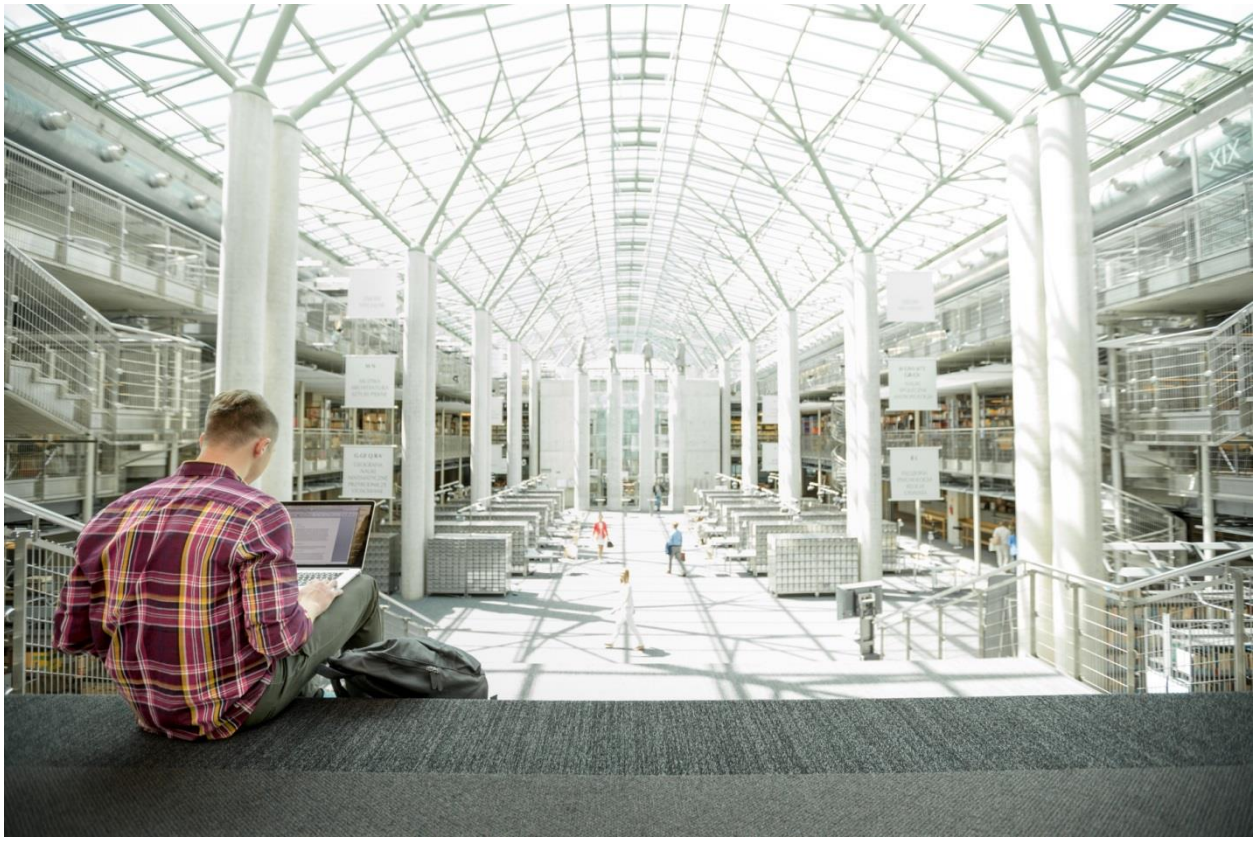

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure

Deployment Guide for Cisco and Hitachi Converged Infrastructure with Cisco UCS Blade Servers, Cisco Nexus 9336C-FX2 Switches, Cisco MDS 9706 Fabric Switches, and Hitachi VSP G1500 and VSP G370 Storage Systems with vSphere 6.5 and vSphere 6.7

Last Updated: April 3, 2019

CISCO VALIDATED DESIGN

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 steps through the deployment of the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a Virtual Server Infrastructure (VSI), as it was described in the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide. 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, Cisco UCS 6400 or 6300 Fabric Interconnects, Cisco Nexus 9000 Series switches, Cisco MDS Fibre channel switches, and Hitachi Virtual Storage Platform (VSP). This architecture is pulled together to support VMware vSphere 6.5 and VMware vSphere 6.7 to support a larger range of customer deployments within vSphere.

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 through the Cisco MDS multilayer switches to Cisco Unified Computing System (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 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
- Cisco Multilayer SAN 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

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(UCS) and the Hitachi Virtual Storage Platform(VSP) connected together by Cisco MDS Multilayer SAN Switches, and further enabled with Cisco Nexus Switches. These components come 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 for two similar topologies featuring differing Cisco UCS Fabric Interconnects as well as differing Hitachi VSP Storage Systems, with both using the same MDS and Nexus switching infrastructure.

The first 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.
- Cisco MDS 9706 – 32Gbps Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center.
- 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

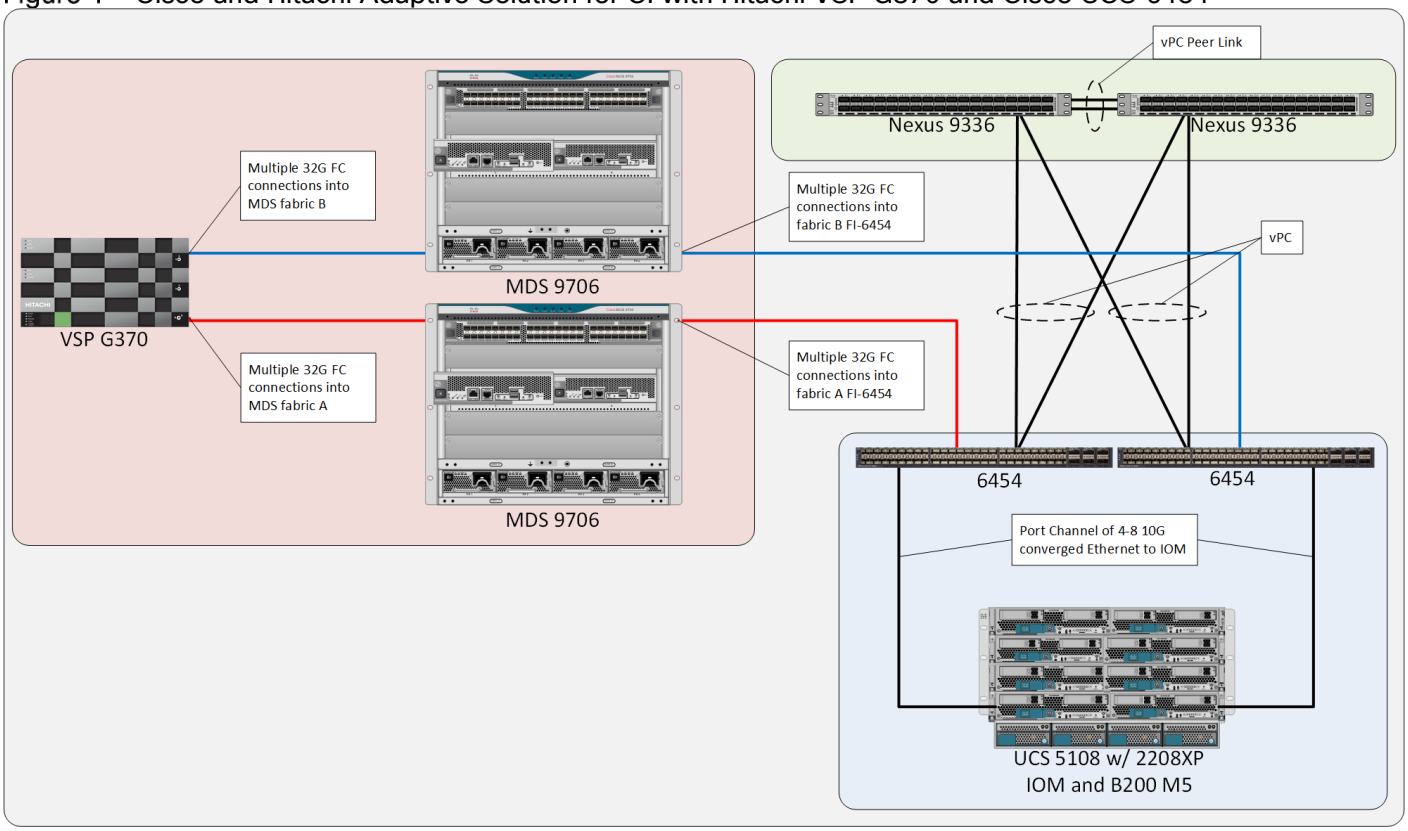

The Cisco UCS B200 M5 blade servers in this topology are hosted within a Cisco UCS 5108 Chassis, and connect into the fabric interconnects from the chassis using Cisco UCS 2208XP I/O Modules (IOM). The 2208XP IOM supports 10G connections into the 10/25G ports of the Cisco UCS 6454 FIs, delivering a high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

The second topology shown in Figure 2 leverages:

- Cisco Nexus 9336C-FX2 – 100Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6332-16UP 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.
- Cisco MDS 9706 – 32Gbps Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center.
- Hitachi VSP G1500 – Enterprise-level, high-performance storage system with optional all-flash configuration

Figure 2 Cisco and Hitachi Adaptive Solution for CI with Hitachi VSP G1500 and Cisco UCS 6332-16UP

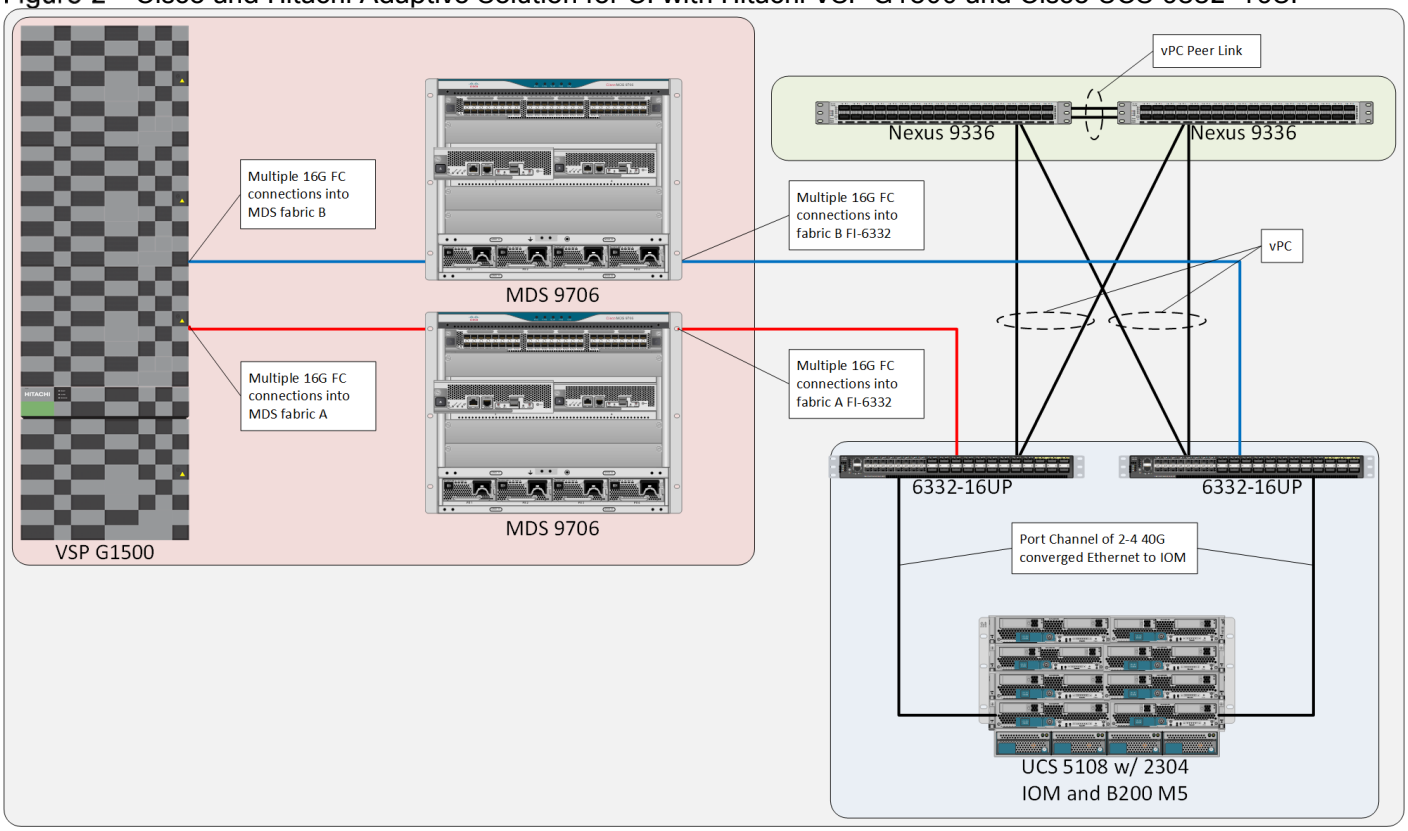

The Cisco UCS B200 M5 servers in this topology are hosted within the same Cisco UCS 5108 Chassis, but connect into the fabric interconnects from the chassis using Cisco UCS 2304 IOM. The Cisco UCS 2304 IOM supports 40G connections going into the Cisco UCS 6332-16UP FIs, delivering a high bandwidth solution that may fit well in a main office type setting.

Management components for both architectures 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.
- Cisco Data Center Network Manager (optional) – Multi-layer network configuration and monitoring.

Both topologies were validated for vSphere 6.5 U2 and vSphere 6.7 U1 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

Cisco MDS Recommended Releases:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b_MDS_NX-OS_Recommended_Releases.html

Hitachi Vantara Interoperability:

https://support.hitachivantara.com/en_us/interoperability.html

In addition, 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 6332	4.0(1b)
	Cisco UCS 2304 IOM	4.0(1b)
	Cisco UCS Fabric Interconnect 6454	4.0(1b)
	Cisco UCS 2208XP IOM	4.0(1b)
	Cisco UCS B200 M5	4.0(1b)

Component		Software Version/Firmware Version
	VMware vSphere	6.7 U1 VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso
	ESXi 6.7 U1 nenic	1.0.25.0
	ESXi 6.7 U1 nfnic	4.0.0.14
	VMware vSphere	6.5 U2 VMware-ESXi-6.5.0-9298722-Custom-Cisco-6.5.2.2.iso
	ESXi 6.5 U2 nenic	1.0.25.0
	ESXi 6.5 U2 fnic	1.6.0.44
	VM Virtual Hardware Version	13(1)
Storage	Hitachi VSP G1500	80-06-42-00/00
	Hitachi VSP G370	88-02-03-60/00
	Cisco MDS 9706 (DS-X97-SF1-K9 & DS-X9648-1536K9)	8.3(1)
	Cisco Data Center Network Manager	11.0(1)

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-Prod-02 to represent infrastructure and production 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-App2	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)	

Variable	Variable Description	Customer Deployed Value
<<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: N0taP4ss)	
<<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_mds_A_mgmt_ip>>	Cisco MDS Management IP address (Example: 192.168.168.18)	
<<var_mds_A_hostname>>	Cisco MDS hostname (Example: aa19-9706-1)	
<<var_mds_B_mgmt_ip>>	Cisco MDS Management IP address (Example: 192.168.168.19)	
<<var_mds_B_hostname>>	Cisco MDS hostname (Example: aa19-9706-2)	
<<var_vsan_a_id>>	VSAN used for the A Fabric between the VSP/MDS/FI (Example: 101)	
<<var_vsan_b_id>>	VSAN used for the A Fabric between the VSP/MDS/FI (Example: 102)	
<<vsp_hostname>> <<vsp-g370>> / <<vsp-g1500>>	Hitachi VSP storage system name (Example g370-[Serial Number])	
<<var_ucs_clustername>> <<var_ucs_6454_clustername>> / <<var_ucs_6332_clustername>>	Cisco UCS Manager cluster host name (Example: AA19-6454)	

Variable	Variable Description	Customer Deployed Value
<<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 two validated topologies 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 UCS 6454 with the VSP G370 Topology

Figure 3 shows the cabling configuration used in the design featuring the Cisco UCS 6454 with the VSP G370.

Figure 3 Cabling Diagram for Cisco and Hitachi Converged Infrastructure featuring Cisco UCS 6454 with the VSP G370

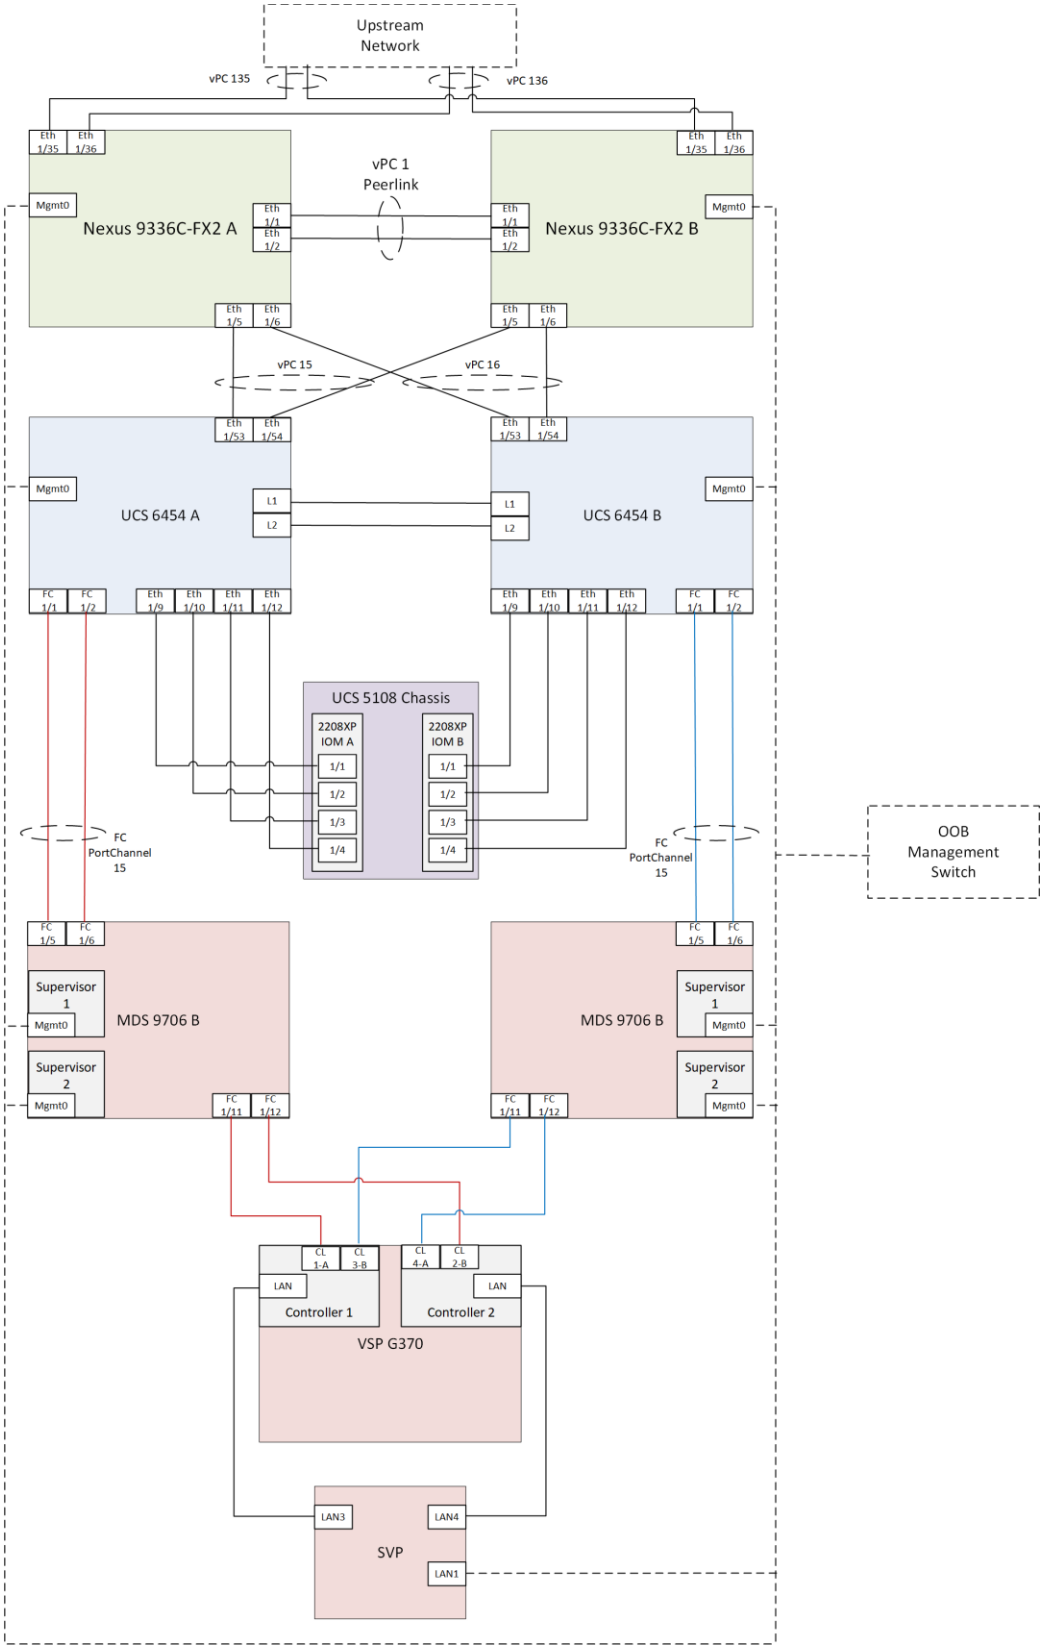

Tables listing the specific port connections with the cables used in the deployment of the Cisco UCS 6454 and the VSP G370 are provided below.

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

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
	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 Topology

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 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6332-16UP FI A	FC 1/1	32Gb FC	MDS 9706 A	FC 1/5
	FC 1/2	32Gb FC	MDS 9706 A	FC 1/6
	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

Local Device	Local Port	Connection	Remote Device	Remote Port
	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 6332-16UP FI B	FC 1/1	32Gb FC	MDS 9706 B	FC 1/5
	FC 1/2	32Gb FC	MDS 9706 B	FC 1/6
	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 Cisco MDS 9706 A Cabling Information for Cisco UCS 6454 to VSP G370 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9706 A	FC 1/5	32Gb FC	Cisco UCS 6454 FI A	FC 1/1
	FC 1/6	32Gb FC	Cisco UCS 6454 FI A	FC 1/2
	FC 1/11	32Gb FC	VSP G370 Controller 1	CL 1-A
	FC 1/12	32Gb FC	VSP G370 Controller 2	CL 2-B
	Sup1 MGMT0	GbE	GbE management switch	Any
	Sup2 MGMT0	GbE	GbE management switch	Any

The MDS DS-X9648-1536K9 4/8/16/32 Gbps Advanced FC Module used in this design does not have port groups with shared bandwidth, so sequential port selection will not impact bandwidth. When looking at substituting a differing MDS switch into the topology from the respective compatibility matrices, care should be given to any port group specifics on how bandwidth may be shared between ports.

Table 9 Cisco MDS 9706 B Cabling Information for Cisco UCS 6454 to VSP G370 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9706 B	FC 1/5	32Gb FC	Cisco UCS 6454 FI B	FC 1/1
	FC 1/6	32Gb FC	Cisco UCS 6454 FI B	FC 1/2
	FC 1/11	32Gb FC	VSP G370 Controller 1	CL 3-B
	FC 1/12	32Gb FC	VSP G370 Controller 2	CL 4-A
	Sup1 MGMT0	GbE	GbE management switch	Any
	Sup2 MGMT0	GbE	GbE management switch	Any

Table 10 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 MDS 9706 A	FC 1/11
	CL 2-B	32Gb FC	Cisco MDS 9706 A	FC 1/12
	CL 3-B	32Gb FC	Cisco MDS 9706 B	FC 1/11
	CL 4-A	32Gb FC	Cisco MDS 9706 B	FC 1/12
	Cont1 LAN	GbE	SVP	LAN3
	Cont2 LAN	GbE	SVP	LAN4

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.

Physical Cabling for the Cisco UCS 6332-16UP with the VSP G1500 Topology

Figure 4 illustrates the cabling configuration used in the design featuring the Cisco UCS 6332-16UP with the VSP G1500.

Figure 4 Cabling Diagram for Cisco and Hitachi Converged Infrastructure Featuring Cisco UCS 6332-16UP with the VSP G1500

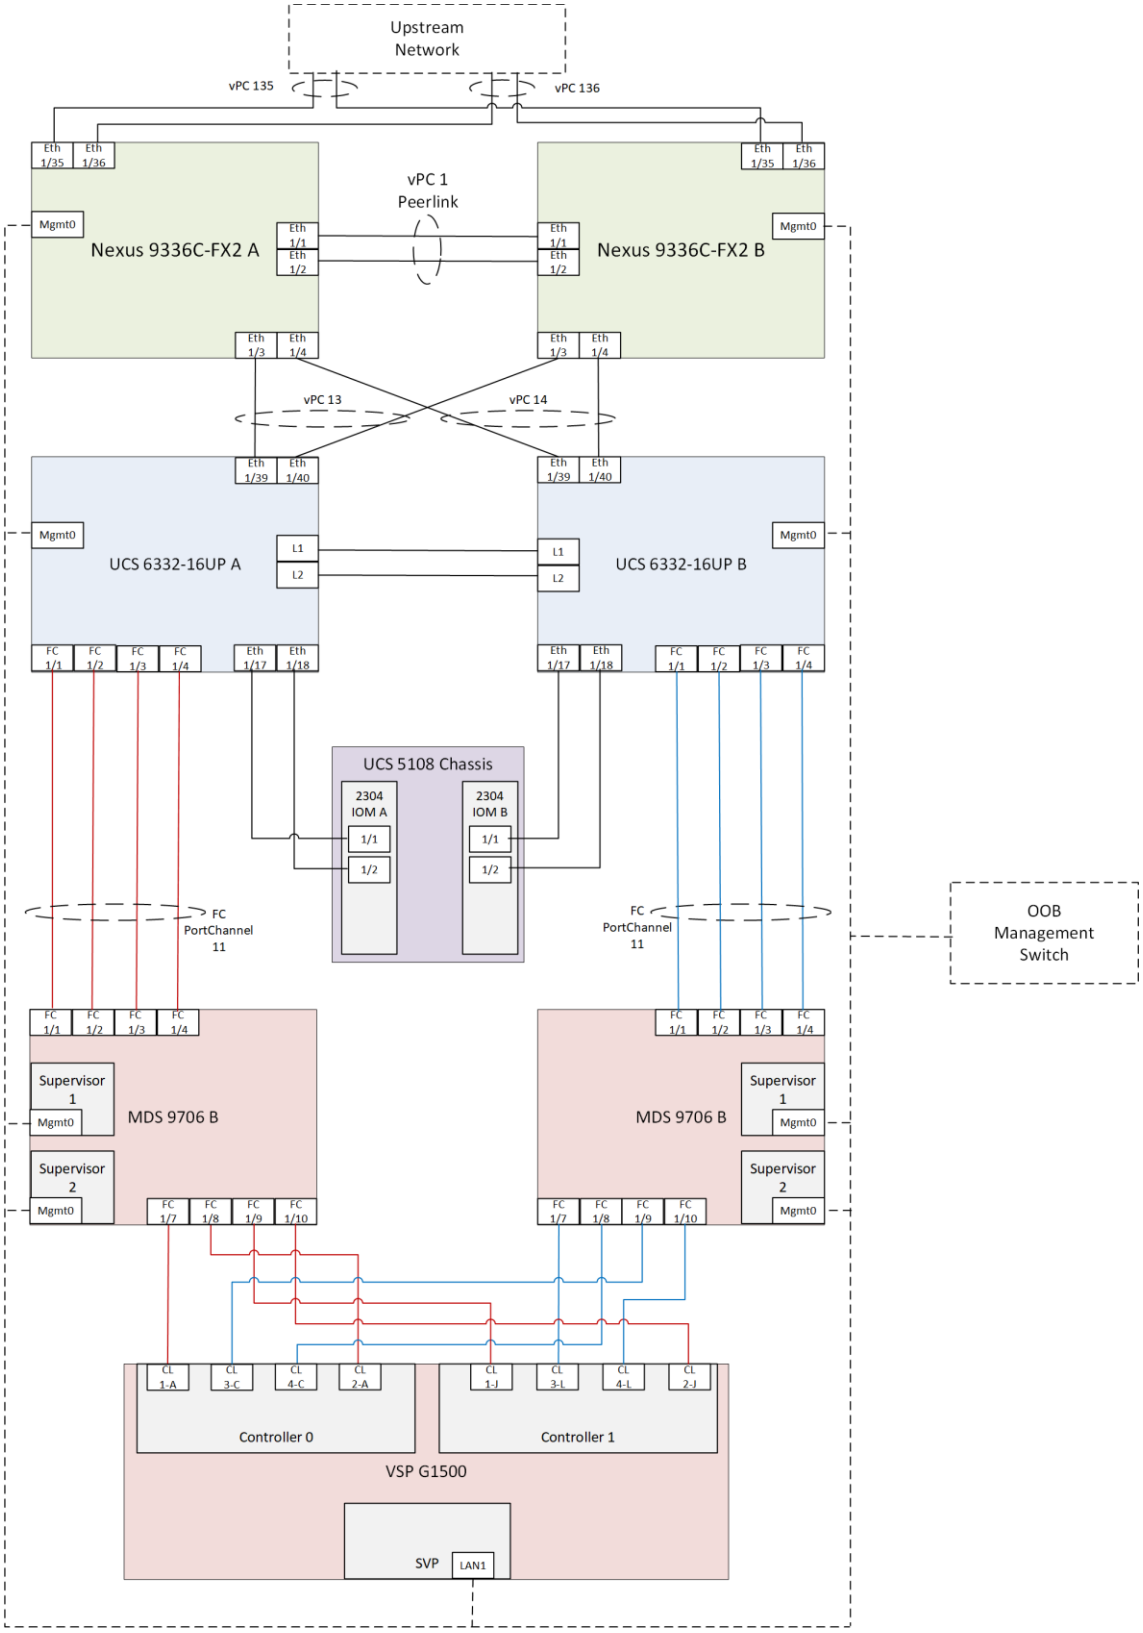

Tables listing the specific port connections with the cables used in the deployment of the Cisco UCS 6332-16UP and the VSP G1500 are below.

Table 11 Cisco Nexus 9336C-FX2 A Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

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
	Eth1/3	40GbE	Cisco UCS 6332-16UP FI A	Eth 1/39
	Eth1/4	40GbE	Cisco UCS 6332-16UP FI B	Eth 1/39
	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 12 Cisco Nexus 9336C-FX2 B Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

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
	Eth1/3	40GbE	Cisco UCS 6332-16UP FI A	Eth 1/40
	Eth1/4	40GbE	Cisco UCS 6332-16UP FI B	Eth 1/40
	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 13 Cisco UCS 6332-16UP A Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6332-16UP FI A	FC 1/1	16Gb FC	MDS 9706 A	FC 1/1
	FC 1/2	16Gb FC	MDS 9706 A	FC 1/2
	FC 1/3	16Gb FC	MDS 9706 A	FC 1/3
	FC 1/4	16Gb FC	MDS 9706 A	FC 1/4
	Eth1/17	40GbE	Cisco UCS Chassis 2304 FEX A	IOM 1/1
	Eth1/18	40GbE	Cisco UCS Chassis 2304 FEX A	IOM 1/2
	Eth1/39	40GbE	Cisco Nexus 9336C-FX2 A	Eth1/3
	Eth1/40	40GbE	Cisco Nexus 9336C-FX2 B	Eth1/3
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6454 FI B	L1

Local Device	Local Port	Connection	Remote Device	Remote Port
	L2	GbE	Cisco UCS 6454 FI B	L2

Ports 1-16 are Universal ports in the UCS 6332-16UP that can be used for Ethernet or Fibre Channel, with ports 17-40 primarily used as server ports either with 40Gbps QSFP+ ports, or breakout cables to support 10Gbps. The last ports of 35-40 are generally used for network uplinks and will not support QSFP copper twinax type cables.

Table 14 Cisco UCS 6332-16UP B Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6332-16UP FI B	FC 1/1	16Gb FC	MDS 9706 B	FC 1/1
	FC 1/2	16Gb FC	MDS 9706 B	FC 1/2
	FC 1/3	16Gb FC	MDS 9706 B	FC 1/3
	FC 1/4	16Gb FC	MDS 9706 B	FC 1/4
	Eth1/17	40GbE	Cisco UCS Chassis 2304 FEX A	IOM 1/1
	Eth1/18	40GbE	Cisco UCS Chassis 2304 FEX A	IOM 1/2
	Eth1/39	40GbE	Cisco Nexus 9336C-FX2 A	Eth1/4
	Eth1/40	40GbE	Cisco Nexus 9336C-FX2 B	Eth1/4
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6454 FI A	L1
	L2	GbE	Cisco UCS 6454 FI A	L2

Table 15 Cisco MDS 9706 A Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9706 A	FC 1/1	16Gb FC	Cisco UCS 6332-16UP FI A	FC 1/1
	FC 1/2	16Gb FC	Cisco UCS 6332-16UP FI A	FC 1/2
	FC 1/3	16Gb FC	Cisco UCS 6332-16UP FI A	FC 1/3
	FC 1/4	16Gb FC	Cisco UCS 6332-16UP FI A	FC 1/4
	FC 1/7	16Gb FC	VSP G1500	CL 1-A
	FC 1/8	16Gb FC	VSP G1500	CL 2-A
	FC 1/9	16Gb FC	VSP G1500	CL 1-J
	FC 1/10	16Gb FC	VSP G1500	CL 2-J
	Sup1 MGMT0	GbE	GbE management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	Sup2 MGMT0	GbE	GbE management switch	Any

The MDS DS-X9648-1536K9 4/8/16/32 Gbps Advanced FC Module used with the MDS 9706 in this design does not have port groups with shared bandwidth, so sequential port selection will not impact bandwidth. When looking at substituting a differing MDS switch into the topology from the respective compatibility matrices, care should be given to any port group specifics on how bandwidth may be shared between ports.

Table 16 Cisco MDS 9706 B Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9706 B	FC 1/1	16Gb FC	Cisco UCS 6332-16UP FI B	FC 1/1
	FC 1/2	16Gb FC	Cisco UCS 6332-16UP FI B	FC 1/2
	FC 1/3	16Gb FC	Cisco UCS 6332-16UP FI B	FC 1/3
	FC 1/4	16Gb FC	Cisco UCS 6332-16UP FI B	FC 1/4
	FC 1/7	16Gb FC	VSP G1500	CL 3-L
	FC 1/8	16Gb FC	VSP G1500	CL 4-C
	FC 1/9	16Gb FC	VSP G1500	CL 3-C
	FC 1/10	16Gb FC	VSP G1500	CL 4-L
	Sup1 MGMT0	GbE	GbE management switch	Any
	Sup2 MGMT0	GbE	GbE management switch	Any

Table 17 Hitachi VSP G1500 Cabling Information for Cisco UCS 6332-16UP to VSP G1500 Topology

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G1500	CL 1-A	16Gb FC	Cisco MDS 9706 A	FC 1/7
	CL 2-A	16Gb FC	Cisco MDS 9706 A	FC 1/8
	CL 1-J	16Gb FC	Cisco MDS 9706 A	FC 1/9
	CL 2-J	16Gb FC	Cisco MDS 9706 A	FC 1/10
	CL 3-L	16Gb FC	Cisco MDS 9706 B	FC 1/7
	CL 4-C	16Gb FC	Cisco MDS 9706 B	FC 1/8
	CL 3-C	16Gb FC	Cisco MDS 9706 B	FC 1/9
	CL 4-L	16Gb FC	Cisco MDS 9706 B	FC 1/10
	SVP	GbE	GbE management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	LAN			

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 explains 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.

Configuration connections for both Fabric Interconnect platforms are listed in these steps and both sets of Cisco UCS vPCs are not necessary in a deployment that will only deploy a single UCS domain.

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 13
AA19-9336-1(config-if)# description vPC UCS 6332-16UP-1 FI
AA19-9336-1(config-if)# interface port-channel 14
AA19-9336-1(config-if)# description vPC UCS 6332-16UP-2 FI
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 Ethernet1/3
AA19-9336-1(config-if)# description vPC 13 connection to UCS 6332-16UP-1 FI Ethernet1/39
AA19-9336-1(config-if)# interface Ethernet1/4
AA19-9336-1(config-if)# description vPC 14 connection to UCS 6332-16UP-2 FI Ethernet1/39
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 13
AA19-9336-2(config-if)# description vPC UCS 6332-16UP-1 FI
AA19-9336-2(config-if)# interface port-channel 14
AA19-9336-2(config-if)# description vPC UCS 6332-16UP-2 FI
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 Ethernet1/3
AA19-9336-2(config-if)# description vPC 13 connection to UCS 6332-16UP-1 FI Ethernet1/40
AA19-9336-2(config-if)# interface Ethernet1/4
AA19-9336-2(config-if)# description vPC 14 connection to UCS 6332-16UP-2 FI Ethernet1/40
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:

Port Channels to both a Cisco UCS 6332-16UP FI pair and Cisco UCS 6454 FI pair are shown below. The specific ports selected for these connections into the Nexus should reflect the cabling implemented for deployed Cisco UCS FI in a customer environment, and both should not be configured as shown in this example unless two UCS FI pairs are being deployed.

Nexus Connection vPC to Cisco UCS 6332-16UP A

```
AA19-9336-1&2 (config-if)# int ethernet 1/3
AA19-9336-1&2 (config-if)# channel-group 13 mode active
AA19-9336-1&2 (config-if)# no shut
AA19-9336-1&2 (config-if)# int port-channel 13
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 13
```

Nexus Connection vPC to Cisco UCS 6332-16UP B

```
AA19-9336-1&2 (config-if)# int ethernet 1/4
AA19-9336-1&2 (config-if)# channel-group 14 mode active
AA19-9336-1&2 (config-if)# no shut
AA19-9336-1&2 (config-if)# int port-channel 14
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 14
```

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 above command block, the “l4port” is 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
```

Cisco MDS Configuration

The MDS configuration implements a common redundant physical fabric design with fabrics represented as "A" and "B". The validating lab provided a basic MDS fabric supporting two VSP Storage Systems that are connected to two differing UCS domains within the SAN environment. Larger deployments may require a multi-tier core-edge or edge-core-edge design with port channels connecting the differing layers of the topology. Further discussion of these kinds of topologies, as well as considerations in implementing more complex SAN environments can be found in this white paper: https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white-paper-c11-729697.pdf

The configuration steps described below are implemented for the Cisco MDS 9706, but are similar to steps required for other Cisco MDS 9000 series switches that may be appropriate for a deployment. When making changes to the design that comply with the compatibility matrices of Cisco and Hitachi, it is required to consult the appropriate configuration documents of the differing equipment to confirm the correct implementation steps.

Physical Connectivity

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

Initial MDS 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:

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


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

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

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

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

Please register Cisco MDS 9000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. MDS 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]: <enter>

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

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

  Enter the switch name : <<var_mds_A_hostname>>|<<var_mds_B_hostname>>

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

```
    Mgmt0 IPv4 address : <<var_mds_A_mgmt_ip>>|<<var_mds_B_mgmt_ip>>

    Mgmt0 IPv4 netmask : <<var_oob_netmask>>

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

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

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

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

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

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

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

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

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

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

    Congestion-drop for logical-type core must be greater than or equal to
    Congestion-drop for logical-type edge. Hence, Congestion drop for
    logical-type core will be set as default.

  Enable the http-server? (yes/no) [y]: <enter>

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

Clock config format [HH:MM:SS Day Mon YYYY] [example: 18:00:00 1 november 2012]: <enter>

Enter clock config :17:26:00 2 january 2019

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

Enter timezone config [PST/MST/CST/EST] :EST
Enter Hrs offset from UTC [-23:+23] : <enter>
Enter Minutes offset from UTC [0-59] : <enter>

 Configure summertime? (yes/no) [n]: <enter>

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

    NTP server IPv4 address : <var_oob_ntp>

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

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

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

  Configure default zone policy (permit/deny) [deny]: <enter>

  Enable full zoneset distribution? (yes/no) [n]: <enter>

  Configure default zone mode (basic/enhanced) [basic]: <enter>

The following configuration will be applied:
  password strength-check
  switchname aa19-9706-1
  interface mgmt0
    ip address 192.168.168.18 255.255.255.0
    no shutdown
  ip default-gateway 192.168.168.254
  ssh key rsa 1024 force
  feature ssh
  no feature telnet
  system timeout congestion-drop default logical-type edge
```

```
  system timeout congestion-drop default logical-type core
  feature http-server
 clock set 17:26:00 2 january 2019
 clock timezone EST 0 0
  ntp server 192.168.168.254
  system default switchport shutdown
  system default switchport trunk mode on
  no system default zone default-zone permit
  no system default zone distribute full
  no system default zone mode enhanced

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

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

Cisco MDS Switch Configuration

Cisco MDS 9706 A and Cisco MDS 9706 B

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

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

```
aa19-9706-1&2# configure terminal
aa19-9706-1&2(config)# feature npiv
aa19-9706-1&2(config)# feature fport-channel-trunk
aa19-9706-1&2(config)# feature lldp
aa19-9706-1&2(config)# device-alias mode enhanced
aa19-9706-1&2(config)# device-alias commit
```

The device-alias commit will trigger a warning that this command will clear existing device aliases on attached fabrics, which should not impact the initial deployment instructions being followed here.

Configure Individual Ports

Cisco MDS 9706 A

To configure individual ports and port-channels for switch A, follow these steps:

In this step and in the following sections, configure the <var_ucs_6454_clustername> / <var_ucs_6332_clustername> and <vsp-g370> / <vsp-g1500> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
aa19-9706-1(config)# interface fc1/1
aa19-9706-1(config-if)# switchport description <var_ucs_6332_clustername>-a:1/1
aa19-9706-1(config-if)# channel-group 11 force
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/2
aa19-9706-1(config-if)# switchport description <var_ucs_6332_clustername>-a:1/2
aa19-9706-1(config-if)# channel-group 11 force
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/3
```

```
aa19-9706-1(config-if)# switchport description <var_ucs_6332_clustername>-a:1/3
aa19-9706-1(config-if)# channel-group 11 force
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/4
aa19-9706-1(config-if)# switchport description <var_ucs_6332_clustername>-a:1/4
aa19-9706-1(config-if)# channel-group 11 force
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/5
aa19-9706-1(config-if)# switchport description <var_ucs_6454_clustername>-a:1/1
aa19-9706-1(config-if)# channel-group 15 force
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/6
aa19-9706-1(config-if)# switchport description <var_ucs_6454_clustername>-a:1/2
aa19-9706-1(config-if)# channel-group 15 force
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/7
aa19-9706-1(config-if)# switchport description <vsp-g1500>-a:CL 1-A
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/8
aa19-9706-1(config-if)# switchport description <vsp-g1500>-a:CL 2-A
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/9
aa19-9706-1(config-if)# switchport description <vsp-g1500>-a:CL 1-J
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/10
aa19-9706-1(config-if)# switchport description <vsp-g1500>-a:CL 2-J
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/11
aa19-9706-1(config-if)# switchport description <vsp-g370>-a:CL 1-A
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface fc1/12
aa19-9706-1(config-if)# switchport description <vsp-g370>-a:CL 2-B
aa19-9706-1(config-if)# no shutdown
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface port-channel 11
aa19-9706-1(config-if)# switchport description <var_ucs_6332_clustername>-portchannel
aa19-9706-1(config-if)# channel mode active
aa19-9706-1(config-if)#
aa19-9706-1(config-if)# interface port-channel 15
aa19-9706-1(config-if)# switchport description <var_ucs_6454_clustername>-portchannel
aa19-9706-1(config-if)# channel mode active
aa19-9706-1(config-if)# exit
```

Create Port Descriptions - Fabric B

To configure individual ports and port-channels for switch B, follow these steps:

From the global configuration mode, run the following commands:

```
aa19-9706-2(config-if)# interface fc1/1
aa19-9706-2(config-if)# switchport description <var_ucs_6332_clustername>-b:1/1
aa19-9706-2(config-if)# channel-group 11 force
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/2
aa19-9706-2(config-if)# switchport description <var_ucs_6332_clustername>-b:1/2
aa19-9706-2(config-if)# channel-group 11 force
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/3
```

```
aa19-9706-2(config-if)# switchport description <var_ucs_6332_clustername>-b:1/3
aa19-9706-2(config-if)# channel-group 11 force
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/4
aa19-9706-2(config-if)# switchport description <var_ucs_6332_clustername>-b:1/4
aa19-9706-2(config-if)# channel-group 11 force
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/5
aa19-9706-2(config-if)# switchport description <var_ucs_6454_clustername>-b:1/1
aa19-9706-2(config-if)# channel-group 15 force
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/6
aa19-9706-2(config-if)# switchport description <var_ucs_6454_clustername>-b:1/2
aa19-9706-2(config-if)# channel-group 15 force
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/7
aa19-9706-2(config-if)# switchport description <vsp-g1500>-a:CL 3-L
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/8
aa19-9706-2(config-if)# switchport description <vsp-g1500>-a:CL 4-C
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/9
aa19-9706-2(config-if)# switchport description <vsp-g1500>-a:CL 3-C
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/10
aa19-9706-2(config-if)# switchport description <vsp-g1500>-a:CL 4-L
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/11
aa19-9706-2(config-if)# switchport description <vsp-g370>-a:CL 3-B
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface fc1/12
aa19-9706-2(config-if)# switchport description <vsp-g370>-a:CL 4-A
aa19-9706-2(config-if)# no shutdown
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface port-channel 11
aa19-9706-2(config-if)# switchport description <var_ucs_6332_clustername>-portchannel
aa19-9706-2(config-if)# channel mode active
aa19-9706-2(config-if)#
aa19-9706-2(config-if)# interface port-channel 15
aa19-9706-2(config-if)# switchport description <var_ucs_6454_clustername>-portchannel
aa19-9706-2(config-if)# channel mode active
aa19-9706-2(config-if)# exit
```

Create VSANs

Cisco MDS 9706 A

To create the necessary VSANs for fabric A and add ports, follow these steps:

From the global configuration mode, run the following commands:

```
aa19-9706-1(config)# vsan database
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id>
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> name Fabric-A
aa19-9706-1(config-vsan-db)# exit
aa19-9706-1(config)# zone smart-zoning enable vsan <var_vsan_a_id>
aa19-9706-1(config)# vsan database
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface fc1/7
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface fc1/8
```

```
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface fc1/9
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface fc1/10
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface fc1/11
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface fc1/12
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface port-channel 11
aa19-9706-1(config-vsan-db)# vsan <var_vsan_a_id> interface port-channel 15
aa19-9706-1(config-vsan-db)# end
aa19-9706-1# copy run start
```

For the fc 1/x vsan assignments above and below, there will be a warning message about traffic impact for these changes, which can be ignored. The option of "y" to continue should be specified if asked.

Cisco MDS 9706 B

To create the necessary VSANs for fabric B and add ports to them, follow these steps:

From the global configuration mode, run the following commands:

```
aa19-9706-2(config)# vsan database
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id>
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> name Fabric-B
aa19-9706-2(config-vsan-db)#exit
aa19-9706-2(config)# zone smart-zoning enable vsan <var_vsan_b_id>
aa19-9706-2(config)# vsan database
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface fc1/7
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface fc1/8
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface fc1/9
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface fc1/10
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface fc1/11
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface fc1/12
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface port-channel 11
aa19-9706-2(config-vsan-db)#vsan <var_vsan_b_id> interface port-channel 15
aa19-9706-2(config-vsan-db)# end
aa19-9706-2# copy run start
```

Configuring Fibre Channel Ports on Hitachi Virtual Storage Platform

In order for Hitachi Virtual Storage Platform fibre channel ports to be exposed properly to the MDS and 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. Note that URLs for VSP F1500 and VSP G1500 are different than those for Hitachi Virtual Storage Platform Fx00 Models and Gx00 Models:

2. VSP F1500 and VSP G1500: https://<IP of Storage System SVP>/sanproject/emergency.do, for example, if Storage System SVP IP address is 10.0.0.1, the URL would be:

   https://10.0.0.1/sanproject/emergency.do

3. VSP Fx00 Models and VSP Gx00 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/88600451200/emergency.do

4. Log into Hitachi Storage Navigator.

5. From the left Explorer pane, select the Storage Systems tab.

6. 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.

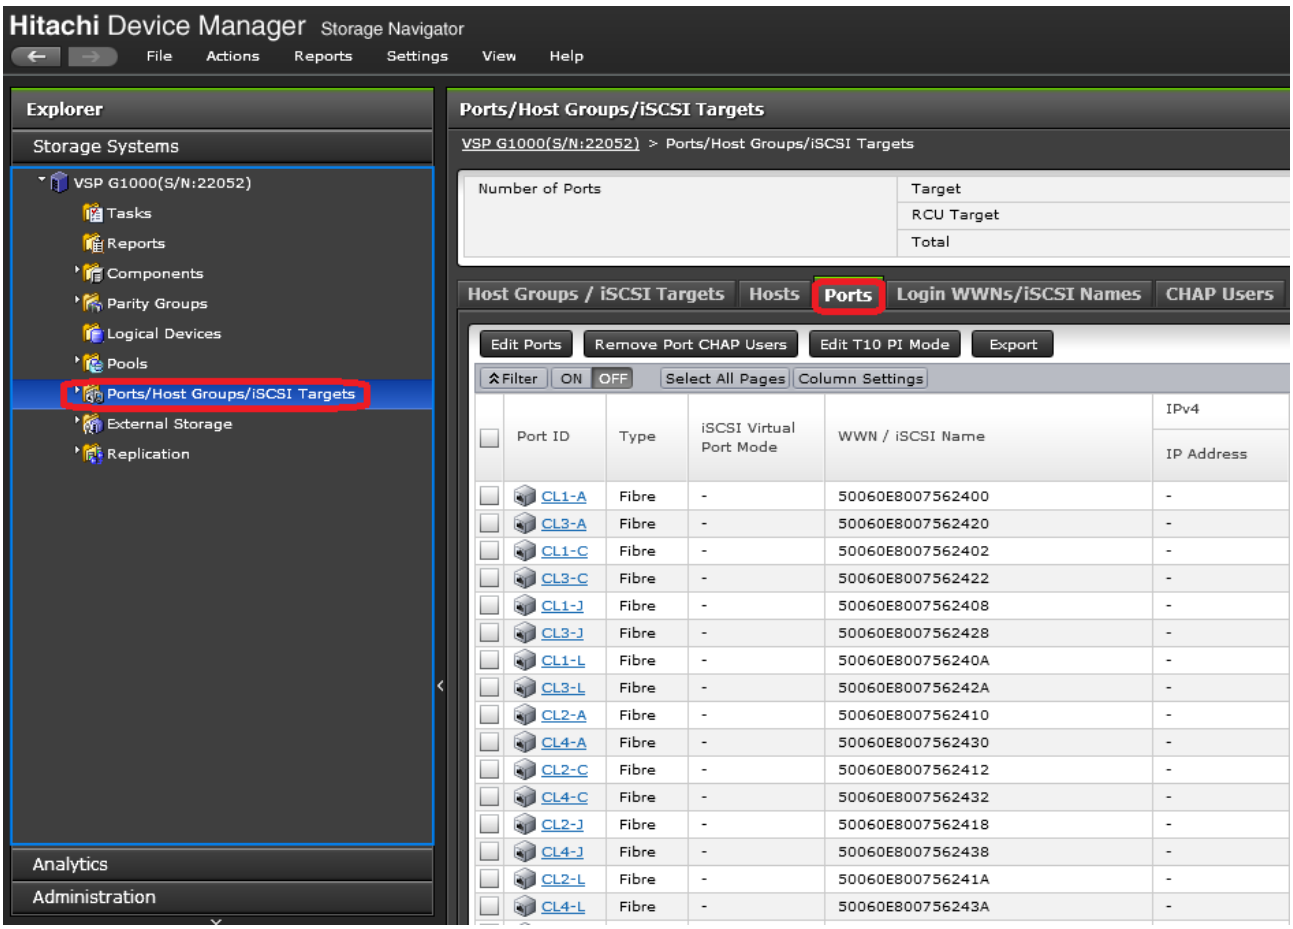

7. Select the checkboxes for the ports being used within the solution, then click the Edit Ports button to instantiate the Edit Ports dialog box.

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

   - Port Attribute: Target
   - Port Security: Enable
   - Port Speed: 16Gbps (G1500), 32Gbps (G350/370/700/900)
   - Fabric: ON
   - Connection Type: P-to-P

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

9. Example ports used in the Cisco UCS 6454 to VSP G370 used in this design are listed in Table 18 .

Table 18 VSP G370 to MDS Ports

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G370	CL 1-A	32Gb FC	Cisco MDS 9706 A	FC 1/11
	CL 2-B	32Gb FC	Cisco MDS 9706 A	FC 1/12

Local Device	Local Port	Connection	Remote Device	Remote Port
	CL 3-B	32Gb FC	Cisco MDS 9706 B	FC 1/11
	CL 4-A	32Gb FC	Cisco MDS 9706 B	FC 1/12

10. Example ports used in the Cisco UCS 6332-16UP to VSP G1500 used in this design are listed in Table 19 .

Table 19 VSP G1500 to MDS Ports

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G1500	CL 1-A	16Gb FC	Cisco MDS 9706 A	FC 1/7
	CL 2-A	16Gb FC	Cisco MDS 9706 A	FC 1/8
	CL 1-J	16Gb FC	Cisco MDS 9706 A	FC 1/9
	CL 2-J	16Gb FC	Cisco MDS 9706 A	FC 1/10
	CL 3-L	16Gb FC	Cisco MDS 9706 B	FC 1/7
	CL 4-C	16Gb FC	Cisco MDS 9706 B	FC 1/8
	CL 3-C	16Gb FC	Cisco MDS 9706 B	FC 1/9
	CL 4-L	16Gb FC	Cisco MDS 9706 B	FC 1/10

Figure 5 VSP G370 Edit Ports Pop-Up Window

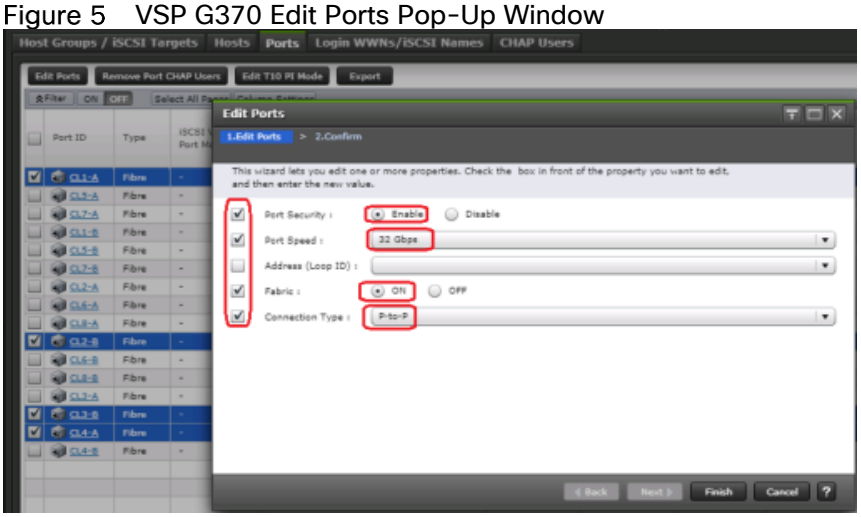

Figure 6 VSP G1500 Edit Ports Pop-Up Window

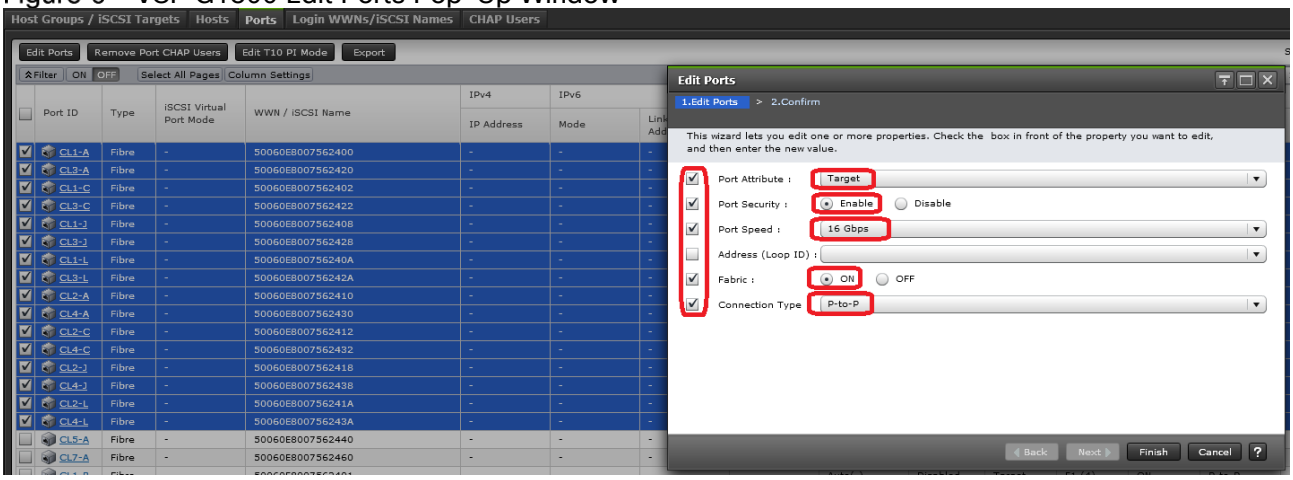

11. Click OK for any warning that appears.

12. Click Finish.

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

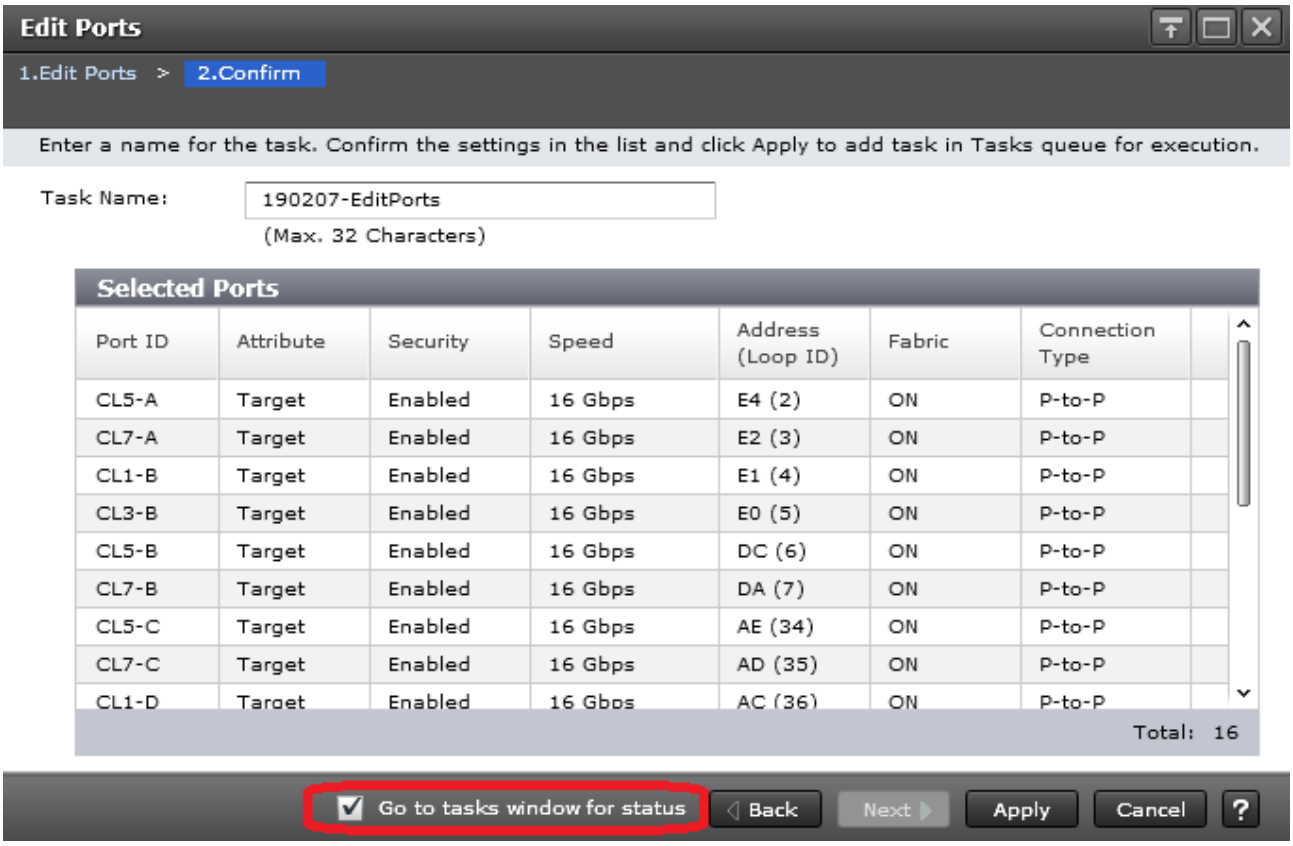

14. The Task view window will appear and show the completion status of the Edit Ports task. Wait until the task status shows Complete and proceed to the next section.

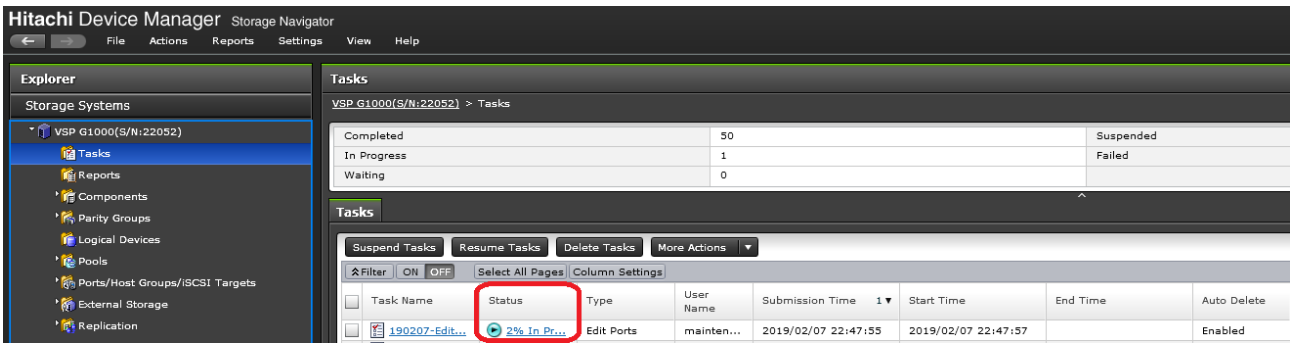
Hitachi Device Manager Storage Navigator
File Actions Reports Settings View Help
Explorer
Storage Systems
VSP G1000(S/N:22052)
Tasks
Reports
Components
Parity Groups
Logical Devices
Pools
Ports/Host Groups/iSCSI Targets
External Storage
Replication
Tasks
VSP G1000(S/N:22052) > Tasks
Completed 50 Suspended
In Progress 1 Failed
Waiting 0
Tasks
Suspend Tasks Resume Tasks Delete Tasks More Actions
Filter ON OFF Select All Pages Column Settings
Task Name Status Type User Name Submission Time Start Time End Time Auto Delete
190207-Edit... 2% In Pr... Edit Ports mainten... 2019/02/07 22:47:55 2019/02/07 22:47:57 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 and MDS Switches covered 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.

The Cisco UCS 6332-16UP was additionally validated in this design, but will not be covered step by step for configuration within this section. Steps that will differ in the configuration of the 6332-16UP versus those used for the 6454, will be called out as to their differences.

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(1b)

This document assumes the use of Cisco UCS 4.0(1b). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.0(1b), 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 (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, complete the following 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:

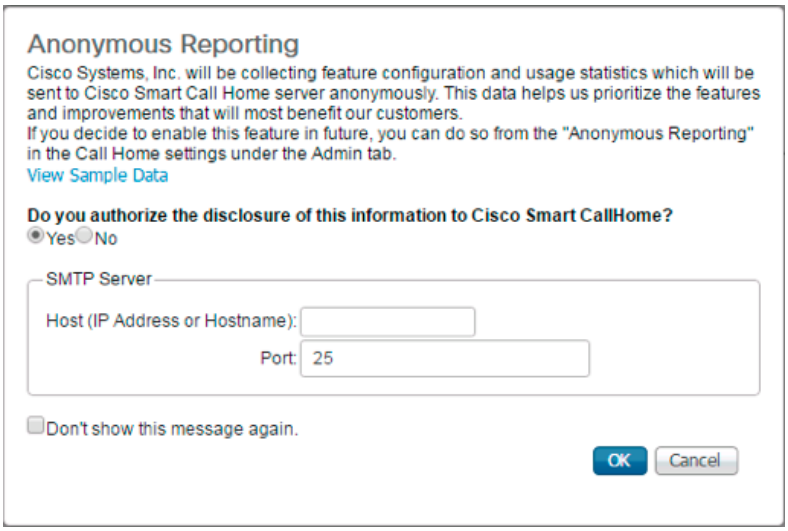

If you want to enable or disable Anonymous Reporting at a later date, it can be found within Cisco UCS Manager under: Admin -> Communication Management -> Call Home, which has a tab on the far right for Anonymous Reporting.

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.

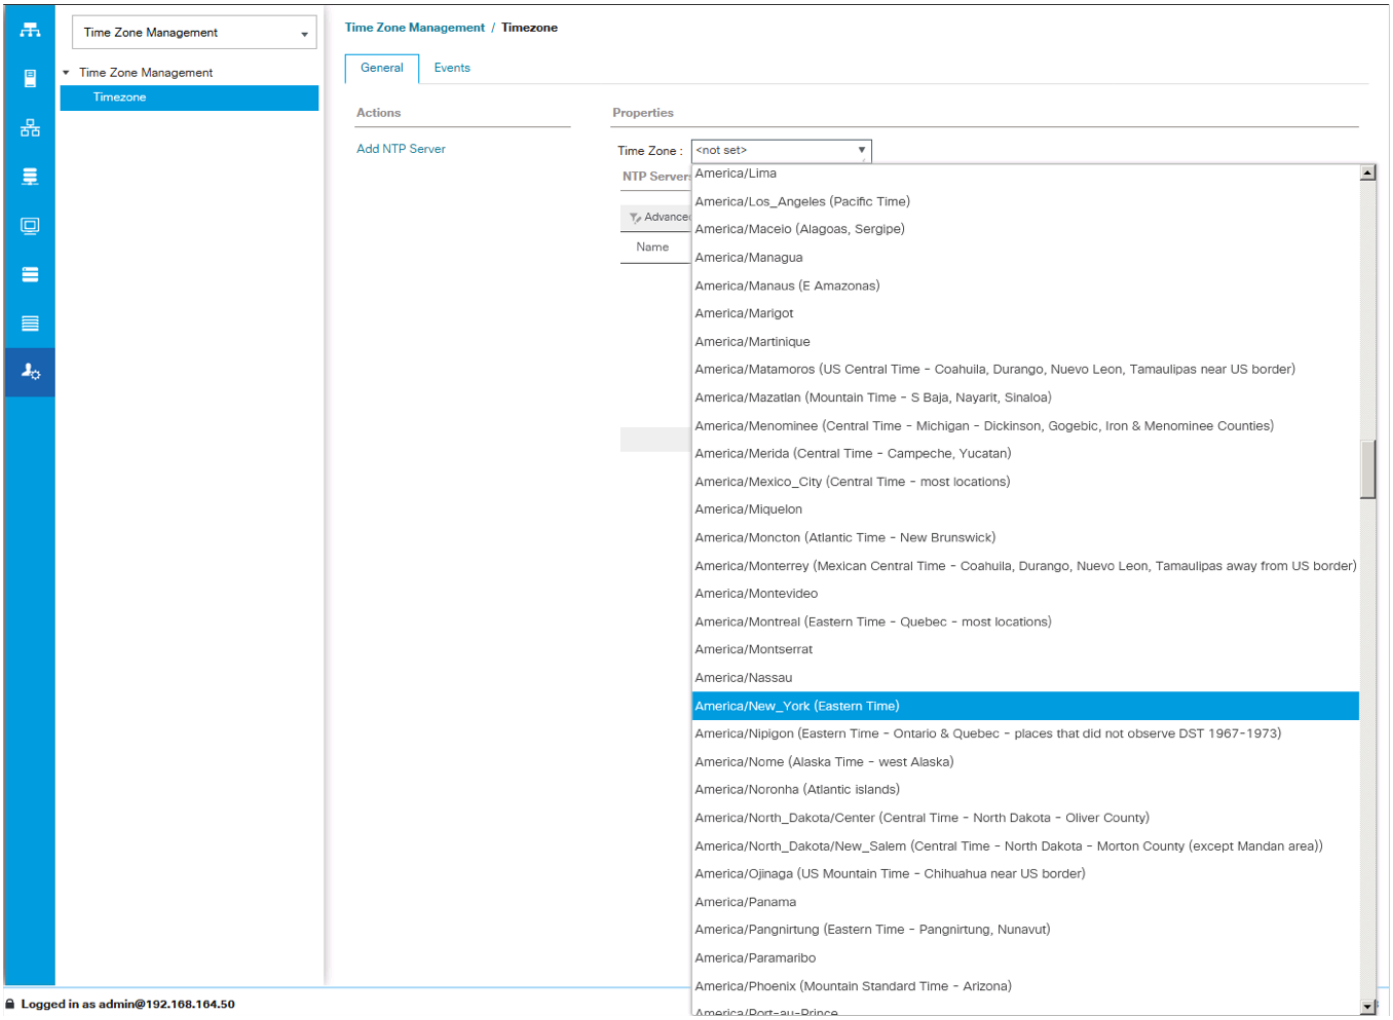

3. In the Properties pane, select the appropriate time zone in the Timezone menu.

4. Click Save Changes, and then click OK.

5. Click Add NTP Server.

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

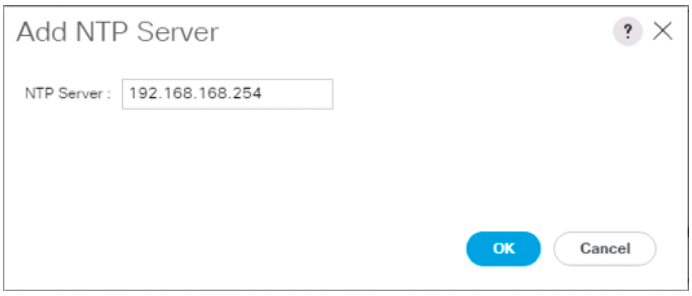

7. Click OK.

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 on the left under the drop-down.

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

3. Set the Link Grouping Preference to Port Channel.

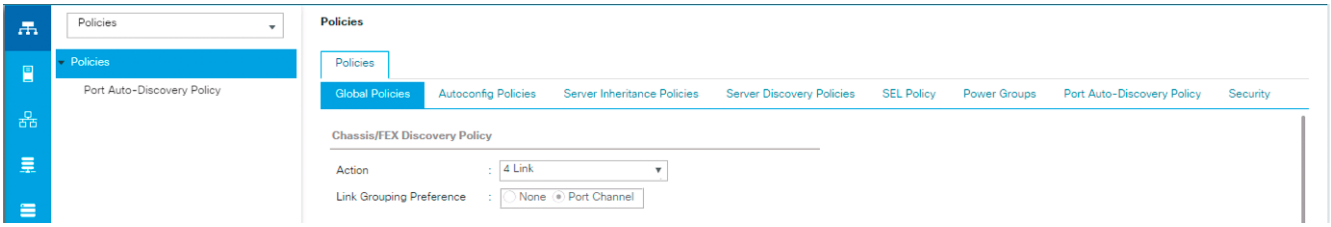

4. Leave other settings alone or change if appropriate to your environment.

5. Click Save Changes.

6. Click OK.

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 Ethernet Ports.

4. Select the ports that are connected to the chassis, right-click them, and select “Configure as Server Port.”

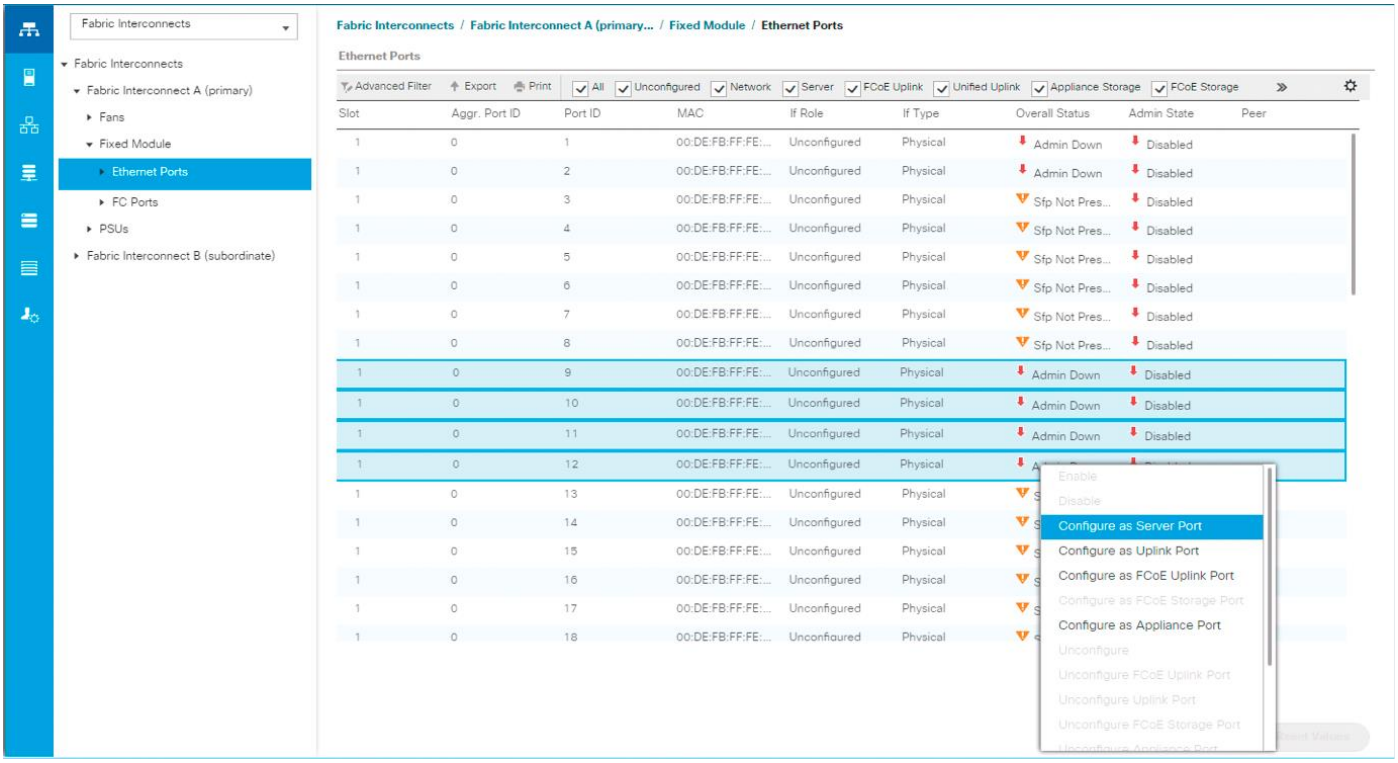

5. Click Yes to confirm server ports and click OK.

6. Verify that the ports connected to the chassis are now configured as server ports.

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

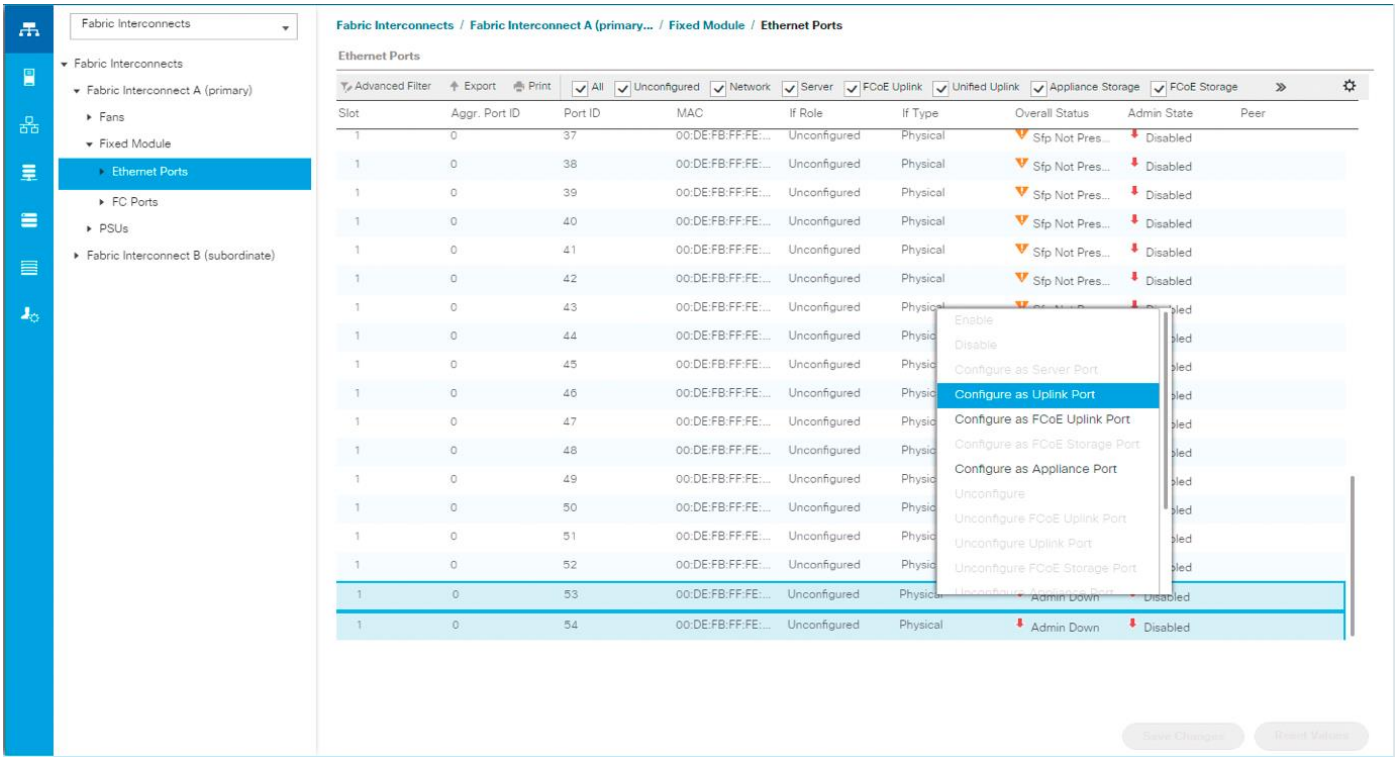

8. Click Yes to confirm uplink ports and click OK.

9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.

10. Expand Ethernet Ports.

11. Select the ports that are connected to the chassis, right-click them and select Configure as Server Port.

12. Click Yes to confirm server ports and click OK.

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

14. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis

To acknowledge all Cisco UCS chassis, follow these steps:

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

2. Expand Chassis and select each chassis that is listed.

3. Right-click each chassis and select Acknowledge Chassis.

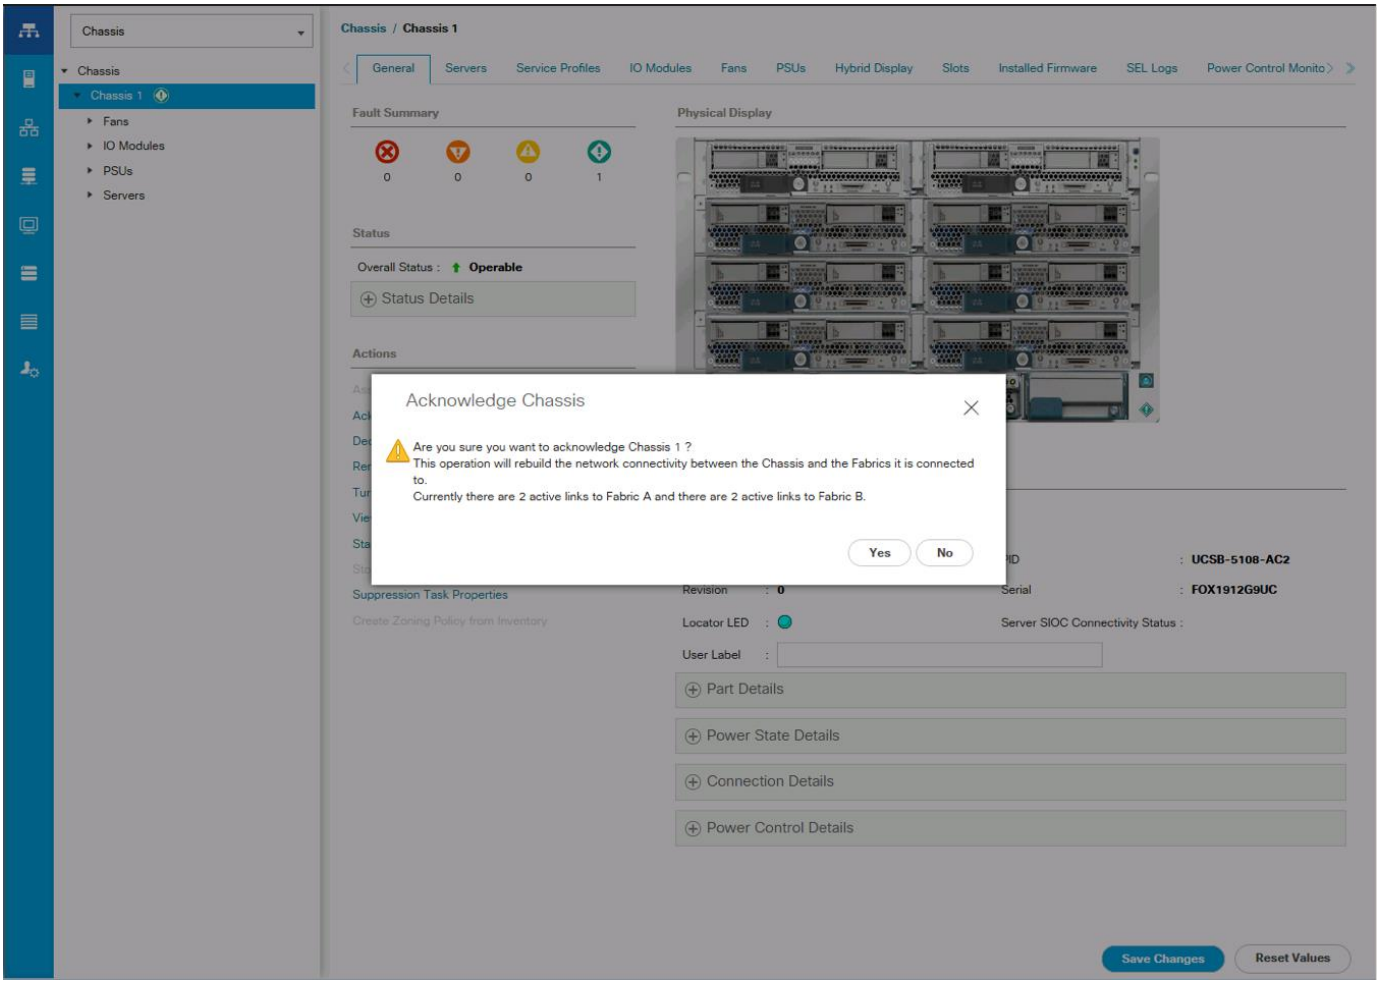

4. Click Yes and then click OK to complete acknowledging the chassis.

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.

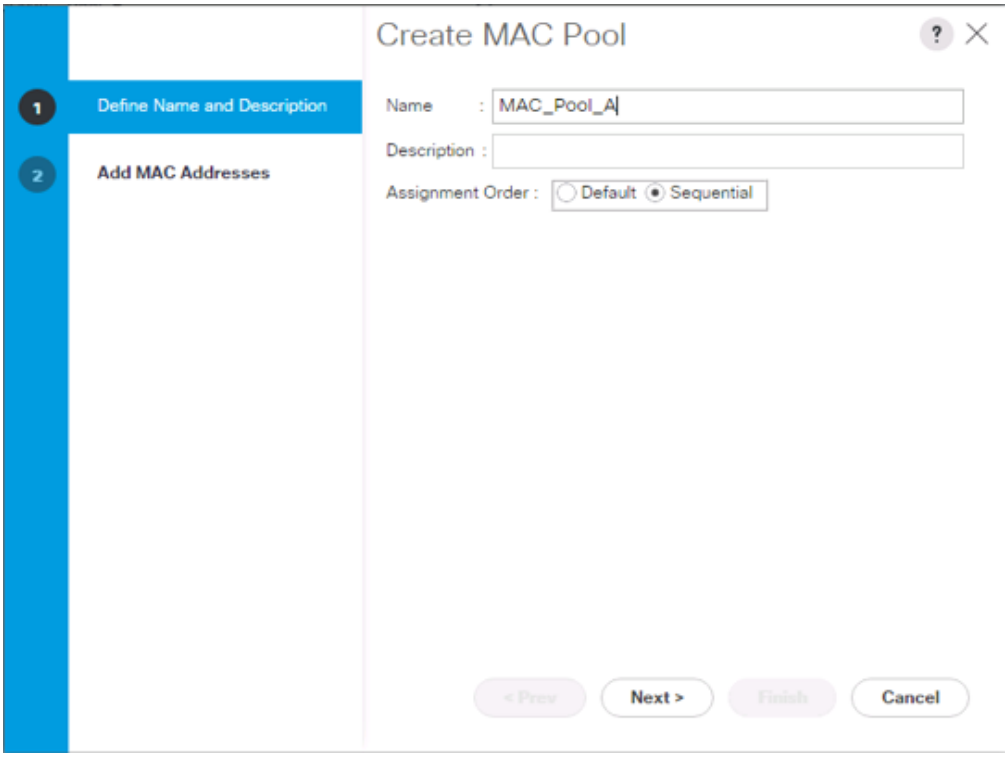

8. Click Next.

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) vs 32(for UCS 6332-16UP) 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.

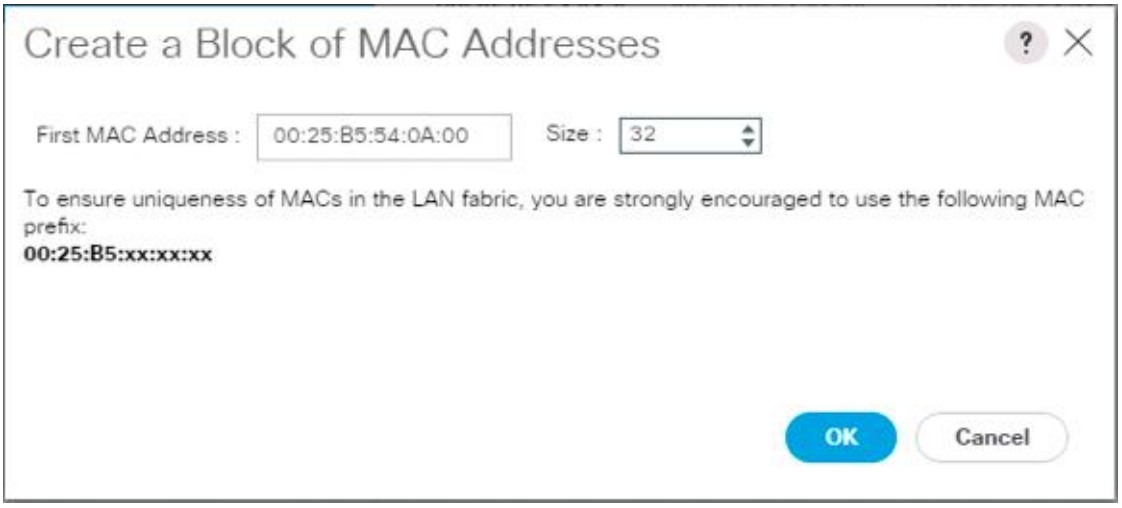

12. Click OK.

13. Click Finish.

14. In the confirmation message, click OK.

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

16. Select Create MAC Pool to create the MAC address pool.

17. Enter MAC_Pool_B as the name of the MAC pool.

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

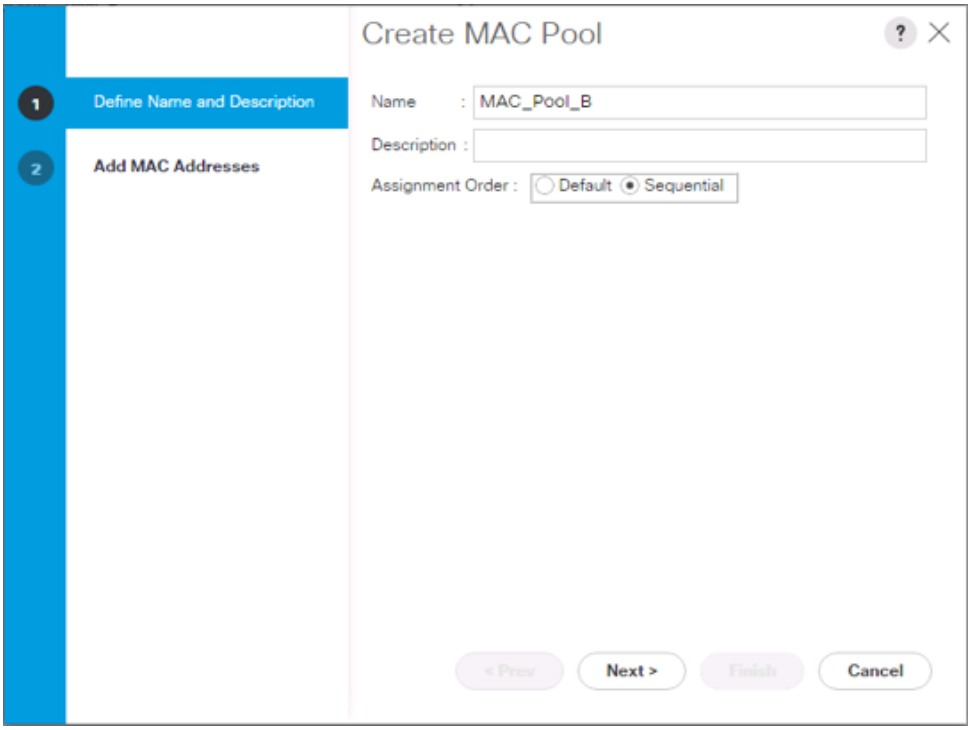

19. Click Next.

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 vs 32 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.

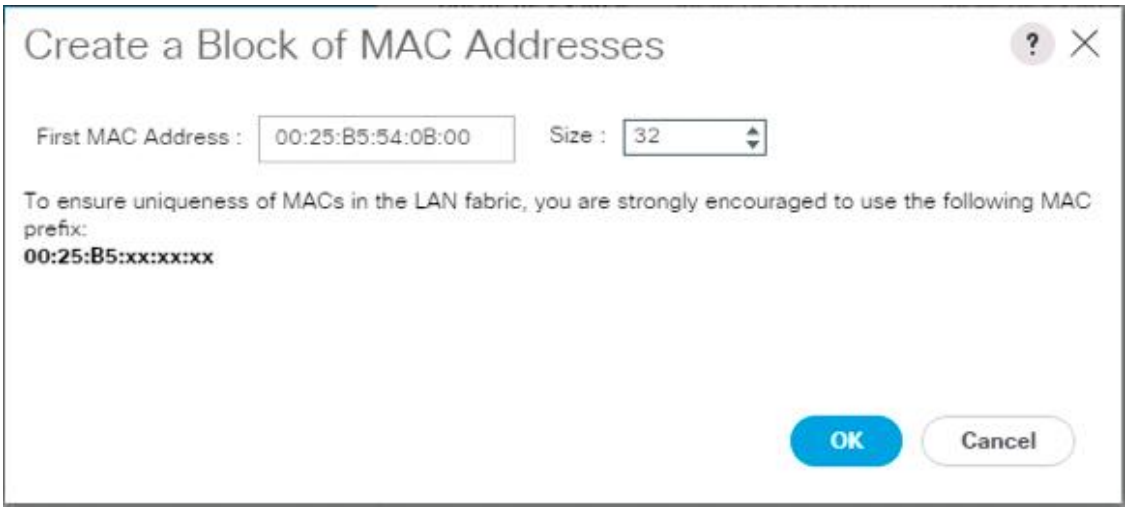

23. Click OK.

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.

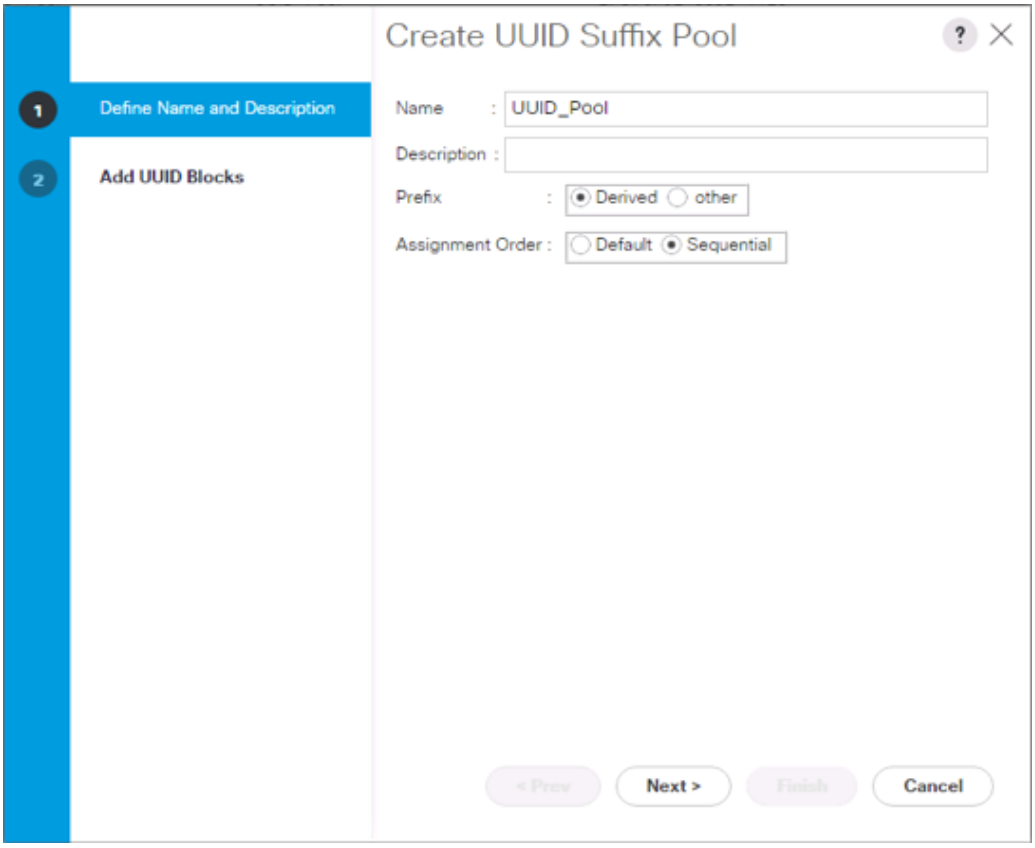

6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.

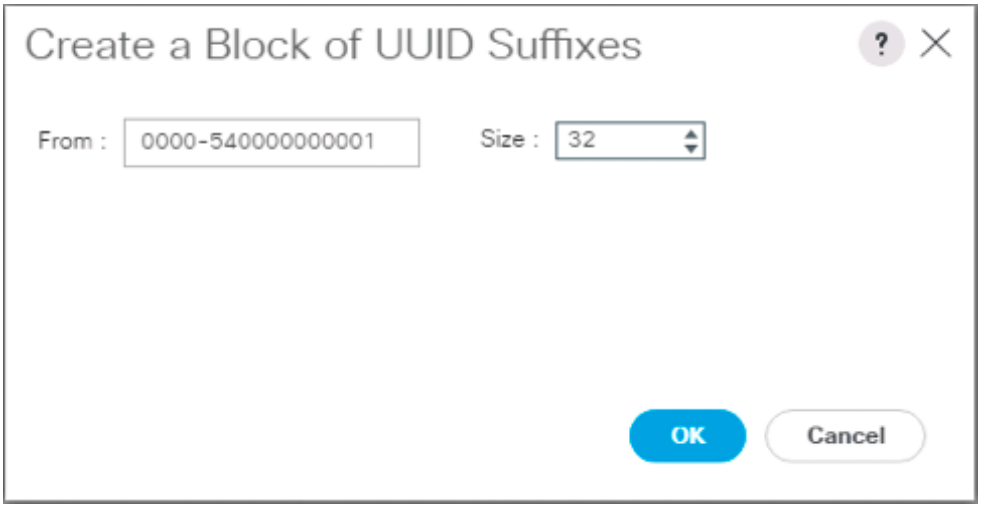

The starting From number (0000-54) has been adjusted to give it a differentiator from other UCS domains that may be adjacent.

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.

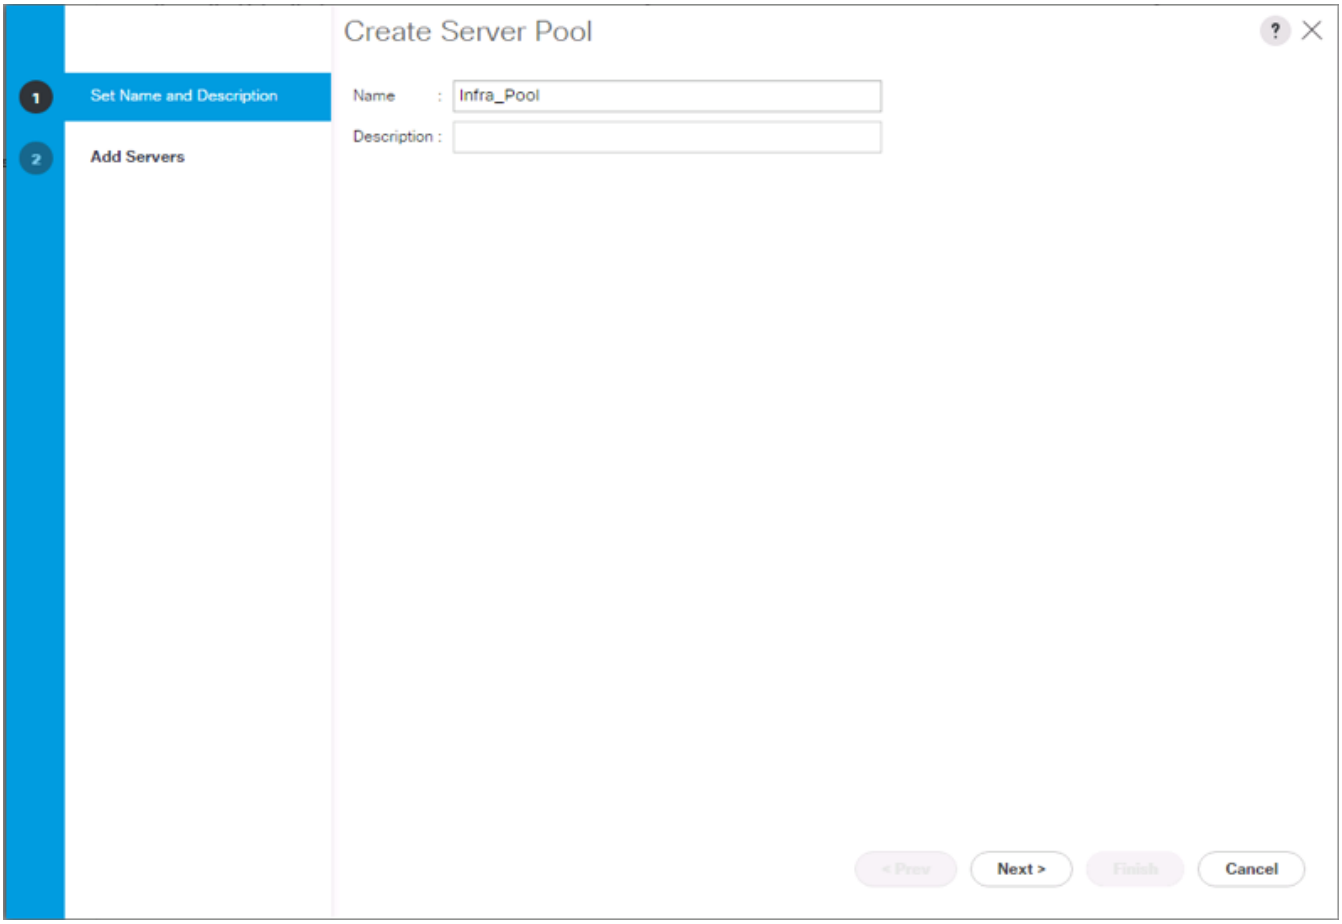

6. Optional: Enter a description for the server pool.

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.

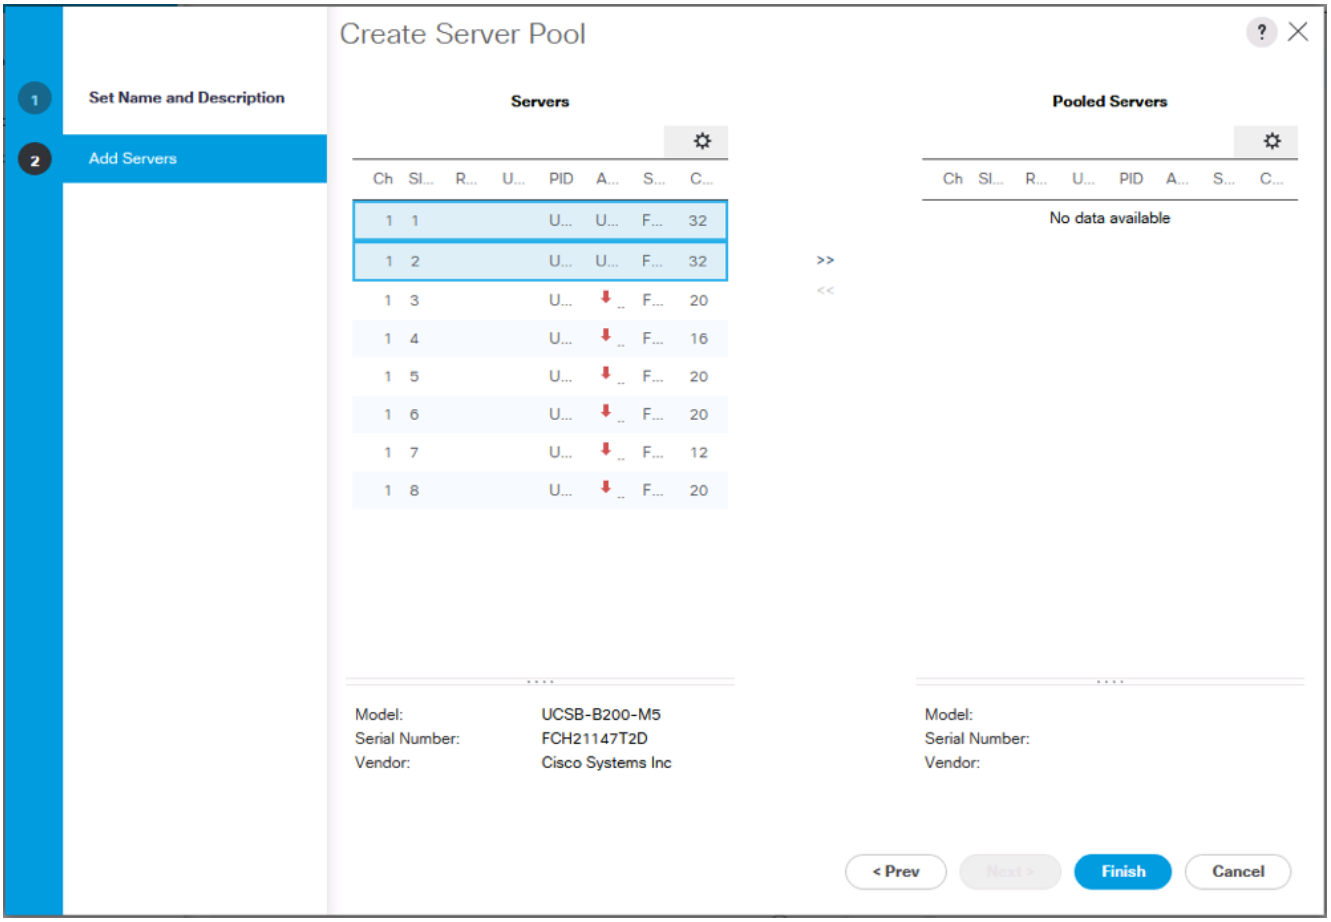

9. Click Finish.

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.

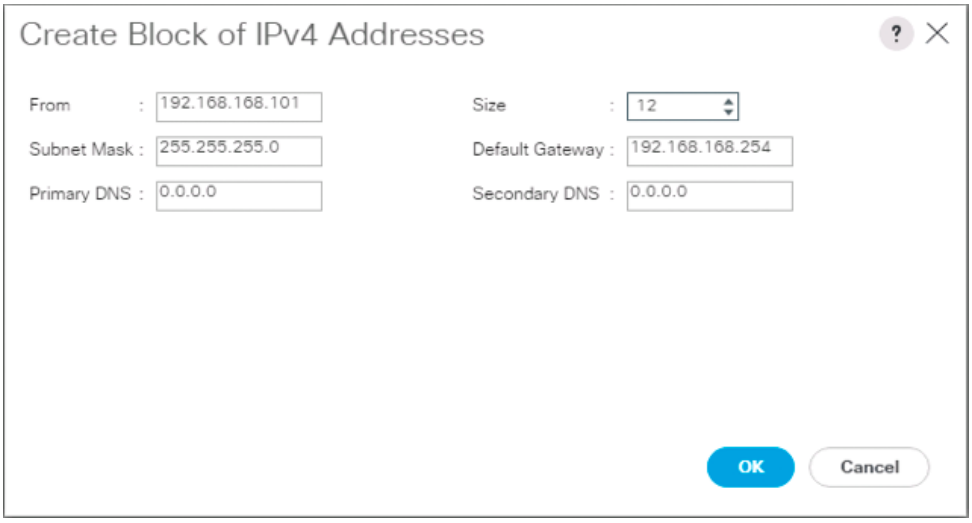

4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

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 on the left.

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.

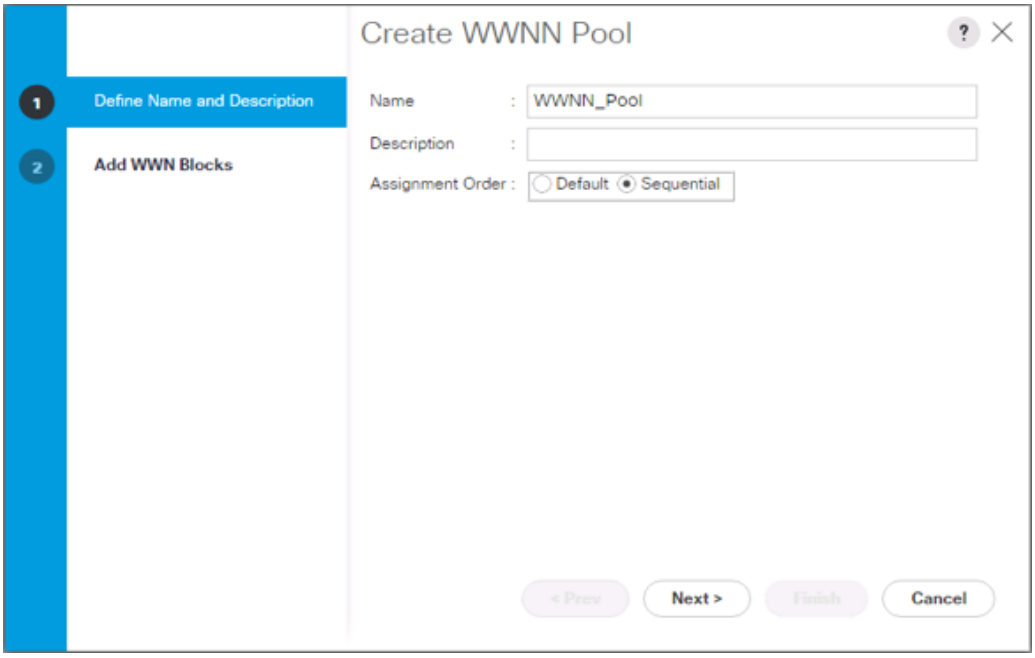

8. Click Next.

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.

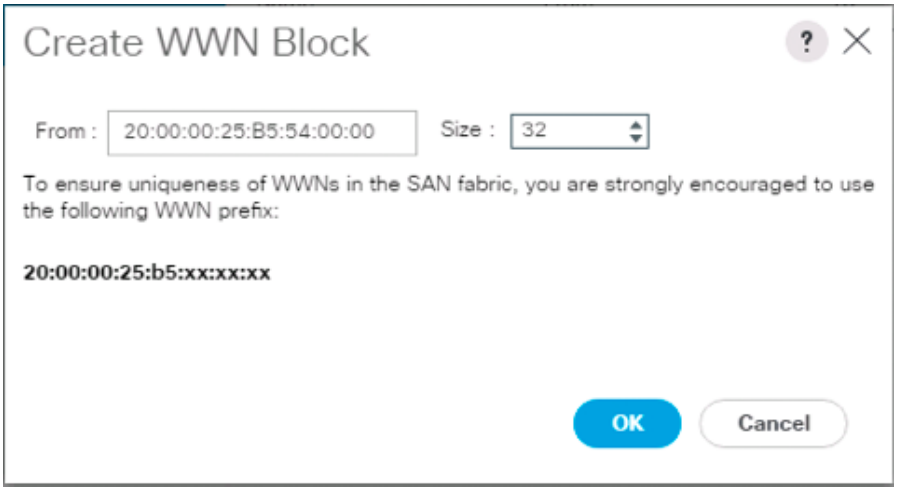

12. Click OK.
13. Click Finish to create the WWNN Pool.
14. Click OK.

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.

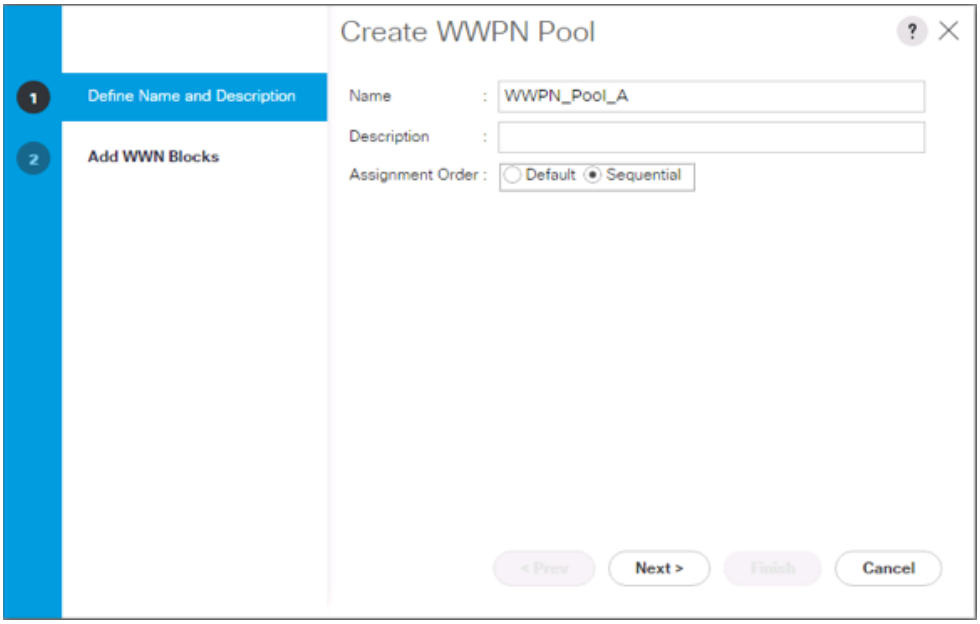

9. Click Next.

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.

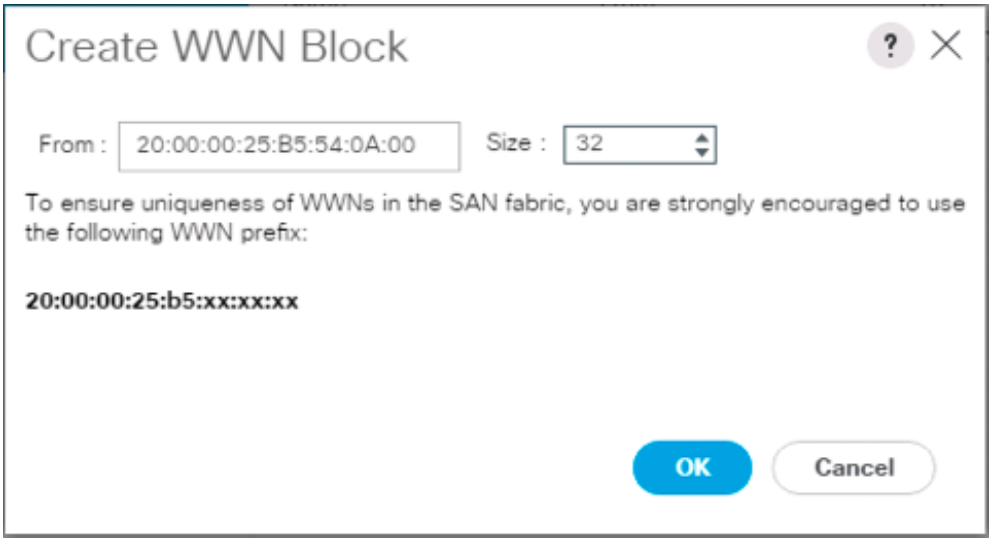

13. Click OK.

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.

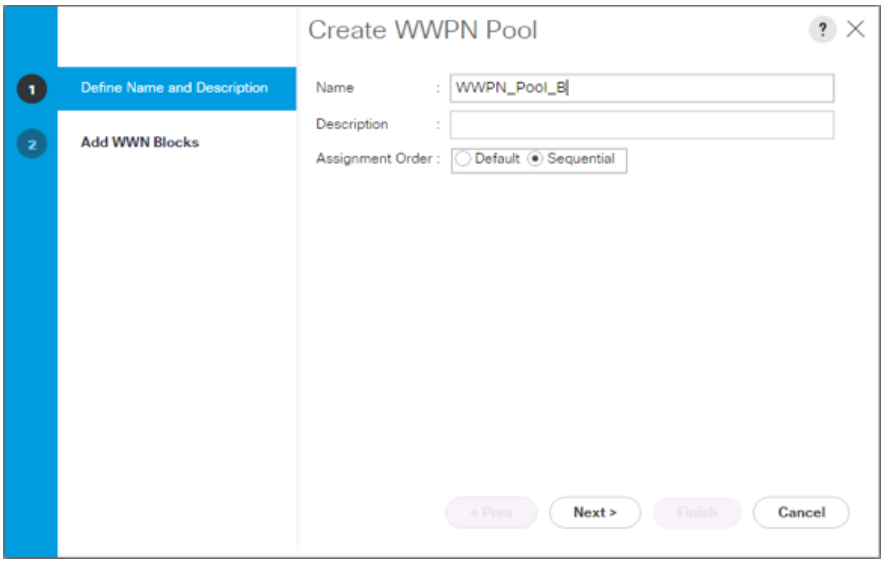

21. Click Next.

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.

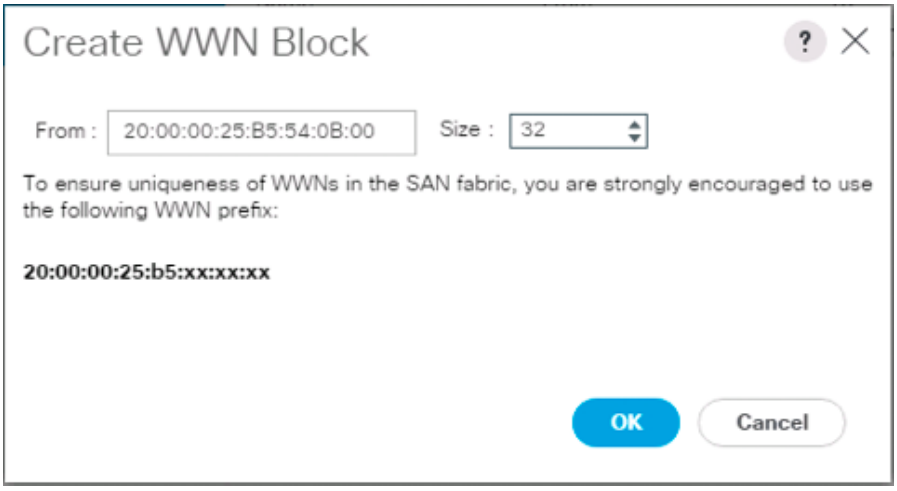

25. Click OK.

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(1b)B for the Blade Package, and optionally set version 4.0(1b)C for the Rack Package.

7. Leave Excluded Components with only Local Disk selected.

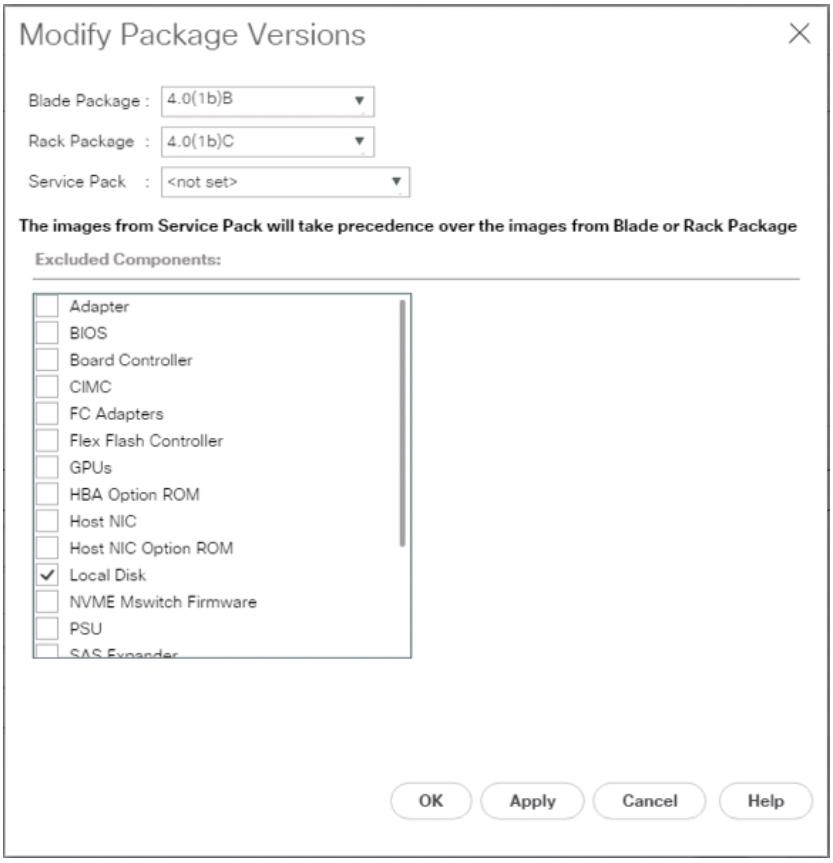

8. Click OK to modify the host firmware package and OK again to acknowledge the changes.

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.

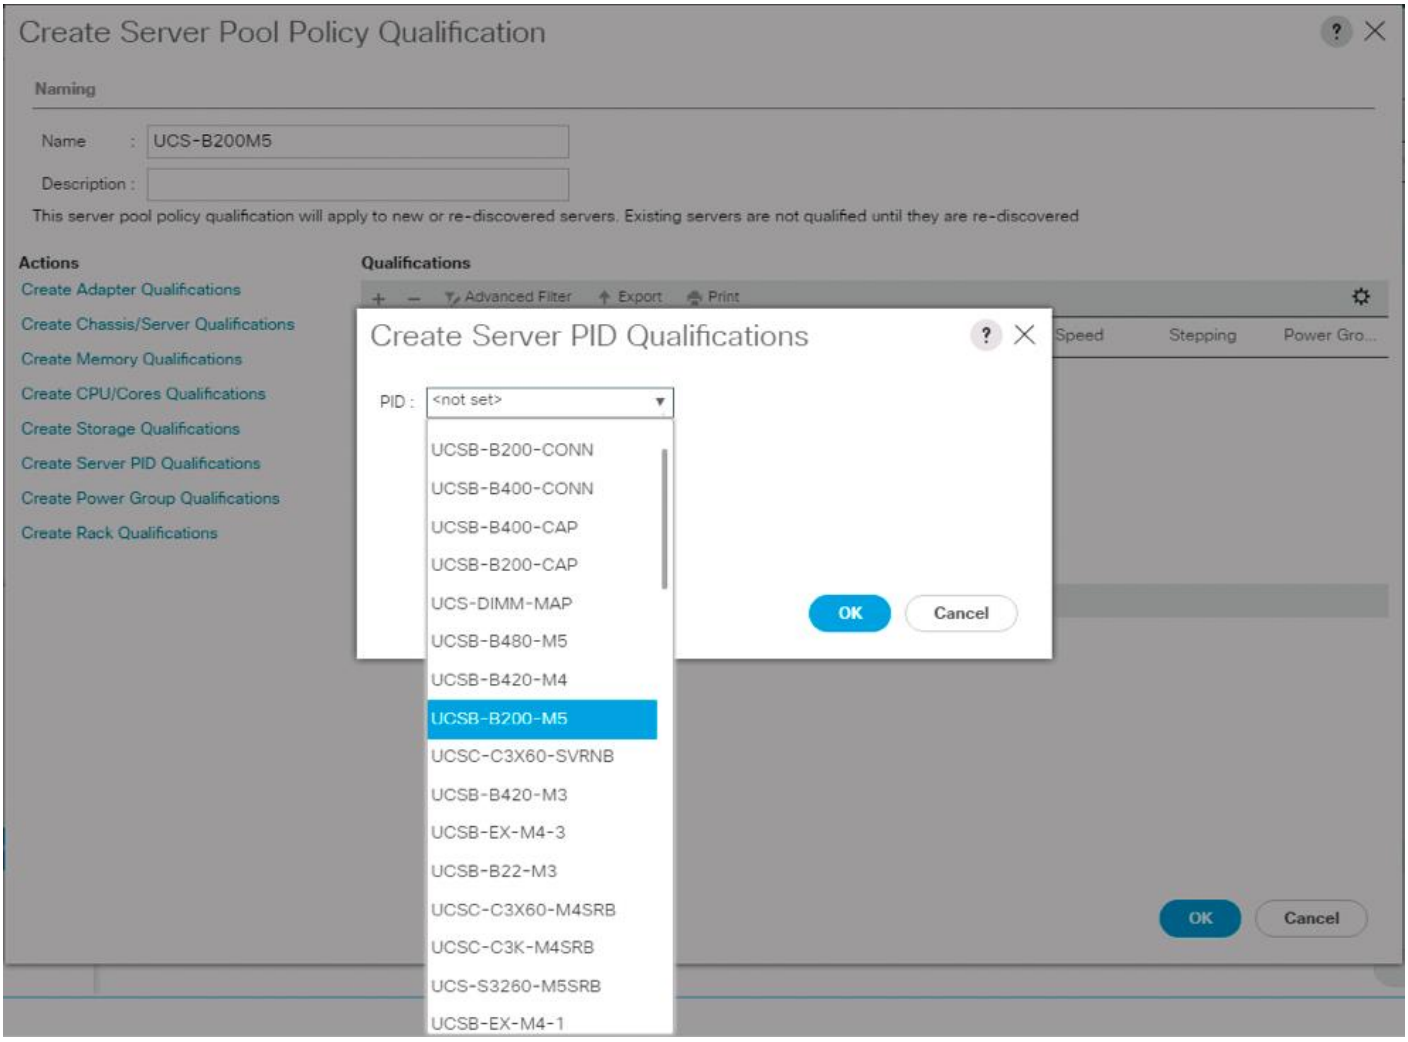

8. Click OK.

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.

Place the Cisco Custom Image VMware ESXi 6.7 U1 ISO on the HTTP server and follow these steps to create a vMedia Policy:

1. In Cisco UCS Manager, select Servers on the left.
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.

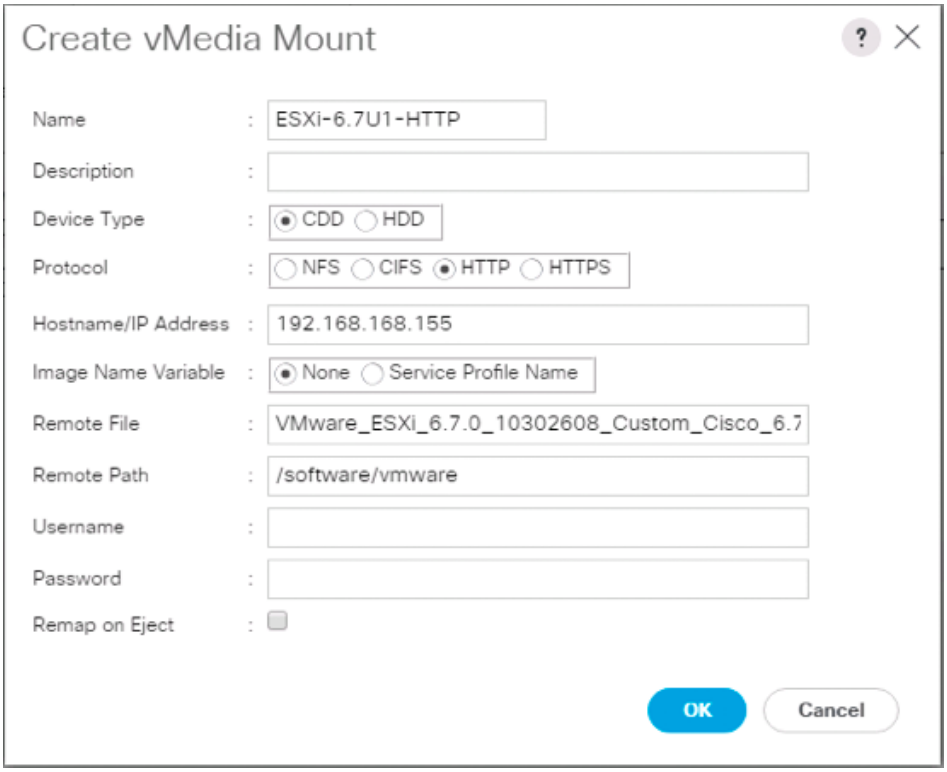

15. Click OK to create the vMedia Mount.

16. Click OK then 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 on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `VM-Host` as the BIOS policy name.

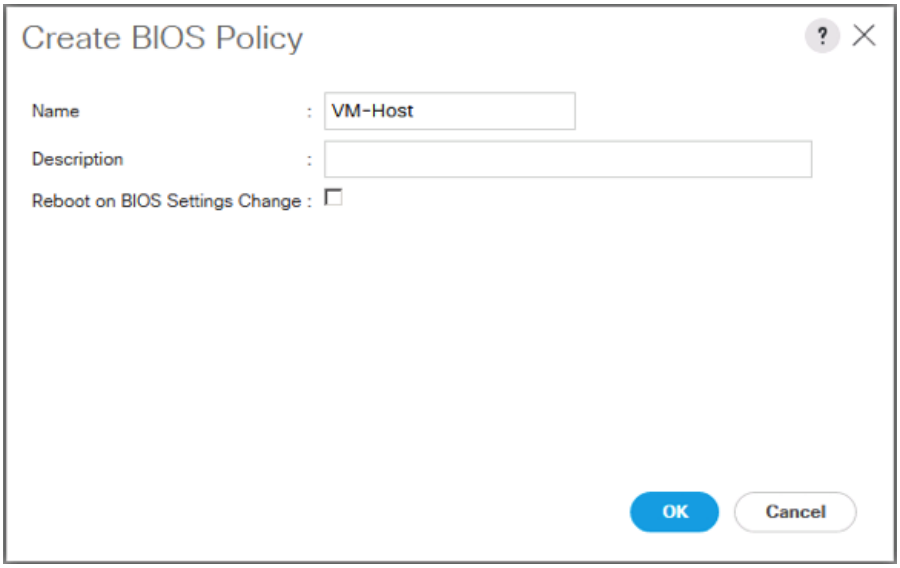

6. Select and right-click the newly created BIOS Policy.

7. Within the Main tab of the Policy:

   a. Change CDN Control to enabled.

   b. Change the Quiet Boot setting to disabled.

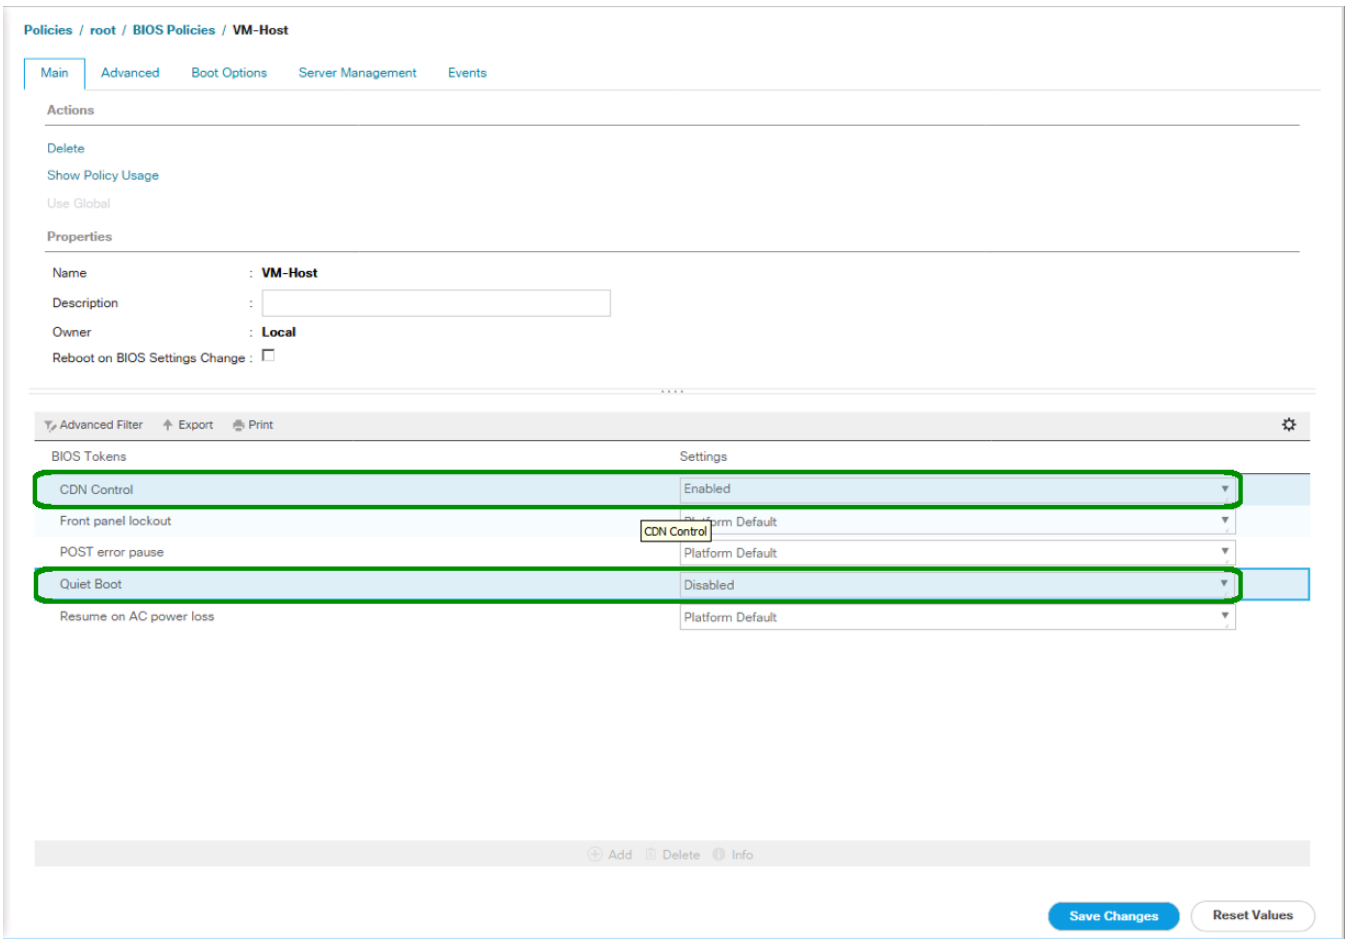

8. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab.

9. Set the following within the Processor tab:

   a. DRAM Clock Throttling -> Performance

   b. Frequency Floor Override -> Enabled

   c. Processor C State -> Disabled

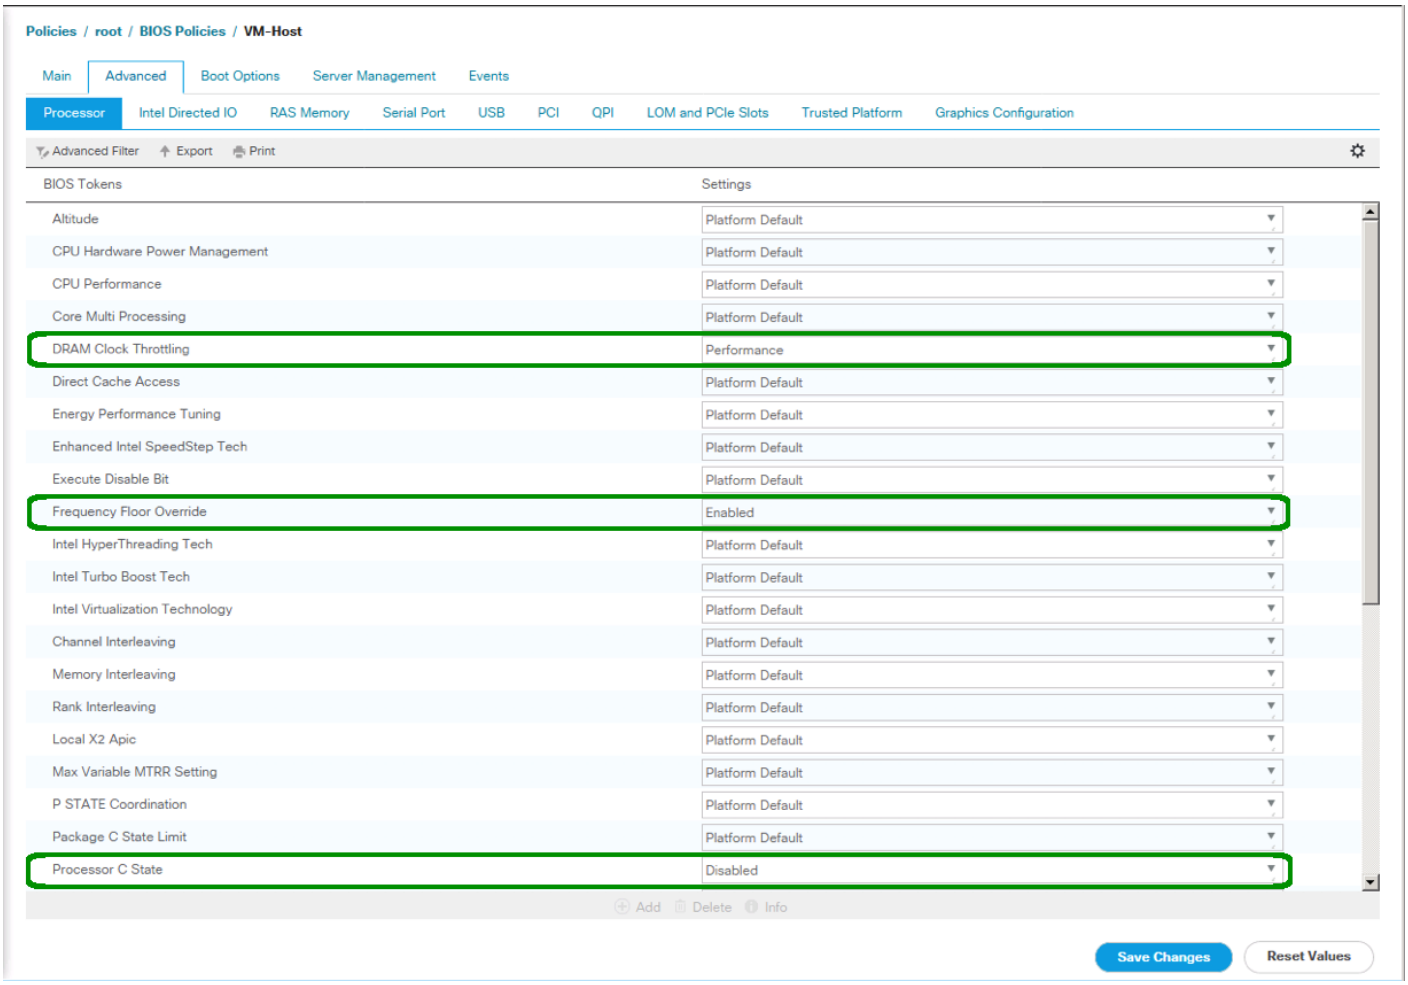

10. Scroll down to the remaining Processor options and select:

   a. Processor C1E -> disabled

   b. Processor C3 Report -> disabled

   c. Processor C7 Report -> disabled

   d. Energy Performance -> performance

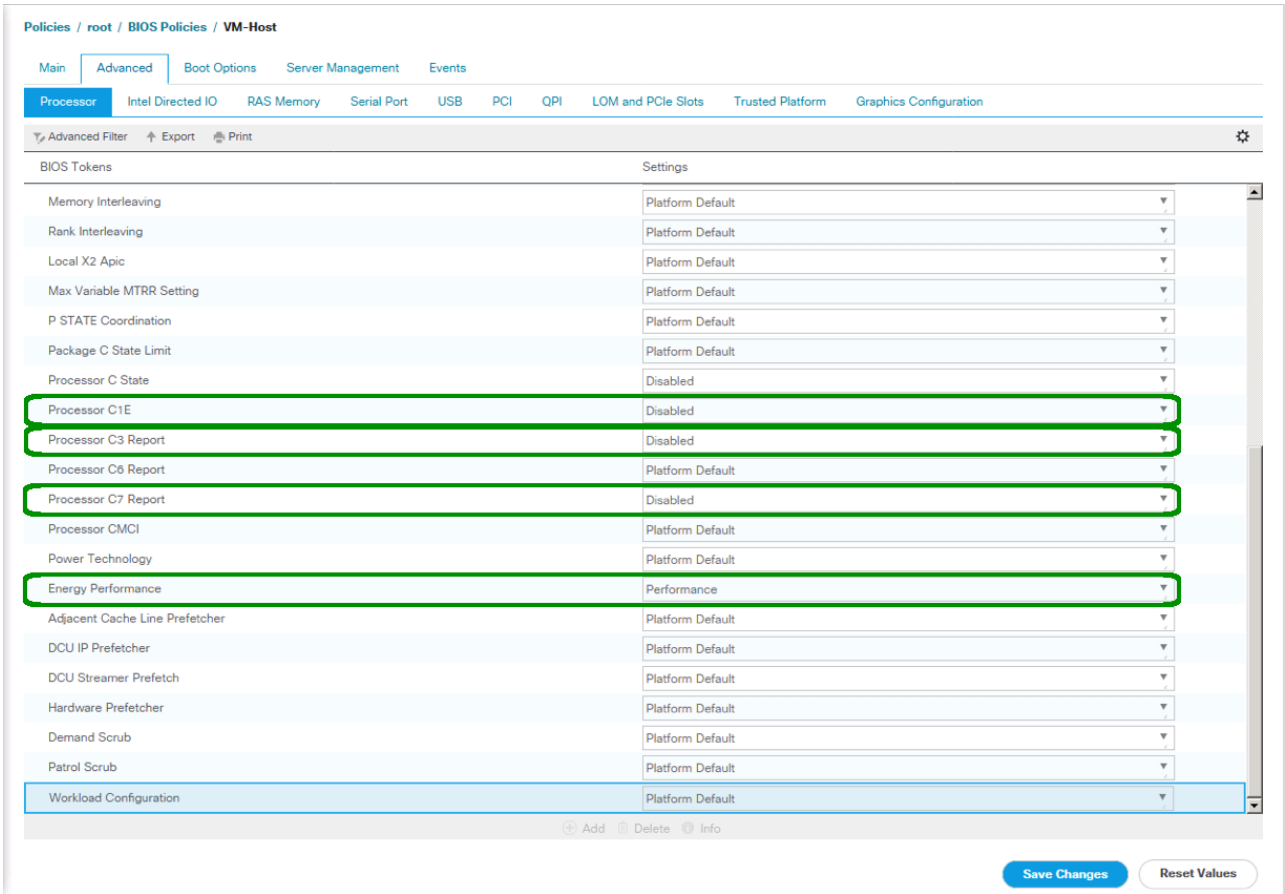

11. Click the RAS Memory tab and select:

    a. LV DDR Mode -> performance-mode

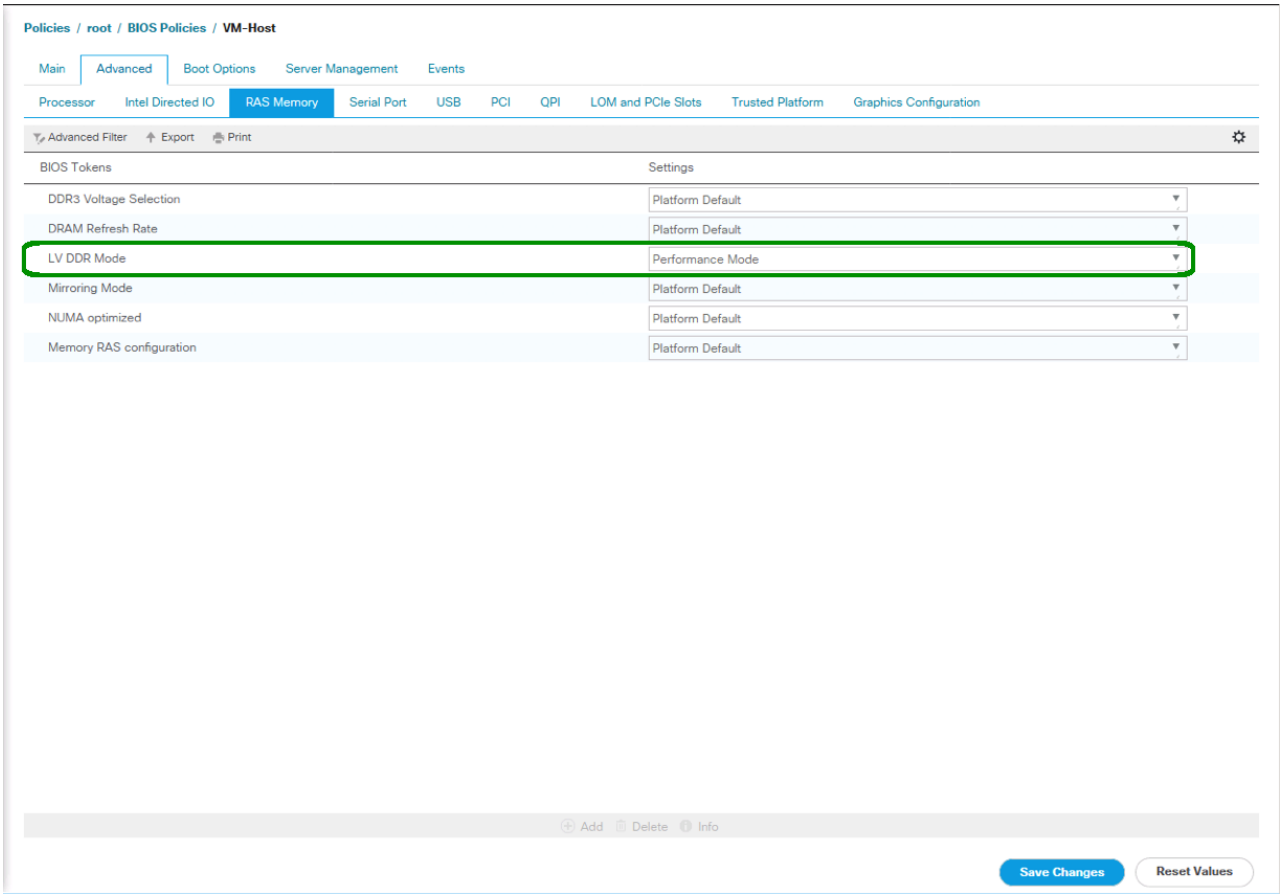

12. Click Save Changes.

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).

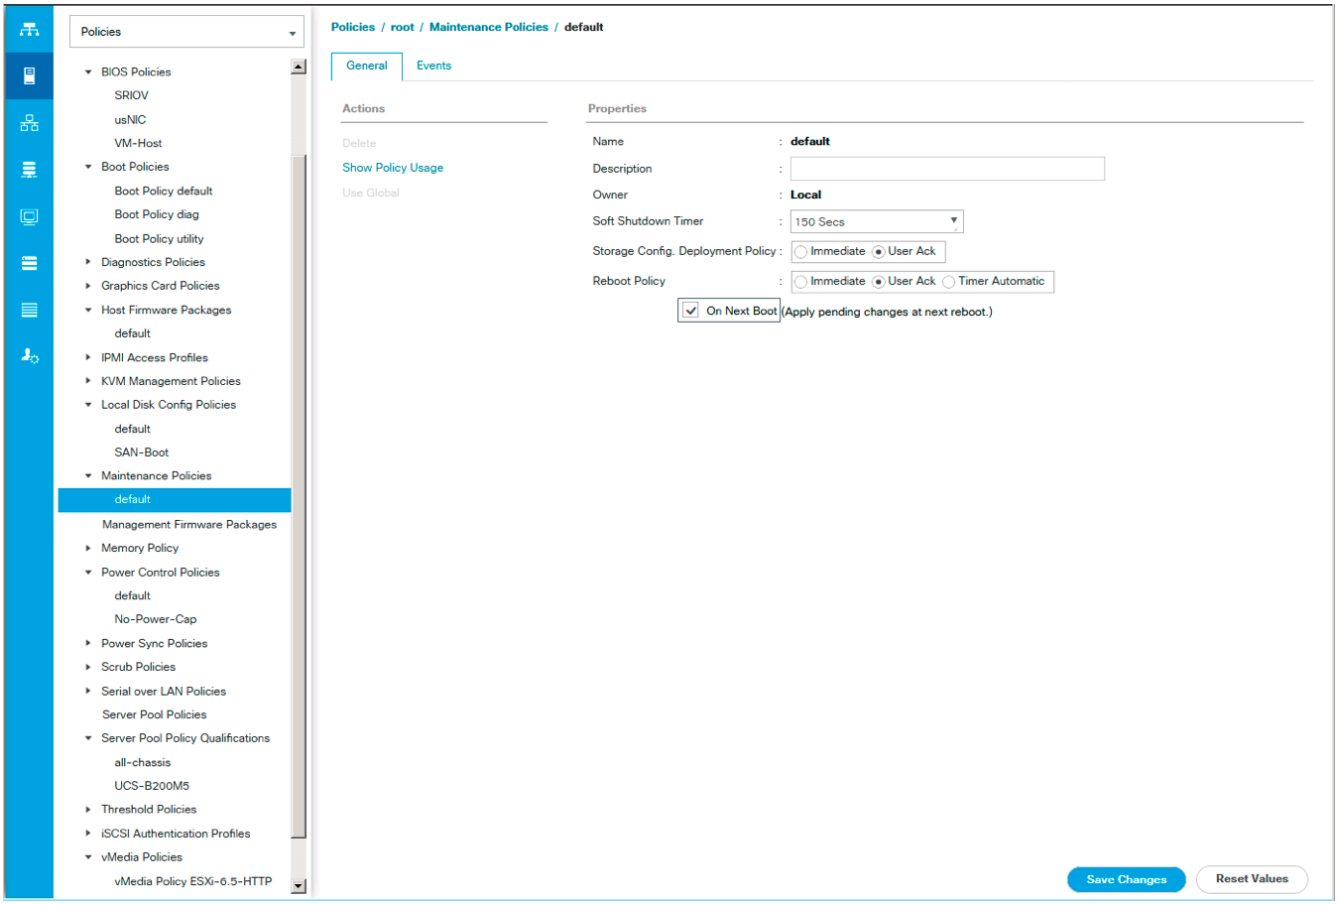

6. Click Save Changes.

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.

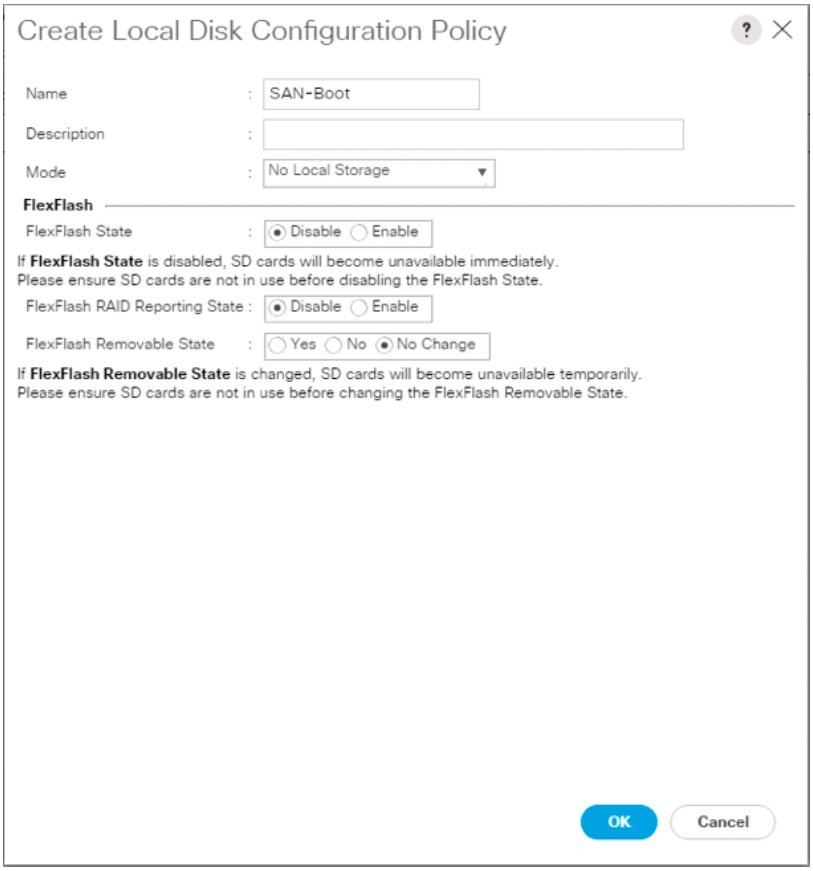

8. Click OK.

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.

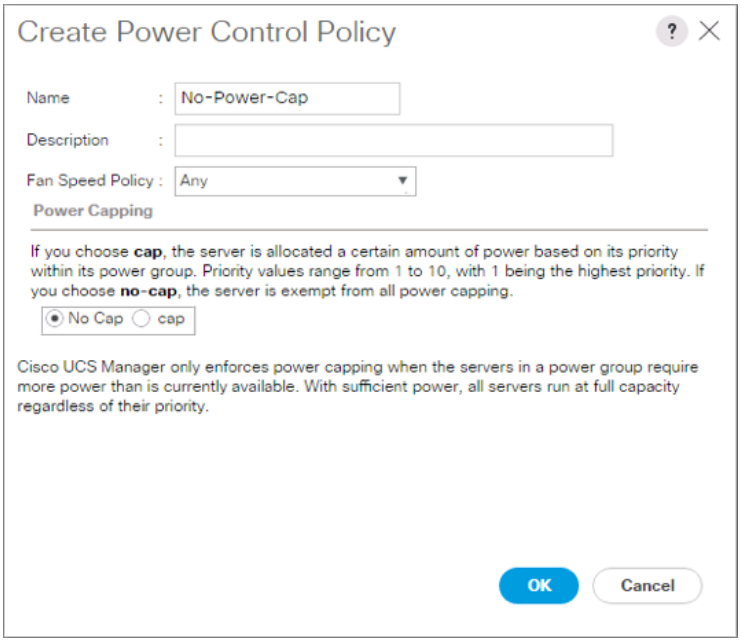

7. Click OK to create the power control policy.

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.

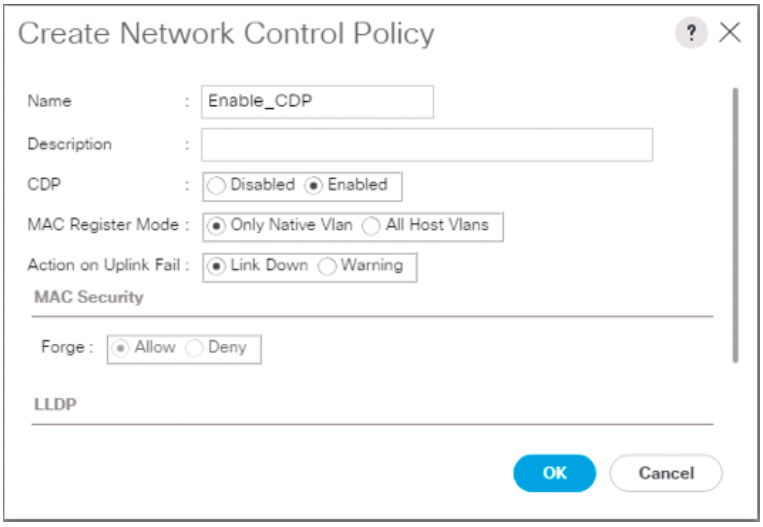

8. Click OK.

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.

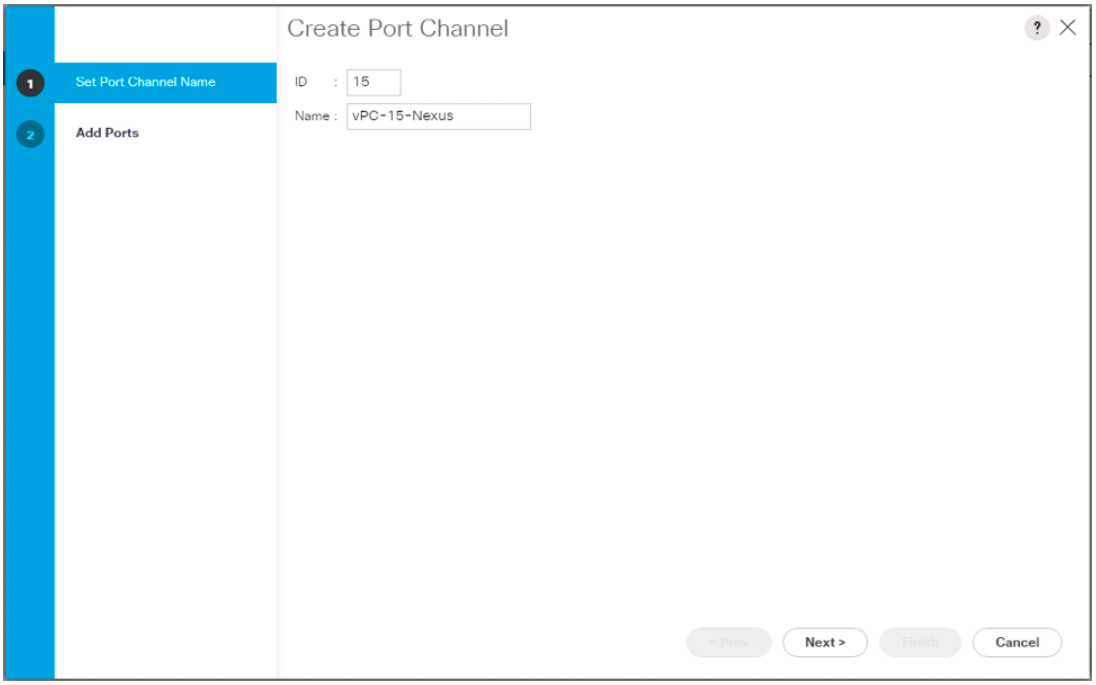

7. Click Next.

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

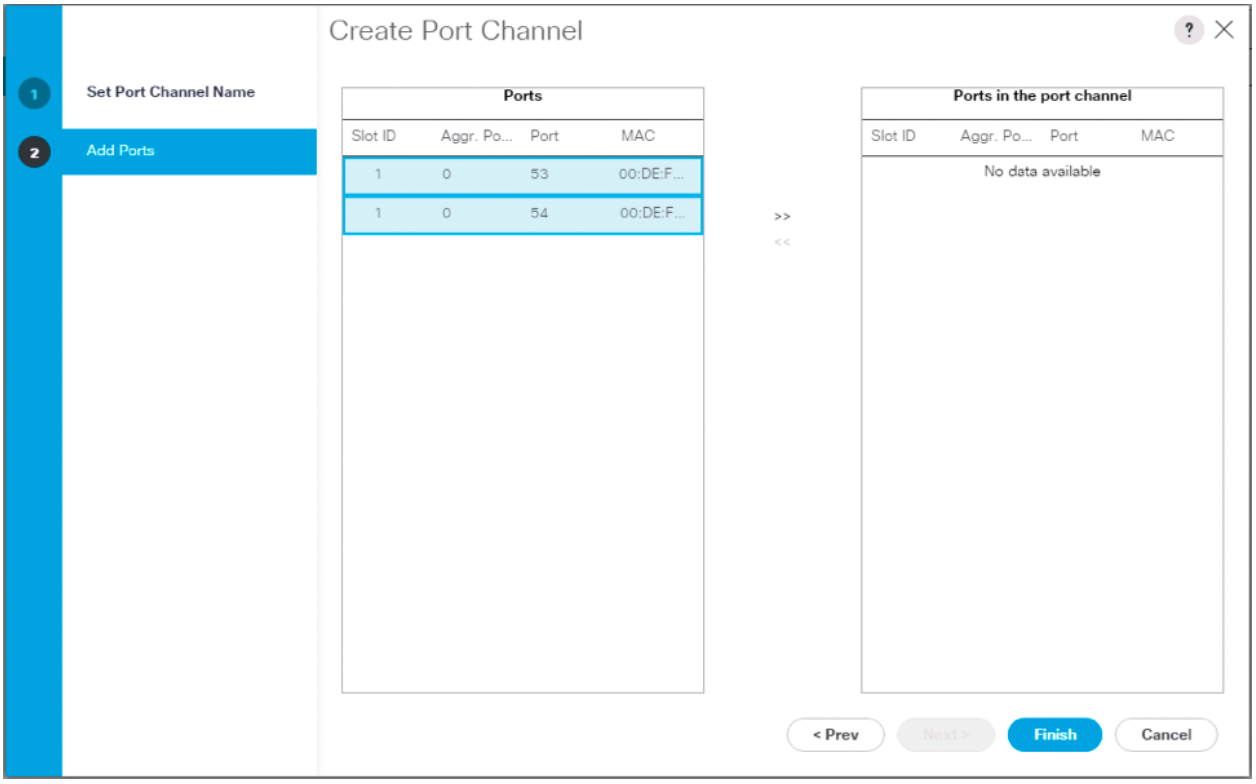

9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter a unique ID for the port channel, (`16` in 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.

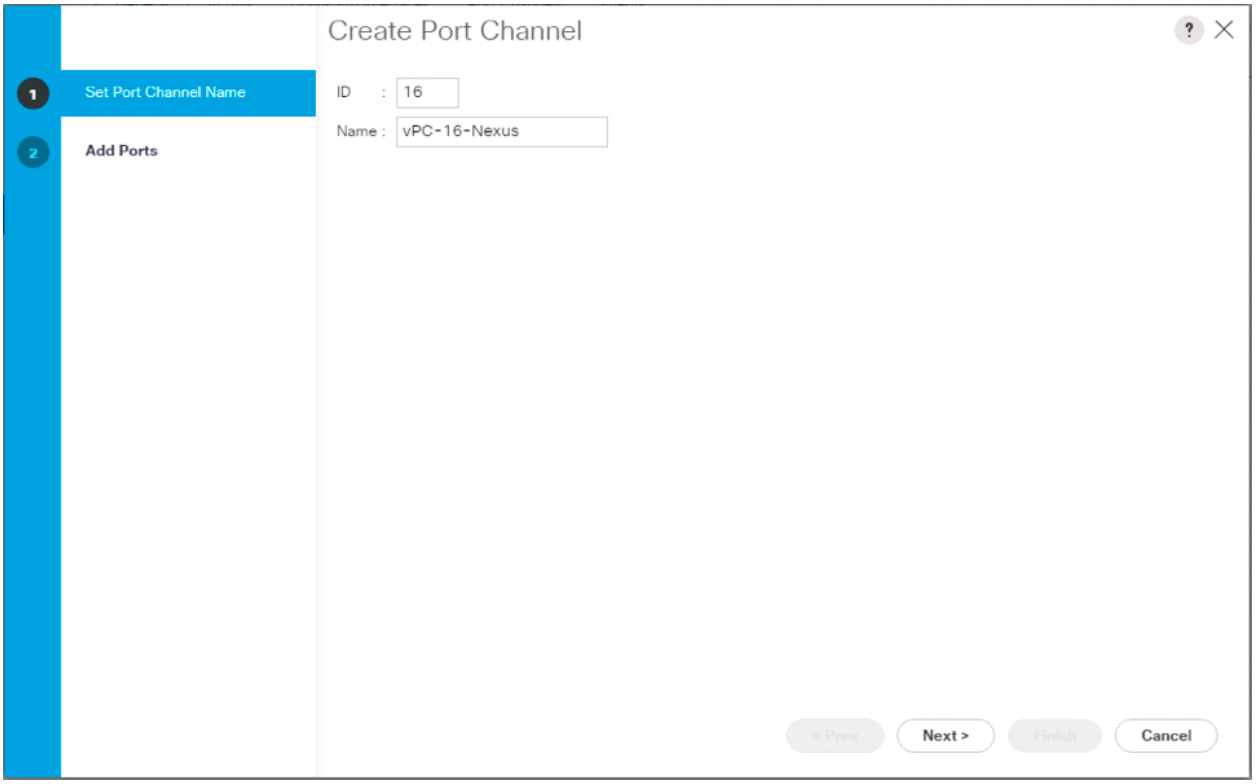

17. Click Next.
18. Select the following ports to be added to the port channel:
    a. Slot ID 1 and port 53
    b. Slot ID 1 and port 54

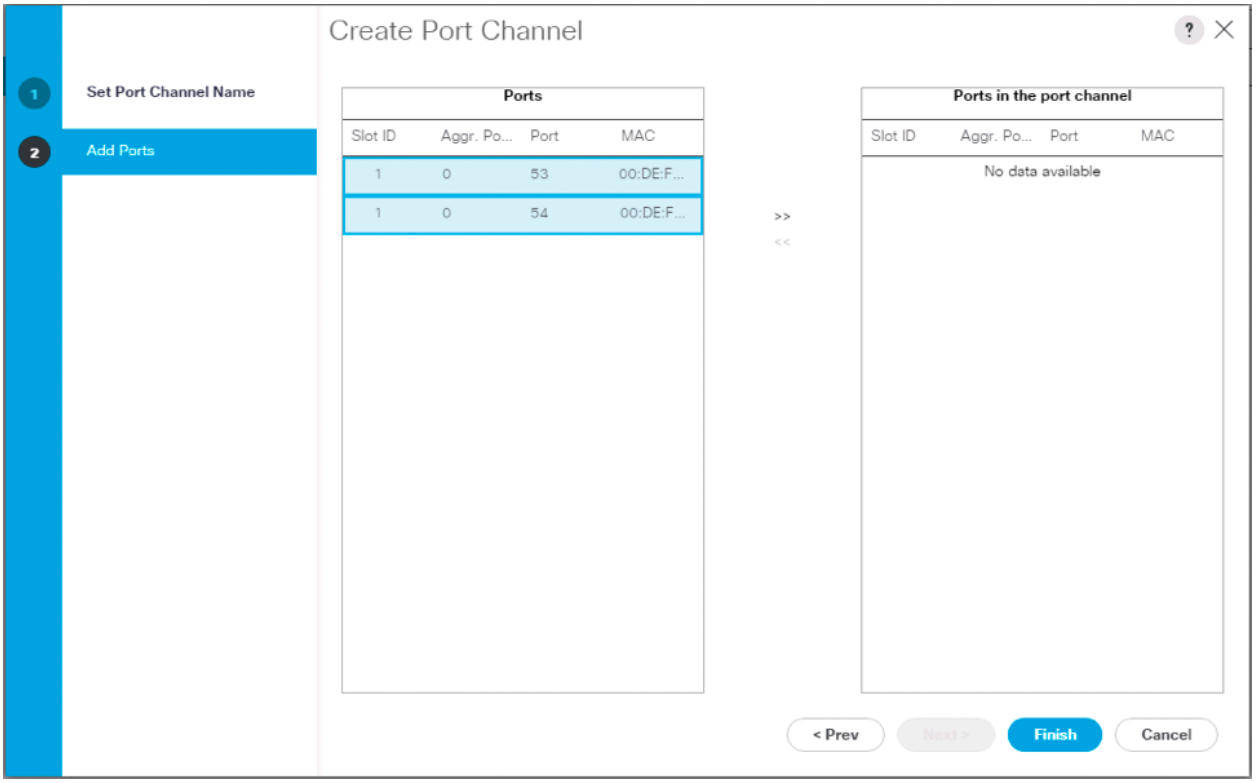

19. Click >> to add the ports to the port channel.

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.

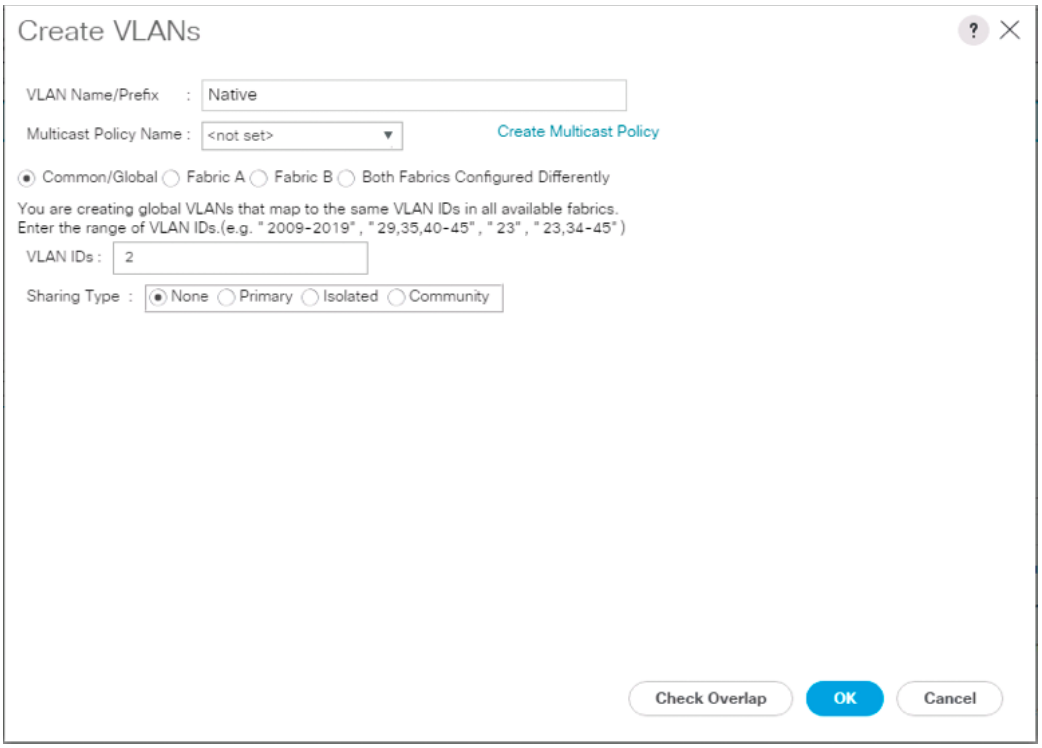

9. Click OK and then click OK again.

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.

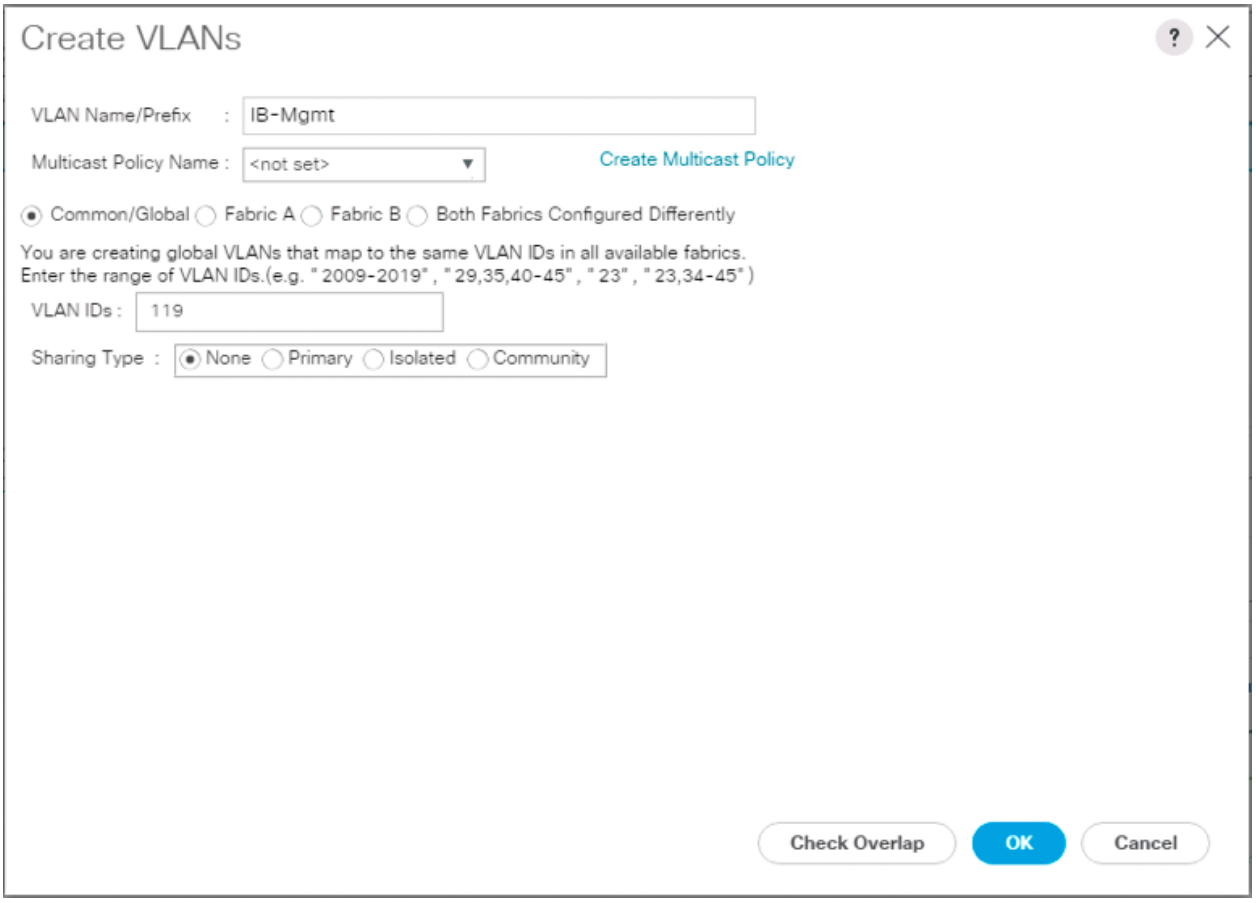

18. Click OK and then click OK again.

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.

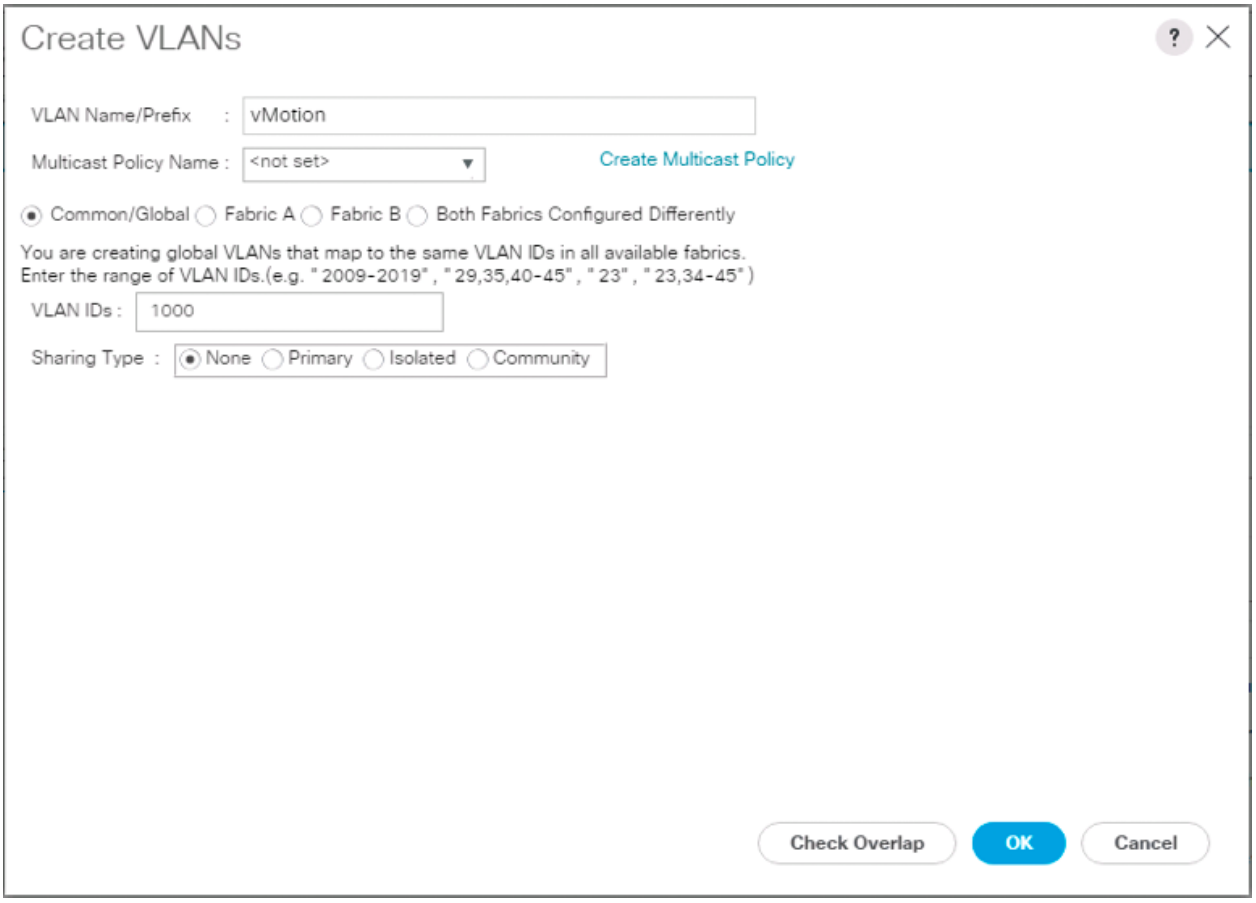

25. Click OK and then click OK again.

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.

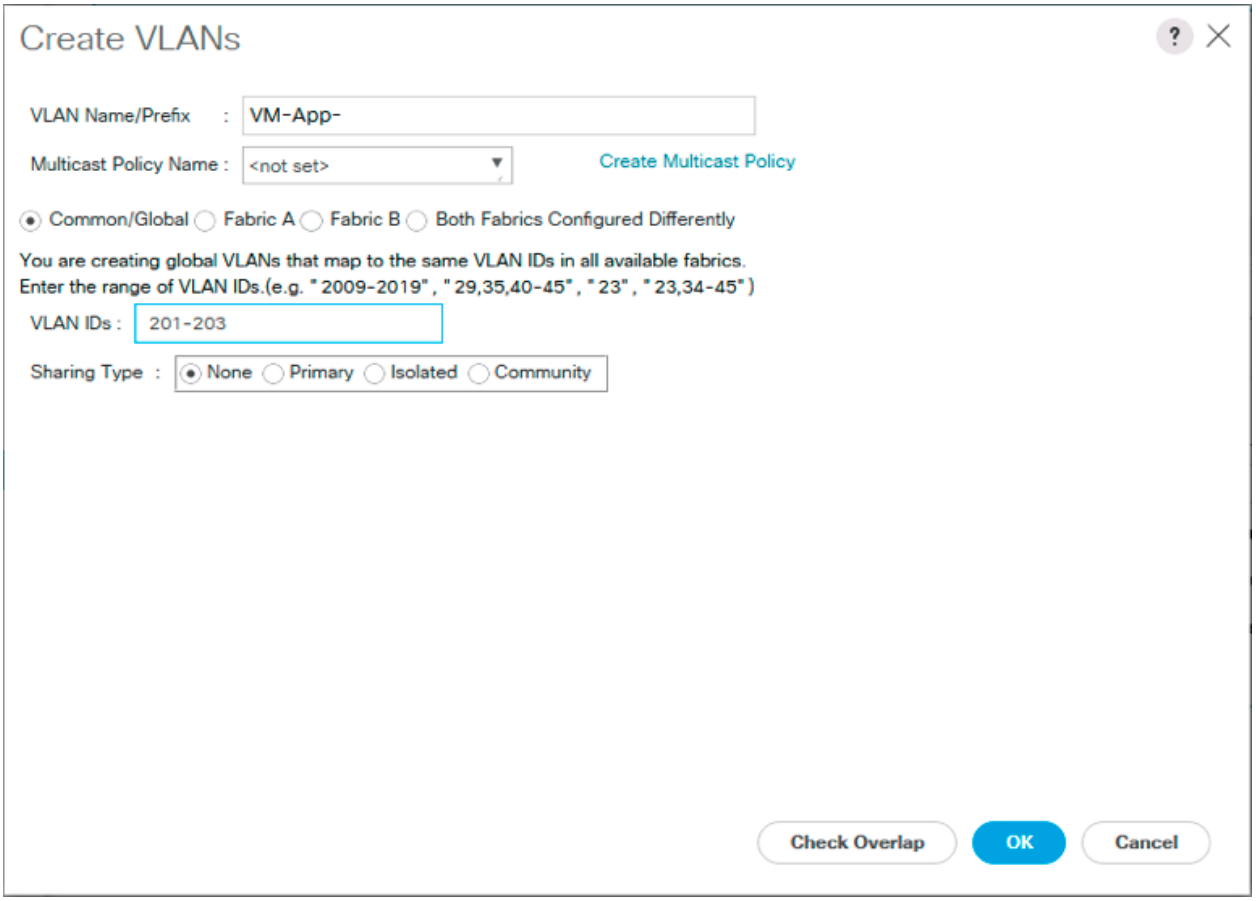

32. Click OK and then click OK again.

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.

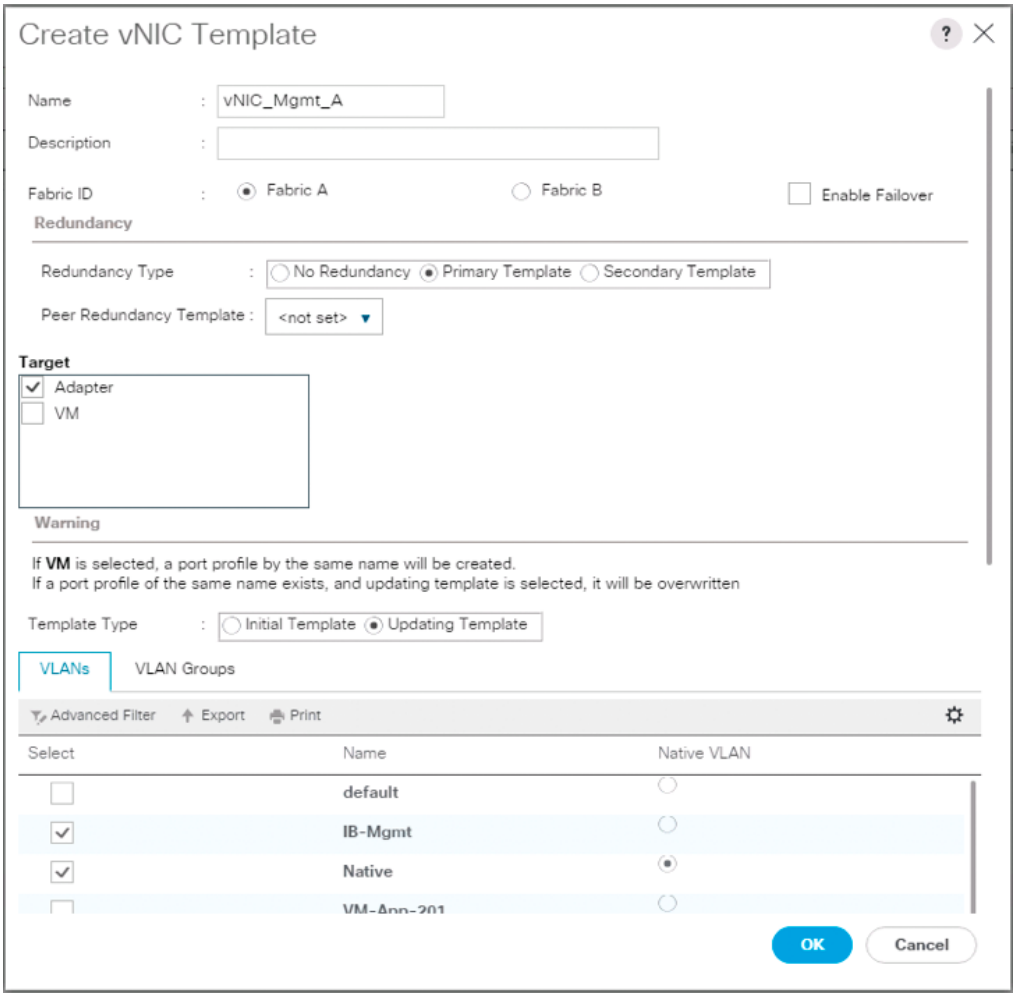

12. Set `Native` as the native VLAN.

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`.

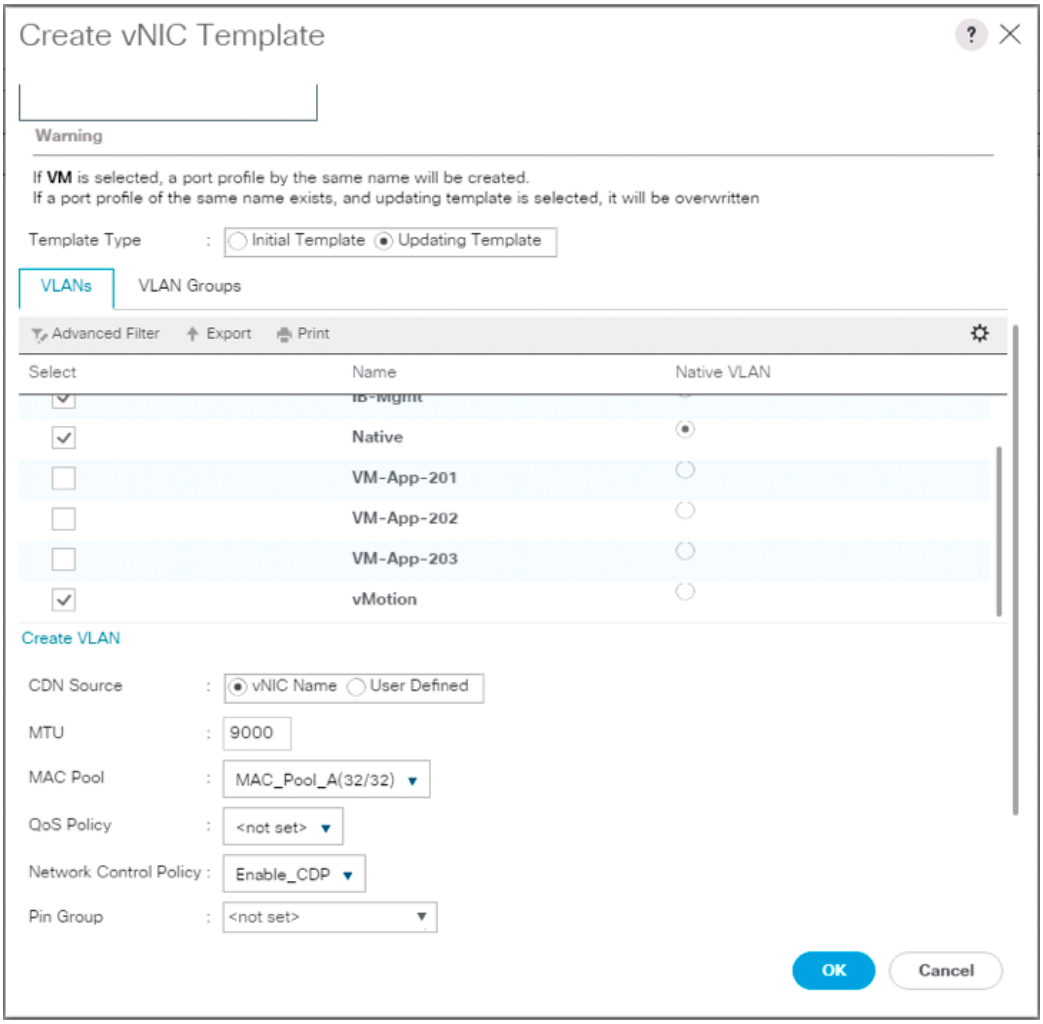

17. Click OK to create the vNIC template.

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.

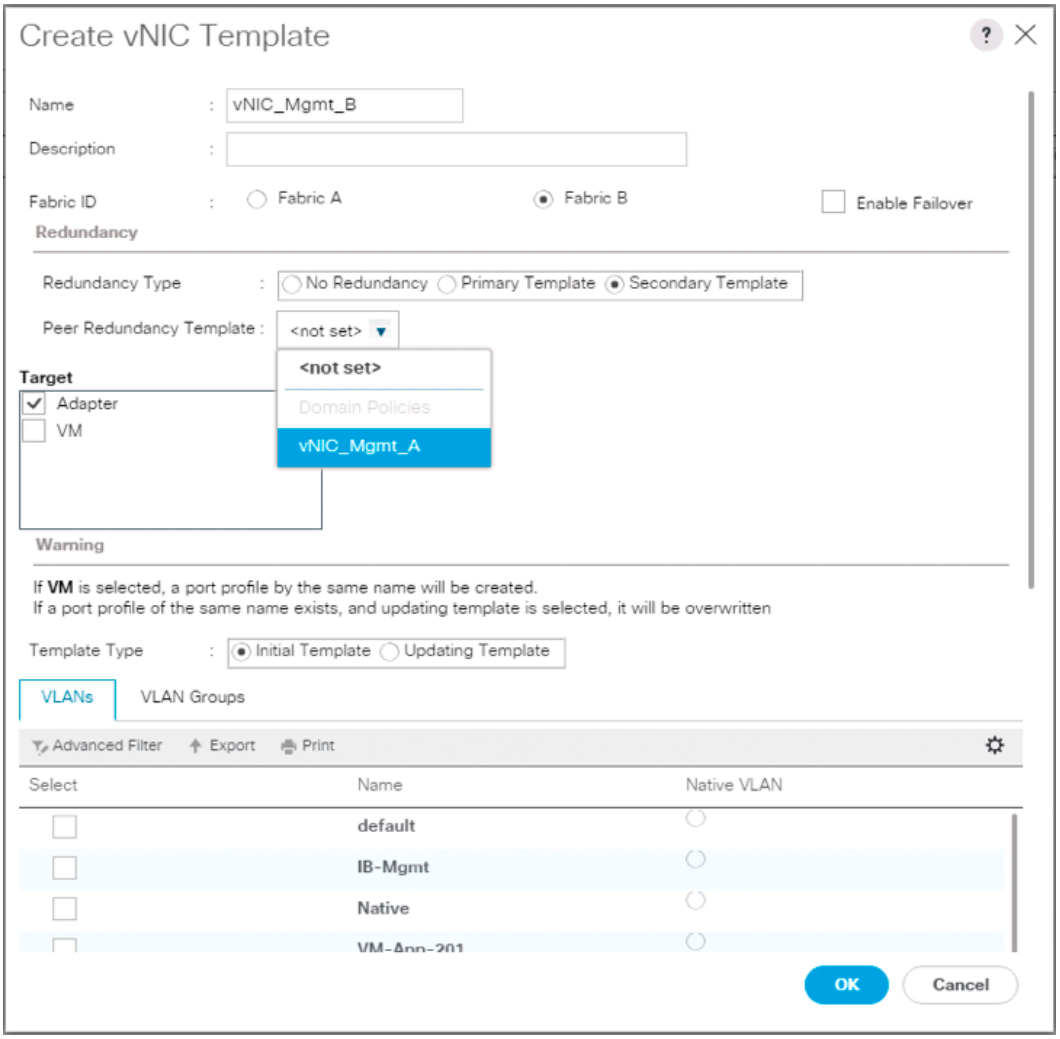

10. In the MAC Pool list, select MAC_Pool_B.

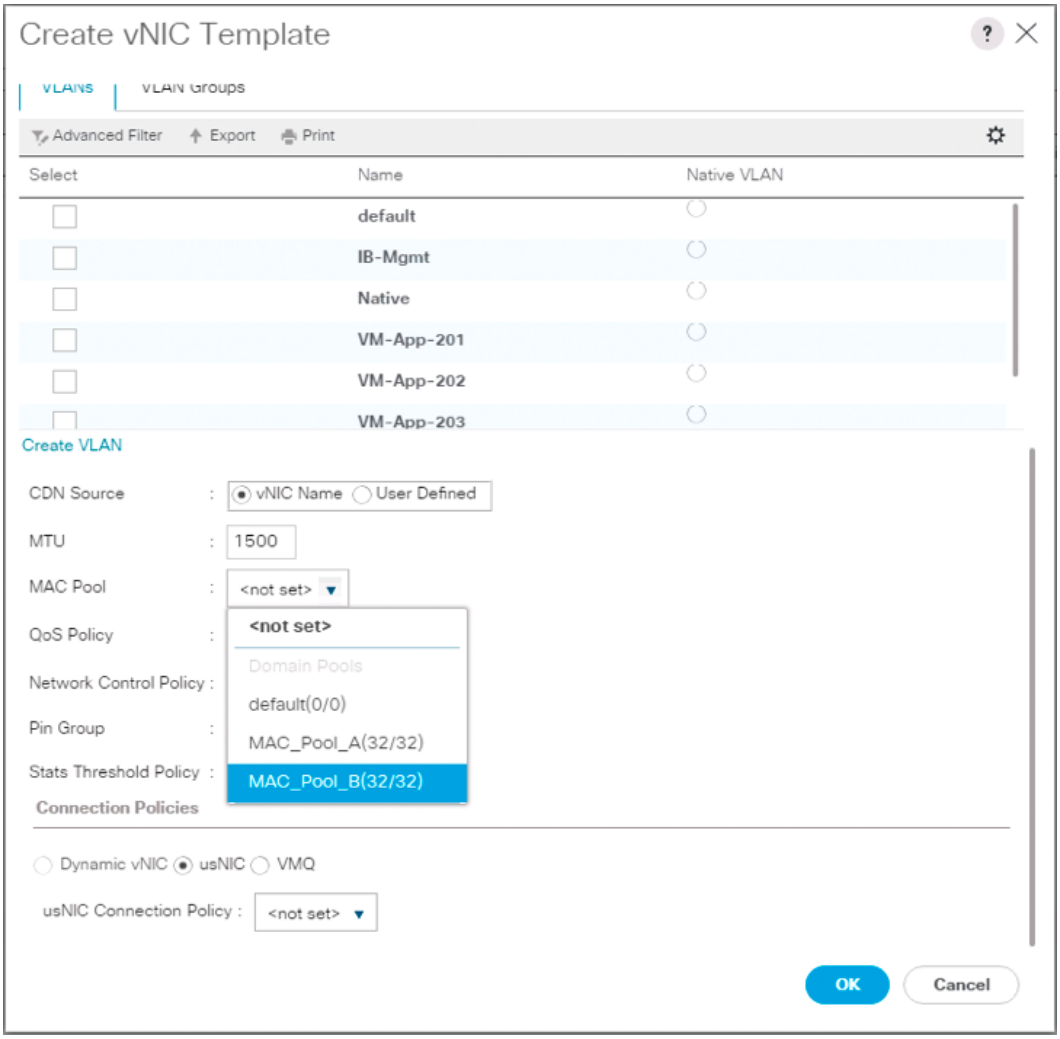

11. Click OK to create the vNIC template.

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.

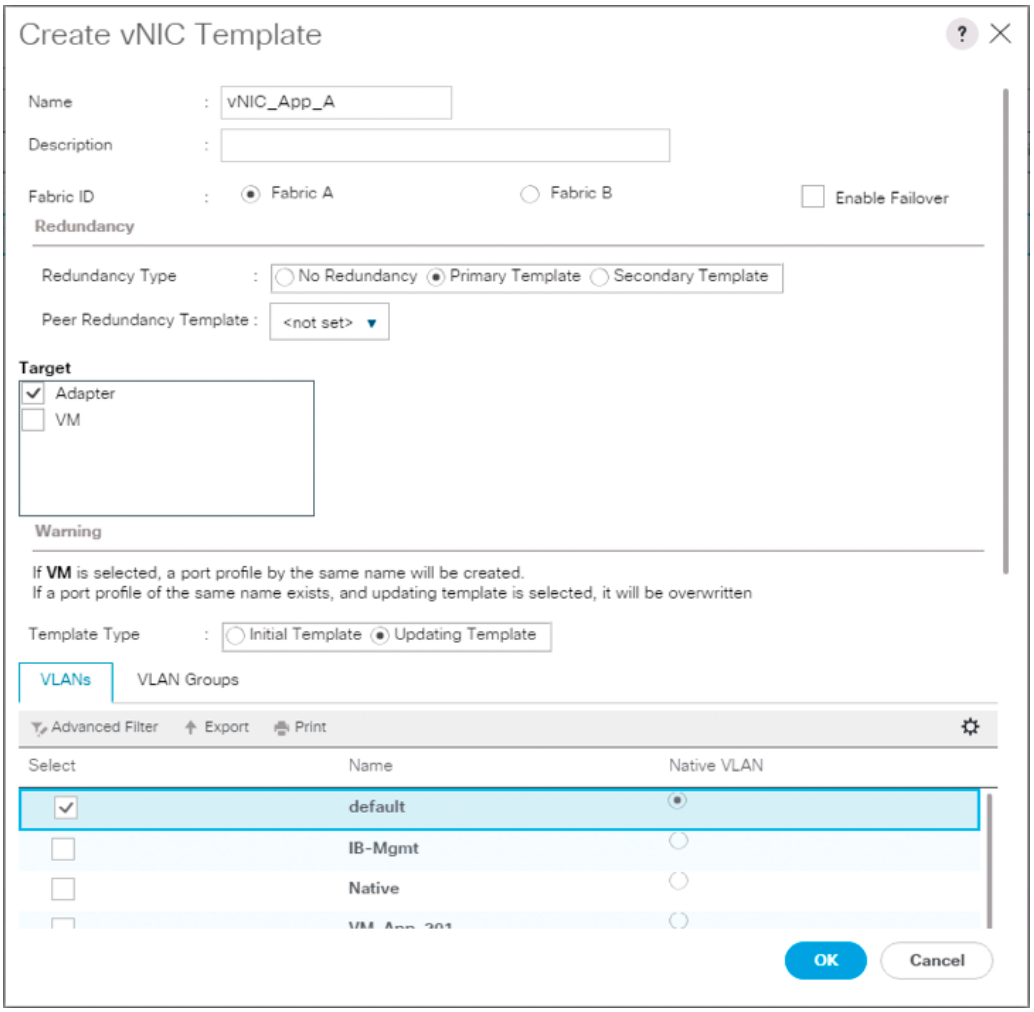

12. Under VLANs, select the checkboxes for any application or production VLANs that should be delivered to the ESXi hosts.

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`.

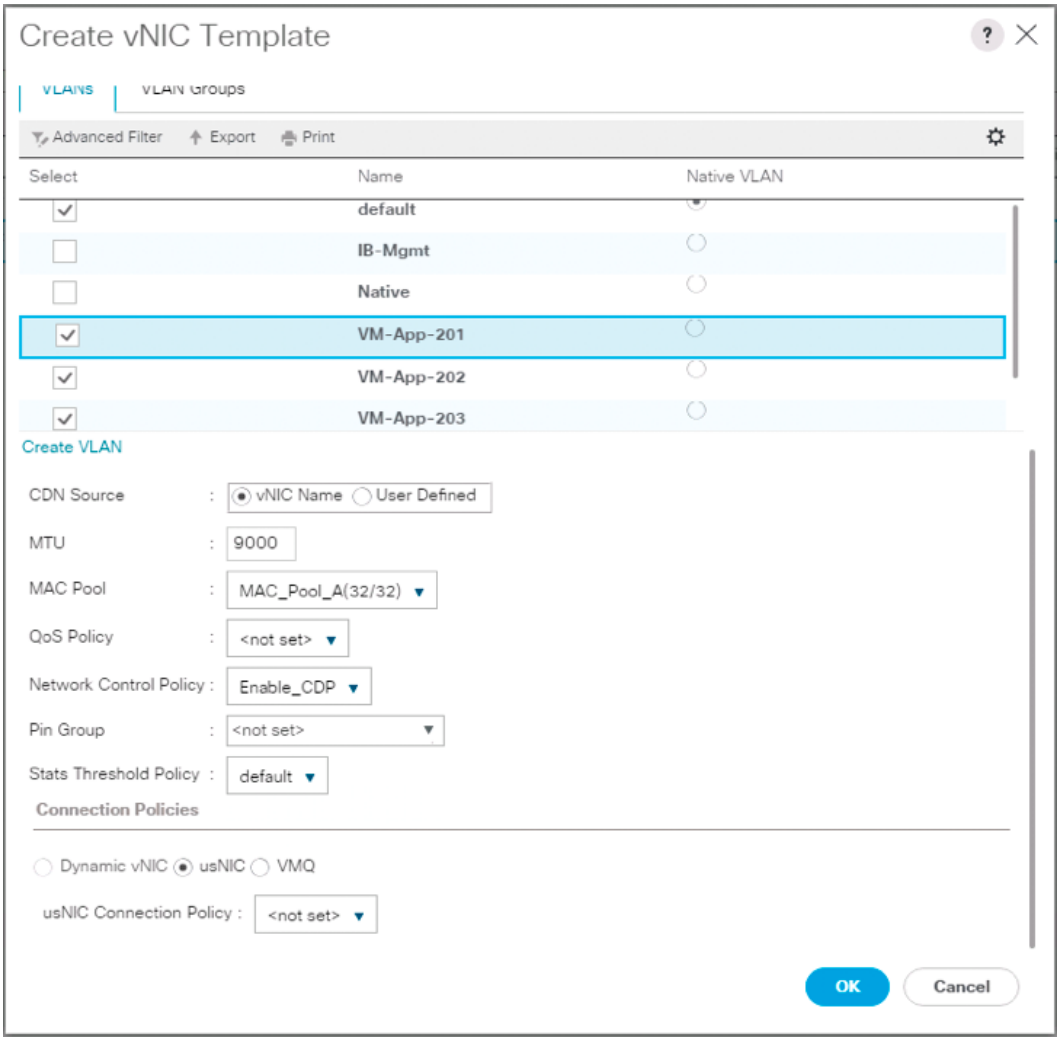

16. Click OK to create the vNIC template.

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.

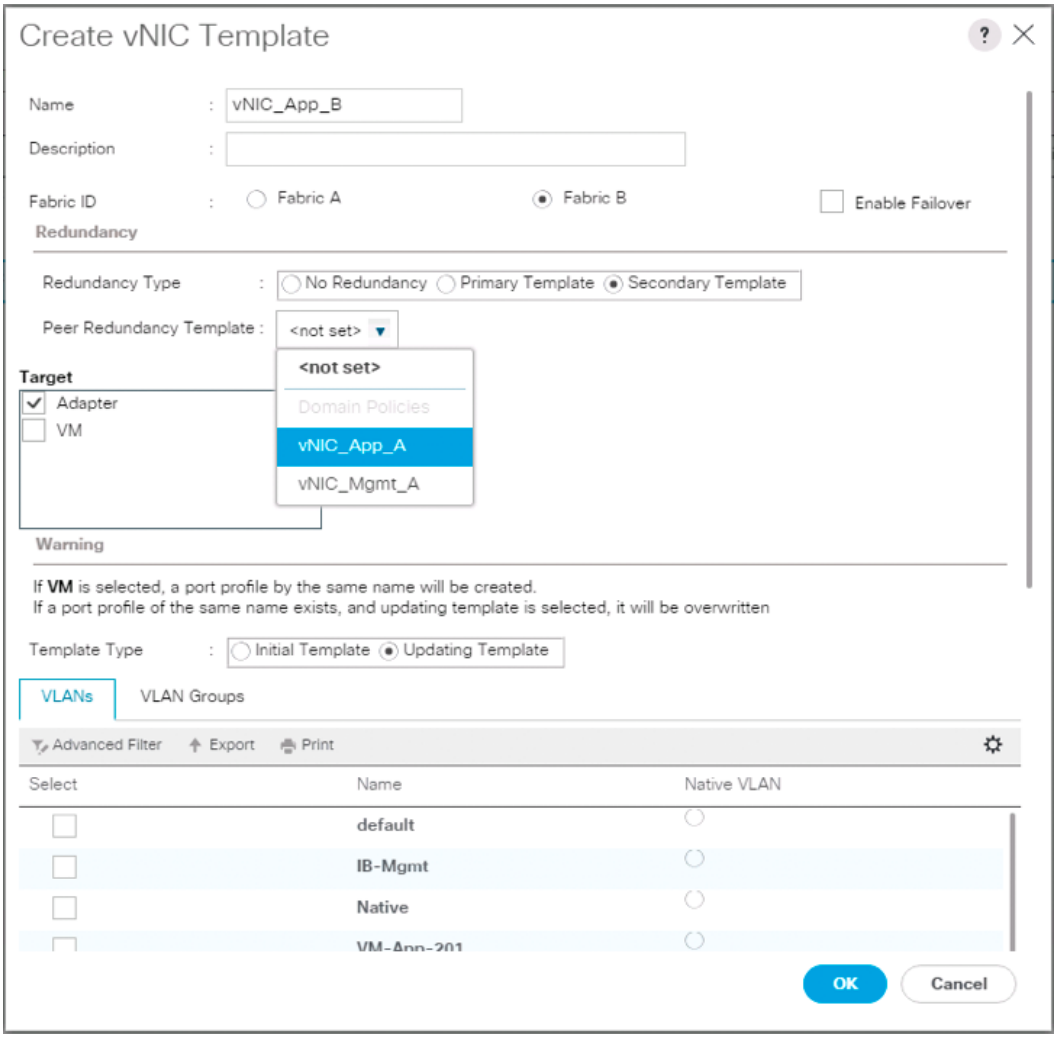

10. In the MAC Pool list, select `MAC_Pool_B`.

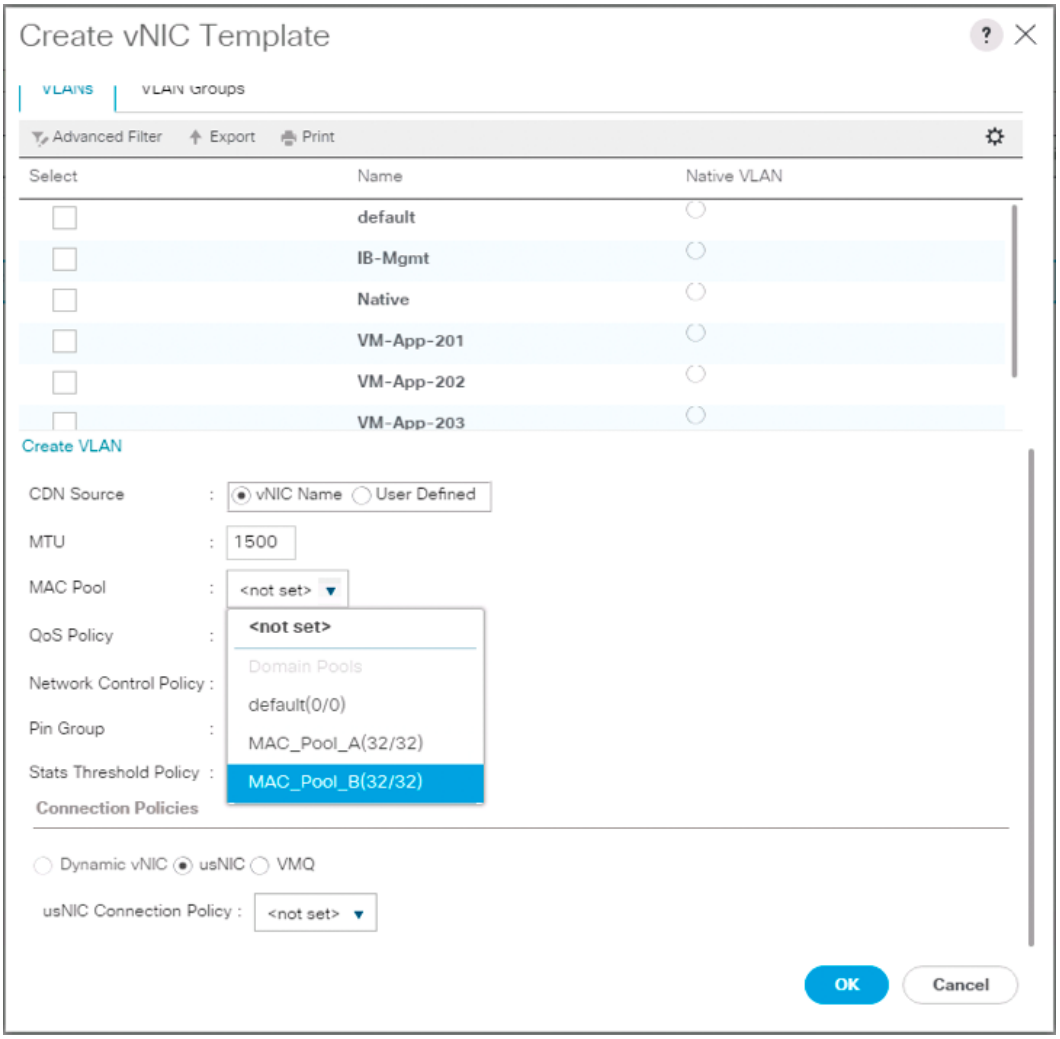

11. Click OK to create the vNIC template.

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.

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.

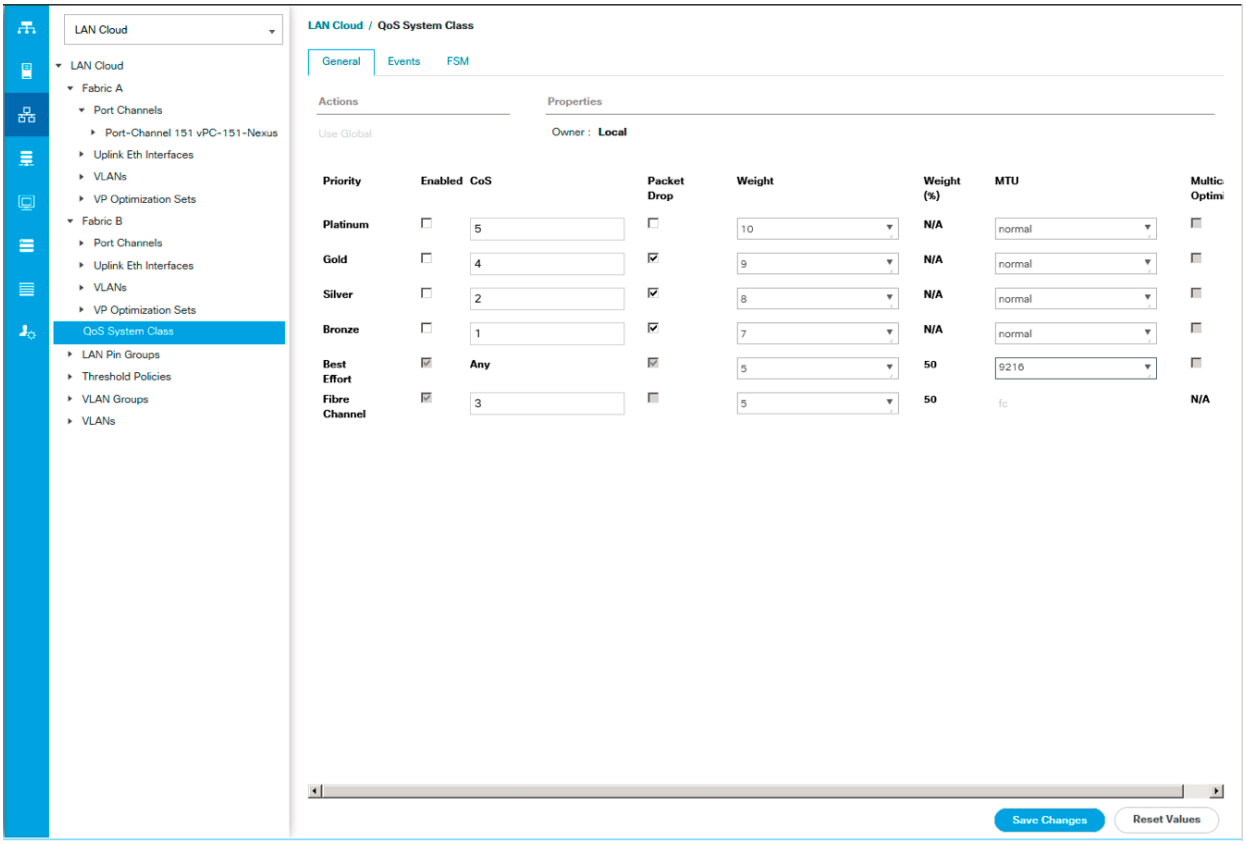

6. Click OK

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.

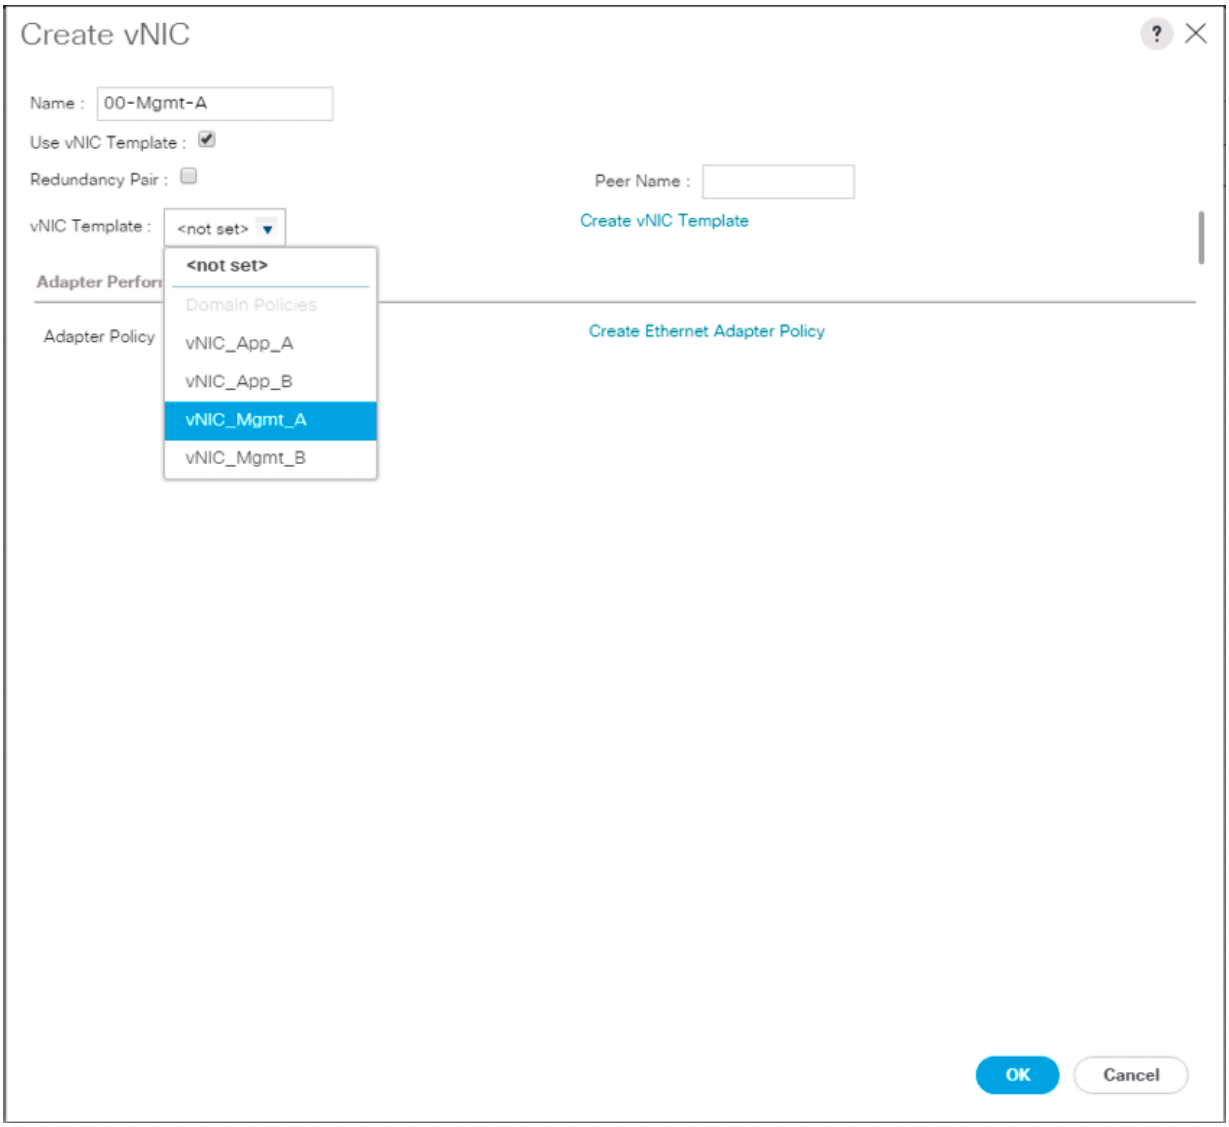

11. Click OK to add this vNIC to the policy.

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.

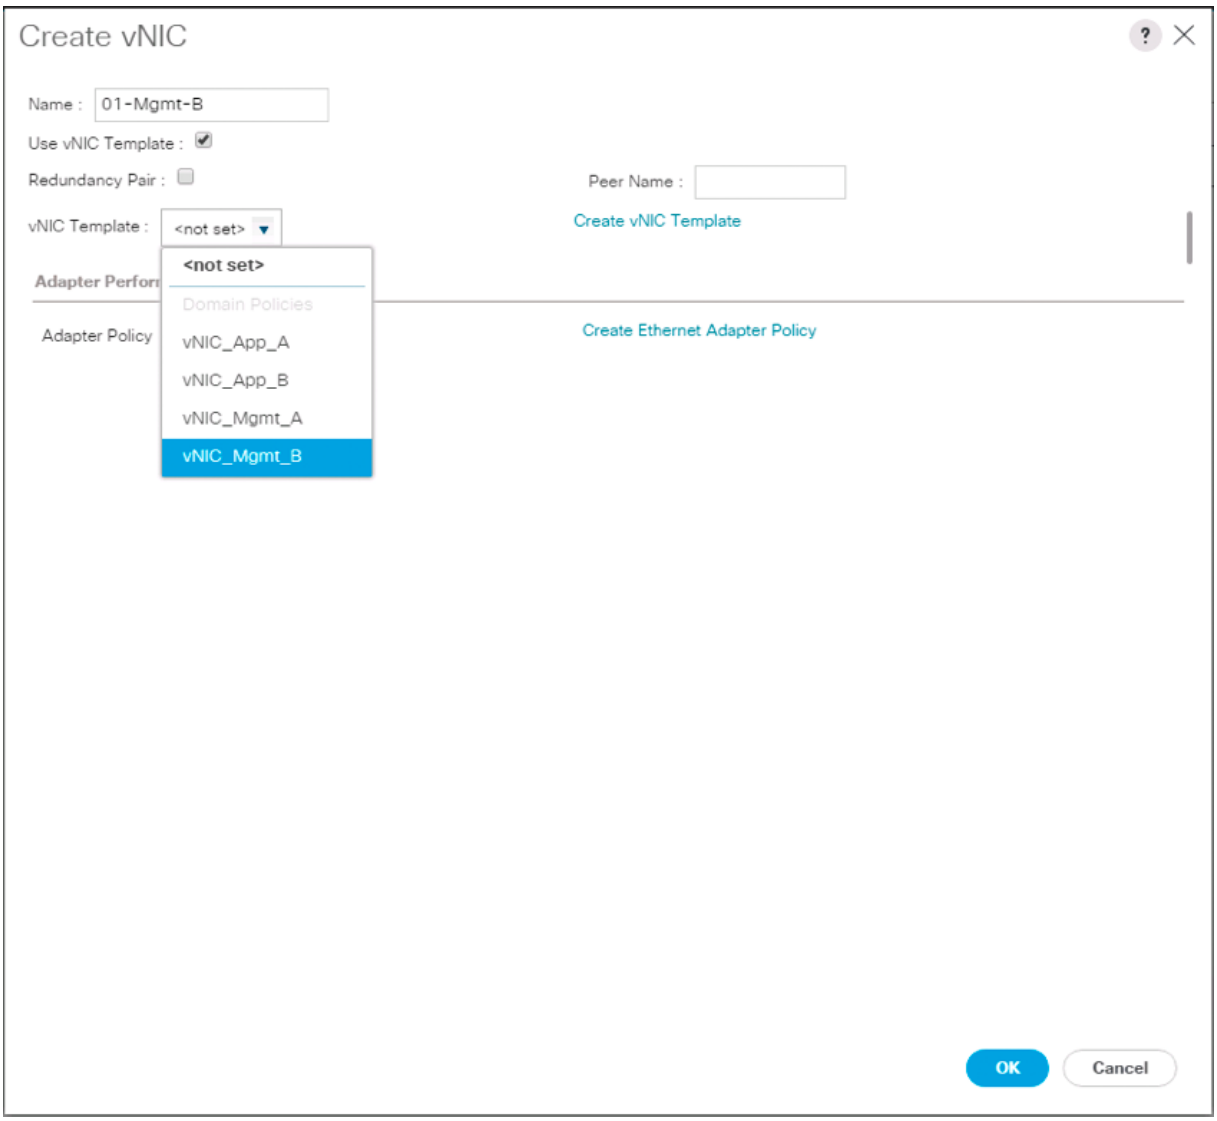

18. Click the upper Add button to add a vNIC.

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.

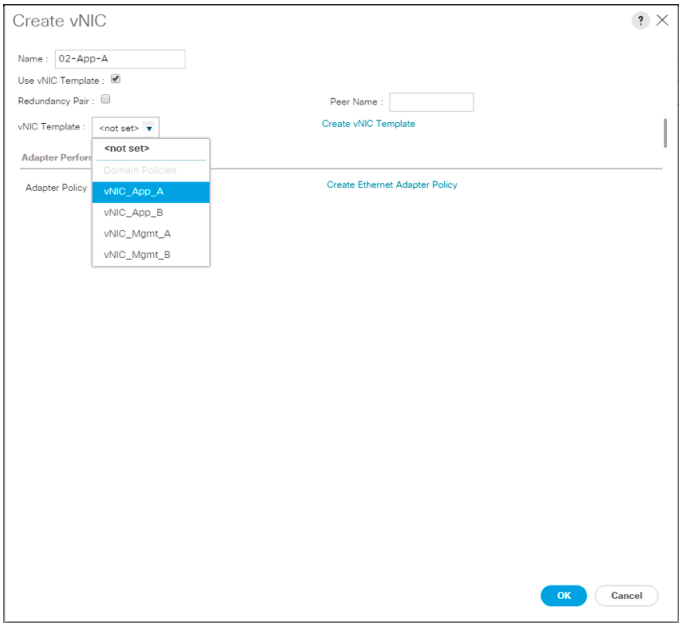

24. Click the upper Add button to add a vNIC to the policy.

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.

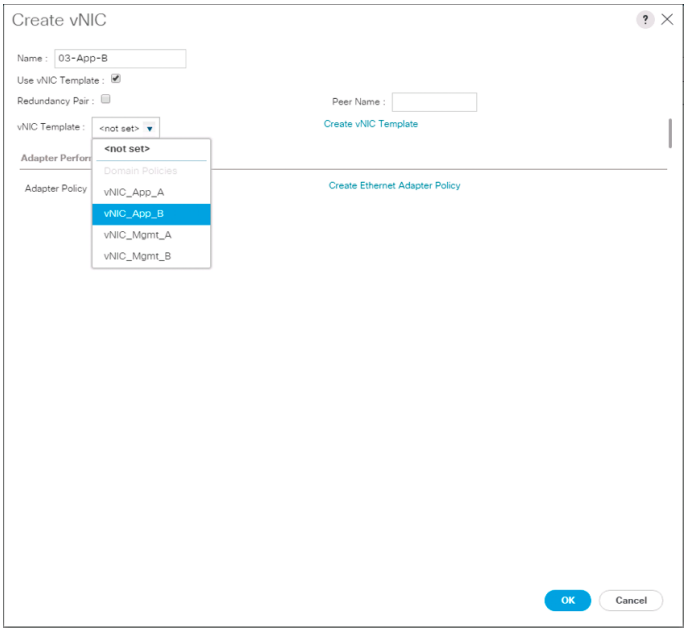

29. Click OK to add this vNIC to the policy.

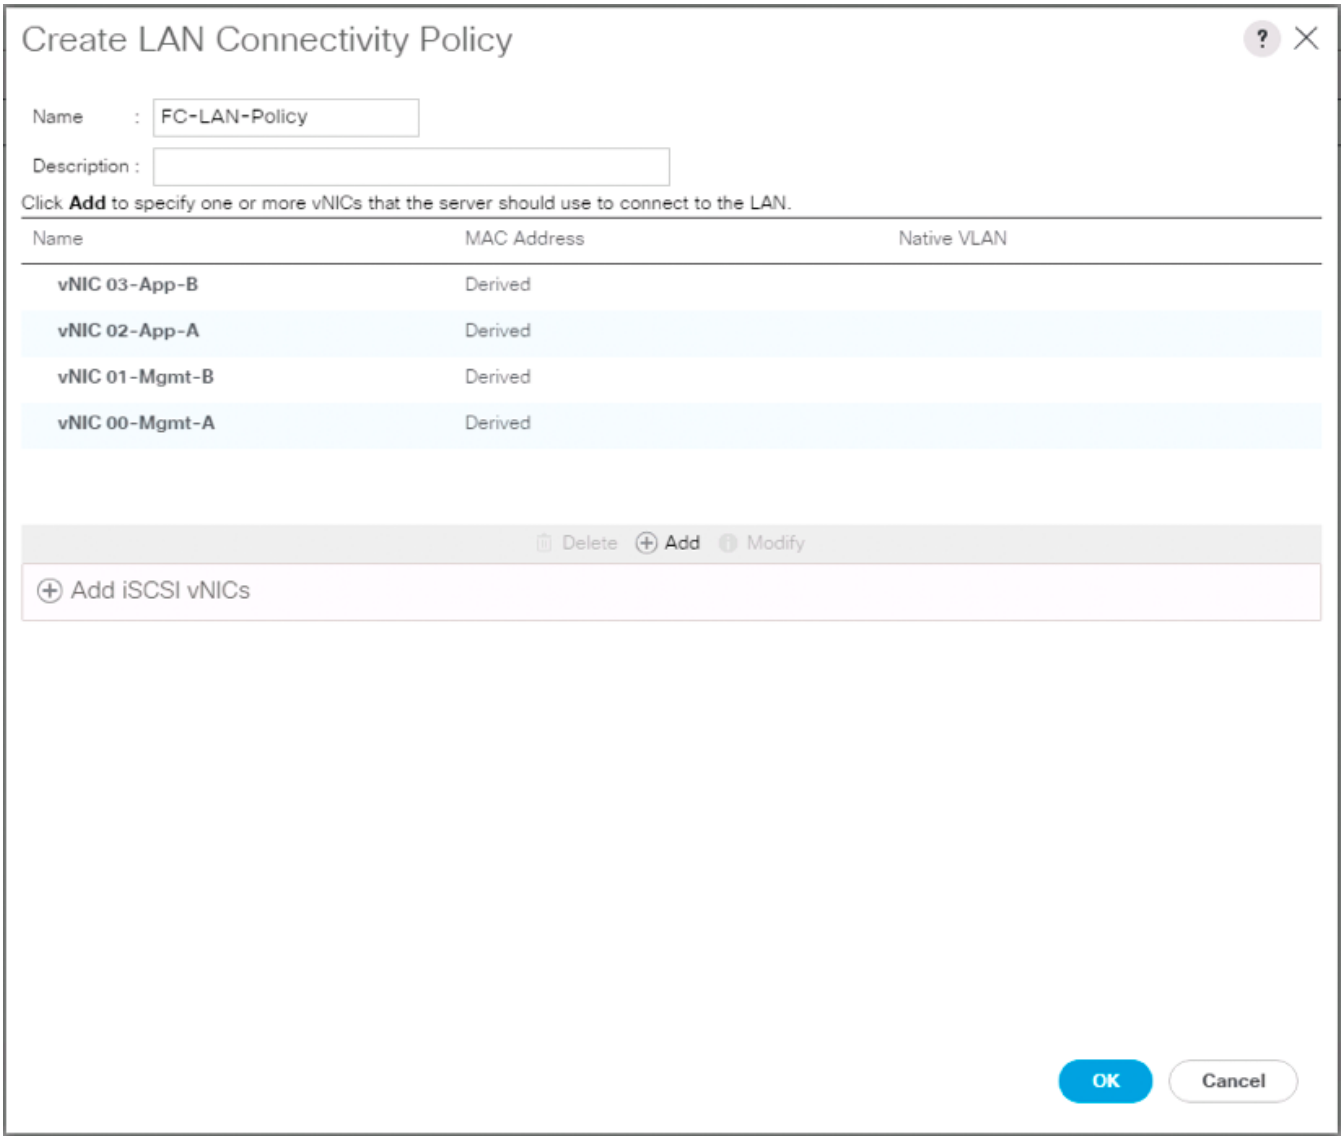

30. Click OK to create the LAN Connectivity Policy.

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.

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. The Cisco UCS 6332-16UP has a similar mechanism controlling the first 16 ports starting from the first port, configuring in increments of the first 6, 12, or all 16 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.

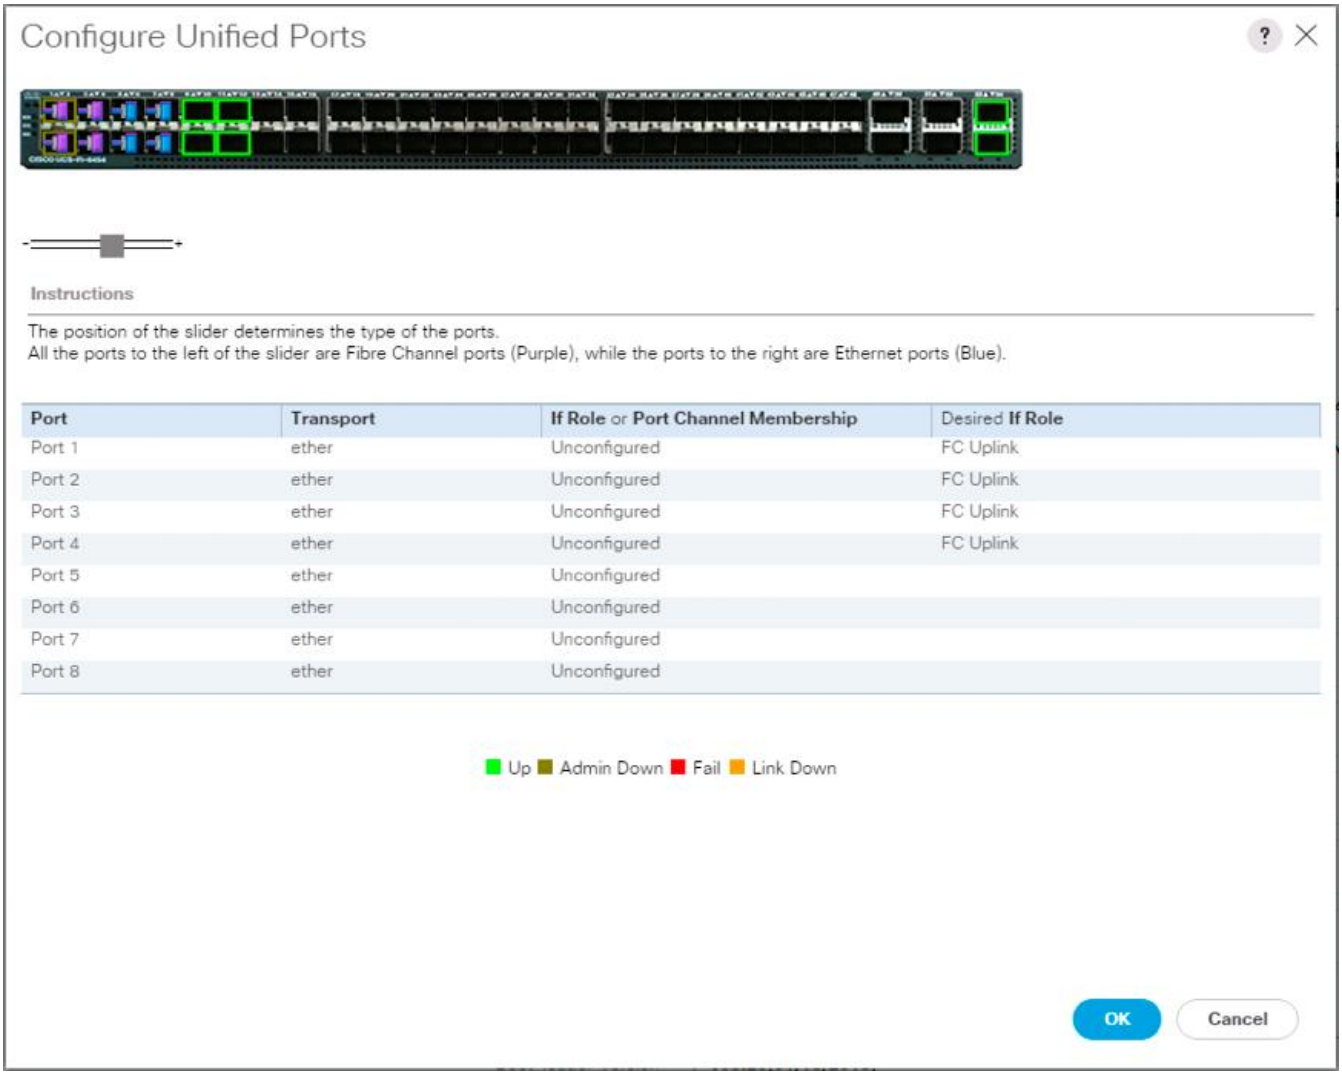

For Cisco UCS 6332-16UP, these fixed ports will be in groups of 6, 12, or 16 ports to be set as FC Uplinks.

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.

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 > SAN Cloud.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A
6. Leave Disabled selected 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.

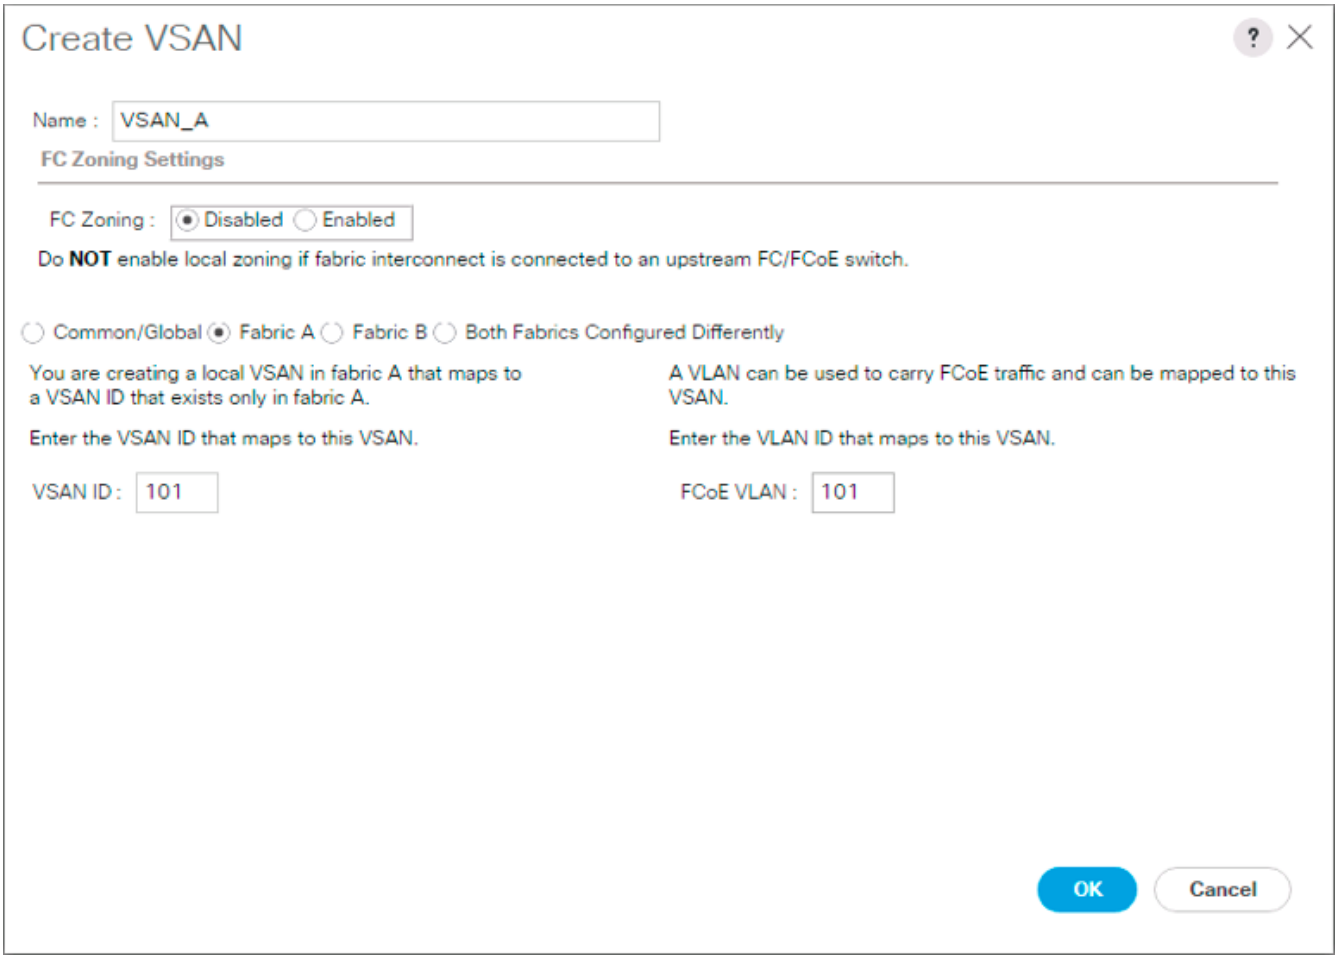

9. Click OK and then click OK again.

10. Under SAN Cloud, right-click VSANs.

11. Select Create VSAN.

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

13. Leave Disabled selected 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.

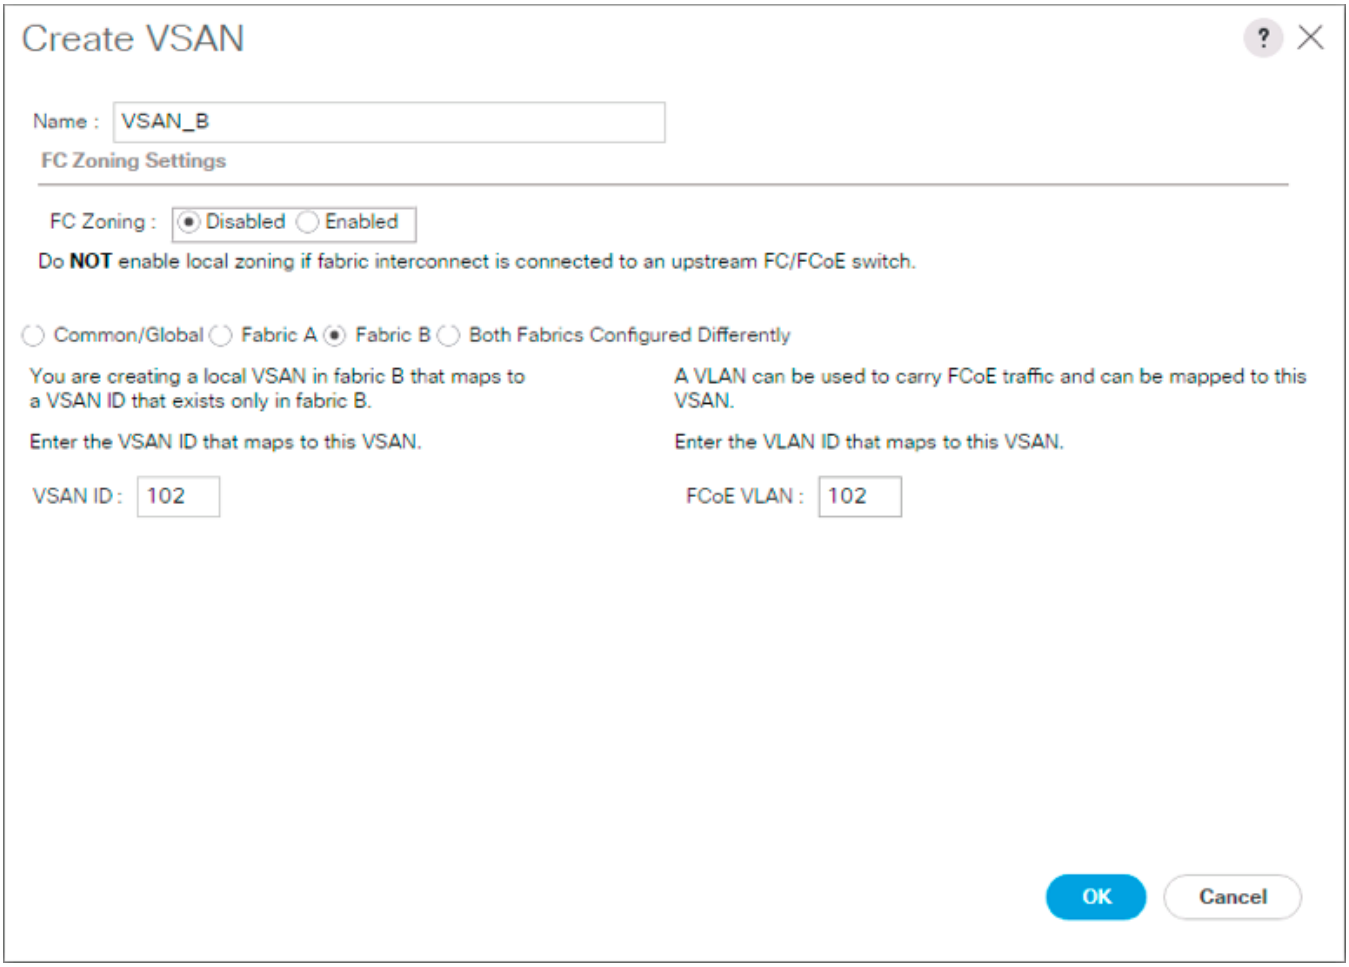

16. Click OK and then click OK again.

Create FC Port Channels

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

Fabric-A

1. In the navigation pane under SAN > SAN Cloud expand the Fabric A tree.
2. Right-click FC Port Channels.
3. Select Create FC Port Channel.
4. Enter 1 for the ID and Po1 for the Port Channel name.

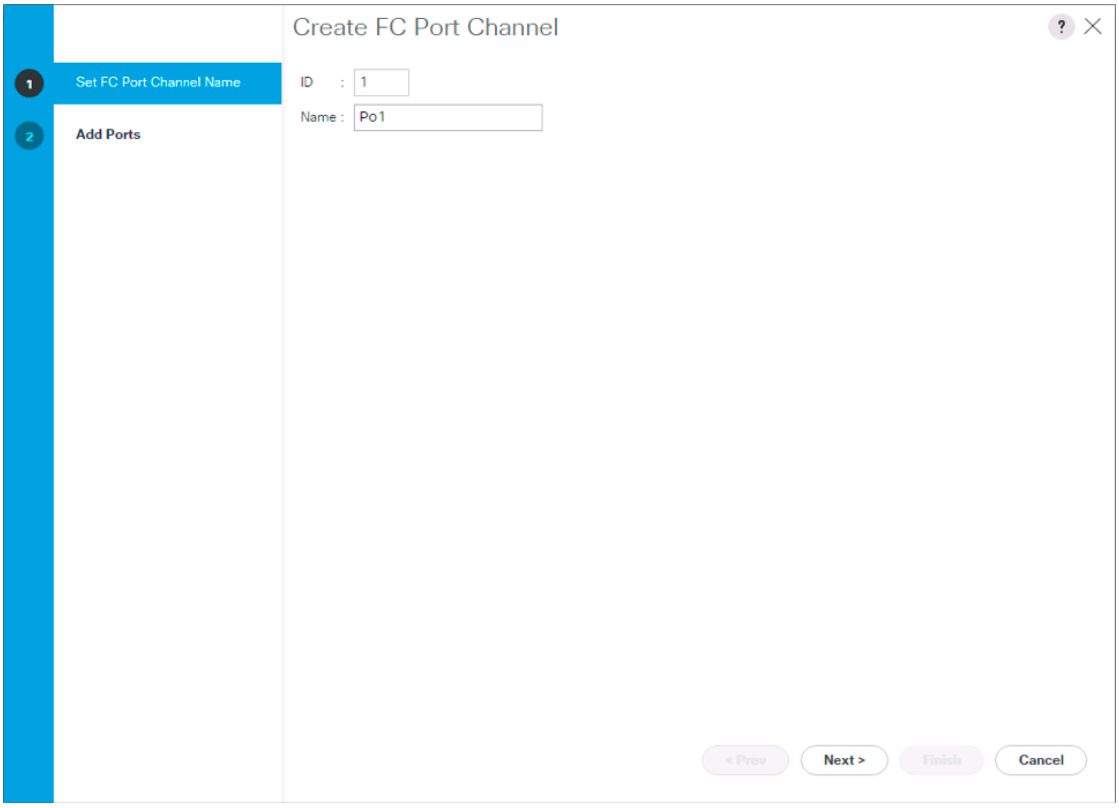

5. Click Next.

6. Set the Port Channel Admin Speed to 32Gbps, or appropriate for the environment, choose connected ports and click >> to add the ports to the port channel.

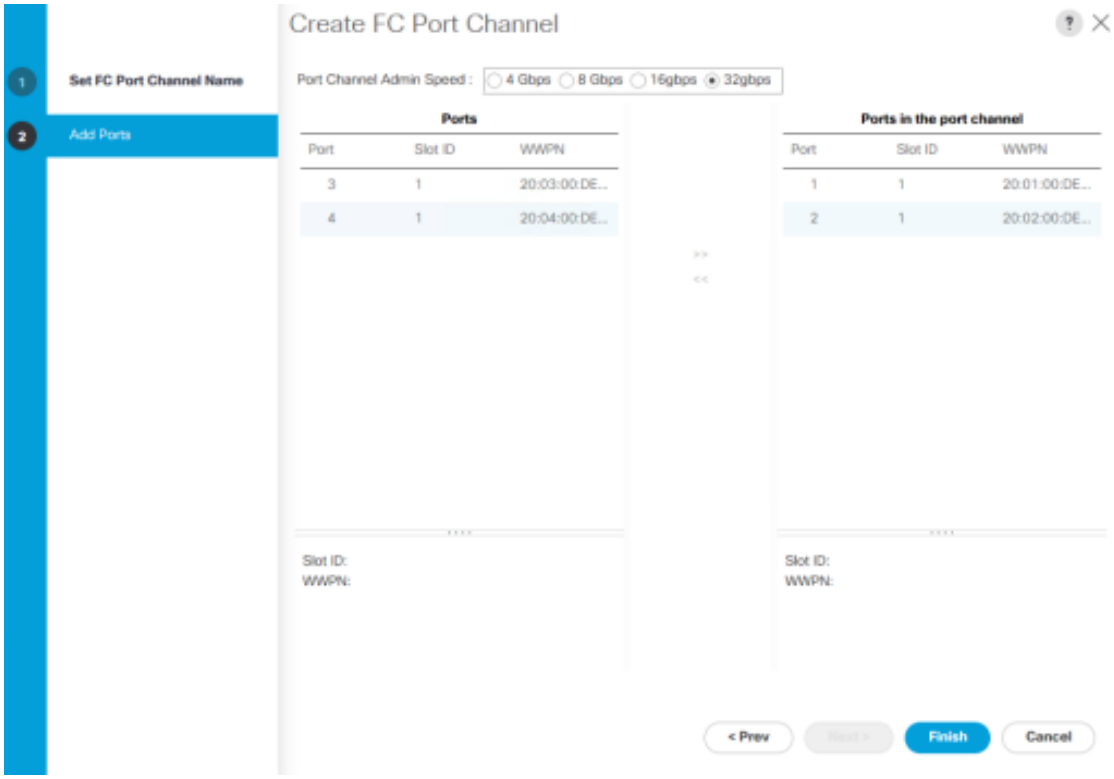

7. Click Finish.

8. Click OK.

9. Select the newly created Port-Channel.

10. Under the VSAN drop-down list for Port-Channel 1, select VSAN_A 101.

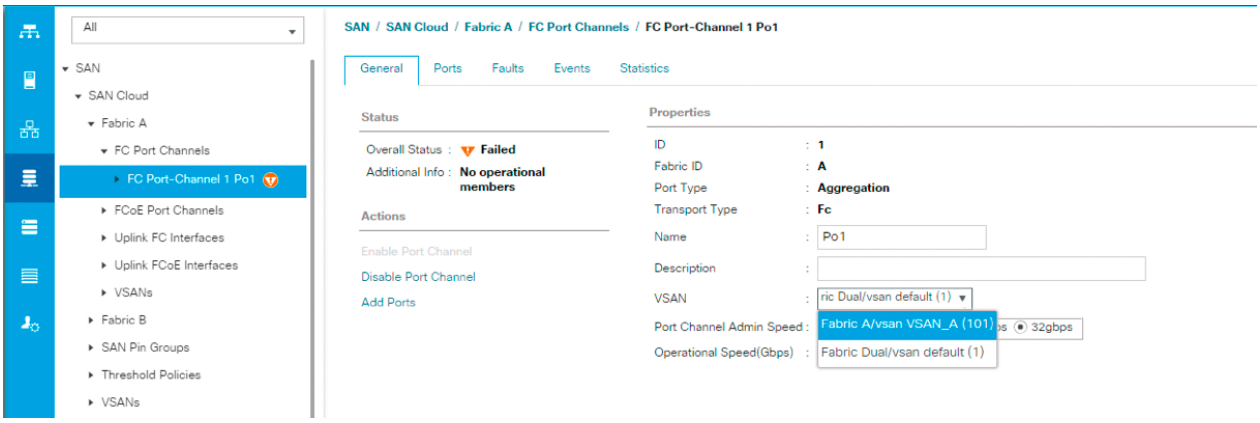

11. Click Save Changes and then click OK.

Fabric-B

1. In the navigation pane, under SAN > SAN Cloud, expand the Fabric B tree.

2. Right-click FC Port Channels.

3. Select Create Port Channel.

4. Enter 2 for the ID and Po2 for the Port Channel name.

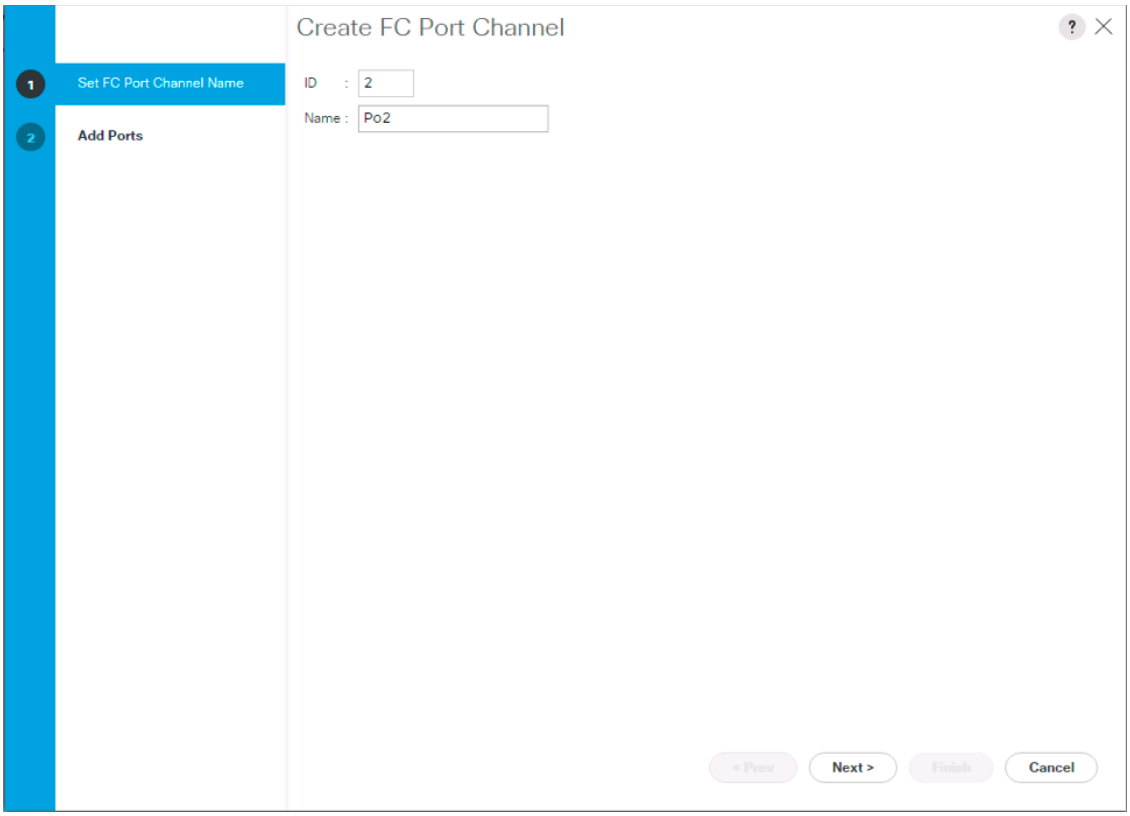

5. Click Next

6. Set the Port Channel Admin Speed to 32Gbps, or appropriate for the environment, choose connected ports and click >> to add the ports to the port channel.

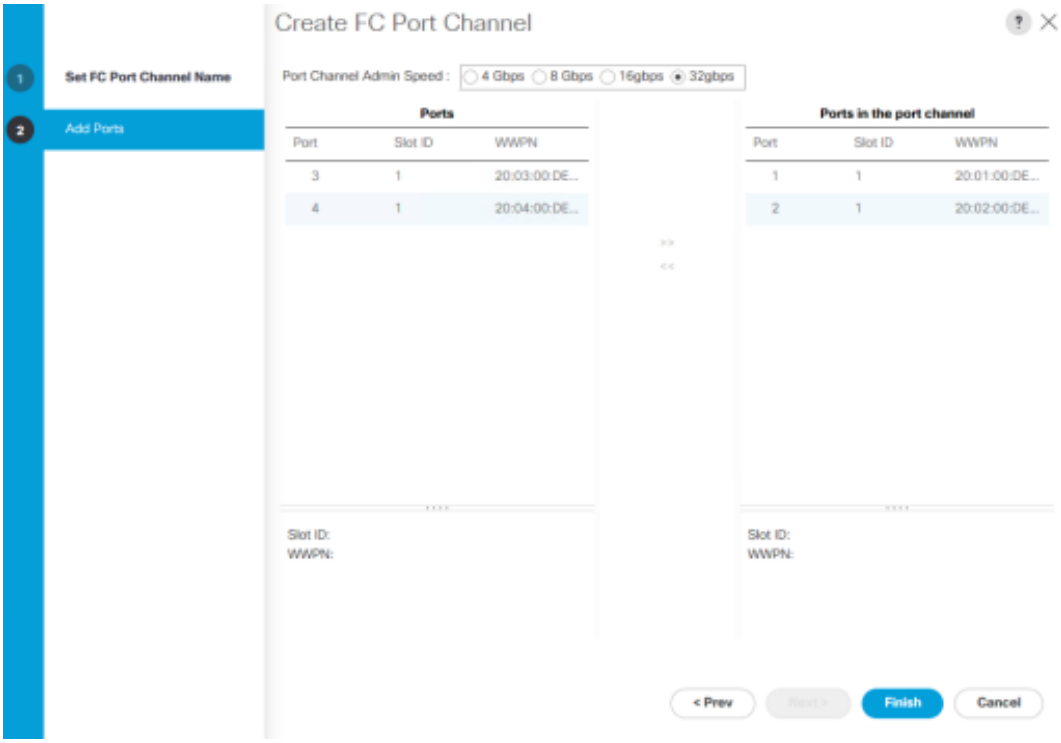

7. Click Finish.

8. Click OK.

9. Select the newly created Port-Channel

10. Under the VSAN drop-down list for Port-Channel 2, select VSAN_B 102.

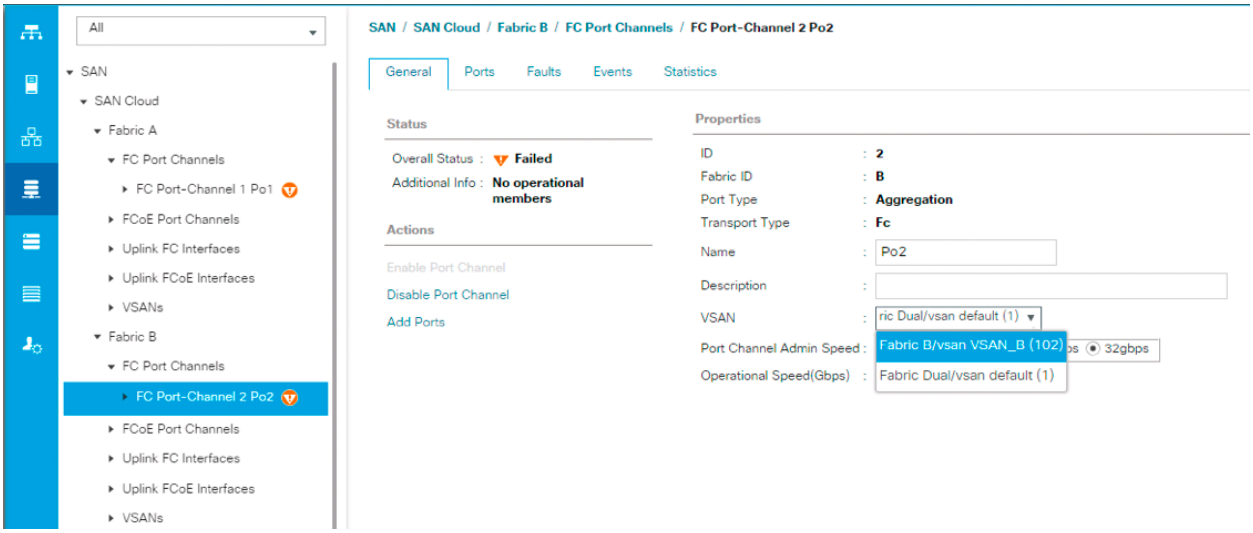

11. Click Save Changes and then click OK.

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.

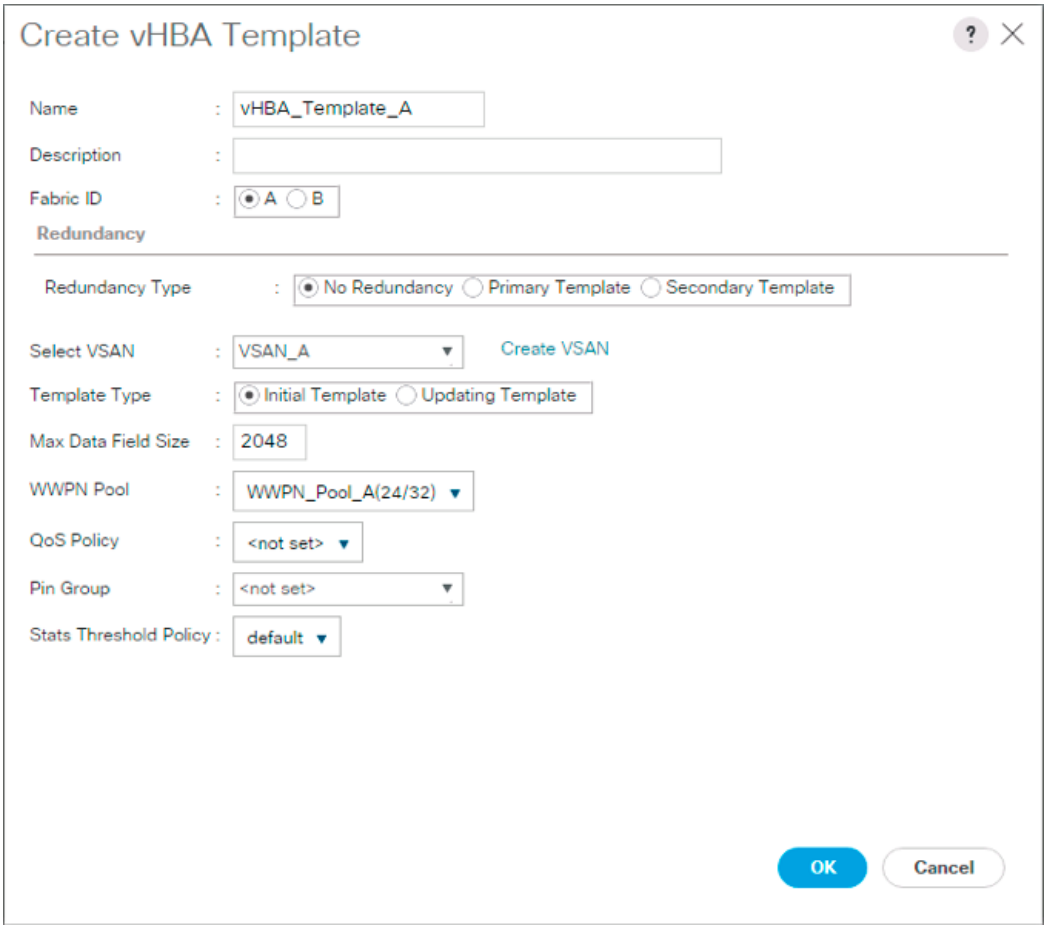

12. Click OK.

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.

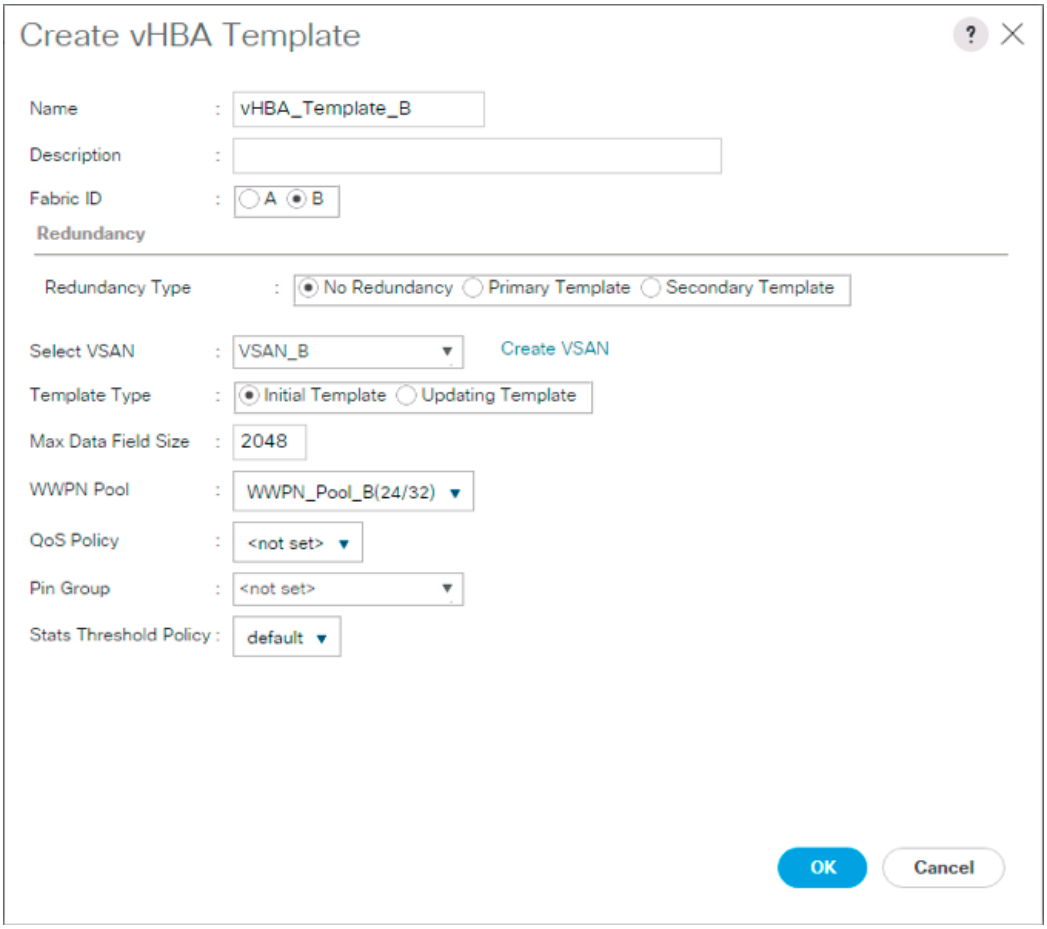

22. Click OK.

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.

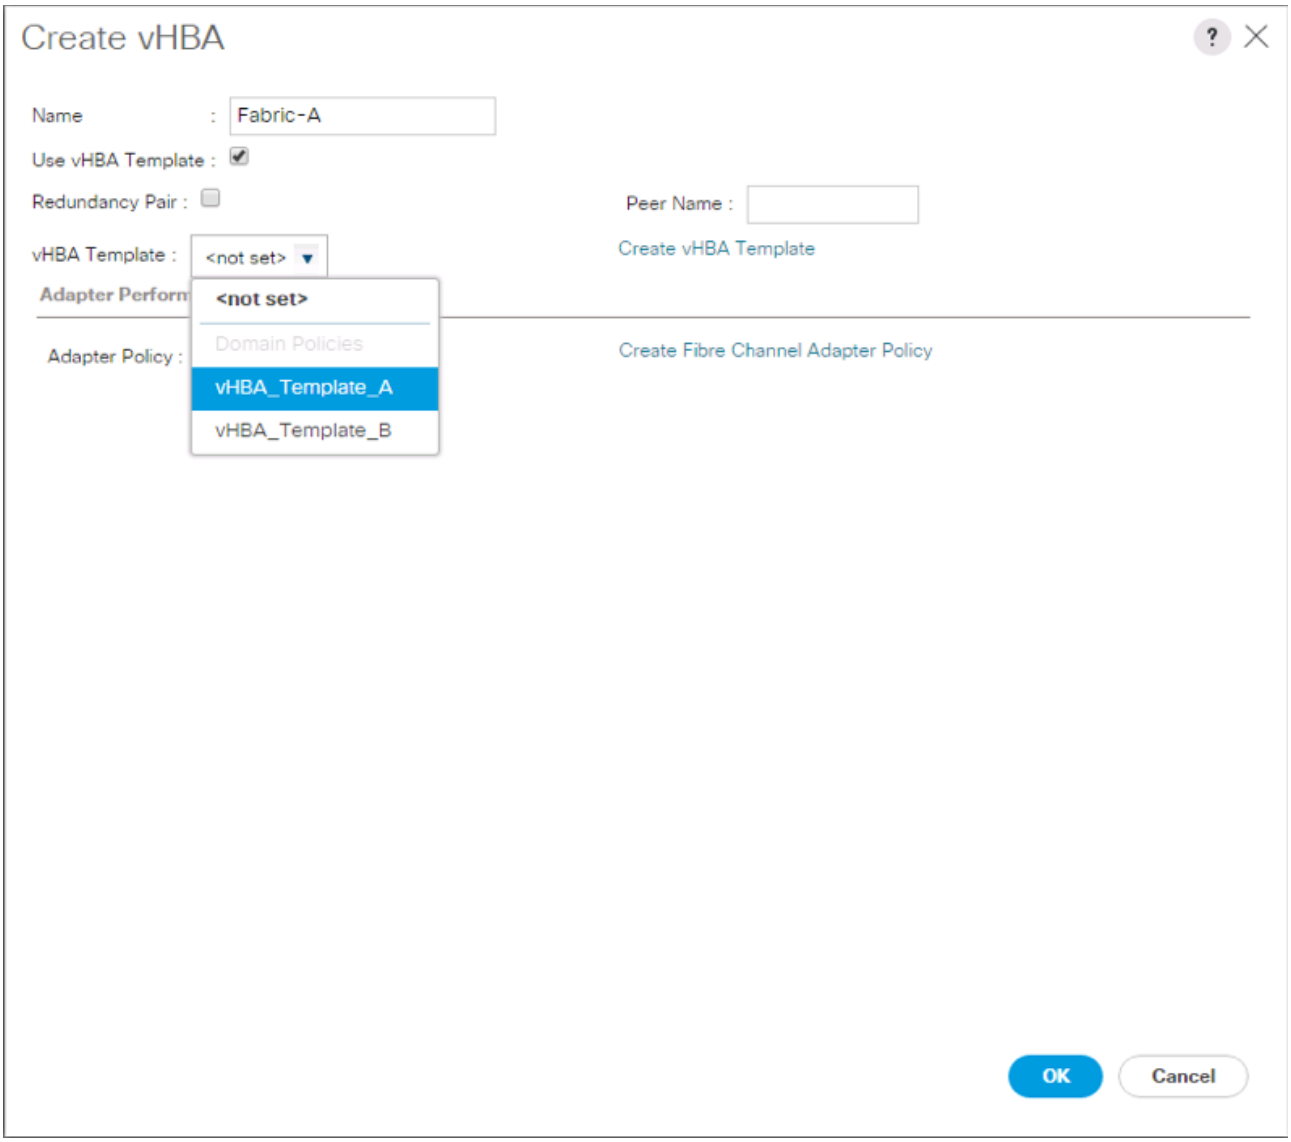

12. In the Adapter Policy list, select VMWare.

13. Click OK.

14. Click the Add button at the bottom 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.

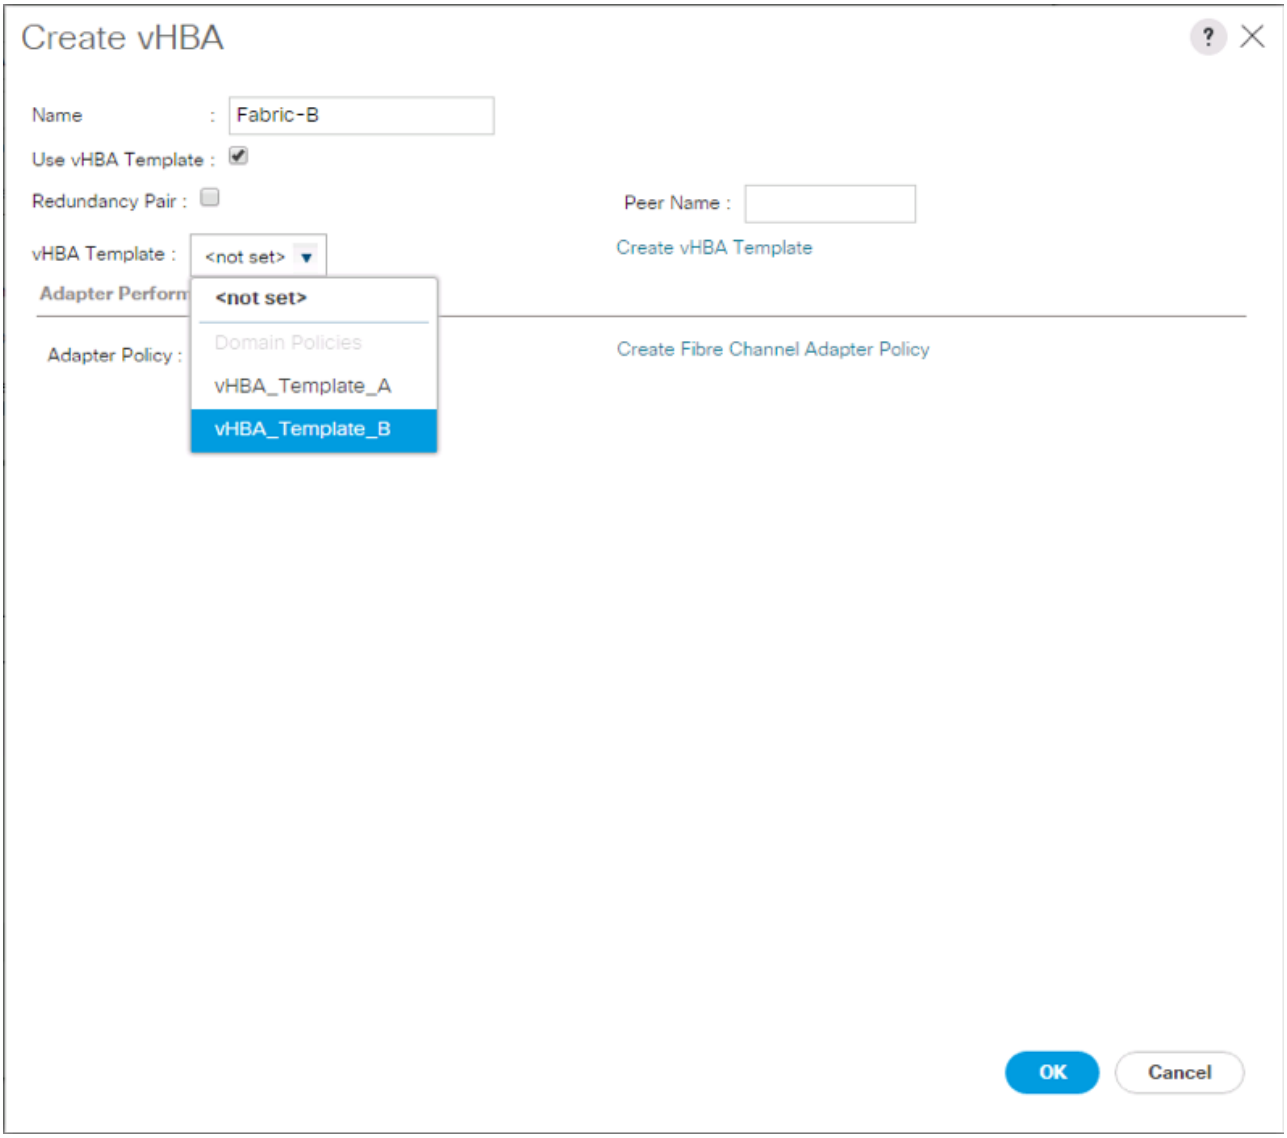

19. In the Adapter Policy list, select VMWare.

20. Click OK.

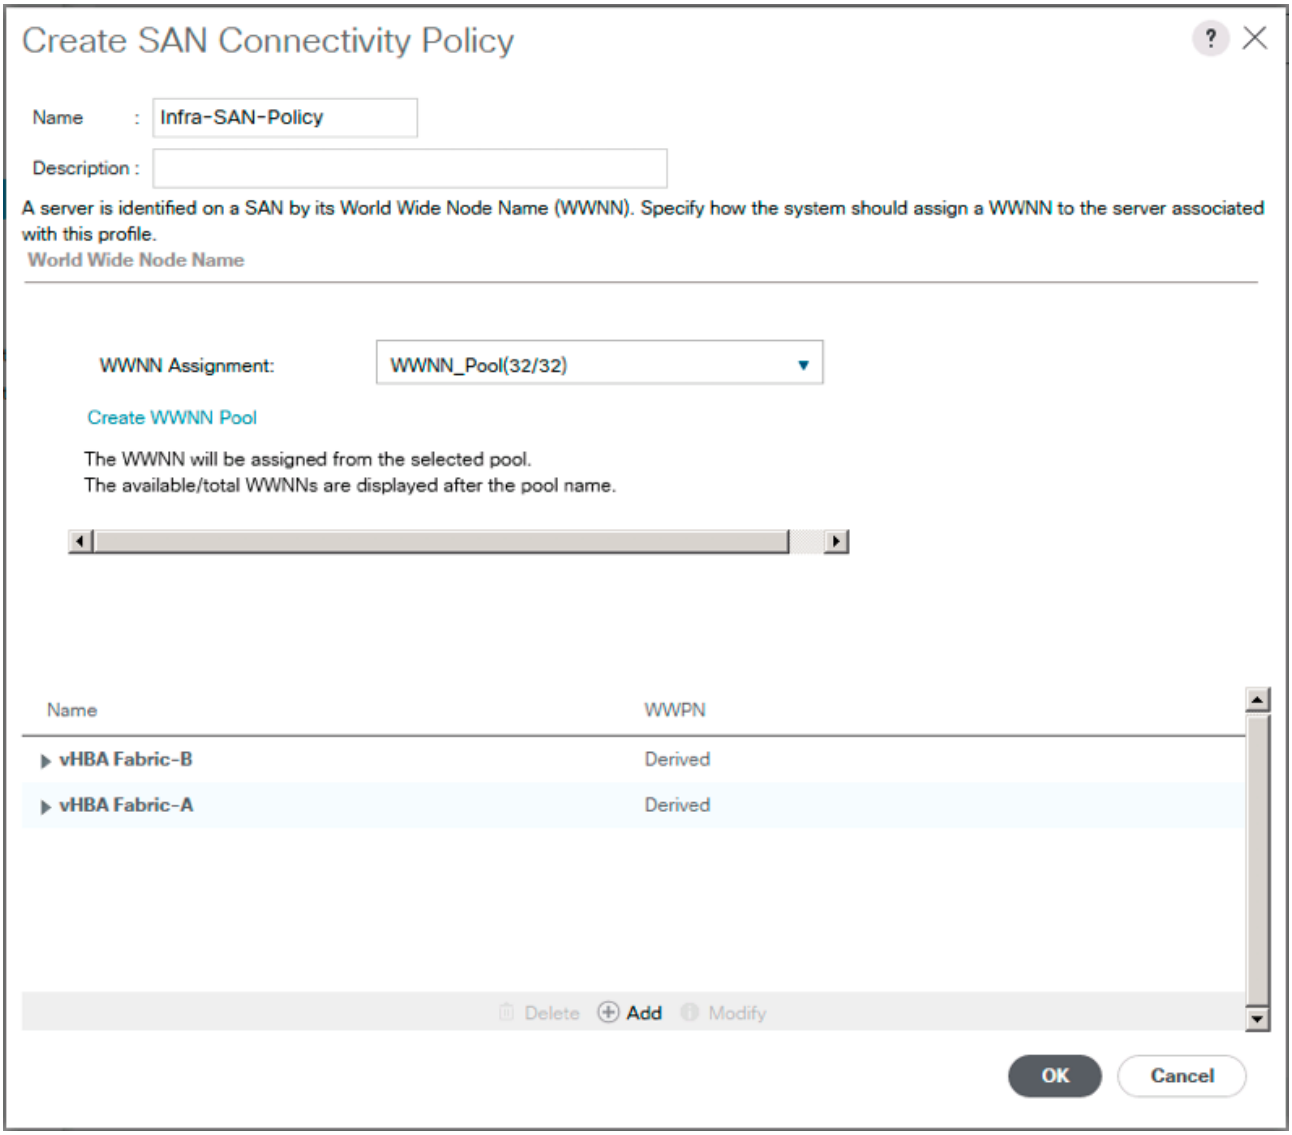

21. Click OK to create the SAN Connectivity Policy.

22. Click OK to confirm creation.

Create Boot Policy

The VSP G370 and/or G1500 target WWPN will need to be collected at this point 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 MDS will be fairly quick provided there is clear identification of the port cabling from the VSP ports to the MDS ports.

Table 20 VSP G370 to MDS Port Information Carried Forward

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G370	CL 1-A	32Gb FC	Cisco MDS 9706 A	FC 1/11

Local Device	Local Port	Connection	Remote Device	Remote Port
	CL 2-B	32Gb FC	Cisco MDS 9706 A	FC 1/12
	CL 3-B	32Gb FC	Cisco MDS 9706 B	FC 1/11
	CL 4-A	32Gb FC	Cisco MDS 9706 B	FC 1/12

Along with the Table 17 information for the VSP G1500:

Table 21 VSP G1500 to MDS Port Information Carried Forward

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP G1500	CL 1-A	16Gb FC	Cisco MDS 9706 A	FC 1/7
	CL 2-A	16Gb FC	Cisco MDS 9706 A	FC 1/8
	CL 1-J	16Gb FC	Cisco MDS 9706 A	FC 1/9
	CL 2-J	16Gb FC	Cisco MDS 9706 A	FC 1/10
	CL 3-L	16Gb FC	Cisco MDS 9706 B	FC 1/7
	CL 4-C	16Gb FC	Cisco MDS 9706 B	FC 1/8
	CL 3-C	16Gb FC	Cisco MDS 9706 B	FC 1/9
	CL 4-L	16Gb FC	Cisco MDS 9706 B	FC 1/10

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

Running the sh flogi database command on MDS A:

