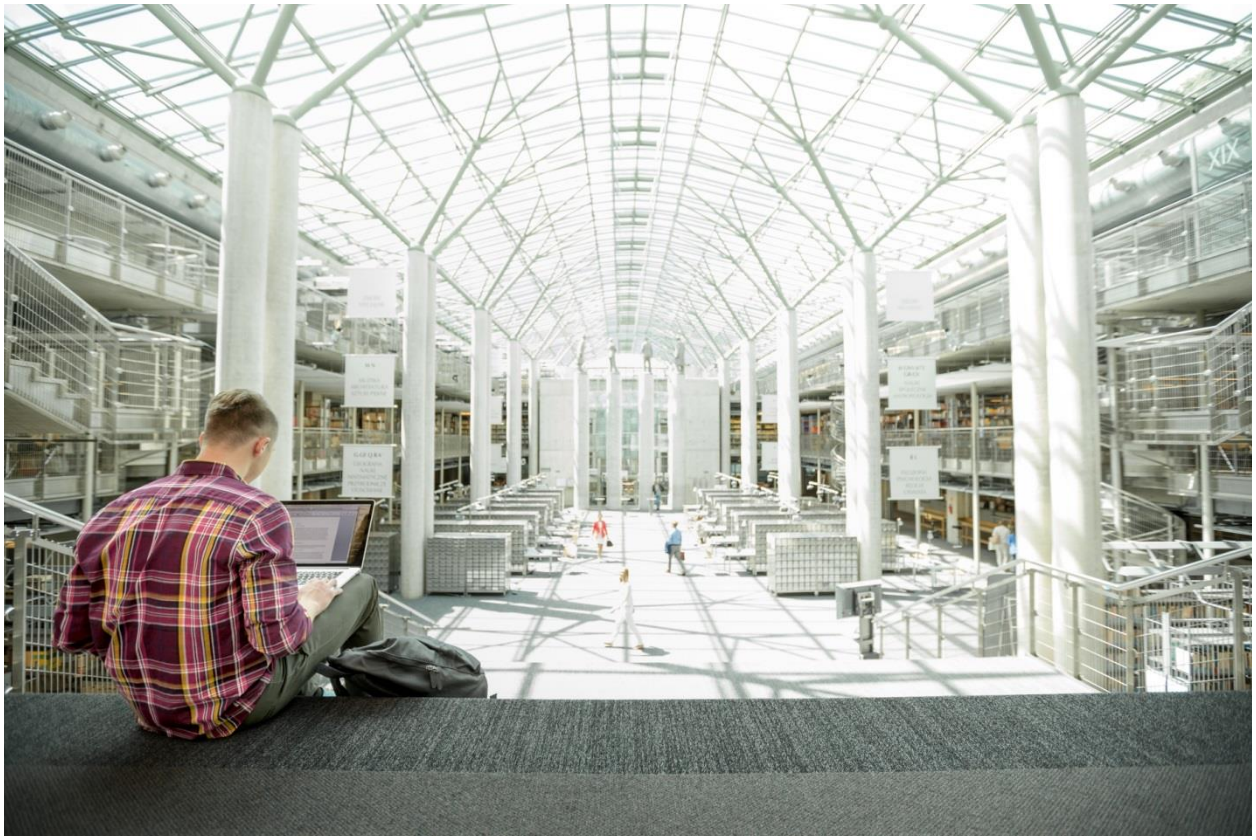

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as Direct Attached Storage

Deployment Guide for Cisco and Hitachi Converged Infrastructure with Cisco UCS Blade Servers, Cisco Nexus 9336C-FX2 Switches, and Hitachi VSP G370 with vSphere 6.5 and vSphere 6.7 Connected as Direct Attached Storage

Last Updated: June 20, 2019

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.

Table of Contents

Executive Summary

Cisco Validated Designs consist of 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 our customers.

Cisco and Hitachi are working together to deliver a converged infrastructure solution that helps enterprise businesses meet the challenges of today and position themselves for the future. Leveraging decades of industry expertise and superior technology, this Cisco CVD offers a resilient, agile, and flexible foundation for today's businesses. In addition, the Cisco and Hitachi partnership extends beyond a single solution, enabling businesses to benefit from their ambitious roadmap of evolving technologies such as advanced analytics, IoT, cloud, and edge capabilities. With Cisco and Hitachi, organizations can confidently take the next step in their modernization journey and prepare themselves to take ad-vantage of new business opportunities enabled by innovative technology.

This document explains the deployment of the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a Virtual Server Infrastructure (VSI). The recommended solution architecture is built on Cisco Unified Computing System (Cisco UCS) using the unified software release to support the Cisco UCS hardware platforms for Cisco UCS B-Series Blade Servers, Cisco UCS 6400 or 6300 Fabric Interconnects, Cisco Nexus 9000 Series switches, and Hitachi Virtual Storage Platform (VSP).

Solution Overview

Introduction

Modernizing your data center can be overwhelming, and it’s vital to select a trusted technology partner with proven expertise. With Cisco and Hitachi as partners, companies can build for the future by enhancing systems of record, supporting systems of innovation, and growing their business. Organizations need an agile solution, free from operational inefficiencies, to deliver continuous data availability, meet SLAs, and prioritize innovation.

Hitachi and Cisco Adaptive Solutions for Converged Infrastructure as a Virtual Server Infrastructure (VSI) is a best practice datacenter architecture built on the collaboration of Hitachi Vantara and Cisco to meet the needs of enterprise customers utilizing virtual server workloads. This architecture is composed of the Hitachi Virtual Storage Platform (VSP) connecting directly to Cisco Unified Computing System (Cisco UCS), and further enabled with the Cisco Nexus family of switches.

These deployment instructions are based on the buildout of the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure validated reference architecture that was added to cover a direct attached storage configuration. The reference architecture covers specifics of products utilized within the Cisco validation lab, but the solution is considered relevant for equivalent supported components listed within Cisco and Hitachi Vantara’s published compatibility matrixes. Supported adjustments from the example validated build must be evaluated with care as their implementation instructions may differ.

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 modernize their infrastructure to meet SLAs and their business needs at any scale.

Purpose of this Document

This document provides a step by step configuration and implementation guide for the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure solution. This solution features a validated reference architecture composed of:

- Cisco UCS Compute
- Cisco Nexus Switches
- Hitachi Virtual Storage Platform

For the design decisions and technology discussion of the solution, please refer to the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci_design.html

The initial work for this CVD using the MDS involved two topologies including the Hitachi VSP G370 paired with the Cisco UCS 6454 Fabric Interconnect, and the Hitachi VSP G1500 paired with the Cisco UCS 6332-16UP Fabric Interconnect, both validating vSphere 6.5 U2 and vSphere 6.7 U1. For the direct attached architecture featured in this deployment guide, the validation set was reduced to the Hitachi VSP G370 paired with the Cisco UCS 6454 Fabric Interconnect validating with vSphere 6.7 U1. The examples shown are from the validated topology, but within the reference architecture, the previously validated components are still considered supported.

Solution Design

Architecture

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure is a validated reference architecture targeting Virtual Server Infrastructure(VSI) implementations. The architecture is built around the Cisco Unified Computing System (Cisco UCS) and the Hitachi Virtual Storage Platform (VSP) connected directly for storage traffic, and further enabled with Cisco Nexus Switches. The Cisco MDS is removed in this design, giving a slightly easier configuration, but adopters should consider that scaling options will be impacted when removing the MDS from the solution. Within this direct attached model, these components are brought together to form a powerful and scalable design, built on the best practices of both companies to create an ideal environment for virtualized systems.

The solution is built and validated featuring the direct connection of the Cisco UCS Fabric Interconnect to the Hitachi VSP Storage System, using the Nexus switching infrastructure to further extend the infrastructure.

The topology shown in Figure 1 leverages:

- Cisco Nexus 9336C-FX2 – 100Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6454 Fabric Interconnect – Unified management of UCS compute, and the compute's access to storage and networks.
- Cisco UCS B200 M5 – High powered, versatile blade server, conceived for virtual computing.
- Hitachi VSP G370 – Mid-range, high-performance storage system with optional all-flash configuration.

Figure 1 Cisco and Hitachi Adaptive Solution for CI with Hitachi VSP G370 and Cisco UCS 6454 Directly Attached

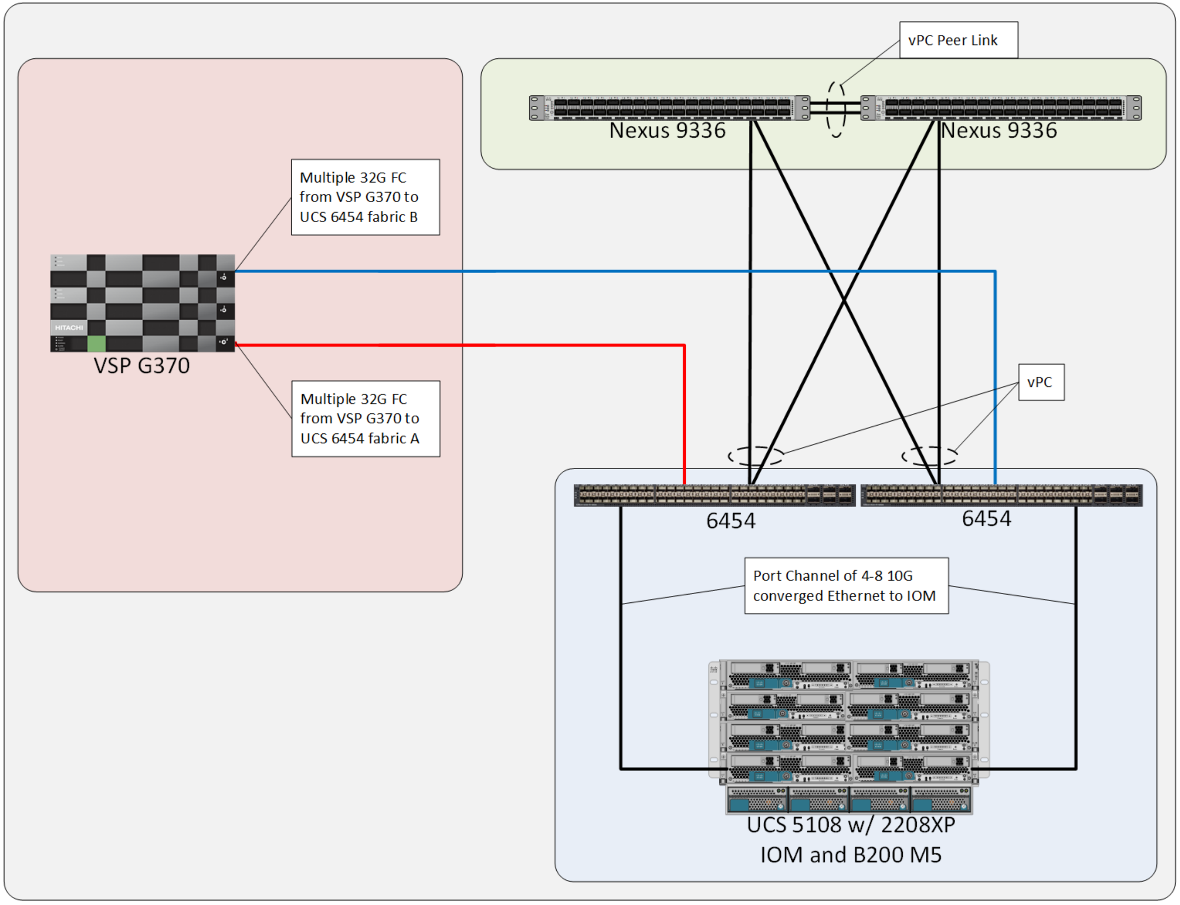
vPC Peer Link
Nexus 9336
Nexus 9336
Multiple 32G FC from VSP G370 to UCS 6454 fabric B
VSP G370
Multiple 32G FC from VSP G370 to UCS 6454 fabric A
vPC
6454
6454
Port Channel of 4-8 10G converged Ethernet to IOM
UCS 5108 w/ 2208XP IOM and B200 M5

The Cisco UCS B200 M5 blade servers in this topology are hosted within a Cisco UCS 5108 Chassis and connect into the fabric interconnects from the chassis using Cisco UCS 2208XP I/O Modules (IOM). The 2208XP IOM supports 10G connections into the 10/25G ports of the Cisco UCS 6454 FIs, delivering a high port availability that may fit well in a branch office setting.

Management components for the architecture additionally include:

- Cisco UCS Manager – Management delivered through the Fabric Interconnect, providing stateless compute, and policy driven implementation of the servers managed by it.
- Cisco Intersight (optional) – Comprehensive unified visibility across UCS domains, along with proactive alerts and enablement of expedited Cisco TAC communications.

The direct attached architecture was validated for vSphere 6.7 U1, with the MDS based topology additionally validated with vSphere 6.5 U2 to accommodate a larger range of expected customer deployments. Previous, and newer versions of vSphere, as well as other vendor hypervisors may be supported. These additional hypervisors must be within the compatibility and interoperability matrices listed at the start of the next section but are not included in this validated design.

Deployment Hardware and Software

Hardware and Software Versions

Table 1 lists the validated hardware and software versions used for this solution. Configuration specifics are given in this deployment guide for the devices and versions listed in the following tables. Component and software version substitution from what is listed is considered acceptable within this reference architecture, but substitution will need to comply with the hardware and software compatibility matrices from both Cisco and Hitachi.

Cisco UCS Hardware Compatibility Matrix:

https://ucshcltool.cloudapps.cisco.com/public/

Cisco Nexus and MDS Interoperability Matrix:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx/Matrix1.html

Cisco Nexus Recommended Releases for Nexus 9K:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended_release/b_Minimum_and_Recommended_Cisco_NX-OS_Releases_for_Cisco_Nexus_9000_Series_Switches.html

Hitachi Vantara Interoperability:

https://support.hitachivantara.com/en_us/interoperability.html sub-page -> (VSP G1X00, F1500, Gxx0, Fxx0, VSP, HUS VM VMWare Support Matrix)

Any substituted hardware or software may have different configurations from what is detailed in this guide and will require a thorough evaluation of the substituted product reference documents.

Table 1 Validated Hardware and Software

Component		Software Version/Firmware Version
Network	Cisco Nexus 9336C-FX2	7.0(3)I7(5a)
Compute	Cisco UCS Fabric Interconnect 6454	4.0(2b)
	Cisco UCS 2208XP IOM	4.0(2b)
	Cisco UCS B200 M5	4.0(2b)
	VMware vSphere	6.7 U1 VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso
	ESXi 6.7 U1 nenic	1.0.27.0
	ESXi 6.7 U1 nfnic	4.0.0.33

Component		Software Version/Firmware Version
	VM Virtual Hardware Version	13(1)
Storage	Hitachi VSP G370	88-02-03-60/00

(1) Hardware Version 13 was kept for initial transfer support between vSphere 6.5 and vSphere 6.7 of the VM test harness used.

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for the Cisco and Hitachi Converged Infrastructure. References are made to which component is being configured with each step, either "-1" or "-2". For example, AA19-9336-1 and AA19-9336-2 are used to identify the two Nexus switches that are provisioned with this document, with AA19-9336-1 and 2 used to represent a command invoked on both Nexus switches. 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 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 of a configuration step for both Nexus switches:

```
AA19-9336-1&2  (config)# ntp server <<var_oob_ntp>> use-vrf management
```

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. The tables provided can be copied or printed for use as a reference to align the appropriate customer deployed values for configuration specifics used within the guide.

Table 2 lists the VLANs necessary for deployment as outlined in this guide.

Table 2 VLANs Used in the Deployment

VLAN Name	VLAN Purpose	ID Used in Validating this Document	Customer Deployed Value
Out of Band Mgmt	VLAN for out-of-band management interfaces	19	
In-Band Mgmt	VLAN for in-band management interfaces	119	
Native	VLAN to which untagged frames are assigned	2	
vMotion	VLAN for VMware vMotion	1000	
VM-App1	VLAN for Production VM Interfaces	201	
VM-App2	VLAN for Production VM Interfaces	202	
VM-App3	VLAN for Production VM Interfaces	203	

Table 3 lists additional configuration variables are used throughout the document as pointers to where a customer provided name, or reference for relevant existing information will be used.

Table 3 Variables for Information Used in the Design

Variable	Variable Description	Customer Deployed Value
<<var_nexus_A_hostname>>	Nexus switch A hostname (Example: b19-93180-1)	
<<var_nexus_A_mgmt_ip>>	Out-of-band management IP for Nexus switch A (Example: 192.168.164.13)	
<<var_nexus_B_hostname>>	Nexus switch B hostname (Example: b19-93180-2)	
<<var_nexus_B_mgmt_ip>>	Out-of-band management IP for Nexus switch B (Example: 192.168.164.14)	
<<var_oob_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)	
<<var_oob_gateway>>	Out-of-band management network gateway (Example: 192.168.164.254)	
<<var_oob_ntp>>	Out-of-band management network NTP server (Example: 192.168.164.254)	
<<var_nexus_A_ib_ip>>	In-band management HSRP network interface Nexus switch A (Example: 10.1.164.252)	
<<var_nexus_B_ib_ip>>	In-band management HSRP network interface for Nexus switch B (Example: 10.1.164.253)	
<<var_nexus_ib_vip>>	In-band management HSRP network VIP (Example: 10.1.164.254)	
<<var_password>>	Administrative password (Example: NotaP4ss)	
<<var_dns_domain_name>>	DNS domain name (Example: ucp.cisco.com)	
<<var_nameserver_ip>>	DNS server IP(s) (Example: 10.1.168.9)	
<<var_timezone>>	Time zone (Example: America/New_York)	
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID (Example: 119)	
<<var_ib_mgmt_vlan_netmask_length>>	Length of IB-MGMT-VLAN Netmask (Example: /24)	
<<var_ib_gateway_ip>>	In-band management network VLAN ID (Example: 10.1.168.1)	
<<var_vmotion_vlan_id>>	vMotion management network VLAN ID (Example: 1000)	
<<var_vmotion_vlan_netmask_length>>	Length of vMotion-VLAN Netmask (Example: /24)	
<<var_vsan_a_id>>	VSAN used for the A Fabric between the VSP /FI (Example: 101)	
<<var_vsan_b_id>>	VSAN used for the A Fabric between the VSP /FI (Example: 102)	

Variable	Variable Description	Customer Deployed Value
<<vsp_hostname>> <<vsp-g370>>	Hitachi VSP storage system name (Example g370-[Serial Number])	
<<var_ucs_clustername>> <<var_ucs_6454_clustername>>	Cisco UCS Manager cluster host name (Example: AA19-6454)	
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address (Example: 192.168.168.16)	
<<var_ucs_mgmt_vip>>	Cisco UCS fabric interconnect (FI) Cluster out-of-band management IP address (Example: 192.168.168.15)	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address (Example: 192.168.168.17)	
<<var_vm_host_infra_01_ip>>	VMware ESXi host 01 in-band management IP (Example: 10.1.168.21)	
<<var_vm_host_infra_02_ip>>	VMware ESXi host 02 in-band management IP (Example: 10.1.168.22)	
<<var_vm_host_infra_vmotion_01_ip>>	VMware ESXi host 01 vMotion IP (Example: 192.168.100.21)	
<<var_vm_host_infra_vmotion_02_ip>>	VMware ESXi host 02 vMotion IP (Example: 192.168.100.22)	
<<var_vmotion_subnet_mask>>	vMotion subnet mask (Example: 255.255.255.0)	
<<var_vcenter_server_ip>>	IP address of the vCenter Server (Example: 10.1.168.100)	

Physical Cabling

This section explains the cabling examples used for the validated topology in the environment. To make connectivity clear in this example, the tables include both the local and remote port locations.

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. The upstream network from the Nexus 9336C-FX2 switches is out of scope of this document, with only the assumption that these switches will connect to the upstream switch or switches with a virtual Port Channel (vPC).

Physical Cabling for the Cisco UCS 6454 with the VSP G370 Topology

Figure 2 shows the cabling configuration used in the design featuring the Cisco UCS 6454 with the VSP G370.

Figure 2 Cabling Diagram for Cisco and Hitachi Converged Infrastructure Featuring Cisco UCS 6454 with the VSP G370

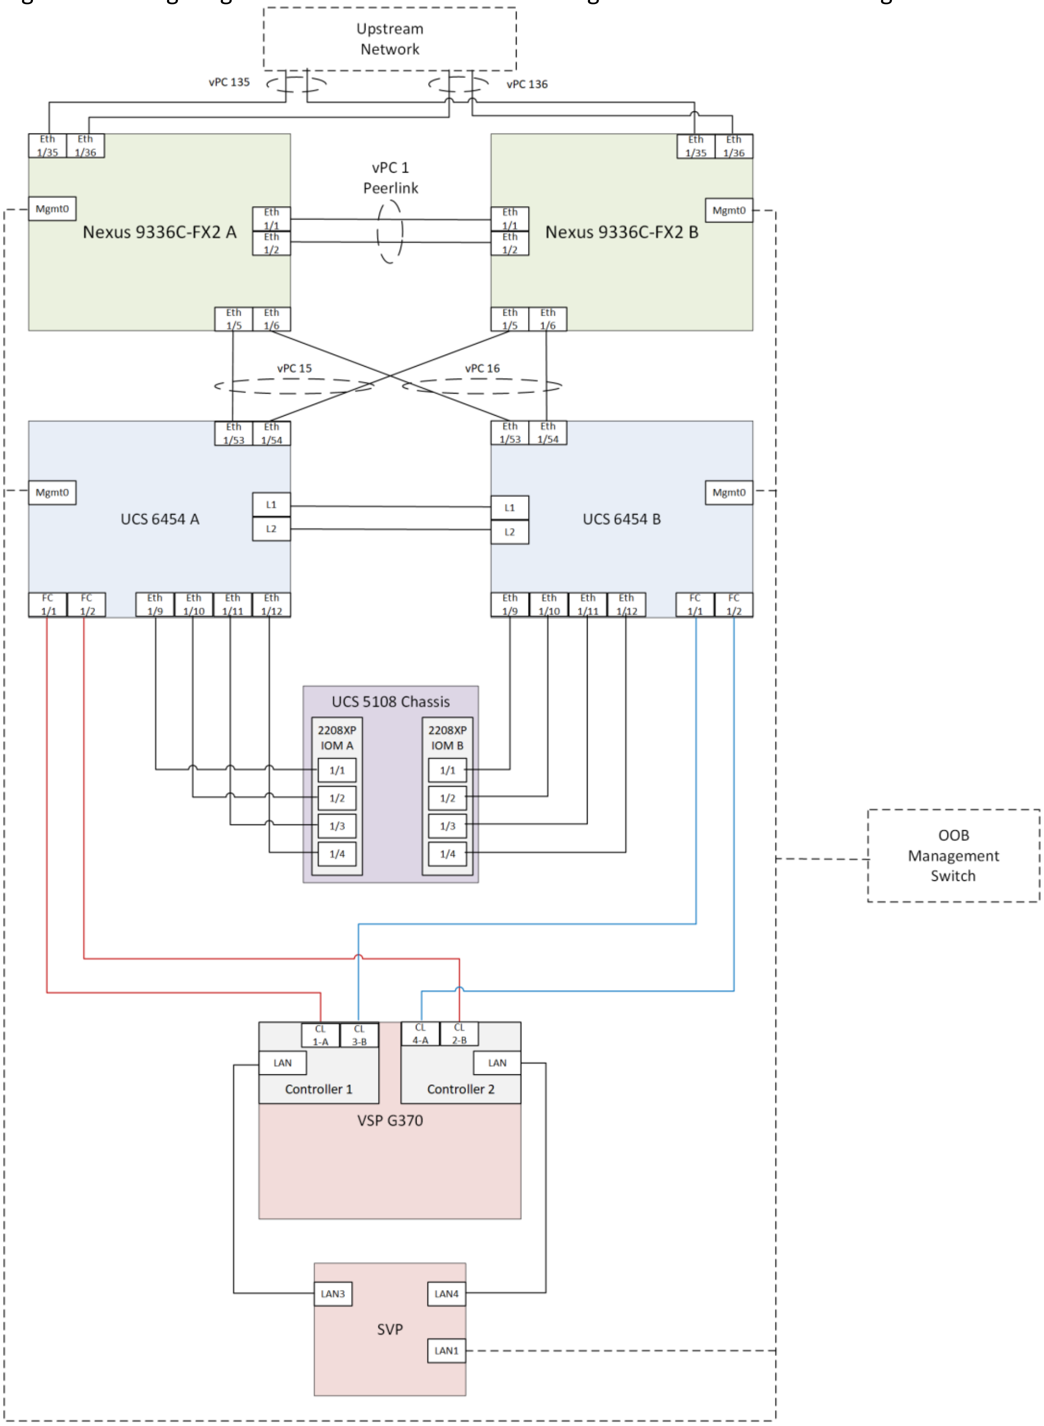

The following tables list the specific port connections with the cables used in the deployment of the Cisco UCS 6454 and the VSP G370.

Table 4 Cisco Nexus 9336C-FX2 A Cabling Information for Cisco UCS 6454 to VSP G370

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336C-FX2 A	Eth1/1	40GbE	Cisco Nexus 9336C-FX2 B	Eth1/1
	Eth1/2	40GbE	Cisco Nexus 9336C-FX2 B	Eth1/2

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/5	100GbE	Cisco UCS 6454 FI A	Eth 1/53
	Eth1/6	100GbE	Cisco UCS 6454 FI B	Eth 1/53
	Eth1/35	40GbE or 100GbE	Upstream Network Switch	Any
	Eth1/36	40GbE or 100GbE	Upstream Network Switch	Any
	MGMT0	GbE	GbE management switch	Any

Selecting 100GbE between the Nexus 9336C-FX2 switches and the Cisco UCS 6454 fabric interconnects is not required but was selected as an available option between the devices.

Table 5 Cisco Nexus 9336C-FX2 B Cabling Information for Cisco UCS 6454 to VSP G370

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336C-FX2 B	Eth1/1	40GbE	Cisco Nexus 9336C-FX2 A	Eth1/1
	Eth1/2	40GbE	Cisco Nexus 9336C-FX2 A	Eth1/2
	Eth1/5	100GbE	Cisco UCS 6454 FI A	Eth 1/54
	Eth1/6	100GbE	Cisco UCS 6454 FI B	Eth 1/54
	Eth1/35	40GbE or 100GbE	Upstream Network Switch	Any
	Eth1/36	40GbE or 100GbE	Upstream Network Switch	Any
	MGMT0	GbE	GbE management switch	Any

Table 6 Cisco UCS 6454 A Cabling Information for Cisco UCS 6454 to VSP G370

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6454 FI A	FC 1/1	32Gb FC	VSP G370	CL1-A
	FC 1/2	32Gb FC	VSP G370	CL2-B
	Eth1/9	10GbE	Cisco UCS Chassis 2208XP FEX A	IOM 1/1
	Eth1/10	10GbE	Cisco UCS Chassis 2208XP FEX A	IOM 1/2
	Eth1/11	10GbE	Cisco UCS Chassis 2208XP FEX A	IOM 1/3
	Eth1/12	10GbE	Cisco UCS Chassis 2208XP FEX A	IOM 1/4
	Eth1/33	40GbE	Cisco Nexus 9336C-FX2 A	Eth1/5
	Eth1/34	40GbE	Cisco Nexus 9336C-FX2 B	Eth1/5
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6454 FI B	L1
	L2	GbE	Cisco UCS 6454 FI B	L2

Ports 1-8 on the Cisco UCS 6454 are unified ports that can be configured as Ethernet or as Fibre Channel ports. Server ports should be initially deployed started with 1/9 to give flexibility for FC port needs, and ports 49-54 are not configurable for server ports. Also, ports 45-48 are the only configurable ports for 1Gbps connections that may be needed to a network switch.

Table 7 Cisco UCS 6454 B Cabling Information for Cisco UCS 6454 to VSP G370 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6454 FI B	FC 1/1	32Gb FC	VSP G370	CL3-B
	FC 1/2	32Gb FC	VSP G370	CL4-A
	Eth1/9	10GbE	Cisco UCS Chassis 2208XP FEX B	IOM 1/1
	Eth1/10	10GbE	Cisco UCS Chassis 2208XP FEX B	IOM 1/2
	Eth1/11	10GbE	Cisco UCS Chassis 2208XP FEX B	IOM 1/3
	Eth1/12	10GbE	Cisco UCS Chassis 2208XP FEX B	IOM 1/4
	Eth1/33	40GbE	Cisco Nexus 9336C-FX2 A	Eth1/6
	Eth1/34	40GbE	Cisco Nexus 9336C-FX2 B	Eth1/6
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6454 FI A	L1
	L2	GbE	Cisco UCS 6454 FI A	L2

Table 8 Hitachi VSP G370 Cabling Information for Cisco UCS 6454 to VSP G370 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G370	CL 1-A	32Gb FC	Cisco UCS 6454 FI A	FC 1/1
	CL 2-B	32Gb FC	Cisco UCS 6454 FI A	FC 1/2
	CL 3-B	32Gb FC	Cisco UCS 6454 FI B	FC 1/1
	CL 4-A	32Gb FC	Cisco UCS 6454 FI B	FC 1/2
	Cont1 LAN	GbE	SVP	LAN3
	Cont2 LAN	GbE	SVP	LAN4

16Gb or 32Gb FC can be used with the VSP G370 and 16Gb FC can be used with the VSP G1500, 8Gb FC cannot be used on either VSP model when using a direct attached topology as appropriate fill patterns are not compatible with 8Gb.

SVP will be configured by a Hitachi Vantara support engineer at the time of initial configuration and is out of scope of the primary deployment.

Cisco Nexus Switch Configuration

The Nexus switch configuration will explain the basic L2 and L3 functionality for the application environment used in the validation environment hosted by the UCS domains. The application gateways are hosted by the pair of Nexus switches, but primary routing is passed onto an existing router that is upstream of the converged infrastructure. This upstream router will need to be aware of any networks created on the Nexus switches, but configuration of an upstream router is beyond the scope of this deployment guide.

Physical Connectivity

Physical cabling should be completed by following the diagram and table references found in section Deployment Hardware and Software.

Initial Nexus Configuration Dialogue

Complete this dialogue on each switch, using a serial connection to the console port of the switch, unless Power on Auto Provisioning is being used.

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
         ---- System Admin Account Setup ----


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

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

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

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

Please register Cisco Nexus9000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus9000 devices must be registered to receive
entitled support services.

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

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


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

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

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

  Enter the switch name : <<var_nexus_A_hostname>>|<<var_nexus_B_hostname>>

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

    Mgmt0 IPv4 address : << var_nexus_A_mgmt_ip>>|<< var_nexus_B_mgmt_ip>>
    Mgmt0 IPv4 netmask : <<var_oob_mgmt netmask>

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

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

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

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

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

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

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

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

  NTP server IPv4 address: <<var_oob_ntp>>

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

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

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

The following configuration will be applied:
  password strength-check
  switchname AA19-9336-1
vrf context management
ip route 0.0.0.0/0 192.168.168.254
exit
  no feature telnet
  ssh key rsa 1024 force
  feature ssh
  system default switchport
  system default switchport shutdown
  copp profile strict
interface mgmt0
ip address 192.168.168.13 255.255.255.0
no shutdown

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

Enable Features and Settings

To enable IP switching features, run the following commands on each Cisco Nexus:

```
AA19-9336-1&2 (config)# feature lacp
AA19-9336-1&2 (config)# feature vpc
AA19-9336-1&2 (config)# feature interface-vlan
AA19-9336-1&2 (config)# feature hsrp
```

The reference of `AA19-9336-1&2` is used to represent a command run on both switches, `AA19-9336-1` represents a command to run only on the first Nexus switch, and `AA19-9336-2` stands for a command that should only be run on the second Nexus switch.

Additionally, configure the spanning tree and save the running configuration to start-up:

```
AA19-9336-1&2 (config)# spanning-tree port type network default
AA19-9336-1&2 (config)# spanning-tree port type edge bpduguard default
AA19-9336-1&2 (config)# spanning-tree port type edge bpdufilter default
```

Create VLANs

Run the following commands on both switches to create VLANs:

```
AA19-9336-1&2 (config)# vlan 119
AA19-9336-1&2 (config-vlan)# name IB-MGMT
AA19-9336-1&2 (config-vlan)# vlan 2
```

```
AA19-9336-1&2 (config-vlan)# name Native
AA19-9336-1&2 (config-vlan)# vlan 1000
AA19-9336-1&2 (config-vlan)# name vMotion
AA19-9336-1&2 (config-vlan)# vlan 201
AA19-9336-1&2 (config-vlan)# name Web
AA19-9336-1&2 (config-vlan)# vlan 202
AA19-9336-1&2 (config-vlan)# name App
AA19-9336-1&2 (config-vlan)# vlan 203
AA19-9336-1&2 (config-vlan)# name DB
AA19-9336-1&2 (config-vlan)# exit
```

Continue adding VLANs as appropriate to your environment.

Add Individual Port Descriptions for Troubleshooting

To add individual port descriptions for troubleshooting activity and verification for switch A, enter the following commands from the global configuration mode:

```
AA19-9336-1(config)# interface port-channel 11
AA19-9336-1(config-if)# description vPC peer-link
AA19-9336-1(config-if)# interface port-channel 15
AA19-9336-1(config-if)# description vPC UCS 6454-1 FI
AA19-9336-1(config-if)# interface port-channel 16
AA19-9336-1(config-if)# description vPC UCS 6454-2 FI
AA19-9336-1(config-if)# interface port-channel 135
AA19-9336-1(config-if)# description vPC Upstream Network Switch A
AA19-9336-1(config-if)# interface port-channel 136
AA19-9336-1(config-if)# description vPC Upstream Network Switch B
```

The port-channel numbers will need to match between the two switches, and while the port numbering can be somewhat arbitrary, a numbering scheme of the first port in the port channel is represented in the numbering scheme used, where port channel 11 has a first port of 1/1, and port channel 136 has a first port of 1/36.

```
AA19-9336-1(config-if)# interface Ethernet1/1
AA19-9336-1(config-if)# description vPC peer-link connection to AA19-9336-2 Ethernet1/1
AA19-9336-1(config-if)# interface Ethernet1/2
AA19-9336-1(config-if)# description vPC peer-link connection to AA19-9336-2 Ethernet1/2
AA19-9336-1(config-if)# interface Ethernet1/5
AA19-9336-1(config-if)# description vPC 15 connection to UCS 6454-1 FI Ethernet1/53
AA19-9336-1(config-if)# interface Ethernet1/6
AA19-9336-1(config-if)# description vPC 16 connection to UCS 6454-2 FI Ethernet1/53
AA19-9336-1(config-if)# interface Ethernet1/35
AA19-9336-1(config-if)# description vPC 135 connection to Upstream Network Switch A
AA19-9336-1(config-if)# interface Ethernet1/36
AA19-9336-1(config-if)# description vPC 136 connection to Upstream Network Switch B
AA19-9336-1(config-if)# exit
```

In these steps, the interface commands for the VLAN interface and Port Channel interfaces, will create these interfaces if they do not already exist.

To add individual port descriptions for troubleshooting activity and verification for switch B, enter the following commands from the global configuration mode:

```
AA19-9336-2(config)# interface port-channel 11
AA19-9336-2(config-if)# description vPC peer-link
AA19-9336-2(config-if)# interface port-channel 15
AA19-9336-2(config-if)# description vPC UCS 6454-1 FI
AA19-9336-2(config-if)# interface port-channel 16
AA19-9336-2(config-if)# description vPC UCS 6454-2 FI
AA19-9336-2(config-if)# interface port-channel 135
AA19-9336-2(config-if)# description vPC Upstream Network Switch A
AA19-9336-2(config-if)# interface port-channel 136
AA19-9336-2(config-if)# description vPC Upstream Network Switch B
```

```
AA19-9336-2(config-if)# interface Ethernet1/1
AA19-9336-2(config-if)# description vPC peer-link connection to AA19-9336-1 Ethernet1/1
AA19-9336-2(config-if)# interface Ethernet1/2
AA19-9336-2(config-if)# description vPC peer-link connection to AA19-9336-1 Ethernet1/2
AA19-9336-2(config-if)# interface Ethernet1/5
AA19-9336-2(config-if)# description vPC 15 connection to UCS 6454-1 FI Ethernet1/54
AA19-9336-2(config-if)# interface Ethernet1/6
AA19-9336-2(config-if)# description vPC 16 connection to UCS 6454-2 FI Ethernet1/54
AA19-9336-2(config-if)# interface Ethernet1/35
AA19-9336-2(config-if)# description vPC 135 connection to Upstream Network Switch A
AA19-9336-2(config-if)# interface Ethernet1/36
AA19-9336-2(config-if)# description vPC 136 connection to Upstream Network Switch B
AA19-9336-2(config-if)# exit
```

Create the vPC Domain

The vPC domain will be assigned a unique number from 1-1000 and will handle the vPC settings specified within the switches. To set the vPC domain configuration on 9336C-FX2 A, run the following commands:

```
AA19-9336-1(config)# vpc domain 10
AA19-9336-1(config-vpc-domain)# peer-switch
AA19-9336-1(config-vpc-domain)# role priority 10
AA19-9336-1(config-vpc-domain)# peer-keepalive destination <<var_nexus_B_mgmt_ip>> source
<<var_nexus_A_mgmt_ip>>
AA19-9336-1(config-vpc-domain)# delay restore 150
AA19-9336-1(config-vpc-domain)# peer-gateway
AA19-9336-1(config-vpc-domain)# auto-recovery
AA19-9336-1(config-vpc-domain)# ip arp synchronize
AA19-9336-1(config-vpc-domain)# exit
```

On the 9336C-FX2 B switch run these slightly differing commands, noting that role priority and peer-keepalive commands will differ from what was previously set:

```
AA19-9336-2(config)# vpc domain 10
AA19-9336-2(config-vpc-domain)# peer-switch
AA19-9336-2(config-vpc-domain)# role priority 20
AA19-9336-2(config-vpc-domain)# peer-keepalive destination <<var_nexus_A_mgmt_ip>> source
<<var_nexus_B_mgmt_ip>>
AA19-9336-2(config-vpc-domain)# delay restore 150
AA19-9336-2(config-vpc-domain)# peer-gateway
AA19-9336-2(config-vpc-domain)# auto-recovery
AA19-9336-2(config-vpc-domain)# ip arp synchronize
AA19-9336-2(config-vpc-domain)# exit
```

Configure Port Channel Member Interfaces

On each switch, configure the Port Channel member interfaces that will be part of the vPC Peer Link and configure the vPC Peer Link:

```
AA19-9336-1&2 (config)# int eth 1/1-2
AA19-9336-1&2 (config-if-range)# channel-group 11 mode active
AA19-9336-1&2 (config-if-range)# no shut
AA19-9336-1&2 (config-if-range)# int port-channel 11
AA19-9336-1&2 (config-if)# switchport mode trunk
AA19-9336-1&2 (config-if)# switchport trunk native vlan 2
AA19-9336-1&2 (config-if)# switchport trunk allowed vlan 119,1000,201-203
AA19-9336-1&2 (config-if)# vpc peer-link
```

Configure Virtual Port Channels

On each switch, configure the Port Channel member interfaces and the vPC Port Channels to the Cisco UCS Fabric Interconnect and the upstream network switches:

Nexus Connection vPC to Cisco UCS 6454 A

```
AA19-9336-1&2 (config-if)# int ethernet 1/5
AA19-9336-1&2 (config-if)# channel-group 15 mode active
AA19-9336-1&2 (config-if)# no shut
AA19-9336-1&2 (config-if)# int port-channel 15
AA19-9336-1&2 (config-if)# switchport mode trunk
AA19-9336-1&2 (config-if)# switchport trunk native vlan 2
AA19-9336-1&2 (config-if)# switchport trunk allowed vlan 119,1000,201-203
AA19-9336-1&2 (config-if)# spanning-tree port type edge trunk
AA19-9336-1&2 (config-if)# mtu 9216
AA19-9336-1&2 (config-if)# load-interval counter 3 60
AA19-9336-1&2 (config-if)# vpc 15
```

Nexus Connection vPC to Cisco UCS 6454 B

```
AA19-9336-1&2 (config-if)# int ethernet 1/6
AA19-9336-1&2 (config-if)# channel-group 16 mode active
AA19-9336-1&2 (config-if)# no shut
AA19-9336-1&2 (config-if)# int port-channel 16
AA19-9336-1&2 (config-if)# switchport mode trunk
AA19-9336-1&2 (config-if)# switchport trunk native vlan 2
AA19-9336-1&2 (config-if)# switchport trunk allowed vlan 119,1000,201-203
AA19-9336-1&2 (config-if)# spanning-tree port type edge trunk
AA19-9336-1&2 (config-if)# mtu 9216
AA19-9336-1&2 (config-if)# load-interval counter 3 60
AA19-9336-1&2 (config-if)# vpc 16
```

Nexus Connection vPC to Upstream Network Switch A

```
AA19-9336-1&2 (config-if)# interface Ethernet1/35
AA19-9336-1&2 (config-if)# channel-group 135 mode active
AA19-9336-1&2 (config-if)# no shut
AA19-9336-1&2 (config-if)# int port-channel 135
AA19-9336-1&2 (config-if)# switchport mode trunk
AA19-9336-1&2 (config-if)# switchport trunk native vlan 2
AA19-9336-1&2 (config-if)# switchport trunk allowed vlan 119
AA19-9336-1&2 (config-if)# vpc 135
```

Nexus Connection vPC to Upstream Network Switch B

```
AA19-9336-1&2 (config-if)# interface Ethernet1/36
AA19-9336-1&2 (config-if)# channel-group 136 mode active
AA19-9336-1&2 (config-if)# no shut
AA19-9336-1&2 (config-if)# int port-channel 136
AA19-9336-1&2 (config-if)# switchport mode trunk
AA19-9336-1&2 (config-if)# switchport trunk native vlan 2
AA19-9336-1&2 (config-if)# switchport trunk allowed vlan 119
AA19-9336-1&2 (config-if)# vpc 136
```

Create Hot Standby Router Protocol (HSRP) Switched Virtual Interfaces (SVI)

These interfaces can be considered optional if the subnets of the VLANs used within the environment are managed entirely by an upstream switch, but if that is the case, all managed VLANs will need to be carried up through the vPC to the Upstream switches.

More advanced Cisco routing protocols can be configured within the Nexus switches but are not covered in this design. Routing between the SVIs is directly connected between them as they reside in the same Virtual Routing and Forwarding instance (VRF), and traffic set to enter and exit the VRF will traverse the default gateway set for the switches.

For 9336C-FX2 A:

Nexus A IB-Mgmt SVI

```
AA19-9336-1(config-if)# int vlan 119
AA19-9336-1(config-if)# no shutdown
AA19-9336-1(config-if)# ip address <<var_nexus_A_ib_ip>>/24
AA19-9336-1(config-if)# hsrp 19
AA19-9336-1(config-if-hsrp)# preempt
AA19-9336-1(config-if-hsrp)# ip <<var_nexus_ib_vip>>
```

When HSRP priority is not set, it defaults to 100. Alternating SVIs within a switch are set to a number higher than 105 to set those SVIs to default to be the standby router for that network. Be careful when the VLAN SVI for one switch is set without a priority (defaulting to 100), the partner switch is set to a priority with a value other than 100.

Nexus A Web SVI

```
AA19-9336-1(config-if-hsrp)# int vlan 201
AA19-9336-1(config-if)# no shutdown
AA19-9336-1(config-if)# ip address 172.18.101.252/24
AA19-9336-1(config-if)# hsrp 101
AA19-9336-1(config-if-hsrp)# preempt
AA19-9336-1(config-if-hsrp)# priority 105
AA19-9336-1(config-if-hsrp)# ip 172.18.101.254
```

Nexus A App SVI

```
AA19-9336-1(config-if-hsrp)# int vlan 202
AA19-9336-1(config-if)# no shutdown
AA19-9336-1(config-if)# ip address 172.18.102.252/24
AA19-9336-1(config-if)# hsrp 102
AA19-9336-1(config-if-hsrp)# preempt
AA19-9336-1(config-if-hsrp)# ip 172.18.102.254
```

Nexus A DB SVI

```
AA19-9336-1(config-if-hsrp)# int vlan 203
AA19-9336-1(config-if)# no shutdown
AA19-9336-1(config-if)# ip address 172.18.103.252/24
AA19-9336-1(config-if)# hsrp 103
AA19-9336-1(config-if-hsrp)# preempt
AA19-9336-1(config-if-hsrp)# priority 105
AA19-9336-1(config-if-hsrp)# ip 172.18.103.254
```

For 9336C-FX2 B:

Nexus B IB-Mgmt SVI

```
AA19-9336-2(config-if)# int vlan 119
AA19-9336-2(config-if)# no shutdown
AA19-9336-2(config-if)# ip address <<var_nexus_B_ib_ip>>/24
AA19-9336-2(config-if)# hsrp 19
AA19-9336-2(config-if-hsrp)# preempt
AA19-9336-2(config-if-hsrp)# priority 105
AA19-9336-2(config-if-hsrp)# <<var_nexus_ib_vip>>
```

Nexus B Web SVI

```
AA19-9336-2(config-if-hsrp)# int vlan 201
AA19-9336-2(config-if)# no shutdown
AA19-9336-2(config-if)# ip address 172.18.101.253/24
AA19-9336-2(config-if)# hsrp 101
AA19-9336-2(config-if-hsrp)# preempt
AA19-9336-2(config-if-hsrp)# ip 172.18.101.254
```

Nexus B App SVI

```
AA19-9336-2(config-if-hsrp)# int vlan 202
AA19-9336-2(config-if)# no shutdown
AA19-9336-2(config-if)# ip address 172.18.102.253/24
AA19-9336-2(config-if)# hsrp 102
AA19-9336-2(config-if-hsrp)# preempt
AA19-9336-2(config-if-hsrp)# priority 105
AA19-9336-2(config-if-hsrp)# ip 172.18.102.254
```

Nexus B DB SVI

```
AA19-9336-2(config-if-hsrp)# int vlan 203
AA19-9336-2(config-if)# no shutdown
AA19-9336-2(config-if)# ip address 172.18.103.253/24
AA19-9336-2(config-if)# hsrp 103
AA19-9336-2(config-if-hsrp)# preempt
AA19-9336-2(config-if-hsrp)# ip 172.18.103.254
```

Set Global Configurations

Run the following commands on both switches to set global configurations:

```
AA19-9336-1&2 (config-if-hsrp)# port-channel load-balance src-dst l4port
AA19-9336-1&2 (config)# ip route 0.0.0.0/0 <<var_ib_gateway_ip>>
AA19-9336-1&2 (config)# ntp server <<var_oob_ntp>> use-vrf management
```

In the above command block, the “l4port” contains the letter L and 4, not the number fourteen.

The ntp server should be an accessible NTP server for use by the switches. In this case, point to an out-of-band source.

```
AA19-9336-1&2 (config)# ntp master 3
AA19-9336-1&2 (config)# ntp source <<var_nexus_ib_vip>>
```

Setting the switches as ntp masters to redistribute as an ntp source is optional here but can be a valuable fix if the tenant networks are not enabled to reach the primary ntp server.

*** Save all configurations to this point on both Nexus Switches ***

```
AA19-9336-1&2 (config)# copy running-config startup-config
```

Configure Fibre Channel Ports on Hitachi Virtual Storage Platform

In order for Hitachi Virtual Storage Platform fibre channel ports to be exposed properly to the Cisco UCS components, modification of the ports from their default values must be performed. Prior to beginning this section, ensure that you have credentials on the Hitachi Virtual Storage Platform that have at least the **Administrator** role permissions within Hitachi Storage Navigator. Your partner or Hitachi services personnel provide credentials to your Hitachi Virtual Storage Platform after initial setup and configuration of the storage system.

To configure the fibre channel ports within the VSP storage system, follow these steps:

1. Access Hitachi Storage Navigator through a web browser.

2. **VSP Fxoo Models and VSP Gxoo Models**: https://<IP of Storage System SVP>/dev/storage/886000<Serial Number of Storage System>/emergency.do – for example, if Storage System SVP IP address is 10.0.0.2 and Serial Number of Storage System is 451200, the URL would be:

   https://10.0.0.2/dev/storage/886000451200/emergency.do

3. Log into Hitachi Storage Navigator.

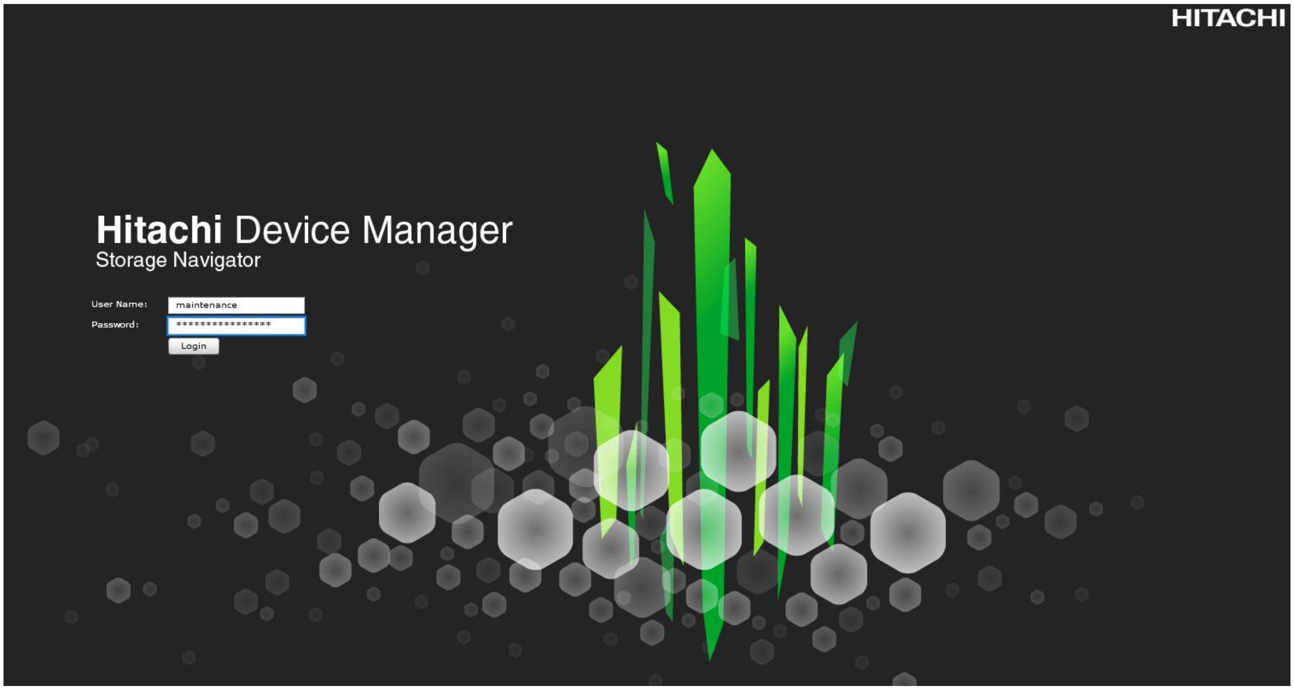

Hitachi Device Manager
Storage Navigator

User Name: maintenance
Password: ****************

4. From the left Explorer pane, select the **Storage Systems** tab.

5. Expand the storage system being configured. Highlight the **Ports/Host Groups/iSCSI Targets** element in the navigation tree, then click on the **Ports** tab in the main configuration pane.

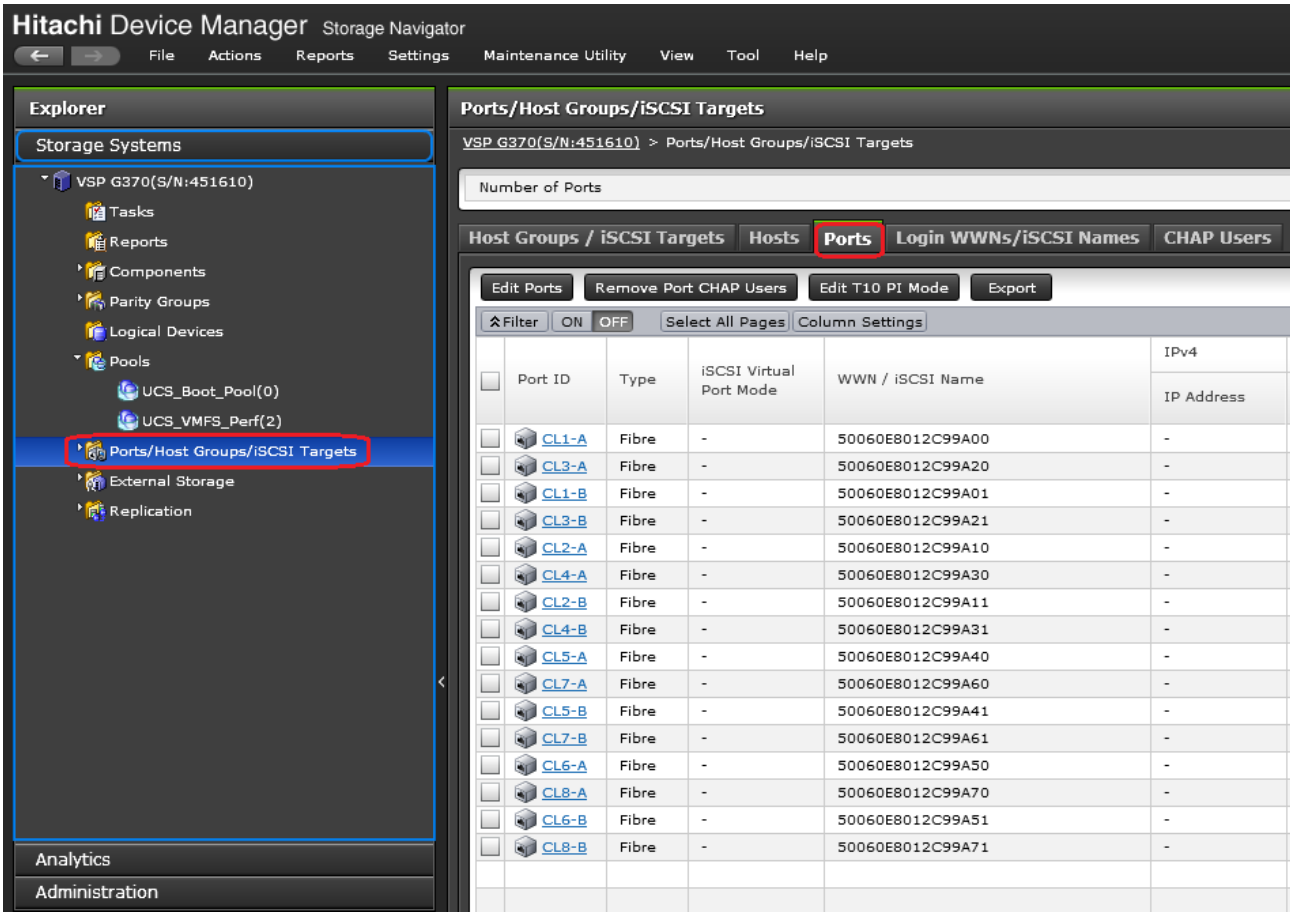

☐	Port ID	Type	iSCSI Virtual Port Mode	WWN / iSCSI Name	IPv4 IP Address
☐	CL1-A	Fibre	-	50060E8012C99A00	-
☐	CL3-A	Fibre	-	50060E8012C99A20	-
☐	CL1-B	Fibre	-	50060E8012C99A01	-
☐	CL3-B	Fibre	-	50060E8012C99A21	-
☐	CL2-A	Fibre	-	50060E8012C99A10	-
☐	CL4-A	Fibre	-	50060E8012C99A30	-
☐	CL2-B	Fibre	-	50060E8012C99A11	-
☐	CL4-B	Fibre	-	50060E8012C99A31	-
☐	CL5-A	Fibre	-	50060E8012C99A40	-
☐	CL7-A	Fibre	-	50060E8012C99A60	-
☐	CL5-B	Fibre	-	50060E8012C99A41	-
☐	CL7-B	Fibre	-	50060E8012C99A61	-
☐	CL6-A	Fibre	-	50060E8012C99A50	-
☐	CL8-A	Fibre	-	50060E8012C99A70	-
☐	CL6-B	Fibre	-	50060E8012C99A51	-
☐	CL8-B	Fibre	-	50060E8012C99A71	-

6. Select the checkboxes for the ports being used within the solution, then click the **Edit Ports** button to instantiate the Edit Ports dialog box.

7. Select checkboxes to edit the following settings to modify the selected ports:

   - Port Attribute: Target
   - Port Security: Enable
   - Port Speed: Auto
   - Fabric: ON
   - Connection Type: P-to-P

Port Attribute will only appear as an option in VSP G1500 Edit Ports dialogue.

8. Example ports used in the Cisco UCS 6454 to VSP G370 used in this design are listed in Table 9 .

Table 9 VSP G370 to UCS Ports

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G370	CL 1-A	32Gb FC	Cisco UCS 6454 FI A	FC 1/1
	CL 2-B	32Gb FC	Cisco UCS 6454 FI A	FC 1/2

Local Device	Local Port	Connection	Remote Device	Remote Port
	CL 3-B	32Gb FC	Cisco UCS 6454 FI B	FC 1/1
	CL 4-A	32Gb FC	Cisco UCS 6454 FI B	FC 1/2

Figure 3 VSP G370 Edit Ports Pop-Up Window

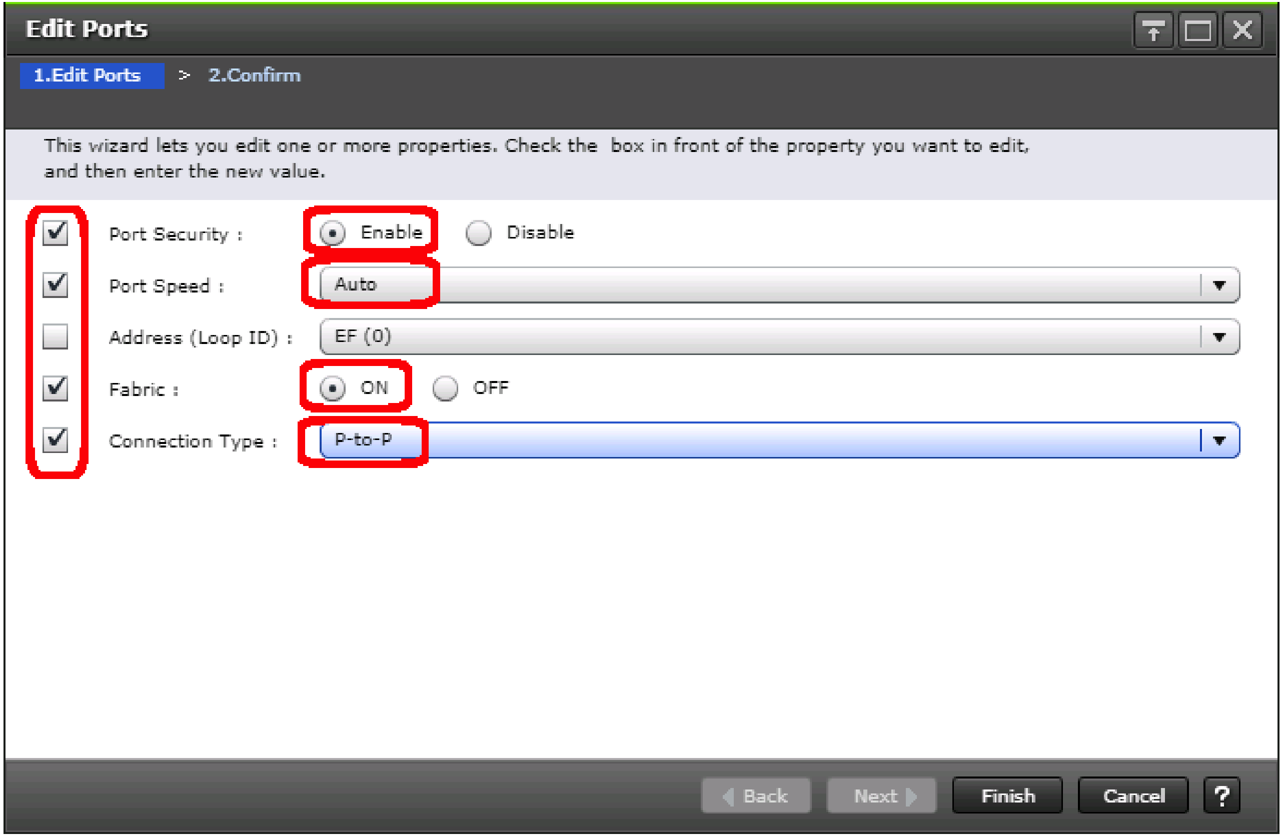
Edit Ports

1.Edit Ports > 2.Confirm

This wizard lets you edit one or more properties. Check the box in front of the property you want to edit, and then enter the new value.

☑ Port Security : ● Enable ○ Disable

☑ Port Speed : Auto

□ Address (Loop ID) : EF (0)

☑ Fabric : ● ON ○ OFF

☑ Connection Type : P-to-P

9. Click **OK** for any warning that appears.

10. Click **Finish**.

11. Review the changes to be made and check the **Go to tasks window for status** box, then click the **Apply** button.

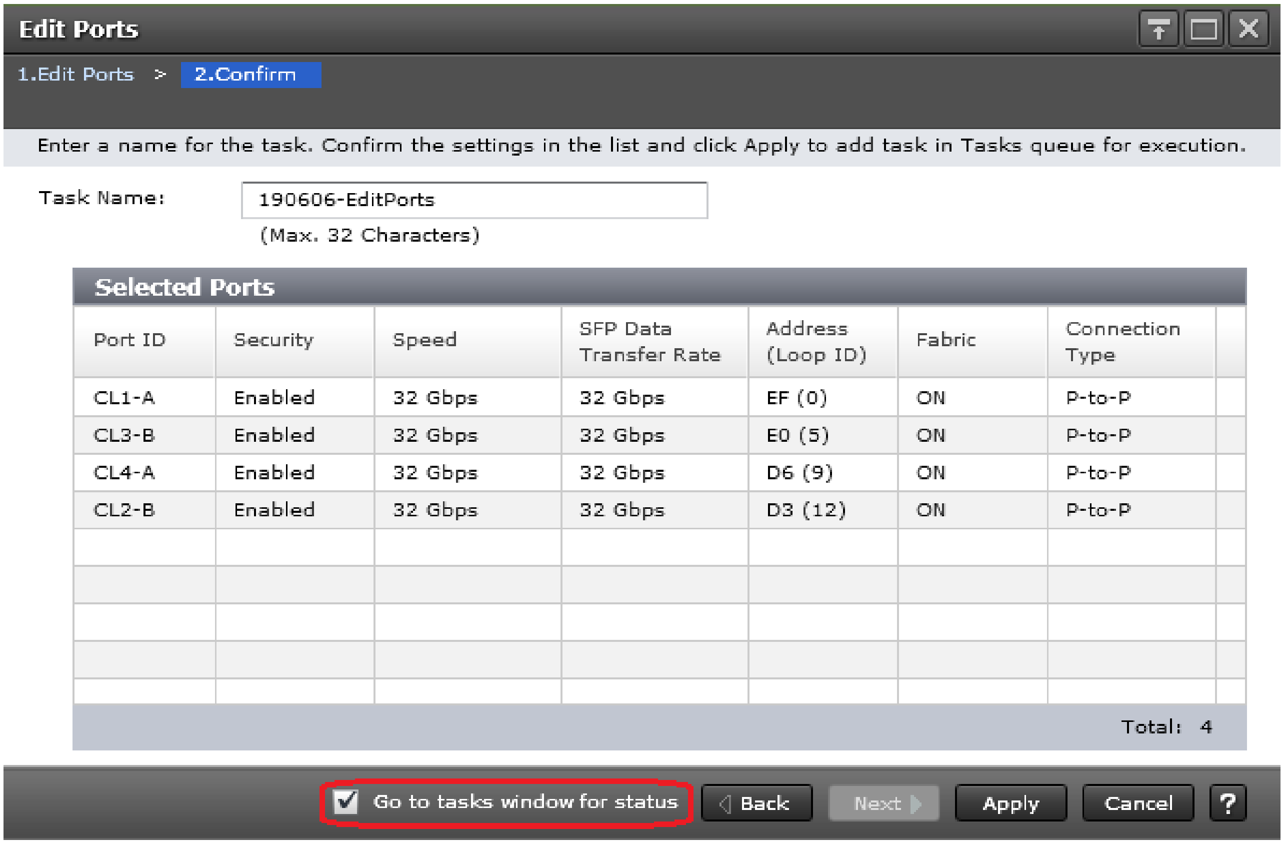

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: 190606-EditPorts

(Max. 32 Characters)

Selected Ports

Port ID	Security	Speed	SFP Data Transfer Rate	Address (Loop ID)	Fabric	Connection Type
CL1-A	Enabled	32 Gbps	32 Gbps	EF (0)	ON	P-to-P
CL3-B	Enabled	32 Gbps	32 Gbps	E0 (5)	ON	P-to-P
CL4-A	Enabled	32 Gbps	32 Gbps	D6 (9)	ON	P-to-P
CL2-B	Enabled	32 Gbps	32 Gbps	D3 (12)	ON	P-to-P

Total: 4

☑ Go to tasks window for status

12. The Task view window will appear and show the completion status of the Edit Ports task. Wait until the task status shows Complete and proceed to the next section.

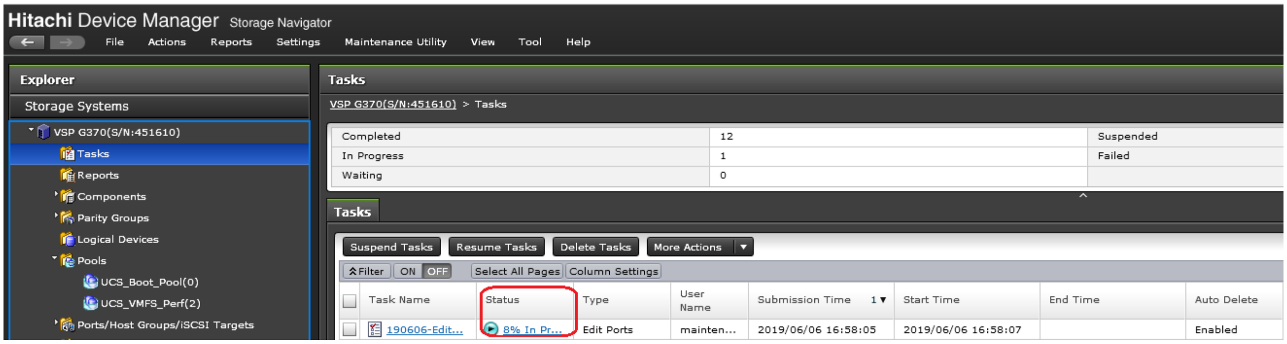

Completed	12	Suspended
In Progress	1	Failed
Waiting	0	

□	Task Name	Status	Type	User Name	Submission Time 1▼	Start Time	End Time	Auto Delete
□	190606-Edit...	8% In Pr...	Edit Ports	mainten...	2019/06/06 16:58:05	2019/06/06 16:58:07		Enabled

Cisco UCS Compute Configuration

This section explains the configuration of the Cisco UCS 6454 Fabric Interconnects used in this UCP solution. As with the Nexus explained beforehand, some changes may be appropriate for a customer's environment, but care should be taken when stepping outside of these instructions as it may lead to an improper configuration.

Physical Connectivity

Physical cabling should be completed by following the diagram and table references in section Deployment Hardware and Software.

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

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

Cisco UCS Base Configuration

The initial configuration dialogue for the Cisco UCS 6454 Fabric Interconnects will be provide the primary information to the first fabric interconnect, with the second taking on most settings after joining the cluster.

To start on the configuration of the Fabric Interconnect A, connect to the console of the fabric interconnect and step through the Basic System Configuration Dialogue:

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

This setup utility will guide you through the basic configuration of
the system. Only minimal configuration including IP connectivity to
the Fabric interconnect and its clustering mode is performed through these steps.

Type Ctrl-C at any time to abort configuration and reboot system.
To back track or make modifications to already entered values,
complete input till end of section and answer no when prompted
to apply configuration.


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": <<var_password>>
Confirm the password for "admin": <<var_password>>

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: y

Enter the switch fabric (A/B) []: A

Enter the system name:  <<var_ucs_6454_clustername>>

Physical Switch Mgmt0 IP address : <<var_ucsa_mgmt_ip>>

Physical Switch Mgmt0 IPv4 netmask : <<var_oob_mgmt_mask>>

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

Cluster IPv4 address : <<var_ucs_mgmt_vip>>
```

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

    DNS IP address : <<var_nameserver_ip>>

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

    Default domain name : <<var_dns_domain_name>>

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

  Following configurations will be applied:

    Switch Fabric=A
    System Name=AA19-6454
    Enforced Strong Password=yes
    Physical Switch Mgmt0 IP Address=192.168.168.16
    Physical Switch Mgmt0 IP Netmask=255.255.255.0
    Default Gateway=192.168.168.254
    Ipv6 value=0
    DNS Server=10.1.168.9
    Domain Name=ucp.cisco.com

    Cluster Enabled=yes
    Cluster IP Address=192.168.168.15
    NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized.
          UCSM will be functional only after peer FI is configured in clustering mode.

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

  Configuration file - Ok
```

Wait for the appearance of a login prompt on UCS FI A before proceeding to B.

Continue the configuration on the console of the Fabric Interconnect B:

```
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:
    Connecting to peer Fabric interconnect... done
    Retrieving config from peer Fabric interconnect... done
    Peer Fabric interconnect Mgmt0 IPv4 Address: 192.168.168.16
    Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
    Cluster IPv4 address          : 192.168.168.15

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

  Physical Switch Mgmt0 IP address : 192.168.164.17
  Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
  Applying configuration. Please wait.
```

Cisco UCS Manager Setup

Log into Cisco UCS Manager

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

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

2. Click the Launch UCS Manager link within the opening page.

3. If prompted to accept security certificates, accept as necessary.

4. When the Cisco UCS Manager login is prompted, enter `admin` as the user name and enter the administrative password.

5. Click Login to log into Cisco UCS Manager.

Anonymous Reporting

During the first connection to the Cisco UCS Manager GUI, a pop-up window will appear to allow for the configuration of Anonymous Reporting to Cisco on use to help with future development. To create anonymous reporting, follow this step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products, and provide the appropriate SMTP server gateway information if configuring:

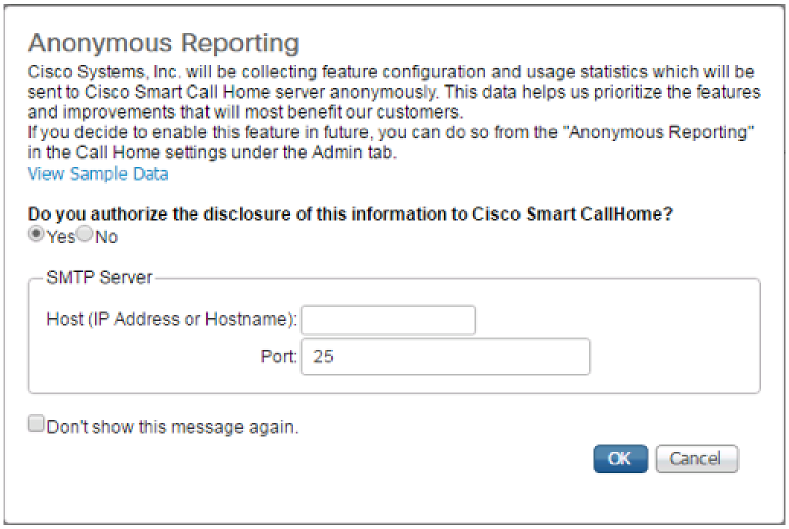

Anonymous Reporting

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: 25

□ Don't show this message again.

If you want to enable or disable Anonymous Reporting at a later date, it can be found within Cisco UCS Manager under: **Admin** -> **Communication Management** -> **Call Home**, which has a tab on the far right for **Anonymous Reporting**.

Place Cisco UCS Fabric Interconnects in Fiber Channel Switching Mode

In order to use Fiber Channel Storage Ports for storage directly connected to the Cisco UCS fabric interconnects, the fabric interconnects must be changed from fiber channel end host mode to fiber channel switching mode.

To place the fabric interconnects in fiber channel switching mode, follow these steps:

1. In Cisco UCS Manager, click Equipment.

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

3. In the center pane, select set FC Switching Mode. Click Yes and OK for the confirmation message.

The standby FI will reboot, after completion the primary will be prompted for a reboot as a pending activity.

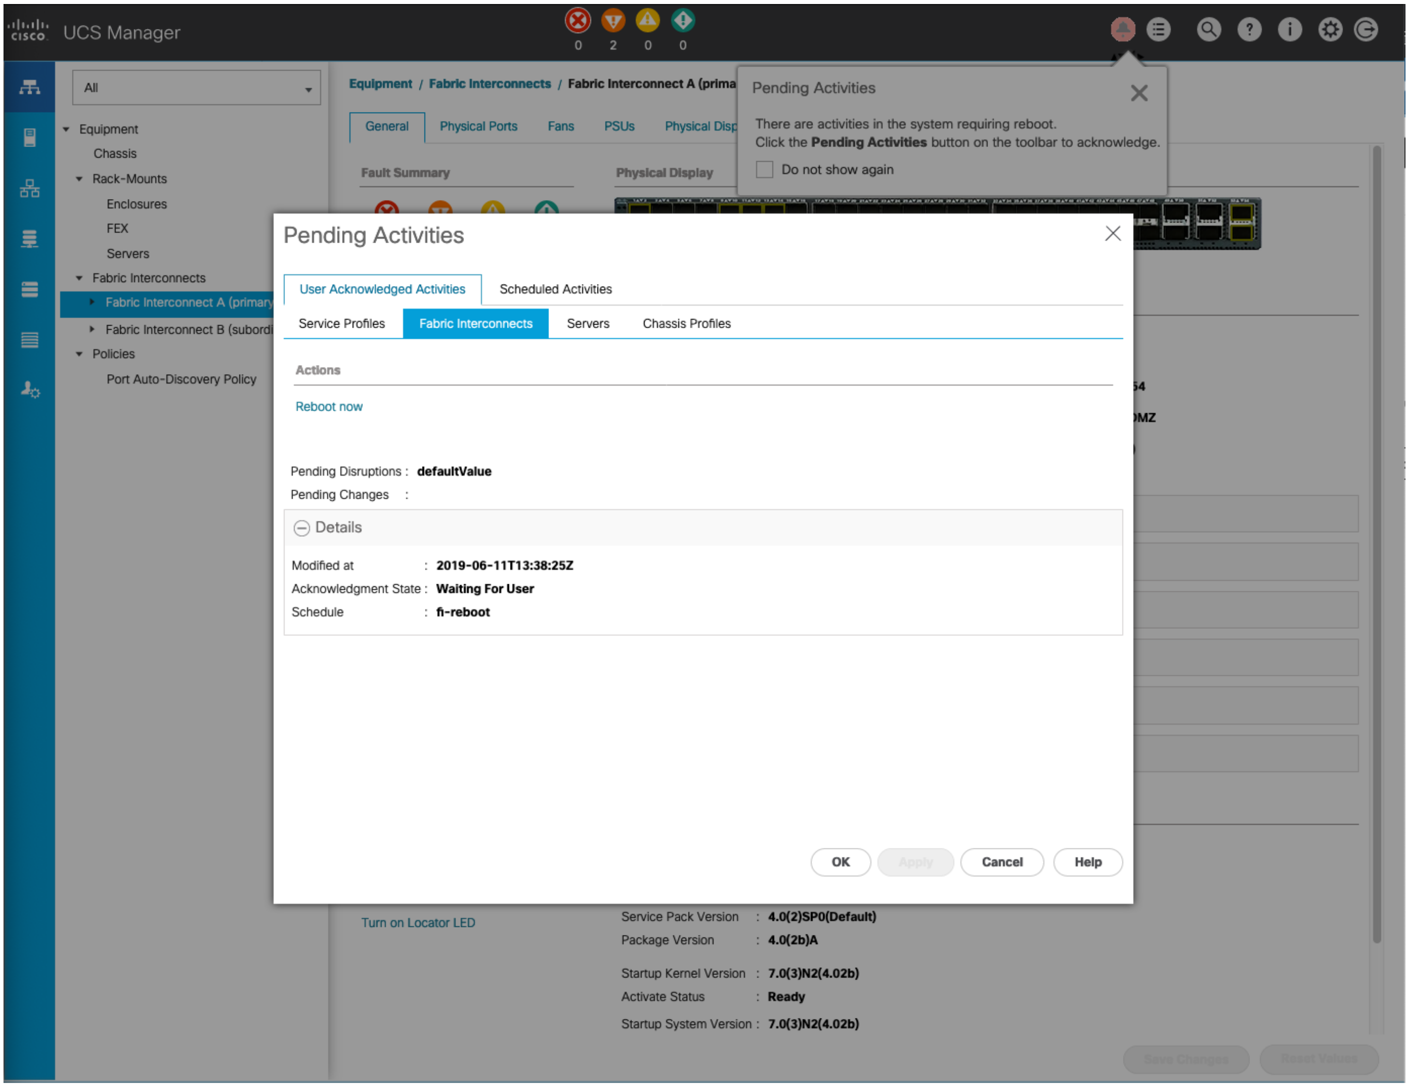

Pending Activities

There are activities in the system requiring reboot.
Click the **Pending Activities** button on the toolbar to acknowledge.

□ Do not show again

Pending Disruptions : **defaultValue**
Pending Changes :

Modified at : **2019-06-11T13:38:25Z**
Acknowledgment State : **Waiting For User**
Schedule : **fi-reboot**

Service Pack Version : **4.0(2)SP0(Default)**
Package Version : **4.0(2b)A**

Startup Kernel Version : **7.0(3)N2(4.02b)**
Activate Status : **Ready**
Startup System Version : **7.0(3)N2(4.02b)**

4. Select Reboot now under Actions. Click Yes at the confirmation prompt.

5. Wait for both Fabric Interconnects to reboot by monitoring the console ports and log back into Cisco UCS Manager.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, follow these steps:

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

2. Select Timezone Management drop-down list and click Timezone.

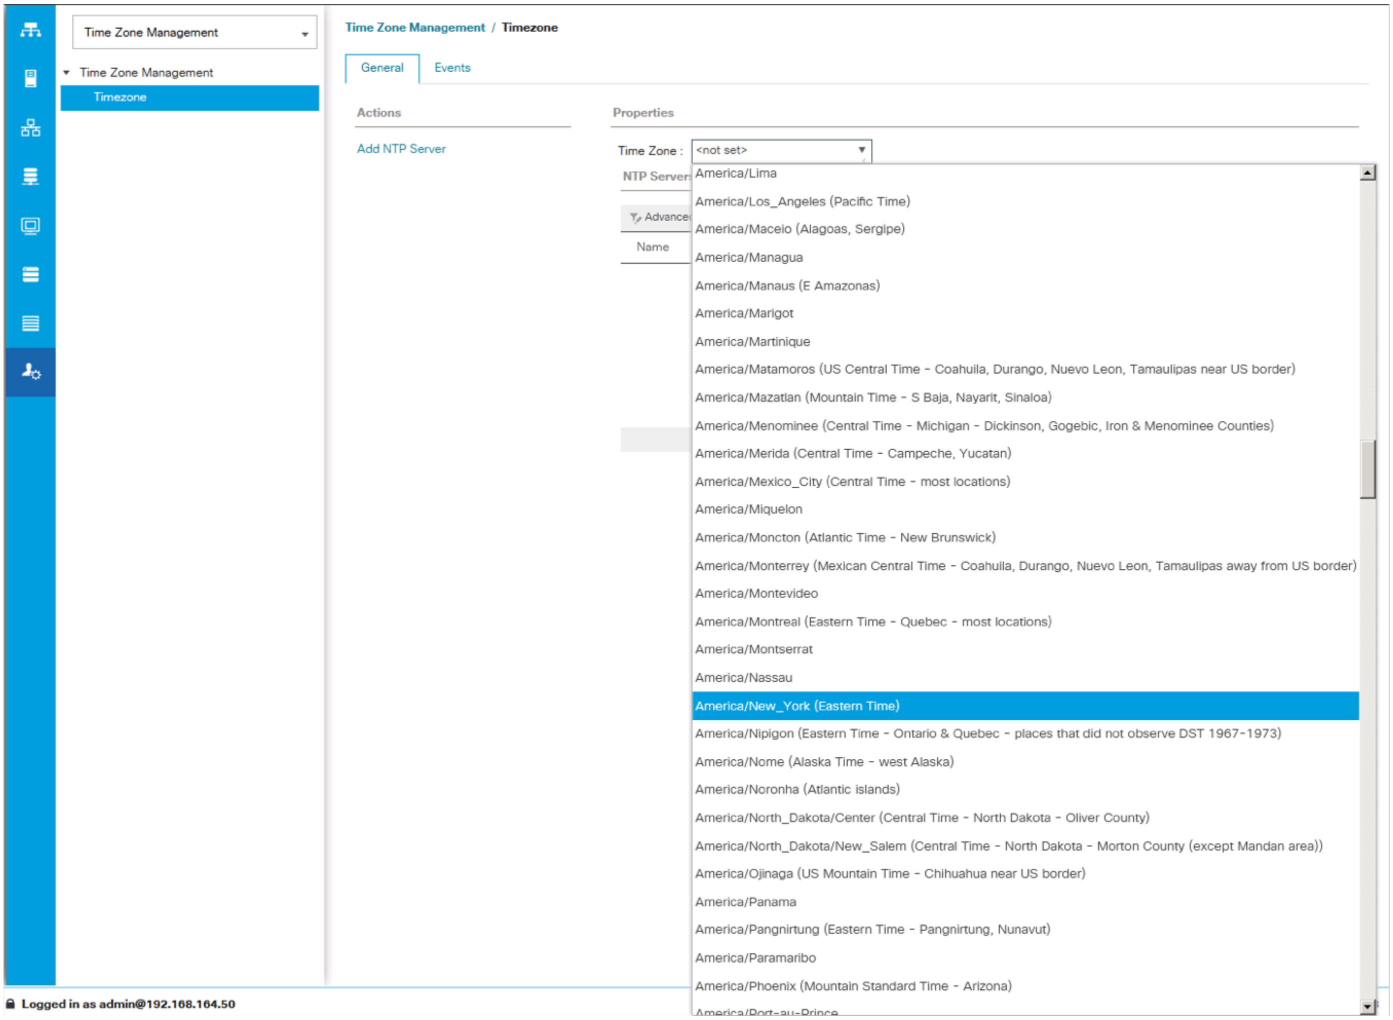
Time Zone Management
Time Zone Management
Timezone
Time Zone Management / Timezone
General
Events
Actions
Add NTP Server
Properties
NTP Server
Advanced
Name
America/Lima
America/Los_Angeles (Pacific Time)
America/Maceio (Alagoas, Sergipe)
America/Managua
America/Manaus (E Amazonas)
America/Marigot
America/Martinique
America/Matamoros (US Central Time - Coahuila, Durango, Nuevo Leon, Tamaulipas near US border)
America/Mazatlan (Mountain Time - S Baja, Nayarit, Sinaloa)
America/Menominee (Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee Counties)
America/Merida (Central Time - Campeche, Yucatan)
America/Mexico_City (Central Time - most locations)
America/Miquelon
America/Moncton (Atlantic Time - New Brunswick)
America/Monterrey (Mexican Central Time - Coahuila, Durango, Nuevo Leon, Tamaulipas away from US border)
America/Montevideo
America/Montreal (Eastern Time - Quebec - most locations)
America/Montserrat
America/Nassau
America/New_York (Eastern Time)
America/Nipigon (Eastern Time - Ontario & Quebec - places that did not observe DST 1967-1973)
America/Nome (Alaska Time - west Alaska)
America/Noronha (Atlantic islands)
America/North_Dakota/Center (Central Time - North Dakota - Oliver County)
America/North_Dakota/New_Salem (Central Time - North Dakota - Morton County (except Mandan area))
America/Ojinaga (US Mountain Time - Chihuahua near US border)
America/Panama
America/Pangnirtung (Eastern Time - Pangnirtung, Nunavut)
America/Paramaribo
America/Phoenix (Mountain Standard Time - Arizona)
America/Port-au-Prince

Time Zone : <not set>

Logged in as admin@192.168.164.50

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

4. Click Save Changes, and then click OK.

5. Click Add NTP Server.

6. Enter <<var_oob_ntp>> and click OK.

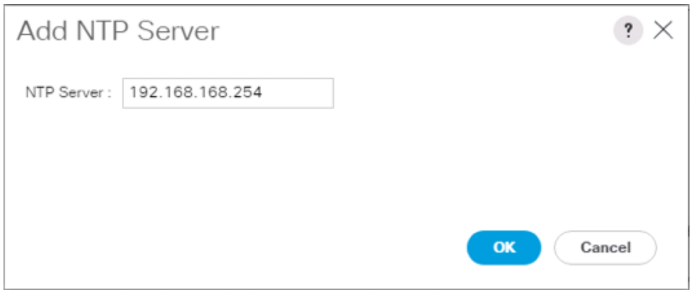

NTP Server : 192.168.168.254

7. Click OK.

Configure Cisco UCS Servers

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis. To modify the chassis discovery policy, follow these steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and select Policies in the list from the drop-down list.

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

3. Set the Link Grouping Preference to Port Channel.

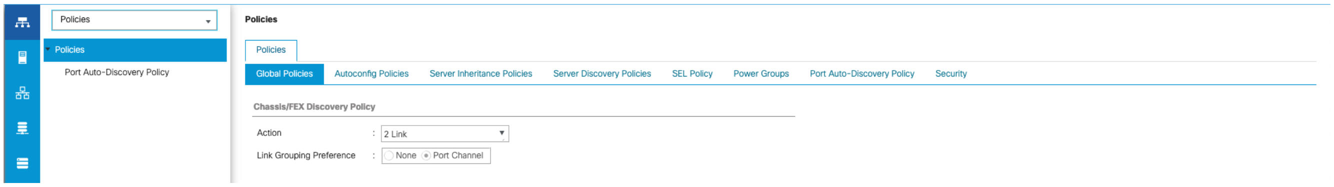
Policies
Policies
Port Auto-Discovery Policy
Policies
Policies
Global Policies
Autoconfig Policies
Server Inheritance Policies
Server Discovery Policies
SEL Policy
Power Groups
Port Auto-Discovery Policy
Security

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : ○ None ● Port Channel

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

4. Leave other settings alone or change if appropriate to your environment.

5. Click Save Changes.

6. Click OK.

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.

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

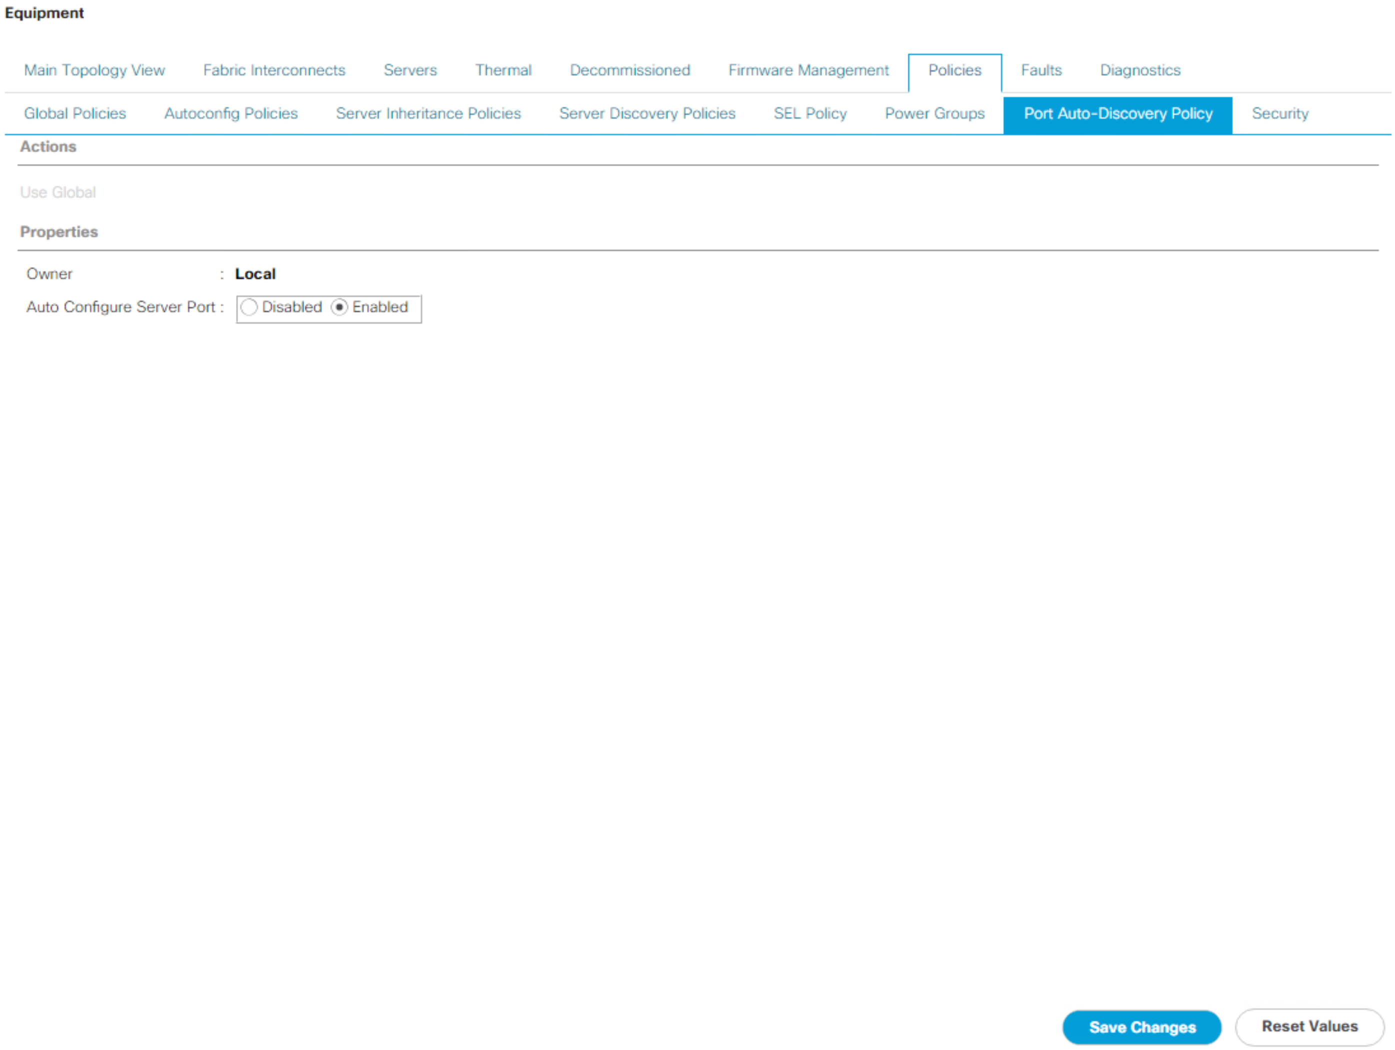

Owner : **Local**

Auto Configure Server Port : ○ Disabled ● Enabled

3. Click Save Changes and then click 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.

2. Within the Global Policies sub-tab, scroll down to Info Policy and select Enabled for Action.

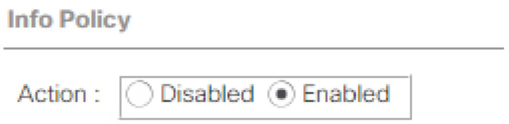

Action : ○ Disabled ● Enabled

3. Click Save Changes and then click OK.

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

Configure Unified Ports

The Cisco UCS 6454 Fabric Interconnects will have a slider mechanism within the Cisco UCS Manager GUI interface that will control the first 8 ports starting from the first port, allowing the selection of the first 4, or all 8 of the unified ports.

To enable the fibre channel ports, follow these steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary)
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 the first 4 or all 8 of the ports to be set as FC Uplinks.

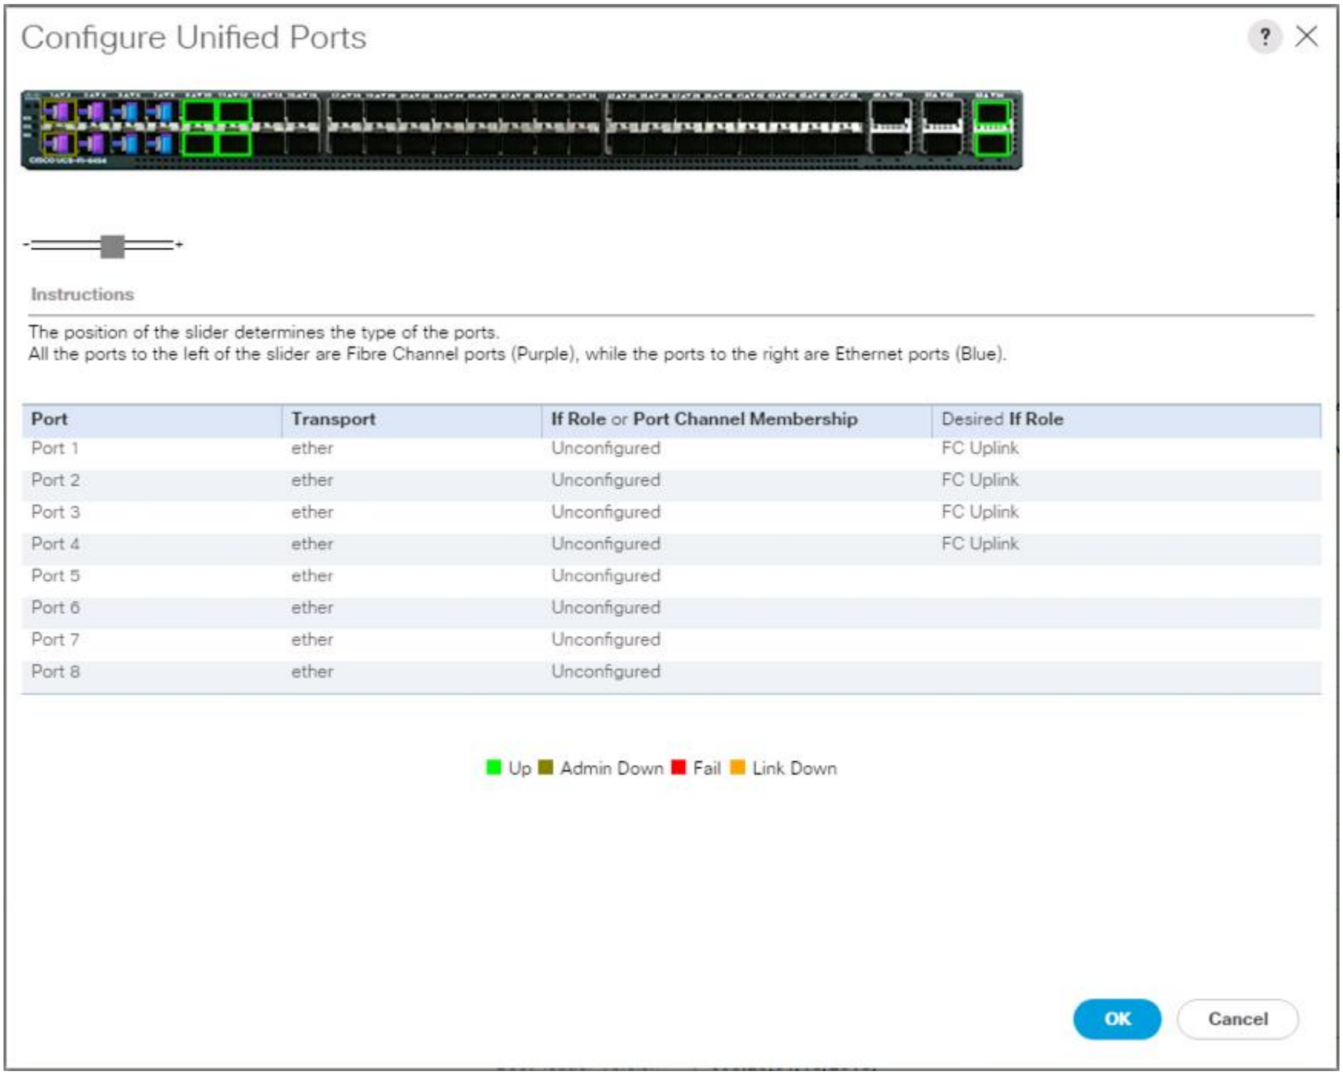

Configure Unified Ports

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 to continue

7. Click Yes within the subsequent warning pop-up and wait for reboot to complete.

8. Log back into UCSM when available.

9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (primary).

10. Select Configure Unified Ports.

11. 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.

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

13. Click OK to continue.

14. Click Yes within the subsequent warning pop-up and wait for reboot to complete.

Enable Server and Uplink Ports

To enable server and uplink ports, follow these steps:

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

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

3. Expand FC Ports

4. Select the ports that are connected to the VSP, right-click them, and select "Configure as FC Storage Port."

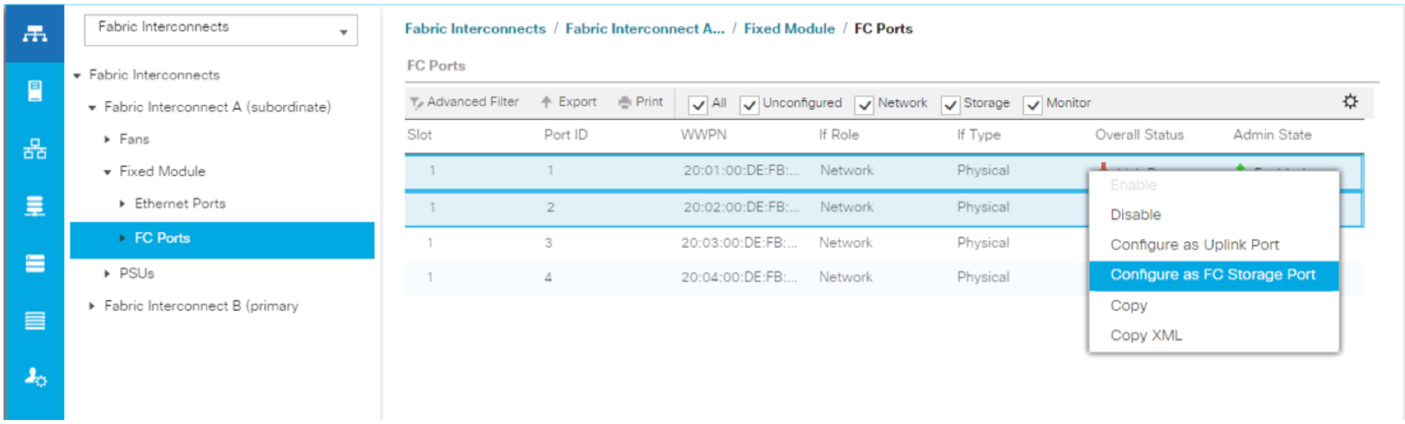

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
1	1	20:01:00:DE:FB:...	Network	Physical		
1	2	20:02:00:DE:FB:...	Network	Physical		
1	3	20:03:00:DE:FB:...	Network	Physical		
1	4	20:04:00:DE:FB:...	Network	Physical		

5. Select ports 53 and 54 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

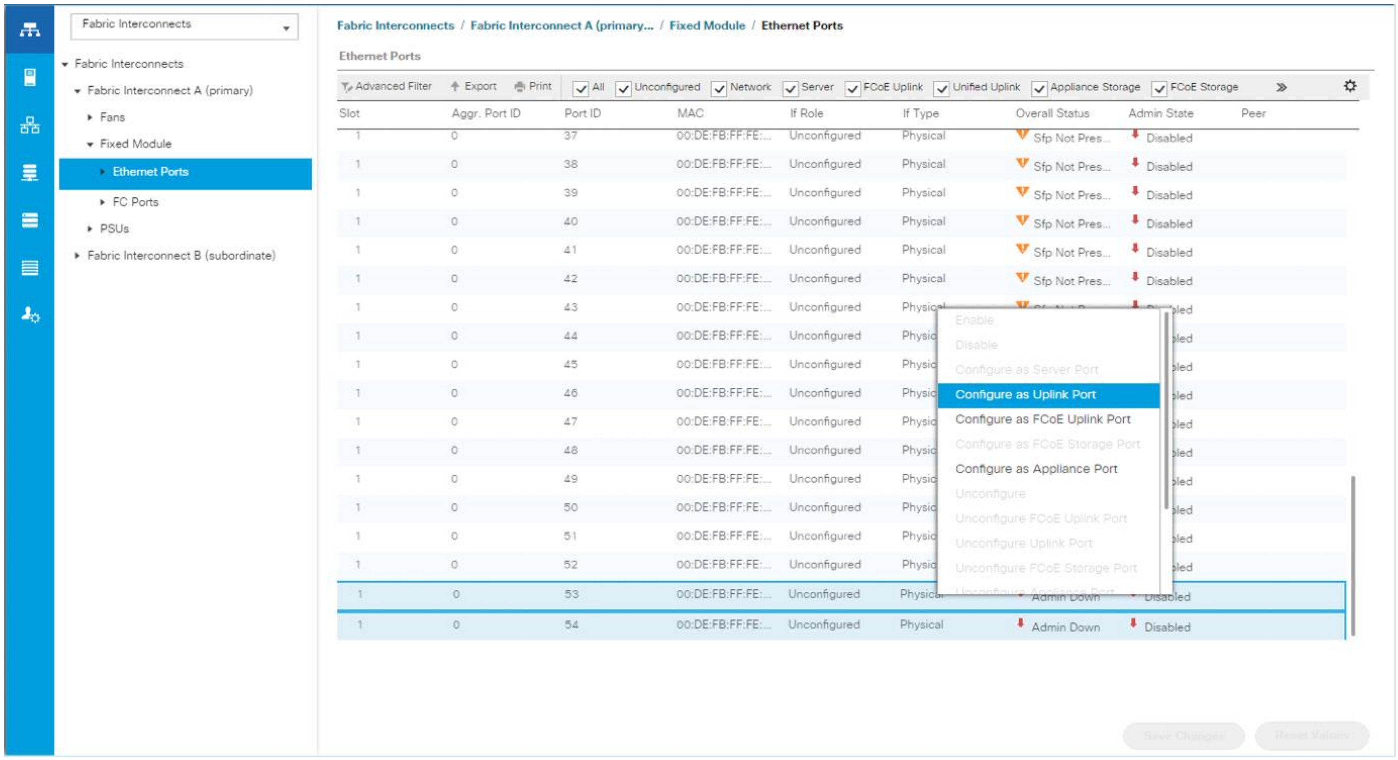

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	37	00:DE:FB:FF:FE:...	Unconfigured	Physical	Sfp Not Pres...	Disabled	
1	0	38	00:DE:FB:FF:FE:...	Unconfigured	Physical	Sfp Not Pres...	Disabled	
1	0	39	00:DE:FB:FF:FE:...	Unconfigured	Physical	Sfp Not Pres...	Disabled	
1	0	40	00:DE:FB:FF:FE:...	Unconfigured	Physical	Sfp Not Pres...	Disabled	
1	0	41	00:DE:FB:FF:FE:...	Unconfigured	Physical	Sfp Not Pres...	Disabled	
1	0	42	00:DE:FB:FF:FE:...	Unconfigured	Physical	Sfp Not Pres...	Disabled	
1	0	43	00:DE:FB:FF:FE:...	Unconfigured	Physical		led	
1	0	44	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	45	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	46	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	47	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	48	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	49	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	50	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	51	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	52	00:DE:FB:FF:FE:...	Unconfigured	Physic		led	
1	0	53	00:DE:FB:FF:FE:...	Unconfigured	Physical	Admin Down	Disabled	
1	0	54	00:DE:FB:FF:FE:...	Unconfigured	Physical	Admin Down	Disabled	

6. Click Yes to confirm uplink ports and click OK.

7. Repeat steps 1-6 for Fabric B ports.

Create Pools

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 the LAN tab in the navigation pane.

