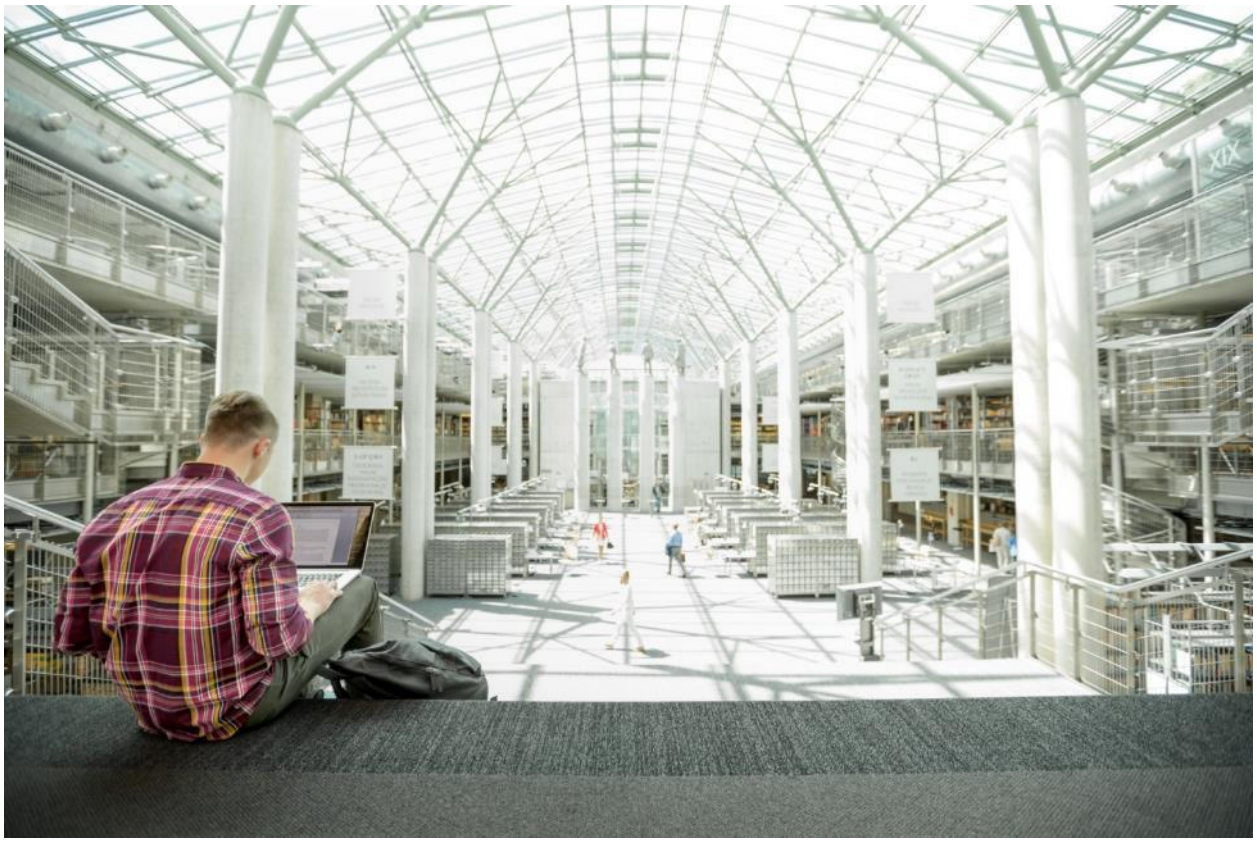

VersaStack with Cisco UCS M5 Servers, IBM SVC, and vSphere 6.5 U1

Deployment Guide for VersaStack with Cisco UCS M5, IBM SVC, and vSphere 6.5 U1

Last Updated: May 31, 2018

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, visit:

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

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series. Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2018 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary 8

Solution Overview 9

Introduction 9

Audience 9

Purpose of this Document 9

Solution Design 11

Architecture 11

Physical Topology 11

Software Revisions 13

Configuration Guidelines 14

Physical Infrastructure 15

Cisco UCS Connectivity to Nexus Switches 15

IBM SAN Volume Controller Connectivity to Nexus Switches 16

Cisco UCS connectivity to SAN Fabric 17

Network Configuration 21

Cisco Nexus 9000 Initial Configuration Setup 21

Cisco Nexus A 21

Cisco Nexus B 23

Enable Appropriate Cisco Nexus 9000 Features and Settings 24

Cisco Nexus 9000 A and Cisco Nexus 9000 B 24

Create VLANs for VersaStack IP Traffic 25

Cisco Nexus 9000 A and Cisco Nexus 9000 B 25

Configure Virtual Port Channel Domain 25

Cisco Nexus 9000 A 25

Cisco Nexus 9000 B 26

Configure Network Interfaces for the vPC Peer Links 26

Cisco Nexus 9000 A 27

Cisco Nexus 9000 B 27

Configure Network Interfaces to Cisco UCS Fabric Interconnect 28

Cisco Nexus 9000 A 28

Cisco Nexus 9000 B 30

Configure Network Interfaces Connected to IBM SVC iSCSI Ports 31

Management Uplink into Existing Network Infrastructure 33

Cisco Nexus 9000 A and B using Port Channel Example....33
Cisco MDS 9396S Initial Configuration Setup....34
Enable Appropriate Cisco MDS Features and Settings....34
Initial Storage Configuration....35
IBM FlashSystem 900 Base Configuration....35
Prerequisites....35
Creating a Replacement USB Key....35
IBM FlashSystem 900 Initial Configuration....36
IBM FlashSystem 900 GUI Setup....41
IBM Storwize V5030 Base Configuration....50
Prerequisites....50
IBM Storwize V5000 Initial Configuration....51
IBM Storwize V5000 GUI Setup....54
IBM San Volume Controller Base Configuration....69
Prerequisites....70
IBM SAN Volume Controller Initial Configuration....70
IBM SAN Volume Controller GUI Setup....74
Cisco MDS Zone Configuration....83
Adding Additional Nodes....87
Adding External Storage to the SVC....92
IBM SVC iSCSI configuration....102
Modifying Interface MTU....107
Cisco UCS Server Configuration....108
Cisco UCS Initial Configuration....108
Cisco UCS 6332-16UP A....108
Cisco UCS 6332-16UP B....109
Cisco UCS Software Upgrade....109
Log in to Cisco UCS Manager....109
Upgrade Cisco UCS Manager Software to Version 3.2(3d)....110
Configure Cisco UCS Call Home....110
Add a Block of Management IP Addresses for KVM Access....110
Synchronize Cisco UCS to NTP....111
Edit Chassis Discovery Policy....112
Enable Server and Uplink Ports....113
Acknowledge Cisco UCS Chassis and FEX....115

Enable Fibre Channel Ports 115
Create VSAN for the Fibre Channel Interfaces 116
Create Port Channels for the Fibre Channel Interfaces 118
Create Port Channels for Ethernet Uplinks 121
Create MAC Address Pools 122
Create UUID Suffix Pool 125
Create Server Pool 126
Create a WWNN Address Pool for FC based Storage Access 127
Create a WWPN Address Pools for FC Based Storage Access 128
Create IQN Pools for iSCSI Boot and LUN Access 130
Create IP Pools for iSCSI Boot and LUN Access 131
Create VLANs 133
Create Host Firmware Package 135
Set Jumbo Frames in Cisco UCS Fabric 136
Create Local Disk Configuration Policy 137
Create Network Control Policy for Link Layer Discovery Protocol 137
Create Power Control Policy 138
Create Server Pool Qualification Policy (Optional) 139
Create Server BIOS Policy 140
Update Default Maintenance Policy 142
Create vNIC/vHBA Placement Policy 142
Create vNIC Templates 143
Create Management vNICs 144
Create vMotion vNICs 148
Create Application vNICs 152
Create iSCSI vNICs 156
Create LAN Connectivity Policy 160
Adding iSCSI vNICs in LAN Policy 169
Create vHBA Templates for FC Connectivity 171
Create FC SAN Connectivity Policies 173
Create iSCSI Boot Policy 175
Create FC Boot Policies 176
Create iSCSI Boot Service Profile Template 182
Configure Storage Provisioning 183
Configure Networking Options 184

Configure Storage Options 185
Configure Zoning Options 185
Configure vNIC/HBA Placement 185
Configure vMedia Policy 185
Configure Server Boot Order 186
Configure Maintenance Policy 191
Configure Server Assignment 192
Configure Operational Policies 193
Create iSCSI Boot Service Profiles 194
Create FC Boot Service Profile Template 195
Configure Storage Provisioning 196
Configure Networking Options 197
Configure SAN Connectivity 197
Configure Zoning 198
Configure vNIC/HBA Placement 198
Configure vMedia Policy 198
Configure Server Boot Order 198
Configure Maintenance Policy 199
Configure Server Assignment 200
Configure Operational Policies 201
Create FC Boot Service Profiles 202
Backup the Cisco UCS Manager Configuration 204
Adding Servers 204
Gather Necessary WWPN Information (FC Deployment) 204
Gather Necessary IQN Information (iSCSI Deployment) 205
IBM SVC iSCSI Storage Configuration 207
IBM SVC iSCSI Configuration 207
Create Volumes on the Storage System 207
Map Volumes to Hosts 209
IBM SVC Fibre Channel Storage Configuration 213
Cisco MDS 9396S SAN Zoning for UCS Hosts 213
Cisco MDS - A Switch 213
Cisco MDS - B Switch 215
IBM SVC FC Configuration 217
Create Volumes on the Storage System 218

Map Volumes to Hosts 219
VMware vSphere Setup for UCS Host Environment 223
VMware ESXi 6.5 U1 223
Log in to Cisco UCS Manager 223
Install ESXi on the UCS Servers 223
Set Up Management Networking for ESXi Hosts 225
VMware vSphere Configuration 227
Download VMware vSphere Client 228
Log in to VMware ESXi Hosts Using VMware vSphere Client 228
Install VMware Drivers for the Cisco Virtual Interface Card (VIC) 228
Mount Required Datastores 229
Configure NTP on ESXi Hosts 230
Move VM Swap File Location 231
Deploy VMware vCenter Appliance 6.5 232
Set Up vCenter Server 232
Setup Datacenter, Cluster, DRS and HA for ESXi Nodes 233
Add the VMware ESXi Hosts Using the VMware vSphere Web Client 234
ESXi Dump Collector Setup for iSCSI Hosts (iSCSI Configuration Only) 239
Configure ESXi Networking 240
Update Management vSwitch0 Configuration 240
Create vSwitch1 for vMotion 244
Configure iSCSI Adapters 253
Create a VMware vDS for Application and Production Networks 263
Add the ESXi Hosts to the vDS 267
About the Authors 275

Executive Summary

Cisco Validated Designs (CVDs) deliver 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 the customers and to guide them from design to deployment.

Customers looking to deploy applications using shared data center infrastructure face a number of challenges. A recurrent infrastructure challenge is to achieve the levels of IT agility and efficiency that can effectively meet the company business objectives. Addressing these challenges requires having an optimal solution with the following key characteristics:

- Availability: Help ensure applications and services availability at all times with no single point of failure
- Flexibility: Ability to support new services without requiring underlying infrastructure modifications
- Efficiency: Facilitate efficient operation of the infrastructure through re-usable policies
- Manageability: Ease of deployment and ongoing management to minimize operating costs
- Scalability: Ability to expand and grow with significant investment protection
- Compatibility: Minimize risk by ensuring compatibility of integrated components

Cisco and IBM have partnered to deliver a series of VersaStack solutions that enable strategic data center platforms with the above characteristics. VersaStack solution delivers an integrated architecture that incorporates compute, storage and network design best practices thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance and support that can be used in various stages (planning, designing and implementation) of a deployment.

The VersaStack solution, described in this CVD, delivers a converged infrastructure platform specifically designed for Virtual Server Infrastructure. In this deployment, SVC standardizes storage functionality across different arrays and provides a single point of control for virtualized storage for greater flexibility and with the addition of Cisco UCS M5 servers and Cisco UCS 6300 series interconnects, the solution provides better compute performance and network throughput with 40G support for ethernet and 16G support for fibre channel connectivity. The design showcases:

- Cisco Nexus 9000 switching architecture running NX-OS mode
- IBM SVC providing single point of management and control for IBM FlashSystem 900 and IBM Storwize V5030
- Cisco Unified Computing System (Cisco UCS) servers with Intel Xeon processors
- Storage designs supporting Fibre Channel and iSCSI based storage access
- VMware vSphere 6.5U1 hypervisor
- Cisco MDS Fibre Channel (FC) switches for SAN connectivity

Solution Overview

Introduction

VersaStack solution is a pre-designed, integrated and validated architecture for the data center that combines Cisco UCS servers, Cisco Nexus family of switches, Cisco MDS fabric switches and IBM SVC, Storwize and FlashSystem storage arrays into a single, flexible architecture. VersaStack is designed for high availability, with no single points of failure, while maintaining cost-effectiveness and flexibility in design to support a wide variety of workloads.

VersaStack design can support different hypervisor options, bare metal servers and can also be sized and optimized based on customer workload requirements. VersaStack design discussed in this document has been validated for resiliency (under fair load) and fault tolerance during system upgrades, component failures, and partial as well as complete loss of power scenarios.

This document discusses the design principles that go into the VersaStack solution, which is a validated Converged Infrastructure (CI) jointly developed by Cisco and IBM. The solution is a predesigned, best-practice data center architecture with VMware vSphere built on the Cisco Unified Computing System (Cisco UCS) with Cisco UCS 6300 series fabric interconnects, Cisco UCS M5 series blade and rack mount servers, the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and IBM SAN Volume Controller (SVC) virtualizing Storwize V5030 and FS900 Storage arrays supporting Fibre Channel and iSCSI based storage access.

The solution architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Server Infrastructure (VSI).

Audience

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

Purpose of this Document

This document provides step-by-step configuration and implementation guidelines for setting up VersaStack. The following design elements distinguish this version of VersaStack from previous models:

- Cisco UCS B200 M5
- Cisco UCS 6300 Fabric Interconnects
- IBM SVC 2145-SV1 release 7.8.1.4
- IBM FlashSystem 900 and IBM Storwize V5030 release 7.8.1.4
- Support for the Cisco UCS release 3.2.3

- Validation of IP-based storage design with Nexus NX-OS switches supporting iSCSI based storage access

For more information on previous VersaStack models, please refer to the VersaStack guides at:

http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/versastack-designs.html

Solution Design

Architecture

VersaStack with Cisco UCS M5 and IBM SVC architecture aligns with the converged infrastructure configurations and best practices as identified in the previous VersaStack releases. The system includes hardware and software compatibility support between all components and aligns to the configuration best practices for each of these components. All the core hardware components and software releases are listed and supported on both the Cisco compatibility list:

http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html

and IBM Interoperability Matrix:

http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss

The system supports high availability at network, compute and storage layers such that no single point of failure exists in the design. The system utilizes 10 and 40 Gbps Ethernet jumbo-frame based connectivity combined with port aggregation technologies such as virtual port-channels (VPC) for non-blocking LAN traffic forwarding. A dual SAN 16Gbps environment provides redundant storage access from compute devices to the storage controllers.

Physical Topology

This VersaStack with Cisco UCS M5 and IBM SVC solution utilizes Cisco UCS platform with Cisco B200 M5 half-width blades connected and managed through Cisco UCS 6332-16UP Fabric Interconnects and the integrated Cisco UCS manager. These high performance servers are configured as stateless compute nodes where ESXi 6.5 U1 hypervisor is loaded using SAN (iSCSI and FC) boot. The boot disks to store ESXi hypervisor image and configuration along with the block and file based datastores to host application Virtual Machines (VMs) are provisioned on the IBM storage devices.

The link aggregation technologies play an important role in VersaStack solution providing improved aggregate bandwidth and link resiliency across the solution stack. Cisco UCS, and Cisco Nexus 9000 platforms support active port channeling using 802.3ad standard Link Aggregation Control Protocol (LACP). In addition, the Cisco Nexus 9000 series features virtual Port Channel (vPC) capability which allows links that are physically connected to two different Cisco Nexus devices to appear as a single "logical" port channel. Each Cisco UCS Fabric Interconnect (FI) is connected to both the Cisco Nexus 93180 switches using virtual port-channel (vPC) enabled 40GbE uplinks for a total aggregate system bandwidth of 80GBps. Additional ports can be easily added to the design for increased throughput. Each Cisco UCS 5108 chassis is connected to the UCS FIs using a pair of 40GbE ports from each IO Module for a combined 80GbE uplink. When Cisco UCS C-Series servers are used, each of the Cisco UCS C-Series servers connects directly into each of the FIs using a 10/40Gbps converged link for an aggregate bandwidth of 20/80Gbps per server.

To provide the compute to storage system connectivity, this design guides highlights two different storage connectivity options:

- Option 1: iSCSI based storage access through Cisco Nexus Fabric

- Option 2: FC based storage access through Cisco MDS 9396S

While storage access from the Cisco UCS compute nodes to IBM SVC storage nodes can be iSCSI or FC based, IBM SVC nodes, FlashSystem 900 and Storwize V5030 communicate with each other on Cisco MDS 9396S based FC SAN fabric only.

Figure 1 illustrates a high-level topology of the system connectivity. The VersaStack infrastructure satisfies the high-availability design requirements and is physically redundant across the network, compute and storage stacks. The integrated compute stack design presented in this document can withstand failure of one or more links as well as the failure of one or more devices.

IBM SVC Nodes, IBM FlashSystem 900 and IBM Storwize V5030 are all connected using a Cisco MDS 9396S based redundant FC fabric. To provide FC based storage access to the compute nodes, Cisco UCS Fabric Interconnects are connected to the same Cisco MDS 9396S switches and zoned appropriately. To provide iSCSI based storage access, IBM SVC is connected directly to the Cisco Nexus 93180 switches. One 10GbE port from each IBM SVC Node is connected to each of the two Cisco Nexus 93180 switches providing an aggregate bandwidth of 40Gbps.

Based on the customer requirements, the compute to storage connectivity in the SVC solution can be deployed as an FC-only option, iSCSI-only option or a combination of both. Figure 1 shows the connectivity option to support both iSCSI and FC.

Figure 1 VersaStack iSCSI and FC Storage Design with IBM SVC

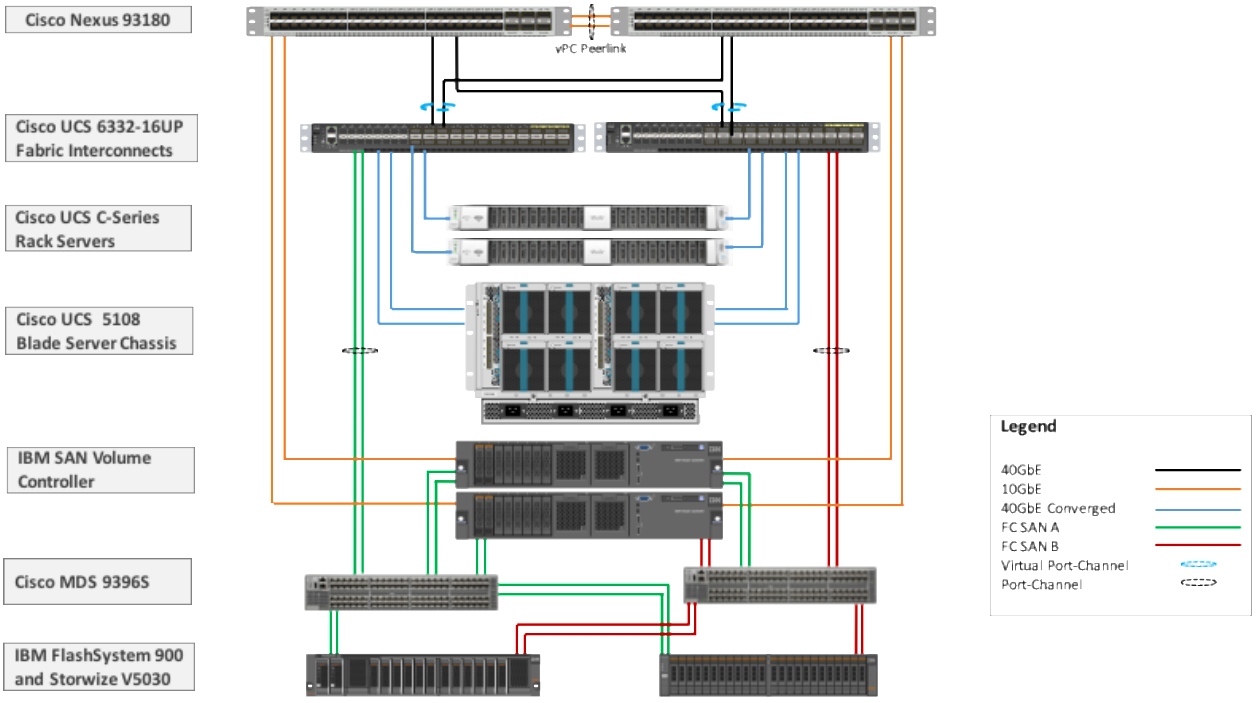

The reference architecture covered in this document leverages:

- One* Cisco UCS 5108 Blade Server chassis with 2304 Series Fabric Extenders (FEX)
- Four* Cisco UCS B200-M5 Blade Servers

- Two Cisco UCS 6332-16UPUP Fabric Interconnects (FI)
- Two Cisco Nexus 93180YC-EX Switches
- Two Cisco MDS 9396S Fabric Switches
- Two** IBM SAN Volume Controller 2145-SV1 nodes
- One IBM FlashSystem 900
- One dual controller IBM Storwize V5030
- VMware vSphere 6.5 update 1
- VMware Virtual Distributed Switch (VDS) ***

* The actual number of servers in customer environment will vary.

** This design guide showcases two 2145-SV1 nodes setup as a two-node cluster. This configuration can be customized for customer specific deployments.

*** This deployment guide covers both VMware Virtual Distributed Switch (VDS) and standard vSwitch.

This document guides customers through the low-level steps for deploying the base architecture. These procedures cover everything from physical cabling to network, compute, and storage device configurations.

For detailed information about the VersaStack design, see:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/versastack_svc_ucsm5_designguide.html

Software Revisions

Table 1 lists the hardware and software versions used for the solution validation.

It is important to note that Cisco, IBM, and VMware have interoperability matrices that should be referenced to determine support for any specific implementation of VersaStack. See the following links for more information:

- IBM System Storage Interoperation Center
- Cisco UCS Hardware and Software Interoperability Tool

Table 1 Hardware and Software Revisions

Layer	Device	Image	Comments

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6300 Series, Cisco UCS B-200 M5, Cisco UCS C-220 M5	3.2(3d)	Includes the Cisco UCS-IOM 2304, Cisco UCS Manager, and Cisco UCS VIC 1340
	Cisco nenic Driver	1.0.16.0	Ethernet driver for Cisco VIC
	Cisco fnic Driver	1.6.0.37	FCoE driver for Cisco VIC
Network	Cisco Nexus Switches	7.0(3)I4(7)	NXOS
	Cisco MDS 9396S	7.3(0)DY(1)	FC switch firmware version
Storage	IBM SVC	7.8.1.4	Software version
	IBM FlashSystem 900	1.4.5.0	Software version
	IBM Storwize V5030	7.8.1.4	Software version
Software	VMware vSphere ESXi	6.5 update 1	Software version
	VMware vCenter	6.5 update 1	Software version

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available VersaStack configuration. Therefore, appropriate references are provided to indicate the component being configured at each step, such as 01 and 02 or A and B. For example, the Cisco UCS fabric interconnects are identified as FI-A or FI-B. This document is intended to enable customers and partners to fully configure the customer environment and during this process, various steps may require the use of customer-specific naming conventions, IP addresses, and VLAN schemes, as well as appropriate MAC addresses.

This document details network (Nexus and MDS), compute (Cisco UCS), virtualization (VMware) and related storage configurations (host to storage system connectivity).

Table 2 lists various VLANs, VSANs and subnets used to setup VersaStack infrastructure to provide connectivity between core elements of the design.

Table 2 VersaStack Infrastructure Configuration

VLAN Name	VLAN	Subnet
IB-MGMT	111	192.168.160.0/22
Infra-iSCSI-A	3161	10.29.161.0/24

VLAN Name	VLAN	Subnet
infra-iSCSI-B	3162	10.29.162.0/24
vMotion	3173	10.29.173.0/24
Native-2	2	N/A
VM Network	3174	10.29.174.0/24
VSAN-A	101	N/A
VSAN-B	102	N/A

Physical Infrastructure

The information in this section is provided as a reference for cabling the equipment in VersaStack environment. To simplify the documentation, the architecture shown in Figure 1 is broken down into network, compute and storage related physical connectivity details.

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

Customers can choose interfaces and ports of their liking but failure to follow the exact connectivity shown in figures below will result in changes to the deployment procedures since specific port information is used in various configuration steps

Cisco UCS Connectivity to Nexus Switches

For physical connectivity details of Cisco UCS to the Cisco Nexus switches, refer to Figure 2 and 0.

Figure 2 Cisco UCS Connectivity to the Nexus Switches

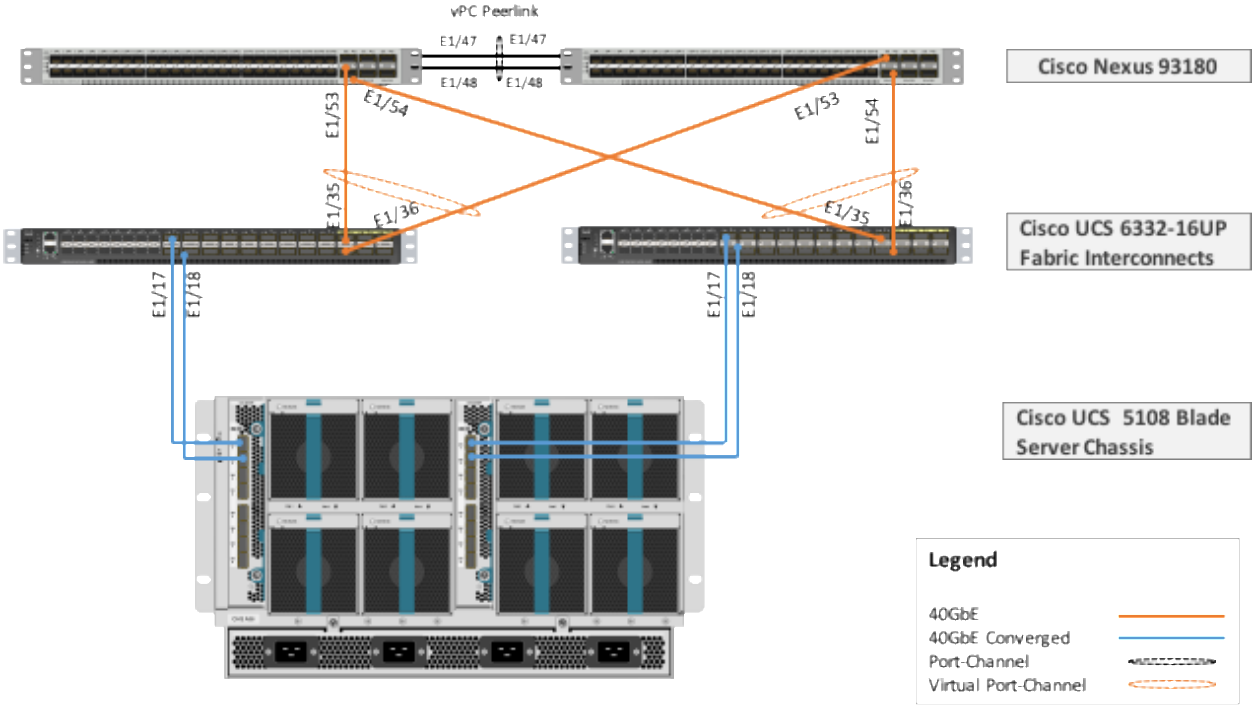

Table 3 Cisco UCS connectivity to Nexus Switches

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/17	10GbE	Cisco UCS Chassis FEX A	IOM 1/1
Cisco UCS Fabric Interconnect A	Eth1/18	10GbE	Cisco UCS Chassis FEX A	IOM 1/2
Cisco UCS Fabric Interconnect A	Eth1/35	10GbE	Cisco Nexus 93180 A	Eth1/53
Cisco UCS Fabric Interconnect A	Eth1/36	10GbE	Cisco Nexus 93180 B	Eth1/53
Cisco UCS Fabric Interconnect B	Eth1/17	10GbE	Cisco UCS Chassis FEX B	IOM 1/1
Cisco UCS Fabric Interconnect B	Eth1/18	10GbE	Cisco UCS Chassis FEX B	IOM 1/2
Cisco UCS Fabric Interconnect B	Eth1/35	10GbE	Cisco Nexus 93180 A	Eth1/54
Cisco UCS Fabric Interconnect B	Eth1/36	10GbE	Cisco Nexus 93180 B	Eth1/54

IBM SAN Volume Controller Connectivity to Nexus Switches

For physical connectivity details of SVC nodes to the Cisco Nexus Switches, refer to Table 3 . This deployment shows connectivity for a pair of IBM 2145-SV1 nodes. Additional nodes can be connected to open ports on Nexus switches as needed.

Figure 3 IBM SVC Connectivity to Nexus 9k Switches

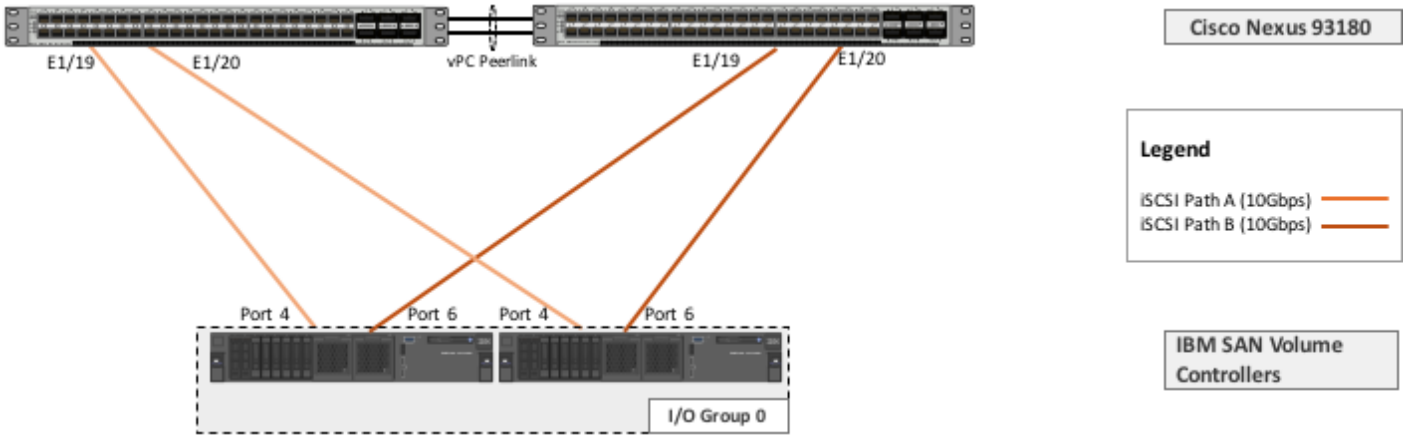

Table 4 IBM SVC Connectivity to the Nexus Switches

Local Device	Local Port	Connection	Remote Device	Remote Port
IBM 2145-SV1 node 1	Port 4	10GbE	Cisco Nexus 93180 A	Eth1/19
IBM 2145-SV1 node 1	Port 6	10GbE	Cisco Nexus 93180 B	Eth1/19
IBM 2145-SV1 node 2	Port 4	10GbE	Cisco Nexus 93180 A	Eth1/20
IBM 2145-SV1 node 2	Port 6	10GbE	Cisco Nexus 93180 B	Eth1/20

Cisco UCS connectivity to SAN Fabric

For physical connectivity details of Cisco UCS to an MDS 9396S based redundant SAN fabric, refer to Figure 4 and 0.

Figure 4 Cisco UCS Connectivity to Cisco MDS Switches

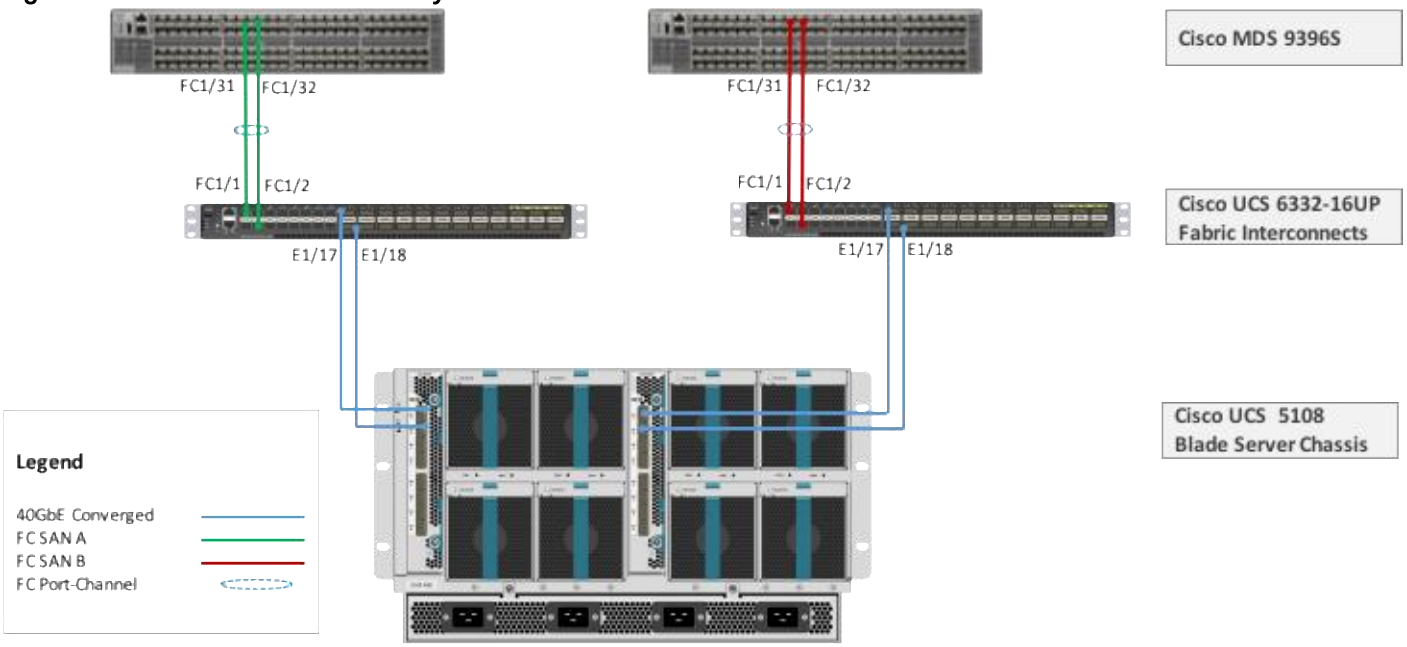

Table 5 Cisco UCS Connectivity to Cisco MDS Switches

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	FC1/1	16Gbps	Cisco MDS 9396S A	FC1/31
Cisco UCS Fabric Interconnect A	FC1/2	16Gbps	Cisco MDS 9396S A	FC1/32
Cisco UCS Fabric Interconnect B	FC1/1	16Gbps	Cisco MDS 9396S B	FC1/31
Cisco UCS Fabric Interconnect B	FC1/2	16Gbps	Cisco MDS 9396S B	FC1/32

Figure 5 illustrates FC connectivity for IBM SVC, IBM Storwize V5030 and IBM FS900. Additional nodes can be connected and configured by following the same design guidelines.

Figure 5 IBM SVC and Storage System FC Connectivity

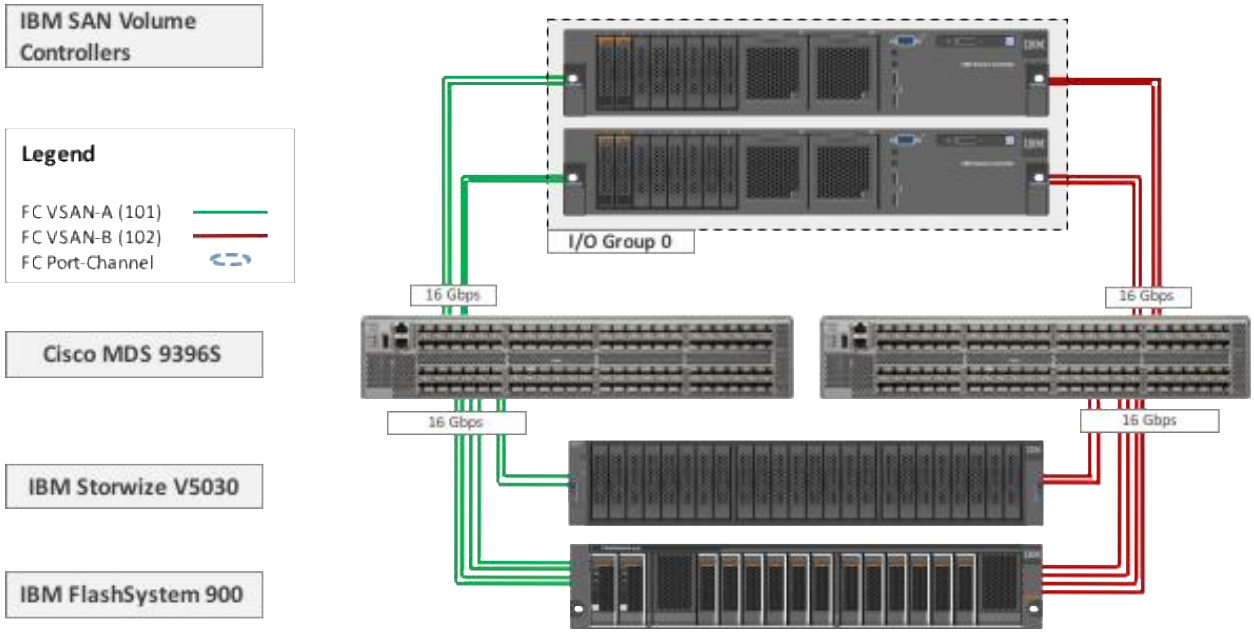

Table 6 IBM SVC and Storage System FC Connectivity

Local Device	Local Ports	Connection	Remote Device	Remote Port
IBM 2145-SV1 Node 1	Ports 1,3	16Gbps	Cisco MDS 9396S A	FC1/17, FC1/18
IBM 2145-SV1 Node 1	Ports 2,4	16Gbps	Cisco MDS 9396S B	FC1/17, FC1/18
IBM 2145-SV1 Node 2	Ports 1,3	16Gbps	Cisco MDS 9396S A	FC1/19, FC1/20
IBM 2145-SV1 Node 2	Ports 2,4	16Gbps	Cisco MDS 9396S B	FC1/19, FC1/20
IBM FS900 Canister 1	Ports 1,3	16Gbps	Cisco MDS 9396S A	FC1/13-14
IBM FS900 Canister 1	Ports 2,4	16Gbps	Cisco MDS 9396S B	FC1/13-14
IBM FS900 Canister 2	Ports 1,3	16Gbps	Cisco MDS 9396S A	FC1/15-16
IBM FS900 Canister 2	Ports 2,4	16Gbps	Cisco MDS 9396S B	FC1/15-16
IBM V5030 Controller 1	Port 3	16Gbps	Cisco MDS 9396S A	FC1/21
IBM V5030 Controller 1	Port 4	16Gbps	Cisco MDS 9396S B	FC1/21
IBM V5030 Controller 2	Port 3	16Gbps	Cisco MDS 9396S A	FC1/22
IBM V5030 Controller 2	Port 4	16Gbps	Cisco MDS 9396S B	FC1/22

IBM SVC nodes can be configured with multiple FC HBA cards, were each FC port can be dedicated for specific traffic. E.g SVC node-to-node communication, SVC node-to-Storage Array, or SVC node-to-UCS

connectivity for servicing host IO. In larger installations, reserving HBA ports for specific tasks is beneficial for optimal performance and balanced load distribution.

Network Configuration

Cisco Nexus 9000 Initial Configuration Setup

The steps provided in this section details for the initial Cisco Nexus 9000 Switch setup. In this case, we are connected using a Cisco 2901 Terminal Server that is connected via the console port on the switch.

Cisco Nexus A

To set up the initial configuration for the first Cisco Nexus switch, complete the following steps:

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

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

         ---- 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): y

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

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

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

  Enter the switch name : <Name of the Switch>

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

```
 Mgmt0 IPv4 address : <Mgmt. IP address for Switch A>
 Mgmt0 IPv4 netmask : <Mgmt. IP Subnet Mask>
 Configure the default gateway? (yes/no) [y]:
 IPv4 address of the default gateway : <Default GW for the Mgmt. IP >
 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]: 2048
  Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address : <NTP IP address>
 Configure default interface layer (L3/L2) [L2]:
  Configure default switchport interface state (shut/noshut) [noshut]:
  Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
The following configuration will be applied:
  password strength-check
switchname <Name of the Switch>
vrf context management
 ip route 0.0.0.0/0 <Default GW for the Mgmt. IP >
exit
 no feature telnet
ssh key rsa 2048 force
 feature ssh
ntp server <NTP Server IP address>
system default switchport
 no system default switchport shutdown
 copp profile strict
interface mgmt0 ip address <Mgmt. IP address for Switch A> <Mgmt. IP Subnet Mask>
no shutdown
Would you like to edit the configuration? (yes/no) [n]:
Use this configuration and save it? (yes/no) [y]:
[########################################] 100% Copy complete.
```

Cisco Nexus B

To set up the initial configuration for the second Cisco Nexus switch, complete the following steps:

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

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

         ---- 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 re-
maining dialogs.

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

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

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

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

  Enter the switch name : <Name of the Switch>

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

    Mgmt0 IPv4 address : <Mgmt. IP address for Switch B>

    Mgmt0 IPv4 netmask : <Mgmt. IP Subnet Mask>

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

    IPv4 address of the default gateway : <Default GW for the Mgmt. IP>

  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]: 2048

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

```
    NTP server IPv4 address : <NTP Server IP address>
  Configure default interface layer (L3/L2) [L2]:
  Configure default switchport interface state (shut/noshut) [noshut]:
  Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
The following configuration will be applied:
password strength-check
switchname <Name of the Switch>
vrf context management
ip route 0.0.0.0/0 <Default GW for the Mgmt. IP>
exit
 no feature telnet
ssh key rsa 2048 force
 feature ssh
ntp server <NTP Server IP address>
system default switchport
 no system default switchport shutdown
 copp profile strict
interface mgmt0 ip address <Mgmt. IP address for Switch B> <Mgmt. IP Subnet Mask>
no shutdown
Would you like to edit the configuration? (yes/no) [n]:
Use this configuration and save it? (yes/no) [y]:
[########################################] 100% Copy complete.
```

Enable Appropriate Cisco Nexus 9000 Features and Settings

Cisco Nexus 9000 A and Cisco Nexus 9000 B

To enable the IP switching feature and set default spanning tree behaviors, complete the following steps:

1. On each Nexus 9000, enter the configuration mode:

```
config terminal
```

2. Use the following commands to enable the necessary features:

```
feature lacp
feature vpc
```

```
feature interface-vlan
```

3. Configure the spanning tree and save the running configuration to start-up:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
copy run start
```

Create VLANs for VersaStack IP Traffic

Cisco Nexus 9000 A and Cisco Nexus 9000 B

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

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

```
vlan <IB-Mgmt VLAN id>
name IB-MGMT-VLAN
vlan <Native VLAN id>
name Native-VLAN
vlan <vMotion VLAN id>
name vMotion-VLAN
vlan <VM Traffic VLAN id>
name VM-Traffic-VLAN
vlan <iSCSI-a_VLAN_id>
name iSCSI-A-VLAN
vlan <iSCSI-B VLAN id>
name iSCSI-B-VLAN
exit
copy run start
```

Configure Virtual Port Channel Domain

Cisco Nexus 9000 A

To configure vPC domain for switch A, complete the following steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain 10
```

2. Make the Nexus 9000A the primary vPC peer by defining a low priority value:

```
role priority 10
```

3. Use the management interfaces on the supervisors of the Nexus 9000s to establish a keepalive link:

```
peer-keepalive destination <Mgmt. IP address for Switch B>  source <Mgmt. IP
address for Switch A>
```

4. Enable the following features for this vPC domain:

```
peer-switch

delay restore 150

peer-gateway

ip arp synchronize

auto-recovery

copy run start
```

Cisco Nexus 9000 B

To configure the vPC domain for switch B, complete the following steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain 10
```

2. Make the Nexus 9000A the primary vPC peer by defining a low priority value:

```
role priority 20
```

3. Use the management interfaces on the supervisors of the Nexus 9000s to establish a keepalive link:

```
peer-keepalive destination <Mgmt. IP address for Switch A> source <Mgmt. IP
address for Switch B>
```

4. Enable the following features for this vPC domain:

```
peer-switch

delay restore 150

peer-gateway

ip arp synchronize

auto-recovery

copy run start
```

Configure Network Interfaces for the vPC Peer Links

To configure the network interfaces for the vPC Peer links, complete the following steps:

Cisco Nexus 9000 A

1. Define a port description for the interfaces connecting to vPC Peer <var_nexus_B_hostname>>.

```
interface Eth1/47
description VPC Peer <Nexus-B Switch Name>:1/47
interface Eth1/48
description VPC Peer <Nexus-B Switch Name>:1/48
```

2. Apply a port channel to both vPC Peer links and bring up the interfaces.

```
interface Eth1/47,Eth1/48
channel-group 10  mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var_nexus_B_hostname>>.

```
interface Po10
description vPC peer-link
```

4. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLAN.

```
switchport
switchport mode trunk
switchport trunk native vlan <Native VLAN id>
switchport trunk allowed vlan <IB-MGMT VLAN id>, <vMotion VLAN id>, <VM
Traffic VLAN id>, <iSCSI-A VLAN id>, <iSCSI-B VLAN id>
```

5. Make this port-channel the VPC peer link and bring it up.

```
vpc peer-link
no shutdown
copy run start
```

Cisco Nexus 9000 B

1. Define a port description for the interfaces connecting to VPC Peer <var_nexus_A_hostname>>.

```
interface Eth1/47
description VPC Peer <Nexus-A Switch Name>:1/47
interface  Eth1/48
description VPC Peer <Nexus-A Switch Name>:1/48
```

2. Apply a port channel to both VPC Peer links and bring up the interfaces.

```
interface Eth1/47,Eth1/48

channel-group 10 mode active

no shutdown
```

3. Define a description for the port-channel connecting to <<var_nexus_A_hostname>>.

```
interface Po10

description vPC peer-link
```

4. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLAN.

```
switchport

switchport mode trunk

switchport trunk native vlan <Native VLAN id>

switchport trunk allowed vlan <IB-MGMT VLAN id>, <vMotion VLAN id>, <VM
Traffic VLAN id>, <iSCSI-A VLAN id>, <iSCSI-B VLAN id>

Make this port-channel the VPC peer link and bring it up.

vpc peer-link no shutdown

copy run start
```

Configure Network Interfaces to Cisco UCS Fabric Interconnect

Cisco Nexus 9000 A

1. Define a description for the port-channel connecting to <<var_ucs_clustername>>-A.

```
interface Po13

description <UCS Cluster Name>-A
```

2. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport

switchport mode trunk

switchport trunk native vlan <Native VLAN id>

switchport trunk allowed vlan <IB-MGMT VLAN id>, <vMotion VLAN id>, <VM
Traffic VLAN id>, <iSCSI-A VLAN id>, <iSCSI-B VLAN id>
```

3. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

4. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

5. Make this a VPC port-channel and bring it up.

```
vpc 13
no shutdown
```

6. Define a port description for the interface connecting to <UCS Cluster Name>-A.

```
interface Eth1/53
description <UCS Cluster Name>-A:35
```

7. Apply it to a port channel and bring up the interface.

```
channel-group 13 force mode active
no shutdown
```

8. Define a description for the port-channel connecting to <UCS Cluster Name>-B.

```
interface Po14
description <UCS Cluster Name>-B
```

9. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport
switchport mode trunk
switchport trunk native vlan <Native VLAN id>
switchport trunk allowed vlan <IB-MGMT VLAN id>, <vMotion VLAN id>, <VM
Traffic VLAN id>, <iSCSI-A VLAN id>, <iSCSI-B VLAN id>
```

10. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

11. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

12. Make this a VPC port-channel and bring it up.

```
vpc 14
no shutdown
```

13. Define a port description for the interface connecting to <UCS Cluster Name>-B.

```
interface Eth1/54
description <UCS Cluster Name>-B:1/35
```

14. Apply it to a port channel and bring up the interface.

```
channel-group 14 force  mode active
no shutdown
copy run start
```

Cisco Nexus 9000 B

1. Define a description for the port-channel connecting to <UCS Cluster Name>-B.

```
interface Po13
description <UCS Cluster Name>-A
```

2. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport
switchport mode trunk
switchport trunk native vlan <Native VLAN id>
switchport trunk allowed vlan <IB-MGMT VLAN id>, <vMotion VLAN id>, <VM
Traffic VLAN id>, <iSCSI-A VLAN id>, <iSCSI-B VLAN id>
```

3. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

4. Set the MTU to 9216 to support jumbo frames.

```
mtu 9216
```

5. Make this a VPC port-channel and bring it up.

```
vpc 13
no shutdown
```

6. Define a port description for the interface connecting to <UCS Cluster Name>-B.

```
interface Eth1/53
description <UCS Cluster Name>-A:1/36
```

7. Apply it to a port channel and bring up the interface.

```
channel-group 13 force mode active
no shutdown
```

8. Define a description for the port-channel connecting to <UCS Cluster Name>-A.

```
interface Po14
```

```
description <UCS Cluster Name>-B
```

9. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport

switchport mode trunk

switchport trunk native vlan <Native VLAN id>

switchport trunk allowed vlan <IB-MGMT VLAN id>, <vMotion VLAN id>, <VM
Traffic VLAN id>, <iSCSI-A VLAN id>, <iSCSI-B VLAN id>
```

10. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

11. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

12. Make this a VPC port-channel and bring it up.

```
vpc 14

no shutdown
```

13. Define a port description for the interface connecting to <UCS Cluster Name>-A.

```
interface Eth1/54

description <UCS Cluster Name>-B:1/36
```

14. Apply it to a port channel and bring up the interface.

```
channel-group 14 force  mode active

no shutdown

copy run start
```

Configure Network Interfaces Connected to IBM SVC iSCSI Ports

Cisco Nexus 9000 A

1. Define a description for the Ethernet port connecting to <SVC Cluster Node1, P1>

```
interface Ethernet1/19

description <SVC-Clus-Node1-iSCSI-P1>
```

2. Make the Interface a access port, and configure the switchport access VLAN.

```
switchport mode access

switchport access vlan <iSCSI-A VLAN id>
```

3. Make the interface spanning normal.

```
spanning-tree port type normal
```

4. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
copy run start
```

5. Define a description for the Ethernet port connecting to <SVC Cluster Node2, P1>

```
interface Ethernet1/20
description <SVC-Clus-Node2-iSCSI-P1>
```

6. Make the Interface a access port, and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-A VLAN id>
```

7. Make the interface spanning normal.

```
spanning-tree port type normal
```

8. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
copy run start
```

Cisco Nexus 9000 B

1. Define a description for the Ethernet port connecting to <SVC Cluster Node1, P2>

```
interface Ethernet1/19
description <SVC-Clus-Node1-iSCSI-P2>
```

2. Make the Interface a access port, and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-B VLAN id>
```

3. Make the interface spanning normal.

```
spanning-tree port type normal
```

4. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

```
no shutdown
copy run start
```

5. Define a description for the Ethernet port connecting to <SVC Cluster Node1, P2>

```
interface Ethernet1/20
description <SVC-Clus-Node2-iSCSI-P2>
```

6. Make the Interface a access port, and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-B VLAN id>
```

7. Make the interface spanning normal.

```
spanning-tree port type normal
```

8. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
copy run start
```

Management Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the VersaStackPod environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the VersaStack environment into the infrastructure. The following procedure can be used to create an uplink vPC to the existing environment.

Cisco Nexus 9000 A and B using Port Channel Example

To enable management access across the IP switching environment leveraging port channel in config mode run the following commands:

1. Define a description for the port-channel connecting to management switch.

```
interface po6
description IB-MGMT
```

2. Configure the port as an access VLAN carrying the InBand management VLAN traffic.

```
switchport
switchport mode access
switchport access vlan <IB-MGMT VLAN id>
```

3. Make the port channel and associated interfaces normal spanning tree ports.

```
spanning-tree port type normal
```

4. Make this a VPC port-channel and bring it up.

```
vpc 6

no shutdown
```

5. Define a port description for the interface connecting to the management plane.

```
interface Eth1/33

description IB-MGMT-SWITCH_uplink
```

6. Apply it to a port channel and bring up the interface.

```
channel-group 6 force mode active

no shutdown
```

7. Save the running configuration to start-up in both Nexus 9000s and run commands to look at port and port channel.

```
Copy run start

sh int eth1/33 br

sh port-channel summary
```

Cisco MDS 9396S Initial Configuration Setup

Setup the initial configuration on Cisco MDS 9393S switches; this procedure is similar to setting up the Nexus switches detailed in the Nexus initial setup section. Configure the switch name, management IP address, Netmask, Gateway and other settings such as NTP and Clock and proceed to the next section.

Enable Appropriate Cisco MDS Features and Settings

Cisco MDS A and B

1. To enable the feature on both switches, enter the following commands:

```
Config
feature npiv
feature fport-channel-trunk

exit

copy running start
```

Initial Storage Configuration

IBM FlashSystem 900 Base Configuration

This section covers initial setup for IBM FS 900. Configuring the IBM FlashSystem FS900 is a two-stage setup. A USB key and the IBM setup software will be used for the initial configuration and IP assignment, and the Spectrum Virtualize web interface will be used to complete the configuration.

Prerequisites

Begin this procedure only after the physical installation of the FlashSystem 900 has been completed. The computer used to initialize the FlashSystem 900 must have a USB port and a network connection to the FlashSystem 900. An USB key was included with the FlashSystem 900 in order to run the setup. This is provided with the original shipment but the software can be download and copied to a blank, replacement USB drive if required, see below.

Creating a Replacement USB Key

In the event, original USB Key is unavailable, a replacement USB key can be recreated using the steps below to download the System Initialization software from IBM Fix Central support web site.

1. Go to https://www.ibm.com/support/fixcentral/.

2. Enter the information below into the Find Product search tool.

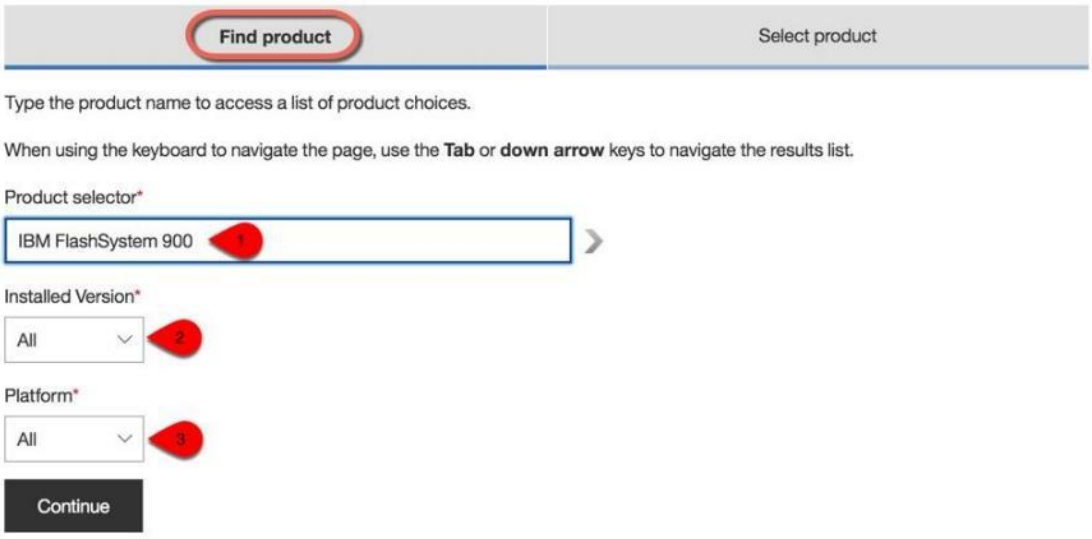

3. Click Continue.

4. Scroll down the results page and locate the latest firmware release for the FlashSystem 900 as depicted below.

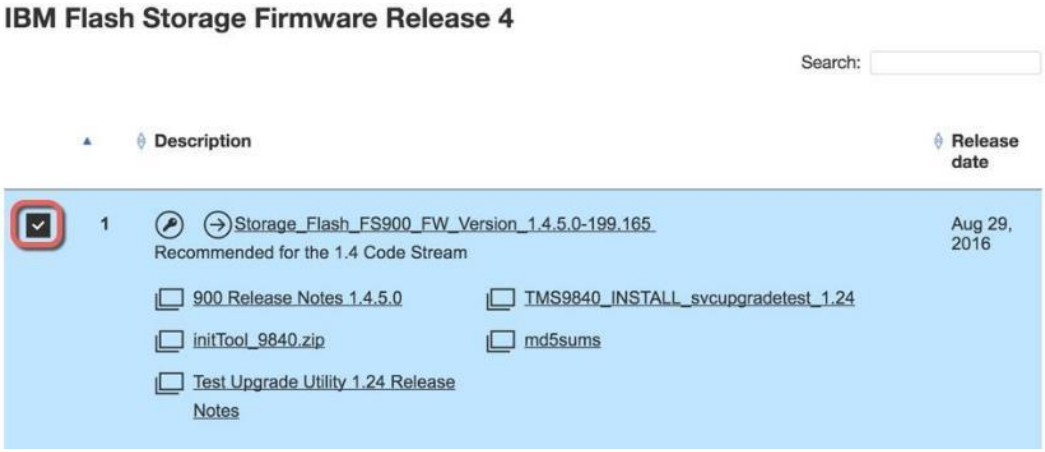

5. Click Continue to proceed to the download page.

You will need your IBM login account to download software.

6. Scroll down the web page and select the hyper link for the InitTool as detailed below.

7. Extract the contents of the zip archive file to the root directory of any USB key formatted with a FAT32, ext2 or ext3 file system.

IBM FlashSystem 900 Initial Configuration

To complete this initial configuration process, access to the powered on FlashSystem 900, the USB flash drive that was shipped with the system, the network credentials (IP Address, subnet and gateway) of the system, and a personal computer is required.

1. Run the System Initialization tool from the USB key. For Windows clients run the InitTool.bat located in the root directory of the USB key. For MAC, Red Hat, and Ubuntu, locate the root directory of the USB flash drive, such as /Volumes/ and type: sh InitTool.sh.

2. Click Next to continue.

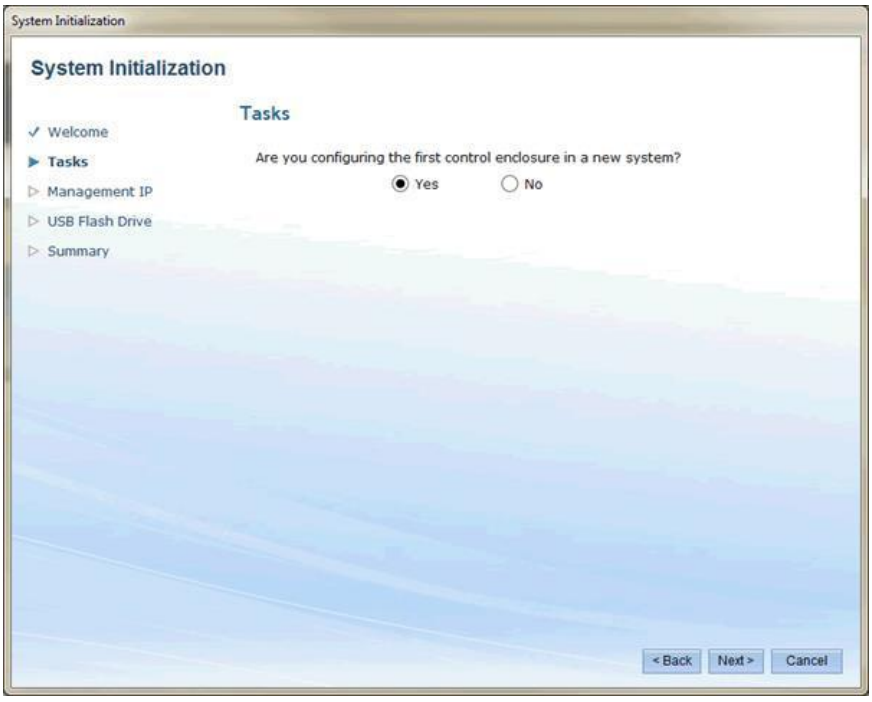

3. Select Yes to Are you configuring the first control enclosure in a new system? and click Next.

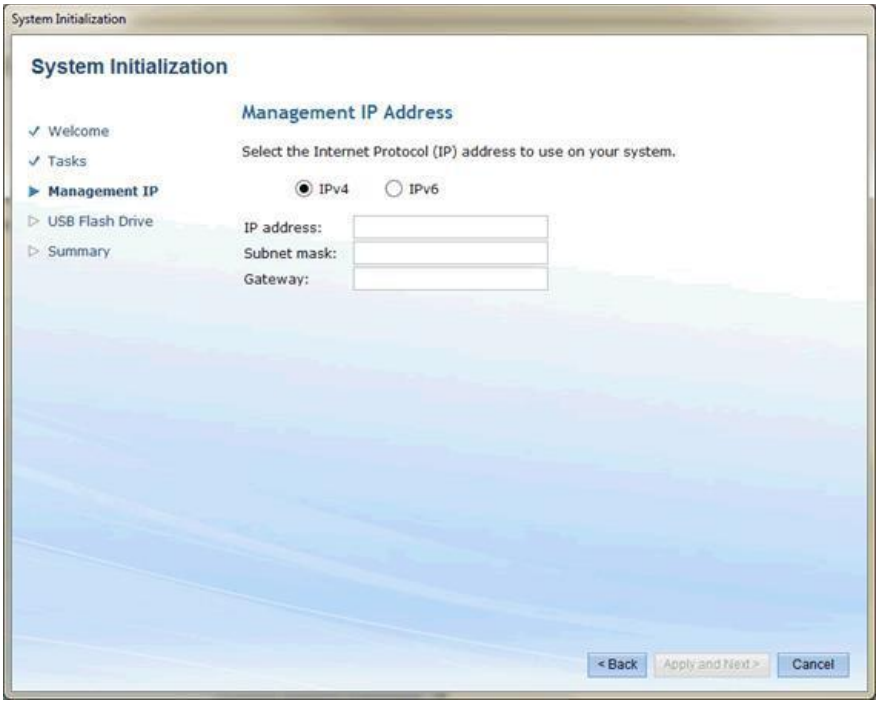

4. Input the IP address for your FlashSystem 900 as well as the required subnet mask and gateway. Click Apply, then click Next.

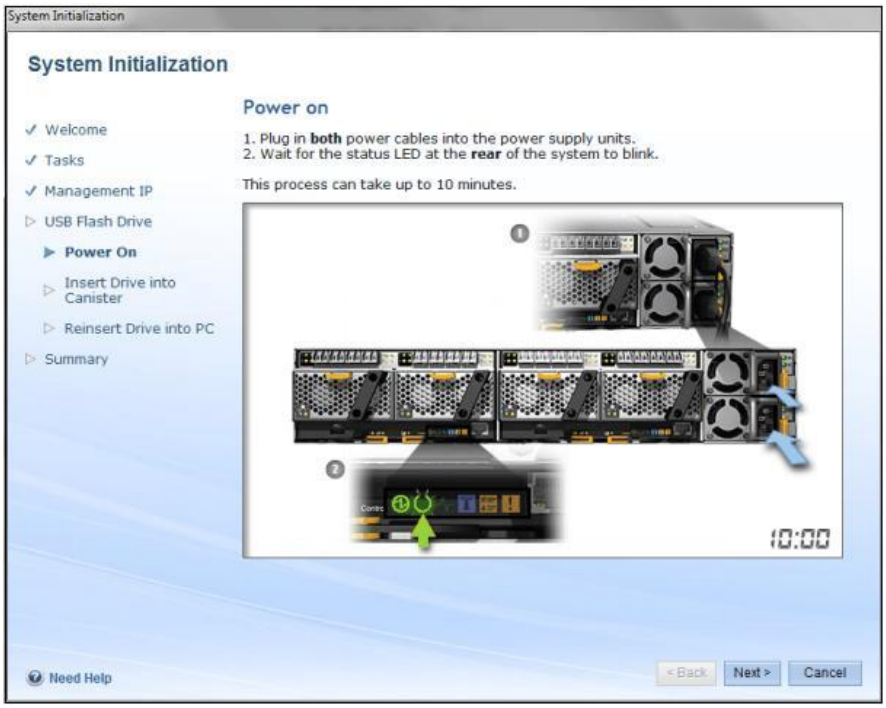

5. Connect-up both power supply units, as depicted above. Wait for the status LED to come on, flash, and then come on solid. The process can take up to 10 minutes.

6. Safely eject/remove the USB key from the computer and insert into the left USB port on the FlashSystem 900 as pictured above. The blue Identity LED will turn on, then off. This process can take up to 3 minutes. Click Next.

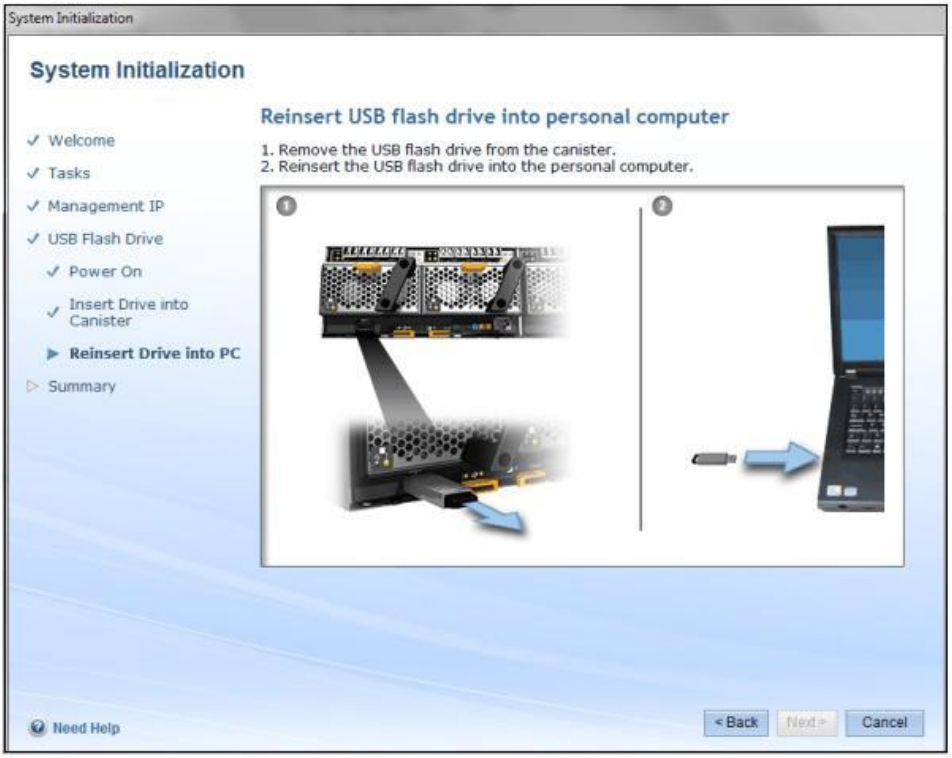

7. Remove the USB key from the FlashSystem 900 and reinsert into the computer.

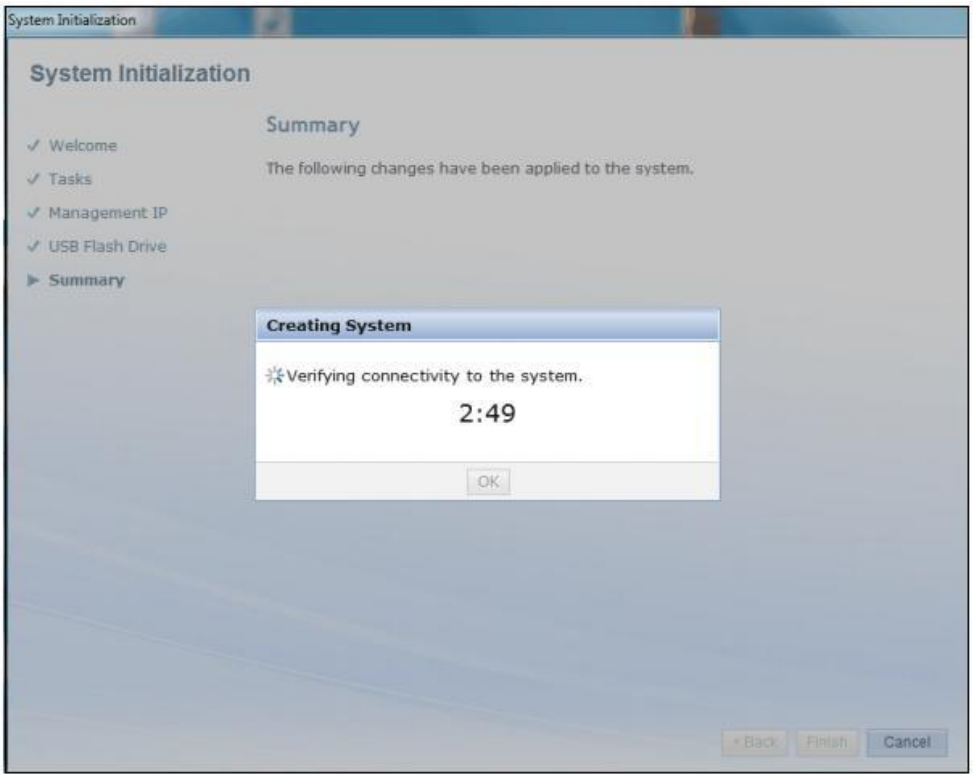

8. If the Ethernet ports of the newly initialized FlashSystem 900 are attached to the same network as the computer where InitTool was run, InitTool checks connectivity to the system and displays the result of the system initialization process.

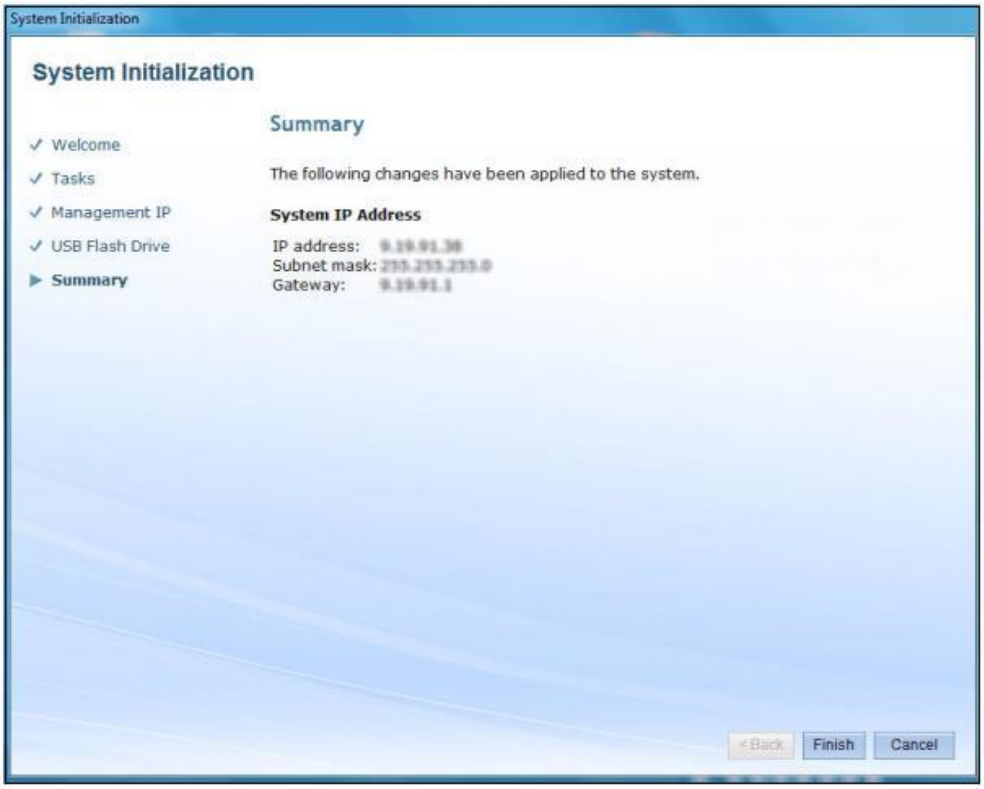

9. The initialization software indicates that the operation has completed successfully. Click Finish. FlashSystem 900 management GUI should now be available for further setup.

IBM FlashSystem 900 GUI Setup

After completing the initial tasks above, launch the management GUI, and configure the IBM FlashSystem 900. At the time of writing, the following browsers (or later) are supported with the management GUI; Firefox 32, Internet Explorer 10 and Google Chrome 37.

The IBM Redbook publication provides in-depth knowledge of the IBM FlashSystem 900 product architecture, software and hardware, implementation, and hints and tips: Implementing IBM FlashSystem 900

1. Log in to the management GUI of the FlashSystem 900 using the IP address provided in the configuration steps above.

2. Log into the management GUI as superuser with the password of passw0rd. Click Log In.

3. The system will prompt to change the password for superuser. Make a note of the new password and then click Log In.

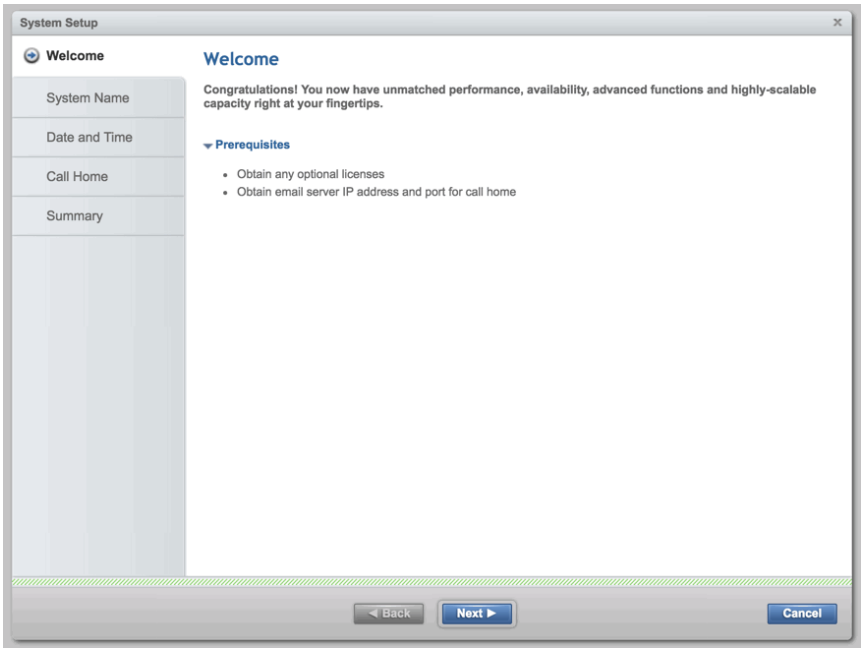

4. On the Welcome to System Setup screen click Next.

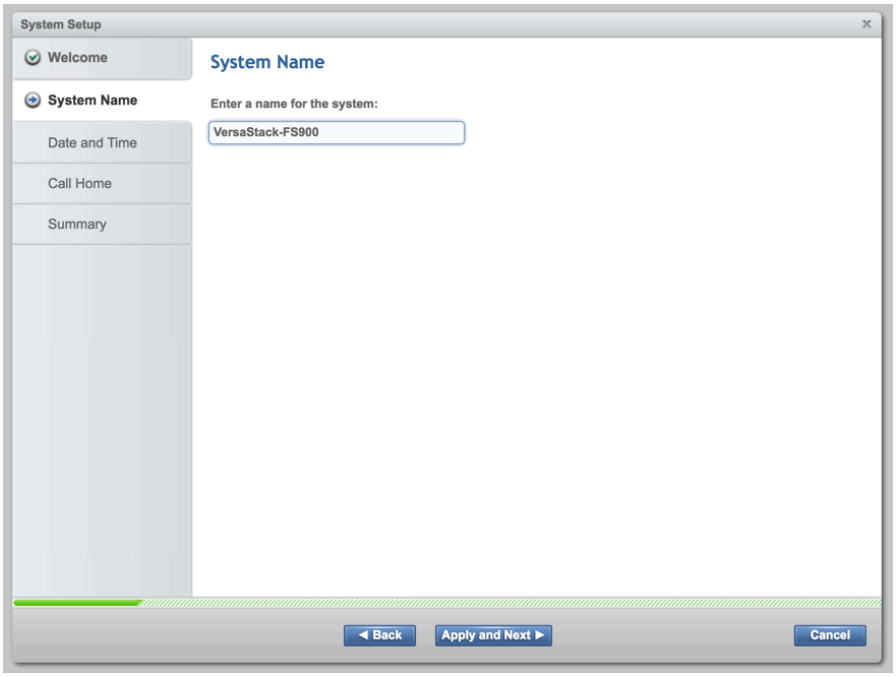

5. Enter the System Name and click Apply and Next to proceed.

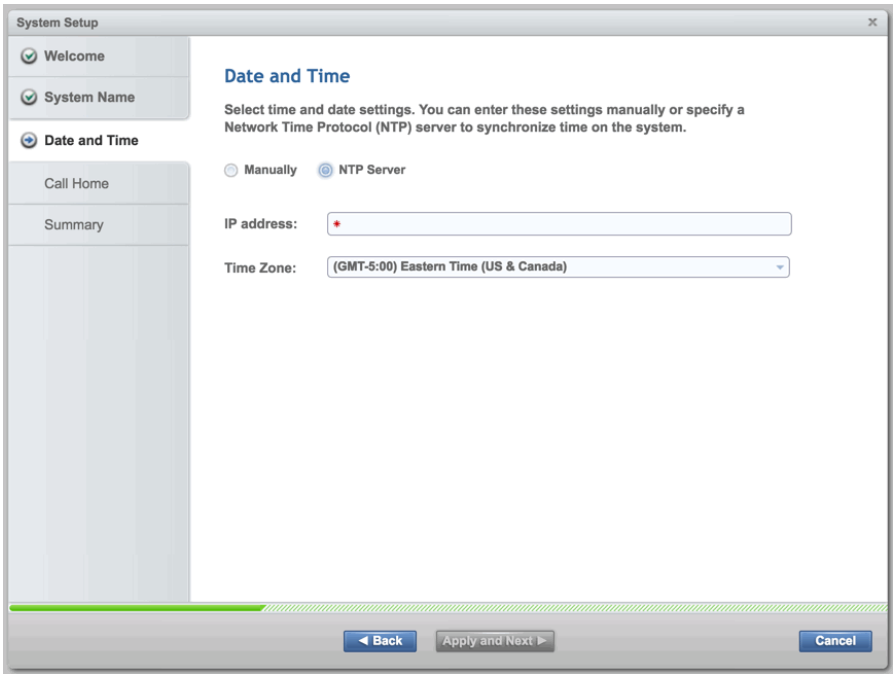

6. Configure the system date and time; where possible it's preferred practice is to configure the system with an NTP server. Click Apply and Next.

It is highly recommended to configure email event notifications, which will automatically notify IBM support centers when problems occur.

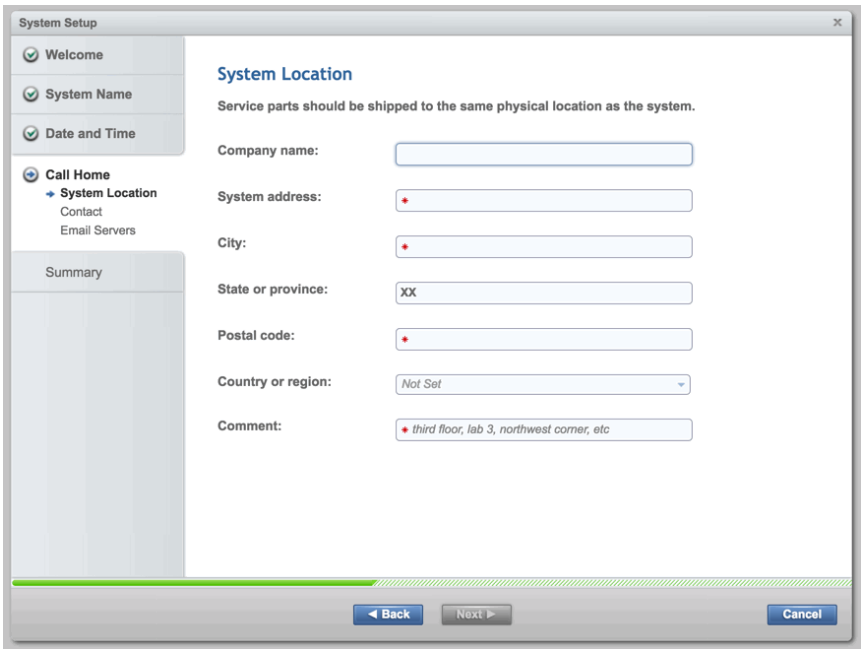

7. Enter the complete company name and address, and then click Next.

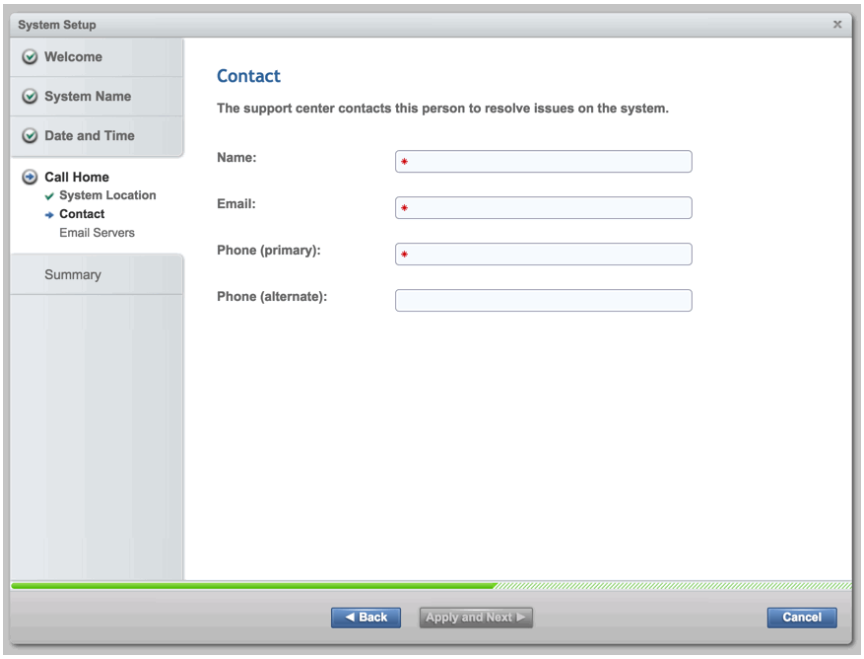

8. Enter the information of the contact person for the support center and click Apply and Next.

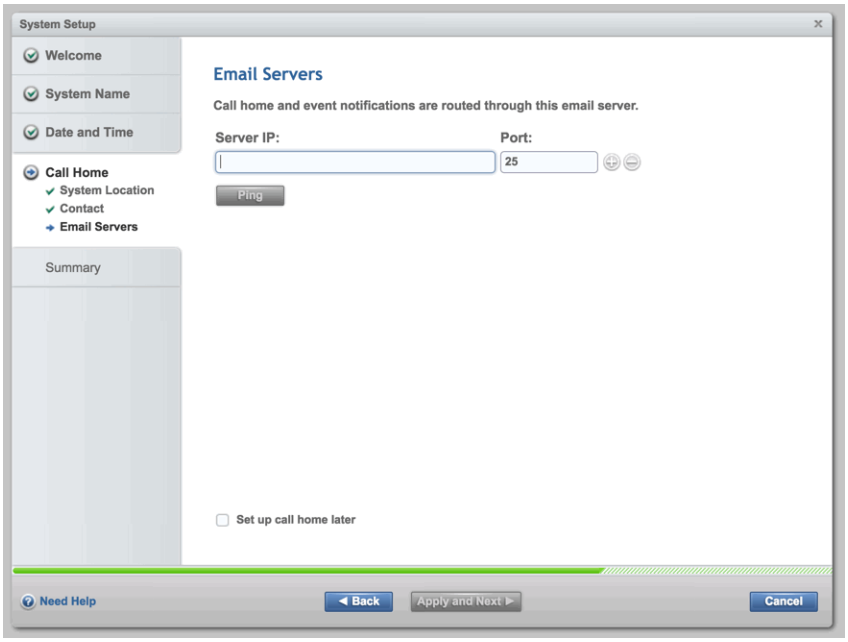

9. Enter the IP address and port for one or more of the email servers for the Call Home email notification.

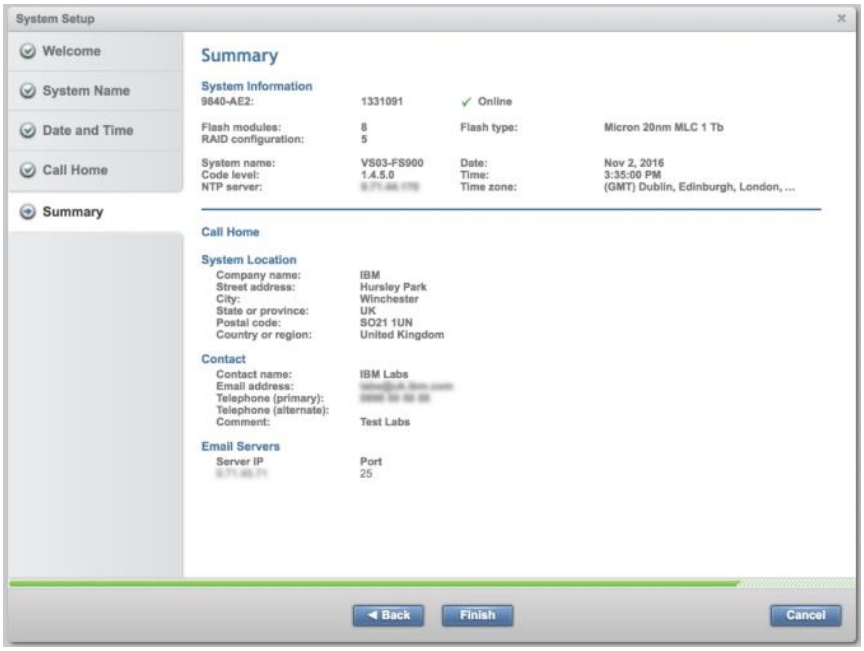

10. Review the final summary page, and click Finish to complete the System Setup wizard.

11. Setup Complete. Click Close

During system setup, FlashSystem 900 discovers the quantity and size of flash modules installed in the system. As the final step of the system setup procedure, on clicking Finish, a single RAID 5 array is created on these flash modules. One flash module is reserved to acts as an active spare while the remaining modules forms the RAID 5 array.

12. The System view for IBM FS900, as shown above, is now available.

13. Along the left side the management GUI, hover over each of the icons on the Navigation Dock to become familiar with the options.

14. Select the Settings icon from the Navigation Dock and choose Network.

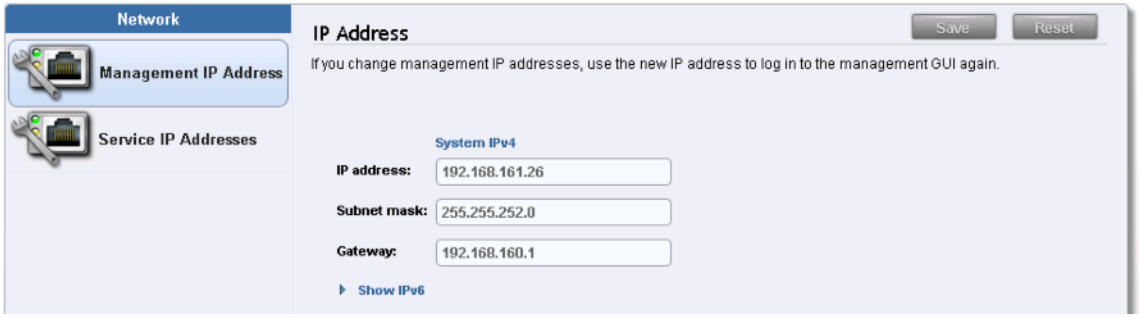

15. On the Network screen, highlight the Management IP Addresses section. Change the IP address if necessary and click OK. The application might need to close and redirect the browser to the new IP address.

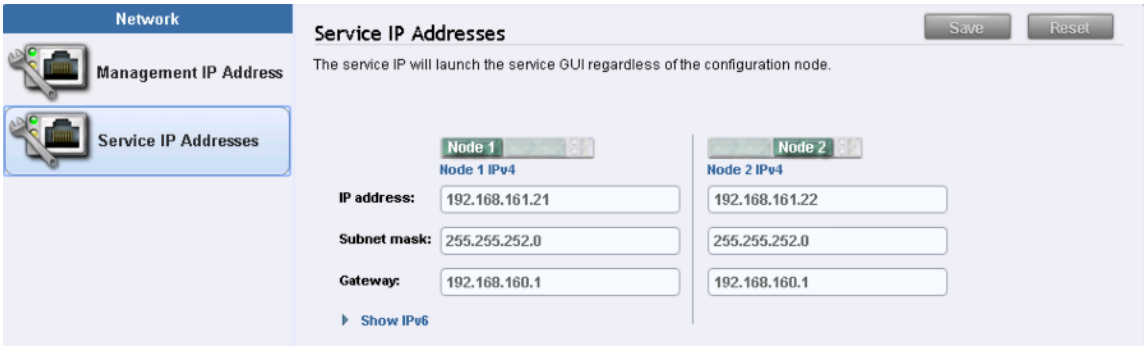

16. While still on the Network screen, select Service IP Addresses from the list on the left and change the service IP addresses for Node 1 and Node 2 as required.

17. Click the Access icon from the Navigation Dock on the left and select Users to access the Users screen.

18. Select Create User.

19. Enter a new name for an alternative admin account. Leave the SecurityAdmin default as the User Group, and input the new password, then click Create. Optionally, an SSH Public Key generated on a Unix server through the command “ssh-keygen -t rsa” can be copied to a public key file and associated with this user through the Choose File button.

20. Logout from the superuser account and log back in as the new account created during the last step.

21. Select Volumes from the Navigation Dock and then select Volumes.

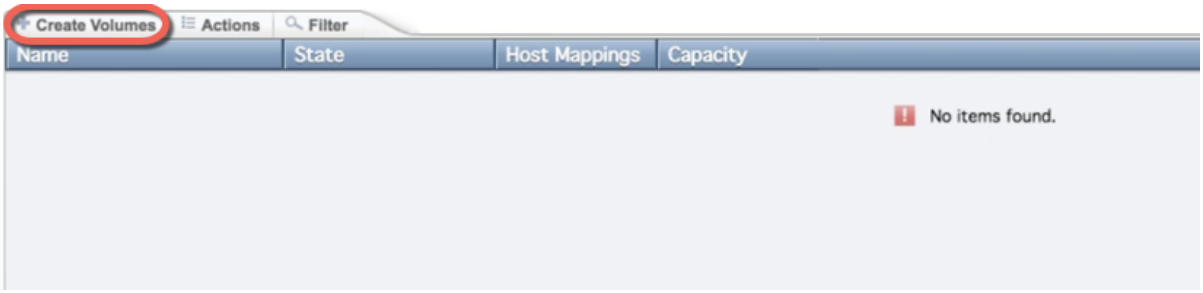

22. Click Create Volumes.

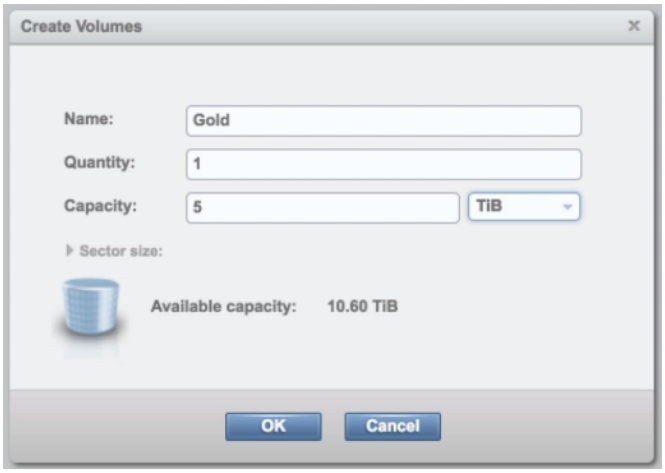

23. Create a volume to be used as the Gold storage tier. This volume will be used by the SAN volume controller to cater for high I/O and low latency, mission critical applications and data. Click OK.

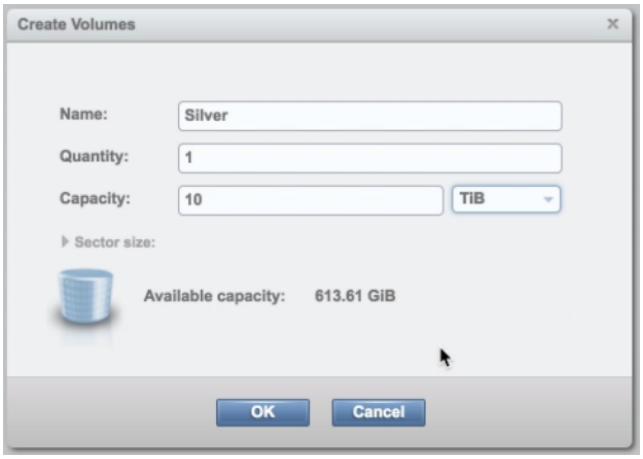

24. Click Create Volumes and repeat the process again to provision some of the flash storage for a hybrid storage pool. Using the SAN Volume Controller, the slower speed Enterprise disks from IBM Storwize V5000 will be combined with the flash modules provisioned here, to create an EazyTier storage pool.

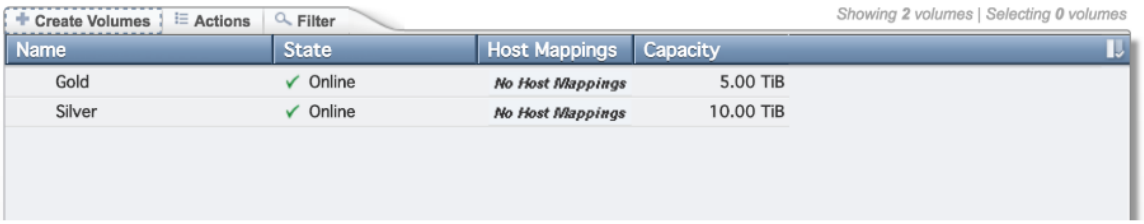

25. Validate the volumes you have created on your FlashSystem 900.

IBM Storwize V5030 Base Configuration

Configuring the IBM Storwize V5000 Second Generation is a two-stage setup. The technician port (T) will be used for the initial configuration and IP assignment, and the management GUI will be used to complete the configuration.

For a more in-depth look at installing the IBM Storwize V5000 Second Generation hardware, refer to the Redbook publication: Implementing the IBM Storwize V5000 Gen2

Prerequisites

Begin this procedure only after the physical installation of the IBM Storwize V5000 has been completed. The computer used to initialize the IBM Storwize V5000 must have an Ethernet cable connected to the technician port of the IBM Storwize V5000 as well as a supported browser installed. At the time of writing, the following browsers or later are supported with the management GUI; Firefox 32, Internet Explorer 10 and Google Chrome 37.

Do not connect the technician port to a switch. If a switch is detected, the technician port connection might shut down, causing a 746 node error.

IBM Storwize V5000 Initial Configuration

1. Power on the IBM Storwize V5000 control enclosure. Use the supplied power cords to connect both power supply units. The enclosure does not have power switches.

When deploying the expansion enclosures, expansion enclosures must be powered on before powering on the control enclosure.

2. From the rear of the control enclosure, check the LEDs on each node canister. The canister is ready with no critical errors when Power is illuminated, Status is flashing, and Fault is off. See figure below for reference.

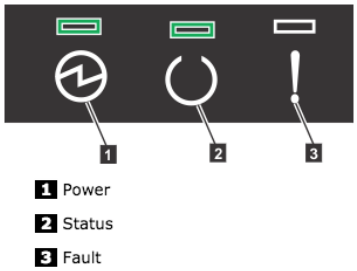

3. Configure an Ethernet port, on the computer used to connect to the control enclosure, to enable Dynamic Host Configuration Protocol (DHCP) configuration of its IP address and DNS settings.

If DHCP cannot be enabled, configure the PC networking as follows: specify the static IPv4 address 192.168.0.2, subnet mask 255.255.255.0, gateway 192.168.0.1, and DNS 192.168.0.1.

4. Locate the Ethernet port that is labelled T on the rear of the IBM Storwize V5000 node canister. For the IBM Storwize V5030 system, there is a dedicated technician port.

5. Connect an Ethernet cable between the port of the computer that is configured in step 3 and the technician port. After the connection is made, the system will automatically configure the IP and DNS settings for the personal computer if DHCP is available.

6. After the Ethernet port of the personal computer is connected, open a supported browser and browse to address http://install. (If DCHP is not enabled, open a supported browser and go to the following static IP address 192.168.0.1.) The browser will automatically be directed to the initialization tool.

In case of a problem due to a change in system states, wait 5 - 10 seconds and then try again.

7. Click Next on the System Initialization welcome message.

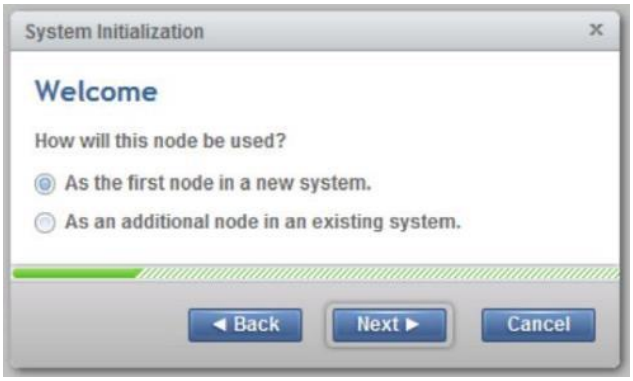

8. Click Next to continue with As the first node in a new system.

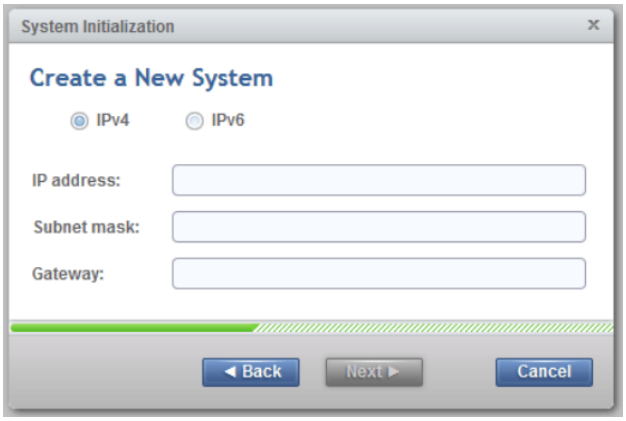

9. Complete all of the fields with the networking details for managing the system. This will be referred to as the System or Cluster IP address. Click Next.

