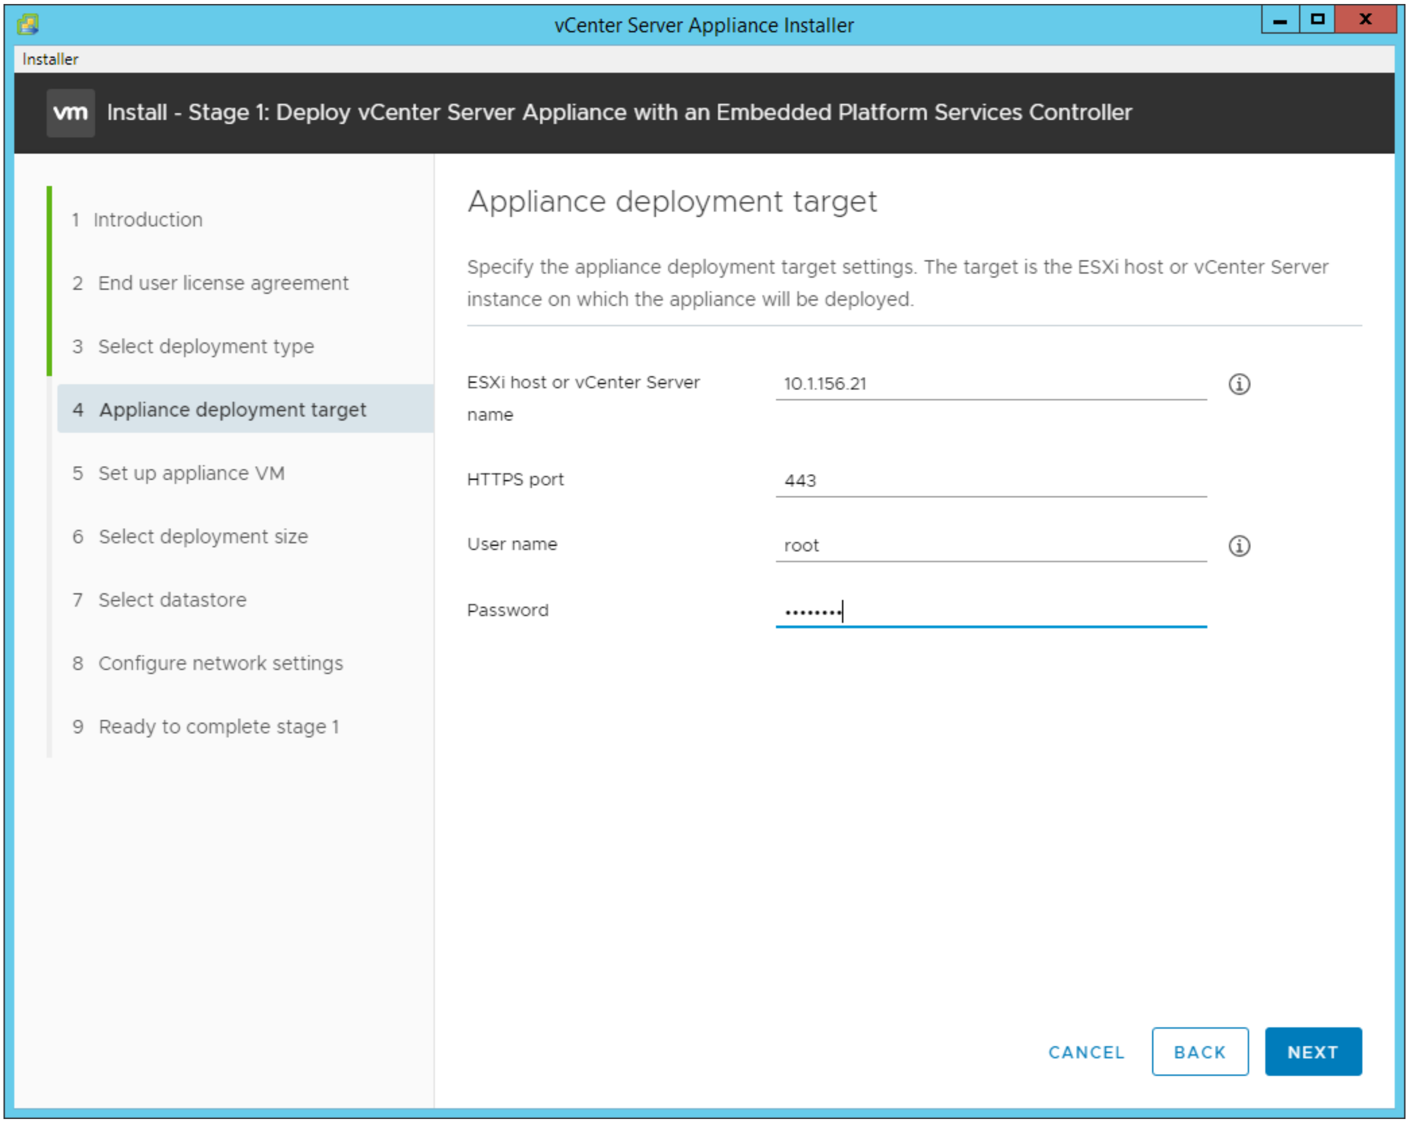

vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore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

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

ESXi host or vCenter Server name: 10.1.156.21

HTTPS port: 443

User name: root

Password: ••••••••

9. Click Yes to accept the certificate.

10. Enter the Appliance VM name and password details in the “Set up appliance VM” section. Click NEXT.

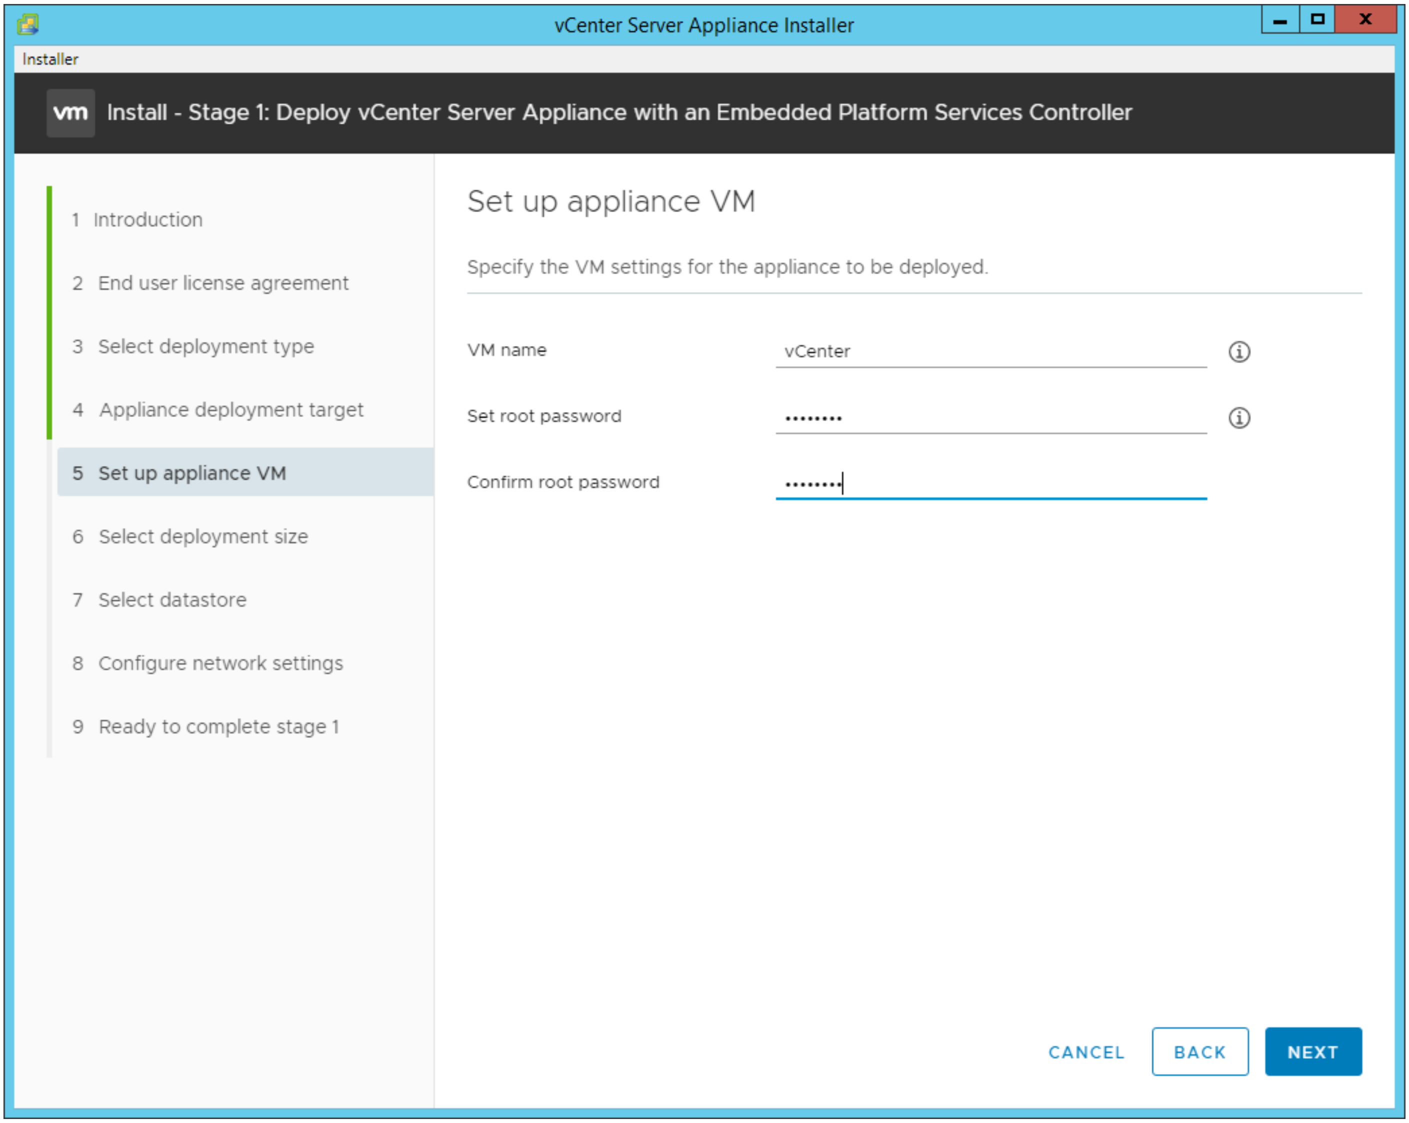
vCenter Server Appliance Installer
Installer
vm
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore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

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name: vCenter

Set root password: ••••••••

Confirm root password: ••••••••

11. In the “Select deployment size” section, Select the deployment size and Storage size. For example, “Small.”

12. Click NEXT.

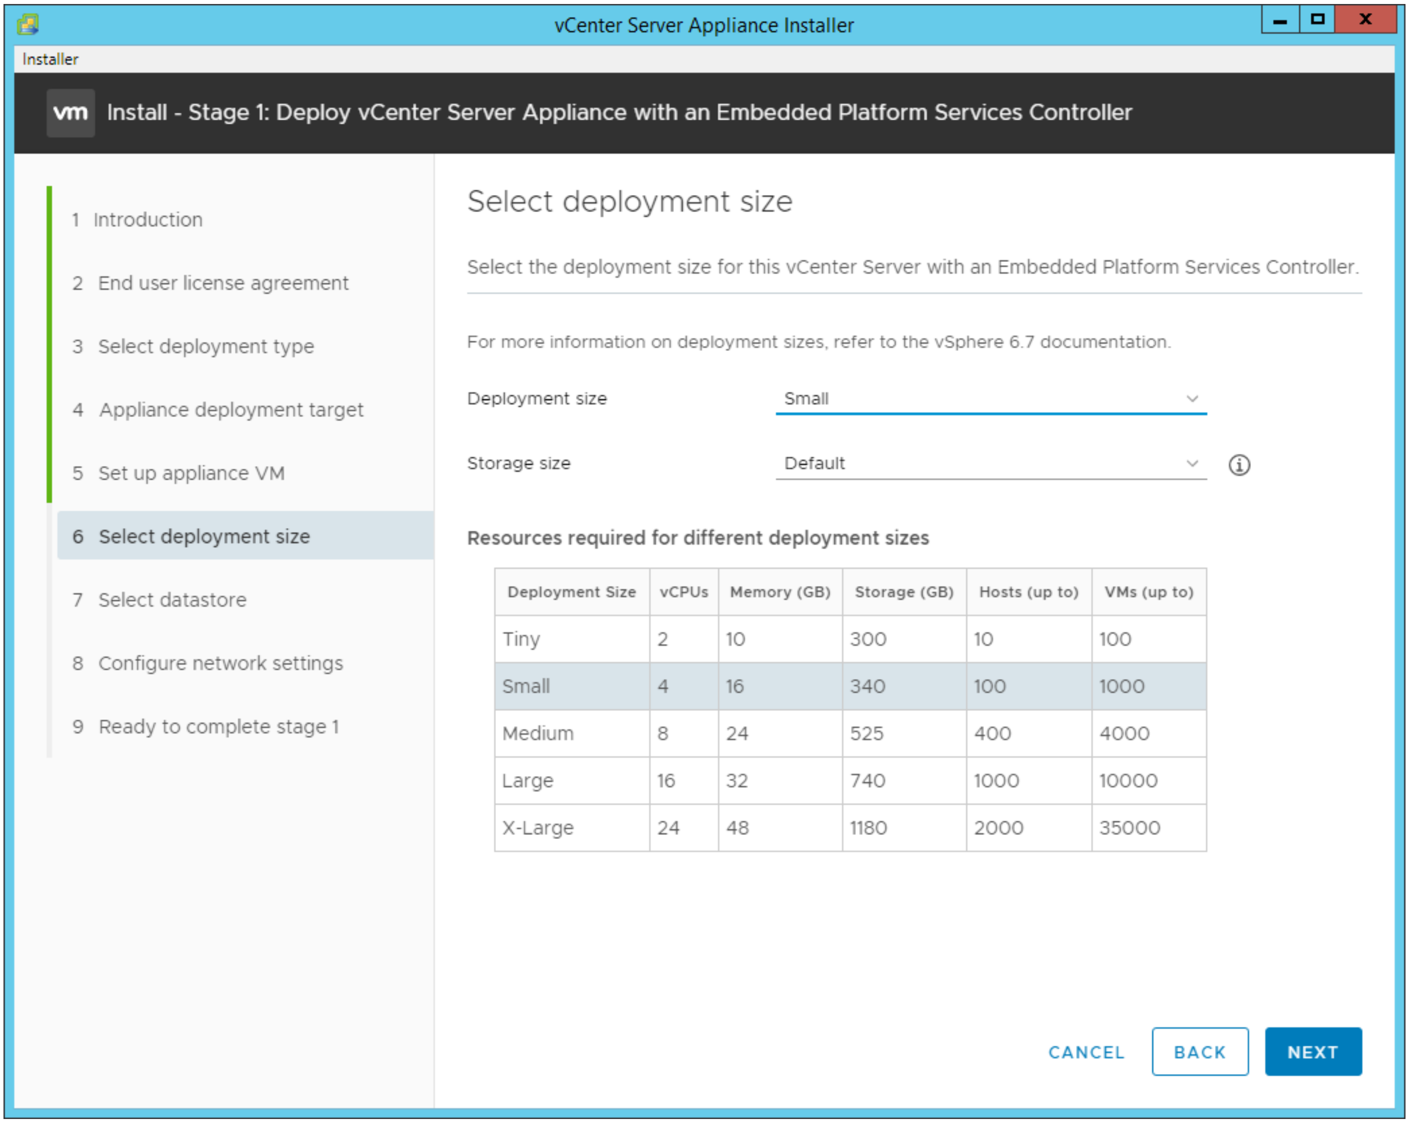

vCenter Server Appliance Installer
Installer
vm
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore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

Select deployment size

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

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

Deployment size: Small

Storage size: Default

Resources required for different deployment sizes

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

13. Select infra_datastore_2. Click NEXT.

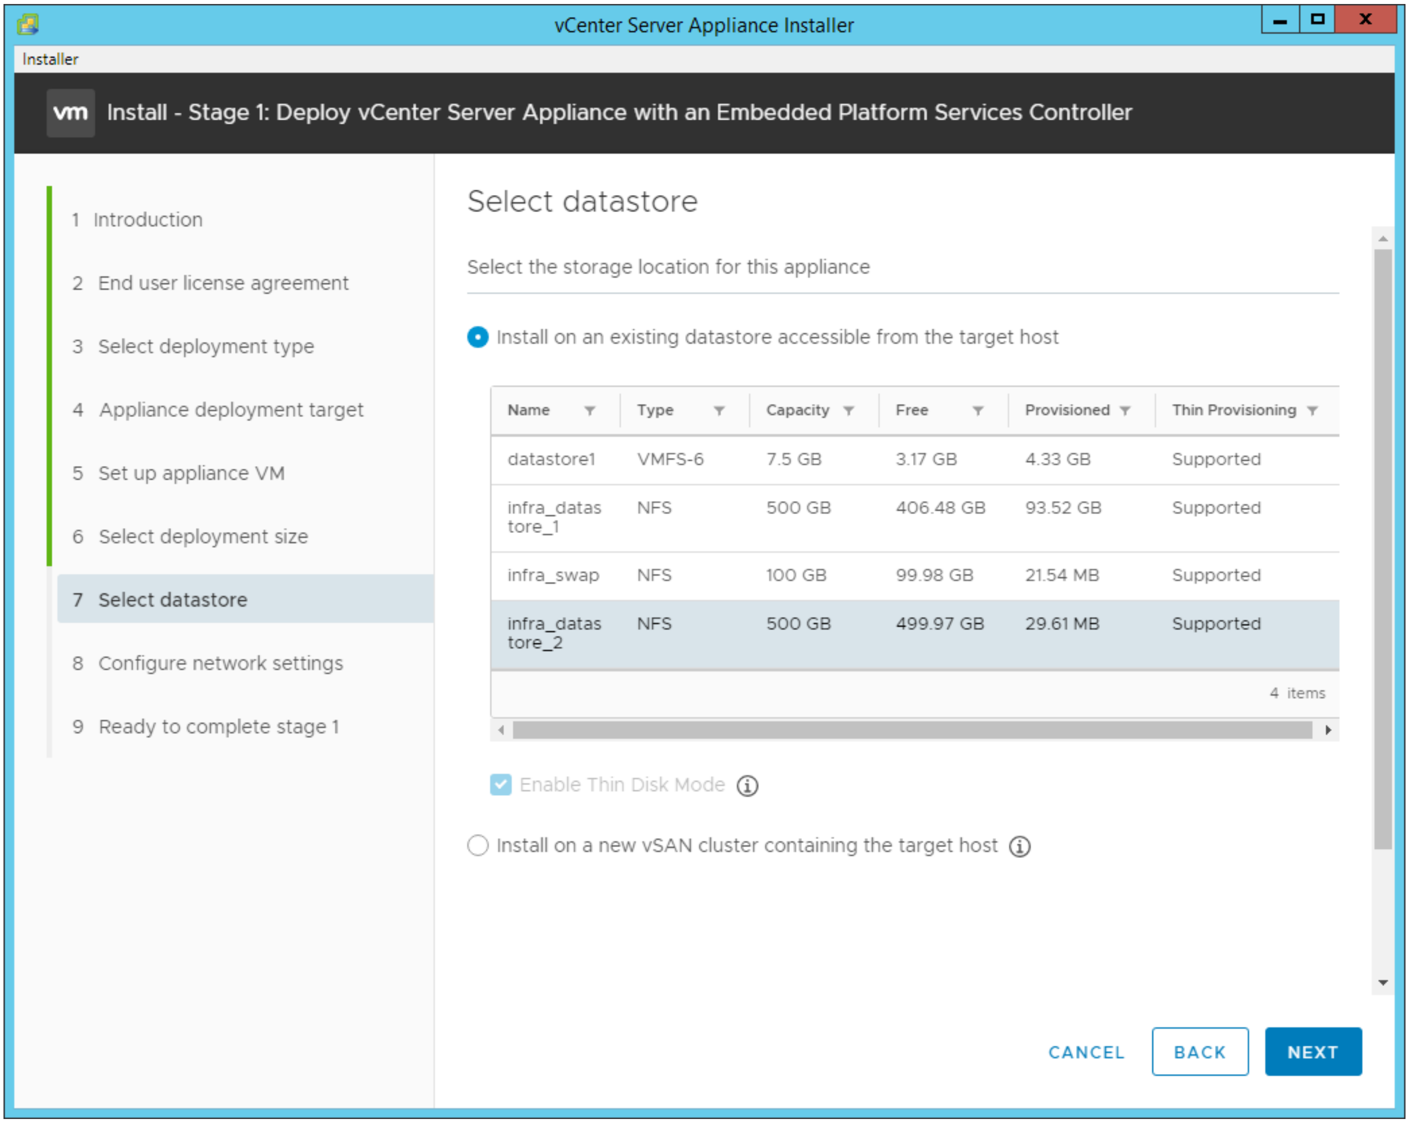
vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore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

Select datastore

Select the storage location for this appliance

● Install on an existing datastore accessible from the target host

Name	Type	Capacity	Free	Provisioned	Thin Provisioning
datastore1	VMFS-6	7.5 GB	3.17 GB	4.33 GB	Supported
infra_datastore_1	NFS	500 GB	406.48 GB	93.52 GB	Supported
infra_swap	NFS	100 GB	99.98 GB	21.54 MB	Supported
infra_datastore_2	NFS	500 GB	499.97 GB	29.61 MB	Supported

4 items

☑ Enable Thin Disk Mode

○ Install on a new vSAN cluster containing the target host

14. In the “Network Settings” section, configure the below settings:

    a. Choose a Network: IB-MGMT Network.

It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it not get moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, and it is attempted to bring up vCenter on a different host than the one it was running on before the shutdown, vCenter will not have a functional network connection. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. Moving vCenter to vSwitch0 to be brought up on a different ESXi host always occurs correctly without requiring vCenter to be up and running.

    b. IP version: IPV4

    c. IP assignment: static

d. System name: <vcenter-fqdn>

e. IP address: <vcenter-ip>

f. Subnet mask or prefix length: <vcenter-subnet-mask>

g. Default gateway: <vcenter-gateway>

h. DNS Servers: <dns-server1>, <dns-server2>

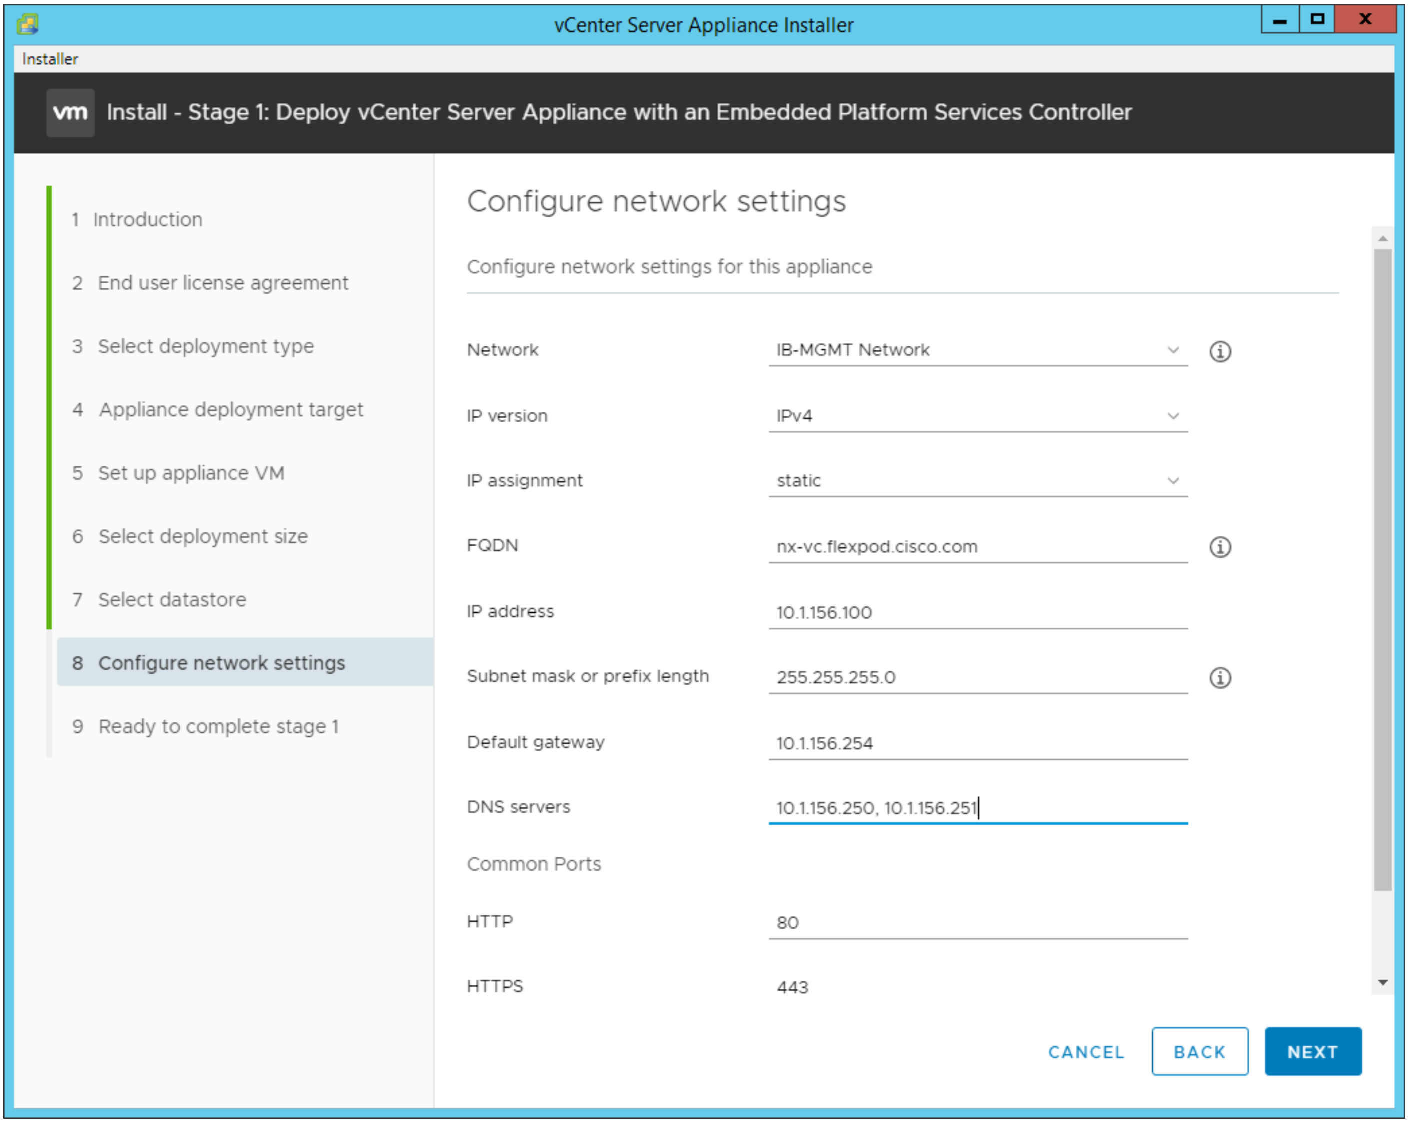
vCenter Server Appliance Installer
Installer
vm
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore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

Configure network settings

Configure network settings for this appliance

Network	IB-MGMT Network
IP version	IPv4
IP assignment	static
FQDN	nx-vc.flexpod.cisco.com
IP address	10.1.156.100
Subnet mask or prefix length	255.255.255.0
Default gateway	10.1.156.254
DNS servers	10.1.156.250, 10.1.156.251

Common Ports

HTTP	80
HTTPS	443

15. Click NEXT.

16. Review all values and click Finish to complete the installation.

17. The vCenter appliance installation will take a few minutes to complete.

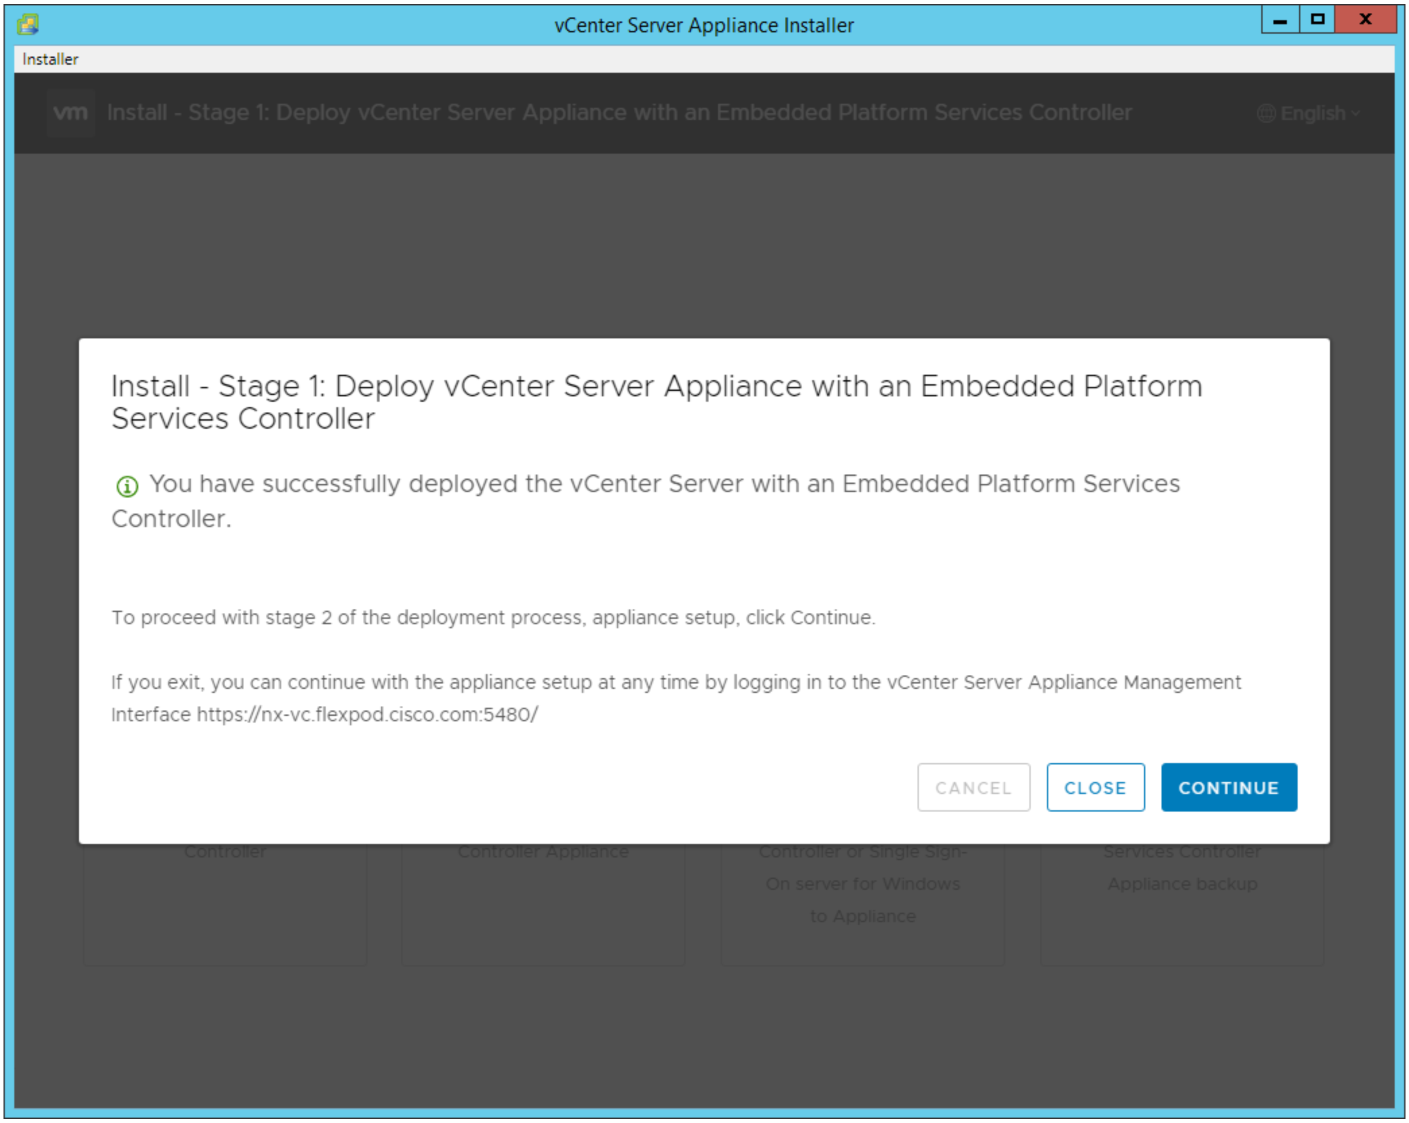
vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller
English

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

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

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

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

18. Click CONTINUE to proceed with stage 2 configuration.

19. Click NEXT.

20. In the Appliance Configuration, configure these settings:

    a. Time Synchronization Mode: Synchronize time with NTP servers.
    b. NTP Servers: <switch-a-ntp-ip>, <switch-b-ntp-ip>
    c. SSH access: Enabled.

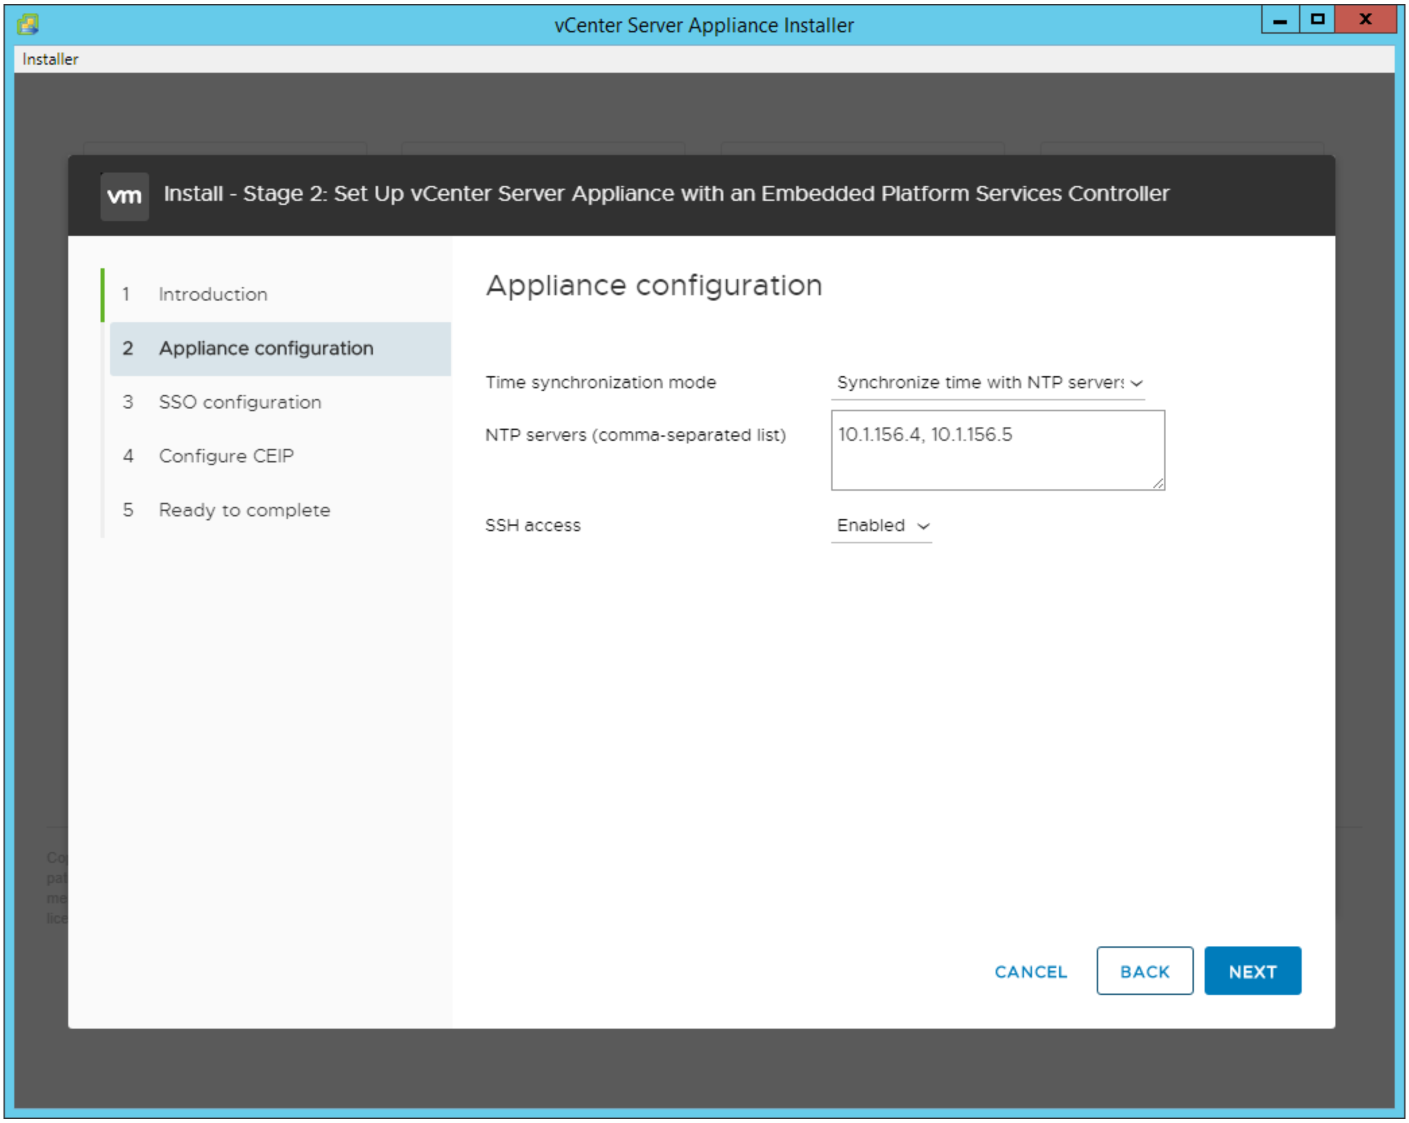
vCenter Server Appliance Installer
Installer
Install - Stage 2: Set Up vCenter Server Appliance with an Embedded Platform Services Controller
1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
CANCEL
BACK
NEXT

Appliance configuration

Time synchronization mode: Synchronize time with NTP servers

NTP servers (comma-separated list): 10.1.156.4, 10.1.156.5

SSH access: Enabled

21. Click NEXT.

22. Complete the SSO configuration as shown below:

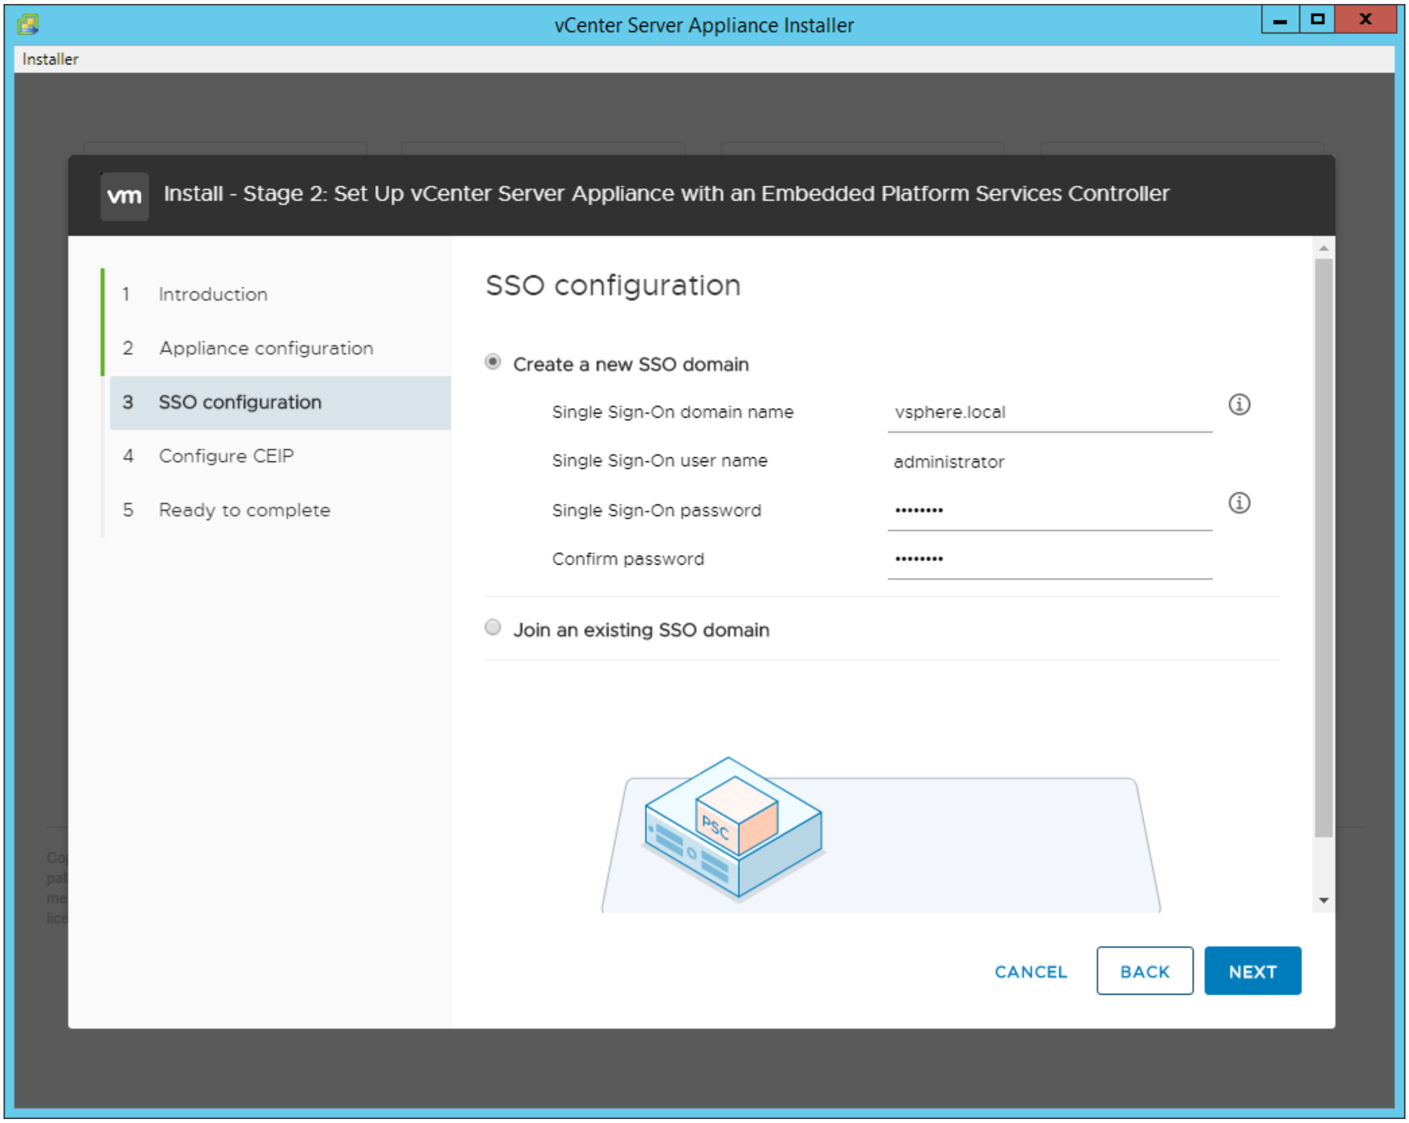

vCenter Server Appliance Installer

Installer

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

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

SSO configuration

● Create a new SSO domain

Single Sign-On domain name: vsphere.local
Single Sign-On user name: administrator
Single Sign-On password: ••••••••
Confirm password: ••••••••

○ Join an existing SSO domain

23. Click NEXT.

24. Decide whether to join VMware’s Customer Experience Improvement Program (CEIP).

25. Click NEXT.

26. Review the configuration and click FINISH.

27. Click OK.

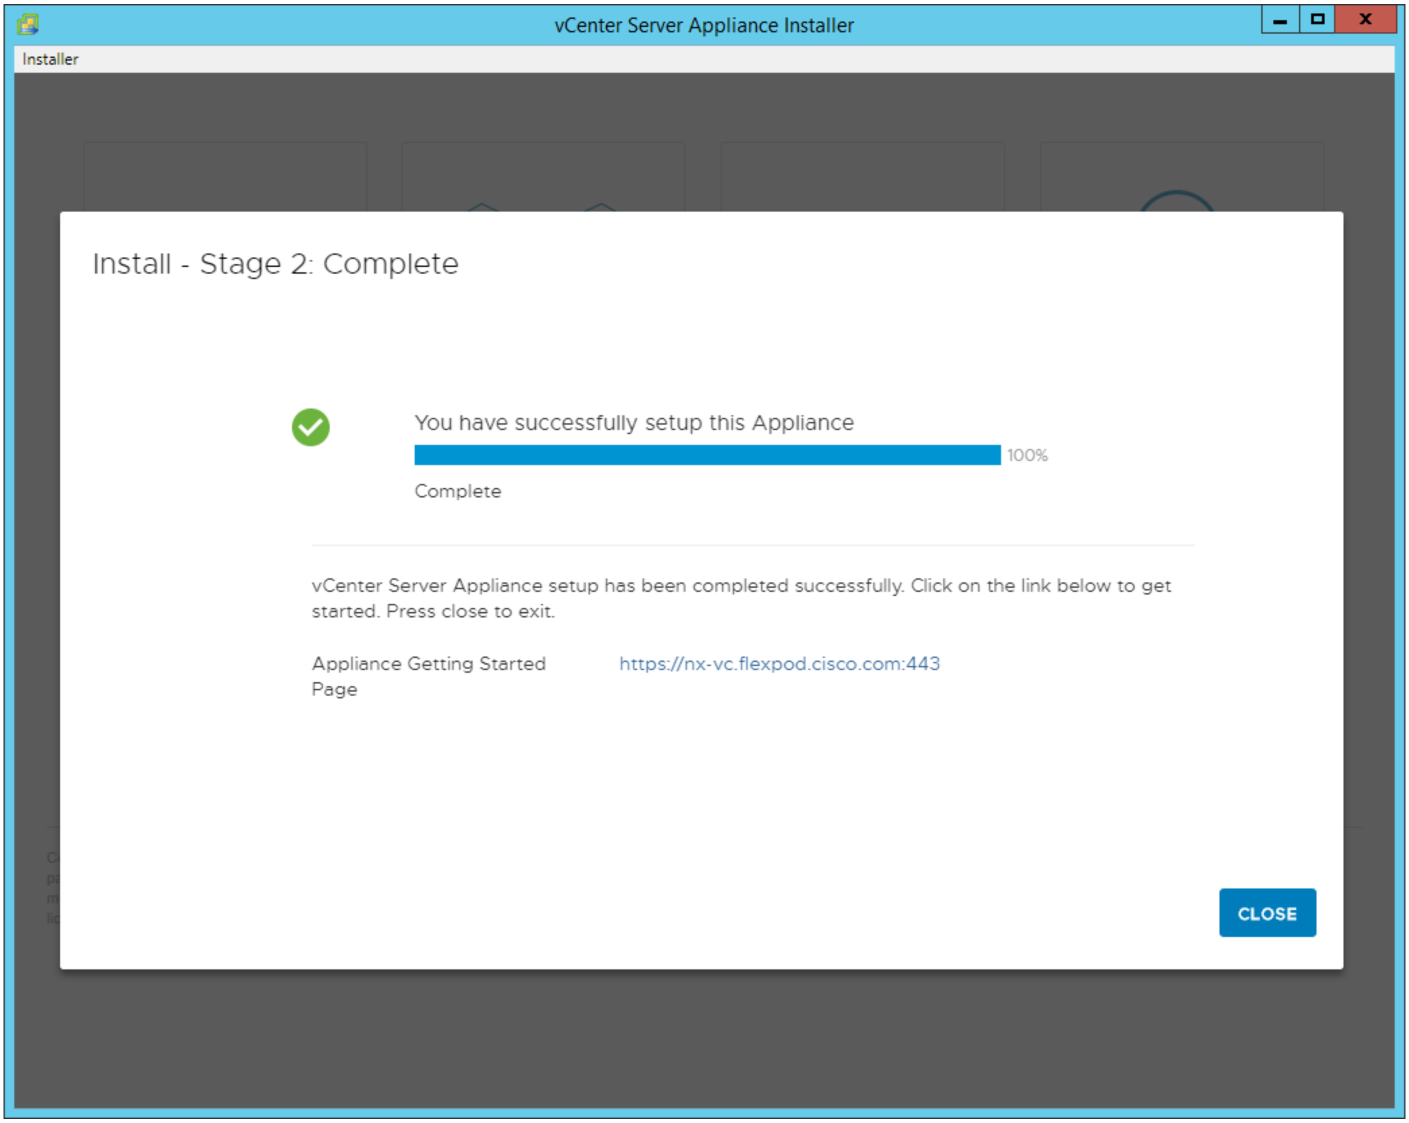

Install - Stage 2: Complete

You have successfully setup this Appliance

100%

Complete

vCenter Server Appliance setup has been completed successfully. Click on the link below to get started. Press close to exit.

Appliance Getting Started Page https://nx-vc.flexpod.cisco.com:443

28. Click CLOSE. Eject or unmount the VCSA installer ISO.

Adjust vCenter CPU Settings

If a vCenter deployment size of Small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the UCS server CPU hardware configuration. Cisco UCS B200 and C220 servers are 2-socket servers. In this validation, the Small deployment size was selected and vCenter was setup for a 4-socket server. This setup will cause issues in the VMware ESXi cluster Admission Control. To resolve the Admission Control issue, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.

2. Click Open the VMware Host Client.

3. Enter root for the user name.

4. Enter the root password.

5. Click Login to connect.

6. In the center pane, right-click the vCenter VM and select Edit settings.

7. In the Edit settings window, expand CPU and check the value of Sockets.

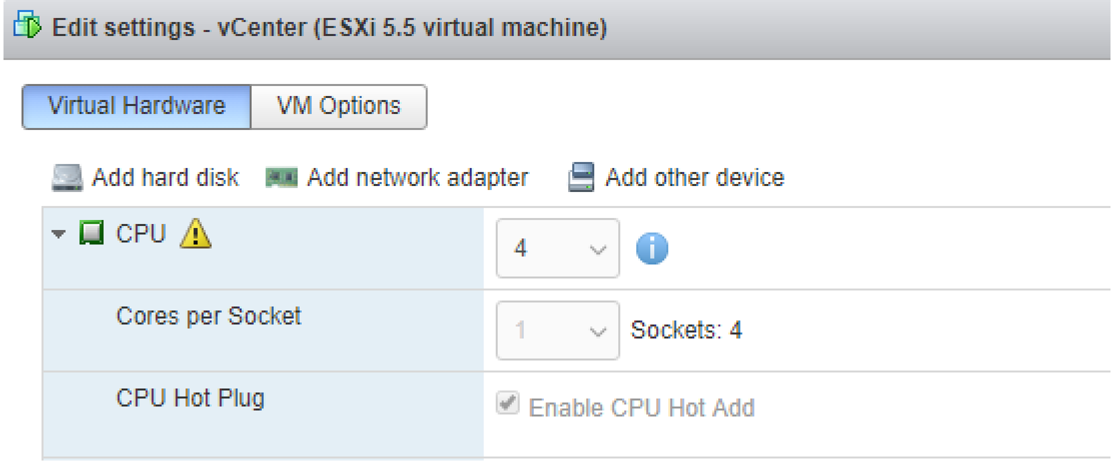

Edit settings - vCenter (ESXi 5.5 virtual machine)

CPU: 4

Cores per Socket: 1 Sockets: 4

CPU Hot Plug: ☑ Enable CPU Hot Add

8. If the number of Sockets is greater than 2, it will need to be adjusted. Click Cancel.

9. If the number of Sockets needs to be adjusted:

   a. Right-click the vCenter VM and select Guest OS > Shut down. Click Yes on the confirmation.

   b. Once vCenter is shut down, right-click the vCenter VM and select Edit settings.

   c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value 2.

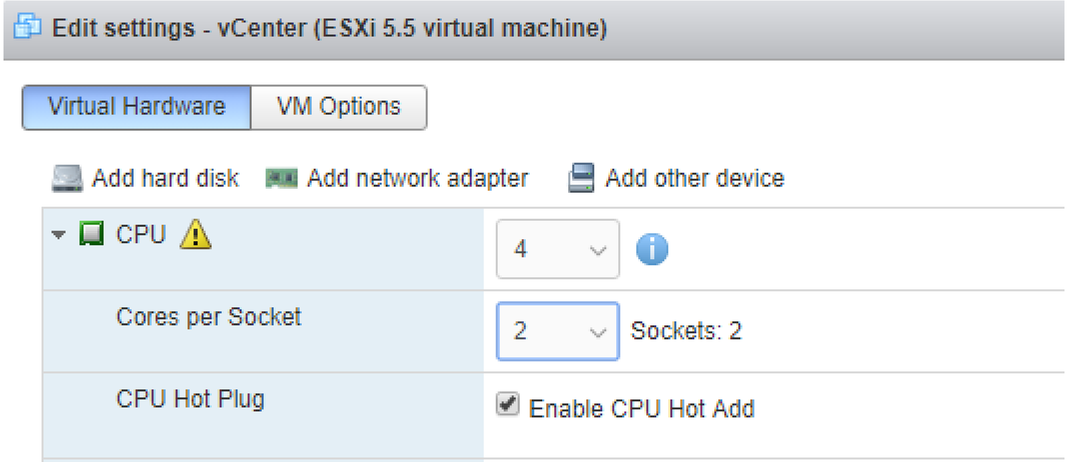

Edit settings - vCenter (ESXi 5.5 virtual machine)

CPU: 4

Cores per Socket: 2 Sockets: 2

CPU Hot Plug: ☑ Enable CPU Hot Add

   d. Click Save.

   e. Right-click the vCenter VM and select Power > Power on. Wait approximately 10 minutes for vCenter to come up.

Setup VMware vCenter Server

To setup the VMware vCenter Server, follow these steps:

1. Using a web browser, navigate to Error! Hyperlink reference not valid..

The VMware vSphere HTML5 Client is fully featured in vSphere 6.7U1. However, vCenter plugins used in this FlexPod deployment are only available with the VMware vSphere Web Client. The Web Client is used in this document.

2. If the link is available, click Download Enhanced Authentication Plugin. Install the same by double-clicking the downloaded file.

3. Log in using the Single Sign-On username (administrator@vsphere.local) and password created during the vCenter installation.