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.

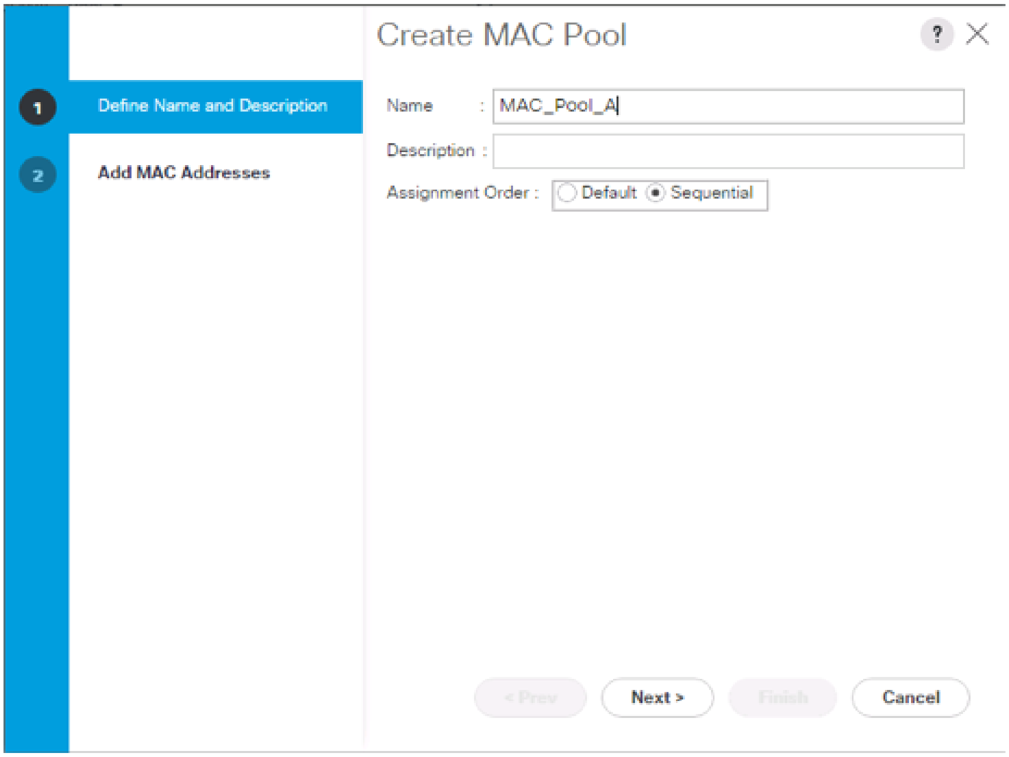

Name : MAC_Pool_A

Description :

Assignment Order : ○ Default ● Sequential

8. Click Next.

9. Click Add.

10. Specify a starting MAC address.

For Cisco UCS deployments, the recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the of information of also embedding and FI number reference of 54(for UCS 6454) giving us `00:25:B5:54:0A: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.

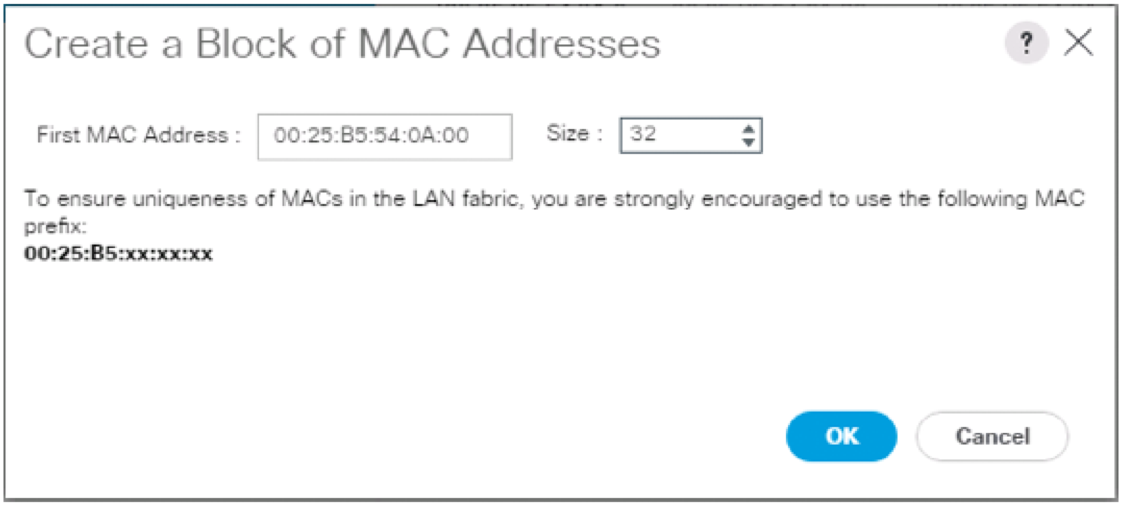

First MAC Address : 00:25:B5:54:0A:00

Size : 32

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.

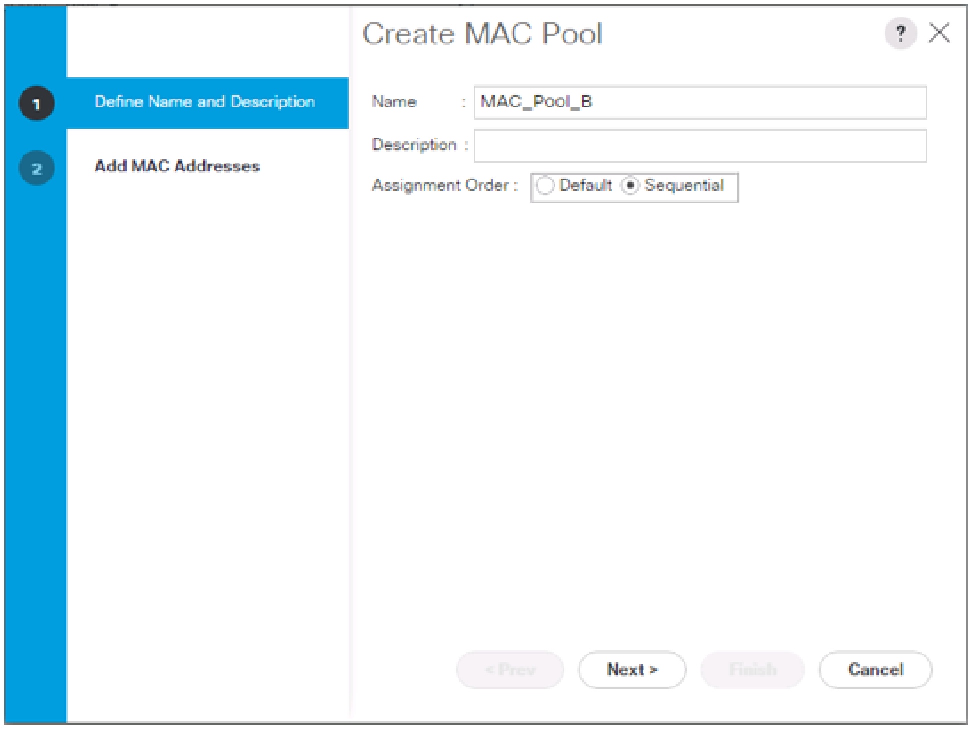

Name : MAC_Pool_B

Description :

Assignment Order : ○ Default ● Sequential

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

For Cisco UCS deployments, it is recommended to place `0B` in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward the of information of also embedding and FI number reference of 54 giving us `00:25:B5:54:0A:00` as our first MAC address.

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

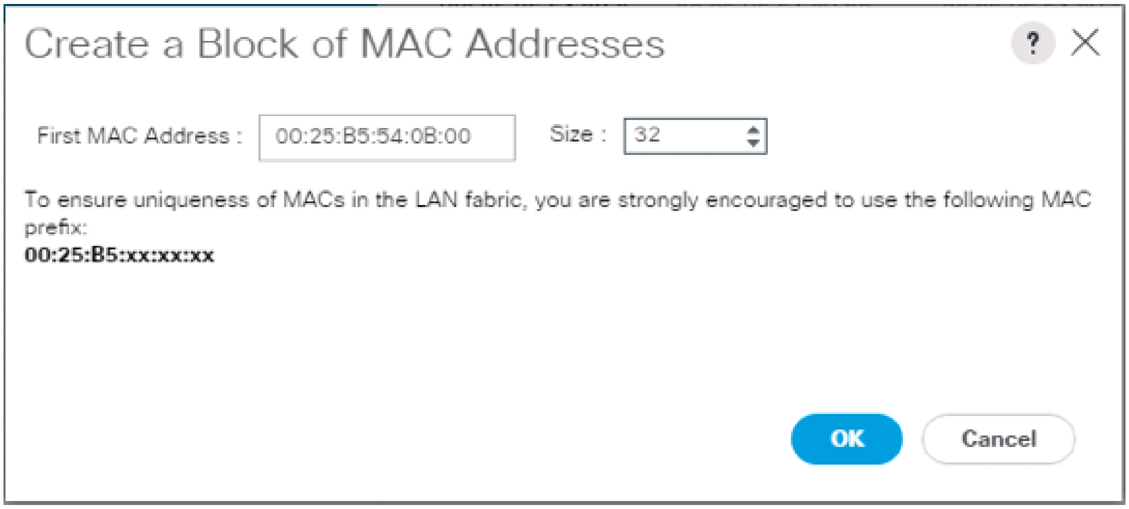

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:54:0B:00

Size : 32

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

23. Click OK.

24. Click Finish.

25. 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 the Servers tab in the navigation pane.

2. Select Pools > root.

3. Right-click UUID Suffix Pools.

4. Select Create UUID Suffix Pool.

5. Enter UUID_Pool as the name of the UUID suffix pool.

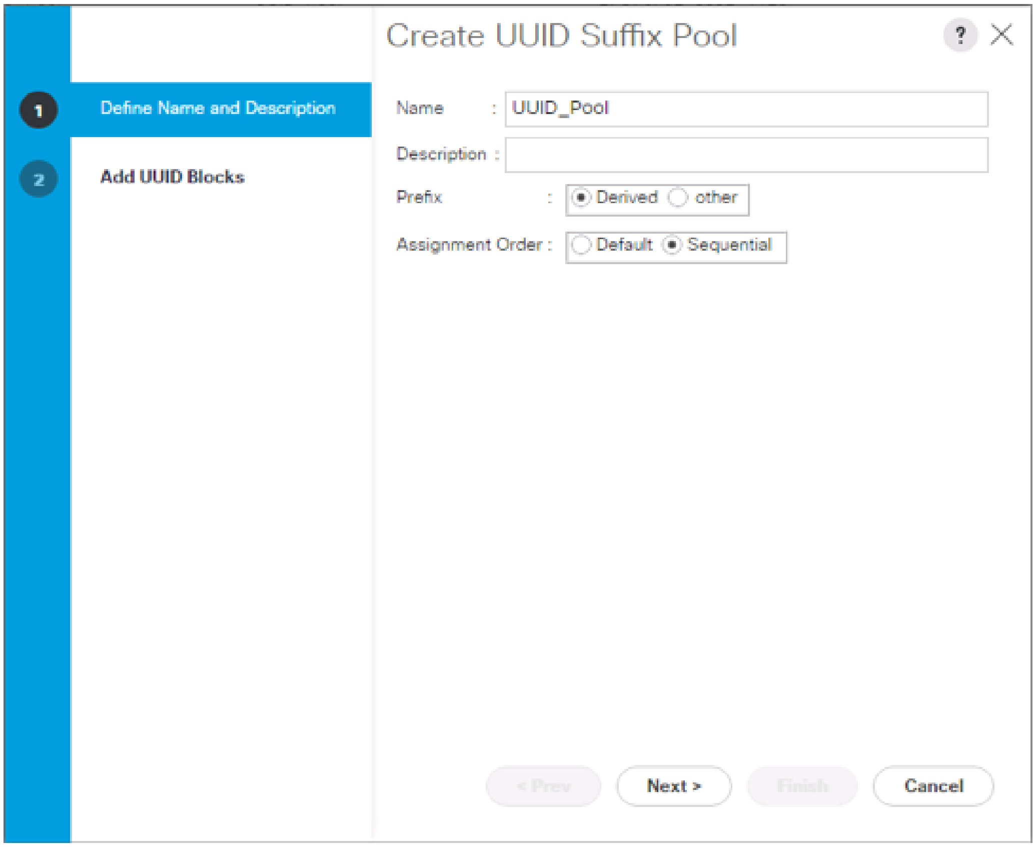

Create UUID Suffix Pool

Name : UUID_Pool

Description :

Prefix : ● Derived ○ other

Assignment Order : ○ Default ● Sequential

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.

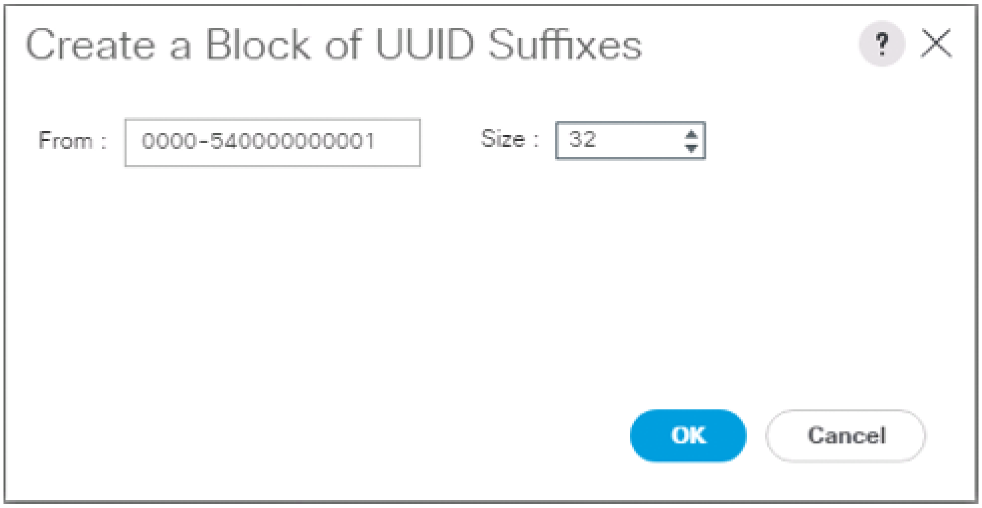

Create a Block of UUID Suffixes

From : 0000-540000000001

Size : 32

The starting From number (0000-54) has been adjusted to give it a differentiator from other UCS domains that may be adjacent.

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

12. Click OK.

13. Click Finish.

14. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, follow these steps:

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

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

2. Select Pools > root.

3. Right-click Server Pools.

4. Select Create Server Pool.

5. Enter Infra_Pool as the name of the server pool.

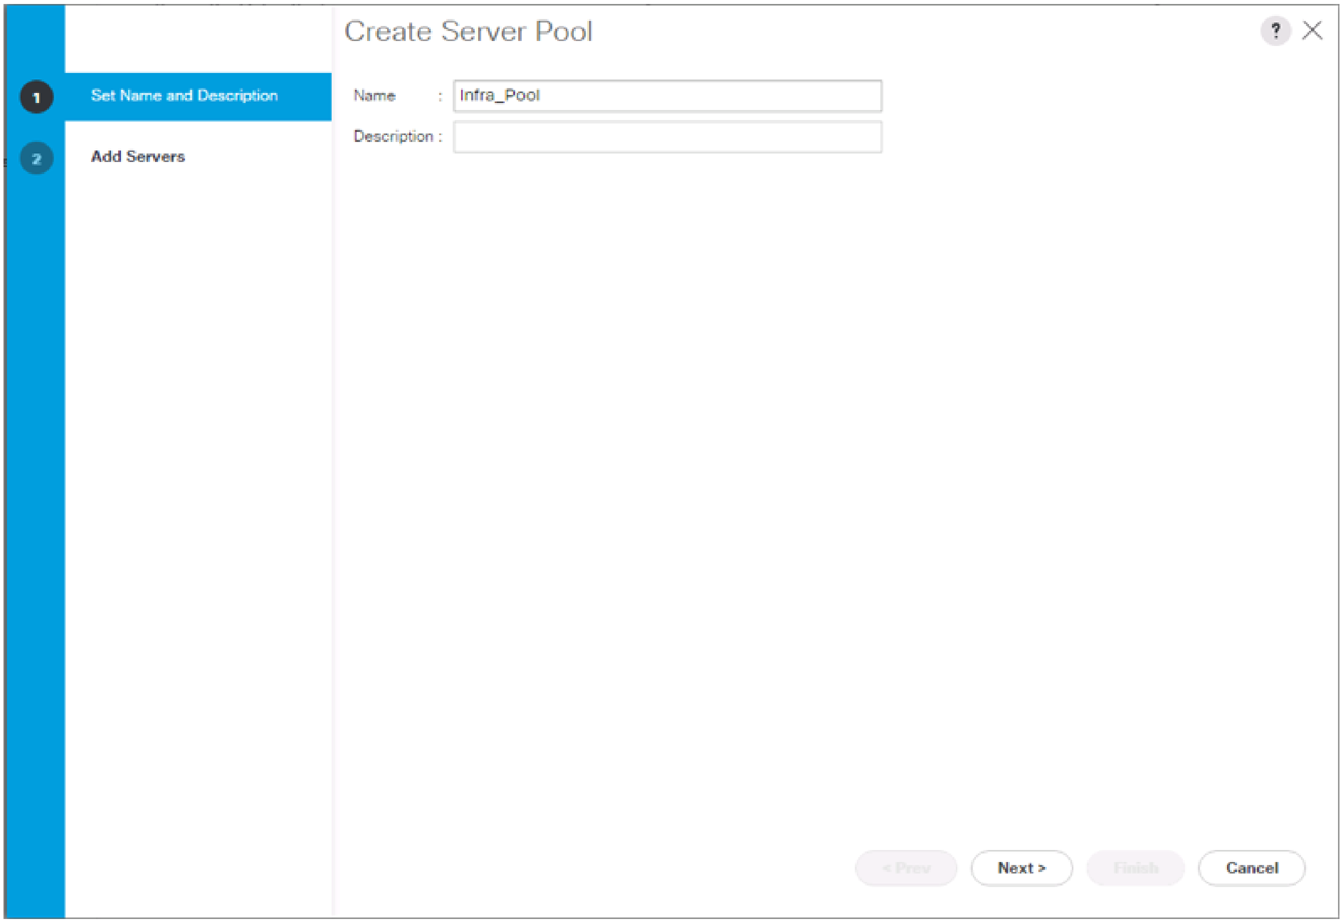

Name : Infra_Pool

Description :

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

7. Click Next.

8. Select two (or more) servers to be used for the VMware cluster and click >> to add them to the Infra_Pool server pool.

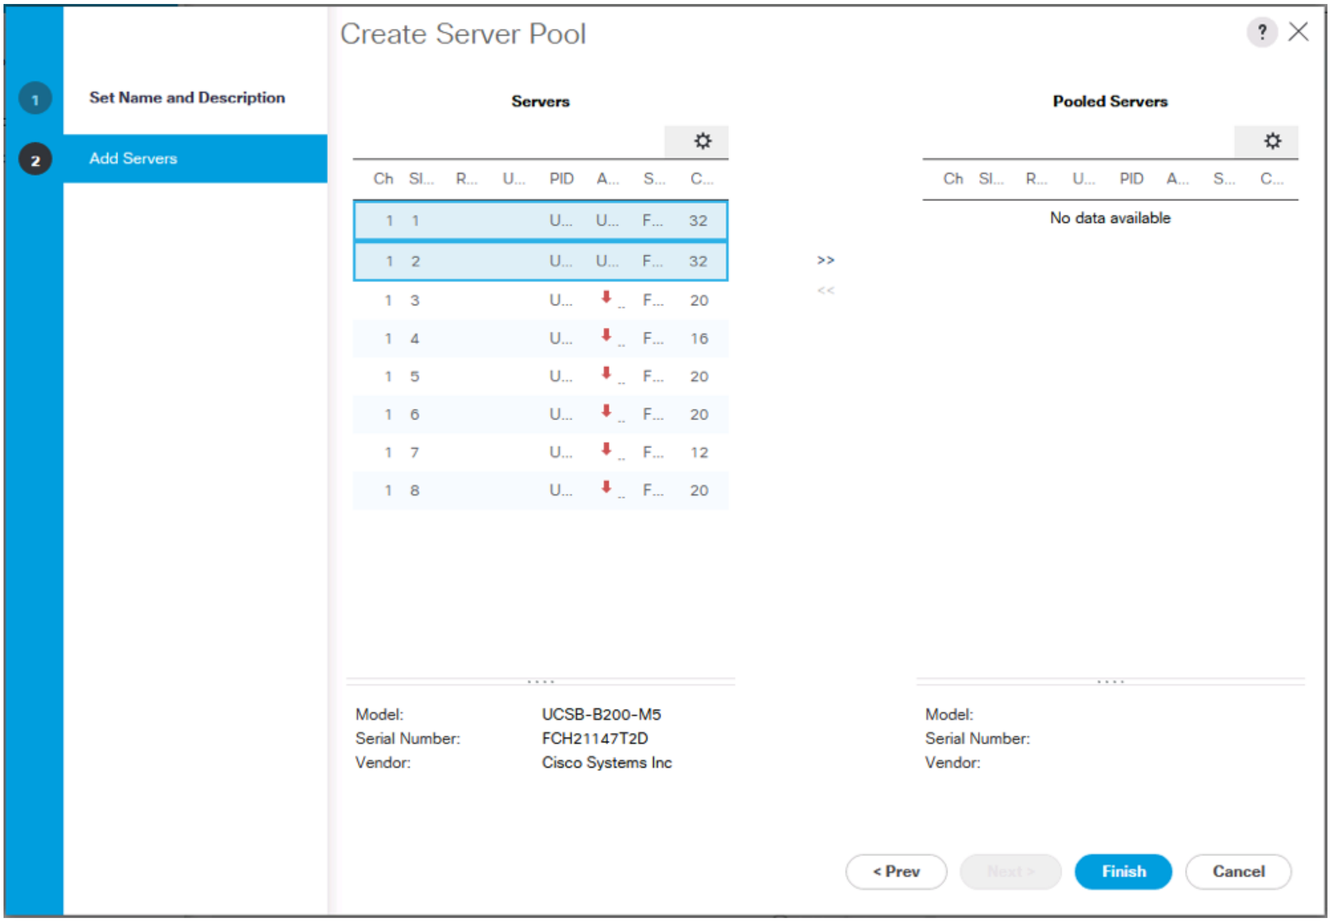

Create Server Pool

Servers

Ch	Sl...	R...	U...	PID	A...	S...	C...
1	1			U...	U...	F...	32
1	2			U...	U...	F...	32
1	3			U...	⬇ ..	F...	20
1	4			U...	⬇ ..	F...	16
1	5			U...	⬇ ..	F...	20
1	6			U...	⬇ ..	F...	20
1	7			U...	⬇ ..	F...	12
1	8			U...	⬇ ..	F...	20

Model: UCSB-B200-M5
Serial Number: FCH21147T2D
Vendor: Cisco Systems Inc

Pooled Servers

Ch	Sl...	R...	U...	PID	A...	S...	C...

No data available

Model:
Serial Number:
Vendor:

9. Click Finish.

10. Click OK.

Add a 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 the LAN tab in the navigation pane.

2. Select Pools > root > IP Pools.

3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.

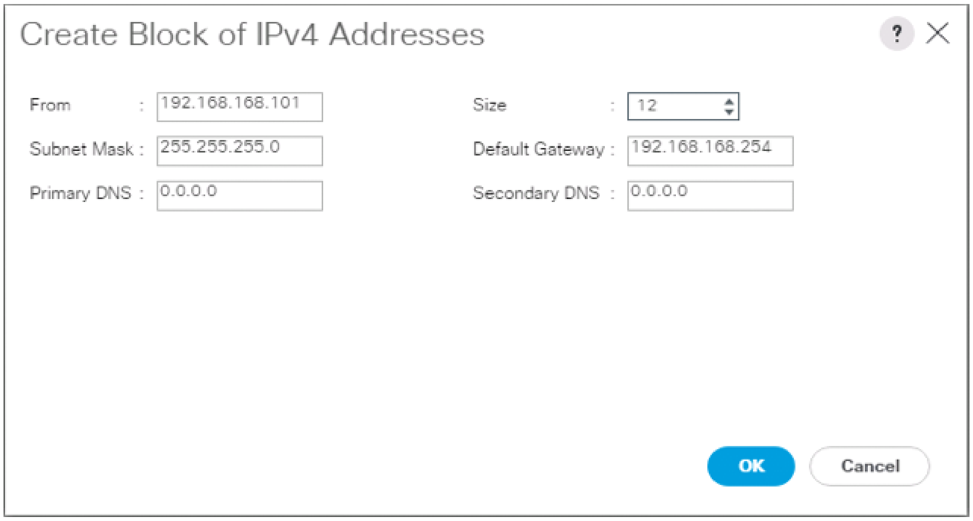

Create Block of IPv4 Addresses

From :	192.168.168.101	Size :	12
Subnet Mask :	255.255.255.0	Default Gateway :	192.168.168.254
Primary DNS :	0.0.0.0	Secondary DNS :	0.0.0.0

4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

5. Click OK to create the block of IPs.

6. Click OK.

Create a WWNN Pool

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

1. Select the SAN tab.

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.

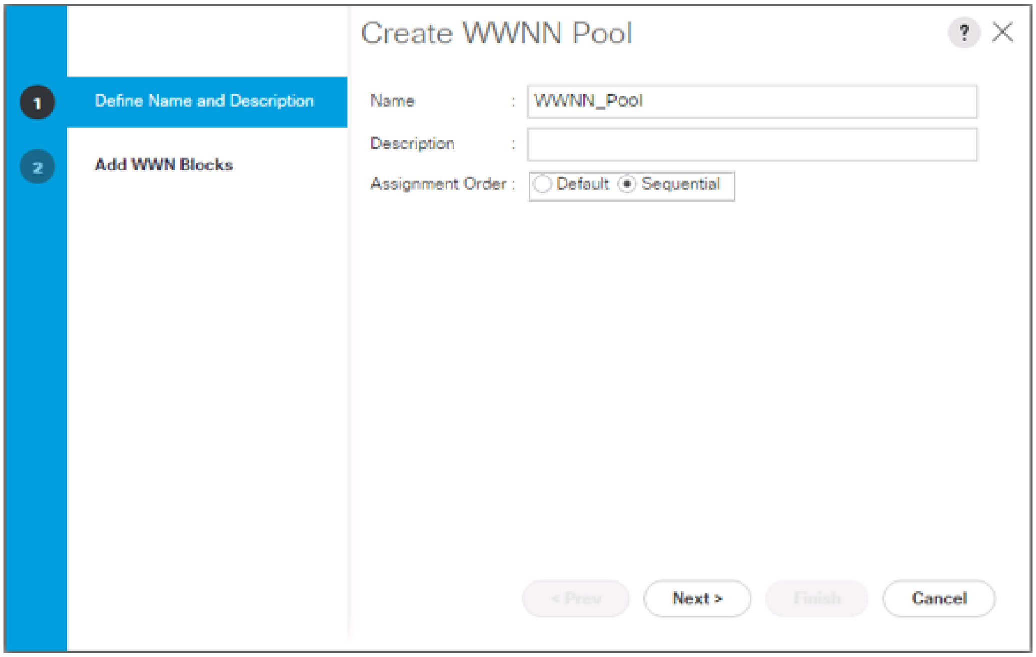

Create WWNN Pool

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 WWN 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 6th octet was changed from 00 to 54 to represent as identifying information for the 6454 Cisco UCS domain.

When you have multiple Cisco 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.

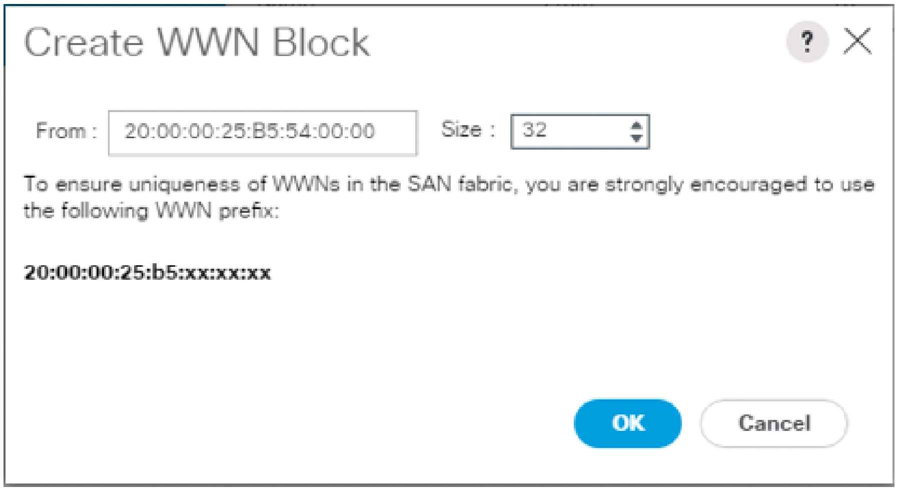

Create WWN Block

From : 20:00:00:25:B5:54:00:00 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 to create the WWNN Pool.

14. Click OK.

Create WWPN Pools

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

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

2. Select Pools > root.

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

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

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

6. Enter WWPN_Pool_A as the name of the WWPN pool.

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

8. Select Sequential for Assignment Order.

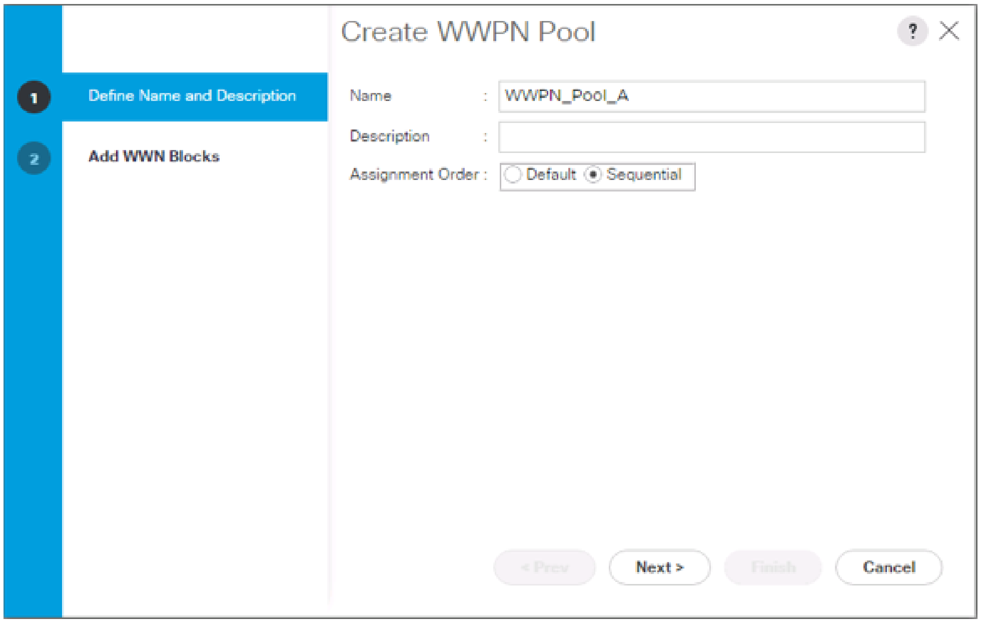

Name : WWPN_Pool_A

Description :

Assignment Order : ○ Default ● Sequential

9. Click Next.

10. Click Add.

11. Specify a starting WWPN.

For the solution, the recommendation is to place `0A` 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:B5:54:0A:00`.

---

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

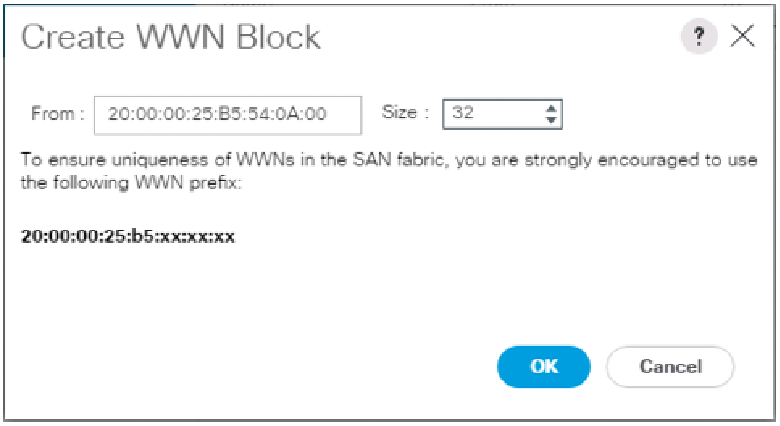

From : 20:00:00:25:B5:54:0A:00 Size : 32

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

13. Click OK.

14. Click Finish.

15. In the confirmation message, click OK.

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

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

18. Enter WWPN_Pool_B as the name of the WWPN pool.

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

20. Select Sequential for Assignment Order.

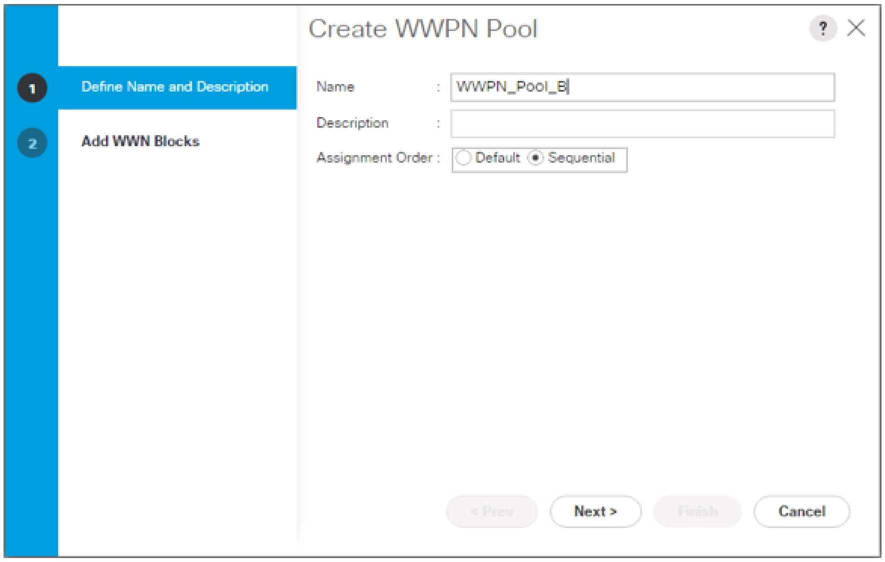

Create WWPN Pool

Name : WWPN_Pool_B

Description :

Assignment Order : ○ Default ● Sequential

21. Click Next.

22. Click Add.

23. Specify a starting WWPN.

For the solution, the recommendation is to place `0B` 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:B5:54:0B:00`.

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

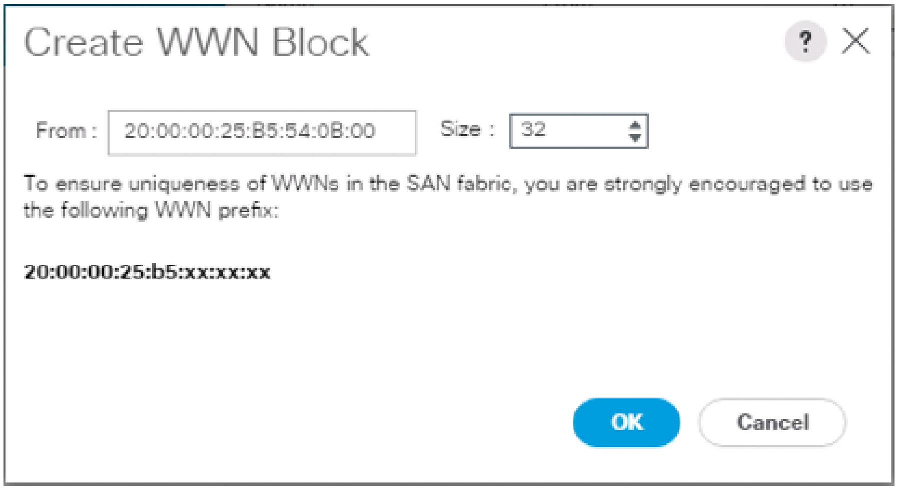

Create WWN Block

From : 20:00:00:25:B5:54:0B: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

25. Click OK.

26. Click Finish.

27. In the confirmation message, click OK.

Set Packages and Policies

Create Host Firmware Package

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

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

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

2. Select Policies > root.

3. Expand Host Firmware Packages.

4. Select default.

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

6. Select the version 4.0(2b)B for the Blade Package, and optionally set version 4.0(2b)C for the Rack Package.

7. Leave Excluded Components with only Local Disk selected.

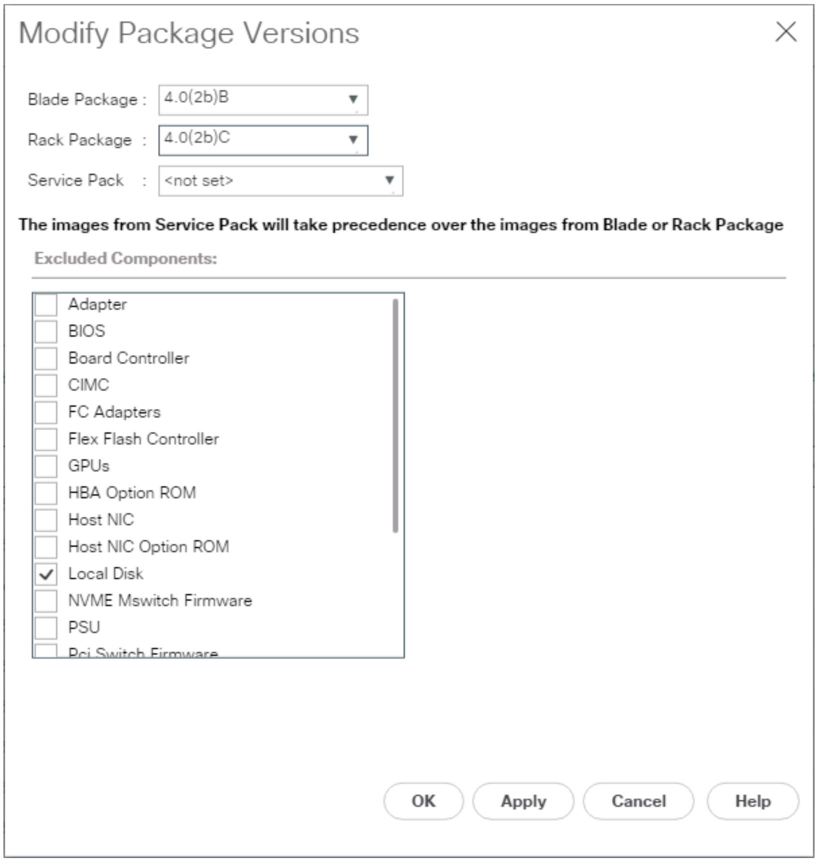

Modify Package Versions

Blade Package : 4.0(2b)B

Rack Package : 4.0(2b)C

Service Pack : <not set>

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

8. Click OK to modify the host firmware package and OK again to acknowledge the changes.

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 the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy `UCS-B200M5`.
6. Select Create Server PID Qualifications.
7. Select Cisco UCS-B200-M5 from the PID drop-down list.

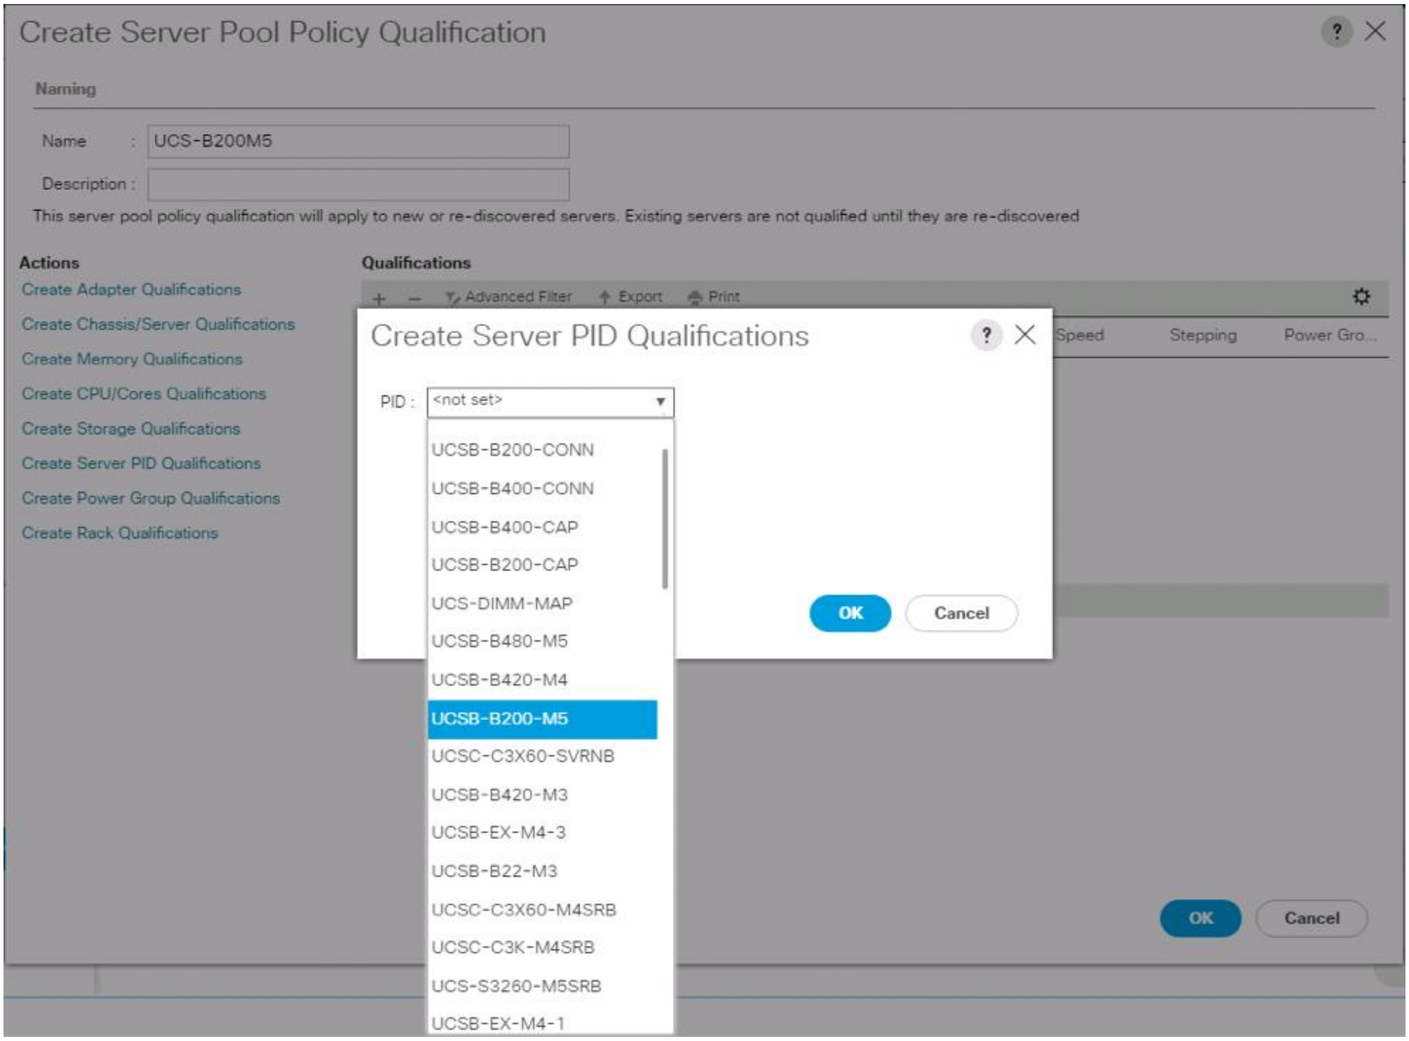

Create Server Pool Policy Qualification

Naming

Name : UCS-B200M5

Description :

This server pool policy qualification will apply to new or re-discovered servers. Existing servers are not qualified until they are re-discovered

Create Server PID Qualifications

PID : <not set>

Actions
Create Adapter Qualifications
Create Chassis/Server Qualifications
Create Memory Qualifications
Create CPU/Cores Qualifications
Create Storage Qualifications
Create Server PID Qualifications
Create Power Group Qualifications
Create Rack Qualifications
Qualifications
+ − Advanced Filter Export Print
Speed
Stepping
Power Gro...
UCSB-B200-CONN
UCSB-B400-CONN
UCSB-B400-CAP
UCSB-B200-CAP
UCS-DIMM-MAP
UCSB-B480-M5
UCSB-B420-M4
UCSB-B200-M5
UCSC-C3X60-SVRNB
UCSB-B420-M3
UCSB-EX-M4-3
UCSB-B22-M3
UCSC-C3X60-M4SRB
UCSC-C3K-M4SRB
UCS-S3260-M5SRB
UCSB-EX-M4-1
OK
Cancel
OK
Cancel

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 click OK for the confirmation.

Download the Image for ESXi 6.7 U1

The VMware Cisco Custom Image will need to be downloaded for use during installation by manual access to the UCS KVM vMedia, or through a vMedia Policy explained in the following subsection.

To download the Cisco Custom Image, follow these steps:

1. Click the following link: [VMware vSphere Hypervisor Cisco Custom Image (ESXi) 6.7 U1.]()

2. You will need a user id and password on vmware.com to download this software.

3. Download the .iso file.

Create vMedia Policy for VMware ESXi 6.7 U1 Install Boot (optional, if manually attaching ISO through KVM)

A separate HTTP web server is required to automate the availability of the ESXi image to each Service Profile on first power on. The creation of this web server is not included in this document but can be any existing web server capable of serving files through HTTP that are accessible on the OOB network that the ESXi image can be placed upon.

To create a vMedia Policy,place the Cisco Custom Image VMware ESXi 6.7 U1 ISO on the HTTP server and 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 ISO for ESXi 6.7 U1`" 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.

> Since DNS server IPs were not entered into the KVM IP earlier, it is necessary to enter the IP of the web server instead of the hostname.

12. Leave "None" selected for Image Name Variable.
13. Enter VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso as the Remote File name.
14. Enter the web server path to the ISO file in the Remote Path field.

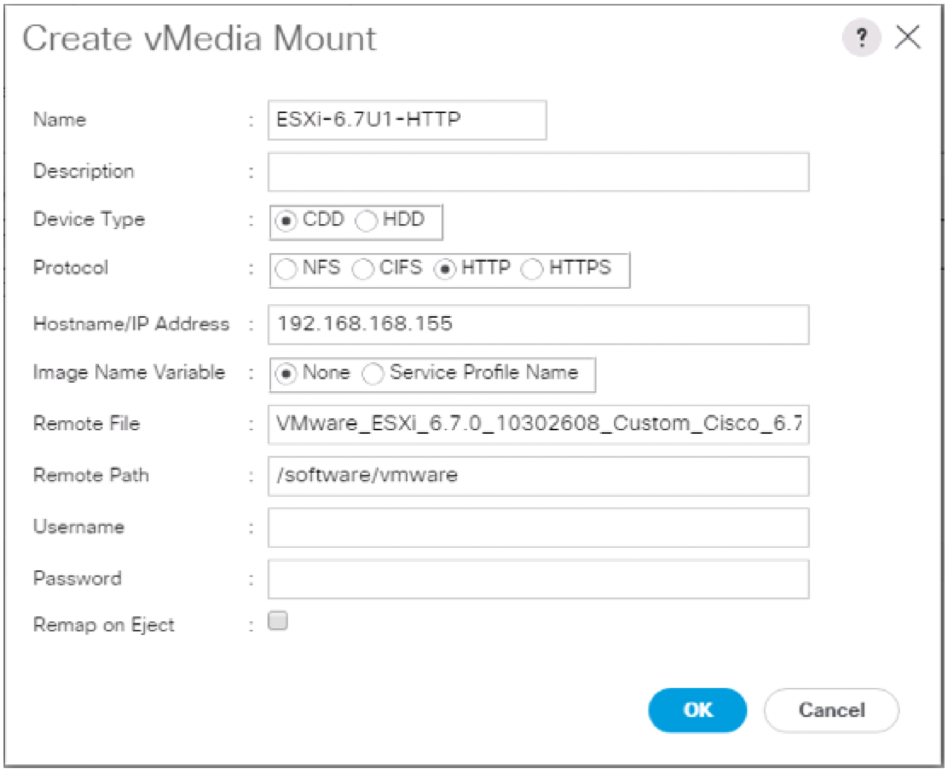
Create vMedia Mount

Name : ESXi-6.7U1-HTTP

Description :

Device Type : ● CDD ○ HDD

Protocol : ○ NFS ○ CIFS ● HTTP ○ HTTPS

Hostname/IP Address : 192.168.168.155

Image Name Variable : ● None ○ Service Profile Name

Remote File : VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7

Remote Path : /software/vmware

Username :

Password :

Remap on Eject : □

15. Click OK to create the vMedia Mount.

16. 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. After ESXi is installed, the vMedia will not be referenced as long as the boot disk is accessible.

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Select Policies > root.

3. Right-click BIOS Policies.

4. Select Create BIOS Policy.

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

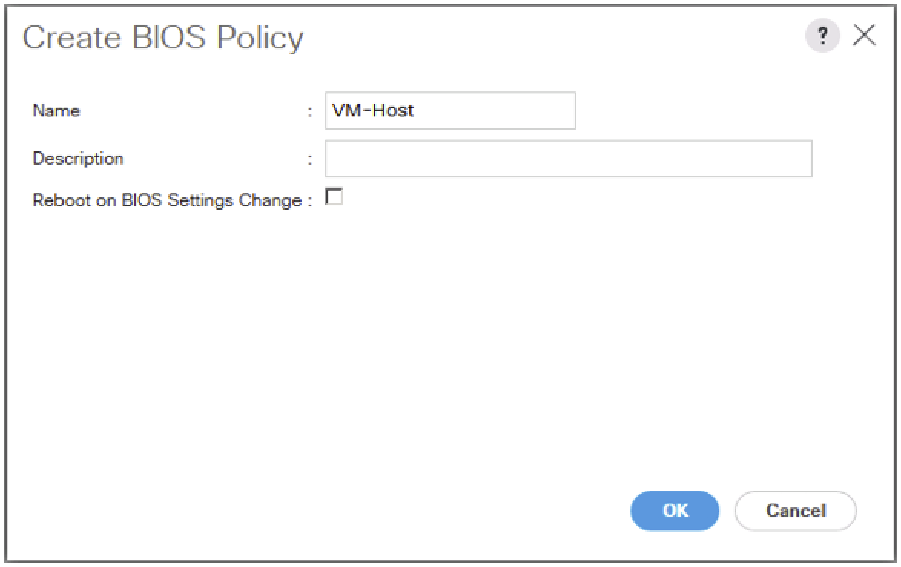

Create BIOS Policy

Name : VM-Host

Description :

Reboot on BIOS Settings Change : □

6. Select and right-click the newly created BIOS Policy.

7. Within the Main tab of the Policy:

    a. Change CDN Control to enabled.

    b. Change the Quiet Boot setting to disabled.

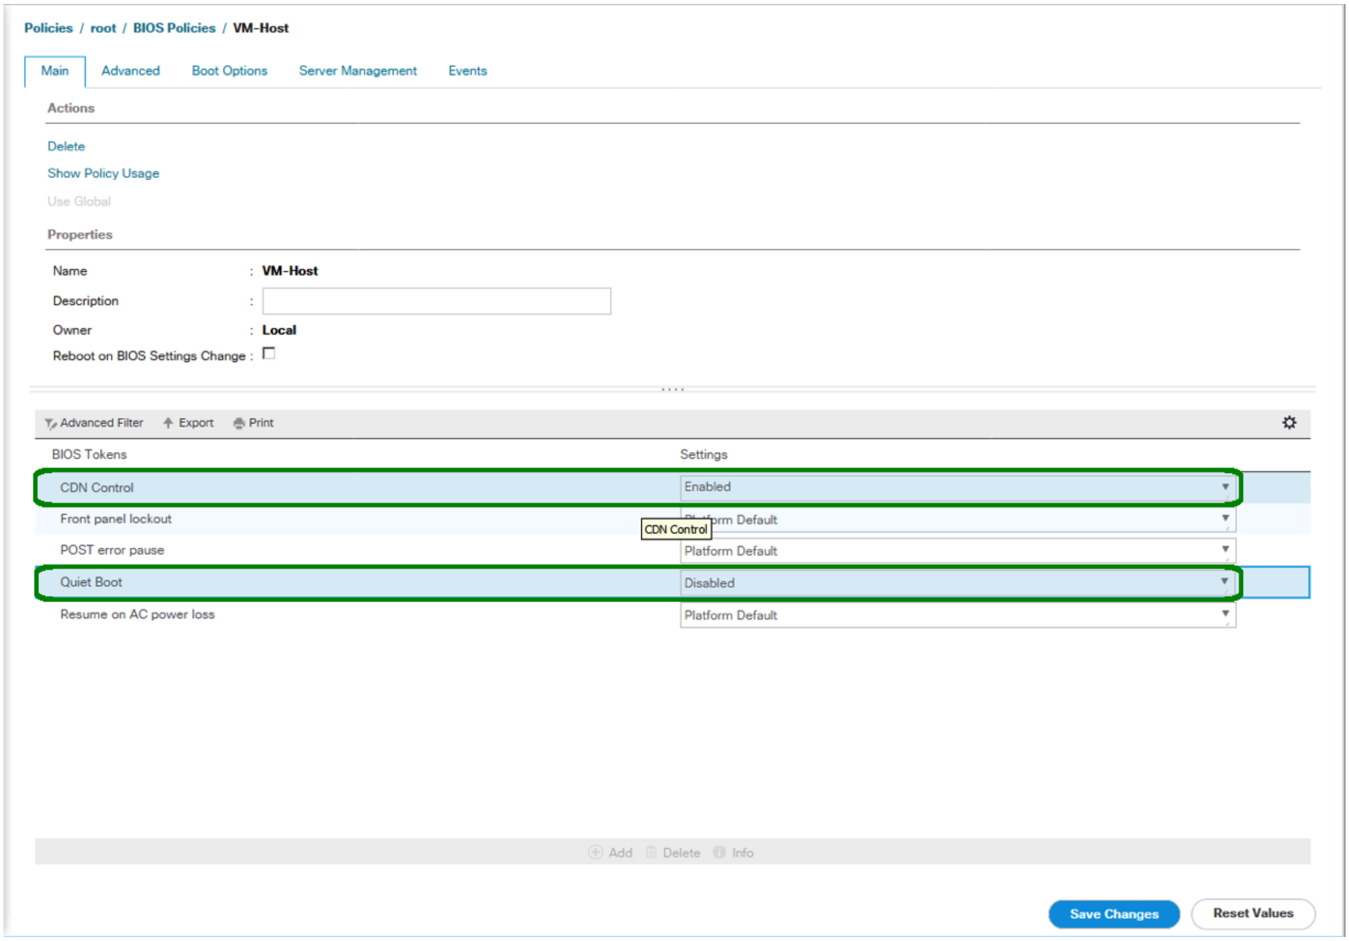

Policies / root / BIOS Policies / VM-Host

Name : VM-Host

Description :

Owner : Local

Reboot on BIOS Settings Change : □

BIOS Tokens	Settings
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.

9. Set the following within the Processor tab:

    a. DRAM Clock Throttling -> Performance

    b. Frequency Floor Override -> Enabled

    c. Processor C State -> Disabled

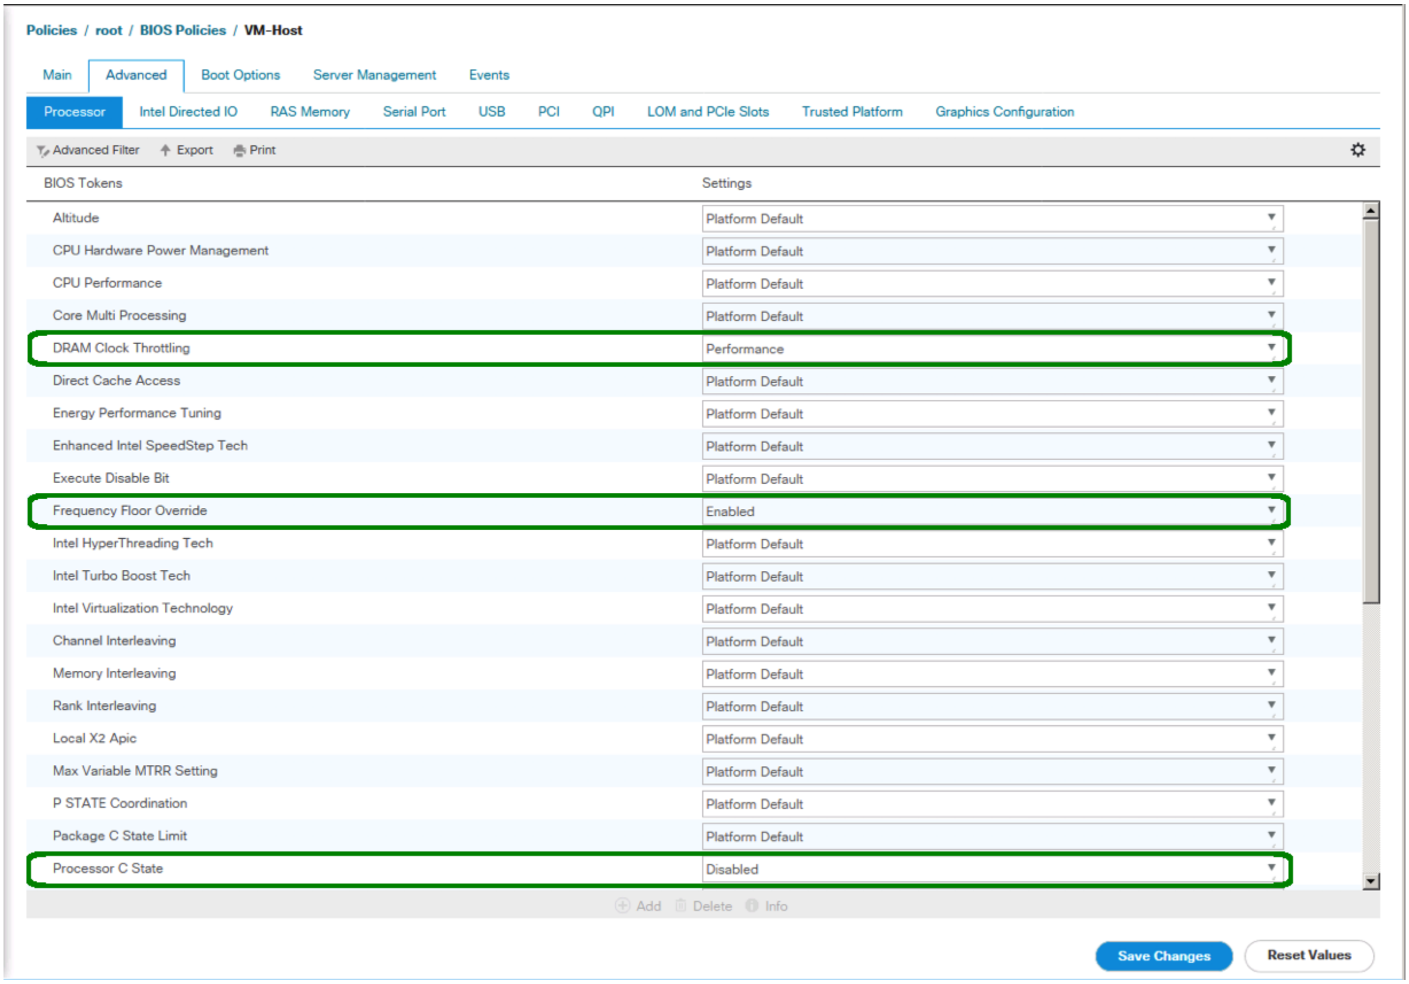

BIOS Tokens	Settings
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
CPU Performance	Platform Default
Core Multi Processing	Platform Default
DRAM Clock Throttling	Performance
Direct Cache Access	Platform Default
Energy Performance Tuning	Platform Default
Enhanced Intel SpeedStep Tech	Platform Default
Execute Disable Bit	Platform Default
Frequency Floor Override	Enabled
Intel HyperThreading Tech	Platform Default
Intel Turbo Boost Tech	Platform Default
Intel Virtualization Technology	Platform Default
Channel Interleaving	Platform Default
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Processor C State	Disabled

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

    a. Processor C1E -> disabled

    b. Processor C3 Report -> disabled

    c. Processor C7 Report -> disabled

    d. Energy Performance -> performance

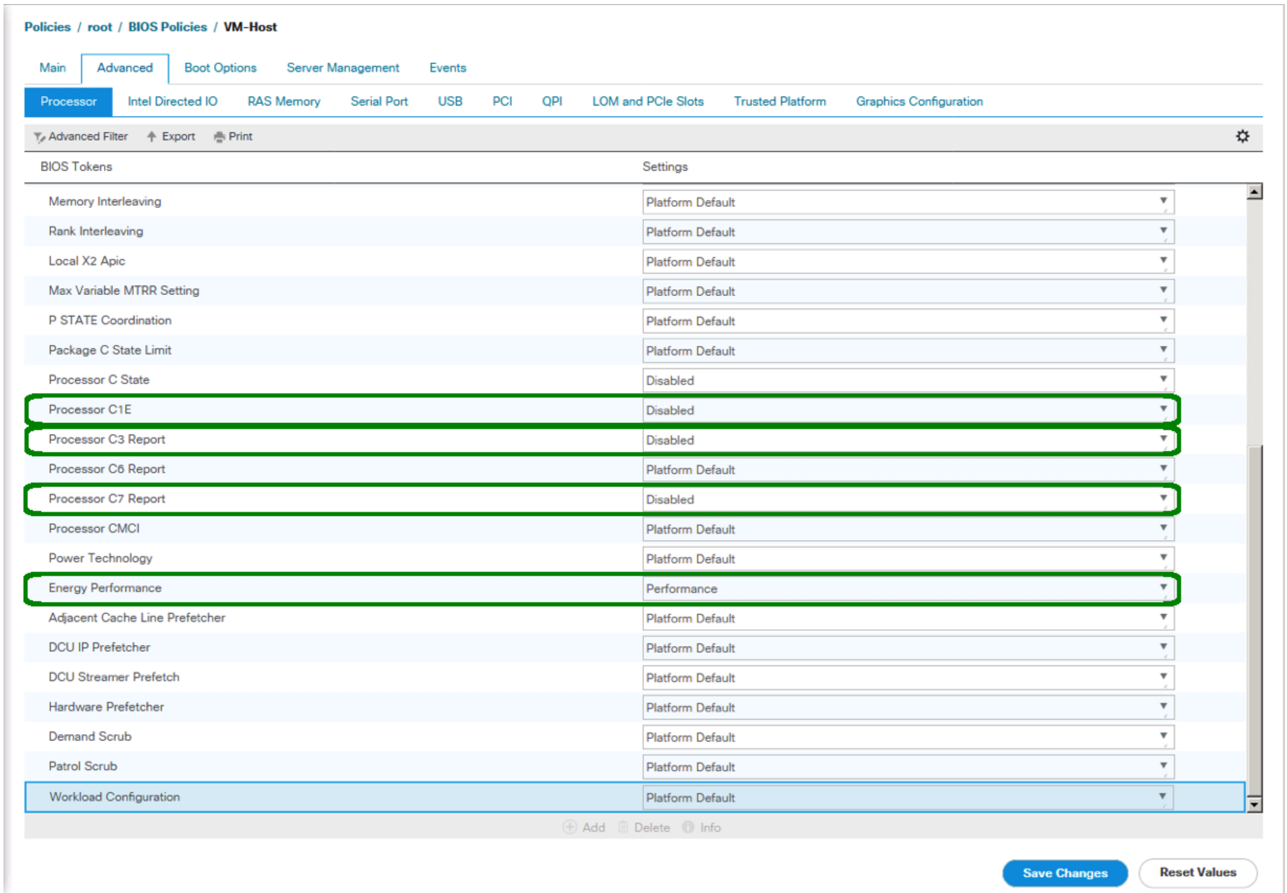

BIOS Tokens	Settings
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Platform Default
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Platform Default
Energy Performance	Performance
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	Platform Default

11. Click the RAS Memory tab and select:

    a. LV DDR Mode -> performance-mode

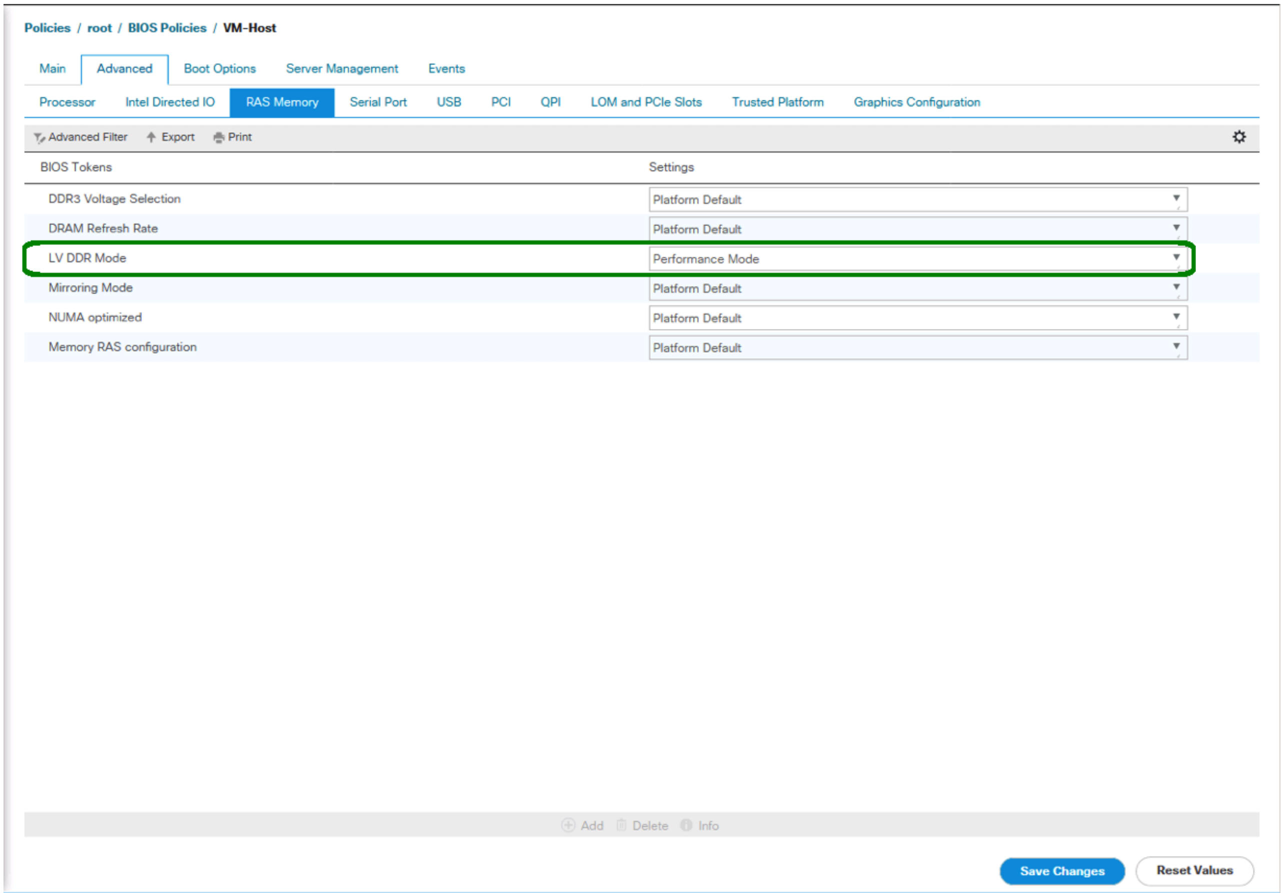

BIOS Tokens	Settings
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

12. Click Save Changes.

13. Click OK.

Update the Default Maintenance Policy

To update the default Maintenance Policy, follow these steps:

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

2. Select Policies > root.

3. Select Maintenance Policies > default.

4. Change the Reboot Policy to User Ack.

5. (Optional: Click “On Next Boot” to delegate maintenance windows to server owners).

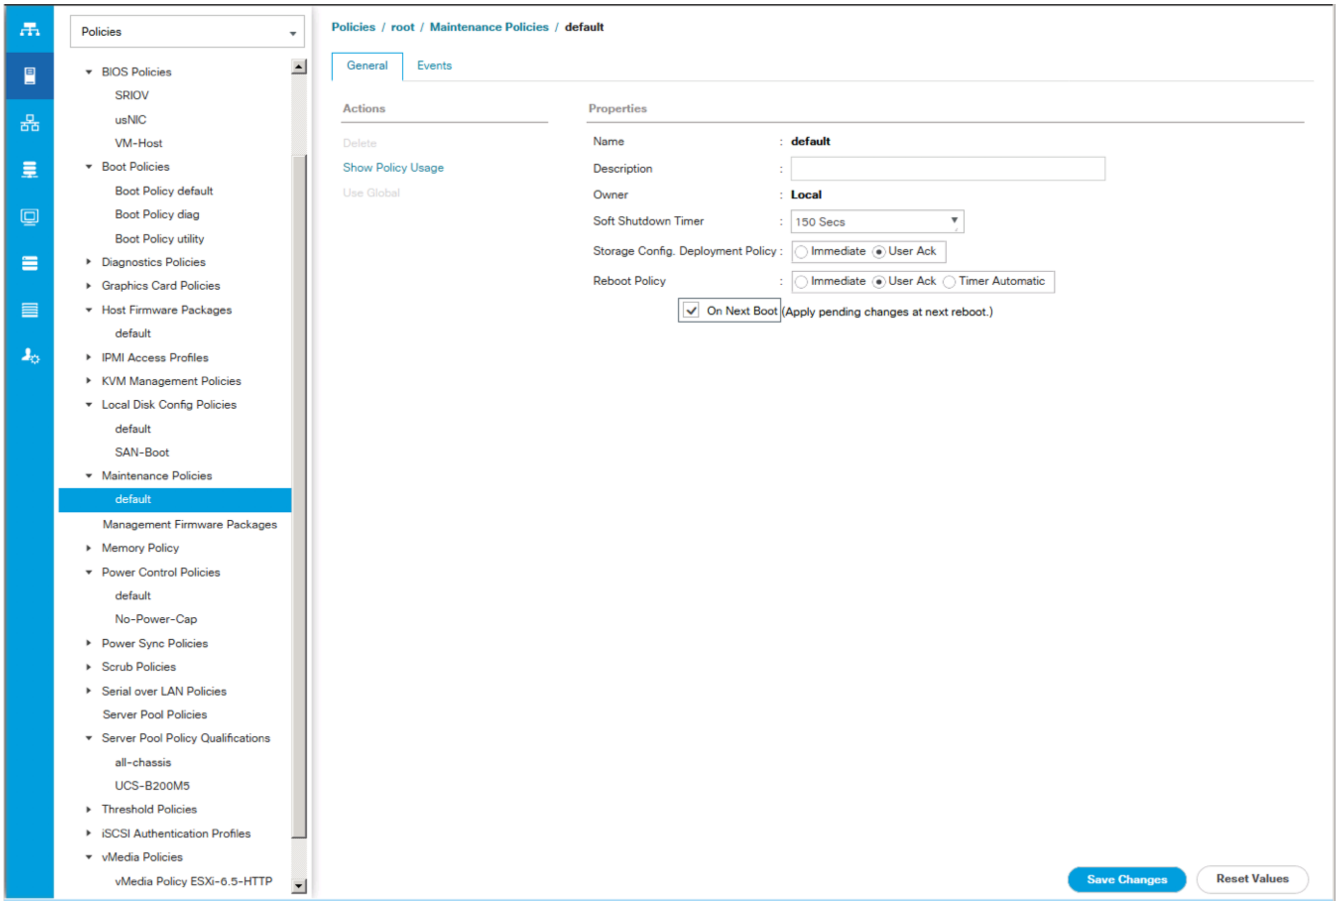

Policies / root / Maintenance Policies / default

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 change.

Create Local Disk Configuration Policy

A local disk configuration 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 the Servers tab in the navigation pane.

2. Select Policies > root.

3. Right-click Local Disk Config Policies.

4. Select Create Local Disk Configuration Policy.

5. Enter SAN-Boot as the local disk configuration policy name.

6. Change the mode to No Local Storage.

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

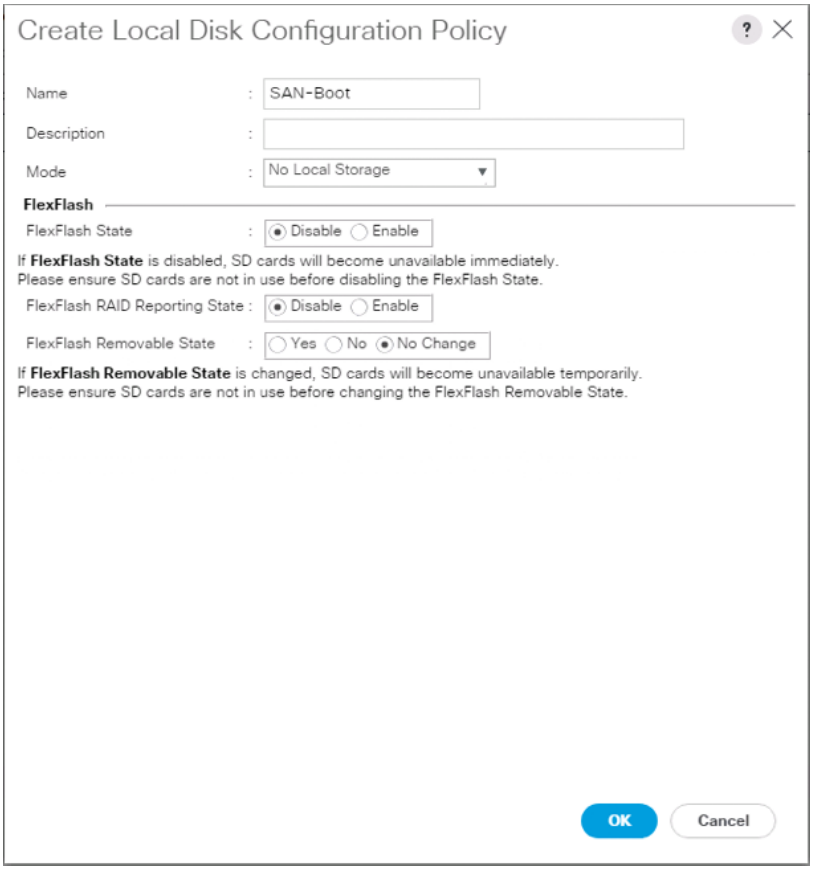

Create Local Disk Configuration Policy

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.

8. 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 the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.

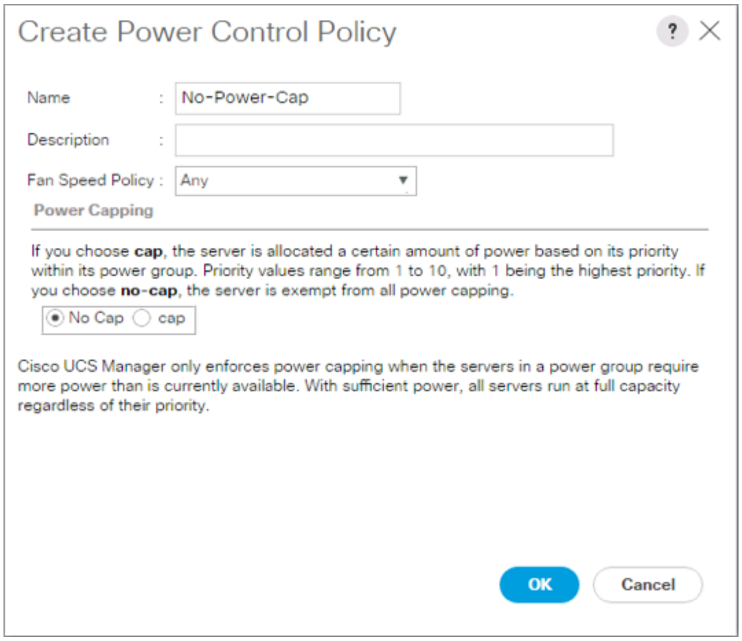

Create Power Control Policy

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 Network Control Policy for Cisco Discovery Protocol

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

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

2. Select Policies > root.

3. Right-click Network Control Policies.

4. Select Create Network Control Policy.

5. Enter `Enable_CDP` as the policy name.

6. For CDP, select the Enabled option.

7. Click OK to create the network control policy.

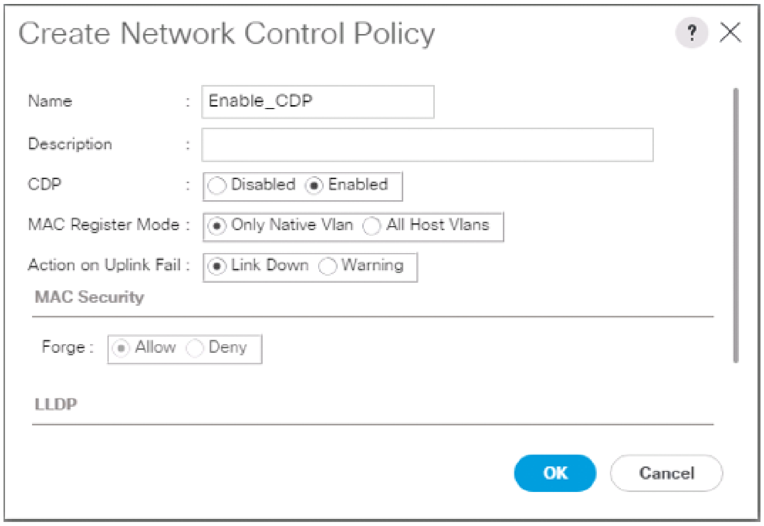

Create Network Control Policy

Name : Enable_CDP

Description :

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

8. Click OK.

Configure Cisco UCS LAN Connectivity

Create Uplink Port Channels

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

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

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

2. Under LAN > LAN Cloud, expand the Fabric A tree.

3. Right-click Port Channels.

4. Select Create Port Channel.

5. Enter a unique ID for the port channel, (`15` in our example to correspond with the upstream Nexus port channel).

6. With 15 selected, enter `vPC-15-Nexus` as the name of the port channel.

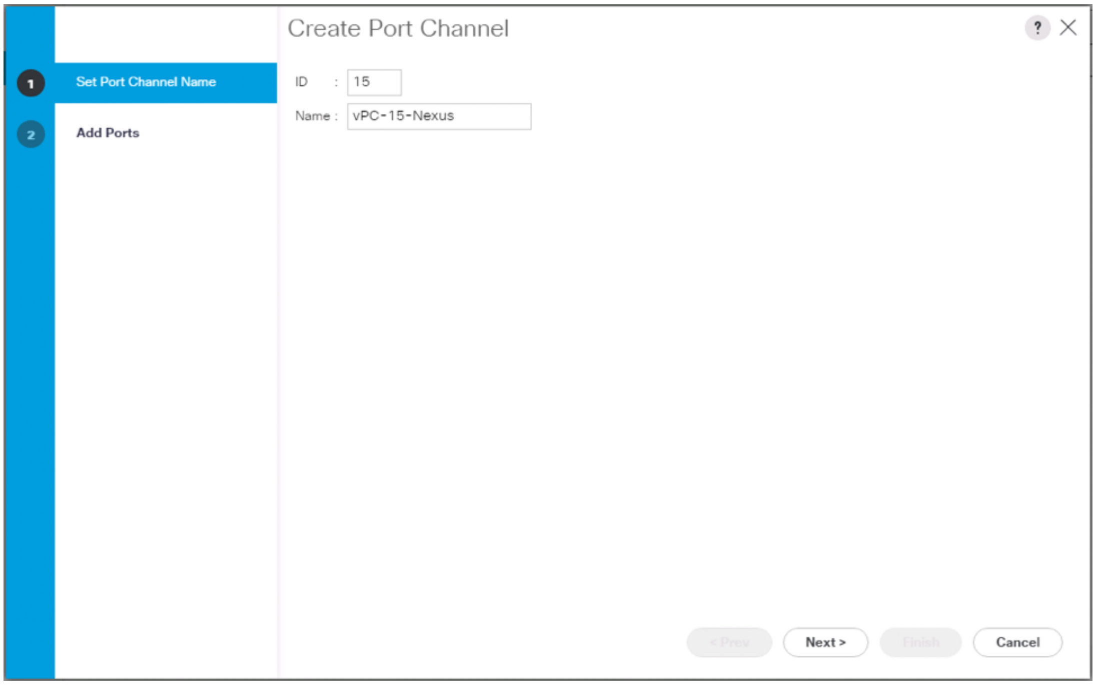

ID : 15

Name : vPC-15-Nexus

7. Click Next.

8. Select the following ports to be added to the port channel:

   a. Slot ID 1 and port 53

   b. Slot ID 1 and port 54

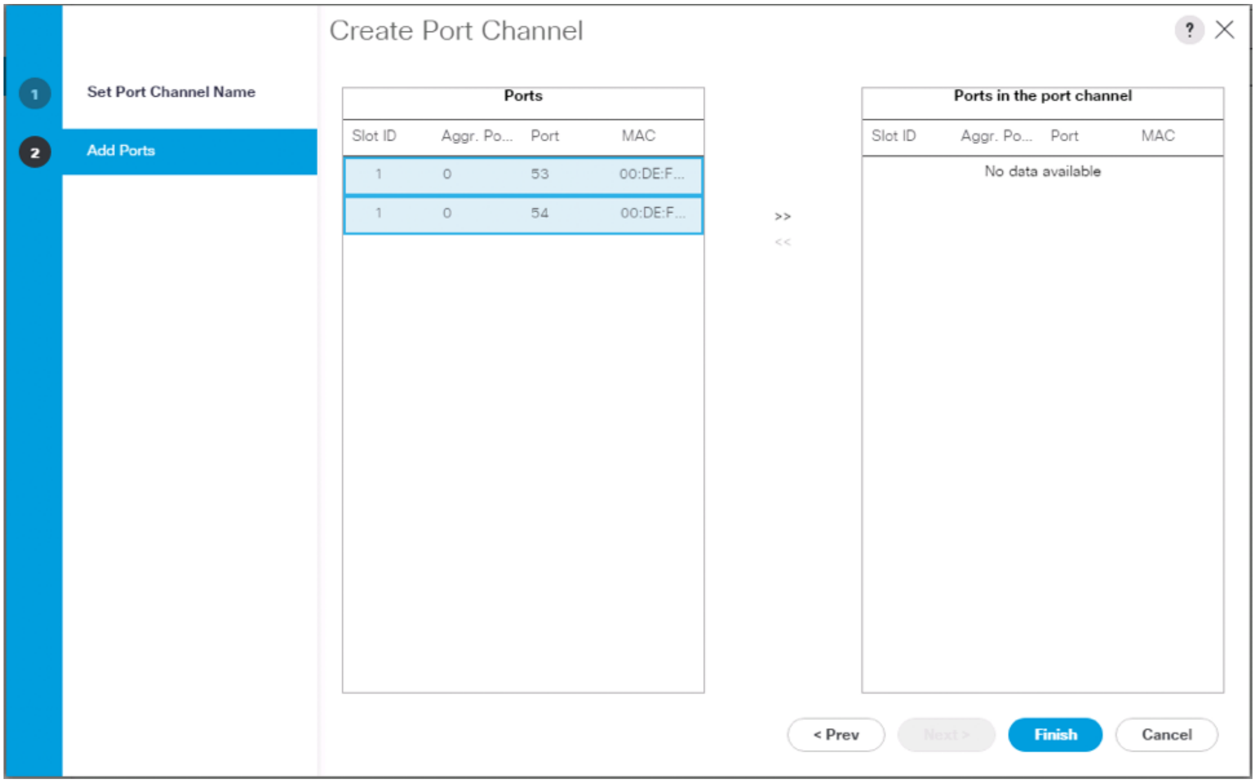

Ports

Slot ID	Aggr. Po...	Port	MAC
1	0	53	00:DE:F...
1	0	54	00:DE:F...

Ports in the port channel

Slot ID	Aggr. Po...	Port	MAC

No data available

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

10. Click Finish to create the port channel.

11. Click OK.

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

13. Right-click Port Channels.

14. Select Create Port Channel.

15. Enter a unique ID for the port channel, (`16` in our example to correspond with the upstream Nexus port channel).

16. With 16 selected, enter `vPC-16-Nexus` as the name of the port channel.

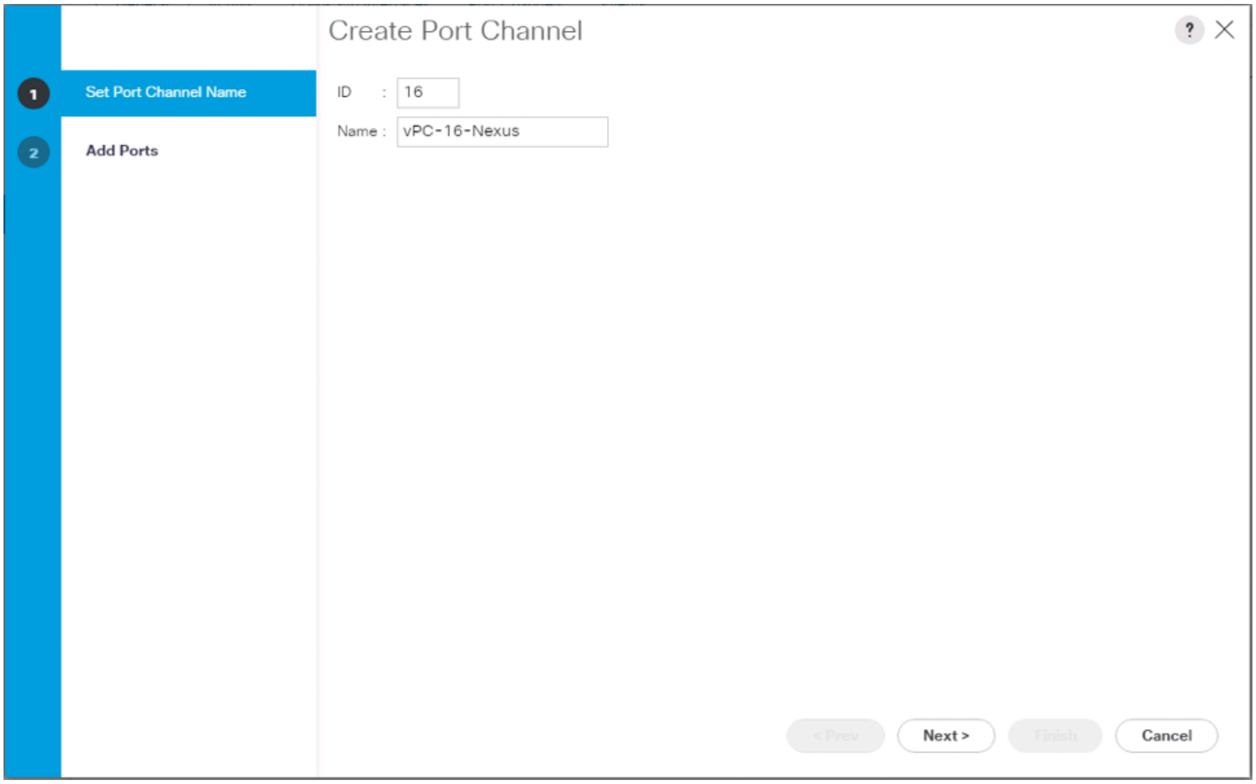

ID : 16

Name : vPC-16-Nexus

17. Click Next.

18. Select the following ports to be added to the port channel:

    a. Slot ID 1 and port 53

    b. Slot ID 1 and port 54

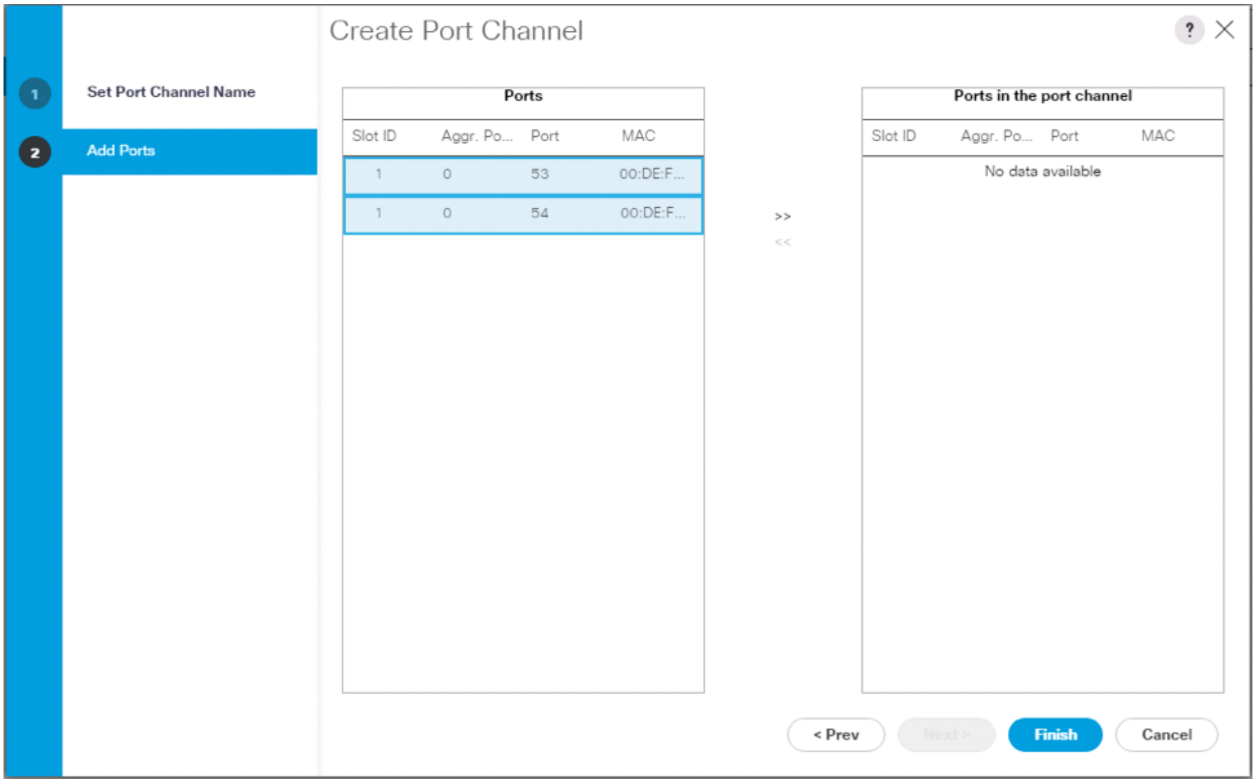

Ports

Slot ID	Aggr. Po...	Port	MAC
1	0	53	00:DE:F...
1	0	54	00:DE:F...

Ports in the port channel

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

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

20. Click Finish to create the port channel.

21. Click OK.

When using QSFP+ passive copper cables (e.g. QSFP-100G-CU1M), setting the appropriate port speed for the configured port channel interfaces may be needed depending upon switch and switch ports used.

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 the LAN tab in the navigation pane.

In this procedure, six unique VLANs are created. See Table 2 for a list of VLANs to be created.

2. Select LAN > LAN Cloud.

3. Right-click VLANs.

4. Select Create VLANs.

5. Enter `Native` 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.

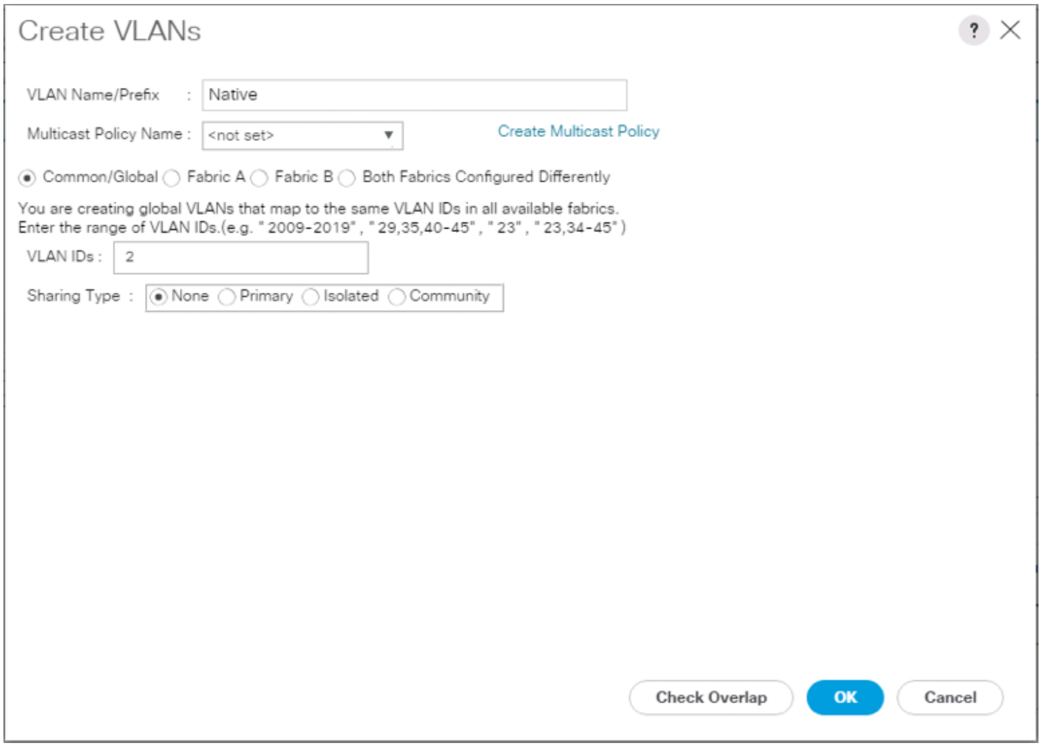

Create VLANs

VLAN Name/Prefix : Native

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

9. Click OK and then click OK again.

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

11. Click Yes and then click OK.

12. Right-click VLANs.

13. Select Create VLANs.

14. Enter `IB-Mgmt` as the name of the VLAN to be used for management traffic.

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.