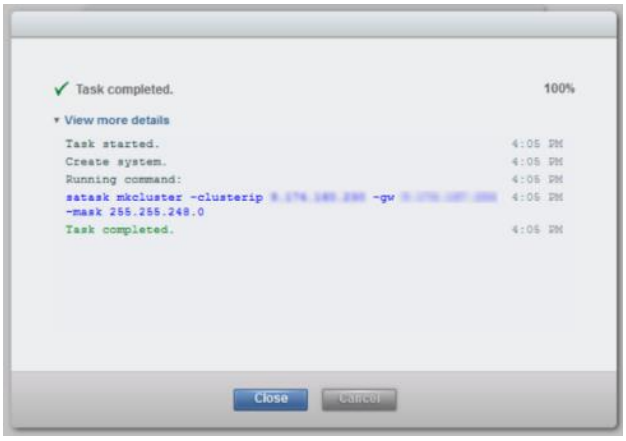

10. The setup task completes and you are provided a view of the generated satask CLI command as show above. Click Close. The storage enclosure will now reboot.

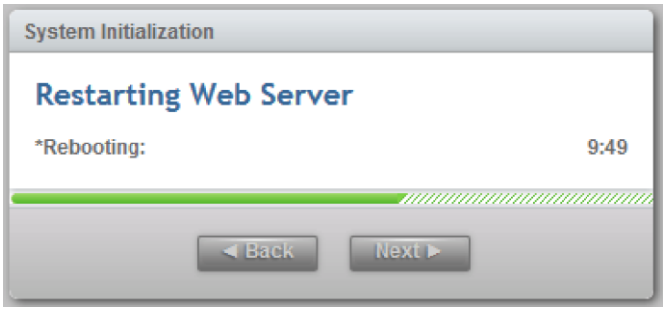

11. The system takes approximately 10 minutes to reboot and reconfigure the Web Server. After this time, click Next to proceed to the final step.

12. After completing the initialization process, disconnect the cable between the computer and the technician port, as instructed above. Re-establish the connection to the customer network and click Finish to be redirected to the management address provided during the configuration.

IBM Storwize V5000 GUI Setup

After completing the initial tasks above, launch the management GUI and continue configuring the IBM Storwize V5000 system.

Following e-Learning module introduces the IBM Storwize V5000 management interface and provides an overview of the system setup tasks, including configuring the system, migrating and configuring storage, creating hosts, creating and mapping volumes, and configuring email notifications: Getting Started

1. Log in to the management GUI using the cluster IP address configured above.

2. Read and accept the license agreement. Click Accept.

3. Login as superuser with the password of passw0rd. Click Log In.

The system will prompt to change the password for superuser. Change the password for superuser and make a note of the password. Click Log In.

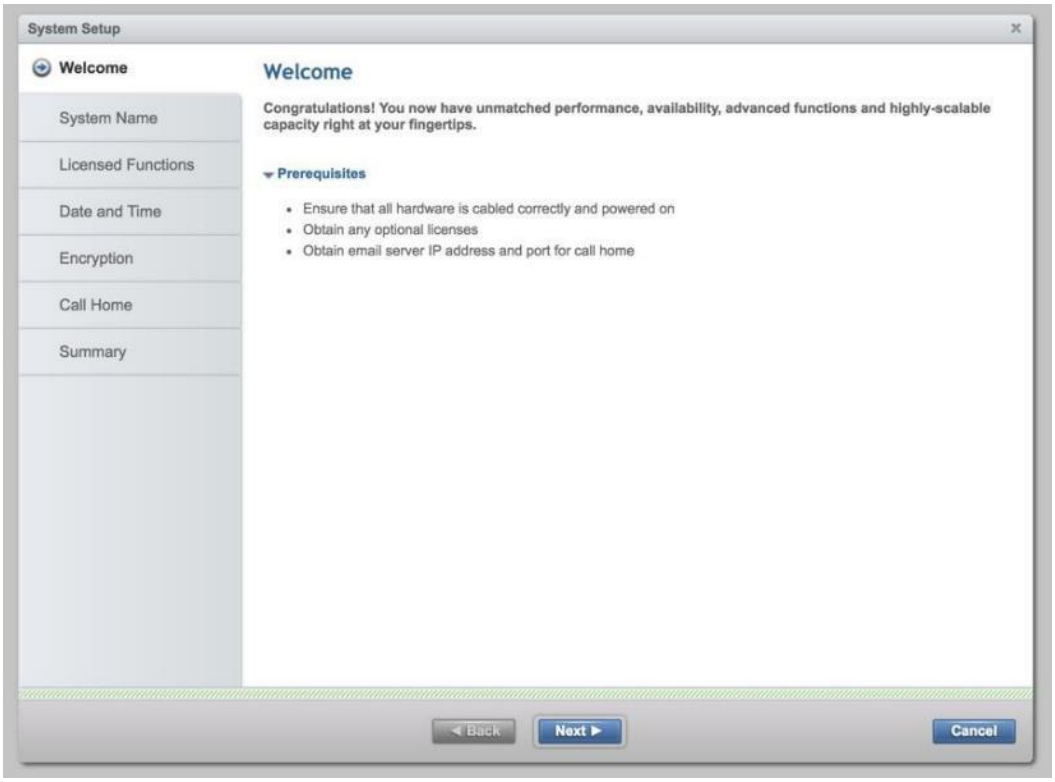

4. On the Welcome to System Setup screen click Next.

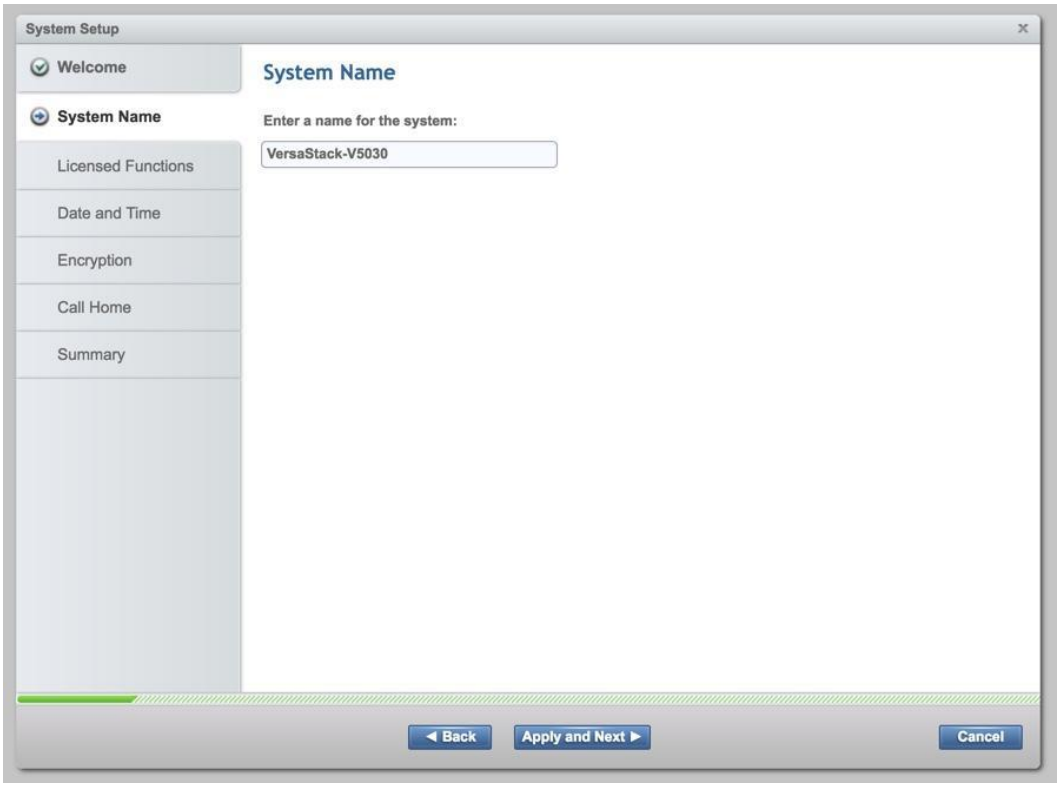

5. Enter the System Name and click Apply and Next to proceed.

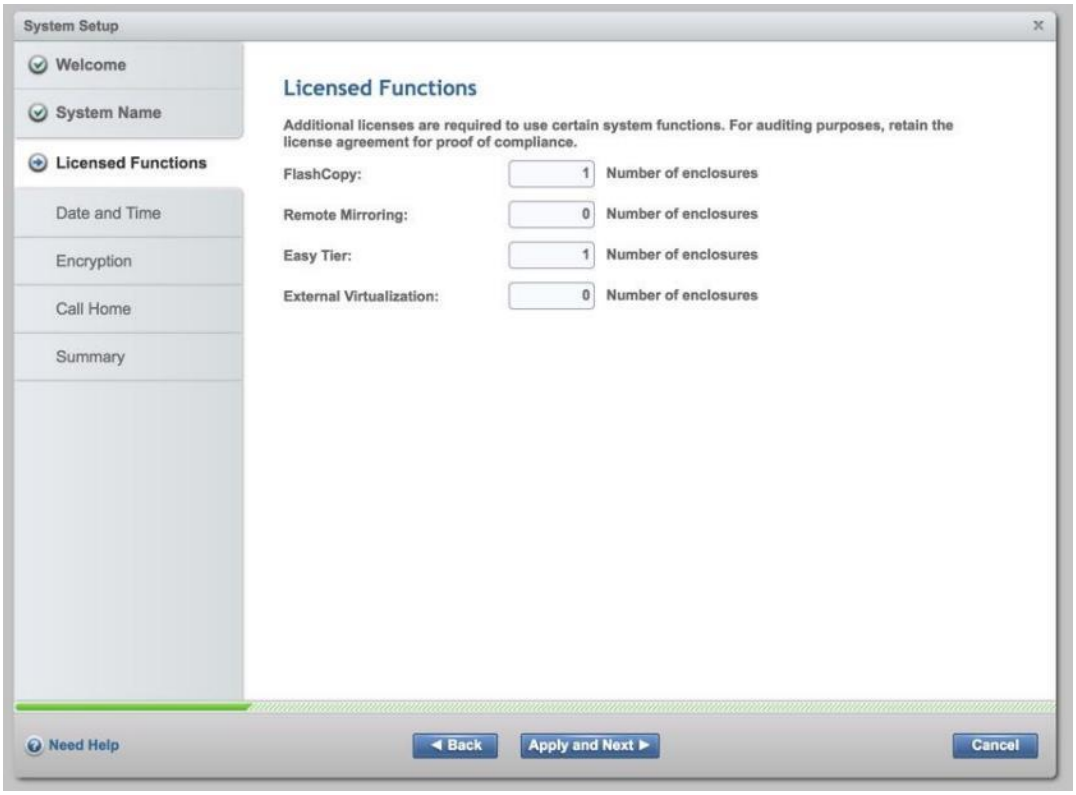

6. Select the license that was purchased, and enter the number of enclosures that will be used for Flash-Copy, Remote Mirroring, Easy Tier, and External Virtualization. Click Apply and Next to proceed.

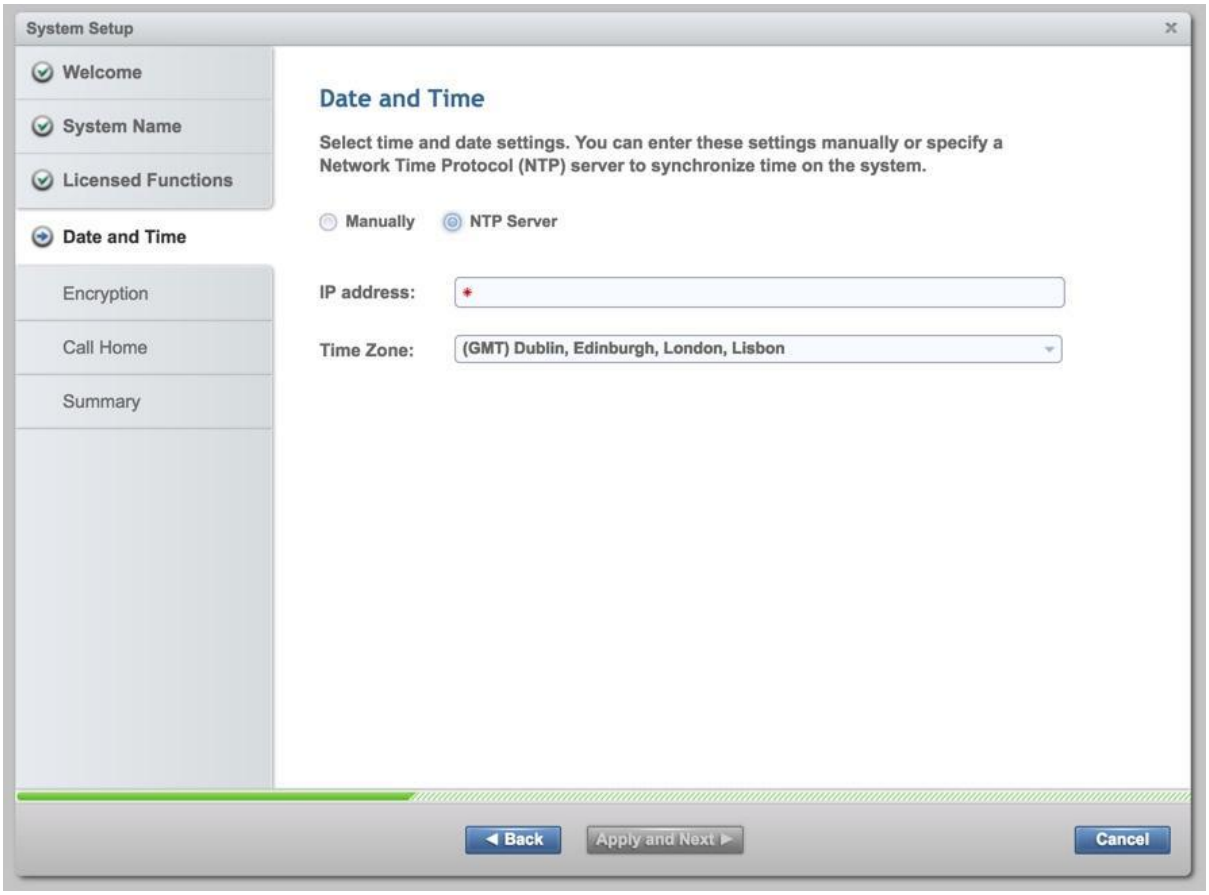

7. Configure the date and time settings, inputting NTP server details if available. Click Apply and Next to proceed.

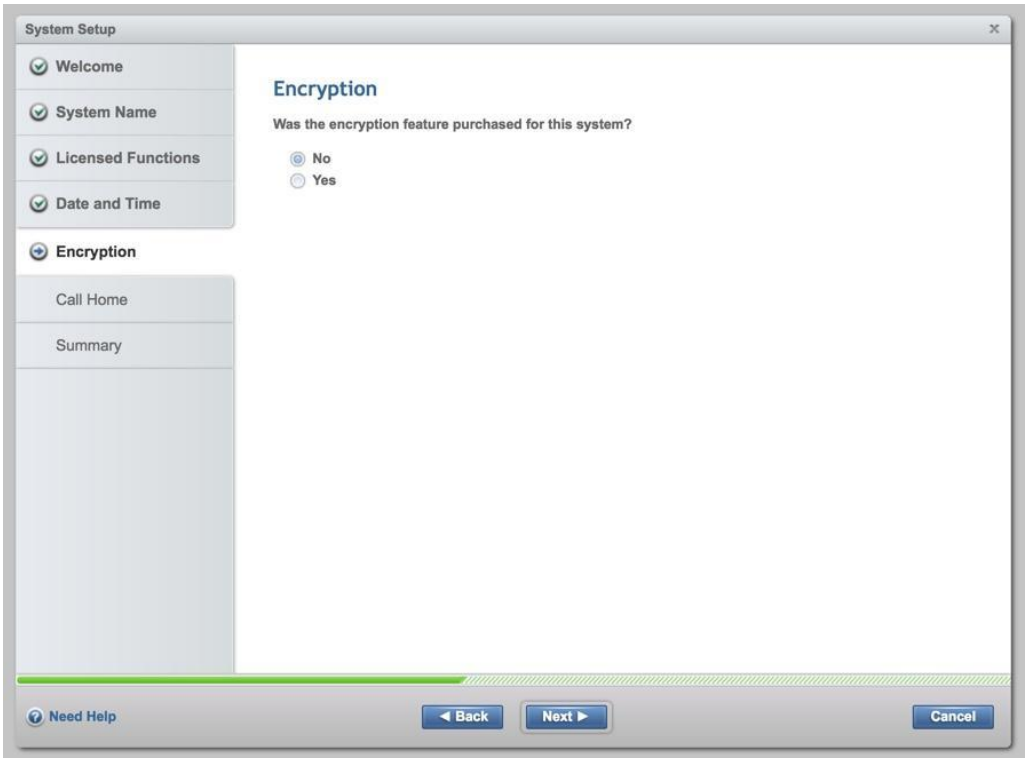

8. Enable the Encryption feature (or leave it disabled). Click Next to proceed.

It is highly recommended to configure email event notifications which will automatically notify IBM support centers when problems occur.

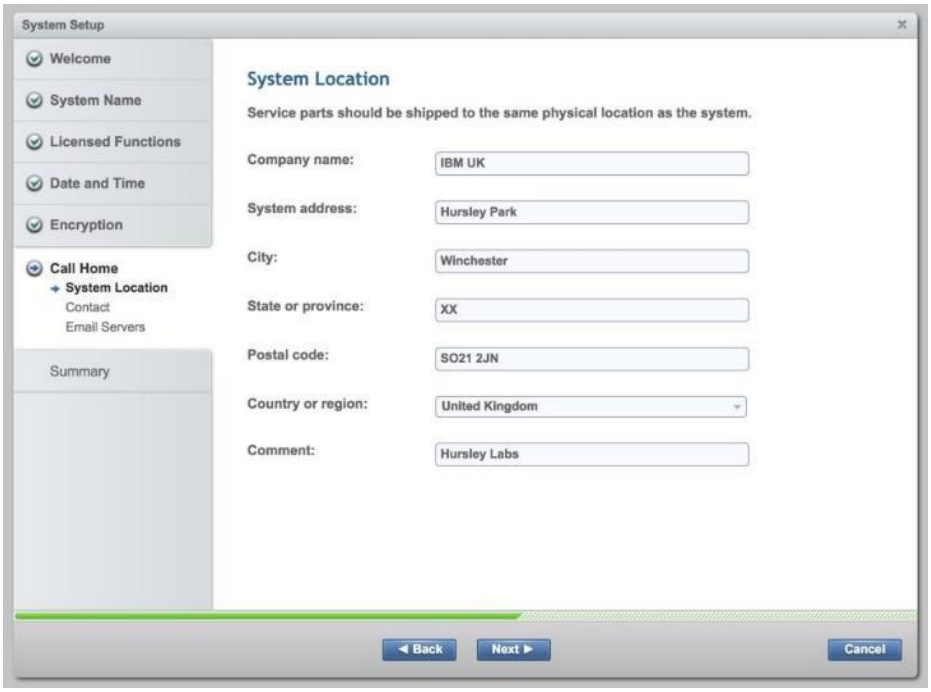

9. Enter the complete company name and address and then click Next.

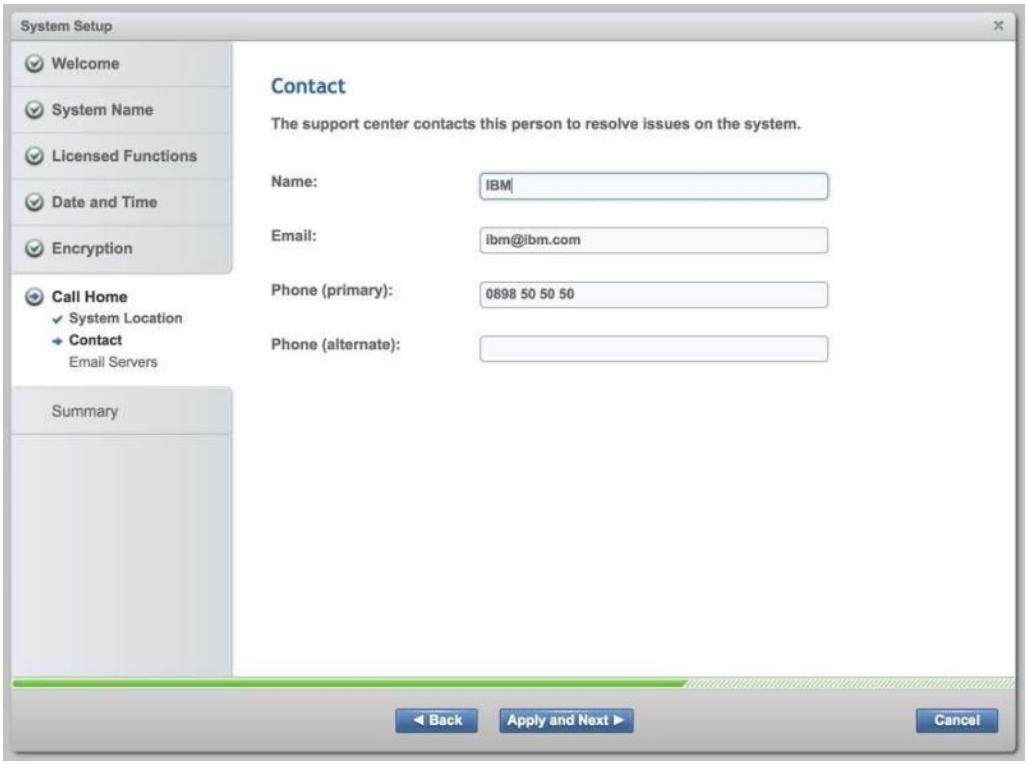

10. Enter the contact person for the support center calls. Click Apply and Next.

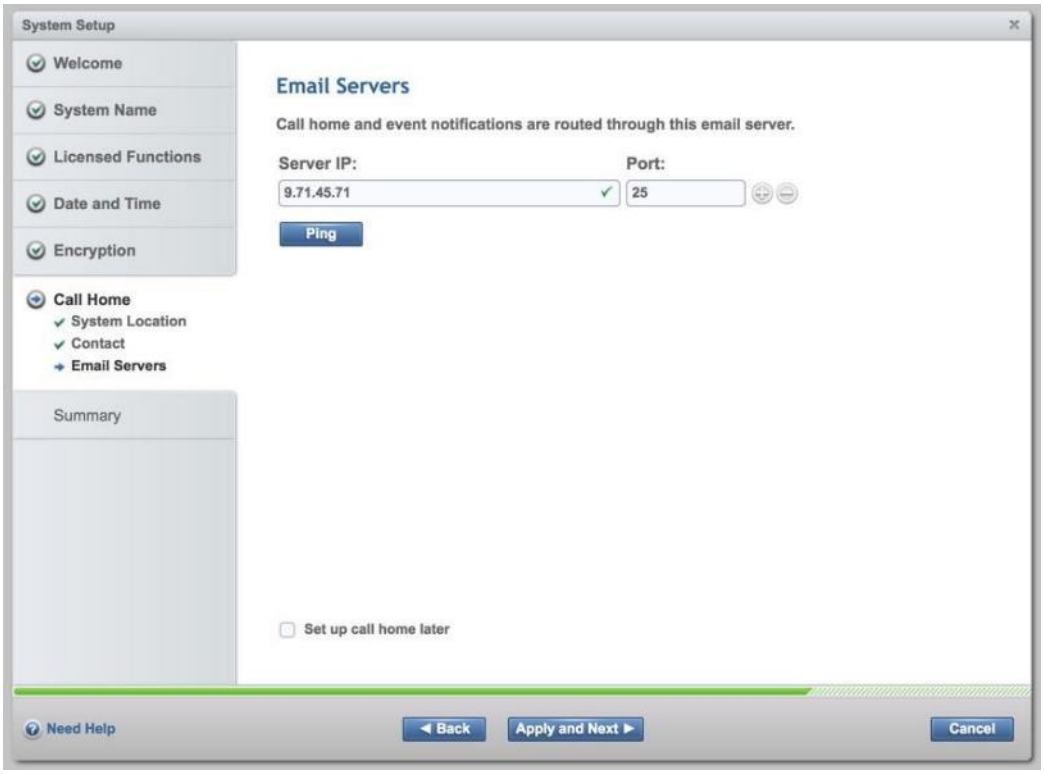

11. Enter the IP address and server port for one or more of the email servers for the Call Home email notification. Click Apply and Next.

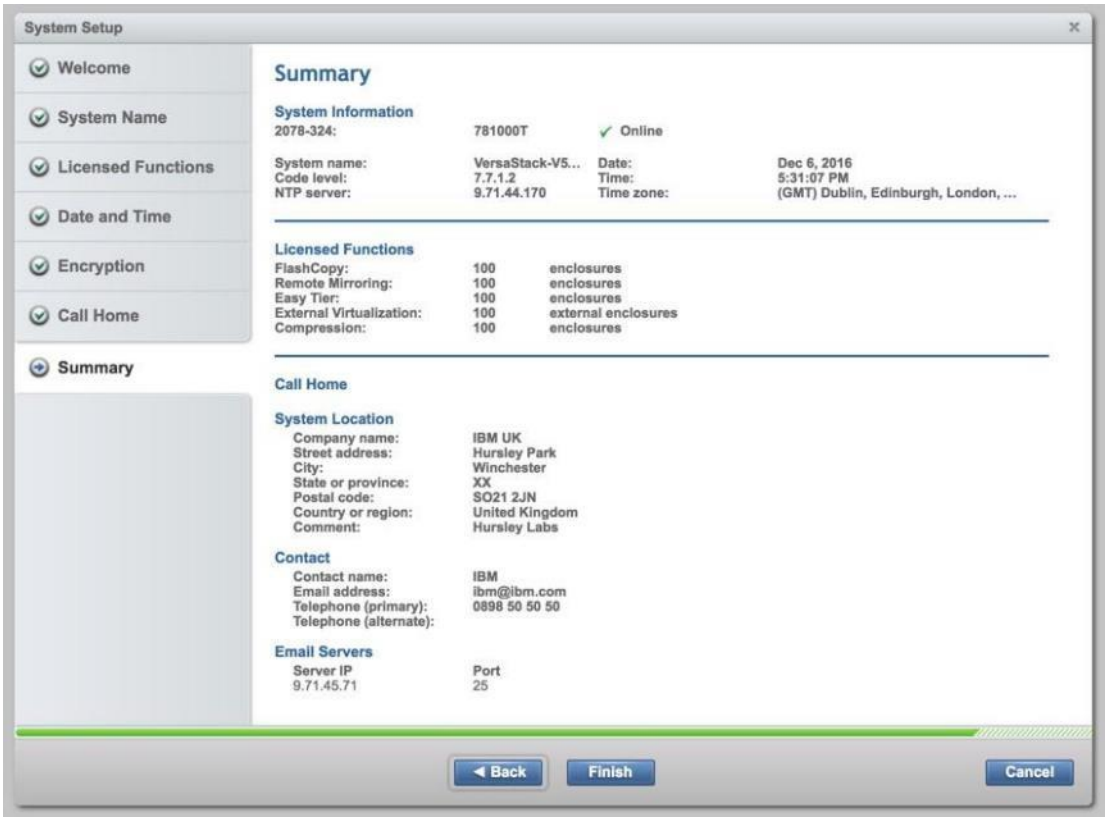

12. Review the final summary page and click Finish to complete the System Setup wizard.

13. Setup Completed. Click Close.

14. The System view of IBM Storwize V5000 is now available, as depicted above.

15. In the left side menu, hover over each of the icons on the Navigation Dock to become familiar with the options.

16. Select the Setting icon from the Navigation Dock and choose Network.

17. On the Network screen, highlight the Management IP Addresses section. Then click the number 1 interface on the left-hand side to bring up the Ethernet port IP menu. Change the IP address if necessary and click OK. The application might need to close and redirect the browser to the new IP address.

18. While still on the Network screen, select 1) Service IP Addresses from the list on the left and 2) Node Canister left then 3) change the IP address for port 1, click OK.

19. Repeat this process for port 1 on Node Canisters right (and port 2 left/right if you have cabled those ports).

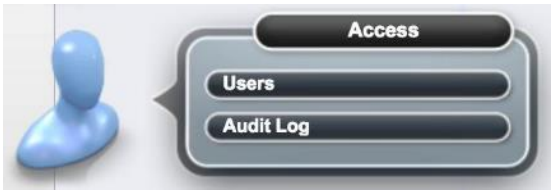

20. Click the Access icon from the Navigation Dock on the left and select Users to access the Users screen.

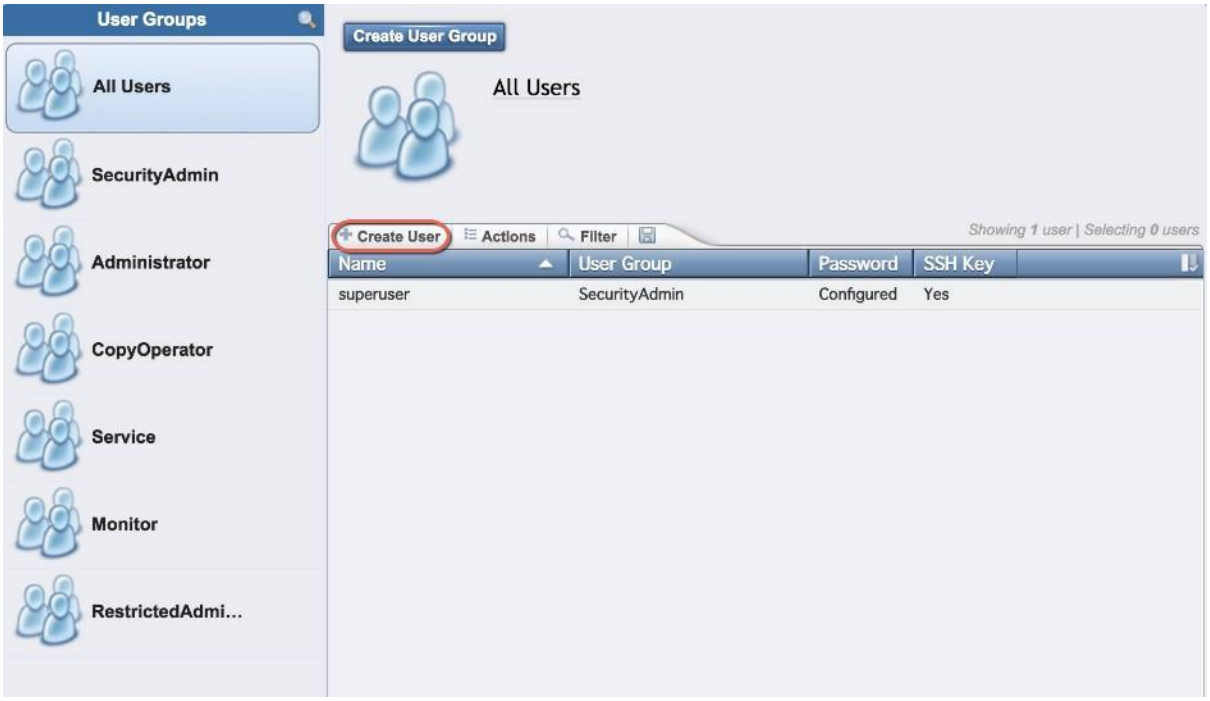

21. Select Create User.

22. Enter a new name for an alternative admin account. Leave the SecurityAdmin default as the User Group, and input the new password, then click Create. Optionally, an SSH Public Key generated on a Unix server through the command “ssh-keygen -t rsa” can be copied to a public key file and associated with this user through the Choose File button.

23. Logout from the superuser account and log back in as the new account.

24. Select Pools from the Navigation Dock and select MDisk by Pools.

25. Click Create Pool, and enter the name of the new storage pool. Click Create.

26. Select Add Storage.

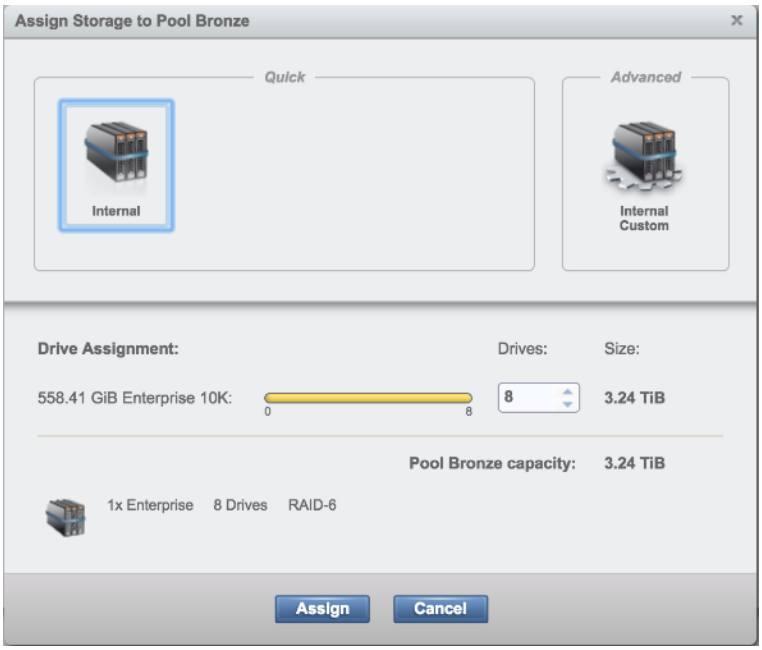

27. Select Internal, review the drive assignments and then select Assign.

Depending on customer configuration, select Internal Custom to manually create tired storage pools grouping together disk by capabilities. In this deployment, Flash and Enterprise class disk are utilized for Silver pool and Nearline disks are utilized for Bronze storage pool.

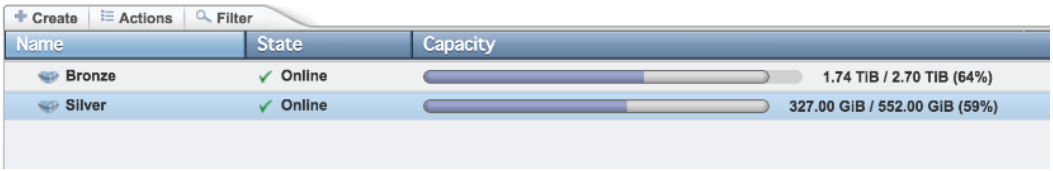

28. Validate the pools are online and have the relevant storage assigned.

29. Select Volumes from the Navigation Dock and then select Volumes.

30. Click Create Volumes.

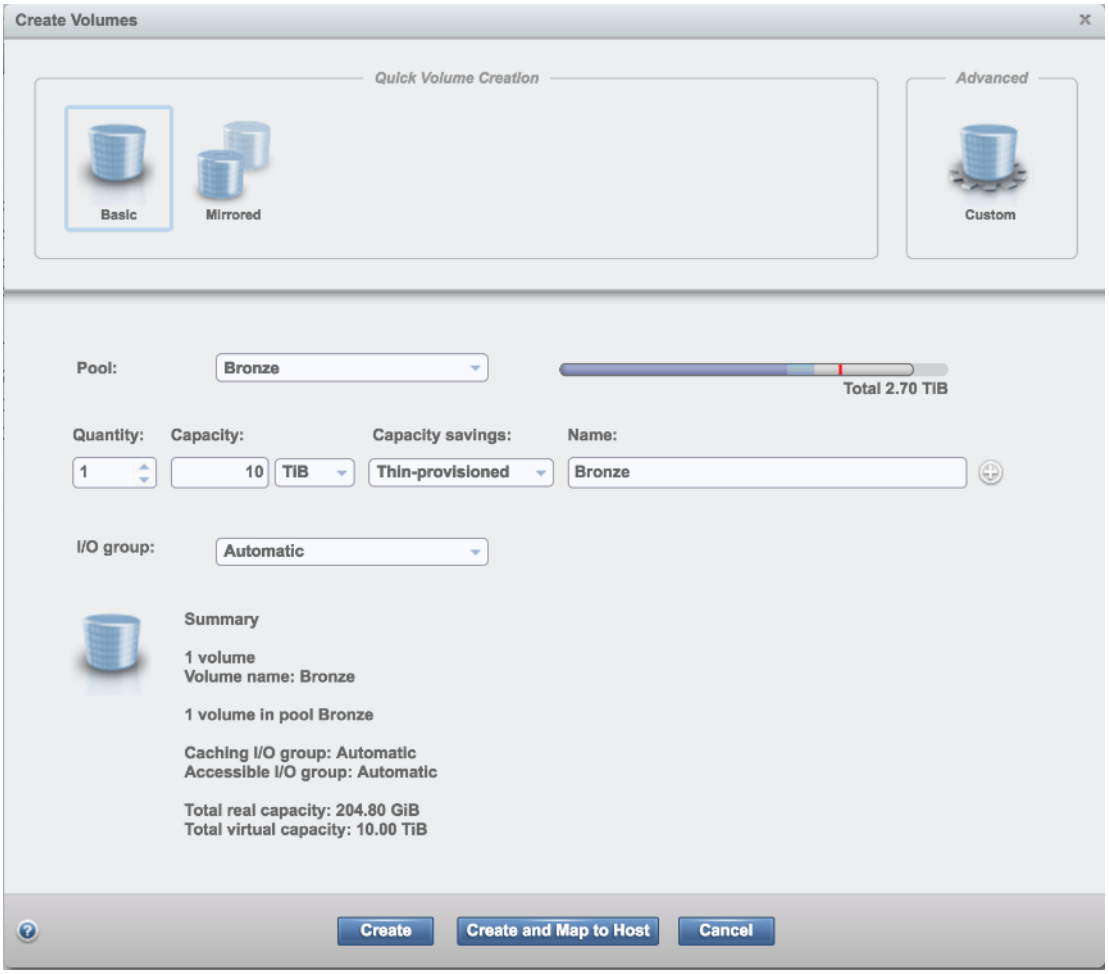

31. Choose Basic or Mirrored volume setting and select the Bronze pool that was previously added. Select I/O group Automatic and input the, capacity, desired capacity savings and name for the volume. Click Create and then click Close.

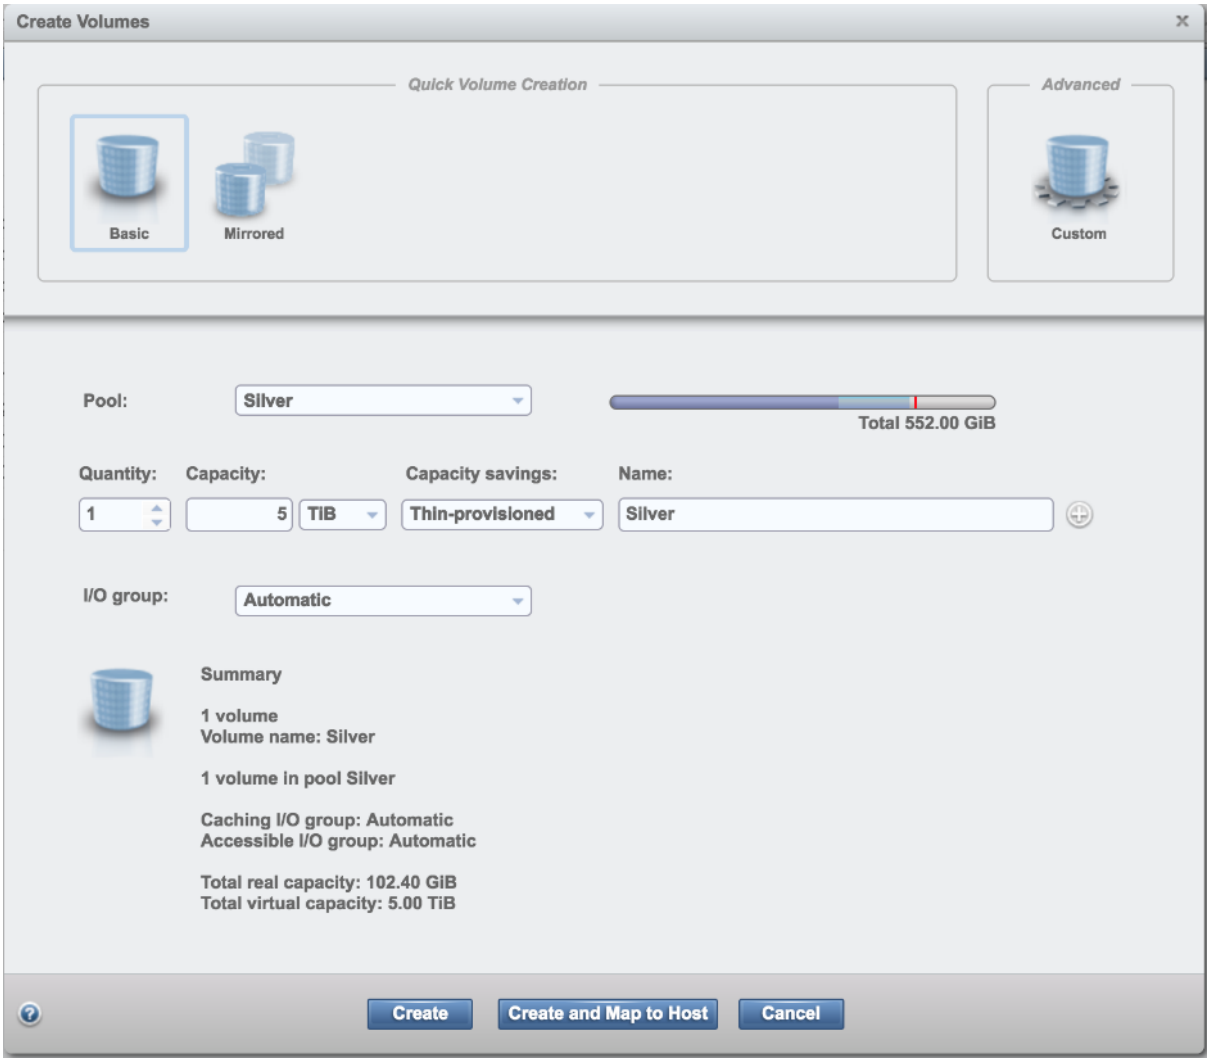

32. Click create volume again, select the storage Pool Silver, and select I/O group Automatic. Input the, capacity, desired capacity savings and the volume name. Click Create and then click Close.

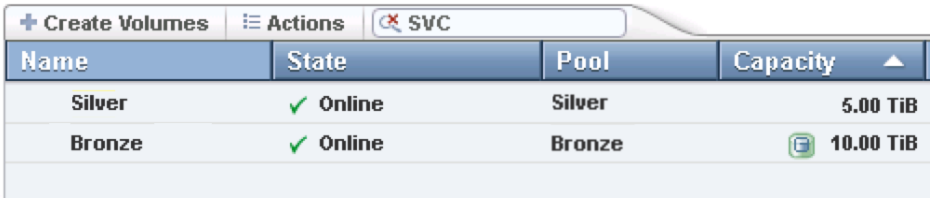

33. Validate the created volumes.

IBM San Volume Controller Base Configuration

Configuring the IBM SAN Volume Controller is a two-stage setup. The technician port (T) will be used for the initial configuration and IP assignment to the configuration node, and the management GUI will be used to complete the configuration.

For an in-depth look at installing and configuring the IBM SAN Volume Controller, refer to Redbook publication: Implementing the IBM System Storage SAN Volume Controller.

Prerequisites

Begin this procedure only after the physical installation of the IBM SAN Volume Controller has been completed. The computer used to initialize the IBM SAN Volume Controller must have an Ethernet cable connecting the personal computer to the technician port of the IBM SAN Volume Controller as well as a supported browser installed. At the time of writing, the following browsers or later are supported with the management GUI; Firefox 45, Internet Explorer 11 and Google Chrome 49.

IBM SAN Volume Controller Initial Configuration

To initialize a new system, you must connect a computer, configured as below, to the technician port on the rear of the IBM SAN Volume Controller node and then run the initialization tool. This node becomes the configuration node and provides access to the initialization GUI. Access the initialization GUI by using the management IP address through your IP network or through the technician port.

Use the initialization GUI to add each additional candidate node to the system.

1. Power on the IBM SAN Volume Controller node. Use the supplied power cords to connect both power supply units.

2. Check the LEDs on each node. The Power LED should be solidly on after a few seconds, but if it continues to blink after one minute, press the power-control button. The LED's are shown below for both the DH8 and SV1 SAN Volume Controller.

The SAN Volume Controller runs an extended series of power-on self-tests. The node might appear to be idle for up to five minutes after powering on.

3. Operator information LEDs on the front of the DH8 SAN Volume Controller:

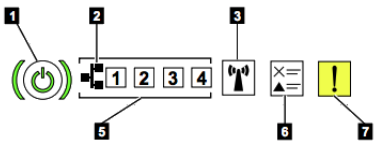

1. Power-control button and power-on LED
2. Ethernet icon
3. System-locator button and LED (blue)
5. Ethernet activity LEDs
7. System-error LED (amber)

4. Operator information LEDS on front of the SV1 SAN Volume Controller:

   1. Power-control button and power-on LED
   2. Identify LED
   3. Node status LED
   4. Node fault LED
   5. Battery status LED

5. Configure an Ethernet port, on the computer used to connect to the SAN Volume Controller, to enable Dynamic Host Configuration Protocol (DHCP) configuration of its IP address and DNS settings.

If DHCP cannot be enabled, configure the PC networking as follows: specify the static IPv4 address 192.168.0.2, subnet mask 255.255.255.0, gateway 192.168.0.1, and DNS 192.168.0.1.

Do not connect the technician port to a switch. If a switch is detected, the technician port connection might shut down, causing a 746 node error.

---

6. Locate the Ethernet port labelled T on the rear of the IBM SAN Volume Controller node. Refer to the appropriate figures below that show the location of the technician port, labelled 1, on each model.

Figure 6 IBM 2145-SV1 SAN Volume Controller

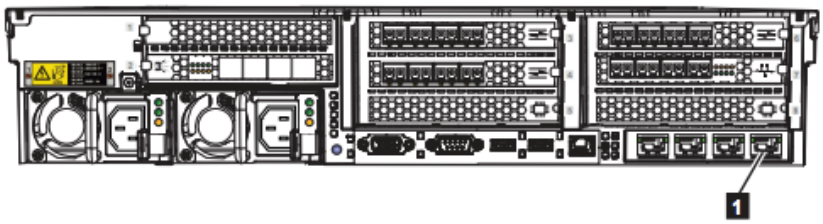

Figure 7 IBM 2145-DH8 SAN Volume Controller

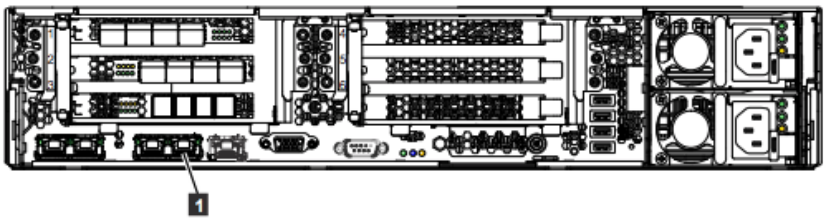

7. Connect an Ethernet cable between the port of the computer that is configured in step 3 and the technician port. After the connection is made, the system will automatically configure the IP and DNS settings for the personal computer if DHCP is available. If it is not available, the system will use the values provided in the steps above.

8. After the Ethernet port of the personal computer is connected, open a supported browser and browse to address http://install. (In case DCHP is not enabled, open a supported browser and go to the following static IP address 192.168.0.1.) The browser will automatically be directed to the initialization tool.

If there is a problem due to a change in system states, wait 5 - 10 seconds and then try again.

9. Click Next on the System Initialization welcome message.

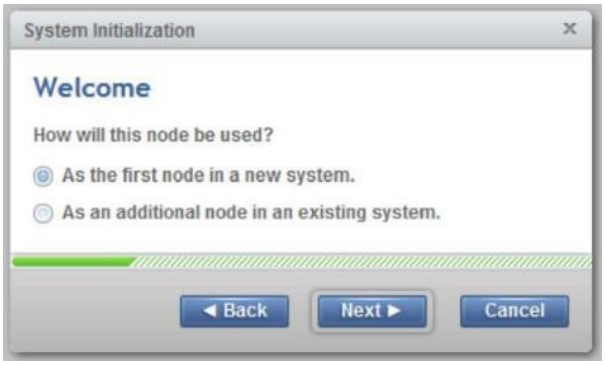

10. Click Next to continue with “As the first node in a new system.”

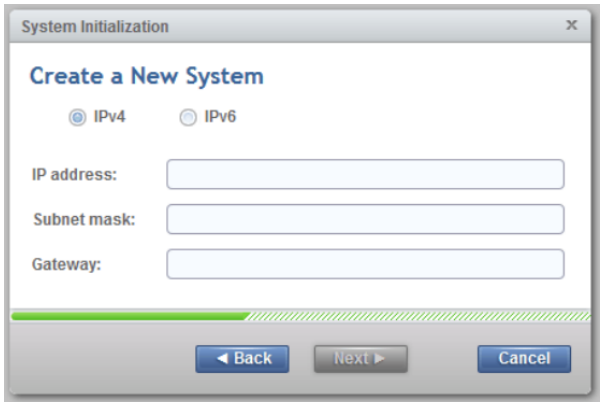

11. Complete all of the fields with the networking details for managing the system. This will be referred to as the System or Cluster IP address. Click Next.

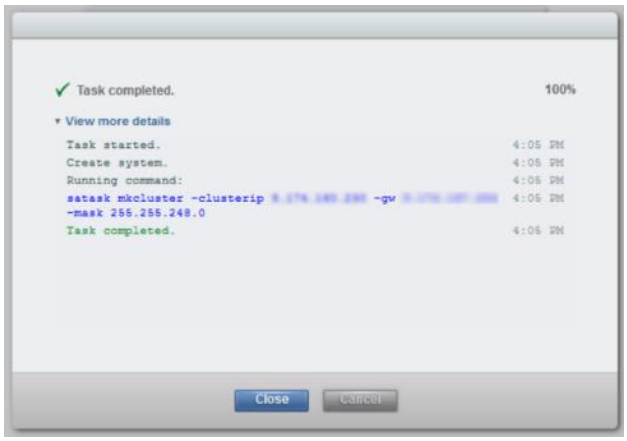

12. The setup task completes and provides a view of the generated satask CLI command as show above. Click Close. The storage enclosure will now reboot.

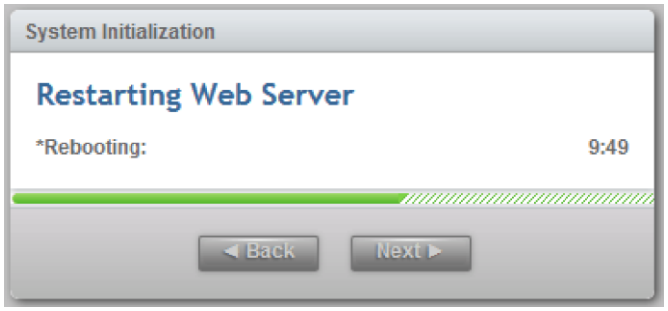

13. The system takes approximately 10 minutes to reboot and reconfigure the Web Server. After this time, click Next to proceed to the final step.

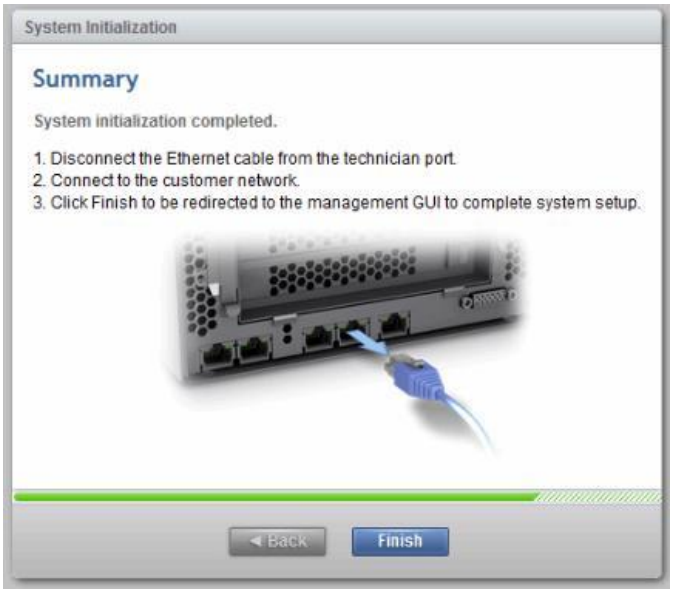

14. After completing the initialization process, disconnect the cable between the computer and the technician port, as instructed above. Re-establish the connection to the customer network and click Finish to be redirected to the management address that was provided while configuring the system.

15. Use the management GUI to continue the system configuration.

IBM SAN Volume Controller GUI Setup

After completing the initial tasks above, launch the management GUI and complete the following steps:

1. Log in to the management GUI using the previously configured cluster IP address.

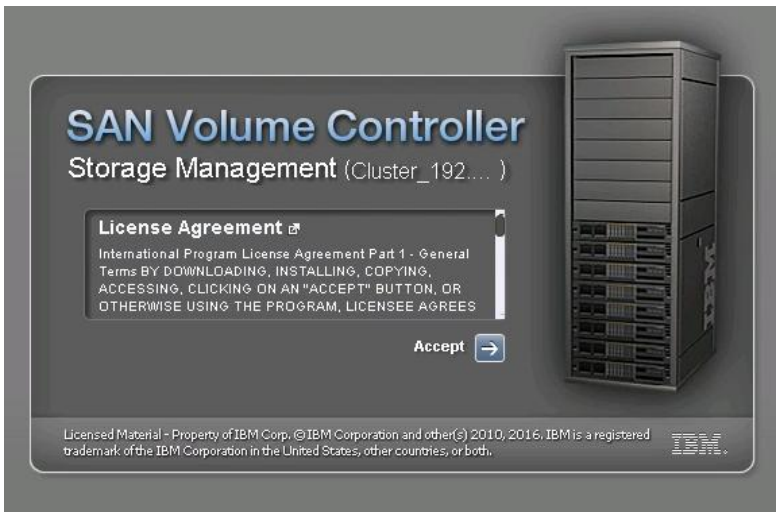

2. Read and accept the license agreement. Click Accept.

3. Resetting the superuser password is required when logging in for the first time. Make a note of the password and then click Log In.

The default password for the superuser account is *passw0rd* (zero and not o).

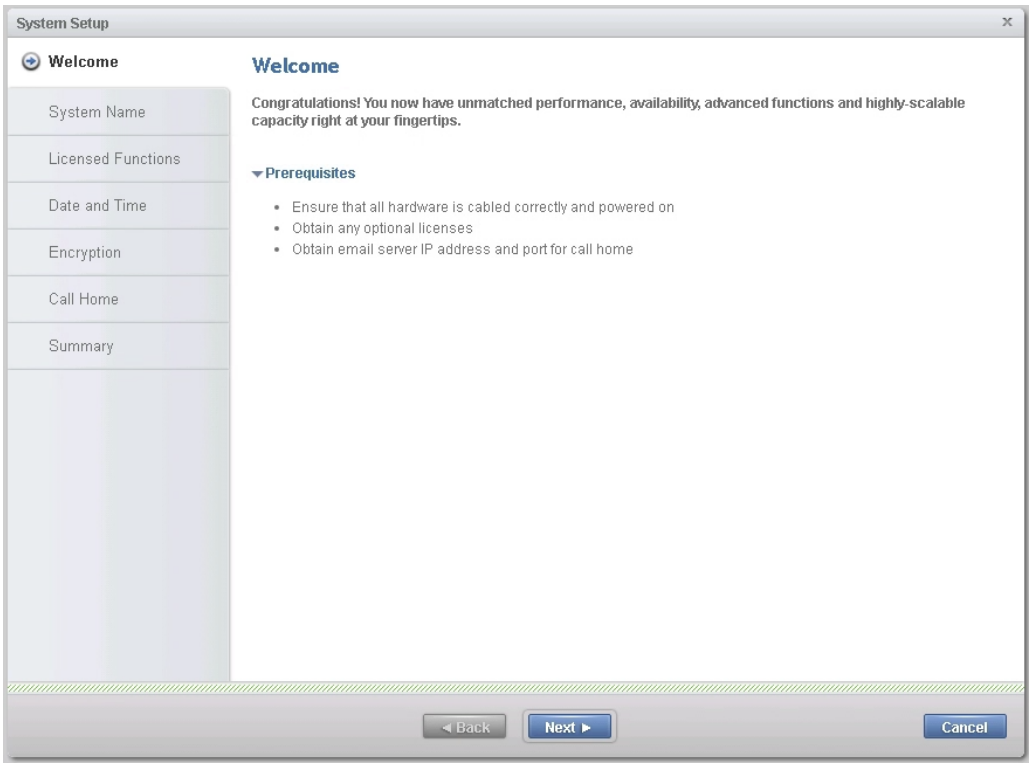

4. On the Welcome to System Setup screen click Next.

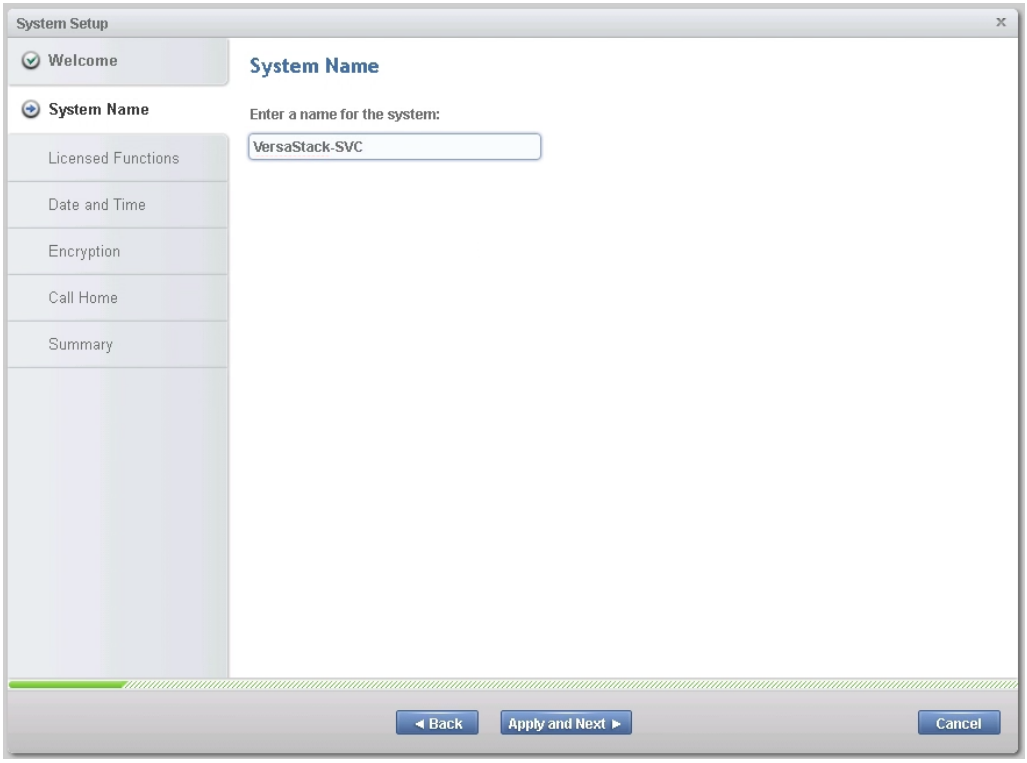

5. Enter the System Name and click Apply and Next to proceed.

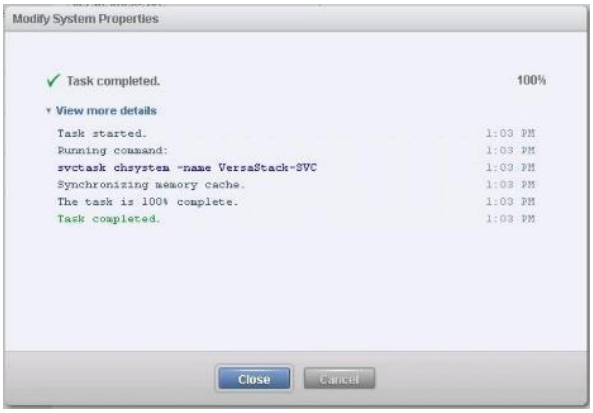

6. The *Modify System Properties* dialogue box shows the CLI command issued when applying the system name. Click Close to continue.

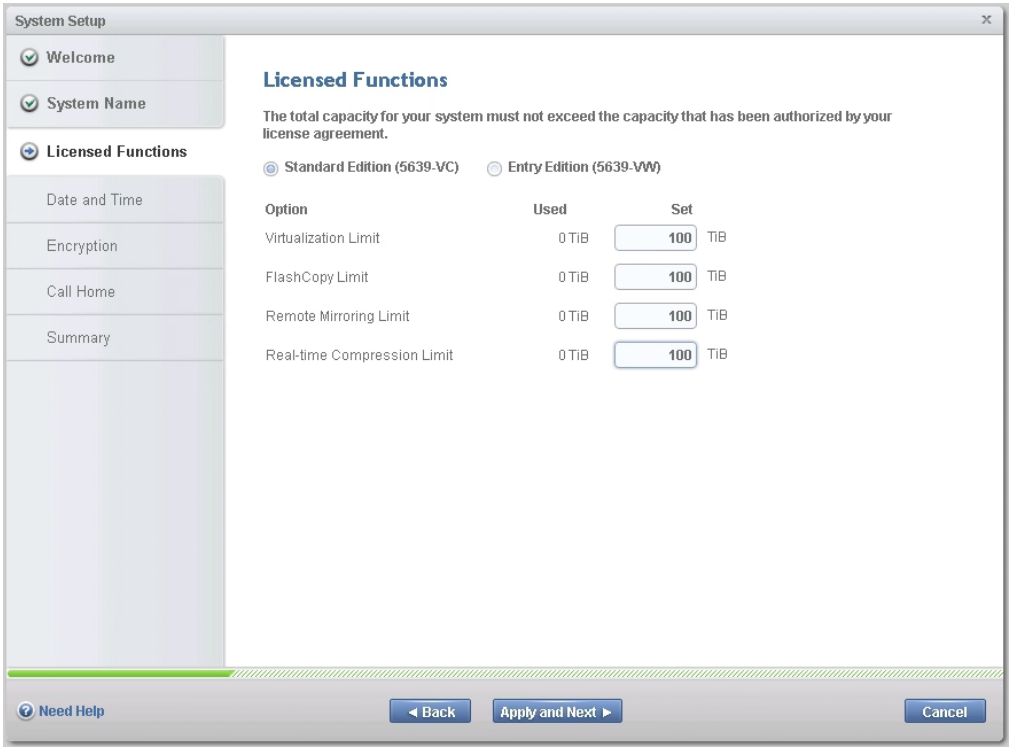

7. Select the licensed functions that were purchased with the system and enter the value (in TB) for Virtualization, FlashCopy, Remote Mirroring, and Real-time Compression. Click Apply and Next to proceed and click Close on the settings confirmation dialouge box.

The above vaules are shown as an example. Consider your license agreement details when populationg this information.

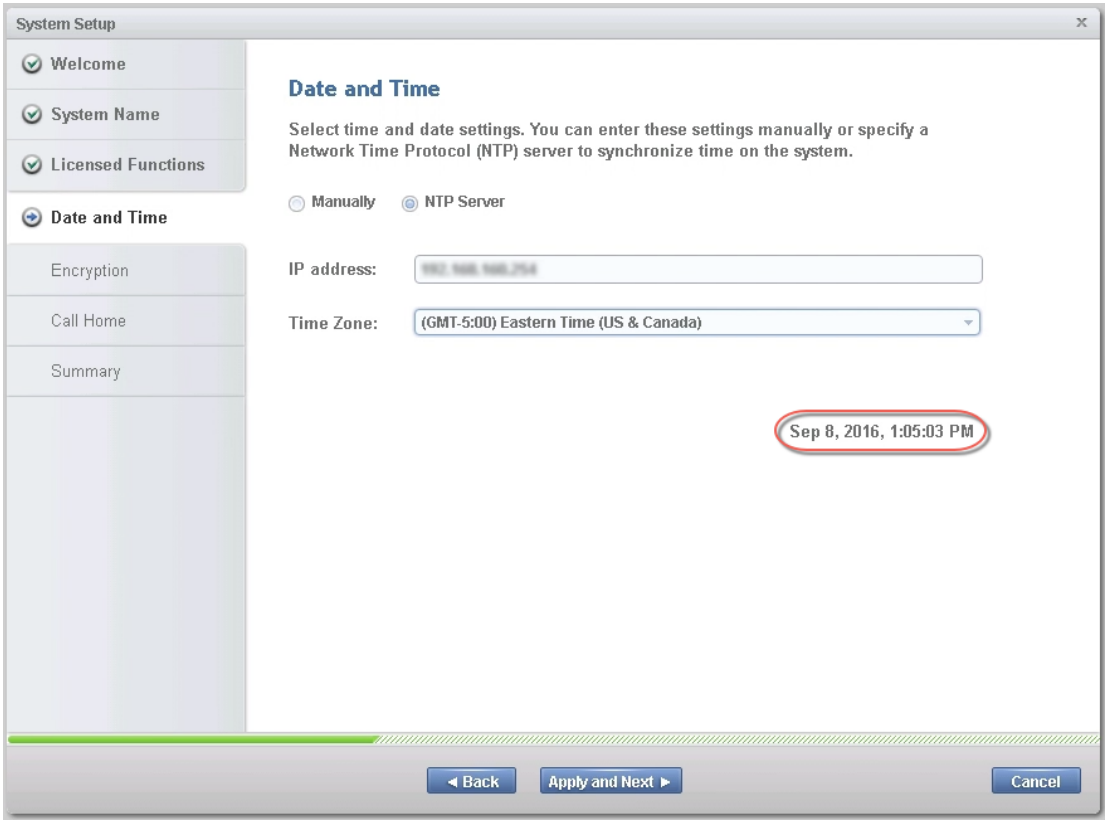

8. Configure the date and time settings and enter the NTP server details. Click Apply and Next to proceed and click Close on the settings confirmation dialouge box.

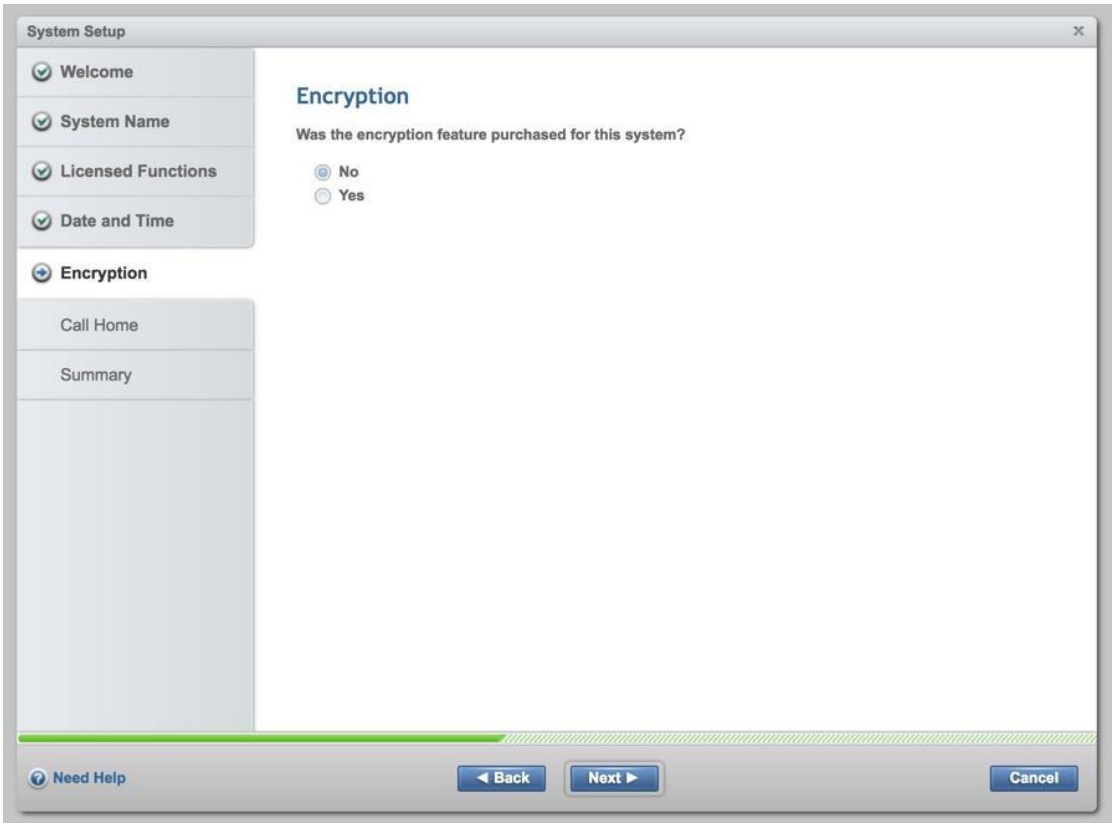

9. Enable the encryption feature (or leave it disabled). Click Next to proceed.

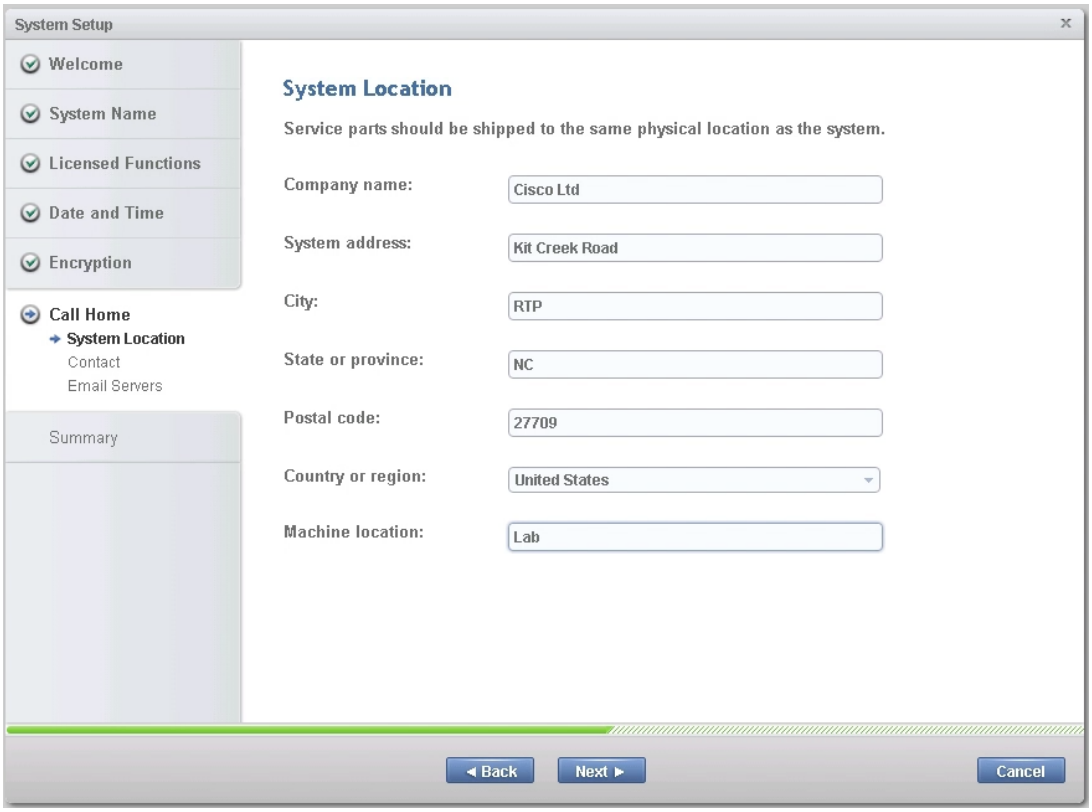

10. Enter the complete company name and address, and then click Next.

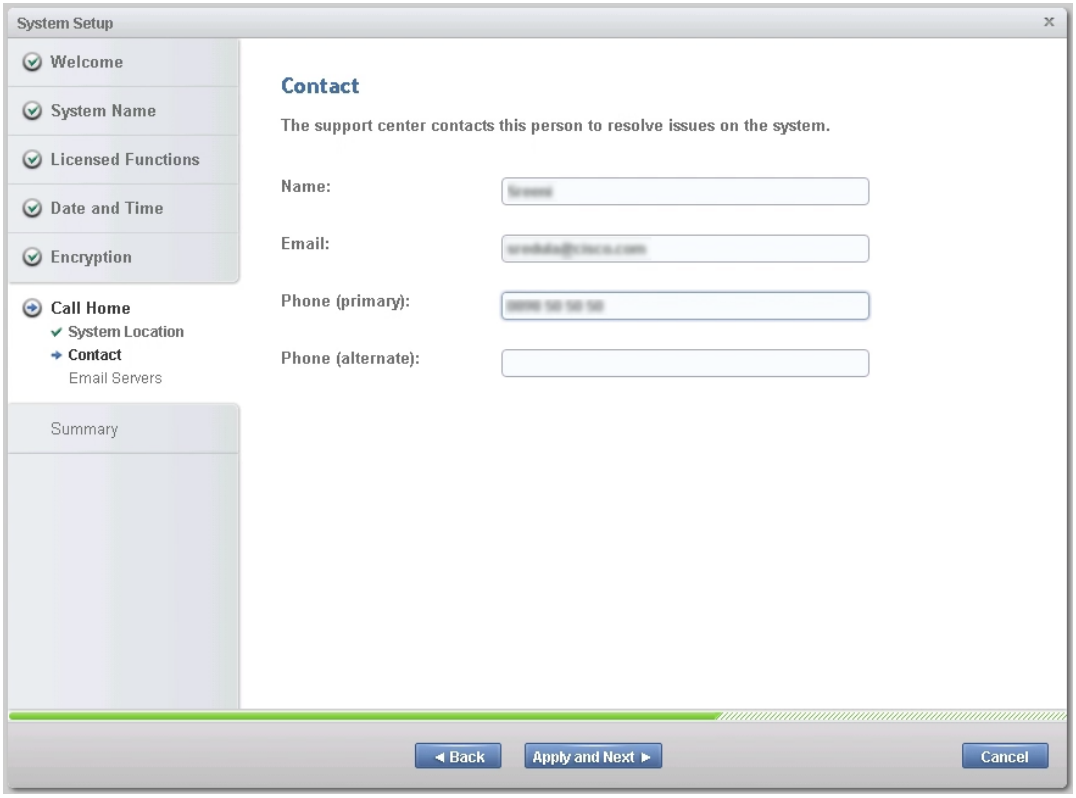

11. Enter the contact information for the support center. Click Apply and Next and click Close on the settings confirmation dialogue box.

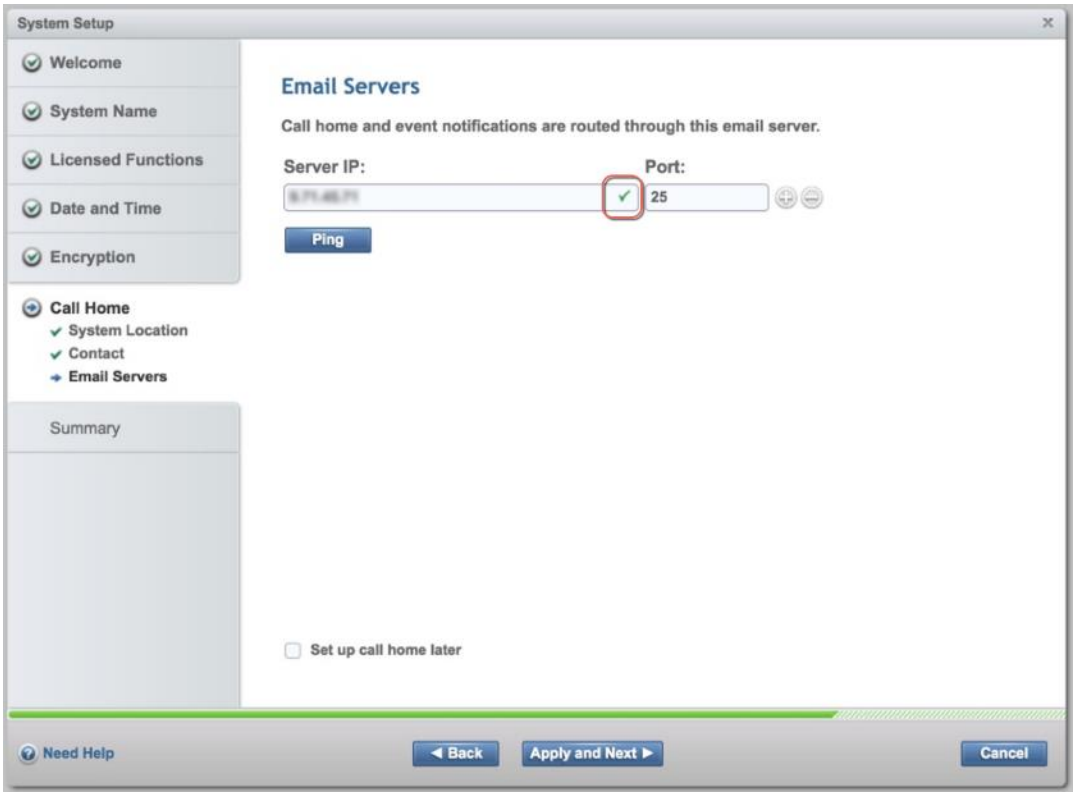

12. Enter the IP address and server port for one or more of the email servers for the Call Home email notification. Click Apply and Next.

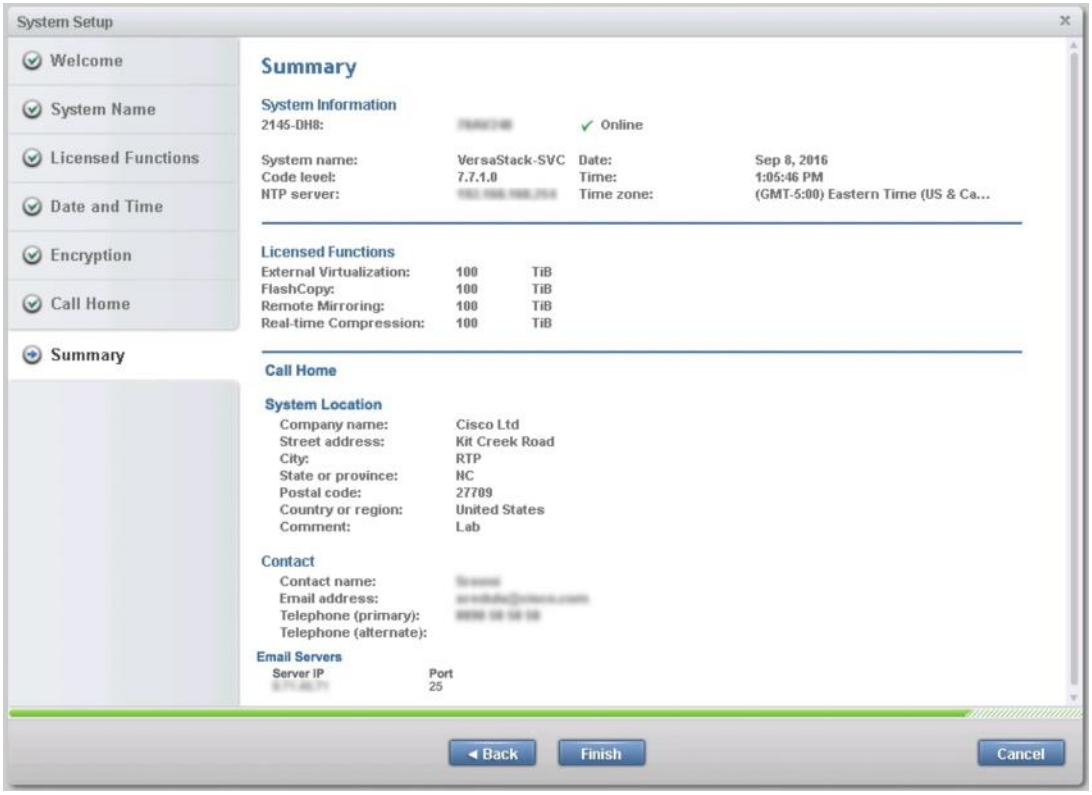

13. Review the final summary page, and click Finish to complete the System Setup wizard and click Close on the settings confirmation dialouge box.

14. Setup Completed. Click Close.

15. The System view for IBM SVC, as shown above, is now available.

Before adding a node to a system, make sure that the switch zoning is configured such that the node being added is in the same zone as all other nodes in the system.

Cisco MDS Zone Configuration

The following steps will configure zoning for the WWPNs for setting up the IBM SVC nodes as well as communication between SVC nodes and the FS900 and V5030 storage systems. WWPN information for various nodes can be easily collected using the command “show flogi database.” Refer to Table 6 to identify the ports where IBM nodes are connected to the MDS switches. In this configuration step, various zones will be created to enable communication between all the IBM nodes.

The configuration below assumes two SVC nodes have been deployed. FC ports 1 and 3 from each node are connected to the MDS-A switch and FC ports 2 and 4 are connected to MDS-B switch. Customers can adjust the configuration according to their deployment size.

Cisco MDS - A Switch

Log in to the MDS switch and complete the following steps.

1. Configure all the relevant ports (Table 6) on Cisco MDS as follows:

```
interface fc1/x
  port-license acquire
  no shutdown
!
```

2. Create the VSAN and add all the ports from Table 6:

```
vsan database
  vsan 101 interface fc1/x
  vsan 101 interface fc1/x
  <..>
```

3. The WWPNs obtained from “show flogi database” will be used in this step. Replace the variables with actual WWPN values.

```
device-alias database
  device-alias name SVC-Clus-Node1-FC1 pwwn <Actual PWWN for Node1 FC1>
  device-alias name SVC-Clus-Node1-FC3 pwwn <Actual PWWN for Node1 FC3>
  device-alias name SVC-Clus-Node2-FC1 pwwn <Actual PWWN for Node2 FC1>
  device-alias name SVC-Clus-Node2-FC3 pwwn <Actual PWWN for Node2 FC3>
  device-alias name FS900-Can1-FC1 pwwn <Actual PWWN for FS900 CAN1 FC1>
  device-alias name FS900-Can1-FC3 pwwn <Actual PWWN for FS900 CAN1 FC3>
  device-alias name FS900-Can2-FC1 pwwn <Actual PWWN for FS900 CAN2 FC1>
  device-alias name FS900-Can3-FC3 pwwn <Actual PWWN for FS900 CAN2 FC3>
  device-alias name V5030-Cont1-FC3 pwwn <Actual PWWN for V5030 Cont1 FC3>
  device-alias name V5030-Cont2-FC3 pwwn < Actual PWWN for V5030 Cont2 FC3>
device-alias commit
```

4. Create the zones and add device-alias members for the SVC inter-node and SVC nodes to storage system configurations.

```
zone name Inter-Node vsan 101
member device-alias SVC-Clus-Node1-FC1
member device-alias SVC-Clus-Node1-FC3
member device-alias SVC-Clus-Node2-FC1
member device-alias SVC-Clus-Node2-FC3

!
zone name SVC-5030 vsan 101
member device-alias SVC-Clus-Node1-FC1
member device-alias SVC-Clus-Node1-FC3
member device-alias SVC-Clus-Node2-FC1
member device-alias SVC-Clus-Node2-FC3
member device-alias V5030-Cont1-FC3
member device-alias V5030-Cont2-FC3
!

zone name SVC-FS900 vsan 101
member device-alias SVC-Clus-Node1-FC1
member device-alias SVC-Clus-Node1-FC3
member device-alias SVC-Clus-Node2-FC1
member device-alias SVC-Clus-Node2-FC3
member device-alias FS900-Can1-FC1
member device-alias FS900-Can1-FC3
member device-alias FS900-Can2-FC1
member device-alias FS900-Can2-FC3
!
```

5. Add zones to zoneset.

```
zoneset name versastackzoneset vsan 101
    member Inter-Node
```

```
    member SVC-V5030
    member SVC-FS900
```

6. Activate the zoneset

```
zoneset activate name versastackzoneset vsan 101
```

Validate all the HBA's are logged into the MDS switch. The SVC nodes and storage systems should be powered on.

7. Validate all the HBAs are logged into the switch using the “show zoneset active” command.

```
VersaStack-SVC-FabA# show zoneset active
zoneset name versastackzoneset vsan 101
  zone name Inter-Node vsan 101
  * fcid 0x400240 [pwwn 50:05:07:68:0c:11:93:c8] [SVC-Clus-Node1-FC1]
  * fcid 0x400260 [pwwn 50:05:07:68:0c:13:93:c8] [SVC-Clus-Node1-FC3]
  * fcid 0x400280 [pwwn 50:05:07:68:0c:11:93:c2] [SVC-Clus-Node2-FC1]
  * fcid 0x4002a0 [pwwn 50:05:07:68:0c:13:93:c2] [SVC-Clus-Node2-FC3]

zone name SVC-V5030 vsan 101
  * fcid 0x400240 [pwwn 50:05:07:68:0c:11:93:c8] [SVC-Clus-Node1-FC1]
  * fcid 0x400260 [pwwn 50:05:07:68:0c:13:93:c8] [SVC-Clus-Node1-FC3]
  * fcid 0x400280 [pwwn 50:05:07:68:0c:11:93:c2] [SVC-Clus-Node2-FC1]
  * fcid 0x4002a0 [pwwn 50:05:07:68:0c:13:93:c2] [SVC-Clus-Node2-FC3]
  * fcid 0x400200 [pwwn 50:05:07:68:0d:0c:58:f0] [V5030-Cont1-FC3]
  * fcid 0x400220 [pwwn 50:05:07:68:0d:0c:58:f1] [V5030-Cont2-FC3]

  zone name SVC-FS900 vsan 101
  * fcid 0x400240 [pwwn 50:05:07:68:0c:11:93:c8] [SVC-Clus-Node1-FC1]
  * fcid 0x400260 [pwwn 50:05:07:68:0c:13:93:c8] [SVC-Clus-Node1-FC3]
  * fcid 0x400280 [pwwn 50:05:07:68:0c:11:93:c2] [SVC-Clus-Node2-FC1]
  * fcid 0x4002a0 [pwwn 50:05:07:68:0c:13:93:c2] [SVC-Clus-Node2-FC3]
  * fcid 0x4000a0 [pwwn 50:05:07:60:5e:83:cc:81] [FS900-Can1-FC1]
  * fcid 0x4000c0 [pwwn 50:05:07:60:5e:83:cc:91] [FS900-Can1-FC3]
  * fcid 0x400080 [pwwn 50:05:07:60:5e:83:cc:a1] [FS900-Can2-FC1]
  * fcid 0x4000e0 [pwwn 50:05:07:60:5e:83:cc:b1] [FS900-Can2-FC3]
```

8. Save the configuration.

```
 copy run start
```

Cisco MDS - B Switch

Log into the MDS switch and complete the following steps:

1. Configure all the relevant ports (Table 6) on Cisco MDS as follows:

```
interface fc1/x
  port-license acquire
  no shutdown
!
```

2. Create the VSAN and add all the ports from Table 6 :

```
vsan database
  vsan 102 interface fc1/x
  vsan 102 interface fc1/x
  <..>
```

3. The WWPNs obtained from “show flogi database” will be used in this step. Replace the variables with actual WWPN values.

```
device-alias database
  device-alias name SVC-Clus-Node1-FC2 pwwn <Actual PWWN for Node1 FC2>
  device-alias name SVC-Clus-Node1-FC4 pwwn <Actual PWWN for Node1 FC4>
  device-alias name SVC-Clus-Node2-FC2 pwwn <Actual PWWN for Node2 FC2>
  device-alias name SVC-Clus-Node2-FC4 pwwn <Actual PWWN for Node2 FC4>
  device-alias name FS900-Can1-FC2 pwwn <Actual PWWN for FS900 CAN1 FC2>
  device-alias name FS900-Can1-FC4 pwwn <Actual PWWN for FS900 CAN1 FC4>
  device-alias name FS900-Can2-FC2 pwwn <Actual PWWN for FS900 CAN2 FC2>
  device-alias name FS900-Can3-FC4 pwwn <Actual PWWN for FS900 CAN2 FC4>
  device-alias name V5030-Cont1-FC4 pwwn <Actual PWWN for V5030 Cont1 FC4>
  device-alias name V5030-Cont2-FC4 pwwn < Actual PWWN for V5030 Cont2 FC4>
device-alias commit
```

4. Create the zones and add device-alias members for the SVC inter-node and SVC nodes to storage system configurations.

```
zone name Inter-Node vsan 102
member device-alias SVC-Clus-Node1-FC2
member device-alias SVC-Clus-Node1-FC4
member device-alias SVC-Clus-Node2-FC2
member device-alias SVC-Clus-Node2-FC4

!
zone name SVC-5030 vsan 102
member device-alias SVC-Clus-Node1-FC2
member device-alias SVC-Clus-Node1-FC4
member device-alias SVC-Clus-Node2-FC2
member device-alias SVC-Clus-Node2-FC4
member device-alias V5030-Cont1-FC4
member device-alias V5030-Cont2-FC4
!

zone name SVC-FS900 vsan 102
member device-alias SVC-Clus-Node1-FC2
member device-alias SVC-Clus-Node1-FC4
member device-alias SVC-Clus-Node2-FC2
member device-alias SVC-Clus-Node2-FC4
member device-alias FS900-Can1-FC2
member device-alias FS900-Can1-FC4
member device-alias FS900-Can2-FC2
member device-alias FS900-Can2-FC4
!
```

5. Add zones to zoneset.

```
zoneset name versastackzoneset vsan 102
    member Inter-Node
    member SVC-V5030
```

```
    member SVC-FS900
```

6. Activate the zoneset

```
zoneset activate  name versastackzoneset vsan 102
```

Validate all the HBA's are logged into the MDS switch. The SVC nodes and storage systems should be powered on.

7. Validate all the HBAs are logged into the switch using the “show zoneset active” command.

```
VersaStack-SVC-FabB# show zoneset active

zoneset name versastackzoneset vsan 102
  zone name Inter-Node vsan 102
  * fcid 0x770240 [pwwn 50:05:07:68:0c:12:93:c8] [SVC-Clus-Node1-FC2]
  * fcid 0x770260 [pwwn 50:05:07:68:0c:14:93:c8] [SVC-Clus-Node1-FC4]
  * fcid 0x770280 [pwwn 50:05:07:68:0c:12:93:c2] [SVC-Clus-Node2-FC2]
  * fcid 0x7702a0 [pwwn 50:05:07:68:0c:14:93:c2] [SVC-Clus-Node2-FC4]



  zone name SVC-V5030 vsan 102
  * fcid 0x770240 [pwwn 50:05:07:68:0c:12:93:c8] [SVC-Clus-Node1-FC2]
  * fcid 0x770260 [pwwn 50:05:07:68:0c:14:93:c8] [SVC-Clus-Node1-FC4]
  * fcid 0x770280 [pwwn 50:05:07:68:0c:12:93:c2] [SVC-Clus-Node2-FC2]
  * fcid 0x7702a0 [pwwn 50:05:07:68:0c:14:93:c2] [SVC-Clus-Node2-FC4]
  * fcid 0x7701c0 [pwwn 50:05:07:68:0d:10:58:f0] [V5030-Cont1-FC4]
  * fcid 0x770220 [pwwn 50:05:07:68:0d:10:58:f1] [V5030-Cont2-FC4]

  zone name SVC-FS900 vsan 102
  * fcid 0x770240 [pwwn 50:05:07:68:0c:12:93:c8] [SVC-Clus-Node1-FC2]
  * fcid 0x770260 [pwwn 50:05:07:68:0c:14:93:c8] [SVC-Clus-Node1-FC4]
  * fcid 0x770280 [pwwn 50:05:07:68:0c:12:93:c2] [SVC-Clus-Node2-FC2]
  * fcid 0x7702a0 [pwwn 50:05:07:68:0c:14:93:c2] [SVC-Clus-Node2-FC4]

  * fcid 0x770080 [pwwn 50:05:07:60:5e:83:cc:82] [FS900-Can1-FC2]
  * fcid 0x7700e0 [pwwn 50:05:07:60:5e:83:cc:92] [FS900-Can1-FC4]
  * fcid 0x7700a0 [pwwn 50:05:07:60:5e:83:cc:a2] [FS900-Can2-FC2]
  * fcid 0x7700c0 [pwwn 50:05:07:60:5e:83:cc:b2] [FS900-Can2-FC4]
```

8. Save the configuration.

```
  copy run start
```

Adding Additional Nodes

To have a fully functional SVC system, a second node should be added to the configuration. The configuration node should not be aware of additional nodes in the system because they became part of the same zone on the fabric. To add a node to a clustered system, complete the following steps.

1. Right-click the empty node space on io_grp0.

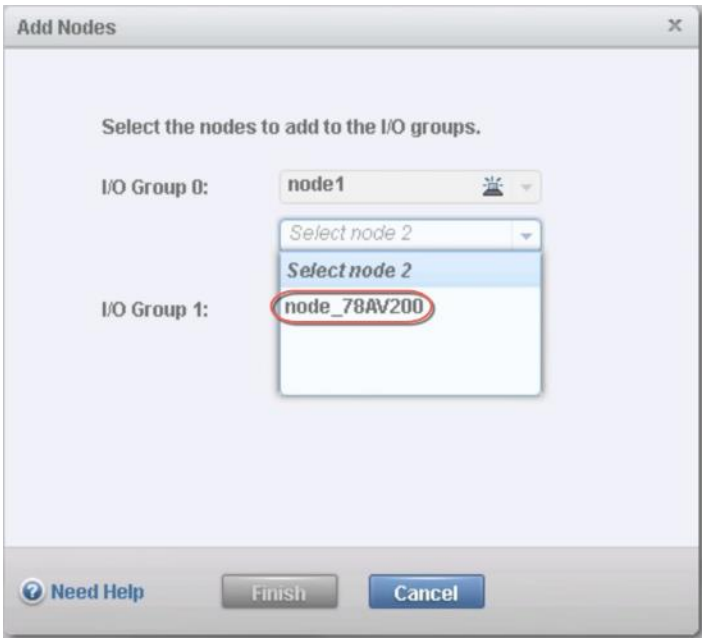

2. Select the 2nd node to be used in io_grp0.

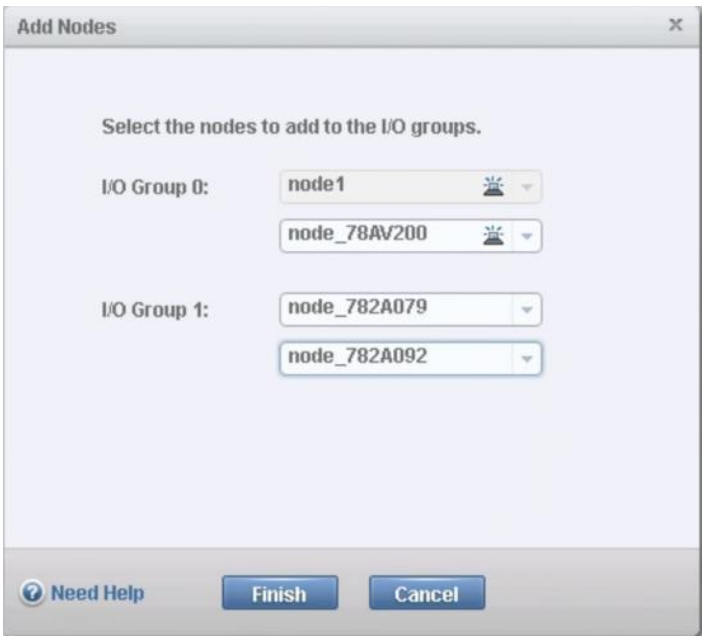

3. Repeat the process for any additional node pairs in order to create additional io_grp (1,2,3). Click Finish and click Close on the settings confirmation dialouge box.

4. On the left side of the management GUI, hover over each of the icons on the Navigation Dock to become familiar with the options.

5. Select the Setting icon from the Navigation Dock and choose Network.

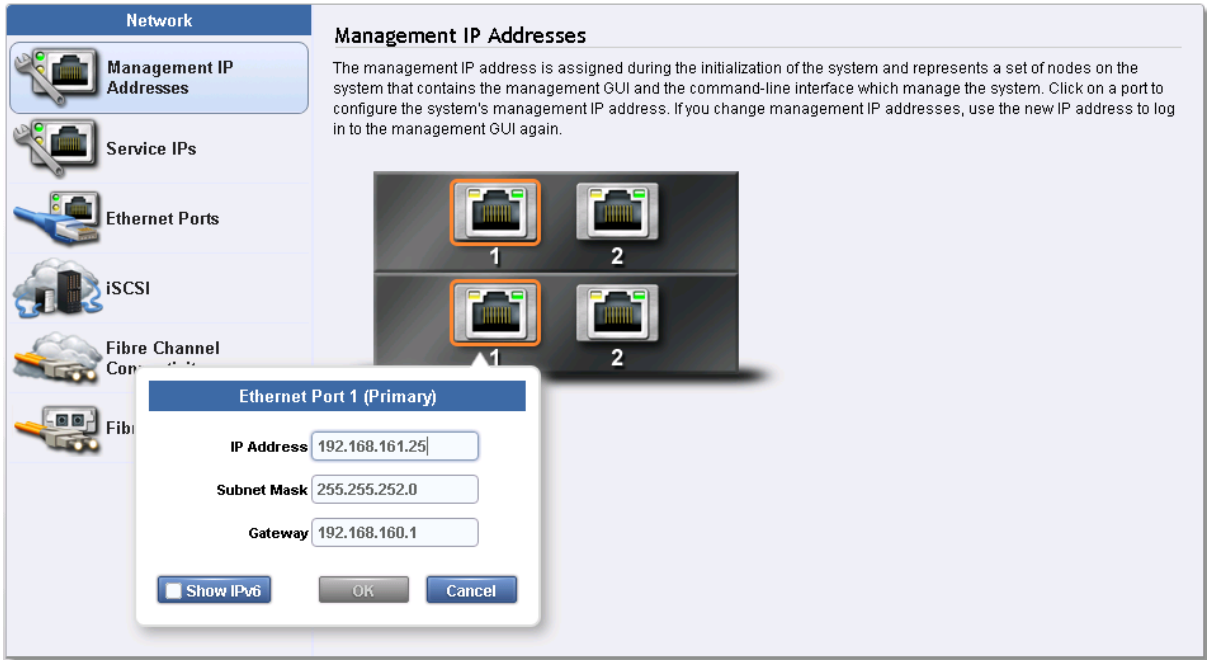

6. On the Network screen, highlight the Management IP Addresses section. Then click the number 1 interface on the left-hand side to bring up the Ethernet port IP menu. If required, change the IP address configured during the system initialization, click OK.

7. While still on the *Network* screen, select (1) Service IP Addresses from the list on the left and (2) Node Name node1 (io_grp0) then (3) change the IP address for port 1 to reflect the service IP allocated to this node, click OK.