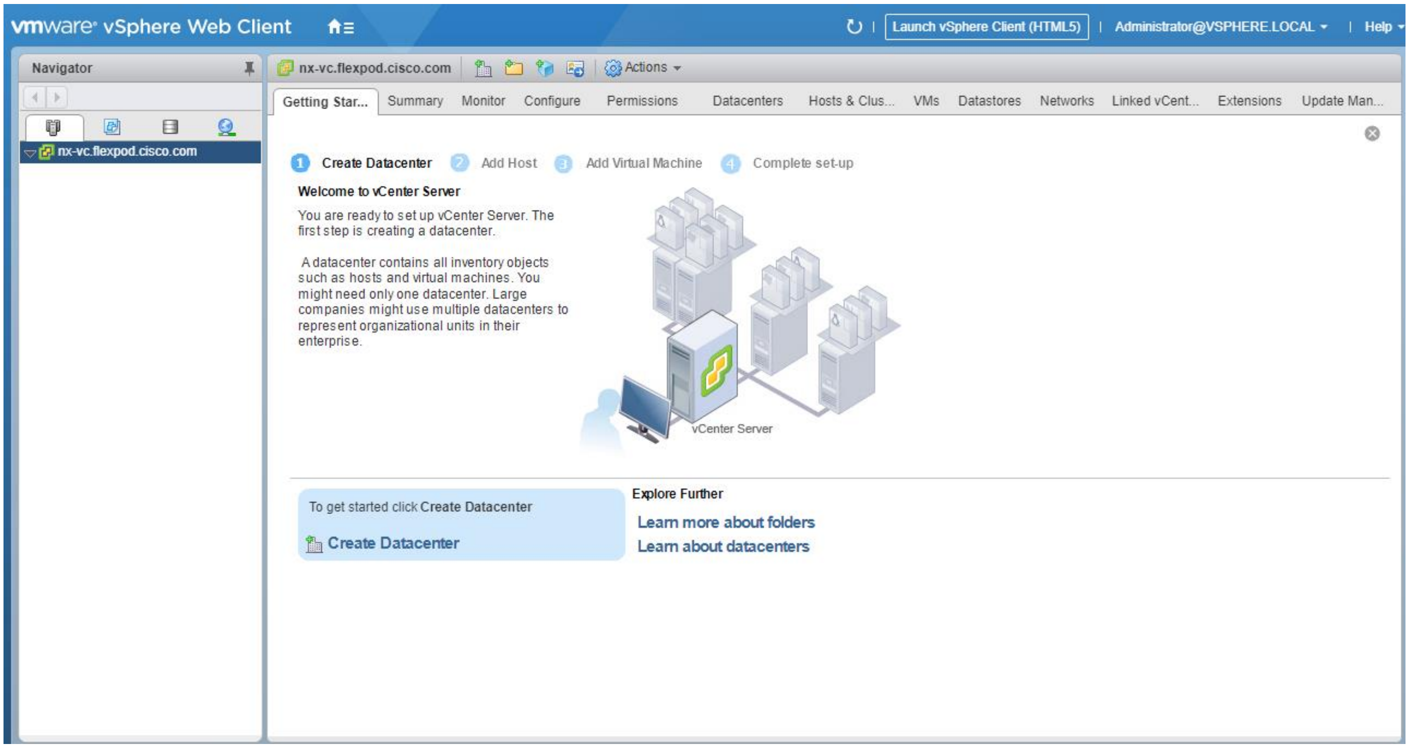

Welcome to vCenter Server

You are ready to set up vCenter Server. The first step is creating a datacenter.

A datacenter contains all inventory objects such as hosts and virtual machines. You might need only one datacenter. Large companies might use multiple datacenters to represent organizational units in their enterprise.

To get started click Create Datacenter

4. Click “Create Datacenter” in the center pane.

5. Type “FlexPod-DC” in the Datacenter name field.

6. Click OK.

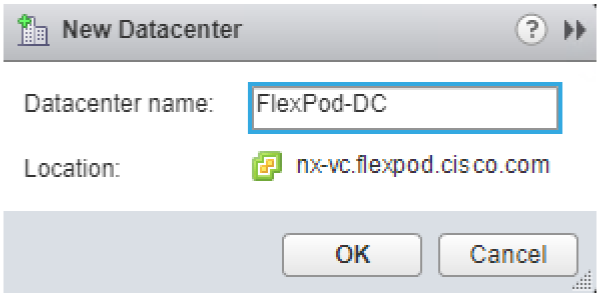

New Datacenter

Datacenter name: FlexPod-DC

Location: nx-vc.flexpod.cisco.com

7. Right-click the data center FlexPod-DC in the list in the left pane. Click New Cluster.

8. Name the cluster FlexPod-Management.

9. Check the box to turn on DRS. Leave the default values.

10. Check the box to turn on vSphere HA. Leave the default values.

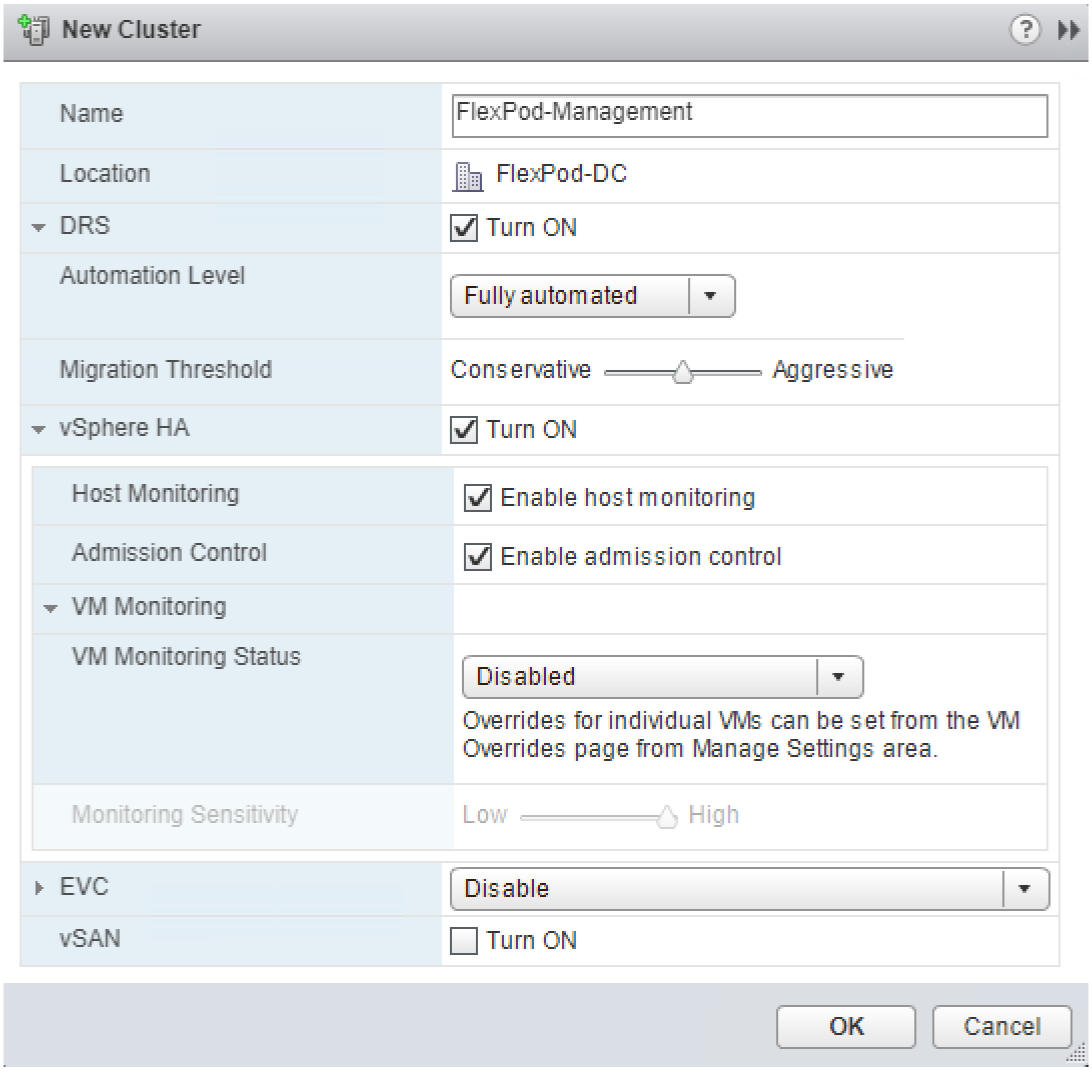

New Cluster

Name	FlexPod-Management
Location	FlexPod-DC
DRS	☑ Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative — Aggressive
vSphere HA	☑ Turn ON
Host Monitoring	☑ Enable host monitoring
Admission Control	☑ Enable admission control
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low — High
EVC	Disable
vSAN	□ Turn ON

11. Click OK to create the new cluster.

12. Expand “FlexPod-DC”.

13. Right-click “FlexPod-Management” and select Settings.

14. Select Configuration > General in the list and select Edit to the right of General.

15. Select Datastore specified by host and click OK.

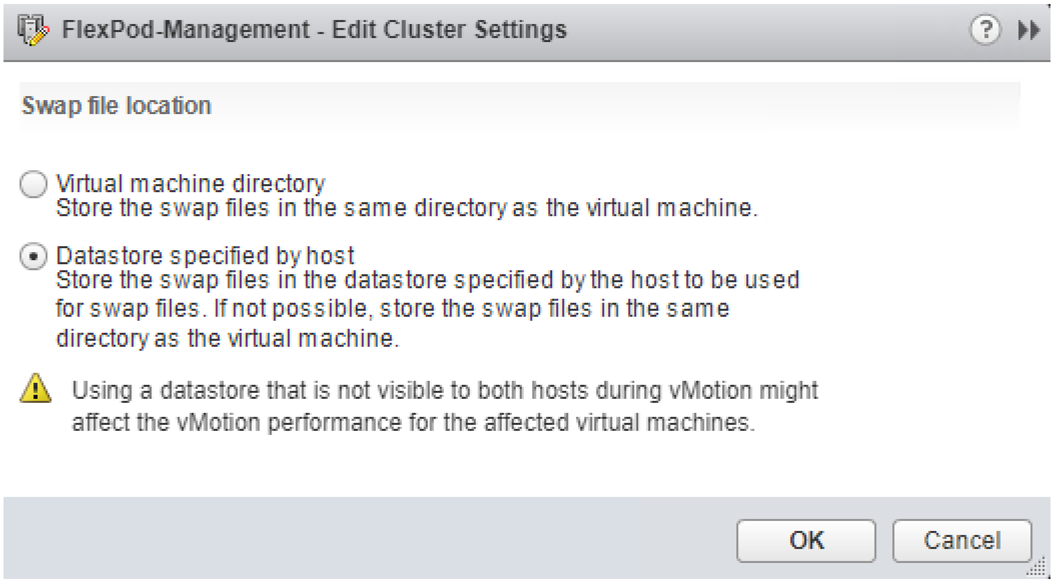

FlexPod-Management - Edit Cluster Settings

Swap file location

○ Virtual machine directory
Store the swap files in the same directory as the virtual machine.

● Datastore specified by host
Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine.

Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

16. Right-click “FlexPod-Management” and click Add Host.

17. In the Host field, enter either the IP address or the FQDN name of one of the VMware ESXi hosts. Click Next.

18. Type root as the user name and the root password. Click Next to continue.

19. Click Yes to accept the certificate.

20. Review the host details and click Next to continue.

21. Assign a license or leave in evaluation mode and click Next to continue.

22. Click Next to continue.

23. Click Next to continue.

24. Review the configuration parameters and click Finish to add the host.

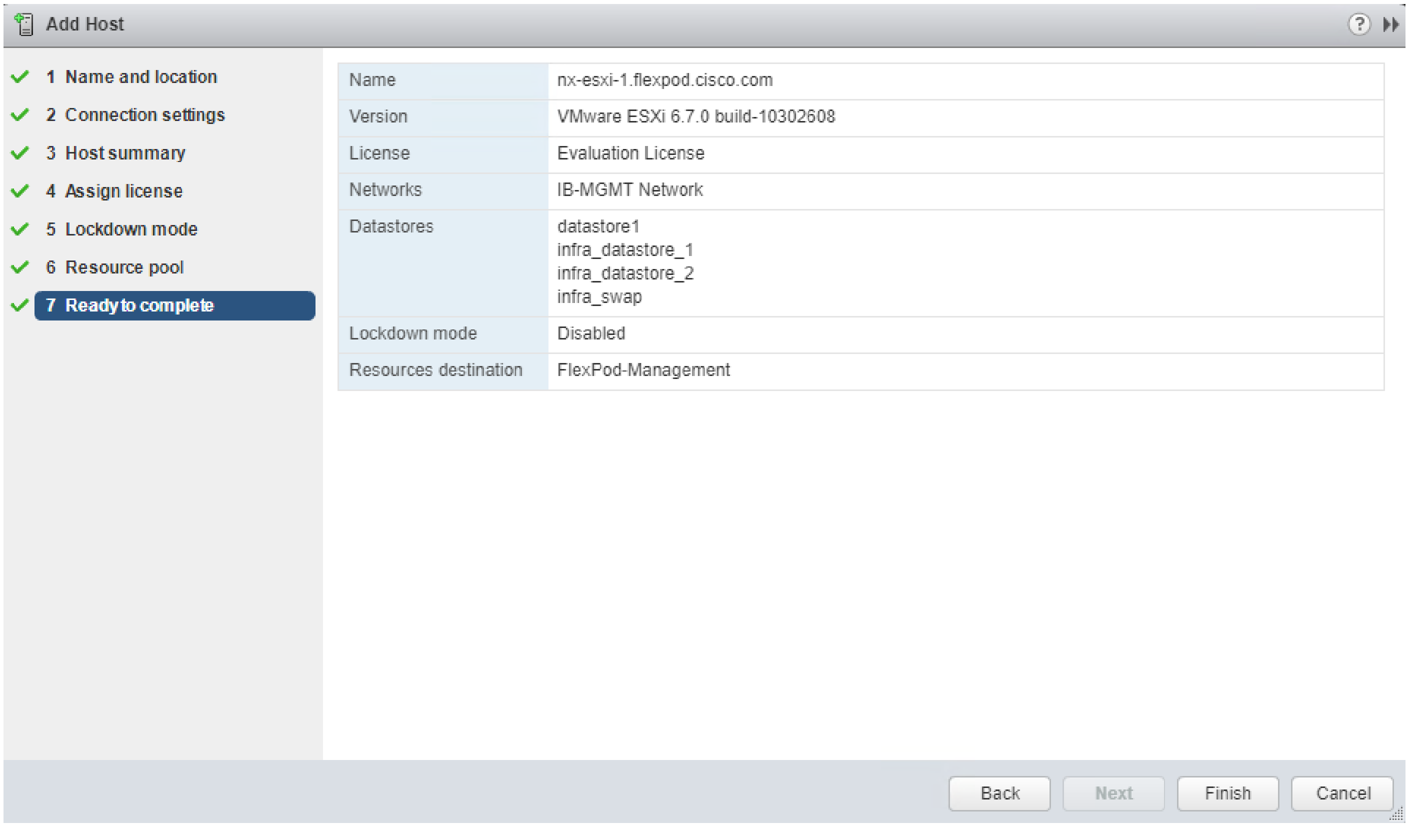

Name	nx-esxi-1.flexpod.cisco.com
Version	VMware ESXi 6.7.0 build-10302608
License	Evaluation License
Networks	IB-MGMT Network
Datastores	datastore1 infra_datastore_1 infra_datastore_2 infra_swap
Lockdown mode	Disabled
Resources destination	FlexPod-Management

25. Repeat steps 16 to 24 to add the remaining VMware ESXi hosts to the cluster.

Two VMware ESXi hosts will be added to the cluster.

26. In the list, right-click the first ESXi host and select Settings.

27. In the center pane under Virtual Machines, select Swap file location.

28. On the right, click Edit.

29. Select the infra_swap datastore and click OK.

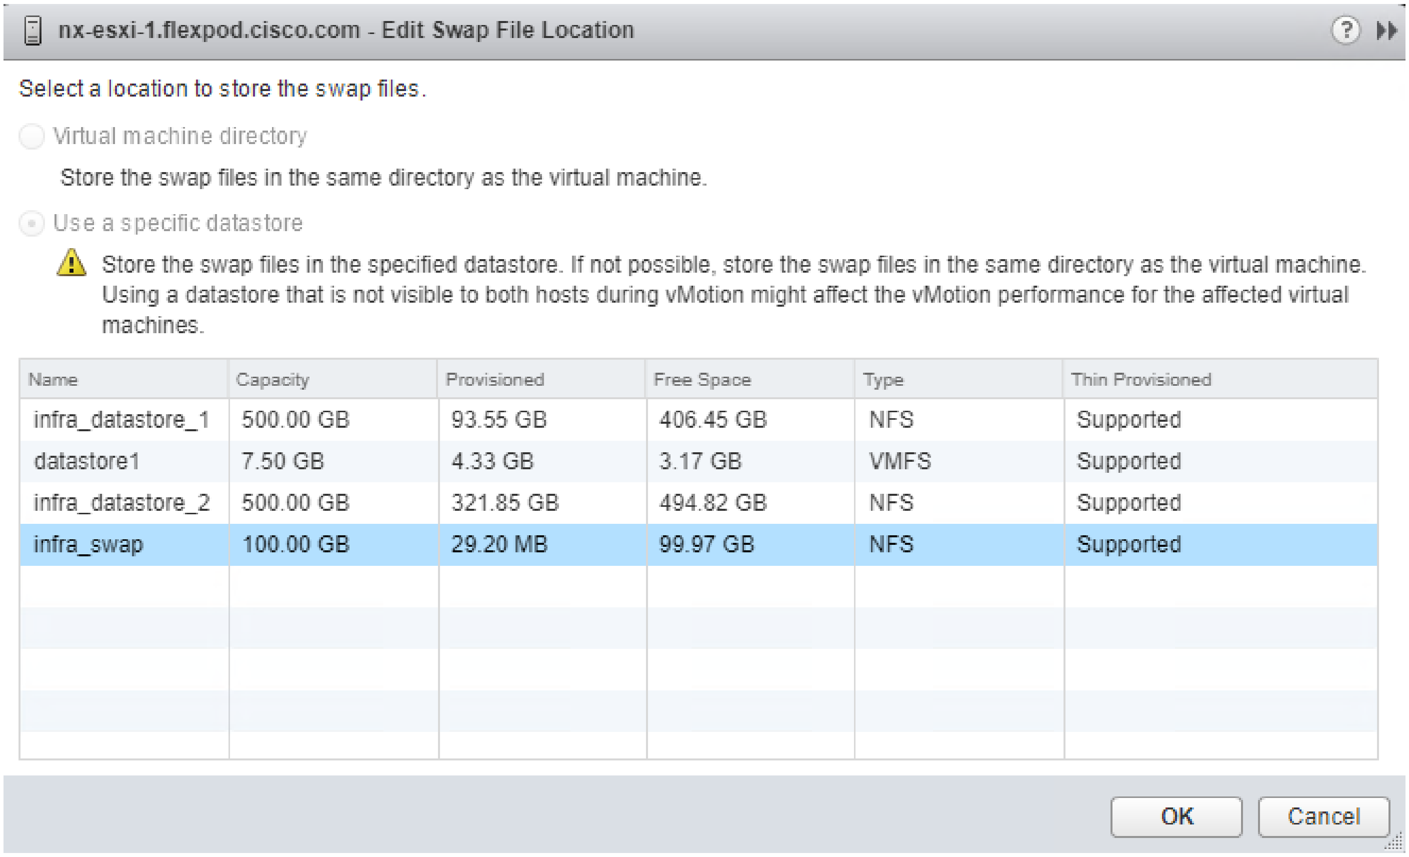

nx-esxi-1.flexpod.cisco.com - Edit Swap File Location

Select a location to store the swap files.

○ Virtual machine directory

Store the swap files in the same directory as the virtual machine.

● Use a specific datastore

⚠ Store the swap files in the specified datastore. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

Name	Capacity	Provisioned	Free Space	Type	Thin Provisioned
infra_datastore_1	500.00 GB	93.55 GB	406.45 GB	NFS	Supported
datastore1	7.50 GB	4.33 GB	3.17 GB	VMFS	Supported
infra_datastore_2	500.00 GB	321.85 GB	494.82 GB	NFS	Supported
infra_swap	100.00 GB	29.20 MB	99.97 GB	NFS	Supported

30. Repeat this process to configure the Swap file location on the second ESXi host.

Add AD User Authentication to vCenter (Optional)

If an AD Infrastructure is set up in this FlexPod environment, you can setup in AD and authenticate from vCenter.

To add an AD user authentication to the vCenter, follow these steps:

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).
2. Connect to https://<vcenter-ip> and select Log into vSphere Web Client.
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. Navigate to Home. In the center pane, select System Configuration under Administration.
5. Select Nodes and under Nodes select the vCenter.
6. In the center pane, select the Manage tab, and within the Settings select Active Directory and click Join.
7. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Click OK.
8. Right-click the vCenter and select Reboot.

9. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.

10. Log back into the vCenter Web Client.

11. Navigate to Home. In the center pane, select System Configuration under Administration.

12. Select Nodes and under Nodes select the vCenter.

13. In the center pane under the Manage tab, select Active Directory. Make sure your Active Directory Domain is listed.

14. Navigate back to the vCenter Home.

15. In the center pane under Administration, select Roles.

16. Under Single Sign-On, select Configuration.

17. In the center pane, select the Identity Sources tab.

18. Click the green + sign to add an Identity Source.

19. Select the Active Directory (Integrated Windows Authentication) Identity source type and click Next.

20. Your AD domain name should be filled in. Leave Use machine account selected and click Next.

21. Click Finish to complete adding the AD Domain.

22. Your AD domain should now appear in the Identity Sources list.

23. Under Access Control, select Global Permissions.

24. In the center pane, click the green + sign to add a Global Permission.

25. In the Global Permission Root – Add Permission window, click Add.

26. In the Select Users/Groups window, select your AD domain for the Domain.

27. Under Users and Groups, select either the FlexPod Admin user or the Domain Admins group.

The FlexPod Admin user was created in the Domain Admins group. The selection here depends on whether the FlexPod Admin user will be the only user used in this FlexPod or you would like to add other users later. By selecting the Domain Admins group, any user placed in that group in the AD domain will be able to login to vCenter as an Administrator.

28. Click Add, then click Check names to verify correctness of the names. Click OK to acknowledge the correctness of the names.

29. Click OK to add the selected User or Group.

30. Verify the added User or Group is listed under Users and Groups and the Administrator role is assigned.

31. Click OK.

32. Log out and log back into the vCenter Web Client as the FlexPod Admin user. You will need to add the domain name to the user, for example, flexadmin@domain.

33. If implementing iSCSI boot, complete the VMware ESXI Dump Collector steps in the iSCSI Addition Appendix.

FlexPod VMware vSphere Distributed Switch (vDS)

This section provides detailed procedures for installing the VMware vDS on the FlexPod ESXi Management Hosts.

In the Cisco UCS setup section of this document two sets of vNICs were setup. The vmnic ports associated with the Infra-vDS-A and B vNICs will be placed on the VMware vDS in this procedure. The vMotion interfaces will be placed on the vDS.

An IB-MGMT, vMotion, and a VM-Traffic port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would need to have the corresponding VLANs added to the Cisco UCS LAN cloud, to the Cisco UCS Infra-vDS-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC peer-link interfaces on the switches.

In this document, the infrastructure ESXi management VMkernel ports, the vCenter management interface, and the infrastructure NFS VMkernel ports are left on vSwitch0 to facilitate bringing the virtual environment back up in the event it needs to be completely shut down. The vMotion VMkernel ports are moved to the vDS to allow QoS marking of vMotion to be done at the VLAN level in the vDS if vMotion needs to have QoS policies applied in the future. The vMotion port group is also pinned to UCS fabric B. Pinning is better done in a vDS to ensure consistency across all ESXi hosts.

Configure the VMware vDS in vCenter

VMware vSphere Web Client

To configure the vDS, follow these steps:

1. After logging into the VMware vSphere Web Client, select Networking under Home.

2. Right-click the FlexPod-DC datacenter and select Distributed Switch > New Distributed Switch.

3. Give the Distributed Switch a descriptive name and click Next.

4. Make sure Distributed switch: 6.6.0 is selected and click Next.

5. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter VM-Traffic for the Port group name. Click Next.

6. Review the information and click Finish to complete creating the vDS.

7. Expand the FlexPod-DC datacenter and the newly created vDS. Select the newly created vDS.

8. Select the VM-Traffic port group. In the center pane, select the Edit distributed port group settings icon . The Edit button can be used to change the number of ports in the port group to a number larger than the default of 8. All of the other properties of the port group can also be changed under Edit.

9. Select VLAN.

10. Select VLAN for VLAN type and enter the VM-Traffic VLAN ID. Click OK.

If the port group is used by an application that could be affected by DDTS id: CSCvp49398 which impacts C-Series servers connected directly to Cisco UCS 6454 Fabric Interconnects using 25GbE and Cisco VIC 1455 or 1457, consider enabling Beacon Probing for Network failure detection on this port group. Due to this bug, data loss for up to 75-90 seconds can be observed when the Fabric Interconnect being actively used to forward server traffic is rebooted. This data loss issue is not observed on B-Series servers and this bug does not impact the traffic forwarding under normal operation. Since a Fabric Interconnect reboot is required when upgrading the UCS system, customers should plan their upgrade during off-peak hours to minimize traffic disruption. Lab testing has shown that this bug does not affect VMware vCenter to ESXi host communication, iSCSI and NFS storage protocols, and vMotion. Beacon Probing should only be enabled on port groups that are used by affected applications. If only Cisco UCS B-Series servers are being used in your implementation, do not enable Beacon Probing.

11. Select the vDS . Click the Edit distributed switch settings icon on the right .

12. in the Edit Settings window, select Advanced.

13. Change the MTU to 9000. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.

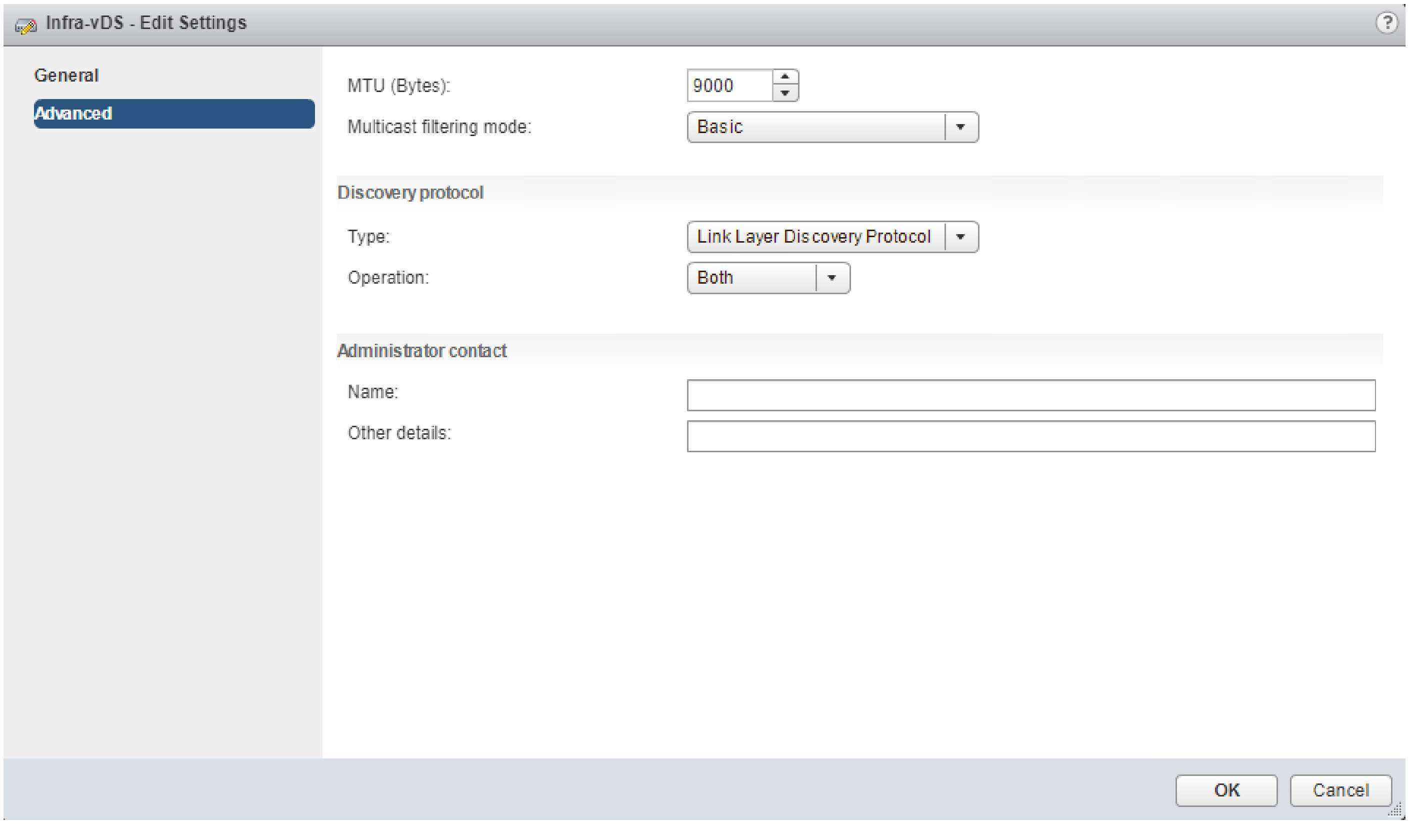

Infra-vDS - Edit Settings

General
Advanced

MTU (Bytes): 9000
Multicast filtering mode: Basic

Discovery protocol

Type: Link Layer Discovery Protocol
Operation: Both

Administrator contact

Name:
Other details:

14. For the first port group used for vMotion, right-click the vDS, select Distributed Port Group, and select New Distributed Port Group.

15. Enter vMotion as the name and click Next.

16. Set the VLAN type to VLAN, enter the VLAN ID used for vMotion, click the Customize default policies configuration check box, and click Next.

17. Leave the Security options set to Reject and click Next.

18. Leave the Ingress and Egress traffic shaping options as Disabled and click Next.

19. Select Uplink 1 from the list of Active uplinks and click the down arrow icon twice to place Uplink 1 in the list of Standby uplinks. This will pin all vMotion traffic to UCS Fabric Interconnect B except when a failure occurs.

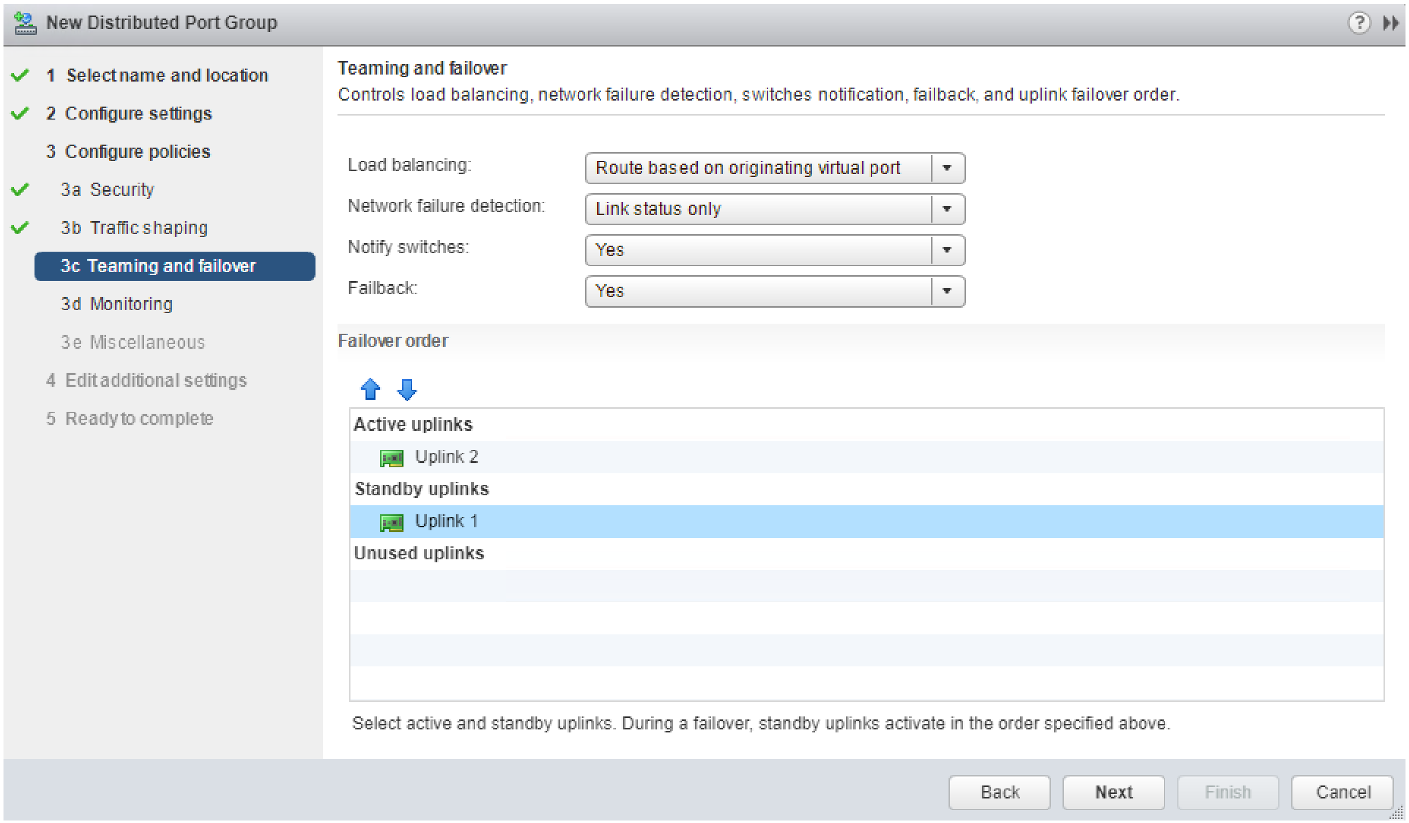

Teaming and failover

Controls load balancing, network failure detection, switches notification, failback, and uplink failover order.

Load balancing: Route based on originating virtual port

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active uplinks
- Uplink 2

Standby uplinks
- Uplink 1

Unused uplinks

Select active and standby uplinks. During a failover, standby uplinks activate in the order specified above.

20. Click Next.

21. Leave NetFlow disabled and click Next.

22. Leave Block all ports set as No and click Next.

23. Leave the additional settings dialogue options as they are shown and click Next.

24. Confirm the options and click Finish to create the port group.

25. For the second port group used for Infrastructure In-Band Management, , right-click the vDS, select Distributed Port Group, and select New Distributed Port Group.

26. Enter IB-MGMT as the name and click Next.

27. Set the VLAN type to VLAN, enter the VLAN ID used for In-Band Management, do not select the Customize default policies configuration check box, and click Next.

28. Confirm the options and click Finish to create the port group.

29. Right-click the vDS and select Add and Manage Hosts.

30. Make sure Add hosts is selected and click Next.

31. Click the green + sign to add New hosts. Select the two FlexPod Management hosts and click OK. Click Next.

32. Leave Manage physical adapters and Manage VMkernel adapters selected. Do not select Migrate virtual machine networking and click Next.

33. Select vmnic2 on the first host and click Assign uplink. Select Uplink 1 and click OK. Select vmnic3 on the first host and click Assign uplink. Select Uplink 2 and click OK. Repeat this process to assign uplinks from both hosts to the vDS.

It is important to assign the uplinks as shown in the vDS. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.

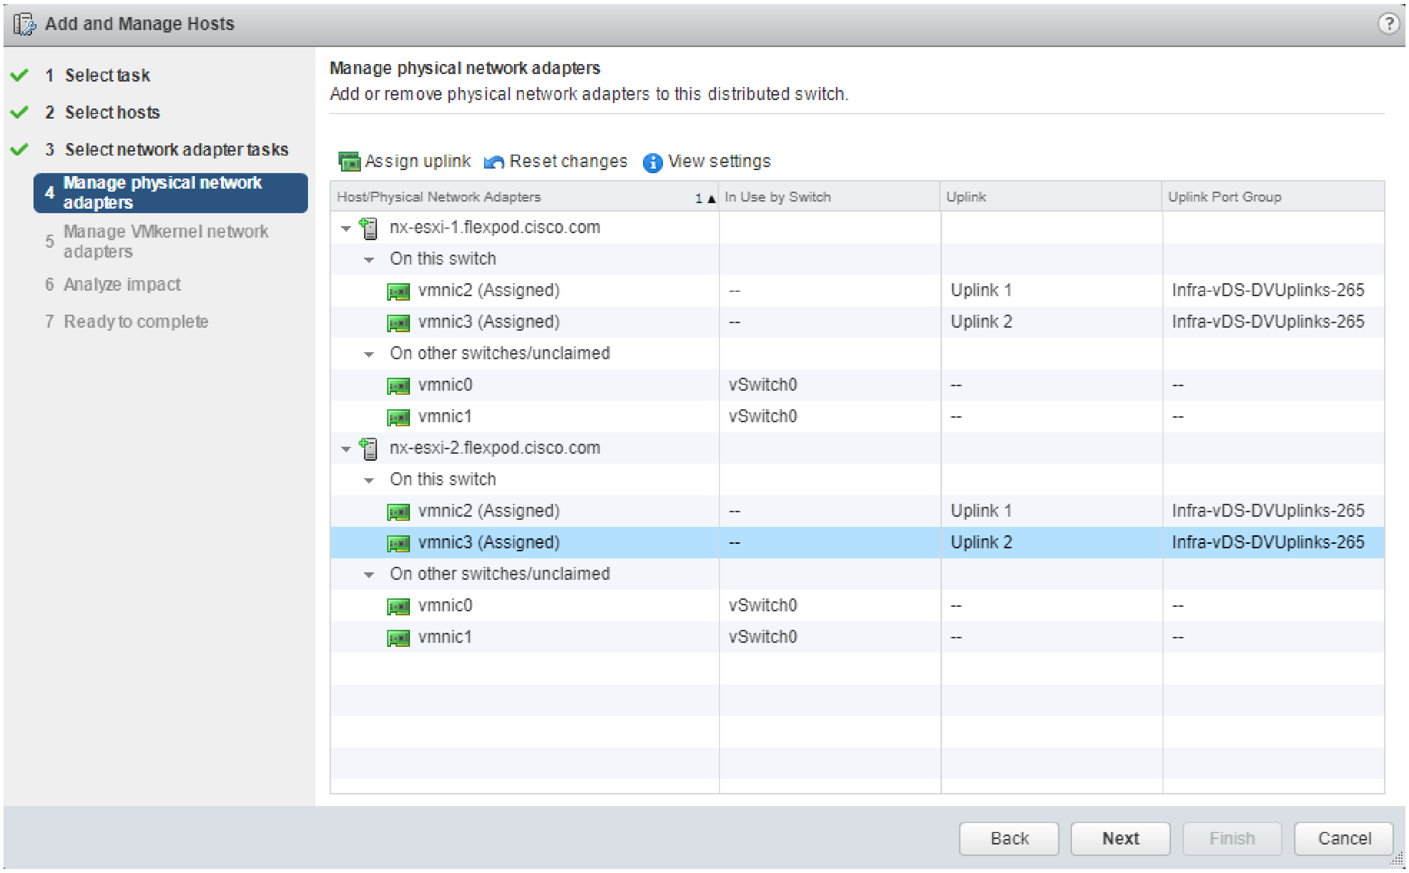

Manage physical network adapters

Add or remove physical network adapters to this distributed switch.

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
nx-esxi-1.flexpod.cisco.com			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	Infra-vDS-DVUplinks-265
vmnic3 (Assigned)	--	Uplink 2	Infra-vDS-DVUplinks-265
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
nx-esxi-2.flexpod.cisco.com			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	Infra-vDS-DVUplinks-265
vmnic3 (Assigned)	--	Uplink 2	Infra-vDS-DVUplinks-265
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--

34. Click Next.

35. Select vmk1 (VMkernel vMotion) on the first host and click Assign port group.

36. Select the vMotion destination port group and click OK.

37. Do not assign the other VMkernel ports.

38. Repeat this process for the second ESXi host.

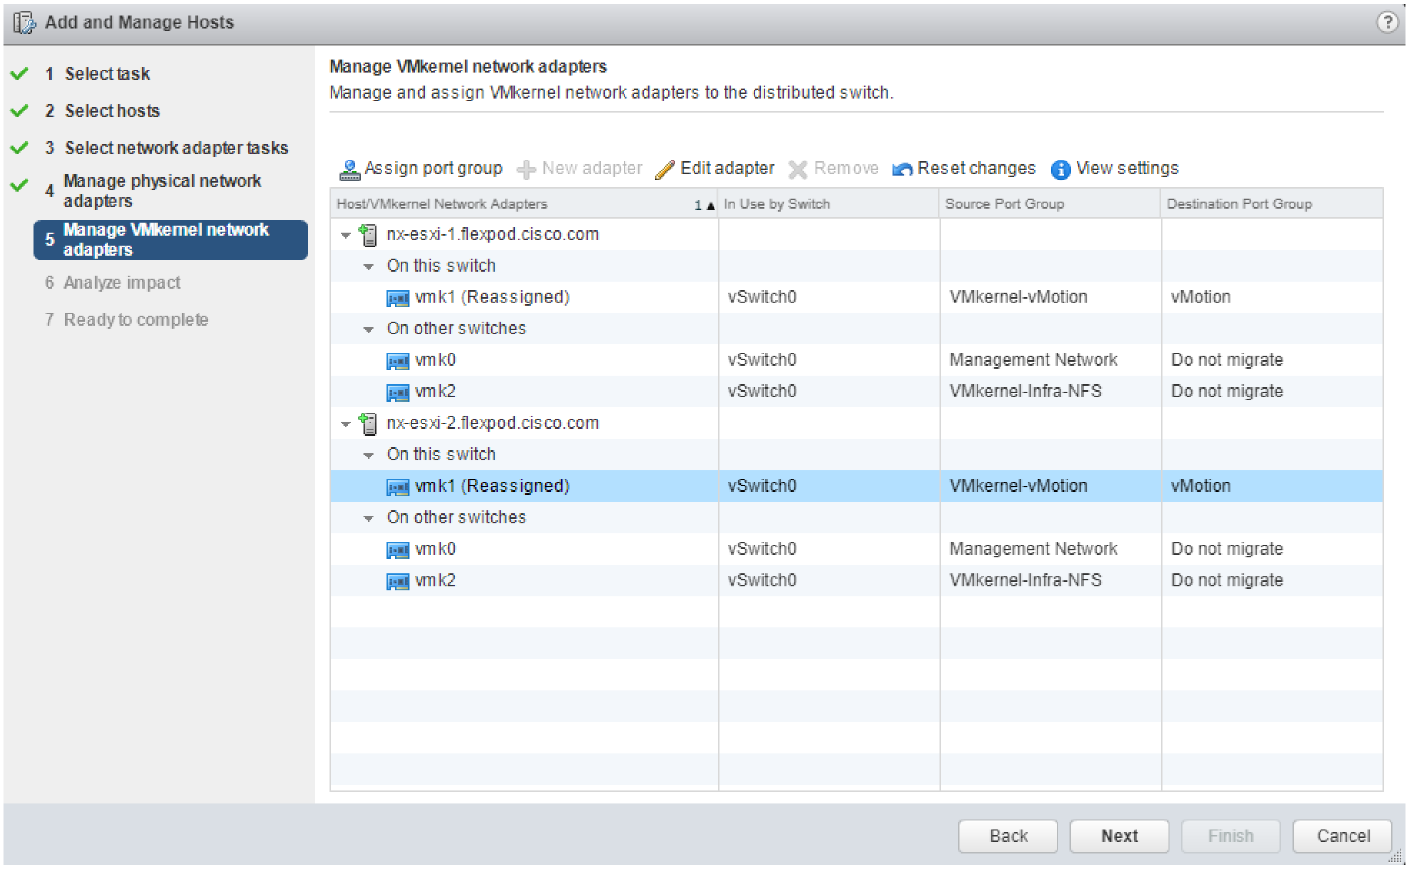

Add and Manage Hosts

✓ 1 Select task
✓ 2 Select hosts
✓ 3 Select network adapter tasks
✓ 4 Manage physical network adapters
5 Manage VMkernel network adapters
6 Analyze impact
7 Ready to complete

Manage VMkernel network adapters

Manage and assign VMkernel network adapters to the distributed switch.

Host/VMkernel Network Adapters 1 ▲	In Use by Switch	Source Port Group	Destination Port Group
nx-esxi-1.flexpod.cisco.com			
On this switch			
vmk1 (Reassigned)	vSwitch0	VMkernel-vMotion	vMotion
On other switches			
vmk0	vSwitch0	Management Network	Do not migrate
vmk2	vSwitch0	VMkernel-Infra-NFS	Do not migrate
nx-esxi-2.flexpod.cisco.com			
On this switch			
vmk1 (Reassigned)	vSwitch0	VMkernel-vMotion	vMotion
On other switches			
vmk0	vSwitch0	Management Network	Do not migrate
vmk2	vSwitch0	VMkernel-Infra-NFS	Do not migrate

39. Confirm the vMotion VMkernel adapter on each host has a valid and correct Destination Port Group and Click Next.

40. Click Next after confirming there is no impact detected in the Analyze impact screen.

41. Click Finish to complete adding the two ESXi hosts to the vDS.

VMware ESXi 6.7U1 TPM Attestation

If your Cisco UCS servers have Trusted Platform Module (TPM) 2.0 modules installed, the TPM can provide assurance that ESXi has booted with UEFI Secure Boot enabled and using only digitally signed code. In the Cisco UCS section of this document, the secure boot was enabled in the boot policy. A server can boot with UEFI Secure Boot with or without a TPM 2.0 module. If it has a TPM, VMware vCenter can attest that the server booted with UEFI Secure Boot. Follow these steps:

1. If you Cisco UCS servers have TPM 2.0 modules installed, TPM Attestation can be verified in the vSphere HTML5 Client. To get to the HTML5 client from the Web Client, click “Launch vSphere Client (HTML5) in the upper center portion of the Web Client window.

2. From the Hosts and Clusters window in the vSphere Client, select the vCenter . In the center pane, select Monitor > Security. The Attestation status will appear as shown below:

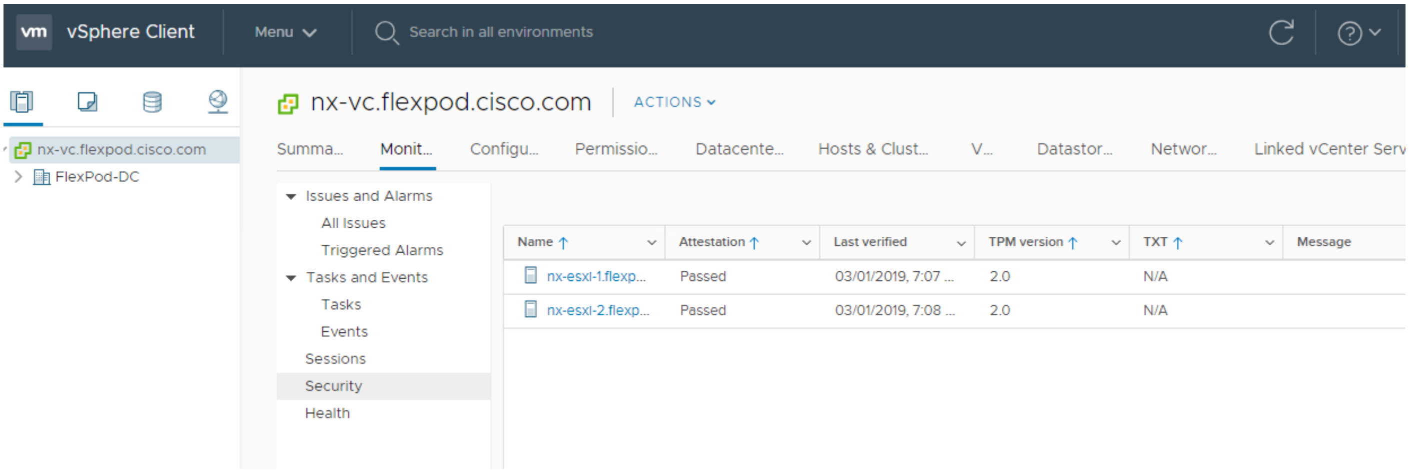

Name	Attestation	Last verified	TPM version	TXT	Message
nx-esxi-1.flexp...	Passed	03/01/2019, 7:07 ...	2.0	N/A	
nx-esxi-2.flexp...	Passed	03/01/2019, 7:08 ...	2.0	N/A	

FlexPod Management Tools Setup

Cisco UCS Manager Plug-in for VMware vSphere Web Client

The Cisco UCS Manager Plug-in for VMware vSphere Web Client allows administration of UCS domains through the VMware vCenter administrative interface. The capabilities of the plug-in include:

- View Cisco UCS physical hierarchy
- View inventory, installed firmware, faults, power and temperature statistics
- Map the ESXi host to the physical server
- Manage firmware for Cisco UCS B and C series servers
- View VIF paths for servers
- Launch the Cisco UCS Manager GUI
- Launch the KVM consoles of UCS servers
- Switch the existing state of the locator LEDs

The installation is only valid for VMware vCenter 5.5 or higher and will require revisions of .NET Framework 4.5 and VMware PowerCLI 5.1 or greater.

Cisco UCS Manager Plug-in Installation

To begin the plug-in installation on a Windows system that meets the previously stated requirements, follow these steps:

1. Download the plugin and registration tool from: https://software.cisco.com/download/home/286282669/type/286282010/release/2.0.4?catid=282558030.
2. Place the downloaded ucs-vcplugin-2.0.4.zip file onto the web server used for hosting the ONTAP software and VMware ESXi ISO.
3. Unzip the Cisco_UCS_Plugin_Registration_Tool_1.2.2.zip and open the executable file within it.
4. Leave Register Plugin selected for the Action, fill in the following vCenter information and click Submit:
   a. IP/Hostname
   b. Username
   c. Password
   d. URL that plugin has been uploaded to

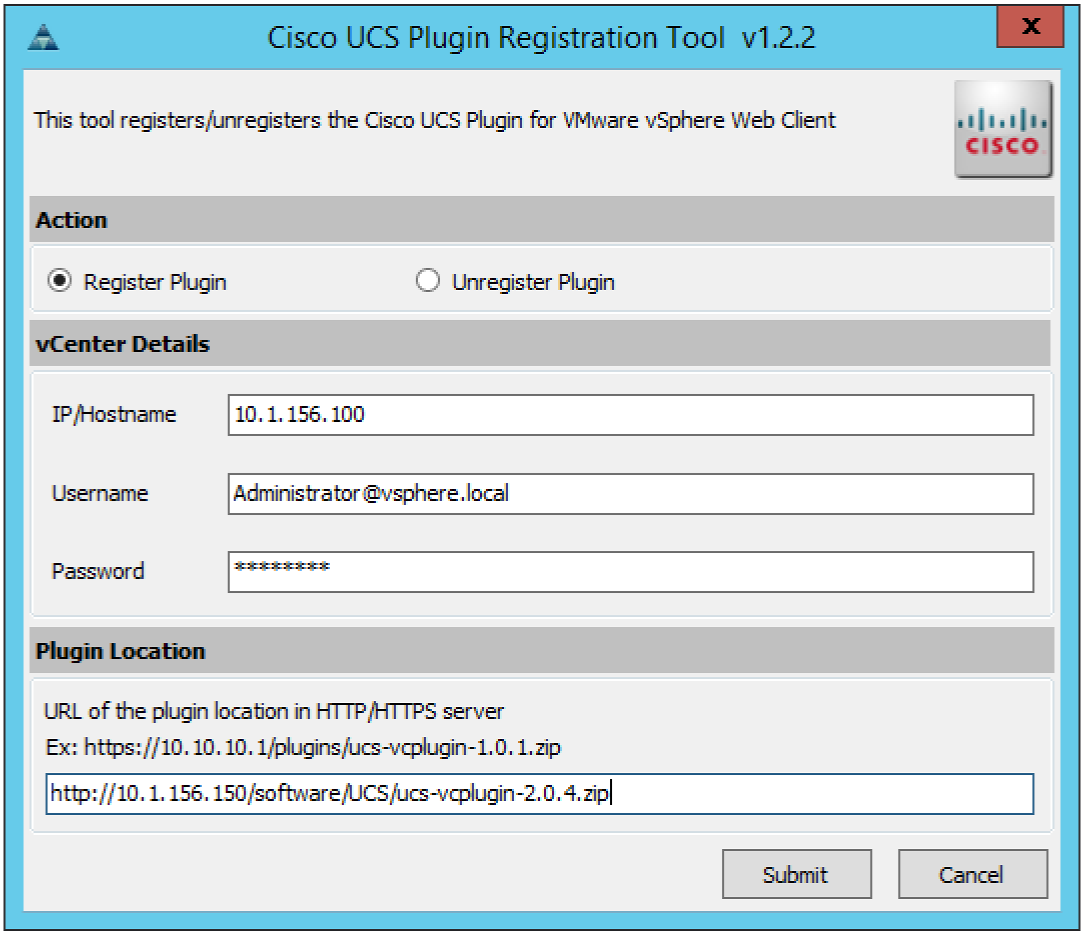

Cisco UCS Plugin Registration Tool v1.2.2

This tool registers/unregisters the Cisco UCS Plugin for VMware vSphere Web Client

Action

● Register Plugin ○ Unregister Plugin

vCenter Details

IP/Hostname	10.1.156.100
Username	Administrator@vsphere.local
Password	*********

Plugin Location

URL of the plugin location in HTTP/HTTPS server

Ex: https://10.10.10.1/plugins/ucs-vcplugin-1.0.1.zip

http://10.1.156.150/software/UCS/ucs-vcplugin-2.0.4.zip

5. A pop-up will appear explaining that ‘allowHttp=true’ will need to be added to the webclient.properties file on the VCSA in the /etc/vmware/vsphere-client directory.