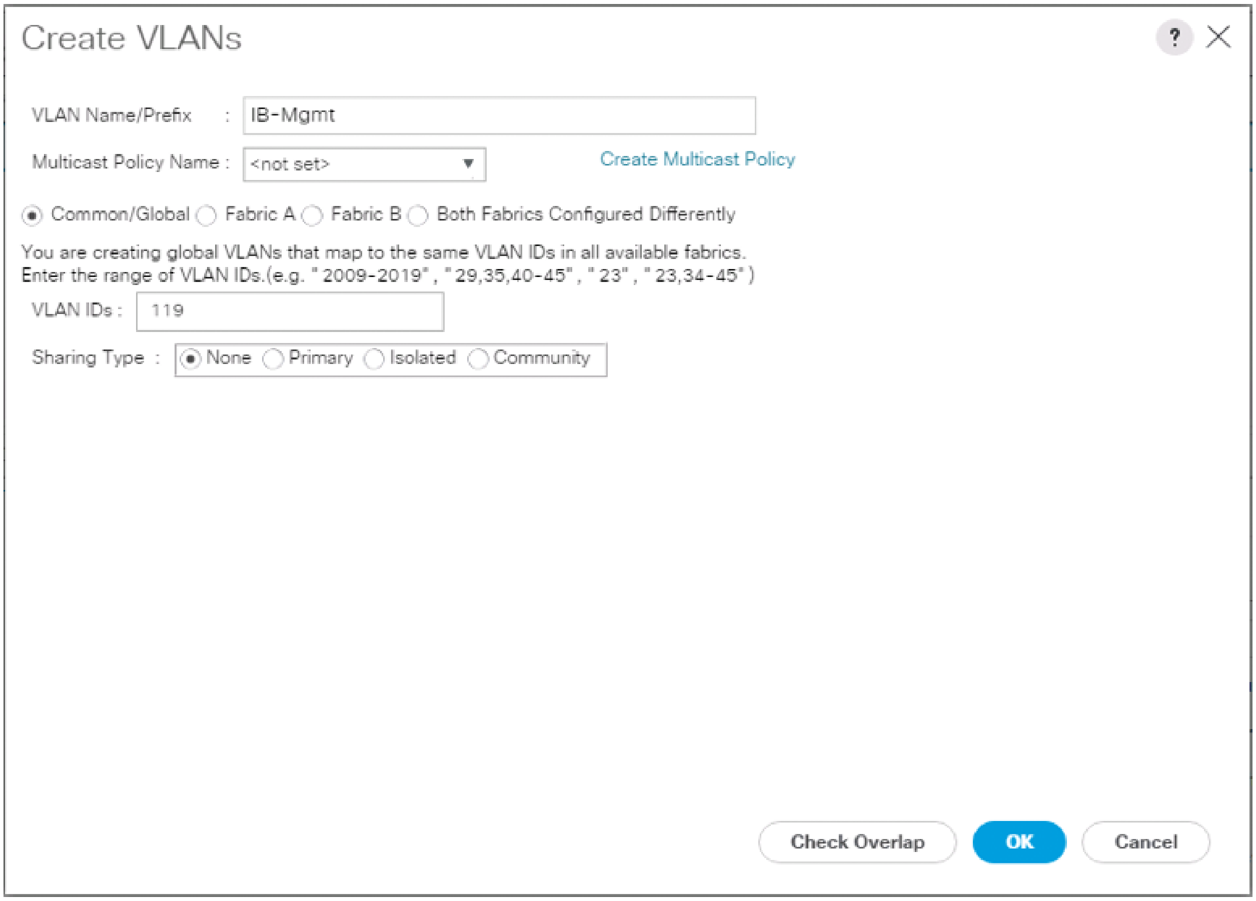

Create VLANs

VLAN Name/Prefix : IB-Mgmt

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 : 119

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

18. Click OK and then click OK again.

19. Right-click VLANs.

20. Select Create VLANs.

21. Enter `vMotion` as the name of the VLAN to be used for vMotion.

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

23. Enter the vMotion VLAN ID.

24. Keep the Sharing Type as None.

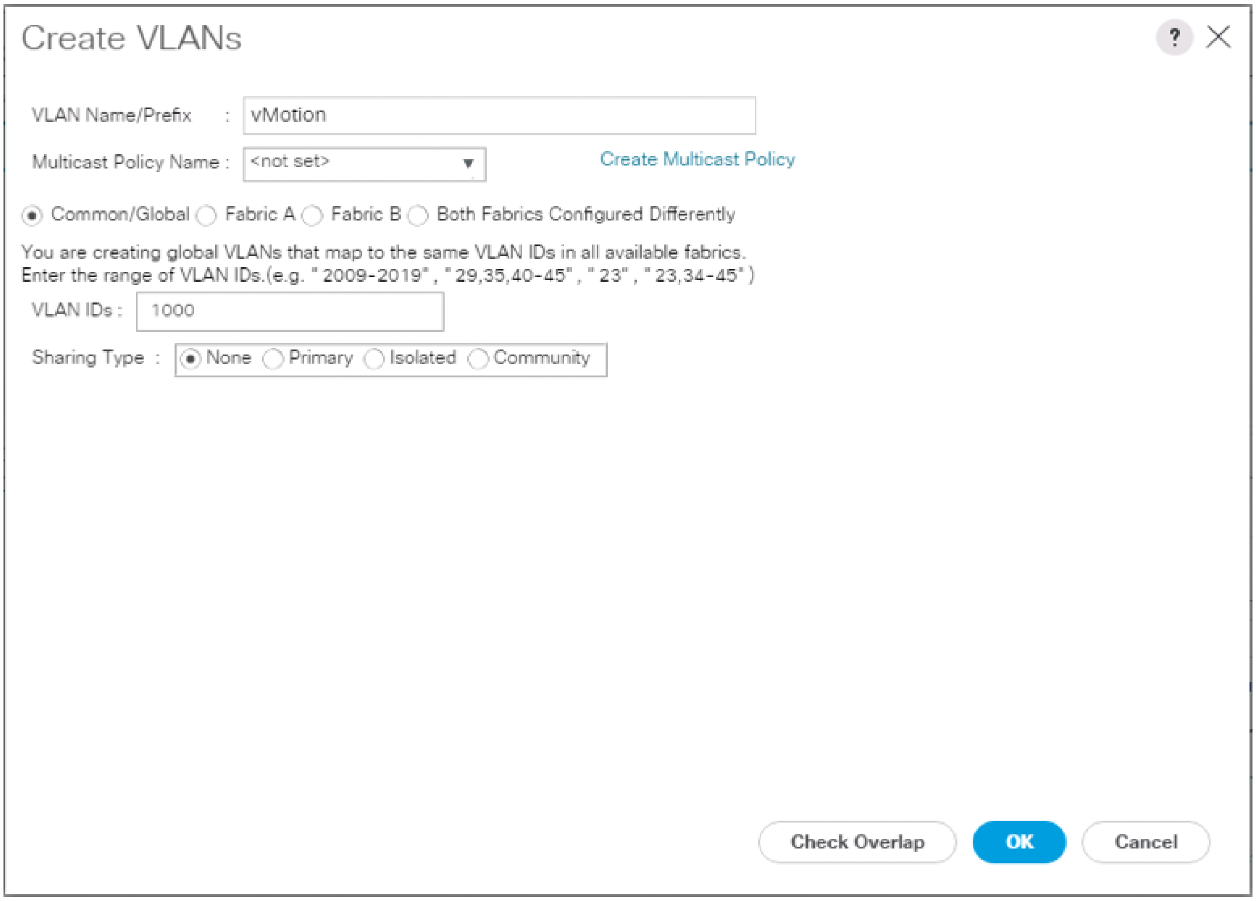

Create VLANs

VLAN Name/Prefix : vMotion

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 : 1000

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

25. Click OK and then click OK again.

26. Right-click VLANs.

27. Select Create VLANs.

28. Enter `VM–App–` as the prefix of the VLANs to be used for VM Traffic.

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

30. Enter the VM-Traffic VLAN ID range.

31. Keep the Sharing Type as None.

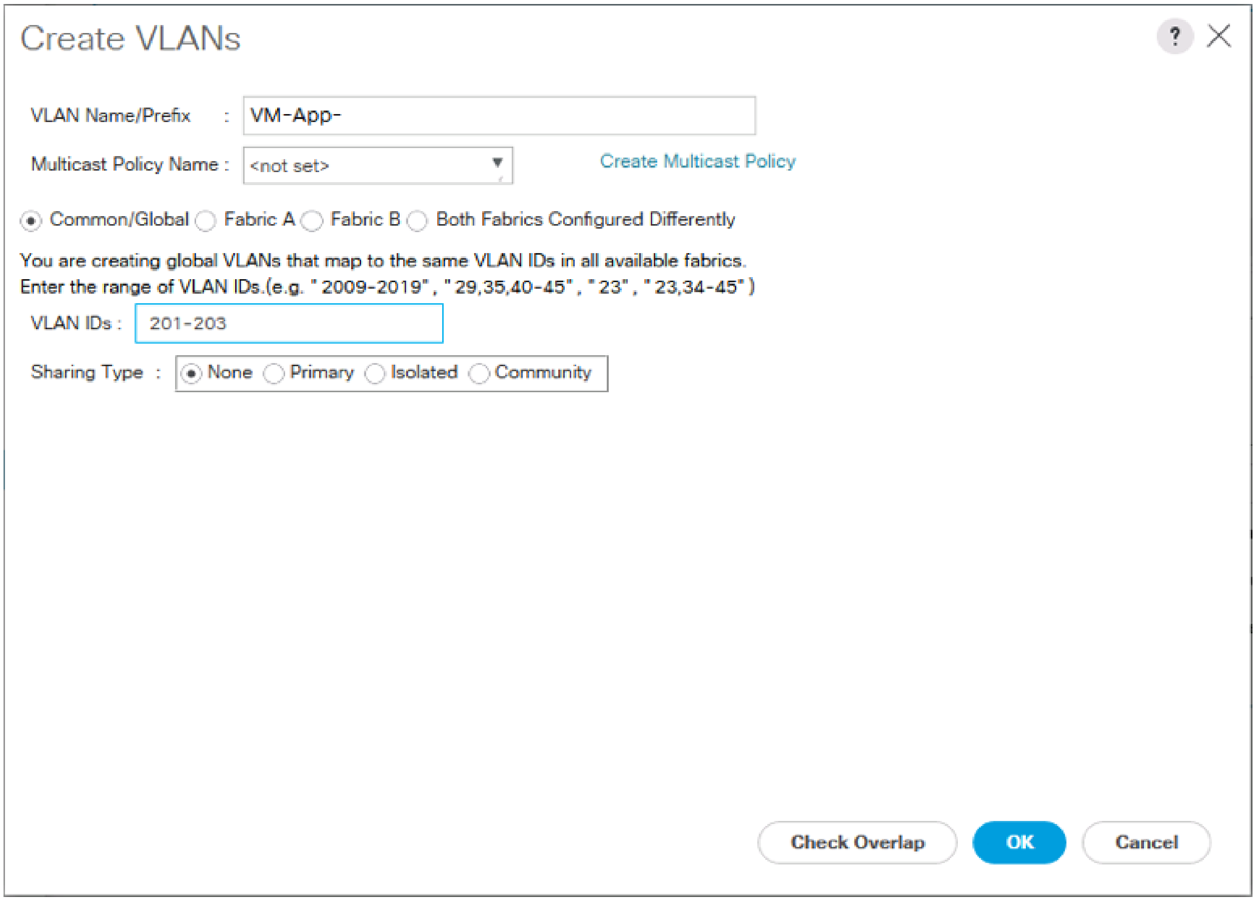

Create VLANs

VLAN Name/Prefix : VM-App-

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 : 201-203

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

32. Click OK and then click OK again.

33. Repeat as needed for any additional VLANs created on the upstream Nexus switches.

Create vNIC Templates

To create the multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, follow the steps in this section.

Create Management vNICs

For the vNIC_Mgmt_A Template, follow these steps:

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

2. Select Policies > root.

3. Right-click vNIC Templates.

4. Select Create vNIC Template.

5. Enter `vNIC_Mgmt_A` as the vNIC template name.

6. Keep Fabric A selected.

7. Select `Primary Template` for the Redundancy Type.

8. Leave Peer Redundancy Template as <not set>

Redundancy Type and specification of Redundancy Template are configuration options to later allow changes to the Primary Template to automatically adjust onto the Secondary Template.

9. Under Target, make sure that the VM checkbox is not selected.

10. Select Updating Template as the Template Type.

11. Under VLANs, select the checkboxes for `IB-Mgmt, vMotion, and Native` VLANs.

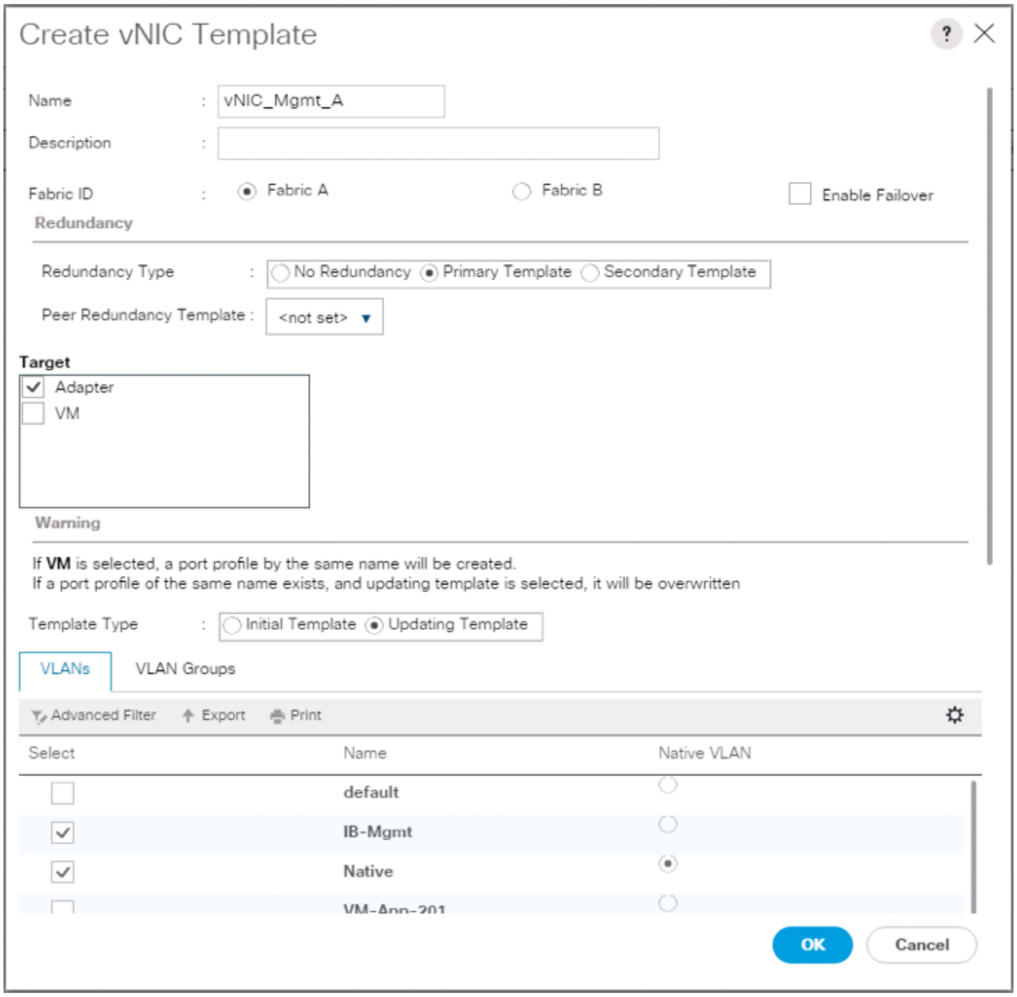

Create vNIC Template

Name : vNIC_Mgmt_A

Description :

Fabric ID : ● Fabric A ○ Fabric B □ Enable Failover

Redundancy

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

Peer Redundancy Template : <not set>

Target

☑ Adapter

□ VM

Warning

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

Template Type : ○ Initial Template ● Updating Template

Select	Name	Native VLAN
□	default	○
☑	IB-Mgmt	○
☑	Native	●
□	VM-App-201	○

12. Set `Native` as the native VLAN.

13. Leave `vNIC Name` selected for the CDN Source.

14. For MTU, enter `9000`.

15. In the MAC Pool list, select `MAC_Pool_A`.

16. In the Network Control Policy list, select `Enable_CDP`.

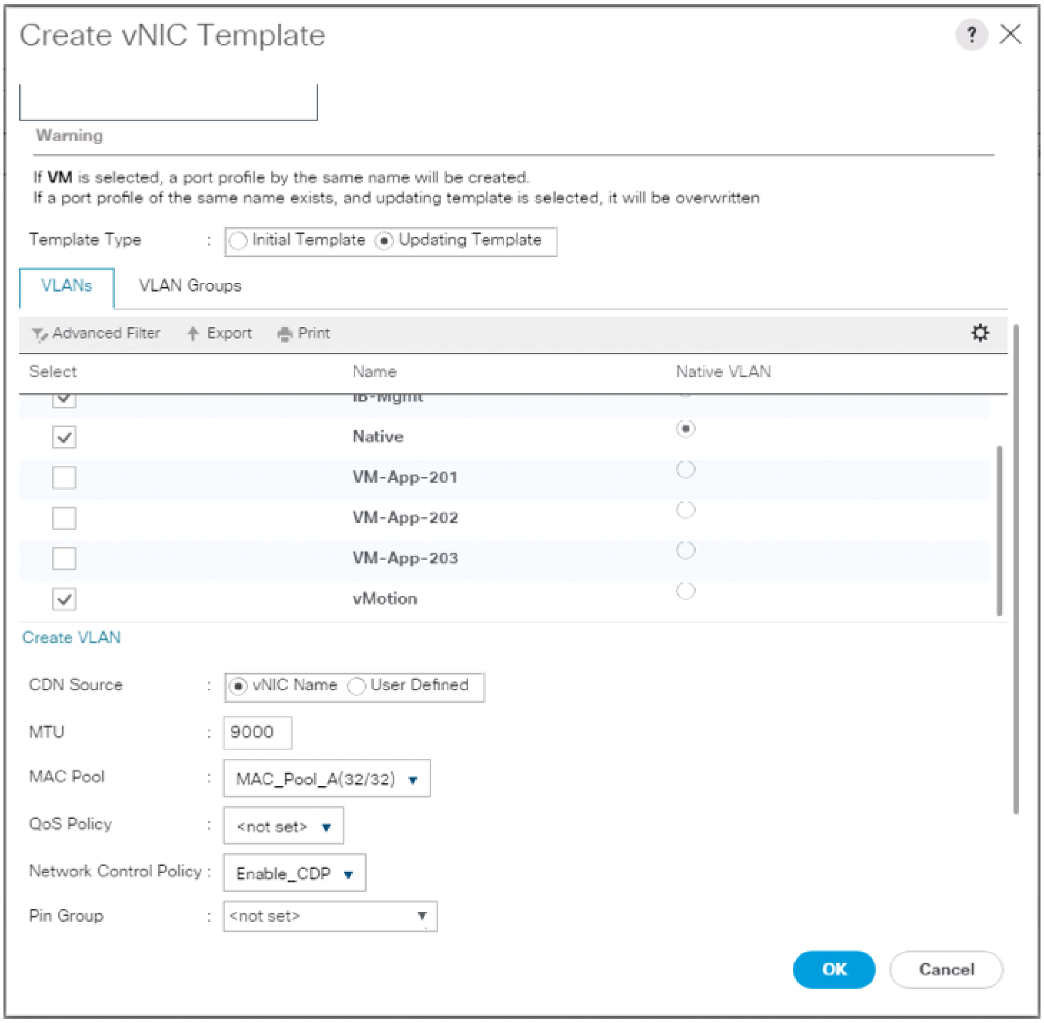

Create vNIC Template

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
☑	IB-Mgmt	○
☑	Native	●
□	VM-App-201	○
□	VM-App-202	○
□	VM-App-203	○
☑	vMotion	○

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC_Pool_A(32/32)

QoS Policy : <not set>

Network Control Policy : Enable_CDP

Pin Group : <not set>

17. Click OK to create the vNIC template.

18. Click OK.

For the vNIC_Mgmt_B Template, follow these steps:

1. In the navigation pane, select the LAN tab.

2. Select Policies > root.

3. Right-click vNIC Templates.

4. Select Create vNIC Template

5. Enter `vNIC_Mgmt_B` as the vNIC template name.

6. Select Fabric B.

7. Select Secondary Template for Redundancy Type.

8. For the Peer Redundancy Template drop-down, select vNIC_Mgmt_A.

With Peer Redundancy Template selected, Template Type, VLANs, CDN Source, MTU, and Network Control Policy are all pulled from the Primary Template.

9. Under Target, make sure the VM checkbox is not selected.

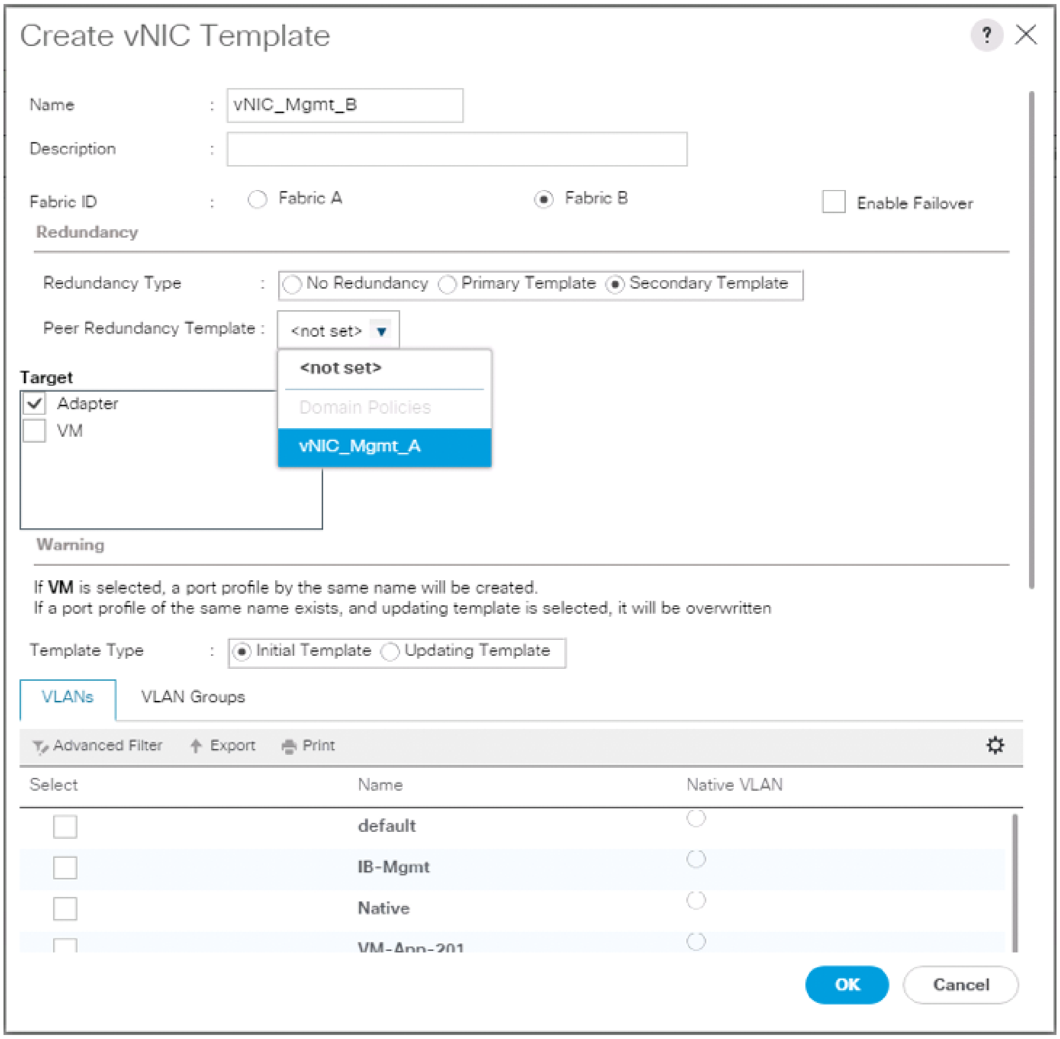

Create vNIC Template

Name : vNIC_Mgmt_B

Description :

Fabric ID : ○ Fabric A ● Fabric B □ Enable Failover

Redundancy

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

Peer Redundancy Template : <not set>

Target

☑ Adapter

□ VM

Warning

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

Template Type : ● Initial Template ○ Updating Template

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Native	○
□	VM-App-201	○

10. In the MAC Pool list, select MAC_Pool_B.

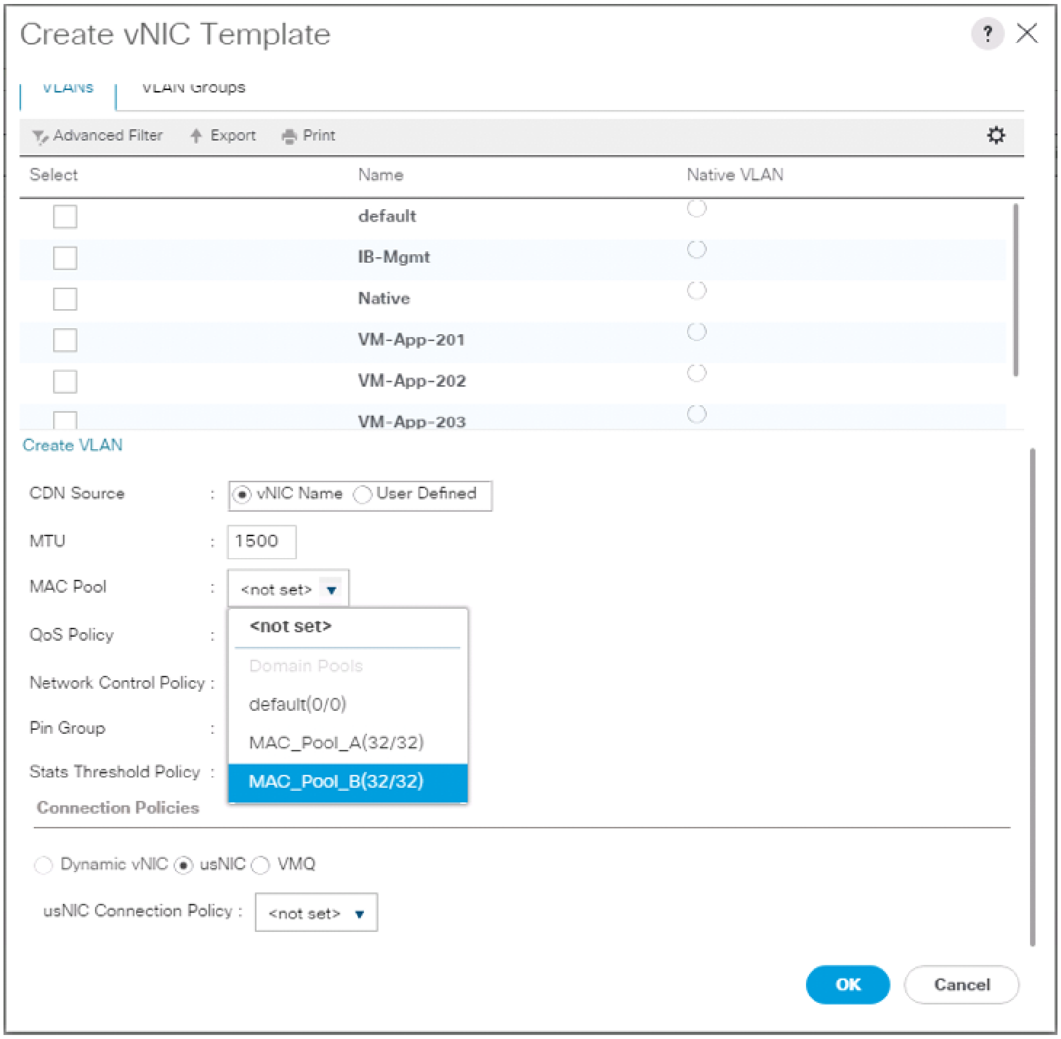

Select	Name	Native VLAN
☐	default	○
☐	IB-Mgmt	○
☐	Native	○
☐	VM-App-201	○
☐	VM-App-202	○
☐	VM-App-203	○

CDN Source : ● vNIC Name ○ User Defined

MTU : 1500

MAC Pool : <not set>

- <not set>
- Domain Pools
- default(0/0)
- MAC_Pool_A(32/32)
- MAC_Pool_B(32/32)

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

○ Dynamic vNIC ● usNIC ○ VMQ

usNIC Connection Policy : <not set>

11. Click OK to create the vNIC template.

12. Click OK.

Create Application vNICs

For the vNIC_App_A Template, follow these steps:

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

2. Select Policies > root.

3. Right-click vNIC Templates.

4. Select Create vNIC Template.

5. Enter `vNIC_App_A` as the vNIC template name.

6. Keep Fabric A selected.

7. Select `Primary Template` for the Redundancy Type.

8. Leave Peer Redundancy Template as <not set>

9. Under Target, make sure that the VM checkbox is not selected.

10. Select Updating Template as the Template Type.

11. Set `default` as the native VLAN.

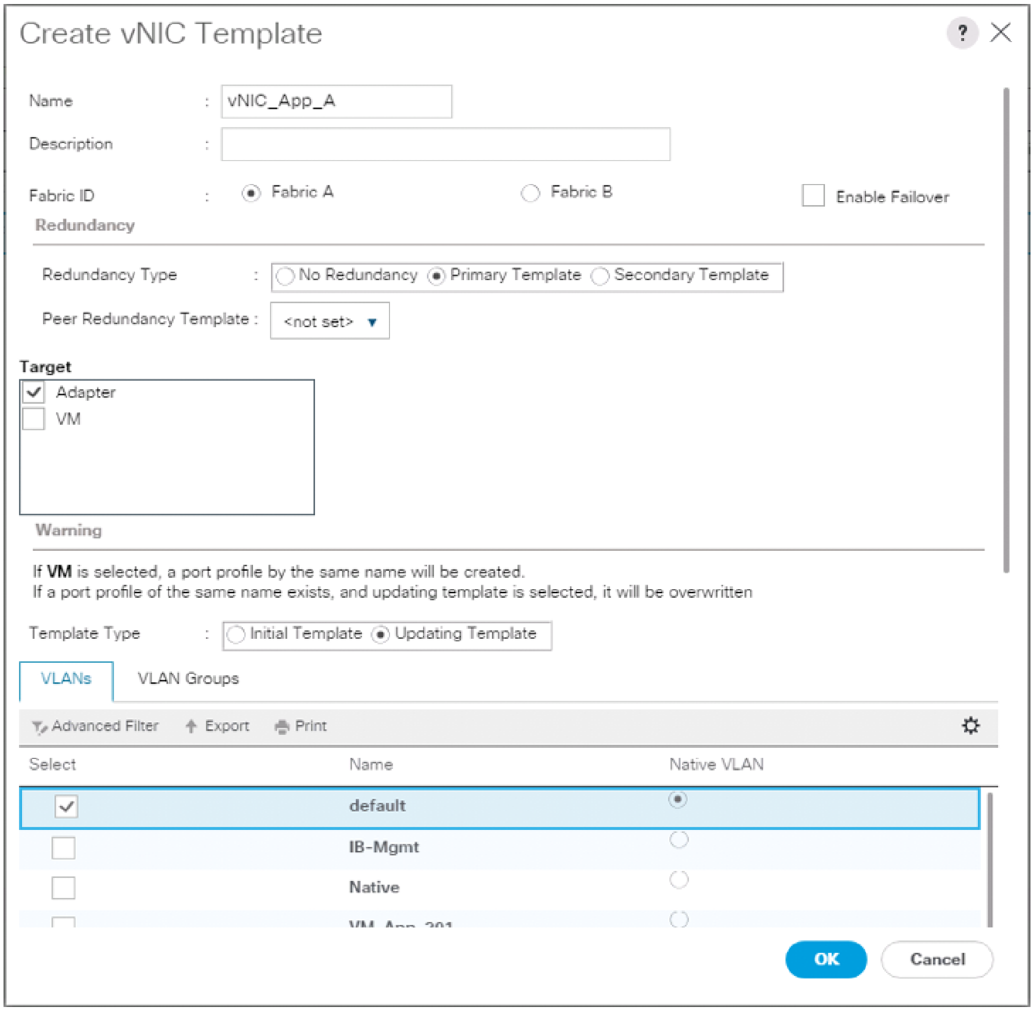

Create vNIC Template

Name : vNIC_App_A

Description :

Fabric ID : ● Fabric A ○ Fabric B □ Enable Failover

Redundancy

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

Peer Redundancy Template : <not set>

Target

☑ Adapter

□ VM

Warning

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

Template Type : ○ Initial Template ● Updating Template

Select	Name	Native VLAN
☑	**default**	●
□	**IB-Mgmt**	○
□	**Native**	○
□	**VM-App-201**	○

12. Under VLANs, select the checkboxes for any application or production VLANs that should be delivered to the ESXi hosts.

13. For MTU, enter `9000`.

14. In the MAC Pool list, select `MAC_Pool_A`.

15. In the Network Control Policy list, select `Enable_CDP`.

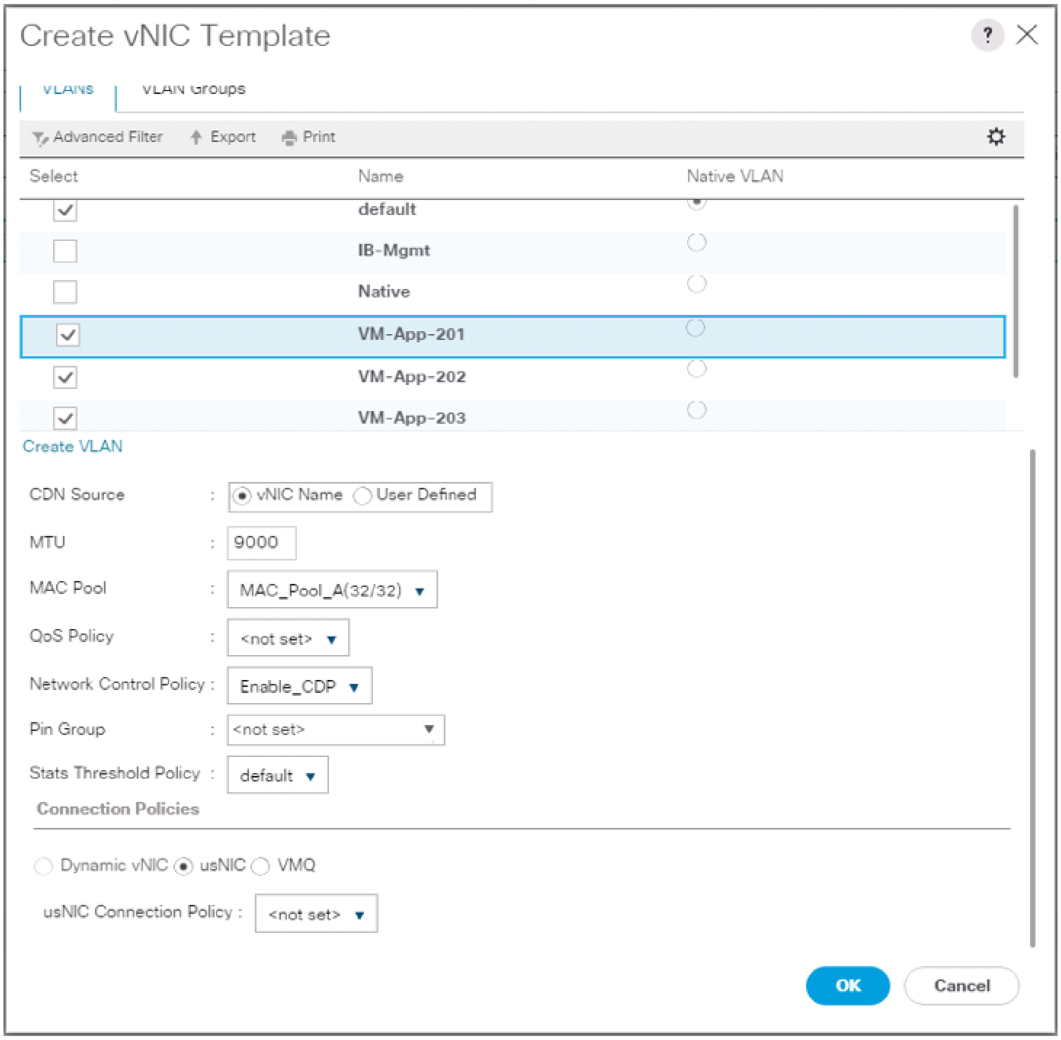

Create vNIC Template

Select	Name	Native VLAN
☑	default	●
□	IB-Mgmt	○
□	Native	○
☑	VM-App-201	○
☑	VM-App-202	○
☑	VM-App-203	○

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC_Pool_A(32/32)

QoS Policy : <not set>

Network Control Policy : Enable_CDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

○ Dynamic vNIC ● usNIC ○ VMQ

usNIC Connection Policy : <not set>

16. Click OK to create the vNIC template.

17. Click OK.

For the vNIC_App_B Template, follow these steps:

1. In the navigation pane, select the LAN tab.

2. Select Policies > root.

3. Right-click vNIC Templates.

4. Select Create vNIC Template

5. Enter `vNIC_App_B` as the vNIC template name.

6. Select Fabric B.

7. Select Secondary Template for Redundancy Type.

8. For the Peer Redundancy Template drop-down, select vNIC_App_A.

With Peer Redundancy Template selected, MAC Pool will be the main configuration option left for this vNIC template.

9. Under Target, make sure the VM checkbox is not selected.

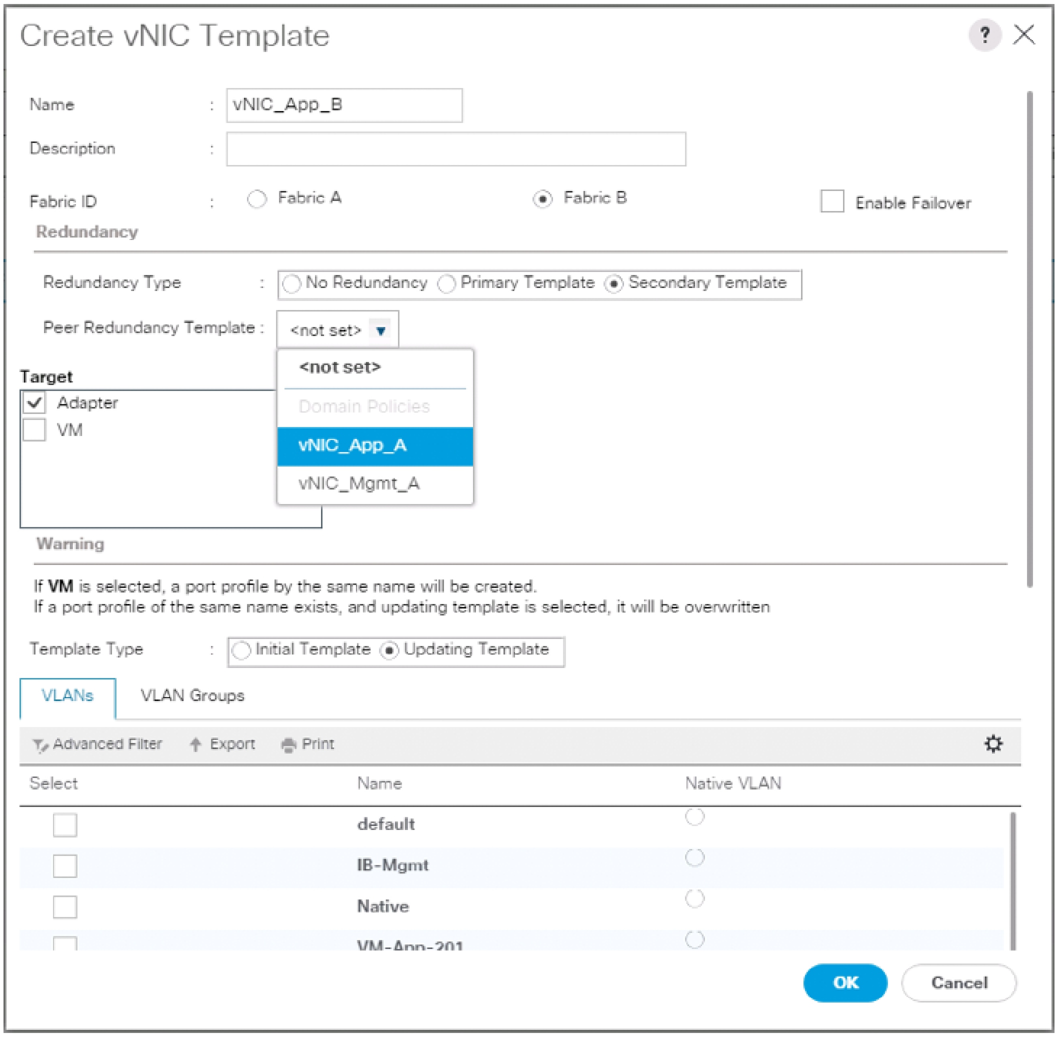

Create vNIC Template

Name : vNIC_App_B

Description :

Fabric ID : ○ Fabric A ● Fabric B □ Enable Failover

Redundancy

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

Peer Redundancy Template : <not set>

<not set>

Domain Policies

vNIC_App_A

vNIC_Mgmt_A

Target

☑ Adapter

□ VM

Warning

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

Template Type : ○ Initial Template ● Updating Template

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Native	○
□	VM-App-201	○

10. In the MAC Pool list, select `MAC_Pool_B`.

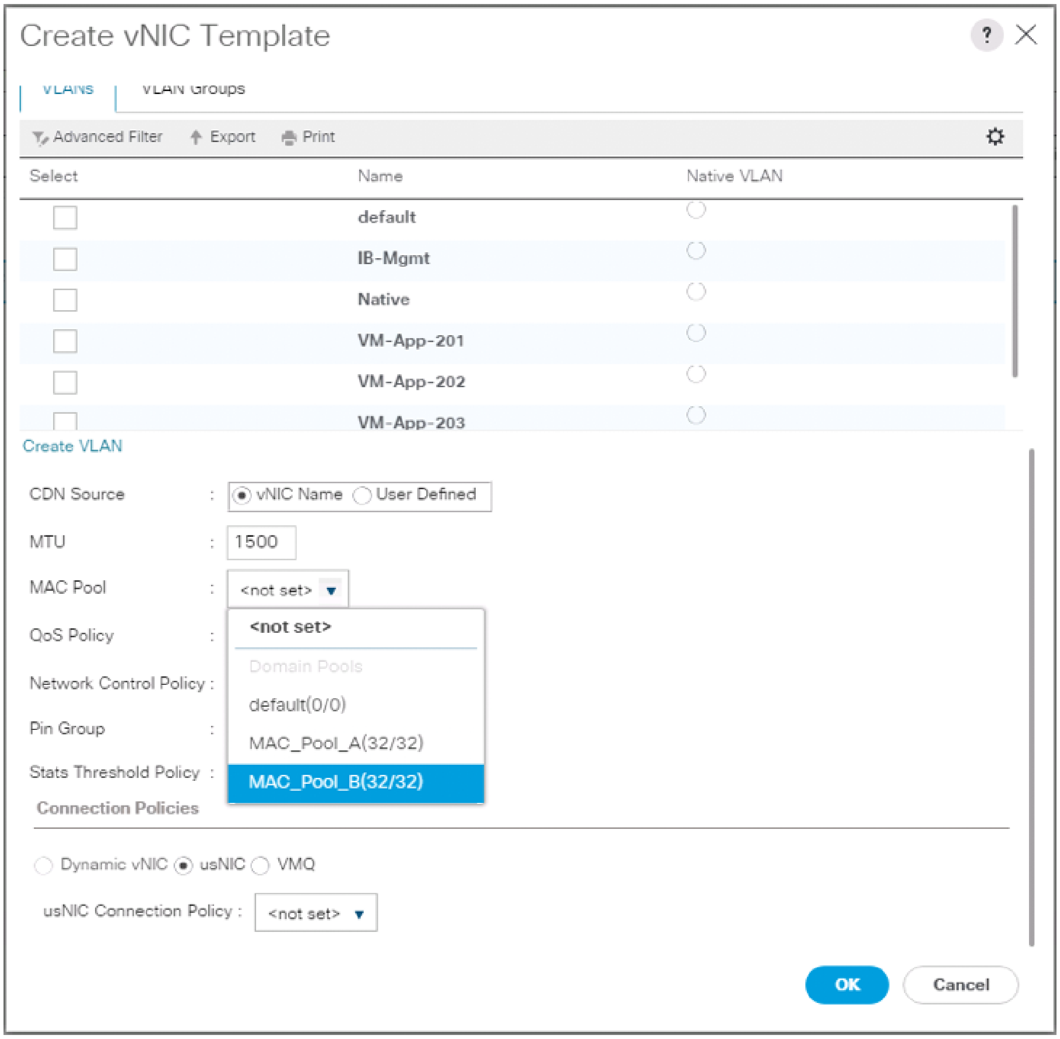

Create vNIC Template

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Native	○
□	VM-App-201	○
□	VM-App-202	○
□	VM-App-203	○

CDN Source : ● vNIC Name ○ User Defined

MTU : 1500

MAC Pool : <not set>

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

○ Dynamic vNIC ● usNIC ○ VMQ

usNIC Connection Policy : <not set>

11. Click OK to create the vNIC template.

12. Click OK.

Set Jumbo Frames in Cisco UCS Fabric

These steps are unnecessary for the Cisco UCS 6454 FIs as they default to jumbo frames but left in as reference for other UCS FI platforms.

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

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

2. Select LAN > LAN Cloud > QoS System Class.

3. In the right pane, click the General tab.

4. On the Best Effort row, enter `9216` in the box under the MTU column.

5. Click Save Changes in the bottom of the window.

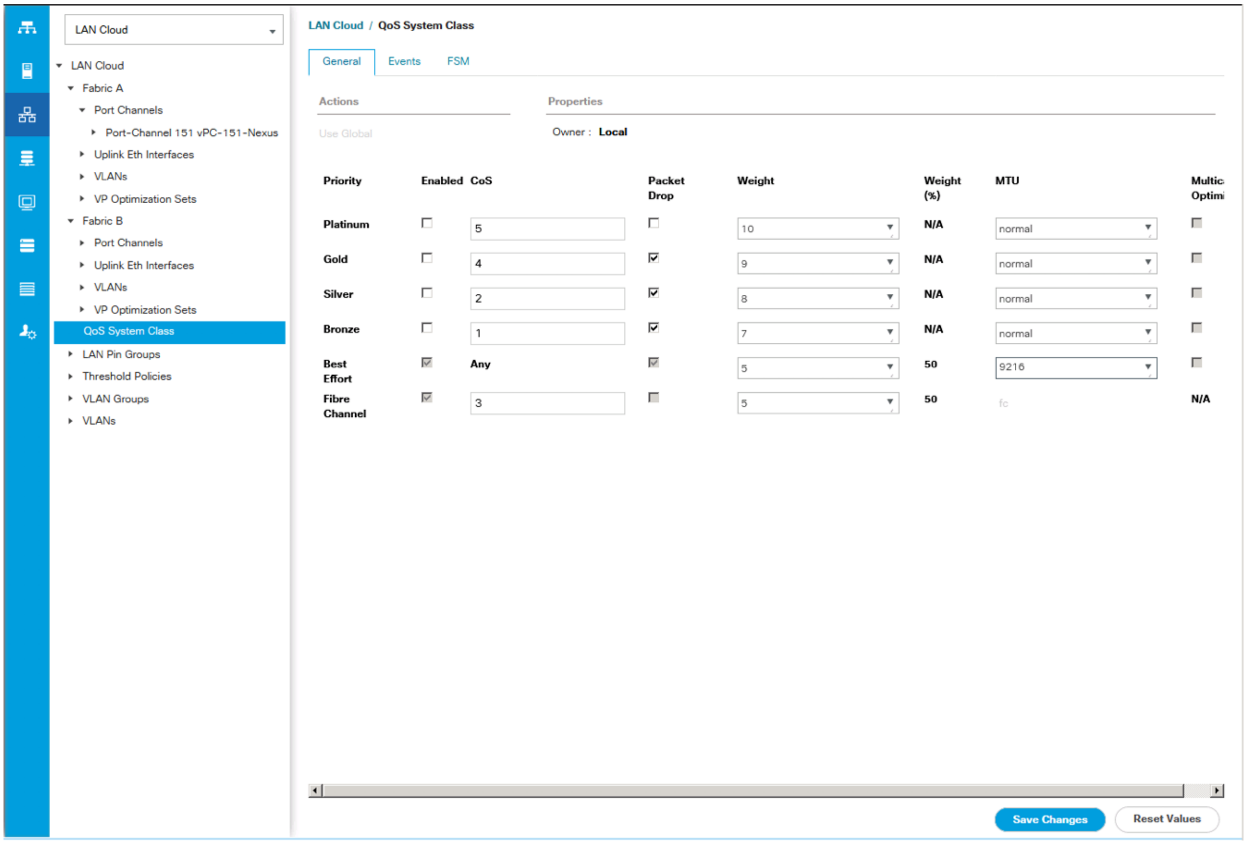

Owner : **Local**

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multic Optimi
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	9216	☐
Fibre Channel	☑	3	☐	5	50	fc	N/A

6. Click OK

Create LAN Connectivity Policy

To configure the necessary Fibre Channel Infrastructure LAN Connectivity Policy, follow these steps:

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

2. Select LAN > Policies > root.

3. Right-click LAN Connectivity Policies.

4. Select Create LAN Connectivity Policy.

5. Enter `FC-LAN-Policy` 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-Mgmt-A` as the name of the vNIC.

The numeric prefix of “00-“ and subsequent increments on the later vNICs are used in the vNIC naming to force the device ordering through Consistent Device Naming (CDN). Without this, some operating systems might not respect the device ordering that is set within Cisco UCS.

8. Select the Use vNIC Template checkbox.

9. In the vNIC Template list, select vNIC_Mgmt_A.

10. In the Adapter Policy list, select VMWare.

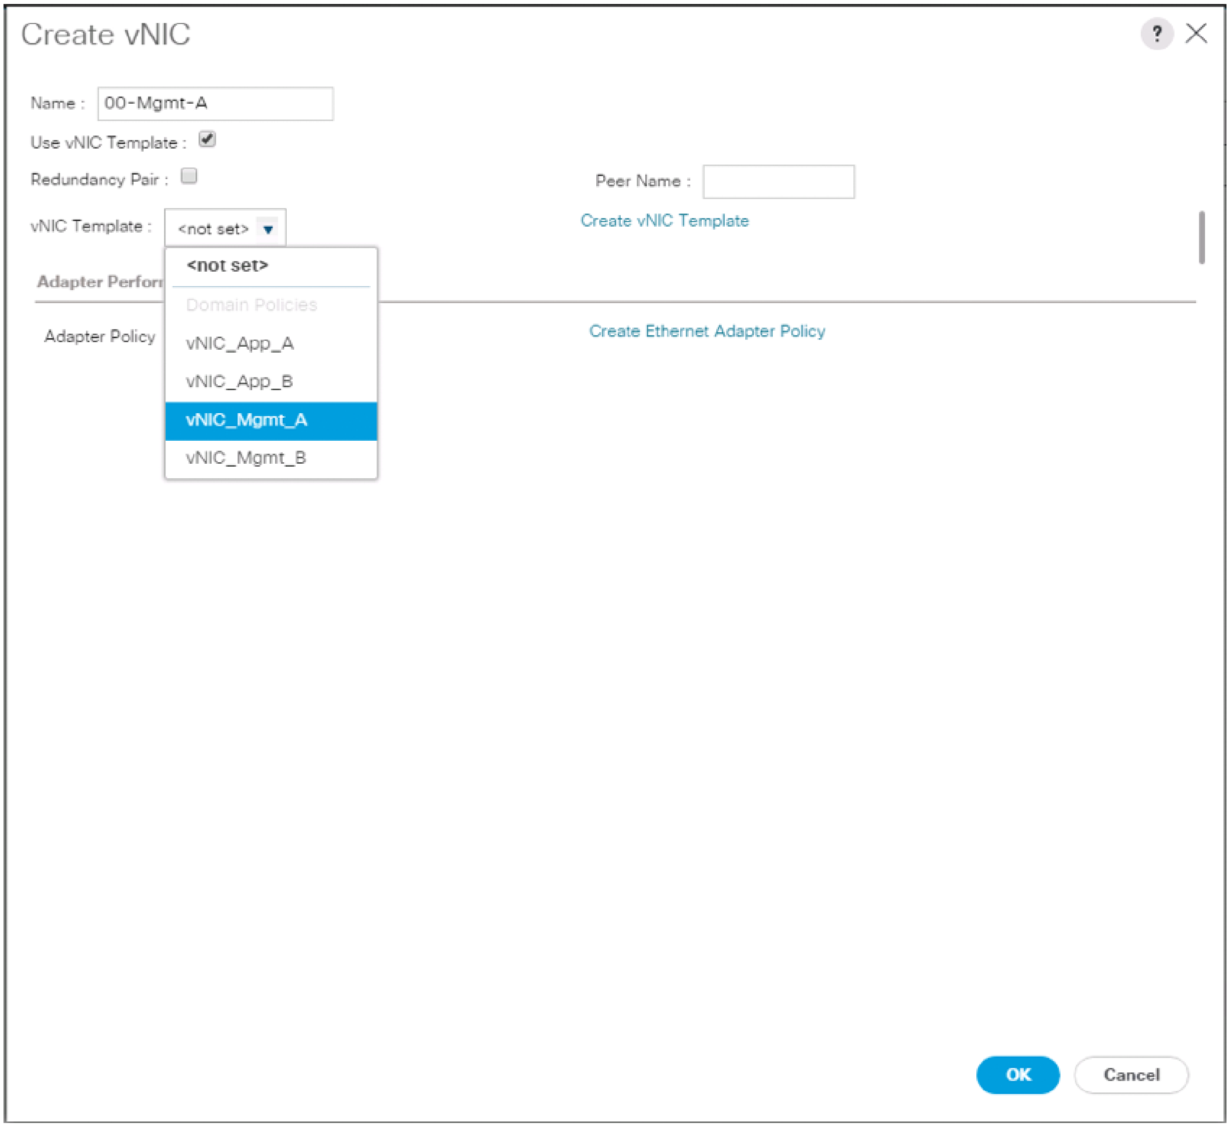

Create vNIC

Name : 00-Mgmt-A

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : <not set>

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-Mgmt-B` as the name of the vNIC.

14. Select the Use vNIC Template checkbox.

15. In the vNIC Template list, select vNIC_Mgmt_B.

16. In the Adapter Policy list, select VMWare.

17. Click OK to add the vNIC to the policy.

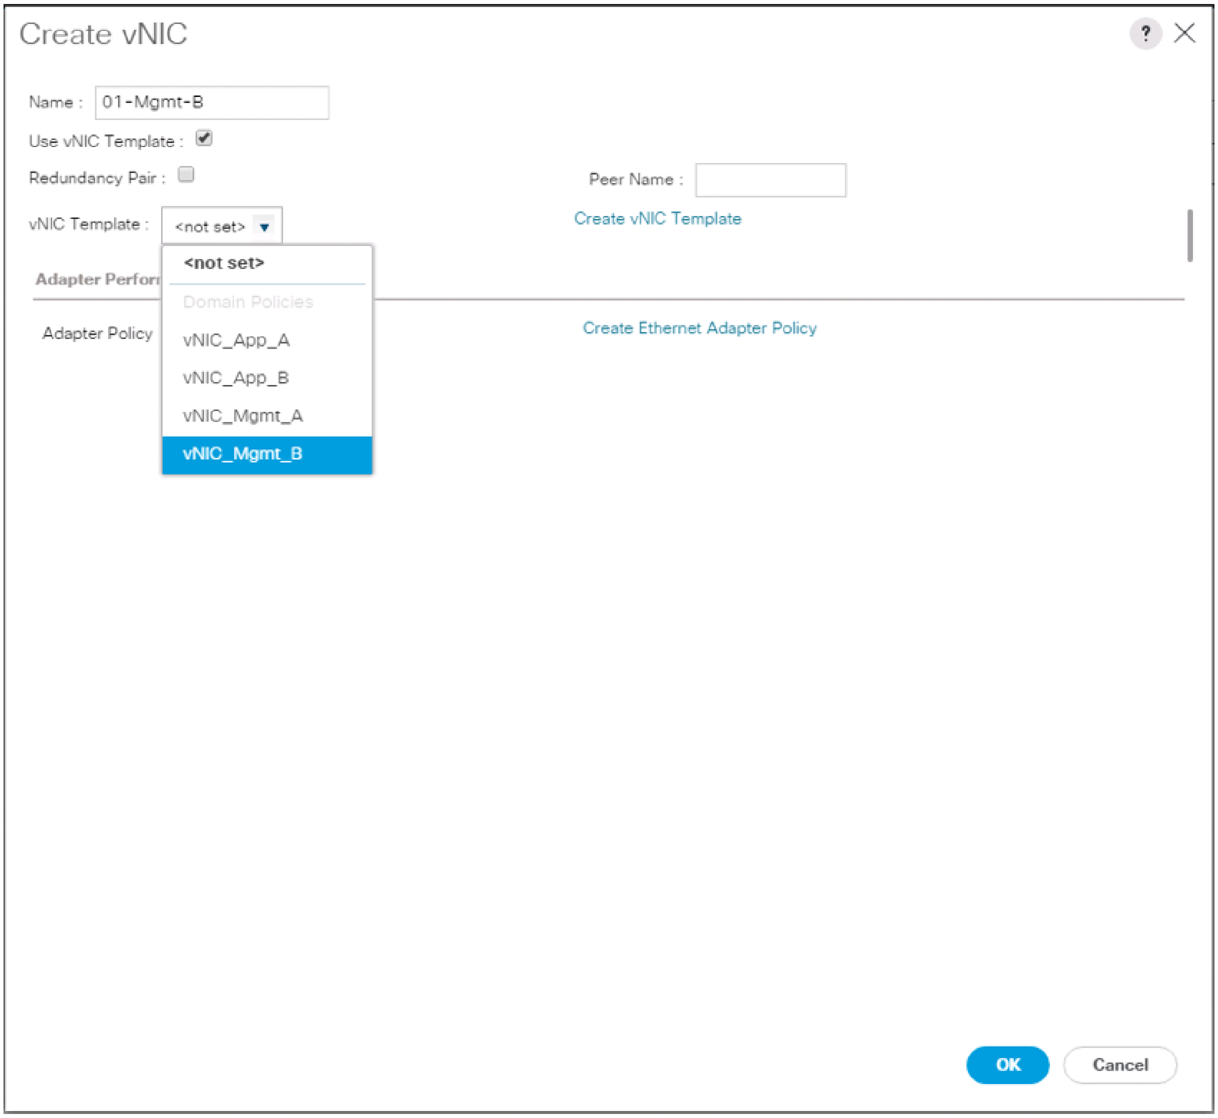

Create vNIC

Name : 01-Mgmt-B

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : <not set>

18. Click the upper Add button to add a vNIC.

19. In the Create vNIC dialog box, enter `02-App-A` as the name of the vNIC.

20. Select the Use vNIC Template checkbox.

21. In the vNIC Template list, select vNIC_App_A.

22. In the Adapter Policy list, select VMWare.

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

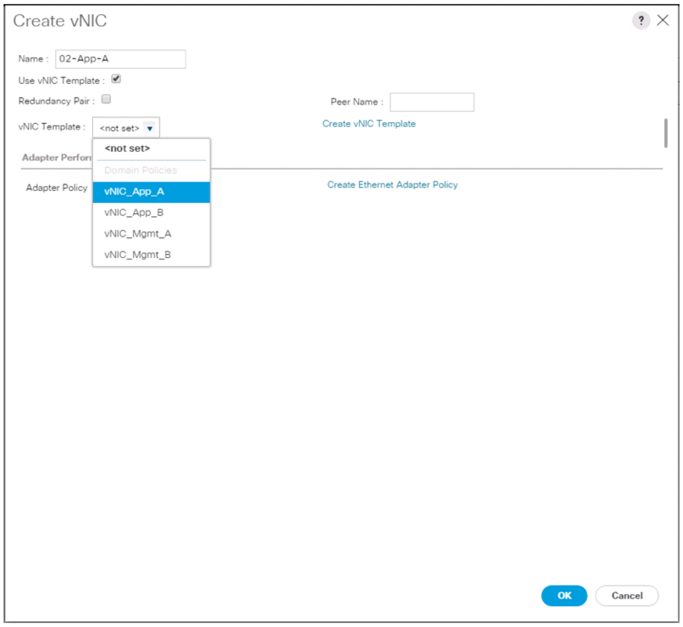

Create vNIC

Name : 02-App-A

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : <not set>

24. Click the upper Add button to add a vNIC to the policy.

25. In the Create vNIC dialog box, enter `03-App-B` as the name of the vNIC.

26. Select the Use vNIC Template checkbox.

27. In the vNIC Template list, select vNIC_App_B.

28. In the Adapter Policy list, select VMWare.

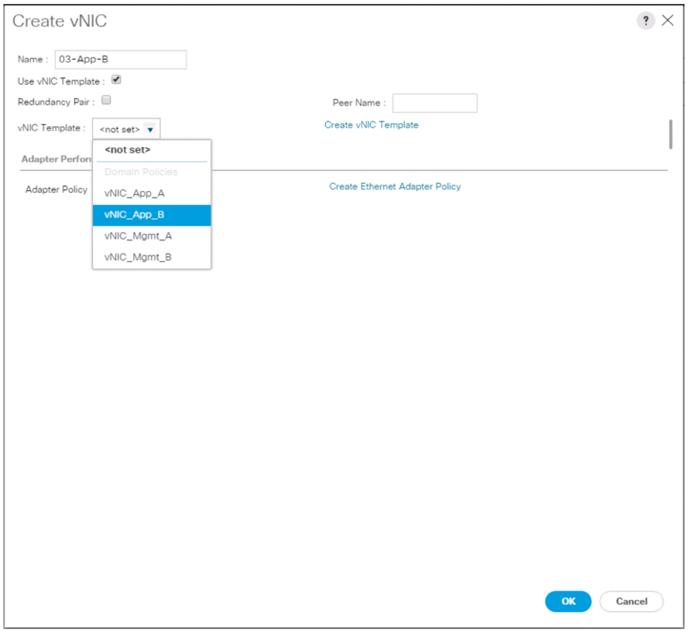

Create vNIC

Name : 03-App-B

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : <not set>

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

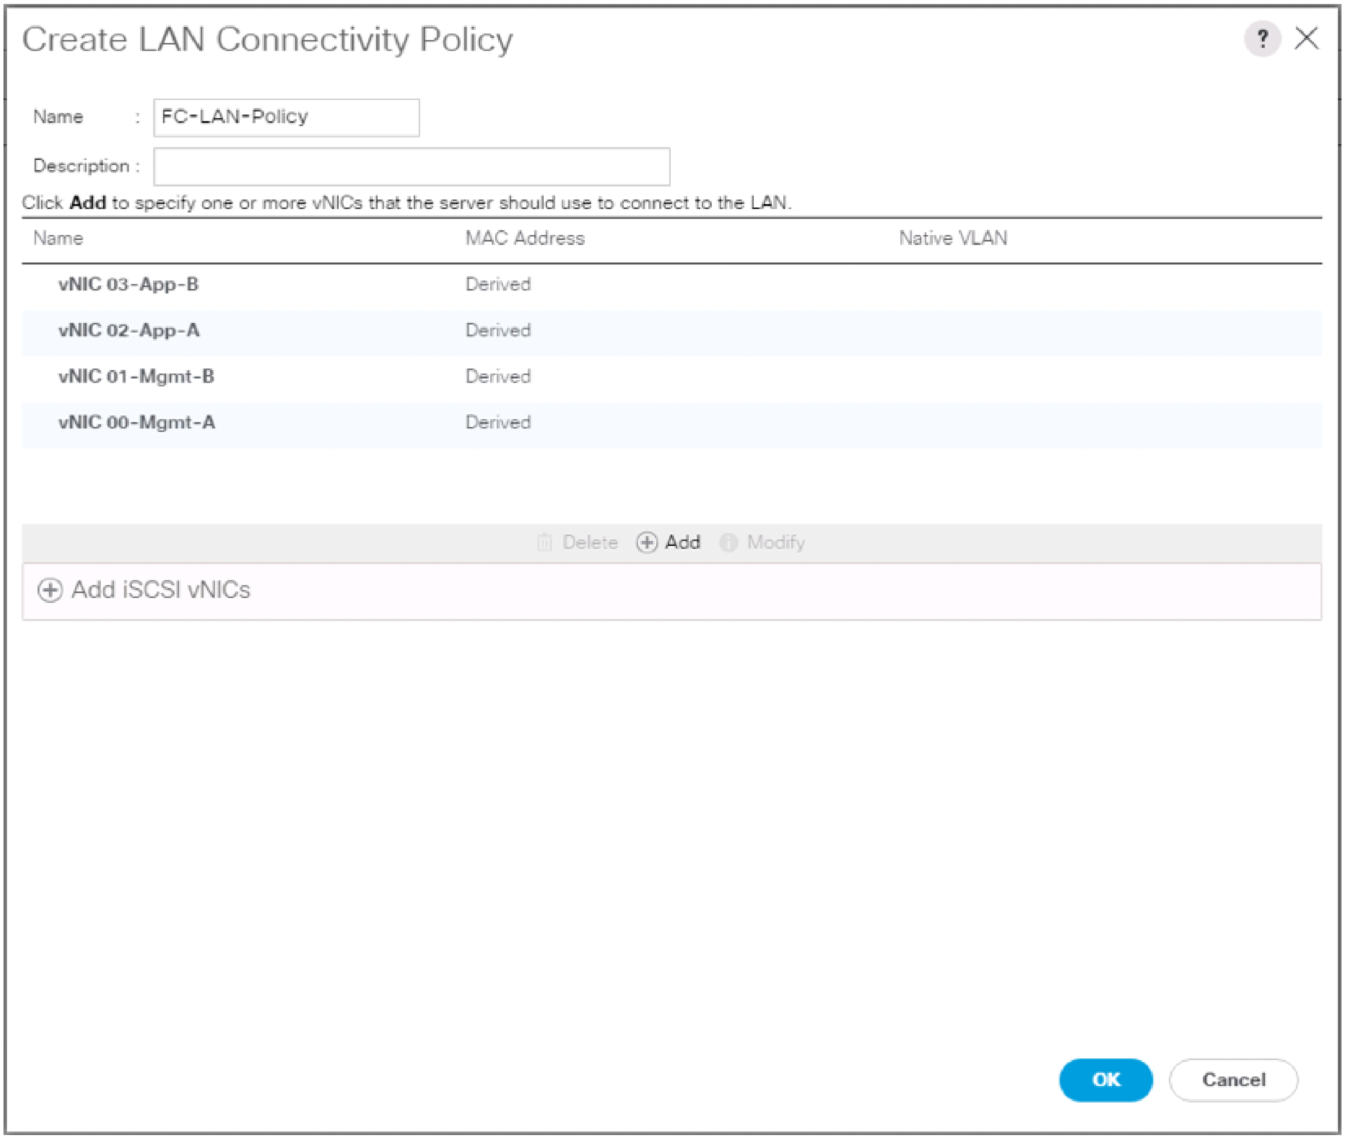

Create LAN Connectivity Policy

Name : FC-LAN-Policy

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-App-B	Derived	
vNIC 02-App-A	Derived	
vNIC 01-Mgmt-B	Derived	
vNIC 00-Mgmt-A	Derived	

30. Click OK to create the LAN Connectivity Policy.

31. Click OK.

Configure FC SAN Connectivity

These Fibre Channel configuration steps will enable the provisioning of volumes to be used as datastores by the vSphere hosts, and the creation of Cisco UCS Service Profiles that will be configured to boot from Fibre Channel LUNs.

Create VSANs

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 tab in the navigation pane.

In this procedure, two VSANs are created.

2. Select SAN > Storage Cloud.

3. Right-click VSANs.

4. Select Create Storage VSAN.

5. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A

6. Select **Enabled** for FC Zoning.

7. Select Fabric A.

8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. It is recommended use the same ID for both parameters and to use something other than 1.

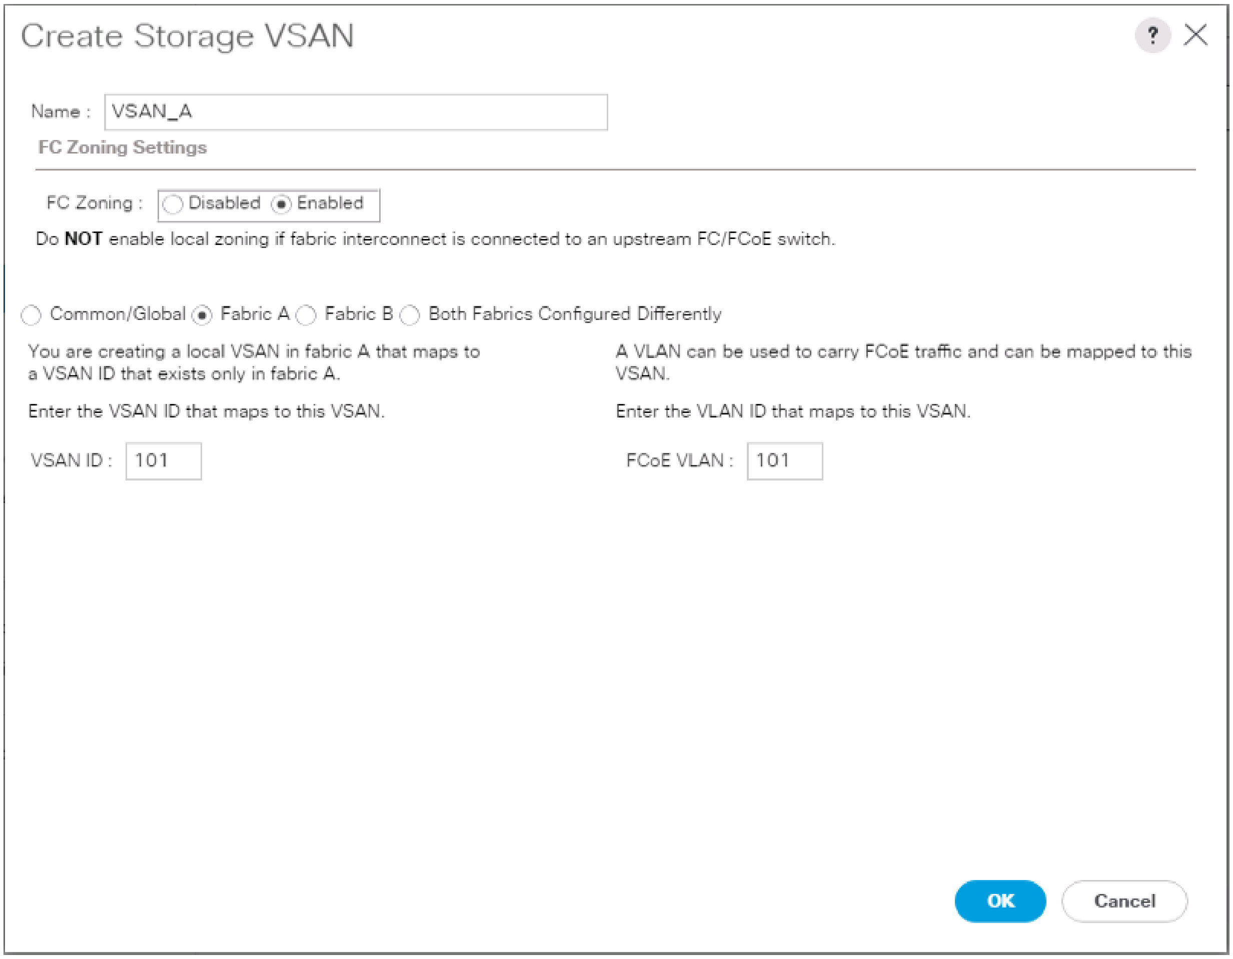

Create Storage VSAN

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 Storage Cloud, right-click VSANs.

11. Select Create Storage VSAN.

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

13. Select **Enabled** for FC Zoning.

14. Select Fabric B.

15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. It is recommended use the same ID for both parameters and to use something other than 1.

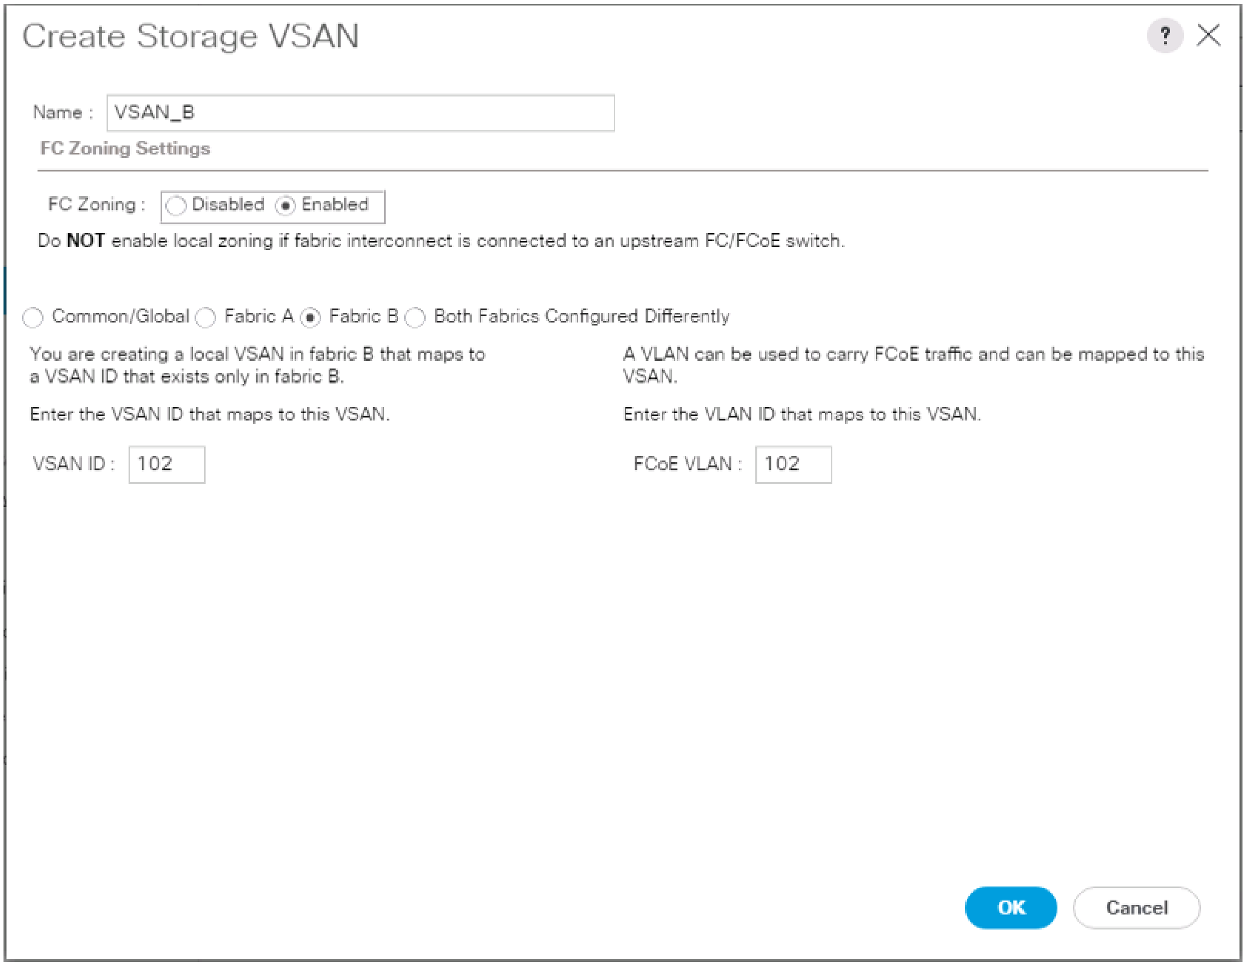

Create Storage VSAN

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.

Assign VSANs to FC Storage Ports

To assign storage VSANs to FC Storage Ports, follow these steps:

1. In Cisco UCS Manager, click SAN.

2. Select SAN > Storage Cloud.

3. Expand Fabric A and Storage FC Interfaces.

4. Select the first FC Interface.

5. For User Label, enter the storage controller name and port.

6. From the drop-down list select VSAN VSAN_A.

7. Adjust Admin Speed as appropriate for connections.

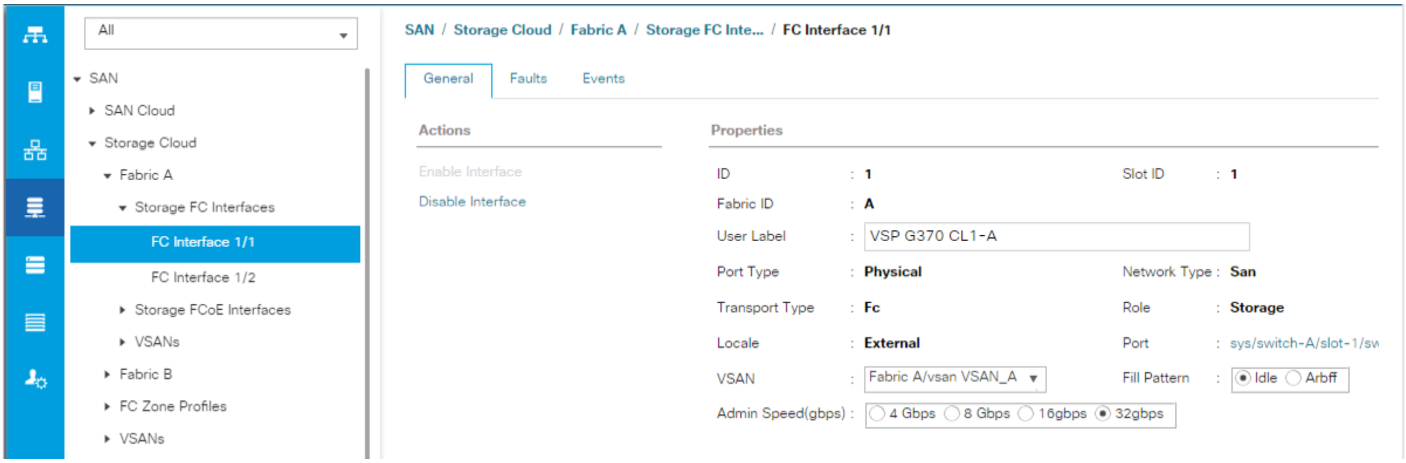

ID : 1 Slot ID : 1
Fabric ID : A
User Label : VSP G370 CL1-A
Port Type : Physical Network Type : San
Transport Type : Fc Role : Storage
Locale : External Port : sys/switch-A/slot-1/sw
VSAN : Fabric A/vsan VSAN_A Fill Pattern : ● Idle ○ Arbff
Admin Speed(gbps) : ○ 4 Gbps ○ 8 Gbps ○ 16gbps ● 32gbps

8. Click Save Changes and OK.

9. Repeat steps 5-8 for the second FC Interface.

10. Within Storage Cloud, expand Fabric B and Storage FC Interfaces.

11. Repeat steps 4 through 9 for the Fabric B interfaces, using VSAN_B.

If the UCS FC ports show as error disabled at this point due to a timing of operations, a disable and subsequent enable of the error disabled port will be needed.

Create vHBA Templates

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

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

2. Select Policies > root.

3. Right-click vHBA Templates.

4. Select Create vHBA Template.

5. Enter `vHBA_Template_A` as the vHBA template name.

6. Keep Fabric A selected.

7. Leave Redundancy Type as No Redundancy.

8. Select VSAN_A.

9. Leave Initial Template as the Template Type.

10. Select WWPN_Pool_A as the WWPN Pool.

11. Click OK to create the vHBA template.

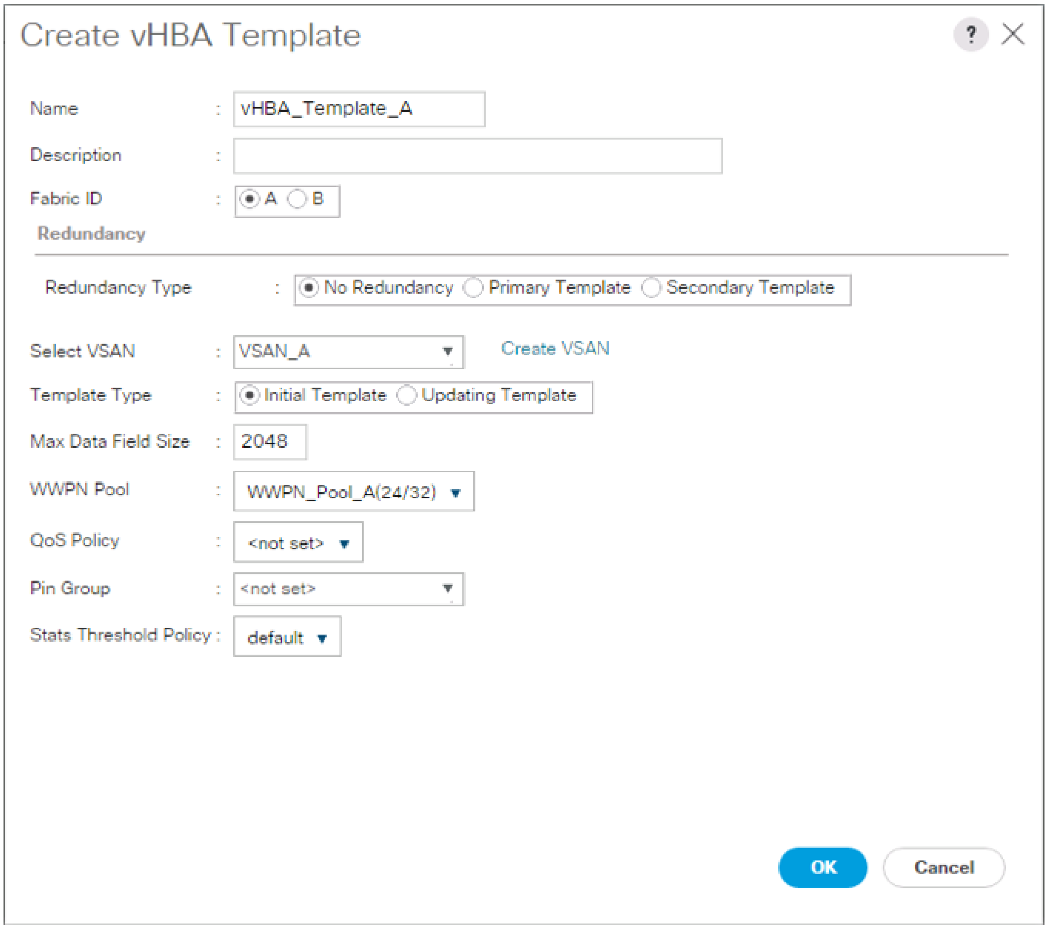

Create vHBA Template

Name : vHBA_Template_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(24/32)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

12. Click OK.

13. Right-click vHBA Templates.

14. Select Create vHBA Template.

15. Enter `vHBA_Template_B` as the vHBA template name.

16. Select Fabric B as the Fabric ID.

17. Leave Redundancy Type as 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.

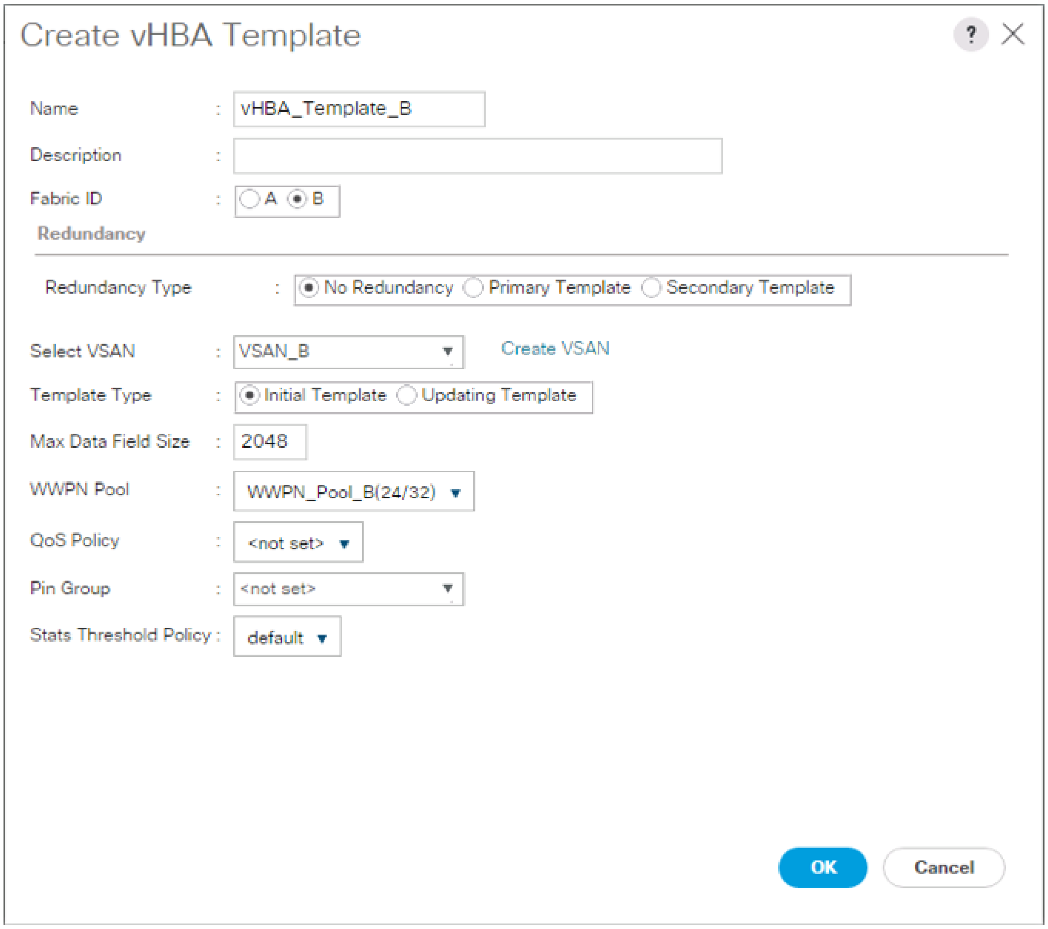

Create vHBA Template

Name : vHBA_Template_B

Description :

Fabric ID : ○ A ● B

Redundancy

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

Select VSAN : VSAN_B

Template Type : ● Initial Template ○ Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN_Pool_B(24/32)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

22. Click OK.

Create SAN Connectivity Policy

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > Policies > root.
3. Right-click SAN Connectivity Policies.
4. Select Create SAN Connectivity Policy.
5. Enter `Infra-SAN-Policy` 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 `Fabric-A` as the name of the vHBA.
9. Select the Use vHBA Template checkbox.

10. Leave Redundancy Pair unselected.

11. In the vHBA Template list, select vHBA_Template_A.

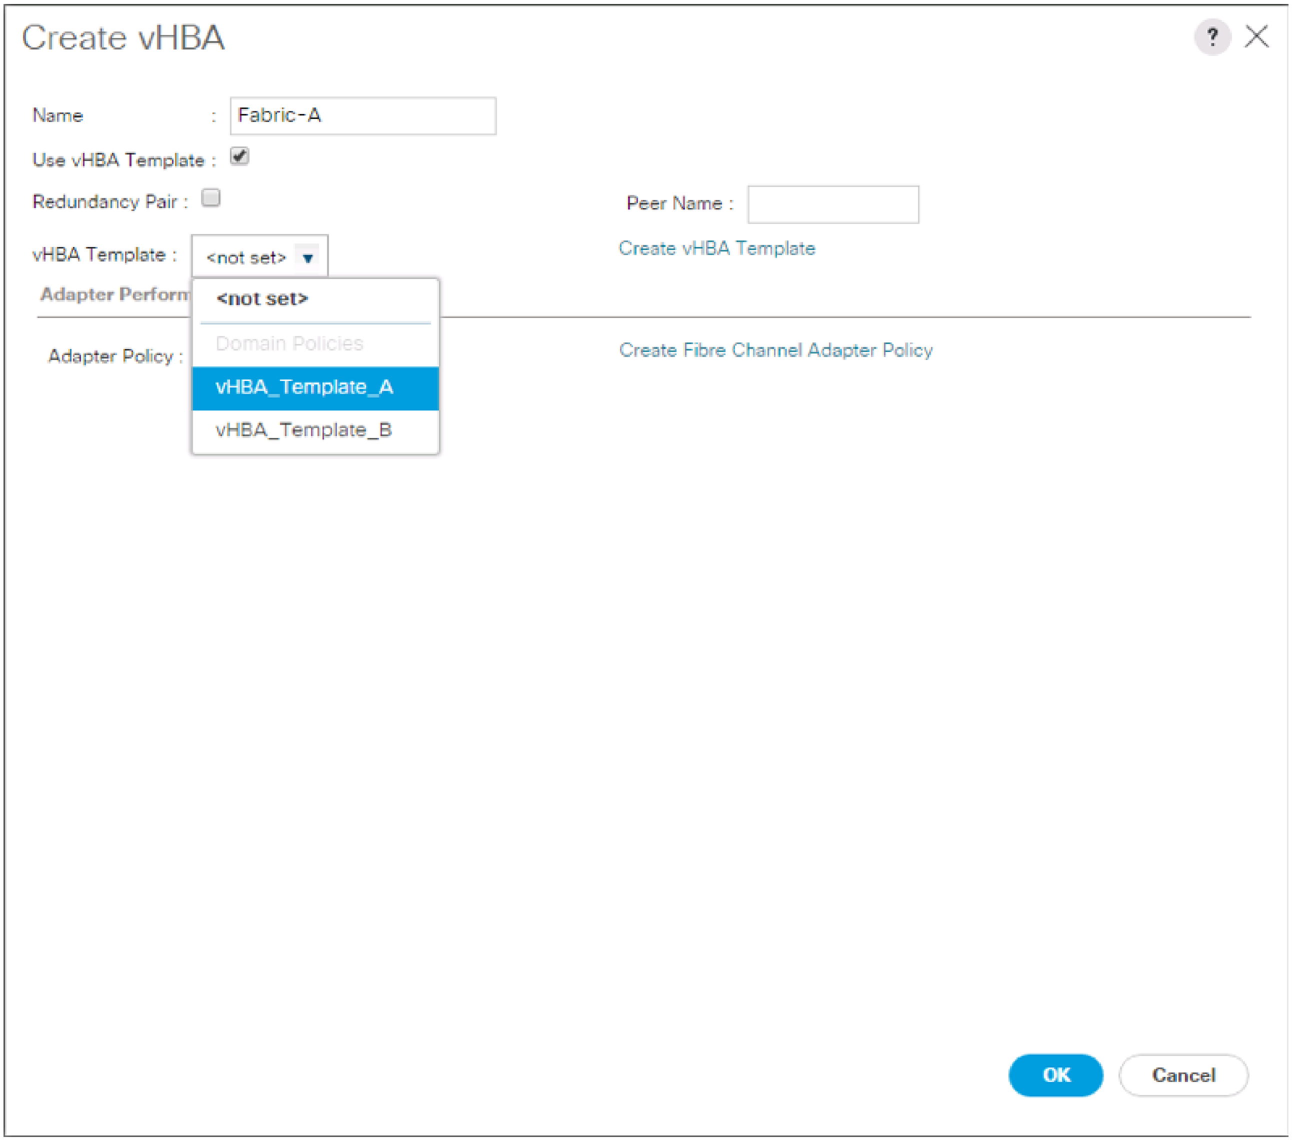

Create vHBA

Name : Fabric-A

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : <not set>

12. In the Adapter Policy list, select VMWare.

13. Click OK.

14. Click Add to add a second vHBA.

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

16. Select the Use vHBA Template checkbox.

17. Leave Redundancy Pair unselected.

18. In the vHBA Template list, select vHBA_Template_B.

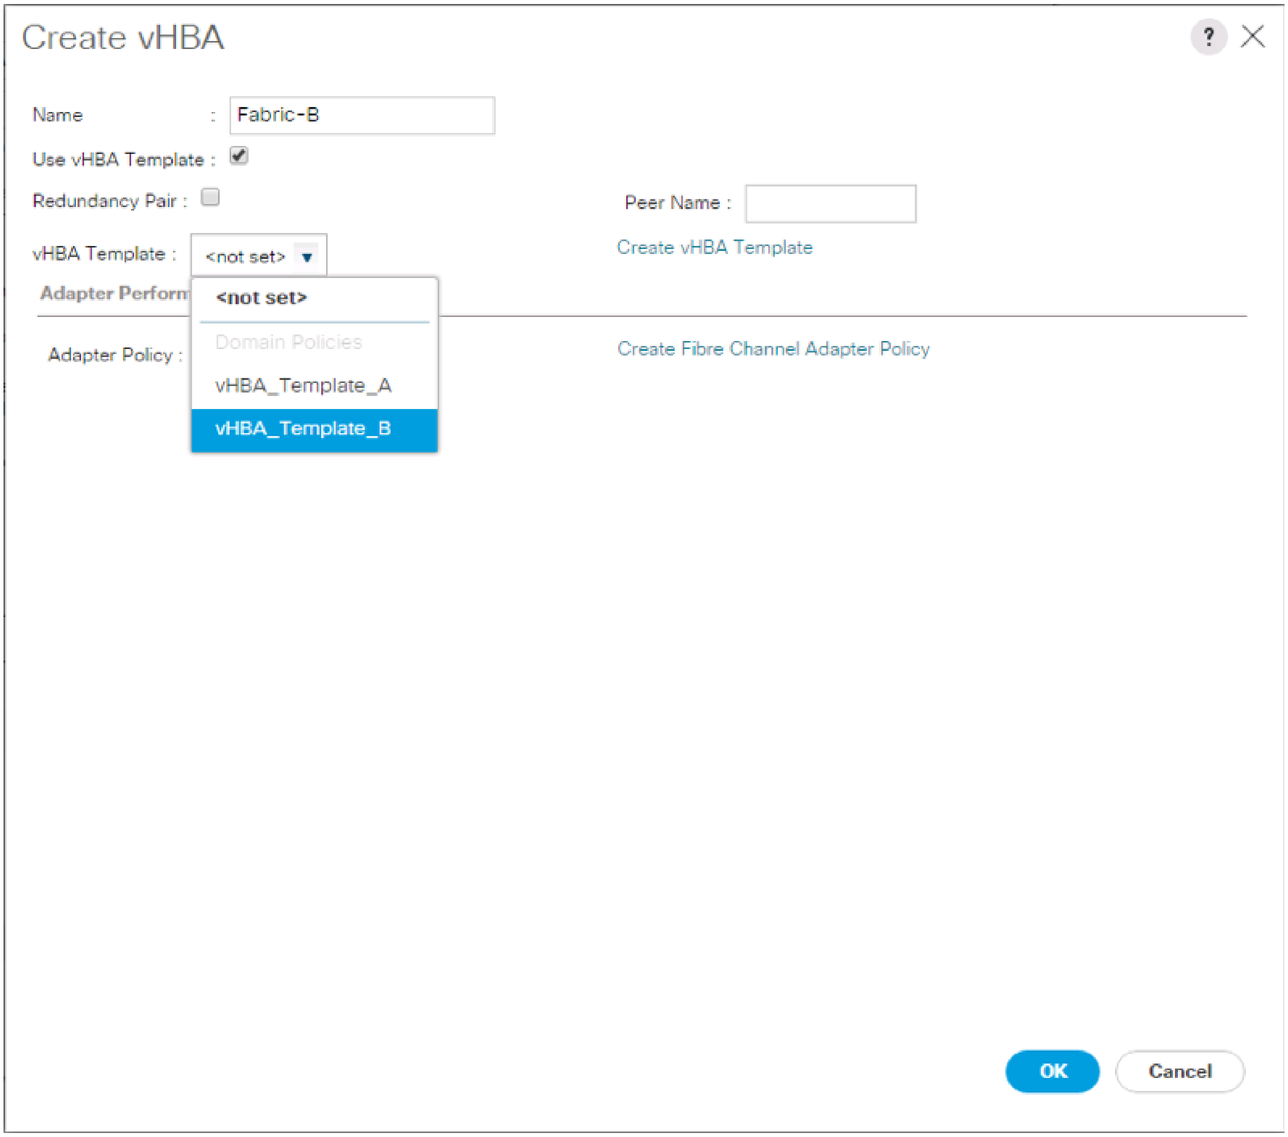

Create vHBA

Name : Fabric-B

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : <not set>

Adapter Policy :

19. In the Adapter Policy list, select VMWare.

20. Click OK.

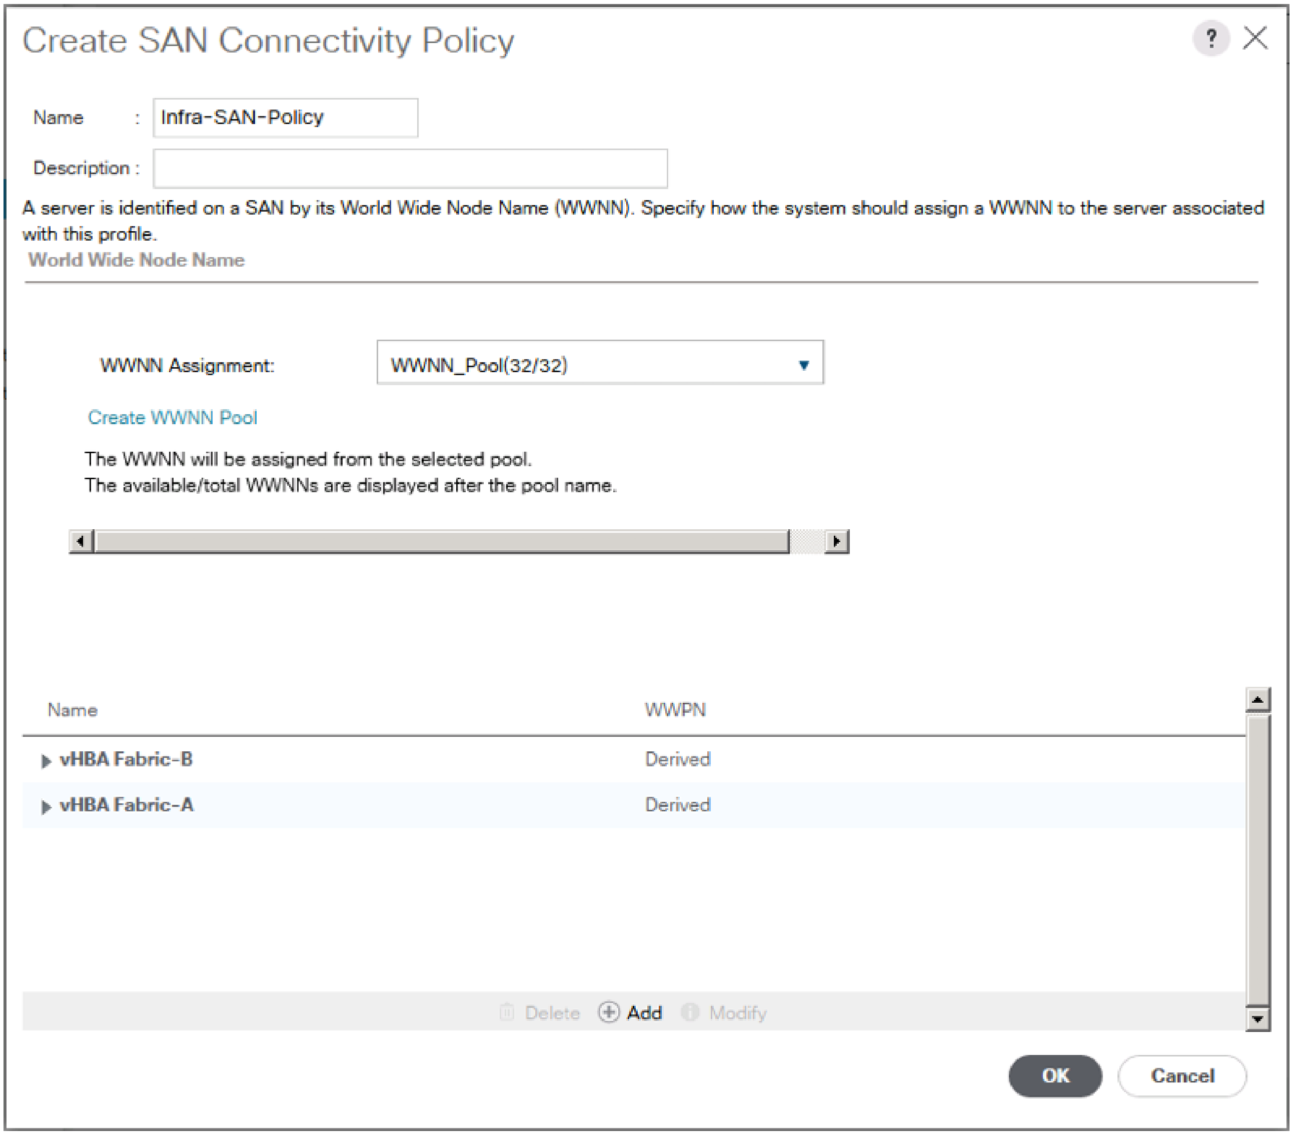

Create SAN Connectivity Policy

Name : Infra-SAN-Policy

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)

Create WWNN Pool

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

Name	WWPN
vHBA Fabric-B	Derived
vHBA Fabric-A	Derived

21. Click OK to create the SAN Connectivity Policy.

22. Click OK to confirm creation.

Create Boot Policy

The VSP G370 target WWPN will need to be collected to provide the Cisco UCS Boot Policy.

These target WWPN can be collected directly from the VSP but running the show flogi database command from each FI will be fairly quick provided there is clear identification of the port cabling from the VSP ports to the FI ports.

Table 10 VSP G370 to FI Port Information Carried Forward

Storage	Local Port	Connection	FI	FI Port
Hitachi VSP G370	CL 1-A	32Gb FC	Cisco UCS 6454 FI A	FC 1/1
	CL 2-B	32Gb FC	Cisco UCS 6454 FI A	FC 1/2
	CL 3-B	32Gb FC	Cisco UCS 6454 FI B	FC 1/11

Storage	Local Port	Connection	FI	FI Port
	CL 4-A	32Gb FC	Cisco UCS 6454 FI B	FC 1/12

Using this table, it is possible to get the expected local port (VSP) to the FI port values. With this information, the WWPN can be pulled out of the flogi to port connections on the respective FI.

Running the sh flogi database command on UCS FI A:

```
UCS-6454-A# connect nxos
UCS-6454-A(nx-os)# sh flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/1            101   0x840040  50:06:0e:80:12:c9:9a:00 50:06:0e:80:12:c9:9a:00
fc1/2            101   0x840060  50:06:0e:80:12:c9:9a:11 50:06:0e:80:12:c9:9a:11
```

Running the sh flogi database command on UCS FI B:

```
UCS-6454-B# connect nxos
UCS-6454-B(nx-os)# sh flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/1            102   0x860040  50:06:0e:80:12:c9:9a:21 50:06:0e:80:12:c9:9a:21
fc1/2            102   0x860060  50:06:0e:80:12:c9:9a:30 50:06:0e:80:12:c9:9a:30
```

Find the appropriate VSP G370 local ports for each fabric and record the values to be used for Primary and Secondary Boot Targets. In the example lab environment flogi output, the FI Interface (FI Port) values in the previous table for this fabric have been cross referenced, and the WWPN(Port Name) for these interfaces are recorded.