```
AA19-9706-1# sh flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/7            101   0xbc0080  50:06:0e:80:07:56:24:00 50:06:0e:80:07:56:24:00
fc1/8            101   0xbc0060  50:06:0e:80:07:56:24:10 50:06:0e:80:07:56:24:10
fc1/9            101   0xbc0000  50:06:0e:80:07:56:24:08 50:06:0e:80:07:56:24:08
fc1/10           101   0xbc0100  50:06:0e:80:07:56:24:18 50:06:0e:80:07:56:24:18
fc1/11           101   0xbc01a0  50:06:0e:80:12:c9:9a:00 50:06:0e:80:12:c9:9a:00
fc1/12           101   0xbc0180  50:06:0e:80:12:c9:9a:11 50:06:0e:80:12:c9:9a:11
```

Running the sh flogi database command on MDS B:

```
aa19-9706-2# sh flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/7            102   0x2800e0  50:06:0e:80:07:56:24:2a 50:06:0e:80:07:56:24:2a
fc1/8            102   0x2800a0  50:06:0e:80:07:56:24:32 50:06:0e:80:07:56:24:32
fc1/9            102   0x280060  50:06:0e:80:07:56:24:22 50:06:0e:80:07:56:24:22
fc1/10           102   0x280020  50:06:0e:80:07:56:24:3a 50:06:0e:80:07:56:24:3a
fc1/11           102   0x2801e0  50:06:0e:80:12:c9:9a:21 50:06:0e:80:12:c9:9a:21
fc1/12           102   0x280180  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 MDS Interface (Remote Port) values in

the previous table for this fabric have been cross referenced, and the WWPN(Port Name) for these interfaces are recorded.