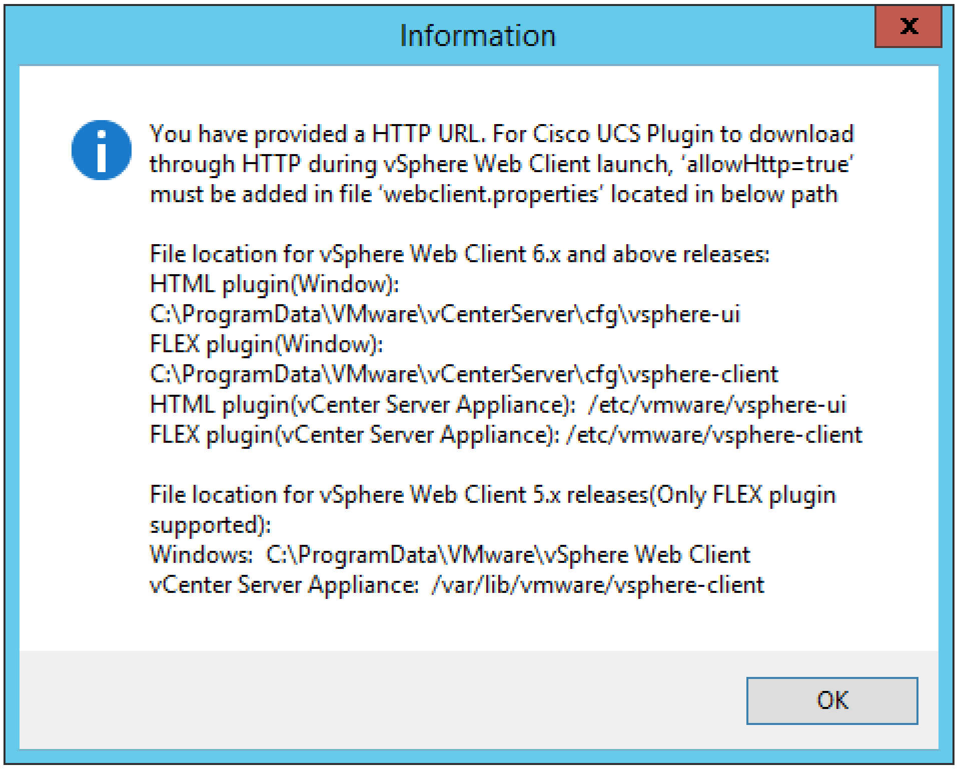

You have provided a HTTP URL. For Cisco UCS Plugin to download through HTTP during vSphere Web Client launch, 'allowHttp=true' must be added in file 'webclient.properties' located in below path

File location for vSphere Web Client 6.x and above releases:
HTML plugin(Window):
C:\ProgramData\VMware\vCenterServer\cfg\vsphere-ui
FLEX plugin(Window):
C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client
HTML plugin(vCenter Server Appliance): /etc/vmware/vsphere-ui
FLEX plugin(vCenter Server Appliance): /etc/vmware/vsphere-client

File location for vSphere Web Client 5.x releases(Only FLEX plugin supported):
Windows: C:\ProgramData\VMware\vSphere Web Client
vCenter Server Appliance: /var/lib/vmware/vsphere-client

6. Take care of this issue after the plugin has been registered, click OK to close the Information dialogue box.

7. Click OK to confirm that the Cisco UCS Plugin registered successfully.

8. To resolve the change needed for the HTTP download of the vSphere Web Client launch, connect to the VCSA with ssh using the root account, open the BASH shell, and type:

```
echo 'allowHttp=true' >> /etc/vmware/vsphere-client/webclient.properties
```

This will add "allowHttp=true" to the end of the webclient.properties file. Make sure to use two greater than symbols ">>" to append to the end of the configuration file, a single greater than symbol will replace the entire pre-existing file with what has been sent with the echo command.

9. Reboot the VCSA.

FlexPod UCS Domain Registration

Registration of the FlexPod UCS Domain can now be performed. The account used will correlate to the permissions allowed to the plugin, admin will be used in our example, but a read only account could be used with the plugin if that was appropriate for the environment.

To register the UCS Domain, follow these steps:

1. Open the vSphere Web Client.

2. Select the Home from the Navigator or drop-down list and select the Cisco UCS icon appearing in the Administration section.

3. Click the Register button and provide the following options in the Register UCS Domain dialogue box that appears:

    a. UCS Hostname/IP

    b. Username

    c. Password

    d. Port (if different than 443)

    e. Leave SSL selected and click the Visible to All users option

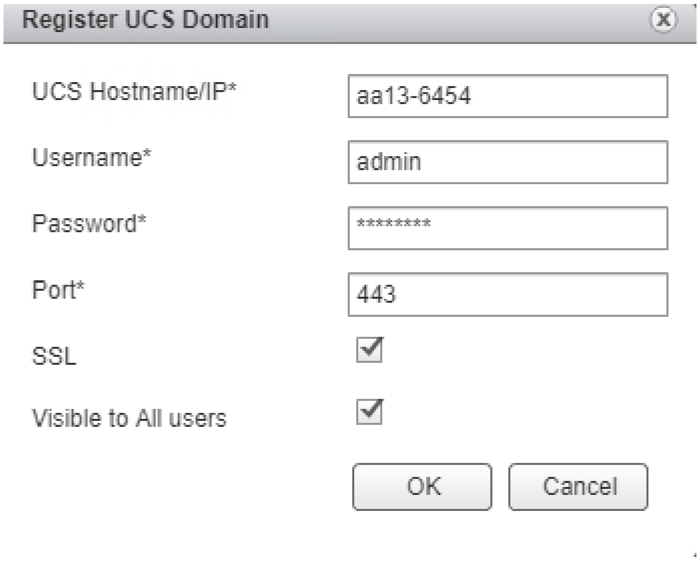

Register UCS Domain

UCS Hostname/IP*	aa13-6454
Username*	admin
Password*	********
Port*	443
SSL	☑
Visible to All users	☑

4. Click OK then click OK again to register the UCS Domain.

Use the Cisco UCS vCenter Plugin

The plugin can now enable the functions described at the start of this section by double-clicking the registered UCS Domain:

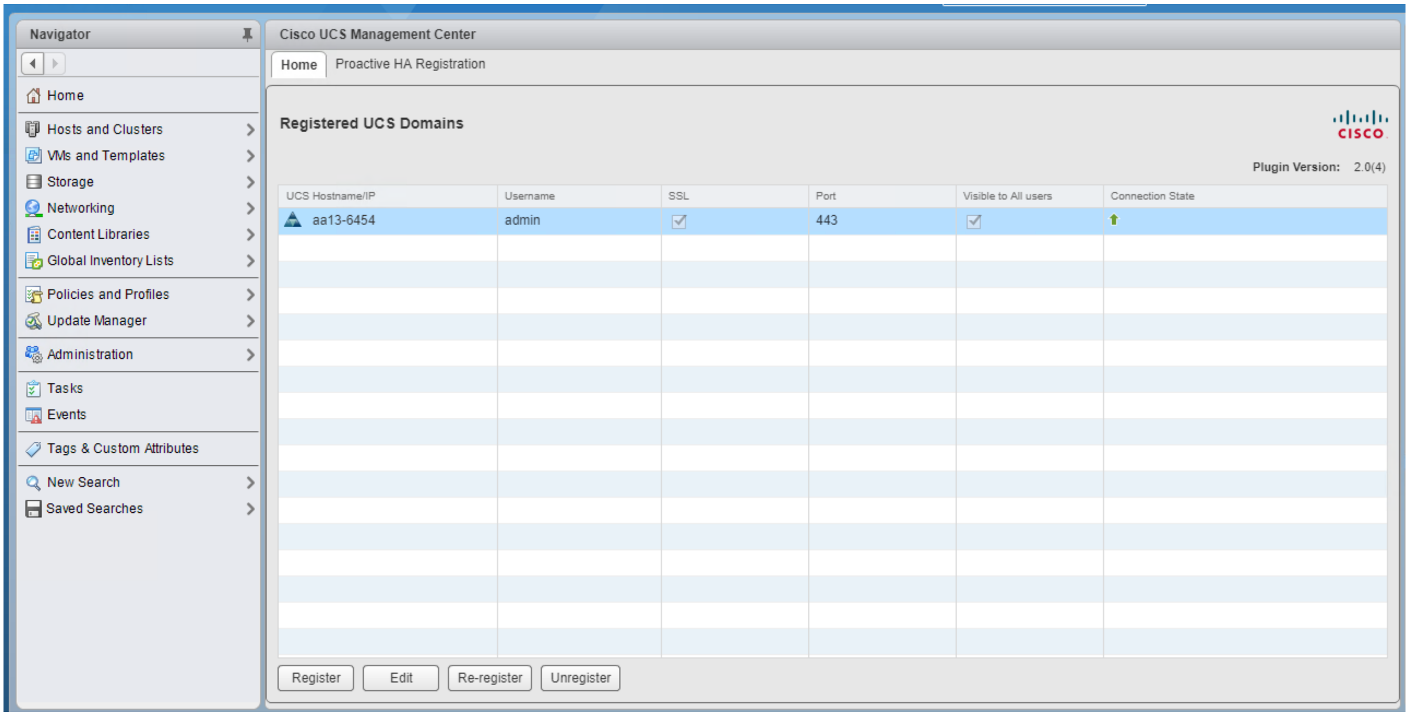

Registered UCS Domains

Plugin Version: 2.0(4)

UCS Hostname/IP	Username	SSL	Port	Visible to All users	Connection State
aa13-6454	admin	✓	443	✓	↑

You can view the components associated to the domain:

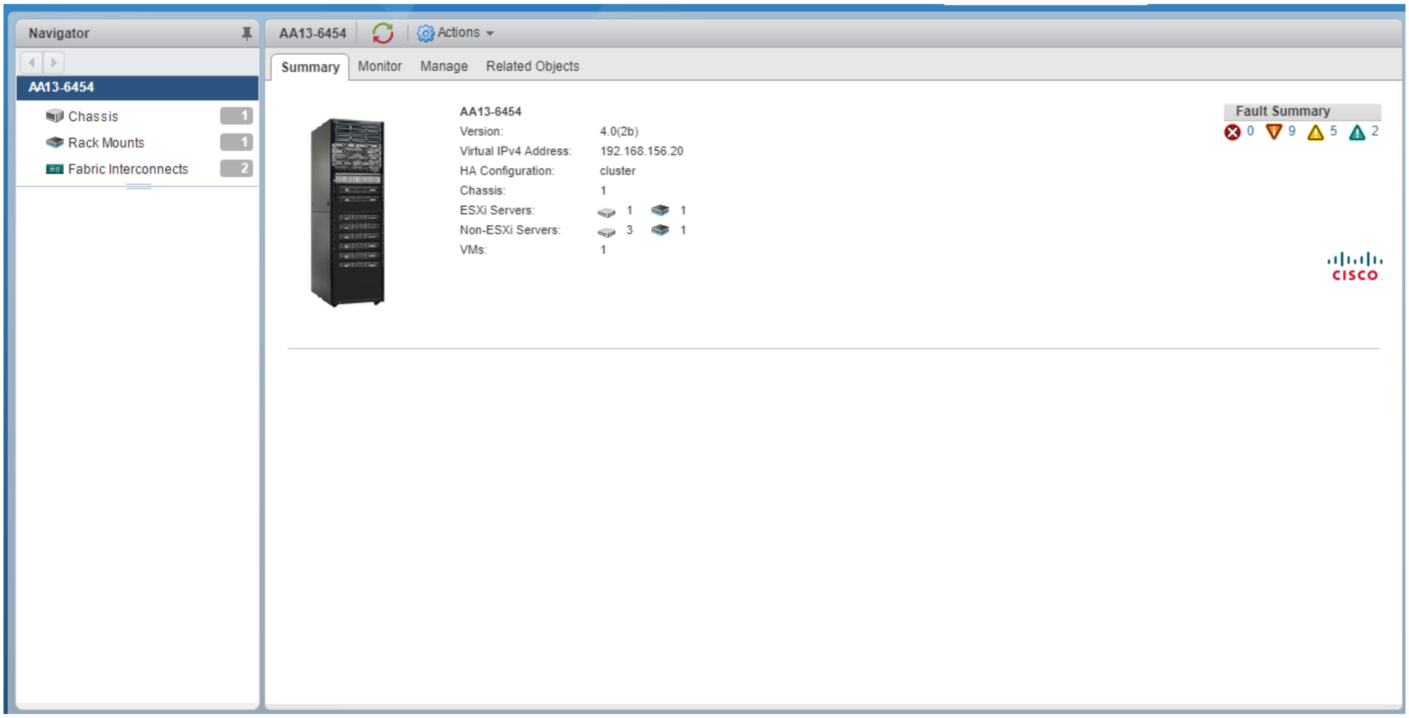

AA13-6454

Version: 4.0(2b)

Virtual IPv4 Address: 192.168.156.20

HA Configuration: cluster

Chassis: 1

ESXi Servers: 1 1

Non-ESXi Servers: 3 1

VMs: 1

Selecting within the chassis or rack mounts will provide a list of ESXi or non-ESXi servers to perform operations on the following:

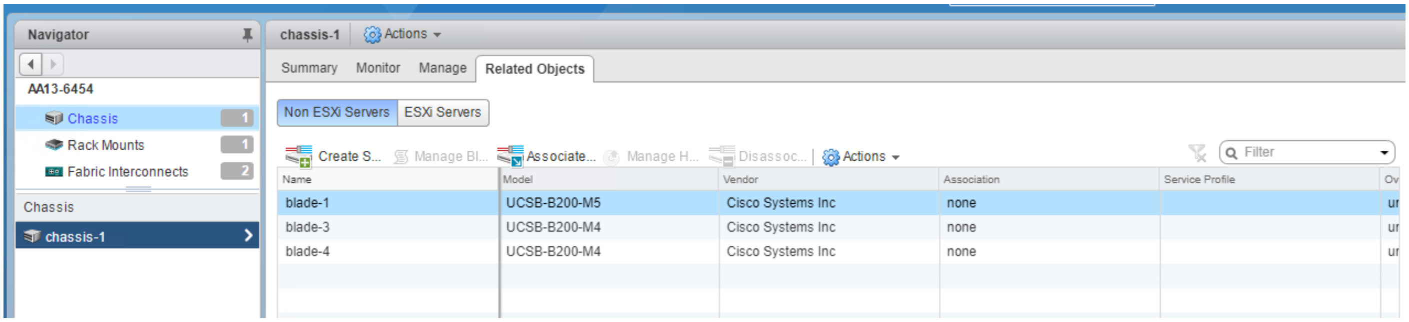

Name	Model	Vendor	Association	Service Profile	Ov
blade-1	UCSB-B200-M5	Cisco Systems Inc	none		ur
blade-3	UCSB-B200-M4	Cisco Systems Inc	none		ur
blade-4	UCSB-B200-M4	Cisco Systems Inc	none		ur

In addition to viewing and working within objects shown in the UCS Plugin's view of the UCS Domain, direct access of UCS functions provided by the plugin can be selected within the drop-down list of hosts registered to vCenter:

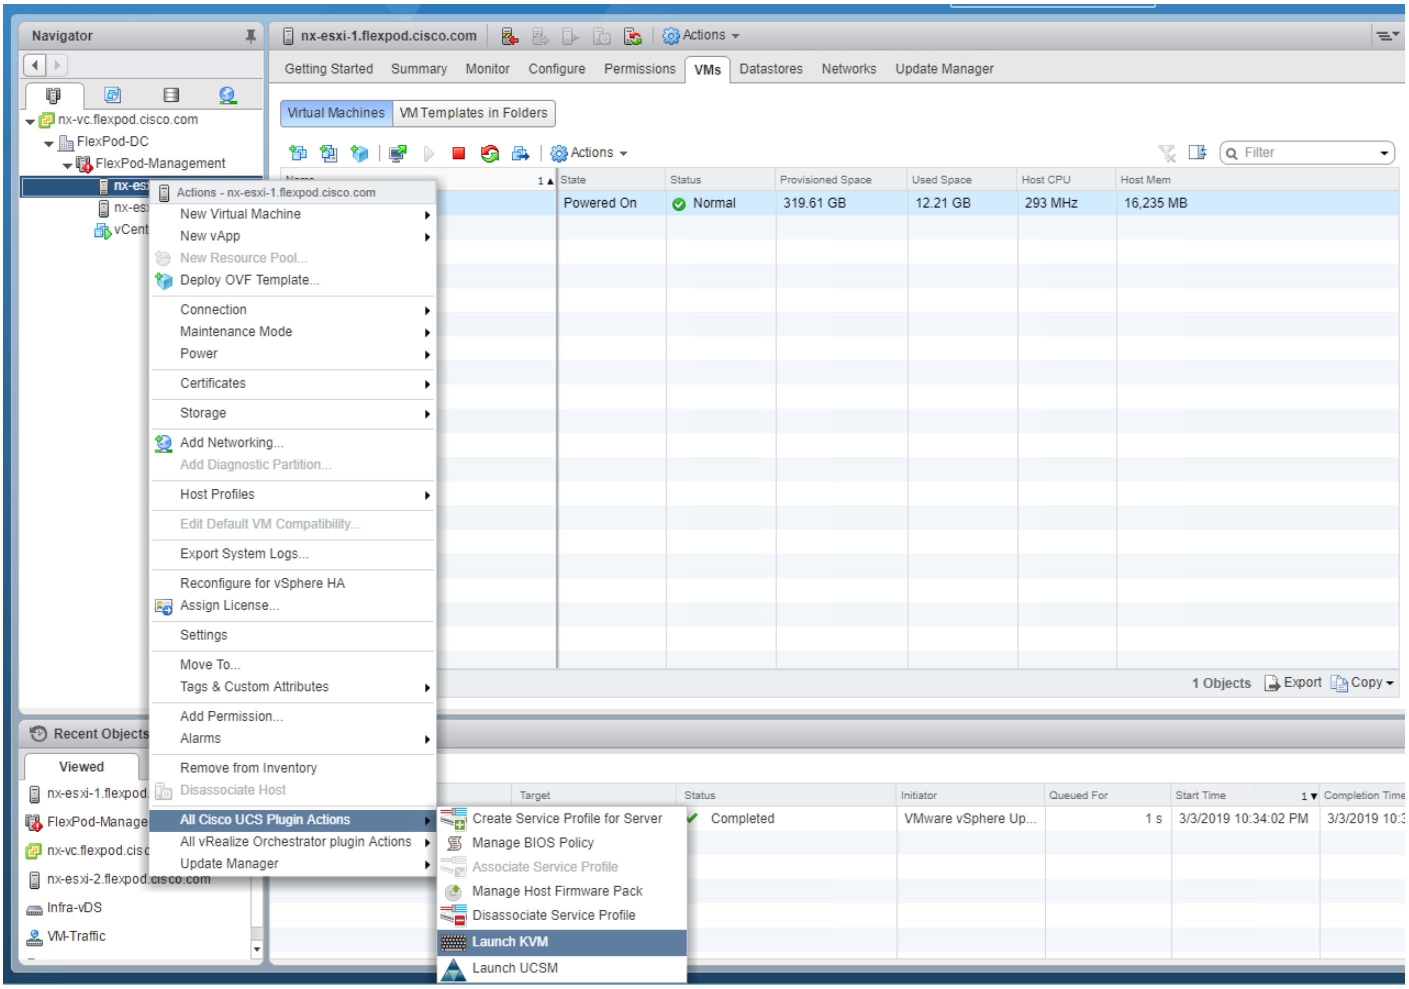

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
	Powered On	Normal	319.61 GB	12.21 GB	293 MHz	16,235 MB

Target	Status	Initiator	Queued For	Start Time	Completion Time
	Completed	VMware vSphere Up...	1 s	3/3/2019 10:34:02 PM	3/3/2019 10:3

For detailed installation instructions and usage information, please refer to the Cisco UCS Manager Plug-in for VMware vSphere Web Client User Guide: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vmware_tools/vCenter/vCenter_Plugin_User_Guide/2x/b_vCenter_2x.html.

NetApp Virtual Storage Console 7.2.1 Deployment Procedure

This section describes the deployment procedures for the NetApp Virtual Storage Console (VSC).

Virtual Storage Console 7.2.1 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.5:

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

Install Virtual Storage Console 7.2.1

The Backup and Recovery capability has been integrated with SnapCenter and requires additional licenses for SnapCenter to perform backup and recovery of virtual machines and applications.

To install the VSC 7.2.1 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

1. Go to vSphere Web Client > Host Cluster > Deploy OVF Template.
2. Browse to the VSC OVF file downloaded from the NetApp Support site.
3. Enter the VM name and select a datacenter or folder in which to deploy. Click Next.
4. Select a host cluster resource in which to deploy OVF. Click Next.
5. Review the details and accept the license agreement.
6. Select Storage.
7. From Select Networks, choose a destination network and click Next.
8. From Customize Template, enter the VSC administrator password, vCenter name or IP address and other configuration details and click Next.

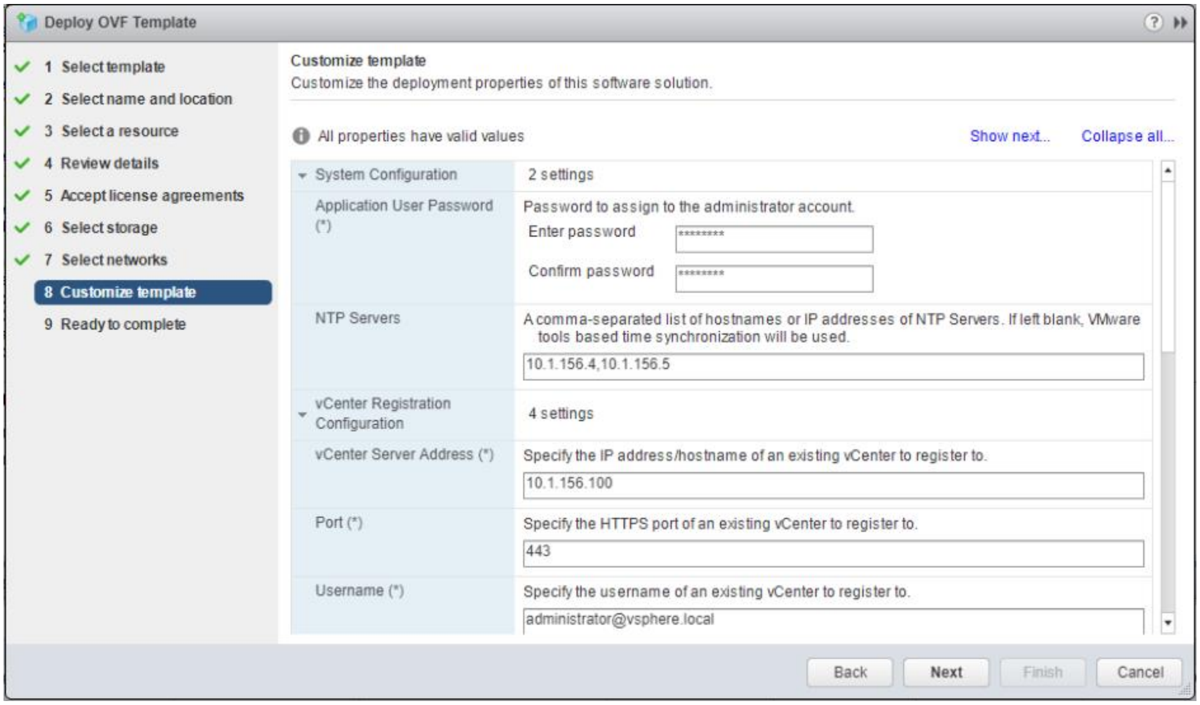

Customize template

Customize the deployment properties of this software solution.

All properties have valid values

System Configuration	2 settings
Application User Password (*)	Password to assign to the administrator account. Enter password ******** Confirm password ********
NTP Servers	A comma-separated list of hostnames or IP addresses of NTP Servers. If left blank, VMware tools based time synchronization will be used. 10.1.156.4,10.1.156.5
vCenter Registration Configuration	4 settings
vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to. 10.1.156.100
Port (*)	Specify the HTTPS port of an existing vCenter to register to. 443
Username (*)	Specify the username of an existing vCenter to register to. administrator@vsphere.local

9. Review the configuration details entered and click Finish to complete the deployment of NetApp-VSC VM.

10. Power on the NetApp-VSC VM and open the VM console.

11. During the NetApp-VSC VM boot process, you see a prompt to install VMware Tools. From vCenter, select NetApp-VSC VM > Guest OS > Install VMware Tools.

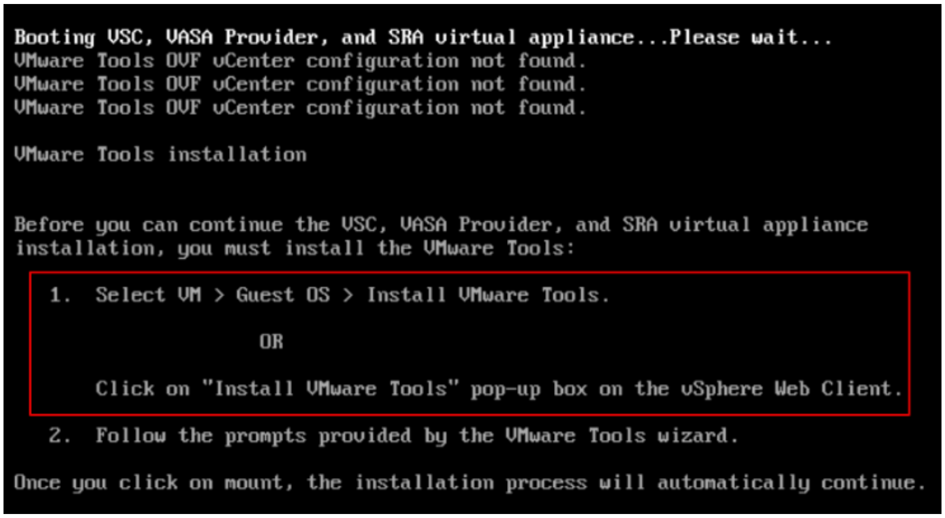

```
Booting VSC, VASA Provider, and SRA virtual appliance...Please wait...
VMware Tools OVF vCenter configuration not found.
VMware Tools OVF vCenter configuration not found.
VMware Tools OVF vCenter configuration not found.

VMware Tools installation


Before you can continue the VSC, VASA Provider, and SRA virtual appliance
installation, you must install the VMware Tools:

   1.  Select VM > Guest OS > Install VMware Tools.

                        OR

       Click on "Install VMware Tools" pop-up box on the vSphere Web Client.

   2.  Follow the prompts provided by the VMware Tools wizard.

Once you click on mount, the installation process will automatically continue.
```

12. Networking configuration and vCenter registration information was provided during the OVF template customization, therefore after NetApp-VSC VM is running, VSC, vSphere API for Storage Awareness (VASA), and VMware Storage Replication Adapter (SRA) are registered with vCenter.

13. Log out of the vCenter Client and log in again. From the Home menu, confirm that the NetApp VSC is installed.

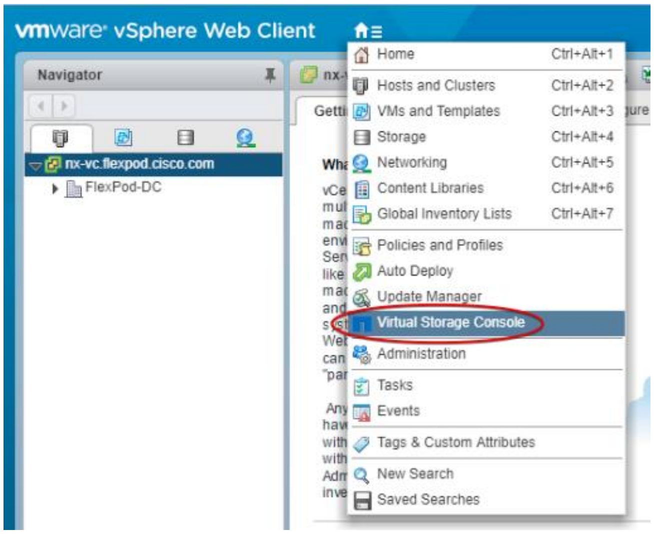

Download the NetApp NFS VAAI Plug-In

To install the NetApp NFS VAAI Plug-In, follow this step:

1. Download the NetApp NFS Plug-In 1.1.2 for VMware .vib file from the NFS Plugin Download page and save it to your local machine or admin host.

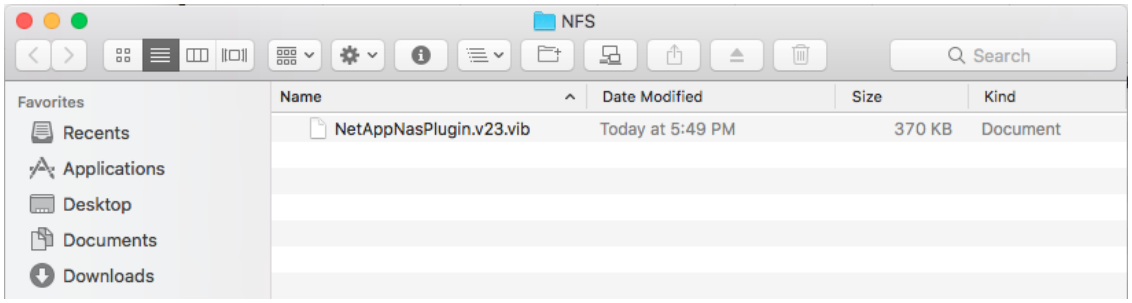

Name	Date Modified	Size	Kind
NetAppNasPlugin.v23.vib	Today at 5:49 PM	370 KB	Document

Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration capability and the Provisioning and Cloning capability, follow these steps:

1. Using the vSphere web client, log in to the vCenter Server as the FlexPod admin user. If the vSphere web client was previously opened, close it and then reopen it.

2. In the Home screen, click the Home tab and click Virtual Storage Console.

You can modify the storage credentials with the vsadmin account or another SVM level account with RBAC privileges.

When using the cluster admin account, add storage from the cluster level.

3. Select Storage Systems. Under the Objects tab, click Actions > Modify.

4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name and the admin password for password. Confirm that Use TLS to Connect to This Storage System is selected. Click OK.

5. Click OK to accept the controller privileges.

6. Wait for the Storage Systems to update. You might need to click Refresh to complete this update.

To add the discovered cluster and SVMs with the cluster admin account, follow these steps:

1. Click Storage Systems and select Add in the tool bar in the Objects tab.

2. Enter the DNS name or IP address of the cluster management LIF.

3. Enter the cluster admin credentials and click OK to add the storage system.

Optimal Storage Settings for ESXi Hosts

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

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

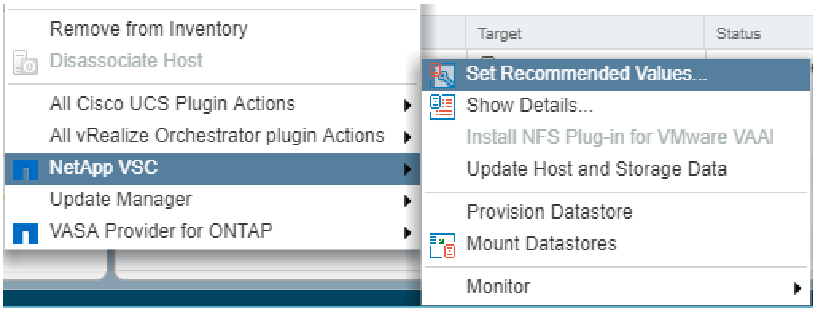
Remove from Inventory
Disassociate Host
All Cisco UCS Plugin Actions
All vRealize Orchestrator plugin Actions
NetApp VSC
Update Manager
VASA Provider for ONTAP
Target
Status
Set Recommended Values...
Show Details...
Install NFS Plug-in for VMware VAAI
Update Host and Storage Data
Provision Datastore
Mount Datastores
Monitor

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

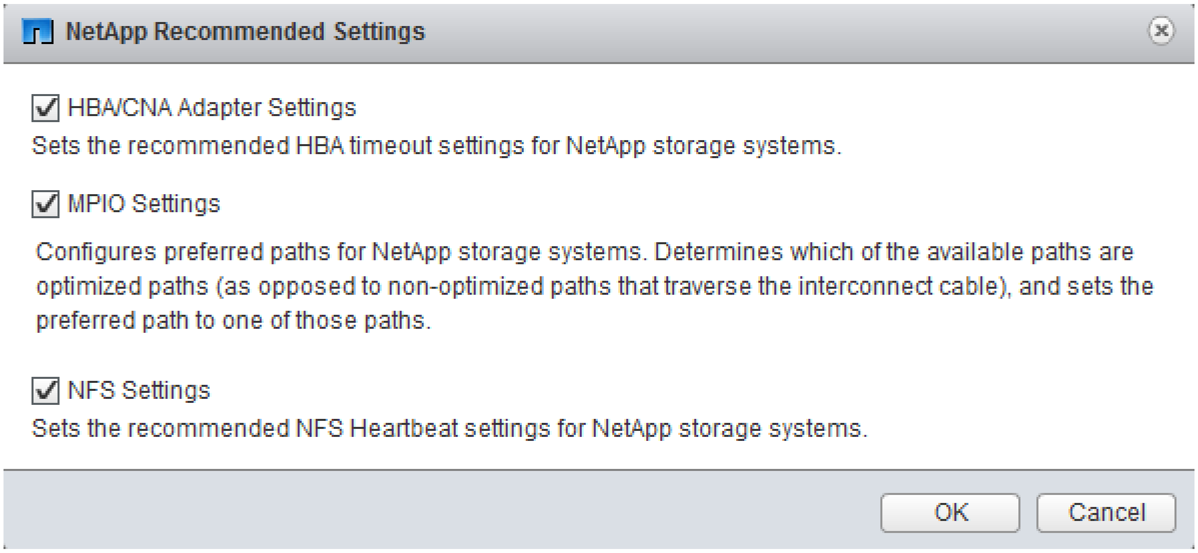

NetApp Recommended Settings

☑ HBA/CNA Adapter Settings

Sets the recommended HBA timeout settings for NetApp storage systems.

☑ MPIO Settings

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

☑ NFS Settings

Sets the recommended NFS Heartbeat settings for NetApp storage systems.

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

3. Click OK.

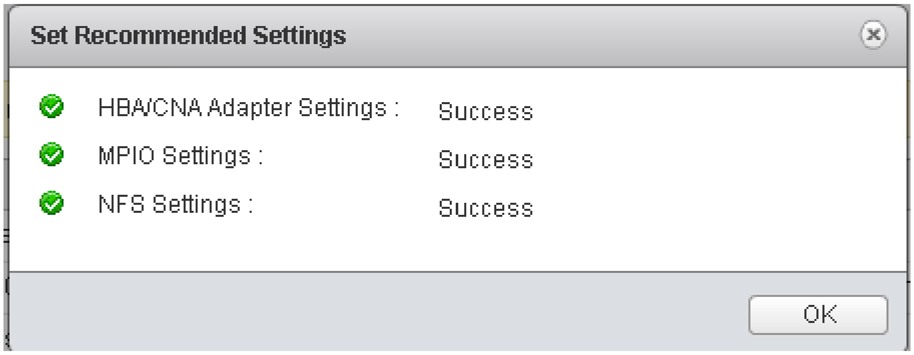

Set Recommended Settings

HBA/CNA Adapter Settings : Success

MPIO Settings : Success

NFS Settings : Success

4. From the Home screen in the vSphere web client, select Virtual Storage Console.

5. Under Virtual Storage Console, select NFS VAAI Tools.

6. Upload the NFS Plug-in by choosing Select File and browsing to the location where the downloaded plug-in is stored.

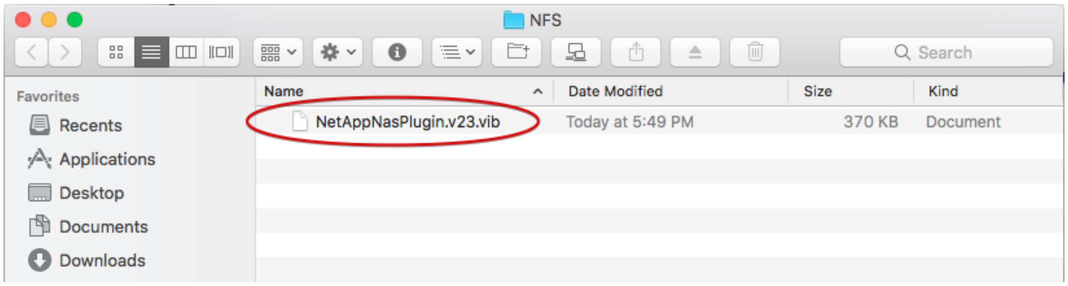

Name	Date Modified	Size	Kind
NetAppNasPlugin.v23.vib	Today at 5:49 PM	370 KB	Document

7. Click Upload to transfer the plug-in to vCenter.

8. Click Install on Host.

9. Select both ESXi hosts and click Install.

10. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.

Virtual Storage Console 7.2.1 Provisioning Datastores

Using VSC, the administrator can provision an NFS, FC or iSCSI datastore and attach it to a single host or multiple hosts in the cluster. The following steps describe provisioning a datastore and attaching it to the cluster.

Provision NFS Datastore

To provision the NFS datastore, follow these steps:

1. From the Home screen of the vSphere web client, right-click the FlexPod-Management cluster and select NetApp VSC > Provision Datastore.

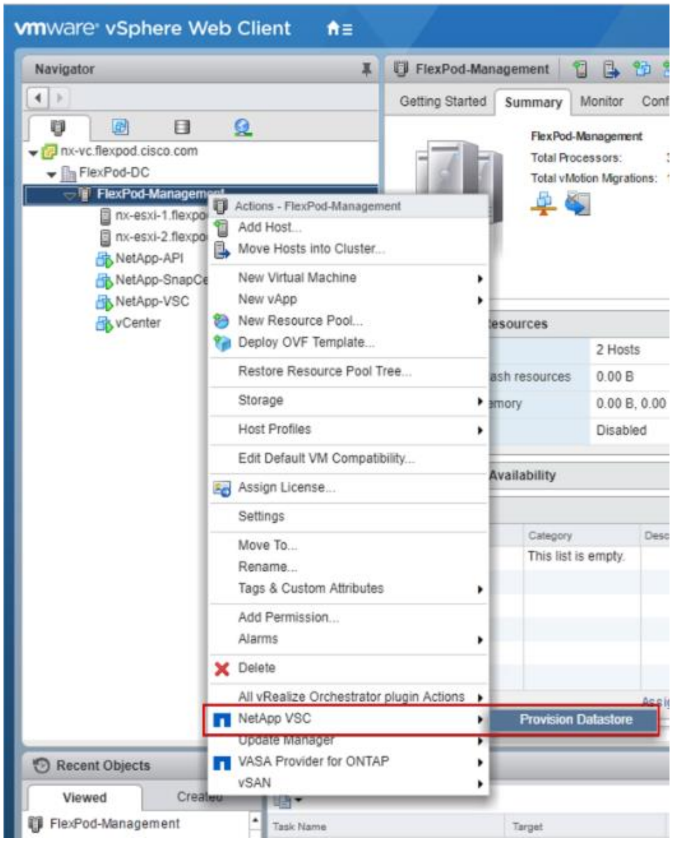

2. Enter the datastore name and select the type as NFS.

3. Click Next.

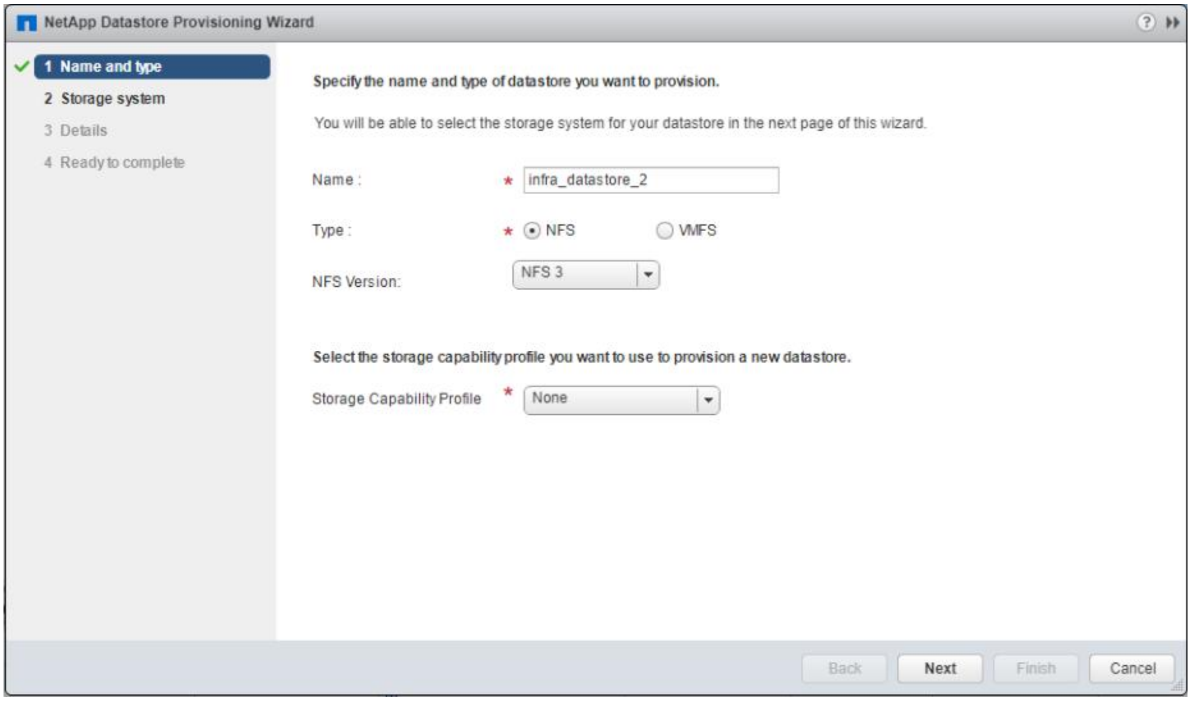

Specify the name and type of datastore you want to provision.

You will be able to select the storage system for your datastore in the next page of this wizard.

Name : * infra_datastore_2

Type : * ● NFS ○ VMFS

NFS Version: NFS 3

Select the storage capability profile you want to use to provision a new datastore.

Storage Capability Profile * None

4. Select the cluster name and the desired SVM to create the datastore. In this example, the Infra-SVM is selected.

5. Click Next.

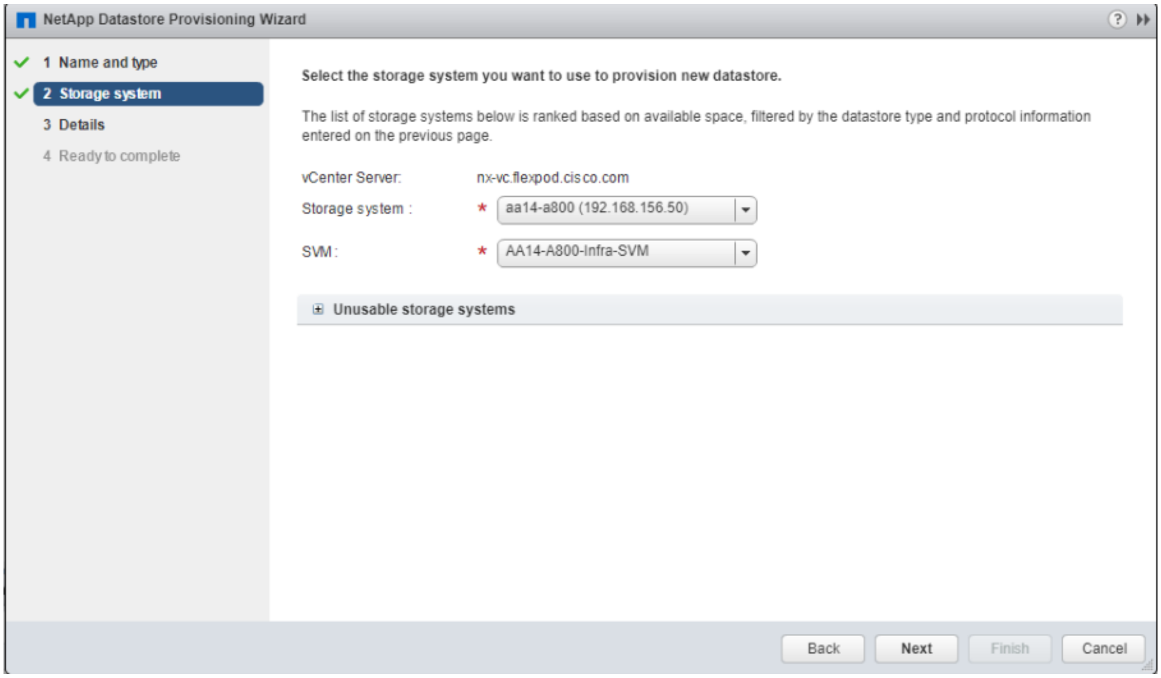

Select the storage system you want to use to provision new datastore.

The list of storage systems below is ranked based on available space, filtered by the datastore type and protocol information entered on the previous page.

vCenter Server: nx-vc.flexpod.cisco.com

Storage system : * aa14-a800 (192.168.156.50)

SVM : * AA14-A800-Infra-SVM

6. Enter the size of the datastore and select the aggregate name.

7. Click Next.

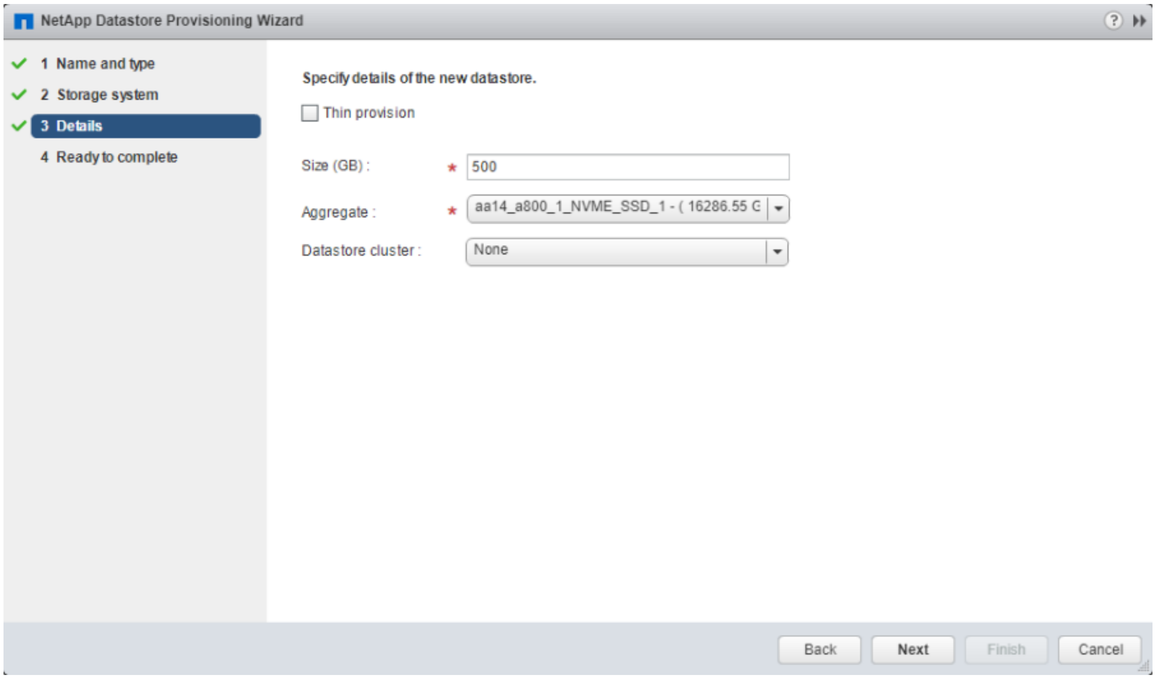
NetApp Datastore Provisioning Wizard

Specify details of the new datastore.

□ Thin provision

Size (GB) : * 500

Aggregate : * aa14_a800_1_NVME_SSD_1 - (16286.55 G

Datastore cluster : None

8. Review the details and click Finish.

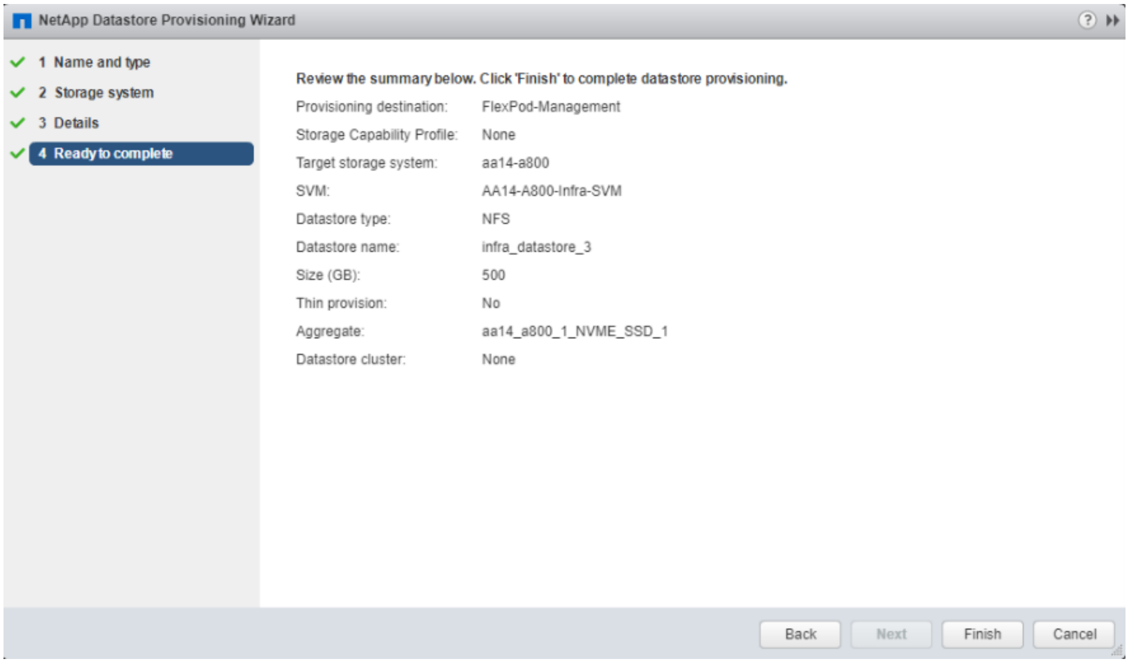
NetApp Datastore Provisioning Wizard

Review the summary below. Click 'Finish' to complete datastore provisioning.

Provisioning destination:	FlexPod-Management
Storage Capability Profile:	None
Target storage system:	aa14-a800
SVM:	AA14-A800-Infra-SVM
Datastore type:	NFS
Datastore name:	infra_datastore_3
Size (GB):	500
Thin provision:	No
Aggregate:	aa14_a800_1_NVME_SSD_1
Datastore cluster:	None

9. Click Ok.

The datastore is created and mounted on all the hosts in the cluster. Click Refresh from the vSphere web client to see the newly created datastore.

Provision FC Datastore

To provision the FC datastore, follow these steps:

1. From the Home screen of the vSphere web client, right-click the FlexPod-Management cluster and select NetApp VSC > Provision Datastore.

2. Enter the datastore name and select the type as VMFS. For the VMFS protocol, select FC/FCoE.

3. Click Next.

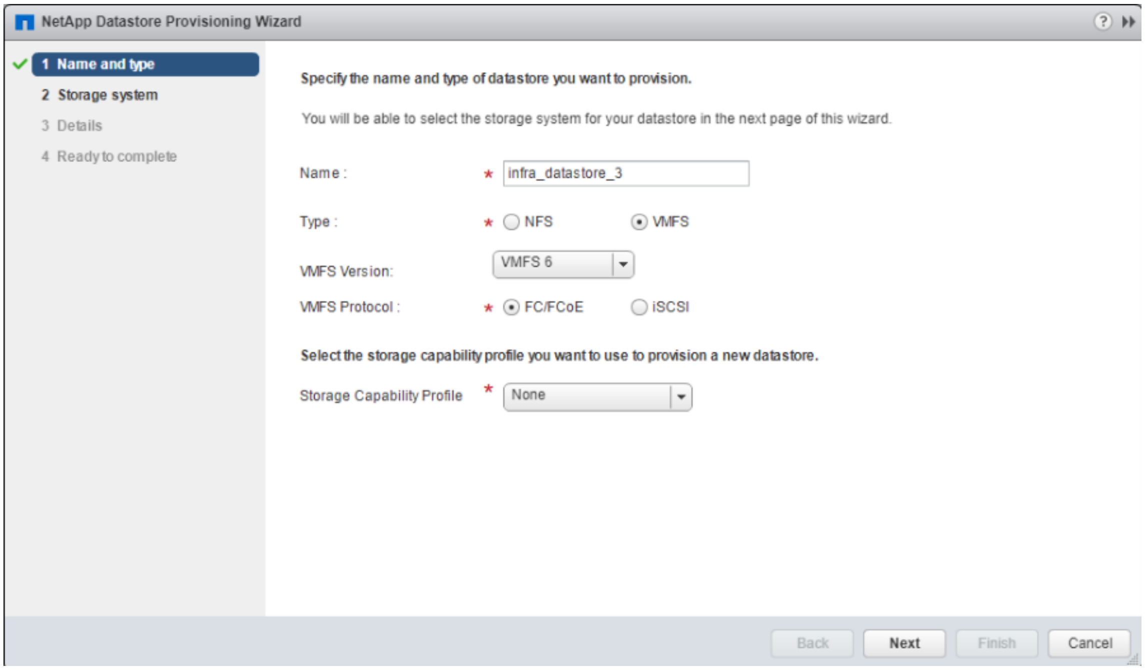

NetApp Datastore Provisioning Wizard

Specify the name and type of datastore you want to provision.

You will be able to select the storage system for your datastore in the next page of this wizard.

Name : * infra_datastore_3

Type : * ○ NFS ● VMFS

VMFS Version: VMFS 6

VMFS Protocol : * ● FC/FCoE ○ iSCSI

Select the storage capability profile you want to use to provision a new datastore.

Storage Capability Profile * None

4. Select the cluster name and the desired SVM to create the datastore. In this example, Infra-SVM is selected.

5. Click Next.

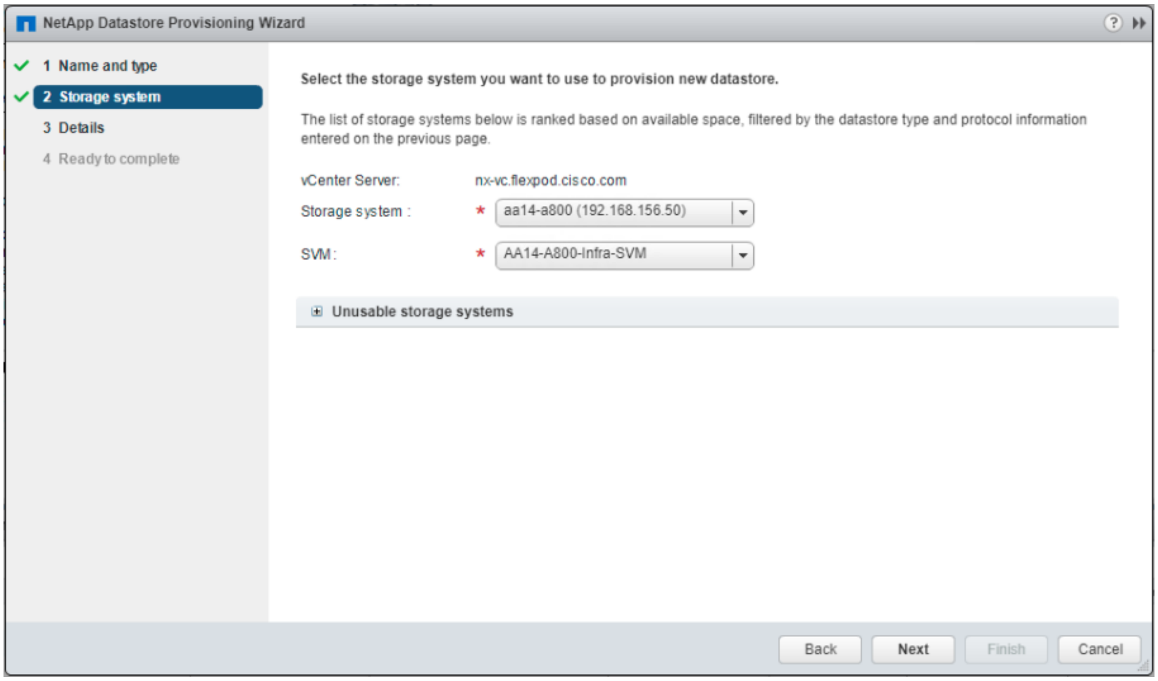

Select the storage system you want to use to provision new datastore.

