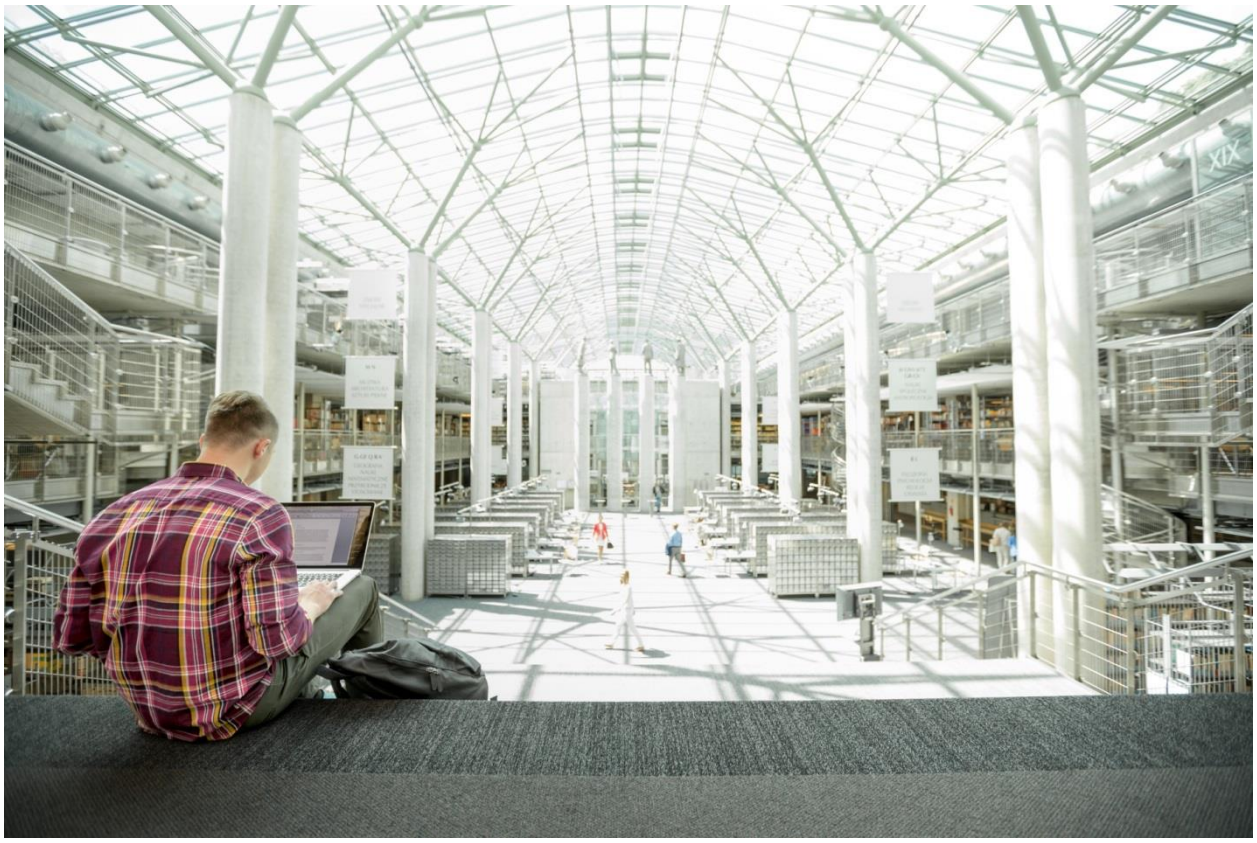

FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1

Deployment Guide for FlexPod Datacenter with Cisco UCS Manager 3.1 and VMware vSphere 6.0 U1

Last Updated: August 26, 2016

About Cisco Validated Designs

The 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.

Table of Contents

About Cisco Validated Designs 2

Executive Summary 9

Solution Overview 10

Introduction 10

Audience 10

Purpose of this Document 10

What's New? 10

Solution Design 11

Architecture 11

Physical Topology 11

Deployment Hardware and Software 13

Software Revisions 13

Configuration Guidelines 13

Physical Infrastructure 19

FlexPod Cabling 19

Network Switch Configuration 28

Physical Connectivity 28

FlexPod Cisco Nexus Base 28

Set Up Initial Configuration 28

FlexPod Cisco Nexus Switch Configuration 30

Enable Licenses 30

Set Global Configurations 31

Create VLANs 31

Add NTP Distribution Interface 32

Add Individual Port Descriptions for Troubleshooting 32

Create Port Channels 35

Configure Port Channel Parameters 37

Configure Virtual Port Channels 39

Uplink into Existing Network Infrastructure 41

Storage Configuration 42

AFF80XX Series Controllers 42

NetApp Hardware Universe 42

Controllers 42

Disk Shelves 42
Clustered Data ONTAP 8.3.2 43
Complete the Configuration Worksheet 43
Configure ONTAP Nodes 43
Log In to the Cluster 52
Zero All Spare Disks 53
Set Onboard Unified Target Adapter 2 Port Personality 53
Set Auto-Revert on Cluster Management 53
Set Up Management Broadcast Domain 54
Set Up Service Processor Network Interface 54
Create Aggregates 54
Verify Storage Failover 55
Disable Flow Control on UTA2 Ports 56
Disable Unused FC Capability on CNA Ports 56
Configure NTP 56
Configure SNMP 57
Configure AutoSupport 57
Enable Cisco Discovery Protocol 58
Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP 58
Create Interface Groups 58
Create VLANs 58
Create Storage Virtual Machine 59
Create Load-Sharing Mirrors of SVM Root Volume 59
Create Block Protocol (iSCSI, FC) Service 59
Configure HTTPS Access 60
Configure NFSv3 61
Create FlexVol Volumes 61
Create Boot LUNs 61
Schedule Deduplication 62
Create iSCSI LIFs 62
Create FCP LIFs 63
Create NFS LIF 63
Add Infrastructure SVM Administrator 63
Server Configuration 65
Cisco UCS Base Configuration 65

Perform Initial Setup of Cisco UCS 6332-16UP and 6248UP Fabric Interconnects for FlexPod Environments 65
Cisco UCS Setup 66
Log in to Cisco UCS Manager 66
Upgrade Cisco UCS Manager Software to Version 3.1(1h) 67
Anonymous Reporting 67
Add Block of IP Addresses for KVM Access 67
Synchronize Cisco UCS to NTP 68
Edit Chassis Discovery Policy 70
Enable FC Switching 71
Configure Unified Ports 72
Enable FC Storage Ports 75
Enable Server and Uplink Ports 75
Acknowledge Cisco UCS Chassis and FEX 76
Create Uplink Port Channels to Cisco Nexus Switches 77
Create a WWNN Pool for FC Boot 79
Create WWPN Pools 81
Create VSANs 85
Assign VSANs to FC Storage Ports 89
Create Storage Connection Policies for FC Zoning 93
Create vHBA Templates 97
Create SAN Connectivity Policy 99
Create MAC Address Pools 102
Create IQN Pools for iSCSI Boot 106
Create IP Pools for iSCSI Boot 108
Create UUID Suffix Pool 111
Create Server Pool 112
Create VLANs 113
Create Host Firmware Package 117
Set Jumbo Frames in Cisco UCS Fabric 118
Create Local Disk Configuration Policy (Optional) 119
Create Network Control Policy for Cisco Discovery Protocol 120
Create Power Control Policy 121
Create Server Pool Qualification Policy (Optional) 122
Create Server BIOS Policy 123
Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts 124

Update the Default Maintenance Policy....125
Create vNIC Templates....126
Create LAN Connectivity Policy....136
Create vMedia Policy for VMware ESXi 6.0 U1b Install Boot....142
Create Boot Policies (iSCSI Boot)....144
Create Boot Policies (FC Boot)....145
Create Service Profile Template (iSCSI Boot)....150
Create Service Profile Template (FC Boot)....162
Create Service Profiles....169
Add More Servers to FlexPod Unit....171
Gather Necessary Information....171
Storage Configuration – Boot LUNs and Igroups....173
Clustered Data ONTAP Boot Storage Setup....173
Create igroups....173
Map Boot LUNs to igroups....173
VMware vSphere 6.0 U1 Setup....174
VMware ESXi 6.0 U1....174
Download Cisco Custom Image for ESXi 6.0 U1....174
Log in to Cisco UCS 6300/6200 Fabric Interconnect....174
Set Up VMware ESXi Installation....175
Install ESXi....175
Set Up Management Networking for ESXi Hosts....176
Download VMware vSphere Client....178
Log in to VMware ESXi Hosts by Using VMware vSphere Client....179
Set Up VMkernel Ports and Virtual Switch....179
Setup iSCSI Multipathing....187
Install VMware Drivers for the Cisco Virtual Interface Card (VIC)....188
Mount Required Datastores....193
Configure NTP on ESXi Hosts....196
Move VM Swap File Location....197
VMware vCenter 6.0 U1b....198
Install the Client Integration Plug-in....198
Building the VMware vCenter Server Appliance....199
Setting Up VMware vCenter Server....208
ESXi Dump Collector Setup for iSCSI-Booted Hosts....218

FlexPod Cisco Nexus 1110-X and 1000V vSphere 219
Configure CIMC Interface on Both Cisco Nexus 1110-Xs 219
Configure Serial over LAN for Both Cisco Nexus 1110-Xs 220
Configure Cisco Nexus 1110-X Virtual Appliances 222
Set Up the Primary Cisco Nexus 1000V VSM 224
Set Up the Secondary Cisco Nexus 1000V VSM 225
Install Cisco Virtual Switch Update Manager 226
Register the Cisco Nexus 1000V in VMware vCenter 230
Perform Base Configuration of the Primary VSM 231
Add VMware ESXi Hosts to Cisco Nexus 1000V 234
Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V 237
Cisco Nexus 1000V vTracker 239
FlexPod Management Tools Setup 241
Cisco UCS Performance Manager 241
Cisco UCS Performance Manager OVA Deployment 241
Cisco UCS Performance Manager Initial Configuration 246
Cisco UCS Performance Manager Deployment 256
Cisco UCS Performance Manager Configuration of FlexPod Infrastructure 262
NetApp Virtual Storage Console 6.2P1 Deployment Procedure 267
Virtual Storage Console 6.2P1 Pre-Installation Considerations 267
Install Virtual Storage Console 6.2P1 268
Register VSC with vCenter Server 271
Discover and Add Storage Resources 272
Optimal Storage Settings for ESXi Hosts 272
VSC 6.2P1 Backup and Recovery 273
OnCommand Unified Manager 6.3P2 277
OnCommand Unified Manager OVF Deployment 277
OnCommand Unified Manager Basic Setup 283
OnCommand Performance Manager 2.0 290
OnCommand Performance Manager OVF Deployment 290
OnCommand Performance Manager Basic Setup 295
Link OnCommand Performance Manager to OnCommand Unified Manager 297
NetApp NFS Plug-In 1.1.0 for VMware VAAI 299
Enable VMware vStorage for NFS in Clustered Data ONTAP 299
Install NetApp NFS Plug-In for VMware VAAI 300

About the Authors....303
Acknowledgements....303

Executive Summary

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

This document describes the Cisco and NetApp® FlexPod Datacenter with Cisco UCS Manager unified software release 3.1 and VMware vSphere 6.0 U1. Cisco UCS Manager (UCSM) 3.1 provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with Cisco UCS unified software release and VMware vSphere 6.0 U1 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, and NetApp AFF.

Solution Overview

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod, which uses best of breed storage, server and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

Audience

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

Purpose of this Document

This document provides a step by step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect, NetApp AFF, and Cisco Nexus 9000 solution. For the design decisions and technology discussion of the solution, please refer to FlexPod Datacenter with Cisco Unified Software Release and VMware vSphere 6 Design Guide:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi60_n9k_design.html

What's New?

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

- Validation of Cisco UCS 6300 Fabric Interconnects
- Support for the Cisco UCS 3.1(1h) unified software release, Cisco UCS B200-M4 servers, and Cisco UCS C220-M4 servers
- Support for the latest release of NetApp Data ONTAP® 8.3.2
- An IP-based storage design, supplemented with direct attached fibre channel connectivity, supporting both NAS datastores, and FC and iSCSI based SAN LUNs
- Validation of VMware vSphere 6.0 U1b
- HTML-based Cisco UCS Manager

Solution Design

Architecture

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere® built on FlexPod includes NetApp All Flash FAS storage, Cisco Nexus® networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

Figure 1 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with IP-based storage. This design uses the Cisco UCS Fabric Interconnect, Cisco Nexus 9000, and Cisco UCS C-Series and B-Series servers and the NetApp AFF family of storage controllers connected in a highly available modular design. This infrastructure is deployed to provide either FC or iSCSI-booted hosts with file-level and block-level access to shared storage. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Physical Topology

Figure 1 illustrates the physical architectures.

Figure 1 FlexPod Design with Cisco Nexus 9000 and NetApp Data ONTAP

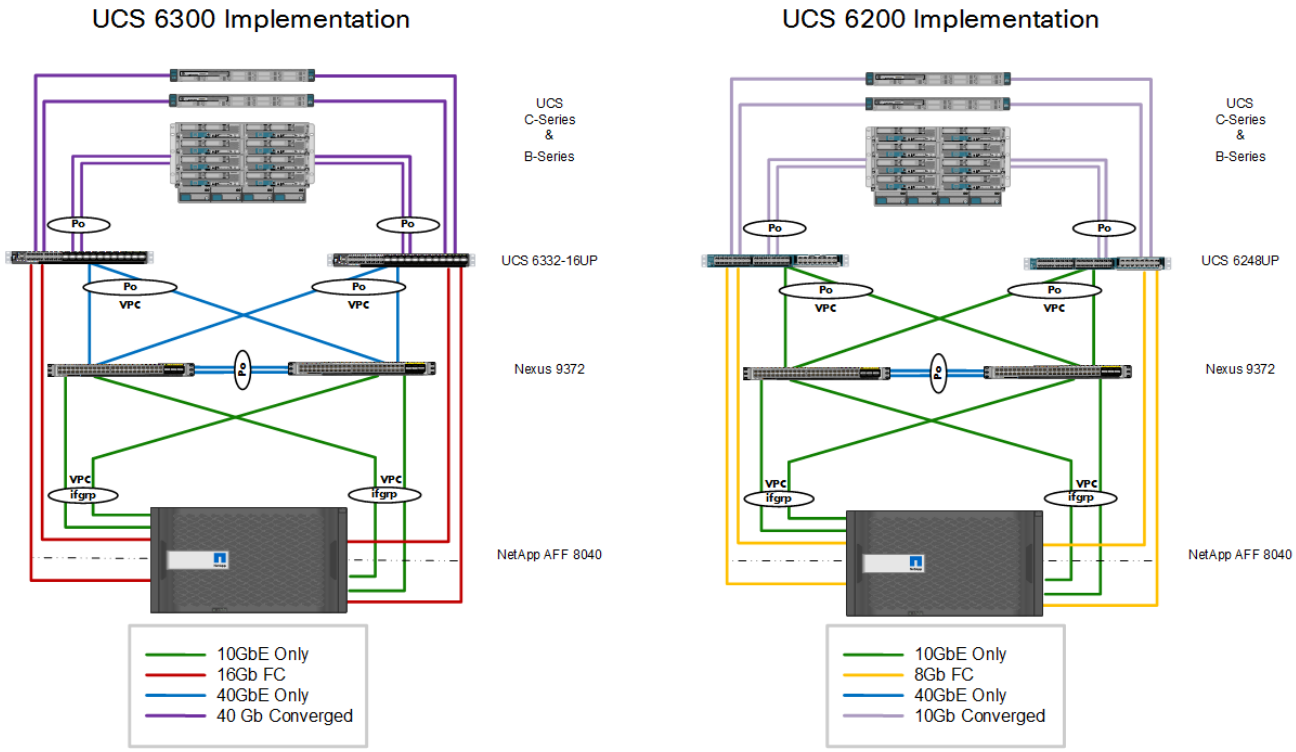

The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches
- Two Cisco UCS 6332-16UP or Two Cisco UCS 6248UP fabric interconnects
- One NetApp AFF8040 (HA pair) running clustered Data ONTAP with Disk shelves and Solid State Drives (SSD)

For server virtualization, the deployment includes VMware vSphere 6.0 U1. Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps for deploying the base architecture, as shown in Figure 1. These procedures cover everything from physical cabling to network, compute and storage device configurations.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6300 Series, UCS B-200 M4, UCS C-220 M4	3.1(1h)	Includes the Cisco UCS-IOM 2304, Cisco UCS Manager, Cisco UCS VIC 1340 and Cisco UCS VIC 1385
	Cisco eNIC	2.3.0.7	
	Cisco fNIC	1.6.0.25	
Network	Cisco Nexus 9000 NX-OS	7.0(3)I1(3)	
	Cisco Nexus 1000V	5.2(1)SV3(1.5b)	
	Cisco Nexus 1110-X	5.2(1)SP1(7.3)	
Storage	NetApp AFF 8040	Data ONTAP 8.3.2	
Software	Cisco UCS Manager	3.1(1h)	
	Cisco UCS Performance Manager	2.0	
	VMware vSphere ESXi	6.0 U1b	
	VMware vCenter	6.0 U1b	
	NetApp Virtual Storage Console (VSC)	6.2	
	OnCommand Performance Manager	2.0	

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-01, VM-Host-Prod-02 to represent infrastructure and production hosts deployed to each of the fabric interconnects in this document. Finally, to

indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?

 [-node] <nodename>                  Node

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

 |  -port {<netport>|<ifgrp>}        Associated Network Port

 [-vlan-id] <integer> }              Network Switch VLAN Identifier
```

Example:

```
network port vlan -node <node01> -vlan-name i0a-<vlan id>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 3 lists the virtual machines (VMs) necessary for deployment as outlined in this guide. Table 2 describes the VLANs necessary for deployment as outlined in this guide.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out of Band Mgmt	VLAN for out-of-band management interfaces	13
In-Band Mgmt	VLAN for in-band management interfaces	113
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for Infrastructure NFS traffic	3170
vMotion	VLAN for VMware vMotion	3173
VM-Traffic	VLAN for Production VM Interfaces	3174
iSCSI-A	VLAN for Fabric A iSCSI	901
iSCSI-B	VLAN for Fabric B iSCSI	902
Packet-Ctrl	VLAN Nexus 1110-X Packet and Control	3176

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

Table 3 Virtual Machines

Virtual Machine Description	Host Name
Active Directory	
vCenter Server	
NetApp VSC	
NetApp OnCommand Unified Manager	

Virtual Machine Description	Host Name
OnCommand Performance Manager	

Table 4 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps.

Table 4 Configuration Variables

Variable	Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01 (Example: 192.168.156.21)
<<var_node01_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_url_boot_software>>	Data ONTAP 8.3.2 URL (Example: http://192.168.156.9/832_q_image.tgz)
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02 (Example: 192.168.156.22)
<<var_node02_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_clustername>>	Storage cluster host name (Example: clus)
<<var_cluster_base_license_key>>	Cluster base license key (Example: 1234567890ABCD1234567890ABCD)
<<var_nfs_license>>	NFS license key (Example: 1234567890ABCD1234567890ABCD)
<<var_fc_license>>	Fiber Channel license key (if using Fiber Channel) (Example: 1234567890ABCD1234567890ABCD)
<<var_iscsi_license>>	iSCSI license key (if using iSCSI) (Example: 1234567890ABCD1234567890ABCD)
<<var_password>>	Global default administrative password (Example: Fl3xP0d9)
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster (Example: 192.168.157.20)
<<var_clustermgmt_mask>>	In-band management network netmask (Example: 255.255.255.0)
<<var_clustermgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.157.1)
<<var_dns_domain_name>>	DNS domain name (Example: flexpod.com)

Variable	Value
<<var_nameserver_ip>>	DNS server IP(s) (Example: 192.168.156.9)
<<var_node_location>>	Node location string for each node (Example: RTP9-D04)
<<var_node01_sp_ip>>	Out-of-band cluster node 01 service processor management IP (Example: 192.168.156.18)
<<var_node01_sp_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node01_sp_gateway>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP (Example: 192.168.156.19)
<<var_node02_sp_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node02_sp_gateway>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_node01>>	Cluster node 01 hostname (Example: clus-01)
<<var_node02>>	Cluster node 02 hostname (Example: clus-02)
<<var_num_disks>>	Number of disks to assign to each storage controller (Example: 5)
<<var_nfs_vlan_id>>	Infrastructure NFS VLAN ID for LIF (Example: 3170)
<<var_iscsi_vlan_A_id>>	Infrastructure iSCSI-A VLAN ID for LIF (Example: 901)
<<var_iscsi_vlan_B_id>>	Infrastructure iSCSI-B VLAN ID for LIF (Example: 902)
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID (Example: 113)
<<var_oob_mgmt_vlan_id>>	Out-of-band management network VLAN ID (Example: 13)
<<var_timezone>>	FlexPod time zone (Example: America/New_York)
<<var_global_ntp_server_ip>>	NTP server IP address for out-of-band mgmt. (Example: 192.168.156.1)
<<var_switch_a_ntp_ip>>	NTP server IP address for Nexus 9372 Switch A (Example: 192.168.156.1)
<<var_switch_b_ntp_ip>>	NTP server IP address for Nexus 9372 Switch B (Example: 192.168.156.1)
<<var_ib-mgmt_vlan_netmask_length>>	Length of IB-MGMT-VLAN Netmask (Example: /24)

Variable	Value
<<var_snmp_contact>>	Administrator e-mail address (Example: admin@flexpod.com)
<<var_snmp_location>>	Cluster location string (Example: RTP9-D04)
<<var_cert_common_name>>	Common name string for certificate (Example: “clus.flexpod.com”)
<<var_cert_country>>	Country for certificate (Example: “USA”)
<<var_cert_state>>	State for certificate (Example: “NC”)
<<var_cert_locality>>	Locality for certificate (Example: “RTP”)
<<var_cert_org>>	Organization for certificate (Example: “FlexPod”)
<<var_cert_unit>>	Organizational Unit for certificate (Example: “Dev”)
<<var_cert_email>>	E-mail address for certificate (Example: “admin@flexpod.com”)
<<var_cert_days>>	Days until certificate expiration (Example: 365)
<<var_oncommand_server_fqdn>>	VSC or OnCommand VM fully qualified domain name (FQDN) (Example: ocum.flexpod.com)
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name (Example: fl3xp0d)
<<var_mailhost>>	Mail server host name (Example: smtp.flexpod.com)
<<var_storage_admin_email>>	Administrator e-mail address (Example: storage@flexpod.com)
<<var_node01_nfs_lif_infra_swap_ip>>	IP address of Infra Swap (Example: 192.168.170.21)
<<var_node01_nfs_lif_infra_swap_mask>>	Subnet Mask of Infra Swap (Example: 255.255.255.0)
<<var_node02_nfs_lif_infra_datastore_1_ip>>	IP address of Datastore 1 (Example: 192.168.170.22)
<<var_node02_nfs_lif_infra_datastore_1_mask>>	Subnet mask of Datastore 1 (Example: 255.255.255.0)
<<var_vserver_mgmt_ip>>	Management IP address for Vserver (Example: 192.168.156.23)
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver (Example: 255.255.255.0)
<<var_vserver_mgmt_gateway>>	Default Gateway for Vserver (Example: 192.168.156.1)

Variable	Value
<<var_vsadmin_password>>	Password for VS admin account (Example: Fl3xP0d)
<<var_ucs_6248_clustername>>	Cisco UCS Manager cluster host name (Example: ucs-6248)
<<var_ucs_6332_clustername>>	Cisco UCS Manager cluster host name (Example: ucs-6332)
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (Fl) A out-of-band management IP address (Example: 192.168.156.51)
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_ucsb_mgmt_ip>>	Cisco UCS Fl B out-of-band management IP address (Example: 192.168.156.52)
<<var_vm_host_infra_01_iqn>>	Cisco UCS Service Profile generated IQN of Infra 02 (Example: iqn.1992-08.com.cisco:ucs-6248-host:1
<<var_vm_host_prod_02_iqn>>	Cisco UCS Service Profile generated IQN of Infra 01 (Example: iqn.1992-08.com.cisco:ucs-6332-host:1
<<var_vm_host_infra_01_ip>>	VMware ESXi host 02 out-of-band management IP (Example: 10.1.156.25)
<<var_vm_host_prod_02_ip>>	VMware ESXi host 01 out-of-band management IP (Example: 10.1.156.28)
<<var_nfs_vlan_ip_host_01>>	ESXi host 1, NFS VLAN IP (Example: 192.168.170.25)
<<var_nfs_vlan_ip_mask_host_01>>	ESXi host1, NFS VLAN subnet mask (Example: 255.255.255.0)
<<var_nfs_vlan_ip_host_02>>	ESXi host 2, NFS VLAN IP (Example: 192.168.170.28)
<<var_nfs_vlan_ip_mask_host_02>>	ESXi host2, NFS VLAN subnet mask (Example: 255.255.255.0)
<<var_vcenter_server_ip>>	IP address of the vCenter Server (Example: 10.1.156.100)

Variable	Value
<<var_svm_mgmt_vlan_id>>	Infrastructure Vserver management VLAN ID (Example: 13)
<<var_node01_iscsi_lif01a_ip>>	iSCSI LIF 01a IP address (Example: 192.168.91.21)
<<var_node01_iscsi_lif01a_mask>>	iSCSI LIF 01a subnet mask (Example: 255.255.255.0)
<<var_node01_iscsi_lif01b_ip>>	iSCSI LIF 01b IP address (Example: 192.168.92.21)
<<var_node01_iscsi_lif01b_mask>>	iSCSI LIF 01b subnet mask (Example: 255.255.255.0)
<<var_node01_iscsi_lif02a_ip>>	iSCSI LIF 02a IP address (Example: 192.168.91.22)
<<var_node01_iscsi_lif02a_mask>>	iSCSI LIF 02a subnet mask (Example: 255.255.255.0)
<<var_node01_iscsi_lif02b_ip>>	iSCSI LIF 02b IP address (Example: 192.168.92.22)
<<var_node01_iscsi_lif02b_mask>>	iSCSI LIF 02b subnet mask (Example: 255.255.255.0)
<<var_node01_fcp_p0rt1 >>	Node 1 FC port 1 (Example: 0a)
<<var_node01_fcp_p0rt2 >>	Node 1 FC port 2 (Example: 0b)
<<var_node02_fcp_p0rt1 >>	Node 2 FC port 1 (Example: 0a)
<<var_node02_fcp_p0rt2 >>	Node 2 FC port 2 (Example: 0b)
<<var_vserver_mgmt_ip>>	Management IP address for Infrastructure Vserver (Example: 192.168.156.23)
<<var_vserver_mgmt_mask>>	Management subnet mask for Infrastructure Vserver (Example: 255.255.255.0)
<<var_oncommand_server_ip>>	IP address of the OnCommand Unified Manager (Example: 10.1.156.10)
<<var_rule_index>>	Rule index number (Example: 1)

Physical Infrastructure

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain the details for the prescribed and supported configuration of the NetApp AFF8040 running clustered Data ONTAP 8.3.2.

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

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

Be sure to follow the cabling directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.

Figure 2 shows a cabling diagram for a FlexPod configuration using the Cisco Nexus 9000 and NetApp storage systems with clustered Data ONTAP connected to the UCS 6332-16UP Fabric Interconnect, and Figure 3 shows that same configuration utilizing the UCS 6248UP Fabric Interconnect. Cabling of both the 6332-16UP and the 6248UP in adjacency are done as an example of interchangeable viability between the two fabric interconnect models, and not intended to imply a requirement of simultaneous deployment.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 2 FlexPod Cabling Diagram (6332-16UP)

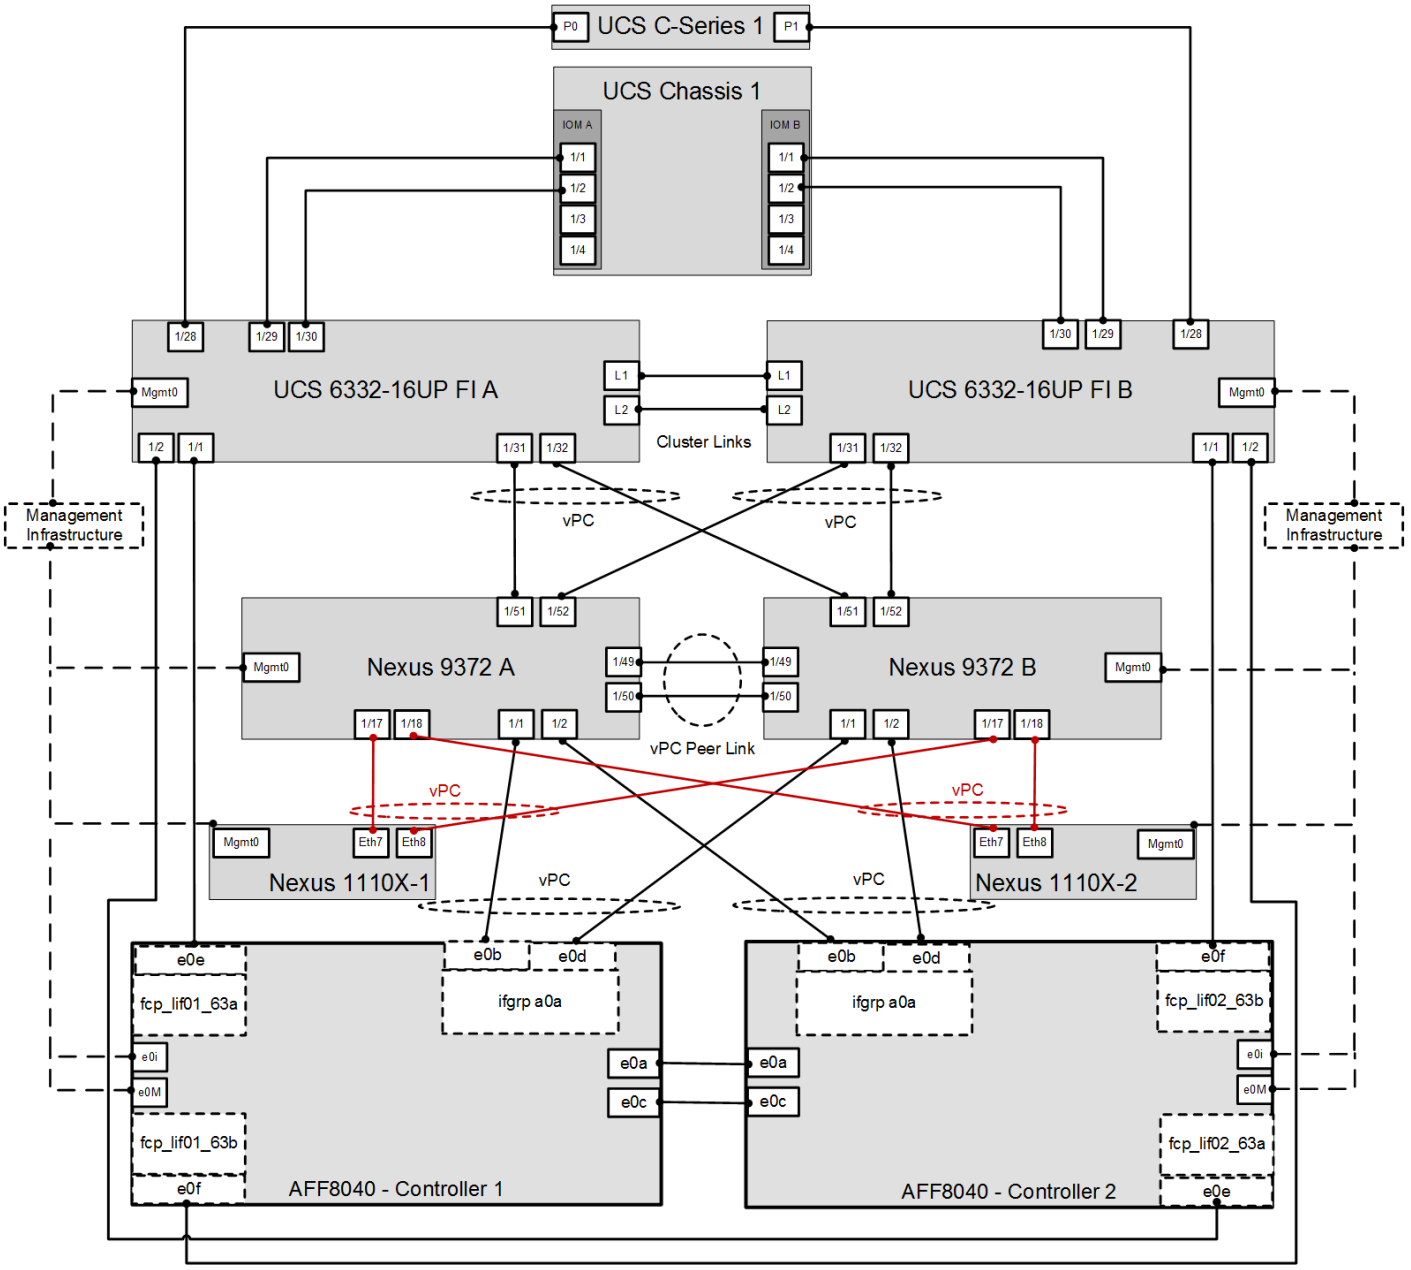
UCS C-Series 1
P0
P1
UCS Chassis 1
IOM A
IOM B
1/1
1/2
1/3
1/4
1/28
1/29
1/30
UCS 6332-16UP FI A
UCS 6332-16UP FI B
Mgmt0
L1
L2
Cluster Links
1/31
1/32
vPC
Management Infrastructure
1/51
1/52
Nexus 9372 A
Nexus 9372 B
1/49
1/50
vPC Peer Link
1/17
1/18
Eth7
Eth8
Nexus 1110X-1
Nexus 1110X-2
e0e
e0b
e0d
e0f
fcp_lif01_63a
ifgrp a0a
fcp_lif02_63b
e0i
e0M
e0a
e0c
fcp_lif01_63b
fcp_lif02_63a
AFF8040 - Controller 1
AFF8040 - Controller 2

Figure 3 FlexPod Cabling Diagram (6248UP)

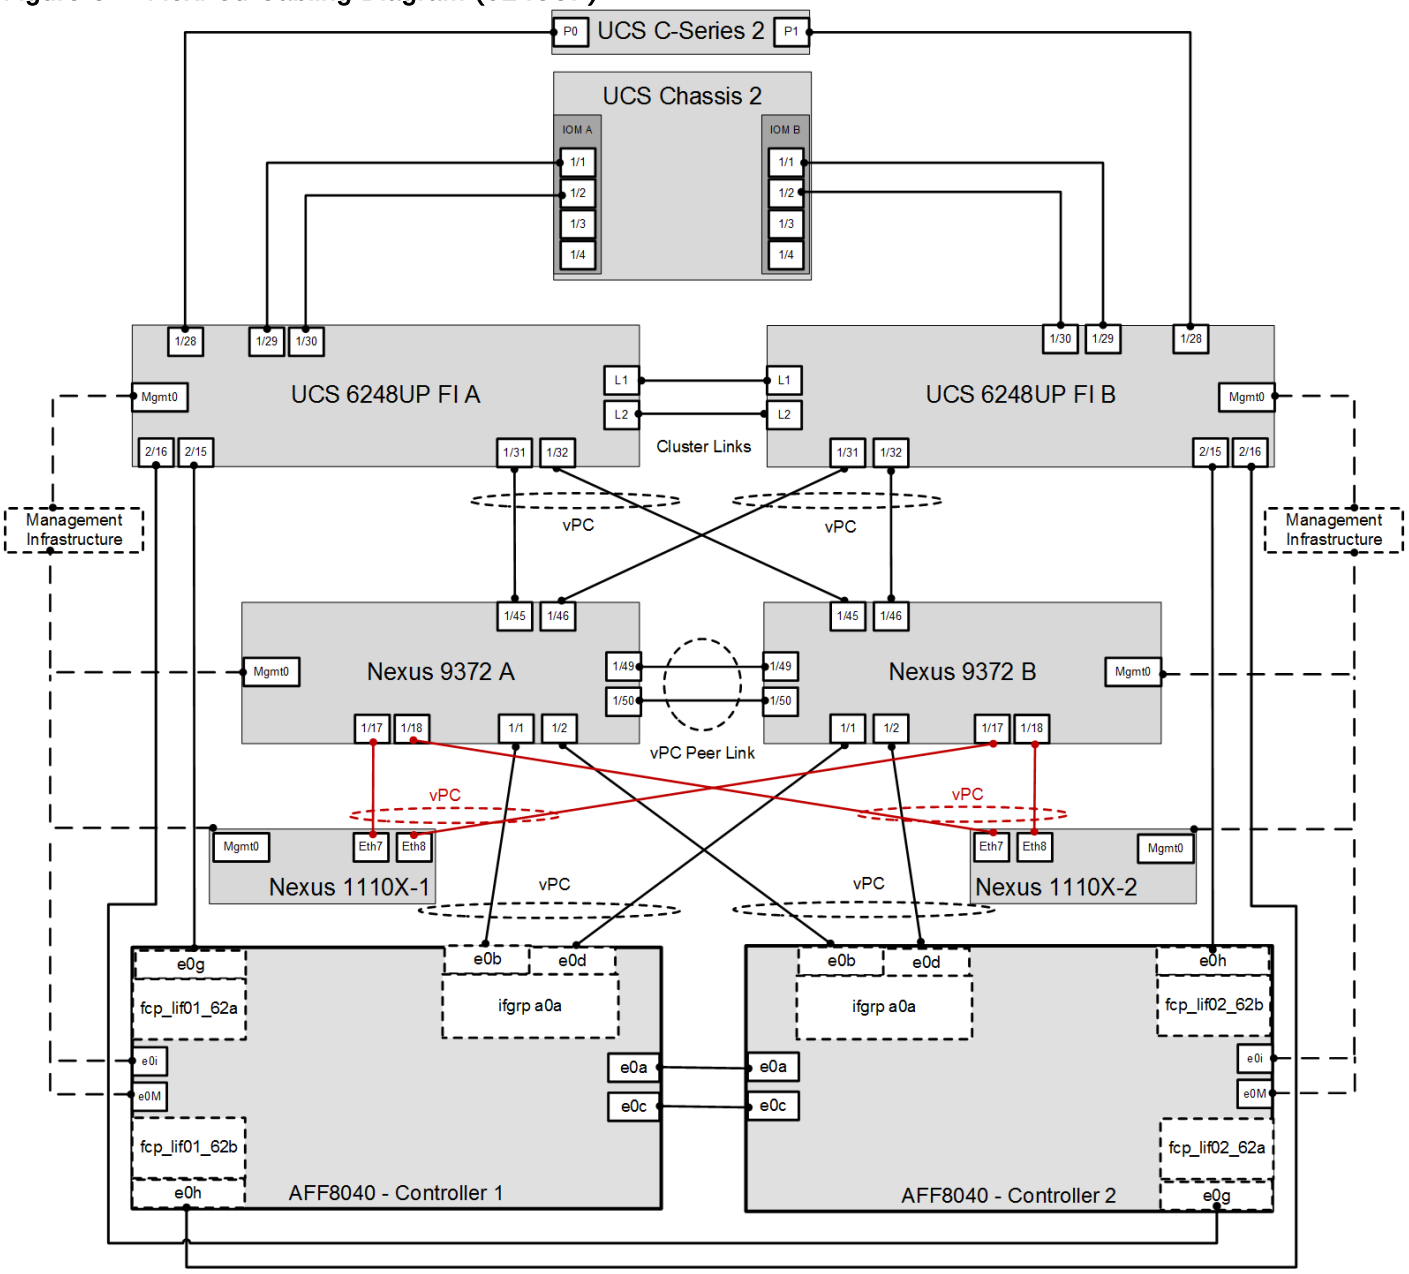

Table 5 through Error! Reference source not found. provide the details of all the connections in use.

Table 5 Cisco Nexus 9372-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 A	Eth1/1	10GbE	NetApp Controller 1	e0b
	Eth1/2	10GbE	NetApp Controller 2	e0b
	Eth1/17	10GbE	Nexus 1110-X 1	Eth7

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/18	10GbE	Nexus 1110-X 2	Eth7
	Eth1/45	10GbE	Cisco UCS 6248UP FI A	Eth1/31
	Eth1/46	10GbE	Cisco UCS 6248UP FI B	Eth1/31
	Eth1/49	40GbE	Cisco Nexus 9372 B	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 B	Eth1/50
	Eth1/51	40GbE	Cisco UCS 6332-16UP FI A	Eth1/31
	Eth1/52	40GbE	Cisco UCS 6332-16UP FI B	Eth1/31
	MGMT0	GbE	GbE management switch	Any

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

Table 6 Cisco Nexus 9372-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 B	Eth1/1	10GbE	NetApp Controller 1	e0d
	Eth1/2	10GbE	NetApp Controller 2	e0d
	Eth1/17	10GbE	Nexus 1110-X 1	Eth8
	Eth1/18	10GbE	Nexus 1110-X 2	Eth8
	Eth1/45	10GbE	Cisco UCS 6248UP FI A	Eth1/32
	Eth1/46	10GbE	Cisco UCS 6248UP FI B	Eth1/32
	Eth1/49	40GbE	Cisco Nexus 9372 A	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 A	Eth1/50
	Eth1/51	40GbE	Cisco UCS 6332-16UP FI A	Eth1/32
	Eth1/52	40GbE	Cisco UCS 6332-16UP FI B	Eth1/32
	MGMT0	GbE	GbE management switch	Any

Table 7 NetApp Controller-1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 2	e0a
	e0b	10GbE	Cisco Nexus 9372 A	Eth1/1
	e0c	10GbE	NetApp Controller 2	e0c
	e0d	10GbE	Cisco Nexus 9372 B	Eth1/1
	0e	16Gb FC	Cisco UCS 6332-16UP FI A (only if using FC connectivity)	FC 1/1
	0f	16Gb FC	Cisco UCS 6332-16UP FI B (only if using FC connectivity)	FC 1/1
	0g	16Gb FC	Cisco UCS 6248UP FI A (only if using FC connectivity)	FC 2/15
	0h	16Gb FC	Cisco UCS 6248UP FI B (only if using FC connectivity)	FC 2/15

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

Table 8 NetApp Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 1	e0a
	e0b	10GbE	Cisco Nexus 9372 A	Eth1/2
	e0c	10GbE	NetApp Controller 1	e0c

Local Device	Local Port	Connection	Remote Device	Remote Port
	e0d	10GbE	Cisco Nexus 9372 B	Eth1/2
	0e	16Gb FC	Cisco UCS 6332-16UP FI A (only if using FC connectivity)	FC 1/2
	0f	16Gb FC	Cisco UCS 6332-16UP FI B (only if using FC connectivity)	FC 1/2
	0g	16Gb FC	Cisco UCS 6248UP FI A (only if using FC connectivity)	FC 2/16
	0h	16Gb FC	Cisco UCS 6248UP FI B (only if using FC connectivity)	FC 2/16

Table 9 Cisco UCS 6332-16UP Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	FC 1/1	16Gb FC	NetApp controller 1 (only if using FC)	0e
	FC 1/2	16Gb FC	NetApp controller 2 (only if using FC)	0e
	Eth1/28	40GbE	Cisco UCS C-Series 1	Port 0
	Eth1/29	40GbE	Cisco UCS Chassis 1 2304 FEX A	IOM 1/1
	Eth1/30	40GbE	Cisco UCS Chassis 1 2304 FEX A	IOM 1/2
	Eth1/31	40GbE	Cisco Nexus 9372 A	Eth1/51
	Eth1/32	40GbE	Cisco Nexus 9372 B	Eth1/51
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6332-16UP FI B	L1
	L2	GbE	Cisco UCS 6332-16UP FI B	L2

Table 10 Cisco UCS 6332-16UP Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	FC 1/1	16Gb FC	NetApp controller 1 (only if using FC)	0f
	FC 1/2	16Gb FC	NetApp controller 2 (only if using FC)	0f

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/28	40GbE	Cisco UCS C-Series 1	Port 1
	Eth1/29	40GbE	Cisco UCS Chassis 1 2304 FEX B	IOM 1/1
	Eth1/30	40GbE	Cisco UCS Chassis 1 2304 FEX B	IOM 1/2
	Eth1/31	40GbE	Cisco Nexus 9372 A	Eth1/52
	Eth1/32	40GbE	Cisco Nexus 9372 B	Eth1/52
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6332-16UP FI B	L1
	L2	GbE	Cisco UCS 6332-16UP FI B	L2

Table 11 Cisco UCS 6248UP Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	FC 2/15	8Gb FC	NetApp controller 1 (only if using FC)	0g
	FC 2/16	8Gb FC	NetApp controller 2 (only if using FC)	0g
	Eth1/28	10GbE	Cisco UCS C-Series 2	Port 0
	Eth1/29	10GbE	Cisco UCS Chassis 2 2204 FEX A	IOM 1/1
	Eth1/30	10GbE	Cisco UCS Chassis 2 2204 FEX A	IOM 1/2
	Eth1/31	10GbE	Cisco Nexus 9372 A	Eth1/45
	Eth1/32	10GbE	Cisco Nexus 9372 B	Eth1/45
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6248UP FI B	L1
	L2	GbE	Cisco UCS 6248UP FI B	L2

Table 12 Cisco UCS 6248UP Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	FC 2/15	8Gb FC	NetApp controller 1 (only if using FC)	0f

Local Device	Local Port	Connection	Remote Device	Remote Port
	FC 2/16	8Gb FC	NetApp controller 2 (only if using FC)	0f
	Eth1/28	10GbE	Cisco UCS C-Series 2	Port 1
	Eth1/29	10GbE	Cisco UCS Chassis 2204 FEX B	IOM 1/1
	Eth1/30	10GbE	Cisco UCS Chassis 2204 FEX B	IOM 1/2
	Eth1/31	10GbE	Cisco Nexus 9372 A	Eth1/46
	Eth1/32	10GbE	Cisco Nexus 9372 B	Eth1/46
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6248UP FI B	L1
	L2	GbE	Cisco UCS 6248UP FI B	L2

Table 13 Cisco UCS C-Series 1

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 1	Port 0	40GbE	Cisco UCS 6332-16UP FI A	Eth1/28
	Port 1	40GbE	Cisco UCS 6332-16UP FI B	Eth1/28

Table 14 Cisco UCS C-Series 2

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 2	Port 0	10GbE	Cisco UCS 6248UP FI A	Eth1/28
	Port 1	10GbE	Cisco UCS 6248UP FI B	Eth1/28

Network Switch Configuration

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

Physical Connectivity

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

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Nexus 9000 7.0(3)I1(3).

The following procedure includes setup of NTP distribution on the In-Band Management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes the default VRF will be used to route the In-Band Management VLAN.

Set Up Initial Configuration

Cisco Nexus 9372PX A

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

1. Configure the switch.

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

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

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

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

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

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

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

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

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

Enter the switch name: <<var_nexus_A_hostname>>

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

Mgmt0 IPv4 address: <<var_nexus_A_mgmt0_ip>>
```

```
Mgmt0 IPv4 netmask: <<var_nexus_A_mgmt0_netmask>>

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

IPv4 address of the default gateway: <<var_nexus_A_mgmt0_gw>>

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

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

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


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

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

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

NTP server IPv4 address: <<var_global_ntp_server_ip>>

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

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

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

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

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus 9372PX B

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

1. Configure the switch.

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

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

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

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

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

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

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

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

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

Enter the switch name: <<var_nexus_B_hostname>>
```

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

Mgmt0 IPv4 address: <<var_nexus_B_mgmt0_ip>>

Mgmt0 IPv4 netmask: <<var_nexus_B_mgmt0_netmask>>

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

IPv4 address of the default gateway: <<var_nexus_B_mgmt0_gw>>

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

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

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


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

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

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

NTP server IPv4 address: <<var_global_ntp_server_ip>>

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

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

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

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

2. Review the configuration summary before enabling the configuration.

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

FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

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

```
config t

feature interface-vlan

feature lacp

feature vpc

feature lldp
```

Set Global Configurations

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

Run the following commands to set global configurations:

```
spanning-tree port type network default

spanning-tree port type edge bpduguard default

spanning-tree port type edge bpdufilter default

port-channel load-balance src-dst l4port

ntp server <<var_global_ntp_server_ip>> use-vrf management

ntp master 3

ip route 0.0.0.0/0 <<var_ib-mgmt-vlan_gateway>>

copy run start
```

Create VLANs

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
vlan <<var_ib-mgmt_vlan_id>>

name IB-MGMT-VLAN

exit

vlan <<var_native_vlan_id>>

name Native-VLAN

exit

vlan <<var_vmotion_vlan_id>>

name vMotion-VLAN

exit

vlan <<var_vm-traffic_vlan_id>>

name VM-Traffic-VLAN

exit

vlan <<var_nfs_vlan_id>>

name NFS-VLAN

exit

vlan <<var_iscsi-a_vlan_id>>
```

```
name iSCSI-A-VLAN

exit

vlan <<var_iscsi-b_vlan_id>>

name iSCSI-B-VLAN

exit

vlan <<var_packet-ctrl_vlan_id>>

name Packet-Ctrl-VLAN

exit
```

Add NTP Distribution Interface

Cisco Nexus 9372PX A

From the global configuration mode, run the following commands:

```
ntp source <<var_switch_a_ntp_ip>>

interface Vlan<<var_ib-mgmt_vlan_id>>

ip address <<var_switch_a_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>

no shutdown

exit
```

Cisco Nexus 9372PX B

From the global configuration mode, run the following commands:

```
ntp source <<var_switch_b_ntp_ip>>

interface Vlan<<var_ib-mgmt_vlan_id>>

ip address <<var_switch_b_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>

no shutdown

exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 9372PX A

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

In this step and in further sections, configure the <<var_ucs_6248_clustername>> and <<var_ucs_6332_clustername>> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface Eth1/1
```

```
description <<var_node01>>:e0b

exit

interface Eth1/2

description <<var_node02>>:e0b

exit

interface Eth1/17

description <<var_n1110-x>>-1:eth 7

exit

interface Eth1/18

description <<var_n1110-x>>-2:eth7

exit

interface Eth1/45

description <<var_ucs_6248_clustername>>-a:1/31

exit

interface Eth1/46

description <<var_ucs_6248_clustername>>-b:1/31

exit

interface Eth1/49

description <<var_nexus_B_hostname>>:1/49

exit

interface Eth1/50

description <<var_nexus_B_hostname>>:1/50

exit

interface Eth1/51

description <<var_ucs_6332_clustername>>-a:1/31

exit

interface Eth1/52

description <<var_ucs_6332_clustername>>-b:1/31

exit
```

Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_node01>>:e0d
exit
interface Eth1/2
description <<var_node02>>:e0d
exit
interface Eth1/17
description <<var_n1110-x>>-1:eth 8
exit
interface Eth1/18
description <<var_n1110-x>>-2:eth 8
exit
interface Eth1/45
description <<var_ucs_6248_clustername>>-a:1/32
exit
interface Eth1/46
description <<var_ucs_6248_clustername>>-b:1/32
exit
interface Eth1/49
description <<var_nexus_A_hostname>>:1/49
exit
interface Eth1/50
description <<var_nexus_A_hostname>>:1/50
exit
interface Eth1/51
description <<var_ucs_6332_clustername>>-a:1/32
exit
interface Eth1/52
description <<var_ucs_6332_clustername>>-b:1/32
exit
```

Create Port Channels

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
exit
interface Eth1/49-50
channel-group 10 mode active
no shutdown
exit
interface Po11
description <<var_node01>>
exit
interface Eth1/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
exit
interface Eth1/2
channel-group 12 mode active
no shutdown
exit
interface Po117
description <<var_n1110-x>>-1
exit
interface Eth1/17
channel-group 117 mode active
no shutdown
exit
interface Po118
```

```
description <<var_n1110-x>>-2
exit
interface Eth1/18
channel-group 118 mode active
no shutdown
exit
copy run start
interface Po145
description <<var_ucs_6248_clustername>>-a
exit
interface Eth1/45
channel-group 145 mode active
no shutdown
exit
interface Po112
description <<var_ucs_6248_clustername>>-b
exit
interface Eth1/46
channel-group 146 mode active
no shutdown
exit
interface Po151
description <<var_ucs_6332_clustername>>-a
exit
interface Eth1/51
channel-group 151 mode active
no shutdown
exit
interface Po152
description <<var_ucs_6332_clustername>>-b
exit
interface Eth1/52
channel-group 152 mode active
```

```
no shutdown

exit
```

Configure Port Channel Parameters

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
interface Po10

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>, <<var_packet-
ctrl_vlan_id>>

spanning-tree port type network

exit

interface Po11

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>,
<<var_iscsi-b_vlan_id>>

spanning-tree port type edge trunk

mtu 9216

exit

interface Po12

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>,
<<var_iscsi-b_vlan_id>>

spanning-tree port type edge trunk

mtu 9216

exit

interface Po117

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_packet-ctrl_vlan_id>>
```

```
spanning-tree port type edge trunk

exit

interface Po118

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_packet-ctrl_vlan_id>>

spanning-tree port type edge trunk

exit

interface Po145

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>

spanning-tree port type edge trunk

mtu 9216

exit

interface Po146

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>

spanning-tree port type edge trunk

mtu 9216

exit

interface Po151

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>

spanning-tree port type edge trunk

mtu 9216

exit

interface Po152

switchport mode trunk
```

```
switchport trunk native vlan 2

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>

spanning-tree port type edge trunk

mtu 9216

exit




copy run start
```

Configure Virtual Port Channels

Cisco Nexus 9372PX A

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

From the global configuration mode, run the following commands:

```
vpc domain <<var_nexus_vpc_domain_id>>

role priority 10

peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source <<var_nexus_A_mgmt0_ip>>

peer-switch

peer-gateway

auto-recovery

delay restore 150

exit

interface Po10

vpc peer-link

exit

interface Po11

vpc 11

exit

interface Po12

vpc 12

exit

interface Po117

vpc 117

exit
```

```
interface Po118

vpc 118

exit

interface Po145

vpc 111

exit

interface Po146

vpc 112

exit

interface Po151

vpc 111

exit

interface Po152

vpc 112

exit

copy run start
```

Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands.

```
vpc domain <<var_nexus_vpc_domain_id>>

role priority 20

peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source <<var_nexus_B_mgmt0_ip>>

peer-switch

peer-gateway

auto-recovery

delay restore 150

exit

interface Po10

vpc peer-link

exit

interface Po11

vpc 11

exit
```

```
interface Po12

vpc 12

exit

interface Po117

vpc 117

exit

interface Po118

vpc 118

exit

interface Po145

vpc 111

exit

interface Po146

vpc 112

exit

interface Po151

vpc 111

exit

interface Po152

vpc 112

exit

copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 9372PX switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Storage Configuration

AFF80XX Series Controllers

See the following sections in the Site Requirements Guide for planning the physical location of the storage systems:

- Site Preparation
- System Connectivity Requirements
- Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements
- 80xx Series Systems

NetApp Hardware Universe

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

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the HWU application at the NetApp Support site.

1. Access the HWU application to view the System Configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers found in the AFF8000 Series product documentation at the NetApp Support site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported by the AFF 80xx is available at the NetApp Support site.

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

Clustered Data ONTAP 8.3.2

Complete the Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the Clustered Data ONTAP 8.3 Software Setup Guide. You must have access to the NetApp Support site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the Clustered Data ONTAP 8.3 Software Setup Guide to learn about configuring ONTAP. Table 15 lists the information that you will need to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 15 ONTAP software installation prerequisites

Cluster Detail	Cluster Detail Value
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.3.2 URL	<<var_url_boot_software>>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If Data ONTAP 8.3.2 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.2 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14.

4. To install new software, select option 7.

```
7
```

5. Enter y to perform an upgrade.

```
y
```

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

```
e0M
```

7. Enter y to reboot now.

```
y
```

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

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

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

This web server must be pingable.

```
<<var_url_boot_software>>
```

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

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

```
y
```

12. Enter y to reboot the node.

```
y
```

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

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

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

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

```
4
```

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

```
y
```

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

```
y
```

The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue with the node 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If Data ONTAP 8.3.2 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.2 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.

```
7
```

5. Enter `y` to perform an upgrade.

```
y
```

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

```
e0M
```

7. Enter `y` to reboot now.

```
y
```

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

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

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

This web server must be pingable.

```
<<var_url_boot_software>>
```

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

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

```
y
```

12. Enter `y` to reboot the node.

```
y
```

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

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

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

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

```
4
```

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

```
y
```

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

```
y
```

The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when Data ONTAP 8.3.2 boots on the node for the first time.

1. Follow the prompts to set up node 01:

```
Welcome to node setup.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
```

```
  "exit" or "quit" - if you want to quit the setup wizard.
     Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.

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

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
Enter the node management interface netmask: <<var_node01_mgmt_mask>>
Enter the node management interface default gateway: <<var_node01_mgmt_gateway>>
A node management interface on port e0M with IP address <<var_node01_mgmt_ip>> has been created

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility from the
NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.
```

2. Press Enter and log in to the node with the admin user ID and no password.

3. At the node command prompt, enter the following commands to set HA mode for storage failover.

 If the node responds that the HA mode was already set, then proceed with step 4.

```
::> storage failover modify -mode ha
Mode set to HA.  Reboot node to activate HA.

::> system node reboot

Warning: Are you sure you want to reboot node "localhost"? {y|n}: y
```

4. After reboot, set up the node with the preassigned values.

```
Welcome to node setup.

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

To accept a default or omit a question, do not enter a value.



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

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility from the
NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.
```

```
Alternatively, you can use the "cluster setup" command to configure the cluster.
```

5. Log in to the node as the admin user with no password.

Repeat this procedure for storage cluster node 02.

Create Cluster on Node 01

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

Table 16 Cluster `create` in ONTAP prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster node01 IP address	<<var_node01_mgmt_ip>>
Cluster node01 netmask	<<var_node01_mgmt_mask>>
Cluster node01 gateway	<<var_node01_mgmt_gateway>>

Run the `cluster setup` command to start the Cluster Setup wizard.

```
cluster setup
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster? {create, join}:
```

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

To create a new cluster, complete the following steps:

1. Run the following command to create a new cluster:

```
create
```

2. Enter `no` for the single-node cluster option.

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

3. Enter `no` for a cluster network using network switches.

```
Will the cluster network be configured to use network switches? [yes]:no
```

4. The system defaults are displayed. Enter `yes` to use the system defaults. Use the following prompts to configure the cluster ports.

```
Existing cluster interface configuration found:

Port    MTU     IP              Netmask
e0a     9000    169.254.118.102 255.255.0.0
e0c     9000    169.254.191.92  255.255.0.0

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

System Defaults:
Private cluster network ports [e0a,e0c].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.

Do you want to use these defaults? {yes, no} [yes]: yes
```

If four ports are being used for the switchless cluster interconnect, enter e0a, e0b, e0c, and e0d for the private cluster network ports above.

5. The steps to create a cluster are displayed.

```
Enter the cluster administrators (username “admin”) password: <<var_password>>
Retype the password: <<var_password>>

It can take several minutes to create cluster interfaces...


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

Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter an additional license key []:<<var_iscsi_license>>
```

The cluster is created. This can take a few minutes.

For this validated architecture, NetApp recommends installing license keys for NetApp SnapRestore® data recovery software, NetApp FlexClone® data replication technology, and the NetApp SnapManager® suite. In addition, install all required storage protocol licenses and all licenses that came with the AFF bundle. After you finish entering the license keys, press Enter.

```
Enter the cluster management interface port [e0e]: e0i
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway: <<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

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

If you have more than one name server IP address, separate the IP addresses with a comma.

7. Set up the node.

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

The node management interface has been modified to use port e0M with IP address <<var_node01_mgmt_ip>>.
This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on
your system.
For further information on AutoSupport, please see: http://support.netapp.com/autosupport/
Press enter to continue: Enter
Cluster "<<var_clustername>>" has been created.
To complete cluster setup, you must join each additional node to the cluster by running "cluster setup" on
each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for
information about additional system configuration tasks.  You can find the Software Setup Guide on the NetApp
Support Site.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line
interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address
(<<var_clustermgmt_ip>>).
To access the command-line interface, connect to the cluster management IP address (for example, ssh
admin@<<var_clustermgmt_ip>>).



<<var_clustername>>::>
```

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

Join Node 02 to Cluster

The first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01, and the node joining the cluster in this example is node 02.

Table 17 Cluster `join` in ONTAP prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster node02 IP address	<<var_node02_mgmt_ip>>
Cluster node02 netmask	<<var_node02_mgmt_mask>>

Cluster Detail	Cluster Detail Value
Cluster node02 gateway	<<var_node02_mgmt_gateway>>

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

1. If prompted, enter `admin` in the login prompt.

```
admin
```

2. Run the `cluster setup` command to start the Cluster Setup wizard.

```
cluster setup

This node's storage failover partner is already a member of a cluster.
Storage failover partners must be members of the same cluster.
The cluster setup wizard will default to the cluster join dialog.

Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{join}:
```

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

3. Run the following command to join a cluster:

```
join
```

4. Data ONTAP detects the existing cluster and agrees to join the same cluster. Follow the prompts to join the cluster.

```
Existing cluster interface configuration found:

Port    MTU     IP              Netmask
e0a     9000    169.254.1.79    255.255.0.0
e0c     9000    169.254.100.157 255.255.0.0
Do you want to use this configuration? {yes, no} [yes]: no

System Defaults:
Private cluster network ports [e0a,e0c].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
```

If four ports are being used for the switchless cluster interconnect, enter e0a, e0b, e0c, and e0d for the private cluster network ports above.

```
Do you want to use these defaults? {yes, no} [yes]:Enter
It can take several minutes to create cluster interfaces...
```

5. The steps to join a cluster are displayed.

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

Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter
Joining cluster <<var_clustername>>
Starting cluster support services ..

This node has joined the cluster <<var_clustername>>.


Step 2 of 3: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.


SFO is enabled.


Step 3 of 3: Set Up the Node
You can type "back", "exit", or "help" at any question.

Notice: HA is configured in management.
```

The node should find the cluster name. Cluster joining can take a few minutes.

6. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node02_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node02_netmask>>]: Enter
Enter the node management interface default gateway [<<var_node02_gw>>]: Enter
The node management interface has been modified to use port e0M with IP address <<var_node02_mgmt_ip>>.
This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on
your system.
For further information on AutoSupport, please see: http://support.netapp.com/autosupport/
Press enter to continue: Enter

This node has been joined to cluster "<<var_clustername>>".
To complete cluster setup, you must join each additional node to the cluster by running "cluster setup" on
each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for
information about additional system configuration tasks.  You can find the Software Setup Guide on the NetApp
Support Site.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line
interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address
(<<var_clustermgmt_ip>>).
To access the command-line interface, connect to the cluster management IP address (for example, ssh
admin@<<var_clustermgmt_ip>>).
```

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

Log In to the Cluster

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

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

Zero All Spare Disks

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

```
disk zerospares
```

Disk autoassign should have assigned half of the connected disks to each node in the HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare disks can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Onboard Unified Target Adapter 2 Port Personality

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

1. Verify the Current Mode and Current Type properties of the ports by running the `ucadmin` show command.

```
ucadmin show
                      Current  Current    Pending  Pending    Admin
Node          Adapter Mode     Type       Mode     Type       Status
------------  ------- -------  ---------  -------  ---------  -----------
<<var_node01>>
              0e      fc       target     -        -          online
<<var_node01>>
              0f      fc       target     -        -          online
<<var_node01>>
              0g      cna      target     -        -          online
<<var_node01>>
              0h      cna      target     -        -          online
<<var_node02>>
              0e      fc       target     -        -          online
<<var_node02>>
              0f      fc       target     -        -          online
<<var_node02>>
              0g      cna      target     -        -          online
<<var_node02>>
              0h      cna      target     -        -          online
8 entries were displayed.
```

2. Verify that the Current Mode and Current Type properties for all ports are set properly. Set ports used for Fibre Channel (FC) connectivity to mode `fc`; otherwise, set them to the mode `cna`. That includes FCoE ports, which should be set to the mode `cna`. The port type for all protocols should be set to `target`. Change the port personality with the following command:

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

The ports must be offline to run this command. To take an adapter offline, run the `fcp adapter modify -node <home node of the port> -adapter <port name> -state down` command. Ports must be converted in pairs (for example, 0e and 0f). After conversion, a reboot is required, and the ports must be brought back to the up state.

Set Auto-Revert on Cluster Management

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

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

Run the following command:

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

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Ports used for data services (for example, e0b, e0d, e0e, e0f, e0g, e0h, e0j, e0k, and e0l) should be removed from the default broadcast domain, leaving just the management network ports (e0i and e0M). To perform this task, run the following commands:

```
broadcast-domain remove-ports –broadcast-domain Default –ports
<<var_node01>>:e0c,<<var_node01>>:e0d,<<var_node01>>:e0e,<<var_node01>>:e0f,<<var_node01>>:e0g,<<var_node01>>
:e0h,<<var_node01>>:e0j,<<var_node01>>:e0k,<<var_node01>>:e0l,<<var_node02>>:e0c,<<var_node02>>:e0d,<<var_nod
e02>>:e0e,<<var_node02>>:e0f,<<var_node02>>:e0g,<<var_node02>>:e0h,<<var_node02>>:e0j,<<var_node02>>:e0k,<<va
r_node02>>:e0l
broadcast-domain show
```

Set Up Service Processor Network Interface

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

```
system service-processor network modify –node <<var_node01>> -address-family IPv4 –enable true –dhcp none –
ip-address <<var_node01_sp_ip>> -netmask <<var_node01_sp_mask>> -gateway <<var_node01_sp_gateway>>

system service-processor network modify –node <<var_node02>> -address-family IPv4 –enable true –dhcp none –
ip-address <<var_node02_sp_ip>> -netmask <<var_node02_sp_mask>> -gateway <<var_node02_sp_gateway>>
```

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

Create Aggregates

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

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <<var_node01>> -diskcount <<var_num_disks>>
aggr create -aggregate aggr1_node02 -node <<var_node02>> -diskcount <<var_num_disks>>
```

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

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

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

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

```
node run <<var_node01>> aggr options aggr1_node01 nosnap on
node run <<var_node02>> aggr options aggr1_node02 nosnap on
```

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

```
node run <<var_node01>> snap delete –A –a –f aggr1_node01
node run <<var_node02>> snap delete –A –a –f aggr1_node02
```

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

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

Verify Storage Failover

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

1. Verify the status of storage failover.

```
storage failover show
```

Both the nodes `<<var_node01>>` and `<<var_node02>>` must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

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

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

4. Continue with step 6 if high availability is configured.

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

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

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

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

Disable Flow Control on UTA2 Ports

NetApp recommends disabling flow control on all of the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps:

1. Run the following commands to configure node 01:

```
network port modify -node <<var_node01>> -port e0b,e0d,e0e,e0f,e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <<var_node02>> -port e0b,e0d,e0e,e0f,e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show –fields flowcontrol-admin
```

Disable Unused FC Capability on CNA Ports

If a UTA2 port is set to CNA mode, and is only expected to handle Ethernet data traffic (e.g., NFS), then the unused FC capability of the port should be disabled by setting the corresponding FCP adapter to state down, with the `fcp adapter modify` command. Here are some examples:

```
fcp adapter modify -node <<var_node01>> -adapter 0g –state down
fcp adapter modify -node <<var_node01>> -adapter 0h –state down
fcp adapter modify -node <<var_node02>> -adapter 0g –state down
fcp adapter modify -node <<var_node02>> -adapter 0h –state down
fcp adapter show –fields state
```

Configure NTP

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

1. Set the time zone for the cluster.

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

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

2. Set the date for the cluster.

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

The format for the date is `<[Century][Year][Month][Day][Hour][Minute].[Second]>` (for example, `201309081735.17`).

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

```
cluster time-service ntp server create -server <<var_switch_a_ntp_ip>>
cluster time-service ntp server create -server <<var_switch_b_ntp_ip>>
```

Configure SNMP

To configure the Simple Network Management Protocol (SNMP), complete the following steps:

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

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

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

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

Configure SNMPv1 Access

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

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

Create SNMPv3 User

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

1. Create a user called `snmpv3user`.

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

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

Configure AutoSupport

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

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

Enable Cisco Discovery Protocol

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

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

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

Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

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

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
```

If you are using iSCSI, run the following commands to create the broadcast domains for iSCSI on ONTAP:

```
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, run the following commands:

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0b
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0d

ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0b
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0d

ifgrp show
```

Create VLANs

To create VLANs, create NFS VLAN ports and add them to the NFS broadcast domain:

```
network port modify –node <<var_node01>> -port a0a –mtu 9000
network port modify –node <<var_node02>> -port a0a –mtu 9000
network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <<var_node01>>:a0a-<<var_nfs_vlan_id>>,
<<var_node02>>:a0a-<<var_nfs_vlan_id>>
```

If you are using iSCSI, create iSCSI VLAN ports and add them to the iSCSI broadcast domains:

```
network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_node01>>:a0a-
<<var_iscsi_vlan_A_id>>,<<var_node02>>:a0a-<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_node01>>:a0a-
<<var_iscsi_vlan_B_id>>,<<var_node02>>:a0a-<<var_iscsi_vlan_B_id>>
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

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

2. Select the SVM data protocols to configure, keeping `fcp`, `iscsi`, and `nfs`.

```
vserver remove-protocols –vserver Infra-SVM -protocols cifs,ndmp
```

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

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

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

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

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plugin.

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

Create Load-Sharing Mirrors of SVM Root Volume

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

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

```
volume create –vserver Infra-SVM –volume rootvol_m01 –aggregate aggr1_node01 –size 1GB –type DP
volume create –vserver Infra-SVM –volume rootvol_m02 –aggregate aggr1_node02 –size 1GB –type DP
```

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

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

3. Create the mirroring relationships.

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

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set –source-path //Infra-SVM/rootvol
snapmirror show
```

Create Block Protocol (iSCSI, FC) Service

If you are using iSCSI, run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

```
iscsi create -vserver Infra-SVM
iscsi show
```

If you are using FC, run the following command to create the FCP service on each SVM. This command also starts the FCP service and sets the WWN for the SVM.

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

Configure HTTPS Access

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

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

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

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

```
security certificate show
```

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

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

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

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

```
security certificate create -common-name <<var_cert_common_name>> -type  server -size 2048 -country
<<var_cert_country>> -state <<var_cert_state>> -locality <<var_cert_locality>> -organization <<var_cert_org>>
-unit <<var_cert_unit>> -email-addr <<var_cert_email>> -expire-days <<var_cert_days>> -protocol SSL -hash-
function SHA256 -vserver Infra-SVM
```

4. To obtain the values for the parameters required in step 6 (<<var_cert_ca>> and <<var_cert_serial>>), run the `security certificate show` command.

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

```
security ssl modify -vserver <<var_clustername>> -server-enabled true -client-enabled false -ca
<<var_cert_ca>> -serial <<var_cert_serial>> -common-name <<var_cert_common_name>>
```

6. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
```

```
Warning: Modifying the cluster configuration will cause pending web service requests to be interrupted as the
web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http –vserver <<var_clustername>>
```

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

7. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set –privilege admin
vserver services web modify –name compat –vserver * -enabled true
```

Configure NFSv3

To configure NFSv3 on the SVM, complete the following steps:

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for each ESXi host created so that each host has its own rule index. For example, the first ESXi host has rule index 1, the second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create –vserver Infra-SVM -policyname default –ruleindex 1 –protocol nfs -
clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys –rwrule sys -superuser sys –allow-suid false
vserver export-policy rule create –vserver Infra-SVM -policyname default –ruleindex 2 –protocol nfs -
clientmatch <<var_esxi_host2_nfs_ip>> -rorule sys –rwrule sys -superuser sys –allow-suid false
vserver export-policy rule show
```

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

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

Create FlexVol Volumes

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

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

To create a FlexVol volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node02 -size 500GB -state online
-policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-snapshot-space 0

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

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

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

Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 15GB -ostype vmware -space-
reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-02 -size 15GB -ostype vmware -space-reserve
disabled
```

Schedule Deduplication

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

1. After the volumes are created, assign a once-a-day dedup schedule to `esxi_boot`:

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

2. Create the Always_On_Deduplication efficiency policy:

```
cron create -name 1min –minute
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31,32,33,34,35,36,37,38,39
,40,41,42,43,44,45,46,47,48,48,50,51,52,53,54,55,56,57,58,59
efficiency policy create -vserver Infra-SVM -policy Always_On_Deduplication -type scheduled -schedule 1min -
qos-policy background -enabled true
```

3. Optionally, assign the Always On Deduplication policy to `infra_datastore_1`:

```
efficiency modify -vserver Infra-SVM -volume infr_datastore_1 -policy Always-On-Deduplication
```

4. If you do not want to assign an Always On Deduplication policy to `infra_datastore_1`, assign the once-a-day deduplication schedule:

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

Create iSCSI LIFs

If you are using iSCSI, run the following commands to create four iSCSI LIFs (two on each node):

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-node
<<var_node01>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address <<var_node01_iscsi_lif01a_ip>> -netmask
<<var_node01_iscsi_lif01a_mask>> -status-admin up –failover-policy disabled –firewall-policy data –auto-
revert false

network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -home-node
<<var_node01>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address <<var_node01_iscsi_lif01b_ip>> -netmask
<<var_node01_iscsi_lif01b_mask>> -status-admin up –failover-policy disabled –firewall-policy data –auto-
revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -home-node
<<var_node02>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address <<var_node02_iscsi_lif01a_ip>> -netmask
<<var_node02_iscsi_lif01a_mask>> -status-admin up –failover-policy disabled –firewall-policy data –auto-
revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-node
<<var_node02>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address <<var_node02_iscsi_lif01b_ip>> -netmask
<<var_node02_iscsi_lif01b_mask>> -status-admin up –failover-policy disabled –firewall-policy data –auto-
revert false
network interface show
```

Create FCP LIFs

If you are using FCP, run the following commands to create four FCP LIFs (two on each node) per attached fabric interconnect:

```
network interface create -vserver Infra-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node
<<var_node01>> -home-port <<var_node01_fcp_port1>> –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node
<<var_node01>> -home-port <<var_node01_fcp_port2>> –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node
<<var_node02>> -home-port <<var_node02_fcp_port1>> –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node
<<var_node02>> -home-port <<var_node02_fcp_port2>> –status-admin up
network interface show
```

Create NFS LIF

To create an NFS LIF, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_infra_swap -role data -data-protocol nfs -home-node
<<var_node01>> -home-port a0a-<<var_nfs_vlan_id>> –address <<var_node01_nfs_lif_infra_swap_ip>> -netmask
<<var_node01_nfs_lif_infra_swap_mask>> -status-admin up –failover-policy broadcast-domain-wide –firewall-
policy data –auto-revert true

network interface create -vserver Infra-SVM -lif nfs_infra_datastore_1 -role data -data-protocol nfs -home-
node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>> –address <<var_node02_nfs_lif_infra_datastore_1_ip>> -
netmask <<var_node02_nfs_lif_infra_datastore_1_mask>> -status-admin up –failover-policy broadcast-domain-wide
–firewall-policy data –auto-revert true

network interface show
```

NetApp recommends creating a new LIF for each datastore.

Add Infrastructure SVM Administrator

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

1. Run the following commands:

```
network interface create –vserver Infra-SVM –lif vsmgmt –role data –data-protocol none –home-node
<<var_node02>> -home-port  e0i –address <<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up –
failover-policy broadcast-domain-wide –firewall-policy mgmt –auto-revert true
```

The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

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

```
network route create –vserver Infra-SVM -destination 0.0.0.0/0 –gateway <<var_svm_mgmt_gateway>>

network route show
```

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

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

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

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

Server Configuration

Cisco UCS Base Configuration

This FlexPod deployment will show configuration steps for both the Cisco UCS 6332-16UP and Cisco UCS 6248UP Fabric Interconnects (FI) in a design that will support iSCSI as well as Fibre Channel direct attached connectivity to the NetApp AFF. Implementation of both of these protocols simultaneously should not be considered mandatory, and the selection of one or the other should be acceptable depending upon your environment and preferences.

Configuration steps will be referenced for both fabric interconnects and will be called out by the specific model where steps have differed.

Perform Initial Setup of Cisco UCS 6332-16UP and 6248UP Fabric Interconnects for FlexPod Environments

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

Cisco UCS 6332-16UP A

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

1. Connect to the console port on the first Cisco UCS 6332-16UP fabric interconnect.

```
Enter the configuration method: 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 passwords? (y/n) [y]: y

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

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

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

Which switch fabric (A|B): A

Enter the system name: <<var_ucs_clustername>>

Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>

Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>

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

Cluster IPv4 address: <<var_ucs_cluster_ip>>

Configure DNS Server IPv4 address? (yes/no) [no]: y

DNS IPv4 address: <<var_nameserver_ip>>
```

```
Configure the default domain name? y

Default domain name: <<var_dns_domain_name>>

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

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

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

Cisco UCS 6332-16UP B

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

1. Connect to the console port on the second Cisco UCS 6332-16UP fabric interconnect.

```
Enter the configuration method: 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: <<var_password>>

Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>

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

(yes/no): y
```

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

3. Repeat these steps for the 6248 Fabric Interconnects.

Cisco UCS Setup

Log in to Cisco UCS Manager

The steps are the same between the UCS 6332-16UP and the UCS 6248UP Fabric Interconnects unless otherwise noted

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 fabric interconnect cluster address.

2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.

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

4. When prompted, enter `admin` as the user name and enter the administrative password.

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

Upgrade Cisco UCS Manager Software to Version 3.1(1h)

This document assumes the use of Cisco UCS 3.1(1h). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.1(1h), refer to Cisco UCS Manager Install and Upgrade Guides.

Anonymous Reporting

To create anonymous reporting, complete the following step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products:

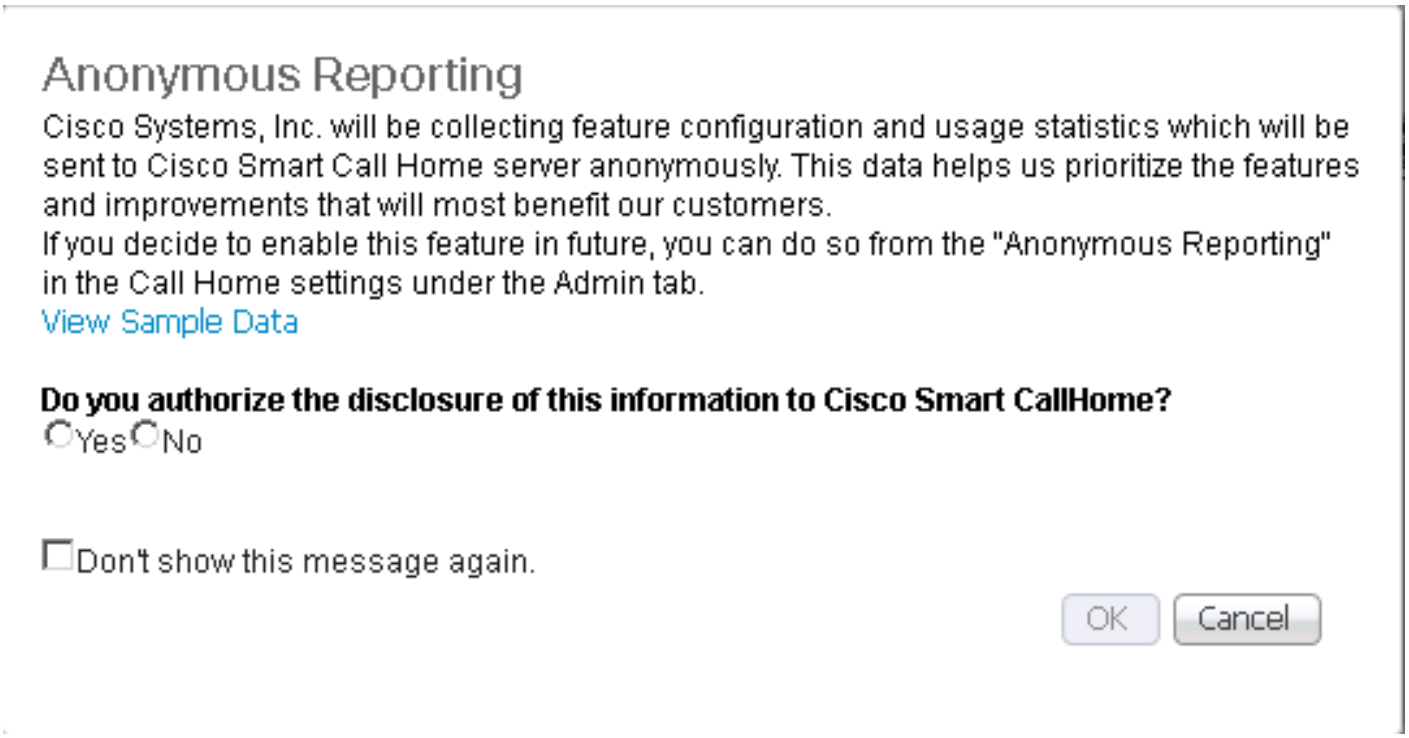

Add Block of IP Addresses for KVM Access

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

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

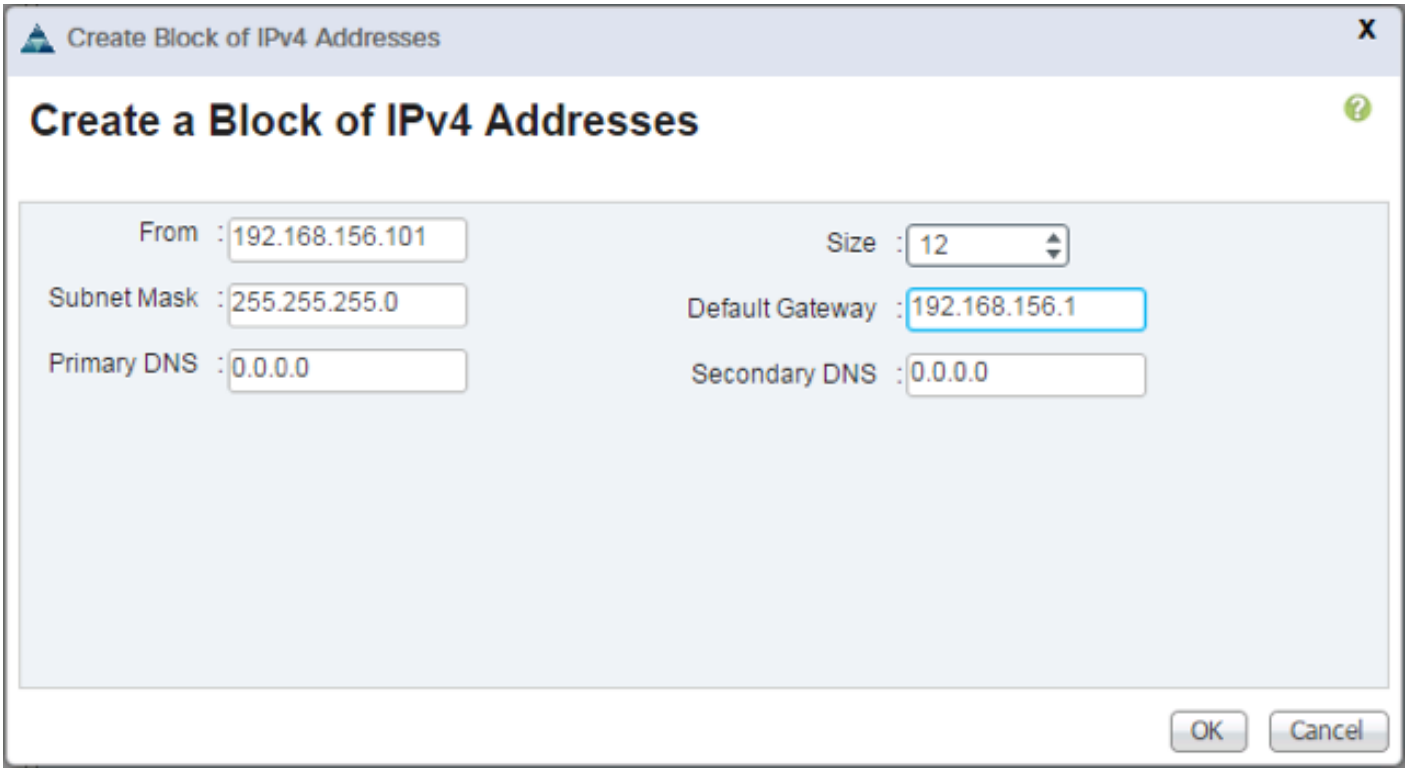

5. Click OK to create.
6. Click OK in the confirmation message.

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.

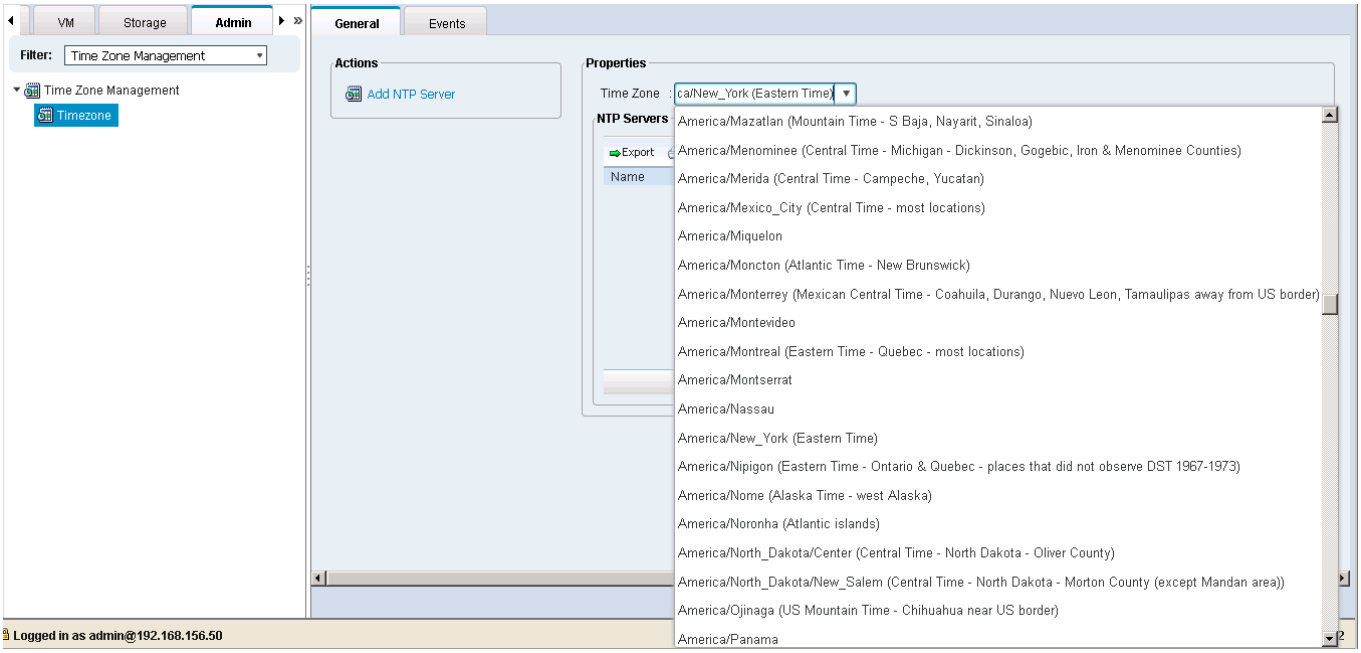

3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.
6. Enter <<var_switch_a_ntp_ip>> and click OK.

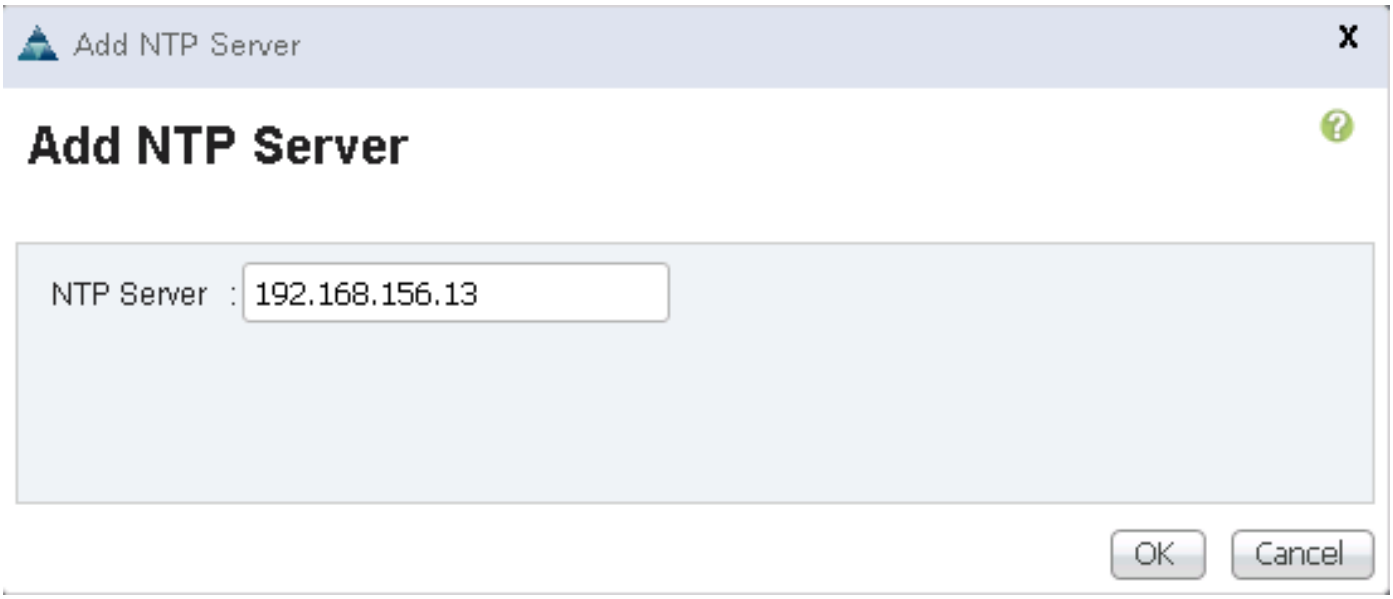

7. Click Add NTP Server.
8. Enter <<var_switch_b_ntp_ip>> and click OK.

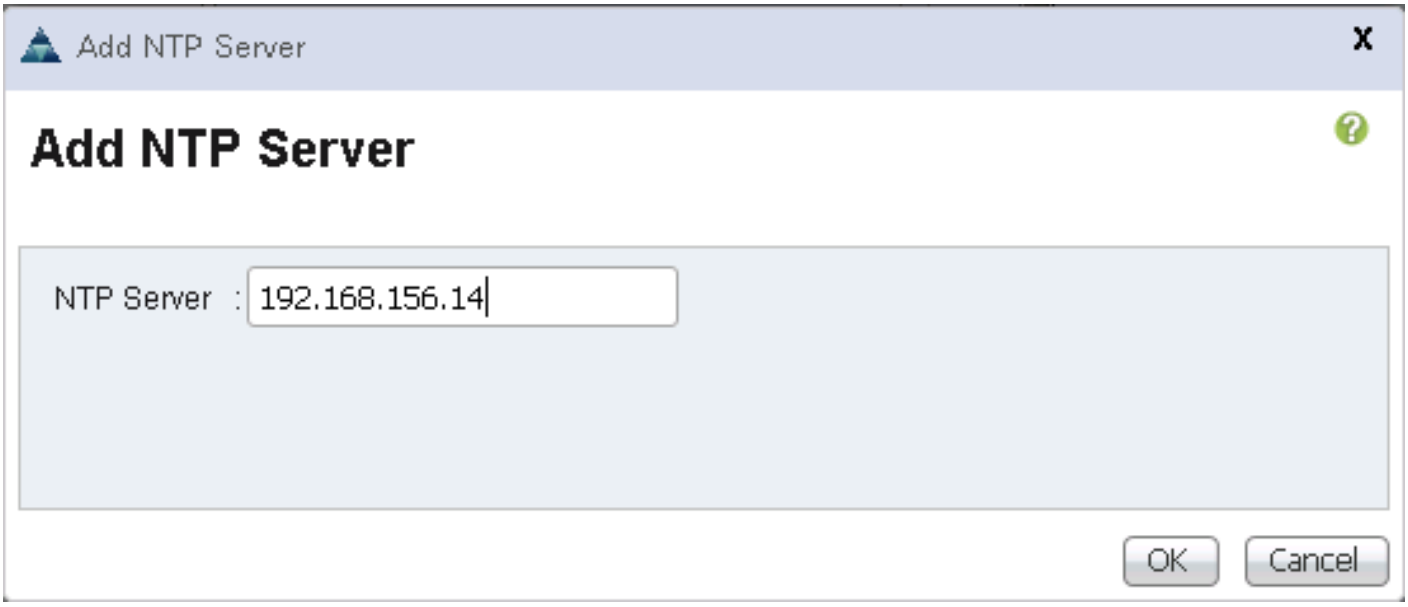

9. Click OK.

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 in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.

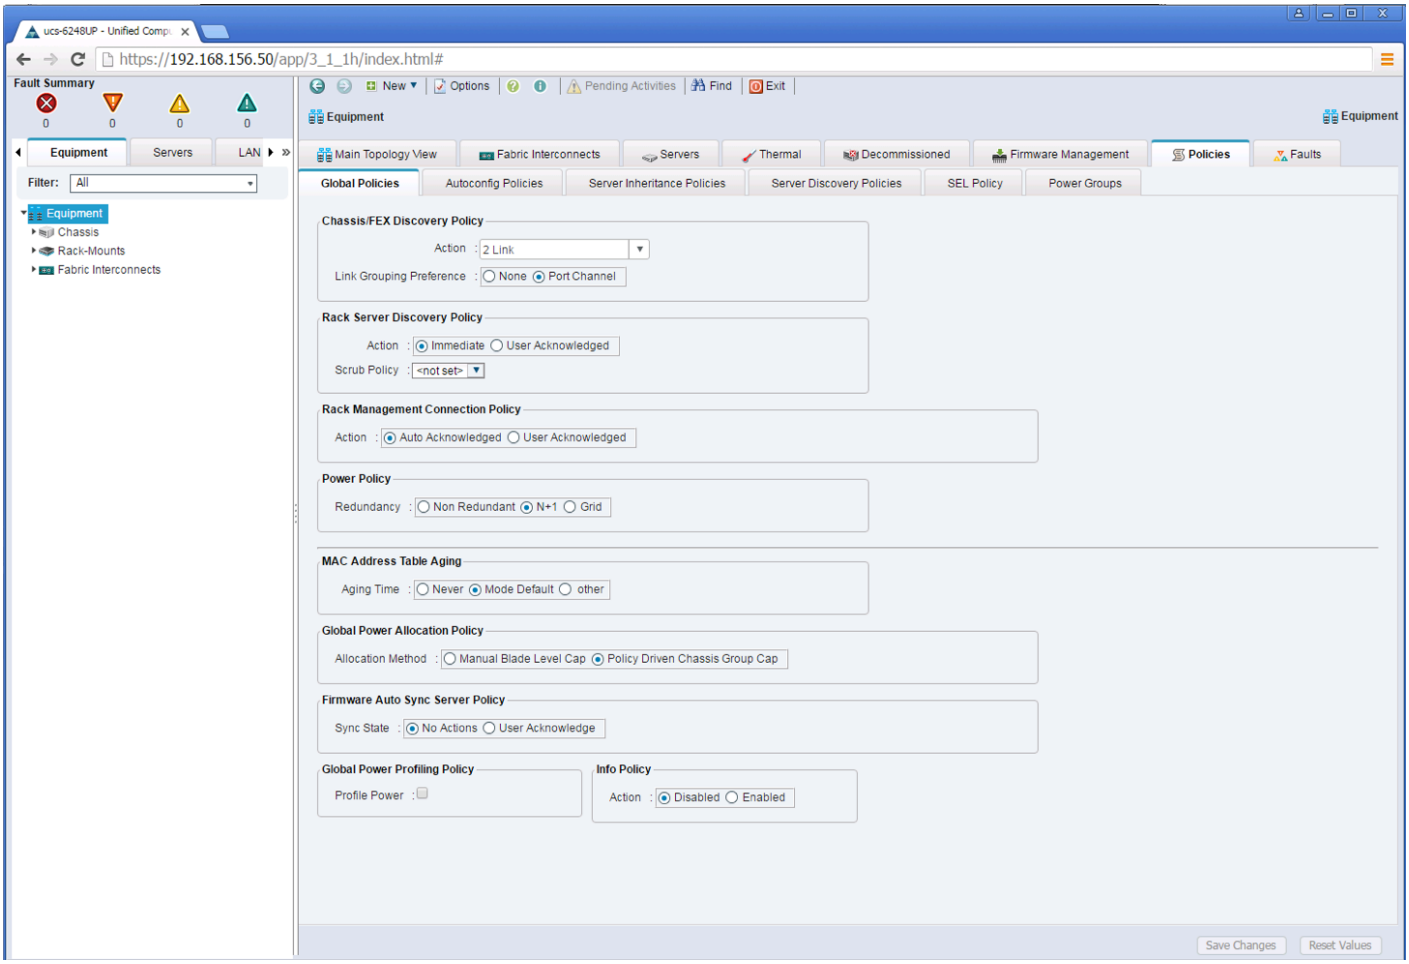

5. Click Save Changes.

6. Click OK.

Enable FC Switching

To use direct attached Fibre Channel connectivity, the fabric interconnects will need to be placed in Fibre Channel Switching mode, by completing the following steps:

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

2. Expand Fabric Interconnects and select either Fabric Interconnect.

This next step will reboot both UCS Fabric Interconnects. If any servers are running on this system, they should be shut down before this step is executed.

3. In the Actions pane, select Set FC Switching Mode. Click Yes. Click OK.

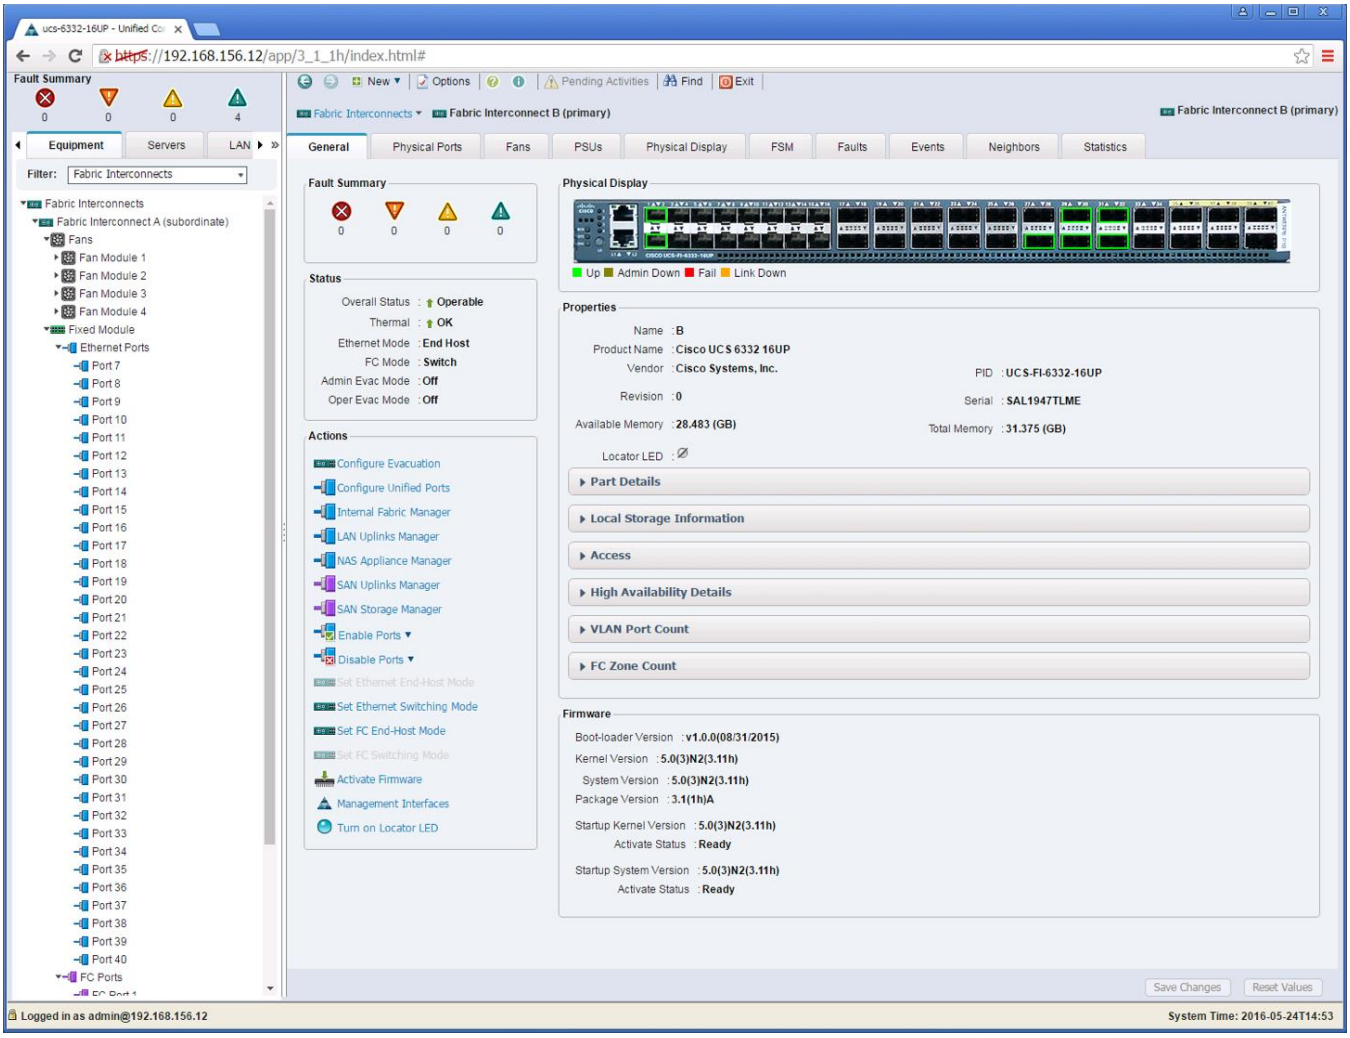

4. After the Fabric Interconnects have rebooted, log back into UCS Manager.

5. Expand Fabric Interconnects and select Fabric Interconnects.

6. For each Fabric Interconnect, verify under Status that the FC Mode is now Switch.

Configure Unified Ports

Fibre Channel port configurations will slightly differ between the 6332-16UP and the 6248UP Fabric Interconnects. Both Fabric Interconnects will have a slider mechanism within the UCSM GUI interface, but the fibre channel port selection options for the 6332-16UP will be from the first 16 ports starting from the first port, and configured in increments of the first 6, 12, or all 16 of the unified ports. With the 6248UP, the port selection options will start from the upper end of the 32 fixed ports, or the upper end of the 16 ports of the expansion module, going down in contiguous increments of 2.

To enable the fibre channel ports, complete the following steps for the 6332-16UP:

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

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

3. Select Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
5. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 6, 12, or 16 ports to be set as FC Uplinks.

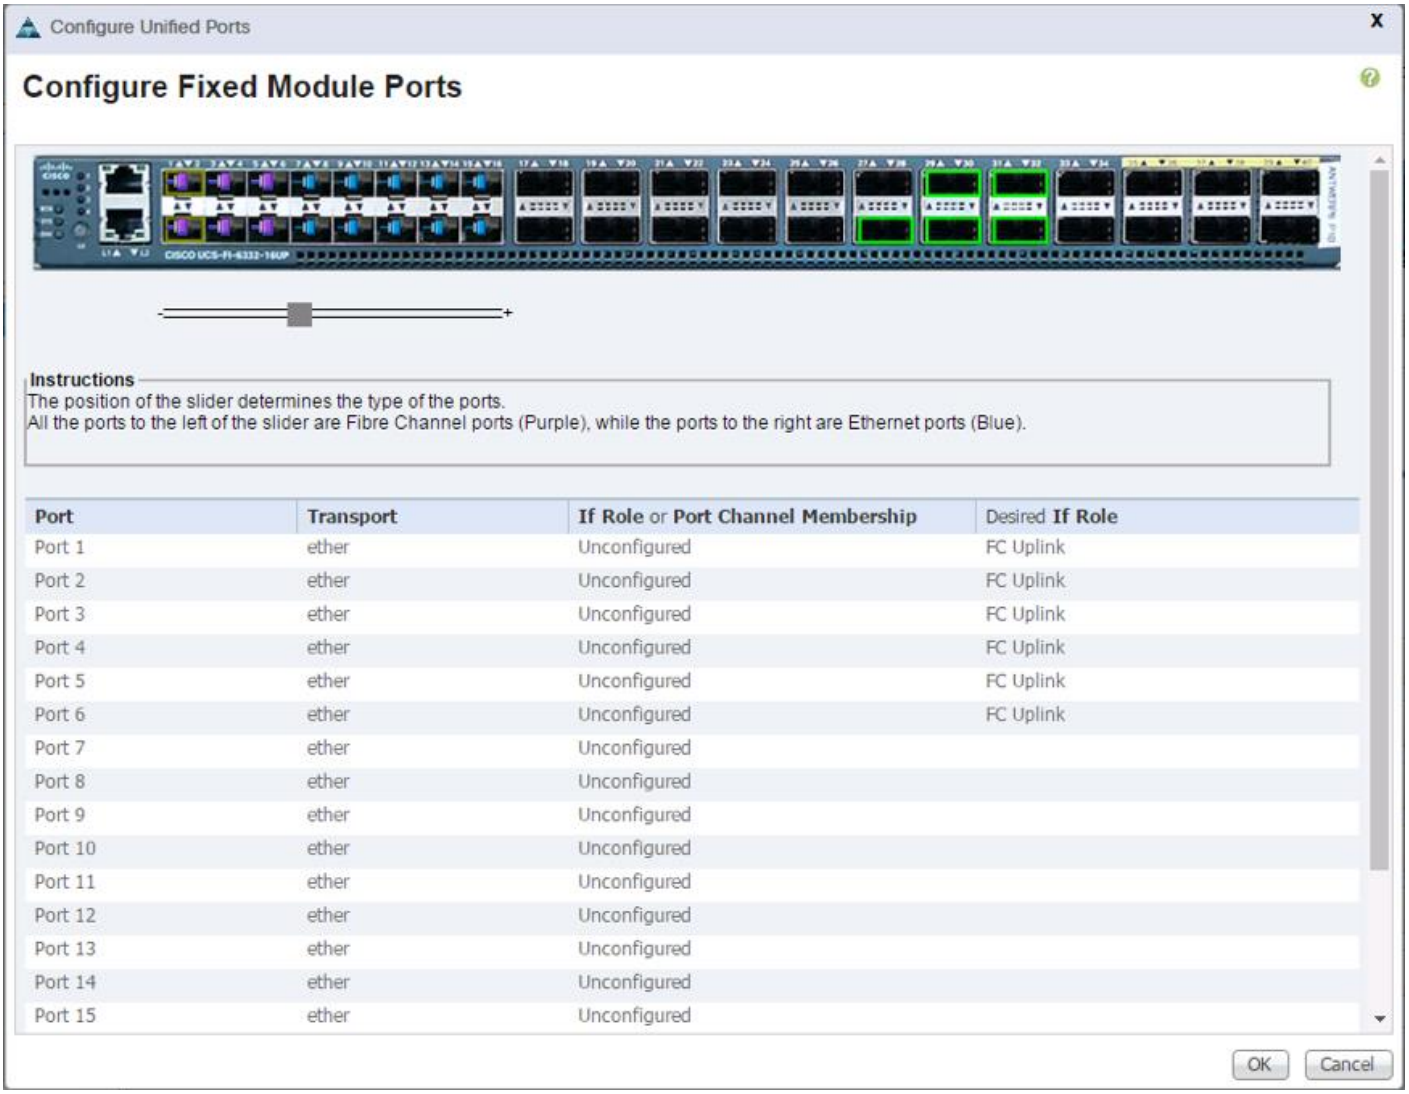

6. Click OK to continue
7. Select Equipment > Fabric Interconnects > Fabric Interconnect B (primary)
8. Select Configure Unified Ports.
9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
10. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 6, 12, or 16 ports to be set as FC Uplinks.

11. Click OK to continue

To enable the fibre channel ports, complete the following steps for the 6248UP:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary)
3. Select Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
5. Configure the Fixed Module Ports from the subsequent Configure Fixed Module Ports pop-up window, or click on the Configure Expansion Module Ports button to select from expansion module ports.
6. Within either option (Expansion Module shown below) move the gray slider bar from the right to the left selecting ports in increments of two to set as FC Uplinks.

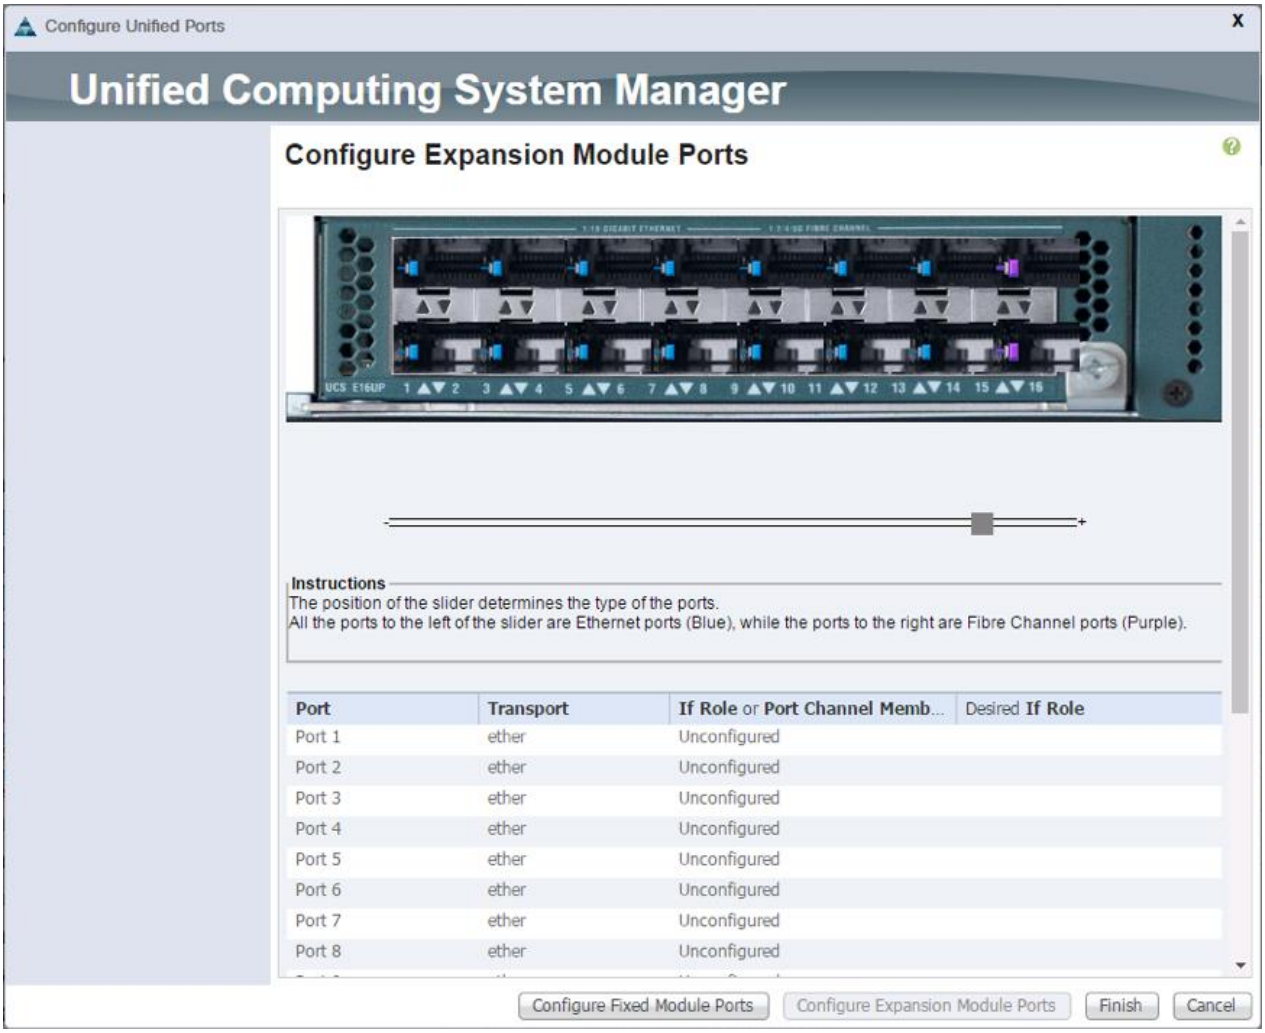

7. Click Finish.

8. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate)
9. Select Configure Unified Ports.
10. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
11. Configure the Fixed Module Ports from the subsequent Configure Fixed Module Ports pop-up window, or click on the Configure Expansion Module Ports button to select from expansion module ports.
12. Within either option move the gray slider bar from the right to the left selecting ports in increments of two to set as FC Uplinks.
13. Click Finish.

Enable FC Storage Ports

To enable FC 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. (or Expansion Module if FC ports were selected from the Expansion Module in the 6248UP)
3. Expand FC Ports.
4. Select ports that are connected to the FC ports on the storage controllers, (this will be ports 1 and 2 in our 6332-16UP example and ports 15 and 16 of the Expansion Module in our 6248UP example), right-click them, and select Configure as FC Storage Port. Click Yes to confirm.
5. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module. (or Expansion Module if FC ports were selected from the Expansion Module in the 6248UP)
6. Expand FC Ports.
7. Select ports that are connected to the FC ports on the storage controllers, (this will be ports 1 and 2 in our 6332-16UP example and ports 15 and 16 of the Expansion Module in our 6248UP example), right-click them, and select Configure as FC Storage Port. Click Yes to confirm.

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 Ethernet Ports.
4. Select the ports that are connected to the chassis, and Cisco FEX, and direct connect UCS C-Series servers, right-click them, and select “Configure as Server Port.”

5. Click Yes to confirm server ports and click OK.
6. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
7. Select ports 31 and 32 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

The last 6 ports of the UCS 6332 and UCS 6332-16UP FIs will only work with optical based QSFP transceivers and AOC cables, so they can be better utilized as uplinks to upstream resources that might be optical only.

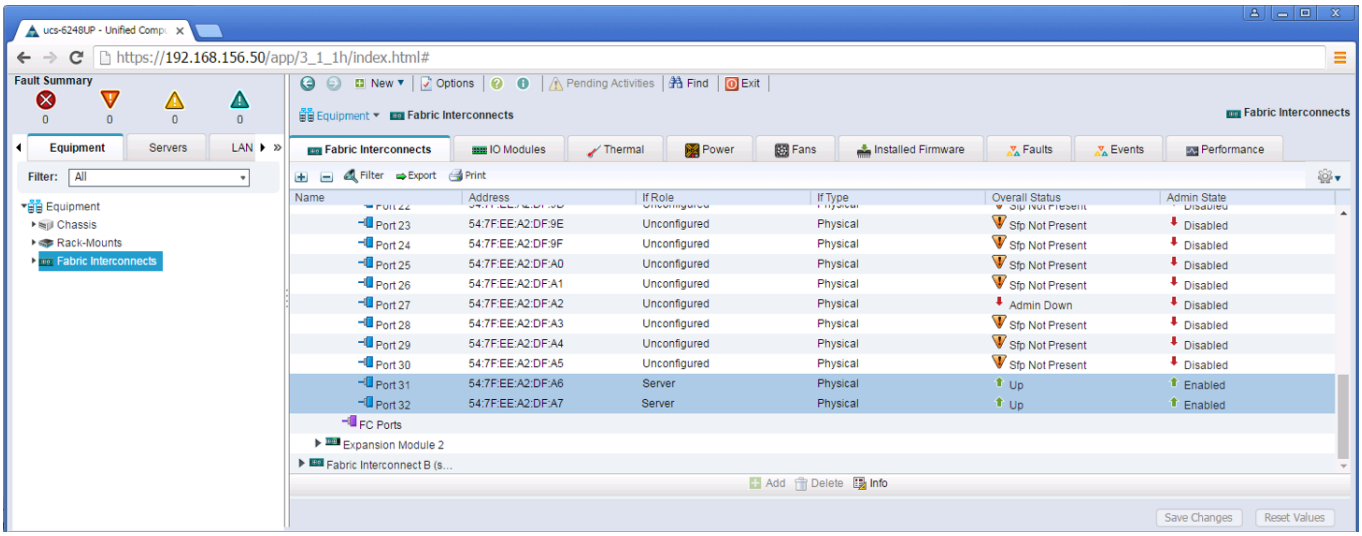

8. Click Yes to confirm uplink ports and click OK.
9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
10. Expand Ethernet Ports.
11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select ports 19 and 20 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external 2232 FEX modules, 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.

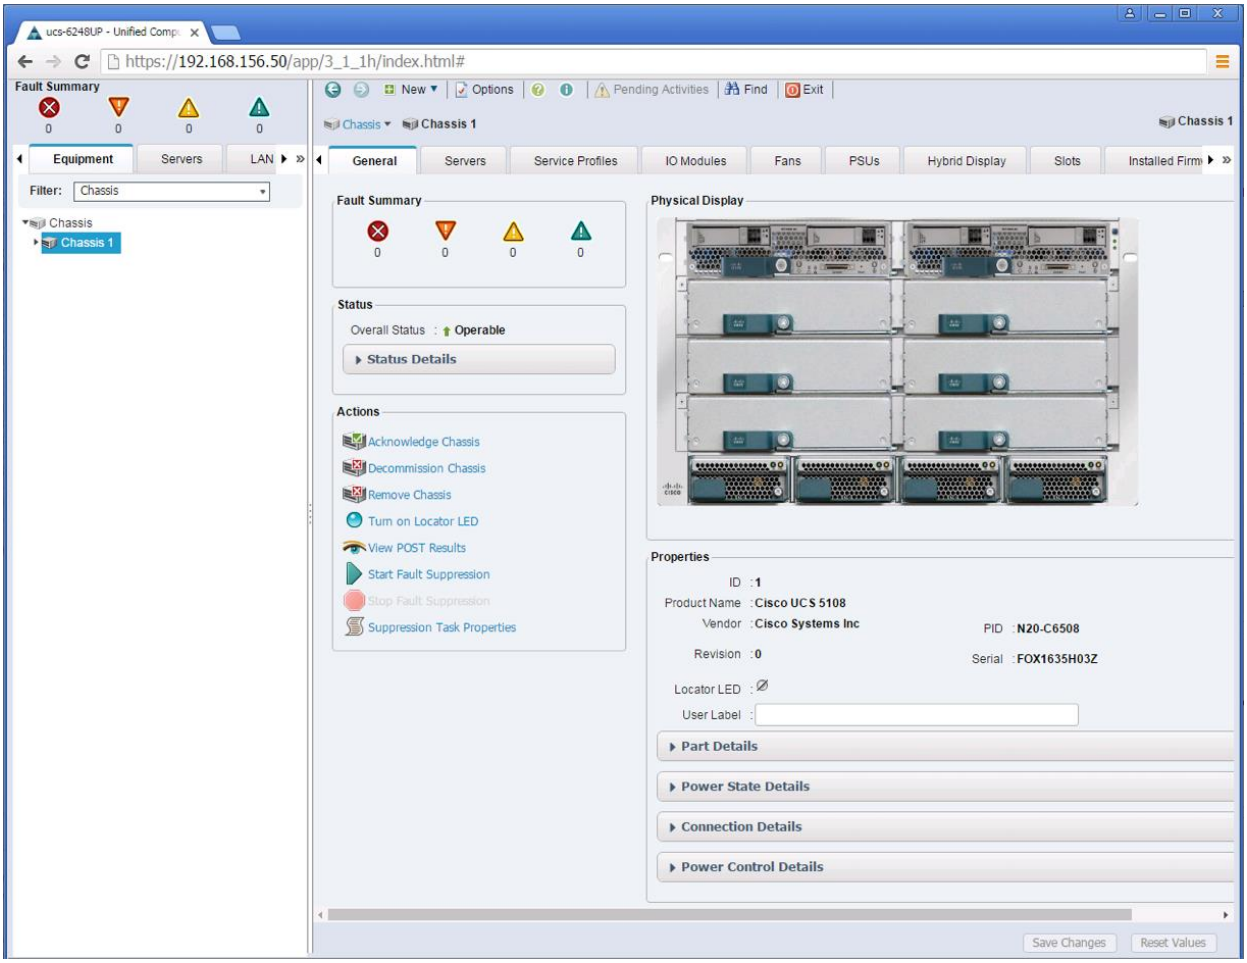

4. Click Yes and then click OK to complete acknowledging the chassis.

5. If the Nexus 2232 FEX is part of the configuration, expand Rack Mounts and FEX.

6. Right-click each FEX that is listed and select Acknowledge FEX.

7. Click Yes and then click OK to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus Switches

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

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

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

---

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

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

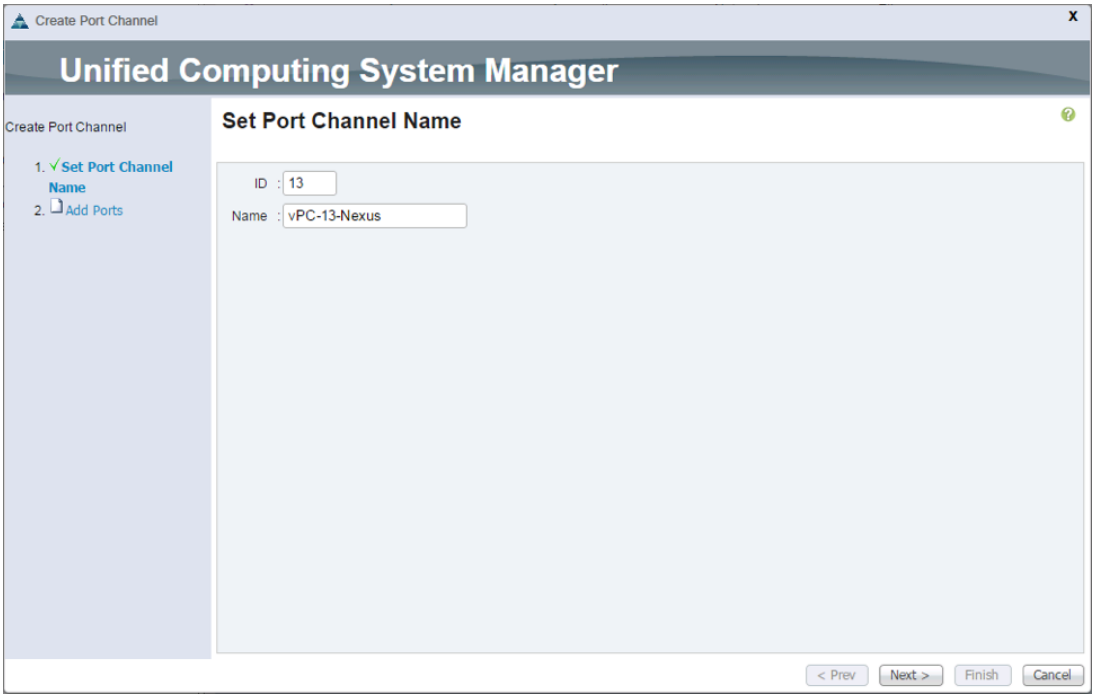

8. Select the following ports to be added to the port channel:
   - Slot ID 1 and port 31
   - Slot ID 1 and port 32
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter `14` as the unique ID of the port channel.
16. Enter `vPC-14-Nexus` as the name of the port channel.

17. Click Next.
18. Select the following ports to be added to the port channel:
    - Slot ID 1 and port 31
    - Slot ID 1 and port 32
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create a WWNN Pool for FC Boot

To configure the necessary WWNN pool for the Cisco UCS environment, complete the following steps on Cisco UCS Manager.

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

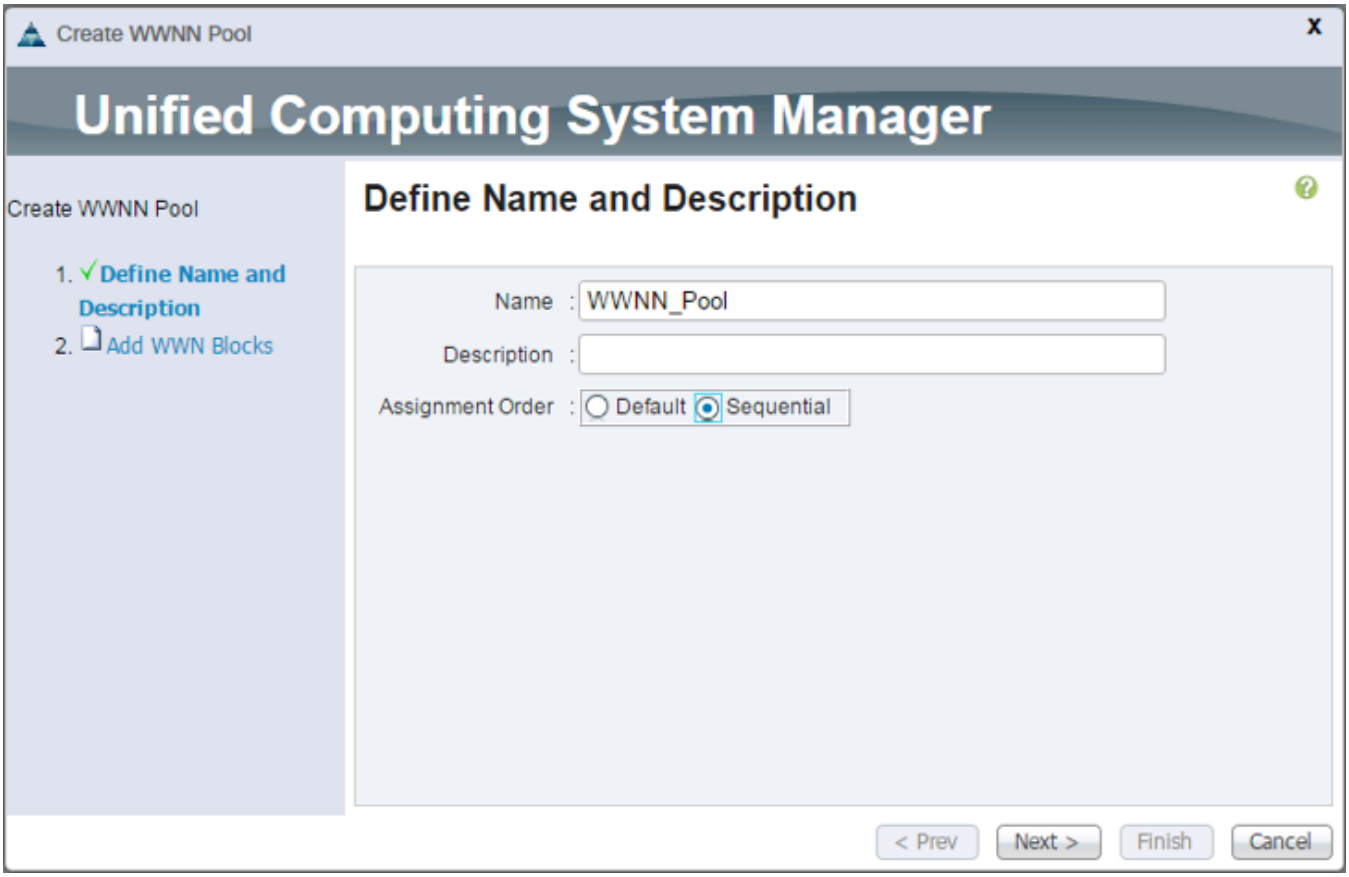

8. Click Next.

9. Click Add.

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

Modifications of the WWN block, as well as the WWPN and MAC Addresses, can convey identifying information for the UCS domain. Within the From field in our example, the 6th octet was changed from 00 to 91 to represent as identifying information for this being in building 9 on the 1st floor, and the 7th octet was changed from 00 to 10 to represent our first UCS domain.

Also, when having multiple UCS domains sitting in adjacency, it is important that these blocks, the WWN, WWPN, and MAC hold differing values between each set.

---

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

12. Click OK.

Create WWPN Pools

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. In this procedure, two WWPN pools are created, one for each switching fabric.
4. Right-click WWPN Pools under the root organization.
5. Select Create WWPN Pool to create the WWPN pool.
6. Enter `WWPN_Pool_A` as the name of the WWPN pool.
7. Optional: Enter a description for the WWPN pool.
8. Select Sequential for Assignment Order.

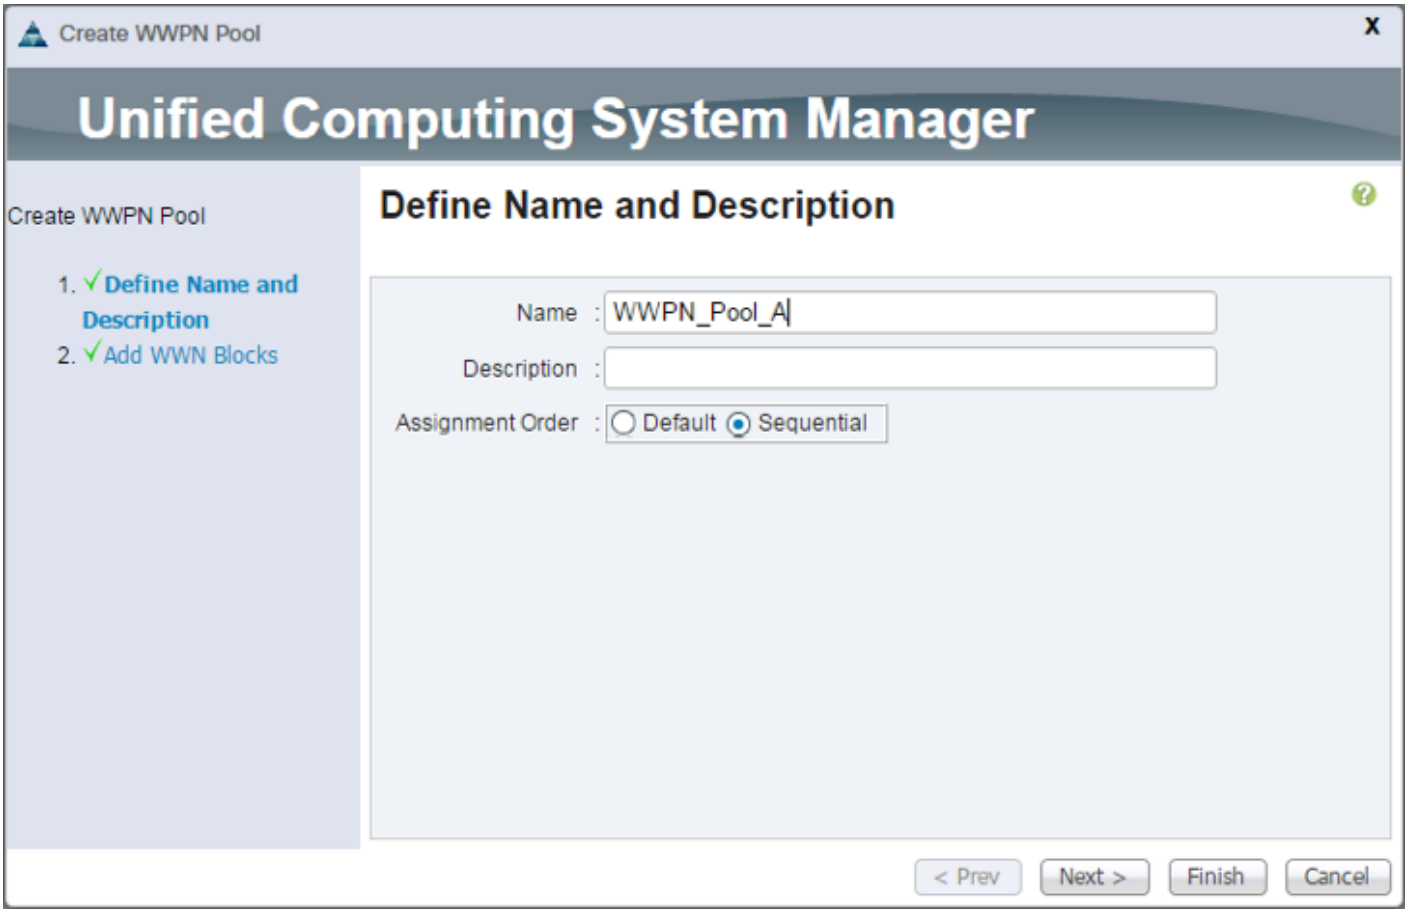

9. Click Next.
10. Click Add.
11. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:91:1A:00`.

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

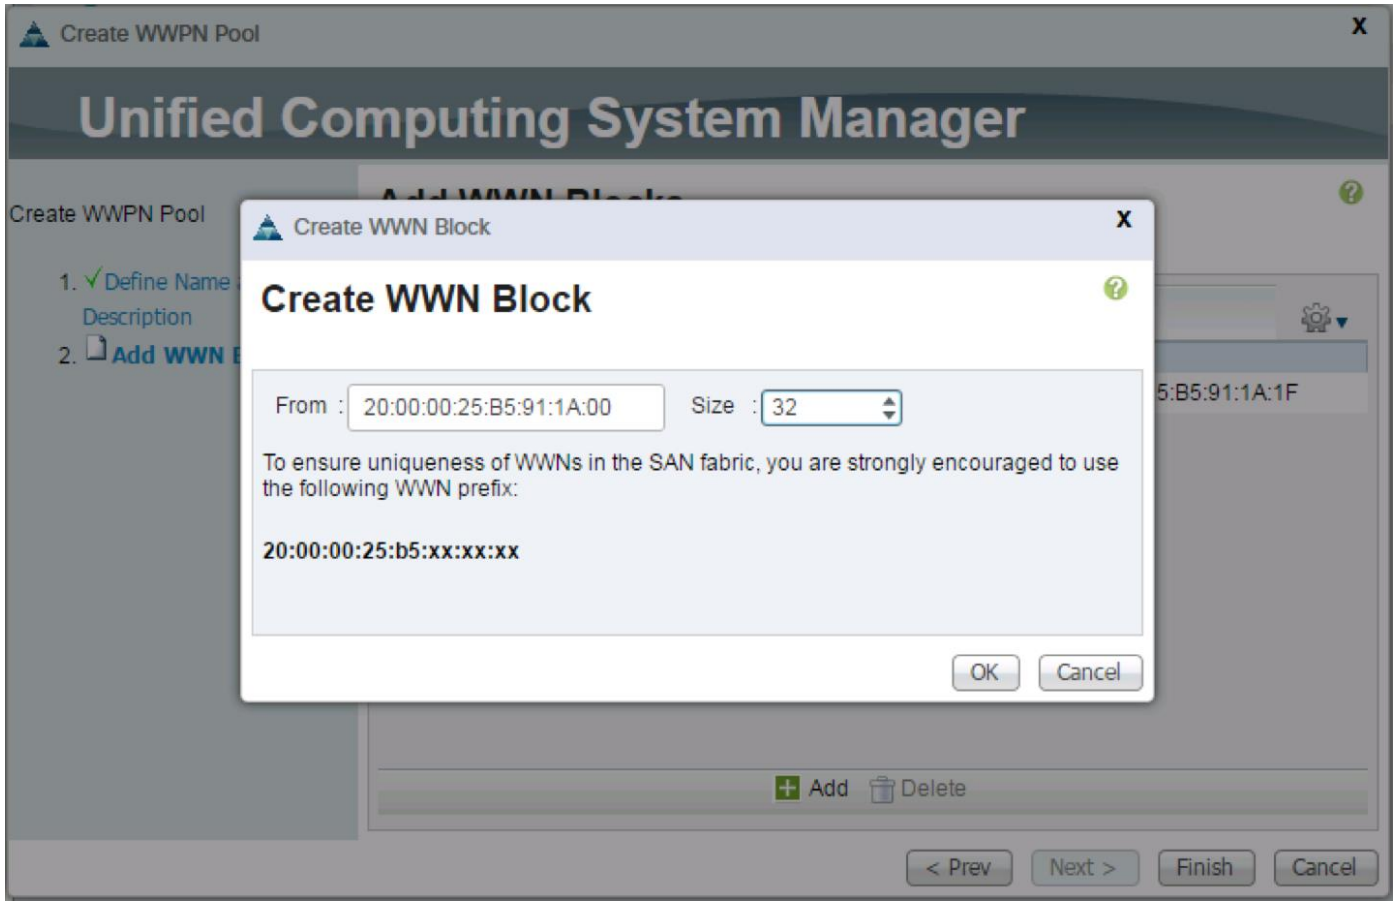

13. Click OK.
14. Click Finish.
15. In the confirmation message, click OK.
16. Right-click WWPN Pools under the root organization.
17. Select Create WWPN Pool to create the WWPN pool.
18. Enter `WWPN_Pool_B` as the name of the WWPN pool.
19. Optional: Enter a description for the WWPN pool.
20. Select Sequential for Assignment Order.

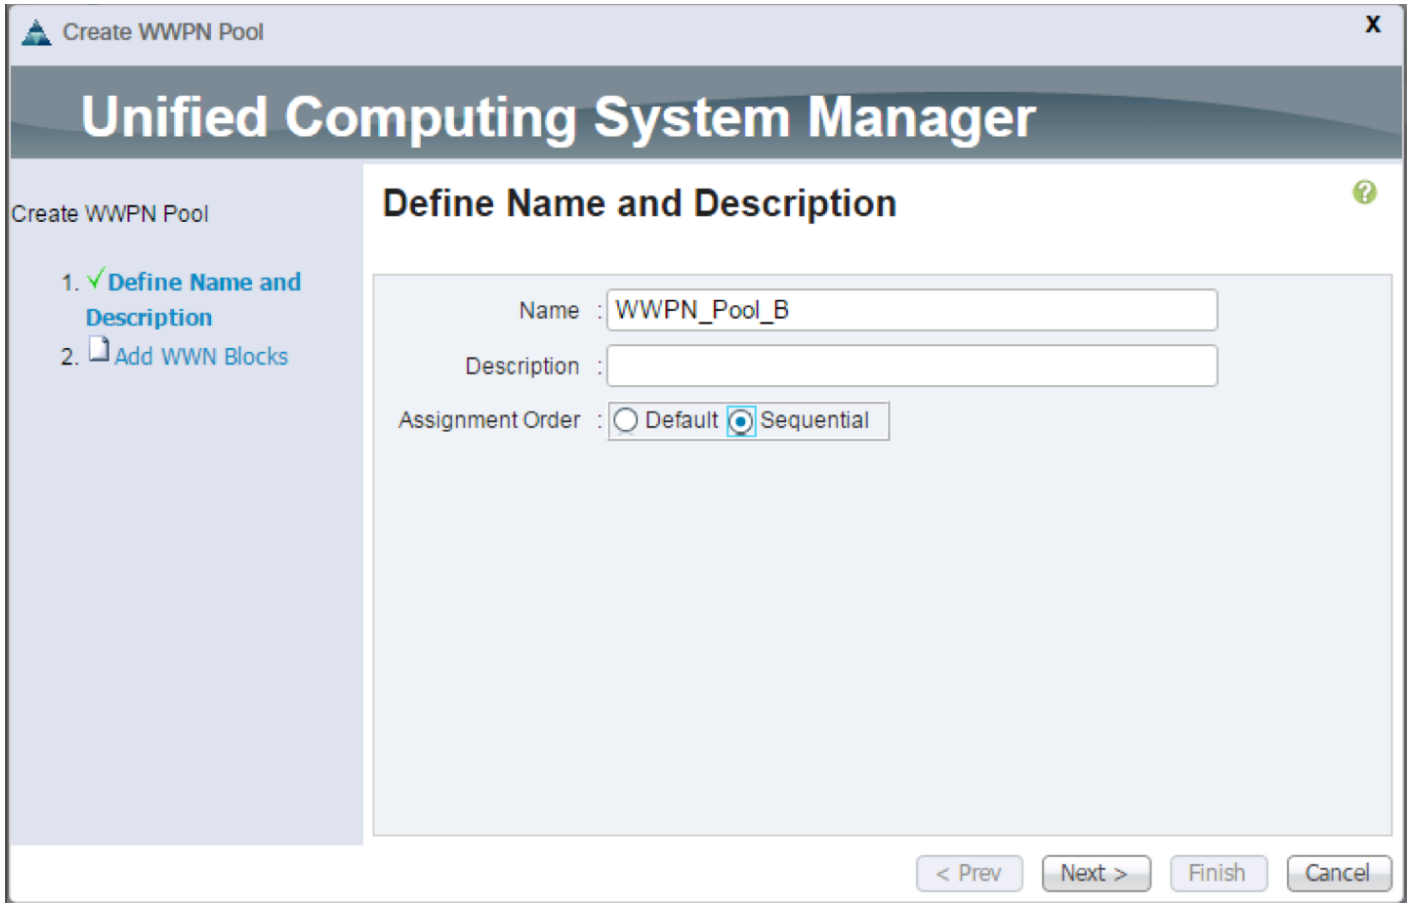

21. Click Next.
22. Click Add.
23. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:91:1AB:00`.

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

25. Click OK.
26. Click Finish.
27. In the confirmation message, click OK.

Create VSANs

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

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

In this procedure, two VSANs are created.

---

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

7. Select Fabric A.

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

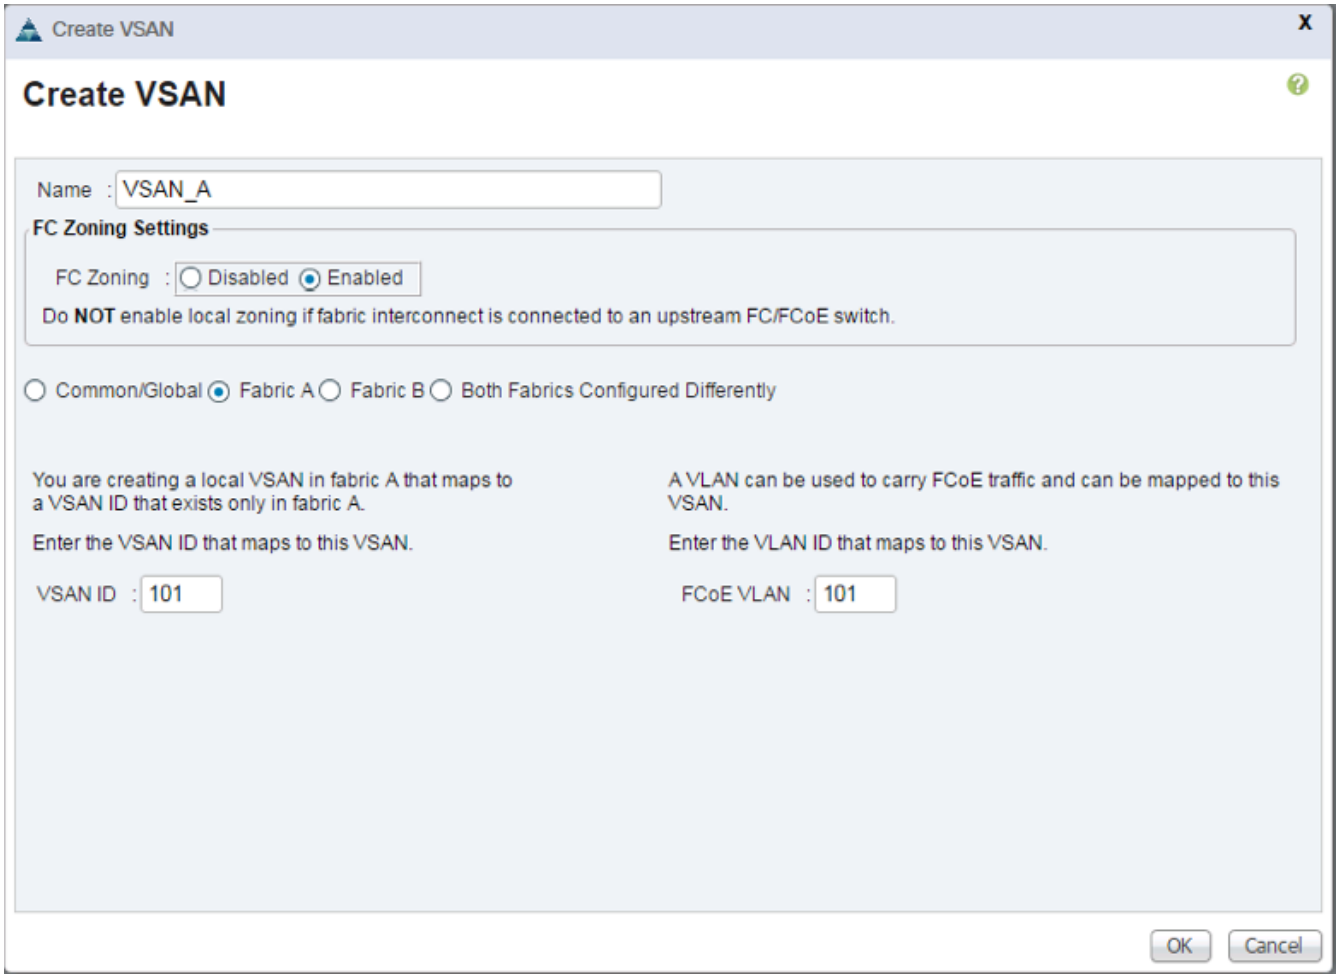

9. Click OK, and then click OK again.

10. Under SAN Cloud, right-click VSANs.

11. Select Create VSAN.

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

13. Select Enabled for FC Zoning.

14. Select Fabric B.

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

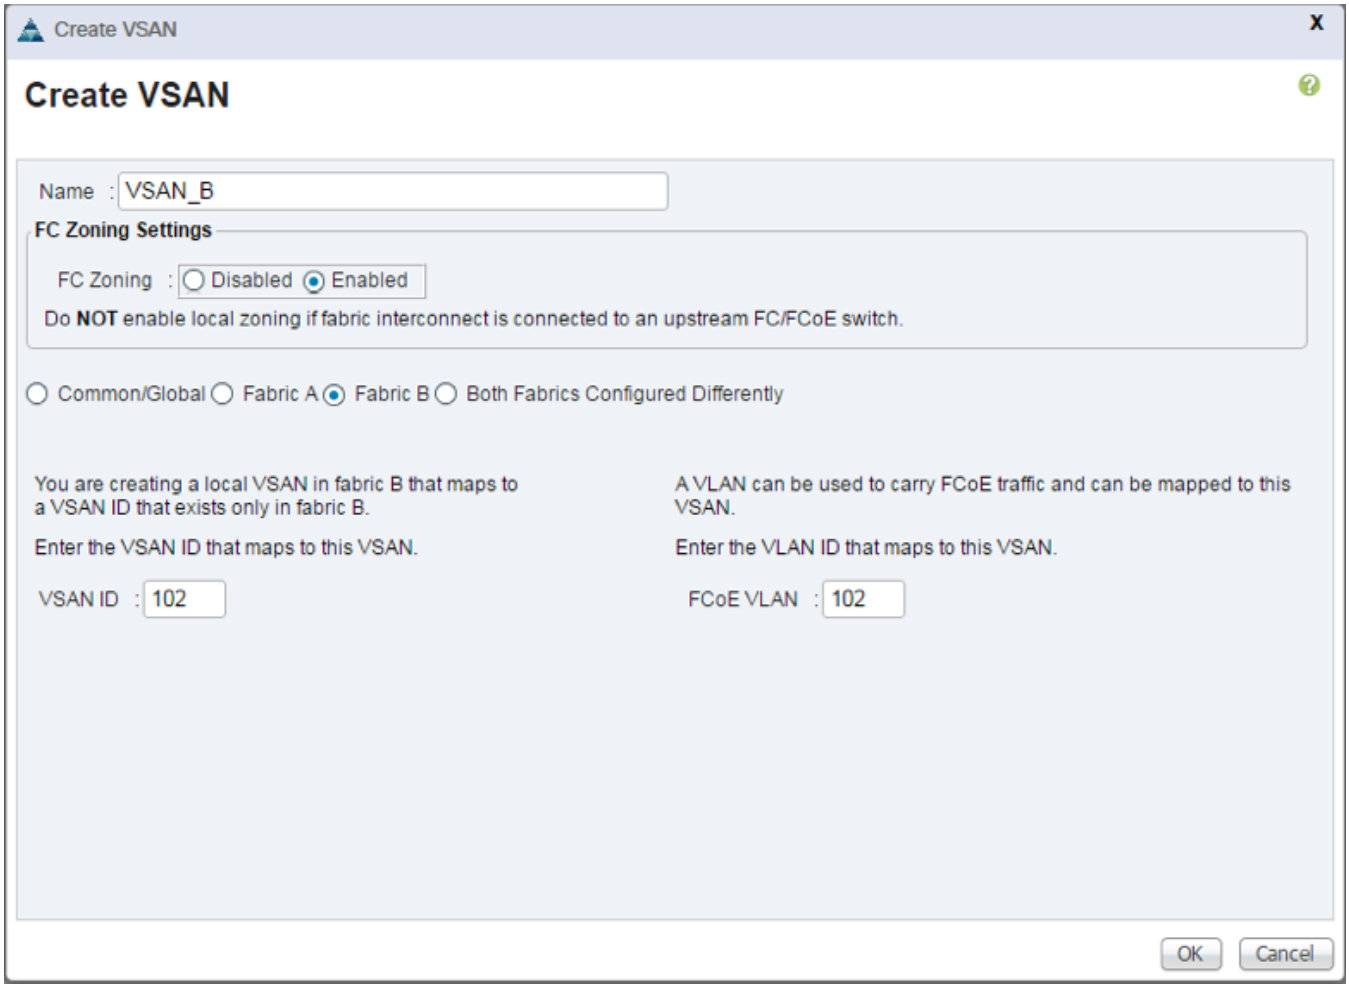

16. Click OK, and then click OK again.
17. Under Storage Cloud, right-click VSANs.
18. Select Create Storage VSAN.
19. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A.
20. Select Enabled for FC Zoning.
21. Select Fabric A.
22. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for Fabric A above.

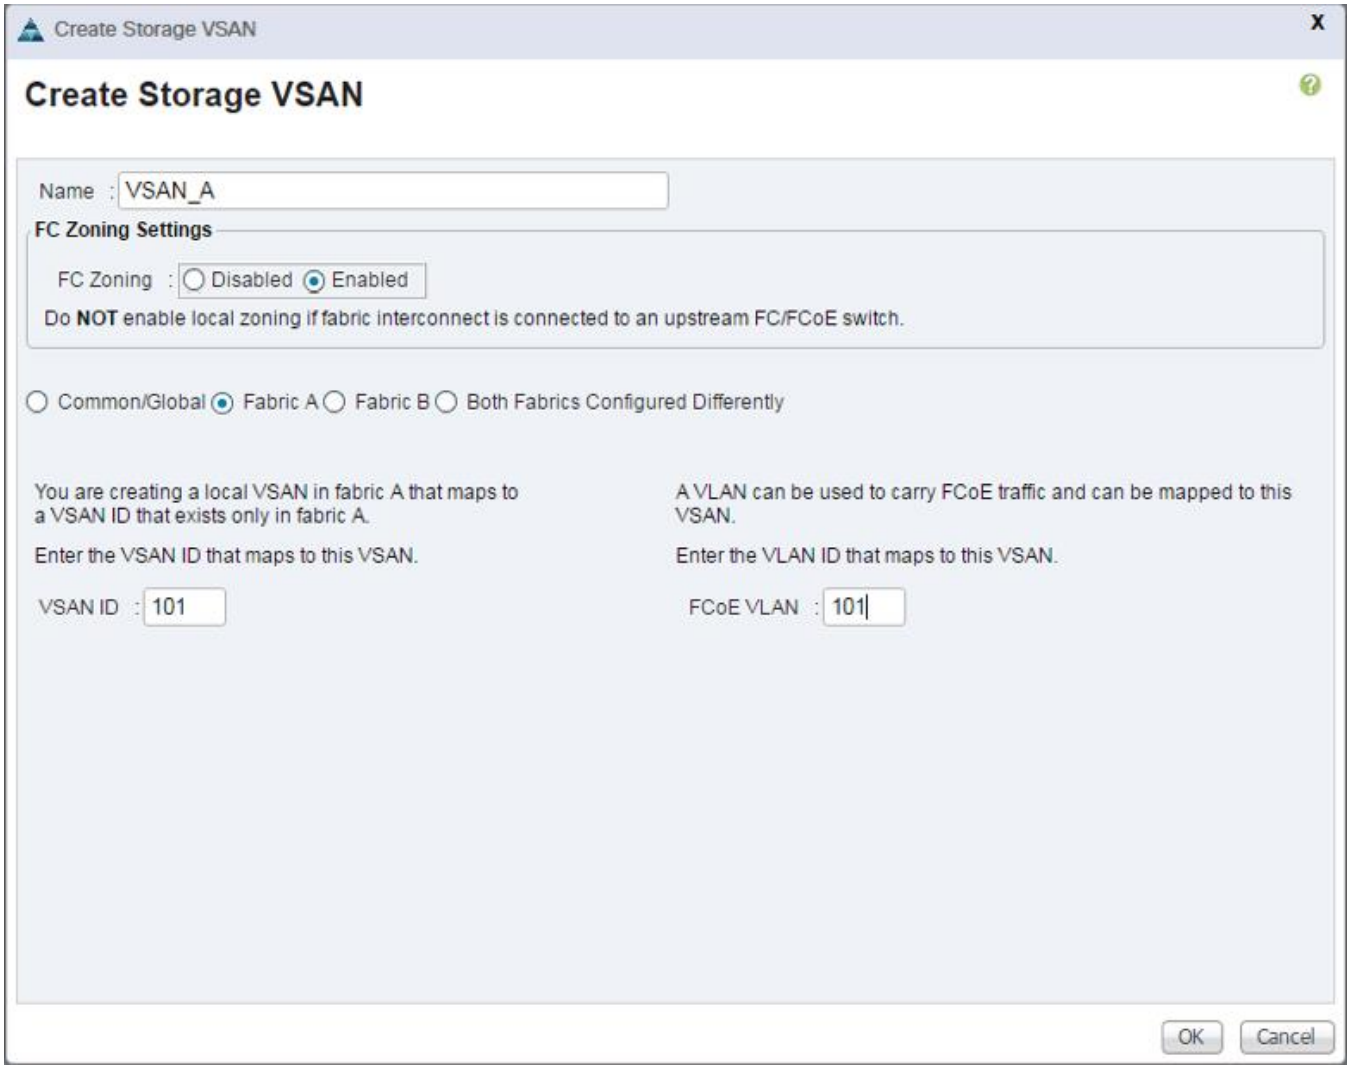

23. Click OK, and then click OK again.
24. Under Storage Cloud, right-click VSANs.
25. Select Create Storage VSAN.
26. Enter `VSAN_B` as the name of the VSAN to be used for Fabric B.
27. Select Enabled for FC Zoning.
28. Select Fabric B.
29. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for Fabric B above.

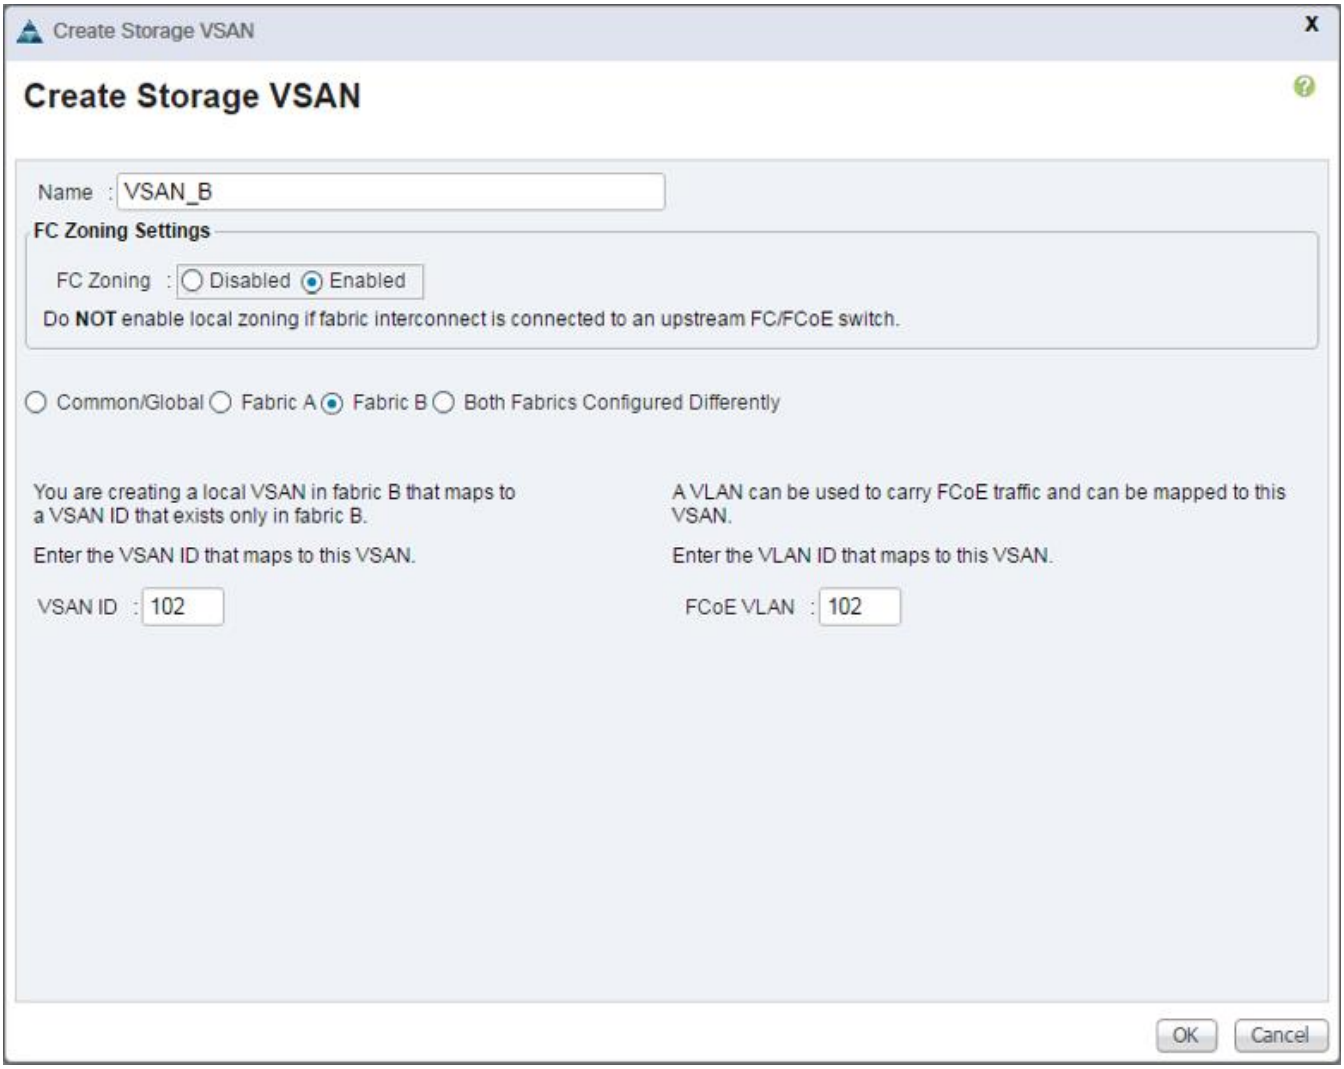

30. Click OK, and then click OK again.

Assign VSANs to FC Storage Ports

To assign the necessary virtual storage area networks (VSANs) to the FC Storage Ports for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > Storage Cloud.
3. Expand Fabric A and Storage FC Interfaces.
4. Right-click FC Interface 1/1 for the 6332-16UP or 2/15 in the 6248UP and select Storage FC Interface.
5. Set the User Label to the storage controller name and port that this interface is connected to.
6. Select `VSAN_A(101)` as the VSAN.

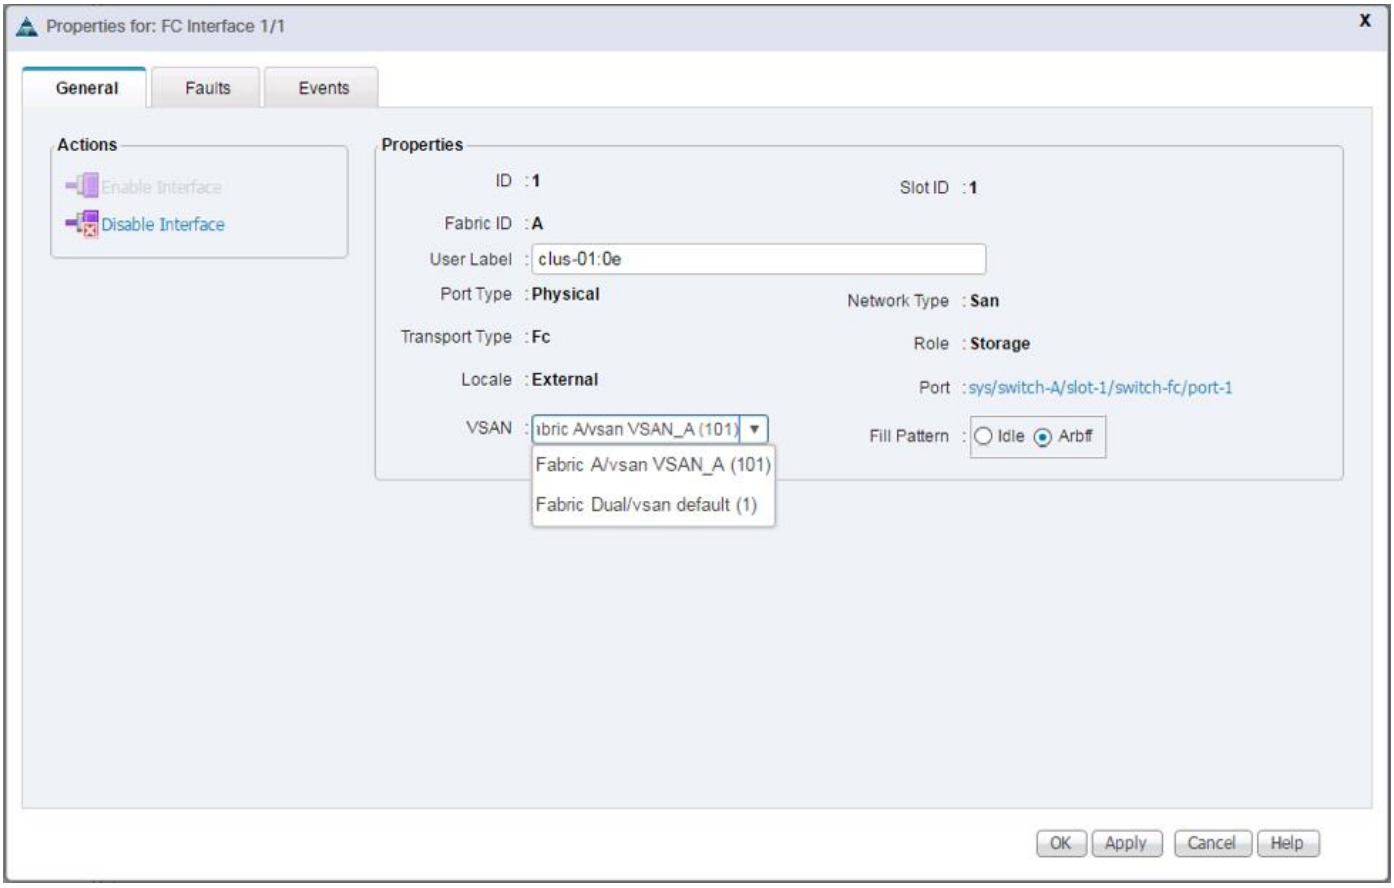

7. Click OK.
8. Expand Fabric A and Storage FC Interfaces.
9. Right-click FC Interface 1/2 for the 6332-16UP or 2/16 in the 6248UP and select Storage FC.
10. Set the User Label to the storage controller name and port that this interface is connected to.
11. Select `VSAN_A(101)` as the VSAN.