Table 11 Fabric A Boot Targets for the VSP G370

	6454 FI Interface	Example Local Port	Target Role	WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
VSP G370 Controller 1	1/1	CL 1-A	Primary/ VMFS	50:06:0e:80:12:c9:9a:00	
VSP G370 Controller 2	1/2	CL 2-B	Secondary/ VMFS	50:06:0e:80:12:c9:9a:11	

Repeat these steps for the VSP G370 Fabric B Primary and Secondary Boot Targets:

Table 12 Fabric B Boot Targets for the VSP G370

	6454 FI Interface	Example Local Port	Target Role	WWN/WWPN Example Environment	WWN/WWPN Customer Environment
VSP G370 Controller 1	1/1	CL 3-B	Primary Boot/ VMFS	50:06:0e:80:12:c9:9a:21	
VSP G370 Controller 2	1/2	CL 4-A	Secondary Boot/ VMFS	50:06:0e:80:12:c9:9a:30	

To create boot policies for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-FC-G370-A` as the name of the boot policy.

6. Optional: Enter a description for the boot policy.

> Do not select the Reboot on Boot Order Change checkbox.

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

8. Expand the vHBAs drop-down list and select Add SAN Boot.

9. In the Add SAN Boot dialog box, enter `Fabric-A` in the vHBA field.

10. Confirm that Primary is selected for the Type option.

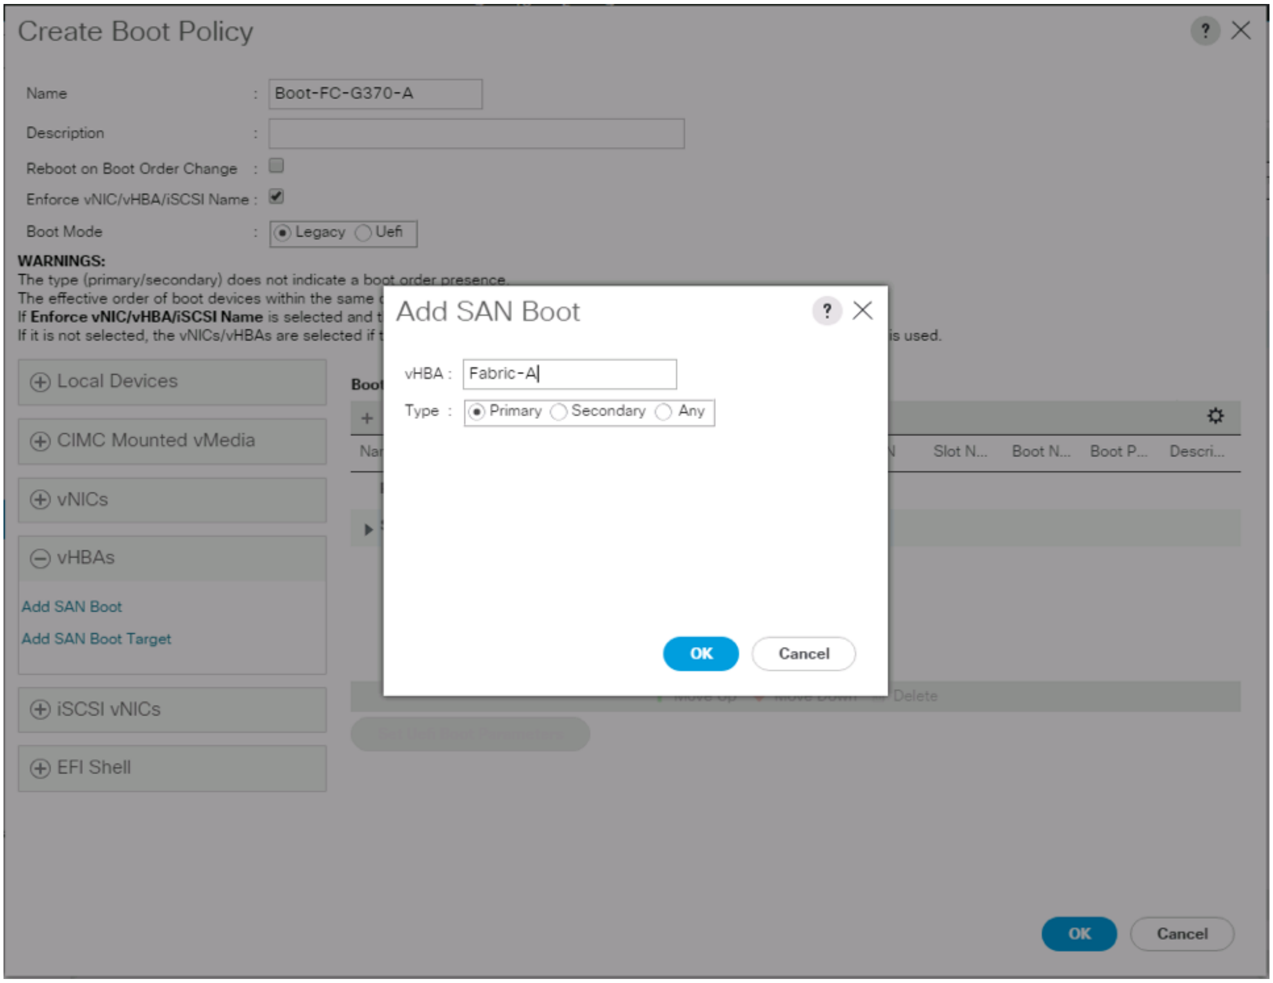

Create Boot Policy

Name : Boot-FC-G370-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
If **Enforce vNIC/vHBA/iSCSI Name** is selected and t ... is used.
If it is not selected, the vNICs/vHBAs are selected if t

Add SAN Boot

vHBA : Fabric-A

Type : ● Primary ○ Secondary ○ Any

11. Click OK to add the SAN boot initiator.

12. From the vHBA drop-down list, select Add SAN Boot Target.

13. Leave `0` as the value for Boot Target LUN.

14. Enter the WWPN for `Controller1 (CL 1A)` recorded in Table 11 .

15. Select Primary for the SAN boot target type.

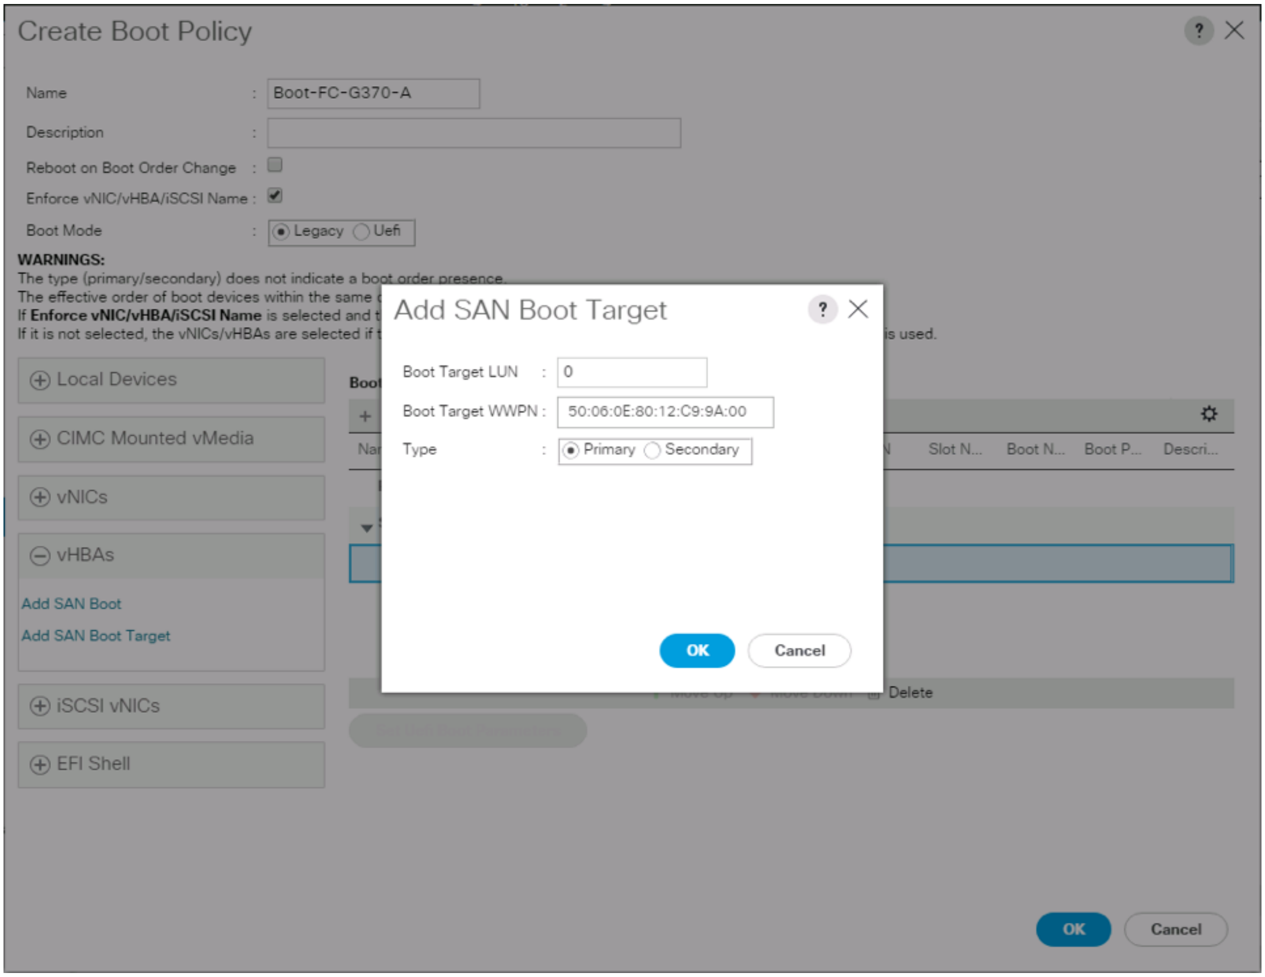

Create Boot Policy

Name : Boot-FC-G370-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
If **Enforce vNIC/vHBA/iSCSI Name** is selected and t
If it is not selected, the vNICs/vHBAs are selected if t … is used.

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 50:06:0E:80:12:C9:9A:00

Type : ● Primary ○ Secondary

16. Click OK to add the SAN boot target.

17. From the vHBA drop-down list, select Add SAN Boot Target.

18. Leave 0 as the value for Boot Target LUN.

19. Enter the WWPN for `Controller2 (CL 2B)` recorded in Table 11 .

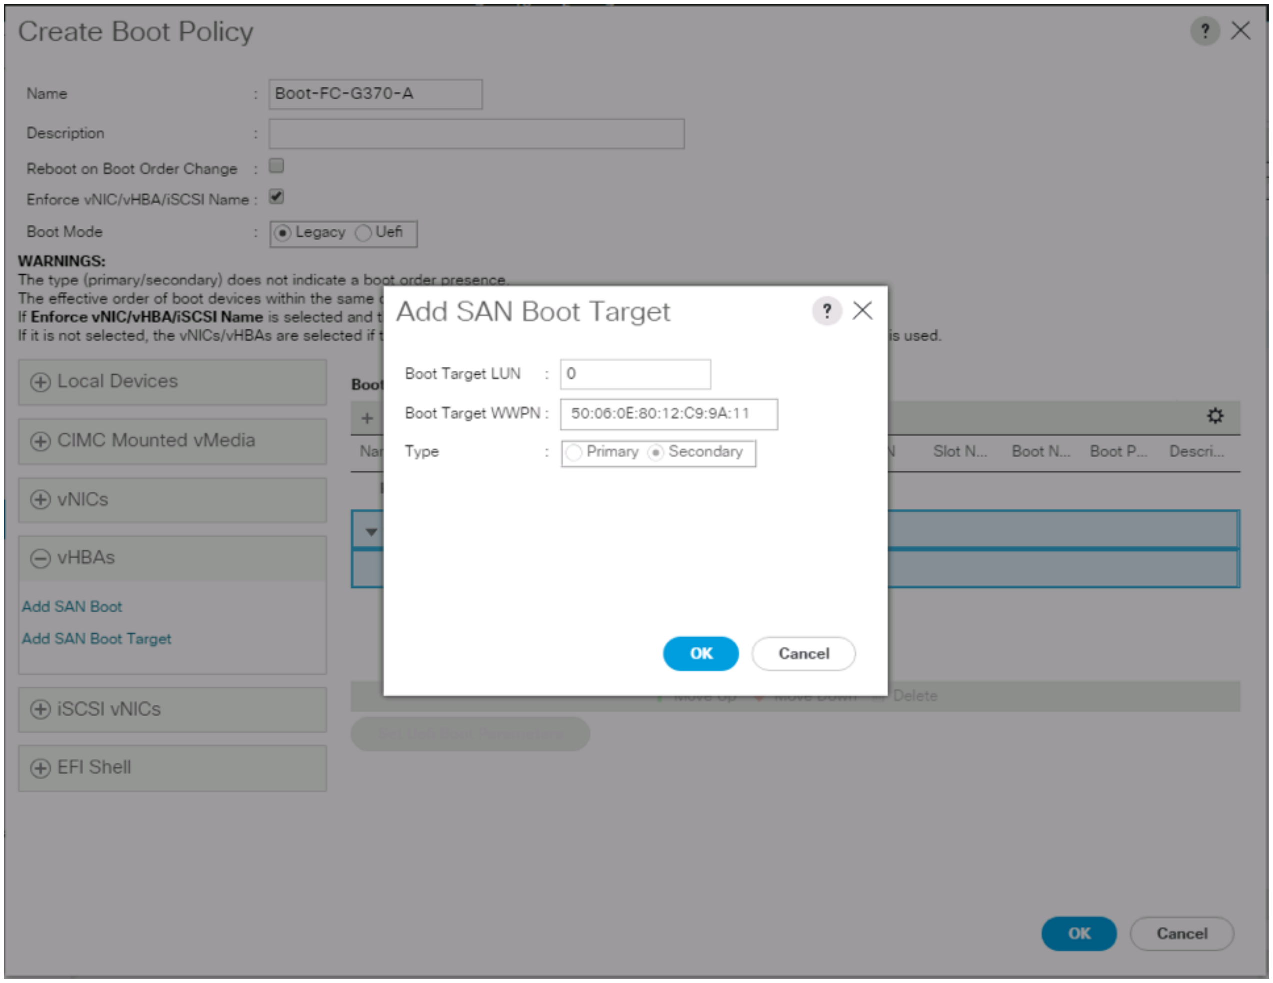

Create Boot Policy

Name : Boot-FC-G370-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
If **Enforce vNIC/vHBA/iSCSI Name** is selected and t
If it is not selected, the vNICs/vHBAs are selected if t … is used.

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 50:06:0E:80:12:C9:9A:11

Type : ○ Primary ● Secondary

20. Click OK to add the SAN boot target.

21. From the vHBA drop-down list, select Add SAN Boot.

22. In the Add SAN Boot dialog box, enter `Fabric-B` in the vHBA box.

The SAN boot type should automatically be set to Secondary and the Type option should be unavailable.

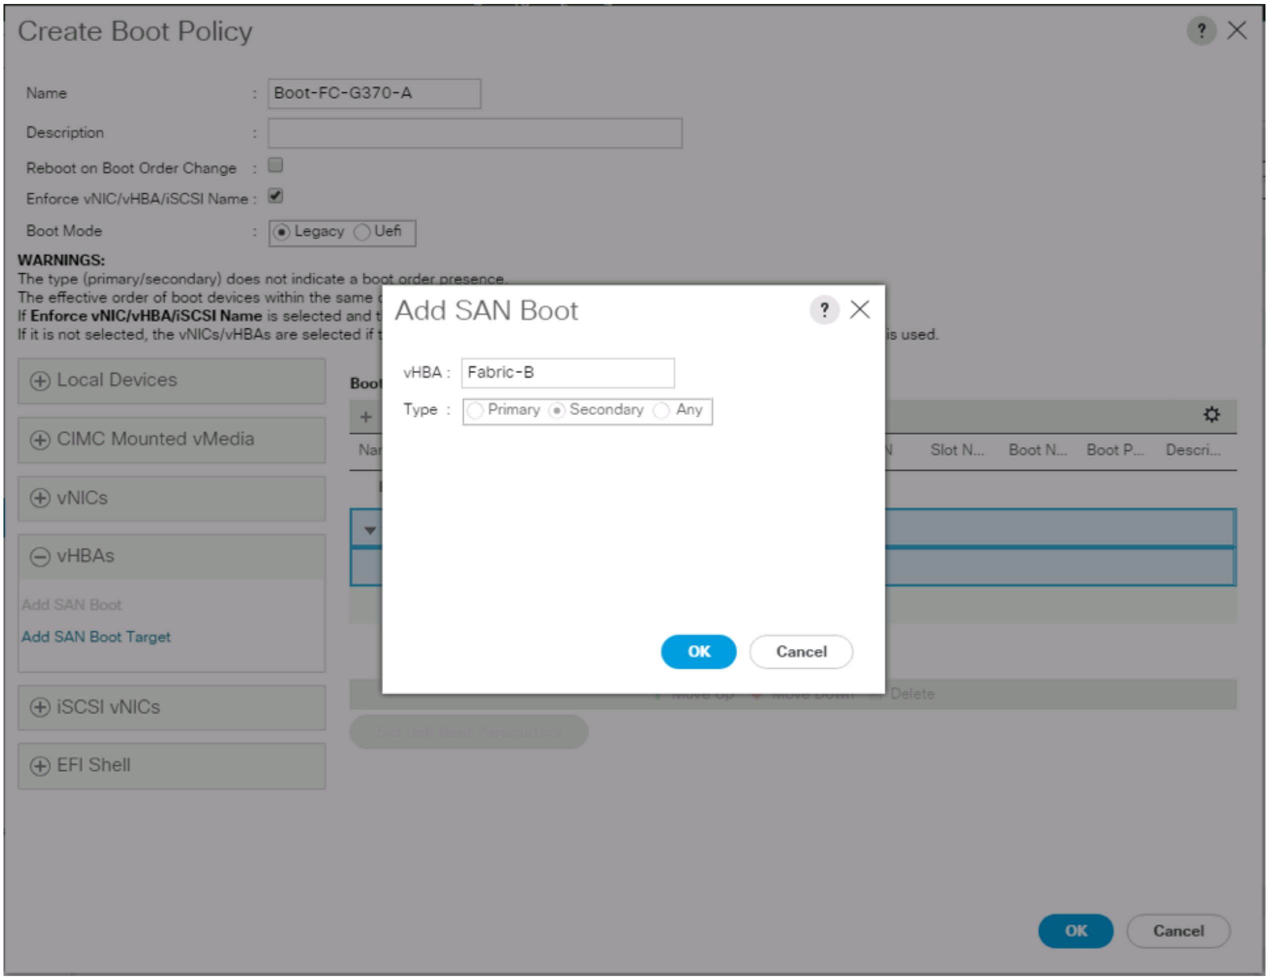

Create Boot Policy

Name : Boot-FC-G370-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 c
If **Enforce vNIC/vHBA/iSCSI Name** is selected and t
If it is not selected, the vNICs/vHBAs are selected if t … is used.

Add SAN Boot

vHBA : Fabric-B

Type : ○ Primary ● Secondary ○ Any

23. Click OK to add the SAN boot initiator.

24. From the vHBA drop-down list, select Add SAN Boot Target.

25. Leave `0` as the value for Boot Target LUN.

26. Enter the WWPN for `Controller1 (CL 3B)` recorded in Table 12 .

27. Select Primary for the SAN boot target type.

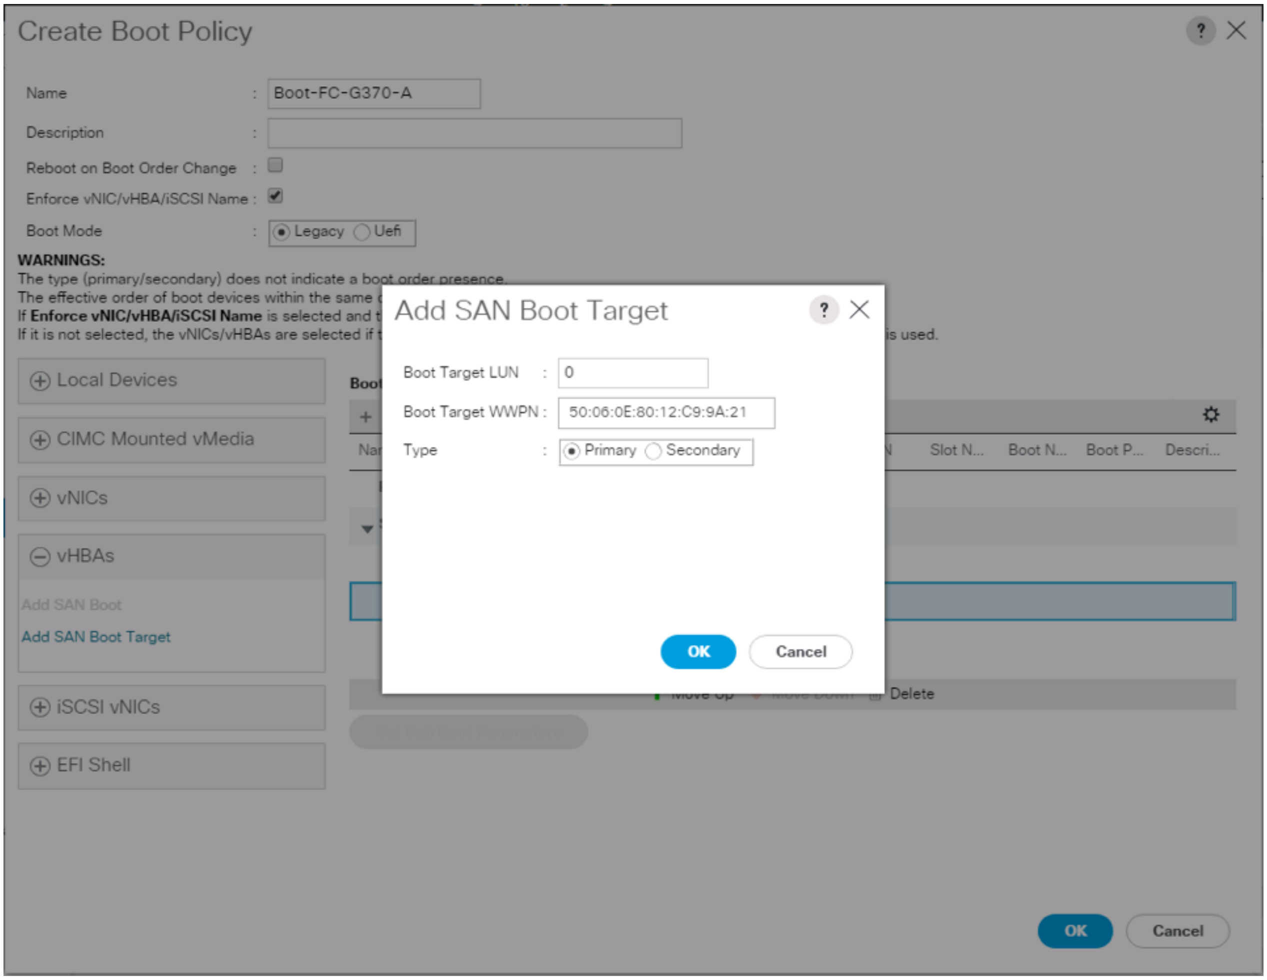

Create Boot Policy

Name : Boot-FC-G370-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 c
If **Enforce vNIC/vHBA/iSCSI Name** is selected and t
If it is not selected, the vNICs/vHBAs are selected if t ... is used.

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 50:06:0E:80:12:C9:9A:21

Type : ● Primary ○ Secondary

28. Click OK to add the SAN boot target.

29. From the vHBA drop-down list, select Add SAN Boot Target.

30. Enter `0` as the value for Boot Target LUN.

31. Enter the WWPN for `Controller2 (CL 4A)` recorded in Table 12 .

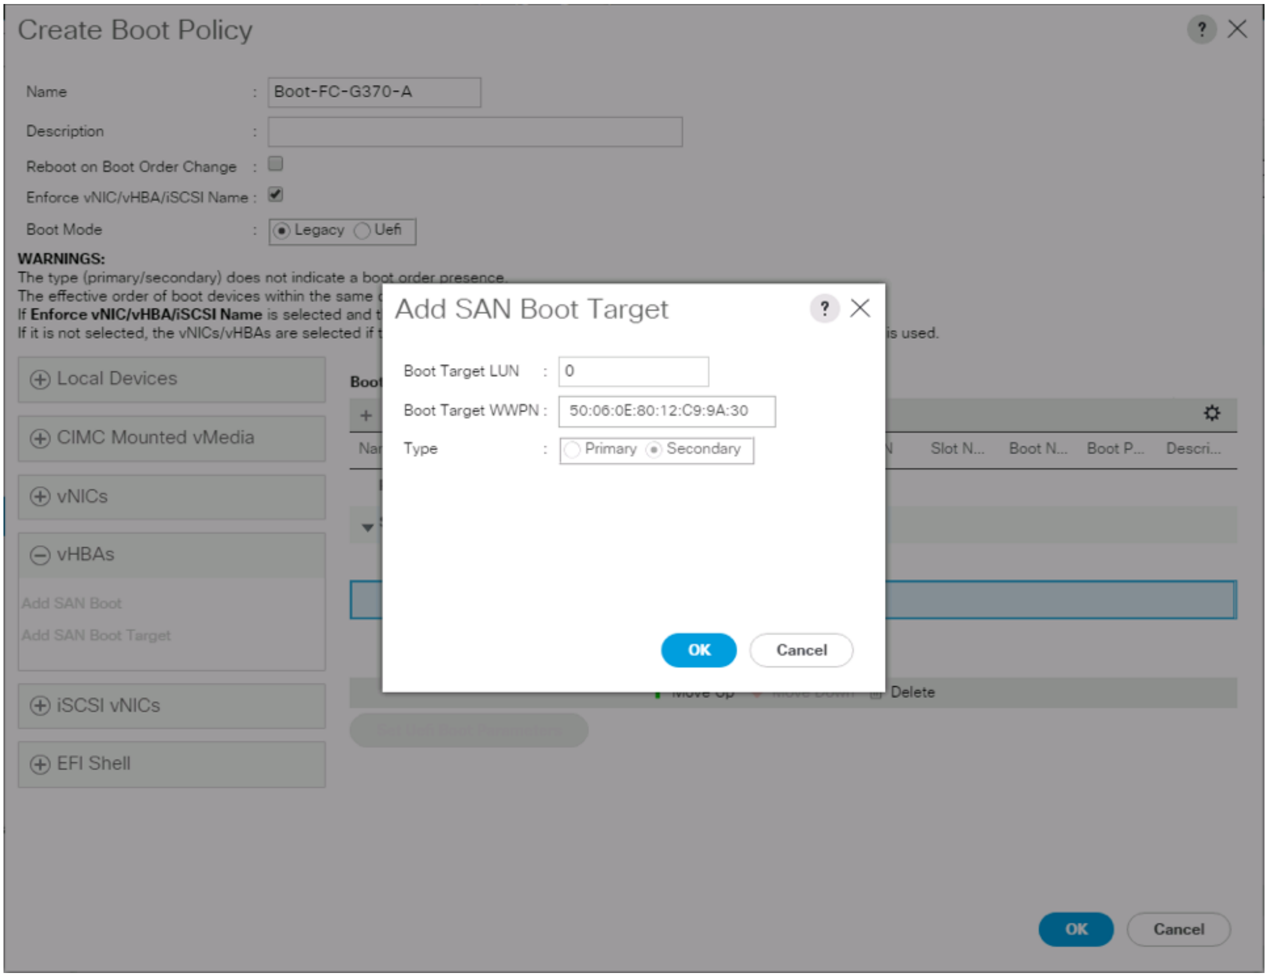

Create Boot Policy

Name : Boot-FC-G370-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
If **Enforce vNIC/vHBA/iSCSI Name** is selected and t
If it is not selected, the vNICs/vHBAs are selected if t … is used.

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 50:06:0E:80:12:C9:9A:30

Type : ○ Primary ● Secondary

32. Click OK to add the SAN boot target.

33. Expand CIMC Mounted vMedia and select Add CIMC Mounted CD/DVD.

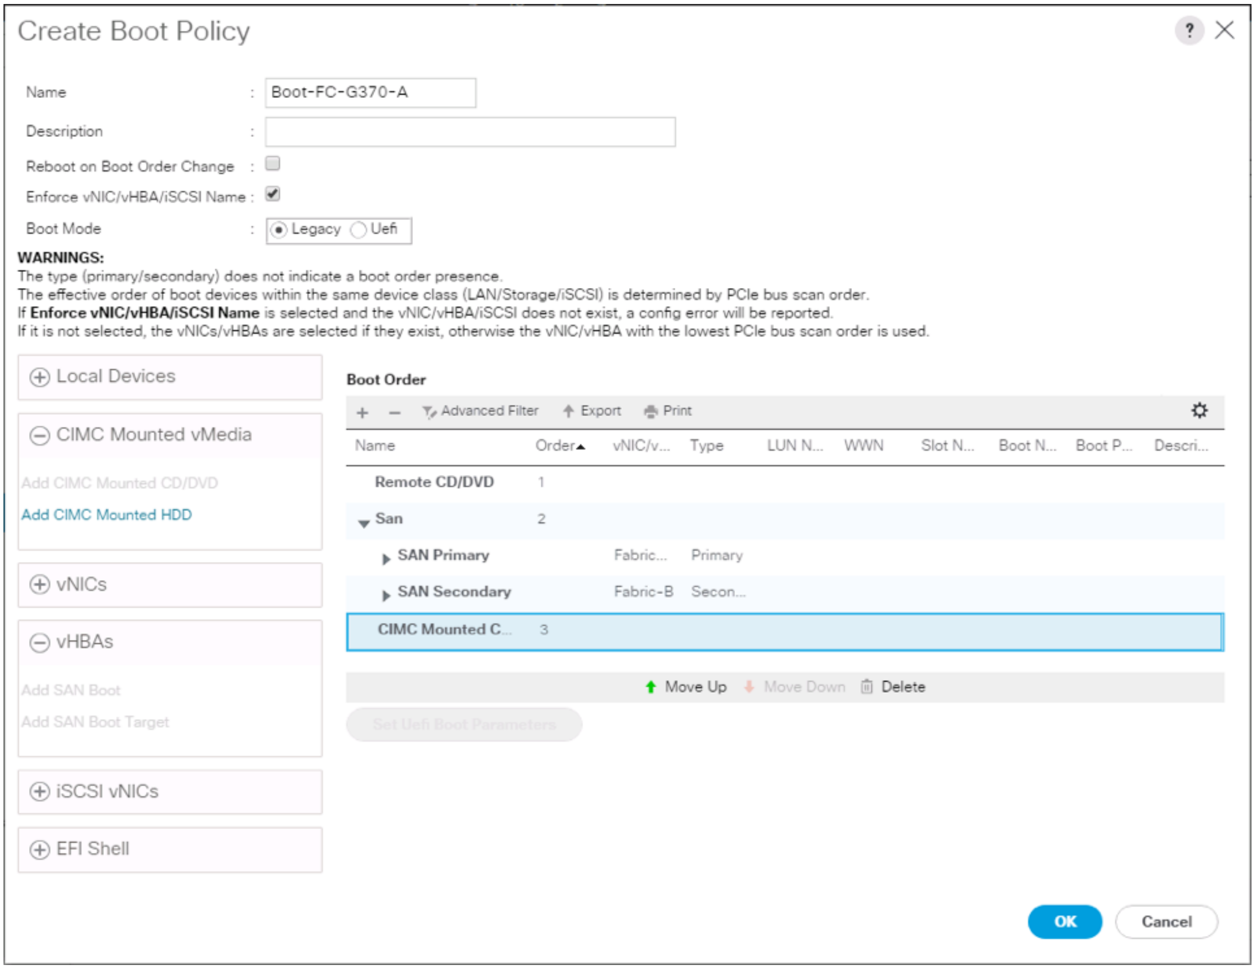

Create Boot Policy

Name : Boot-FC-G370-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	Order ▲	vNIC/v...	Type	LUN N...	WWN	Slot N...	Boot N...	Boot P...	Descri...
Remote CD/DVD	1								
▾ San	2								
▸ SAN Primary		Fabric...	Primary						
▸ SAN Secondary		Fabric-B	Secon...						
CIMC Mounted C...	3								

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

In this design, the boot targets specified aligns with the number of ports used from the VSP. If more VSP target ports are used to support VMFS bandwidth above what is available from the configured boot targets, a Storage Connection Policy will need to be created and used within a vHBA Initiator Group as specified in this Configuration Example: https://www.cisco.com/c/en/us/support/docs/servers-unified-computing/ucs-infrastructure-ucs-manager-software/116082-config-ucs-das-00.html

Create Service Profile Template

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

To create the service profile template, follow these steps:

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

2. Select Service Profile Templates > root.

3. Right-click root.

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

5. Enter `VSI-FC-G370-A` as the name of the service profile template. This service profile template is configured to boot from VSP G370 controller 1 on fabric A.

6. Select the “Updating Template” option.

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

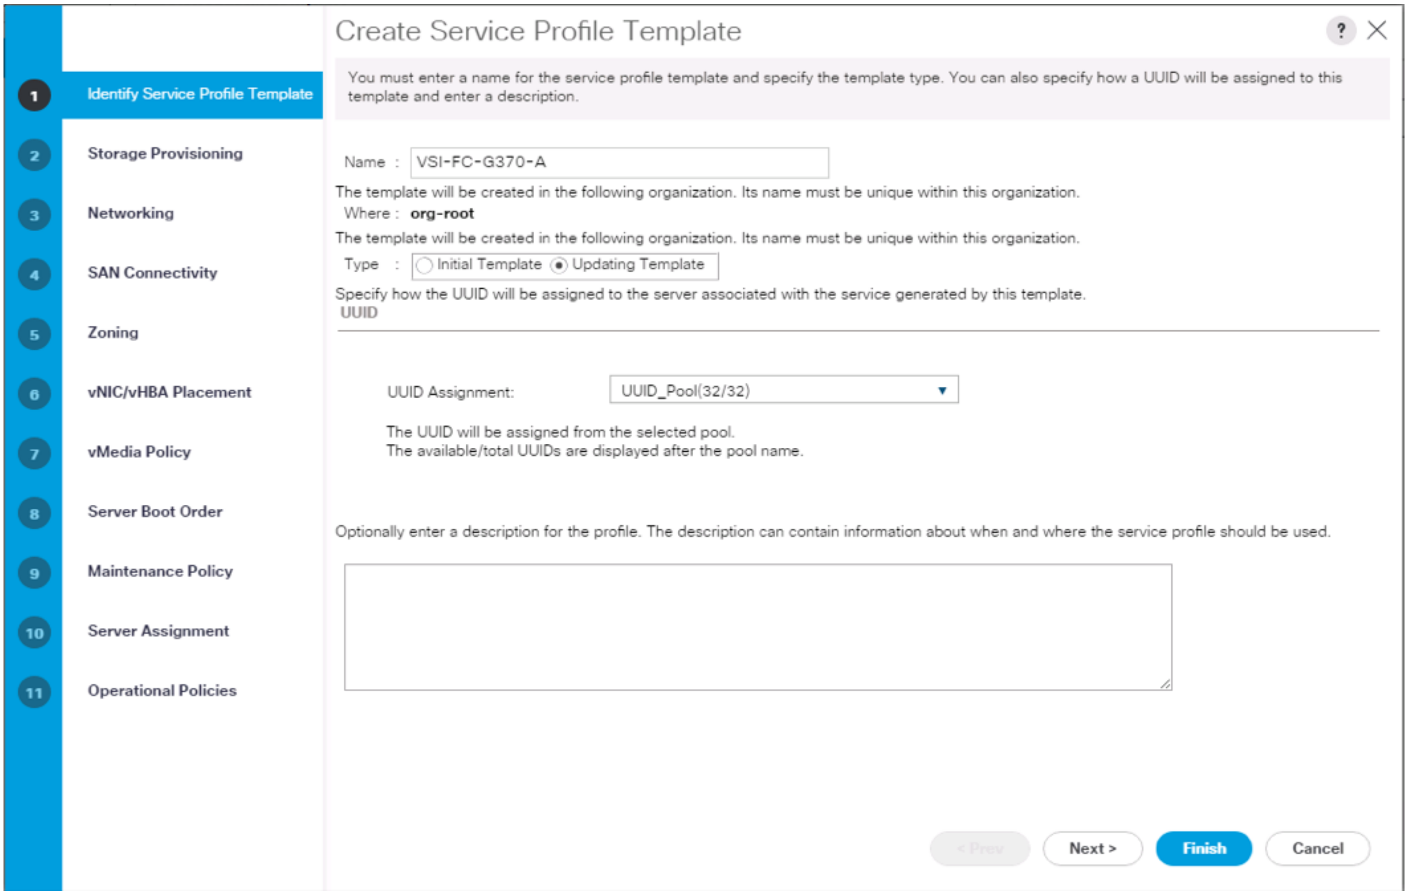

Create Service Profile Template

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 : VSI-FC-G370-A

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

Where : **org-root**

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(32/32)

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 the Local Disk Configuration Policy tab and select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.

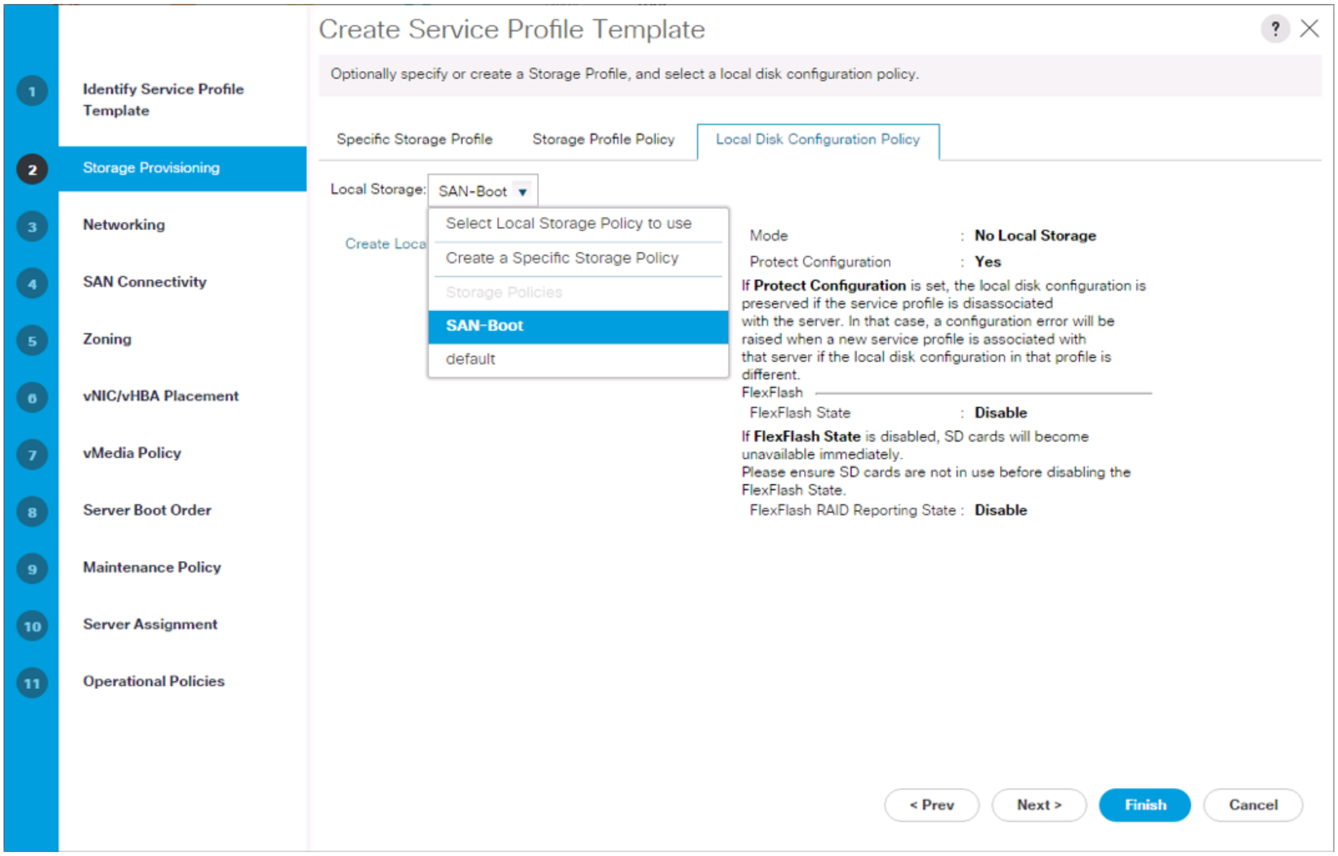

Create Service Profile Template

Optionally specify or create a Storage Profile, and select a local disk configuration policy.

Local Storage: SAN-Boot

Mode : No Local Storage

Protect Configuration : Yes

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State : Disable

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

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 FC-LAN-Policy from the LAN Connectivity Policy drop-down list.

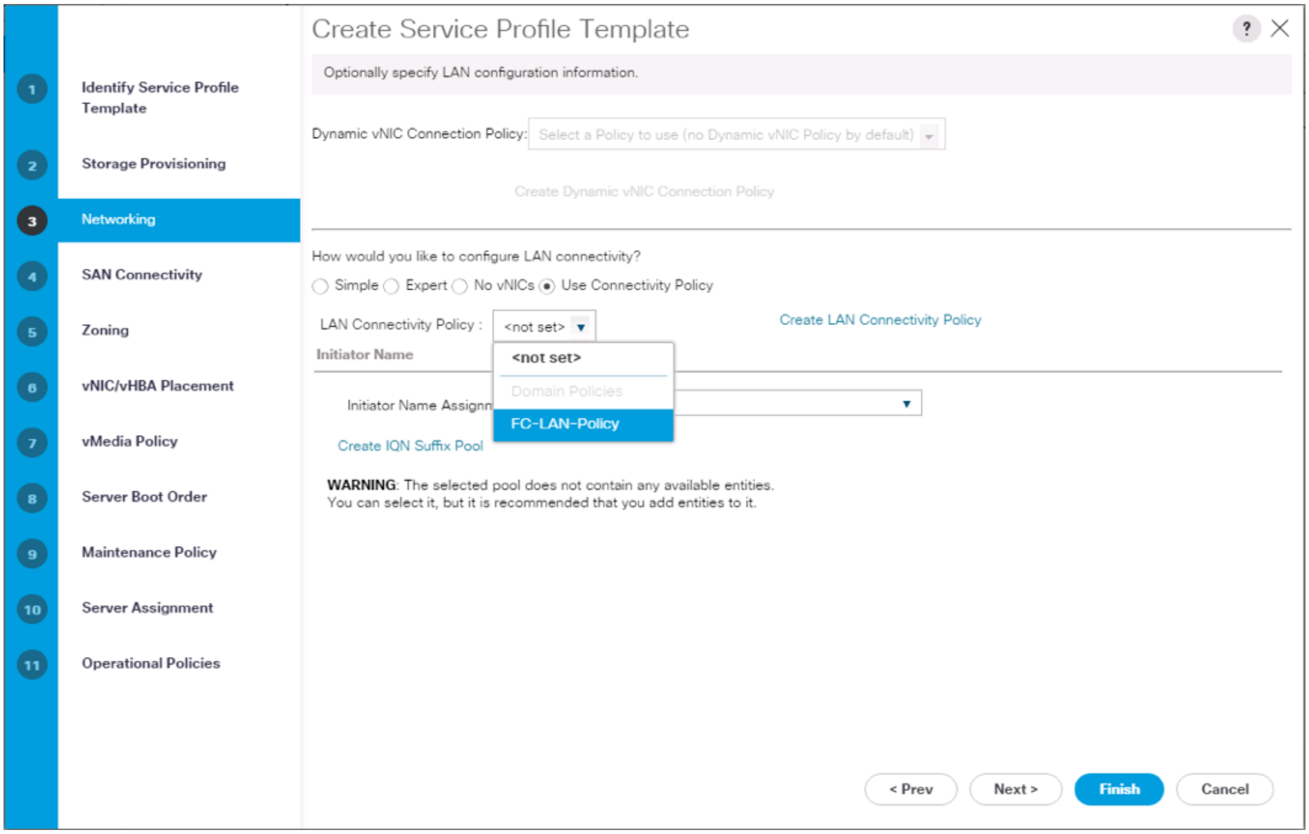

Create Service Profile Template

Optionally specify LAN configuration information.

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

How would you like to configure LAN connectivity?

○ Simple ○ Expert ○ No vNICs ● Use Connectivity Policy

LAN Connectivity Policy : <not set>

Initiator Name

Initiator Name Assignn

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

4. Click Next.

Configure Storage Options

To configure the storage options, follow these steps:

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

2. Pick the Infra-SAN-Policy option from the SAN Connectivity Policy drop-down list.

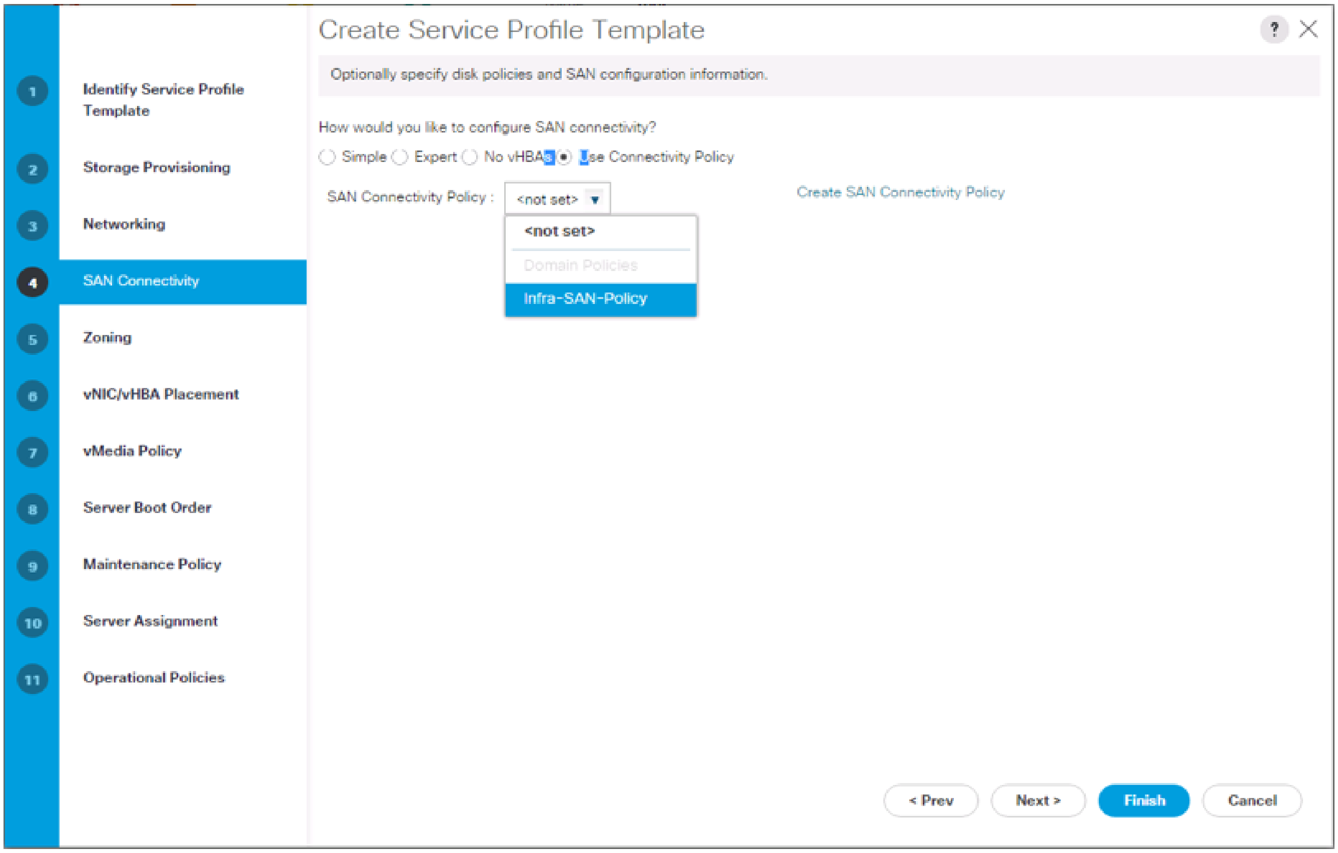

Create Service Profile Template

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 : <not set>

3. Click Next.

Configure Zoning Options

1. Leave Zoning configuration unspecified unless VSP ports exceed the number of boot targets created and a Storage Connection Policy has been created, click Next.

Configure vNIC/HBA Placement

1. In the “Select Placement” list, leave the placement policy as “Let System Perform Placement.”

2. Click Next.

Configure vMedia Policy

1. Do not select a vMedia Policy.

2. Click Next.

Configure Server Boot Order

1. Select `Boot-FC-G370-A` for Boot Policy.

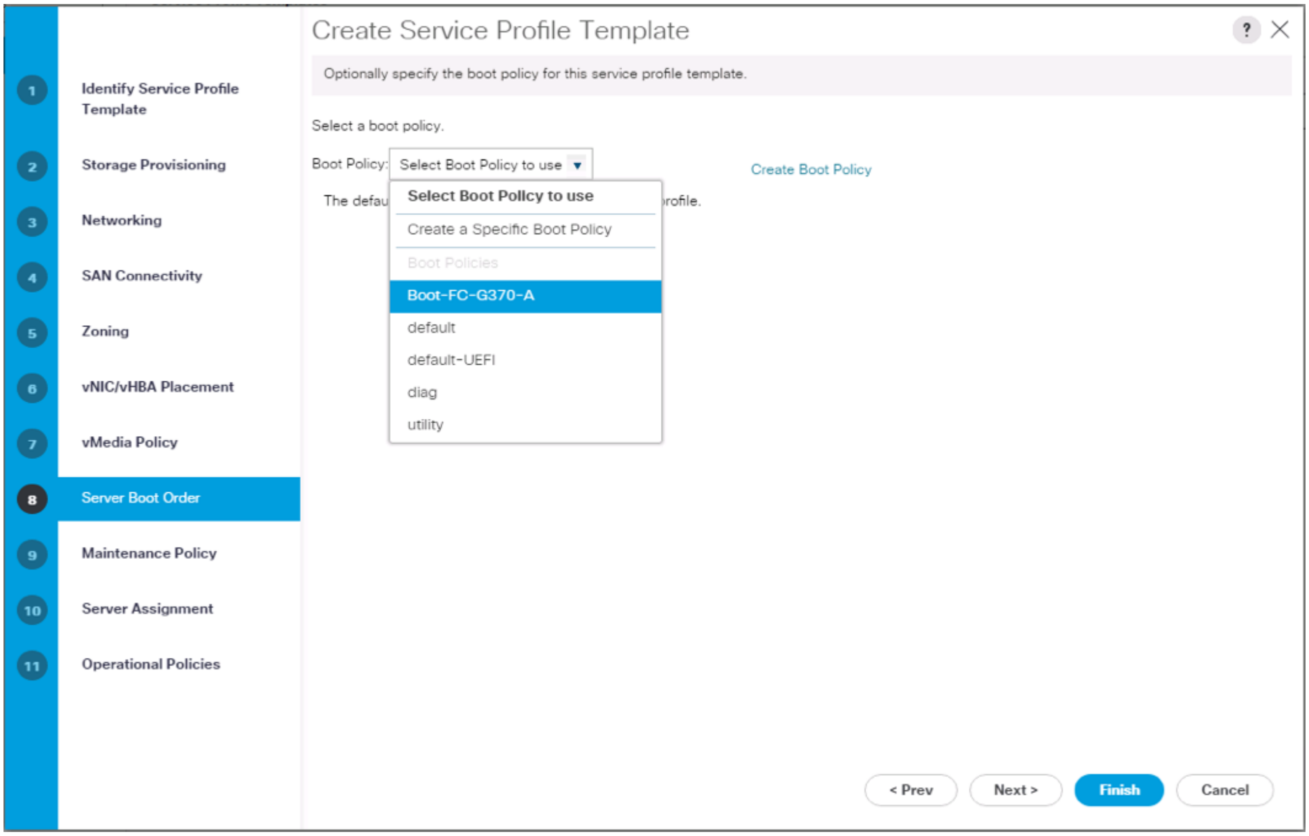

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Select Boot Policy to use

The defau... profile.

2. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

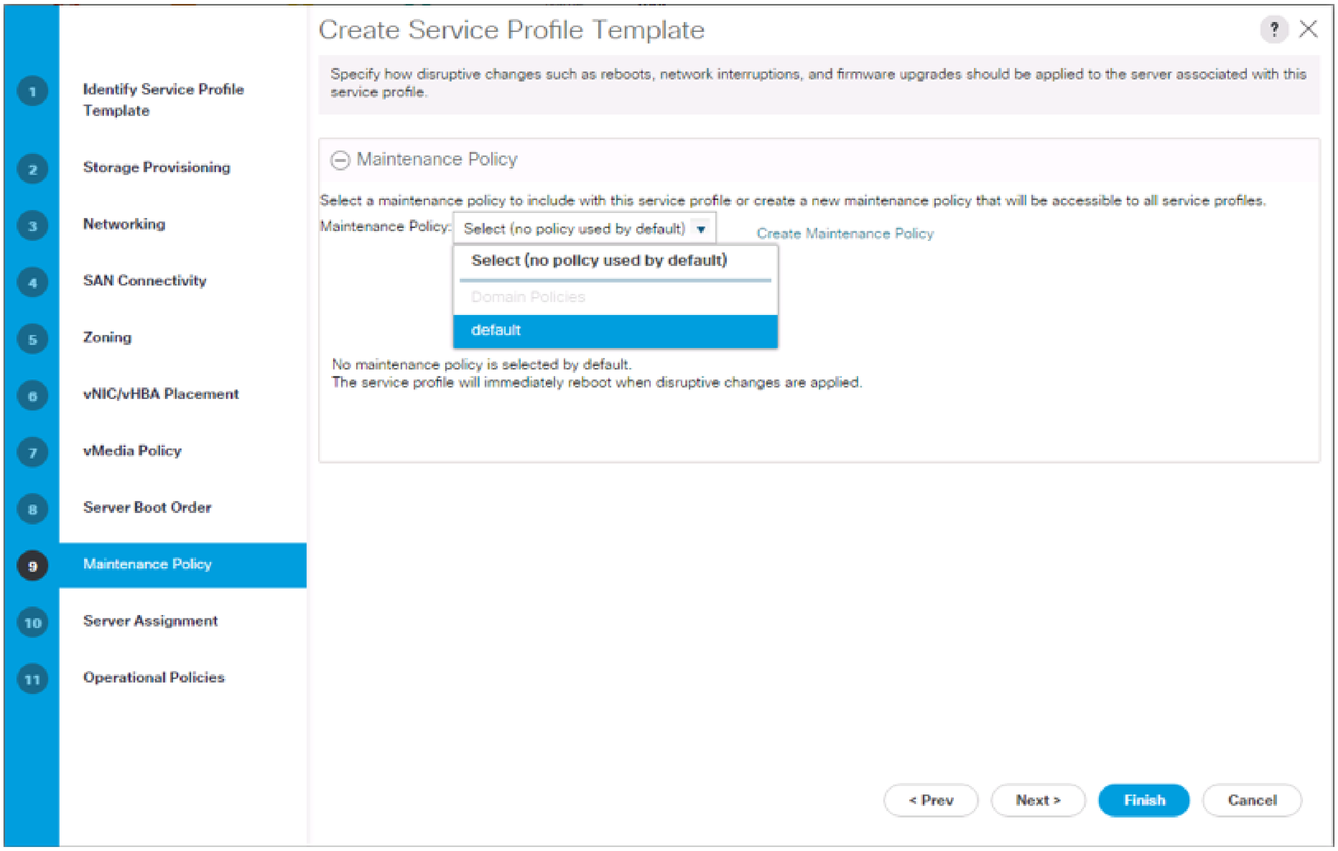

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

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: Select (no policy used by default)

No maintenance policy is selected by default.
The service profile will immediately reboot when disruptive changes are applied.

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, select `Infra_Pool`.
2. Optional: Select a Server Pool Qualification policy.
3. Select Up as the power state to be applied when the profile is associated with the server.
4. Optional: Select “UCS-B200M5” for the Server Pool Qualification.

Firmware Management at the bottom of the page can be left alone as it will use default from the Host Firmware list.

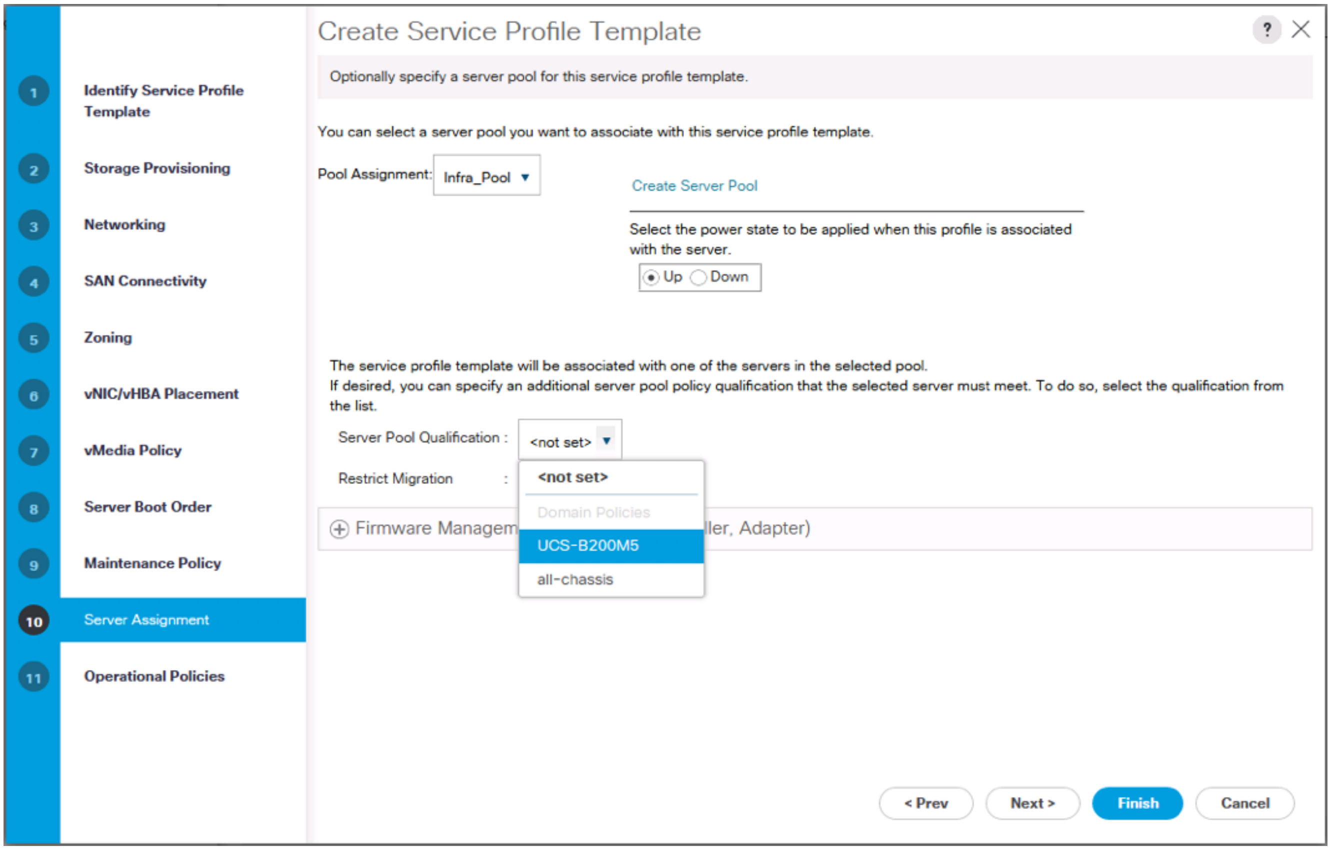

Create Service Profile Template

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 :

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.

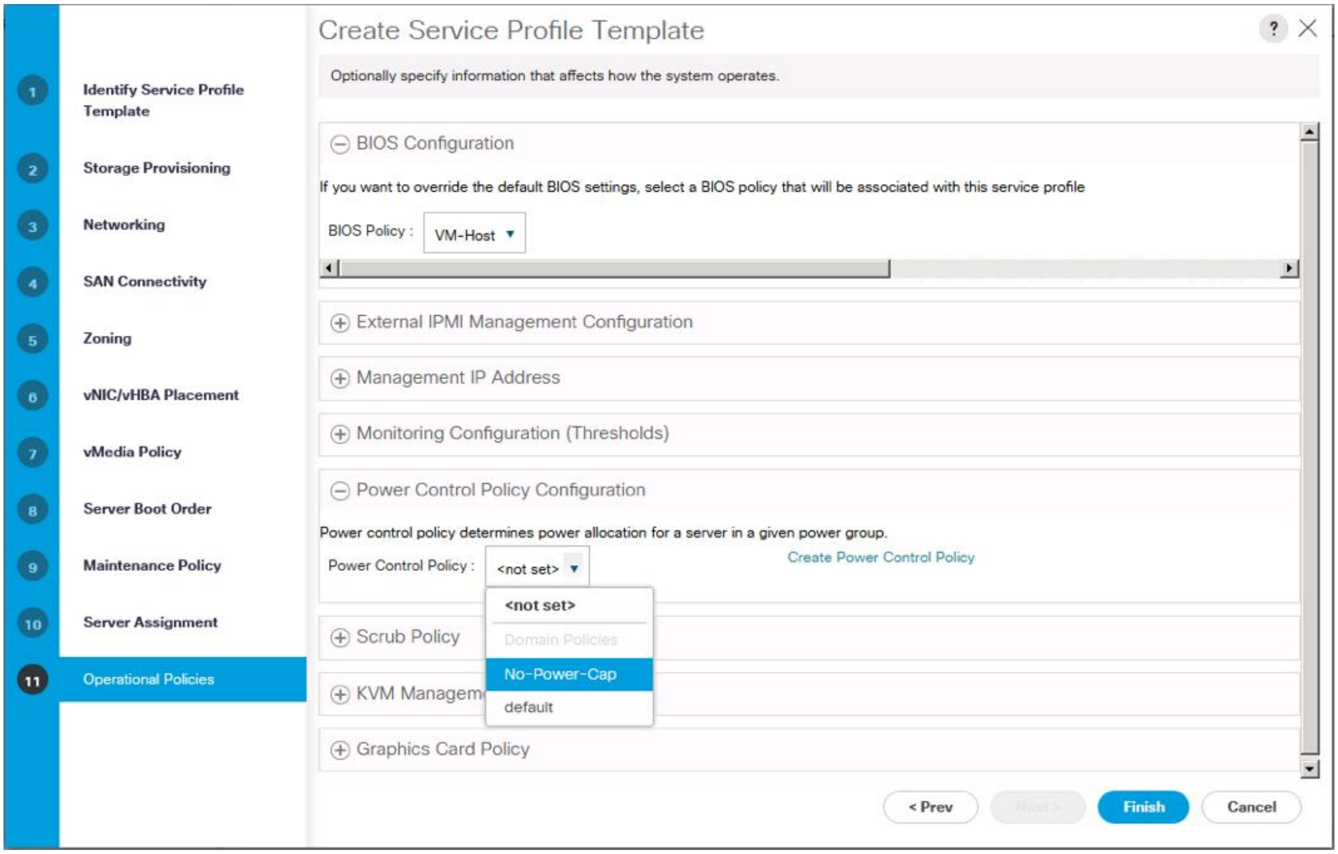

Create Service Profile Template

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 : <not set>

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create vMedia Service Profile Template

If the optional vMedia Policy is being used, a clone of the service profile template created above will be made to reference this vMedia Policy in these steps. The clone of the service profile template will have the vMedia Policy configured for it, and service profiles created from it, will be unbound and re-associated to the original service profile template after ESXi installation.

To create a clone of the VSI-FC-G370-A service profile template, and associate the vMedia Policy to it, follow these steps:

1. Connect to Cisco UCS Manager, click Servers.

2. Select Service Profile Templates > root > Service Template VSI-FC-G370-A.

3. Right-click Service Template VM-Host-FC-A and select Create a Clone.

4. Name the clone VSI-FC-G370-A-vM and click OK.

5. Select Service Template VSI-FC-G370-A-vM.

6. Select the vMedia Policy tab.

7. Under Actions, select Modify vMedia Policy.

8. From the drop-down list, select the ESXi-6.7U1-HTTP vMedia Policy.

9. Click OK then click OK again to complete modifying the Service Profile Template.

Create Service Profiles

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

1. Connect to the UCS 6454 Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.

2. Select Service Profile Templates > root > Service Template VSI-FC-G370-A-vM.

3. Right-click VSI-FC-G370-A-vM and select Create Service Profiles from Template.

4. Enter VSI-G370-0 as the service profile prefix.

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

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

7. Click OK to create the service profiles.

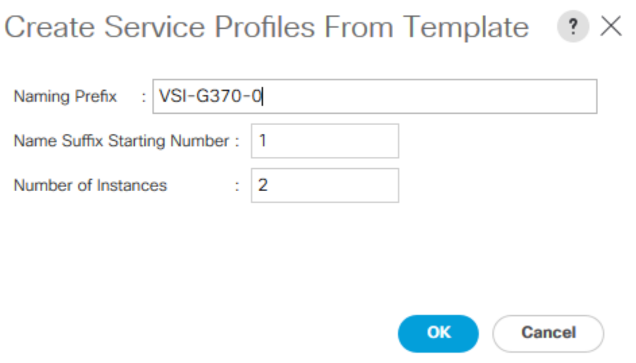

Create Service Profiles From Template

Naming Prefix : VSI-G370-0

Name Suffix Starting Number : 1

Number of Instances : 2

8. Click OK in the confirmation message to provision two Service Profiles.

When VMware ESXi 6.7 U1 has been installed on the hosts, the host Service Profiles can be unbound from the VM-Host-FC-A-vM and rebound to the VM-Host-FC-A Service Profile Template to remove the vMedia mapping from the host, to prevent issues at boot time if the HTTP source for the ESXi ISO is somehow not available.

Configuring Host Connectivity and Presentation of Storage on Hitachi Virtual Storage Platform

Configuration steps in this section assume that parity groups and LDEVs have been configured on the Hitachi VSP as part of the solution build/configuration by a partner or Hitachi professional services. If parity groups have not been configured on the Hitachi VSP, please reference the Hitachi Storage Virtualization Operating System documentation for creating parity groups before continuing with this section.

Ensure that you have planned which parity groups and LDEVs to use for specific storage requirements. Your configuration may vary based on the types of drives ordered with your VSP and with its configured parity groups.

Create a Hitachi Dynamic Provisioning Pool for UCS Server Boot LDEVs

To begin the provisioning process to create the Boot LDEVs that will be used as boot LUNs, follow these steps:

1. Log into Hitachi Storage Navigator.

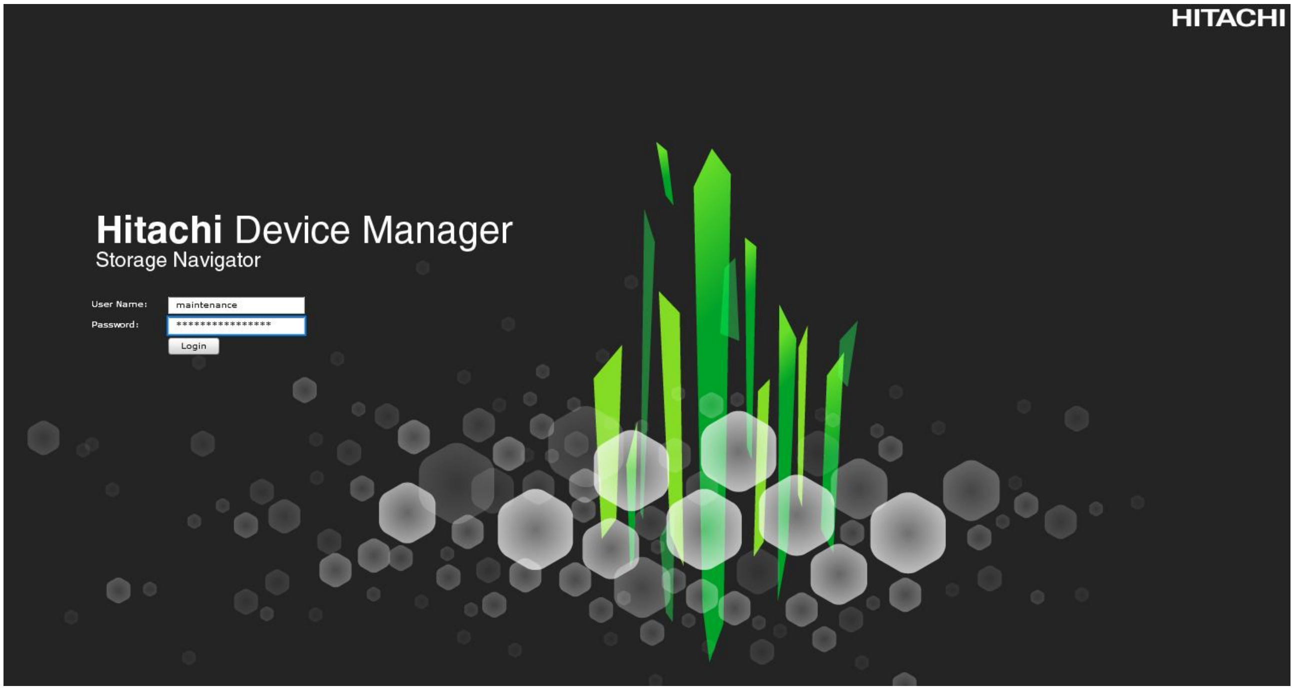

User Name: maintenance

Password: ****************

2. From the left Explorer pane select the **Storage Systems** tab.

3. Expand the storage system being configured. Highlight the **Pools** element in the navigation tree and click **Create Pools** to instantiate the Create Pools dialog box.

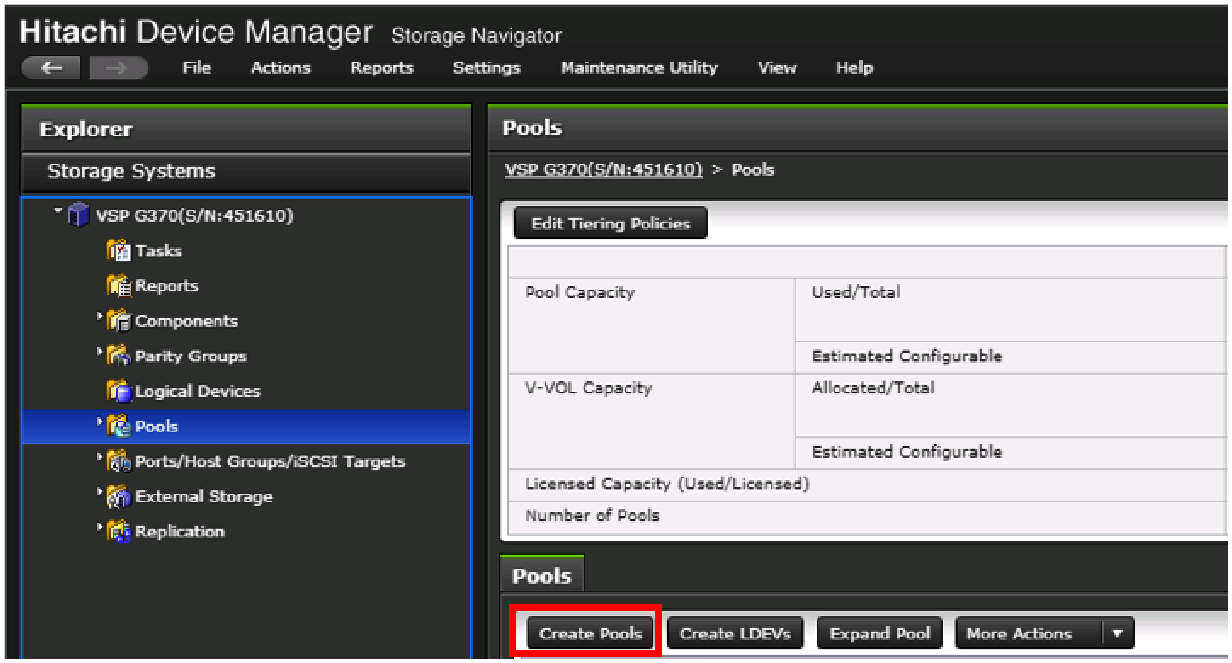

4. Configure the following items in the left pane of the Create Pools dialog box:

    a. Pool Type: Dynamic Provisioning

    b. System Type: Open [Only an option when configuring the G1500]

    c. Multi-Tier Pool: Disable

    d. Data Direct Mapping: Disable

    e. Pool Volume Selection: Manual

5. Select the **Drive Type/RPM** and **RAID Level** desired for the UCS server boot LDEV backing pool using the drop-down lists and click **Select Pool VOLs** to instantiate the Select Pool VOLs dialog box.

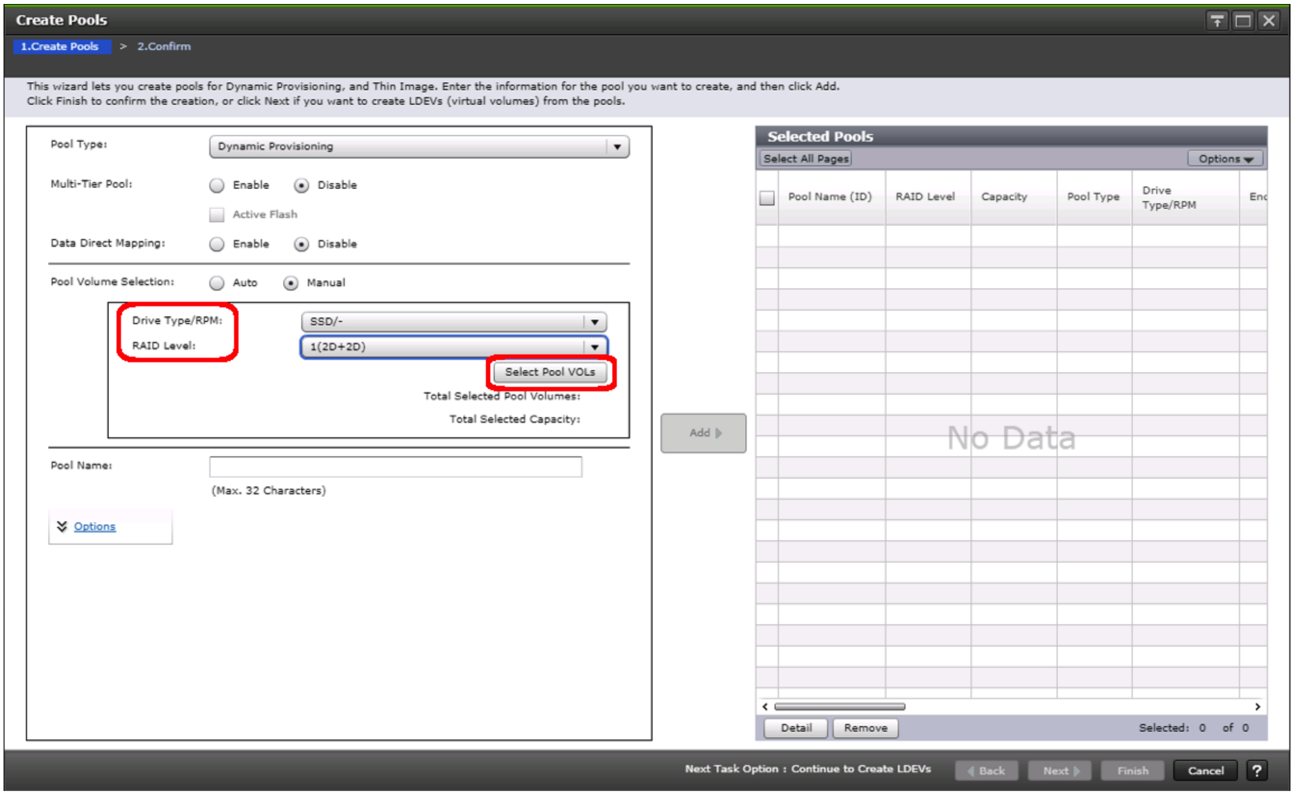

Create Pools

1.Create Pools > 2.Confirm

This wizard lets you create pools for Dynamic Provisioning, and Thin Image. Enter the information for the pool you want to create, and then click Add.
Click Finish to confirm the creation, or click Next if you want to create LDEVs (virtual volumes) from the pools.

Pool Type: Dynamic Provisioning

Multi-Tier Pool: ○ Enable ● Disable

□ Active Flash

Data Direct Mapping: ○ Enable ● Disable

Pool Volume Selection: ○ Auto ● Manual

Drive Type/RPM: SSD/-

RAID Level: 1(2D+2D)

Total Selected Pool Volumes:

Total Selected Capacity:

Pool Name:

(Max. 32 Characters)

Selected Pools

□	Pool Name (ID)	RAID Level	Capacity	Pool Type	Drive Type/RPM	Enc

No Data

Selected: 0 of 0

Next Task Option : Continue to Create LDEVs

6. Within the left pane of the Select Pool VOLs dialog box, select the checkbox next to the LDEVs to be used for the UCS server boot LDEV dynamic provisioning pool.

7. Click **Add** to move the selected LDEV to the right pane of the dialog, then click **OK** to return to the Create Pools dialog box.

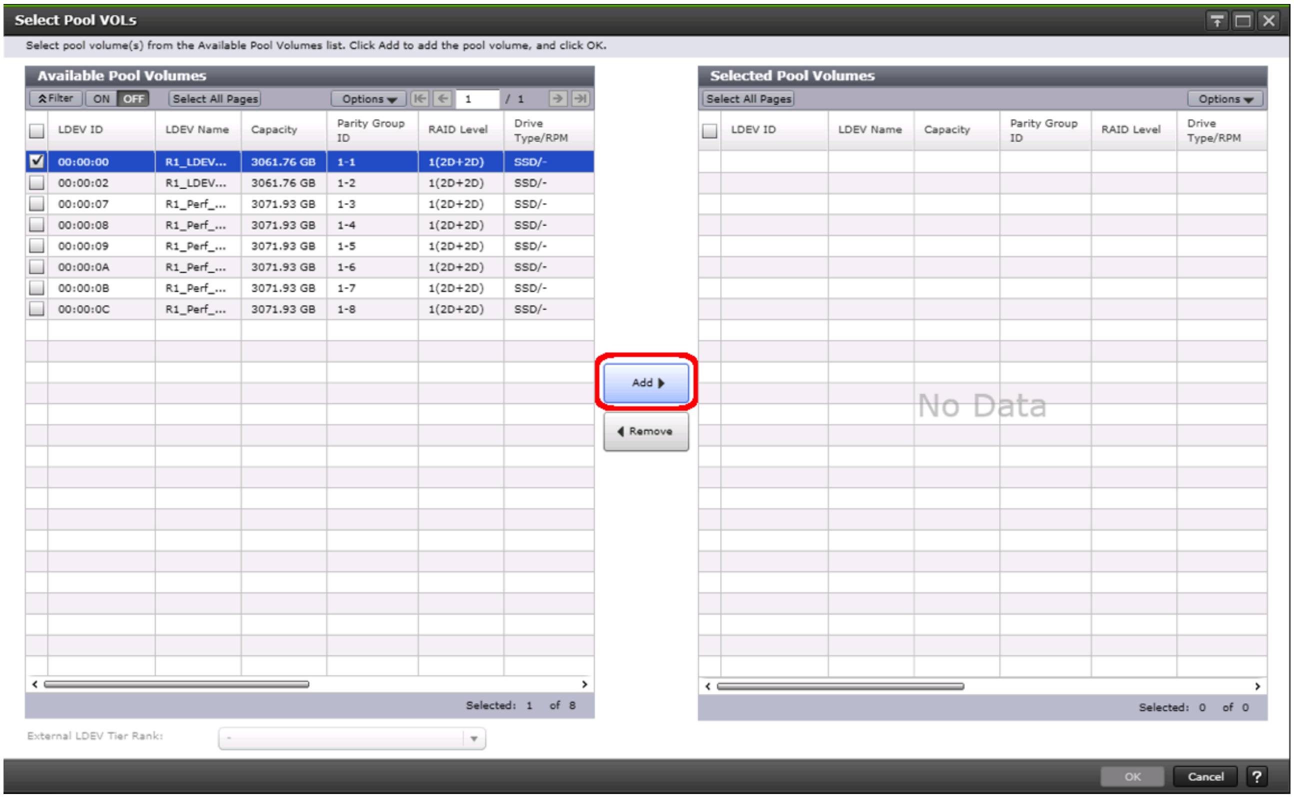

Select Pool VOLs

Select pool volume(s) from the Available Pool Volumes list. Click Add to add the pool volume, and click OK.

Available Pool Volumes

	LDEV ID	LDEV Name	Capacity	Parity Group ID	RAID Level	Drive Type/RPM
☑	00:00:00	R1_LDEV...	3061.76 GB	1-1	1(2D+2D)	SSD/-
□	00:00:02	R1_LDEV...	3061.76 GB	1-2	1(2D+2D)	SSD/-
□	00:00:07	R1_Perf_...	3071.93 GB	1-3	1(2D+2D)	SSD/-
□	00:00:08	R1_Perf_...	3071.93 GB	1-4	1(2D+2D)	SSD/-
□	00:00:09	R1_Perf_...	3071.93 GB	1-5	1(2D+2D)	SSD/-
□	00:00:0A	R1_Perf_...	3071.93 GB	1-6	1(2D+2D)	SSD/-
□	00:00:0B	R1_Perf_...	3071.93 GB	1-7	1(2D+2D)	SSD/-
□	00:00:0C	R1_Perf_...	3071.93 GB	1-8	1(2D+2D)	SSD/-

Selected: 1 of 8

Selected Pool Volumes

	LDEV ID	LDEV Name	Capacity	Parity Group ID	RAID Level	Drive Type/RPM

No Data

Selected: 0 of 0

External LDEV Tier Rank: -

8. You should now see values for **Total Selected Pool Volumes** and **Total Selected Capacity** shown under the **Select Pool VOLs** button. Give the dynamic provisioning pool a descriptive **Pool Name**, then click **Add** to add the pool to be created to the **Selected Pools** pane in the dialog.

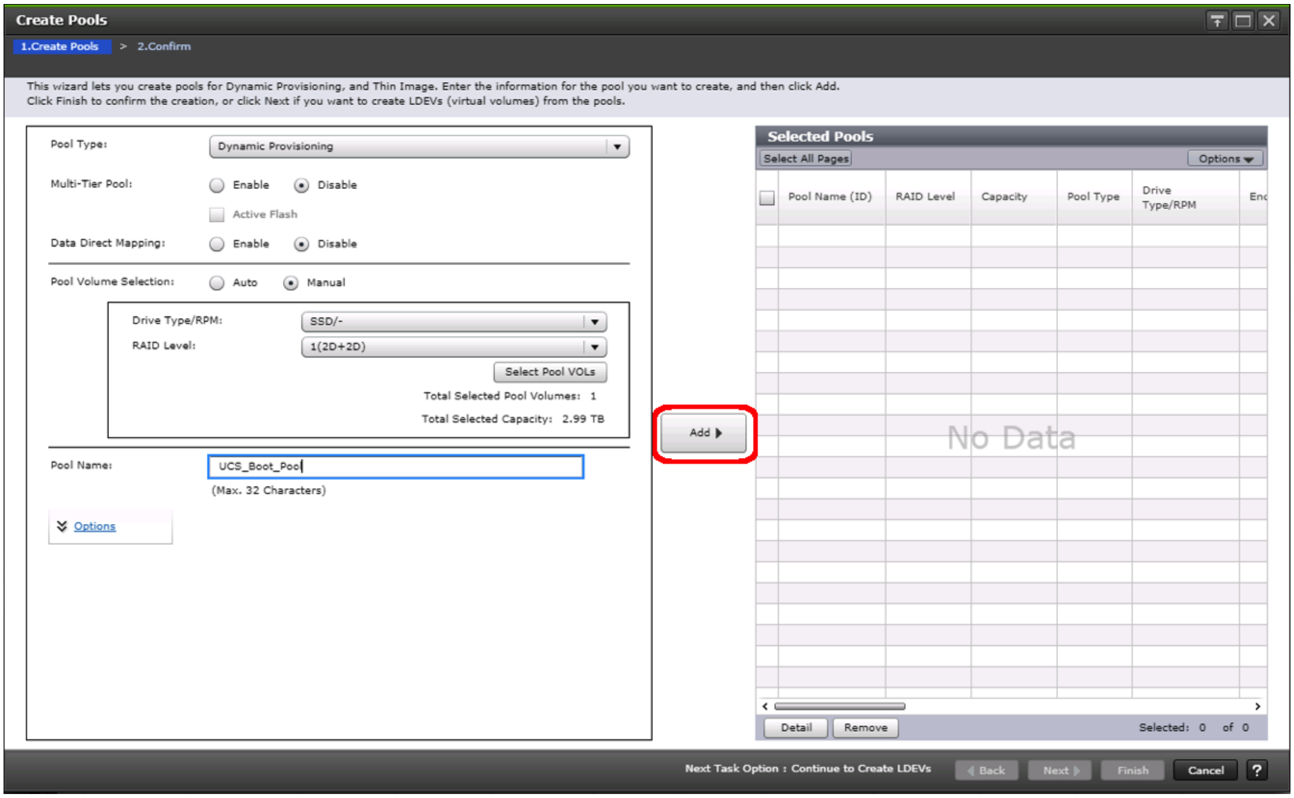

Create Pools

1.Create Pools > 2.Confirm

This wizard lets you create pools for Dynamic Provisioning, and Thin Image. Enter the information for the pool you want to create, and then click Add. Click Finish to confirm the creation, or click Next if you want to create LDEVs (virtual volumes) from the pools.

Pool Type: Dynamic Provisioning

Multi-Tier Pool: ○ Enable ● Disable

□ Active Flash

Data Direct Mapping: ○ Enable ● Disable

Pool Volume Selection: ○ Auto ● Manual

Drive Type/RPM: SSD/-

RAID Level: 1(2D+2D)

Total Selected Pool Volumes: 1

Total Selected Capacity: 2.99 TB

Pool Name: UCS_Boot_Pool

(Max. 32 Characters)

Selected Pools

	Pool Name (ID)	RAID Level	Capacity	Pool Type	Drive Type/RPM	Enc

No Data

Selected: 0 of 0

Next Task Option : Continue to Create LDEVs

9. Click Finish.

10. Review the configuration for the pool to be created in the Create Pools confirmation dialog box and ensure the **Go to tasks window for status** checkbox is checked, then click **Apply**.

11. The tasks status window will appear, wait for the task status to show complete before moving onto the next step.

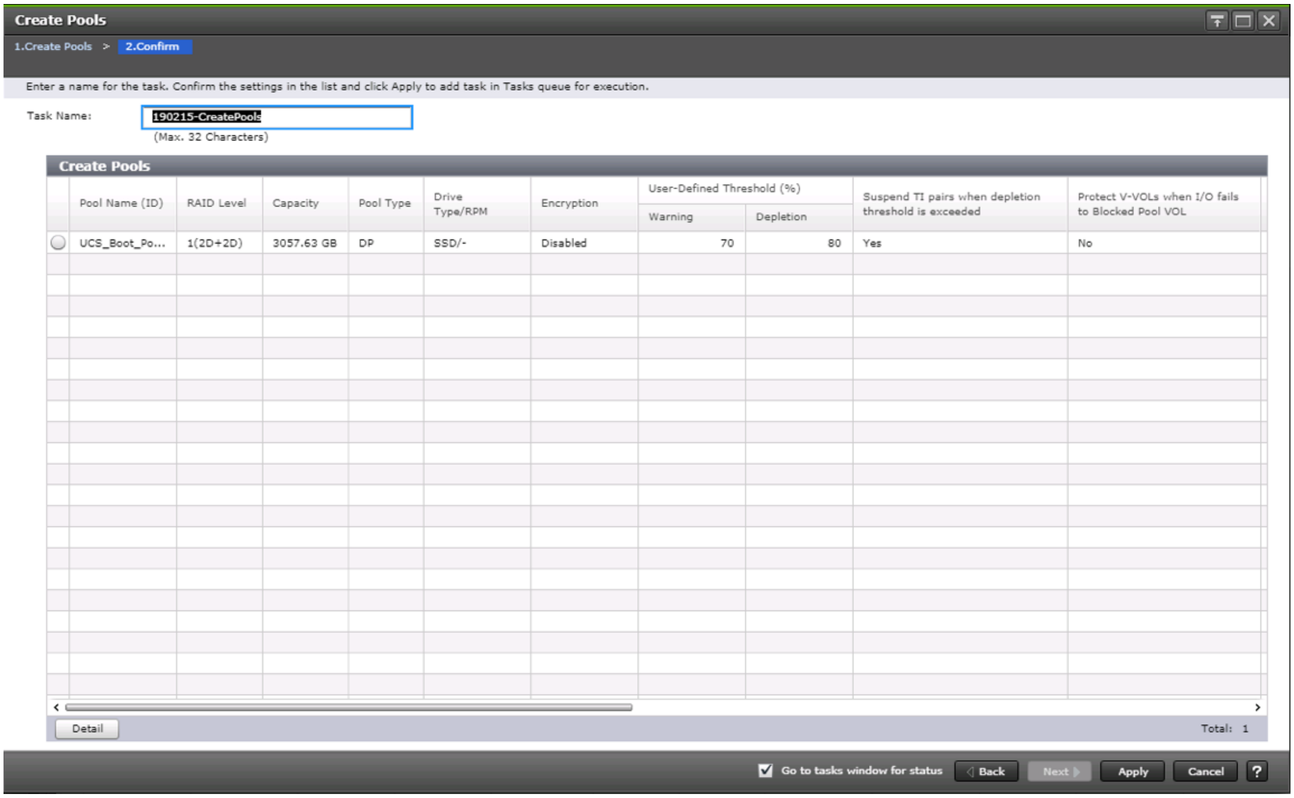

Create Pools

1.Create Pools > 2.Confirm

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: 190215-CreatePools
(Max. 32 Characters)

Create Pools

	Pool Name (ID)	RAID Level	Capacity	Pool Type	Drive Type/RPM	Encryption	User-Defined Threshold (%)		Suspend TI pairs when depletion threshold is exceeded	Protect V-VOLs when I/O fails to Blocked Pool VOL
							Warning	Depletion		
○	UCS_Boot_Po...	1(2D+2D)	3057.63 GB	DP	SSD/-	Disabled	70	80	Yes	No

Total: 1

☑ Go to tasks window for status

Create a Hitachi Dynamic Provisioning Pool for UCS Server VMFS Volume LDEVs

Follow the steps in section Create a Hitachi Dynamic Provisioning Pool for UCS Server Boot LDEVs to create the dynamic provisioning pool for the UCS Server VMFS volume LDEVs, selecting the **Drive Type/RPM**, **RAID Level**, and number of **Pool VOLs** desired for the pool backing the VMFS volumes in the solution.

Create Host Groups for Cisco UCS Server vHBAs on Each Fabric

An individual host group must be created on each physical fibre channel port on the VSP for each vHBA attached to its respective fabric. The number of host groups created will depend on the number of paths per LDEV. Ensure you have documented the specific ports on each fabric being used on the VSP, their WWNs, and each vHBA WWPN before you proceed with this section and ensure that all initiators for the UCS Service Profiles you will be creating host groups for are showing as logged into the respective VSP fibre channel ports by following the steps below.

To create Host Groups for UCS server vHBAs on each fabric, follow these steps:

1. From the Explorer pane within Hitachi Storage Navigator, select the **Storage Systems** tab and expand the storage system being configured.

2. Highlight the **Ports/Host Groups/iSCSI Targets** element in the navigation tree and select the **Login WWNs/iSCSI Names** tab.

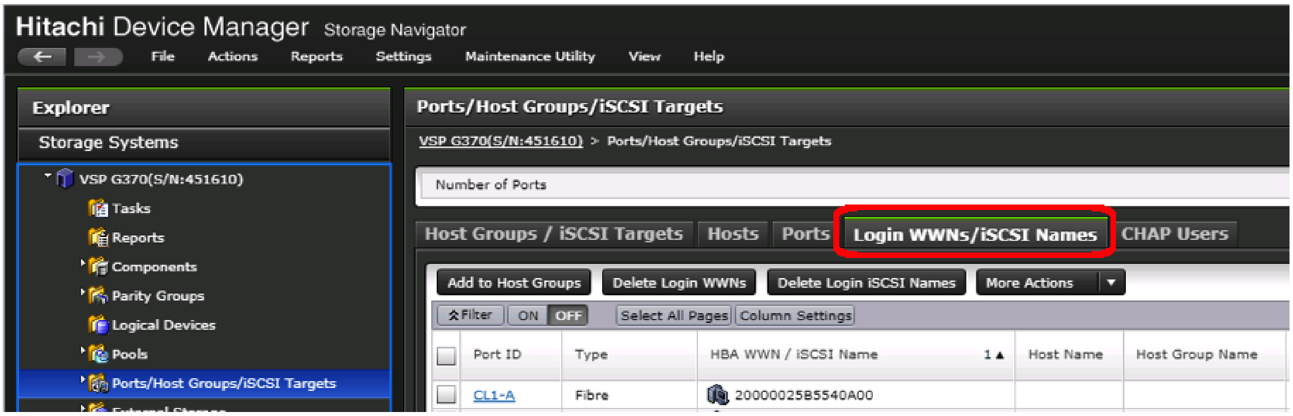

□	Port ID	Type	HBA WWN / iSCSI Name 1▲	Host Name	Host Group Name
□	CL1-A	Fibre	20000025B5540A00		

3. Review the list of WWNs and associated ports. You should be able to see each vHBA assigned to each fabric associated with each port on the VSP that it is zoned to.

4. Click the column names to sort the information to make this task easier or utilize the **Filter** feature to limit the number of records displayed. If any vHBA WWNs do not show in the list, go back and double check the zoning configuration on the FI.

5. With the **Ports/Host Groups/iSCSI Targets** element in the navigation tree still selected, click the **Host Groups/iSCSI Targets** tab.

6. Click **Create Host Groups** to instantiate the Create Host Groups dialog box.

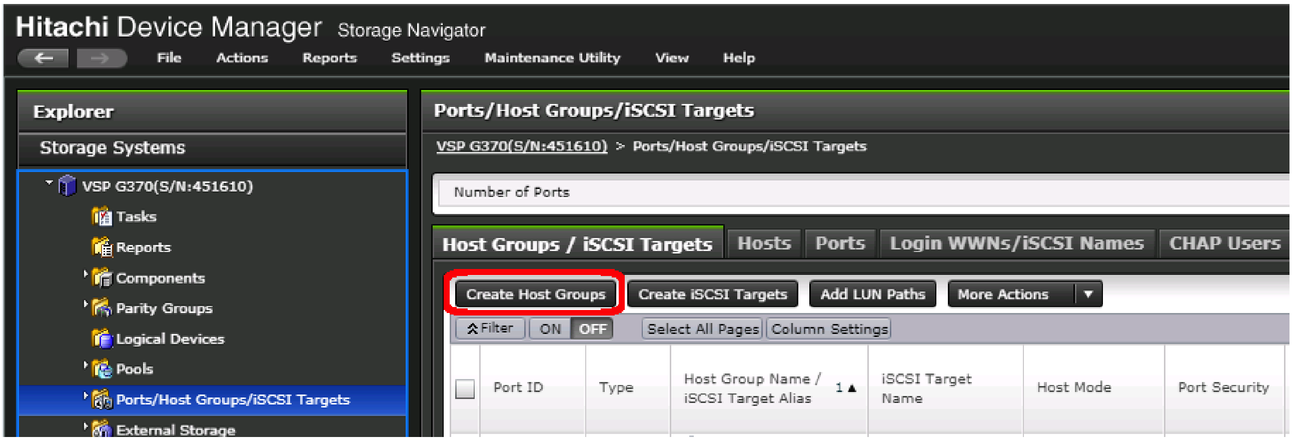

□	Port ID	Type	Host Group Name / iSCSI Target Alias 1▲	iSCSI Target Name	Host Mode	Port Security

7. Host groups will be created separately for fabric A and fabric B vHBAs. Start with the fabric A host group for an individual UCS Service Profile and modify the following within the Create Host Groups dialog box:

   a. **Host Group Name**: Provide a descriptive name for the host and ensure there is an identifier for the fabric you are configuring (i.e., VSI-G370-1_Fab_A)

   b. **Host Mode**: Select 21 [VMware Extension] from the drop-down list.

   c. **Host Mode Options**: For each of the following Host Mode Options, find the Mode Number in the pane, select the checkbox, and click the **Enable** button:

      i. 54 – (VAAI) Support Option for the EXTENDED COPY command
      ii. 63 – (VAAI) Support option for vStorage APIs based on T10 standards
      iii. 114 – The automatic asynchronous reclamation on ESXi6.5 or later

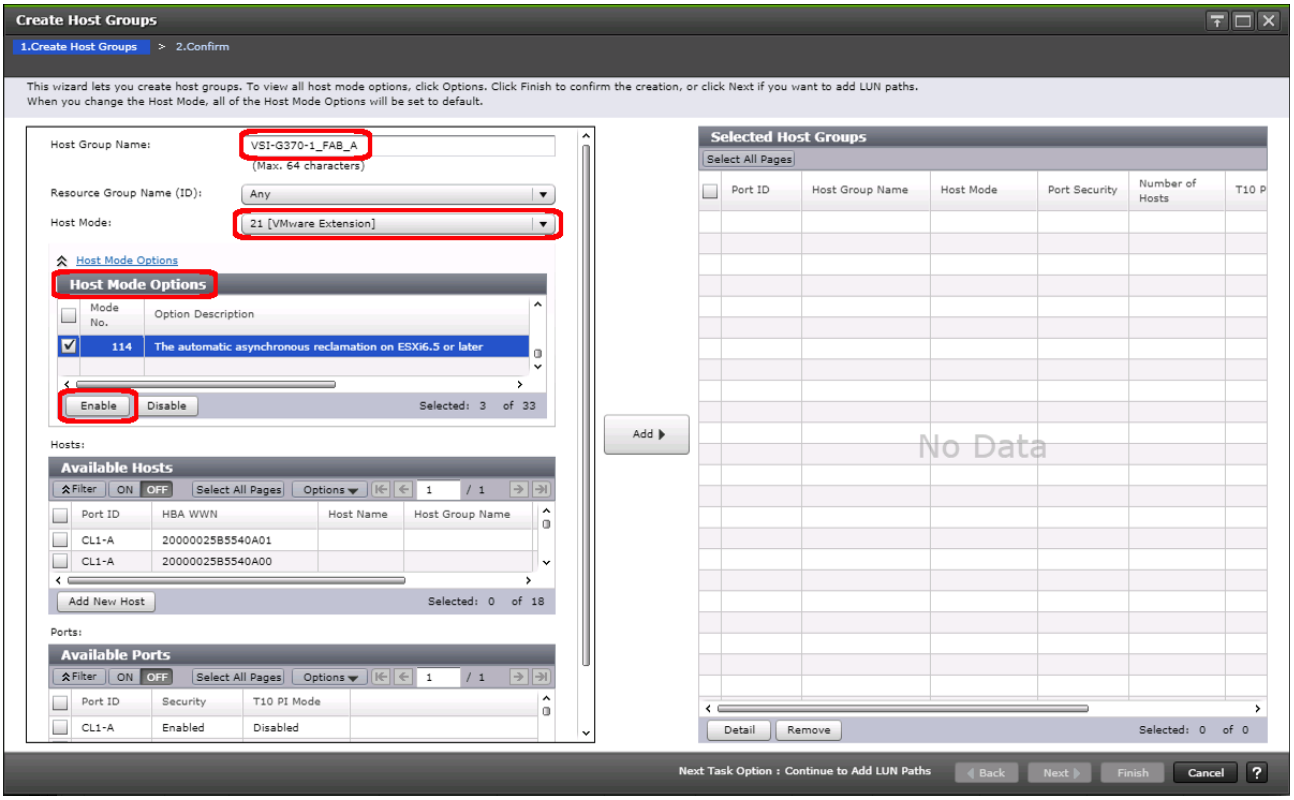

This wizard lets you create host groups. To view all host mode options, click Options. Click Finish to confirm the creation, or click Next if you want to add LUN paths.
When you change the Host Mode, all of the Host Mode Options will be set to default.