Table 22 Fabric A Boot Targets for the VSP G370

	MDS Interface	Example Local Port	Target Role	WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
VSP G370 Controller 1	1/11	CL 1-A	Primary/ VMFS	50:06:0e:80:12:c9:9a:00	
VSP G370 Controller 2	1/12	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 23 Fabric B Boot Targets for the VSP G370

	MDS Interface	Example Local Port	Target Role	WWN/WWPN Example Environment	WWN/WWPN Customer Environment
VSP G370 Controller 1	1/11	CL 3-B	Primary Boot/ VMFS	50:06:0e:80:12:c9:9a:21	
VSP G370 Controller 2	1/12	CL 4-A	Secondary Boot/ VMFS	50:06:0e:80:12:c9:9a:30	

This equivalent information gathering for the VSP G1500 will be:

Table 24 Fabric A Boot Targets for the VSP G1500

	MDS Interface	Example Local Port	Target Role	WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
VSP G1500 Controller 0 Cluster 1	1/7	CL 1-A	Primary Boot/ VMFS	50:06:0e:80:07:56:24:00	
VSP G1500 Controller 0 Cluster 2	1/8	CL 2-A	VMFS	50:06:0e:80:07:56:24:10	
VSP G1500 Controller 1 Cluster 1	1/9	CL 1-J	VMFS	50:06:0e:80:07:56:24:08	
VSP G1500 Controller 1 Cluster 2	1/10	CL 2-J	Secondary Boot / VMFS	50:06:0e:80:07:56:24:18	

Repeat these steps for the VSP G1500 Fabric B Primary and Secondary Boot Targets:

Table 25 Fabric B Boot Targets for the VSP G1500

	MDS Interface	Example Local Port	Target Role	WWN/WWPN Example Environment (Port Name)	WWN/WWPN Customer Environment
VSP G1500 Controller 1 Cluster 1	1/7	CL 3-L	Primary Boot/ VMFS	50:06:0e:80:07:56:24:2a	
VSP G1500 Controller 0 Cluster 2	1/8	CL 4-C	Secondary Boot/ VMFS	50:06:0e:80:07:56:24:32	
VSP G1500 Controller 0 Cluster 1	1/9	CL 3-C	VMFS	50:06:0e:80:07:56:24:22	
VSP G1500 Controller 1 Cluster 2	1/10	CL 4-L	VMFS	50:06:0e:80:07:56:24:3a	

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

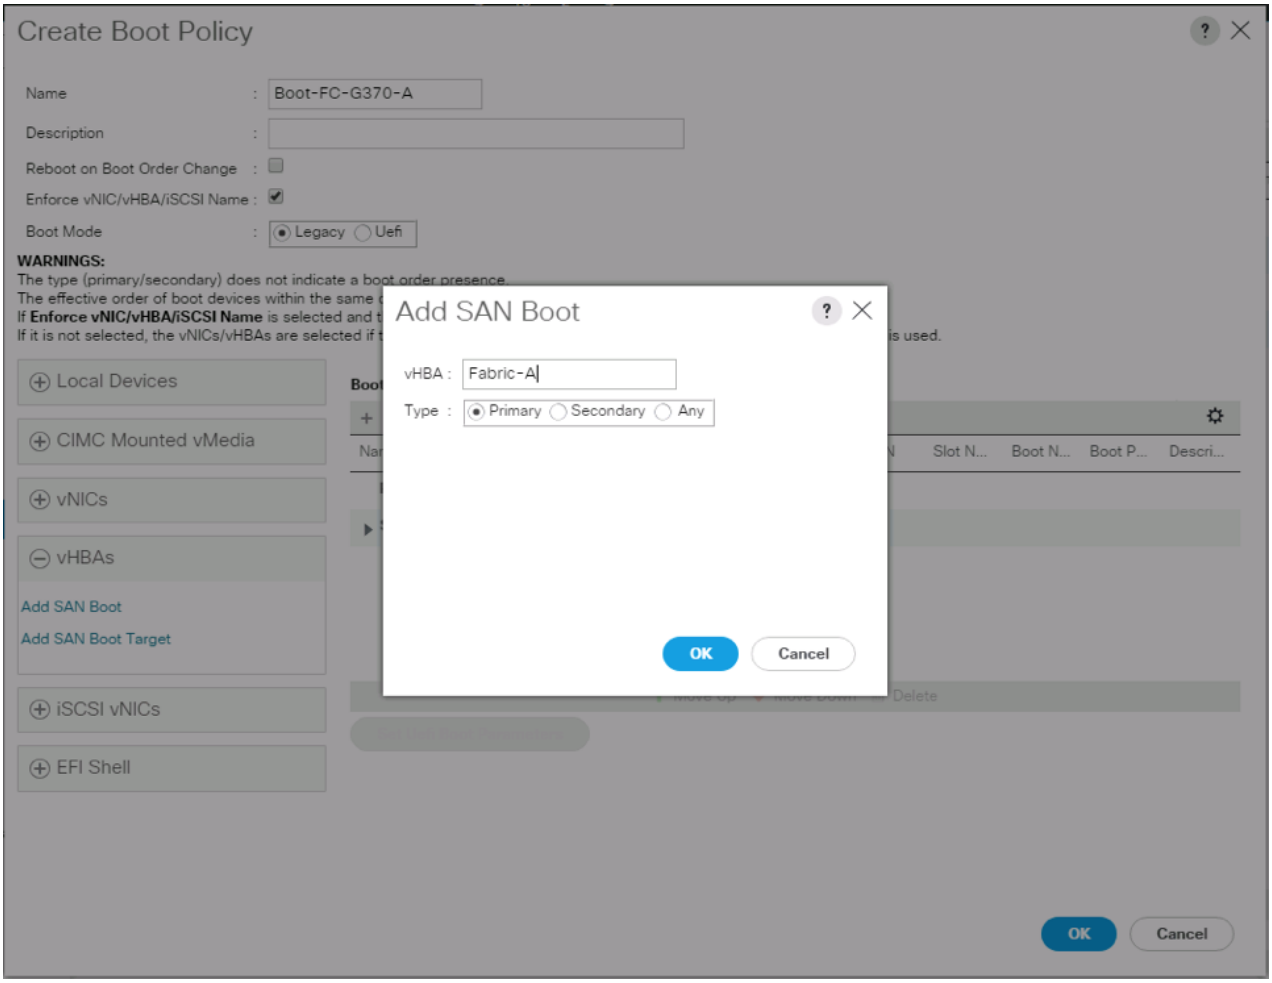

11. Click OK to add the SAN boot initiator.

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

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

14. Enter the WWPN for `Controller1 (CL 1A)` recorded in Error! Reference source not found..

15. Select Primary for the SAN boot target type.

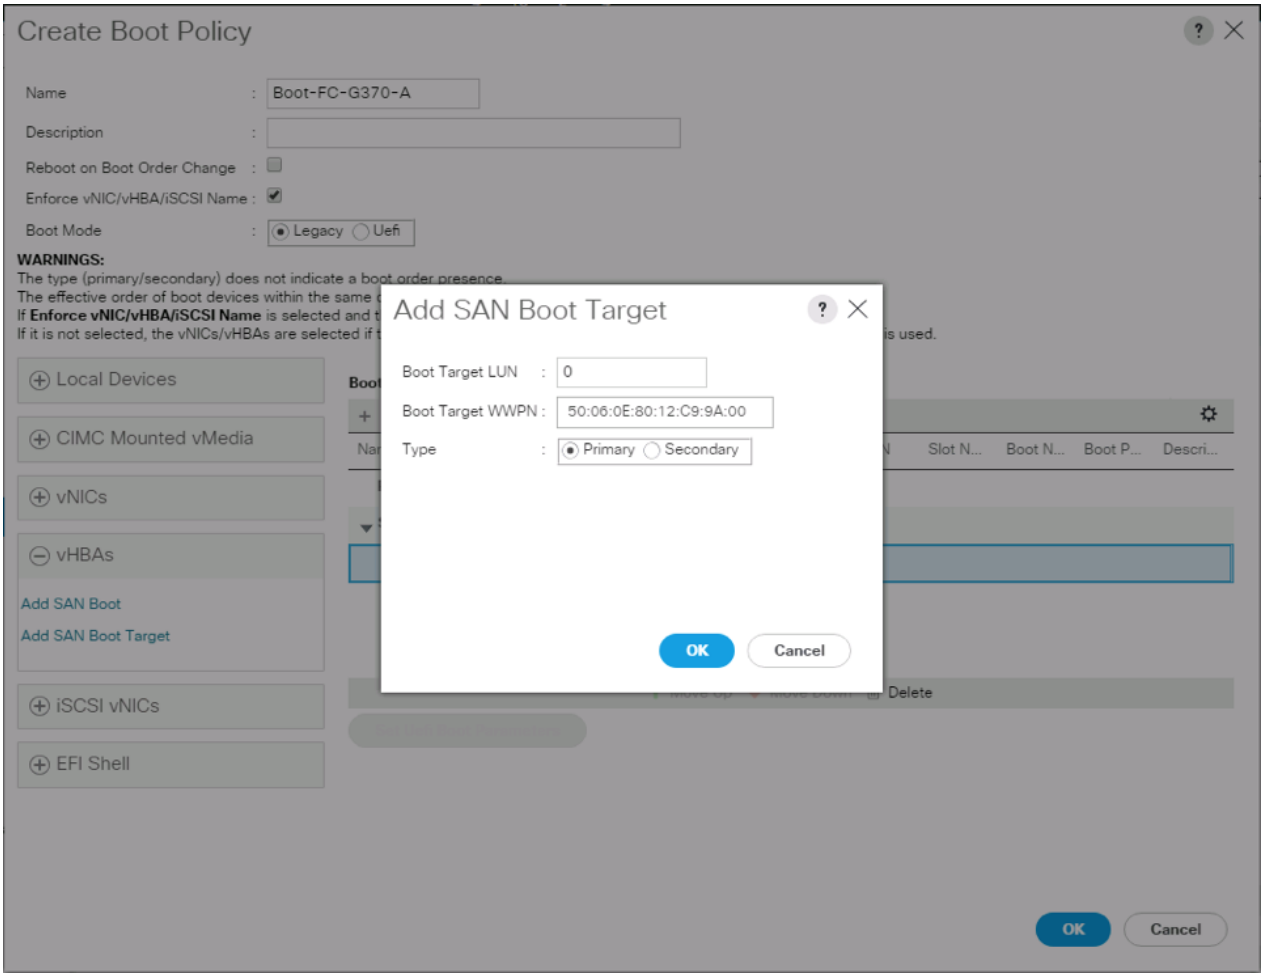

16. Click OK to add the SAN boot target.

17. From the vHBA drop-down menu, 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 25

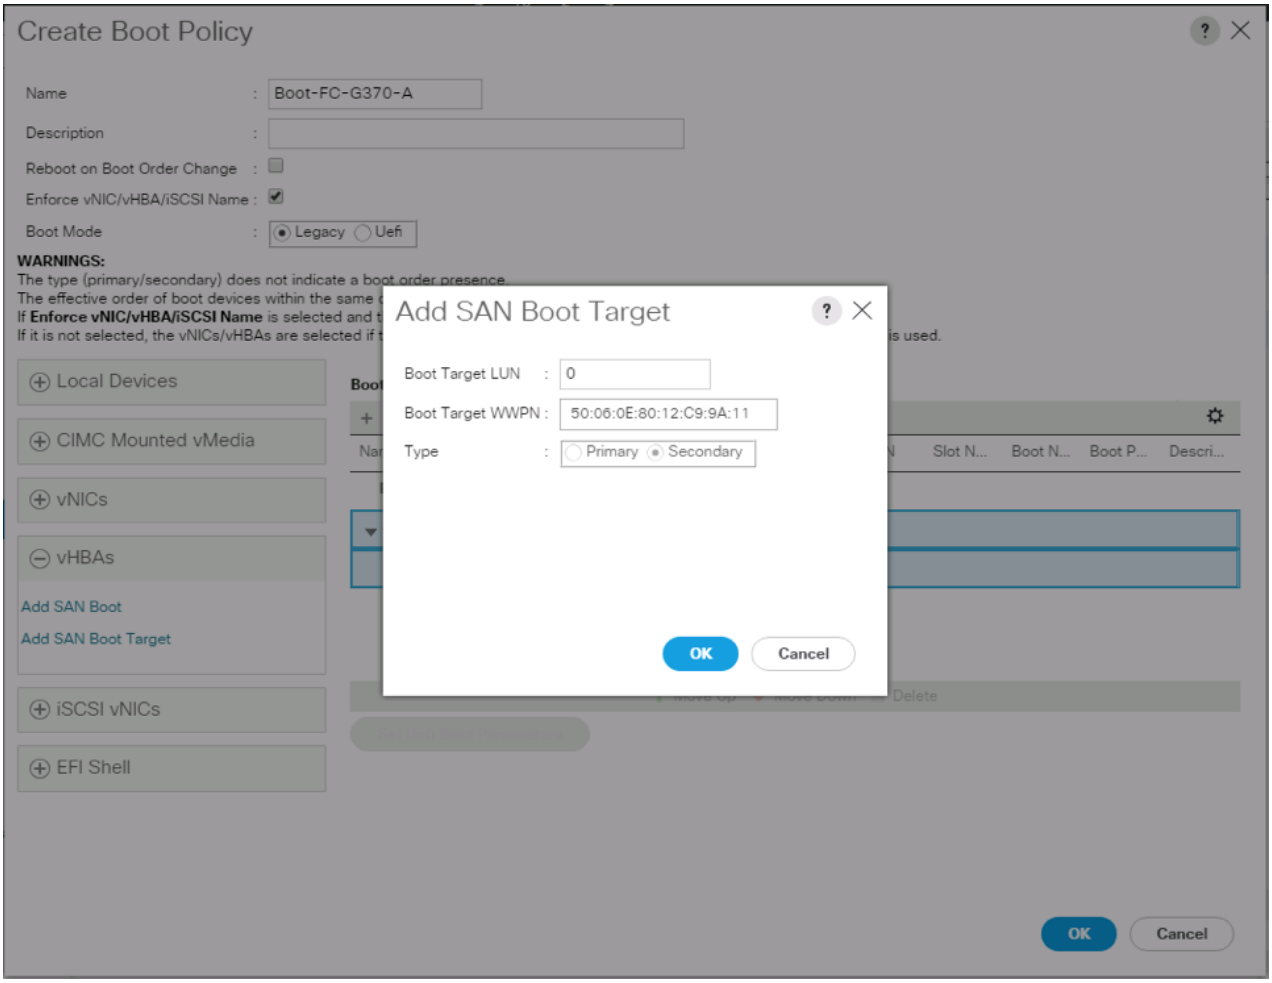

20. Click OK to add the SAN boot target.

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

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

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

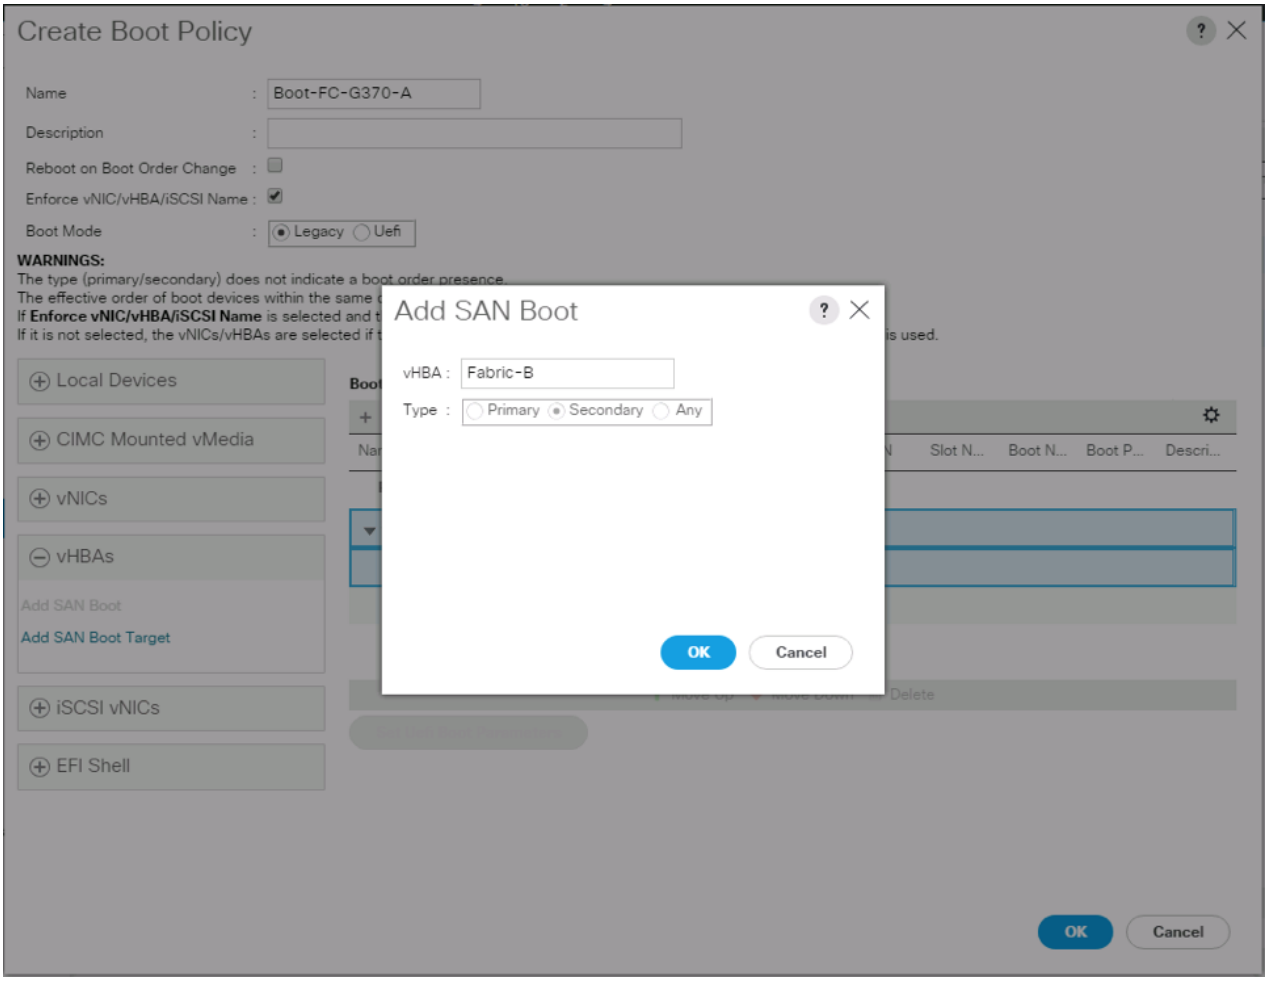

23. Click OK to add the SAN boot initiator.

24. From the vHBA drop-down menu, 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 23

27. Select Primary for the SAN boot target type.

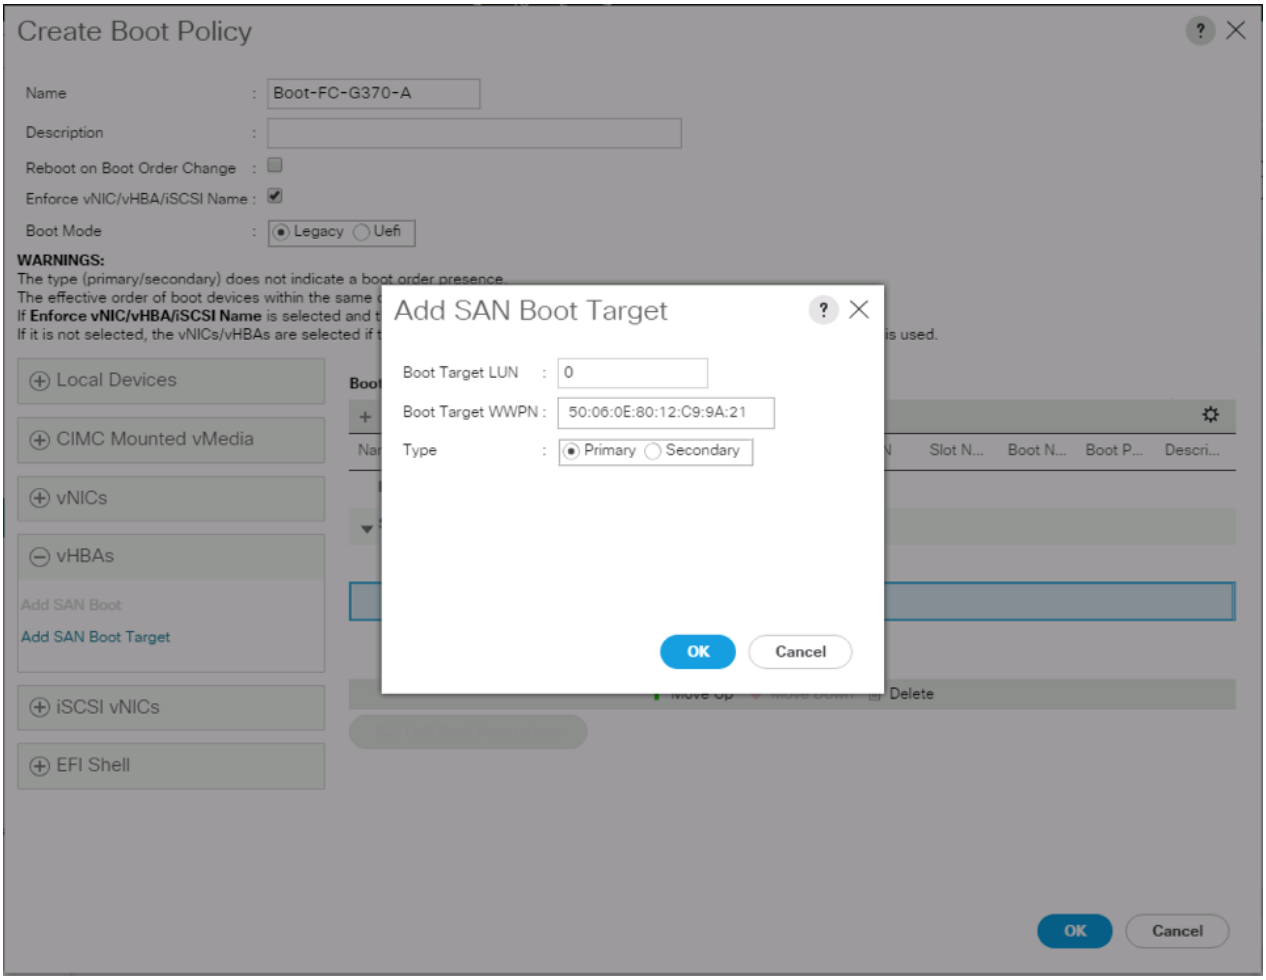

28. Click OK to add the SAN boot target.

29. From the vHBA drop-down menu, 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 24

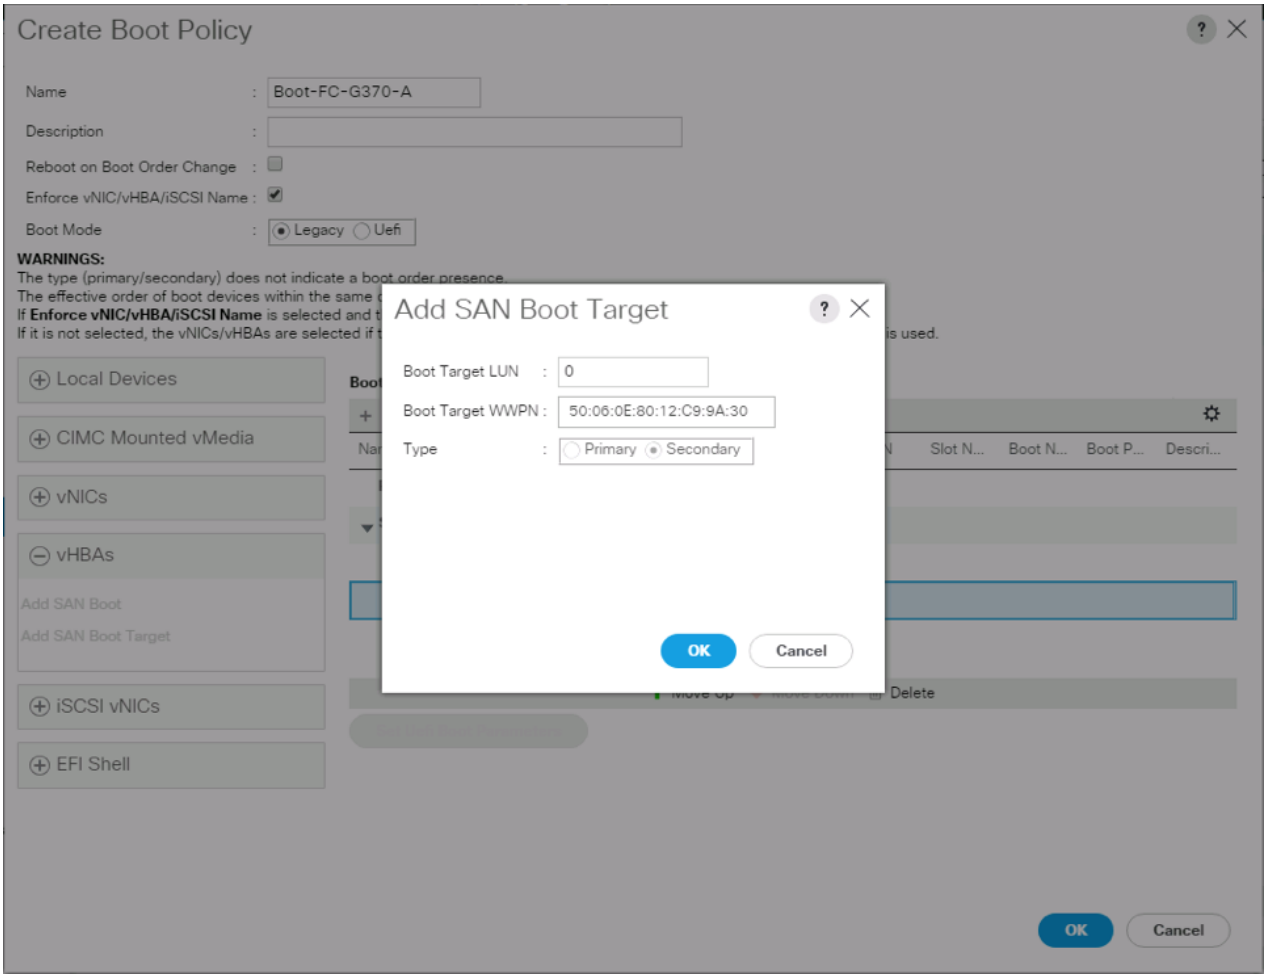

32. Click OK to add the SAN boot target.

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

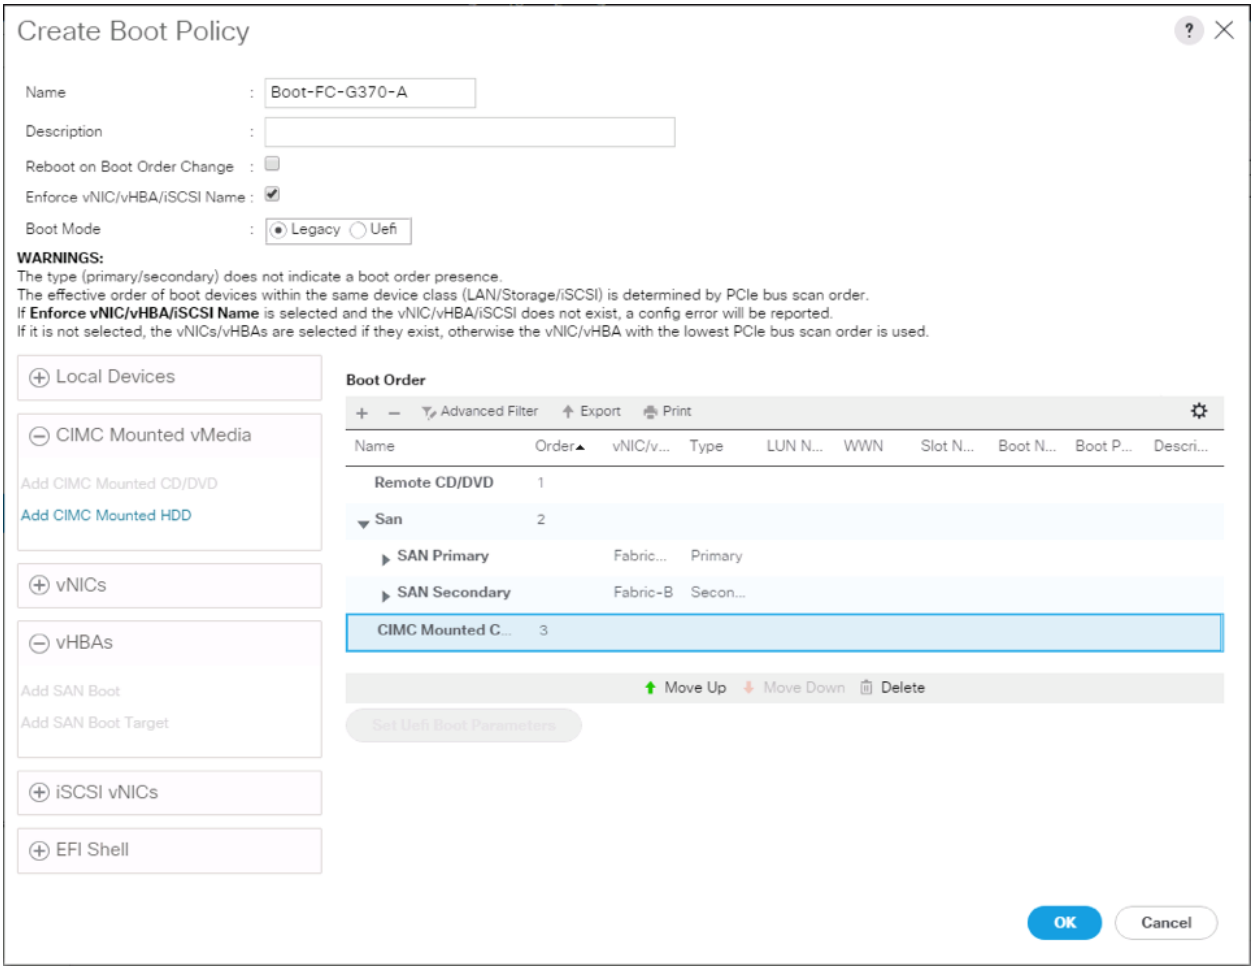

34. Click OK, then click OK again to create the boot policy.

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.

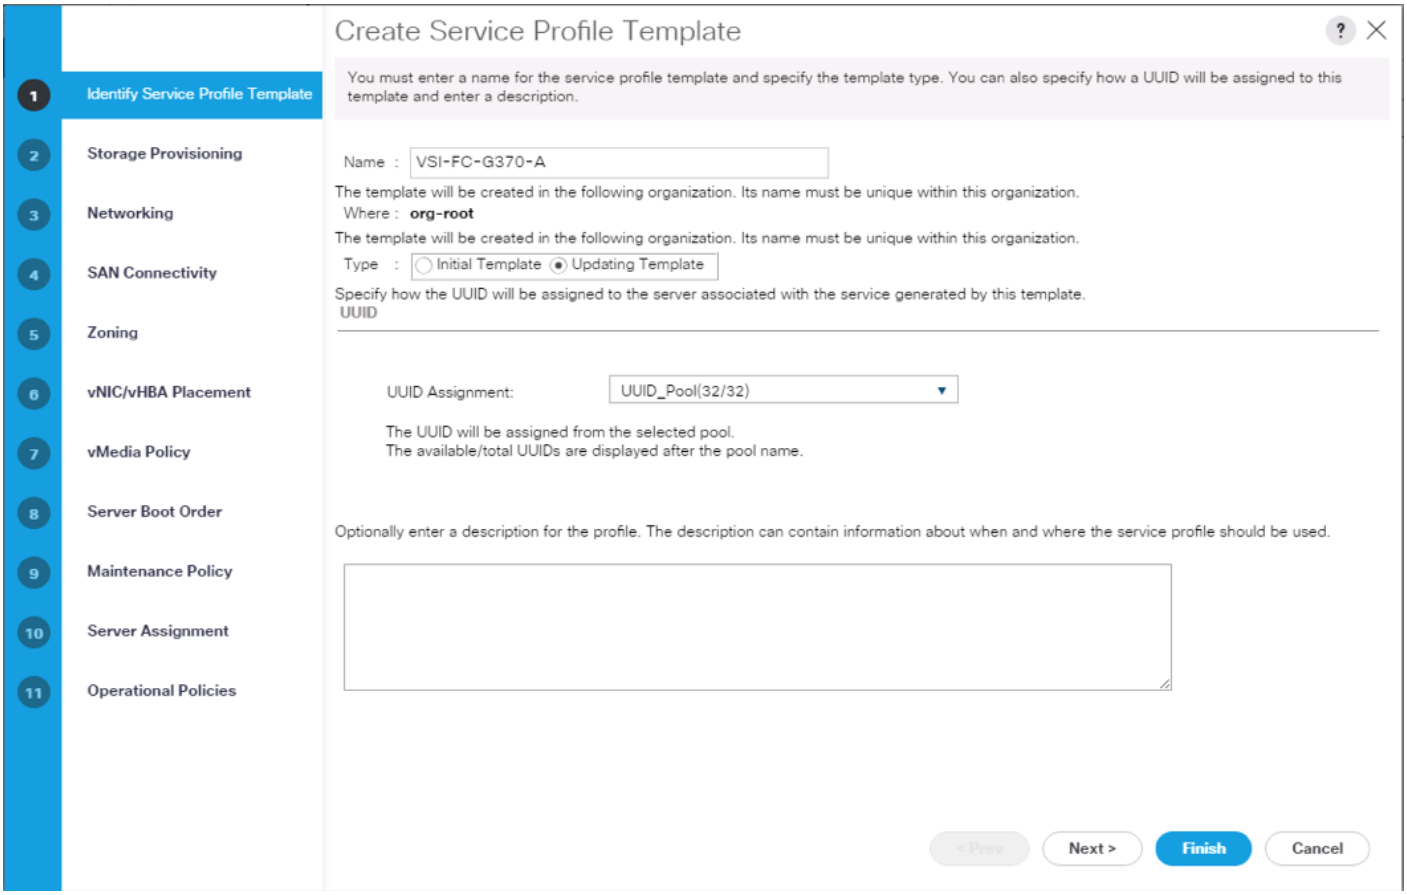

8. Click Next.

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.

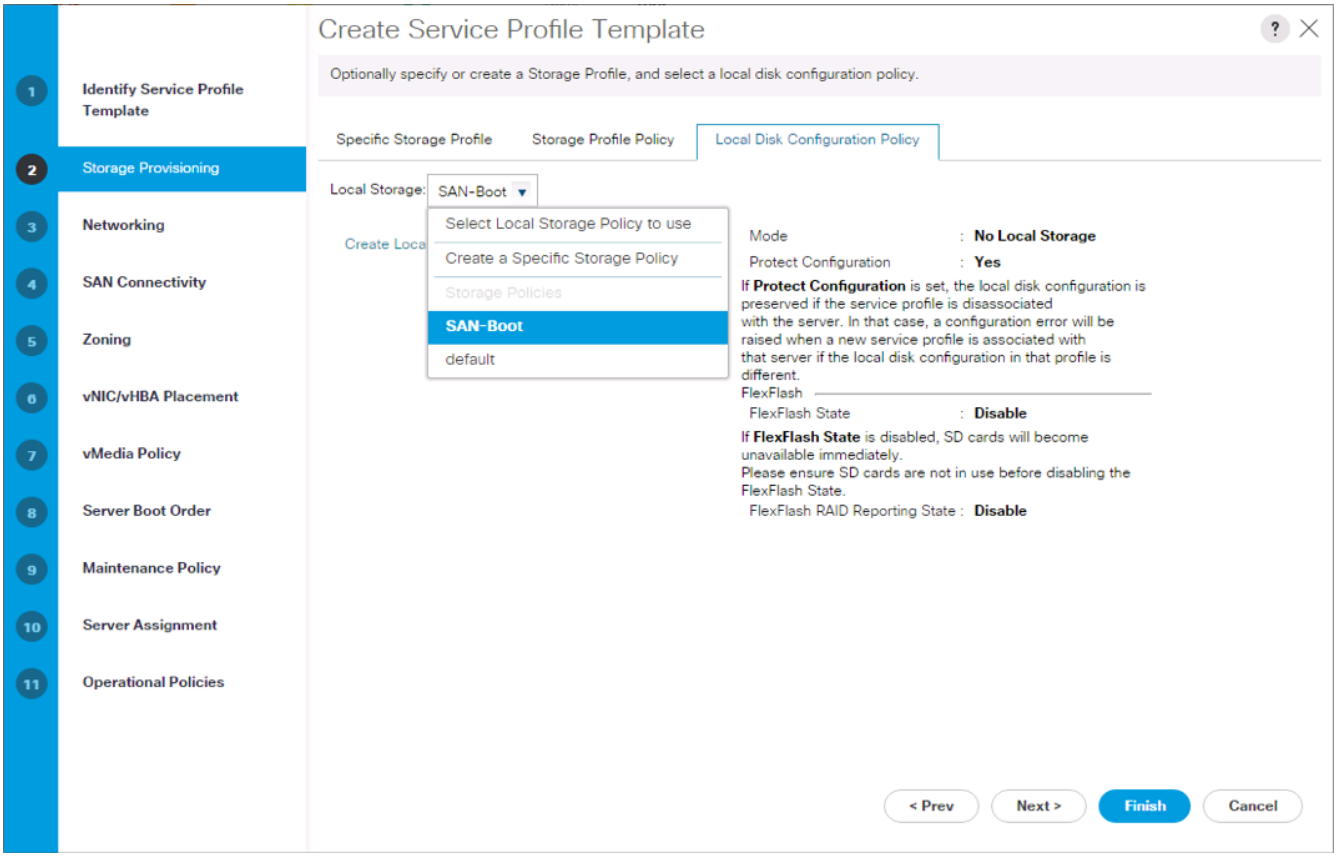

2. Click Next.

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.

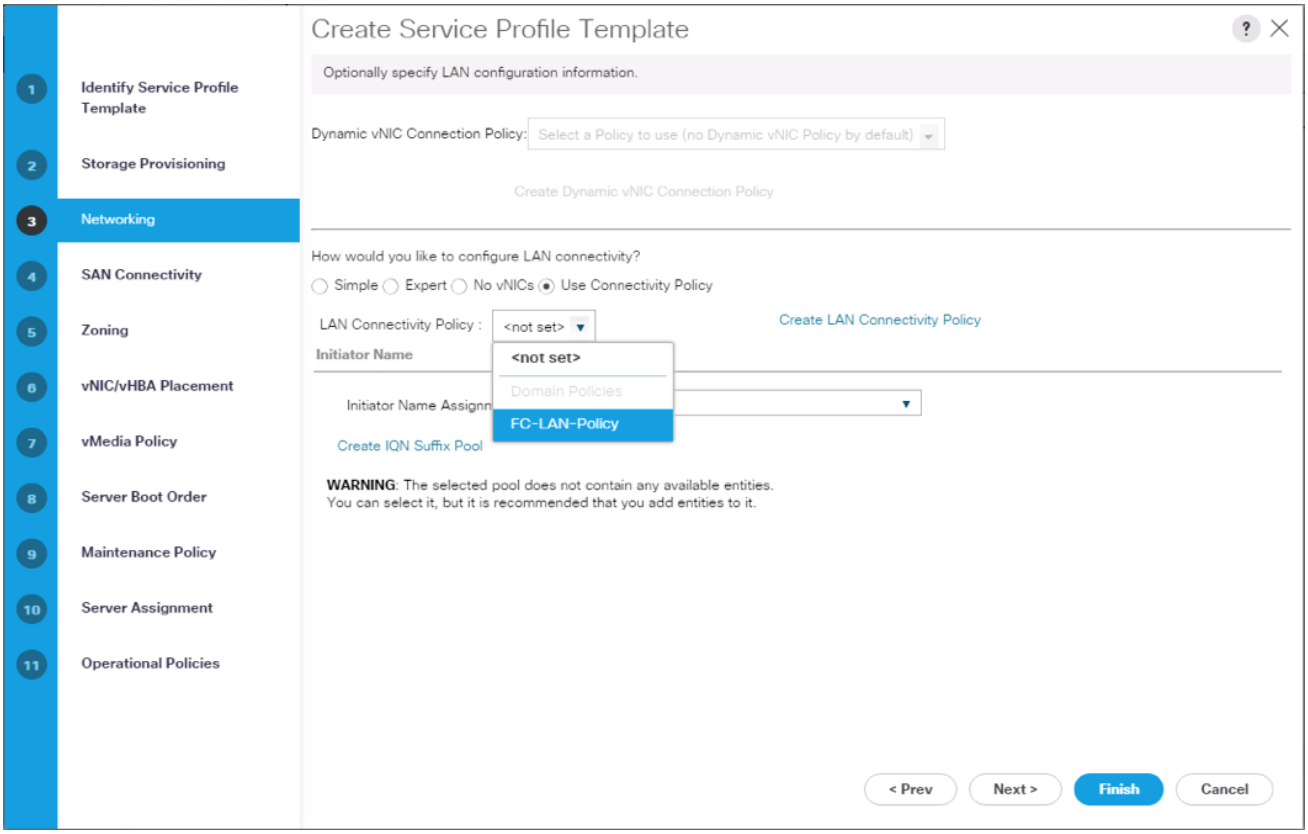

4. Click Next.

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.

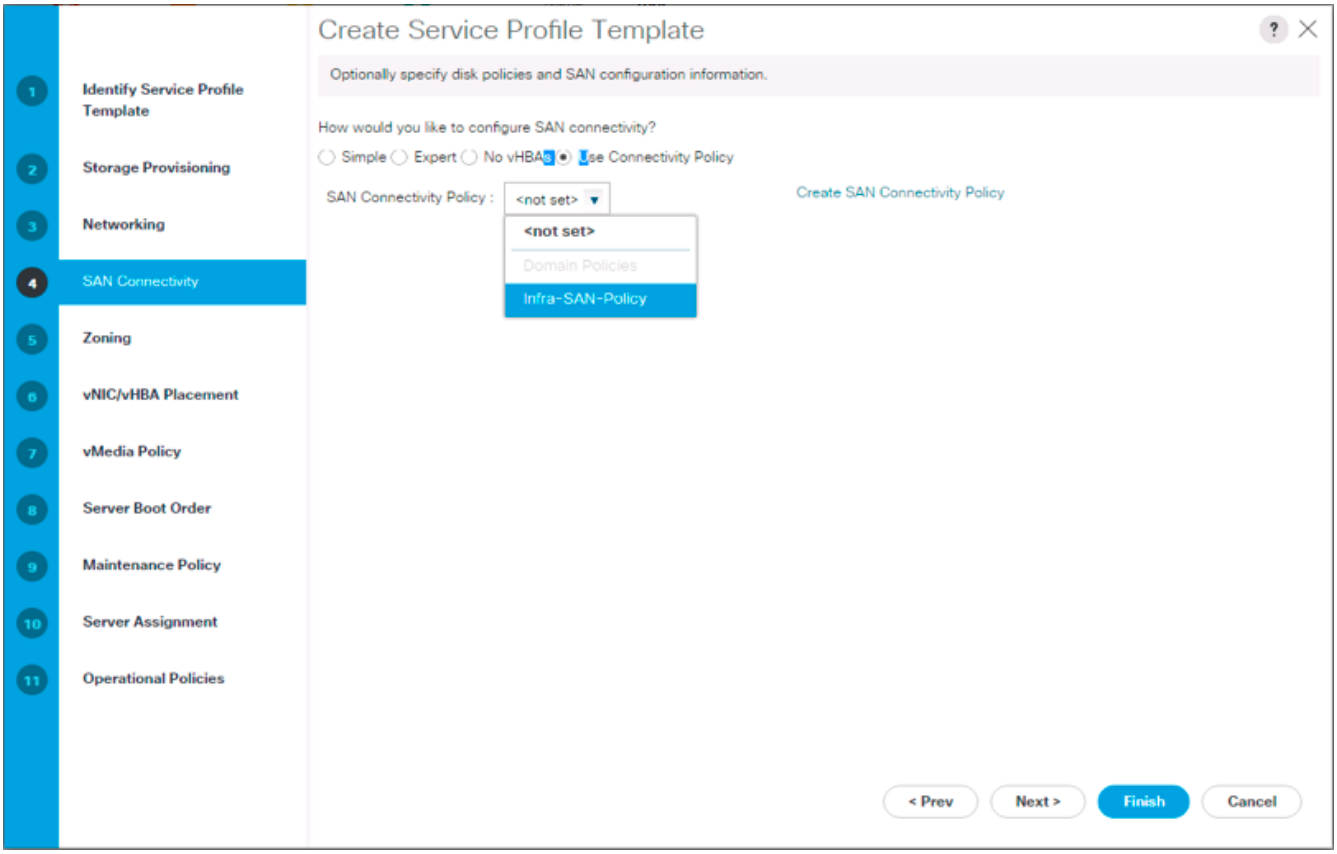

3. Click Next.

Configure Zoning Options

1. Leave Zoning configuration unspecified, and 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.

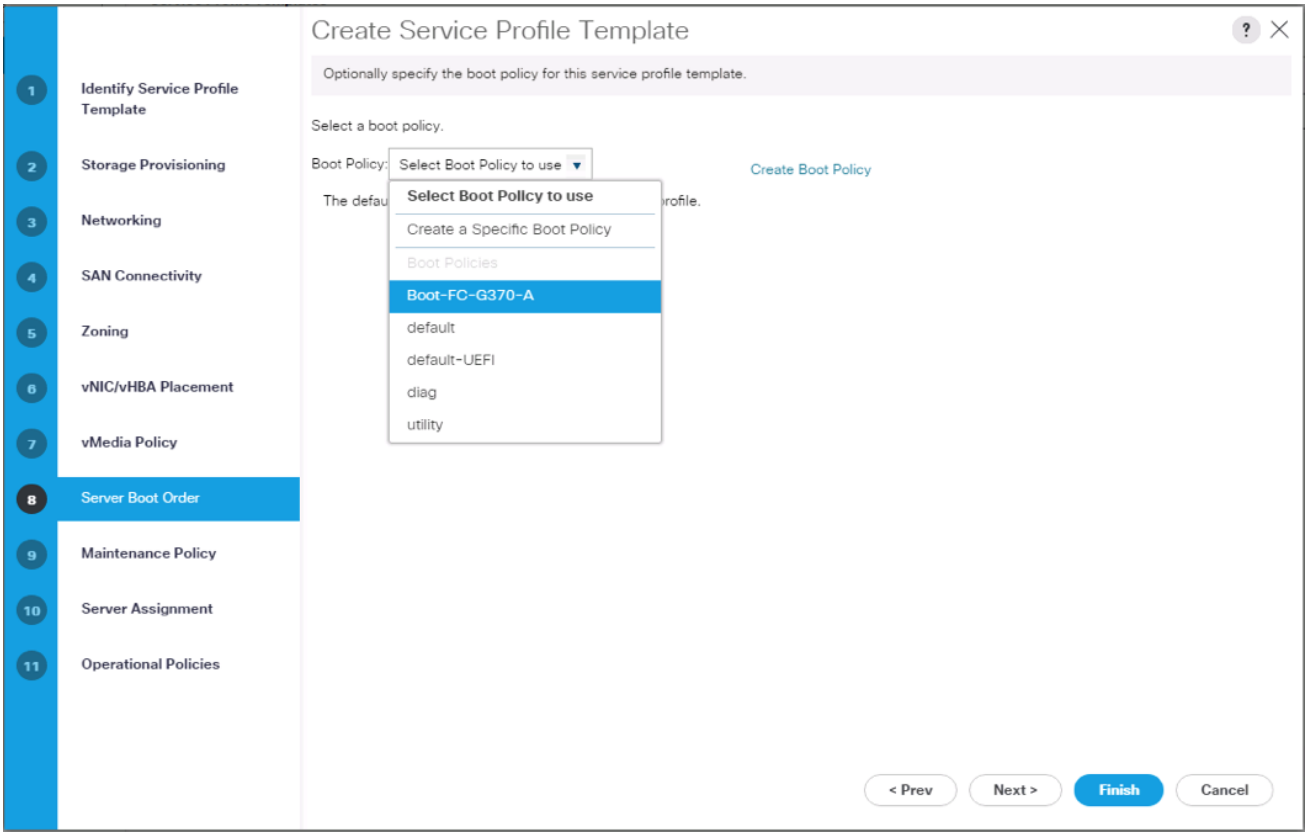

2. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

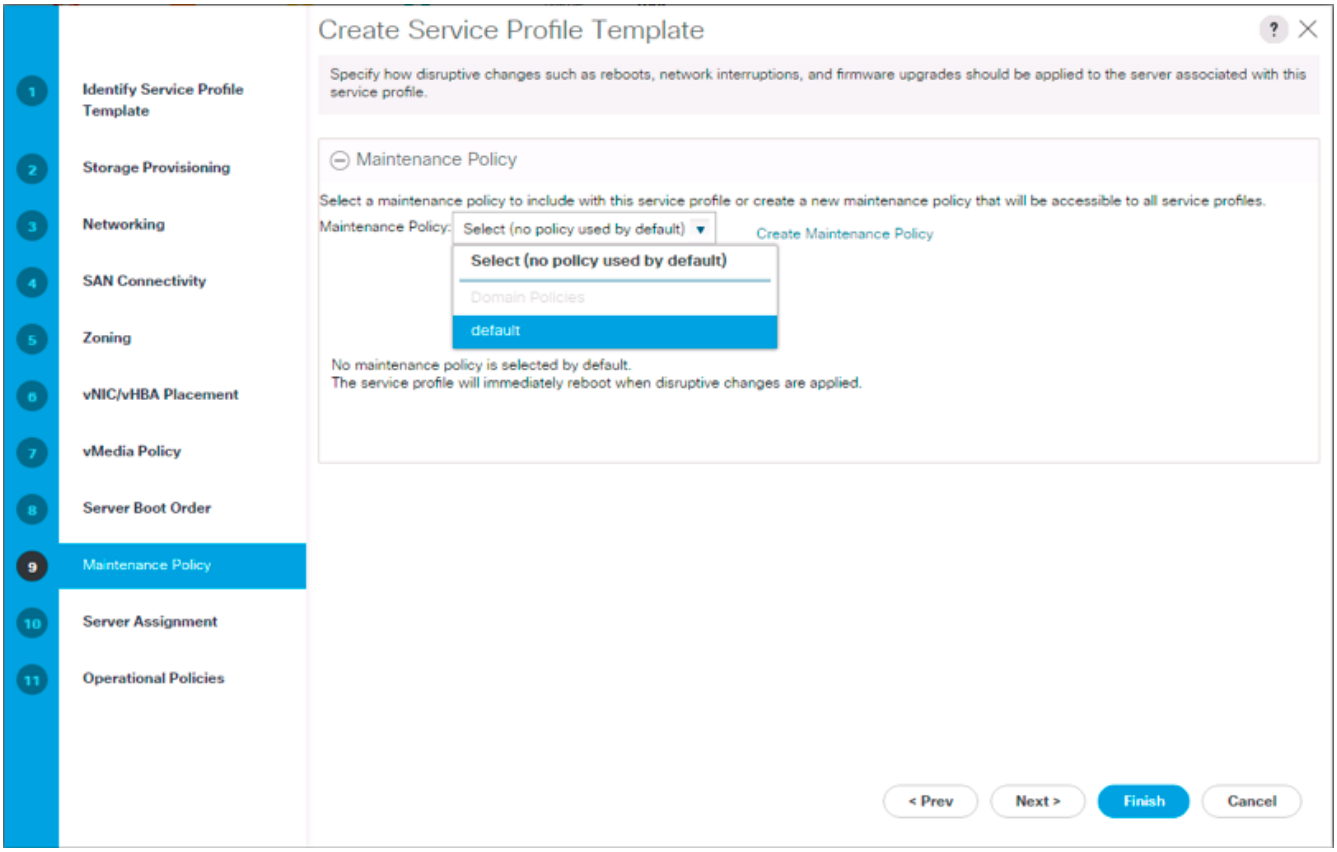

2. Click Next.

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.

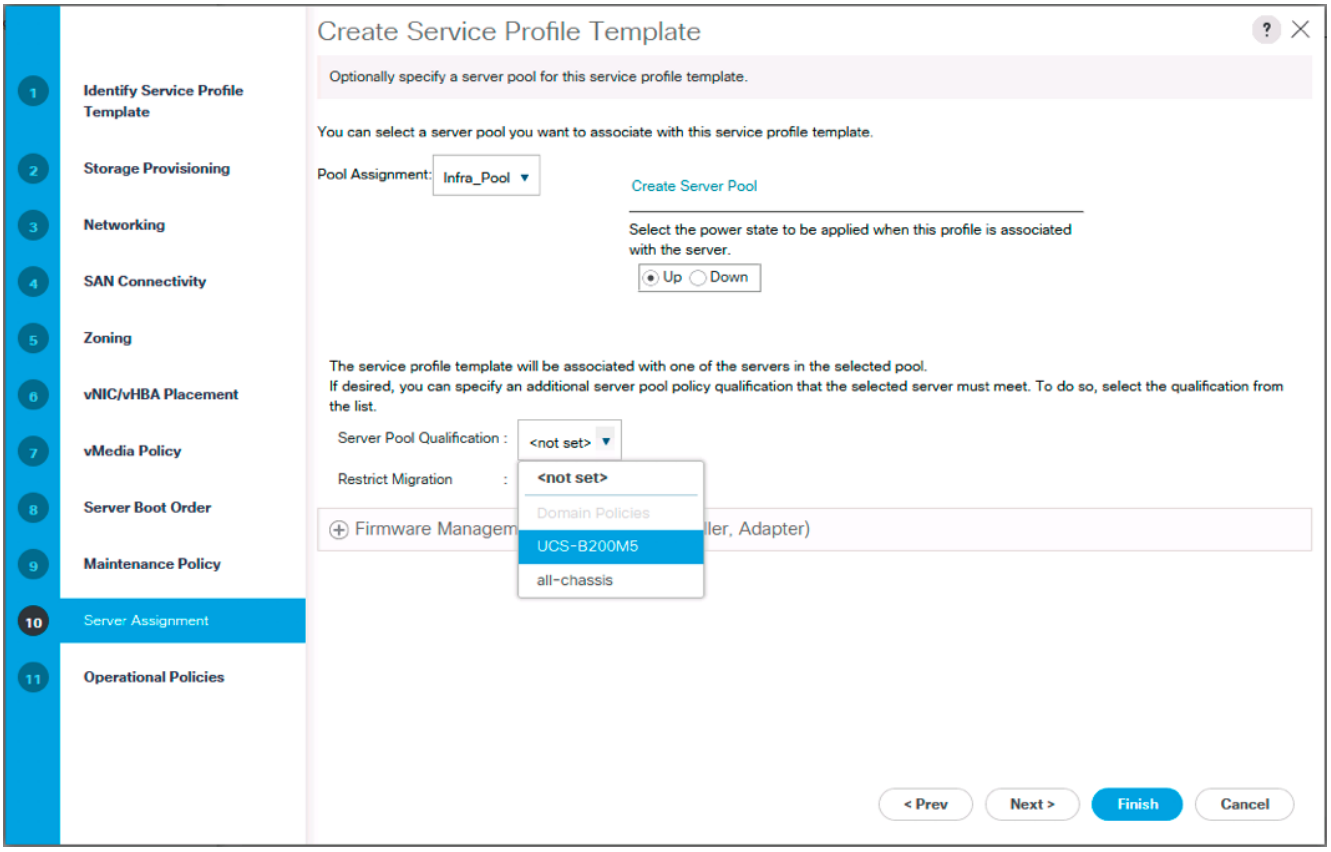

5. Click Next.

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.

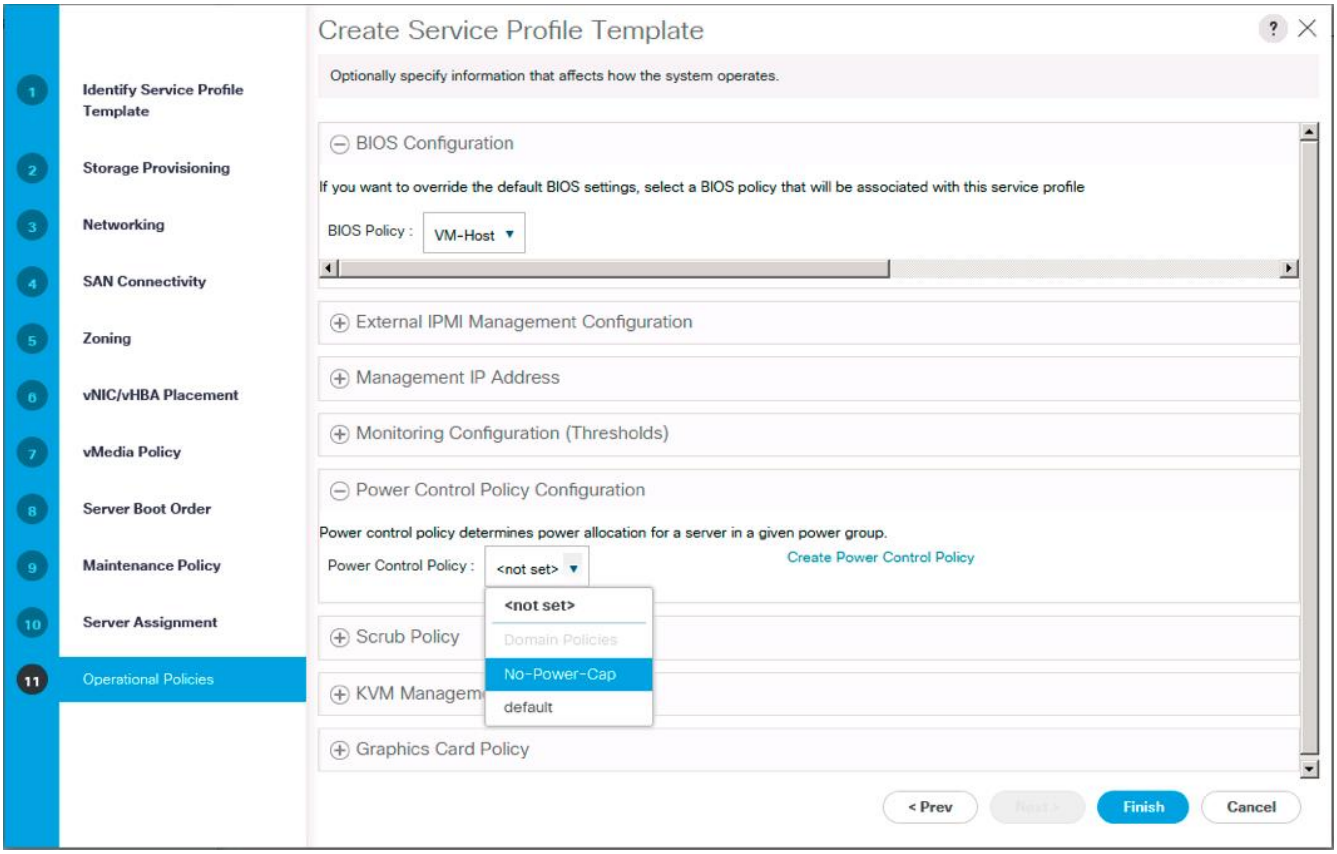

3. Click Finish to create the service profile template.

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 on the left.

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. In the right pane, select the vMedia Policy tab.

7. Under Actions, select Modify vMedia Policy.

8. Using the drop-down, select the ESXi-6.7U1-HTTP vMedia Policy.
9. Click OK then 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.

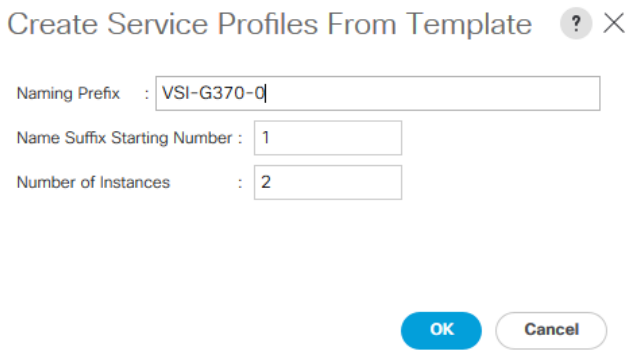

8. Click OK in the confirmation message to provision two Service Profiles.

When VMware ESXi 6.7 U1 or 6.5 U2 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.

Collect UCS Host vHBA Information for Zoning

The VSP Targets that will be used were collected from the flogi database of each MDS fabric. This is not a clear option for the UCS Server host Initiators as each Initiator WWPN will show up within the configured port-channel of the fabric without any specifics of origin that came from the MDS interface ports used to identify the VSP Targets. UCSM will be used to collect the vHBA WWPNs used as the Initiators for the provisioned Service Profiles.

To collect UCS host vHBA information for zoning, follow these steps:

1. To collect the WWPNs, follow these steps with UCSM:
2. Click the SAN icon from the Navigation pane.
3. Select Pools from the drop-down.
4. Expand WWPN Pools and select the WWPN_Pool_A.

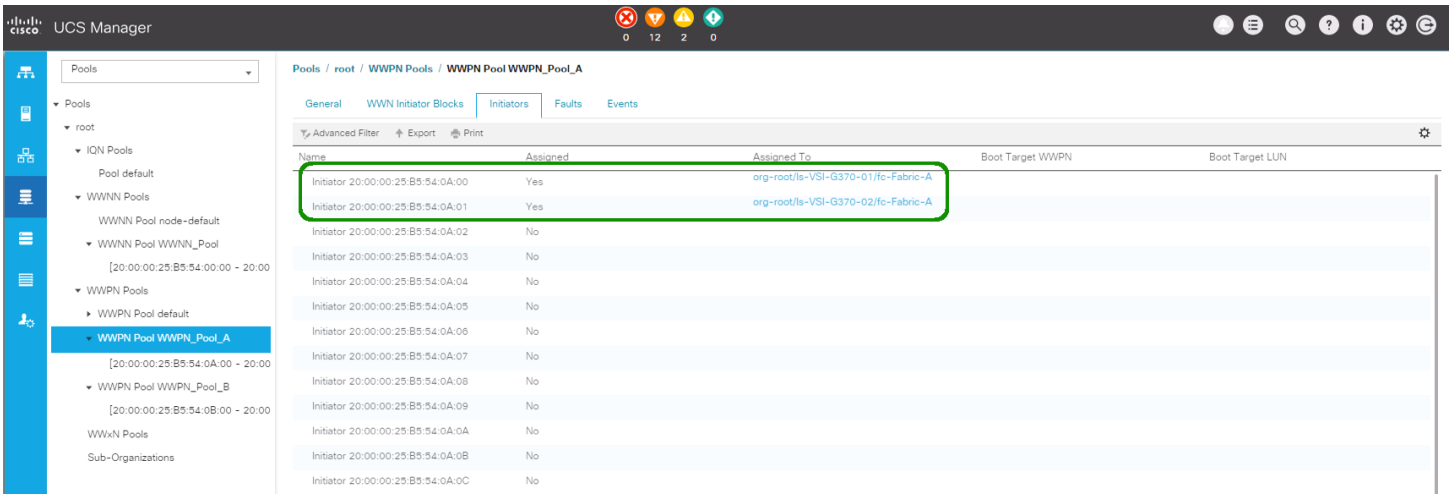

5. Identify the Fabric A Initiators assigned to the provisioned Service Profiles and add them to Table 26 .

Table 26 Fabric A G370 Service Profile Initiators

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

6. Select WWPN_Pool_B.

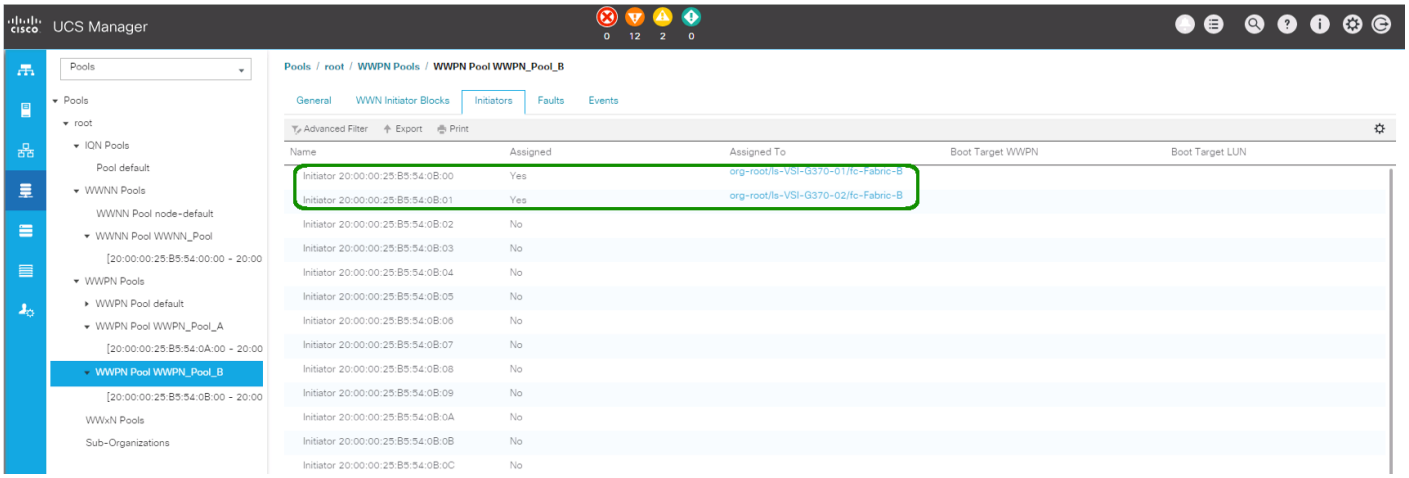

7. Identify the Fabric B Initiators assigned to the provisioned Service Profiles and add them to Table 27 .

Table 27 Fabric B G370 Service Profile Initiators

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

WWPN assignment is set to be sequential, so in most cases it can be extrapolated at initial provisioning based on WWPN Pool Suffix used, but confirmation is recommended.

DCNM Switch Registration and Zoning(Optional)

The Cisco MDS zoning used to connect the Cisco UCS and the Hitachi VSP will be configured using DCNM, which was deployed on resources independent of the Adaptive Solutions for CI data center. Deployment of DCNM is not covered in this document, instructions for the deployment of DCNM can be found here: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/11_0_1/installation/san/b_dcnm_installation_guide_for_san_11_0_1.html

If DCNM is not used in the customer environment, this section should be skipped, and the appendix covering MDS device alias creation and zoning setup through the CLI should be followed from <appendix>

Connecting to DCNM and Registering Switches

Registering of the Nexus switches is optional, but will provide enhanced port visibility as well as the option to gather performance monitoring of the Ethernet traffic. The MDS switches will need to be registered to be able to implement the device alias creation and zoning shown below.

To register the switches, follow these steps:

1. Log into the DCNM installation (URL will be https://<<DCNM_IP>> with the admin account, or provisioned account with appropriate credentials:

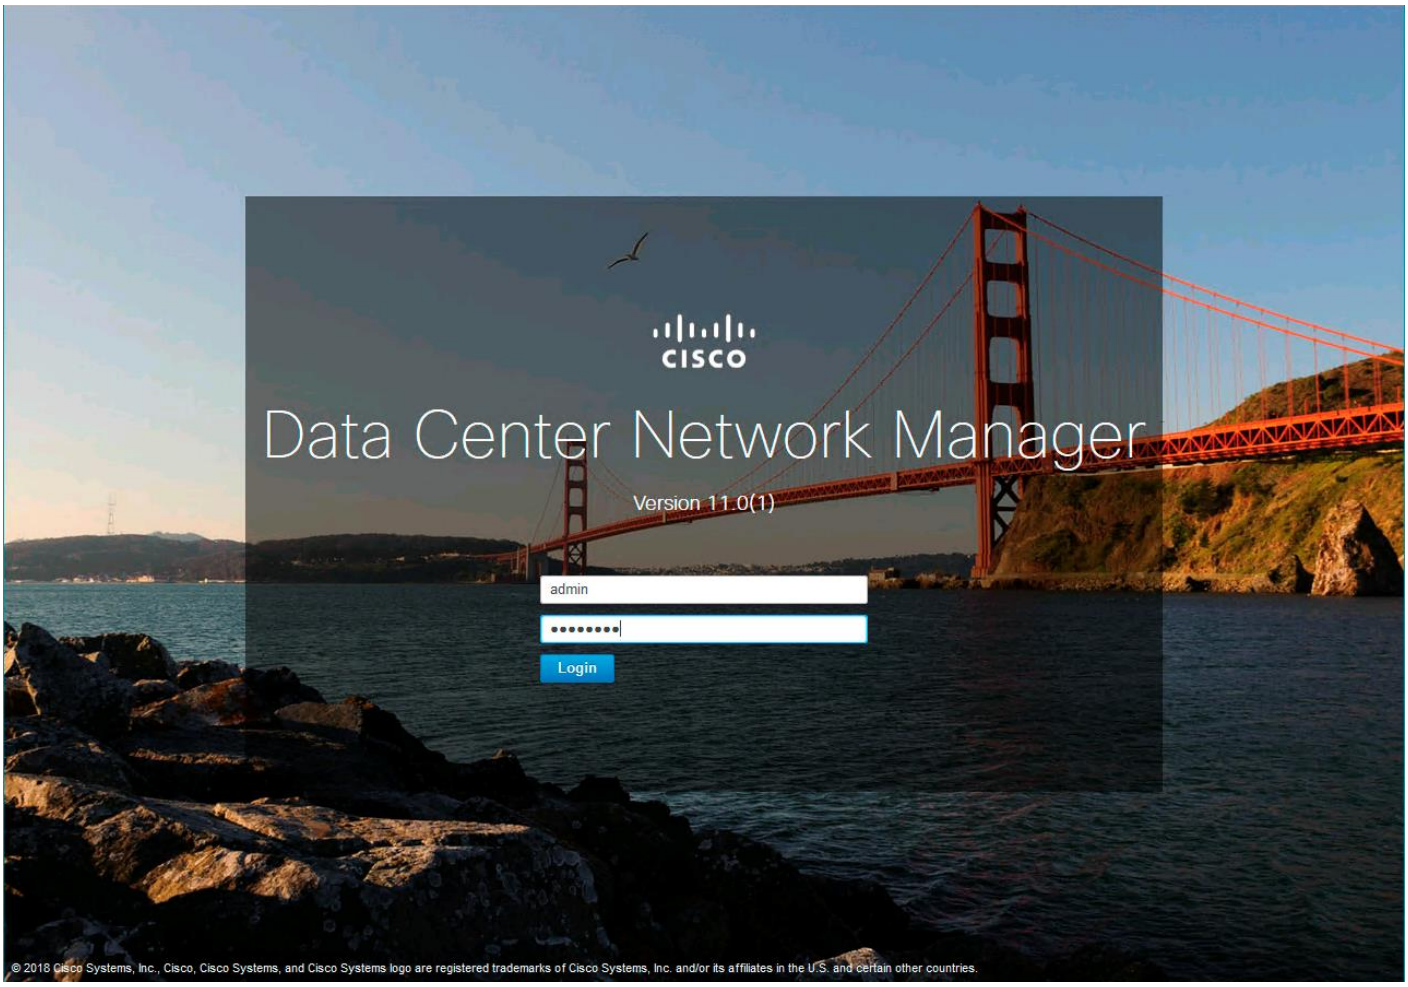

2. Provide LAN switch credentials from the initial dialogue if prompted, by clicking Yes.

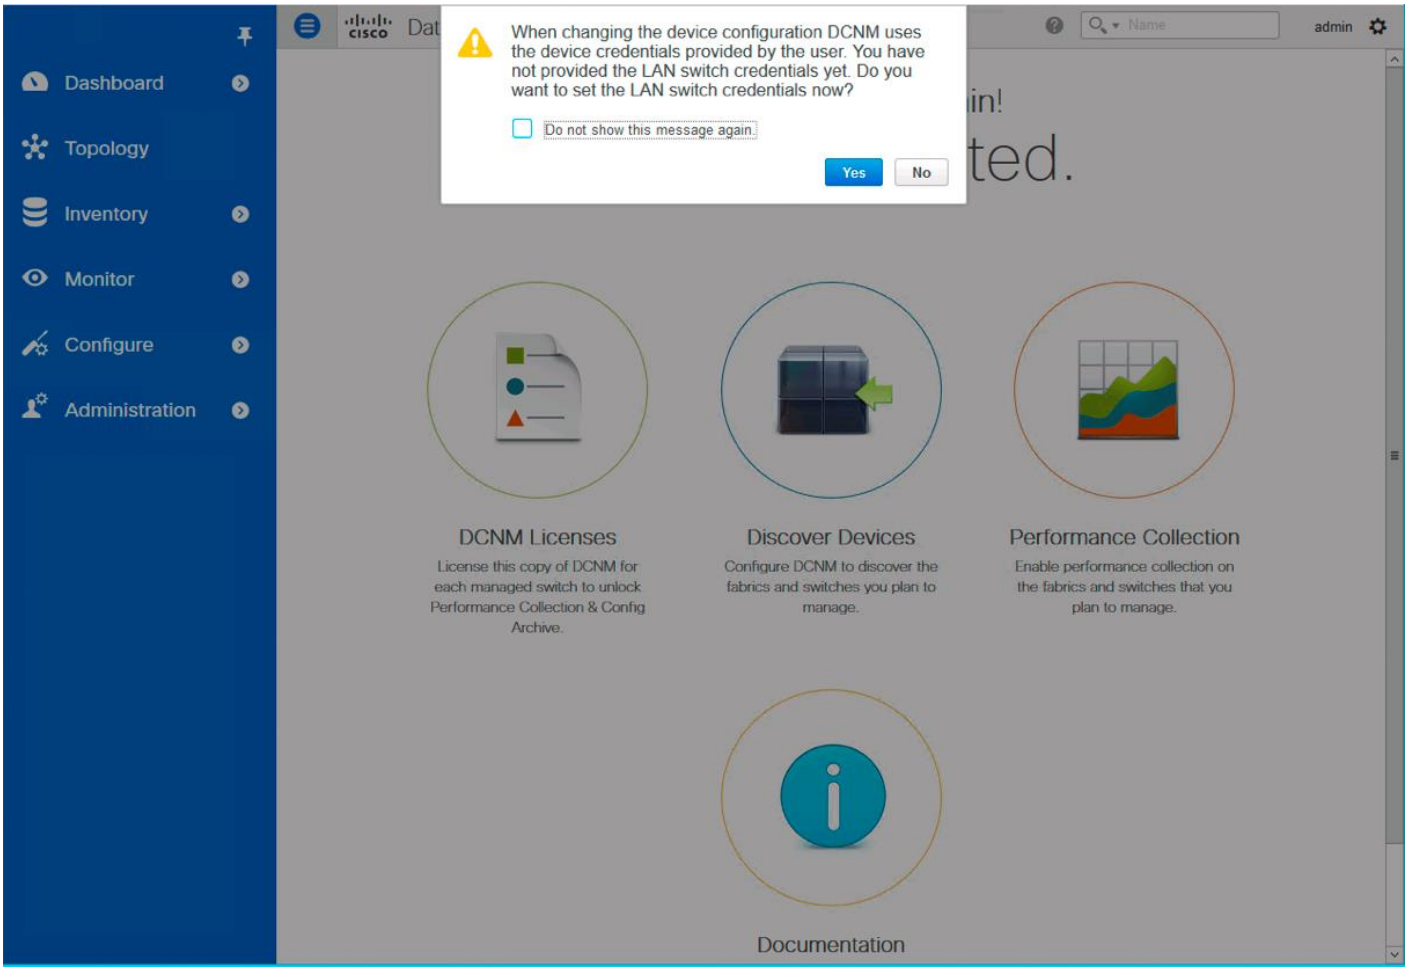

3. If not prompted, add credentials within Administration -> Credentials Management -> LAN Credentials.

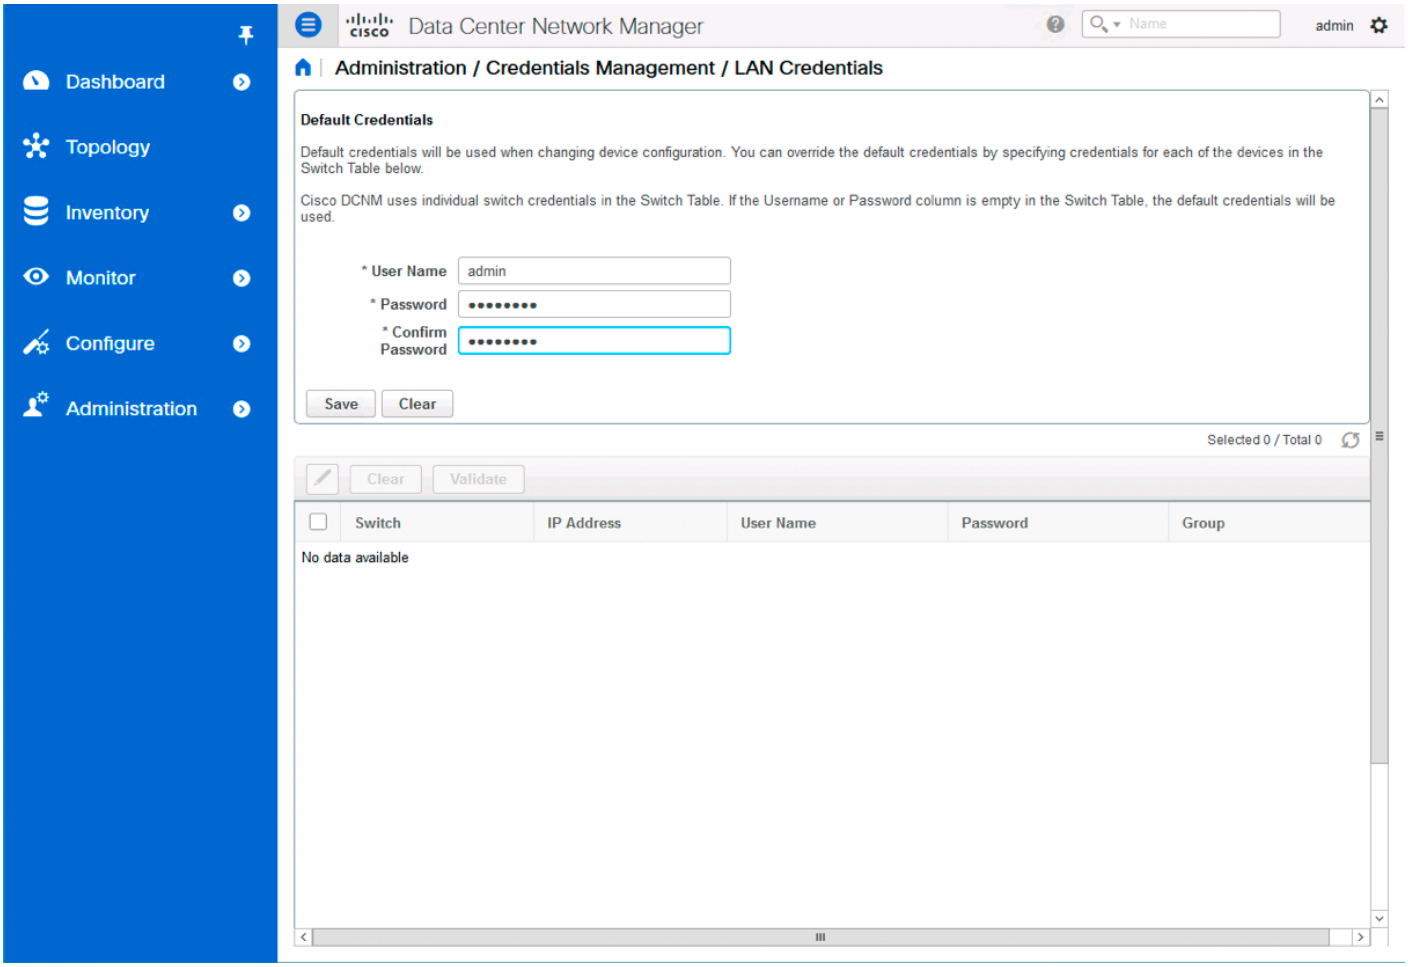

4. Click Save.

5. Click OK.

6. Add the MDS into DCNM by selecting Inventory -> Discovery -> SAN Switches.

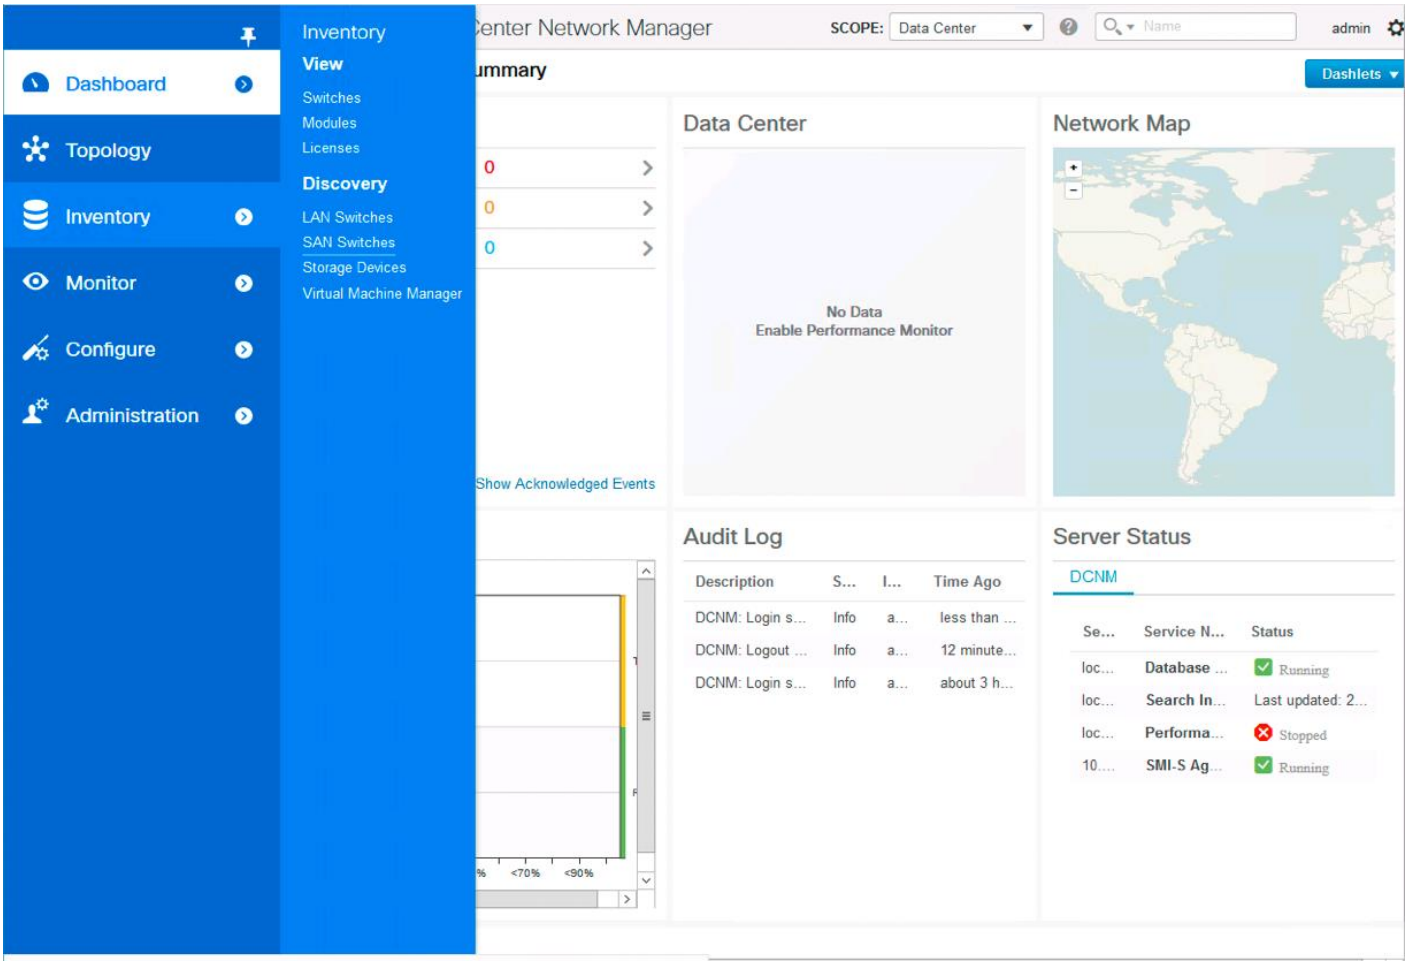

7. Click the + icon on the top left of the Inventory/Discovery/SAN Switches screen and enter the IP and credentials for the first MDS switch:

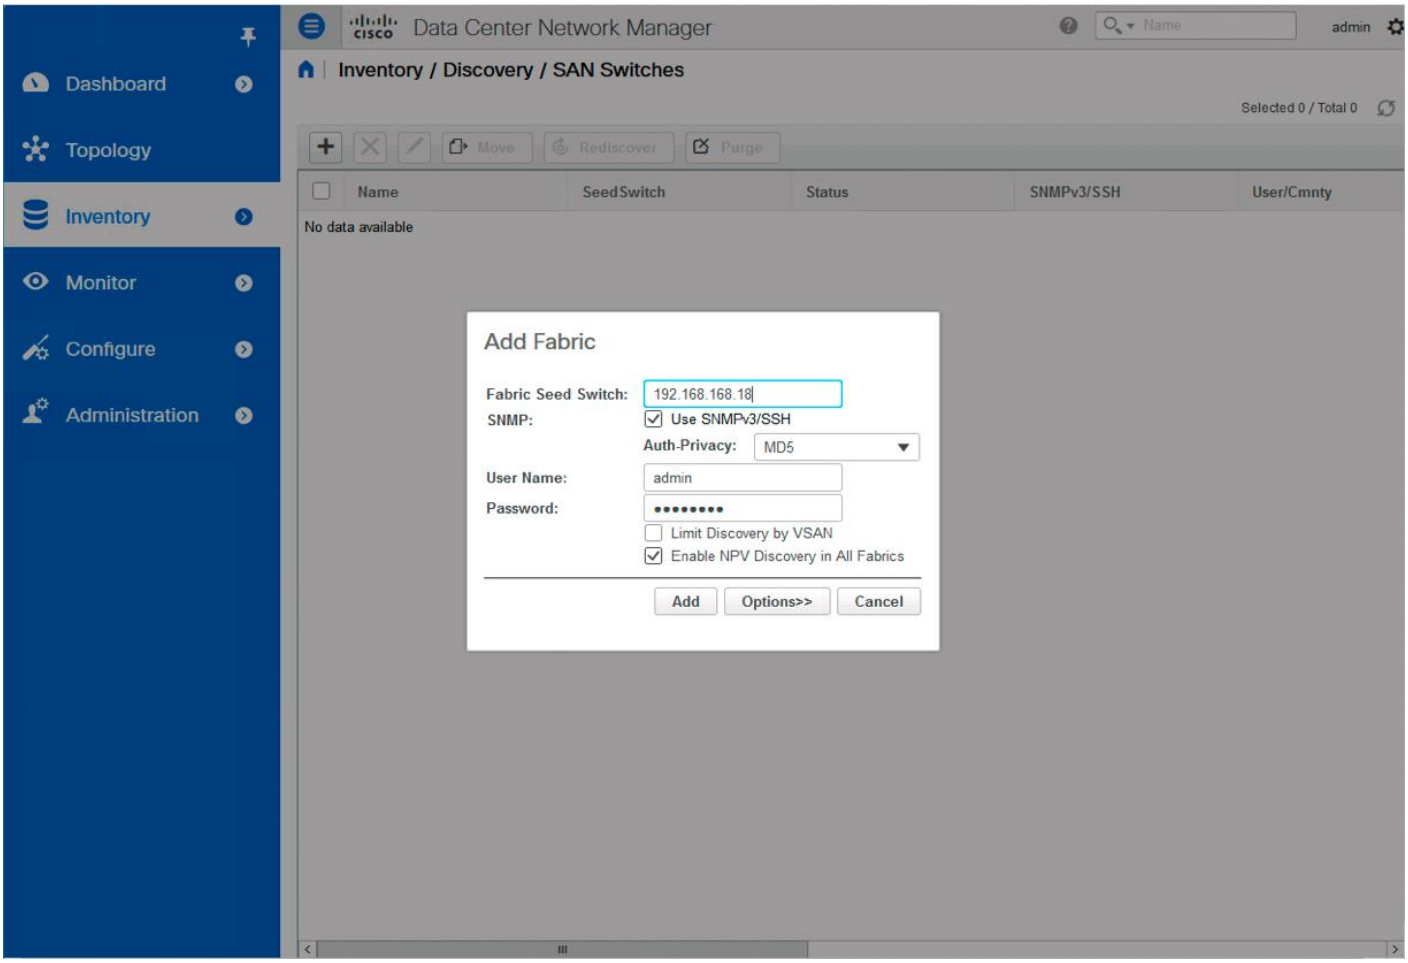

8. Click Add.

9. Repeat steps 1-8 to add the second MDS switch.

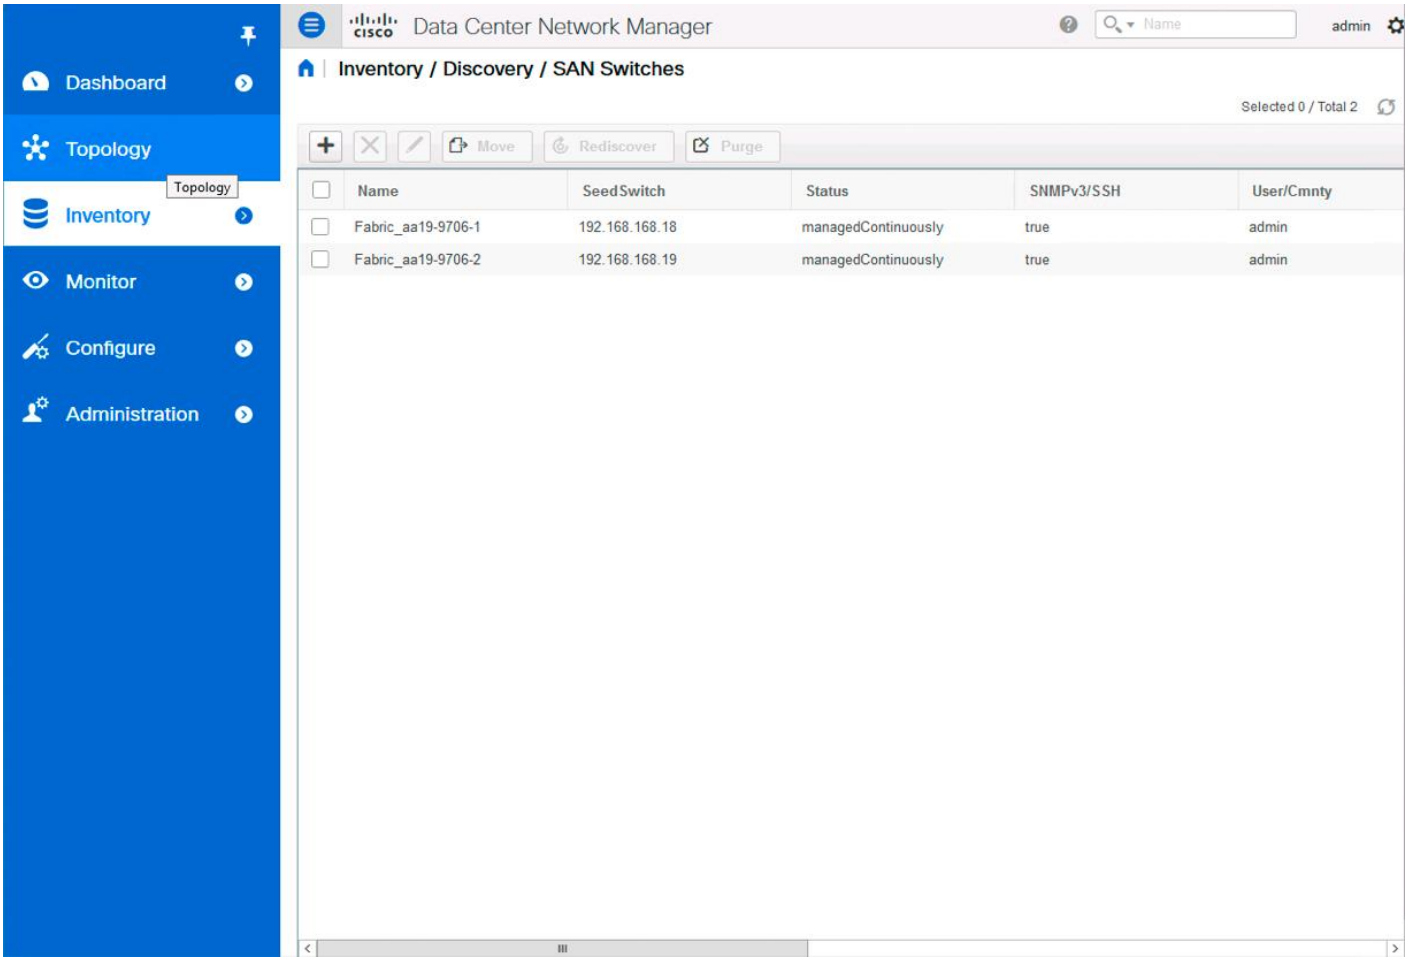

Configuring Device Aliases for the VSP and ESXi hosts

The device aliases for the MDS fabrics will be created before the zoning can occur. To create the device aliases, follow these steps:

1. Select Configure -> Device Aliases.

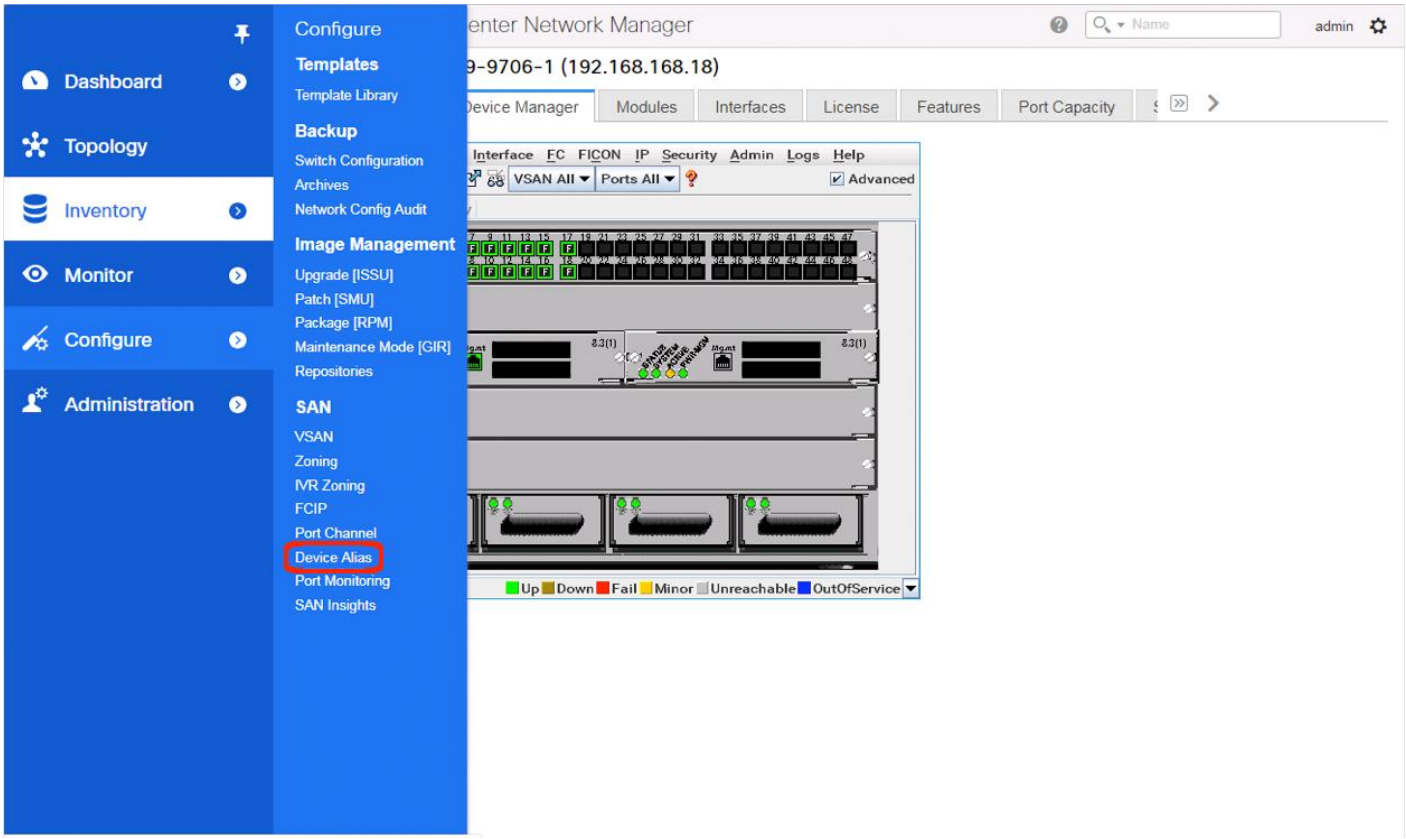

2. Select the appropriate MDS from the Fabric drop-down list and click Create to specify device aliases.

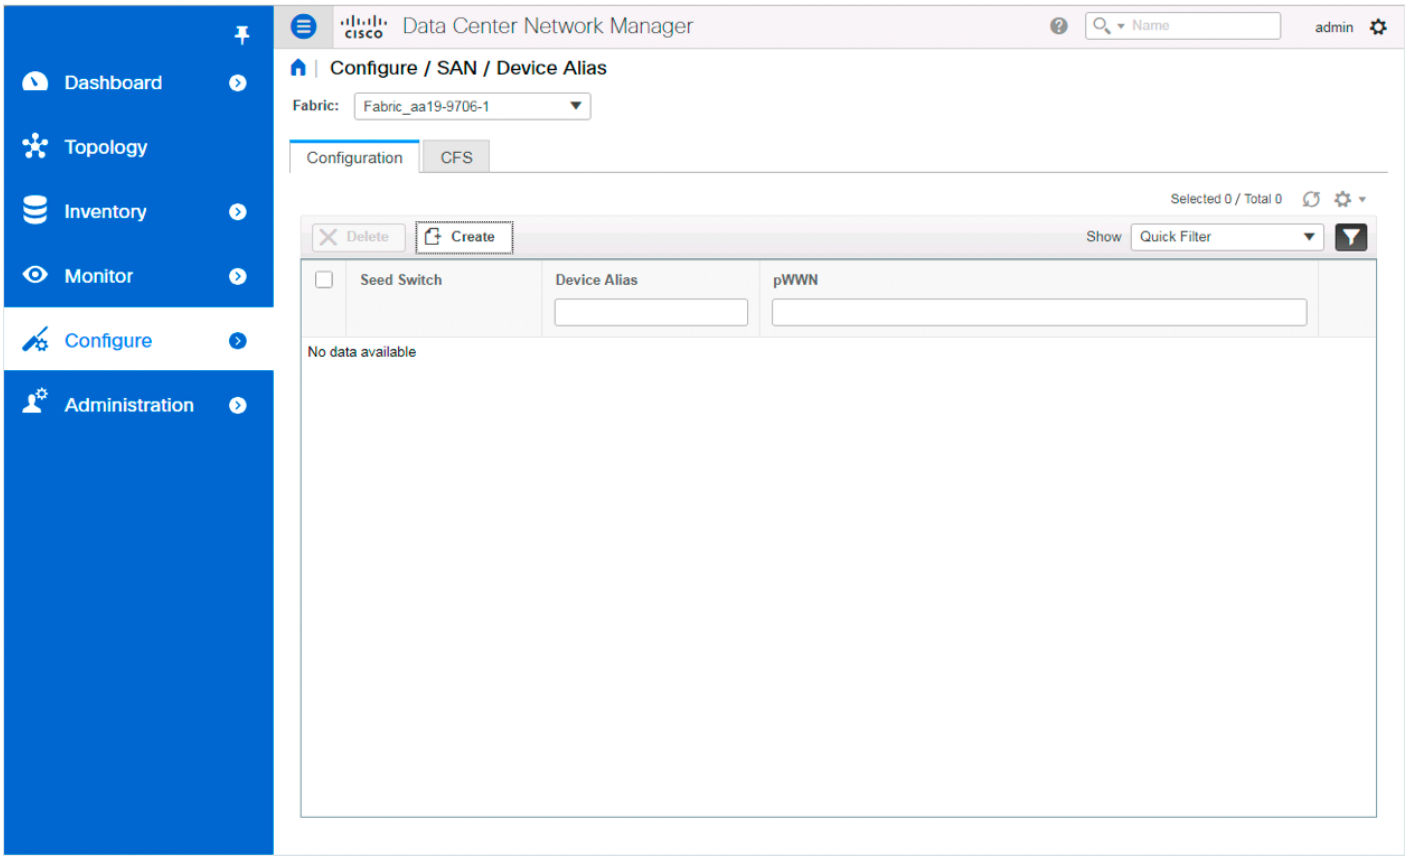

3. Continuing to use the VSP G370 to the UCS 6454 as an example, populate the following tables with data from the Table 22 and Table 23 of targets and the Table 28 and Table 29 of initiators:

Table 28 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 29 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	

4. Select a pWWN.

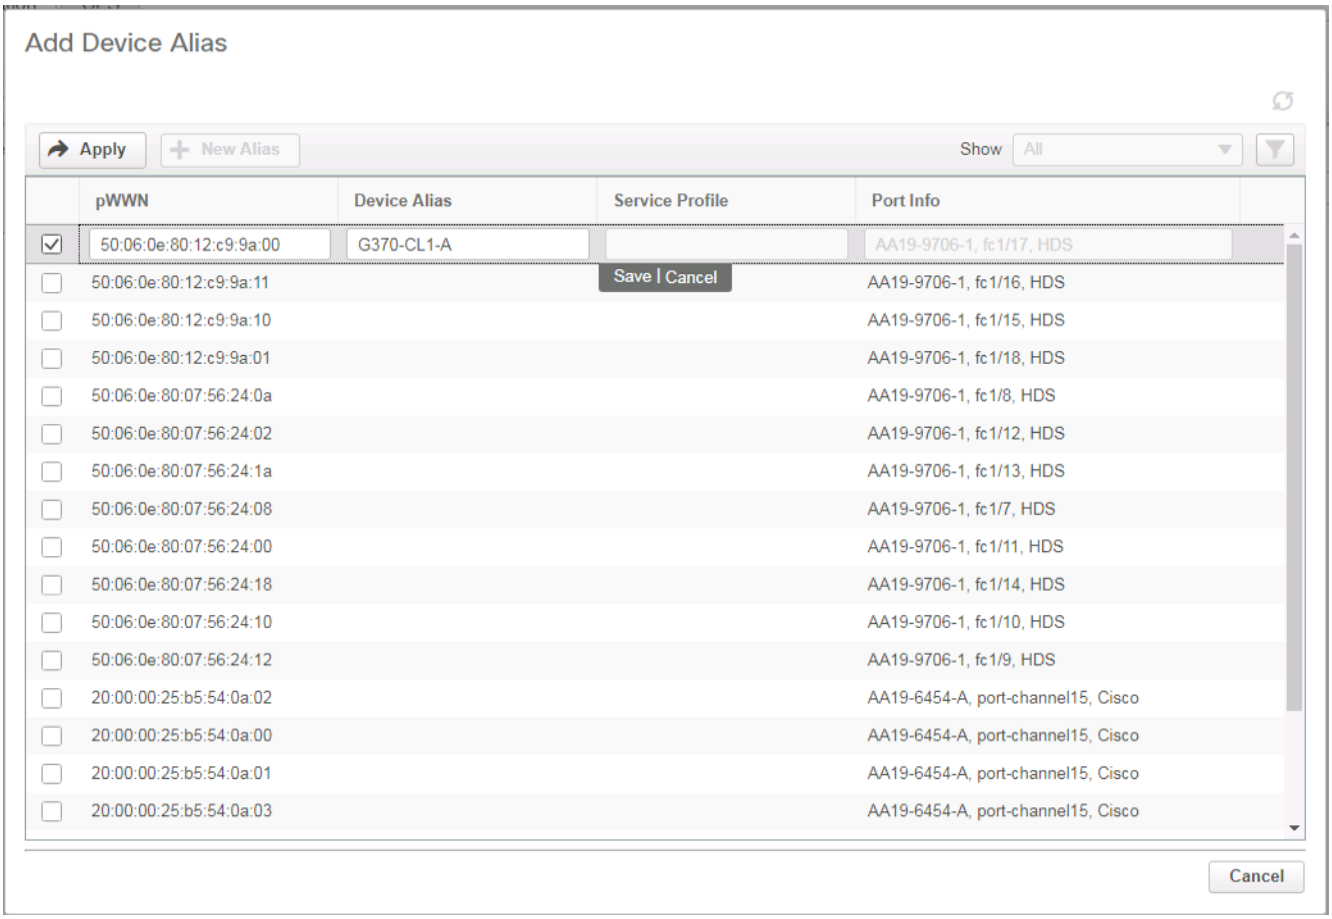

5. Click the Device Alias column and provide an appropriate alias.

6. Click Save.

7. Repeat for each VSP target and Service Profile initiator entry listed, for both MDS fabrics.

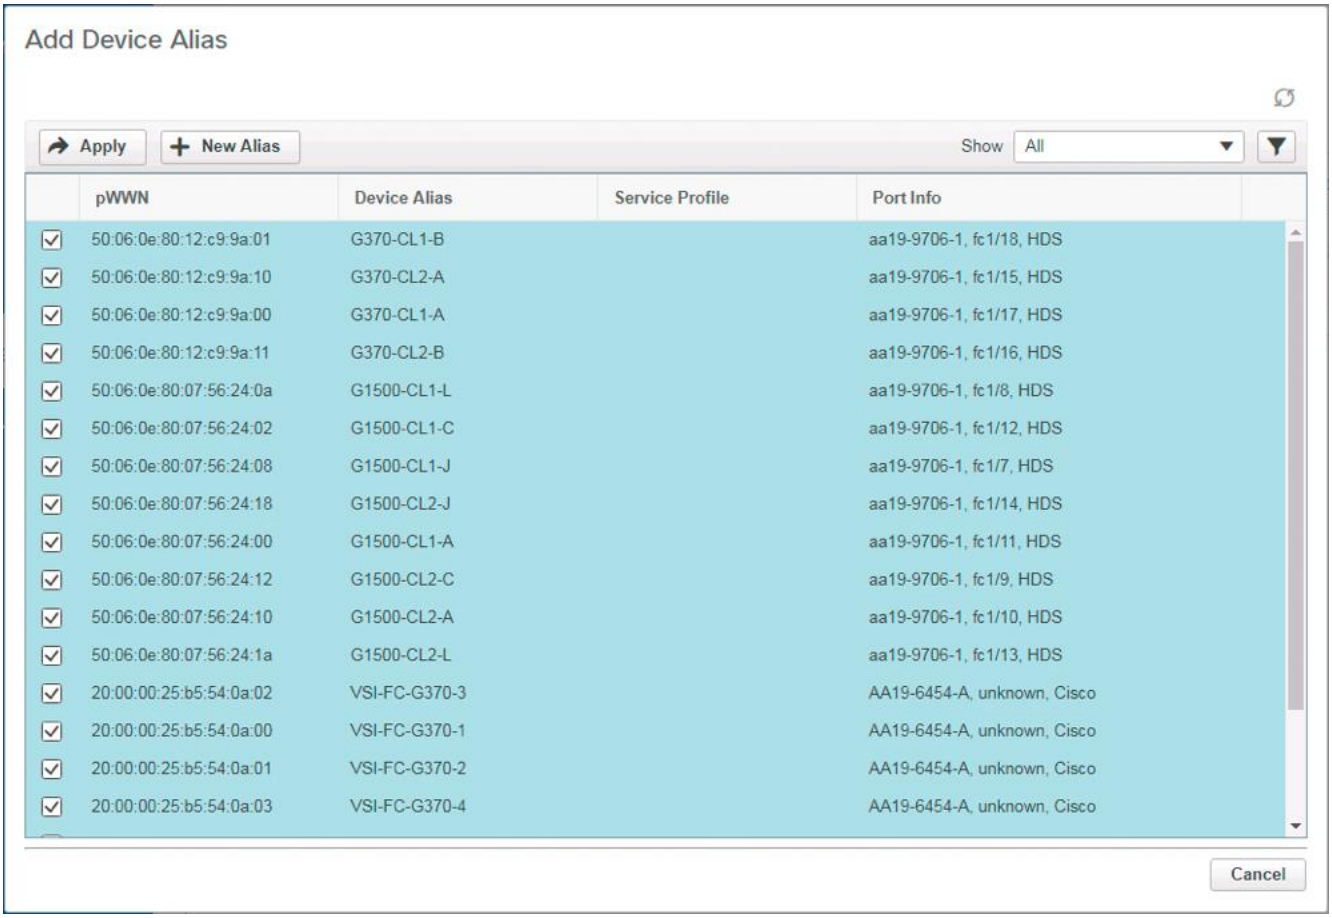

pWWN	Device Alias	Service Profile	Port Info
50:06:0e:80:12:c9:9a:01	G370-CL1-B		aa19-9706-1, fc1/18, HDS
50:06:0e:80:12:c9:9a:10	G370-CL2-A		aa19-9706-1, fc1/15, HDS
50:06:0e:80:12:c9:9a:00	G370-CL1-A		aa19-9706-1, fc1/17, HDS
50:06:0e:80:12:c9:9a:11	G370-CL2-B		aa19-9706-1, fc1/16, HDS
50:06:0e:80:07:56:24:0a	G1500-CL1-L		aa19-9706-1, fc1/8, HDS
50:06:0e:80:07:56:24:02	G1500-CL1-C		aa19-9706-1, fc1/12, HDS
50:06:0e:80:07:56:24:08	G1500-CL1-J		aa19-9706-1, fc1/7, HDS
50:06:0e:80:07:56:24:18	G1500-CL2-J		aa19-9706-1, fc1/14, HDS
50:06:0e:80:07:56:24:00	G1500-CL1-A		aa19-9706-1, fc1/11, HDS
50:06:0e:80:07:56:24:12	G1500-CL2-C		aa19-9706-1, fc1/9, HDS
50:06:0e:80:07:56:24:10	G1500-CL2-A		aa19-9706-1, fc1/10, HDS
50:06:0e:80:07:56:24:1a	G1500-CL2-L		aa19-9706-1, fc1/13, HDS
20:00:00:25:b5:54:0a:02	VSI-FC-G370-3		AA19-6454-A, unknown, Cisco
20:00:00:25:b5:54:0a:00	VSI-FC-G370-1		AA19-6454-A, unknown, Cisco
20:00:00:25:b5:54:0a:01	VSI-FC-G370-2		AA19-6454-A, unknown, Cisco
20:00:00:25:b5:54:0a:03	VSI-FC-G370-4		AA19-6454-A, unknown, Cisco

8. Click Apply.
9. Repeat this process for the other fabric, creating device aliases for the attached devices.

Create Host Zoning

To create host zoning, follow these steps:

1. Select Configure and pick Zoning within the SAN sub-section.

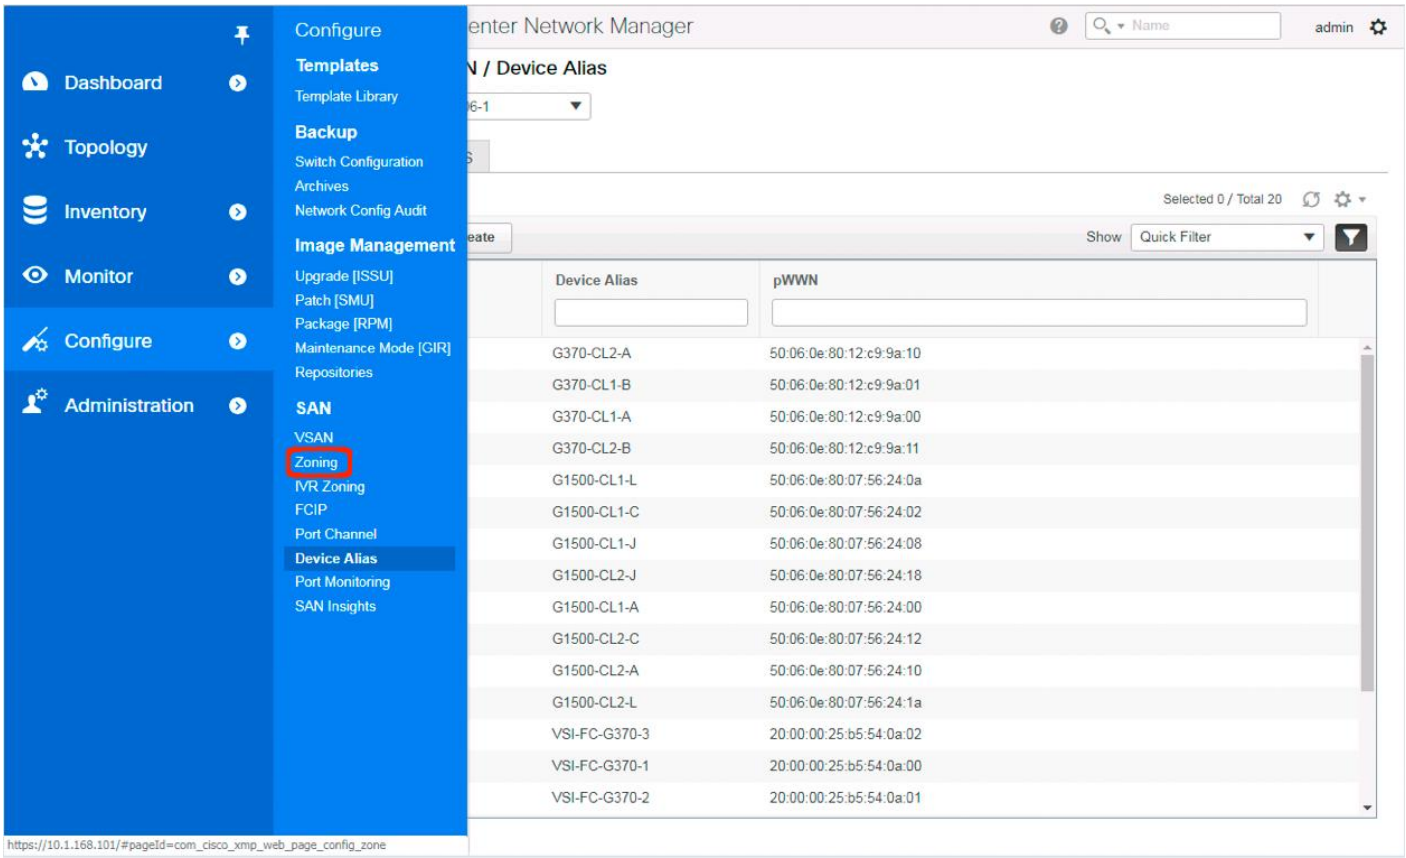

2. Adjust the VSAN to be appropriate for the zoning between hosts and the VSP.

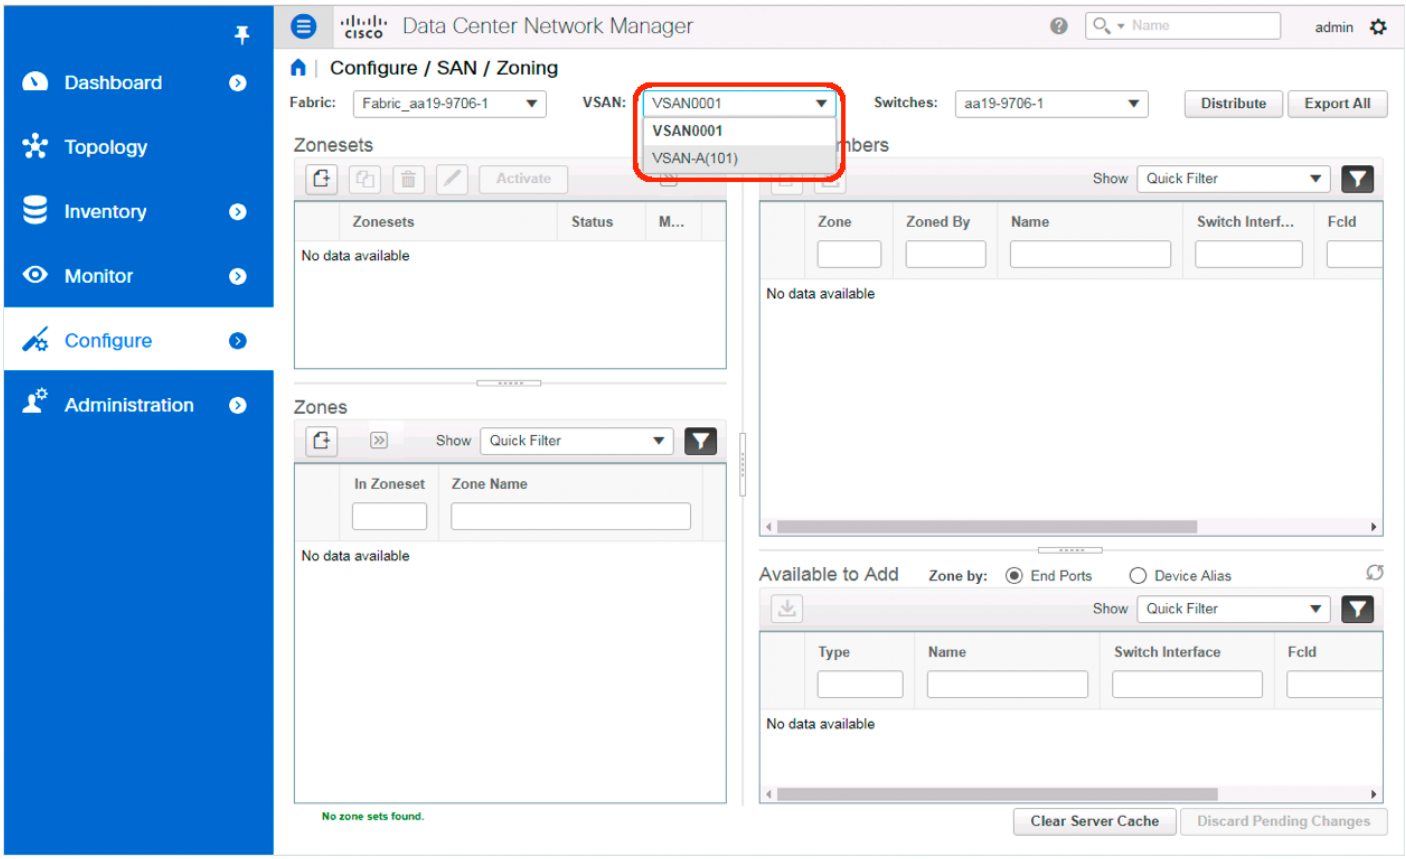

3. Click the Create Zoneset button within Zonesets.

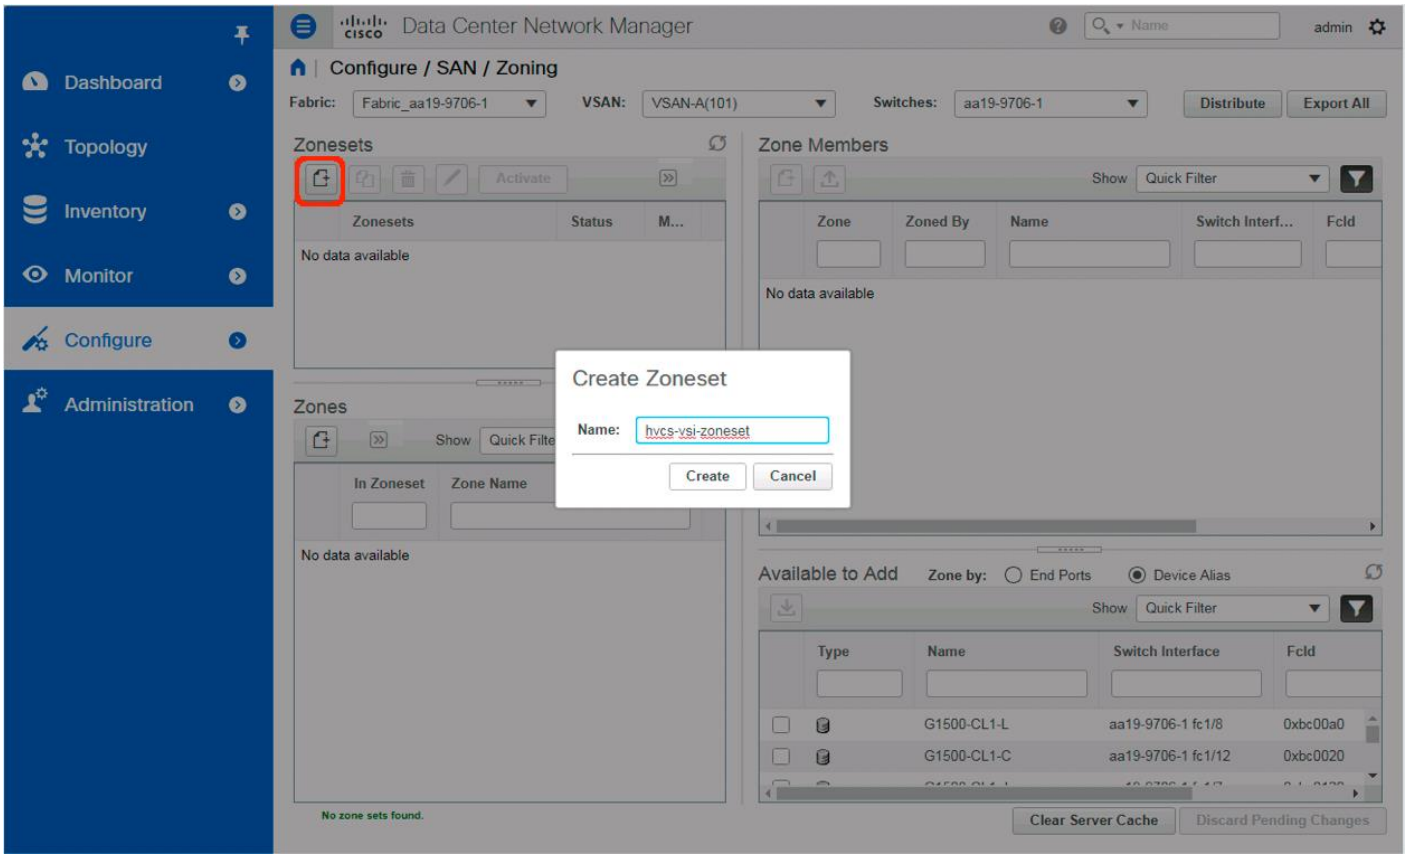

4. With the new zoneset selected, click the Create Zone button within Zones.

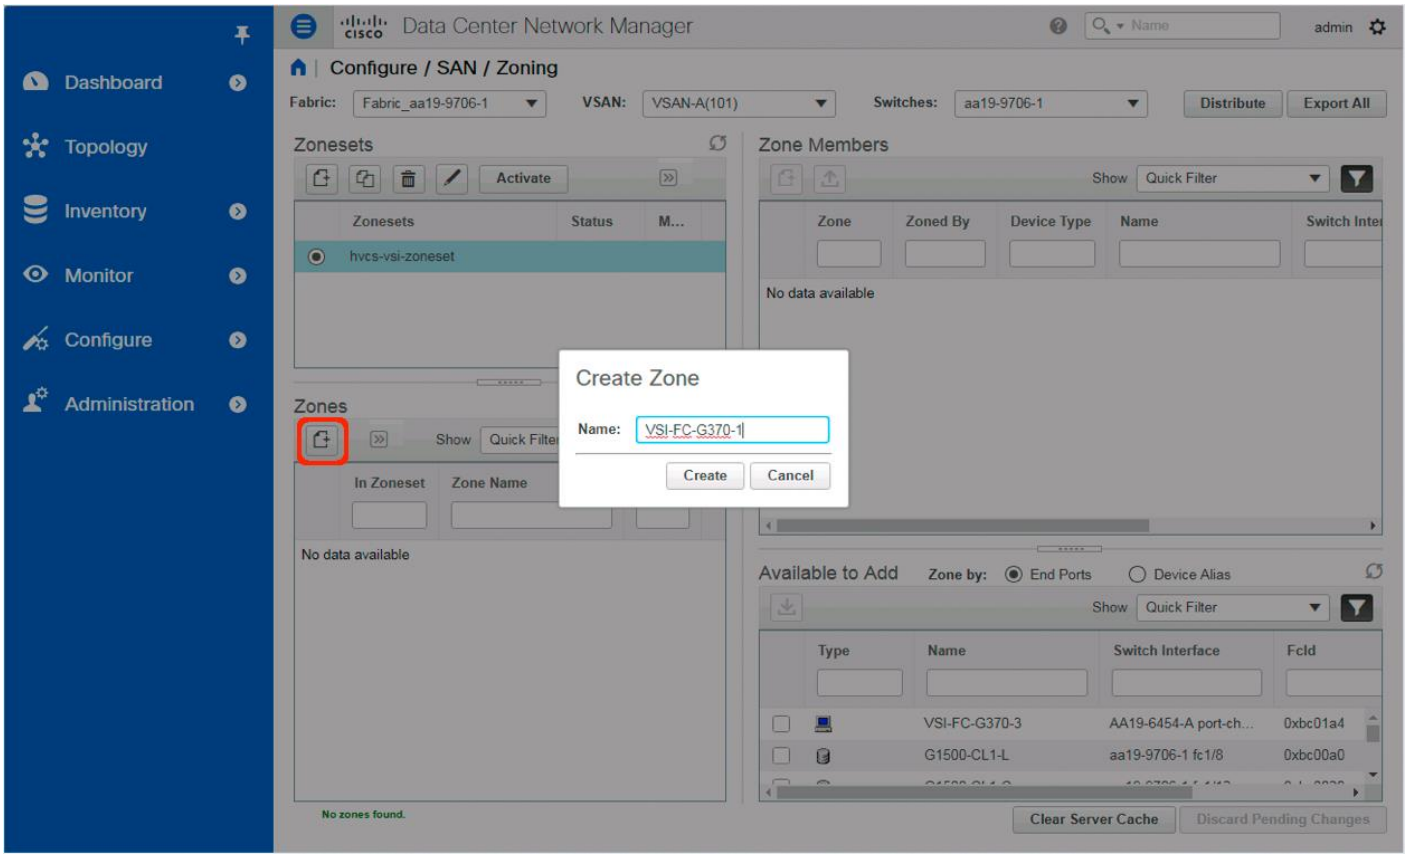

5. Specify an appropriate name for the zone.

6. Click Create.

7. Select the newly created zone and select the appropriate host device alias from the bottom right with the Available to Add section:

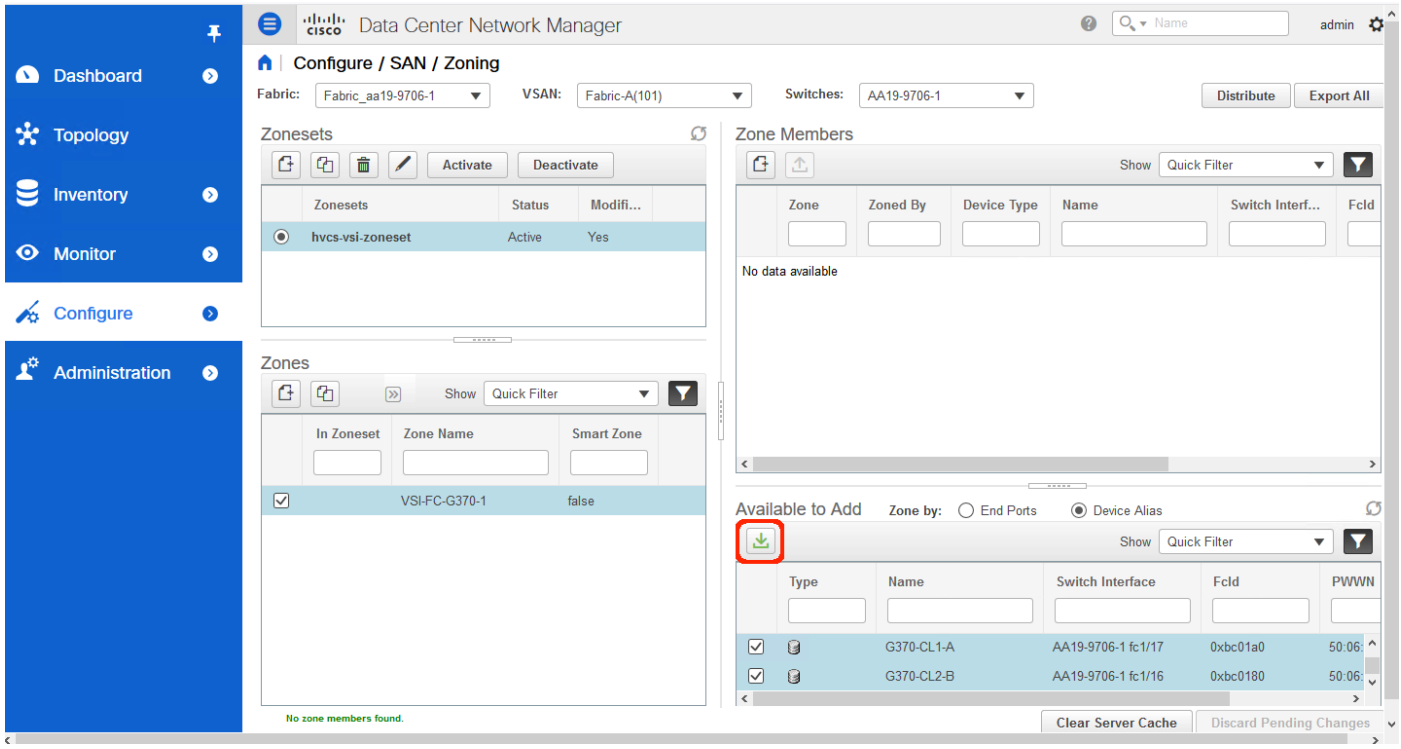

8. Click the green Add Member button within Available to Add.

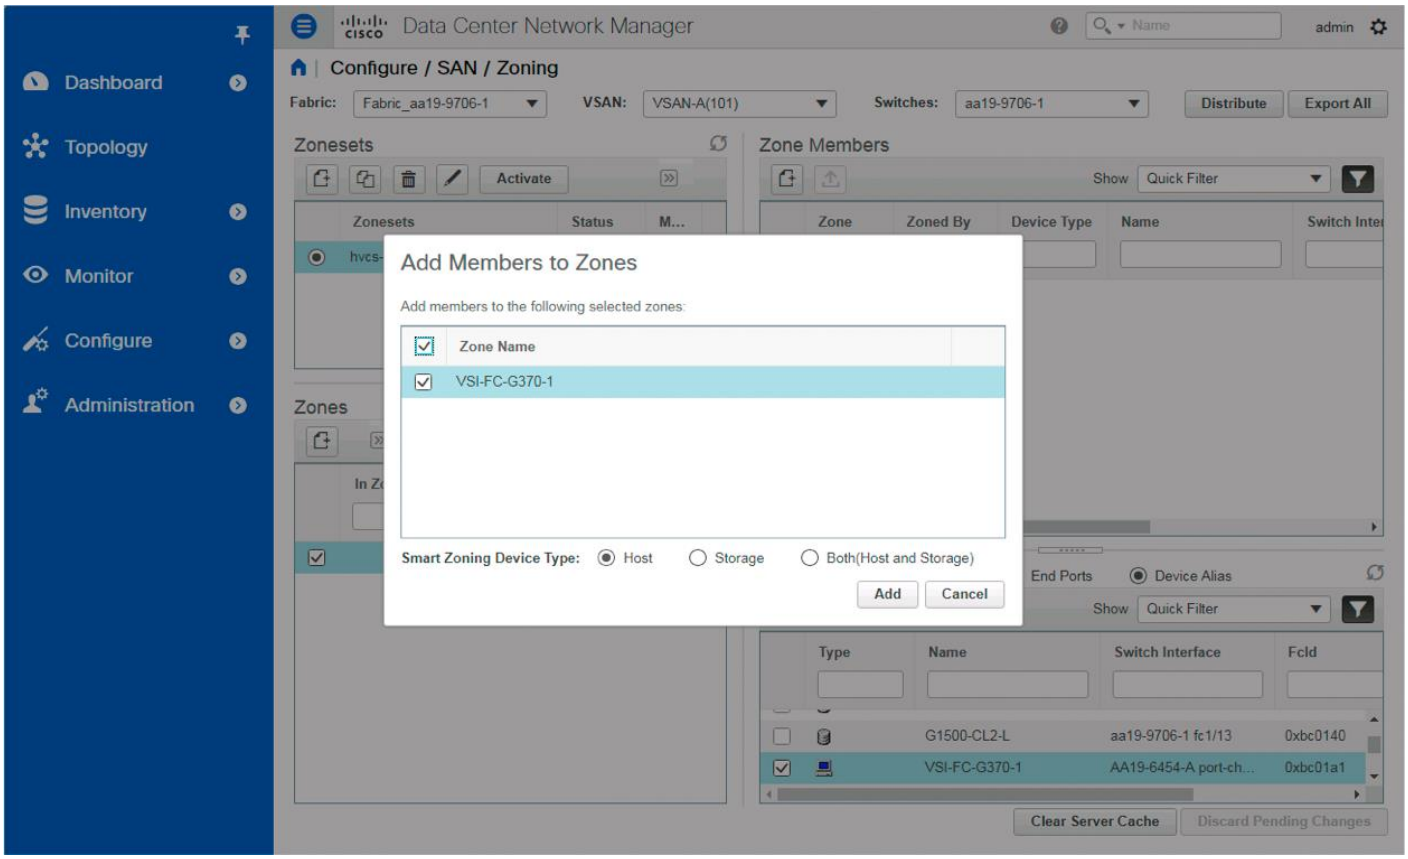

9. Make sure that Host is selected for the host being added and click Add.

10. With the zone still selected, and pick the appropriate VSP device aliases from the bottom right with the Available to Add section:

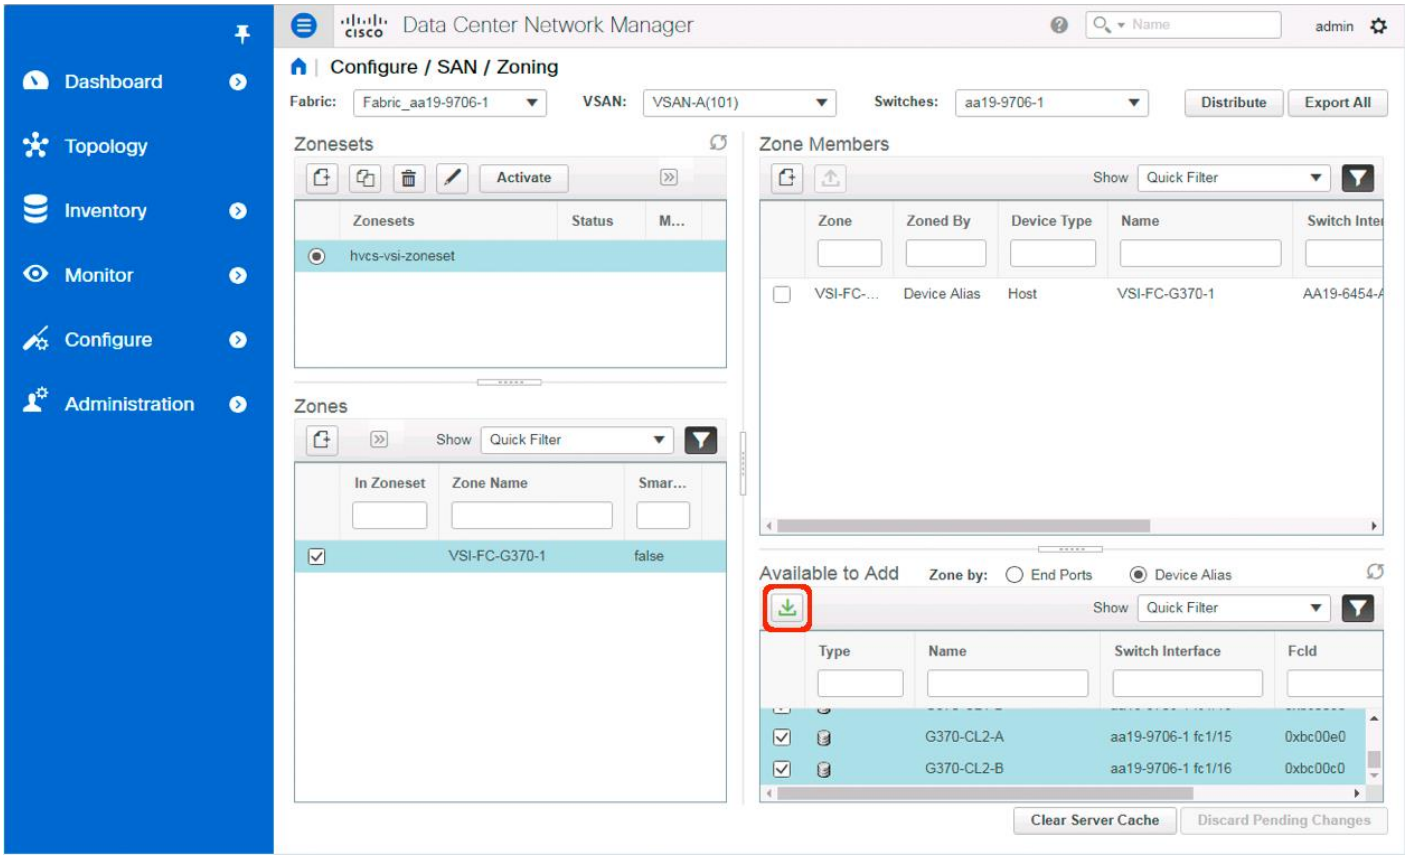

11. Make sure that Storage is selected for the VSP devices being added and click Add.

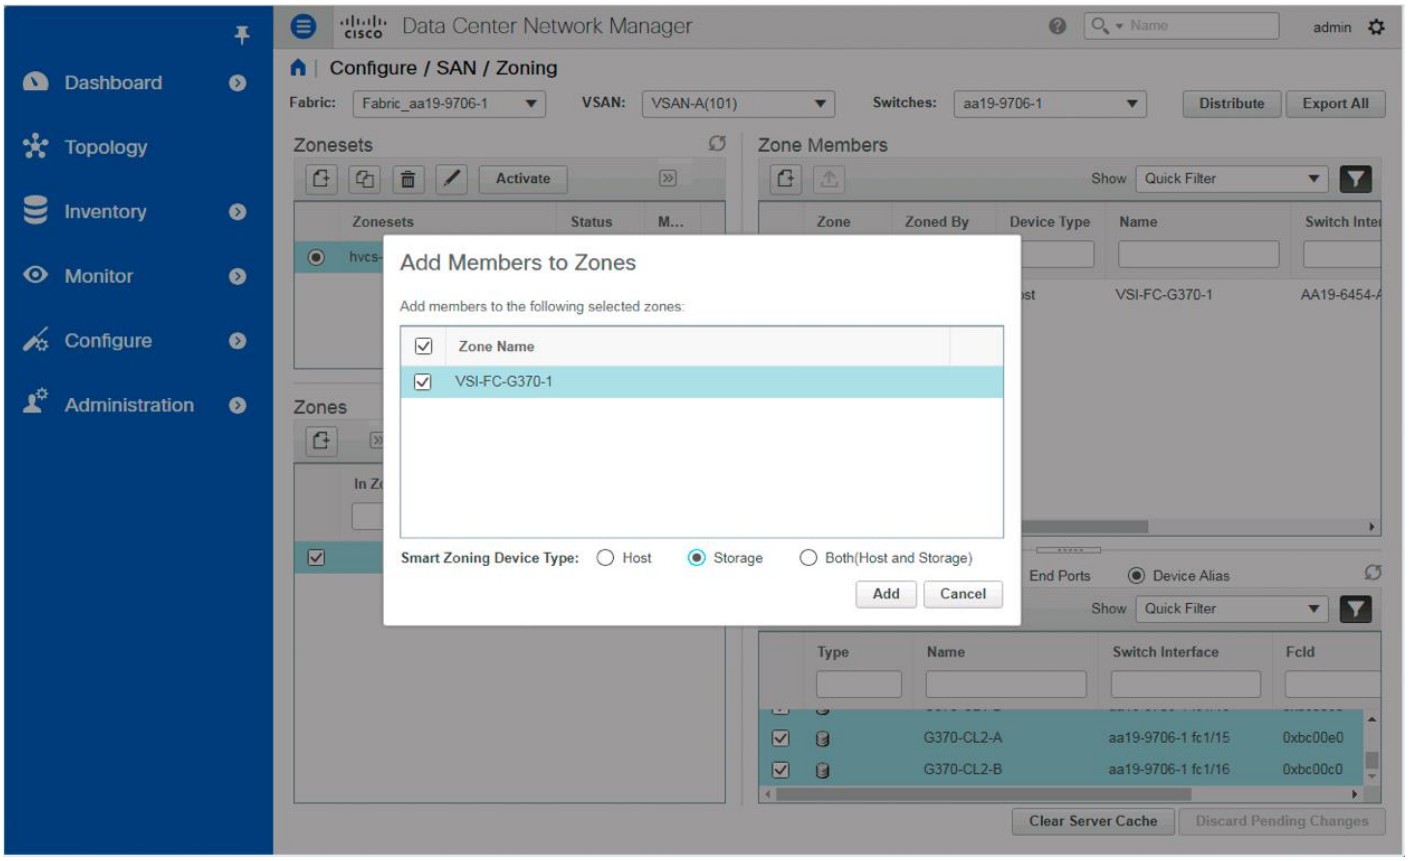

12. With the new zone still selected, click the drop-down list within Zones and select Enable Smart Zoning.

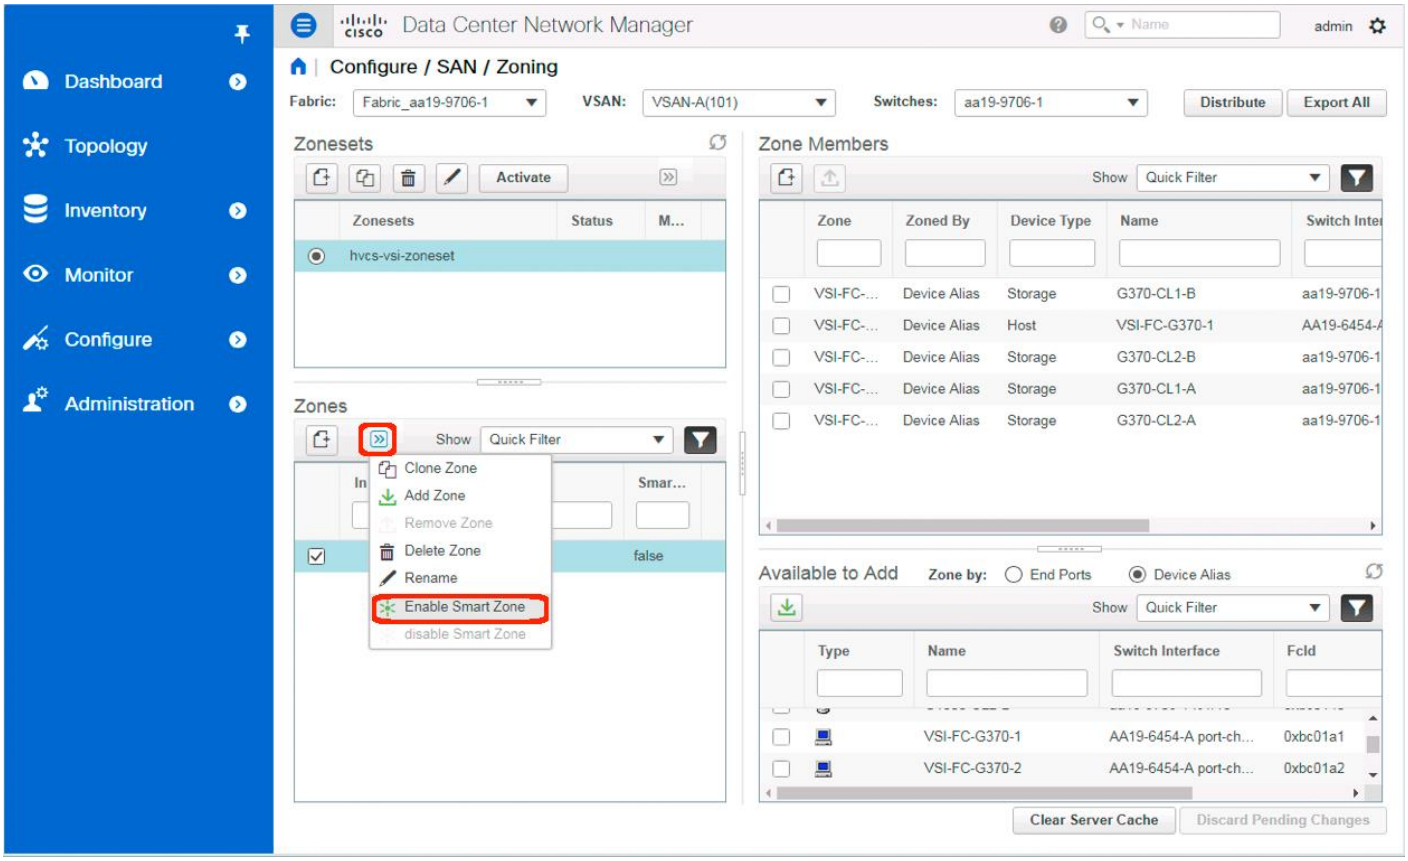

13. Additional zones for hosts associated to the same VSP can be created in the same manner or by selecting the first zone created and selecting Clone Zone from the drop-down options.

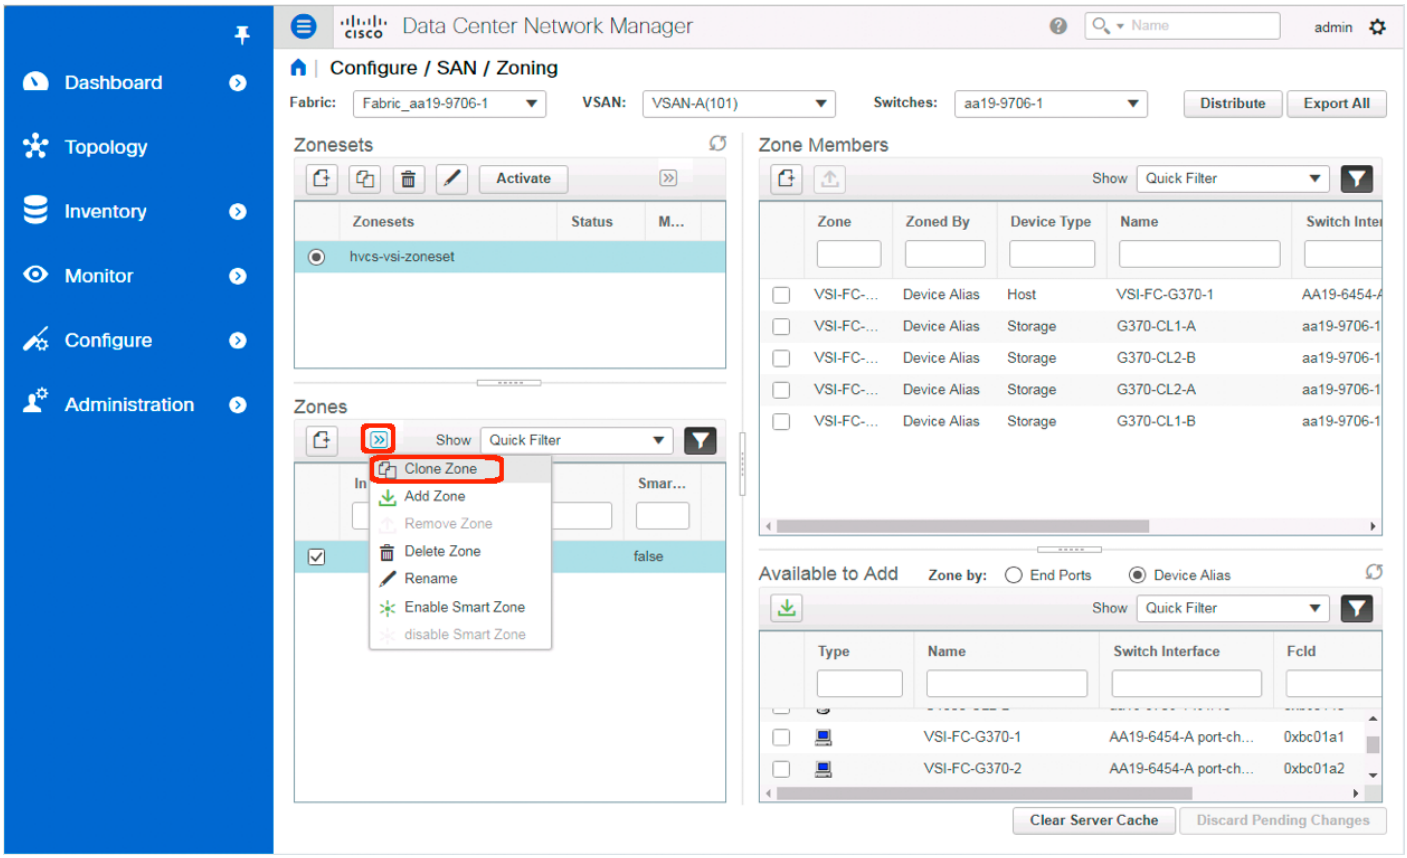

14. Specify the new host to be associated with the cloned zone:

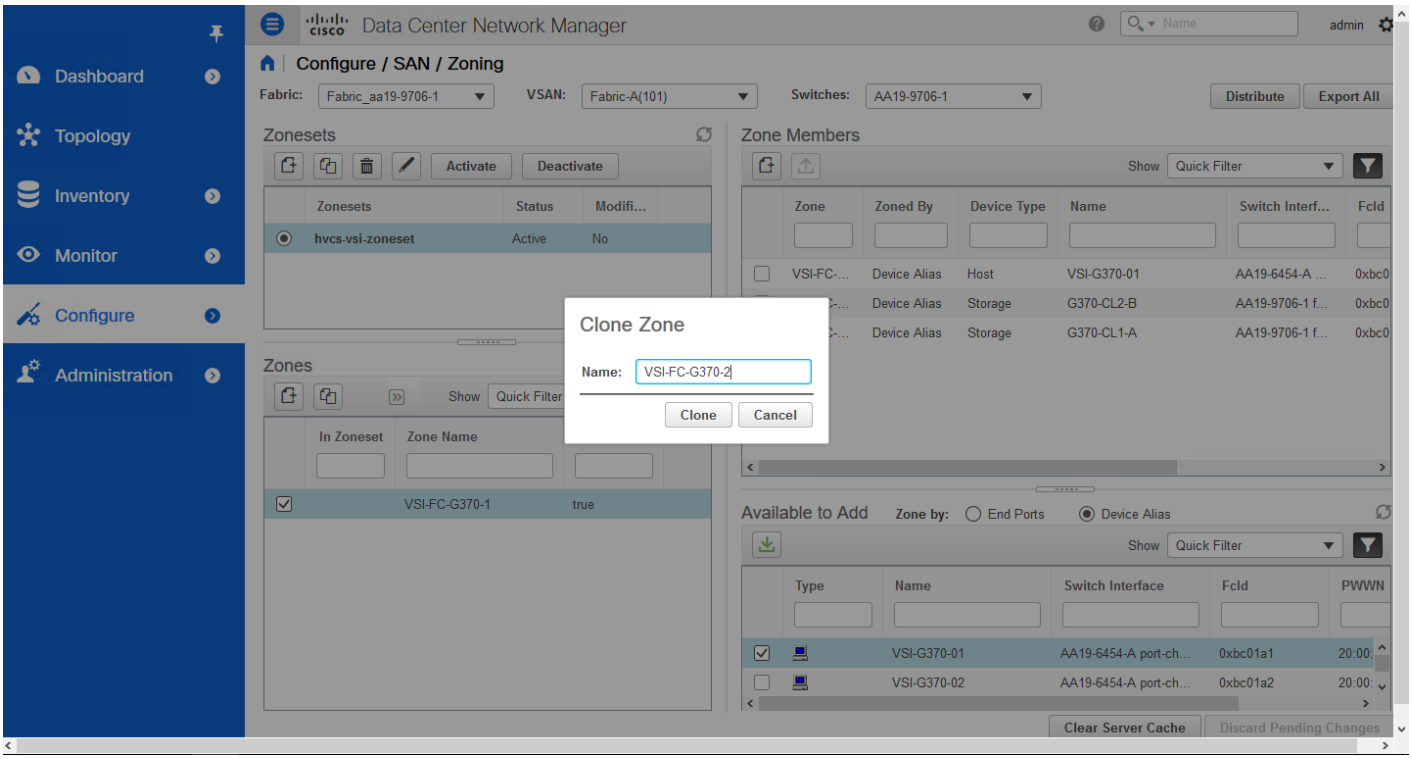

15. Click Clone.

16. Un-select the original zone and select the cloned zone.

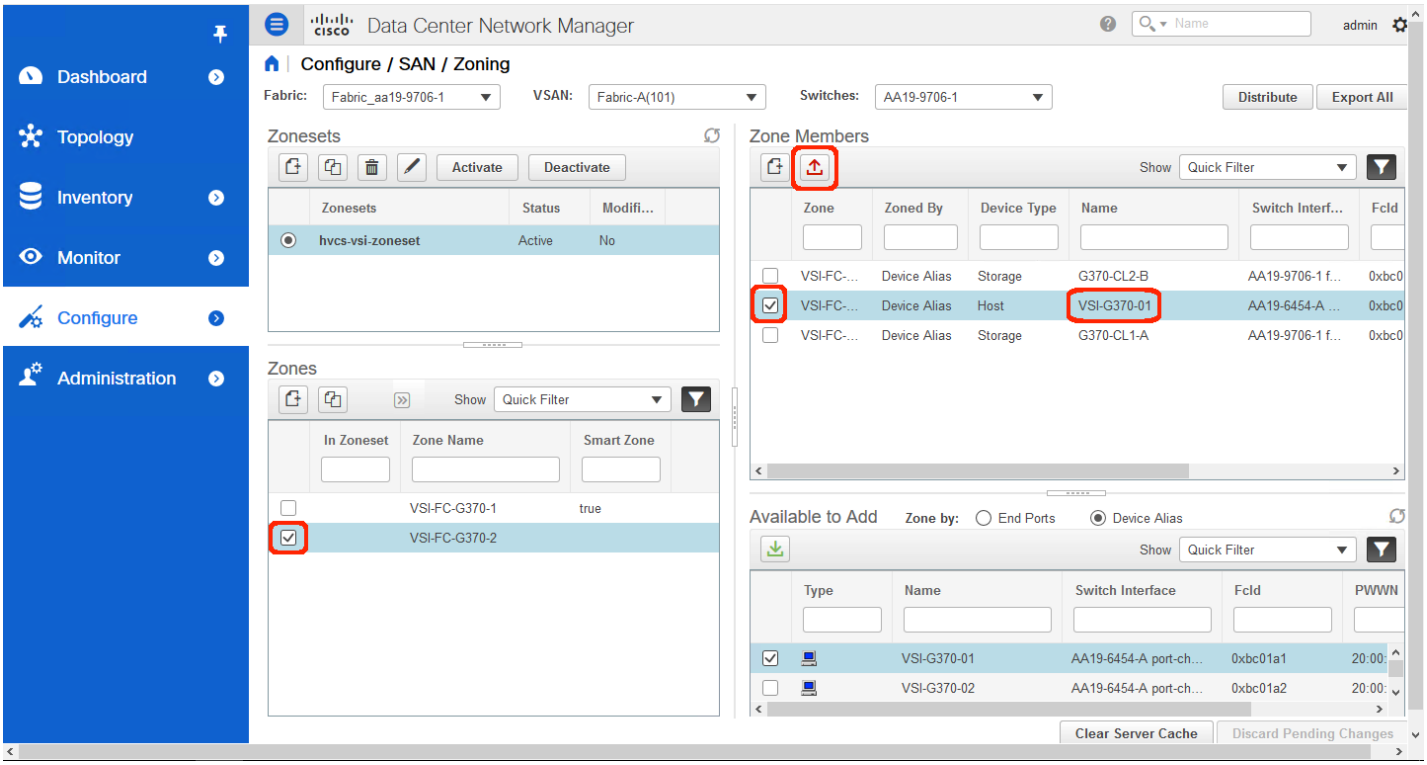

17. Select the host carried over from the cloning operation within the Zone Members section and click Remove Member.

18. Repeat the process of selecting the Enable Smart Zone for the new zone.

19. Re-Select the new zone.

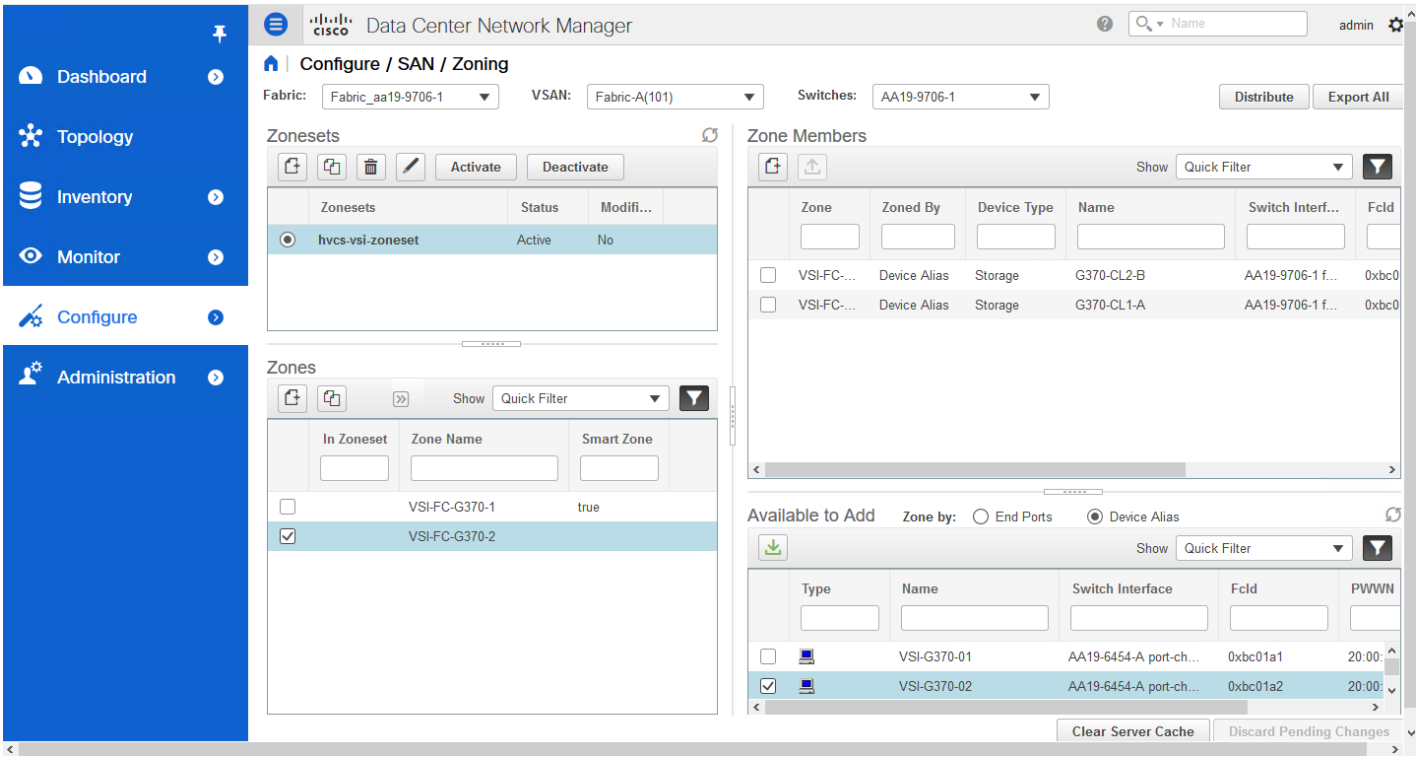

20. Find the host intended for this new zone within the Available to Add section.

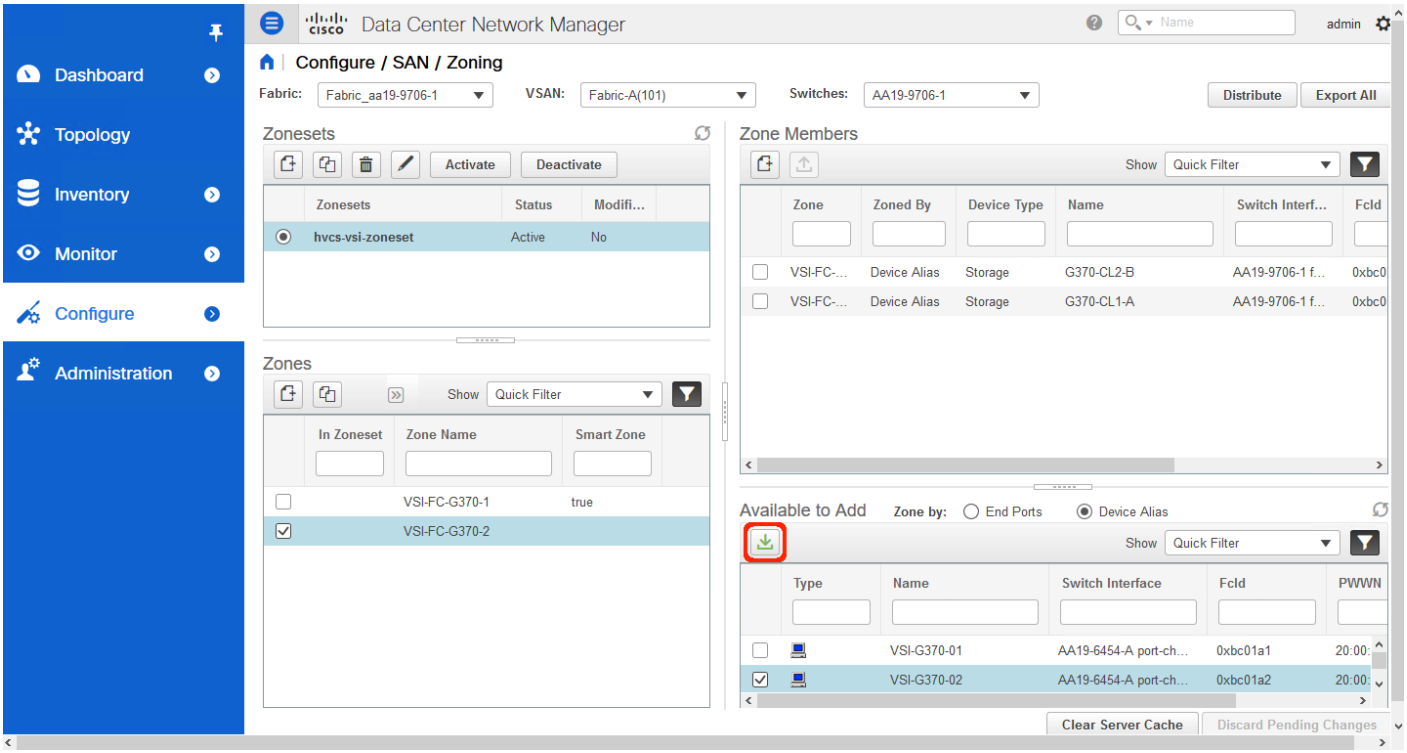

21. Click Add Member.

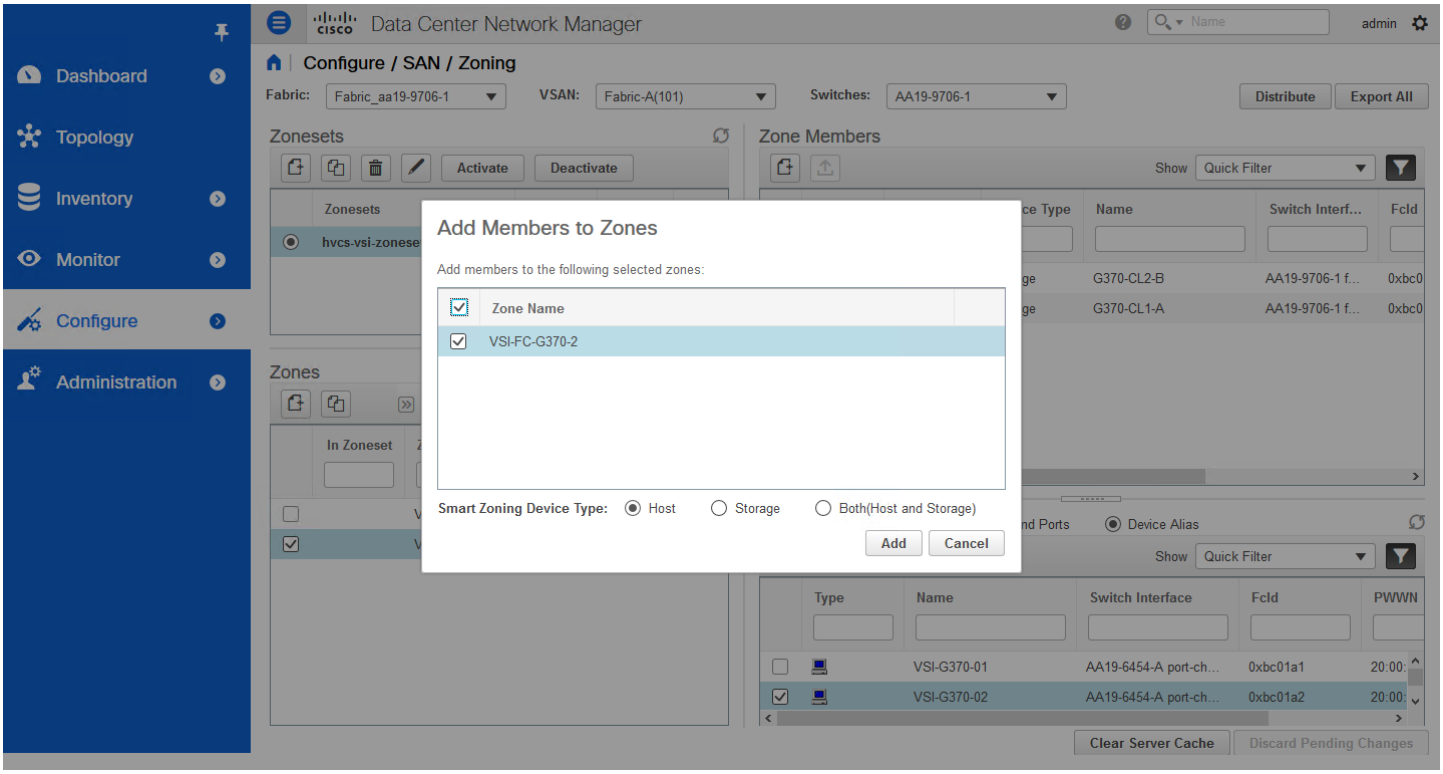

22. Ensure that Host is selected for Smart Zoning Device Type and click Add.

23. Repeat these steps to add zones for all additional hosts.

24. Select all created Zones, and find Add Zone from the drop-down list to add to the zoneset.

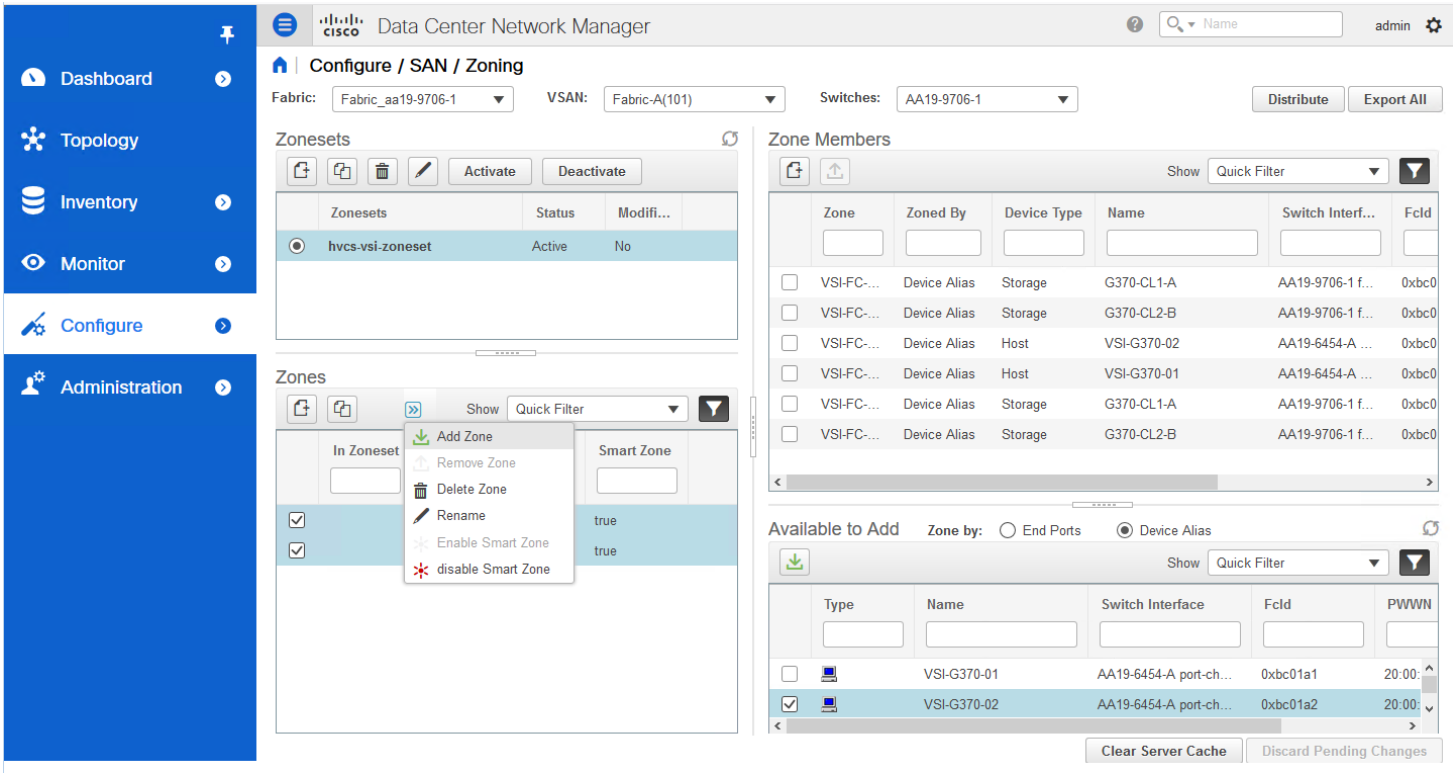

25. Zones will now show checkmarks as In Zoneset.

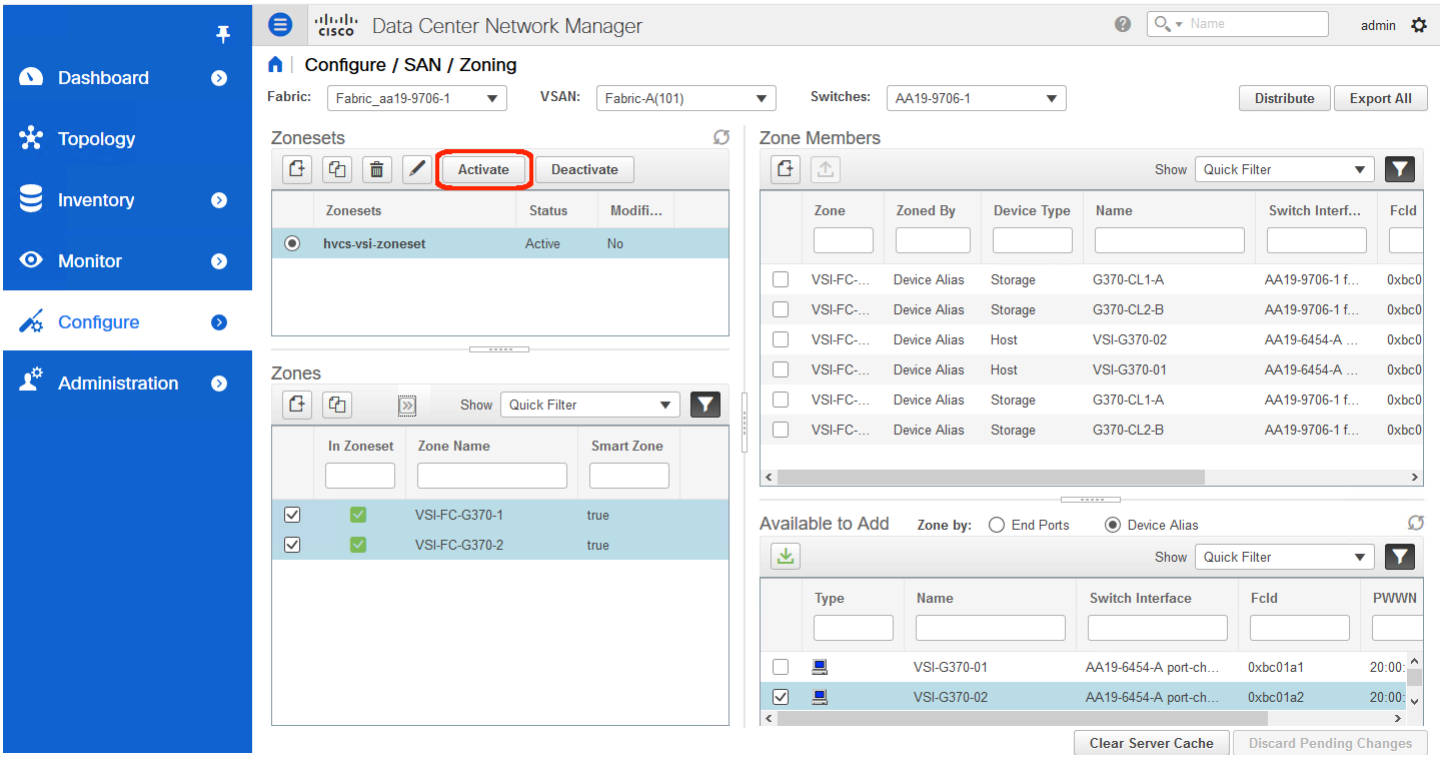

26. Click Activate to activate the zoneset.

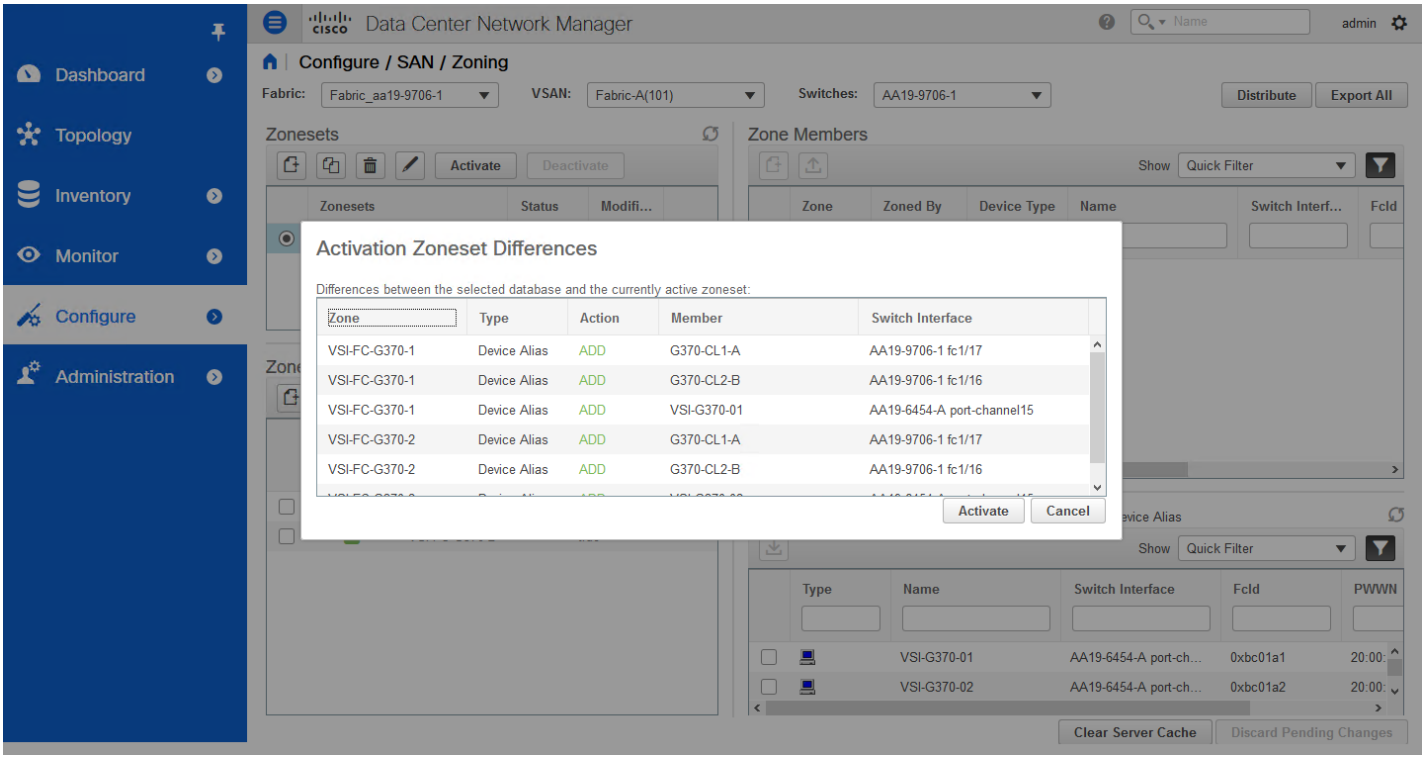

27. Click Activate.

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 the parity groups configured on it.

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.

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.

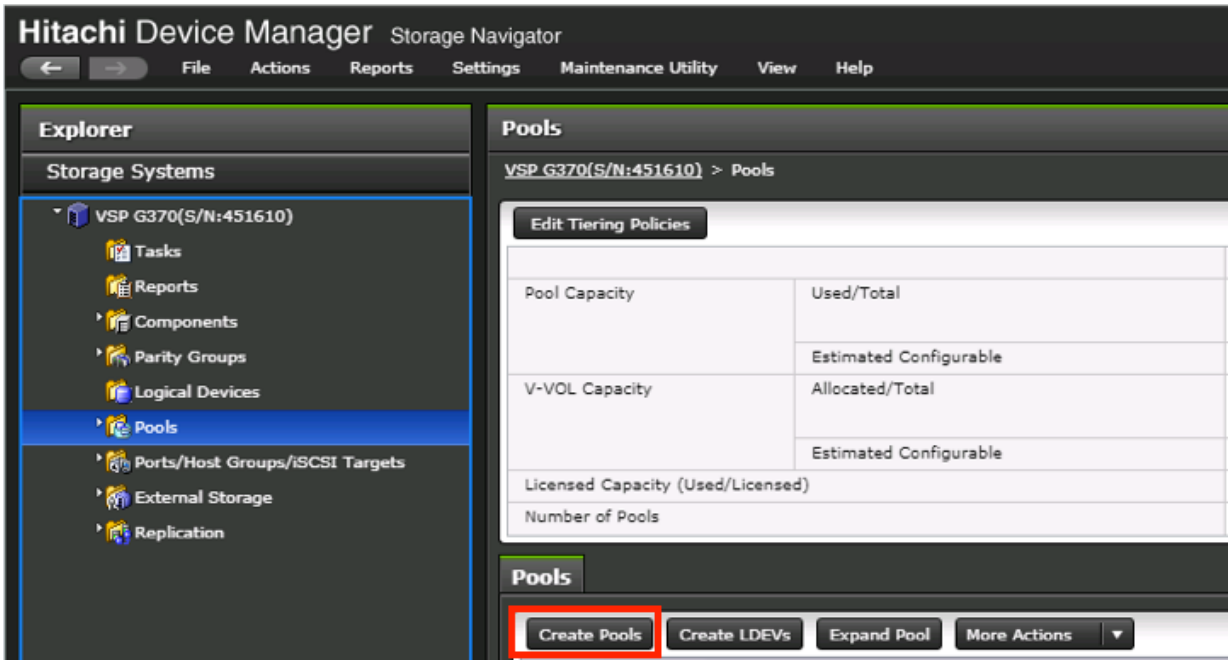

4. Configure the following items in the left pane of the Create Pools dialog box:

   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.

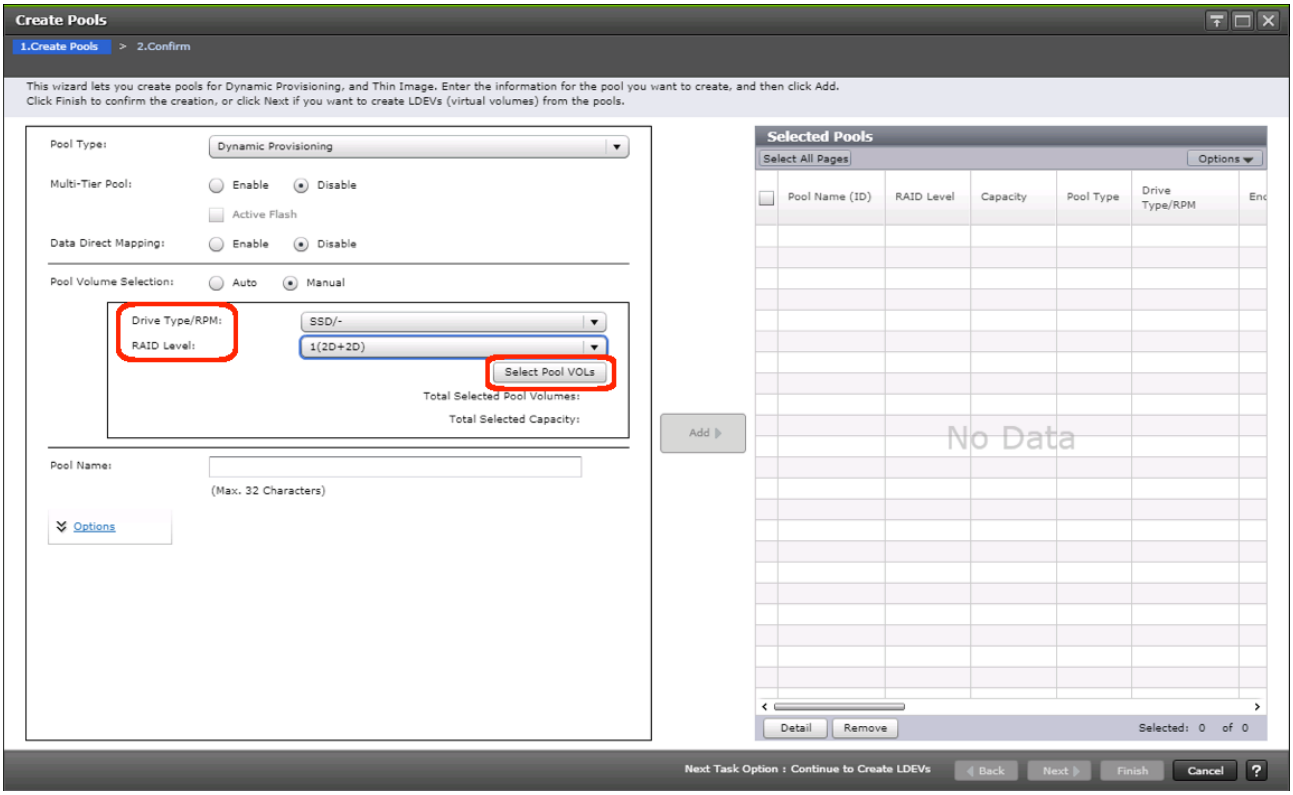

6. Within the left pane of the Select Pool VOLs dialog box, select the checkbox next to the LDEVs to be used for the UCS server boot LDEV dynamic provisioning pool.

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.

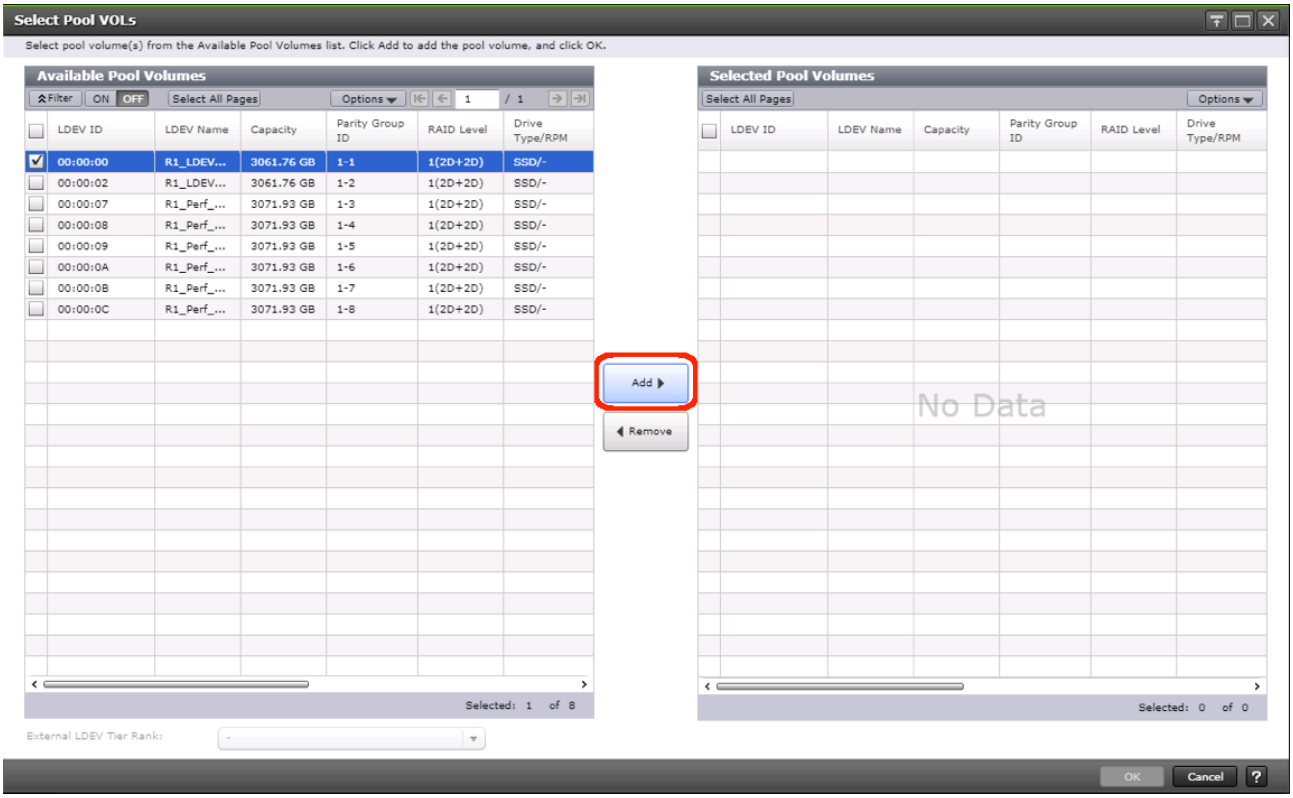

8. You should now see values for Total Selected Pool Volumes and Total Selected Capacity shown under the Select Pool VOLs button. Give the dynamic provisioning pool a descriptive Pool Name, then click Add to add the pool to be created to the Selected Pools pane in the dialog.