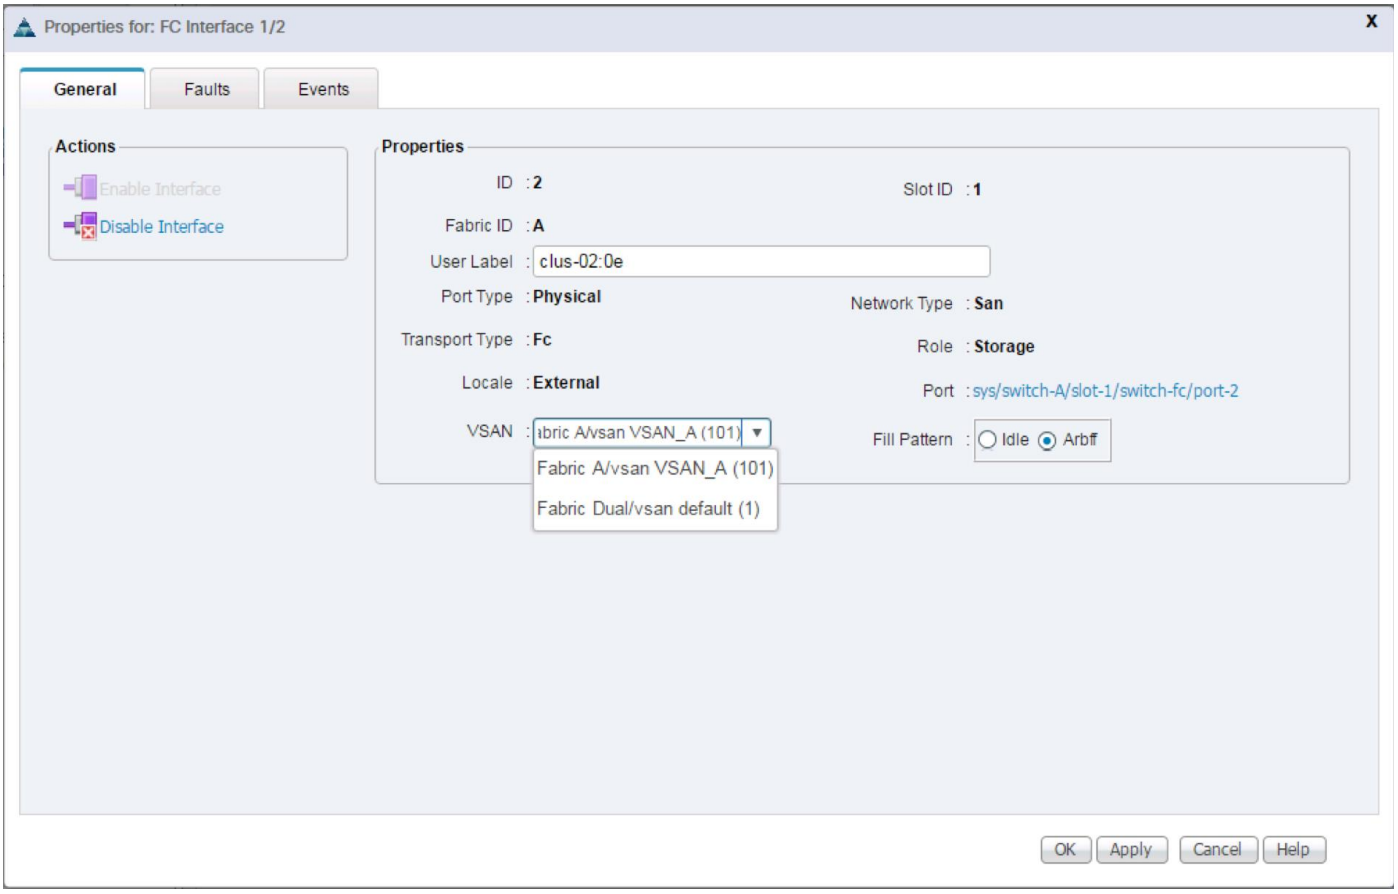

12. Click OK.
13. Expand Fabric B and Storage FC Interfaces.
14. Right-click FC Interface 1/1 for the 6332-16UP or 2/15 in the 6248UP and select Storage FC.
15. Set the User Label to the storage controller name and port that this interface is connected to.
16. Select `VSAN_B(102)` as the VSAN.

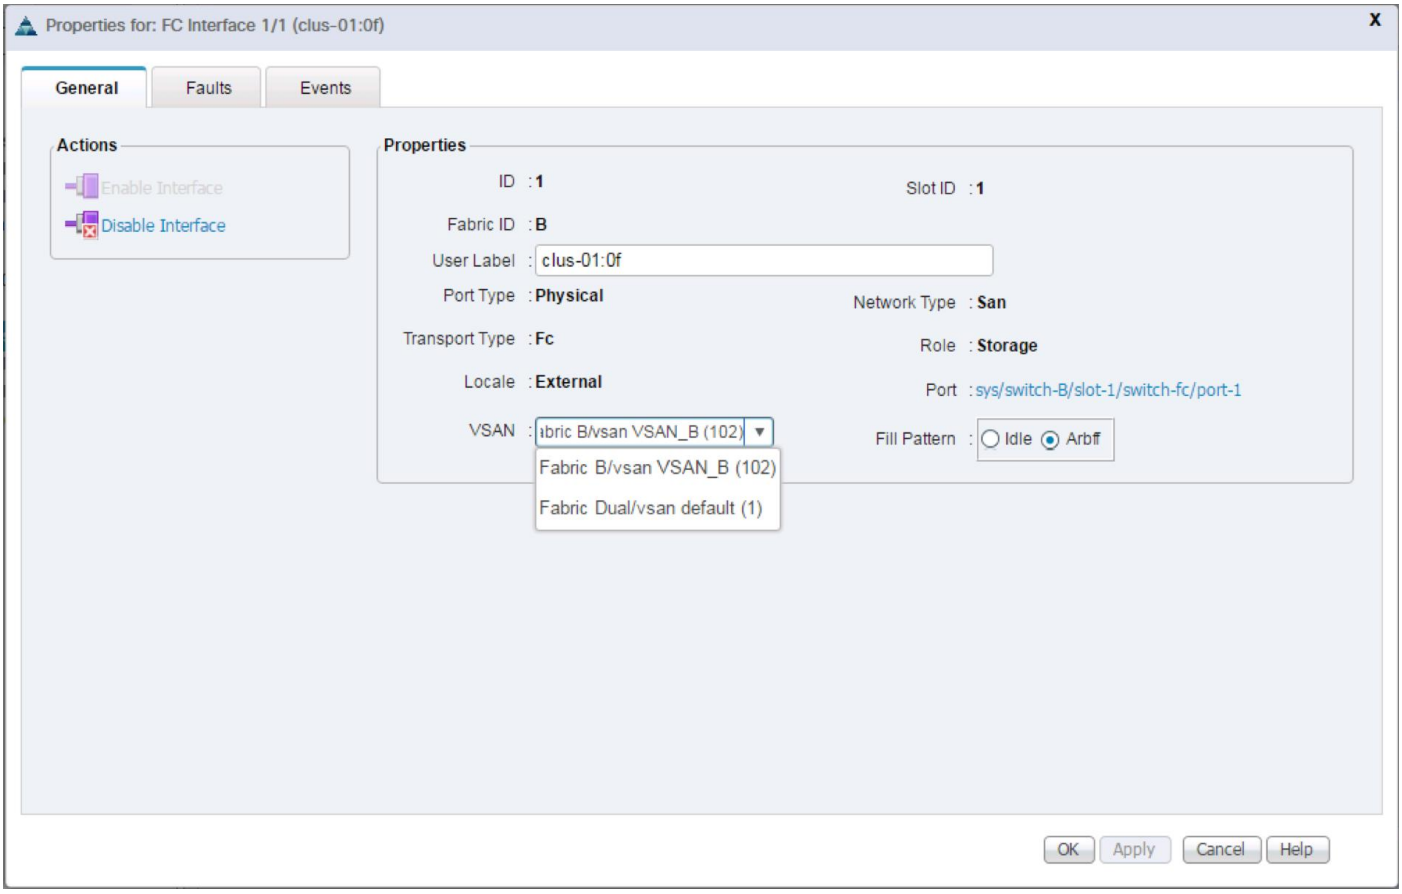

17. Click OK.
18. Expand Fabric B and Storage FC Interfaces.
19. Right-click FC Interface 1/2 for the 6332-16UP or 2/16 in the 6248UP and select Storage FC.
20. Set the User Label to the storage controller name and port that this interface is connected to.
21. Select `VSAN_B(102)` as the VSAN.

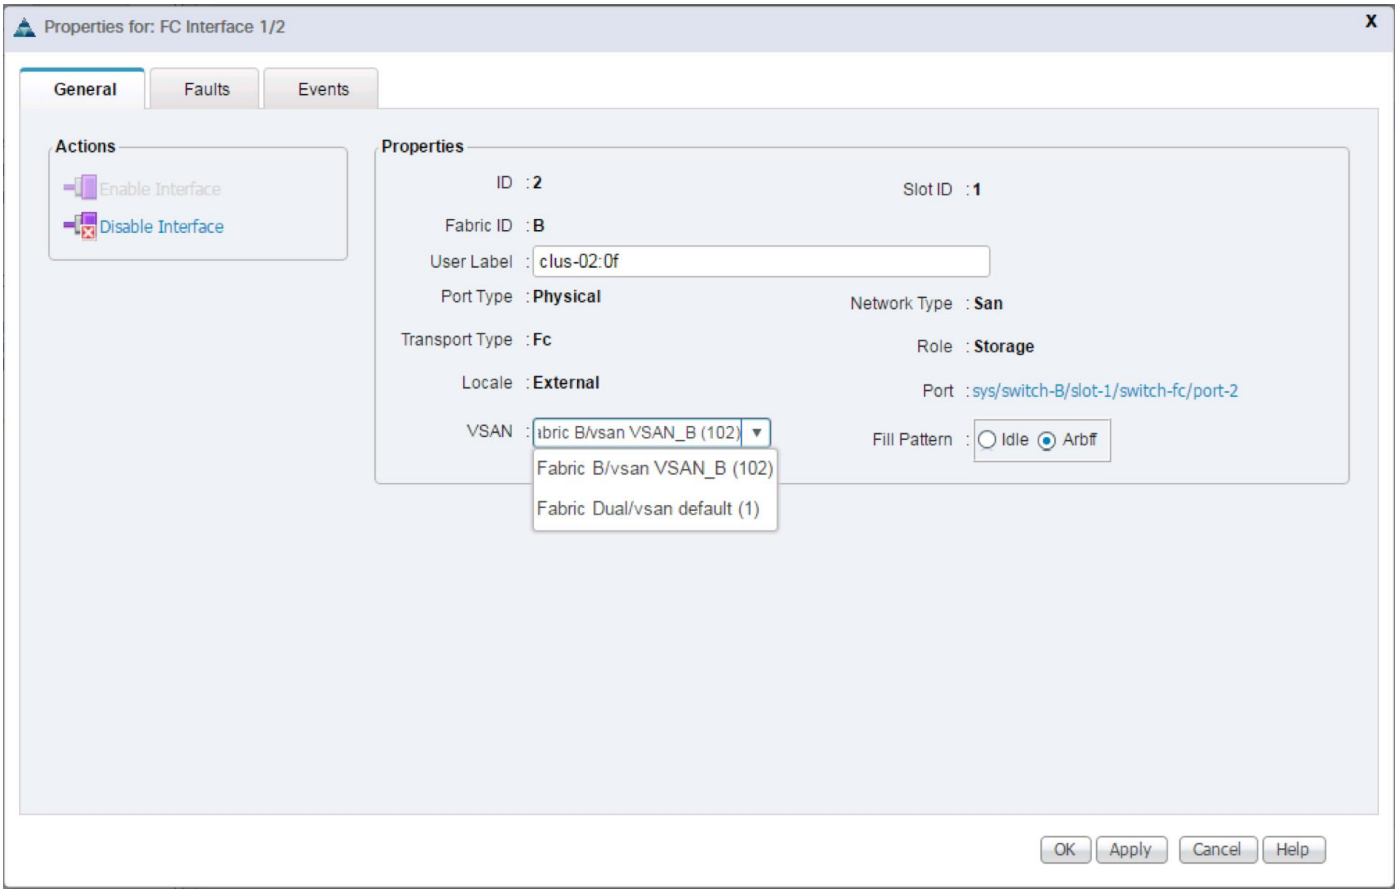

22. Click OK.

Create Storage Connection Policies for FC Zoning

To create Storage Connection Policies for the FC Zoning, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Storage Connection Policies.
4. Select Create Storage Connection Policy.
5. Enter `Infra-Fabric-A` as the name of the policy.
6. Select the Single Initiator Multiple Targets Zoning Type.
7. Click the Plus Sign on the right to add a zoning target.
8. Enter the WWPN for [fcp_lif01_63a or fcp_lif01_62a] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.
9. Select Path A and VSAN_A.

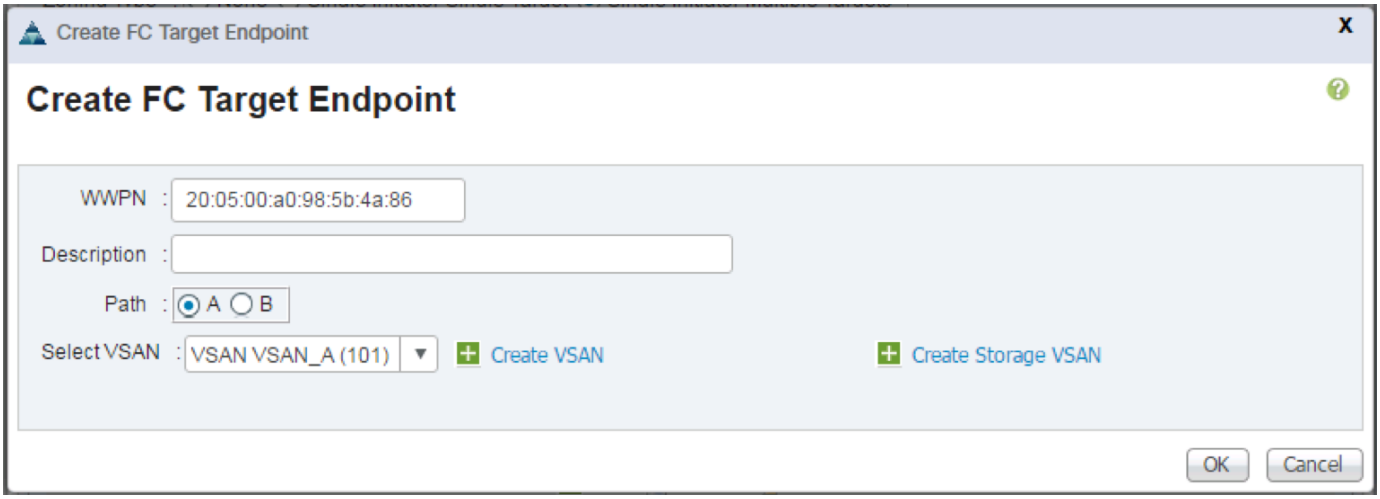

10. Click OK.

11. Click the Plus Sign on the right to add a second zoning target.

12. Enter the WWPN for [fcp_lif02_63a or fcp_lif02_62a] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.

13. Select Path A and VSAN_A.

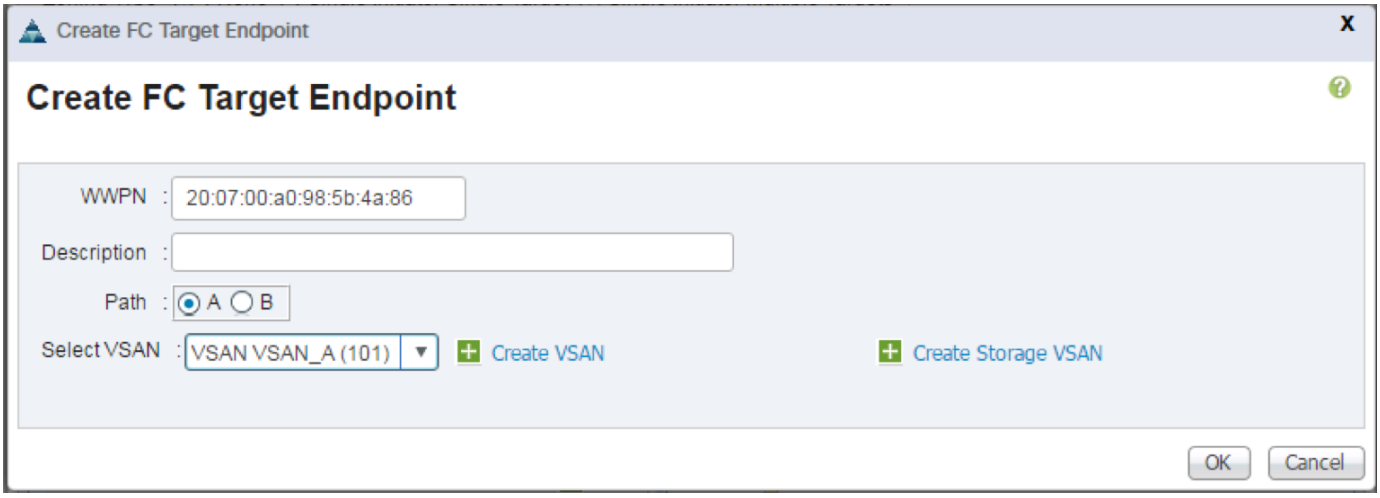

14. Click OK.

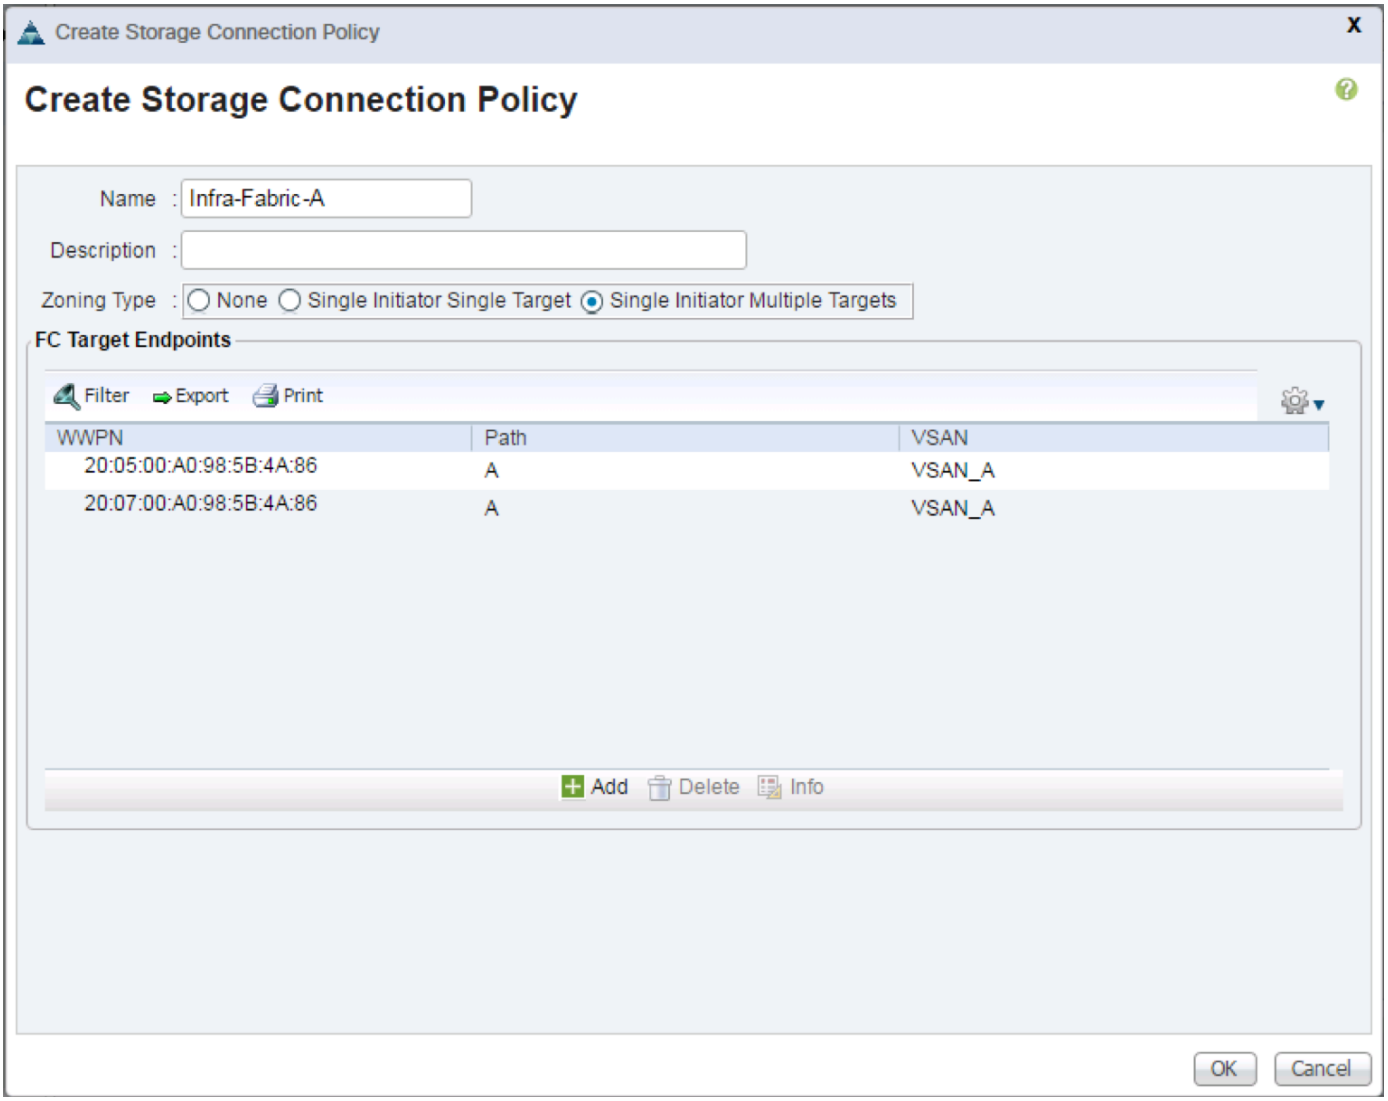

15. Click OK, and then click OK again.
16. Right-click Storage Connection Policies.
17. Select Create Storage Connection Policy
18. Enter `Infra-Fabric-B` as the name of the policy.
19. Select the Single Initiator Multiple Targets Zoning Type.
20. Click the Plus Sign on the right to add a zoning target.
21. Enter the WWPN for [fcp_lif01_63b or fcp_lif01_62b] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.
22. Select Path B and VSAN_B.

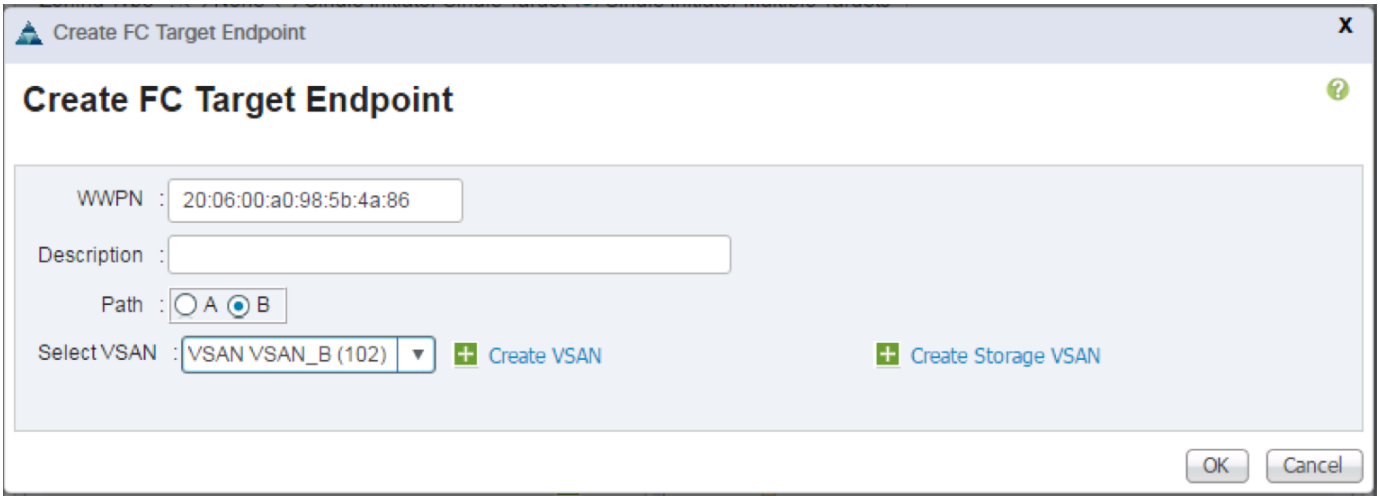

23. Click OK.

24. Click the Plus Sign on the right to add a second zoning target.

25. Enter the WWPN for [fcp_lif02_63b or fcp_lif02_62b] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.

26. Select Path B and VSAN_B.

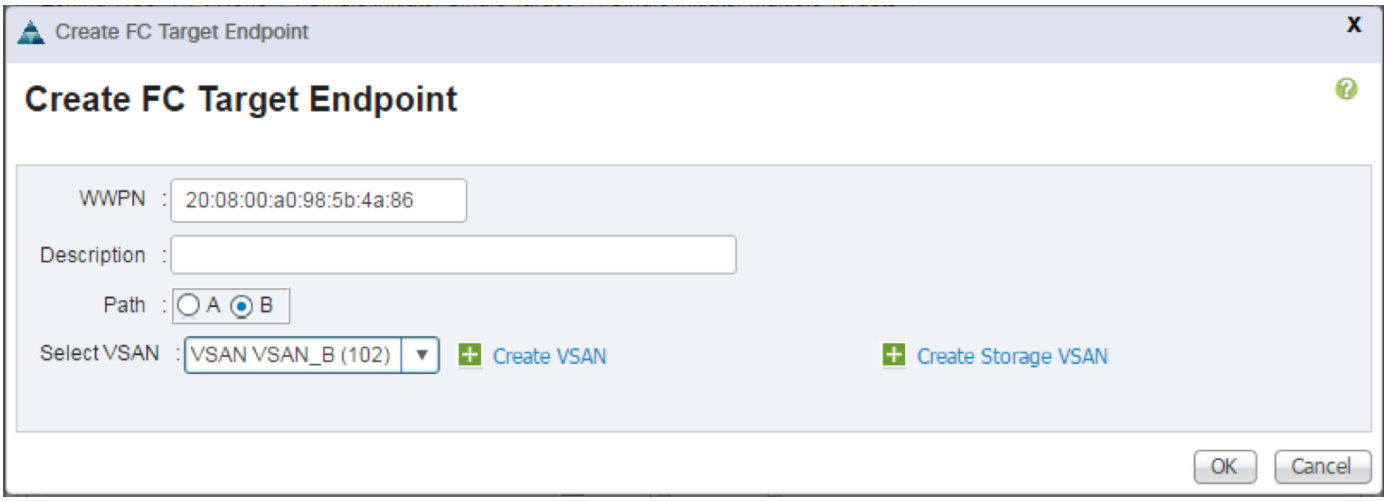

27. Click OK.

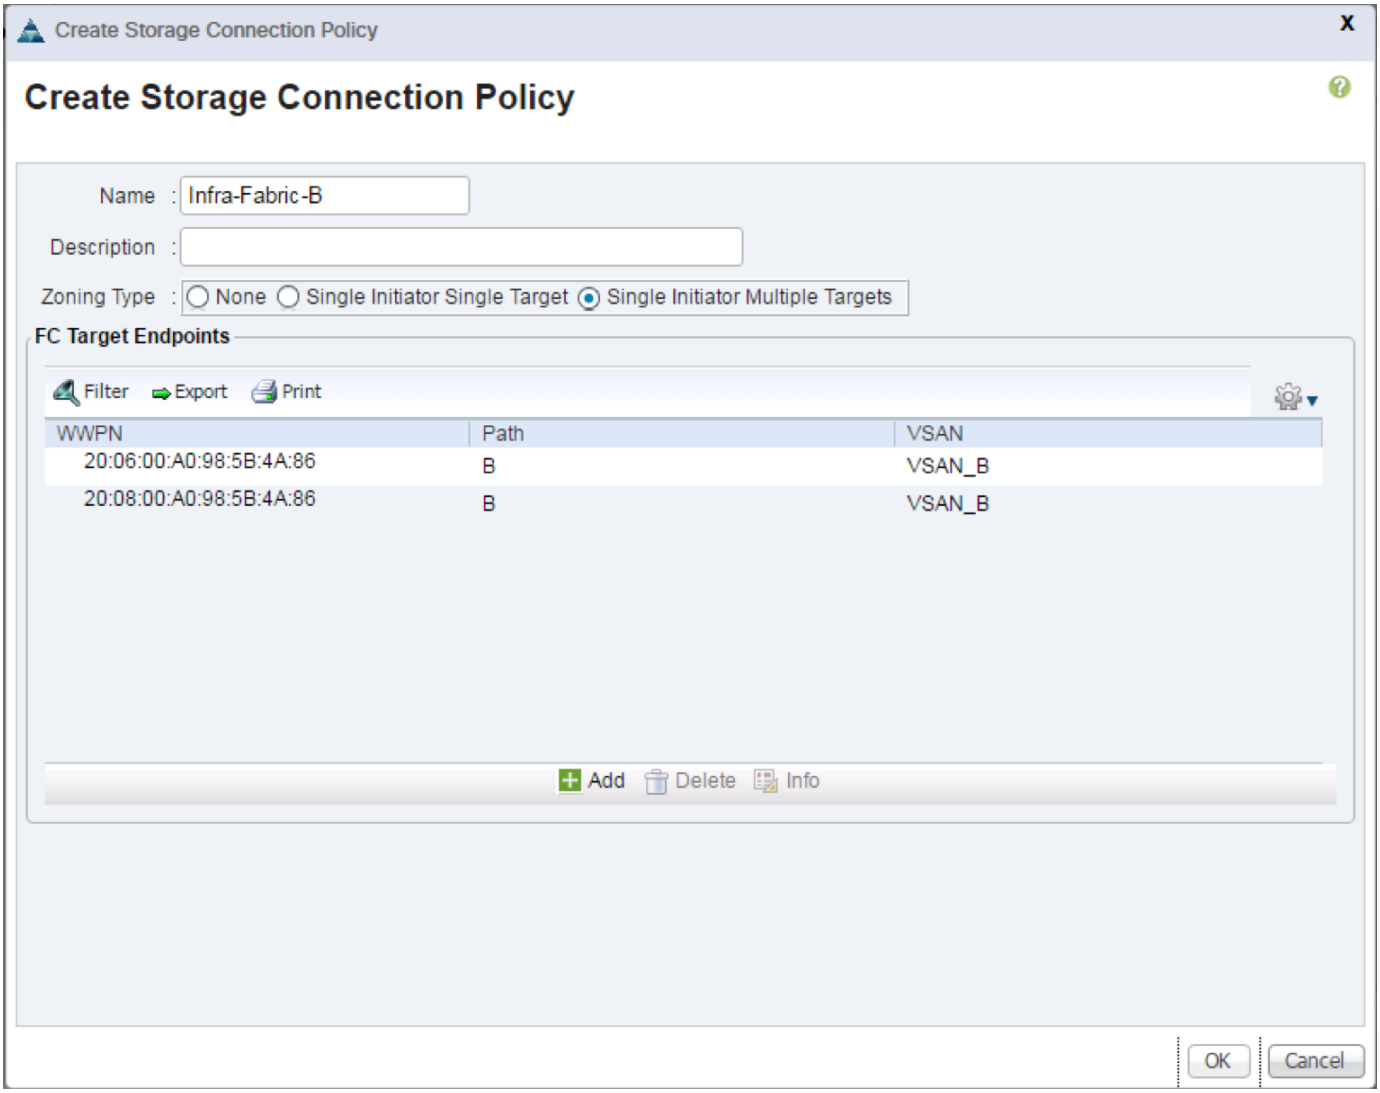

28. Click OK, and then click OK again.

Create vHBA Templates

To create the necessary 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.
4. Select Create vHBA Template.
5. Enter `vHBA_Template_A` as the vHBA template name.
6. Keep Fabric A selected.

7. Select VSAN_A.
8. Leave Initial Template as the Template Type.
9. Select `WWPN_Pool_A` as the WWPN Pool.
10. Click OK to create the vHBA template.
11. Click OK.

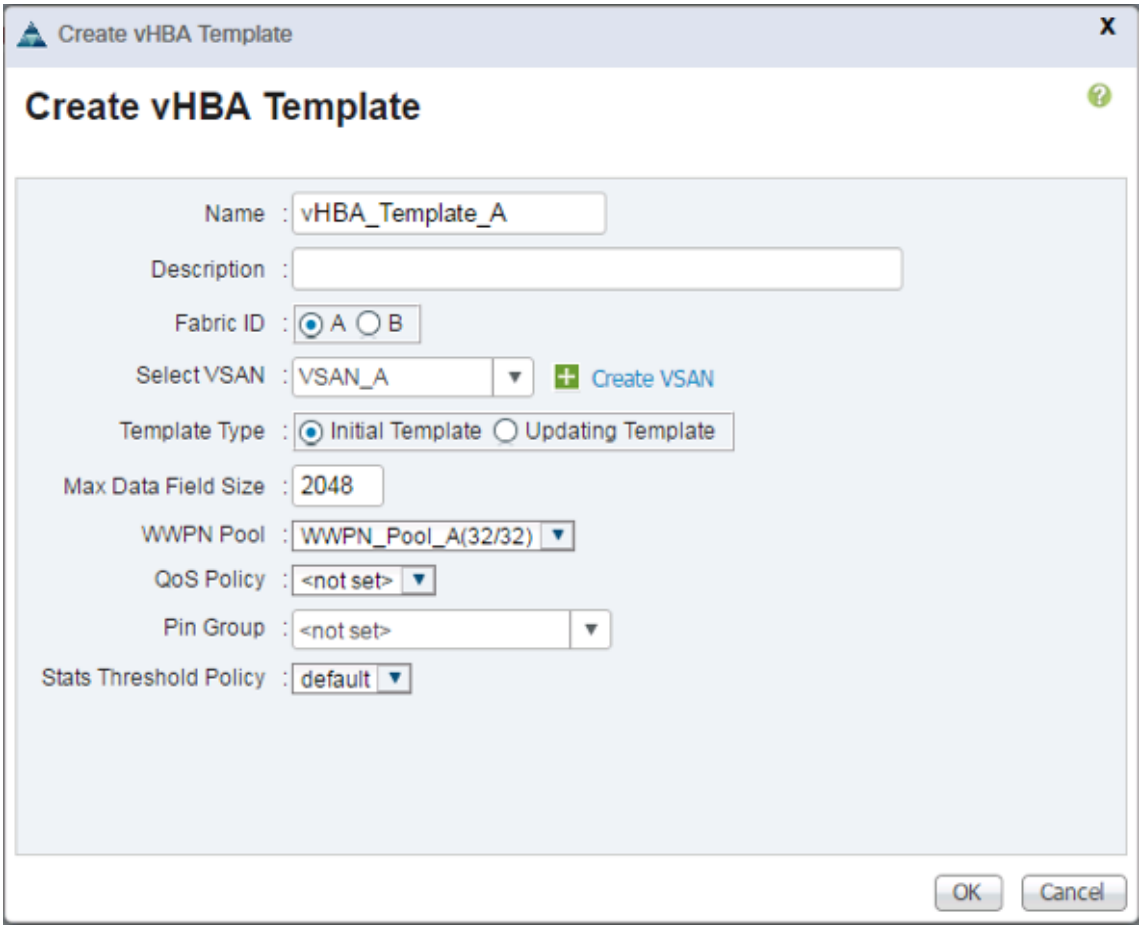

12. Right-click vHBA Templates.
13. Select Create vHBA Template.
14. Enter `vHBA_Template_B` as the vHBA template name.
15. Select Fabric B as the Fabric ID.
16. Select VSAN_B.
17. Leave Initial Template as the Template Type.
18. Select `WWPN_Pool_B` as the WWPN Pool.

19. Click OK to create the vHBA template.

20. Click OK.

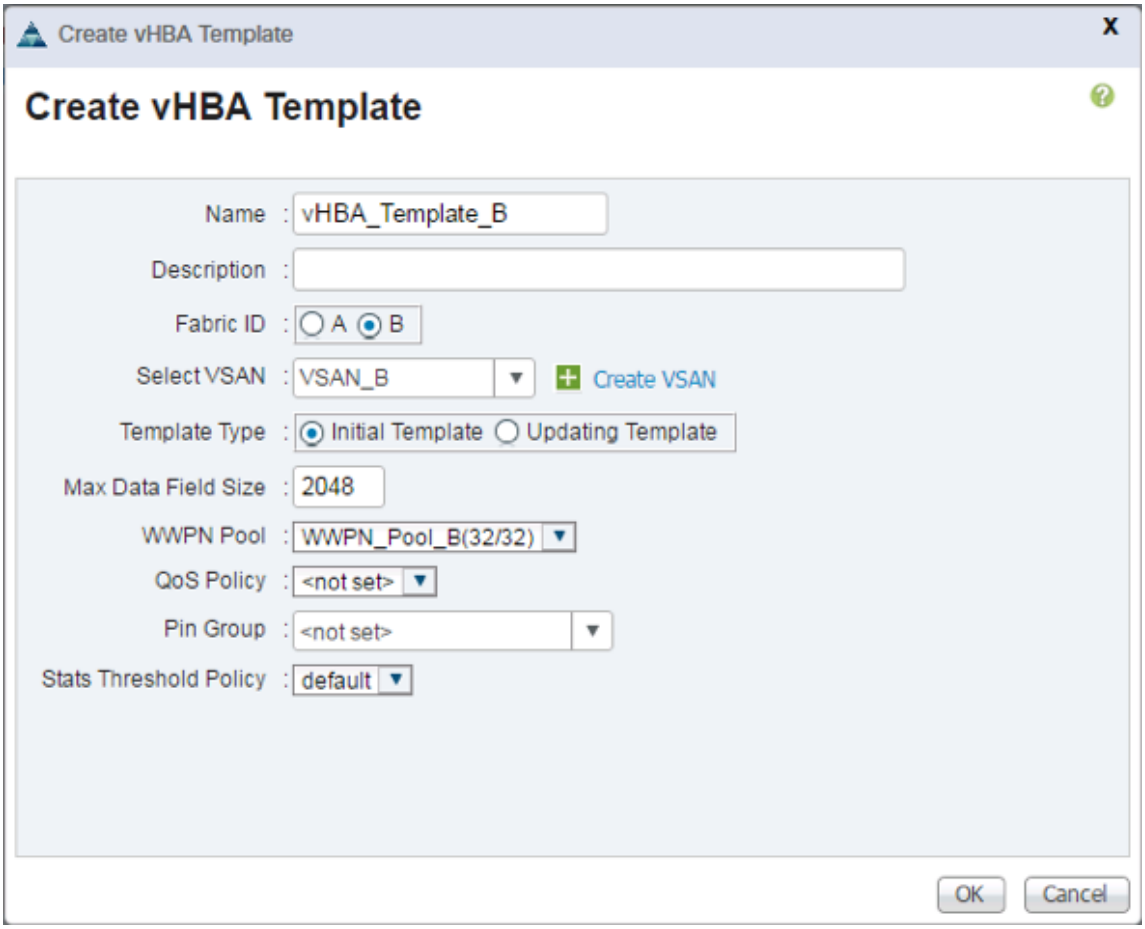

Create SAN Connectivity Policy

During testing, iSCSI vNICs and FC vHBAs were deployed to all hosts regardless of boot policy used. The FC vHBA interfaces can be left out for environments configured to utilize iSCSI for boot and data.

To configure the necessary Infrastructure SAN Connectivity Policy, 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.
4. Select Create SAN Connectivity Policy.
5. Enter Infra-SAN-Policy as the name of the policy.
6. Select the previously created WWNN_Pool for the WWNN Assignment.

7. Click the Add button at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter Fabric-A as the name of the vHBA.
9. Select the Use vHBA Template checkbox.
10. In the vHBA Template list, select vHBA_Template_A.

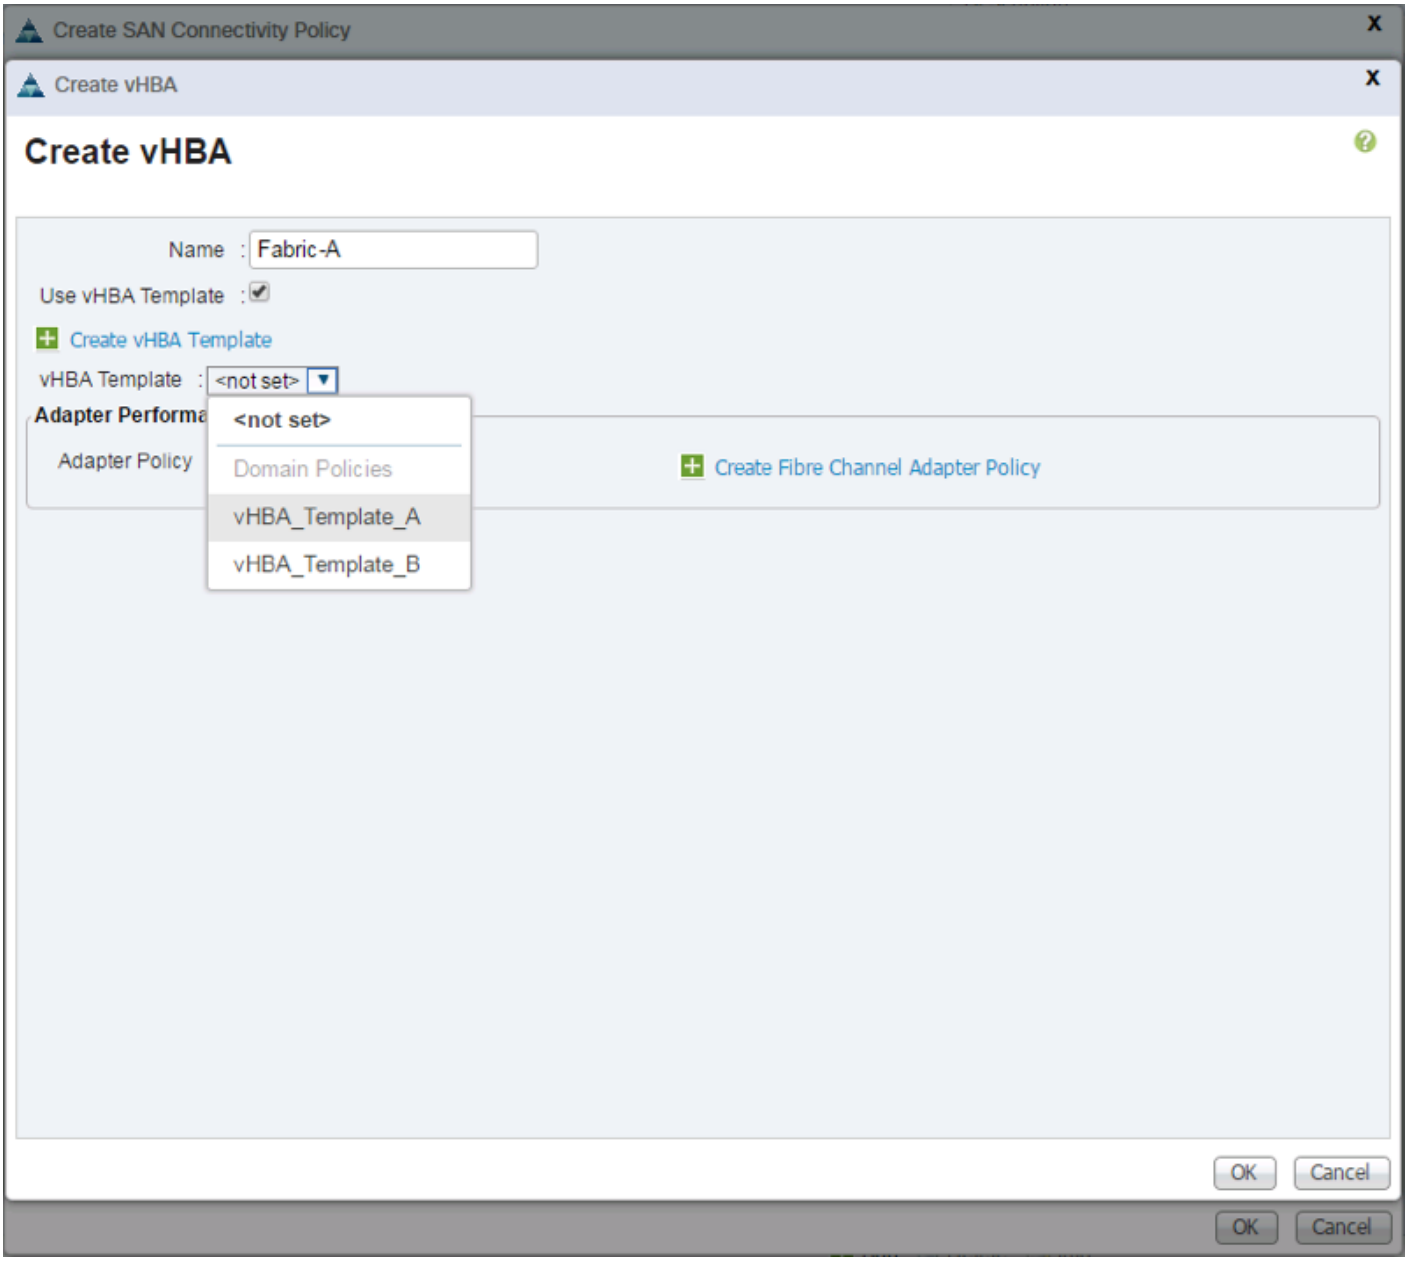

11. In the Adapter Policy list, select VMWare.
12. Click OK.
13. Click the Add button at the bottom to add a second vHBA.

14. In the Create vHBA dialog box, enter Fabric-B as the name of the vHBA.
15. Select the Use vHBA Template checkbox.
16. In the vHBA Template list, select vHBA_Template_B.

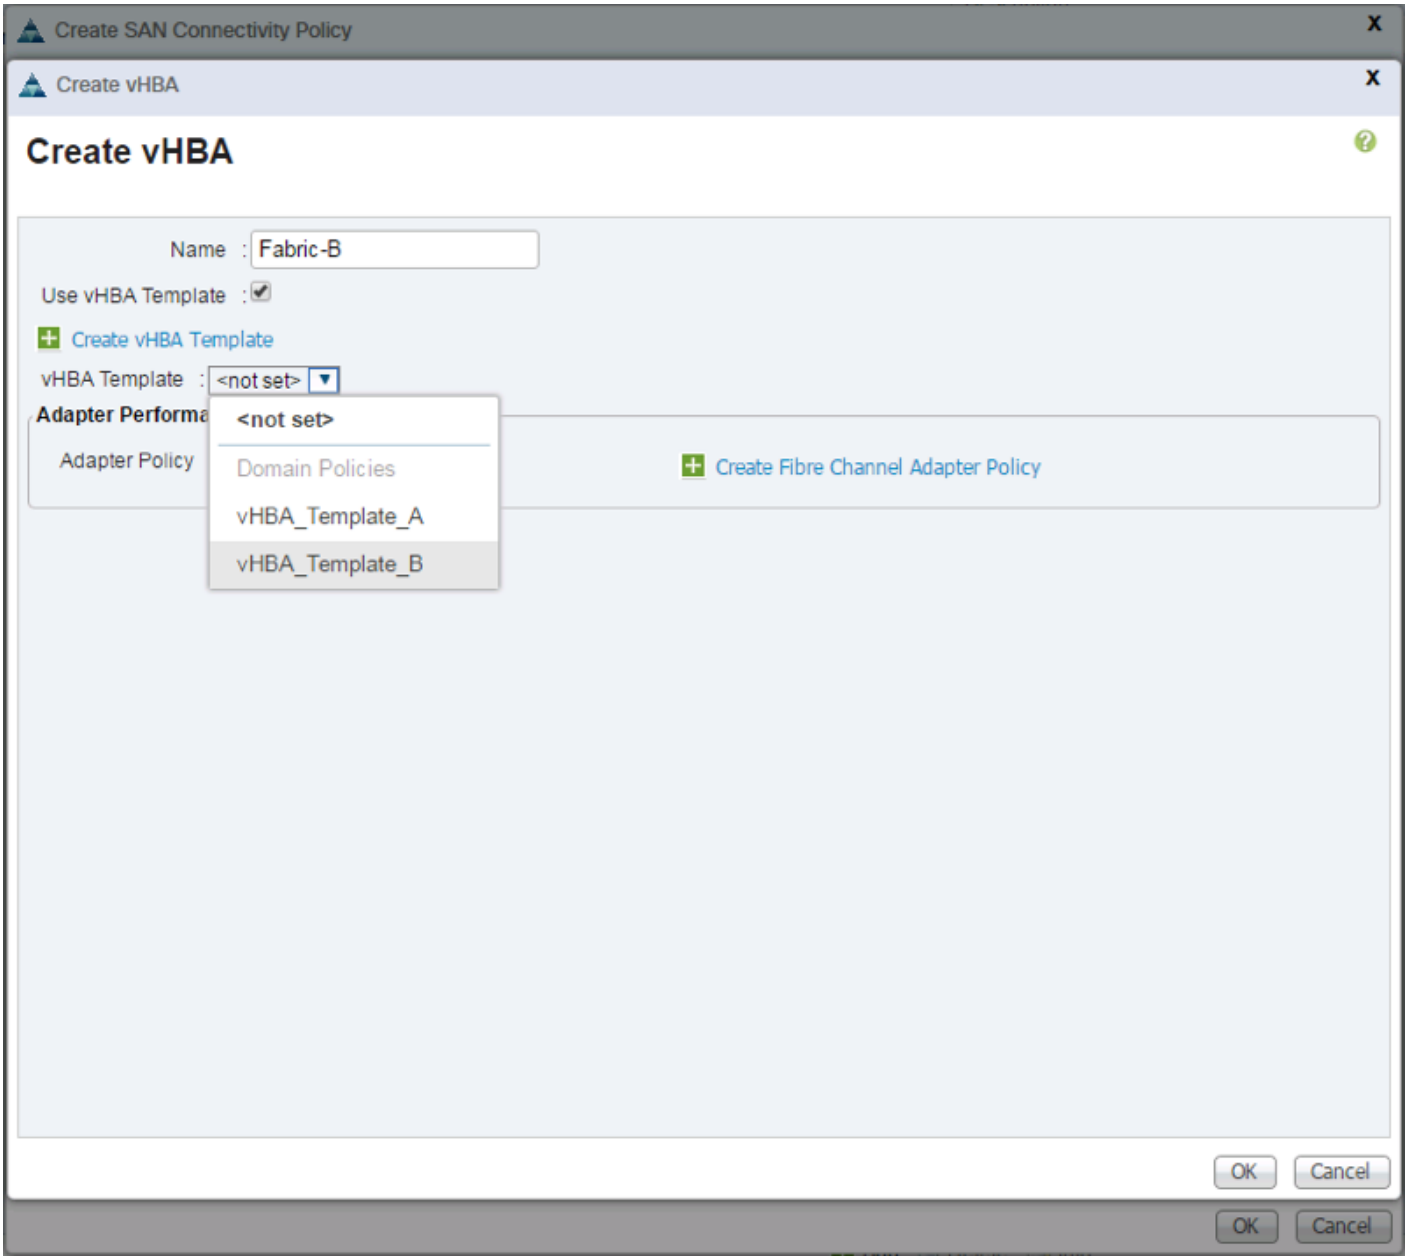

17. In the Adapter Policy list, select VMWare.
18. Click OK.

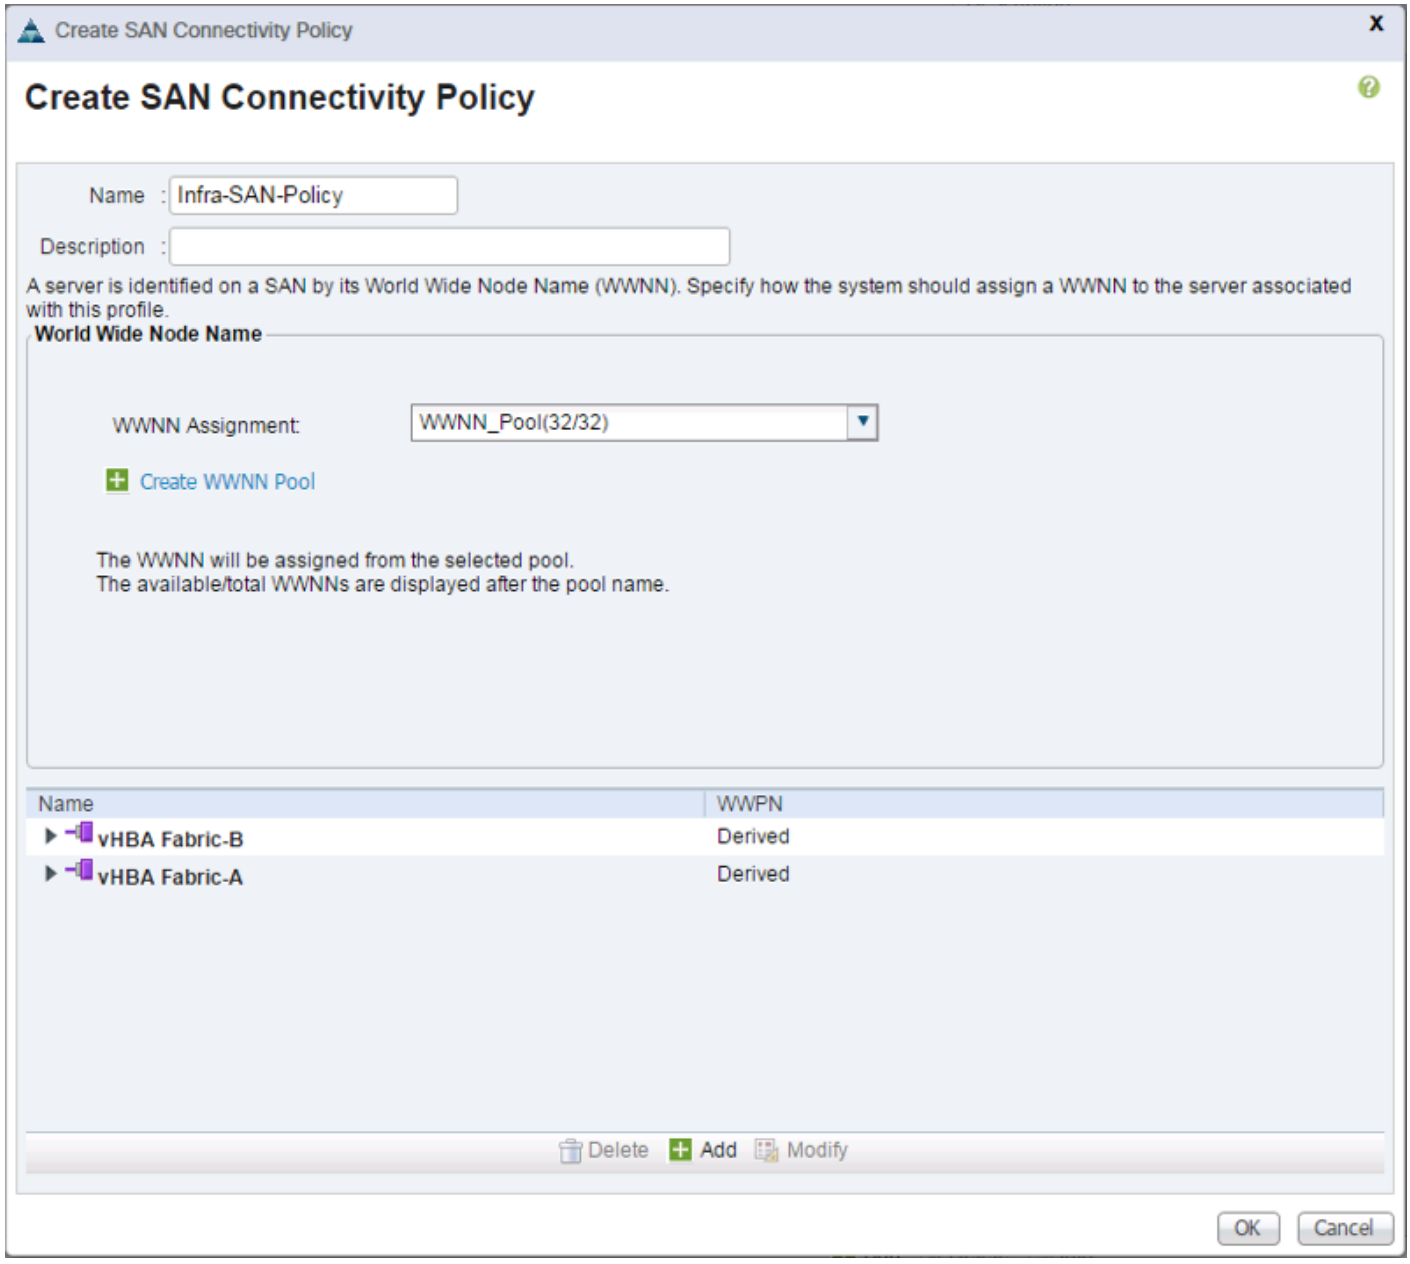

19. Click OK to create the SAN Connectivity Policy.

20. Click OK to confirm creation.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, 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 Sequential as the option for Assignment Order.

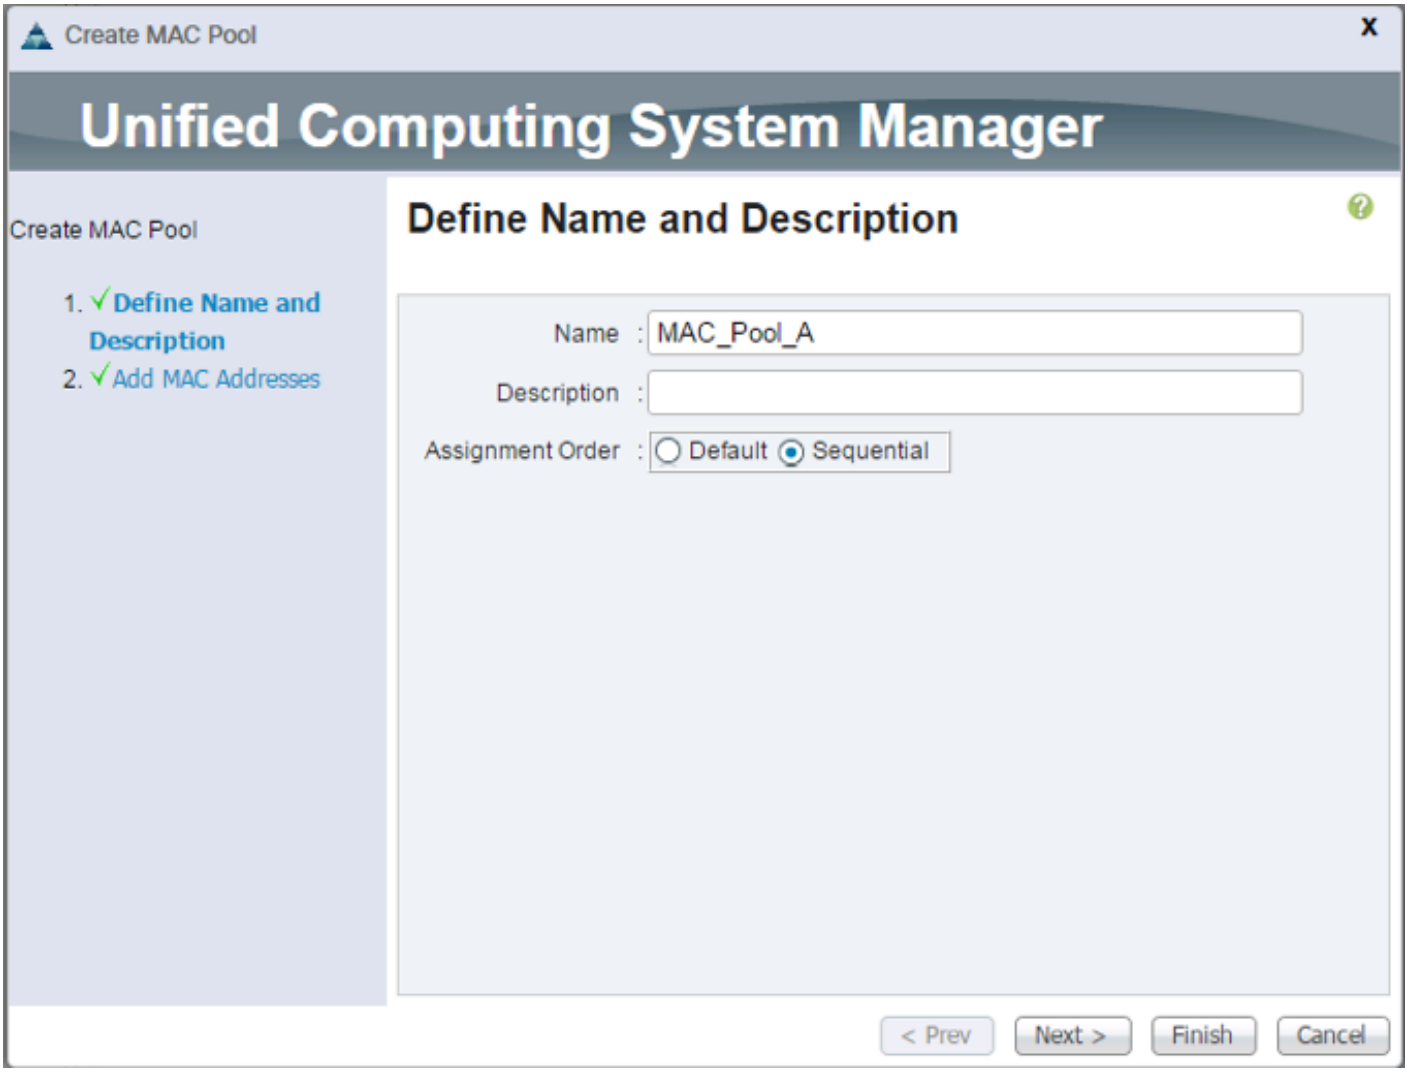

8. Click Next.
9. Click Add.
10. Specify a starting MAC address.

For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the of also embedding the extra building, floor and UCS domain number information giving us `00:25:B5:91:1A:00` as our first MAC address.

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

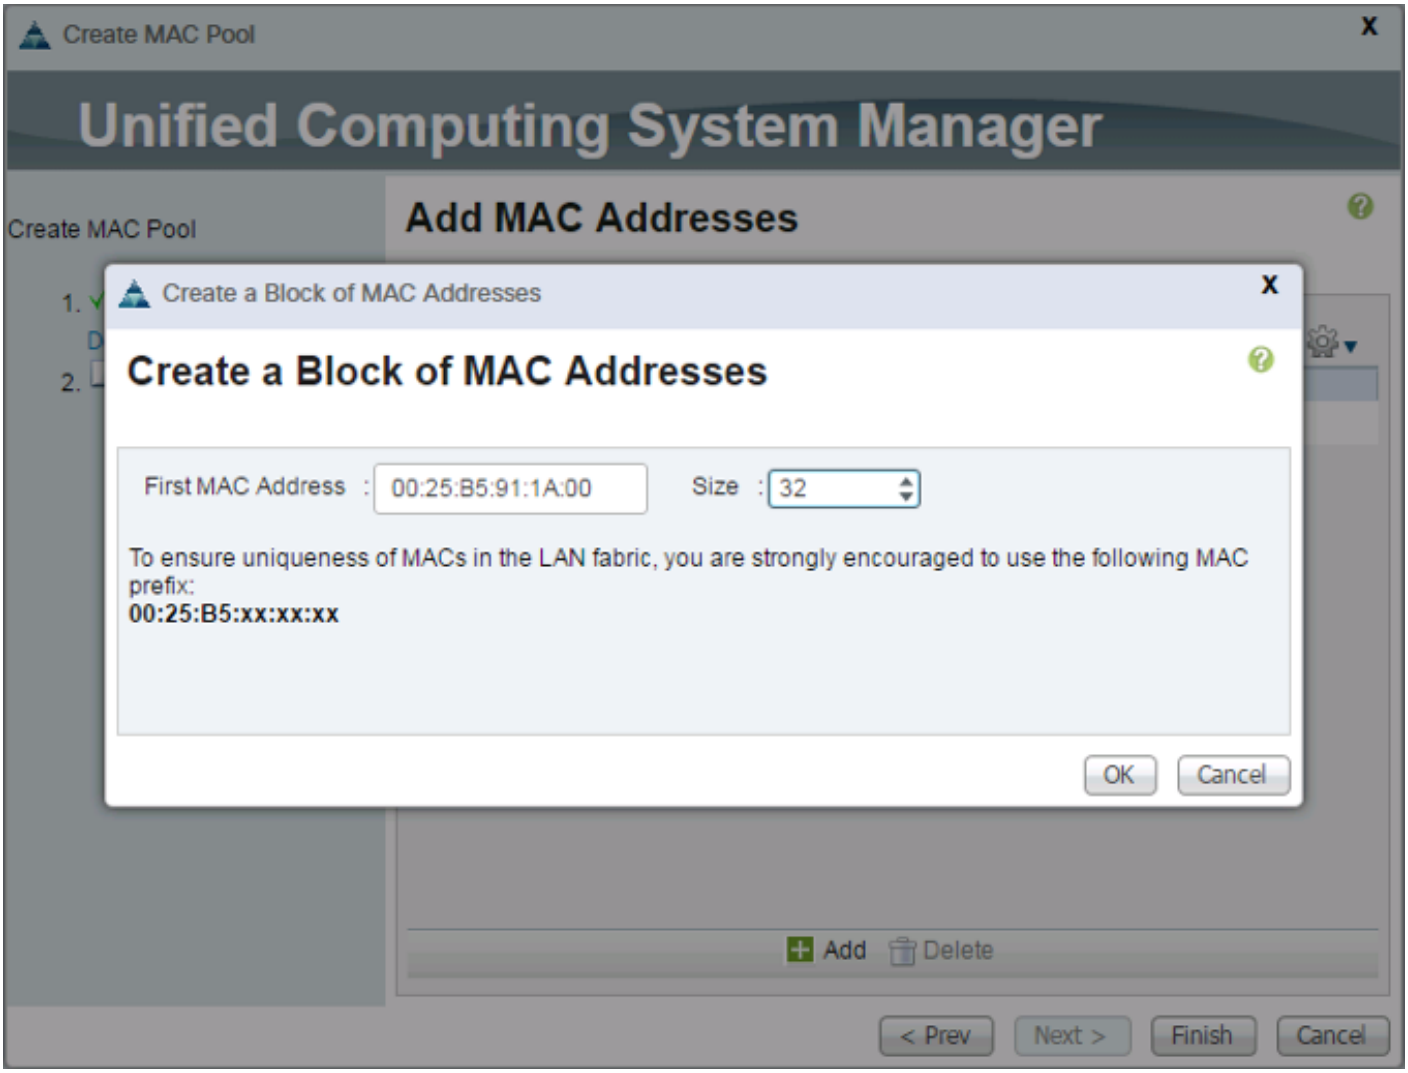

12. Click OK.

13. Click Finish.

14. In the confirmation message, click OK.

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

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

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

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

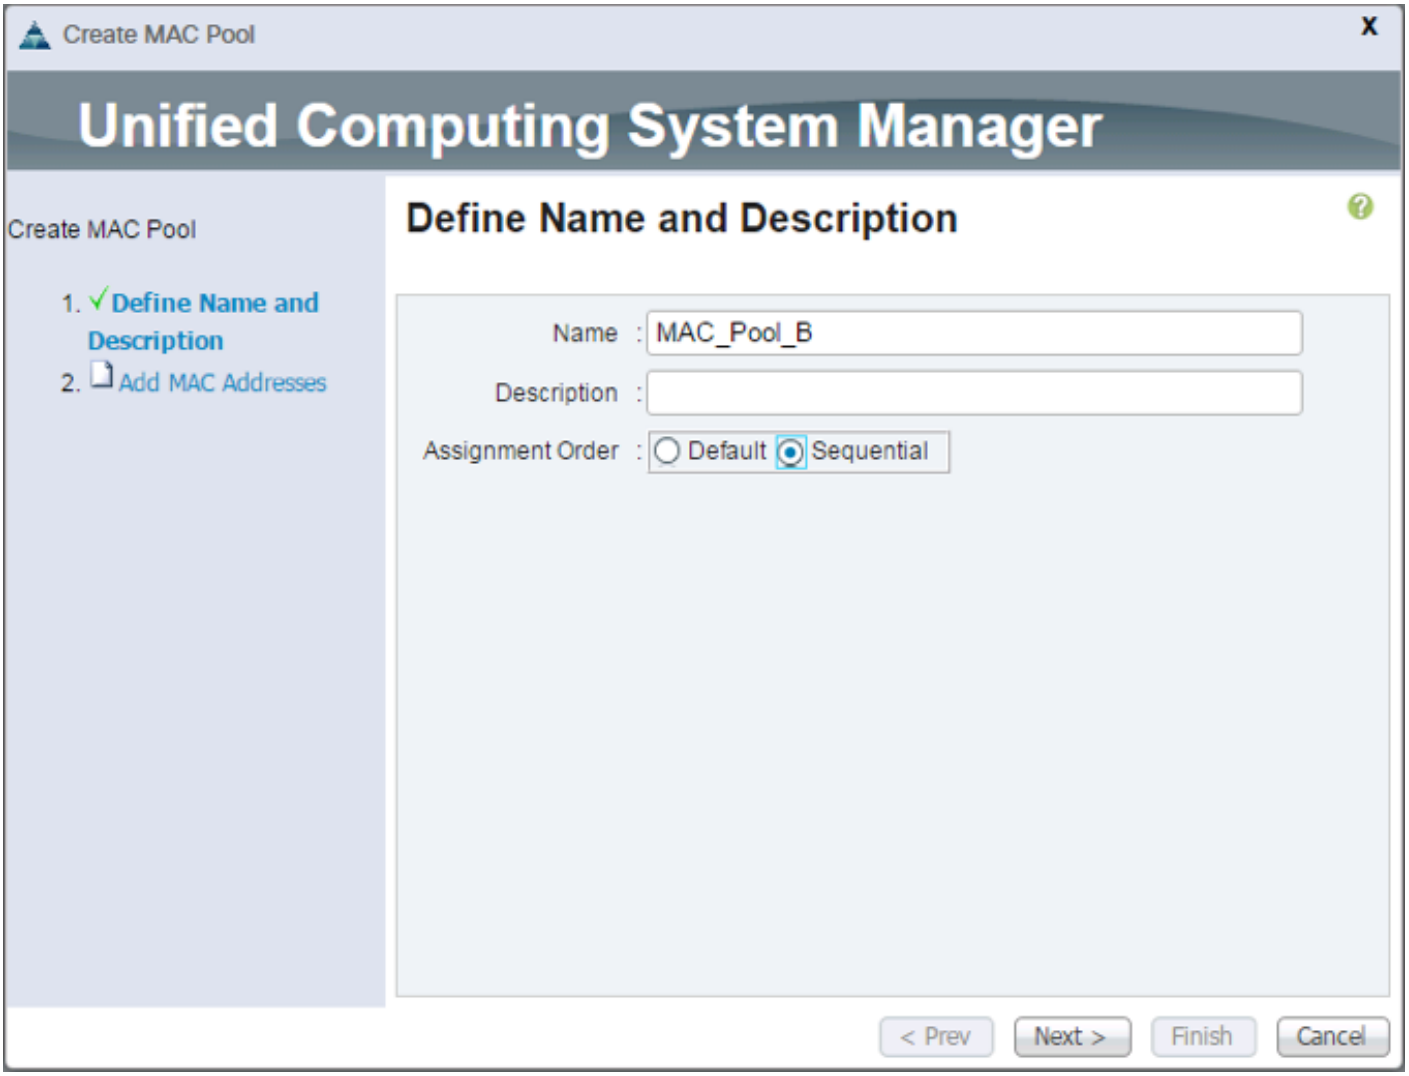

19. Click Next.

20. Click Add.

21. Specify a starting MAC address.

For the FlexPod solution, it is recommended to place `0B` in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the extra building, floor and UCS domain number information giving us `00:25:B5:91:1B:00` as our first MAC address.

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

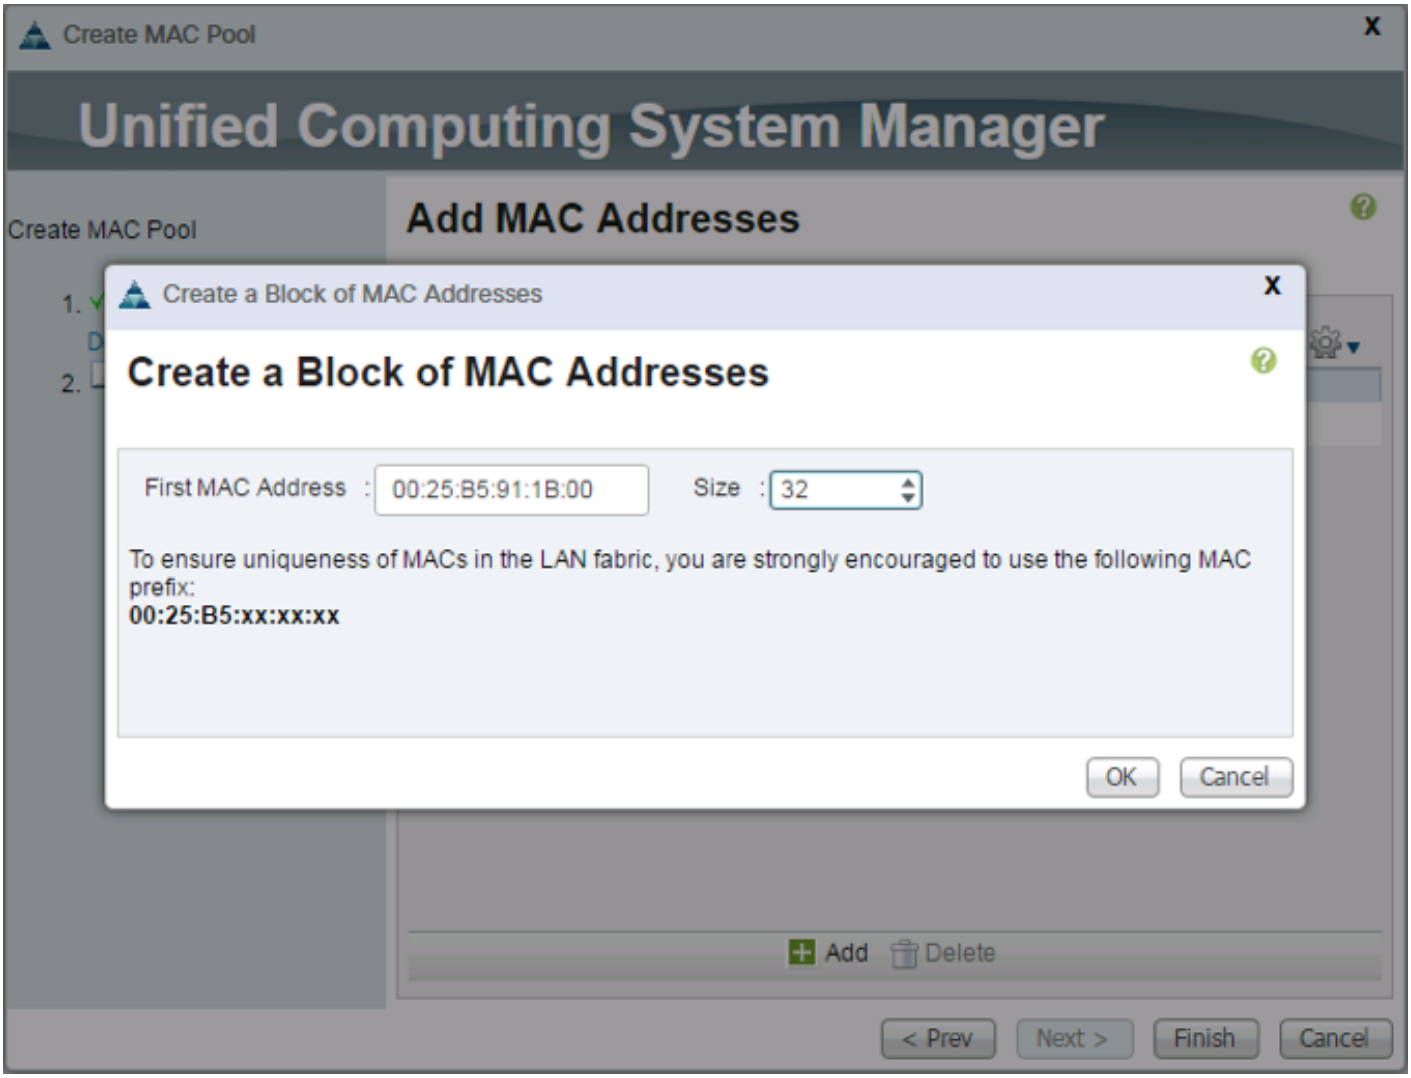

23. Click OK.
24. Click Finish.
25. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps.

1. In the UCS Manager, select the SAN tab on the left.
2. Select Pools > root.
3. Right-click IQN Pools under the root organization.

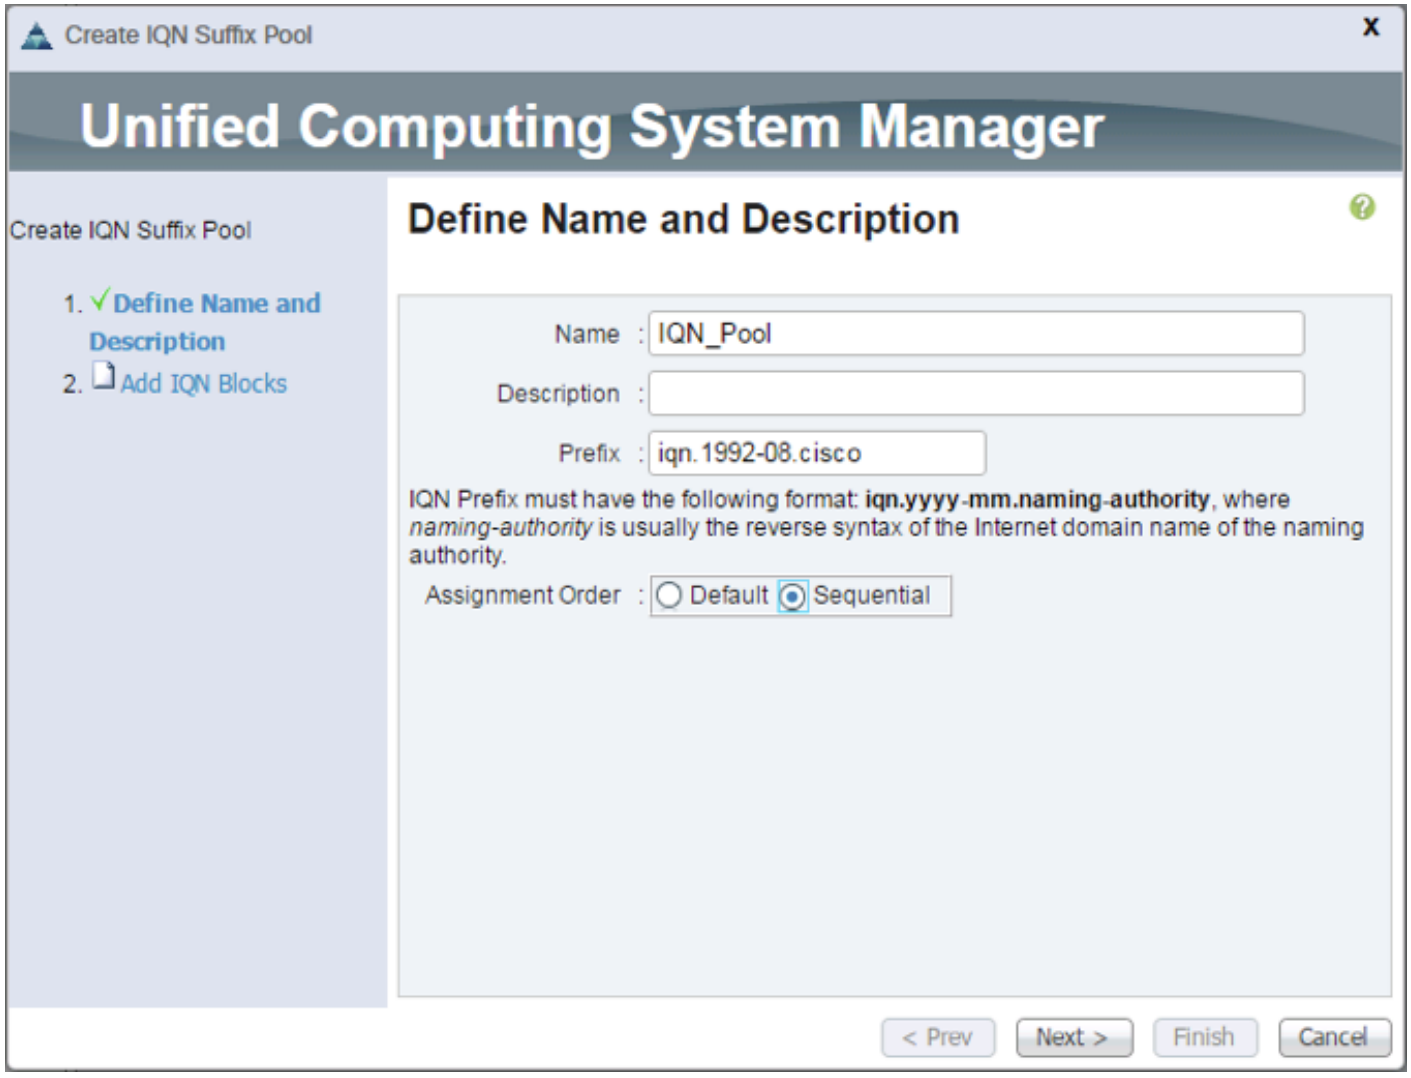

4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter `IQN_Pool` for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter `iqn.1992-08.com.cisco` as the prefix
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter `ucs-host` as the suffix.
12. Enter `1` in the From field.
13. Specify a size of the IQN block sufficient to support the available server resources.

14. Click OK.

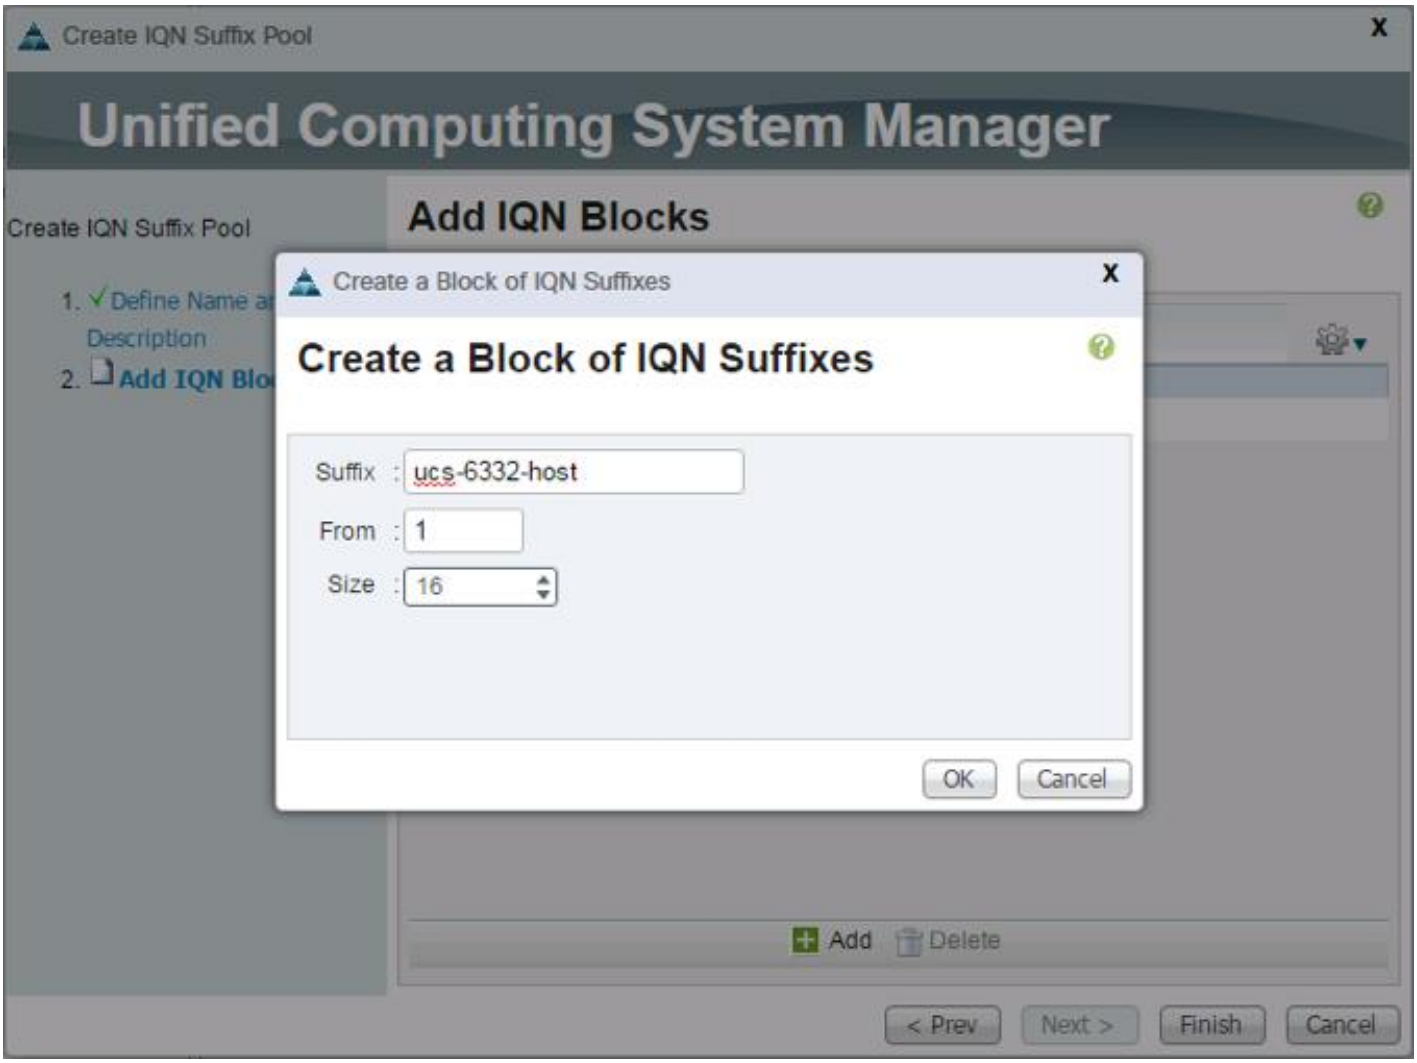

15. Click Finish.

16. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot

These steps provide details for configuring the necessary IP pools iSCSI boot for the Cisco UCS environment.

1. In Cisco UCS Manager, select the LAN tab on the left.

2. Select Pools > root.

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

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

4. Select Create IP Pool to create the IP pool.
5. Enter `iSCSI_IP_Pool_A` for the name of the IP pool.
6. Optional: Enter a description of the IP pool.
7. Select Sequential for Assignment Order.

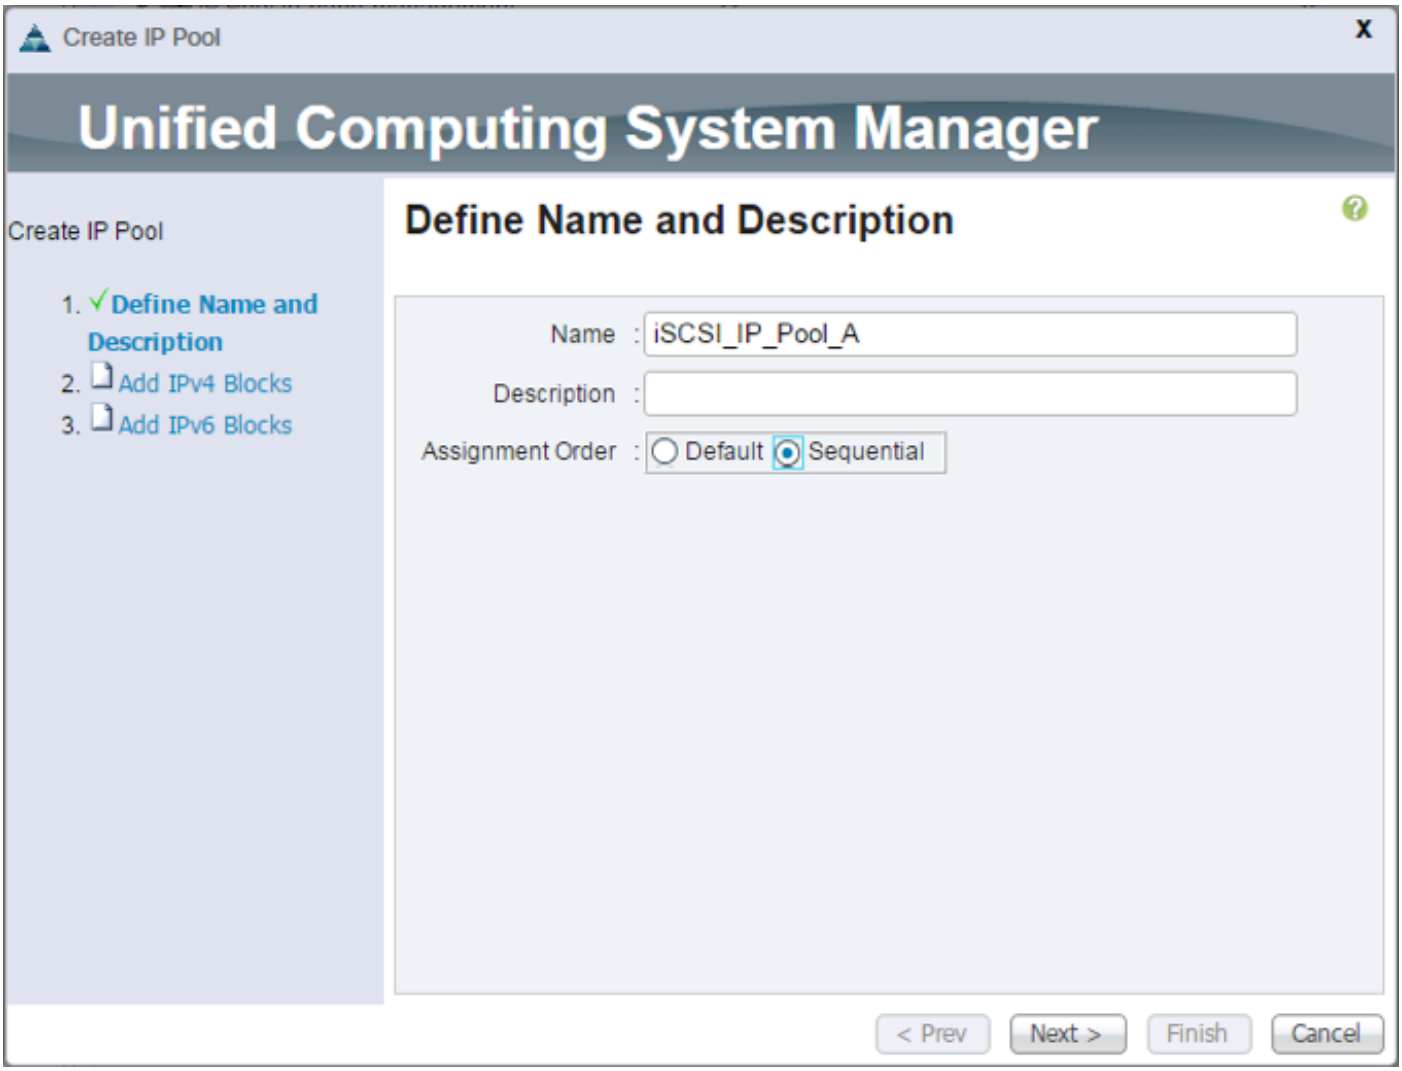

8. Click Next.
9. Click Add.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Set the size to enough addresses to accommodate the servers.

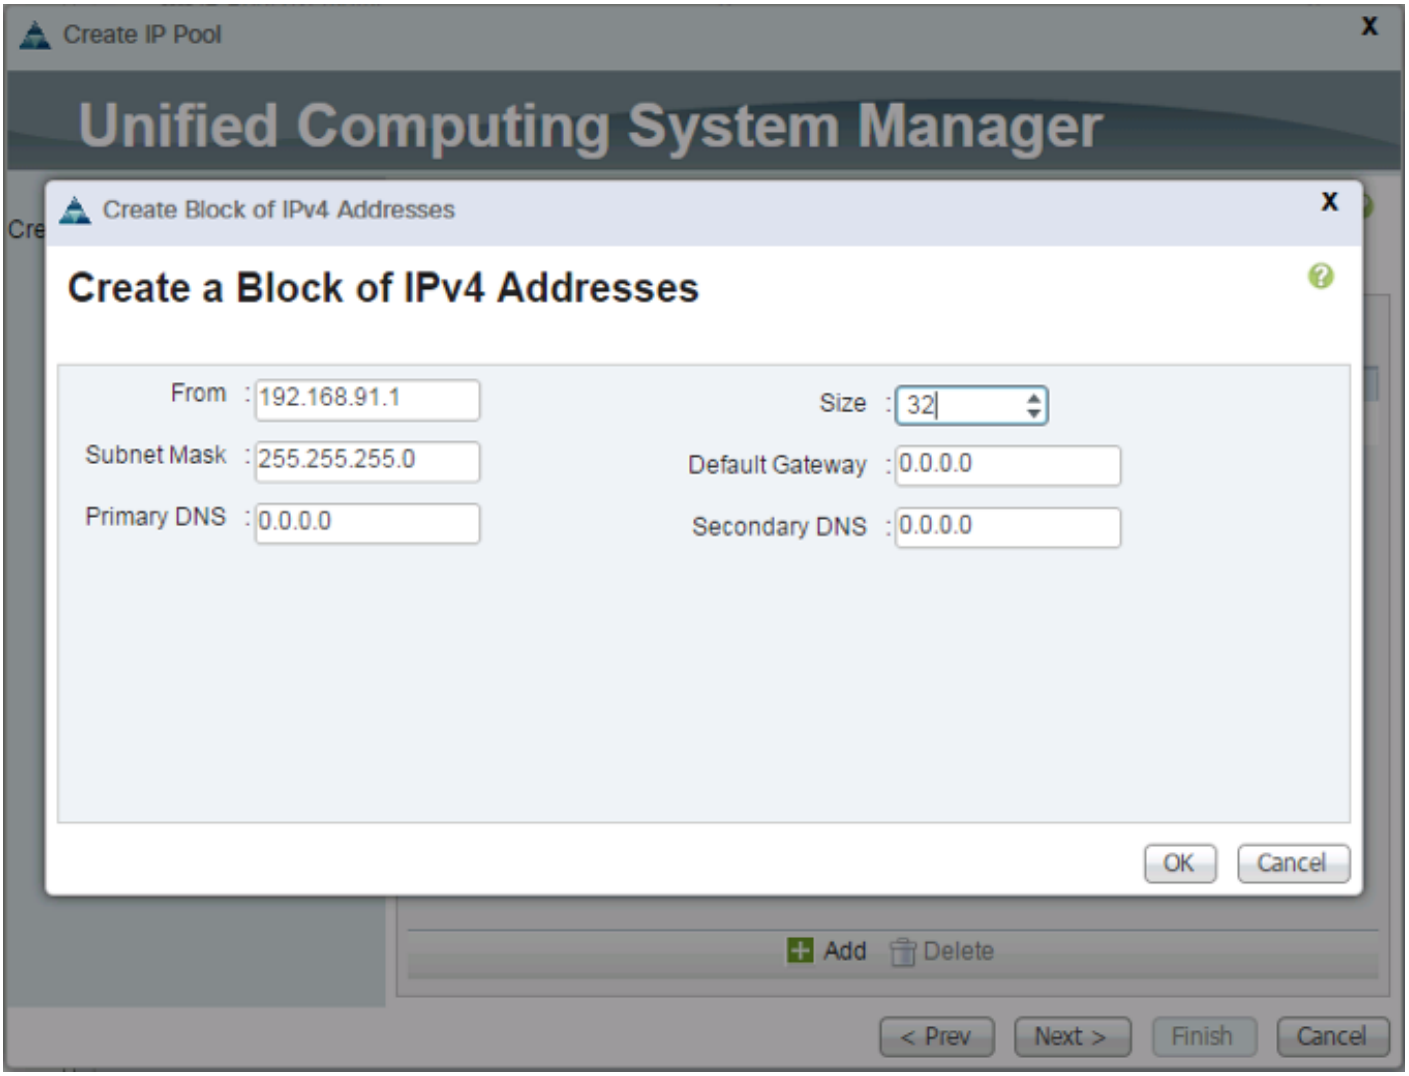

12. Click OK.
13. Click Finish.
14. Repeat these steps for an iSCSI_IP_Pool_B:
15. Right-click IP Pools under the root organization.
16. Select Create IP Pool to create the IP pool.
17. Enter `iSCSI_IP_Pool_B` for the name of the IP pool.
18. Optional: Enter a description of the IP pool.
19. Select Sequential for Assignment Order.
20. Click Next.
21. Click Add.

22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
23. Set the size to enough addresses to accommodate the servers.
24. Click OK.
25. Click Finish.

Create UUID Suffix Pool

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

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

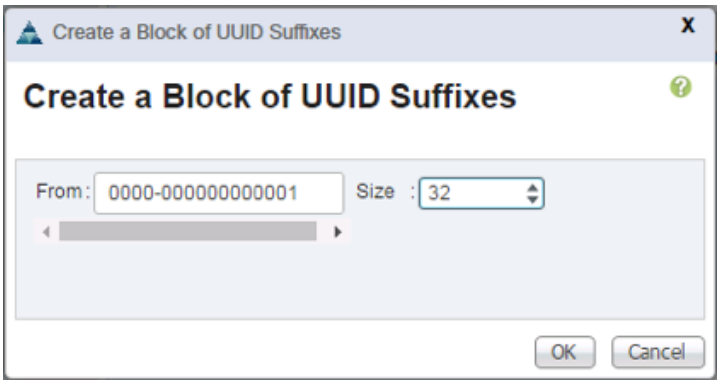

13. Click OK.
14. Click Finish.

15. Click OK.

Create Server Pool

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `Infra_Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the `Infra_Pool` server pool.

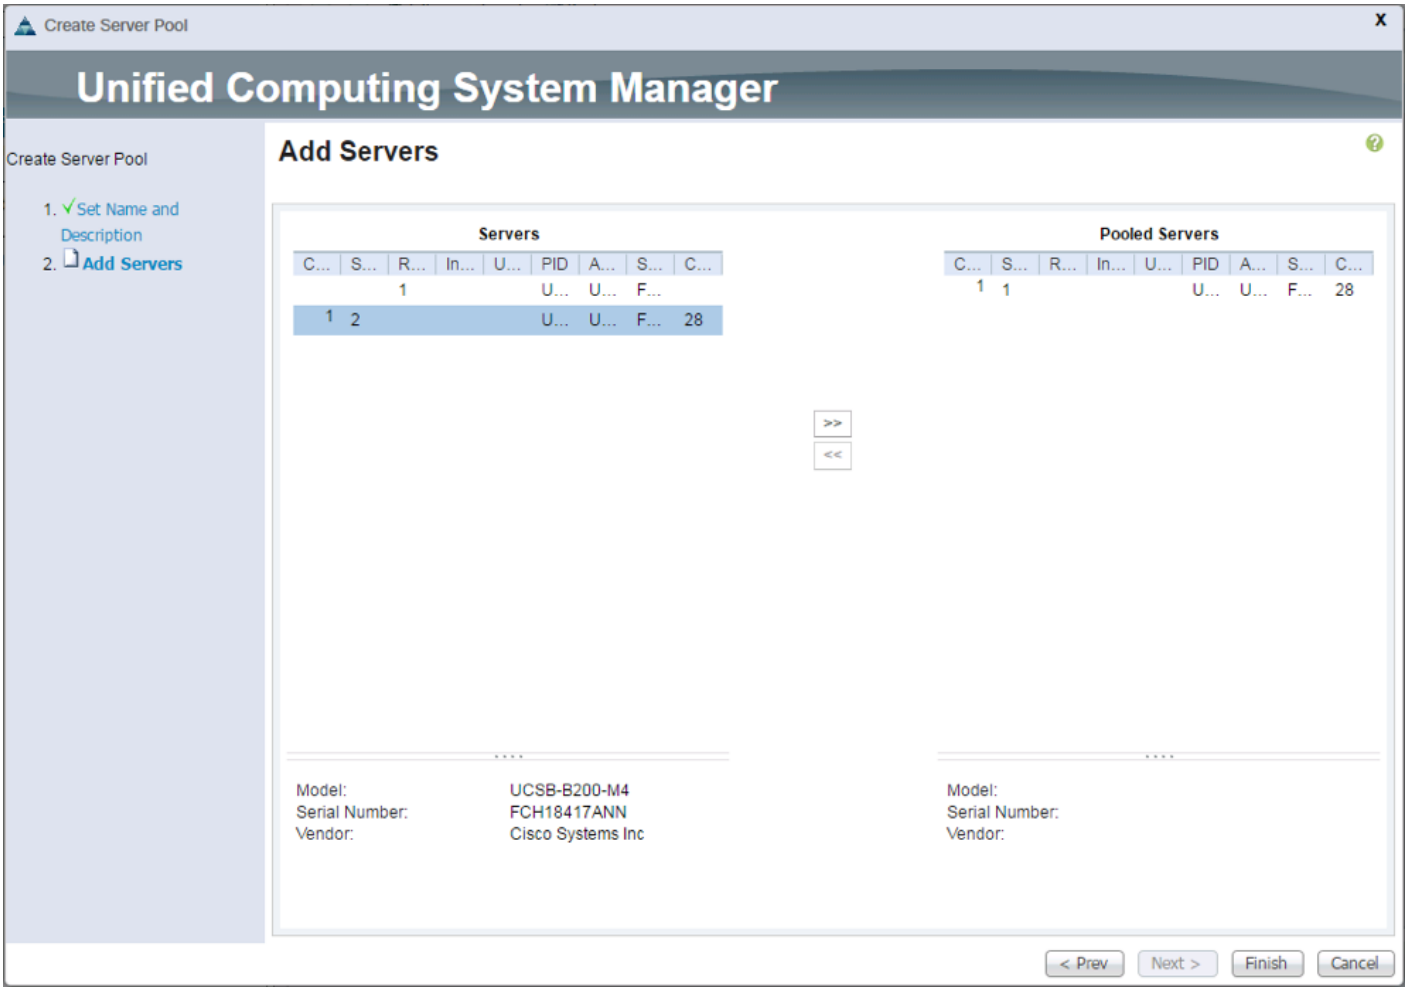

9. Click Finish.
10. Click OK.

Create VLANs

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

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

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

---

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `Native-VLAN` as the name of the VLAN to be used as the native VLAN.

6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.

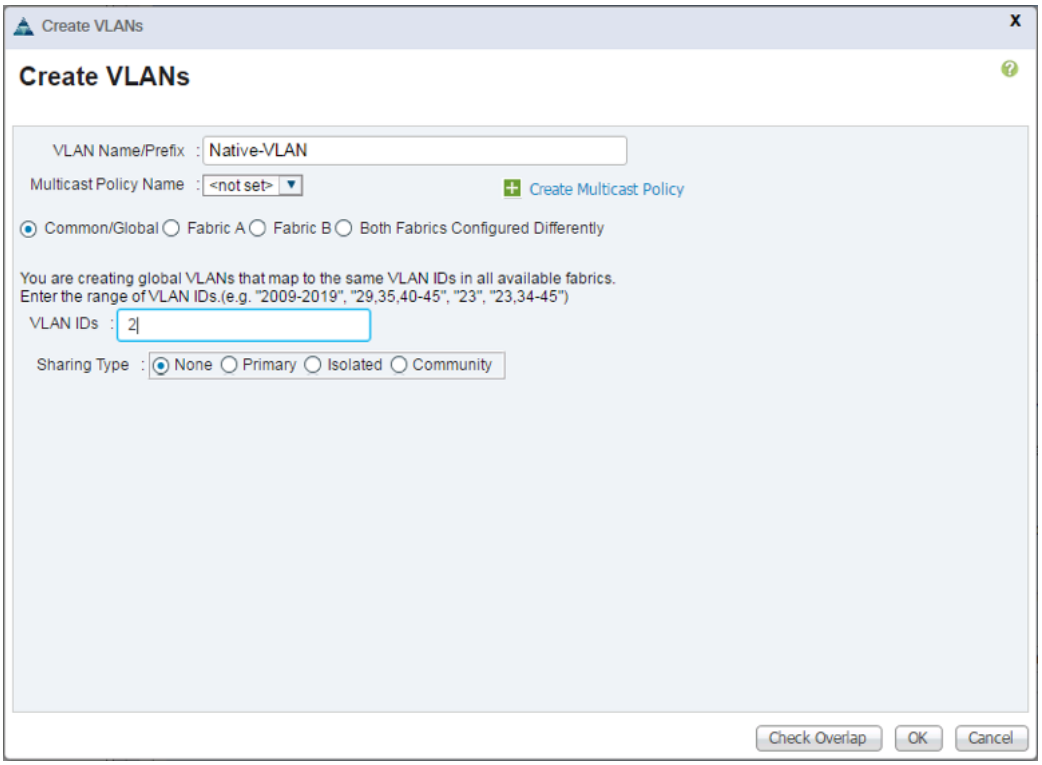

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs.
14. Enter `iSCSI-A-VLAN` as the name of the VLAN to be used for the first iSCSI VLAN.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the VLAN ID for the first iSCSI VLAN.
17. Click OK, then OK.

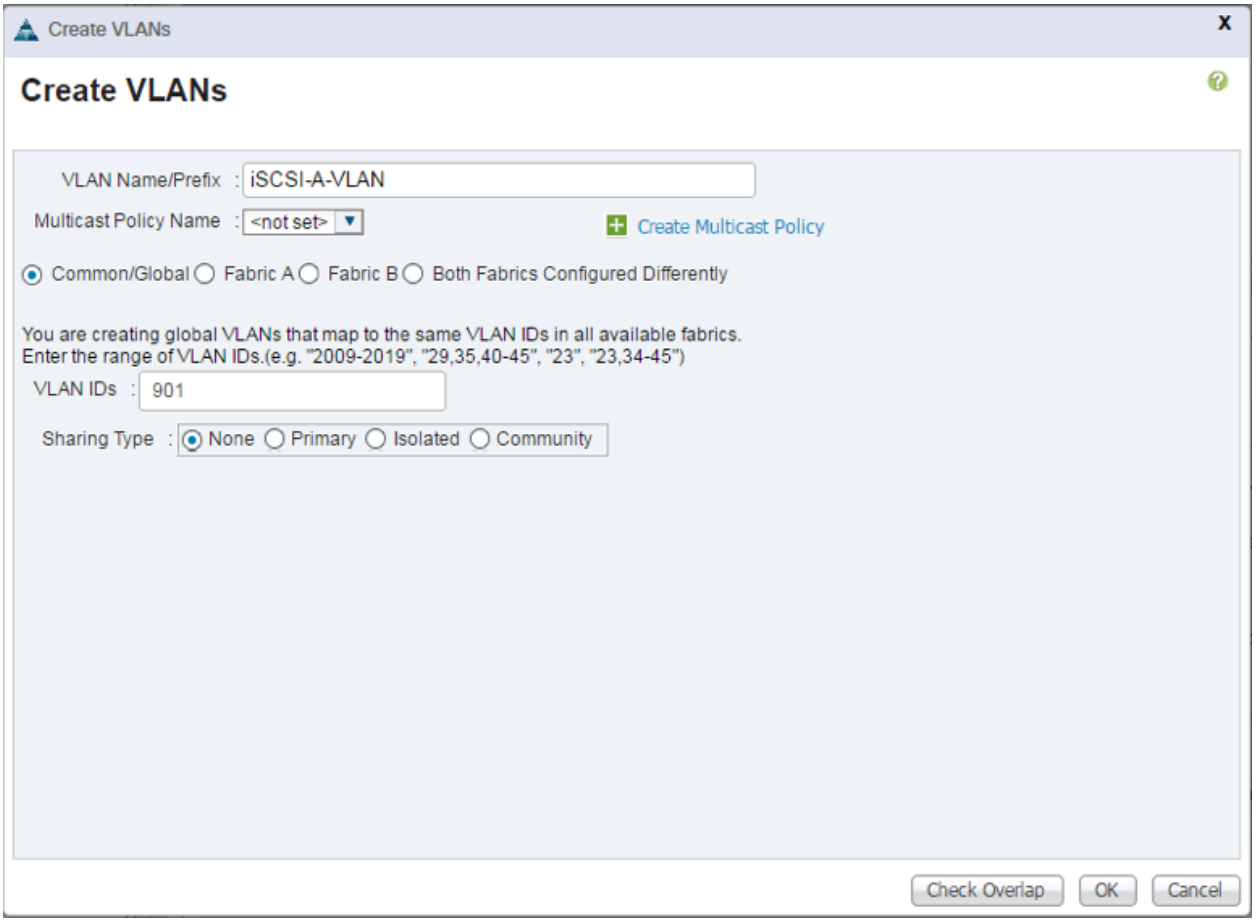

18. Right-click VLANs.
19. Select Create VLANs.
20. Enter `iSCSI-B-VLAN` as the name of the VLAN to be used for the second iSCSI VLAN.
21. Keep the Common/Global option selected for the scope of the VLAN.
22. Enter the VLAN ID for the second iSCSI VLAN.
23. Click OK, then OK.

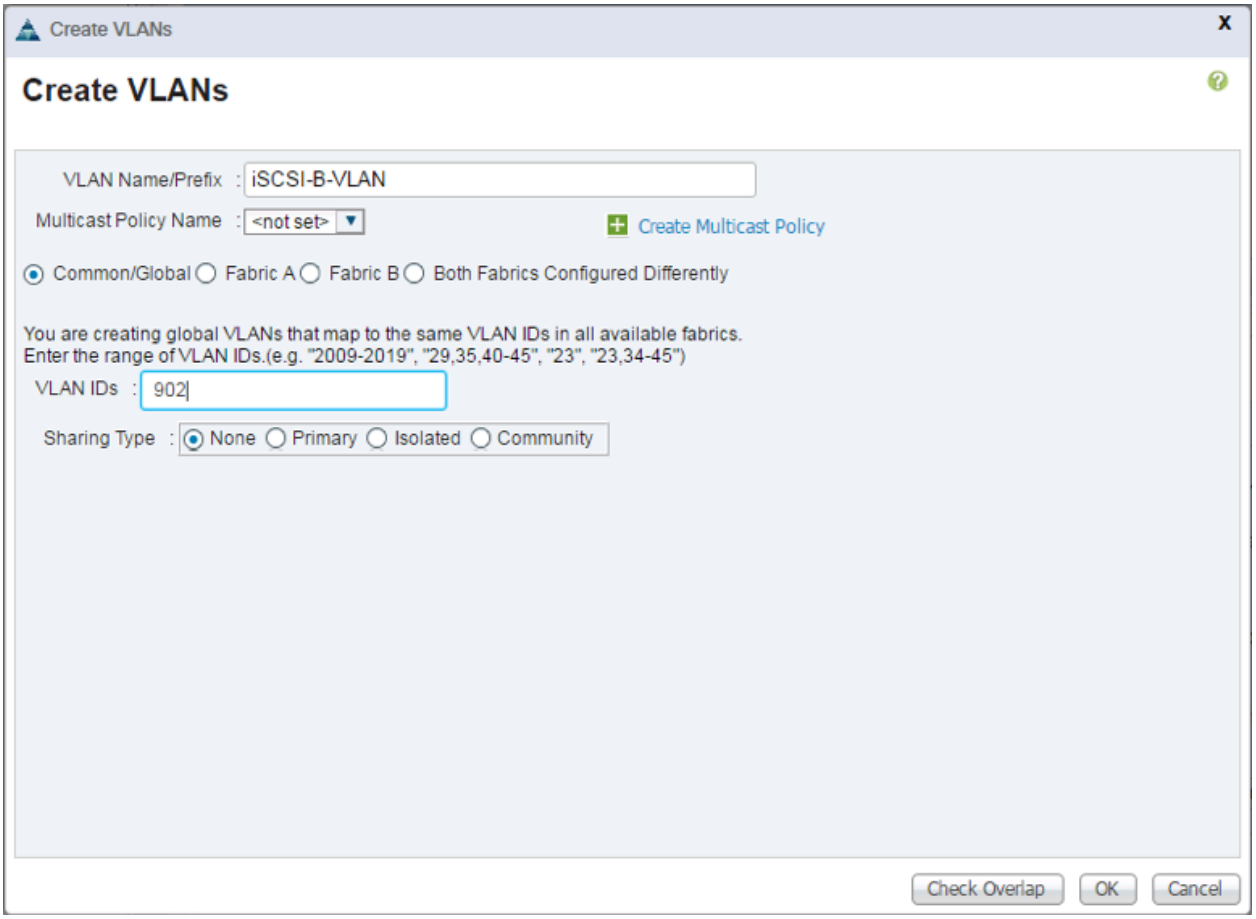

24. Right-click VLANs.
25. Select Create VLANs
26. Enter `IB-Mgmt` as the name of the VLAN to be used for management traffic.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the In-Band management VLAN ID.
29. Keep the Sharing Type as None.
30. Click OK, and then click OK again.
31. Right-click VLANs.
32. Select Create VLANs.
33. Enter `Infra-NFS` as the name of the VLAN to be used for NFS.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the NFS VLAN ID.

36. Keep the Sharing Type as None.
37. Click OK, and then click OK again.
38. Right-click VLANs.
39. Select Create VLANs.
40. Enter `vMotion` as the name of the VLAN to be used for vMotion.
41. Keep the Common/Global option selected for the scope of the VLAN.
42. Enter the vMotion VLAN ID.
43. Keep the Sharing Type as None.
44. Click OK, and then click OK again.
45. Right-click VLANs.
46. Select Create VLANs.
47. Enter `VM-Traffic` as the name of the VLAN to be used for VM Traffic.
48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VM-Traffic VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK, and then click OK again.

Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.1(1h) for both the Blade and Rack Packages.

7. Leave M-Series Package as <not set> and leave Excluded Components with only Local Disk selected.

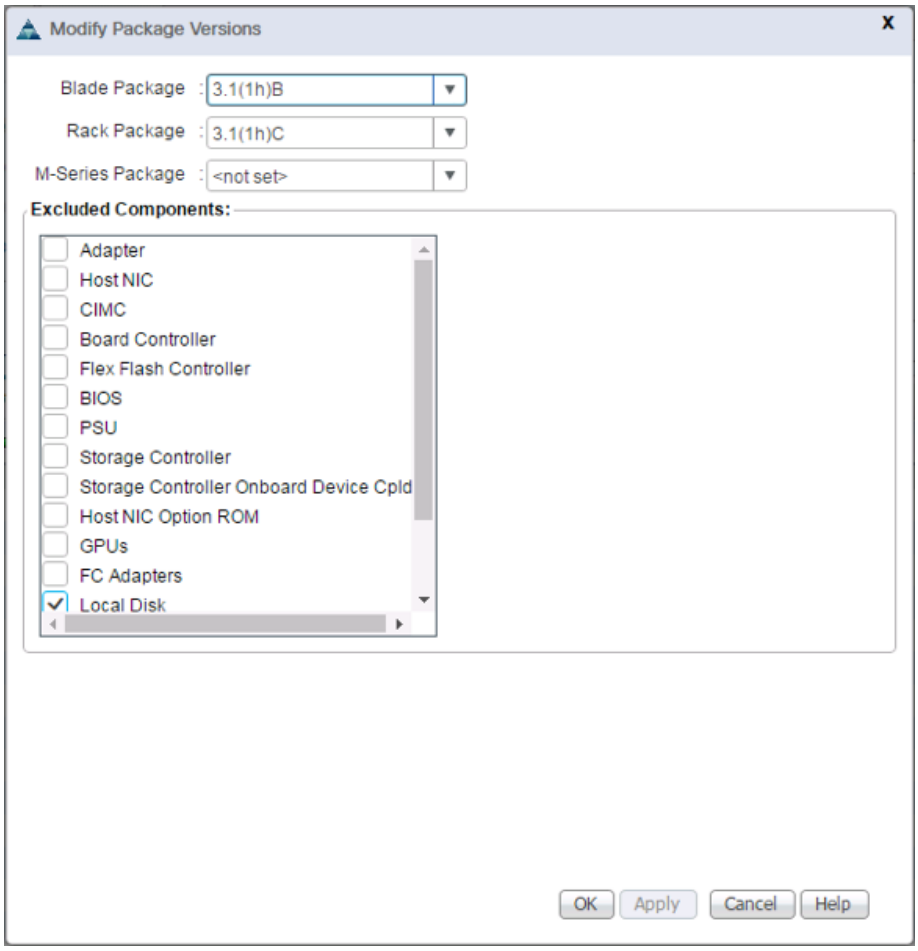

8. Click OK to modify the host firmware package.

Set Jumbo Frames in Cisco UCS Fabric

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

1. In Cisco UCS Manager, 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

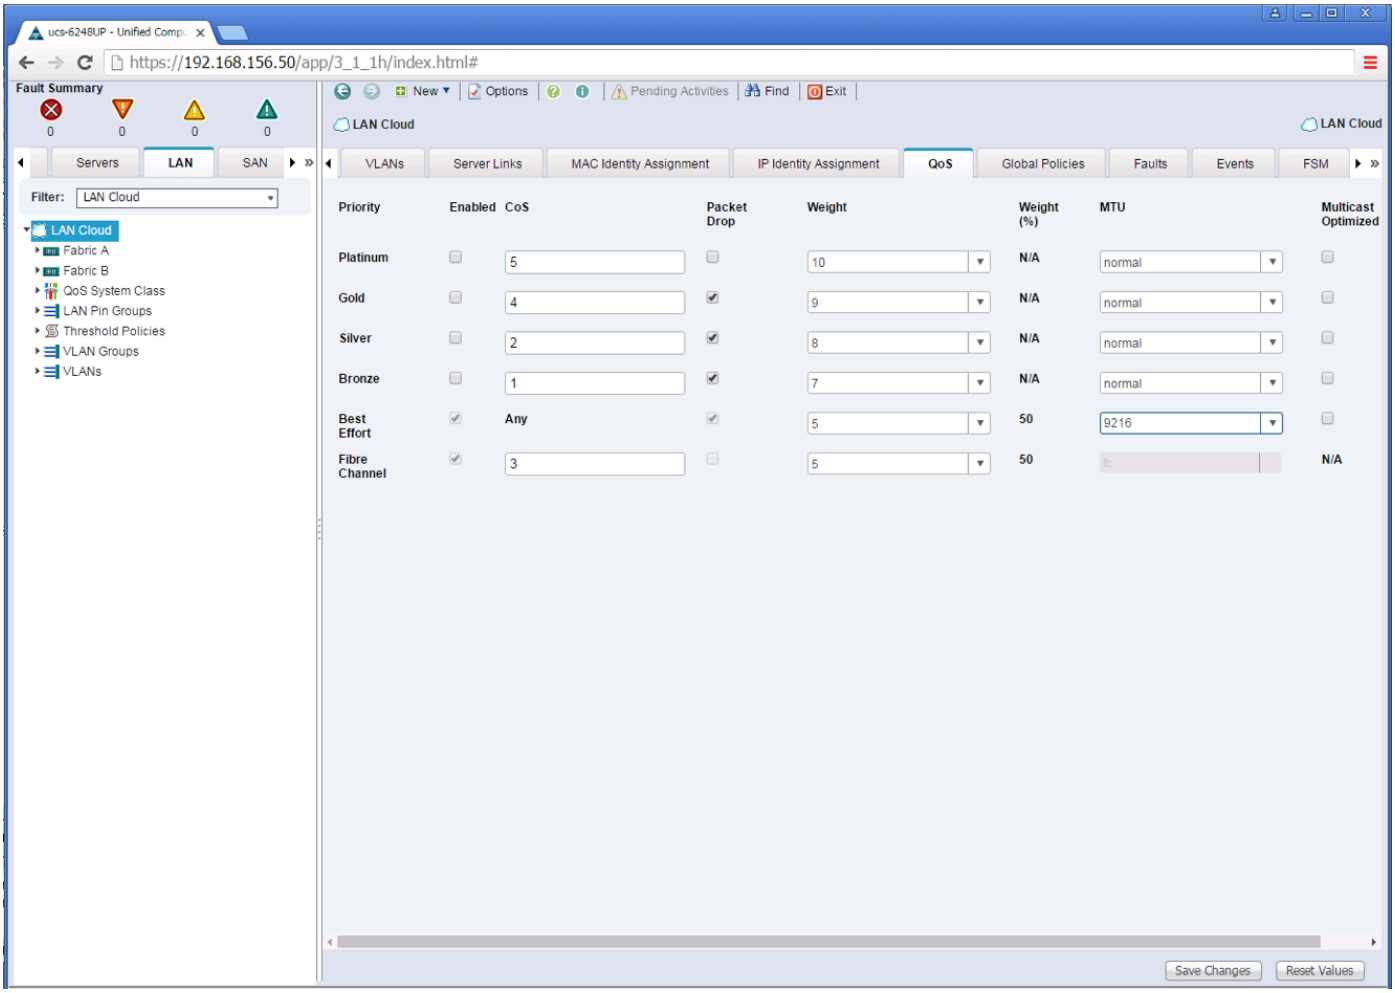

Create Local Disk Configuration Policy (Optional)

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

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

To create a local disk configuration policy, 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.
4. Select Create Local Disk Configuration Policy.
5. Enter `SAN-Boot` as the local disk configuration policy name.
6. Change the mode to No Local Storage.

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

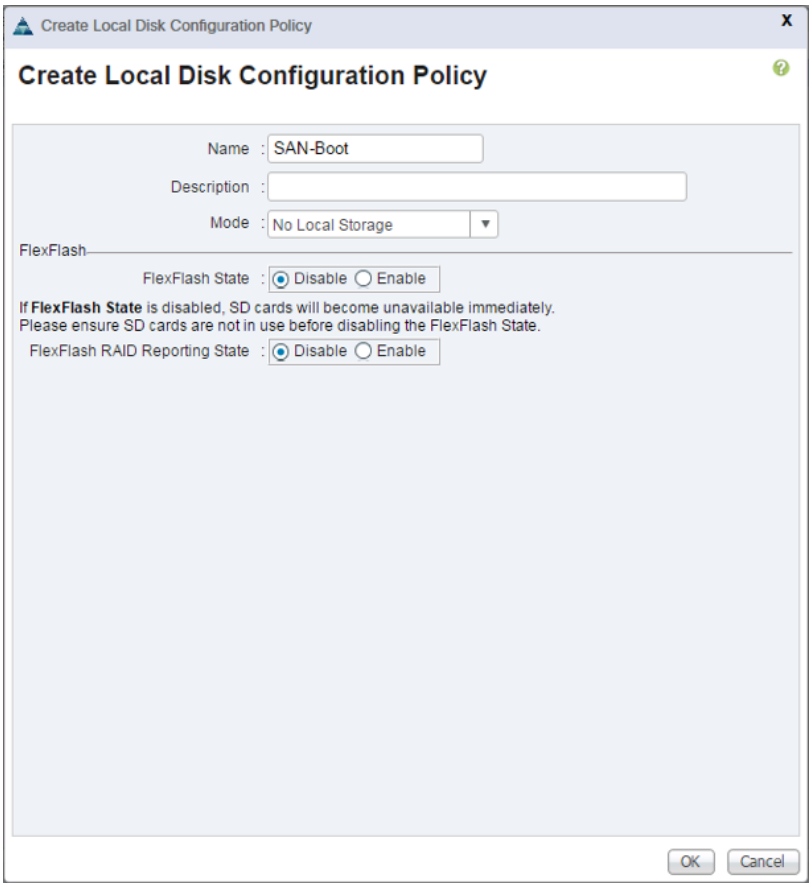

8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable_CDP` as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.

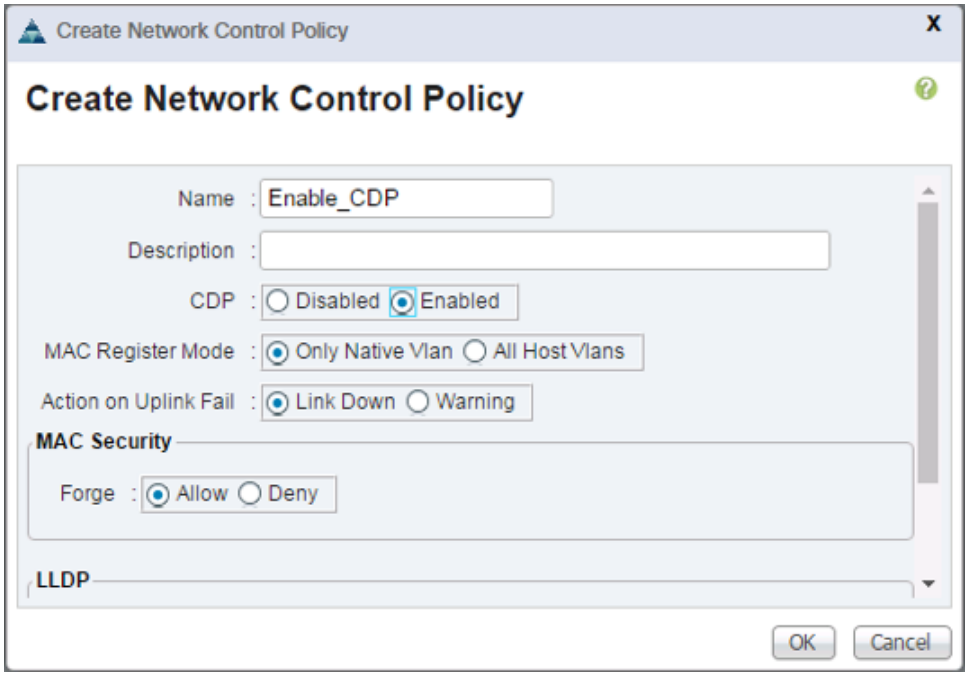

8. Click OK.

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.
4. Select Create Power Control Policy.
5. Enter `No-Power-Cap` as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

Create Server Pool Qualification Policy (Optional)

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

This example creates a policy for Cisco UCS B-Series and Cisco UCS C-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy `UCS-Broadwell`.
6. Select Create CPU/Core Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Select `UCS-CPU-E52660E` as the PID.
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy then OK for the confirmation.

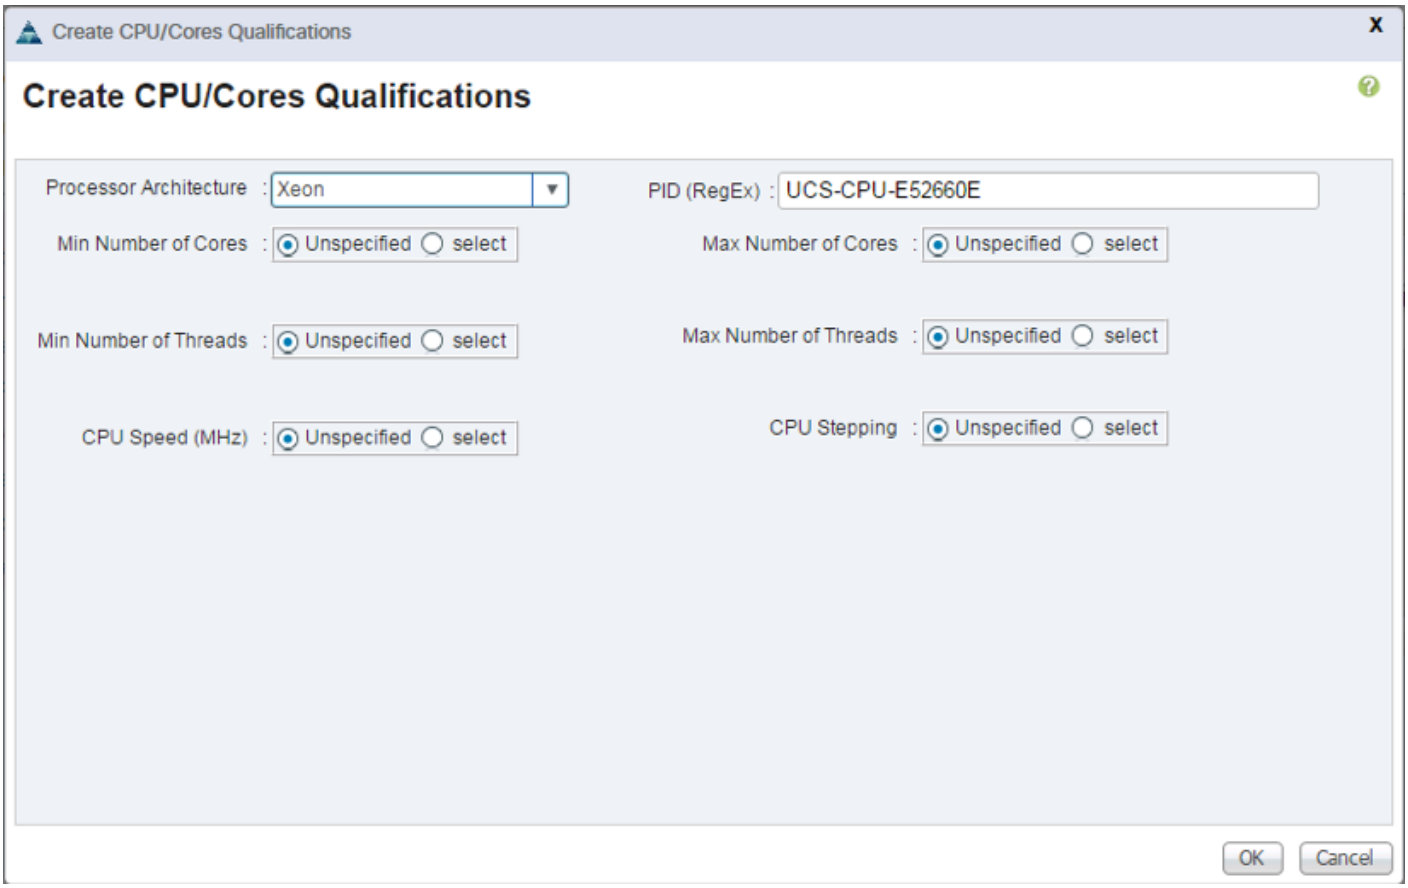

Create Server BIOS Policy

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.
4. Select Create BIOS Policy.
5. Enter `VM-Host-Infra` as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.
8. Click Finish to create the BIOS policy.

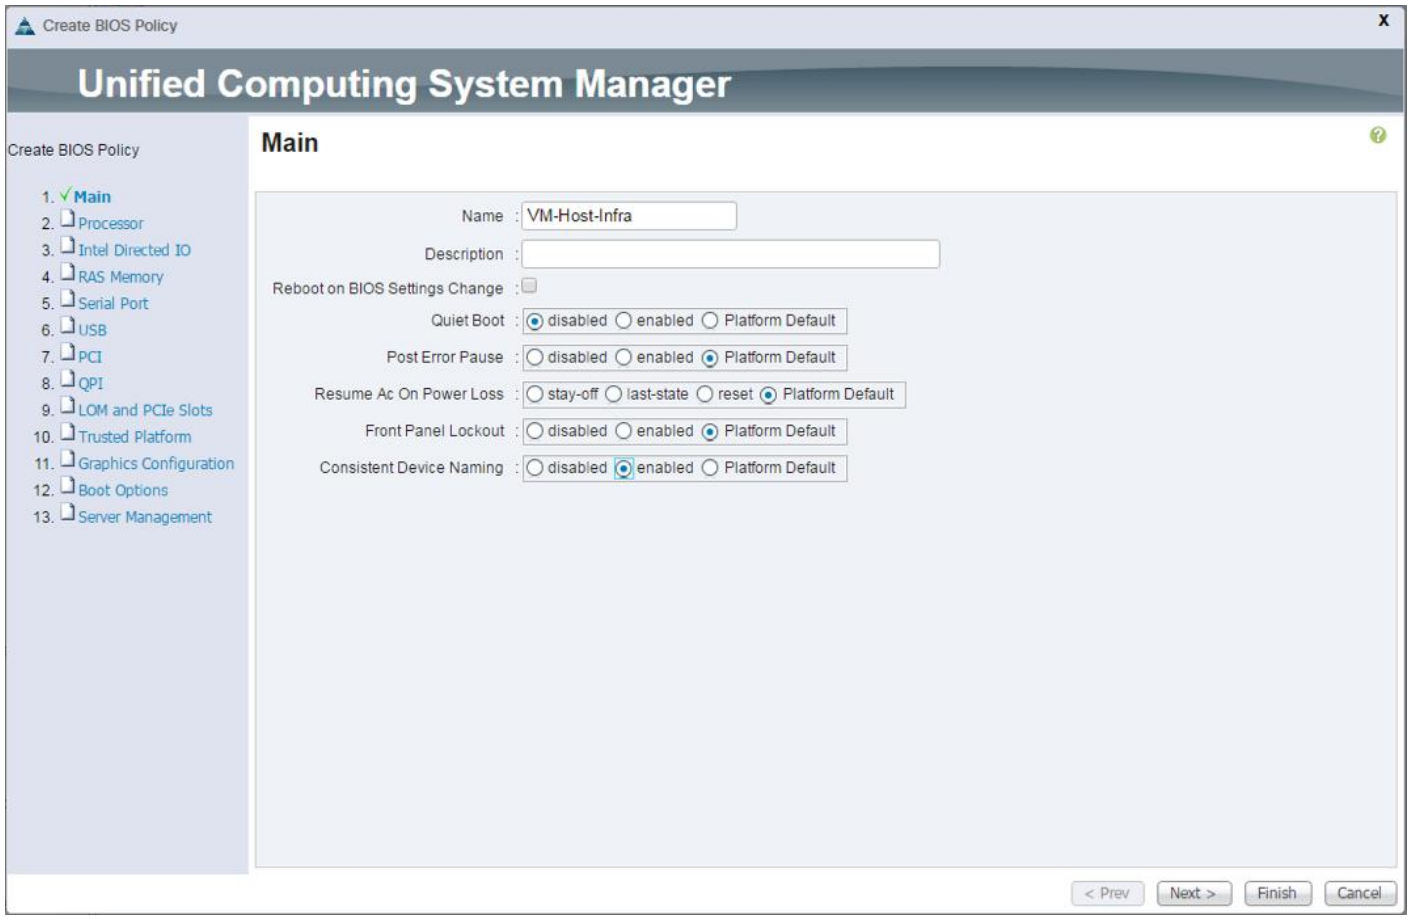

9. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

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