Host Group Name: VSI-G370-1_FAB_A (Max. 64 characters)

Resource Group Name (ID): Any

Host Mode: 21 [VMware Extension]

	Mode No.	Option Description
☑	114	The automatic asynchronous reclamation on ESXi6.5 or later

Selected: 3 of 33

	Port ID	HBA WWN	Host Name	Host Group Name
□	CL1-A	200000025B5540A01		
□	CL1-A	200000025B5540A00		

Selected: 0 of 18

	Port ID	Security	T10 PI Mode
□	CL1-A	Enabled	Disabled

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	T10 P

Selected: 0 of 0

Next Task Option : Continue to Add LUN Paths

8. Write down the WWN information from Table 11 and **Error! Reference source not found.**in the previous Create Device Aliases section:

Table 13 Fabric A Targets and Initiators

	Name	WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
Target	G370-CL1-A	50:06:0e:80:12:c9:9a:00	
Target	G370-CL2-B	50:06:0e:80:12:c9:9a:11	
Initiator	VSI-G370-01	20:00:00:25:B5:54:0A:00	
Initiator	VSI-G370-02	20:00:00:25:B5:54:0A:01	

Table 14 Fabric B Targets and Initiators

		WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
Target	G370-CL3-B	50:06:0e:80:12:c9:9a:21	
Target	G370-CL4-A	50:06:0e:80:12:c9:9a:30	
Initiator	VSI-G370-01	20:00:00:25:B5:54:0B:00	

		WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
Initiator	VSI-G370-02	20:00:00:25:B5:54:0B:01	

9. Scroll down in the Create Host Groups dialog.

10. Within the **Available Hosts** section, click Filter.

11. Create an Attribute/Value filter of:

    – HBA WWN

    – Using "contains" as a qualifier

    – Using the last four characters of the Fabric A initiator for the host

This will be without ":" characters from the above table, and assuming that the last four characters is sufficient to produce a unique matching value. If necessary, use a larger identifying character string.

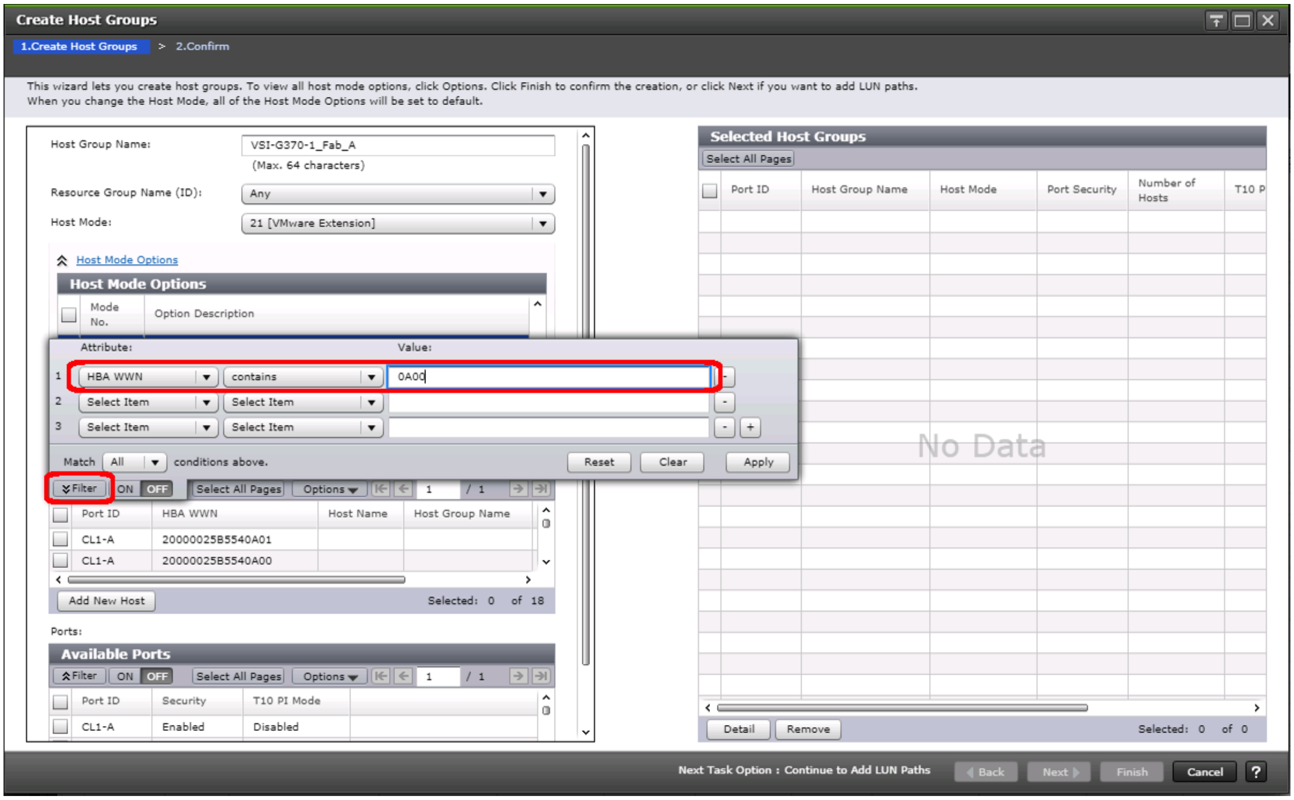

Create Host Groups

1.Create Host Groups > 2.Confirm

This wizard lets you create host groups. To view all host mode options, click Options. Click Finish to confirm the creation, or click Next if you want to add LUN paths.
When you change the Host Mode, all of the Host Mode Options will be set to default.

Host Group Name: VSI-G370-1_Fab_A (Max. 64 characters)

Resource Group Name (ID): Any

Host Mode: 21 [VMware Extension]

Host Mode Options

□	Mode No.	Option Description

	Attribute:		Value:
1	HBA WWN	contains	0A00
2	Select Item	Select Item	
3	Select Item	Select Item	

Match All conditions above.

□	Port ID	HBA WWN	Host Name	Host Group Name
□	CL1-A	20000025B5540A01		
□	CL1-A	20000025B5540A00		

Selected: 0 of 18

Ports:

Available Ports

□	Port ID	Security	T10 PI Mode
□	CL1-A	Enabled	Disabled

Selected Host Groups

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	T10 P

No Data

Selected: 0 of 0

Next Task Option : Continue to Add LUN Paths

12. Click **Apply**.

13. Click Filter again to hide the filter rules dialog box.

14. Select the checkbox for the first port shown in the filtered list within the **Available Hosts** section.

15. Within the **Available Ports** section, check the checkboxes for all ports zoned to the host within Fabric A only.

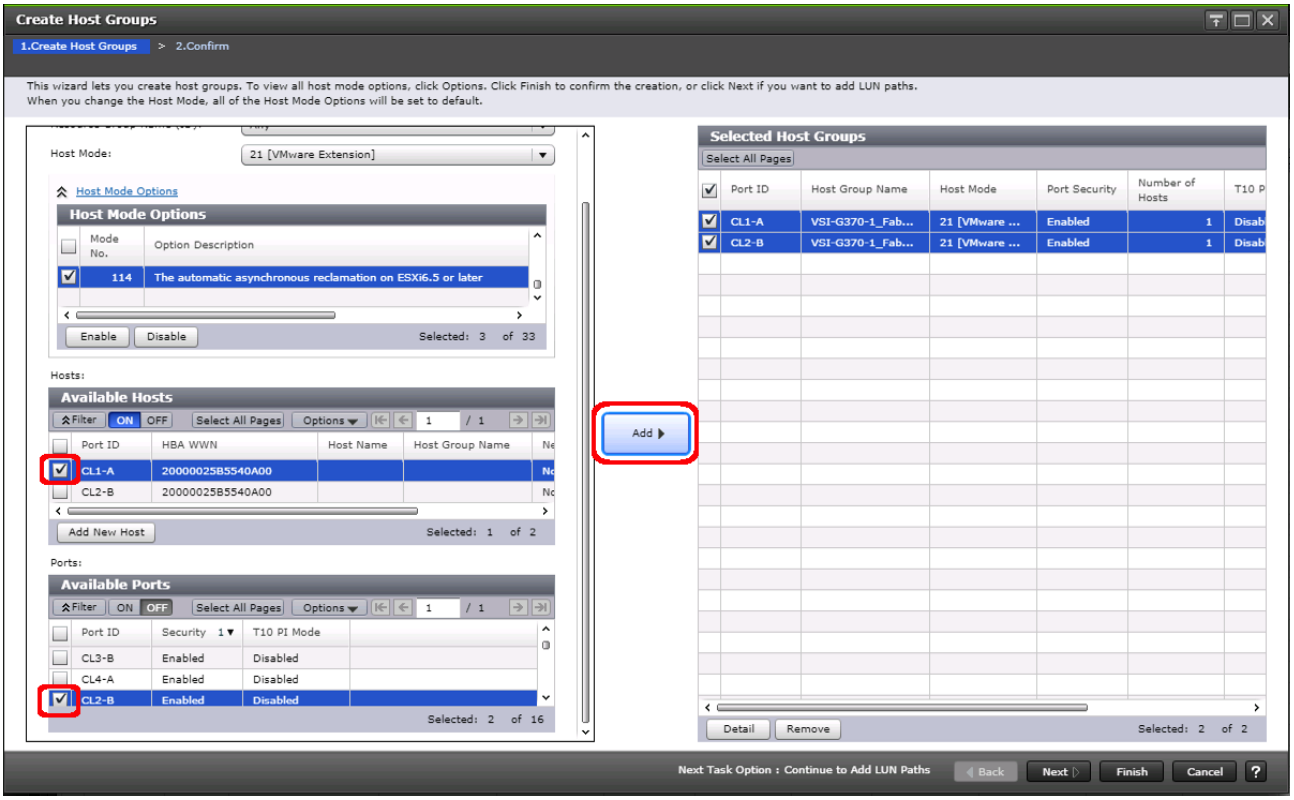

Create Host Groups

1.Create Host Groups > 2.Confirm

This wizard lets you create host groups. To view all host mode options, click Options. Click Finish to confirm the creation, or click Next if you want to add LUN paths.
When you change the Host Mode, all of the Host Mode Options will be set to default.

Host Mode: 21 [VMware Extension]

Host Mode Options

□	Mode No.	Option Description
☑	114	The automatic asynchronous reclamation on ESXi6.5 or later

Selected: 3 of 33

Hosts:

Available Hosts

□	Port ID	HBA WWN	Host Name	Host Group Name	Ne
☑	CL1-A	20000025B5540A00			Nc
□	CL2-B	20000025B5540A00			Nc

Selected: 1 of 2

Ports:

Available Ports

□	Port ID	Security 1▼	T10 PI Mode
□	CL3-B	Enabled	Disabled
□	CL4-A	Enabled	Disabled
☑	CL2-B	Enabled	Disabled

Selected: 2 of 16

Selected Host Groups

☑	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	T10 P
☑	CL1-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Disab
☑	CL2-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Disab

Selected: 2 of 2

Next Task Option : Continue to Add LUN Paths

In the picture above, the CL1-A port entry was also selected within the **Available Ports** section.

16. Click **Add**, then click **Finish**.

17. Review the host group configuration for the Fabric A host groups for the UCS Service Profile being configured.

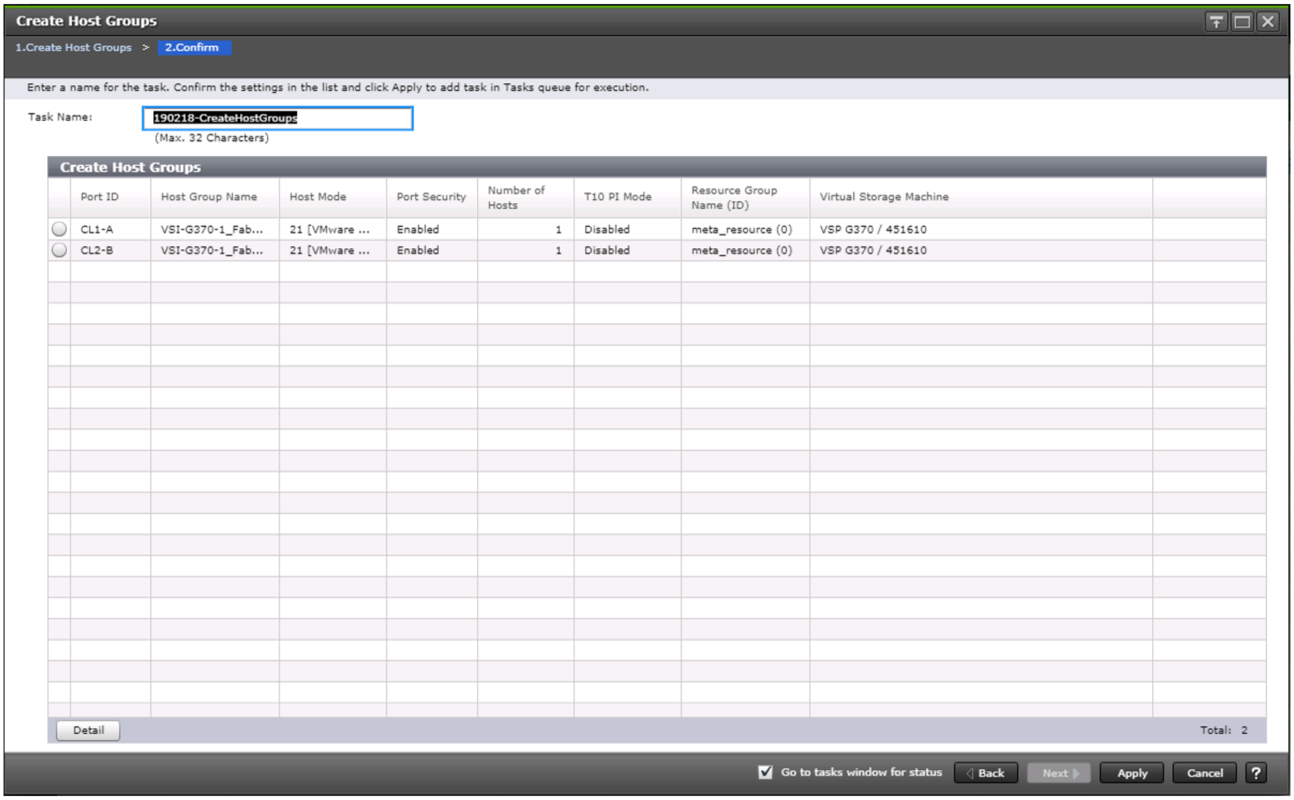

Create Host Groups

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: 190218-CreateHostGroups
(Max. 32 Characters)

Create Host Groups

	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	T10 PI Mode	Resource Group Name (ID)	Virtual Storage Machine
○	CL1-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Disabled	meta_resource (0)	VSP G370 / 451610
○	CL2-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Disabled	meta_resource (0)	VSP G370 / 451610

Total: 2

18. Click **Apply**.

19. Repeat steps 1-18 to create the host groups for all remaining initiator WWN from the Fabric A and Fabric B tables above, using a descriptive name for the host on Fabric A/B, the vHBA WWN on Fabric A/B for the UCS Service Profile, and the associated Fabric A/B ports on the Hitachi VSP.

Create Boot LDEVs for Each UCS Service Profile and Add LDEV Paths

Individual boot LDEVs must be created for each UCS Service Profile for the ESXi hypervisor to be installed onto. Prior to beginning these steps, ensure you have identified the fibre channel ports on the Hitachi VSP that will be used for presentation of the boot LDEVs to the UCS servers. Please note that a maximum of four paths can be used within the UCS Service Profile (two on each fabric) as boot targets.

To create boot LDEVs for each UCS service profile and add LDEV paths, follow these steps:

1. From the left Explorer pane within Hitachi Storage Navigator, select the **Storage Systems** tab and expand the storage system being configured.

2. Expand the **Pools** element in the navigation tree and highlight the UCS Boot pool previously created for use as the backing storage for the UCS boot LDEVs.

3. Select the **Virtual Volumes** tab and click **Create LDEVs** to instantiate the Create LDEVs dialog.

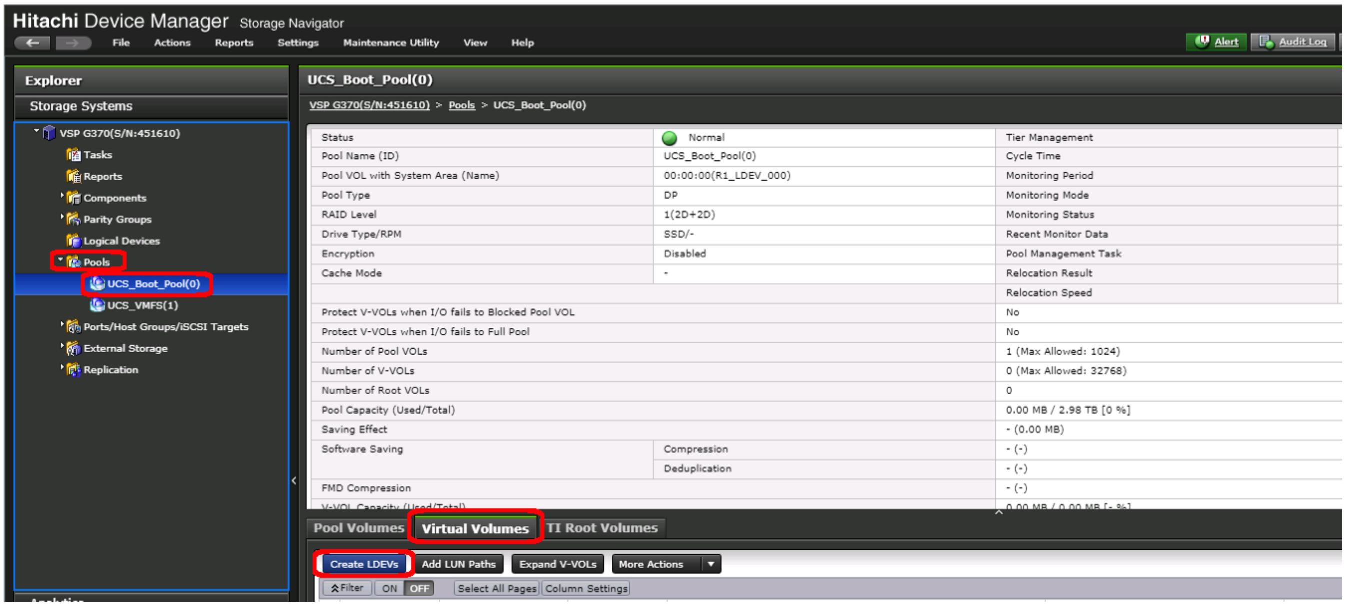

UCS_Boot_Pool(0)

VSP G370(S/N:451610) > Pools > UCS_Boot_Pool(0)

Status	Normal	Tier Management
Pool Name (ID)	UCS_Boot_Pool(0)	Cycle Time
Pool VOL with System Area (Name)	00:00:00(R1_LDEV_000)	Monitoring Period
Pool Type	DP	Monitoring Mode
RAID Level	1(2D+2D)	Monitoring Status
Drive Type/RPM	SSD/-	Recent Monitor Data
Encryption	Disabled	Pool Management Task
Cache Mode	-	Relocation Result
		Relocation Speed
Protect V-VOLs when I/O fails to Blocked Pool VOL		No
Protect V-VOLs when I/O fails to Full Pool		No
Number of Pool VOLs		1 (Max Allowed: 1024)
Number of V-VOLs		0 (Max Allowed: 32768)
Number of Root VOLs		0
Pool Capacity (Used/Total)		0.00 MB / 2.98 TB [0 %]
Saving Effect		- (0.00 MB)
Software Saving	Compression	- (-)
	Deduplication	- (-)
FMD Compression		- (-)
V-VOL Capacity (Used/Total)		0.00 MB / 0.00 MB [- %]

4. Modify the following within the Create LDEVs dialog:

   - **LDEV Capacity**: Enter the capacity desired for the UCS Service Profile boot LDEV. Note that ESXi requires a minimum of 5.2GB for a boot LDEV as documented by VMware.

   - Number of LDEVs: 1

   - **LDEV Name**: Provide a descriptive name and numeric identifier for the boot LDEV. For ease of identification, it is recommended that the server name or other identifier specific to the service profile being configured be entered in the **Prefix** field.

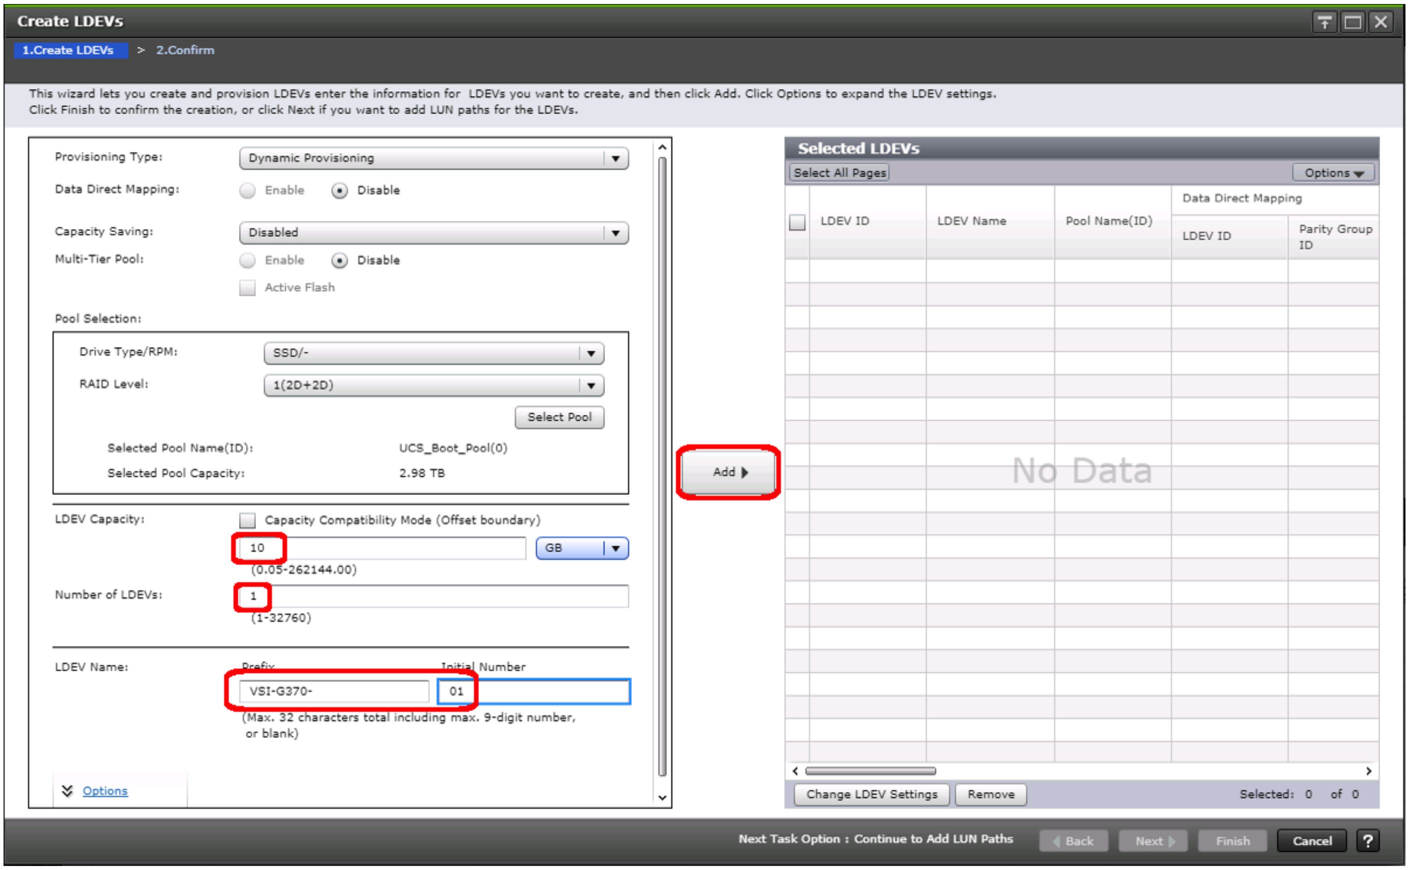

Create LDEVs

1.Create LDEVs > 2.Confirm

This wizard lets you create and provision LDEVs enter the information for LDEVs you want to create, and then click Add. Click Options to expand the LDEV settings.
Click Finish to confirm the creation, or click Next if you want to add LUN paths for the LDEVs.

Provisioning Type: Dynamic Provisioning

Data Direct Mapping: ○ Enable ● Disable

Capacity Saving: Disabled

Multi-Tier Pool: ○ Enable ● Disable

□ Active Flash

Pool Selection:

Drive Type/RPM: SSD/-

RAID Level: 1(2D+2D)

Selected Pool Name(ID): UCS_Boot_Pool(0)

Selected Pool Capacity: 2.98 TB

LDEV Capacity: □ Capacity Compatibility Mode (Offset boundary)

10 GB
(0.05-262144.00)

Number of LDEVs: 1
(1-32760)

LDEV Name: Prefix VSI-G370- Initial Number 01
(Max. 32 characters total including max. 9-digit number, or blank)

□	LDEV ID	LDEV Name	Pool Name(ID)	Data Direct Mapping	
				LDEV ID	Parity Group ID

No Data

Selected: 0 of 0

Next Task Option : Continue to Add LUN Paths

5. Click Add and verify that the boot LDEV is listed in the right-hand Selected LDEVs pane, then click Next.

The **Select LDEVs** screen shows the selected LDEVs to which the paths will be added.

6. Ensure the newly created boot LDEV is the only LDEV in the **Selected LDEVs** pane, then click **Next**.

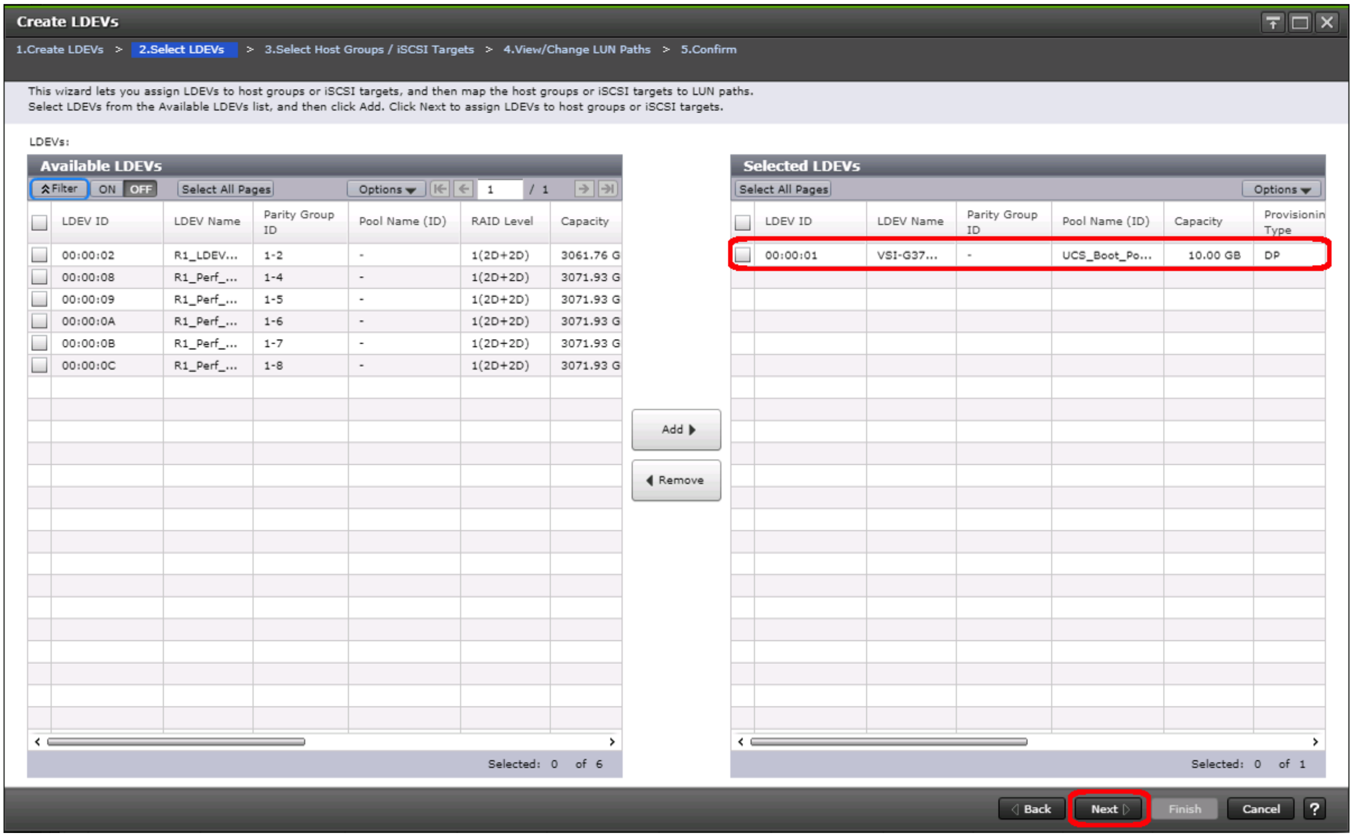
Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

This wizard lets you assign LDEVs to host groups or iSCSI targets, and then map the host groups or iSCSI targets to LUN paths.
Select LDEVs from the Available LDEVs list, and then click Add. Click Next to assign LDEVs to host groups or iSCSI targets.

LDEVs:

Available LDEVs

Filter ON OFF Select All Pages Options 1 / 1

□	LDEV ID	LDEV Name	Parity Group ID	Pool Name (ID)	RAID Level	Capacity
□	00:00:02	R1_LDEV...	1-2	-	1(2D+2D)	3061.76 G
□	00:00:08	R1_Perf_...	1-4	-	1(2D+2D)	3071.93 G
□	00:00:09	R1_Perf_...	1-5	-	1(2D+2D)	3071.93 G
□	00:00:0A	R1_Perf_...	1-6	-	1(2D+2D)	3071.93 G
□	00:00:0B	R1_Perf_...	1-7	-	1(2D+2D)	3071.93 G
□	00:00:0C	R1_Perf_...	1-8	-	1(2D+2D)	3071.93 G

Selected: 0 of 6

Selected LDEVs

□	LDEV ID	LDEV Name	Parity Group ID	Pool Name (ID)	Capacity	Provisionin Type
□	00:00:01	VSI-G37...	-	UCS_Boot_Po...	10.00 GB	DP

Selected: 0 of 1

7. The **Select Host Groups/iSCSI Targets** screen shows all of the host groups that can be assigned to the boot LDEV as a path.

8. Click Filter, then create an Attribute/Value filter:

   – Host Group Name

   – Using "contains" as a qualifier

   – <value which contains text unique to UCS server profile>

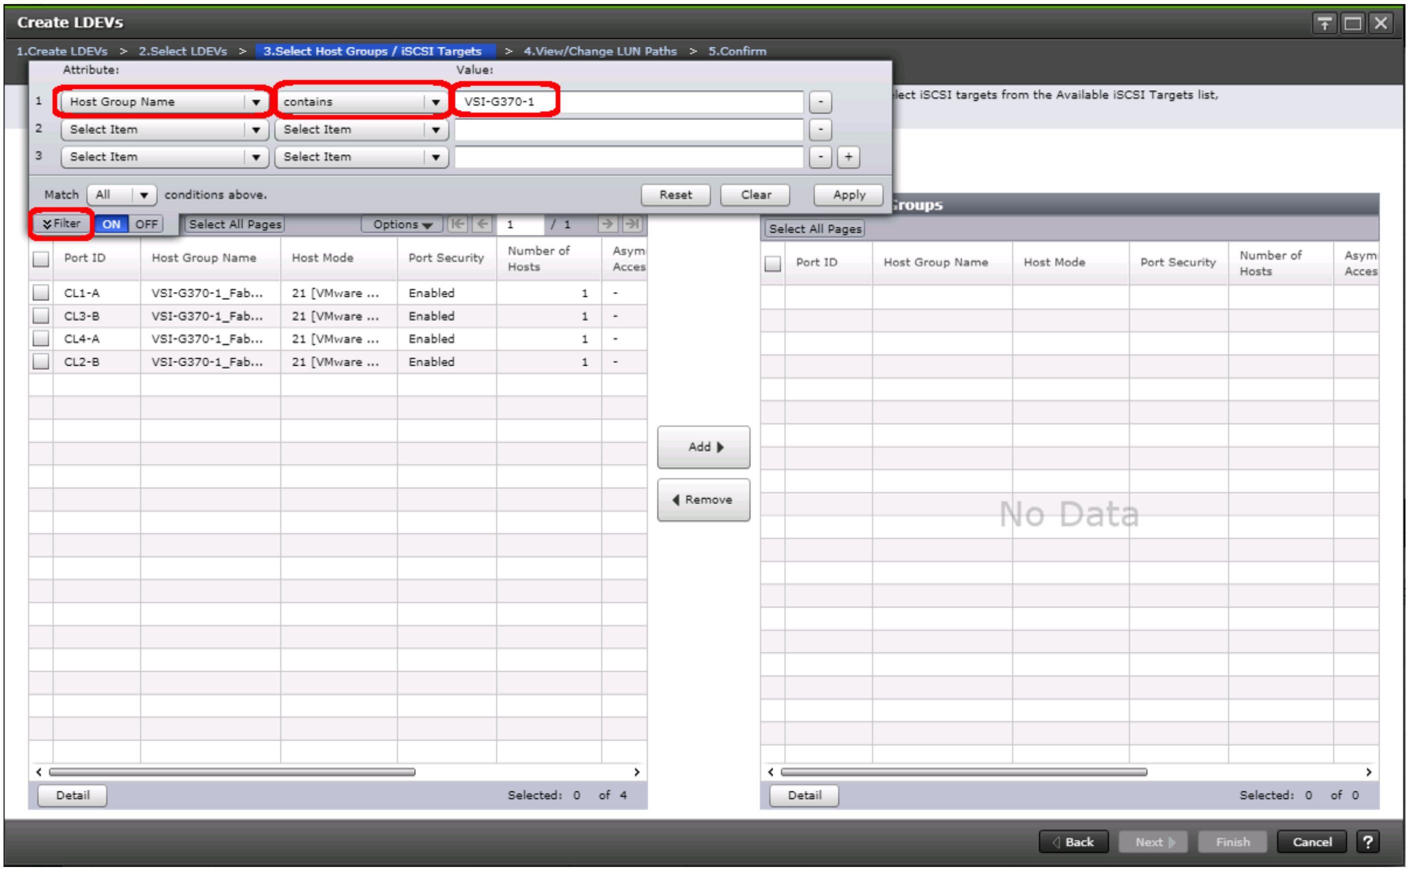

Create LDEVs

1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

ect iSCSI targets from the Available iSCSI Targets list,

	Attribute:		Value:
1	Host Group Name	contains	VSI-G370-1
2	Select Item	Select Item	
3	Select Item	Select Item	

Match All conditions above.

Reset Clear Apply

Filter ON OFF Select All Pages Options 1 / 1

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces
□	CL1-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-
□	CL3-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-
□	CL4-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-
□	CL2-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-

Detail Selected: 0 of 4

Add ▶ ◀ Remove

roups

Select All Pages

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces

No Data

Detail Selected: 0 of 0

Back Next Finish Cancel ?

9. Click **Apply**.

10. Click Filter again to hide the filter rules dialog box.

11. Select the checkboxes for the ports being used as boot LDEV paths in your configuration. Depending on the pathing design used, you may have fewer than four paths for the boot LDEV, but there should be a minimum of one path per fabric used.

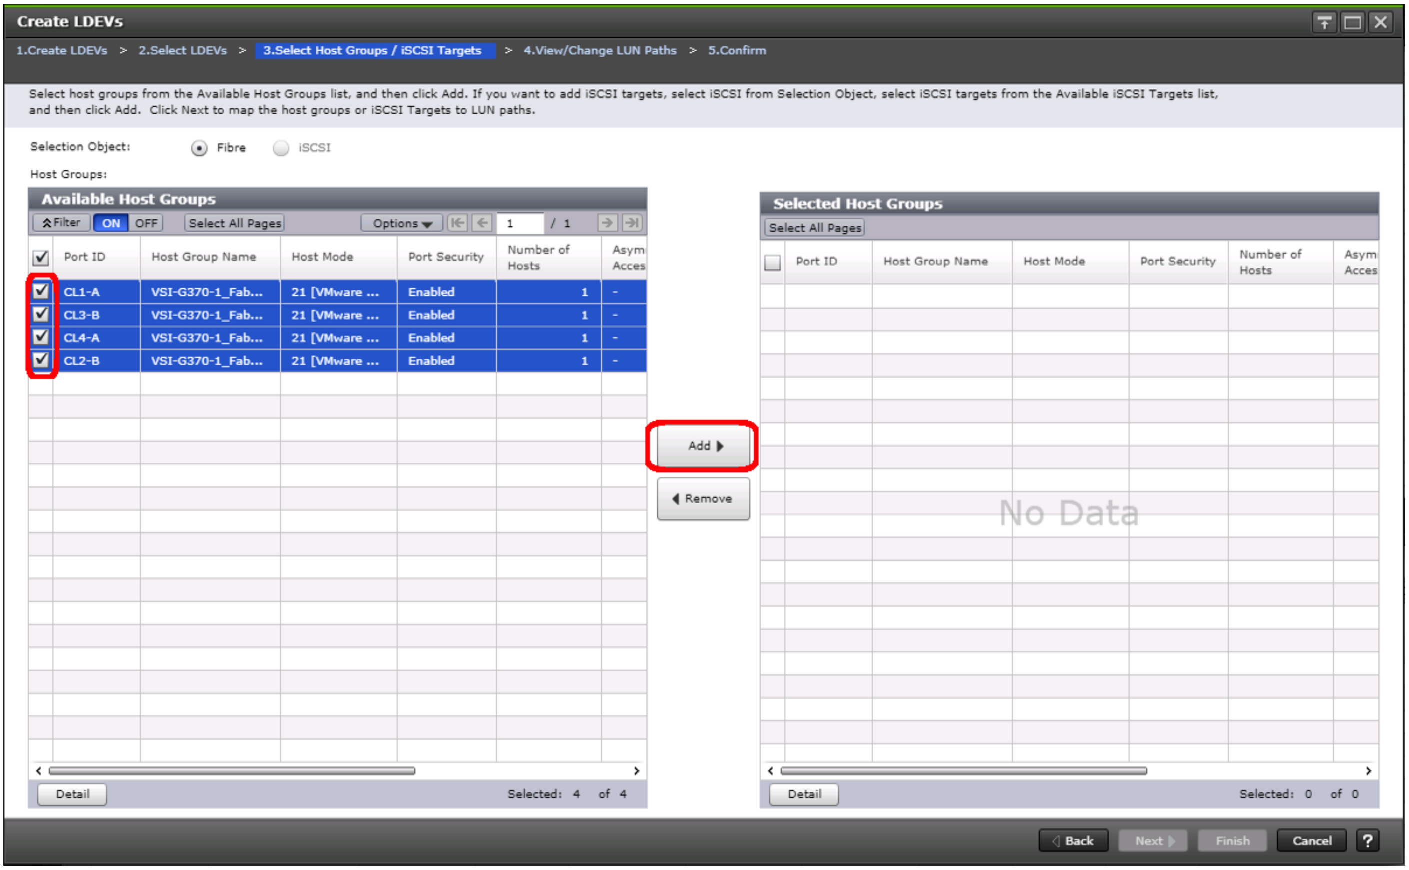

Create LDEVs

1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

Select host groups from the Available Host Groups list, and then click Add. If you want to add iSCSI targets, select iSCSI from Selection Object, select iSCSI targets from the Available iSCSI Targets list, and then click Add. Click Next to map the host groups or iSCSI Targets to LUN paths.

Selection Object: ● Fibre ○ iSCSI

Host Groups:

Available Host Groups

Filter ON OFF | Select All Pages | Options | 1 / 1

☑	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces
☑	CL1-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-
☑	CL3-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-
☑	CL4-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-
☑	CL2-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	-

Detail Selected: 4 of 4

Selected Host Groups

Select All Pages

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces

No Data

Detail Selected: 0 of 0

12. Click the **Add** to populate the **Selected Host Groups** pane with the selected host groups, then click **Next**.

13. The **View/Change LUN Paths** screen shows the LDEV you are adding paths to and the associated host LUN ID that will be presented to the host on a per path basis.

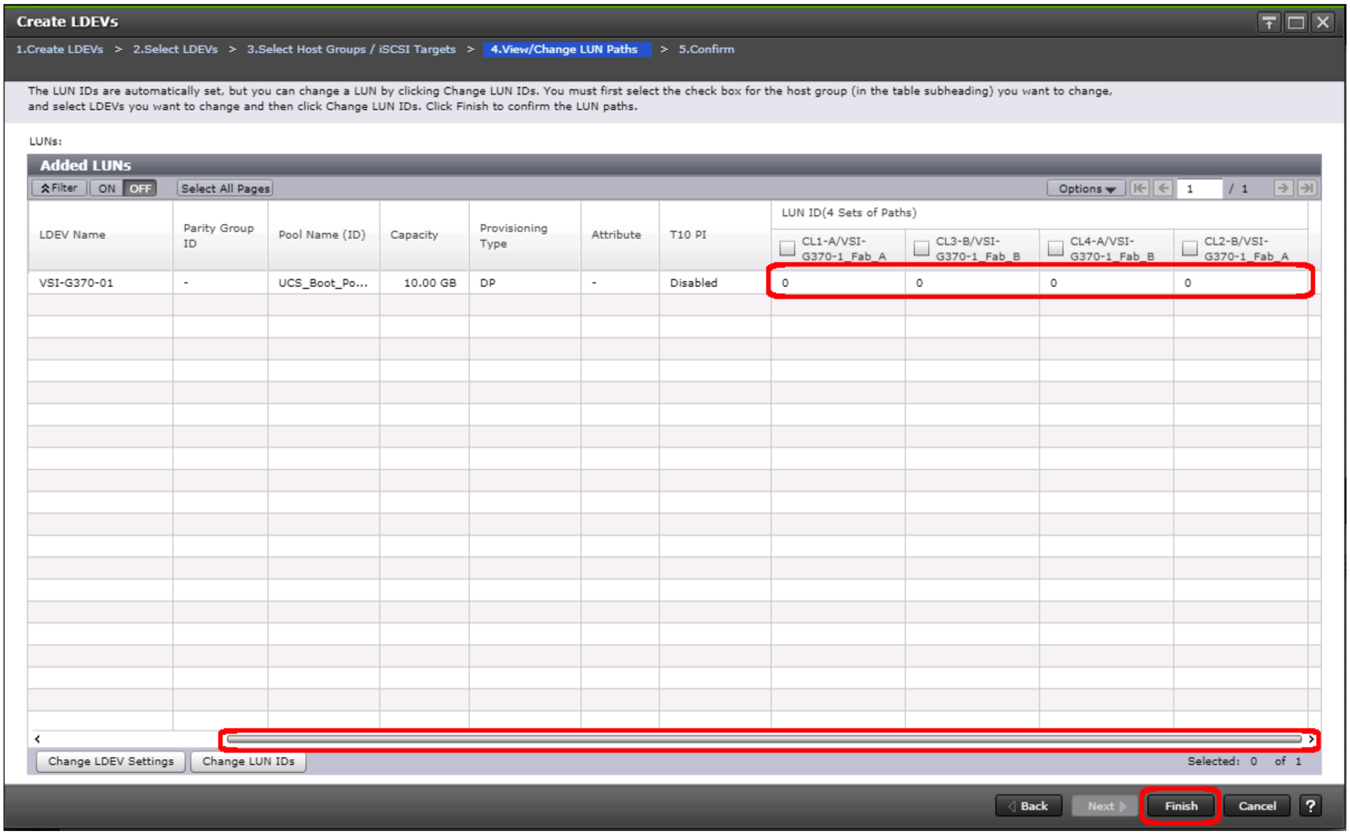

Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

The LUN IDs are automatically set, but you can change a LUN by clicking Change LUN IDs. You must first select the check box for the host group (in the table subheading) you want to change, and select LDEVs you want to change and then click Change LUN IDs. Click Finish to confirm the LUN paths.

LUNs:

Added LUNs

Filter
ON
OFF
Select All Pages
Options
1
/ 1

LDEV Name	Parity Group ID	Pool Name (ID)	Capacity	Provisioning Type	Attribute	T10 PI	LUN ID(4 Sets of Paths)			
							□ CL1-A/VSI-G370-1_Fab_A	□ CL3-B/VSI-G370-1_Fab_B	□ CL4-A/VSI-G370-1_Fab_B	□ CL2-B/VSI-G370-1_Fab_A
VSI-G370-01	-	UCS_Boot_Po...	10.00 GB	DP	-	Disabled	0	0	0	0

14. Use the scrollbar at the bottom of this screen to review the LUN ID assigned and ensure that all LUN IDs are set to zero, then click **Finish**.

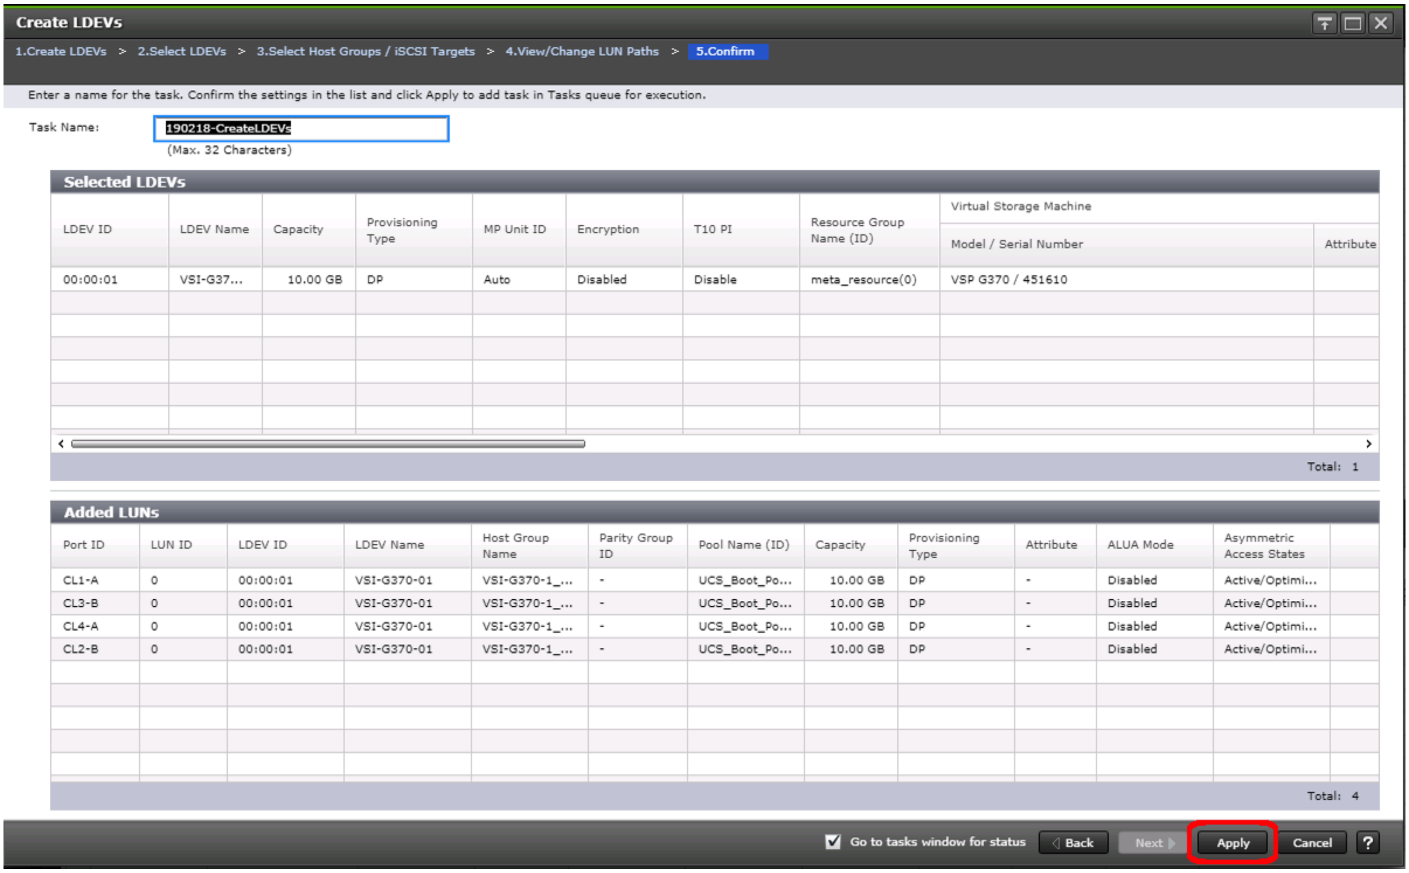

Create LDEVs

1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: 190218-CreateLDEVs
(Max. 32 Characters)

Selected LDEVs

LDEV ID	LDEV Name	Capacity	Provisioning Type	MP Unit ID	Encryption	T10 PI	Resource Group Name (ID)	Virtual Storage Machine	
								Model / Serial Number	Attribute
00:00:01	VSI-G37...	10.00 GB	DP	Auto	Disabled	Disable	meta_resource(0)	VSP G370 / 451610	

Total: 1

Added LUNs

Port ID	LUN ID	LDEV ID	LDEV Name	Host Group Name	Parity Group ID	Pool Name (ID)	Capacity	Provisioning Type	Attribute	ALUA Mode	Asymmetric Access States
CL1-A	0	00:00:01	VSI-G370-01	VSI-G370-1_...	-	UCS_Boot_Po...	10.00 GB	DP	-	Disabled	Active/Optimi...
CL3-B	0	00:00:01	VSI-G370-01	VSI-G370-1_...	-	UCS_Boot_Po...	10.00 GB	DP	-	Disabled	Active/Optimi...
CL4-A	0	00:00:01	VSI-G370-01	VSI-G370-1_...	-	UCS_Boot_Po...	10.00 GB	DP	-	Disabled	Active/Optimi...
CL2-B	0	00:00:01	VSI-G370-01	VSI-G370-1_...	-	UCS_Boot_Po...	10.00 GB	DP	-	Disabled	Active/Optimi...

Total: 4

15. Review the LDEV details and LUN ID configuration of the boot LDEV being created, then click **Apply** to create the LDEV and add paths to the UCS Service Profile.

16. Repeat steps 1-15 to create the boot LDEVs and to assign paths for all other UCS Service Profiles, using a unique LDEV name and associated Host Group Name associated to each UCS Service Profile.

Create Shared VMFS LDEVs and Add LDEV Paths

VMFS LDEVs need to be created for shared VMFS volumes used for virtual machine storage across multiple ESXi servers which share resources within a vSphere cluster. Prior to beginning these steps, ensure you have identified the fibre channel ports on the Hitachi VSP that will be used for presentation of the VMFS LDEVs to the UCS servers. Depending on the pathing design you are using, additional or fewer paths may be configured as compared to the steps below.

A minimum of two paths should be used for shared VMFS LDEVs (one path per fabric).

To create shared VMFS LDEVs and add LDEV paths, follow these steps:

1. From the left Explorer pane within Hitachi Storage Navigator, select the **Storage Systems** tab and expand the storage system being configured.

2. Expand the **Pools** element in the navigation tree and highlight the pool previously created for use as the backing storage for VMFS volumes, select the **Virtual Volumes** tab and click **Create LDEVs** to instantiate the Create LDEVs dialog.

3. Modify the following within the Create LDEVs dialog:

- **LDEV Capacity**: Enter the capacity desired for the VMFS LDEV.
- Number of LDEVs: 1
- **LDEV Name**: Provide a descriptive name and numeric identifier for the VMFS LDEV. For ease of identification, it is recommended that the cluster name or other identifier specific to the VMFS volume being configured be entered in the **Prefix** field.

4. Click **Add** and verify that the VMFS LDEV is listed in the right-hand **Selected LDEVs** pane, then click **Next**.

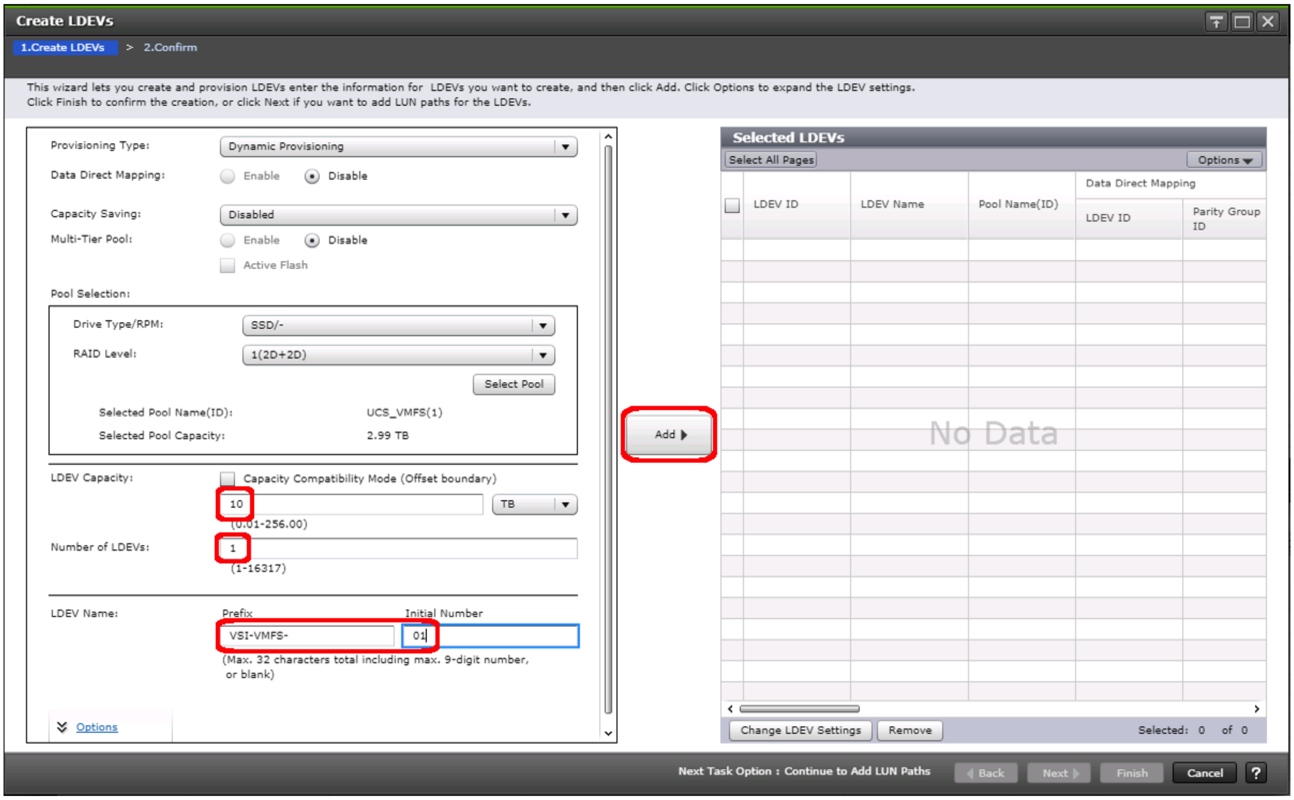

Create LDEVs

1.Create LDEVs > 2.Confirm

This wizard lets you create and provision LDEVs enter the information for LDEVs you want to create, and then click Add. Click Options to expand the LDEV settings.
Click Finish to confirm the creation, or click Next if you want to add LUN paths for the LDEVs.

Provisioning Type: Dynamic Provisioning

Data Direct Mapping: ○ Enable ● Disable

Capacity Saving: Disabled

Multi-Tier Pool: ○ Enable ● Disable

□ Active Flash

Pool Selection:

Drive Type/RPM: SSD/-

RAID Level: 1(2D+2D)

Selected Pool Name(ID): UCS_VMFS(1)

Selected Pool Capacity: 2.99 TB

LDEV Capacity: □ Capacity Compatibility Mode (Offset boundary)

10 TB
(0.01-256.00)

Number of LDEVs: 1
(1-16317)

LDEV Name: Prefix VSI-VMFS- Initial Number 01
(Max. 32 characters total including max. 9-digit number, or blank)

Selected LDEVs

□	LDEV ID	LDEV Name	Pool Name(ID)	Data Direct Mapping	
				LDEV ID	Parity Group ID

No Data

Selected: 0 of 0

Next Task Option : Continue to Add LUN Paths

5. The **Select LDEVs** screen shows the selected LDEVs to which the paths will be added.

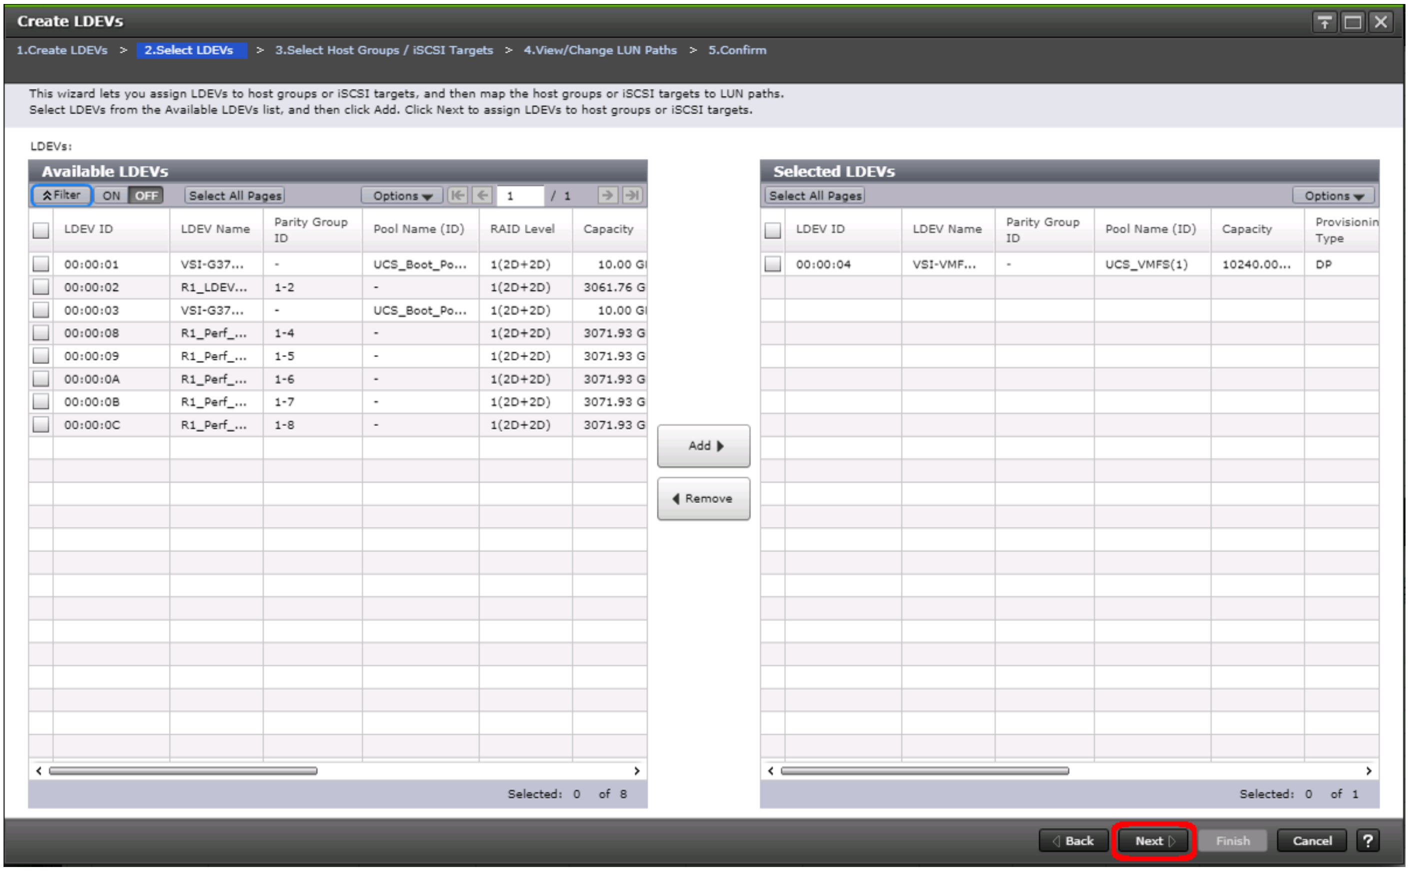
Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

This wizard lets you assign LDEVs to host groups or iSCSI targets, and then map the host groups or iSCSI targets to LUN paths.
Select LDEVs from the Available LDEVs list, and then click Add. Click Next to assign LDEVs to host groups or iSCSI targets.

LDEVs:

Available LDEVs

Filter ON OFF Select All Pages Options 1 / 1

□	LDEV ID	LDEV Name	Parity Group ID	Pool Name (ID)	RAID Level	Capacity
□	00:00:01	VSI-G37...	-	UCS_Boot_Po...	1(2D+2D)	10.00 GI
□	00:00:02	R1_LDEV...	1-2	-	1(2D+2D)	3061.76 G
□	00:00:03	VSI-G37...	-	UCS_Boot_Po...	1(2D+2D)	10.00 GI
□	00:00:08	R1_Perf_...	1-4	-	1(2D+2D)	3071.93 G
□	00:00:09	R1_Perf_...	1-5	-	1(2D+2D)	3071.93 G
□	00:00:0A	R1_Perf_...	1-6	-	1(2D+2D)	3071.93 G
□	00:00:0B	R1_Perf_...	1-7	-	1(2D+2D)	3071.93 G
□	00:00:0C	R1_Perf_...	1-8	-	1(2D+2D)	3071.93 G

Selected: 0 of 8

Selected LDEVs

□	LDEV ID	LDEV Name	Parity Group ID	Pool Name (ID)	Capacity	Provisionin Type
□	00:00:04	VSI-VMF...	-	UCS_VMFS(1)	10240.00...	DP

Selected: 0 of 1

6. Ensure the newly created VMFS LDEV is the only LDEV in the **Selected LDEVs** pane, then click **Next**.

7. The **Select Host Groups/iSCSI Targets** screen shows all of the host groups that can be assigned to the VMFS LDEV as a path.

8. Click Filter, then create multiple Attribute/Value:

   – Host Group Name

   – Using “contains” as a qualifier

   – <value which contains text unique to UCS server profiles to use the VMFS volume>

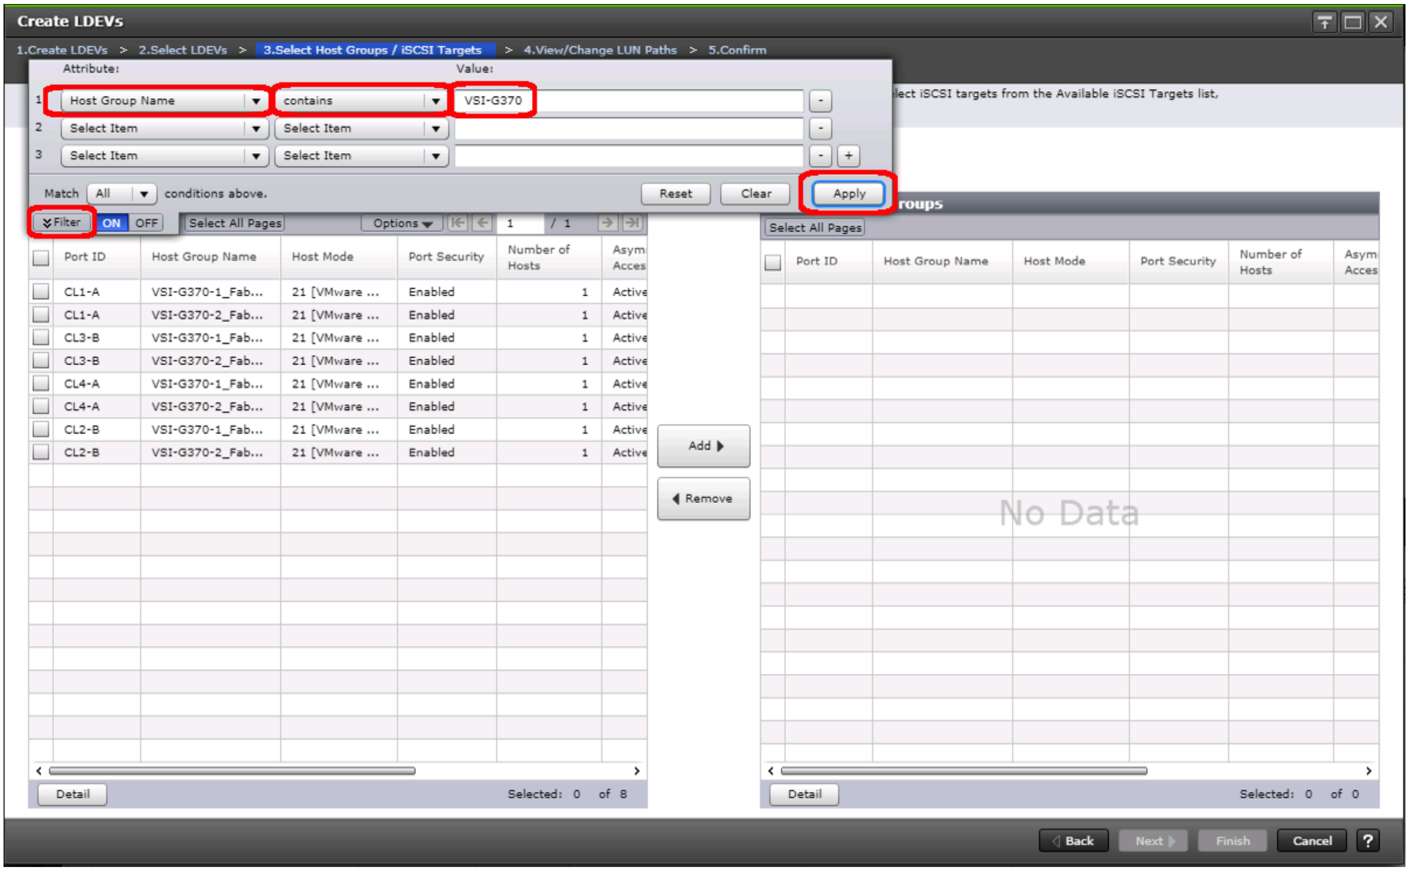
Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm
Reset
Clear
Apply
Filter ON OFF
Select All Pages
Options
Add
Remove
Detail
Back
Next
Finish
Cancel
?

Attribute: Value:

1 Host Group Name | contains | VSI-G370

2 Select Item | Select Item |

3 Select Item | Select Item |

Match All conditions above.

…lect iSCSI targets from the Available iSCSI Targets list,

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces
□	CL1-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
□	CL1-A	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active
□	CL3-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
□	CL3-B	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active
□	CL4-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
□	CL4-A	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active
□	CL2-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
□	CL2-B	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active

Selected: 0 of 8

…roups

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces

No Data

Selected: 0 of 0

9. Click **Apply**.

10. Click Filter again to hide the filter rules dialog box.

11. Select the checkboxes for the ports being used as VMFS LDEV paths in your configuration.

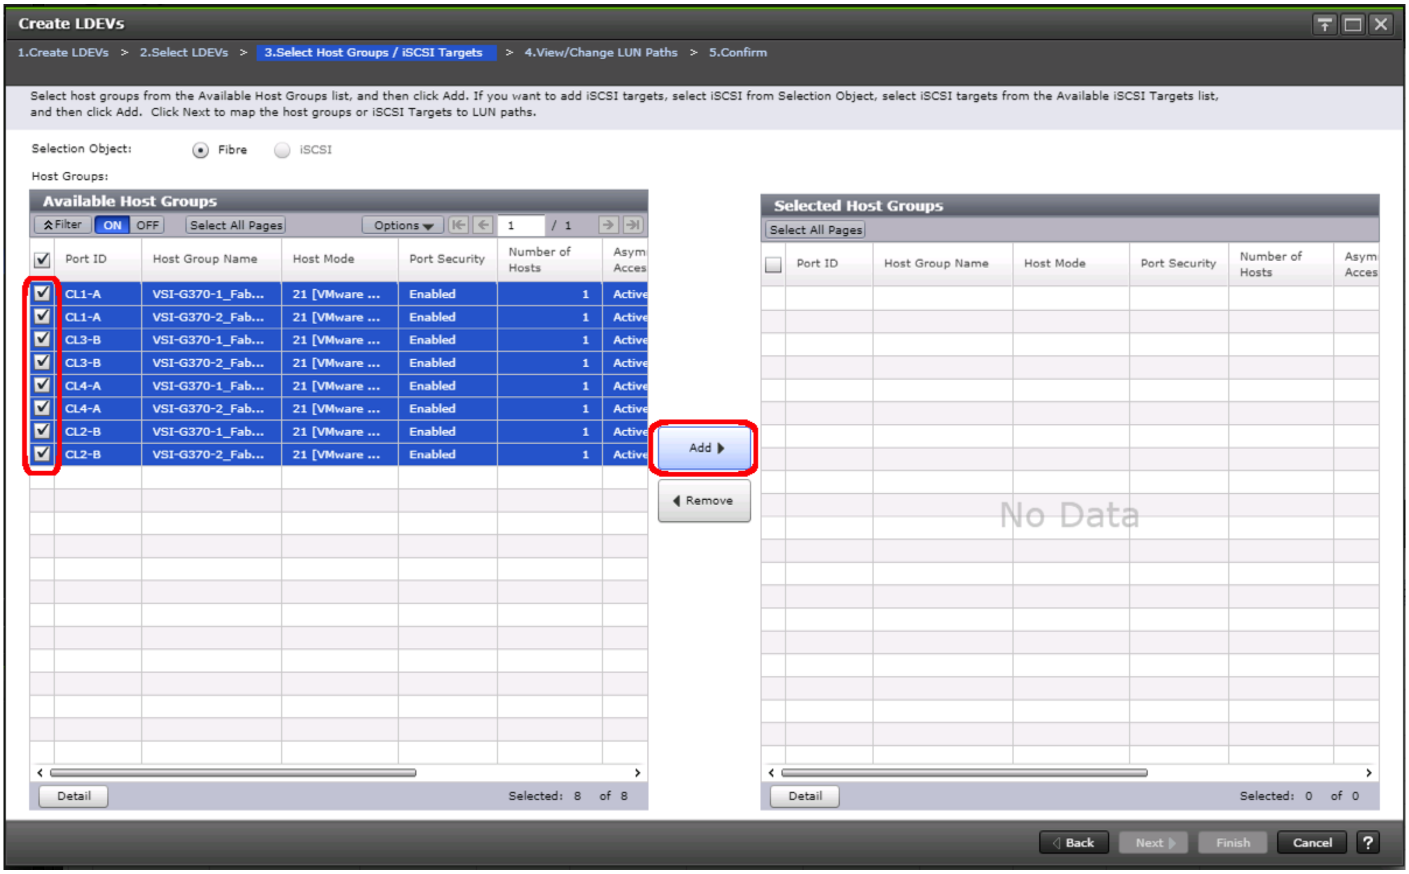

Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

Select host groups from the Available Host Groups list, and then click Add. If you want to add iSCSI targets, select iSCSI from Selection Object, select iSCSI targets from the Available iSCSI Targets list, and then click Add. Click Next to map the host groups or iSCSI Targets to LUN paths.

Selection Object: ● Fibre ○ iSCSI

Host Groups:

Available Host Groups

Filter ON OFF Select All Pages Options 1 / 1

☑	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces
☑	CL1-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL1-A	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL3-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL3-B	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL4-A	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL4-A	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL2-B	VSI-G370-1_Fab...	21 [VMware ...	Enabled	1	Active
☑	CL2-B	VSI-G370-2_Fab...	21 [VMware ...	Enabled	1	Active

Selected: 8 of 8

Selected Host Groups

□	Port ID	Host Group Name	Host Mode	Port Security	Number of Hosts	Asym Acces

No Data

Selected: 0 of 0

Depending on the pathing design used, you may have additional or fewer than four paths for the VMFS LDEV, but there should be a minimum of one path per fabric used.

12. Click **Add** to populate the **Selected Host Groups** pane with the selected host groups, then click **Next**.

13. The **View/Change LUN Paths** screen shows the LDEV you are adding paths to and the associated host LUN ID that will be presented to the host on a per path basis.

14. Use the scrollbar at the bottom of this screen to review the LUN ID assigned and ensure that all LUN IDs are set to a consistent value other than zero for all paths.

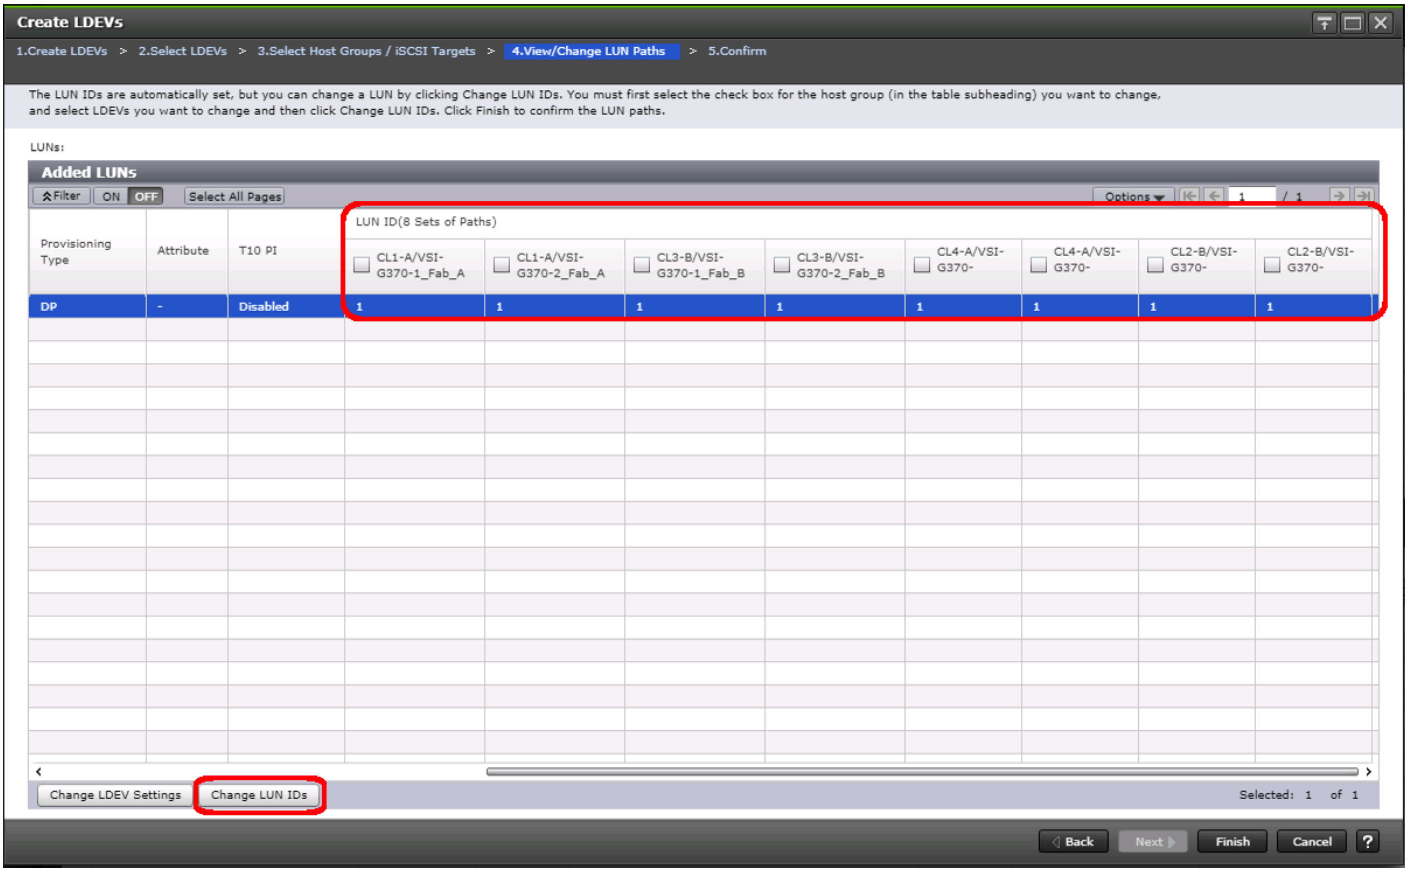
Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

The LUN IDs are automatically set, but you can change a LUN by clicking Change LUN IDs. You must first select the check box for the host group (in the table subheading) you want to change, and select LDEVs you want to change and then click Change LUN IDs. Click Finish to confirm the LUN paths.

LUNs:

Added LUNs

Filter ON OFF Select All Pages Options 1 / 1

Provisioning Type	Attribute	T10 PI	LUN ID(8 Sets of Paths)							
			□ CL1-A/VSI-G370-1_Fab_A	□ CL1-A/VSI-G370-2_Fab_A	□ CL3-B/VSI-G370-1_Fab_B	□ CL3-B/VSI-G370-2_Fab_B	□ CL4-A/VSI-G370-	□ CL4-A/VSI-G370-	□ CL2-B/VSI-G370-	□ CL2-B/VSI-G370-
DP	-	Disabled	1	1	1	1	1	1	1	1

If other LDEVs have been assigned to one host but not others, you will need to modify the Host LUN ID assignment to the next Host LUN ID that is consecutive across all hosts/paths.

15. Ensure you use the scrollbar at the bottom of the dialog to double-check that all Host LUN IDs are set consistently across all paths.

16. To do this, select the checkbox for all ports/paths listed, select the checkbox for the LDEV ID, then click **Change LUN IDs**.

17. The **Change LUN IDs** dialog will appear; enter the next Host LUN ID available across all paths, then click **Finish**.

18. Review the LDEV details and LUN ID configuration of the VMFS LDEV being created.

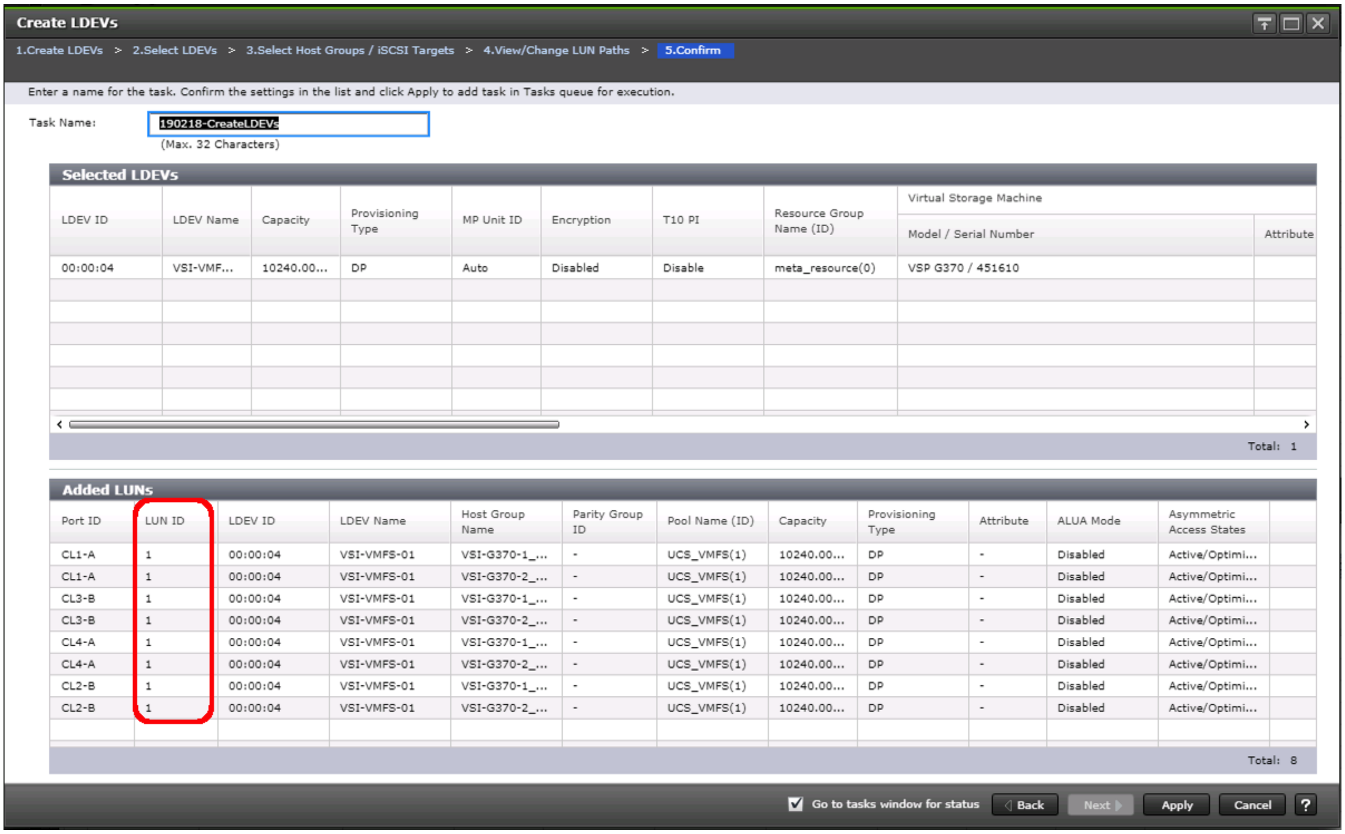

Create LDEVs
1.Create LDEVs > 2.Select LDEVs > 3.Select Host Groups / iSCSI Targets > 4.View/Change LUN Paths > 5.Confirm

Enter a name for the task. Confirm the settings in the list and click Apply to add task in Tasks queue for execution.

Task Name: 190218-CreateLDEVs
(Max. 32 Characters)

Selected LDEVs

LDEV ID	LDEV Name	Capacity	Provisioning Type	MP Unit ID	Encryption	T10 PI	Resource Group Name (ID)	Virtual Storage Machine	
								Model / Serial Number	Attribute
00:00:04	VSI-VMF...	10240.00...	DP	Auto	Disabled	Disable	meta_resource(0)	VSP G370 / 451610	

Total: 1

Added LUNs

Port ID	LUN ID	LDEV ID	LDEV Name	Host Group Name	Parity Group ID	Pool Name (ID)	Capacity	Provisioning Type	Attribute	ALUA Mode	Asymmetric Access States
CL1-A	1	00:00:04	VSI-VMFS-01	VSI-G370-1_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL1-A	1	00:00:04	VSI-VMFS-01	VSI-G370-2_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL3-B	1	00:00:04	VSI-VMFS-01	VSI-G370-1_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL3-B	1	00:00:04	VSI-VMFS-01	VSI-G370-2_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL4-A	1	00:00:04	VSI-VMFS-01	VSI-G370-1_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL4-A	1	00:00:04	VSI-VMFS-01	VSI-G370-2_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL2-B	1	00:00:04	VSI-VMFS-01	VSI-G370-1_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...
CL2-B	1	00:00:04	VSI-VMFS-01	VSI-G370-2_...	-	UCS_VMFS(1)	10240.00...	DP	-	Disabled	Active/Optimi...

Total: 8

☑ Go to tasks window for status

If the output is long enough, use the scrollbar of the **Added LUNs** window to ensure the **LUN ID** column contains the same LUN ID for each port listed.

19. Click **Apply** to create the LDEV and add paths to the UCS Service Profiles which will share this LDEV as a VMFS volume.

20. Repeat steps 1-18 to create additional shared VMFS LDEVs and to assign paths for all UCS Service Profiles which will share access to the VMFS LDEVs used for VMFS volumes.

ESXi Installation

This section explains how to install VMware ESXi 6.7 U1 in the environment.

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 Cisco Custom Image for ESXi 6.7 U1

The VMware Cisco Custom Image is required during the installation by manual access to the UCS KVM vMedia, or through a vMedia Policy covered in a previous subsection. If the Cisco Custom Image was not downloaded during the vMedia Policy setup, download it by following these steps:

1. To download the image, click this link: VMware vSphere Hypervisor Cisco Custom Image (ESXi) 6.7 U1.
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

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

To log into the Cisco UCS environment, follow these steps:

1. Open a web browser to https:// <<var_ucs_mgmt_vip>>
2. Select the Launch UCS Manager Section in the HTML section to pull up the UCSM HTML5 GUI.
3. Enter `admin` for the **Username**, and provide the password used during setup.
4. Click **Servers** -> **Service Profiles** and select the first host provisioned, which should be named `VSI-FC-G370-1`.
5. Click Reset to ensure that the boot LUN is properly recognized by the UCS Service Profile.
6. Click the **KVM Console** option within **Actions** and accept the KVM server certificate in the new window or browser tab that is spawned for the KVM session.
7. Click the link within the new window or browser tab to load the KVM client application.

Set Up VMware ESXi Installation

Skip this step if you are 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 icon in the upper right of the screen.
2. Click Activate Virtual Devices

3. Click Virtual Media again and select Map CD/DVD.

4. Browse to the ESXi installer ISO image file and click Open.

5. Click Map Device.

6. Click the KVM tab to monitor the server boot.

7. Boot the server by selecting Boot Server and clicking OK, then click OK again.

Install ESXi

To install VMware ESXi to the FC bootable LUN of the hosts, follow these steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.

2. After the installer is finished loading, press Enter to continue with the installation.

3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

4. Select the Boot LUN (10.00 GiB) that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

5. Select the appropriate keyboard layout and press Enter.

6. Enter and confirm the root password and press Enter.

7. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.

8. After the installation is complete, if using locally mapped Virtual Media, click the Virtual Media tab and clear the checkmark next to the ESXi installation media. Click Yes.

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer. If using a vMedia Policy, this will be unnecessary as the vMedia will appear after the installed OS.

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

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, follow these steps on each ESXi host:

1. After the server has finished rebooting, press F2 to customize the system.

2. Log in as `root`, enter the corresponding password, and press Enter to log in.

3. (Optional)Select **Troubleshooting Options** and press Enter.

4. (Optional)Press Enter for Enable ESXi Shell.

5. (Optional)Scroll to **Enable SSH** and press Enter.

6. (Optional)Press Esc to return to the main menu.
7. Select the **Configure Management Network** option and press Enter.
8. Select **Network Adapters** option leave vmnic0 selected, arrow down to vmnic1 and press space to select vmnic1 as well and press Enter.
9. Select the **VLAN (Optional)** option and press Enter.
10. Enter the `<<var_ib_mgmt_vlan_id>>` and press Enter.
11. From the Configure Management Network menu, select **IPv4 Configuration** and press Enter.
12. Select the Set Static IP Address and Network Configuration option by using the space bar.
13. Enter `<<var_vm_host_infra_01_ip>>` for the **IPv4 Address** for managing the first ESXi host.
14. Enter `<<var_ib_mgmt_vlan_netmask_length>>` for the **Subnet Mask** for the first ESXi host.
15. Enter `<<var_ib_gateway_ip>>` for the **Default Gateway** for the first ESXi host.
16. Press Enter to accept the changes to the IPv4 configuration.
17. Select the **DNS Configuration** option and press Enter.

Since the IP address is assigned manually, the DNS information must also be entered manually.

18. Enter the IP address of `<<var_nameserver_ip>>` for the **Primary DNS Server**.
19. Optional: Enter the IP address of the **Secondary DNS Server**.
20. Enter the fully qualified domain name (FQDN) for the first ESXi host.
21. Press Enter to accept the changes to the DNS configuration.
22. Select the **IPv6 Configuration** option and press Enter.
23. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
24. Press Esc to exit the Configure Management Network submenu.
25. Press Y to confirm the changes and return to the main menu.
26. The ESXi host reboots. After reboot, press F2 and log back in as root.
27. Select Test Management Network to verify that the management network is set up correctly and press Enter.
28. Press Enter to run the test.
29. Press Enter to exit the window, and press Esc to log out of the VMware console.
30. Repeat steps 1-29 for additional hosts provisioned, using appropriate values.

Log into VMware ESXi Hosts by Using VMware Host Client

To log into the `esxi-x` (x is server number 1-8) ESXi host by using the VMware Host Client, follow these steps:

1. Open a web browser on the management workstation and navigate to the `esxi-x` management IP address. Respond to any security prompts.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.
5. Repeat steps 1-4 to log into all the ESXi hosts in a separate browser tabs or windows.

The first host will need to go through the initial configuration using the VMware Host Client if a vCenter Appliance is being installed to the VSI cluster. Subsequent hosts can be configured directly to the vCenter Server after it is installed to the first ESXi host, or all hosts can be configured directly within the vCenter if a pre-existing server is used that is outside of the deployed converged infrastructure.

Set Up VMkernel Ports and Virtual Switch

To set up the VMkernel ports and the virtual switches on all the ESXi hosts, follow these steps:

1. From the Host Client, select **Networking** within the Navigator window.
2. In the center pane, select the Port groups tab.
3. Right-click the *VM Network* port group and select the **Remove** option.
4. Right-click the *Management Network* and select **Edit Settings**.
5. Expand NIC teaming and select vmnic1 within the Failover order section.
6. Click the Mark standby option.
7. Click Save.
8. Click on the **Add port group** option.
9. Name the port group *IB-Mgmt*.
10. Set the VLAN ID to <<IB-Mgmt VLAN ID>>.
11. Click **Add**.
12. Right-click the *IB-Mgmt* port group and select the Edit Settings option.
13. Expand NIC teaming and select **Yes** within the Override failover order section.
14. Select vmnic1 within the Failover order section.
15. Click on the **Mark standby** option.

16. Click Save.

17. Select the Virtual switches tab.

18. Right-click vSwitch0 and select **Edit settings**.

19. Change the MTU to 9000.

20. Expand NIC teaming and highlight vmnic1. Select **Mark active**.

21. Click Save.

22. Select the VMkernel NICs tab in the center pane.

23. Select Add VMkernel NIC.

24. Enter vMotion within the New port group section.

25. Set the VLAN ID to <<vMotion VLAN ID>>

26. Change the MTU to 9000.

27. Click the Static option within IPv4 settings and expand the section.

28. Enter the Address and Subnet mask to be used for the ESXi vMotion IP.

29. Change the TCP/IP stack to vMotion stack.

30. Click **Create**.

Optionally, with 40GE vNICs, you can create two additional vMotion VMkernel NICs in the same subnet and VLAN to take advantage of the bandwidth. These need to be in new dedicated port groups for the new vMotion VMkernels.

31. Select the Port groups tab.

32. Right-click the vMotion port group and select the **Edit settings** option.

33. Expand the NIC Teaming section and select **Yes** for Override failover order.

34. Highlight vmnic0 and select **Mark standby**.

35. Highlight vmnic1 and select **Mark active**.

36. Click Save.

37. Repeat steps 32-36 if additional vMotion port groups were created.

Create the VSI Datacenter

If a new datacenter is needed, follow these steps on the vCenter:

1. Connect to the vSphere Web Client and right-click the vCenter icon in the top left under the Hosts and Clusters tab, selecting the New Datacenter option from the drop-down list, or directly connect the Create Datacenter from the Getting Started page.

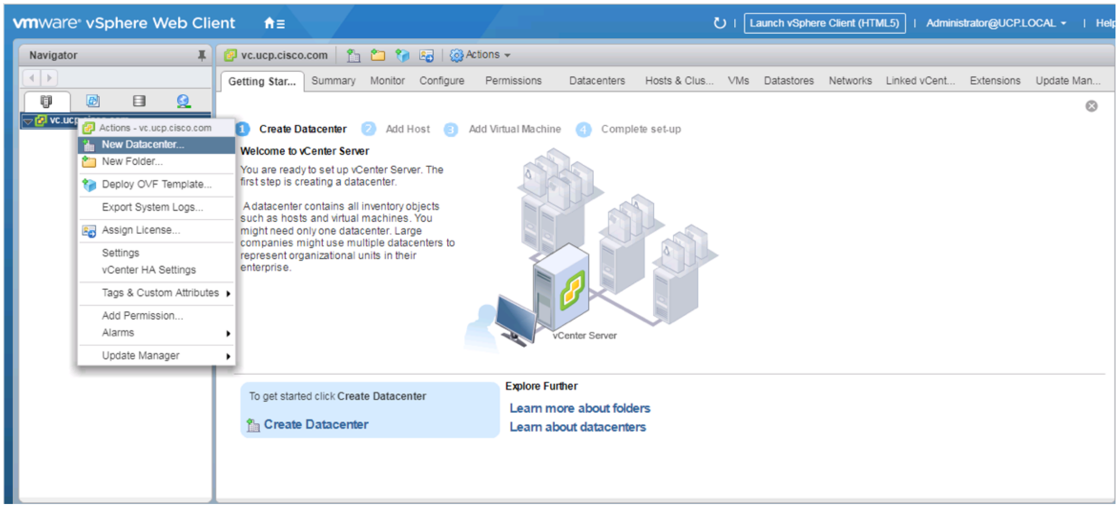

Welcome to vCenter Server

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

2. From the New Datacenter pop-up dialogue enter in a datacenter name and click OK.

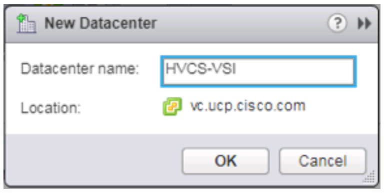

Datacenter name: HVCS-VSI

Location: vc.ucp.cisco.com

Add the VMware ESXi Hosts Using the VMware vSphere Web Client

To add the VMware ESXi Hosts using the VMware vSphere Web Client, follow these steps:

1. From the Hosts and Clusters tab, right-click the new or existing Datacenter within the Navigation window and select **New Cluster…** from the drop-down list.

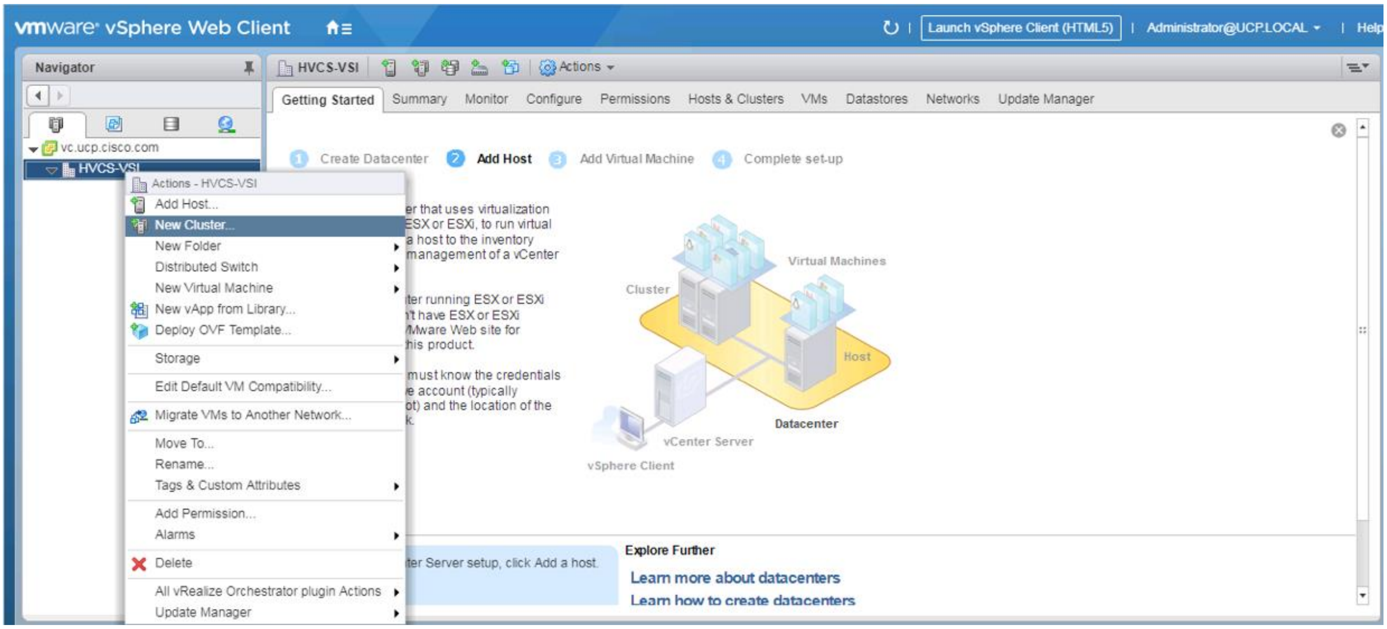

2. Enter a name for the new cluster, select the DRS and HA checkboxes, leaving all other options with the defaults.

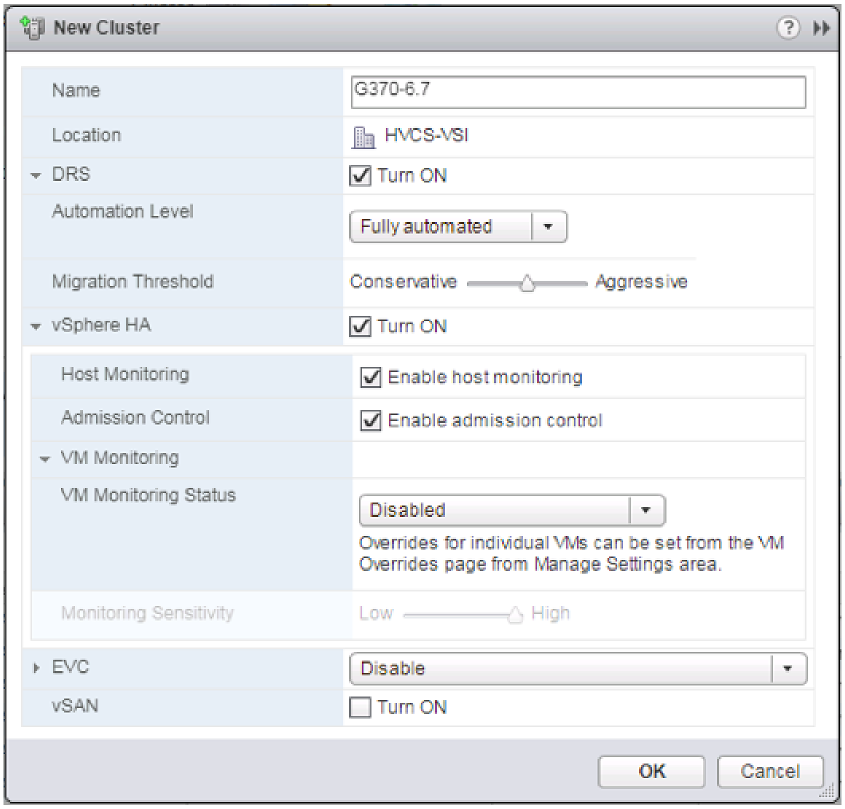

Name	G370-6.7
Location	HVCS-VSI
DRS	☑ Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative — Aggressive
vSphere HA	☑ Turn ON
Host Monitoring	☑ Enable host monitoring
Admission Control	☑ Enable admission control
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low — High
EVC	Disable
vSAN	□ Turn ON

3. Click OK to create the cluster.

4. Right-click the newly created cluster and select the **Add Host…** drop-down option.

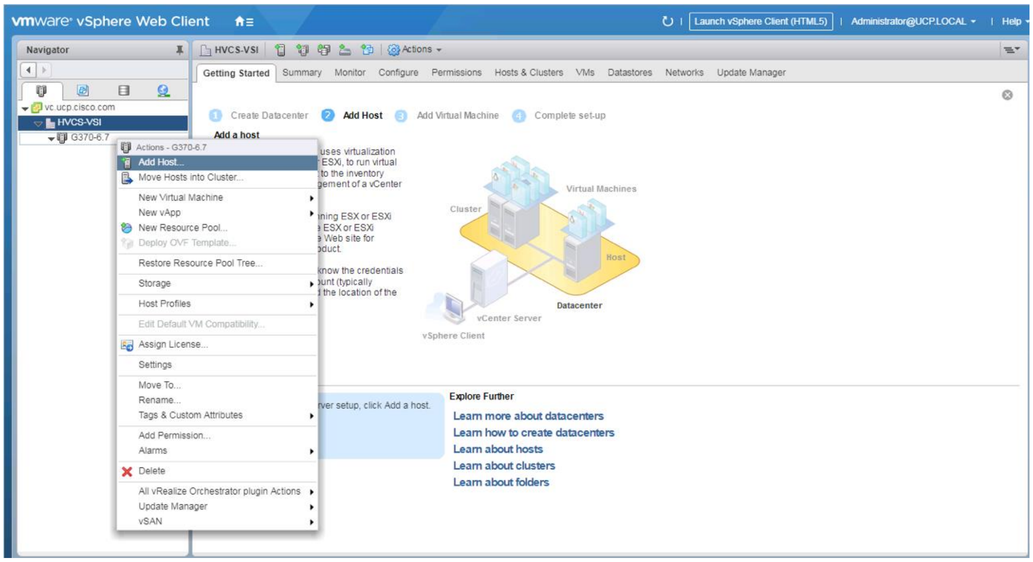

5. Enter the IP or FQDN of the first ESXi host and click Next.

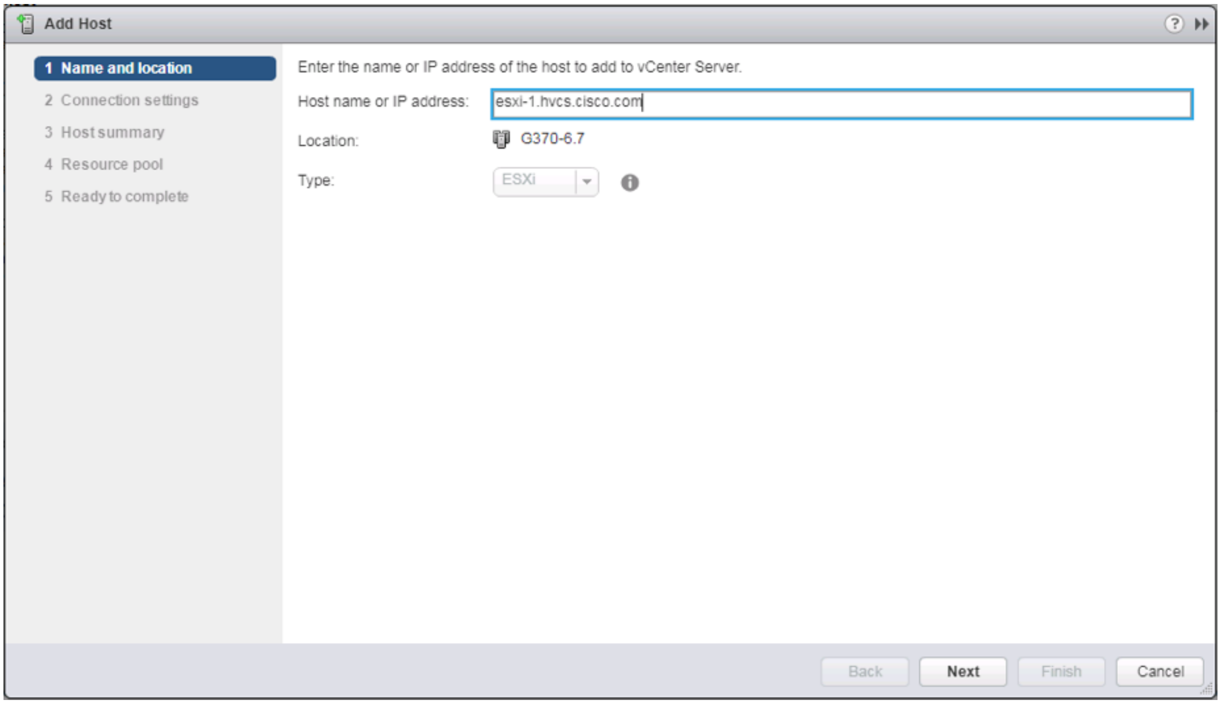

Enter the name or IP address of the host to add to vCenter Server.