8. Repeat this process for port 1 on the remaining nodes (and port 2 if cabled).

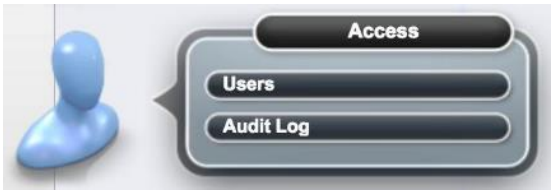

9. Click the Access icon from the Navigation Dock on the left and select Users to access the Users screen.

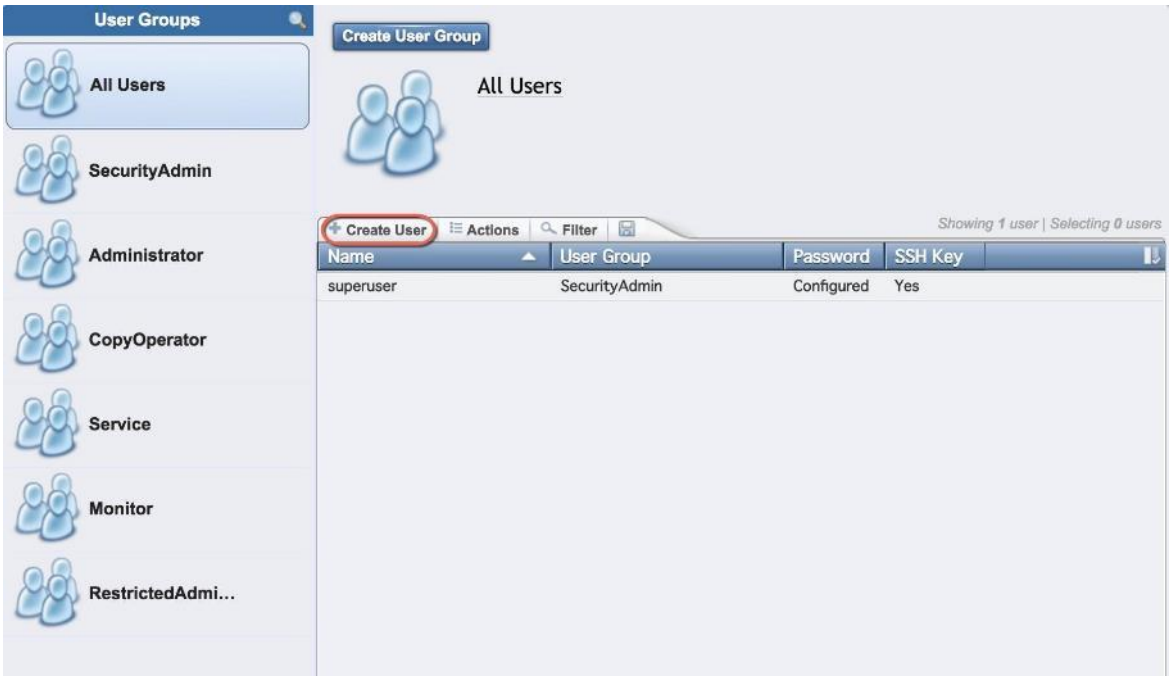

10. Select Create User.

11. Enter a new name for an alternative admin account. Leave the SecurityAdmin default as the User Group, and input the new password, then click Create. Optionally, save the SSH Public Key generate on a Unix server through the command “ssh-keygen -t rsa” to a file and associate the file to this user through the Choose File button.

12. Logout from the superuser account and log back in as the new account just created.

Adding External Storage to the SVC

To configure the IBM FlashSystem 900 and IBM Storwize V5030 as backend storage, complete the following steps:

IBM FlashSystem 900

1. Log into the management GUI for IBM FlashSystem 900.

2. From the management GUI for the IBM FlashSystem 900, select Hosts from the Navigation Dock, then click Hosts.

3. Click on Add Host.

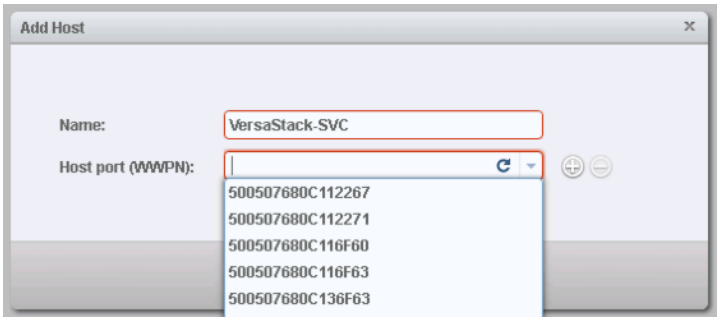

4. Create a host with a name representing the SAN Volume Controller and add all (16 in this example) the WWPNs of the SAN Volume Controller nodes to the new host.

5. Validate that host has been added with the correct number of host ports and that the state is Online

6. From the Navigation Dock, highlight Volumes, and select Volumes.

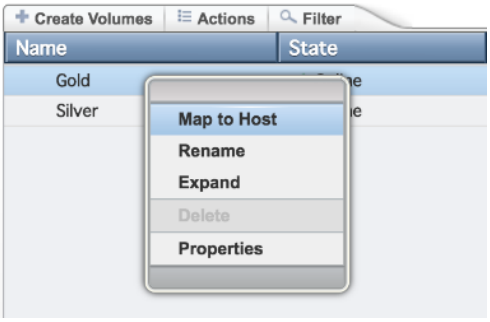

7. Select a volume by right-clicking, and then select Map to Host. In this example, the Gold volume was selected.

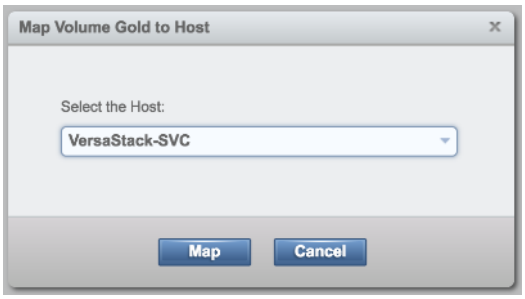

8. Select the host and click on Map and click Close to the Modify Mappings dialogue box.

Switch to IBM SVC management GUI.

9. Using the IBM SAN Volume Controller management GUI, select Pools and then click External Storage.

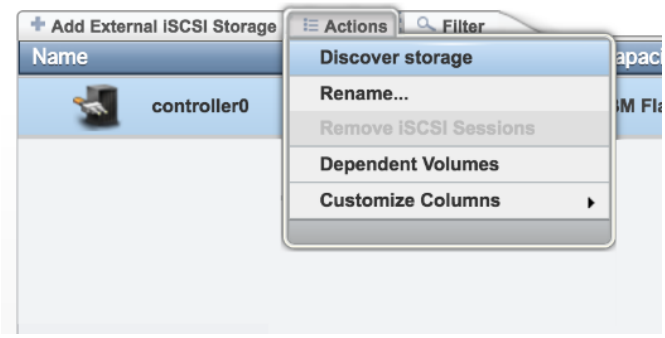

10. Select Actions and Discover storage to refresh any changes to the external storage.

11. Right-click the newly presented controller0 and select Rename. Enter the name for the IBM FlashSystem 900 (FlashSystem 900 in this example). Click Close to the Rename Storage System dialogue box.

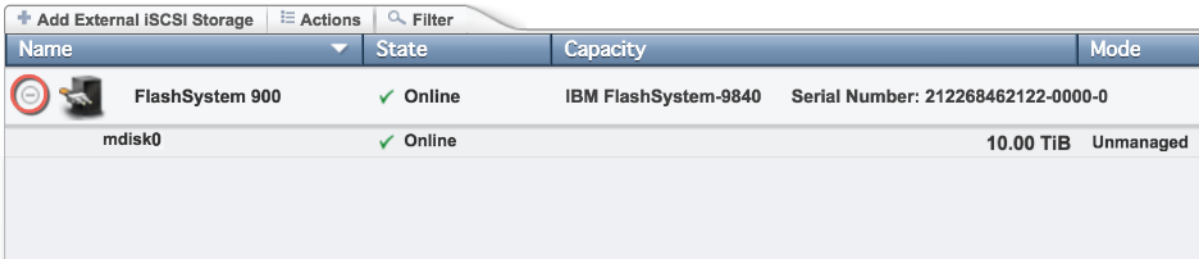

12. Select the FlashSystem 900 and click the plus sign (+), to expand the view. The volumes that we created and mapped on the backend FlashSystem 900 system will be displayed in the SVC MDisk view.

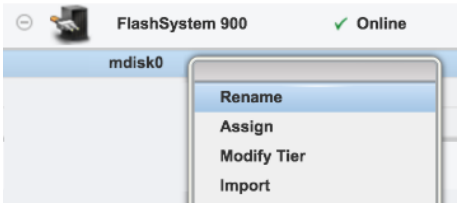

13. Right-click on new mdisk0 and rename to FS900-Gold. Click Rename and then click Close on the Rename MDiks dialogue box.

14. Return to the IBM FlashSystem 900 management GUI, and using the steps above, map the Silver volume to the IBM SAN Volume Controller host. Name the next Volume FS900-Siver.

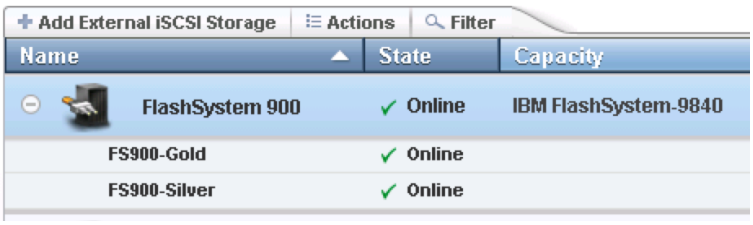

15. Validate that there are 2 MDisks on IBM SAN Volume Controller provided by the IBM FlashSystem 900 backend storage.

IBM Storwize V5030

1. Log into the management GUI for IBM Storwize V5030.

2. From the management GUI for the IBM Storwize V5030, select Hosts from the Navigation Dock, then click Hosts.

3. Click Add Host.

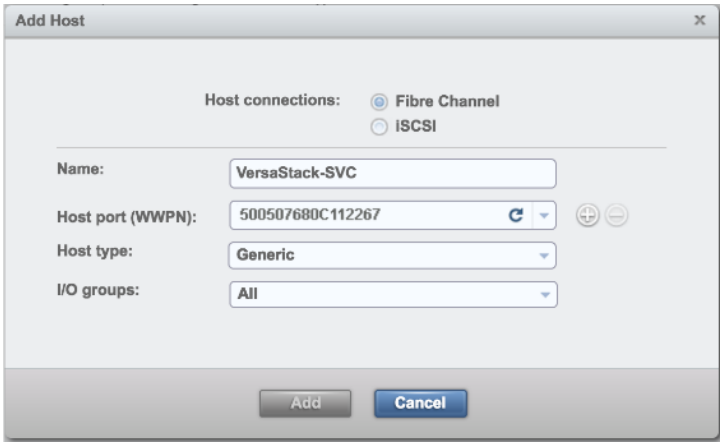

4. Create a host for the SAN Volume Controller and add all the WWPNs (16 in this example) of the IBM SAN Volume Controller nodes to the new host. Leave Host type and I/O groups as default. Click Add.

5. Validate that host has been added with the correct number of host ports and that the state is Online.

6. From the Navigation Dock, highlight Volumes, and select Volumes.

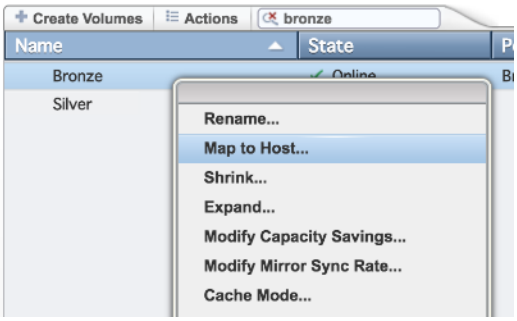

7. Select a volume by right-clicking, and then select Map to Host. In this example, the Bronze volume has been selected.

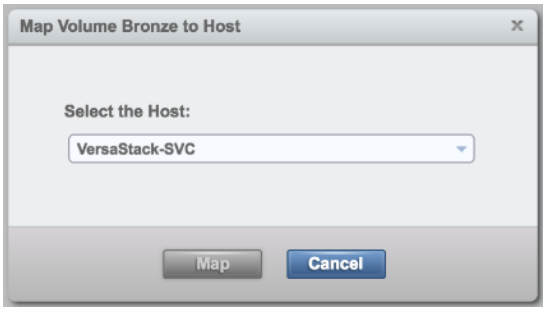

8. Select the host and click on Map and click Close to the Modify Mappings dialogue box.

9. Repeat the above steps to map the Silver volume to the host .

10. Using the IBM SAN Volume Controller management GUI, select Pools and then click External Storage.

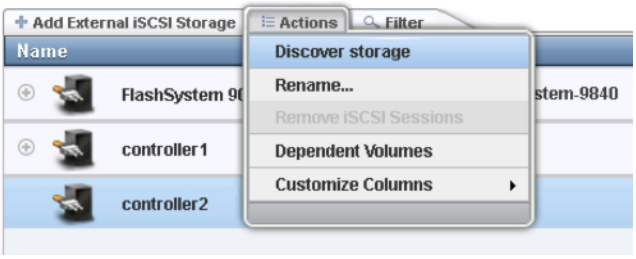

11. Select Actions and Discover storage to refresh any changes to the external storage.

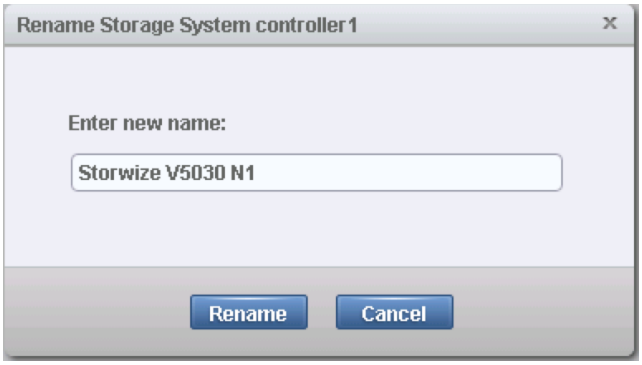

12. Right-click the newly presented controller1 and select Rename. Enter the name for IBM Storzwize V5030 (Storwize V5030 N1 in example above). Click Close.

13. Repeat above step to add controller2 and rename it to Storwize V5030 N2.

The Storwize V5030 is presented as 2 storage controllers, each storage controller represents a canister of the V5030 control enclosure. In this deployment, **N1** and **N2** was used to differentiate the two controllers.

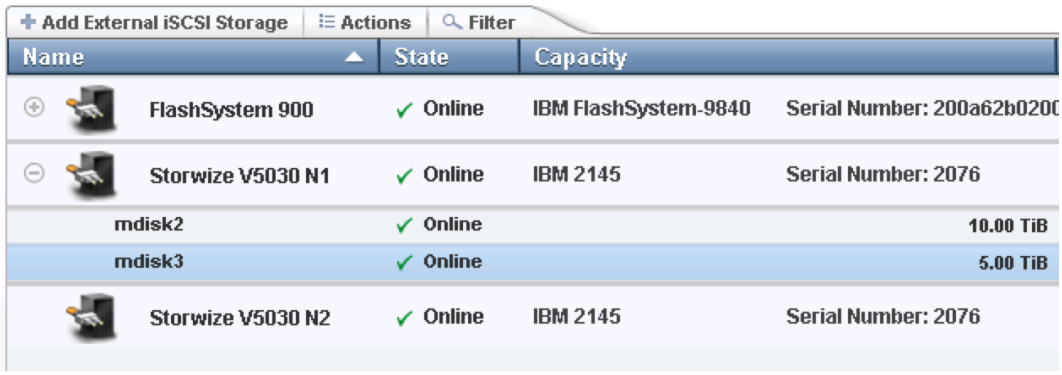

14. Select the Storwize V5030 controller with the Plus sign (+) and expand the view. The volumes that were previously created and mapped should be displayed in the SAN Volume Controller MDisk view.

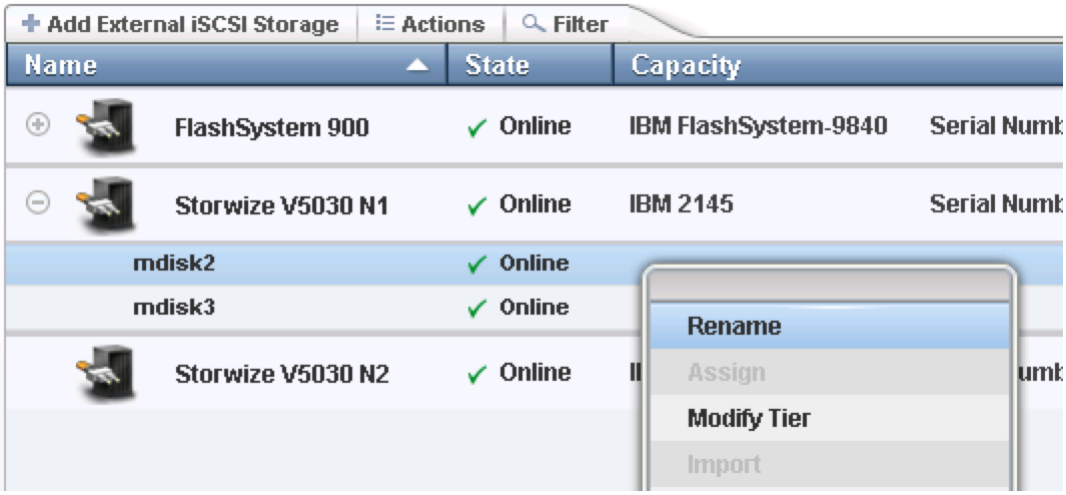

15. Using the volume size as a guide, right-click on new mdisk2 and rename to V5030-Bronze. Click Rename and then click Close. Repeat the process for mdisk3, renaming it to V5030-Silver.

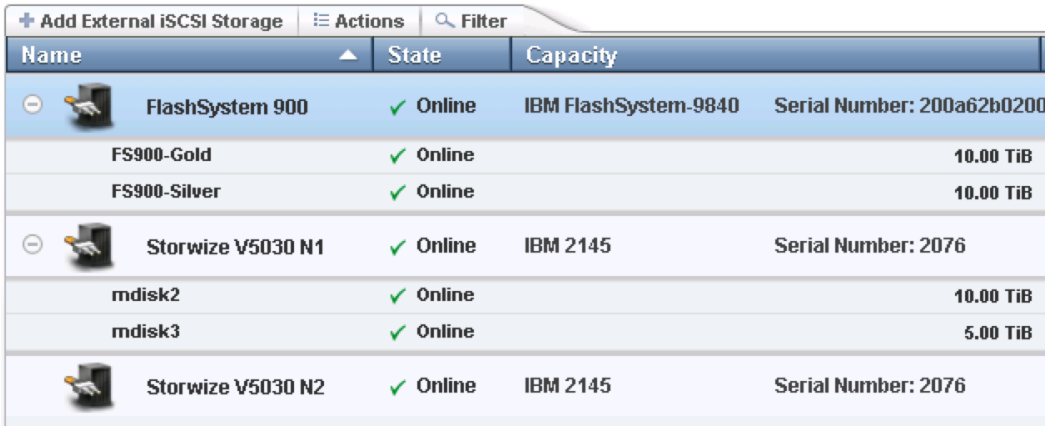

16. Validate that there are 4 MDisks on the IBM SAN Volume Controller provided by both IBM FlashSystem 900 and Storwize V5030 backend storage systems.

17. Highlight Pools from the Navigation Dock, and select MDisks by Pools.

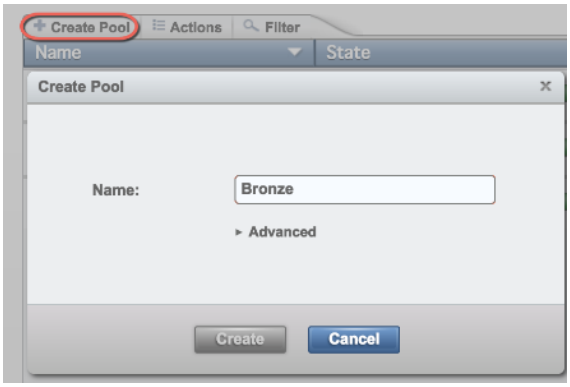

18. Click Create Pool and enter the name of your first pool: Bronze. Repeat this process for the Silver and Gold pools.

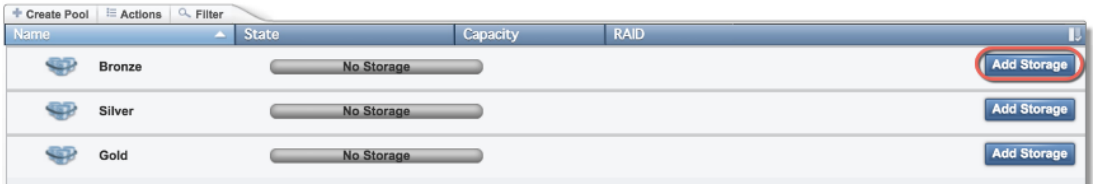

19. Select Add Storage on the Bronze pool.

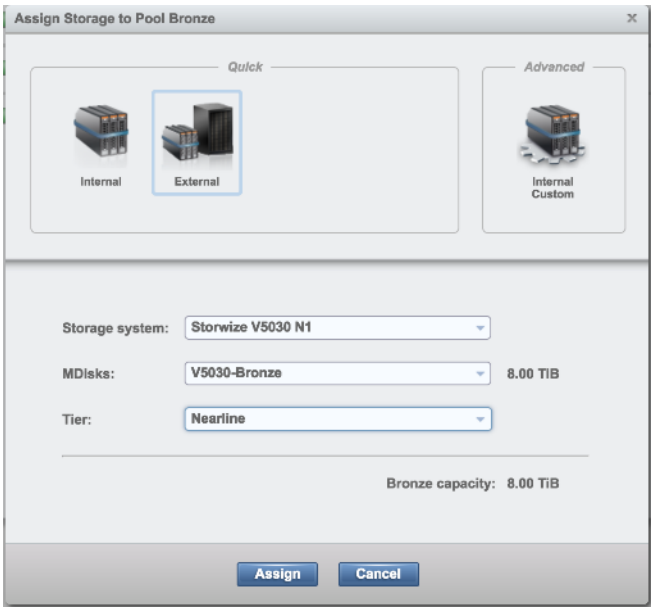

20. Select the Storwize V5030 storage system and the MDisks V5030-Bronze. Set the tier type to conform with the type of provisioned storage.

21. Repeat the process for the Gold pool, selecting the FlashSystem 900 storage system and the MDisks FS900-Gold. Set the tier type to Flash.

22. Repeat the process for the Silver pool, selecting the FlashSystem 900 storage system and the MDisks FS900-Silver. Set the tier type to Flash.

23. Repeat the process again for the Silver pool, selecting the Storwize V5030 storage system and the MDisks V5030-Silver. Set the tier type to Enterprise.

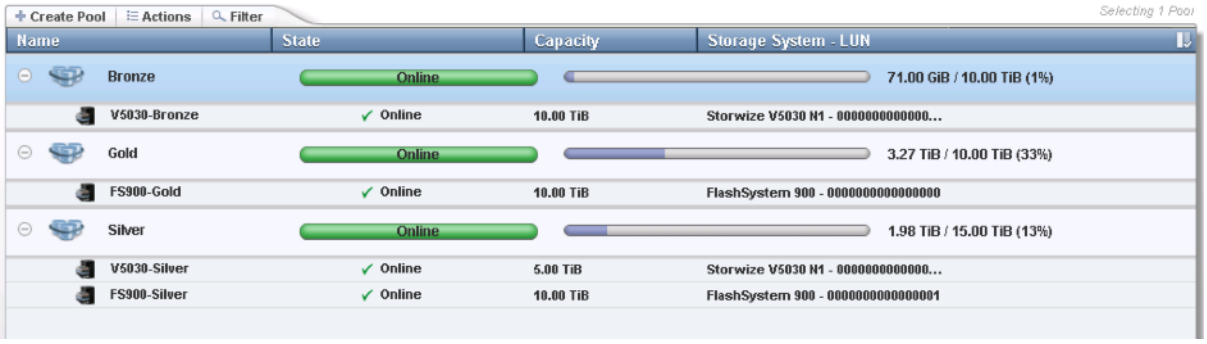

24. Validate that storage pools are as shown above.

IBM SVC iSCSI configuration

Cisco UCS configuration requires information about the iSCSI IQNs on IBM SVC. Therefore, as part of the initial storage configuration, iSCSI ports are configured on IBM SVC

Two 10G ports from each of the IBM SVC nodes are connected to each of Nexus 93180 switches. These ports are configured as shown in Table 7 .

Table 7 IBM SVC iSCSI Interface Configuration

System	Port	Path	VLAN	IP address
Node 1	4	iSCSI-A	3161	10.29.161.249/24
Node 1	6	iSCSI-B	3162	10.29.162.249/24
Node 2	4	iSCSI-A	3161	10.29.161.250/24
Node 2	6	iSCSI-B	3162	10.29.162.250/24

To configure the IBM SVC system for iSCSI storage access, complete the following steps:

1. Log into the IBM SVC GUI and navigate to Settings > Network.
2. Click the iSCSI icon

and enter the system and node names as shown:

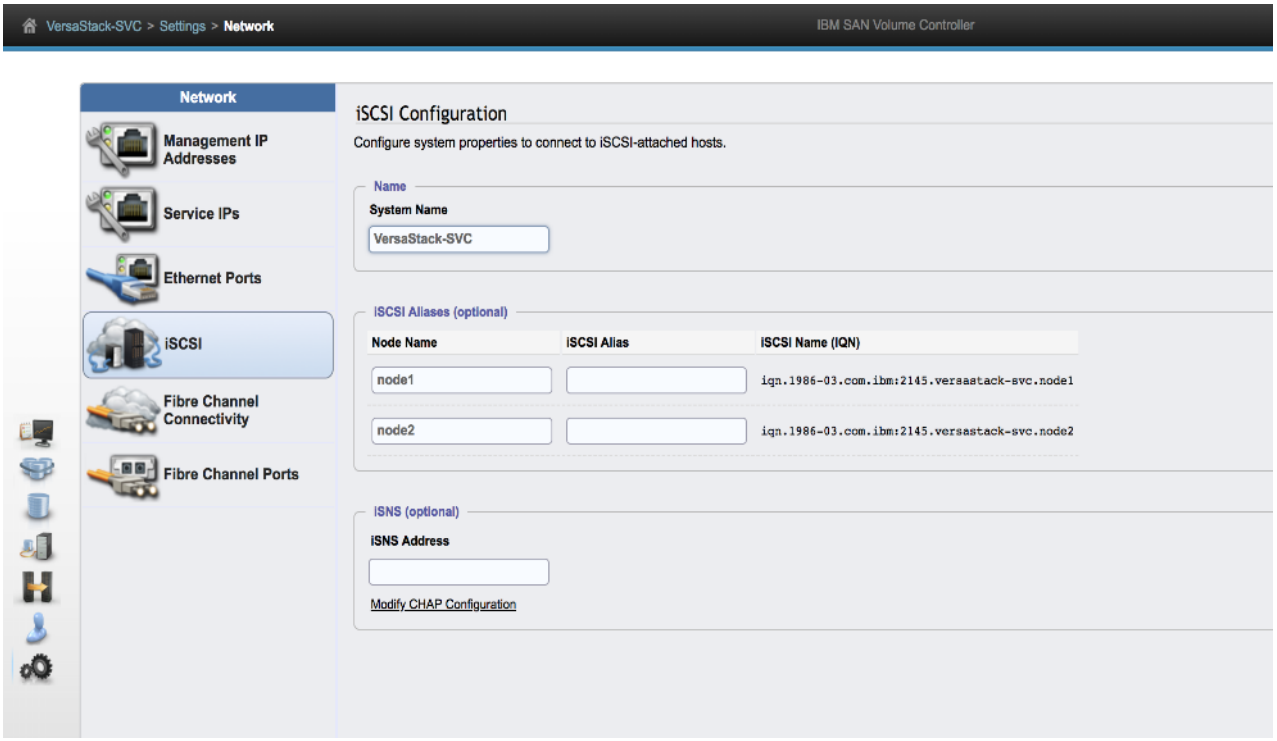

3. Note the resulting iSCSI Name (IQN) in the Table 8 to be used later in the configuration procedure

Table 8 IBM SVC IQN

Node	IQN
Node 1	
Node 2	

4. Click the Ethernet Ports icon

.

5. Click Actions and choose Modify iSCSI Hosts.

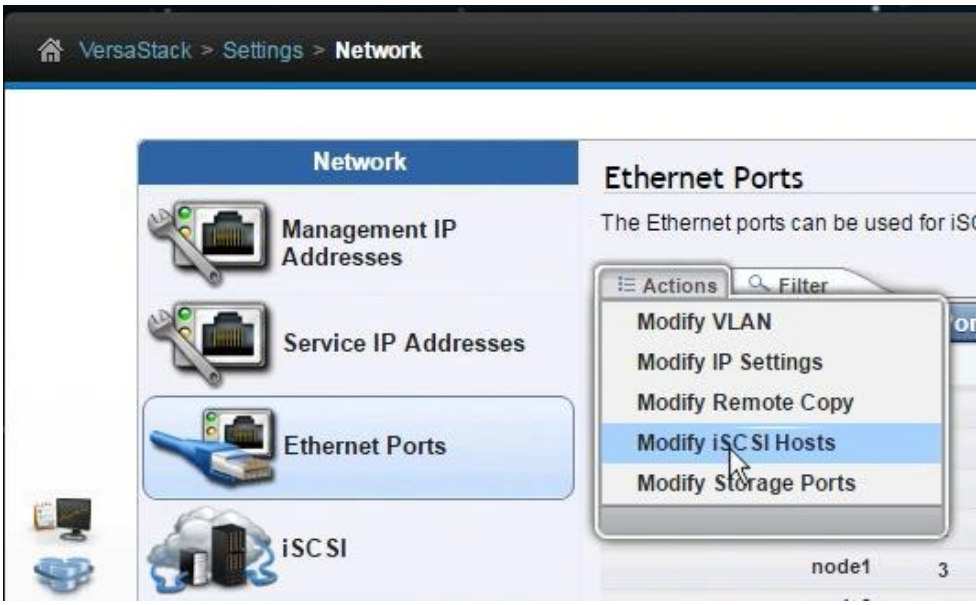

6. Make sure IPv4 iSCSI hosts field is set to enable – if not, change the setting to Enabled and click Modify.

7. If already set, click Cancel to close the configuration box.

8. For each of the four ports listed in Table 7 , repeat the steps 9 – 17.

9. Right-click the appropriate port and choose Modify IP Settings.

10. Enter the IP address, Subnet Mask and Gateway information in Table 7 .

11. Click Modify.

12. Right-click the newly updated port and choose Modify VLAN.

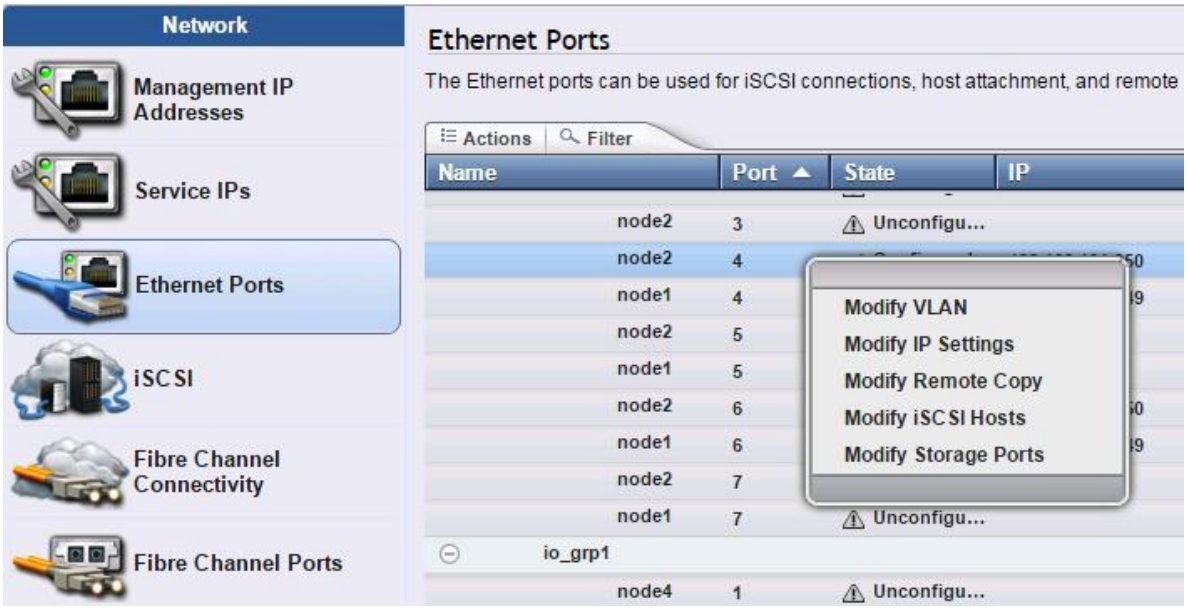

13. Check the check box to Enable VLAN.

14. Enter the appropriate VLAN from Table 7 .

This is only needed if the VLAN is not set as native VLAN in the UCS, do not enable VLAN if the iSCSI VLAN is set as native VLAN.

15. Keep the Apply change to the failover port too check box checked.
16. Click Modify.
17. Repeat the steps for all for iSCSI ports listed in Table 7 .
18. Verify all ports are configured as shown below. The output below shows configuration for two SVC node pairs.

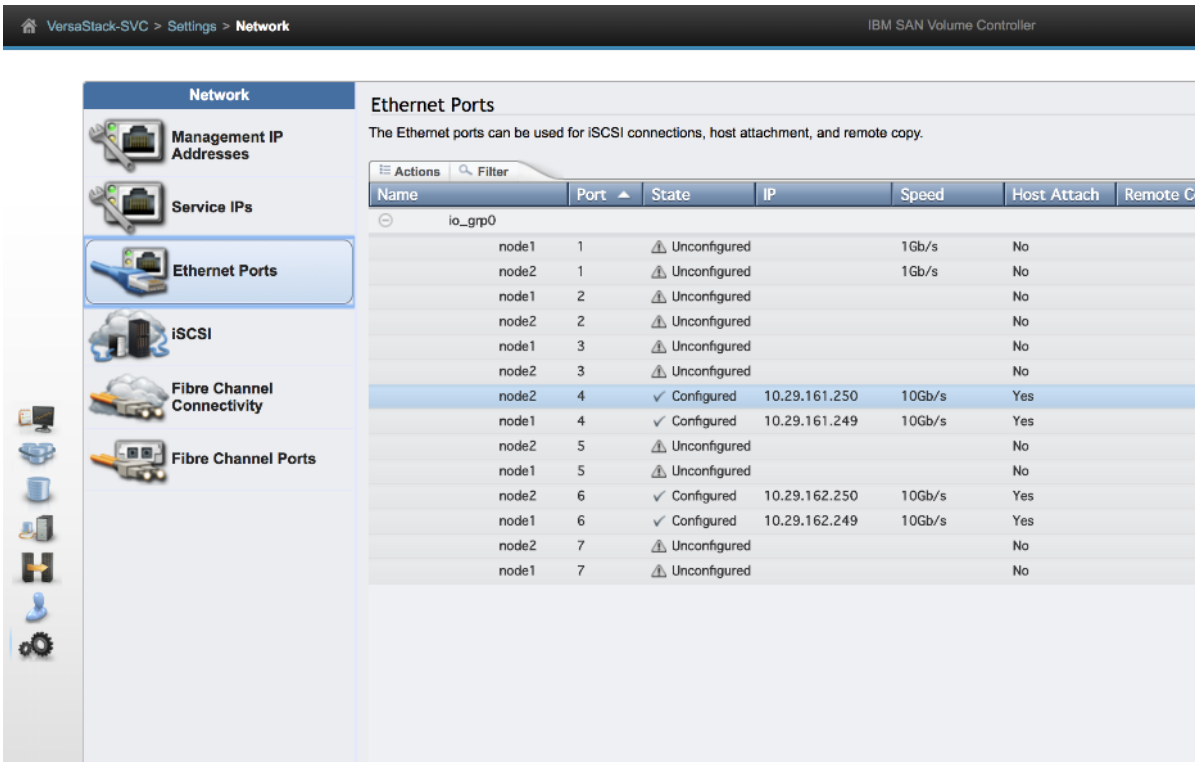

Modifying Interface MTU

Use the cfgportip CLI command to set Jumbo Frames (MTU 9000). The default value of port MTU is 1500. An MTU of 9000 (jumbo frames) provides improved CPU utilization and increased efficiency by reducing the overhead and increasing the size of the payload.

1. In the IBM SVC management GUI, from the Settings options, select System.

2. Select I/O Groups and identify the I/O Group IDs. In this deployment, io_grp0 (ID 0) has being utilized.

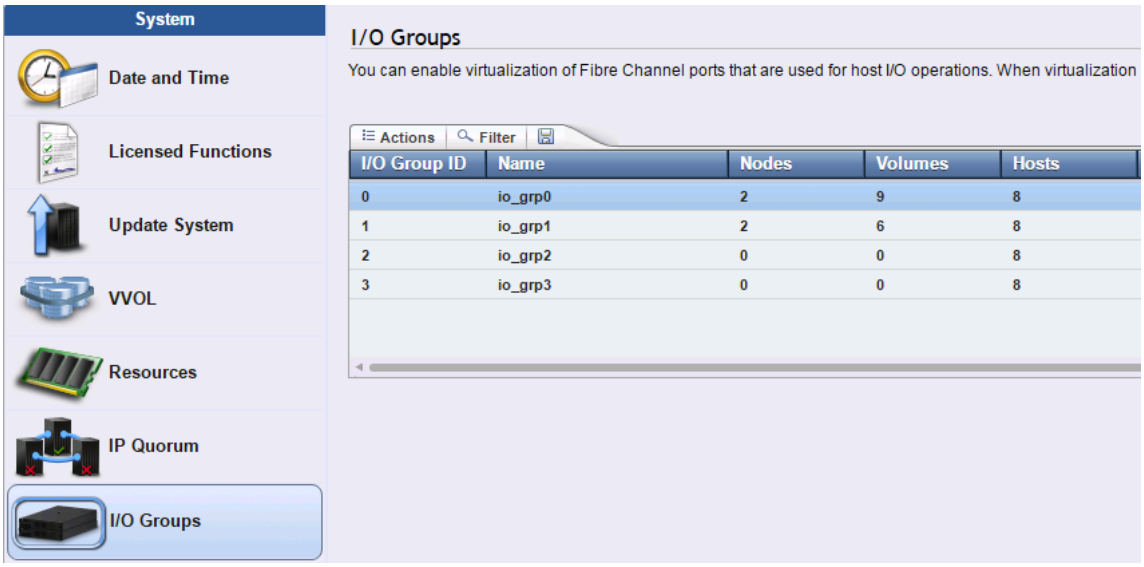

3. SSH to the IBM SVC management IP address and use following CLI command to set the MTU for ports 4 and 6:

```
VersaStack-SVC:admin>cfgportip -mtu 9000 -iogrp 0 4
VersaStack-SVC:admin>cfgportip -mtu 9000 -iogrp 0 6
```

The MTU configuration can be verified using the command: svcinfo lsportip <port number> | grep mtu

This completes the initial configuration of the IBM systems. The next section covers the Cisco UCS configuration.

Cisco UCS Server Configuration

This section covers the Cisco UCS setup for VersaStack infrastructure. This section includes setup for both iSCSI as well as FC SAN boot and storage access.

If a customer environment does require implementing some of the storage protocols covered in this deployment guide, the relevant configuration sections can be skipped.

Cisco UCS Initial Configuration

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a VersaStack environment. The steps are necessary to provision the Cisco UCS C-Series and B-Series servers and should be followed precisely to avoid configuration errors.

Cisco UCS 6332-16UP A

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

1. Connect to the console port on the first Cisco UCS 6332 fabric interconnect.

```
Enter the configuration method. (console/gui) ? console

Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup

You have chosen to setup a new Fabric interconnect? Continue? (y/n): y

Enforce strong password? (y/n) [y]: y

Enter the password for "admin": <password>

Confirm the password for "admin": <password>

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

Which switch fabric (A/B)[]: A

Enter the system name: <Name of the System>

Physical Switch Mgmt0 IP address: <Mgmt. IP address for Fabric A>

Physical Switch Mgmt0 IPv4 netmask: <Mgmt. IP Subnet Mask>

IPv4 address of the default gateway: <Default GW for the Mgmt. IP >

Cluster IPv4 address: <Cluster Mgmt. IP address>

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

DNS IP address: <DNS IP address>

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

Default domain name: <DNS Domain Name>
```

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

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

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

Cisco UCS 6332-16UP B

To configure the second Cisco UCS Fabric Interconnect for use in a VersaStack environment, complete the following steps:

1. Connect to the console port on the second Cisco UCS 6332 fabric interconnect.

```
Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect. This

Fabric interconnect will be added to the cluster.  Continue (y|n)? y

Enter the admin password for the peer Fabric interconnect: <Admin Password>

 Connecting to peer Fabric interconnect... done

 Retrieving config from peer Fabric interconnect... done

 Peer Fabric interconnect Mgmt0 IPv4 Address: <Address provided in last step>

 Peer Fabric interconnect Mgmt0 IPv4 Netmask: <Mask provided in last step>

 Cluster IPv4 address          : <Cluster IP provided in last step>


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

Physical switch Mgmt0 IP address: < Mgmt. IP address for Fabric B>

Apply and save the configuration (select no if you want to re-enter)?

(yes/no): yes
```

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

Cisco UCS Software Upgrade

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6332 fabric interconnect cluster address.
2. Under HTML, click the Launch UCS Manager link to launch the Cisco UCS Manager HTML5 User Interface.
3. When prompted, enter admin as the user name and enter the administrative password.

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

5. Respond to the pop-up on Anonymous Reporting and click OK.

Upgrade Cisco UCS Manager Software to Version 3.2(3d)

This document assumes the use of Cisco UCS 3.2(3d). To upgrade the Cisco UCS Manager software and the UCS 6332 Fabric Interconnect software to version 3.2(3d), refer to Cisco UCS Manager Install and Upgrade Guides.

Configure Cisco UCS Call Home

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

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

2. Select All > Communication Management > Call Home.

3. Change the State to On.

4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

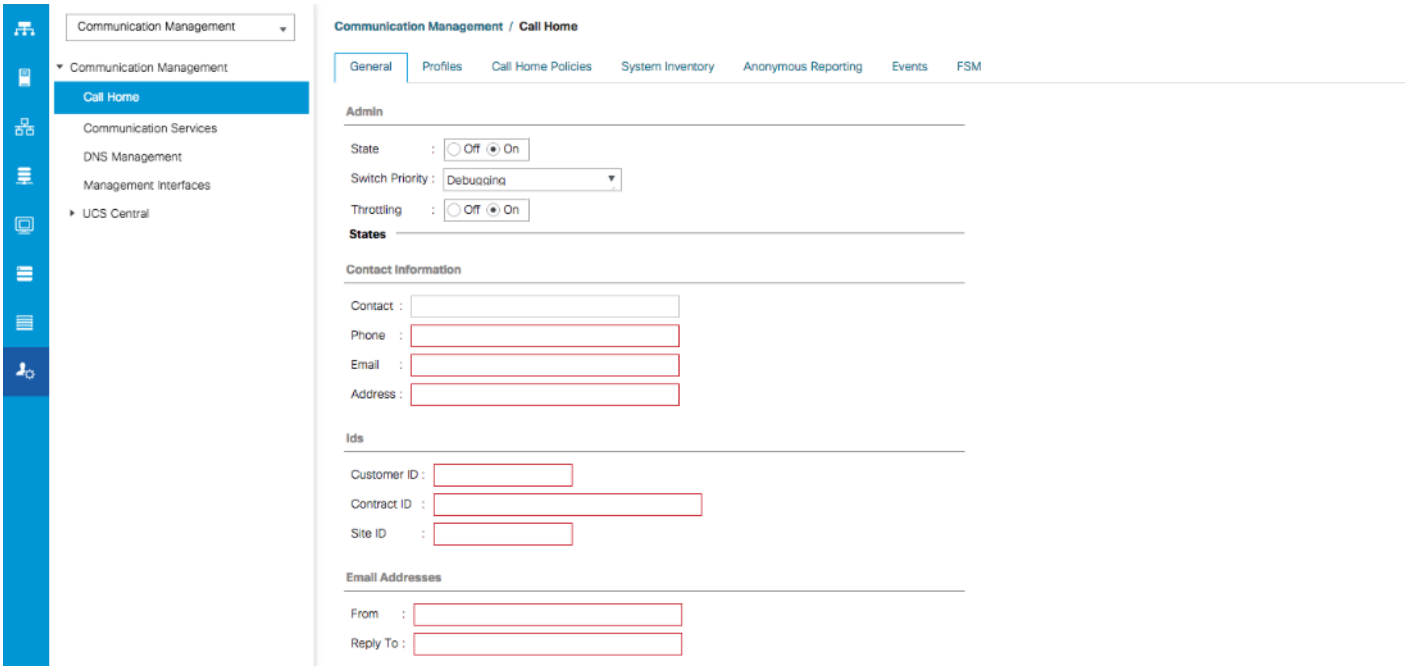

Add a Block of Management IP Addresses for KVM Access

To create a block of IP addresses for out of band (mgmt0) server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Expand Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and choose Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block, the number of IP addresses required, and the subnet and gateway information. Click OK.

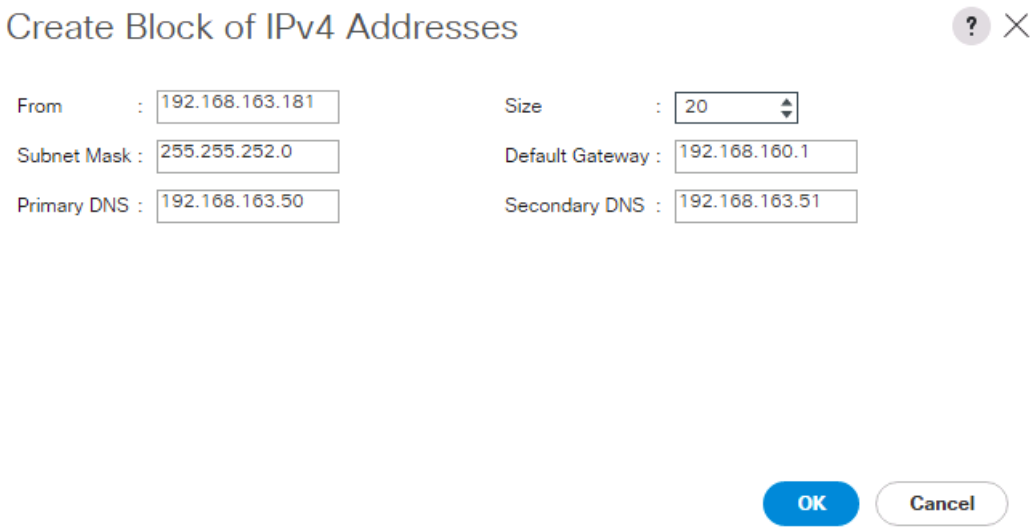

This block of IP addresses should be in the out of band management subnet.

5. Click OK.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, click the Admin tab in the navigation pane.
2. Select All > Timezone Management > 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 <NTP Server IP Address> and click OK.
7. Click OK.

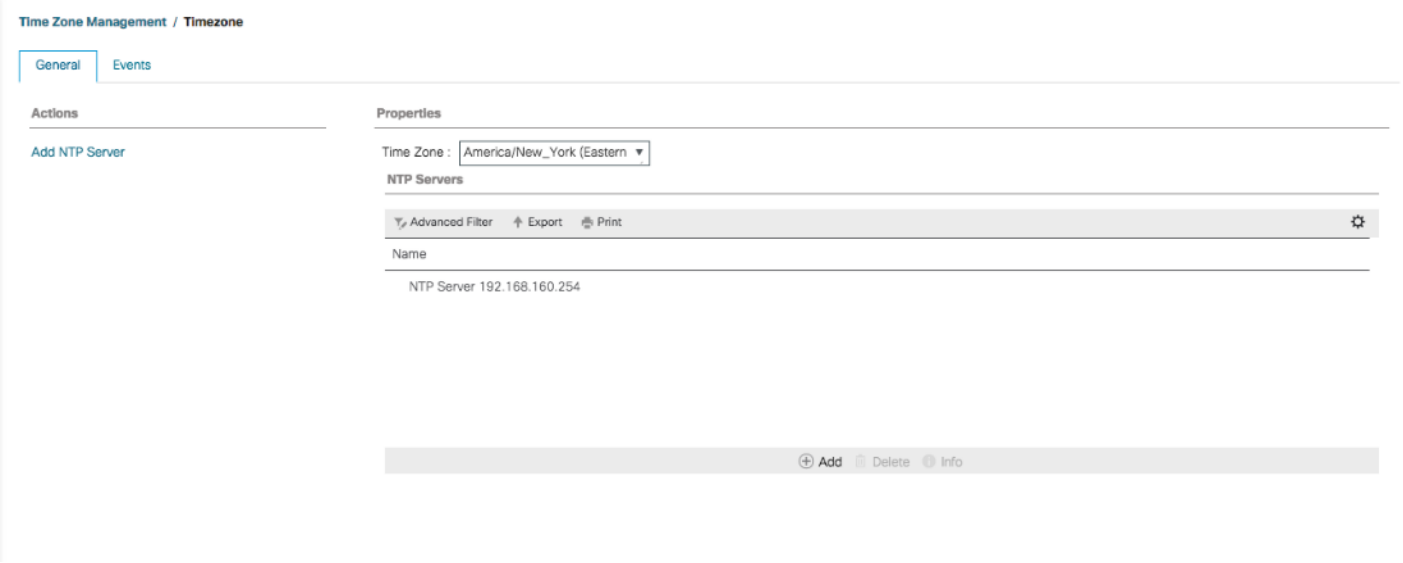

Edit Chassis Discovery Policy

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and select Equipment from the list in the left pane.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum the number of uplink ports that are cabled between any chassis IOM or fabric extender (FEX) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.
5. Click Save Changes.
6. Click OK.

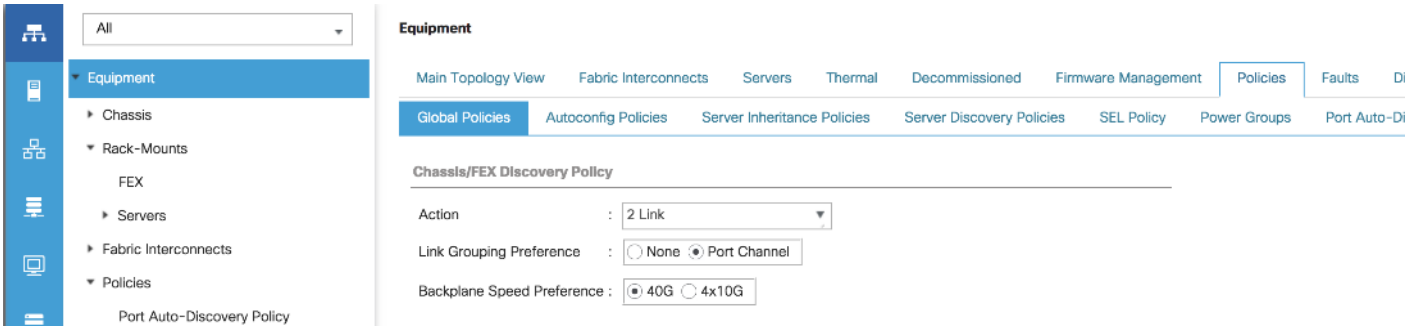

Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Fixed Module.
4. Expand and select Ethernet Ports.
5. Select the ports that are connected to the Cisco UCS 5108 chassis and UCS C-Series servers, one by one, right-click and select Configure as Server Port.

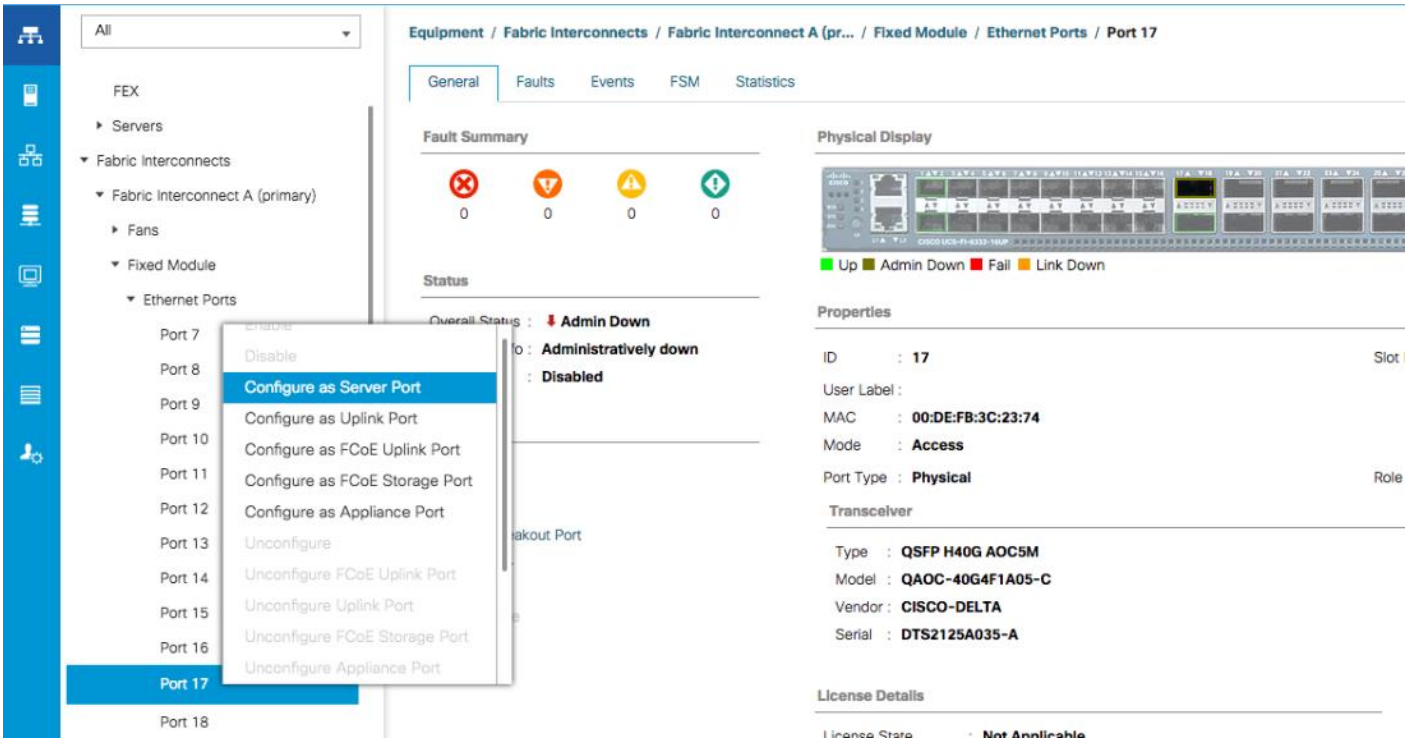

6. Click Yes to confirm server ports and click OK.
7. Verify that the ports connected to the UCS 5108 chassis and C-series servers are now configured as Server ports by selecting Fabric Interconnect A in the left and Physical Ports tab in the right pane.

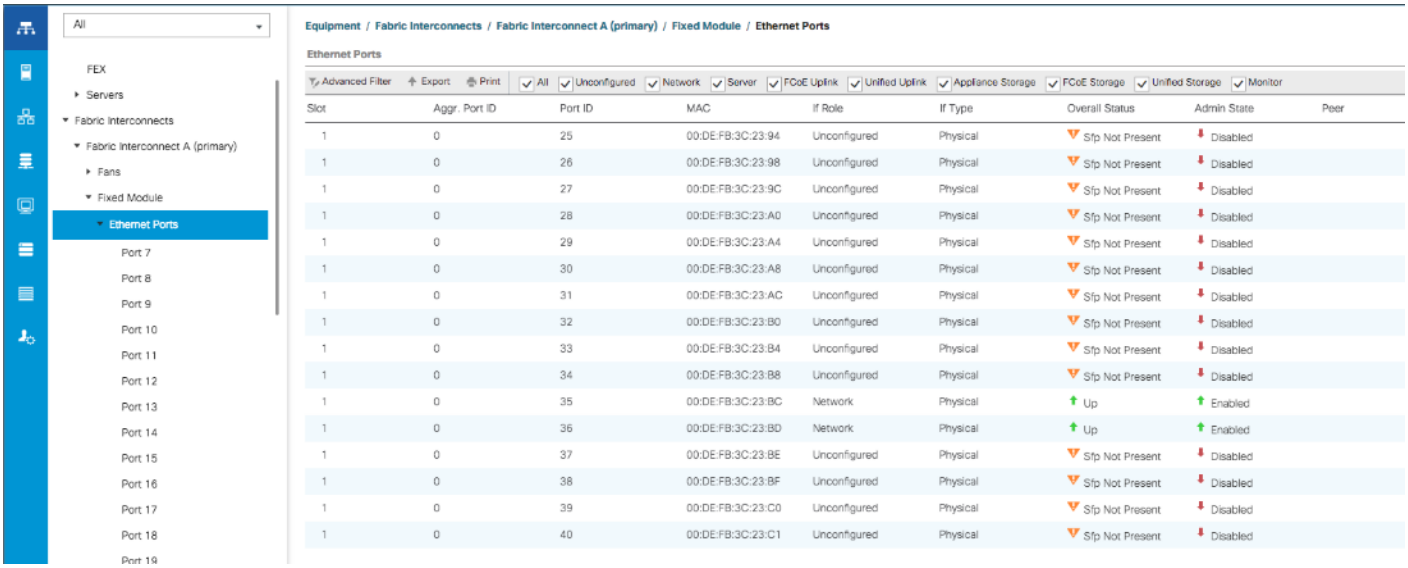

8. Select the ports that are connected to the Cisco Nexus 93180 switches, one by one, right-click and select Configure as Uplink Port.

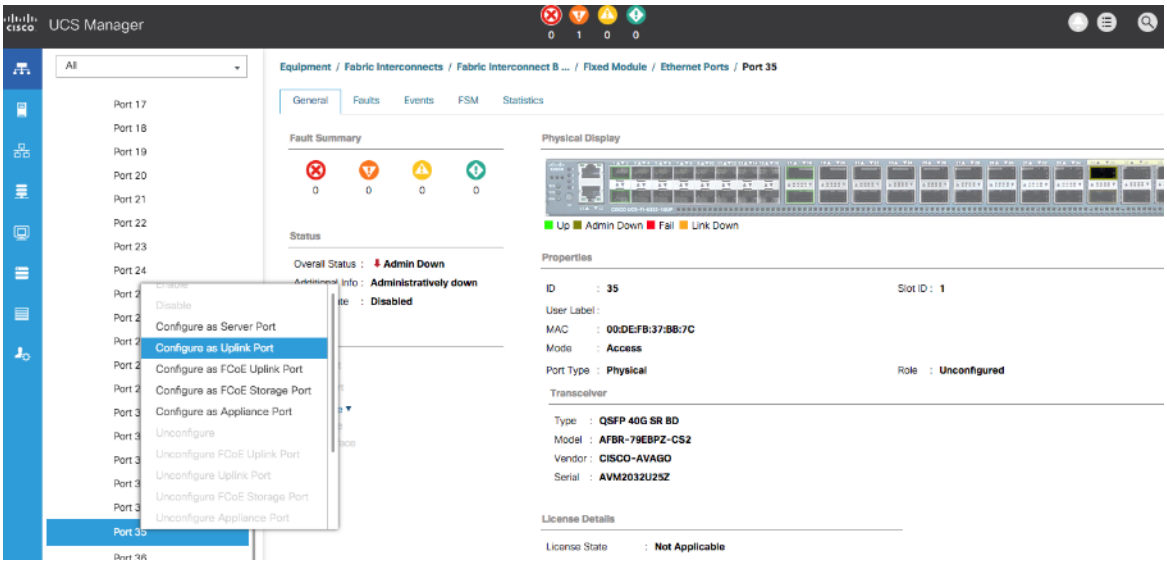

9. Click Yes to confirm uplink ports and click OK.

10. Verify that the uplink ports are now configured as Network ports by selecting Fabric Interconnect A in the left and Physical Ports tab in the right pane.

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

12. Repeat steps 3-10 to configure server and uplink ports on Fabric Interconnect B.

Acknowledge Cisco UCS Chassis and FEX

When the UCS FI ports are configured as server ports, UCS chassis is automatically discovered and may need to be acknowledged. To acknowledge all Cisco UCS chassis, complete the following 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.

If Cisco Nexus 2232PP FEXes are part of the configuration, expand Rack-Mounts and FEX and acknowledge the FEXes one by one.

Enable Fibre Channel Ports

The FC port and uplink configurations can be skipped if the UCS environment does not need access to storage environment using FC

To enable FC uplink ports, complete the following steps.

This step requires a reboot. To avoid an unnecessary switchover, configure the subordinate Fabric Interconnect first.

1. In the Equipment tab, select the Fabric Interconnect B (subordinate FI in this example), and in the Actions pane, select Configure Unified Ports, and click Yes on the splash screen.

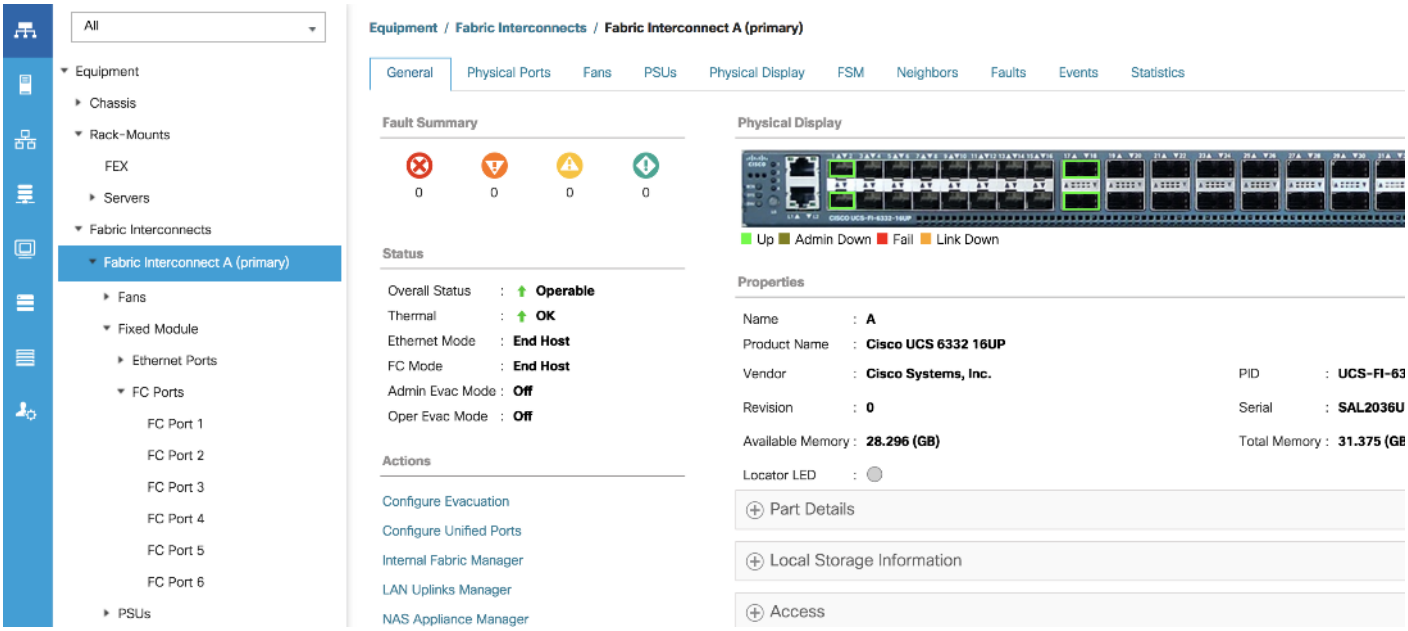

2. Slide the lever to change the ports 31-32 to Fiber Channel. Click Finish followed by Yes to the reboot message. Click OK.

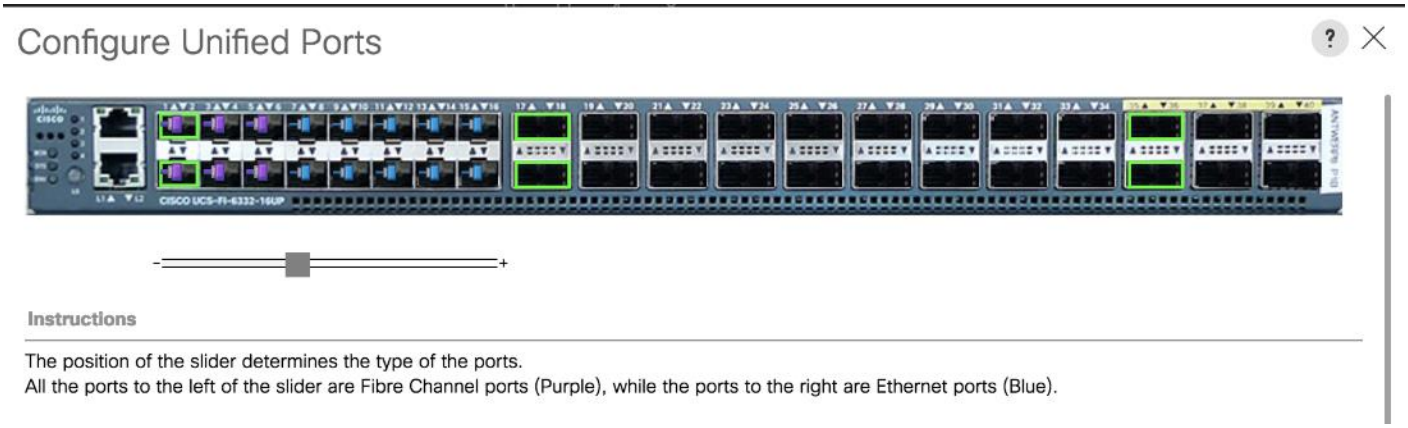

3. When the subordinate has completed reboot, repeat the procedure to configure FC ports on primary Fabric Interconnect. As before, the Fabric Interconnect will reboot after the configuration is complete.

Create VSAN for the Fibre Channel Interfaces

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

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

2. Expand the SAN > SAN Cloud and select Fabric A.

3. Right-click VSANs and choose Create VSAN.

4. Enter VSAN-A as the name of the VSAN for fabric A.
5. Keep the Disabled option selected for FC Zoning.
5. Click the Fabric A radio button.
6. Enter 101 as the VSAN ID for Fabric A.
7. Enter 101 as the FCoE VLAN ID for fabric A. Click OK twice.

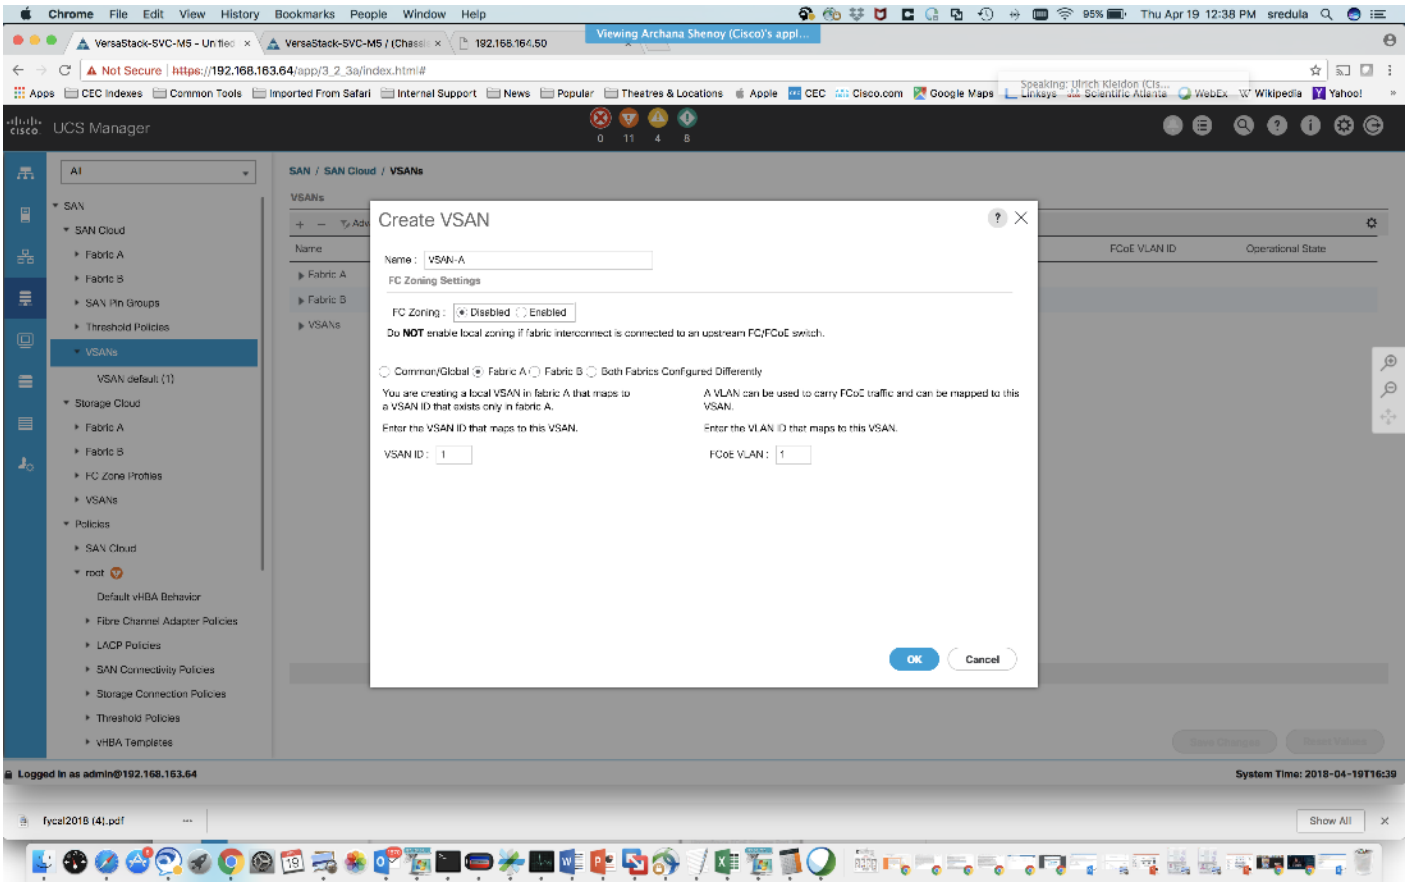

8. In the SAN tab, expand SAN > SAN Cloud > Fabric-B.
9. Right-click VSANs and choose Create VSAN.
10. Enter VSAN-B as the name of the VSAN for fabric B.
11. Keep the Disabled option selected for FC Zoning.
12. Click the Fabric B radio button.
13. Enter 102 as the VSAN ID for Fabric B. Enter 102 as the FCoE VLAN ID for Fabric B. Click OK twice.

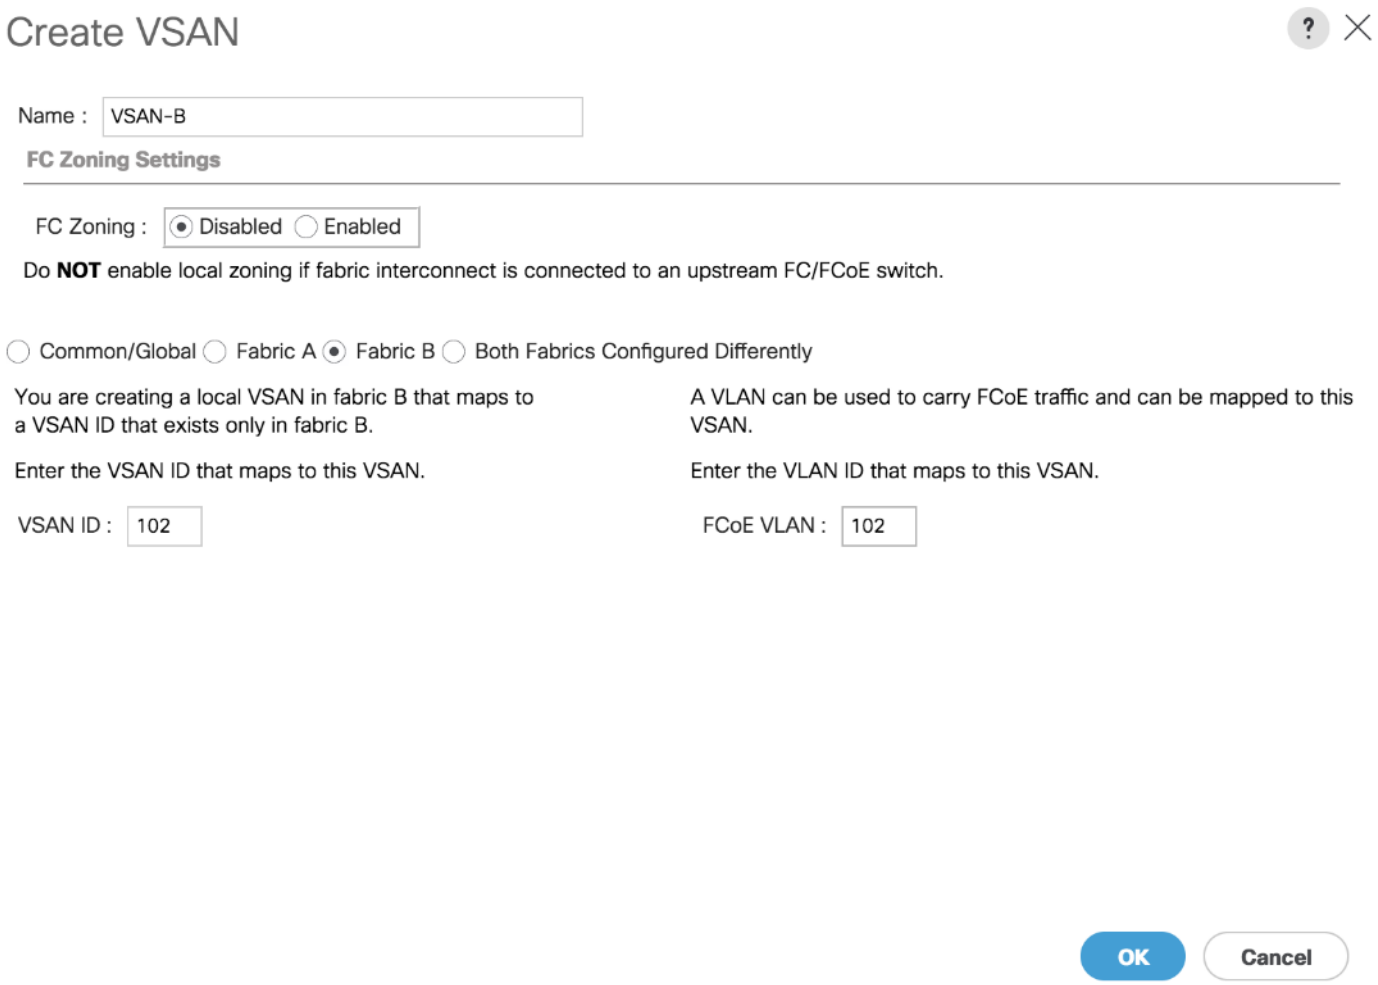

Create Port Channels for the Fibre Channel Interfaces

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

Fabric-A

1. In the navigation pane, under SAN > SAN Cloud, expand the Fabric A tree.
2. Right-click FC Port Channels and choose Create Port Channel.
3. Enter 6 for the port channel ID and Po6 for the port channel name.
4. Click Next then choose ports 31 and 32 and click >> to add the ports to the port channel. Click Finish.

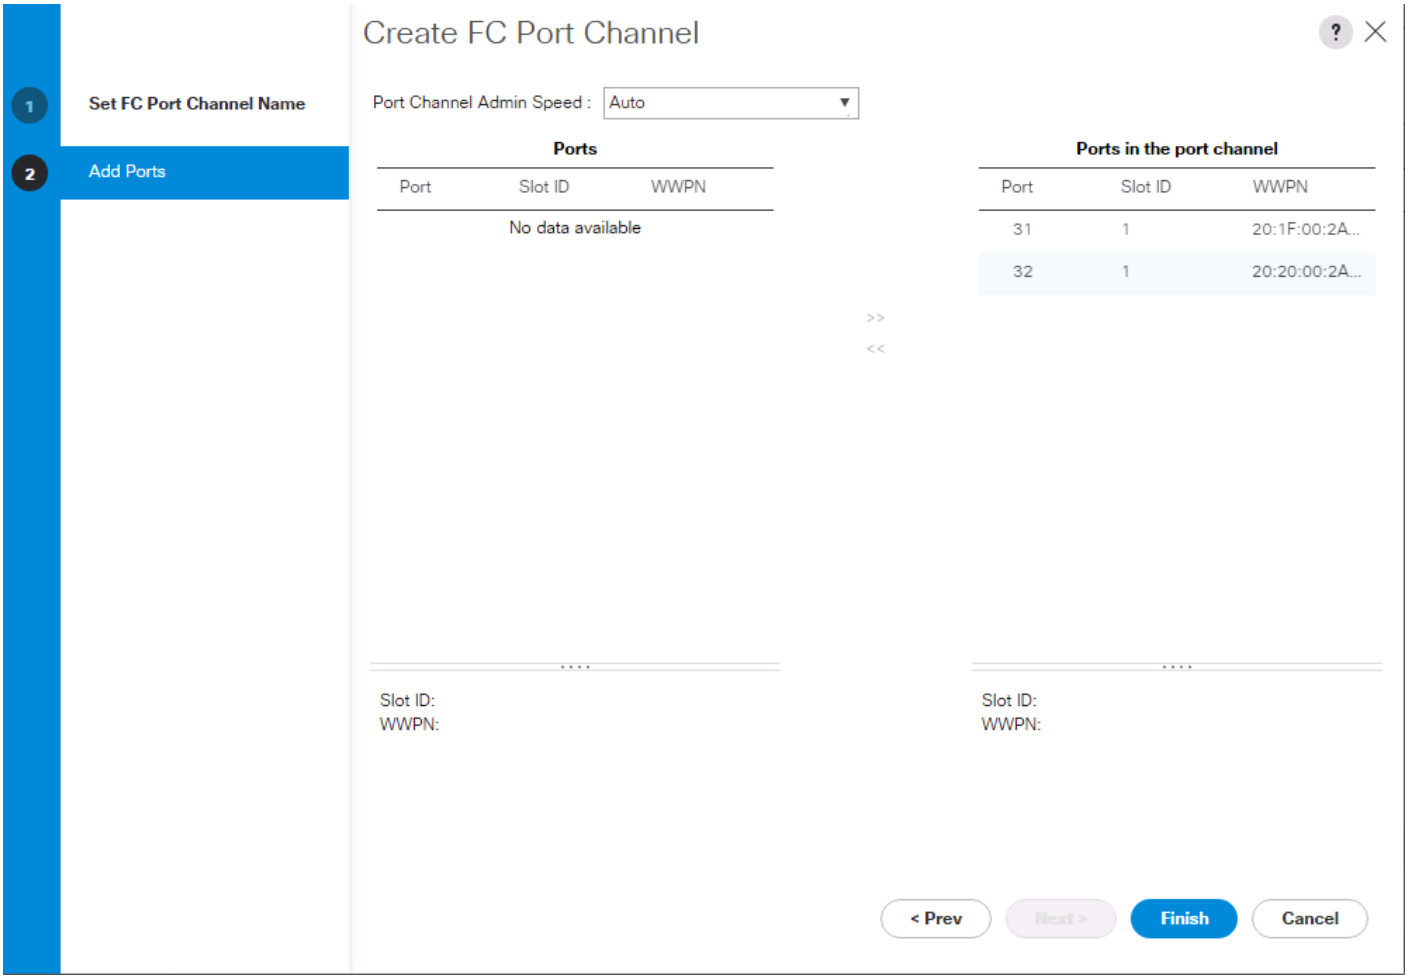

5. Click OK.

6. Select FC Port-Channel 6 from the menu in the left pane and from the VSAN drop-down field, select VSAN 101 in the right pane.

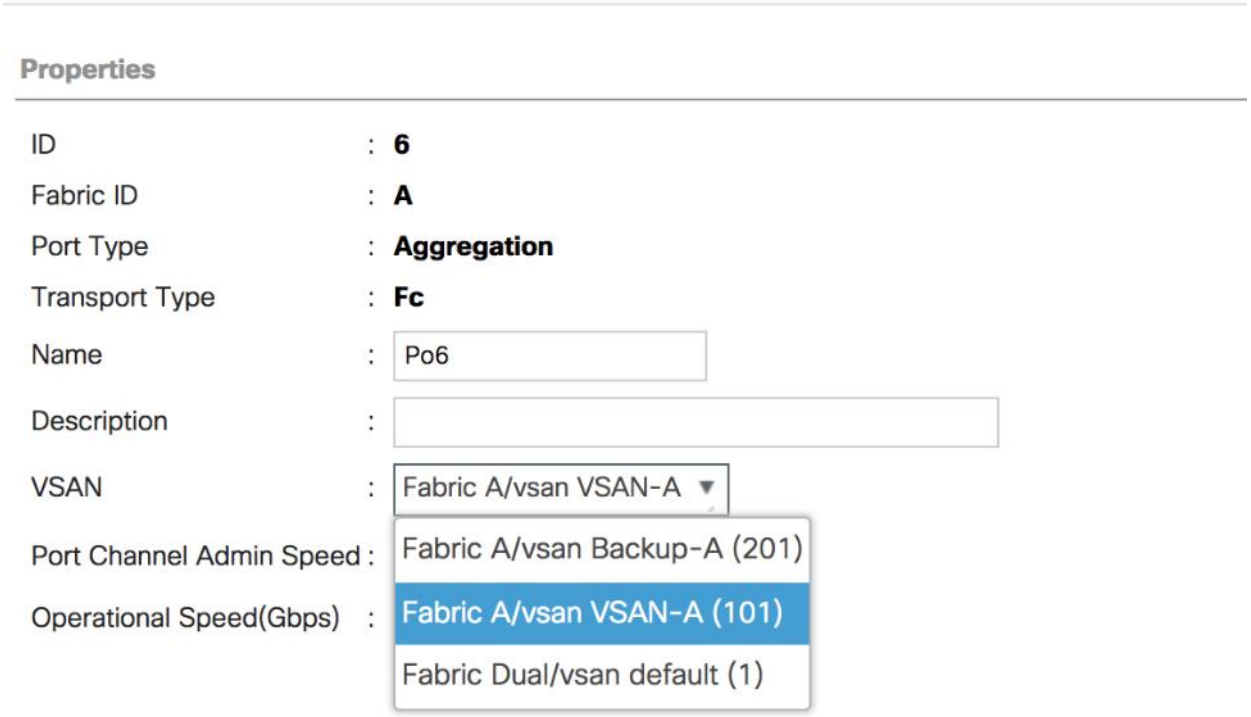

7. Click Save Changes and then click OK.

Fabric-B

1. Click the SAN tab. In the navigation pane, under SAN > SAN Cloud, expand the Fabric B.
2. Right-click FC Port Channels and choose Create Port Channel.
3. Enter 7 for the port channel ID and Po7 for the port channel name. Click Next.
4. Choose ports 31 and 32 and click >> to add the ports to the port channel.
5. Click Finish, and then click OK.
6. Select FC Port-Channel 7 from the menu in the left pane and from the VSAN drop-down list, select VSAN 102 in the right pane.
7. Click Save Changes and then click OK.

To initialize a quick sync of the connections to the MDS switch, right-click the recently created port channels, disable port channel and then re-enable the port channel.

Create Port Channels for Ethernet Uplinks

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

In this procedure, two port channels are created one from each Fabric Interconnect (A and B) to both the Cisco Nexus 93180 switches.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels and choose Create Port Channel.
4. Enter 13 as the unique ID of the port channel.
5. Enter Po13 as the name of the port channel and click Next.

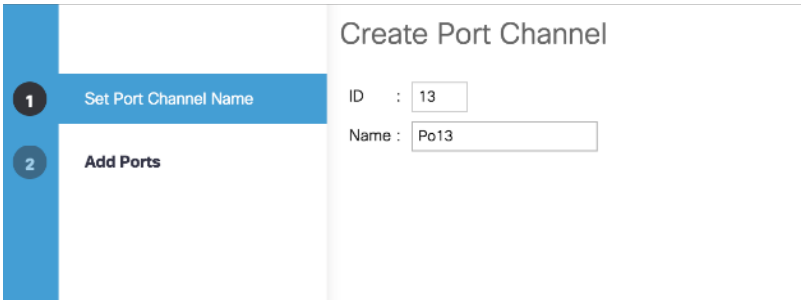

6. Select the network uplink ports to be added to the port channel.

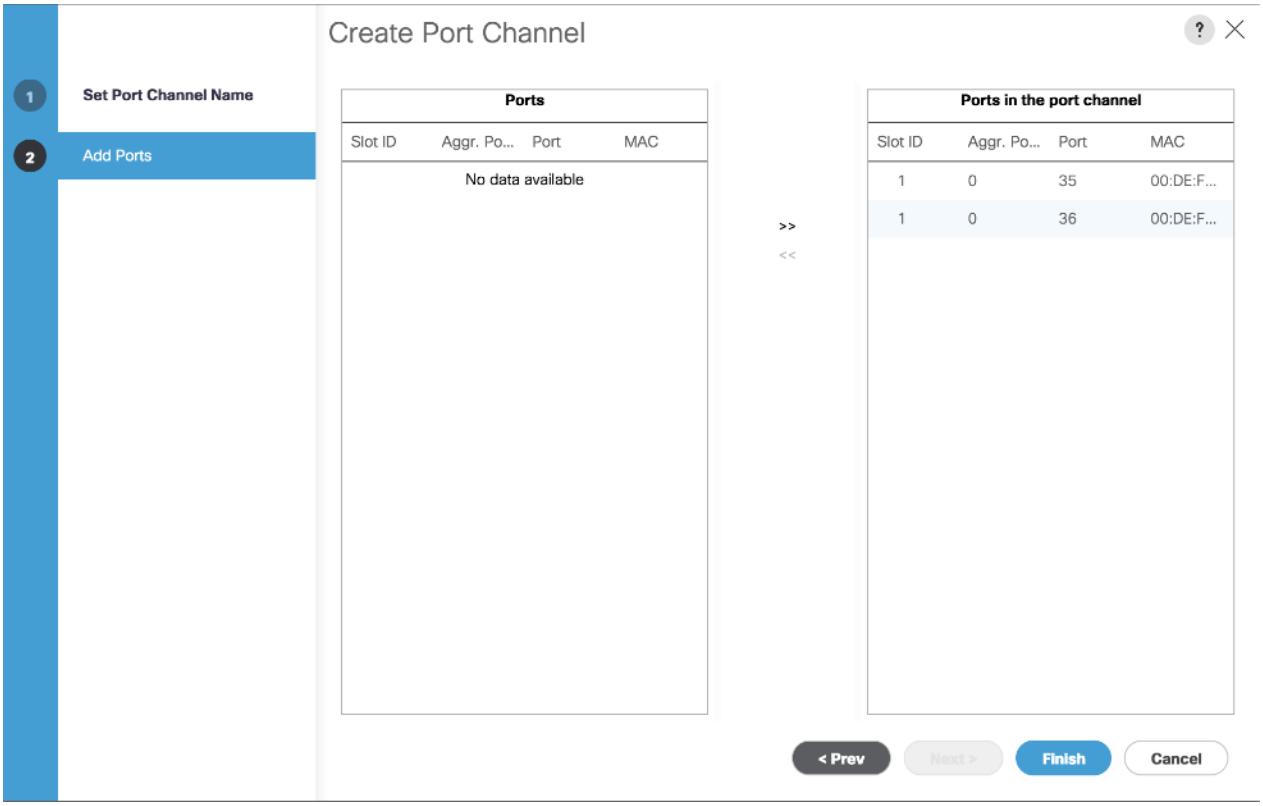

7. Click >> to add the ports to the port channel (35 and 36 in this design).
8. Click Finish to create the port channel and then click OK.
9. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.
10. Right-click Port Channels and choose Create Port Channel.
11. Enter 14 as the unique ID of the port channel.
12. Enter Po14 as the name of the port channel and click Next.
13. Select the network uplink ports (35 and 36 in this design) to be added to the port channel.
14. Click >> to add the ports to the port channel.
15. Click Finish to create the port channel and click OK.

Since the ACI fabric is not configured as yet, the port channels will remain in down state.

Create MAC Address Pools

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

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

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 the option Sequential for the Assignment Order field and click Next.

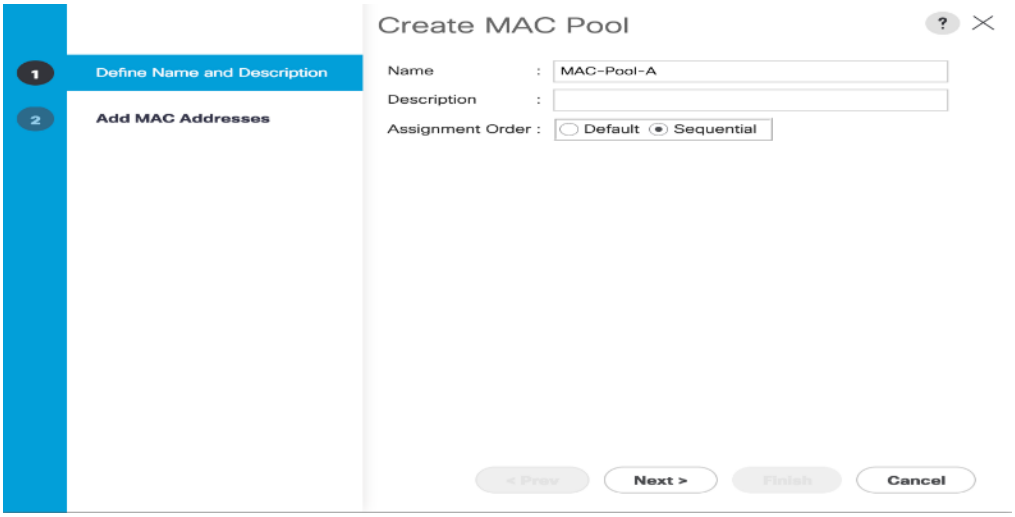

8. Click Add.
9. Specify a starting MAC address.

It is recommended to place 0A in the second last octet of the starting MAC address to identify all of the MAC addresses as Fabric A addresses. It is also recommended to not change the first three octets of the MAC address.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or rack server resources. Remember that multiple Cisco VIC vNICs will be created on each server and each vNIC will be assigned a MAC address.

11. Click OK and then click Finish.
12. In the confirmation message, click OK.
13. Right-click MAC Pools under the root organization.
14. Select Create MAC Pool to create the MAC address pool.
15. Enter MAC-Pool-B as the name of the MAC pool.
16. Optional: Enter a description for the MAC pool.
17. Select the Sequential Assignment Order and click Next.
18. Click Add.
19. Specify a starting MAC address.

It is recommended to place 0B in the second last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. It is also recommended to not change the first three octets of the MAC address.

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

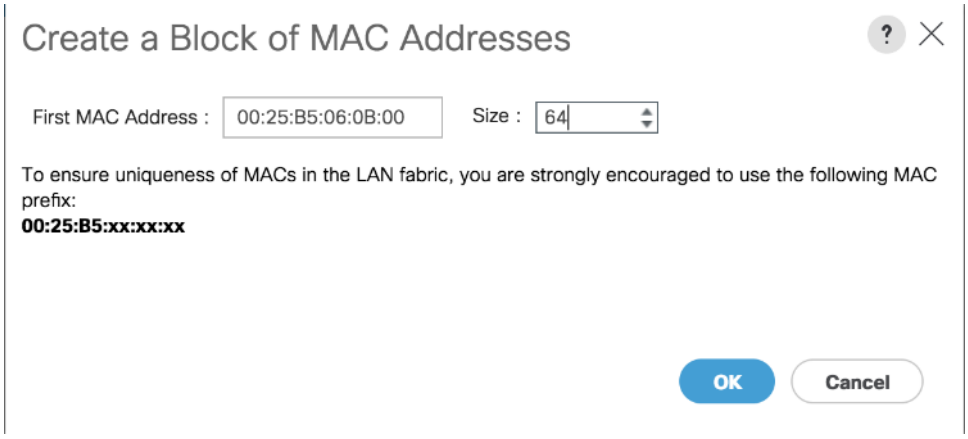

21. Click OK and then click Finish.

22. 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, complete the following steps:

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

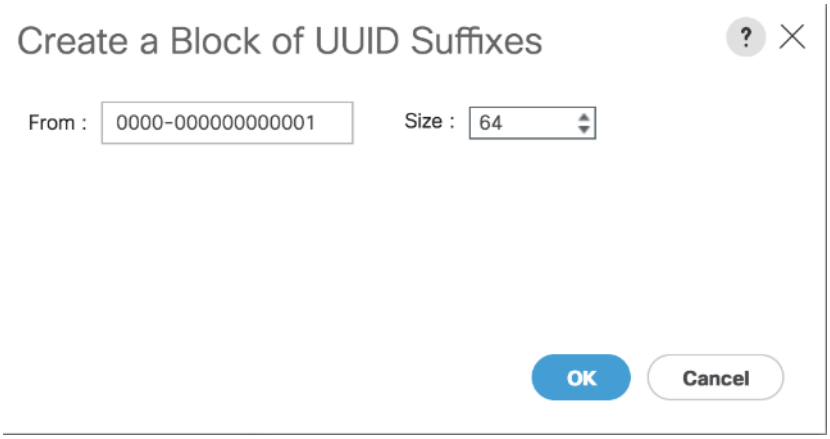

12. Click OK. Click Finish and then click OK.

Create Server Pool

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

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

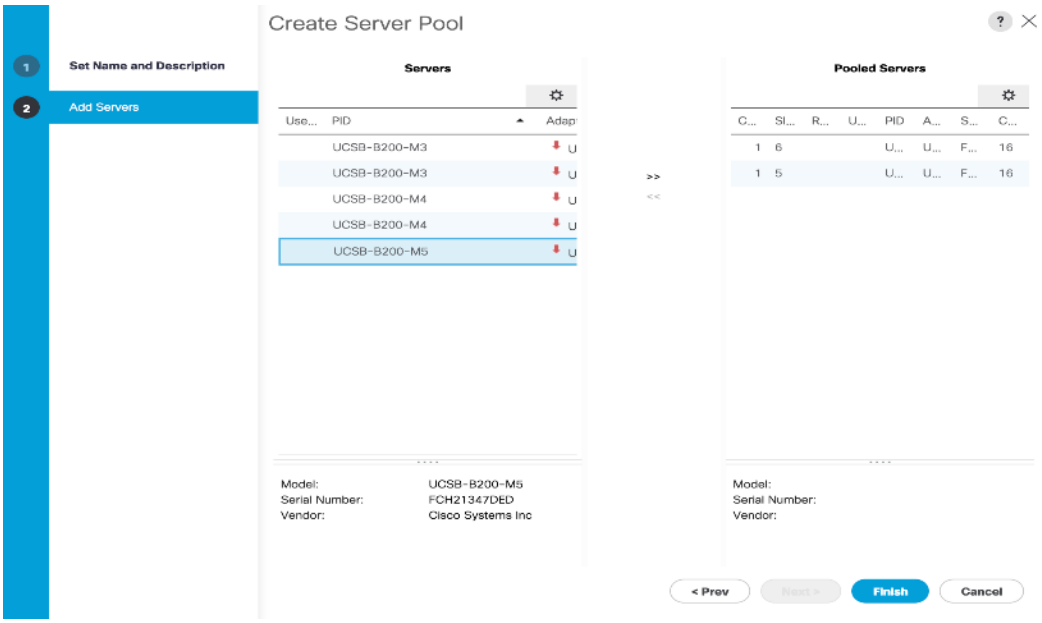

7. Select at least two (or more) servers to be used for the setting up the VMware environment and click >> to add them to the Infra-Server-Pool server pool.

8. Click Finish and click OK.

Create a WWNN Address Pool for FC based Storage Access

This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

For FC boot as well as access to FC LUNs, create a World Wide Node Name (WWNN) pool by completing the following steps:

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

2. Select Pools > root.

3. Right-click WWNN Pools under the root organization and choose Create WWNN Pool to create the WWNN address pool.

4. Enter WWNN-Pool as the name of the WWNN pool.

5. Optional: Enter a description for the WWNN pool.

6. Select the Sequential Assignment Order and click Next.

7. Click Add.

8. Specify a starting WWNN address.

9. Specify a size for the WWNN address pool that is sufficient to support the available blade or rack server resources. Each server will receive one WWNN.

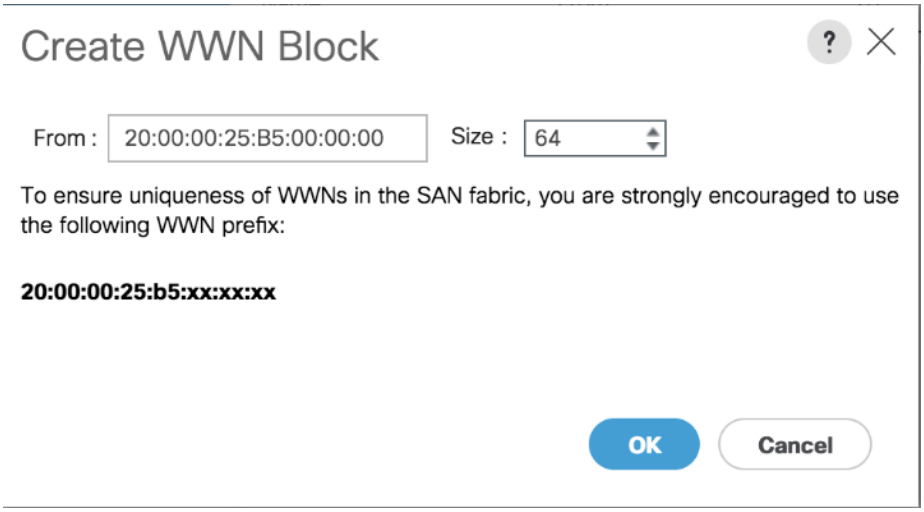

10. Click OK and click Finish.

11. In the confirmation message, click OK.

Create a WWPN Address Pools for FC Based Storage Access

This configuration step can be skipped if the UCS environment does not need access to storage environment using FC.

If you are providing FC boot or access to FC LUNs, create a World Wide Port Name (WWPN) pool for each SAN switching fabric by completing the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. Right-click WWPN Pools under the root organization and choose Create WWPN Pool to create the first WWPN address pool.
4. Enter WWPN-Pool-A as the name of the WWPN pool.
5. Optional: Enter a description for the WWPN pool.
6. Select the Sequential Assignment Order and click Next.
7. Click Add.
8. Specify a starting WWPN address.

It is recommended to place 0A in the second last octet of the starting WWPN address to identify all of the WWPN addresses as Fabric A addresses.