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

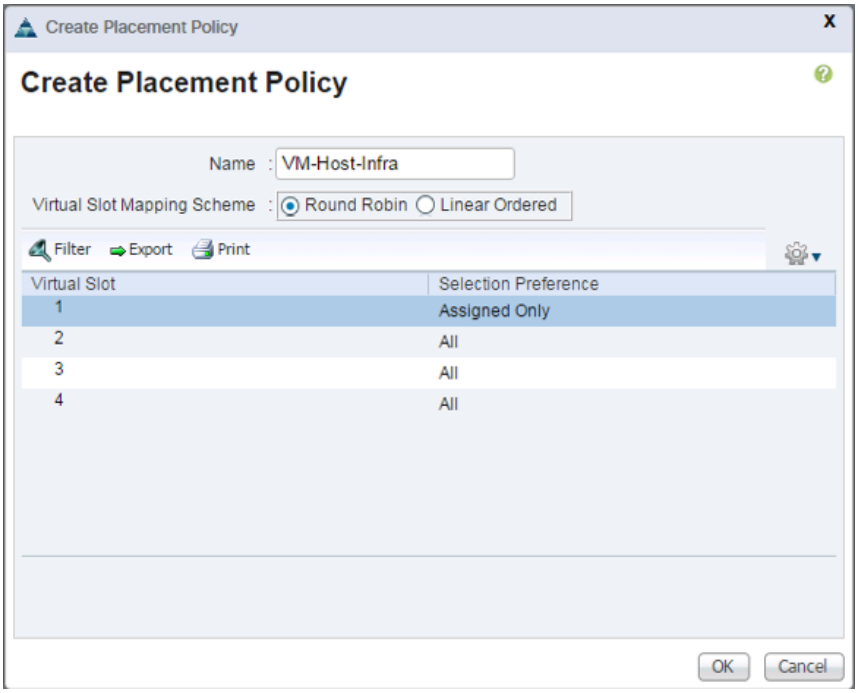

Update the Default Maintenance Policy

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.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. (Optional: Click “On Next Boot” to delegate maintenance windows to server owners)

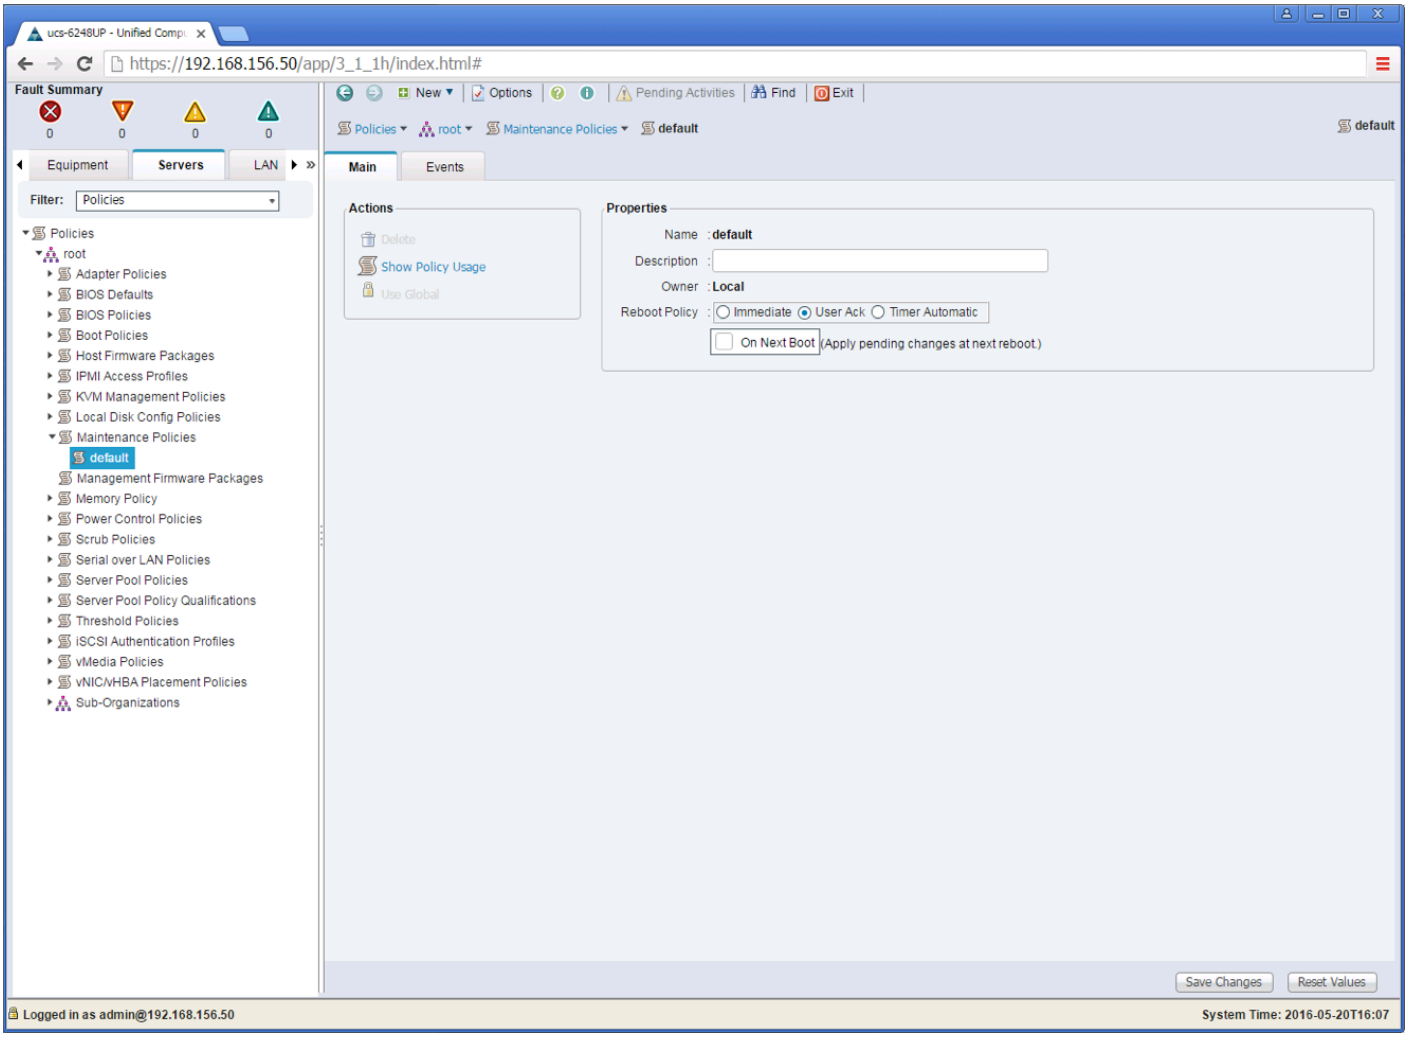

6. Click Save Changes.

7. Click OK to accept the change.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 4 vNIC Templates will be created.

The same Infra VLANs were used on both the infrastructure (Infra) and production (Prod) hosts deployed in this environment. If production networks were to differ in the VLANs used for infrastructure networks, differing vNIC Templates should be created for each.

Create Data vNICs

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_Template_A` as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.
10. Under VLANs, select the checkboxes for `IB-MGMT, Infra-NFS, Native-VLAN,VM-Traffic, and vMotion` VLANs.

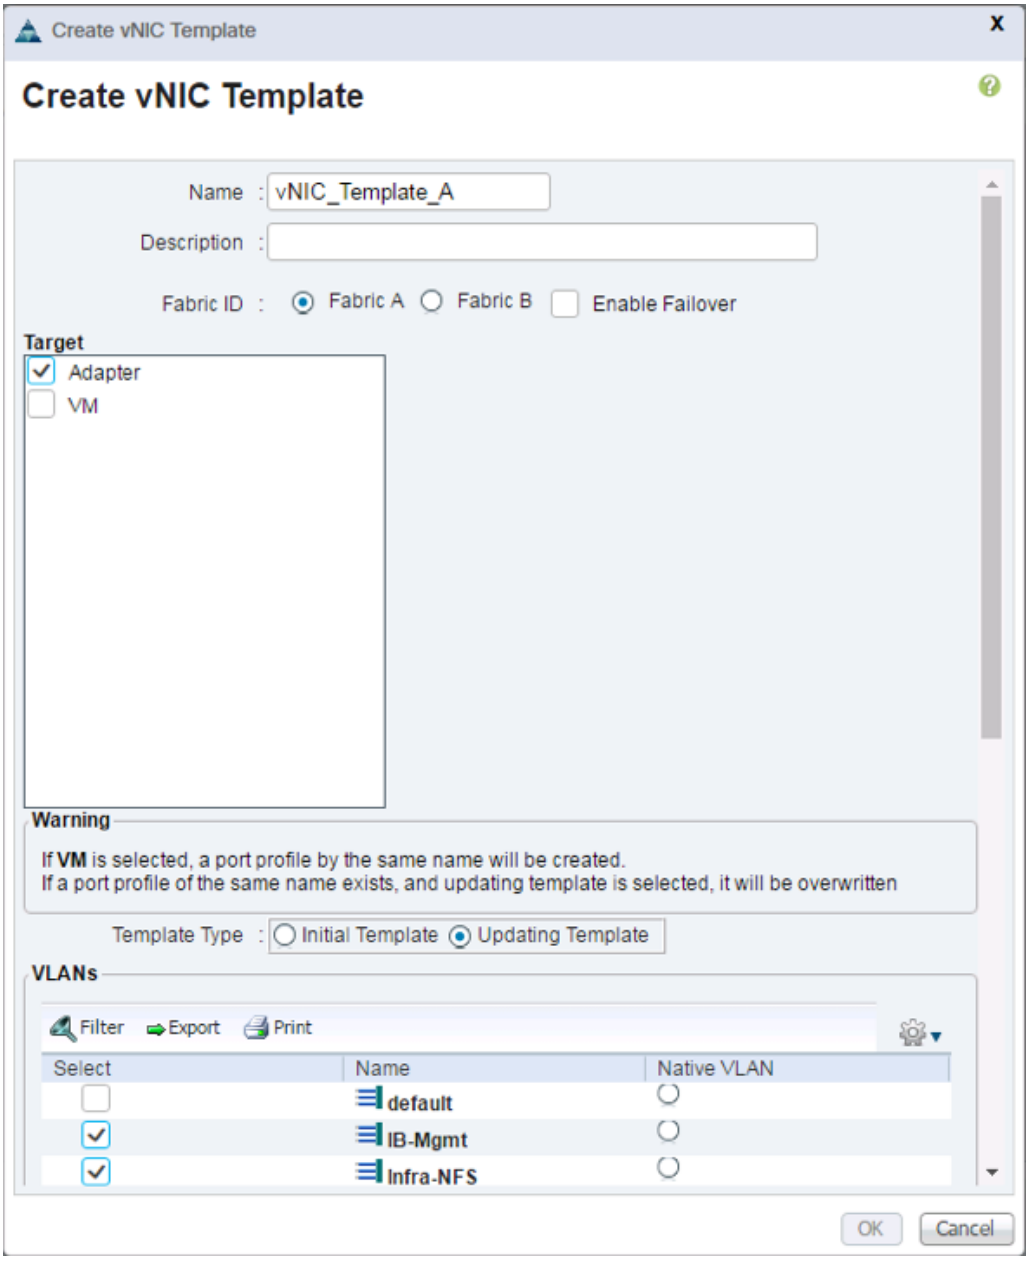

11. Set `Native-VLAN` as the native VLAN.
12. For MTU, enter `9000`.
13. In the MAC Pool list, select `MAC_Pool_A`.
14. In the Network Control Policy list, select `Enable_CDP`.

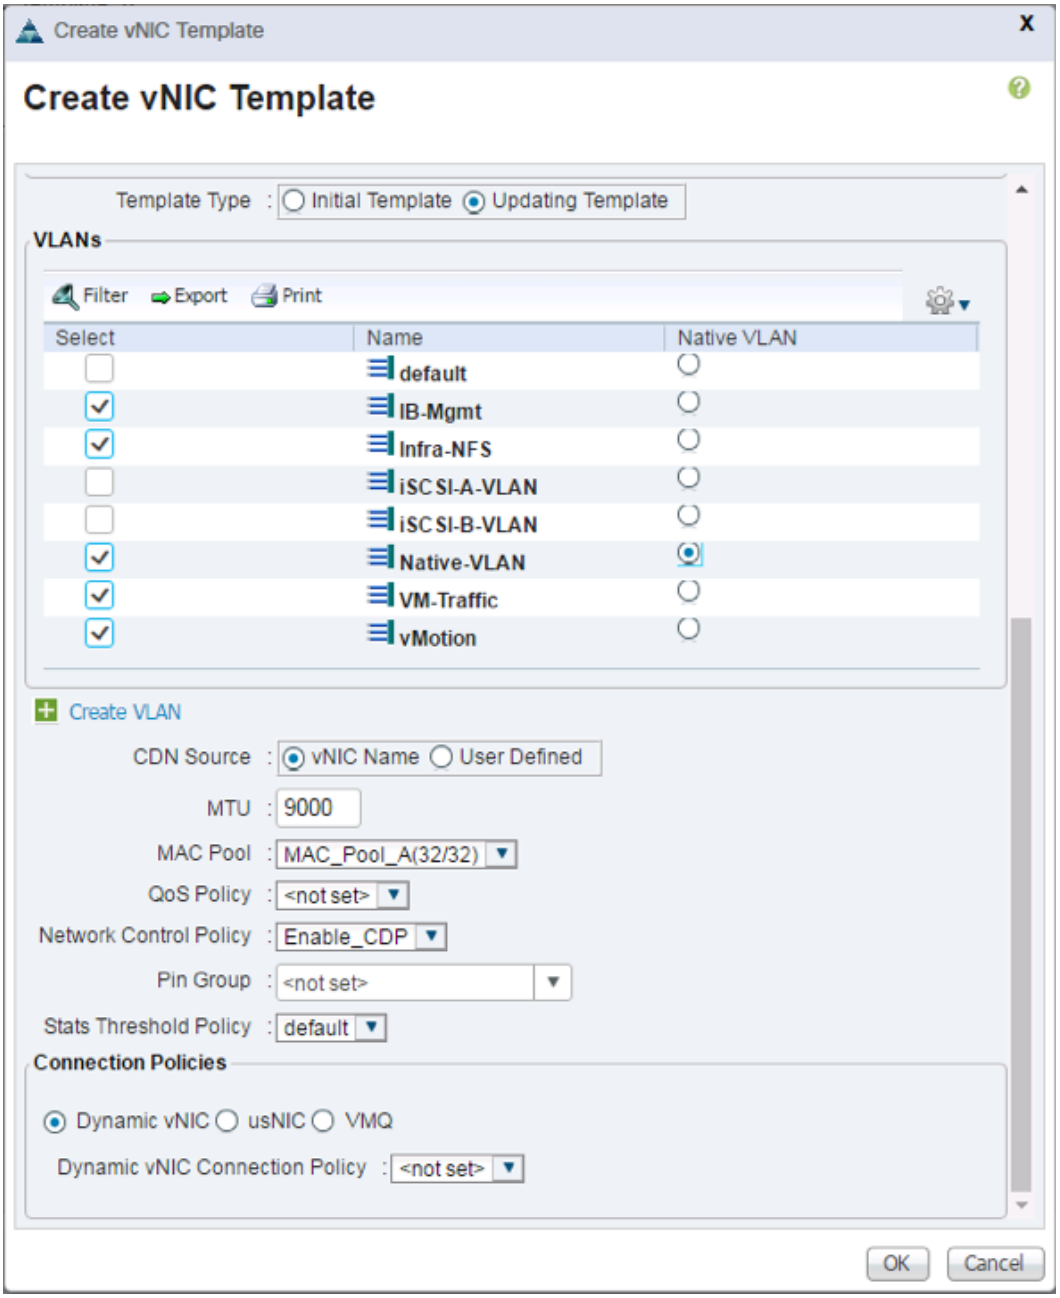

15. Click OK to create the vNIC template.
16. Click OK.

Repeat these equivalent steps for vNIC_Template_B:

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_Template_B` as the vNIC template name.
6. Select Fabric B.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure the VM checkbox is not selected.
9. Select Updating Template as the template type.
10. Under VLANs, select the checkboxes for `IB-MGMT`, `INFRA-NFS`, `Native-VLAN`, and `vMotion` VLANs.

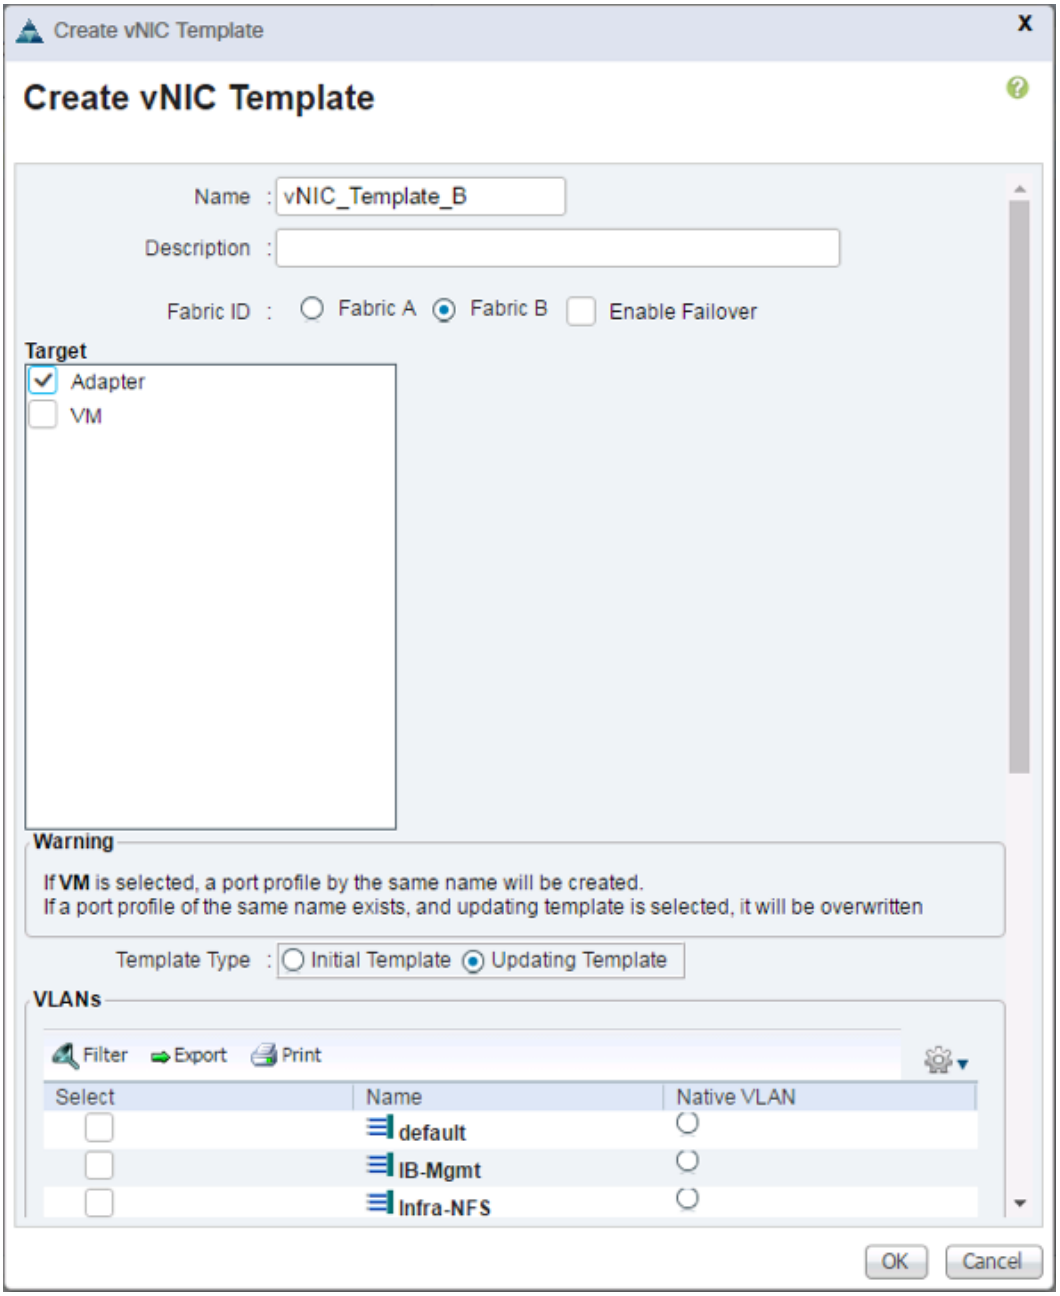

11. Set `default` as the native VLAN.
12. Select `vNIC Name` for the CDN Source.
13. For MTU, enter `9000`.
14. In the MAC Pool list, select `MAC_Pool_B`.
15. In the Network Control Policy list, select `Enable_CDP`.

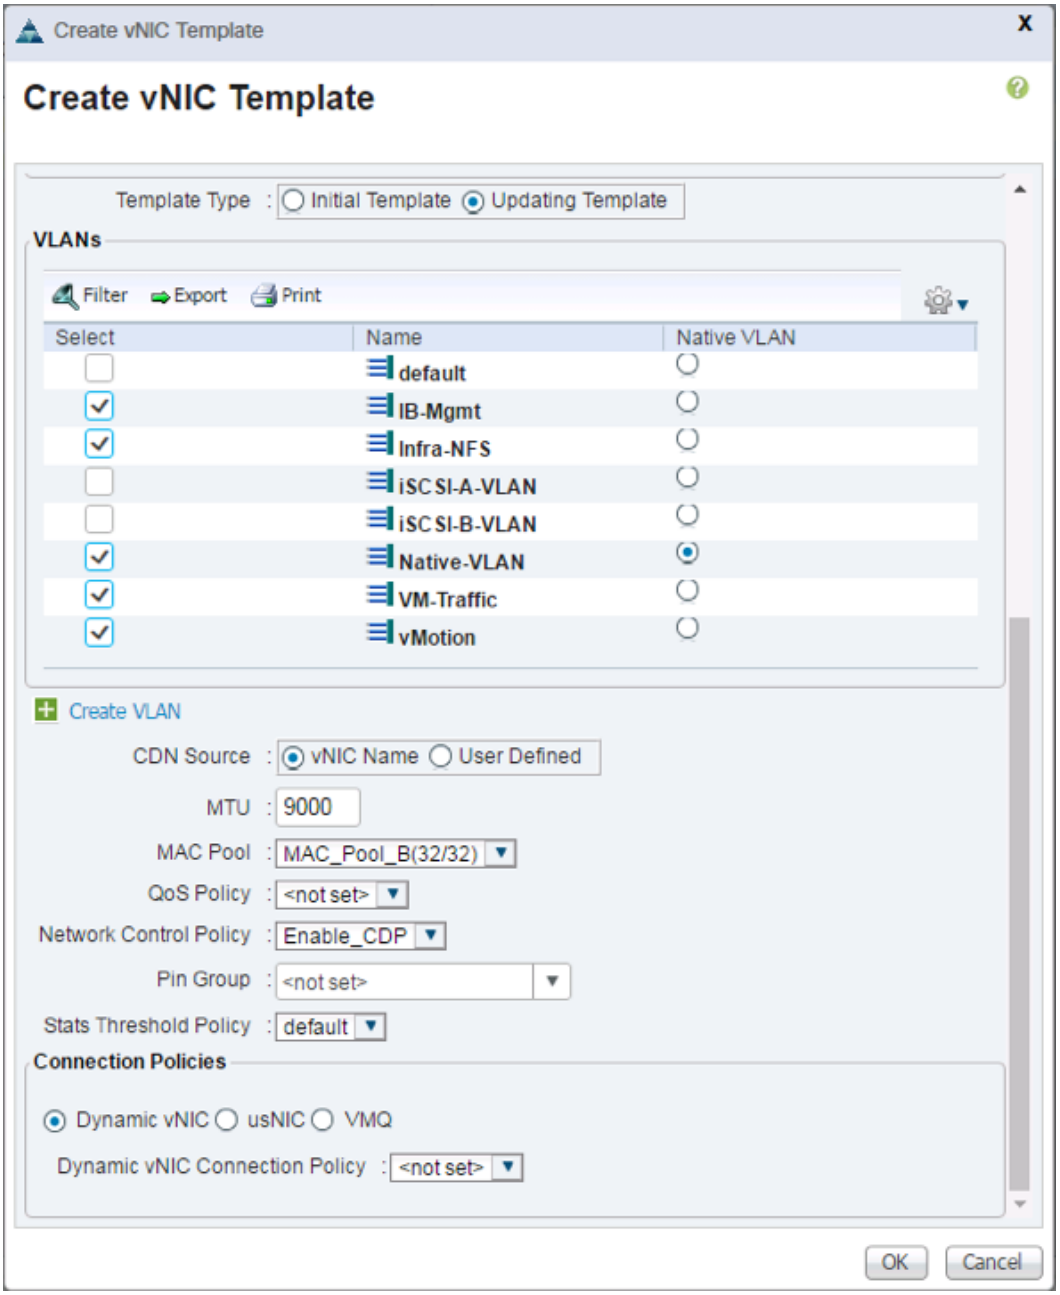

16. Click OK to create the vNIC template.
17. Click OK.

Create iSCSI vNICs

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.

5. Enter `iSCSI_Template_A` as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.

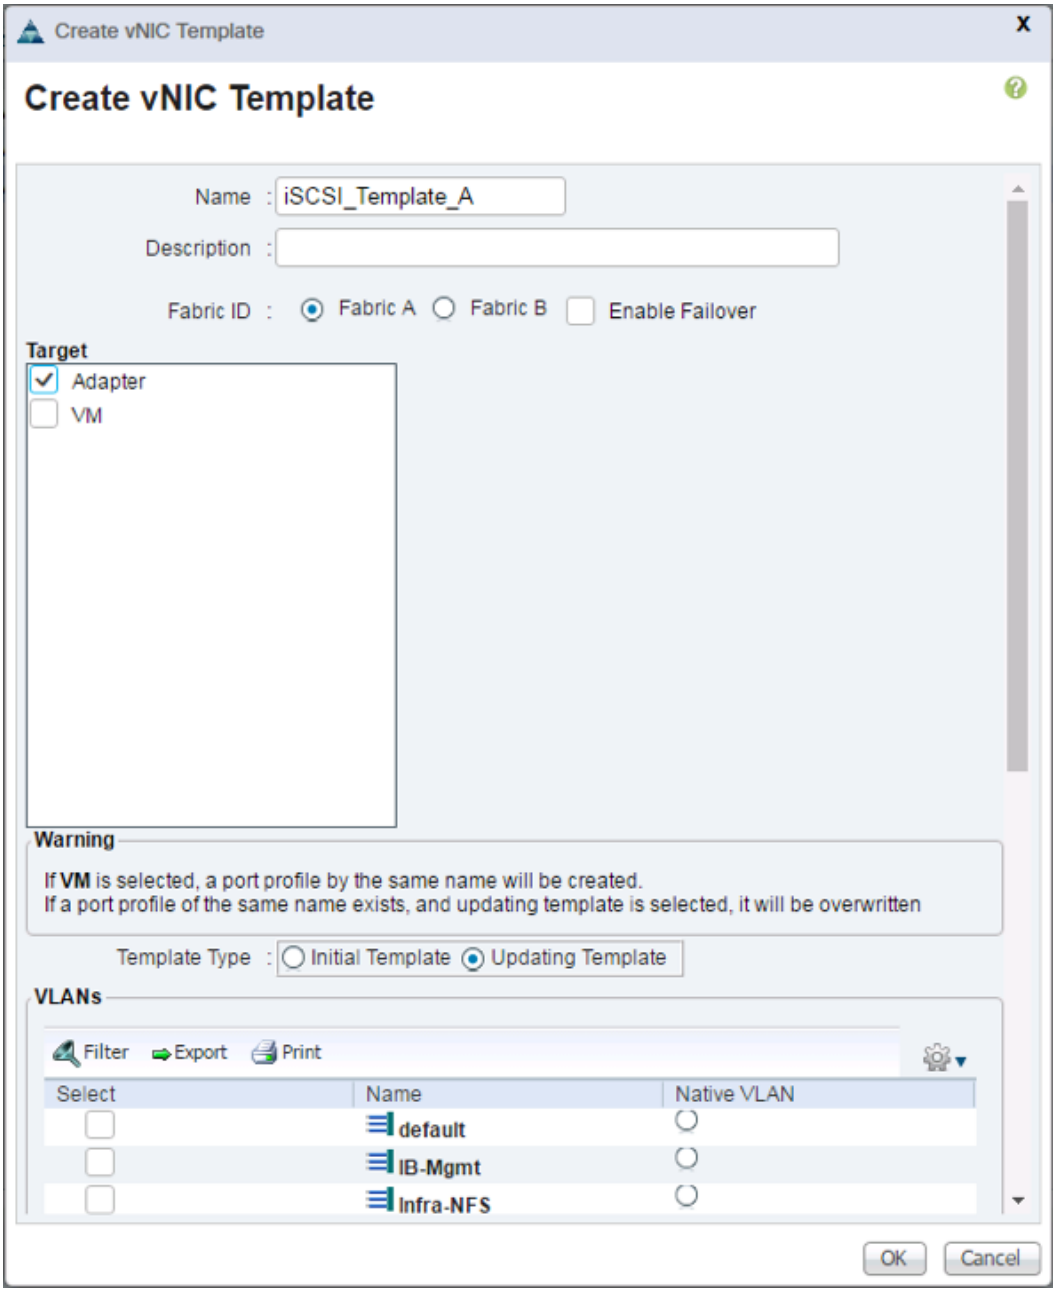

9. Under VLANs, select `iSCSI-A-VLAN`.
10. Set `iSCSI-A-VLAN` as the native VLAN.
11. Under MTU, enter 9000.

12. From the MAC Pool list, select `MAC_Pool_A`.

13. From the Network Control Policy list, select `Enable_CDP`.

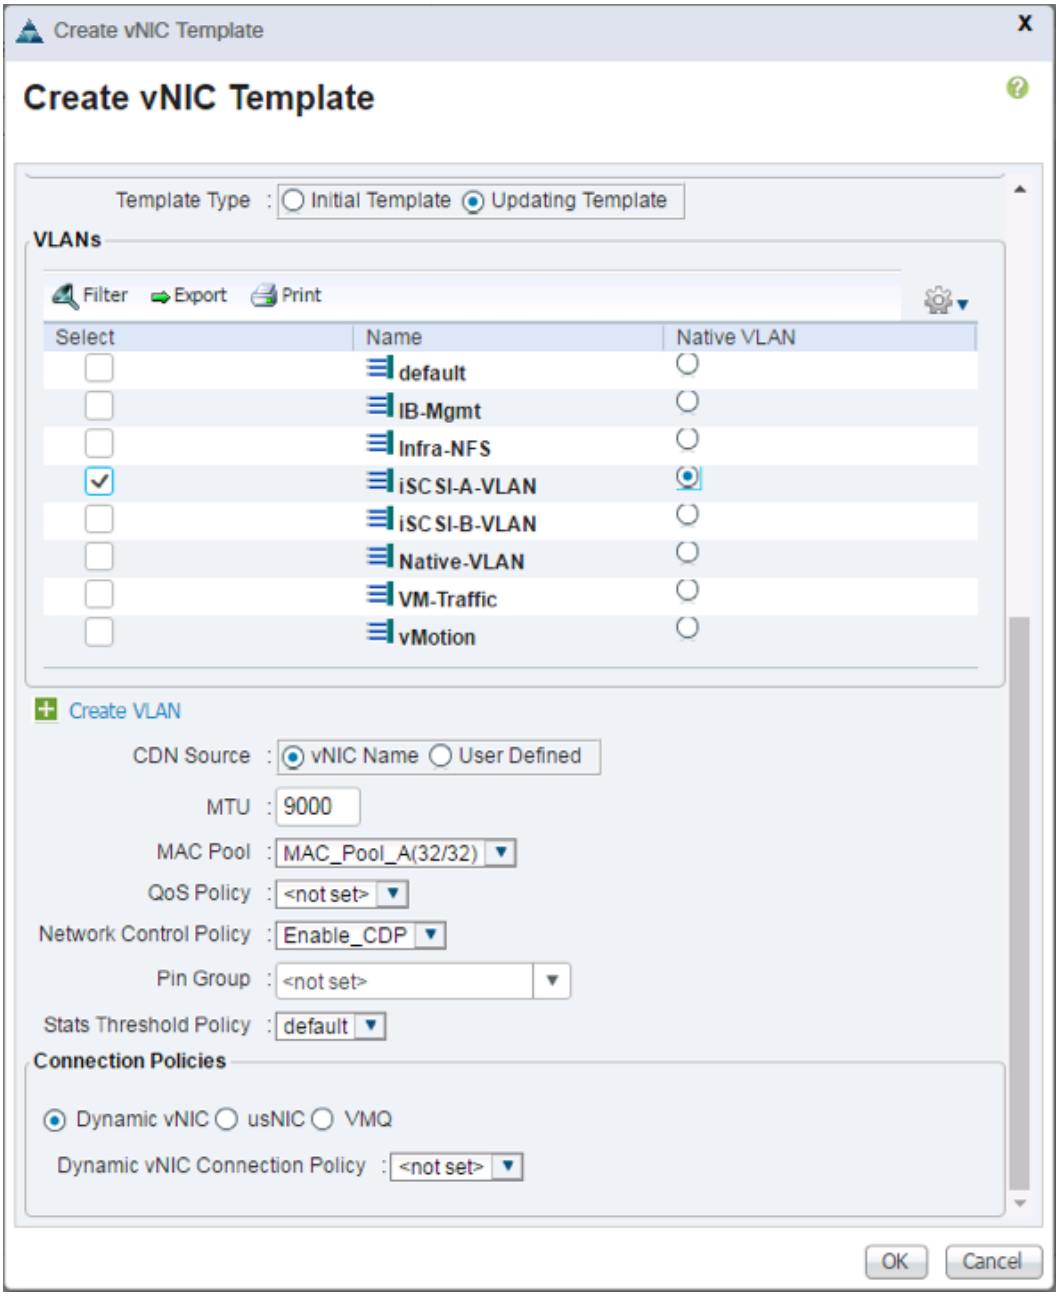

14. Click OK to complete creating the vNIC template.

15. Click OK.

Repeat these equivalent steps for iSCSI_Template_B:

1. Select the LAN tab on the left.

2. Select Policies > root.

3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `iSCSI_Template_B` as the vNIC template name.
6. Select Fabric B. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.

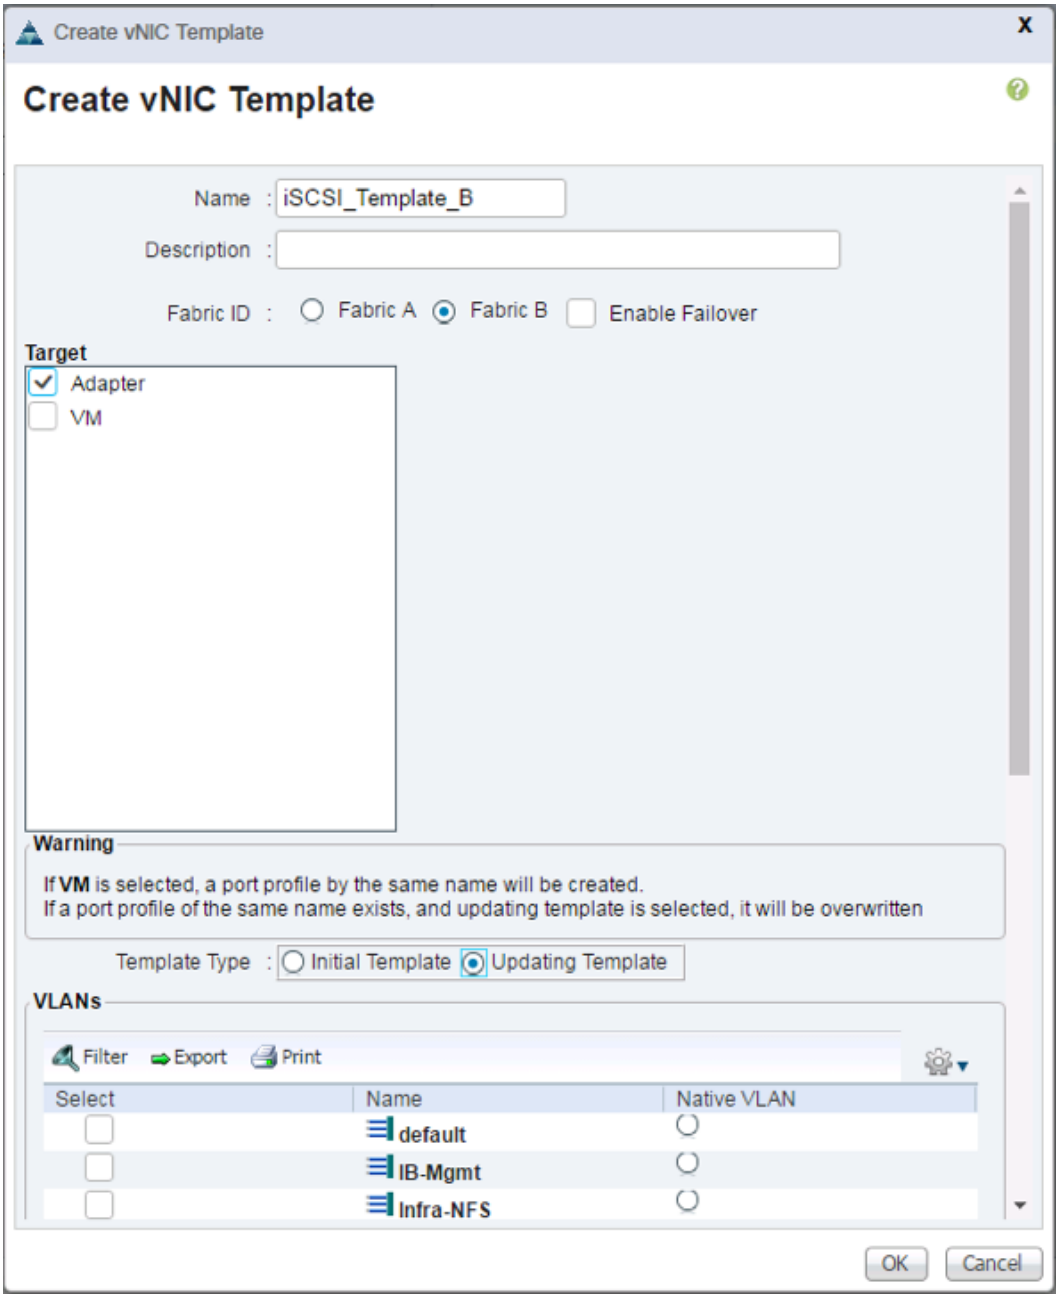

9. Under VLANs, select `iSCSI-B-VLAN`.

10. Set `iSCSI-B-VLAN` as the native VLAN.

11. Under MTU, enter 9000.

12. From the MAC Pool list, select `MAC_Pool_B`.

13. From the Network Control Policy list, select `Enable_CDP`.

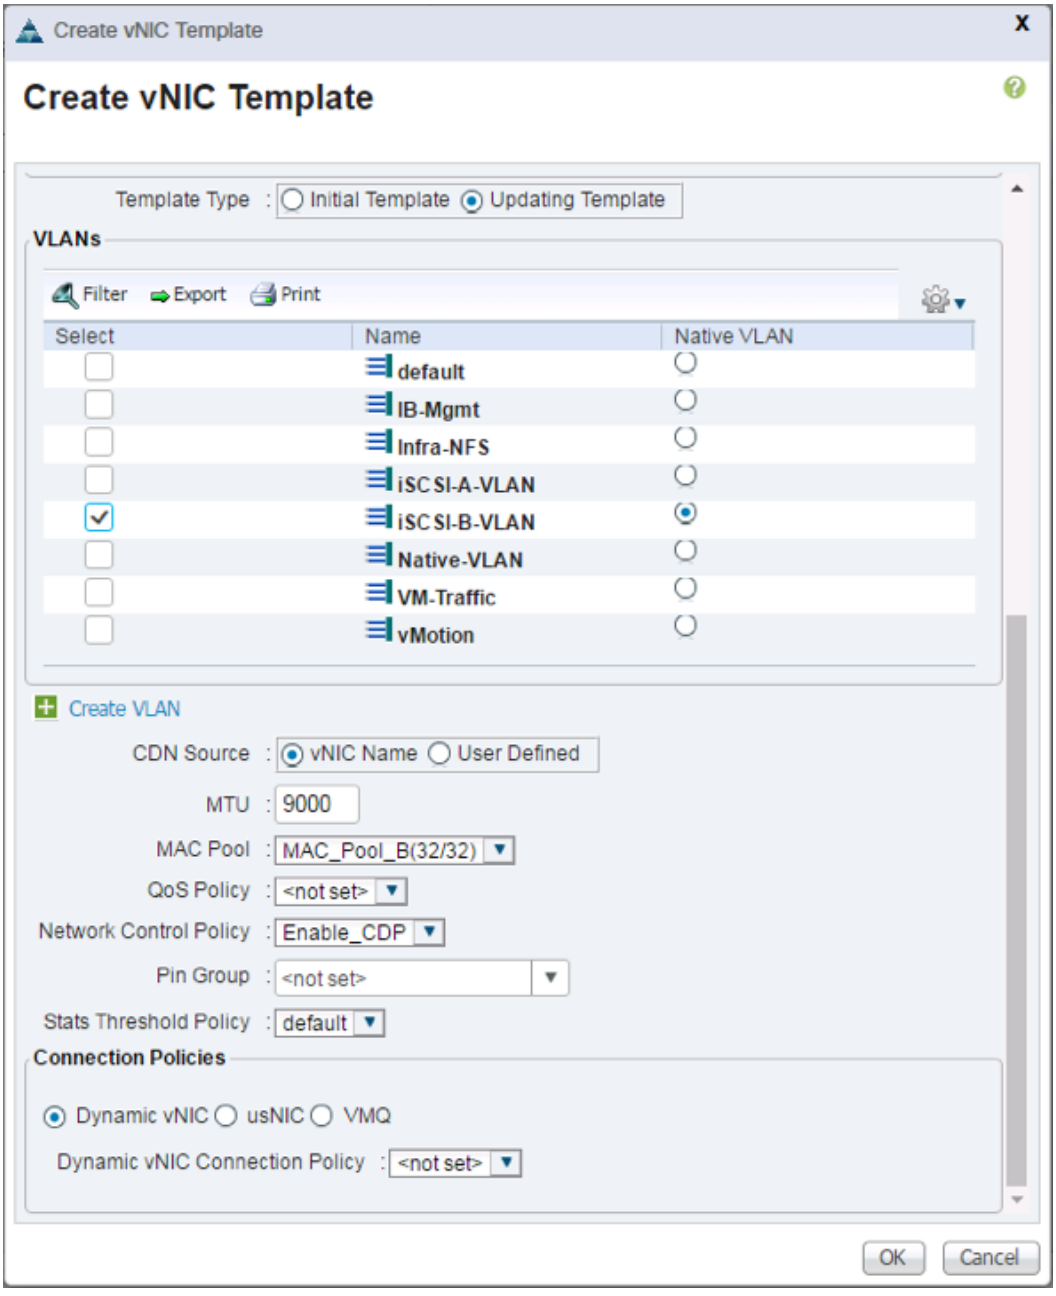

14. Click OK to complete creating the vNIC template.

15. Click OK.

Create LAN Connectivity Policy

During testing, iSCSI vNICs and FC vHBAs were deployed to all hosts regardless of boot policy used. The iSCSI interfaces can be left out for environments configured to utilize fibre channel for boot and data.

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter Infra-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-Infra-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select vNIC_Template_A.
10. In the Adapter Policy list, select VMWare.
11. Click OK to add this vNIC to the policy.

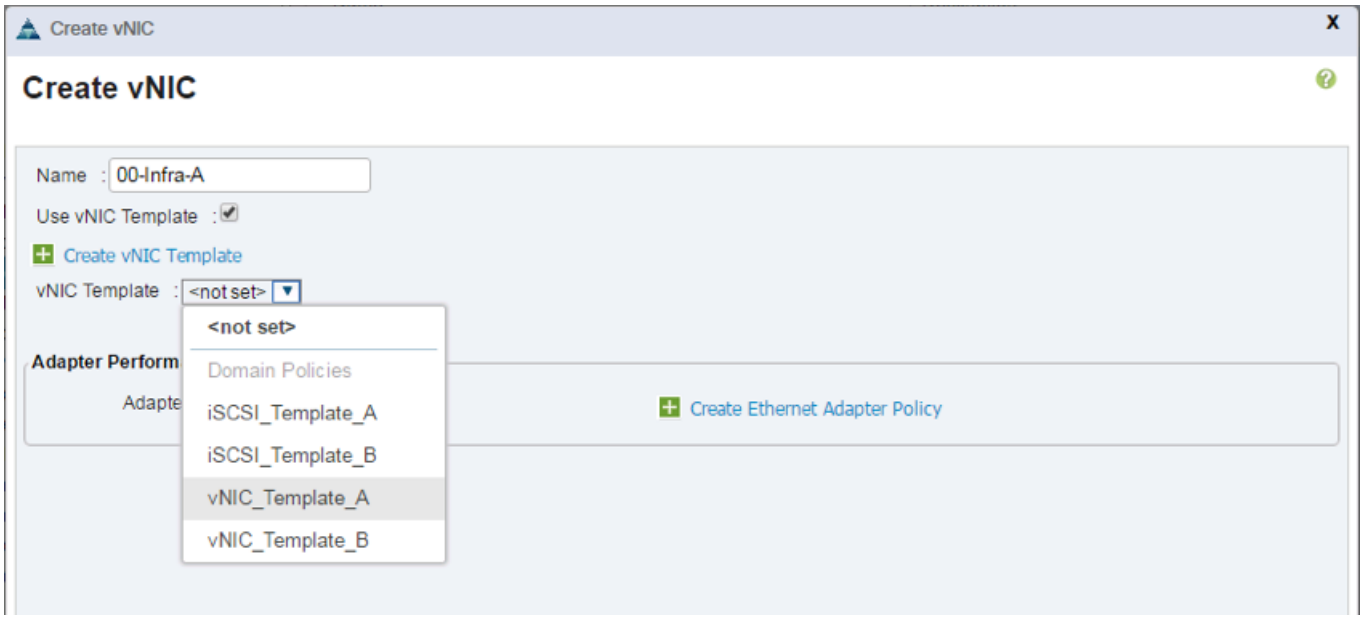

12. Click the upper Add button to add another vNIC to the policy.
13. In the Create vNIC box, enter vNIC-01-Infra-B as the name of the vNIC.

14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select vNIC_Template_B.
16. In the Adapter Policy list, select VMWare.
17. Click OK to add the vNIC to the policy.

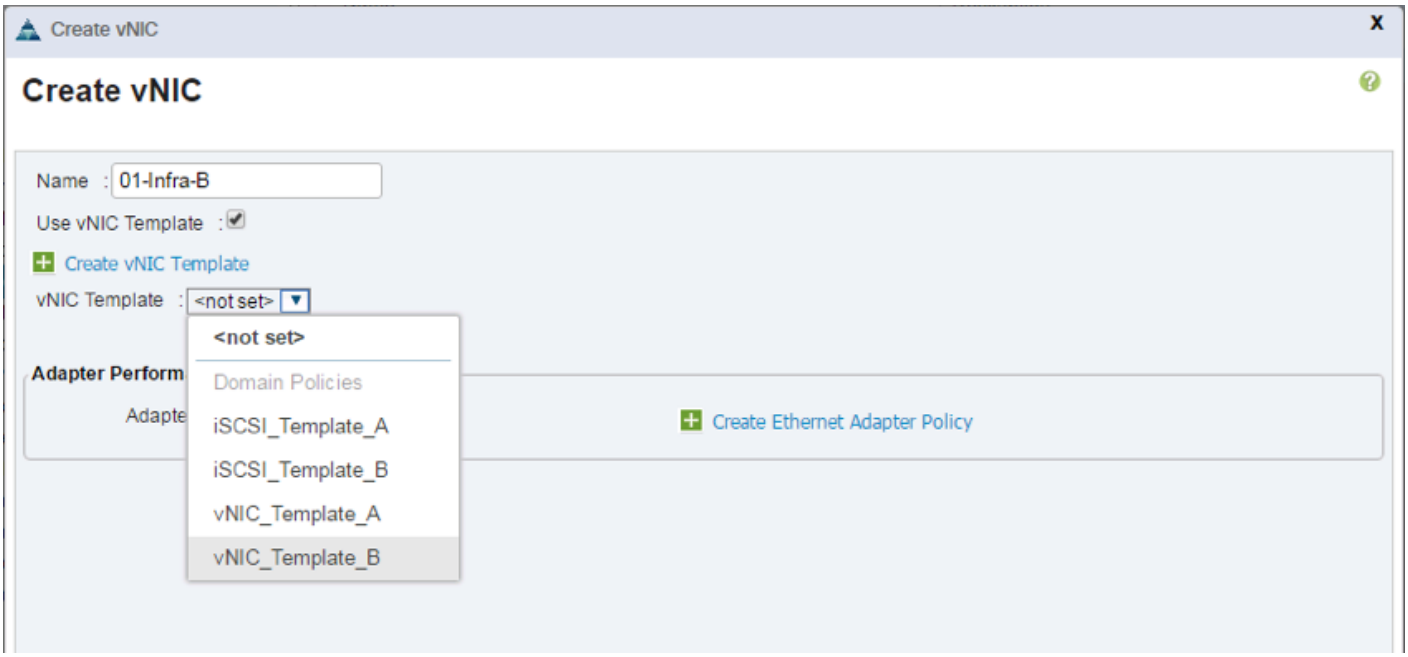

18. Click the upper Add button to add a vNIC.
19. In the Create vNIC dialog box, enter vNIC-02-iSCSI-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select iSCSI_Template_A.
22. In the Adapter Policy list, select VMWare.
23. Click OK to add this vNIC to the policy.

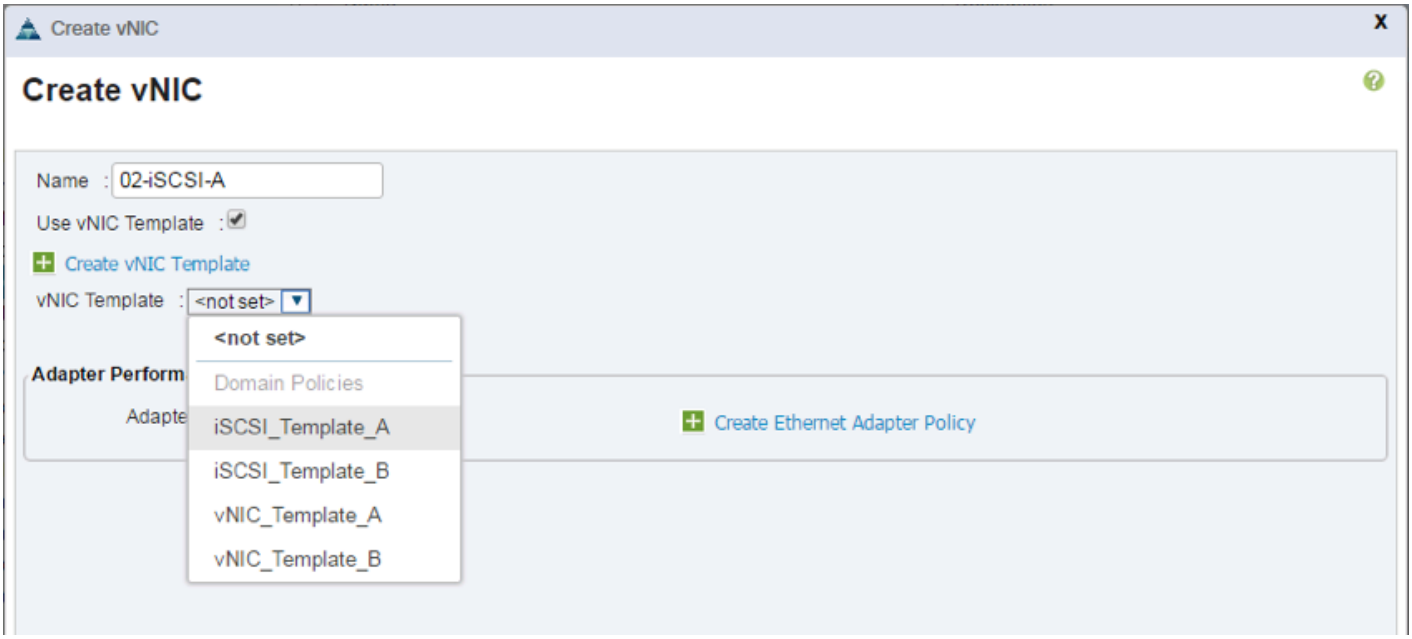

24. Click the upper Add button to add a vNIC to the policy.
25. In the Create vNIC dialog box, enter vNIC-03-iSCSI-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select iSCSI_Template_B.
28. In the Adapter Policy list, select VMWare.

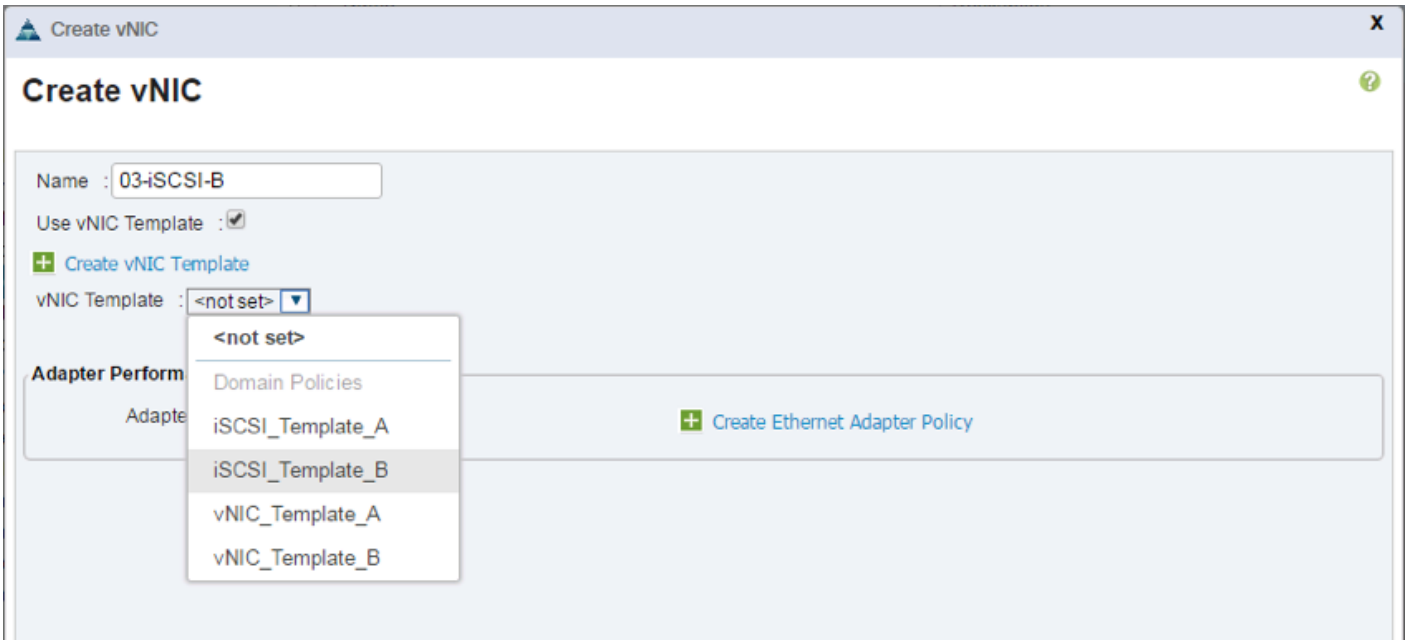

29. Click OK to add this vNIC to the policy.

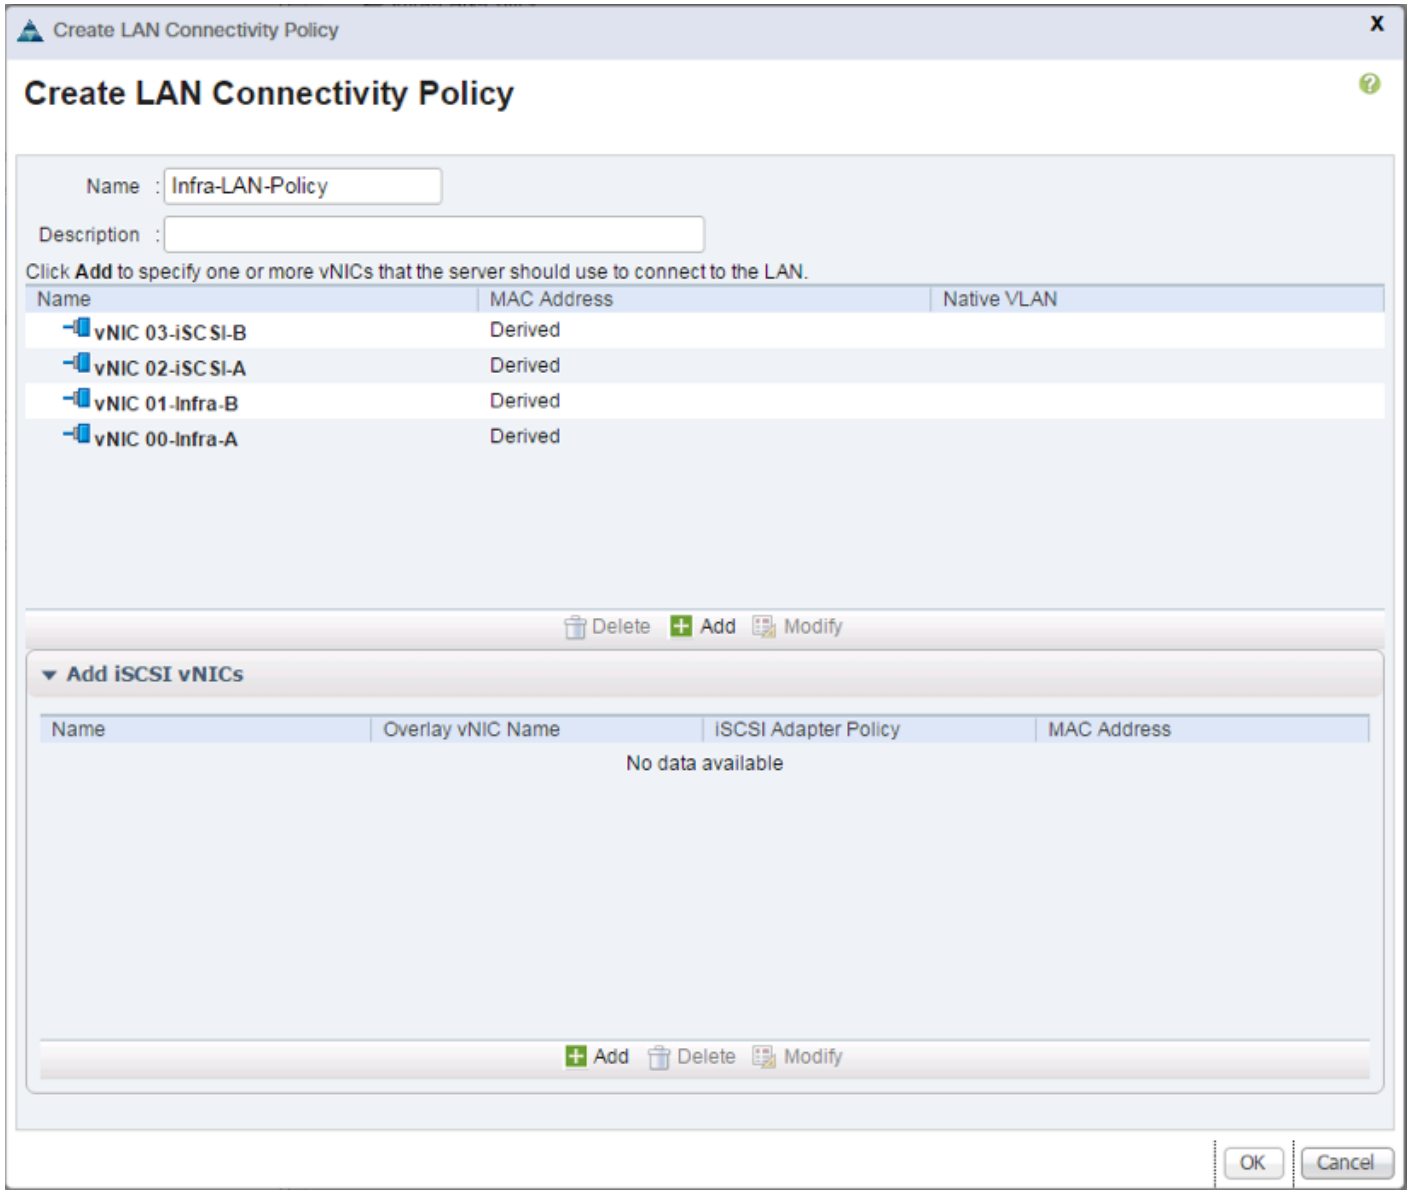

30. Expand the Add iSCSI vNICs section.

31. Click the lower Add button in the iSCSI vNIC section to define an iSCSI boot vNIC.

32. Enter iSCSI-A-vNIC as the name of the vNIC.

33. Select vNIC-02-iSCSI-A for Overlay vNIC.

34. Set the iSCSI Adapter Policy to default.

35. Set the VLAN to Infra-iSCSI-A.

36. Leave the MAC Address set to None.

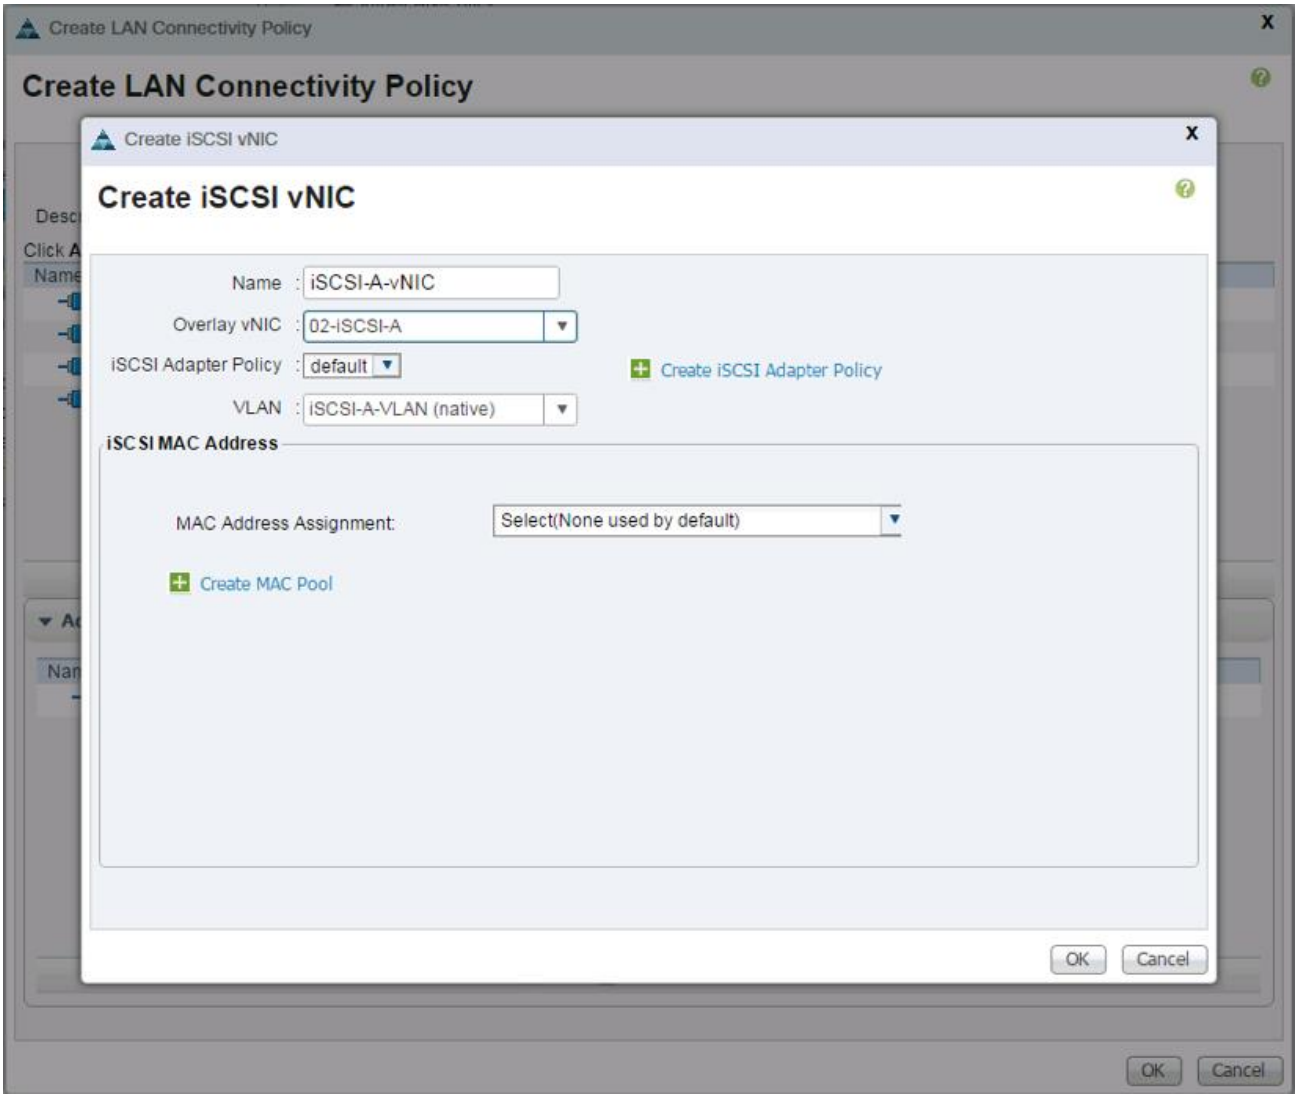

37. Click OK.
38. Click the lower Add button in the iSCSI vNIC section to define an iSCSI boot vNIC.
39. Enter iSCSI-B-vNIC as the name of the vNIC.
40. Set the Overlay vNIC to vNIC-03-iSCSI-B.
41. Set the iSCSI Adapter Policy to default.
42. Set the VLAN to Infra-iSCSI-B.
43. Leave the MAC Address set to None.

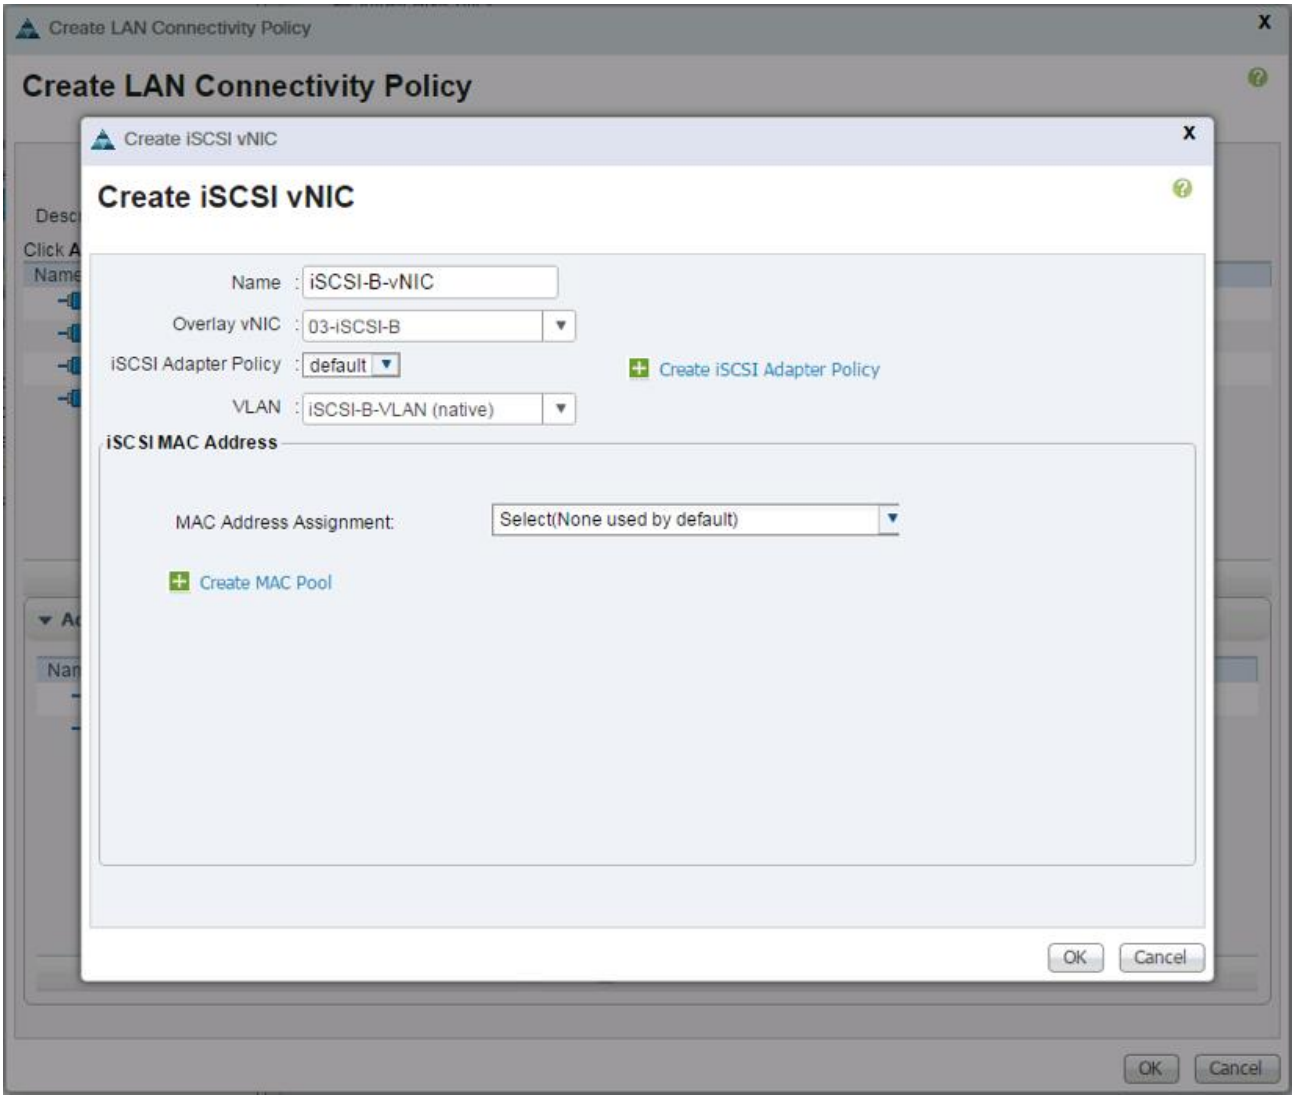
Create LAN Connectivity Policy
Create LAN Connectivity Policy
Create iSCSI vNIC
Create iSCSI vNIC
Name : iSCSI-B-vNIC
Overlay vNIC : 03-iSCSI-B
iSCSI Adapter Policy : default
Create iSCSI Adapter Policy
VLAN : iSCSI-B-VLAN (native)
iSCSI MAC Address
MAC Address Assignment: Select(None used by default)
Create MAC Pool
OK
Cancel
OK
Cancel

44. Click OK.

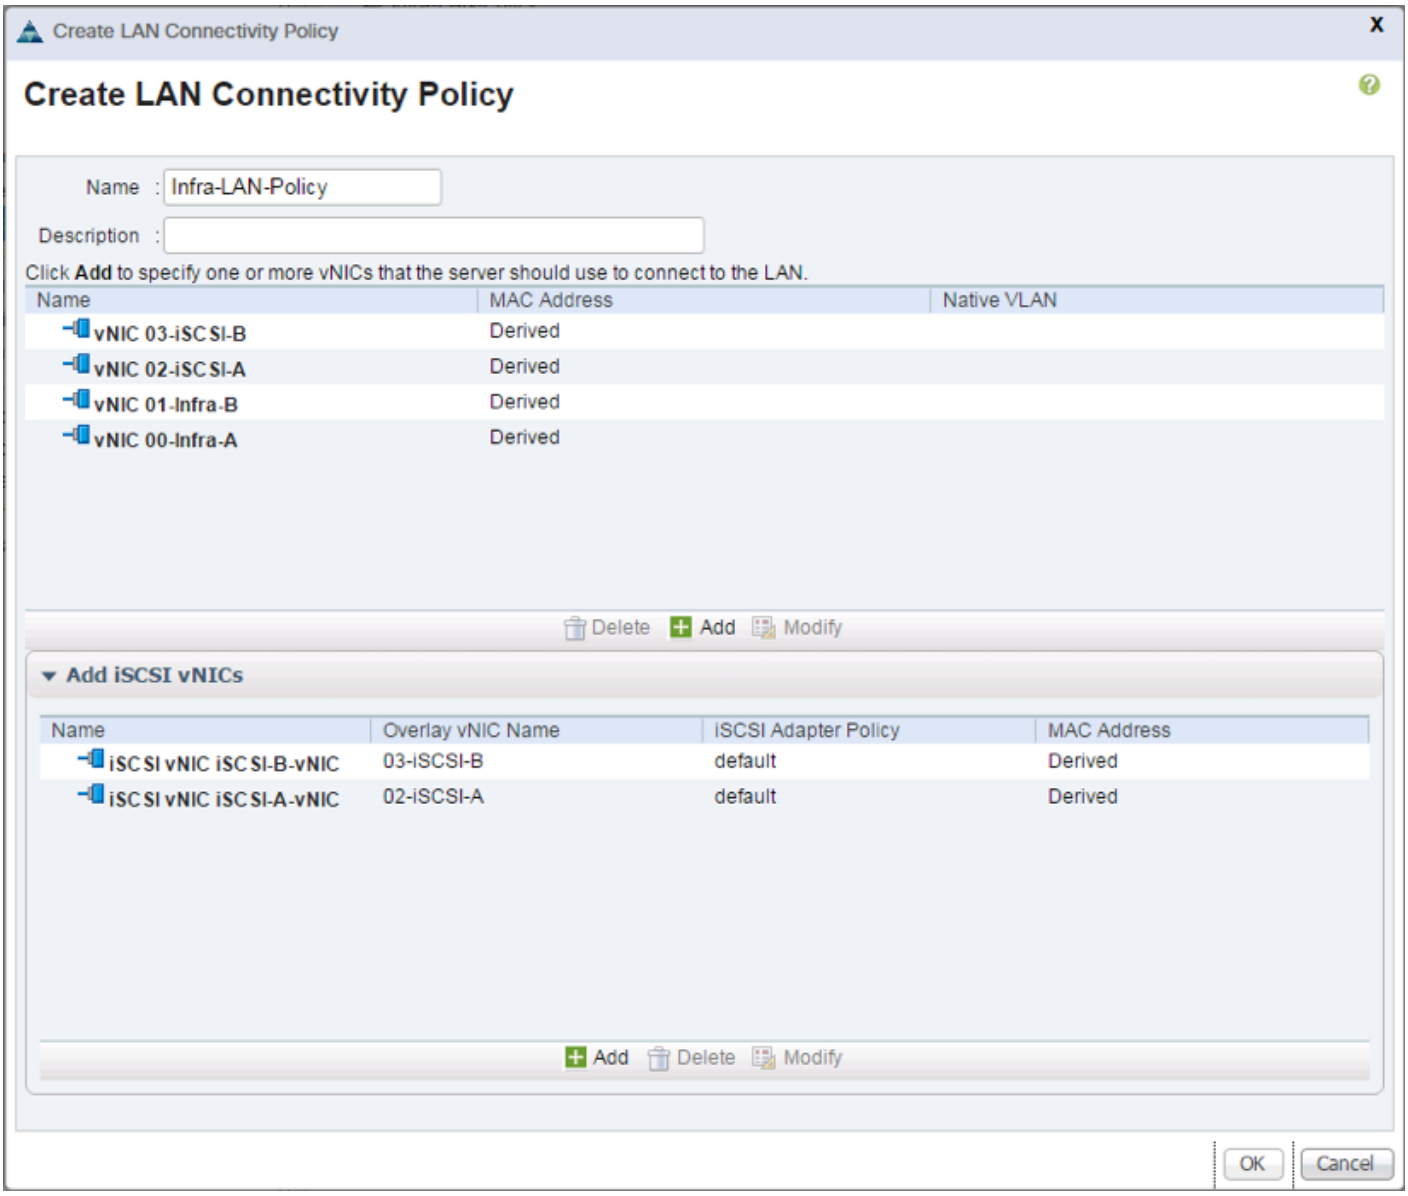

45. Click OK to create the LAN Connectivity Policy.

Create vMedia Policy for VMware ESXi 6.0 U1b Install Boot

In the NetApp Data ONTAP setup steps an HTTP web server is required, which will be used for hosting NetApp Data ONTAP as well as VMware software. The vMedia Policy created here will map the VMware ESXi 6.0u1b ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, complete the following steps:

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

4. Select Create vMedia Policy.
5. Name the policy ESXi-6.0U1b-HTTP.
6. Enter “Mounts Cisco Custom ISO for ESXi 6.0U1b” in the Description field.
7. Click Add.
8. Name the mount ESXi-6.0U1b-HTTP.
9. Select the CDD Device Type.
10. Select the HTTP Protocol.
11. Enter the IP Address of the web server.

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

12. Enter Vmware-ESXi-6.0.0-3380124-Custom-Cisco-6.0.1.2.iso as the Remote File name.
13. Enter the web server path to the ISO file in the Remote Path field.

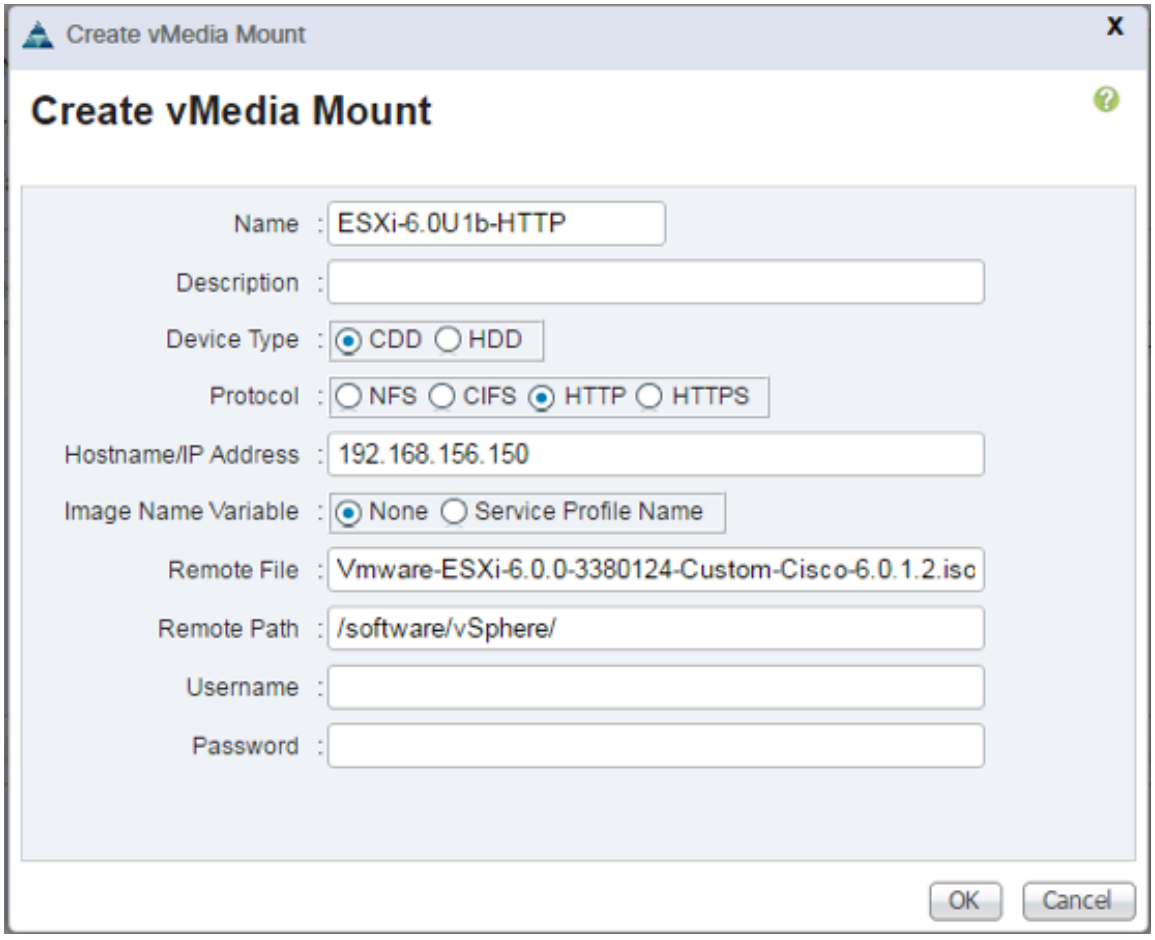

14. Click OK to create the vMedia Mount.

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

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

Create Boot Policies (iSCSI Boot)

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (`iscsi lif01a` and `iscsi lif01b`) and two iSCSI LIFs are on cluster node 2 (`iscsi lif02a` and `iscsi lif02b`). One boot policy is configured in this procedure. This policy configures the primary target to be iscsi_lif01a.

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. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-iSCSI-A` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Add Remote CD/DVD`.
9. Expand the `iSCSI vNICs` section and select `Add iSCSI Boot`.
10. In the `Add iSCSI Boot` dialog box, enter `iSCSI-A-vNIC`.
11. Click OK.
12. Select `Add iSCSI Boot`.
13. In the `Add iSCSI Boot` dialog box, enter `iSCSI-B-vNIC`.
14. Click OK.
15. Expand `CIMC Mounted vMedia`.
16. Select `Add CIMC Mounted CD/DVD`.

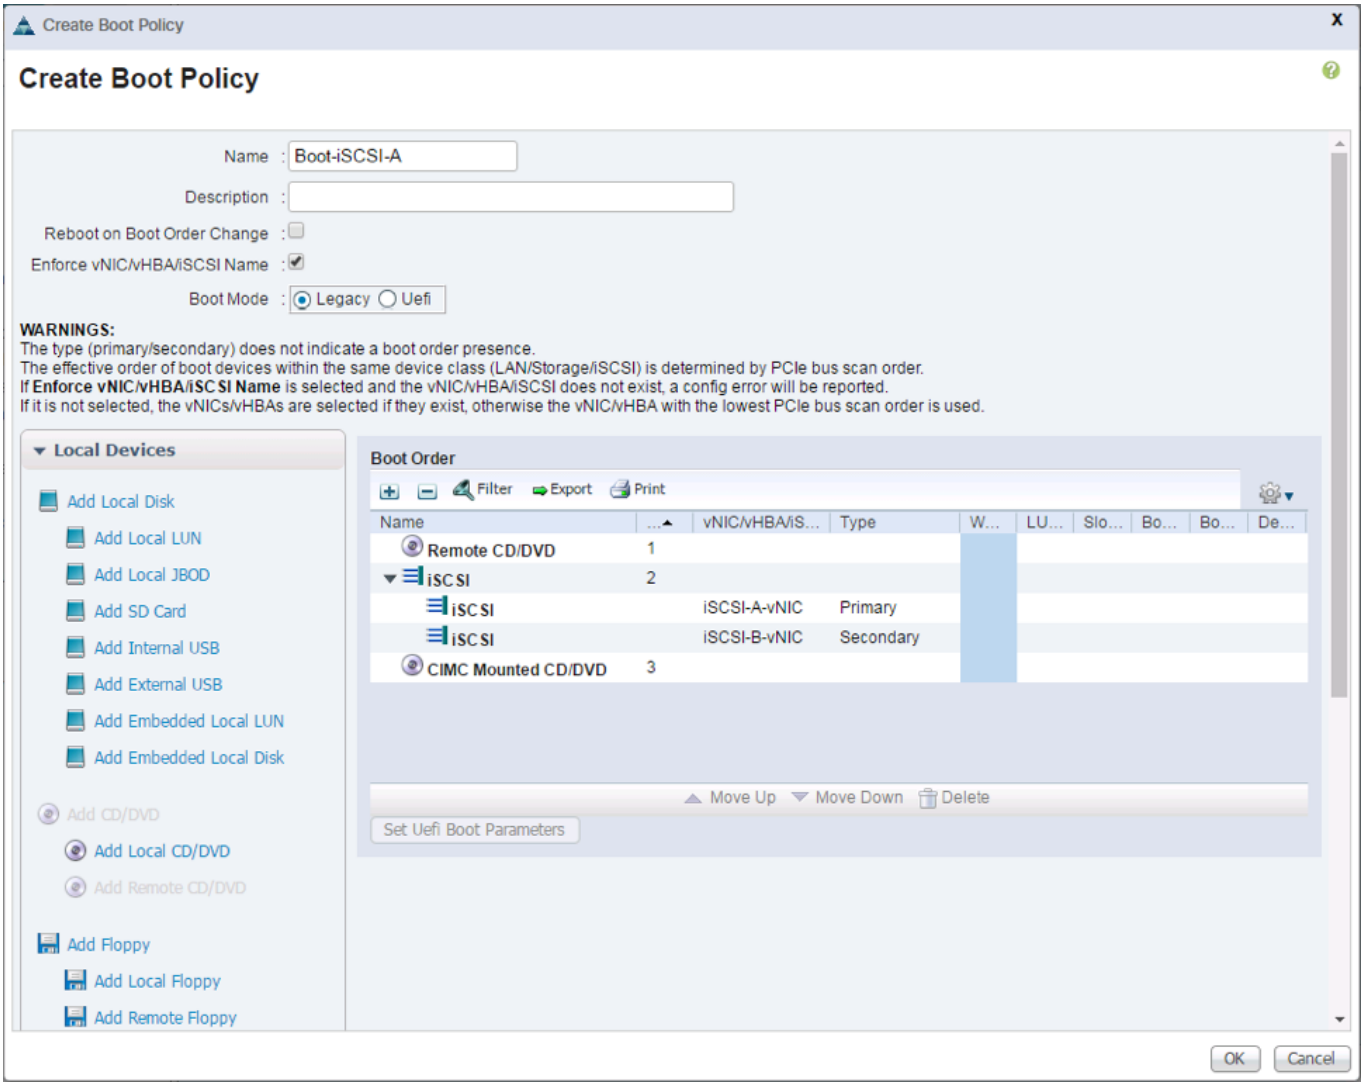

17. Click OK to save the boot policy. Click OK to close the Boot Policy window.

Create Boot Policies (FC Boot)

This procedure applies to a Cisco UCS environment in which two FC logical interfaces (LIFs) are on cluster node 1 and two FC LIFs are on cluster node 2 for each Cisco UCS Fabric Interconnect:

	6332-16UP Fabric A	6332-16UP Fabric B	6248UP Fabric A	6248UP Fabric B
AFF Cluster Node 1 LIF	fcp_lif01_63a (fcp_lif01a)	fcp_lif01_63b (fcp_lif01b)	fcp_lif01_62a (fcp_lif01a)	fcp_lif01_63b (fcp_lif01b)
AFF Cluster Node 2 LIF	fcp_lif02_63a (fcp_lif02a)	fcp_lif02_63b (fcp_lif02b)	fcp_lif02_62a (fcp_lif02a)	fcp_lif02_62a (fcp_lif02b)

Deployments utilizing iSCSI and not FC should ignore these steps.

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. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-FC-A` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.

Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Add Remote CD/DVD`.
9. Expand the vHBAs drop-down menu and select Add SAN Boot.
10. In the Add SAN Boot dialog box, enter `Fabric-A` in the vHBA field.
11. Confirm that Primary is selected for the Type option.

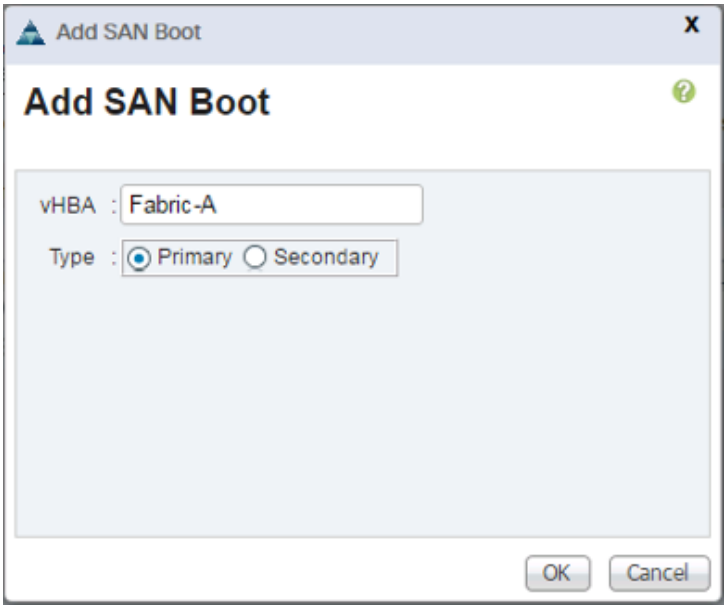

12. Click OK to add the SAN boot initiator.
13. From the vHBA drop-down menu, select Add SAN Boot Target.
14. Keep `0` as the value for Boot Target LUN.

15. Enter the WWPN for `fcp_lif01a`.

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

16. Select Primary for the SAN boot target type.

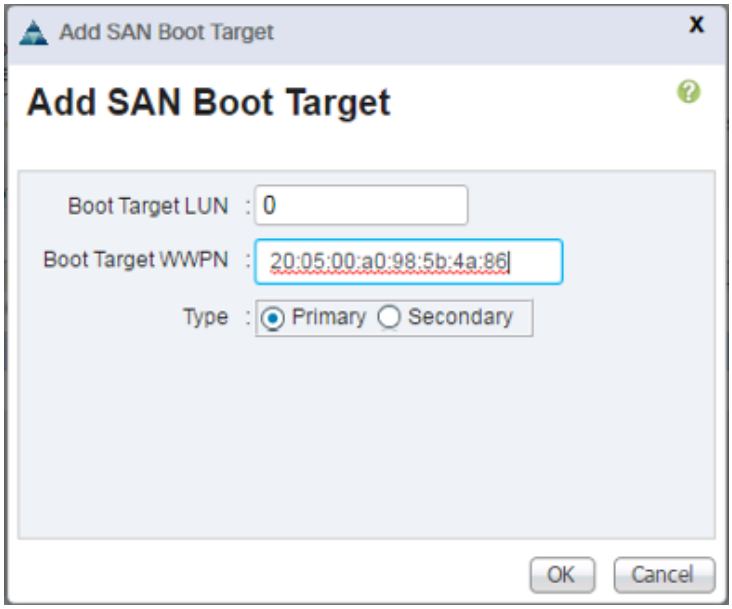

17. Click OK to add the SAN boot target.

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

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

20. Enter the WWPN for `fcp_lif02a`.

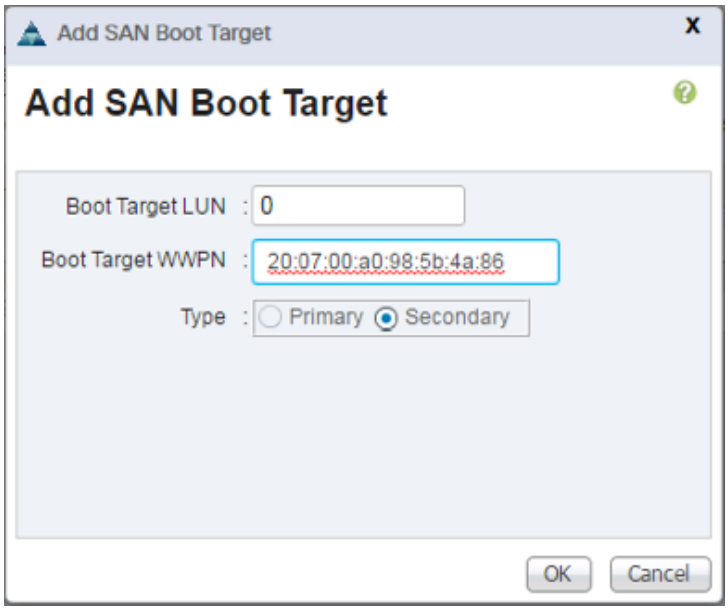

21. Click OK to add the SAN boot target.

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

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

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

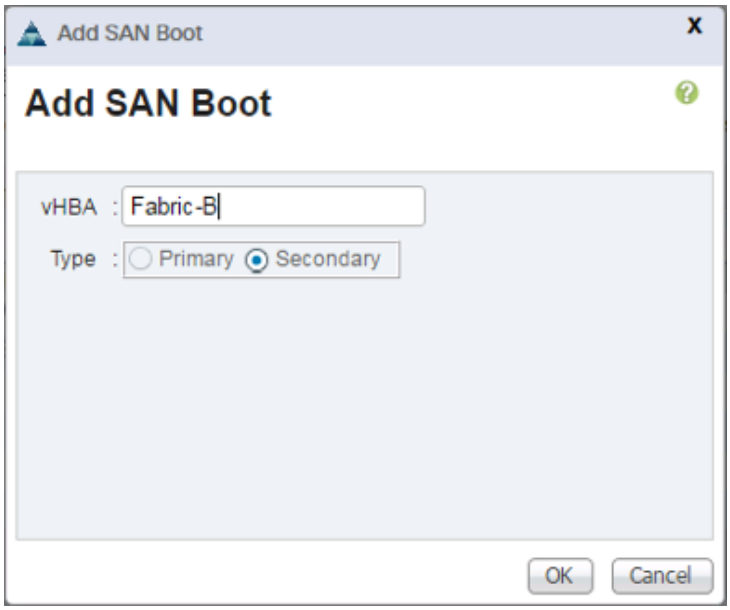

25. Click OK to add the SAN boot initiator.

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

27. Keep `0` as the value for Boot Target LUN.

28. Enter the WWPN for `fcp_lif01b`.

29. Select Primary for the SAN boot target type.

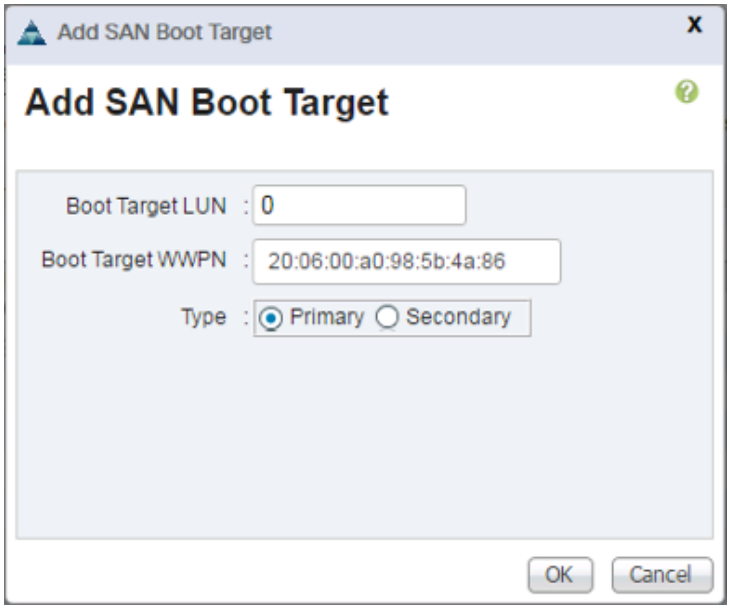

30. Click OK to add the SAN boot target.

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

32. Keep `0` as the value for Boot Target LUN.

33. Enter the WWPN for `fcp_lif02b`.

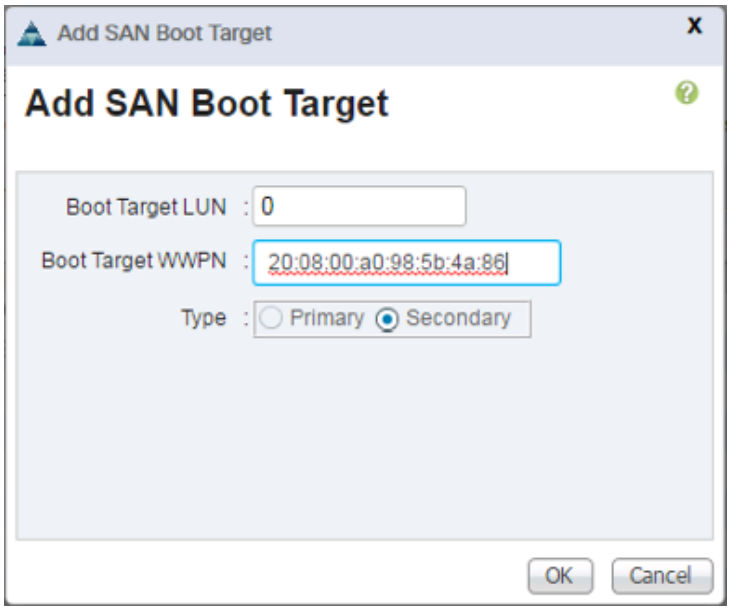

34. Click OK to add the SAN boot target.

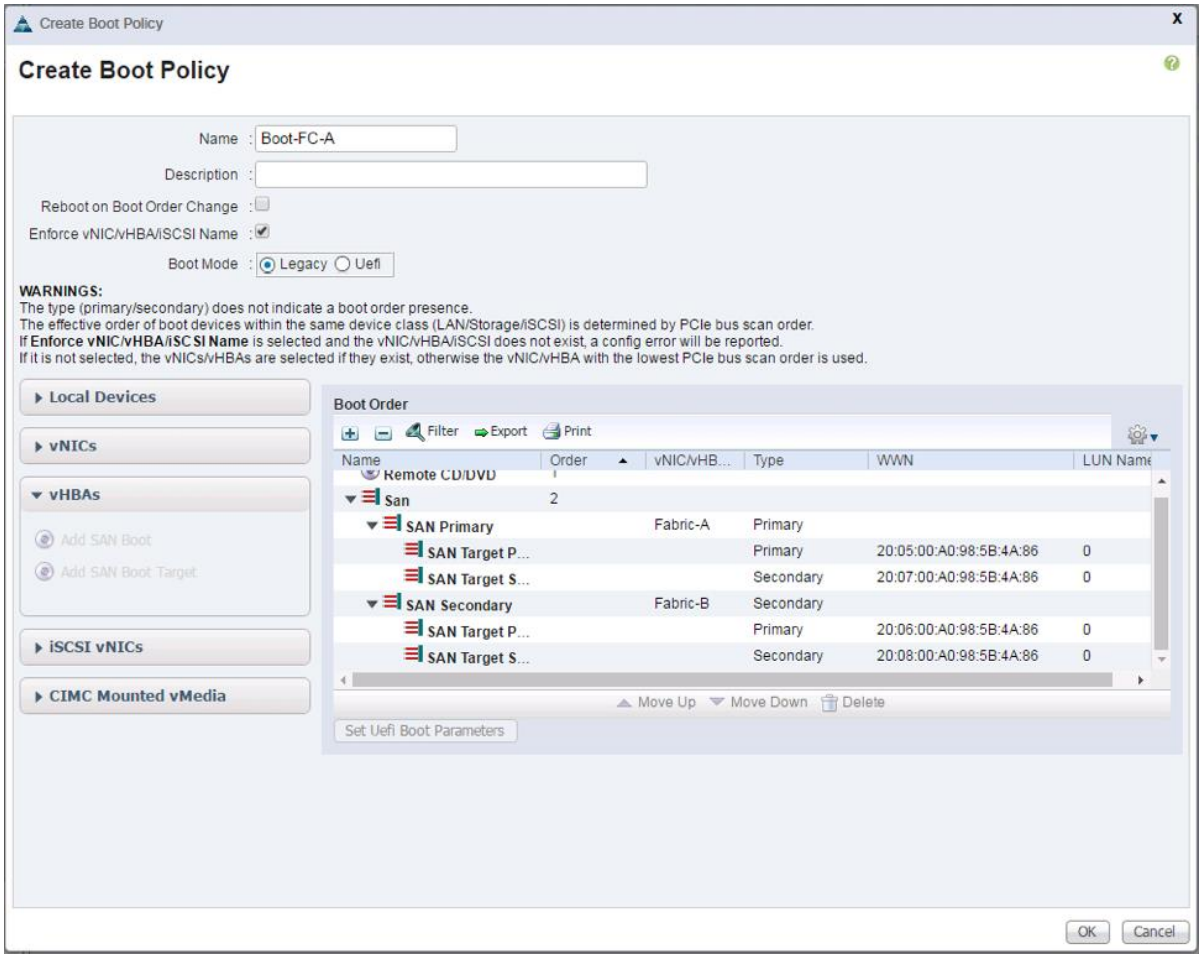

18. Click OK, then click OK again to create the boot policy.

Create Service Profile Template (iSCSI Boot)

Service Profile Templates can be created for Fibre Channel (Fabric) boot or iSCSi boot, with VLANs appropriate infrastructure (Infra) or production (Prod) workloads differentiate by vNIC Templates allowing for VLAN presentation to the interfaces appropriate to the workload. In our example environment, the Service Profile Templates created were:

- VM-Host-Infra-Fabric-A
- VM-Host-Infra-iSCSI-A
- VM-Host-Prod-Fabric-A
- VM-Host-Prod-iSCSI-A

Examples shown in this section are primarily detail one Infra host that was provisioned from a B-Series server on the UCS 6248UP Fabric Interconnect, and one Prod host that was provisioned from a C-Series server on the UCS 6332-16UP Fabric Interconnect.

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

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 VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.

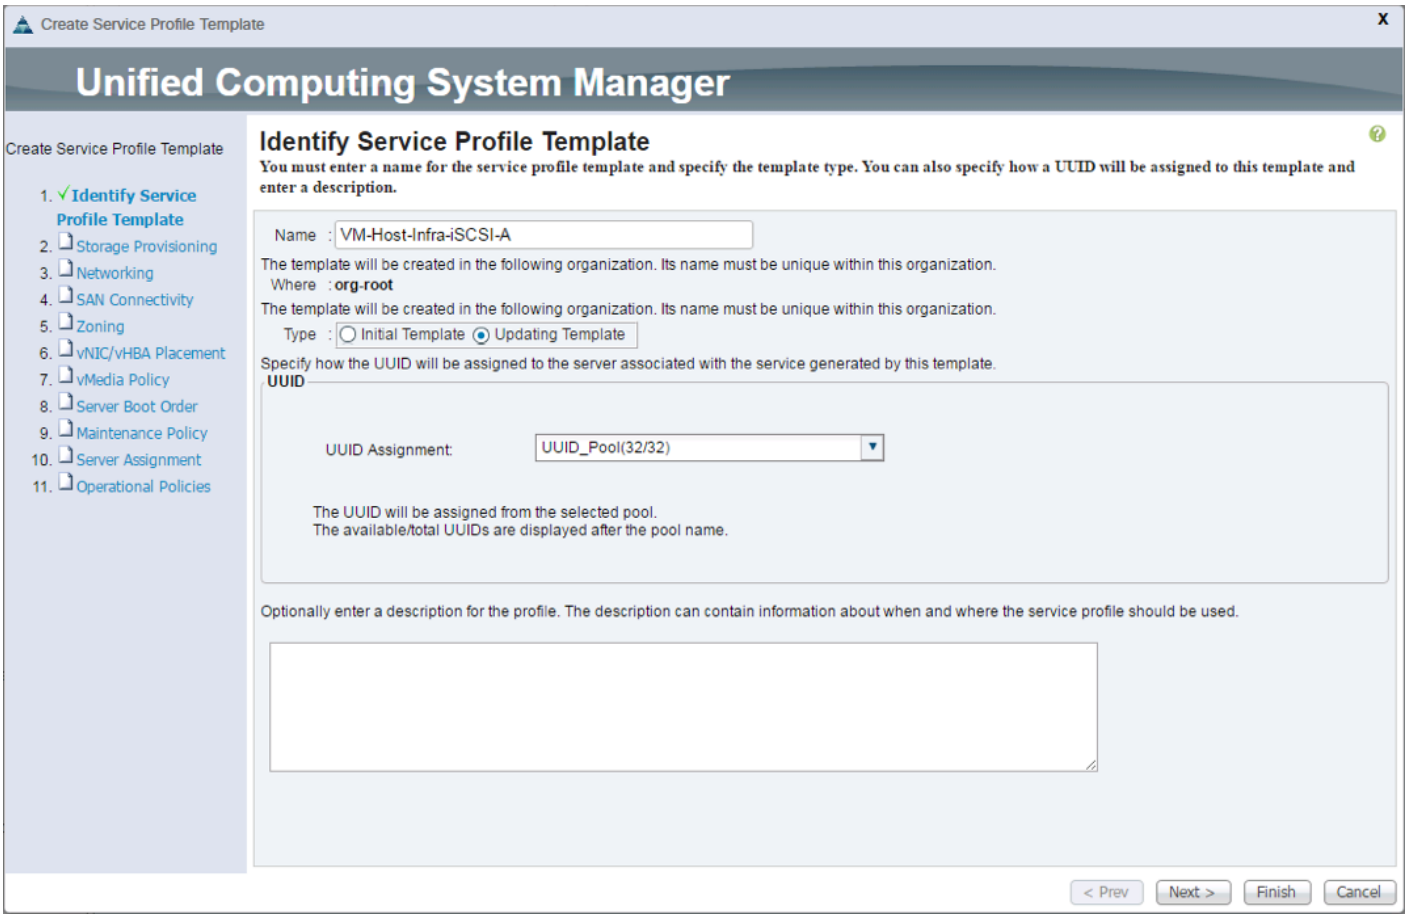

8. Click Next.

Configure Storage Provisioning

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

Configure Networking 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 Infra-LAN-Policy from the LAN Connectivity Policy pull-down.

4. Select IQN_Pool within the Initiator Name Assignment pull-down.

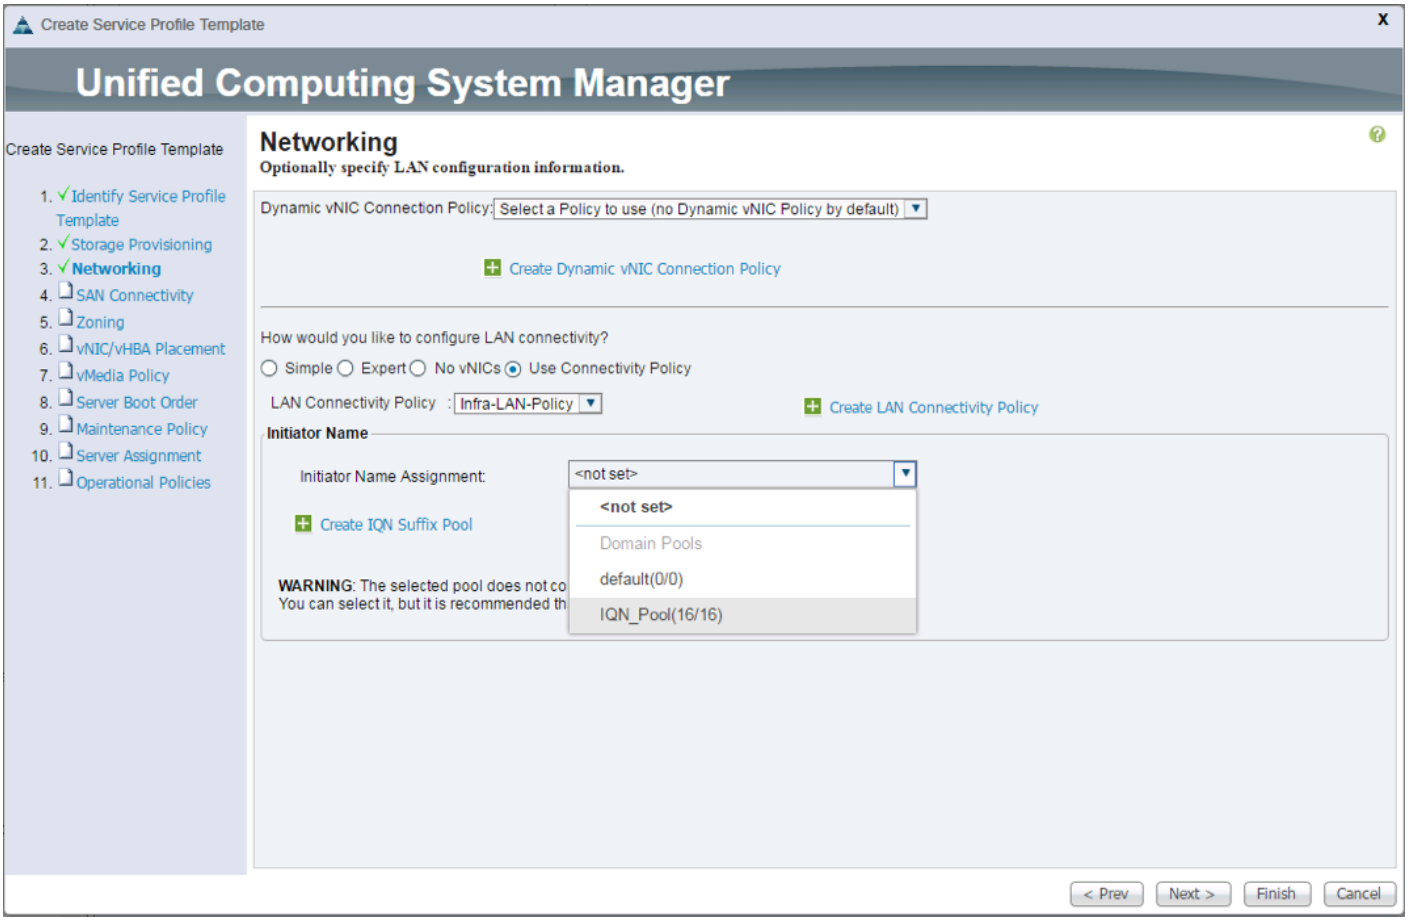

5. Click Next.

Configure Storage Options

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

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

6. Click Next.

Configure Zoning Options

Set no Zoning options 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. From the vMedia Policy pulldown select “ESXi-6.0U1b-HTTP”
2. Click Next.

Configure Server Boot Order

1. Select Boot-iSCSI-A for Boot Policy.
2. In the Boot Order pane, select iSCSI-A-vNIC.
3. Click the “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. Keep the “iSCSI Static Target Interface” button selected and click the  button at the bottom right.

8. Log in to the storage cluster management interface and run the following command:

```
iscsi show
            Target                          Target                        Status
Vserver     Name                            Alias                         Admin
----------  ------------------------------- ----------------------------- ------
Infra-SVM   iqn.1992-08.com.netapp:sn.cbc5f0dff5b911e5aaa600a0985b4a74:vs.3
                                            Infra-SVM                     up
```

9. Note or copy the iSCSI target name for Infra-SVM shown in highlight above.

10. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra-SVM.

11. Enter the IP address of `iSCSI_lif02a` for the IPv4 Address field.

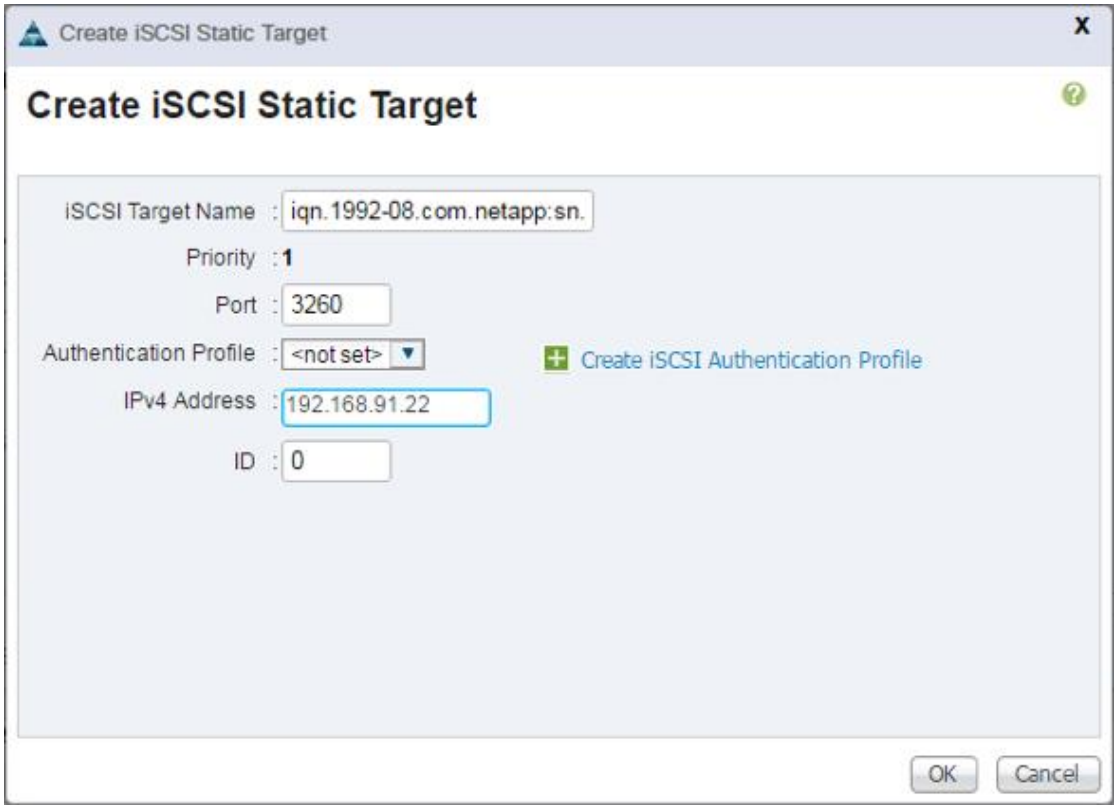

12. Click OK to add the iSCSI static target.

13. Keep the iSCSI Static Target Interface option selected and click the button.

14. In the Create iSCSI Static Target window, paste the iSCSI target node name from `Infra-SVM` into the iSCSI Target Name field.

15. Enter the IP address of `iscsi_lif01a` in the IPv4 Address field.

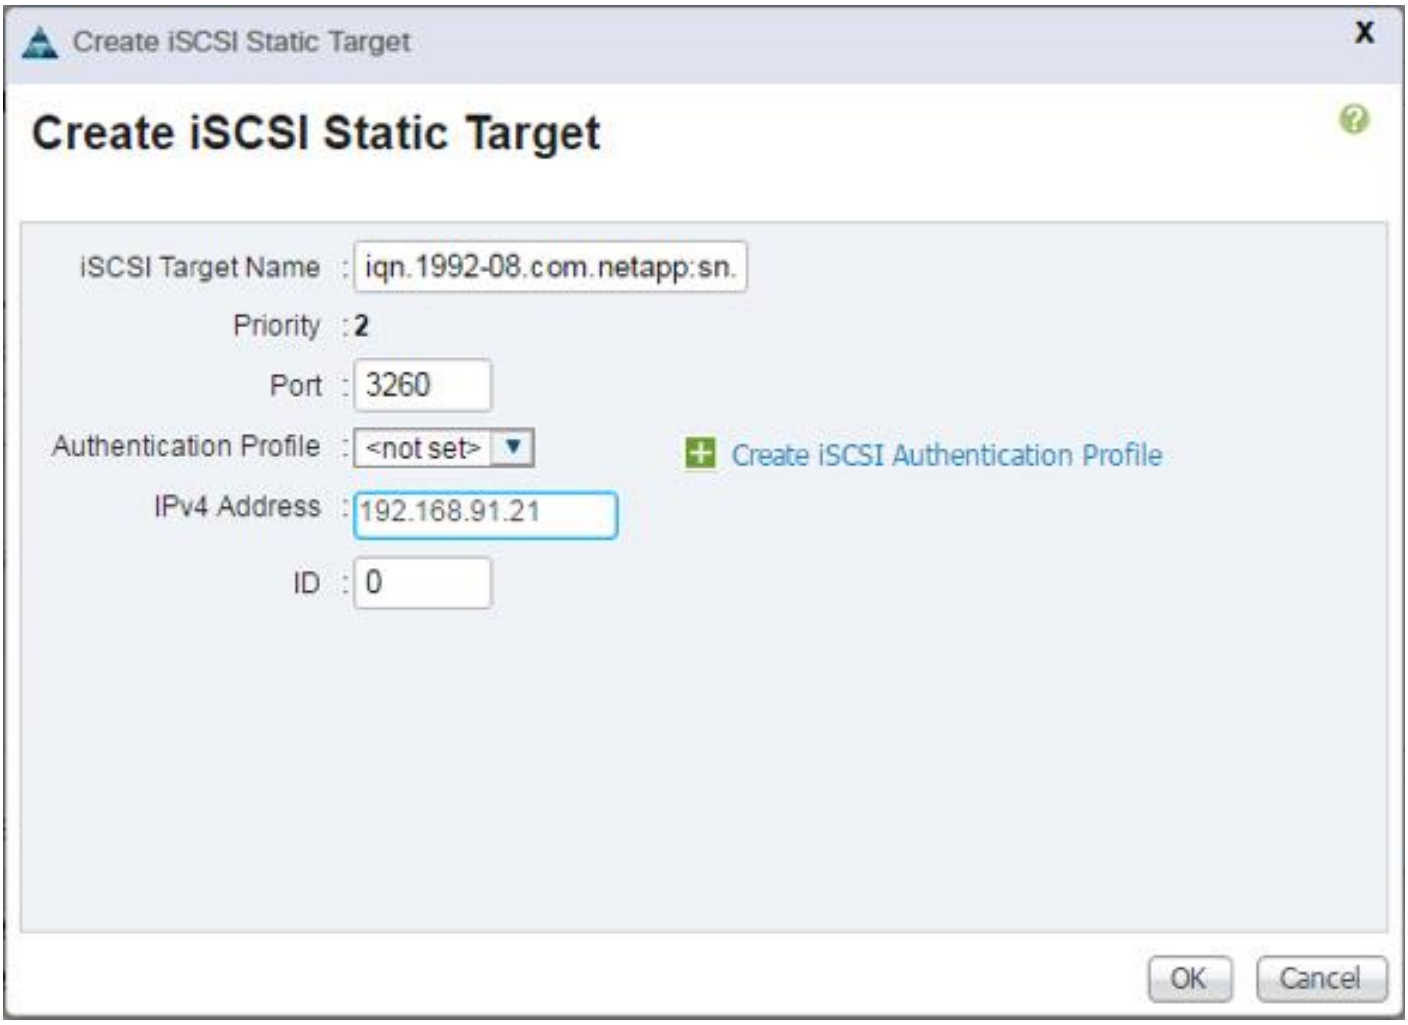

16. Click OK.

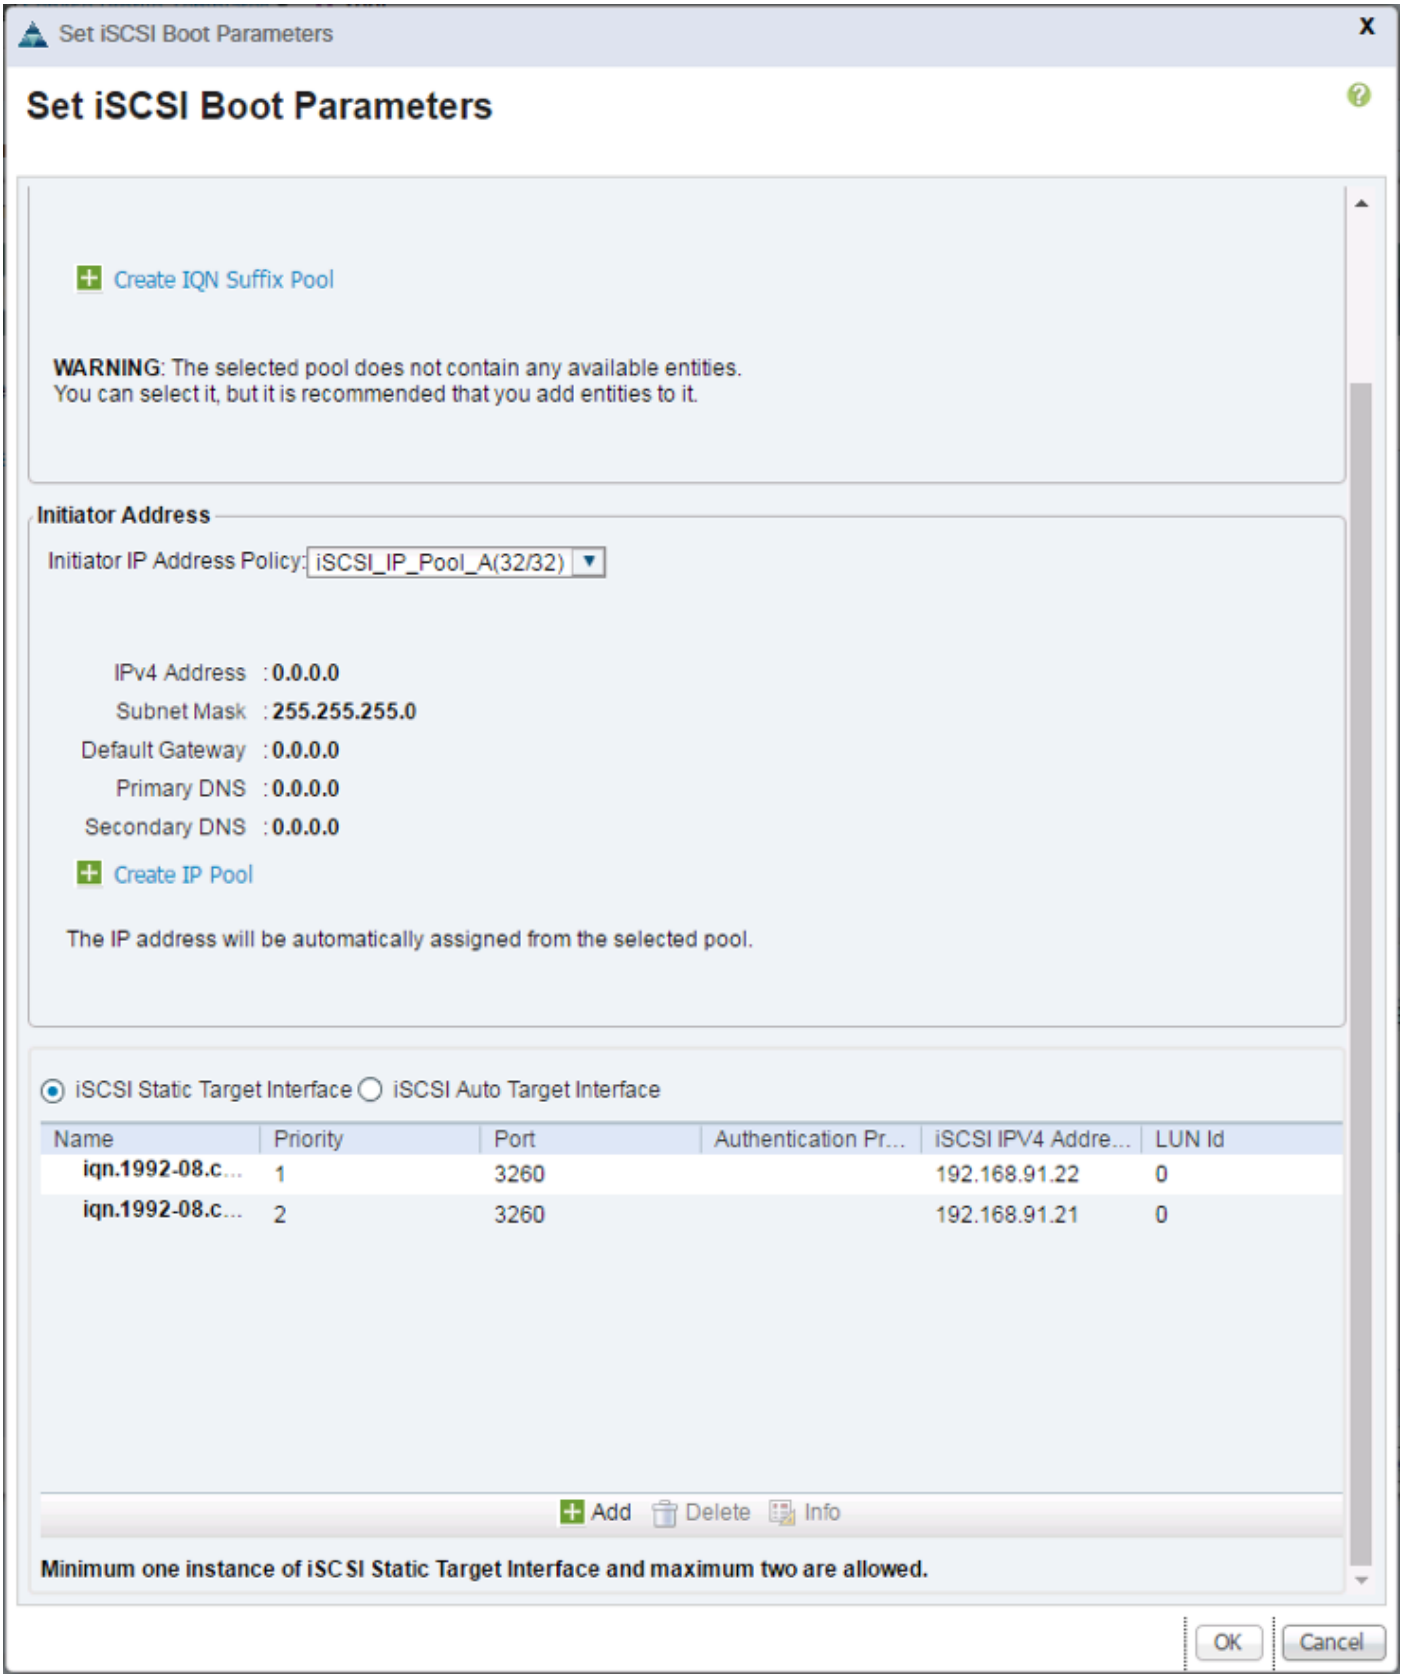
Set iSCSI Boot Parameters
Set iSCSI Boot Parameters
Create IQN Suffix Pool
WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.
Initiator Address
Initiator IP Address Policy: iSCSI_IP_Pool_A(32/32)
IPv4 Address : 0.0.0.0
Subnet Mask : 255.255.255.0
Default Gateway : 0.0.0.0
Primary DNS : 0.0.0.0
Secondary DNS : 0.0.0.0
Create IP Pool
The IP address will be automatically assigned from the selected pool.
iSCSI Static Target Interface
iSCSI Auto Target Interface
Name
Priority
Port
Authentication Pr...
iSCSI IPV4 Addre...
LUN Id
iqn.1992-08.c...
1
3260
192.168.91.22
0
iqn.1992-08.c...
2
3260
192.168.91.21
0
Add
Delete
Info
Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.
OK
Cancel

17. Click OK.

18. In the Boot Order pane, select iSCSI-vNIC-B.
19. Click the Set iSCSI Boot Parameters button.
20. In the Set iSCSI Boot Parameters dialog box, set the leave the “Initiator Name Assignment” to <not set>.
21. In the Set iSCSI Boot Parameters dialog box, set the initiator IP address policy to iSCSI_IP_Pool_B.
22. Keep the iSCSI Static Target Interface option selected and click the button at the bottom right.
23. In the Create iSCSI Static Target window, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field (same target name as above).
24. Enter the IP address of iscsi_lif02b in the IPv4 address field.

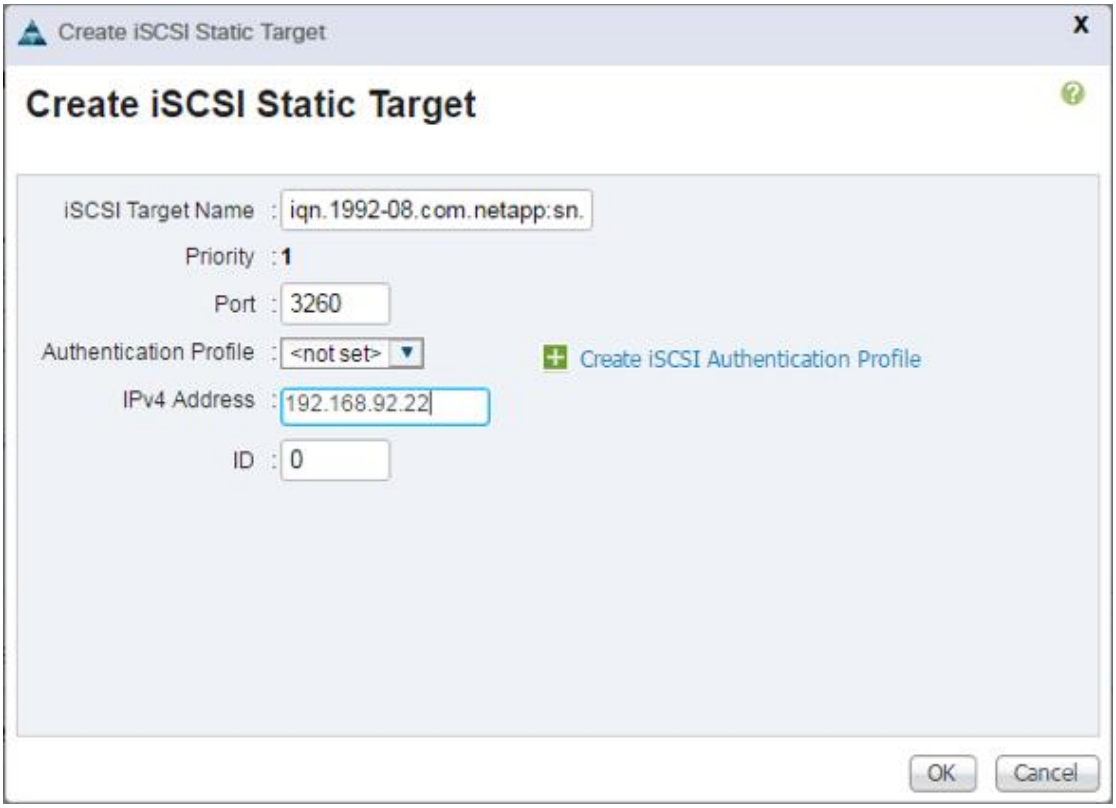

25. Click OK to add the iSCSI static target.
26. Keep the iSCSI Static Target Interface option selected and click the button.
27. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field.
28. Enter the IP address of iscsi_lif01b in the IPv4 Address field.

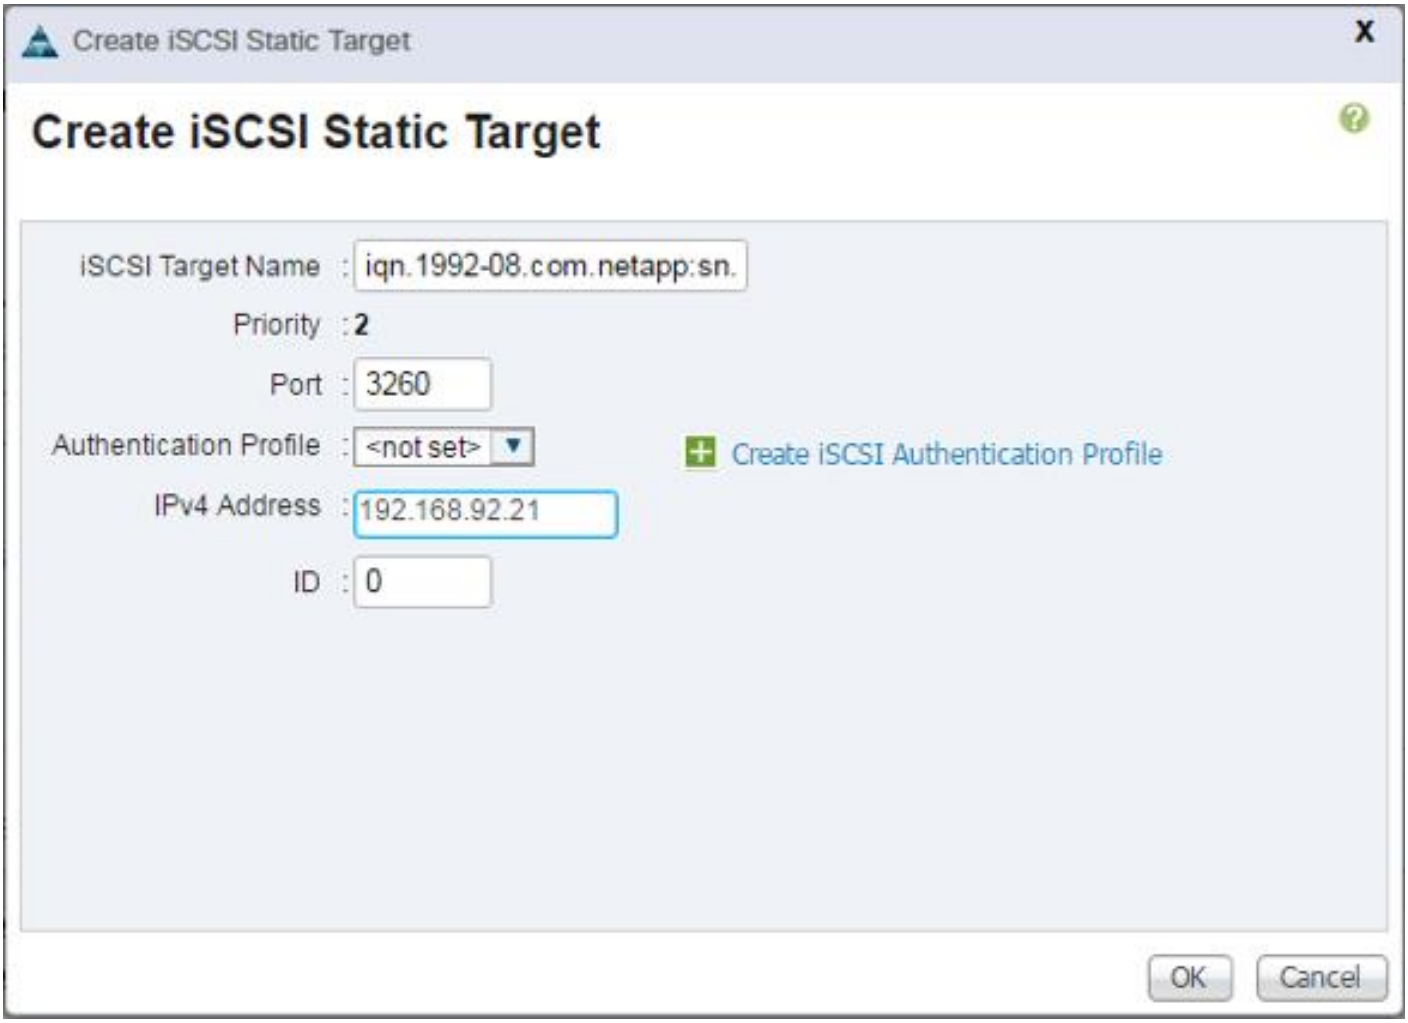
Create iSCSI Static Target
Create iSCSI Static Target
iSCSI Target Name : iqn.1992-08.com.netapp:sn.
Priority : 2
Port : 3260
Authentication Profile : <not set>
Create iSCSI Authentication Profile
IPv4 Address : 192.168.92.21
ID : 0
OK
Cancel

29. Click OK.

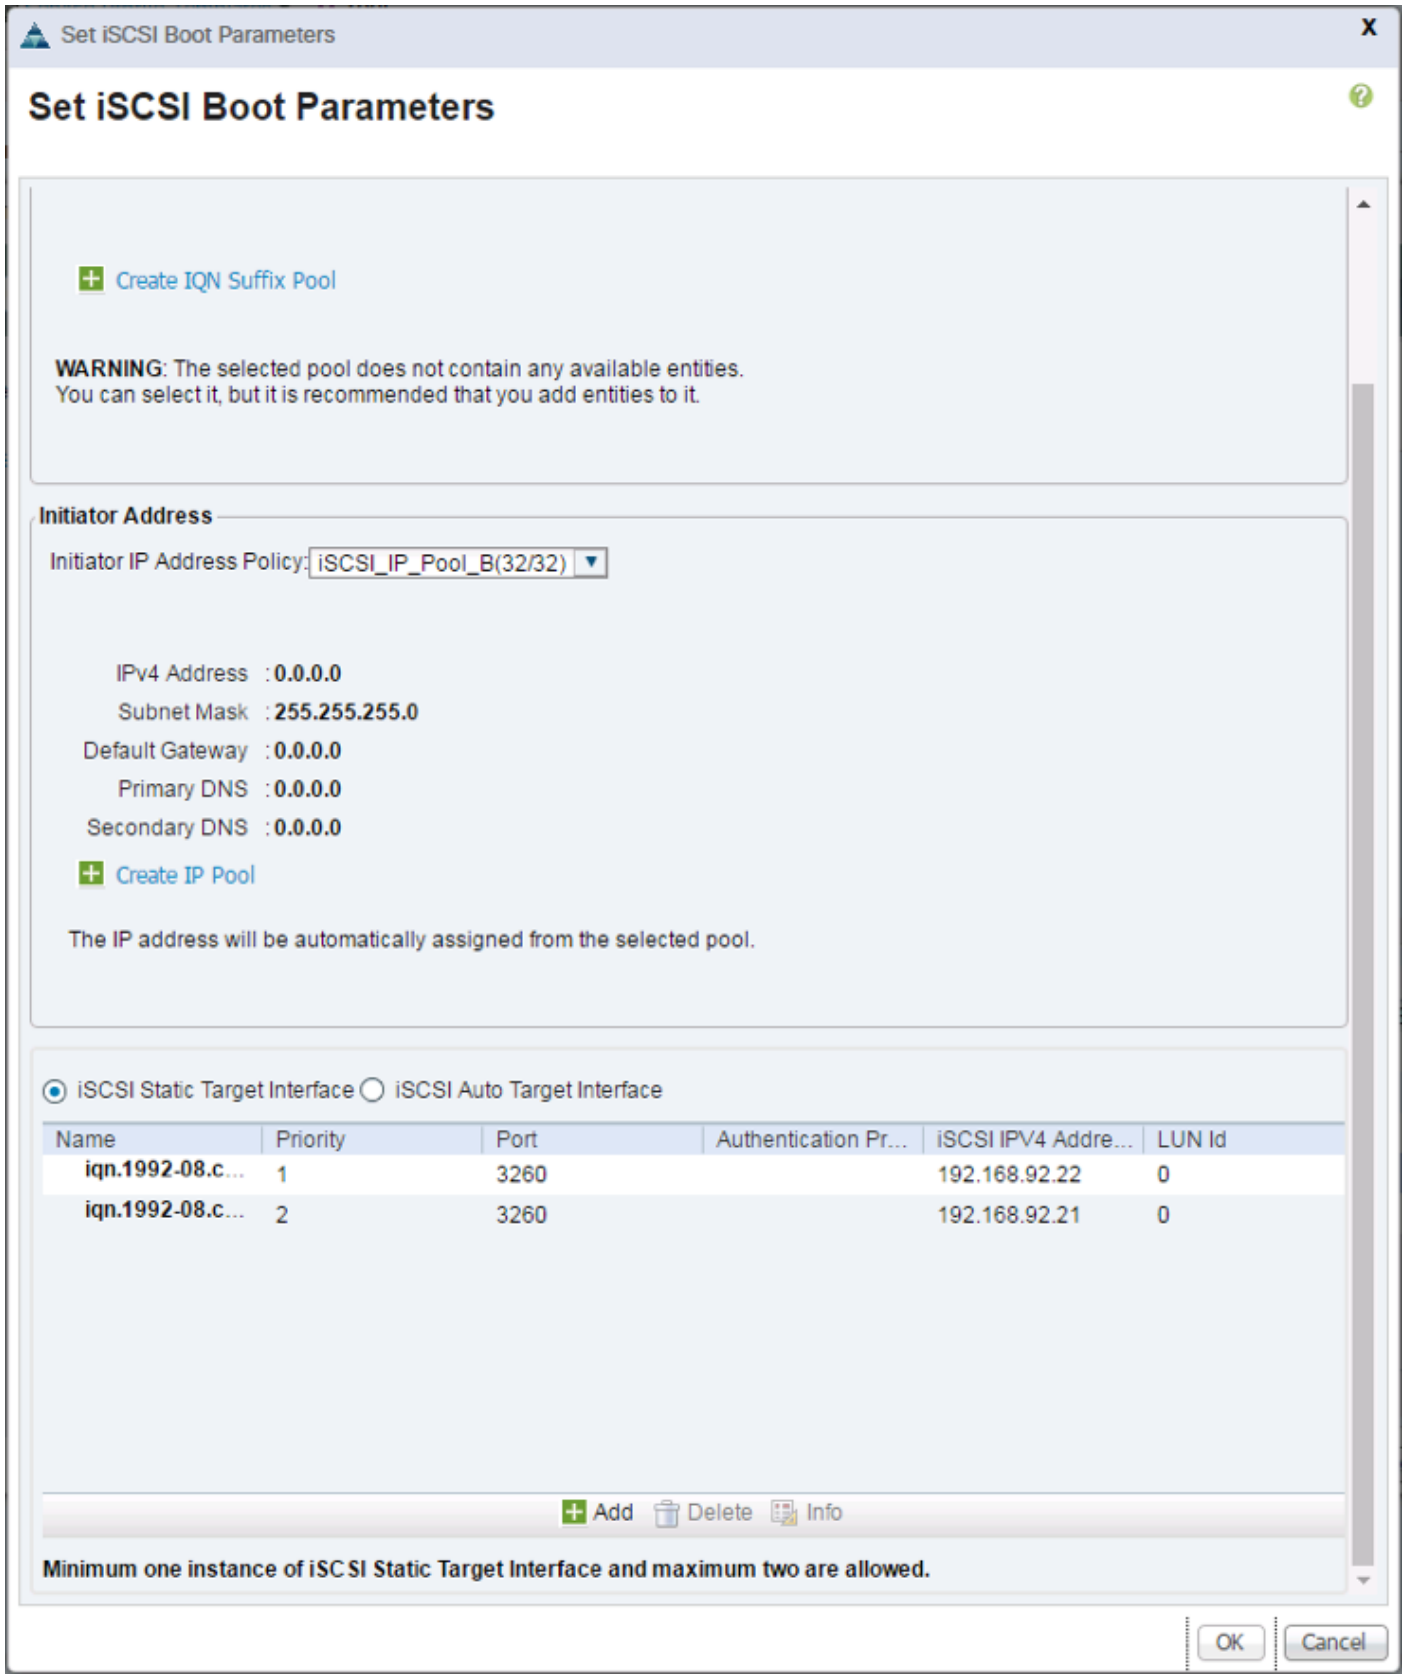
Set iSCSI Boot Parameters
Set iSCSI Boot Parameters
Create IQN Suffix Pool
WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.
Initiator Address
Initiator IP Address Policy: iSCSI_IP_Pool_B(32/32)
IPv4 Address : 0.0.0.0
Subnet Mask : 255.255.255.0
Default Gateway : 0.0.0.0
Primary DNS : 0.0.0.0
Secondary DNS : 0.0.0.0
Create IP Pool
The IP address will be automatically assigned from the selected pool.
iSCSI Static Target Interface
iSCSI Auto Target Interface
Name
Priority
Port
Authentication Pr...
iSCSI IPV4 Addre...
LUN Id
iqn.1992-08.c...
1
3260
192.168.92.22
0
iqn.1992-08.c...
2
3260
192.168.92.21
0
Add
Delete
Info
Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.
OK
Cancel

30. Click OK.

31. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

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_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. Select “UCS-Broadwell” for the Server Pool Qualification.

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

6. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

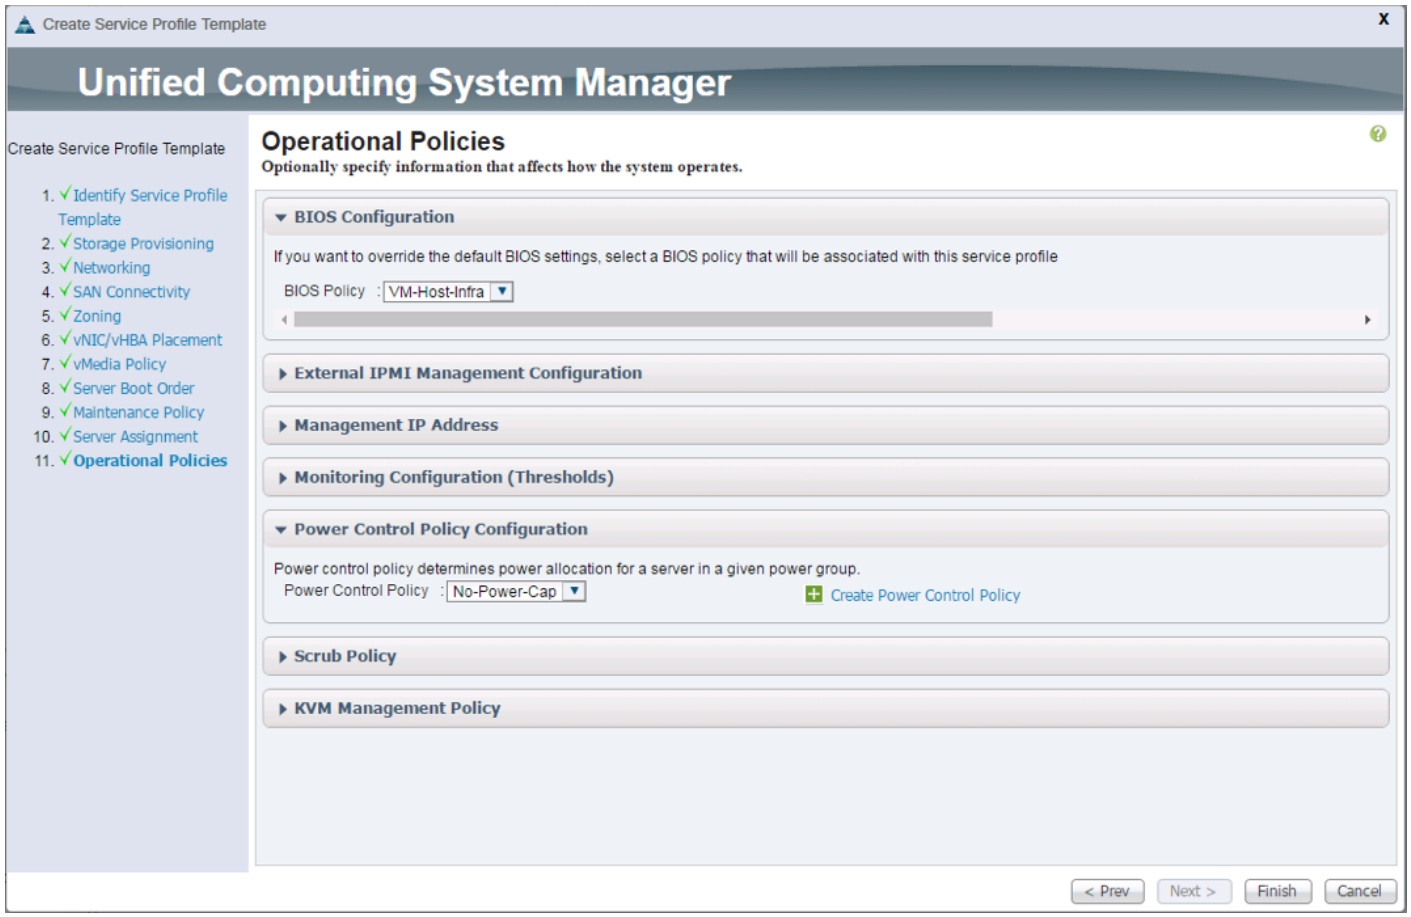

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Service Profile Template (FC Boot)

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

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 VM-Host-Prod-FC-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.

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

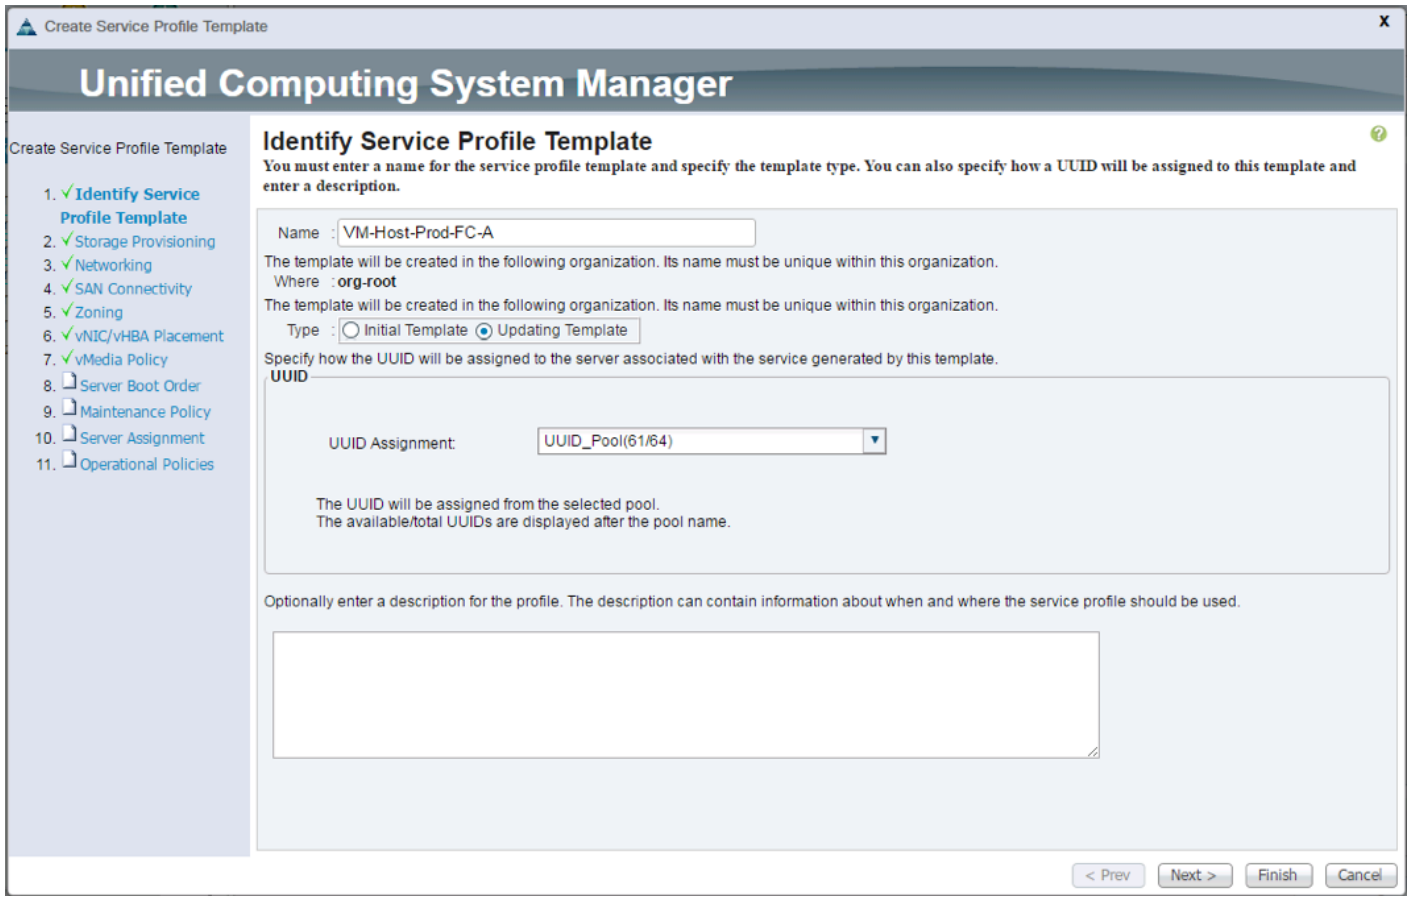

8. Click Next.

Configure Storage Provisioning

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

2. Click Next.

Configure Networking 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 Infra-LAN-Policy from the LAN Connectivity Policy pull-down.

4. Select IQN_Pool within the Initiator Name Assignment pull-down.

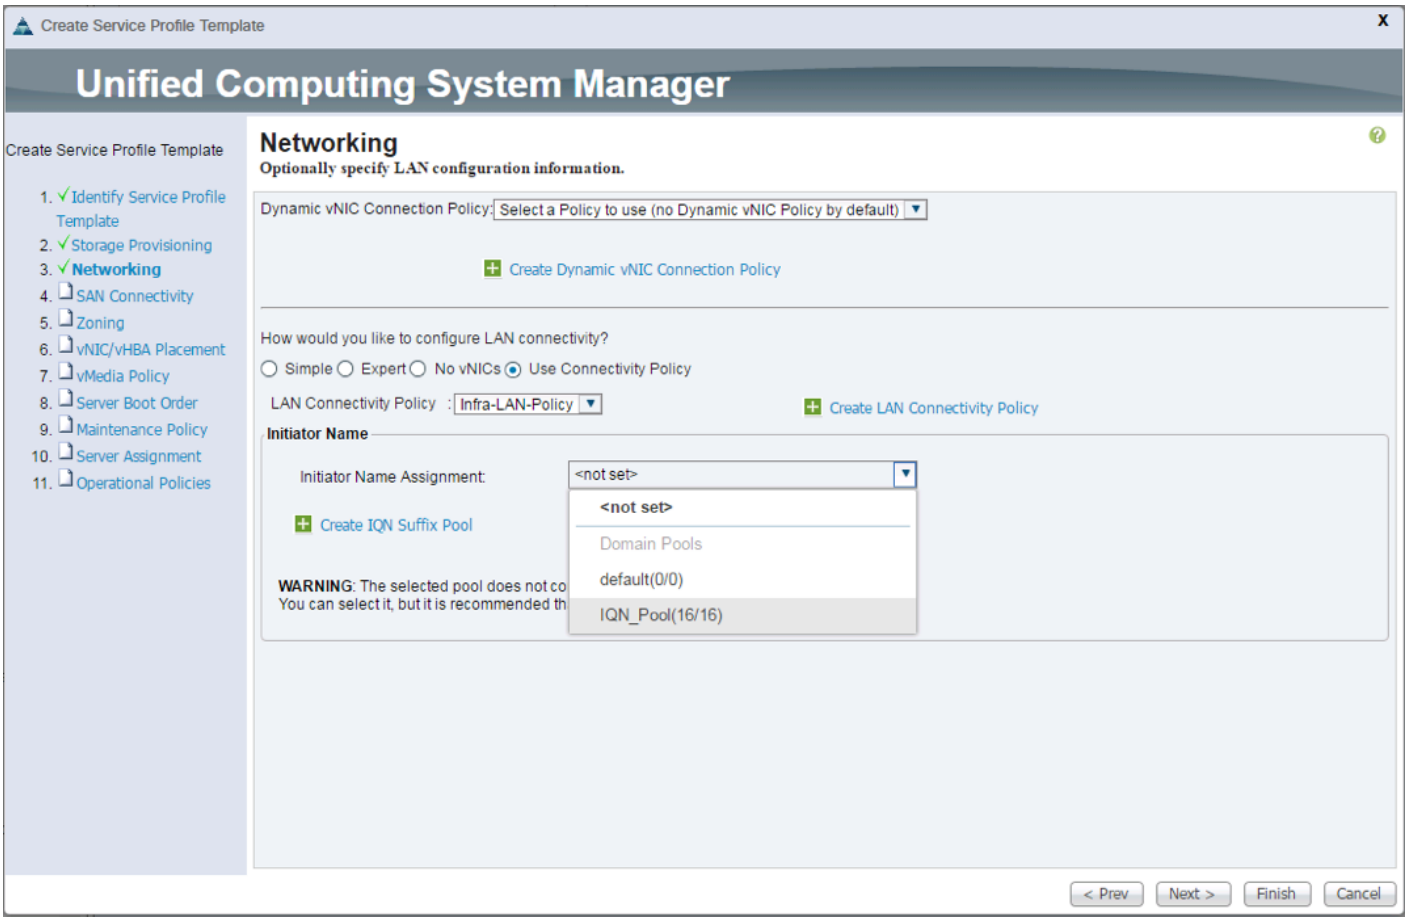

5. Click Next.

Configure Storage Options

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

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

The SAN Connectivity policy created earlier will work for FC or iSCSI environments and should be changed if one of these protocols is not being used.

3. Click Next.

Configure Zoning Options

Set no Zoning options 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. From the vMedia Policy pulldown select “ESXi-6.0U1b-HTTP”
2. Click Next.

Configure Server Boot Order

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

2. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select `Infra_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. Select “UCS-Broadwell” for the Server Pool Qualification.
5. Firmware Management at the bottom of the page can be left alone as it will use default from the Host Firmware list.

6. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

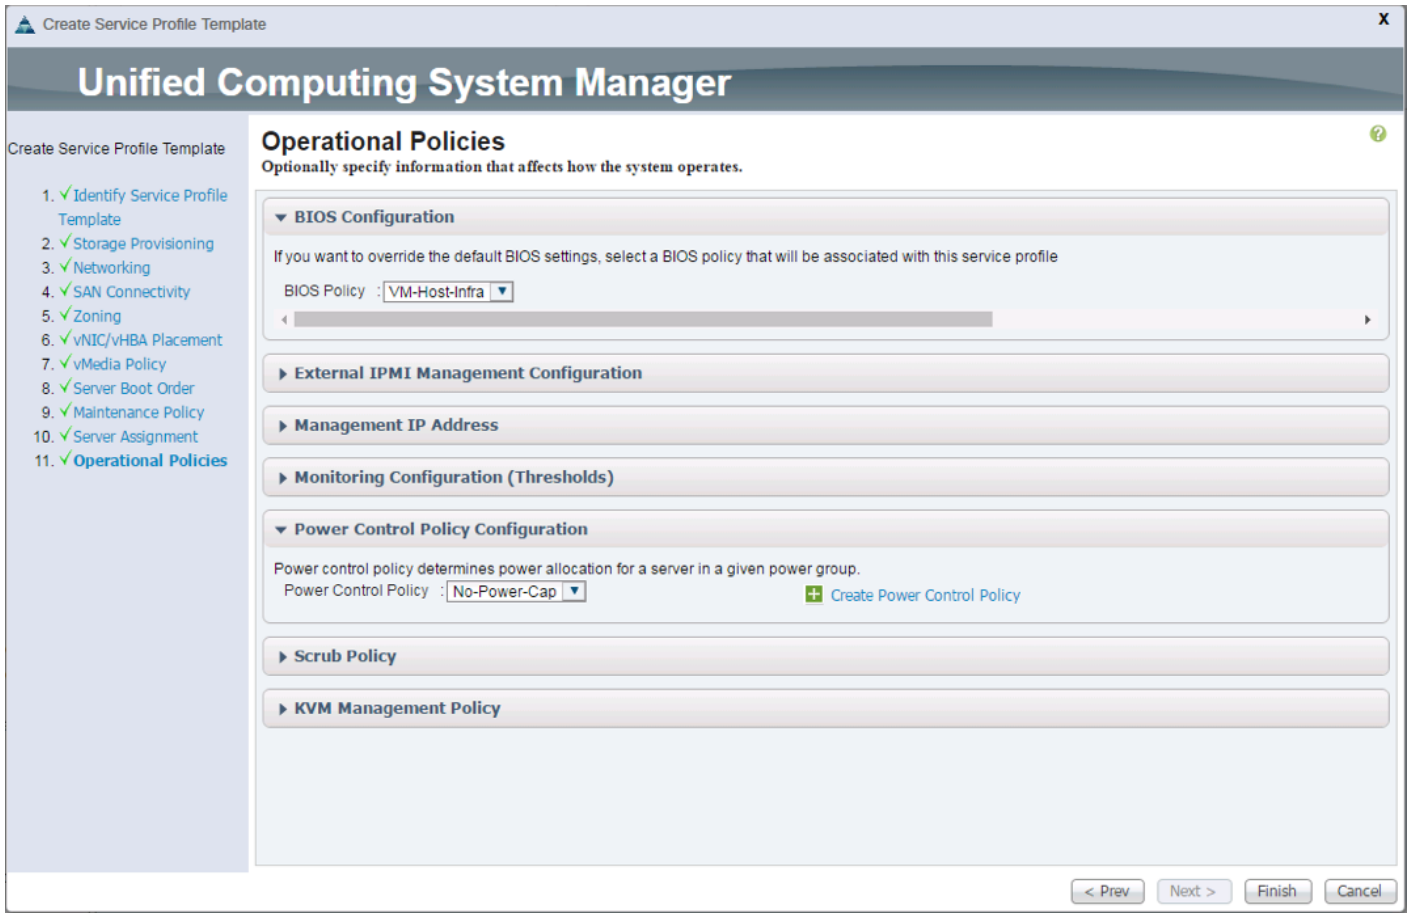

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Service Profiles

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

1. Connect to the UCS 6248UP Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Enter 1 as “Name Suffix Starting Number.”
6. Enter 1 as the “Number of Instances.”
7. Click OK to create the service profile.

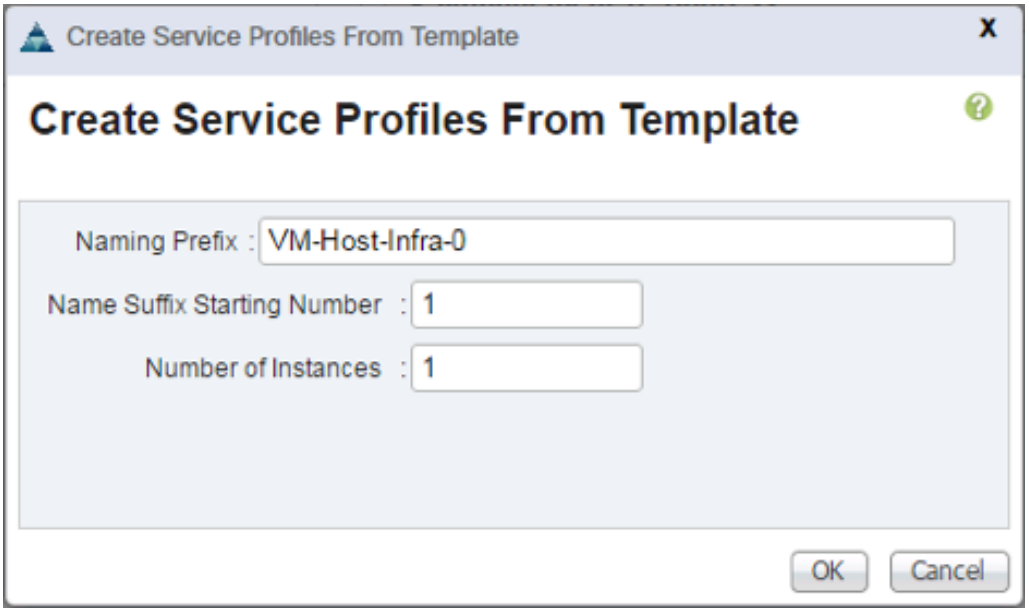

8. Click OK in the confirmation message.
9. Connect to the UCS 6332-16UP Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.
10. Select Service Profile Templates > root > Service Template VM-Host-Prod-Fabric-A.
11. Right-click `VM-Host-Prod-Fabric-A` and select Create Service Profiles from Template.
12. Enter `VM-Host-Prod-0` as the service profile prefix.
13. Enter `1` as “Name Suffix Starting Number.”
14. Enter `1` as the “Number of Instances.”
15. Click OK to create the service profile.

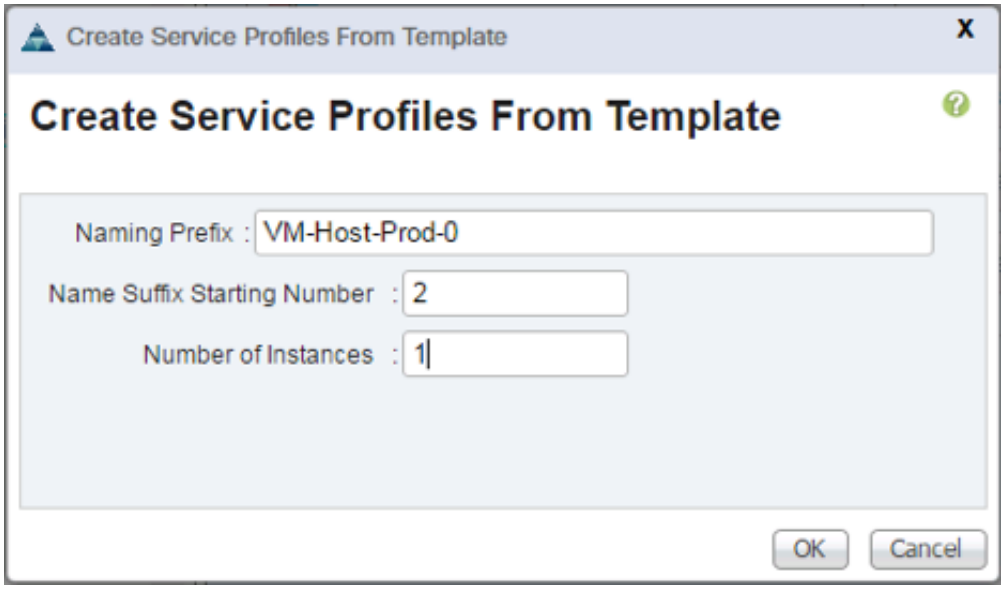

16. Click OK in the confirmation message.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into Table 18 and Table 19 .

Table 18 iSCSI LIFs for iSCSI IQN

Vserver	Target: WWPN (FC), or IQN (iSCSI)
Infra-SVM	

To obtain the FC WWPN, run the `fcp show` command on the storage cluster management interface.

To obtain the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface.

Table 19 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	Initiator: WWPNs (FC) or IQN (iSCSI)	Variables
VM-Host-Infra-01		<<var_vm_host_infra_01_wwpn1>> and <<var_vm_host_infra_01_wwpn2>>; or <<var_vm_host_infra_01_iqn>>

Cisco UCS Service Profile Name	Initiator: WWPNs (FC) or IQN (iSCSI)	Variables
VM-Host-Prod-02		<<var_vm_host_prod_02_wwpn1>> and <<var_vm_host_prod_02_wwpn2>>; or <<var_vm_host_prod_02_iqn>>

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

To obtain the iSCSI vNIC IQN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the “iSCSI vNICs” tab on the right. The “Initiator Name” is displayed at the top of the page under the “Service Profile Initiator Name.”

Storage Configuration – Boot LUNs and Igroups

Clustered Data ONTAP Boot Storage Setup

Create igroups

If you are using FC connectivity, create igroups by entering the following commands from the cluster management node SSH connection:

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol fc –ostype vmware –initiator
<<var_vm_host_infra_01_wwpn1>>, <<var_vm_host_infra_01_wwpn2>>
igroup create –vserver Infra-SVM –igroup VM-Host-Prod-02 –protocol fc –ostype vmware –initiator
<<var_vm_host_prod_02_wwpn1>>, <<var_vm_host_prod_02_wwpn2>>
igroup create –vserver Infra-SVM –igroup MGMT-Hosts –protocol fc –ostype vmware –initiator
<<var_vm_host_infra_01_wwpn1>>, <<var_vm_host_infra_01_wwpn2>>, <<var_vm_host_prod_02_wwpn1>>,
<<var_vm_host_prod_02_wwpn2>>
```