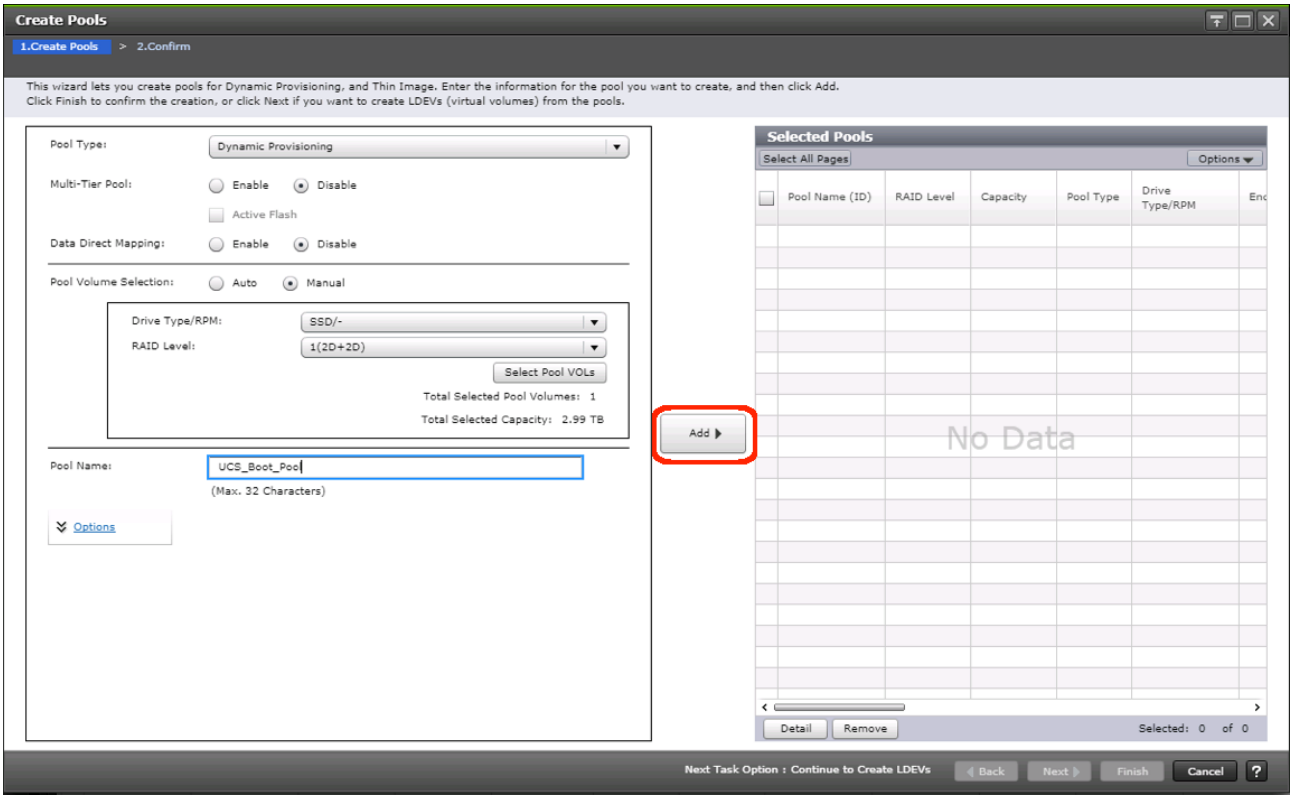

9. Click Finish.

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.

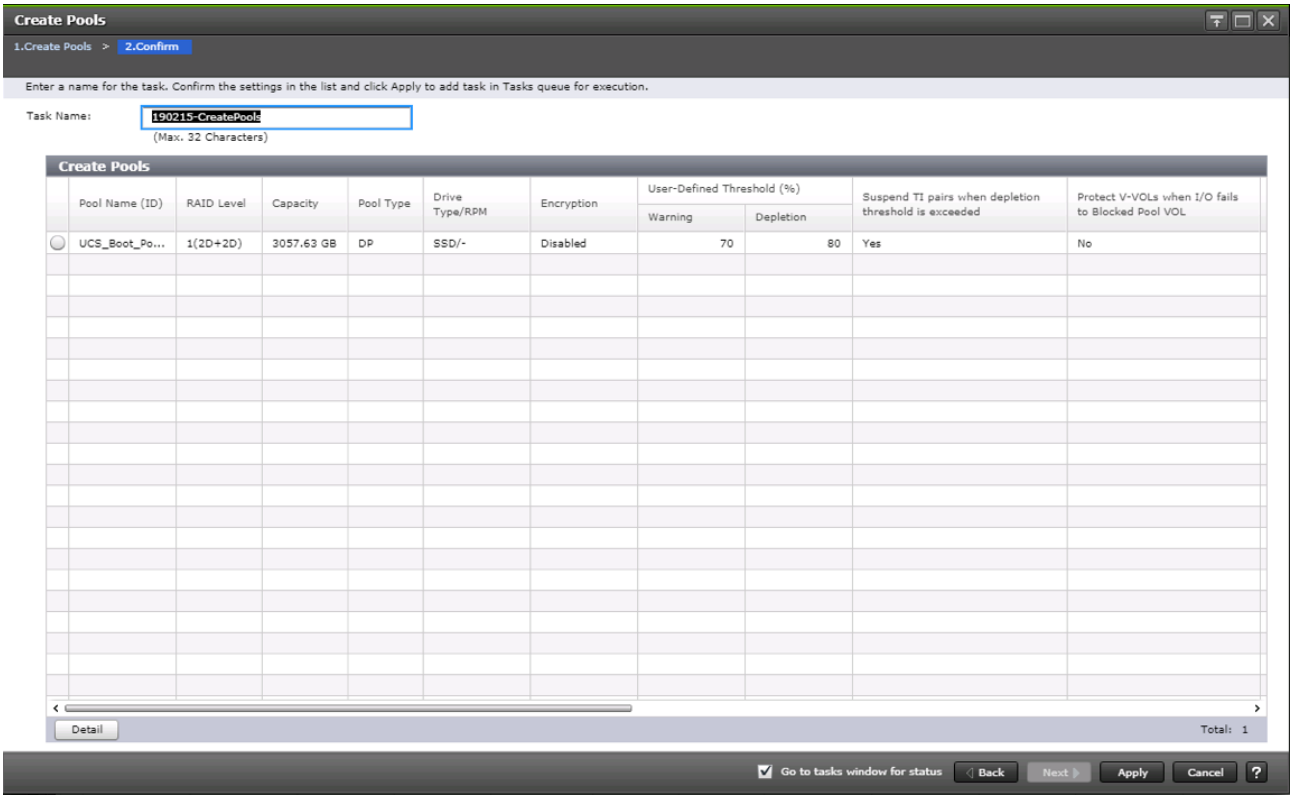

Create a Hitachi Dynamic Provisioning Pool for UCS Server VMFS Volume LDEVs

Follow the same steps as 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 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 left 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.

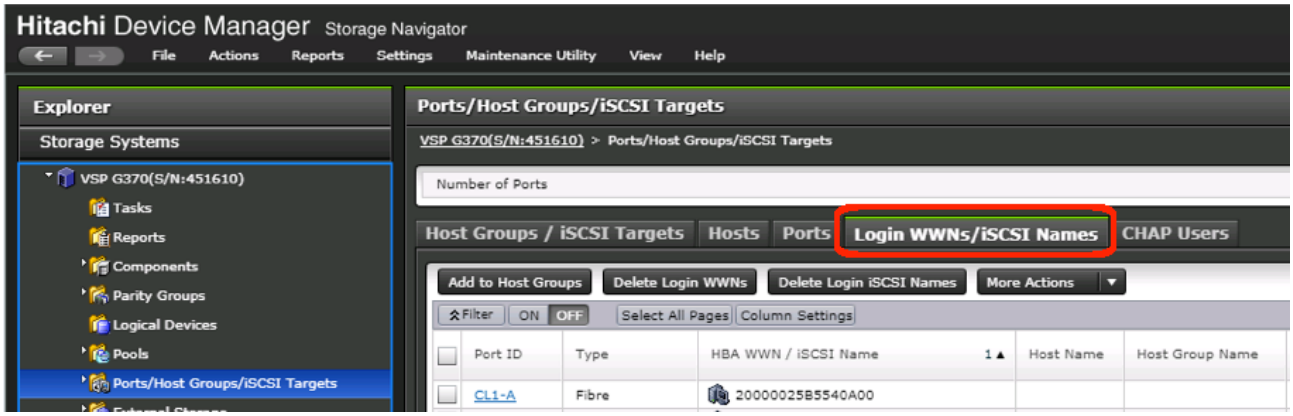

3. Review the list of WWNs and associated ports. You should be able to see each vHBA assigned to each fabric associated with each port on the VSP that it is zoned to.

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 MDS.

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

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

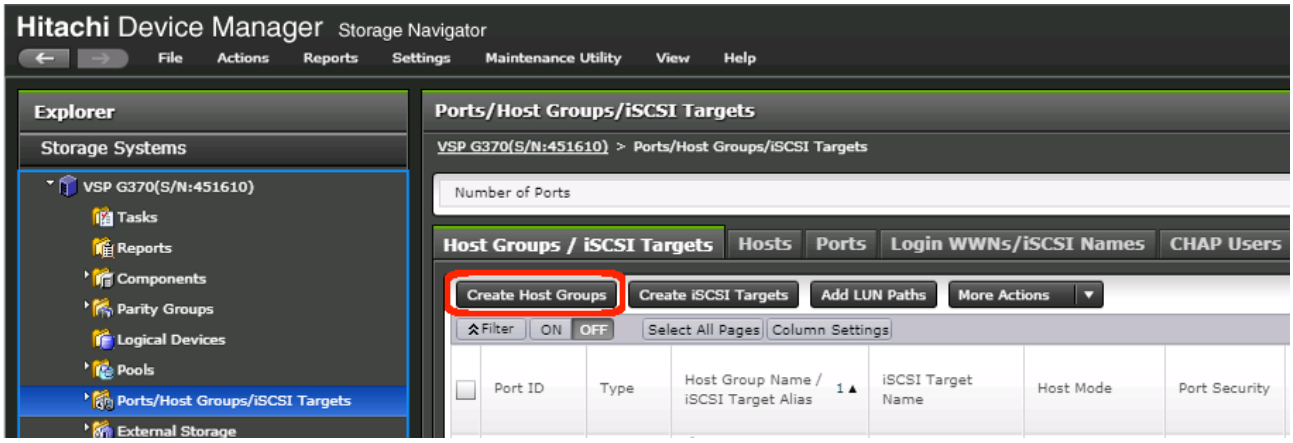

7. Host groups will be created separately for fabric A and fabric B vHBAs. Start with the fabric A host group for an individual UCS Service Profile and modify the following within the Create Host Groups dialog box:

   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

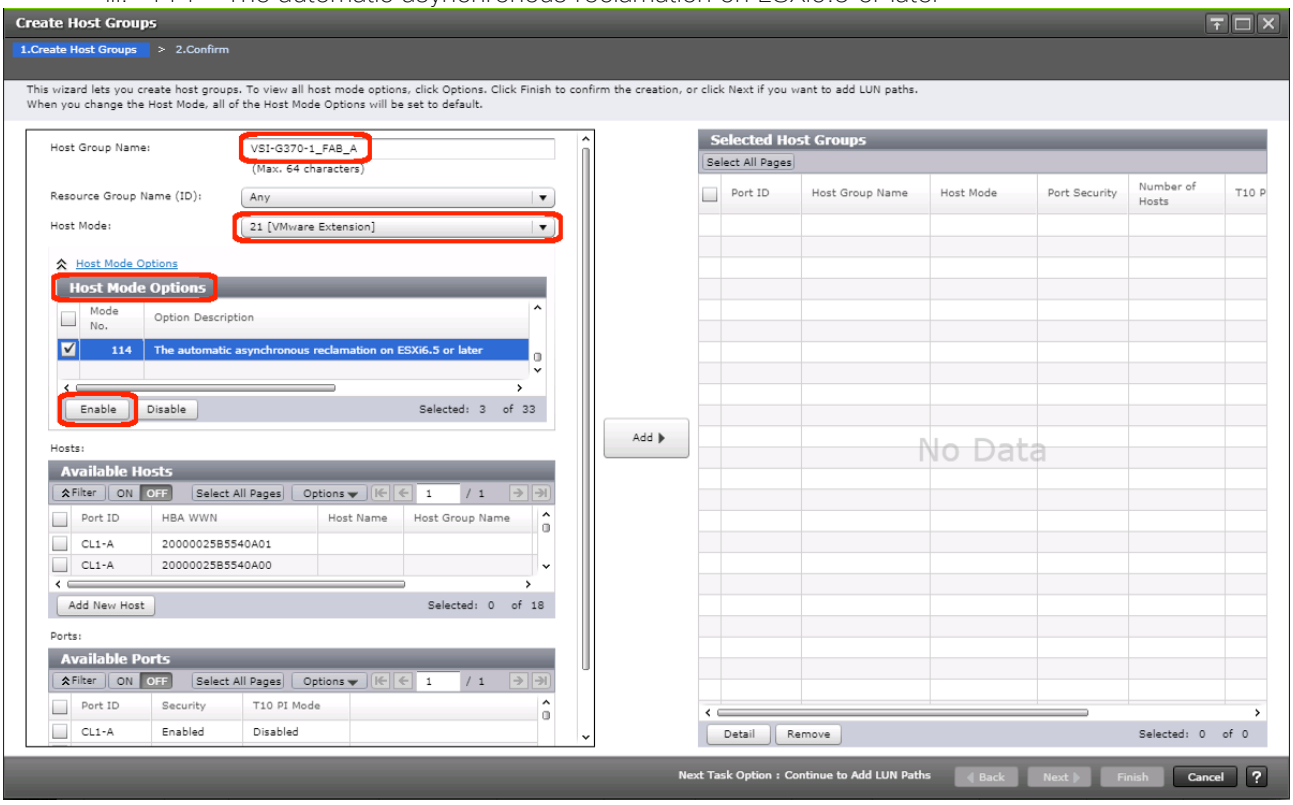

8. Bring down the WWN information from Table 28 and Table 29 in the previous Create Device Aliases section:

Table 30 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 31 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 left pane of 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.

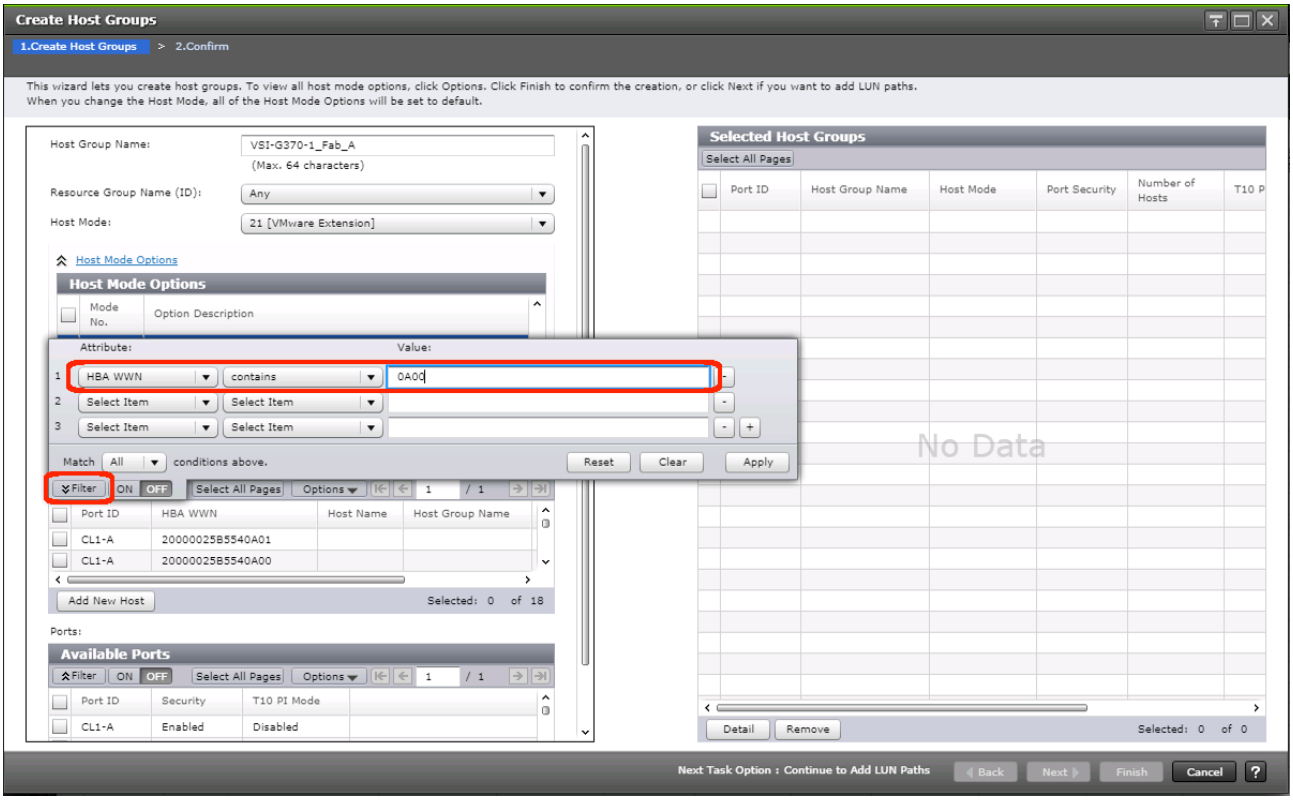

12. Click Apply.

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.

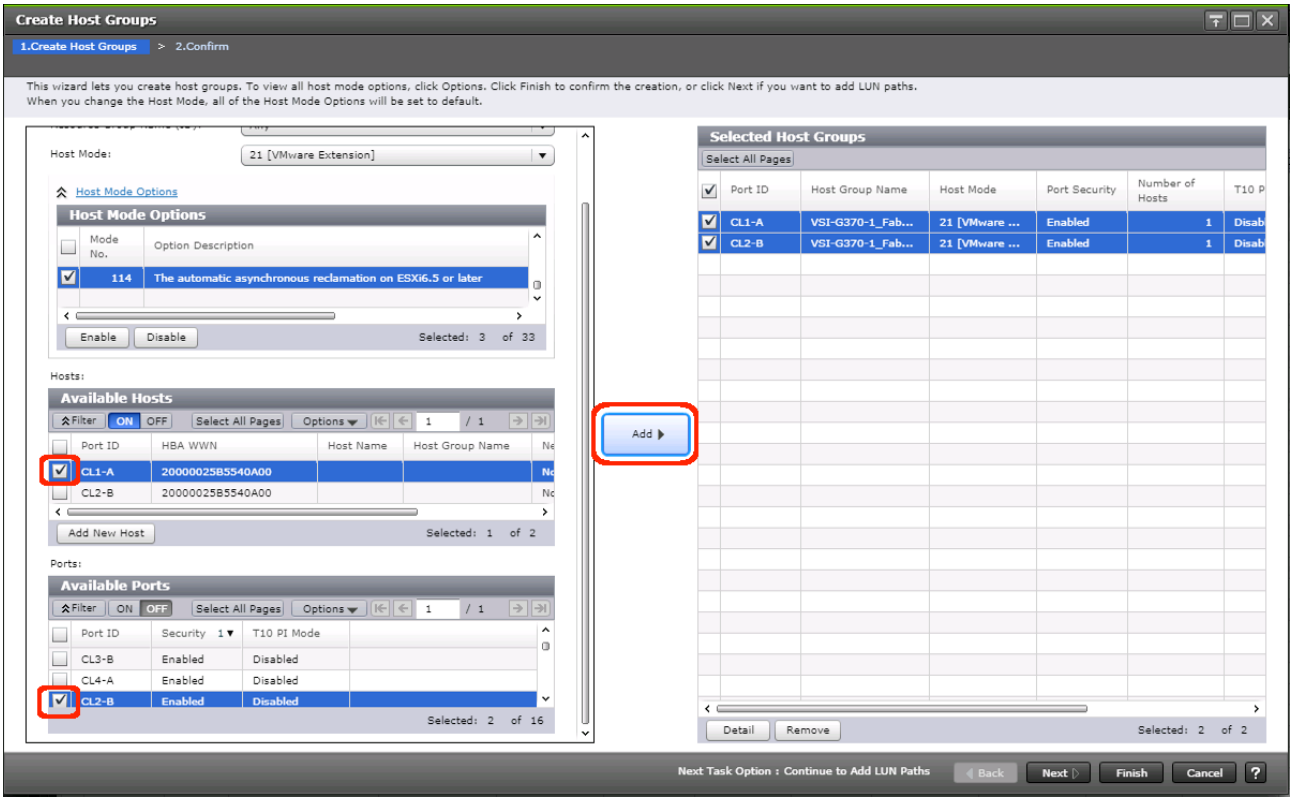

In the picture above, the CL1-A port entry was also selected within the **Available Ports** section.

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.

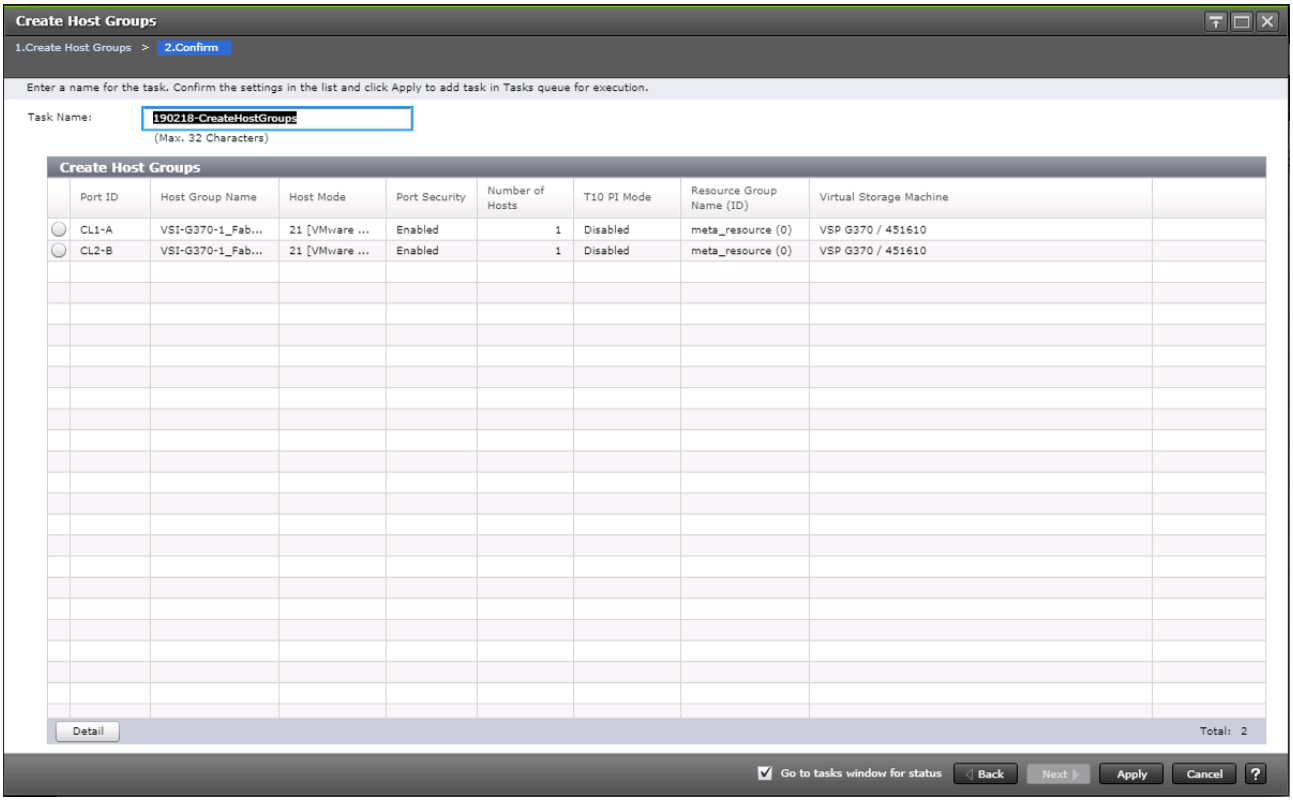

18. Click Apply.

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 in the right hand pane, and click Create LDEVs to instantiate the Create LDEVs dialog.

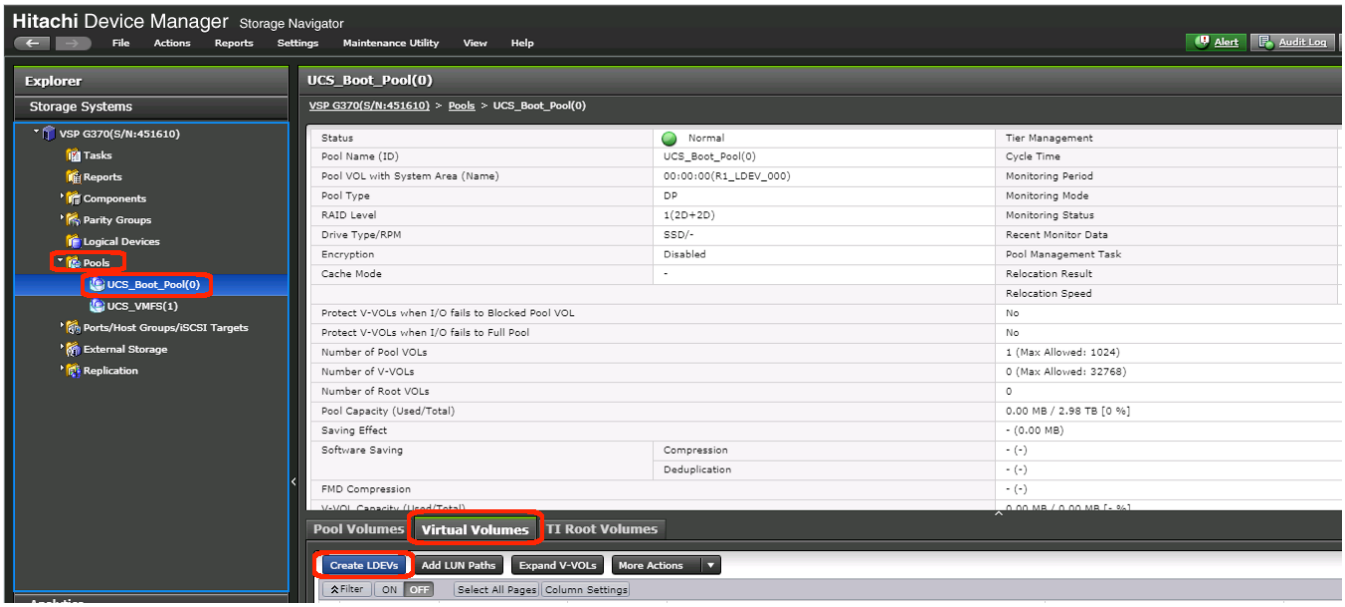

4. Modify the following within the Create LDEVs dialog:

   – 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.

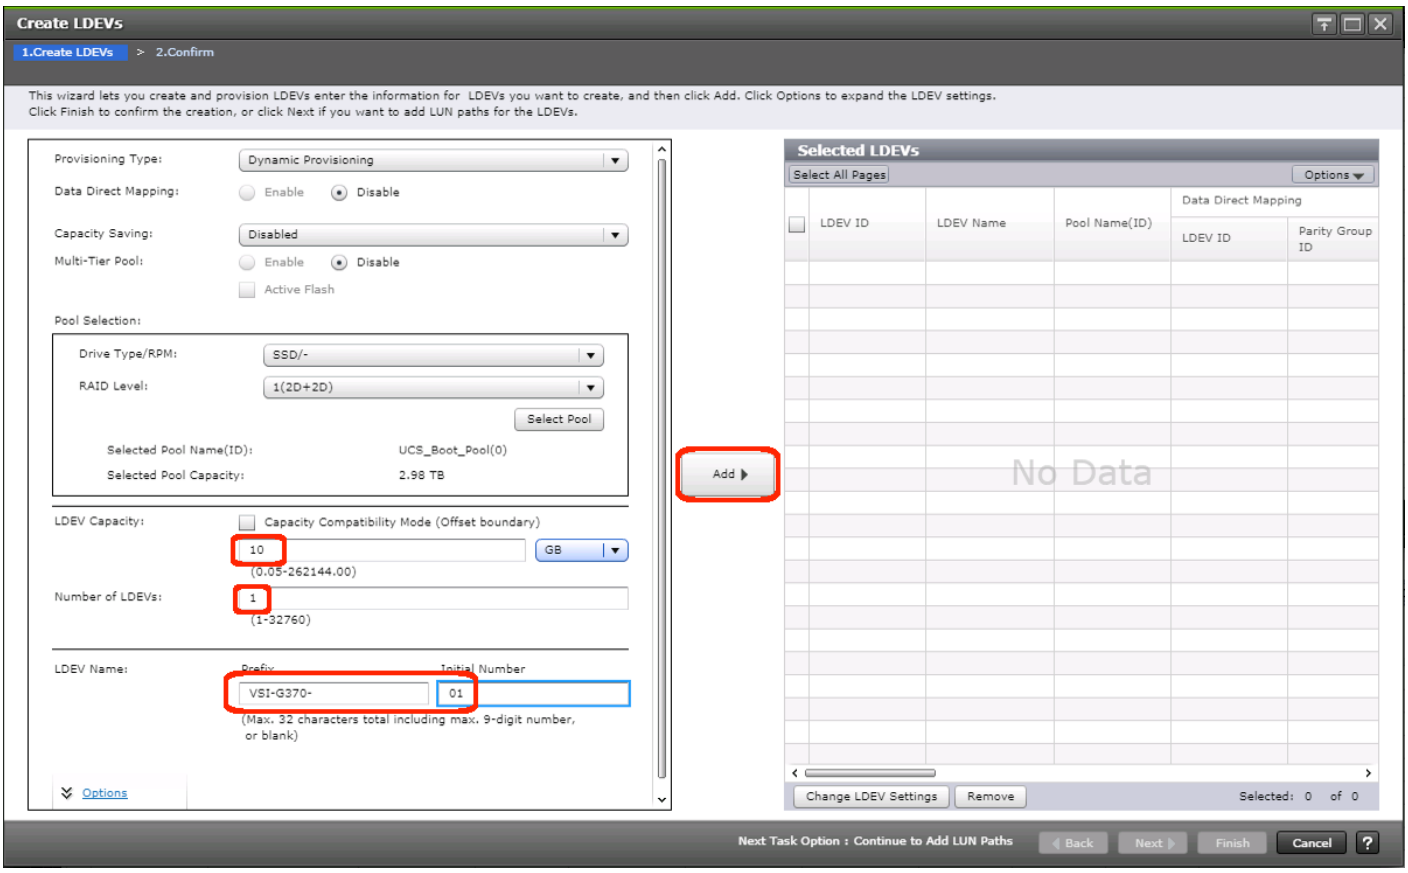

5. Click Add and verify that the boot LDEV is listed in the right-hand Selected LDEVs pane, then click Next.

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

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

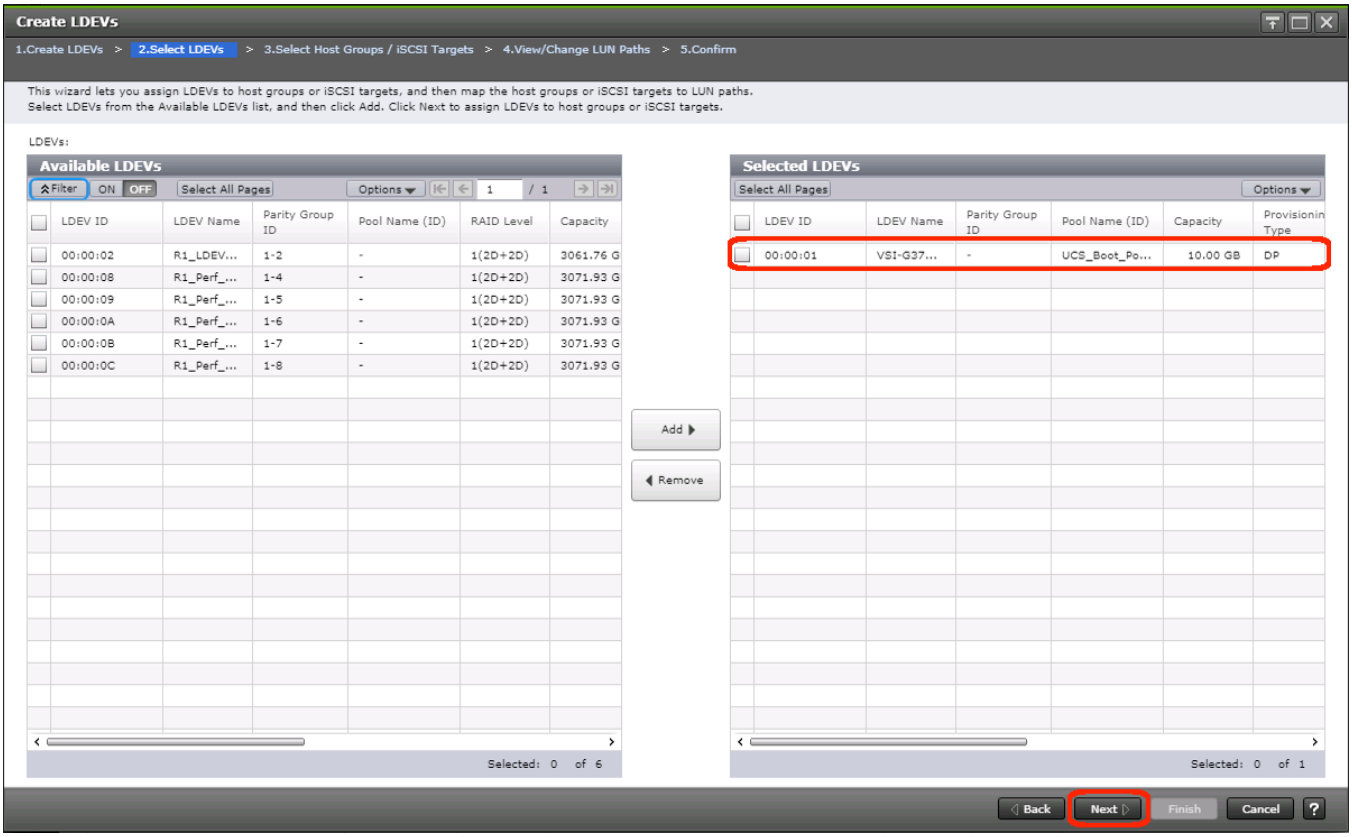

8. The Select Host Groups/iSCSI Targets screen shows all of the host groups that can be assigned to the boot LDEV as a path.

9. 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>

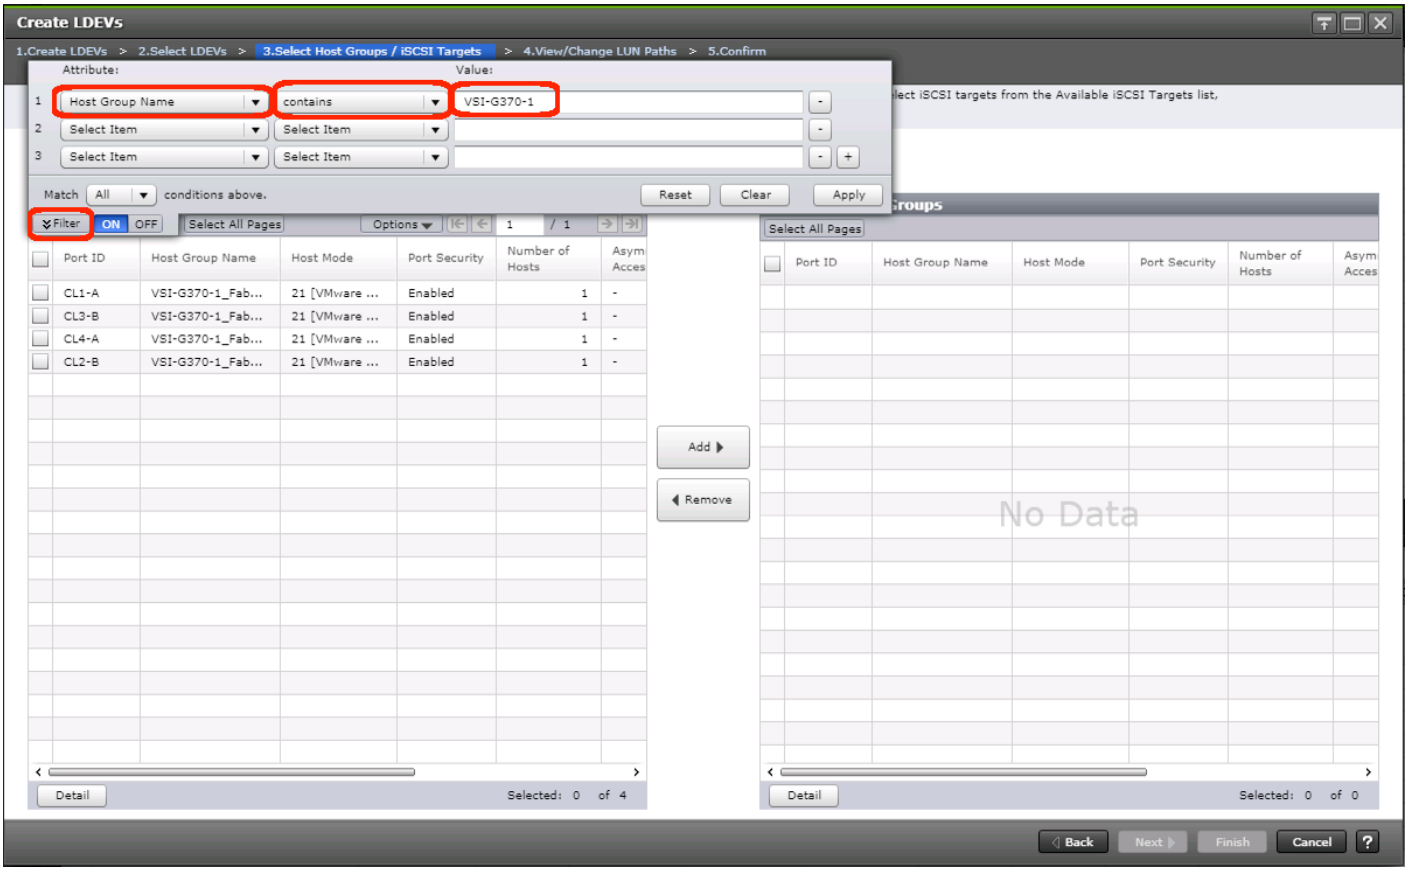

10. Click Apply.

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

12. 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.

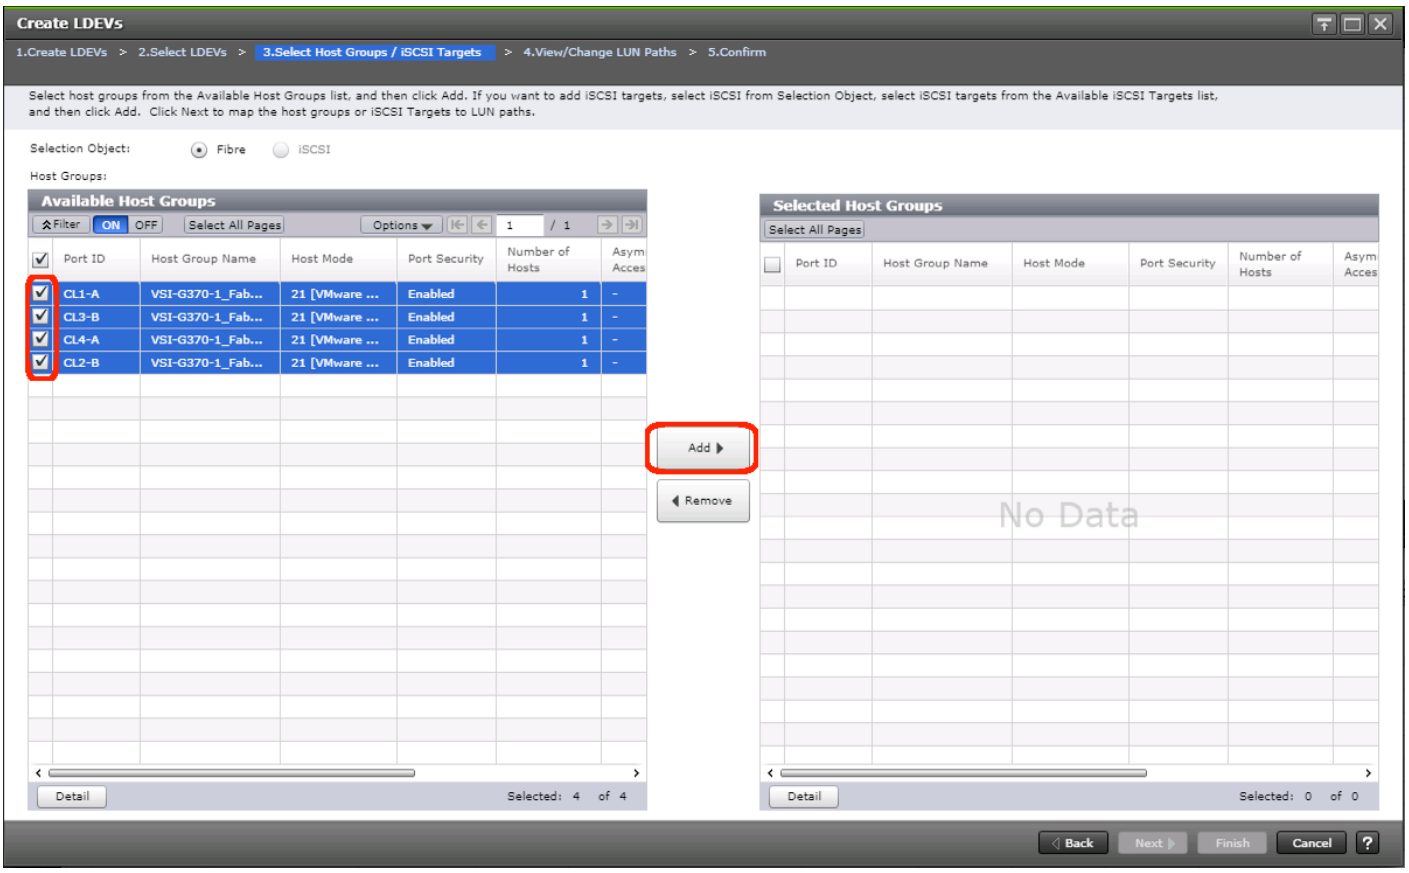

13. Click the Add to populate the Selected Host Groups pane with the selected host groups, then click Next.

14. 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.

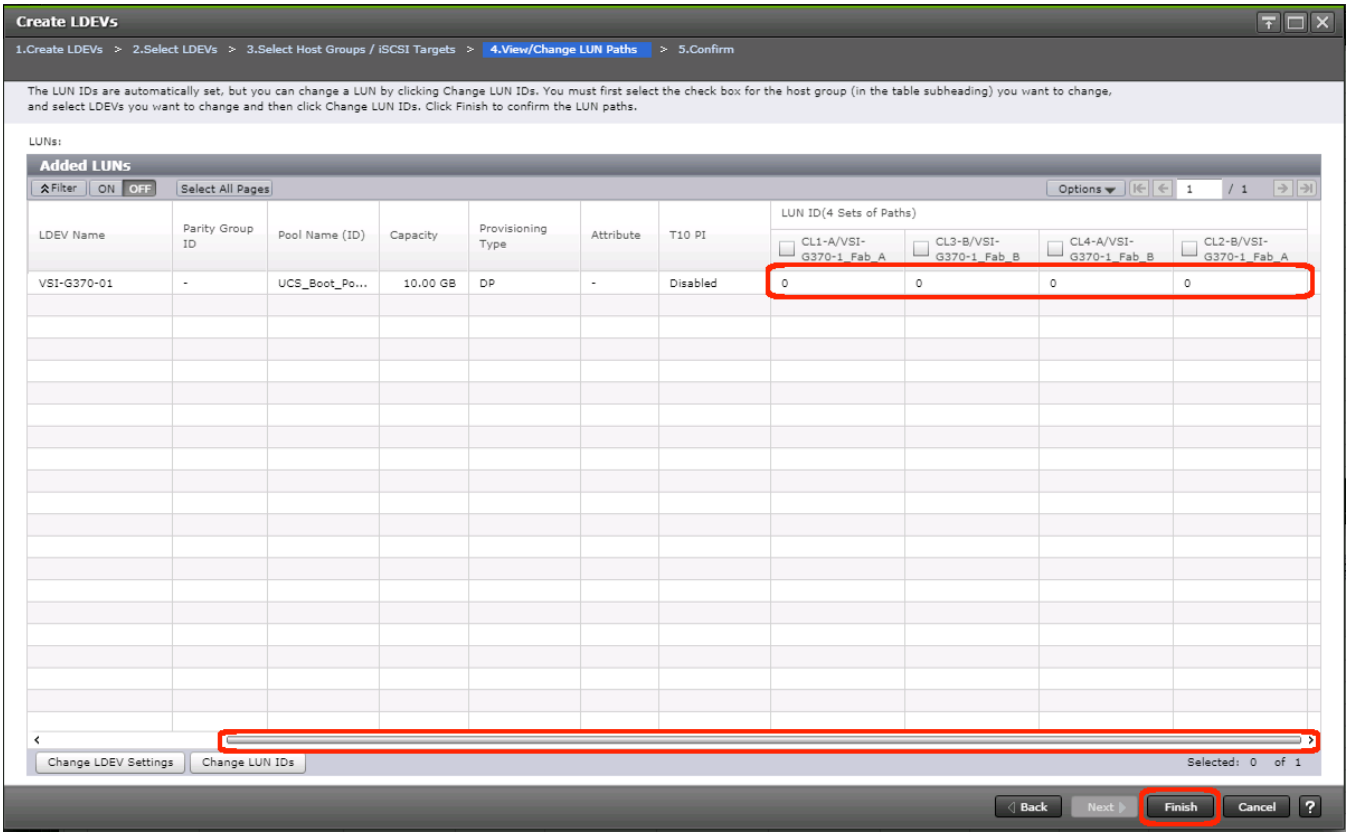

15. Use the scrollbar at the bottom of this screen to review the LUN ID assigned and ensure that all LUN IDs are set to zero, then click Finish.

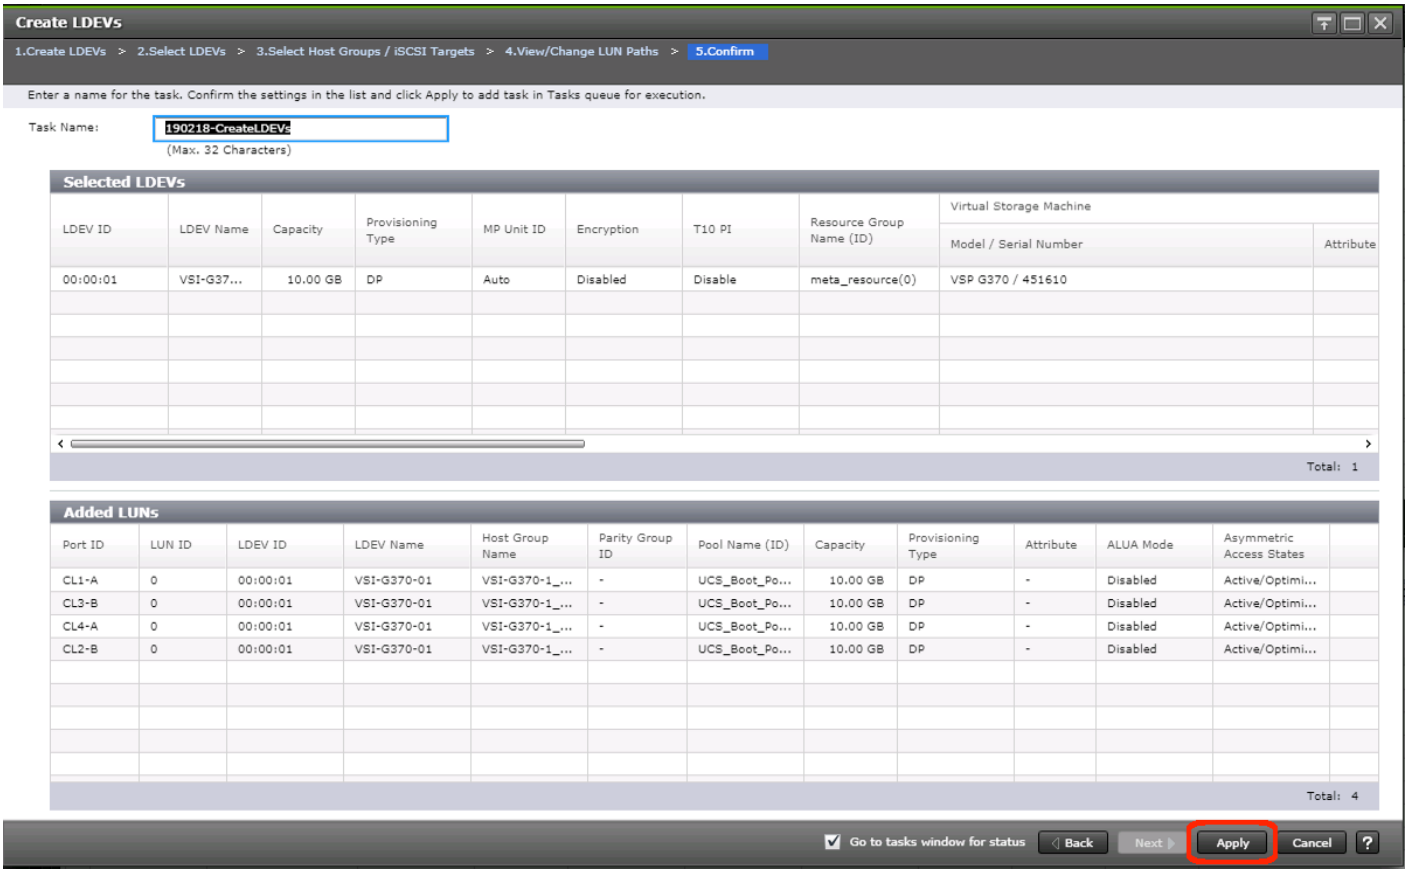

16. Review the LDEV details and LUN ID configuration of the boot LDEV being created, then click Apply to create the LDEV and add paths to the UCS Service Profile.

17. Repeat steps 1–16 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 in the right hand pane, 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.

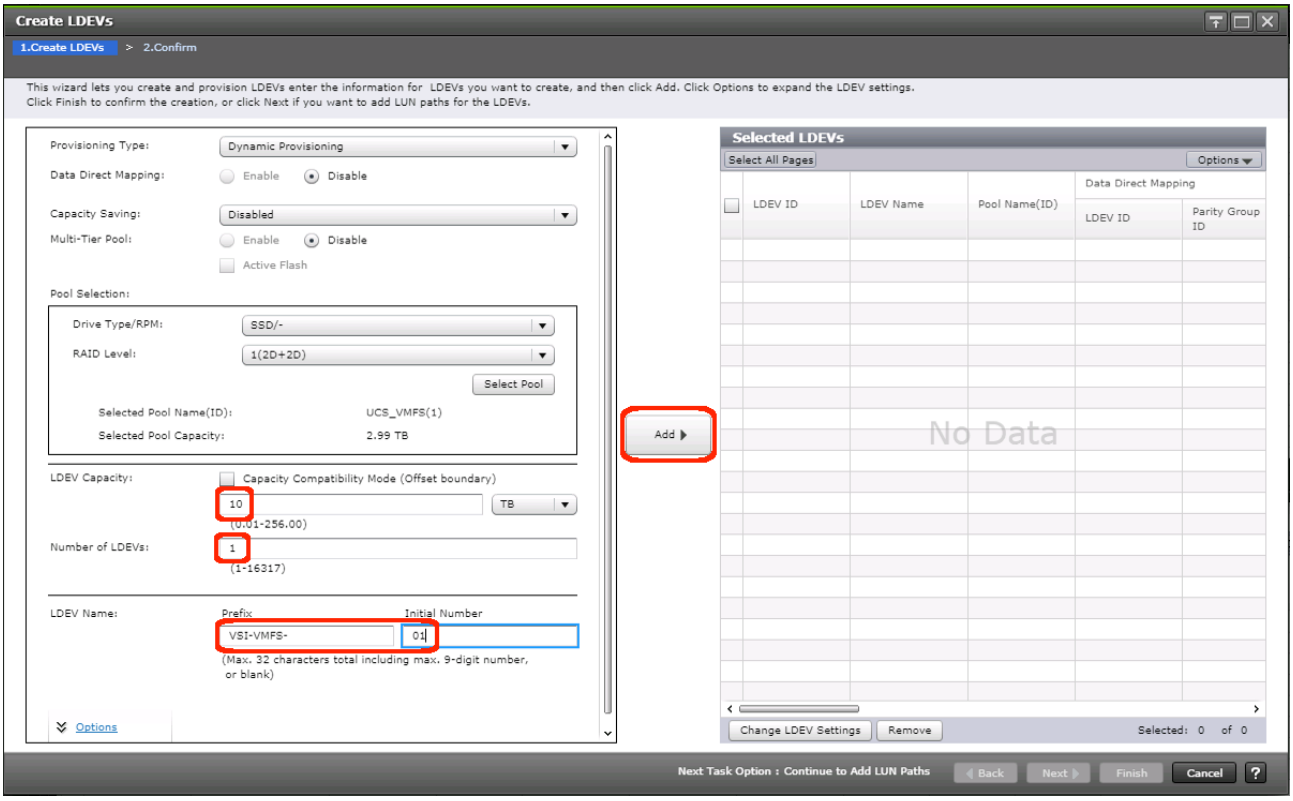

5. The Select LDEVs screen shows the selected LDEVs to which the paths will be added.

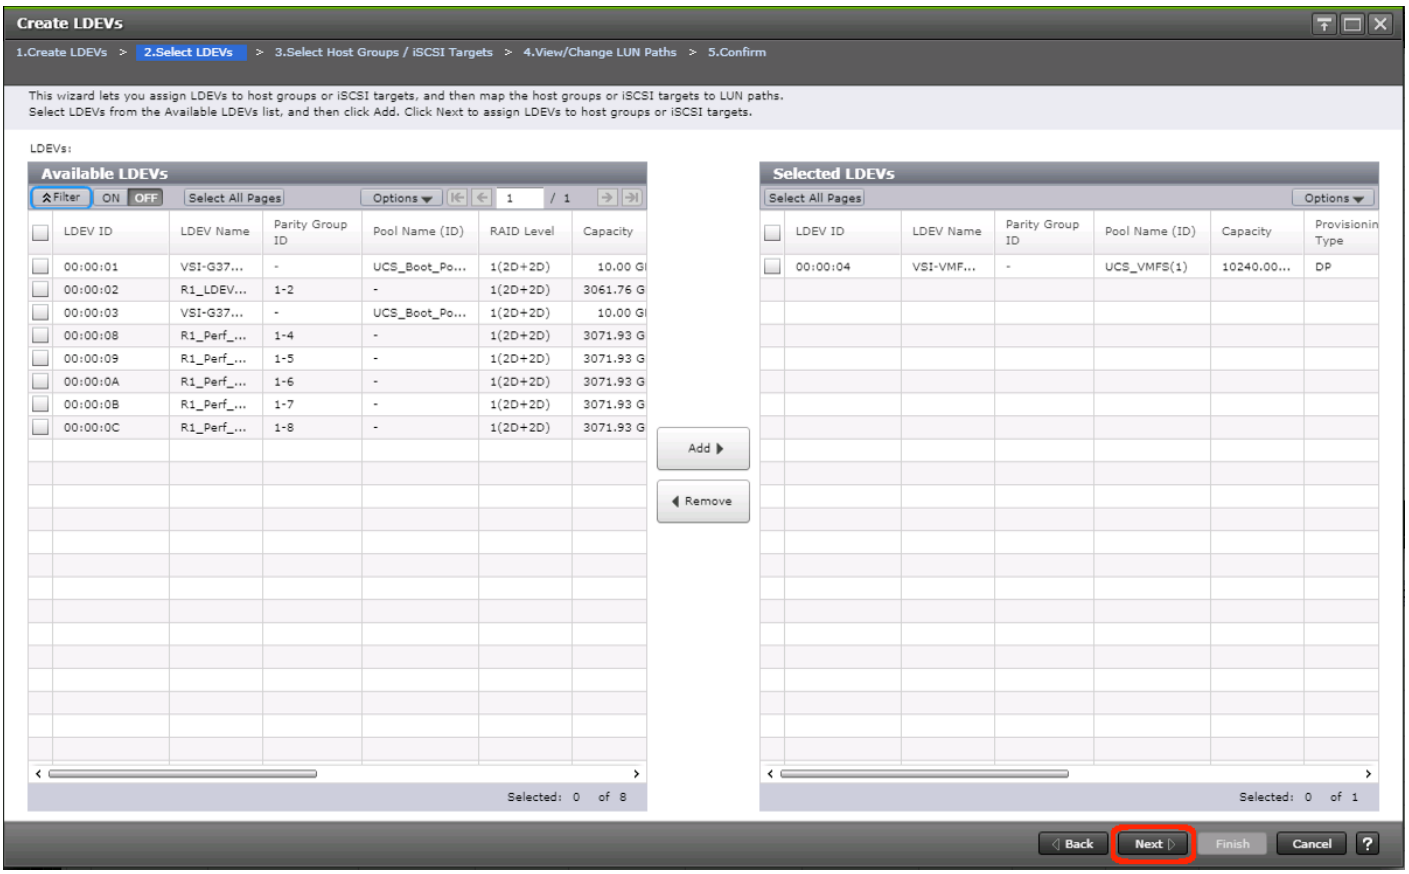

6. Ensure the newly created VMFS LDEV is the only LDEV in the Selected LDEVs pane, then click Next.

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>

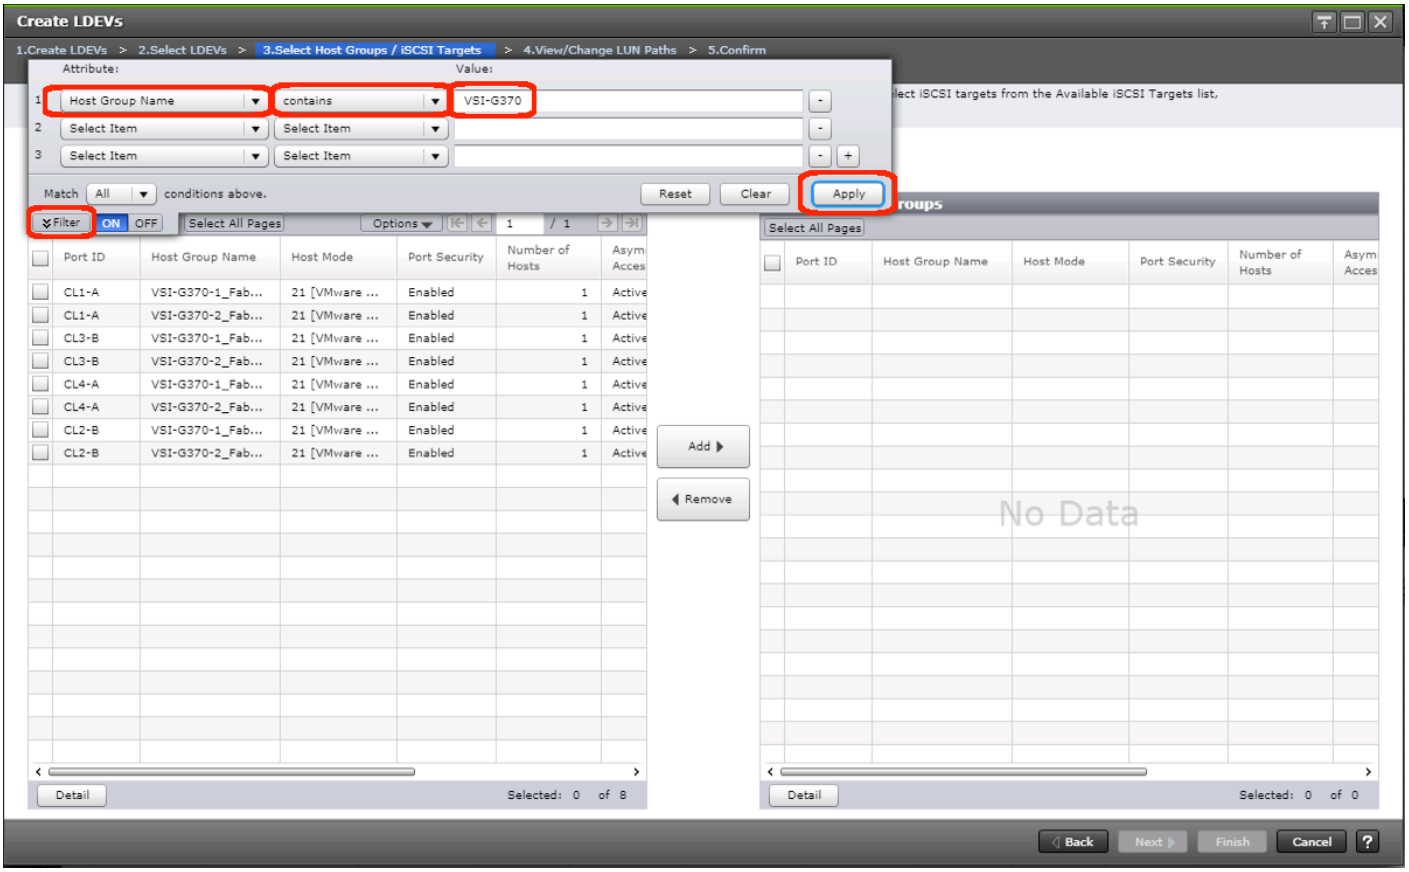

9. Click Apply.

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.

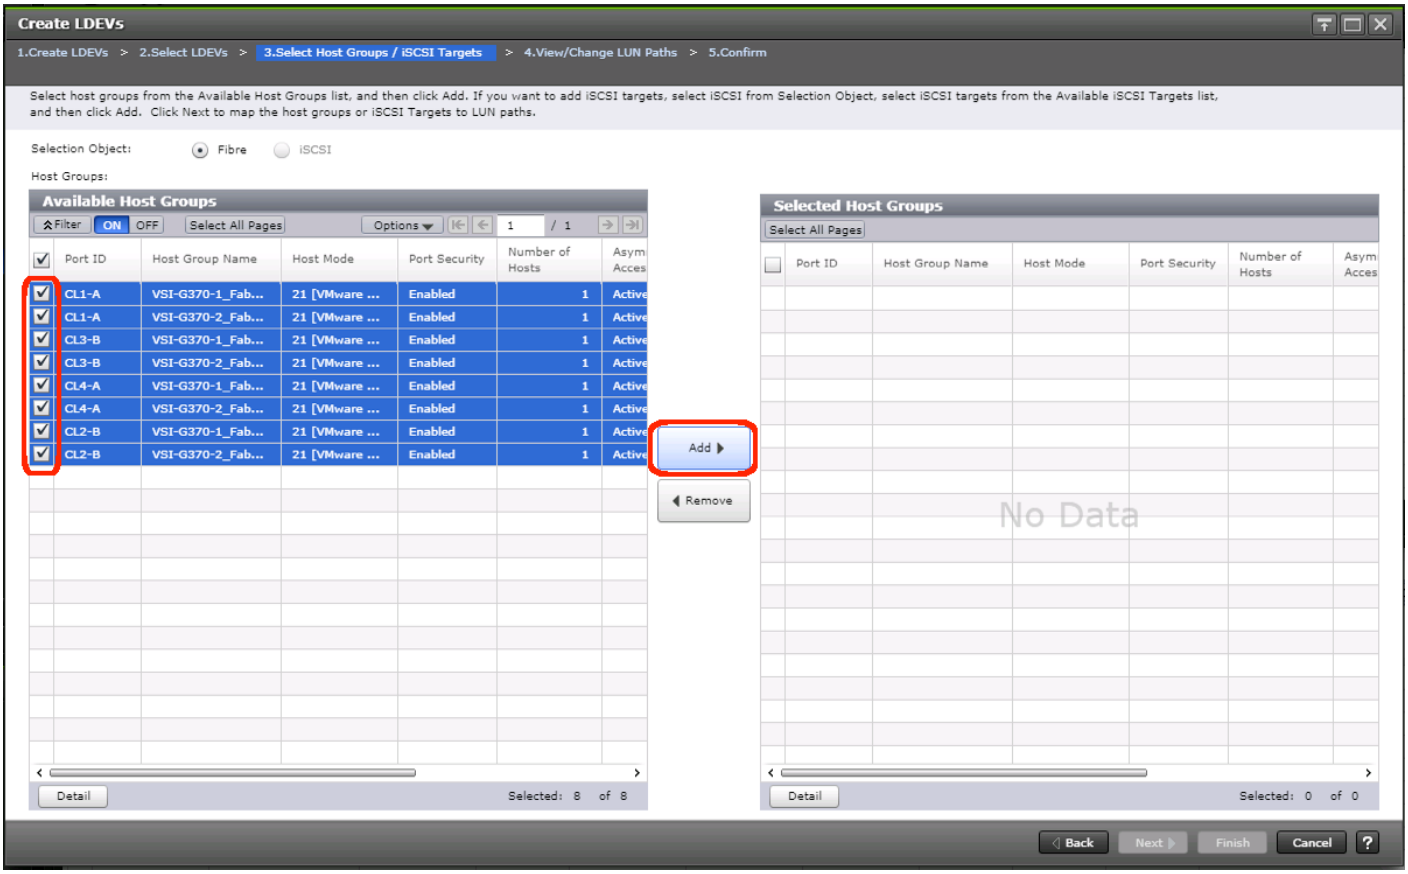

Depending on the pathing design used, you may have additional or fewer than four paths for the VMFS LDEV, but there should be a minimum of one path per fabric used.

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.

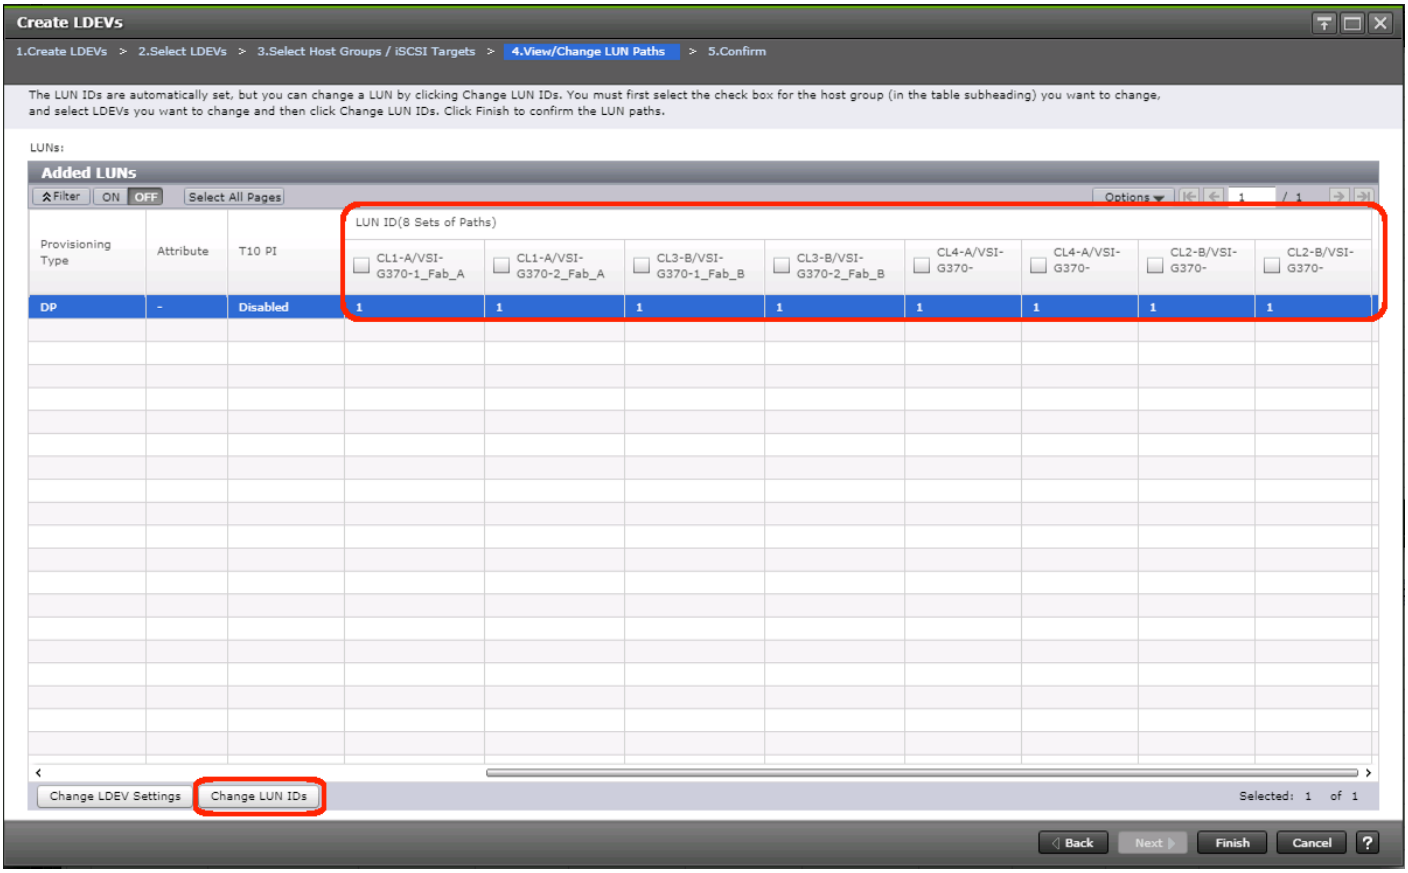

If other LDEVs have been assigned to one host but not others, you will need to modify the Host LUN ID assignment to the next Host LUN ID that is consecutive across all hosts/paths.

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 on the left side of the pane, 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.

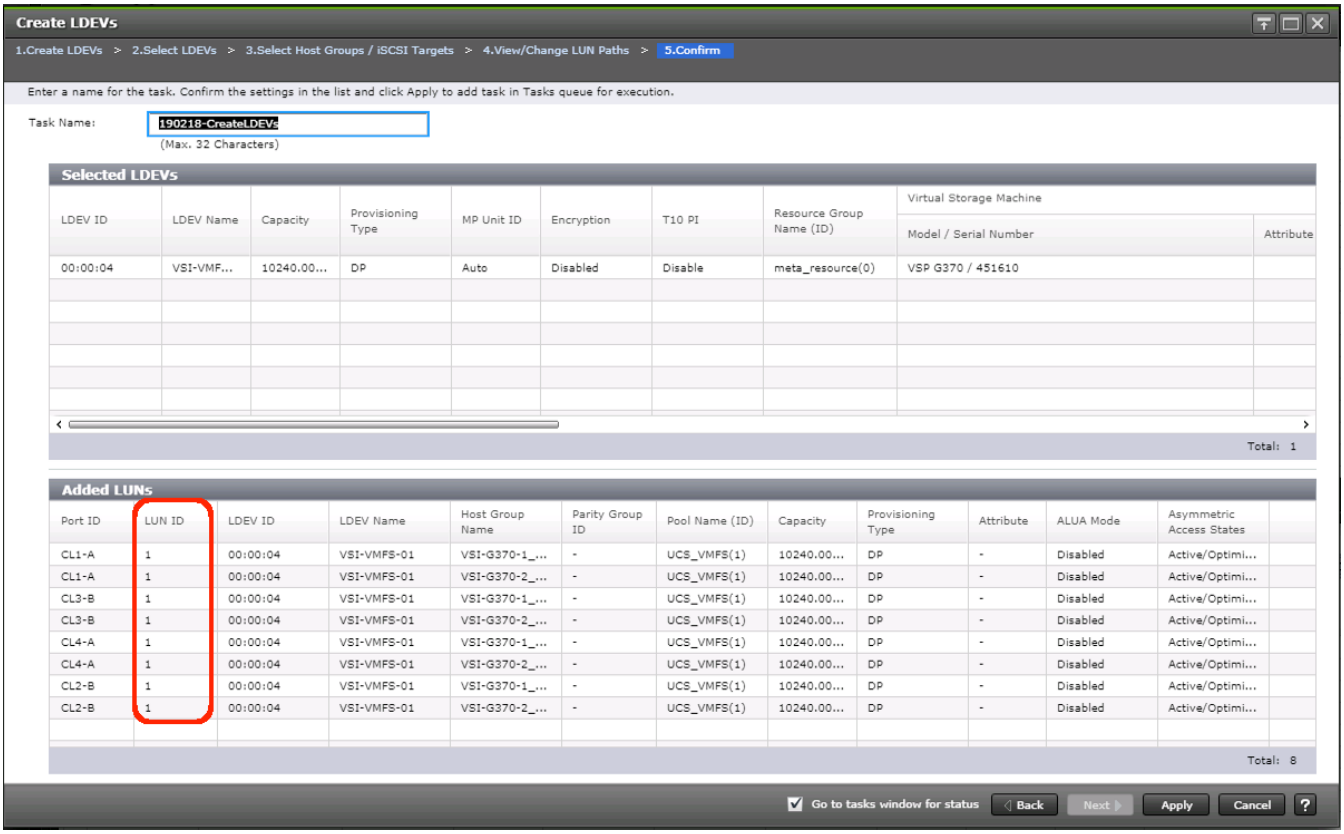

If the output is long enough, use the scrollbar on the right side of the **Added LUNs** window to ensure the **LUN ID** column contains the same LUN ID for each port listed.

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. VMware ESXi 6.5 U2 was also used in the validation, but the deployment steps are not included since they are similar.

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.

The Cisco Custom Image for ESXi 6.5 U2 is available here.

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 on the left.
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. In the center pane, 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.

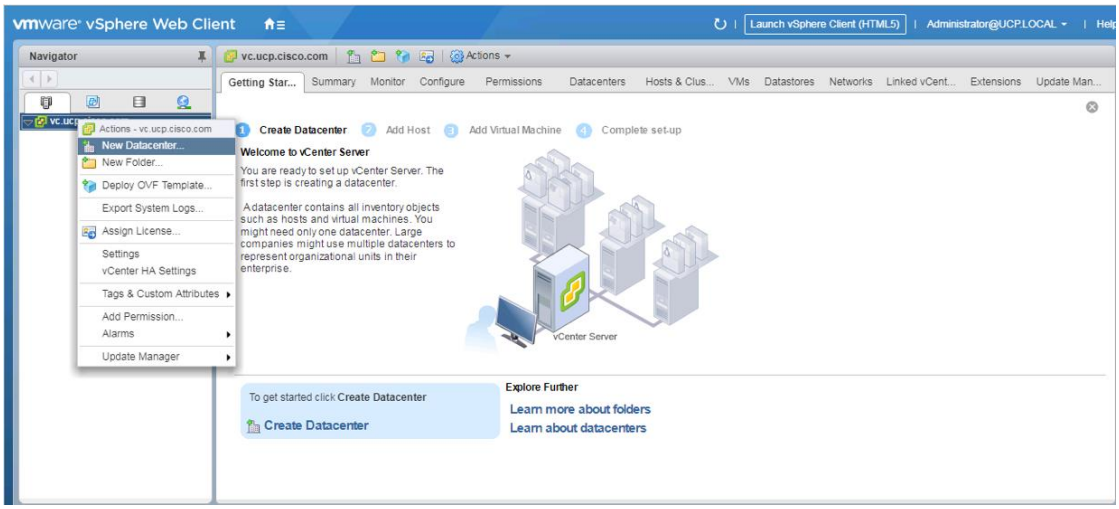

2. From the New Datacenter pop-up dialogue enter in a datacenter name and click OK.

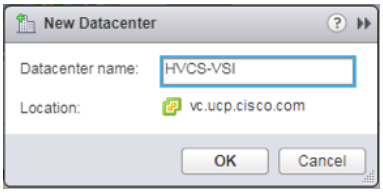

Add the VMware ESXi Hosts Using the VMware vSphere Web Client

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.

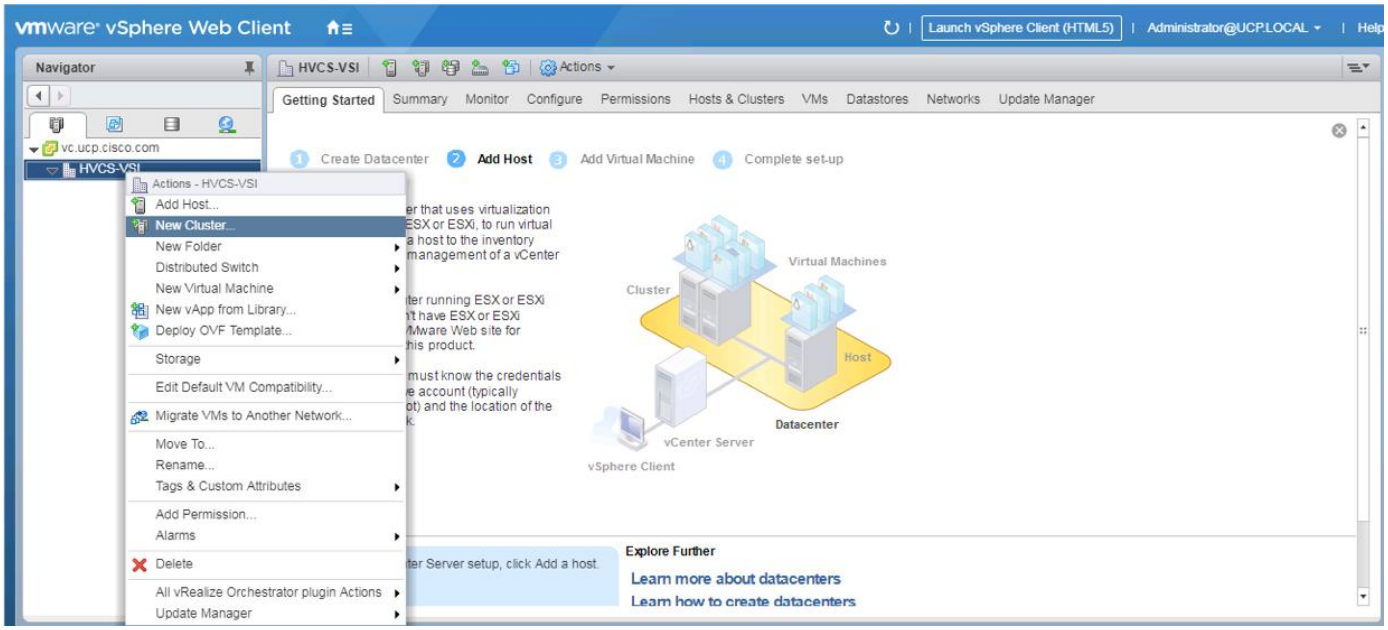

2. Enter a name for the new cluster, select the DRS and HA checkboxes, leaving all other options with the defaults.

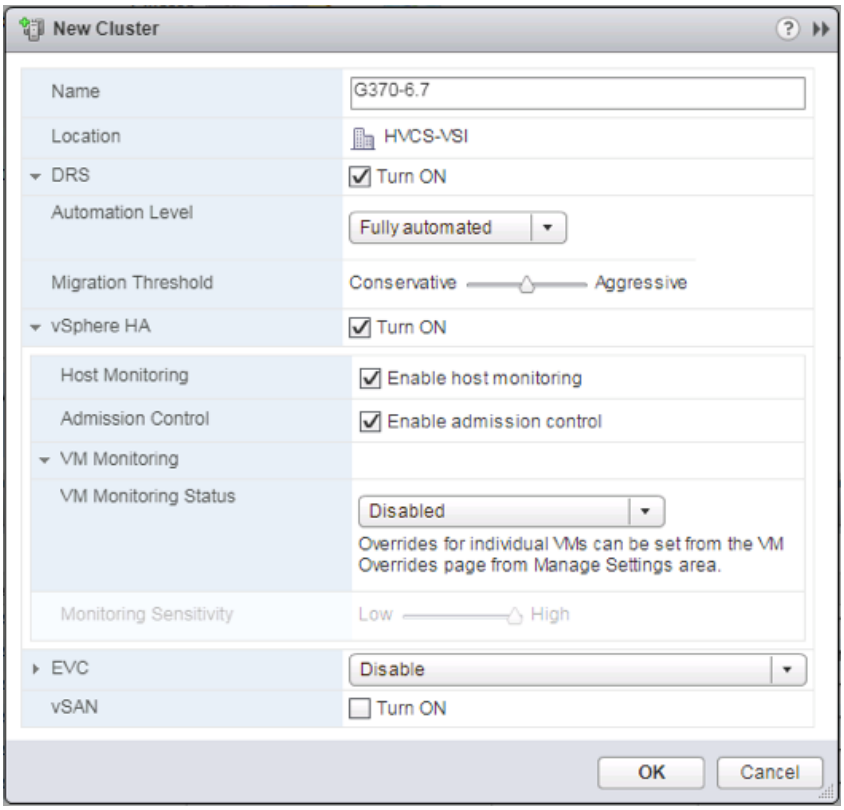

3. Click OK to create the cluster.

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

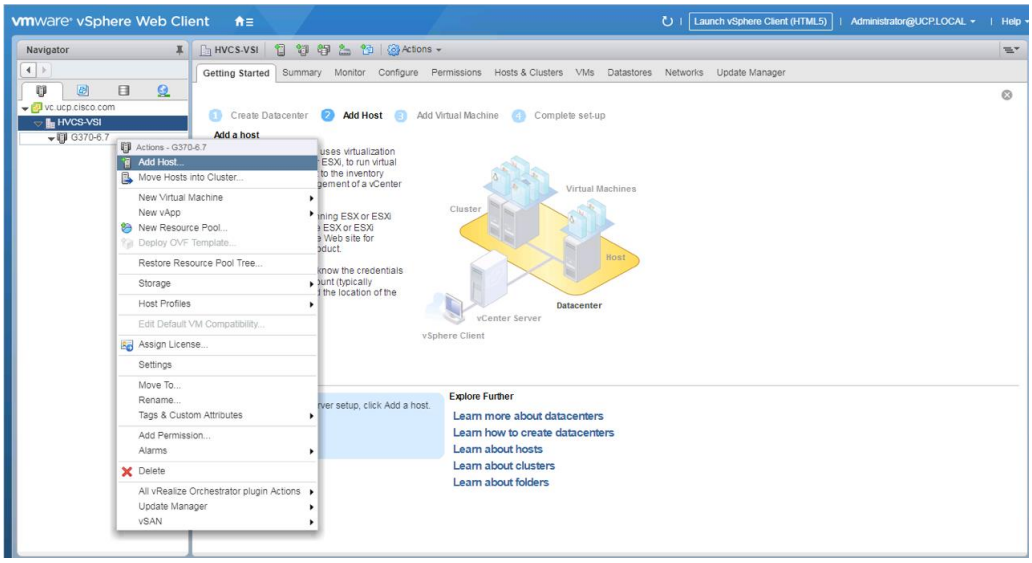

5. Enter the IP or FQDN of the first ESXi host and click Next.

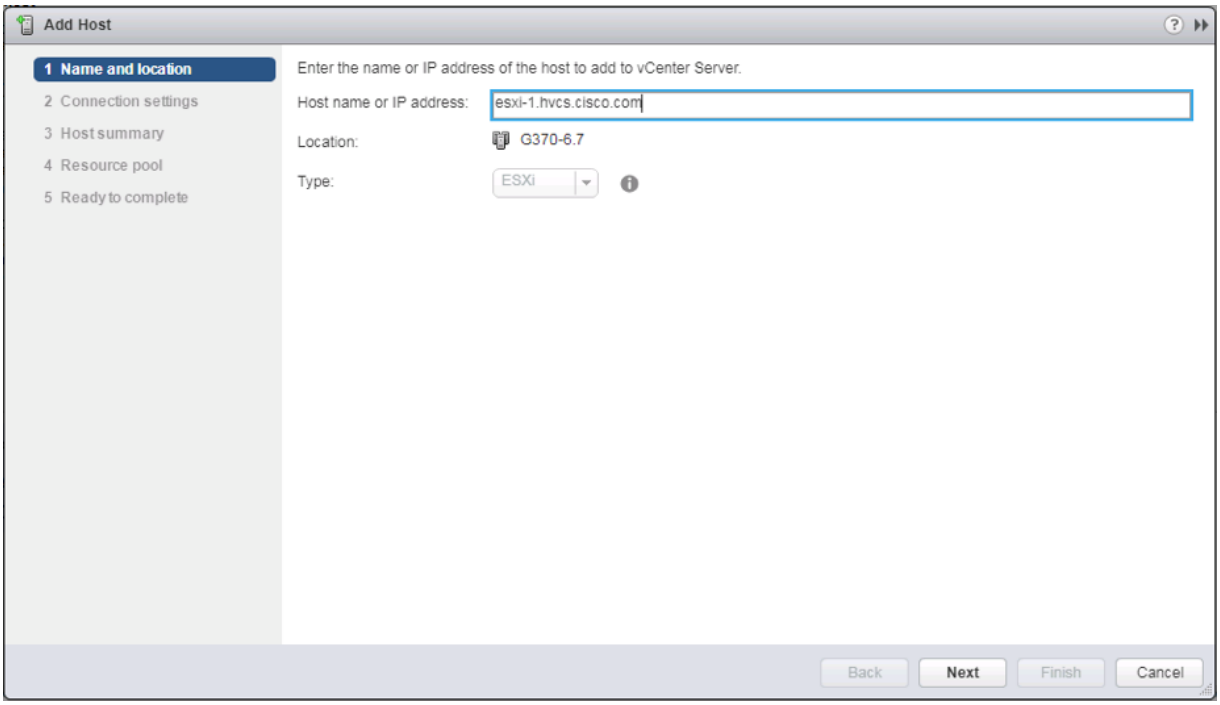

6. Enter `root` for the User Name, provide the password set during initial setup and click Next.

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.

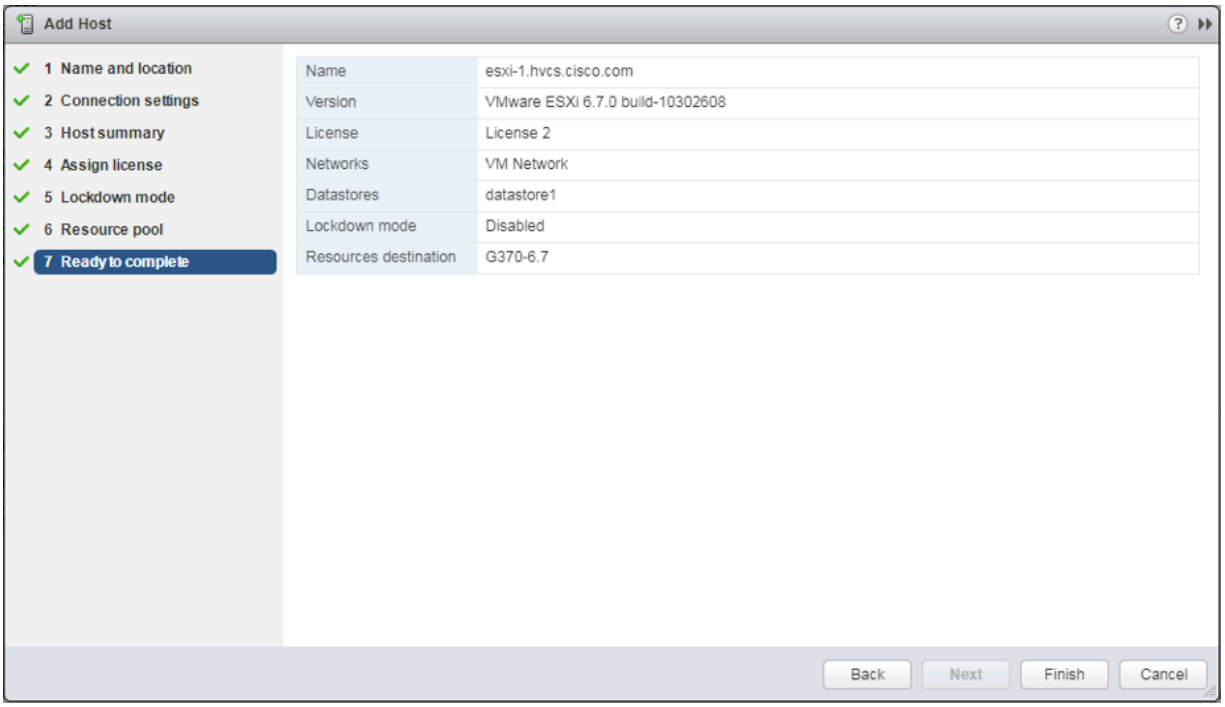

13. Repeat steps 1-12 for each ESXi host to be added to the cluster.

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…

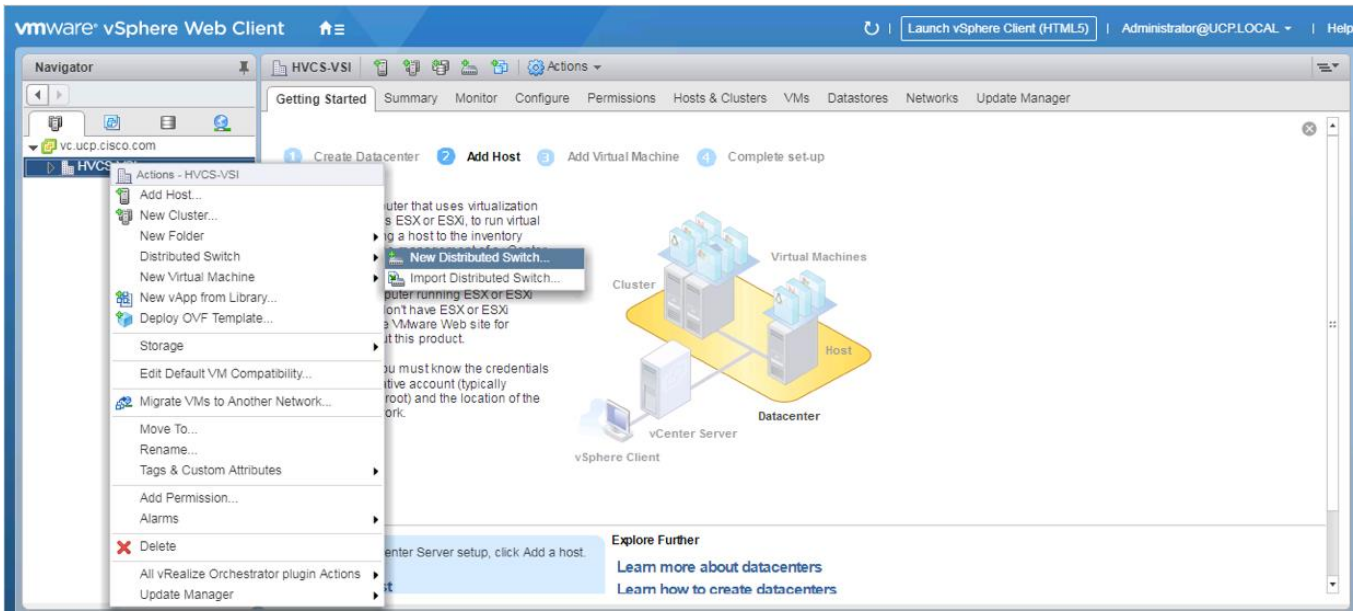

2. Give the Distributed Switch a descriptive name and click Next.

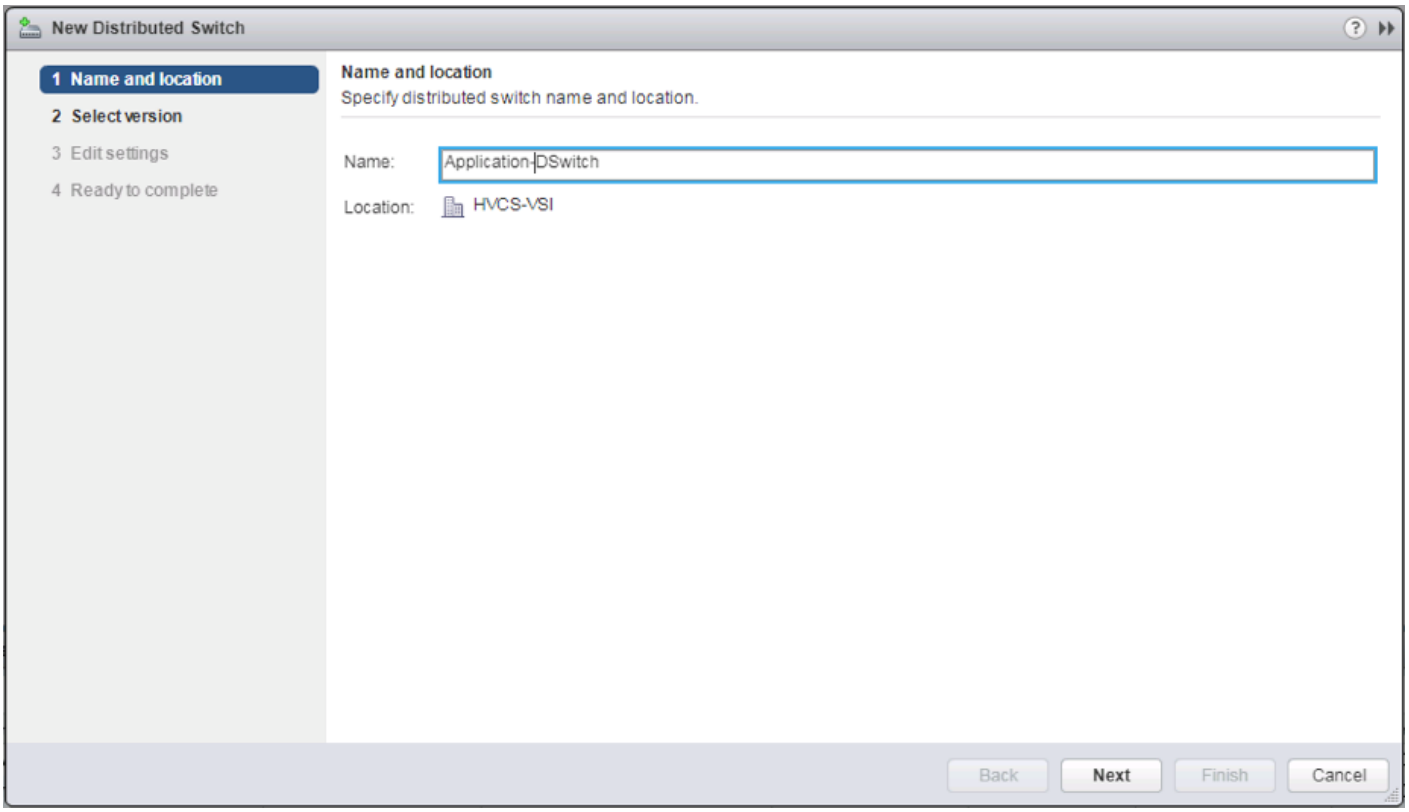

3. Make sure Distributed switch: 6.5.0 is selected if supporting vSphere 6.5 hosts and click Next.

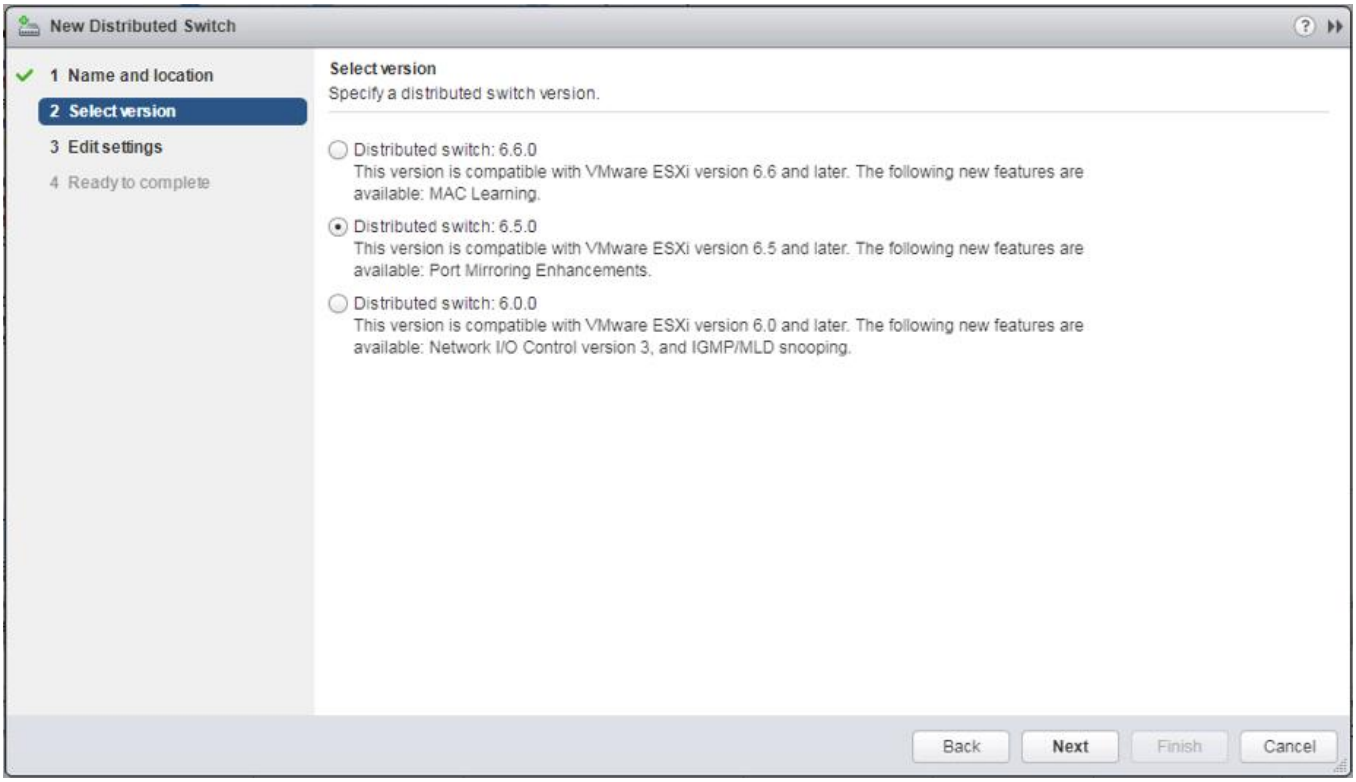

4. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter App-201 for the name of the default Port group to be created. Click Next.