9. Specify a size for the WWPN address pool that is sufficient to support the available blade or rack server resources. Each server's Fabric A vHBA will receive one WWPN from this pool.

10. Click OK and click Finish.
11. In the confirmation message, click OK.
12. Right-click WWPN Pools under the root organization and choose Create WWPN Pool to create the second WWPN address pool.
13. Enter WWPN-Pool-B as the name of the WWPN pool.
14. Optional: Enter a description for the WWPN pool.
15. Select the Sequential Assignment Order and click Next.
16. Click Add.
17. Specify a starting WWPN address.

It is recommended to place 0B in the second last octet of the starting WWPN address to identify all of the WWPN addresses as Fabric B addresses.

18. Specify a size for the WWPN address pool that is sufficient to support the available blade or rack server resources. Each server's Fabric B vHBA will receive one WWPN from this pool.

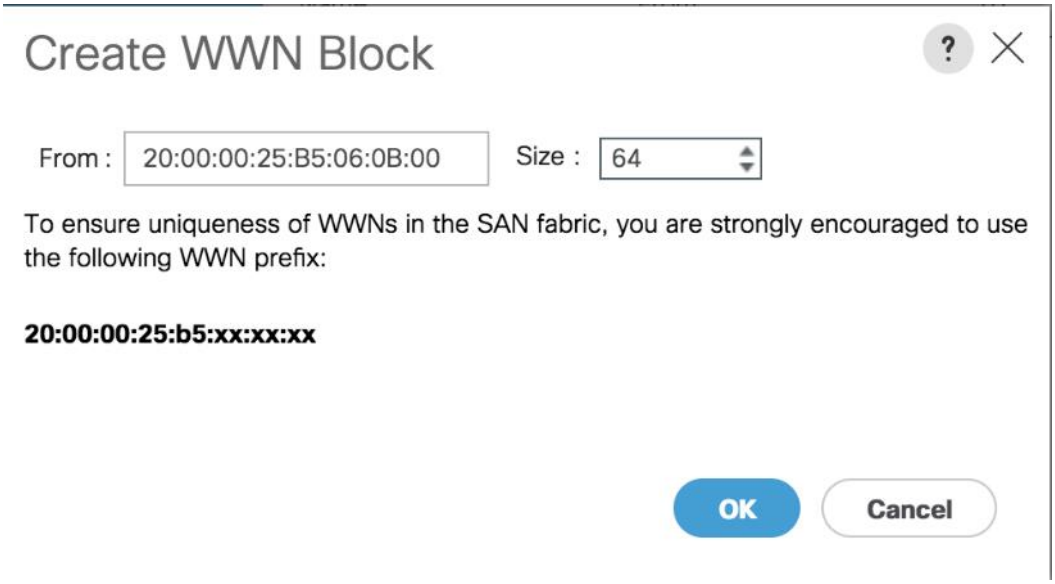

19. Click OK and click Finish.
20. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot and LUN Access

This configuration step can be skipped if the UCS environment does not need access to storage environment using iSCSI.

To enable iSCSI boot and provide access to iSCSI LUNs, configure the necessary IQN pools in the Cisco UCS Manager by completing the following steps:

1. In the UCS Manager, select the SAN tab.
2. Select Pools > root.
3. Right-click IQN Pools under the root organization and choose Create IQN Suffix Pool to create the IQN pool.
4. Enter Infra-IQN-Pool for the name of the IQN pool.
5. Optional: Enter a description for the IQN pool.
6. Enter iqn.1992-08.com.cisco as the prefix
7. Select the option Sequential for Assignment Order field. Click Next.

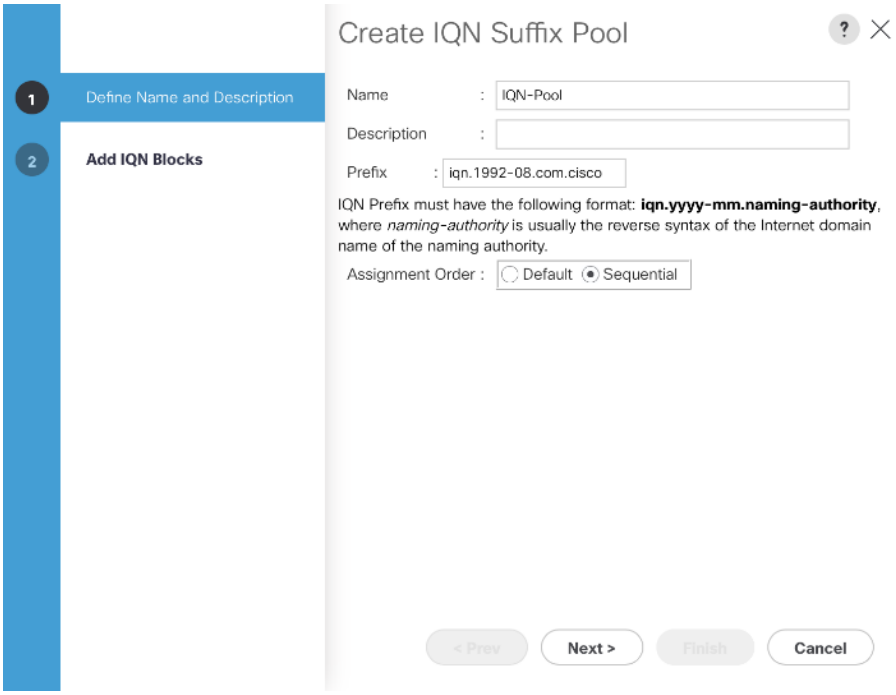

8. Click Add.
9. Enter an identifier with ucs-host as the suffix. A rack number can be added to the suffix to make the IQN unique within a DC.
10. Enter 1 in the From field.

11. Specify a size of the IQN block sufficient to support the available server resources. Each server will receive one IQN.
12. Click OK.

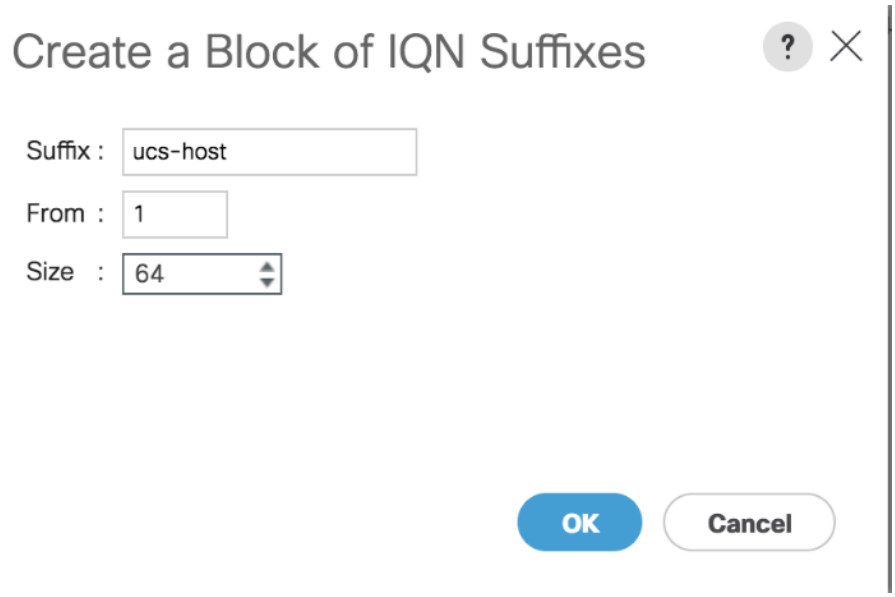

13. Click Finish. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot and LUN Access

This configuration step can be skipped if the UCS environment does not need access to storage environment using iSCSI.

For enabling iSCSI storage access, these steps provide details for configuring the necessary IP pools in the Cisco UCS Manager:

Two IP pools are created, one for each switching fabric.

1. In Cisco UCS Manager, select the LAN tab.
2. Select Pools > root.
3. Right-click IP Pools under the root organization and choose Create IP Pool to create the IP pool.
4. Enter iSCSI-initiator-A for the name of the IP pool.
5. Optional: Enter a description of the IP pool.
6. Select the option Sequential for the Assignment Order field. Click Next.
7. Click Add.

8. In the From field, enter the beginning of the range to assign an iSCSI IP addresses. These addresses are covered in Table 2 .
9. Enter the Subnet Mask.
10. Set the size with sufficient address range to accommodate the servers. Click OK.

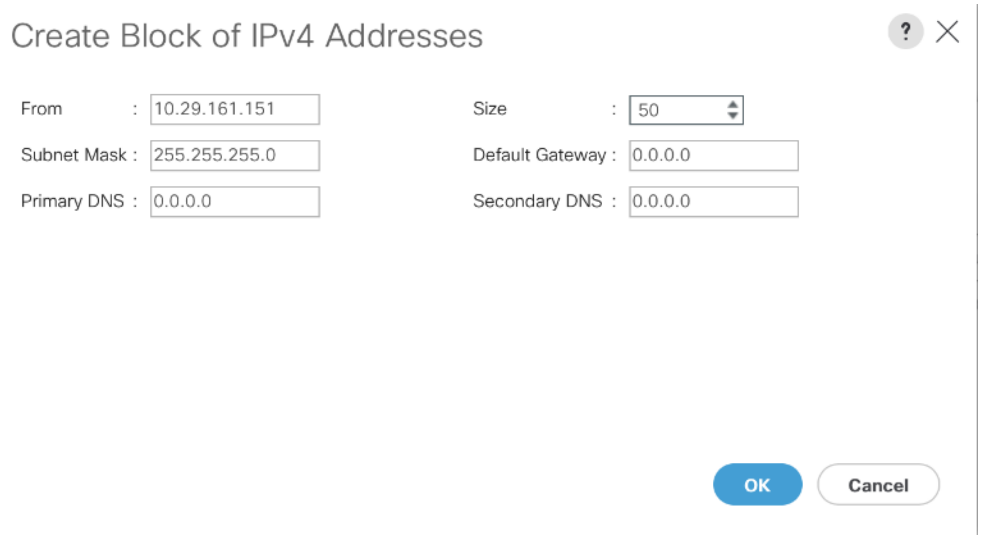

11. Click Next and then click Finish.
12. Click OK in the confirmation message.
13. Right-click IP Pools under the root organization and choose Create IP Pool to create the IP pool.
14. Enter iSCSI-initiator-B for the name of the IP pool.
15. Optional: Enter a description of the IP pool.
16. Select the Sequential option for the Assignment Order field. Click Next.
17. Click Add.
18. In the From field, enter the beginning of the range to assign an iSCSI IP addresses. These addresses are covered in Table 2 .
19. Enter the Subnet Mask.
20. Set the size with sufficient address range to accommodate the servers. Click OK.

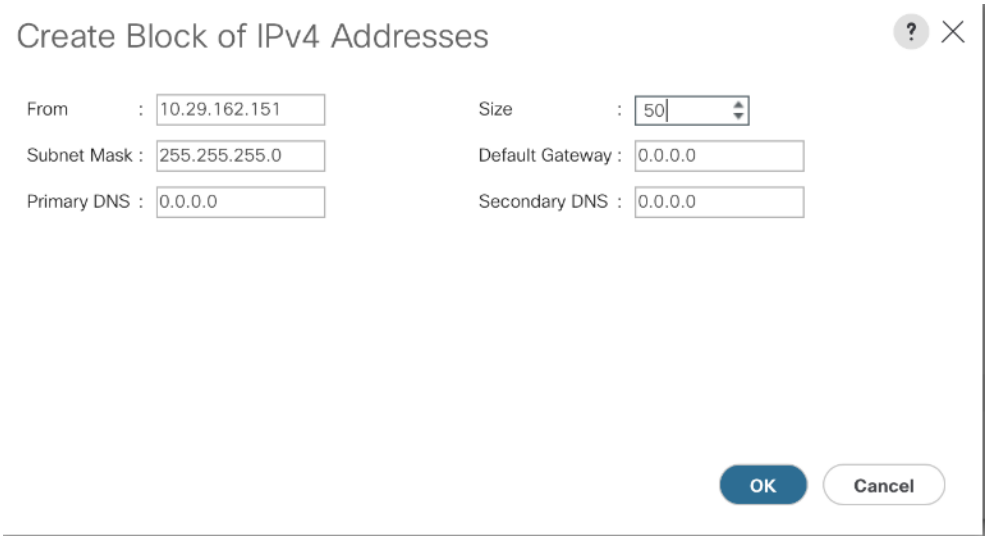

21. Click Next and then click Finish.

22. Click OK in the confirmation message.

Create VLANs

To configure the necessary VLANs in the Cisco UCS Manager, complete the following steps for all the VLANs listed in Table 9 :

Table 9 VLANs on Cisco UCS

VLAN Name	VLAN
IB-Mgmt	11
Infra-iSCSI-A*	3161
infra-iSCSI-B*	3162
Out of Band Mgmt	3171
VM Traffic	3174
vMotion	3173
Native-2	2

* Infra-iSCSI-A/B VLANs are required for iSCSI deployments only.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLANs and choose Create VLANs.

4. Enter name from the VLAN Name column.
5. Keep the Common/Global option selected for the scope of the VLAN.
6. Enter the VLAN ID associated with the name.
7. Keep the Sharing Type as None.
8. Click OK, and then click OK again.

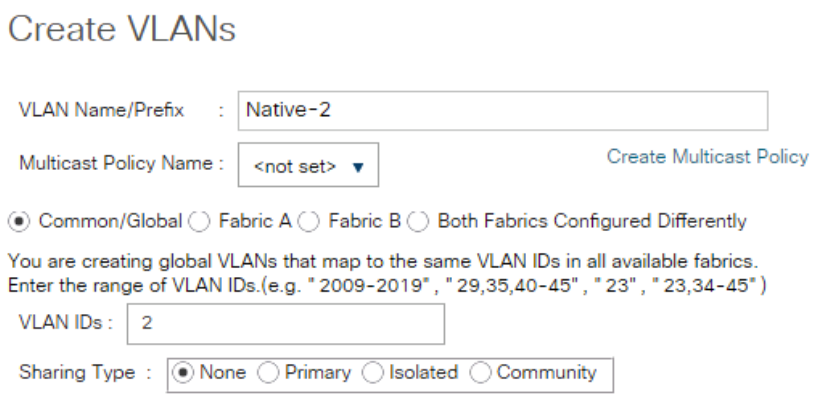

9. Click Yes, and then click OK twice.
10. Repeat these steps for all the VLAN in Table 9 .

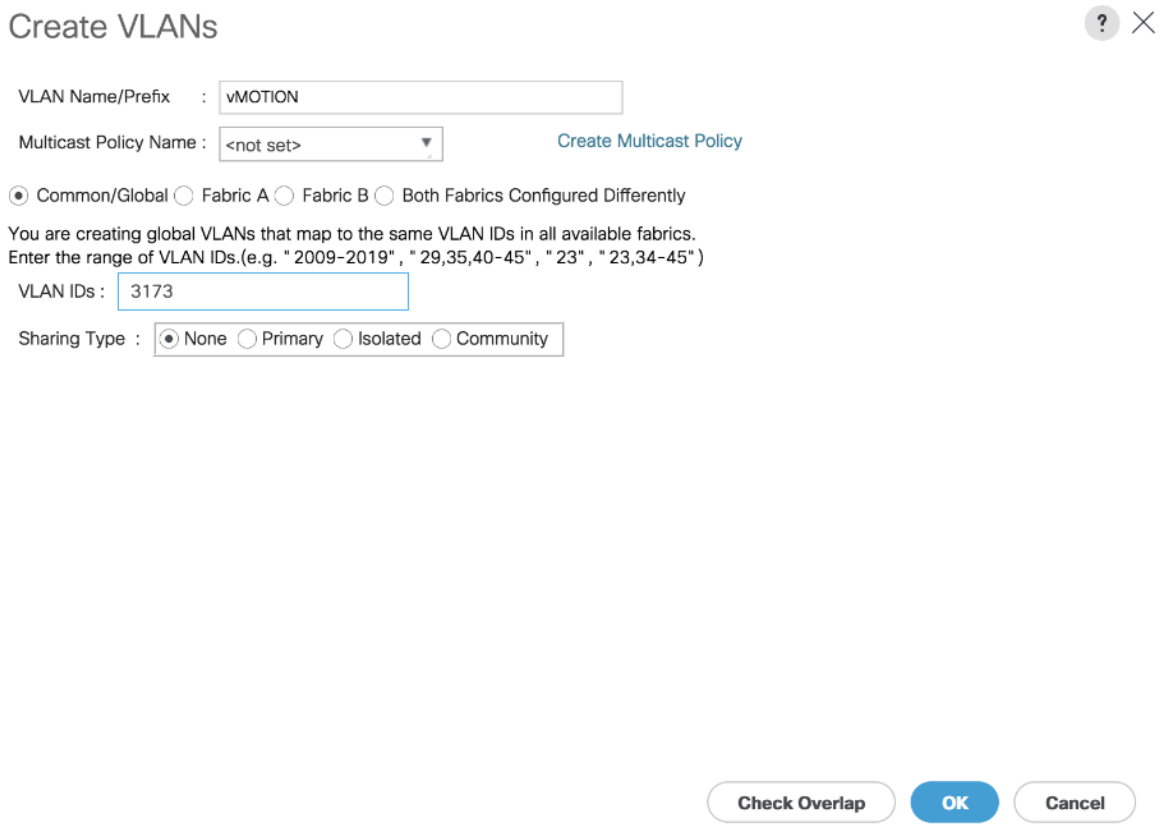

Create Host Firmware Package

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.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages and choose Create Host Firmware Package.
4. Enter Infra-FW-Pack as the name of the host firmware package.
5. Keep the Host Firmware Package as Simple.
6. Select the version 3.2(3d) for both the Blade and Rack Packages.
7. Click OK to create the host firmware package.
8. Click OK.

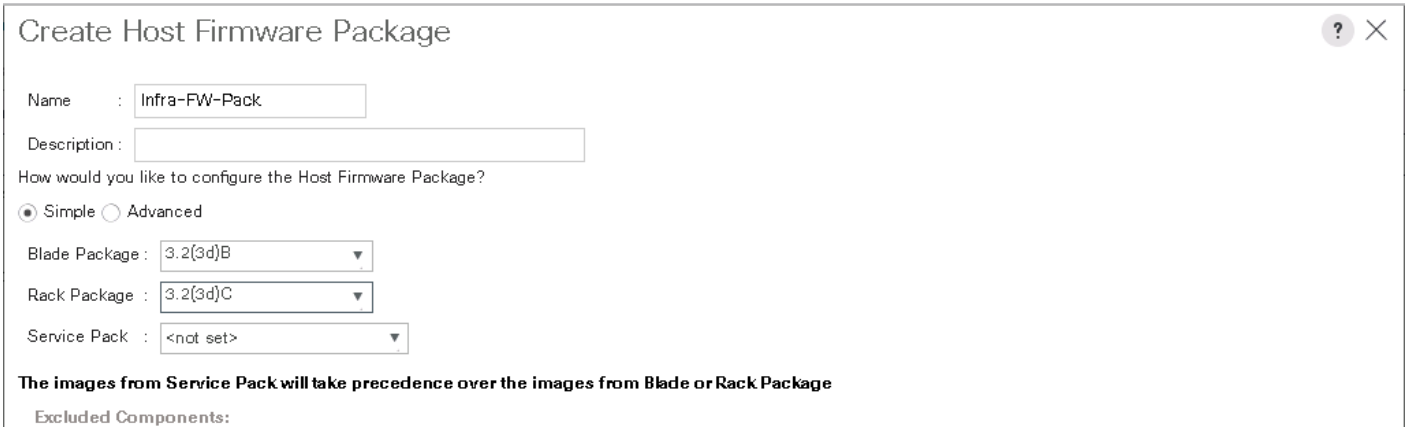

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames in the Cisco UCS fabric, complete the following 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.

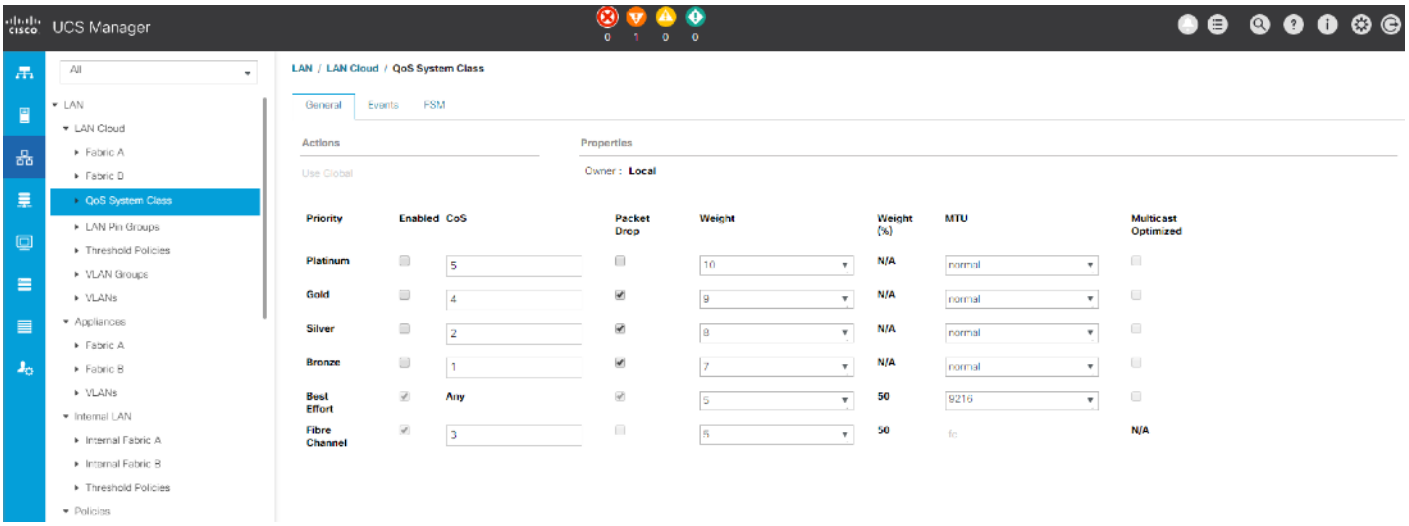

Create Local Disk Configuration Policy

When using an external storage system, a local disk configuration for the Cisco UCS environment is necessary because the servers in the environment will not contain a local disk.

This policy should not be applied to the servers that contain local disks.

To create a local disk configuration policy for no local disks, complete the following 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 and choose Create Local Disk Configuration Policy.
4. Enter SAN-Boot as the local disk configuration policy name.
5. Change the mode to No Local Storage.
6. Click OK to create the local disk configuration policy.

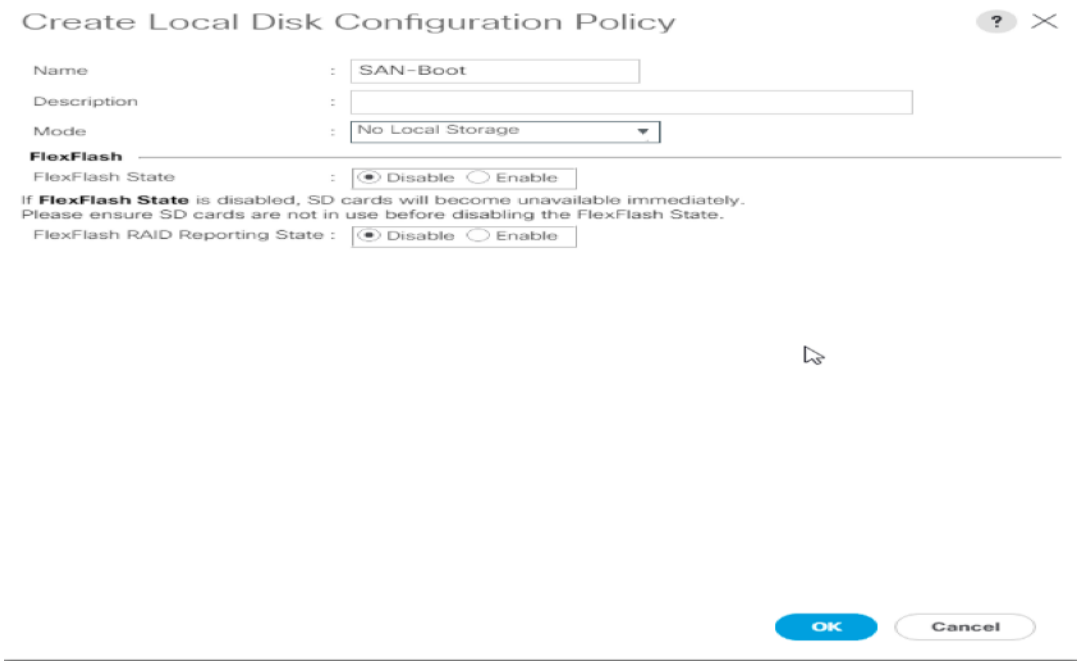

7. Click OK again.

Create Network Control Policy for Link Layer Discovery Protocol

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

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

2. Select Policies > root.
3. Right-click Network Control Policies and choose Create Network Control Policy.
4. Enter Enable-CDP as the policy name.
5. For CDP, select Enabled option.
6. Click OK to create the network control policy.

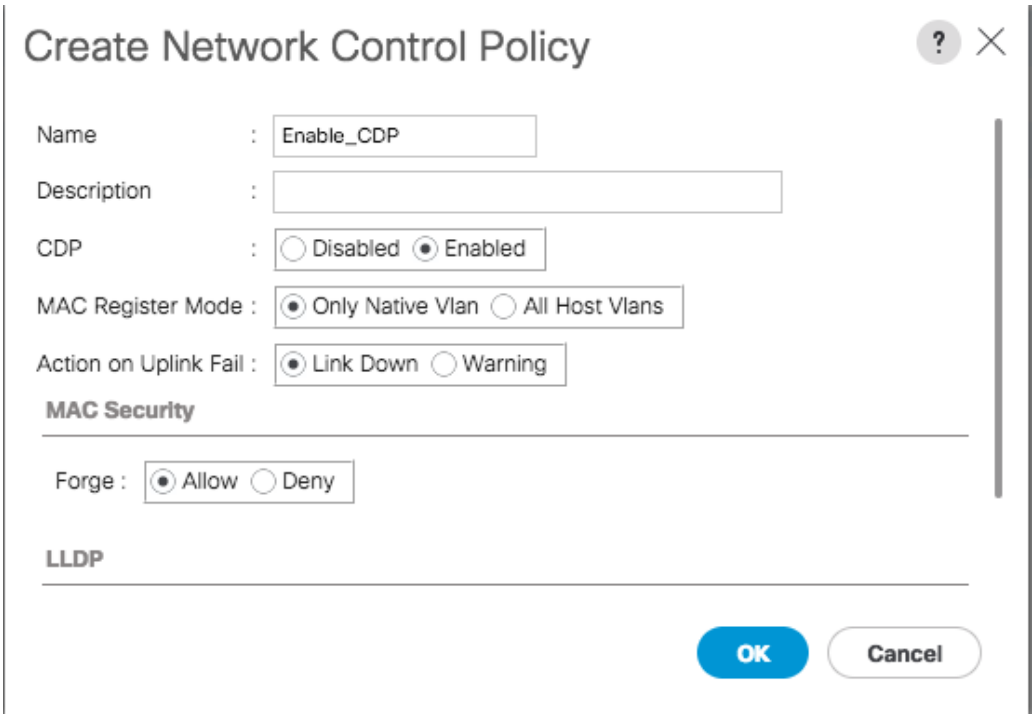

7. Click OK.

Create Power Control Policy

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

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

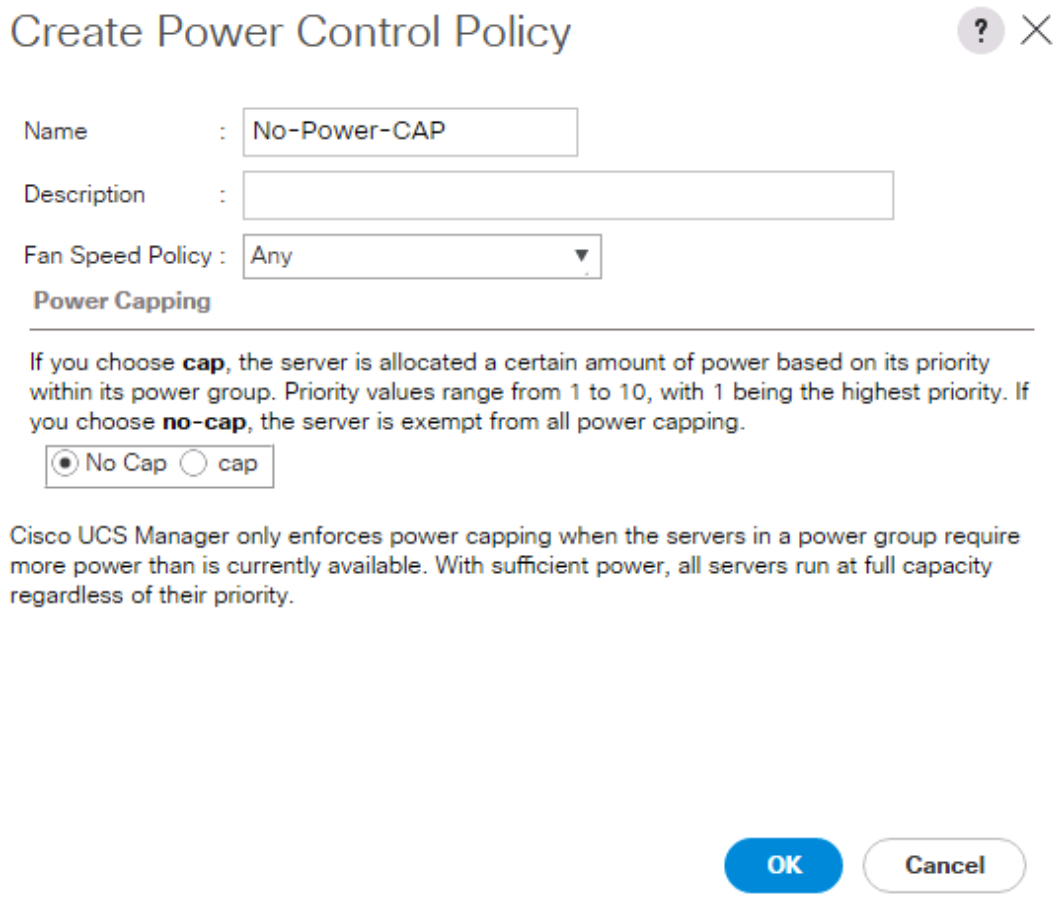

Create Server Pool Qualification Policy (Optional)

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

This example creates a policy for selecting a Cisco UCS B200-M5 server.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications and choose Create Server Pool Policy Qualification.
4. Enter UCSB-B200-M5 as the name for the policy.
5. Choose Create Server PID Qualifications.
6. Select UCSB-B200-M5 as the PID.

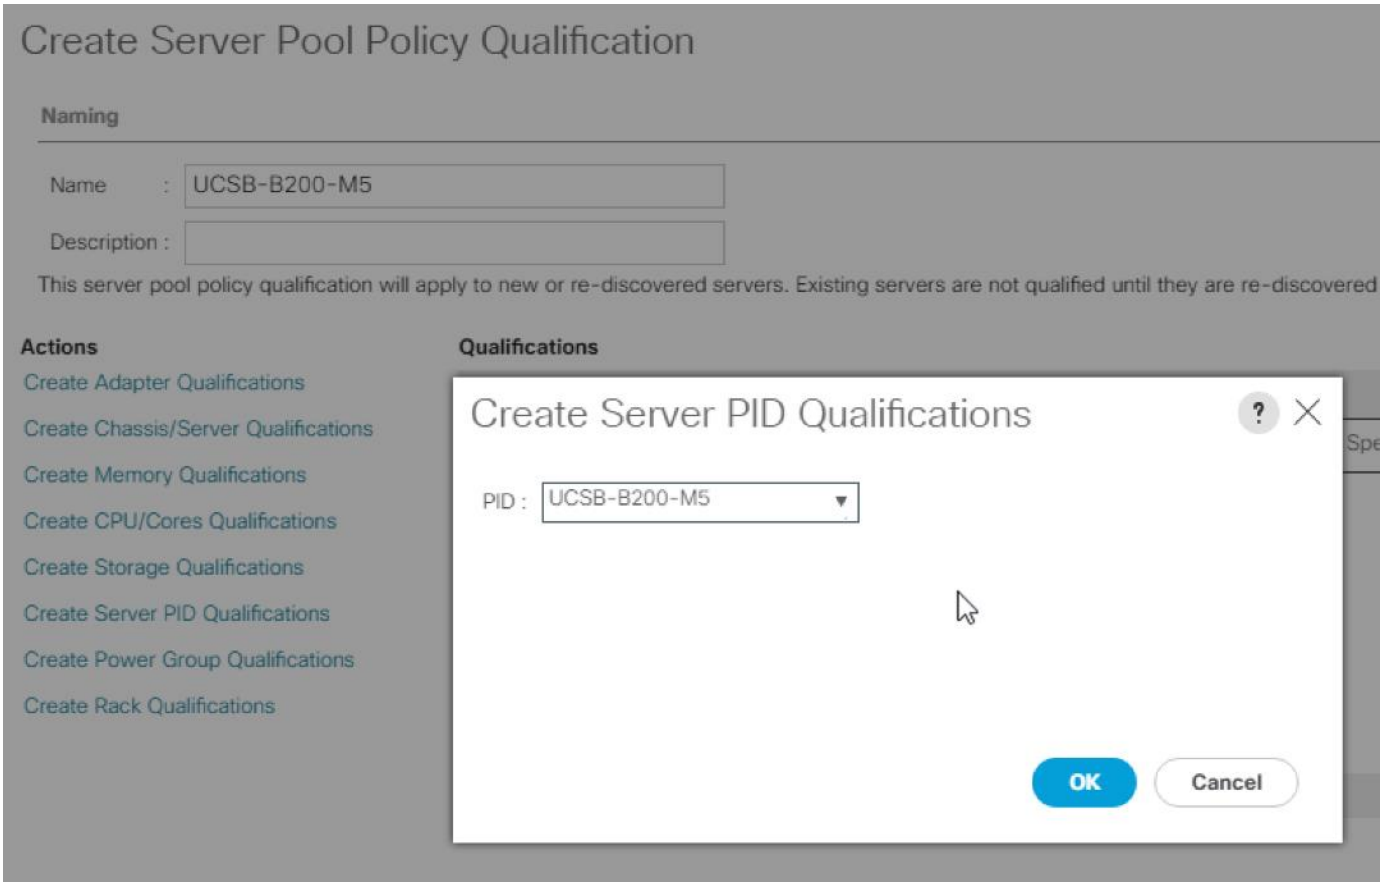

7. Click OK.
8. Click OK to create the server pool policy qualification.

Create Server BIOS Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click BIOS Policies and choose Create BIOS Policy.
4. Enter Infra-Host-BIOS as the BIOS policy name.
5. Select the newly created BIOS Policy.
6. Within the Advanced tab of the Policy:
7. Make the following selections recommended for virtualized workloads on Cisco UCS M5 platforms.
    a. Processor Settings

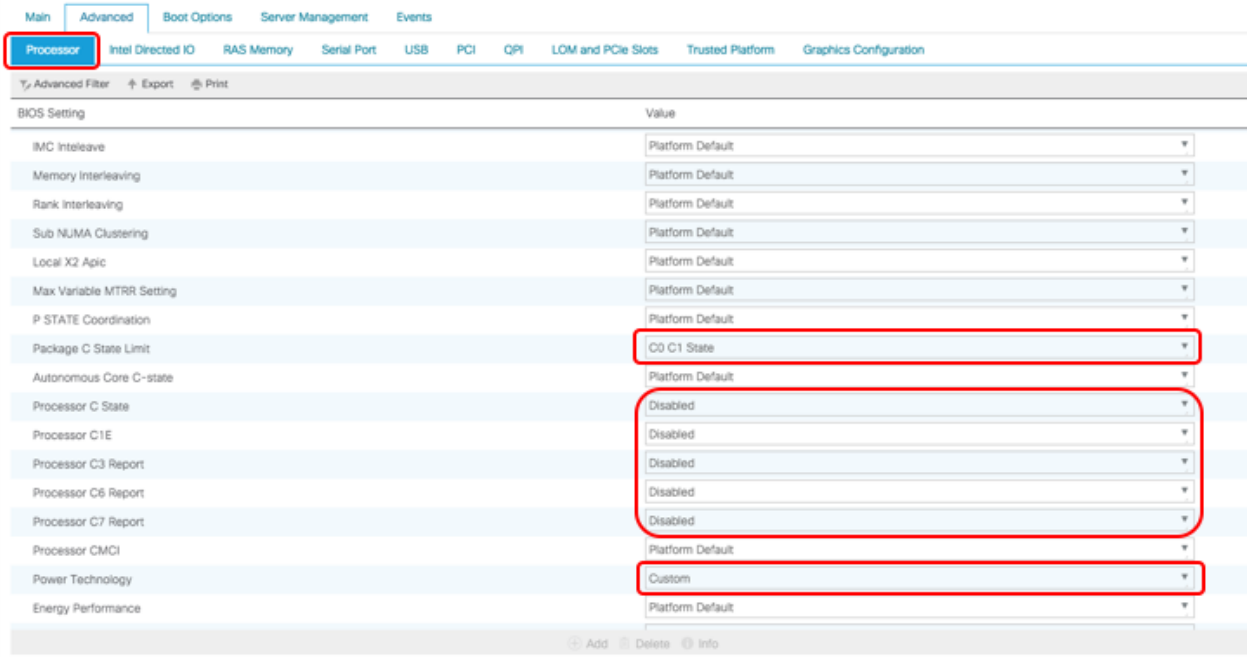

b. Intel Directed IO Settings

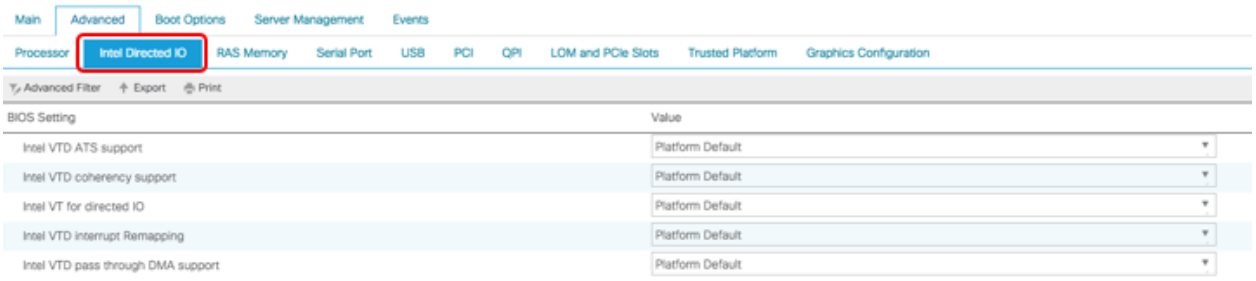

c. Memory settings

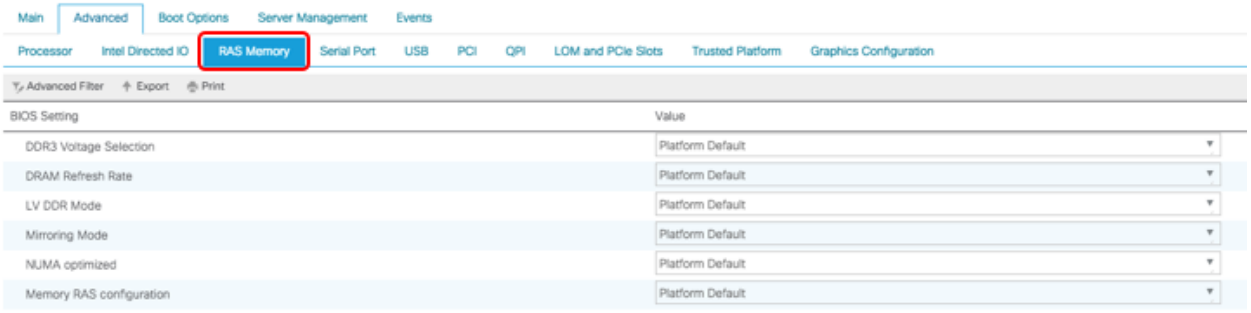

8. Click Save Changes to modify the BIOS policy.

9. Click OK.

Update Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root and then select Maintenance Policies > default.
3. Change the Reboot Policy to User Ack.
4. Check the box to enable On Next Boot
5. Click Save Changes.
6. Click OK to accept the change.

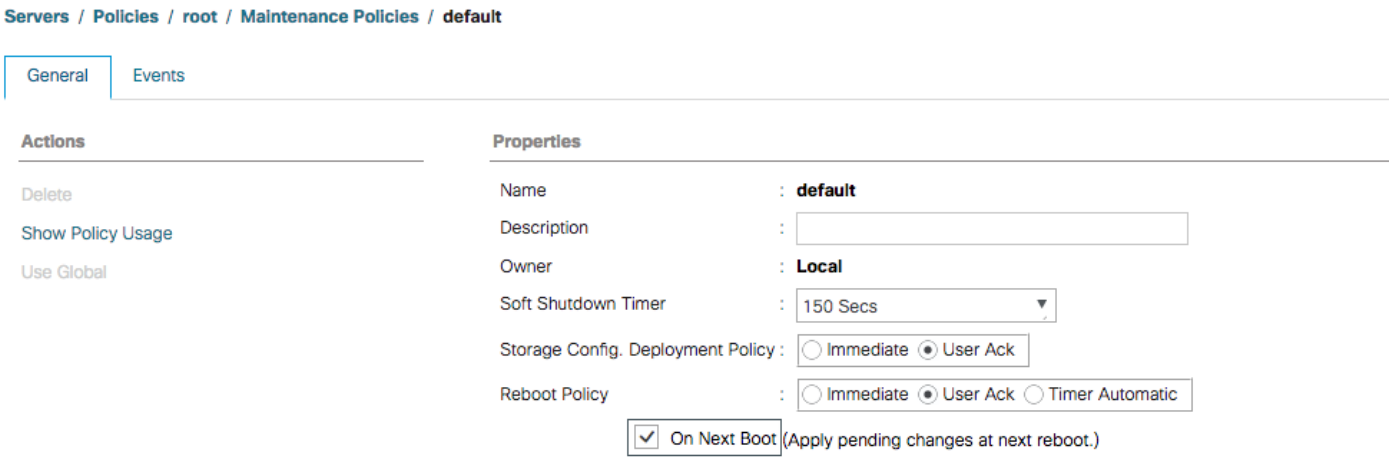

Create vNIC/vHBA Placement Policy

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

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

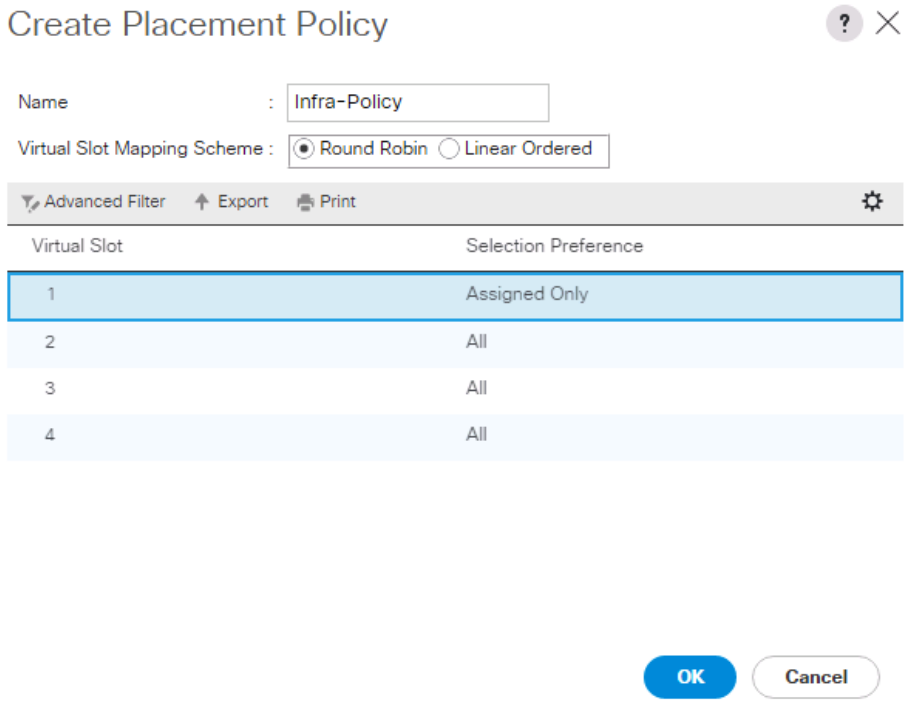

Create vNIC Templates

Eight different vNIC Templates are covered in Table 10 below. Not all VNICs need to be created in all deployments. The vNICs templates covered below are for iSCSI vNICs, infrastructure (management, vMotion etc.) vNICs, and data vNICs (VM traffic) for VMware VDS. Refer to Usage column in Table 10 below to see if a vNIC is needed for a particular ESXi host.

Table 10 NIC Templates and Associated VLANs

Name	Fabric ID	VLANs	Native VLAN	MAC Pool	Usage
vNIC_Mgmt_A	A	IB-Mgmt, Native-2	Native-2	MAC-Pool-A	All ESXi Hosts;
vNIC_Mgmt_B	B	IB-Mgmt, Native-2	Native-2	MAC-Pool-B	All ESXi Hosts;
vNIC_vMotion_A	A	vMotion, Native-2	Native-2	MAC-Pool-A	All ESXi Hosts;
vNIC_vMotion_B	A	vMotion, Native-2	Native-2	MAC-Pool-B	All ESXi Hosts;
vNIC_VM_A	A	VM Network	Native-2	MAC-Pool-A	All ESXi Hosts;
vNIC_VM_B	A	VM Network	Native-2	MAC-Pool-B	All ESXi Hosts;

Name	Fabric ID	VLANs	Native VLAN	MAC Pool	Usage
vNIC_iSCSI_A	A	Infra-iSCSI-A	Infra-iSCSI-A	MAC-Pool-A	iSCSI hosts only
vNIC_iSCSI_B	B	Infra-iSCSI-B	Infra-iSCSI-B	MAC-Pool-B	iSCSI hosts only

Create Management vNICs

For the vNIC_Mgmt_A Template, complete the following 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. Optional: select the Enable Failover checkbox.

Selecting Failover can improve link failover time by handling it at the hardware level, and can guard against any potential for NIC failure not being detected by the virtual switch.

8. Select Primary Template for the Redundancy Type.
9. 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.

10. Under Target, make sure that the VM checkbox is not selected.
11. Select Updating Template as the Template Type.
12. Under VLANs, select the checkboxes for IB-Mgmt and Native-VLAN VLANs.

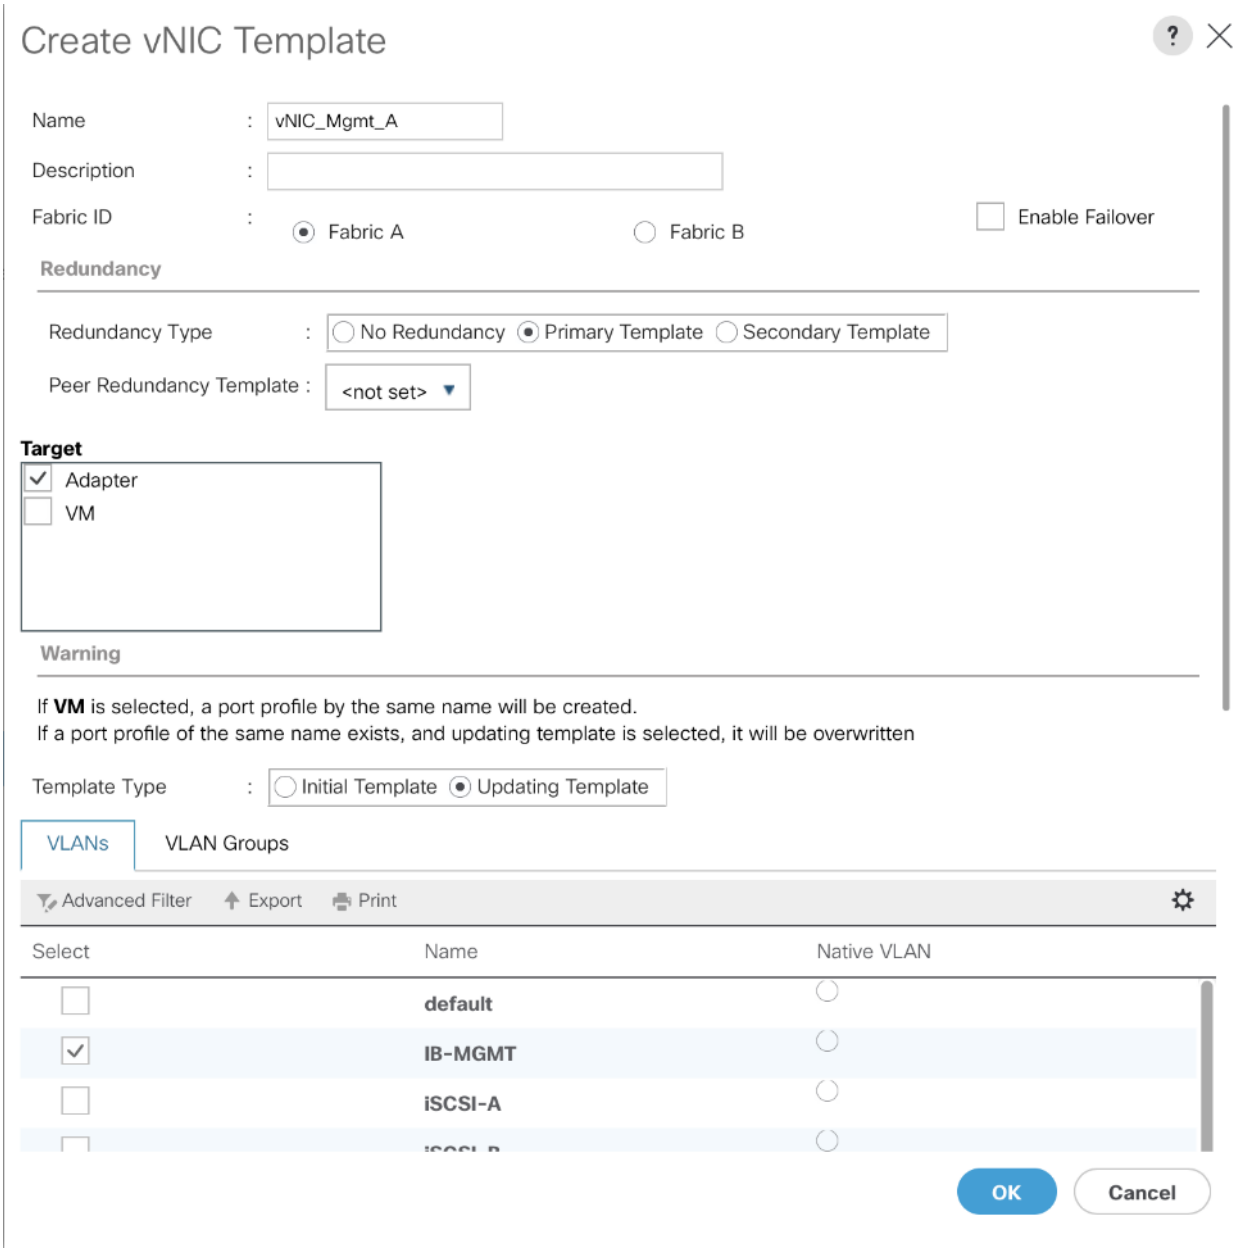

13. Set Native-VLAN as the native VLAN.

14. Leave vNIC Name selected for the CDN Source.

15. Leave 1500 for the MTU.

16. In the MAC Pool list, select MAC_Pool_A.

17. In the Network Control Policy list, select Enable_CDP.

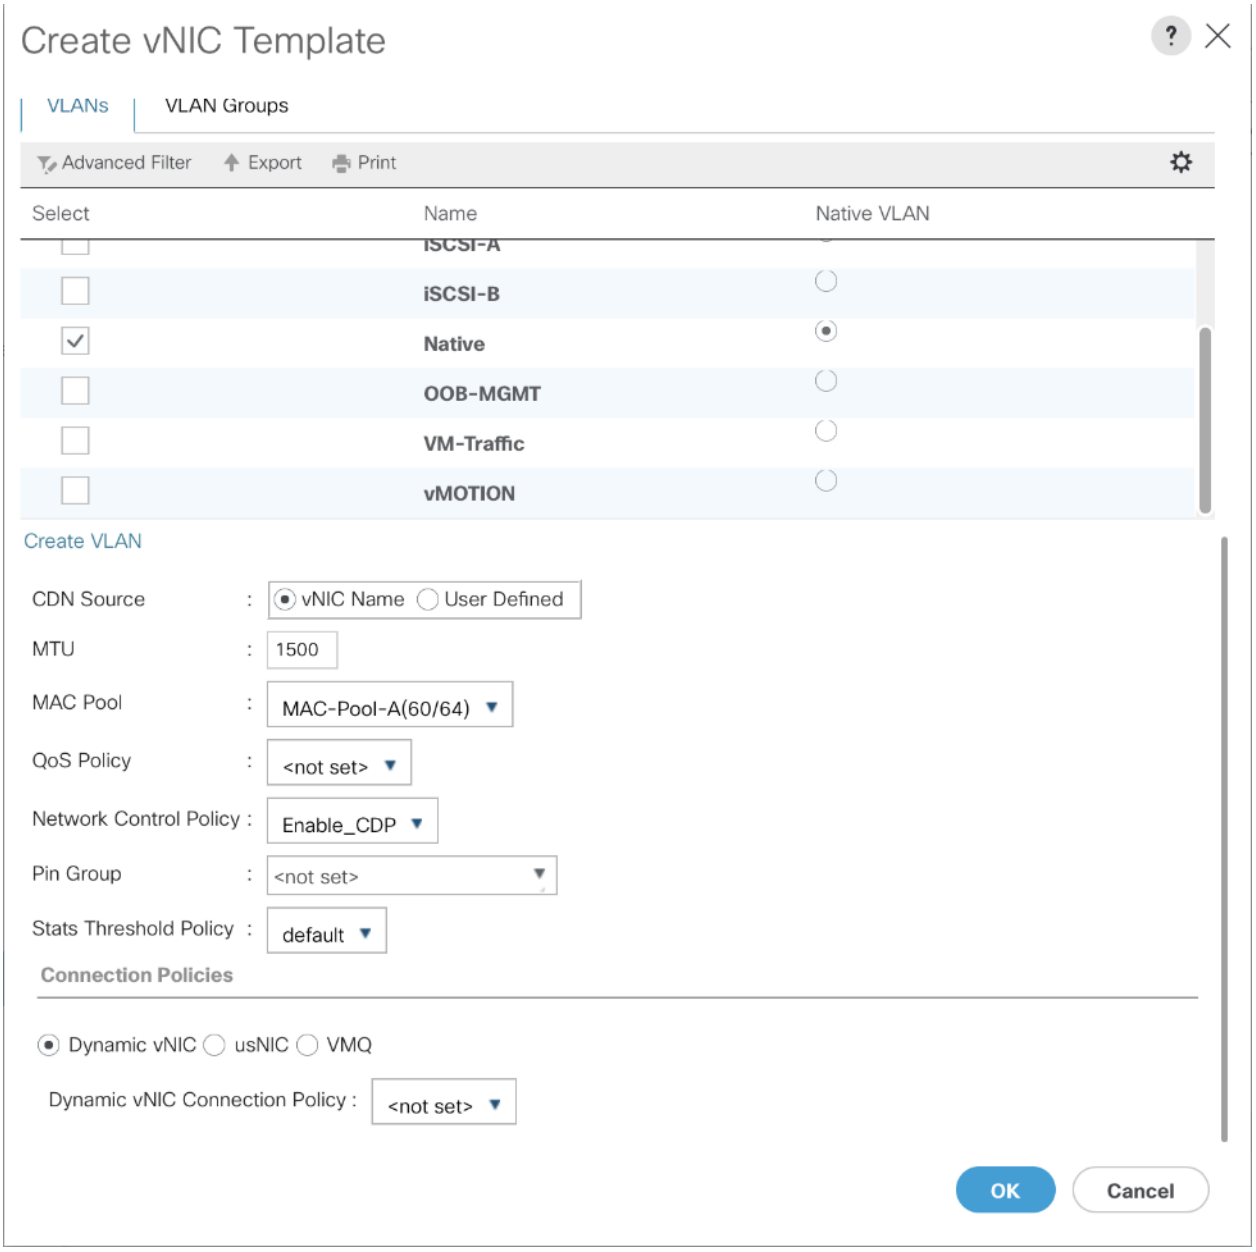

18. Click OK to create the vNIC template.

19. Click OK.

For the vNIC_Mgmt_B Template, complete the following 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, Failover specification, 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.

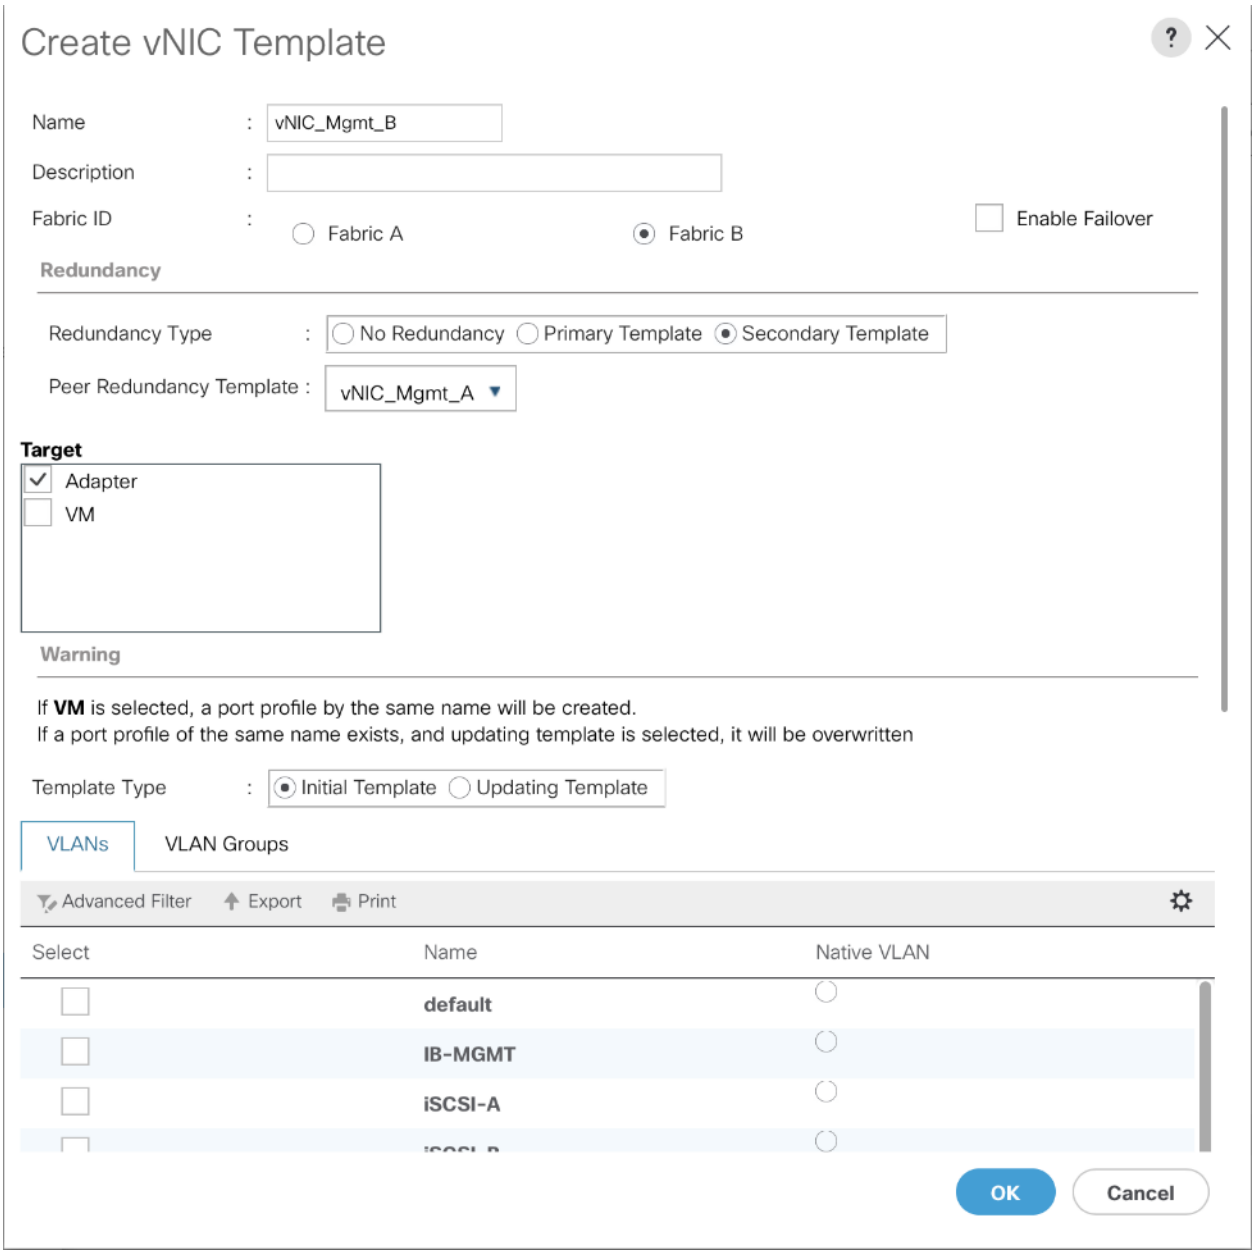

10. In the MAC Pool list, select MAC_Pool_B.

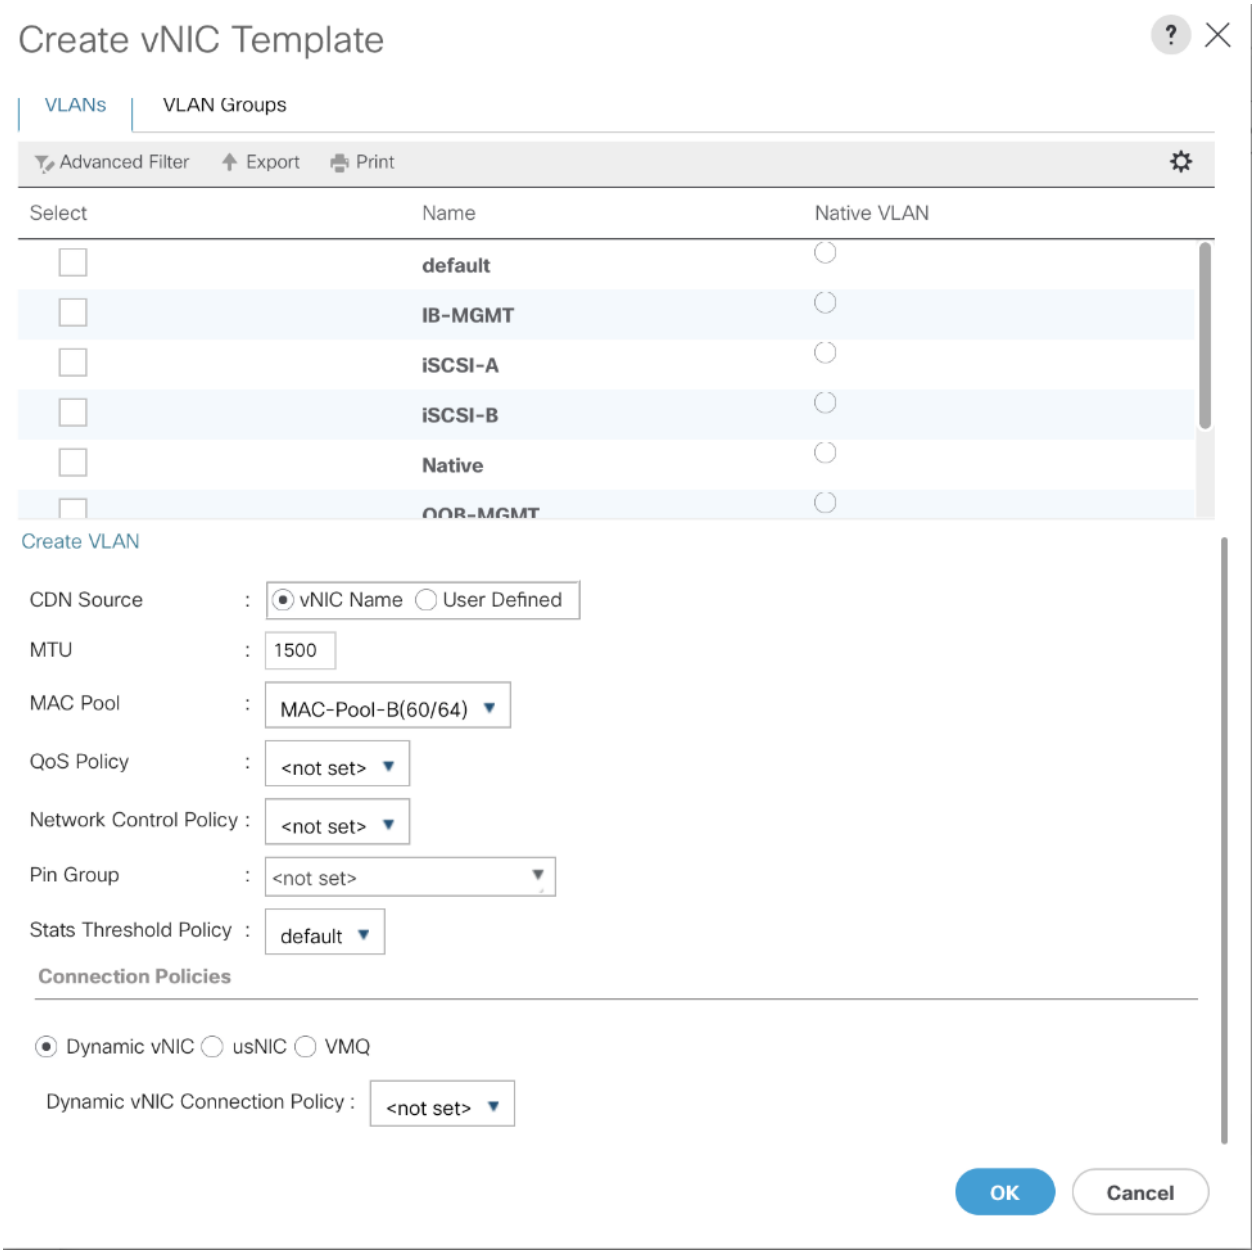

11. Click OK to create the vNIC template.

12. Click OK.

Create vMotion vNICs

For the vNIC_vMotion_A Template, complete the following 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_vMotion_A as the vNIC template name.
6. Keep Fabric A selected.
7. Optional: select the Enable Failover checkbox.
8. Select Primary Template for the Redundancy Type.
9. Leave Peer Redundancy Template as <not set>
10. Under Target, make sure that the VM checkbox is not selected.
11. Select Updating Template as the Template Type.

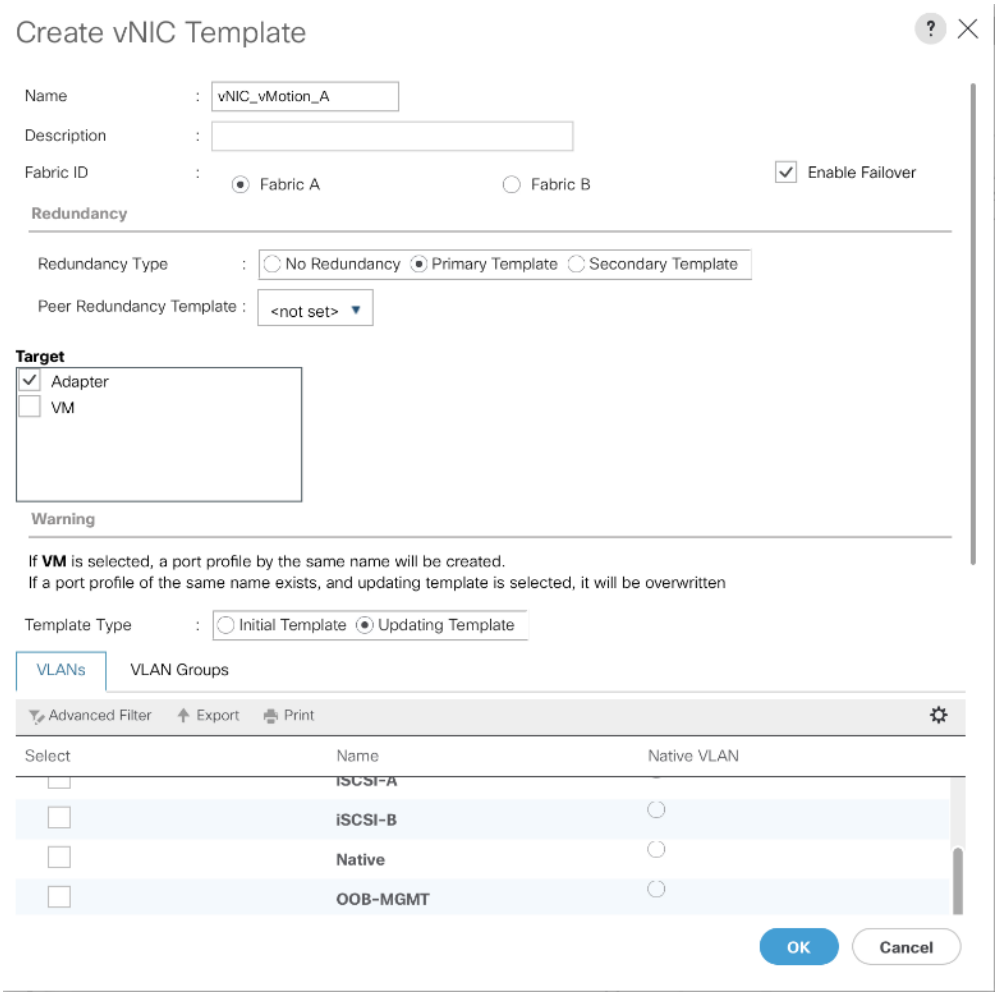

12. Under VLANs, select the checkboxes vMotion as the only VLAN.
13. Set vMotion as the native VLAN.

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.

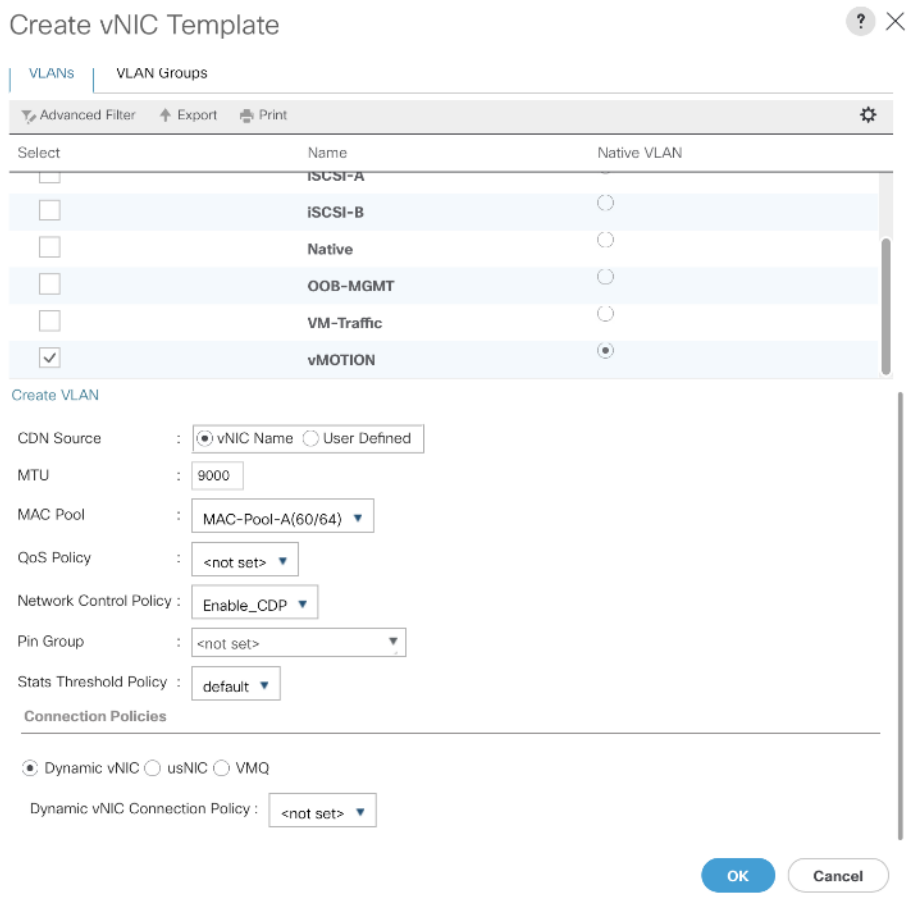

17. Click OK to create the vNIC template.

18. Click OK.

For the vNIC_vMotion_B Template, complete the following 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_vMotion_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_vMotion_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.

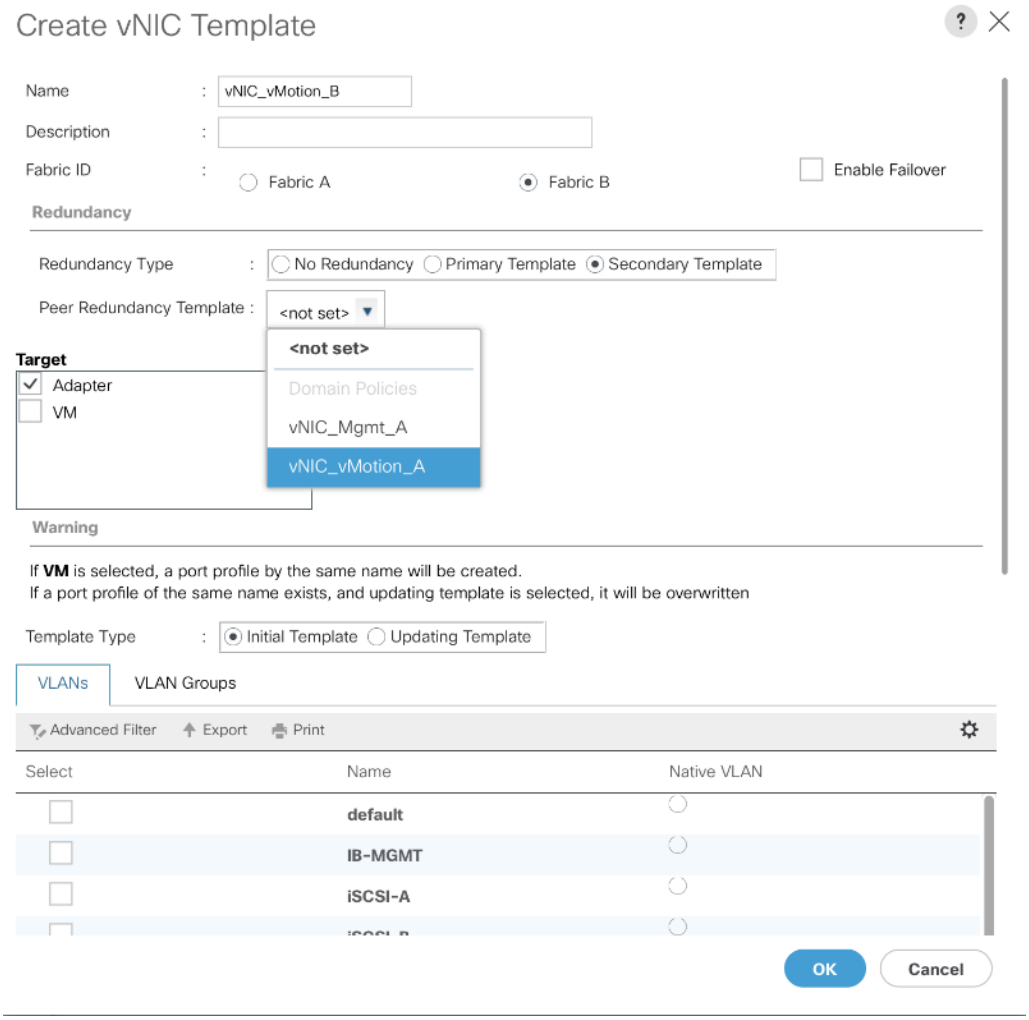

10. the MAC Pool list, select MAC_Pool_B.

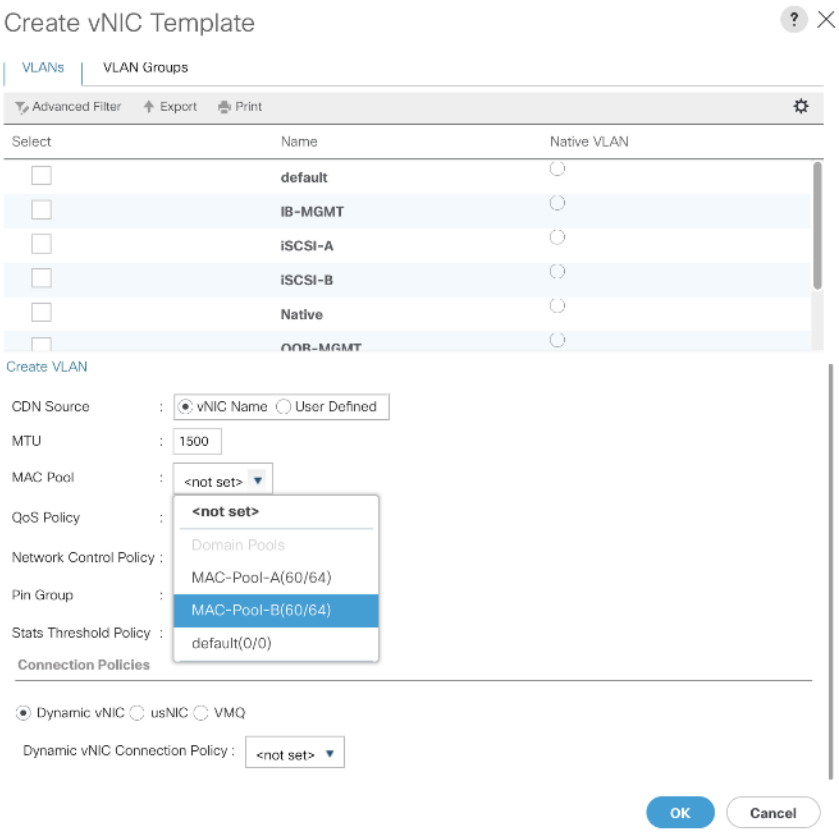

11. Click OK to create the vNIC template.
12. Click OK.

Create Application vNICs

To create the vNIC_VM_A Template, complete the following 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_VM_A as the vNIC template name.
6. Keep Fabric A selected.
7. Optional: select the Enable Failover checkbox.
8. Select Primary Template for the Redundancy Type.
9. Leave Peer Redundancy Template as <not set>.

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

11. Select Updating Template as the Template Type.

12. Set default as the native VLAN.

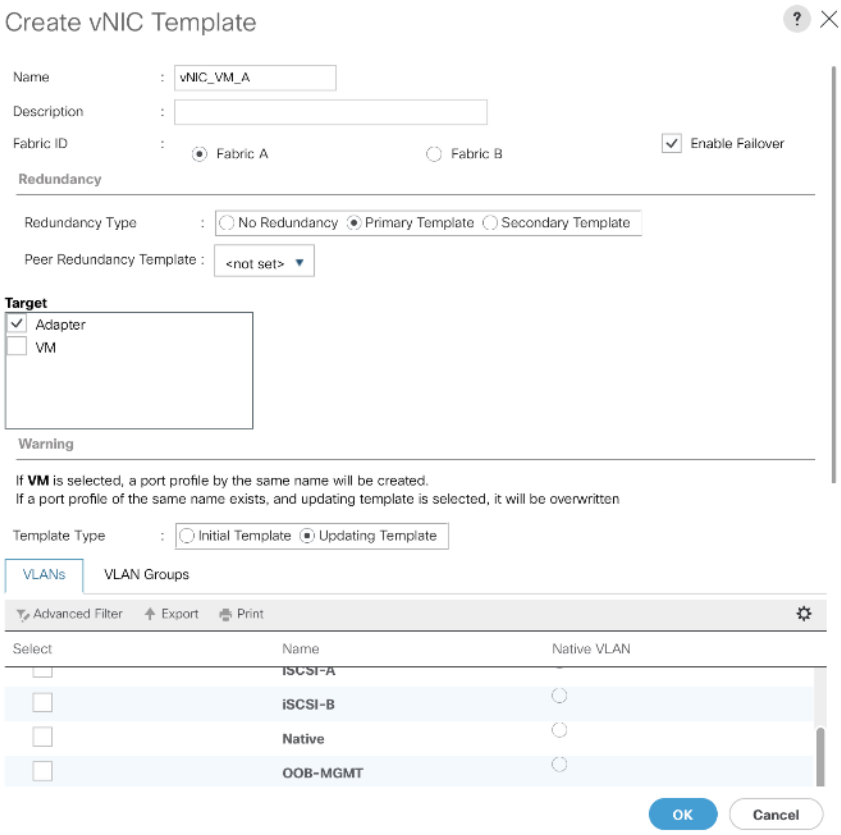

13. Under VLANs, select the checkboxes for any application or production VLANs that should be delivered to the ESXi hosts.

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.

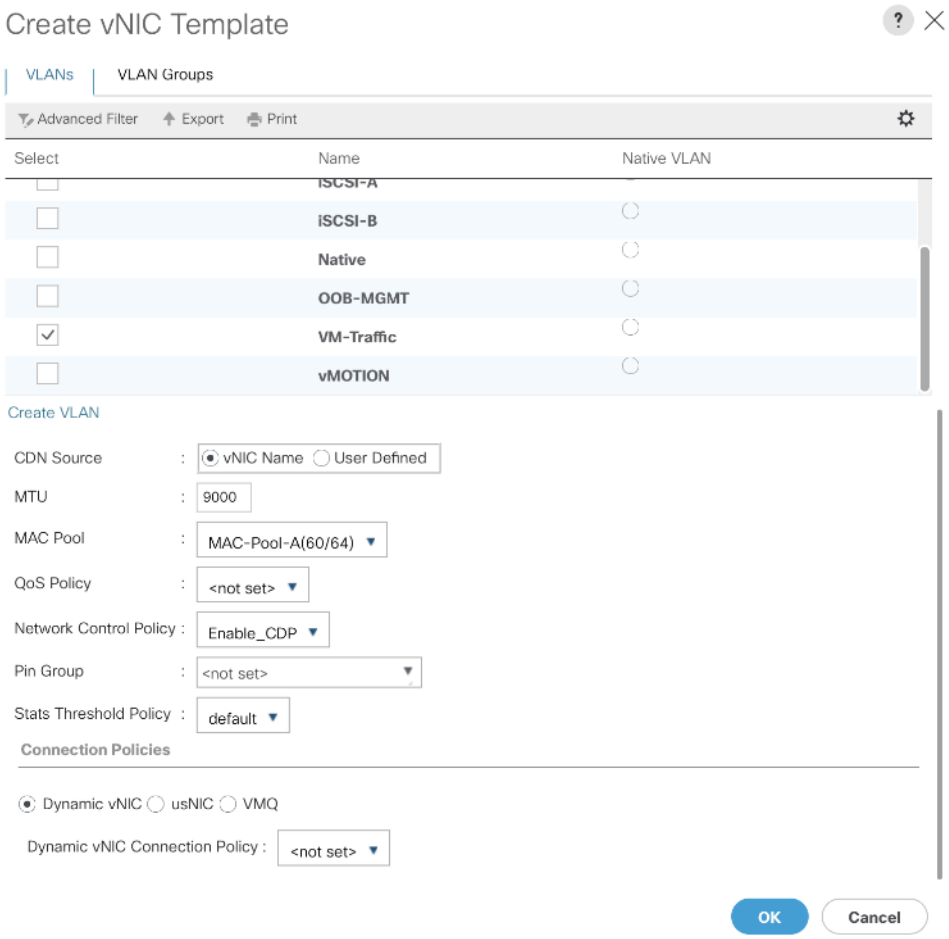

17. Click OK to create the vNIC template.

18. Click OK.

To create the vNIC_VM_B Templates, complete the following 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_VM_B as the vNIC template name.
6. Select Fabric B.
7. Select Secondary Template for Redundancy Type.
8. From the Peer Redundancy Template drop-down list, select vNIC_VM_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.

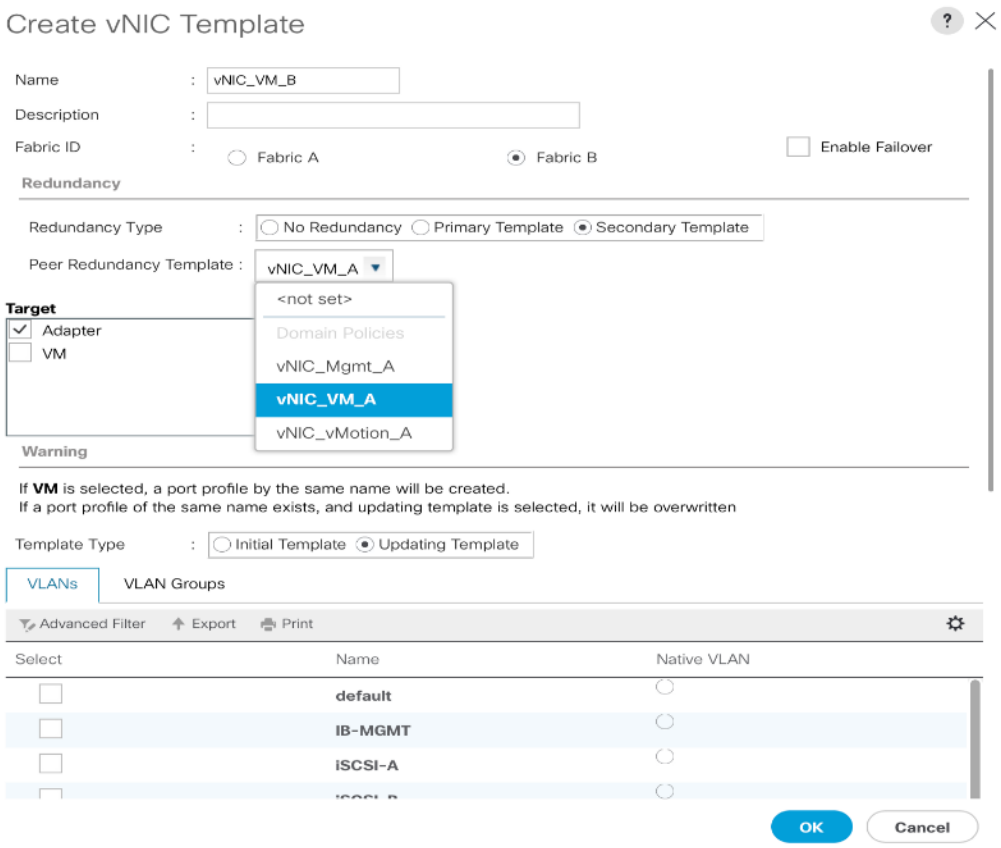

10. In the MAC Pool list, select MAC_Pool_B.

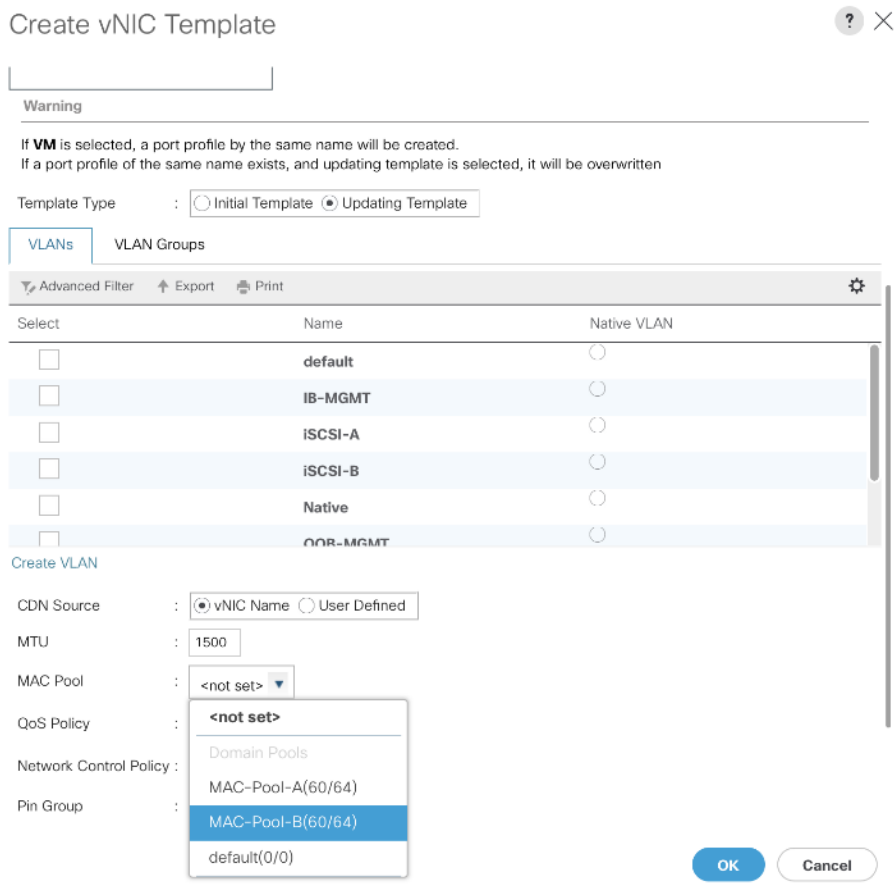

11. Click OK to create the vNIC template.
12. Click OK.

Create iSCSI vNICs

The configuration steps to create iSCSI vNICs can be skipped if the UCS environment does not need to access storage environment using iSCSI.

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_iSCSI_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.

8. Keep the No Redundancy options selected for the Redundancy Type.
9. Under Target, make sure that the Adapter checkbox is selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select iSCSI-A-VLAN as the only VLAN and set it as the Native VLAN.

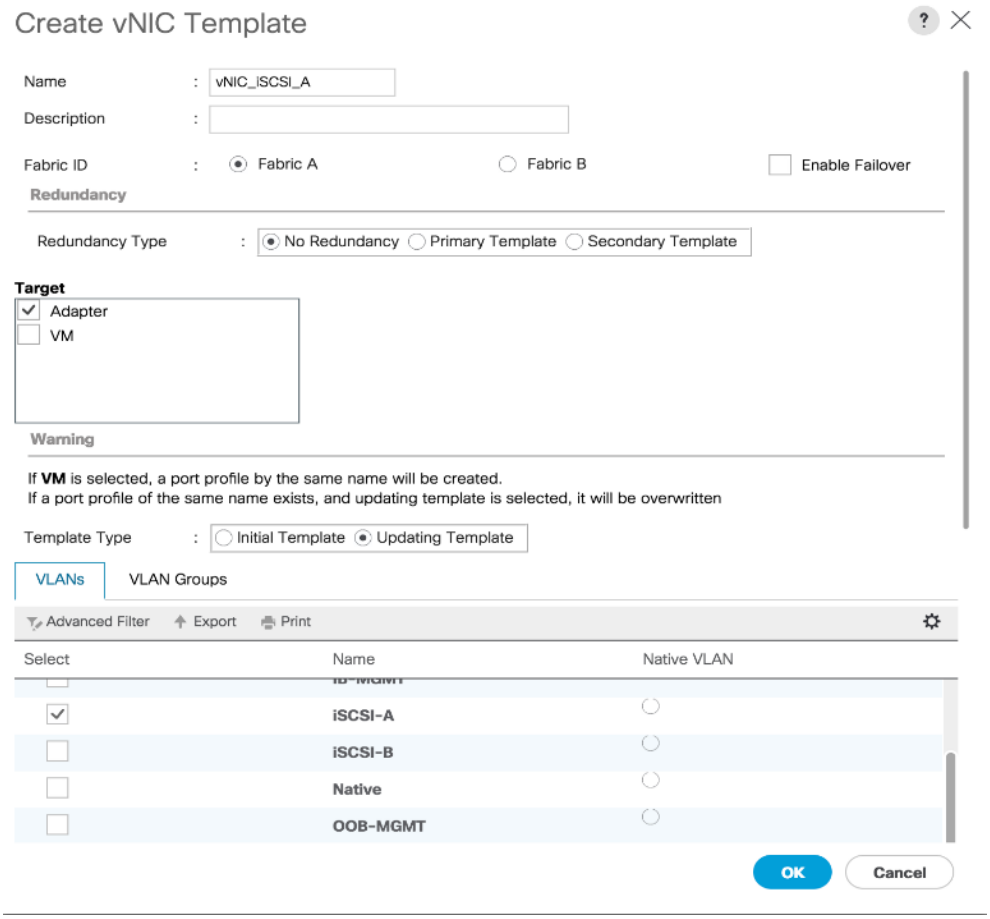

12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_A.
14. In the Network Control Policy list, select Enable_CDP.

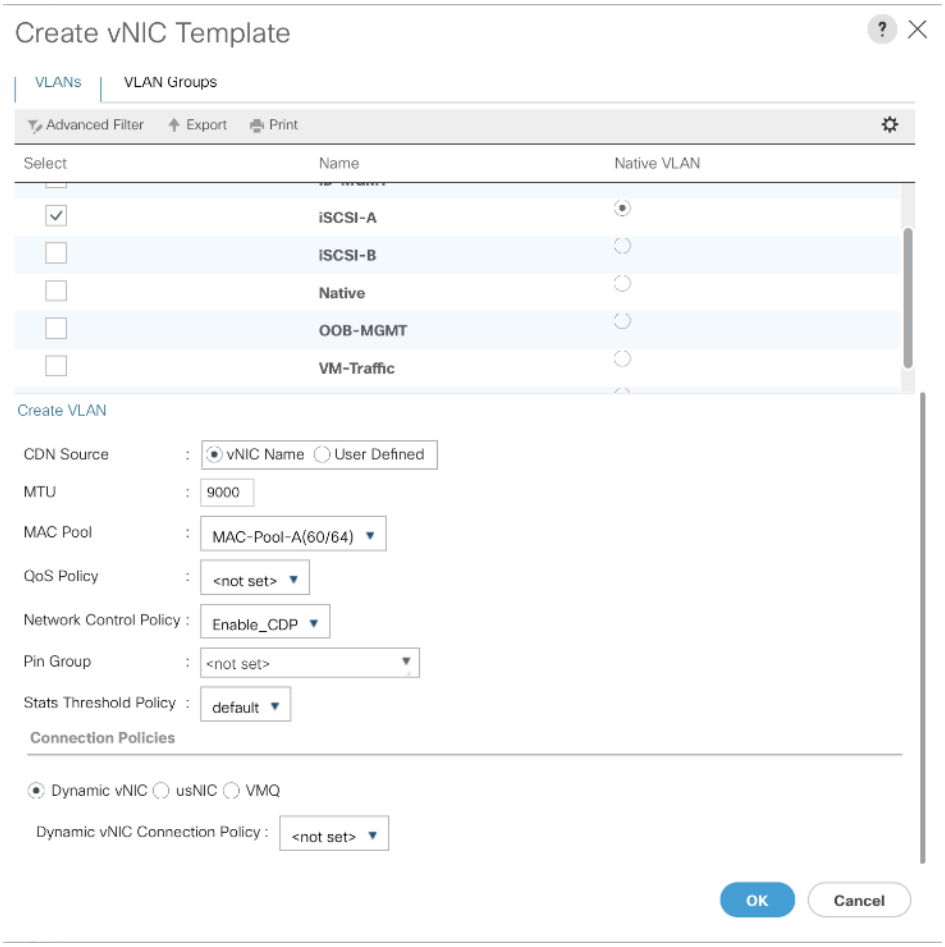

15. Click OK to create the vNIC template.

16. Click OK.

To create the vNIC_iSCSI_B Template, complete the following 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_iSCSI_B as the vNIC template name.
6. Keep Fabric B selected.
7. Do not select the Enable Failover checkbox.
8. Keep the No Redundancy options selected for the Redundancy Type.