If you are using iSCSI connectivity, create igroups by entering the following commands from the cluster management node SSH connection:

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol iscsi –ostype vmware –initiator
<<var_vm_host_infra_01_iqn>>
igroup create –vserver Infra-SVM –igroup VM-Host-Prod-02 –protocol iscsi –ostype vmware –initiator
<<var_vm_host_prod_02_iqn>>
igroup create –vserver Infra-SVM –igroup MGMT-Hosts –protocol iscsi –ostype vmware –initiator
<<var_vm_host_infra_01_iqn>>, <<var_vm_host_prod_02_iqn>>
```

Use the values listed in Table 18 and Table 19 for the WWPN and IQN information.

To view the three igroups just created, type `igroup show`.

Map Boot LUNs to igroups

From the storage cluster management SSH connection, enter the following commands:

```
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-01 –igroup VM-Host-Infra-01 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Prod-02 –igroup VM-Host-Prod-02 –lun-id 0
```

VMware vSphere 6.0 U1 Setup

VMware ESXi 6.0 U1

This section provides detailed instructions for installing VMware ESXi 6.0 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 6.0 U1

1. Click the following link vmware login page.
2. Type your email or customer number and the password and then click Log in.
3. Click on the following link CiscoCustomImage6.0 U1b.
4. Click Download Now.
5. Save it to your destination folder.

This ESXi 6.0 U1b Cisco custom image includes updates for the fnic and eNIC drivers. The versions that are part of this image are: eNIC: 2.3.0.6; fNIC: 1.6.0.24

Log in to Cisco UCS 6300/6200 Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the 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. To download the Cisco UCS Manager software, click the Launch UCS Manager link.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click the Servers tab.

7. Select Servers > Service Profiles > root > `VM-Host-Infra-01`.
8. Right-click `VM-Host-Infra-01` and select KVM Console.
9. If prompted to accept an Unencrypted KVM session, accept as necessary.
10. Select Servers > Service Profiles > root > `VM-Host-Prod-02`.
11. Right-click `VM-Host-Prod-02`. and select KVM Console.
12. If prompted to accept an Unencrypted KVM session, accept as necessary.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

 Skip this step if using vMedia policies. ISO file will already be connected to KVM.

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click Virtual Media.
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. Boot the server by selecting Boot Server and clicking OK. Then click OK again.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To install VMware ESXi to the iSCSI-bootable or FC-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

4. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, click on the Virtual Media tab and clear the ✓ mark next to the ESXi installation media. Click Yes.

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

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

Set Up Management Networking for ESXi Hosts

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

ESXi Host VM-Host-Infra-01

To configure the `VM-Host-Infra-01` ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the `<<var_ib_mgmt_vlan_id>>` and press Enter.
6. Select Network Adapters option and select vmnic04 and press Enter.
7. From the Configure Management Network menu, select IP Configuration and press Enter.
8. Select the Set Static IP Address and Network Configuration option by using the space bar.
9. Enter the IP address for managing the first ESXi host: `<<var_vm_host_infra_01_ip>>`.
10. Enter the subnet mask for the first ESXi host.
11. Enter the default gateway for the first ESXi host.
12. Press Enter to accept the changes to the IP configuration.

13. Select the IPv6 Configuration option and press Enter.
14. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
15. Select the DNS Configuration option and press Enter.

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

16. Enter the IP address of the primary DNS server.
17. Optional: Enter the IP address of the secondary DNS server.
18. Enter the fully qualified domain name (FQDN) for the first ESXi host.
19. Press Enter to accept the changes to the DNS configuration.
20. Press Esc to exit the Configure Management Network submenu.
21. Press Y to confirm the changes and return to the main menu.
22. The ESXi host reboots. After reboot, press F2 and log back in as root.
23. Select Test Management Network to verify that the management network is set up correctly and press Enter.
24. Press Enter to run the test.
25. Press Enter to exit the window.
26. Press Esc to log out of the VMware console.

ESXi Host VM-Host-Prod-02

To configure the `VM-Host-Prod-02` ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root` and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the `<<var_ib-mgmt_vlan_id>>` and press Enter.
6. Select Network Adapters option and select vmnic4 (defined earlier as OOB vNIC) and press Enter.
7. From the Configure Management Network menu, select IP Configuration and press Enter.
8. Select the Set Static IP Address and Network Configuration option by using the space bar.

9. Enter the IP address for managing the second ESXi host: <<var_vm_host_prod_02_ip>>.
10. Enter the subnet mask for the second ESXi host.
11. Enter the default gateway for the second ESXi host.
12. Press Enter to accept the changes to the IP configuration.
13. Select the IPv6 Configuration option and press Enter.
14. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
15. Select the DNS Configuration option and press Enter.

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

16. Enter the IP address of the primary DNS server.
17. Optional: Enter the IP address of the secondary DNS server.
18. Enter the FQDN for the second ESXi host.
19. Press Enter to accept the changes to the DNS configuration.
20. Press Esc to exit the Configure Management Network submenu.
21. Press Y to confirm the changes and return to the main menu.
22. The ESXi host reboots. After reboot, press F2 and log back in as root.
23. Select Test Management Network to verify that the management network is set up correctly and press Enter.
24. Press Enter to run the test.
25. Press Enter to exit the window.
26. Press Esc to log out of the VMware console.

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 VM-Host-Infra-01 management IP address.
2. Download and install the vSphere Client.

This application is downloaded from the VMware website and Internet access is required on the management workstation.

Log in to VMware ESXi Hosts by Using VMware vSphere Client

ESXi Host VM-Host-Infra-01

To log in to the `VM-Host-Infra-01` ESXi host (which was provisioned from an iSCSI boot Service Profile Template) by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-Infra-01` as the host you are trying to connect to: `<<var_vm_host_infra_01_ip>>`.
2. Enter `root` for the user name.
3. Enter the root password.
4. Click Login to connect.

ESXi Host VM-Host-Prod-02

To log in to the `VM-Host-Prod-02` ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-Prod-02` as the host you are trying to connect to: `<<var_vm_host_prod_02_ip>>`.
2. Enter `root` for the user name.
3. Enter the root password.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

To set up the VMkernel ports and the virtual switches on the `VM-Host-Infra-01`. ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of `vSwitch0`, click Properties.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to `9000`.
7. Click OK.

8. Select the Management Network configuration and click Edit.
9. Change the network label to `VMkernel-MGMT` and select the Management Traffic checkbox.
10. Click OK to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to `MGMT Network` and enter `<<var_ib-mgmt_vlan_id>>` in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.
14. Click Close.
15. On the right side of `iScsiBootvSwitch`, click Properties.
16. Select the vSwitch configuration and click Edit.
17. Change the MTU to 9000.
18. Click OK.
19. Select iScsiBootPG and click Edit.
20. Change the Network Label to `VMkernel-iSCSI-A`.
21. Change the MTU to 9000.
22. Click OK.
23. Click Close.
24. In the vSphere Standard Switch view, click Add Networking.
25. Select VMkernel and click Next.
26. Select Create a vSphere standard switch to create a new vSphere standard switch.
27. Select the check boxes for the network adapter vmnic3.
28. Click Next.
29. Change the network label to `VMkernel-iSCSI-B`.
30. Click Next.
31. Enter the IP address and the subnet mask for the iSCSI VLAN B interface for `VM-Host-Infra-01`.

To obtain the iSCSI IP address information; login to the Cisco UCS Manager, in the servers tab select the corresponding service profiles. In the right pane, click the boot order and select the iSCSI-B-vNIC; click set iSCSI boot parameters; the IP address should appear as the initiator IP address.

32. Click Next.
33. Click Finish.
34. On the right side of `vSwitch1`, click Properties.
35. Select the vSwitch configuration and click Edit.
36. Change the MTU to 9000.
37. Click OK.
38. Select VMkernel-iSCSI-B and click Edit.
39. Change the MTU to 9000.
40. Click OK.
41. Click Close.
42. On the right side of `vSwitch0`, click Properties.
43. Click Add.
44. Change the network label to `VMkernel-NFS` and enter `<<var_nfs_vlan_id>>` in the VLAN ID (Optional) field.
45. Click Next.
46. Enter the IP address `<<var_nfs_vlan_ip_host_01>>` and the subnet mask `<<var_nfs_vlan_ip_mask_host_01>>` for the NFS VLAN interface for `VM-Host-Infra-01`.
47. To continue with the NFS VMkernel creation, click Next.
48. To finalize the creation of the NFS VMkernel interface, click Finish.
49. Select the `VMkernel-NFS` configuration and click Edit.
50. Change the MTU to `9000`.
51. Click OK to finalize the edits for the VMkernel-NFS network.
52. Click Add.
53. Change the network label to `VMkernel-vMotion` and enter `<<var_vmotion_vlan_id>>` in the VLAN ID (Optional) field.

54. Click Next.

55. Enter the IP address <<var_vmotion_vlan_ip_host_01>> and the subnet mask <<var_vmotion_vlan_ip_mask_host_01>> for the vMotion VLAN interface for VM-Host-Infra-01.

56. To continue with the vMotion VMkernel creation, click Next.

57. To finalize the creation of the vMotion VMkernel interface, click Finish.

58. Select the `VMkernel-vMotion` configuration and click Edit.

59. Change the MTU to `9000`.

60. Click OK to finalize the edits for the VMkernel-vMotion network.

61. The properties for vSwitch0 should be similar to the following example:

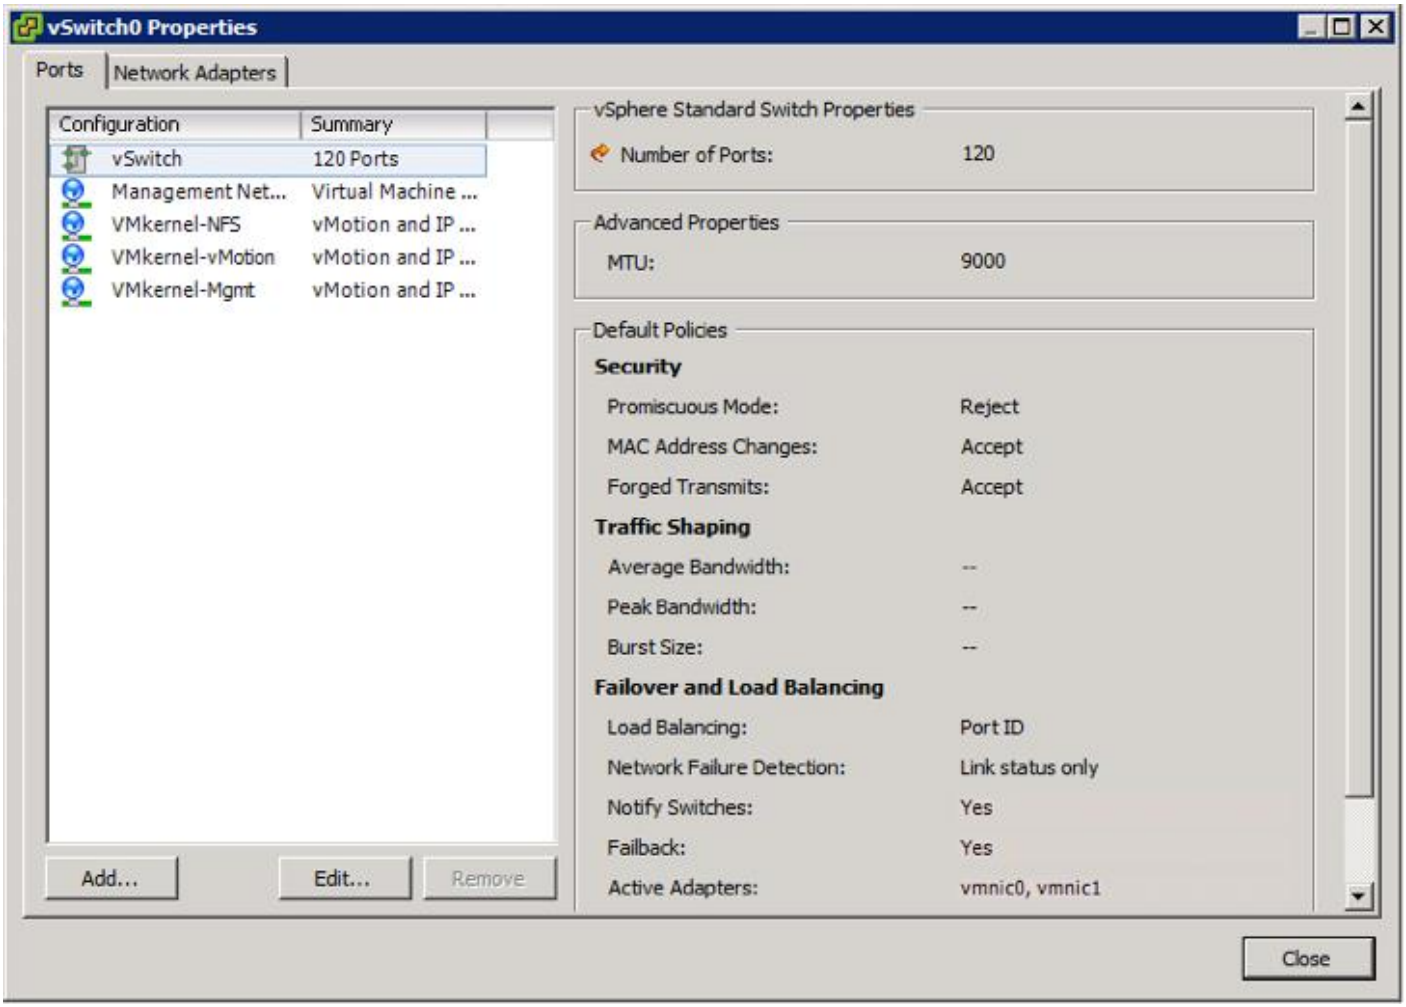

62. To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:

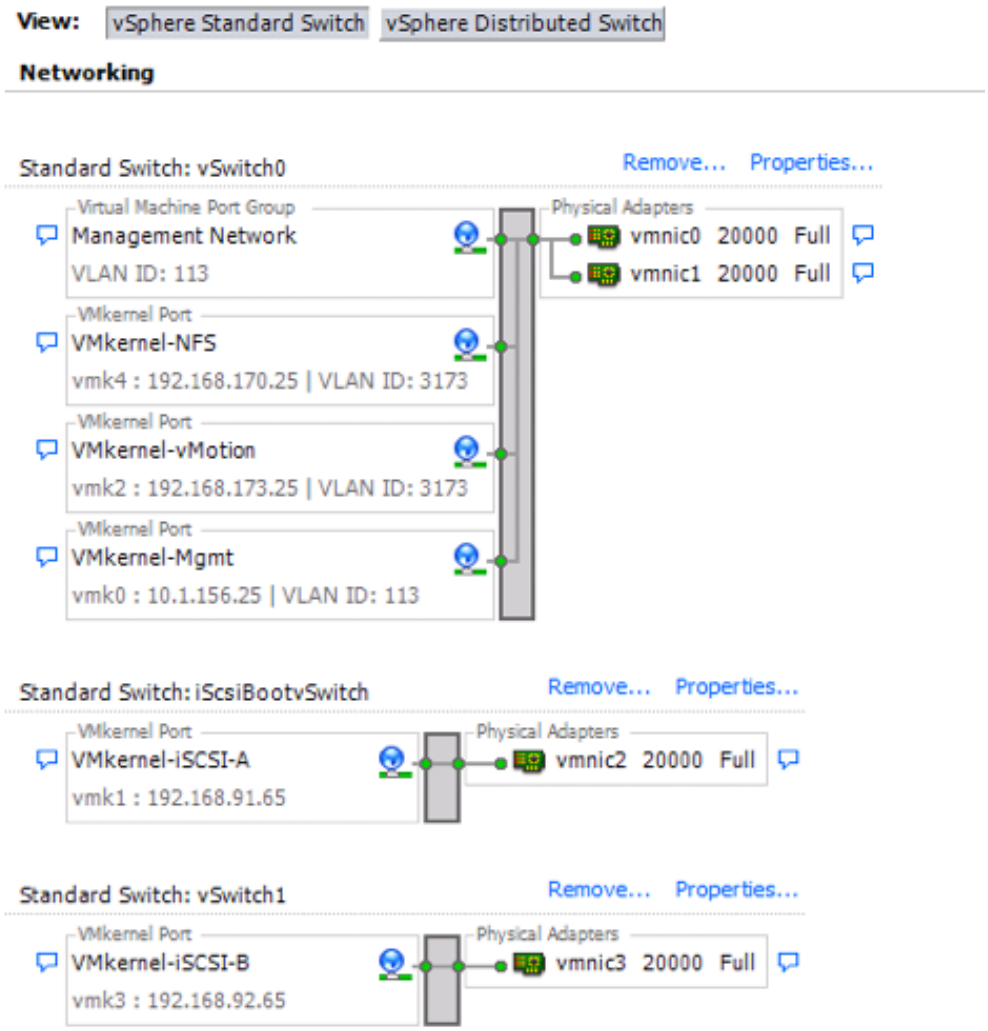

ESXi Host VM-Host-Prod-02

To set up the VMkernel ports and the virtual switches on the `VM-Host-Prod-02, which was provisioned as a Fibre Channel booted` ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of `vSwitch0`, click Properties.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to `9000`.
7. Click OK.

8. Select the Management Network configuration and click Edit.
9. Change the network label to <VMkernel-MGMT> and select the Management Traffic checkbox.
10. Click OK to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to <MGMT Network> and enter <<var_ib-mgmt_vlan_id>> in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.
14. Click Close.
15. In the vSphere Standard Switch view, click Add Networking.
16. Select VMkernel and click Next.
17. Select Create a vSphere standard switch to create a new vSphere standard switch.
18. Select the check boxes for the network adapter vmnic2.
19. Click Next.
20. Change the network label to <VMkernel-iSCSI-A>.
21. Enter the IP address and the subnet mask for the iSCSI VLAN interface for VM-Host-Prod-02.
22. Click Next.
23. Click Finish.
24. On the right side of vSwitch1, click Properties.
25. Change the MTU to 9000.
26. Click OK.
27. In the vSphere Standard Switch view, click Add Networking.
28. Select VMkernel and click Next.
29. Select Create a vSphere standard switch to create a new vSphere standard switch.
30. Select the check boxes for the network adapter vmnic3.
31. Click Next.
32. Change the network label to <VMkernel-iSCSI-B>.

33. Click Next.
34. Enter the IP address and the subnet mask for the iSCSI VLAN interface for `VM-Host-Prod-02`.
35. Click Next.
36. Click Finish.
37. On the right side of `vSwitch2`, click Properties.
38. Select the vSwitch configuration and click Edit.
39. Change the MTU to 9000.
40. Click OK.
41. Select VMkernel-iSCSI-B and click Edit.
42. Change the MTU to 9000.
43. Click OK.
44. Click Close.
45. On the right side of `vSwitch0`, click Properties.
46. Click Add.
47. Select VMkernel and click Next.
48. Change the network label to `VMkernel-NFS` and enter `<<var_nfs_vlan_id>>` in the VLAN ID (Optional) field.
49. Click Next.
50. Enter the IP address <<var_nfs_vlan_ip_host_02>> and the subnet mask <<var_nfs_vlan_ip_mask_host_02>> for the NFS VLAN interface for VM-Host-Prod-02.
51. To continue with the NFS VMkernel creation, click Next.
52. To finalize the creation of the NFS VMkernel interface, click Finish.
53. Select the `VMkernel-NFS` configuration and click Edit.
54. Change the MTU to `9000`.
55. Click OK to finalize the edits for the VMkernel-NFS network.
56. Click Add.

57. Change the network label to `VMkernel-vMotion` and enter `<<var_vmotion_vlan_id>>` in the VLAN ID (Optional) field.

58. Click Next.

59. Enter the IP address <<var_vmotion_vlan_ip_host_02>> and the subnet mask <<var_vmotion_vlan_ip_mask_host_02>> for the vMotion VLAN interface for VM-Host-Prod-02.

60. To continue with the vMotion VMkernel creation, click Next.

61. To finalize the creation of the vMotion VMkernel interface, click Finish.

62. Select the `VMkernel-vMotion` configuration and click Edit.

63. Change the MTU to `9000`.

64. Click OK to finalize the edits for the VMkernel-vMotion network.

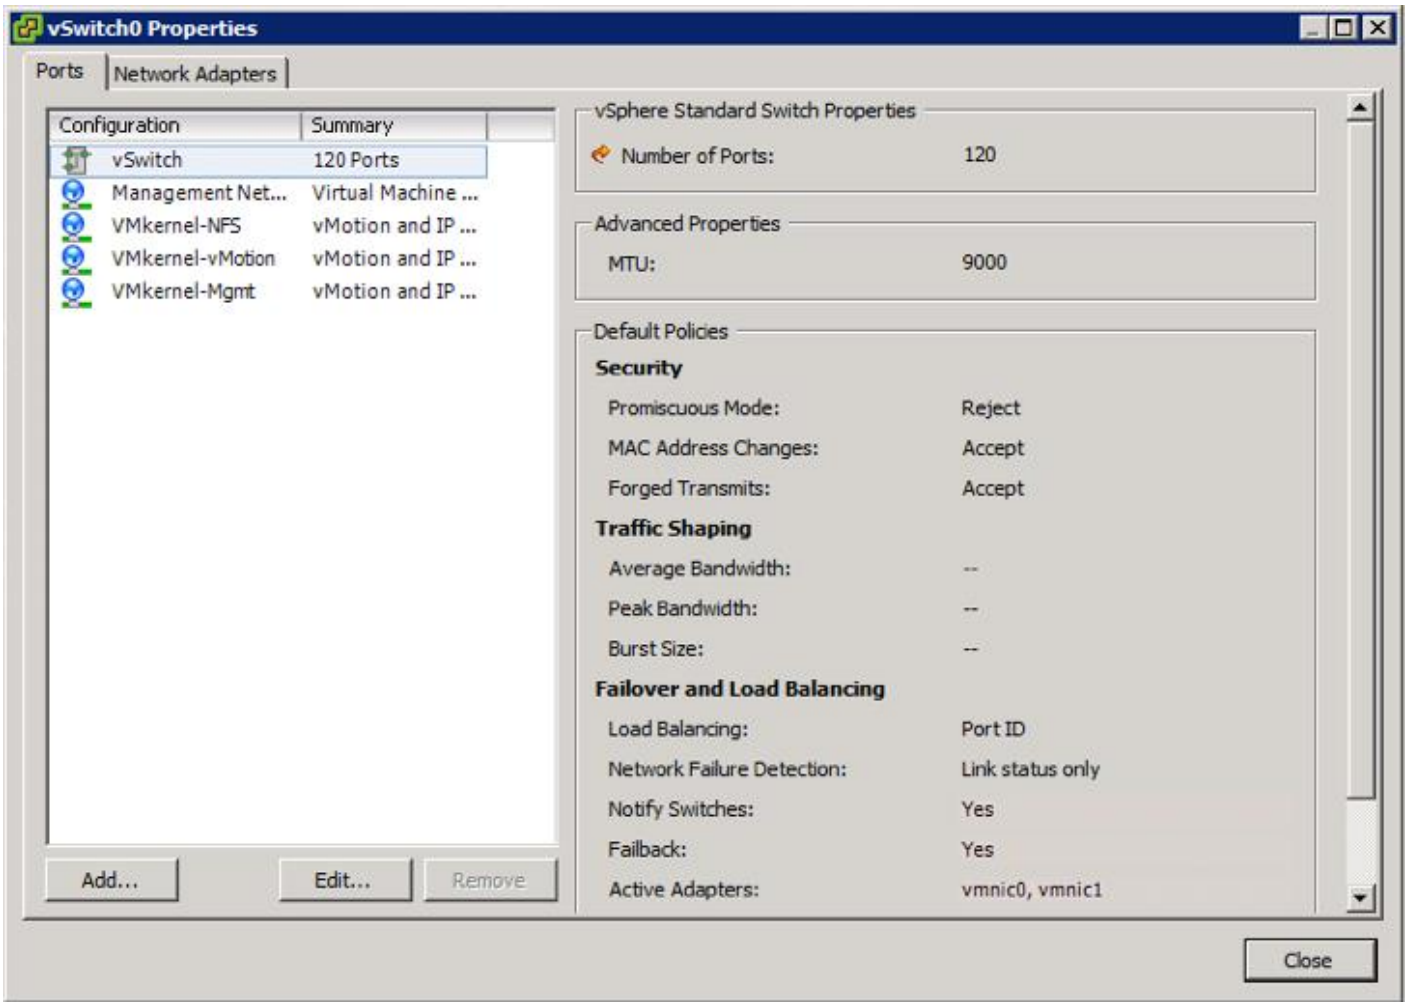

65. To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:

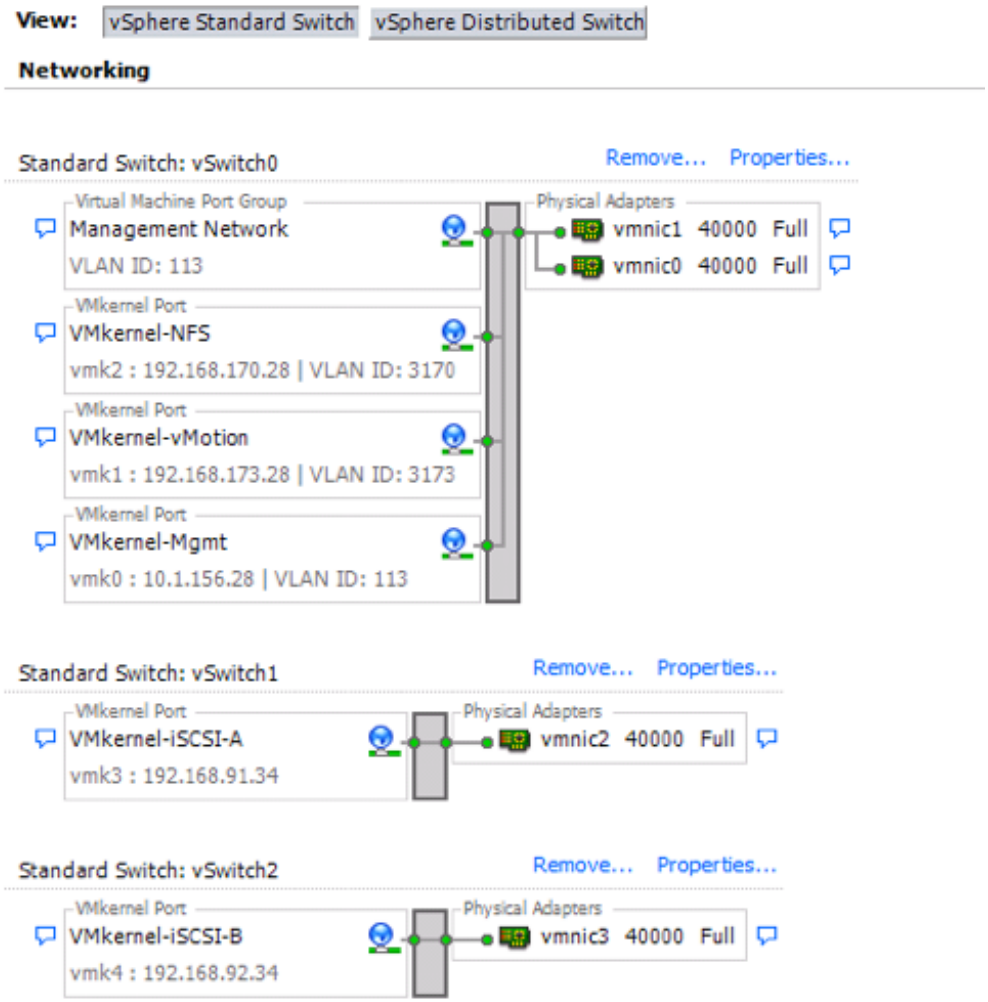

Setup iSCSI Multipathing

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To setup 4 iSCSI paths between storage and the ESXi host, complete the following steps on each ESXi host:

1. For the FC booted host, connect to the vSphere Client, select the configuration tab of the host.
2. Click Storage Adapters in the Hardware pane.
3. Select Add to create an iSCSI Software Adapter.
4. Specify the iqn assigned to the host by the Service Profile and create the iSCSI Software Adapter.
5. For both hosts, continue within the Storage Adapters section in the Hardware pane of the host Configuration tab.
6. Select the iSCSI Software Adapter and click Properties.

7. Select the Dynamic Discovery tab and click Add.
8. Enter the IP address of iscsi_lif01a.
9. Click OK.
10. Repeat putting in the IP addresses of iscsi_lif01b, iscsi_lif02a and iscsi_lif02b.

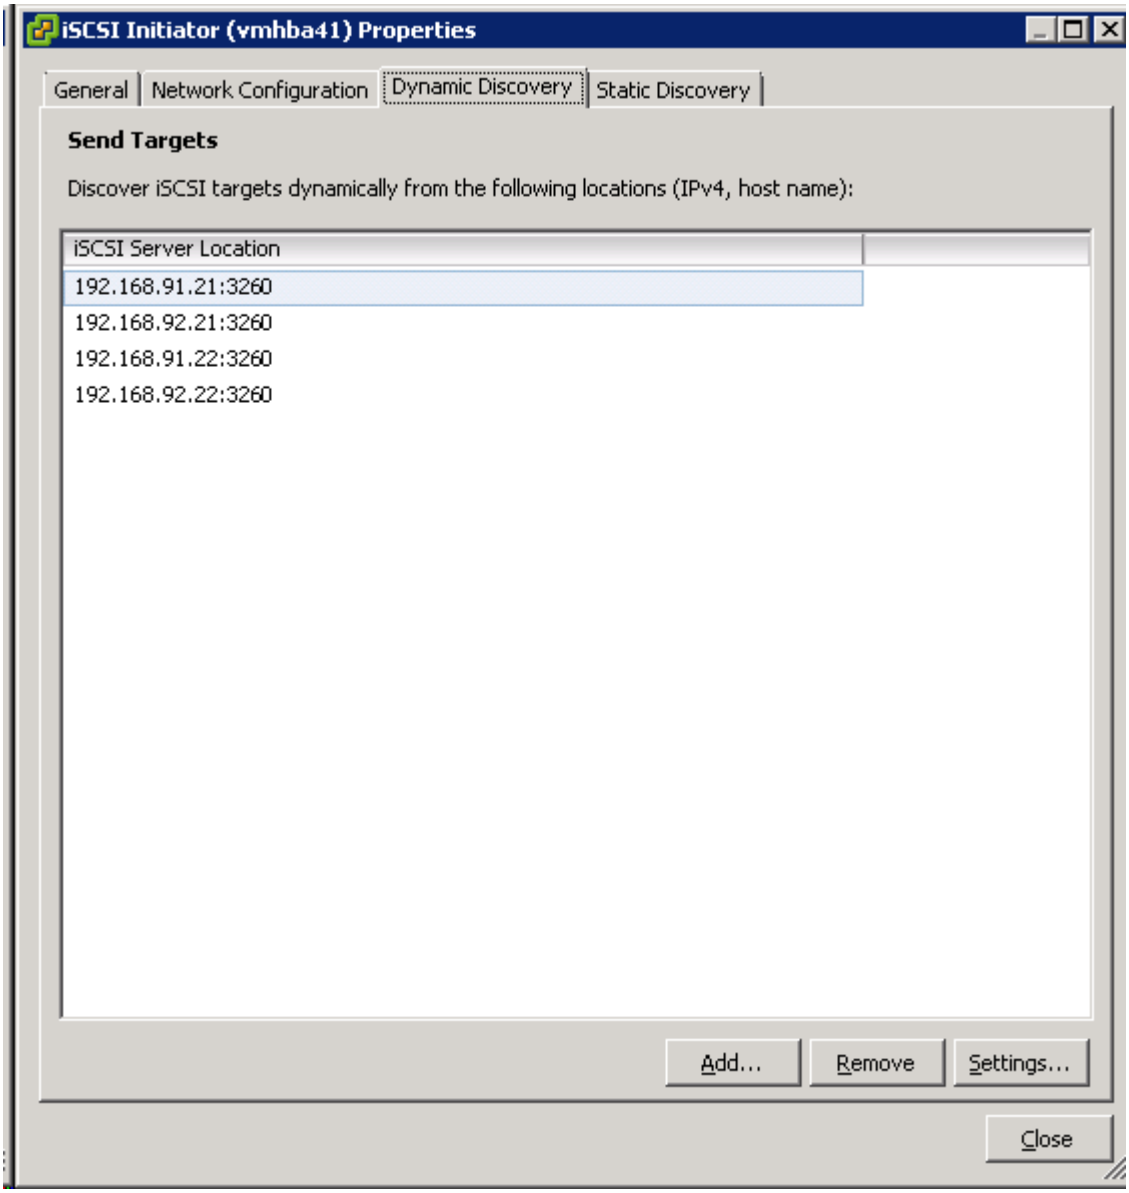

11. Click Close and then click yes to rescan the host bus adapter.
12. You should now see 4 connected paths in the Details pane.

Install VMware Drivers for the Cisco Virtual Interface Card (VIC)

Download and extract the following VMware VIC Drivers to the Management workstation:

fnic Driver version 1.6.0.25
enic Driver version 2.3.0.7

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Prod-02, 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. Click Enter Maintenance Mode within the Commands section of the Summary tab.
4. Click Yes for any dialogue box presented.
5. From Resources > Storage, right-click datastore1 and select Browse Datastore.
6. Click the fourth button and select Upload File.
7. Navigate to the saved location for the downloaded VIC drivers and select fnic_driver_1.6.0.25-3741467.zip.
8. Click Open and Yes to upload the file to datastore1.
9. Click the fourth button and select Upload File.
10. Navigate to the saved location for the downloaded VIC drivers and select ESXi60-enic-2.3.0.7-3642661.zip.
11. Click Open and Yes to upload the file to datastore1.
12. Make sure the files have been uploaded to both ESXi hosts.
13. Within the vSphere Client select the Configuration tab and click Security Profile within the Software section.

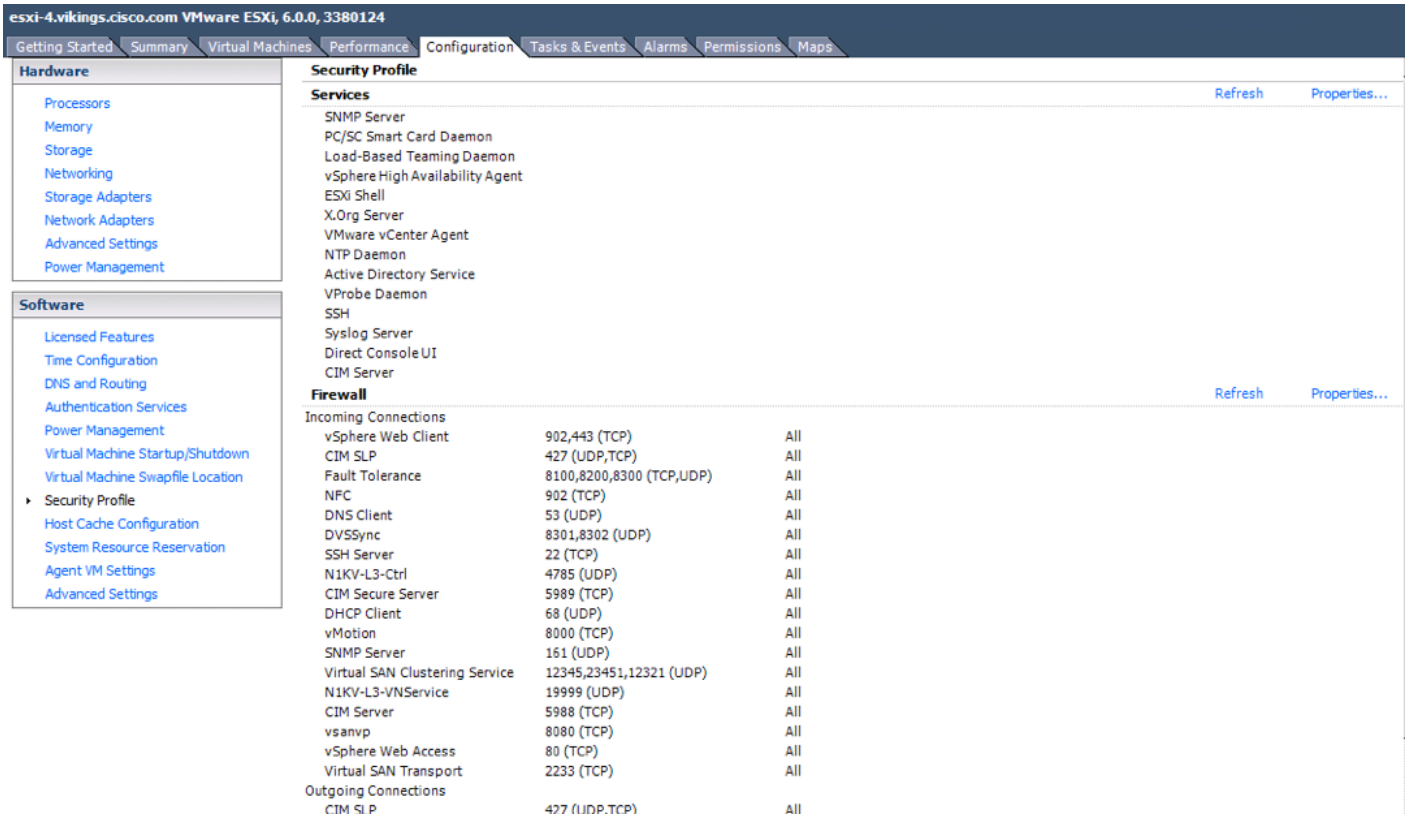

14. Click Properties within the Services section at the top.

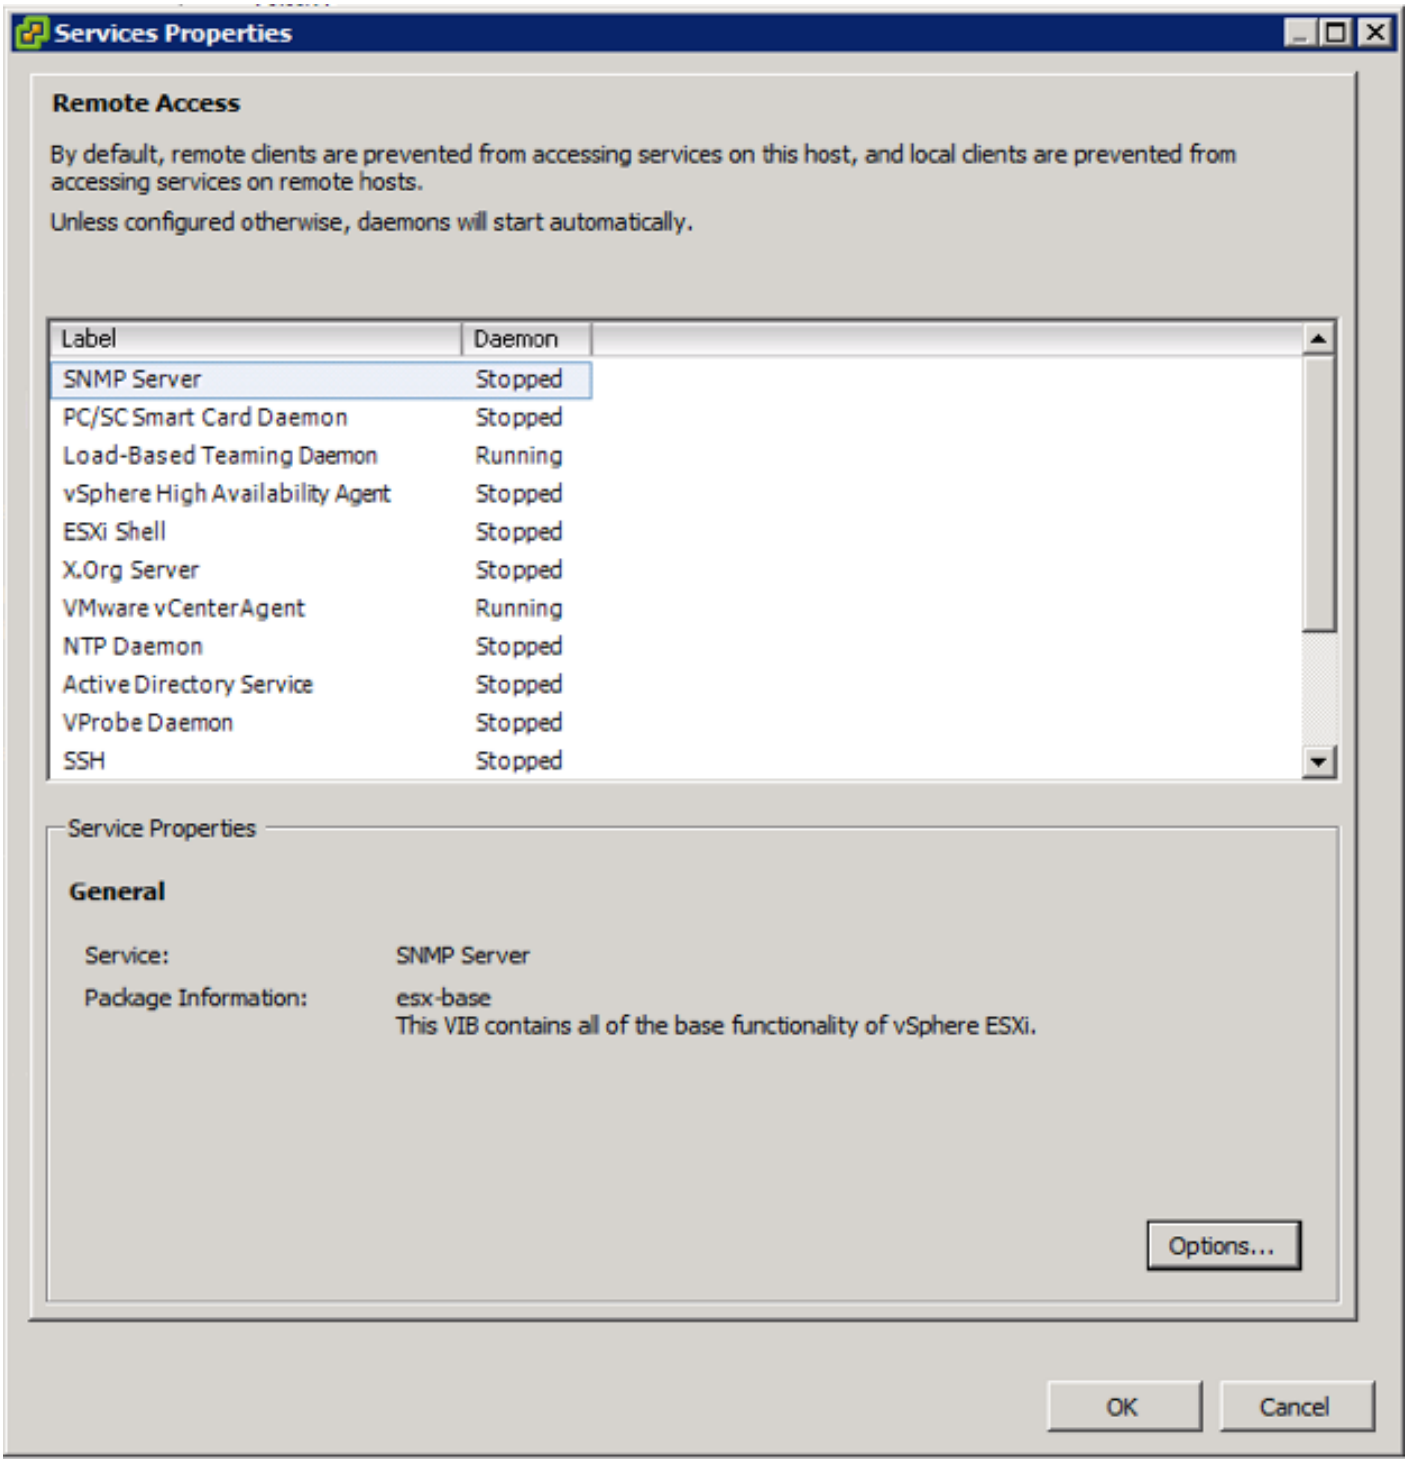

15. Select ESXi Shell and click the Options... button.

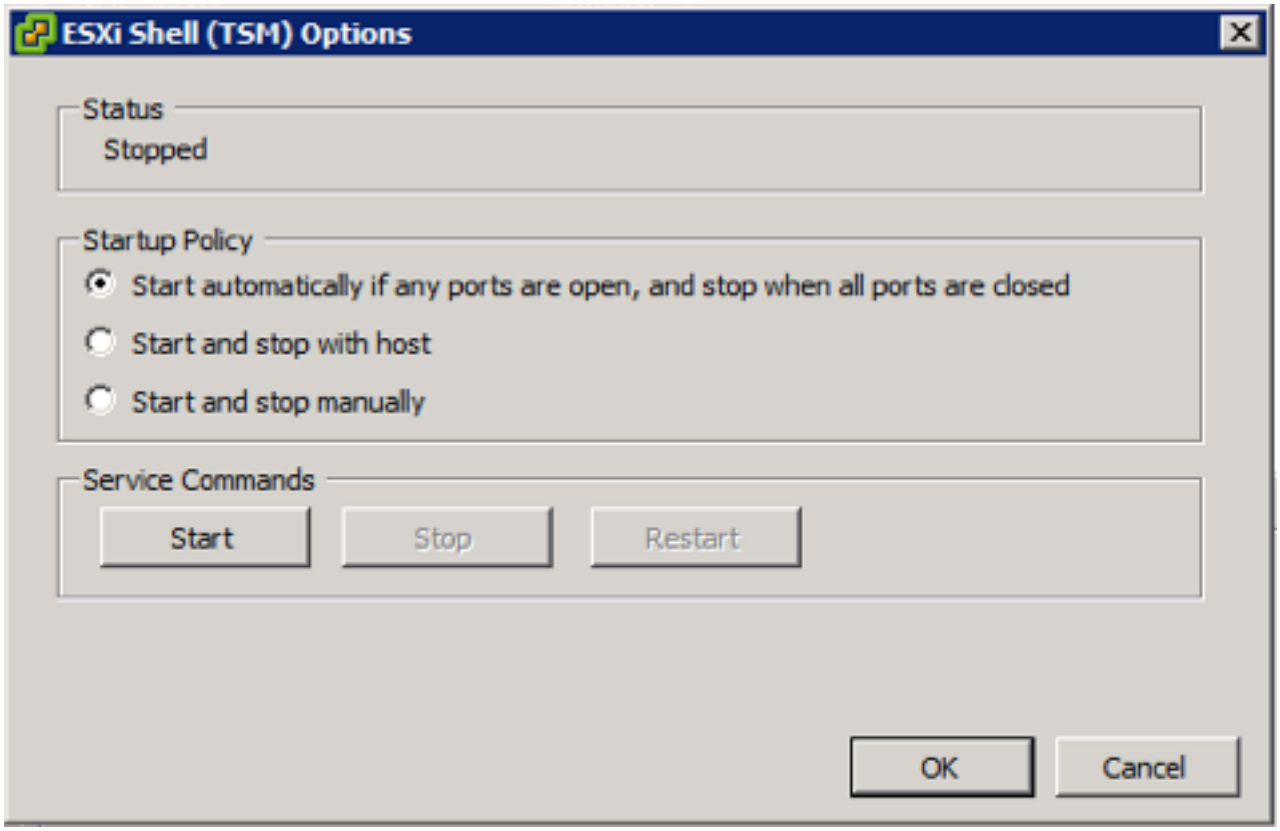

16. Select Start automatically if any ports are open, and stop when all ports are closed within the Startup Policy section.
17. Click Start, and OK.
18. Select SSH and click the Options... button.
19. Select Start automatically if any ports are open, and stop when all ports are closed within the Startup Policy section.
20. Click Start, and OK.
21. Click OK to exit from the Service Properties configuration window.

Enabling ssh can be considered optional if the VMware vSphere Remote CLI is installed and used.

22. Connect to each ESXi hosts through ssh from a shell connection or putty terminal.
23. Login as root with the password specified for <<var_password>>.
24. Run the following commands on each host:
25. esxcli software vib update -d /vmfs/volumes/datastore/fnic_driver_1.6.0.25-offline_bundle-3642682.zip

26. esxcli software vib install -d /vmfs/volumes/datastore/ESXi60-enic-2.3.0.7-offline_bundle-3642661.zip

27. Reboot each host after both commands have been run.

Mount Required Datastores

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

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. To enable configurations, click the Configuration tab.

3. Click Storage in the Hardware pane.

4. From the Datastores area, click Add Storage to open the Add Storage wizard.

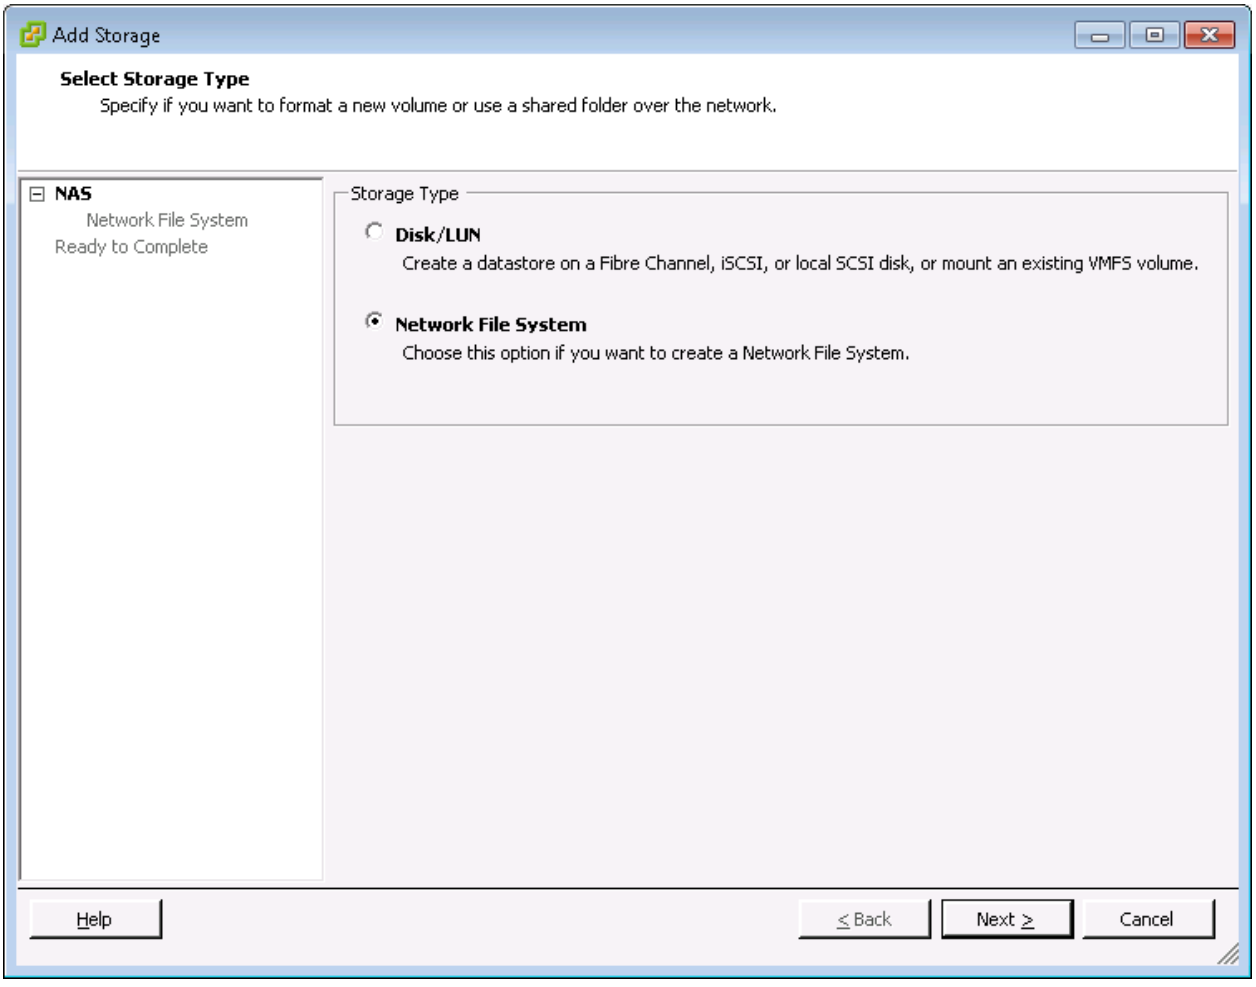

5. Select Network File System and click Next.

6. The wizard prompts for the location of the NFS export. Enter `<<var_node02_nfs_lif_infra_datastore_1_ip>>` as the IP address for `nfs_lif_infra_datastore_1`.

7. Enter `/infra_datastore_1` as the path for the NFS export.

8. Confirm that the Mount NFS read only checkbox is not selected.

9. Enter `infra_datastore_1` as the datastore name.

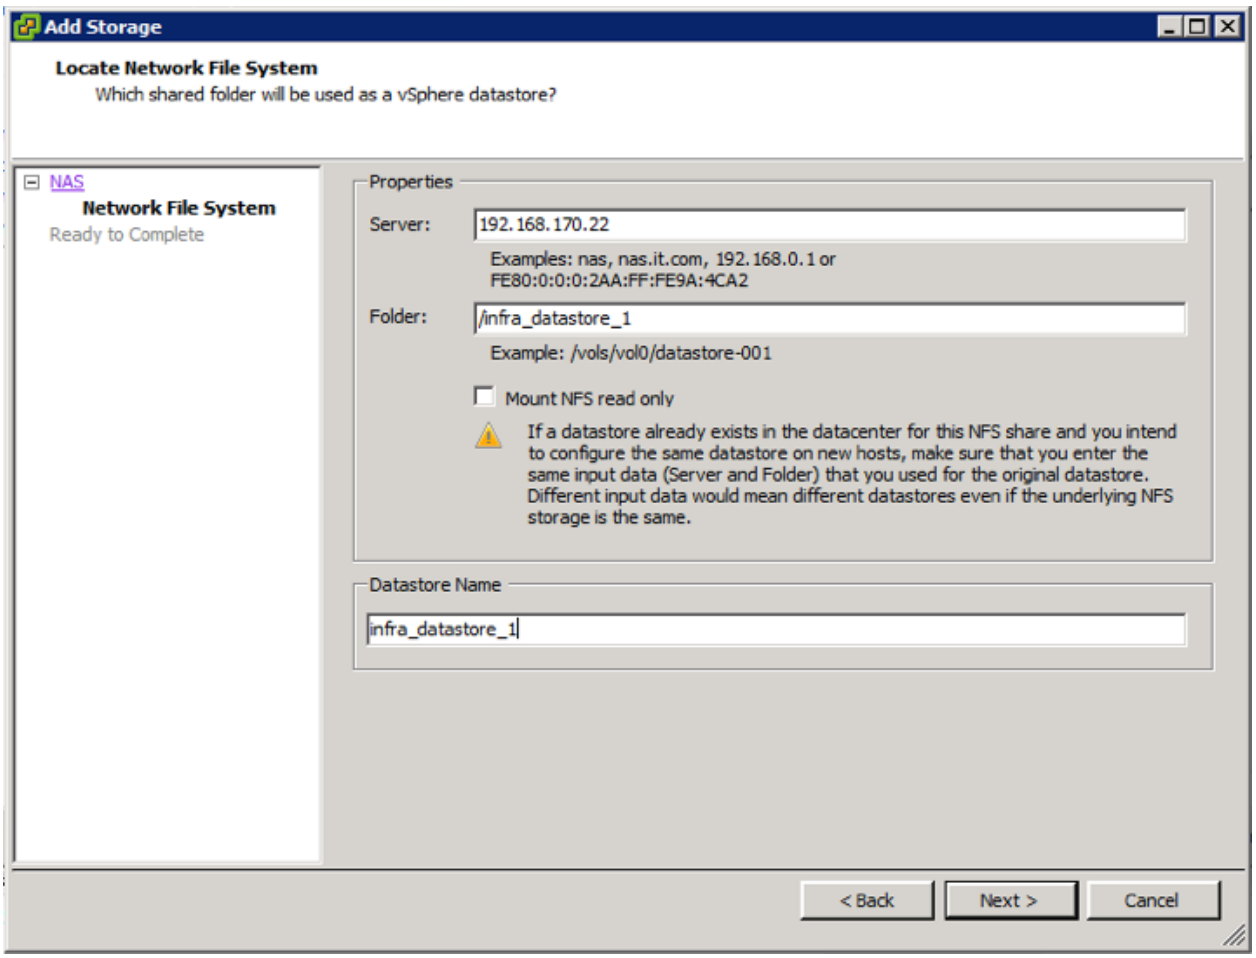

10. To continue with the NFS datastore creation, click Next.

11. To finalize the creation of the NFS datastore, click Finish.

12. From the Datastores area, click Add Storage to open the Add Storage wizard.

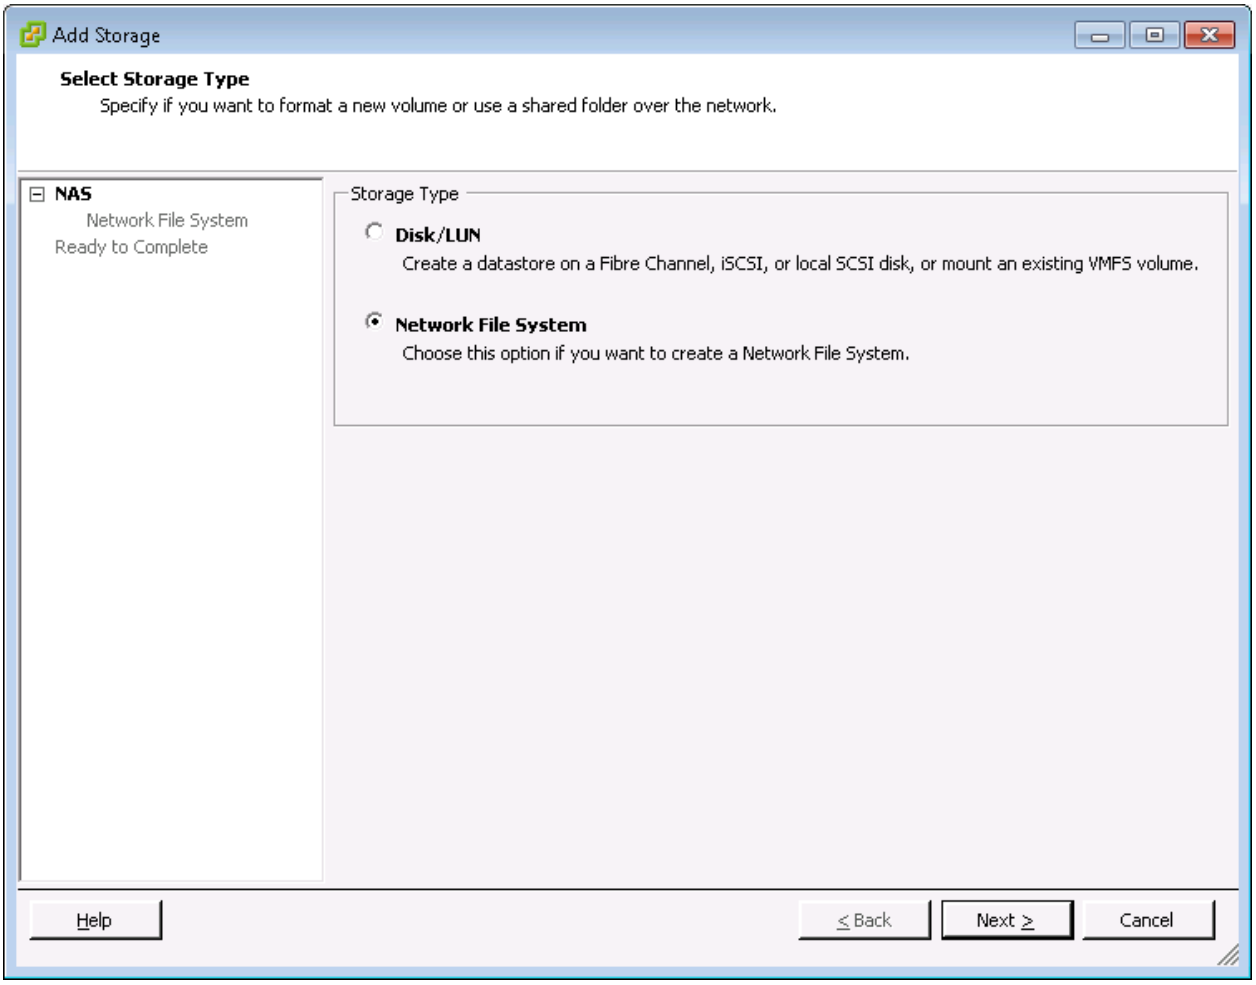

13. Select Network File System and click Next.

14. The wizard prompts for the location of the NFS export. Enter `<<var_node01_nfs_lif_infra_swap_ip>>` as the IP address for `nfs_lif_infra_swap`.

15. Enter `/infra_swap` as the path for the NFS export.

16. Confirm that the Mount NFS read only checkbox is not selected.

17. Enter `infra_swap` as the datastore name.

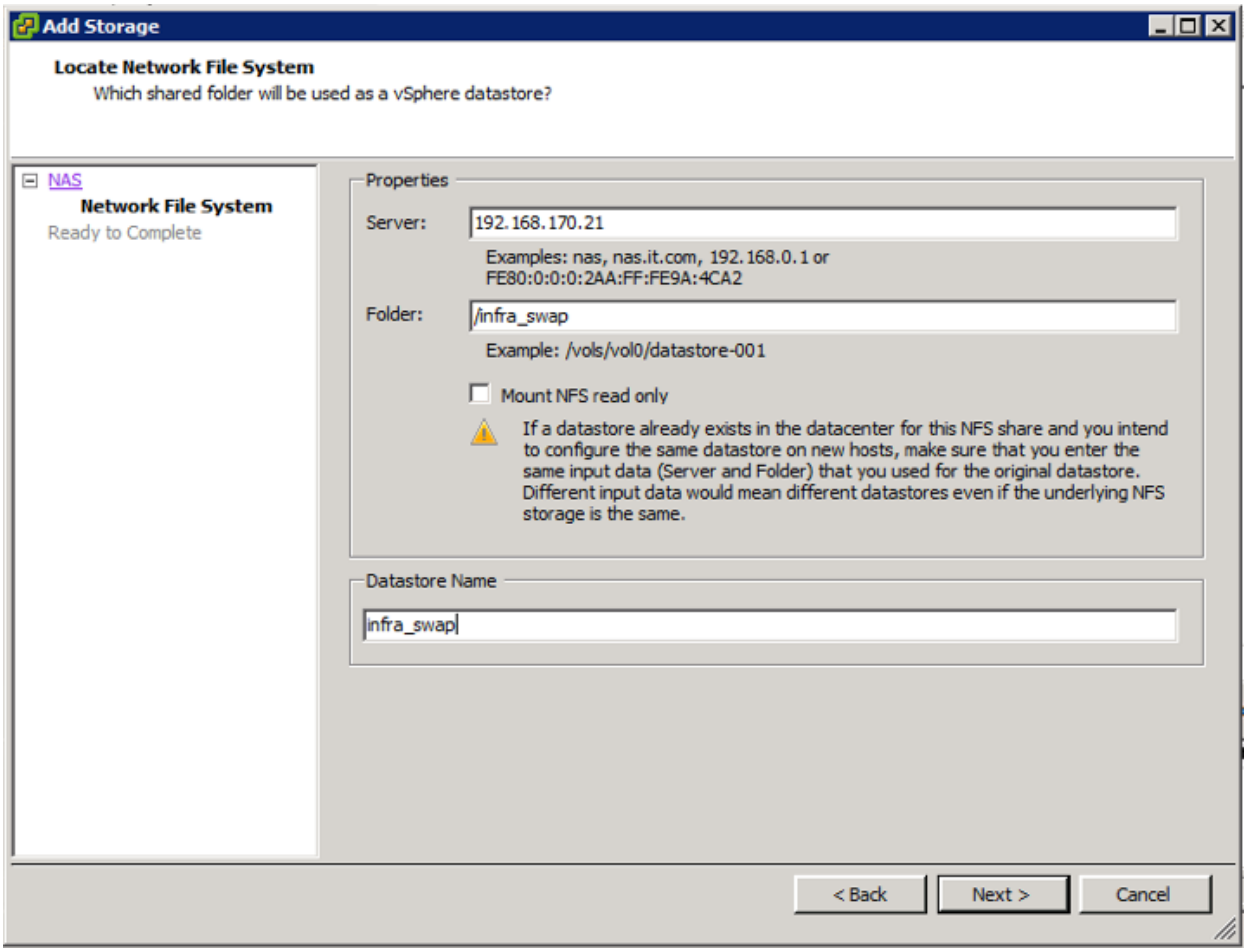

18. To continue with the NFS datastore creation, click Next.

19. To finalize the creation of the NFS datastore, click Finish.

Configure NTP on ESXi Hosts

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Time Configuration in the Software pane.
4. Click Properties at the upper-right side of the window.
5. At the bottom of the Time Configuration dialog box, click Options.
6. In the NTP Daemon (ntpd) Options dialog box, complete the following steps:
    a. Click General in the left pane and select Start and stop with host.

b. Click NTP Settings in the left pane and click Add.

7. In the Add NTP Server dialog box, enter <<var_switch_a_ntp_ip>> as the IP address of the NTP server and click OK.

8. Click Add.

9. In the Add NTP Server dialog box, enter <<var_switch_b_ntp_ip>> as the IP address of the NTP server and click OK.

10. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes checkbox and click OK.

11. In the Time Configuration dialog box, complete the following steps:

    a. Select the NTP Client Enabled checkbox and click OK.

    b. Verify that the clock is now set to approximately the correct time.

 The NTP server time may vary slightly from the host time.

Move VM Swap File Location

ESXi VM-Host-Infra-01 and VM-Host-Prod-02

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. To enable configurations, 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 “Store the swapfile in a swapfile datastore selected below.”

6. Select the <infra_swap> datastore in which to house the swap files.

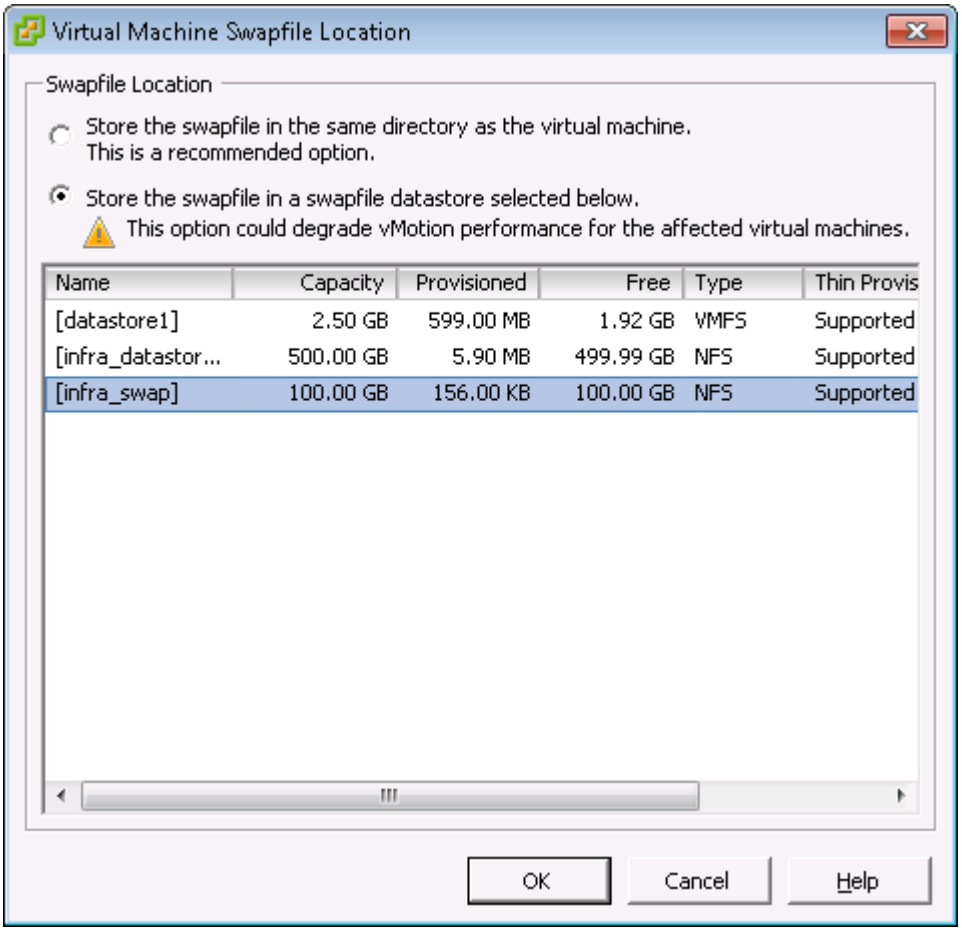

7. Click OK to finalize moving the swap file location.

VMware vCenter 6.0 U1b

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.0 U1b Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Install the Client Integration Plug-in

To install the client integration plug-in, complete the following steps:

1. Download the .iso installer for the vCenter Server Appliance and Client Integration Plug-in.

2. Mount the ISO image to the Windows virtual machine or physical server on which you want to install the Client Integration Plug-In to deploy the vCenter Server Appliance.

3. In the software installer directory, navigate to the vcsa directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe. The Client Integration Plug-in installation wizard appears.

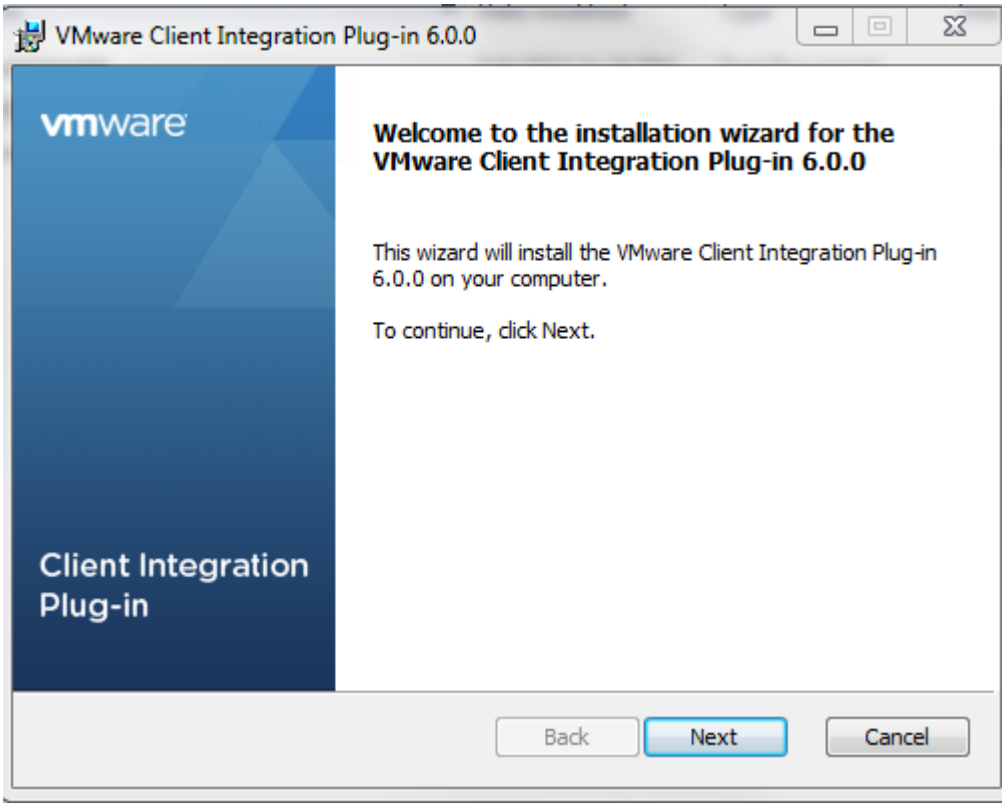

4. On the Welcome page, click Next.
5. Read and accept the terms in the End-User License Agreement and click Next.
6. Click Next.
7. Click Install.

Building the VMware vCenter Server Appliance

To build the VMware vCenter virtual machine, complete the following steps:

1. In the software installer directory, double-click vcsa-setup.html.
2. Allow the plug-in to run on the browser when prompted.

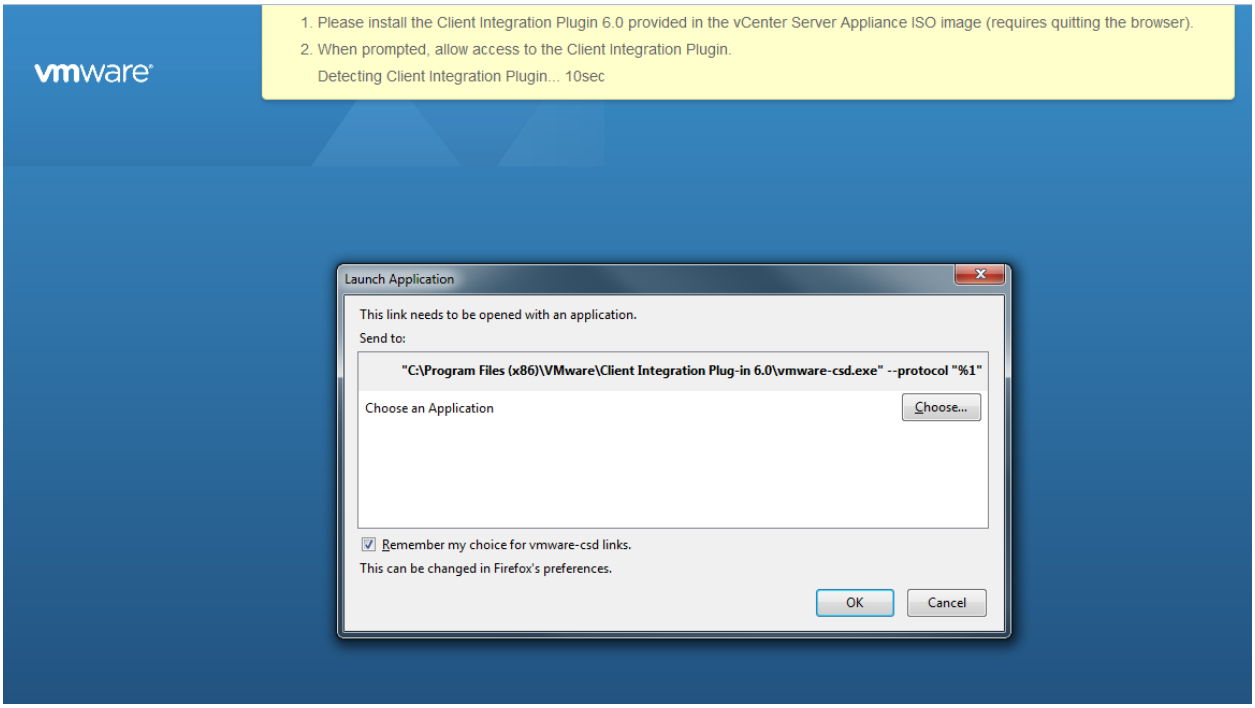

3. On the Home page, click Install to start the vCenter Server Appliance deployment wizard.

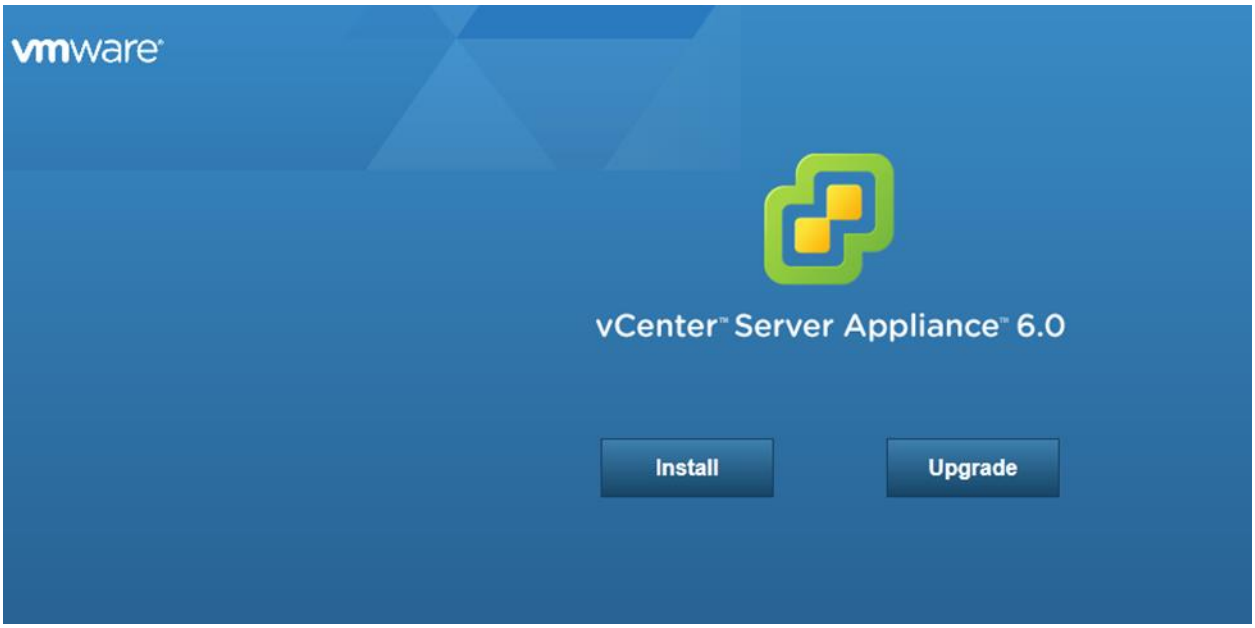

4. Read and accept the license agreement, and click Next.

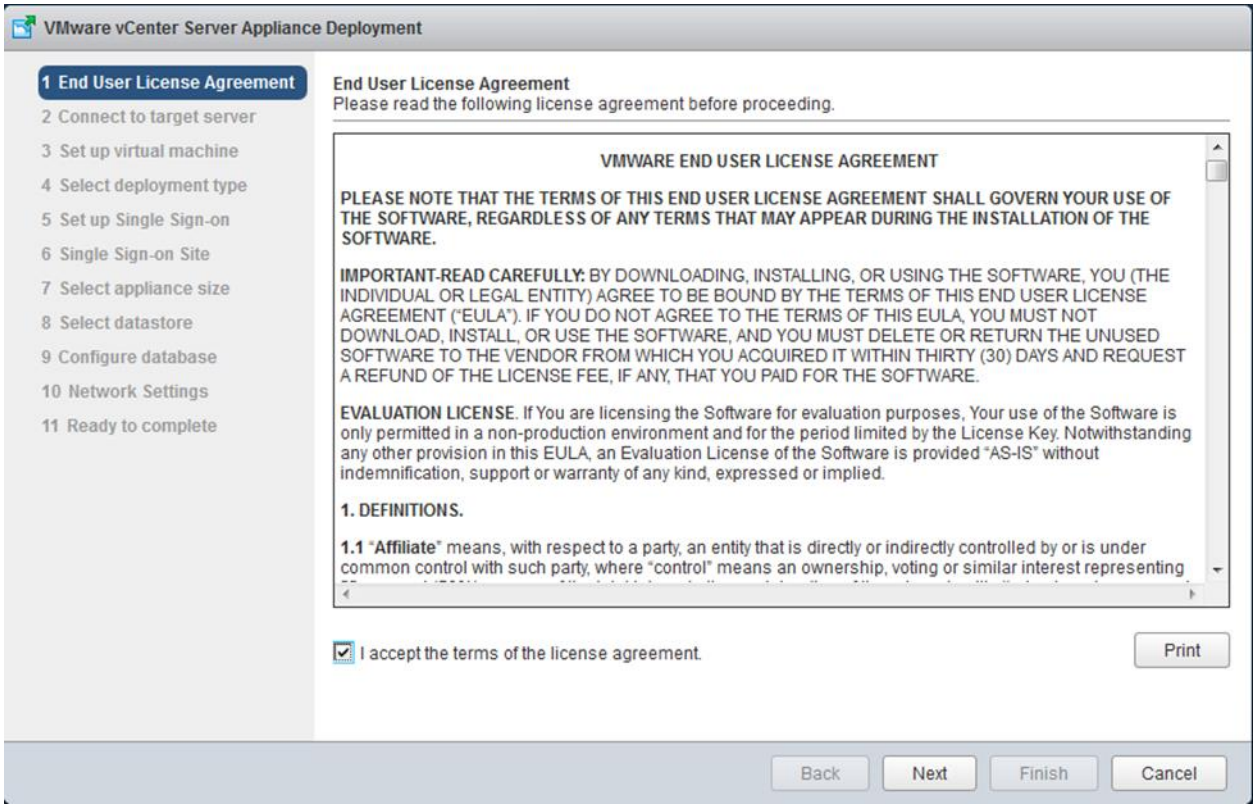

5. In the "Connect to target server" page, enter the ESXi host name, User name and Password.

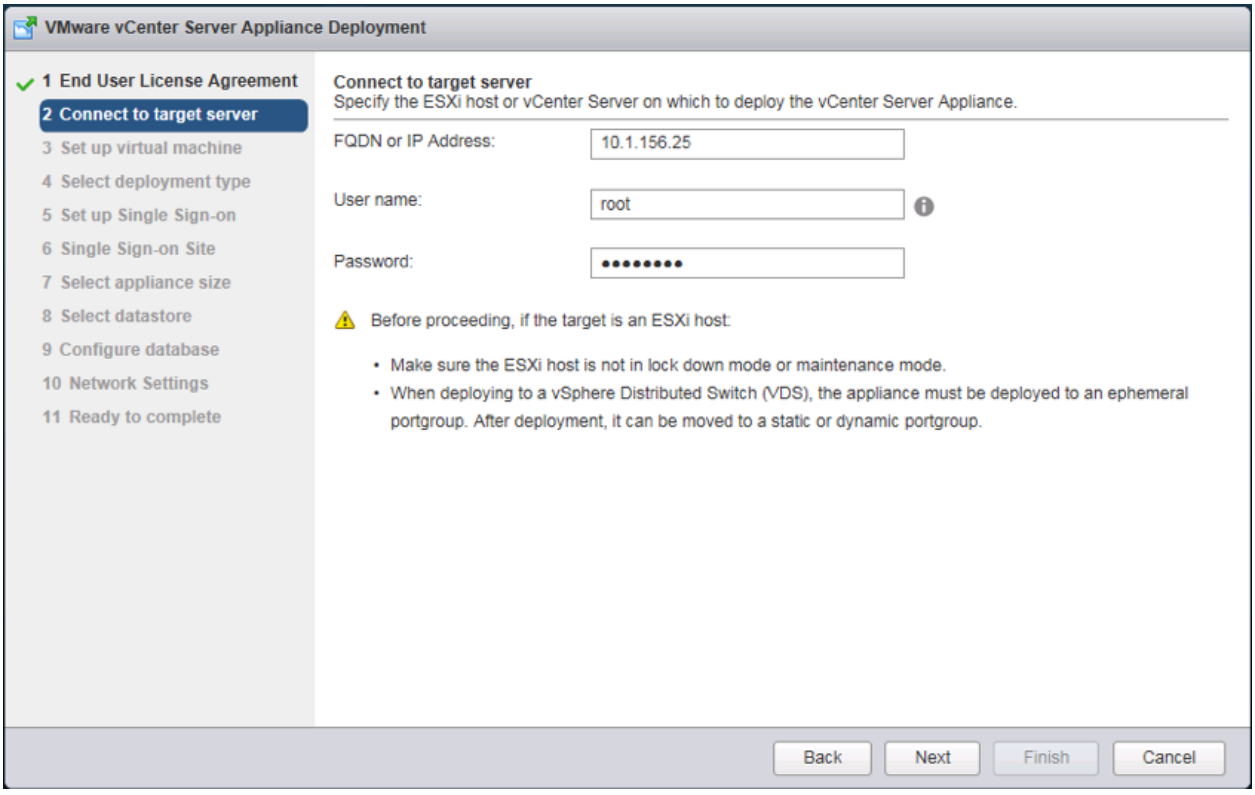

6. Click Yes to accept the certificate.

7. Enter the Appliance name and password details in the “Set up virtual machine” page.

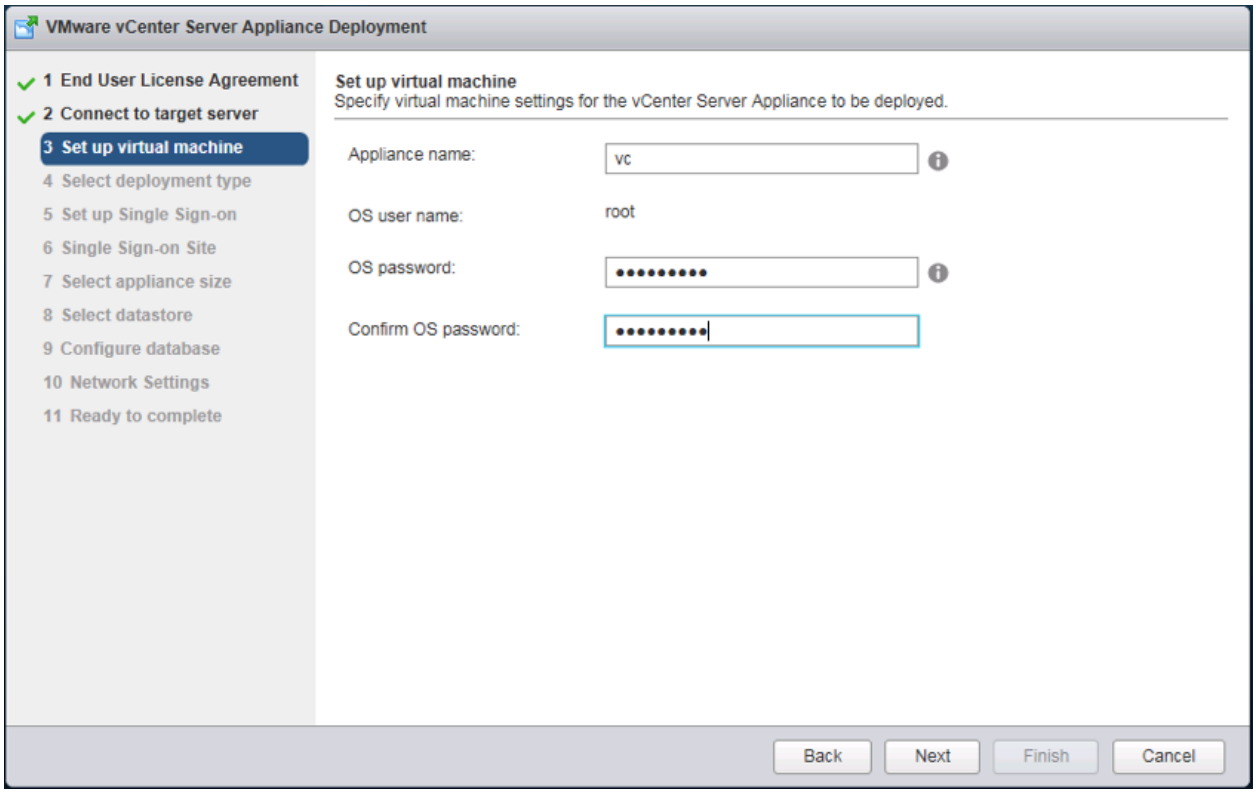

8. In the “Select deployment type” page, choose “Install vCenter Server with an embedded Platform Services Controller.”

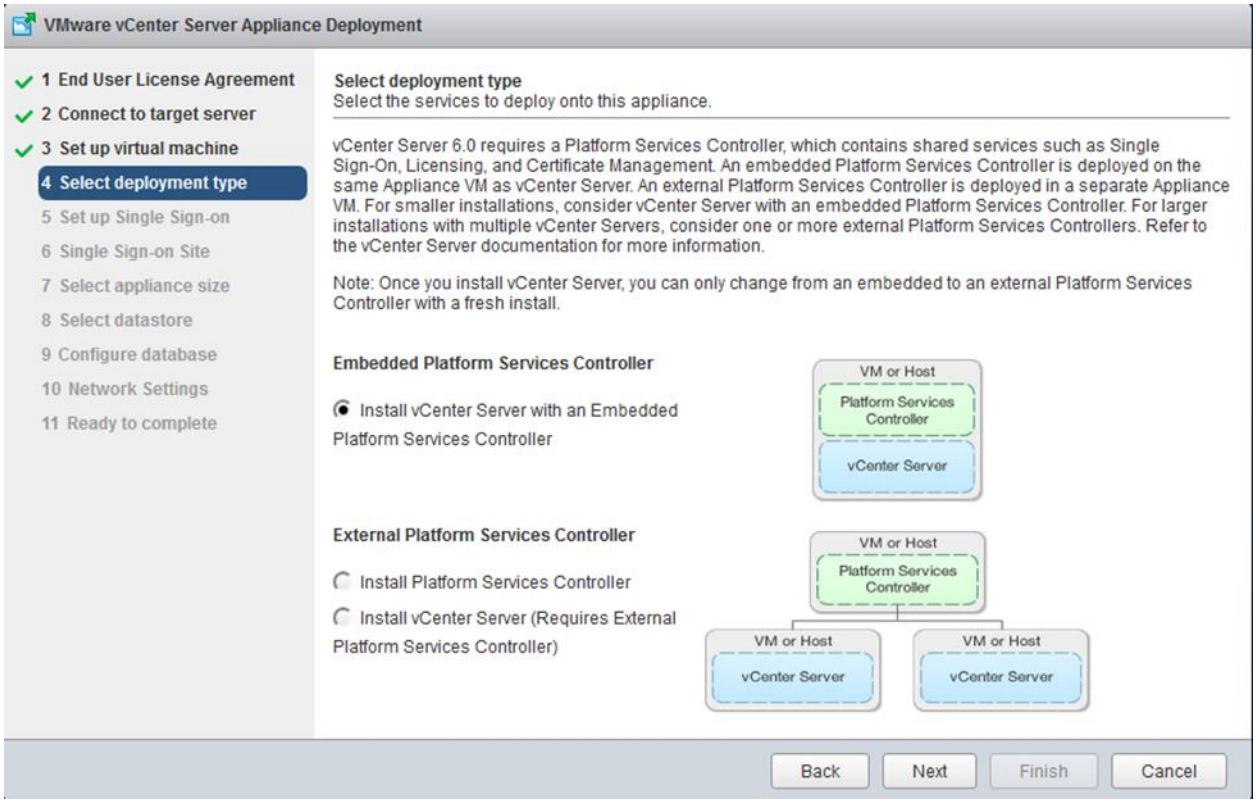

9. Click Next.

10. In the “Set up Single Sign-On” page, select “Create a new SSO domain.”

11. Enter the SSO password, Domain name and Site name.

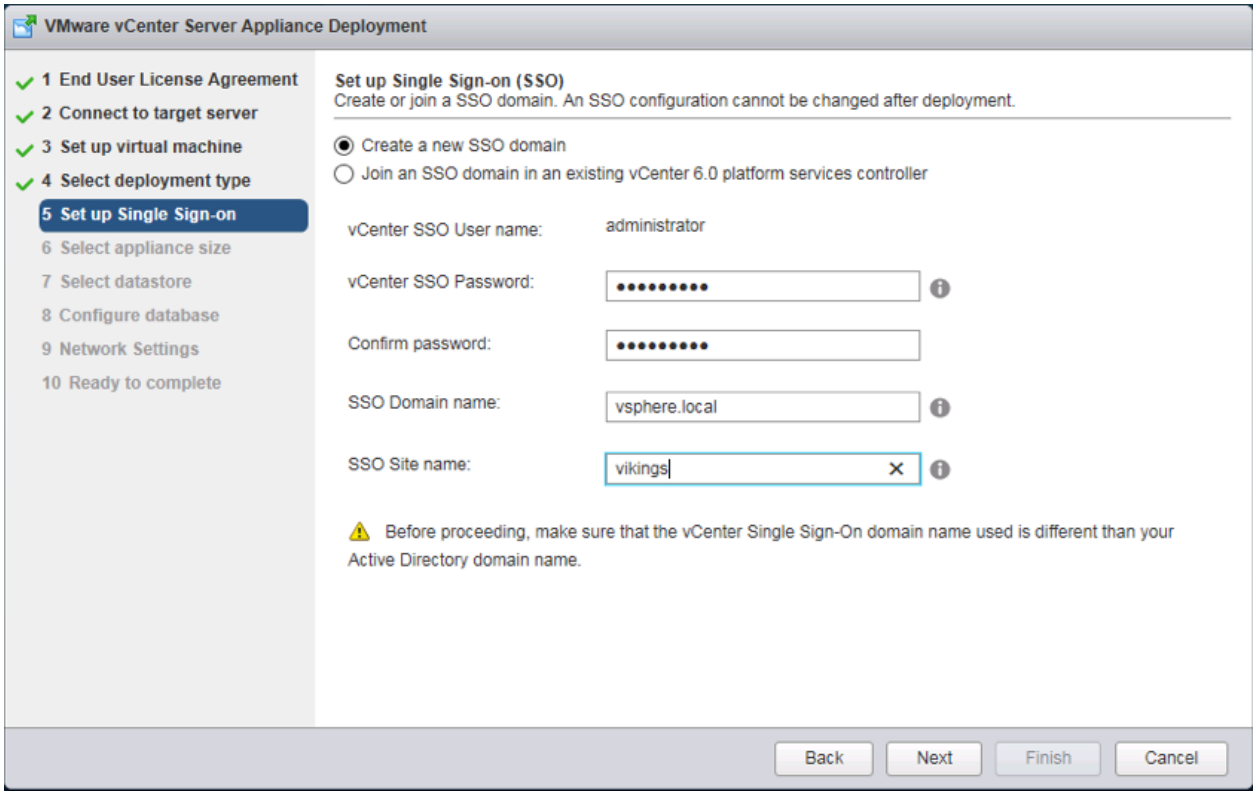

12. Click Next.

13. Select the appliance size. For example, “Tiny (up to 10 hosts, 100 VMs).”

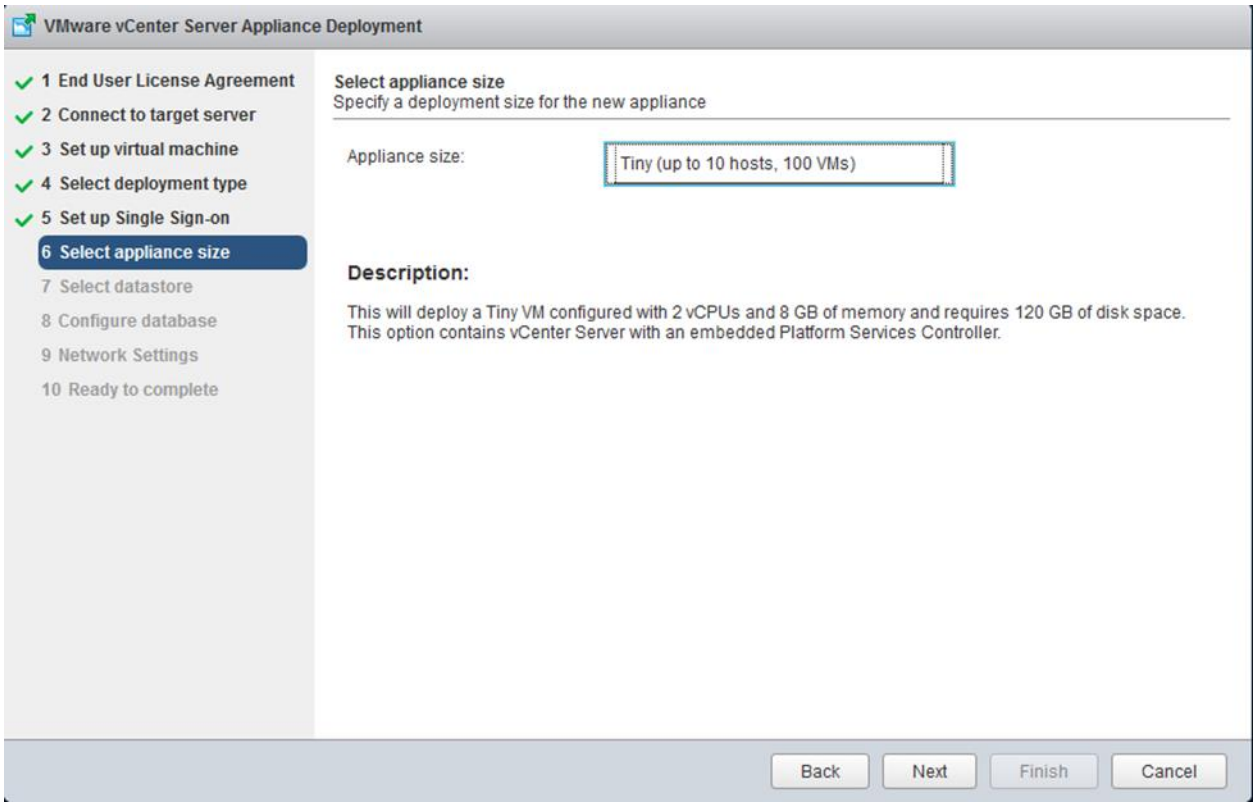

14. Click Next.

15. In the “Select datastore” page, choose infra_datastore_1.

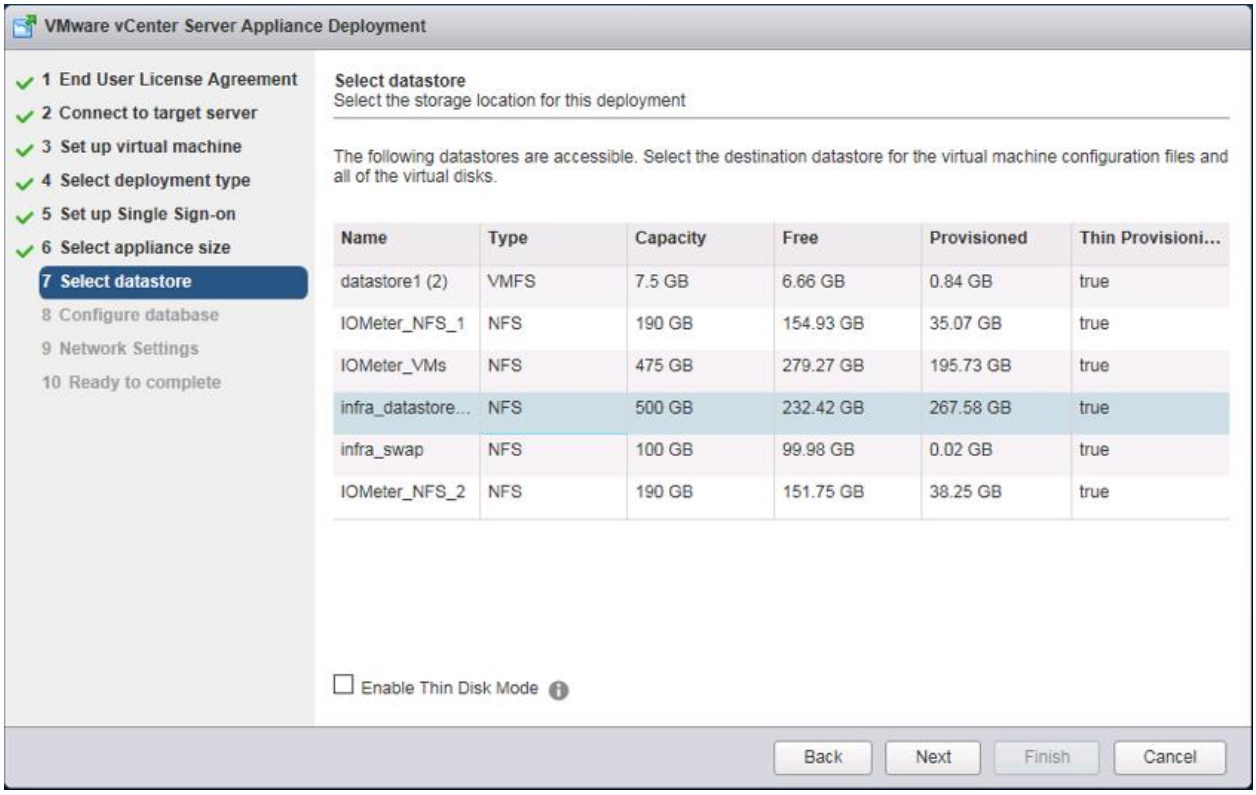

16. Click Next.

17. Select embedded database in the “Configure database” page. Click Next.

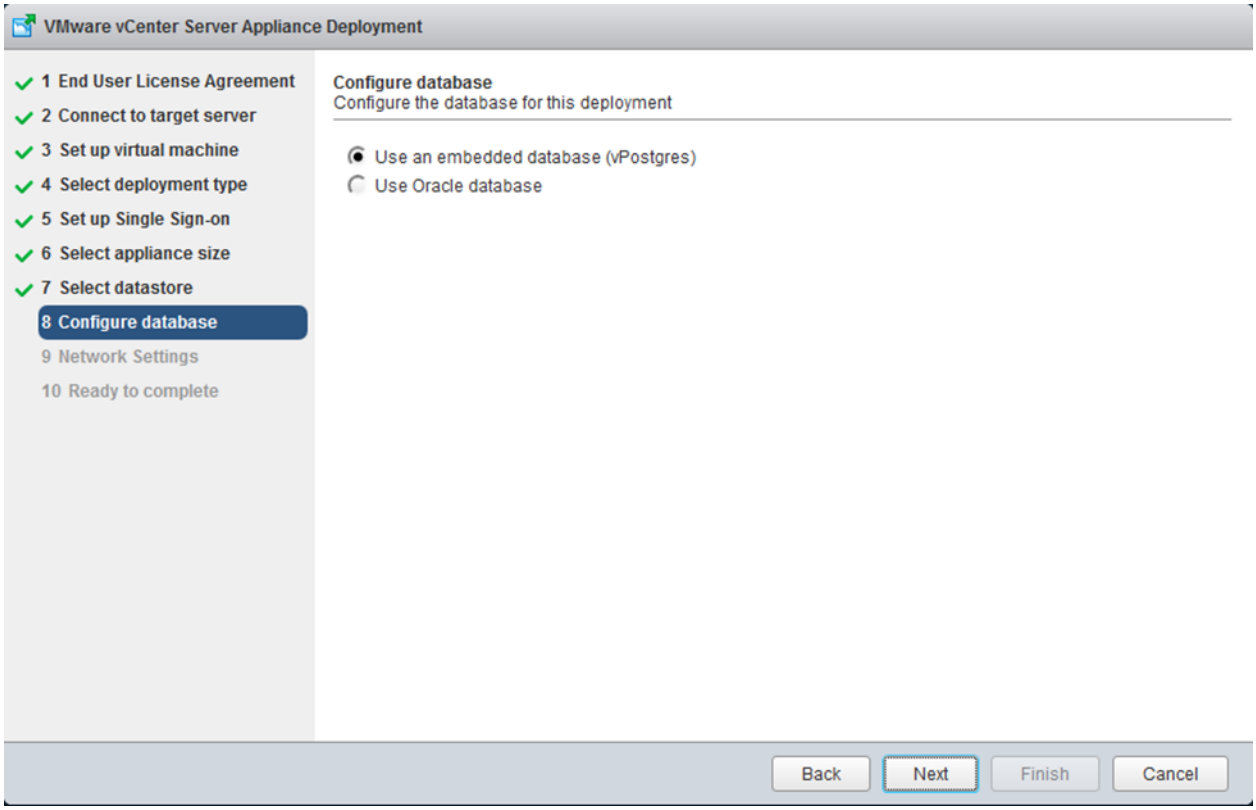

18. In the “Network Settings” page, configure the below settings:

    a. Choose a Network: MGMT-Network
    b. IP address family: IPV4
    c. Network type: static
    d. Network address: <<var_vcenter_ip>>
    e. System name: <<var_vcenter_fqdn>>
    f. Subnet mask: <<var_vcenter_subnet_mask>>
    g. Network gateway: <<var_vcenter_gateway>>
    h. Network DNS Servers: <<var_dns_server>>
    i. Configure time sync: Use NTP servers
    j. (Optional). Enable SSH

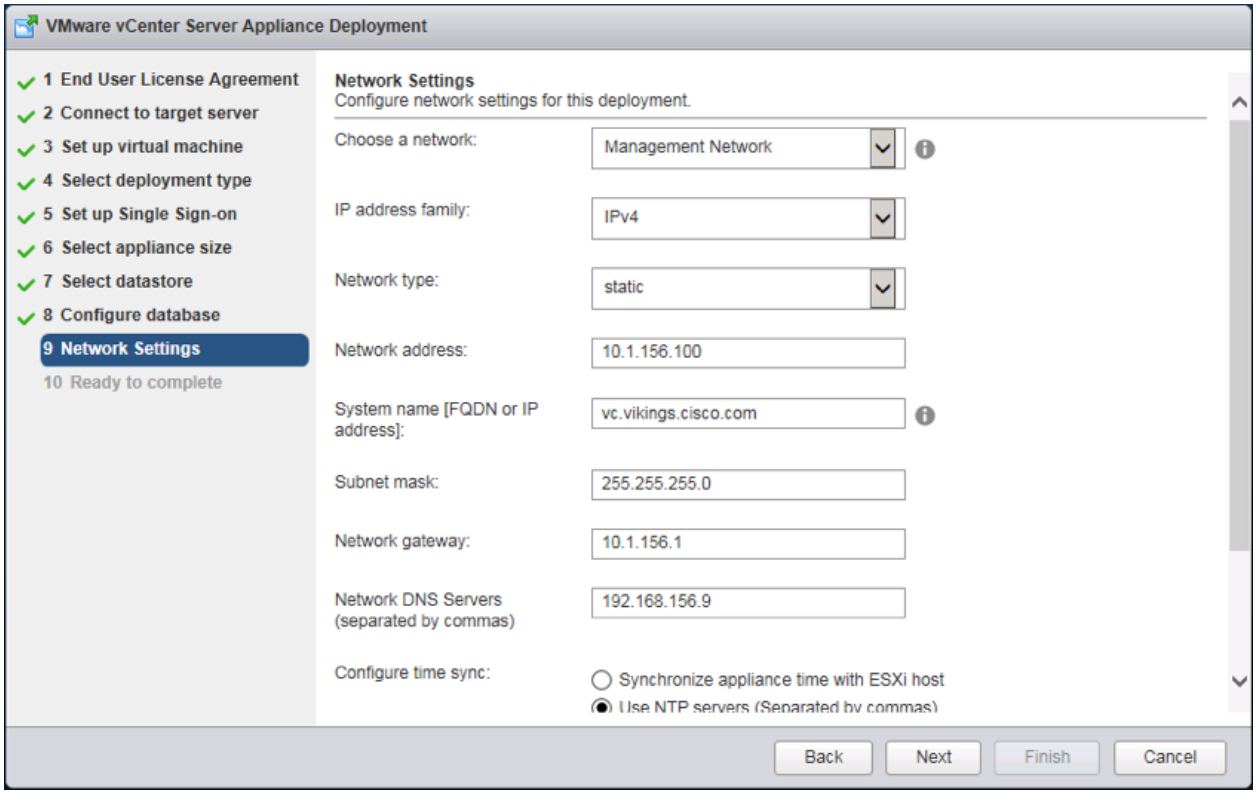

19. Review the configuration and click Finish.

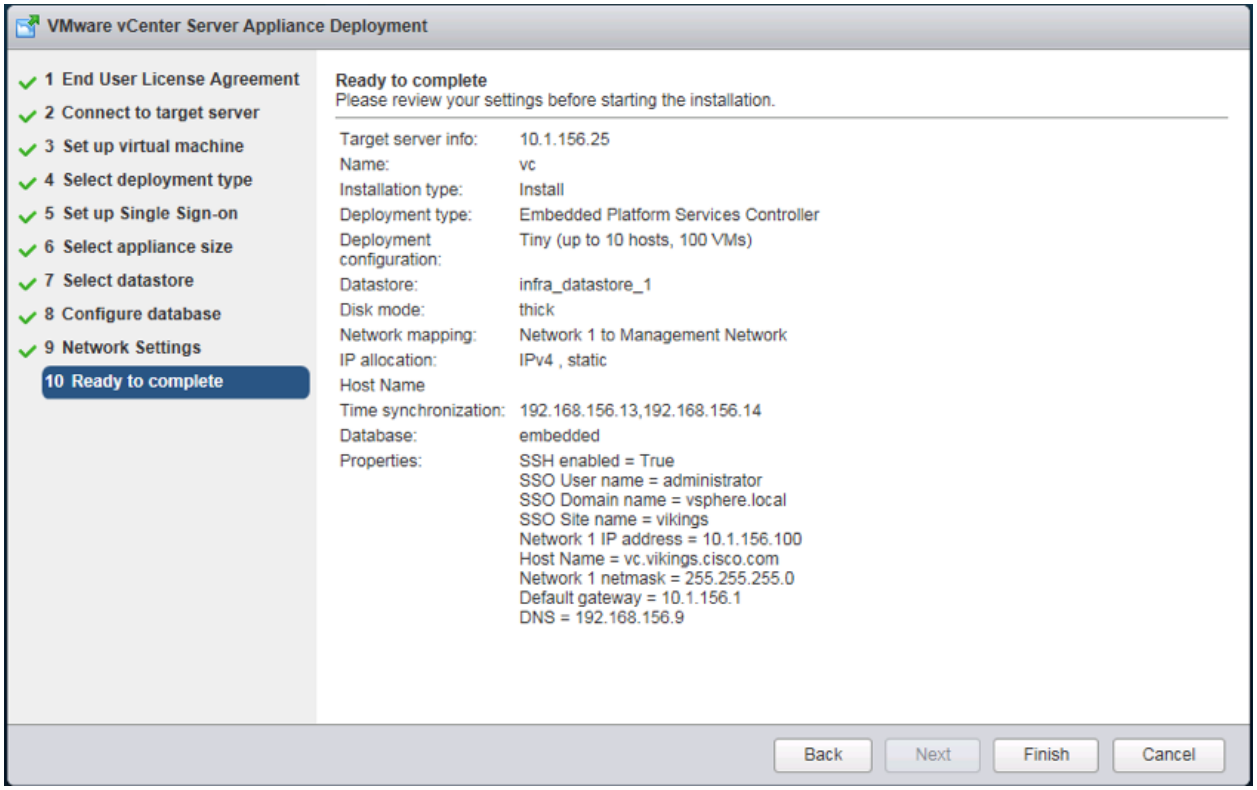

20. The vCenter appliance installation will take few minutes to complete.

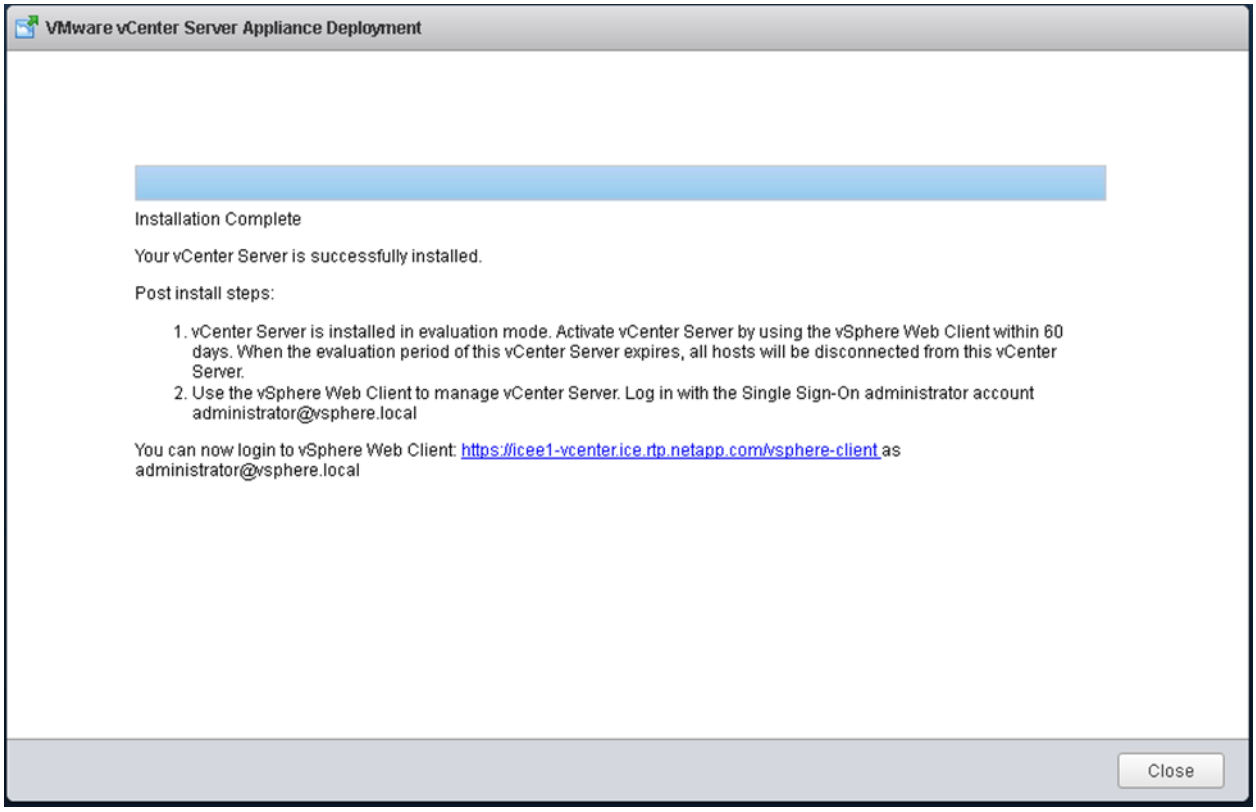

Setting Up VMware vCenter Server

1. Using a web browser, navigate to https://<<var_vcenter_ip>.

2. Click Log in to vSphere Web Client.

3. Click OK if “Launch Application” window appears.

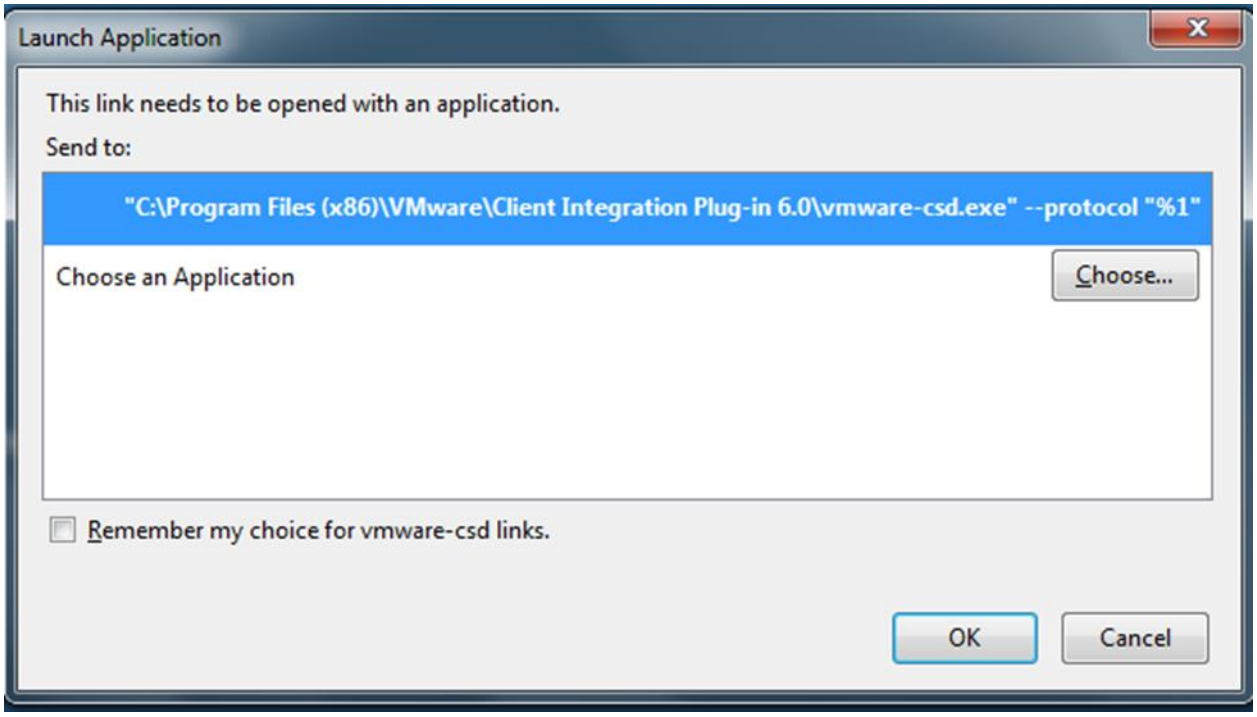

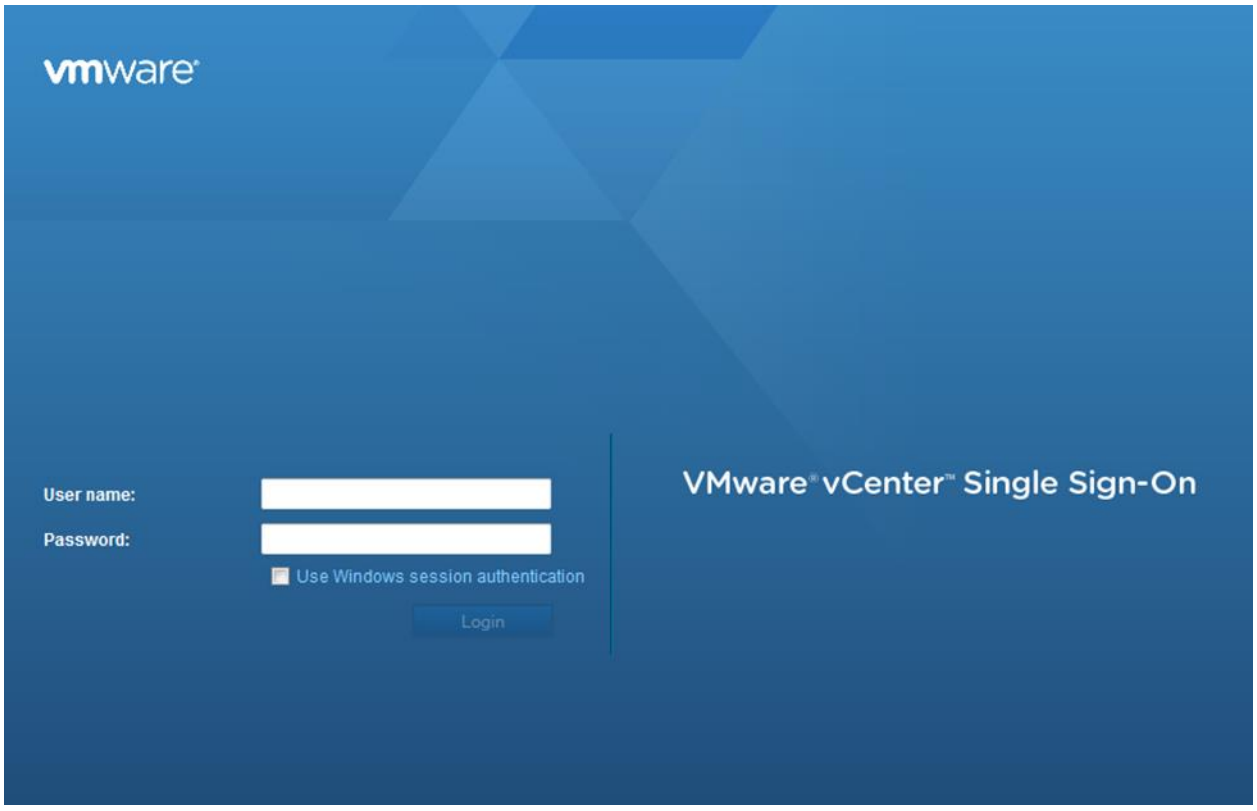

4. Log in using Single Sign-On username and password created during the vCenter installation.

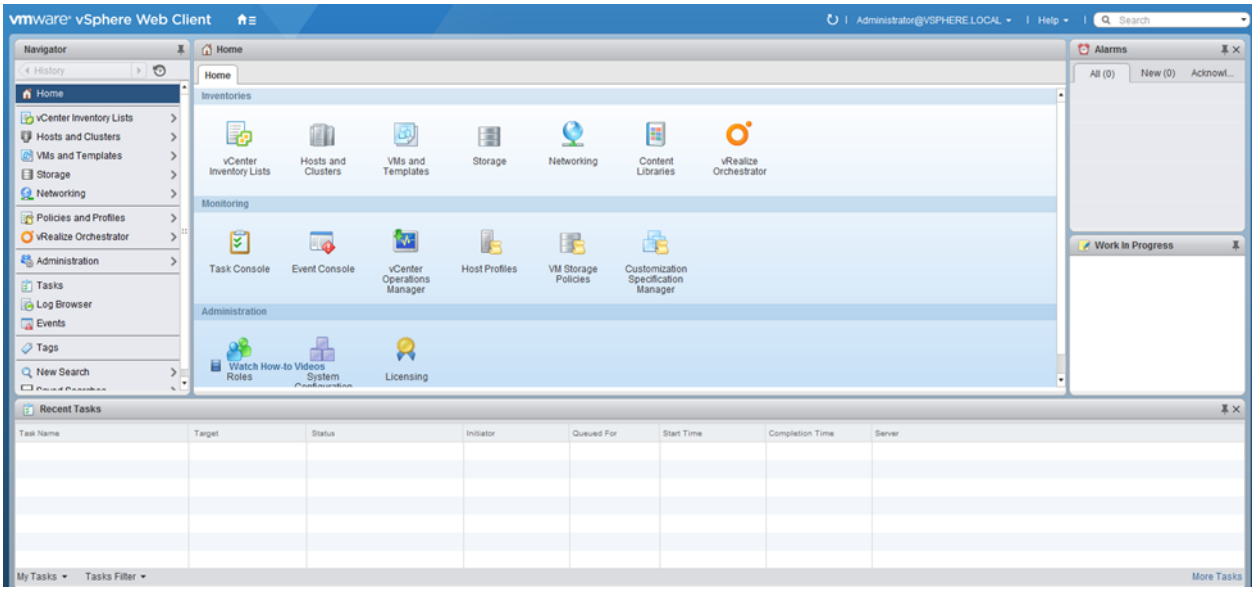

5. Navigate to vCenter Inventory Lists on the left pane.

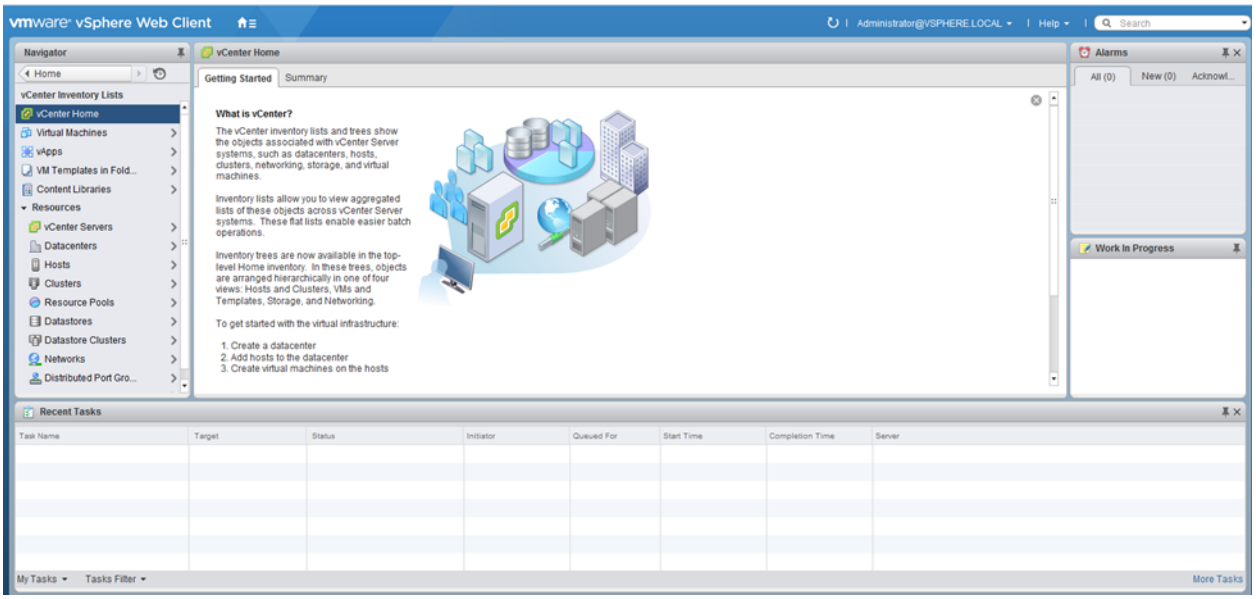

6. Under Resources, click Datacenters in the left plane.

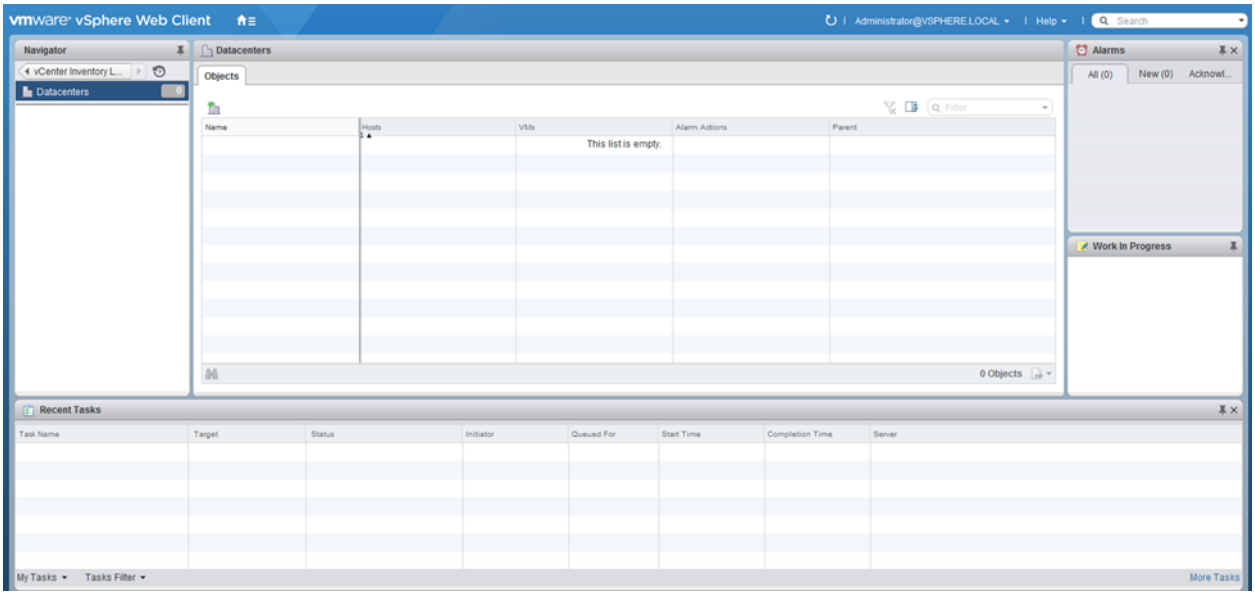

7. To create a Data center, click the leftmost icon in the center pane that has a green plus symbol above it.

8. Type “FlexPod_DC” in the Datacenter name field.

9. Select the vCenter Name/IP option.

10. Click OK.

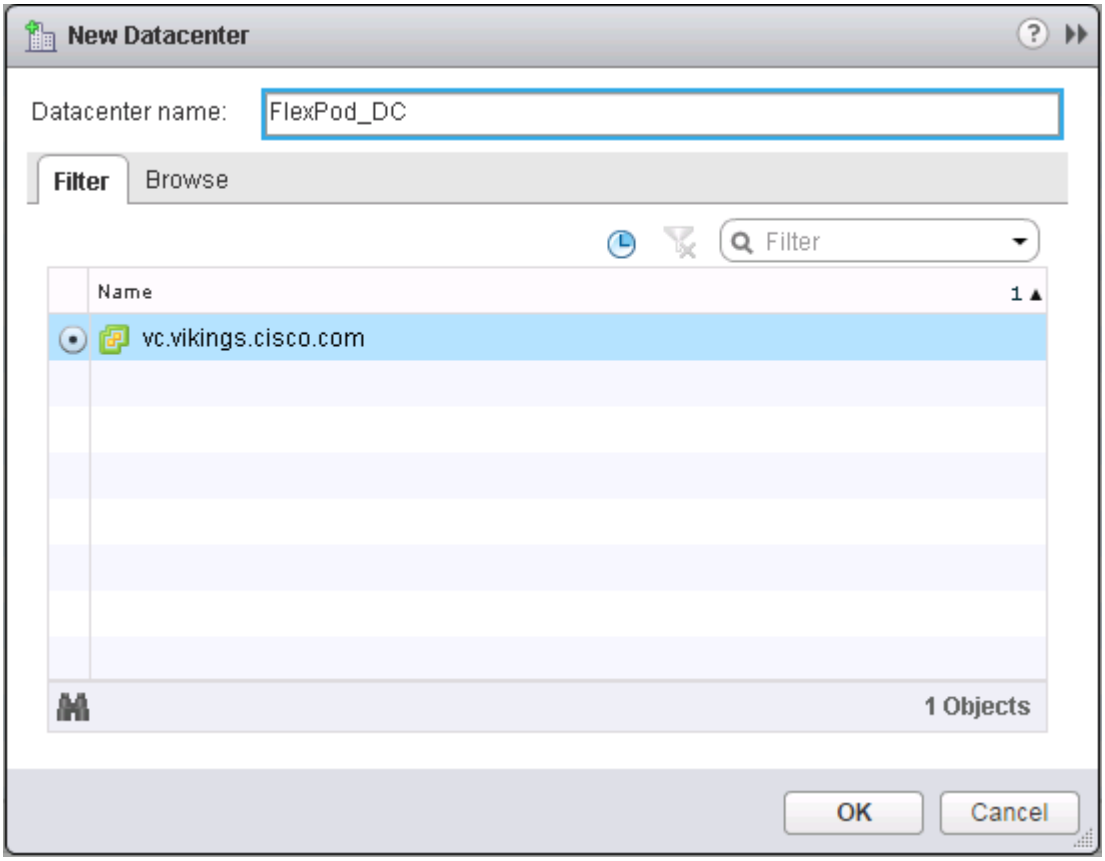

11. Right-click the data center FlexPod_DC in the list in the center pane. Click New Cluster.

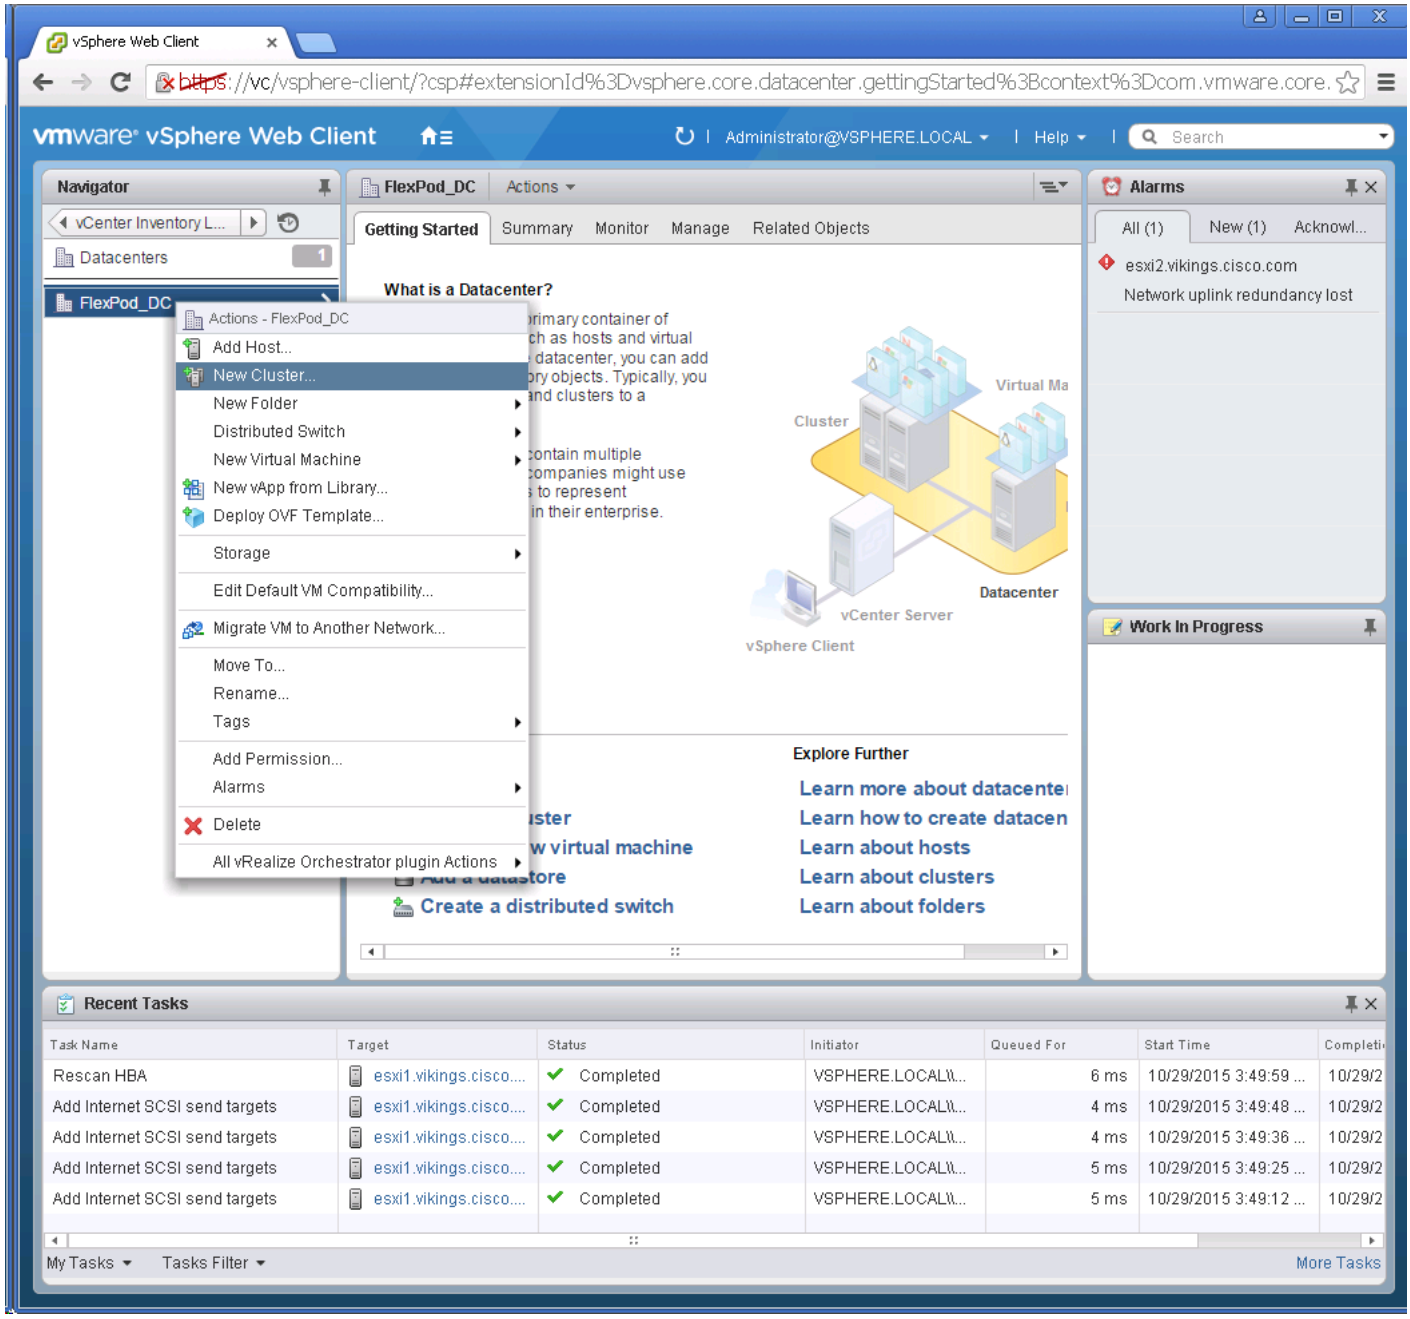

12. Name the cluster FlexPod_Management.

13. Check the box beside DRS. Leave the default values.

14. Check the box beside vSphere HA. Leave the default values.

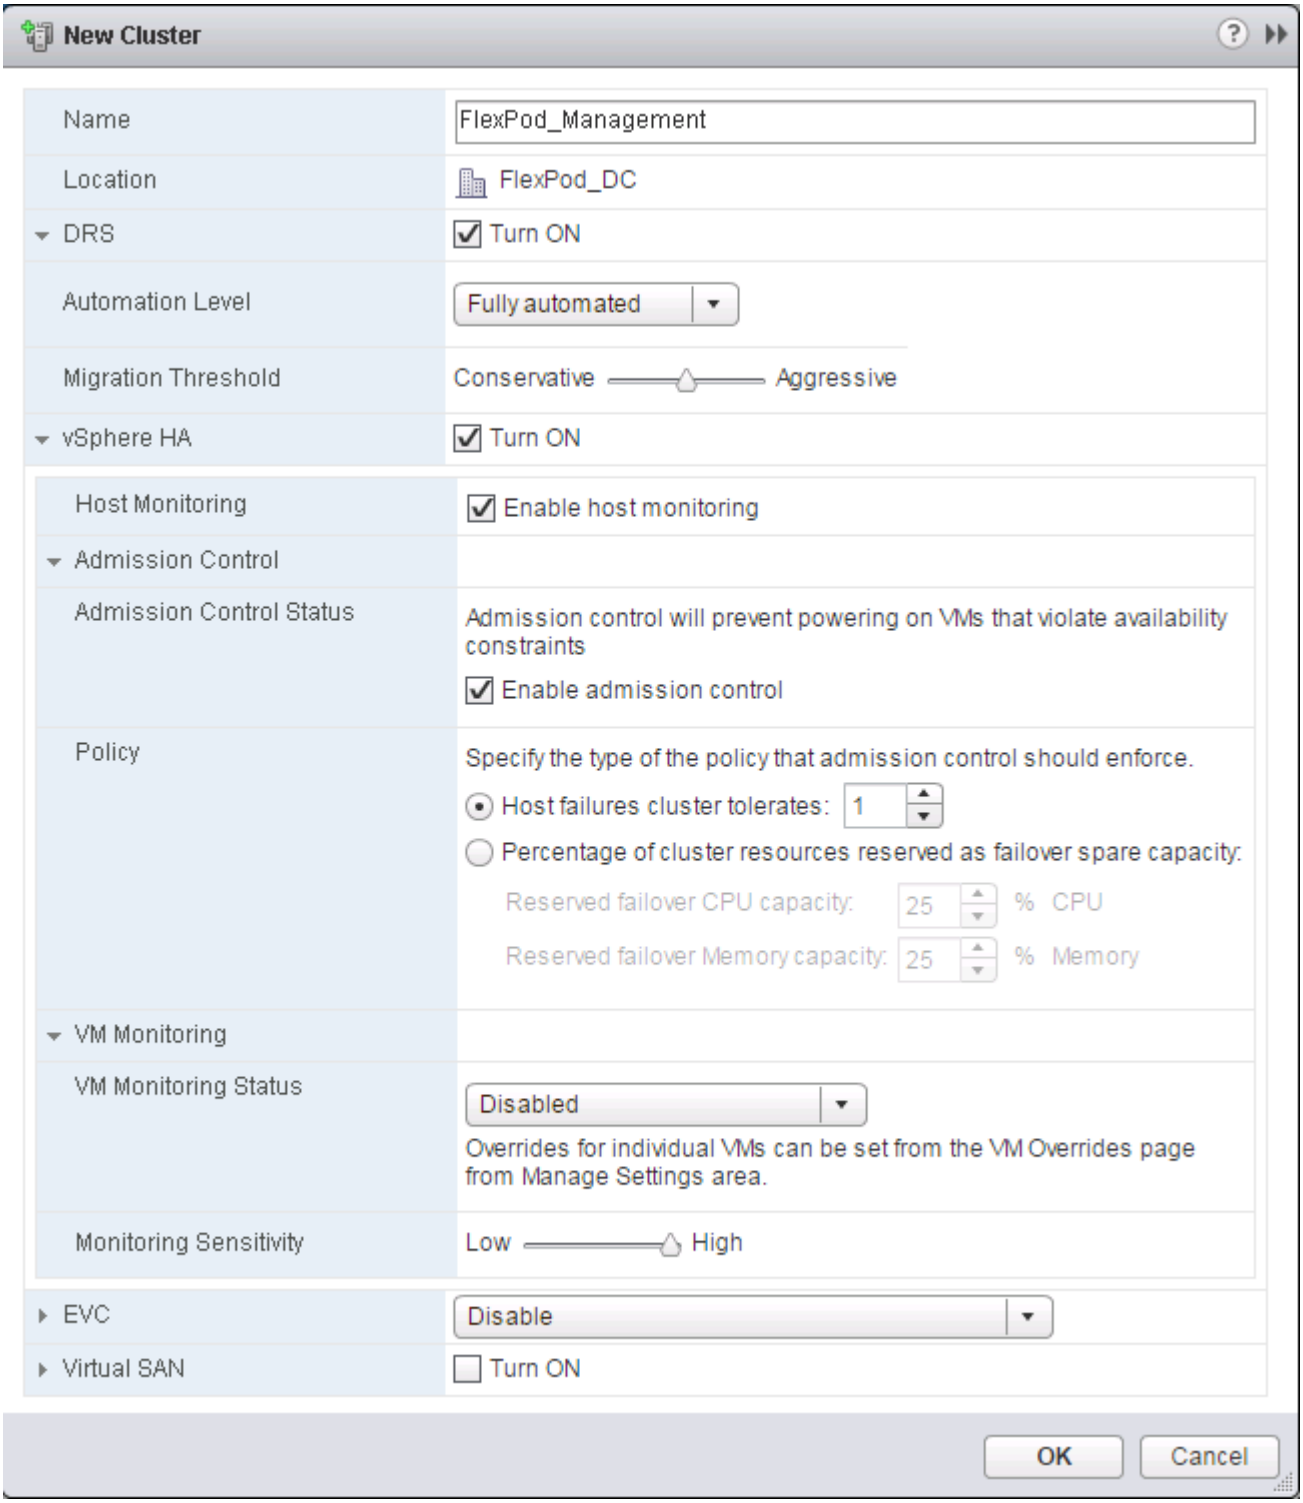

15. Click OK to create the new cluster.

16. On the left pane, double click the “FlexPod_DC”.

17. Click Clusters.

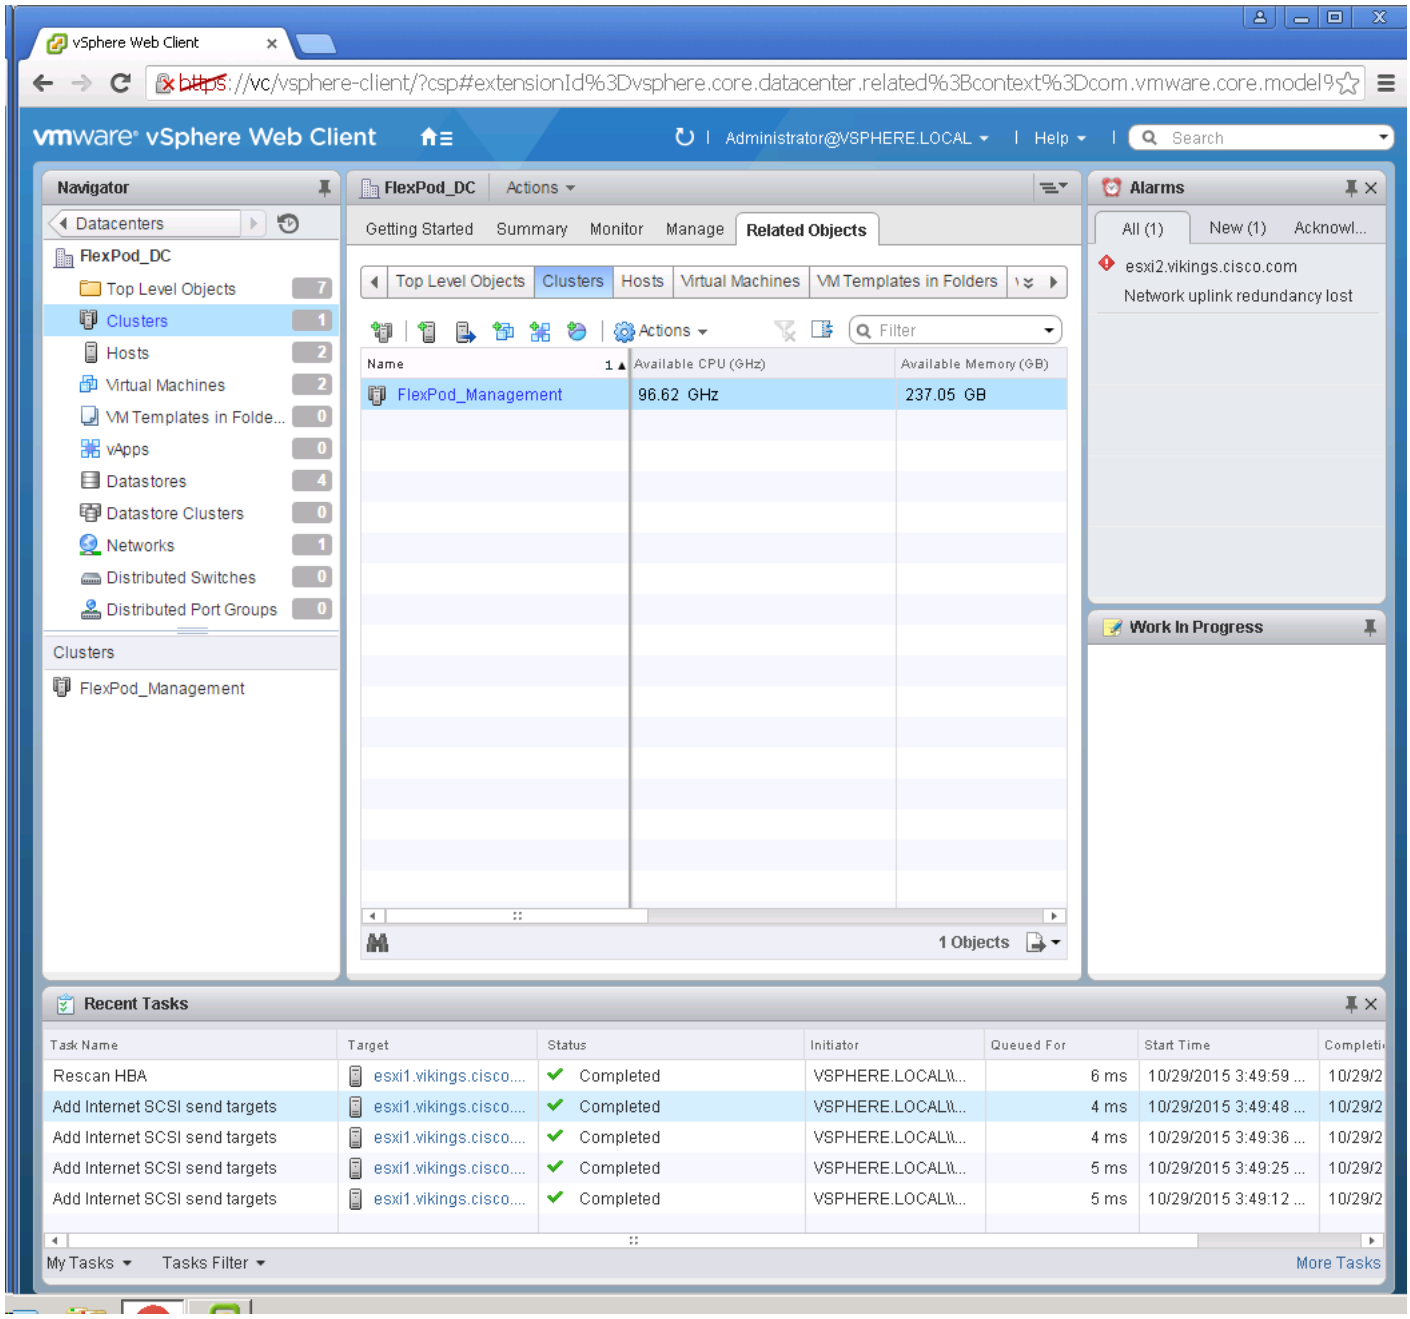

18. Under the Clusters pane, right click the “FlexPod_Management” and select Settings.

19. Select Configuration > General in the list on the left and select Edit to the right of General.

20. Select Datastore specified by host and click OK.

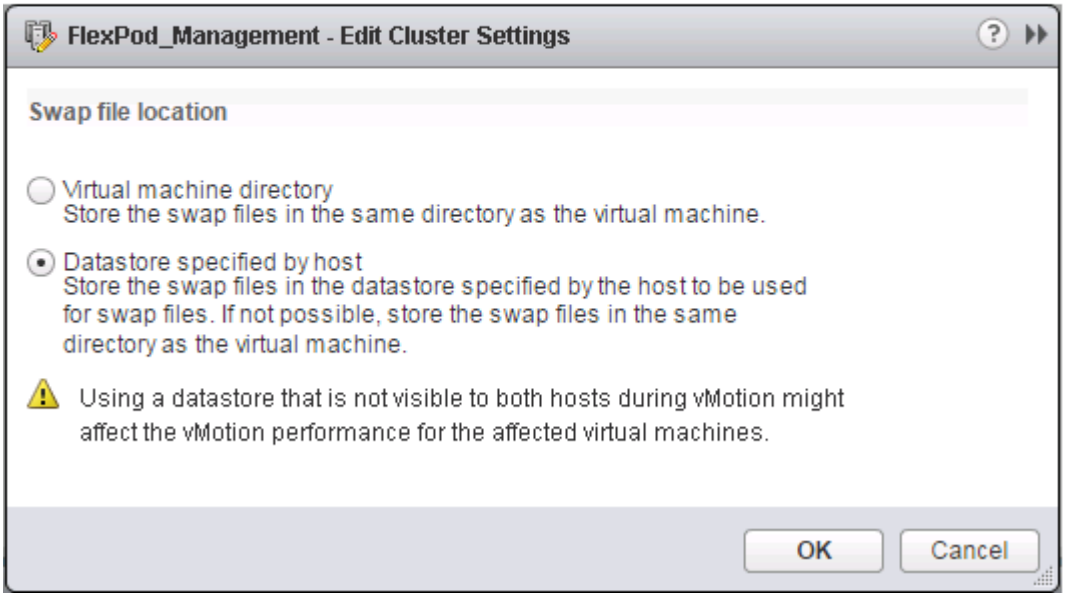

21. Under the Clusters pane, right click the “FlexPod_Management” and click Add Host.

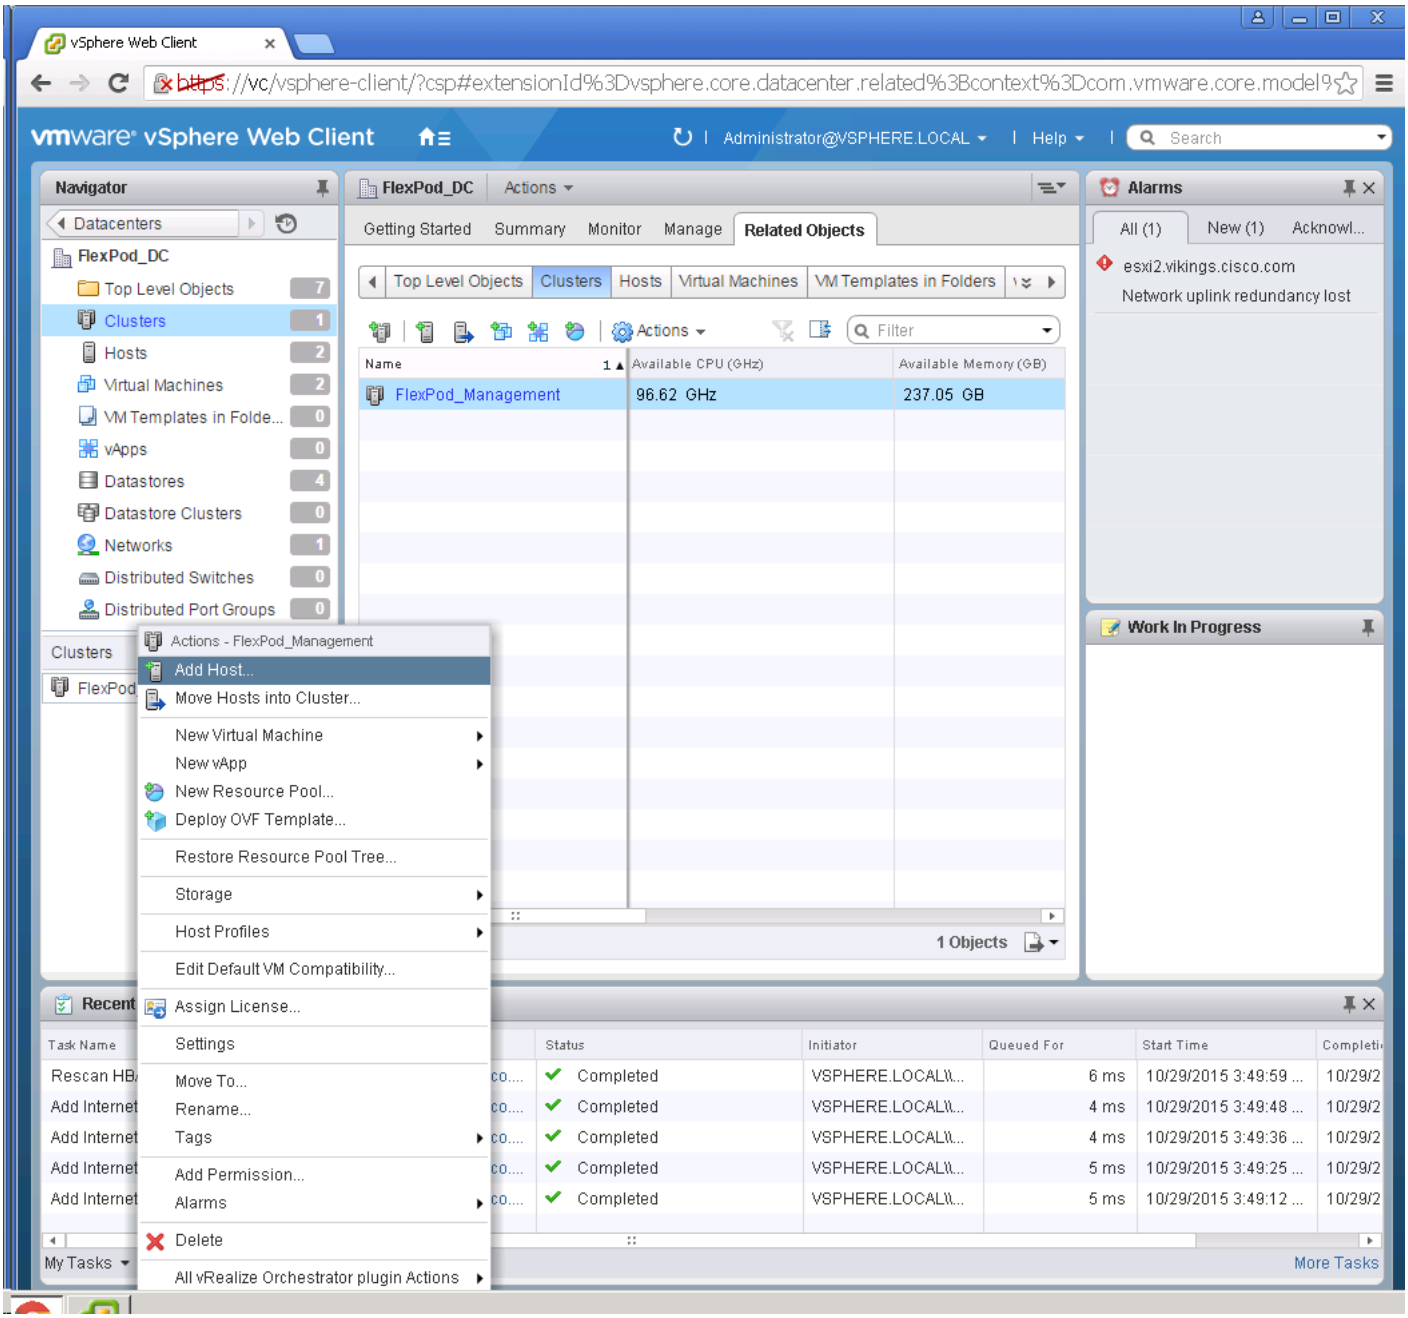

22. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.

23. Type root as the user name and the root password. Click Next to continue.

24. Click Yes to accept the certificate.

25. Review the host details and click Next to continue.

26. Assign a license and click Next to continue.

27. Click Next to continue.

28. Click Next to continue.

29. Review the configuration parameters. Then click Finish to add the host.

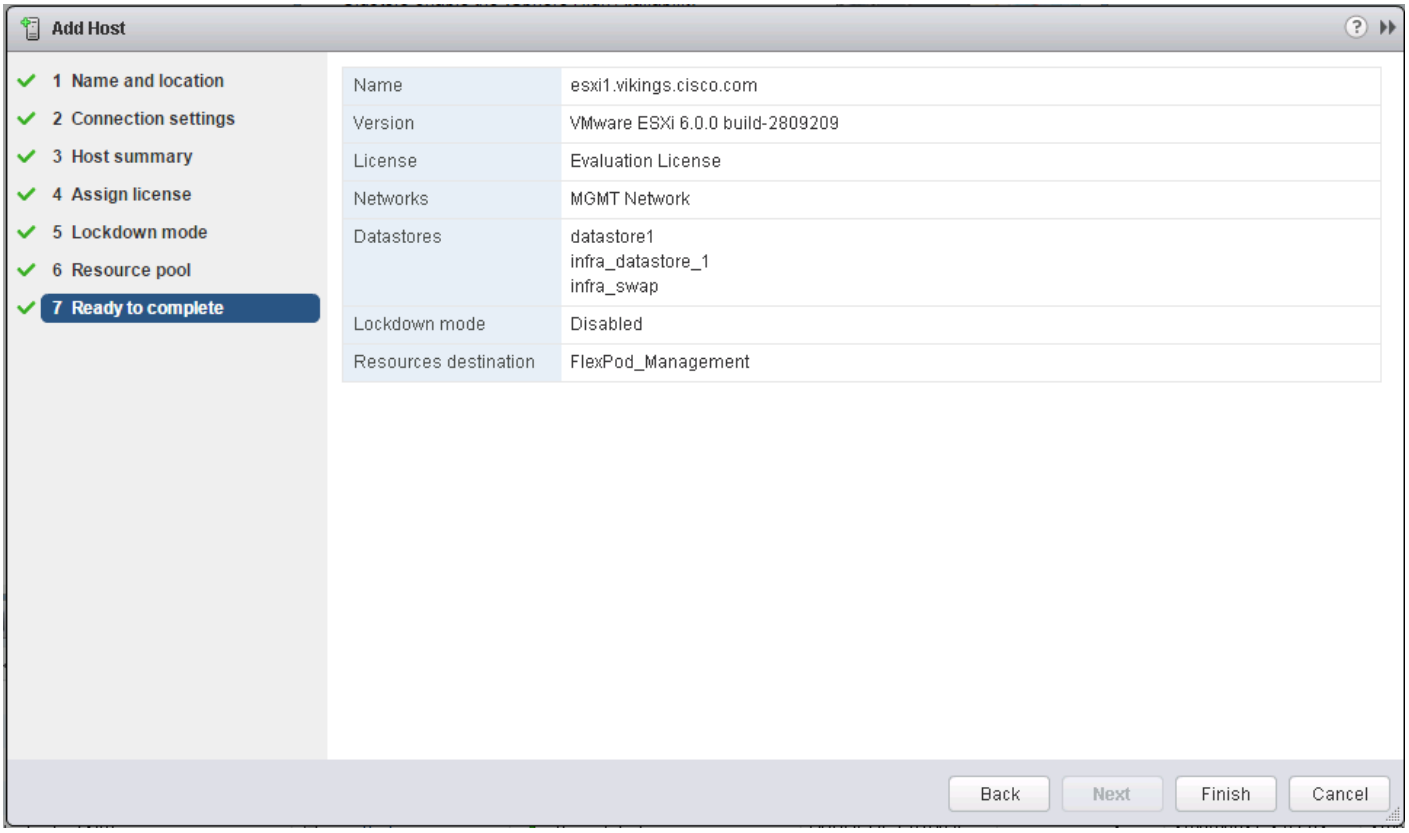

30. Repeat the steps 21 to 29 to add the remaining VMware ESXi hosts to the cluster.

Two VMware ESXi hosts will be added to the cluster.

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. To setup the ESXi Dump Collector, complete the following steps:

1. In the vSphere web client, select Home.
2. In the center pane, click System Configuration.
3. In the left hand pane, click 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. Connect to each ESXi host via ssh as root
9. Run the following commands:

```
esxcli system coredump network set --interface-name=vmk0 - -server-ipv4=10.1.156.100 --server-port=6500

esxcli  system coredump network set --enable=true

esxcli system coredump network check
```

FlexPod Cisco Nexus 1110-X and 1000V vSphere

This section provides detailed procedures for installing a pair of high-availability (HA) Cisco Nexus 1110-X Virtual Services Appliances (VSAs) in a FlexPod configuration. This validation effort used a preexisting management infrastructure to support the VSA devices and therefore does not document the cabling configuration.

Primary and standby Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) are installed on the 1110-Xs and Cisco Nexus 1000V distributed virtual switch (DVS) will be provisioned. This procedure assumes that the Cisco Nexus 1000V software version 5.2(1)SV3(1.5b) has been downloaded from Cisco Nexus 1000V Download Link and expanded. It is recommended to install software version 5.2(1)SP1(7.3) on the Nexus 1110-Xs using Cisco Nexus Cloud Services Platform Software Installation and Upgrade Guide. Additionally, this procedure assumes that Cisco Virtual Switch Update Manager (VSUM) version 1.5.3 has been downloaded from Cisco VSUM Download Link and expanded. This procedure also assumes that VMware vSphere 6.0 Enterprise Plus licensing is installed.

Configure CIMC Interface on Both Cisco Nexus 1110-Xs

Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure the Cisco Integrated Management Controller (CIMC) interface on the Cisco Nexus 1110-X VSAs, complete the following steps:

1. Using the supplied dongle, connect a monitor and USB keyboard to the KVM console port on the front of the Cisco Nexus 1110-X virtual appliance.
2. Reboot the virtual appliance.
3. Press F8 when prompted to configure the CIMC interface.
4. Using the spacebar, set the NIC mode to Dedicated.
5. Clear the checkbox for DHCP enabled.
6. Set the CIMC IP address (<<var_cimc_ip>>) in the out-of-band management VLAN.
7. Set the CIMC subnet mask (<<var_cimc_mask>>).

8. Set the CIMC gateway (<<var_cimc_gateway>>).
9. Set the NIC redundancy to None.
10. Set and reenter the CIMC default password (<<var_password>>).
11. Press F10 to save the configuration.
12. Continue pressing F5 until Network settings configured is shown.
13. Press Esc to reboot the virtual appliance.

Configure Serial over LAN for Both Cisco Nexus 1110-Xs

Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure serial over LAN on the Cisco Nexus 1110-X VSAs, complete the following steps:

1. Use a Web browser to open the URL at http://<<var_cimc_ip>>.
2. Log in to the CIMC with the admin user id and the CIMC default password (<<var_password>>).
3. In the left column, click Remote Presence.
4. Click the Serial over LAN tab.
5. Select the Enabled checkbox for Serial over LAN Properties.
6. From the Baud Rate drop-down menu, select 9600 bps.
7. Click Save Changes.

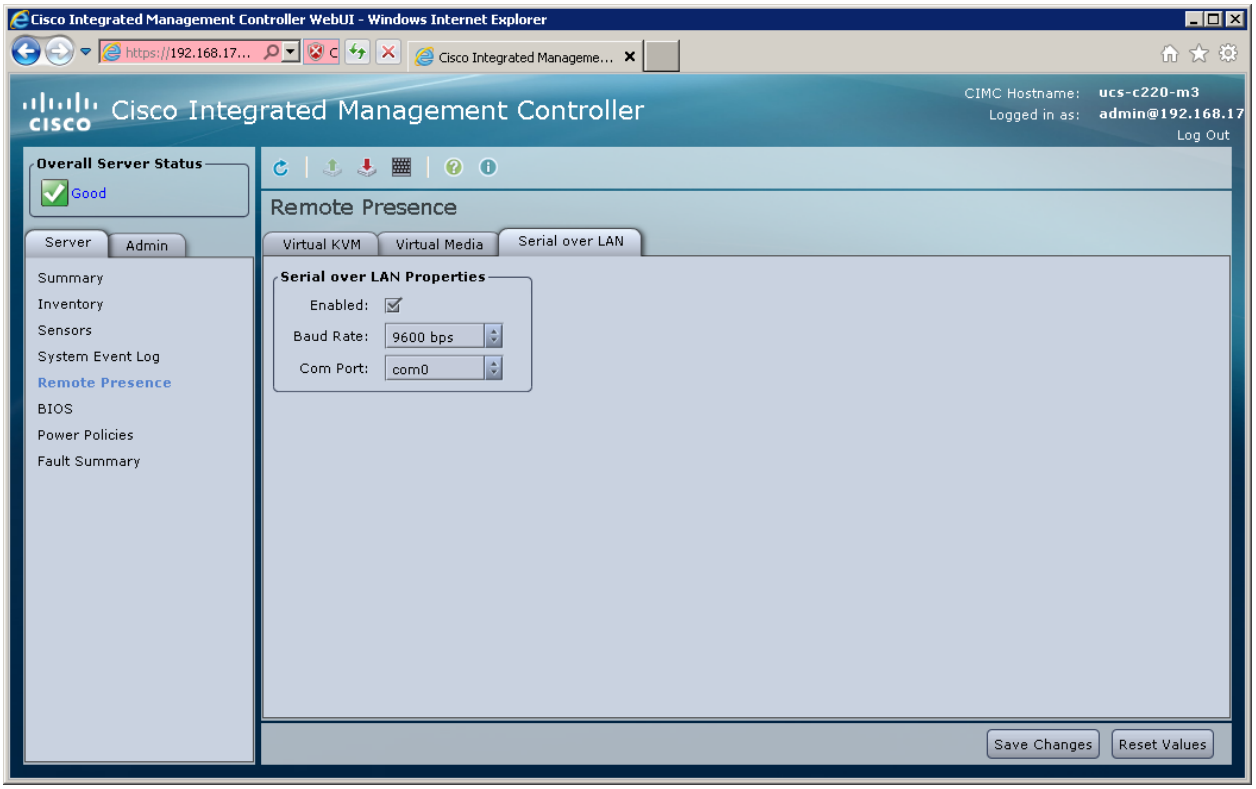

8. Log out of the CIMC Web interface.
9. Use an SSH client to connect to <<var_cimc_ip>> with the default CIMC user name and password.
10. Enter “connect host.”

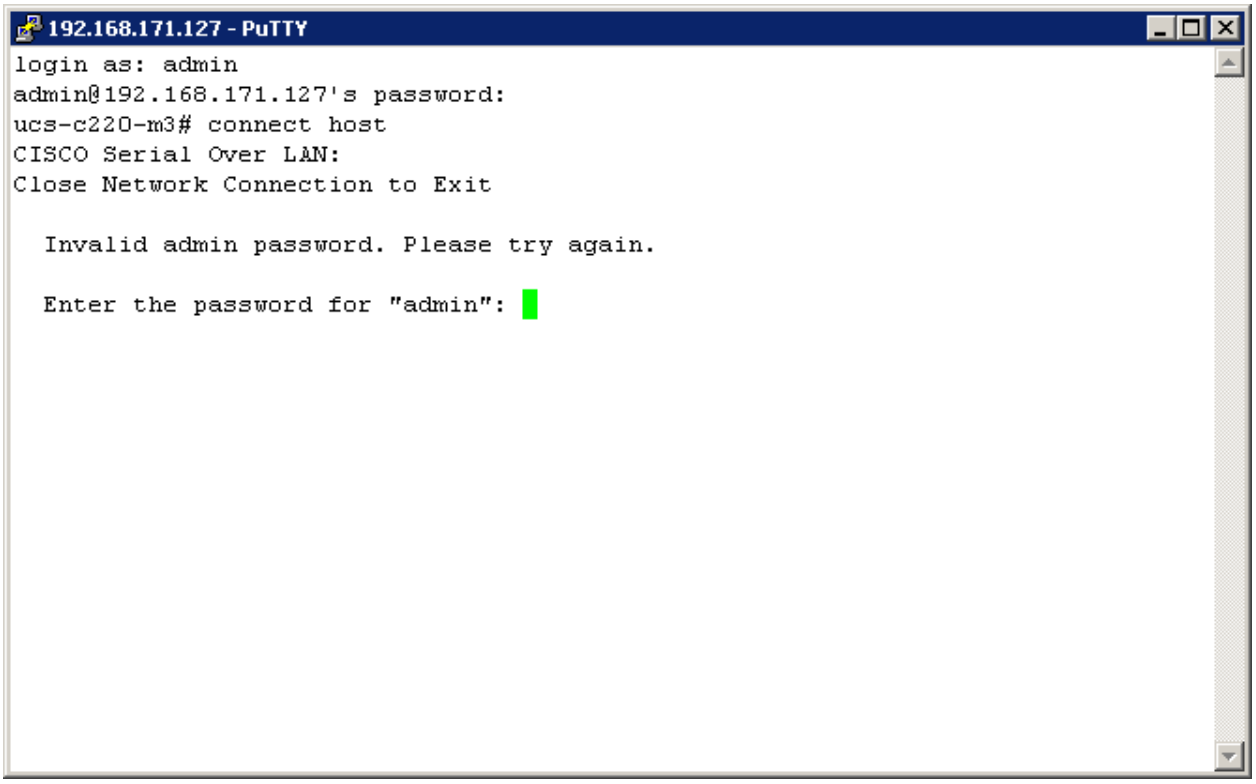

Configure Cisco Nexus 1110-X Virtual Appliances

Cisco Nexus 1110-X A

To configure Cisco Nexus 1110-X A, complete the following steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.

```
Enter the password for “admin”: <<var_password>>

Confirm the password for “admin”: <<var_password>>

Enter HA role[primary/secondary]: primary

Enter the domain id<1-4095>: <<var_vsa_domain_id>>

Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>

 Control Channel Setup.

Choose Uplink: < Gig:1,2 10Gig:7,8 NewPortChannel:0 >[0]: Enter

 Choose type of portchannel <ha/lacp>[ha]: lacp

PortChannel1 - Choose uplinks < Gig:1,2 10Gig:7,8 >[1,2]: 7,8

Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>

  Management Channel setup

  Choose Uplink: < Gig:1,2 Po1:9 NewPortChannel:0 >[9]: Enter

Would you like to enter the basic system configuration dialogue (yes/no): yes
```

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

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

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

Enter the VSA name : <<var_1110x_vsa>>

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

Mgmt0 IP address type V4/V6? (V4): V4

Mgmt0 IPv4 address : <<var_1110x_vsa_ip>>

Mgmt0 IPv4 netmask : <<var_1110x_vsa_mask>>

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

IPv4 address of the default gateway : <<var_1110x_vsa_gateway>>

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

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

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

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

Number of rsa key bits <768-2048> [1024]: Enter

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

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

  NTP server IPv4 address: <<var_switch_a_ntp_ip>>
```

2. Review the configuration summary. If everything is correct, enter no to skip editing the configuration.

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

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

3. The Cisco Nexus 1110-X saves the configuration and reboots. After reboot, log back in as admin.

Cisco Nexus 1110-X B

To configure the Cisco Nexus 1110-X B, complete the following steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.3

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

This is the same password that you entered on the primary Cisco Nexus 1110-X.