Host name or IP address: esxi-1.hvcs.cisco.com

Location: G370-6.7

Type: ESXi

6. Enter `root` for the User Name, provide the password set during initial setup and click Next.

7. Click Yes in the Security Alert pop-up to confirm the host’s certificate.

8. Click Next past the Host summary dialogue.

9. Provide a license by clicking the green + icon under the License title, select an existing license or skip past the Assign license dialogue by clicking Next.

10. Leave the lockdown mode Disabled within the Lockdown mode dialogue window and click Next.

11. Skip past the Resource pool dialogue by clicking Next.

12. Confirm the Summary dialogue and add the ESXi host to the cluster by clicking Next.

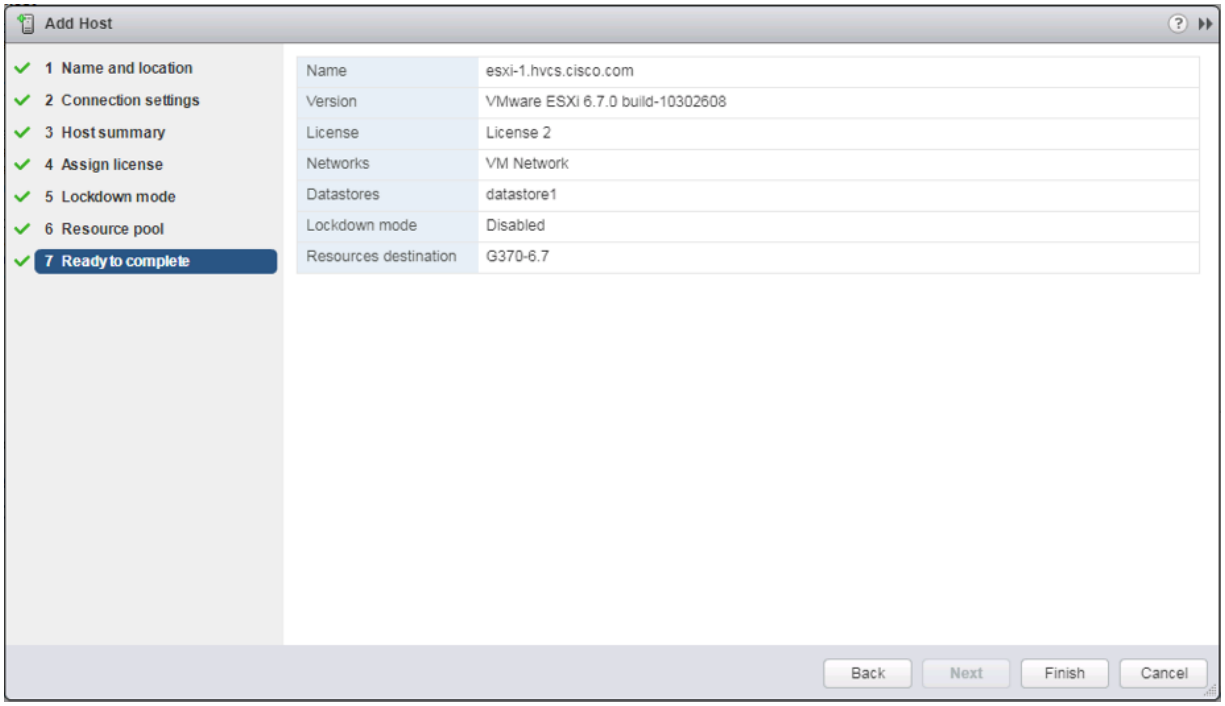

Add Host

✓ 1 Name and location
✓ 2 Connection settings
✓ 3 Host summary
✓ 4 Assign license
✓ 5 Lockdown mode
✓ 6 Resource pool
✓ 7 Ready to complete

Name	esxi-1.hvcs.cisco.com
Version	VMware ESXi 6.7.0 build-10302608
License	License 2
Networks	VM Network
Datastores	datastore1
Lockdown mode	Disabled
Resources destination	G370-6.7

13. Repeat steps 1-12 for each ESXi host to be added to the cluster.

Create VMware vDS for Application Traffic

The VMware vDS setup will consist a single vDS for Application traffic.

To configure the first VMware vDS, follow these steps:

1. Right-click the HVCS-VSI datacenter and select **Distributed Switch** > **New Distributed Switch**…

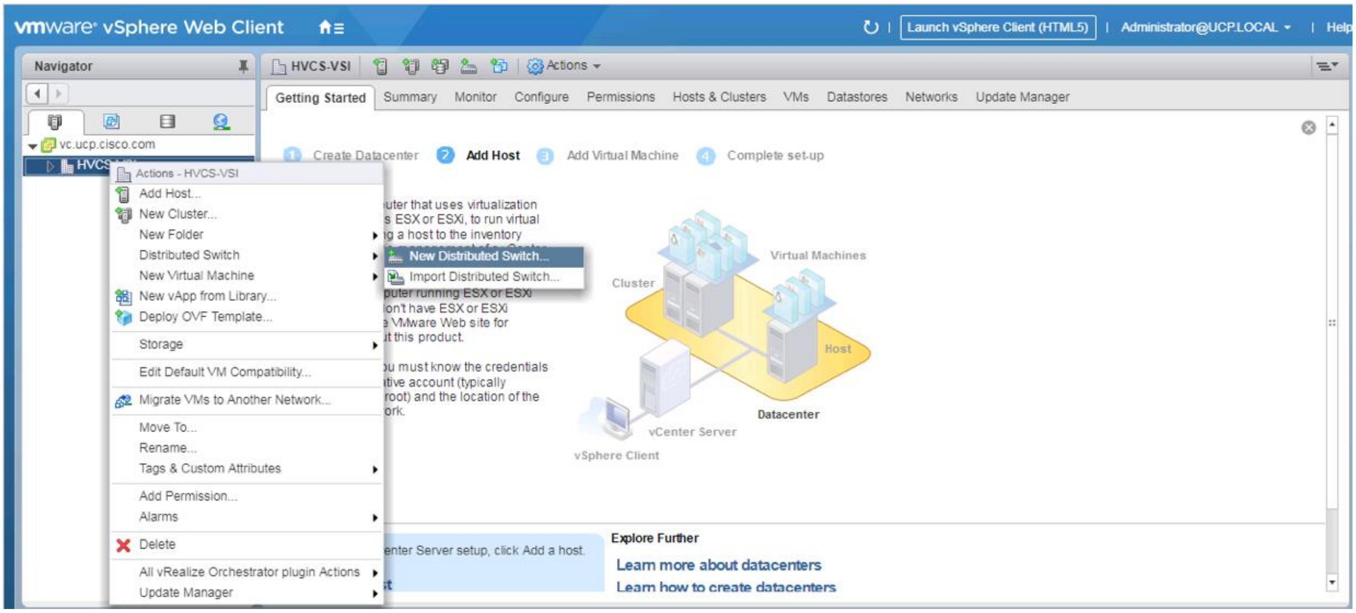

2. Give the Distributed Switch a descriptive name and click Next.

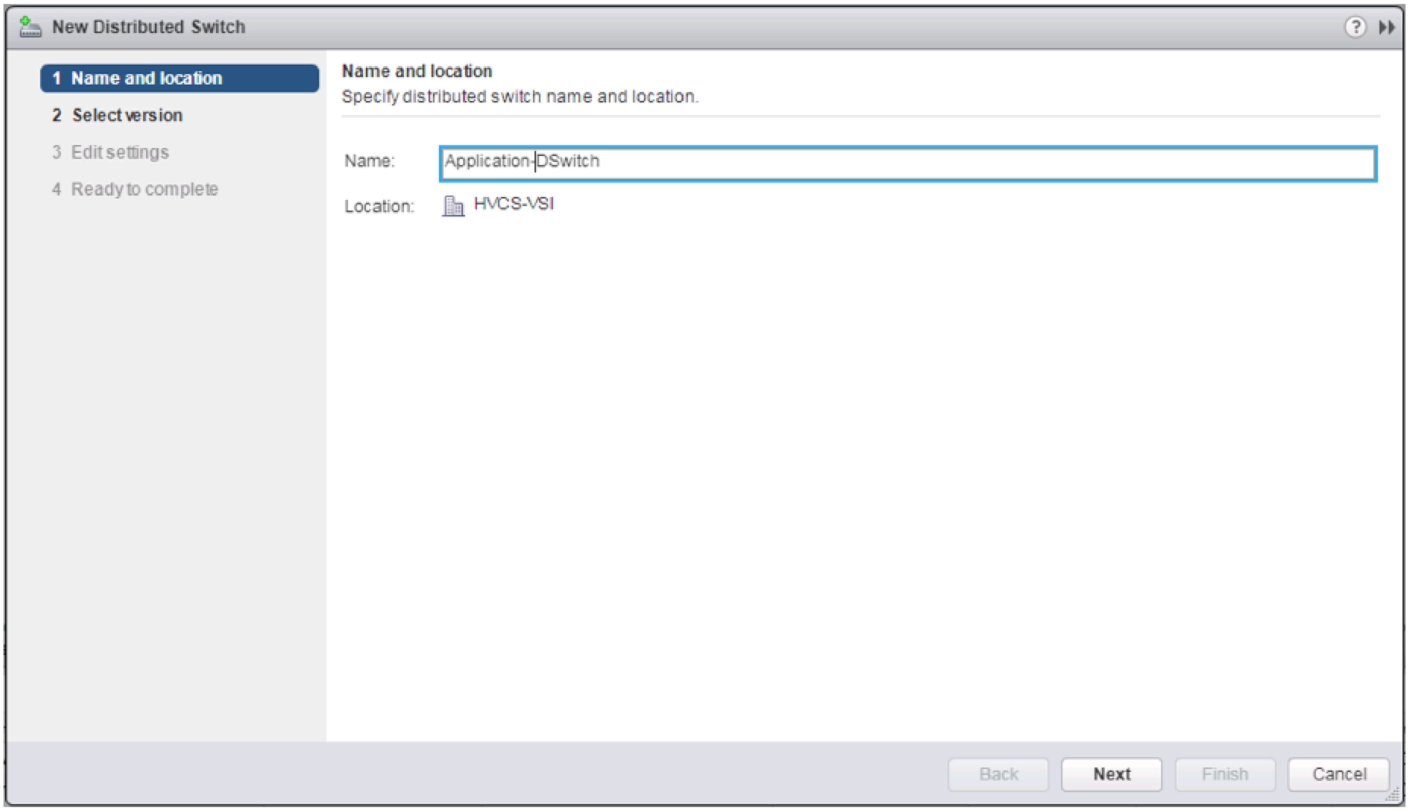

Name and location

Specify distributed switch name and location.

Name: Application-DSwitch

Location: HVCS-VSI

3. Make sure Distributed switch: 6.5.0 is selected if supporting vSphere 6.5 hosts and click Next.

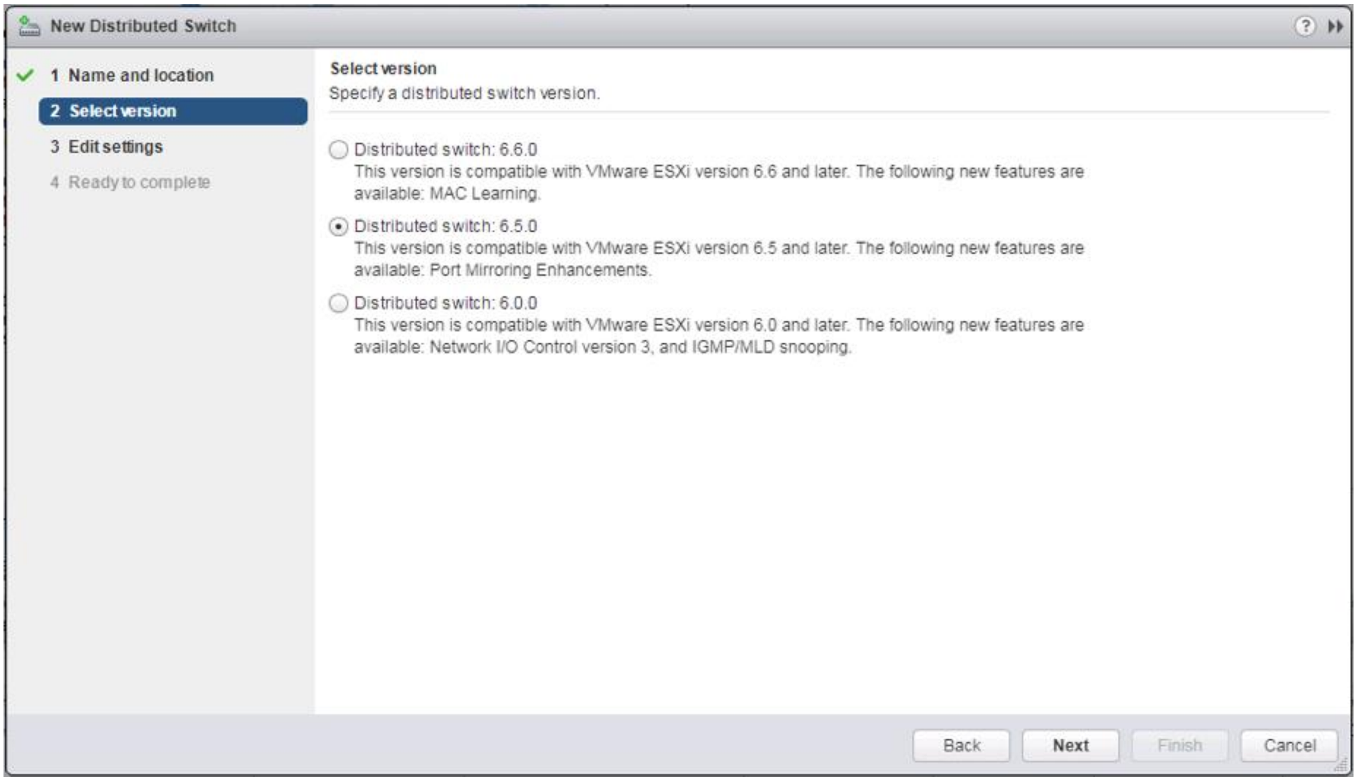

Select version

Specify a distributed switch version.

○ Distributed switch: 6.6.0
This version is compatible with VMware ESXi version 6.6 and later. The following new features are available: MAC Learning.

● Distributed switch: 6.5.0
This version is compatible with VMware ESXi version 6.5 and later. The following new features are available: Port Mirroring Enhancements.

○ Distributed switch: 6.0.0
This version is compatible with VMware ESXi version 6.0 and later. The following new features are available: Network I/O Control version 3, and IGMP/MLD snooping.

4. 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 App-201 for the name of the default Port group to be created. Click Next.

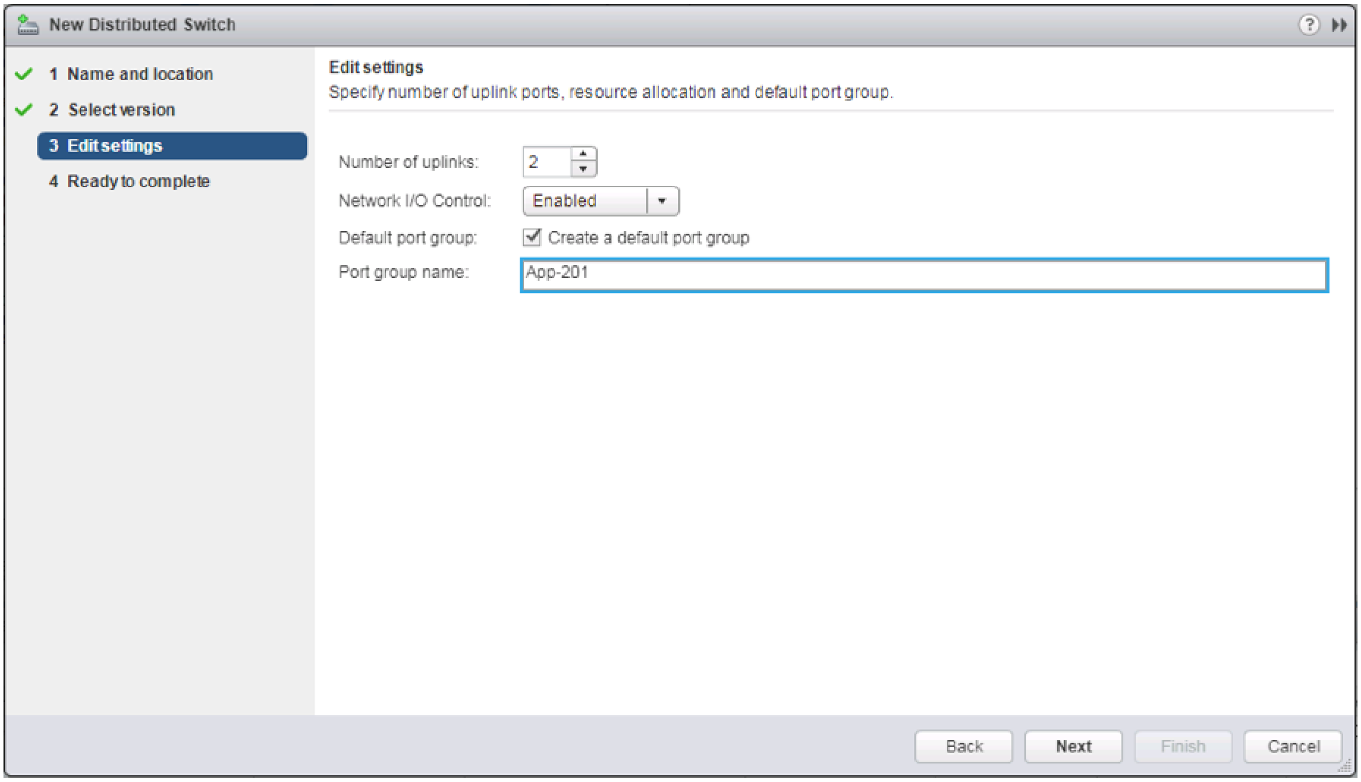

Edit settings

Specify number of uplink ports, resource allocation and default port group.

Number of uplinks: 2

Network I/O Control: Enabled

Default port group: ☑ Create a default port group

Port group name: App-201

5. Review the information and click Finish to complete creating the vDS.

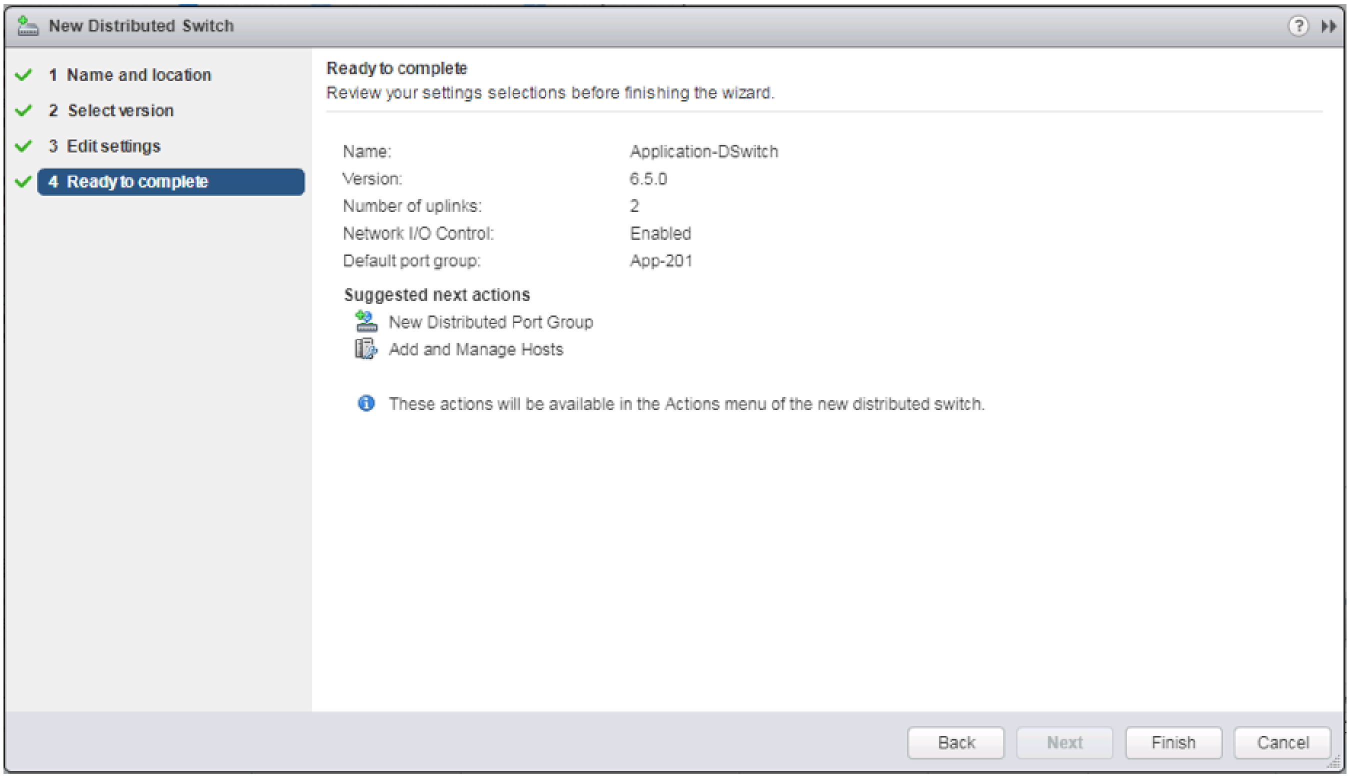

Ready to complete

Review your settings selections before finishing the wizard.

Name:	Application-DSwitch
Version:	6.5.0
Number of uplinks:	2
Network I/O Control:	Enabled
Default port group:	App-201

Suggested next actions

New Distributed Port Group

Add and Manage Hosts

These actions will be available in the Actions menu of the new distributed switch.

6. Right-click the newly created App-DSwitch vDS, and select **Settings** -> **Edit Settings…**

7. Click the **Advanced** option for the Edit Settings window and change the MTU from 1500 to 9000.

8. Click OK to save the changes.

9. Right-click the App-201 Distributed Port Group, and select **Edit Settings…**

10. Click **VLAN**, changing **VLAN type** from None to VLAN, and enter in the appropriate VLAN number for the first application network.

> The application Distributed Port Groups will not need to adjust their NIC Teaming as they will be Active/Active within the two vNICs uplinks associated to the Application-DSwitch, using the default VMware Route based on originating virtual port load balancing algorithm.

11. Click OK to save the changes.

12. Right-click the Application-DSwitch, selecting Distributed Port Group -> New Distributed Port Group… for any additional application networks to be created, setting the appropriate VLAN for each new Distributed Port Group.

Add Hosts to Application-DSwitch

To add the hosts to the newly created vDS from the Navigator, follow these steps:

1. Select the newly created Application-DSwitch.

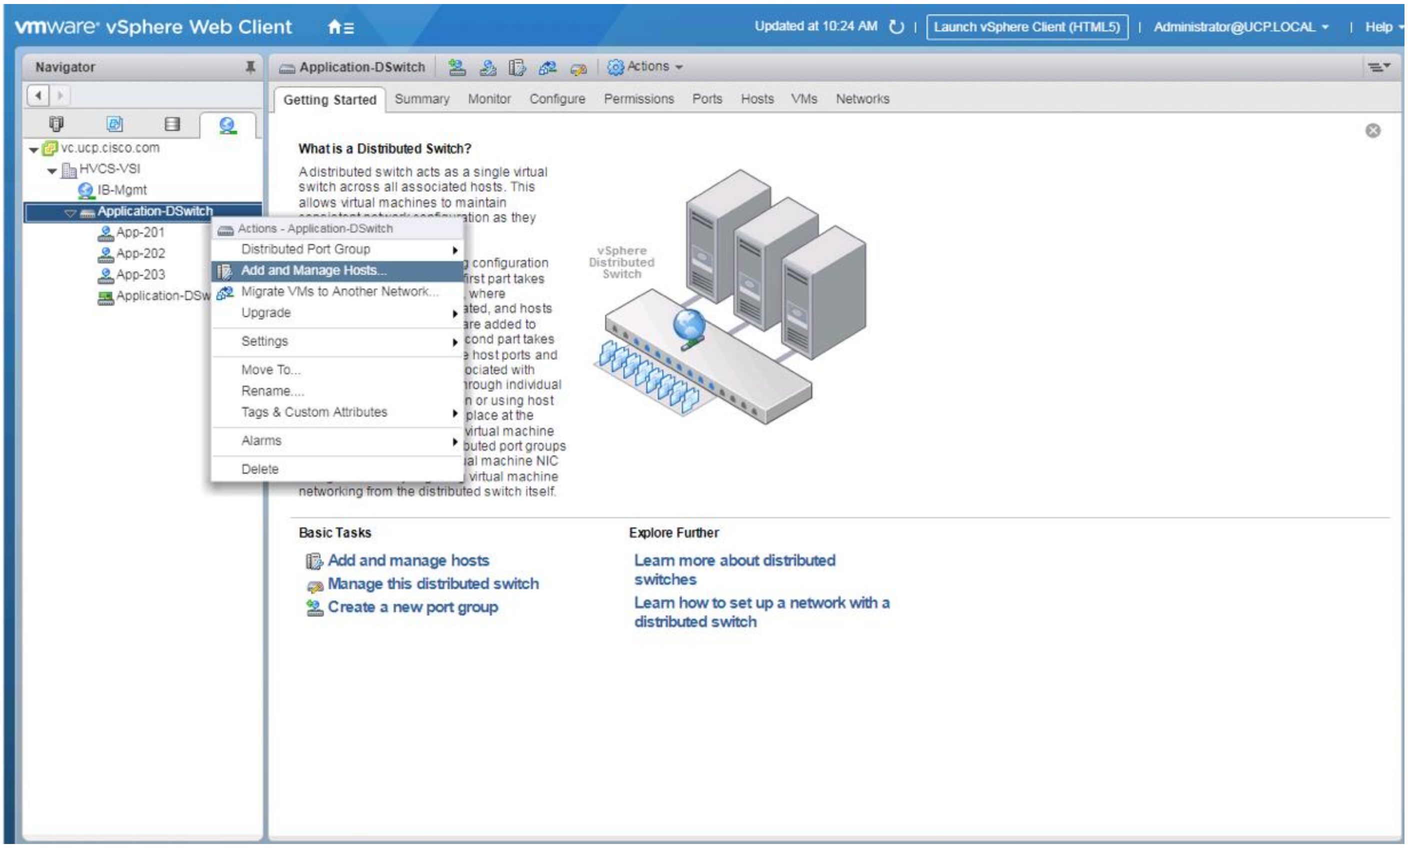

What is a Distributed Switch?

A distributed switch acts as a single virtual switch across all associated hosts. This allows virtual machines to maintain consistent network configuration as they … configuration … first part takes … where … ated, and hosts … are added to … cond part takes … host ports and … ociated with … rough individual … n or using host … place at the … virtual machine … buted port groups … al machine NIC … virtual machine networking from the distributed switch itself.

Basic Tasks

- Add and manage hosts
- Manage this distributed switch
- Create a new port group

Explore Further

- Learn more about distributed switches
- Learn how to set up a network with a distributed switch

2. Right-click it and select the Add and Manage Hosts… option.

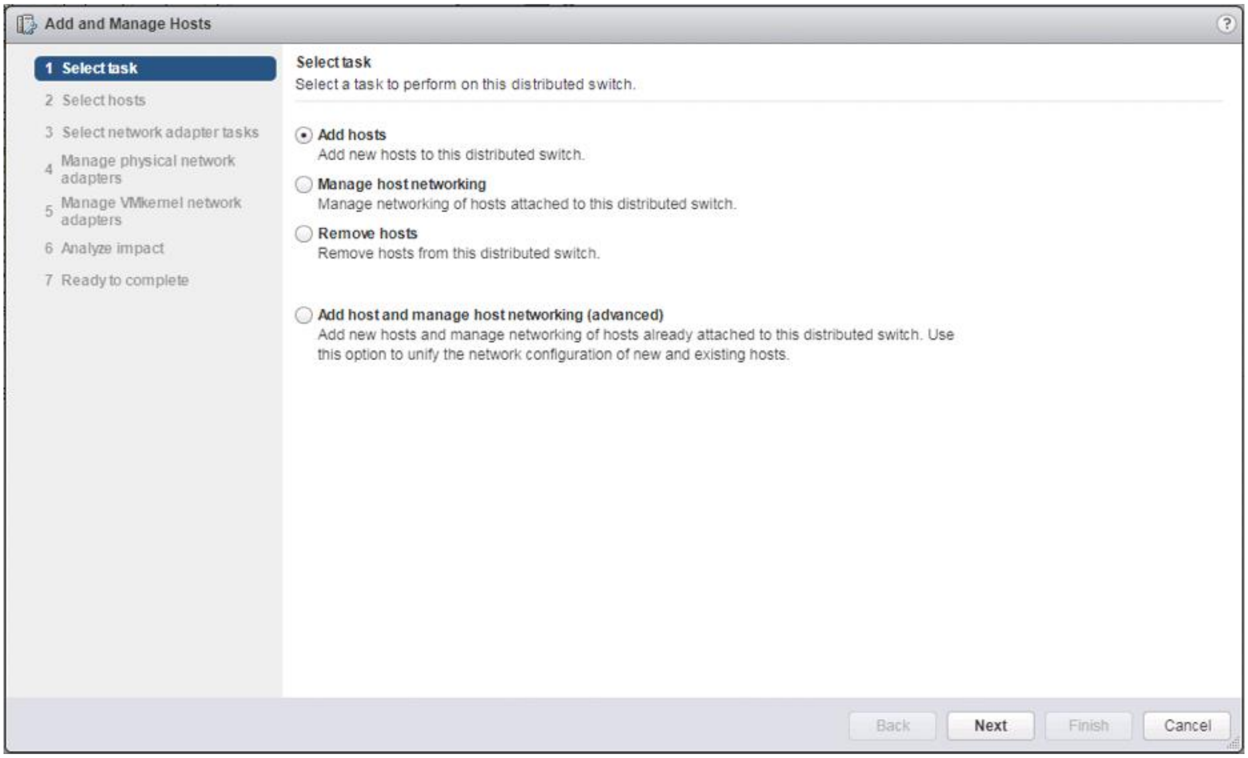

Select task

Select a task to perform on this distributed switch.

● Add hosts
Add new hosts to this distributed switch.

○ Manage host networking
Manage networking of hosts attached to this distributed switch.

○ Remove hosts
Remove hosts from this distributed switch.

○ Add host and manage host networking (advanced)
Add new hosts and manage networking of hosts already attached to this distributed switch. Use this option to unify the network configuration of new and existing hosts.

3. Leave Add hosts selected and click Next.

4. Click the + New hosts… option.

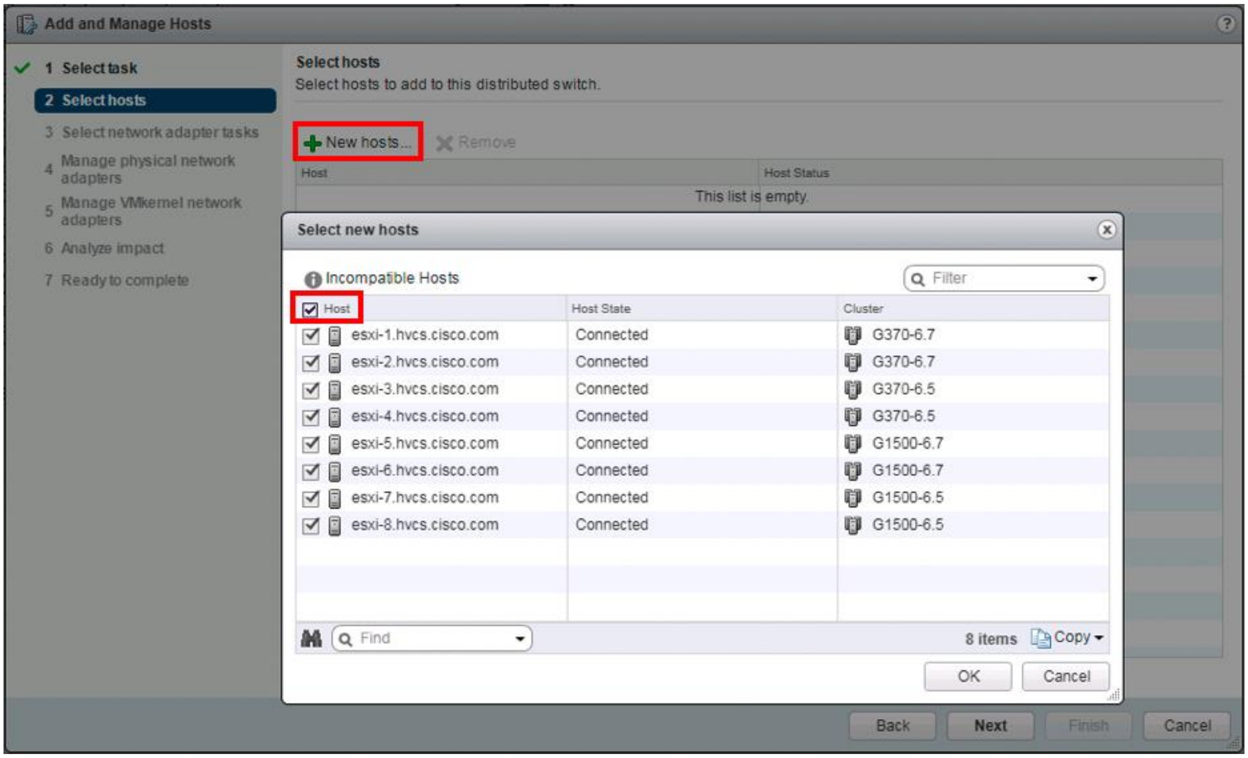

Add and Manage Hosts

✓ 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

Select hosts

Select hosts to add to this distributed switch.

Host	Host Status
This list is empty.	

Select new hosts

Incompatible Hosts

☑ Host	Host State	Cluster
☑ esxi-1.hvcs.cisco.com	Connected	G370-6.7
☑ esxi-2.hvcs.cisco.com	Connected	G370-6.7
☑ esxi-3.hvcs.cisco.com	Connected	G370-6.5
☑ esxi-4.hvcs.cisco.com	Connected	G370-6.5
☑ esxi-5.hvcs.cisco.com	Connected	G1500-6.7
☑ esxi-6.hvcs.cisco.com	Connected	G1500-6.7
☑ esxi-7.hvcs.cisco.com	Connected	G1500-6.5
☑ esxi-8.hvcs.cisco.com	Connected	G1500-6.5

8 items

5. Select the Hosts checkbox near the top of the pop-up window to select all hosts and click OK.

6. Select the Configure identical network settings on multiple hosts (template mode) checkbox.

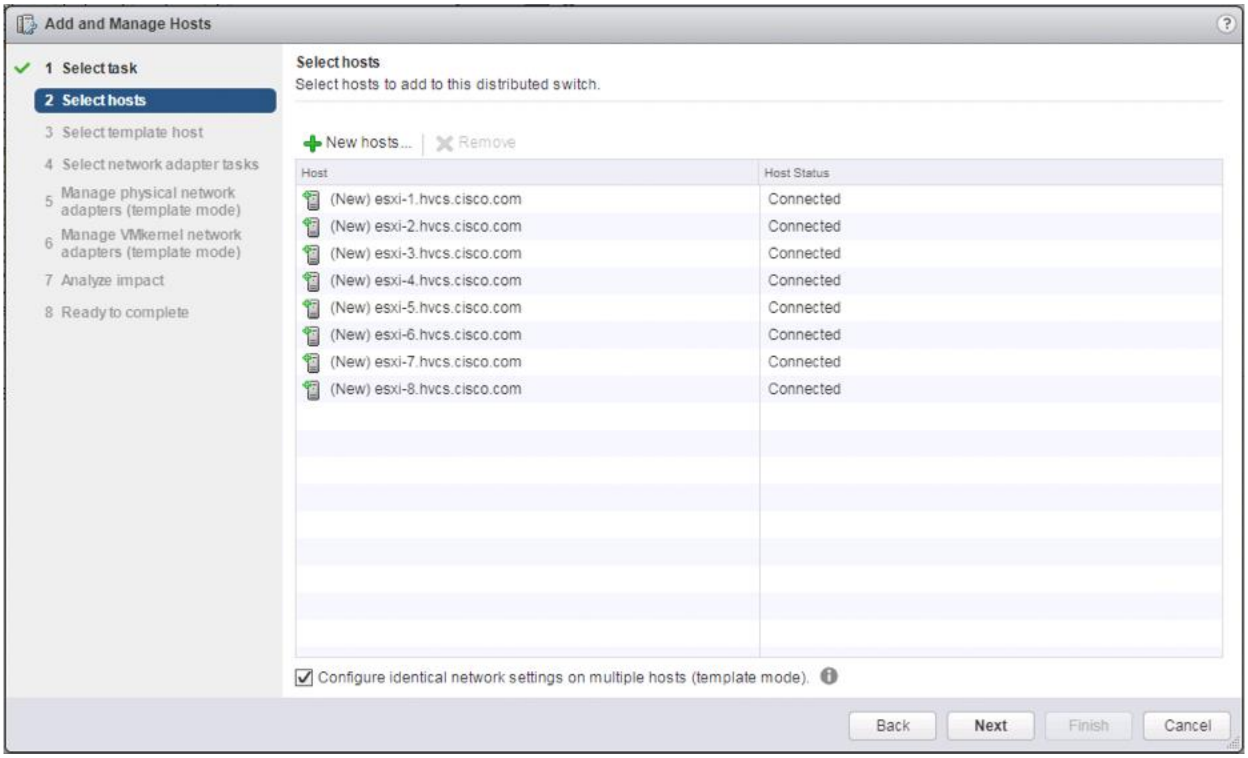

Add and Manage Hosts

✓ 1 Select task
2 Select hosts
3 Select template host
4 Select network adapter tasks
5 Manage physical network adapters (template mode)
6 Manage VMkernel network adapters (template mode)
7 Analyze impact
8 Ready to complete

Select hosts

Select hosts to add to this distributed switch.

Host	Host Status
(New) esxi-1.hvcs.cisco.com	Connected
(New) esxi-2.hvcs.cisco.com	Connected
(New) esxi-3.hvcs.cisco.com	Connected
(New) esxi-4.hvcs.cisco.com	Connected
(New) esxi-5.hvcs.cisco.com	Connected
(New) esxi-6.hvcs.cisco.com	Connected
(New) esxi-7.hvcs.cisco.com	Connected
(New) esxi-8.hvcs.cisco.com	Connected

☑ Configure identical network settings on multiple hosts (template mode).

7. Click Next.

8. Select one of the hosts from the list shown below.

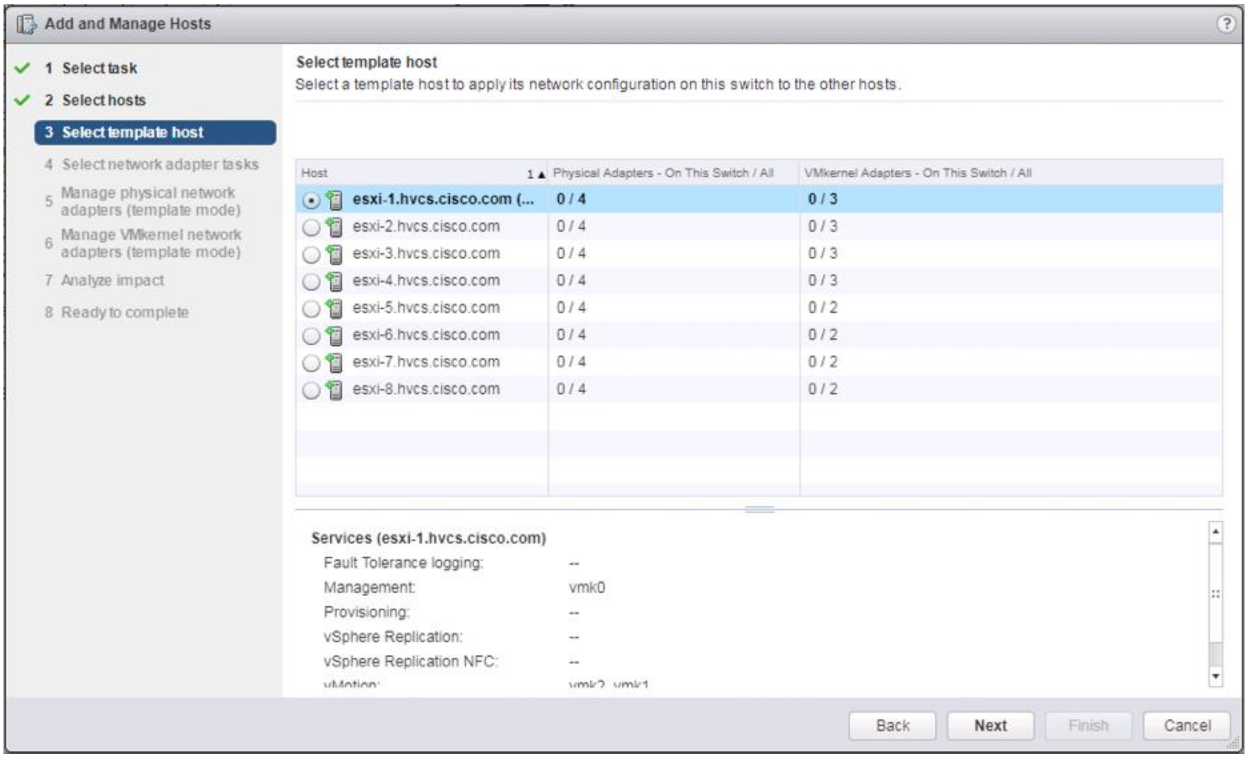
Add and Manage Hosts
✓ 1 Select task
✓ 2 Select hosts
3 Select template host
4 Select network adapter tasks
5 Manage physical network adapters (template mode)
6 Manage VMkernel network adapters (template mode)
7 Analyze impact
8 Ready to complete
Back
Next
Finish
Cancel

Select template host

Select a template host to apply its network configuration on this switch to the other hosts.

Host	Physical Adapters - On This Switch / All	VMkernel Adapters - On This Switch / All
● esxi-1.hvcs.cisco.com (...	0 / 4	0 / 3
○ esxi-2.hvcs.cisco.com	0 / 4	0 / 3
○ esxi-3.hvcs.cisco.com	0 / 4	0 / 3
○ esxi-4.hvcs.cisco.com	0 / 4	0 / 3
○ esxi-5.hvcs.cisco.com	0 / 4	0 / 2
○ esxi-6.hvcs.cisco.com	0 / 4	0 / 2
○ esxi-7.hvcs.cisco.com	0 / 4	0 / 2
○ esxi-8.hvcs.cisco.com	0 / 4	0 / 2

Services (esxi-1.hvcs.cisco.com)

Fault Tolerance logging:	--
Management:	vmk0
Provisioning:	--
vSphere Replication:	--
vSphere Replication NFC:	--
vMotion:	vmk2, vmk1

9. Click Next.

10. Deselect the Manage VMkernel adapters (template mode) option.

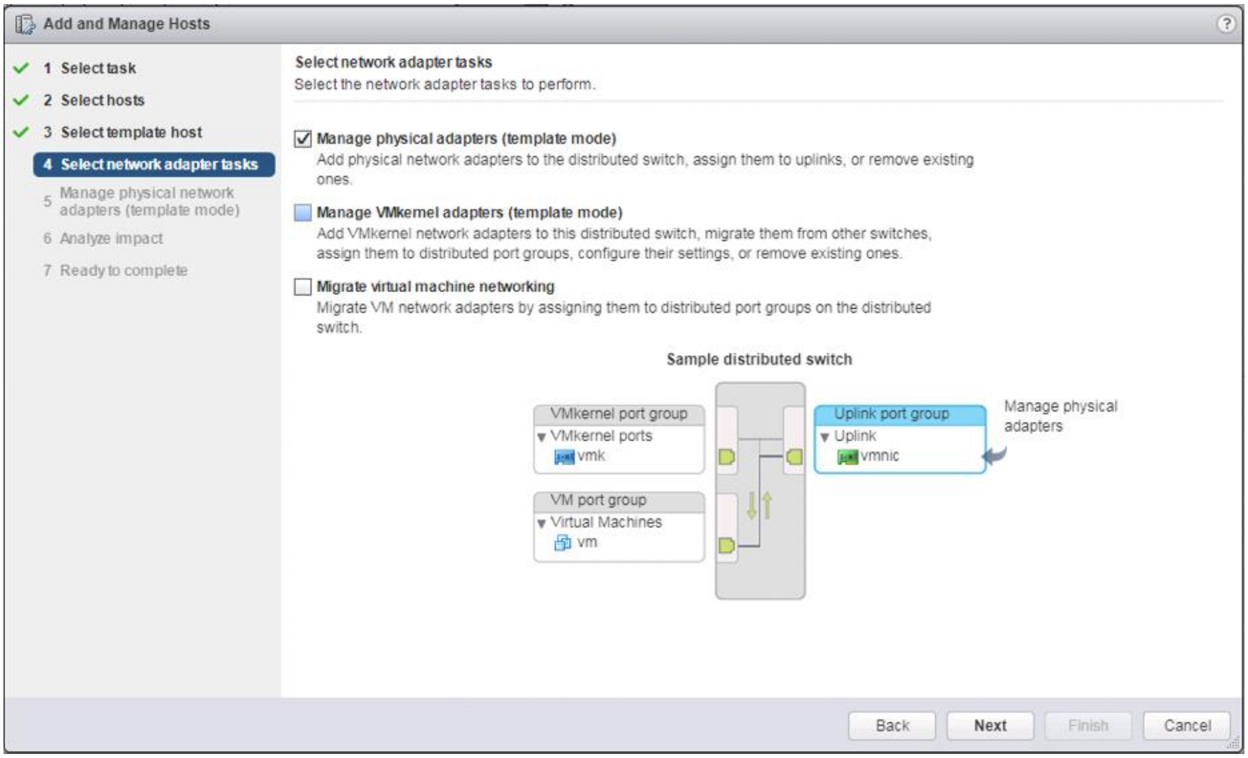

Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- ✓ 3 Select template host
- 4 Select network adapter tasks
- 5 Manage physical network adapters (template mode)
- 6 Analyze impact
- 7 Ready to complete

Select network adapter tasks

Select the network adapter tasks to perform.

☑ **Manage physical adapters (template mode)**
Add physical network adapters to the distributed switch, assign them to uplinks, or remove existing ones.

□ **Manage VMkernel adapters (template mode)**
Add VMkernel network adapters to this distributed switch, migrate them from other switches, assign them to distributed port groups, configure their settings, or remove existing ones.

□ **Migrate virtual machine networking**
Migrate VM network adapters by assigning them to distributed port groups on the distributed switch.

Sample distributed switch

11. Click Next.

12. Select vmnic2.

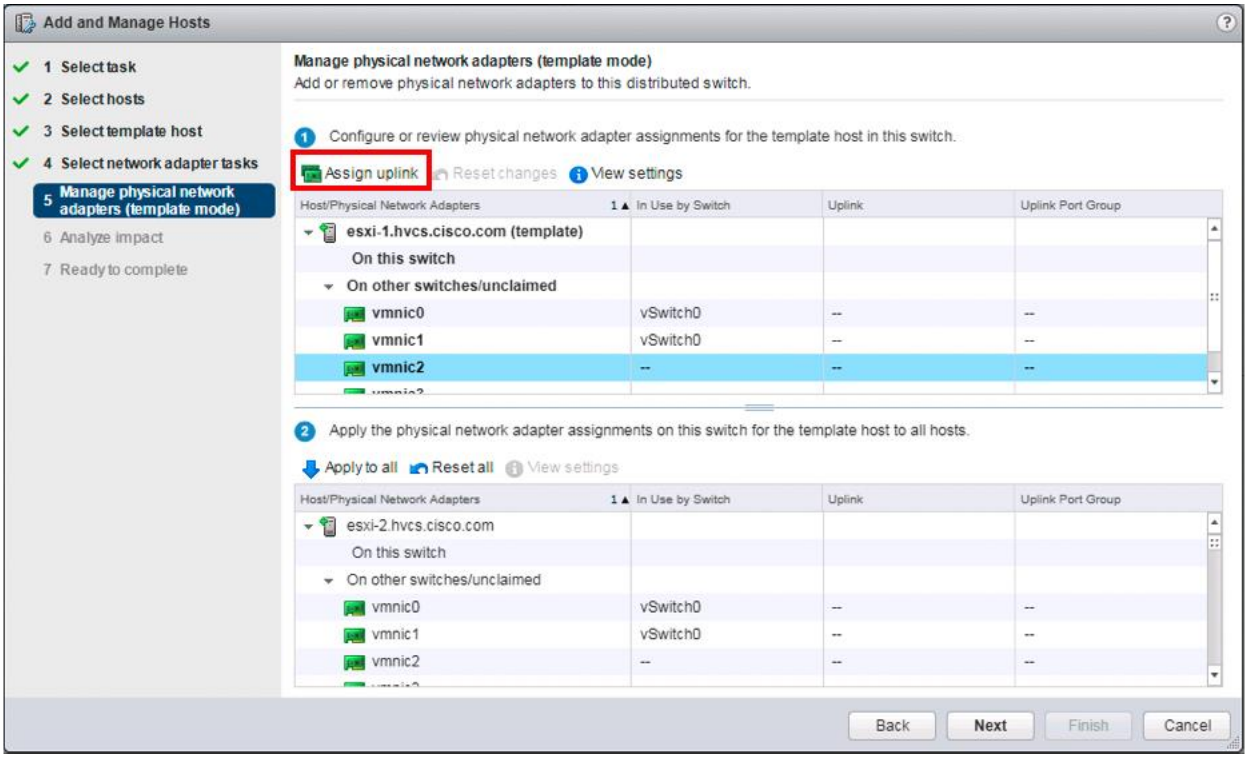

Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- ✓ 3 Select template host
- ✓ 4 Select network adapter tasks
- 5 Manage physical network adapters (template mode)
- 6 Analyze impact
- 7 Ready to complete

Manage physical network adapters (template mode)

Add or remove physical network adapters to this distributed switch.

1 Configure or review physical network adapter assignments for the template host in this switch.

Host/Physical Network Adapters	1 ▲ In Use by Switch	Uplink	Uplink Port Group
esxi-1.hvcs.cisco.com (template)			
On this switch			
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	--	--	--

2 Apply the physical network adapter assignments on this switch for the template host to all hosts.

Host/Physical Network Adapters	1 ▲ In Use by Switch	Uplink	Uplink Port Group
esxi-2.hvcs.cisco.com			
On this switch			
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	--	--	--

13. Click the Assign uplink option.

14. Leave Uplink 1 selected.

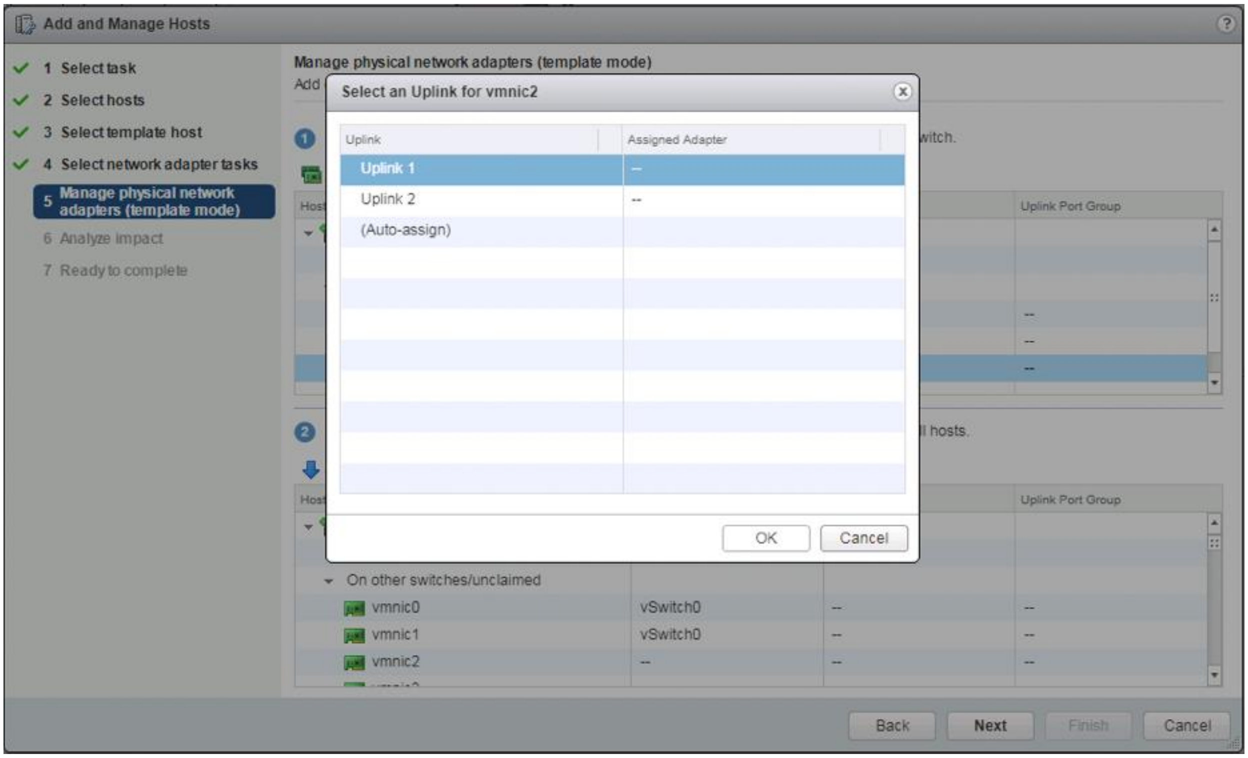

Add and Manage Hosts

1 Select task
2 Select hosts
3 Select template host
4 Select network adapter tasks
5 Manage physical network adapters (template mode)
6 Analyze impact
7 Ready to complete

Manage physical network adapters (template mode)

Select an Uplink for vmnic2

Uplink	Assigned Adapter
Uplink 1	--
Uplink 2	--
(Auto-assign)	

OK Cancel

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	--	--	--

15. Click OK.

16. Select vmnic3 and click Assign uplink to select Uplink 2.

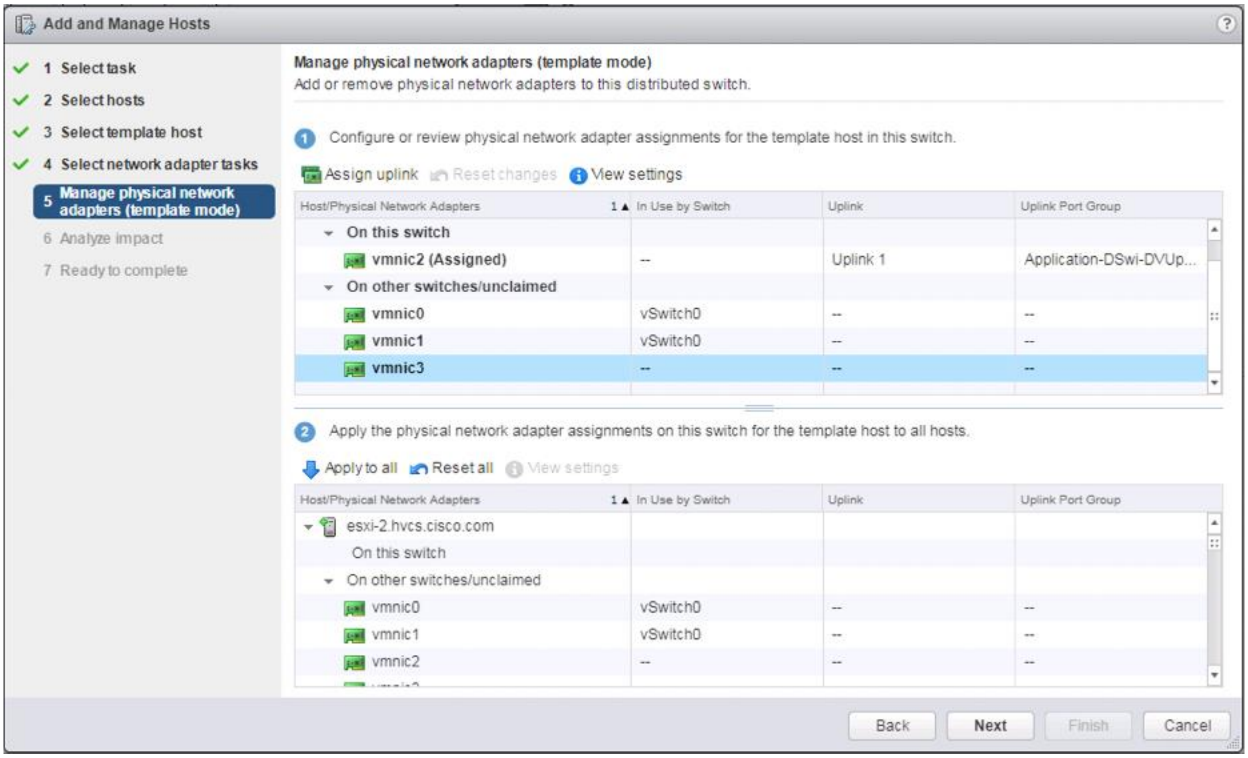

Add and Manage Hosts

1 Select task
2 Select hosts
3 Select template host
4 Select network adapter tasks
5 Manage physical network adapters (template mode)
6 Analyze impact
7 Ready to complete

Manage physical network adapters (template mode)
Add or remove physical network adapters to this distributed switch.

1 Configure or review physical network adapter assignments for the template host in this switch.

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
On this switch			
vmnic2 (Assigned)	--	Uplink 1	Application-DSwi-DVUp...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic3	--	--	--

2 Apply the physical network adapter assignments on this switch for the template host to all hosts.

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
esxi-2.hvcs.cisco.com			
On this switch			
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	--	--	--

17. Click Apply to all.

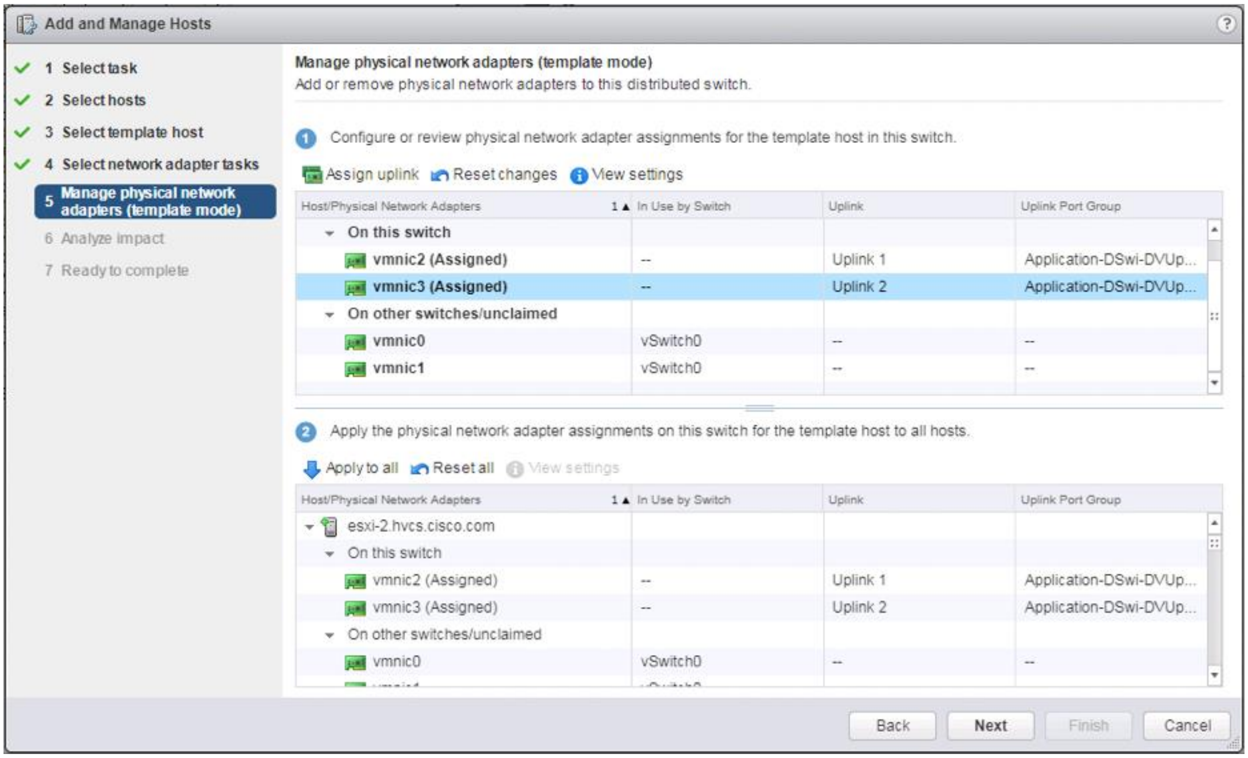

Add and Manage Hosts

Manage physical network adapters (template mode)
Add or remove physical network adapters to this distributed switch.

1 Configure or review physical network adapter assignments for the template host in this switch.

Host/Physical Network Adapters	1 ▲ In Use by Switch	Uplink	Uplink Port Group
On this switch			
vmnic2 (Assigned)	--	Uplink 1	Application-DSwi-DVUp...
vmnic3 (Assigned)	--	Uplink 2	Application-DSwi-DVUp...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--

2 Apply the physical network adapter assignments on this switch for the template host to all hosts.

Host/Physical Network Adapters	1 ▲ In Use by Switch	Uplink	Uplink Port Group
esxi-2.hvcs.cisco.com			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	Application-DSwi-DVUp...
vmnic3 (Assigned)	--	Uplink 2	Application-DSwi-DVUp...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--

18. Click Next.

19. Click Next past Analyze Impact.

20. Verify the summary in the Ready to complete screen.

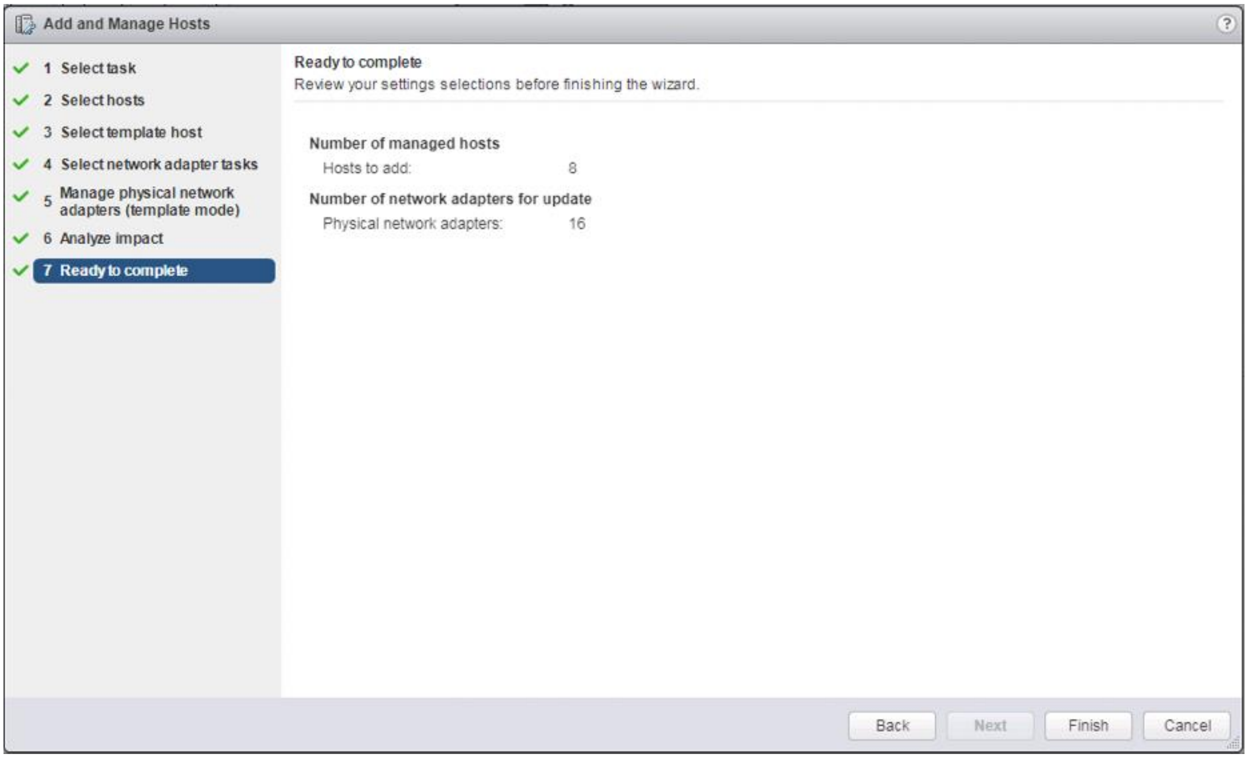

Ready to complete

Review your settings selections before finishing the wizard.

Number of managed hosts

Hosts to add: 8

Number of network adapters for update

Physical network adapters: 16

21. Click Finish to add the hosts.

Add Datastores to Hosts

Datastores have been provisioned and zoned for the associated clusters. To add the datastores to the clusters, follow these steps:

1. From the Hosts tab of the Navigator. Right-click one of the hosts and select Storage -> New Datastore…

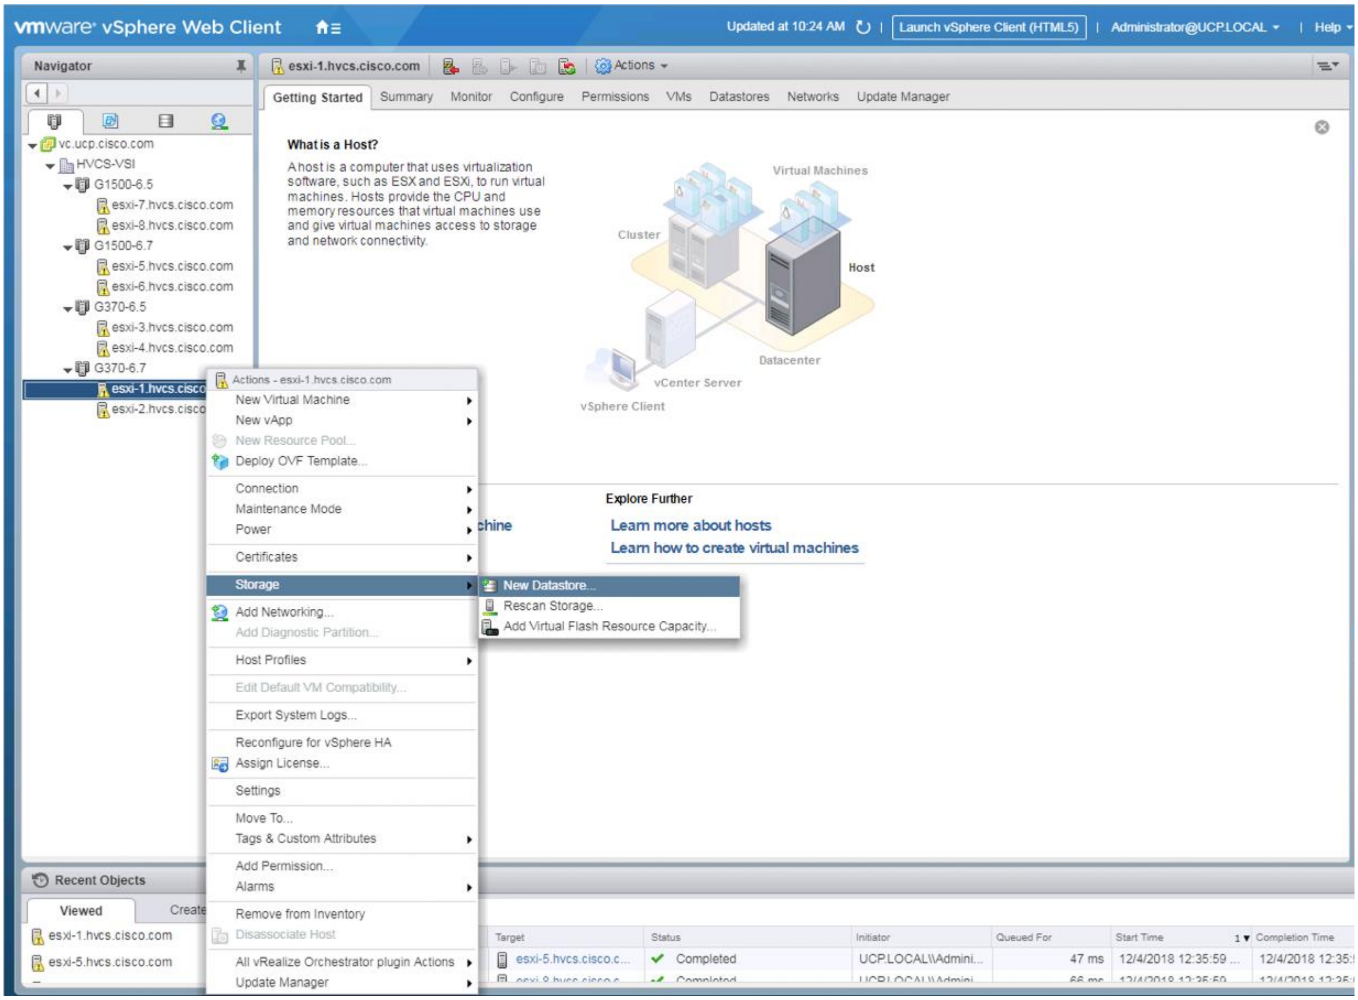

What is a Host?

A host is a computer that uses virtualization software, such as ESX and ESXi, to run virtual machines. Hosts provide the CPU and memory resources that virtual machines use and give virtual machines access to storage and network connectivity.

Target	Status	Initiator	Queued For	Start Time	Completion Time
esxi-5.hvcs.cisco.c...	✓ Completed	UCP.LOCAL\\Admini...	47 ms	12/4/2018 12:35:59 ...	12/4/2018 12:35:

2. Leave VMFS selected and click Next.

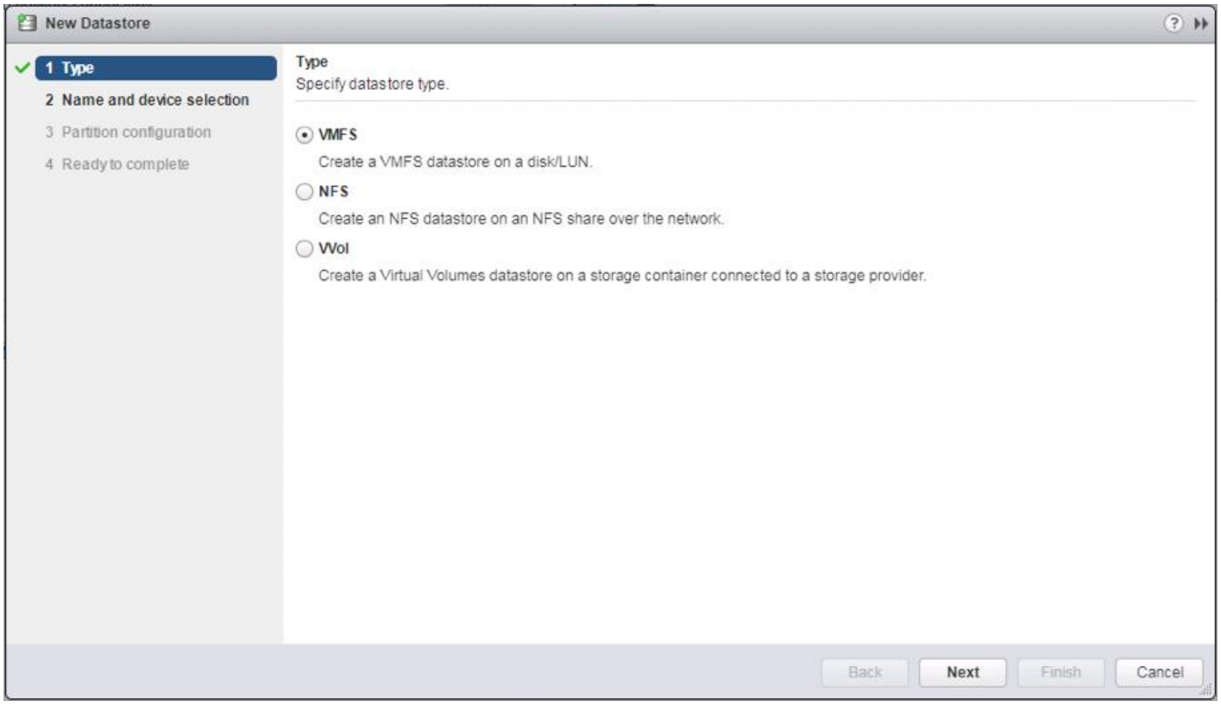

New Datastore

- ✓ 1 Type
- 2 Name and device selection
- 3 Partition configuration
- 4 Ready to complete

Type
Specify datastore type.

● VMFS
Create a VMFS datastore on a disk/LUN.

○ NFS
Create an NFS datastore on an NFS share over the network.

○ VVol
Create a Virtual Volumes datastore on a storage container connected to a storage provider.

3. Provide an appropriate Datastore name and select the appropriate LUN.

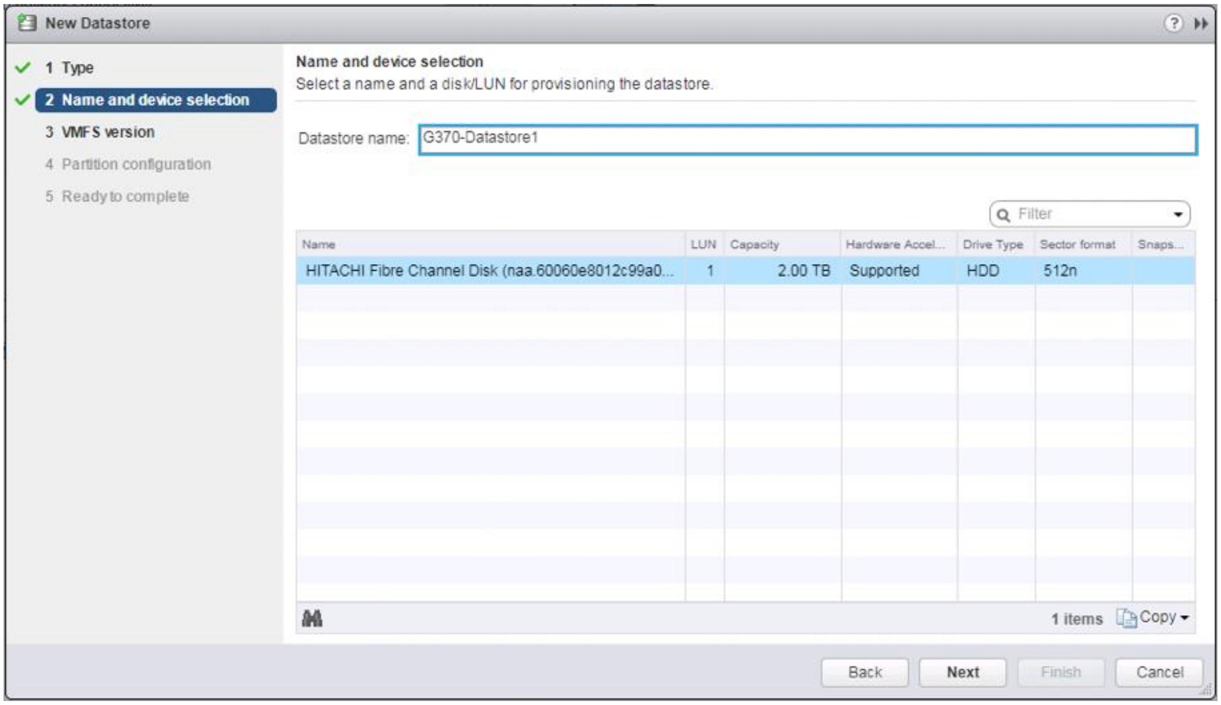

New Datastore

- ✓ 1 Type
- ✓ 2 Name and device selection
- 3 VMFS version
- 4 Partition configuration
- 5 Ready to complete

Name and device selection
Select a name and a disk/LUN for provisioning the datastore.

Datastore name: G370-Datastore1

Name	LUN	Capacity	Hardware Accel...	Drive Type	Sector format	Snaps...
HITACHI Fibre Channel Disk (naa.60060e8012c99a0...	1	2.00 TB	Supported	HDD	512n	

1 items

4. Click Next.

5. Leave VMFS 6 selected.

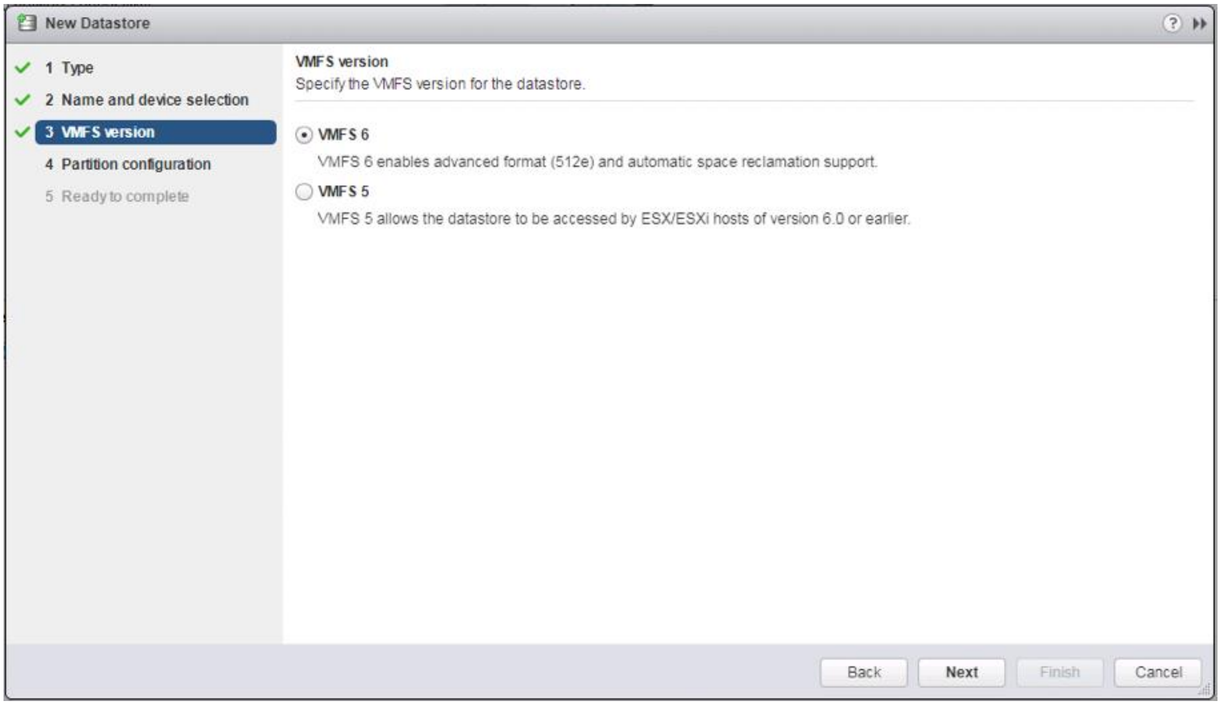

VMFS version

Specify the VMFS version for the datastore.

● VMFS 6

VMFS 6 enables advanced format (512e) and automatic space reclamation support.

○ VMFS 5

VMFS 5 allows the datastore to be accessed by ESX/ESXi hosts of version 6.0 or earlier.

6. Click Next.

7. Leave the defaults for Partition configuration.

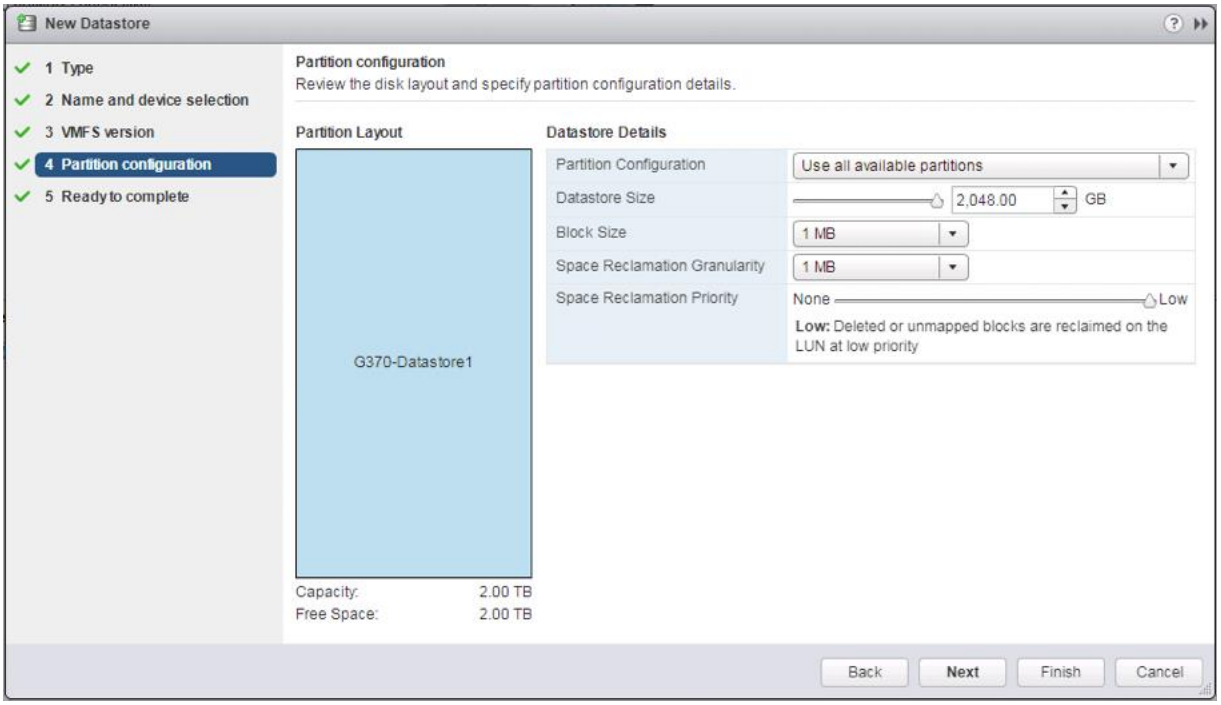

Partition configuration

Review the disk layout and specify partition configuration details.

Partition Layout

G370-Datastore1

Capacity:	2.00 TB
Free Space:	2.00 TB

Datastore Details

Partition Configuration	Use all available partitions
Datastore Size	2,048.00 GB
Block Size	1 MB
Space Reclamation Granularity	1 MB
Space Reclamation Priority	None — Low **Low:** Deleted or unmapped blocks are reclaimed on the LUN at low priority

8. Click Next.

9. Review the settings.

10. Click Finish to create.

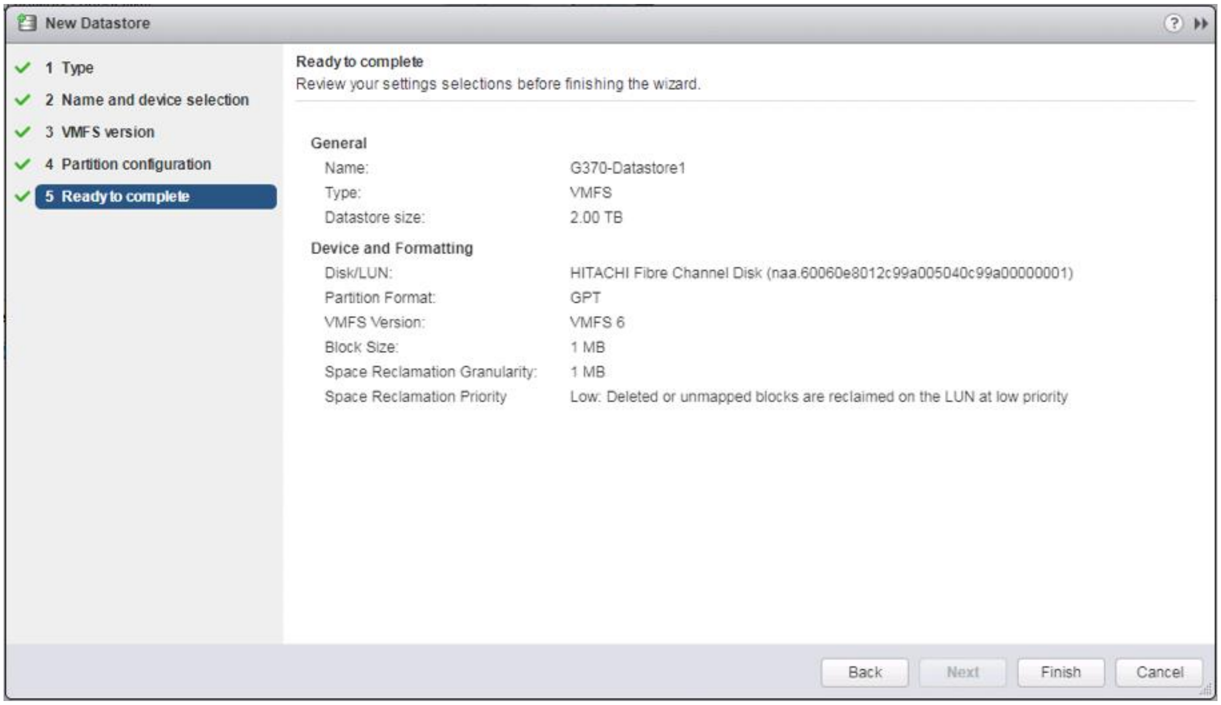

Ready to complete

Review your settings selections before finishing the wizard.

General

Name:	G370-Datastore1
Type:	VMFS
Datastore size:	2.00 TB

Device and Formatting

Disk/LUN:	HITACHI Fibre Channel Disk (naa.60060e8012c99a005040c99a00000001)
Partition Format:	GPT
VMFS Version:	VMFS 6
Block Size:	1 MB
Space Reclamation Granularity:	1 MB
Space Reclamation Priority	Low: Deleted or unmapped blocks are reclaimed on the LUN at low priority

11. Check each host in the clusters associated to the same VSP the datastore was provisioned from; all should show the configured datastore as available:

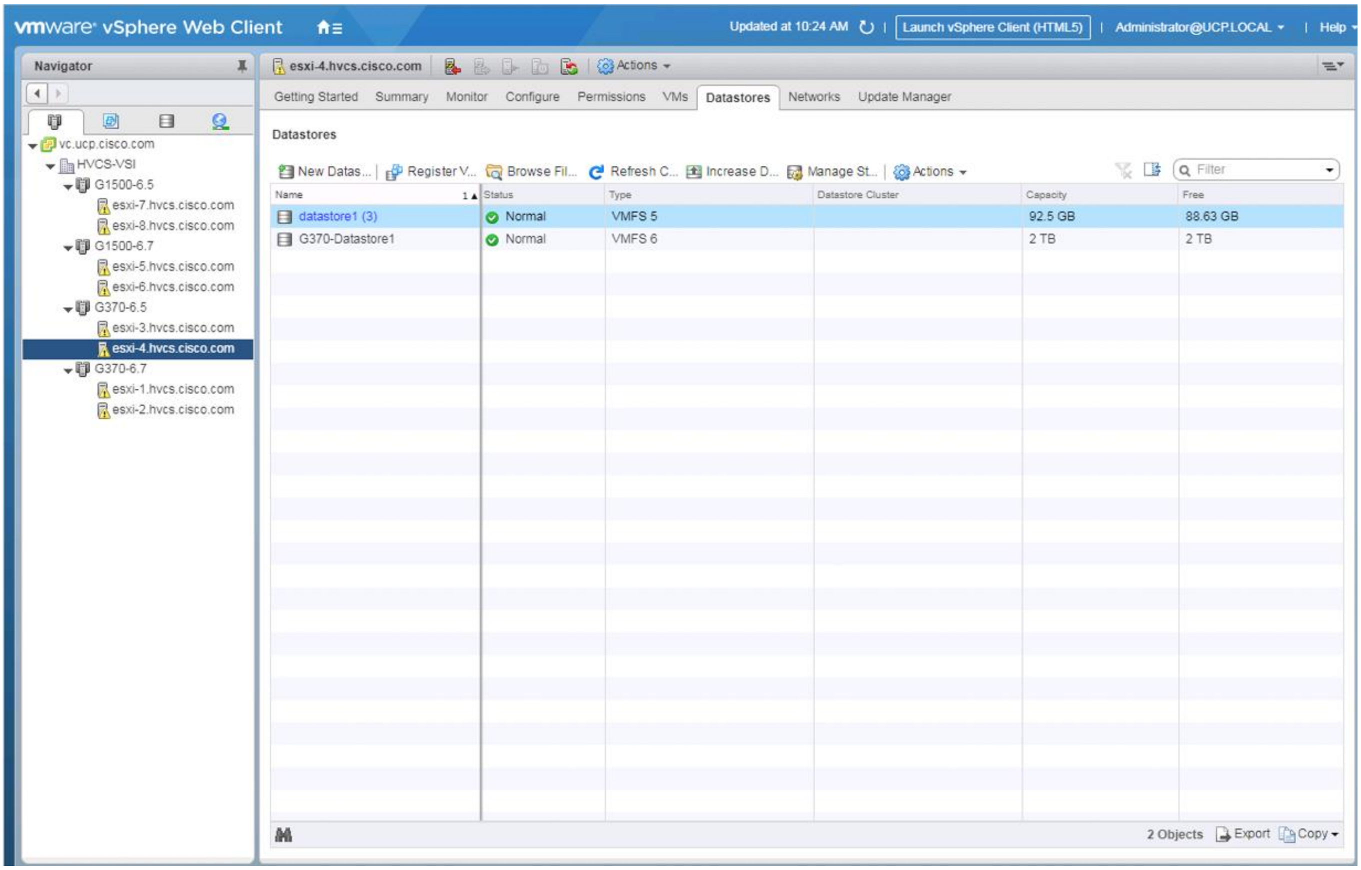

Name	Status	Type	Datastore Cluster	Capacity	Free
datastore1 (3)	Normal	VMFS 5		92.5 GB	88.63 GB
G370-Datastore1	Normal	VMFS 6		2 TB	2 TB

12. If multiple VMFS datastore LUNs were deployed, repeat these steps on a host they are associated to by a Hitachi Host Group.

Configure NTP on ESXi Hosts

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

1. From the Configure tab, select the Time Configuration section under System.

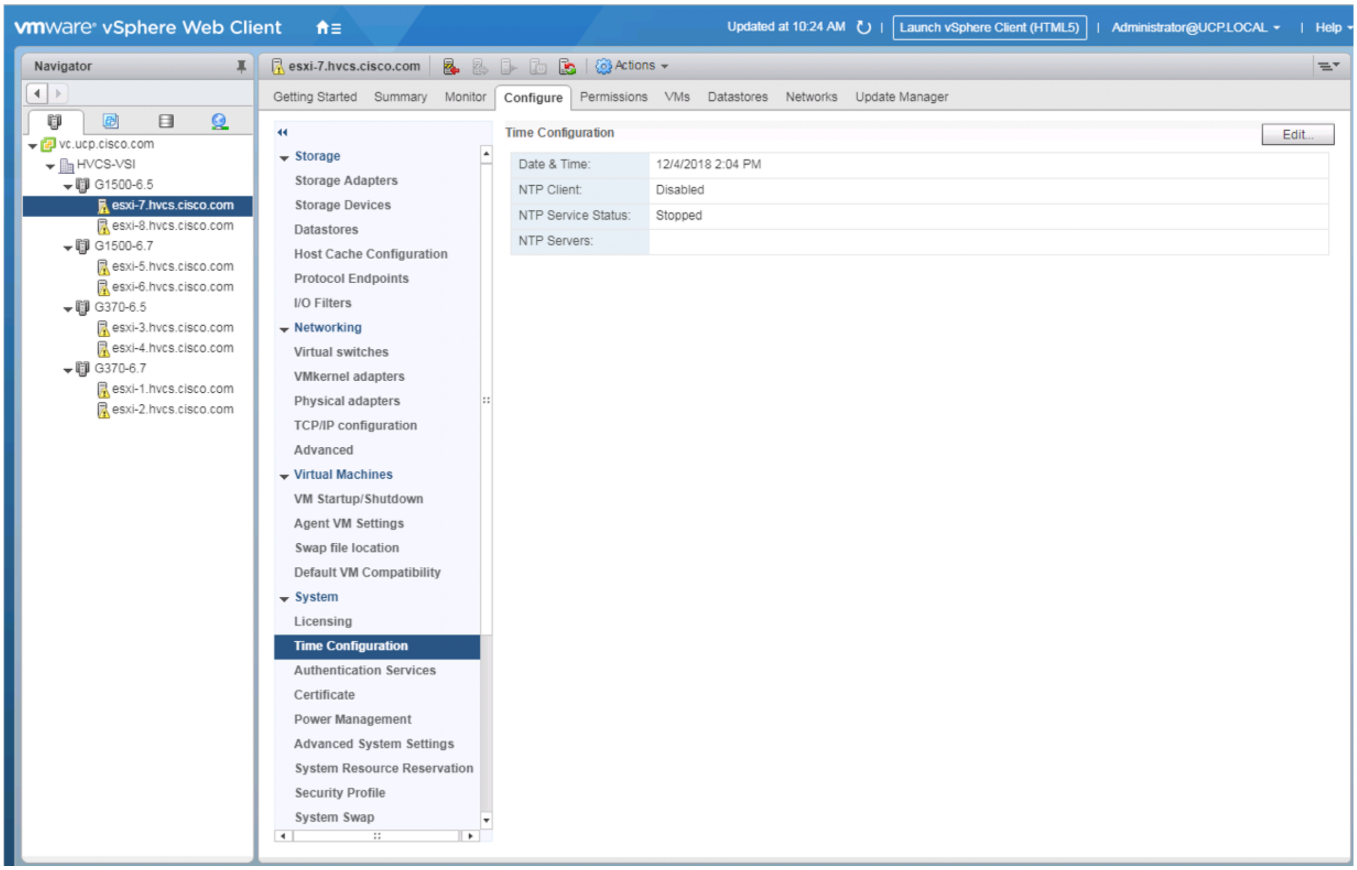

Time Configuration

Date & Time:	12/4/2018 2:04 PM
NTP Client:	Disabled
NTP Service Status:	Stopped
NTP Servers:	

2. Click Edit.

3. Select the Use Network Time Protocol (Enable NTP client) option.

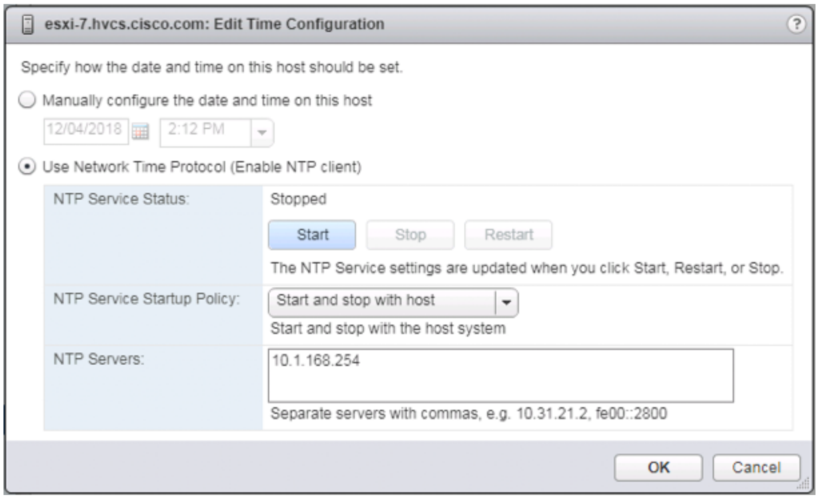

esxi-7.hvcs.cisco.com: Edit Time Configuration

Specify how the date and time on this host should be set.

○ Manually configure the date and time on this host

12/04/2018 2:12 PM

● Use Network Time Protocol (Enable NTP client)

NTP Service Status: Stopped

The NTP Service settings are updated when you click Start, Restart, or Stop.

NTP Service Startup Policy: Start and stop with host

Start and stop with the host system

NTP Servers: 10.1.168.254

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

4. Enter an appropriate NTP server within the NTP Servers box, change NTP Service Startup Policy to Start and stop with host, and click Start.

5. Verify that NTP service is now running and the clock is now set to approximately the correct time.

The NTP server time sync make take a few minutes.

Create and Apply Patch Baselines with VUM

Critical patches are automatically available within VMware Update Manager (VUM) when using current versions of vCenter Server. A Patch Baseline will be made for the deployed vSphere release(s) and applied to each host to install appropriate patches.