9. Under Target, make sure that the Adapter checkbox is selected.

10. Select Updating Template as the Template Type.

11. Under VLANs, select iSCSI-B-VLAN as the only VLAN and set it as the Native VLAN.

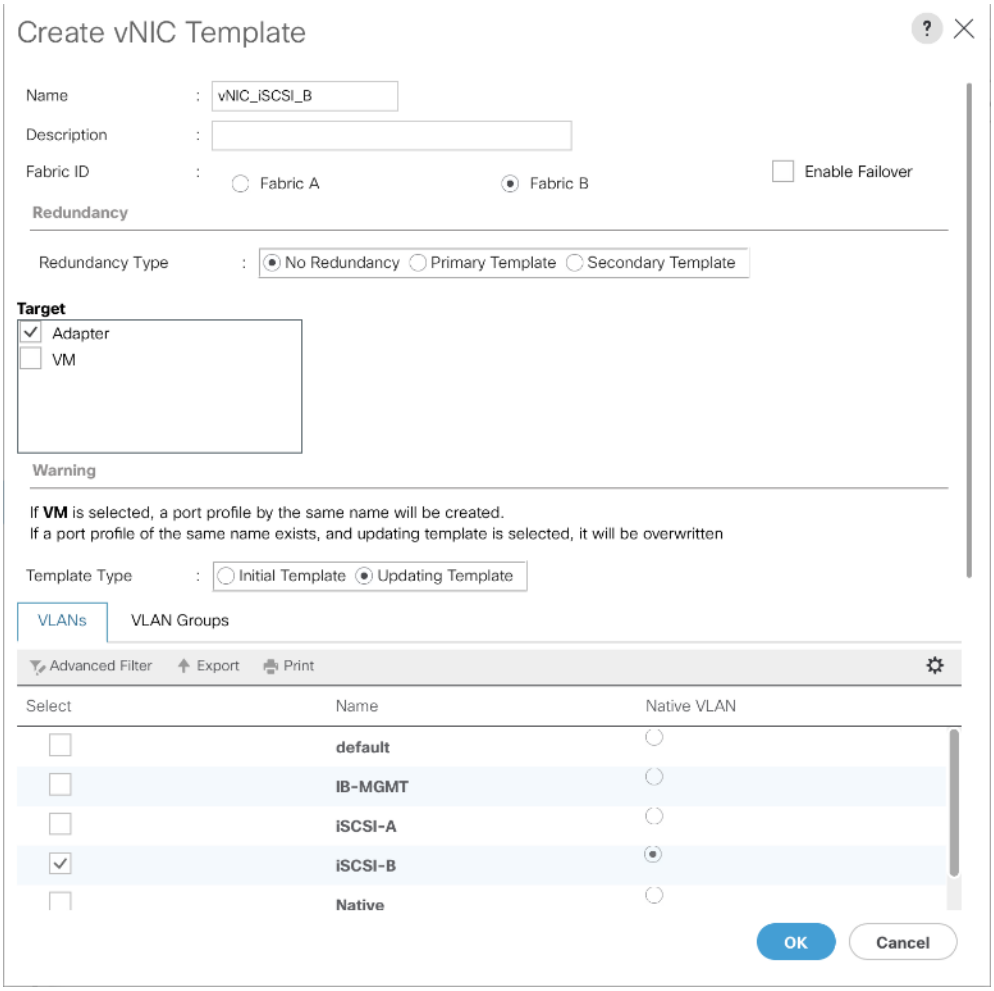

12. For MTU, enter 9000.

13. In the MAC Pool list, select MAC_Pool_A.

14. In the Network Control Policy list, select Enable_CDP.

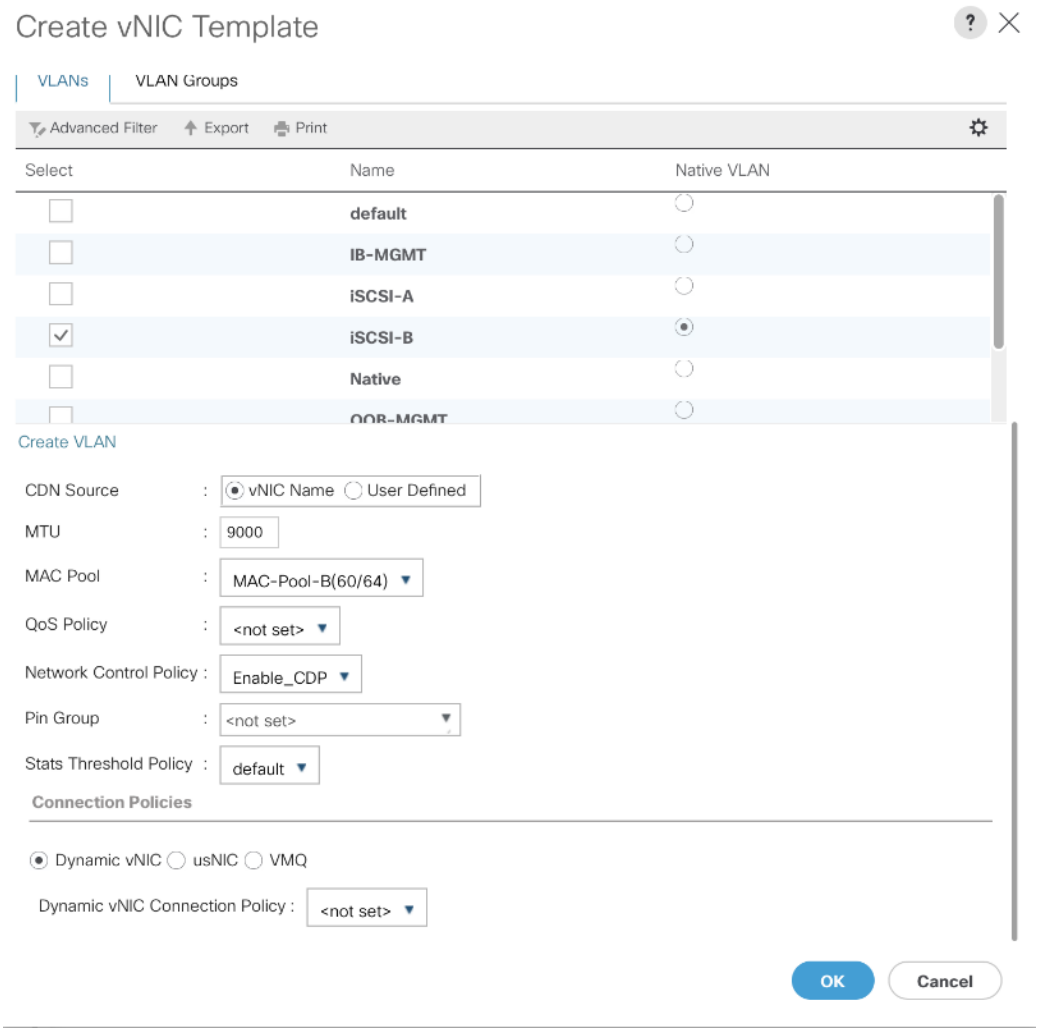

15. Click OK to create the vNIC template.
16. Click OK.

Create LAN Connectivity Policy

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.

Use Infra-LAN-Pol as the name if hosts boot from FC only.

5. Enter iSCSI-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 00-Mgmt-A.
10. In the Adapter Policy list, select VMWare.
11. Click OK to add this vNIC to the policy.

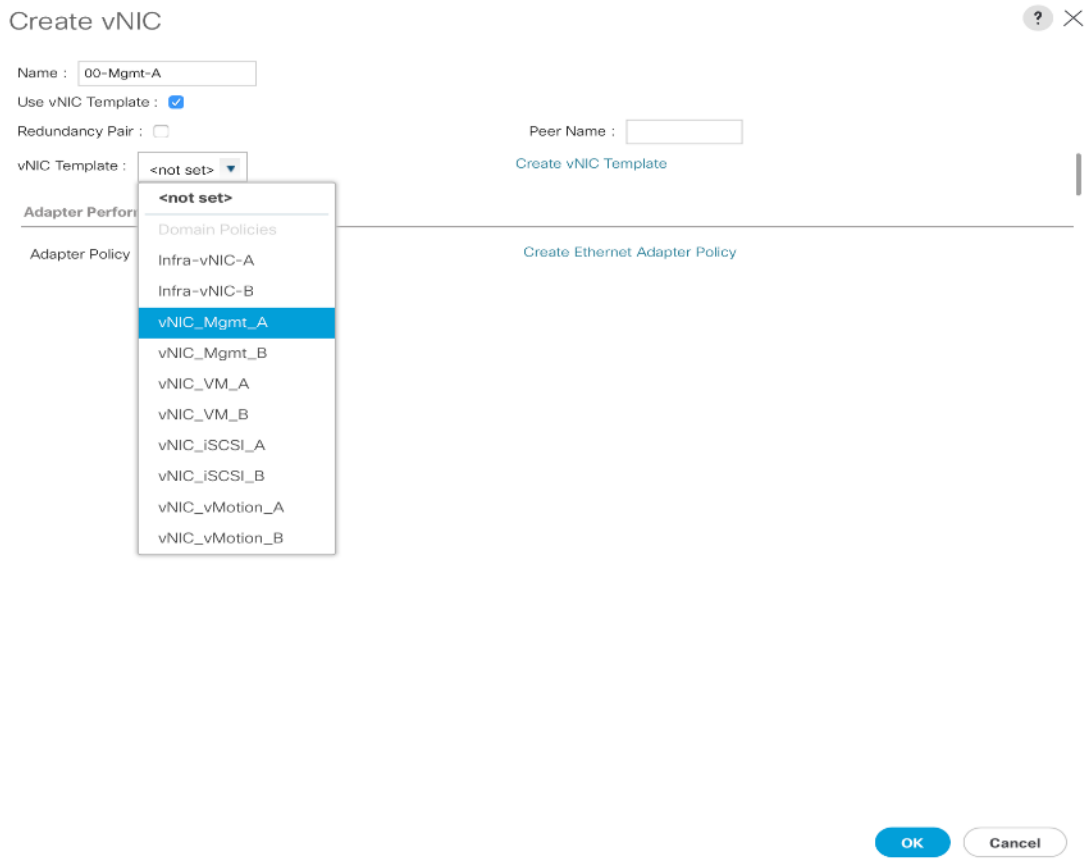

12. Click the upper Add button to add another vNIC to the policy.
13. In the Create vNIC box, enter 01-Mgmt-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select 01-Mgmt-B.

16. In the Adapter Policy list, select VMWare.

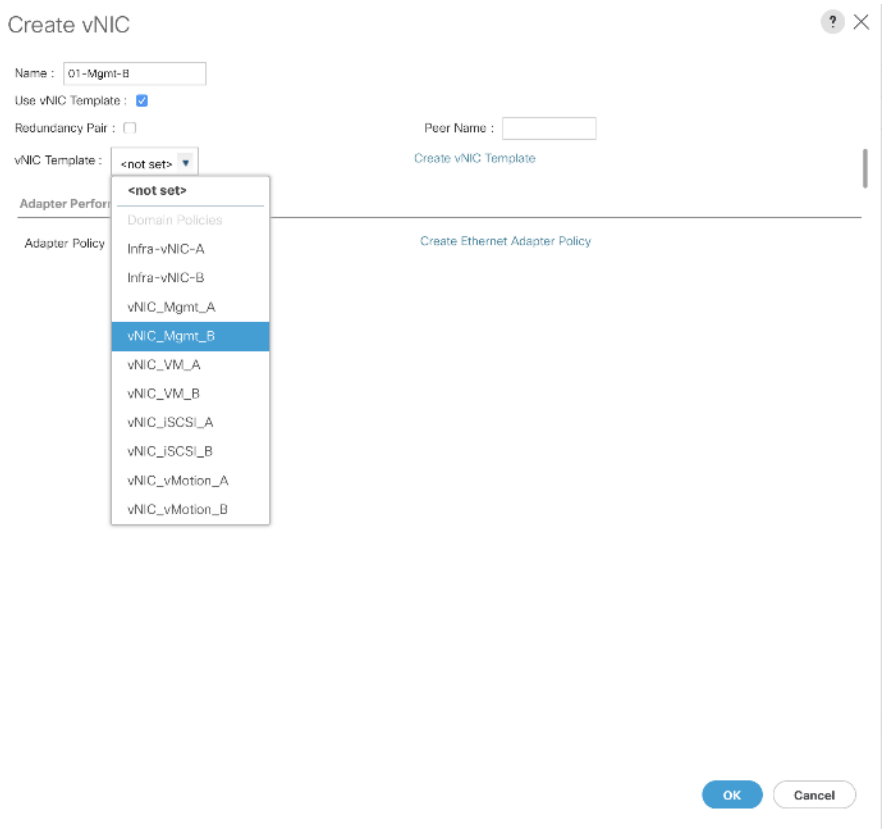

17. Click OK to add the vNIC to the policy.

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

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

20. Select the Use vNIC Template checkbox.

21. In the vNIC Template list, select vNIC_vMotion_A.

22. In the Adapter Policy list, select VMWare.

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

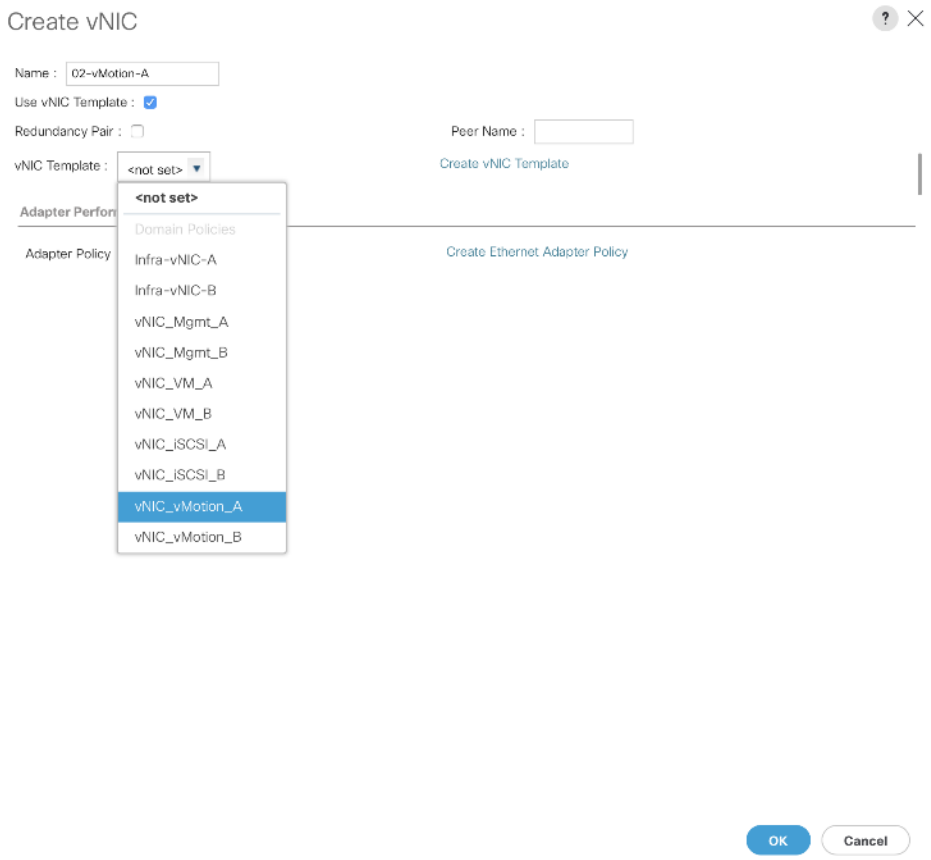

24. Click the upper Add button to add a vNIC to the policy.

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

26. Select the Use vNIC Template checkbox.

27. In the vNIC Template list, select vNIC_vMotion_B.

28. In the Adapter Policy list, select VMWare.

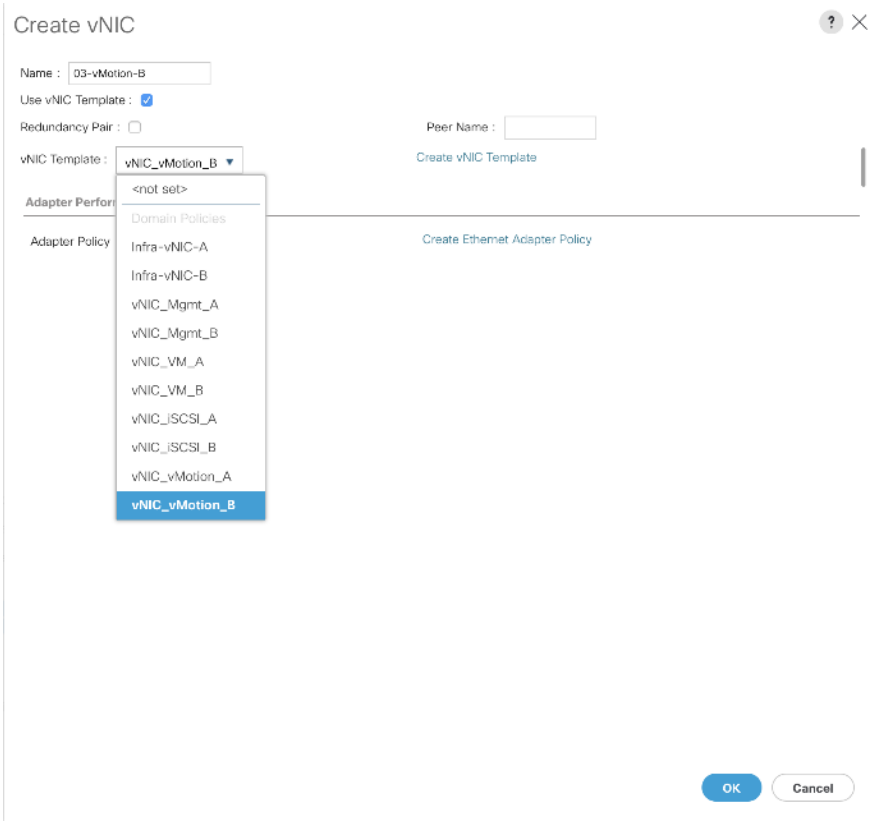

29. Click OK to add this vNIC to the policy.
30. Click the upper Add button to add a vNIC.
31. In the Create vNIC dialog box, enter 04-VM-A as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select vNIC_VM_A.
34. In the Adapter Policy list, select VMWare.
35. Click OK to add this vNIC to the policy.

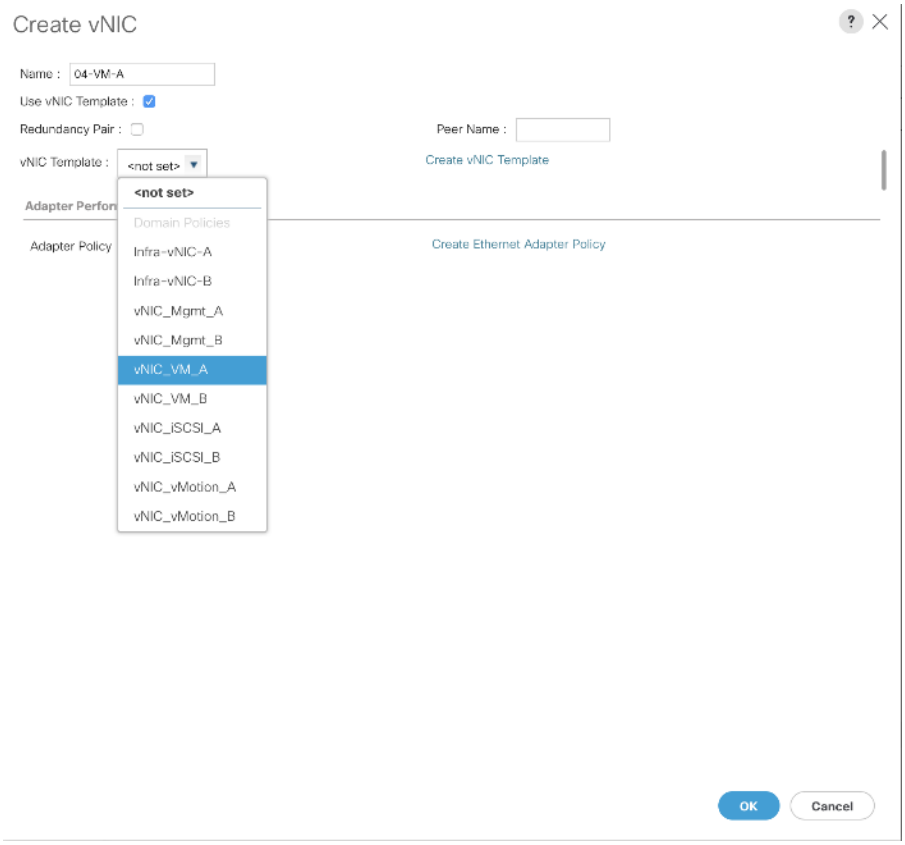

36. Click the upper Add button to add a vNIC to the policy.

37. In the Create vNIC dialog box, enter 05-VM-B as the name of the vNIC.

38. Select the Use vNIC Template checkbox.

39. In the vNIC Template list, select vNIC_VM_B.

40. In the Adapter Policy list, select VMWare.

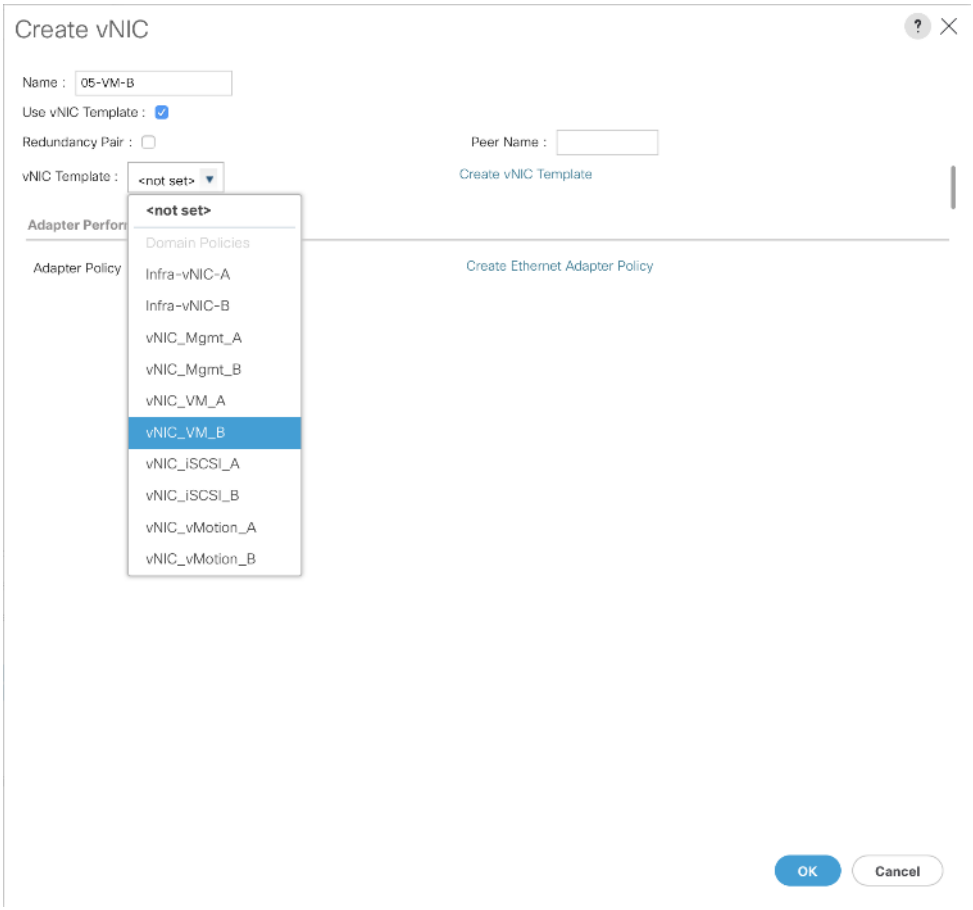

41. Click OK to add this vNIC to the policy.

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

The iSCSI vNICs can be skipped if hosts need FC storage access only.

---

43. In the Create vNIC dialog box, enter 06-iSCSI-A as the name of the vNIC.

44. Select the Use vNIC Template checkbox.

45. In the vNIC Template list, select iSCSI-Template-A.

46. In the Adapter Policy list, select VMWare.

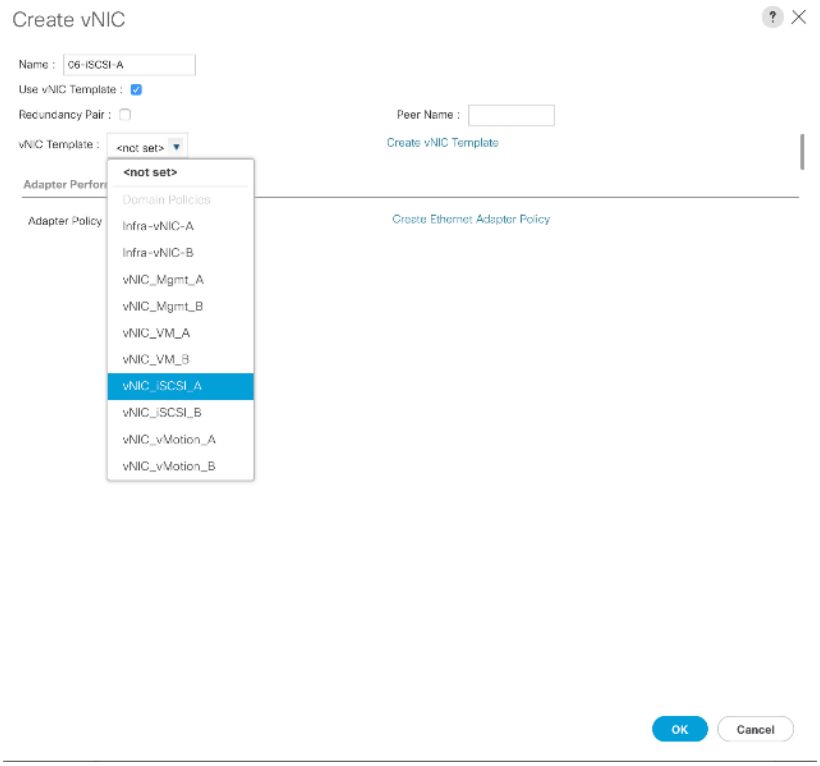

47. Click OK to add this vNIC to the policy.

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

49. In the Create vNIC dialog box, enter 07-iSCSI-B as the name of the vNIC.

50. Select the Use vNIC Template checkbox.

51. In the vNIC Template list, select iSCSI-Template-B.

52. In the Adapter Policy list, select VMWare.

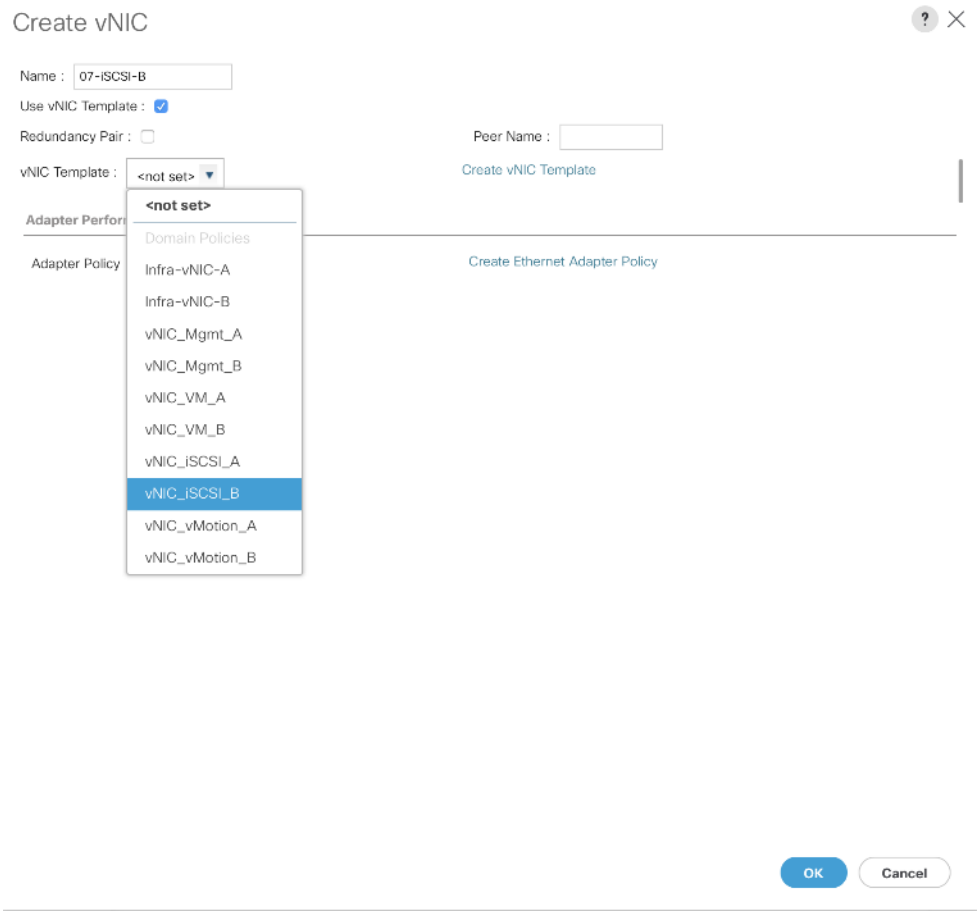

53. Click OK to add this vNIC to the policy.

54. Expand the Add iSCSI vNICs.

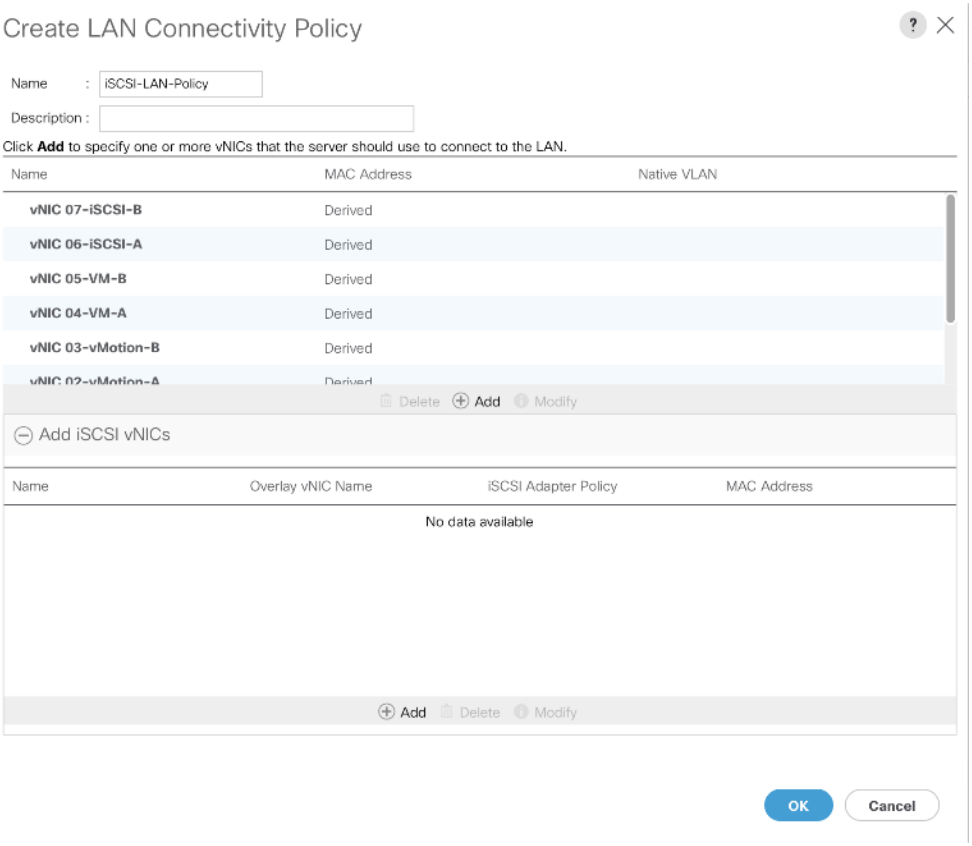

Adding iSCSI vNICs in LAN Policy

This configuration step can be skipped if the UCS environment does not need to access storage environment using iSCSI.

Complete the following steps only if you are using iSCSI SAN access:

1. Verify the iSCSI base vNICs are already added as part of vNIC implementation.
2. Expand the Add iSCSI vNICs section to add the iSCSI boot vNICs.
3. Select Add in the Add iSCSI vNICs section.
4. Set the name to iSCSI–A-vNIC.
5. Select the 06-iSCSI-A as Overlay vNIC.
6. Set the VLAN to iSCSI-A-VLAN (native).
7. Set the iSCSI Adapter Policy to default
8. Leave the MAC Address set to None.

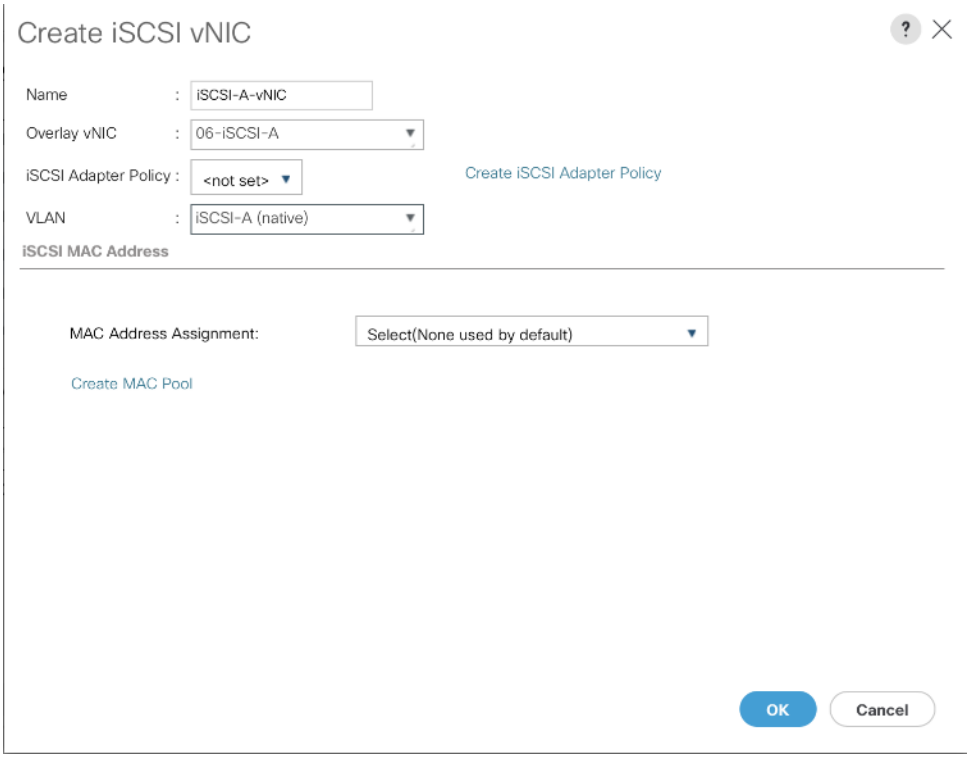

9. Click OK.
10. Select Add in the Add iSCSI vNICs section.
11. Set the name to iSCSI-B-vNIC.
12. Select the 07-iSCSI-A as Overlay vNIC.
13. Set the VLAN to iSCSI-B-VLAN.
14. Set the iSCSI Adapter Policy to default.
15. Leave the MAC Address set to None.

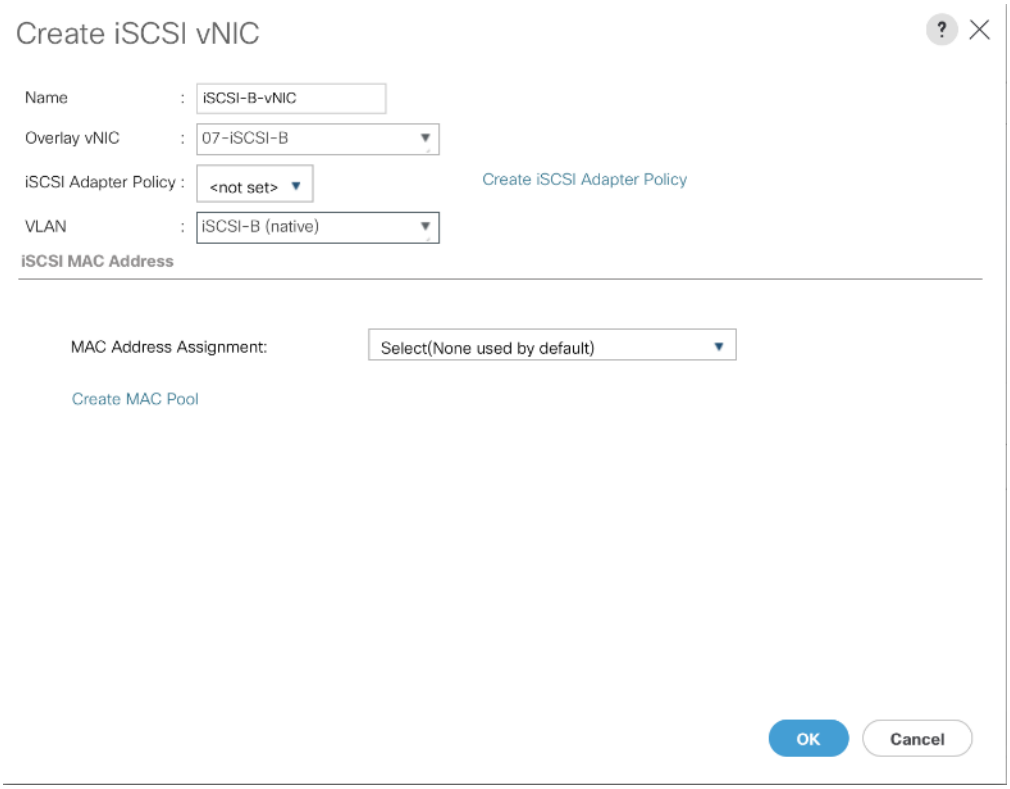

16. Click OK, then click OK again to create the LAN Connectivity Policy.

Create vHBA Templates for FC Connectivity

This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates and choose Create vHBA Template.
4. Enter Infra-vHBA-A as the vHBA template name.
5. Click the radio button to select Fabric A.
6. In the Select VSAN list, Choose VSAN-A.
7. In the WWPN Pool list, Choose WWPN-Pool-A.

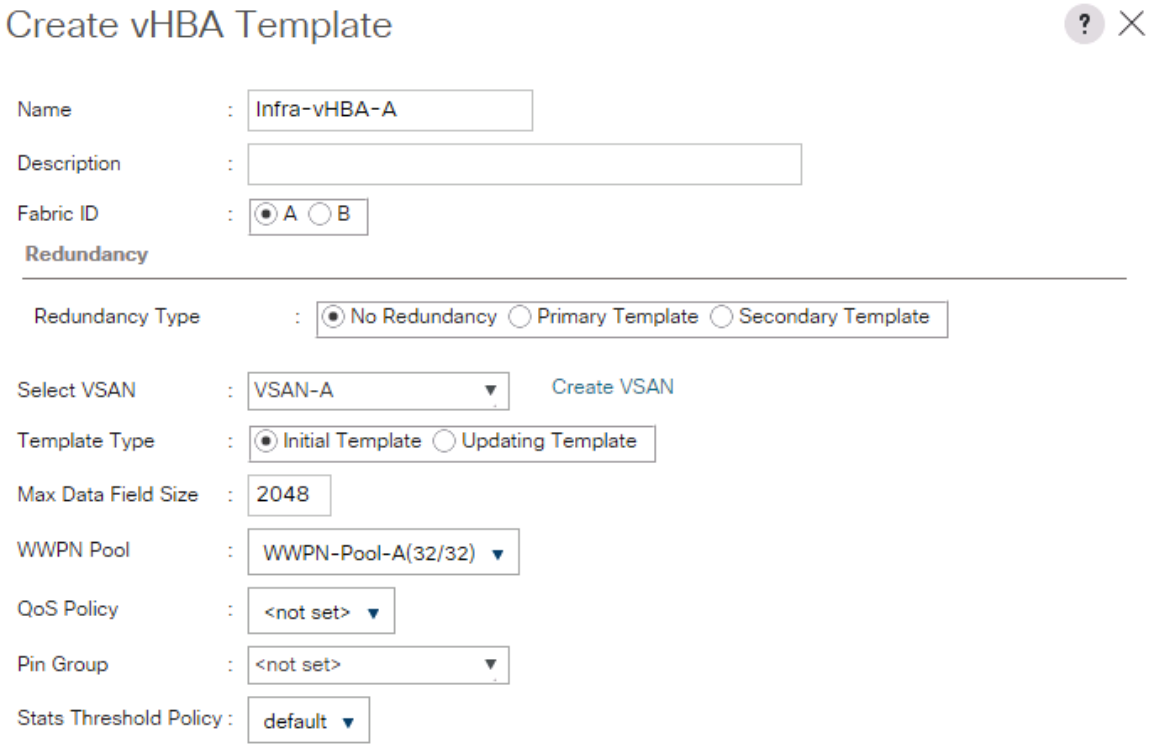

8. Click OK to create the vHBA template.
9. Click OK.
10. Right-click vHBA Templates again and choose Create vHBA Template.
11. Enter Infra-vHBA-B as the vHBA template name.
12. Click the radio button to select Fabric B.
13. In the Select VSAN list, Choose VSAN-B.
14. In the WWPN Pool, Choose WWPN-Pool-B.

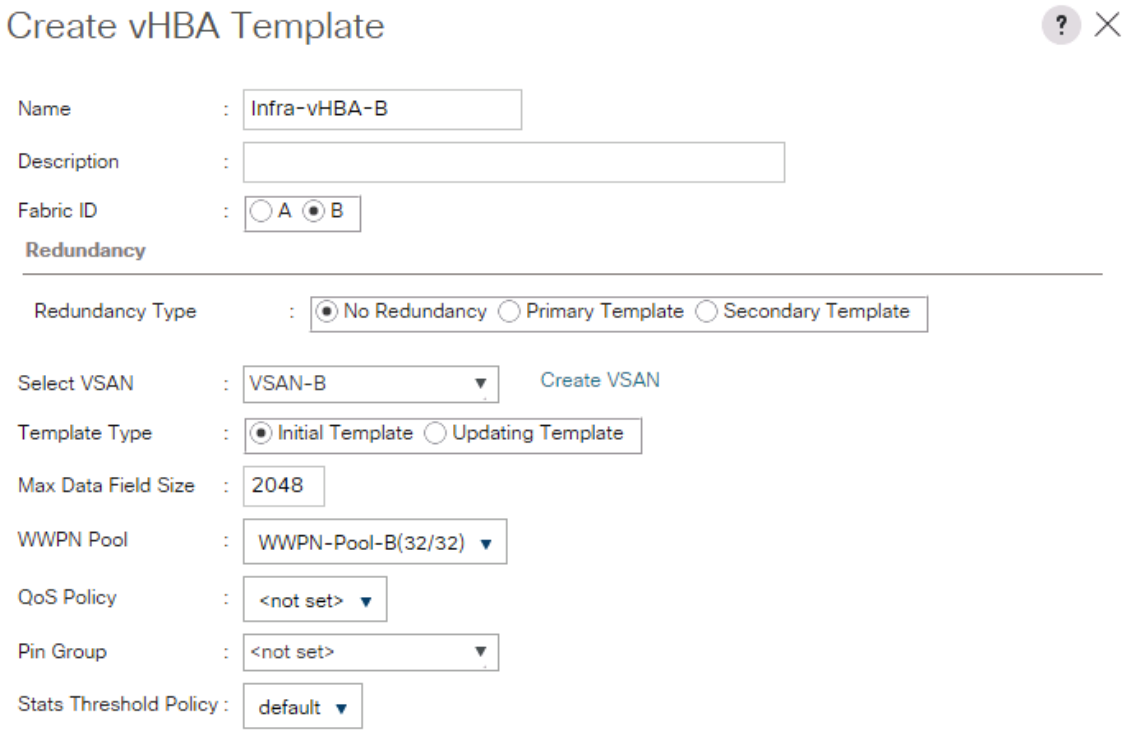

15. Click OK to create the vHBA template.

16. Click OK.

Create FC SAN Connectivity Policies

This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

A SAN connectivity policy defines the vHBAs that will be created as part of a service profile deployment.

To configure the necessary FC SAN Connectivity Policies, complete the following 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 and choose Create SAN Connectivity Policy.
4. Enter Infra-FC-pol as the name of the policy.
5. Select WWNN-Pool from the drop-down list under World Wide Node Name.

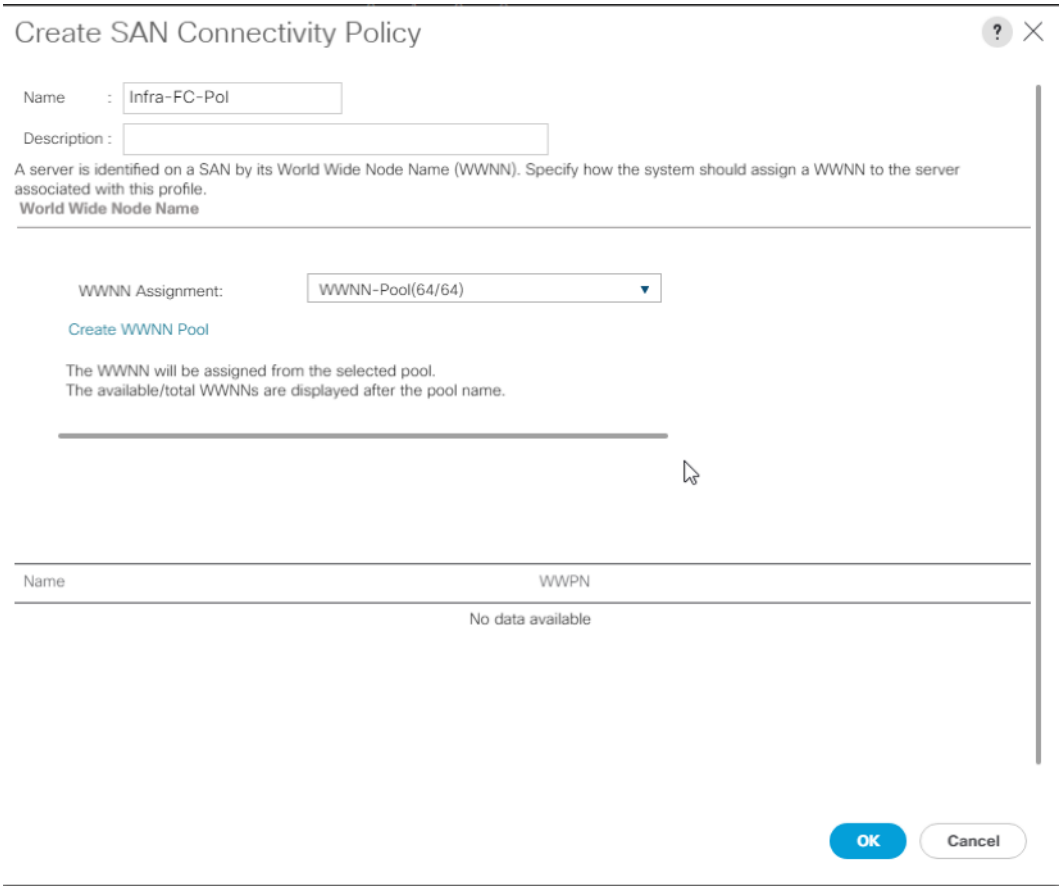

6. Click Add. You might have to scroll down the screen to see the Add link.
7. Under Create vHBA, enter vHBA-A in the Name field.
8. Check the check box Use vHBA Template.
9. From the vHBA Template drop-down list, select Infra-vHBA-A.
10. From the Adapter Policy drop-down list, select VMWare.

11. Click OK.

12. Click Add.
13. Under Create vHBA, enter vHBA-B in the Name field.
14. Check the check box next to Use vHBA Template.
15. From the vHBA Template drop-down list, select infra-vHBA-B.
16. From the Adapter Policy drop-down list, select VMWare.

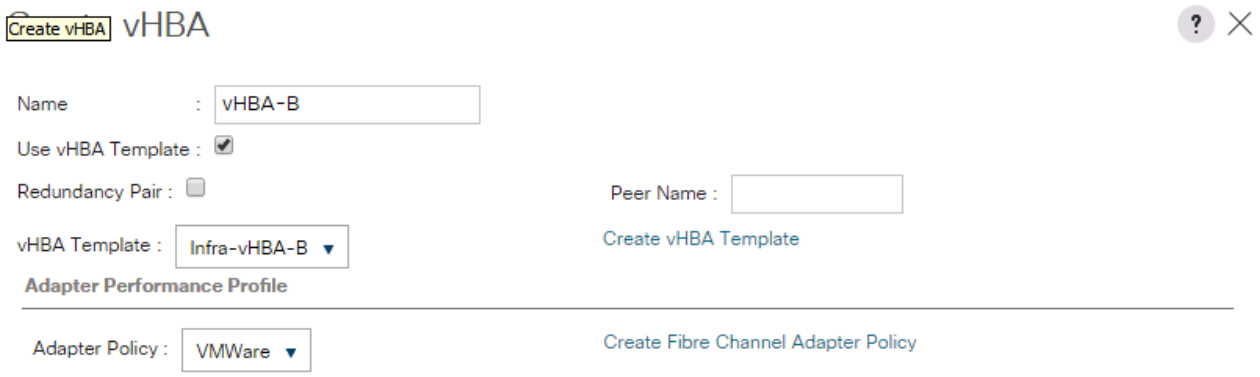

17. Click OK.
18. Click OK again to accept creating the SAN connectivity policy.

Create iSCSI Boot Policy

This configuration step can be skipped if the UCS environment does not need to access storage environment using iSCSI.

This procedure applies to a Cisco UCS environment in which iSCSI interface on Controller A is chosen as the primary target.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies and choose Create Boot Policy.
4. Enter Boot-iSCSI-X-A as the name of the boot policy.
5. Optional: Enter a description for the boot policy.
6. Keep the Reboot on Boot Order Change option cleared.

7. Expand the Local Devices drop-down list and select Add CD/DVD.
8. Expand the iSCSI vNICs section and select Add iSCSI Boot.
9. In the Add iSCSI Boot dialog box, enter iSCSI-A.
10. Click OK.
11. Select Add iSCSI Boot.
12. In the Add iSCSI Boot dialog box, enter iSCSI-B.
13. Click OK.

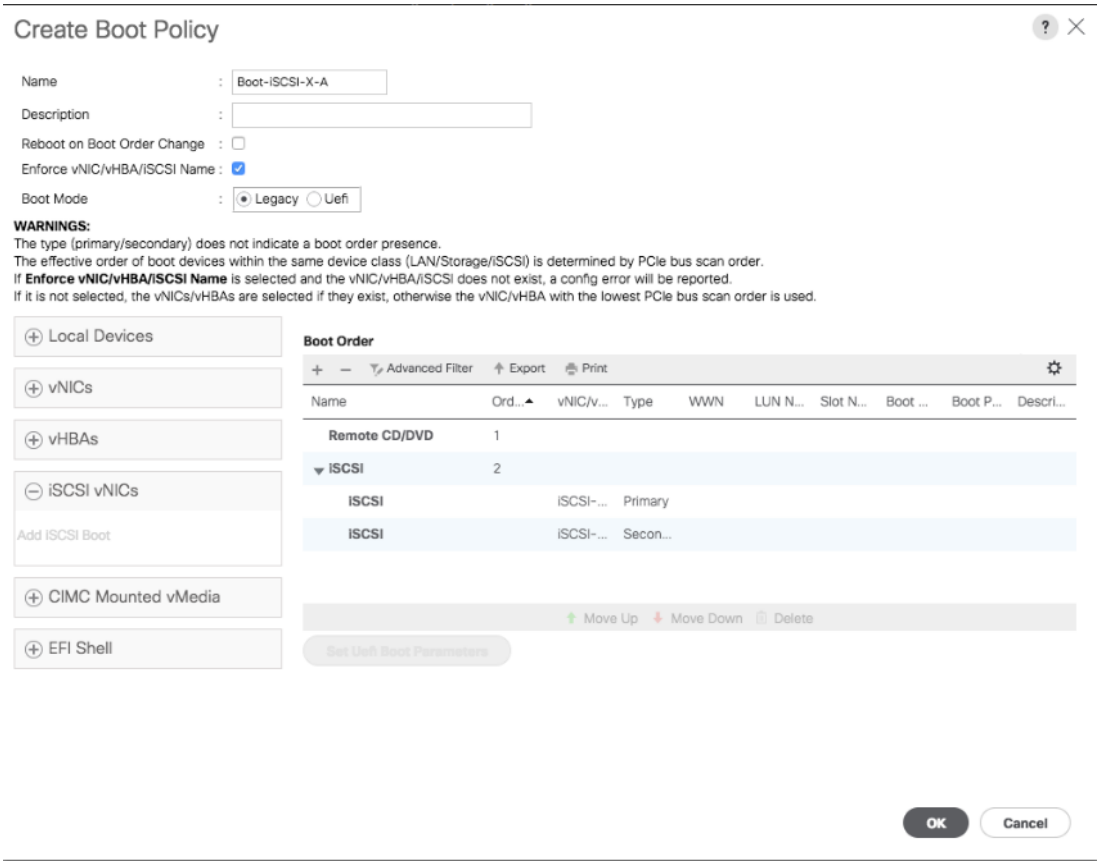

14. Click OK then OK again to save the boot policy.

Create FC Boot Policies

This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

This procedure applies to a Cisco UCS environment in which two FC interfaces are used on each of the SVC nodes for host connectivity. This procedure captures a single boot policy which defines Fabric-A as the

primary fabric. Customer can choose to create a second boot policy which can use Fabric-B as primary fabric to spread the boot-from-san traffic load on both the nodes in case of disaster recovery.

WWPN information from the IBM SVC nodes is required to complete this section. This information can be found by logging into the IBM SVC management address using SSH and issuing the commands as captured below. The information can be recorded in Table 11 .

Since NPIV feature is enabled on the IBM SVC systems, the WWPN permitted for host communication can be different from the physical WWPN. Refer to the example below.

1. Verify the node_id of the SVC nodes using the following command (node1 through node 2 in this example):

   svcinfo lsportfc | grep fc

```
IBM_2145:VersaStack-SVC:admin>svcinfo lsportfc |grep fc
id fc_io_port_id port_id type     port_speed node_id node_name WWPN            nportid status               attachment cluster_use   adapter_location adapter_port_id
0  1            1       fc       16Gb       13      node1     500507680C1193C2 400280  active               switch     local_partner 3                1
1  2            2       fc       16Gb       13      node1     500507680C1293C2 770280  active               switch     local_partner 3                2
2  3            3       fc       16Gb       13      node1     500507680C1393C2 4002A0  active               switch     local_partner 3                3
3  4            4       fc       16Gb       13      node1     500507680C1493C2 7702A0  active               switch     local_partner 3                4
4  5            5       fc       N/A        13      node1     500507680C2193C2 000000  inactive_unconfigured none       local_partner 4                1
5  6            6       fc       N/A        13      node1     500507680C2293C2 000000  inactive_unconfigured none       local_partner 4                2
6  7            7       fc       N/A        13      node1     500507680C2393C2 000000  inactive_unconfigured none       local_partner 4                3
7  8            8       fc       N/A        13      node1     500507680C2493C2 000000  inactive_unconfigured none       local_partner 4                4
16 1            1       fc       16Gb       12      node2     500507680C1193C8 400240  active               switch     local_partner 3                1
17 2            2       fc       16Gb       12      node2     500507680C1293C8 770240  active               switch     local_partner 3                2
18 3            3       fc       16Gb       12      node2     500507680C1393C8 400260  active               switch     local_partner 3                3
19 4            4       fc       16Gb       12      node2     500507680C1493C8 770260  active               switch     local_partner 3                4
20 5            5       fc       N/A        12      node2     500507680C2193C8 000000  inactive_unconfigured none       local_partner 4                1
21 6            6       fc       N/A        12      node2     500507680C2293C8 000000  inactive_unconfigured none       local_partner 4                2
22 7            7       fc       N/A        12      node2     500507680C2393C8 000000  inactive_unconfigured none       local_partner 4                3
23 8            8       fc       N/A        12      node2     500507680C2493C8 000000  inactive_unconfigured none       local_partner 4                4
```

2. Use the following command to record the WWPN corresponding to ports connected to the SAN fabric:

   lstargetportfc -filtervalue host_io_permitted=yes

```
IBM_2145:VersaStack-SVC:admin>lstargetportfc -filtervalue host_io_permitted=yes
id WWPN             WWNN             port_id owning_node_id current_node_id nportid host_io_permitted virtualized
2  500507680C1593C2 500507680C0093C2 1       13            13              400281  yes               yes
4  500507680C1693C2 500507680C0093C2 2       13            13              770281  yes               yes
6  500507680C1793C2 500507680C0093C2 3       13            13              4002A1  yes               yes
8  500507680C1893C2 500507680C0093C2 4       13            13              7702A1  yes               yes
10 500507680C2593C2 500507680C0093C2 5       13                            000000  yes               yes
12 500507680C2693C2 500507680C0093C2 6       13                            000000  yes               yes
14 500507680C2793C2 500507680C0093C2 7       13                            000000  yes               yes
16 500507680C2893C2 500507680C0093C2 8       13                            000000  yes               yes
34 500507680C1593C8 500507680C0093C8 1       12            12              400241  yes               yes
36 500507680C1693C8 500507680C0093C8 2       12            12              770241  yes               yes
38 500507680C1793C8 500507680C0093C8 3       12            12              400261  yes               yes
40 500507680C1893C8 500507680C0093C8 4       12            12              770261  yes               yes
42 500507680C2593C8 500507680C0093C8 5       12                            000000  yes               yes
44 500507680C2693C8 500507680C0093C8 6       12                            000000  yes               yes
46 500507680C2793C8 500507680C0093C8 7       12                            000000  yes               yes
48 500507680C2893C8 500507680C0093C8 8       12                            000000  yes               yes
```

Table 11 IBM SVC – WWPN Information

Node	Port ID	WWPN	Variable	Fabric
SVC Node 1	1		WWPN-SVC-Clus-Node1-FC1-NPIV	A
SVC Node 1	3		WWPN-SVC-Clus-Node1-FC3-NPIV	A
SVC Node 1	2		WWPN-SVC-Clus-Node1-FC2-NPIV	B
SVC Node 1	4		WWPN-SVC-Clus-Node1-FC4-NPIV	B
SVC Node 2	1		WWPN-SVC-Clus-Node2-FC1-NPIV	A
SVC Node 2	3		WWPN-SVC-Clus-Node2-FC3-NPIV	A
SVC Node 2	2		WWPN-SVC-Clus-Node2-FC2-NPIV	B
SVC Node 2	4		WWPN-SVC-Clus-Node2-FC4-NPIV	B

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root.
3. Right-click Boot Policies and choose Create Boot Policy.
4. Enter Boot-Fabric-A as the name of the boot policy.
5. Optional: Enter a description for the boot policy.
6. Keep the Reboot on the Boot Order Change check box unchecked.
7. Expand the Local Devices drop-down list and Choose Add CD/DVD.
8. Expand the vHBAs drop-down list and Choose Add SAN Boot.

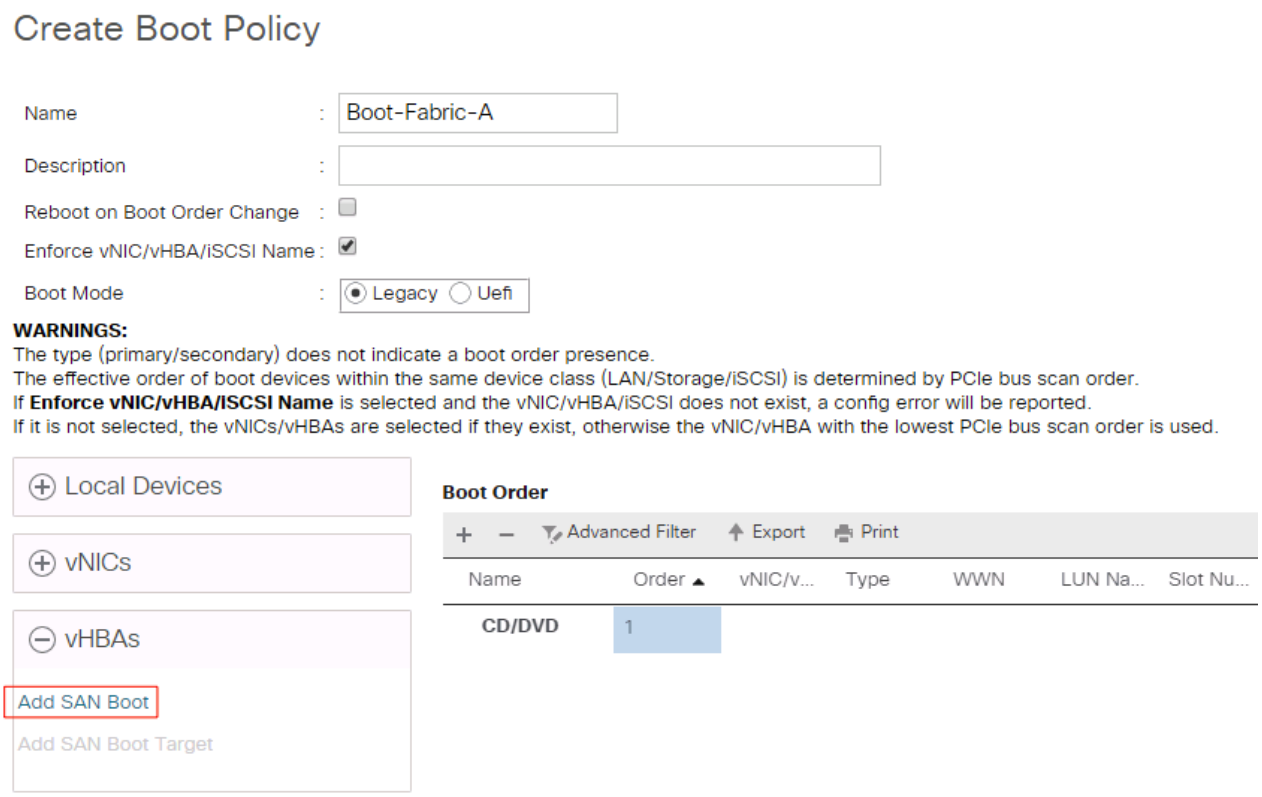

9. Make sure to select Primary radio button as the Type.

10. Enter Fabric-A in the vHBA field.

11. Click OK to add the SAN boot initiator.

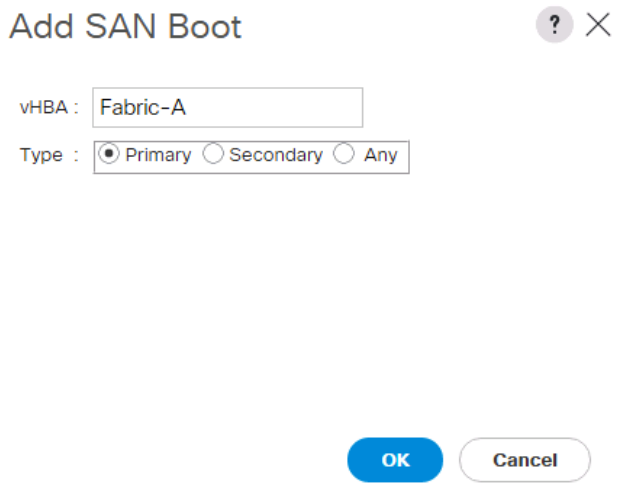

12. From the vHBA drop-down list, choose Add SAN Boot Target.

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

14. Enter the WWPN <WWPN-Node-1-Fabric-A> from Table 11 .

15. Keep the Primary radio button selected as the SAN boot target type.

16. Click OK to add the SAN boot target.

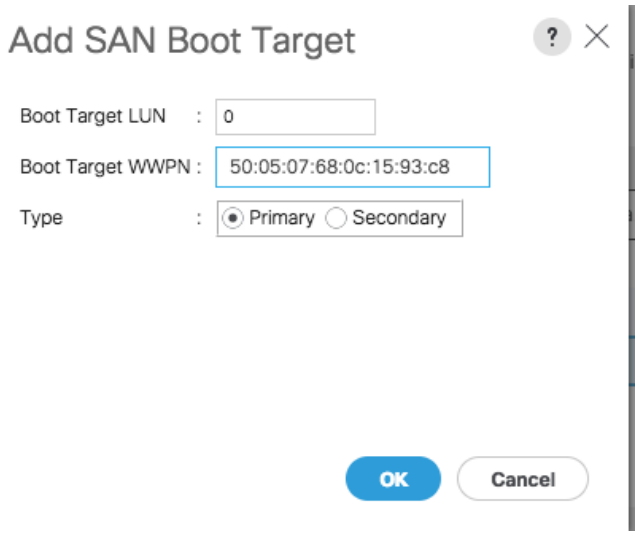

17. From the vHBA drop-down menu, choose Add SAN Boot Target.

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

19. Enter the WWPN <WWPN-Node-2-Fabric-A> from Table 11 .

20. Click OK to add the SAN boot target.

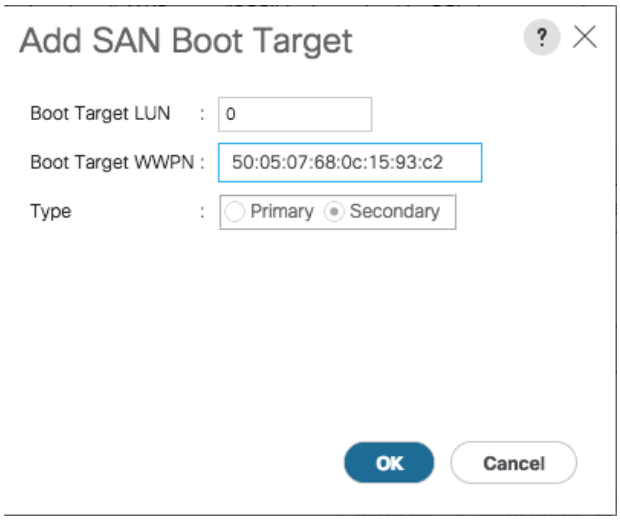

21. From the vHBA drop-down list, choose Add SAN Boot.

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

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

24. Click OK to add the SAN boot initiator.

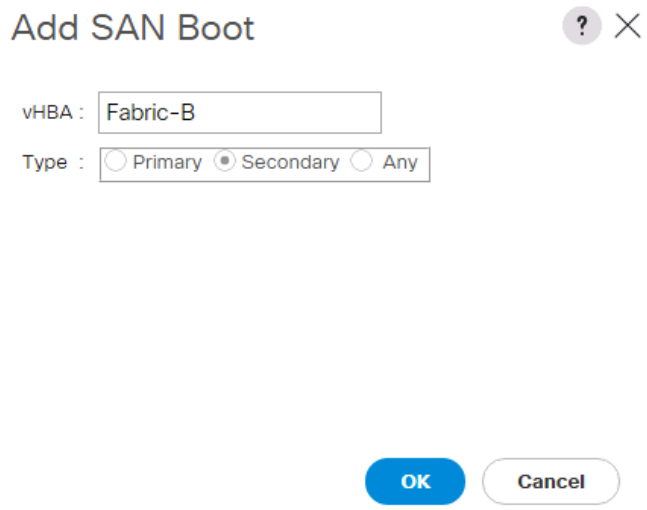

25. From the vHBA drop-down list, choose Add SAN Boot Target.

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

27. Enter the WWPN <WWPN-Node-1-Fabric-B> from Table 11 .

28. Keep Primary as the SAN boot target type.

29. Click OK to add the SAN boot target.

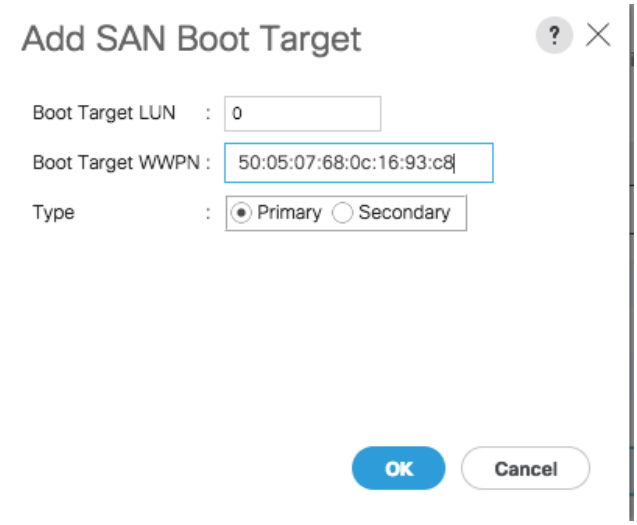

30. From the vHBA drop-down list, choose Add SAN Boot Target.

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

32. Enter the WWPN <WWPN-Node-2-Fabric-B> from Table 11 .

33. Click OK to add the SAN boot target.

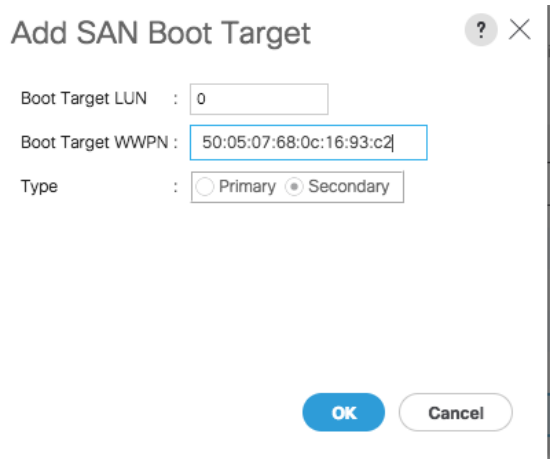

34. Click OK, and then click OK again to create the boot policy.

35. Verify that your SAN boot configuration looks similar to the screenshot below.

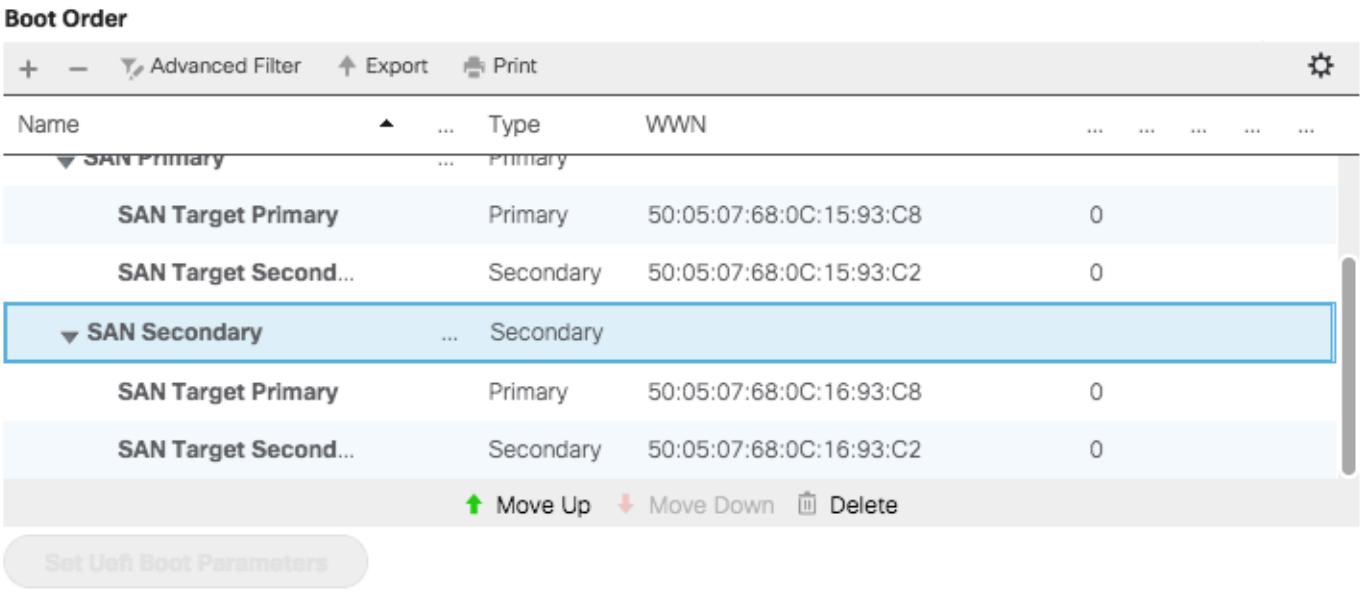

Create iSCSI Boot Service Profile Template

Service profile template configuration for the iSCSI based SAN access is covered in this section.

This section can be skipped if iSCSI boot is not implemented in the customer environment.

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

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

2. Select Service Profile Templates > root.

3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter Infra-ESXi-iSCSI-Host as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.

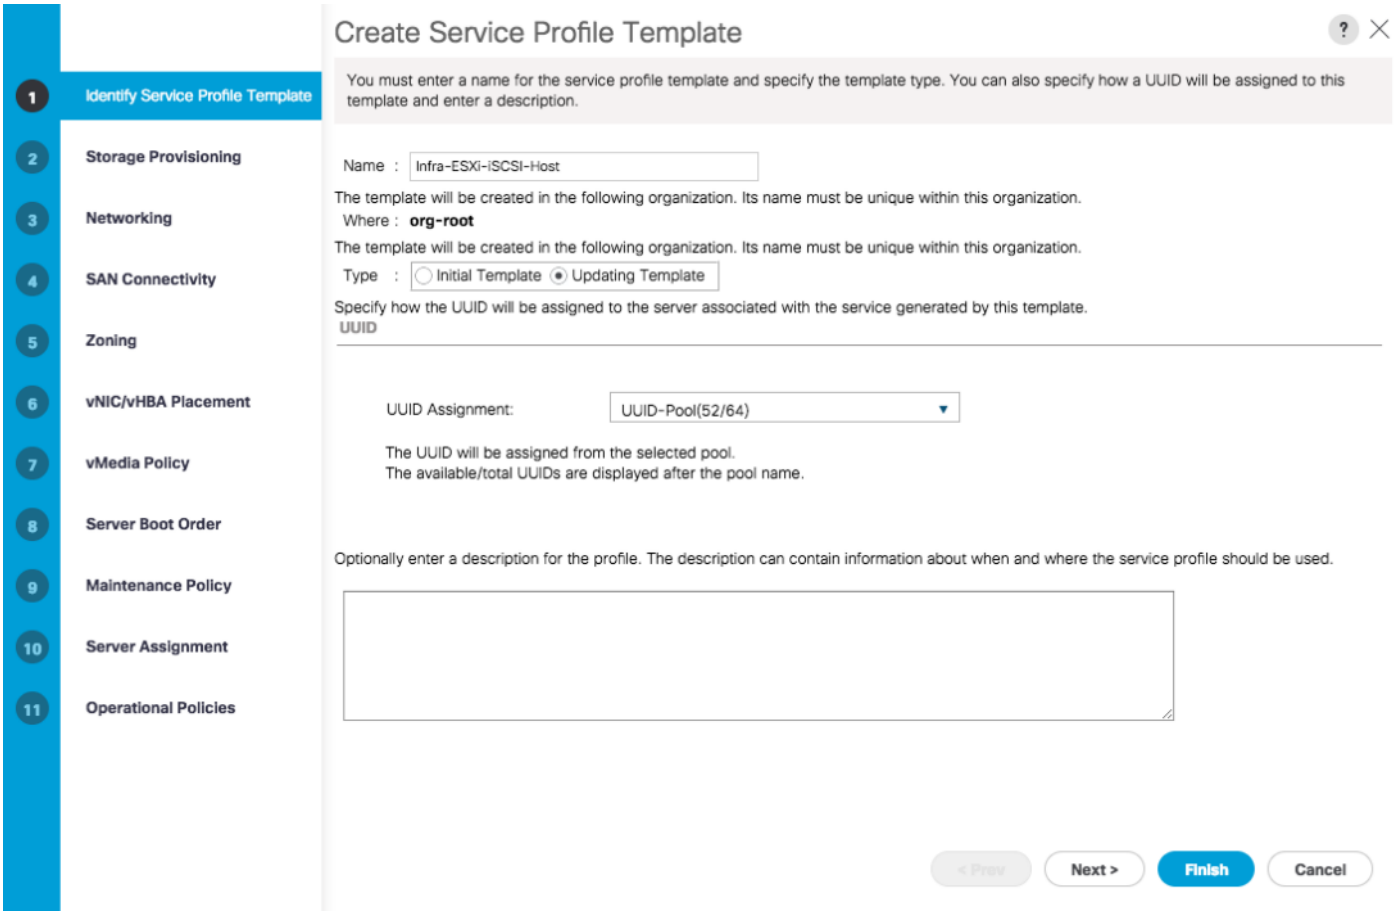

8. Click Next.

Configure Storage Provisioning

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

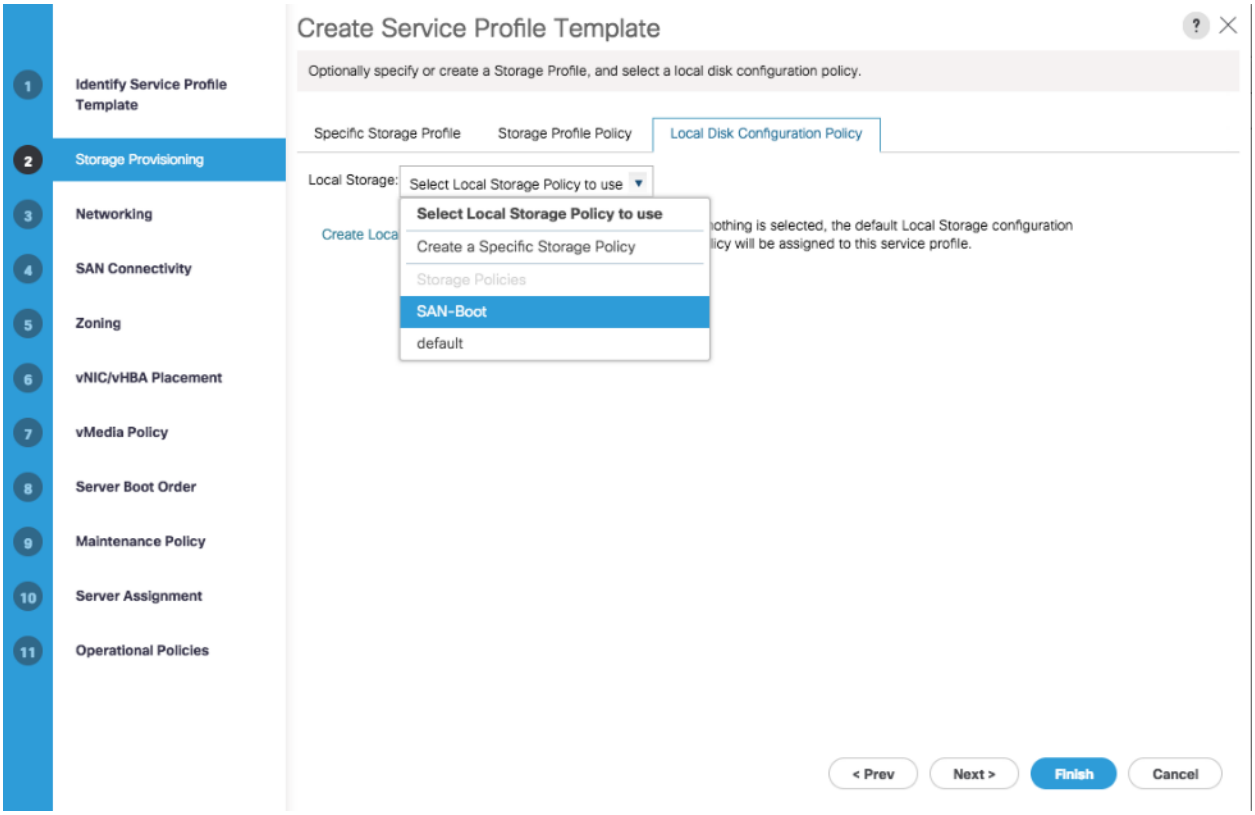

2. Click Next.

Configure Networking Options

To configure the network options, complete the following steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select iSCSI-LAN-Policy from the LAN Connectivity Policy drop-down list.
4. Select IQN_Pool in Initiator Name Assignment.

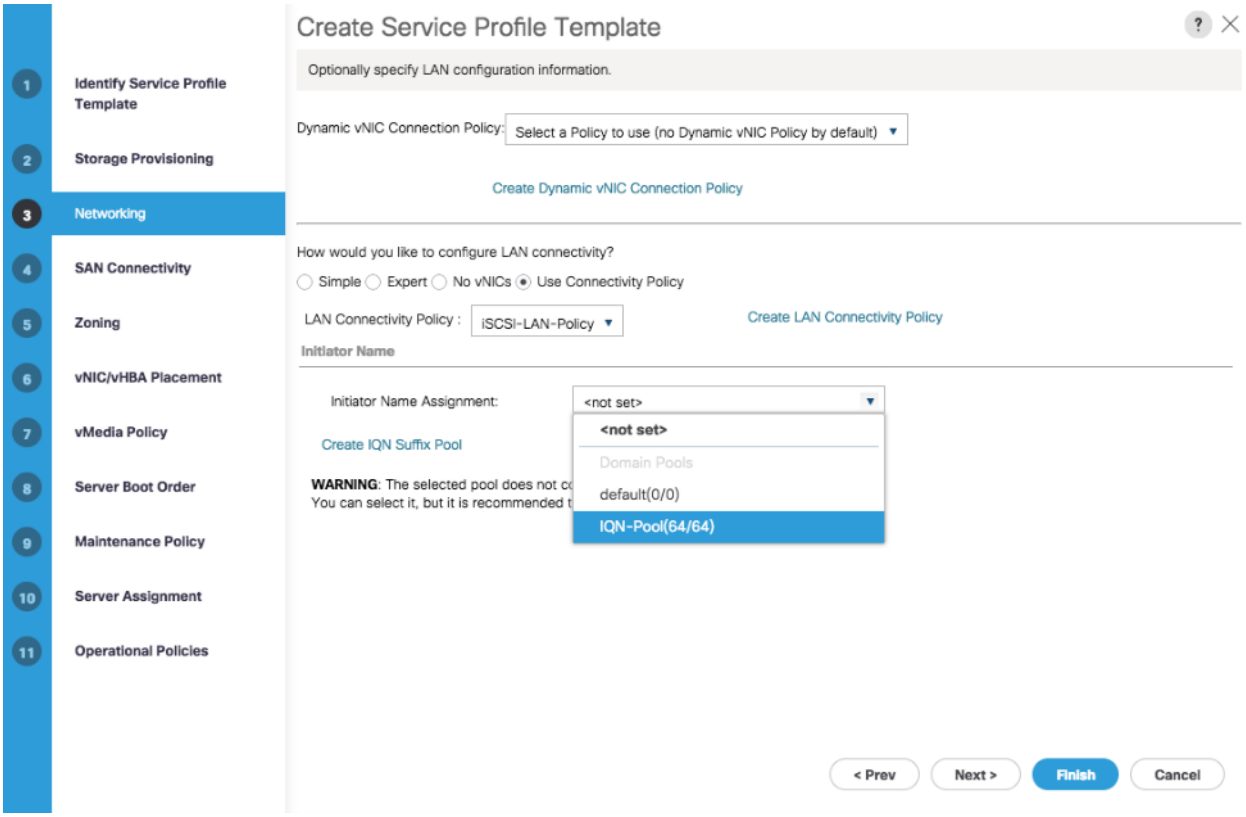

5. Click Next.

Configure Storage Options

1. Select the No vHBA option for the “How would you like to configure SAN connectivity?” field.
2. 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-iSCSI-X-A for Boot Policy.

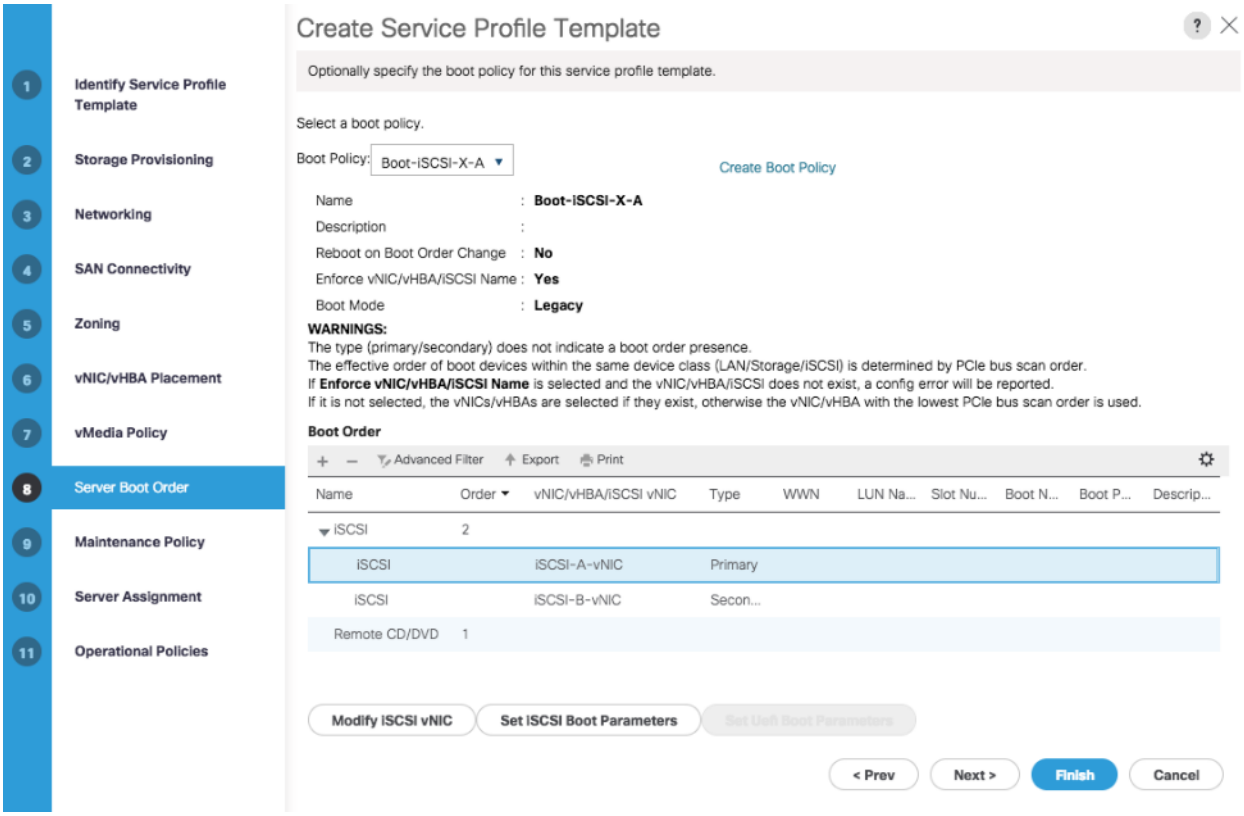

2. In the Boor order, select iSCSI-A-vNIC.
3. Click Set iSCSI Boot Parameters button.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.
5. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
6. Set iSCSI_IP_Pool_A as the “Initiator IP address Policy.”
7. Select iSCSI Static Target Interface option.
8. Click Add.
9. Enter the IP address of Node 1 iSCSI-A interface from Table 7 .

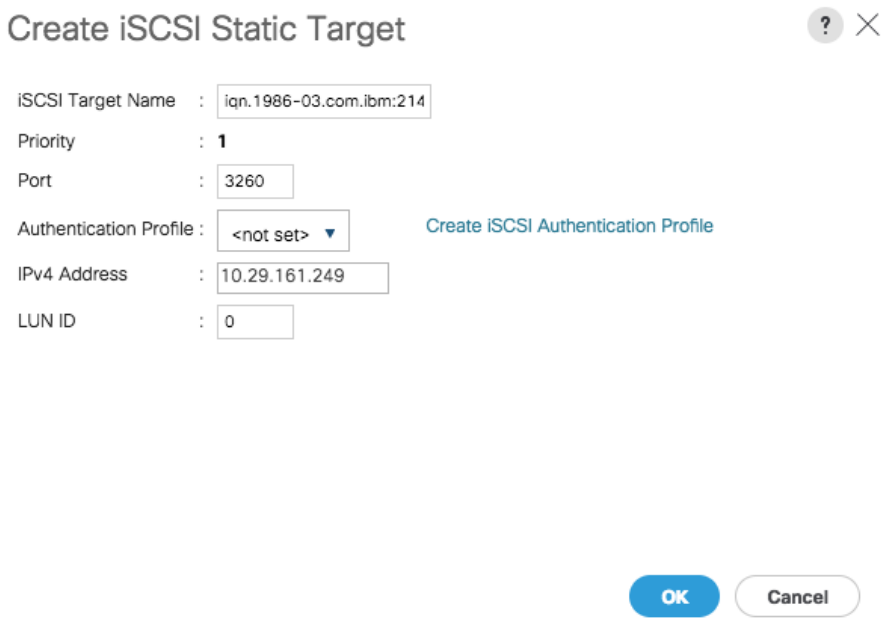

10. Click OK to add the iSCSI Static Target.
11. Keep the iSCSI Static Target Interface option selected and click Add.
12. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 2 (IQN) from Table 8 .
13. Enter the IP address of Node 2 iSCSI-A interface from Table 7 .

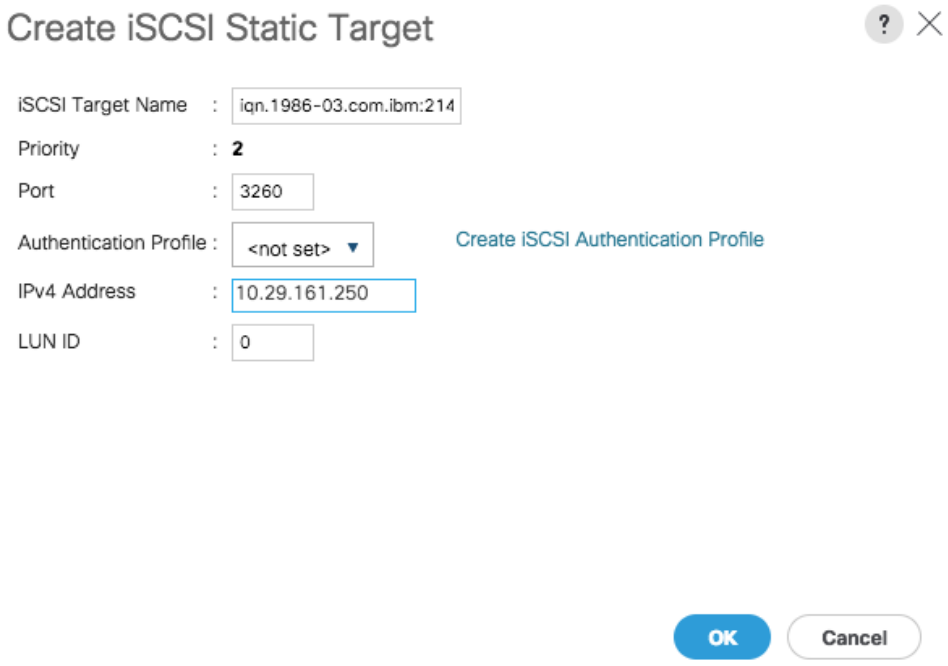

14. Click OK to add the iSCSI Static Target.

15. Verify both the targets on iSCSI Path A as shown below:

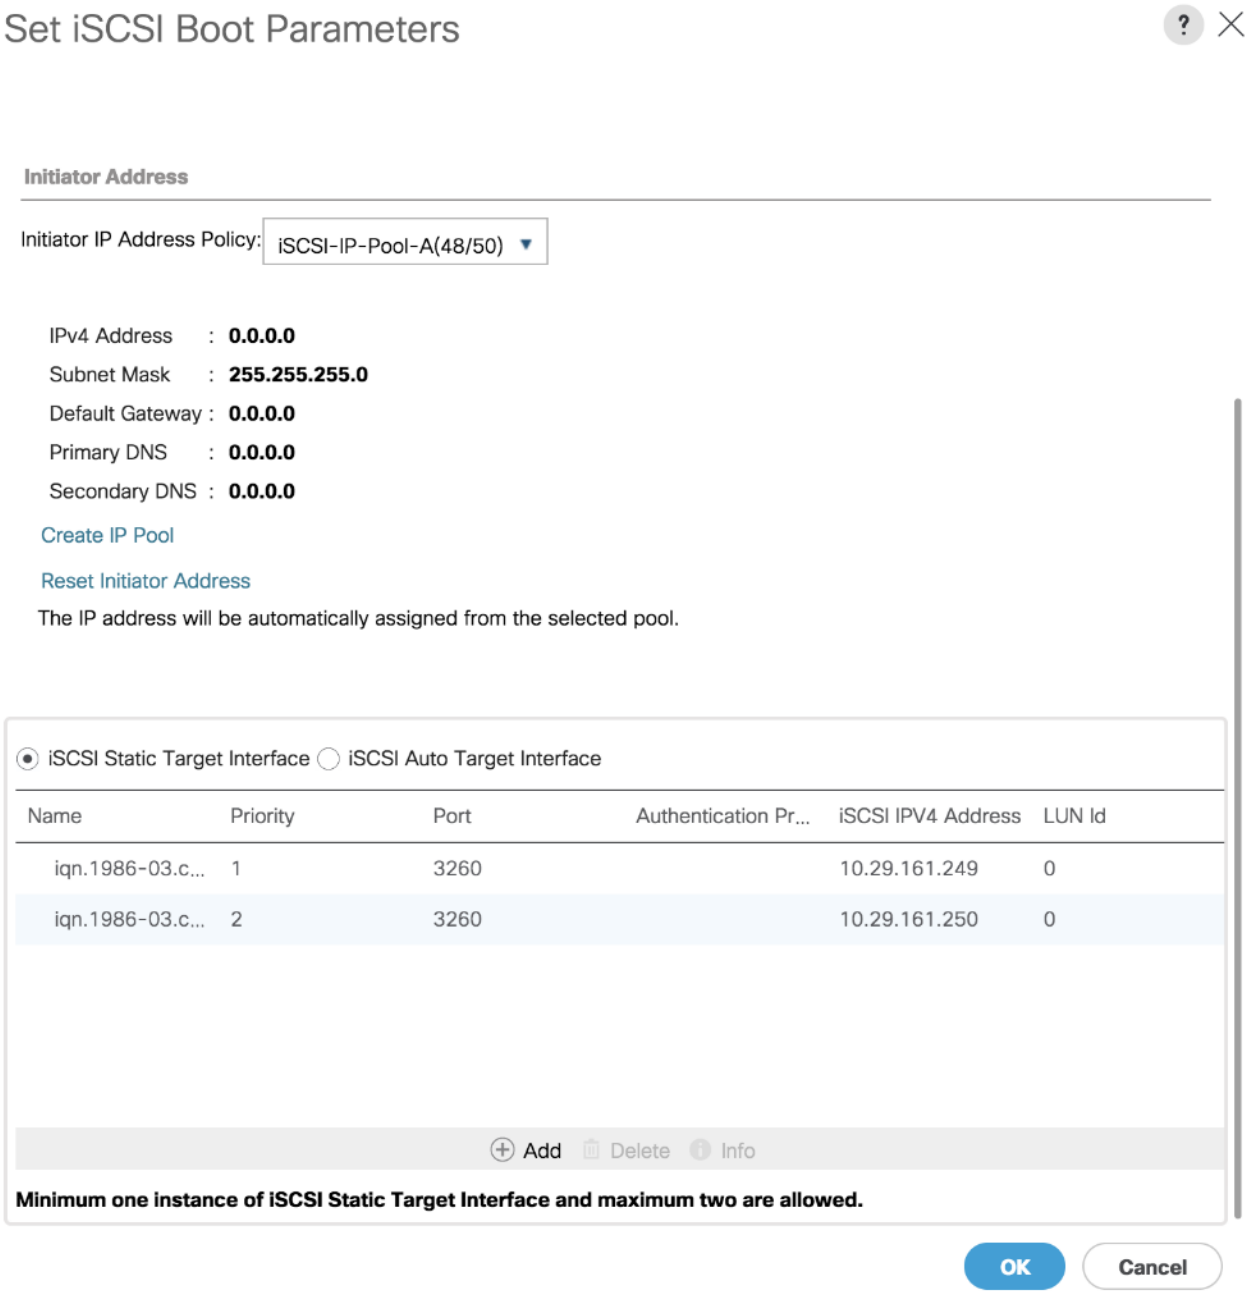

16. Click OK to set the iSCSI-A-vNIC ISCSI Boot Parameters.

17. In the Boor order, select iSCSI-B-vNIC.

18. Click Set iSCSI Boot Parameters button.

19. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.

20. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.

21. Set iSCSI_IP_Pool_B as the “Initiator IP address Policy”.

22. Select iSCSI Static Target Interface option.

23. Click Add.

24. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 1 (IQN) from Table 8 .

25. Enter the IP address of Node 1 iSCSI-B interface from Table 7 .

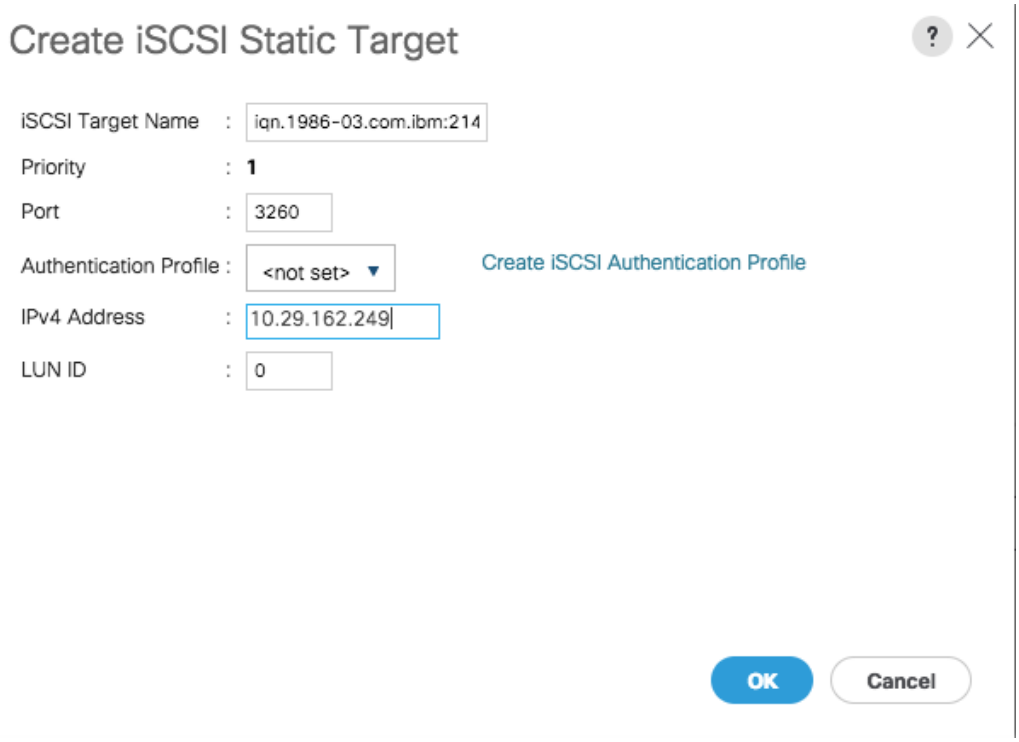

26. Click OK to add the iSCSI Static Target.

27. Keep the iSCSI Static Target Interface option selected and click Add.

28. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 2 (IQN) from Table 8 .