2. Enter the admin password again to confirm: <<var_password>>.

```
Enter HA role[primary/secondary]: secondary

Enter the domain id<1-4095>: <<var_vsa_domain_id>>
```

 This is the same domain id that you entered on the primary Cisco Nexus 1110-X.

```
Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>

  Control Channel Setup.

 Choose Uplink: < Gig:1,2 10Gig:7,8 NewPortChannel:0 >[0]: Enter

  Choose type of portchannel <ha/lacp>[ha]: lacp

 PortChannel1 - Choose uplinks < Gig:1,2 10Gig:7,8 >[1,2]: 7,8

Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>

  Management Channel setup

  Choose Uplink: < Gig:1,2 Po1:9 NewPortChannel:0 >[9]: Enter
```

3. The Cisco Nexus 1110-X saves the configuration and reboots.

Set Up the Primary Cisco Nexus 1000V VSM

Cisco Nexus 1110-X A

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, complete the following steps:

 These steps are completed from the primary Nexus 1110-X A

1. Continue periodically running the following command until module 2 (Cisco Nexus 1110-X B) has a status of ha-standby.

```
show module
```

2. Enter the global configuration mode and create a virtual service blade.

```
config t

virtual-service-blade VSM-1

dir /repository
```

3. If the desired Cisco Nexus 1000V ISO file (n1000v-dk9.5.2.1.SV3.1.5b.iso) is not present on the Cisco Nexus 1110-X, run the copy command to copy it to the Cisco Nexus 1110-X disk. You must place the file either on an FTP server or on a UNIX or Linux® machine (using scp) that is accessible from the Cisco Nexus 1110-X management interface. An example copy command from an FTP server is copy ftp://<<var_ftp_server>>/n1000v-dk9.5.2.1.SV3.1.5b.iso /repository/.

```
virtual-service-blade-type new n1000v-dk9.5.2.1.SV3.1.5b.iso

interface control vlan <<var_pkt-ctrl_vlan_id>>

interface packet vlan <<var_pkt-ctrl_vlan_id>>

enable primary

Enter vsb image:[n1000v-dk9.5.2.1.SV3.1.5b.iso] Enter
```

```
Enter domain id[1-4095]: <<var_vsm_domain_id>>
```

This domain ID should be different than the VSA domain ID.

```
Enter SVS Control mode (L2 / L3): [L3] Enter

Management IP version [V4/V6]: [V4] Enter

Enter Management IP address: <<var_vsm_mgmt_ip>>

Enter Management subnet mask: <<var_vsm_mgmt_mask>>

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

Enter HostName: <<var_vsm_hostname>>

Enter the password for 'admin': <<var_password>>
```

This password must be entered with only uppercase and lowercase letters. No special characters can be used in this password.

```
Do you want to continue with installation with entered details (Y/N)? [Y} Enter

copy run start
```

4. Run show virtual-service-blade summary. Continue periodically entering this command until the primary VSM-1 has a state of VSB POWERED ON.

5. Modify the management, control and packet interface and set PortChannel 1 as the uplink interface (if needed):

```
virtual-service-blade VSM-1

  interface control uplink PortChannel1

  interface management uplink PortChannel1

  interface packet uplink PortChannel1
```

Set Up the Secondary Cisco Nexus 1000V VSM

To set up the secondary Cisco Nexus 1000V VSM on Cisco Nexus 1110-X B, complete the steps in the following two subsections:

Cisco Nexus 1110-X A

```
enable secondary

Enter vsb image: [n1000v-dk9.5.2.1.SV3.1.5b.iso] Enter

Enter domain id[1-4095]: <<var_vsm_domain_id>>

Enter SVS Control mode (L2 / L3): [L3] Enter

Management IP version [V4/V6]: [V4] Enter

Enter Management IP address: <<var_vsm_ mgmt_ip>>

Enter Management subnet mask: <<var_vsm_ mgmt_mask>>
```

```
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>

Enter HostName: <<var_vsm_hostname>>

Enter the password for 'admin': : <<var_password>>

This password must be entered with only uppercase and lowercase letters. No special characters can be
used in this password. Do you want to continue installation with entered details (Y/N)? [Y}
```

Type show virtual-service-blade summary. Continue periodically entering this command until both the primary and secondary VSM-1s have a state of VSB POWERED ON and Roles are correctly identified.

```
copy run start
```

Install Cisco Virtual Switch Update Manager

VMware vSphere Web Client

To install the Cisco Virtual Switch Upgrade Manager from OVA in the VMware virtual environment, complete the following steps:

1. Log into the VMware vSphere Web Client.
2. In the pane on the right, click VMs and Templates.
3. In the center pane, select Actions > Deploy OVF Template.
4. Select Browse and browse to and select the Nexus1000v-vsum.2.0.ova file.
5. Click Open.
6. Click Next.

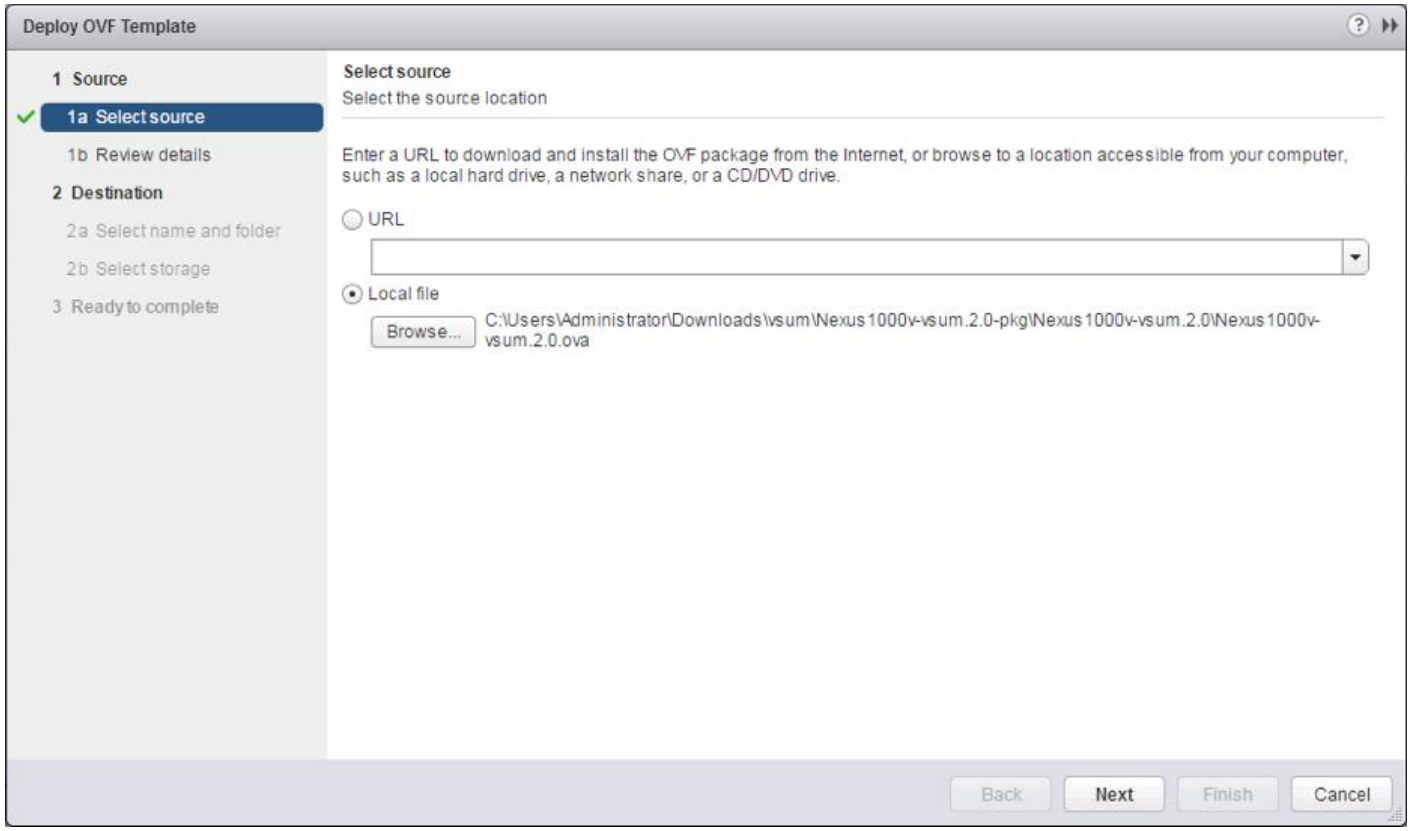

7. Review the details and click Next.

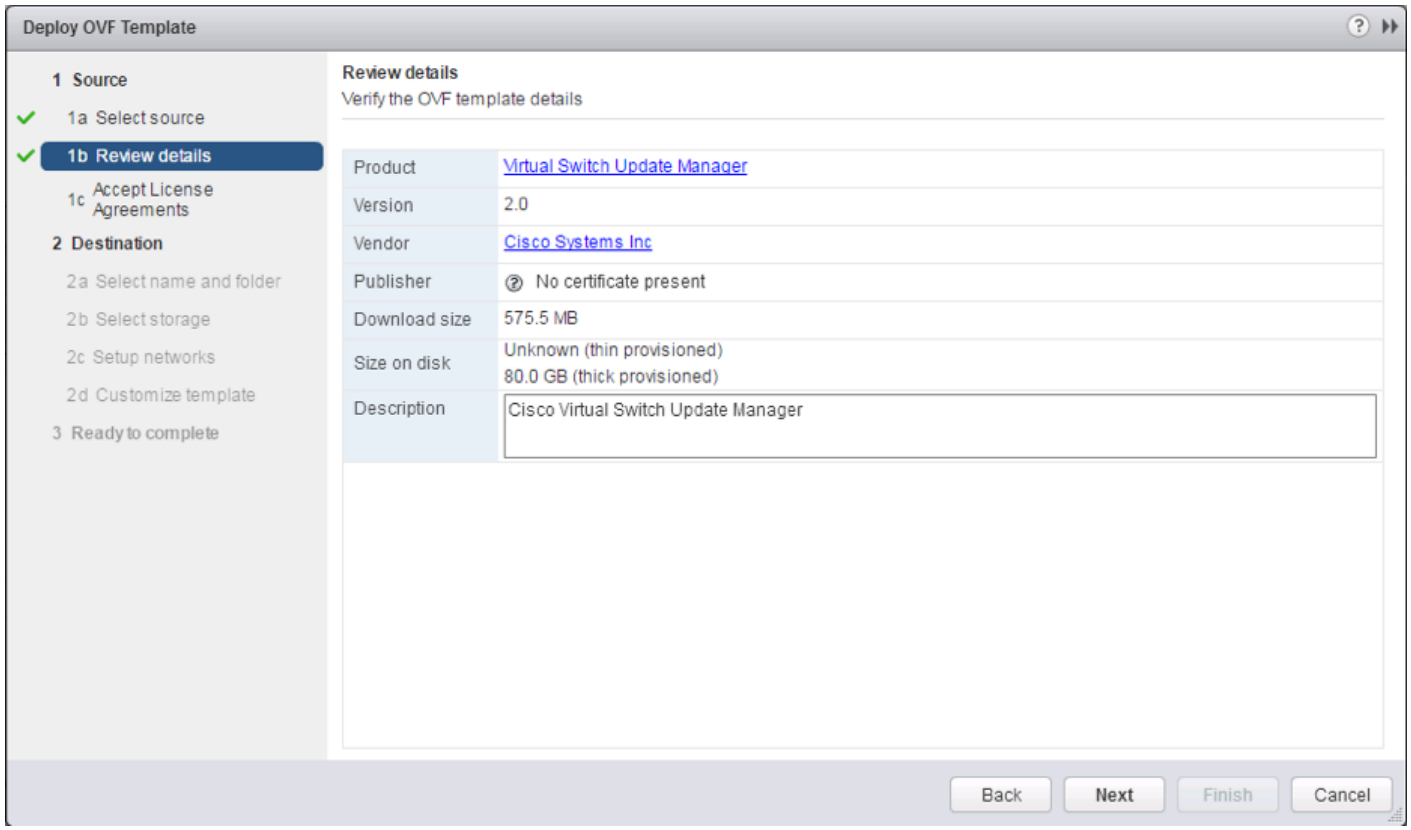

8. Click Accept to accept the License Agreement and click Next.

9. Name the Virtual Machine, select the FlexPod_DC datacenter and click Next.

10. Select the FlexPod_Management cluster and click Next.

11. Select infra_datastore_1 and the Thin Provision virtual disk format and click Next.

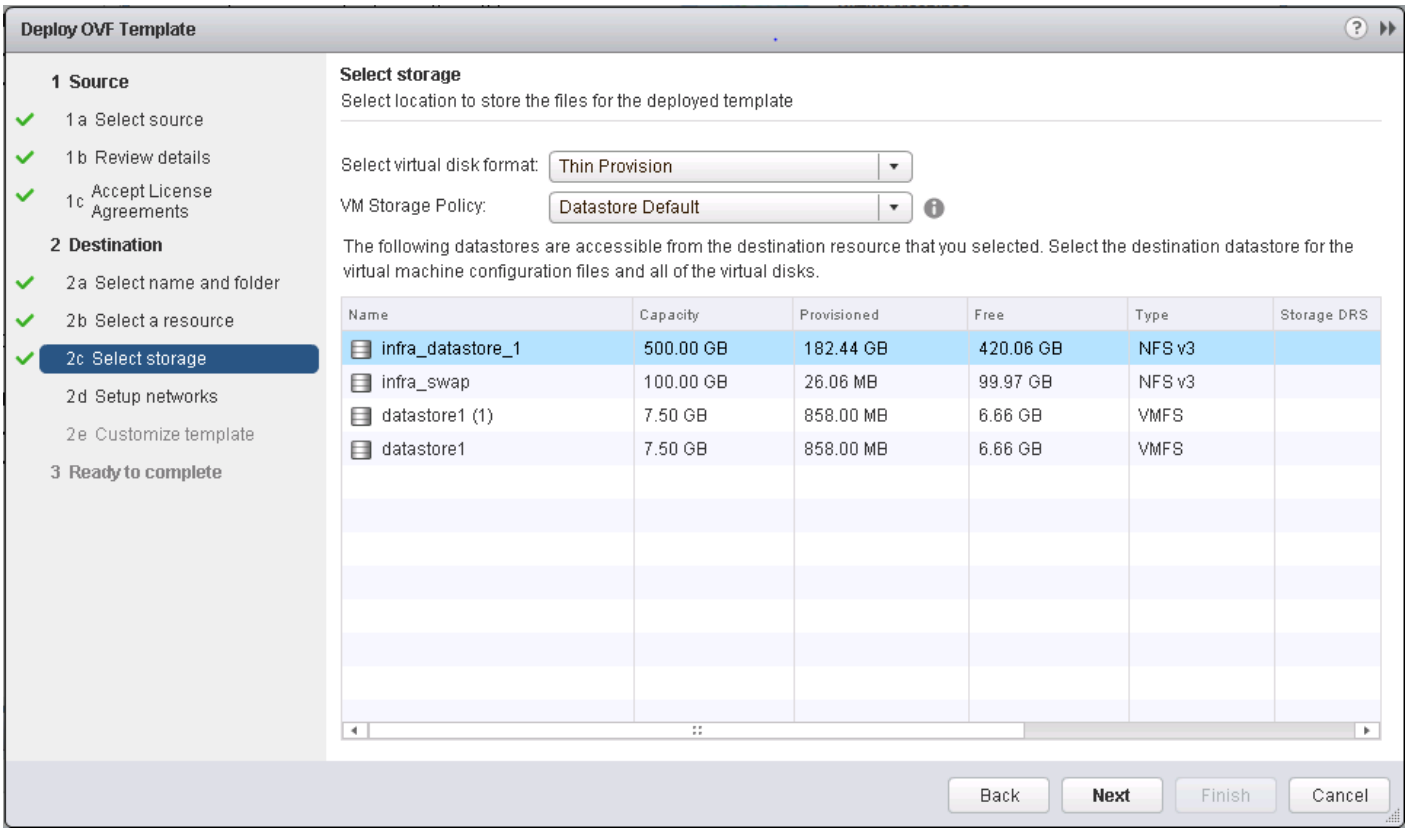

12. Select the IB-MGMT-VLAN Network and click Next.

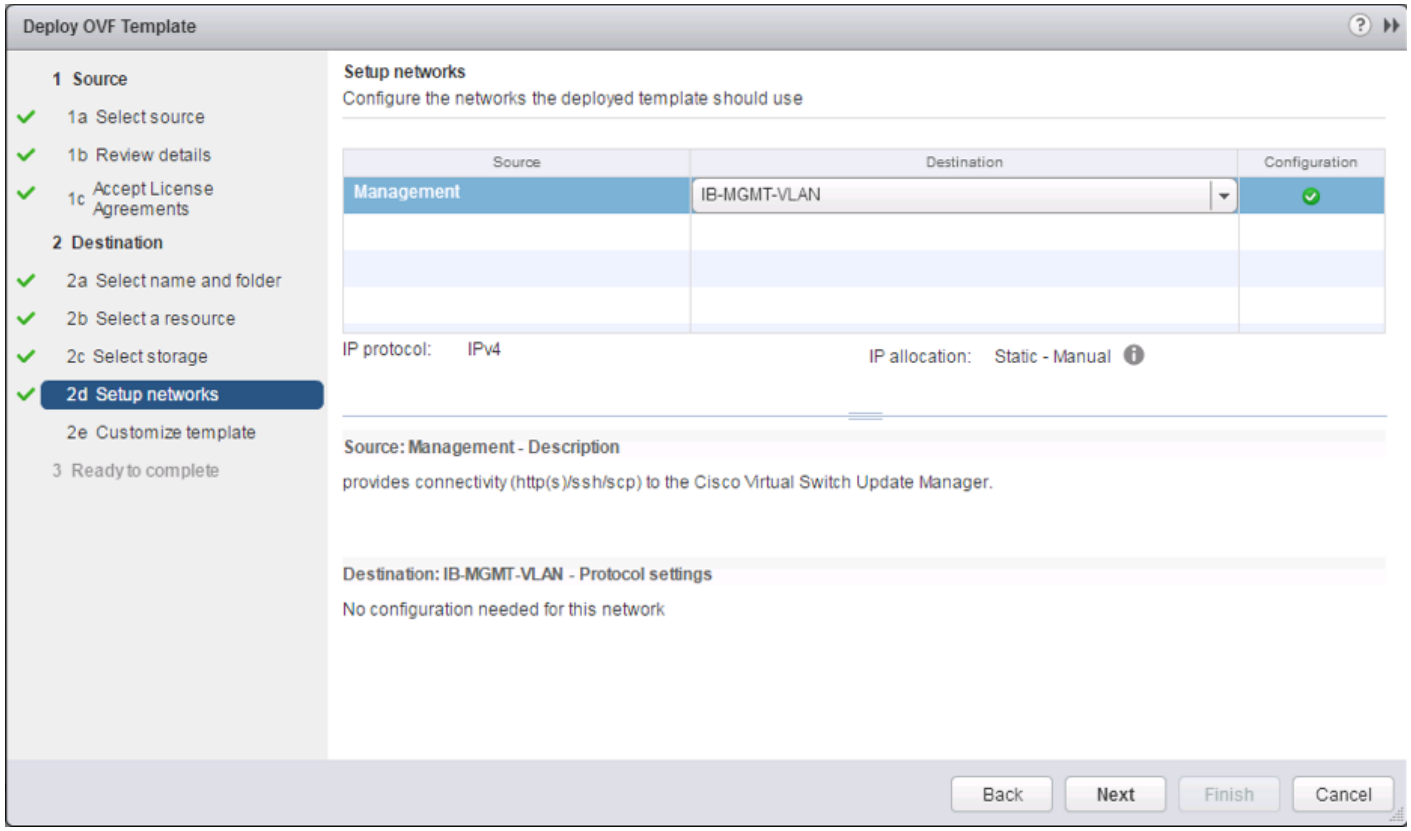

13. Fill in the Networking Properties within the Customize template dialog.

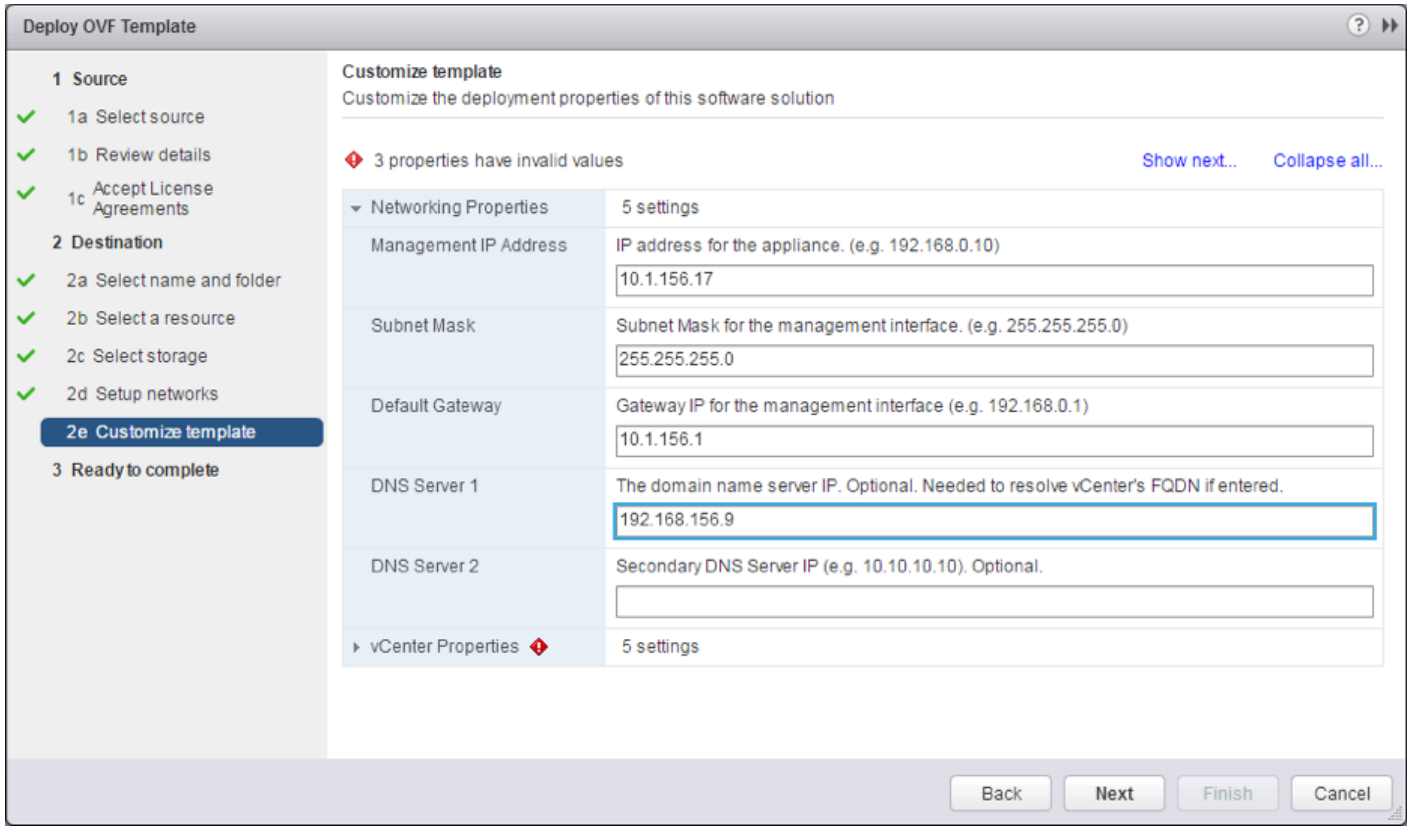

14. Expand the vCenter Properties and fill in vCenter address, Username, and Password fields.

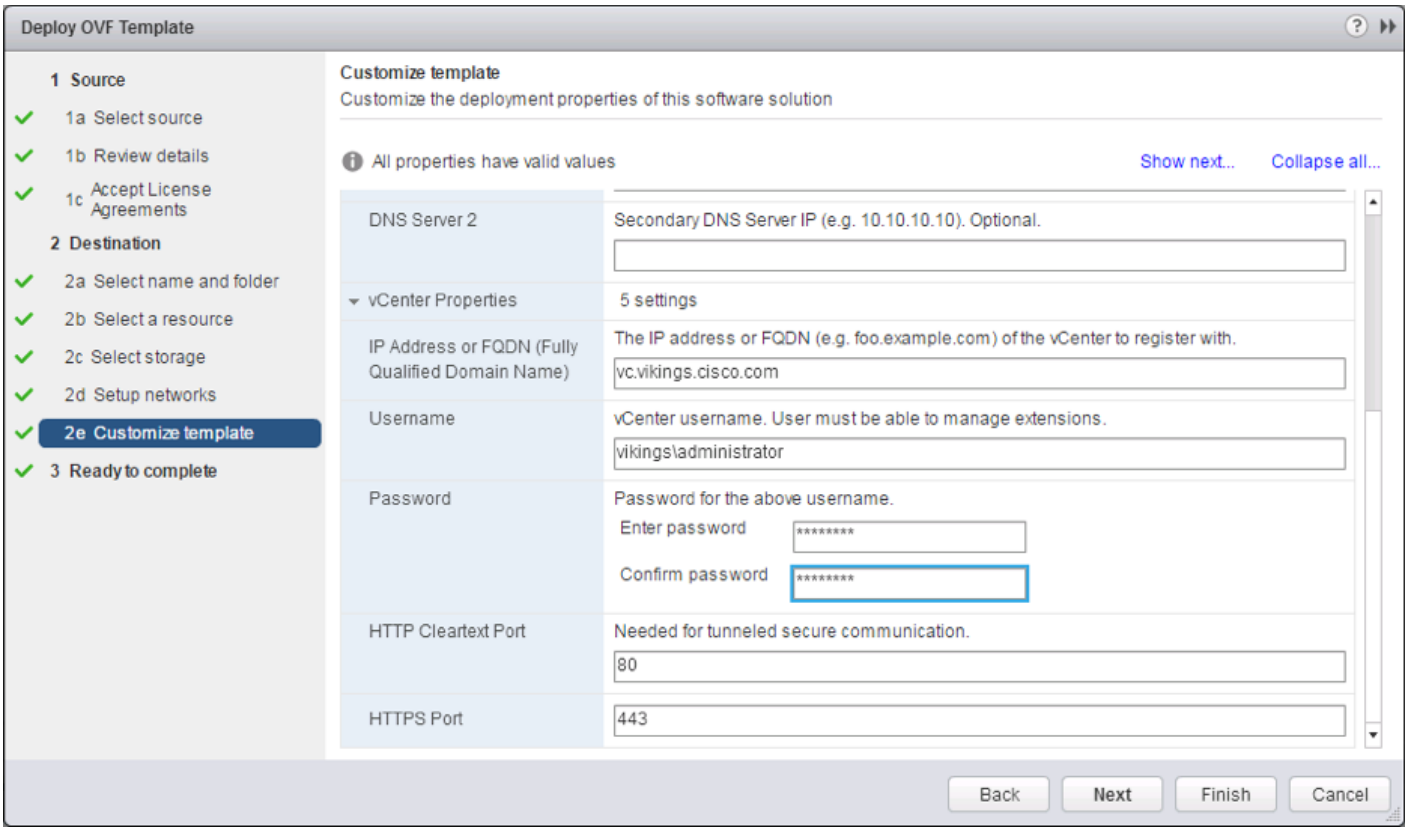

15. Click Next.

16. Review all settings and click Finish.

17. Wait for the Deploy OVF template task to complete.

18. Select the Home button in VMware vSphere Web Client and select Hosts and Clusters.

19. Expand the FlexPod_Management cluster and select the Virtual Switch Update Manager VM from the Summary tab.

20. Click the Console graphic at the top left of the Summary tab. If a security warning pops up, click Allow.

21. If a security certificate warning pops up, click Connect Anyway.

22. Power on the Virtual Switch Update Manager VM.

23. Once the VM has completely booted up, log out and log back into the VMware vSphere Web Client.

Register the Cisco Nexus 1000V in VMware vCenter

VMware vSphere Web Client

To register the Cisco Nexus 1000V, complete the following steps:

1. After logging back into the VMware vSphere Web Client, Cisco Virtual Switch Update Manager should now appear under the Home tab. Select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Click Install.
4. In the pane on the right, select FlexPod_DC.
5. Under Nexus1000v Switch Deployment Process, select I already have a control place (VSM) deployed.
6. Enter the IP Address of the VSM and the admin password.
7. Click Finish.
8. Click the Home button.
9. Select Cisco Virtual Switch Update manager.
10. Under Basic tasks, select Nexus 1000v.
11. Click Configure.
12. In the pane on the right, select FlexPod_DC.
13. The Nexus 1000v Switch should appear under the *Choose an associated Distributed Virtual Switch* section.

Perform Base Configuration of the Primary VSM

SSH Connection to Primary VSM

To perform the base configuration of the primary VSM, complete the following steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands.

```
config t

ntp server <<var_switch_a_ntp_ip>> use-vrf management

ntp server <<var_switch_b_ntp_ip>> use-vrf management

vlan <<var_ib-mgmt_vlan_id>>

name IB-MGMT-VLAN

vlan <<var_nfs_vlan_id>>

name NFS-VLAN

vlan <<var_vmotion_vlan_id>>
```

```
name vMotion-VLANvlan <<var_vm-traffic_vlan_id>>

name VM-Traffic-VLAN

vlan <<var_native_vlan_id>>

name Native-VLAN

vlan <<var_iscsi_a_vlan_id>>

name iSCSI-A-VLAN

vlan <<var_iscsi_b_vlan_id>>

name iSCSI-B-VLAN

vlan <<var_pkt-ctrl_vlan_id>>

name Pkt-Ctrl-VLAN

exit

port-profile type ethernet system-uplink

vmware port-group

switchport mode trunk

switchport trunk native vlan <<var_native_vlan_id>>

switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>

channel-group auto mode on mac-pinning

no shutdown

system vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>
```

Any VLAN that has a VMKernal port should be assigned as a system vlan on both the uplink and the vEthernet ports of the virtual switch.

```
system mtu 9000

state enabled

port-profile type ethernet iscsi-a-uplink

vmware port-group

switchport mode trunk

switchport trunk native vlan <<var_iscsi_a_vlan_id>>

switchport trunk allowed vlan <<var_iscsi_a_vlan_id>>

no shutdown

system vlan <<var_iscsi_a_vlan_id>>

system mtu 9000
```

```
state enabled
port-profile type ethernet iscsi-b-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_b_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
system mtu 9000
state enabled
port-profile type vethernet IB-MGMT-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet NFS-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_nfs_vlan_id>>
no shutdown
system vlan <<var_nfs_vlan_id>>
state enabled
port-profile type vethernet vMotion-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vmotion_vlan_id>>
no shutdown
system vlan <<var_vmotion_vlan_id>>
state enabled
port-profile type vethernet VM-Traffic-VLAN
vmware port-group
```

```
switchport mode access
switchport access vlan <<var_vm-traffic_vlan_id>>
no shutdown
system vlan <<var_vm-traffic_vlan_id>>
state enabled
port-profile type vethernet n1kv-L3
capability l3control
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet iSCSI-A-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_a_vlan_id>>
no shutdown
system vlan <<var_iscsi_a_vlan_id>>
state enabled
port-profile type vethernet iSCSI-B-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
state enabled
exit
copy run start
```

Add VMware ESXi Hosts to Cisco Nexus 1000V

VMware vSphere Web Client

To and VMware ESXi hosts, complete the following steps:

1. Back in the Vmware vSphere Web Client, from the Home tab, select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Select Configure.
4. Select the FlexPod_DC datacenter on the right.
5. Select the VSM on the lower right.
6. Click Manage.
7. In the center pane, select the Add Host tab.
8. Expand the FlexPod_Management ESXi Cluster and select both FlexPod Management Hosts.
9. Click Suggest.
10. Scroll down to Physical NIC Migration and expand each ESXi host.
11. On both hosts, unselect vmnic0, and select vmnic1. For vmnic1, select the system-uplink Profile. Select vmnic2 and select the iscsi-a-uplink Profile. Select vmnic3 and select the iscsi-b-uplink Profile.

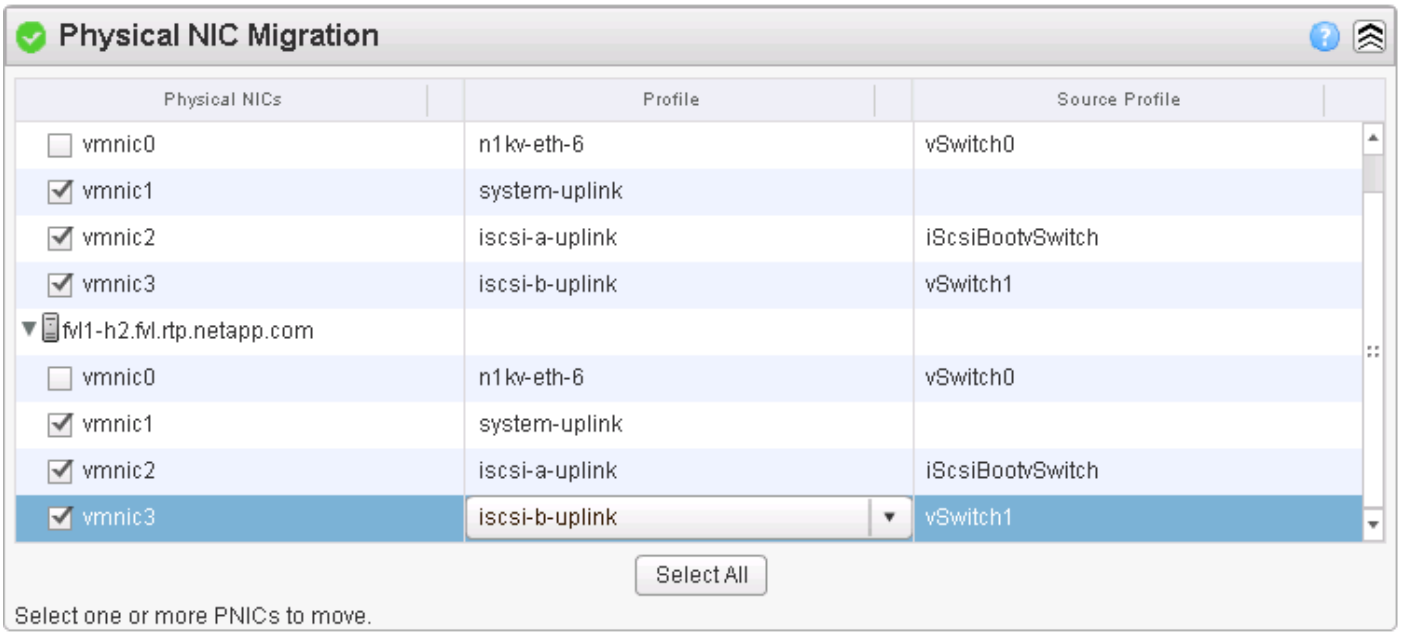

12. Scroll down to VM Kernel NIC Setup and expand both ESXi hosts.
13. All VMkernel ports should already have the appropriate checkboxes selected.

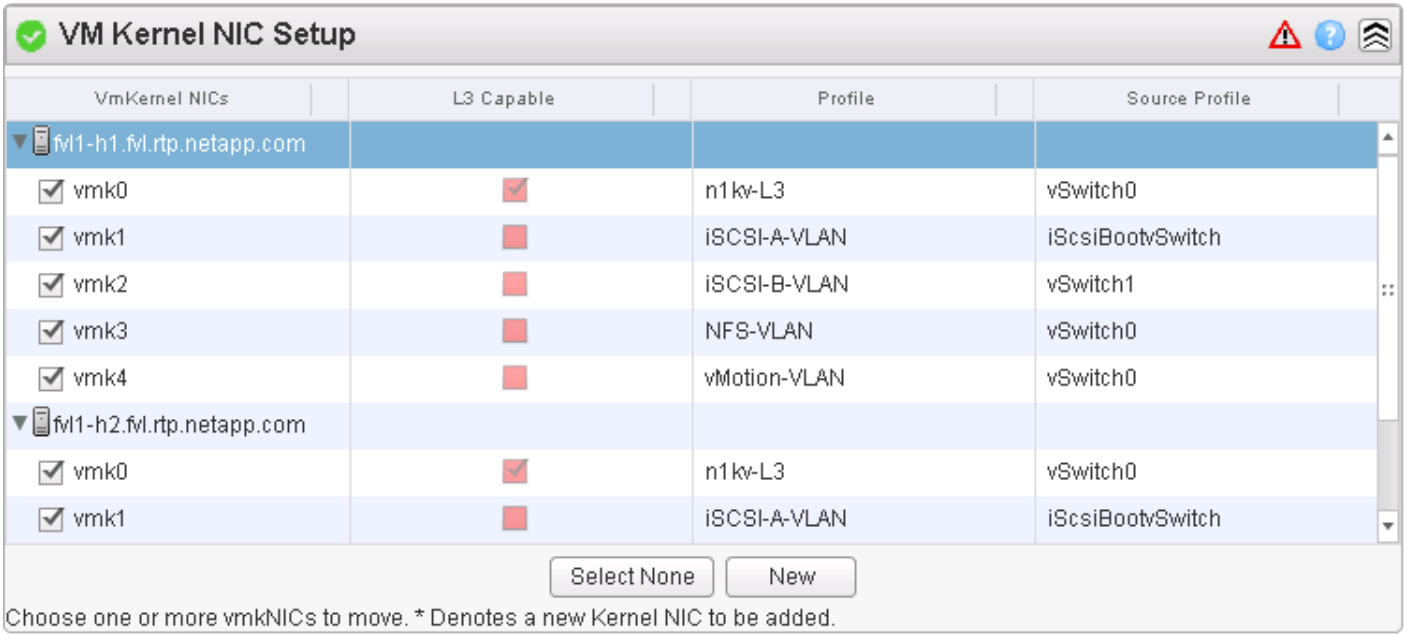

14. Scroll down to VM Migration and expand both ESXi hosts.

15. Select the IB-MGMT-VLAN profile for the VSUM and vCenter Virtual Machines.

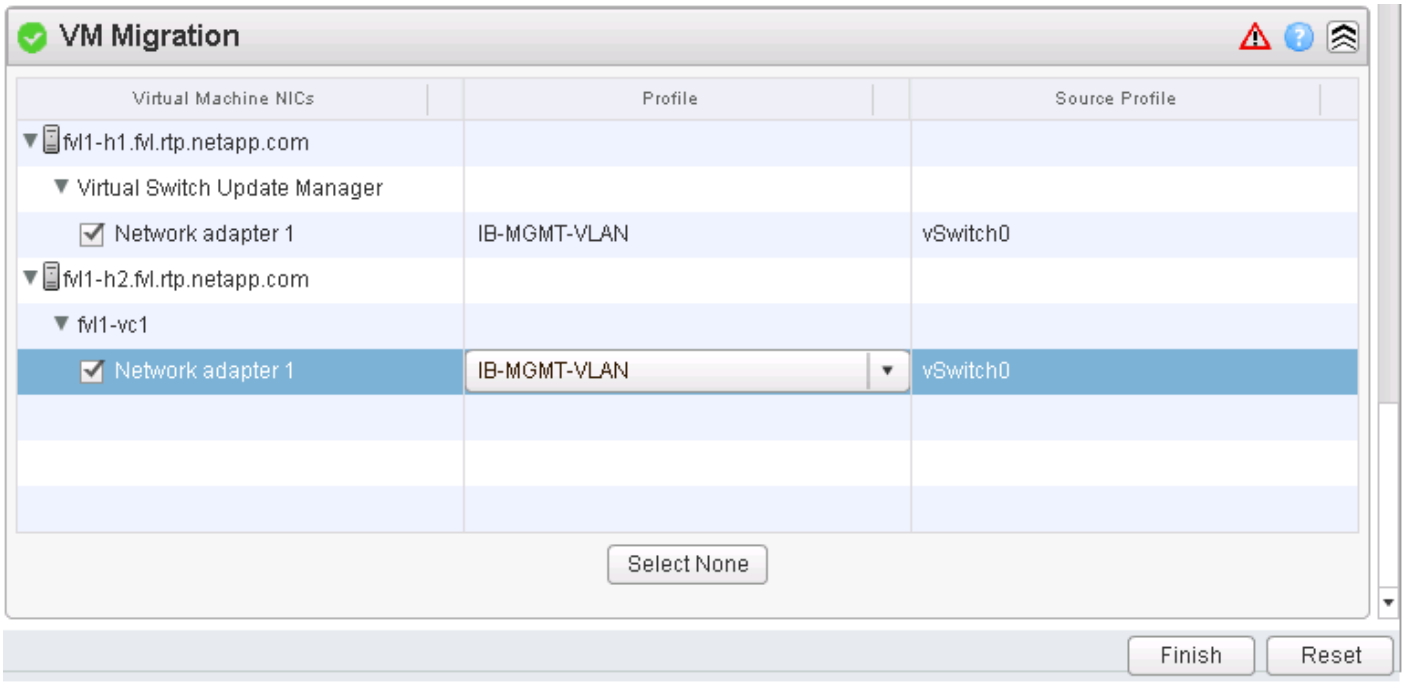

16. Click Finish.

The progress of the virtual switch installation can be monitored from the c# interface.

Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V

To migrate the ESXi host redundant network ports, complete the following steps:

1. In the VMware vSphere Web Client window, select Home > Hosts and Clusters.
2. On the left expand the Datacenter and cluster, and select the first VMware ESXi host.
3. In the center pane, select the Manage tab, then select Networking.
4. Select vSwitch0. All of the port groups on vSwitch0 should be empty. Click the red X under Virtual switches to delete vSwitch0.
5. Click Yes to remove vSwitch0. It may be necessary to refresh the Web Client to see the deletion.
6. Delete iScsiBootvSwitch and vSwitch1.
7. The Nexus 1000V VSM should now be the only virtual switch. Select it and select the third icon above it under Virtual switches (Manage the physical network adapters connected to the selected switch).
8. Click the green plus sign to add an adapter.
9. For UpLink03, select the system-uplink port group and make sure vmnic0 is the Network adapter. Click OK.
10. Click OK to complete adding the Uplink. It may be necessary to refresh the Web Client to see the addition.
11. Repeat this procedure for the second ESXi host.
12. From the SSH client that is connected to the Cisco Nexus 1000V, run show interface status to verify that all interfaces and port channels have been correctly configured.

```
fvl1-vsm.fvl.rtp.netapp.com - PuTTY

--------------------------------------------------------------------------------
Port          Name               Status    Vlan/     Duplex  Speed   Type
                                           Segment
--------------------------------------------------------------------------------
mgmt0         --                 connected routed    full    1000    --
Eth3/1        --                 connected trunk     full    10G     --
Eth3/2        --                 connected trunk     full    10G     --
Eth3/3        --                 connected trunk     full    10G     --
Eth3/4        --                 connected trunk     full    10G     --
Eth4/1        --                 connected trunk     full    10G     --
Eth4/2        --                 connected trunk     full    10G     --
Eth4/3        --                 connected trunk     full    10G     --
Eth4/4        --                 connected trunk     full    10G     --
Po1           --                 connected trunk     full    10G     --
Po2           --                 connected trunk     full    10G     --
Veth1         VMware VMkernel, v connected 3071      auto    auto    --
Veth2         VMware VMkernel, v connected 3076      auto    auto    --
Veth3         VMware VMkernel, v connected 3077      auto    auto    --
Veth4         VMware VMkernel, v connected 3074      auto    auto    --
Veth5         VMware VMkernel, v connected 3072      auto    auto    --
Veth6         Virtual Switch...e connected 3071      auto    auto    --
Veth7         VMware VMkernel, v connected 3071      auto    auto    --
--More--
```

13. Run show module and verify that the two ESXi hosts are present as modules.

```
fvl1-vsm.fvl.rtp.netapp.com - PuTTY

fvl1-vsm# sho module
Mod  Ports  Module-Type                      Model              Status
---  -----  --------------------------------  -----------------  ------------
1    0      Virtual Supervisor Module         Nexus1000V         active *
2    0      Virtual Supervisor Module         Nexus1000V         ha-standby
3    1022   Virtual Ethernet Module           NA                 ok
4    1022   Virtual Ethernet Module           NA                 ok

Mod  Sw                 Hw
---  -----------------  ------------------------------------------------
1    5.2(1)SV3(1.5b)    0.0
2    5.2(1)SV3(1.5b)    0.0
3    5.2(1)SV3(1.5b)    VMware ESXi 6.0.0 Releasebuild-2494585 (6.0)
4    5.2(1)SV3(1.5b)    VMware ESXi 6.0.0 Releasebuild-2494585 (6.0)

Mod  Server-IP        Server-UUID                          Server-Name
---  ---------------  -----------------------------------  --------------------
1    172.20.71.44     NA                                   NA
2    172.20.71.44     NA                                   NA
3    172.20.71.26     722201d0-e027-e511-0000-000000110001 fvl1-h1
4    172.20.71.27     722201d0-e027-e511-0000-000000110002 fvl1-h2

* this terminal session
fvl1-vsm#
```

14. Run copy run start.

Cisco Nexus 1000V vTracker

SSH Connection to Primary VSM

The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network. environment. To connect SSH to the primary VSM, complete the following steps:

1. From an ssh interface connected to the Cisco Nexus 1000V VSM, enter the following:

```
config t

feature vtracker

copy run start

show vtracker upstream-view

show vtracker vm-view vnic

show vtracker vm-view info

show vtracker module-view pnic

show vtracker vlan-view
```

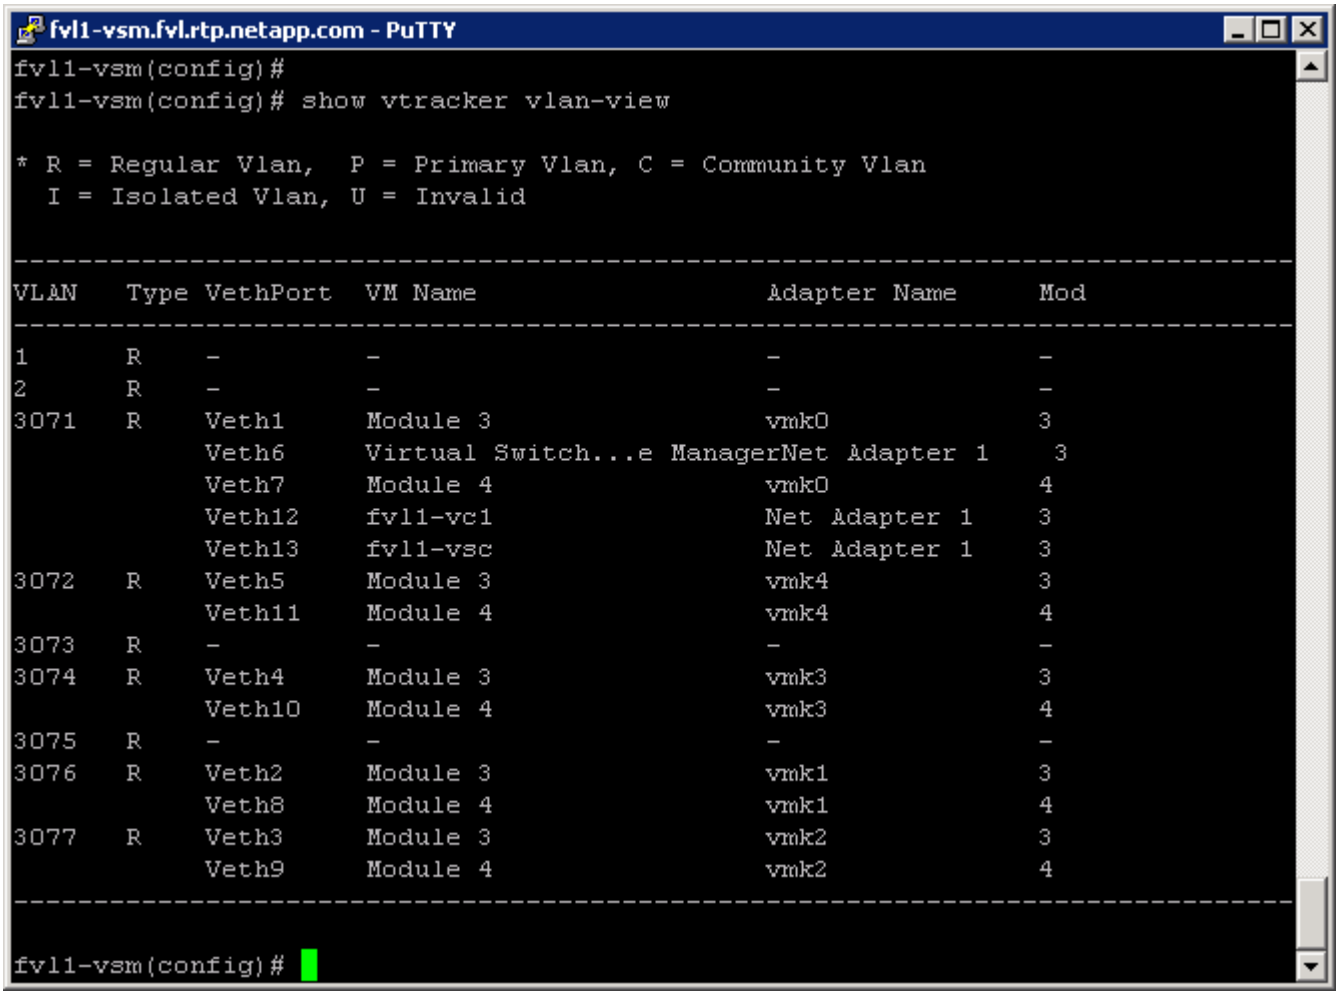
fvl1-vsm.fvl.rtp.netapp.com - PuTTY
fvl1-vsm(config)#
fvl1-vsm(config)# show vtracker vlan-view
* R = Regular Vlan, P = Primary Vlan, C = Community Vlan
I = Isolated Vlan, U = Invalid
VLAN Type VethPort VM Name Adapter Name Mod
1 R - - - -
2 R - - - -
3071 R Veth1 Module 3 vmk0 3
Veth6 Virtual Switch...e ManagerNet Adapter 1 3
Veth7 Module 4 vmk0 4
Veth12 fvl1-vc1 Net Adapter 1 3
Veth13 fvl1-vsc Net Adapter 1 3
3072 R Veth5 Module 3 vmk4 3
Veth11 Module 4 vmk4 4
3073 R - - - -
3074 R Veth4 Module 3 vmk3 3
Veth10 Module 4 vmk3 4
3075 R - - - -
3076 R Veth2 Module 3 vmk1 3
Veth8 Module 4 vmk1 4
3077 R Veth3 Module 3 vmk2 3
Veth9 Module 4 vmk2 4
fvl1-vsm(config)#

FlexPod Management Tools Setup

Cisco UCS Performance Manager

This section describes the deployment and initial configuration of Cisco UCS Performance Manager within a FlexPod.

For full requirements and installation options, download and review the Cisco UCS Performance Manager Installation Guide, Release 2.0.0.

Cisco UCS Performance Manager OVA Deployment

1. Download the Cisco UCS Performance Manager OVA file from the Cisco UCS Performance Manager site to your workstation.
2. Use the VMware vSphere Web Client to log in to vCenter as root, or as a user with superuser privileges, and then select Hosts and Clusters from the Home view.

The standard vSphere Client can be used for installation, but instructions will vary slightly.

3. Right-click the Cluster to deploy to, and select Deploy OVF Template from the pulldown options.

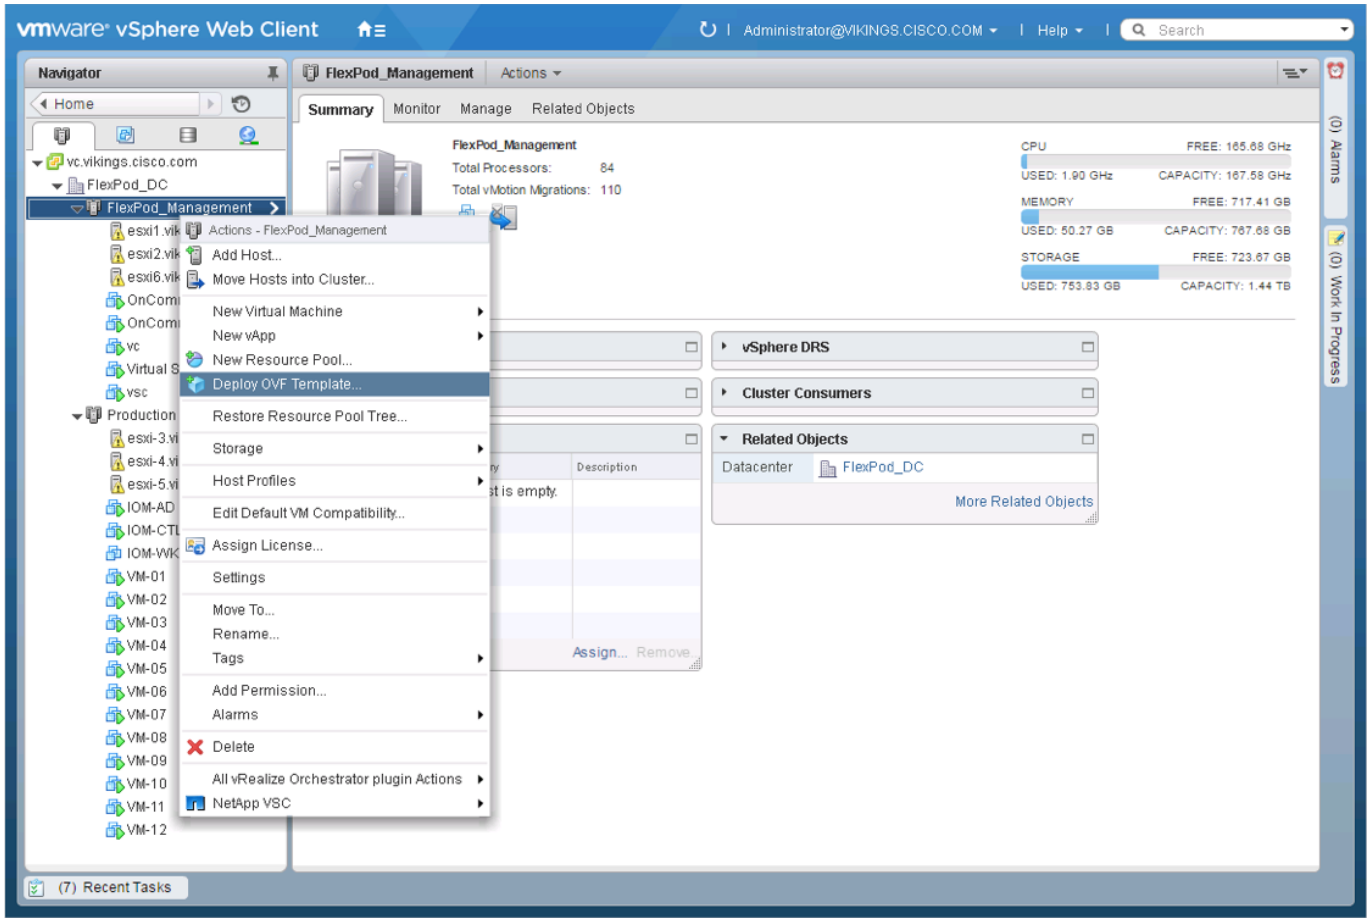

4. In the Selectsource panel, specify the path of the Cisco UCS Performance Manager package as a Local file and browse to find it's location, and then click Next.

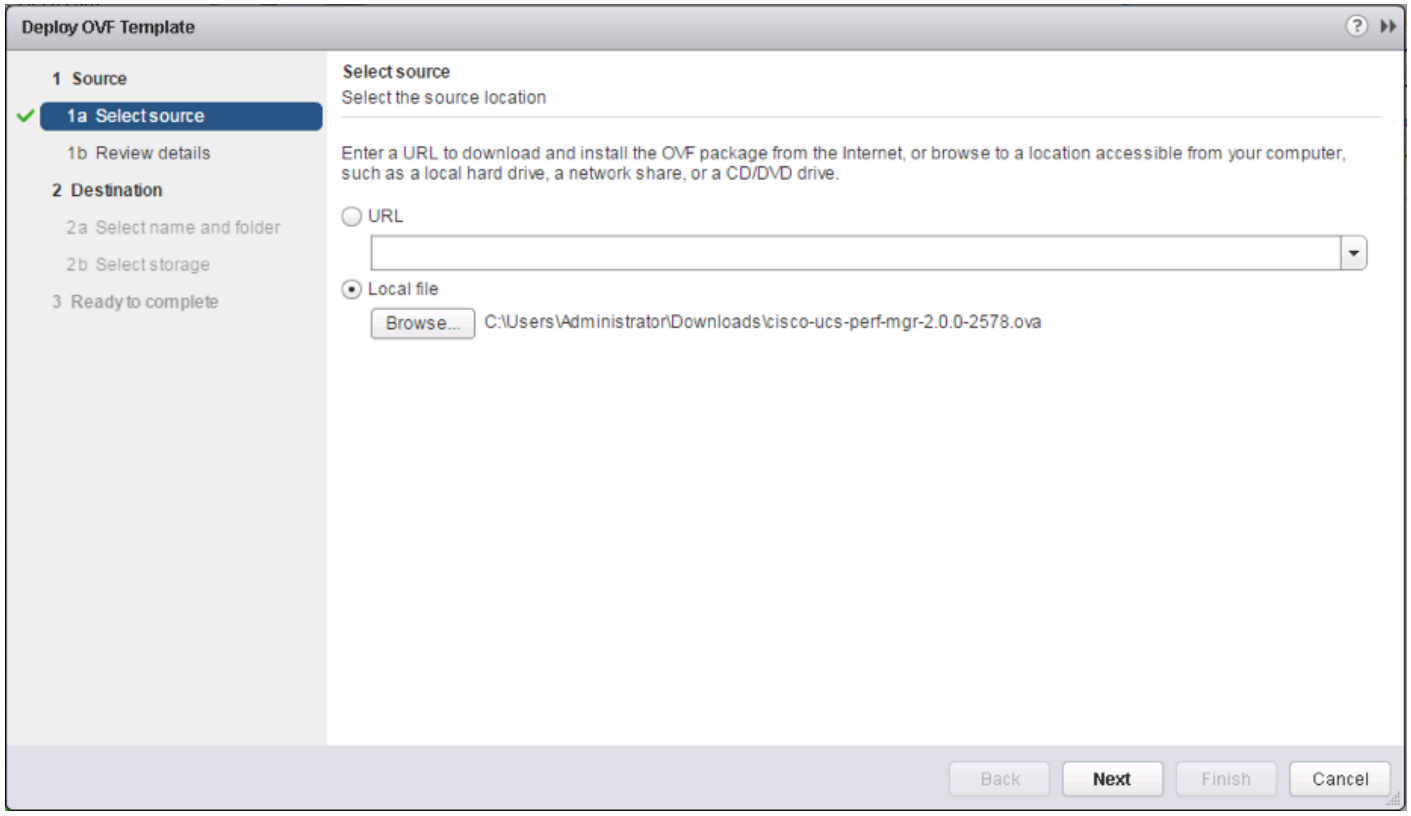

5. Click Next to continue past the Review details pane.
6. In the Select name and folder pane, accept the default name, or specify one appropriate to your environment, pick the data center to deploy to and alternately a folder within that datacenter.

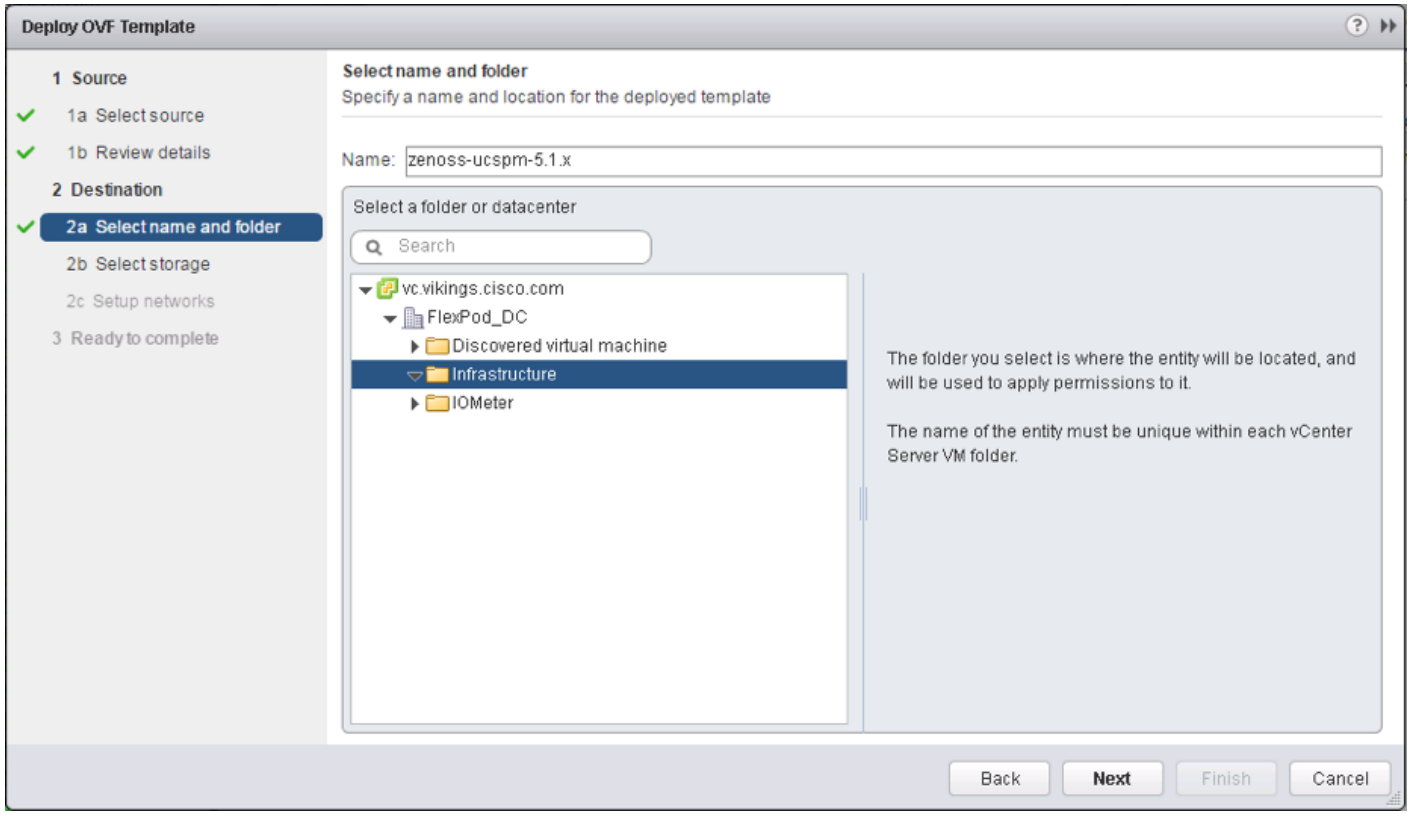

7. Click Next.
8. Specify the datastore to deploy to within the Select storage pane

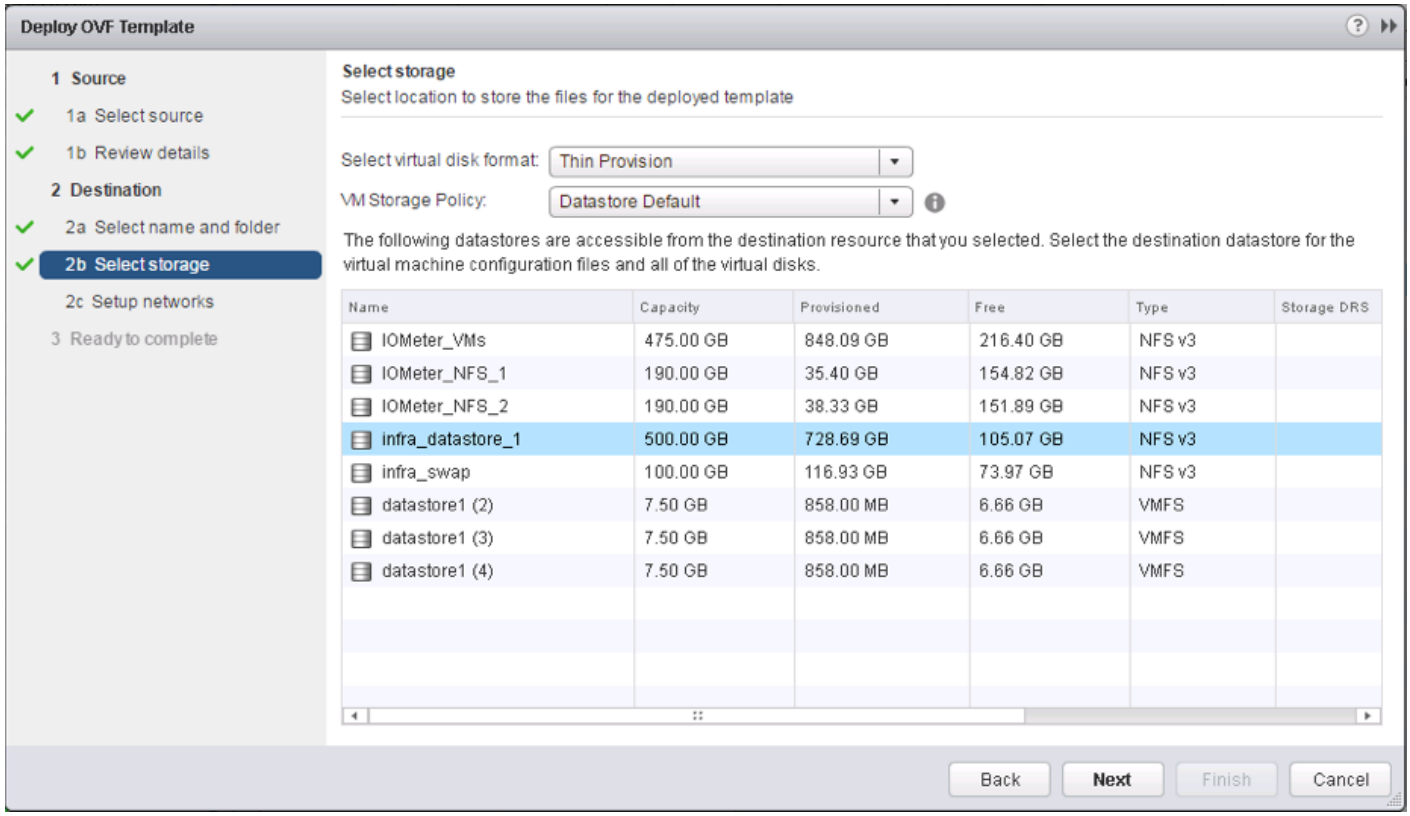

9. Click Next.
10. Specify the Destination port group within the Setup Networks pane.

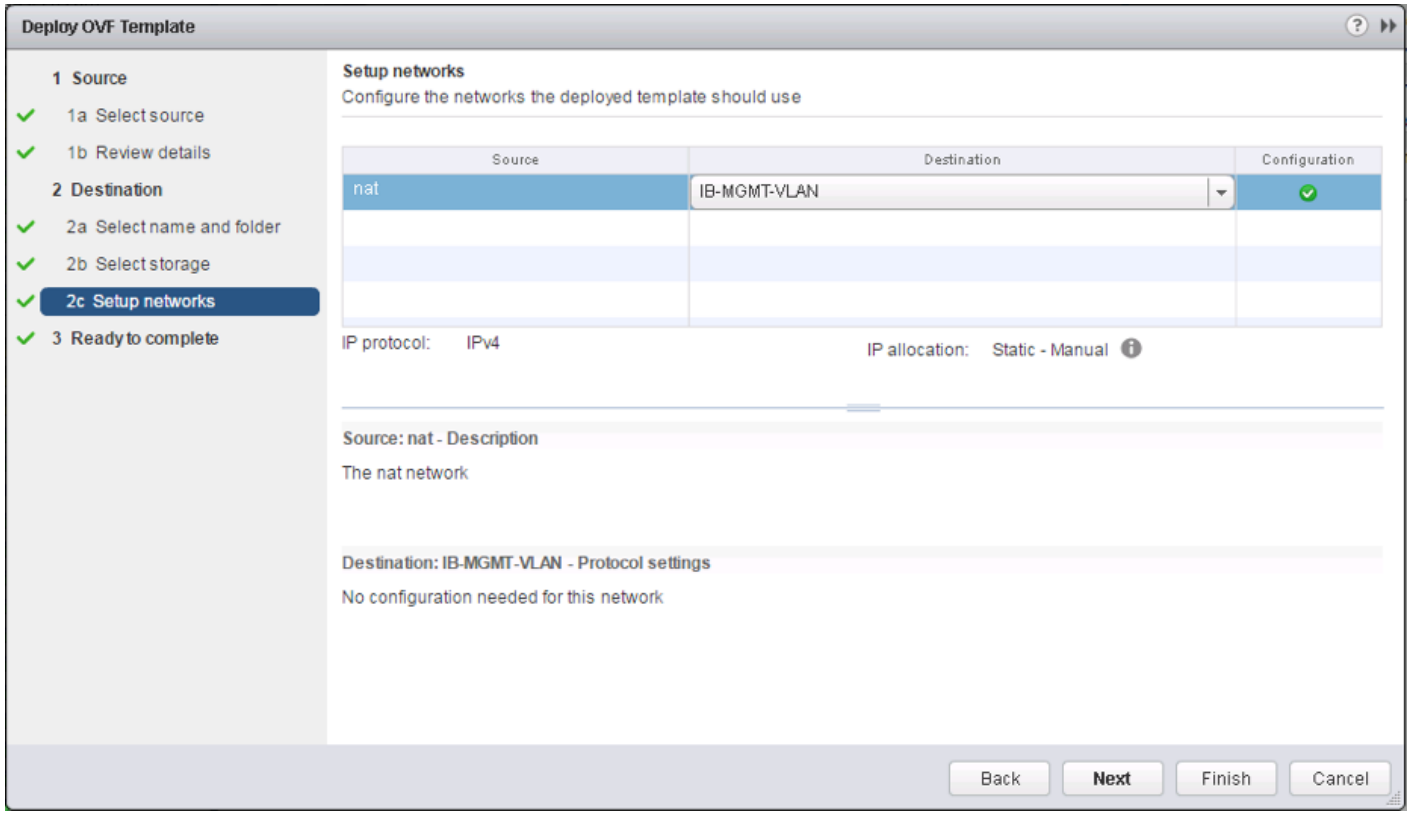

11. Click Next
12. Review he options in the Ready to complete pane, and make sure that “Power on after deployment” is not selected.

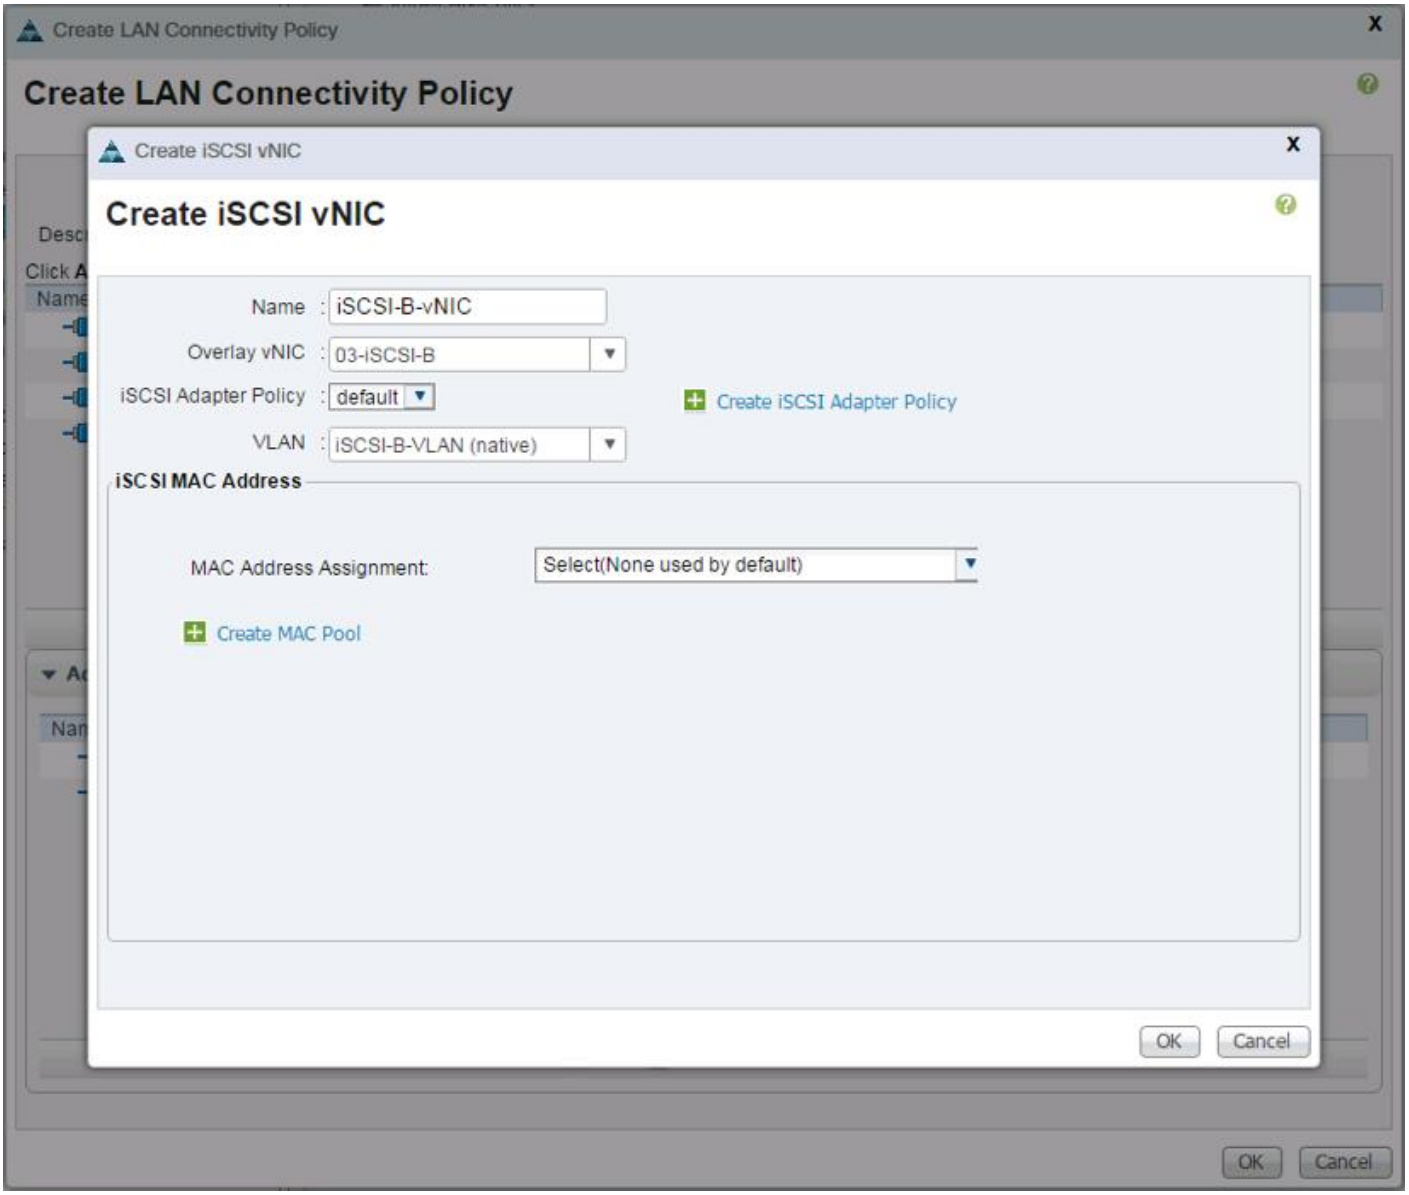

13. Click Finish to deploy the appliance from the Deploy OVF Template wizard.
14. Once the appliance has finished deploying, right click the VM from within the Hosts and Clusters or VMs and Templates section of the vSphere Web Client and select Edit Settings from the pulldown options.
15. Adjust the allocated Memory from 40G to 64G.

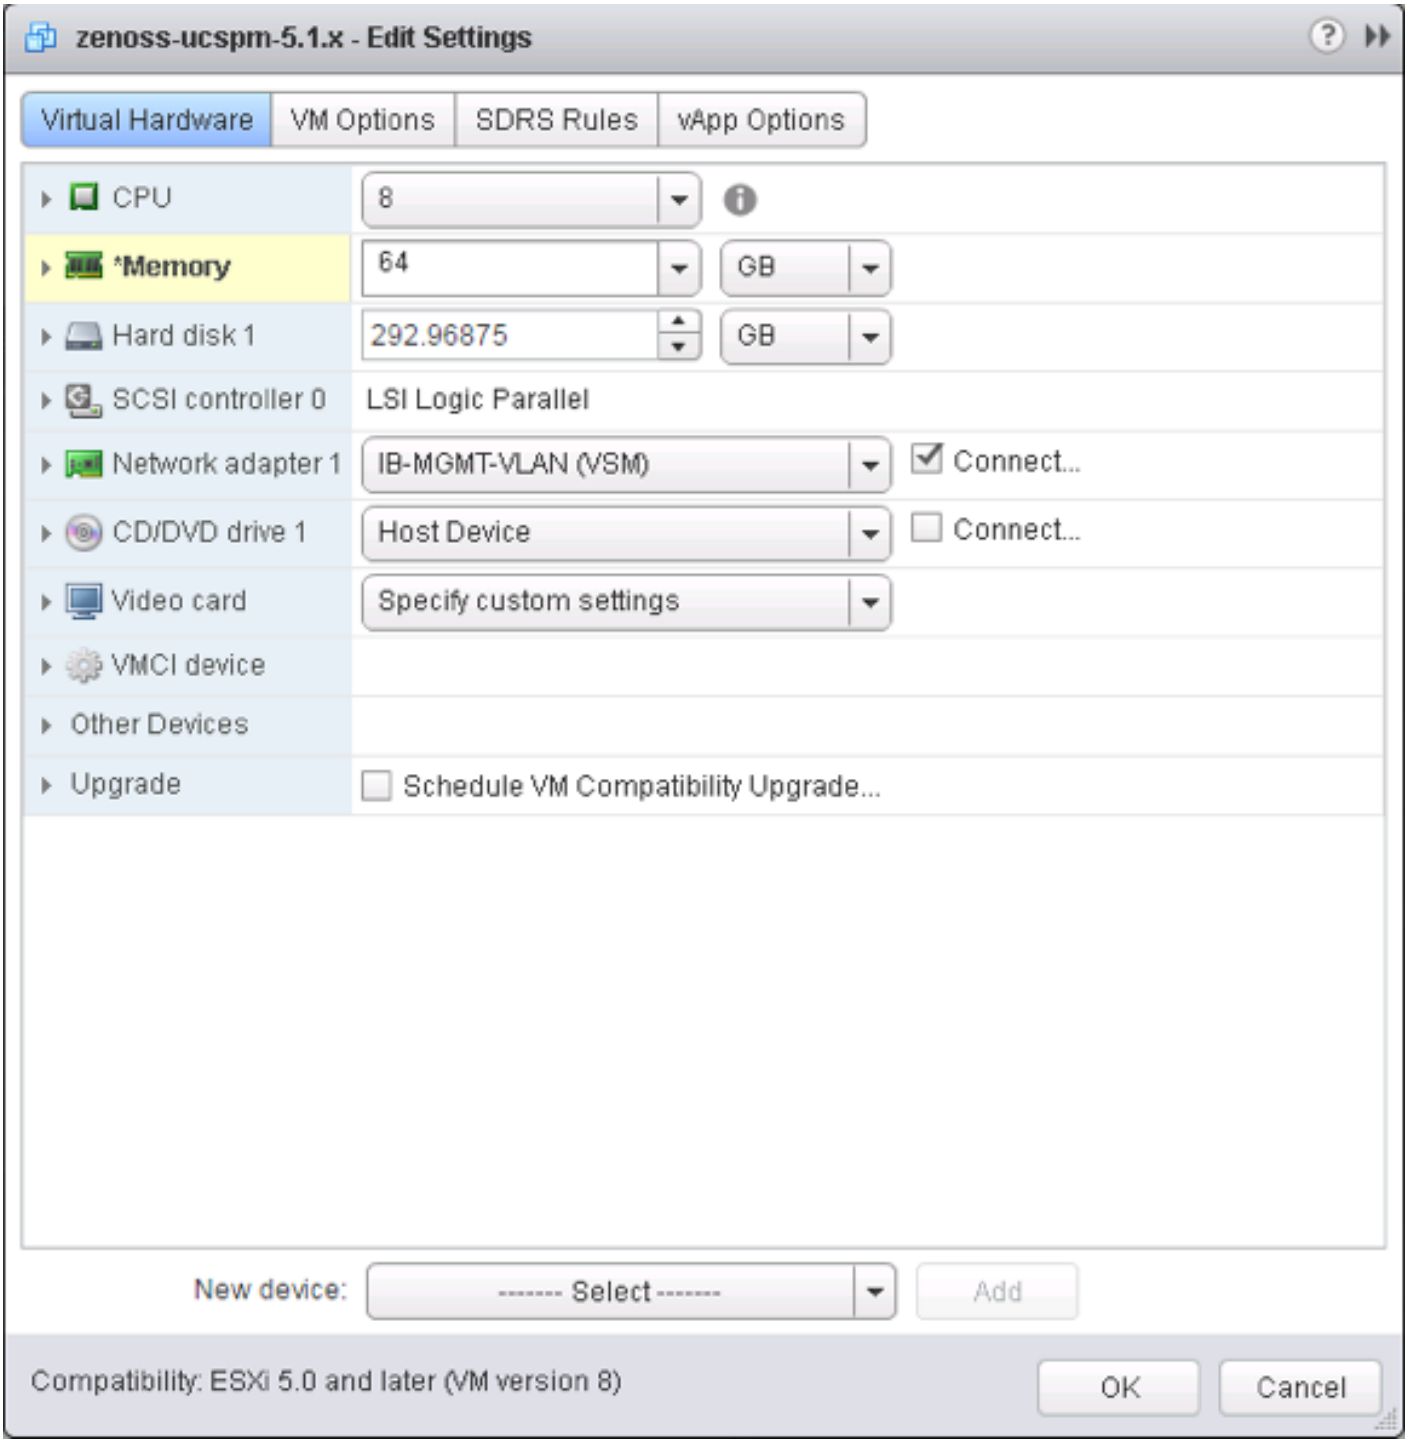

16. Click OK and power on the VM.

Cisco UCS Performance Manager Initial Configuration

1. Open up the Console of the newly provisioned UCS Performance Manager VM from the vSphere web client.

2. Login to the root account with the password “ucsm”

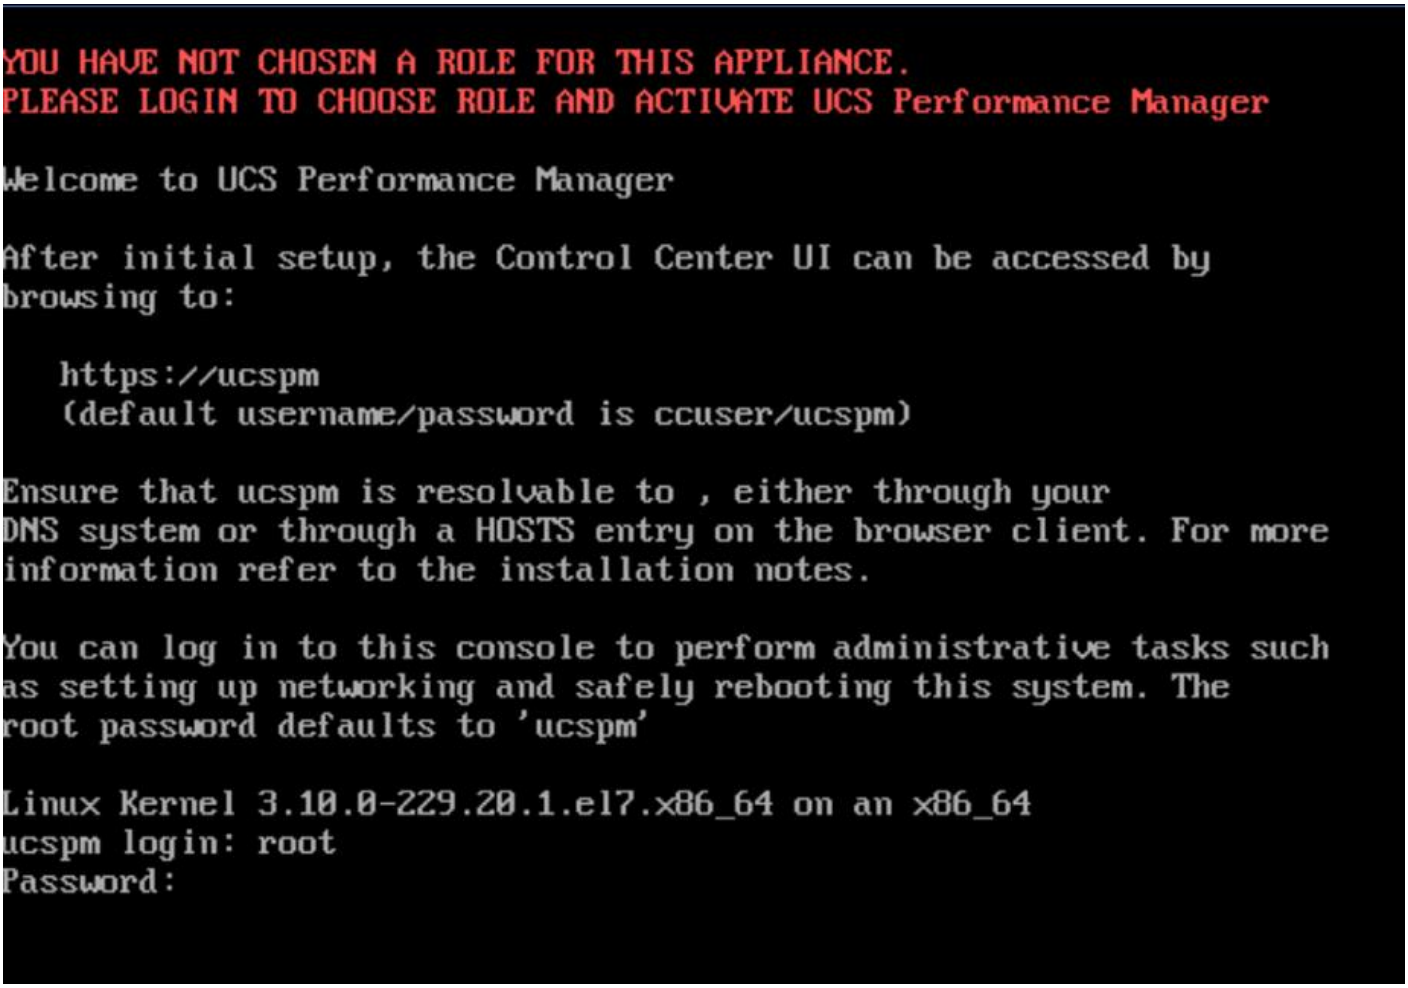

3. Set a new root password and a password for the account “ccuser” when prompted.

4. Within the following screen, leave Master selected to configure the appliance as as the Control Center master host.

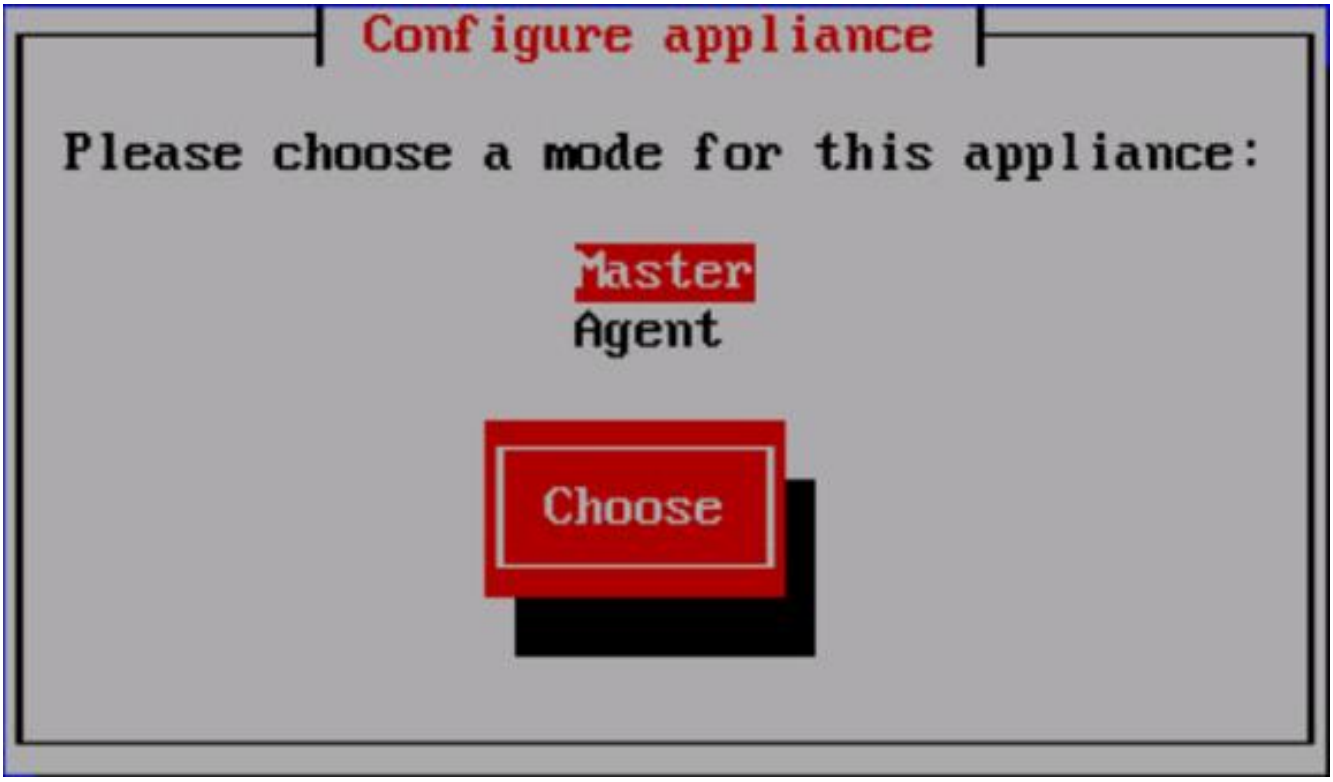

5. Press Enter to initiate a reboot and log back in with the root account after the appliance is back up.
6. Press Return from the Appliance Administration screen with Configure Network and DNS selected.

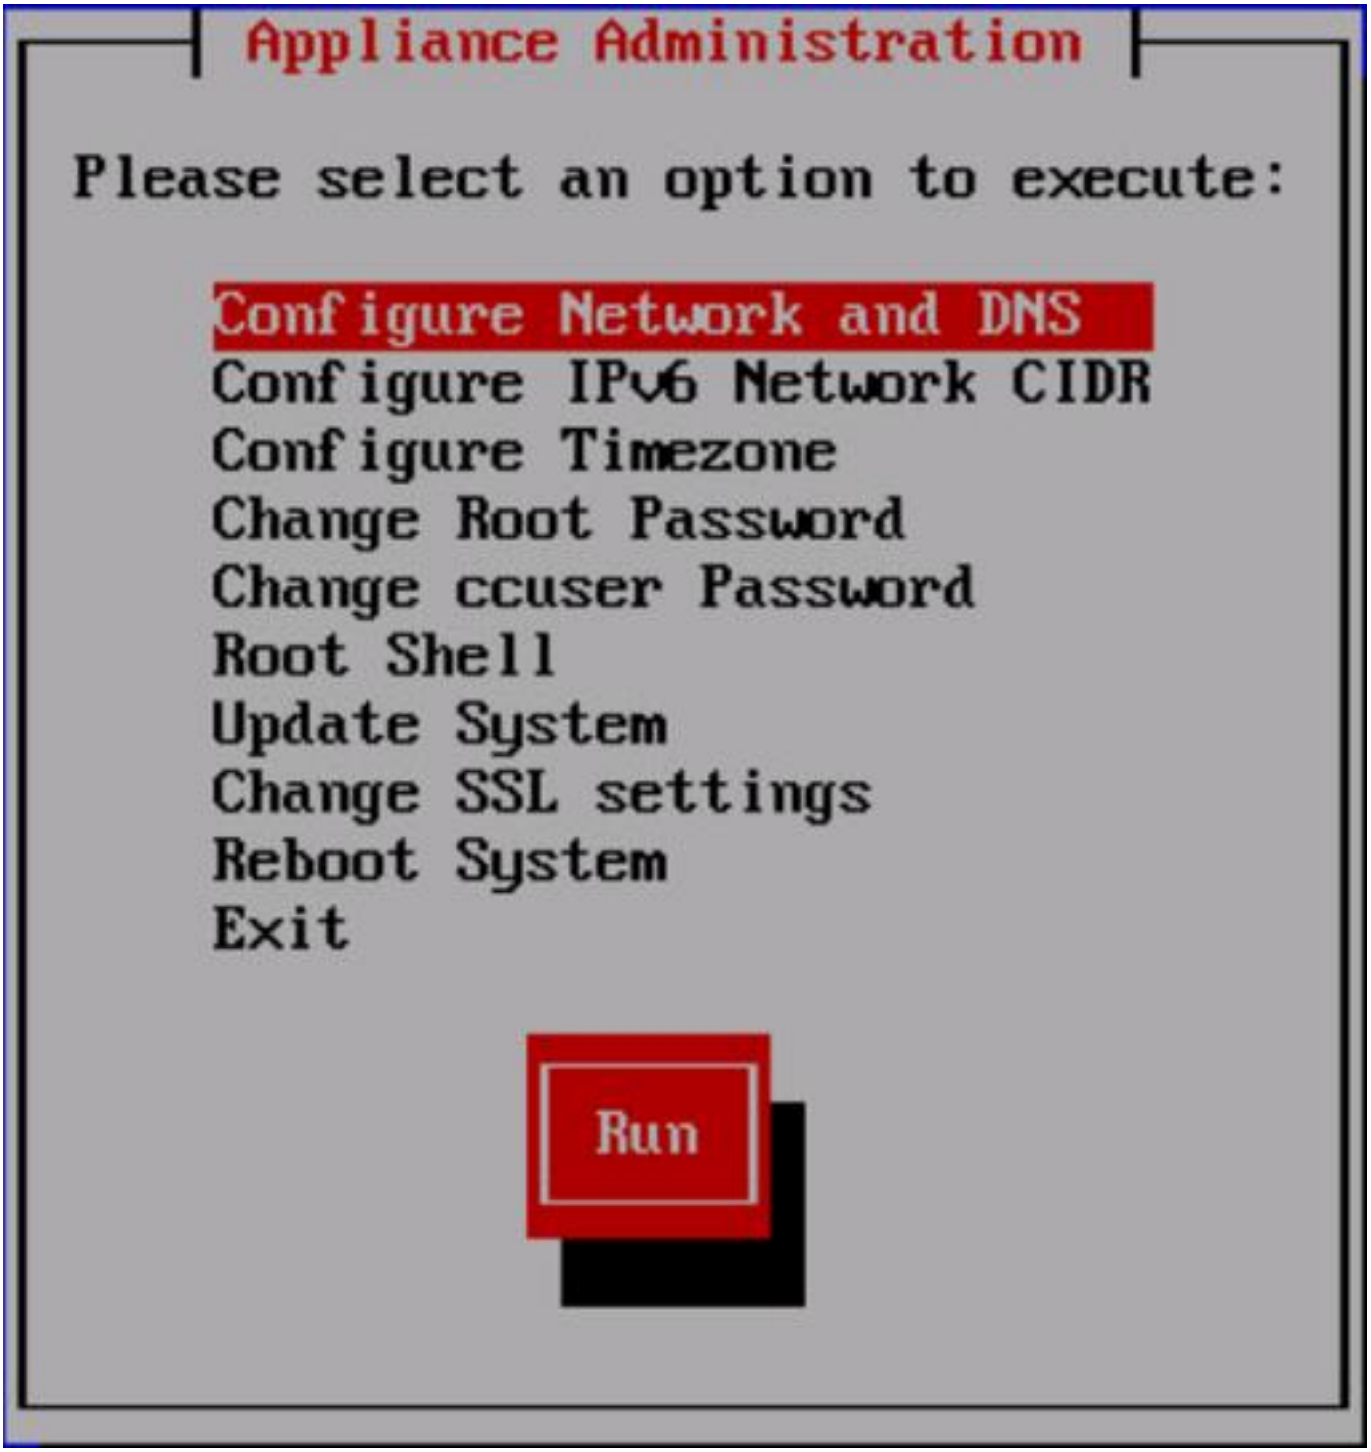

7. Leave Edit a connection selected with the NetworkManager TUI screen and hit return.

8. Select Wired connection 1, and press return to edit.

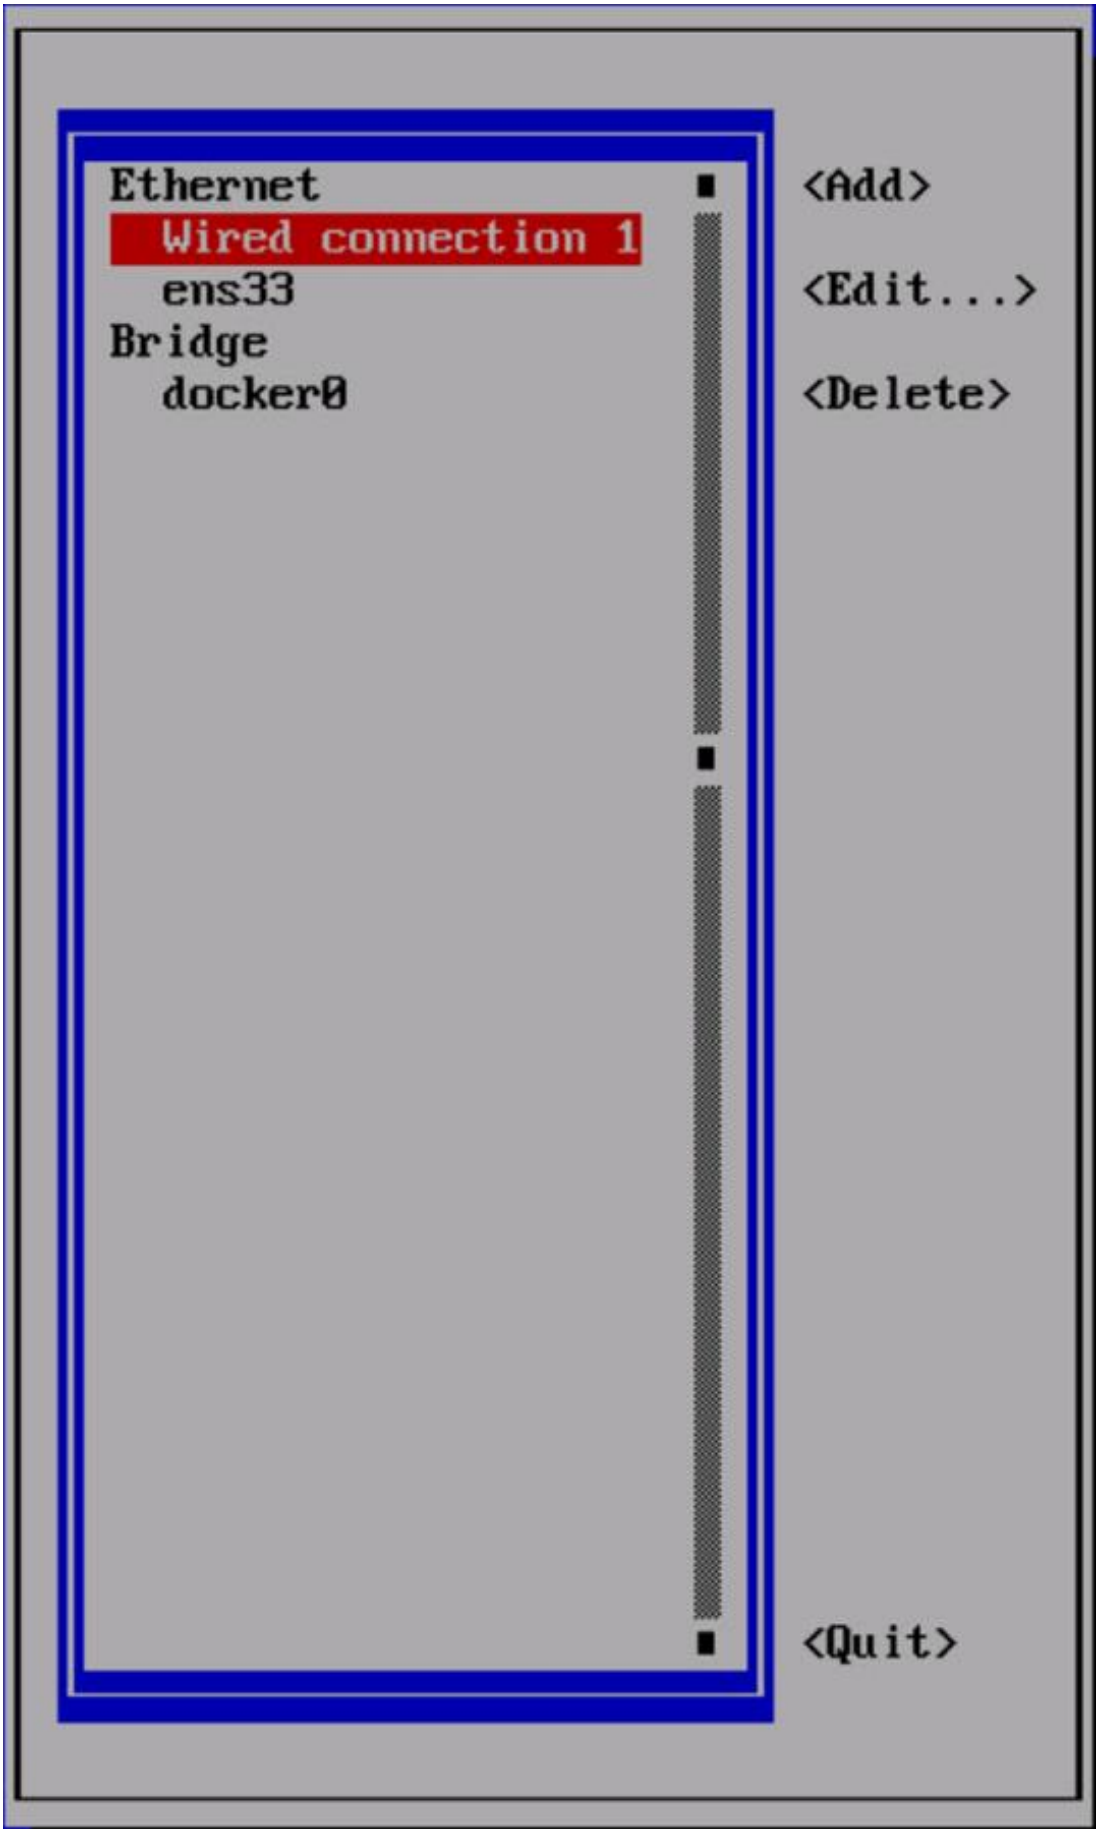
Ethernet
Wired connection 1
ens33
Bridge
docker0
<Add>
<Edit...>
<Delete>
<Quit>

9. Arrow down to the IPV4 CONFIGURATION option and hit enter to change the option from Automatic to Manual.

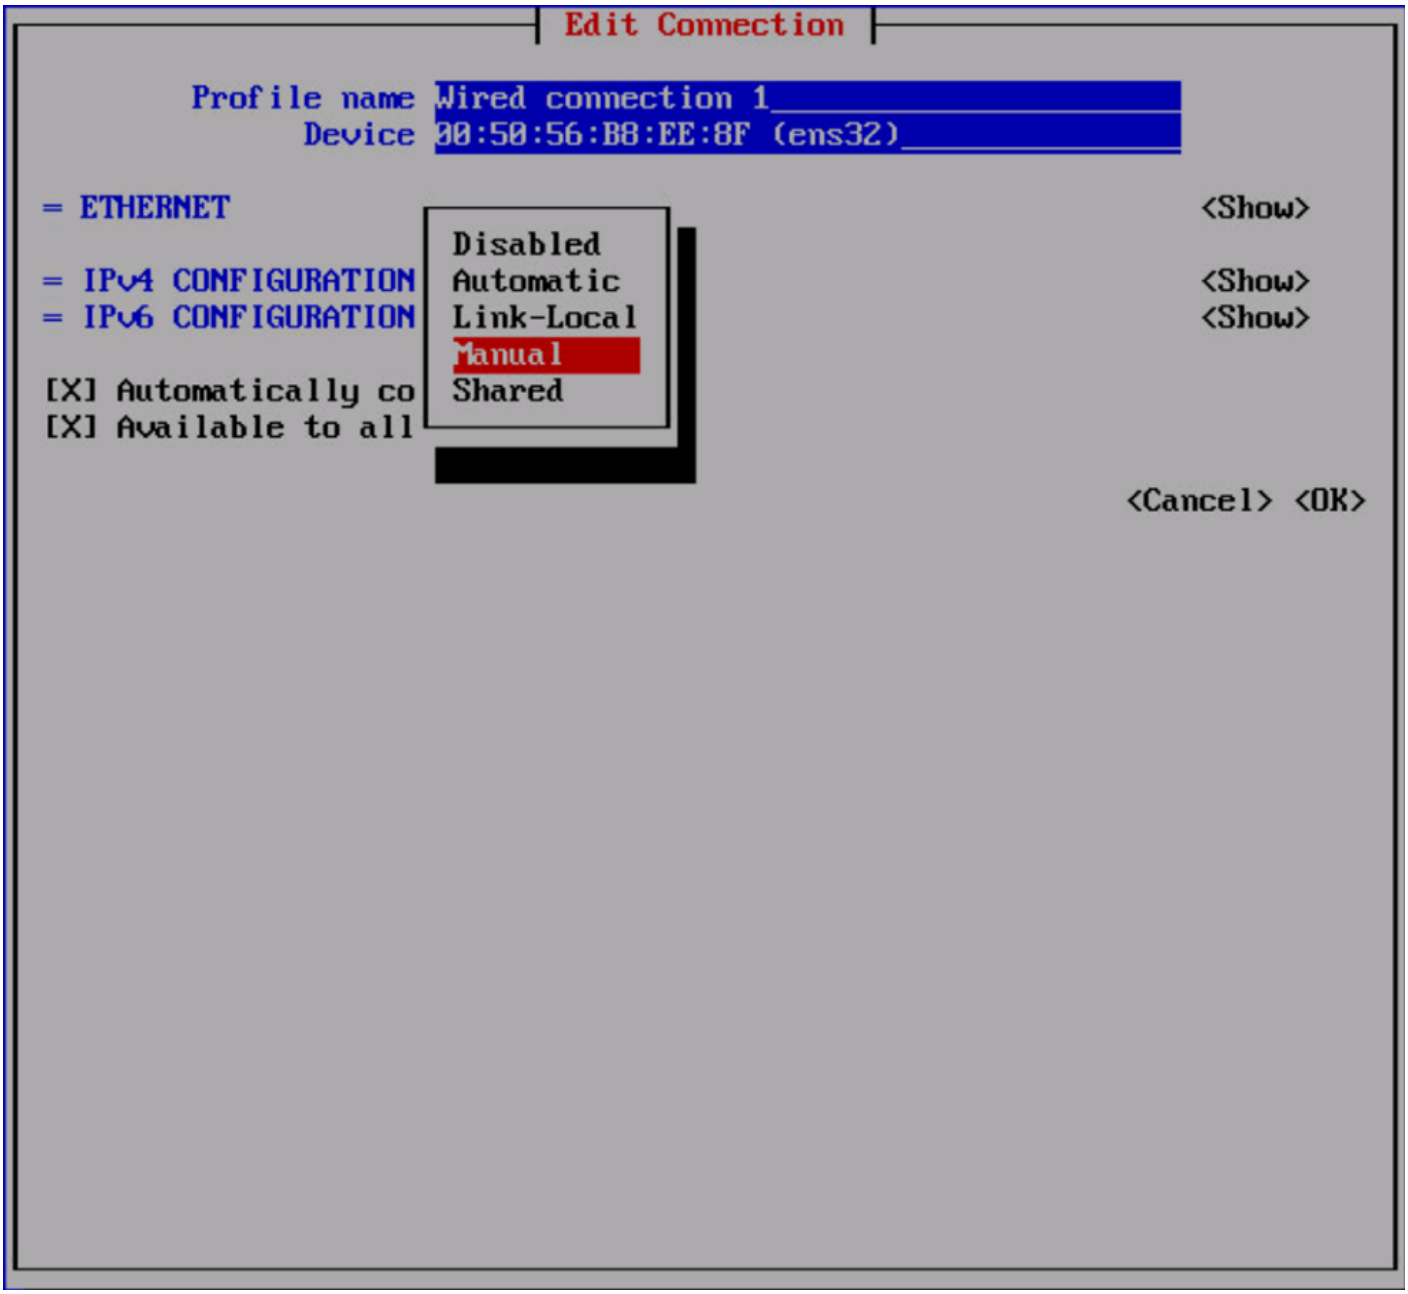

10. Right arrow over to <Show> next to the IPV4 CONFIGURATION and hit return to bring up the options.

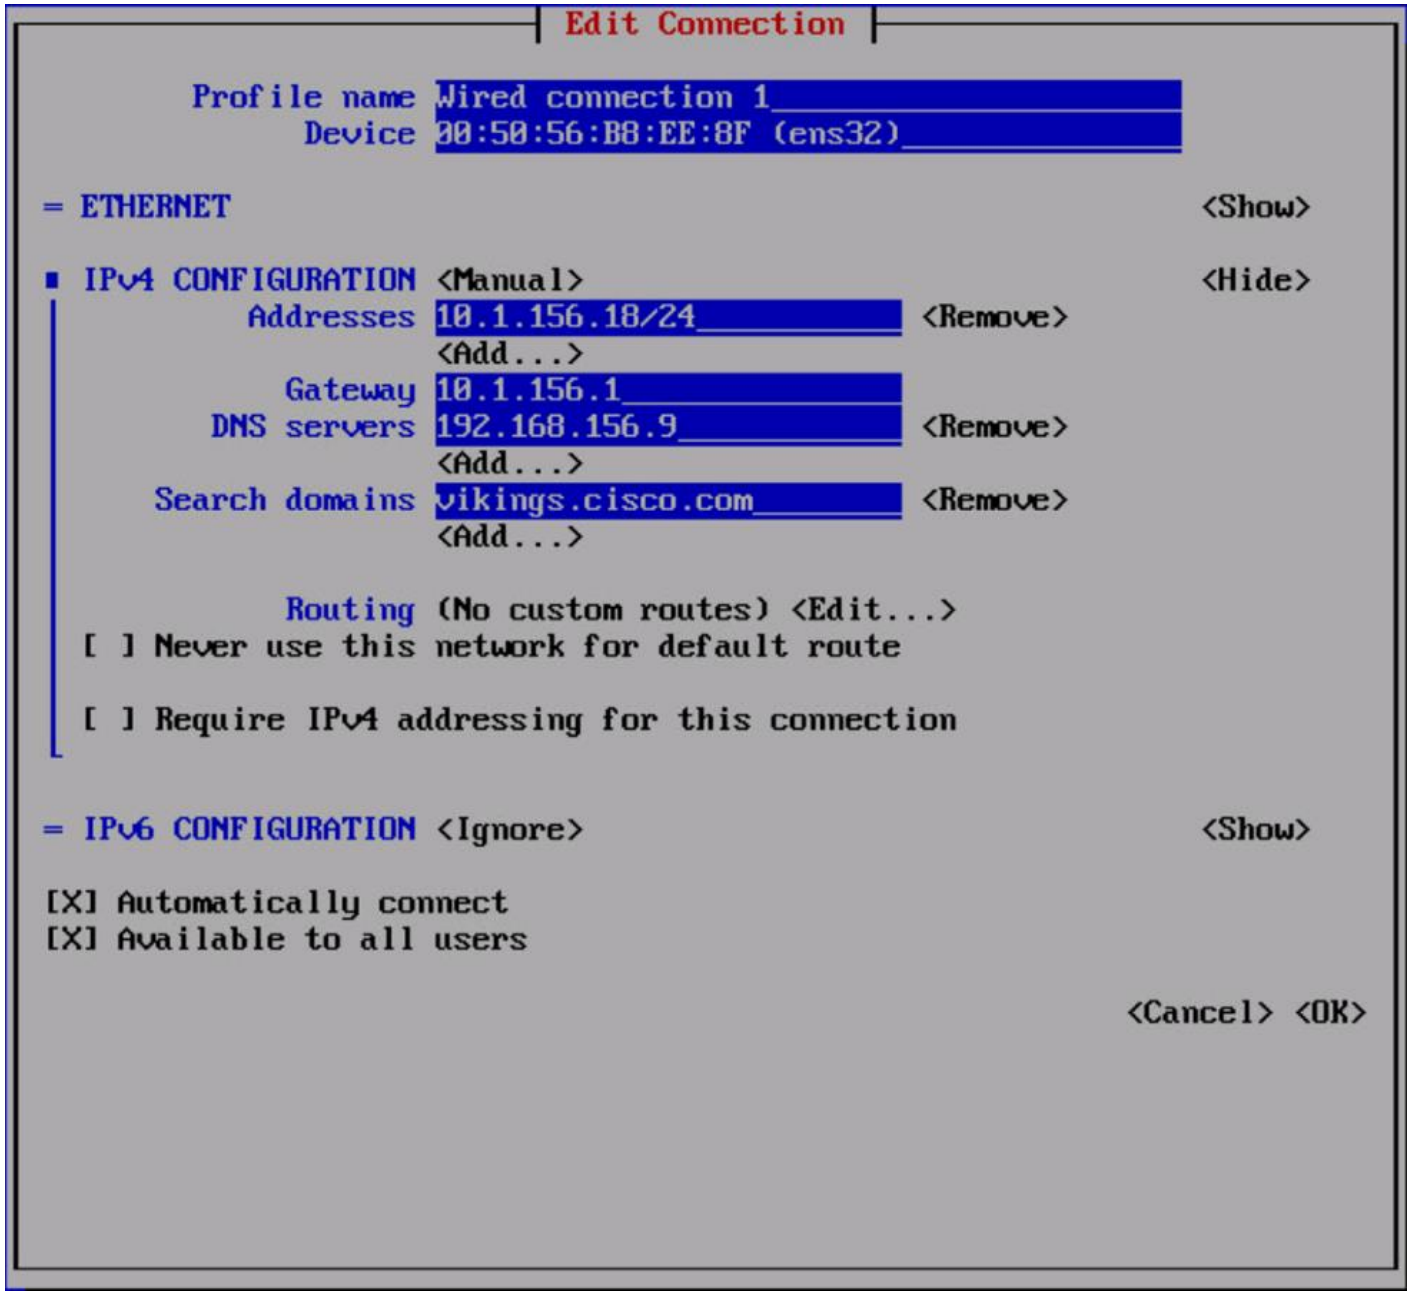

11. Enter the assigned IP address/netmask, Gateway, DNS server(s), and any Search Domains.
12. Optionally change the IPV6 CONFIGURATION option from Automatic to Ignore.
13. Arrow down till <OK> is highlighted and hit return.
14. Hit the right arrow within the interface selection menu you are returned to, arrow further down until <Quit> is selected and hit return.
15. Arrow down to have Reboot highlighted and hit return to initiate a reboot.