To create the baselines and patch the new ESXi hosts, follow these steps:

1. From the Hosts tab select the vCenter and go to the Update Manager tab.

2. Click Go to Admin View.

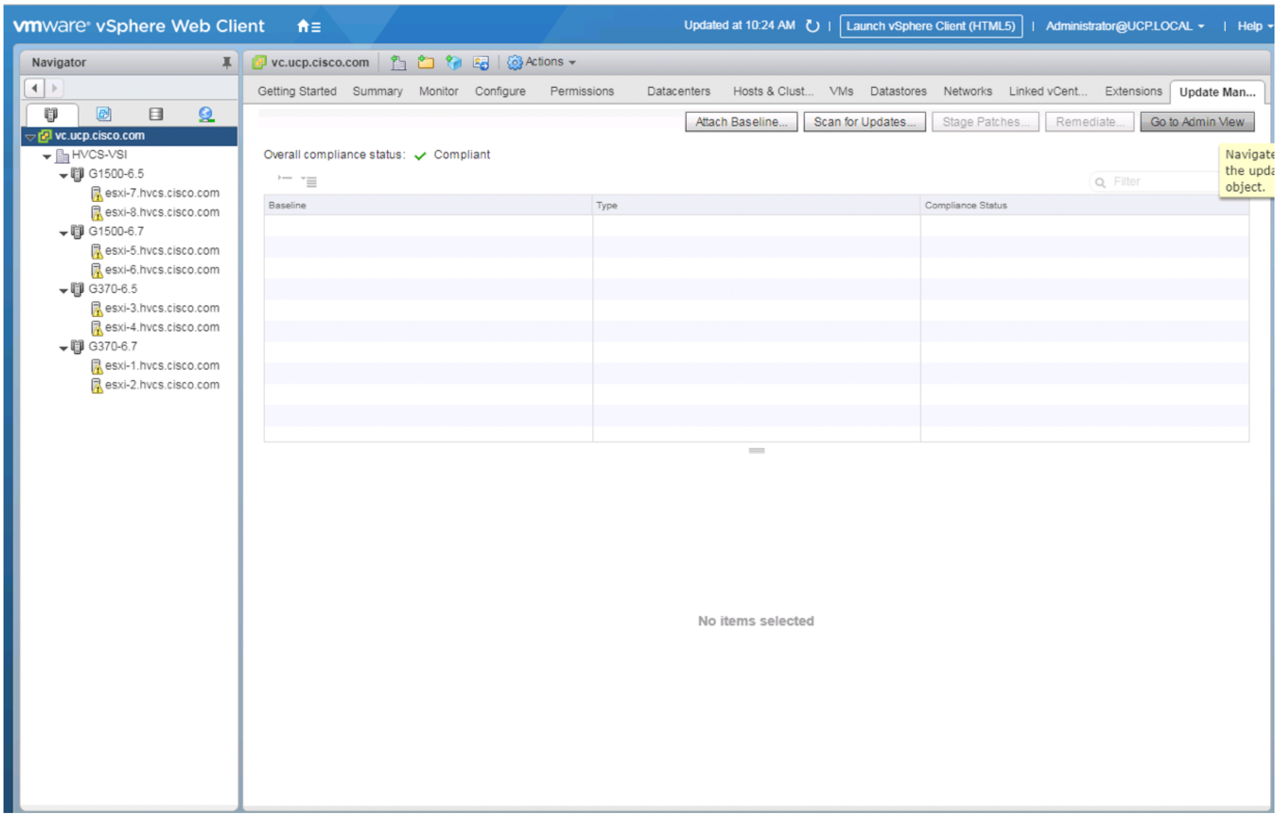

Overall compliance status: ✓ Compliant

No items selected

3. From the Manage tab select Hosts Baselines.

4. Click +New Baseline...

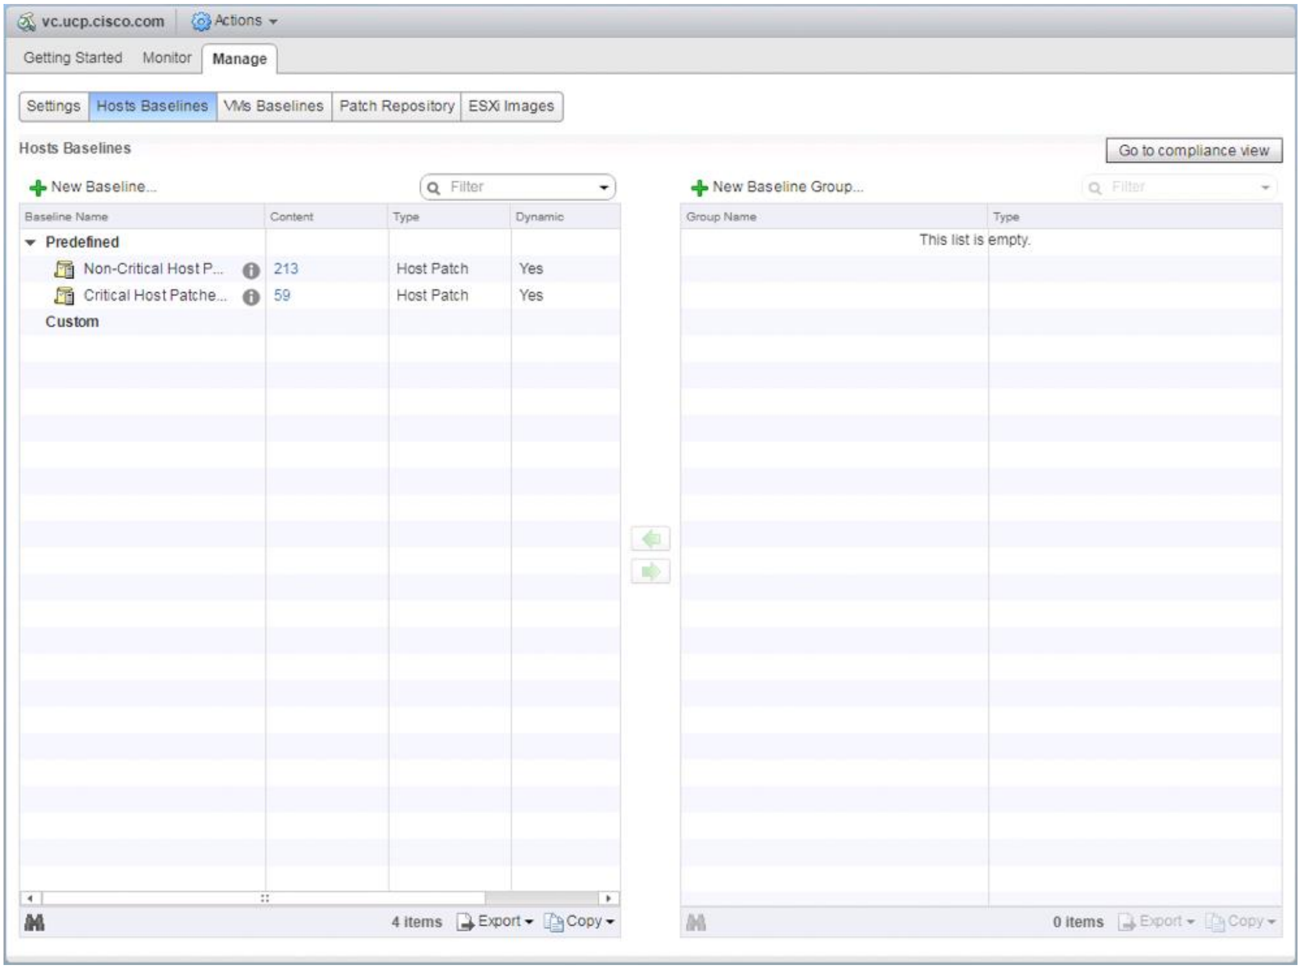

Baseline Name	Content	Type	Dynamic
Predefined			
Non-Critical Host P...	213	Host Patch	Yes
Critical Host Patche...	59	Host Patch	Yes
Custom			

Group Name	Type
This list is empty.	

5. Provide a name for the Baseline, leave the Baseline type selected as Host Patch, and click Next.

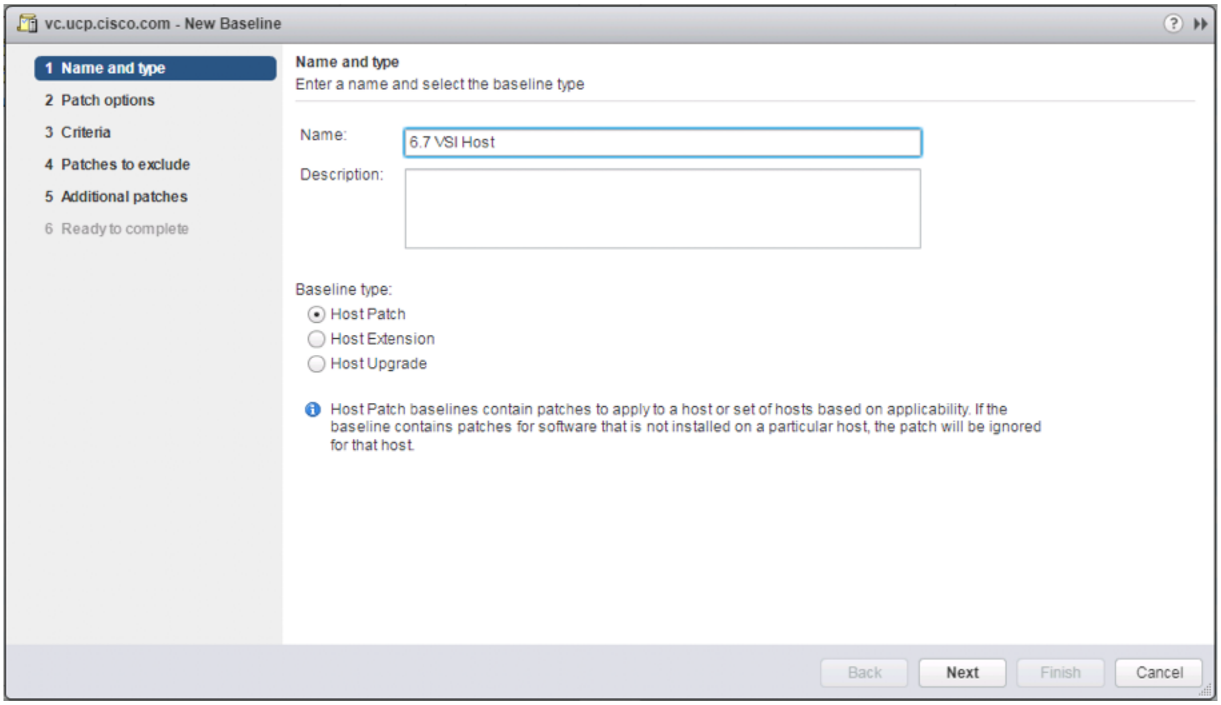

vc.ucp.cisco.com - New Baseline

Name and type
Enter a name and select the baseline type

Name: 6.7 VSI Host

Description:

Baseline type:
● Host Patch
○ Host Extension
○ Host Upgrade

Host Patch baselines contain patches to apply to a host or set of hosts based on applicability. If the baseline contains patches for software that is not installed on a particular host, the patch will be ignored for that host.

6. Leave the Patch options set as Dynamic and click Next.

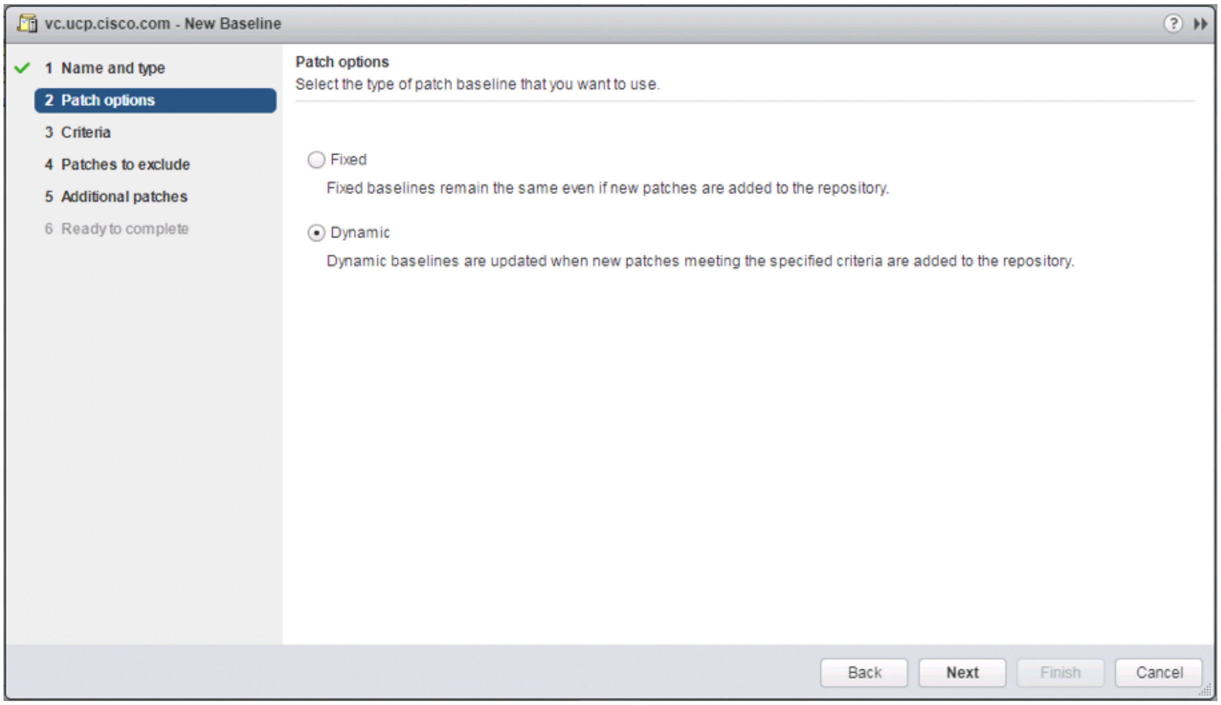

vc.ucp.cisco.com - New Baseline

Patch options
Select the type of patch baseline that you want to use.

○ Fixed
Fixed baselines remain the same even if new patches are added to the repository.

● Dynamic
Dynamic baselines are updated when new patches meeting the specified criteria are added to the repository.

7. Under Product, select the target vSphere release, and under Severity select Critical.

8. Click Next.

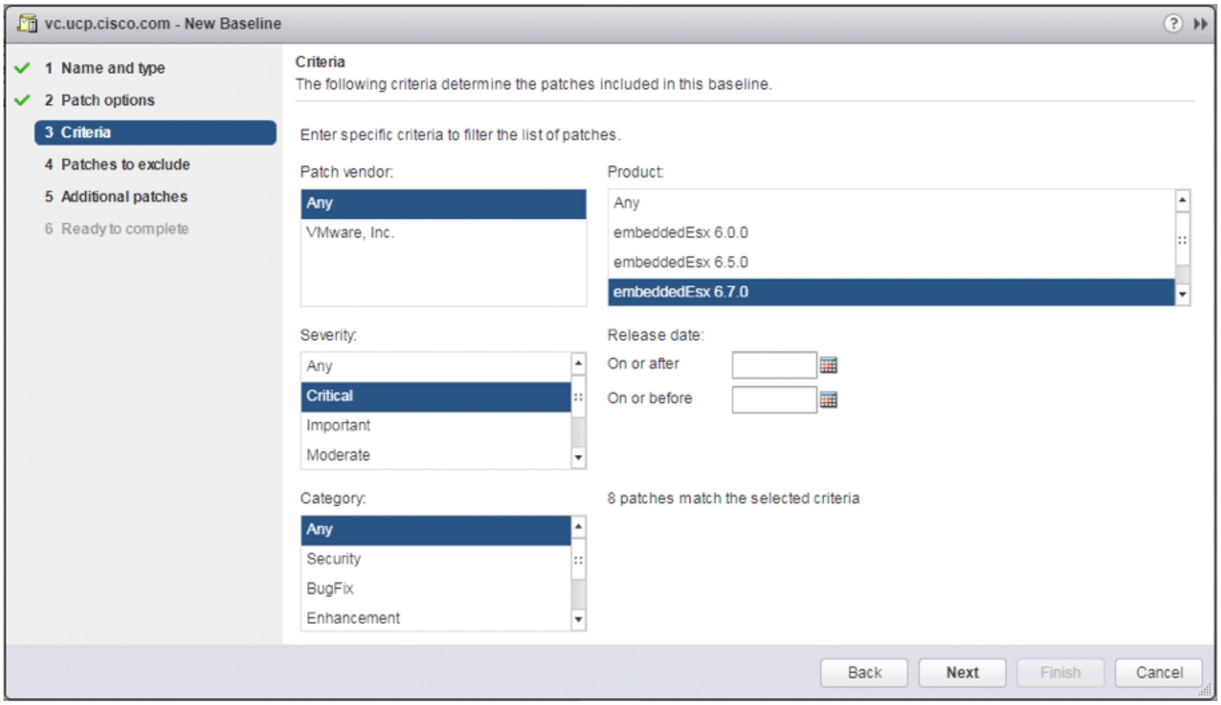

vc.ucp.cisco.com - New Baseline

- ✓ 1 Name and type
- ✓ 2 Patch options
- 3 Criteria
- 4 Patches to exclude
- 5 Additional patches
- 6 Ready to complete

Criteria

The following criteria determine the patches included in this baseline.

Enter specific criteria to filter the list of patches.

Patch vendor: **Any** (selected), VMware, Inc.

Product: Any, embeddedEsx 6.0.0, embeddedEsx 6.5.0, **embeddedEsx 6.7.0** (selected)

Severity: Any, **Critical** (selected), Important, Moderate

Release date: On or after ☐ On or before ☐

Category: **Any** (selected), Security, BugFix, Enhancement

8 patches match the selected criteria

9. Exclude any patches if appropriate and click Next.

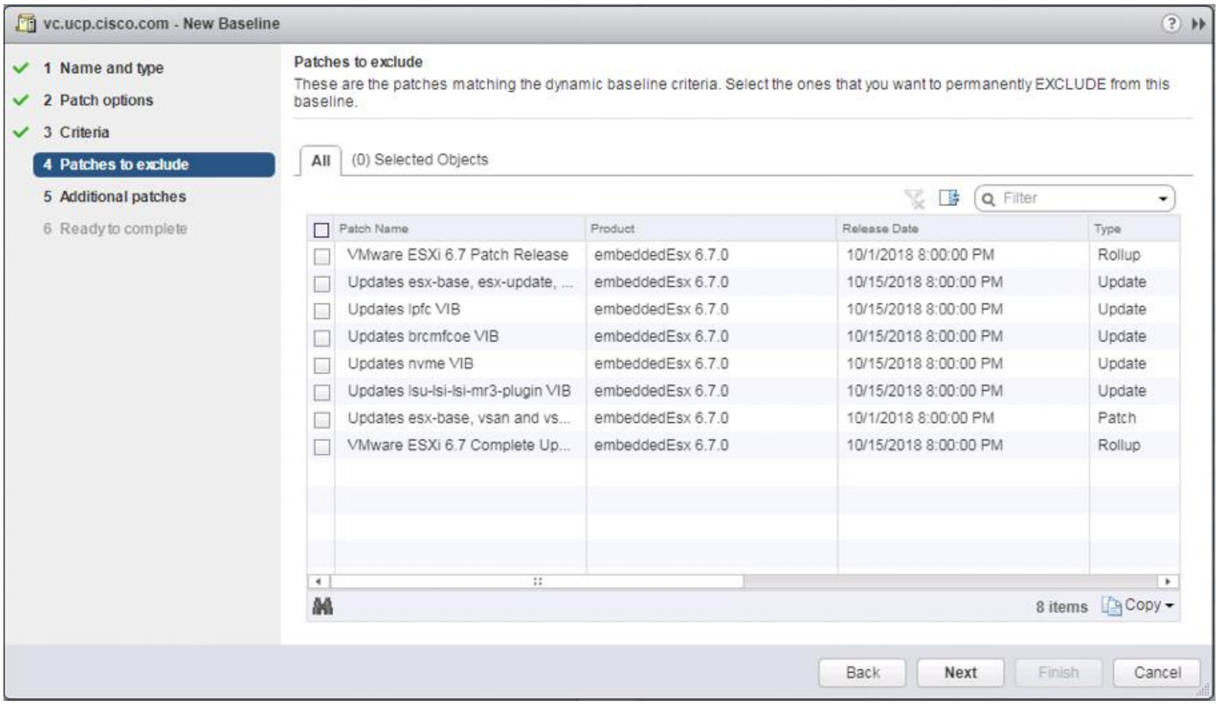

vc.ucp.cisco.com - New Baseline

- ✓ 1 Name and type
- ✓ 2 Patch options
- ✓ 3 Criteria
- 4 Patches to exclude
- 5 Additional patches
- 6 Ready to complete

Patches to exclude

These are the patches matching the dynamic baseline criteria. Select the ones that you want to permanently EXCLUDE from this baseline.

All | (0) Selected Objects

□	Patch Name	Product	Release Date	Type
□	VMware ESXi 6.7 Patch Release	embeddedEsx 6.7.0	10/1/2018 8:00:00 PM	Rollup
□	Updates esx-base, esx-update, ...	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Update
□	Updates lpfc VIB	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Update
□	Updates brcmfcoe VIB	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Update
□	Updates nvme VIB	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Update
□	Updates lsu-lsi-lsi-mr3-plugin VIB	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Update
□	Updates esx-base, vsan and vs...	embeddedEsx 6.7.0	10/1/2018 8:00:00 PM	Patch
□	VMware ESXi 6.7 Complete Up...	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Rollup

8 items

10. Select any additional patches if appropriate and click Next.

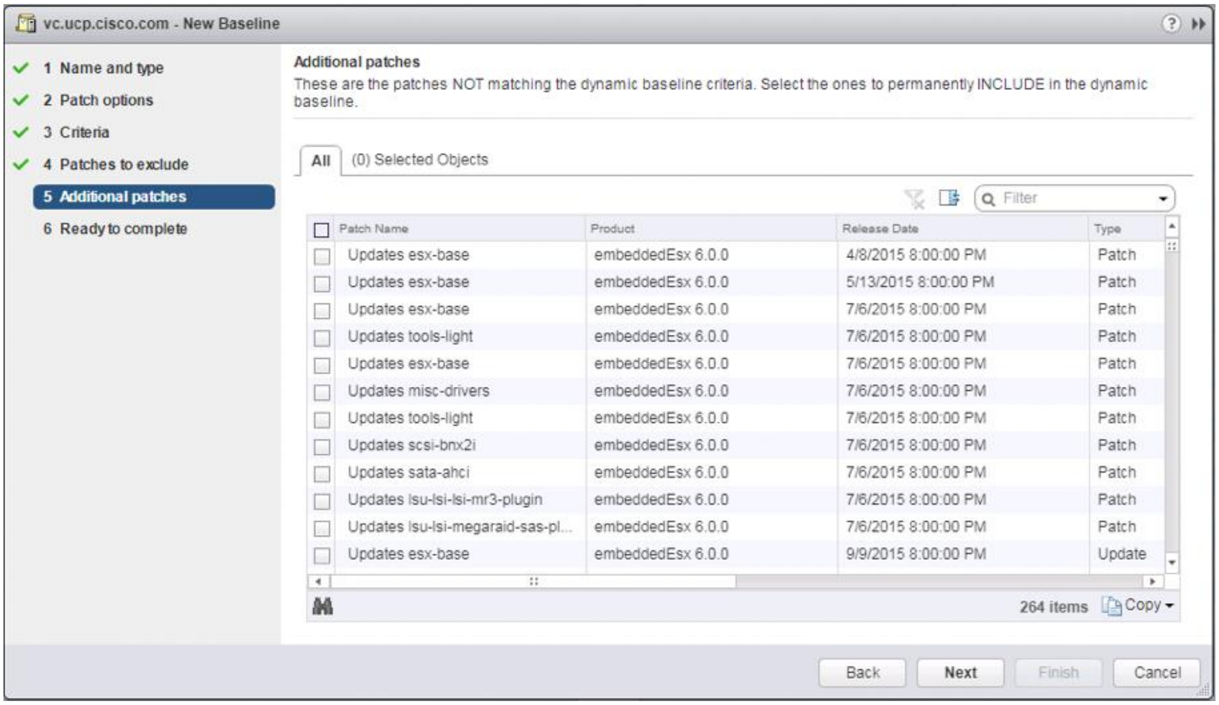

vc.ucp.cisco.com - New Baseline

✓ 1 Name and type
✓ 2 Patch options
✓ 3 Criteria
✓ 4 Patches to exclude
5 Additional patches
6 Ready to complete

Additional patches

These are the patches NOT matching the dynamic baseline criteria. Select the ones to permanently INCLUDE in the dynamic baseline.

All | (0) Selected Objects

□	Patch Name	Product	Release Date	Type
□	Updates esx-base	embeddedEsx 6.0.0	4/8/2015 8:00:00 PM	Patch
□	Updates esx-base	embeddedEsx 6.0.0	5/13/2015 8:00:00 PM	Patch
□	Updates esx-base	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates tools-light	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates esx-base	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates misc-drivers	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates tools-light	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates scsi-bnx2i	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates sata-ahci	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates lsu-lsi-lsi-mr3-plugin	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates lsu-lsi-megaraid-sas-pl...	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch
□	Updates esx-base	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update

264 items

11. Review the selections and click Finish.

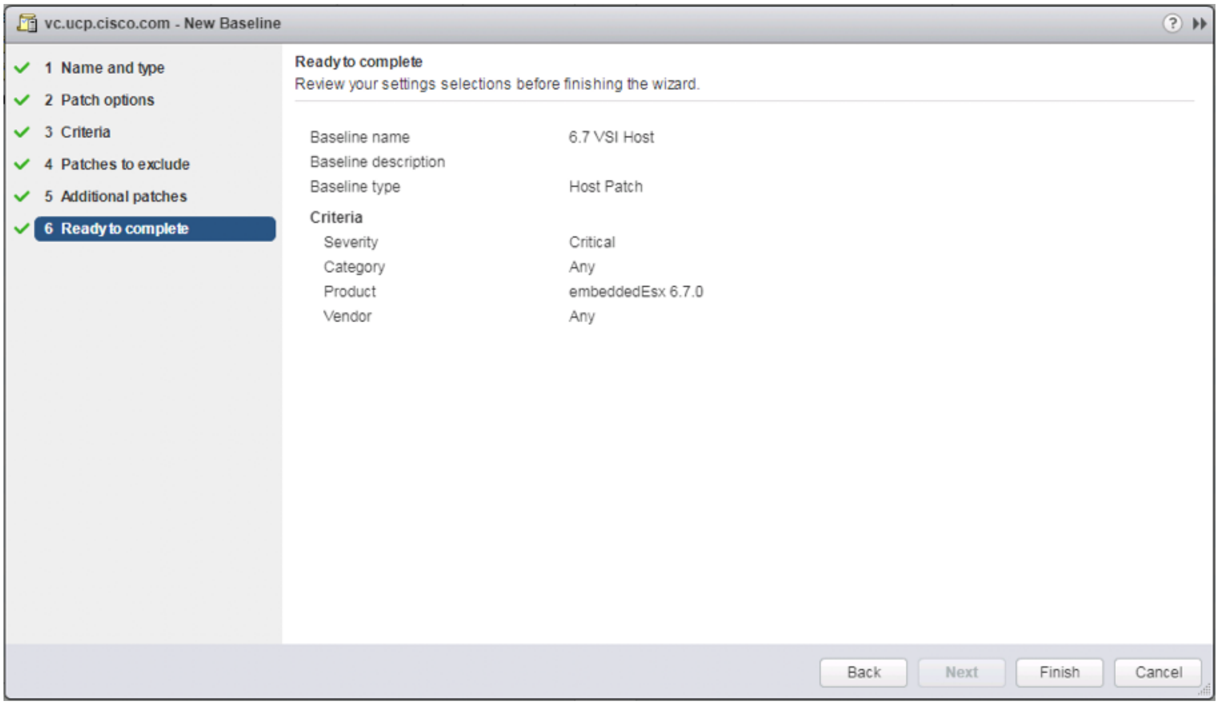

vc.ucp.cisco.com - New Baseline

✓ 1 Name and type
✓ 2 Patch options
✓ 3 Criteria
✓ 4 Patches to exclude
✓ 5 Additional patches
✓ 6 Ready to complete

Ready to complete

Review your settings selections before finishing the wizard.

Baseline name: 6.7 VSI Host
Baseline description:
Baseline type: Host Patch

Criteria
Severity: Critical
Category: Any
Product: embeddedEsx 6.7.0
Vendor: Any

12. Go back to the Hosts view within Navigator.

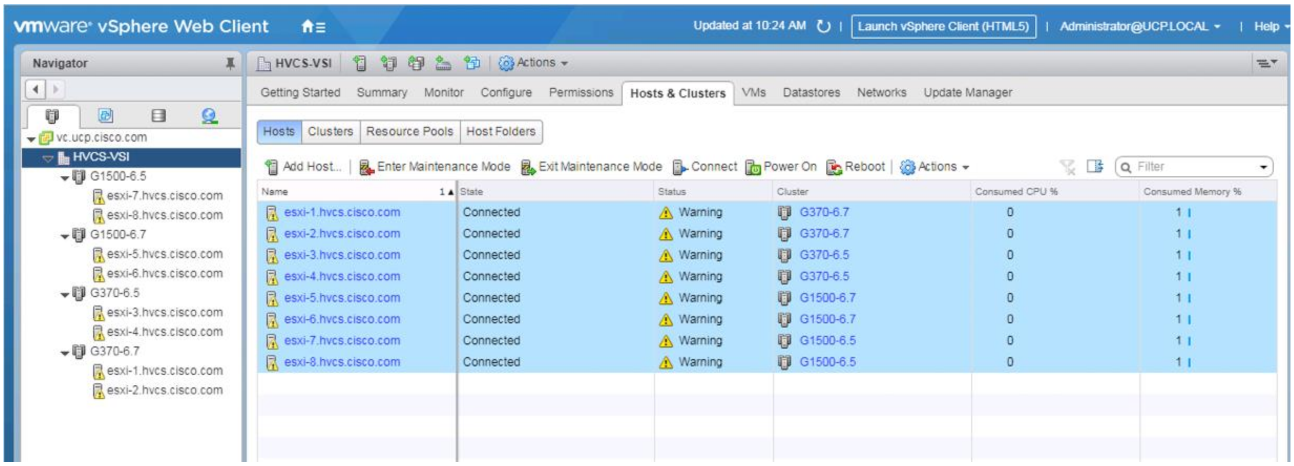

Name	State	Status	Cluster	Consumed CPU %	Consumed Memory %
esxi-1.hvcs.cisco.com	Connected	Warning	G370-6.7	0	1
esxi-2.hvcs.cisco.com	Connected	Warning	G370-6.7	0	1
esxi-3.hvcs.cisco.com	Connected	Warning	G370-6.5	0	1
esxi-4.hvcs.cisco.com	Connected	Warning	G370-6.5	0	1
esxi-5.hvcs.cisco.com	Connected	Warning	G1500-6.7	0	1
esxi-6.hvcs.cisco.com	Connected	Warning	G1500-6.7	0	1
esxi-7.hvcs.cisco.com	Connected	Warning	G1500-6.5	0	1
esxi-8.hvcs.cisco.com	Connected	Warning	G1500-6.5	0	1

13. Select the Datacenter level of HVCS-VSI and within Hosts of the Hosts and Clusters tab select all hosts and click Enter Maintenance Mode.

14. For the first host associated with the vSphere release of the baseline, select the host and the Update Manager tab for that host.

15. Click Attach Baseline….

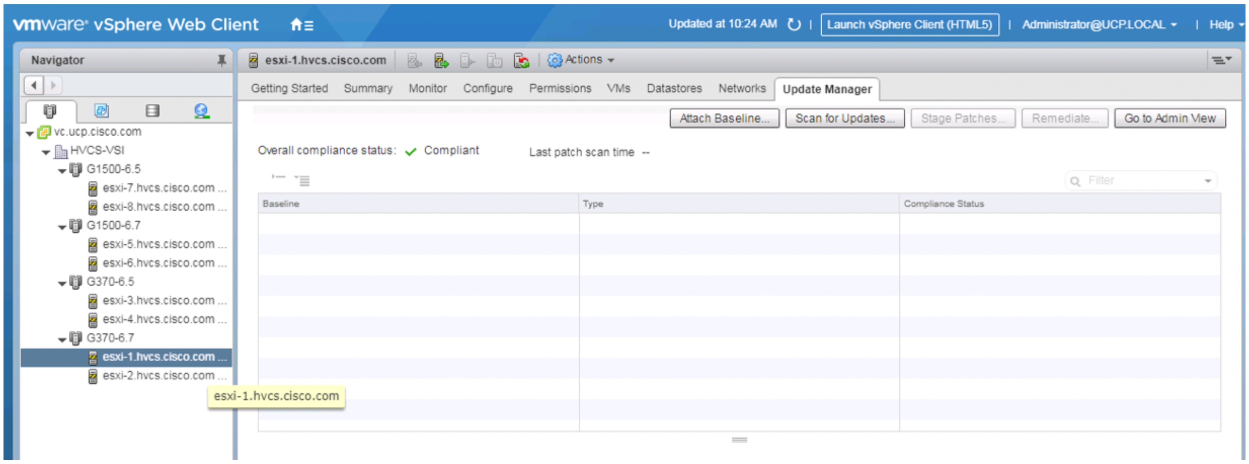

Overall compliance status: ✓ Compliant   Last patch scan time --

Baseline	Type	Compliance Status

16. Select the appropriate Patch Baseline and click OK.

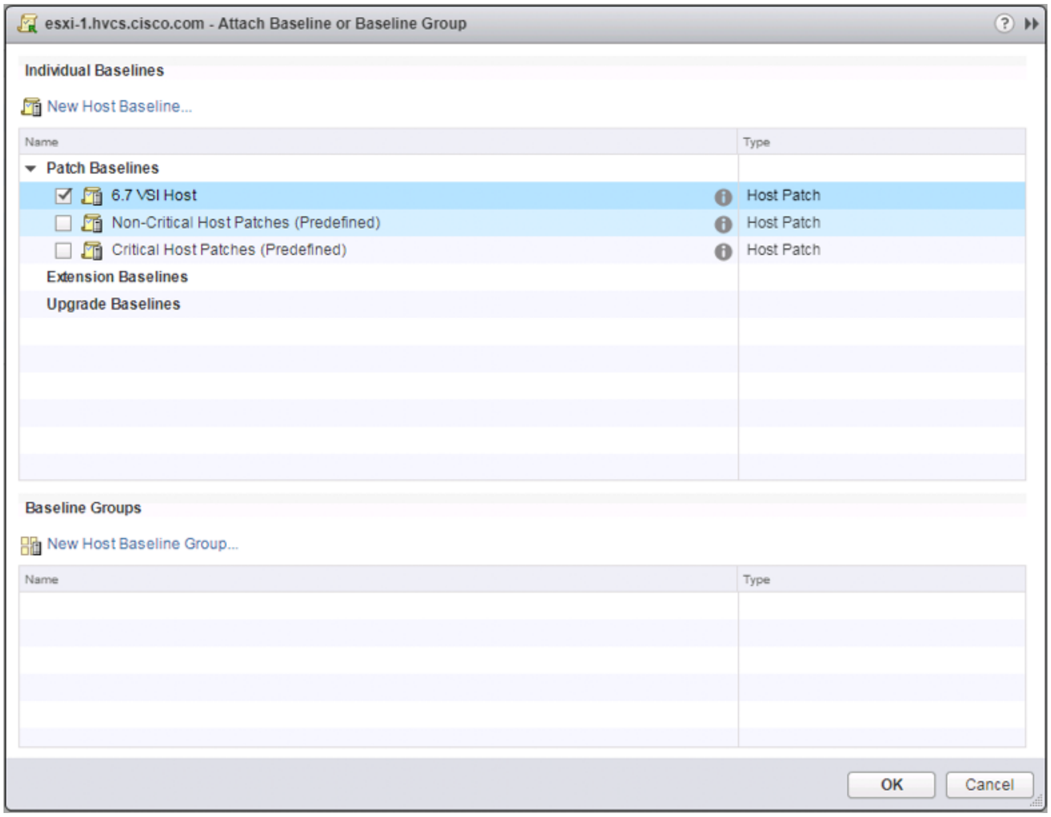
esxi-1.hvcs.cisco.com - Attach Baseline or Baseline Group
Individual Baselines
New Host Baseline...
Baseline Groups
New Host Baseline Group...
OK
Cancel

Name	Type
Patch Baselines	
☑ 6.7 VSI Host	Host Patch
□ Non-Critical Host Patches (Predefined)	Host Patch
□ Critical Host Patches (Predefined)	Host Patch
Extension Baselines	
Upgrade Baselines	

Name	Type

17. Click Remediate.

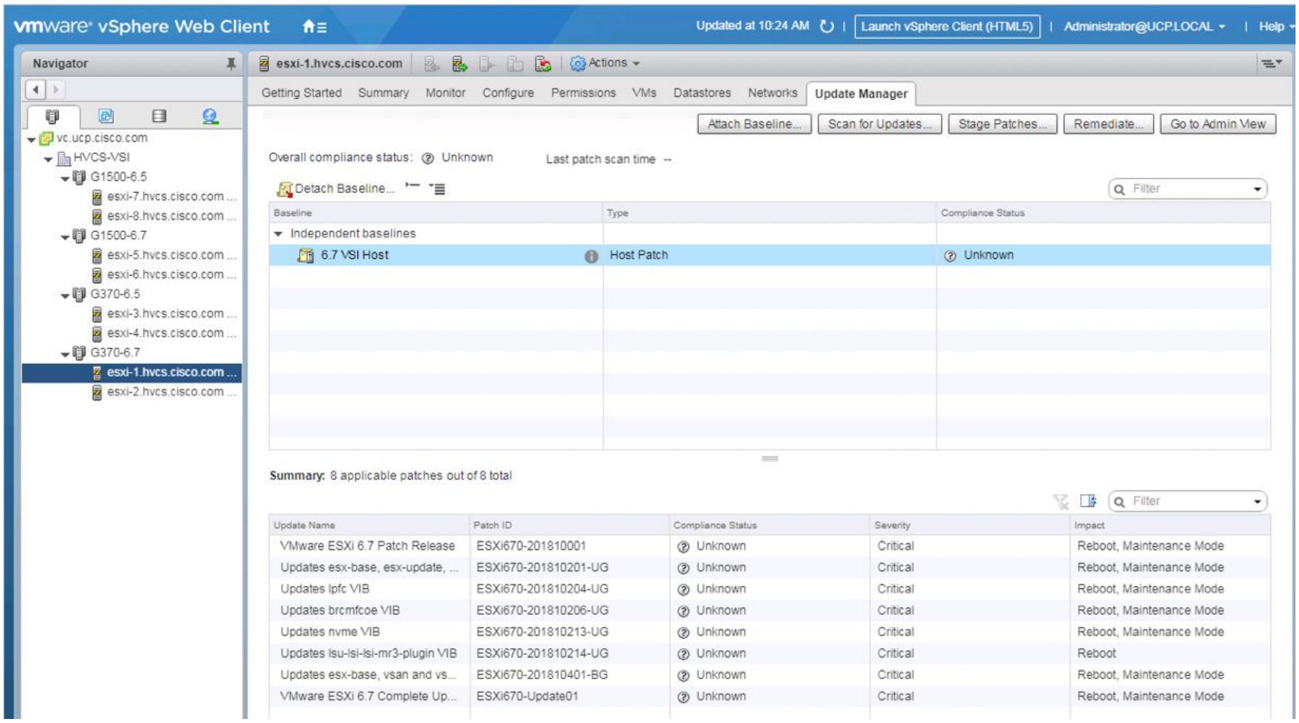

Overall compliance status: Unknown   Last patch scan time --

Baseline	Type	Compliance Status
Independent baselines		
6.7 VSI Host	Host Patch	Unknown

Summary: 8 applicable patches out of 8 total

Update Name	Patch ID	Compliance Status	Severity	Impact
VMware ESXi 6.7 Patch Release	ESXi670-201810001	Unknown	Critical	Reboot, Maintenance Mode
Updates esx-base, esx-update, ...	ESXi670-201810201-UG	Unknown	Critical	Reboot, Maintenance Mode
Updates lpfc VIB	ESXi670-201810204-UG	Unknown	Critical	Reboot, Maintenance Mode
Updates brcmfcoe VIB	ESXi670-201810206-UG	Unknown	Critical	Reboot, Maintenance Mode
Updates nvme VIB	ESXi670-201810213-UG	Unknown	Critical	Reboot, Maintenance Mode
Updates lsu-lsi-lsi-mr3-plugin VIB	ESXi670-201810214-UG	Unknown	Critical	Reboot
Updates esx-base, vsan and vs...	ESXi670-201810401-BG	Unknown	Critical	Reboot, Maintenance Mode
VMware ESXi 6.7 Complete Up...	ESXi670-Update01	Unknown	Critical	Reboot, Maintenance Mode

18. Leave the baseline selected and click Next.

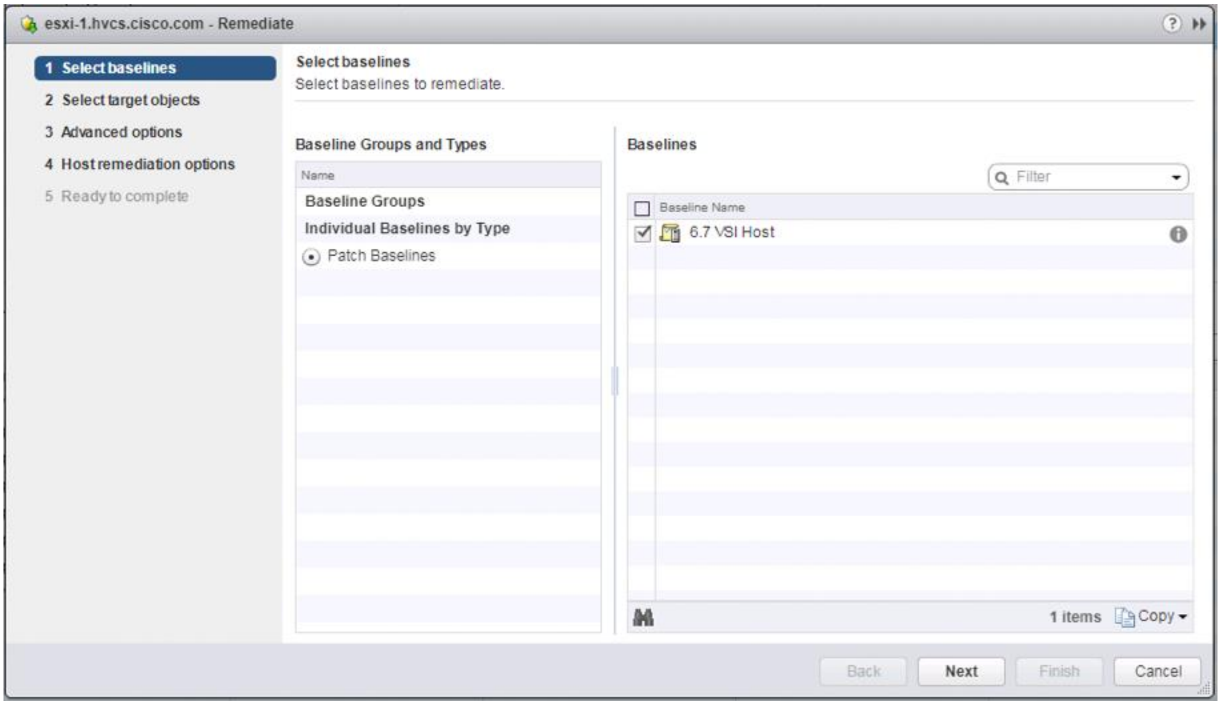

Select baselines

Select baselines to remediate.

Baseline Groups and Types

Name
Baseline Groups
Individual Baselines by Type
● Patch Baselines

Baselines

□	Baseline Name
☑	6.7 VSI Host

1 items

19. Select the hosts appropriate to the vSphere release specified for the baseline and click Next.

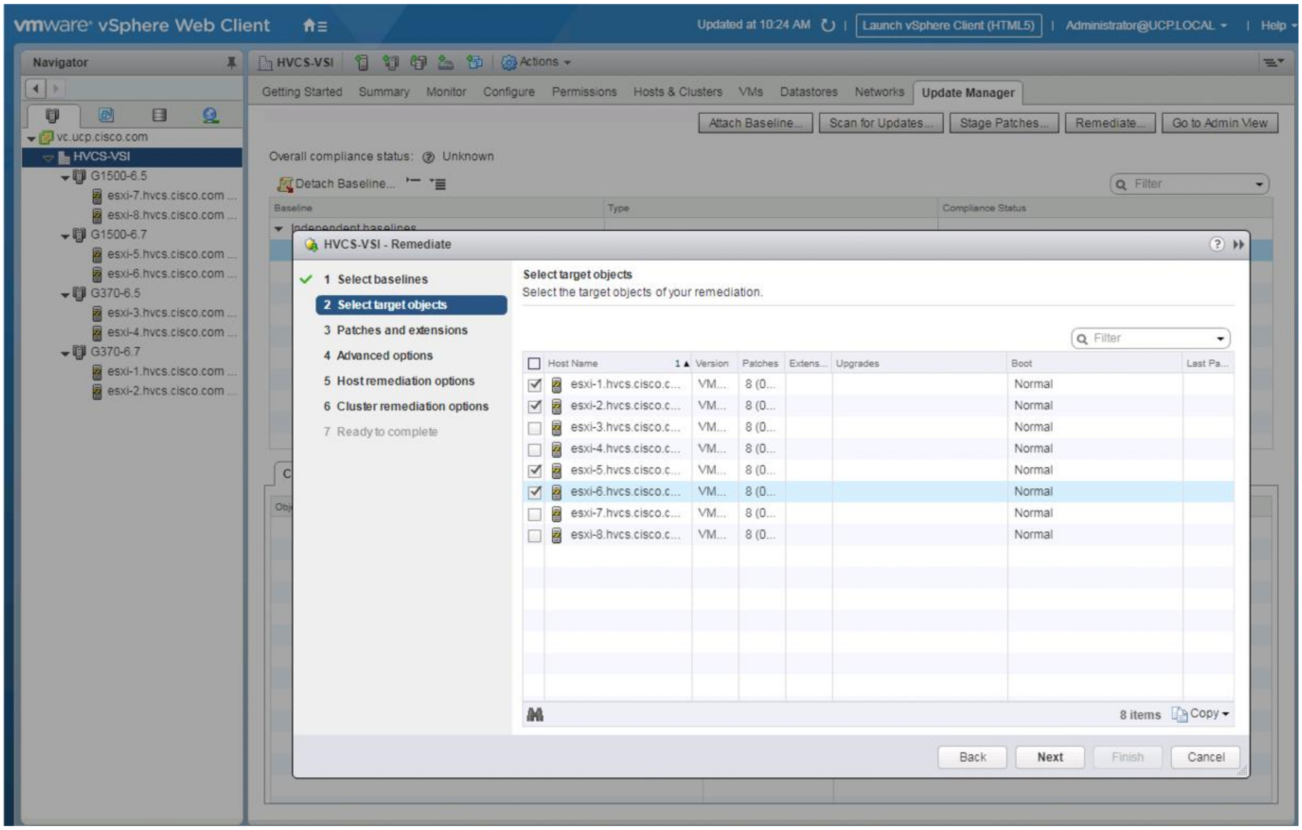

Select target objects

Select the target objects of your remediation.

□	Host Name 1 ▲	Version	Patches	Extens...	Upgrades	Boot	Last Pa...
☑	esxi-1.hvcs.cisco.c...	VM...	8 (0...			Normal	
☑	esxi-2.hvcs.cisco.c...	VM...	8 (0...			Normal	
□	esxi-3.hvcs.cisco.c...	VM...	8 (0...			Normal	
□	esxi-4.hvcs.cisco.c...	VM...	8 (0...			Normal	
☑	esxi-5.hvcs.cisco.c...	VM...	8 (0...			Normal	
☑	esxi-6.hvcs.cisco.c...	VM...	8 (0...			Normal	
□	esxi-7.hvcs.cisco.c...	VM...	8 (0...			Normal	
□	esxi-8.hvcs.cisco.c...	VM...	8 (0...			Normal	

8 items

20. Deselect any patches that should not be applied and click Next.

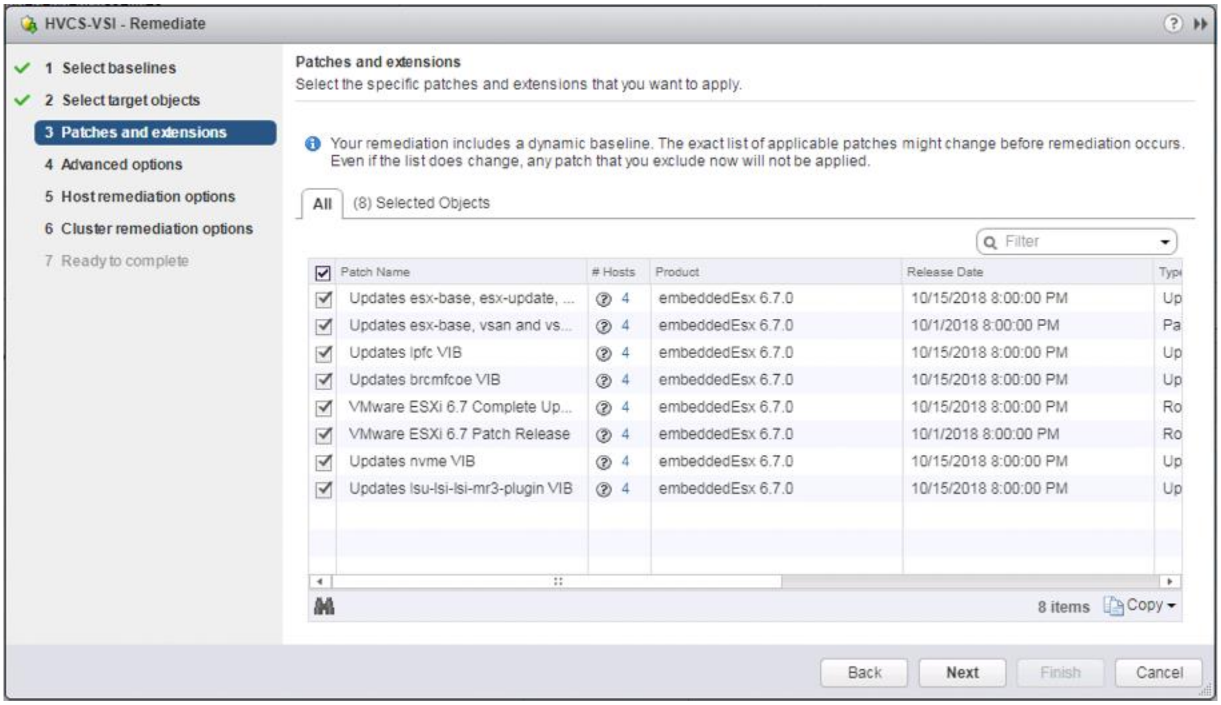

HVCS-VSI - Remediate

- ✓ 1 Select baselines
- ✓ 2 Select target objects
- 3 Patches and extensions
- 4 Advanced options
- 5 Host remediation options
- 6 Cluster remediation options
- 7 Ready to complete

Patches and extensions

Select the specific patches and extensions that you want to apply.

Your remediation includes a dynamic baseline. The exact list of applicable patches might change before remediation occurs. Even if the list does change, any patch that you exclude now will not be applied.

All | (8) Selected Objects

☑	Patch Name	# Hosts	Product	Release Date	Typ
☑	Updates esx-base, esx-update, ...	4	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Up
☑	Updates esx-base, vsan and vs...	4	embeddedEsx 6.7.0	10/1/2018 8:00:00 PM	Pa
☑	Updates lpfc VIB	4	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Up
☑	Updates brcmfcoe VIB	4	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Up
☑	VMware ESXi 6.7 Complete Up...	4	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Ro
☑	VMware ESXi 6.7 Patch Release	4	embeddedEsx 6.7.0	10/1/2018 8:00:00 PM	Ro
☑	Updates nvme VIB	4	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Up
☑	Updates lsu-lsi-lsi-mr3-plugin VIB	4	embeddedEsx 6.7.0	10/15/2018 8:00:00 PM	Up

8 items

21. Click Next past the Advanced options screen.

22. Click Next past the Host remediation options screen.

23. Click Next past the Cluster remediation options.

24. Review the settings and click Finish to run the patch baseline.

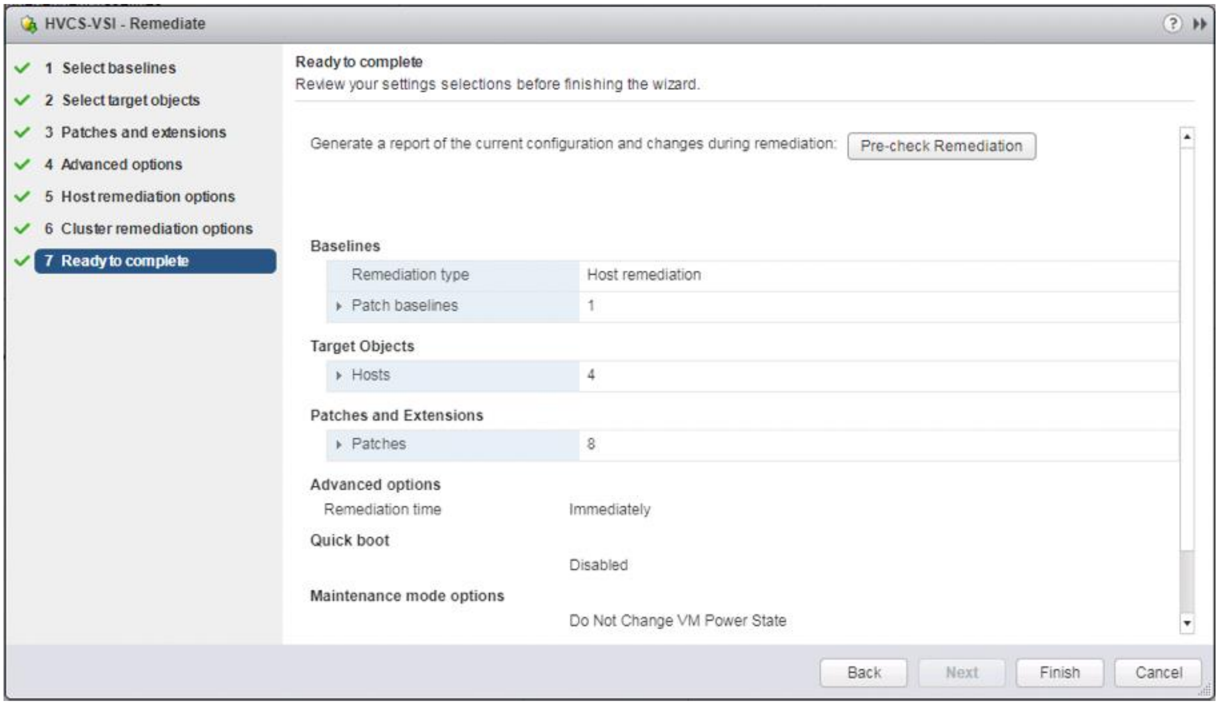

HVCS-VSI - Remediate

- ✓ 1 Select baselines
- ✓ 2 Select target objects
- ✓ 3 Patches and extensions
- ✓ 4 Advanced options
- ✓ 5 Host remediation options
- ✓ 6 Cluster remediation options
- ✓ 7 Ready to complete

Ready to complete

Review your settings selections before finishing the wizard.

Generate a report of the current configuration and changes during remediation: Pre-check Remediation

Baselines

Remediation type	Host remediation
Patch baselines	1

Target Objects

Hosts	4

Patches and Extensions

Patches	8

Advanced options

Remediation time: Immediately

Quick boot

Disabled

Maintenance mode options

Do Not Change VM Power State

vSphere nfnic and nenic Baseline Creation through VUM

With the current Cisco custom ISO for 6.7 U1 and the 4.0(2b) UCSM code, there may be some FC path inconsistencies if a newer nfnic driver is not applied to the ESXi hosts. The following steps will show the creation of a baseline within VUM for these updated drivers, and the CLI update option will be explained as well following this section.

The two drivers updated within this validation can be downloaded from the VMware site:

nfnic - https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI67-CISCO-NFNIC-40033&productId=742

nenic - https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI67-CISCO-NENIC-10270&productId=742

Each of these downloads will come as a zip file to be extracted:

- VMW-ESX-6.7.0-nfnic-4.0.0.33-13031113.zip
- VMW-ESX-6.7.0-nenic-1.0.27.0-11271332.zip

The folders created from each extracted bundle will contain an offline_bundle zip file that will stay compressed:

- VMW-ESX-6.7.0-nfnic-4.0.0.33-offline_bundle-13031113.zip
- VMW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-11271332.zip

To create the driver baselines and patch the new ESXi hosts within VUM, follow these steps:

1. From the Hosts tab select the vCenter and go to the Update Manager tab.
2. Click Go to Admin View.

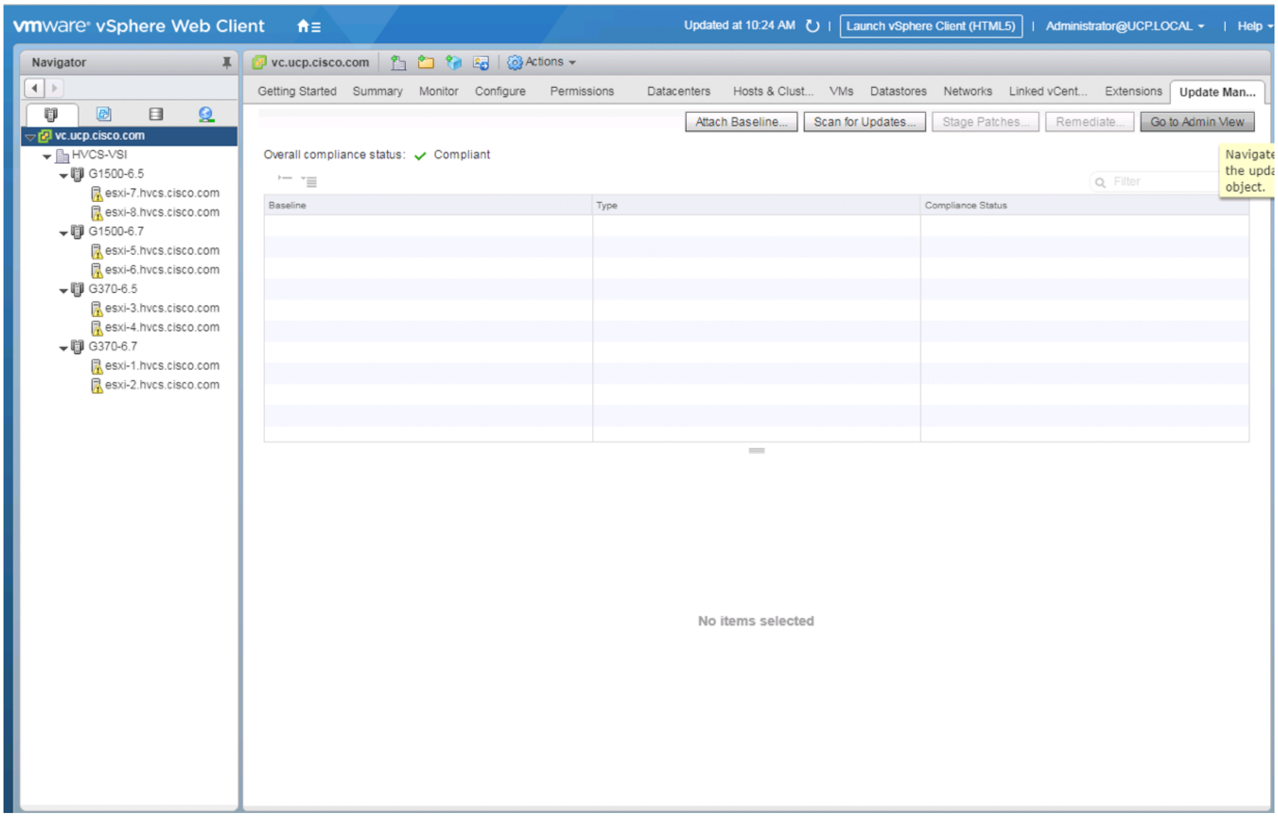

Overall compliance status: ✓ Compliant

No items selected

3. From the Manage tab select Hosts Baselines.

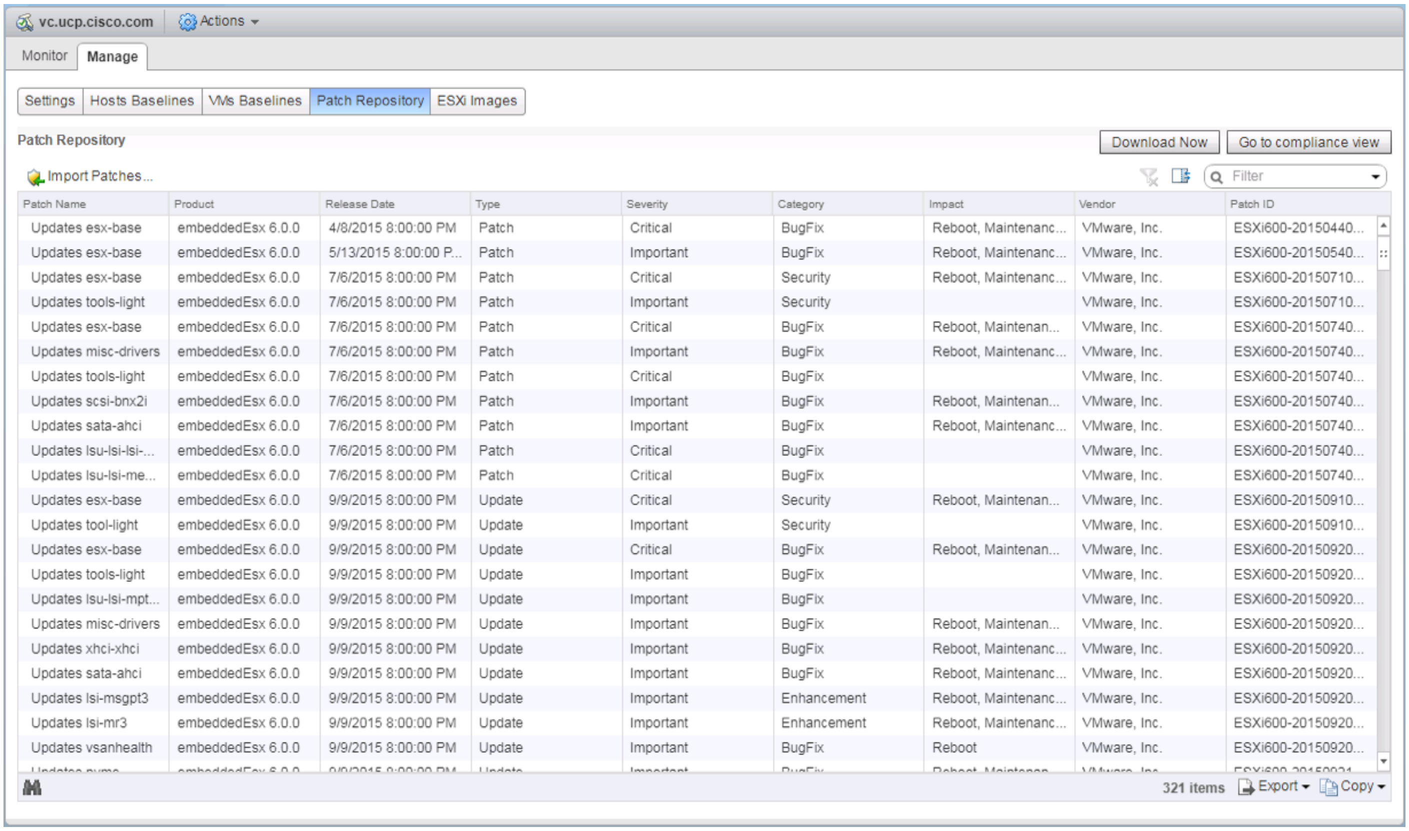

Patch Repository

Patch Name	Product	Release Date	Type	Severity	Category	Impact	Vendor	Patch ID
Updates esx-base	embeddedEsx 6.0.0	4/8/2015 8:00:00 PM	Patch	Critical	BugFix	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150440...
Updates esx-base	embeddedEsx 6.0.0	5/13/2015 8:00:00 P...	Patch	Important	BugFix	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150540...
Updates esx-base	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Critical	Security	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150710...
Updates tools-light	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Important	Security		VMware, Inc.	ESXi600-20150710...
Updates esx-base	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Critical	BugFix	Reboot, Maintenan...	VMware, Inc.	ESXi600-20150740...
Updates misc-drivers	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Important	BugFix	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150740...
Updates tools-light	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Critical	BugFix		VMware, Inc.	ESXi600-20150740...
Updates scsi-bnx2i	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Important	BugFix	Reboot, Maintenan...	VMware, Inc.	ESXi600-20150740...
Updates sata-ahci	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Important	BugFix	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150740...
Updates lsu-lsi-lsi-...	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Critical	BugFix		VMware, Inc.	ESXi600-20150740...
Updates lsu-lsi-me...	embeddedEsx 6.0.0	7/6/2015 8:00:00 PM	Patch	Critical	BugFix		VMware, Inc.	ESXi600-20150740...
Updates esx-base	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Critical	Security	Reboot, Maintenan...	VMware, Inc.	ESXi600-20150910...
Updates tool-light	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	Security		VMware, Inc.	ESXi600-20150910...
Updates esx-base	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Critical	BugFix	Reboot, Maintenan...	VMware, Inc.	ESXi600-20150920...
Updates tools-light	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	BugFix		VMware, Inc.	ESXi600-20150920...
Updates lsu-lsi-mpt...	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	BugFix		VMware, Inc.	ESXi600-20150920...
Updates misc-drivers	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	BugFix	Reboot, Maintenan...	VMware, Inc.	ESXi600-20150920...
Updates xhci-xhci	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	BugFix	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150920...
Updates sata-ahci	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	BugFix	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150920...
Updates lsi-msgpt3	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	Enhancement	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150920...
Updates lsi-mr3	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	Enhancement	Reboot, Maintenanc...	VMware, Inc.	ESXi600-20150920...
Updates vsanhealth	embeddedEsx 6.0.0	9/9/2015 8:00:00 PM	Update	Important	BugFix	Reboot	VMware, Inc.	ESXi600-20150920...

321 items

4. Click the **Import Patches…** under the Patch Repository heading.

5. Click the **Browse…** button and find the first offline_bundle zip file to be uploaded.

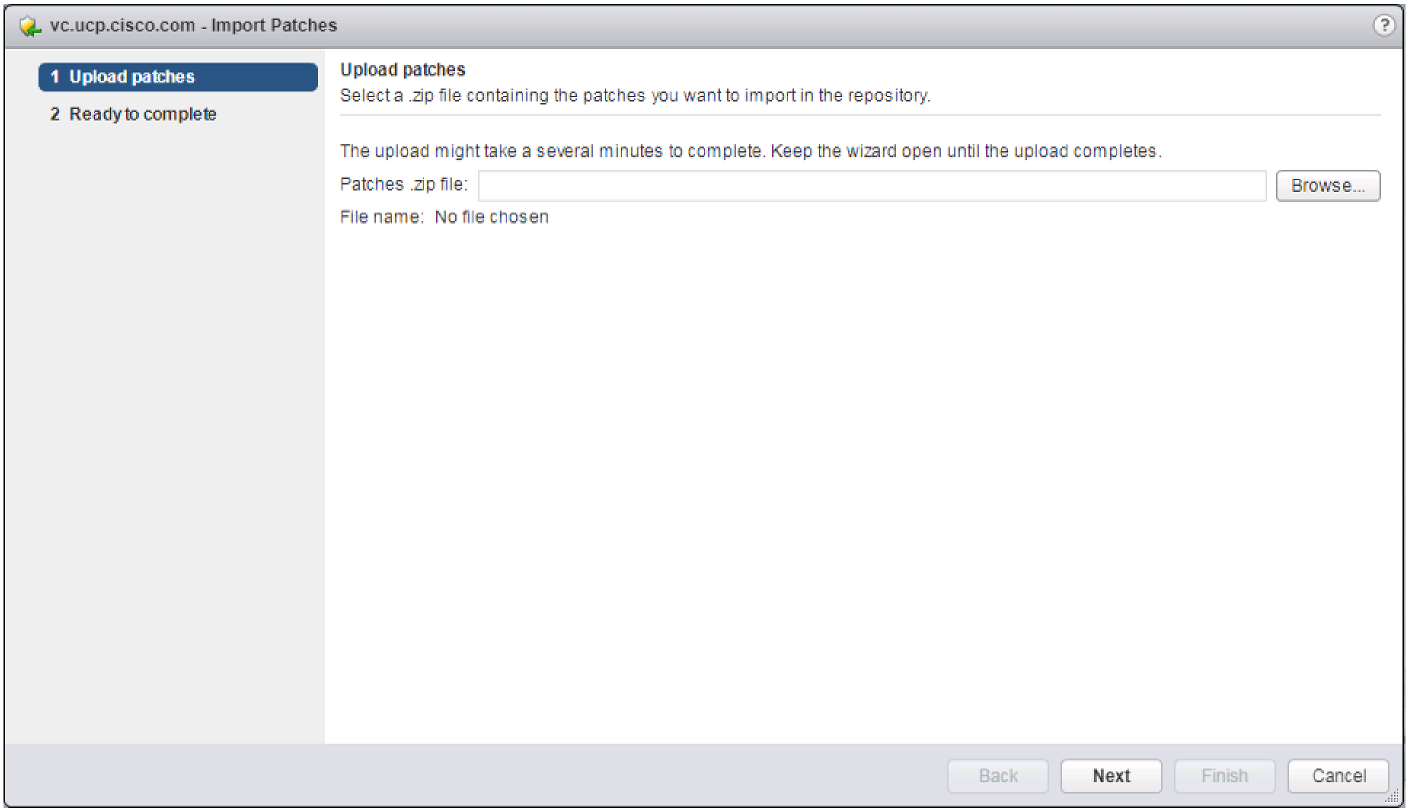

vc.ucp.cisco.com - Import Patches

Upload patches

Select a .zip file containing the patches you want to import in the repository.

The upload might take a several minutes to complete. Keep the wizard open until the upload completes.

Patches .zip file:

File name: No file chosen

6. Click Next after the file has uploaded.

7. Review the summary from the Ready to complete screen and click Finish.

8. Repeat steps 4-7 for the second offline_bundle zip.

9. From the Manage tab select Hosts Baselines.

10. Click +New Baseline…

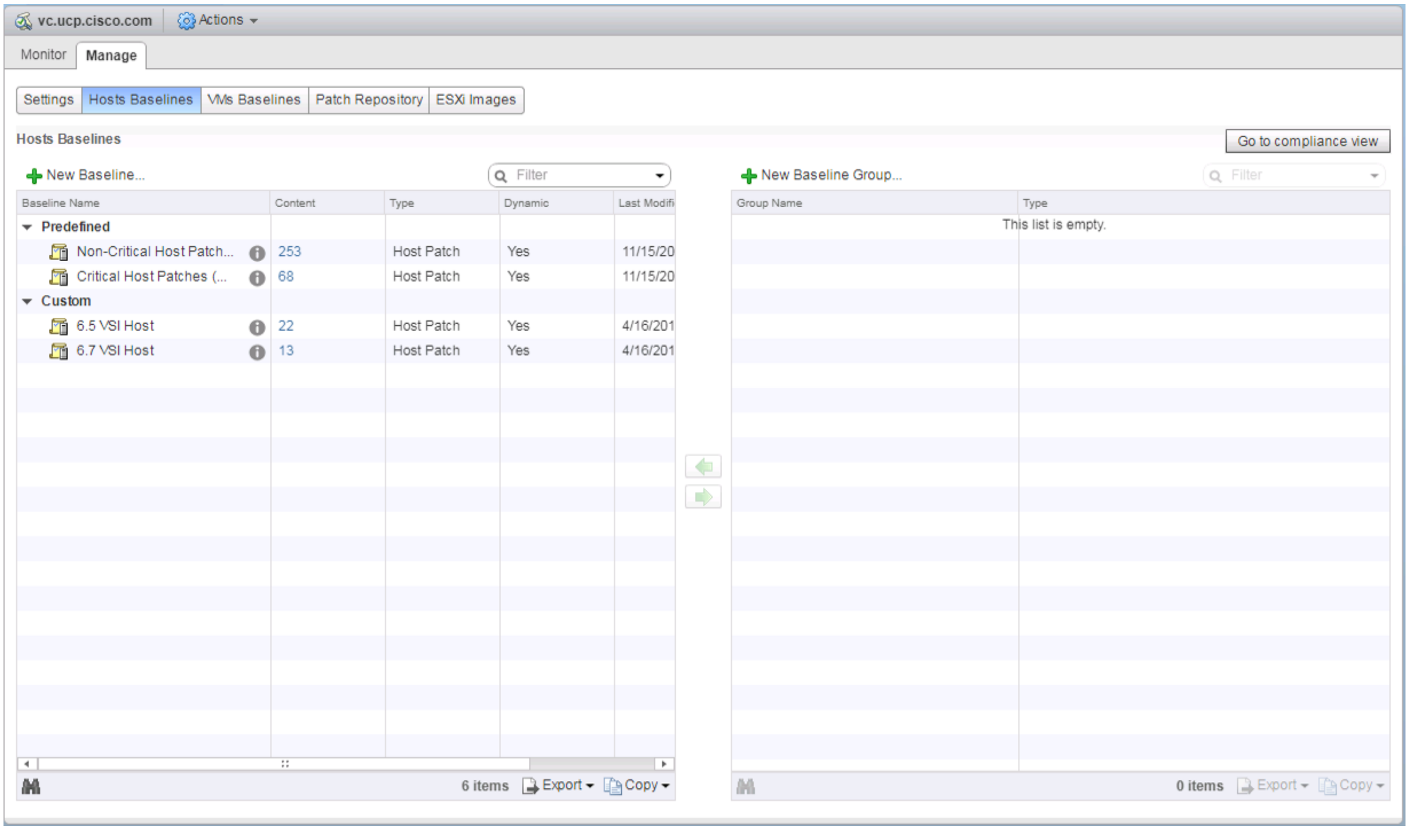

Baseline Name	Content	Type	Dynamic	Last Modifi
Predefined				
Non-Critical Host Patch...	253	Host Patch	Yes	11/15/20
Critical Host Patches (...	68	Host Patch	Yes	11/15/20
Custom				
6.5 VSI Host	22	Host Patch	Yes	4/16/201
6.7 VSI Host	13	Host Patch	Yes	4/16/201

6 items

Group Name	Type

This list is empty.

0 items

11. Provide a name for the baseline and leave Host Patch selected for Baseline type.

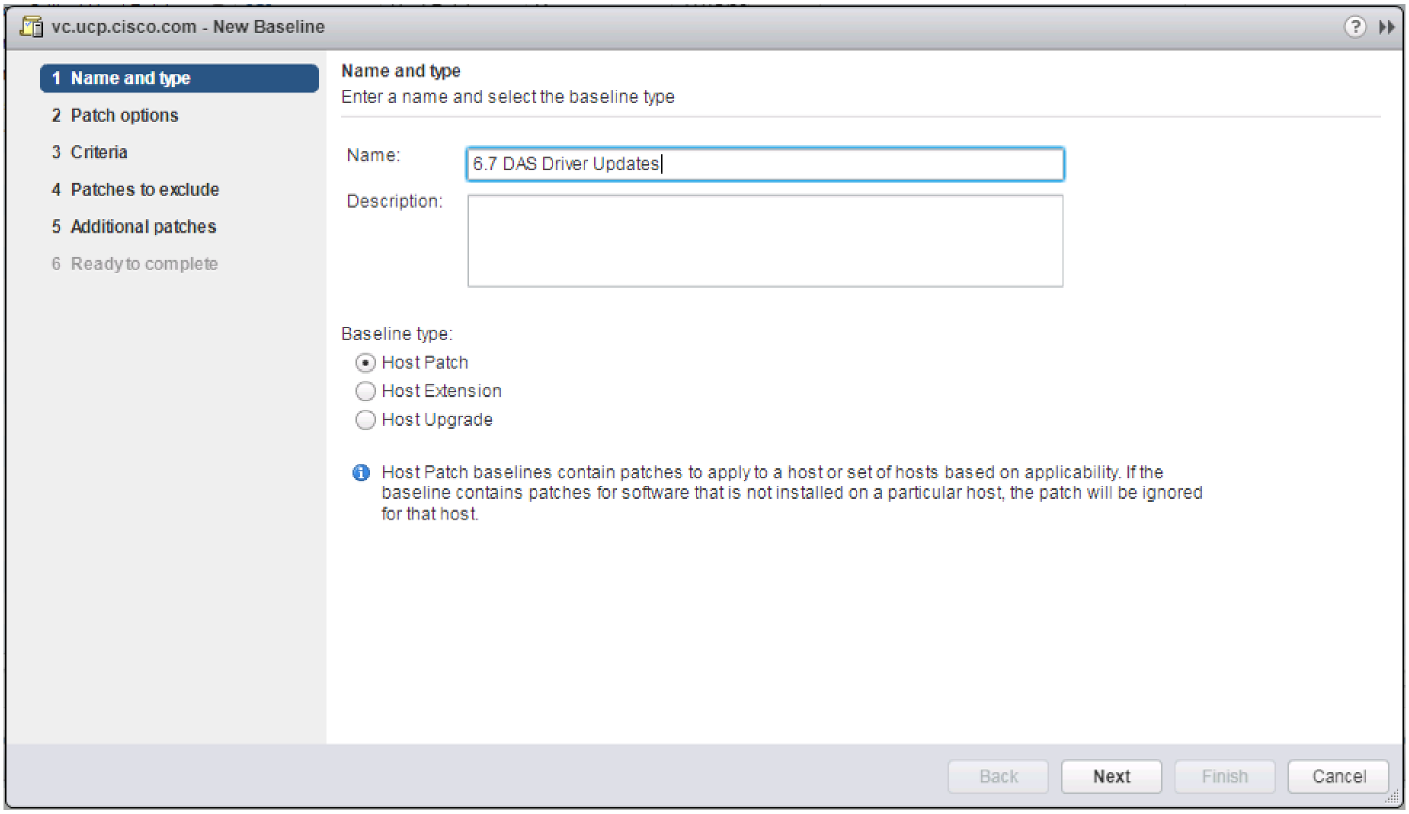

Name and type

Enter a name and select the baseline type

Name: 6.7 DAS Driver Updates

Description:

Baseline type:

● Host Patch

○ Host Extension

○ Host Upgrade

Host Patch baselines contain patches to apply to a host or set of hosts based on applicability. If the baseline contains patches for software that is not installed on a particular host, the patch will be ignored for that host.

12. Click **Next**.

13. Select Fixed for the Patch options.

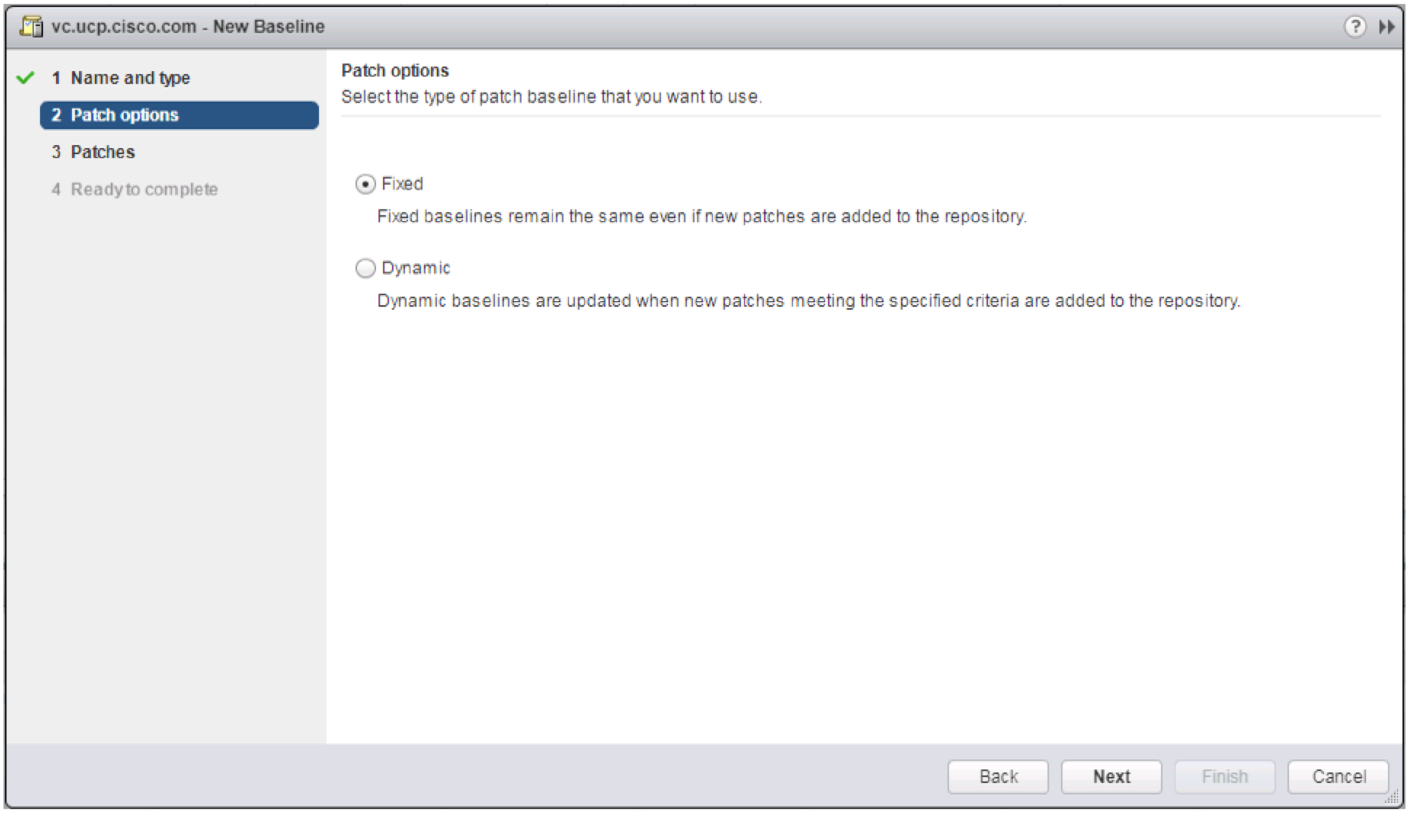

vc.ucp.cisco.com - New Baseline

Patch options

Select the type of patch baseline that you want to use.

◉ Fixed

Fixed baselines remain the same even if new patches are added to the repository.

○ Dynamic

Dynamic baselines are updated when new patches meeting the specified criteria are added to the repository.

14. Click **Next**.

15. Enter “cisco” in the search filter and select the appropriate patches that come up in the results.

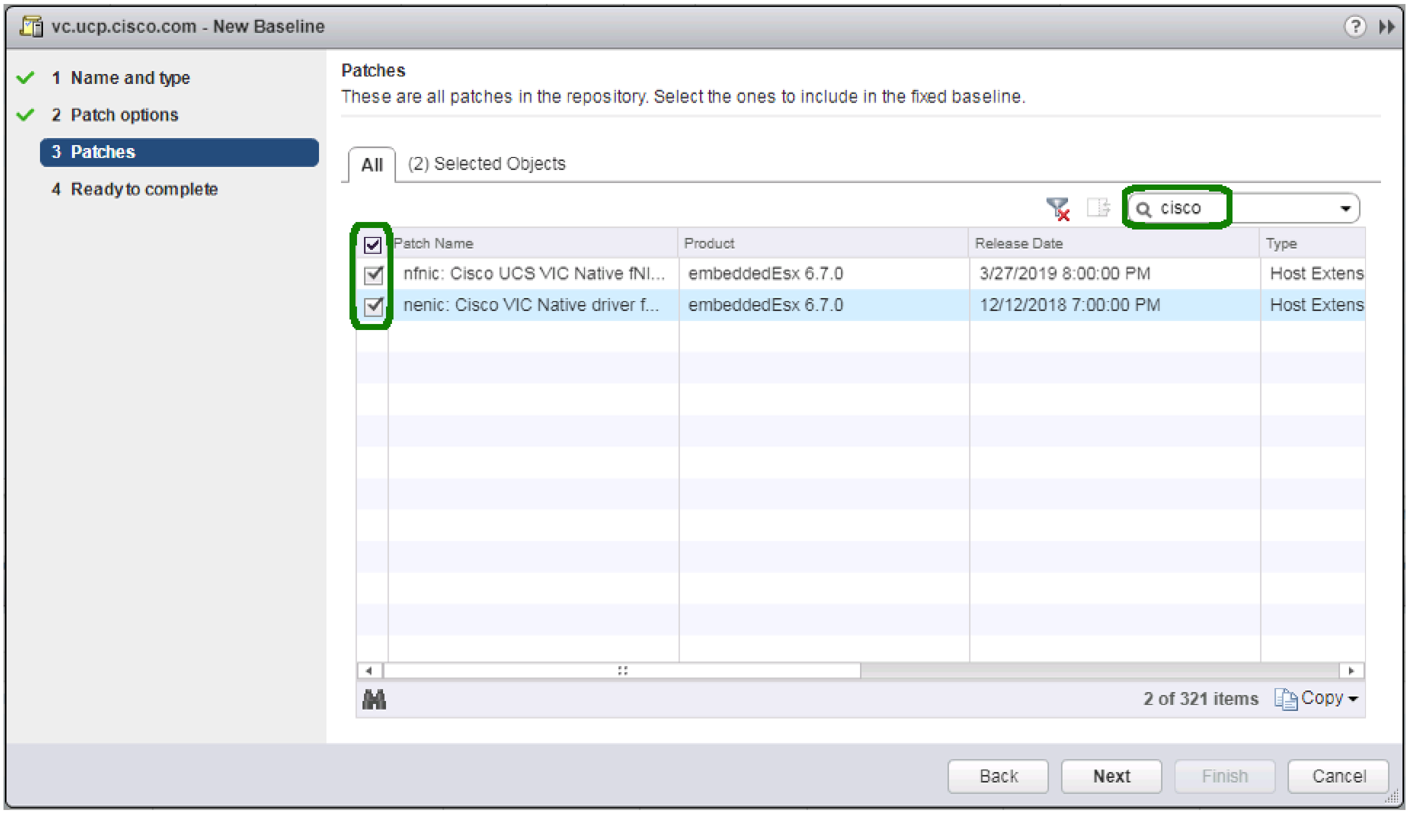

Patches

These are all patches in the repository. Select the ones to include in the fixed baseline.

☑	Patch Name	Product	Release Date	Type
☑	nfnic: Cisco UCS VIC Native fNI...	embeddedEsx 6.7.0	3/27/2019 8:00:00 PM	Host Extens
☑	nenic: Cisco VIC Native driver f...	embeddedEsx 6.7.0	12/12/2018 7:00:00 PM	Host Extens

2 of 321 items

16. Click **Next**.

17. Review the Ready to complete summary and click **Finish**.

18. Return to the Datacenter host view and select the host(s) to patch.

19. Place the hosts in maintenance mode and select the Update Manager tab for the host.

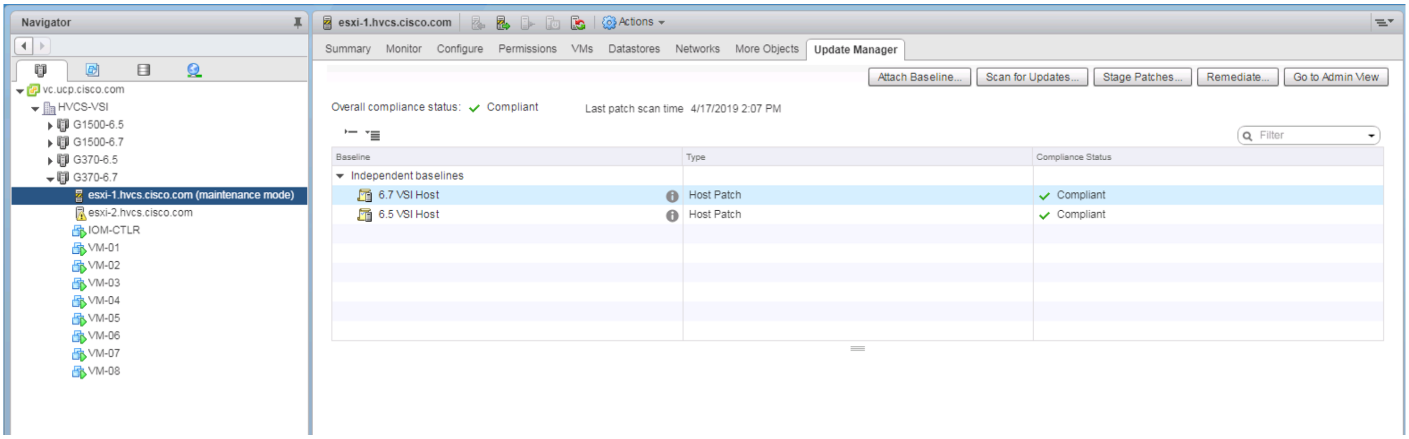

Overall compliance status: ✓ Compliant Last patch scan time 4/17/2019 2:07 PM

Baseline	Type	Compliance Status
Independent baselines		
6.7 VSI Host	Host Patch	✓ Compliant
6.5 VSI Host	Host Patch	✓ Compliant

20. Click the Attach Baseline… button.

21. Select the created driver update baseline.

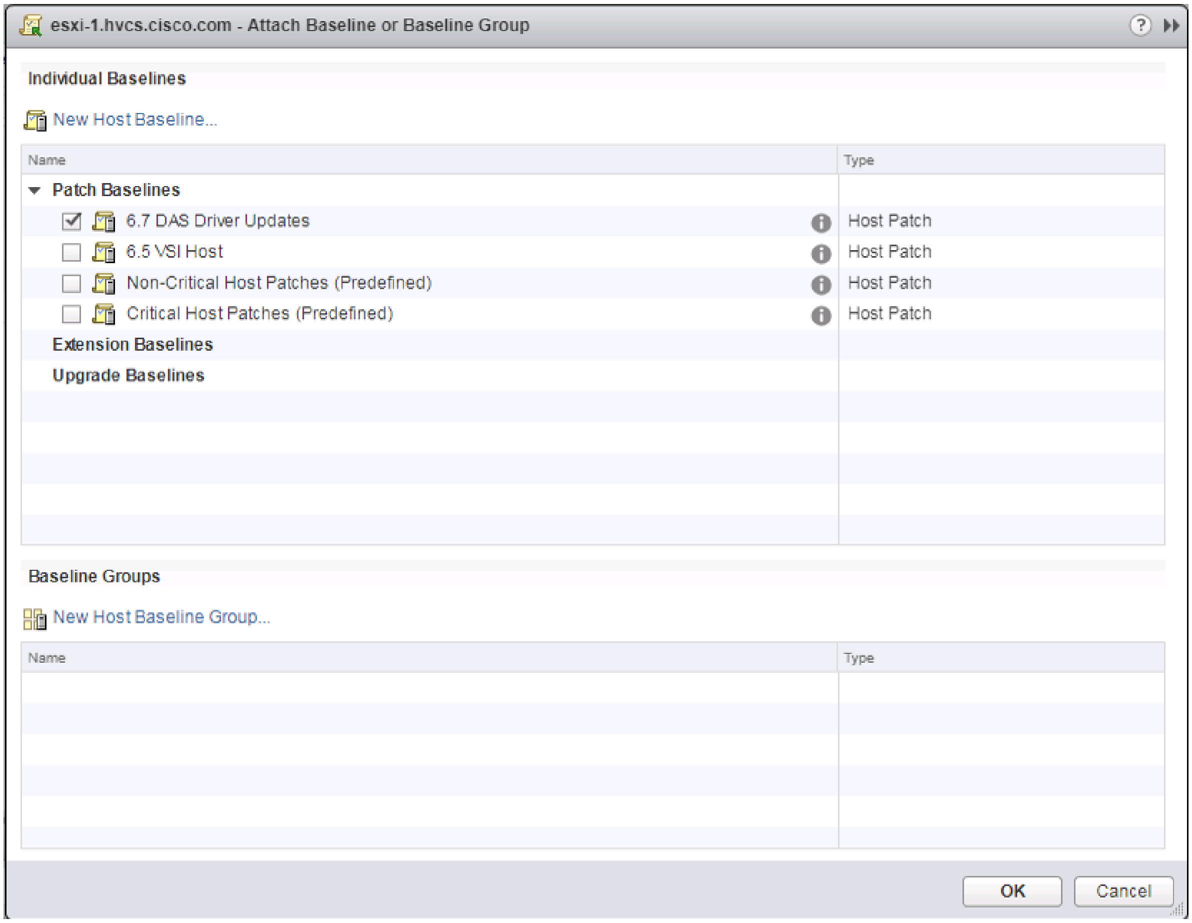

esxi-1.hvcs.cisco.com - Attach Baseline or Baseline Group

Individual Baselines

New Host Baseline...

Name	Type
Patch Baselines	
☑ 6.7 DAS Driver Updates	Host Patch
□ 6.5 VSI Host	Host Patch
□ Non-Critical Host Patches (Predefined)	Host Patch
□ Critical Host Patches (Predefined)	Host Patch
Extension Baselines	
Upgrade Baselines	

Baseline Groups

New Host Baseline Group...

Name	Type

22. Click **OK** to attach the baseline.

23. Click the **Remediate…** button.

24. Select the driver updates baseline.

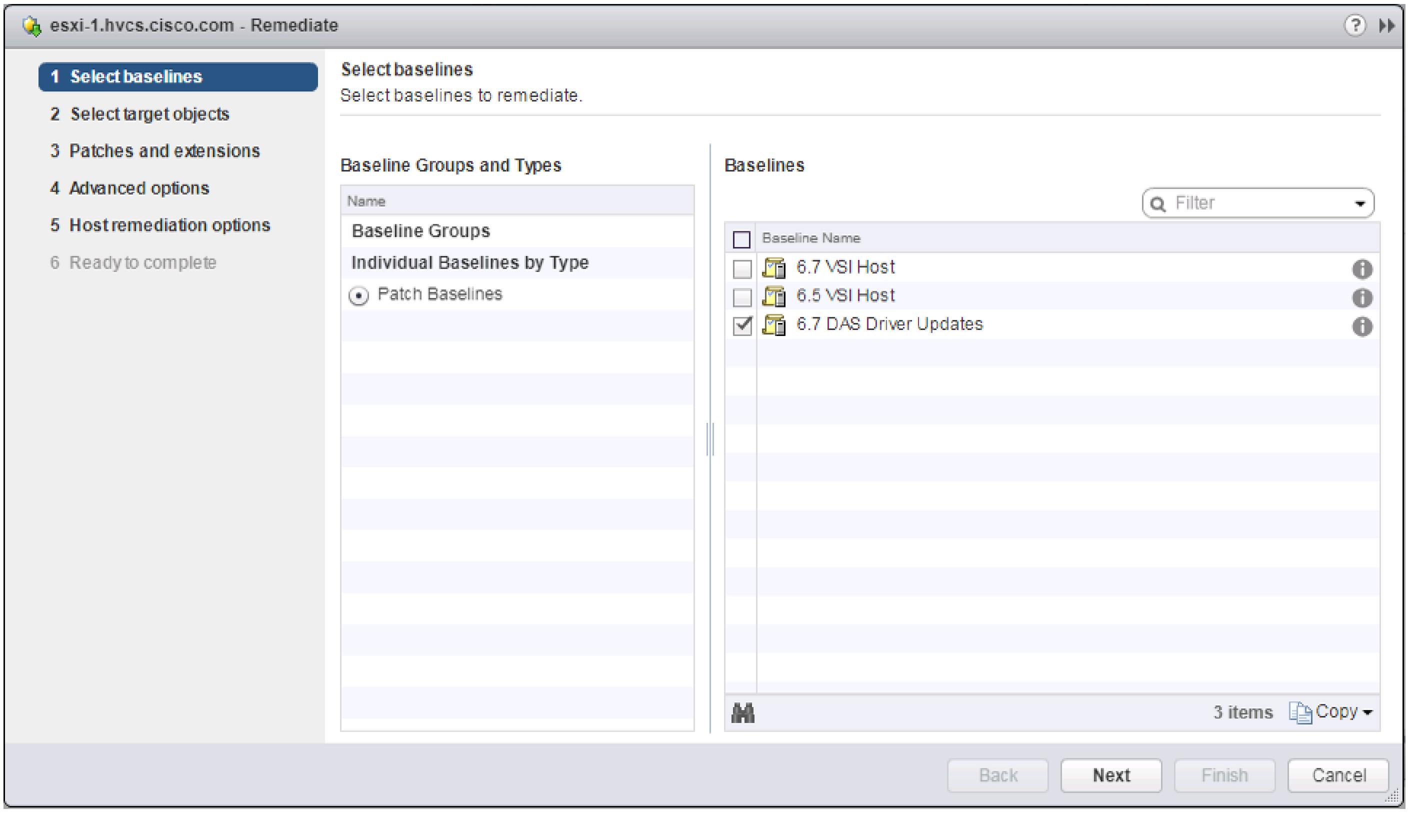

esxi-1.hvcs.cisco.com - Remediate

1 Select baselines
2 Select target objects
3 Patches and extensions
4 Advanced options
5 Host remediation options
6 Ready to complete

Select baselines

Select baselines to remediate.

Baseline Groups and Types

Name
Baseline Groups
Individual Baselines by Type
● Patch Baselines

Baselines

□	Baseline Name
□	6.7 VSI Host
□	6.5 VSI Host
☑	6.7 DAS Driver Updates

3 items

25. Click **Next**.

26. Click **Next** past the Select target objects screen.

27. Click **Next** past the Patches and extensions screen.

28. Click **Next** past the Advanced options screen.

29. Click **Next** past the Host remediation options.

30. Click **Next** past the Cluster remediation options.

31. Review the Ready to complete summary and click **Finish** to apply the patches.

32. Repeat steps 19 to 31 for each additional host.

vSphere nfnic and nenic Patching through esxcli

With the current Cisco custom ISO for 6.7 U1 and the 4.0(2b) UCSM code, there may be some FC path inconsistencies if a newer nfnic driver is not applied to the ESXi hosts. The following steps will show the installation of these updated drivers using the CLI update option.

The two drivers updated within this validation can be downloaded from the VMware site:

nfnic - https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI67-CISCO-NFNIC-40033&productId=742

nenic - https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI67-CISCO-NENIC-10270&productId=742

Each of these downloads will come as a zip file to be extracted:

- VMW-ESX-6.7.0-nfnic-4.0.0.33-13031113.zip
- VMW-ESX-6.7.0-nenic-1.0.27.0-11271332.zip

The folders created from each extracted bundle will contain an offline_bundle zip file that will stay compressed:

- VMW-ESX-6.7.0-nfnic-4.0.0.33-offline_bundle-13031113.zip
- VMW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-11271332.zip

To install VMware VIC Drivers on the ESXi hosts using the esxcli, follow these steps:

1. Upload the offline_bundles to a commonly accessible datastore using the vSphere Web Client.
2. Within the vSphere Web Client, select one of the datastore and click on the Files tab.