29. Enter the IP address of Node 2 iSCSI-B interface from Table 7 .

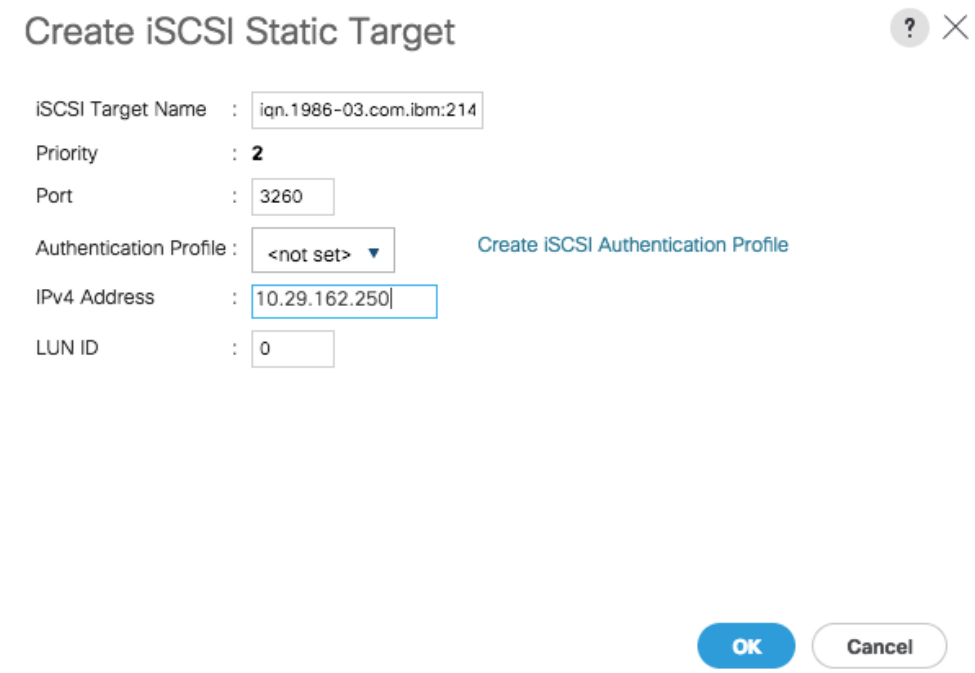

30. Click OK to add the iSCSI Static Target.

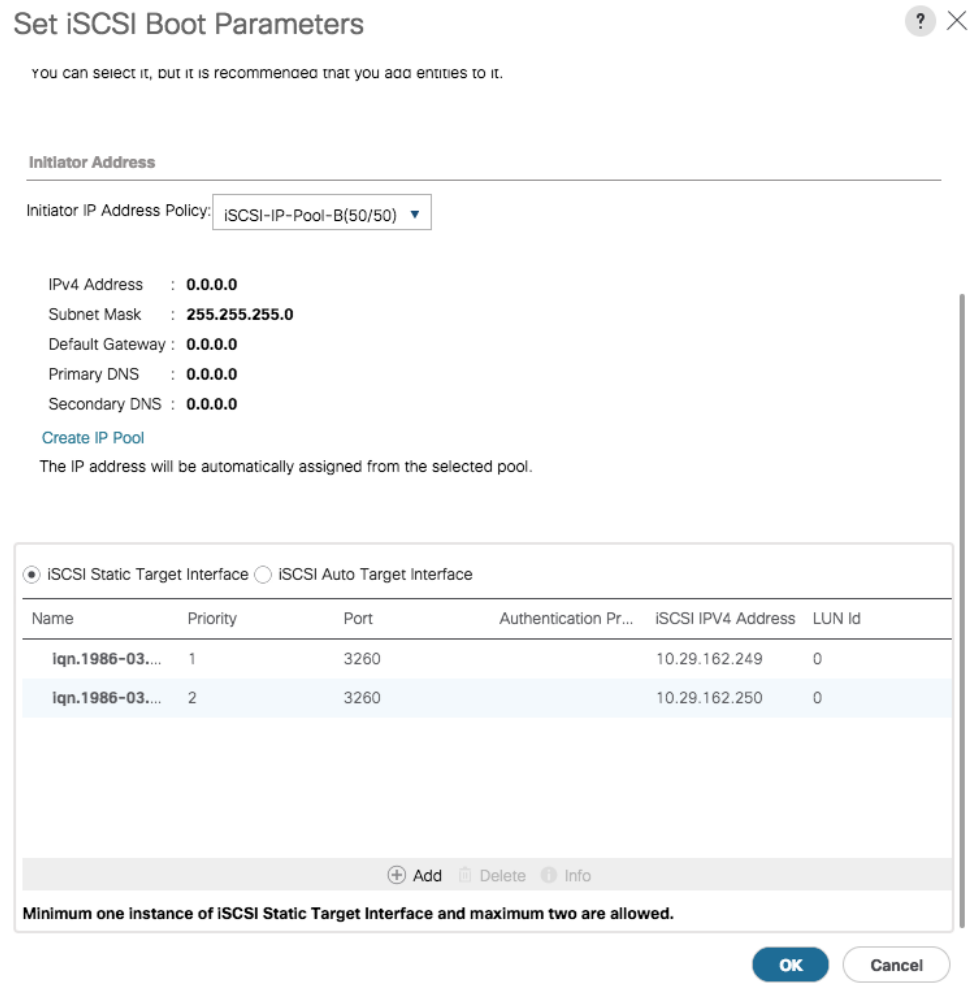

31. Click OK to set the iSCSI-B-vNIC ISCSI Boot Parameters.

32. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

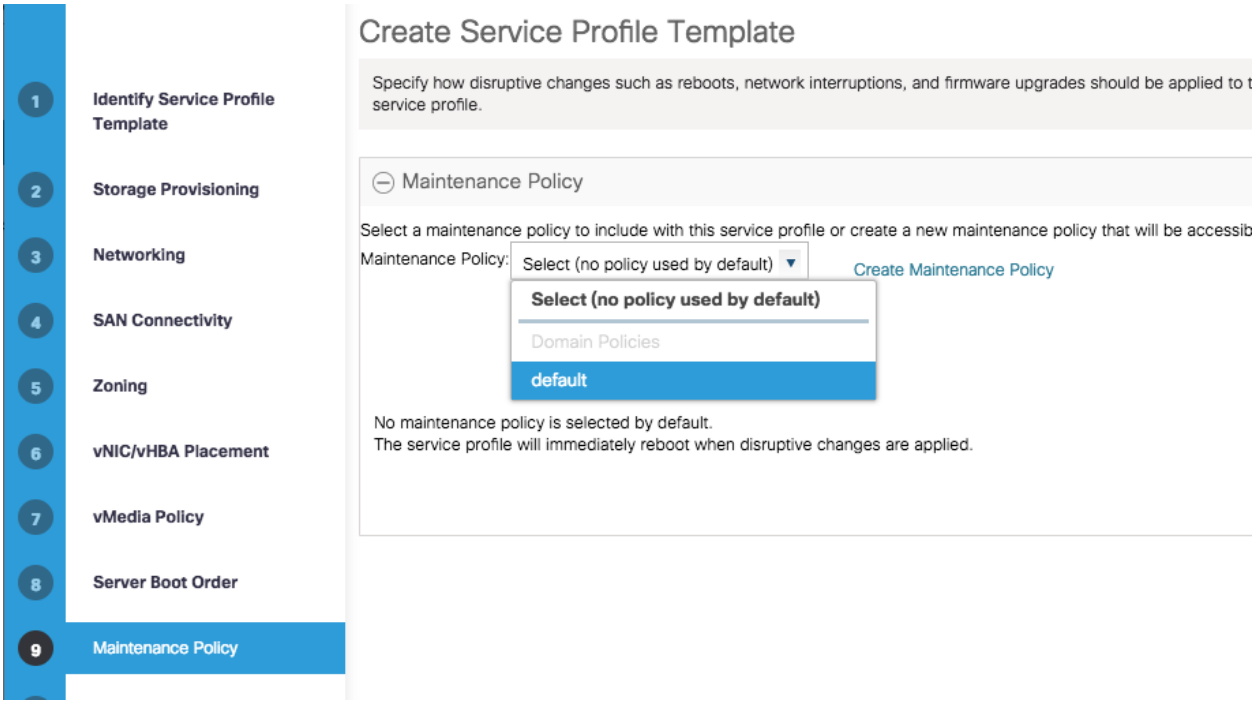

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select Infra-Server-Pool.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. 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.

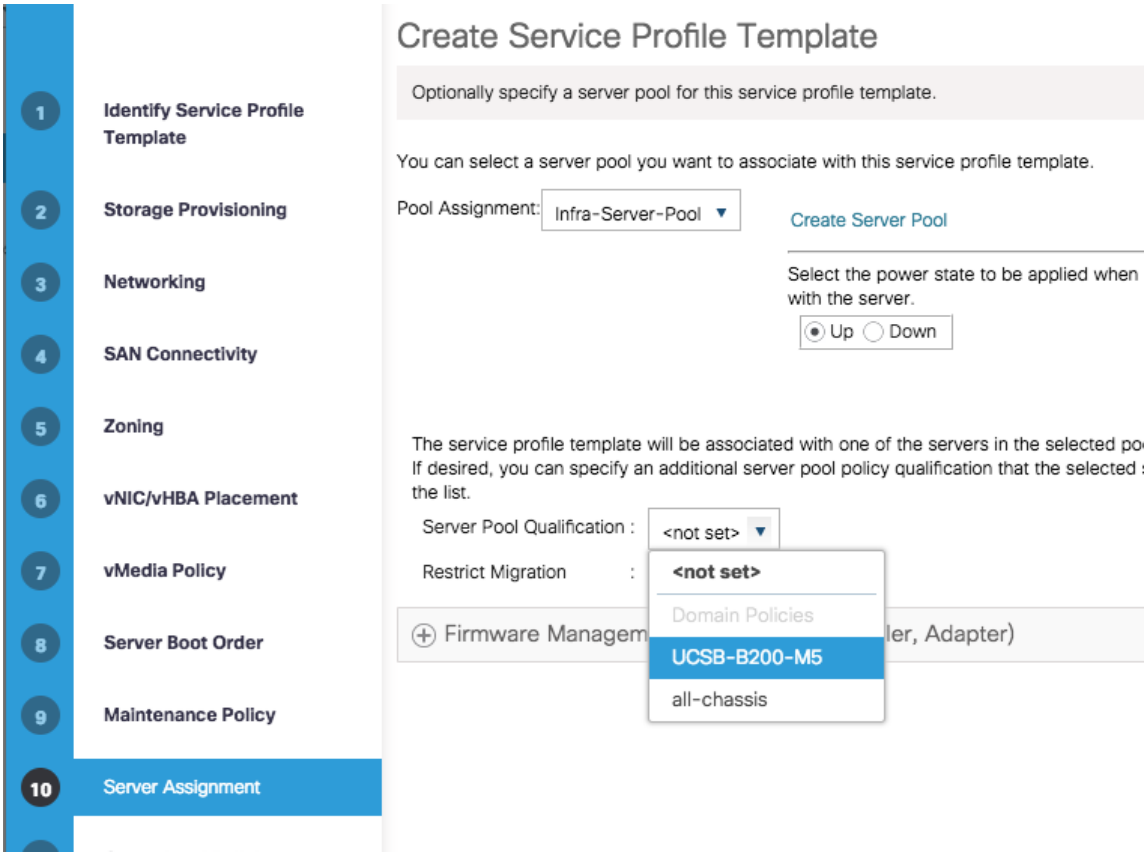

5. Click Next

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

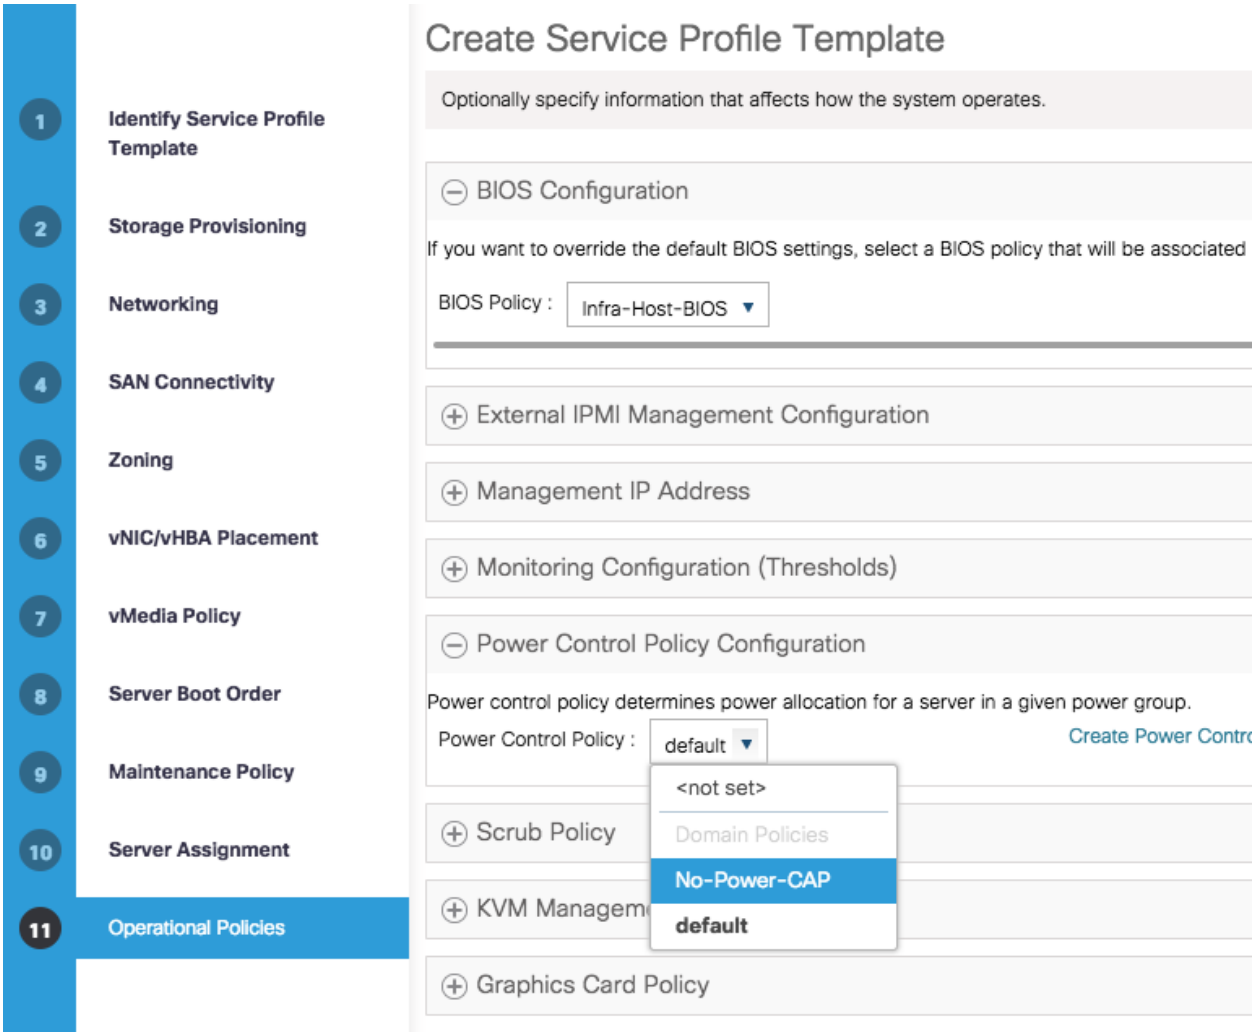

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create iSCSI Boot Service Profiles

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

1. Connect to the UCS 6332-16UP Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Template VM-Host-iSCSI-A.
3. Right-click Infra-ESXi-iSCSI-Host and select Create Service Profiles from Template.
4. Enter Infra-ESXi-Host-iSCSI-Host- for iSCSI deployment as the service profile prefix
5. Enter 1 as the Name Suffix Staring Number.

6. Enter the Number of servers to be deploy in the Number of Instances field.
7. Four service profiles were deployed during this validation
8. Click OK to create the service profile.

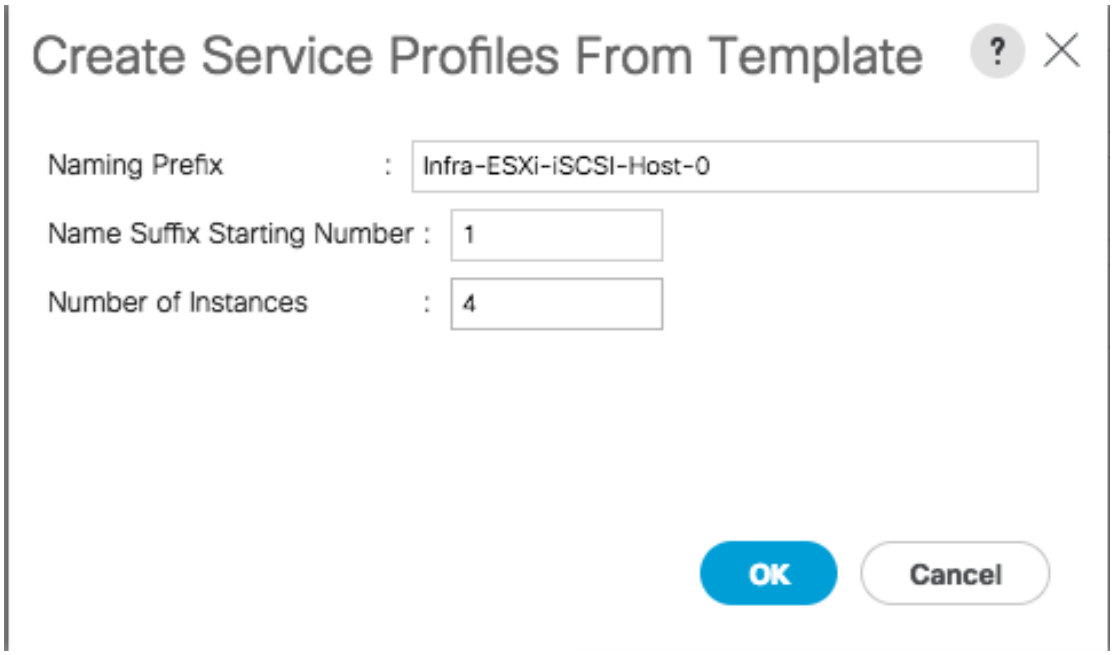

9. Click OK in the confirmation message to provision four VersaStack Service Profiles.

Create FC Boot Service Profile Template

In this procedure, a service profile template is created to use FC Fabric A as primary boot path.

This section can be skipped if FC boot is not implemented in the customer environment.

To create service profile templates, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root.
3. Right-click root and choose Create Service Profile Template. This opens the Create Service Profile Template wizard.
4. Enter Infra-ESXi-Host as the name of the service profile template.
5. Select the Updating Template option.
6. Under UUID, select UUID-Pool as the UUID pool.
7. Click Next.

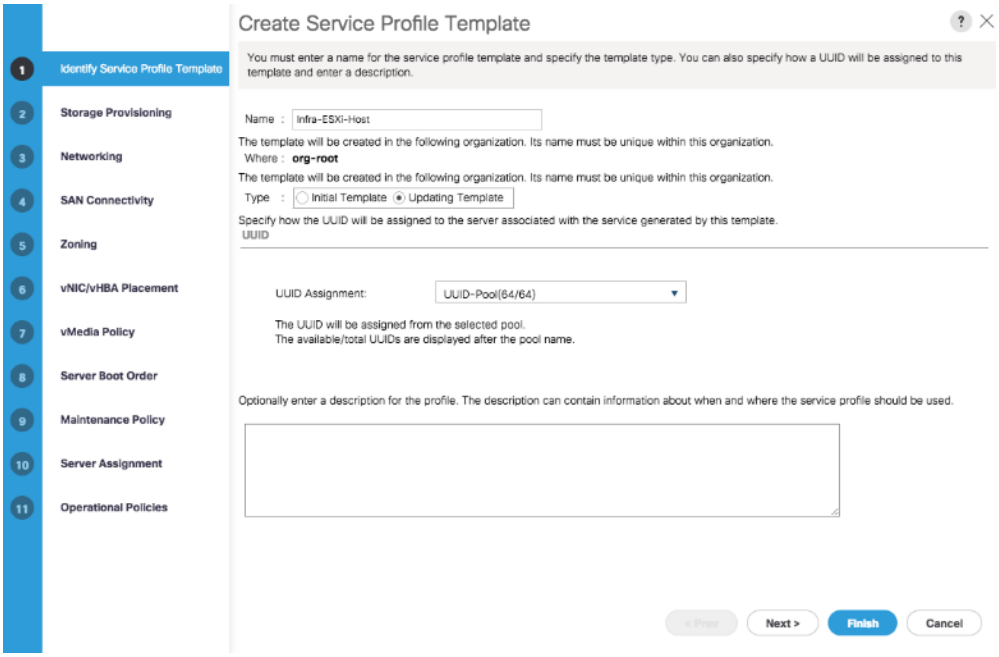

Configure Storage Provisioning

1. Select the Local Disk Configuration Policy tab.
2. Select the SAN-Boot Local Storage Policy. This policy usage requires servers with no local HDDs.
3. Click Next.

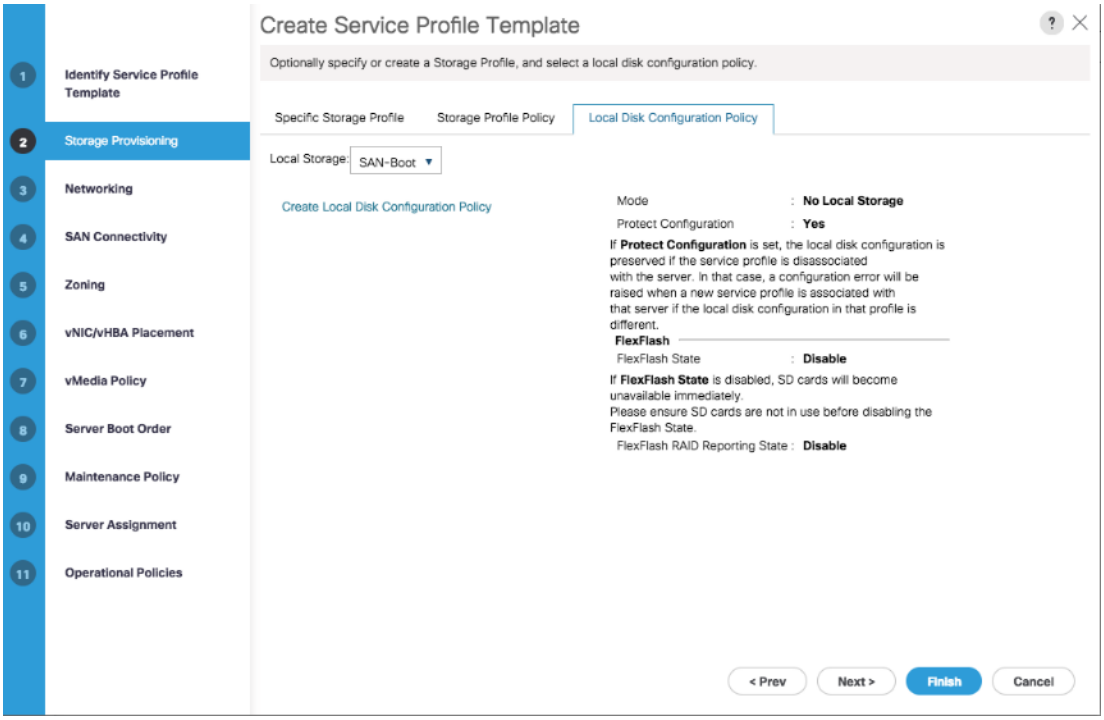

Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Use Connectivity Policy option to configure the LAN connectivity.
3. Select the Infra-LAN-Pol as the LAN Connectivity Policy.
4. Click Next.

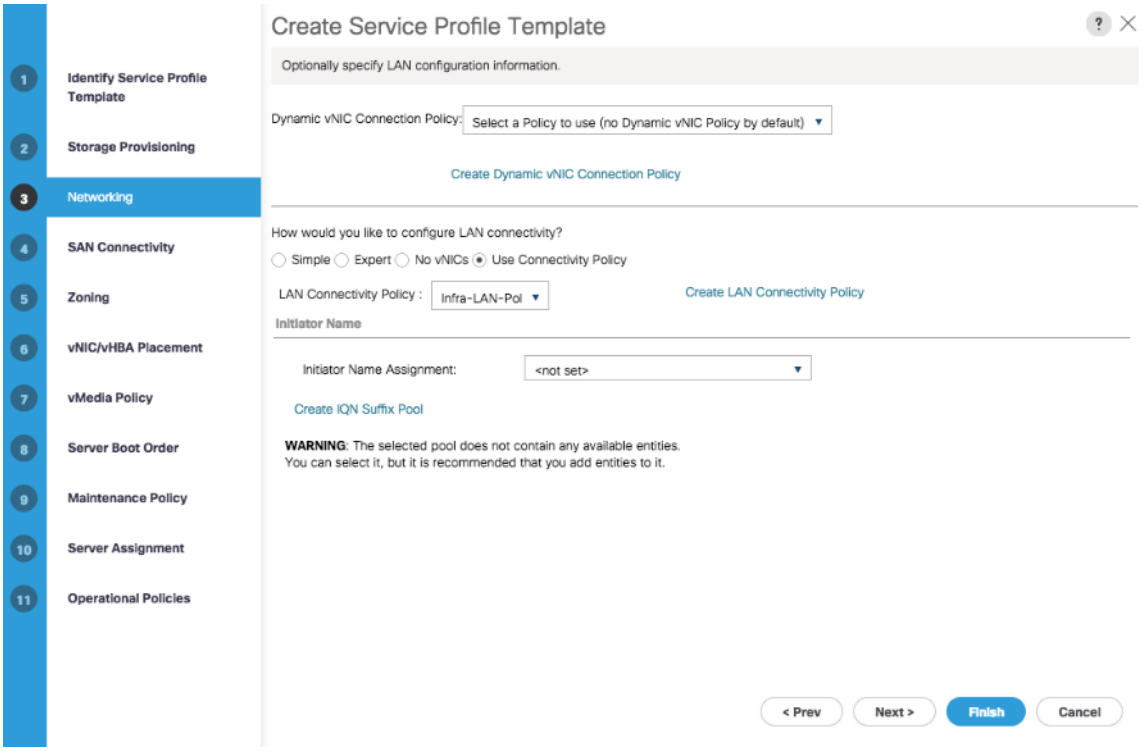

Configure SAN Connectivity

1. Select the Use Connectivity Policy option to configure the SAN connectivity.
2. Select the Infra-FC-Pol as the SAN Connectivity Policy.
3. Click Next.

Configure Zoning

1. It is not necessary to configure any Zoning options.
2. 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. There is no need to set a vMedia Policy.
2. Click Next.

Configure Server Boot Order

1. Select Boot-Fabric-A as the Boot Policy
2. Verify all the boot devices are listed correctly
3. Click Next.

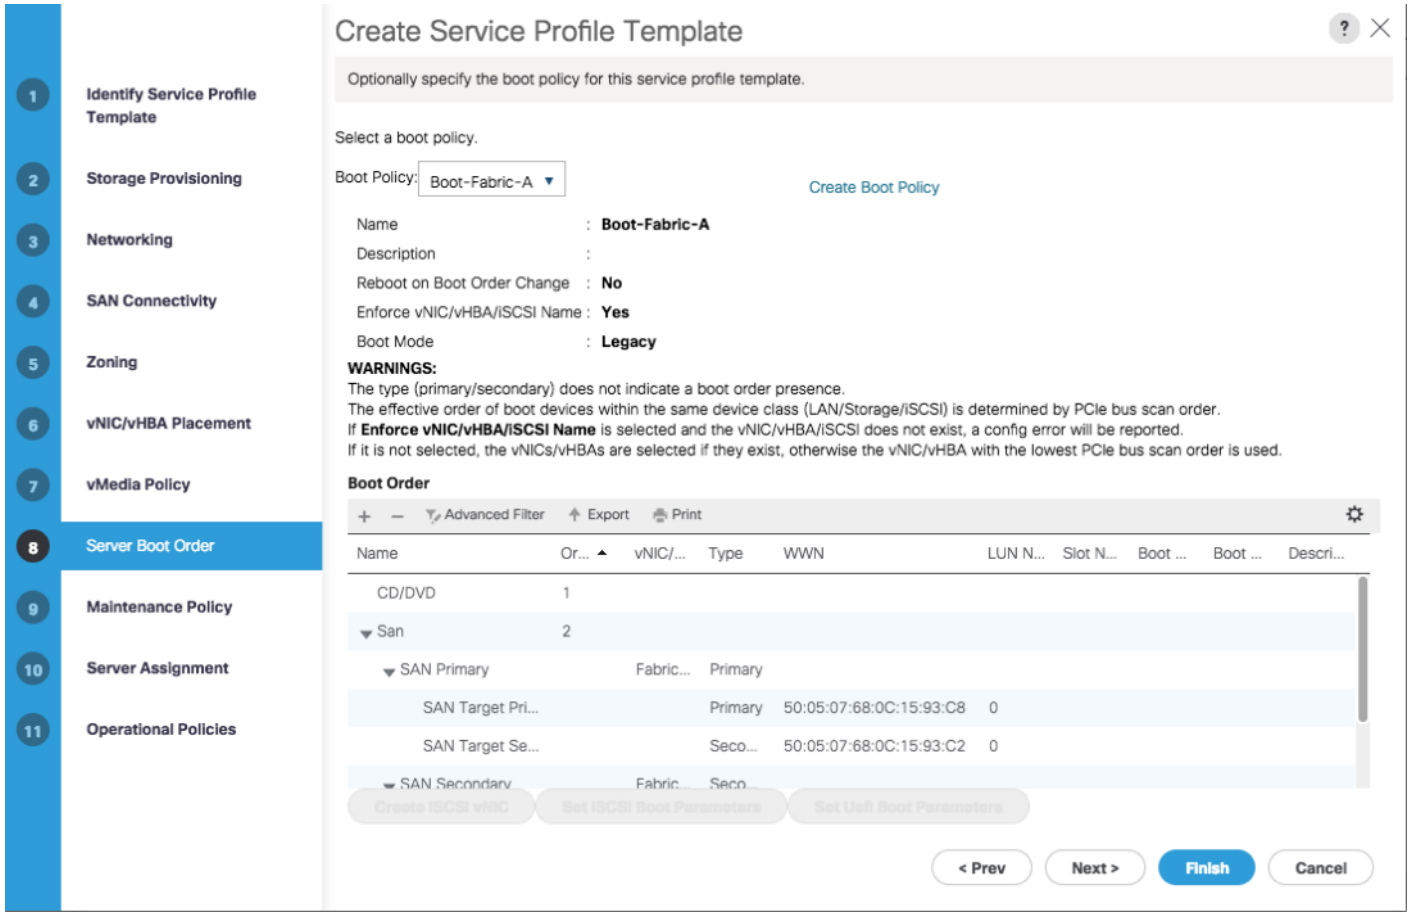

Configure Maintenance Policy

1. Choose the default Maintenance Policy.
2. Click Next.

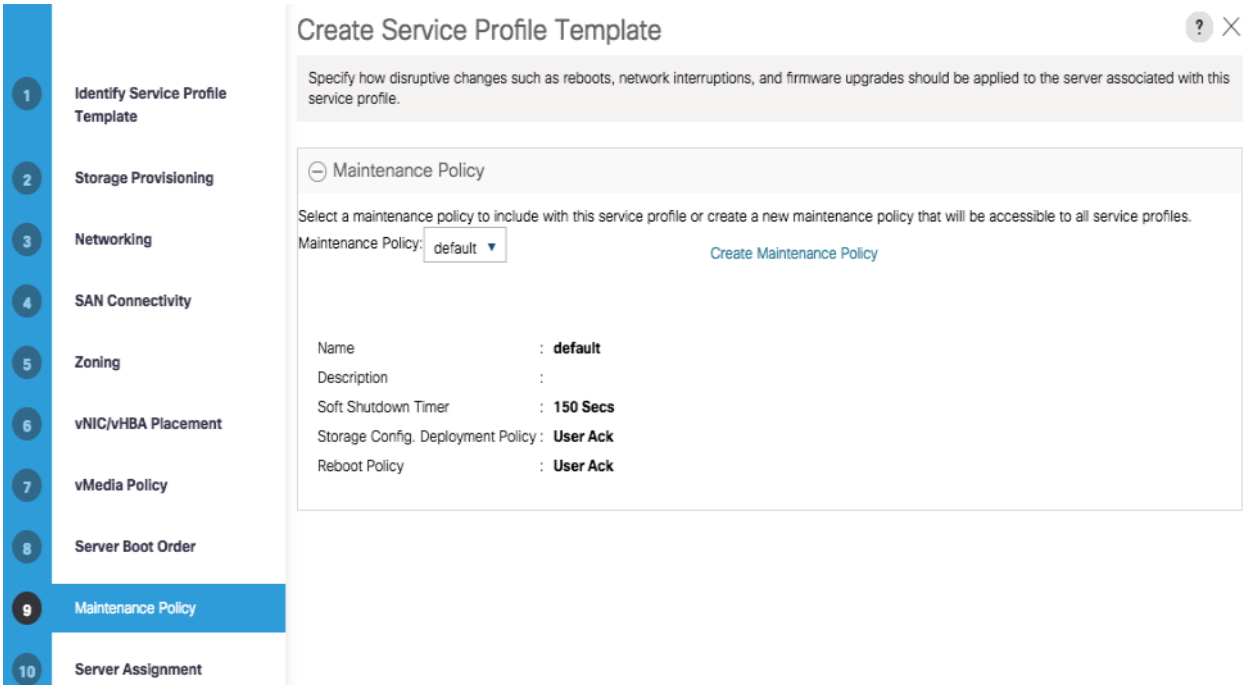

Configure Server Assignment

1. For the Pool Assignment field, select Infra-Server-Pool.
2. Optional: Select a Server Pool Qualification policy.
3. Select the option Up for the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management and select Infra-FW-Pack from the Host Firmware list.
5. Click Next.

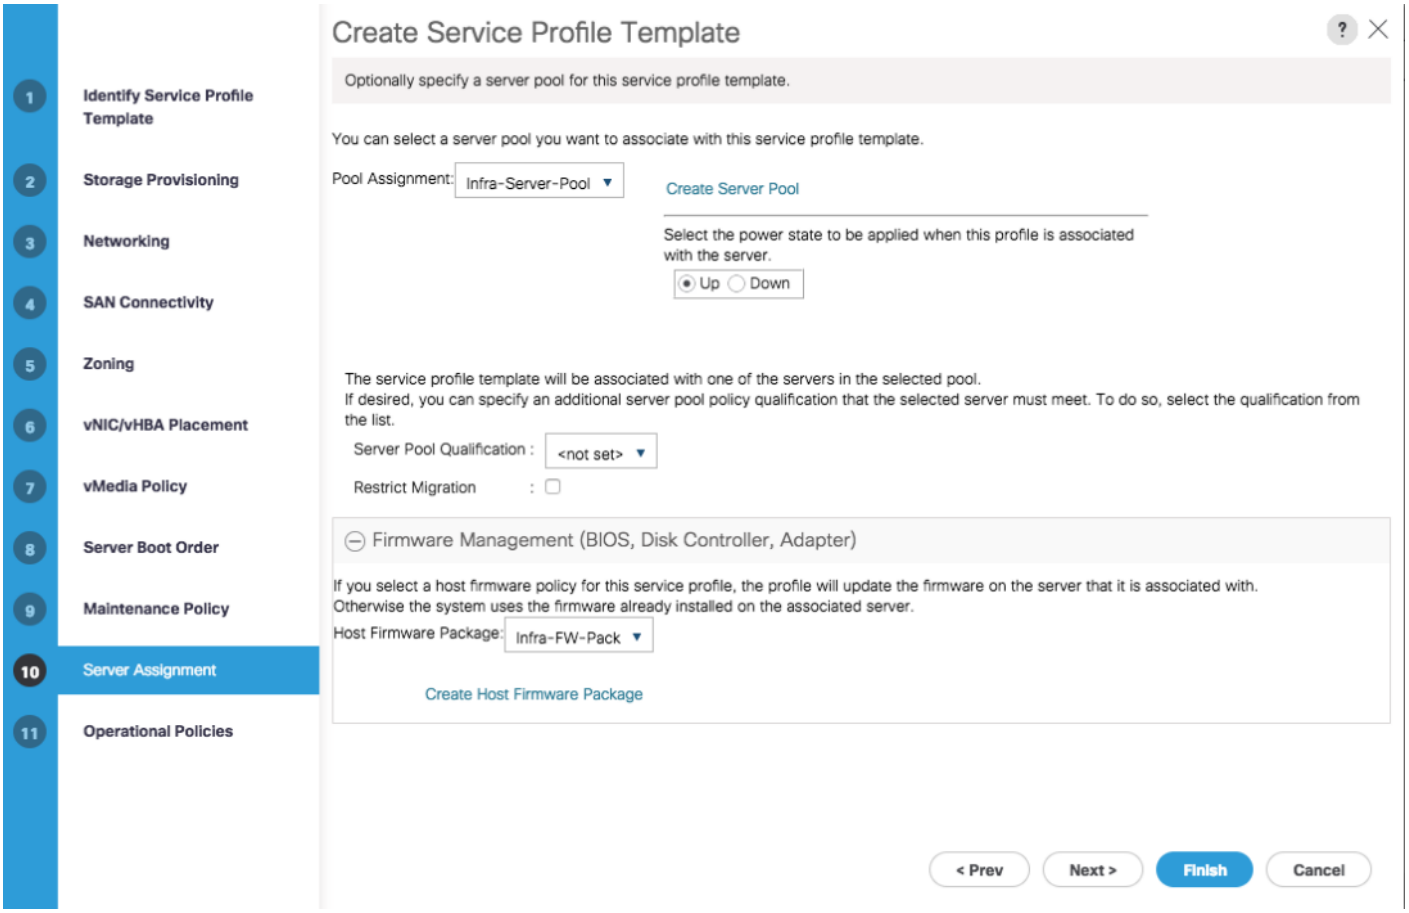

Configure Operational Policies

1. For the BIOS Policy field, select Infra-Host-BIOS.
2. Expand Power Control Policy Configuration and select No-Power-Cap for the Power Control Policy field.

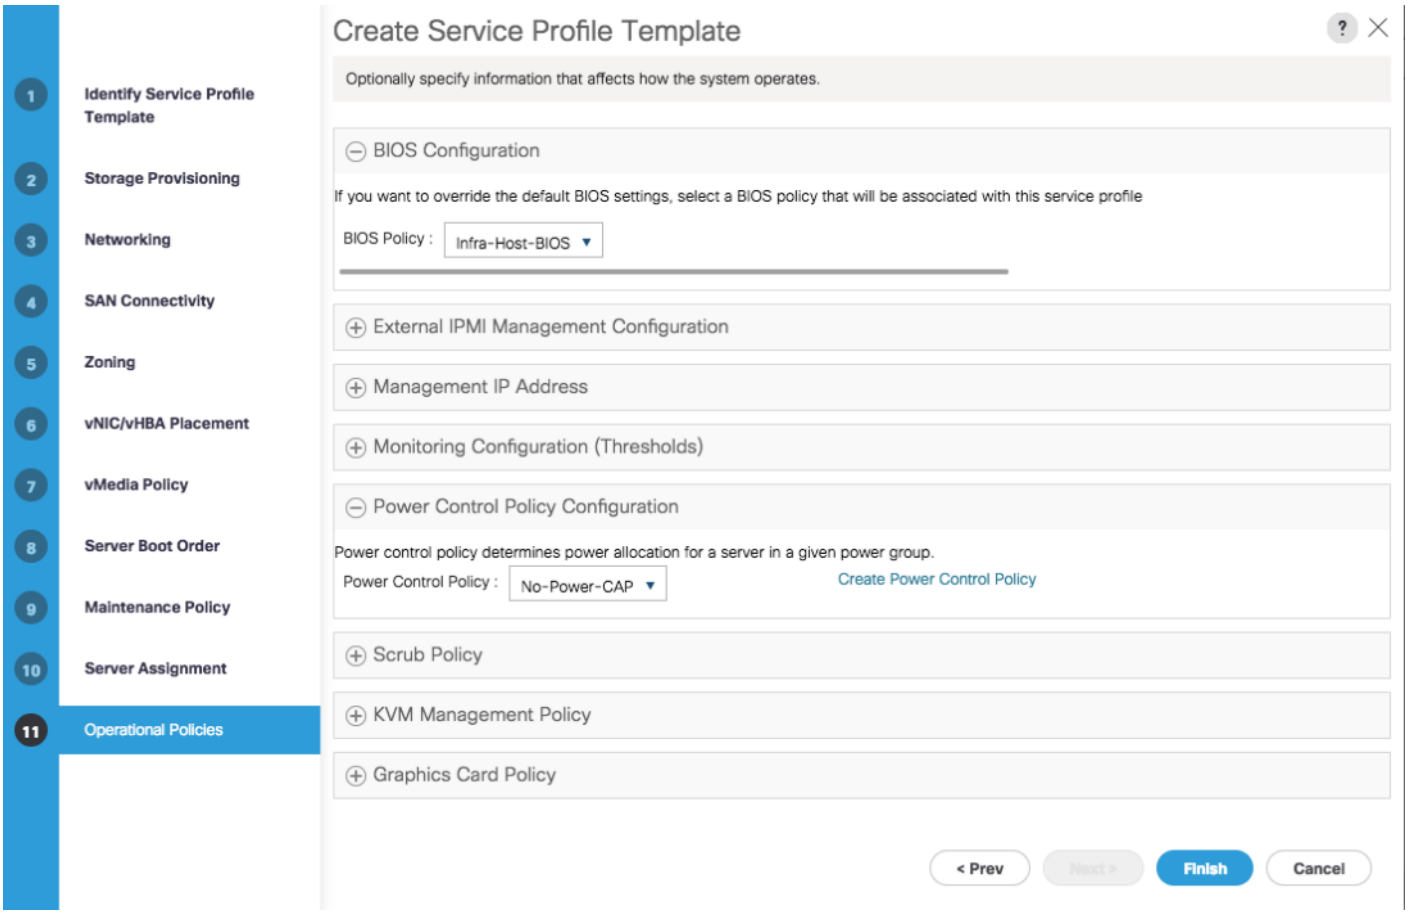

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create FC Boot Service Profiles

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Service Template Infra-ESXi-Host (Infra-ESXi-iSCSI-Host for iSCSI Deployment).
3. Right-click and choose Create Service Profiles from Template.

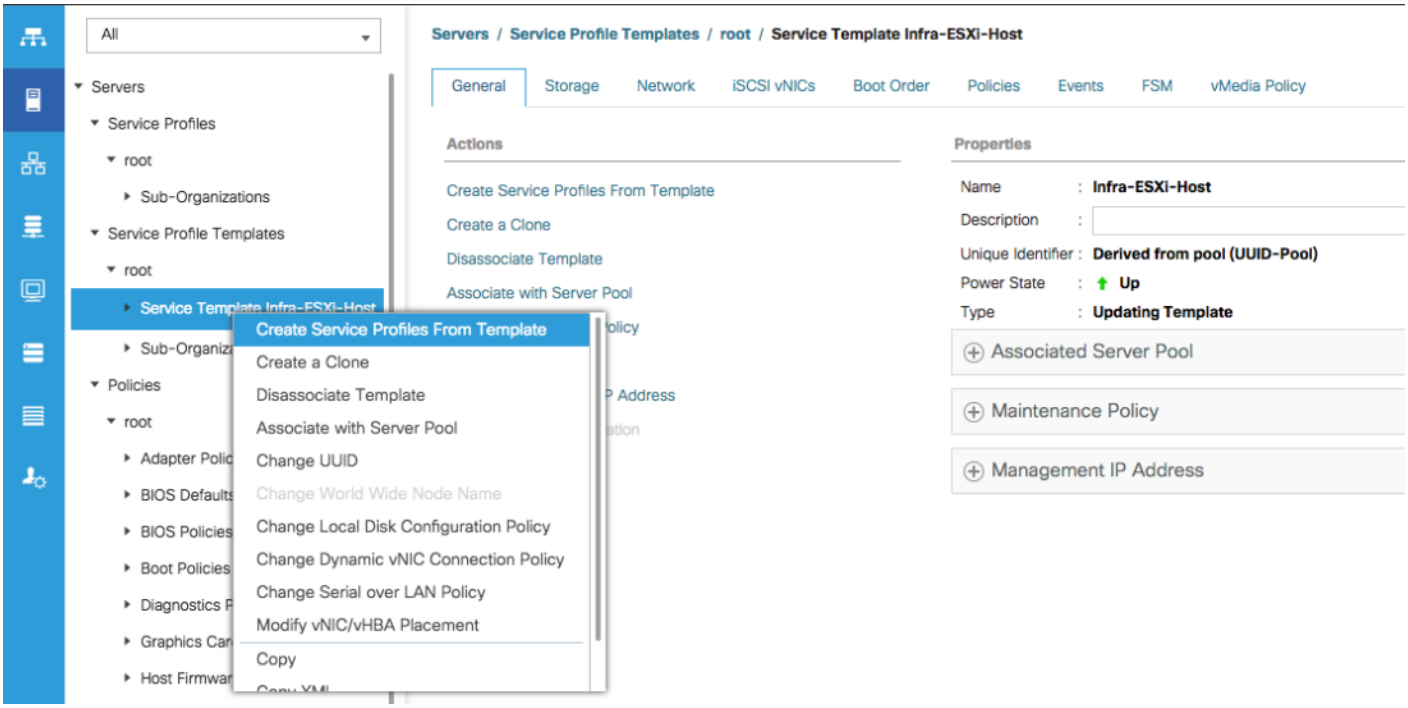

4. Enter Infra-ESXi-Host- as the service profile prefix.
5. Enter 1 as the Name Suffix Staring Number.
6. Enter the Number of servers to be deploy in the Number of Instances field.
7. Four service profiles were deployed during this validation
8. Click OK to create the service profile.

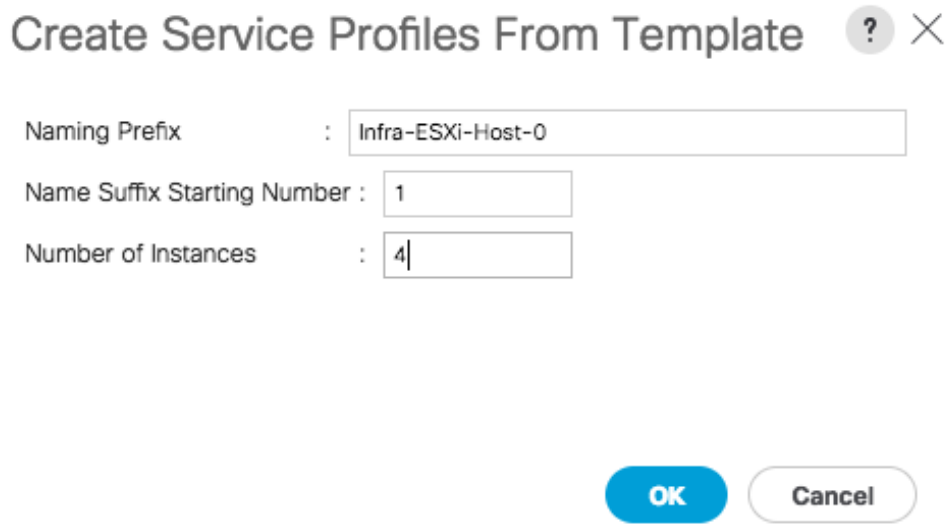

9. Click OK in the confirmation message.

10. Verify that the service profiles are successfully created and automatically associated with the servers from the pool.

Backup the Cisco UCS Manager Configuration

It is recommended to back up the Cisco UCS Configuration. Refer to the link below for additional information:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Admin-Management/3-2/b_Cisco_UCS_Admin_Mgmt_Guide_3_2.html

Adding Servers

Additional server pools, service profile templates, and service profiles can be created under root or in organizations under the root. All the policies at the root level can be shared among the organizations. Any new physical blades can be added to the existing or new server pools and associated with the existing or new service profile templates.

Gather Necessary WWPN Information (FC Deployment)

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will be assigned certain unique configuration parameters. To proceed with the SAN configuration, this deployment specific information must be gathered from each Cisco UCS blade. Complete the following steps:

1. To gather the vHBA WWPN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root. Select each service profile and expand to see the vHBAs.

2. Click vHBAs to see the WWPNs for both HBAs.

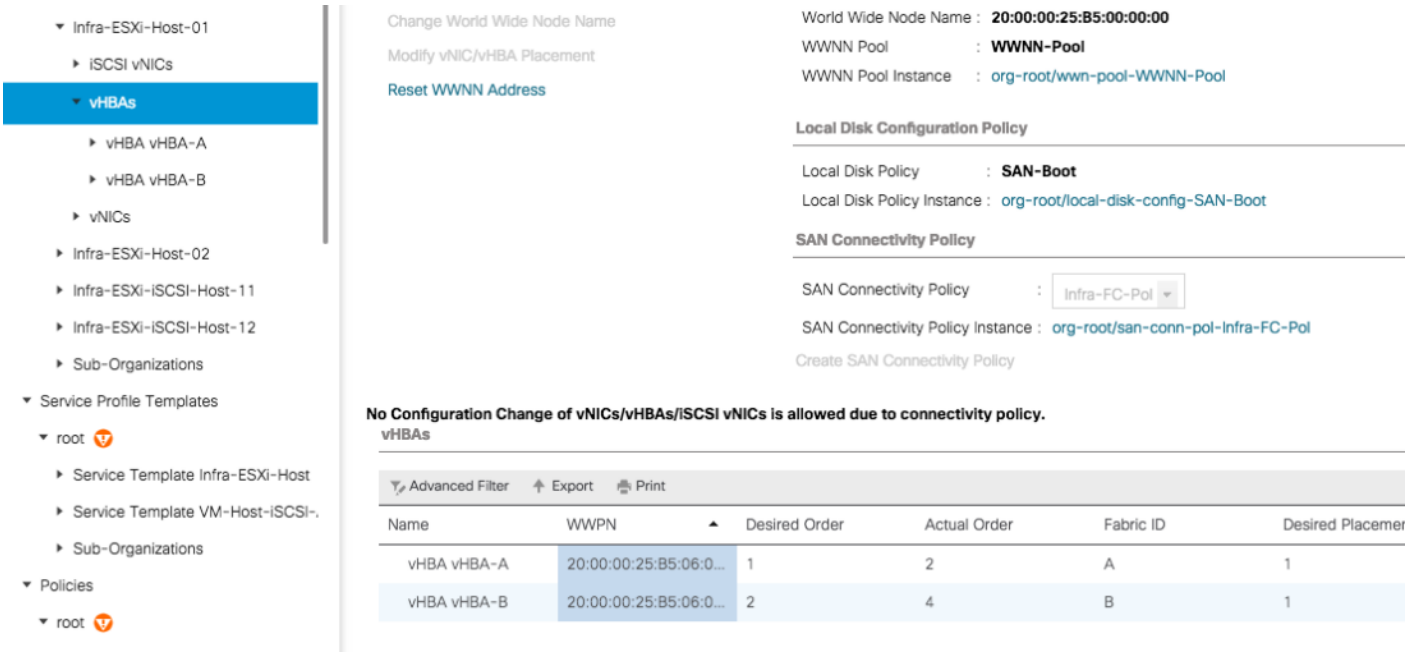

3. Record the WWPN information that is displayed for both the Fabric A vHBA and the Fabric B vHBA for each service profile into the WWPN variable in Table 12 . Add or remove rows from the table depending on the number of ESXi hosts.

Table 12 UCS WWPN Information

Host	vHBA		Value
Infra-ESXi-Host-1	Fabric-A	WWPN-Infra-ESXi-Host-1-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-1-B	20:00:00:25:b5:
Infra-ESXi-Host-2	Fabric-A	WWPN-Infra-ESXi-Host-2-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-2-B	20:00:00:25:b5:
Infra-ESXi-Host-3	Fabric-A	WWPN-Infra-ESXi-Host-3-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-3-B	20:00:00:25:b5:
Infra-ESXi-Host-4	Fabric-A	WWPN-Infra-ESXi-Host-4-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-4-B	20:00:00:25:b5:

Gather Necessary IQN Information (iSCSI Deployment)

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will be assigned certain unique configuration parameters. To proceed with the SAN configuration, this deployment specific information must be gathered from each Cisco UCS blade. Complete the following steps:

1. To gather the vNIC IQN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root.

2. Click each service profile and then click the “iSCSI vNICs” tab on the right. Note “Initiator Name” displayed at the top of the page under “Service Profile Initiator Name

Table 13 Cisco UCS iSCSI IQNs

Cisco UCS Service Profile Name	iSCSI IQN
Infra-ESXi-iSCSI-Host-01	iqn.1992-08.com.cisco:ucs-host____

Cisco UCS Service Profile Name	iSCSI IQN
Infra-ESXi-iSCSI-Host-02	iqn.1992-08.com.cisco:ucs-host_____
Infra-ESXi-iSCSI-Host-03	iqn.1992-08.com.cisco:ucs-host_____
Infra-ESXi–iSCSI-Host-04	iqn.1992-08.com.cisco:ucs-host_____

IBM SVC iSCSI Storage Configuration

IBM SVC iSCSI Configuration

As part of IBM SVC iSCSI configuration, complete the following steps:

- Setup Volumes
- Map Volumes to Hosts

Table 14 List of Volumes on iSCSI IBM SVC*

Volume Name	Capacity (GB)	Purpose	Mapping
Infra-ESXi-iSCSI-Host-01	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-01
Infra-ESXi-iSCSI-Host-02	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-02
Infra-ESXi-iSCSI-Host-03	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-03
Infra-ESXi-iSCSI-Host-04	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-04
Infra-iSCSI-datastore-1	1000**	Shared volume to host VMs	All ESXi hosts: Infra-ESXi-iSCSI-Host-01 to Infra-ESXi-iSCSI-Host-04
Infra-iSCSI-swap	300**	Shared volume to host VMware VM swap directory	All ESXi hosts: Infra-ESXi-iSCSI-Host-01 to Infra-ESXi-Host-04

* Customers should adjust the names and values used for server and volumes names based on their deployment

** The volume size can be adjusted based on customer requirements

Create Volumes on the Storage System

1. Log into the IBM SVC GUI and select the Volumes icon one the left screen and select Volumes

You will repeat the following steps to create and map all the volumes shown in Table 14 .

1. Click Create Volumes as shown below.

2. Click Basic and the select the pool (Bronze in this example) from the drop-down list.

3. Input quantity 1 and the capacity and name from Table 14 . Select Thin-provisioned for Capacity savings and enter the Name of the volume. Select I/O group io_grp0.

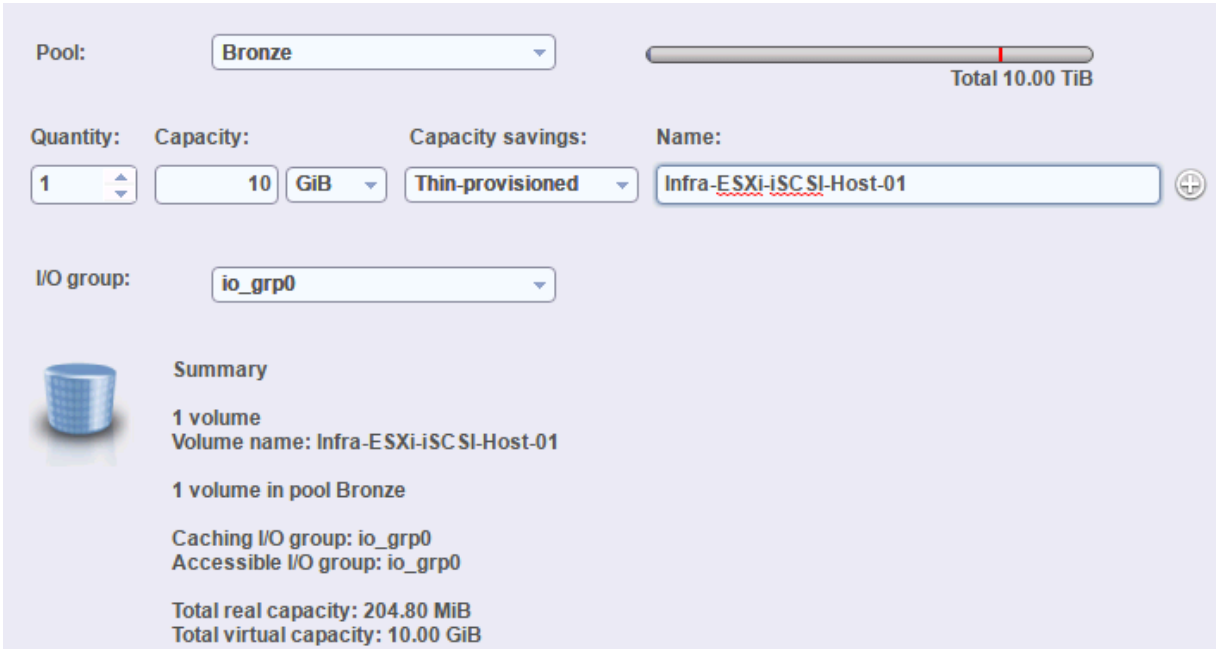

4. Click Create.

5. Repeat the steps above to create all the required volumes and verify all the volumes have successfully been created as shown in the sample output below.

Map Volumes to Hosts

1. Click Hosts.

To add all ESXi hosts (Table 13) to the IBM SVC system, complete the following steps:

1. Click Add Host.

2. Select iSCSI Host.
3. Add the name of the host to match the ESXi service profile name from Table 14 .
4. Type the IQN corresponding to the ESXi host from Table 13 and click Add.

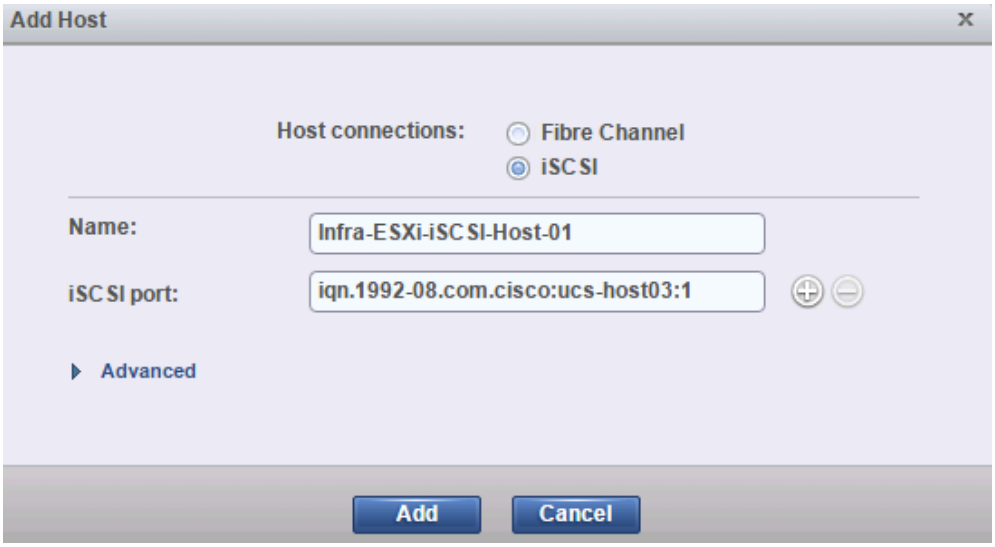

5. Click Close.

6. Click Volumes.

7. Right-click the Boot LUN for the ESXi host and choose Map to Host.

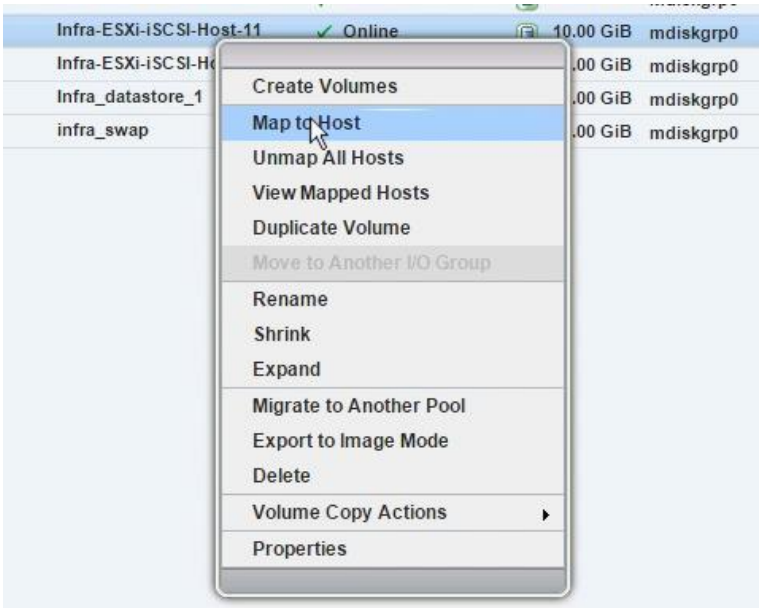

8. From the drop-down list, select the newly created iSCSI host.

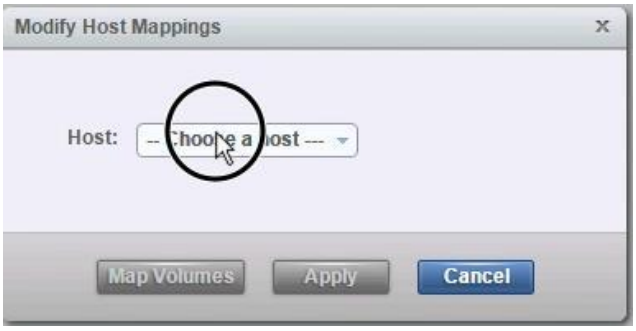

9. Click Map Volumes and when the process is complete, click Close.

10. Repeat the above steps to map shared volumes from Table 14 to the host as well.

11. Repeat the steps outlined in this procedure to add all the ESXi hosts to the storage system and modify their volume mappings to add both the boot LUN as well as the shared volumes to the host.

IBM SVC Fibre Channel Storage Configuration

As part of IBM SVC Fibre Channel configuration, complete the following steps:

- Setup Zoning on Cisco MDS switches
- Setup Volumes on IBM SVC
- Map Volumes to Hosts

Cisco MDS 9396S SAN Zoning for UCS Hosts

The following steps will configure zoning for the WWPNs for the UCS hosts and the IBM SVC nodes. WWPN information collected from the previous steps will be used in this section. Multiple zones will be created for servers in VSAN 101 on Switch A and VSAN 102 on Switch B.

The configuration below assumes 4 UCS services profiles have been deployed in this example. Customers can adjust the configuration according to their deployment size.

Cisco MDS - A Switch

Log in to the MDS switch and complete the following steps.

1. Configure the ports and the port-channel for UCS.

```
interface port-channel1 (For UCS)
  channel mode active
  switchport rate-mode dedicated
!
interface fc1/31 (UCS Fabric A)
  port-license acquire
  channel-group 1 force
  no shutdown
!
interface fc1/32 (UCS Fabric A)
  port-license acquire
  channel-group 1 force
  no shutdown
```

2. Create the VSAN.

```
vsan database
  vsan 101 interface port-channel1
```

3. The WWPNs recorded in Table 12 will be used in the next step. Replace the variables with actual WWPN values.

```
device-alias database
device-alias name Infra-ESXi-Host-01 pwwn <WWPN-Infra-ESXi-Host-1-A>
device-alias name Infra-ESXi-Host-02 pwwn <WWPN-Infra-ESXi-Host-2-A>
device-alias name Infra-ESXi-Host-03 pwwn <WWPN-Infra-ESXi-Host-3-A>
device-alias name Infra-ESXi-Host-04 pwwn <WWPN-Infra-ESXi-Host-4-A>
```

```
device-alias name SVC-Clus-Node1-FC1-NPIV pwwn <WWPN-SVC-Clus-Node1-FC1-NPIV>
device-alias name SVC-Clus-Node1-FC3-NPIV pwwn <WWPN-SVC-Clus-Node1-FC3-NPIV>
device-alias name SVC-Clus-Node2-FC1-NPIV pwwn <WWPN-SVC-Clus-Node2-FC1-NPIV>
device-alias name SVC-Clus-Node2-FC3-NPIV pwwn <WWPN-SVC-Clus-Node2-FC3-NPIV>
device-alias commit
```

4. Create the zones and add device-alias members for the 4 blades.

```
zone name Infra-ESXi-Host-01 vsan 101
member device-alias  Infra-ESXi-Host-01
member device-alias SVC-Clus-Node1-FC1-NPIV
member device-alias SVC-Clus-Node1-FC3-NPIV
member device-alias SVC-Clus-Node2-FC1-NPIV
member device-alias SVC-Clus-Node2-FC3-NPIV
!

zone name Infra-ESXi-Host-02 vsan 101
member device-alias  Infra-ESXi-Host-02
member device-alias SVC-Clus-Node1-FC1-NPIV
member device-alias SVC-Clus-Node1-FC3-NPIV
member device-alias SVC-Clus-Node2-FC1-NPIV
member device-alias SVC-Clus-Node2-FC3-NPIV
!

zone name Infra-ESXi-Host-03 vsan 101
member device-alias  Infra-ESXi-Host-03
member device-alias SVC-Clus-Node1-FC1-NPIV
member device-alias SVC-Clus-Node1-FC3-NPIV
member device-alias SVC-Clus-Node2-FC1-NPIV
member device-alias SVC-Clus-Node2-FC3-NPIV
!

zone name Infra-ESXi-Host-04 vsan 101
member device-alias  Infra-ESXi-Host-04
member device-alias SVC-Clus-Node1-FC1-NPIV
member device-alias SVC-Clus-Node1-FC3-NPIV
member device-alias SVC-Clus-Node2-FC1-NPIV
member device-alias SVC-Clus-Node2-FC3-NPIV
!
```

5. Add zones to zoneset.

```
zoneset name versastackzoneset vsan 101
    member Infra-ESXi-Host-01
    member Infra-ESXi-Host-02
    member Infra-ESXi-Host-03
    member Infra-ESXi-Host-04
```

6. Activate the zoneset.

```
zoneset activate  name versastackzoneset vsan 101
```

Validate all the HBA's are logged into the MDS switch. The SVC nodes and the Cisco servers should be powered on. To start the Cisco servers from Cisco UCS Manager, select the server tab, then click Servers>Service>Profiles>root, and right-click service profile then select boot server.

7. Validate that all the powered on system's HBAs are logged into the switch through the show zoneset command.

show zoneset active

```
MDS-9396S-A# show zoneset active

<SNIP>

  zone name Infra-ESXi-Host-01 vsan 101
  * fcid 0x400004 [pwwn 20:00:00:25:b5:00:0a:00] [Infra-ESXi-Host-01]

  * fcid 0x400241 [pwwn 50:05:07:68:0c:15:93:c8] [SVC-Clus-Node1-FC1-NPIV]

  * fcid 0x400261 [pwwn 50:05:07:68:0c:17:93:c8] [SVC-Clus-Node1-FC3-NPIV]

  * fcid 0x400281 [pwwn 50:05:07:68:0c:15:93:c2] [SVC-Clus-Node2-FC1-NPIV]

  * fcid 0x4002a1 [pwwn 50:05:07:68:0c:17:93:c2] [SVC-Clus-Node2-FC3-NPIV]


<SNIP>
```

8. Save the configuration.

```
 copy run start
```

Cisco MDS - B Switch

The configuration below assumes 4 UCS services profiles have been deployed. Customers can adjust the configuration according to their deployment.

Log in to the MDS switch and complete the following steps.

1. Configure the ports and the port channel for UCS.

```
interface port-channel2 (For UCS)
  channel mode active
  switchport rate-mode dedicated
!
interface fc1/31 (UCS Fabric B)
  port-license acquire
  channel-group 2 force
  no shutdown
!
interface fc1/32 (UCS Fabric B)
  port-license acquire
  channel-group 2 force
  no shutdown
```

2. Create the VSAN.

```
vsan database
  vsan 102 interface port-channel2
```

3. The WWPNs recorded in Table 12 will be used in the next step. Replace the variables with actual WWPN values.

```
device-alias database
device-alias name Infra-ESXi-Host-01 pwwn <WWPN-Infra-ESXi-Host-1-B>
device-alias name Infra-ESXi-Host-02 pwwn <WWPN-Infra-ESXi-Host-2-B>
device-alias name Infra-ESXi-Host-03 pwwn <WWPN-Infra-ESXi-Host-3-B>
device-alias name Infra-ESXi-Host-04 pwwn <WWPN-Infra-ESXi-Host-4-B>
device-alias name SVC-Clus-Node1-FC2-NPIV pwwn <WWPN-SVC-Clus-Node1-FC2-NPIV>
device-alias name SVC-Clus-Node1-FC4-NPIV pwwn <WWPN-SVC-Clus-Node1-FC4-NPIV>
device-alias name SVC-Clus-Node2-FC2-NPIV pwwn <WWPN-SVC-Clus-Node2-FC2-NPIV>
device-alias name SVC-Clus-Node2-FC4-NPIV pwwn <WWPN-SVC-Clus-Node2-FC4-NPIV>
device-alias commit
```

4. Create the zones and add device-alias members for the 4 blades.

```
zone name Infra-ESXi-Host-01 vsan 102
member device-alias  Infra-ESXi-Host-01
member device-alias SVC-Clus-Node1-FC2-NPIV
member device-alias SVC-Clus-Node1-FC4-NPIV
member device-alias SVC-Clus-Node2-FC2-NPIV
member device-alias SVC-Clus-Node2-FC4-NPIV
!

zone name Infra-ESXi-Host-02 vsan 102
member device-alias  Infra-ESXi-Host-02
member device-alias SVC-Clus-Node1-FC2-NPIV
member device-alias SVC-Clus-Node1-FC4-NPIV
member device-alias SVC-Clus-Node2-FC2-NPIV
member device-alias SVC-Clus-Node2-FC4-NPIV
!

zone name Infra-ESXi-Host-03 vsan 102
member device-alias  Infra-ESXi-Host-03
member device-alias SVC-Clus-Node1-FC2-NPIV
member device-alias SVC-Clus-Node1-FC4-NPIV
member device-alias SVC-Clus-Node2-FC2-NPIV
member device-alias SVC-Clus-Node2-FC4-NPIV
!
zone name Infra-ESXi-Host-04 vsan 102
member device-alias  Infra-ESXi-Host-04
member device-alias SVC-Clus-Node1-FC2-NPIV
member device-alias SVC-Clus-Node1-FC4-NPIV
member device-alias SVC-Clus-Node2-FC2-NPIV
member device-alias SVC-Clus-Node2-FC4-NPIV
!
```

5. Add zones to zoneset.

```
zoneset name versastackzoneset vsan 102
    member Infra-ESXi-Host-01
    member Infra-ESXi-Host-02
    member Infra-ESXi-Host-03
    member Infra-ESXi-Host-04
```

6. Activate the zoneset.

```
zoneset activate  name versastackzoneset vsan 102
```

Validate all the HBA's are logged into the MDS switch. The SVC nodes and the Cisco servers should be powered on. To start the Cisco servers from Cisco UCS Manager, select the server tab, then click Servers→Service→Profiles→root, and right-click service profile then select boot server.

7. Validate that all the powered on system's HBAs are logged into the switch through the show zoneset command.

```
show zoneset active

MDS-9396S-B# show zoneset active

<SNIP>

  zone name Infra-ESXi-Host-01 vsan 102
  * fcid 0x770002 [pwwn 20:00:00:25:b5:00:0b:00] [Infra-ESXi-Host-01]

  * fcid 0x770241 [pwwn 50:05:07:68:0c:16:93:c8] [SVC-Clus-Node1-FC2-NPIV]

  * fcid 0x770261 [pwwn 50:05:07:68:0c:18:93:c8] [SVC-Clus-Node1-FC4-NPIV]

  * fcid 0x770281 [pwwn 50:05:07:68:0c:16:93:c2] [SVC-Clus-Node2-FC2-NPIV]

  * fcid 0x7702a1 [pwwn 50:05:07:68:0c:18:93:c2] [SVC-Clus-Node2-FC4-NPIV]

<SNIP>
```

8. Save the configuration.

```
copy run start
```

IBM SVC FC Configuration

As part of IBM SVC FC configuration, complete the following steps:

- Create ESXi boot Volumes (Boot LUNs for all the ESXi hosts)
- Create Share Storage Volumes (for hosting VMs)
- Map Volumes to Hosts

In this deployment example, there are four ESXi hosts – following volumes will be created in this process:

Table 15 List of FC volumes on IBM SVC*

Volume Name	Capacity (GB)	Purpose	Mapping
Infra-ESXi-Host-01	10	Boot LUN for the Host	Infra-ESXi-Host-01
Infra-ESXi-Host-02	10	Boot LUN for the Host	Infra-ESXi-Host-02
Infra-ESXi-Host-03	10	Boot LUN for the Host	Infra-ESXi-Host-03
Infra-ESXi-Host-04	10	Boot LUN for the Host	Infra-ESXi-Host-04

Volume Name	Capacity (GB)	Purpose	Mapping
Infra-datastore-1	1000*	Shared volume to host VMs	All ESXi hosts: Infra-ESXi-Host-01 to Infra-ESXi-Host-04
Infra-swap	300*	Shared volume to host VMware VM swap directory	All ESXi hosts: Infra-ESXi-Host-01 to Infra-ESXi-Host-04

* Customers should adjust the names and values based on their environment.

Create Volumes on the Storage System

1. Log into the IBM SVC GUI and select the Volumes icon one the left screen and select Volumes

You will repeat the following steps to create and map all the volumes shown in Table 15 .

2. Click Create Volumes as shown in below.

3. Click Basic and the select the pool (Bronze in this example) from the drop-down list.

4. Input quantity 1 and the capacity and name from Table 15 . Select Thin-provisioned for Capacity savings and enter the Name of the volume. Select io_grp0 for the I/O group.

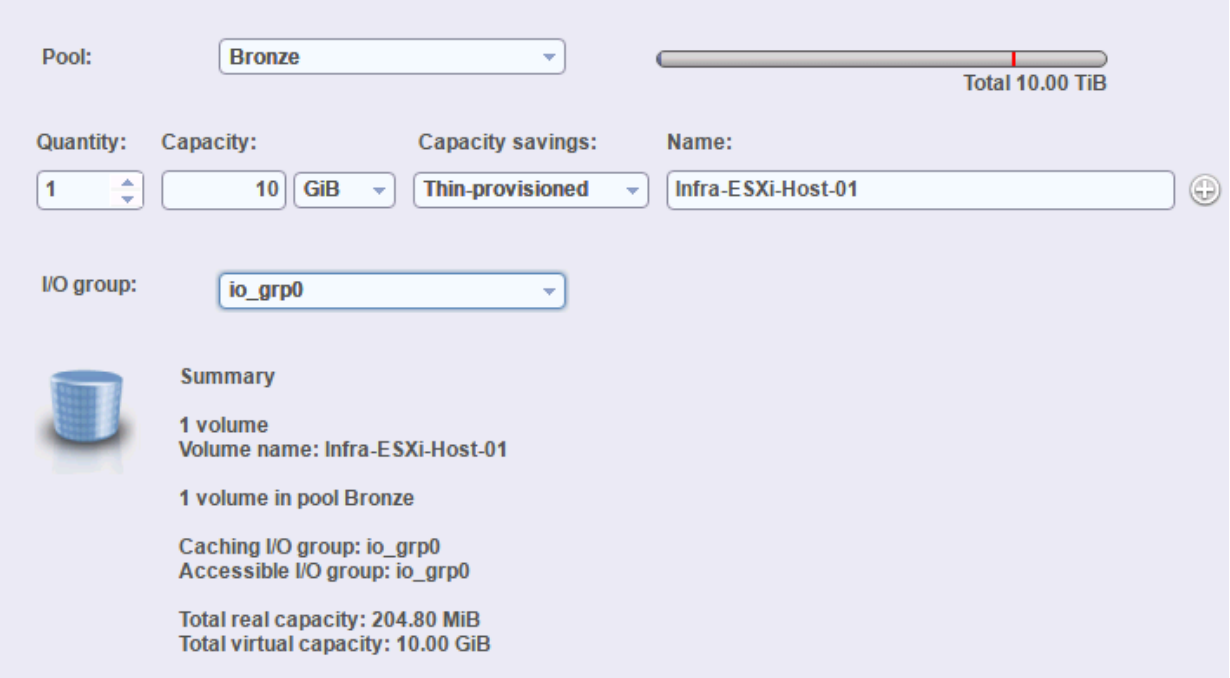

5. Click Create.

6. Repeat the steps above to create all the required volumes and verify all the volumes have successfully been created as shown in the sample output below.

Name	State	Capacity	Pool	Host Mappings	UID
Infra-ESXi-Host-01	Online	10.00 GiB	mdiskgrp0	No	600507640090800588000000000000004
Infra-ESXi-Host-02	Online	10.00 GiB	mdiskgrp0	No	600507640090800588000000000000005
Infra-ESXi-Host-03	Online	10.00 GiB	mdiskgrp0	No	600507640090800588000000000000006
Infra-ESXi-Host-04	Online	10.00 GiB	mdiskgrp0	No	600507640090800588000000000000007
Infra_datastore_1	Online	500.00 GiB	mdiskgrp0	No	600507640090800588000000000000002
infra_swap	Online	100.00 GiB	mdiskgrp0	No	600507640090800588000000000000003

Map Volumes to Hosts

1. Select Host Icon in the left pane and click Hosts.

2. Follow the procedure below to add all ESXi hosts (Table 15) to the IBM SVC system.
3. Click Add Host.

4. Select the Fibre Channel Host.
5. Add the name of the host to match the ESXi service profile name from Table 15 .
6. From the drop-down list, select both (Fabric A and B) WWPNs corresponding to the host in Table 15 .

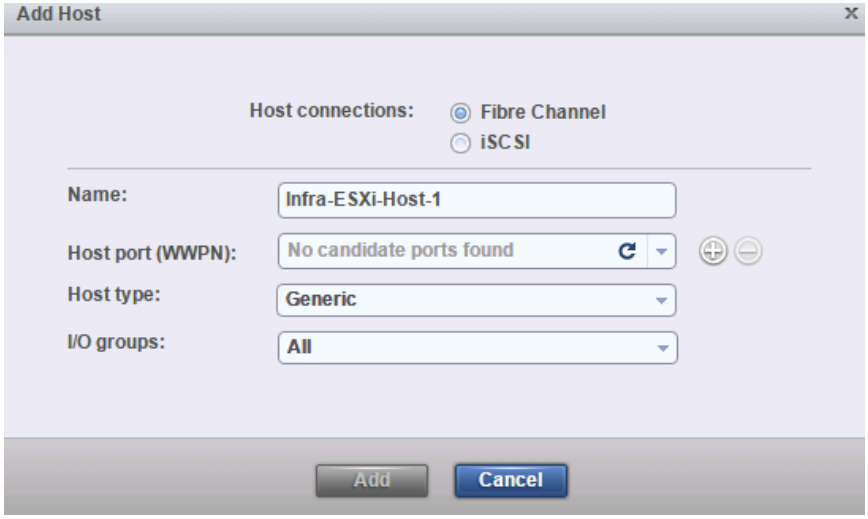

7. Select Host Type Generic and I/O groups All.

8. Click Add.

9. Right-click on the newly created host and select Modify VOLUME Mappings...

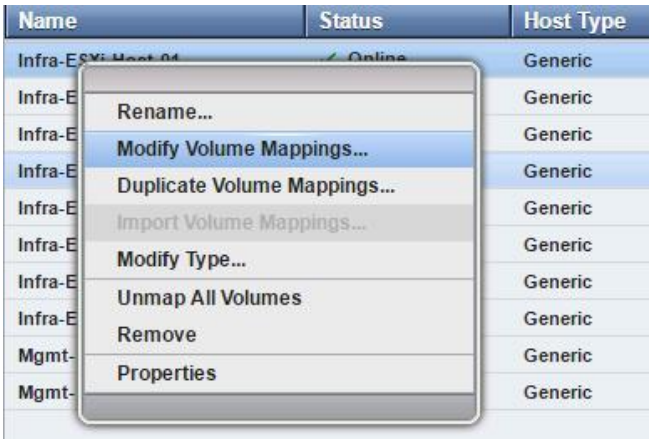

10. Select the Boot LUN corresponding to the host and the shared volumes to the column on the right labelled as Volume Mapped to the Host.

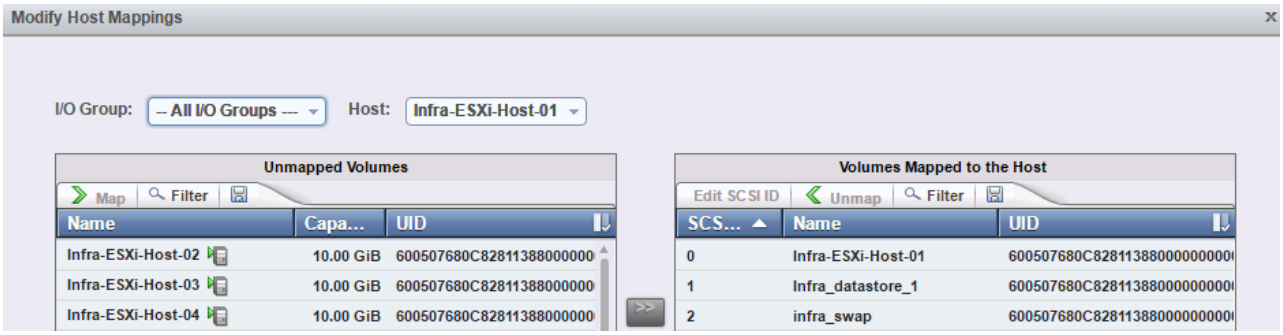

11. Click Map Volumes. Once the process is complete, the Host Mappings column should show Yes as shown in the screenshot below.

12. Repeat the steps above to add all the ESXi hosts in the environment and modify their mappings.

VMware vSphere Setup for UCS Host Environment

VMware ESXi 6.5 U1

This section provides detailed instructions for installing VMware ESXi 6.5 U1 in the VersaStack UCS environment. After the procedures are completed, multiple ESXi hosts will be provisioned to host customer workloads.

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).

Log in to Cisco UCS Manager

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

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

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Under HTML, click the Launch UCS Manager link.
3. When prompted, enter admin as the user name and enter the administrative password.
4. To log in to Cisco UCS Manager, click Login.
5. From the main menu, click the Servers tab.
6. Select Servers > Service Profiles > root > Infra-ESXi-Host-01.

For iSCSI setup, the name of the profile will be Infra-ESXi-iSCSI-Host-01

7. Right-click Infra-ESXi-Host-01 and select KVM Console.
8. If prompted to accept an Unencrypted KVM session, accept as necessary.
9. Open KVM connection to all the hosts by right-clicking the Service Profile and launching the KVM console
10. Boot each server by selecting Boot Server and clicking OK. Then click OK again.

Install ESXi on the UCS Servers

To install VMware ESXi to the boot LUN of the hosts, complete the following steps on each host. The Cisco customer VMware ESXi image can be downloaded from:

https://my.vmware.com/group/vmware/details?downloadGroup=OEM-ESXI65U1-CISCO&productId=614

1. In the KVM windows, click Virtual Media in the upper right of the screen.
2. Click Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.
8. Reset the server by clicking Reset button. Click OK.
9. Select Power Cycle on the next window and click OK and OK again.
10. On reboot, the machine detects the presence of the boot LUNs (sample output below).
11. From the ESXi Boot Menu, select the ESXi installer.

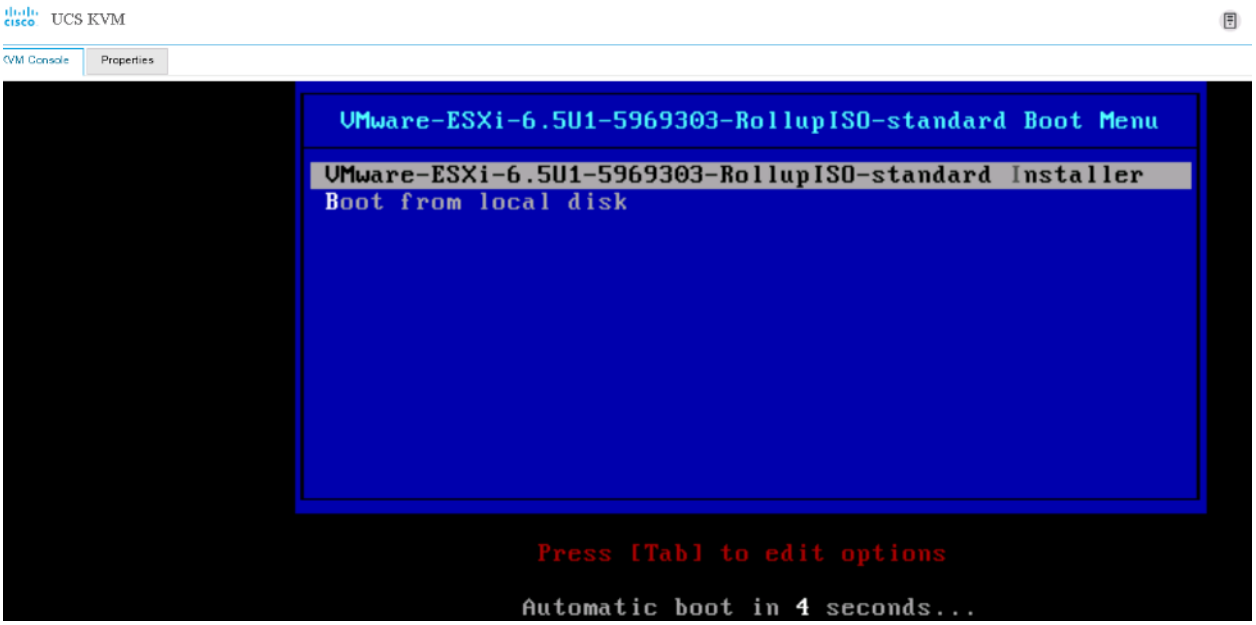

12. After the installer has finished loading, press Enter to continue with the installation.
13. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
14. Select the LUN that was previously set up and discovered as the installation disk for ESXi and press Enter to continue with the installation.
15. Select the appropriate keyboard layout and press Enter.

16. Enter and confirm the root password and press Enter.
17. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
18. After the installation is complete, press Enter to reboot the server.
19. Repeat the ESXi installation process for all the Service Profiles.

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, complete the following steps on each ESXi host.

To configure the ESXi hosts with access to the management network, complete the following steps:

1. After the server has finished post-installation rebooting, press F2 to customize the system.
2. Log in as root, enter the password chosen during the initial setup, and press Enter to log in.
3. Select the Configure Management Network option and press Enter.
4. Select Network Adapters
5. Select vmnic0 (if it is not already selected) by pressing the Space Bar.

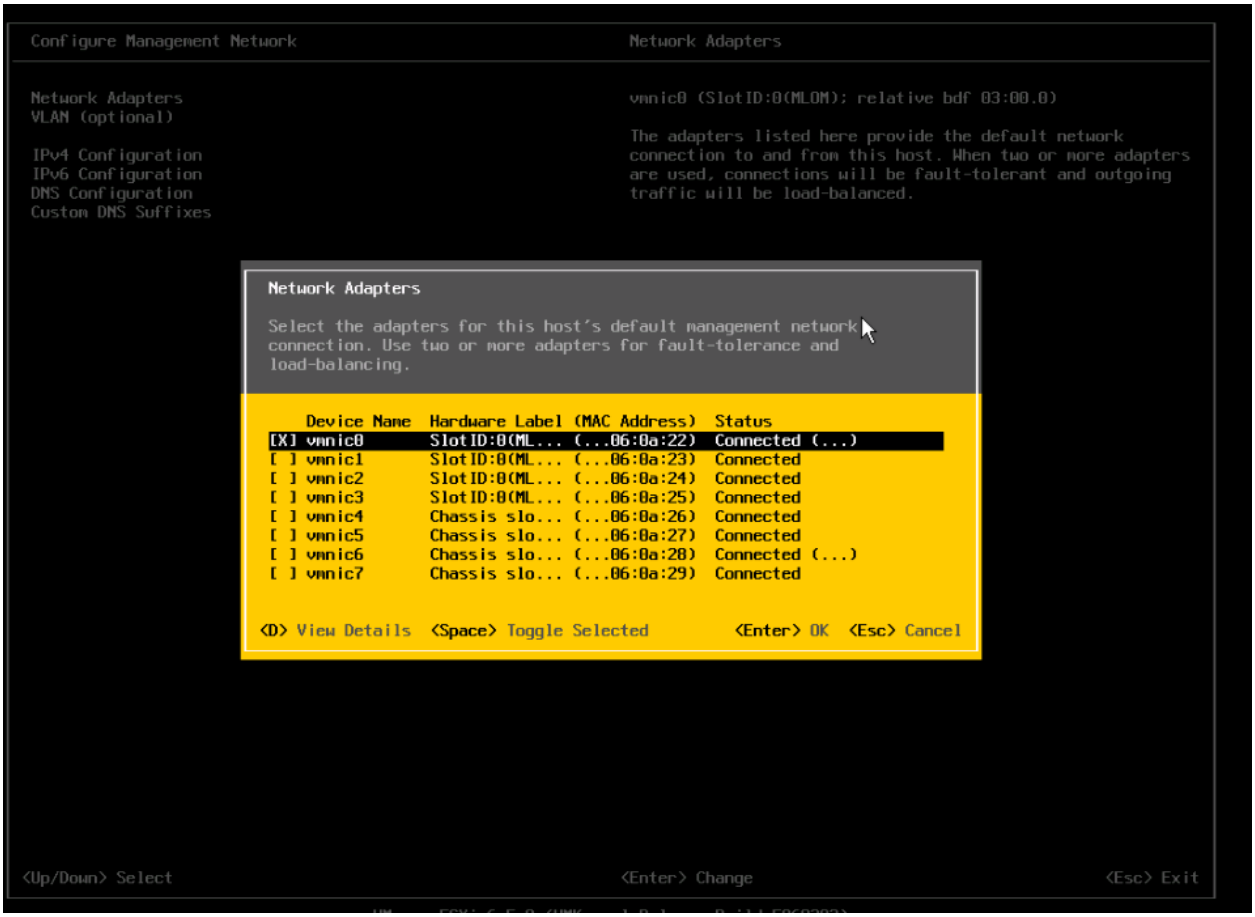

6. Press Enter to save and exit the Network Adapters window.
7. Select the VLAN (Optional) and press Enter.
8. Enter the <IB Mgmt VLAN> (11) and press Enter.

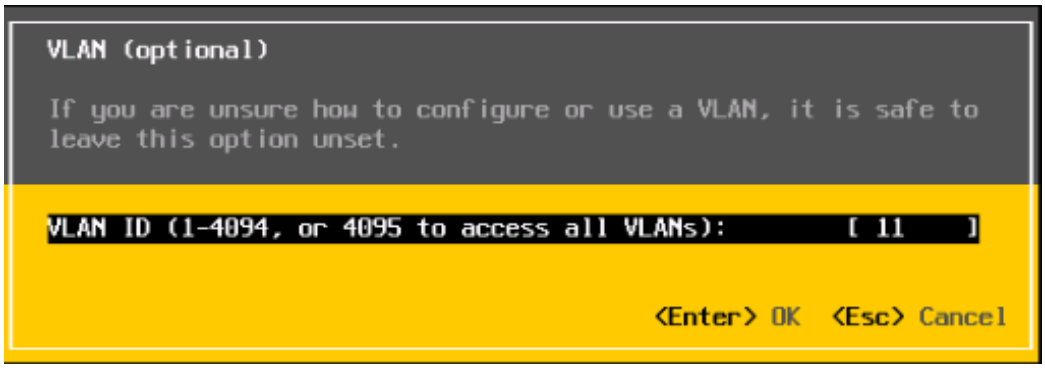

9. Select IPv4 Configuration and press Enter.
10. Select the Set Static IP Address and Network Configuration option by using the Space Bar.
11. Enter the IP address for managing the ESXi host.

12. Enter the subnet mask for the management network of the ESXi host.

13. Enter the default gateway for the ESXi host.

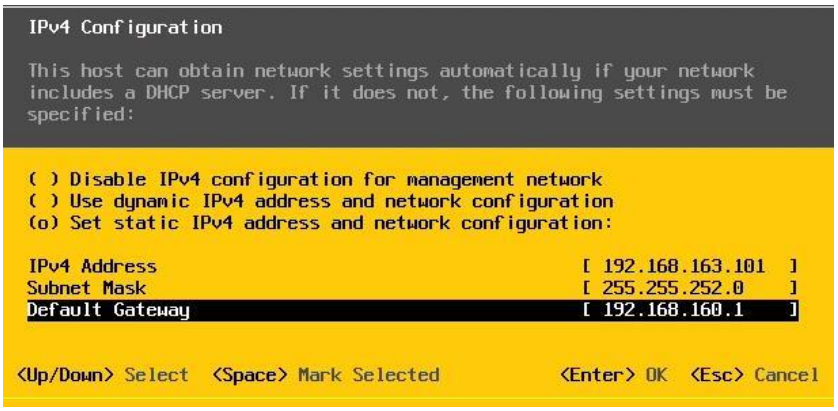

14. Press Enter to accept the changes to the IP configuration.

15. Select the IPv6 Configuration option and press Enter.

16. Using the Space Bar, select Disable IPv6 (restart required) and press Enter.

17. Select the DNS Configuration option and press Enter.

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

18. Enter the IP address of 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 ESXi host.

21. Press Enter to accept the changes to the DNS configuration.

22. Press Esc to exit the Configure Management Network submenu.

23. Press Y to confirm the changes and reboot the host.

24. Repeat this procedure for all the ESXi hosts in the setup.

VMware vSphere Configuration

The vSphere configuration covered in this section is common to all the ESXi servers. In the procedure below, two shared datastores, one for hosting the VMs and another to host the VM swap files, will be mounted to all the ESXi servers. Customers can adjust the number and size of the shared datastores based on their particular deployments.

Download VMware vSphere Client

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

1. Open a web browser on the management workstation and navigate to the management IP address an any ESXi servers.
2. Download and install the vSphere Client for Windows.

Log in to VMware ESXi Hosts Using VMware vSphere Client

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

1. Open the recently downloaded VMware vSphere Client and enter the management IP address of the host.
2. Enter root for the user name.
3. Enter the root password configured during the installation process.
4. Click Login to connect.
5. Repeat this process to log into all the ESXi hosts.

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

For the most recent versions, please refer to Cisco UCS HW and SW Availability Interoperability Matrix. If a more recent driver is made available that is appropriate for VMware vSphere 6.5 U1, download and install the latest drivers.

To install VMware VIC Drivers on the ESXi hosts using esxcli, complete the following steps:

1. Download and extract the following VIC Drivers to the Management workstation:

FNIC Driver version 1.6.0.37:

https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI60-CISCO-FNIC-16037&productId=718

NENIC Driver version 1.0.16.0:

https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESXI65-CISCO-NENIC-10160&productId=614

To install VIC Drivers on ALL the ESXi hosts, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. From Resources > Storage, right-click datastore1 and select Browse Datastore.

4. Click the fourth button and select Upload File.
5. Navigate to the saved location for the downloaded VIC drivers and select fnic_driver_1.6.0.37-offline_bundle-7765239.zip.
6. Click Open and Yes to upload the file to datastore1.
7. Click the fourth button and select Upload File.
8. Navigate to the saved location for the downloaded VIC drivers and select VMW-ESX-6.5.0-nenic-1.0.16.0-offline_bundle-7643104.zip.
9. Click Open and Yes to upload the file to datastore1.
10. Make sure the files have been uploaded to both ESXi hosts.
11. In the ESXi host vSphere Client, select the Configuration tab.
12. In the Software pane, select Security Profile.
13. To the right of Services, click Properties.
14. Select SSH and click Options.
15. Click Start and OK.

The step above does not enable SSH service and the service will not be restarted when ESXi host reboots.

16. Click OK to close the window.
17. Make sure SSH is started on each host.
18. From the management workstation, start an ssh session to each ESXi host. Login as root with the root password.
19. At the command prompt, run the following commands to account for each host:

```
esxcli software vib update -d /vmfs/volumes/datastore1/fnic_driver_1.6.0.37-offline_bundle-7765239.zip

esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.5.0-nenic-1.0.16.0-offline_bundle-7643104.zip

reboot
```

20. After each host has rebooted, log back into each host with vSphere Client.

Mount Required Datastores

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

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Storage in the Hardware window.
4. From the Datastore area, click Add Storage to open the Add Storage wizard.
5. Select Disk/LUN and click Next.
6. Verifying by using the size of the datastore LUN, select the LUN configured for VM hosting and click Next.
7. Accept default VMFS setting and click Next.
8. Click Next for the disk layout.
9. Enter infra-datastore-1 as the datastore name.
10. Click Next to retain maximum available space.
11. Click Finish.
12. Select the second LUN configured for swap file location and click Next.
13. Accept default VMFS setting and click Next.
14. Click Next for the disk layout.
15. Enter infra_swap as the datastore name.
16. Click Next to retain maximum available space.
17. Click Finish.
18. The storage configuration should look similar to figure shown below.
19. Repeat these steps on all the ESXi hosts.

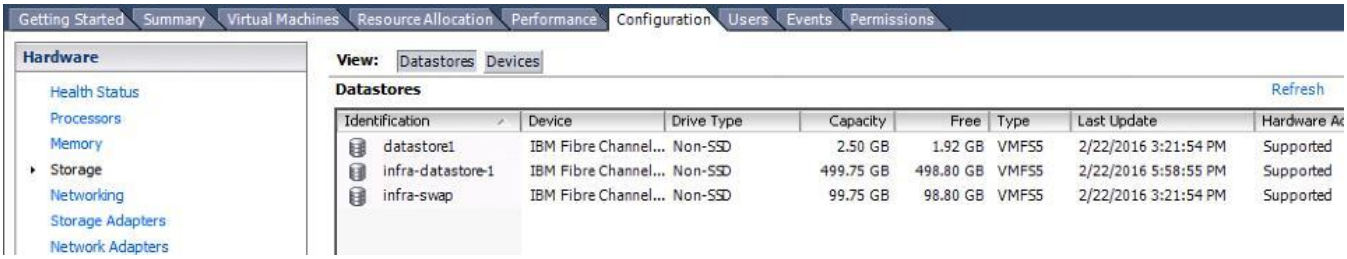

Configure NTP on ESXi Hosts

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

2. From the vSphere Client, select the host in the inventory.
3. Click the Configuration tab.
4. Click Time Configuration in the Software pane.
5. Click Properties.
6. At the bottom of the Time Configuration dialog box, click NTP Client Enabled.
7. At the bottom of the Time Configuration dialog box, click Options.
8. In the NTP Daemon (ntpd) Options dialog box, complete the following steps:
9. Click General in the left pane and select Start and stop with host.
10. Click NTP Settings in the left pane and click Add.
11. In the Add NTP Server dialog box, enter <NTP Server IP Address> as the IP address of the NTP server and click OK.
12. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes checkbox and click OK.
13. Click OK.
14. In the Time Configuration dialog box, verify that the clock is now set to approximately the correct time.

Move VM Swap File Location

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

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Virtual Machine Swapfile Location in the Software pane.
4. Click Edit at the upper-right side of the window.
5. Select the option Store the swapfile in a swapfile datastore selected below.
6. Select the infra-swap datastore to house the swap files.
7. Click OK to finalize the swap file location.

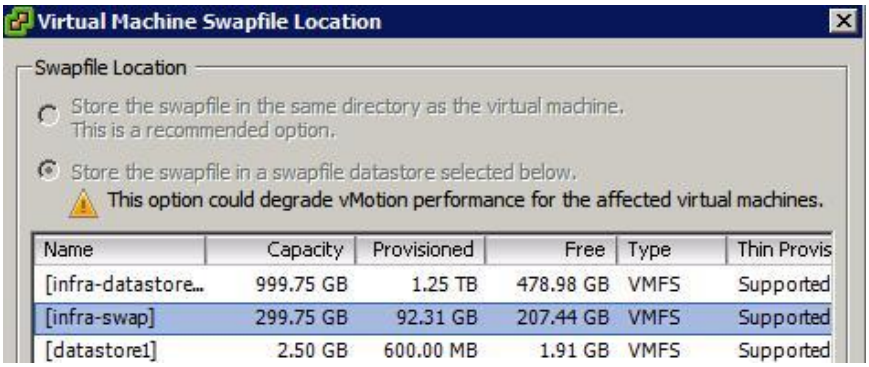

Deploy VMware vCenter Appliance 6.5

This guide recommends following the VMware vSphere documentation to deploy and configure VMware vCenter Appliance.

Set Up vCenter Server

1. Using a web browser, navigate to <vCenter IP Address>.

2. Click the link Log in to vSphere Web Client.

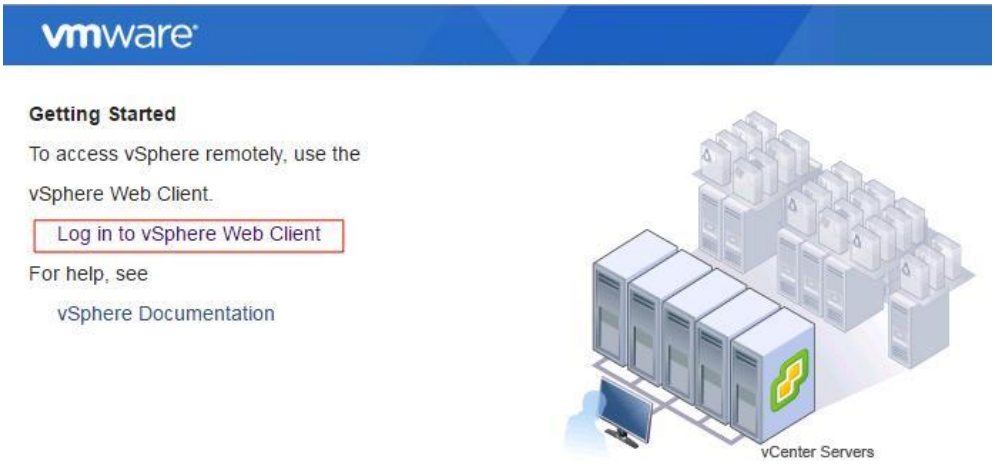

If prompted, run and install the VMWare Remote Console Plug-in.