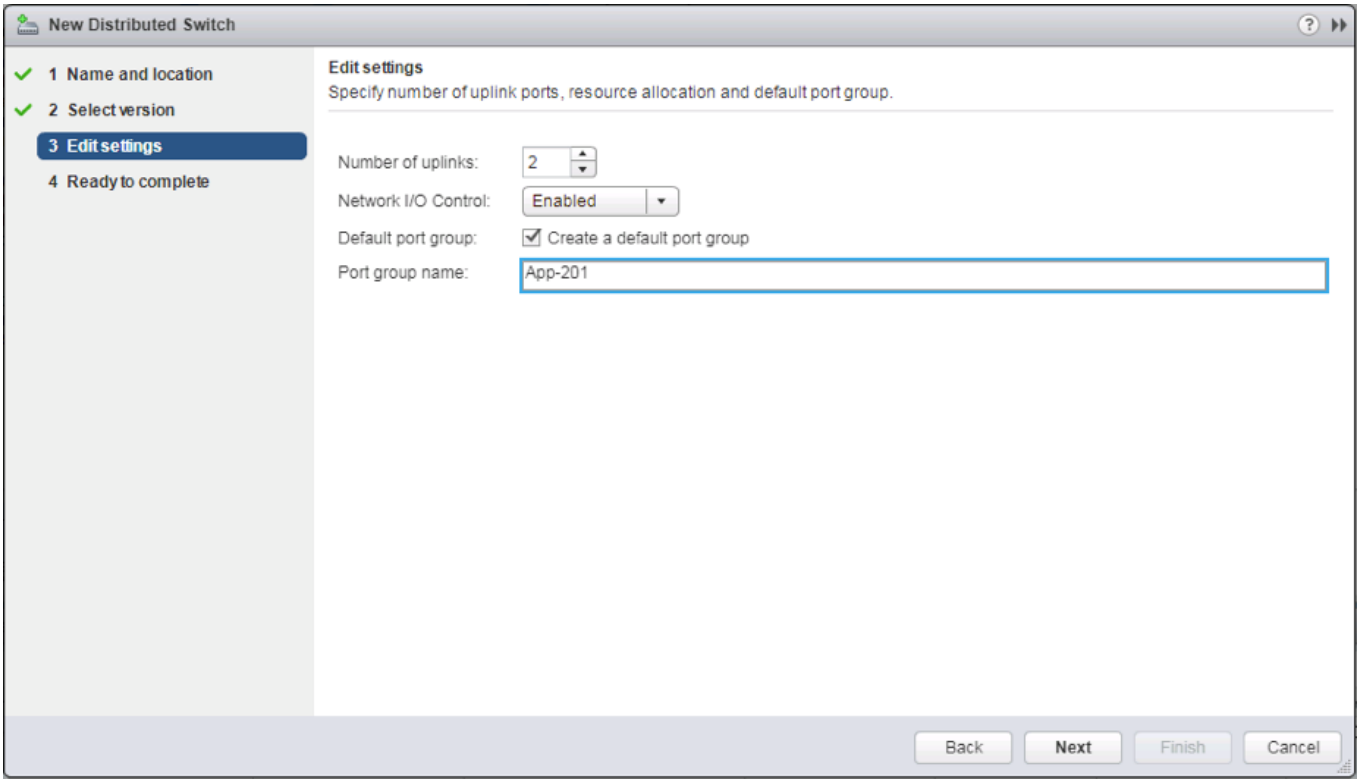

5. Review the information and click Finish to complete creating the vDS.

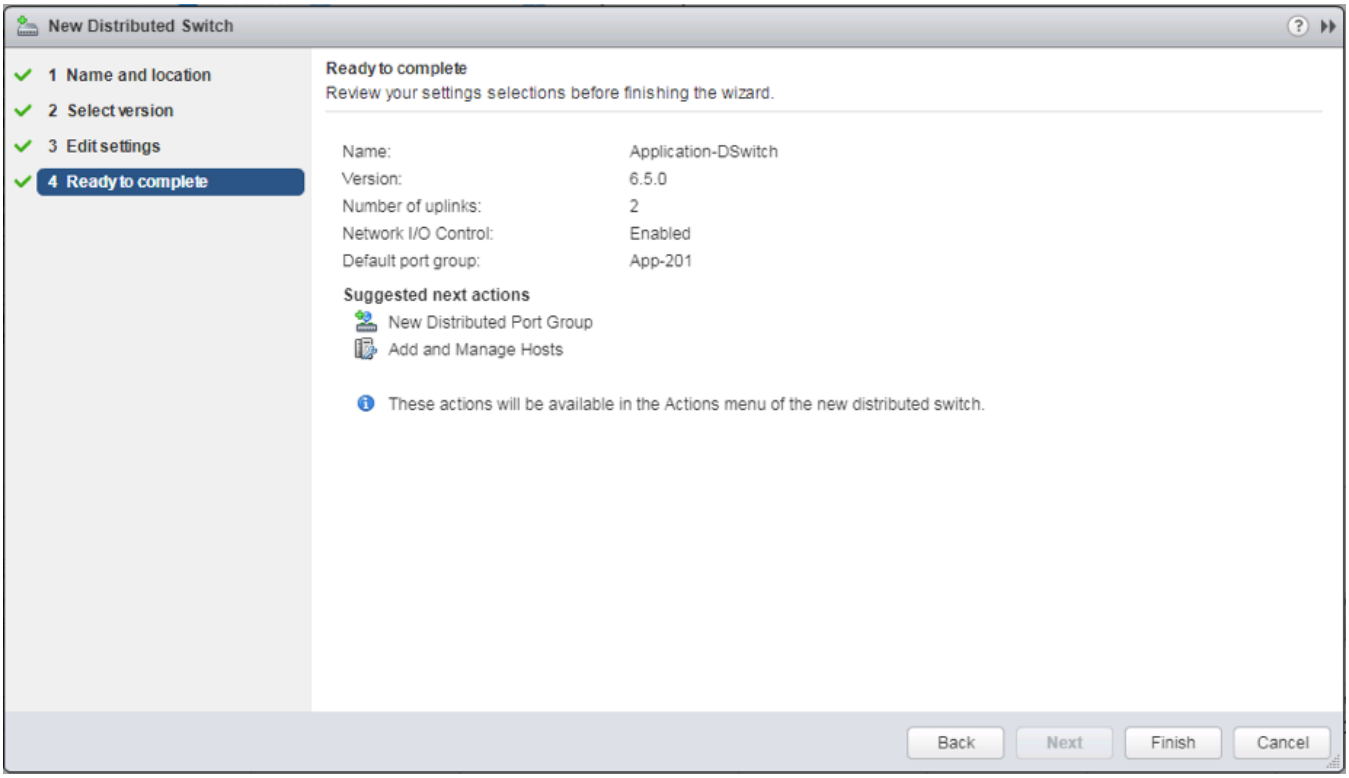

6. Right-click the newly created App-DSwitch vDS, and select Settings -> Edit Settings…

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.

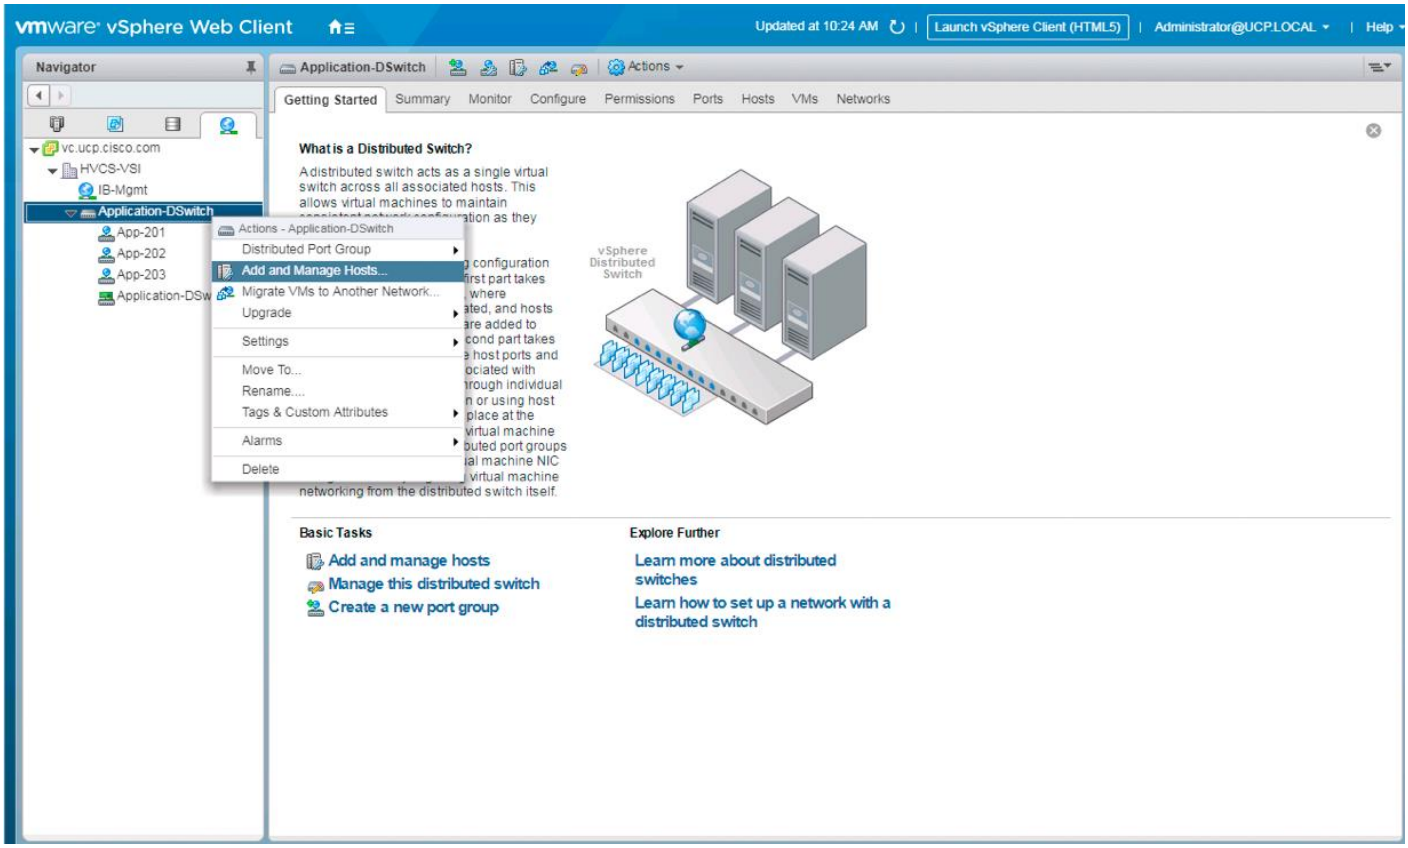

2. Right-click it and select the Add and Manage Hosts… option.

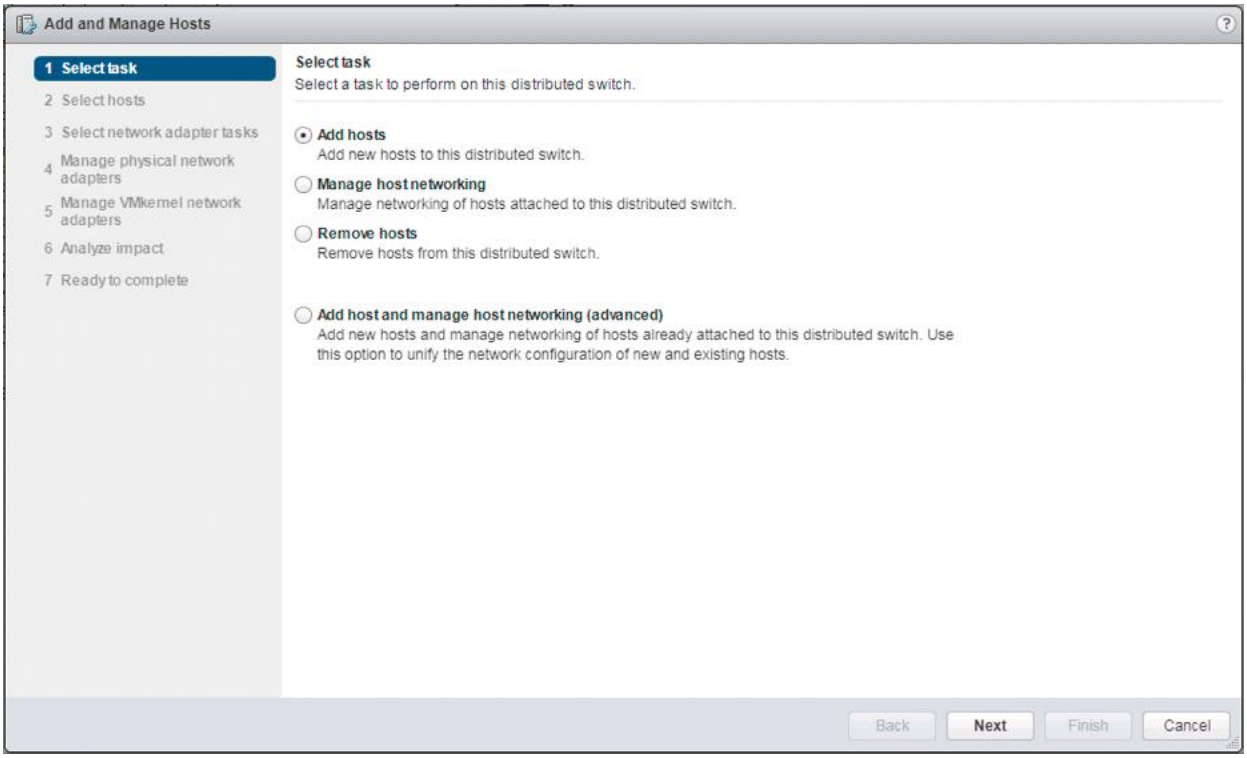

3. Leave Add hosts selected, and click Next.

4. Click the + New hosts… option.

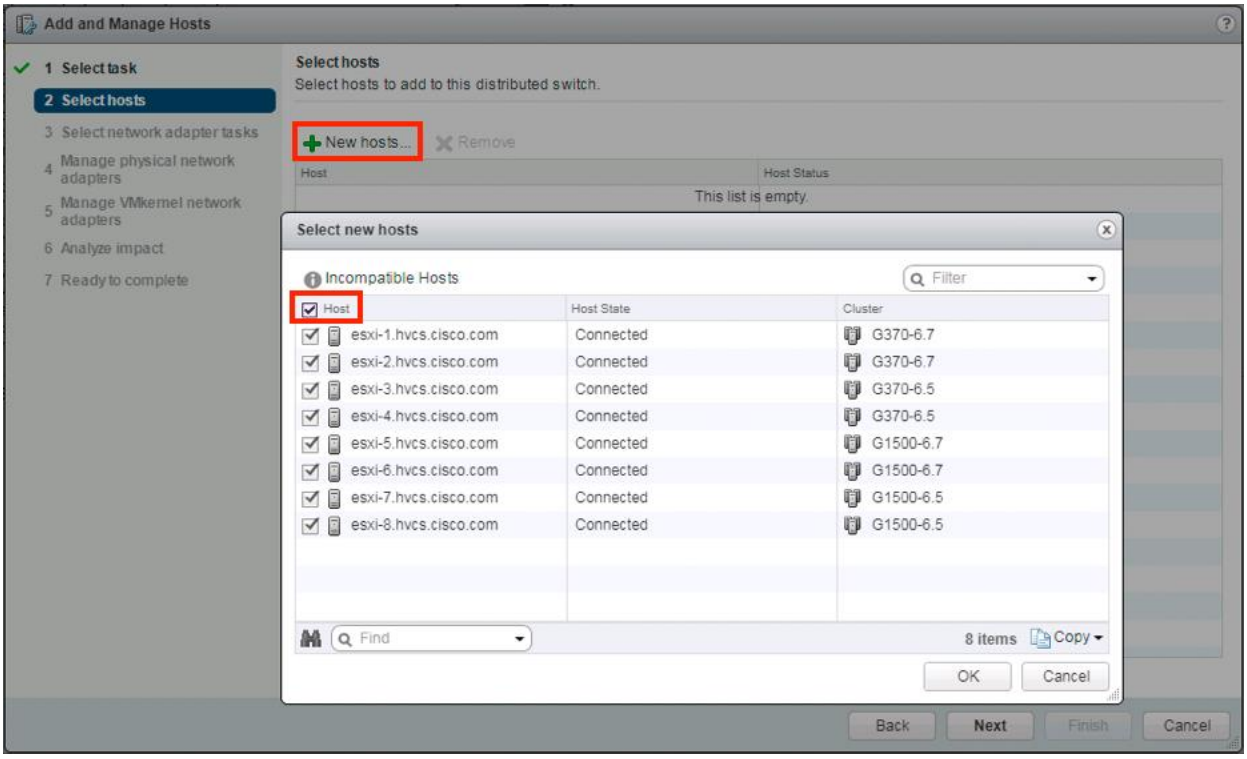

5. Select the Hosts checkbox near the top of the pop-up window to select all hosts and click OK.

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

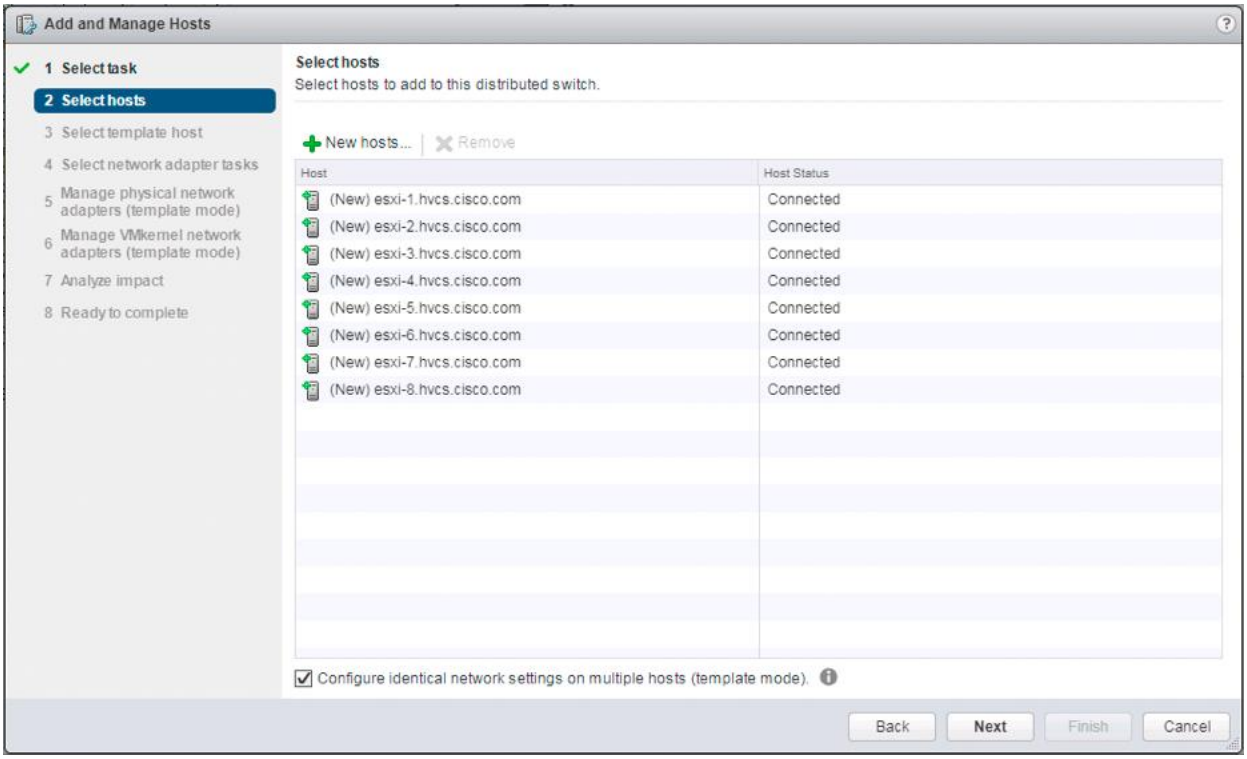

7. Click Next.

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

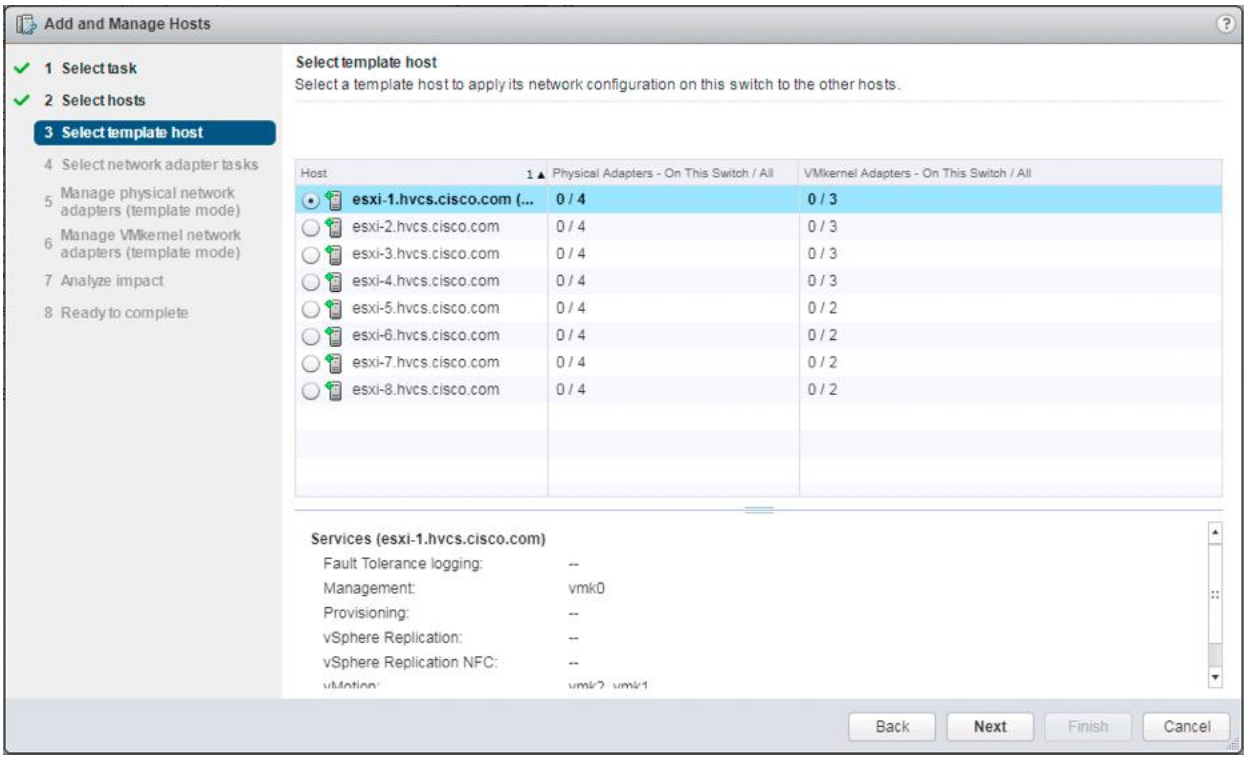

9. Click Next.

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

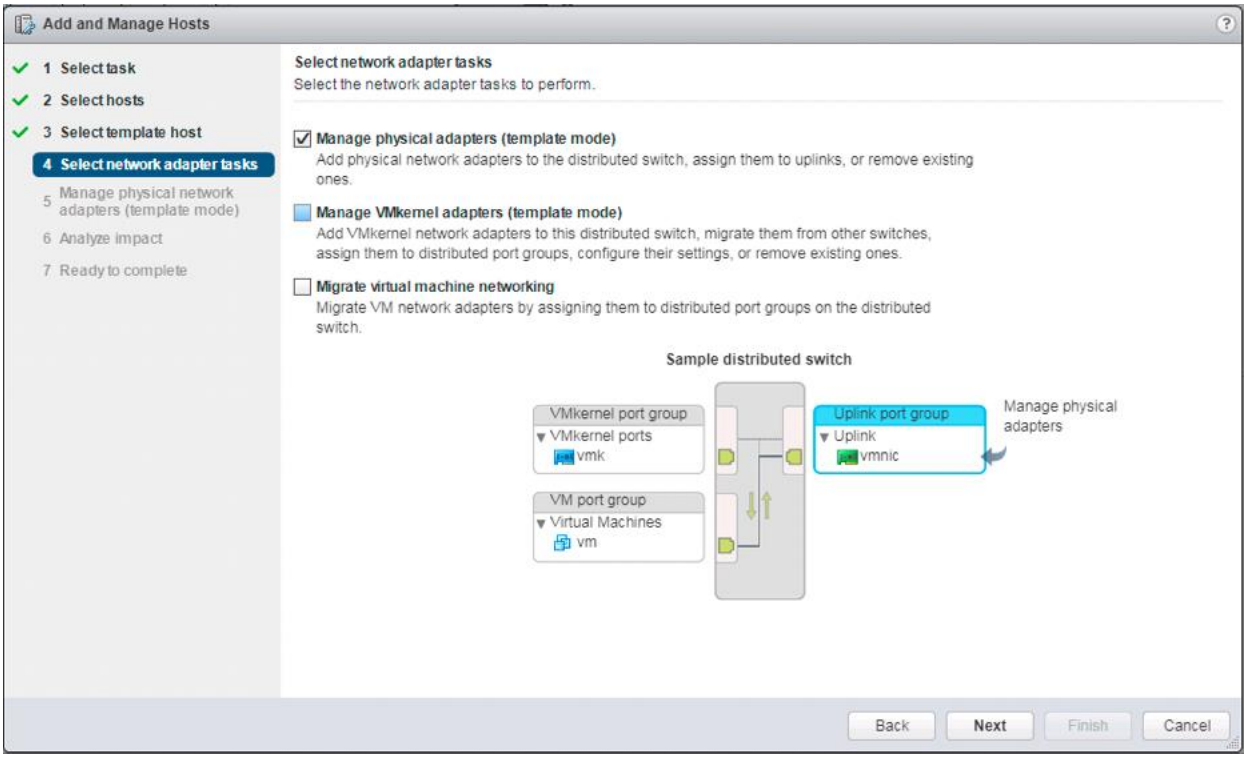

11. Click Next.

12. Select vmnic2.

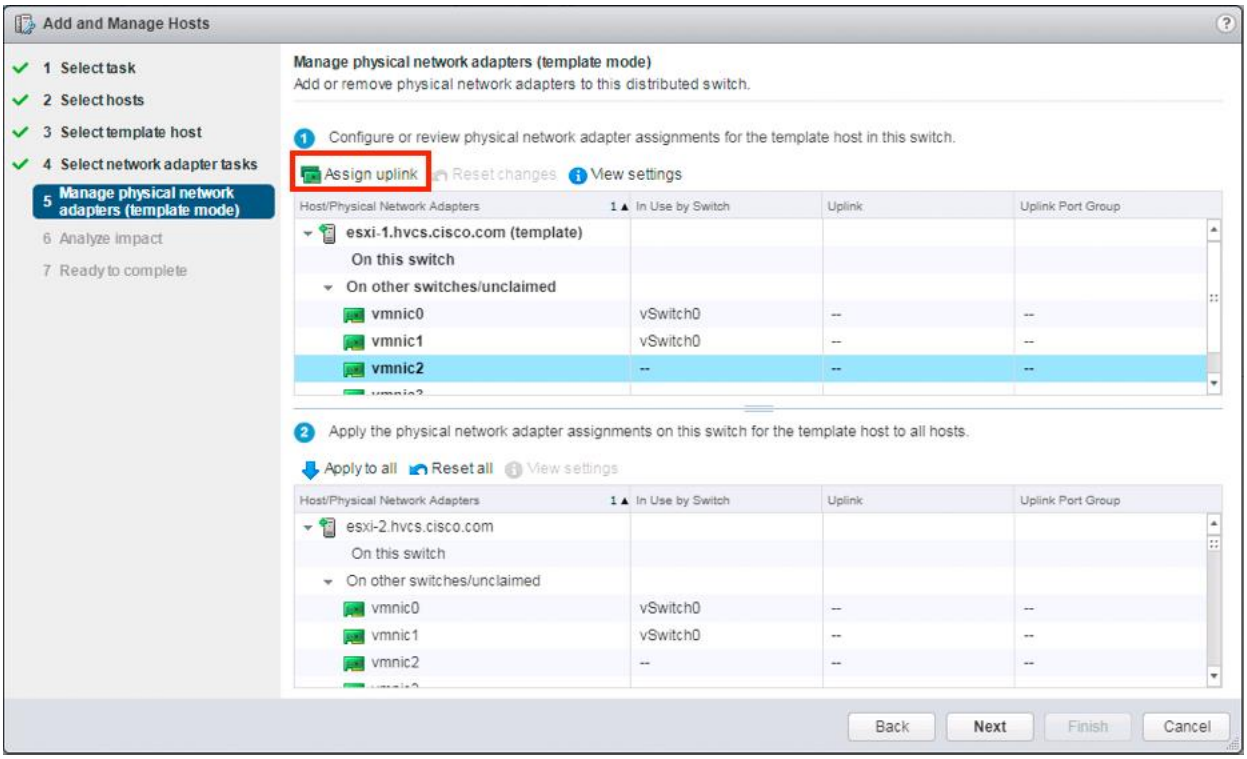

13. Click the Assign uplink option.

14. Leave Uplink 1 selected

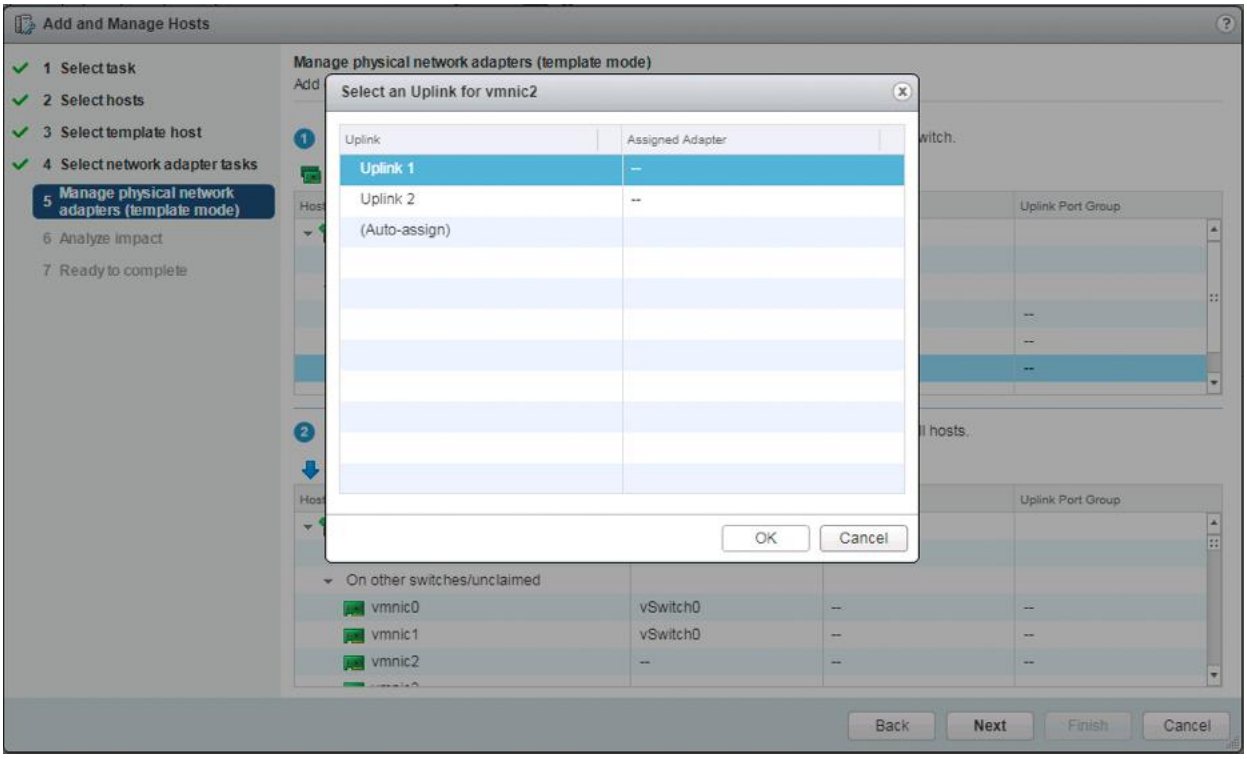

15. Click OK.

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

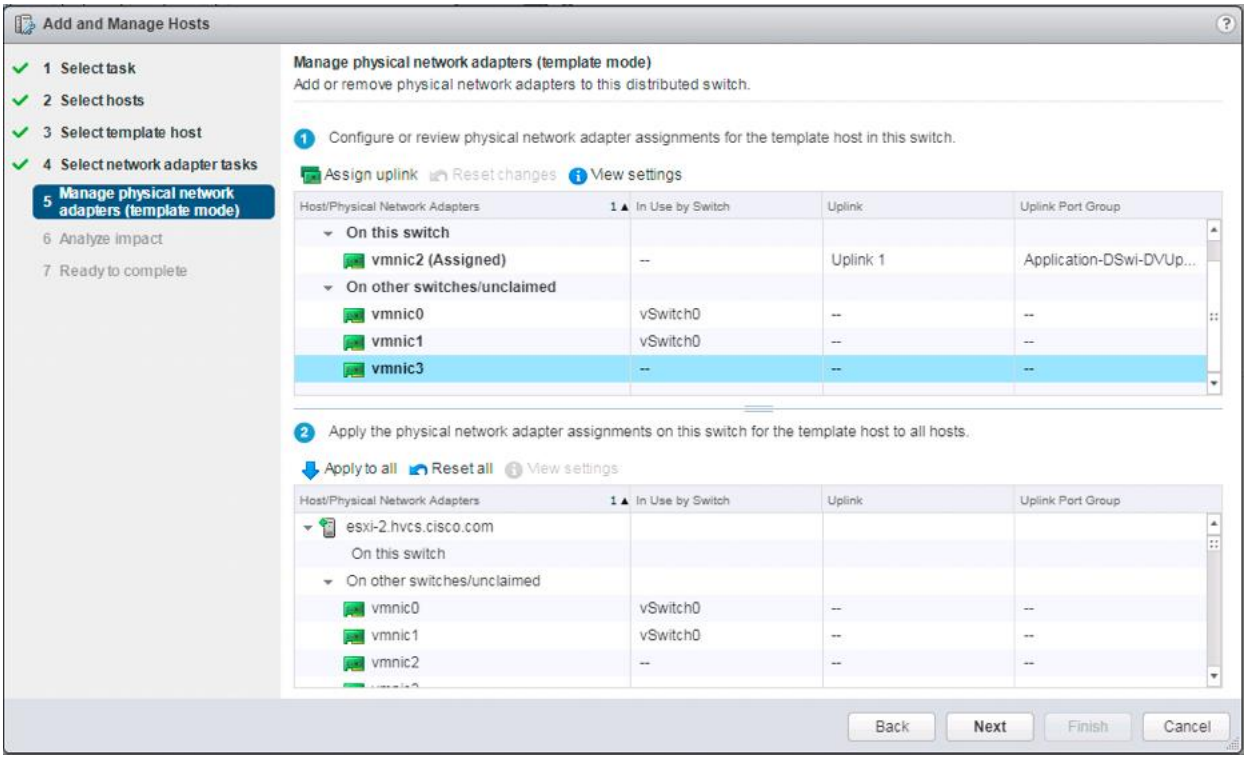

17. Click Apply to all.

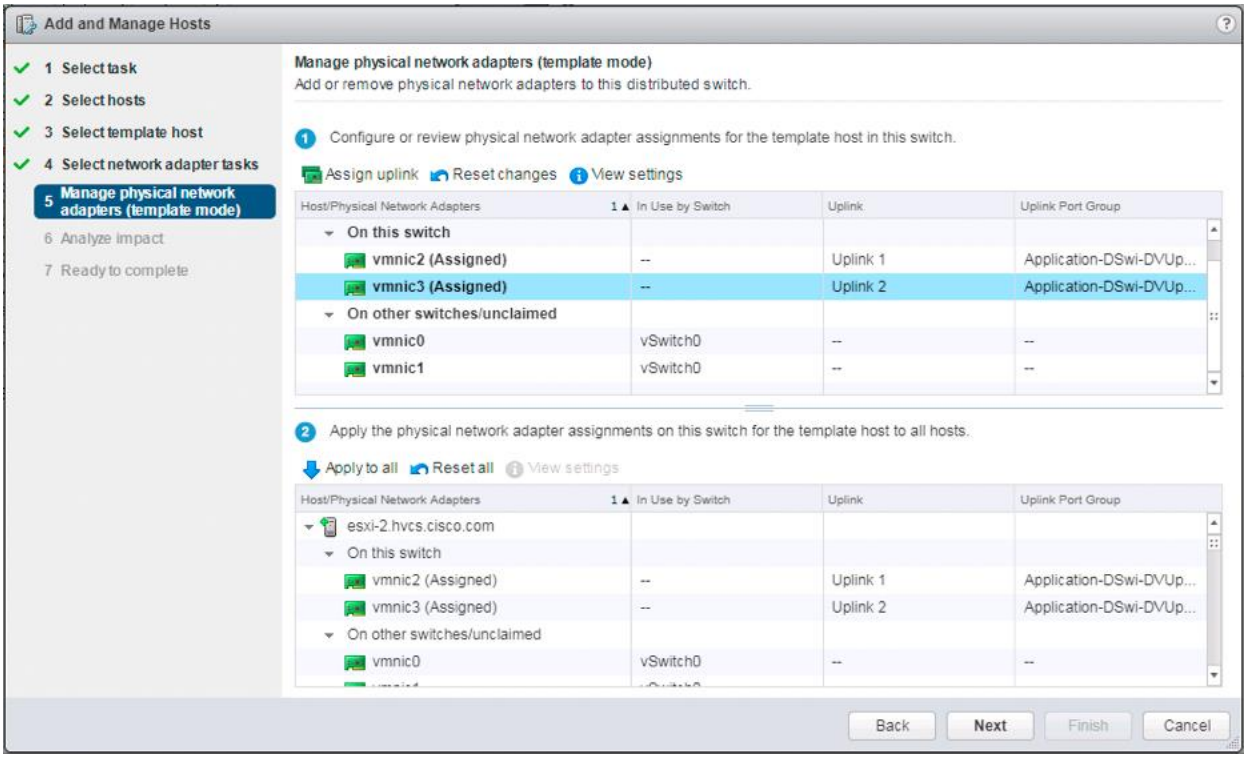

18. Click Next.

19. Click Next past Analyze Impact.

20. Verify the summary in the Ready to complete screen.

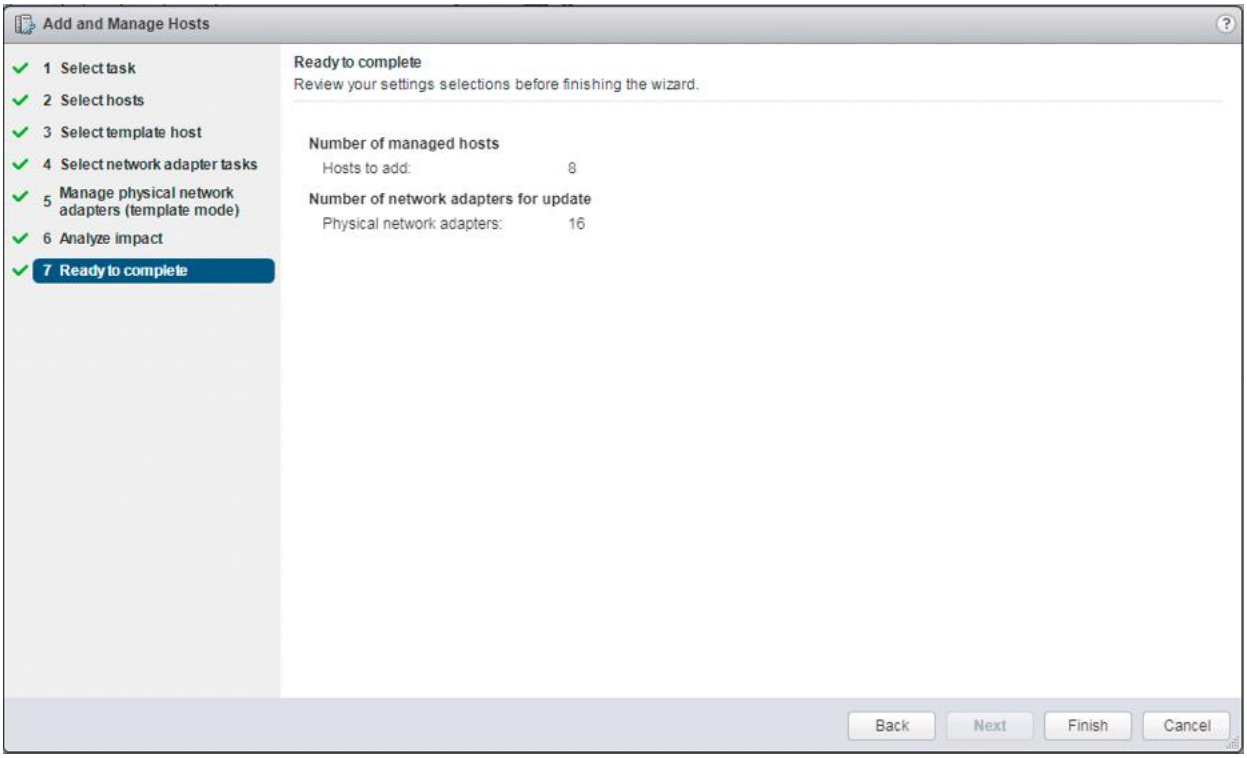

21. Click Finish to add the hosts.

Add Datastores to Hosts

Datastores have been provisioned and zoned for each of the respective VSP associated clusters (G370 versus G1500). 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…

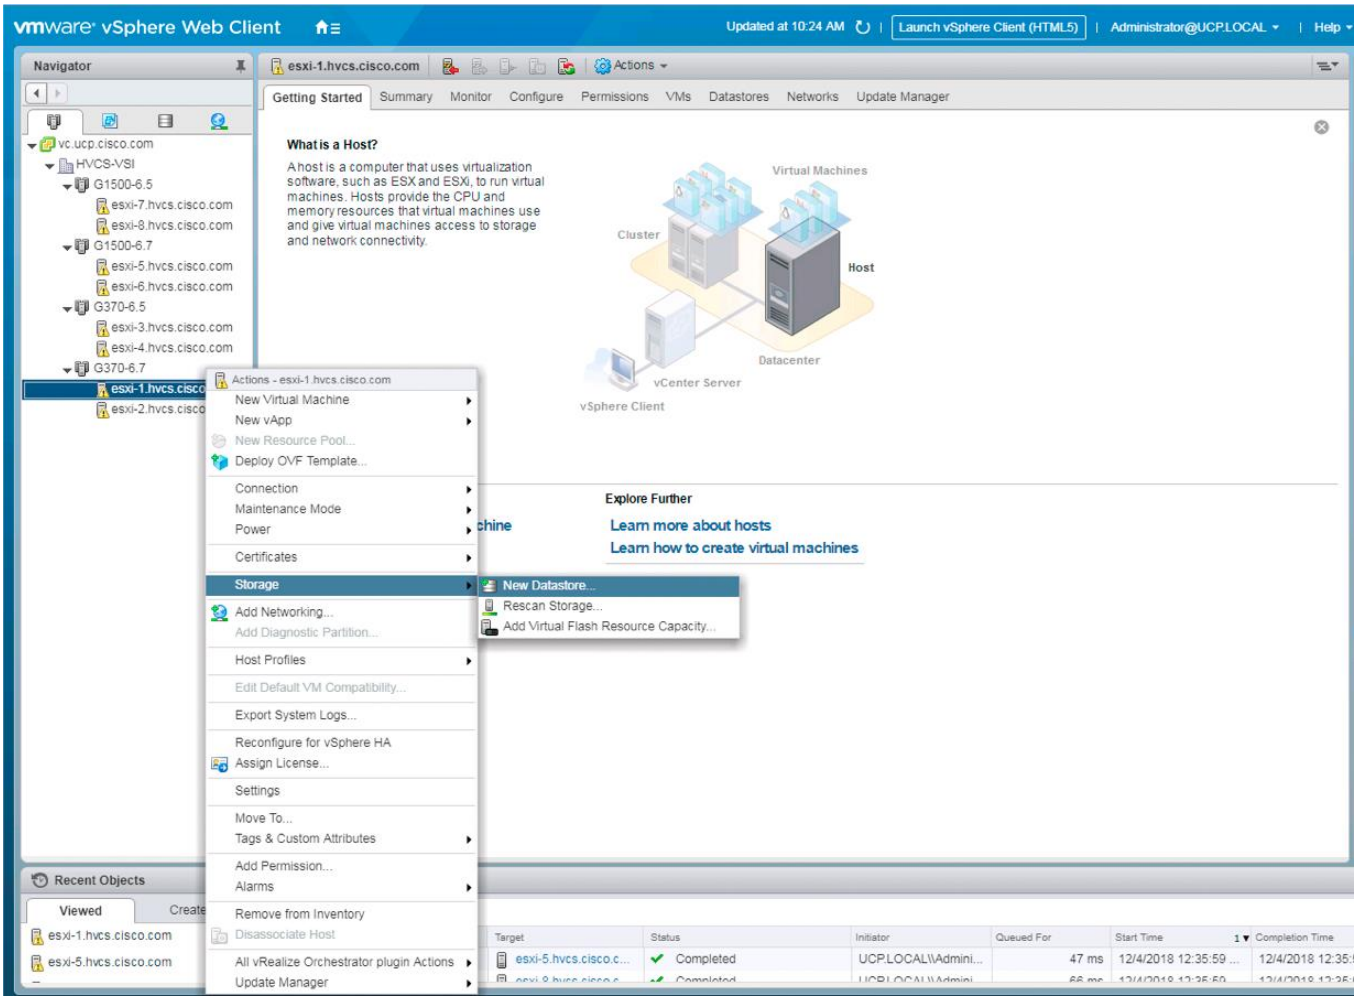

2. Leave VMFS selected and click Next.

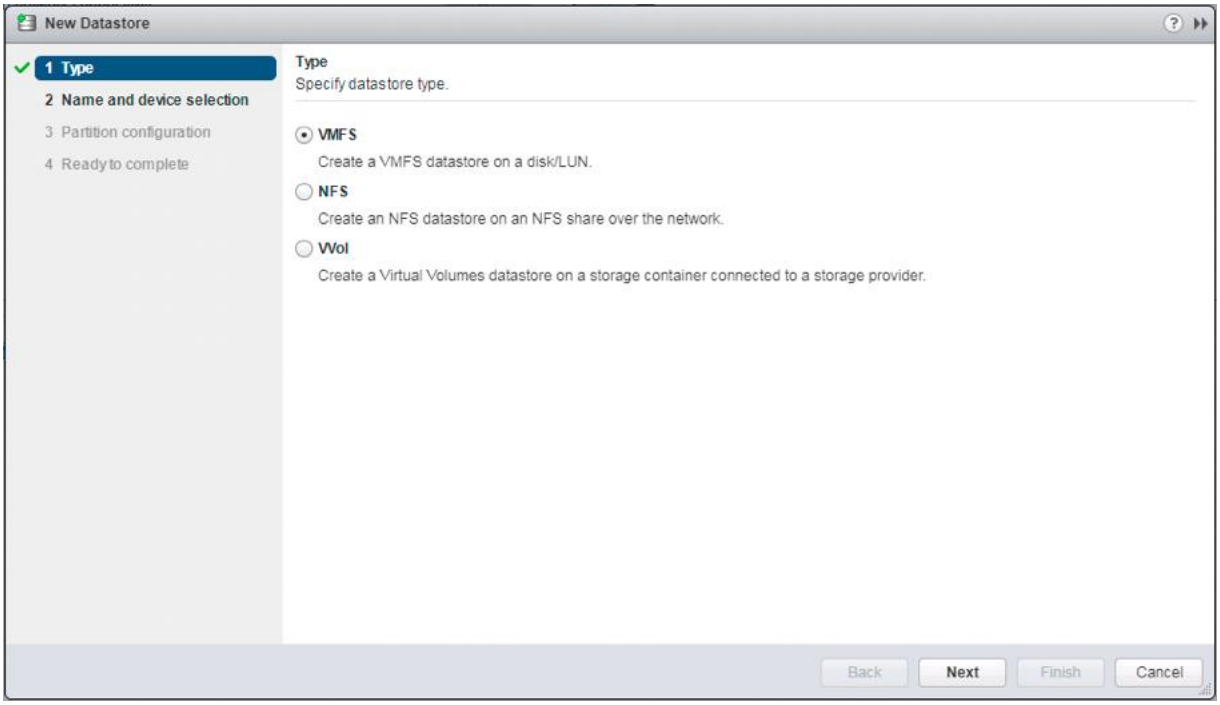

3. Provide an appropriate Datastore name and select the appropriate LUN.

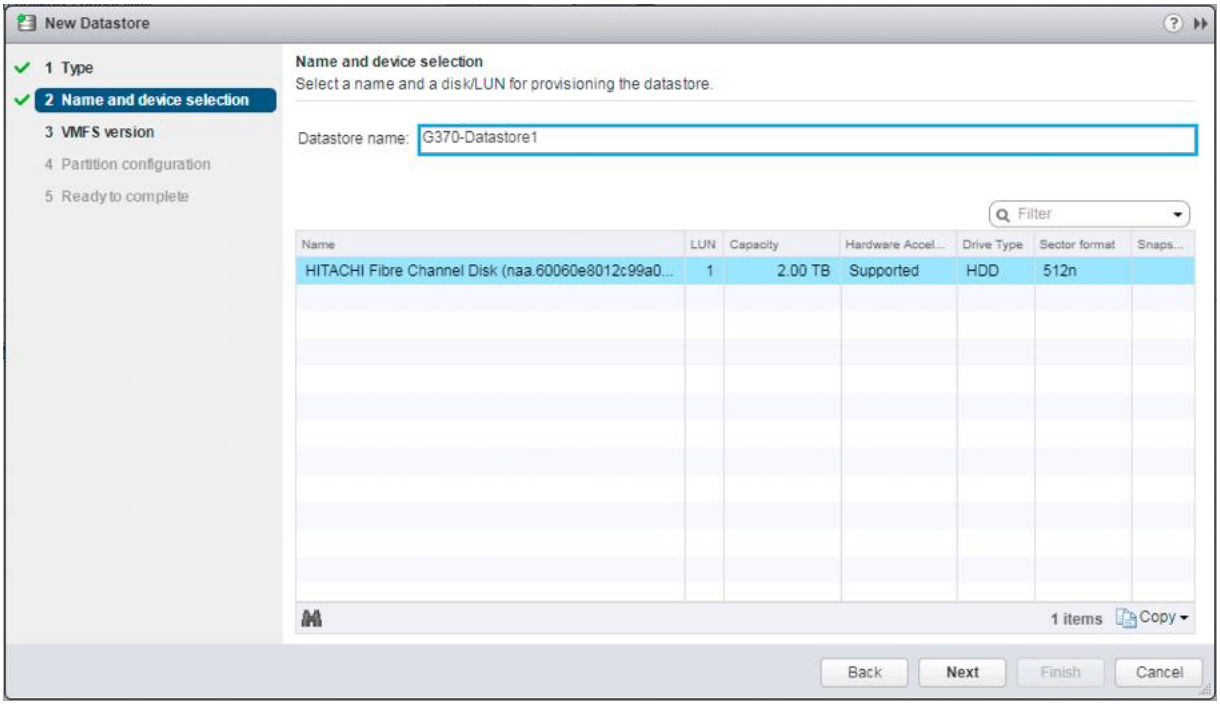

4. Click Next.

5. Leave VMFS 6 selected.

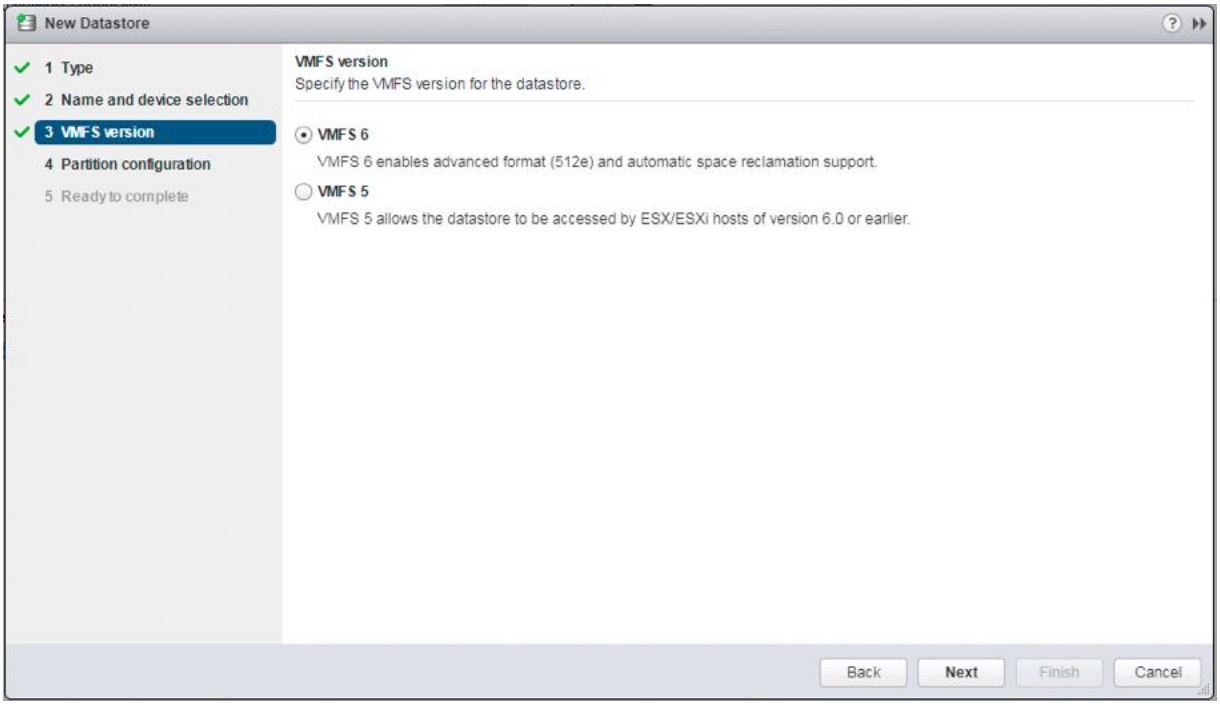

6. Click Next.

7. Leave the defaults for Partition configuration.

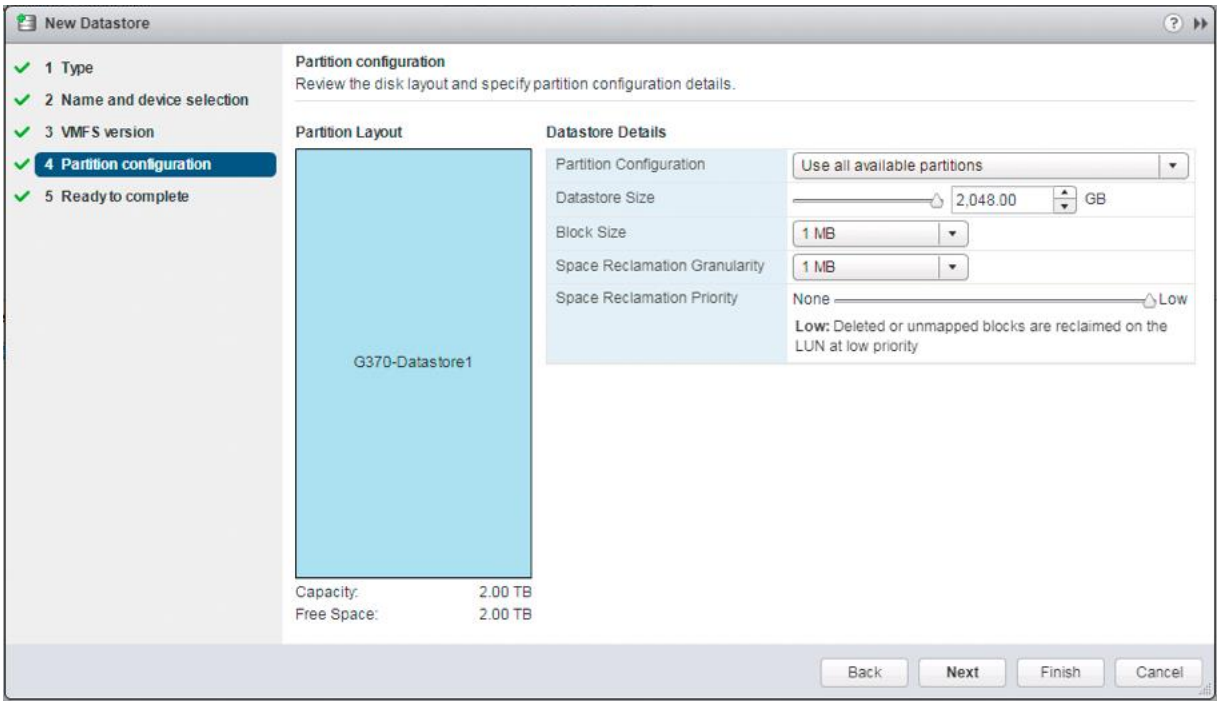

8. Click Next.

9. Review the settings.

10. Click Finish to create.

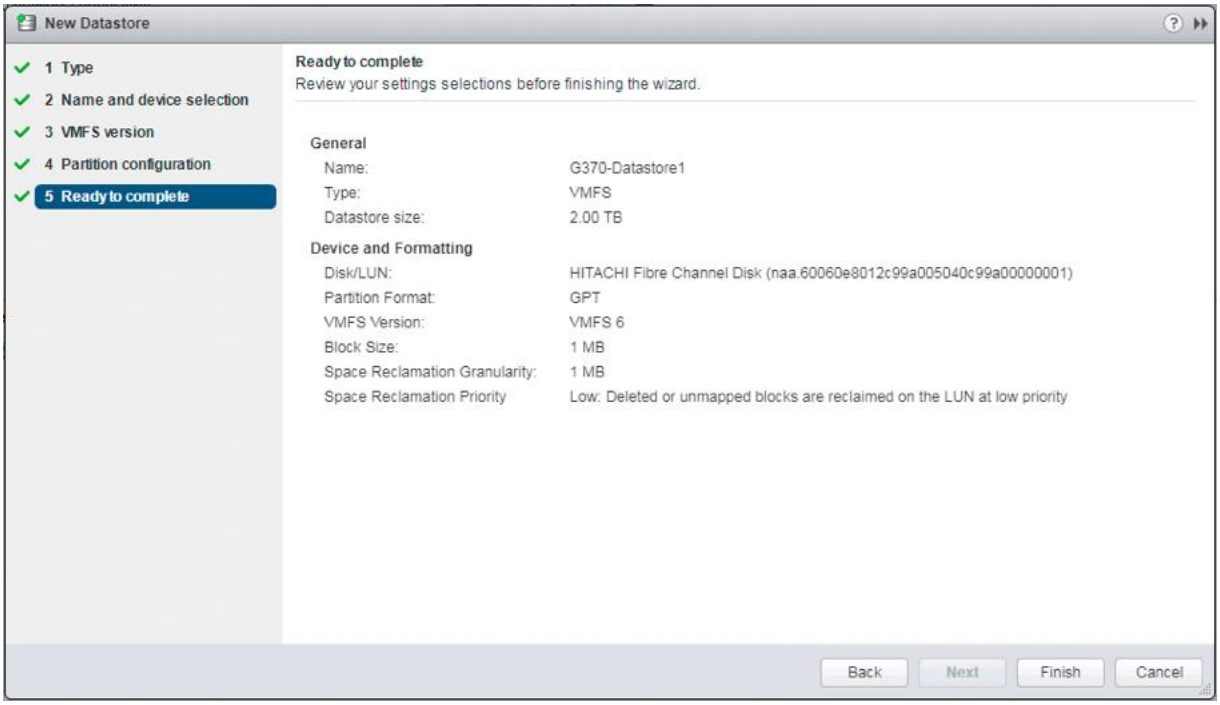

11. Check each host in the clusters associated to the same VSP the datastore was provisioned from; all should show the configured datastore as available:

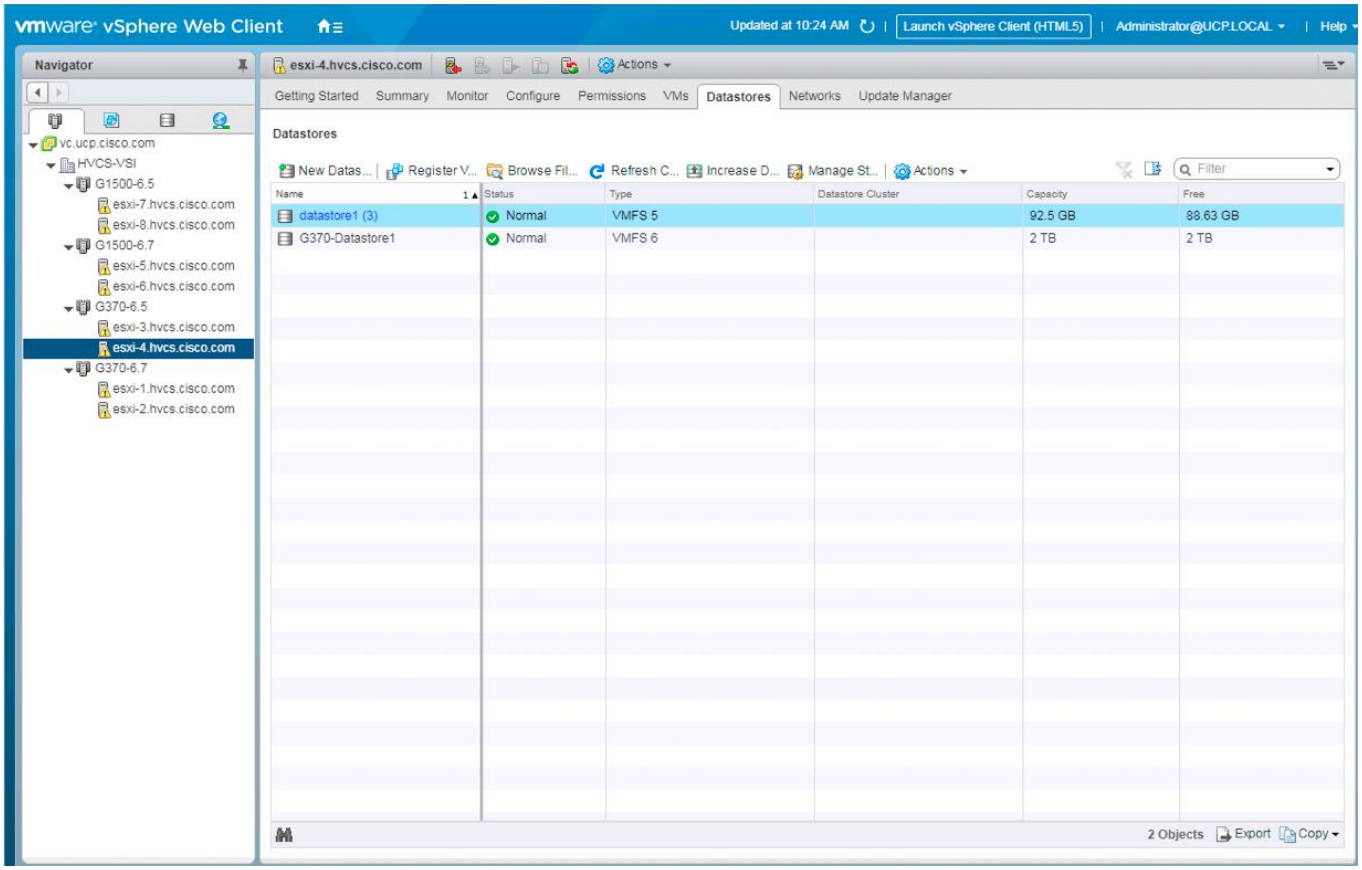

12. If multiple VMFS datastore LUNs were deployed, repeat these steps on a host they are associated to by a Hitachi Host Group.

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.

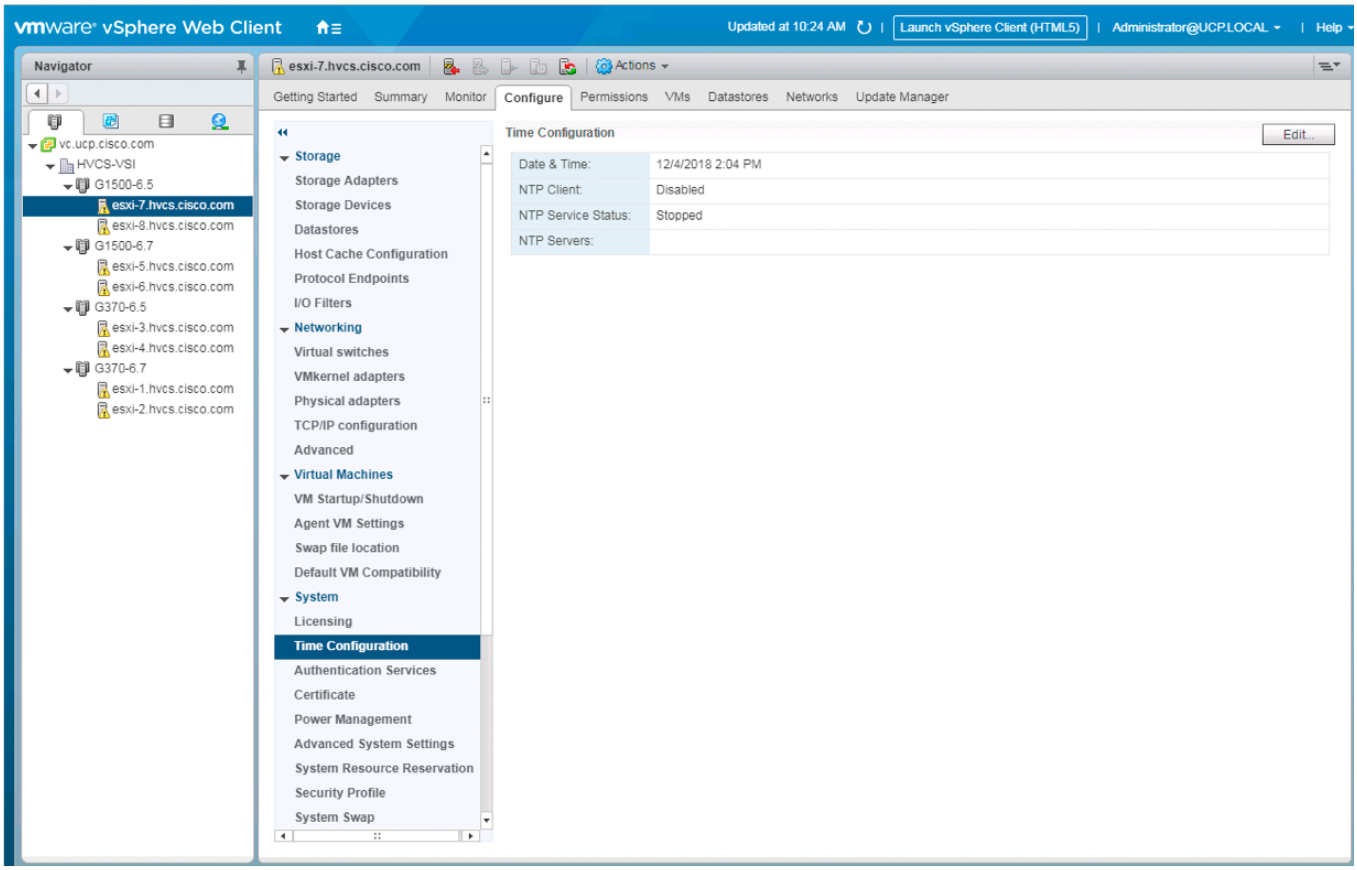

2. Click Edit.

3. Select the Use Network Time Protocol (Enable NTP client) option.

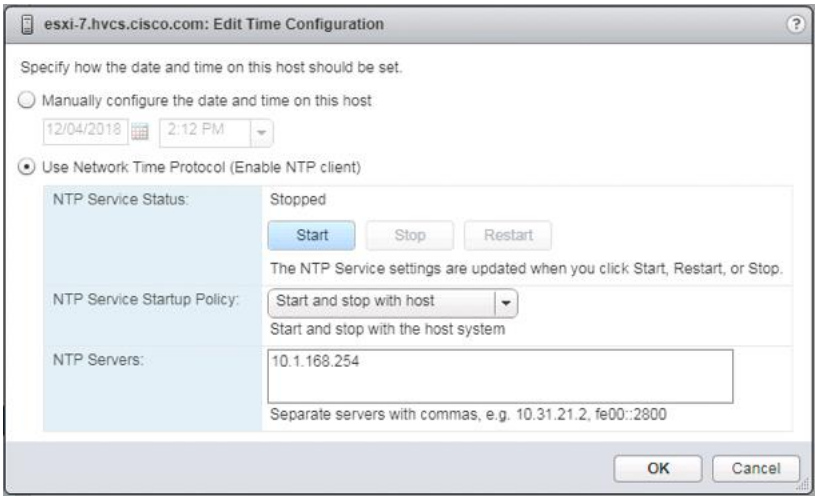

4. Enter an appropriate NTP server within the NTP Servers box, change NTP Service Startup Policy to Start and stop with host, and click Start.

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.

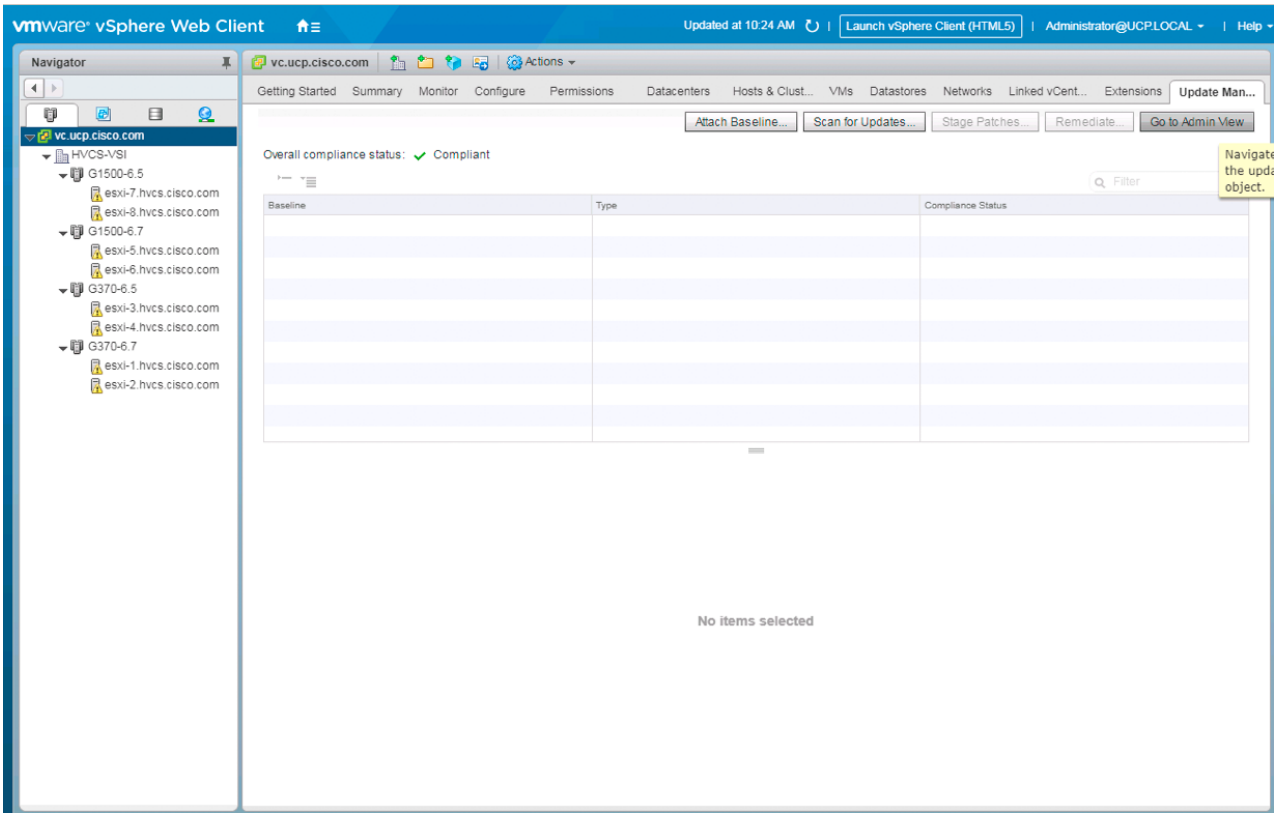

3. From the Manage tab select Hosts Baselines.

4. Click +New Baseline…

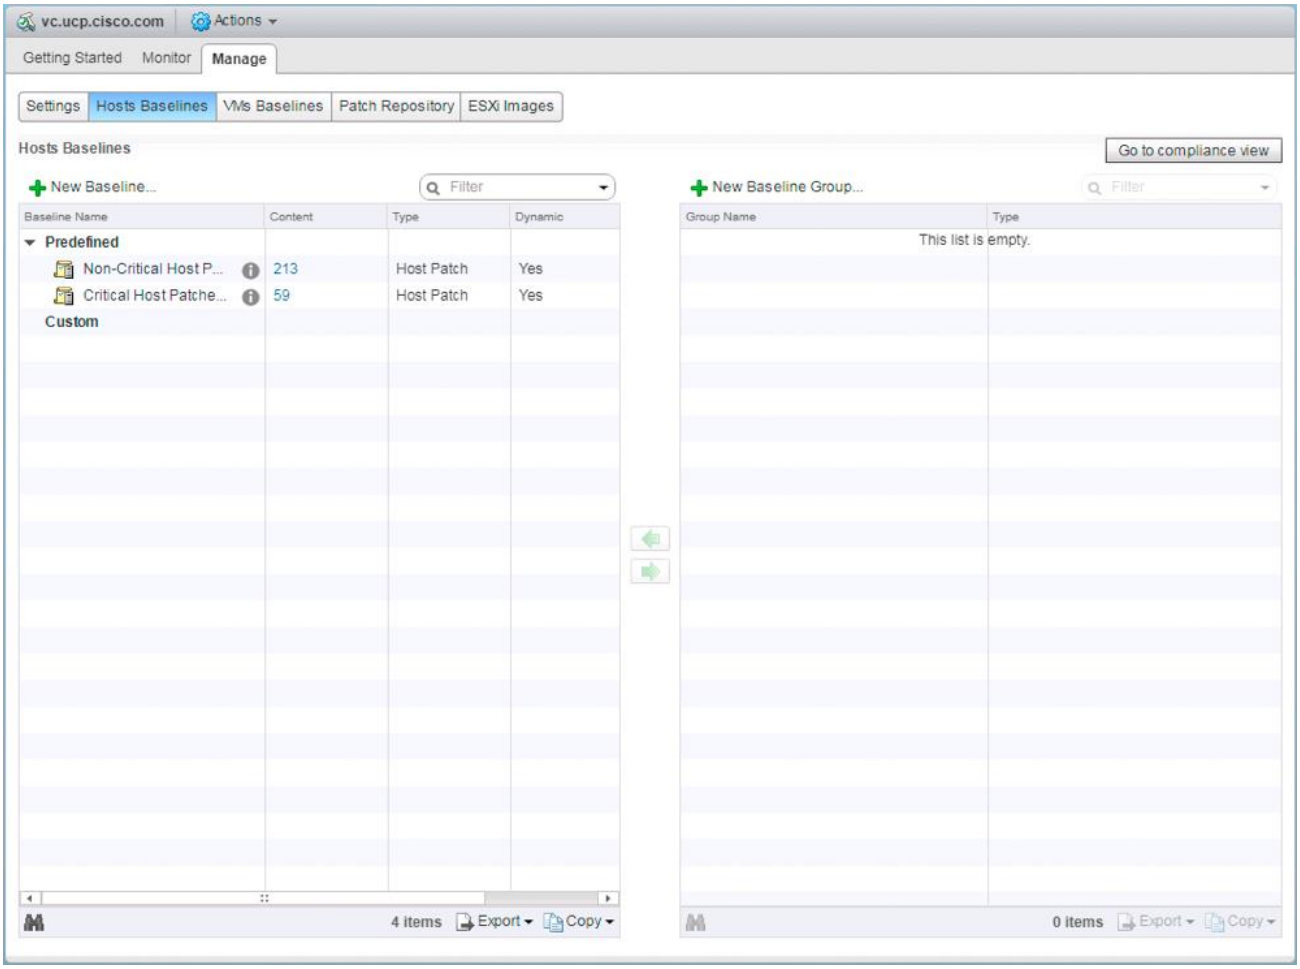

5. Provide a name for the Baseline, leave the Baseline type selected as Host Patch, and click Next.

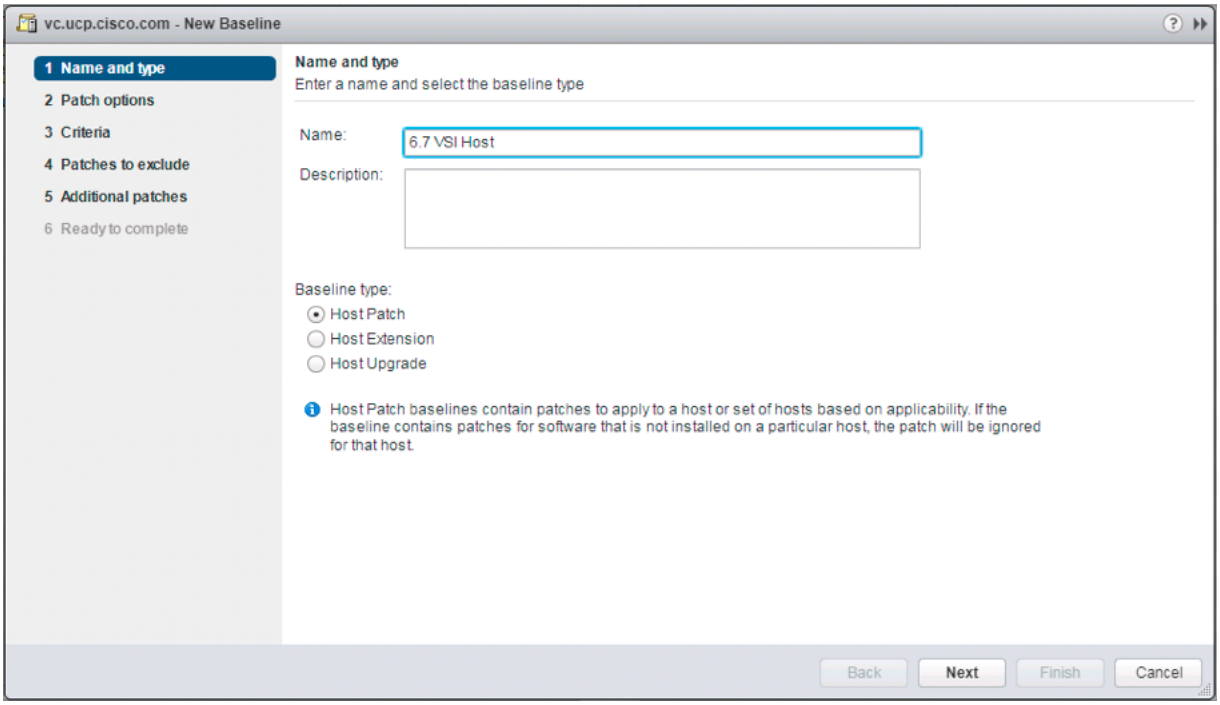

6. Leave the Patch options set as Dynamic and click Next.

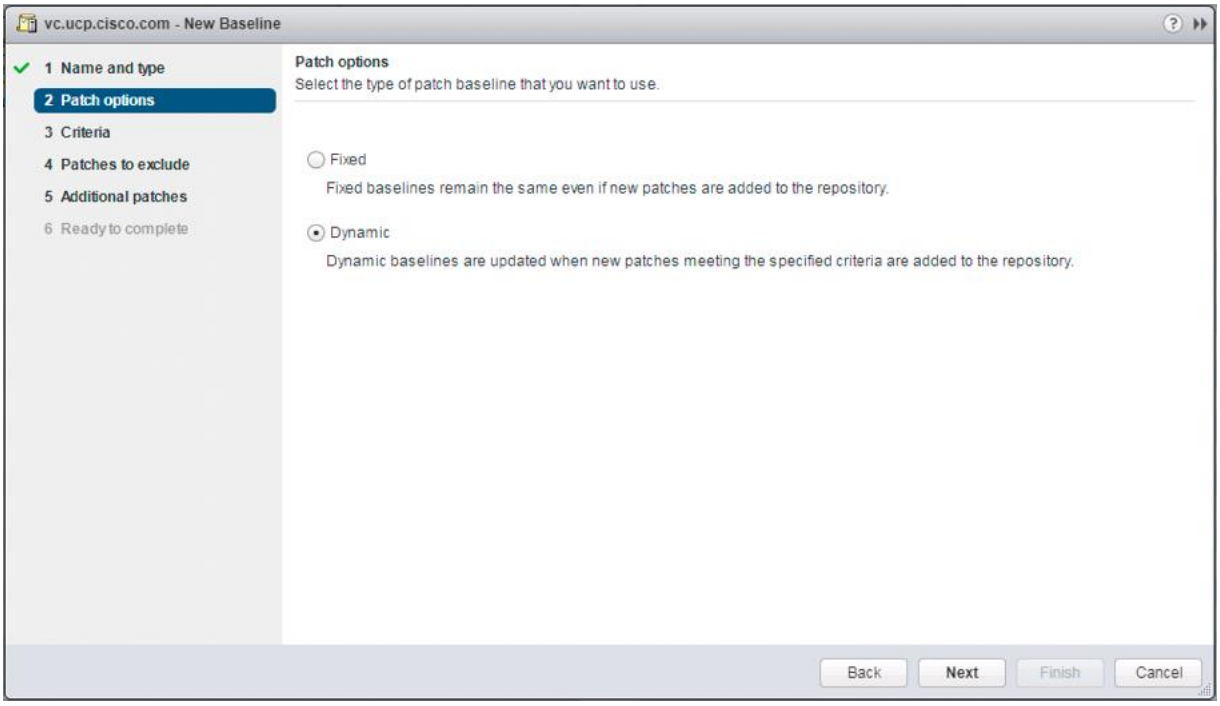

7. Under Product, select the target vSphere release, and under Severity select Critical.

8. Click Next.

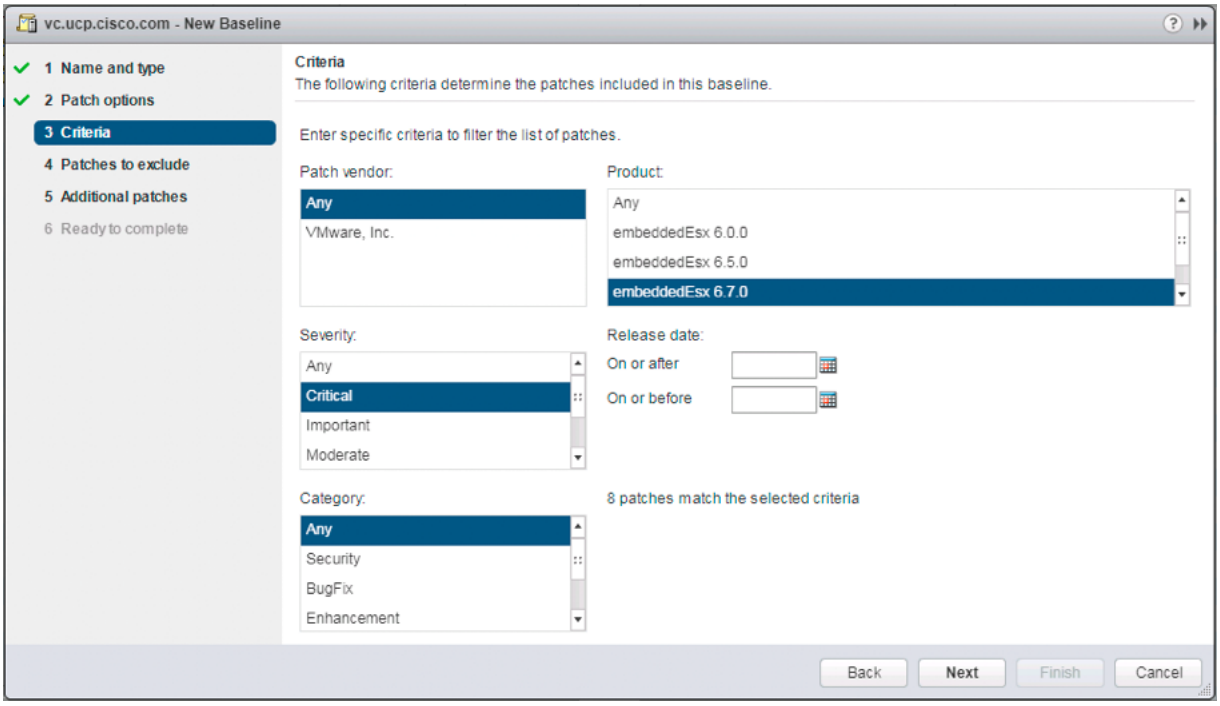

9. Exclude any patches if appropriate and click Next.

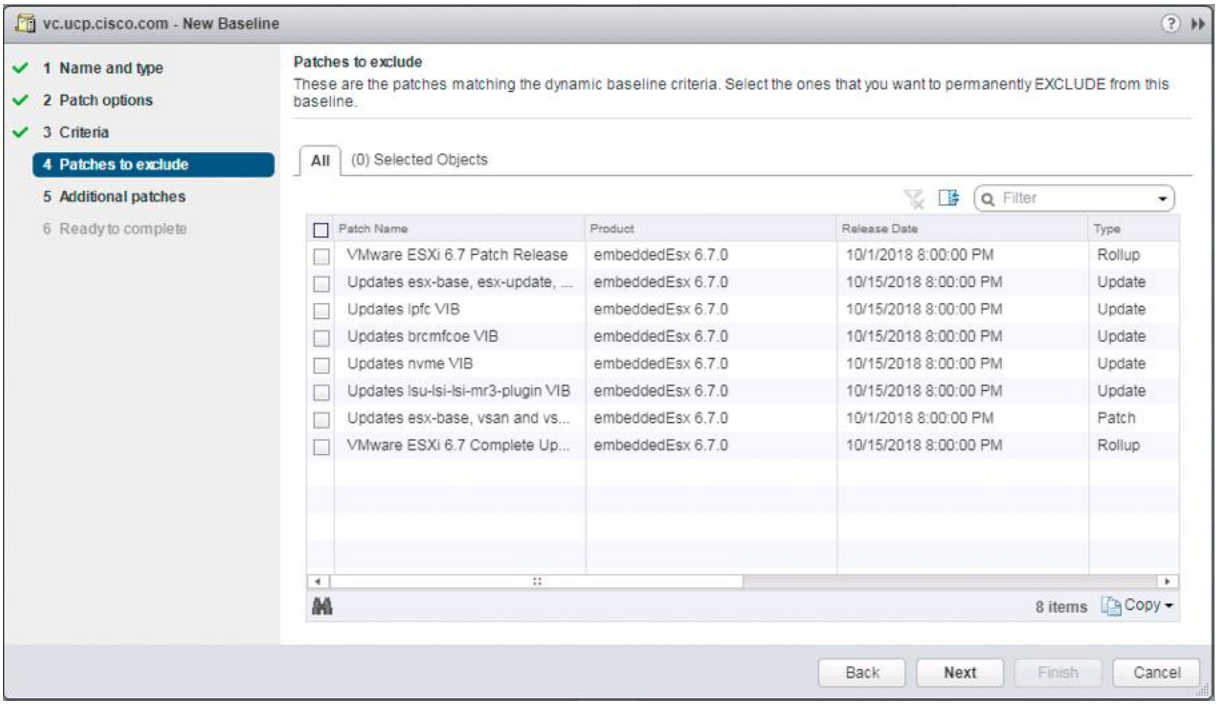

10. Select any additional patches if appropriate and click Next.

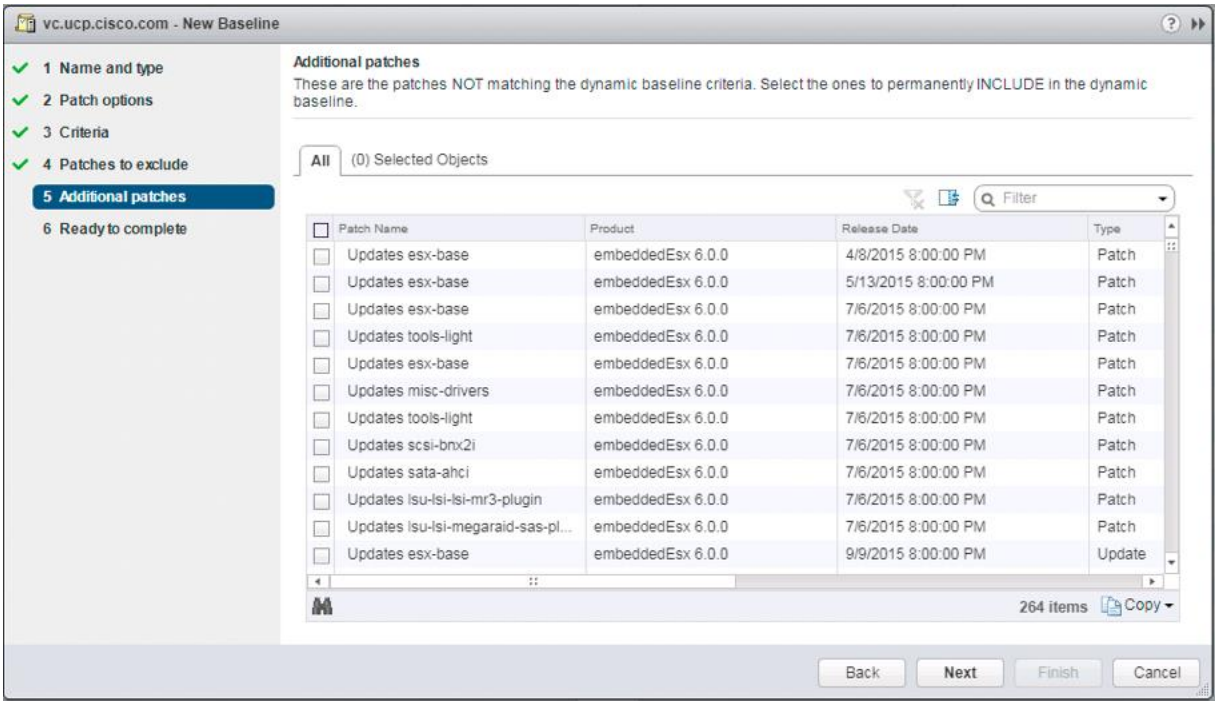

11. Review the selections and click Finish.

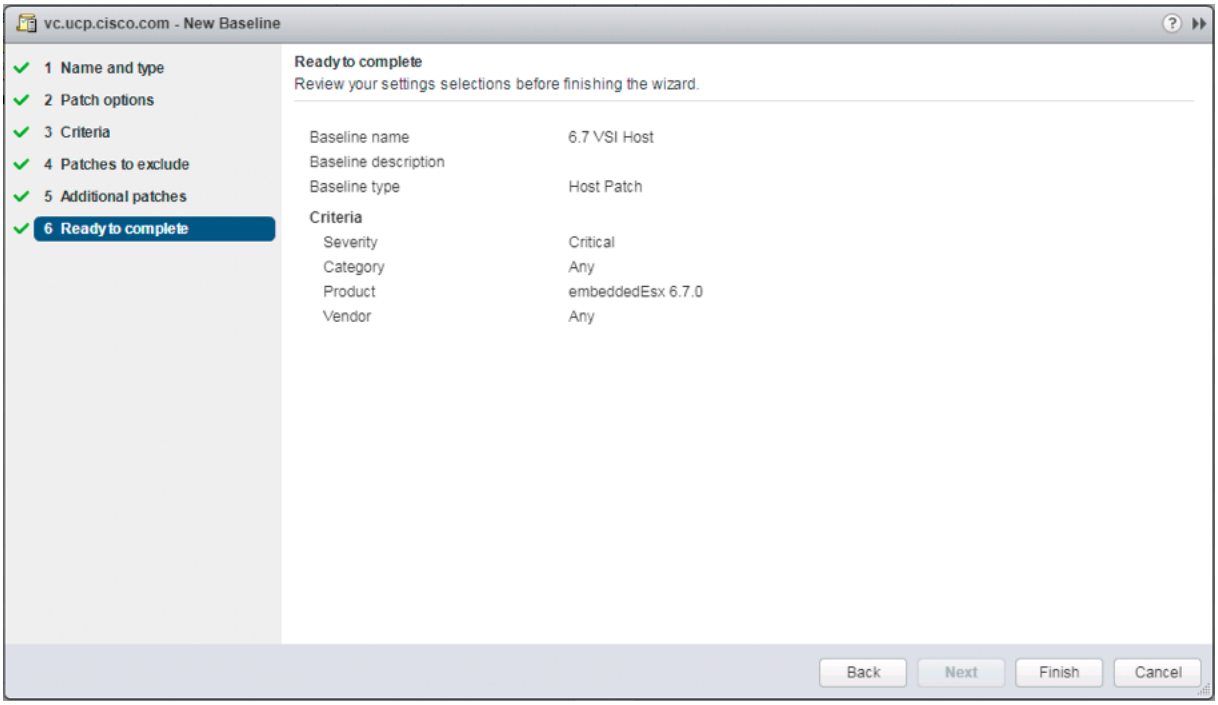

12. Go back to the Hosts view within Navigator.

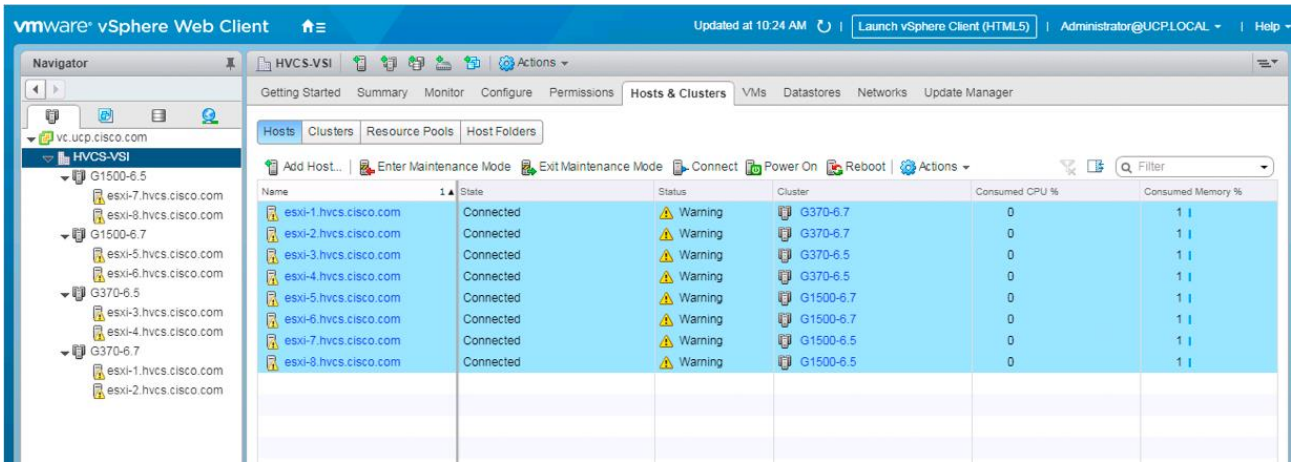

13. Select the Datacenter level of HVCS-VSI and within Hosts of the Hosts and Clusters tab select all hosts and click Enter Maintenance Mode.