The list of storage systems below is ranked based on available space, filtered by the datastore type and protocol information entered on the previous page.

vCenter Server: nx-vc.flexpod.cisco.com

Storage system : * aa14-a800 (192.168.156.50)

SVM : * AA14-A800-Infra-SVM

⊞ Unusable storage systems

6. Enter the size of the datastore. Select the Create New Volume checkbox and select the aggregate name.

7. Click Next.

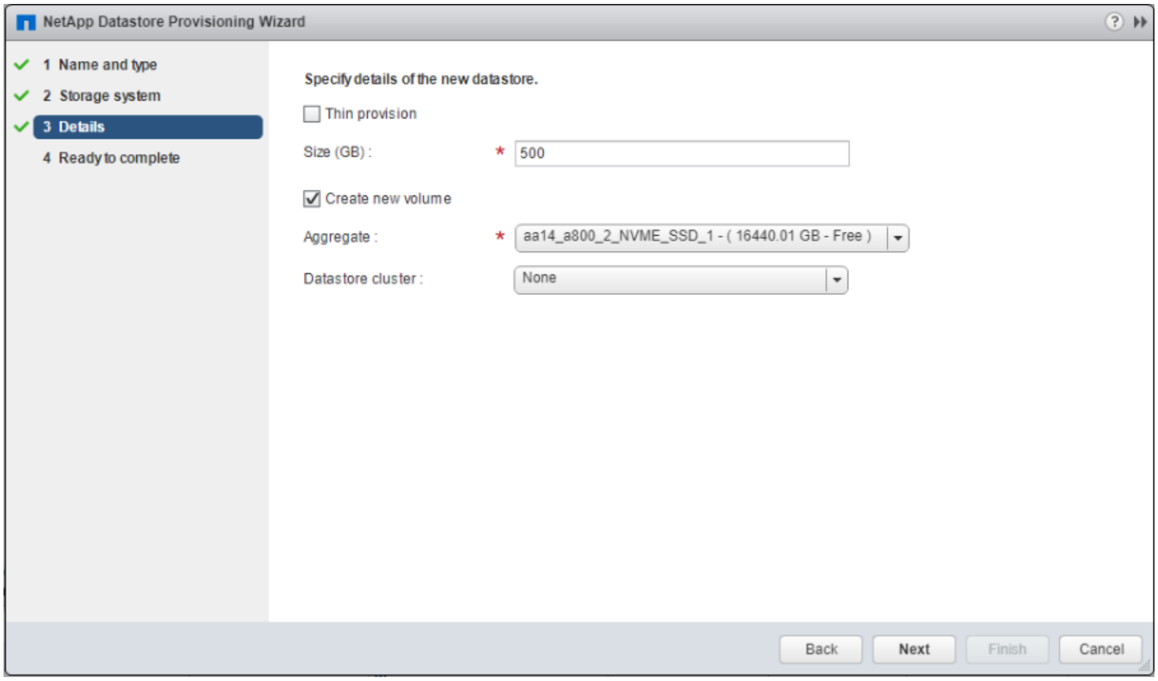

Specify details of the new datastore.

□ Thin provision

Size (GB) : * 500

☑ Create new volume

Aggregate : * aa14_a800_2_NVME_SSD_1 - (16440.01 GB - Free)

Datastore cluster : None

8. Review the details and click Finish.

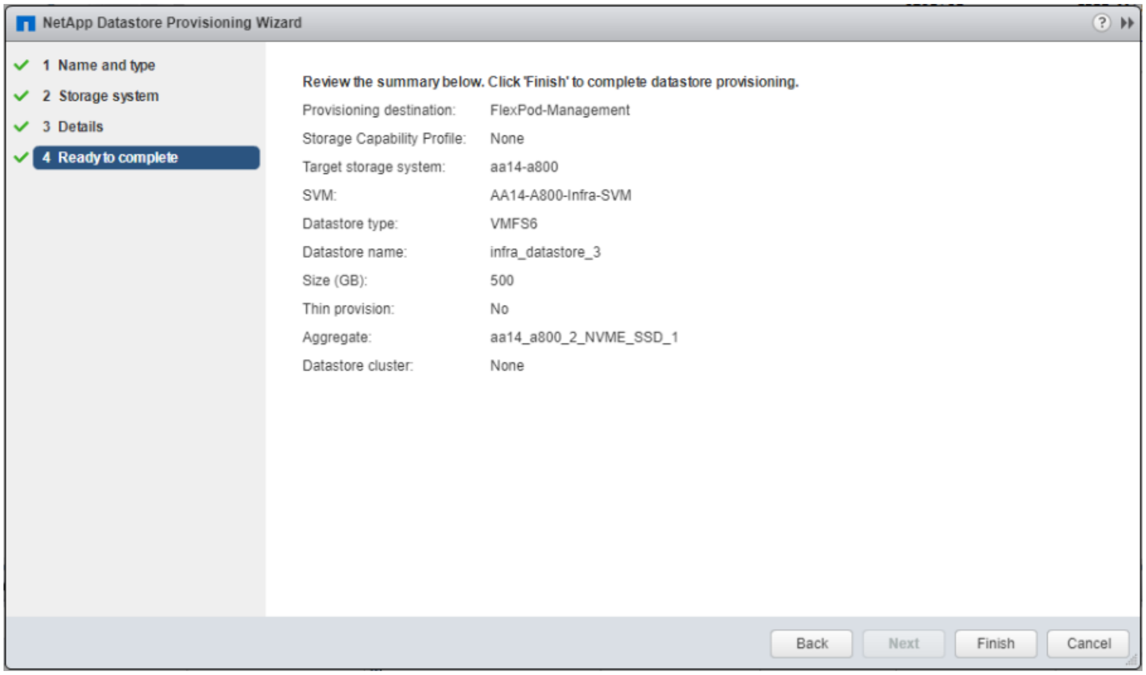

Review the summary below. Click 'Finish' to complete datastore provisioning.

Provisioning destination: FlexPod-Management
Storage Capability Profile: None
Target storage system: aa14-a800
SVM: AA14-A800-Infra-SVM
Datastore type: VMFS6
Datastore name: infra_datastore_3
Size (GB): 500
Thin provision: No
Aggregate: aa14_a800_2_NVME_SSD_1
Datastore cluster: None

9. Click Ok.

The datastore is created and mounted on all the hosts in the cluster. Click Refresh from the vSphere web client to see the newly created datastore.

NetApp SnapCenter

Deploy NetApp SnapCenter

The SnapCenter platform is based on a multitier architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter host agent.

SnapCenter Server Requirements

Table 8 lists the minimum requirements for installing SnapCenter Server and the SnapCenter plug-in on a Windows server. For the latest version compatibility and other plug-in information, refer to the NetApp Interoperability Matrix Tool.

See the SnapCenter 4.1.1 Installation and Setup Guide for the complete documentation.

Table 8 SnapCenter Server Requirements

Component	Requirements
Minimum CPU count	4 cores/vCPUs
Memory	Minimum: 8GB Recommended: 32GB

Component	Requirements
Storage space	Minimum space for installation: 10GB Minimum space for repository: 20GB
Supported operating systems	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016
Software packages	.NET 4.5.2 or later Windows Management Framework 4.0 or later PowerShell 4.0 or later Java 1.8 (64-bit)
Active Directory domain membership	Windows Host must be joined to an Active Directory domain.
Database for SnapCenter repository	MySQL Server 5.7.22 (installed as part of the SnapCenter installation)

Table 9 Port Requirements

Port	Requirement
443	vCenter Server to SnapCenter Server API Access over HTTPS
8144	SnapCenter GUI to SnapCenter Plug-in for VMWare
8145	SnapCenter Server Core API
8146	For SnapCenter server REST API
3306	MySQL

Figure 3 NetApp Snapshot Deployment

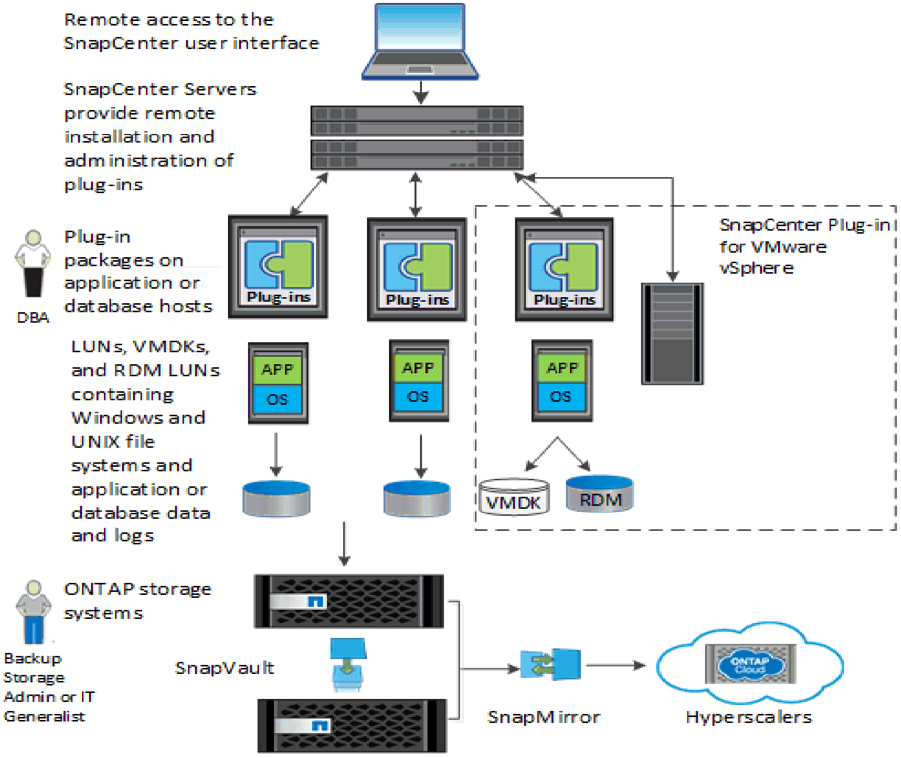
Remote access to the SnapCenter user interface
SnapCenter Servers provide remote installation and administration of plug-ins
Plug-in packages on application or database hosts
DBA
Plug-ins
Plug-ins
Plug-ins
SnapCenter Plug-in for VMware vSphere
APP
OS
APP
OS
APP
OS
LUNs, VMDKs, and RDM LUNs containing Windows and UNIX file systems and application or database data and logs
VMDK
RDM
ONTAP storage systems
Backup Storage Admin or IT Generalist
SnapVault
SnapMirror
ONTAP Cloud
Hyperscalers

SnapCenter 4.1.1 License Requirements

The following licenses are required for SnapCenter on storage systems that run ONTAP 9.5. These licenses do not need to be added into the SnapCenter GUI:

- Protocol licenses (NFS, FCP, or iSCSI as per protocol selection for deployment)
- NetApp FlexClone (for provisioning and cloning only)
- NetApp SnapRestore (for backup and recovery)
- The NetApp SnapManager Suite

SnapCenter Server

The SnapCenter Server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, APIs, and the SnapCenter repository. SnapCenter enables load balancing, high availability, and horizontal scaling across multiple SnapCenter Servers within a single user interface. You can accomplish high availability by using network load balancing (NLB) and Application Request Routing (ARR) with SnapCenter. For larger environments with thousands of hosts, adding multiple SnapCenterServers can help balance the load.

SnapCenter Plug-In for VMware vSphere

The Plug-in for VMware vSphere is a host-side component of the NetApp storage solution. It provides a vSphere web client GUI on vCenter to protect VMware VMs and datastores, and it helps SnapCenter application-specific plug-ins to protect virtualized databases and file systems.

Support for Virtualized Databases and File Systems

The Plug-in for VMware vSphere provides native backup, recovery, and cloning of virtualized applications (virtualized SQL and Oracle databases and Windows file systems) when using the SnapCenter GUI.

SnapCenter natively leverages the Plug-in for VMware vSphere for all SQL, Oracle, and Windows file system data protection operations on VM disks (VMDKs), raw device mappings (RDMs), and NFS datastores.

SnapCenter Installation

Before you install SnapCenter, the SnapCenter Server host system must be up-to-date with Windows updates and no pending system restarts.

1. Download the SnapCenter Server installation package from the NetApp Support site.
2. Install the SnapCenter Server by double-clicking the downloadable .exe file to launch the SnapCenter Server installer.
3. On the Prerequisites Validation screen, the host is validated to see if it meets the requirements to install the SnapCenter Server. If the minimum requirements are not met appropriate error or warning messages are displayed. If the restart of the host is pending, a warning message is displayed.
4. In the Network Load Balancing screen, enable and configure NLB on the host if desired.
   a. Select Create New NLB Cluster and enter the details for first node.
   b. After creating the NLB cluster on the first node, when you run the installer on a second node, select Join Existing NLB Cluster and enter the details.
5. On the Credentials screen, enter the credentials that you want to use to log in to SnapCenter as the administrator.
   a. The SnapCenter Server web component and SnapCenter SMCore Service are installed in the corresponding folders at the default location C:\Program Files\NetApp.
   b. The repository component is installed at the default location C:\ProgramData\NetApp\SnapCenter.
6. On the SnapCenter Ports Configuration screen, enter the port details. The default ports are auto populated but you can specify a custom port. In a NLB setup for the second node, the ports used while installing on the first node are auto populated and are not configurable.
7. On the MySQL Database Connection screen, enter the MySQL database password.
8. On the Ready to Install screen, click Install. Log files are listed (oldest first) in the %temp% folder.

SnapCenter Configuration

The default GUI URL is a secure connection to the default port 8146 on the server where the SnapCenter Server is installed (https://server:8146). If you provided a different server port during the SnapCenter installation, that port is used instead.

For NLB deployment, you must access SnapCenter using the NLB cluster IP (https://NLB_Cluster_IP:8146). If you do not see the SnapCenter UI when you navigate to https://NLB_Cluster_IP:8146 in IE, you must add the NLB IP address as a trusted site in IE on each plug-in host, or you must disable IE Enhanced Security on each plug-in host.

Add a User or Group to a Role

To add a user or group to a role, follow these steps:

1. In the left navigation pane, click Settings.
2. In the Settings page, click Roles.
3. In the Roles page, select the role to which you want to add the user.
4. Click Modify.
5. Click Next until you reach the Users/Groups page of the wizard.
6. Select Domain or Workgroup as the Authentication type. For Domain authentication type, you should specify the domain name of the user or group to which you want to add the user to a role.
7. Select either User or Group.
8. In the User Name or Group Name field, enter a user or group name, and then click Add.
9. Repeat this step to add additional users or groups to the selected role.
10. Click Next to view the summary and then click Finish.

Configure Storage System Connections

To perform data protection and provisioning operations with SnapCenter, you must first set up the storage system connections that give SnapCenter access to ONTAP storage. Storage systems can also be added thru SnapCenter Plugin for vCenter. Both methods require the same set of information.

Install SnapCenter Plug-In for VMware

Host and Privilege Requirements for the Plug-In for VMware vSphere

Review the following requirements before you install the SnapCenter Plug-In for VMware vSphere:

- You must have SnapCenter admin privileges to install and manage the SnapCenter GUI.
- You must install the Plug-In for VMware vSphere on a Windows host (virtual host or physical host). The Plug-In for VMware vSphere must be installed on a Windows host regardless of whether you use the plug-In to protect data on Windows systems or Linux systems.
- When installing a plug-in on a Windows host, if you specify a Run As account that is not built-in or if the Run As user belongs to a local workgroup user, you must disable UAC on the host.
- Do not install the Plug-In for VMware vSphere on the vCenter Server appliance, which is a Linux host. You can only install the Plug-In for VMware vSphere on a Windows hosts.
- You must not install other plug-ins on the host on which the Plug-In for VMware vSphere is installed.

- You must install and register a separate, unique instance of the Plug-In for VMware vSphere for each vCenter Server.
  - Each vCenter Server, whether or not it is in Linked mode, must be paired with a separate instance of the Plug-In for VMware vSphere.
  - Each instance of the Plug-In for VMware vSphere must be installed on a separate Windows host. One instance can be installed on the SnapCenter Server host.
  - vCenter instances in Linked mode must all be paired with the same SnapCenter Server.

For example, if you want to perform backups from six different instances of the vCenter Server, then you must install the Plug-In for VMware vSphere on six hosts (one host can be the SnapCenter Server host) and each vCenter Server must paired with a unique instance of the Plug-In for VMware vSphere.

Run As Credentials

Before you can perform data protection operations, you must set up the SVM connections and add Run As credentials that the SnapCenter Server and the SnapCenter plug-ins use.

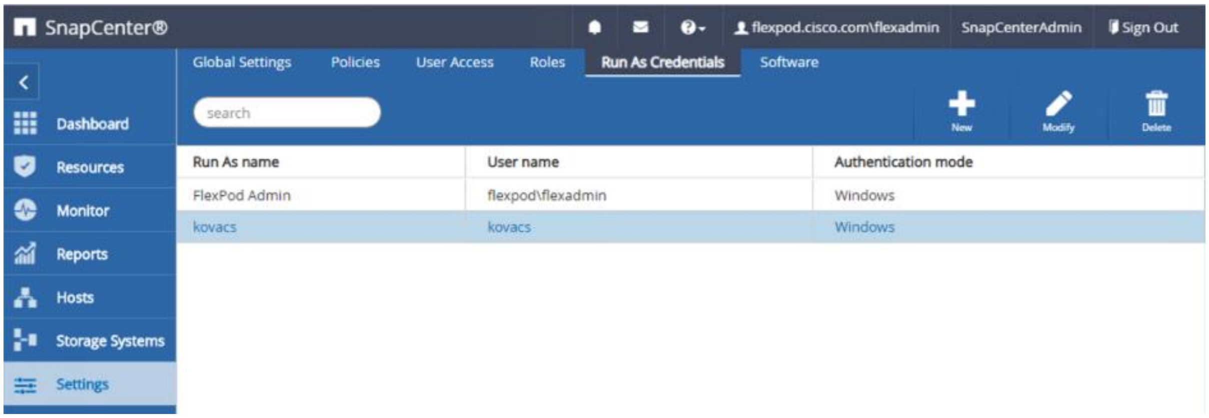

Run As name	User name	Authentication mode
FlexPod Admin	flexpod\flexadmin	Windows
kovacs	kovacs	Windows

1. Domain administrator or any member of the administrator group. Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the ***Username*** field are as follows:

```
NetBIOS\UserName
Domain FQDN\UserName
UserName@upn
```

2. Local administrator (for workgroups only). For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrator's group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the ***username field is: UserName***

3. Run As credentials for individual resource groups. If you set up Run As credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

To perform data protection and provisioning operations with SnapCenter, first you must set up the storage system connections using the SnapCenter Plug-In for vCenter GUI that, to give SnapCenter access to ONTAP storage as described in next section.

Install the Plug-in for VMware vSphere from the SnapCenter GUI

To install the plug-in for VMware vSphere from the SnapCenter GUI, follow these steps:

1. In the left navigation pane, click Hosts. Verify that Managed Hosts is selected at the top.
2. Click Add.
3. On the Hosts page, specify:
   a. Host OS. vSphere
   b. Host Name. Host on which you want to install the Plug-in for VMware
   c. Run As Name. Run As account configured with user credentials with access to the host
   d. Port. Leave the default of 8145
4. On the Plug-Ins to Install page, specify the vCenter information
5. Review the summary, and then click Finish.
6. SnapCenter Server performs the following tasks during the installation process:
   a. Adds the host to the SnapCenter registry.
   b. Installs the Plug-In for VMware vSphere, if the plug-in is not already installed on the host.
   c. If the Plug-In for VMware vSphere is installed on the SnapCenter Server host, it also installs the SnapCenter Plug-In for Microsoft Windows to support SnapCenter repository backup operations by using PowerShell cmdlets. You cannot use any other features of this instance of the Plug-In for Windows.
   d. Adds the Plug-In for VMware vSphere web client to vCenter.

Configure SnapCenter Plug-In for vCenter

After you have successfully installed the Plug-in for VMware vSphere, to configure SnapCenter and make it ready to backup virtual machines, follow these steps:

1. In your browser, navigate to VMware vSphere web client URL https://<vCenter Server>/vsphere-client/?csp.

If currently logged into vCenter, logoff and sign-on again to access the SnapCenter Plug-In for VMware vSphere.

2. On the VMware vCenter Single Sign-On page, login.
3. On the VMware vSphere web client page, click Home and select SnapCenter Plug-In for VMware vSphere to bring up the SnapCenter Plug-In for VMware GUI. It will take you to the SnapCenter Plug-In for VMware dashboard.

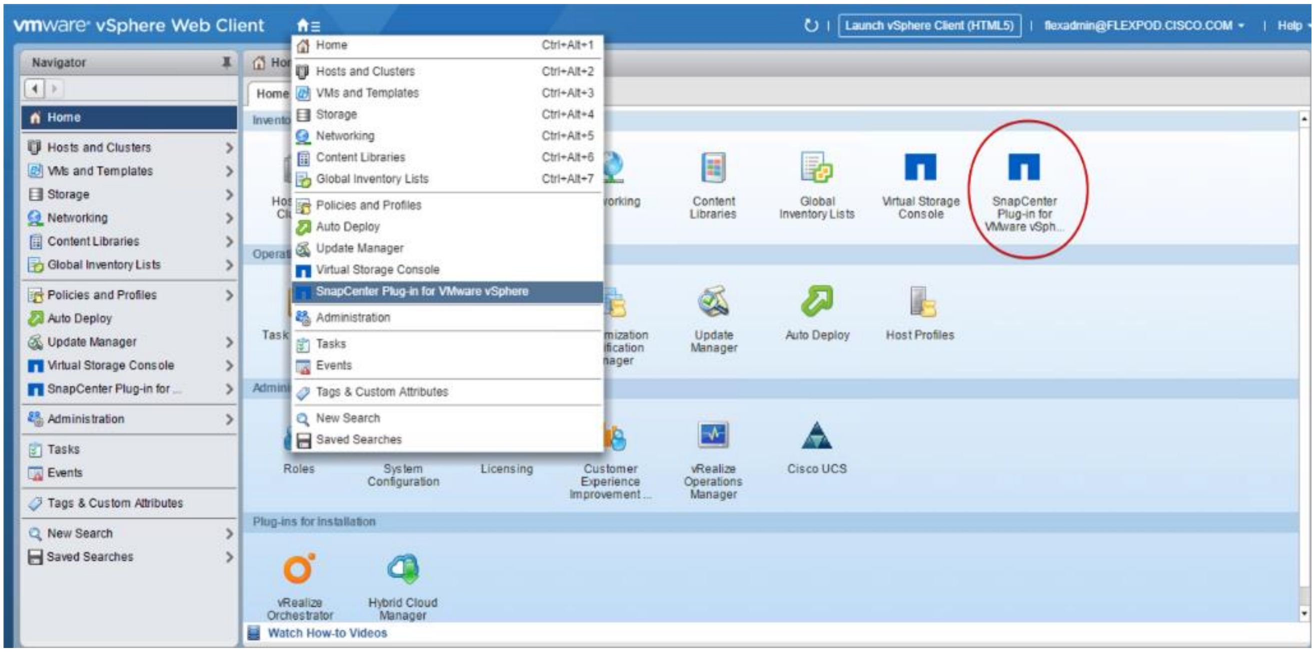

Add Storage Systems (SVM)

To add storage systems, follow these steps:

1. Go to the Storage Systems page.

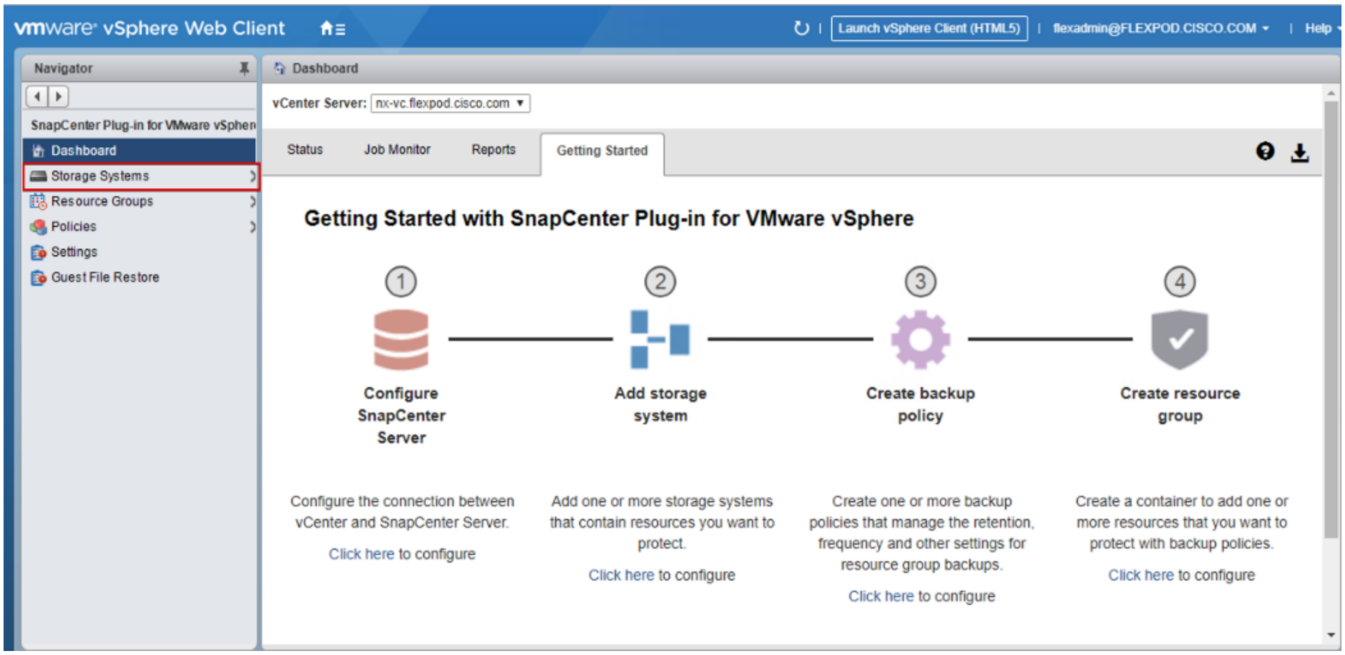
vmware vSphere Web Client
Launch vSphere Client (HTML5)
flexadmin@FLEXPOD.CISCO.COM
Help
Navigator
SnapCenter Plug-in for VMware vSphere
Dashboard
Storage Systems
Resource Groups
Policies
Settings
Guest File Restore
Dashboard
vCenter Server: nx-vc.flexpod.cisco.com
Status
Job Monitor
Reports
Getting Started

Getting Started with SnapCenter Plug-in for VMware vSphere

1 Configure SnapCenter Server

Configure the connection between vCenter and SnapCenter Server.

Click here to configure

2 Add storage system

Add one or more storage systems that contain resources you want to protect.

Click here to configure

3 Create backup policy

Create one or more backup policies that manage the retention, frequency and other settings for resource group backups.

Click here to configure

4 Create resource group

Create a container to add one or more resources that you want to protect with backup policies.

Click here to configure

2. Click + Add Storage System to add an SVM.

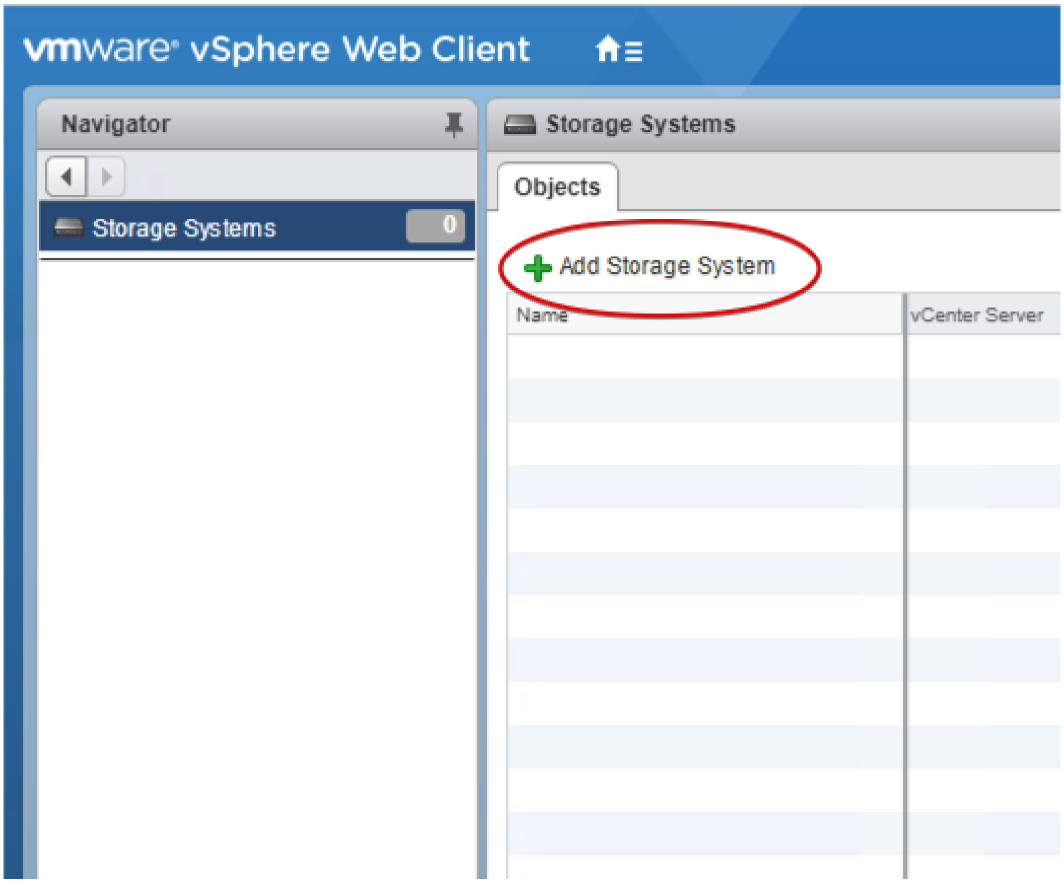

3. Enter vCenter, Storage System, user credentials, and other required information in following dialog box.

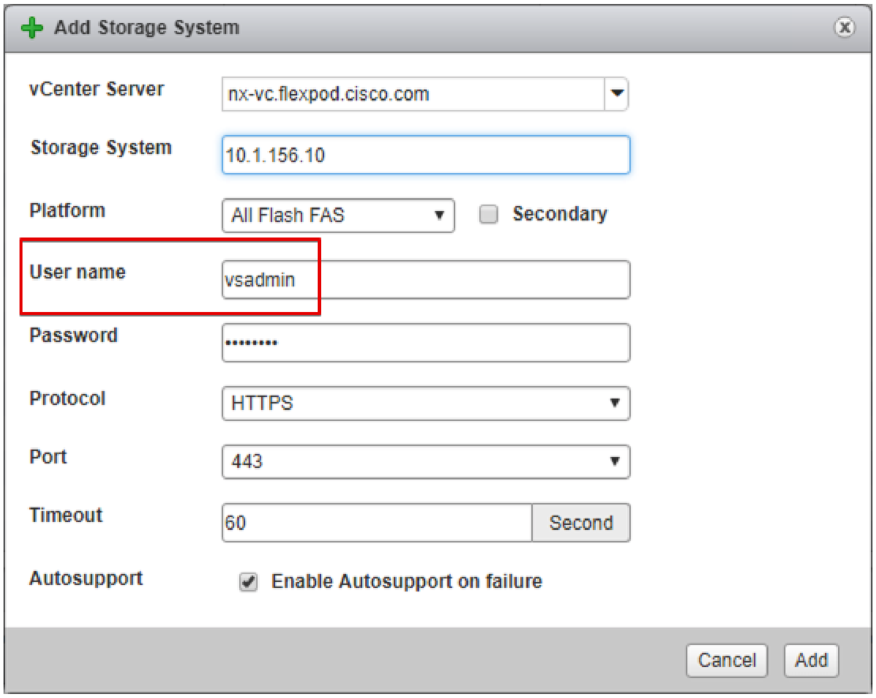
Add Storage System

vCenter Server	nx-vc.flexpod.cisco.com
Storage System	10.1.156.10
Platform	All Flash FAS ☐ Secondary
User name	vsadmin
Password	••••••••
Protocol	HTTPS
Port	443
Timeout	60 Second
Autosupport	☑ Enable Autosupport on failure

You must login to the SVM with the `vsadmin` or another user with vserver admin privileges not the cluster admin user.

Create Backup Policies for Virtual Machines and Datastores

To create backup policies for VMs and datastores, follow these steps:

1. In the left Navigator pane of the VMware vSphere web client, click Policies.

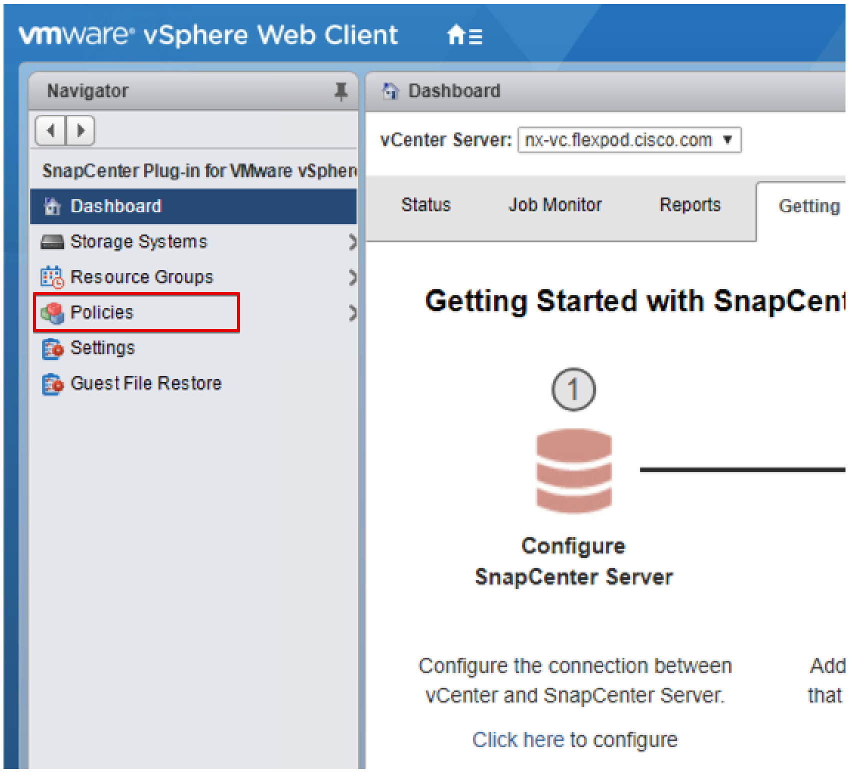
vmware® vSphere Web Client
Navigator
Dashboard
vCenter Server: nx-vc.flexpod.cisco.com ▾
SnapCenter Plug-in for VMware vSpher
Dashboard
Storage Systems
Resource Groups
Policies
Settings
Guest File Restore
Status
Job Monitor
Reports
Getting
Getting Started with SnapCent
1
Configure
SnapCenter Server
Configure the connection between vCenter and SnapCenter Server.
Add
that
Click here to configure

4. On the Policies page, click + New Policy in the toolbar.

5. On the New Backup Policy page, follow these steps:

   a. Enter the policy name and a description.

   b. Enter the backups to keep.

   c. From the Frequency drop-down list, select the backup frequency (hourly, daily, weekly, monthly, and on-demand only).

   d. Click Add.

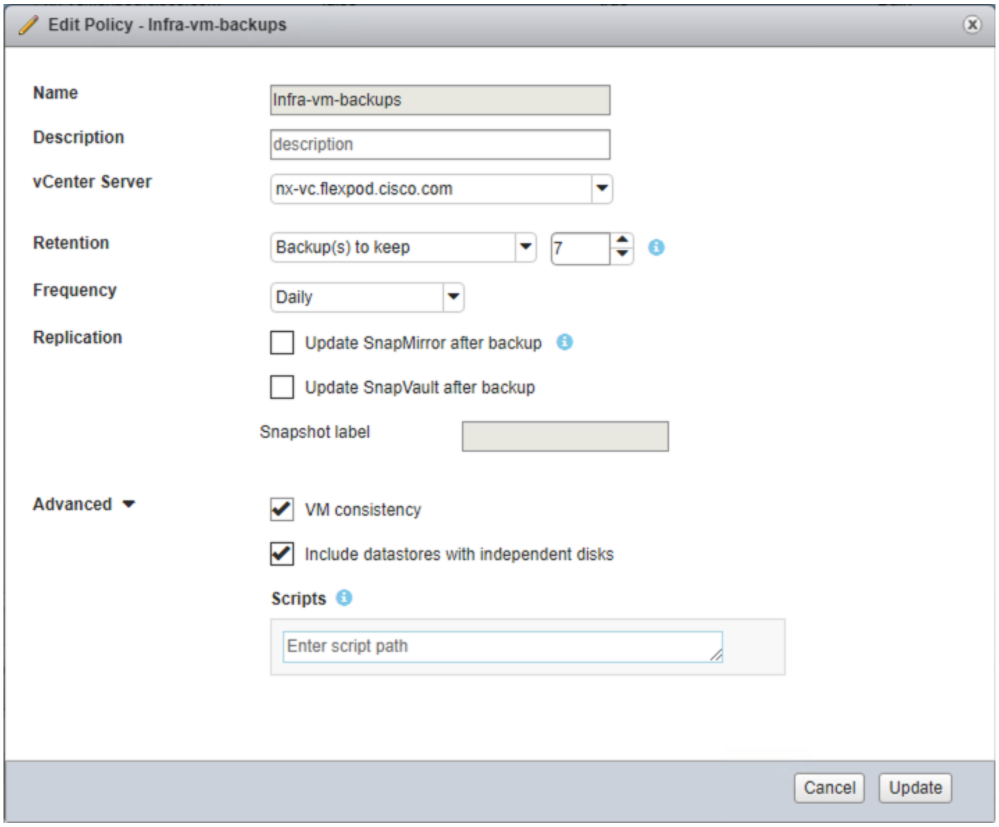

Edit Policy - Infra-vm-backups

Name	Infra-vm-backups
Description	description
vCenter Server	nx-vc.flexpod.cisco.com
Retention	Backup(s) to keep 7
Frequency	Daily
Replication	□ Update SnapMirror after backup □ Update SnapVault after backup Snapshot label
Advanced	☑ VM consistency ☑ Include datastores with independent disks Scripts Enter script path

6. Create multiple policies as required for different sets of VMs or datastores.

Create Resource Groups

Resource groups are groups of virtual machine or datastores that are backed up together. A backup policy is associated with the resource group to back up the virtual machines and retain a certain number of backups as defined in the policy.

To create resource groups, follow these steps:

1. In the left Navigator pane of the VMware vSphere web client, click Resource Groups and then click Create Resource Group. This is the easiest way to create a resource group. However, you can also create a resource group with one resource by performing one of the following steps:

   a. To create a resource group for one virtual machine, click VMs and Templates, right-click a virtual machine, select NetApp SnapCenter from the drop-down list, and then select Create Resource Group from the secondary drop-down list.

   b. To create a resource group for one datastore, click Storage, right-click a datastore, select NetApp SnapCenter from the drop-down list, and then select Create Resource Group from the secondary drop-down list.

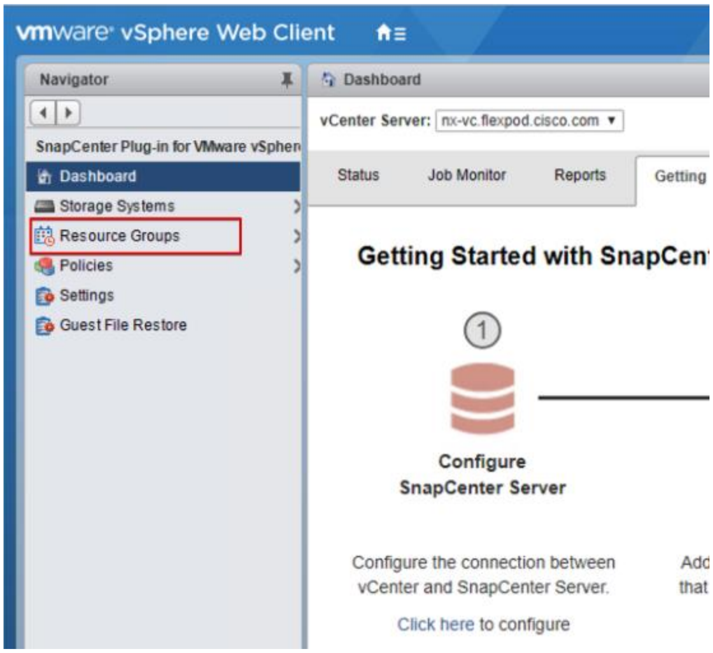
vmware vSphere Web Client
Navigator
Dashboard
vCenter Server: nx-vc.flexpod.cisco.com
SnapCenter Plug-in for VMware vSphere
Dashboard
Storage Systems
Resource Groups
Policies
Settings
Guest File Restore
Status
Job Monitor
Reports
Getting
Getting Started with SnapCen
1
Configure
SnapCenter Server
Configure the connection between
vCenter and SnapCenter Server.
Click here to configure
Add
that

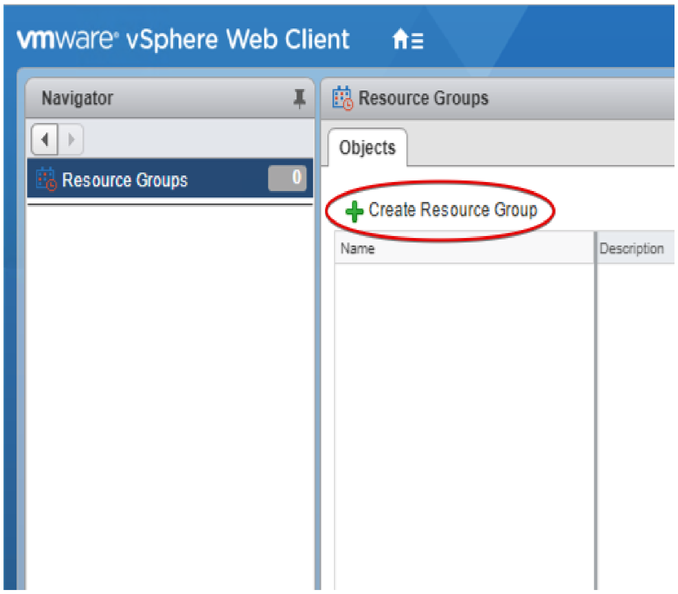

2. On the General Info & Notification page, enter the resource group name and complete the notification settings. Click Next.

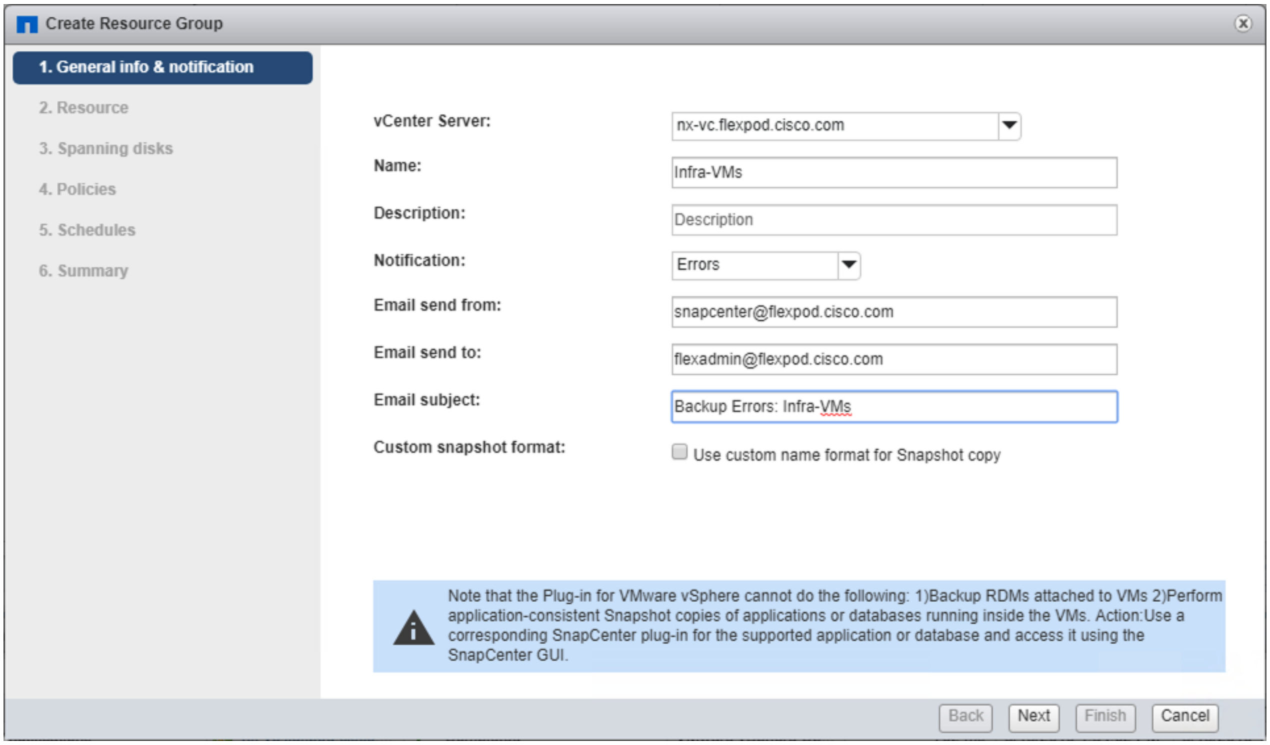

vCenter Server: nx-vc.flexpod.cisco.com

Name: Infra-VMs

Description: Description

Notification: Errors

Email send from: snapcenter@flexpod.cisco.com

Email send to: flexadmin@flexpod.cisco.com

Email subject: Backup Errors: Infra-VMs

Custom snapshot format: □ Use custom name format for Snapshot copy

Note that the Plug-in for VMware vSphere cannot do the following: 1)Backup RDMs attached to VMs 2)Perform application-consistent Snapshot copies of applications or databases running inside the VMs. Action:Use a corresponding SnapCenter plug-in for the supported application or database and access it using the SnapCenter GUI.

3. Select a datastore as a parent entity to create a resource group of virtual machines, and then select the virtual machines from the available list. Click Next.

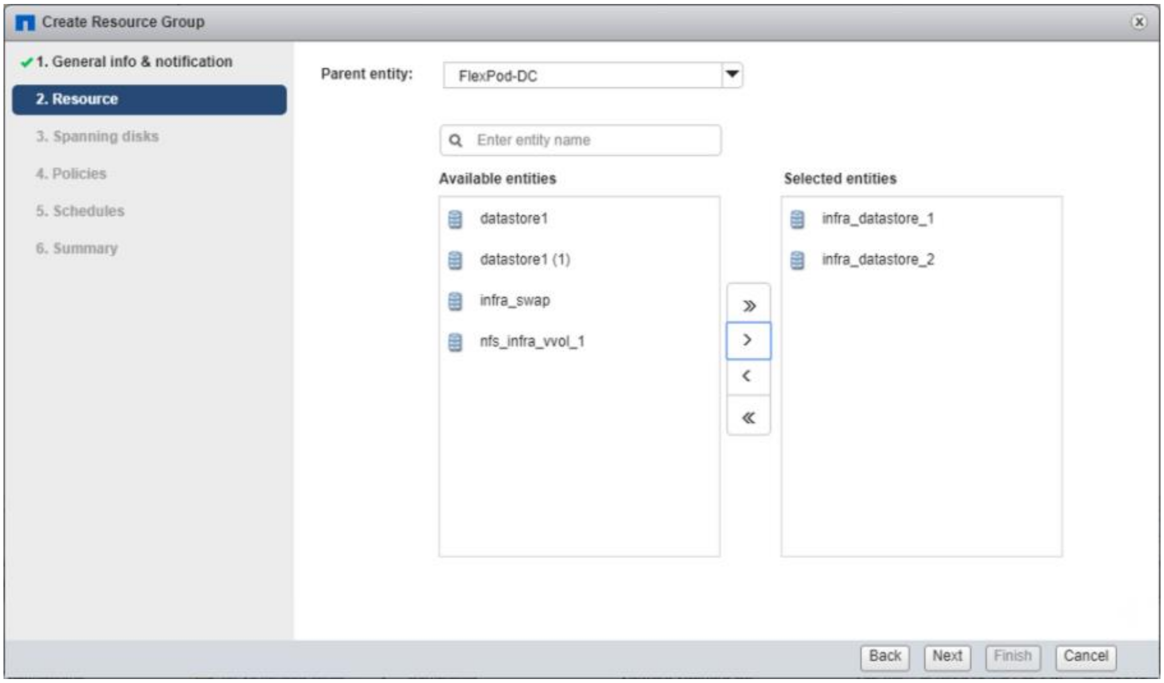

Parent entity: FlexPod-DC

Available entities: datastore1, datastore1 (1), infra_swap, nfs_infra_vvol_1

Selected entities: infra_datastore_1, infra_datastore_2

4. From the Spanning Disks options, select the Always Include all Spanning Datastores option.

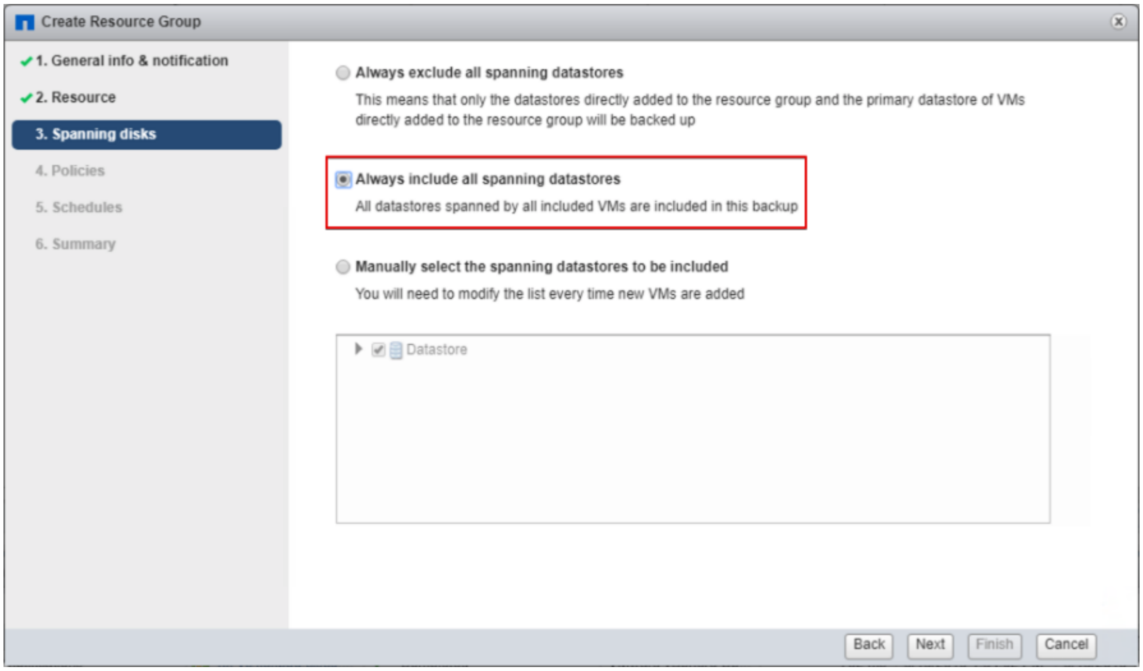

○ **Always exclude all spanning datastores**
This means that only the datastores directly added to the resource group and the primary datastore of VMs directly added to the resource group will be backed up

● **Always include all spanning datastores**
All datastores spanned by all included VMs are included in this backup

○ **Manually select the spanning datastores to be included**
You will need to modify the list every time new VMs are added

5. From the Policies options, select the policies (from those already created earlier) that you want to associate with the resource group. Click Next.

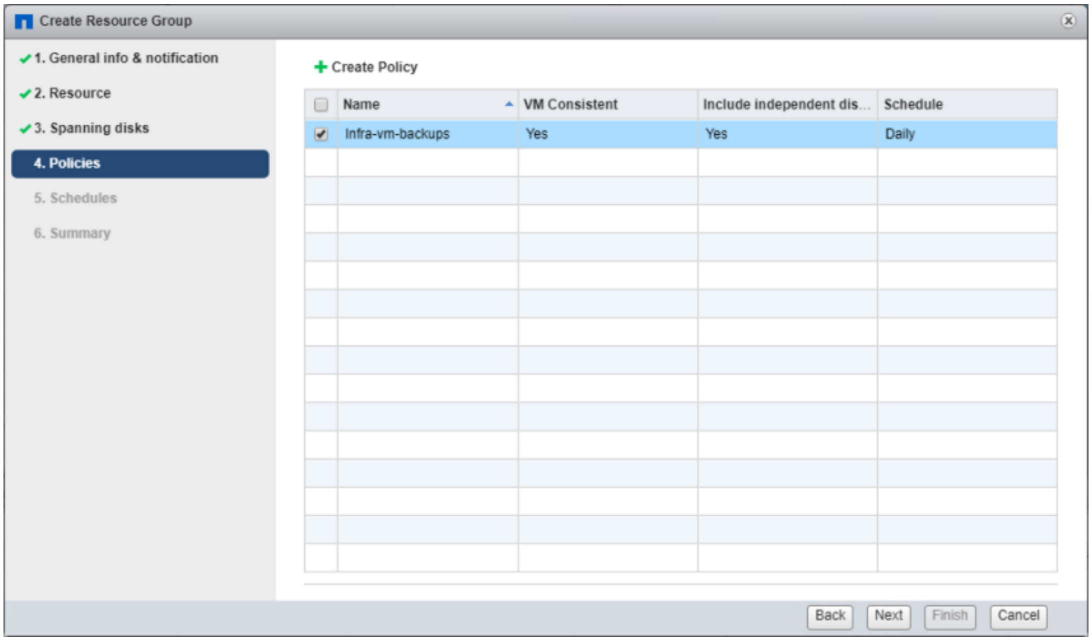

□	Name	VM Consistent	Include independent dis...	Schedule
☑	Infra-vm-backups	Yes	Yes	Daily

6. From the Schedules options, select the schedule for backups for each selected policy. Click Next.

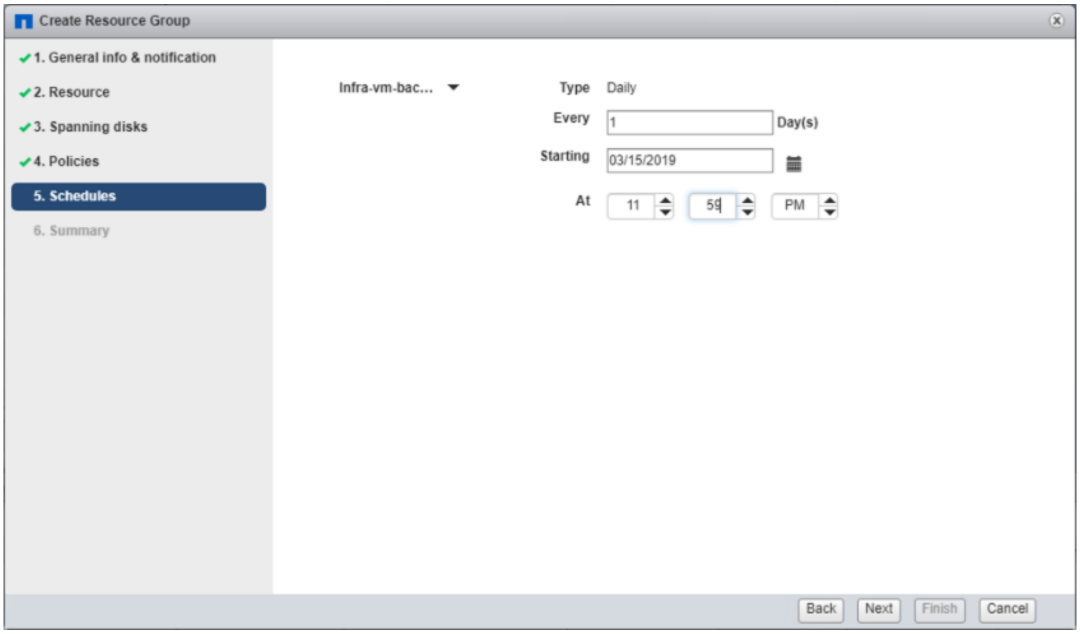

Infra-vm-bac...