3. Log in as root, with the root password entered during the vCenter installation.

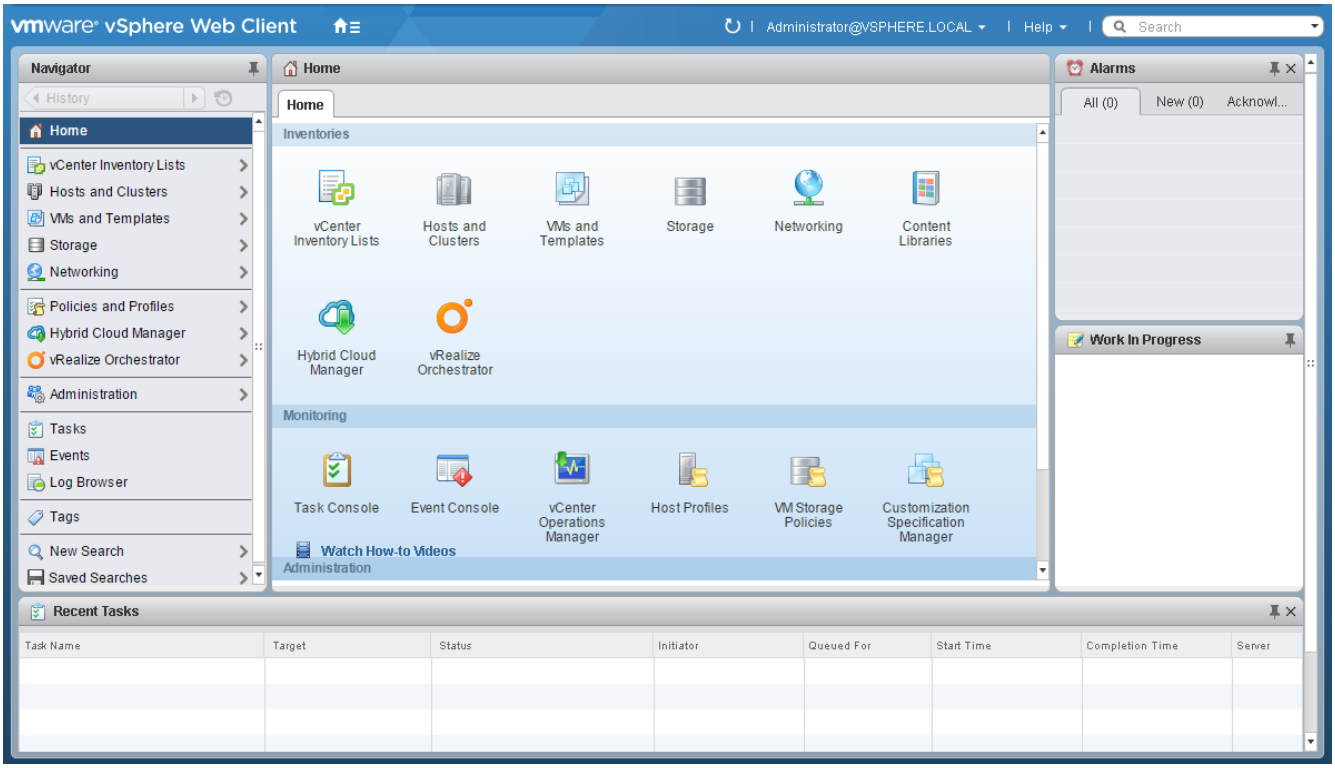

Setup Datacenter, Cluster, DRS and HA for ESXi Nodes

If a new datacenter is needed for the VersaStack, complete the following steps on the vCenter:

1. Connect to the vSphere Web Client and click Hosts and Clusters from the left side Navigator window or the Hosts and Clusters icon from the Home center window

2. From Hosts and Clusters:

3. Right-click the vCenter icon and from the drop-down list select New Datacenter.

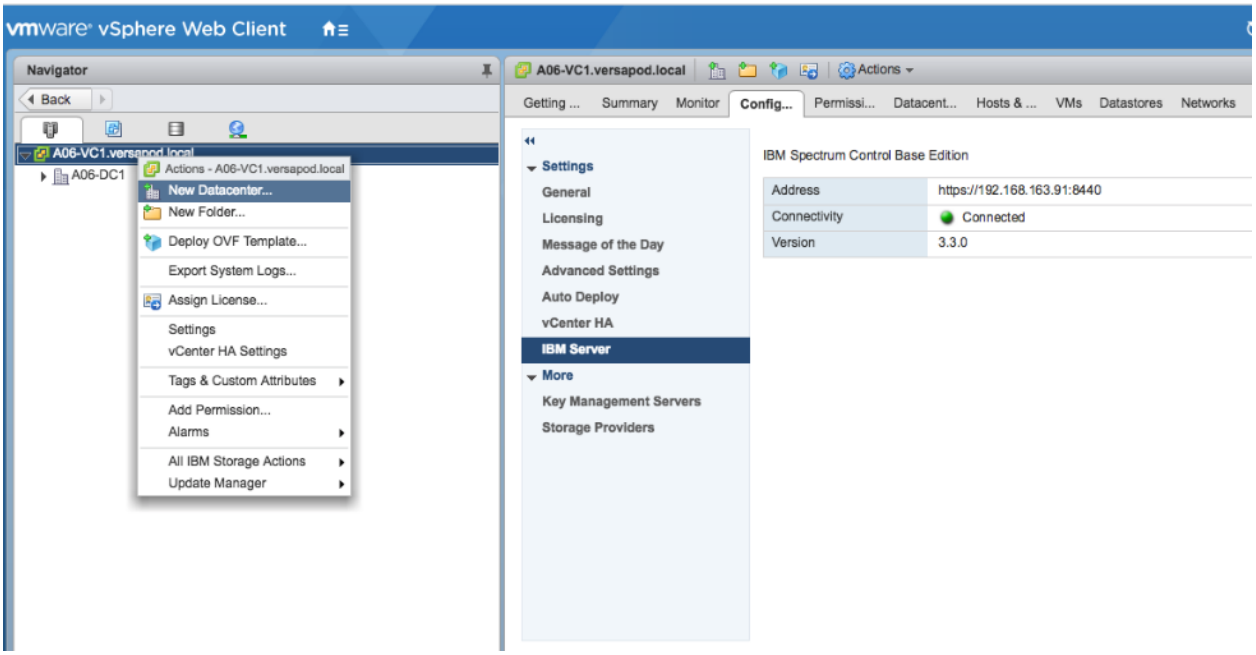

4. From the New Datacenter pop-up dialogue enter in a Datacenter name and click OK.

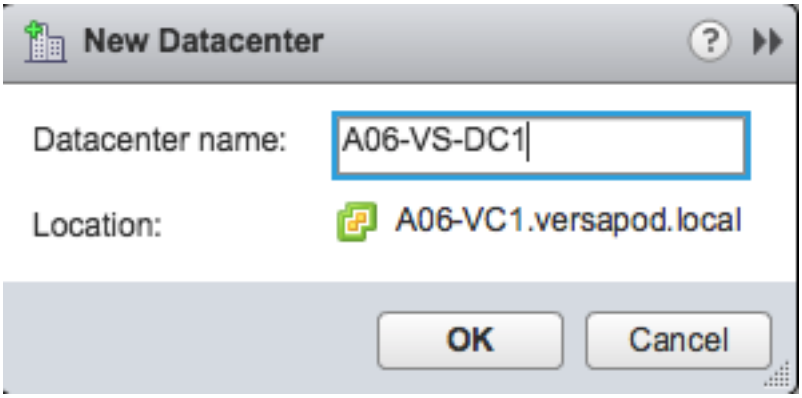

Add the VMware ESXi Hosts Using the VMware vSphere Web Client

To add the VMware ESXi Hosts using the VMware vSphere Web Client, complete the following steps:

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

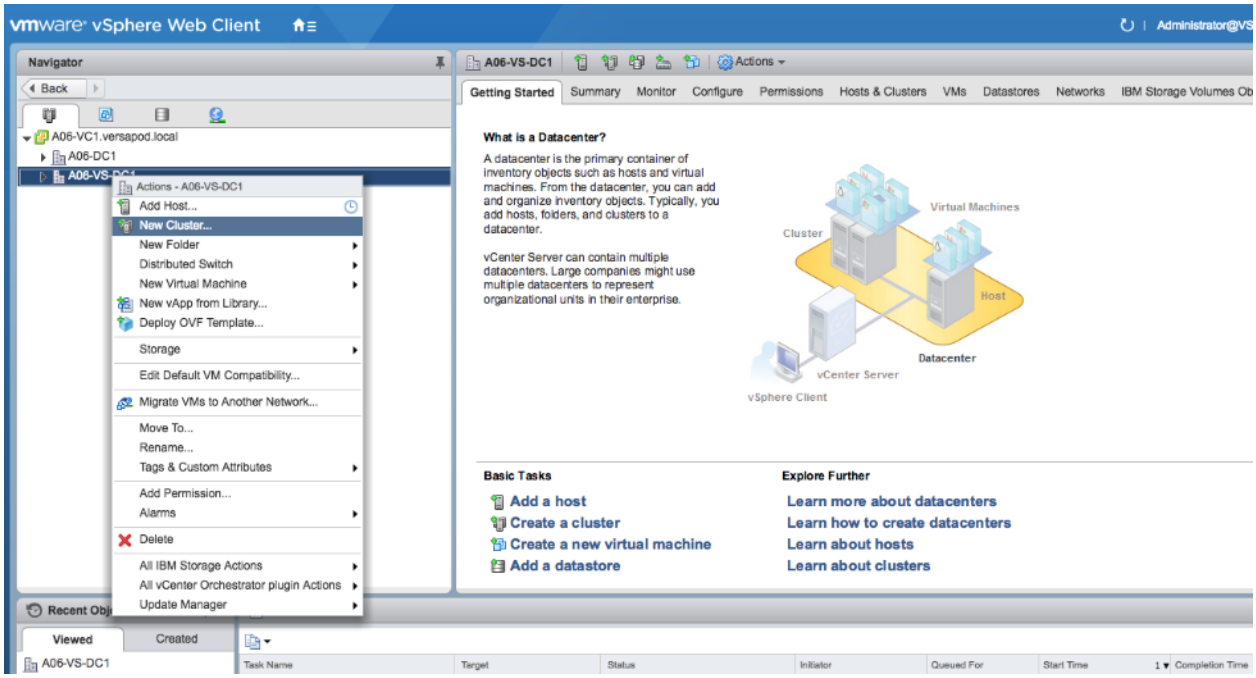

2. Enter a name for the new cluster, select the DRS and HA checkmark boxes, leaving all other options with defaults.

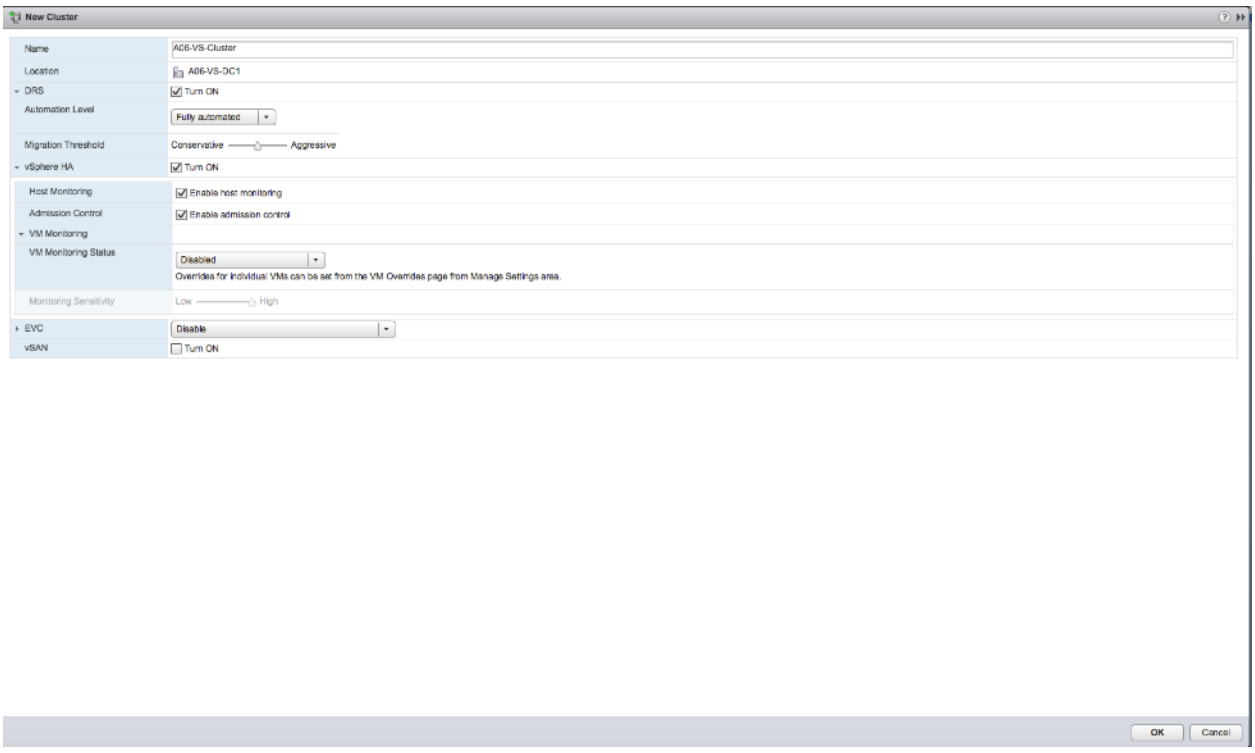

If mixing Cisco UCS B or C-Series M2, M3 or M4 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to Enhanced vMotion Compatibility (EVC) Processor Support.

3. Click OK to create the cluster.

4. Right-click the newly created cluster and from the drop-down list select the Add Host.

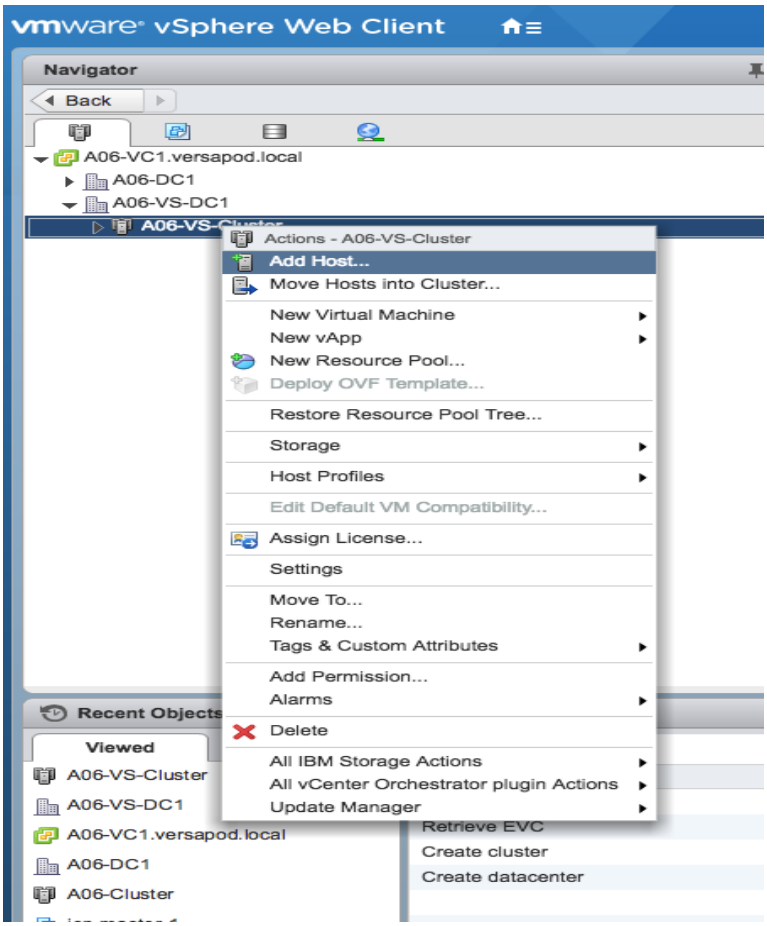

5. Enter the IP or FQDN of the first ESXi host and click Next.

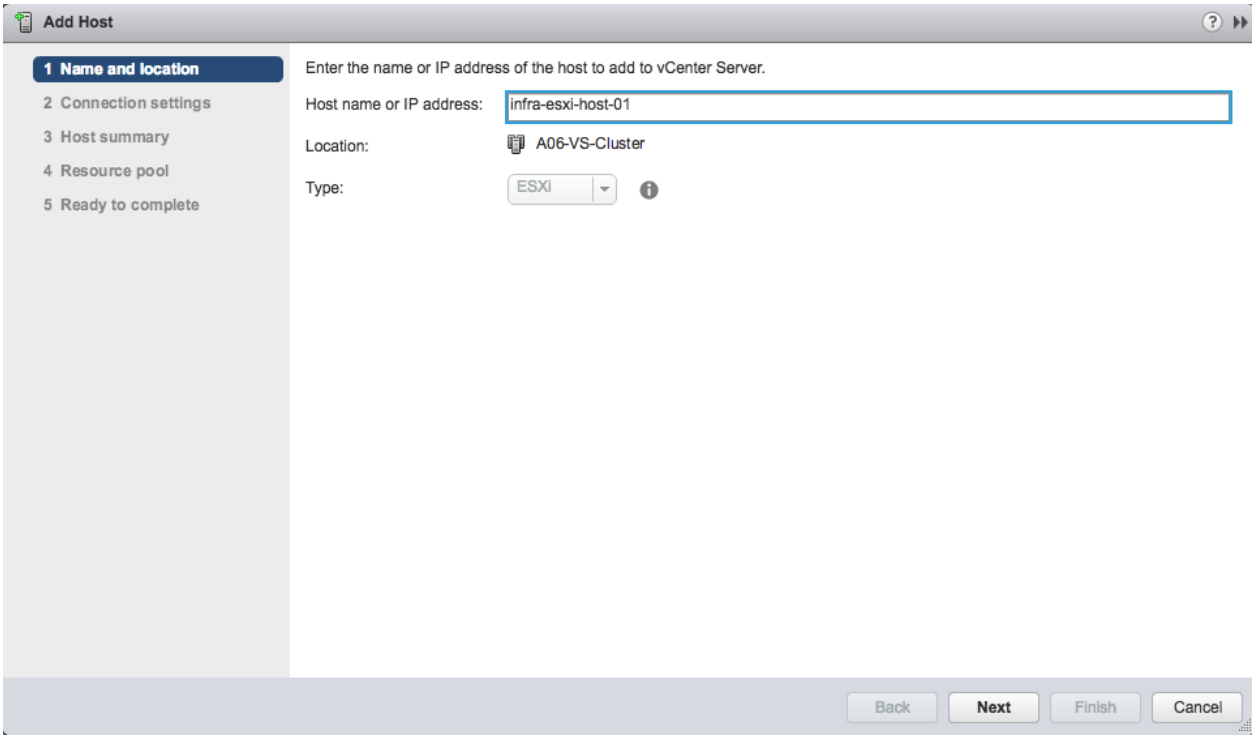

6. Enter root for the User Name, provide the password set during initial setup and click Next.
7. Click Yes in the Security Alert pop-up to confirm the host's certificate.
8. Click Next past the Host summary dialogue.
9. Provide a license by clicking the green + icon under the License title, select an existing license, or skip past the Assign license dialogue by clicking Next.
10. Leave 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.

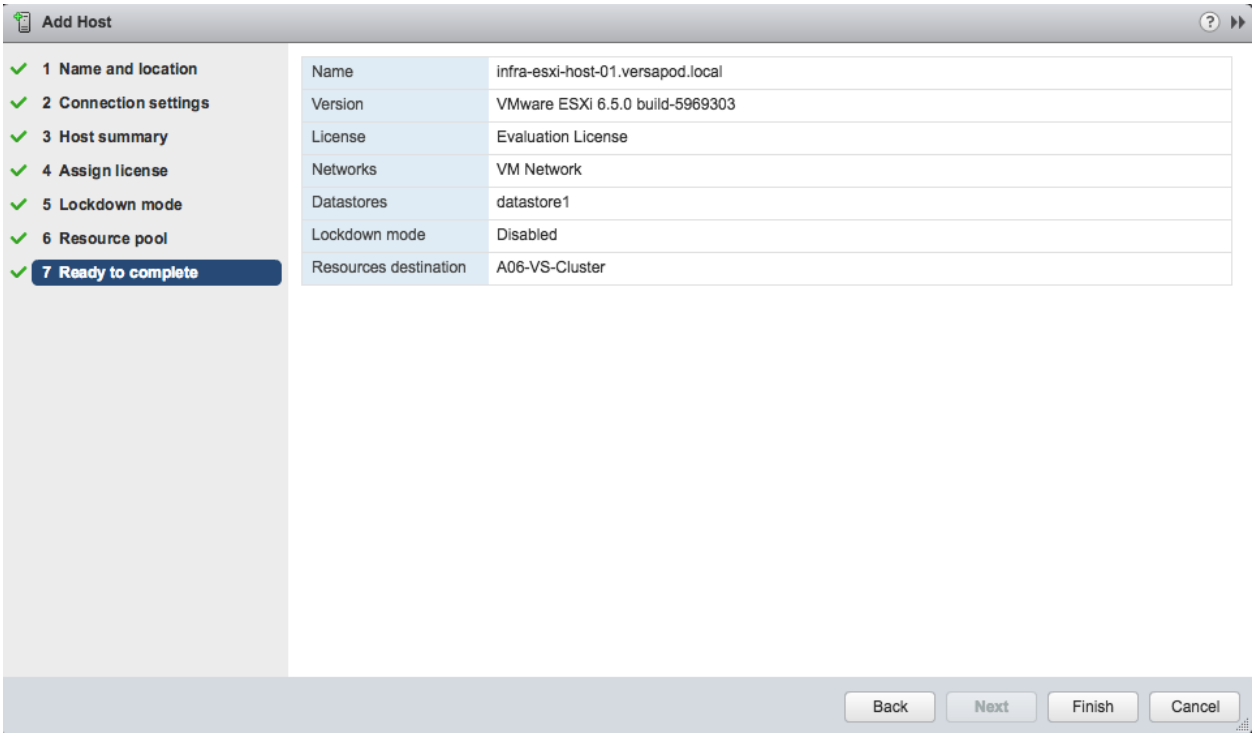

13. Repeat these steps for each ESXi host to be added to the cluster.

14. In vSphere, in the left pane right-click the newly created cluster, and under Storage click Rescan Storage.

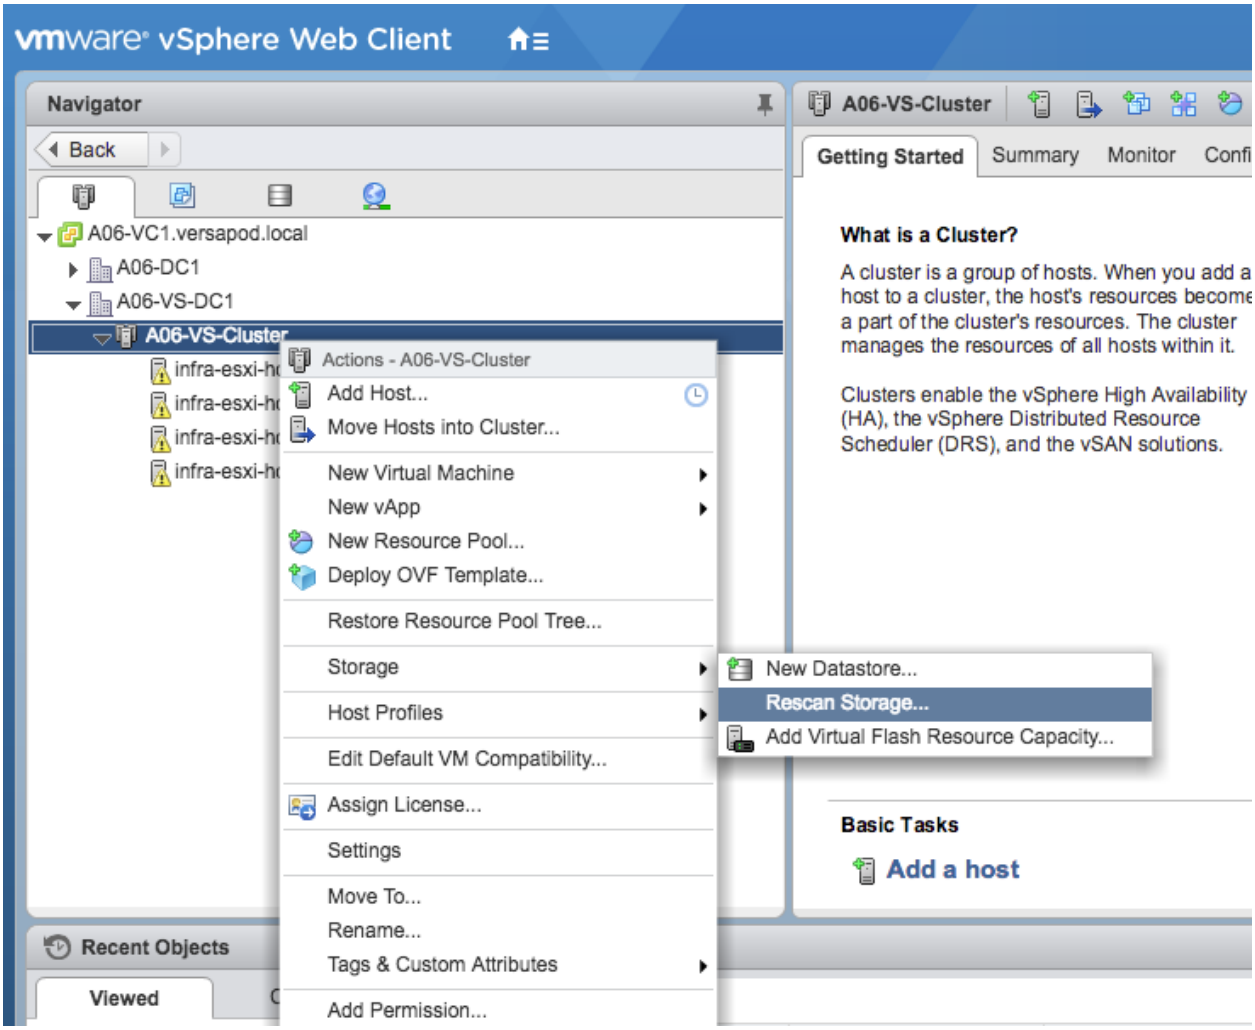

ESXi Dump Collector Setup for iSCSI Hosts (iSCSI Configuration Only)

ESXi hosts booted with iSCSI need to be configured with ESXi dump collection. The Dump Collector functionality is supported by the vCenter but is not enabled by default on the vCenter Appliance.

Make sure the account used to login is Administrator@vsphere.local (or a system admin account)

1. In the vSphere web client, select Home.
2. In the center pane, click System Configuration.
3. In the left hand pane, select Services and select VMware vSphere ESXi Dump Collector.
4. In the Actions menu, choose Start.
5. In the Actions menu, click Edit Startup Type.
6. Select Automatic.

7. Click OK.

8. Select Home > Hosts and Clusters.

9. Expand the DataCenter and Cluster.

10. For each ESXi host, right-click the host and select Settings. Scroll down and select Security Profile. Scroll down to Services and select Edit. Select SSH and click Start. Click OK.

11. SSH to each ESXi hosts and use root for the user id and the associated password to log into the system. Type the following commands to enable dump collection:

```
esxcli system coredump network set --interface-name vmk0 --server-ipv4
<vcenter-ip> --server-port 6500
esxcli system coredump network set --enable true
esxcli system coredump network check
```

12. Optional: Turn off SSH on the host servers.

Configure ESXi Networking

This section covers the virtual switch (vSwitch) setup for Management, vMotion and iSCSI storage traffic and vSphere Distributed Switch (vDS) for application traffic.

Update Management vSwitch0 Configuration

This design uses a separate vSwitch for management using two uplink vNICs with Active/Passive Failover for the port group and routing based on originating port ID for load balancing at the vSwitch level. Traffic from each vNIC take different paths across fabric to enable redundancy and load balancing.

1. Using a web browser, browse to vCenter's IP address. Login to vCenter. From vSphere Web Client, navigate to the datacenter and cluster where the host resides.

2. Select the ESXi host infra-esxi-host-01. On the right window pane, click the Configure Tab. Navigate to Networking > Virtual switches and select vSwitch0 from the Virtual Switches list. Click the Manage physical adapter (third icon) to open the Manage Physical Network Adapters for vSwitch0 window.

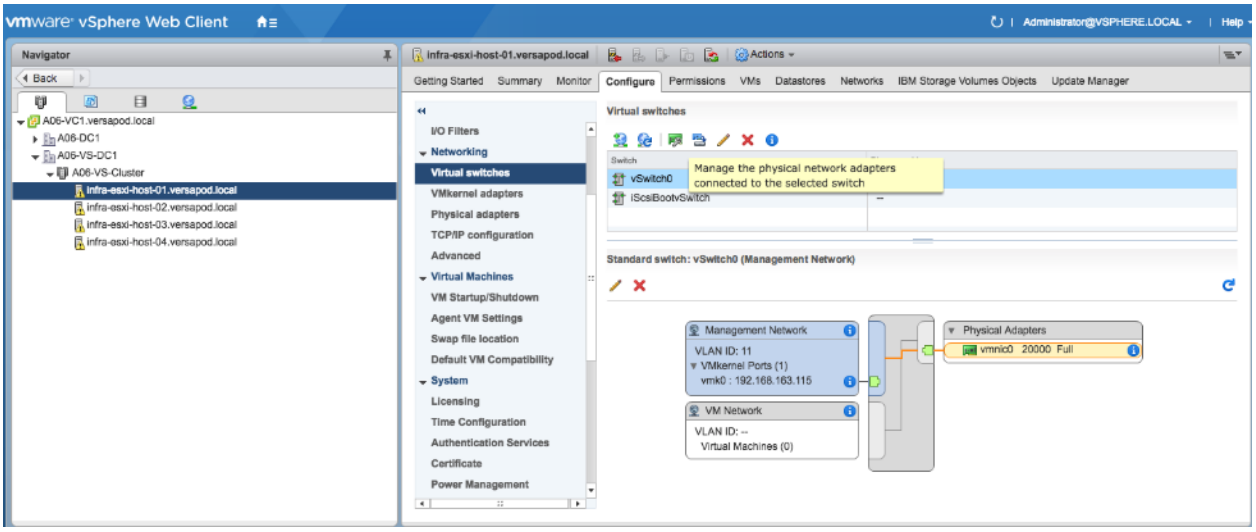

3. Click the green [+] to add a second Adapter. Select an unused vmnic from the list of Network Adapters.

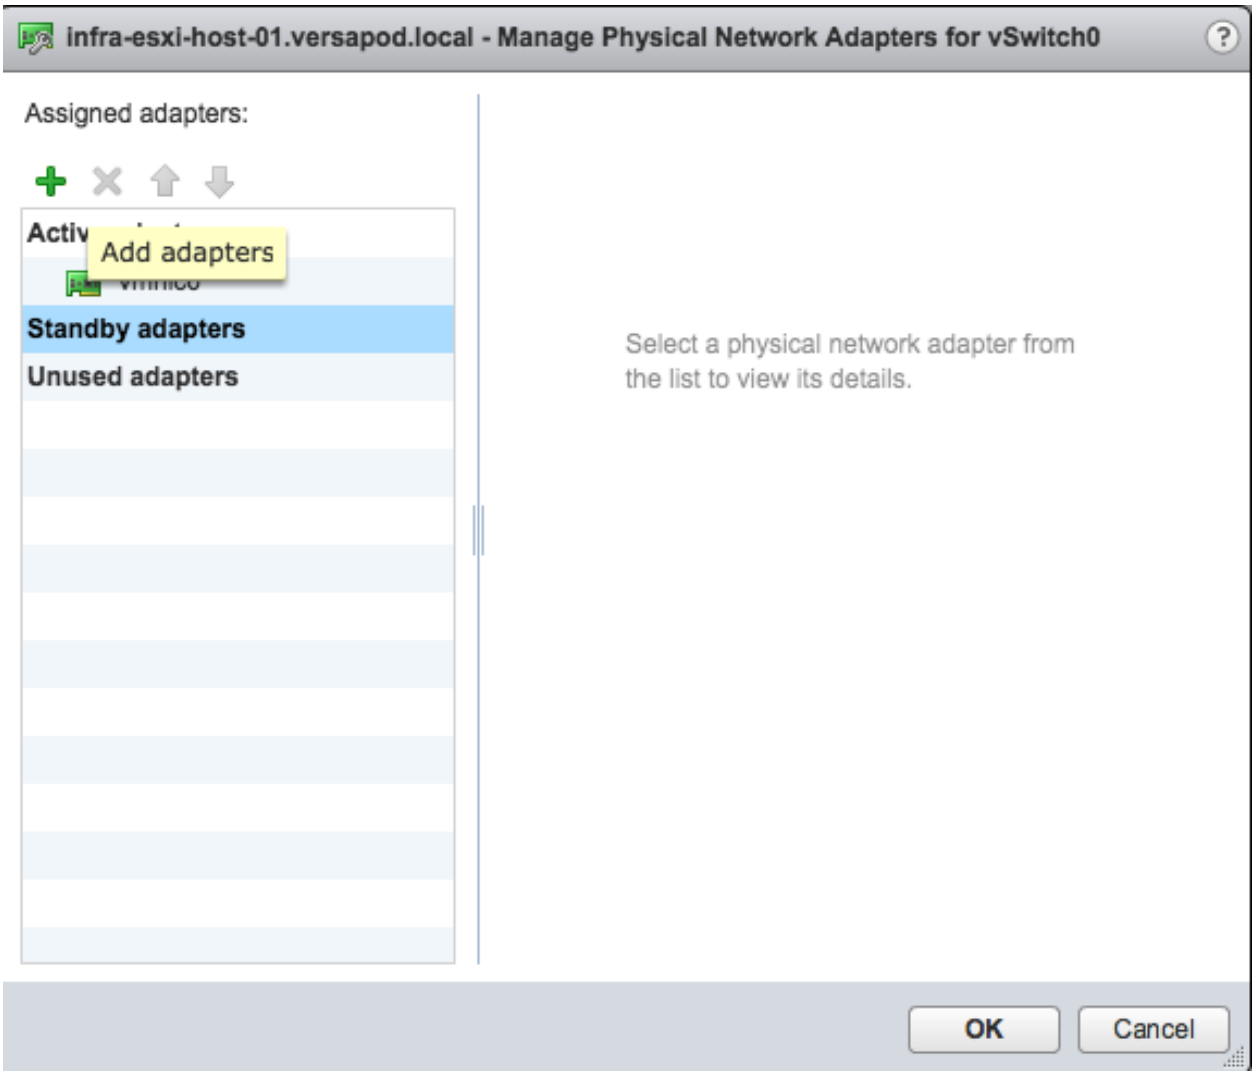

4. Select vmnic1 and click OK to add the vmnic as a second Active adapter to vSwitch0.

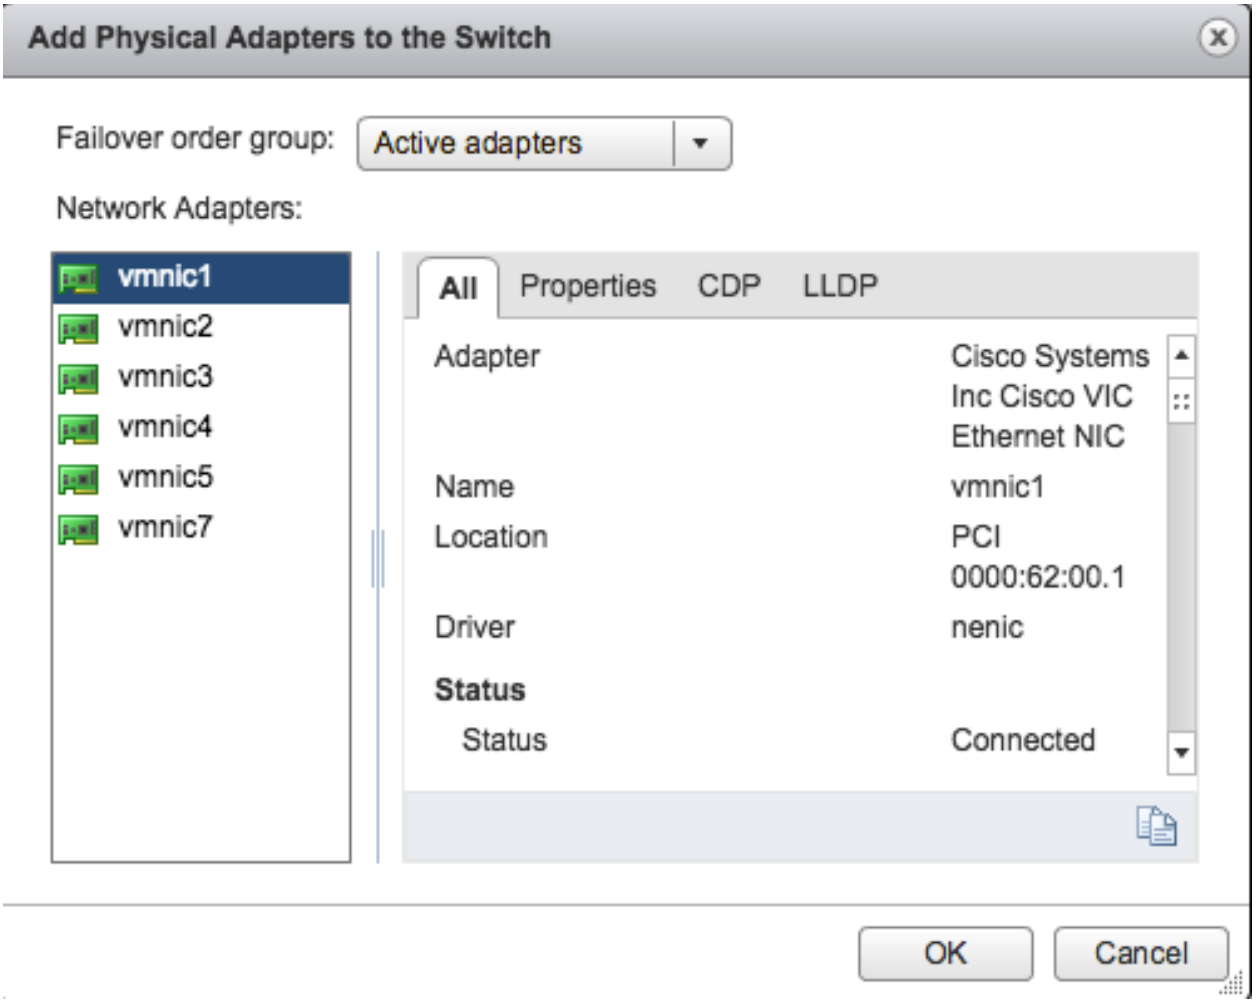

5. Click OK to commit the change.

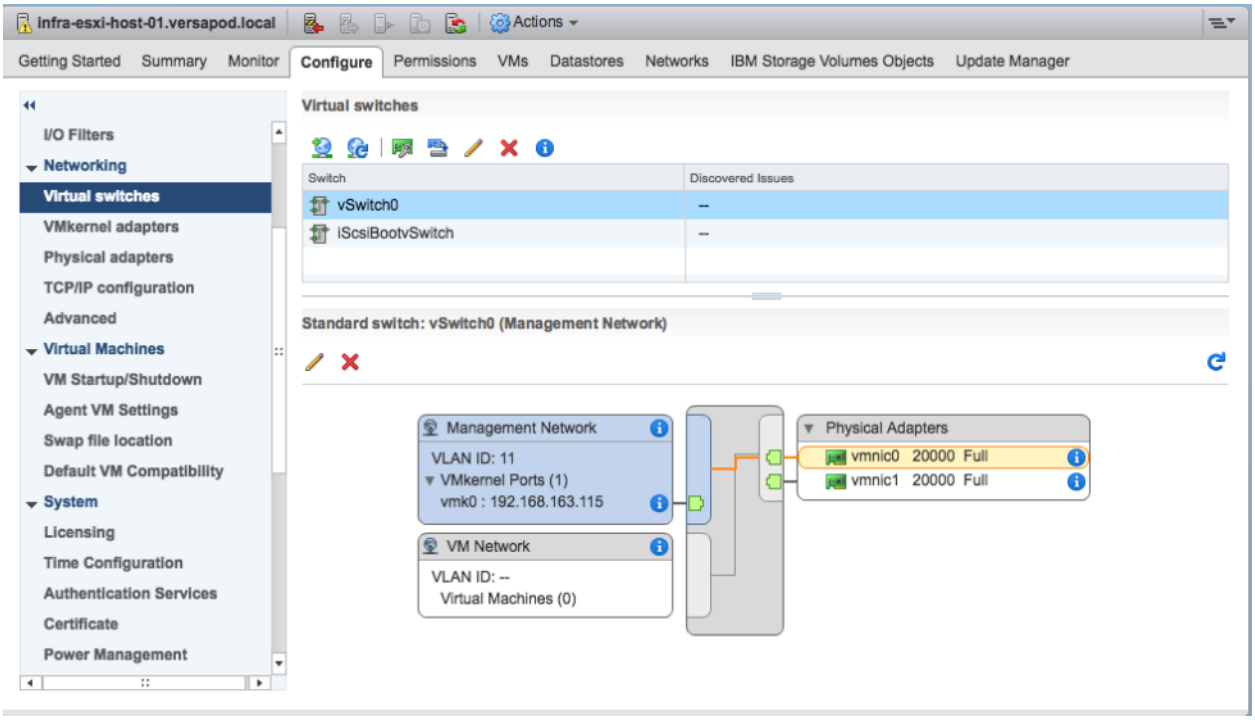

6. While the vSwitch0 of the host still selected, click the Edit Settings icon (5th icon) to open the Edit settings window.

7. Under Teaming and failover, verify the load balancing and failover configuration. Both vmnics should be listed as uplinks under Active adapters.

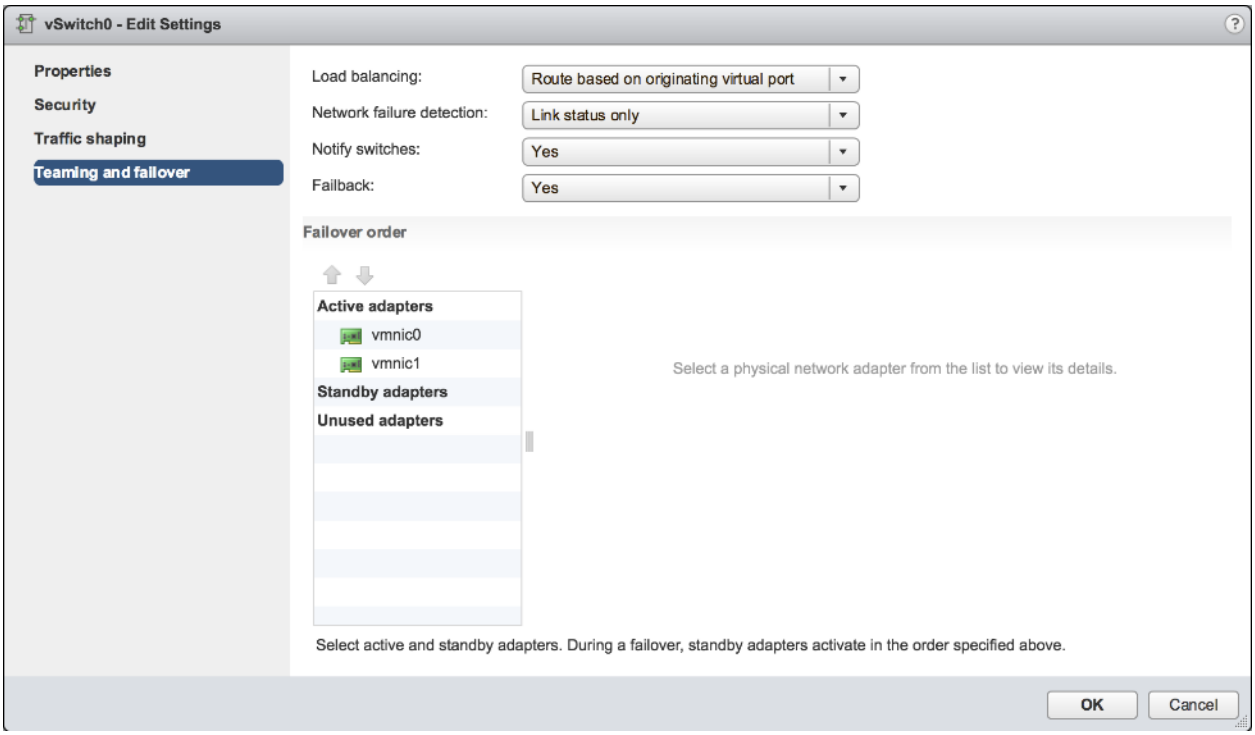

8. Repeat the above procedure for all the ESXi hosts in the Cluster.

Create vSwitch1 for vMotion

This design uses a separate vSwitch for vMotion with two uplink vNICs. To create and setup vSwitch1 for vMotion, complete the following steps.

1. Using a web browser, browse to vCenter's IP address. Login to vCenter. From vSphere Web Client, navigate to the datacenter and cluster where the host resides

2. Select the host infra-esxi-host-01. On the right window pane, select the Configure Tab. Navigate to Networking > Virtual Switches. Click the Add host networking icon (1st icon).

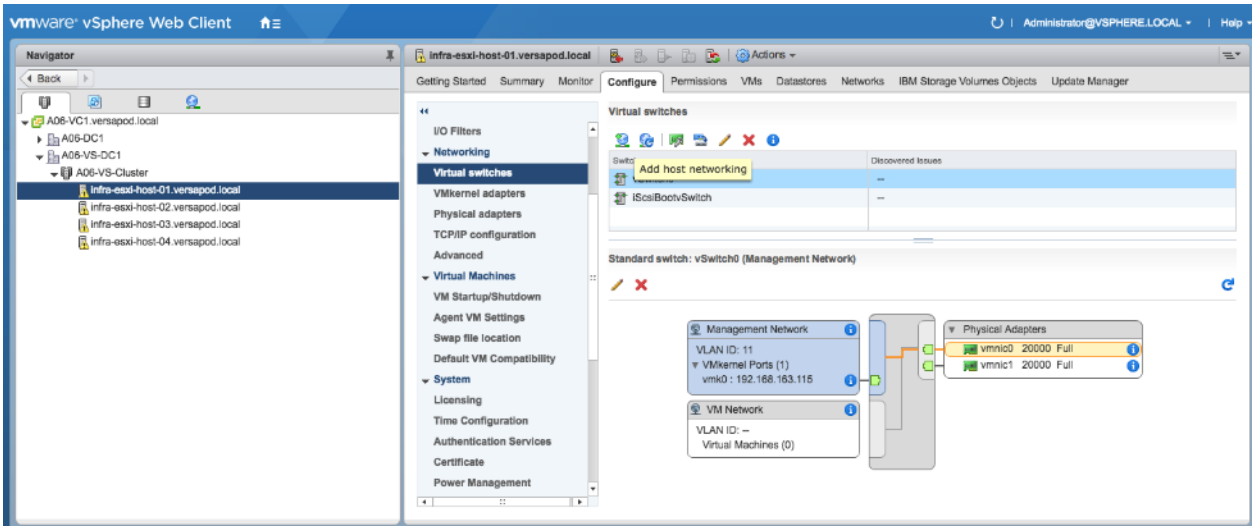

3. Leave VMkernel Network Adapter selected within Select connection type of the Add Networking pop-up window that is generated, and click Next.

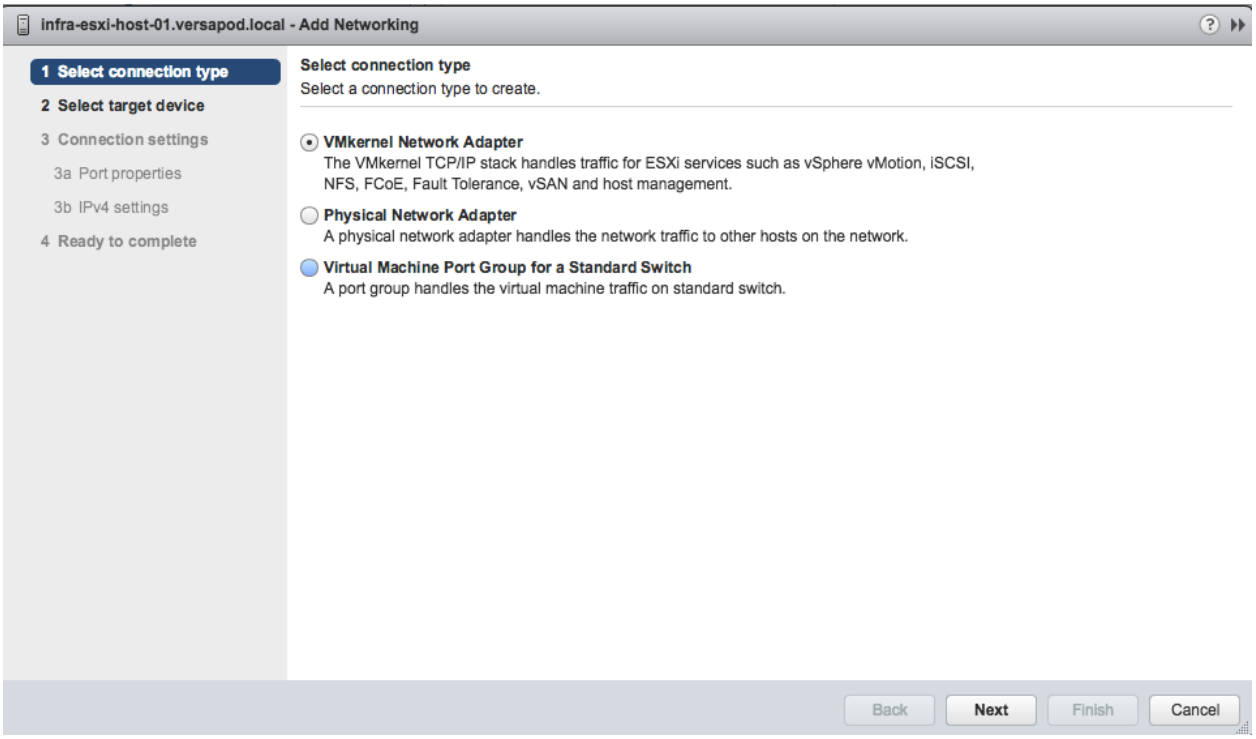

4. Select New Standard switch and click Next.

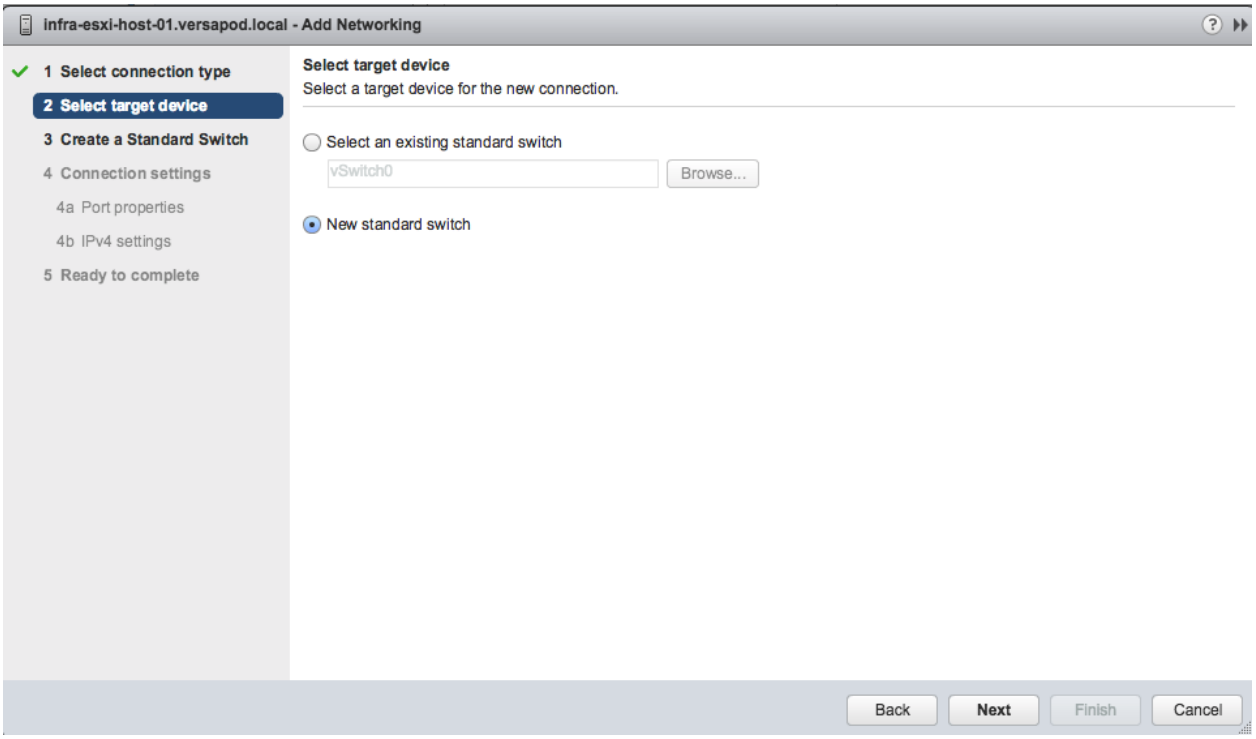

5. Within Select target device, click the New standard switch option, and click Next.

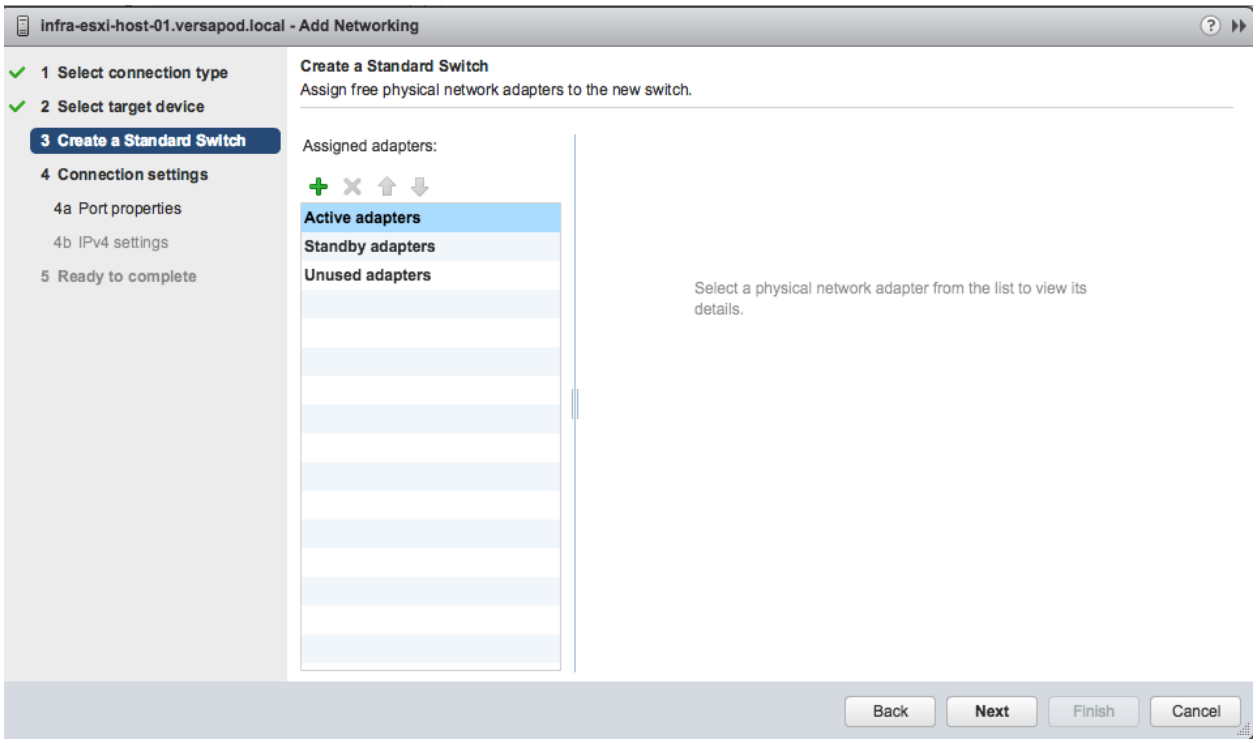

6. Within the Create Standard Switch dialogue press the green + icon below Assigned adapters.

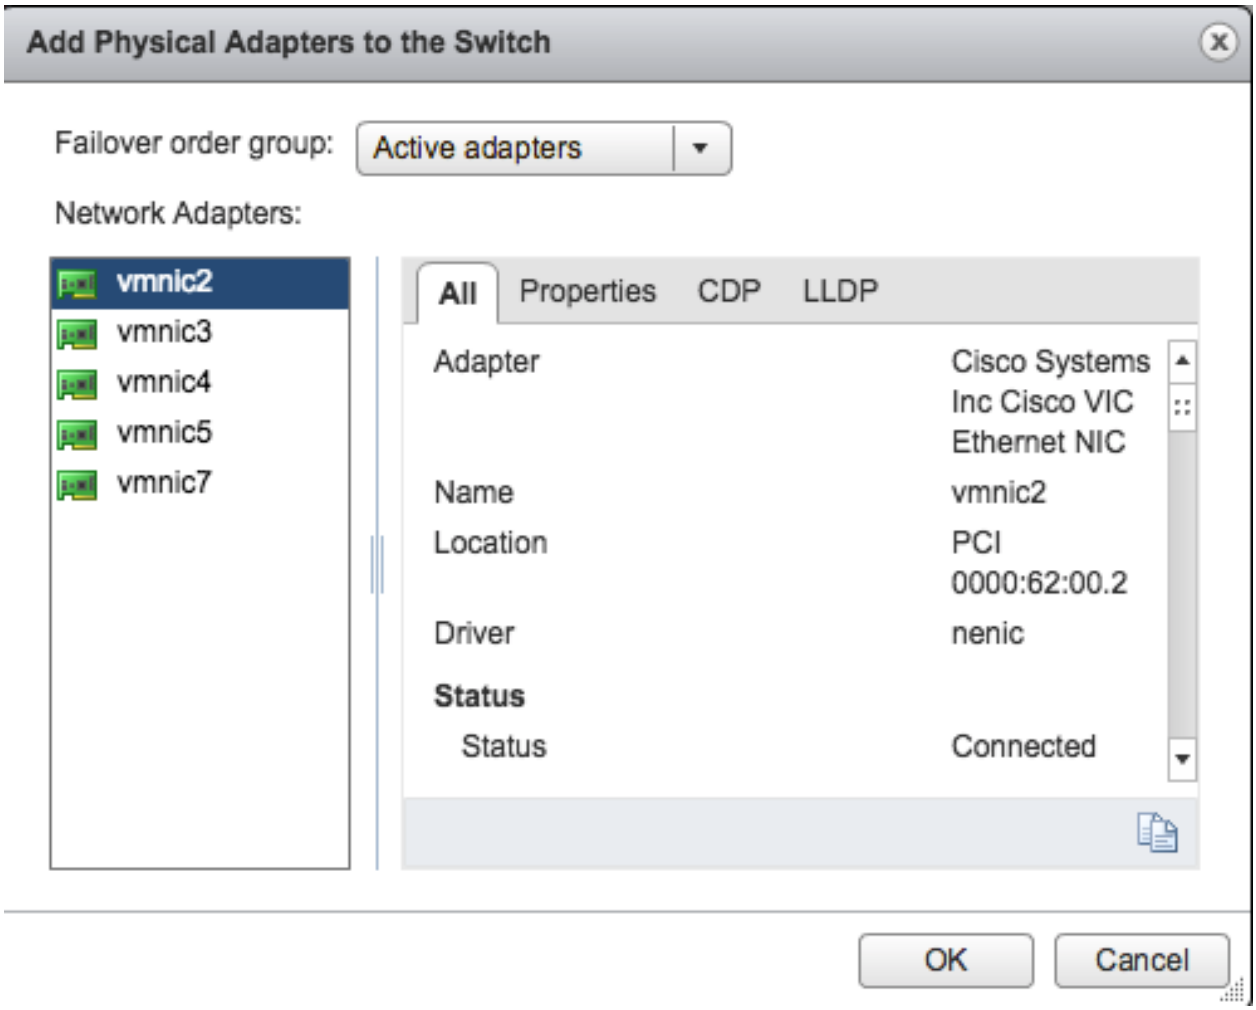

7. Select vmnic2 within the Network Adapters and click OK.

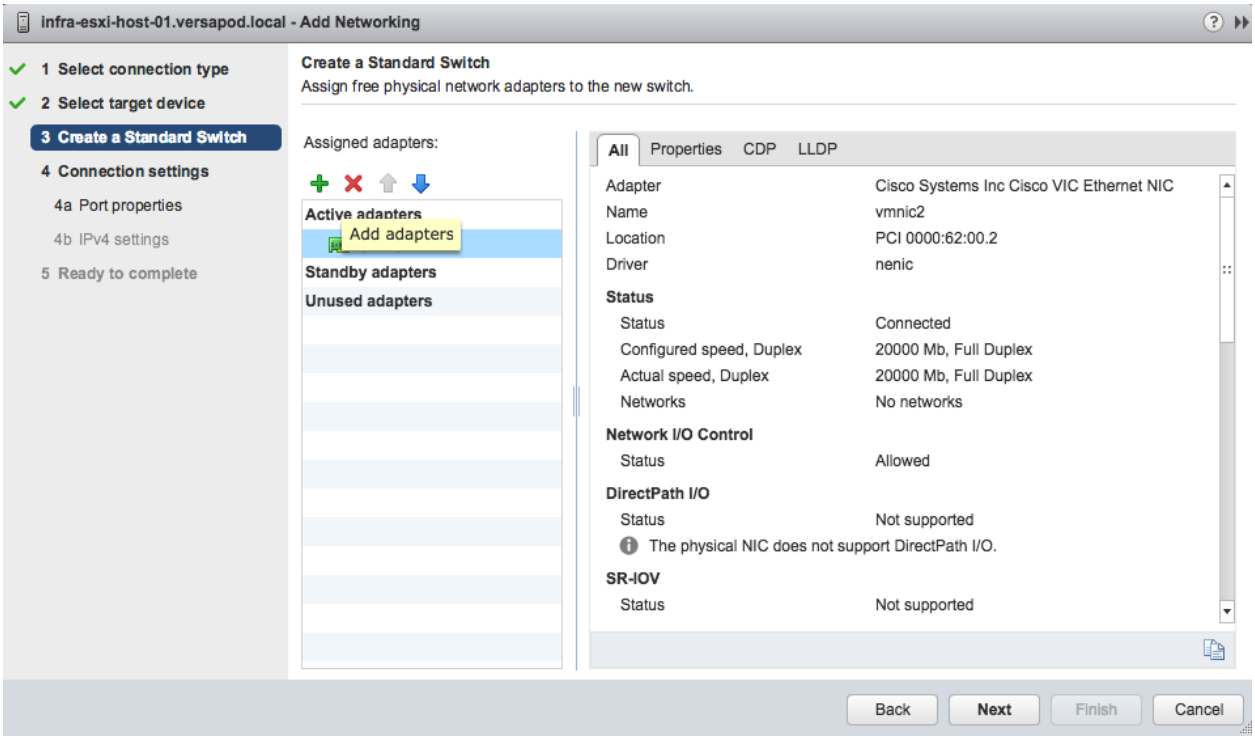

8. While still in the Create a Standard Switch dialogue, click the green + icon one more time.

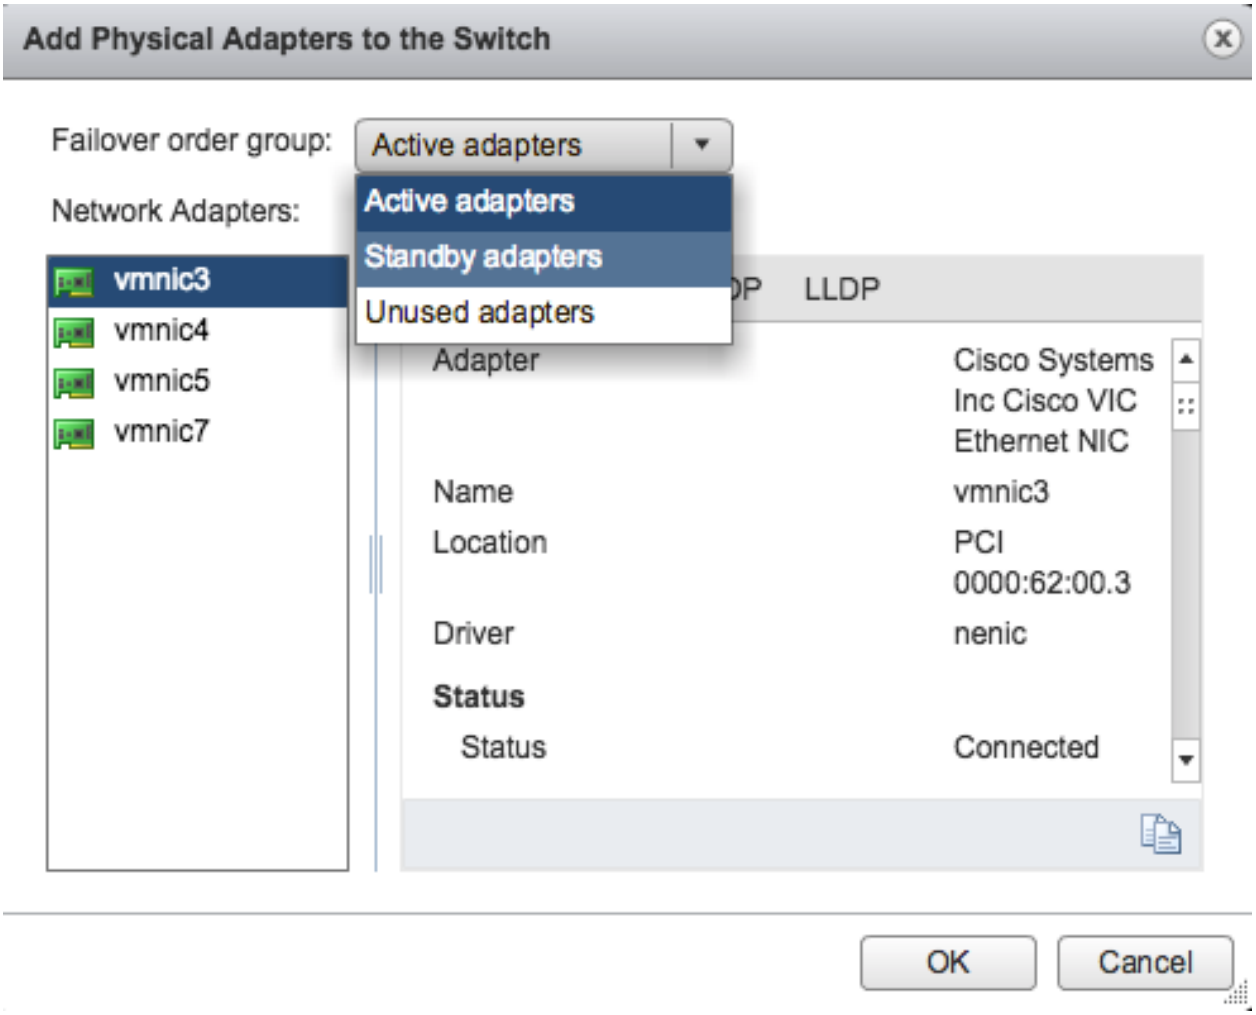

9. Select vmnic3, and from the Failover order group drop-down list, select Standby adapters. Click OK.

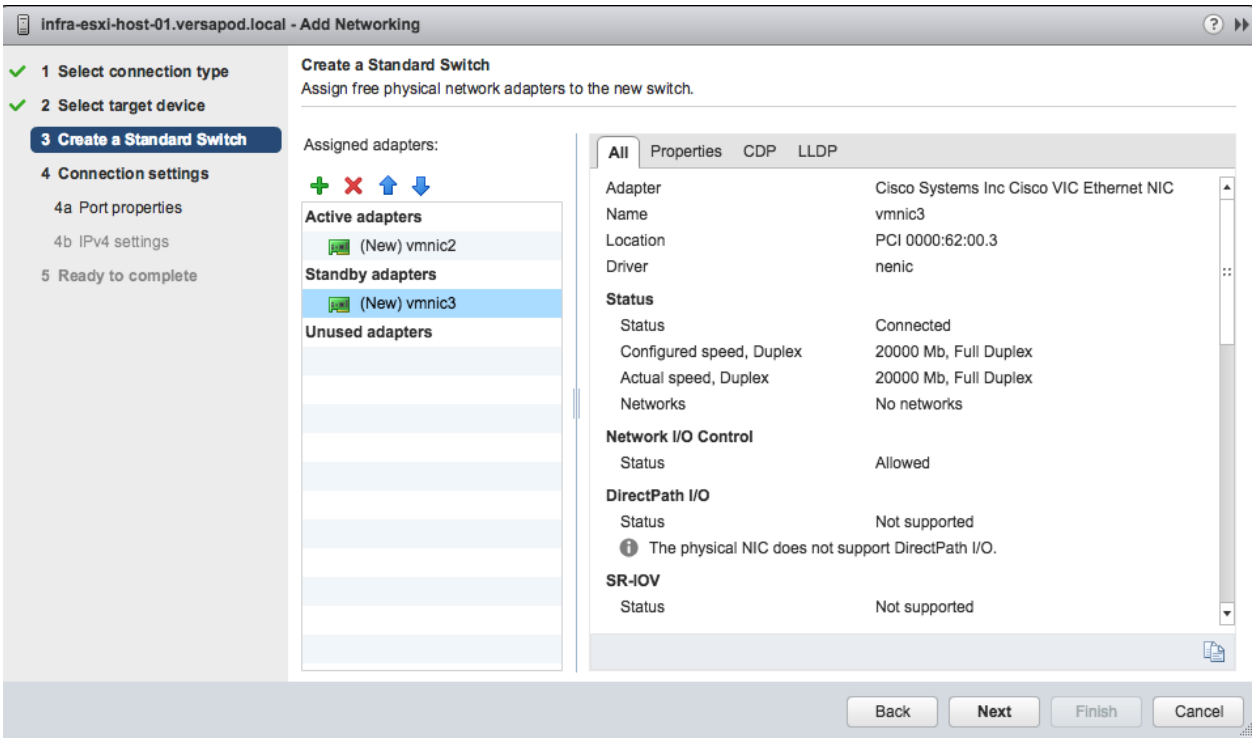

10. Click Next.

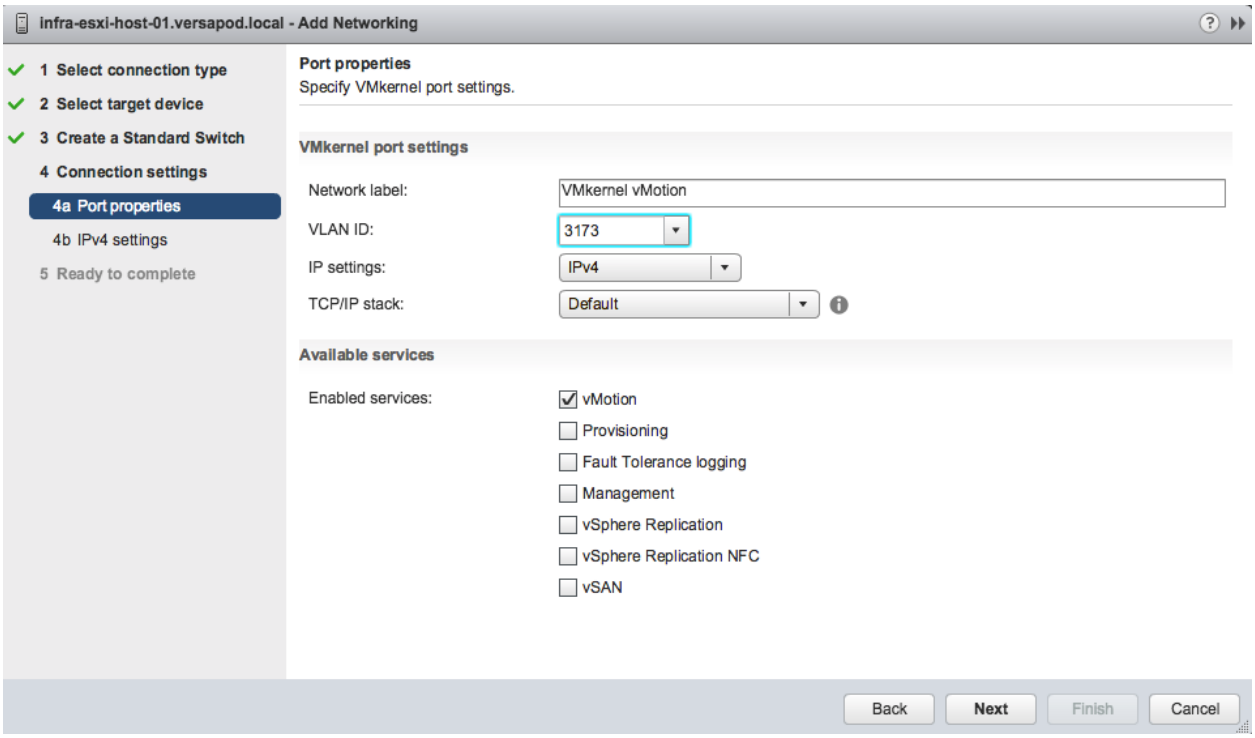

11. Within Port properties under Connection settings, set the Network label to be VMkernel vMotion, set the VLAN ID to the value for <vMotion VLAN id>, and checkmark vMotion traffic under Available services. Click Next.

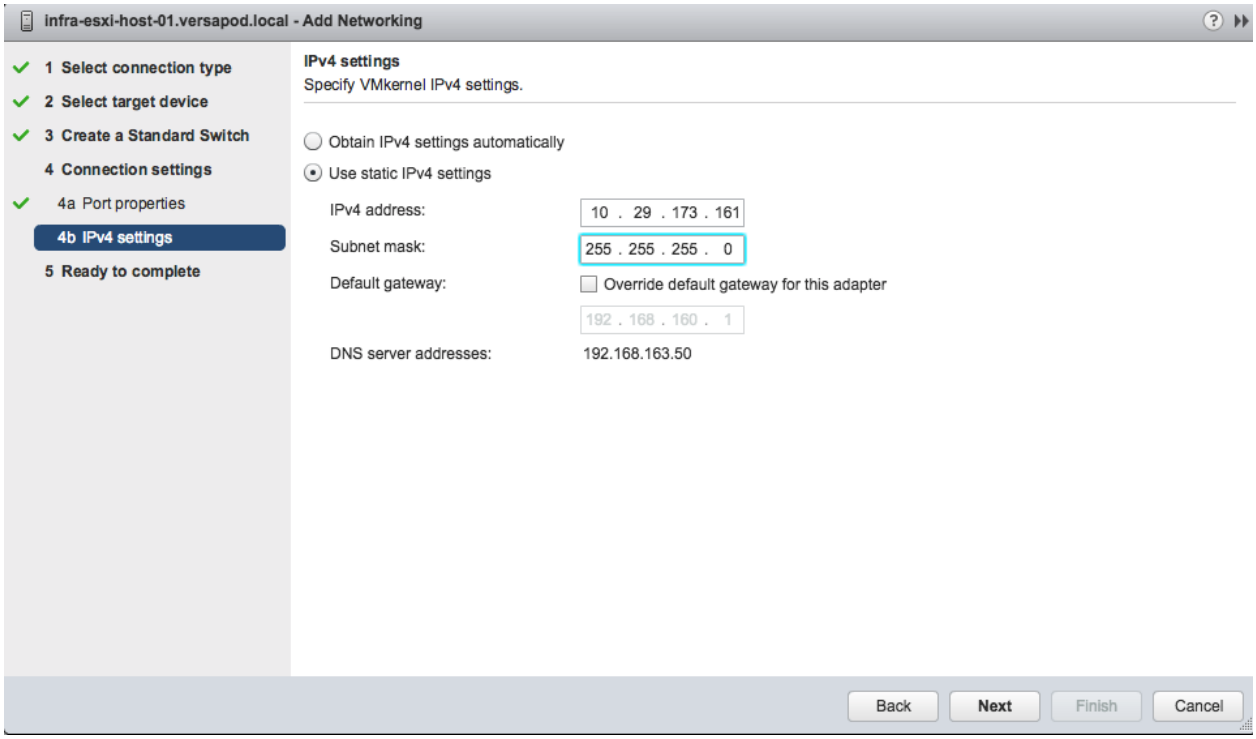

12. Enter <<vMotion IP Address>> in the filed for IPv4 address, and <vMotion Subnet Mask> for the Subnet mask. Click Next.

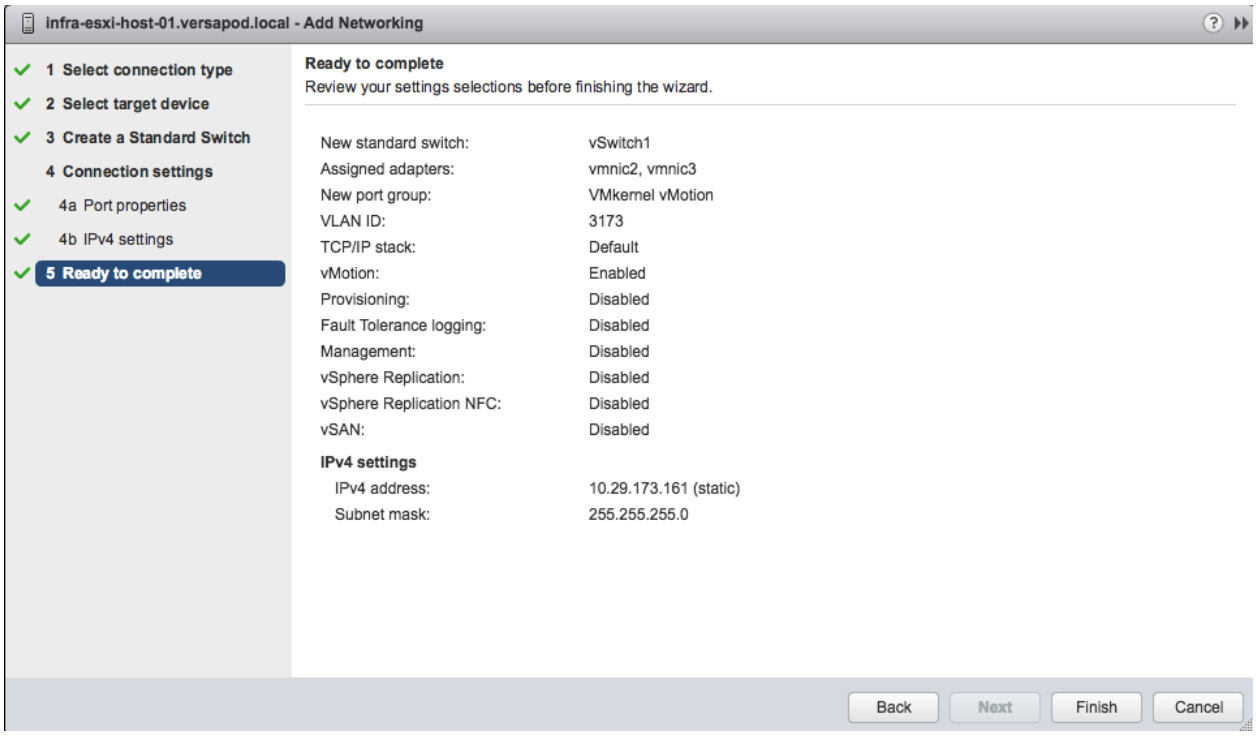

13. Confirm the values shown on the Ready to complete summary page, and click Finish to create the vSwitch and VMkernel for vMotion.

14. Still within the Configure tab for the host, under Networking -> Virtual switches, make sure that vSwitch1 is selected, and click on the pencil icon under the Virtual Switches title to edit the vSwitch properties to adjust the MTU for the vMotion vSwitch.

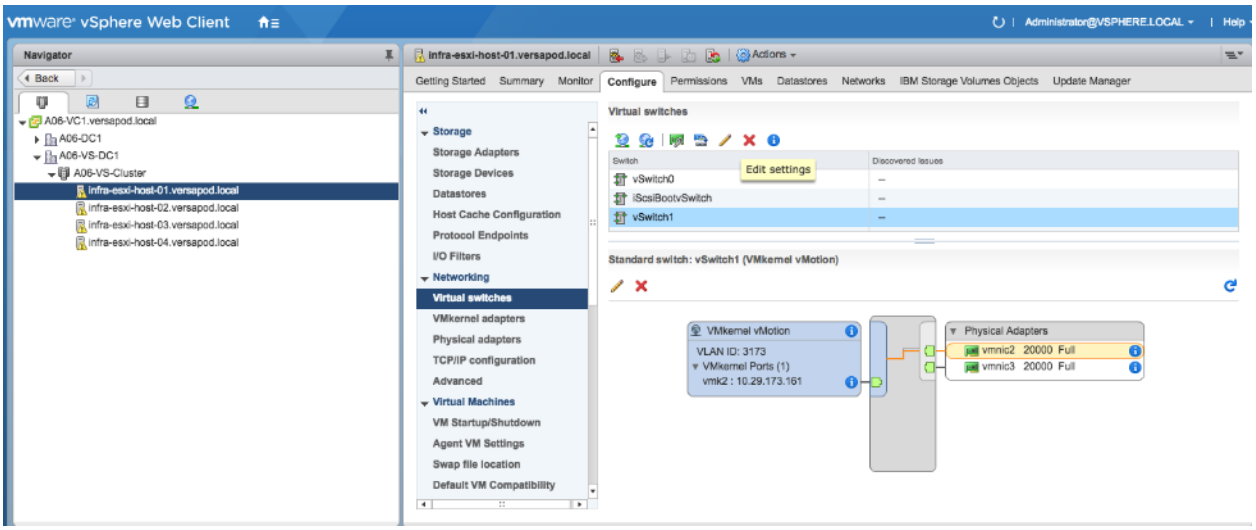

15. Enter 9000 in the Properties dialogue for the vSwitch1 – Edit Settings pop-up that appears. Click OK to apply the change.

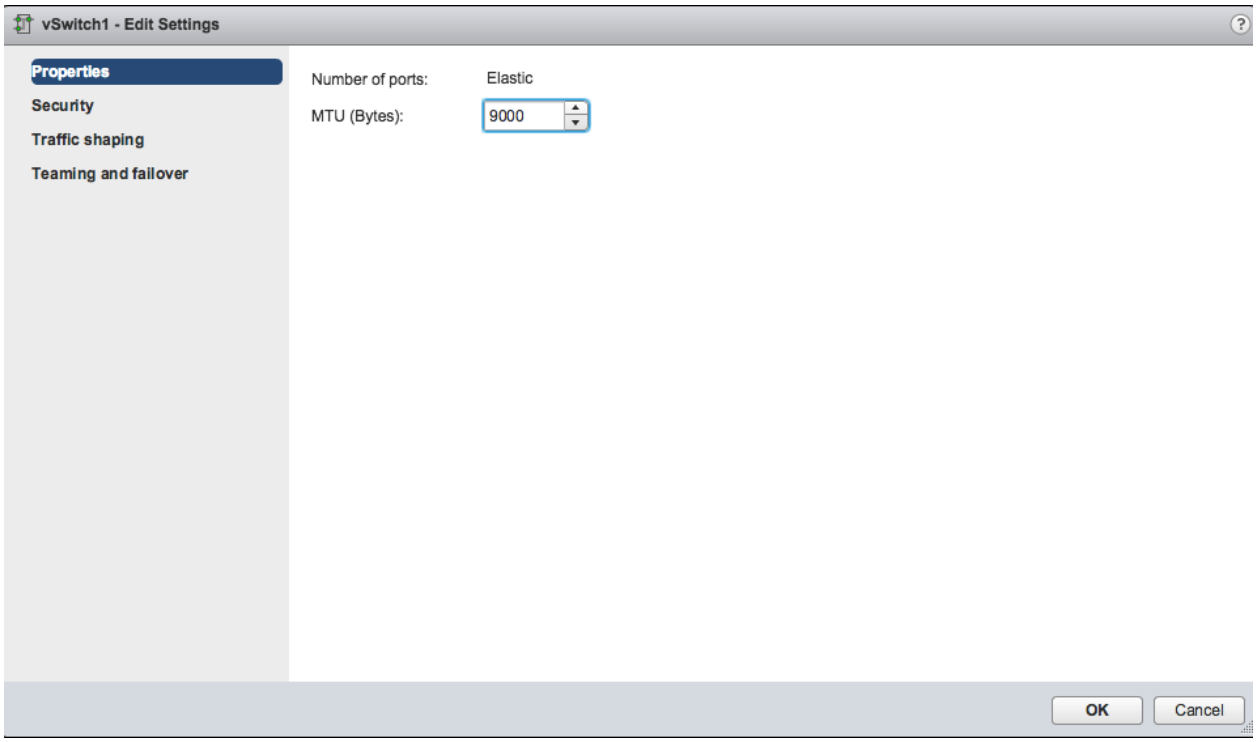

16. Click the VMkernel adapters under Networking for the host, and with the VMkernel for vMotion (vmk2) selected, click the pencil icon to edit the VMkernel settings.

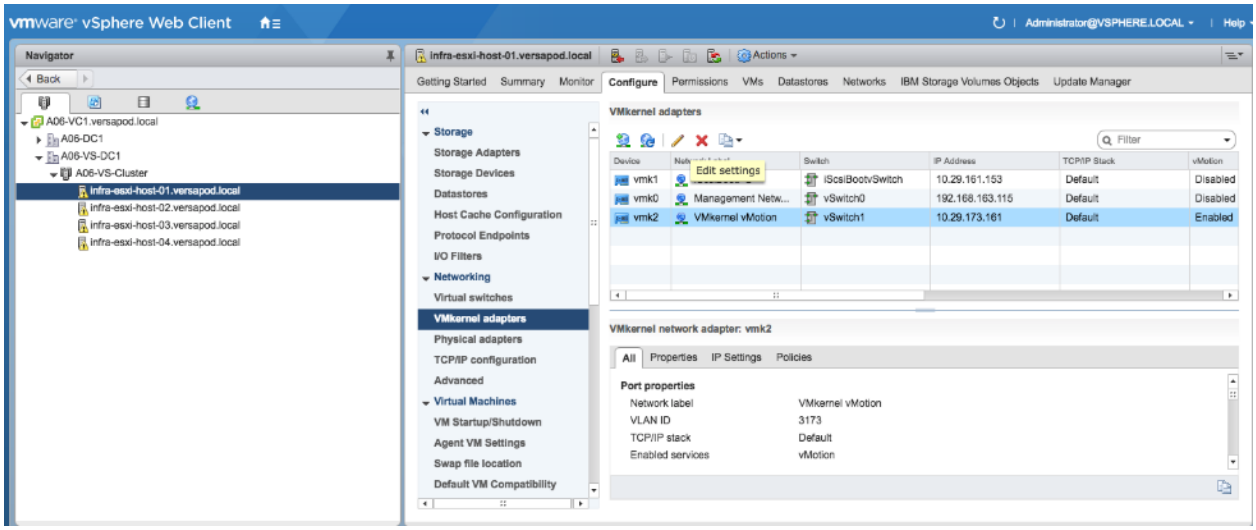

17. Click the NIC settings in the vmk2 – Edit Settings pop-up window that appears, and enter 9000 for the MTU value to use for the VMkernel. Click OK to apply the change.

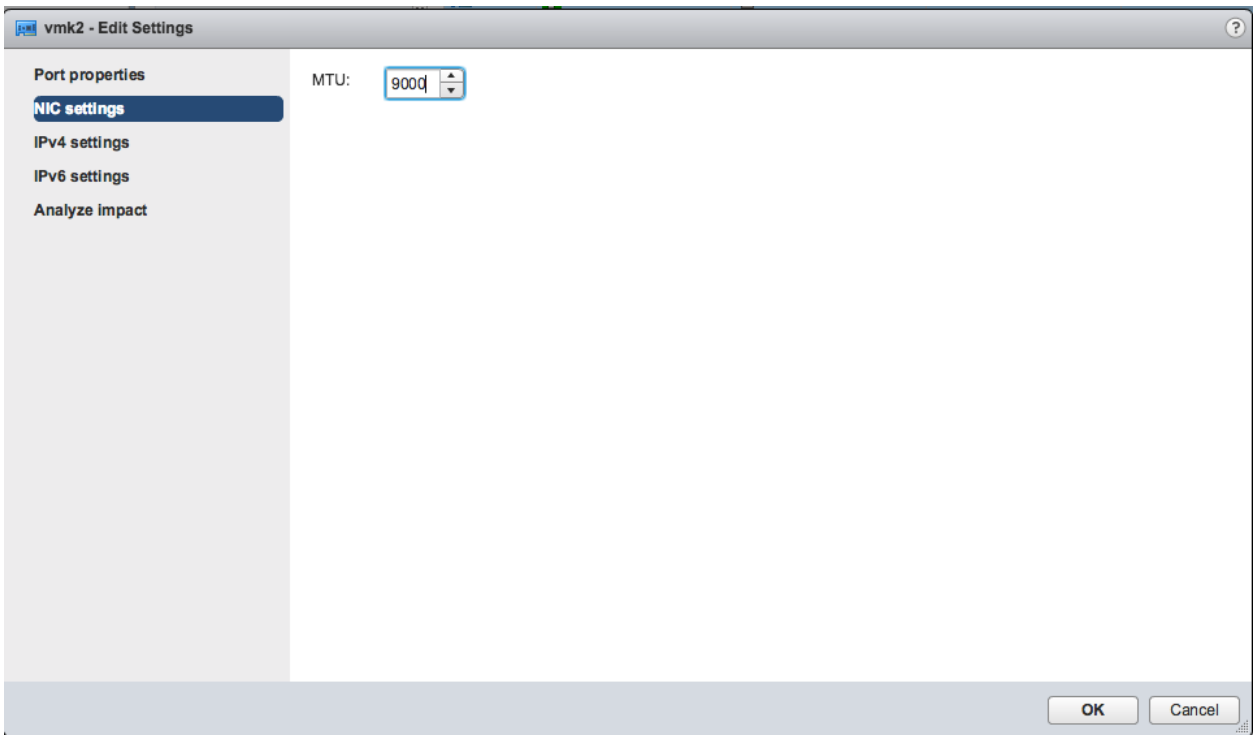

18. Repeat these steps for each host being added to the cluster, changing the vMotion VMkernel IP to an appropriate unique value for each host.

Configure iSCSI Adapters

The base ESXi installation will set up one vmkernel adapter for the iSCSI boot, with a generated vSwitch named iScsiBootvSwitch. vSwitch changes will be needed, as well as the creation of a second vmkernel adapter used for the B side iSCSI boot. To make the vSwitch changes and create the vmkernel adapter, complete the following steps for each host:

Adjust iSCSI A vSwitch MTU

1. From the vSphere Web Client, select the installed iSCSI host, click the Configure tab, and select the Virtual switches section from the Networking section on the left.

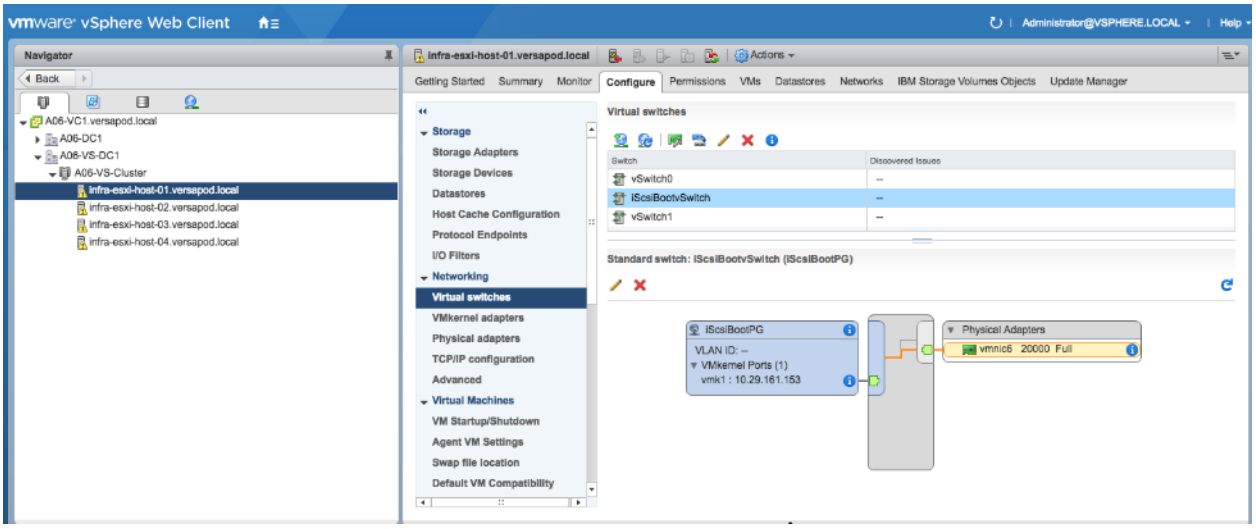

2. Select the iScsiBootvSwitch and click the pencil icon to open up Edit settings for the vSwitch.

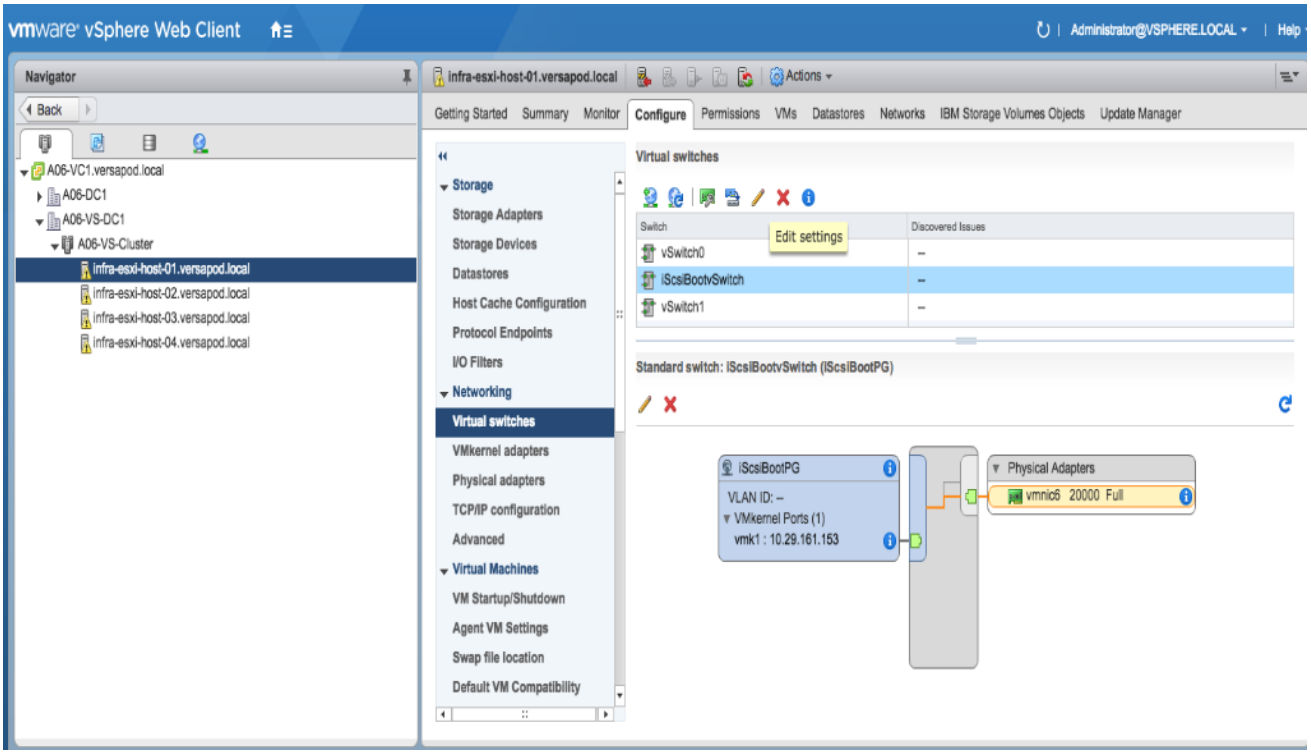

3. Within the Properties section, change the MTU from 1500 to 9000 and click OK to save the changes.

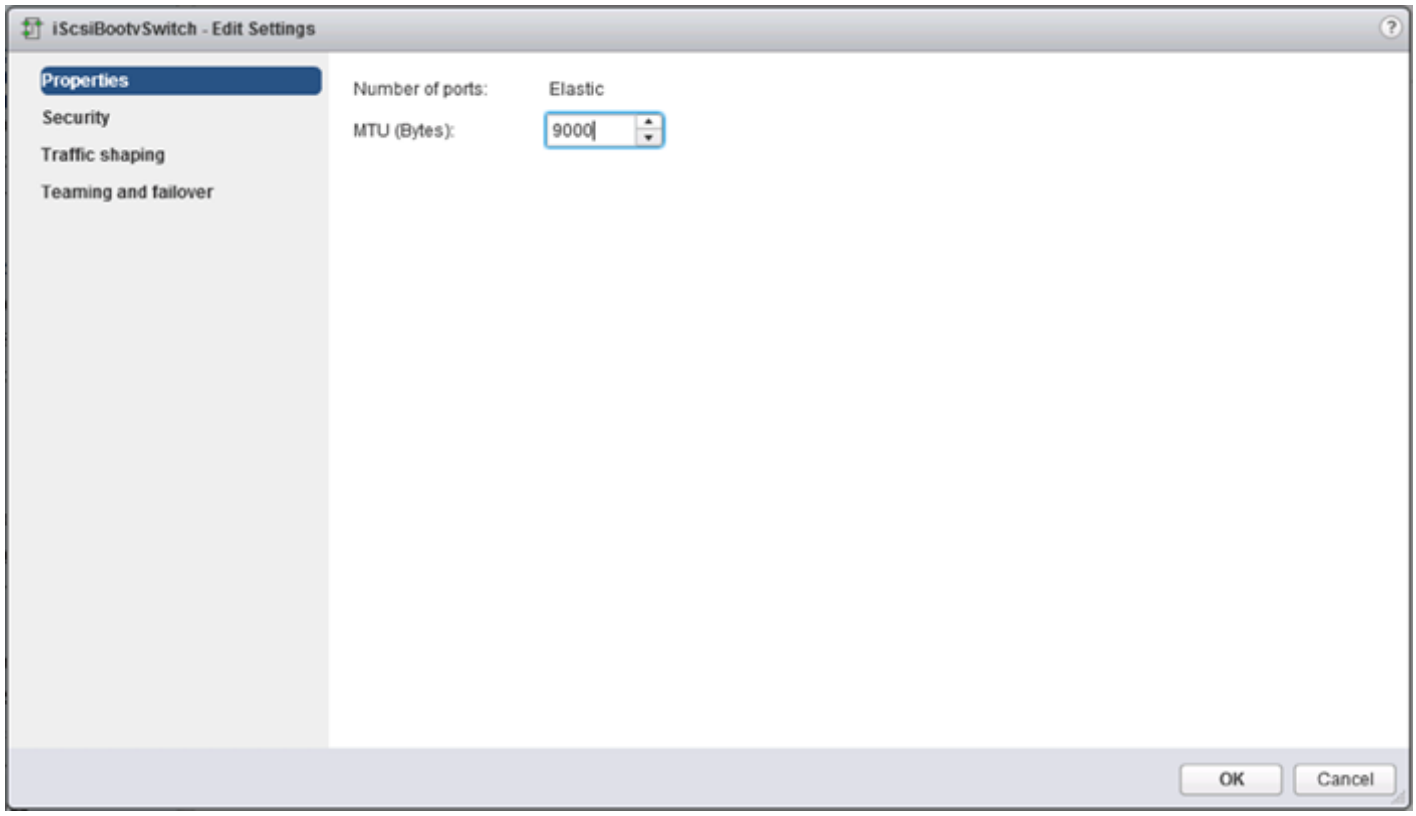

4. Click the vmk1 entry within the iScsiBootPG and select the pencil icon on the left to edit the settings of the vmkernel adapter.