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….

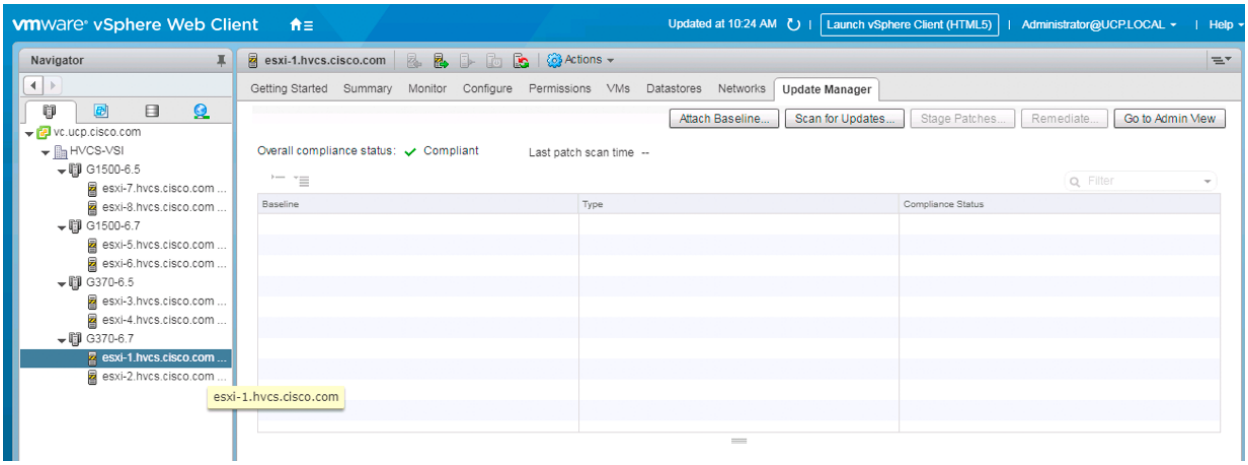

16. Select the appropriate Patch Baseline and click OK.

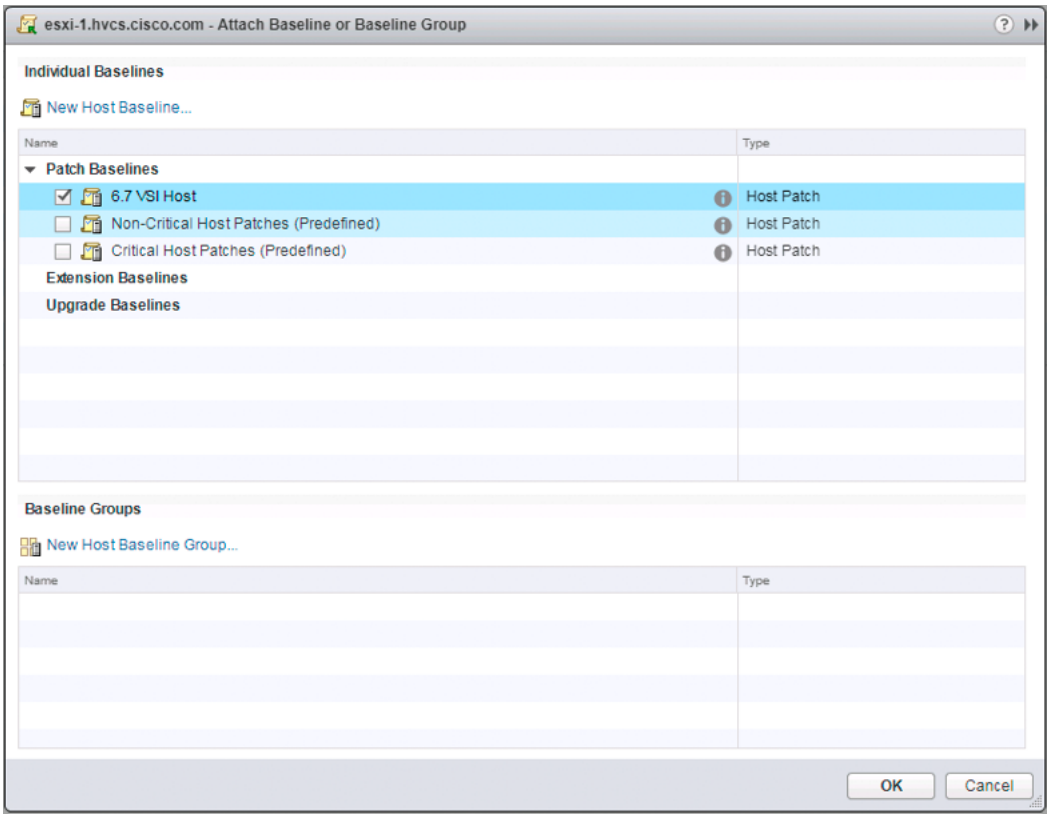

17. Click Remediate.

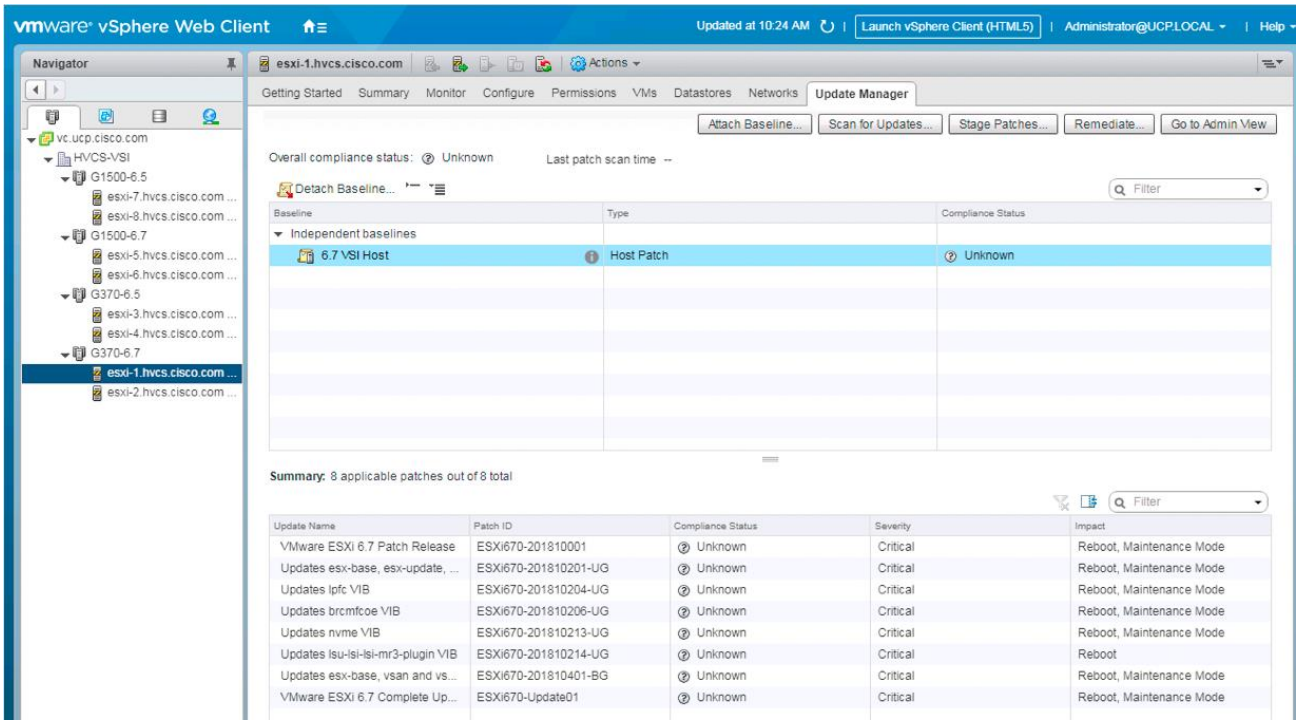

18. Leave the baseline selected and click Next.

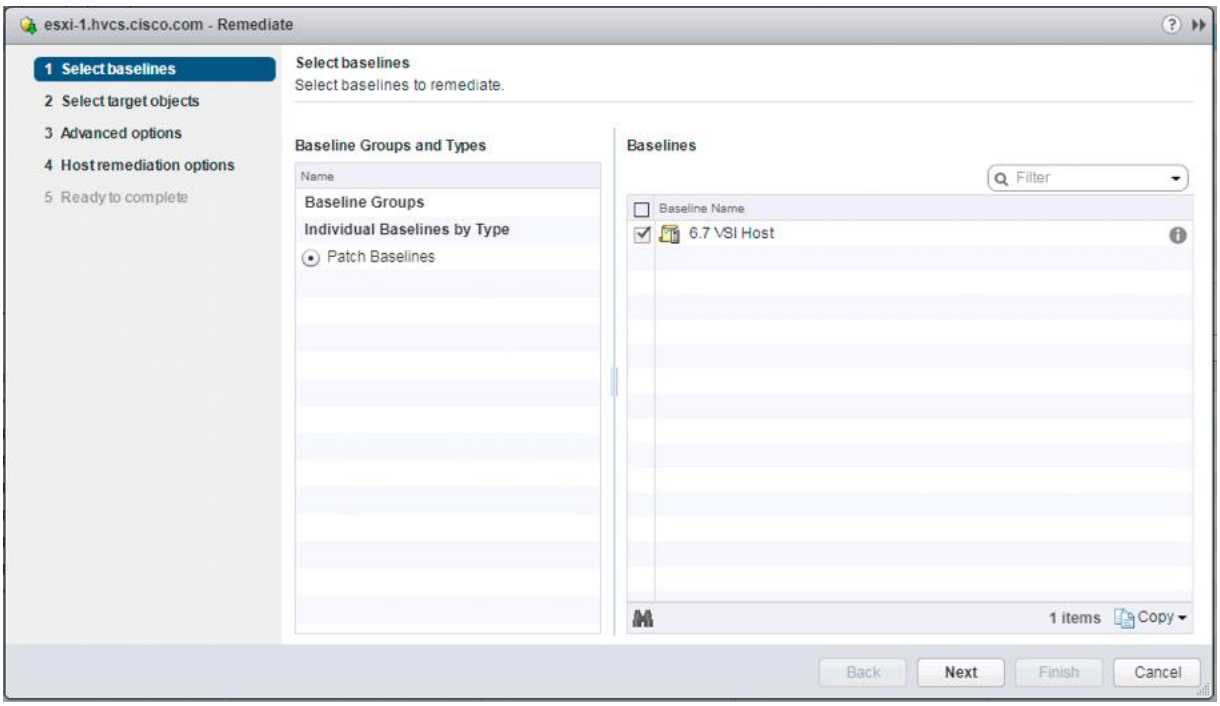

19. Select the hosts appropriate to the vSphere release specified for the baseline and click Next.

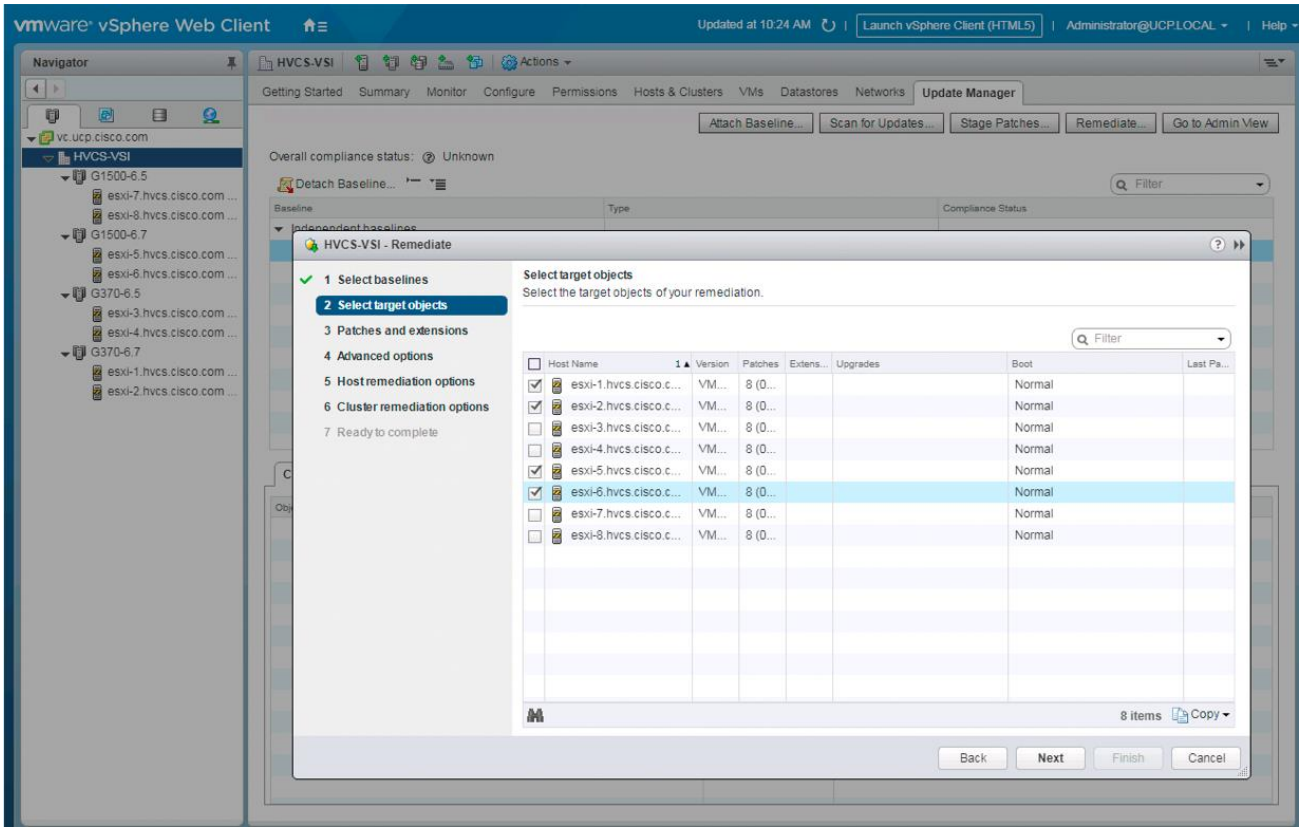

20. Deselect any patches that should not be applied and click Next.

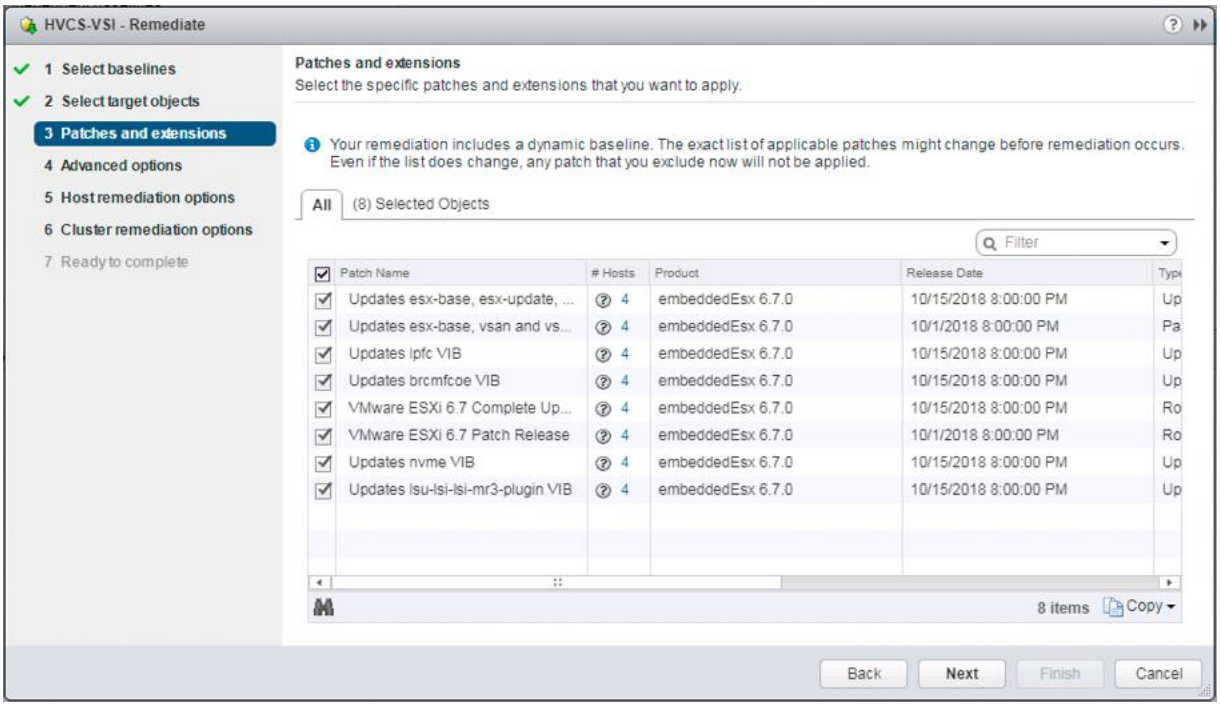

21. Click Next past the Advanced options screen.

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.

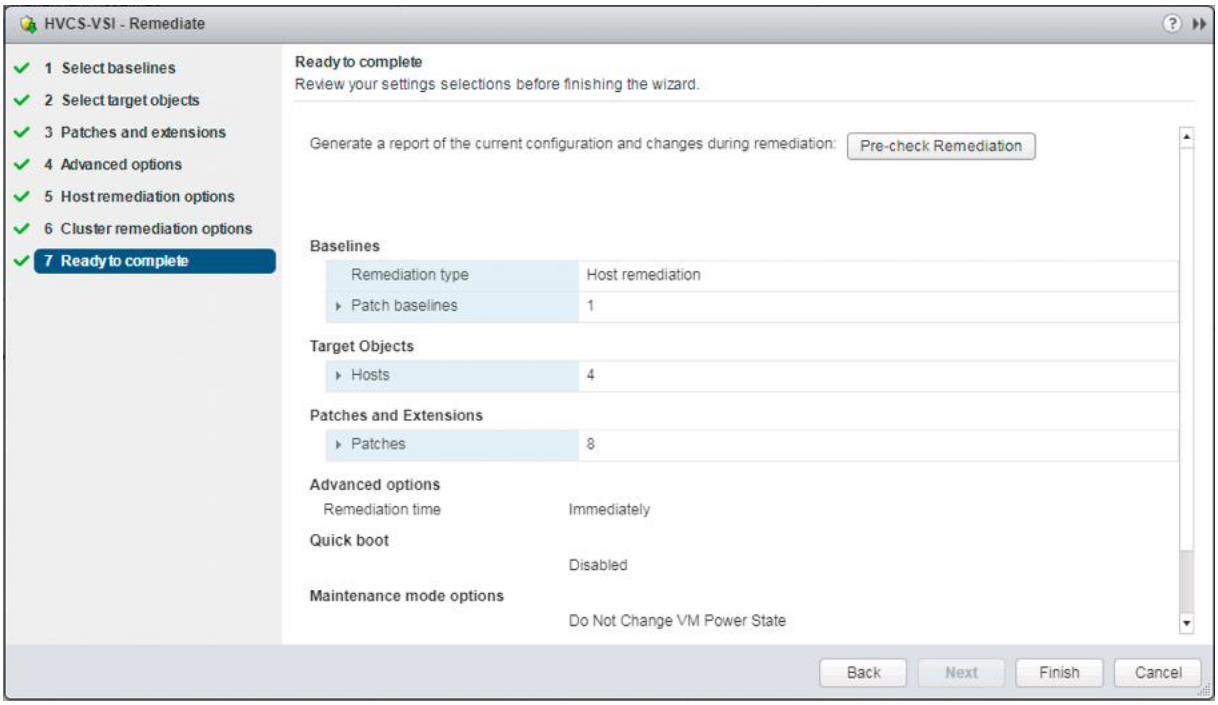

Remediation of L1 Terminal Fault – VMM (L1TF) Security Vulnerability (Optional)

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. More information about the specific impact and VMware's recommendations for remediation of these vulnerabilities in a VMware vSphere environment can be found at 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.

1. Planning and Utilization of the HTAware Mitigation Tool (if analyzing existing workloads).

2. Enablement of the ESXi Side-Channel-Aware Scheduler.

Updating VMware vCenter and VMware ESXi Software

VMware vCenter must be running at specific patch levels prior to mitigation of the L1TF-VMM vulnerabilities. Table 32 lists the release version of VMware vCenter and the specific patch level that needs to be running on the vCenter managing the environment.

Table 32 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 33 lists the minimum Cisco ISO version that must be used to ensure no patching of the ESXi servers is necessary.

Table 33 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 excellent article regarding capacity planning considerations and tested performance degradation when the L1TF-VMM vulnerabilities are remediated at 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:

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.

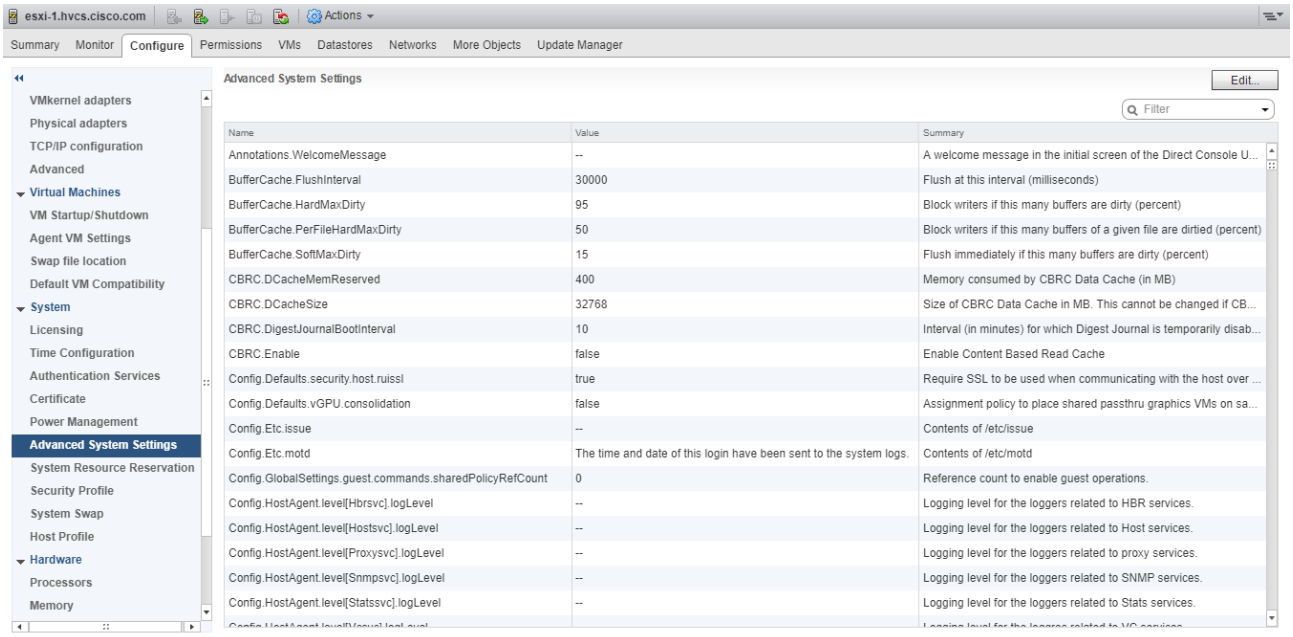

3. Click Edit in the upper right corner of the Advanced System Settings pane, then use the Filter box to search for VMkernel.Boot.hyperthreadingMitigation. Check the Enabled checkbox for the VMkernel.Boot.hyperthreadingMitigation system setting and click OK.

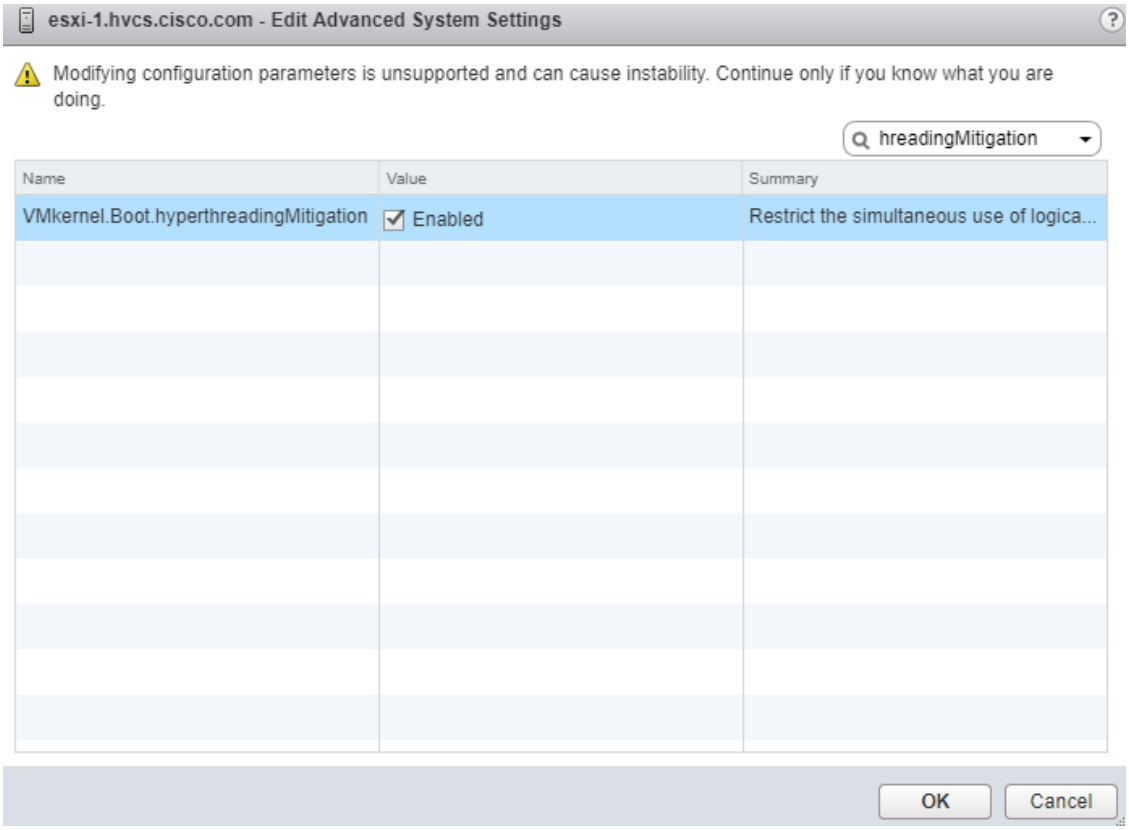

4. Reboot the host and exit Maintenance Mode. The warning message for CVE-2018-3646 should no longer appear in the host Summary tab.

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.

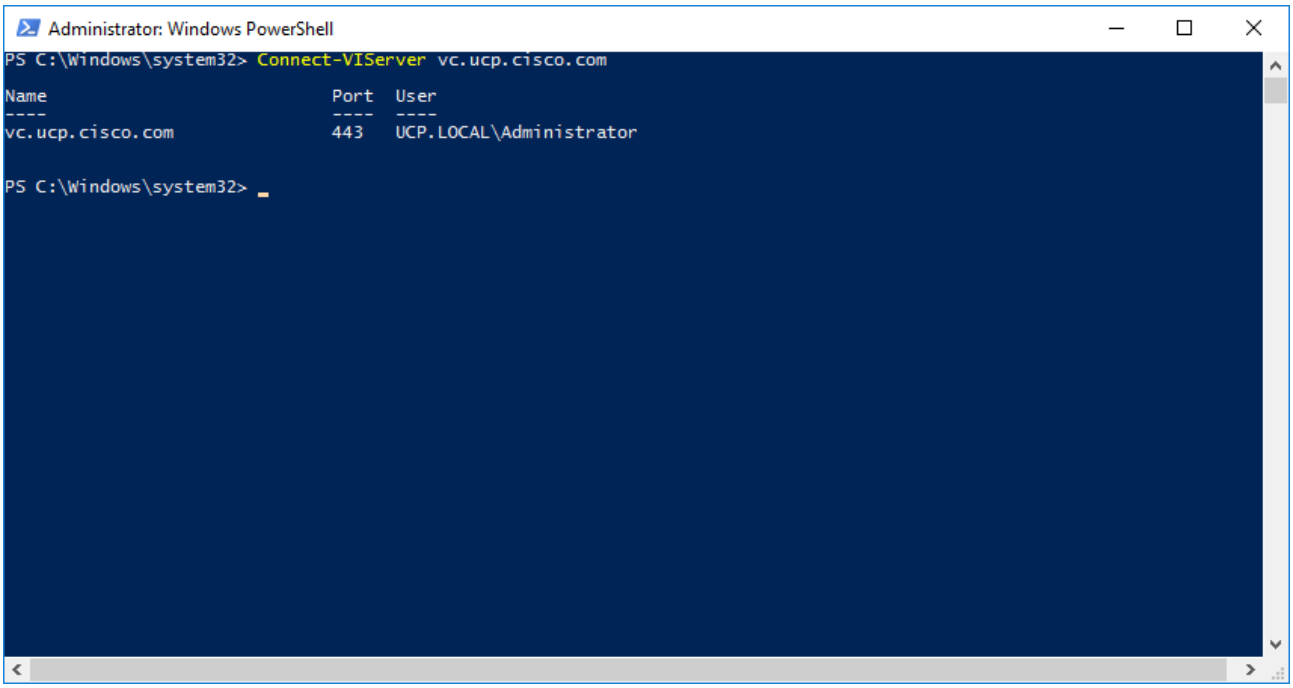

4. Query the remediation status of the hosts in a specific cluster by using the "Get-HTAwareMitigationConfig -ClusterName <name of cluster>" cmdlet. ConfiguredHTAMSetting and RuntimeHTAMSetting should both be false on non-remediated hosts.

```
Administrator: Windows PowerShell
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.

```
Administrator: Windows PowerShell
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 on them. 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% 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 via 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.

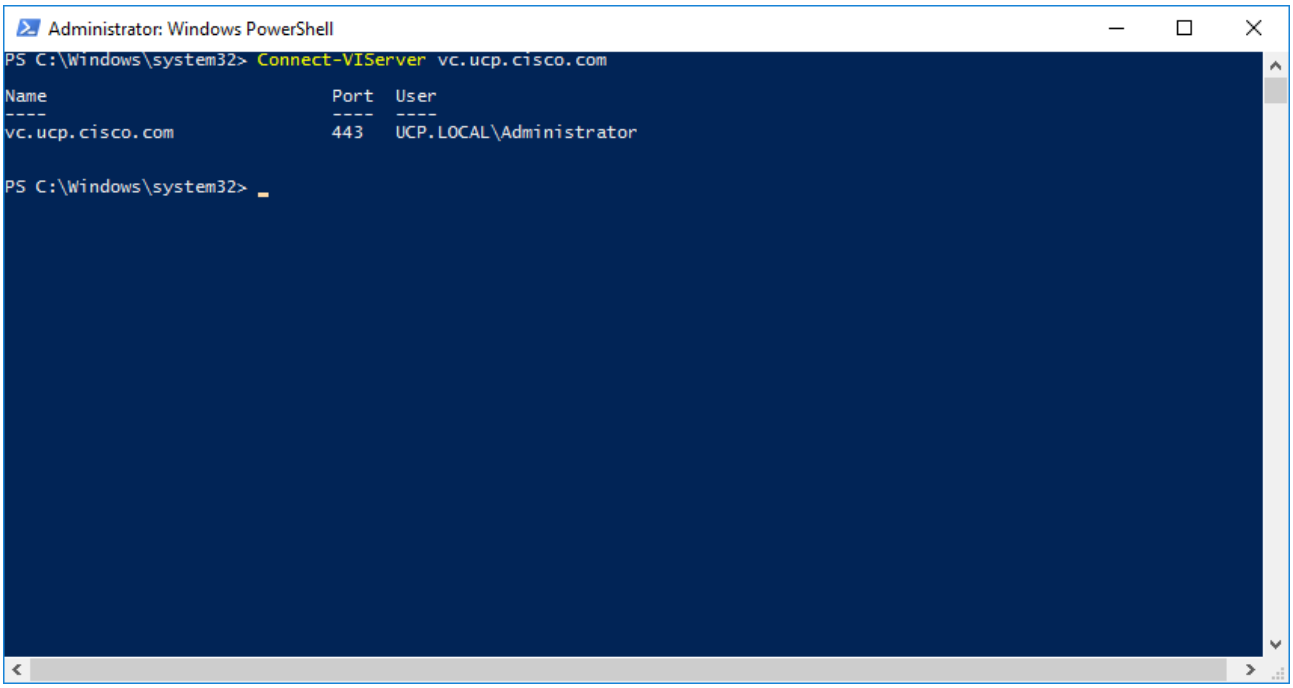

2. Create a variable which will contain all of the ESXi hosts within the cluster you are targeting. In the example shown below, we are creating a variable named “UCSHosts” which contains the host objects within the cluster G1500-6.7 by running the PowerShell command “$UCSHosts = Get-Cluster “G1500-6.7” | Get-VMHost”. Replace G1500-6.7 with the name of the cluster you are targeting.

```
Administrator: Windows PowerShell
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.60060e8007562400003056240000024      RoundRobin                 1000
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000025      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.60060e8007562400003056240000032      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000024      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000026      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.60060e8007562400003056240000031      RoundRobin                 1000
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000032      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.

```
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'} | Set-ScsiLun -CommandsToSwitchPath 20 | Selec
t-Object VMHost, CanonicalName, CommandsToSwitchPath}

VMHost                CanonicalName                        CommandsToSwitchPath
------                -------------                        --------------------
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000024                   20
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000025                   20
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000002b                   20
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000002d                   20
esxi-5.hvcs.cisco.com naa.60060e8007562400003056240000032                   20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000024                   20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000026                   20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000002b                   20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000002d                   20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000031                   20
esxi-6.hvcs.cisco.com naa.60060e8007562400003056240000032                   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 via 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.

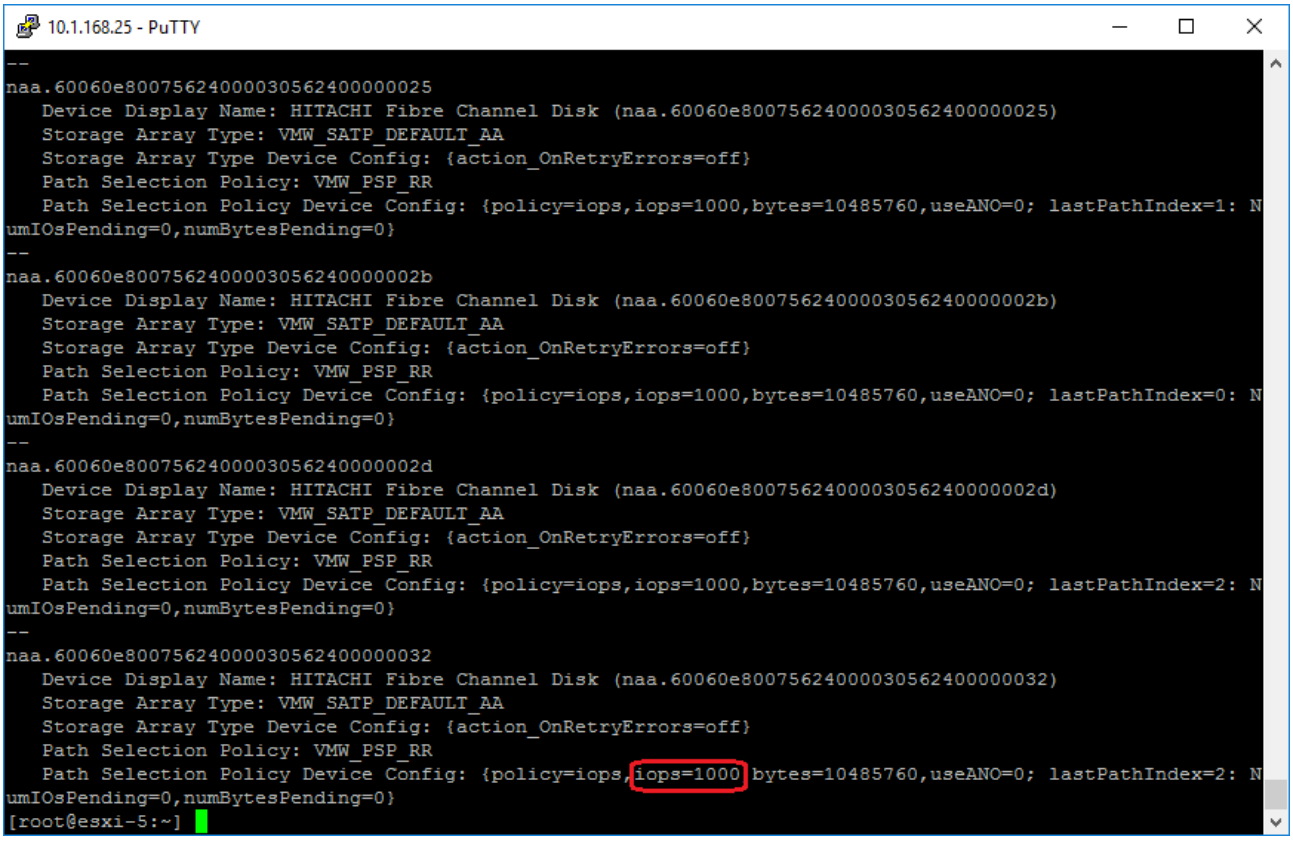

3. To set the RR PSP IOPS limit to 20 on all Hitachi LUNs presented to the host, you may issue the command "for i in `esxcfg-scsidevs -c | awk '{print $1}' | grep naa.60060e80`; do esxcli storage nmp psp roundrobin deviceconfig set --type=iops --iops=20 --device=$i; done" as shown in the example shown below.

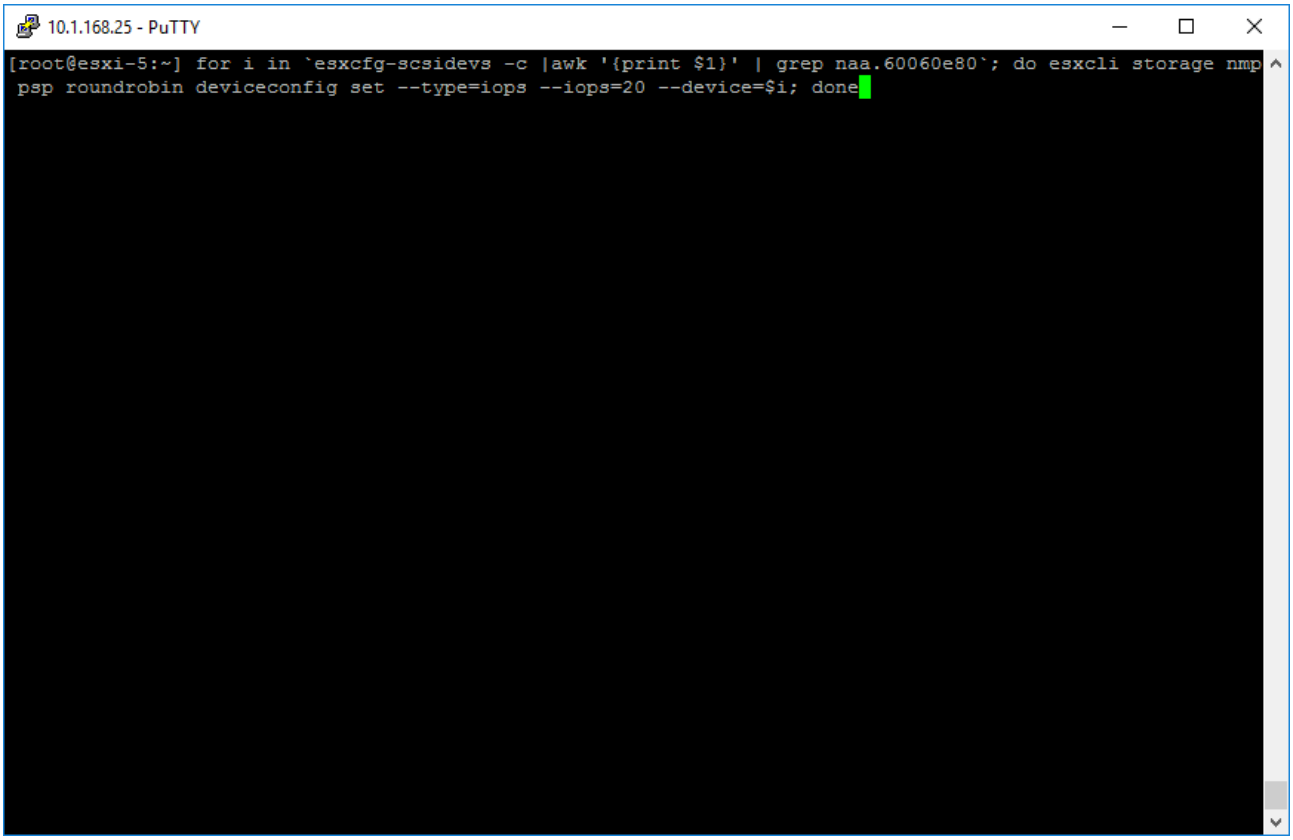

4. Run the command “esxcli storage nmp device list | grep HITACHI -A4 -B1” and ensure the IOPS limit has changed to 20 for Hitachi-presented LUNs as shown in the example below.

```
10.1.168.25 - PuTTY
--
naa.60060e8007562400003056240000025
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e8007562400003056240000025)
   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.60060e8007562400003056240000032
   Device Display Name: HITACHI Fibre Channel Disk (naa.60060e8007562400003056240000032)
   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.

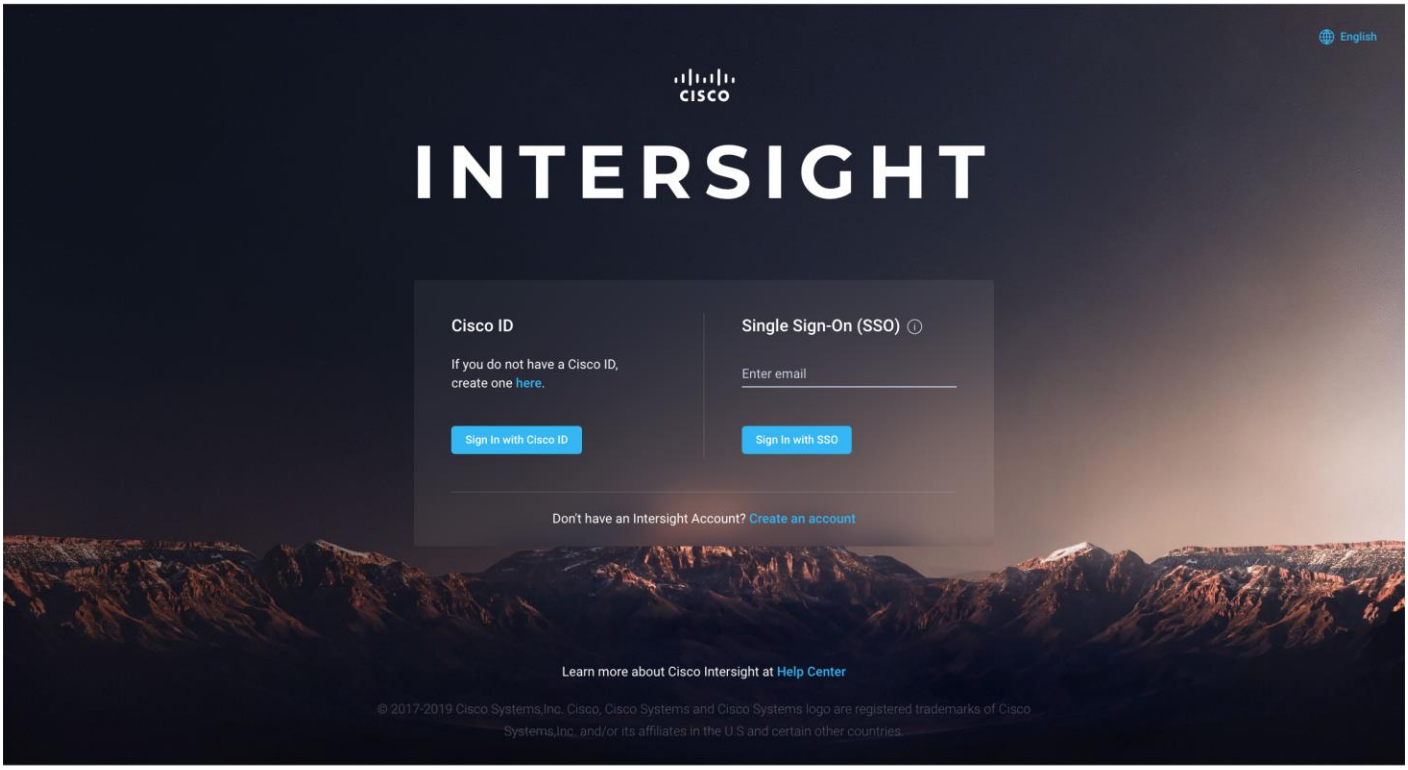

Prerequisites

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 Cisco Intersight Licensing section below 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/.

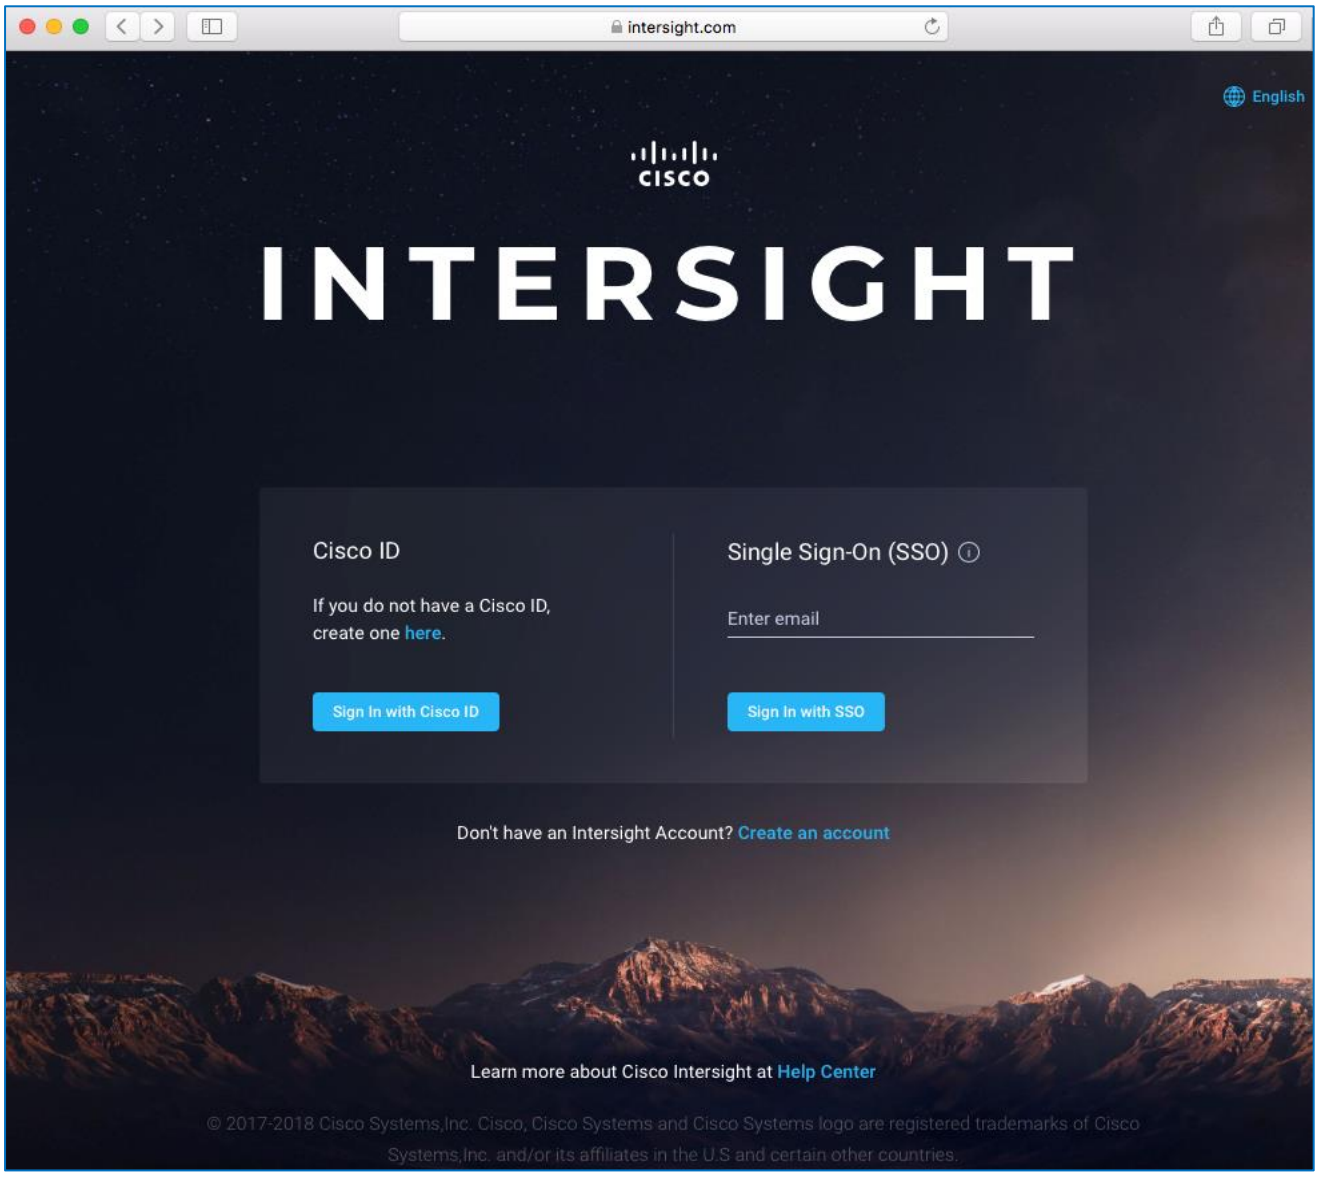

2. Login with a valid cisco.com account or single sign-on using your corporate authentication.

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 left navigation menu, select the Admin icon.
3. From the left navigation pane, select All > Device Connector.
4. In the right window pane, for Intersight Management, click Enabled to enable Intersight management.

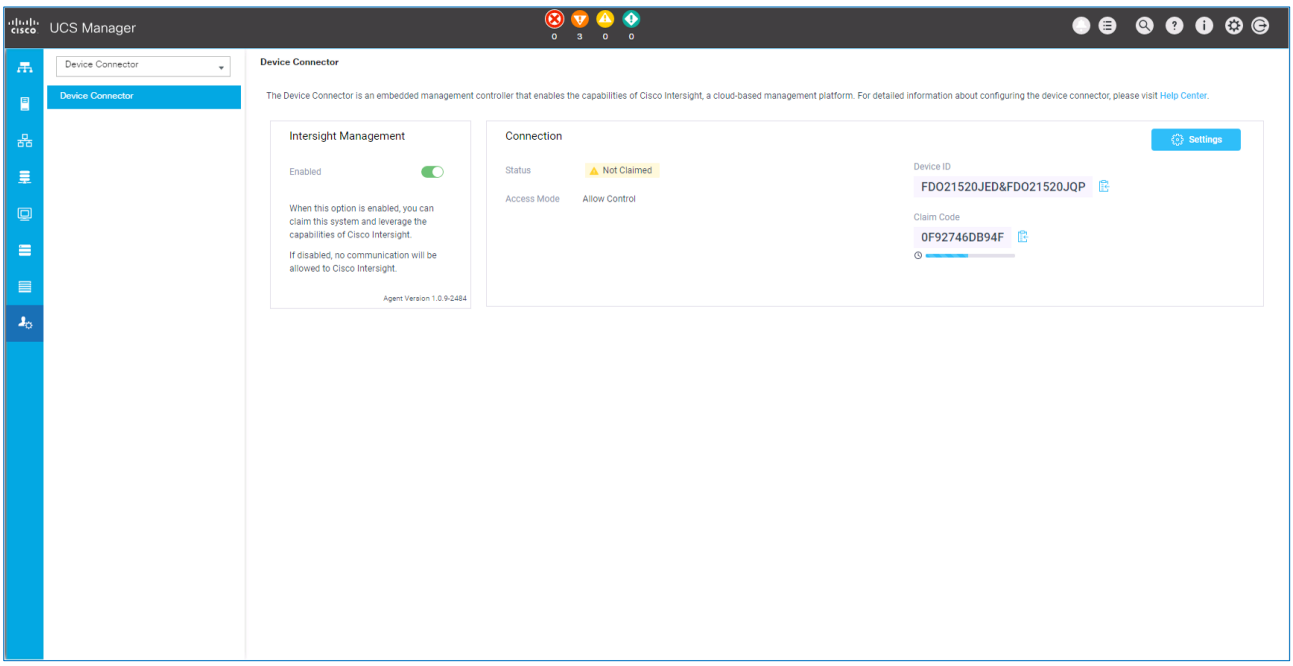

5. From the Connection section, copy the Device ID and Claim ID information. This information will be required to add this device to Cisco Intersight.

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 left navigation menu, select Devices.

2. Click the Claim a New Device button in the top right-hand corner.

3. In the Claim a New Device pop-up window, paste the Device ID and Claim Code collected in the previous section.

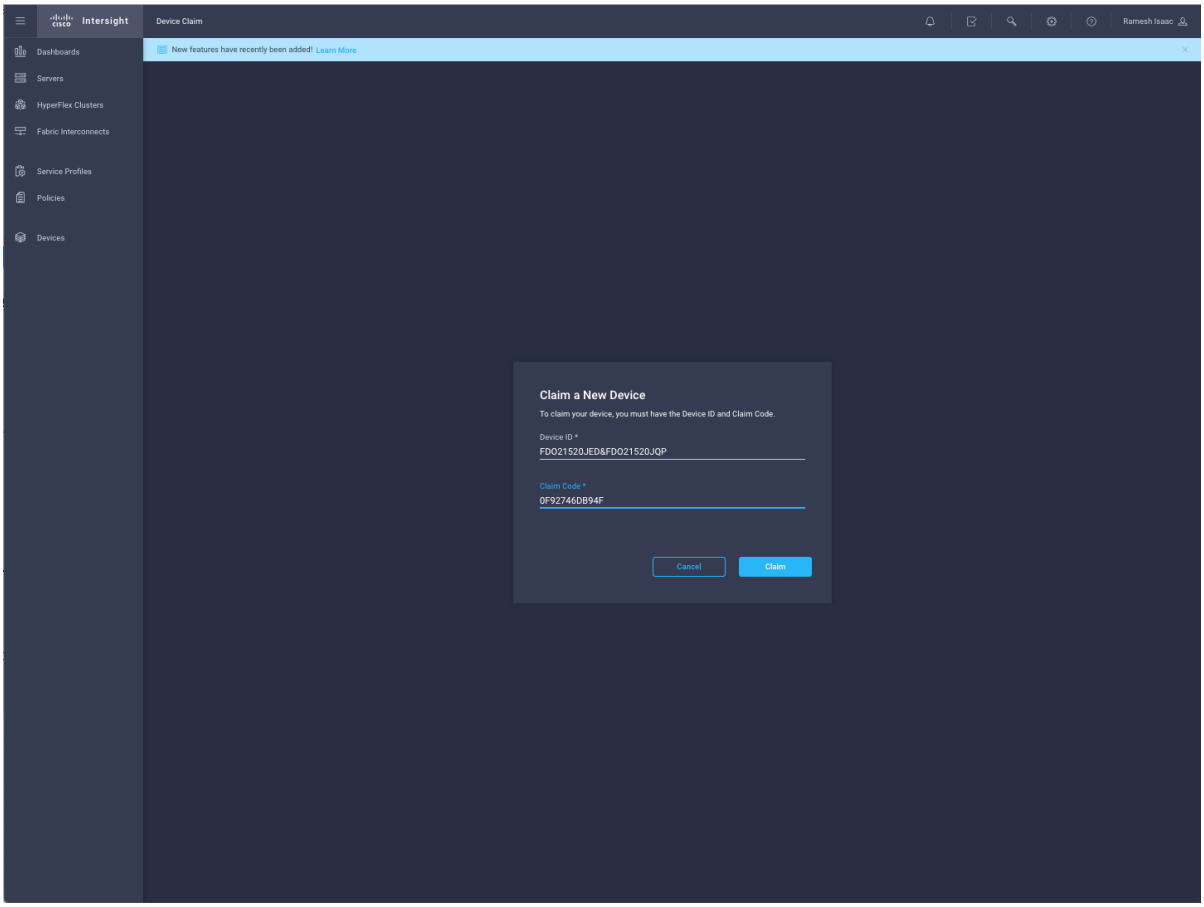

4. Click Claim.
5. On Cisco Intersight, the newly added UCS domain should now have a Status of Connected.
6. On Cisco UCS Manager, the Device Connector should now have a Status of Claimed.
7. Dashboard will present an overview of the managed UCS domains:

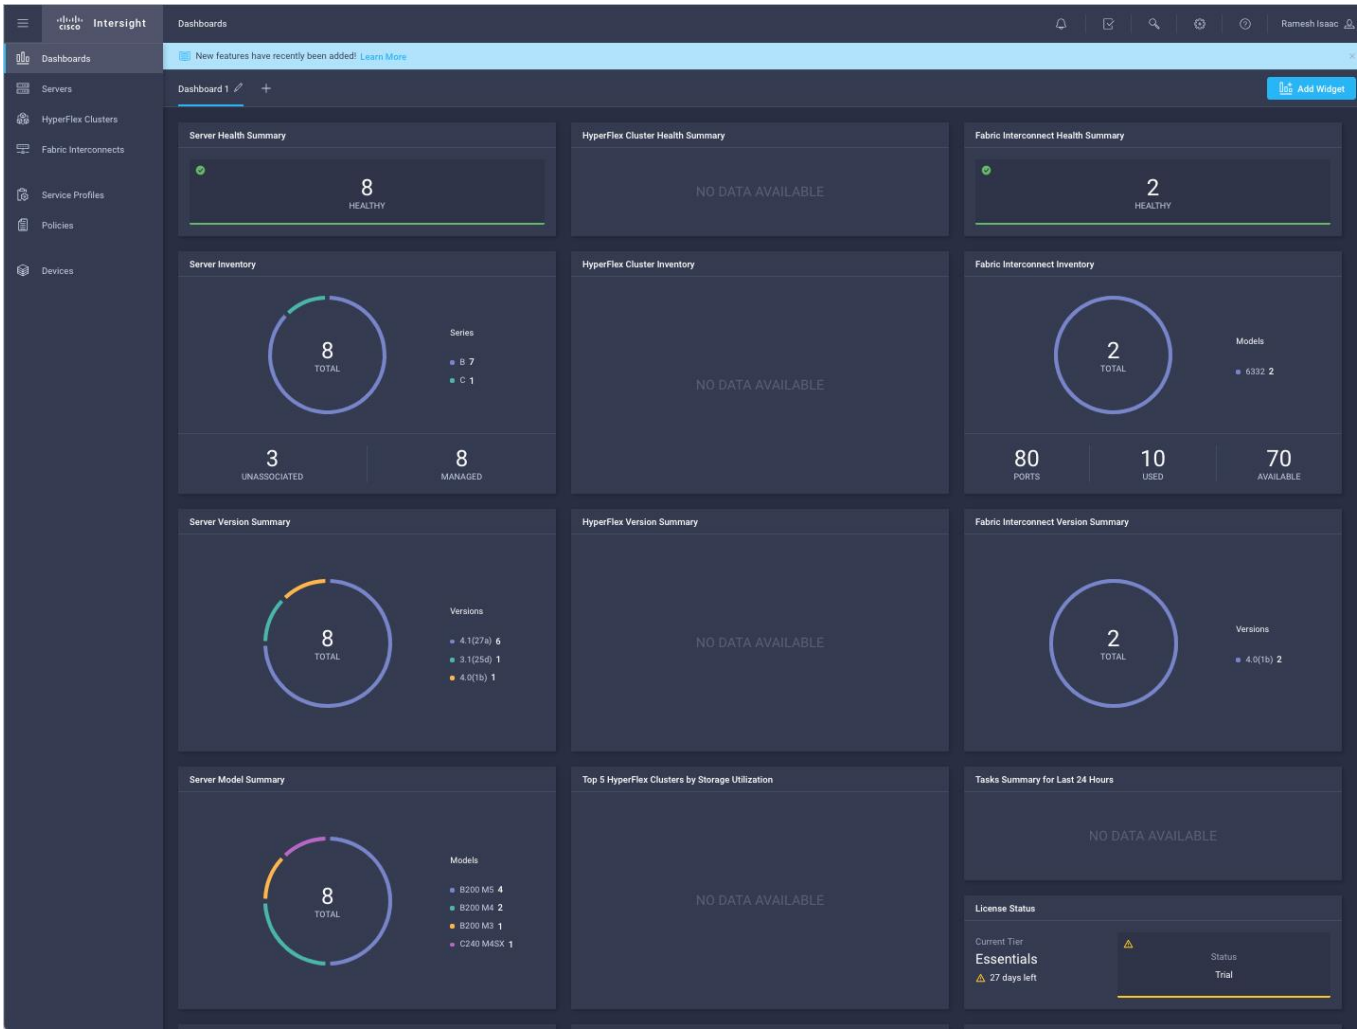
Intersight
Dashboards
Ramesh Isaac
Dashboards
Servers
HyperFlex Clusters
Fabric Interconnects
Service Profiles
Policies
Devices
New features have recently been added! Learn More
Dashboard 1
Add Widget
Server Health Summary
8
HEALTHY
HyperFlex Cluster Health Summary
NO DATA AVAILABLE
Fabric Interconnect 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(2a) 6
3.1(2d) 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

MDS Best Practices: https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white-paper-c11-738426.html

Cisco MDS 9000 Series Interfaces Configuration Guide, Release: 8.xhttps://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/interfaces/cisco_mds9000_interfaces_config_guide_8x.html

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 Cisco MDS switches
- Multi-mode Fibre (MMF) cabling between Cisco MDS switches 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 34 lists the BOM for Cisco UCS 6454 Fabric Interconnect with Hitachi VSP G370 design.

Table 34 Bill of Materials for 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

Vendor	Part Number/Order Code	Description	Quantity
Cisco	UCS-PSU-6332-AC	UCS 6332 Power Supply/100-240VAC	4
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-C9706	MDS 9706 Chassis No Power Supplies, Fans Included	2

Vendor	Part Number/Order Code	Description	Quantity
Cisco	DS-X9706-FAB1	MDS 9706 Crossbar Switching Fabric-1 Module	12
Cisco	DS-9706-KIT-CCO	MDS 9706 Accessory Kit for Cisco	2
Cisco	DS-C9700-LC-BL	Blank Filler Card for Line Card slot in MDS9700 Chassis	6
Cisco	DS-X97-SF1-K9	MDS 9700 Series Supervisor-1	4
Cisco	DS-X9648-1536K9	MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module	2
Cisco	DS-CAC97-3KW	MDS 9700 3000W AC power supply	8
Cisco	M97S3K9-8.2.1	MDS 9700 Supervisor/Fabric-3, NX-OS Software Release 8.2(1)	2
Cisco	CAB-9K16A-US2	Power Cord 250VAC 16A, US/Japan, Src Plug NEMA L6-20	8
Cisco	DS-SFP-FC32G-SW	32 Gbps Fibre Channel SW SFP+, LC	12
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
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

Table 35 lists the BOM for the Cisco UCS 6332-16UP Fabric Interconnect with Hitachi VSP G1500 design.

Table 35 Bill of Materials for UCS 6332-16UP and Hitachi VSP G1500 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

Vendor	Part Number/Order Code	Description	Quantity
Cisco	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	4
Cisco	QSFP-40G-SR-BD	QSFP40G BiDi Short-reach Transceiver	4
Cisco	QSFP-40G-CSR4	QSFP 4x10GBASE-SR Transceiver Module, MPO, 300M	4
Cisco	QSFP-H40G-ACU7M	40GBASE-CR4 Active Copper Cable, 7m	4
Cisco	QSFP-H40G-CU1M	40GBASE-CR4 Passive Copper Cable, 1m	2
Cisco	UCS-FI-6332-16UP-U	UCS 6332-16UP 1RU FI/No PSU/24 QSFP+ 16UP/4x40G Lic/8xUP Lic	2
Cisco	N10-MGT016	UCS Manager v4.0	2
Cisco	QSFP-H40G-CU3M	40GBASE-CR4 Passive Copper Cable, 3m	4
Cisco	UCS-ACC-6332	UCS 6332/ 6454 Chassis Accessory Kit	2
Cisco	UCS-FAN-6332	UCS 6332/ 6454 Fan Module	8
Cisco	DS-SFP-FC16G-SW	16 Gbps Fibre Channel SW SFP+, LC	8
Cisco	UCS-PSU-6332-AC	UCS 6332 Power Supply/100-240VAC	4
Cisco	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	4
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	UCS-MR-X16G1RS-H	16GB DDR4-2666-MHz RDIMM/PC4-21300/single rank/x4/1.2v	96
Cisco	UCSB-MLOM-40G-04	Cisco UCS VIC 1440 modular LOM for Blade Servers	4
Cisco	UCSB-MLOM-PT-01	Cisco UCS Port Expander Card (mezz) for VIC	4
Cisco	UCS-SID-INFR-OI	Other Infrastructure	4

Vendor	Part Number/Order Code	Description	Quantity
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-2304	UCS 2304XP I/O Module (4 External, 8 Internal 40Gb Ports)	2
Cisco	UCSB-PSU-2500ACDV	2500W Platinum AC Hot Plug Power Supply - DV	4
Cisco	CAB-C19-CBN	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	4
Cisco	UCSB-5108-PKG-HW	UCS 5108 Packaging for chassis with half width blades.	1
Cisco	DS-C9706	MDS 9706 Chassis No Power Supplies, Fans Included	2
Cisco	CON-SNT-C9706	SNTC-8X5XNBD MDS 9706 Chassis No Power Supplies, Fans	2
Cisco	CON-SNT-97FA0	SNTC-8X5XNBD MDS 9706 Crossbar Sw	6
Cisco	DS-9706-KIT-CCO	MDS 9706 Accessory Kit for Cisco	2
Cisco	DS-C9700-LC-BL	Blank Filler Card for Line Card slot in MDS9700 Chassis	6
Cisco	DS-X97-SF1-K9	MDS 9700 Series Supervisor-1	4
Cisco	DS-X9648-1536K9	MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module	2
Cisco	DS-SFP-FC16G-SW	16 Gbps Fibre Channel SW SFP+, LC	24
Cisco	DS-CAC97-3KW	MDS 9700 3000W AC power supply	8
Cisco	CAB-9K16A-US2	Power Cord 250VAC 16A, US/Japan, Src Plug NEMA L6-20	8
Cisco	M97S3K9-8.2.1	MDS 9700 Supervisor/Fabric-3, NX-OS Software Release 8.2(1)	2
Hitachi	VSPG1500-SOLUTION.S	Virtual Storage Platform G1500	1
Hitachi	VSPG1500-A0001.S	VSP G1500 Hardware Product	1
Hitachi	DTI4GL.P	4GB USB memory stick with lanyard	1
Hitachi	IP0662-14.P	LAN Cable 14ft	1
Hitachi	IP0665-45.P	RJ-45 Modular In-Line Coupler 6 Conductor	1

Vendor	Part Number/Order Code	Description	Quantity
Hitachi	041-100108-01.P	Virtual Storage Platform G1x00 Microcode Kit	1
Hitachi	DKC810I-CBXE.P	Primary Controller Chassis	1
Hitachi	DKC-F810I-BKML.P	Cache Backup Module Kit for Large Memory	4
Hitachi	DKC-F810I-BMM256.P	Cache Flash Memory (256GB)	8
Hitachi	DKC-F810I-CC2.P	DKU Interface Copper Cable 2m (DKC/DKU to DKU)	4
Hitachi	DKC-F810I-CM32G.P	Cache Memory Module (32GB)	36
Hitachi	DKC-F810I-CPEX.P	Cache Path Control Adapter	2
Hitachi	DKC-F810I-FBX.P	Flash Module Drive Chassis	1
Hitachi	DKC-F810I-HUB.P	Hub	2
Hitachi	DKC-F810I-MFC5.P	Inter-controller connecting cable 5m	1
Hitachi	DKC-F810I-MOD5.P	Inter-controller connecting kit 5m	1
Hitachi	DKC-F810I-PFUC.P	Flash Module Chassis Power Cord Kit(Americas/APAC)	1
Hitachi	DKC-F810I-PHUC.P	LFF/SFF Drive Chassis Power Cord Kit(Americas/APAC)	1
Hitachi	DKC-F810I-PLUC.P	Controller Chassis Power Cord Kit (Americas/APAC)	2
Hitachi	DKC-F810I-SBX.P	SFF Drive Chassis	1
Hitachi	DKC-F810I-SCA.P	Standard Back-end Director	4
Hitachi	DKC-F810I-CBXF.P	Secondary Controller Chassis	1
Hitachi	3287294-001.P	Name Plate (HTC-R8)	2
Hitachi	DKC-F810I-16FC16.P	16 port, 16Gbps Fibre Channel Front-end Director	4
Hitachi	DKC-F810I-MP2.P	Virtual Storage Director Pair	6
Hitachi	3292099-001.P	Name Plate (G1500)	2
Hitachi	DKC-F810I-7R0FP.P	7TB High Density Flash Module	12
Hitachi	DKC-F810I-1R9MGM.P	1.9TB SFF Solid State Drive	24
Hitachi	DKC-F810I-1R8JGM.P	1800GB, 10K rpm SFF Disk Drive	24
Hitachi	044-235001-01.P	VSP G1x00 SVOS Media	1
Hitachi	044-235001-01B.P	VSP G1x00 SVOS Base (Incl 1 VSD Pair)	1
Hitachi	044-235001-01BU.P	VSP G1x00 SVOS Add'l VSD Pair	7
Hitachi	044-232355-01.P	Hitachi Command Suite Media	1
Hitachi	044-235037-01.P	VSP G1500 VSD Activation Media Kit	1

Vendor	Part Number/Order Code	Description	Quantity
Hitachi	044-235037-03.P	VSP G1500 VSD Activation License	8
Hitachi	044-235001-080I.P	VSP G1x00 SVOS 80TB Initial Purchase Block	1
Hitachi	044-235001-020.P	VSP G1x00 SVOS 20TB Block	1

Appendix: MDS Device Alias and Zoning through CLI

Create Device Aliases

Using the WWPN target and initiator table created earlier in Table 28 and Table 29 collect the information for the device aliases to be created on each fabric into the following tables:

Table 36 Fabric A Targets and Initiators

	Name	pWWN/WWPN Example Environment (Port Name)	pWWN/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 37 Fabric B Targets and Initiators

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

With the appropriate information collected, proceed to create the device aliases on the MDS fabrics using the following as an example:

Fabric A Device Aliases

```
aa19-9706-1(config)# device-alias database
aa19-9706-1(config-device-alias-db)# device-alias name VSI-G370-1 pwwn 20:00:00:25:B5:54:0A:00
aa19-9706-1(config-device-alias-db)# device-alias name VSI-G370-2 pwwn 20:00:00:25:B5:54:0A:01
aa19-9706-1(config-device-alias-db)# device-alias name G370-CL1-A pwwn 50:06:0E:80:12:C9:9A:00
aa19-9706-1(config-device-alias-db)# device-alias name G370-CL2-B pwwn 50:06:0E:80:12:C9:9A:11
aa19-9706-1(config-device-alias-db)# exit
aa19-9706-1(config)# device-alias commit
```

Fabric B Device Aliases

```
aa19-9706-2(config)# device-alias database
aa19-9706-2(config-device-alias-db)# device-alias name VSI-G370-1 pwwn 20:00:00:25:B5:54:0B:00
aa19-9706-2(config-device-alias-db)# device-alias name VSI-G370-2 pwwn 20:00:00:25:B5:54:0B:01
aa19-9706-2(config-device-alias-db)# device-alias name G370-CL3-B pwwn 50:06:0E:80:12:C9:9A:21
```

```
aa19-9706-2(config-device-alias-db)# device-alias name G370-CL4-A pwwn 50:06:0E:80:12:C9:9A:30
aa19-9706-2(config-device-alias-db)# exit
aa19-9706-2(config)# device-alias commit
```

Create Zoning

With the device alias created, the following will occur on each fabric in this example:

- Smart-zoning will be set for the fabric
- Zones for each host initiator to storage targets will be created
- A zoneset will be created
- Created zones will be added to the zone
- The zoneset will be activated
- The configuration will be saved to allow the device aliases and the zoning to persist on each fabric

Fabric A Zoning

```
aa19-9706-1(config)# zone name VSI-G370-1 vsan 101
aa19-9706-1(config-zone)# member device-alias VSI-G370-1 init
aa19-9706-1(config-zone)# member device-alias G370-CL1-A target
aa19-9706-1(config-zone)# member device-alias G370-CL2-B target
aa19-9706-1(config-zone)# zone name VSI-G370-2 vsan 101
aa19-9706-1(config-zone)# member device-alias VSI-G370-2 init
aa19-9706-1(config-zone)# member device-alias G370-CL1-A target
aa19-9706-1(config-zone)# member device-alias G370-CL2-B target
aa19-9706-1(config-zone)# zoneset name vsi-zoneset vsan 101
aa19-9706-1(config-zoneset)# member VSI-G370-1
aa19-9706-1(config-zoneset)# member VSI-G370-2
aa19-9706-1(config-zoneset)# zoneset activate name vsi-zoneset vsan 101
Zoneset activation initiated. check zone status
aa19-9706-1(config)# copy run start
[##########################################] 100%
Copy complete.
```

Fabric B Zoning

```
aa19-9706-2(config)# zone name VSI-G370-1 vsan 102
aa19-9706-2(config-zone)# member device-alias VSI-G370-1 init
aa19-9706-2(config-zone)# member device-alias G370-CL3-B target
aa19-9706-2(config-zone)# member device-alias G370-CL4-A target
aa19-9706-2(config-zone)# zone name VSI-G370-2 vsan 102
aa19-9706-2(config-zone)# member device-alias G370-CL3-B target
aa19-9706-2(config-zone)# member device-alias G370-CL4-A target
aa19-9706-2(config-zone)# zoneset name vsi-zoneset vsan 102
aa19-9706-2(config-zoneset)# member VSI-G370-1
aa19-9706-2(config-zoneset)# member VSI-G370-2
aa19-9706-2(config-zoneset)# zoneset activate name vsi-zoneset vsan 102
Zoneset activation initiated. check zone status
aa19-9706-2(config)# copy run start
[##########################################] 100%
Copy complete.
```

Appendix: MDS Example startup-configuration File

MDS A Configuration

MDS B is identical except for the VSAN used and the example wwn/pwwn.

```
version 8.3(1)
power redundancy-mode ps-redundant
power redundancy-mode ps-redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $5$ZUTjKo32$/.kdz66FvCWJctfVT0cQwlBqHzKkIhWrIc86gXv6aX8  role network-admin
ip domain-lookup
ip host AA19-9706-1  192.168.168.18
aaa group server radius radius
class-map type qos match-all copp-s-selfIp
snmp-server user admin network-admin auth md5 0x267983550d7062ba68d3e8acb81bb136 priv
0x267983550d7062ba68d3e8acb81bb136 localizedkey
snmp-server host 10.1.168.101 traps version 2c public udp-port 2162
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 192.168.168.254
vlan 1
vlan 1
vsan database
  vsan 101 name "Fabric-A"
device-alias mode enhanced
device-alias database
  device-alias name G370-CL1-A pwwn 50:06:0e:80:12:c9:9a:00
  device-alias name G370-CL1-B pwwn 50:06:0e:80:12:c9:9a:01
  device-alias name G370-CL2-A pwwn 50:06:0e:80:12:c9:9a:10
  device-alias name G370-CL2-B pwwn 50:06:0e:80:12:c9:9a:11
  device-alias name G1500-CL1-A pwwn 50:06:0e:80:07:56:24:00
  device-alias name G1500-CL1-C pwwn 50:06:0e:80:07:56:24:02
  device-alias name G1500-CL1-J pwwn 50:06:0e:80:07:56:24:08
  device-alias name G1500-CL1-L pwwn 50:06:0e:80:07:56:24:0a
  device-alias name G1500-CL2-A pwwn 50:06:0e:80:07:56:24:10
  device-alias name G1500-CL2-C pwwn 50:06:0e:80:07:56:24:12
  device-alias name G1500-CL2-J pwwn 50:06:0e:80:07:56:24:18
  device-alias name G1500-CL2-L pwwn 50:06:0e:80:07:56:24:1a
  device-alias name VSI-FC-G370-1 pwwn 20:00:00:25:b5:54:0a:00
  device-alias name VSI-FC-G370-2 pwwn 20:00:00:25:b5:54:0a:01
  device-alias name VSI-FC-G370-3 pwwn 20:00:00:25:b5:54:0a:02
  device-alias name VSI-FC-G370-4 pwwn 20:00:00:25:b5:54:0a:03
  device-alias name VSI-FC-G1500-1 pwwn 20:00:00:25:b5:32:0a:00
  device-alias name VSI-FC-G1500-2 pwwn 20:00:00:25:b5:32:0a:01
  device-alias name VSI-FC-G1500-3 pwwn 20:00:00:25:b5:32:0a:02
  device-alias name VSI-FC-G1500-4 pwwn 20:00:00:25:b5:32:0a:03
```

```
device-alias commit

fcdomain fcid database
```

```
  vsan 101 wwn 50:06:0e:80:07:56:24:18 fcid 0xbc0000 dynamic
    !          [G1500-CL2-J]
  vsan 101 wwn 50:06:0e:80:07:56:24:02 fcid 0xbc0020 dynamic
    !          [G1500-CL1-C]
  vsan 101 wwn 50:06:0e:80:07:56:24:10 fcid 0xbc0040 dynamic
    !          [G1500-CL2-A]
  vsan 101 wwn 50:06:0e:80:07:56:24:12 fcid 0xbc0060 dynamic
    !          [G1500-CL2-C]
  vsan 101 wwn 50:06:0e:80:12:c9:9a:01 fcid 0xbc0080 dynamic
    !          [G370-CL1-B]
  vsan 101 wwn 50:06:0e:80:07:56:24:0a fcid 0xbc00a0 dynamic
    !          [G1500-CL1-L]
  vsan 101 wwn 50:06:0e:80:12:c9:9a:11 fcid 0xbc00c0 dynamic
    !          [G370-CL2-B]
  vsan 101 wwn 50:06:0e:80:12:c9:9a:10 fcid 0xbc00e0 dynamic
    !          [G370-CL2-A]
  vsan 101 wwn 50:06:0e:80:07:56:24:00 fcid 0xbc0100 dynamic
    !          [G1500-CL1-A]
  vsan 101 wwn 50:06:0e:80:07:56:24:08 fcid 0xbc0120 dynamic
    !          [G1500-CL1-J]
  vsan 101 wwn 50:06:0e:80:07:56:24:1a fcid 0xbc0140 dynamic
    !          [G1500-CL2-L]
  vsan 101 wwn 50:06:0e:80:12:c9:9a:00 fcid 0xbc0160 dynamic
    !          [G370-CL1-A]
  vsan 101 wwn 24:01:00:de:fb:d6:3b:40 fcid 0xbc0180 dynamic
  vsan 101 wwn 20:00:00:25:b5:32:0a:01 fcid 0xbc0181 dynamic
    !          [VSI-FC-G1500-2]
  vsan 101 wwn 20:00:00:25:b5:32:0a:00 fcid 0xbc0182 dynamic
    !          [VSI-FC-G1500-1]
  vsan 101 wwn 20:00:00:25:b5:32:0a:02 fcid 0xbc0183 dynamic
    !          [VSI-FC-G1500-3]
  vsan 101 wwn 20:00:00:25:b5:32:0a:03 fcid 0xbc0184 dynamic
    !          [VSI-FC-G1500-4]
  vsan 101 wwn 24:01:00:de:fb:ff:fe:00 fcid 0xbc01a0 dynamic
  vsan 101 wwn 20:00:00:25:b5:54:0a:00 fcid 0xbc01a1 dynamic
    !          [VSI-FC-G370-1]
  vsan 101 wwn 20:00:00:25:b5:54:0a:01 fcid 0xbc01a2 dynamic
    !          [VSI-FC-G370-2]
  vsan 101 wwn 20:00:00:25:b5:54:0a:03 fcid 0xbc01a3 dynamic
    !          [VSI-FC-G370-4]
  vsan 101 wwn 20:00:00:25:b5:54:0a:02 fcid 0xbc01a4 dynamic
    !          [VSI-FC-G370-3]
zone smart-zoning enable vsan 101
!Active Zone Database Section for vsan 101
zone name VSI-FC-G370-1 vsan 101
    member device-alias VSI-FC-G370-1 init
    member device-alias G370-CL1-A target
    member device-alias G370-CL1-B target
    member device-alias G370-CL2-A target
    member device-alias G370-CL2-B target

zone name VSI-FC-G370-2 vsan 101
    member device-alias VSI-FC-G370-2 init
    member device-alias G370-CL1-A target
    member device-alias G370-CL1-B target
    member device-alias G370-CL2-A target
    member device-alias G370-CL2-B target

zone name VSI-FC-G370-3 vsan 101
    member device-alias VSI-FC-G370-3 init
    member device-alias G370-CL1-A target
    member device-alias G370-CL1-B target
    member device-alias G370-CL2-A target
    member device-alias G370-CL2-B target

zone name VSI-FC-G370-4 vsan 101
    member device-alias VSI-FC-G370-4 init
    member device-alias G370-CL1-A target
    member device-alias G370-CL1-B target
```

```
    member device-alias G370-CL2-A target
    member device-alias G370-CL2-B target

zone name VSI-FC-G1500-1 vsan 101
    member device-alias G1500-CL1-A target
    member device-alias G1500-CL1-C target
    member device-alias G1500-CL1-J target
    member device-alias G1500-CL1-L target
    member device-alias G1500-CL2-A target
    member device-alias G1500-CL2-C target
    member device-alias G1500-CL2-J target
    member device-alias G1500-CL2-L target
    member device-alias VSI-FC-G1500-1 init

zone name VSI-FC-G1500-2 vsan 101
    member device-alias G1500-CL1-A target
    member device-alias G1500-CL1-C target
    member device-alias G1500-CL1-J target
    member device-alias G1500-CL1-L target
    member device-alias G1500-CL2-A target
    member device-alias G1500-CL2-C target
    member device-alias G1500-CL2-J target
    member device-alias G1500-CL2-L target
    member device-alias VSI-FC-G1500-2 init

zone name VSI-FC-G1500-3 vsan 101
    member device-alias G1500-CL1-A target
    member device-alias G1500-CL1-C target
    member device-alias G1500-CL1-J target
    member device-alias G1500-CL1-L target
    member device-alias G1500-CL2-A target
    member device-alias G1500-CL2-C target
    member device-alias G1500-CL2-J target
    member device-alias G1500-CL2-L target
    member device-alias VSI-FC-G1500-3 init

zone name VSI-FC-G1500-4 vsan 101
    member device-alias G1500-CL1-A target
    member device-alias G1500-CL1-C target
    member device-alias G1500-CL1-J target
    member device-alias G1500-CL1-L target
    member device-alias G1500-CL2-A target
    member device-alias G1500-CL2-C target
    member device-alias G1500-CL2-J target
    member device-alias G1500-CL2-L target
    member device-alias VSI-FC-G1500-4 init

zoneset name hvcs-vsi-zoneset vsan 101
    member VSI-FC-G370-1
    member VSI-FC-G370-2
    member VSI-FC-G370-3
    member VSI-FC-G370-4
    member VSI-FC-G1500-1
    member VSI-FC-G1500-2
    member VSI-FC-G1500-3
    member VSI-FC-G1500-4

zoneset activate name hvcs-vsi-zoneset vsan 101
do clear zone database vsan 101
!Full Zone Database Section for vsan 101
zoneset name hvcs-vsi-zoneset vsan 101


interface mgmt0
  ip address 192.168.168.18 255.255.255.0

interface port-channel11
  channel mode active
  switchport description UCS-6332-portchannel
  switchport speed auto max 32000
  switchport rate-mode dedicated
```

```
interface port-channel15
  channel mode active
  switchport description UCS-6454-portchannel
  switchport speed auto max 32000
  switchport rate-mode dedicated
vsan database
  vsan 101 interface port-channel11
  vsan 101 interface port-channel15
  vsan 101 interface fc1/7
  vsan 101 interface fc1/8
  vsan 101 interface fc1/9
  vsan 101 interface fc1/10
  vsan 101 interface fc1/11
  vsan 101 interface fc1/12
  vsan 101 interface fc1/13
  vsan 101 interface fc1/14
  vsan 101 interface fc1/15
  vsan 101 interface fc1/16
  vsan 101 interface fc1/17
  vsan 101 interface fc1/18

interface fc1/1
  switchport description UCS-6332-A:1/1
  channel-group 11 force
  no shutdown

interface fc1/2
  switchport description UCS-6332-A:1/2
  channel-group 11 force
  no shutdown

interface fc1/3
  switchport description UCS-6332-A:1/3
  channel-group 11 force
  no shutdown

interface fc1/4
  switchport description UCS-6332-A:1/4
  channel-group 11 force
  no shutdown

interface fc1/5
  switchport description UCS-6454-A:1/1
  channel-group 15 force
  no shutdown

interface fc1/6
  switchport description UCS-6454-A:1/2
  channel-group 15 force
  no shutdown

interface fc1/7
  switchport description G1500-A:CL 1-A
  no shutdown

interface fc1/8
  switchport description G1500-A:CL 2-A
  no shutdown

interface fc1/9
  switchport description G1500-A:CL 1-J
  no shutdown

interface fc1/10
  switchport description G1500-A:CL 2-J
  no shutdown

interface fc1/11
  switchport description G370-A:CL 1-A
  no shutdown

interface fc1/12
```

```
  switchport description G370-A:CL 2-B
  no shutdown

interface fc1/13
  no shutdown

interface fc1/14
  no shutdown

interface fc1/15
  no shutdown

interface fc1/16
  no shutdown

interface fc1/17
  no shutdown

interface fc1/18
  no shutdown

interface fc1/19

interface fc1/20

interface fc1/21

interface fc1/22

interface fc1/23

interface fc1/24

interface fc1/25

interface fc1/26

interface fc1/27

interface fc1/28

interface fc1/29

interface fc1/30

interface fc1/31

interface fc1/32

interface fc1/33

interface fc1/34

interface fc1/35

interface fc1/36

interface fc1/37

interface fc1/38

interface fc1/39

interface fc1/40

interface fc1/41

interface fc1/42

interface fc1/43

interface fc1/44
```

```
interface fc1/45

interface fc1/46

interface fc1/47

interface fc1/48
clock timezone EST 0 0
switchname AA19-9706-1
line console
line vty
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/6
interface fc1/1
interface fc1/5
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
ip default-gateway 192.168.168.254
```

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-channel13
  description vPC UCS 6332-16UP-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 13

interface port-channel14
  description vPC UCS 6332-16UP-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 14

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 Ethernet1/3
  description vPC 13 connection to UCS 6332-16UP-1 FI Ethernet1/39
  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 13 mode active

interface Ethernet1/4
  description vPC 14 connection to UCS 6332-16UP-2 FI Ethernet1/39
  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 14 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

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-channel13
  description vPC UCS 6332-16UP-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 13

interface port-channel14
  description vPC UCS 6332-16UP-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 14

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 Ethernet1/3
  description vPC 13 connection to UCS 6332-16UP-1 FI Ethernet1/40

interface Ethernet1/4
  description vPC 14 connection to UCS 6332-16UP-2 FI Ethernet1/40
  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 14 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
  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 13 mode active

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

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
```