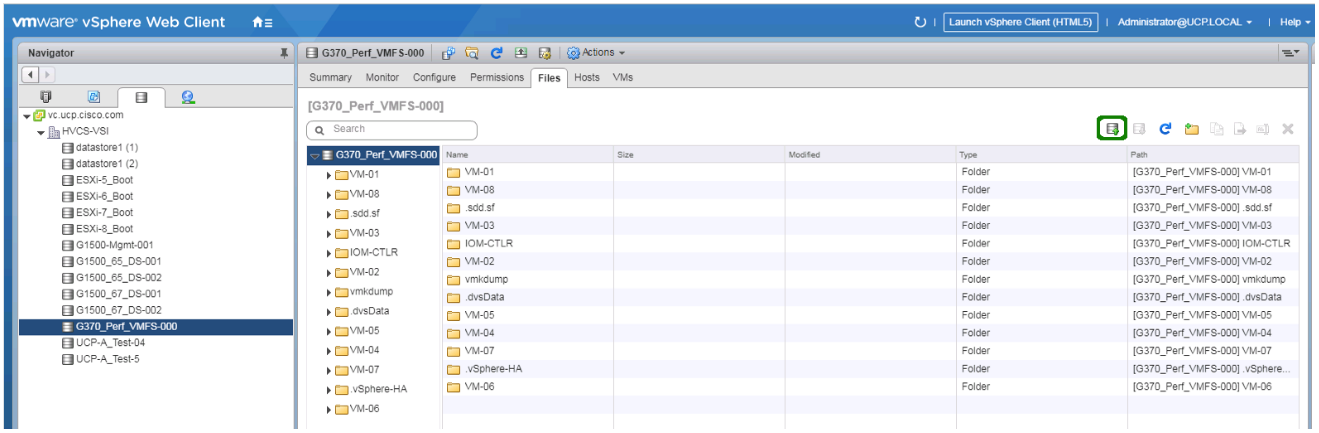

Name	Size	Modified	Type	Path
VM-01			Folder	[G370_Perf_VMFS-000] VM-01
VM-08			Folder	[G370_Perf_VMFS-000] VM-08
.sdd.sf			Folder	[G370_Perf_VMFS-000] .sdd.sf
VM-03			Folder	[G370_Perf_VMFS-000] VM-03
IOM-CTLR			Folder	[G370_Perf_VMFS-000] IOM-CTLR
VM-02			Folder	[G370_Perf_VMFS-000] VM-02
vmkdump			Folder	[G370_Perf_VMFS-000] vmkdump
.dvsData			Folder	[G370_Perf_VMFS-000] .dvsData
VM-05			Folder	[G370_Perf_VMFS-000] VM-05
VM-04			Folder	[G370_Perf_VMFS-000] VM-04
VM-07			Folder	[G370_Perf_VMFS-000] VM-07
.vSphere-HA			Folder	[G370_Perf_VMFS-000] .vSphere...
VM-06			Folder	[G370_Perf_VMFS-000] VM-06

3. Click the Upload a file to the Datastore button.
4. Select and upload the nfnic offline_bundle (VMW-ESX-6.7.0-nfnic-4.0.0.33-offline_bundle-13031113.zip) from the extracted driver download.
5. Select and upload the nenic offline_bundle (VMW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-11271332.zip) from the extracted driver download.
6. Place all hosts in Maintenance mode requiring update.
7. Connect to each ESXi host through ssh from a shell connection or putty terminal.
8. Login as root with the root password.
9. Run the following commands (substituting the appropriate datastore directory as needed) on each host:

```
esxcli software vib update -d /vmfs/volumes/G370_Perf_VMFS-000/VMW-ESX-6.7.0-
nfnic-4.0.0.33-offline_bundle-13031113.zip
esxcli software vib update -d /vmfs/volumes/G370_Perf_VMFS-000/V
MW-ESX-6.7.0-nenic-1.0.27.0-offline_bundle-11271332.zip
```

10. Reboot each host by typing `reboot` from the SSH connection after the command has been run.

Remediation of L1 Terminal Fault – VMM (L1TF) Security Vulnerability (Optional)

CVE-2018-3646 describes a new class of CPU speculative-execution vulnerabilities on Intel processors manufactured from 2009 to 2018. While optional, it is strongly recommended that these vulnerabilities be patched.

Multiple attack vectors are exposed through these vulnerabilities, and separate mitigation steps for each attack vector are necessary for complete mitigation. For more information about the specific impact and VMware's recommendations for remediation of these vulnerabilities in a VMware vSphere environment, refer to this VMware Knowledge Base article: https://kb.vmware.com/s/article/55806.

The mitigation for L1TF-VMM as recommended by VMware is broken up into three distinct phases:

1. Updating VMware vCenter and VMware ESXi software.

2. Planning and Utilization of the HTAware Mitigation Tool (if analyzing existing workloads).

3. Enablement of the ESXi Side-Channel-Aware Scheduler.

Update VMware vCenter and VMware ESXi Software

VMware vCenter must be running at specific patch levels prior to mitigation of the L1TF-VMM vulnerabilities. Table 15 lists the release version of VMware vCenter and the specific patch level that needs to be running on the vCenter managing the environment.

Table 15 VMware vCenter Versions Required for L1TF-VMM Mitigation

VMware vCenter Version	Patch Level Required for L1TF-VMM Mitigation
6.7	6.7.0d
6.5	6.5u2c

VMware ESXi must also be running at specific patch levels prior to mitigation of the L1TF-VMM vulnerabilities. If you use the Cisco custom ISOs for VMware vSphere 6.5 and 6.7 described in this guide to install the hypervisor, no action is necessary to update the ESXi servers in the environment. Table 16 lists the minimum Cisco ISO version that must be used to ensure no patching of the ESXi servers is necessary.

Table 16 Minimum Cisco ISO Versions Required for L1TF-VMM Mitigation

VMware ESXi Version	Minimum Cisco ISO Version Required
6.7 U1	VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso
6.5 U2	VMware-ESXi-6.5.0-9298722-Custom-Cisco-6.5.2.2.iso

Planning and Utilization of the HTAware Mitigation Tool

It is important to understand the impact to maximum performance on a host when comparing L1TF-VMM non-mitigated and mitigated environments. You must take these impacts into consideration whether you are deploying a greenfield environment, or simply adding capacity to an existing environment. VMware provides an article regarding capacity planning considerations and tested performance degradation when the L1TF-VMM vulnerabilities are remediated: https://kb.vmware.com/s/article/55767.

Consider using the PowerShell HTAware Mitigation Tool from VMware to analyze existing non-mitigated environments that you may be migrating virtual infrastructure from. This allows you to understand if there are any virtual machine configurations that may be impacted when moved to a mitigated environment.

Enablement of the ESXi Side-Channel-Aware Scheduler (vSphere Web Client Method)

When a non-mitigated host is running a patched version of ESXi, a suppressible warning message is displayed in the host Summary tab as shown below:

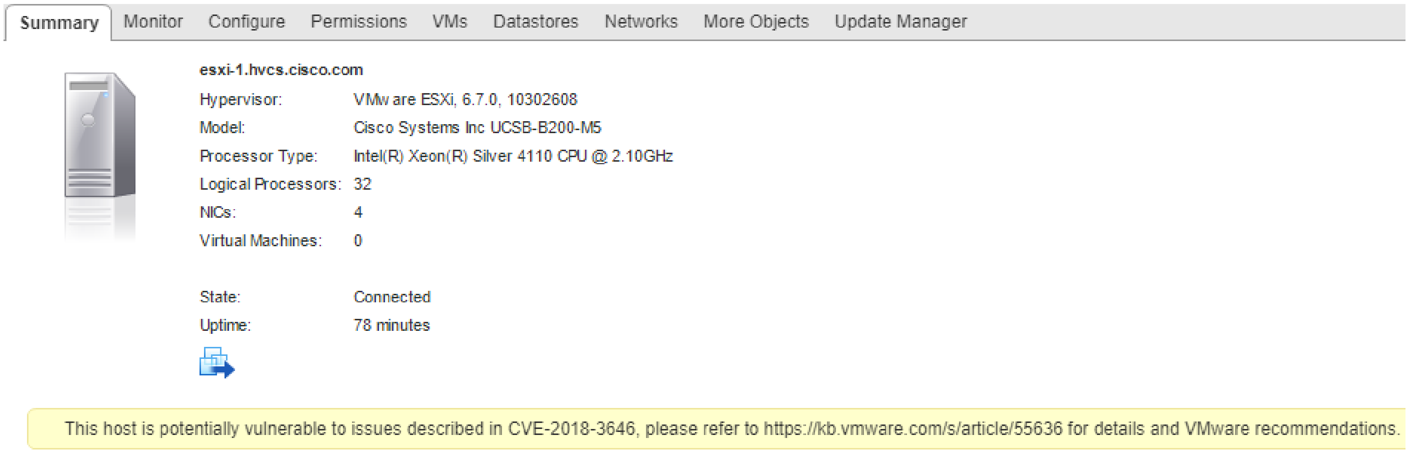

esxi-1.hvcs.cisco.com

Hypervisor: VMware ESXi, 6.7.0, 10302608
Model: Cisco Systems Inc UCSB-B200-M5
Processor Type: Intel(R) Xeon(R) Silver 4110 CPU @ 2.10GHz
Logical Processors: 32
NICs: 4
Virtual Machines: 0

State: Connected
Uptime: 78 minutes

This host is potentially vulnerable to issues described in CVE-2018-3646, please refer to https://kb.vmware.com/s/article/55636 for details and VMware recommendations.

To use the vSphere Web Client to remediate a host, follow these steps:

1. Place the host to be mitigated into Maintenance Mode.

2. Select the ESXi host from the inventory, click the Configure tab, then click System-> Advanced System Settings.

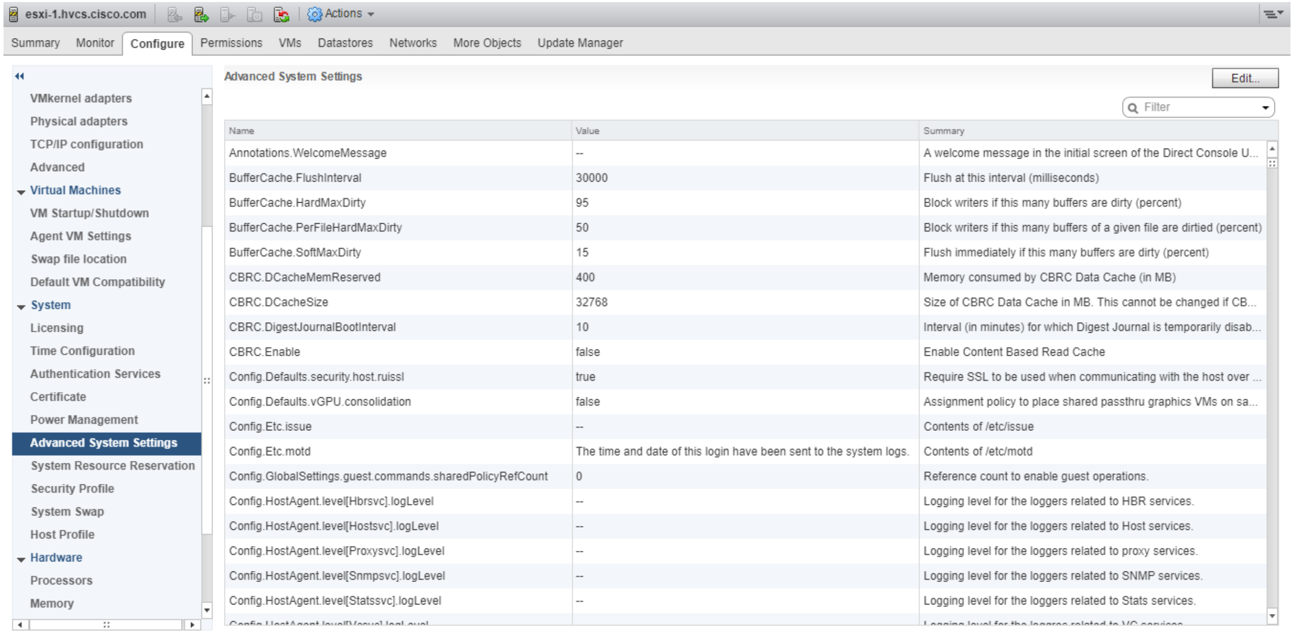

Advanced System Settings

Name	Value	Summary
Annotations.WelcomeMessage	--	A welcome message in the initial screen of the Direct Console U...
BufferCache.FlushInterval	30000	Flush at this interval (milliseconds)
BufferCache.HardMaxDirty	95	Block writers if this many buffers are dirty (percent)
BufferCache.PerFileHardMaxDirty	50	Block writers if this many buffers of a given file are dirtied (percent)
BufferCache.SoftMaxDirty	15	Flush immediately if this many buffers are dirty (percent)
CBRC.DCacheMemReserved	400	Memory consumed by CBRC Data Cache (in MB)
CBRC.DCacheSize	32768	Size of CBRC Data Cache in MB. This cannot be changed if CB...
CBRC.DigestJournalBootInterval	10	Interval (in minutes) for which Digest Journal is temporarily disab...
CBRC.Enable	false	Enable Content Based Read Cache
Config.Defaults.security.host.ruissl	true	Require SSL to be used when communicating with the host over ...
Config.Defaults.vGPU.consolidation	false	Assignment policy to place shared passthru graphics VMs on sa...
Config.Etc.issue	--	Contents of /etc/issue
Config.Etc.motd	The time and date of this login have been sent to the system logs.	Contents of /etc/motd
Config.GlobalSettings.guest.commands.sharedPolicyRefCount	0	Reference count to enable guest operations.
Config.HostAgent.level[Hbrsvc].logLevel	--	Logging level for the loggers related to HBR services.
Config.HostAgent.level[Hostsvc].logLevel	--	Logging level for the loggers related to Host services.
Config.HostAgent.level[Proxysvc].logLevel	--	Logging level for the loggers related to proxy services.
Config.HostAgent.level[Snmpsvc].logLevel	--	Logging level for the loggers related to SNMP services.
Config.HostAgent.level[Statssvc].logLevel	--	Logging level for the loggers related to Stats services.

3. Click Edit in the Advanced System Settings pane, then use the Filter box to search for **VMkernel.Boot.hyperthreadingMitigation**. Check the Enabled checkbox for the **VMkernel.Boot.hyperthreadingMitigation** system setting and click OK.

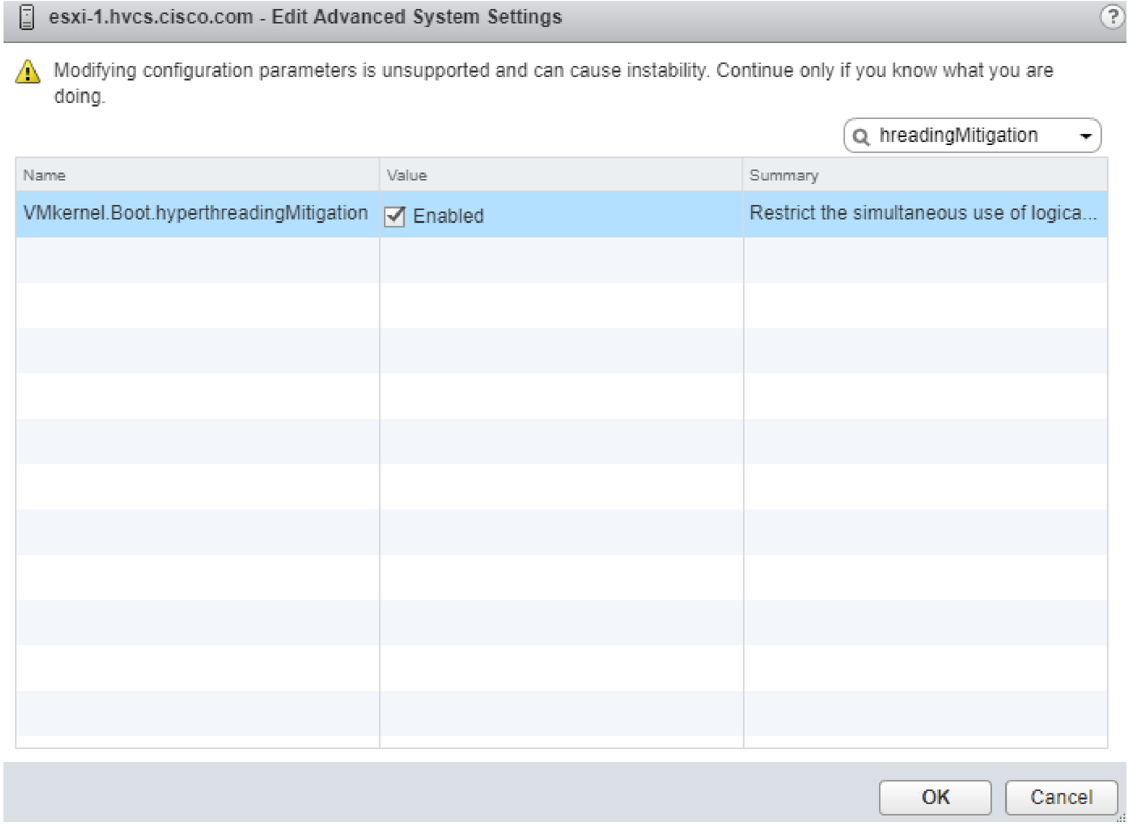

esxi-1.hvcs.cisco.com - Edit Advanced System Settings

Modifying configuration parameters is unsupported and can cause instability. Continue only if you know what you are doing.

hreadingMitigation

Name	Value	Summary
VMkernel.Boot.hyperthreadingMitigation	☑ Enabled	Restrict the simultaneous use of logica...

4. Reboot the host and exit Maintenance Mode. The warning message for CVE-2018-3646 should no longer appear in the host Summary tab.

Enablement of the ESXi Side-Channel-Aware Scheduler (PowerShell HTAware Mitigation Tool Method)

Additional benefits of the HTAware Mitigation Tool are its capabilities to analyze and mitigate hosts in a batch fashion on a per-cluster or per-host basis. This is particularly convenient if you have just deployed multiple hosts that have not been put into production and are not yet servicing workloads.

To use the HTAware Mitigation Tool to remediate an idle cluster of hosts, follow these steps:

1. Place all hosts in the cluster into maintenance mode.

2. If you have not already installed the HTAware Mitigation Tool PowerShell cmdlets, follow the instructions within VMware KB 56931 to setup and import them.

3. Open a Windows PowerShell command window as Administrator and connect to your vCenter server managing the cluster to be remediated with the "**Connect-VIServer**" cmdlet.

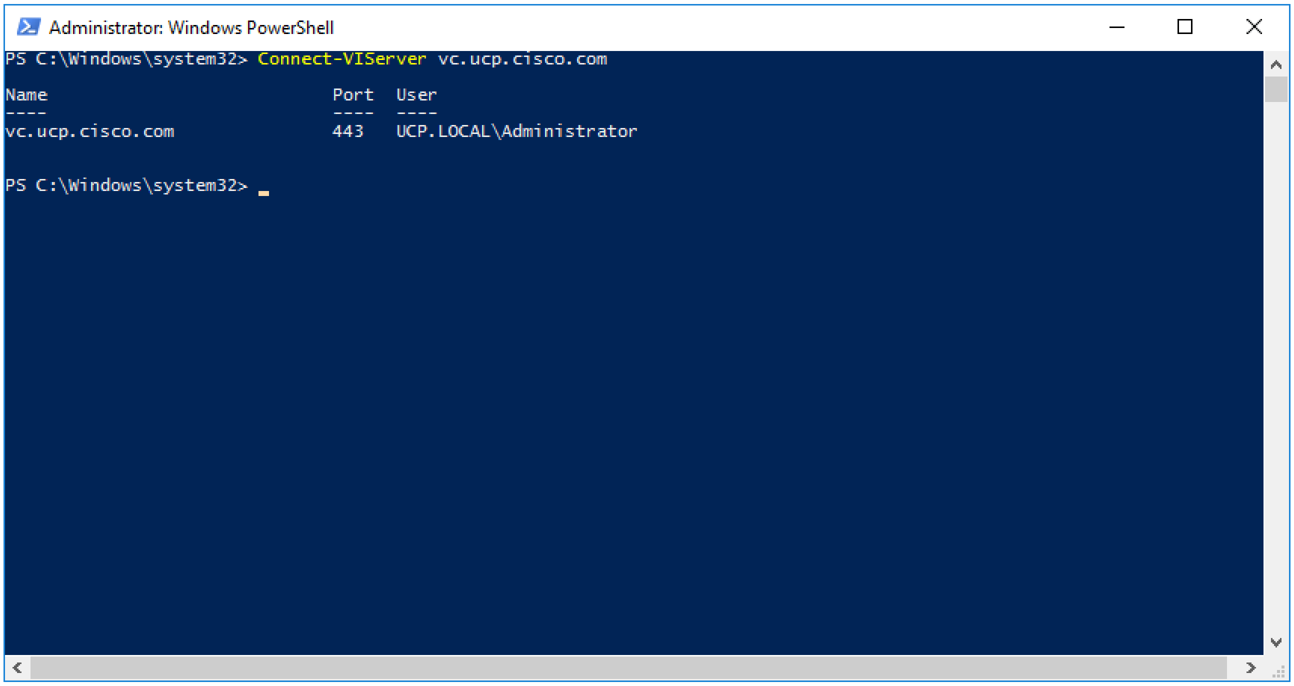

```
PS C:\Windows\system32> Connect-VIServer vc.ucp.cisco.com

Name                           Port  User
----                           ----  ----
vc.ucp.cisco.com               443   UCP.LOCAL\Administrator


PS C:\Windows\system32> _
```

4. Query the remediation status of the hosts in a specific cluster by using the "**Get-HTAwareMitigationConfig -ClusterName <name of cluster>**" cmdlet. **ConfiguredHTAMSetting** and **RuntimeHTAMSetting** should both be false on non-remediated hosts.

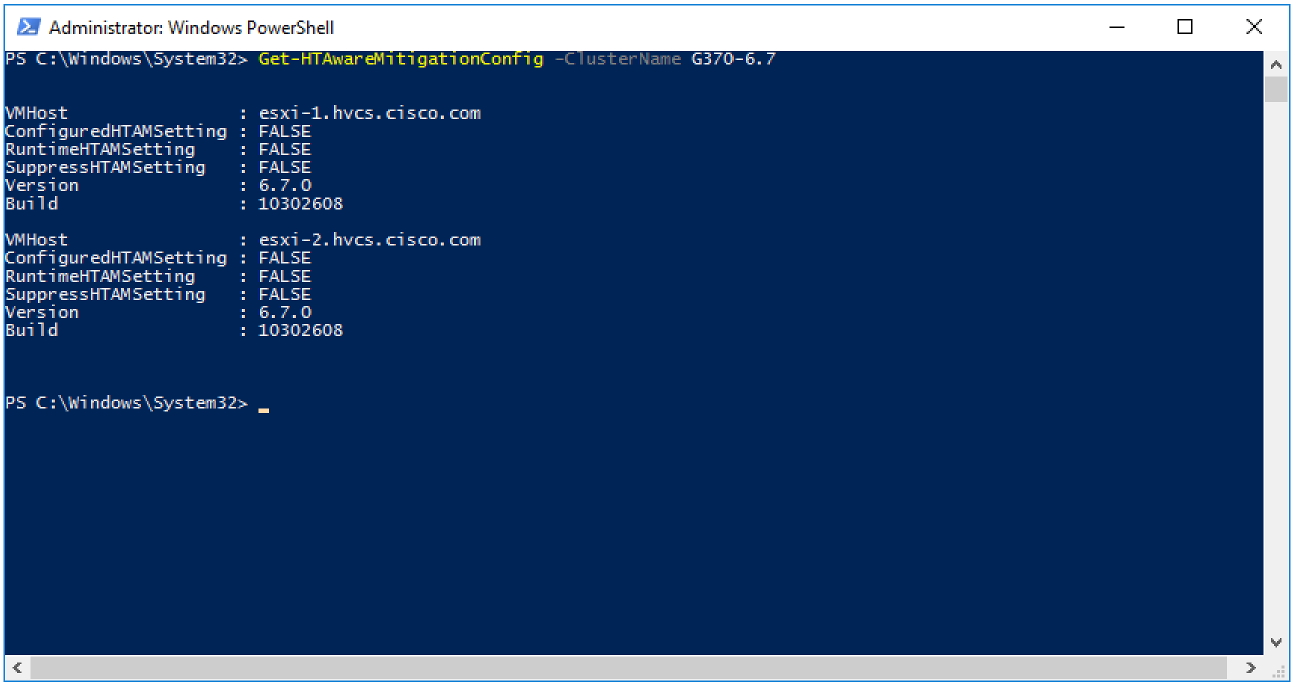

```
PS C:\Windows\System32> Get-HTAwareMitigationConfig -ClusterName G370-6.7


VMHost                : esxi-1.hvcs.cisco.com
ConfiguredHTAMSetting : FALSE
RuntimeHTAMSetting    : FALSE
SuppressHTAMSetting   : FALSE
Version               : 6.7.0
Build                 : 10302608

VMHost                : esxi-2.hvcs.cisco.com
ConfiguredHTAMSetting : FALSE
RuntimeHTAMSetting    : FALSE
SuppressHTAMSetting   : FALSE
Version               : 6.7.0
Build                 : 10302608


PS C:\Windows\System32> _
```

If **ConfiguredHTAMSetting** and **RuntimeHTAMSetting** values are "N/A", then the host is not running a patched version of ESXi that supports remediation. Ensure the host is running a version of ESXi which contains the patches necessary to remediate the L1TF-VMM vulnerability.

5. Enable the ESXi Side-Channel-Aware Scheduler by using the "**Set-HTAwareMitigationConfig -ClusterName <name of cluster> -Enable**" cmdlet.

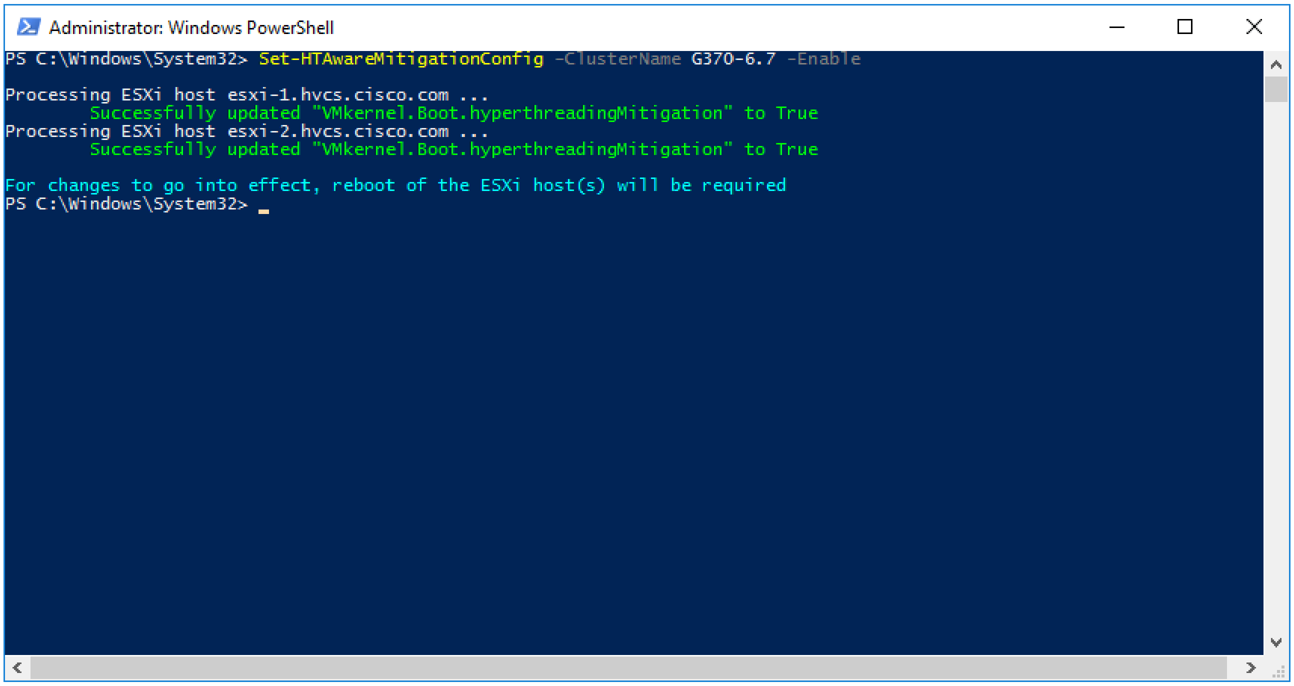

```
PS C:\Windows\System32> Set-HTAwareMitigationConfig -ClusterName G370-6.7 -Enable

Processing ESXi host esxi-1.hvcs.cisco.com ...
        Successfully updated "VMkernel.Boot.hyperthreadingMitigation" to True
Processing ESXi host esxi-2.hvcs.cisco.com ...
        Successfully updated "VMkernel.Boot.hyperthreadingMitigation" to True

For changes to go into effect, reboot of the ESXi host(s) will be required
PS C:\Windows\System32> _
```

6. Reboot the hosts and exit Maintenance Mode. The warning message for CVE-2018-3646 should no longer appear in the host Summary tabs.

Configuration of VMware Round Robin Path Selection Policy IOPS Limit

Hitachi best practices show that performance and total IOPS throughput can be increased by 3-5 percent on Hitachi Virtual Storage Platform by setting the IOPS limit for the VMware Round Robin Path Selection Policy (RR PSP) from the default value of 1,000 to 20. This causes ESXi to switch to the next available path for a LUN after 20 IO instead of after 1,000 IO. This setting is configurable via ESXCLI on a host-by-host basis, or PowerCLI may be used to apply this setting across multiple hosts in a cluster.

Change the Round Robin Path Selection Policy through PowerCLI for Multiple Hosts in a Cluster

This method will allow you to configure the RR PSP IOPS limit through PowerCLI on a per-cluster basis within your environment. Perform the following steps to change the IOPS limit value from 1,000 to 20 on all Hitachi-presented LUNs.

To change the round robin path selection policy through PowerCLI for multiple hosts in a cluster, follow these steps:

1. Open a Windows PowerShell command window as Administrator and connect to your vCenter server managing the cluster to be modified with the "**Connect-VIServer**" cmdlet.

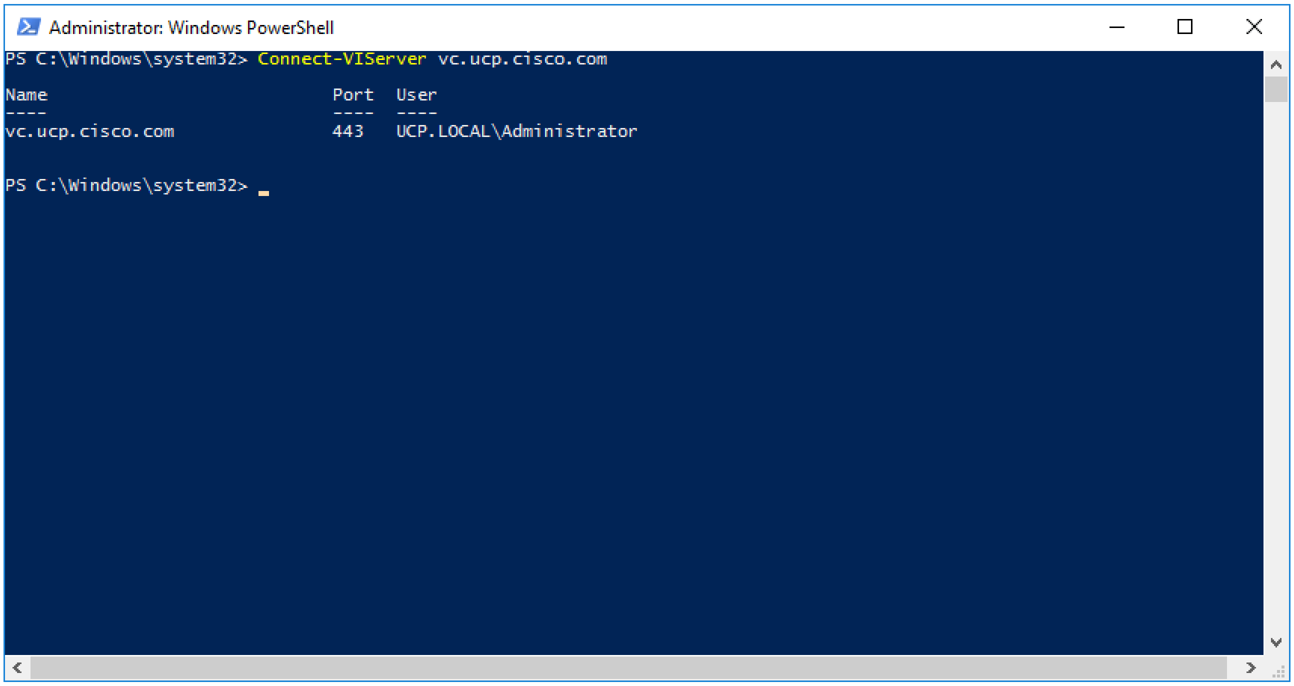

```
PS C:\Windows\system32> Connect-VIServer vc.ucp.cisco.com

Name                           Port  User
----                           ----  ----
vc.ucp.cisco.com               443   UCP.LOCAL\Administrator


PS C:\Windows\system32> _
```

2. Create a variable which will contain all of the ESXi hosts within the cluster you are targeting. In the example shown below, we are creating a variable named "UCSHosts" which contains the host objects within the cluster G1500-6.7 by running the PowerShell command "**$UCSHosts = Get-Cluster "G1500-6.7" | Get-VMHost**". Replace G1500-6.7 with the name of the cluster you are targeting.

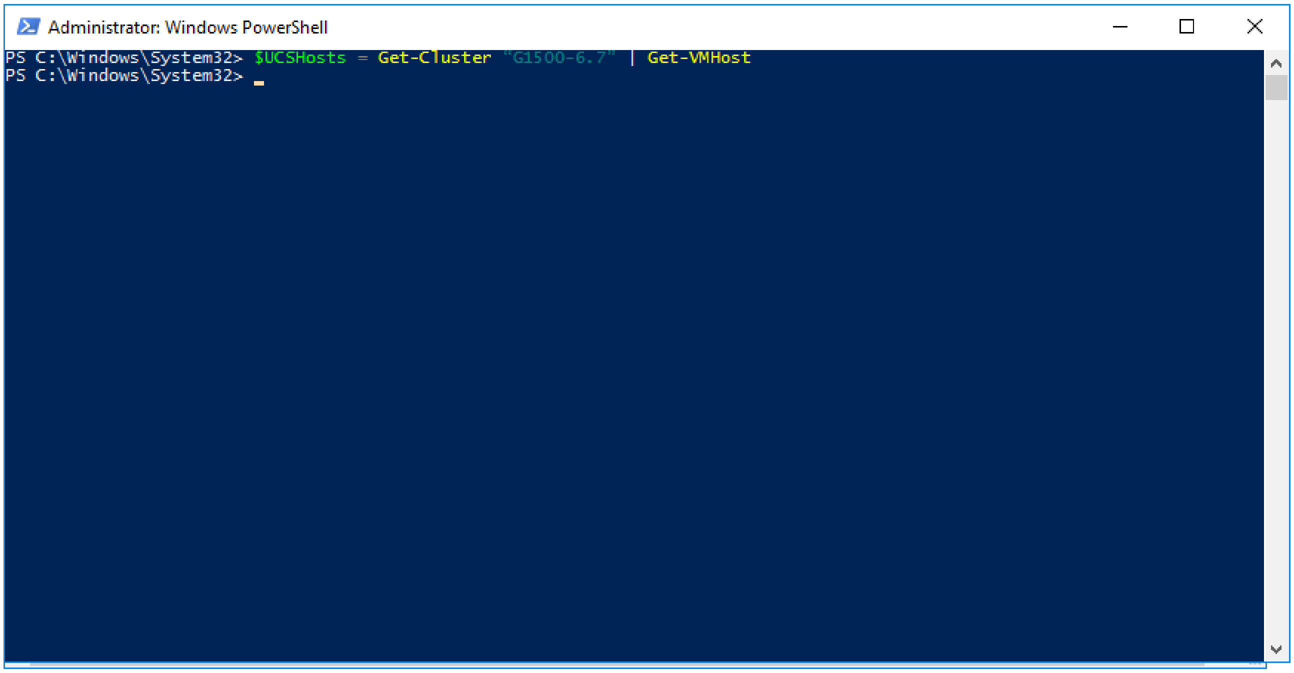

```
PS C:\Windows\System32> $UCSHosts = Get-Cluster "G1500-6.7" | Get-VMHost
PS C:\Windows\System32> _
```

3. To verify that the UCSHosts variable contains the host objects within your cluster, you may issue the command "**echo $UCSHosts**". You should see your individual hosts listed in the PowerShell output similar to what is shown below.

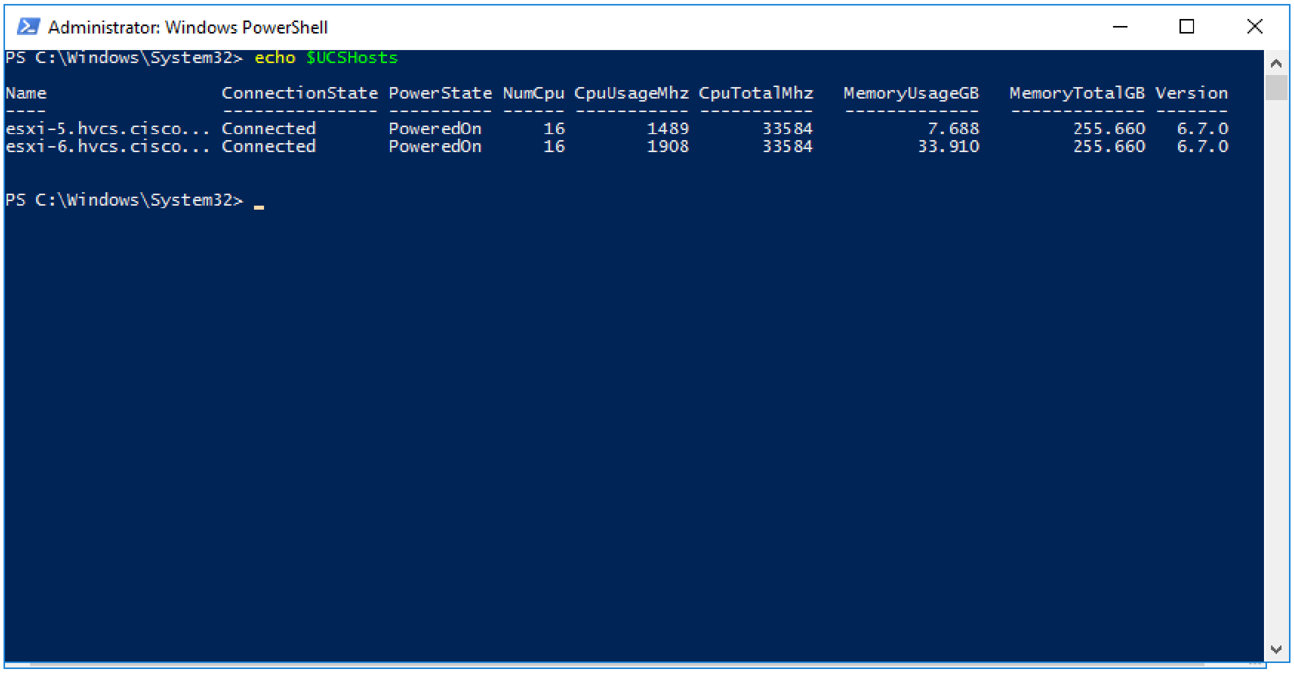

Administrator: Windows PowerShell

```
PS C:\Windows\System32> echo $UCSHosts

Name                 ConnectionState PowerState NumCpu CpuUsageMhz CpuTotalMhz   MemoryUsageGB   MemoryTotalGB Version
----                 --------------- ---------- ------ ----------- -----------   -------------   ------------- -------
esxi-5.hvcs.cisco... Connected       PoweredOn      16        1489       33584           7.688         255.660   6.7.0
esxi-6.hvcs.cisco... Connected       PoweredOn      16        1908       33584          33.910         255.660   6.7.0


PS C:\Windows\System32> _
```

4. To list all Hitachi LUNs presented to the cluster hosts and show the current RR PSP IOPS limit, you may issue the command “foreach ($UCS in $UCSHosts) {Get-VMHost $UCS | Get-ScsiLun -LunType Disk | Where-Object {$_.CanonicalName -like ‘naa.60060e80*’ -and $_.MultiPathPolicy -like ‘RoundRobin’} | Select-Object VMHost, CanonicalName, MultipathPolicy, CommandsToSwitchPath}”. Note that we are filtering on the wildcarded NAA ID “naa.60060e80*” so that only Hitachi-presented LUNs are listed in the output, which should look similar to what is shown below.

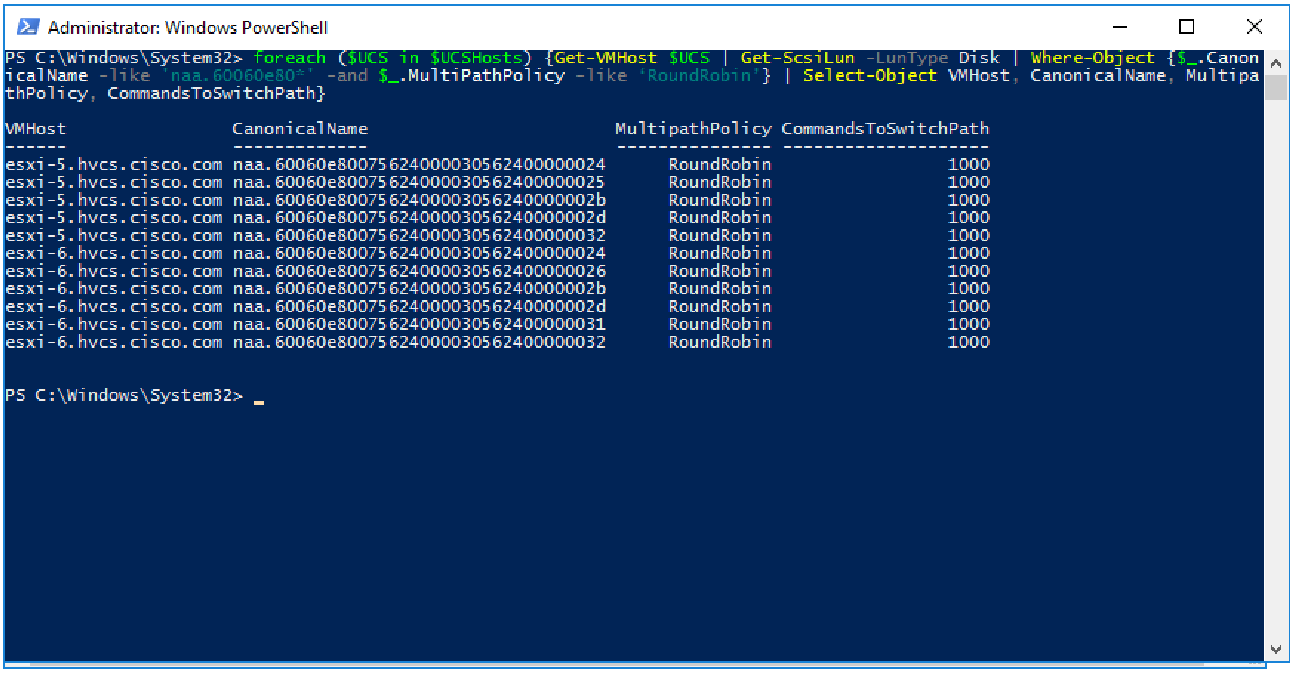

Administrator: Windows PowerShell

```
PS C:\Windows\System32> foreach ($UCS in $UCSHosts) {Get-VMHost $UCS | Get-ScsiLun -LunType Disk | Where-Object {$_.Canon
icalName -like 'naa.60060e80*' -and $_.MultiPathPolicy -like 'RoundRobin'} | Select-Object VMHost, CanonicalName, Multipa
thPolicy, CommandsToSwitchPath}

VMHost                CanonicalName                        MultipathPolicy CommandsToSwitchPath
------                -------------                        --------------- --------------------
esxi-5.hvcs.cisco.com naa.60060e80075624000030562400000024      RoundRobin                 1000
esxi-5.hvcs.cisco.com naa.60060e80075624000030562400000025      RoundRobin                 1000
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000002b      RoundRobin                 1000
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000002d      RoundRobin                 1000
esxi-5.hvcs.cisco.com naa.60060e80075624000030562400000032      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000024      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000026      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000002b      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000002d      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000031      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000032      RoundRobin                 1000


PS C:\Windows\System32> _
```

5. To set the RR PSP IOPS limit to 20 on all Hitachi LUNs presented to the cluster hosts, you may issue the command “foreach ($UCS in $UCSHosts) {Get-VMHost $UCS | Get-ScsiLun -LunType Disk | Where-Object {$_.CanonicalName -like ‘naa.60060e80*’ -and $_.MultipathPolicy -like ‘RoundRobin’} | Set-ScsiLun -CommandsToSwitchPath 20 | Select-

Object VMHost, CanonicalName, CommandsToSwitchPath}". You should see the CommandsToSwitchPath item change to 20 for each Hitachi-presented LUN, similar to the output shown below.

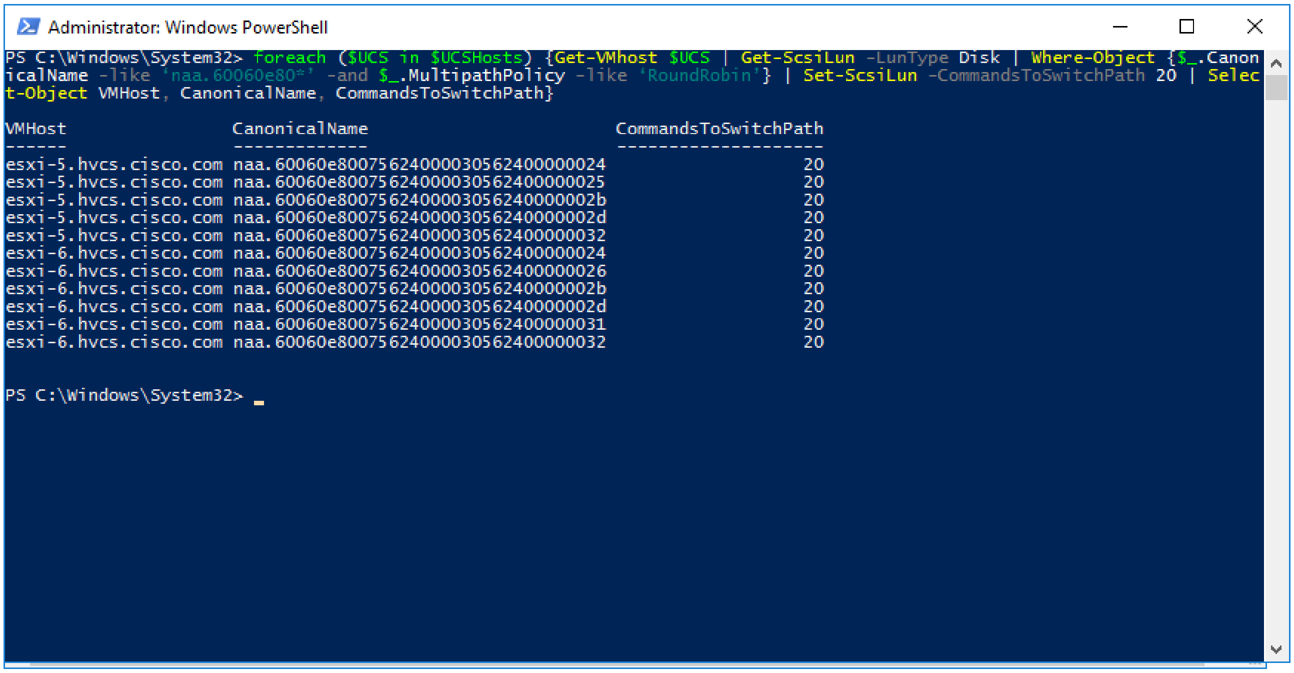

```
PS C:\Windows\System32> foreach ($UCS in $UCSHosts) {Get-VMhost $UCS | Get-ScsiLun -LunType Disk | Where-Object {$_.Canon
icalName -like 'naa.60060e80*' -and $_.MultipathPolicy -like 'RoundRobin'} | Set-ScsiLun -CommandsToSwitchPath 20 | Selec
t-Object VMHost, CanonicalName, CommandsToSwitchPath}

VMHost                CanonicalName                         CommandsToSwitchPath
------                -------------                         --------------------
esxi-5.hvcs.cisco.com naa.60060e80075624000030562400000024                    20
esxi-5.hvcs.cisco.com naa.60060e80075624000030562400000025                    20
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000002b                    20
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000002d                    20
esxi-5.hvcs.cisco.com naa.60060e80075624000030562400000032                    20
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000024                    20
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000026                    20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000002b                    20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000002d                    20
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000031                    20
esxi-6.hvcs.cisco.com naa.60060e80075624000030562400000032                    20


PS C:\Windows\System32> _
```

6. Repeat steps 1-5 for each cluster that you would like to change the VMware Round Robin Path Selection Policy IOPS limit.

Change the Round Robin Path Selection Policy through ESXCLI for a Single Host

This method will allow you to configure the RR PSP IOPS limit through ESXCLI on a per-host basis within your environment. Perform the following steps to change the IOPS limit value from 1,000 to 20 on all Hitachi-presented LUNs.

To change the round robin path selection policy through ESXCLI for a single host, follow these steps:

1. Enable ESXi shell and/or SSH for the host to be configured. Follow the instructions in VMware KB 2004746 if not already enabled in your environment and login to the ESXi shell either locally on the host or through SSH.

2. To list all Hitachi LUNs presented to the cluster hosts and show the current RR PSP IOPS limit, you may issue the command "**esxcli storage nmp device list | grep HITACHI -A4 -B1**". This will show all Hitachi-presented LUNs and the current Path Selection Policy Device Config which includes the IOPS limit, as shown in the example below.

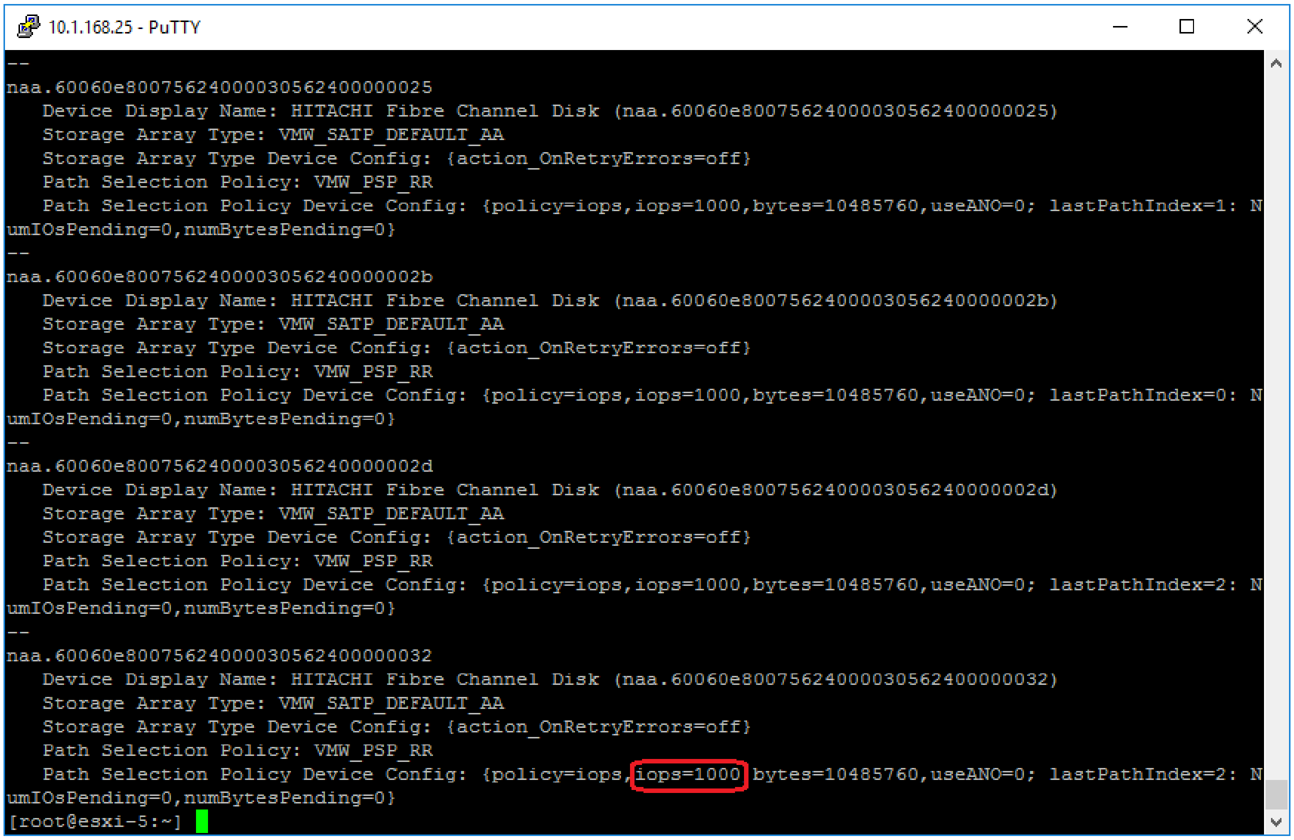

```
--
naa.60060e80075624000030562400000025
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e80075624000030562400000025)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1: N
umIOsPending=0,numBytesPending=0}
--
naa.60060e8007562400003056240000002b
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e8007562400003056240000002b)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=1000,bytes=10485760,useANO=0; lastPathIndex=0: N
umIOsPending=0,numBytesPending=0}
--
naa.60060e8007562400003056240000002d
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e8007562400003056240000002d)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=1000,bytes=10485760,useANO=0; lastPathIndex=2: N
umIOsPending=0,numBytesPending=0}
--
naa.60060e80075624000030562400000032
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e80075624000030562400000032)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=1000,bytes=10485760,useANO=0; lastPathIndex=2: N
umIOsPending=0,numBytesPending=0}
[root@esxi-5:~]
```

3. To set the RR PSP IOPS limit to 20 on all Hitachi LUNs presented to the host, you may issue the command "for i in `esxcfg-scsidevs -c | awk '{print $1}' | grep naa.60060e80`; do esxcli storage nmp psp roundrobin deviceconfig set --type=iops --iops=20 --device=$i; done" as shown in the example shown below.

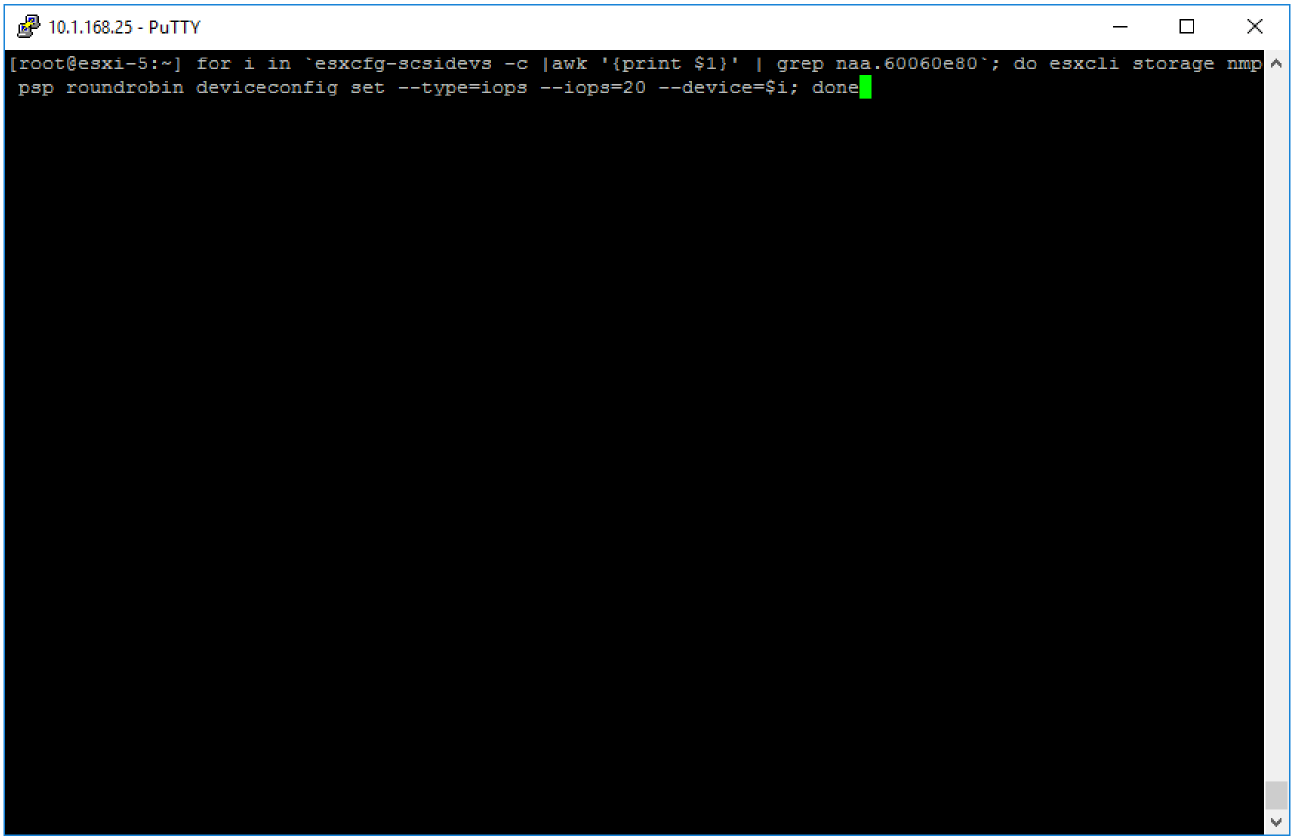

```
[root@esxi-5:~] for i in `esxcfg-scsidevs -c |awk '{print $1}' | grep naa.60060e80`; do esxcli storage nmp
 psp roundrobin deviceconfig set --type=iops --iops=20 --device=$i; done
```

4. Run the command "**esxcli storage nmp device list | grep HITACHI -A4 -B1**" and ensure the IOPS limit has changed to 20 for Hitachi-presented LUNs as shown in the example below.

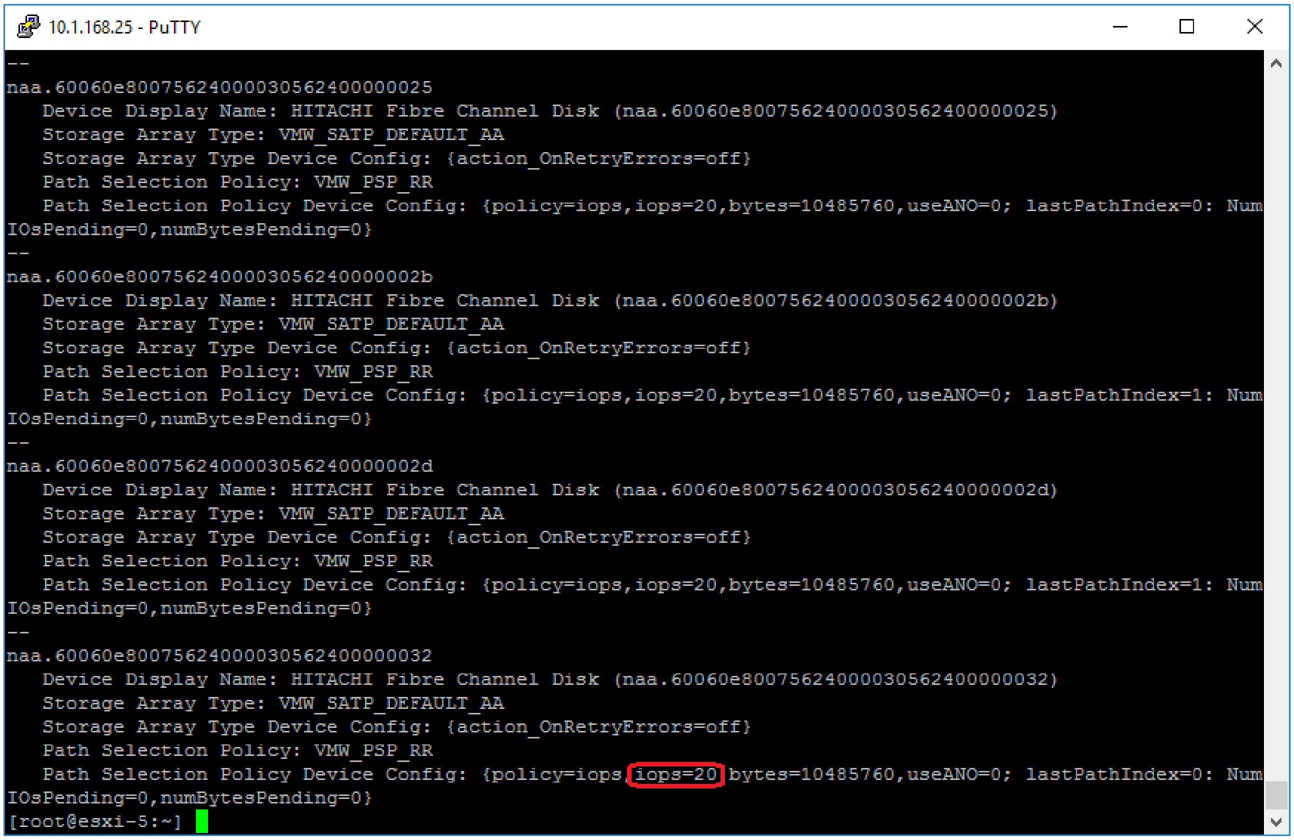

```
--
naa.60060e80075624000030562400000025
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e80075624000030562400000025)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=20,bytes=10485760,useANO=0; lastPathIndex=0: Num
IOsPending=0,numBytesPending=0}
--
naa.60060e8007562400003056240000002b
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e8007562400003056240000002b)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=20,bytes=10485760,useANO=0; lastPathIndex=1: Num
IOsPending=0,numBytesPending=0}
--
naa.60060e8007562400003056240000002d
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e8007562400003056240000002d)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=20,bytes=10485760,useANO=0; lastPathIndex=1: Num
IOsPending=0,numBytesPending=0}
--
naa.60060e80075624000030562400000032
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e80075624000030562400000032)
   Storage Array Type: VMW_SATP_DEFAULT_AA
   Storage Array Type Device Config: {action_OnRetryErrors=off}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=iops,iops=20,bytes=10485760,useANO=0; lastPathIndex=0: Num
IOsPending=0,numBytesPending=0}
[root@esxi-5:~]
```

5. Repeat steps 1-4 on each host that you would like to change the VMware Round Robin Path Selection Policy IOPS limit.

Cisco Intersight Registration

Cisco Intersight gives manageability and visibility to multiple UCS domains through a common interface, regardless of location. The Base addition is available for UCSM starting at release 3.2(1) at no additional cost.

To add the Cisco UCS Fabric Interconnects into Intersight, follow these steps:

1. Connect to https://www.intersight.com.

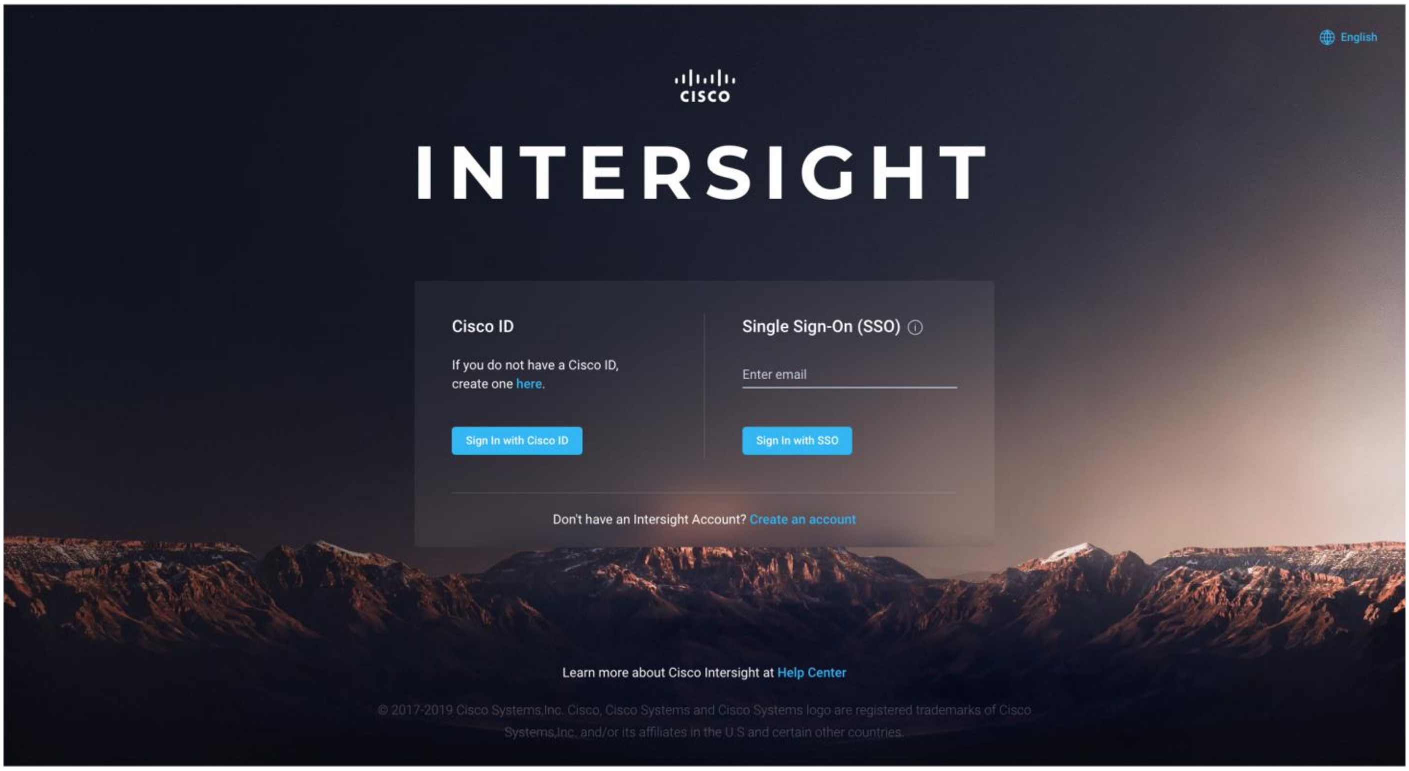

If you do not have a Cisco ID, create one here.

Don't have an Intersight Account? Create an account

Learn more about Cisco Intersight at Help Center

© 2017-2019 Cisco Systems,Inc. Cisco, Cisco Systems and Cisco Systems logo are registered trademarks of Cisco Systems,Inc. and/or its affiliates in the U.S and certain other countries.

Prerequisites

The following prerequisites are necessary to setup access to Cisco Intersight:

1. An account on **cisco.com**.

2. A valid Cisco Intersight account. This can be created by navigating to https://intersight.com and following the instructions for creating an account. The account creation requires at least one device to be registered in Intersight and requires **Device ID** and **Claim ID** information from the device. See Collecting Information From Cisco UCS Domain for an example of how to get **Device ID** and **Claim ID** from Cisco UCS Fabric Interconnect devices.

3. Valid License on Cisco Intersight – see the Cisco Intersight Licensing section for more information.

4. Cisco UCS Fabric Interconnects must be able to do a DNS lookup to access Cisco Intersight.

5. Device Connectors on Fabric Interconnects must be able to resolve *svc.ucs-connect.com*.

6. Allow outbound HTTPS connections (port 443) initiated from the Device Connectors on Fabric Interconnects to Cisco Intersight. HTTP Proxy is supported.

Setup Information

To setup access to Cisco Intersight, the following information must be collected from the Cisco UCS Domain. The deployment steps provided below will show how to collect this information.

- Device ID
- Claim Code

Cisco Intersight Licensing

Cisco Intersight is offered in two editions:

- Base license which is free to use, and offers a large variety of monitoring, inventory and reporting features.
- Essentials license, at an added cost but provides advanced monitoring, server policy and profile configuration, firmware management, virtual KVM features, and more. A 90-day trial of the Essentials license is available for use as an evaluation period.

New features and capabilities will be added to the different licensing tiers in future release.

Deployment Steps

To setup access to Cisco Intersight from a Cisco UCS domain, complete the steps outlined in this section.

Connect to Cisco Intersight

To connect and access Cisco Intersight, follow these steps:

1. Use a web browser to navigate to Cisco Intersight at https://intersight.com/.

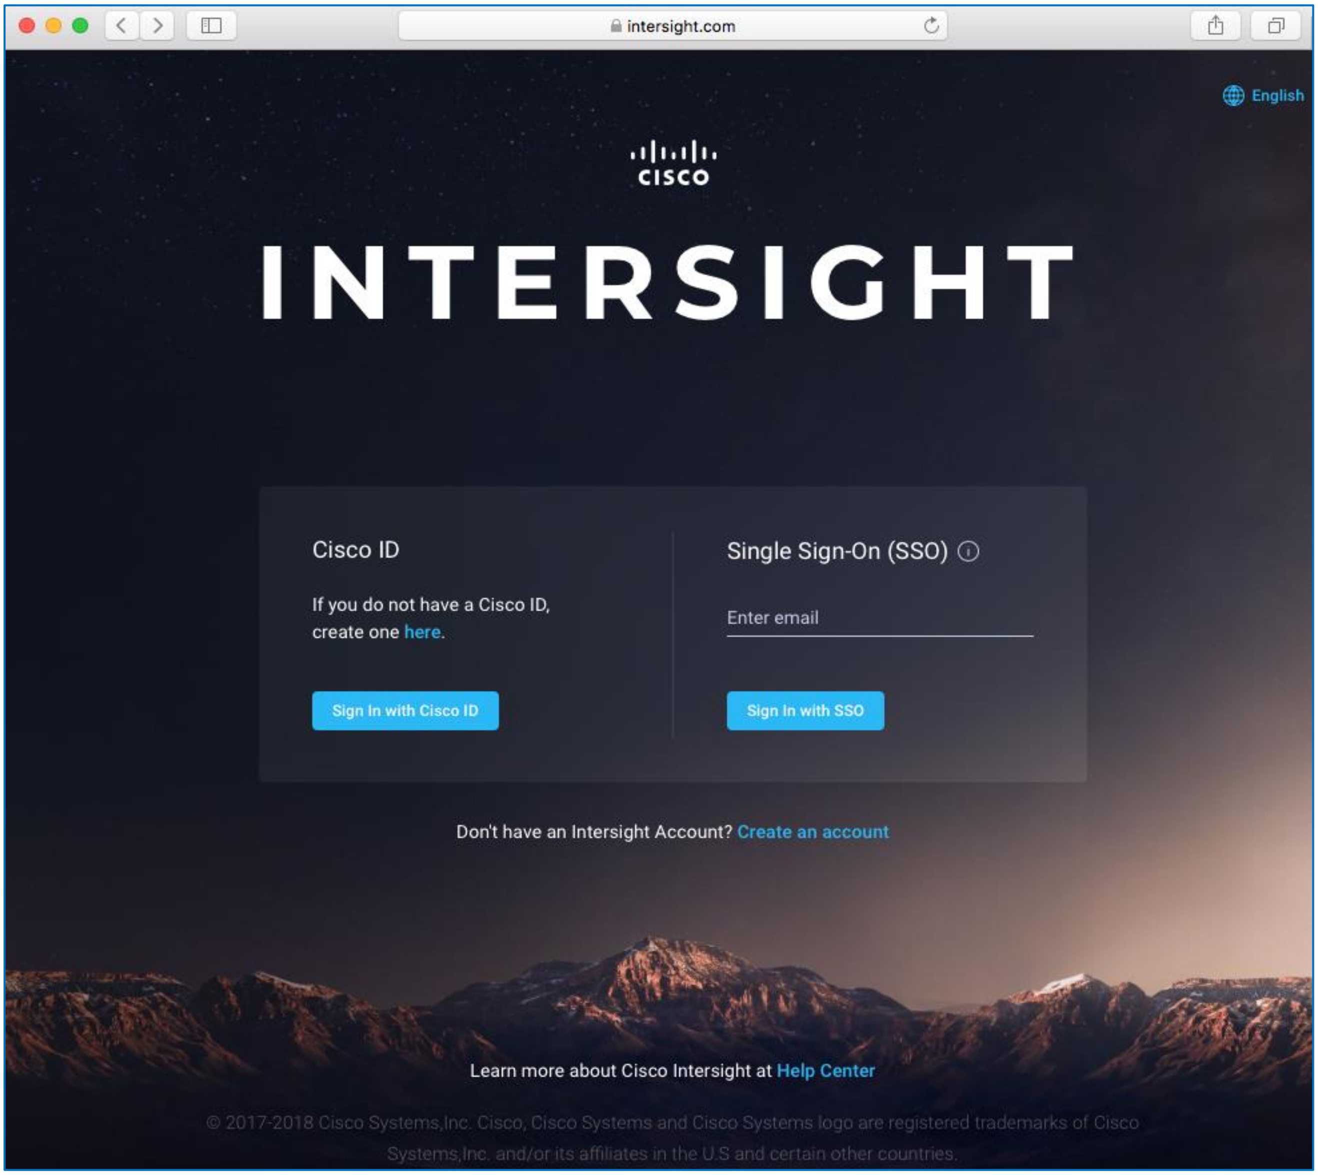

If you do not have a Cisco ID, create one here.

Don't have an Intersight Account? Create an account

Learn more about Cisco Intersight at Help Center

© 2017-2018 Cisco Systems,Inc. Cisco, Cisco Systems and Cisco Systems logo are registered trademarks of Cisco Systems,Inc. and/or its affiliates in the U.S and certain other countries.

2. Login with a valid cisco.com account or single sign-on using your corporate authentication.

Collect Information from UCS Domain

To collect information from Cisco UCS Fabric Interconnects to setup access to Cisco Intersight, follow these steps:

1. Use a web browser to navigate to the UCS Manager GUI. Login using the **admin** account.

2. From the navigation menu, select the **Admin** icon.

3. Select **All > Device Connector**.

4. From the **Intersight Management** pane, click **Enabled** to enable Intersight management.

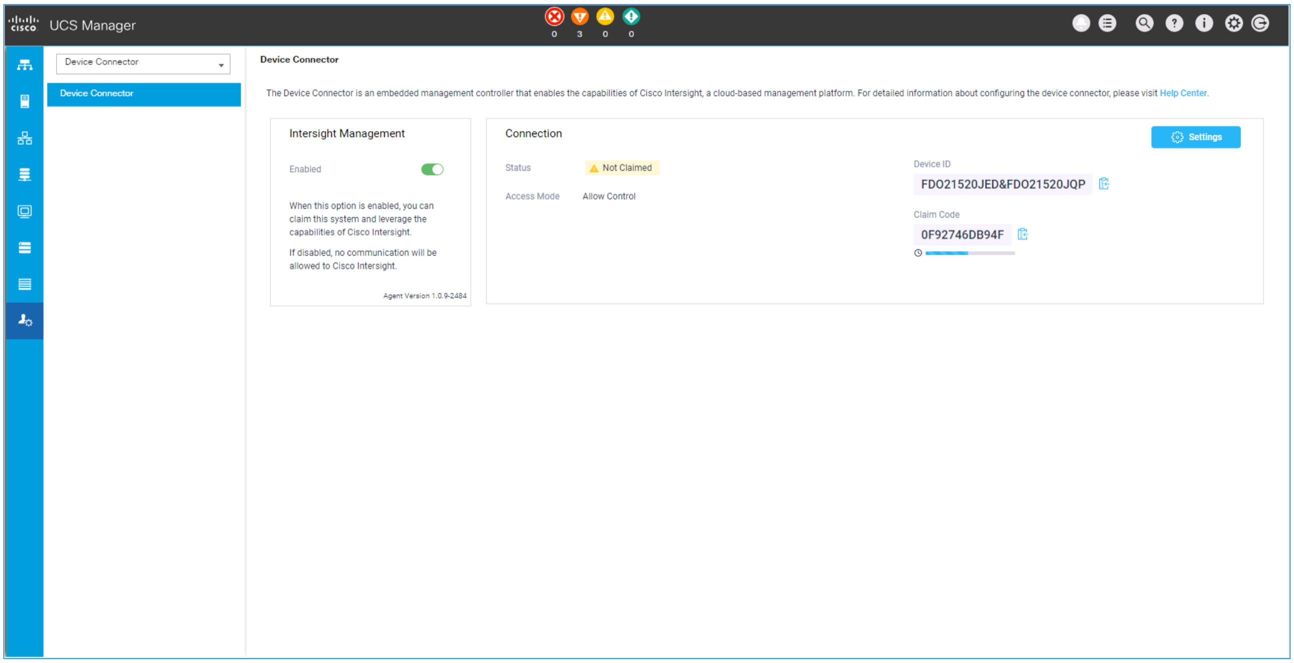

Device Connector

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center.

Intersight Management

Enabled

When this option is enabled, you can claim this system and leverage the capabilities of Cisco Intersight.

If disabled, no communication will be allowed to Cisco Intersight.

Agent Version 1.0.9-2484

Connection

Status: Not Claimed

Access Mode: Allow Control

Device ID: FDO21520JED&FDO21520JQP

Claim Code: 0F92746DB94F

5. From the **Connection** pane, copy the **Device ID** and **Claim ID** information. This information will be required to add this device to Cisco Intersight.

6. (Optional) Click Settings to change **Access Mode** and to configure **HTTPS Proxy**.

Add Cisco UCS Domain to Cisco Intersight

To add Cisco UCS Fabric Interconnects to Cisco Intersight to manage the UCS domain, follow these steps:

1. From Cisco Intersight, in the navigation menu, select **Devices**.

2. Click **Claim a New Device**.

3. In the **Claim a New Device** pop-up window, paste the **Device ID** and **Claim Code** collected in the previous section.

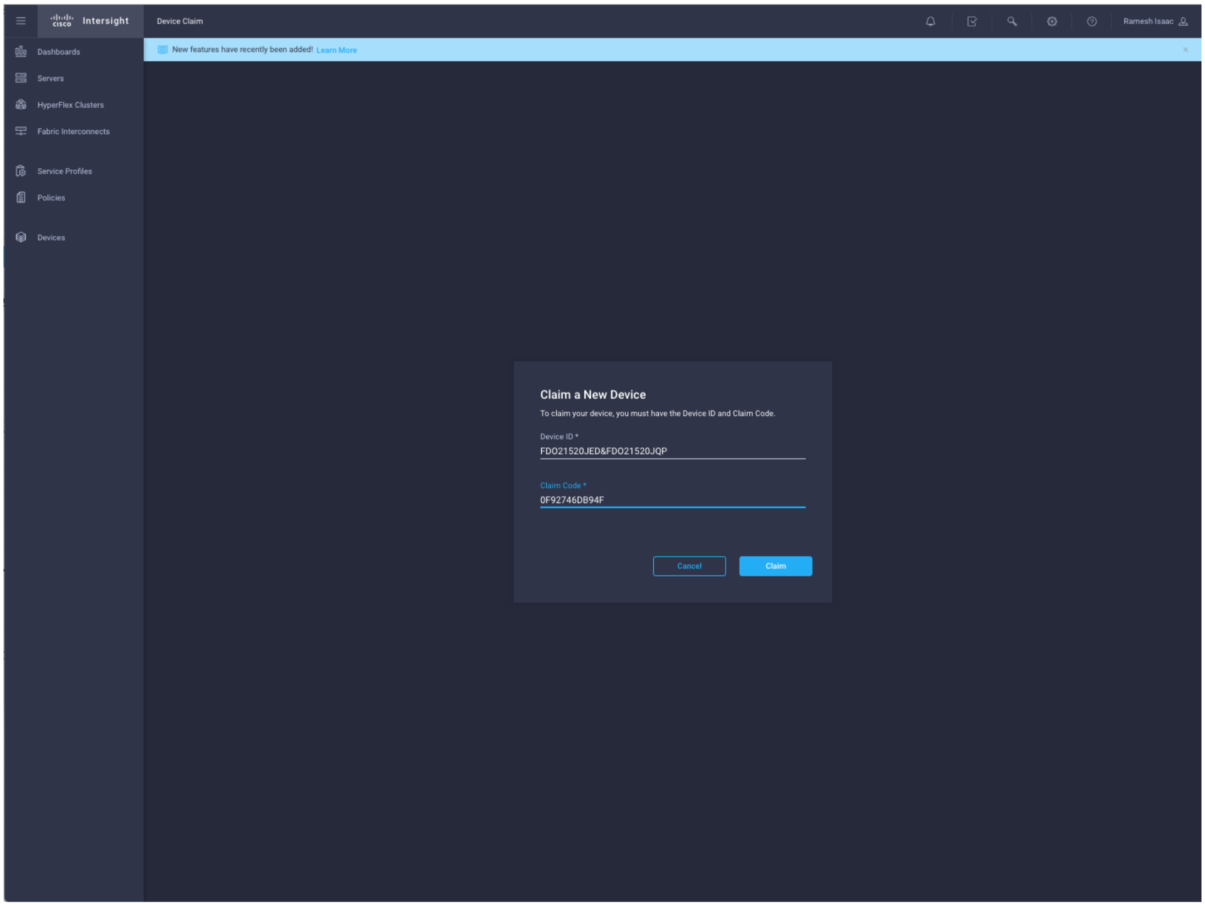

Claim a New Device

To claim your device, you must have the Device ID and Claim Code.

Device ID *
FDO21520JED&FDO21520JQP

Claim Code *
0F92746DB94F

4. Click **Claim**.

On Cisco Intersight, the newly added UCS domain should now have a **Status** of **Connected**.

On Cisco UCS Manager, the **Device Connector** should now have a **Status** of **Claimed**.

The Dashboard will present an overview of the managed UCS domains:

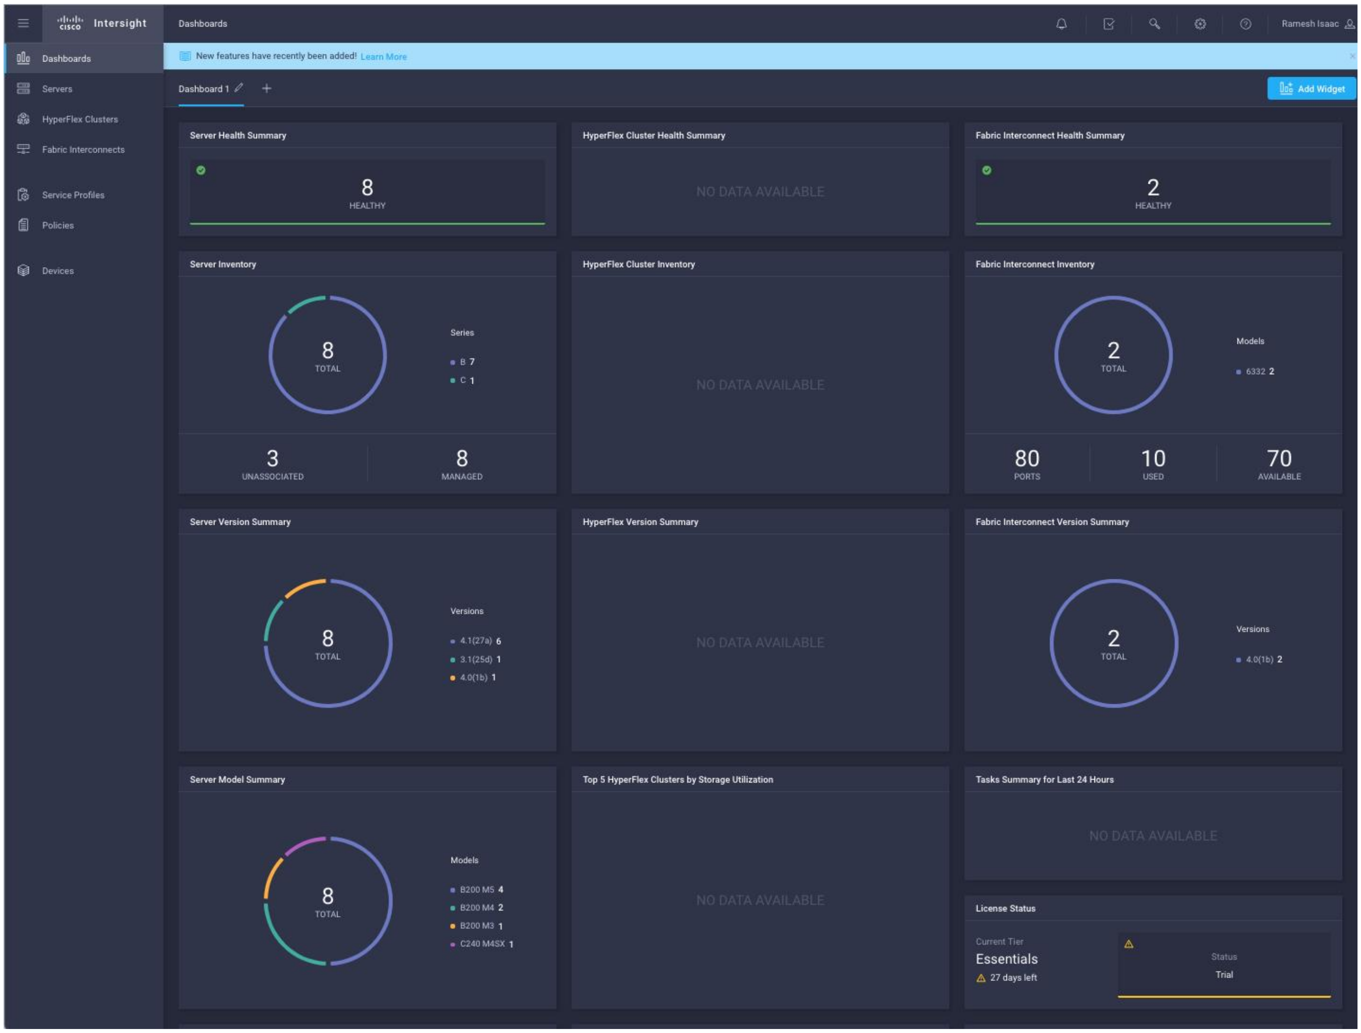

Server Health Summary
8 HEALTHY

HyperFlex Cluster Health Summary
NO DATA AVAILABLE

Fabric Interconnect Health Summary
2 HEALTHY

Server Inventory
8 TOTAL
Series
B 7
C 1
3 UNASSOCIATED
8 MANAGED

HyperFlex Cluster Inventory
NO DATA AVAILABLE

Fabric Interconnect Inventory
2 TOTAL
Models
6332 2
80 PORTS
10 USED
70 AVAILABLE

Server Version Summary
8 TOTAL
Versions
4.1(27a) 6
3.1(25d) 1
4.0(1b) 1

HyperFlex Version Summary
NO DATA AVAILABLE

Fabric Interconnect Version Summary
2 TOTAL
Versions
4.0(1b) 2

Server Model Summary
8 TOTAL
Models
B200 M5 4
B200 M4 2
B200 M3 1
C240 M4SX 1

Top 5 HyperFlex Clusters by Storage Utilization
NO DATA AVAILABLE

Tasks Summary for Last 24 Hours
NO DATA AVAILABLE

License Status
Current Tier
Essentials
⚠ 27 days left
⚠
Status
Trial

About the Authors

Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.

Ramesh Isaac is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Group. Ramesh has worked in the data center and mixed-use lab settings since 1995. He started in information technology supporting UNIX environments and focused on designing and implementing multi-tenant virtualization solutions in Cisco labs before entering Technical Marketing where he has supported converged infrastructure and virtual services as part of solution offerings as Cisco. Ramesh has certifications from Cisco, VMware, and Red Hat.

Tim Darnell, Master Solutions Architect and Product Owner, Hitachi Vantara

Tim Darnell is a Master Solutions Architect and Product Owner in the Hitachi Vantara Converged Product Engineering Group. Tim has worked on data center and virtualization technologies since 1997. He started his career in systems administration and has worked in a multitude of roles since, from technical reference authoring to consulting in large, multi-national corporations as a technical advisor. He is currently a Product Owner at Hitachi Vantara, responsible for the Unified Compute Platform Converged Infrastructure line of products that focus on VMware vSphere product line integrations. Tim holds multiple VCAP and VCP certifications from VMware and is a RedHat Certified Engineer.

Appendix: References

Cisco

Nexus vPC Best Practices: https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf

Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide, Release 7.x: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/interfaces/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Interfaces_Configuration_Guide_7x.html

Cisco UCS Best Practices: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-manager/whitepaper_c11-697337.html

Cisco UCS Performance and Tuning: https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/whitepaper_c11-740098.pdf

Cisco UCS 6454 Spec Sheet https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/ucs-6454-fab-int-specsheet.pdf

Cisco UCS 6300 Spec Sheet https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/6332-specsheet.pdf

Hitachi

Hitachi Provisioning Guide for VSP G130, G/F350, G/F370, G/F700, G/F900: https://knowledge.hitachivantara.com/@api/deki/files/55795/SVOS_RF_v8_3_1_Provisioning_Guide_VSP_Gx00_Fx00_MK-97HM85026-03.pdf?revision=1

Hitachi Provisioning Guide for Open Systems for VSP G1x00 and F1500: https://knowledge.hitachivantara.com/@api/deki/files/50437/SVOS_RF_v8_3_Provisioning_Guide_VSP_G1x00_F1500_MK-92RD8014-20.pdf?revision=1

Appendix: Bill of Materials

Bill of Materials

For each design tested in this solution, a bill of materials (BOM) was generated. Please note that the following are not included in the BOMs below and will need to be identified separately depending on your specific configuration:

- Racks for both Cisco and Hitachi components
- Power distribution units (PDUs)
- Multi-mode Fibre (MMF) cabling between Cisco Fabric Interconnects and Hitachi VSP storage systems
- Power cables and rail kits for Hitachi VSP storage systems
- Services, Maintenance, and Support plans for each component

The BOMs below are representative of the equipment used in Cisco Systems lab environments to certify each design. Components, interconnect cabling, and quantities may differ depending on your specific configuration needs. It is important to note that any component changes must be referenced against both Cisco and Hitachi compatibility matrices to ensure proper support is available.

Table 17 lists the BOM for Cisco UCS 6454 Fabric Interconnect with Hitachi VSP G370 design.

Table 17 Bill of Materials for Direct Attached UCS 6454 and Hitachi VSP G370 Design

Vendor	Part Number/Order Code	Description	Quantity
Cisco	N9K-C9336C-FX2	Nexus 9300 Series, 36p 40/100G QSFP28	2
Cisco	NXOS-9.2.2	Nexus 9500, 9300, 3000 Base NX-OS Software Rel 9.2.2	2
Cisco	N3K-C3064-ACC-KIT	Nexus 3K/9K Fixed Accessory Kit	2
Cisco	NXA-PAC-1100W-PE2	Nexus AC 1100W PSU - Port Side Exhaust	4
Cisco	NXA-FAN-65CFM-PE	Nexus Fan, 65CFM, port side exhaust airflow	6
Cisco	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	4
Cisco	QSFP-100G-AOC1M	100GBASE QSFP Active Optical Cable, 1m	4
Cisco	QSFP-40G-SR-BD	QSFP40G BiDi Short-reach Transceiver	4
Cisco	QSFP-H40G-CU1M	40GBASE-CR4 Passive Copper Cable, 1m	2
Cisco	UCS-FI-6454-U	UCS Fabric Interconnect 6454	2
Cisco	N10-MGT016	UCS Manager v4.0	2
Cisco	UCS-ACC-6332	UCS 6332/ 6454 Chassis Accessory Kit	2
Cisco	UCS-FAN-6332	UCS 6332/ 6454 Fan Module	8
Cisco	UCS-PSU-6332-AC	UCS 6332 Power Supply/100-240VAC	4

Vendor	Part Number/Order Code	Description	Quantity
Cisco	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	4
Cisco	SFP-H10GB-CU2-5M	10GBASE-CU SFP+ Cable 2.5 Meter	8
Cisco	UCSB-5108-AC2-UPG	UCS 5108 Blade Server AC2 Chassis/0 PSU/8 fans/0 FEX	1
Cisco	N20-FW016	UCS 5108 Blade Chassis FW Package 4.0	1
Cisco	N20-FAN5	Fan module for UCS 5108	8
Cisco	N01-UAC1	Single phase AC power module for UCS 5108	1
Cisco	N20-CBLKB1	Blade slot blanking panel for UCS 5108/single slot	4
Cisco	N20-CAK	Accessory kit for UCS 5108 Blade Server Chassis	1
Cisco	UCSB-B200-M5	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz	4
Cisco	UCS-CPU-6140	2.3 GHz 6140/140W 18C/24.75MB Cache/DDR4 2666MHz	8
Cisco	UCSB-MLOM-40G-04	Cisco UCS VIC 1440 modular LOM for Blade Servers	4
Cisco	UCS-SID-INFR-OI	Other Infrastructure	4
Cisco	UCSB-HS-M5-R	CPU Heat Sink for UCS B-Series M5 CPU socket (Rear)	4
Cisco	UCSB-LSTOR-BK	FlexStorage blanking panels w/o controller, w/o drive bays	8
Cisco	UCSB-HS-M5-F	CPU Heat Sink for UCS B-Series M5 CPU socket (Front)	4
Cisco	UCS-SID-WKL-OW	Other Workload	4
Cisco	UCS-IOM-2208XP	UCS 2208XP I/O Module (8 External, 32 Internal 10Gb Ports)	2
Cisco	UCSB-PSU-2500ACDV	2500W Platinum AC Hot Plug Power Supply - DV	4
Cisco	UCSB-5108-PKG-HW	UCS 5108 Packaging for chassis with half width blades.	1
Cisco	CAB-US515P-C19-US	NEMA 5-15 to IEC-C19 13ft US	4
Cisco	UCS-MR-X32G2RS-H	32GB DDR4-2666-MHz RDIMM/PC4-21300/dual rank/x4/1.2v	48
Cisco	UCS-DIMM-BLK	UCS DIMM Blanks	48
Cisco	DS-SFP-FC32G-SW	32 Gbps Fibre Channel SW SFP+, LC	8
Hitachi	VSP-G-SOLUTION.S	VSP G Unified Platform	1
Hitachi	VSP-G370-A0008.S	VSP G370 Covered Product Unified (FC/iSCSI)	1
Hitachi	G370-F-BASE-S.P	VSP G370 Foundation Base Package	1
Hitachi	GXX0-4X1R9TB.P	VSP GXX0 Flash Pack 4 x 1.9TB SSD Package	12
Hitachi	VSP-G370-A0001.S	VSP G370 Product Unified (FC/iSCSI)	1

Vendor	Part Number/Order Code	Description	Quantity
Hitachi	FD221577-001.P	SVP Bezel ASM (including brackets)	1
Hitachi	HDW2-F850-1PS32.P	VSP G SFP for 32Gbps Shortwave	16
Hitachi	HDW2-F850-DBSC.P	VSP G/F XX0 Drive Box (SFF)	1
Hitachi	HDW-F850-SCQ1.P	VSP G SAS Cable 1m	2
Hitachi	HDW2-F850-SVP.P	VSP G/FXX0 SVP - Service Processor	1
Hitachi	HDW2-F850-4HF32R.P	VSP G/FXX0 Host I/O Module FC 16/32G 4port	4

Appendix: Nexus A Configuration Example

```
version 7.0(3)I7(5a) Bios:version 05.31
switchname AA19-9336-1
vdc AA19-9336-1 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature vpc

username admin password 5 $5$wsy2Bp4V$stK.pozTENuOUwnW8Y0/TMGz/CauQYUfwlxBR2EugI7  role network-admin
ip domain-lookup
system default switchport
copp profile strict
snmp-server user admin network-admin auth md5 0xba69923b15f9f03d162b30bb91e7785b priv
0xba69923b15f9f03d162b30bb91e7785b localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 192.168.168.254 use-vrf management
ntp source 10.1.168.1
ntp master 3

ip route 0.0.0.0/0 10.1.168.254
vlan 1-2,119,201-203,1000
vlan 2
  name Native
vlan 119
  name IB-MGMT
vlan 201
  name Web
vlan 202
  name App
vlan 203
  name DB
vlan 1000
  name vMotion

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
vrf context management
  ip route 0.0.0.0/0 192.168.168.254
port-channel load-balance src-dst l4port
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 192.168.168.14 source 192.168.168.13
  delay restore 150
  peer-gateway
  auto-recovery
  ip arp synchronize


interface Vlan1
interface Vlan119
  no shutdown
  no ip redirects
  ip address 10.1.168.2/24
  no ipv6 redirects
  hsrp 19
    preempt
    ip 10.1.168.1

interface Vlan201
  no shutdown
  no ip redirects
  ip address 172.18.101.252/24
  no ipv6 redirects
  hsrp 101
    preempt
    priority 105
    ip 172.18.101.254

interface Vlan202
  no shutdown
  no ip redirects
  ip address 172.18.102.252/24
  no ipv6 redirects
  hsrp 102
    preempt
    ip 172.18.102.254

interface Vlan203
  no shutdown
  no ip redirects
  ip address 172.18.103.252/24
  no ipv6 redirects
  hsrp 103
    preempt
    priority 105
    ip 172.18.103.254

interface port-channel11
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  spanning-tree port type network
  vpc peer-link

interface port-channel15
  description vPC UCS 6454-1 FI
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  spanning-tree port type edge trunk
  mtu 9216
  load-interval counter 3 60
  vpc 15

interface port-channel16
  description vPC UCS 6454-2 FI
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  spanning-tree port type edge trunk
  mtu 9216
  load-interval counter 3 60
  vpc 16

interface port-channel135
  description vPC Upstream Network Switch A

interface port-channel136
  description vPC Upstream Network Switch B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119
  vpc 136

interface Ethernet1/1
  description vPC peer-link connection to AA19-9336-2 Ethernet1/1
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  channel-group 11 mode active

interface Ethernet1/2
  description vPC peer-link connection to AA19-9336-2 Ethernet1/2
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  channel-group 11 mode active

interface Ethernet1/5
  description vPC 15 connection to UCS 6454-1 FI Ethernet1/53
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  mtu 9216
  load-interval counter 3 60
  channel-group 15 mode active

interface Ethernet1/6
  description vPC 16 connection to UCS 6454-2 FI Ethernet1/53
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  mtu 9216
  load-interval counter 3 60
  channel-group 16 mode active

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23
interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35
  description vPC 135 connection to Upstream Network Switch A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119
  channel-group 135 mode active

interface Ethernet1/36
  description vPC 136 connection to Upstream Network Switch B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119
  channel-group 136 mode active

interface mgmt0
  vrf member management
  ip address 192.168.168.13/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.5a.bin
no system default switchport shutdown
```

Appendix: Nexus B Configuration Example

```
version 7.0(3)I7(5a) Bios:version 05.31
switchname AA19-9336-2
vdc AA19-9336-2 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature vpc

username admin password 5 $5$c0gohGBw$09At8vxbCEsH8R6nXJhJe0AAE83XfK1rQHZ9/Stg6x1  role network-admin
ip domain-lookup
system default switchport
copp profile strict
snmp-server user admin network-admin auth md5 0x3aca90a8ed874105ac3e972e2b7d68fe priv
0x3aca90a8ed874105ac3e972e2b7d68fe localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 192.168.168.254 use-vrf management
ntp source 10.1.168.1
ntp master 3

ip route 0.0.0.0/0 10.1.168.254
vlan 1-2,119,201-203,1000
vlan 2
  name Native
vlan 119
  name IB-MGMT
vlan 201
  name Web
vlan 202
  name App
vlan 203
  name DB
vlan 1000
  name vMotion

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
vrf context management
  ip route 0.0.0.0/0 192.168.168.254
port-channel load-balance src-dst l4port
vpc domain 10
  peer-switch
  role priority 20
  peer-keepalive destination 192.168.168.13 source 192.168.168.14
  delay restore 150
  peer-gateway
  auto-recovery
  ip arp synchronize


interface Vlan1
interface Vlan119
  no shutdown
  no ip redirects
  ip address 10.1.168.3/24
  no ipv6 redirects
  hsrp 19
    preempt
    priority 105
    ip 10.1.168.1

interface Vlan201
  no shutdown
  no ip redirects
  ip address 172.18.101.253/24
  no ipv6 redirects
  hsrp 101
    preempt
    ip 172.18.101.254

interface Vlan202
  no shutdown
  no ip redirects
  ip address 172.18.102.253/24
  no ipv6 redirects
  hsrp 102
    preempt
    priority 105
    ip 172.18.102.254

interface Vlan203
  no shutdown
  no ip redirects
  ip address 172.18.103.253/24
  no ipv6 redirects
  hsrp 103
    preempt
    ip 172.18.103.254

interface port-channel11
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  spanning-tree port type network
  vpc peer-link

interface port-channel15
  description vPC UCS 6454-1 FI
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  spanning-tree port type edge trunk
  mtu 9216
  load-interval counter 3 60
  vpc 15

interface port-channel16
  description vPC UCS 6454-2 FI
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  spanning-tree port type edge trunk
  mtu 9216
  load-interval counter 3 60
  vpc 16

interface port-channel135
  description vPC Upstream Network Switch A

interface port-channel136
  description vPC Upstream Network Switch B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119
  vpc 136

interface Ethernet1/1
  description vPC peer-link connection to AA19-9336-1 Ethernet1/1
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  channel-group 11 mode active

interface Ethernet1/2
  description vPC peer-link connection to AA19-9336-1 Ethernet1/2
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  channel-group 11 mode active

interface Ethernet1/5
  description vPC 15 connection to UCS 6454-1 FI Ethernet1/54
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  mtu 9216
  load-interval counter 3 60
  channel-group 15 mode active

interface Ethernet1/6
  description vPC 16 connection to UCS 6454-2 FI Ethernet1/54
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119,201-203,1000
  mtu 9216
  load-interval counter 3 60
  channel-group 16 mode active

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23
interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35
  interface Ethernet1/35
  description vPC 135 connection to Upstream Network Switch A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119
  channel-group 135 mode active

interface Ethernet1/36
  description vPC 136 connection to Upstream Network Switch B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 119
  channel-group 136 mode active

interface mgmt0
  vrf member management
  ip address 192.168.168.14/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.5a.bin
no system default switchport shutdown
```