Type Daily

Every 1 Day(s)

Starting 03/15/2019

At 11 59 PM

7. Review the summary and click Finish to complete the creation of the resource group.

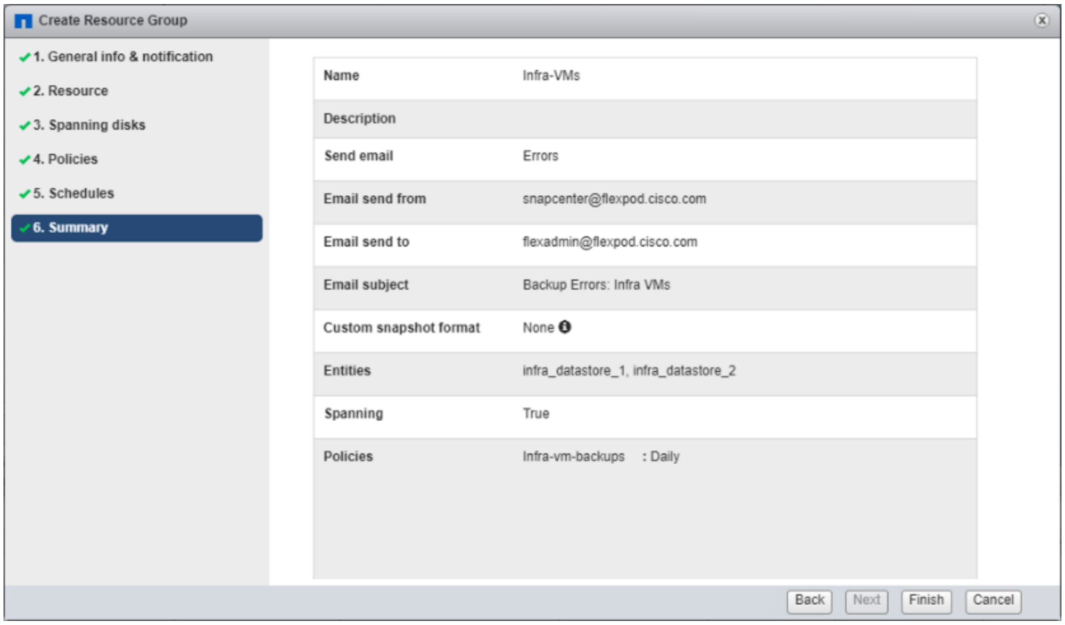

Name	Infra-VMs
Description	
Send email	Errors
Email send from	snapcenter@flexpod.cisco.com
Email send to	flexadmin@flexpod.cisco.com
Email subject	Backup Errors: Infra VMs
Custom snapshot format	None
Entities	infra_datastore_1, infra_datastore_2
Spanning	True
Policies	Infra-vm-backups : Daily

View Virtual Machine Backups and Restore from vCenter by Using SnapCenter Plug-In

View Backups

Backups of the resources included in the resource group occurs according to the schedule of the policies associated with the resource group. To view the backups associated with each schedule, follow these steps:

1. Go to any virtual machine that has been added to Resource Group > More Objects > Backups. It shows all the backups for that resource.

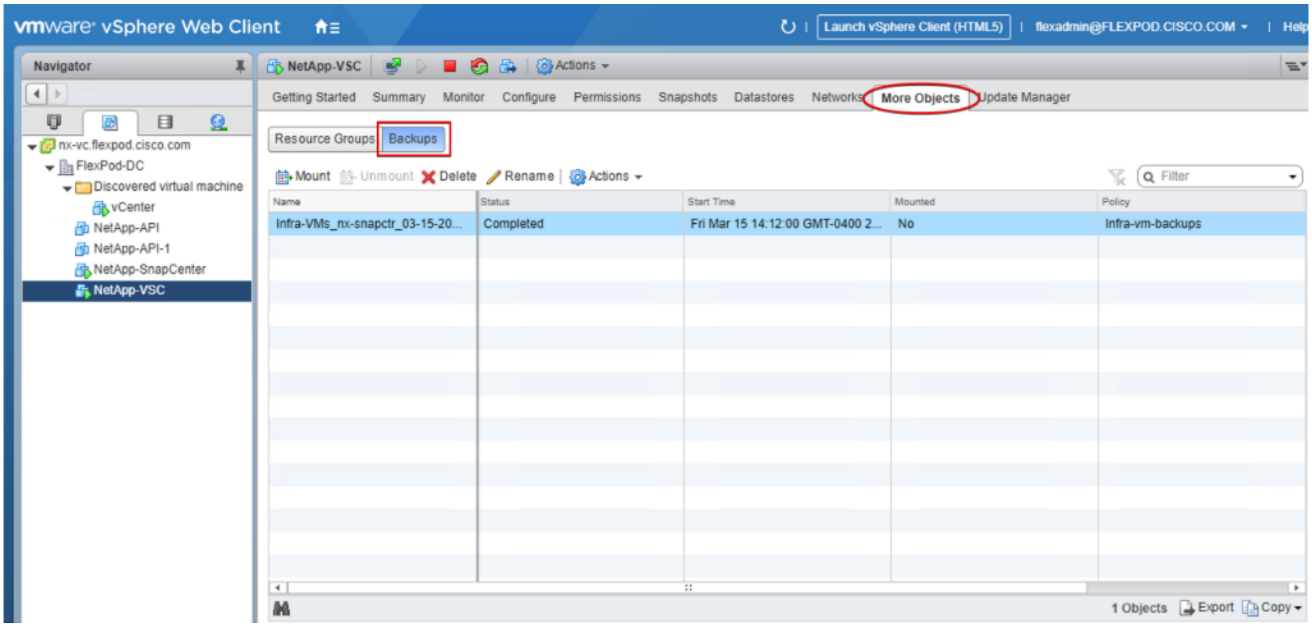

Name	Status	Start Time	Mounted	Policy
Infra-VMs_nx-snapctr_03-15-20...	Completed	Fri Mar 15 14:12:00 GMT-0400 2...	No	Infra-vm-backups

2. On the VMware vSphere Web Client page, click Home and select SnapCenter Plug-In for VMware vSphere to bring up the SnapCenter Plug-In for VMware GUI. Click Resource Groups and select any resource group. In the right pane, the completed backups are displayed.

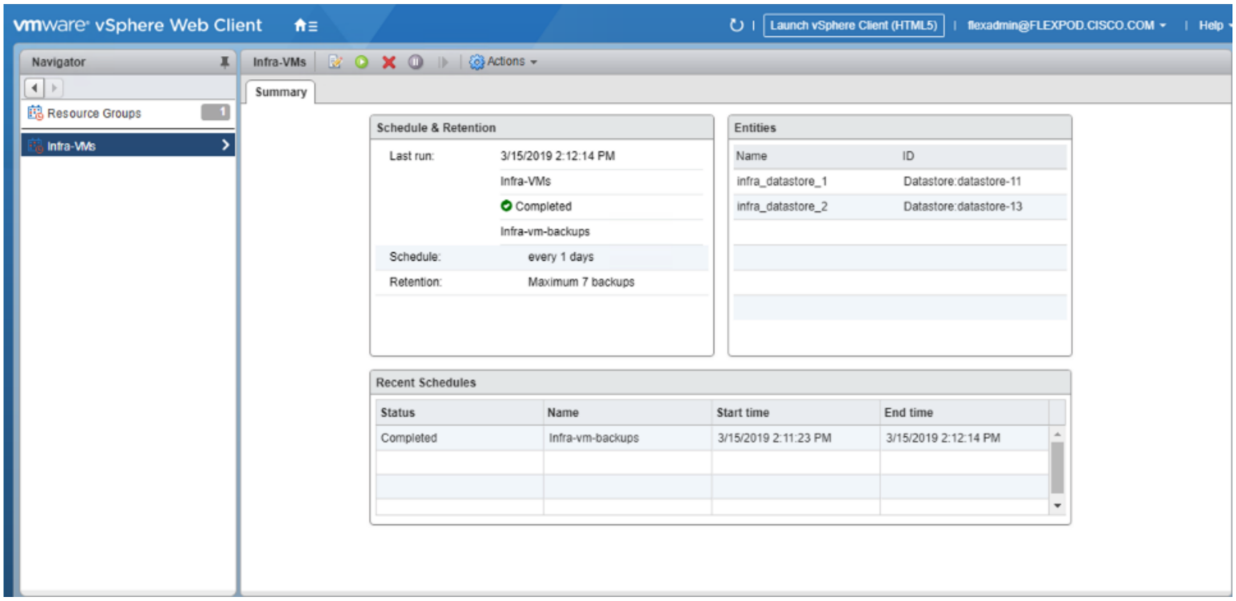
vmware vSphere Web Client
Launch vSphere Client (HTML5)
flexadmin@FLEXPOD.CISCO.COM
Help
Navigator
Resource Groups
1
Infra-VMs
Actions
Summary

Schedule & Retention

Last run: 3/15/2019 2:12:14 PM
Infra-VMs
Completed
Infra-vm-backups

Schedule: every 1 days

Retention: Maximum 7 backups

Entities

Name	ID
infra_datastore_1	Datastore:datastore-11
infra_datastore_2	Datastore:datastore-13

Recent Schedules

Status	Name	Start time	End time
Completed	Infra-vm-backups	3/15/2019 2:11:23 PM	3/15/2019 2:12:14 PM

3. An on-demand backup can also be created for any resource group by following these steps:

   a. Click the virtual machine or datastore resource > More Objects > Resource Groups. Select the Resource Group and click Run Now. This process triggers the backup of the resource group containing a particular resource.

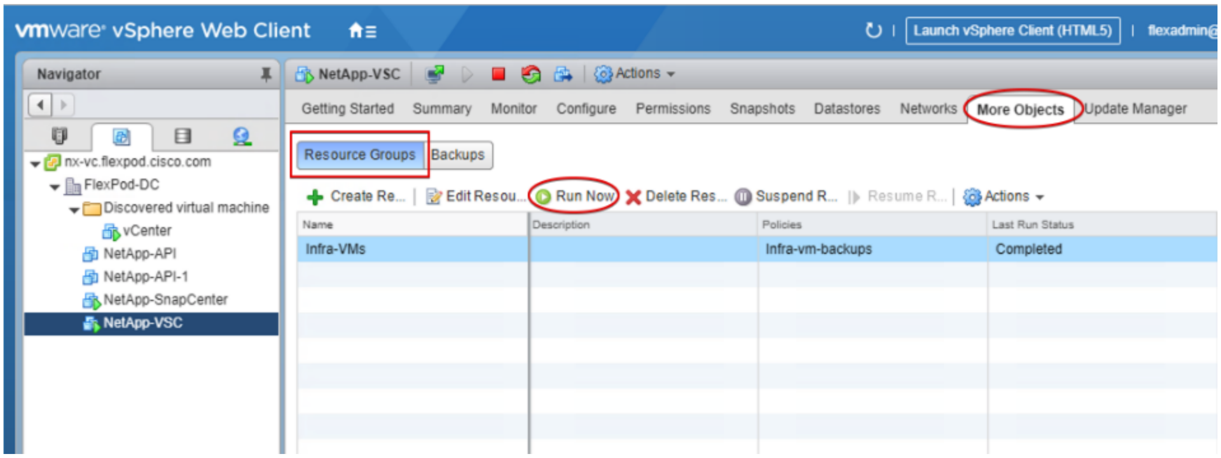

Name	Description	Policies	Last Run Status
Infra-VMs		Infra-vm-backups	Completed

Restore from vCenter by Using SnapCenter Plug-In

To restore from vCenter by using SnapCenter Plug-In, follow these steps:

1. The Plug-In for VMware vSphere provides native backup, recovery, and cloning of virtualized applications. Go to vCenter Web Client GUI > Select a VM > right-click and select NetApp SnapCenter.

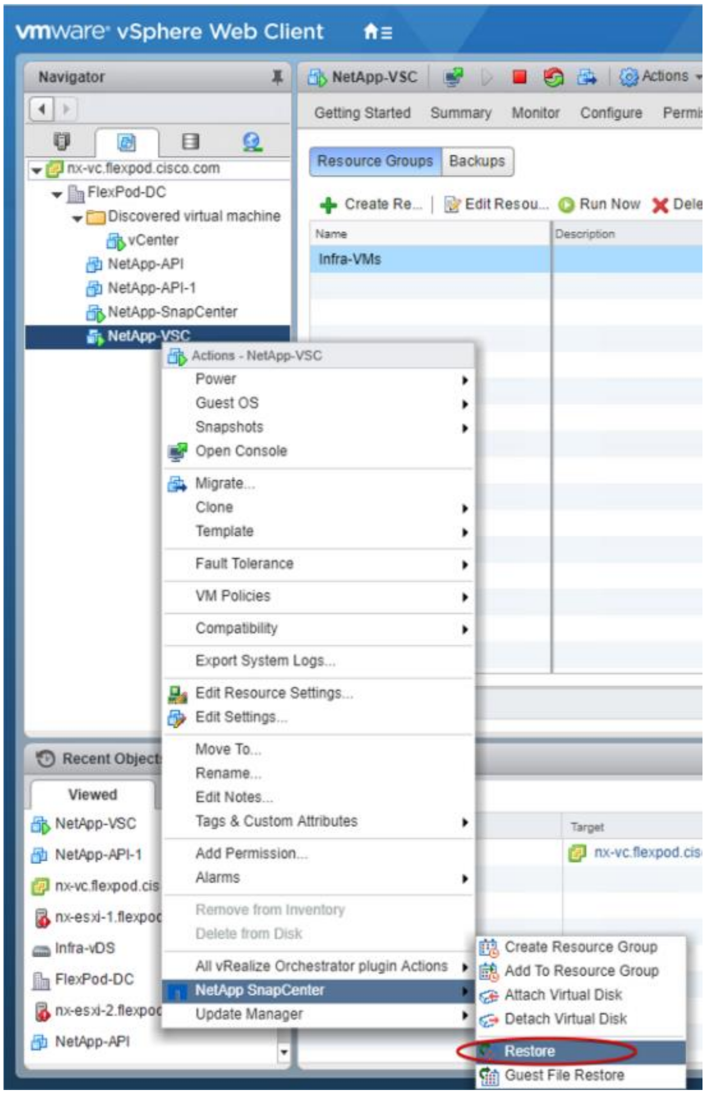

2. Select a backup from which to restore. Click Next.

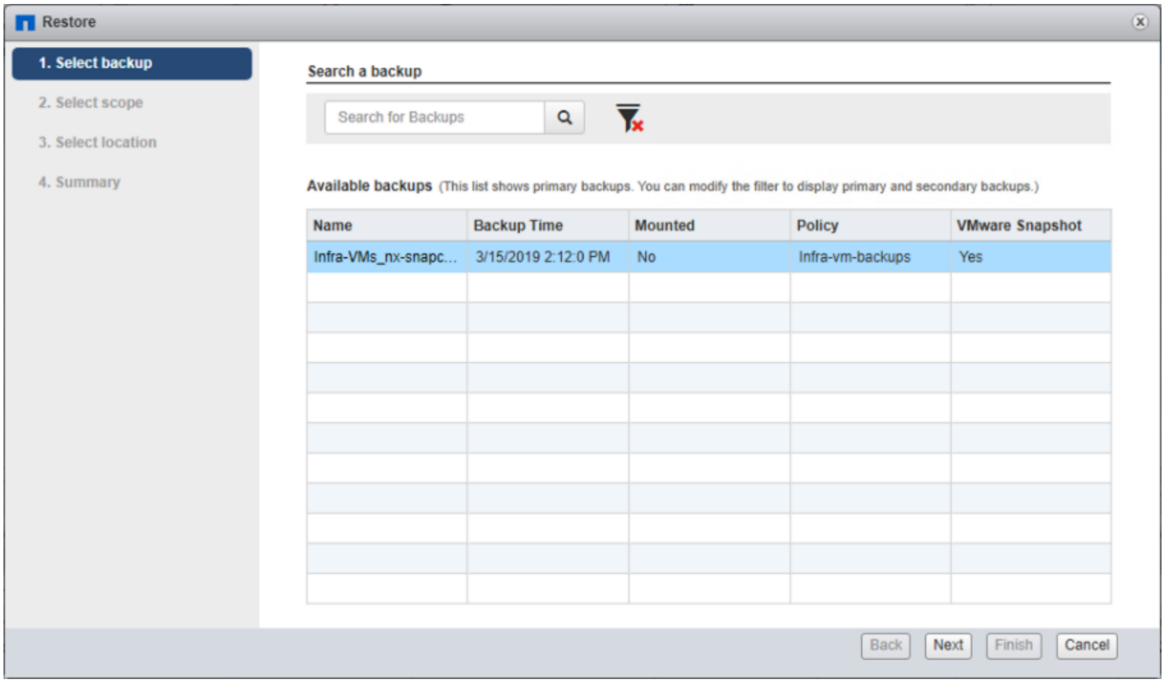

Available backups (This list shows primary backups. You can modify the filter to display primary and secondary backups.)

Name	Backup Time	Mounted	Policy	VMware Snapshot
Infra-VMs_nx-snapc...	3/15/2019 2:12:0 PM	No	Infra-vm-backups	Yes

3. From the Restore Scope drop-down list, select either Entire virtual machine to restore the virtual machine with all VMDKs, or select Particular Virtual Disk to restore the VMDK without affecting the virtual machine configuration and other VMDKs.

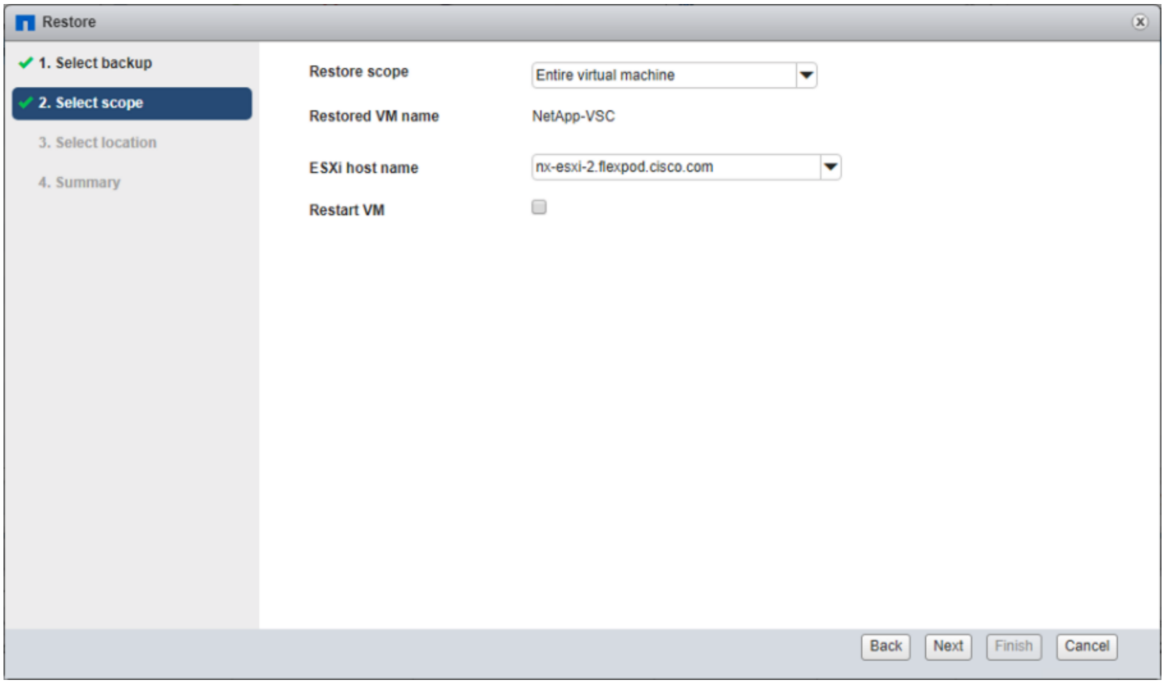

Restore scope	Entire virtual machine
Restored VM name	NetApp-VSC
ESXi host name	nx-esxi-2.flexpod.cisco.com
Restart VM	□

4. Select the destination datastore. Click Next.

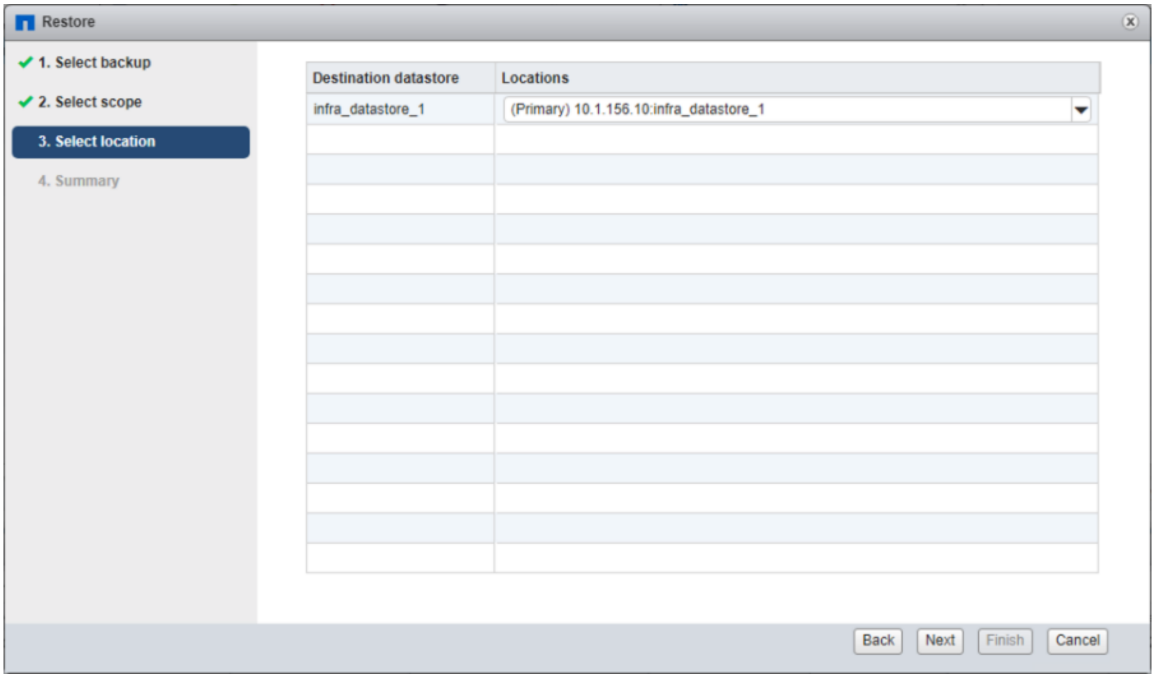

Destination datastore	Locations
infra_datastore_1	(Primary) 10.1.156.10:infra_datastore_1

5. Review the Summary and click Finish to complete the restore process.

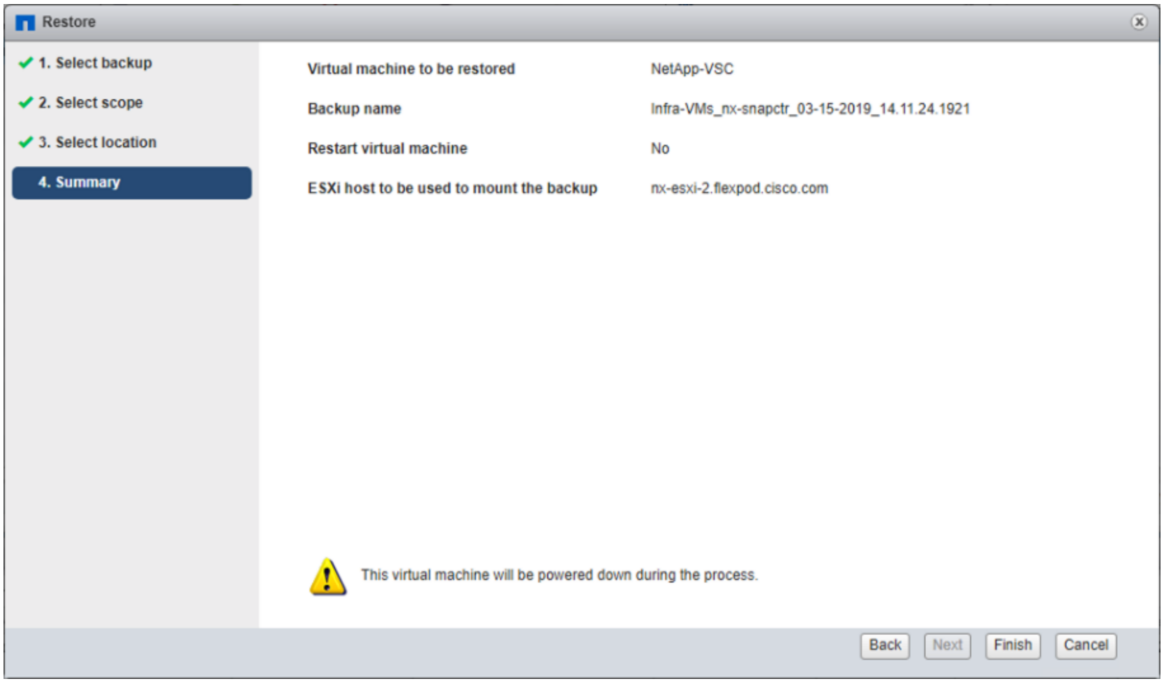

Virtual machine to be restored	NetApp-VSC
Backup name	Infra-VMs_nx-snapctr_03-15-2019_14.11.24.1921
Restart virtual machine	No
ESXi host to be used to mount the backup	nx-esxi-2.flexpod.cisco.com

This virtual machine will be powered down during the process.

Install OnCommand Unified Manager 9.5

This section describes the steps to deploy NetApp OnCommand Unified Manager 9.5. These steps will demonstrate deploying OnCommand Unified Manager 9.5 on a CentOS Linux 7.6 virtual machine. The following table lists the recommended configuration for the virtual machine to install and run OnCommand Unified Manager (OCUM) to ensure acceptable performance.

Table 10 Virtual Machine Configuration

Hardware Configuration	Recommended Settings
RAM	16 GB
Processors	4 CPUs/ vCPUs
CPU Cycle Capacity	9572 MHz total
Free Disk Space	Red Hat or CentOS • 50 GB allotted to the root partition • 100 GB of free disk space allotted to the /opt/netapp/data directory, which is mounted on an LVM drive on a separate local disk attached to the target system

The /tmp directory should have at least 10 GB of free space and the /var/log directory should have at least 16GB of free space.

To install OnCommand Unified Manager 9.5, follow these steps:

1. Download the Minimal ISO image of CentOS 7 Linux from the centos.org mirror site closest to your location. Save the ISO image to one of the infra datastores created in a previous step to a folder called vmimages.

2. Log into vSphere and right-click the FlexPod-DC object in the Navigator panel and select New Virtual Machine -> New Virtual Machine.

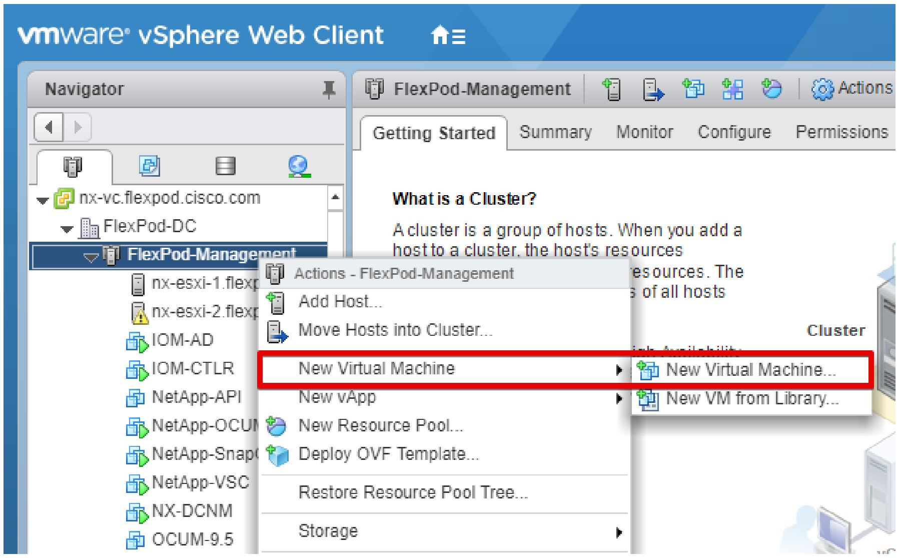

3. Choose the FlexPod-DC datacenter location and provide a name for the virtual machine.

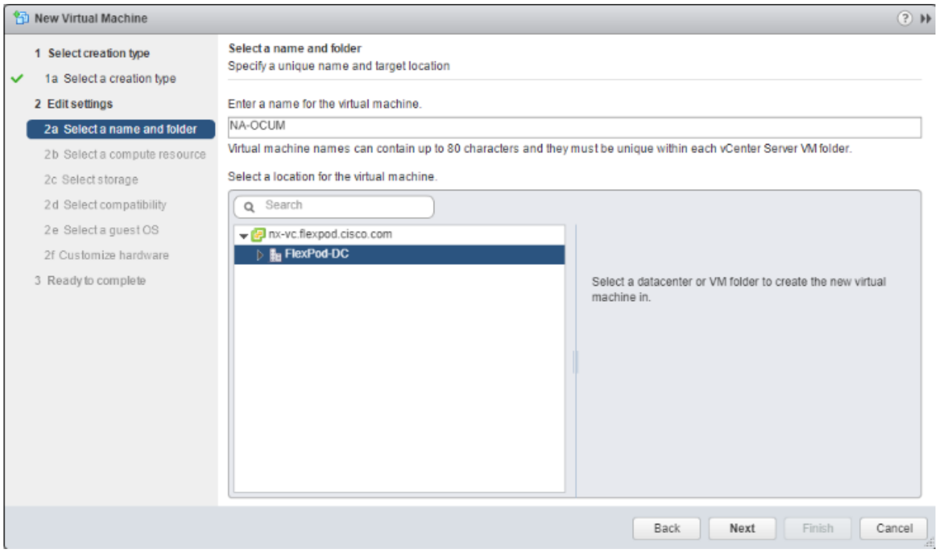

Select a name and folder

Specify a unique name and target location

Enter a name for the virtual machine.

NA-OCUM

Virtual machine names can contain up to 80 characters and they must be unique within each vCenter Server VM folder.

Select a location for the virtual machine.

nx-vc.flexpod.cisco.com
FlexPod-DC

Select a datacenter or VM folder to create the new virtual machine in.

4. Select the FlexPod-Management cluster and choose next to proceed to the next screen.

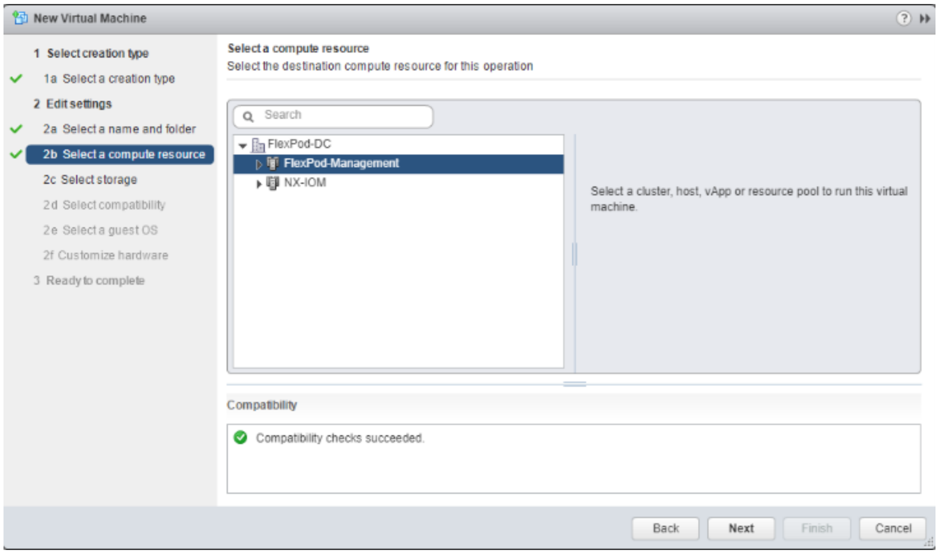

Select a compute resource

Select the destination compute resource for this operation

FlexPod-DC
FlexPod-Management
NX-IOM

Select a cluster, host, vApp or resource pool to run this virtual machine.

Compatibility

Compatibility checks succeeded.

5. Choose the datastore where the virtual machine will be deployed and leave the default VM storage policy unchanged.

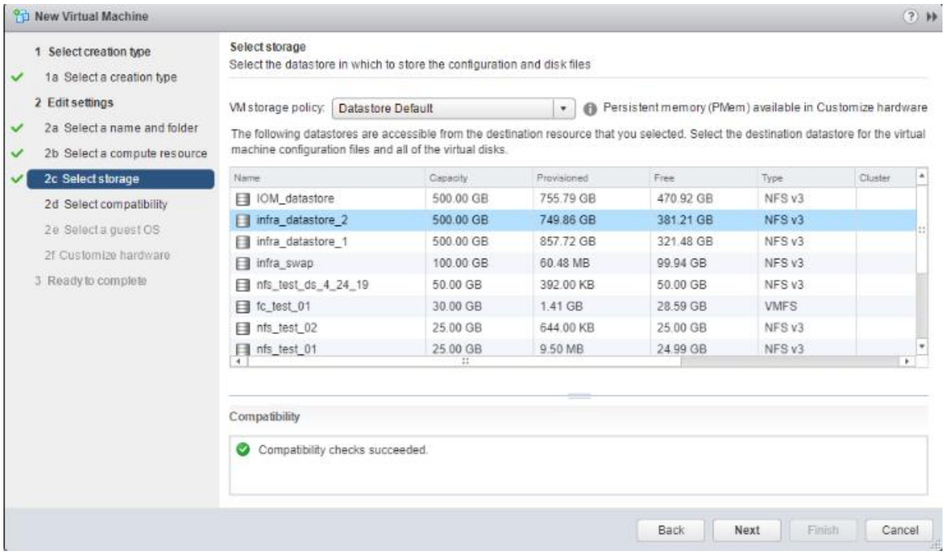

Select storage
Select the datastore in which to store the configuration and disk files

VM storage policy: Datastore Default

Persistent memory (PMem) available in Customize hardware

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Cluster
IOM_datastore	500.00 GB	755.79 GB	470.92 GB	NFS v3	
infra_datastore_2	500.00 GB	749.86 GB	381.21 GB	NFS v3	
infra_datastore_1	500.00 GB	857.72 GB	321.48 GB	NFS v3	
infra_swap	100.00 GB	60.48 MB	99.94 GB	NFS v3	
nfs_test_ds_4_24_19	50.00 GB	392.00 KB	50.00 GB	NFS v3	
fc_test_01	30.00 GB	1.41 GB	28.59 GB	VMFS	
nfs_test_02	25.00 GB	644.00 KB	25.00 GB	NFS v3	
nfs_test_01	25.00 GB	9.50 MB	24.99 GB	NFS v3	

Compatibility

✓ Compatibility checks succeeded.

6. Select next to advance the following screen.

7. Set the Guest OS Family to Linux and the Guest OS Version to CentOS 7 (64-bit).

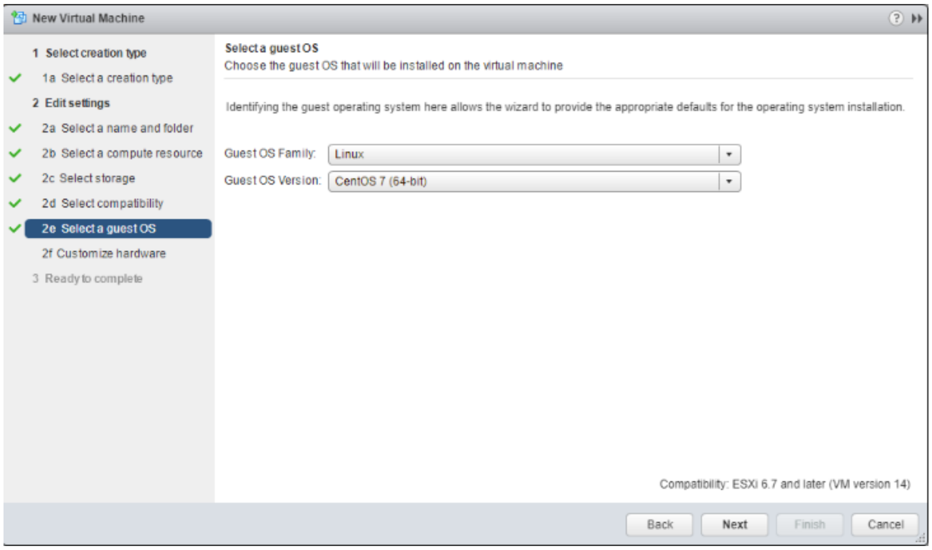

Select a guest OS
Choose the guest OS that will be installed on the virtual machine

Identifying the guest operating system here allows the wizard to provide the appropriate defaults for the operating system installation.

Guest OS Family: Linux

Guest OS Version: CentOS 7 (64-bit)

Compatibility: ESXi 6.7 and later (VM version 14)

8. Edit the settings of the virtual machine to the following specifications:

   a. CPU: 4 cores

   b. Memory: 16GB

   c. Hard disk 1: 200GB

   d. New Network: IB-MGMT Network

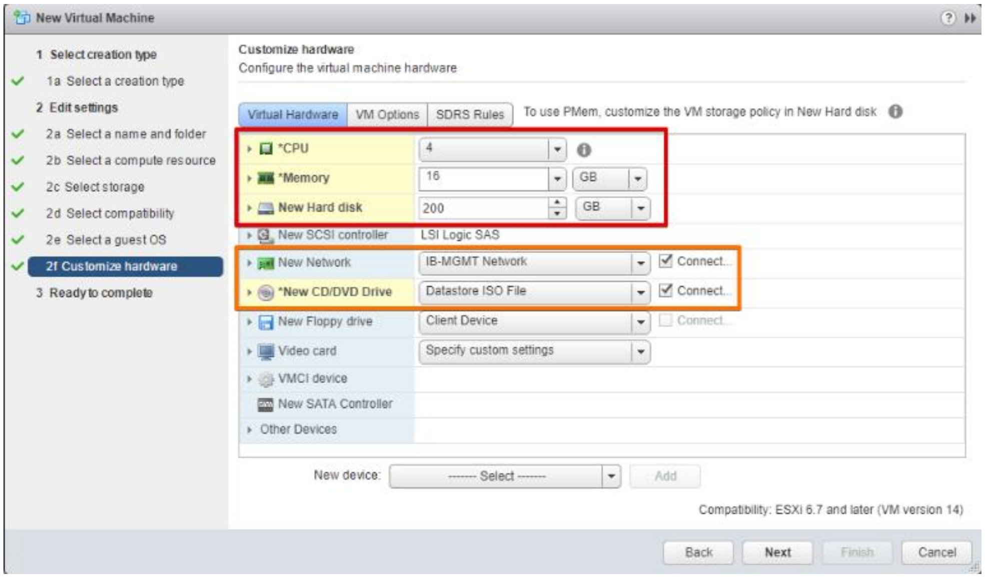

New Virtual Machine

1 Select creation type
✓ 1a Select a creation type
2 Edit settings
✓ 2a Select a name and folder
✓ 2b Select a compute resource
✓ 2c Select storage
✓ 2d Select compatibility
✓ 2e Select a guest OS
✓ 2f Customize hardware
3 Ready to complete

Customize hardware
Configure the virtual machine hardware

Virtual Hardware | VM Options | SDRS Rules — To use PMem, customize the VM storage policy in New Hard disk

*CPU	4		
*Memory	16	GB	
New Hard disk	200	GB	
New SCSI controller	LSI Logic SAS		
New Network	IB-MGMT Network		☑ Connect...
*New CD/DVD Drive	Datastore ISO File		☑ Connect...
New Floppy drive	Client Device		□ Connect...
Video card	Specify custom settings		
VMCI device			
New SATA Controller			
Other Devices			

New device: ------- Select ------- Add

Compatibility: ESXi 6.7 and later (VM version 14)

9. Connect the CD/DVD drive 1 by selecting the Connect check box and browse the host datastore for the CentOS 7 Linux ISO image uploaded to the datastore.

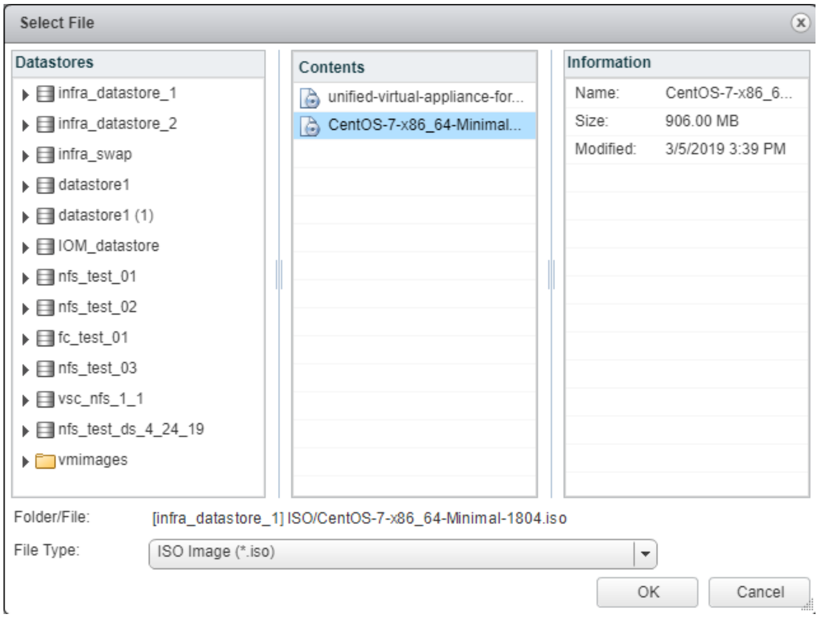

Select File

Datastores
- infra_datastore_1
- infra_datastore_2
- infra_swap
- datastore1
- datastore1 (1)
- IOM_datastore
- nfs_test_01
- nfs_test_02
- fc_test_01
- nfs_test_03
- vsc_nfs_1_1
- nfs_test_ds_4_24_19
- vmimages

Contents
- unified-virtual-appliance-for...
- CentOS-7-x86_64-Minimal...

Information

Name:	CentOS-7-x86_6...
Size:	906.00 MB
Modified:	3/5/2019 3:39 PM

Folder/File: [infra_datastore_1] ISO/CentOS-7-x86_64-Minimal-1804.iso

File Type: ISO Image (*.iso)

10. Choose next and review the virtual machine summary and select Finish.

11. Power-on the new virtual machine and open the console to observe the virtual machine booting to the CentOS 7 setup screen. Press Enter to start the installation.

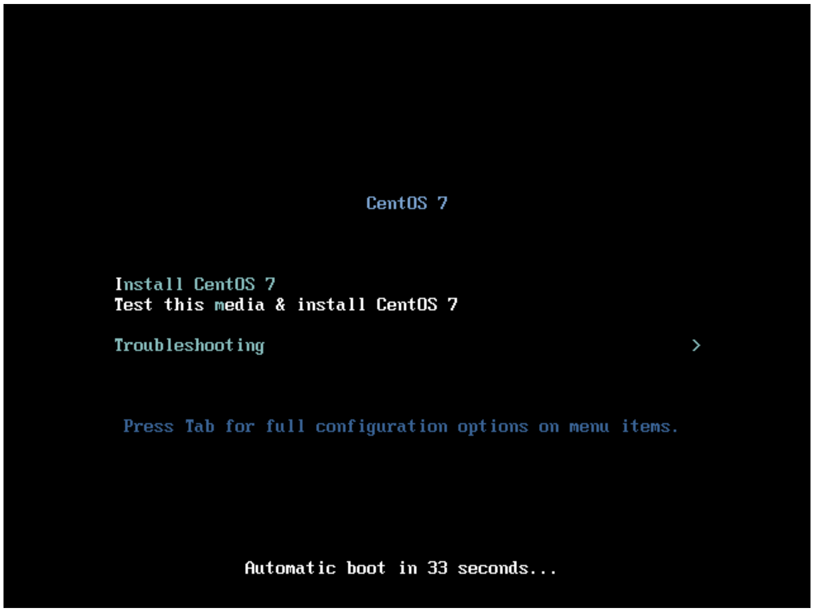

```
CentOS 7

Install CentOS 7
Test this media & install CentOS 7

Troubleshooting                                    >

Press Tab for full configuration options on menu items.

Automatic boot in 33 seconds...
```

12. Select the desired language and continue.

13. When prompted, enter the time zone information to specify the geographic area and region.

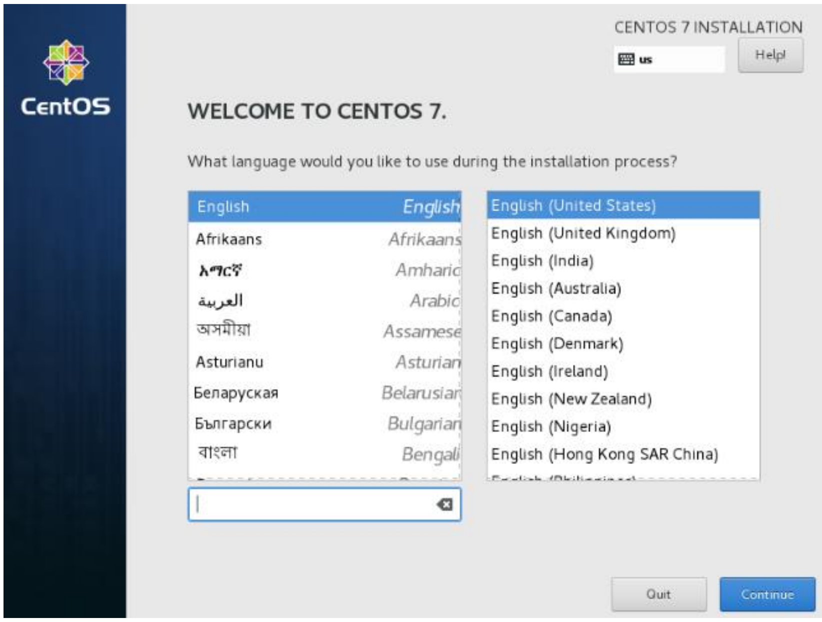

WELCOME TO CENTOS 7.

What language would you like to use during the installation process?

English	English
Afrikaans	Afrikaans
አማርኛ	Amharic
العربية	Arabic
অসমীয়া	Assamese
Asturianu	Asturian
Беларуская	Belarusian
Български	Bulgarian
বাংলা	Bengali

English (United States)
English (United Kingdom)
English (India)
English (Australia)
English (Canada)
English (Denmark)
English (Ireland)
English (New Zealand)
English (Nigeria)
English (Hong Kong SAR China)

14. Complete the required installation sections to specify the following information:

    a. Date & Time

    b. Keyboard layout

    c. Installation Destination

    d. Network & Hostname

15. When finished, begin the installation.

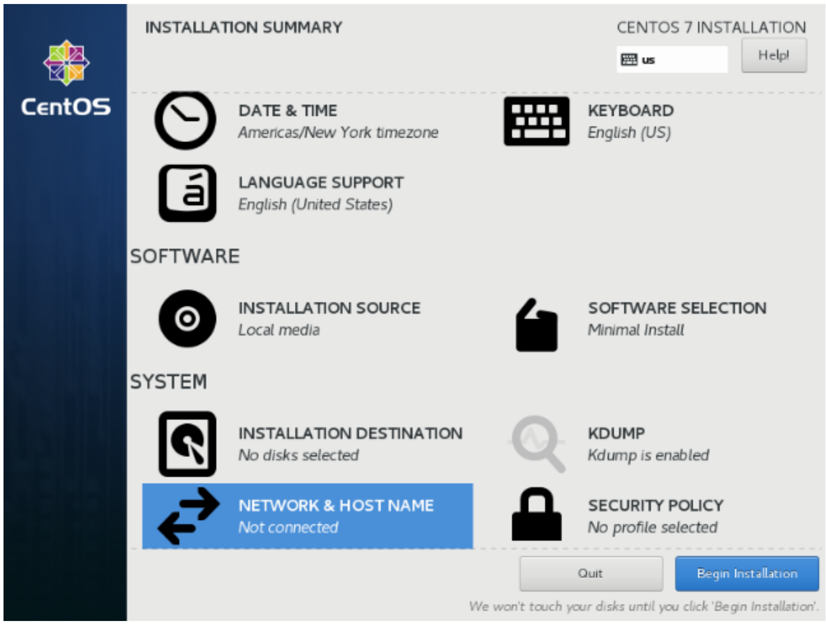

INSTALLATION SUMMARY

CENTOS 7 INSTALLATION

DATE & TIME
Americas/New York timezone

KEYBOARD
English (US)

LANGUAGE SUPPORT
English (United States)

SOFTWARE

INSTALLATION SOURCE
Local media

SOFTWARE SELECTION
Minimal Install

SYSTEM

INSTALLATION DESTINATION
No disks selected

KDUMP
Kdump is enabled

NETWORK & HOST NAME
Not connected

SECURITY POLICY
No profile selected

We won't touch your disks until you click 'Begin Installation'.

16. While the installation is occurring, set the password for the root user account.

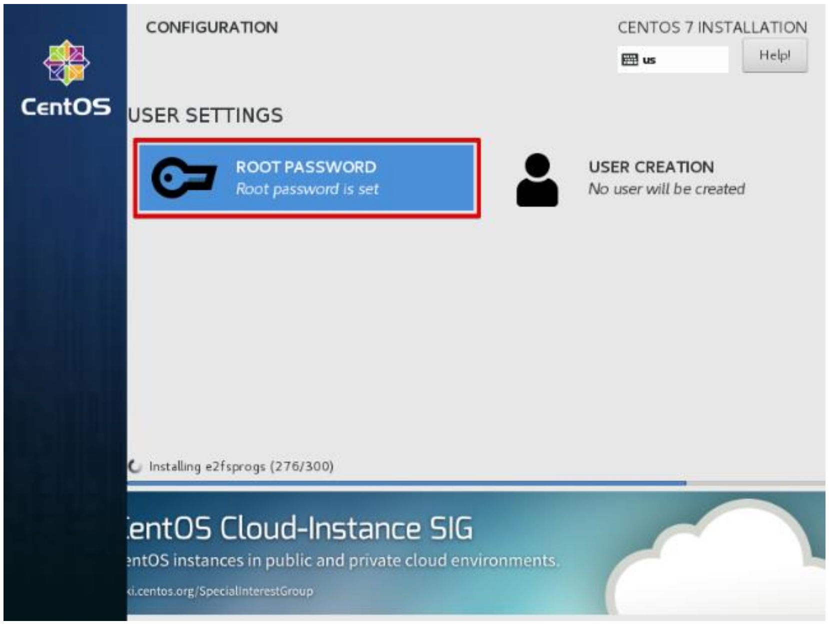

CONFIGURATION

CENTOS 7 INSTALLATION

USER SETTINGS

ROOT PASSWORD
Root password is set

USER CREATION
No user will be created

Installing e2fsprogs (276/300)

CentOS
us
Help!
entOS Cloud-Instance SIG
entOS instances in public and private cloud environments.
i.centos.org/SpecialInterestGroup

17. When the installation script is finished, a summary screen will display. Reboot the virtual machine to complete the installation and boot into CentOS Linux.

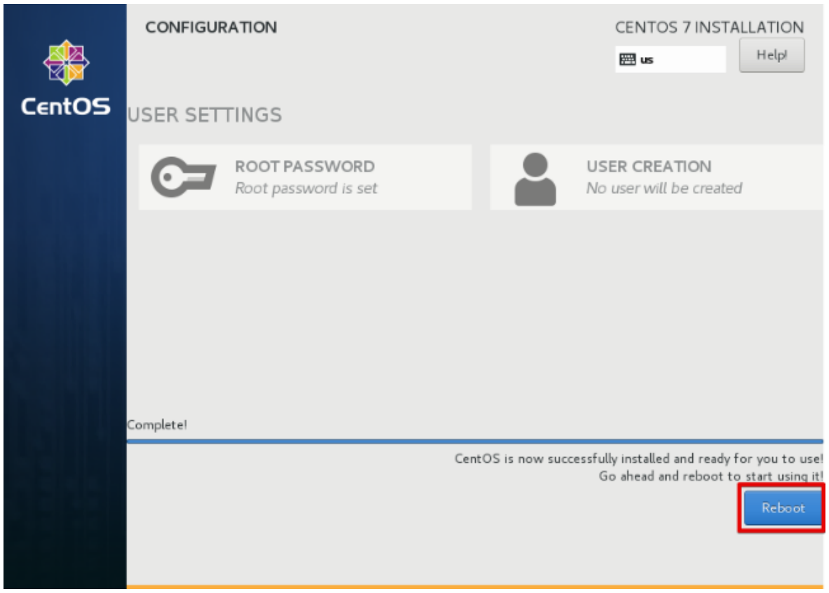

ROOT PASSWORD
Root password is set

USER CREATION
No user will be created

Complete!

CentOS is now successfully installed and ready for you to use!
Go ahead and reboot to start using it!