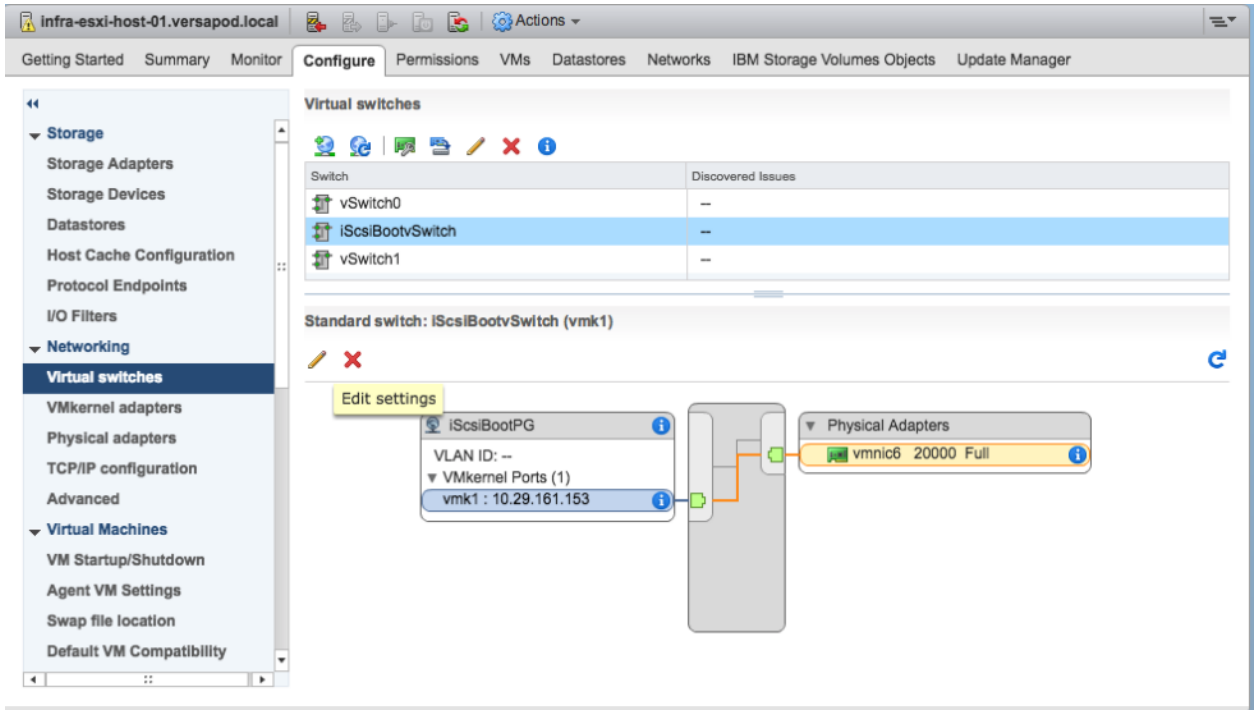

5. Select NIC settings on the left side of the Edit Settings window and adjust the MTU from 1500 to 9000.

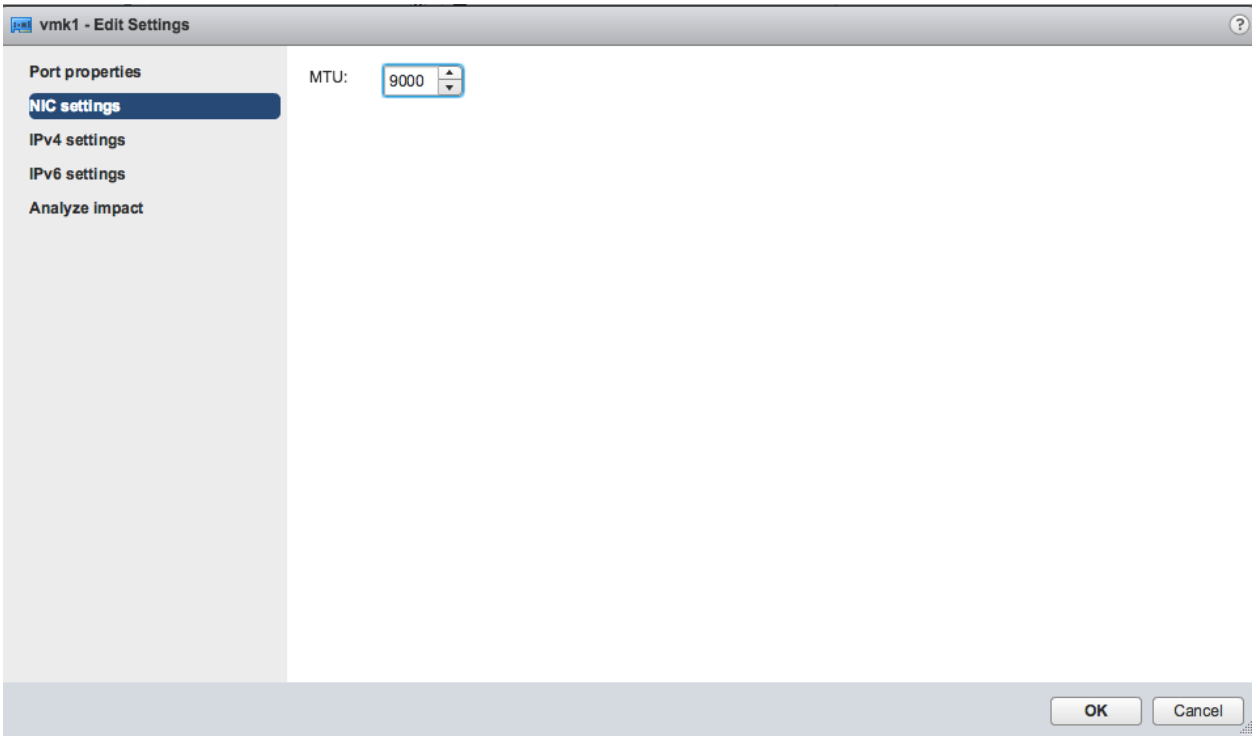

6. Click the IPv4 settings for vmk1, and change the IPv4 settings from the Cisco UCS Manager iSCSI-A-Pool assigned IP to one that is not in the IP block.

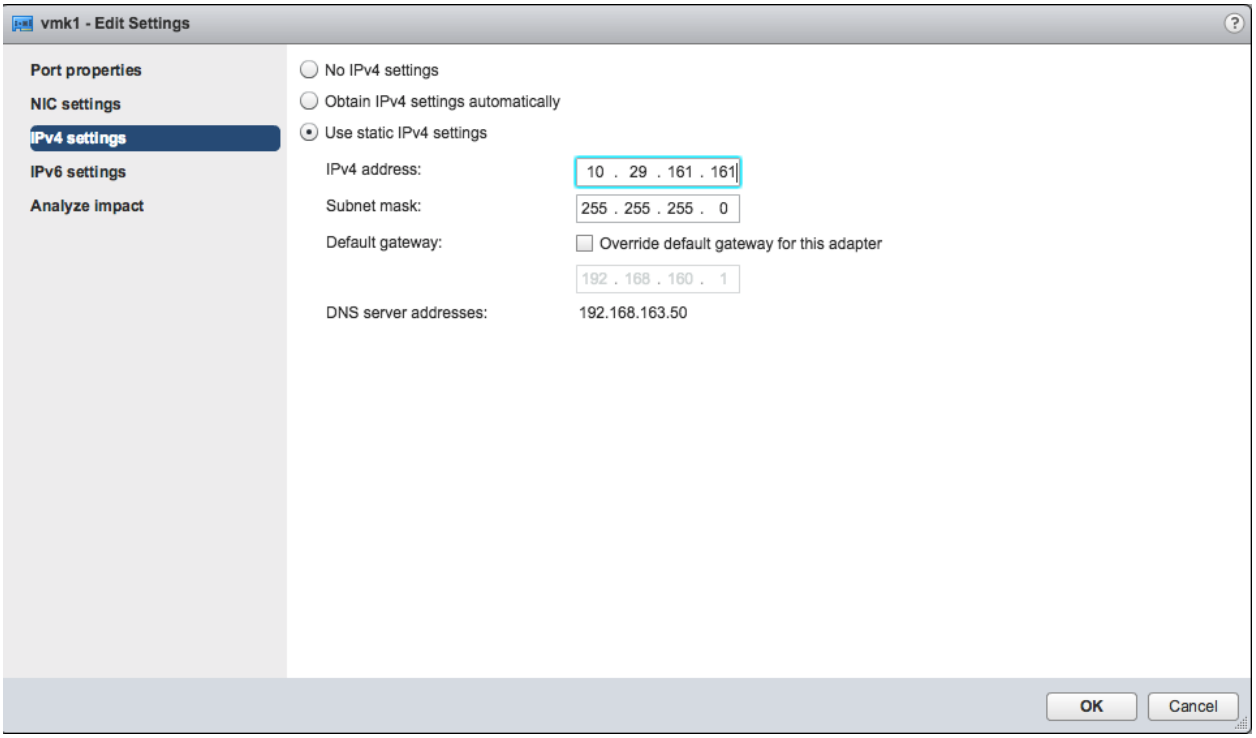

7. Click OK to apply the changes to the vmkernel adapter.

Create iSCSI B vSwitch and vmkernel Adapter

1. Click the Add host networking icon under Virtual switches.

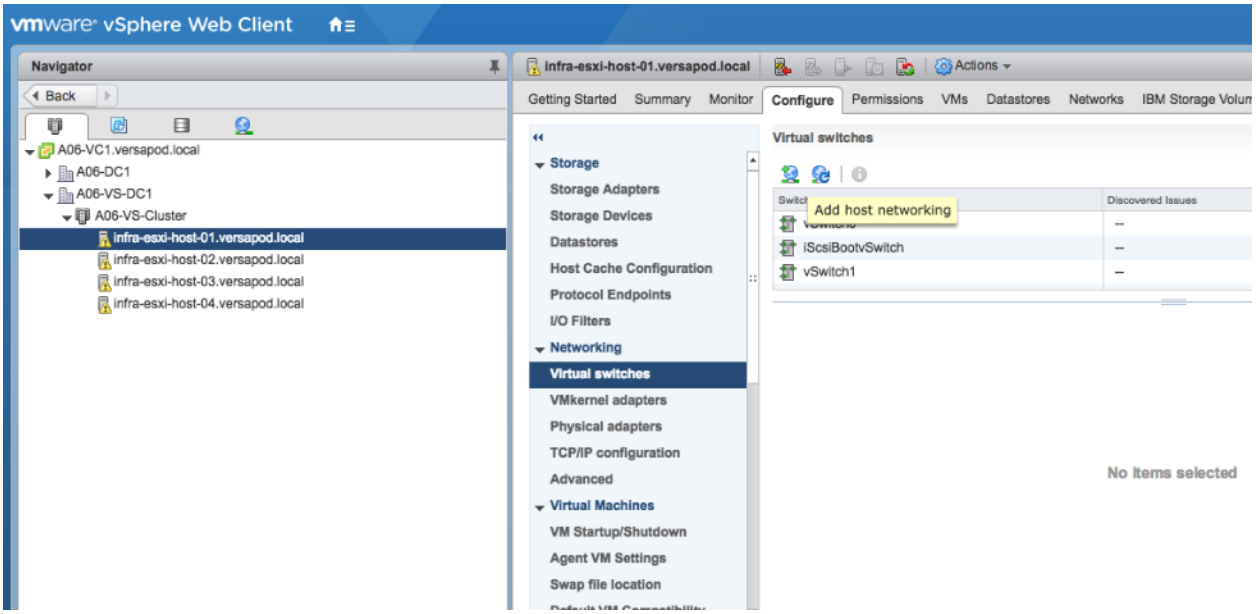

2. Leave VMkernel Network Adapter selected and click Next.

3. Change the Select target device option to New standard switch and click Next.

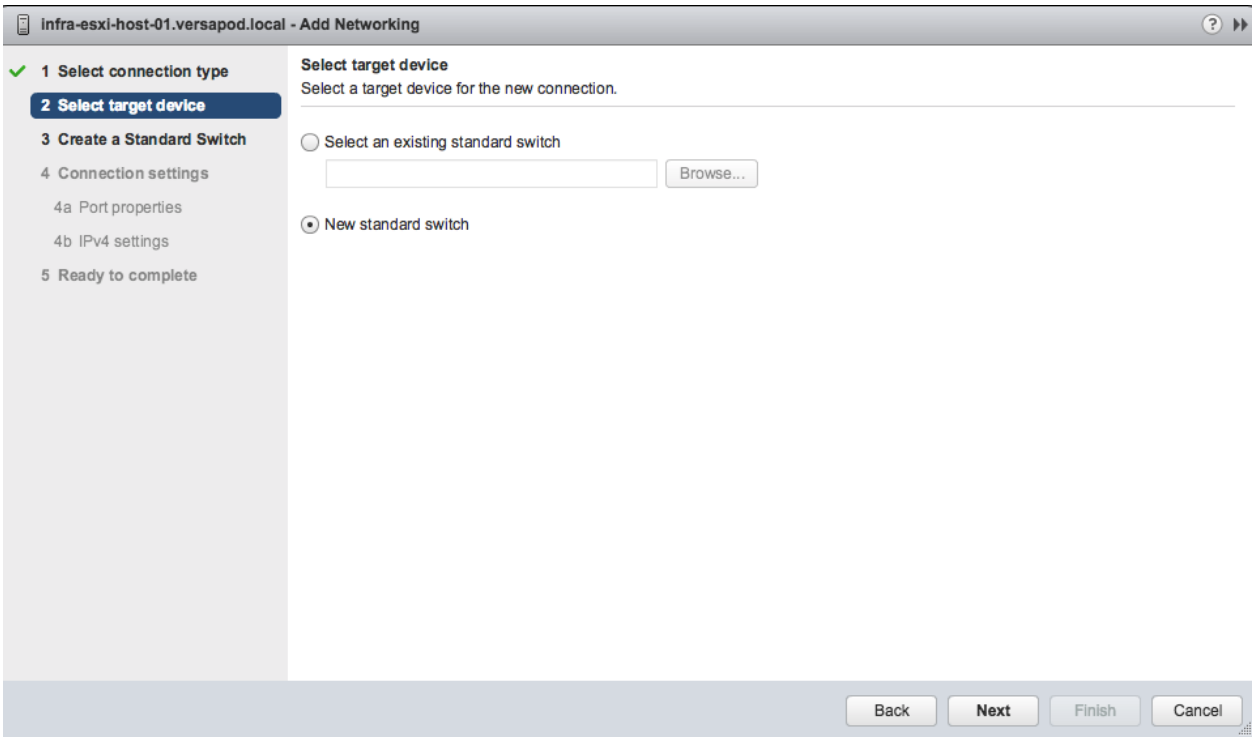

4. Click the green plus icon under Assigned adapters and select vmnic7 from the listed adapters in the resulting window.

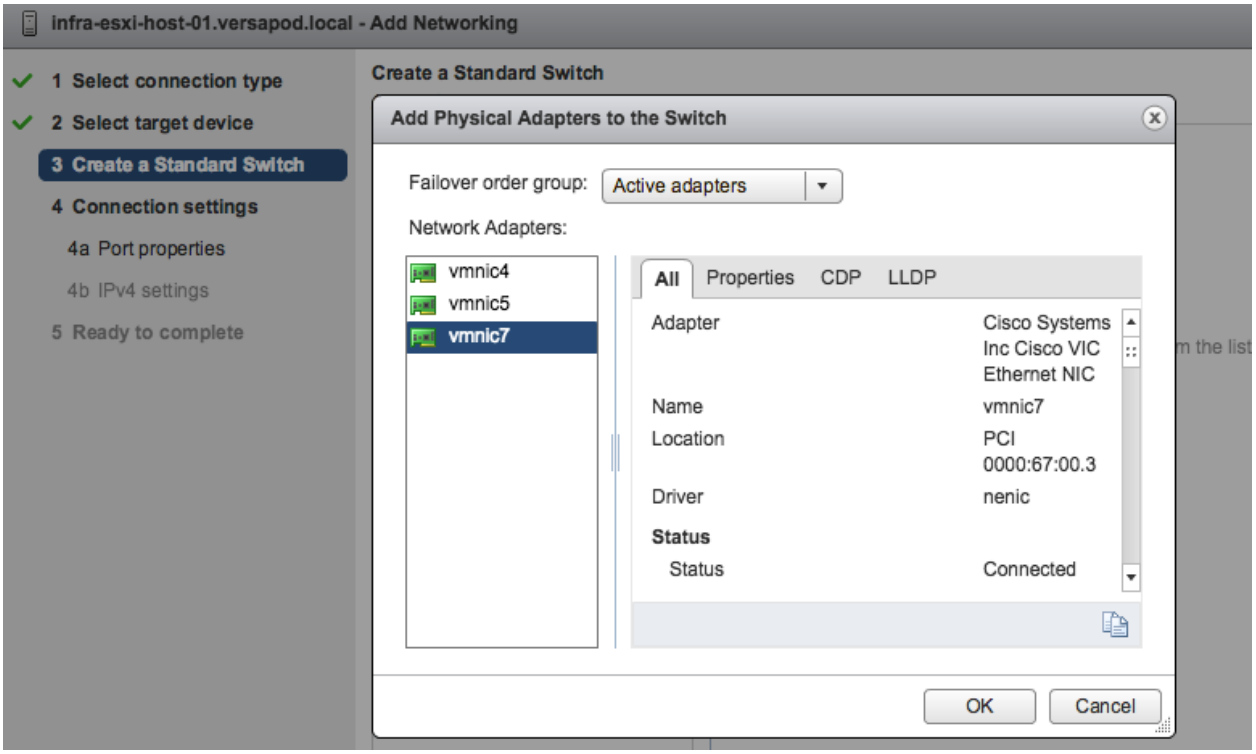

5. Click OK to add the vmnic to the vSwitch and click Next.

6. (Optional) Enter a relevant name for the Network label.

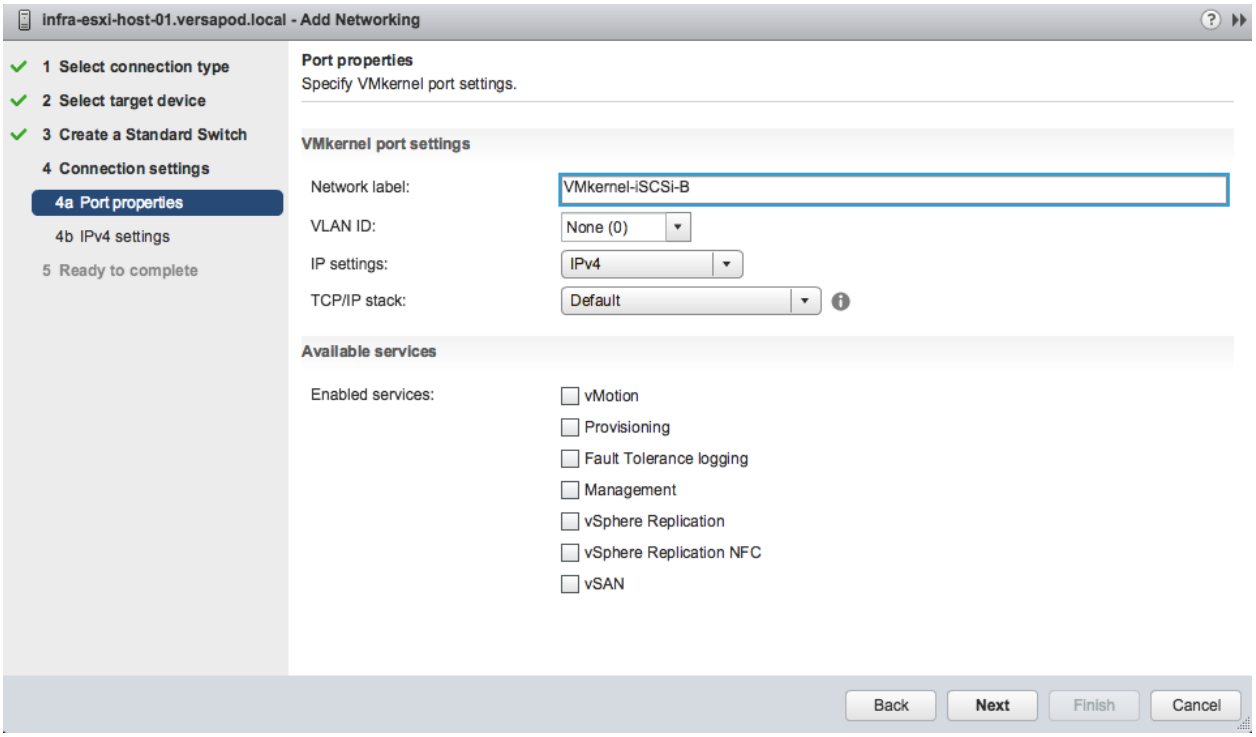

7. Click Next.

8. Change the option for IPv4 settings to Use static IPv4 settings and enter valid IP and subnet mask information that is outside of the UCS iSCSI Pool B.

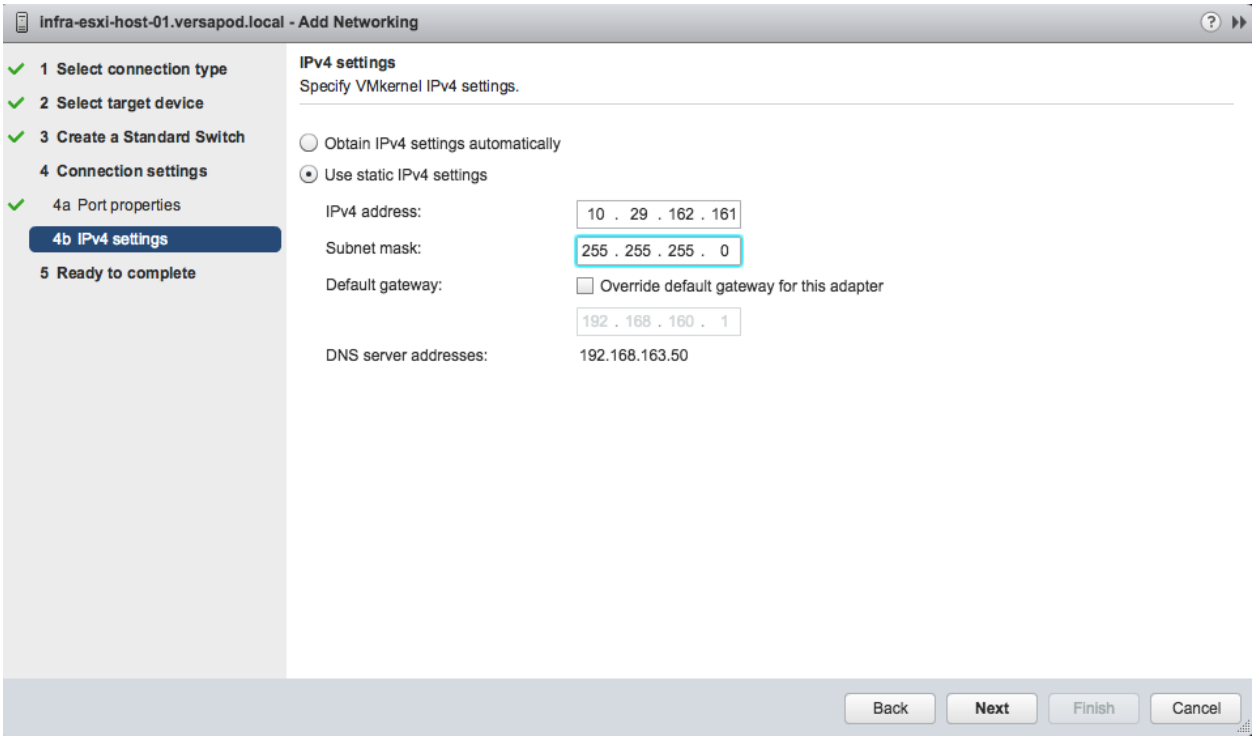

9. Click Next and click Finish in the resulting Summary window.

Setup iSCSI Multipathing

To setup the iSCSI multipathing on the ESXi hosts, complete the following steps:

1. From the vSphere Web Client, select the host and select the Configure tab within the host view.

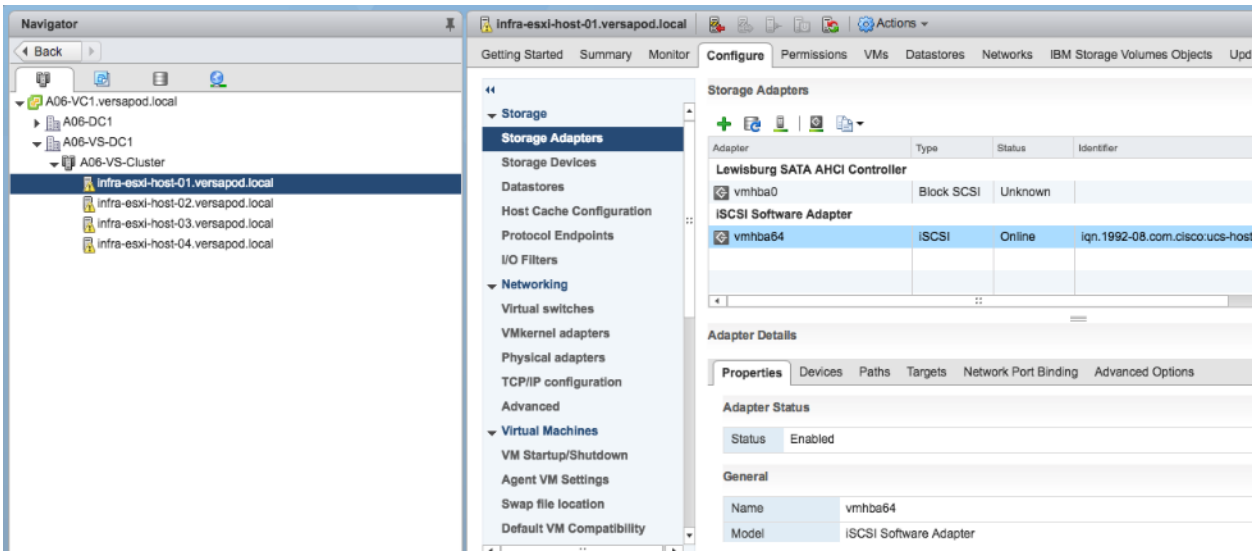

2. Select Storage Adapters from within the Storage section and vmhba64 under the iSCSI Software Adapter listing.

3. Select the Targets tab under the Adapter Details.

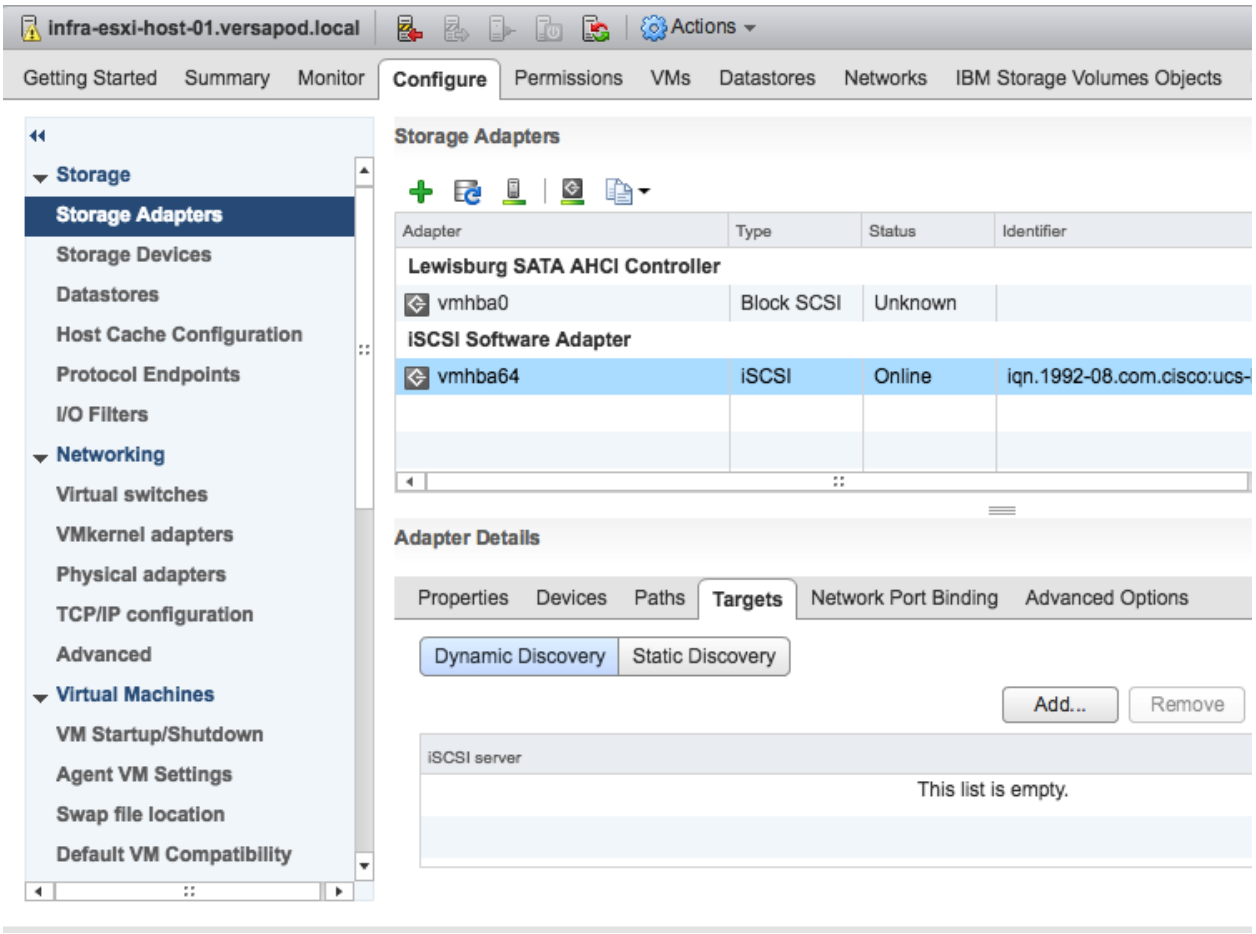

4. With Dynamic Discovery selected, click **Add...**

5. Click Add. Enter the first iSCSI interface IP address for IBM SVC Node 1 storage from Table 7 and click OK.

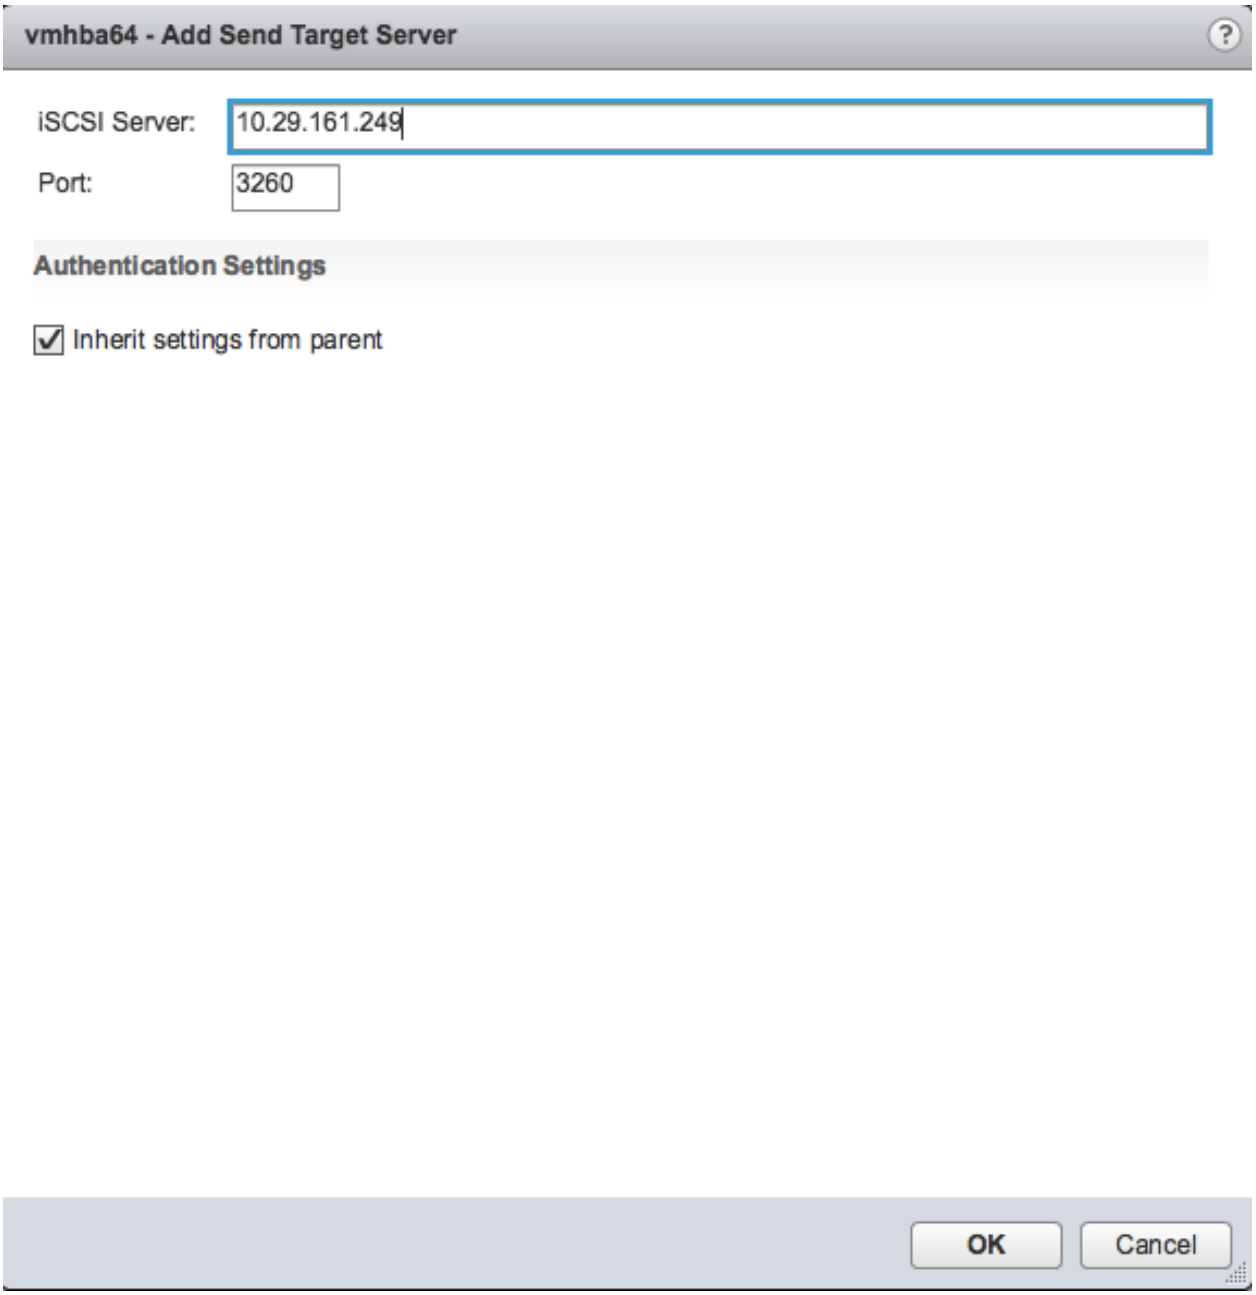

6. Click OK and repeat the previous step to add all IP addresses for all the nodes.

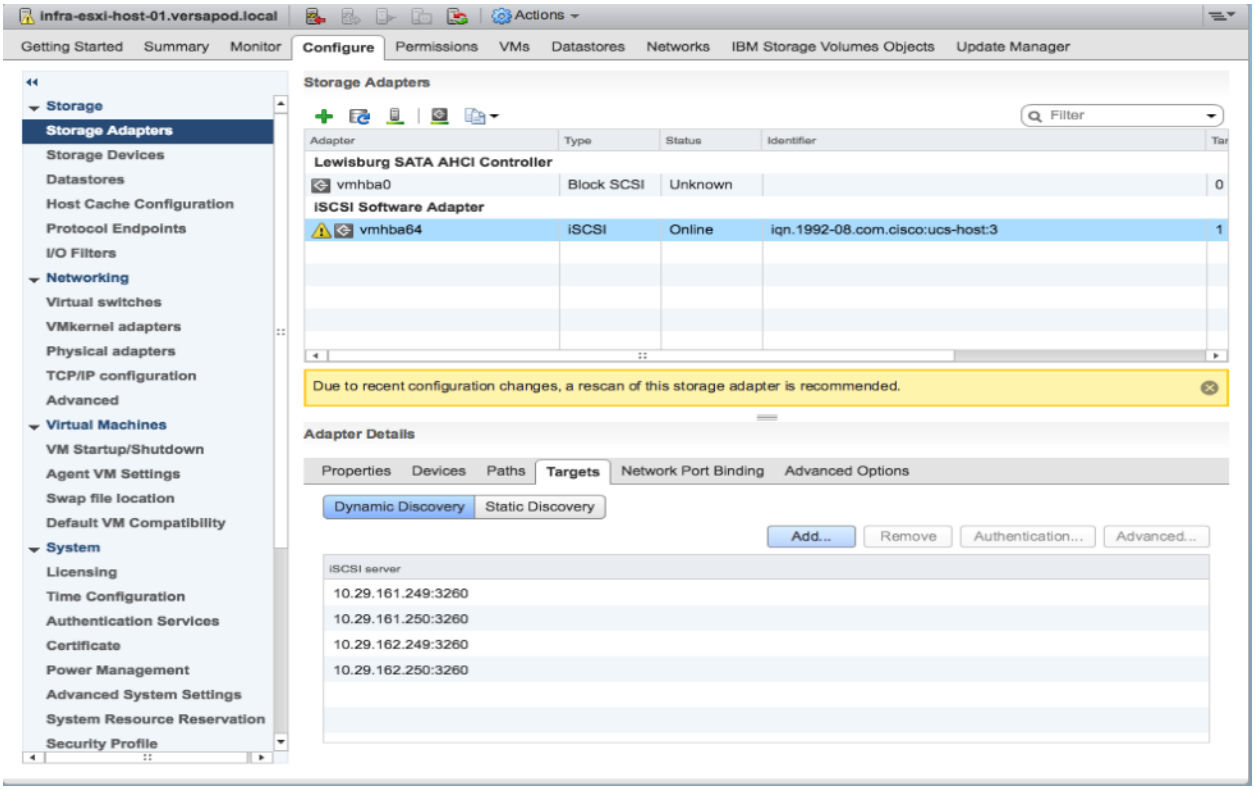

7. Rescan the storage adapters with the third icon at the top of the page.

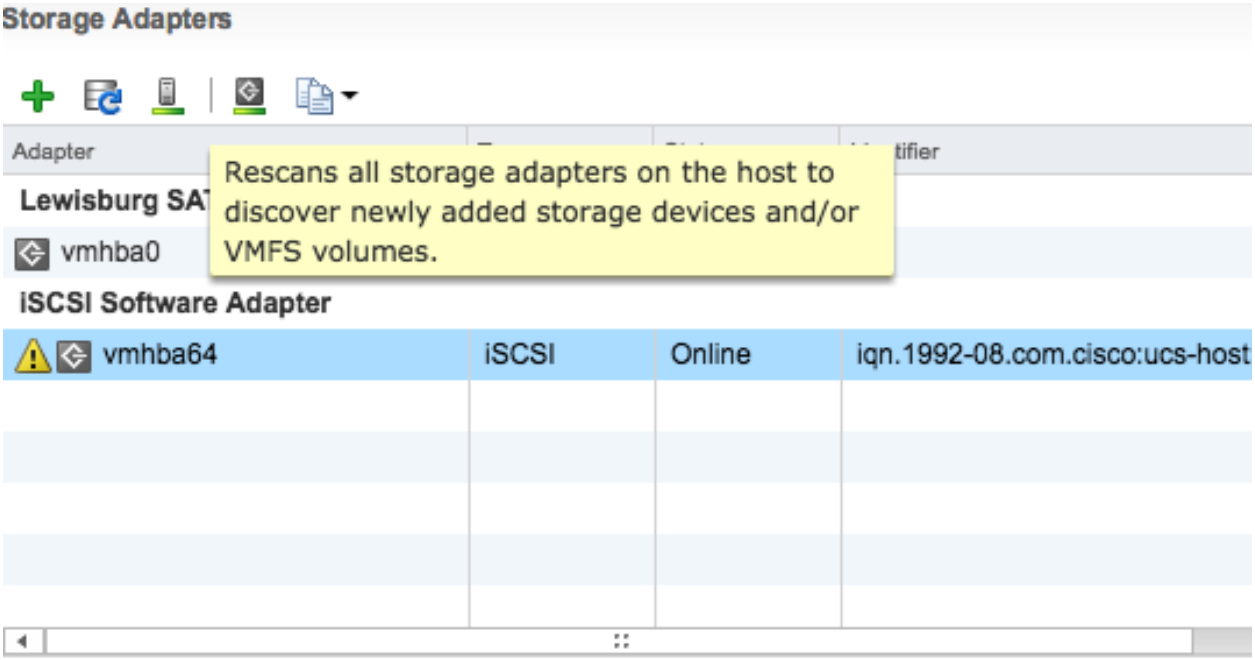

8. Observed Paths should now be four times what it previously was.

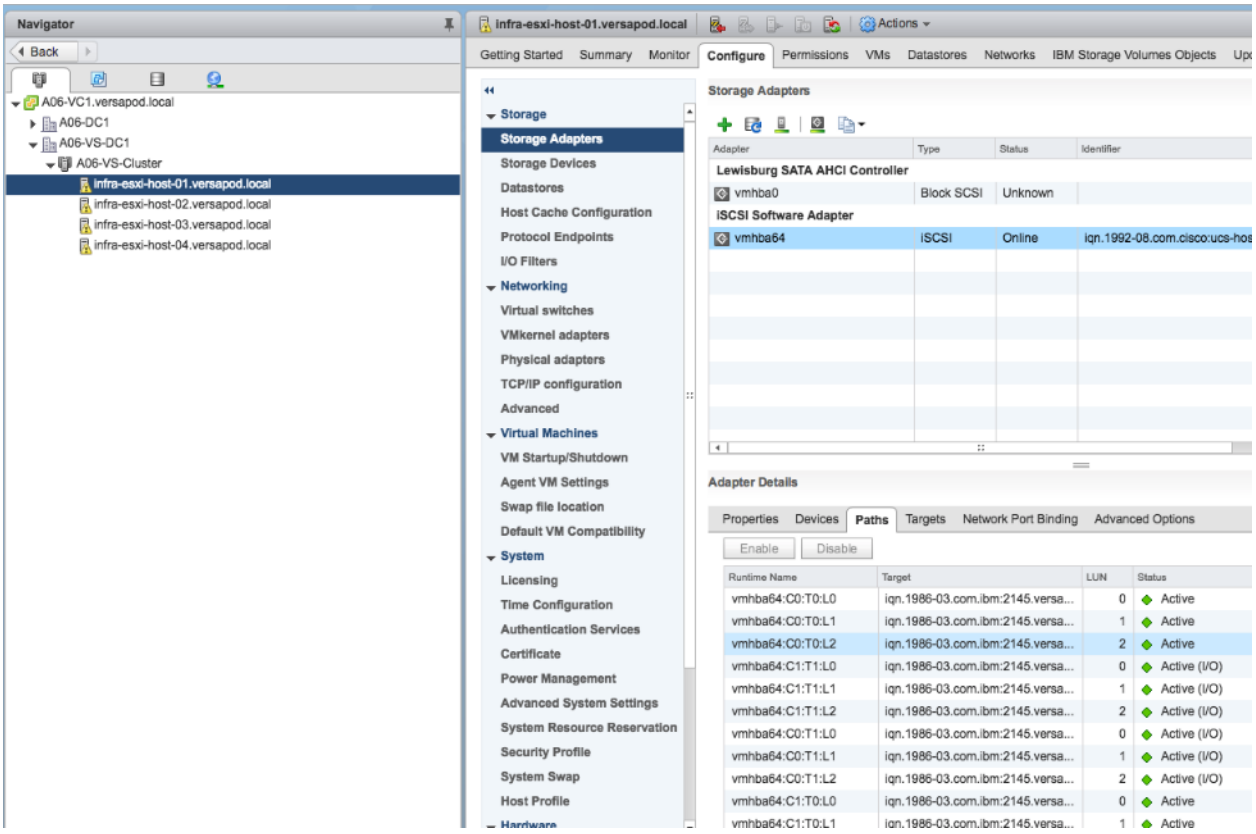

Create a VMware vDS for Application and Production Networks

Production networks will be configured on a VMware vDS to allow additional configuration, as well as consistency between hosts. To configure the VMware vDS, click the right-most icon within the Navigation window, and complete the following steps:

1. Right-click the Datacenter (A06-VS-DC1 in the example screenshot below), select from the drop-down list Distributed Switch > New Distributed Switch.

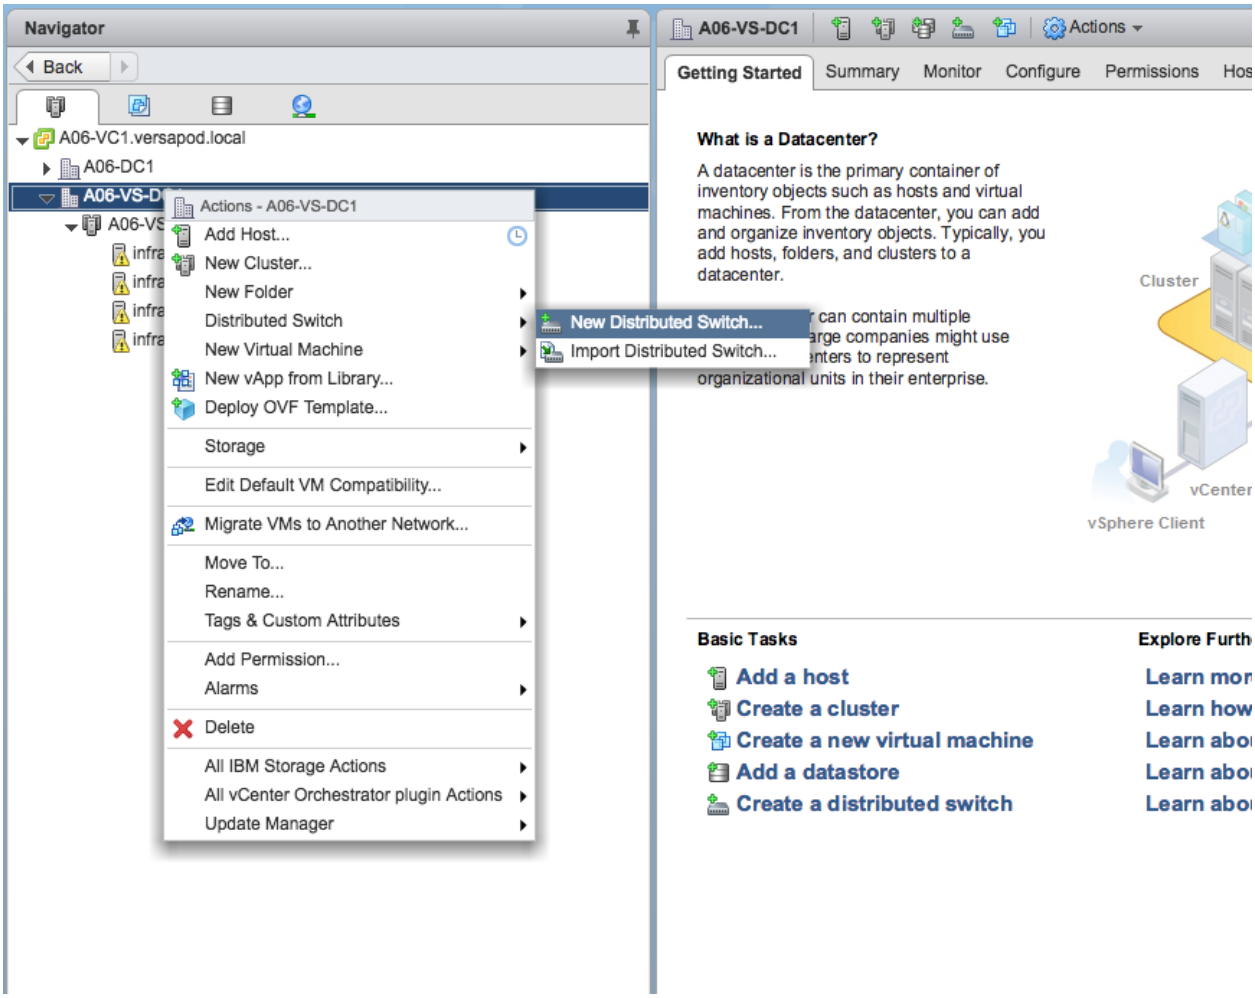

2. Provide a relevant name for the Name field, and click Next.

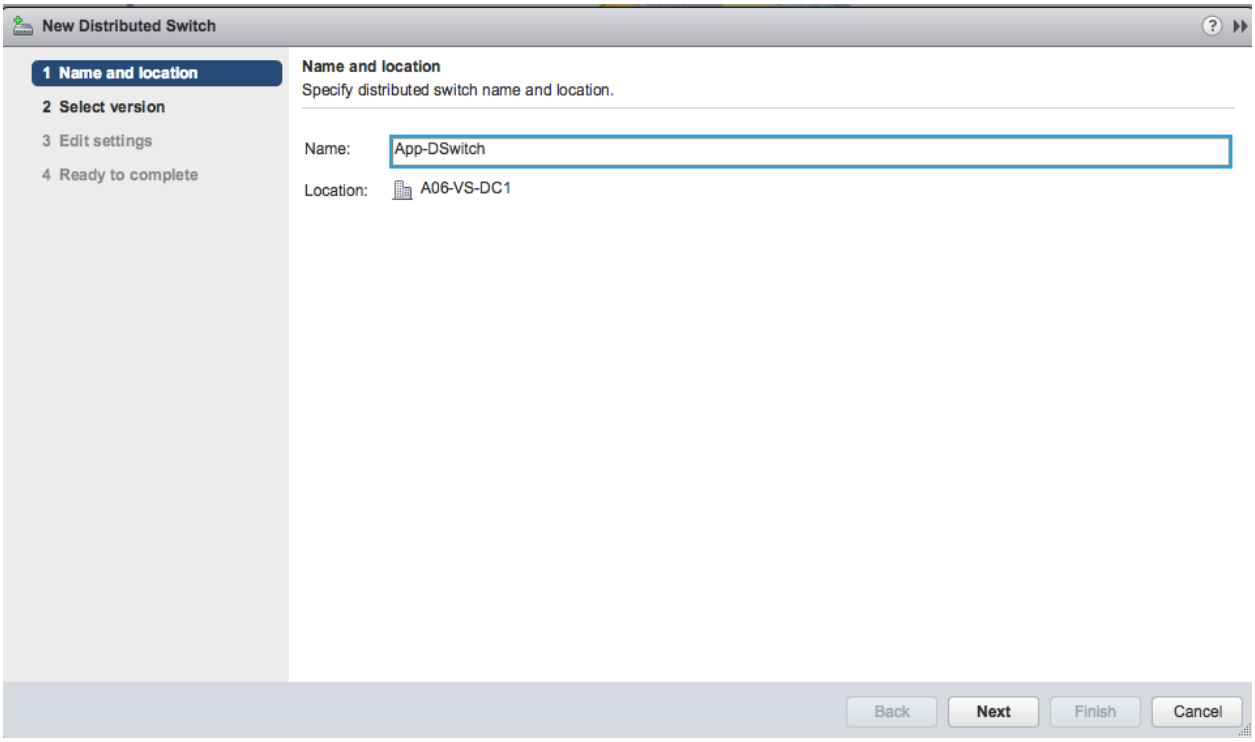

3. Leave the version selected as Distributed switch: 6.0.0, and click Next.

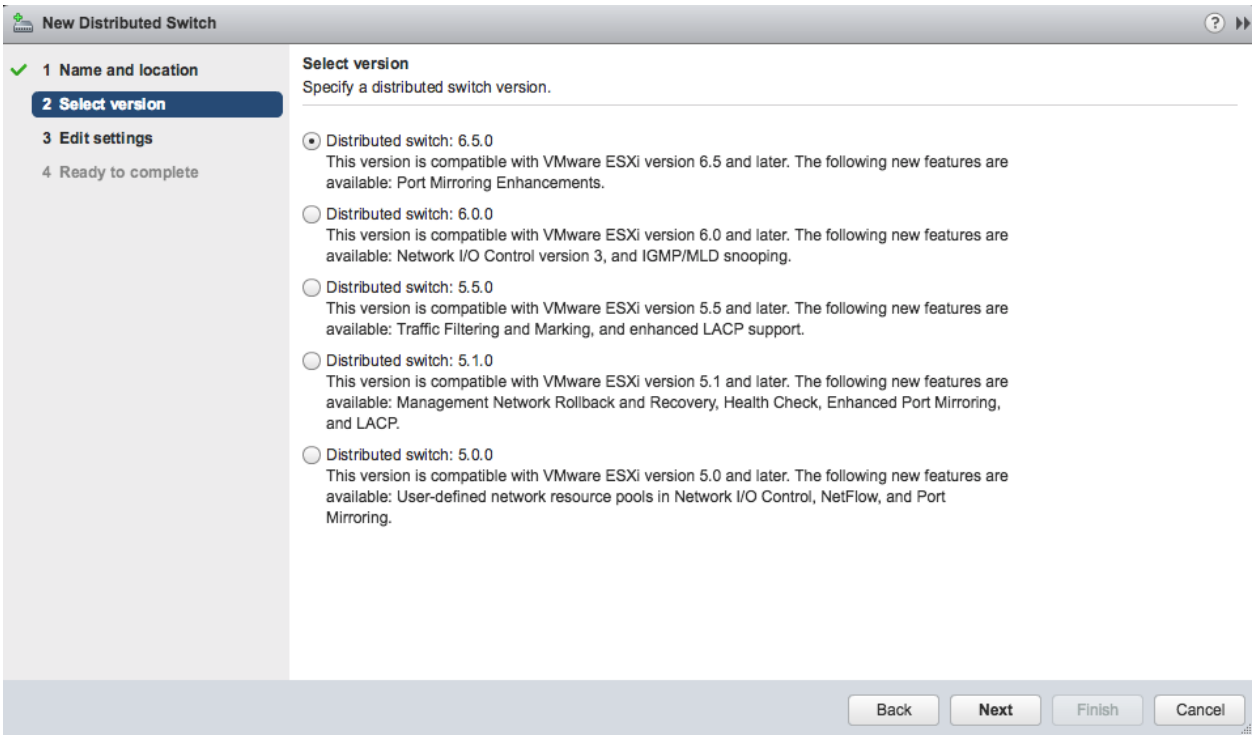

4. Change the Number of uplinks from 4 to 2. If VMware Network I/O Control is to be used for Quality of Service, Leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter VM-Net-3174 for the name of the default Port group to be created. Click Next.

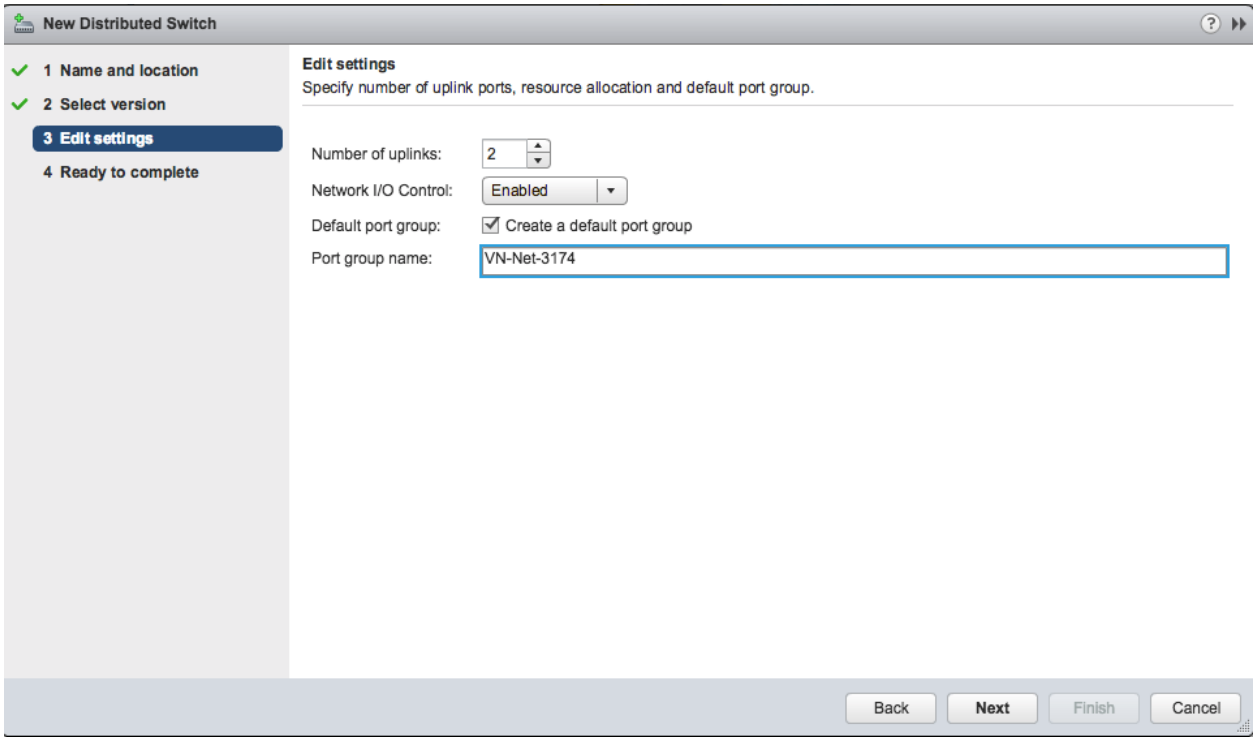

5. Review the summary in the Ready to complete page and click Finish to create the vDS.

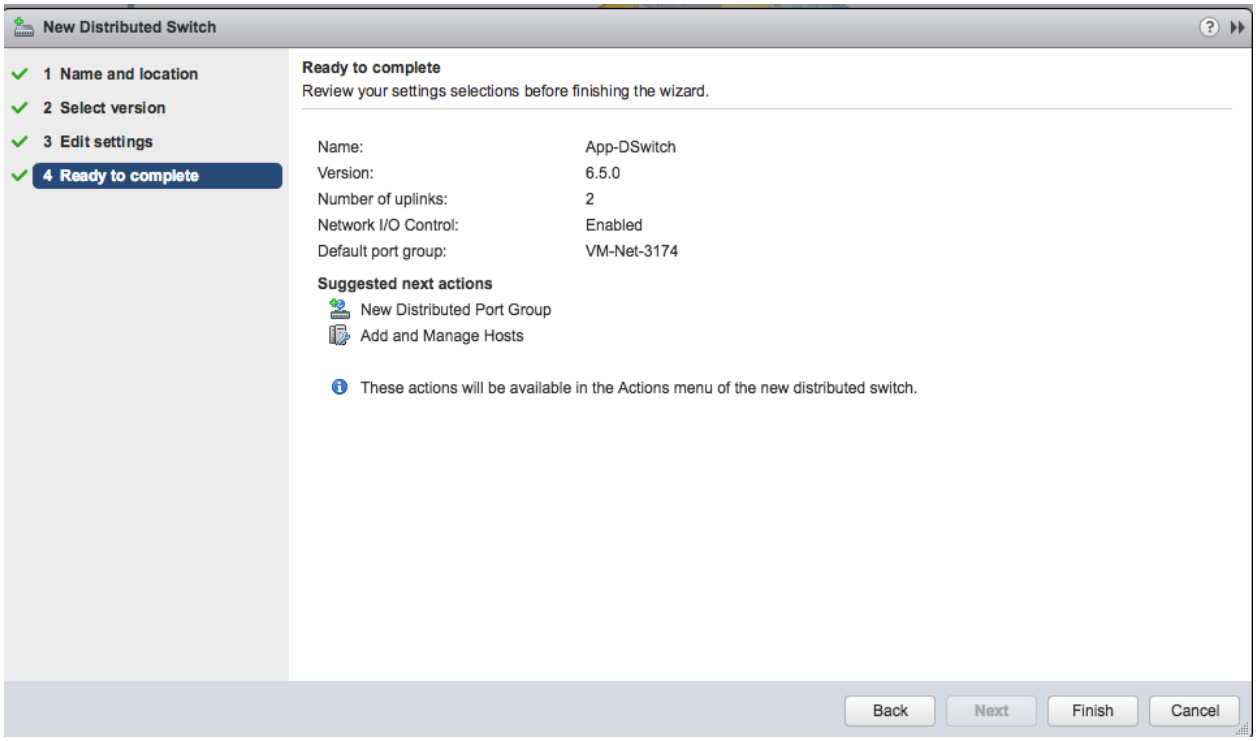

6. Right-click the newly created App-DSwitch vDS, and select Settings > Edit Settings...

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 VM-Net-3174 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 App-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 App-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 the ESXi Hosts to the vDS

With the vDS and the distributed port groups created within the vDS in place, the ESXi hosts will be added to the vDS.

To add the ESXi Hosts to the vDS, complete the following steps:

1. Within the Networking sub-tab of Hosts and Clusters of the Navigator window, right-click the vDS and select **Add and Manage Hosts...**

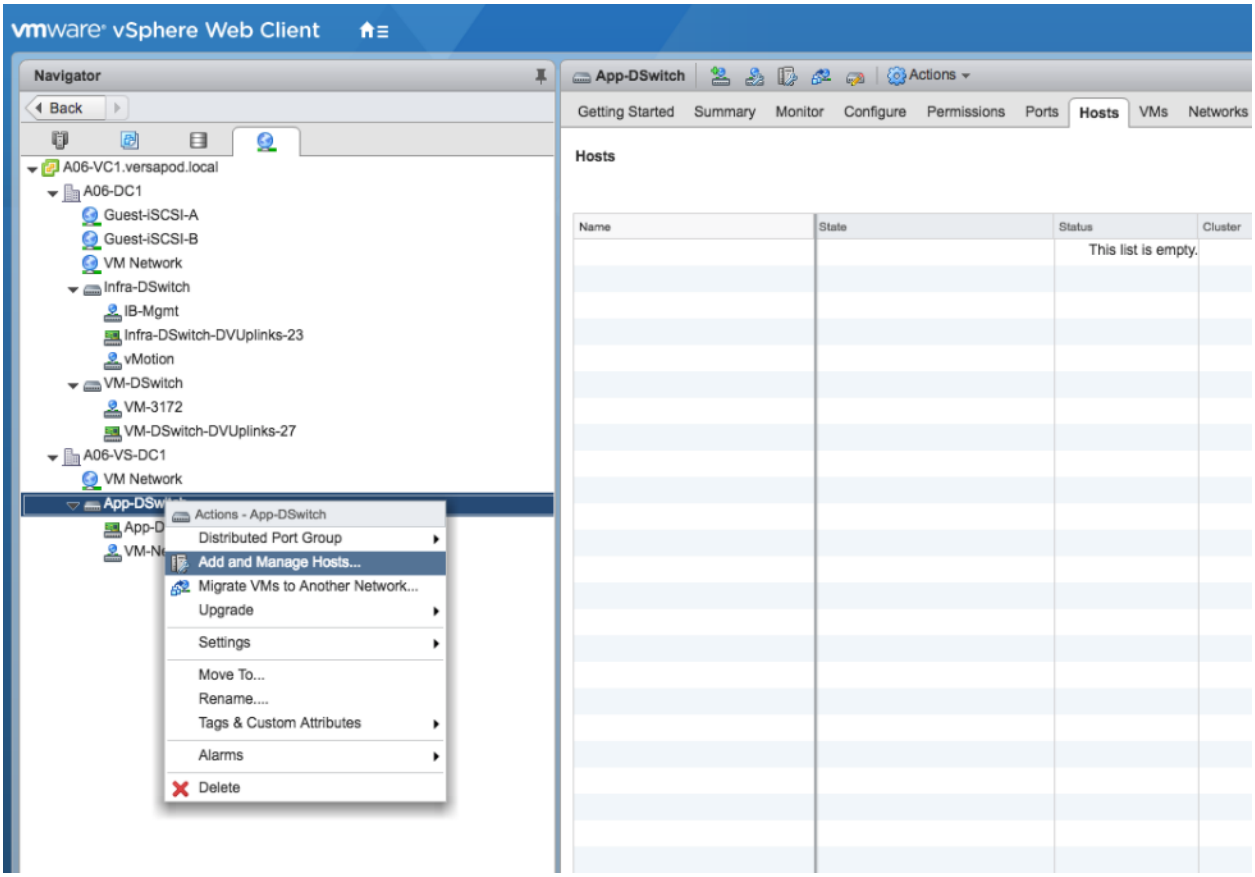

2. Leave Add hosts selected and click Next.

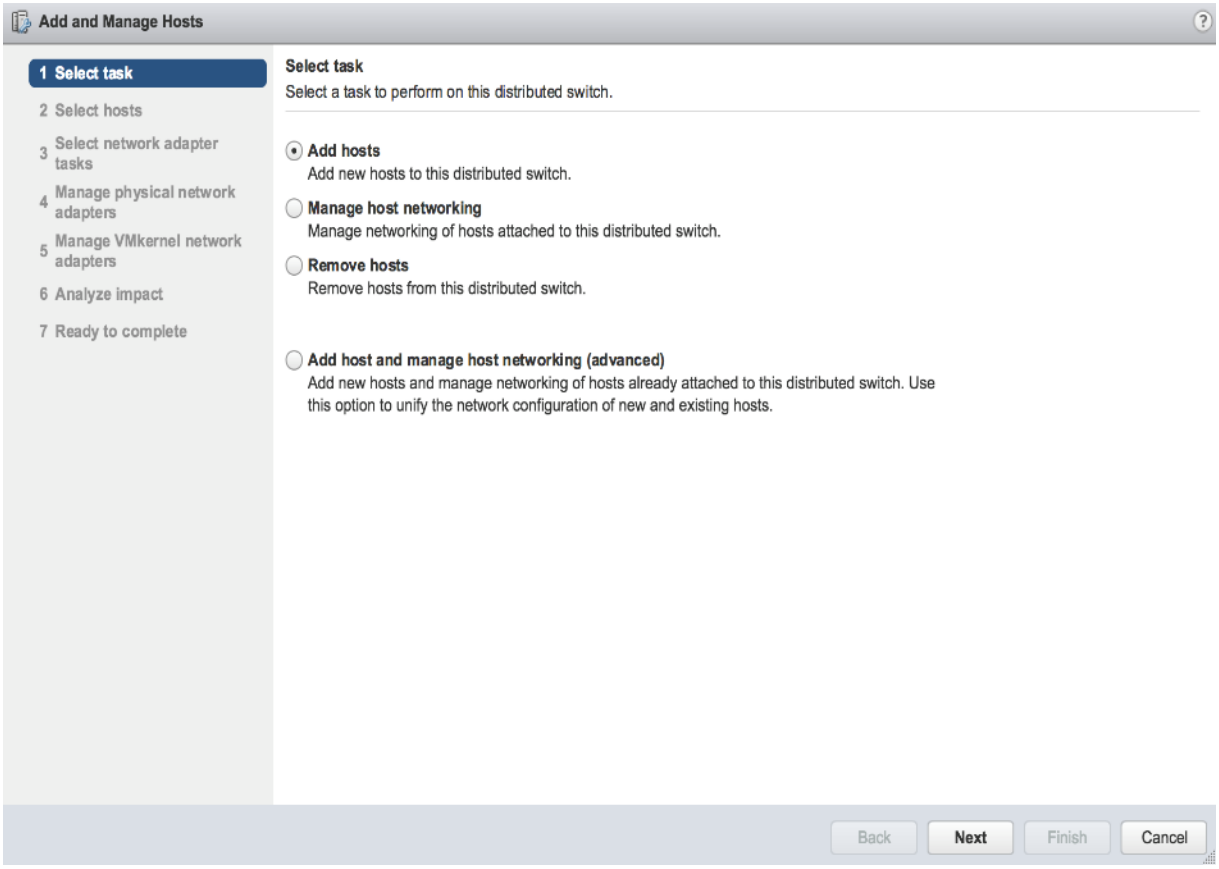

3. Click the green + icon next to New hosts...

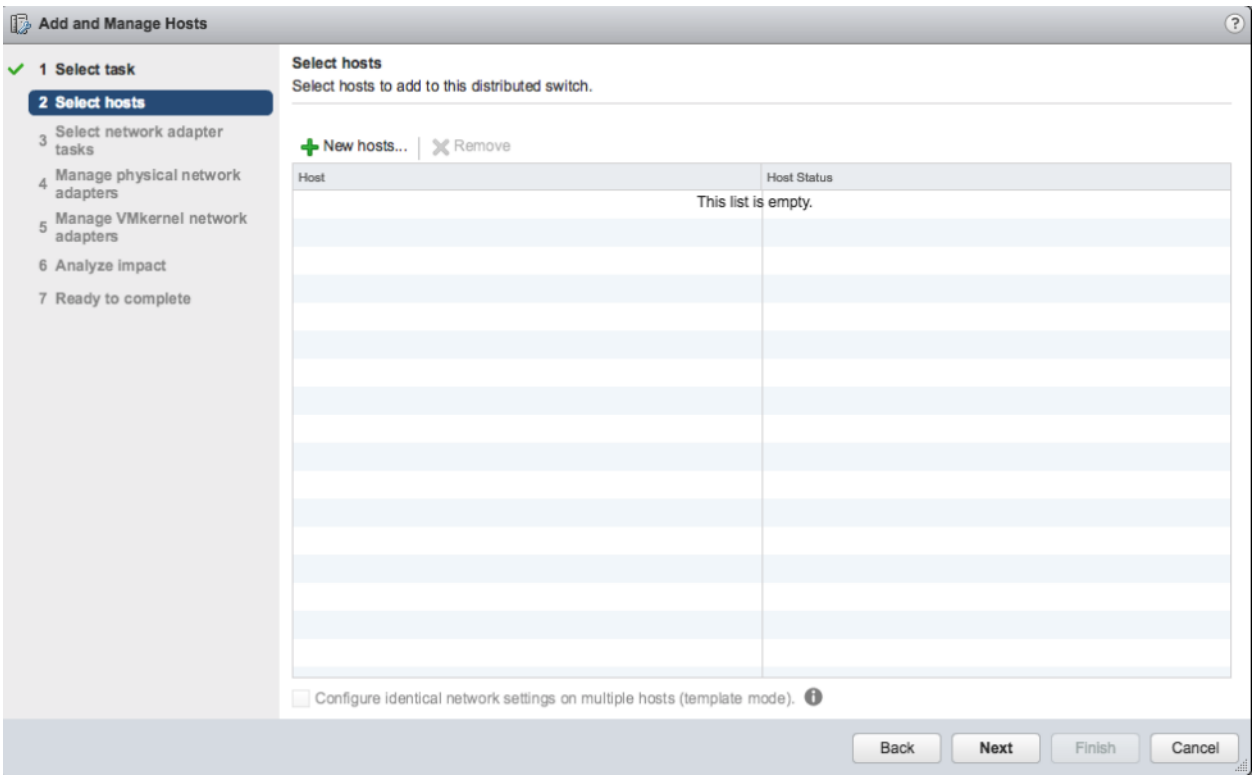

4. In the Select new hosts pop-up that appears, select the hosts to be added, and click OK to begin joining them to the vDS.

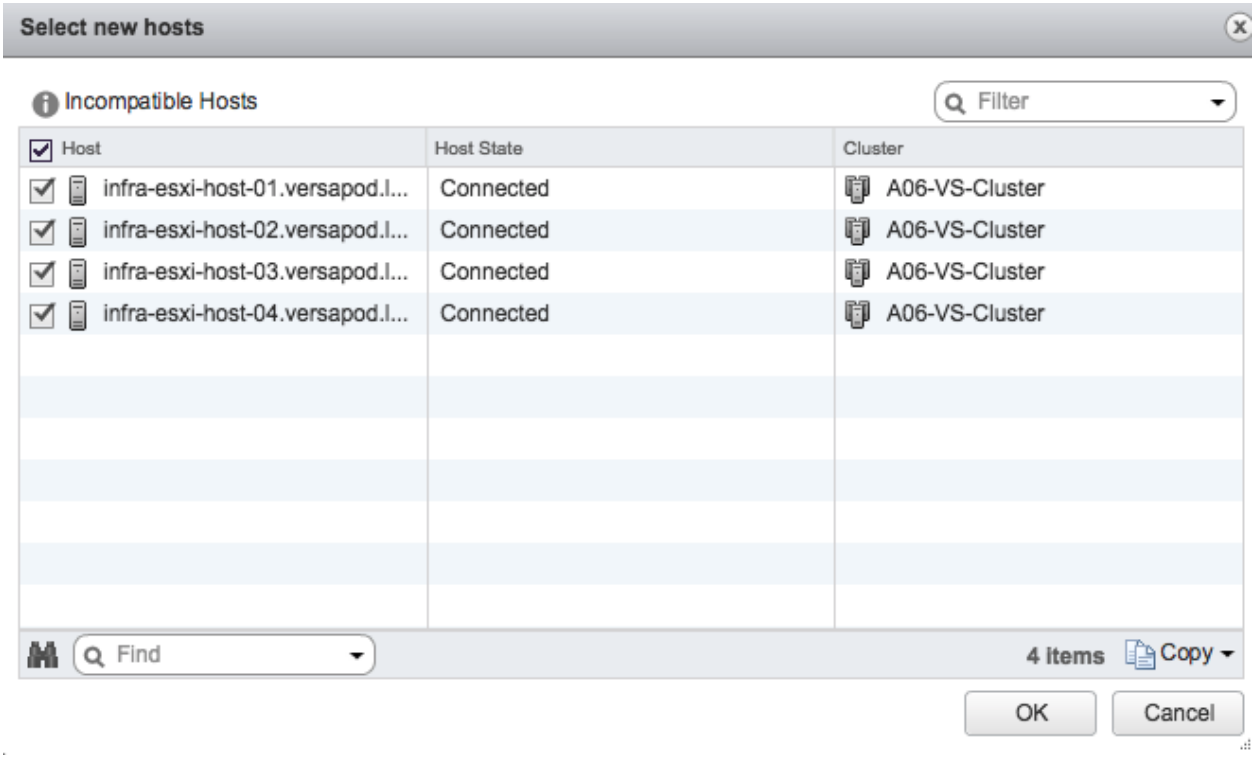

5. Click Next.

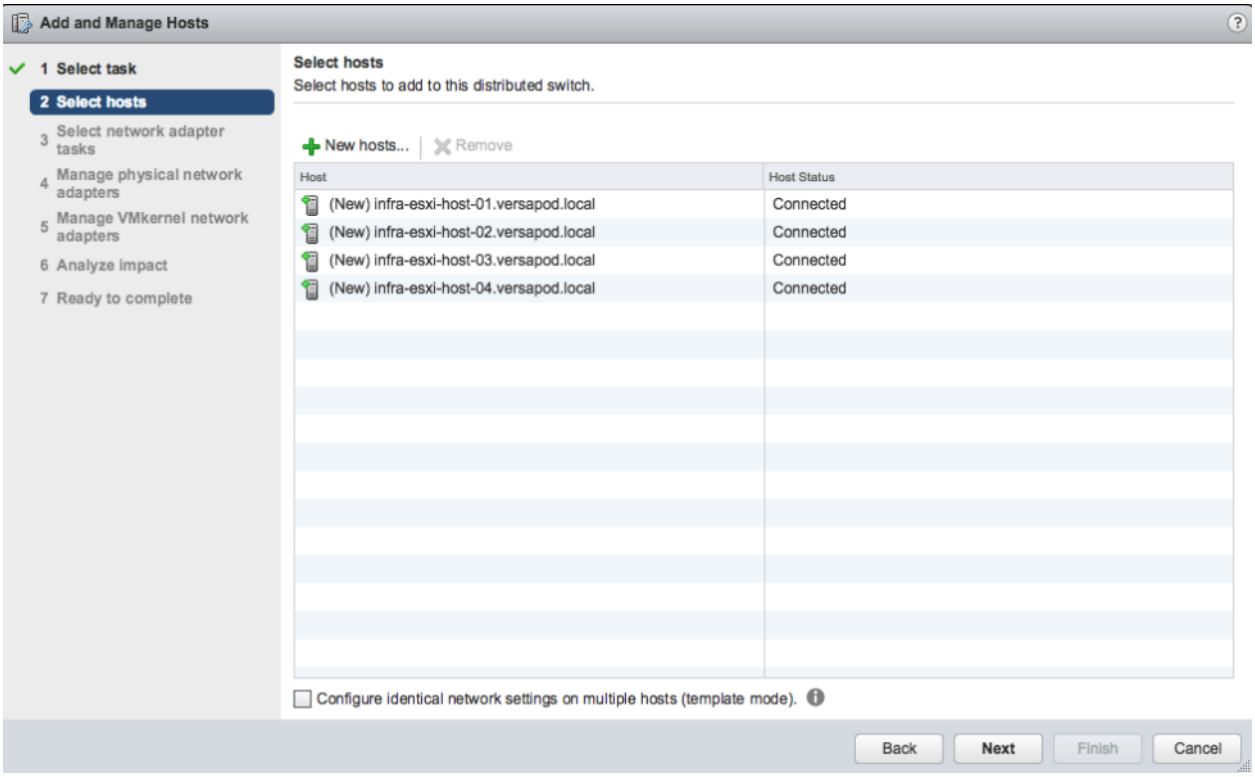

6. Unselect Manage VMkernel adapters (template mode) if it is selected, and click Next.

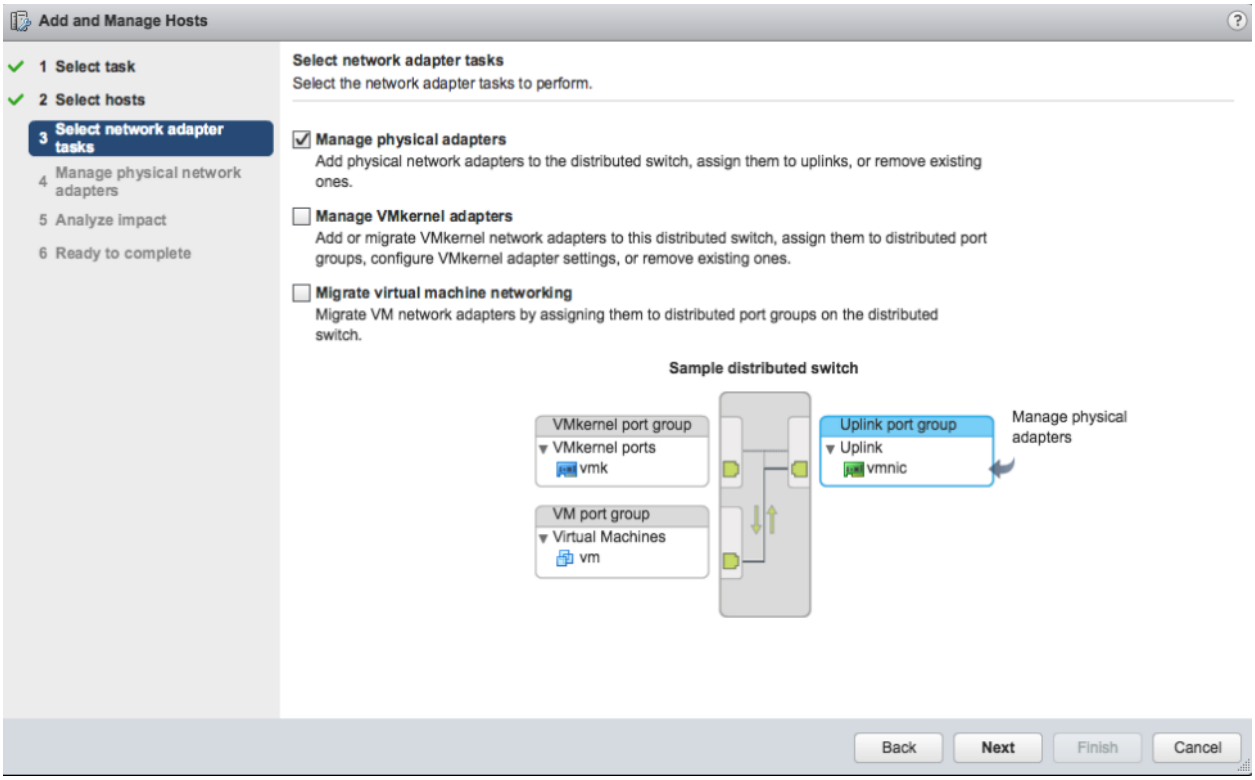

7. For each vmnic (vmnic4 and vmnic5) to be assigned from the Host/Physical Network Adapters column, select the vmnic and click the Assign uplink.

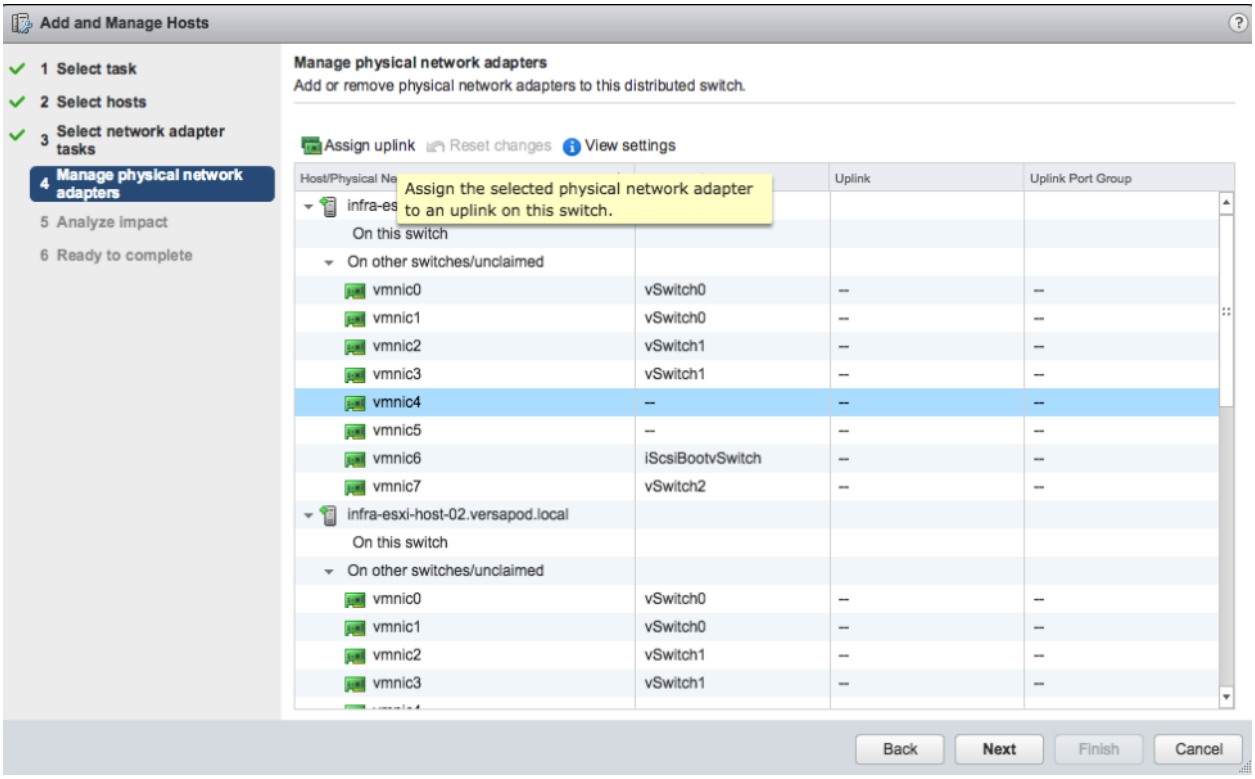

8. Assign the first to Uplink 1 and assign the second to Uplink 2.

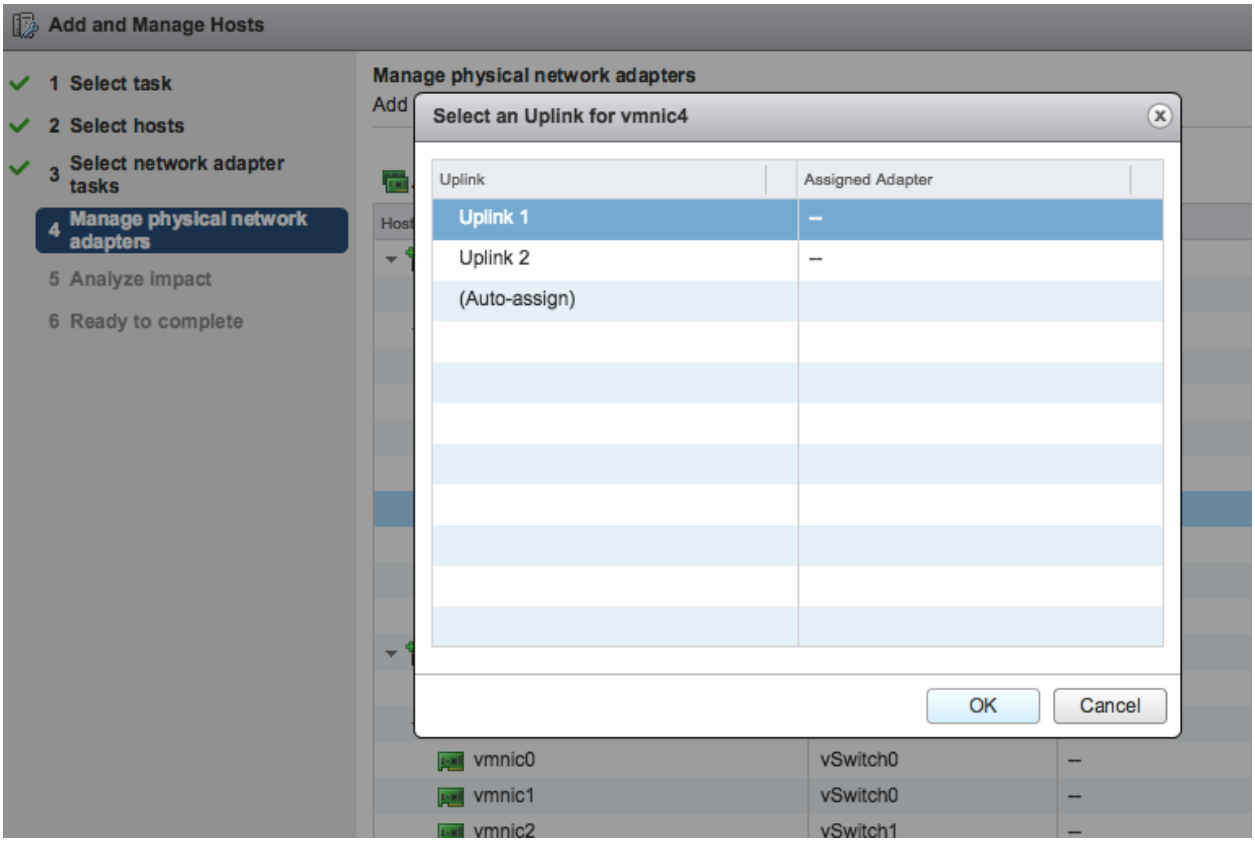

9. Repeat this process until all vmnics (2 per host) have been assigned.

10. Click Next.

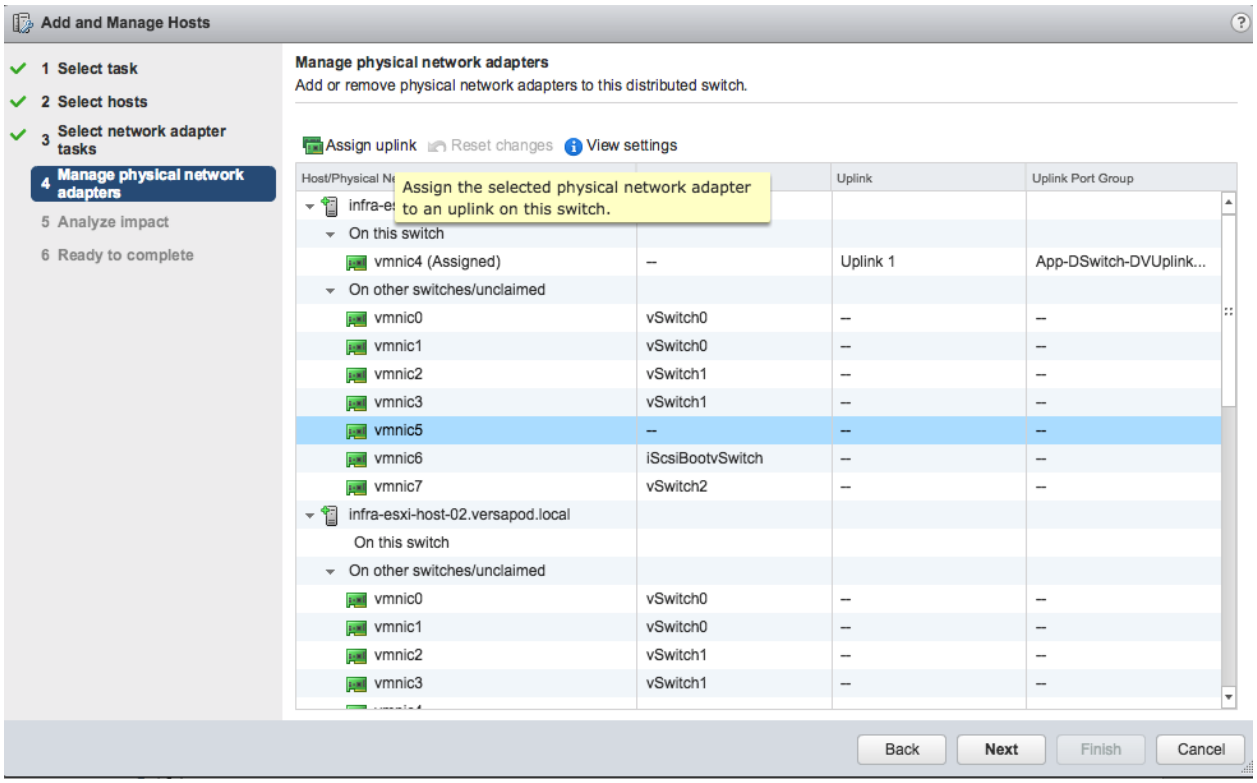

11. Proceed past the Analyze impact screen if no issues appear.

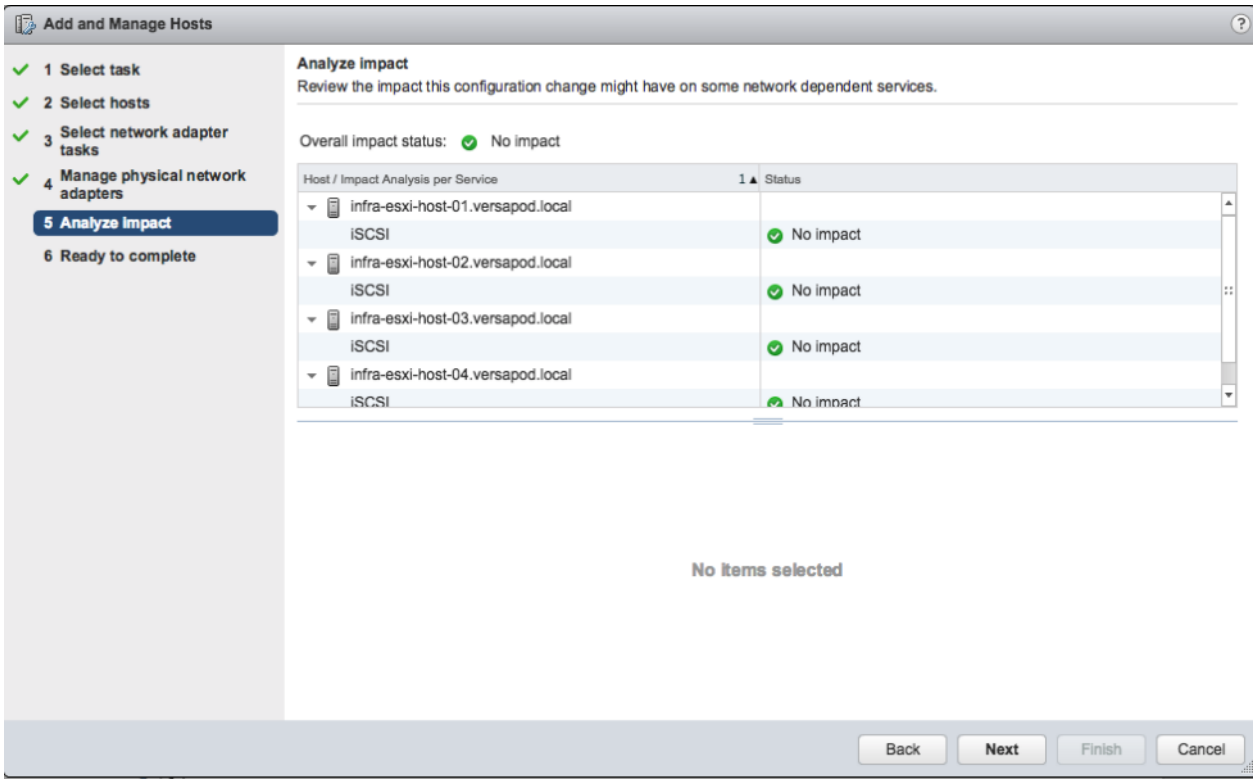

12. Review the Ready to complete summary and click Finish to add the hosts to the vDS.

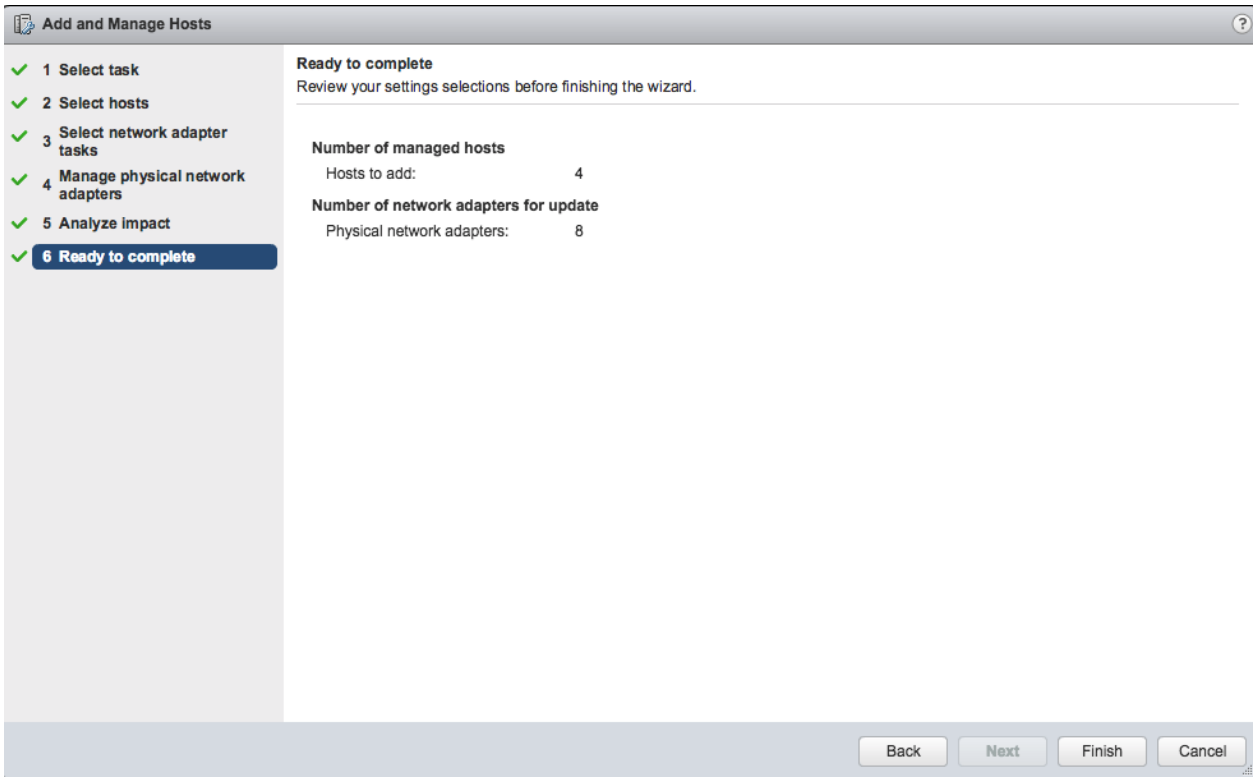

About the Authors

Sreenivasa Edula, Technical Marketing Engineer, UCS Data Center Solutions Engineering, Cisco Systems, Inc.

Sreeni has over 17 years of experience in Information Systems with expertise across Cisco Data Center technology portfolio, including DC architecture design, virtualization, compute, network, storage and cloud computing.

Warren Hawkins, Virtualization Test Specialist for IBM Spectrum Virtualize, IBM

Working as part of the development organization within IBM Storage, Warren Hawkins is also a speaker and published author detailing best practices for integrating IBM Storage offerings into virtualized infrastructures. Warren has a background in supporting Windows and VMware environments working in second-line and third-line support in both public and private sector organizations. Since joining IBM in 2013, Warren has played a crucial part in customer engagements and, using his field experience, has established himself as the Test Lead for the IBM Spectrum Virtualize™ product family, focusing on clustered host environments.