18. Log into the console of the CentOS virtual machine as the root user.

19. Verify IP connectivity and update the installation with the latest patches available for CentOS.

```
yum update
```

20. Download required tools needed to install the EPEL and MySQL repositories with the following steps.

```
yum install wget nano
```

Install the Extra Packages for Enterprise Linux (EPEL) and MySQL Repositories

To install the extra packages for EPEL and MySQL repositories, follow this step:

1. Download and install the EPEL and MySQL repositories with the following command:

```
wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
wget https://repo.mysql.com/yum/mysql-5.7-community/el/7/x86_64/mysql57-community-release-el7-
7.noarch.rpm

yum install epel-release-latest-7.noarch.rpm
yum install mysql57-community-release-el7-7.noarch.rpm
```

Install Open VM Tools on the CentOS Virtual Machine

To install the Open VM tools on the CentOS virtual machine, follow this step:

1. Run the following command from the console as the root user to install open-vm-tools:

```
yum install open-vm-tools
```

Configure CentOS Firewall

To configure the CentOS firewall, follow this step:

1. Run the following command to permit HTTPS traffic over TCP port 443 for secure OCUM communication.

```
firewall-cmd --zone=public --add-port=443/tcp --permanent; firewall-cmd --reload
```

Install OnCommand Unified Manager

To install the OnCommand Unified Manager, follow these steps:

1. Browse to the NetApp support website to download OnCommand Unified Manager 9.5 and review the security vulnerability disclosures concerning MySQL components in OnCommand Unified Manager.

2. Accept the license agreement and choose Accept & Continue to proceed. Select the *Red Hat and CentOS Linux Installation and Upgrade file* option from the list. Download the installation package .zip file and the associated SHA256 checksum file:

   a. OnCommandUnifiedManager-rhel7-9.5.zip

   b. OnCommandUnifiedManager-rhel7-9.5.zip.sha256.

3. Copy the OnCommand Unified Manager package and the OnCommand Unified Manager SHA256 checksum files to the /tmp directory of the CentOS virtual machine via SCP.

4. Compute the hash value of the OnCommand package and verify the file hash to ensure the integrity of the OnCommand Unified Manager installation package.

```
sha256sum OnCommandUnifiedManager-rhel7-9.5.zip
cat OnCommandUnifiedManager-rhel7-9.5.zip.sha256
```

5. Unzip the OnCommand Unified installation package and verify the contents are all present.

```
unzip OnCommandUnifiedManager-rhel7-9.5.zip
ls *.rpm
```

6. Review the output and ensure the following files were extracted to the /tmp directory:

   a. ocie-au-9.5.0-2018.11.J119.x86_64.rpm

   b. ocie-server-9.5.0-2018.11.J119.x86_64.rpm

   c. ocie-serverbase-9.5.0-2018.11.J53.x86_64.rpm

   d. netapp-application-server-9.5.0-2018.11.J53.x86_64.rpm

   e. netapp-platform-base-9.5.0-2018.11.J53.x86_64.rpm

   f. netapp-ocum-9.5-x86_64.rpm

7. Run the pre-installation script to ensure there are no system configuration settings that will conflict with the installation of OnCommand Unified Manager.

```
./pre_install_check.sh
```

8. If no issues are identified continue with the installation as show below:

```
yum install *.rpm
```

Do not attempt installation by using alternative commands (such as rpm -ivh ...). The successful installation of Unified Manager on a Red Hat Enterprise Linux or CentOS Linux system requires that all Unified Manager files and related files are installed in a specific order into a specific directory structure that is enforced automatically by the yum install *.rpm command.

9. After the installation messages are complete, navigate to the web UI of OnCommand Unified Manager with the IP address or the FQDN of the host via HTTPS.

10. Log into the OnCommand Unified Manager web UI with the maintenance user name (umadmin), and *admin* for the default password.

Add Storage Systems to OnCommand Unified Manager

To configure OnCommand Unified Manager and add a storage system for monitoring, follow these steps:

1. In a web browser navigate to the IP address or FQDN of the OnCommand Unified Manager installation and login with the maintenance user that was created in the previous section.

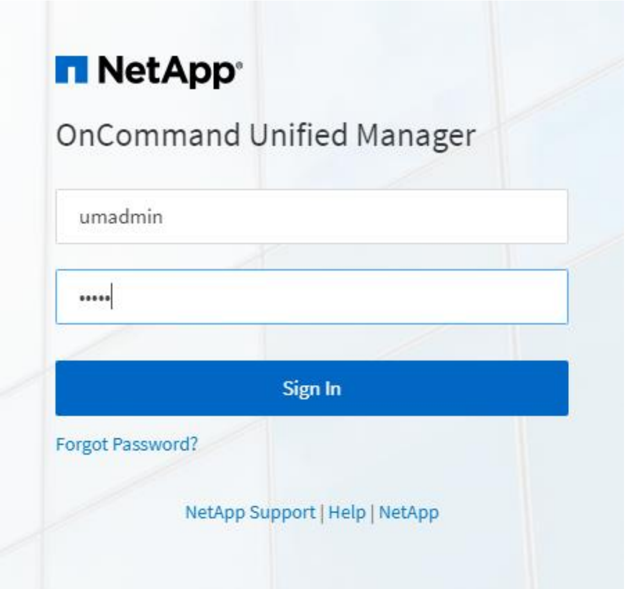

umadmin

•••••

2. Enter the email address that OnCommand will use to send alerts, enter the mail server configuration, and the IP address or hostname of the NTP server. Select Next and complete the AutoSupport configuration.

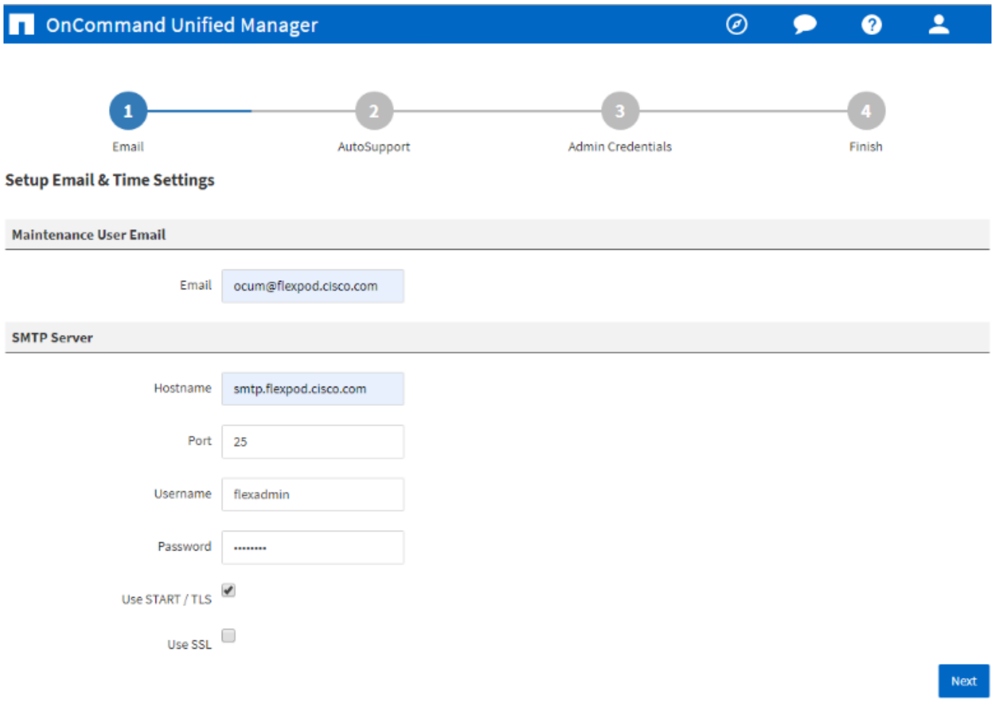

Setup Email & Time Settings

Maintenance User Email

Email: ocum@flexpod.cisco.com

SMTP Server

Hostname: smtp.flexpod.cisco.com

Port: 25

Username: flexadmin

Password: ••••••••

Use START / TLS ☑

Use SSL □

3. Change the default umadmin password by choosing the change password link on the Change Admin Credentials page.

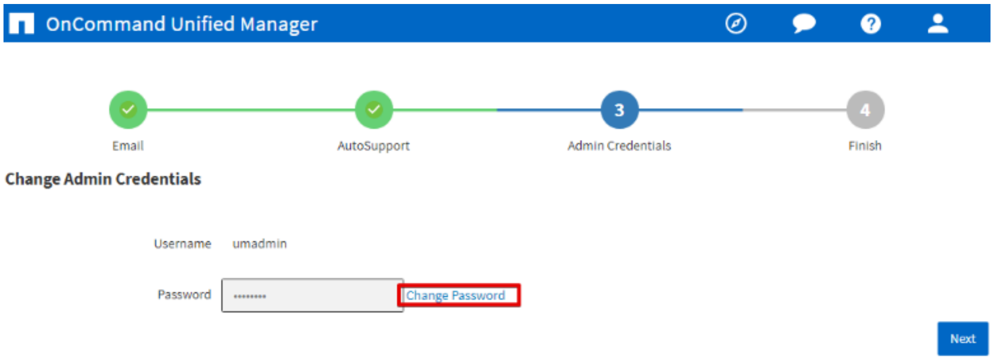

Change Admin Credentials

Username: umadmin

Password: ••••••••

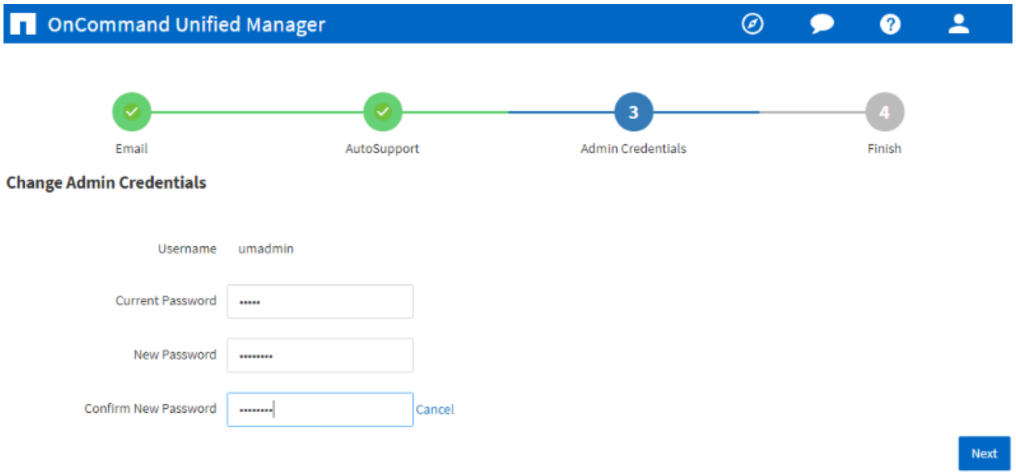

Change Admin Credentials

Username umadmin

Current Password •••••

New Password ••••••••

Confirm New Password ••••••••

4. Save and complete the configuration and continue to the next step.

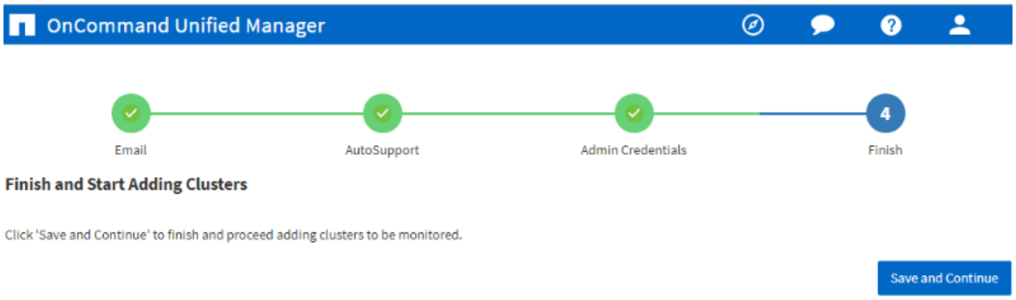

Finish and Start Adding Clusters

Click 'Save and Continue' to finish and proceed adding clusters to be monitored.

5. Click Add on the Configuration/Cluster Data Sources screen to add the storage cluster.

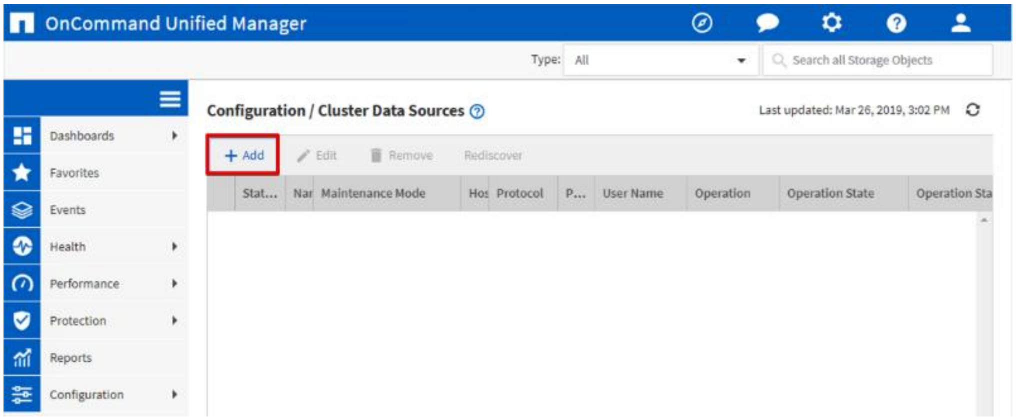

Configuration / Cluster Data Sources

Last updated: Mar 26, 2019, 3:02 PM

6. Enter the hostname and login credentials to the storage cluster.

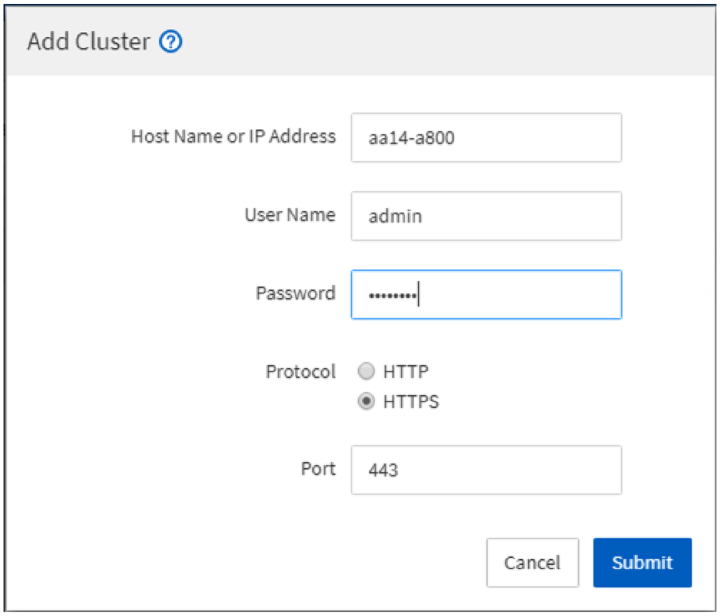

Add Cluster

Host Name or IP Address: aa14-a800

User Name: admin

Password: ••••••••

Protocol: ○ HTTP ● HTTPS

Port: 443

7. When prompted to trust the self-signed certificate from OnCommand Unified Manager, select Yes to finish and add the storage system. The initial discovery process can take up to 15 minutes to complete.

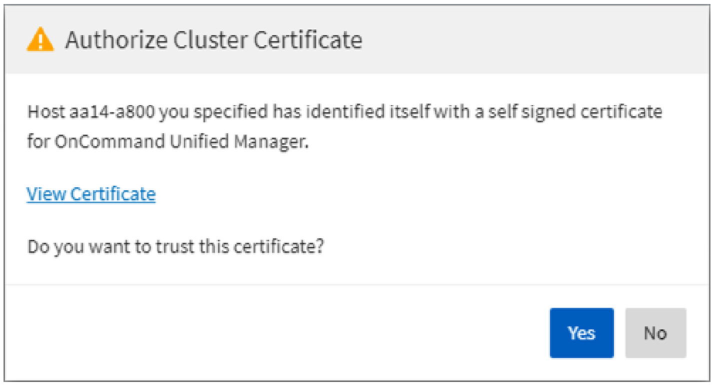

Authorize Cluster Certificate

Host aa14-a800 you specified has identified itself with a self signed certificate for OnCommand Unified Manager.

View Certificate

Do you want to trust this certificate?

Sample Tenant Provisioning

Provision a Sample Application Tenant

This section describes a sample procedure for provisioning an application tenant. The procedure refers to previous sections of this document and can be used as a guide and modified as needed when provisioning an application tenant.

1. Plan your application tenant and determine what storage protocols will be provided in the tenant. In the architecture explained in this document, fibre channel, NFS, iSCSI, and CIFS/SMB can be provided to the tenant. Also, plan what network VLANs the tenant will use. It is recommended to have a VLAN for virtual machine management traffic. One or two VLANs (iSCSI needs two if VMware RDM LUNs or iSCSI datastores will be provisioned) are also needed for each storage protocol used except fibre channel. Fibre channel will have new storage LIFs defined with the same VSANs configured for the FlexPod Infrastructure.

2. In the Nexus switches, declare all added VLANs and configure the VM VLAN as an allowed VLAN on the UCS port channels and the vPC peer link. Also, Layer 3 with HSRP or VRRP can be configured in the Nexus switches to provide this VLAN access to the outside. Layer 3 setup is not explained in this document but is explained in the Nexus 9000 documentation. Configure the storage VLANs on the UCS and storage port channels, and on the vPC peer link. The VM VLAN can also be added to the storage port channels in order to configure the tenant SVM management interface on this VLAN.

3. In the storage cluster:

    a. Create a broadcast domain with MTU 1500 for the tenant SVM management interface. Create a broadcast domain with MTU 9000 for each tenant storage protocol except fibre channel.

    b. Create VLAN interface ports on the node interface group on each node for tenant SVM management (VM VLAN) and for the VLAN for each storage protocol except fibre channel. Add these VLAN ports to the appropriate broadcast domains.

    c. Create the tenant SVM and follow all procedures in that section.

    d. Create Load-Sharing Mirrors for the tenant SVM.

    e. Create the FC service for the tenant SVM if fibre channel is being deployed in this tenant.

    f. Optionally, create a self-signed security certificate for the tenant SVM.

    g. Configure NFSv3 for the tenant SVM.

    h. Create a VM datastore volume in the tenant SVM.

    i. Create a once-a-day deduplication schedule on the VM datastore volume.

    j. If fibre channel is being deployed in this tenant, configure four FCP LIFs in the tenant SVM on the same fibre channel ports as in the Infrastructure SVM.

    k. Create an NFS LIF in the tenant SVM on each storage node.

    l. Create a boot LUN in the esxi_boot volume in the Infra-SVM for each tenant VMware ESXi host.

    m. Add the tenant SVM Administrator, SVM management LIF on the SVM management VLAN port, and default route for the SVM.

4. In Cisco UCS, one method of tenant setup is to dedicate a VMware ESXi cluster and set of UCS servers to each tenant. Service profiles will be generated for at least two tenant ESXi hosts. These hosts can boot from LUNs from the esxi_boot volume in the Infra-SVM but will also have access to FC storage in the tenant SVM.

    a. Create a Server Pool for the tenant ESXi host servers.
    b. Create all tenant VLANs in the LAN Cloud.
    c. Add the tenant VLANs to the vDS vNIC templates.
    d. Generate service profiles from the service profile template with the vMedia policy for the tenant ESXi hosts. Remember to bind these service profiles to the service profile template without the vMedia policy after VMware ESXi installation.

5. In the Cisco MDS 9132T switches:

    a. Create device aliases for the tenant ESXi host vHBAs and the FC LIFs in the tenant storage SVM.
    b. Create zones for the tenant ESXi hosts with fibre channel targets from both the storage Infra-SVM and the tenant SVM.
    c. Add these zones to the Fabric zoneset and activate the zoneset.

6. In the storage cluster:

    a. Create igroups for the tenant ESXi hosts in both the Infra-SVM and tenant SVM. Also, create an igroup in the tenant SVM that includes the WWPNs for all tenant ESXi hosts to support shared storage from the tenant SVM.
    b. In Infra-SVM, map the boot LUNs created earlier to the tenant ESXi hosts. Tenant FC storage can be created later using NetApp VSC.

7. Install and configure VMware ESXi on all tenant host servers. It is not necessary to map infra_datastore_1 or infra_datastore_2 unless you want the tenant ESXi hosts to have access to VMs or VM templates in these datastores.

8. In VMware vCenter, create a cluster for the tenant ESXi hosts. Add the hosts to the cluster.

9. Using the vCenter Web Client, add the tenant hosts to the infrastructure vDS or create a tenant vDS and add the hosts to it. In the VMware vDS, add port-profiles for the tenant VLANs. When migrating the hosts to the vDS, leave only the ESXi management interfaces on vSwitch0.

10. Back in vCenter, add in any necessary VMkernel ports for storage interfaces remembering to set the MTU correctly on these interfaces. Mount the tenant NFS datastore on the tenant cluster if one was created. Tenant iSCSI VMkernel ports can be created on the vDS with the port groups pinned to the appropriate fabric.

11. Using the NetApp VSC plugin to the vCenter Web Client, set recommended values for all tenant ESXi hosts. Install the NetApp NFS Plug-in for VMware VAAI for all tenant hosts and reboot each host.

12. You can now begin provisioning virtual machines on the tenant cluster. The NetApp VSC plugin can be used to provision fibre channel, iSCSI, and NFS datastores.

13. Optionally, use NetApp SnapCenter to provision backups of tenant virtual machines.

Appendix

FlexPod iSCSI Addition

Cisco Nexus Switch Configuration

This section is a delta section for adding infrastructure iSCSI to the Nexus switches. This section should be executed after the Cisco Nexus Switch Configuration section in the main document is completed.

Create Infrastructure iSCSI VLANs on Cisco Nexus A and Cisco Nexus B

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

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

```
vlan <infra-iscsi-a-vlan-id>
name Infra-iSCSI-A-VLAN
vlan <infra-iscsi-b-vlan-id>
name Infra-iSCSI-B-VLAN
exit
```

Add Infrastructure iSCSI VLANs to Port-Channels on Cisco Nexus A and Cisco Nexus B

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

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

```
interface Po10,Po113,Po114,Po15,Po16
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
copy run start
```

NetApp Storage Configuration

When using the iSCSI protocol for SAN boot connectivity, use the Storage Configuration steps outlined in the body of this document. Where appropriate, replace the Fibre Channel configuration steps with the steps listed here.

If the iSCSI license was not installed during the cluster configuration, make sure to install the license before creating the iSCSI service.

Create Block Protocol (iSCSI) Service

If the FCP protocol is not being used in the environment it should be removed from the vserver configured in a previous step. If FCP will be used in addition to iSCSI or in the future, step 1 can be omitted.

1. Remove FCP protocol from the vserver.

```
vserver remove-protocols -vserver Infra-SVM -protocols fcp
```

2. Create the iSCSI block service.

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

Create iSCSI Broadcast Domains

To create the broadcast domains for each of the iSCSI VLANs, run the following commands:

```
network port broadcast-domain create -broadcast-domain Infra-iSCSI-A -mtu 9000
network port broadcast-domain create -broadcast-domain Infra-iSCSI-B -mtu 9000
```

Create iSCSI VLANs

To create iSCSI VLANs, follow these steps:

1. Modify the MTU size on the parent interface group hosting the iSCSI traffic using the following commands:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

2. Create VLAN ports for the iSCSI LIFs on each storage controller.

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-b-vlan-id>

network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-b-vlan-id>
```

Add VLANs to iSCSI Broadcast Domains

To add each of the iSCSI VLAN ports to the corresponding broadcast domain, run the following commands:

```
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-node01>:a0a-
<infra-iscsi-a-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-node01>:a0a-
<infra-iscsi-b-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-node02>:a0a-
<infra-iscsi-a-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-node02>:a0a-
<infra-iscsi-b-vlan-id>

network port broadcast-domain show
```

Create iSCSI LIFs

To create four iSCSI LIFs, run the following commands (two on each node):

```
network interface create -vserver Infra-SVM -lif iscsi-lif01a -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<infra-iscsi-a-vlan-id> -status-admin up

network interface create -vserver Infra-SVM -lif iscsi-lif01b -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<infra-iscsi-b-vlan-id> -status-admin up

network interface create -vserver Infra-SVM -lif iscsi-lif02a -role data -data-protocol iscsi -home-
node <st-node02> -home-port a0a-<infra-iscsi-a-vlan-id> -status-admin up

network interface create -vserver Infra-SVM -lif iscsi-lif02b -role data -data-protocol iscsi -home-
node <st-node02> -home-port a0a-<infra-iscsi-b-vlan-id> -status-admin up
```

```
network interface show
```

Create igroups

Table 11 iSCSI IQN for SVM

SVM Name	SVM Target IQN
Infra-SVM	

Table 12 iSCSI vNIC IQN Configuration

Cisco UCS Service Profile Name	iSCSI IQN	Variable
vm-host-infra-01		<vm-host-infra-01-iqn>
vmhost-infra-02		<vm-host-infra-02-iqn>

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

Create igroups by entering the following commands from the storage cluster management LIF SSH connection:

```
lun igroup create –vserver Infra-SVM –igroup vm-host-infra-01 –protocol iscsi –ostype vmware –
initiator <vm-host-infra-01-iqn>

lun igroup create –vserver Infra-SVM –igroup vm-host-infra-02 –protocol iscsi –ostype vmware –
initiator <vm-host-infra-02-iqn>
```

Use the values listed in Table 8 and Table 9 for the IQN information.

To view the two igroups just created, use the command lun igroup show.

```
lun igroup show -protocol iscsi
```

Map Boot LUNs to igroups

From the storage cluster management LIF SSH connection, run the following commands:

```
lun mapping create –vserver Infra-SVM –path /vol/esxi_boot/vm-host-infra-01 –igroup vm-host-infra-01
–lun-id 0

lun mapping create –vserver Infra-SVM –path /vol/esxi_boot/vm-host-infra-02 –igroup vm-host-infra-02
–lun-id 0
```

Cisco UCS iSCSI Configuration

The following subsections can be completed to add infrastructure iSCSI to the Cisco UCS. These subsections can be completed in place of the subsections in the Cisco UCS Configuration section of this document labeled (FCP), or they can be completed in addition to the FCP sections to have the option of FCP or iSCSI boot.

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Expand Pools > root.
3. Right-click IQN Pools.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-Pool for the name of the IQN pool
6. Optional: Enter a description for the IQN pool
7. Enter iqn.2010-11.com.flexpod as the prefix.
8. Select Sequential for Assignment Order
9. Click Next.
10. Click Add.
11. Enter ucs-host as the suffix.

If multiple Cisco UCS domains are being used, a more specific IQN suffix may need to be used.

12. Enter 1 in the From field.
13. Specify the size of the IQN block sufficient to support the available server resources.

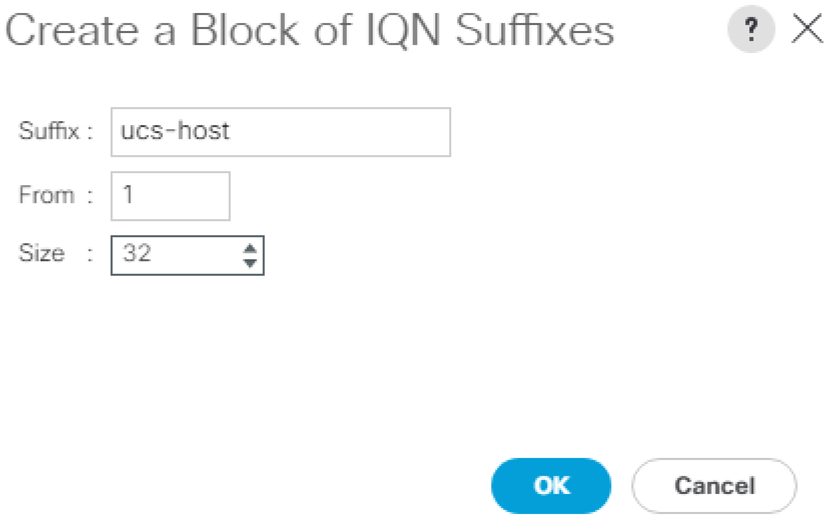
Create a Block of IQN Suffixes

Suffix : ucs-host

From : 1

Size : 32

14. Click OK.
15. Click Finish and OK to complete creating the IQN pool.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment within the FlexPod Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Pools > root > Sub-Organizations > FlexPod Organization.
3. Right-click IP Pools under the FlexPod Organization.
4. Select Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Select Sequential for the assignment order.
8. Click Next.
9. Click Add to add a block of IP addresses.
10. In the From field, enter the beginning of the range to assign as iSCSI boot IP addresses on Fabric A.
11. Set the size to enough addresses to accommodate the servers.
12. Enter the appropriate Subnet Mask.
13. Click OK.
14. Click Next.
15. Click Finish and OK to complete creating the Fabric A iSCSI IP Pool.
16. Right-click IP Pools under the FlexPod Organization.
17. Select Create IP Pool.
18. Enter iSCSI-IP-Pool-B as the name of IP pool.
19. Optional: Enter a description for the IP pool.
20. Select Sequential for the assignment order.
21. Click Next.
22. Click Add to add a block of IP addresses.
23. In the From field, enter the beginning of the range to assign as iSCSI IP addresses on Fabric B.
24. Set the size to enough addresses to accommodate the servers.

25. Enter the appropriate Subnet Mask.

26. Click OK.

27. Click Next.

28. Click Finish and OK to complete creating the Fabric B iSCSI IP Pool.

Create iSCSI VLANs

To configure the necessary iSCSI virtual local area networks (VLANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand LAN > LAN Cloud.

3. Right-click VLANs.

4. Select Create VLANs.

5. Enter Infra-iSCSI-A as the name of the VLAN to be used for iSCSI-A.

6. Keep the Common/Global option selected for the scope of the VLAN.

7. Enter the Infra-iSCSI-A VLAN ID.

8. Keep the Sharing Type as None.

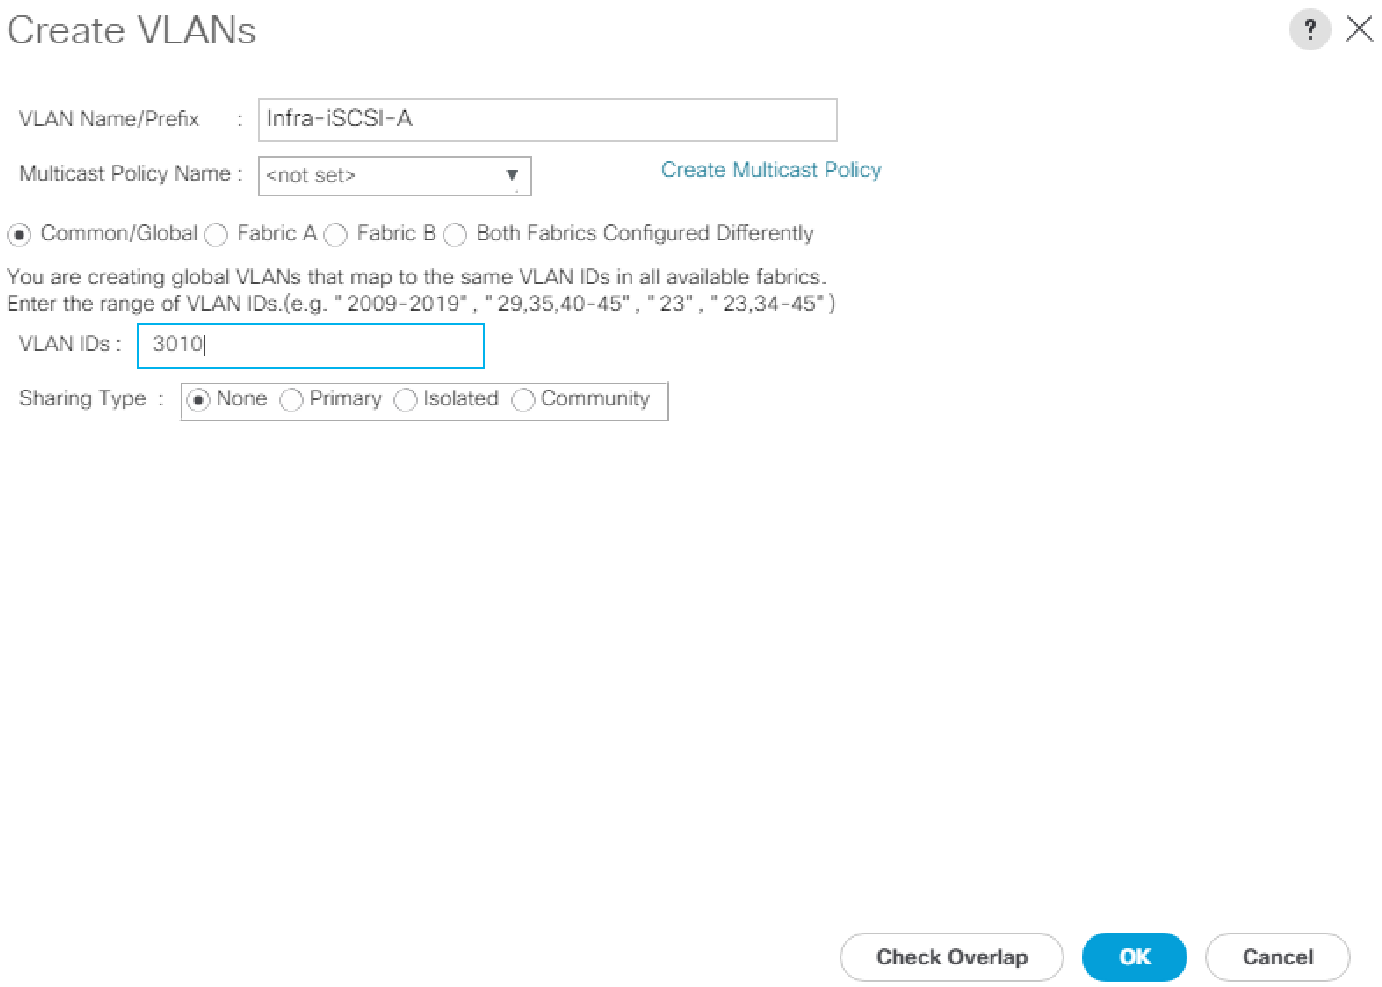

Create VLANs

VLAN Name/Prefix : Infra-iSCSI-A

Multicast Policy Name : <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. " 2009-2019" , " 29,35,40-45" , " 23" , " 23,34-45")

VLAN IDs : 3010

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

9. Click OK and then click OK again.

10. Right-click VLANs.

11. Select Create VLANs.

12. Enter Infra-iSCSI-B as the name of the VLAN to be used for iSCSI-B.

13. Keep the Common/Global option selected for the scope of the VLAN.

14. Enter the iSCSI-B VLAN ID.

15. Keep the Sharing Type as None.

16. Click OK and then click OK again.

Create iSCSI vNIC Templates

To create iSCSI virtual network interface card (vNIC) templates for the Cisco UCS environment within the FlexPod Organization, follow these steps:

1. Select LAN.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click vNIC Templates under the FlexPod Organization.
4. Select Create vNIC Template.
5. Enter iSCSI-Template-A as the vNIC template name.
6. Select Fabric A. Do not select the Enable Failover checkbox.
7. Leave Redundancy Type set at No Redundancy.
8. Under Target, make sure that only the Adapter checkbox is selected.
9. Select Updating Template for Template Type.
10. Under VLANs, select only Infra-iSCSI-A.
11. Select Infra-iSCSI-A as the native VLAN.
12. Leave vNIC Name set for the CDN Source.
13. Under MTU, enter 9000.
14. From the MAC Pool list, select MAC-Pool-A.
15. From the Network Control Policy list, select Enable-CDP-LLDP.

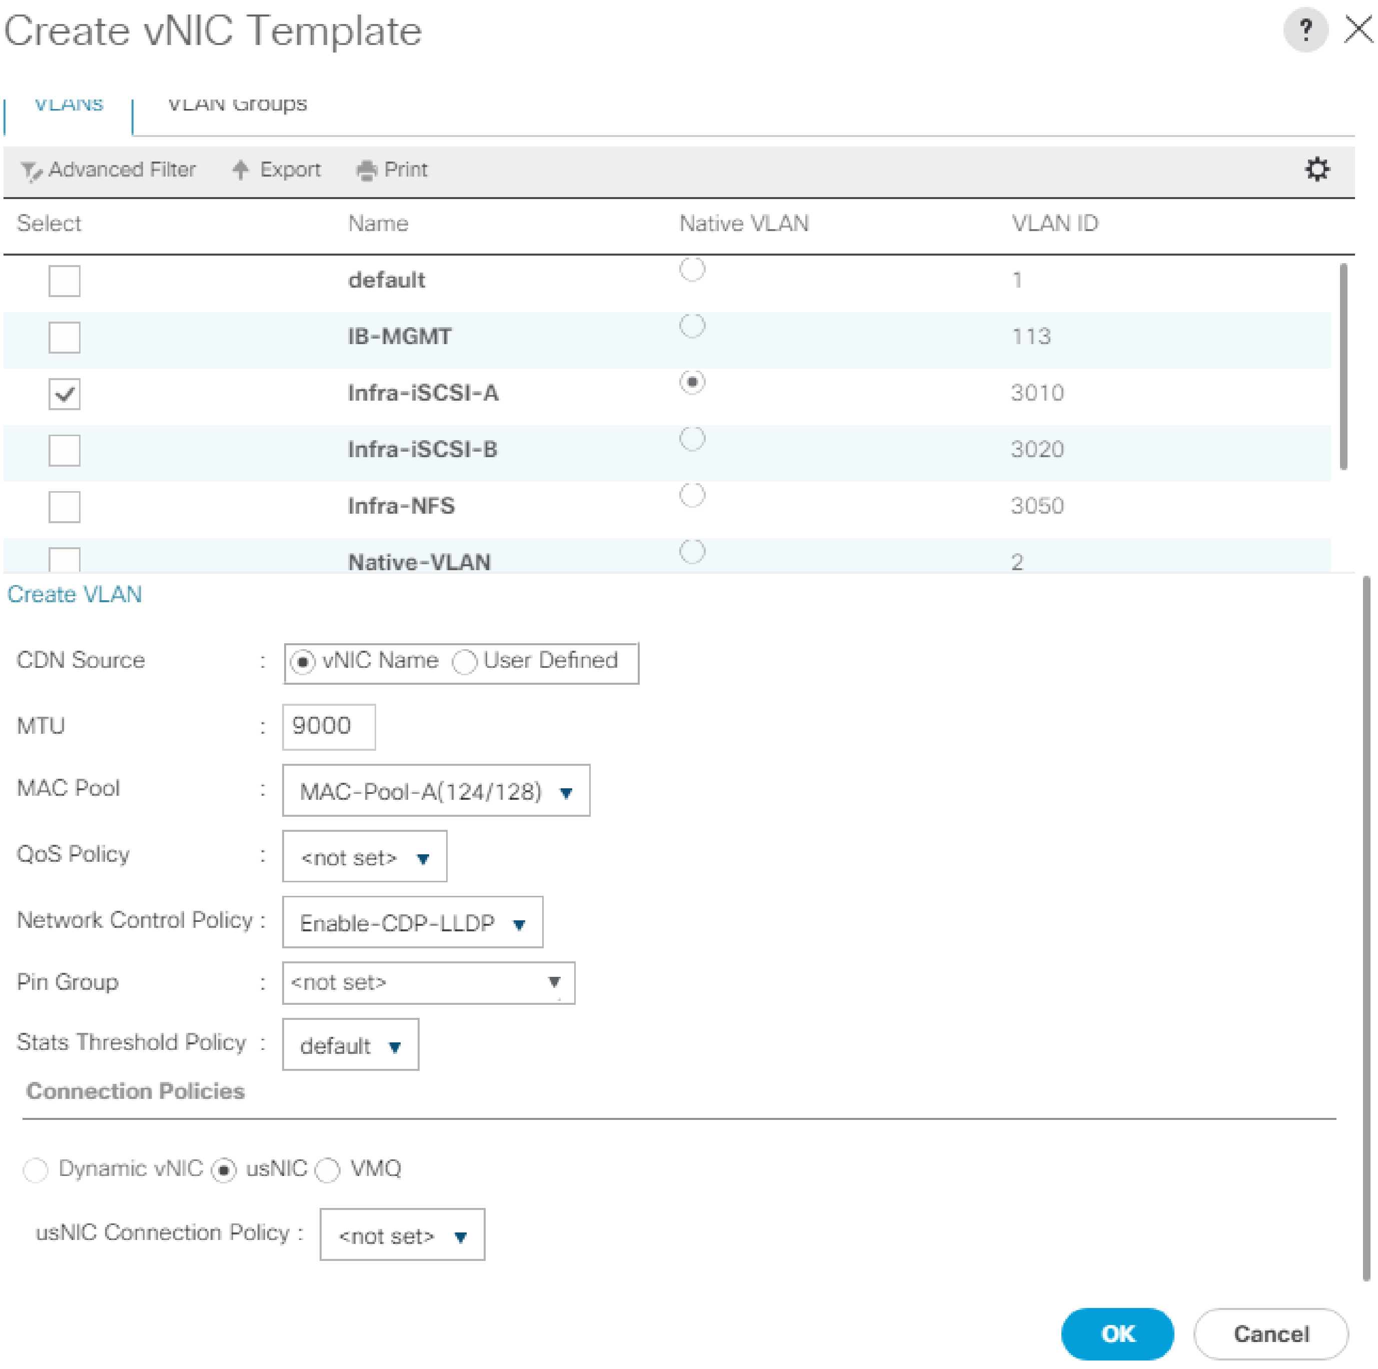

Create vNIC Template

Select	Name	Native VLAN	VLAN ID
□	default	○	1
□	IB-MGMT	○	113
☑	Infra-iSCSI-A	●	3010
□	Infra-iSCSI-B	○	3020
□	Infra-NFS	○	3050
□	Native-VLAN	○	2

Create VLAN

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(124/128)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

○ Dynamic vNIC ● usNIC ○ VMQ

usNIC Connection Policy : <not set>

16. Click OK to complete creating the vNIC template.

17. Click OK.

18. Right-click vNIC Templates.

19. Select Create vNIC Template.

20. Enter iSCSI-Template-B as the vNIC template name.

21. Select Fabric B. Do not select the Enable Failover checkbox.

22. Leave Redundancy Type set at No Redundancy.

23. Under Target, make sure that only the Adapter checkbox is selected.

24. Select Updating Template for Template Type.

25. Under VLANs, select only Infra-iSCSI-B.

26. Select Infra-iSCSI-B as the native VLAN.

27. Leave vNIC Name set for the CDN Source.

28. Under MTU, enter 9000.

29. From the MAC Pool list, select MAC-Pool-B.

30. From the Network Control Policy list, select Enable-CDP-LLDP.

31. Click OK to complete creating the vNIC template.

32. Click OK.

Create LAN Connectivity Policy for iSCSI Boot

To configure the necessary Infrastructure LAN Connectivity Policy within the FlexPod Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand LAN > Policies > root > Sub-Organizations > FlexPod Organization.

3. Right-click LAN Connectivity Policies under the FlexPod Organization.

4. Select Create LAN Connectivity Policy.

5. Enter iSCSI-Boot as the name of the policy.

6. Click the upper Add button to add a vNIC.

7. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.

8. Select the Use vNIC Template checkbox.

9. In the vNIC Template list, select Infra-A.

10. In the Adapter Policy list, select VMWare.

11. Click OK to add this vNIC to the policy.

12. Click the upper Add button to add another vNIC to the policy.
13. In the Create vNIC box, enter 01-Infra-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select Infra-B.
16. In the Adapter Policy list, select VMWare.
17. Click OK to add the vNIC to the policy.
18. Click the upper Add button to add a vNIC.
19. In the Create vNIC dialog box, enter 02-Infra-vDS-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select Infra-vDS-A.
22. In the Adapter Policy list, select VMWare.
23. Click OK to add this vNIC to the policy.
24. Click the upper Add button to add another vNIC to the policy.
25. In the Create vNIC box, enter 03-Infra-vDS-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select Infra-vDS-B.
28. In the Adapter Policy list, select VMWare.
29. Click OK to add the vNIC to the policy.
30. Click the upper Add button to add a vNIC.
31. In the Create vNIC dialog box, enter 04-iSCSI-A as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select iSCSI-Template-A.
34. In the Adapter Policy list, select VMWare.
35. Click OK to add this vNIC to the policy.
36. Click Add to add a vNIC to the policy.
37. In the Create vNIC dialog box, enter 05-iSCSI-B as the name of the vNIC.

38. Select the Use vNIC Template checkbox.

39. In the vNIC Template list, select iSCSI-Template-B.

40. In the Adapter Policy list, select VMWare.

41. Click OK to add this vNIC to the policy.

42. Expand Add iSCSI vNICs.

43. Select Add in the Add iSCSI vNICs section.

44. Set the name to iSCSI-Boot-A.

45. Select 04-iSCSI-A as the Overlay vNIC.

46. Set the iSCSI Adapter Policy to default.

47. Leave the VLAN set to Infra-iSCSI-A (native).

48. Leave the MAC Address set to None.

49. Click OK.

50. Select Add in the Add iSCSI vNICs section.

51. Set the name to iSCSI-Boot-B.

52. Select 05-iSCSI-B as the Overlay vNIC.

53. Set the iSCSI Adapter Policy to default.

54. Leave the VLAN set to Infra-iSCSI-B (native).

55. Leave the MAC Address set to None.

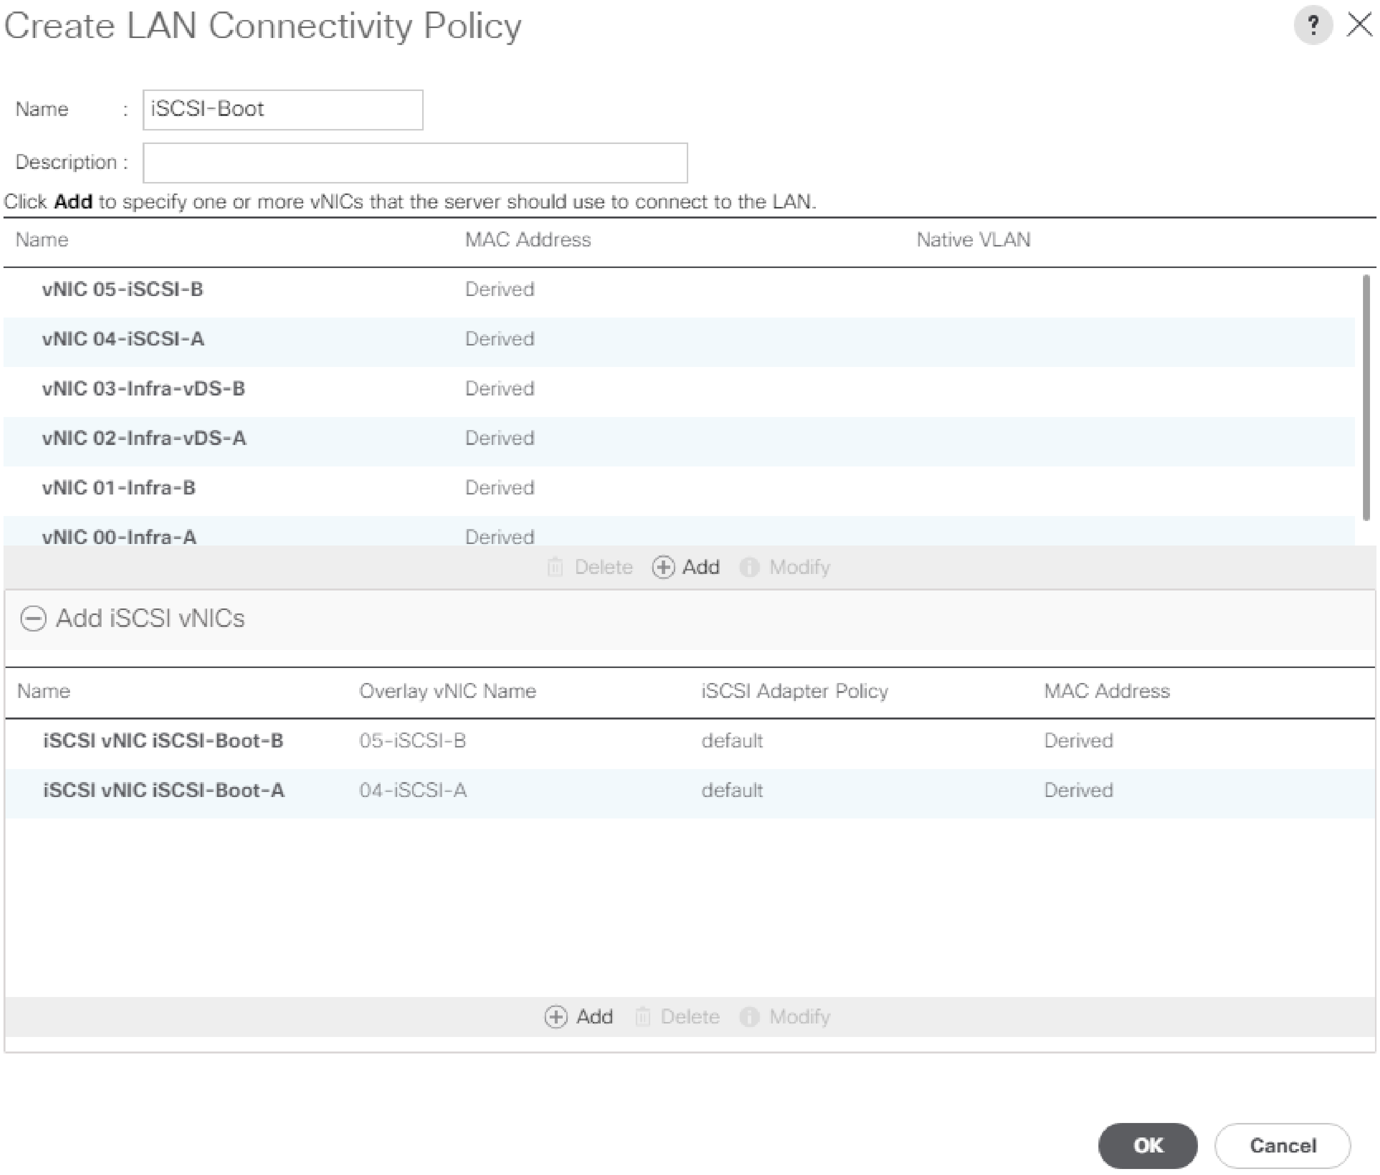

Create LAN Connectivity Policy

Name : iSCSI-Boot

Description :

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Native VLAN
vNIC 05-iSCSI-B	Derived	
vNIC 04-iSCSI-A	Derived	
vNIC 03-Infra-vDS-B	Derived	
vNIC 02-Infra-vDS-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-B	05-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-Boot-A	04-iSCSI-A	default	Derived

56. Click OK, then click OK again to create the LAN Connectivity Policy.

Create iSCSI Boot Policy

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi-lif01a and iscsi-lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi-lif02a and iscsi-lif02b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).

One boot policy is configured in this procedure. The policy configures the primary target to be iscsi-lif01a.

It is recommended not use UEFI Secure Boot for iSCSI boot of UCS servers at this time. UEFI Secure Boot will not be set up in this procedure.

To create a boot policy for the Cisco UCS environment within the FlexPod Organization, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click Boot Policies under the FlexPod Organization.
4. Select Create Boot Policy.
5. Enter Boot-iSCSI-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Leave the Legacy Boot Mode selected.
9. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
10. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
11. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.
12. Click OK.
13. Select Add iSCSI Boot.
14. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-B.
15. Click OK.
16. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

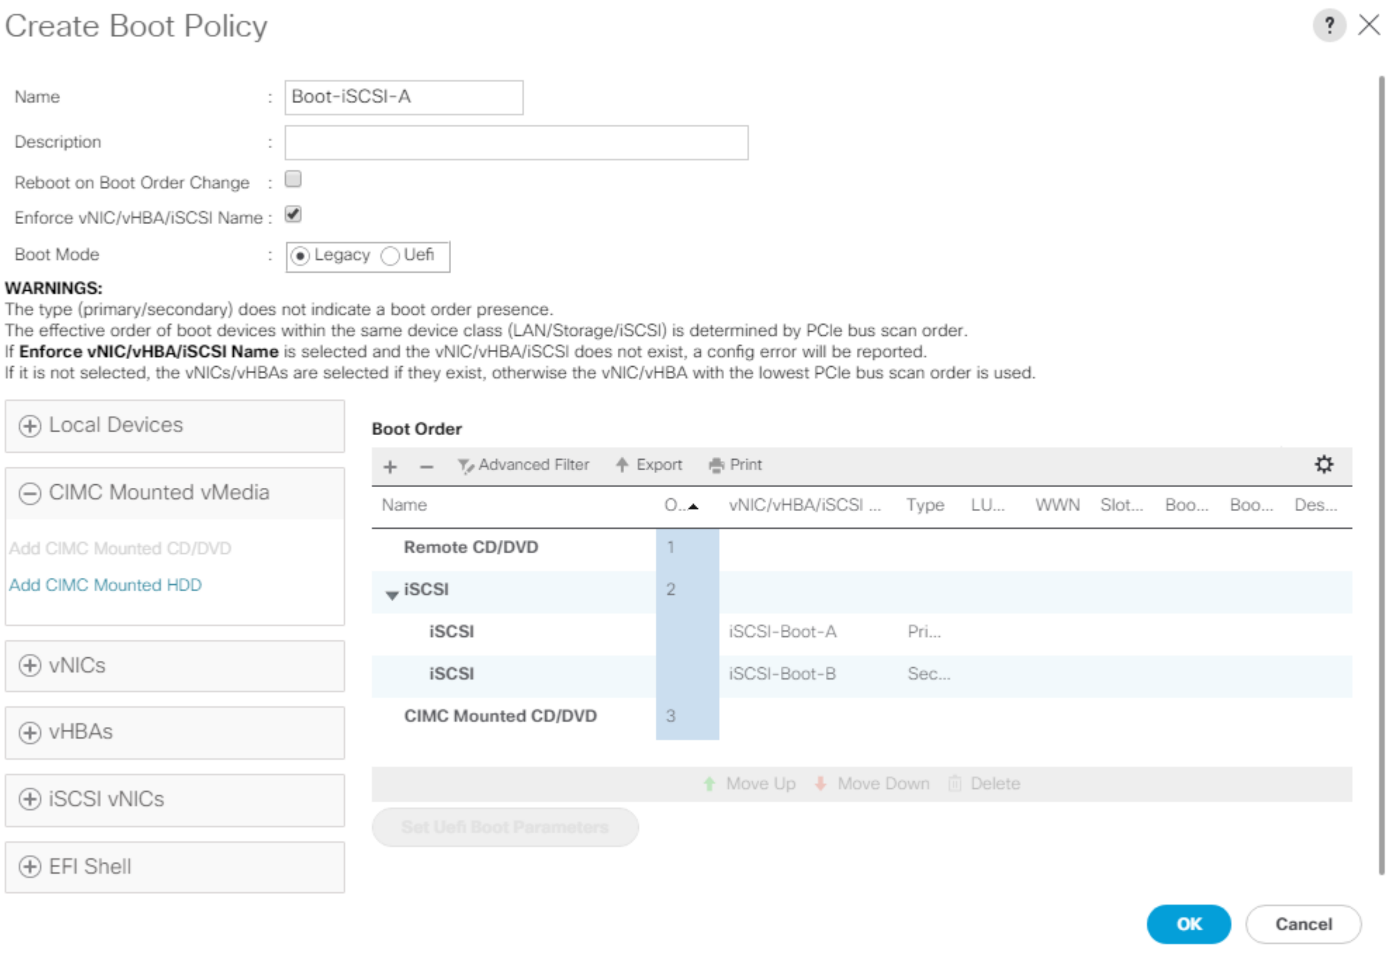

Create Boot Policy

Name : Boot-iSCSI-A

Description :

Reboot on Boot Order Change : □

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ● Legacy ○ Uefi

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	O..	vNIC/vHBA/iSCSI ...	Type	LU...	WWN	Slot...	Boo...	Boo...	Des...
Remote CD/DVD	1								
▾ iSCSI	2								
iSCSI		iSCSI-Boot-A	Pri...						
iSCSI		iSCSI-Boot-B	Sec...						
CIMC Mounted CD/DVD	3								

17. Click OK then click OK again to create the policy.

Create iSCSI Boot Service Profile Template

In this procedure, one service profile template for Infrastructure ESXi hosts within the FlexPod Organization is created for Fabric A boot.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > FlexPod Organization.
3. Right-click the FlexPod Organization.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.

6. Select the "Updating Template" option.

7. Under UUID Assignment, select UUID_Pool.

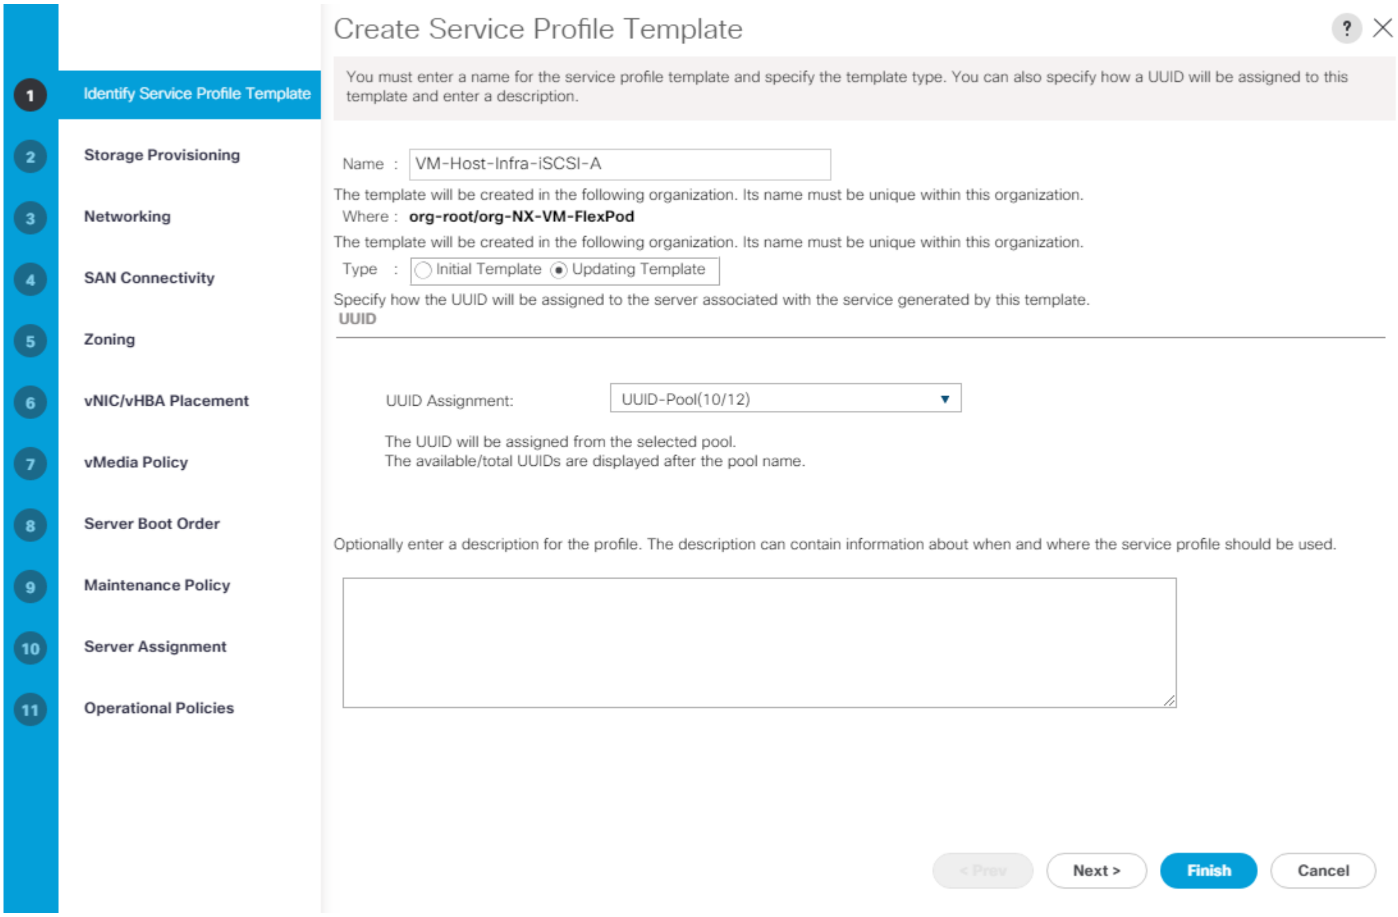

Create Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name : VM-Host-Infra-iSCSI-A

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-NX-VM-FlexPod**

The template will be created in the following organization. Its name must be unique within this organization.

Type : ○ Initial Template ● Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

UUID

UUID Assignment: UUID-Pool(10/12)

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

8. Click Next.

Configure Storage Provisioning

To configure the storage provisioning, follow these steps:

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy tab and select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.

2. Click Next.

Configure Networking Options

To configure the network options, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.

2. Select the "Use Connectivity Policy" option to configure the LAN connectivity.

3. Select iSCSI-Boot from the LAN Connectivity Policy drop-down list.

4. Select IQN_Pool in Initiator Name Assignment.

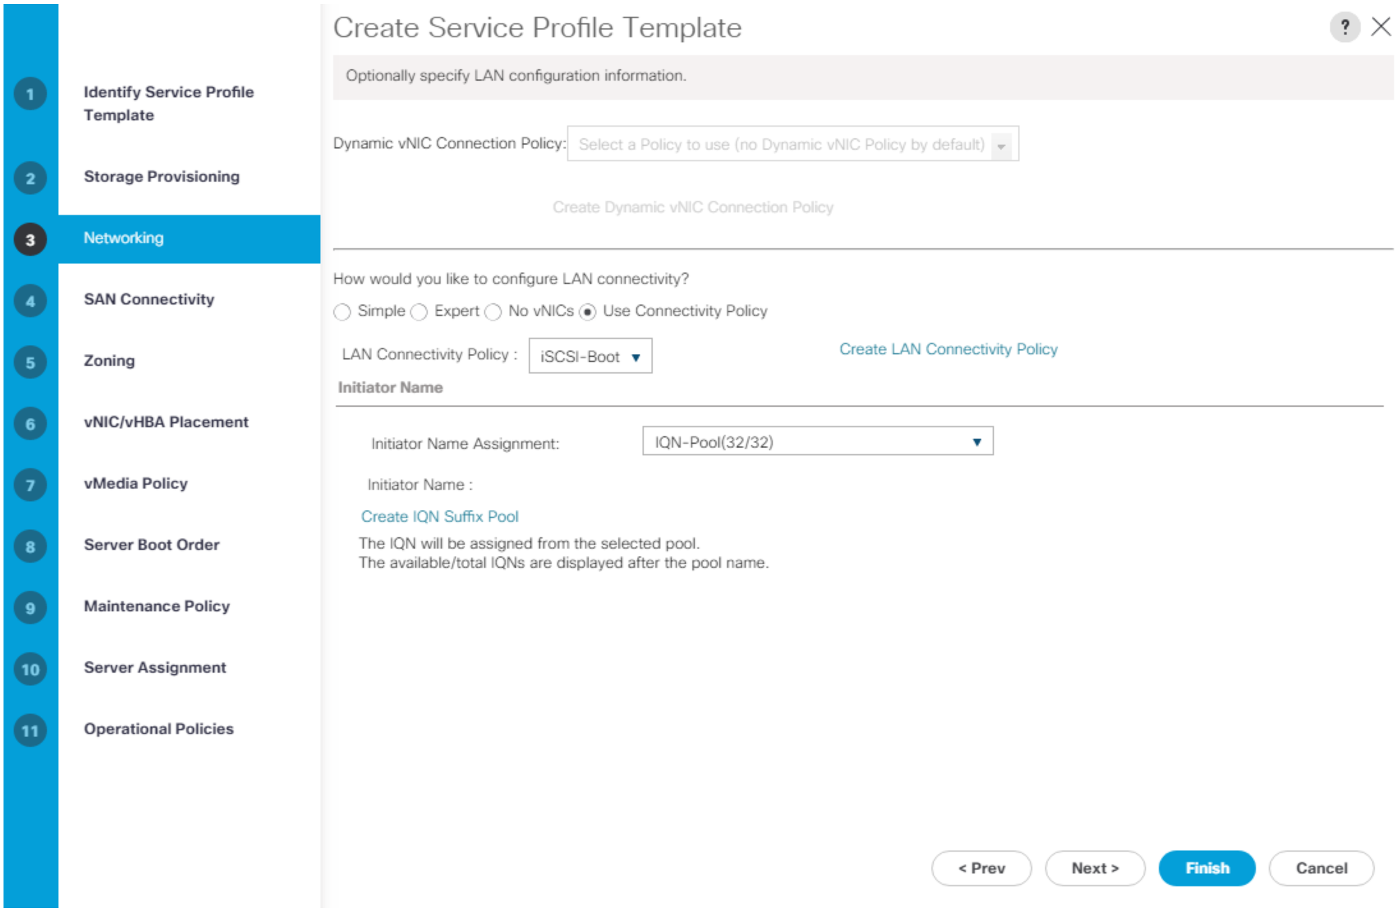

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default)

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

○ Simple ○ Expert ○ No vNICs ● Use Connectivity Policy

LAN Connectivity Policy : iSCSI-Boot

Create LAN Connectivity Policy

Initiator Name

Initiator Name Assignment: IQN-Pool(32/32)

Initiator Name :

Create IQN Suffix Pool

The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name.

5. Click Next.

Configure Storage Options

To configure the storage options, follow these steps:

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

2. Click Next.

Configure Zoning Options

To configure the zoning options, follow this step:

1. Make no changes and click Next.

Configure vNIC/HBA Placement

To configure the vNIC/HBA placement, follow these steps:

1. In the “Select Placement” list, leave the placement policy as “Let System Perform Placement”.

2. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

3. Do not select a vMedia Policy.

4. Click Next.

Configure Server Boot Order

To configure the server boot orders, follow these steps:

1. Select Boot-iSCSI-A for Boot Policy.

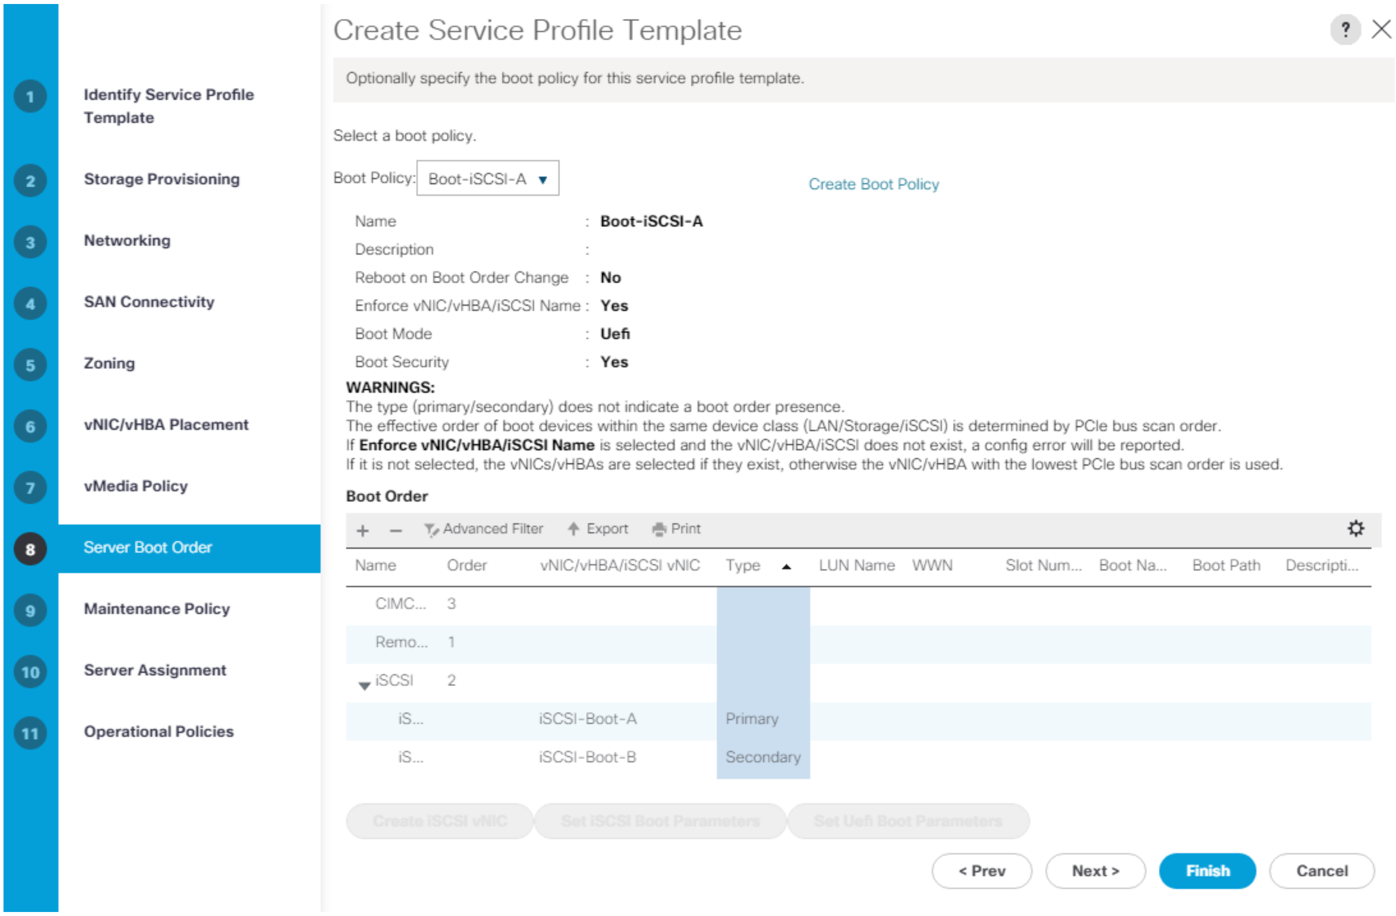

Create Service Profile Template

1 Identify Service Profile Template
2 Storage Provisioning
3 Networking
4 SAN Connectivity
5 Zoning
6 vNIC/vHBA Placement
7 vMedia Policy
8 Server Boot Order
9 Maintenance Policy
10 Server Assignment
11 Operational Policies

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Boot-iSCSI-A

Create Boot Policy

Name : **Boot-iSCSI-A**
Description :
Reboot on Boot Order Change : **No**
Enforce vNIC/vHBA/iSCSI Name : **Yes**
Boot Mode : **Uefi**
Boot Security : **Yes**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN	Slot Num...	Boot Na...	Boot Path	Descripti...
CIMC...	3								
Remo...	1								
iSCSI	2								
iS...		iSCSI-Boot-A	Primary						
iS...		iSCSI-Boot-B	Secondary						

2. In the Boor order, select iSCSI-Boot-A.

3. Click Set iSCSI Boot Parameters.

4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.

5. Leave the "Initiator Name Assignment" dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.

6. Set iSCSI-IP-Pool-A as the "Initiator IP address Policy."

7. Select iSCSI Static Target Interface option.

8. Click Add.

9. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, log into the storage cluster management interface and run the “iscsi show” command”.

10. Enter the IP address of iscsi-lif01a for the IPv4 Address field.

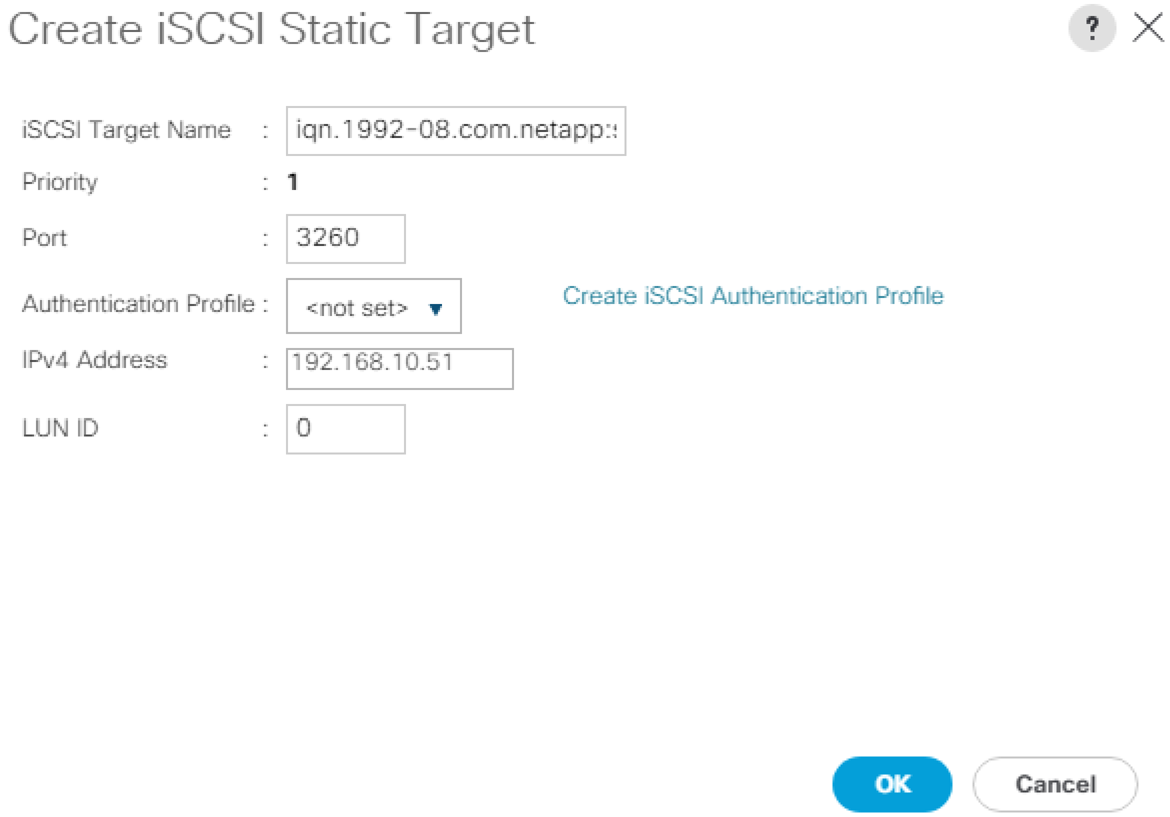

Create iSCSI Static Target

iSCSI Target Name : iqn.1992-08.com.netapp::

Priority : 1

Port : 3260

Authentication Profile : <not set>

IPv4 Address : 192.168.10.51

LUN ID : 0

11. Click OK to add the iSCSI static target.

12. Click Add.

13. Enter the iSCSI Target Name.

14. Enter the IP address of iscsi-lif02a for the IPv4 Address field.

15. Click OK to add the iSCSI static target.

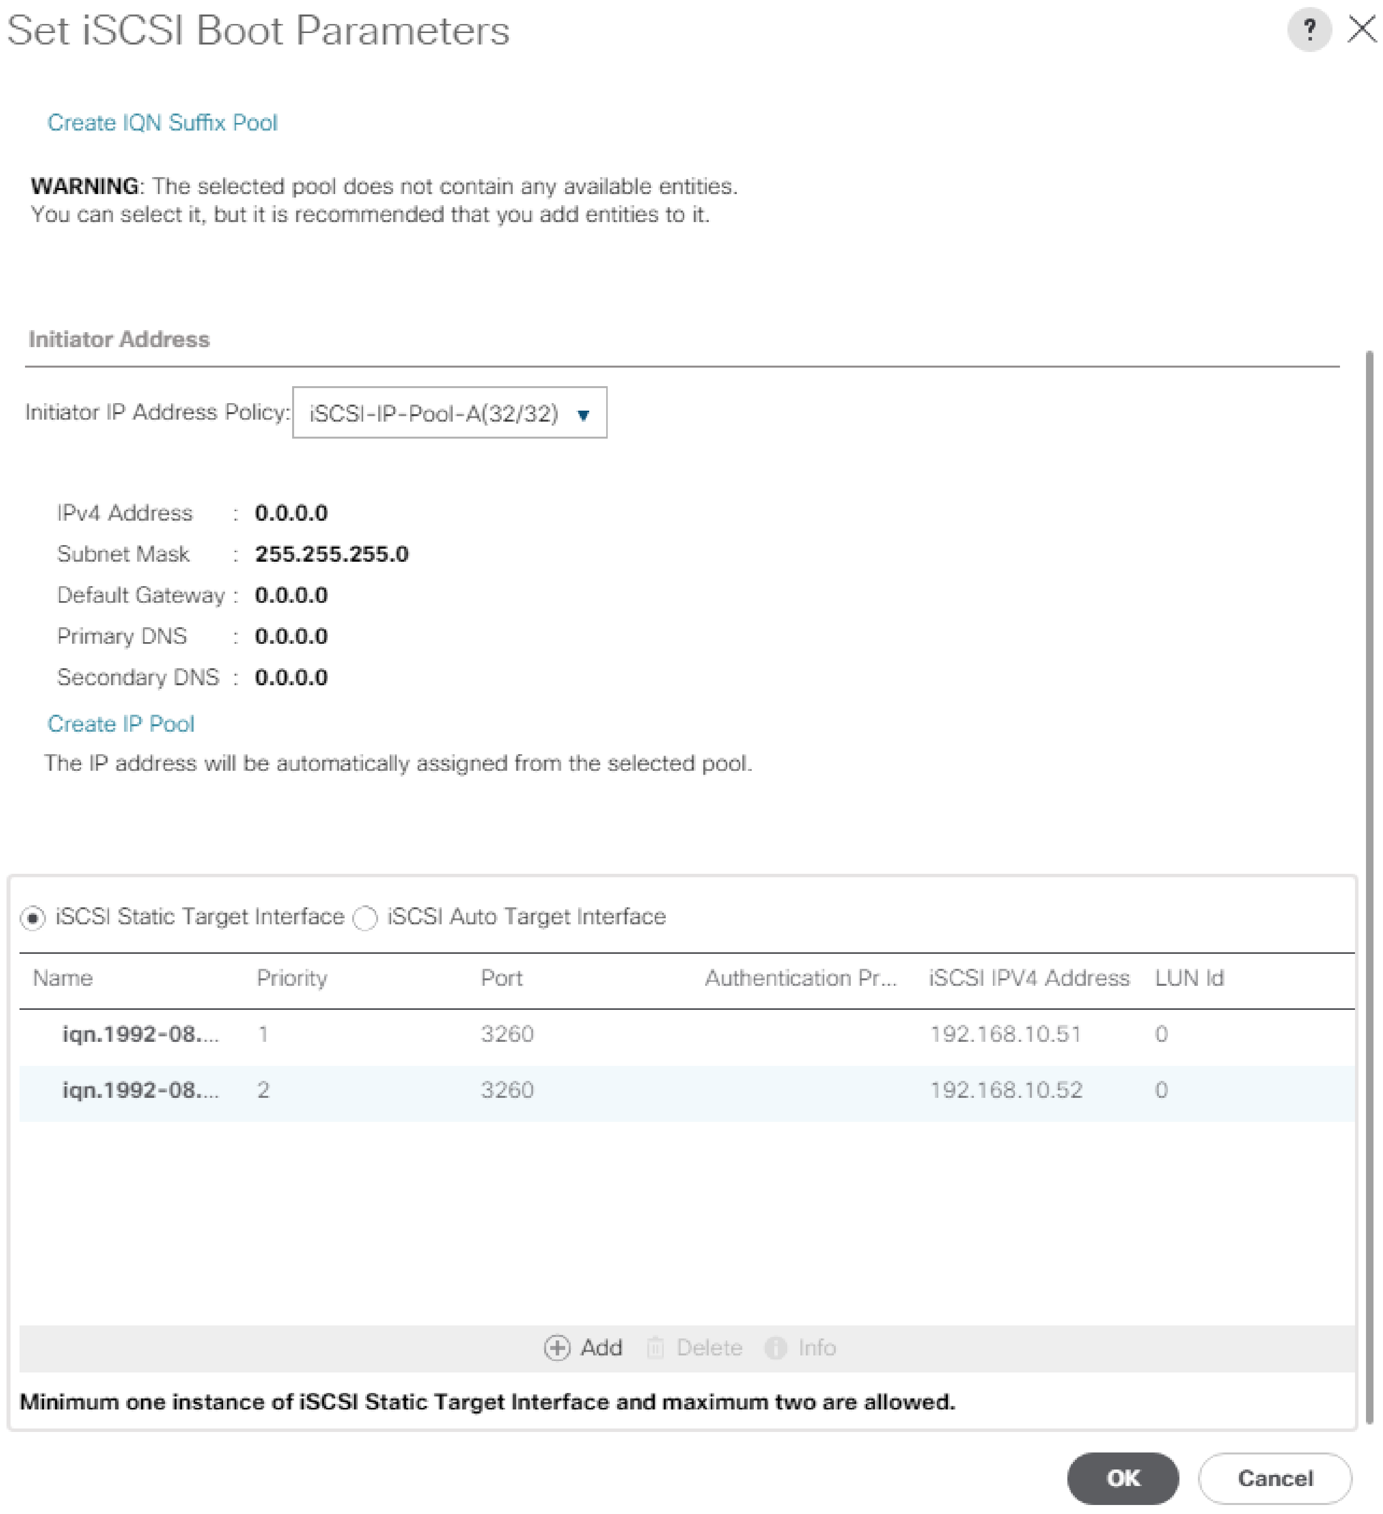

Set iSCSI Boot Parameters

Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-A(32/32) ▾

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

Create IP Pool

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

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.10.51	0
iqn.1992-08....	2	3260		192.168.10.52	0

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

11. Click OK to complete setting the iSCSI Boot Parameters.

12. In the Boot order, select iSCSI-Boot-B.

13. Click Set iSCSI Boot Parameters.

14. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.

15. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.

16. Set iSCSI-IP-Pool-B as the “Initiator IP address Policy”.

17. Select the iSCSI Static Target Interface option.

18. Click Add.

19. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, login into storage cluster management interface and run “iscsi show” command”.

20. Enter the IP address of iscsi-lif01b for the IPv4 Address field.

21. Click OK to add the iSCSI static target.

22. Click Add.

23. Enter the iSCSI Target Name.

24. Enter the IP address of iscsi-lif02b for the IPv4 Address field.

25. Click OK to add the iSCSI static target.

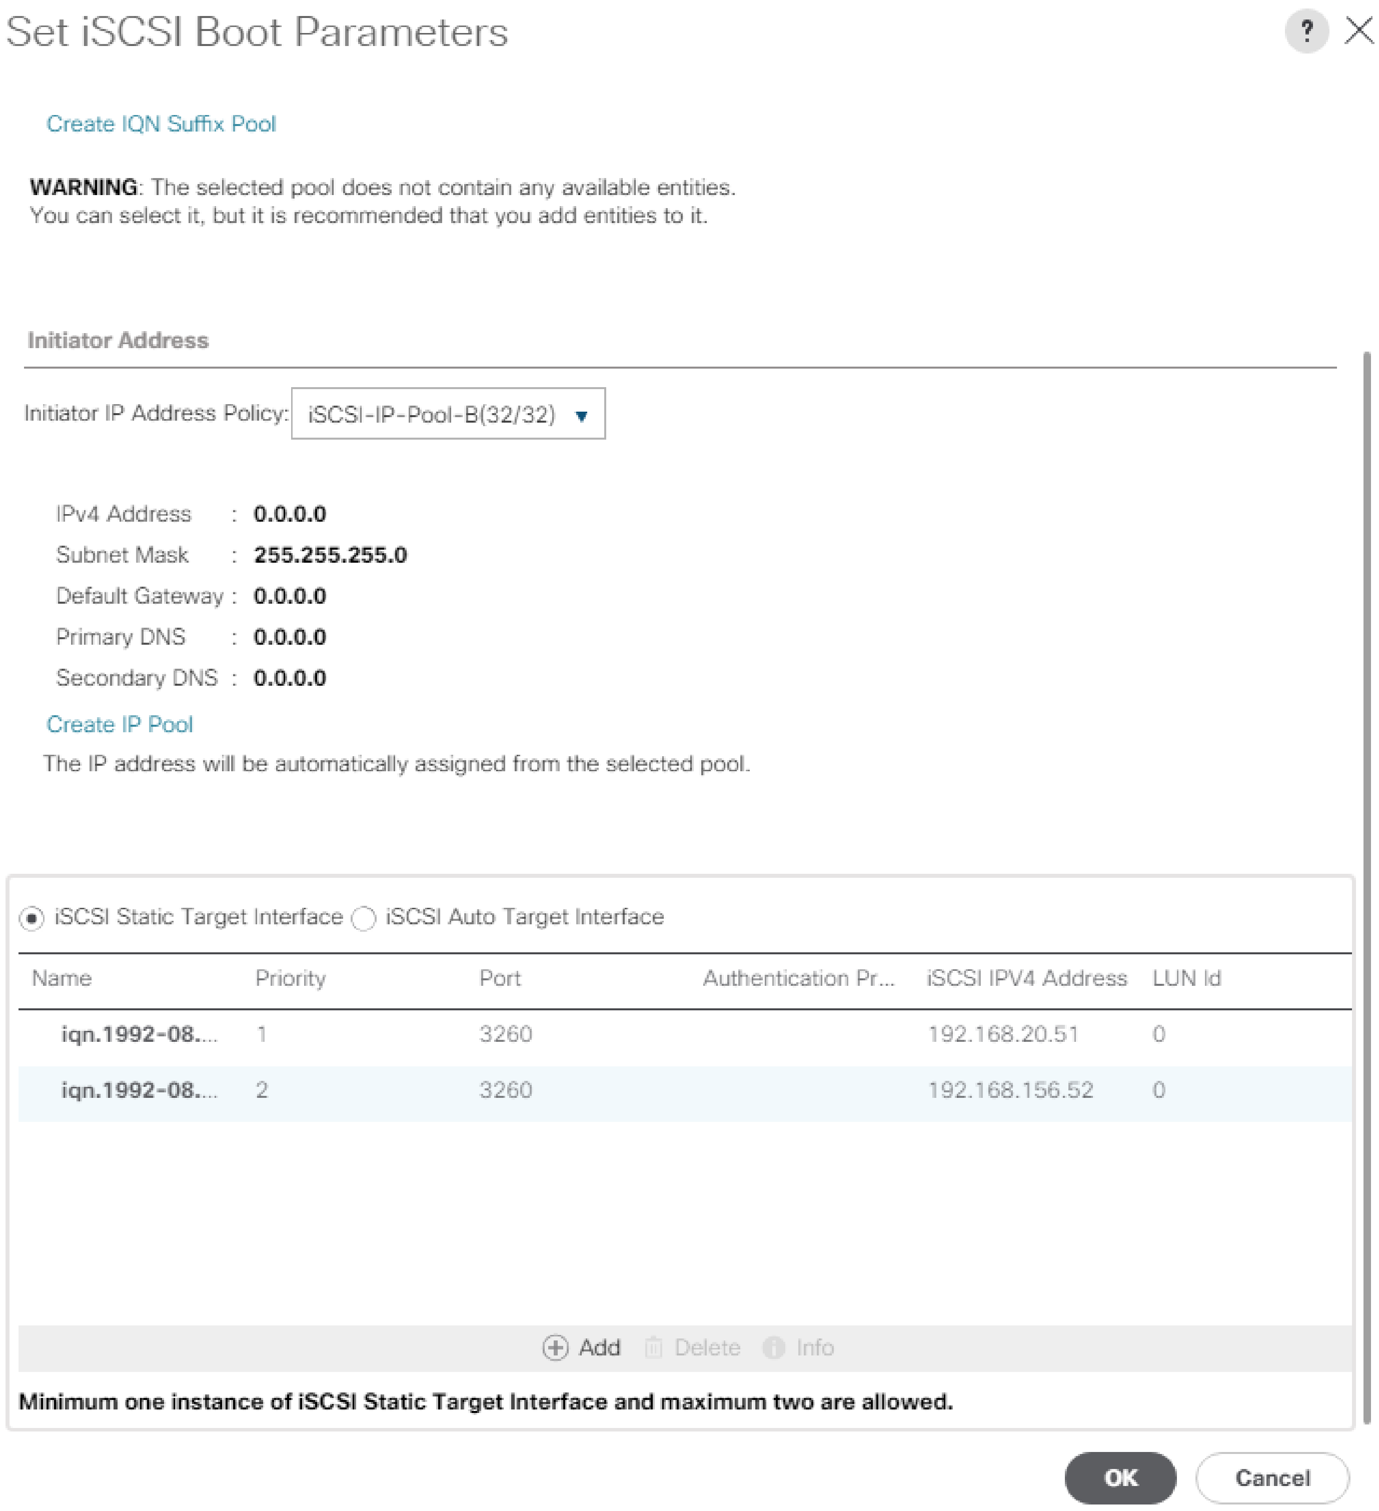

Set iSCSI Boot Parameters

Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-B(32/32) ▾

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

Create IP Pool

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

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.20.51	0
iqn.1992-08....	2	3260		192.168.156.52	0

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

26. Click OK to complete setting the iSCSI Boot Parameters.

27. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.

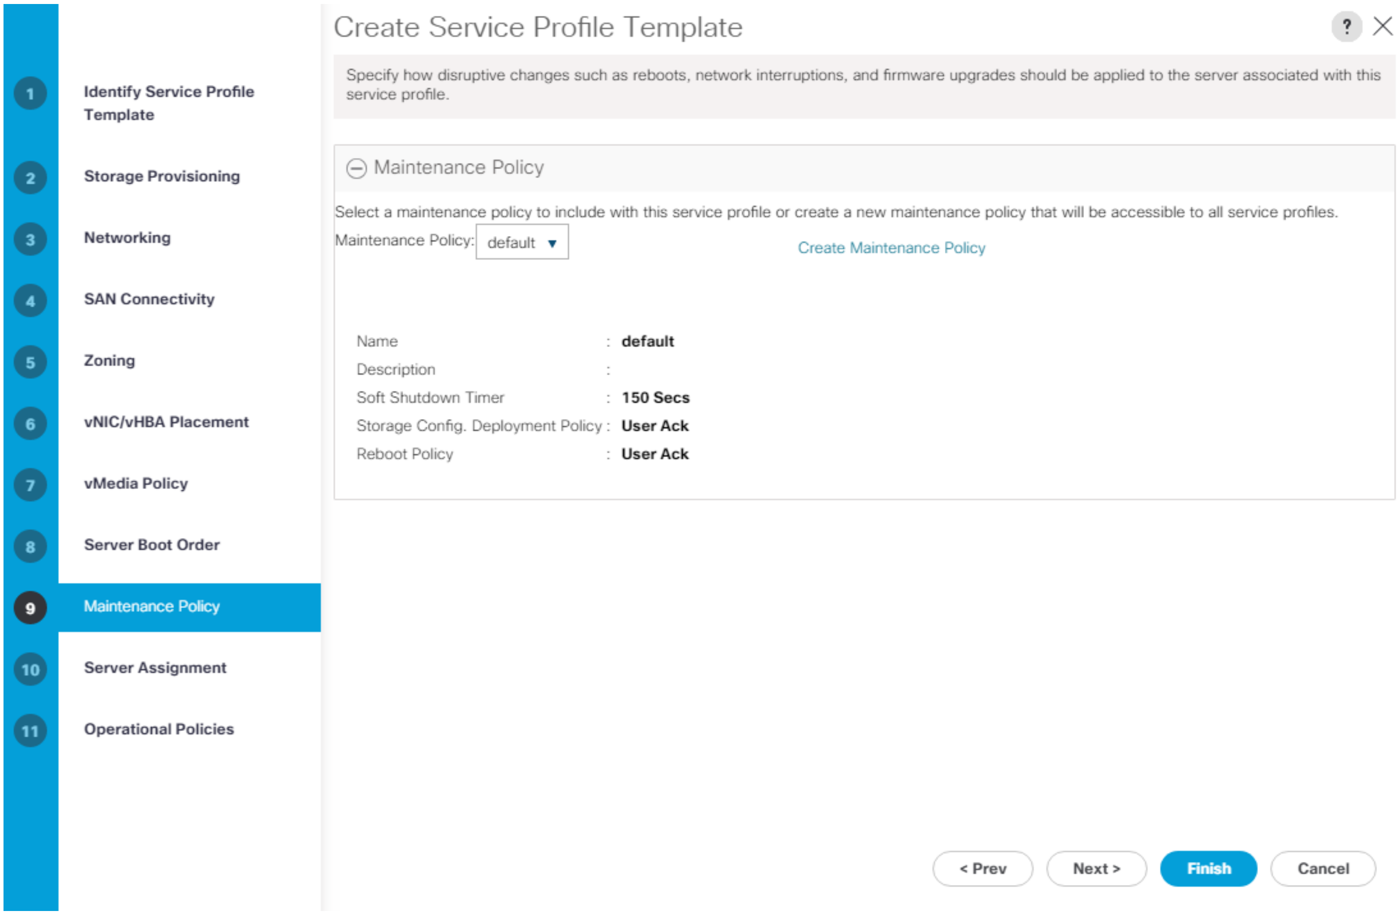

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: default

Name	default
Description	
Soft Shutdown Timer	150 Secs
Storage Config. Deployment Policy	User Ack
Reboot Policy	User Ack

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, select Infra-Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select “UCS-B200-M5” for the Server Pool Qualification to select only B200 M5 servers in the pool.
4. Expand Firmware Management at the bottom of the page and select the default policy.

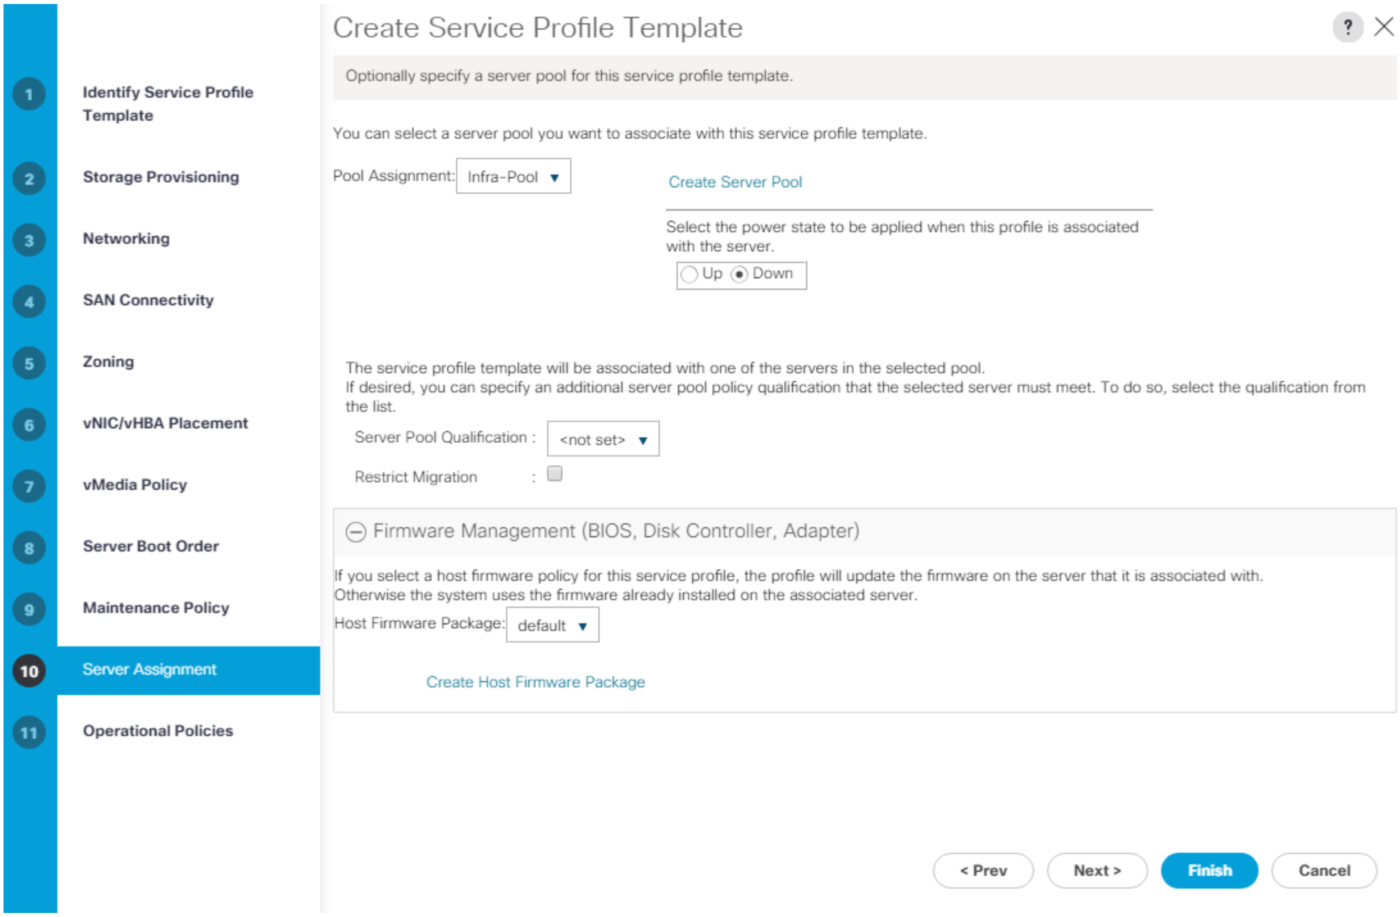

Create Service Profile Template

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: Infra-Pool ▾

Select the power state to be applied when this profile is associated with the server.

○ Up ● Down

The service profile template will be associated with one of the servers in the selected pool.
If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification : <not set> ▾

Restrict Migration : □

⊖ Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with.
Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: default ▾

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, select VM-Host.

2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

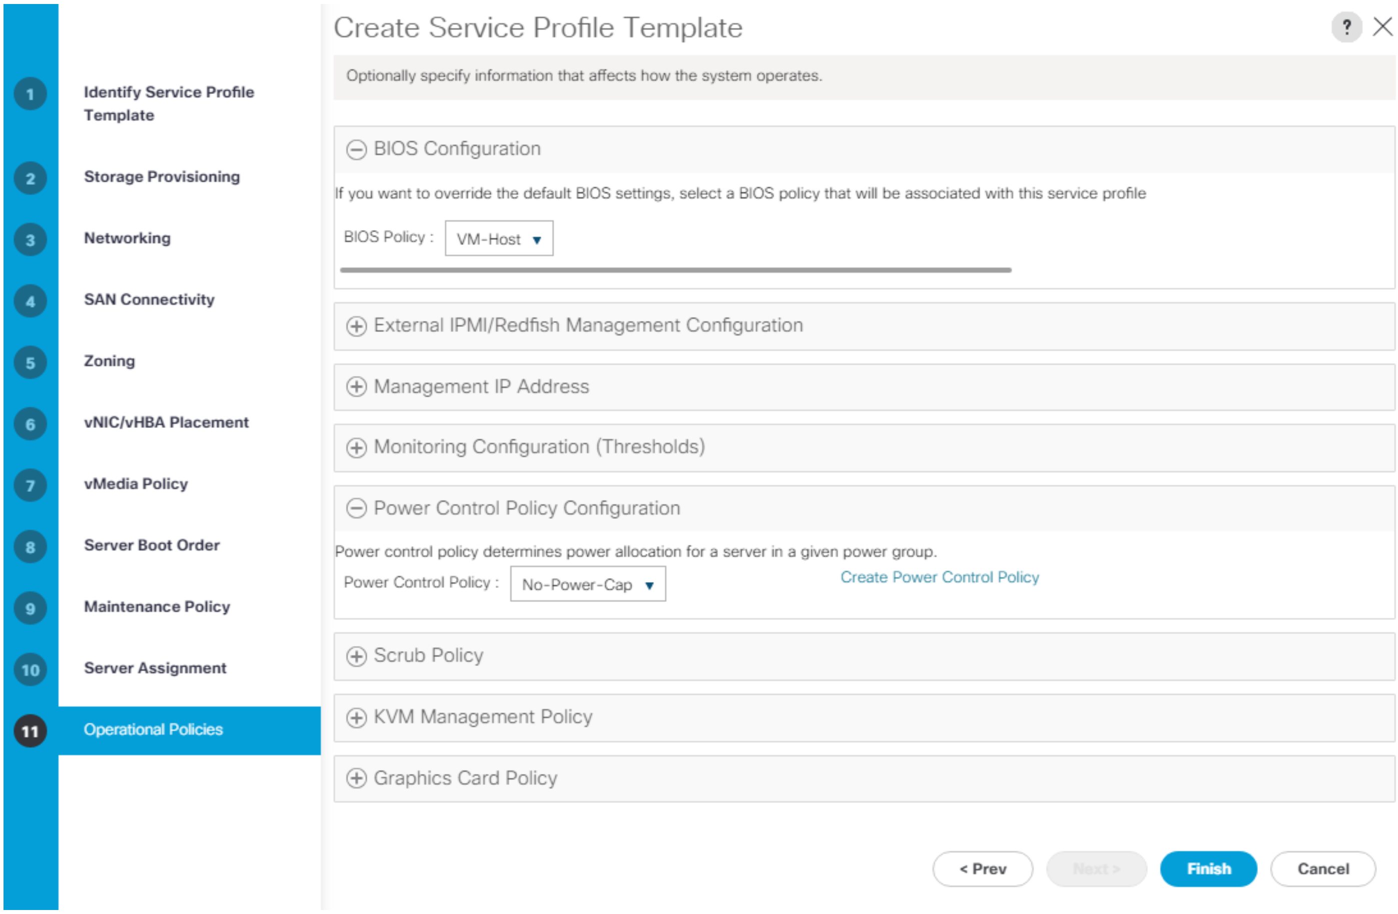

Create Service Profile Template

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy : VM-Host

Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : No-Power-Cap

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create vMedia-Enabled Service Profile Template

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.

2. Select Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template VM-Host-Infra-iSCSI-A.

3. Right-click VM-Host-Infra-iSCSI-A and select Create a Clone.

4. Name the clone VM-Host-Infra-iSCSI-A-vM and click OK then click OK again to create the clone.

5. Select the newly-created VM-Host-Infra-iSCSI-A-vM and select the vMedia Policy tab.

6. Click Modify vMedia Policy.

7. Select the ESXi-6.7U1-HTTP vMedia Policy and click OK.

8. Click OK to confirm.

Create Service Profiles

To create service profiles from the service profile template, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template VM-Host-Infra-iSCSI-A-vM.
3. Right-click VM-Host-Infra-iSCSI-A-vM and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Enter 1 as "Name Suffix Starting Number."
6. Enter 2 as the "Number of Instances."

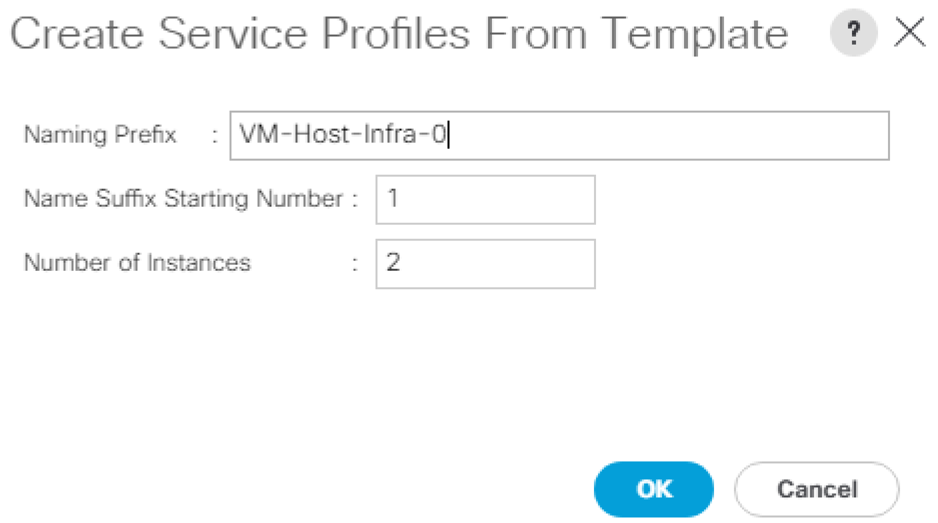

Create Service Profiles From Template

Naming Prefix : VM-Host-Infra-0

Name Suffix Starting Number : 1

Number of Instances : 2

7. Click OK to create the service profiles.
8. Click OK in the confirmation message.

Once VMware ESXi 6.7 U1 has been installed on the hosts, the host Service Profiles can be bound to the VM-Host-Infra-iSCSI-A Service Profile Template to remove the vMedia Mapping from the host.

VMware vSphere Configuration

Set Up VMkernel Ports and Virtual Switch on ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To add the iSCSI networking configuration on the ESXi hosts, follow the steps at the end of section Set Up VMkernel Ports and Virtual Switch:

1. From the Host Client, select Networking.

2. In the center pane, select the Virtual switches tab.
3. Highlight the iScsiBootvSwitch line.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Click Save.
7. Select Add standard virtual switch.
8. Name the virtual switch iScsiBootvSwitch.
9. Change the MTU to 9000.
10. Select vmnic5 for Uplink 1.

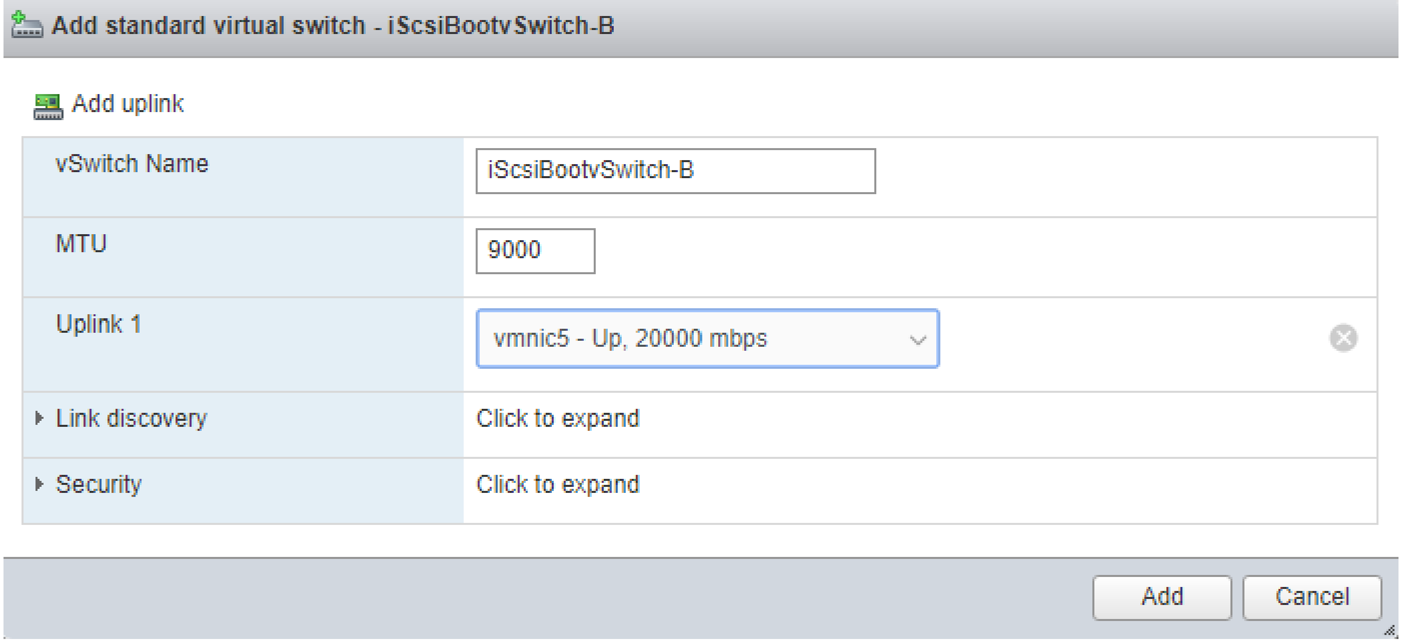

Add standard virtual switch - iScsiBootvSwitch-B

vSwitch Name	iScsiBootvSwitch-B
MTU	9000
Uplink 1	vmnic5 - Up, 20000 mbps
Link discovery	Click to expand
Security	Click to expand

11. Click Add to complete adding the virtual switch.
12. At the top, select the VMkernel NICs tab.
13. Select the iScsiBootPG line and select Edit Settings.
14. Change the MTU to 9000.
15. Expand IPv4 Settings and enter a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.

It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

16. Click Save to save the changes to the VMkernel port.

17. Click Add VMkernel NIC.

18. For New port group, enter iScsiBootPG-B

19. For Virtual switch, select iScsiBootvSwitch-B.

20. Leave the VLAN ID set at 0.

21. Change the MTU to 9000.

22. Select Static IPv4 settings and expand IPv4 settings.

23. Enter a unique IP address in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B.

It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

24. Do not select any of the Services.

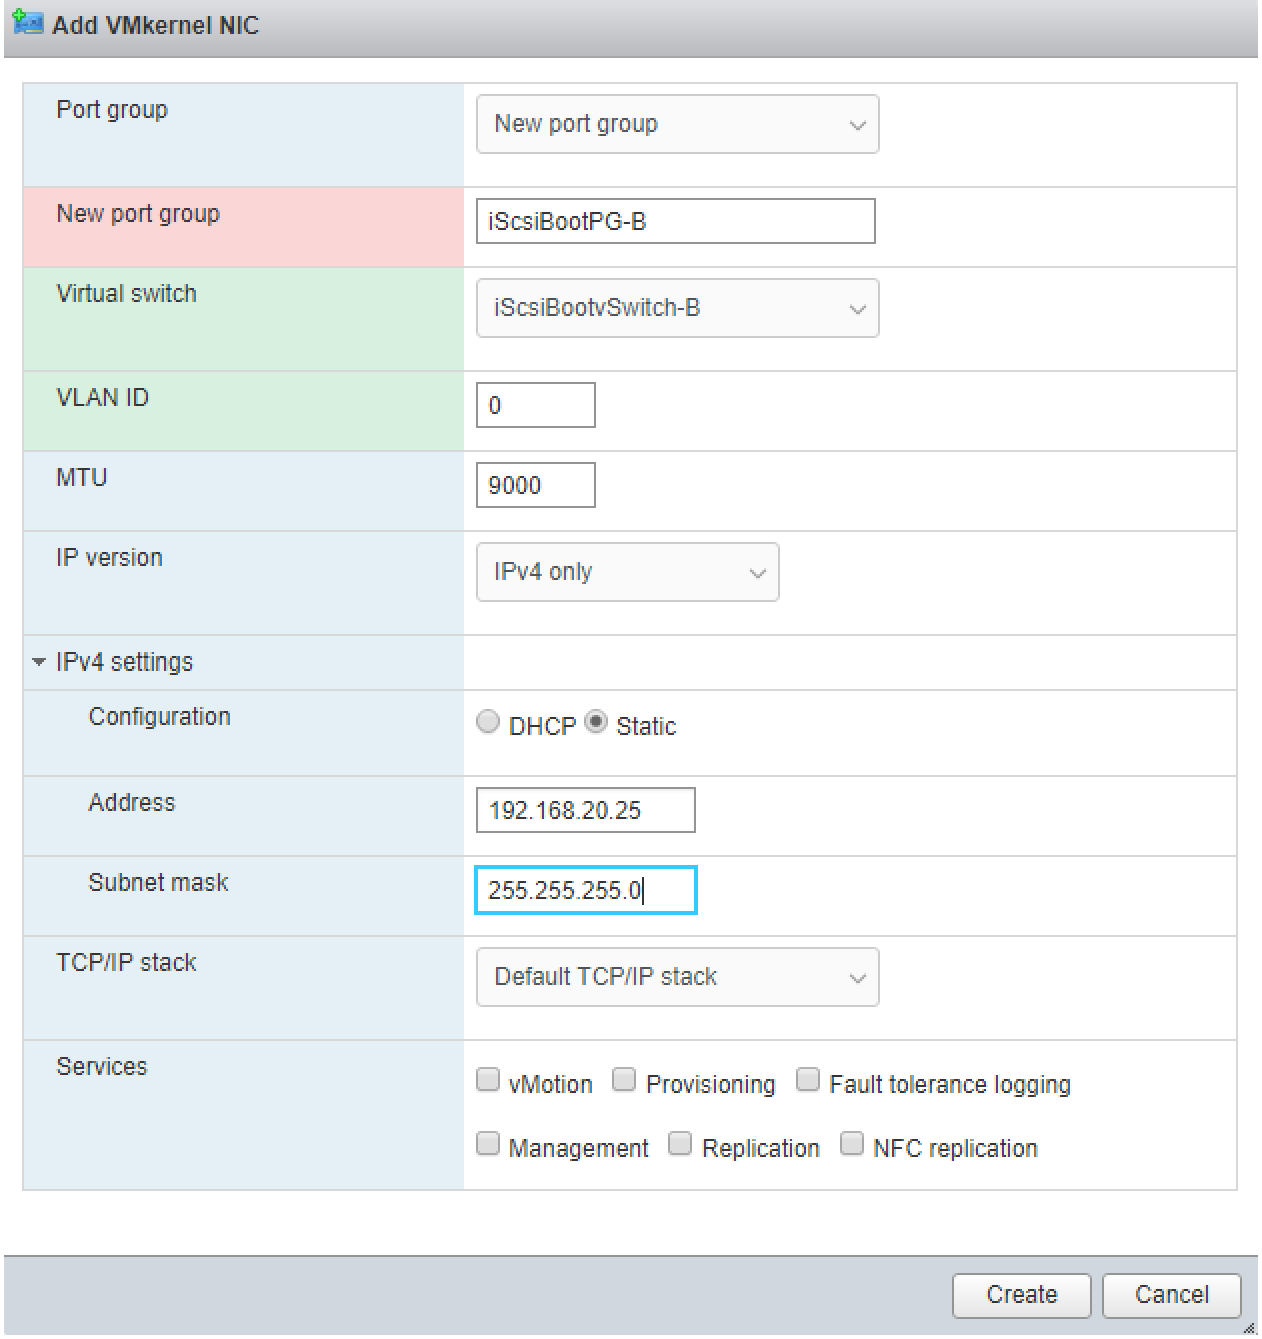

Add VMkernel NIC

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	iScsiBootvSwitch-B
VLAN ID	0
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	○ DHCP ● Static
Address	192.168.20.25
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	□ vMotion □ Provisioning □ Fault tolerance logging □ Management □ Replication □ NFC replication

25. Click Create.

26. On the left select Storage, then in the center pane select the Adapters tab.

27. Click Software iSCSI to configure software iSCSI for the host.

28. In the Configure iSCSI window, under Dynamic targets, select Add dynamic target.

29. Select to add address and enter the IP address of iscsi-lif01a from storage SVM Infra-SVM. Click Return.

30. Repeat this process to add the IP addresses for iscsi-lif02a, iscsi-lif01b, and iscsi-lif01b.

31. Click Save configuration.

32. Click Software iSCSI to configure software iSCSI for the host.

33. Verify that four static targets and four dynamic targets are listed for the host.

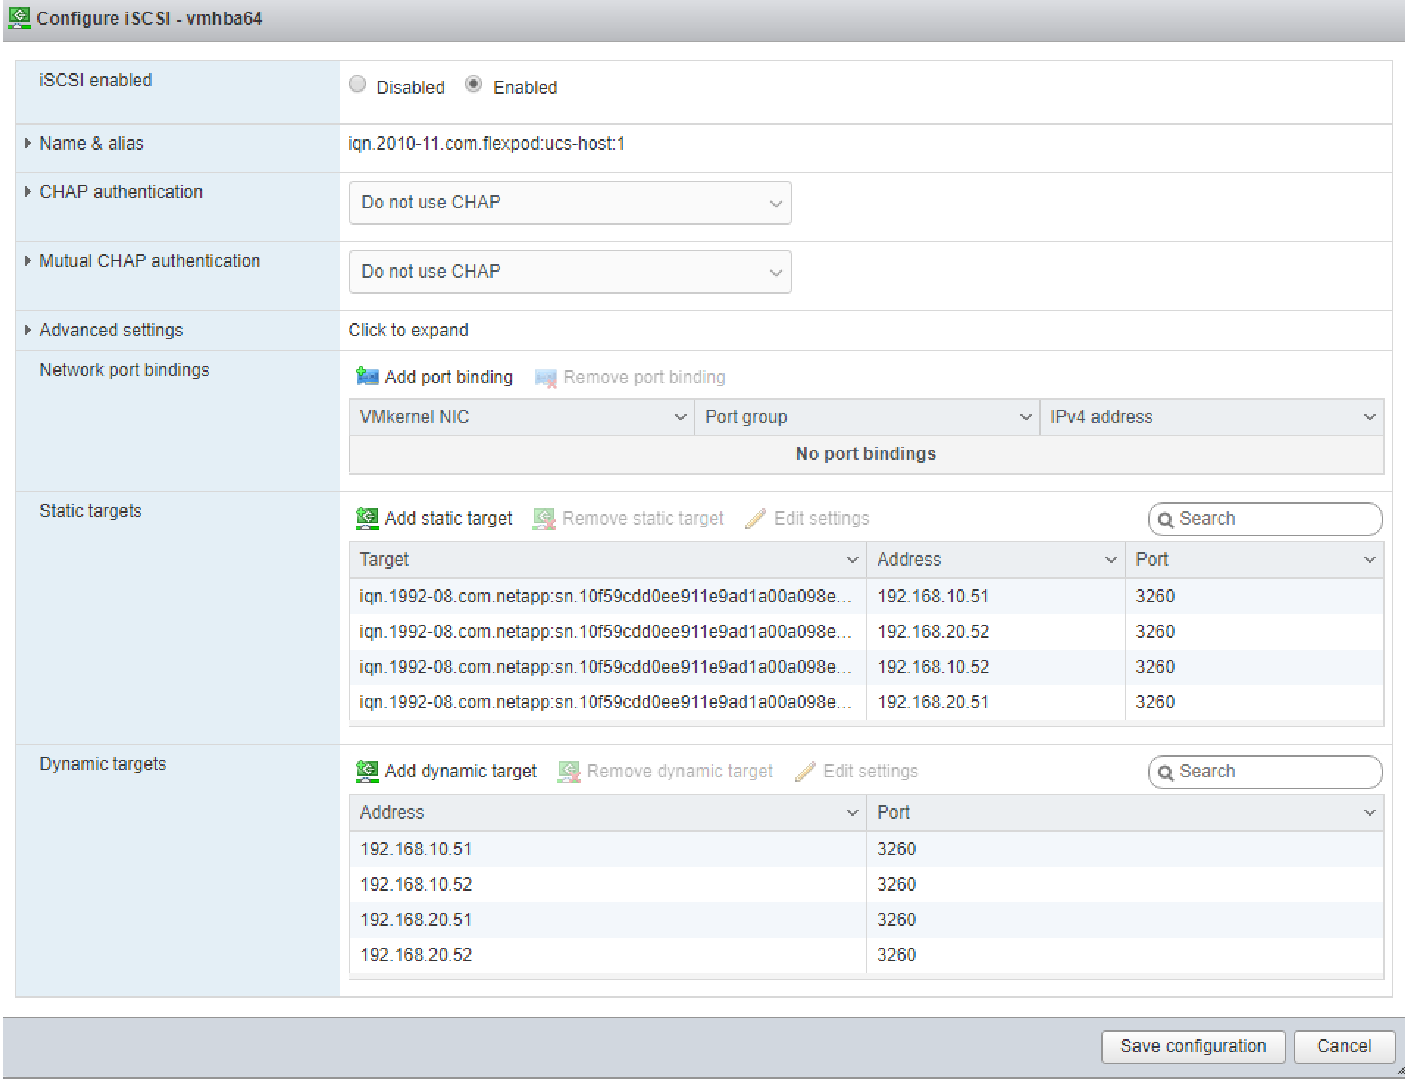

Configure iSCSI - vmhba64

iSCSI enabled: ○ Disabled ● Enabled

Name & alias: iqn.2010-11.com.flexpod:ucs-host:1

CHAP authentication: Do not use CHAP

Mutual CHAP authentication: Do not use CHAP

Advanced settings: Click to expand

Network port bindings:

VMkernel NIC	Port group	IPv4 address
No port bindings		

Static targets:

Target	Address	Port
iqn.1992-08.com.netapp:sn.10f59cdd0ee911e9ad1a00a098e...	192.168.10.51	3260
iqn.1992-08.com.netapp:sn.10f59cdd0ee911e9ad1a00a098e...	192.168.20.52	3260
iqn.1992-08.com.netapp:sn.10f59cdd0ee911e9ad1a00a098e...	192.168.10.52	3260
iqn.1992-08.com.netapp:sn.10f59cdd0ee911e9ad1a00a098e...	192.168.20.51	3260

Dynamic targets:

Address	Port
192.168.10.51	3260
192.168.10.52	3260
192.168.20.51	3260
192.168.20.52	3260

34. Click Cancel to close the window.

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. This procedure should be run at the end of the vCenter deployment section. To setup the ESXi Dump Collector, follow these steps:

1. Log into the vSphere Web Client as administrator@vsphere.local and select Home.

2. In the center pane, click System Configuration.
3. In the left pane, select Services.
4. Under services, click VMware vSphere ESXi Dump Collector.
5. In the center pane, click the green start icon ▶ to start the service.
6. In the Actions menu, click Edit Startup Type.
7. Select Automatic.
8. Click OK.
9. Connect to each ESXi host via ssh as root
10. Run the following commands:

```
esxcli system coredump network set -v vmk0 -j <vcenter-ip>
esxcli system coredump network set -e true
esxcli system coredump network check
```

Create a FlexPod ESXi Custom ISO using VMware vCenter

In this validation document, the Cisco Custom Image for ESXi 6.7 U1 GA Install CD ISO was used to install VMware ESXi. After this installation the Cisco VIC nfnic and nenic drivers had to be updated and the NetApp NFS Plug-in for VMware VAAI had to be installed during the FlexPod deployment. vCenter 6.7 U1 can be used to produce a FlexPod custom ISO containing the updated VIC drivers and the NetApp NFS Plug-in for VMware VAAI. This ISO can be used to install VMware ESXi 6.7 U1 without having to do any addition driver updates. This ISO can be produced by following these steps:

1. Download the Cisco Custom Image for ESXi 6.7 U1 Offline Bundle. This file (VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1_Bundle.zip) can be used to produce the Cisco Custom Image for ESXi 6.7 U1 GA Install CD ISO.
2. Download the following driver images and extract the listed .zip file:
   - VMware ESXi 6.7 nfnic 4.0.0.24 FC Driver for Cisco nfnic - VMW-ESX-6.7.0-nfnic-4.0.0.24-offline_bundle-12271979
   - VMware ESXi 6.7 nenic 1.0.27.0 NIC Driver for Cisco nenic - VMW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-11271332
   - NetApp NFS Plug-in for VMware VAAI 1.1.2 – NetAppNasPlugin.v23.zip
3. Connect to the VMware vCenter Web Client and login as an Administrator@vsphere.local.
4. From the Home view, select System Configuration.
5. In the Navigator pane, select Services. Under Services, select ImageBuilder Service.
6. In the center pane, to the right of ImageBuilder Service, select Actions > Edit Startup Type.

7. Select Automatic to start the service with vCenter starts and click OK.

8. Select Actions > Start to start the ImageBuilder Service.

9. Log out and log back into the VMware vCenter Web Client.

10. From the Home view, select Auto Deploy.

11. Select the Software Depots tab.

12. Click the ⬆ icon to upload a software depot.

13. Name the depot ESXi-6.7U1-Cisco-Custom. Click Browse and browse to the VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1_Bundle.zip file. Click Open.

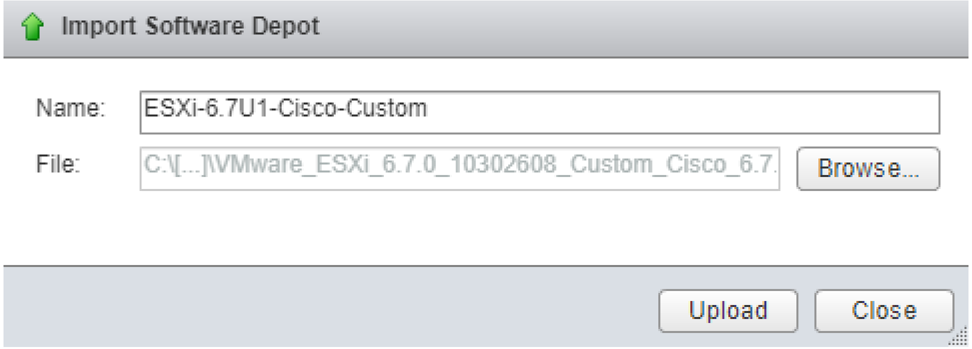

Import Software Depot

Name: ESXi-6.7U1-Cisco-Custom

File: C:\[...]\VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7

14. Click Upload to upload the software depot. Click Close when the upload is complete.

15. Repeat steps 12-14 to add software depots for nfnic-4.0.0.24, nenic-1.0.27.0, and NetAppNasPlugin-v23.

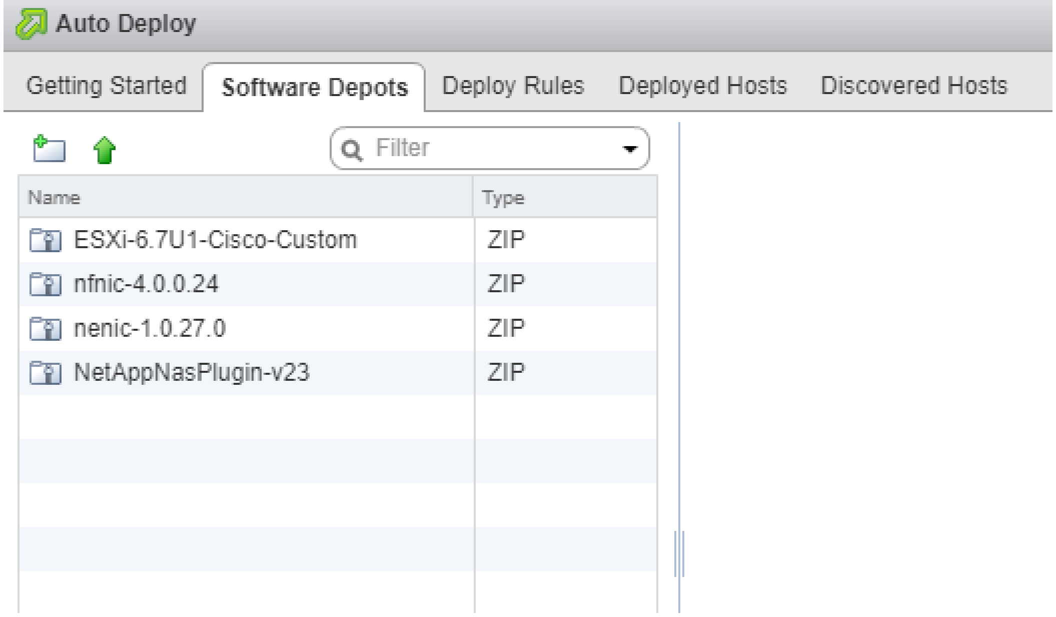

Name	Type
ESXi-6.7U1-Cisco-Custom	ZIP
nfnic-4.0.0.24	ZIP
nenic-1.0.27.0	ZIP
NetAppNasPlugin-v23	ZIP

16. Click 📁 to add a custom software depot.

17. Select Custom depot and name the custom depot FlexPod-ESXi-6.7U1.

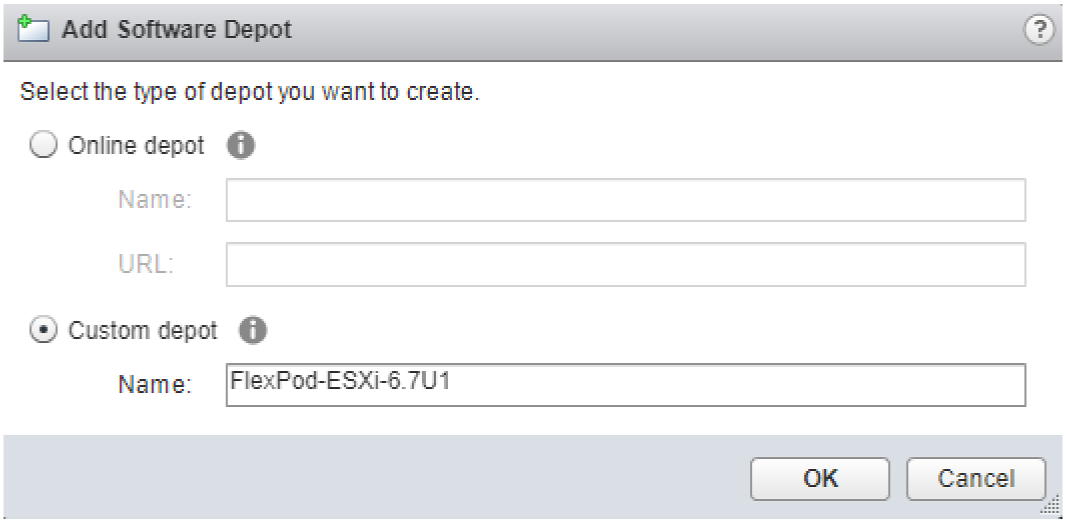

Add Software Depot

Select the type of depot you want to create.

○ Online depot

Name:

URL:

● Custom depot

Name: FlexPod-ESXi-6.7U1

18. Click OK to add the custom software depot.

19. Select the ESXi-6.7U1-Cisco-Custom software depot. To the right, select Image Profiles and then select the VMware-ESXi-6.7.0-10302608-Custom-Cisco-6.7.1.1 image profile. Click to clone the image profile.

20. Name the clone FlexPod-ESXi-6.7U1. For Vendor, enter Cisco-NetApp. For Description, enter “Cisco Custom ISO for ESXi 6.7U1 GA with nfnic-4.0.0.24, nenic-1.0.27.0, and NetAppNasPlugin-v23”. Select FlexPod-ESXi-6.7U1 for Software depot.

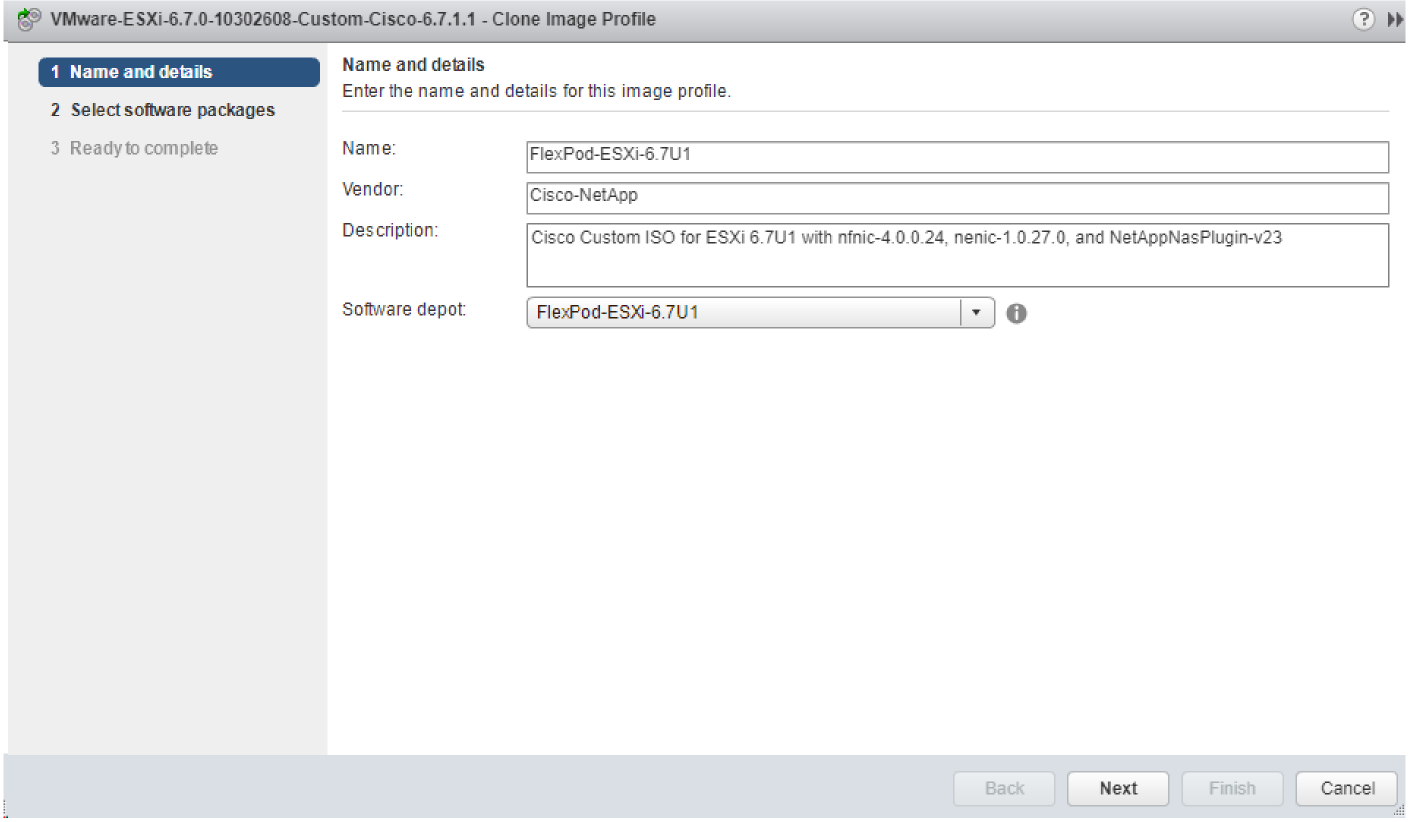

VMware-ESXi-6.7.0-10302608-Custom-Cisco-6.7.1.1 - Clone Image Profile

Name and details

Enter the name and details for this image profile.

Name:	FlexPod-ESXi-6.7U1
Vendor:	Cisco-NetApp
Description:	Cisco Custom ISO for ESXi 6.7U1 with nfnic-4.0.0.24, nenic-1.0.27.0, and NetAppNasPlugin-v23
Software depot:	FlexPod-ESXi-6.7U1

21. Click Next.

22. Under Available software packages, check NetAppNasPlugin 1.1.2-3, uncheck nenic 1.0.25.0, check nenic 1.0.27.0, uncheck nfnic 4.0.0.14, and check nfnic 4.0.0.24. Leave the remaining selections unchanged.

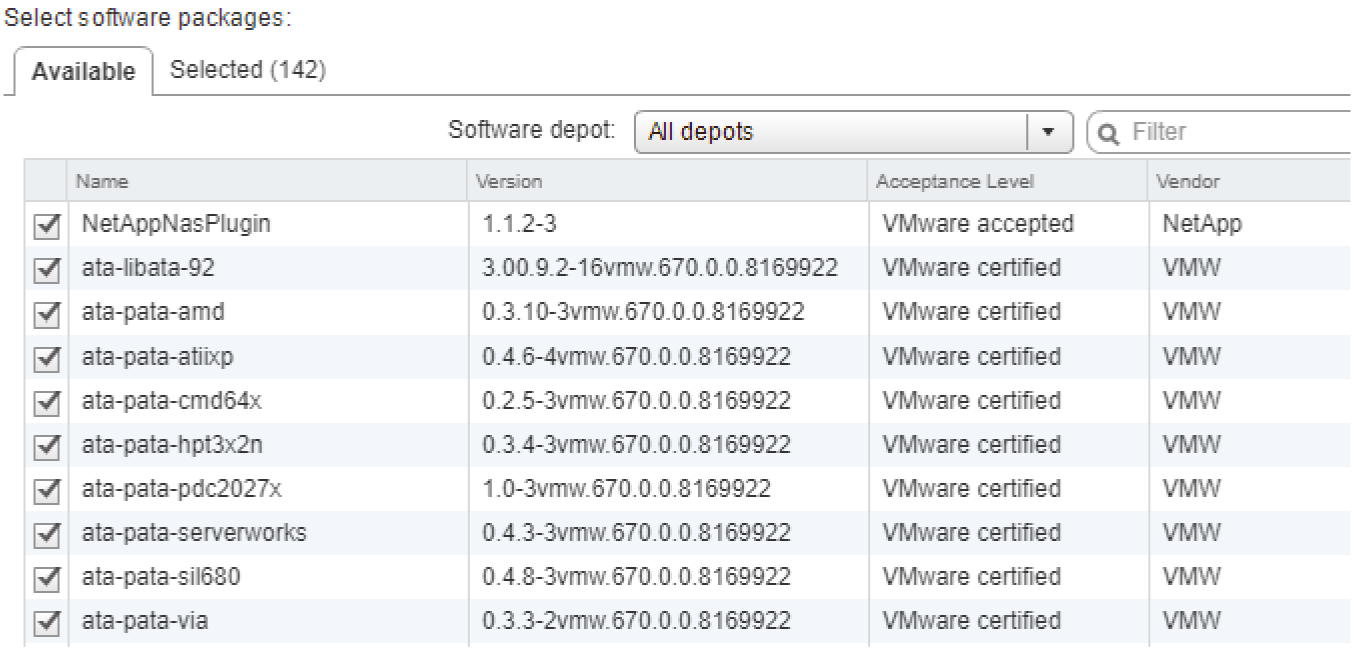

Select software packages:

Software depot: All depots

	Name	Version	Acceptance Level	Vendor
☑	NetAppNasPlugin	1.1.2-3	VMware accepted	NetApp
☑	ata-libata-92	3.00.9.2-16vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-amd	0.3.10-3vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-atiixp	0.4.6-4vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-cmd64x	0.2.5-3vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-hpt3x2n	0.3.4-3vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-pdc2027x	1.0-3vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-serverworks	0.4.3-3vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-sil680	0.4.8-3vmw.670.0.0.8169922	VMware certified	VMW
☑	ata-pata-via	0.3.3-2vmw.670.0.0.8169922	VMware certified	VMW

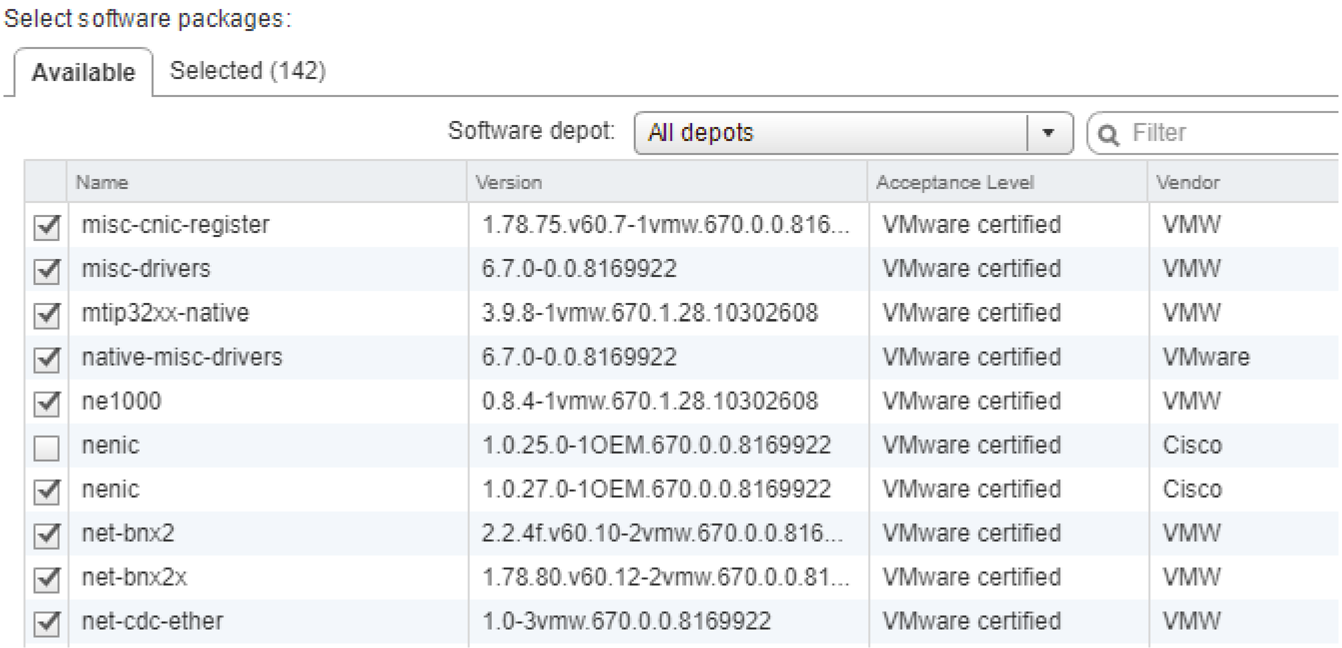

Select software packages:

Software depot: All depots

	Name	Version	Acceptance Level	Vendor
☑	misc-cnic-register	1.78.75.v60.7-1vmw.670.0.0.816...	VMware certified	VMW
☑	misc-drivers	6.7.0-0.0.8169922	VMware certified	VMW
☑	mtip32xx-native	3.9.8-1vmw.670.1.28.10302608	VMware certified	VMW
☑	native-misc-drivers	6.7.0-0.0.8169922	VMware certified	VMware
☑	ne1000	0.8.4-1vmw.670.1.28.10302608	VMware certified	VMW
□	nenic	1.0.25.0-1OEM.670.0.0.8169922	VMware certified	Cisco
☑	nenic	1.0.27.0-1OEM.670.0.0.8169922	VMware certified	Cisco
☑	net-bnx2	2.2.4f.v60.10-2vmw.670.0.0.816...	VMware certified	VMW
☑	net-bnx2x	1.78.80.v60.12-2vmw.670.0.0.81...	VMware certified	VMW
☑	net-cdc-ether	1.0-3vmw.670.0.0.8169922	VMware certified	VMW

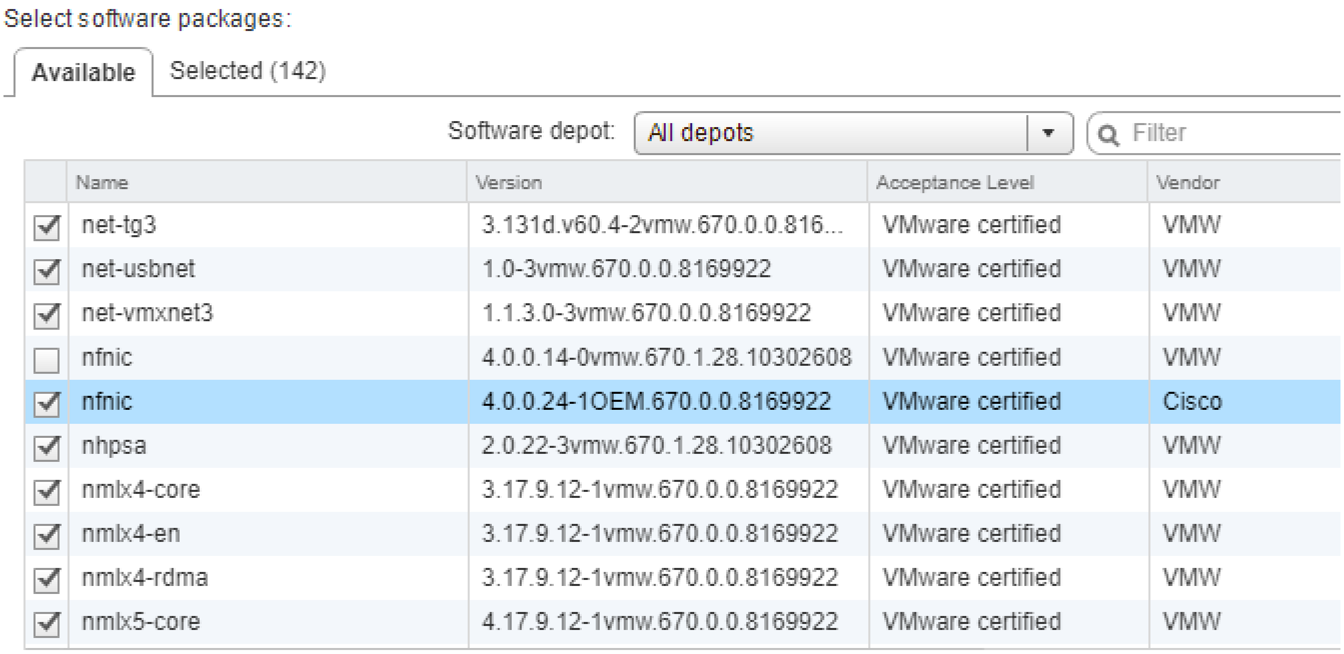

Select software packages:

Software depot: All depots

	Name	Version	Acceptance Level	Vendor
☑	net-tg3	3.131d.v60.4-2vmw.670.0.0.816...	VMware certified	VMW
☑	net-usbnet	1.0-3vmw.670.0.0.8169922	VMware certified	VMW
☑	net-vmxnet3	1.1.3.0-3vmw.670.0.0.8169922	VMware certified	VMW
□	nfnic	4.0.0.14-0vmw.670.1.28.10302608	VMware certified	VMW
☑	nfnic	4.0.0.24-1OEM.670.0.0.8169922	VMware certified	Cisco
☑	nhpsa	2.0.22-3vmw.670.1.28.10302608	VMware certified	VMW
☑	nmlx4-core	3.17.9.12-1vmw.670.0.0.8169922	VMware certified	VMW
☑	nmlx4-en	3.17.9.12-1vmw.670.0.0.8169922	VMware certified	VMW
☑	nmlx4-rdma	3.17.9.12-1vmw.670.0.0.8169922	VMware certified	VMW
☑	nmlx5-core	4.17.9.12-1vmw.670.0.0.8169922	VMware certified	VMW

23. Click Next.

24. Click Finish.

25. Select the FlexPod-ESXi-6.7U1 custom software depot. On the right, select Image Profiles and select the FlexPod-ESXi-6.7U1 image profile. Click ⬇ to export an image profile. ISO should be selected. Click Generate image to generate a bootable ESXi installable image.

26. Once “Image generated successfully” appears, click Download image to download the ISO.

27. Optionally, generate the ZIP archive to generate an offline bundle for the FlexPod image.

FlexPod Backups

Cisco UCS Backup

Automated backup of the UCS domain is important for recovery of the UCS Domain from issues ranging catastrophic failure to human error. There is a native backup solution within Cisco UCS that allows local or remote backup using FTP/TFTP/SCP/SFTP as options.

Backups created can be a binary file containing the Full State, which can be used for a restore to the original or a replacement pair of fabric interconnects. Alternately this XML configuration file consisting of All configurations, just System configurations, or just Logical configurations of the UCS Domain. For scheduled backups, options will be Full State or All Configuration, backup of just the System or Logical configurations can be manually initiated.

Specification of the backup can be done by following these steps within the Cisco UCS Manager GUI:

1. Select Admin within the Navigation pane and select All.

2. Click the Policy Backup & Export tab within All.

3. For a Full State Backup, All Configuration Backup, or both, specify the following:

   a. Hostname : <IP or FQDN of host that will receive the backup>
   b. Protocol: [FTP/TFTP/SCP/SFTP]
   c. User: <account on host to authenticate>
   d. Password: <password for account on host>
   e. Remote File: <full path and filename prefix for backup file>

Admin State must be Enabled to fill in the Remote File field.

   f. Admin State: <select Enable to activate the schedule on save, Disable to disable schedule on Save>
   g. Schedule: [Daily/Weekly/Bi Weekly]

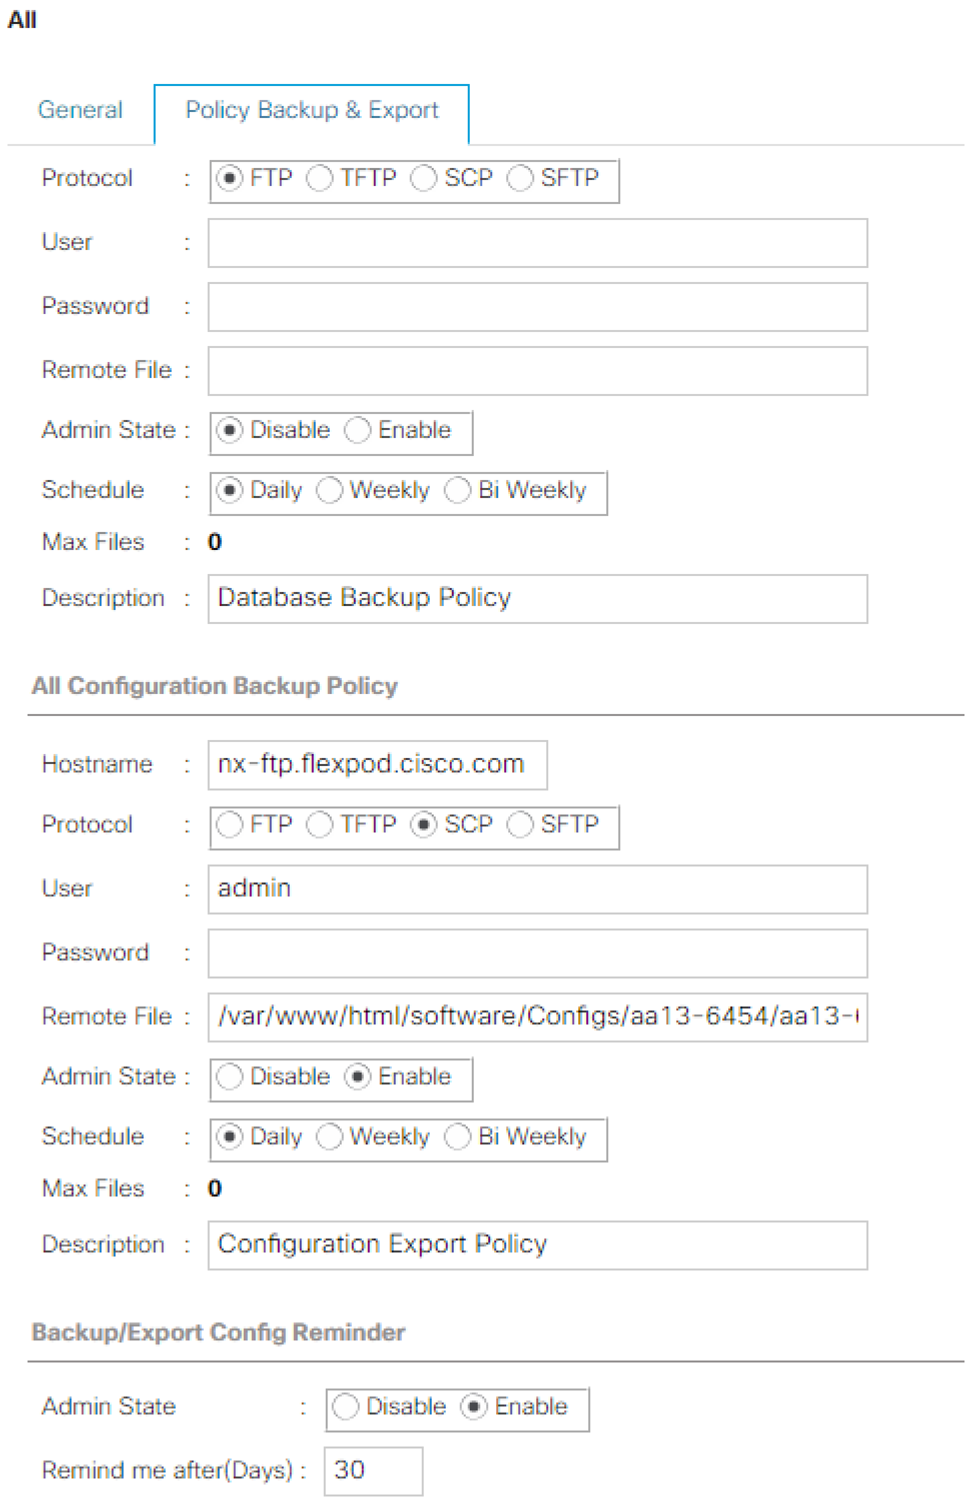

All

Protocol : ● FTP ○ TFTP ○ SCP ○ SFTP

User :

Password :

Remote File :

Admin State : ● Disable ○ Enable

Schedule : ● Daily ○ Weekly ○ Bi Weekly

Max Files : **0**

Description : Database Backup Policy

All Configuration Backup Policy

Hostname : nx-ftp.flexpod.cisco.com

Protocol : ○ FTP ○ TFTP ● SCP ○ SFTP

User : admin

Password :

Remote File : /var/www/html/software/Configs/aa13-6454/aa13-

Admin State : ○ Disable ● Enable

Schedule : ● Daily ○ Weekly ○ Bi Weekly

Max Files : **0**

Description : Configuration Export Policy

Backup/Export Config Reminder

Admin State : ○ Disable ● Enable

Remind me after(Days) : 30

4. Click Save Changes to create the Policy.

Cisco Nexus and MDS Backups

The configuration of the Cisco Nexus 9000 and Cisco MDS 9132T switches can be backed up manually at any time with the copy command, but automated backups can be put in place with the NX-OS feature scheduler. An example of setting up automated configuration backups of one of the FlexPod 9336C-FX2 switches is shown below:

```
conf t
feature scheduler
scheduler logfile size 1024
scheduler job name backup-cfg
copy running-config tftp://<server-ip>/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
exit
scheduler schedule name daily
job name backup-cfg
time daily 2:00
end
```

On the Cisco MDS 9132T, remove “vrf management” from the copy command.

Show the job that has been setup:

```
sh scheduler job
Job Name: backup-cfg
--------------------
copy running-config tftp://10.1.156.150/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management

==============================================================================

show scheduler schedule
Schedule Name       : daily
---------------------------
User Name           : admin
Schedule Type       : Run every day at 2 Hrs 0 Mins
Last Execution Time : Sun Apr  9 02:00:00 2017
Last Completion Time: Sun Apr  9 02:00:01 2017
Execution count     : 3
-----------------------------------------------
     Job Name            Last Execution Status
-----------------------------------------------
backup-cfg                        Success  (0)
==============================================================================
```

The documentation for the feature scheduler can be found here: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_01010.html

VMware VCSA Backup

Basic backup of the vCenter Server Appliance is also available within the native capabilities of the VCSA, though within the default solution this is manually initiated provide a scheduled solution. To create a backup, follow these steps:

1. Connect to the VCSA Console at https://<VCSA IP>:5480.

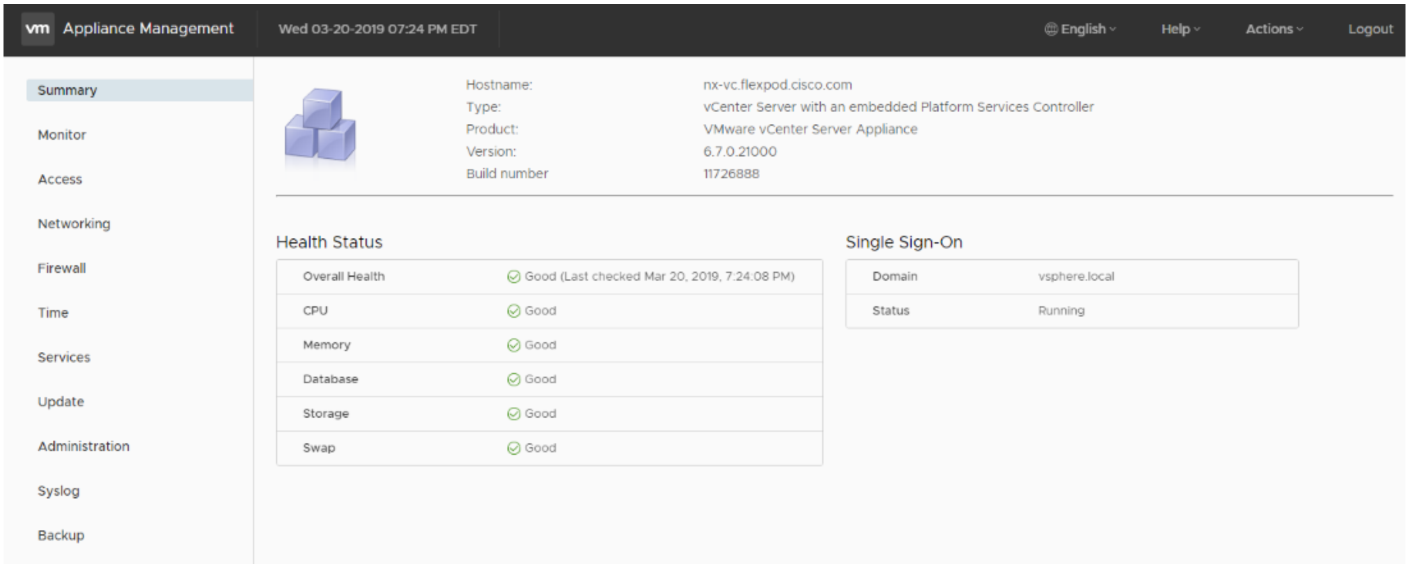

Hostname:	nx-vc.flexpod.cisco.com
Type:	vCenter Server with an embedded Platform Services Controller
Product:	VMware vCenter Server Appliance
Version:	6.7.0.21000
Build number	11726888

Health Status

Overall Health	✓ Good (Last checked Mar 20, 2019, 7:24:08 PM)
CPU	✓ Good
Memory	✓ Good
Database	✓ Good
Storage	✓ Good
Swap	✓ Good

Single Sign-On

Domain	vsphere.local
Status	Running

2. Click Backup in the list to open up the Backup Appliance Dialogue.

3. To the right of Backup Schedule, click CONFIGURE.

4. Specify:

   a. The Backup location with the protocol to use [HTTPS/HTTP/SCP/FTPS/FTP]

   b. The User name and password.

   c. The Number of backups to retain.

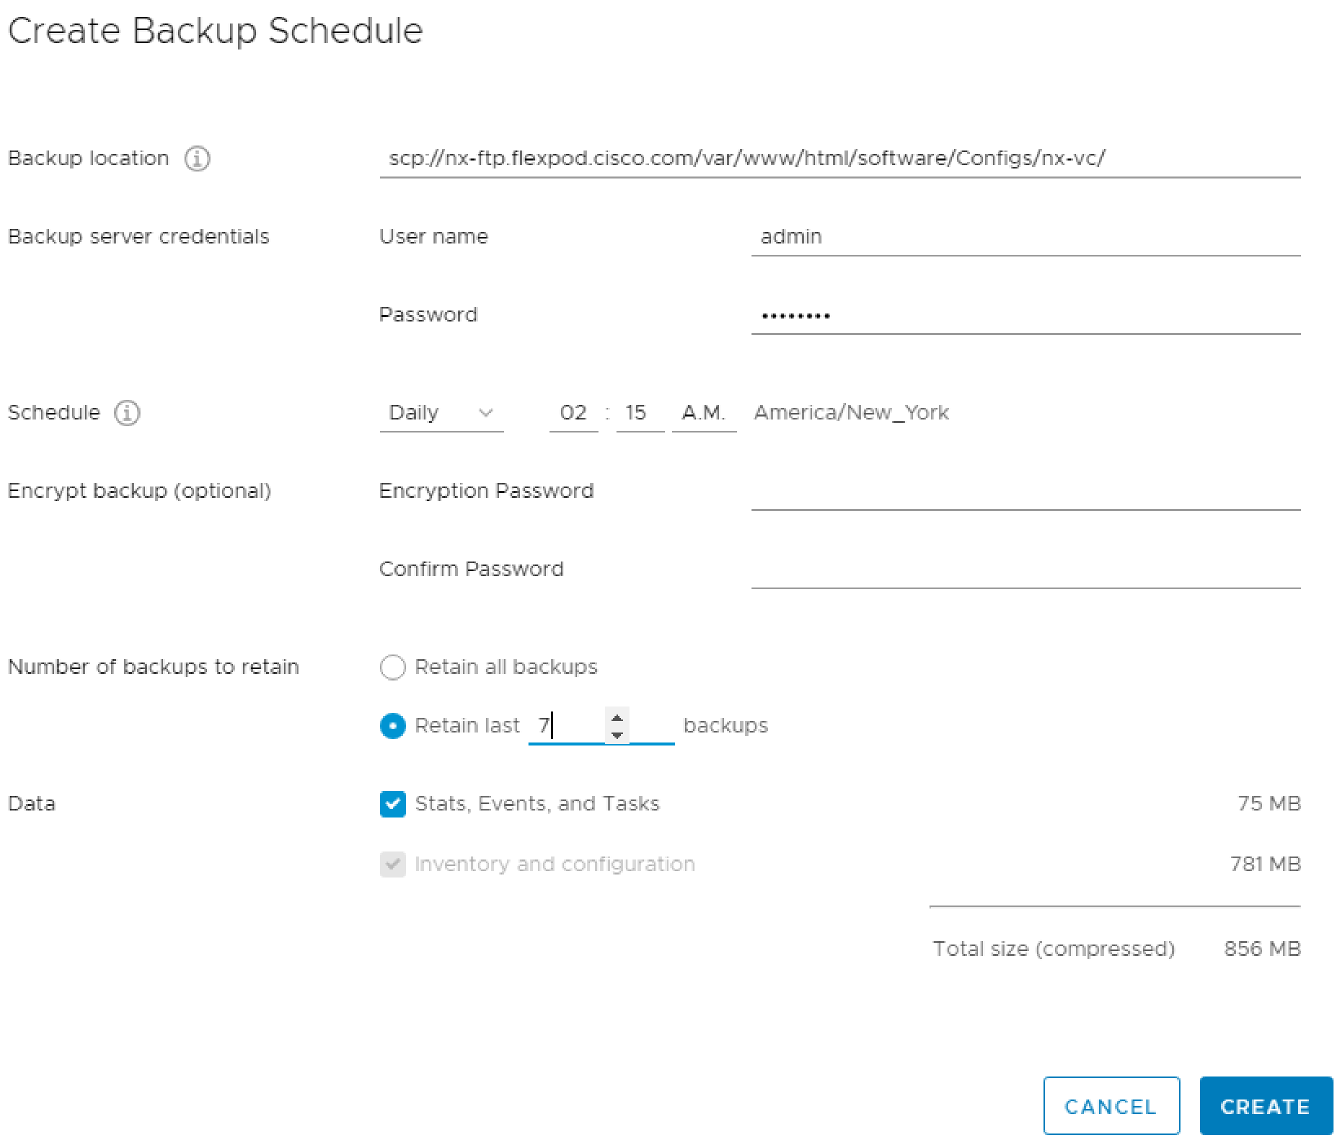

Create Backup Schedule

Backup location: scp://nx-ftp.flexpod.cisco.com/var/www/html/software/Configs/nx-vc/

Backup server credentials
- User name: admin
- Password: ••••••••

Schedule: Daily 02 : 15 A.M. America/New_York

Encrypt backup (optional)
- Encryption Password:
- Confirm Password:

Number of backups to retain
- ○ Retain all backups
- ● Retain last 7 backups

Data
- ☑ Stats, Events, and Tasks 75 MB
- ☑ Inventory and configuration 781 MB

Total size (compressed) 856 MB

5. Click Create.

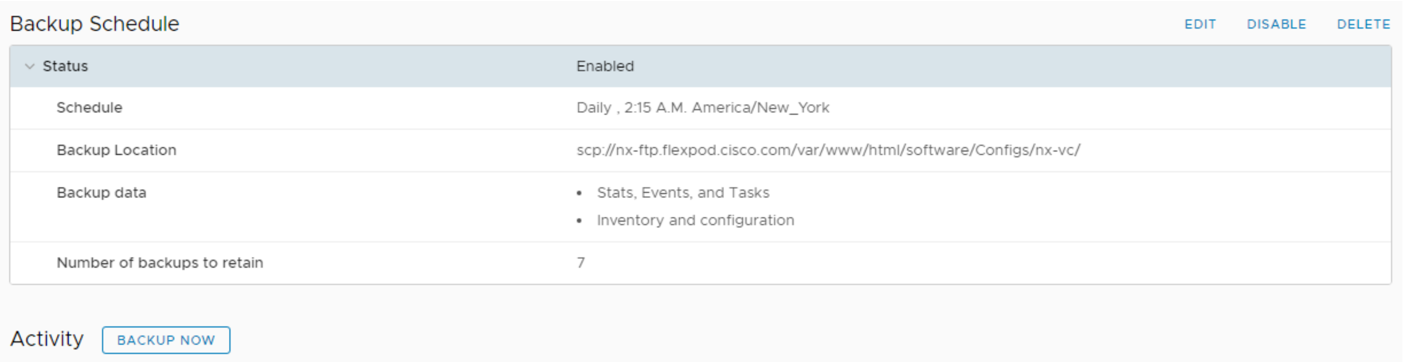

Backup Schedule

Status	Enabled
Schedule	Daily , 2:15 A.M. America/New_York
Backup Location	scp://nx-ftp.flexpod.cisco.com/var/www/html/software/Configs/nx-vc/
Backup data	• Stats, Events, and Tasks • Inventory and configuration
Number of backups to retain	7

Activity

6. The Backup Schedule should now show a Status of Enabled.
7. Restoration can be initiated with the backed-up files using the Restore function of the VCSA 6.7 Installer.

About the Authors

John George, Technical Marketing Engineer, Data Center Solutions Engineering, Cisco Systems, Inc.

John has been involved in designing, developing, validating, and supporting the FlexPod Converged Infrastructure since it was developed over eight years ago. Before his roles with FlexPod, he supported and administered a large worldwide training network and VPN infrastructure. John holds a Master's degree in Computer Engineering from Clemson University.

Scott Kovacs, Technical Marketing Engineer, Infrastructure and Cloud Engineering, NetApp

Scott is a Technical Marketing Engineer in the Converged Infrastructure Engineering team at NetApp. He has been with NetApp since 2007, serving in a variety of Technical Support, Professional Services and engineering roles. Scott has over 20 years of experience in the IT industry specializing in data management, Fibre Channel networking, and security.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.
- Atul Bhalodia, Sr. Technical Marketing Engineer, NetApp