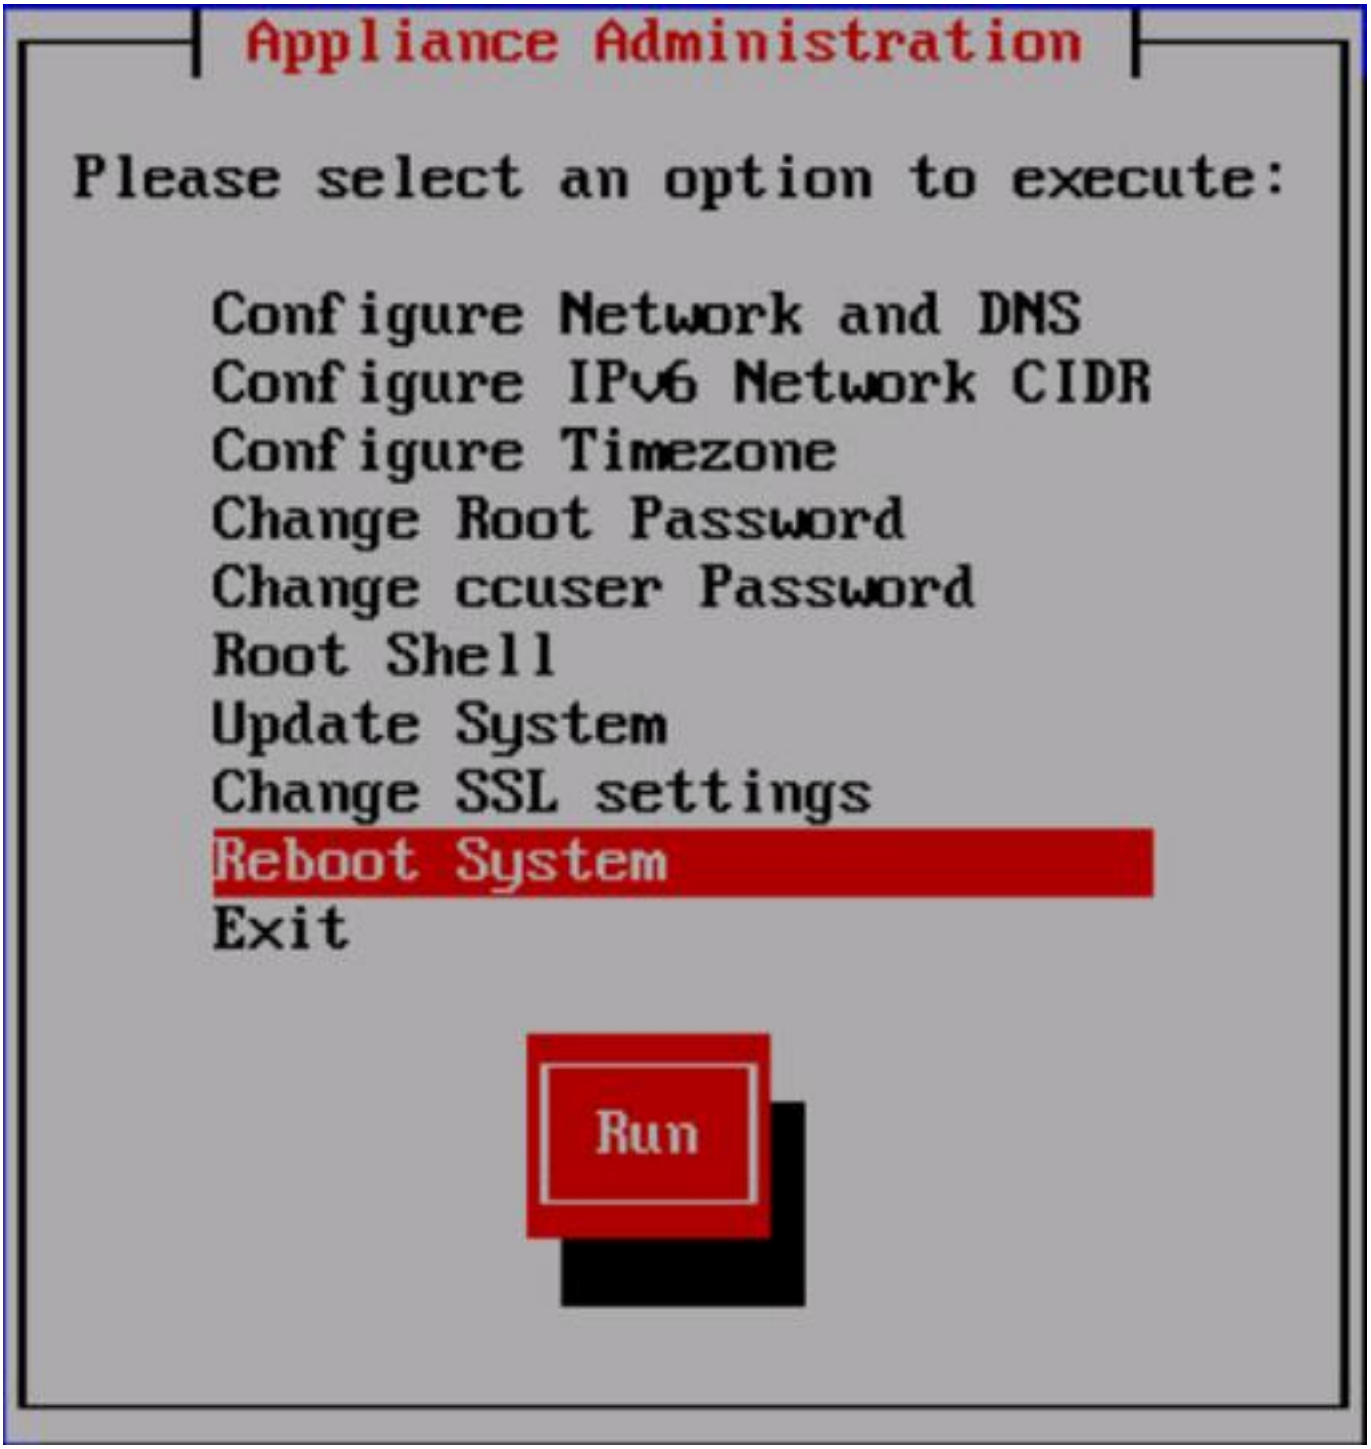

At this point further console operations can continue from with the remote console of the vSphere Web Client, or can be initiated through an ssh connection via shell or putty.

16. Re-login as root once the appliance as finished rebooting and select Configure Network and DNS from the Appliance Administration menu

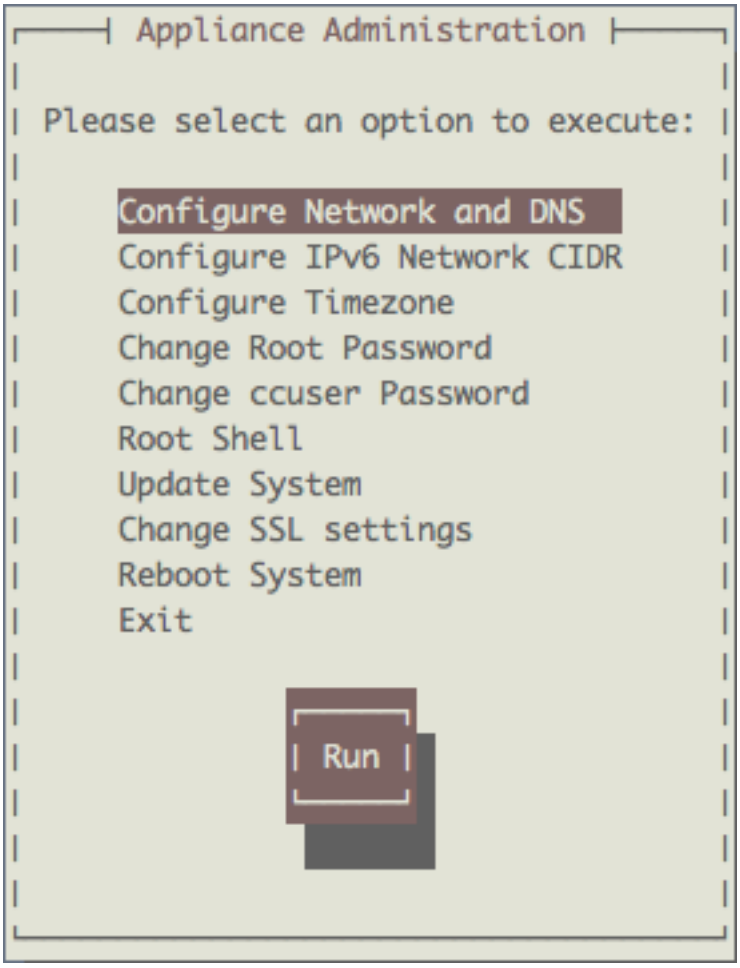

17. Within the NetworkManager TUI menu select Set system hostname and hit return.

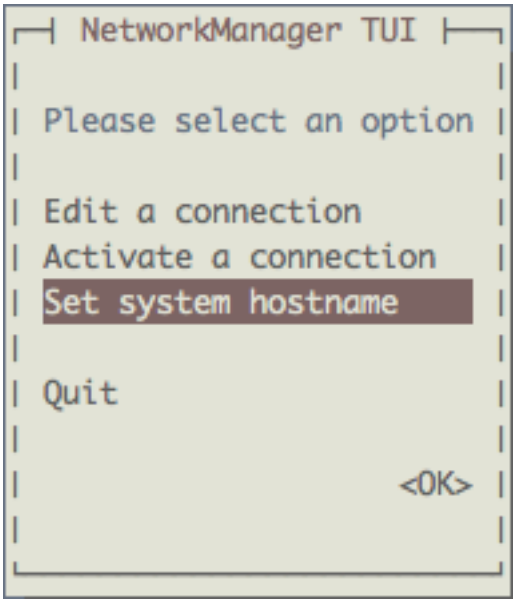

18. Enter the name desired for the UCS Performance Manager master host and press return.

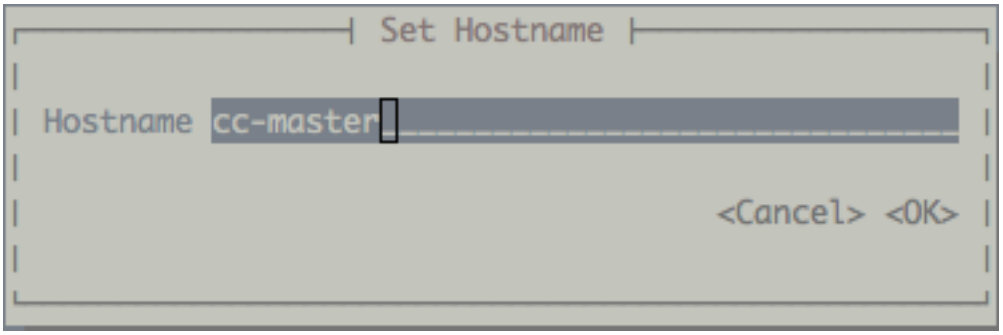

19. Hit return to move past the confirmation screen.

20. Arrow down within the Appliance Administration menu to the Reboot System option and press return to change the system hostname.

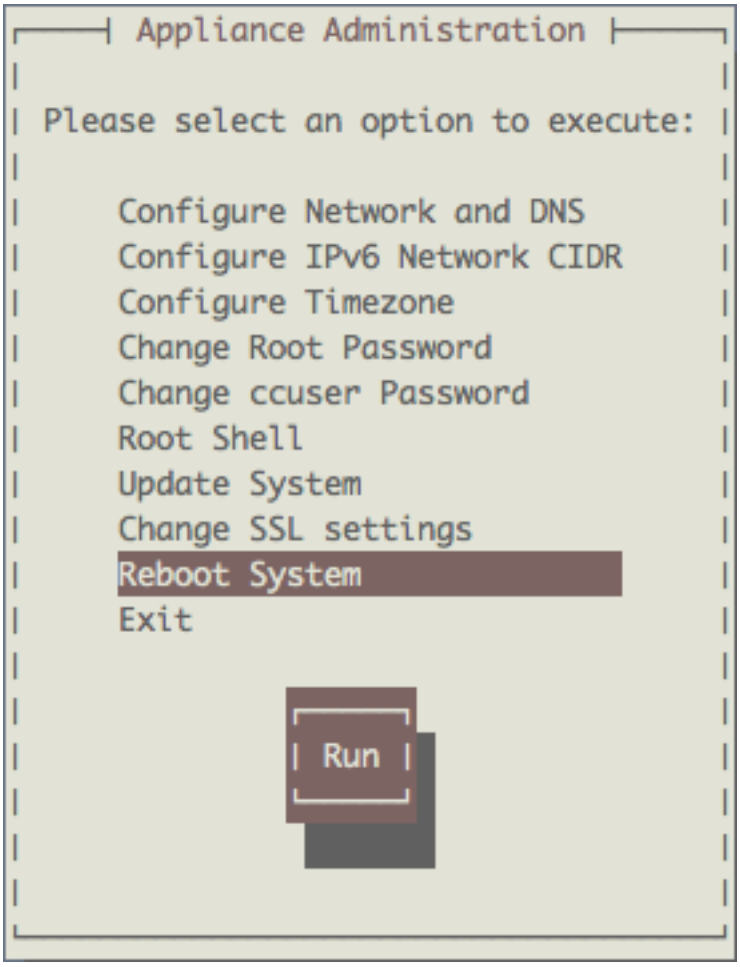

Cisco UCS Performance Manager Deployment

The base IP or hostname of the UCS Performance Manager master host will allow a connection the Control Center that the UCS Performance Manager virtual host is provisioned from.

In our example *ControlCenter* was registered in our DNS as cc-master.vikings.cisco.com, and was accessible as https://cc-master.vikings.cisco.com or via the IP assigned to it. The UCS Performance

Manger will need to to resolve in DNS as a CNAME alias or local /etc/hosts entry from the browsing system as ucspm.*ControlCenter* which was “ucspm.cc-master.vikings.cisco.com” in our example.

1. Login to the Control Center page with the ccuser account, confirm and Security Exceptions identified when connecting.

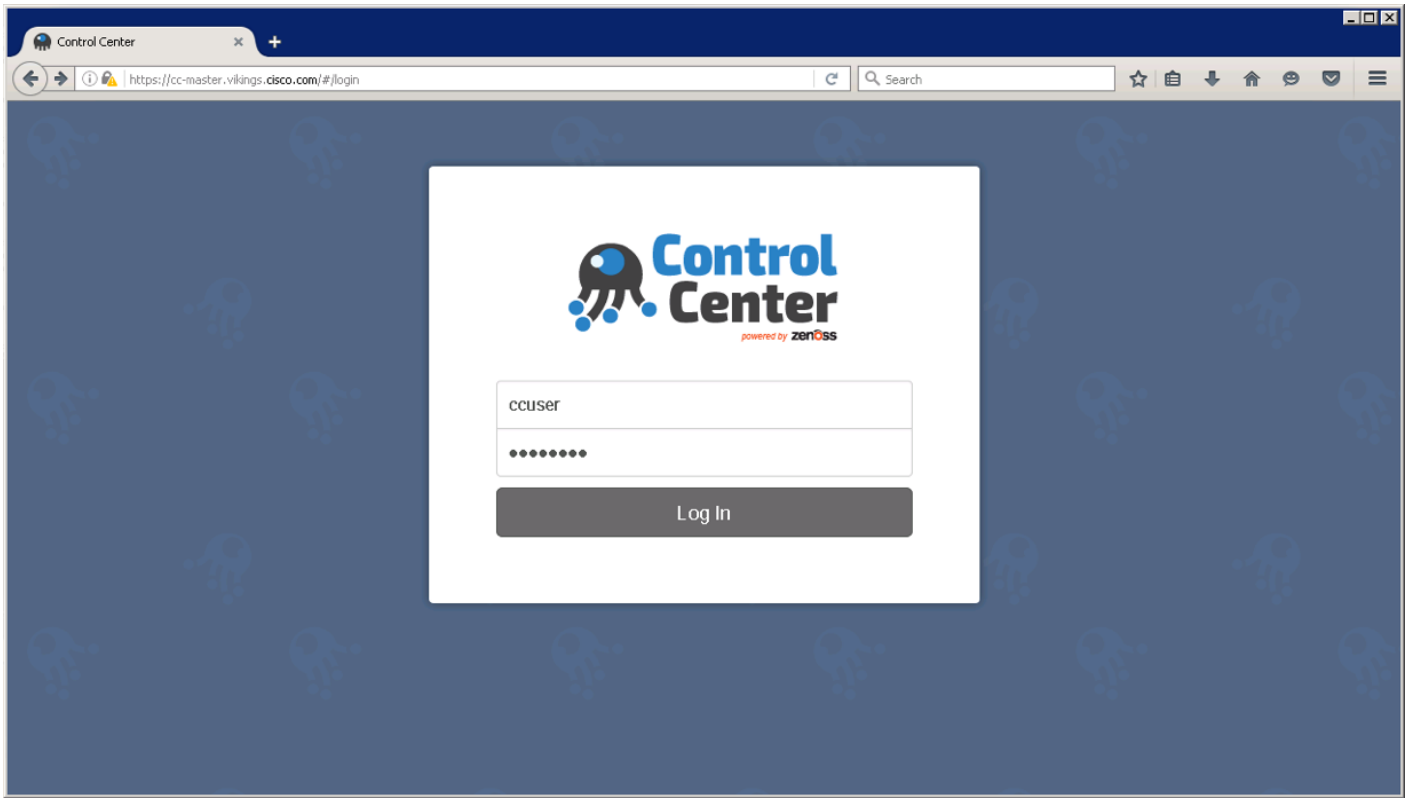

2. The initial login will bring up a Deployment Wizard to provision the UCS Performance Manager virtual host. If this does not come up, it can be initiated by selecting the +Application button in the mid to upper right of the screen.

3. Enter the Host and port values (ucspm.cc-master.vikings.cisco.com:4979 in our example), select default for the Resource Pool ID, and enter 75% for the RAM Commitment.

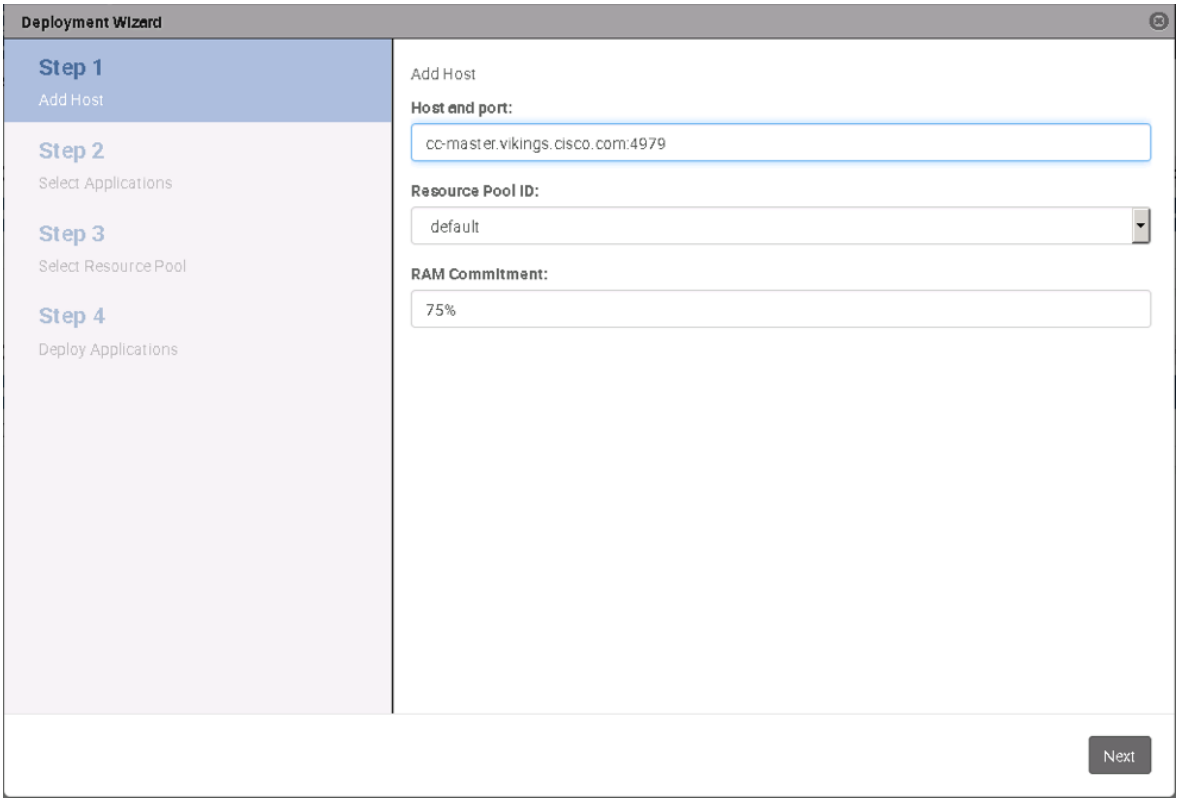

4. Click Next.

5. Select ucspm (v2.0.0) as the application to install.

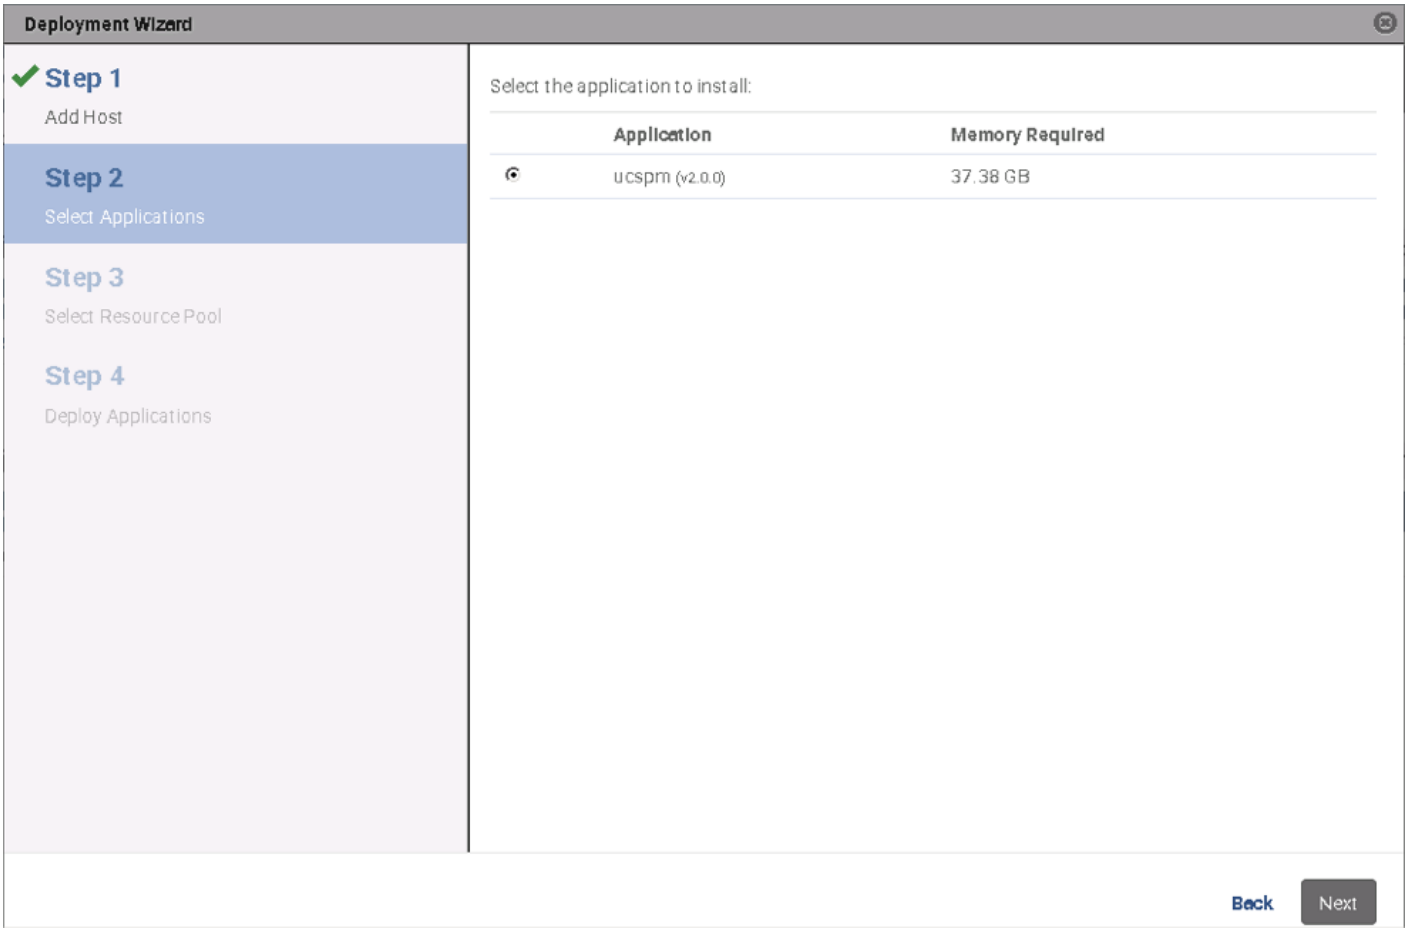

6. Click Next.

7. Select default as the resource pool.

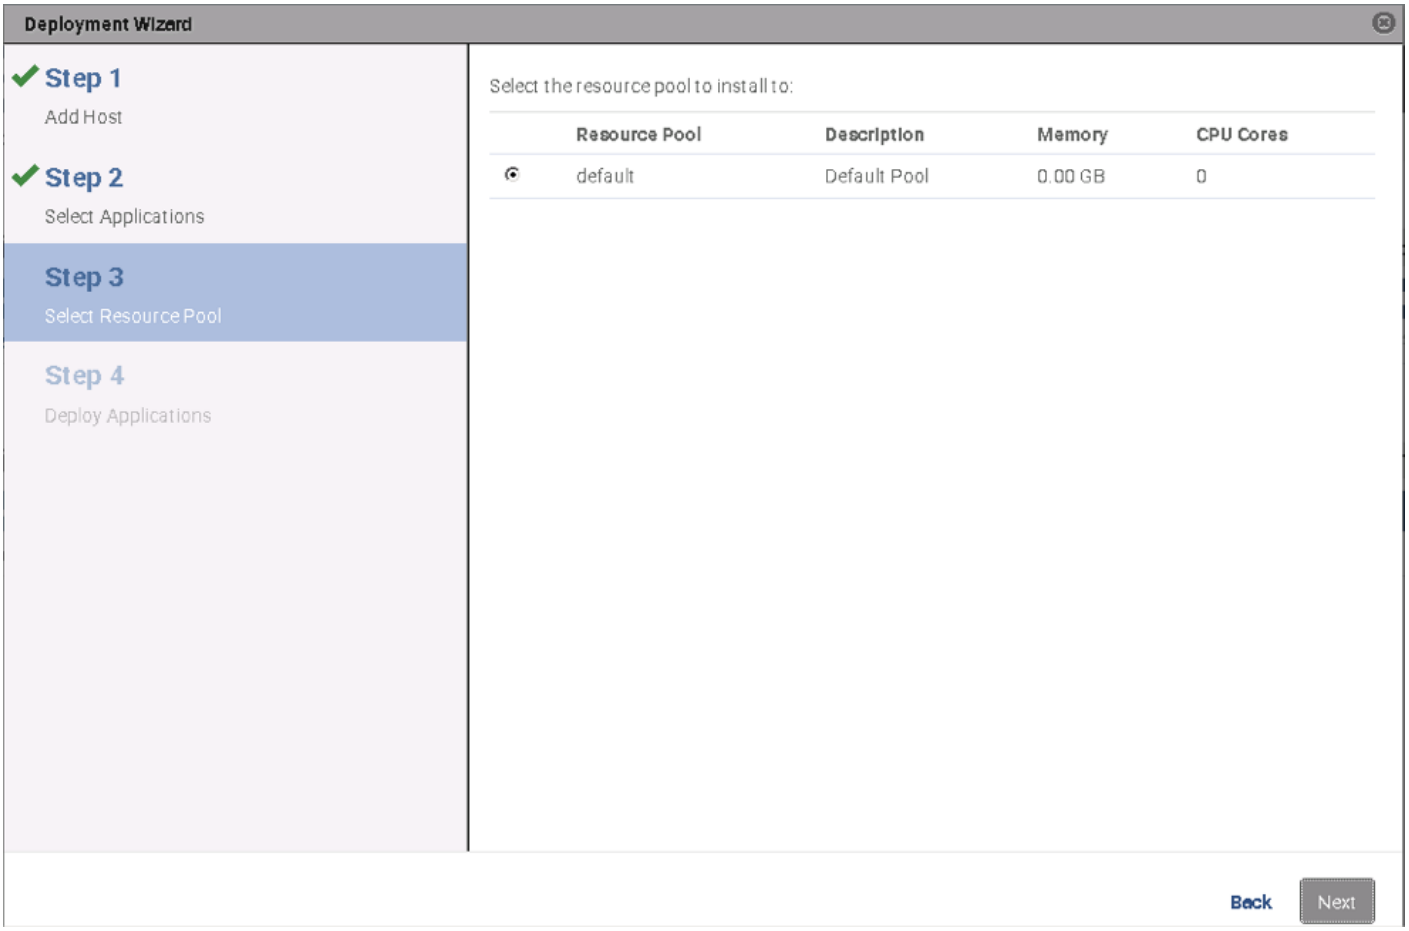

8. Click Next.

9. Specify a Deployment ID name to provision the application.

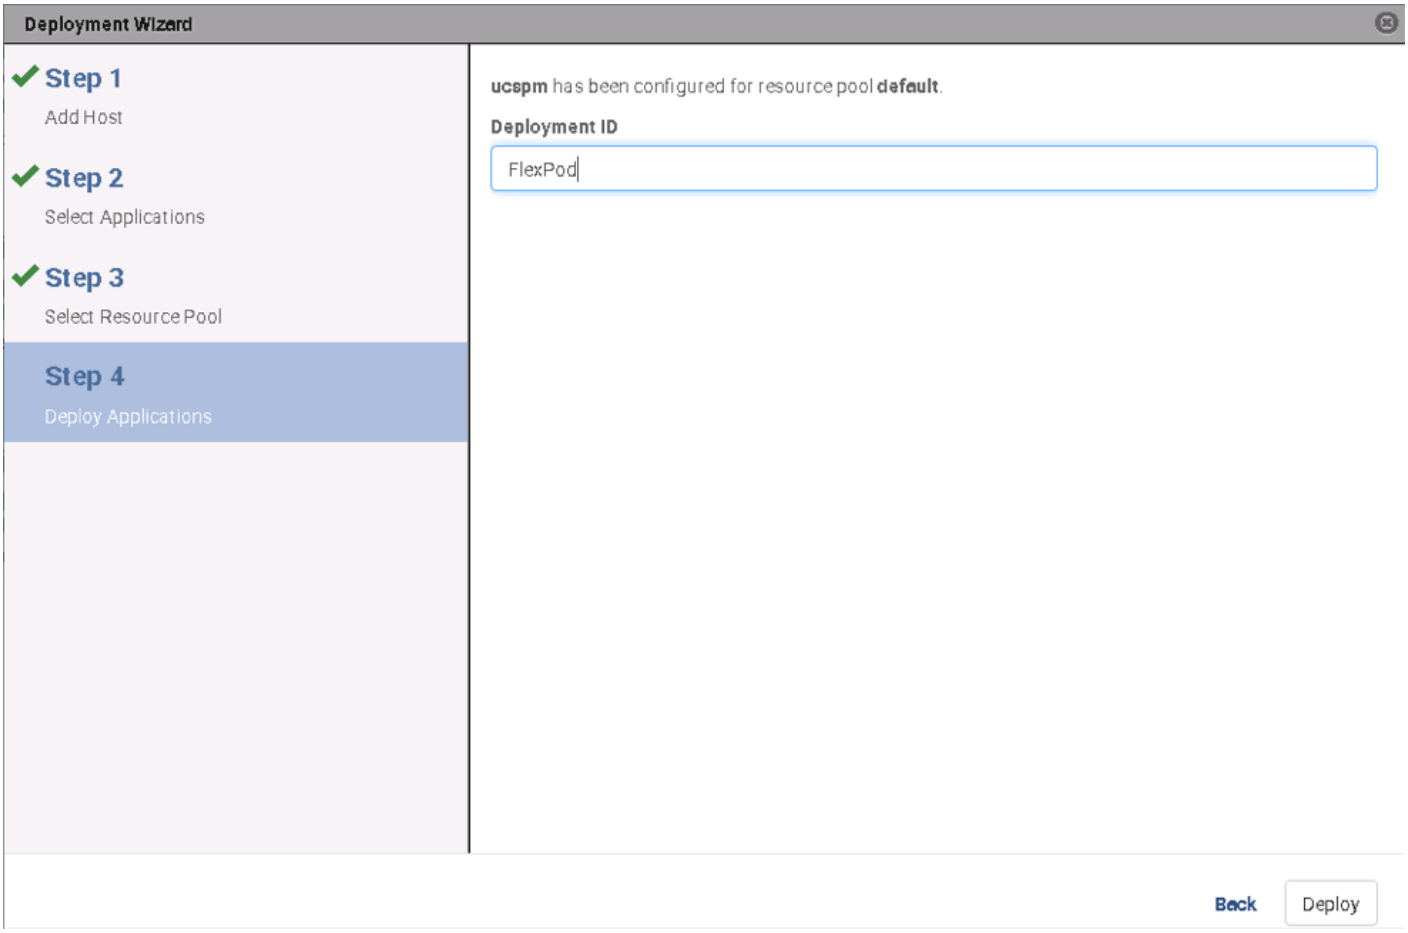

10. Click Deploy to provision.

11. In the Actions column of the Applications table, click the Start option of the ucspm (v2.0.0) row.

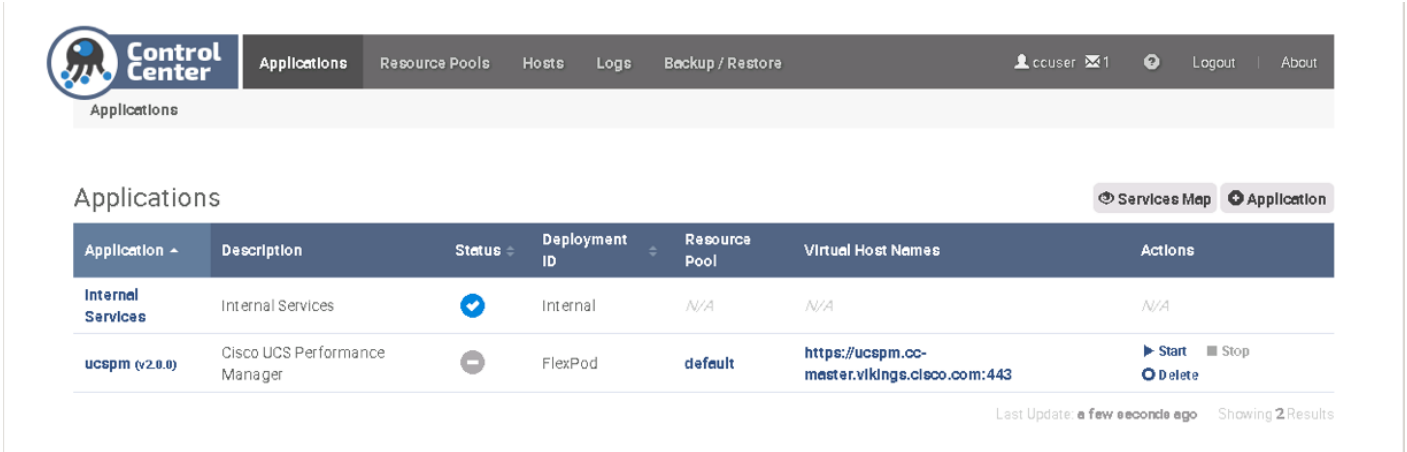

12. A Start Service dialog window will pop up.

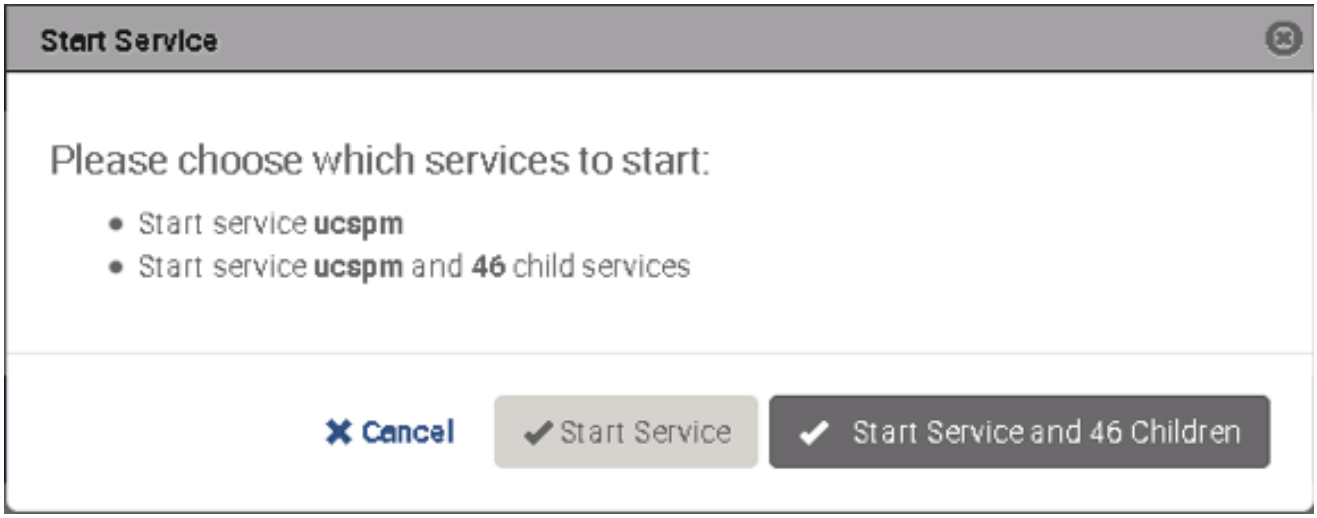

13. Select Start Service and 46 Children.

14. In the Application column of the Applications table, click ucspm in the ucspm row

15. Scroll down to watch child services starting. Typically, child services take 4-5 minutes to start. When no child service shows a red exclamation point icon, Cisco UCS Performance Manager is running.

Cisco UCS Performance Manager Configuration of FlexPod Infrastructure

1. Using another web browser session, connect to the Cisco UCS Performance Manager virtual host using the DNS CNAME or local entry configured within the host the browser is running from (https://ucspm.vikings.cisco.com/).

2. Scroll down to the licensing screen that first appears, select the “Click this box to verify you agree to the License.” in the bottom left, and click Accept License in the bottom right of the page.

3. Click the “Get Started!” option in the initial screen.

4. Click Add License File in the following screen if you are adding one, or go directly to Click Next on the Licensing screen if you are going to run with the 30 day trial.

5. Specify an admin password and create a local account on the following screen.

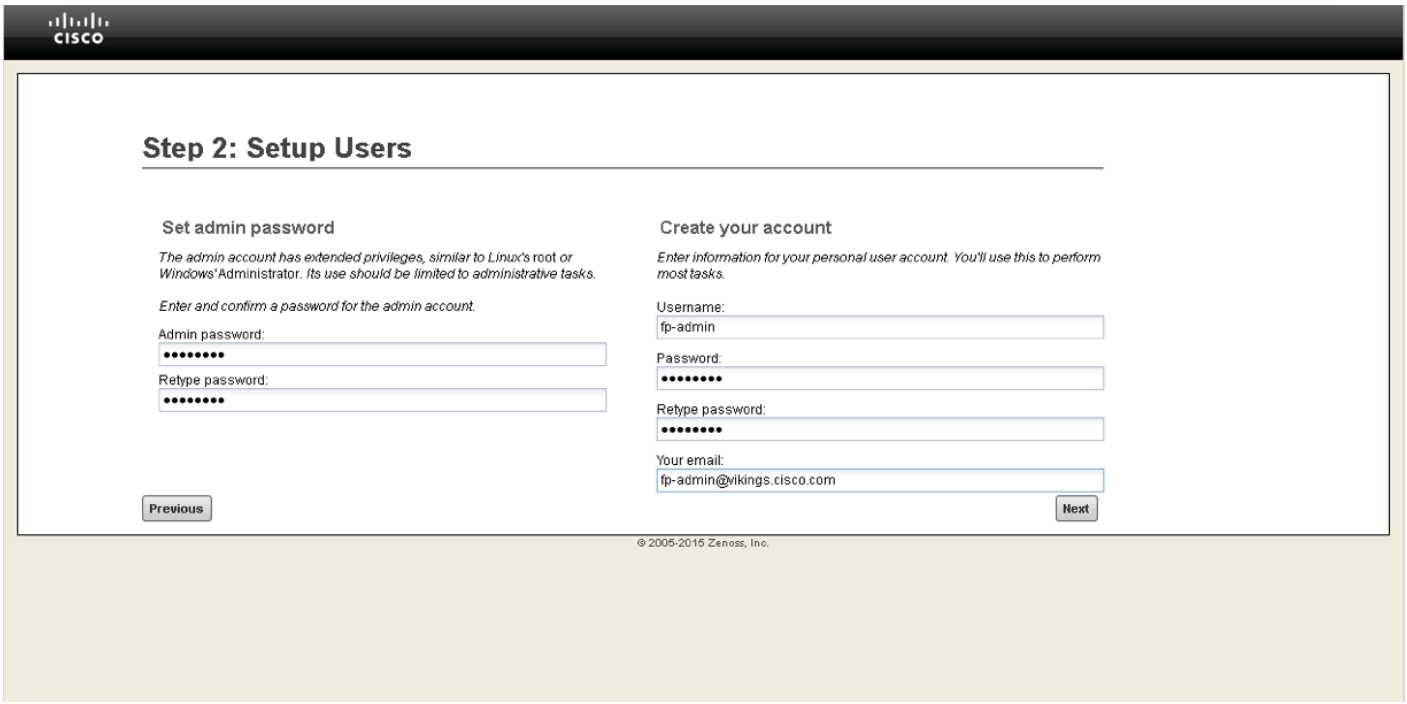

6. Click Next.

Administrator and admin accounts were used as the access accounts in this section. In a production environment it may be more appropriate to use dedicated read only accounts.

7. Add a UCS Central IP, administrative account and password if you have one independently deployed in your environment (UCS Central was not covered in this document.)

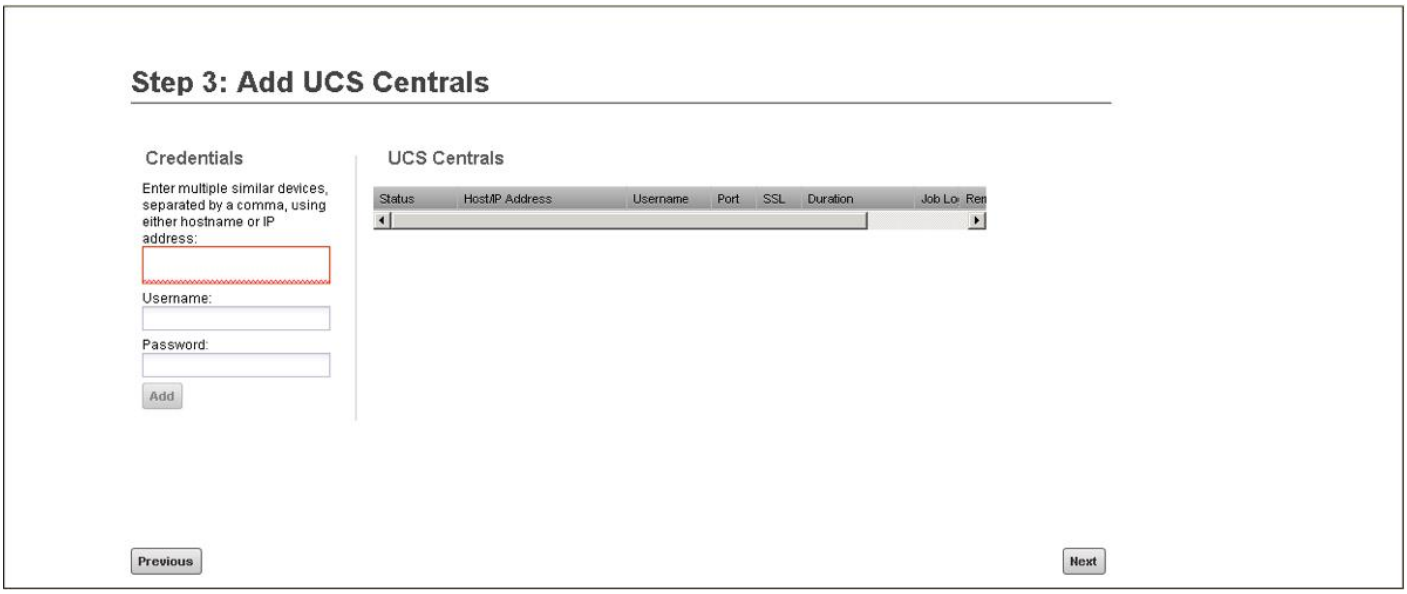

8. Click Next.

9. Add in the UCS Fabric Interconnect(s) virtual IPs for UCS Domains to be monitored in your environment.

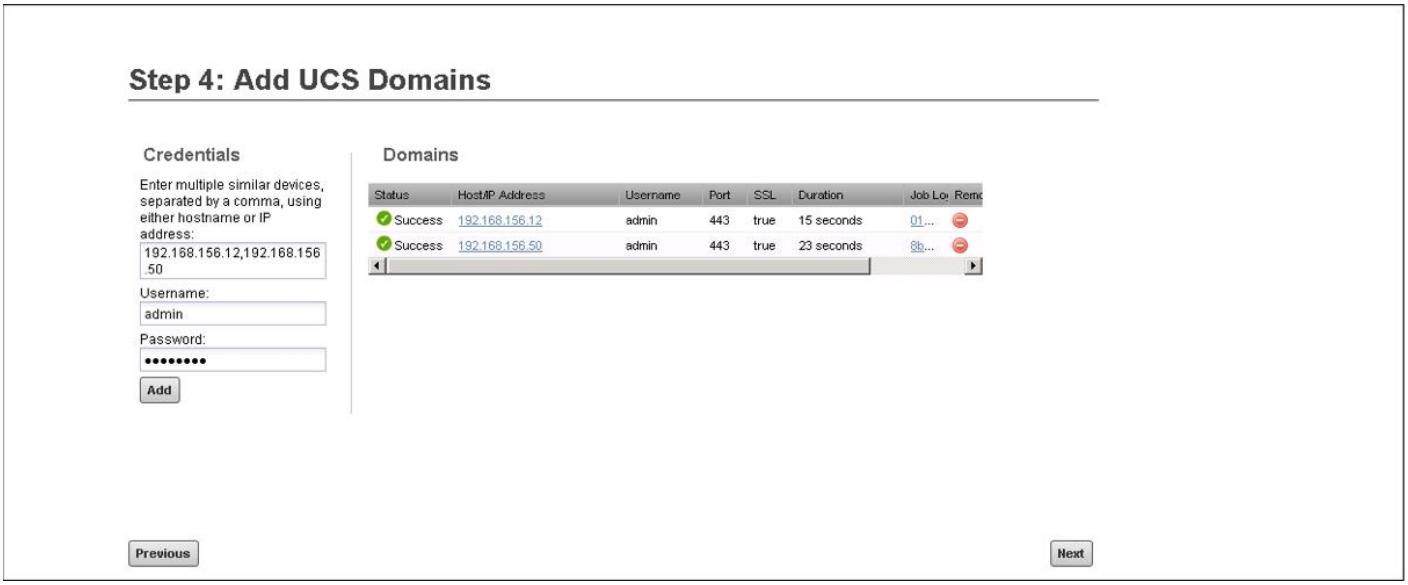

10. Click Add, then click Next.

11. Select Network from the Category column to add in the Nexus Switches to be monitored.

12. Select Cisco Nexus 9000 and enter the IP and credentials for access.

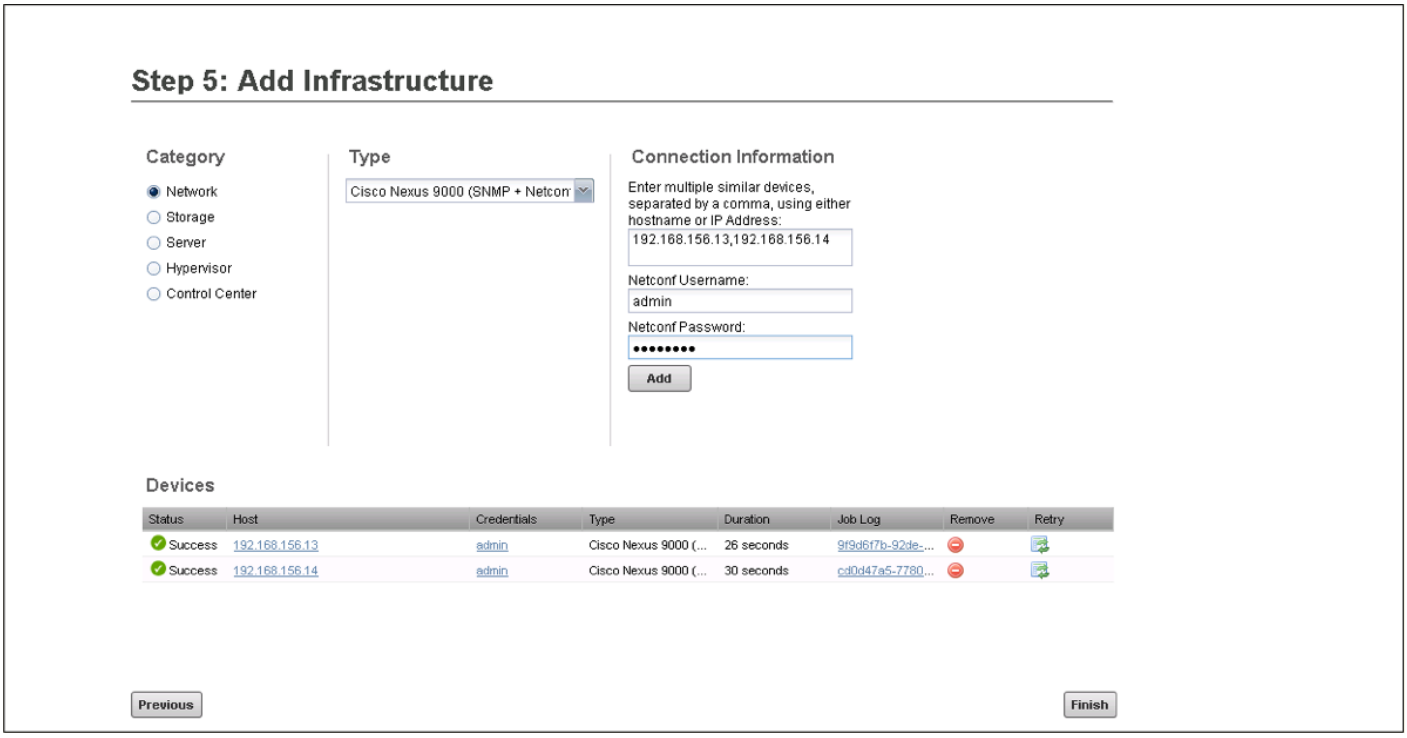

13. Click Add.
14. Click Storage from the Category column to add in the NetApp AFF.
15. Select NetApp C-Mode Filer (ZAPI) and enter the IP and credentials for access.
16. Click the “Use SSL?” checkmark box.

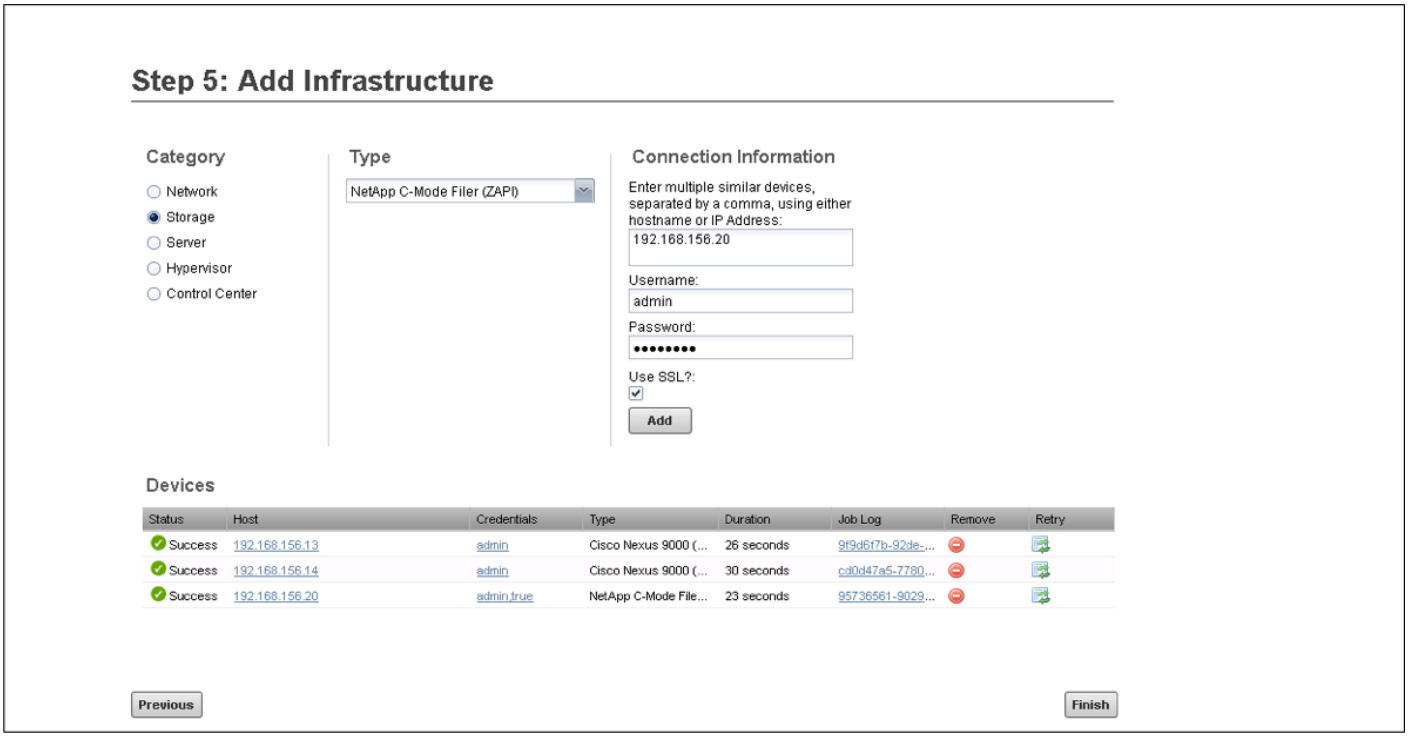

17. Click Add.

18. Click Hypervisor from the Category column to add in the vCenter Server.

19. Select vSphere EndPoint (SOAP) and enter the IP and credentials for access vCenter.

20. Click the “Use SSL?” checkmark box.

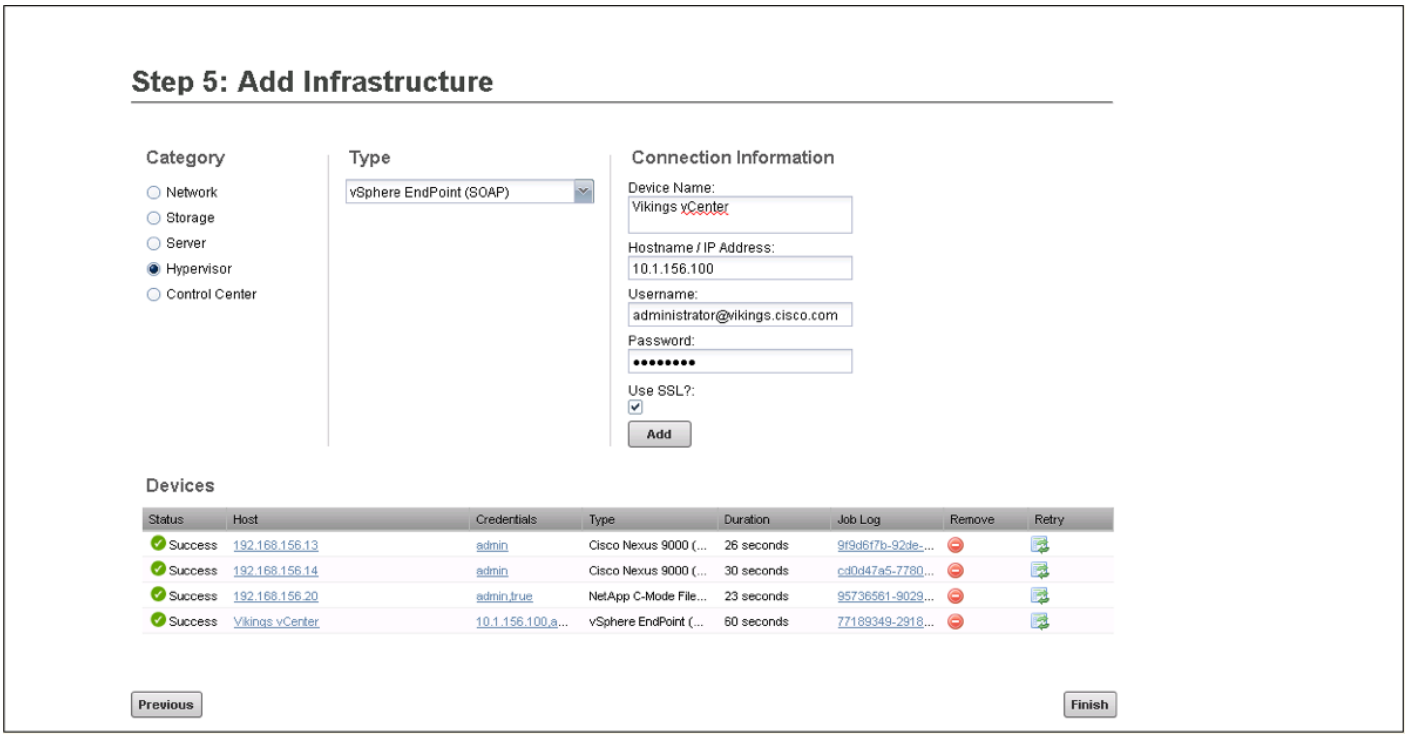

21. Click Finish.

The initial configuration of UCS Performance Manger to access the FlexPod environment is now complete.

Reference the Cisco UCS Performance Manager Administration Guide, Release 2.0.0 for further use and configuration of Cisco UCS Performance Manager.

NetApp Virtual Storage Console 6.2P1 Deployment Procedure

This section describes the deployment procedures for the NetApp VSC.

Virtual Storage Console 6.2P1 Pre-Installation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.3.2:

- Protocol licenses (NFS and FCP)
- FlexClone (for provisioning and cloning only)
- SnapRestore (for backup and recovery)
- SnapManager suite

Install Virtual Storage Console 6.2P1

To install the VSC 6.2P1 software, complete the following steps:

1. Build a VSC VM with Windows Server 2012 R2, 4GB RAM, two CPUs, and one virtual network interface in the <<var_ib_mgmt_vlan_id>> VLAN. The virtual network interface should be a VMXNET 3 adapter.

2. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain.

3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the User Interfaces and Infrastructure Feature on the VM.

4. Install all Windows updates on the VM.

5. Log in to the VSC VM as FlexPod admin user.

6. Download the x64 version of the Virtual Storage Console 6.2P1 from the NetApp Support site.

7. From the VMware Console, right-click the VSC .exe file downloaded in step 6 and select Run as administrator.

8. Select the appropriate language and click OK.

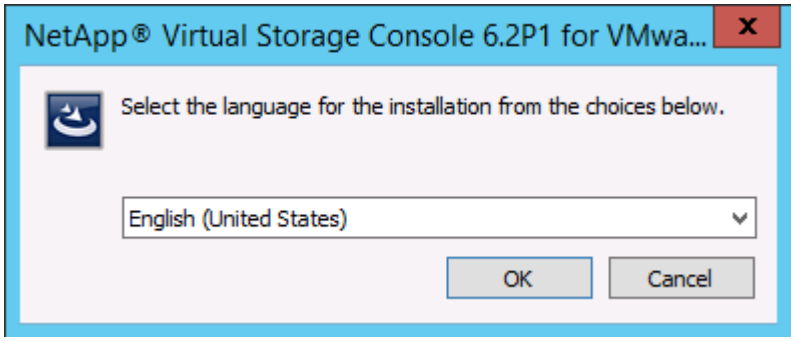

9. On the Installation wizard Welcome page, click Next.

10. Select the checkbox to accept the message, click Next.

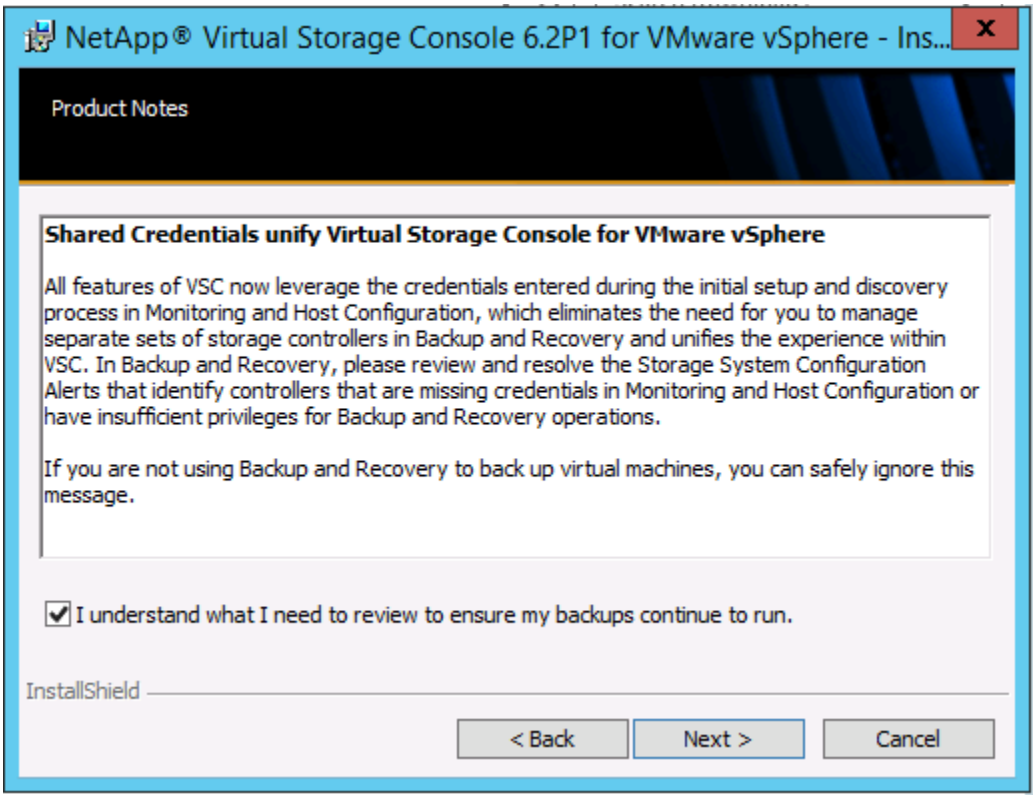

The Backup and Recovery capability requires an additional license.

11. Click Next to accept the default installation location.

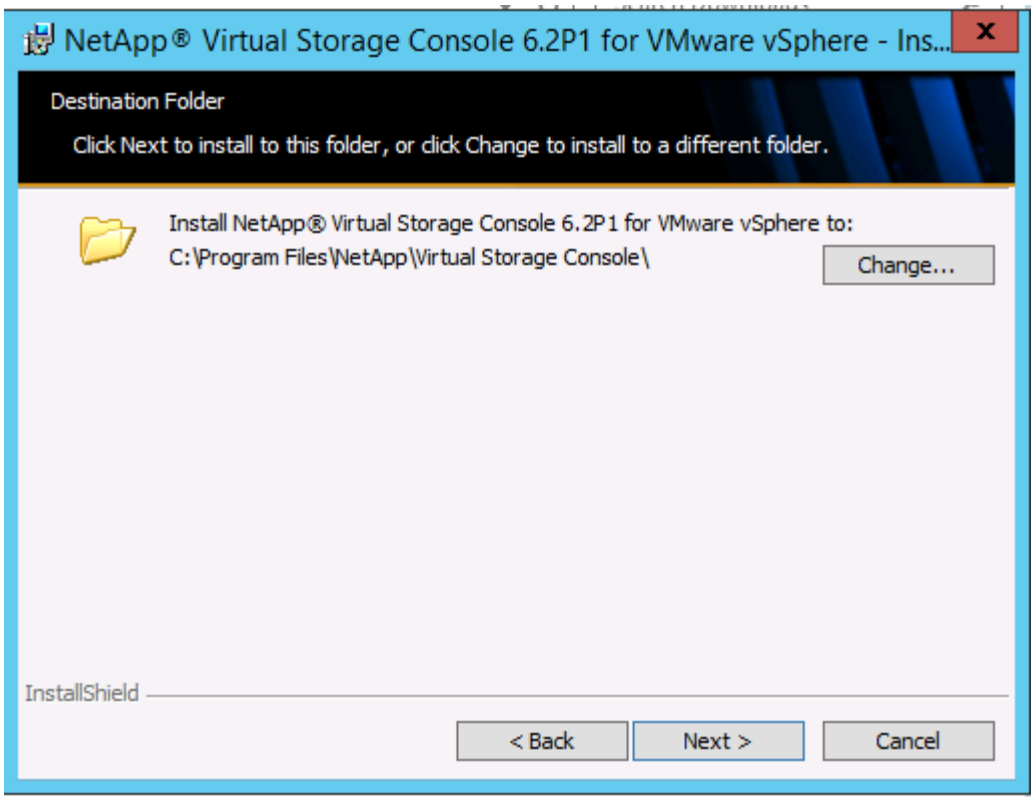

12. Click Install.

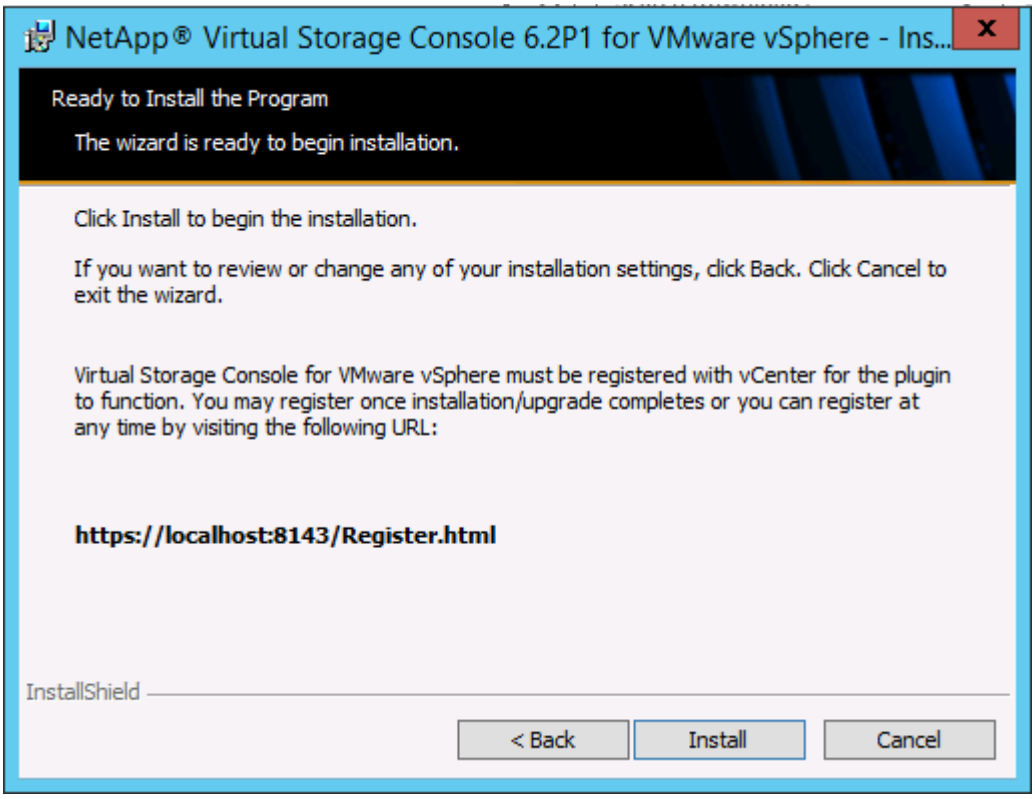

13. Click Finish.

Register VSC with vCenter Server

To register the VSC with the vCenter server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open https://localhost:8143/Register.html in Internet Explorer.

2. Click Continue to This Website (Not Recommended).

3. In the Plug-in Service Information section, select the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.

4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user or root), and user password for the vCenter server. Click Register to complete the registration.

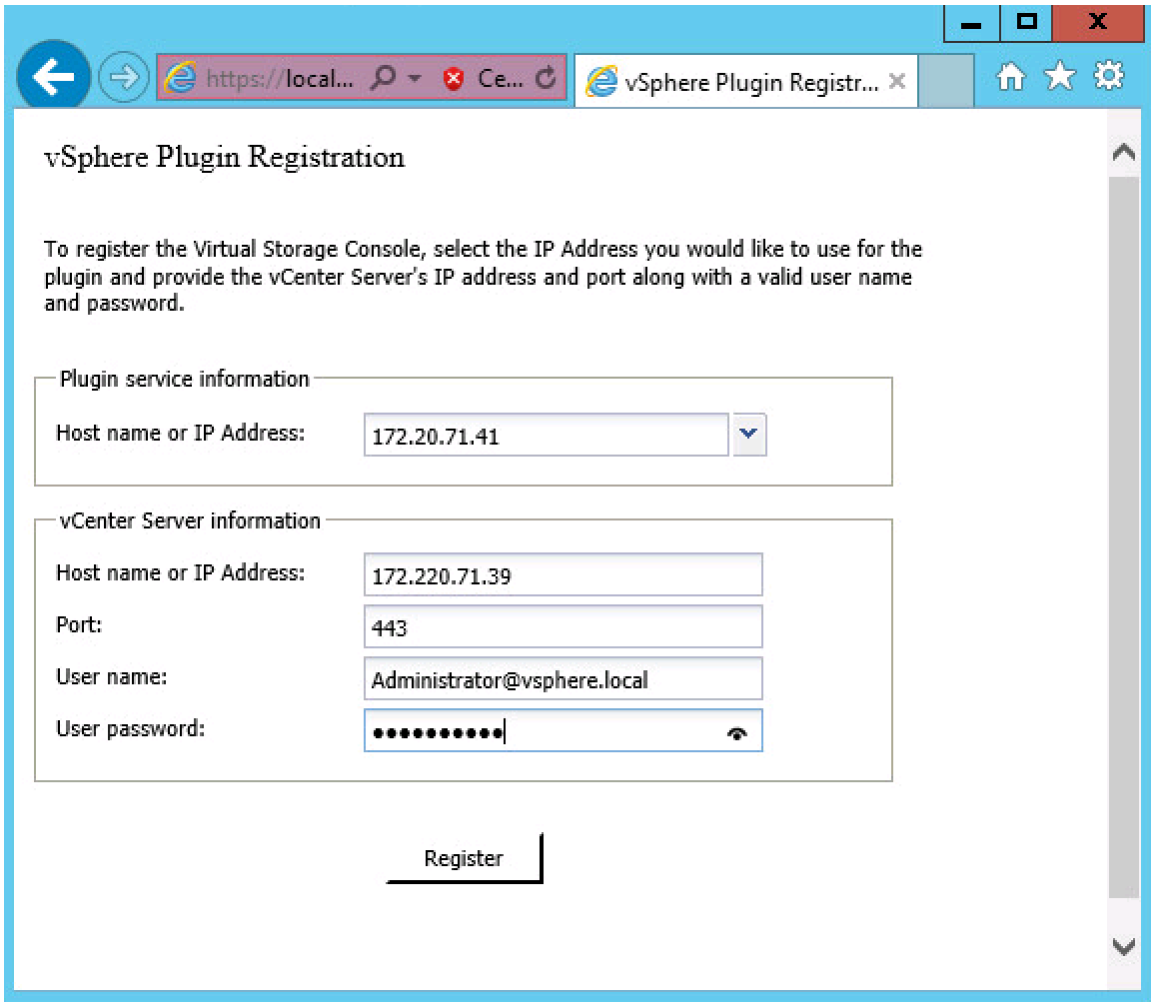

5. After successful registration, the storage controllers are discovered automatically.

The storage discovery process may take some time.

Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, complete the following steps:

1. Using the vSphere web client, log in to the vCenter Server as FlexPod admin user or root. If the vSphere web client was previously opened, close it and then reopen it.

2. In the Home screen, click the Home tab and click Virtual Storage Console.

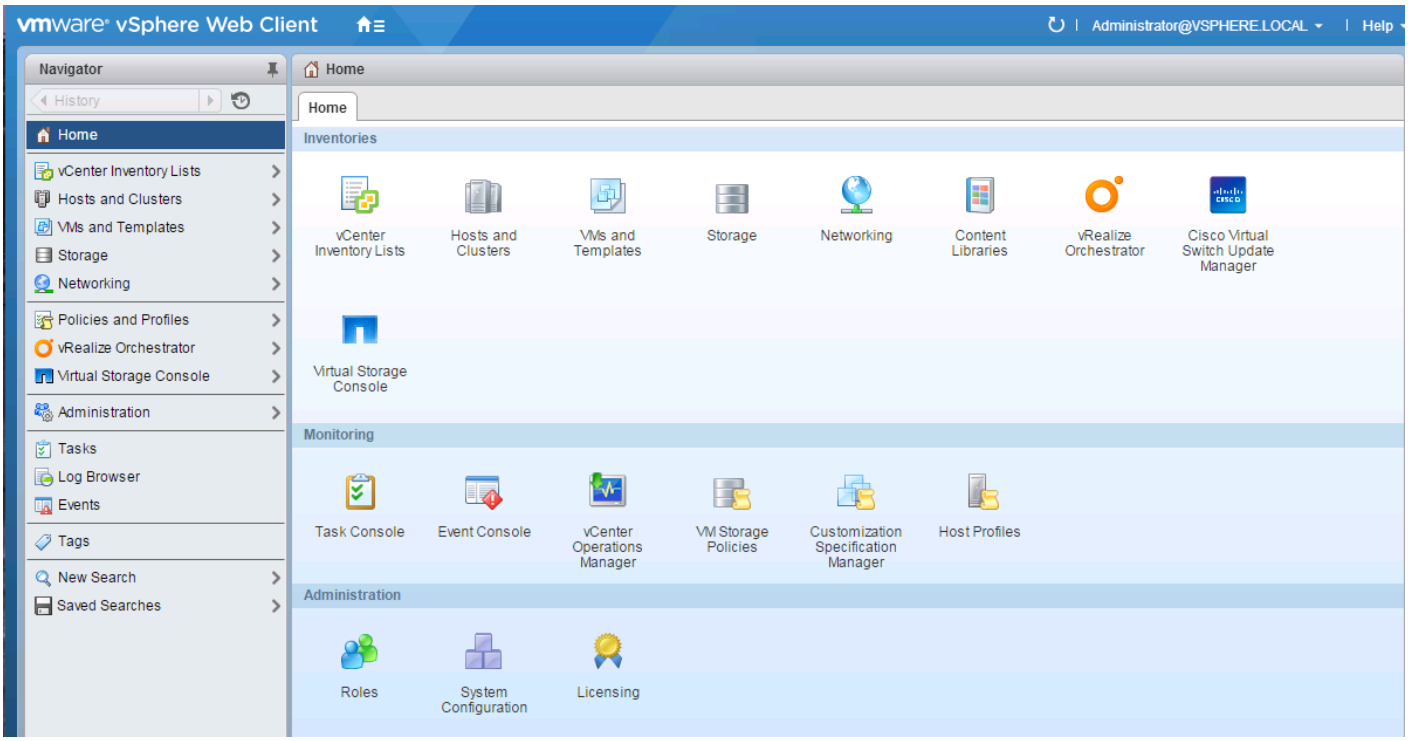

3. Select Storage Systems. Under the Objects tab, click Actions > Modify.

4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name and the admin password for password. Confirm that Use SSL to connect to this storage system is selected. Click OK.

5. Click OK to accept the controller privileges.

Optimal Storage Settings for ESXi Hosts

VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the Home screen, click vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for These Hosts.

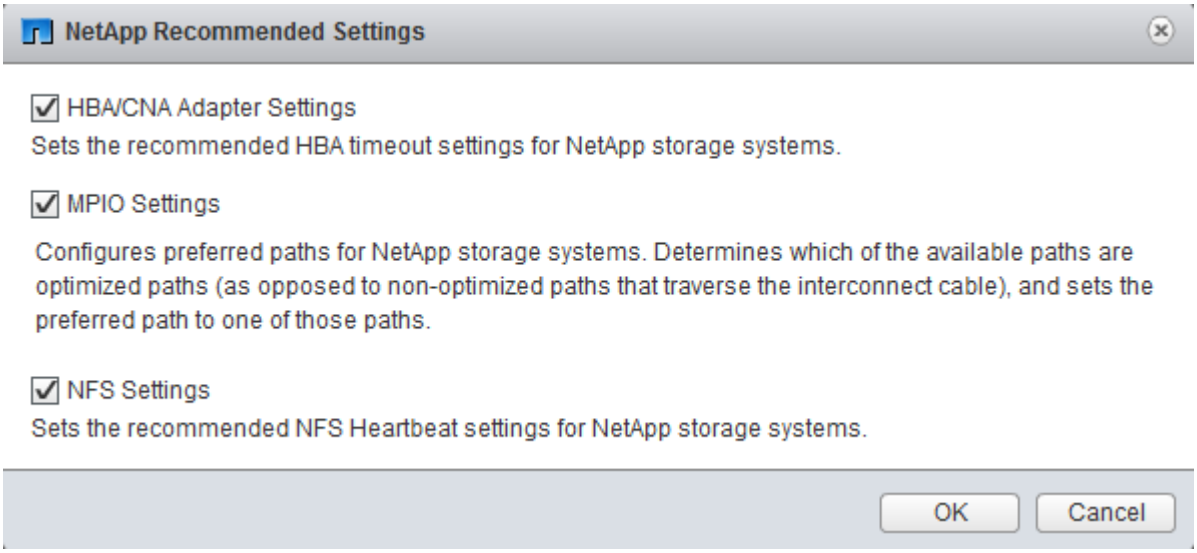

2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings. This functionality sets values for HBAs and CNAs, sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).

3. Click OK.

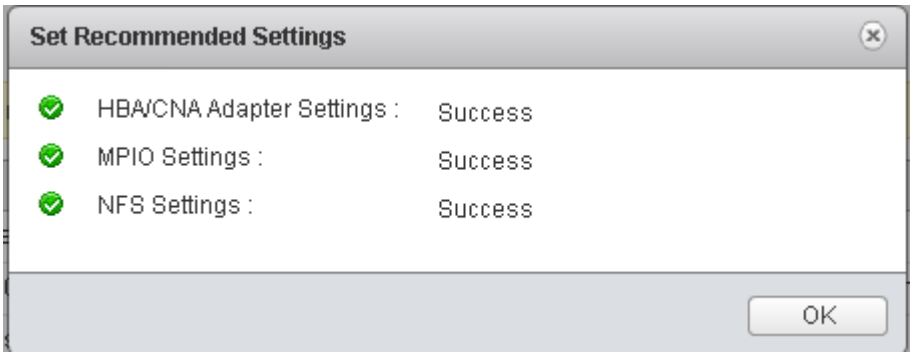

4. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.

VSC 6.2P1 Backup and Recovery

Prerequisites to Use Backup and Recovery Capability

Before you begin using the backup and recovery capability to schedule backups and restore your datastores, VMs, or virtual disk files, you must confirm that the storage systems that contain the datastores and VMs for which you are creating backups have valid storage credentials.

If you plan to use the SnapMirror update option, add all of the destination storage systems with valid storage credentials.

Backup and Recovery Configuration

1. From the Home screen, select the Home tab and click Storage.

2. On the left, expand the Datacenter and select Datastores.

3. Right-click the datastore that you want to backup. Select NetApp VSC > Backup > Schedule Backup Job.

If you prefer a one-time backup, then select Backup Now instead of Schedule Backup.

4. Type a backup job name and description.

If you want to create a VMware snapshot for each backup, select Perform VMware Consistency Snapshot in the options pane.

5. Click Next.

6. Select any options to include in the backup.

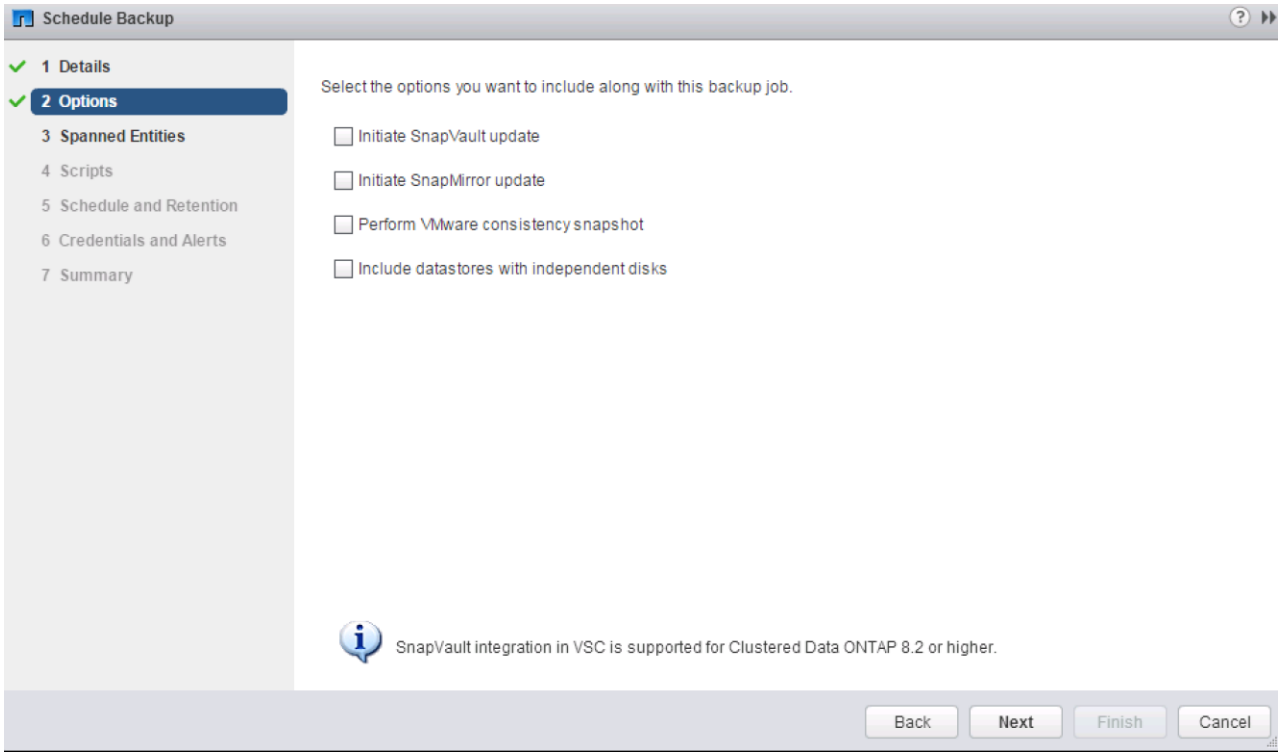

7. Click Next on the Options screen.

8. Click Next on the Spanned Entities screen.

9. Select one or more backup scripts if available and click Next in the Scripts screen.

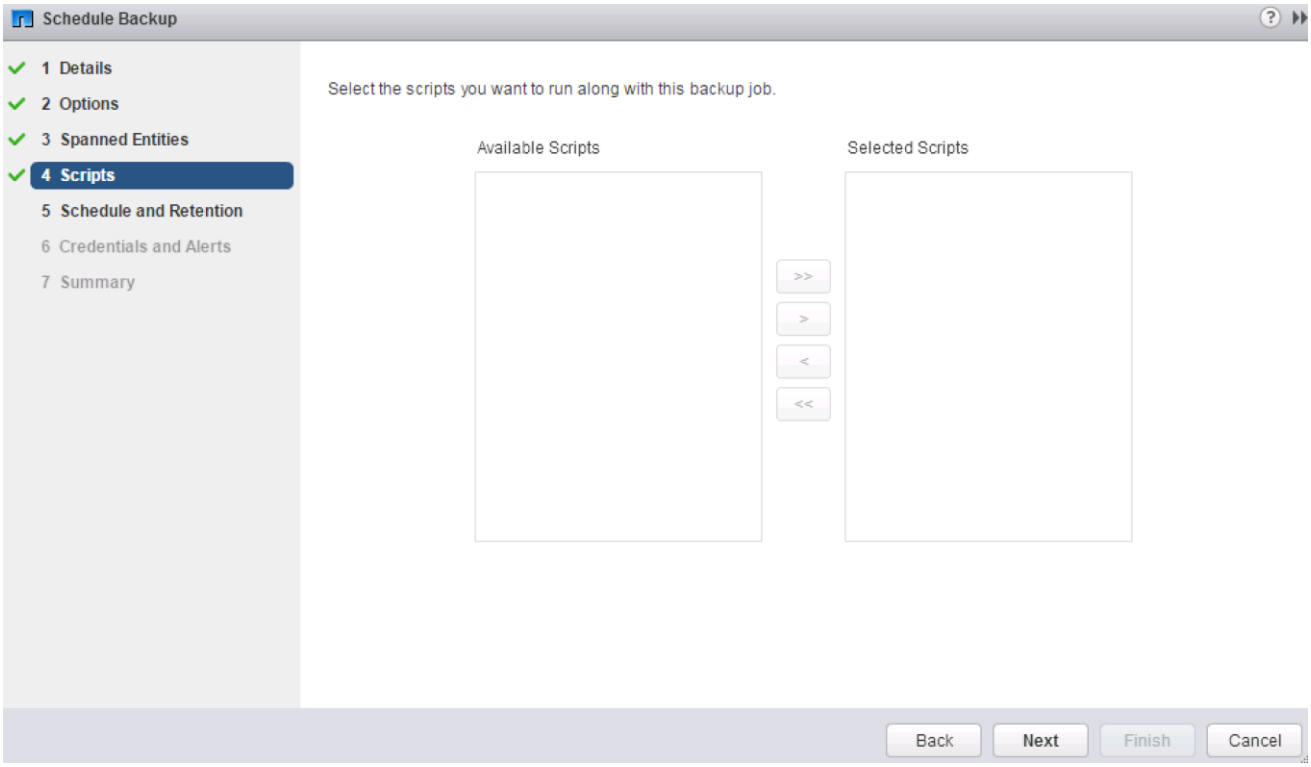

10. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click Next.

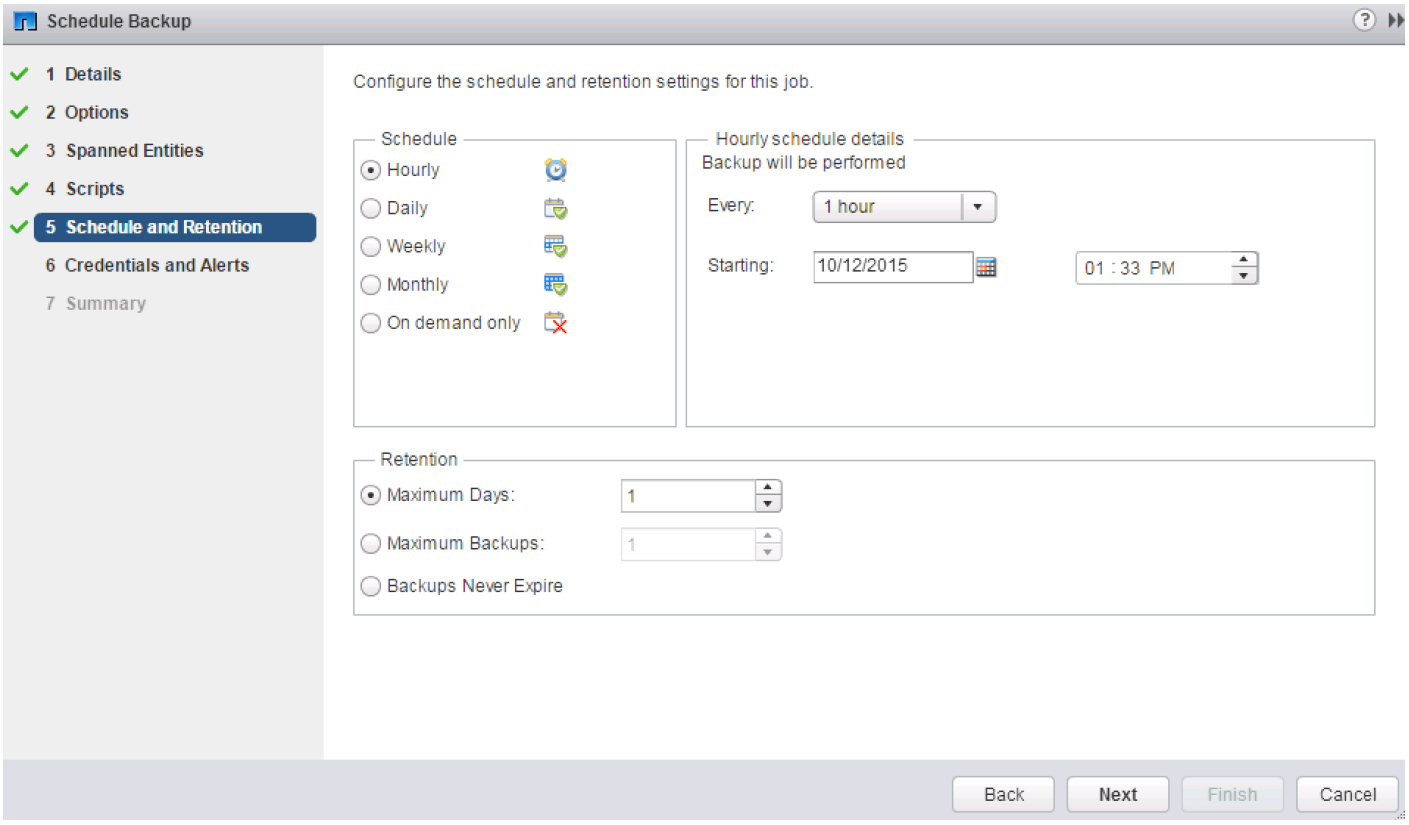

11. Use the default vCenter credentials or type the user name and password for the vCenter server and click Next.

12. Specify backup retention details as per requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate them. Click Next.

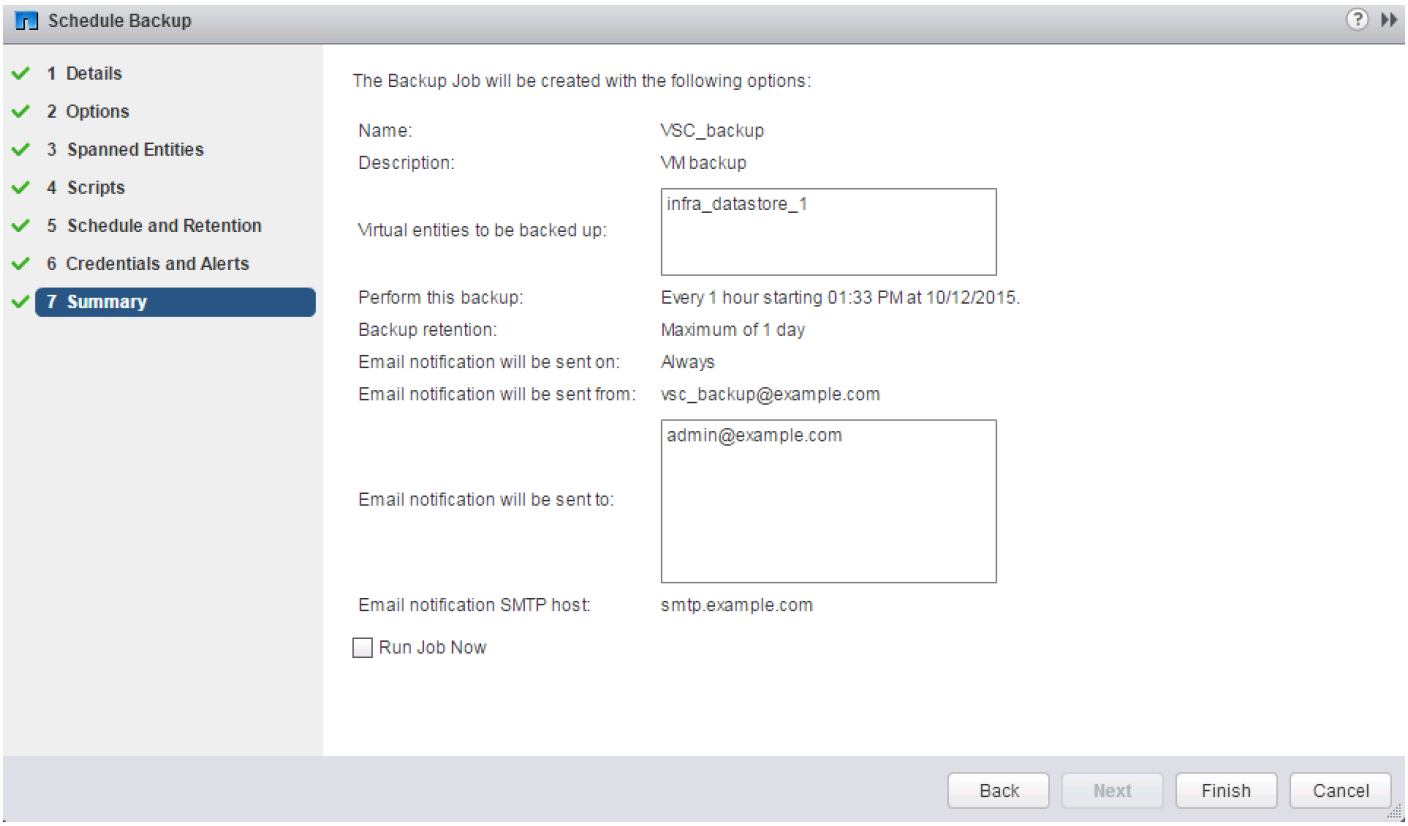

13. Review the summary page and click Finish. If you want to run the job immediately, select the Run Job Now option and then click Finish.

14. Click OK.

15. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by entering the following command:

```
volume modify -volume infra_datastore_1 -snapshot-policy none
```

16. Also, to delete any existing automatic Snapshot copies that have been created on the volume, enter the following command:

```
volume snapshot show -volume infra_datastore_1
volume snapshot delete -volume infra_datastore_1 -vserver Infra-SVM -snapshot <snapshot name>
```

The wildcard character (*) can be used in snapshot names in the previous command.

OnCommand Unified Manager 6.3P2

OnCommand Unified Manager OVF Deployment

To install the OnCommand Unified Manager, complete the following steps:

Download and review the OnCommand Unified Manager Installation and Setup Guide.

1. Download OnCommand Unified Manager version 6.3P2 (OnCommandUnifiedManager-6.3P2.ova) from the OnCommand download site.
2. Log in to the vSphere web client. Go to vCenter > VMs and Templates.
3. At the top of the center pane, select Actions > Deploy OVF Template.

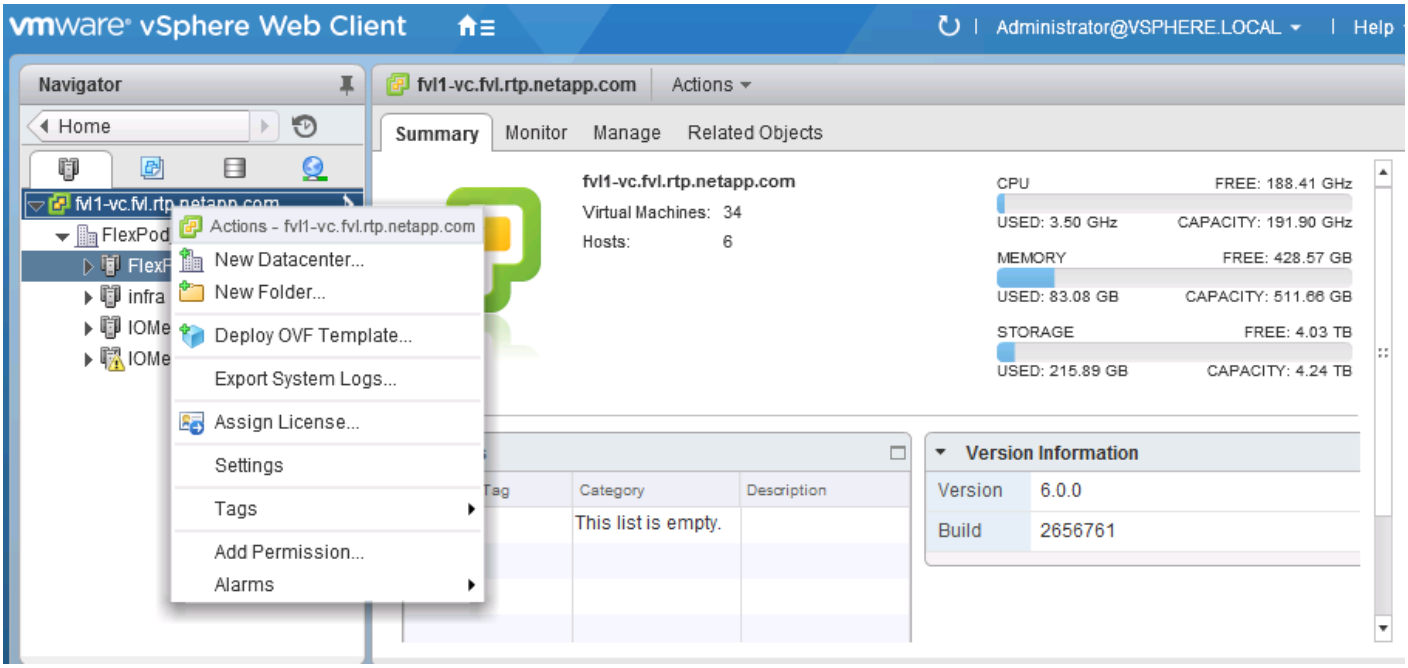

4. Browse the `.ova` file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.
5. Select Accept Extra Configuration Options, and click Next.
6. Review the deployment details, and click Next.

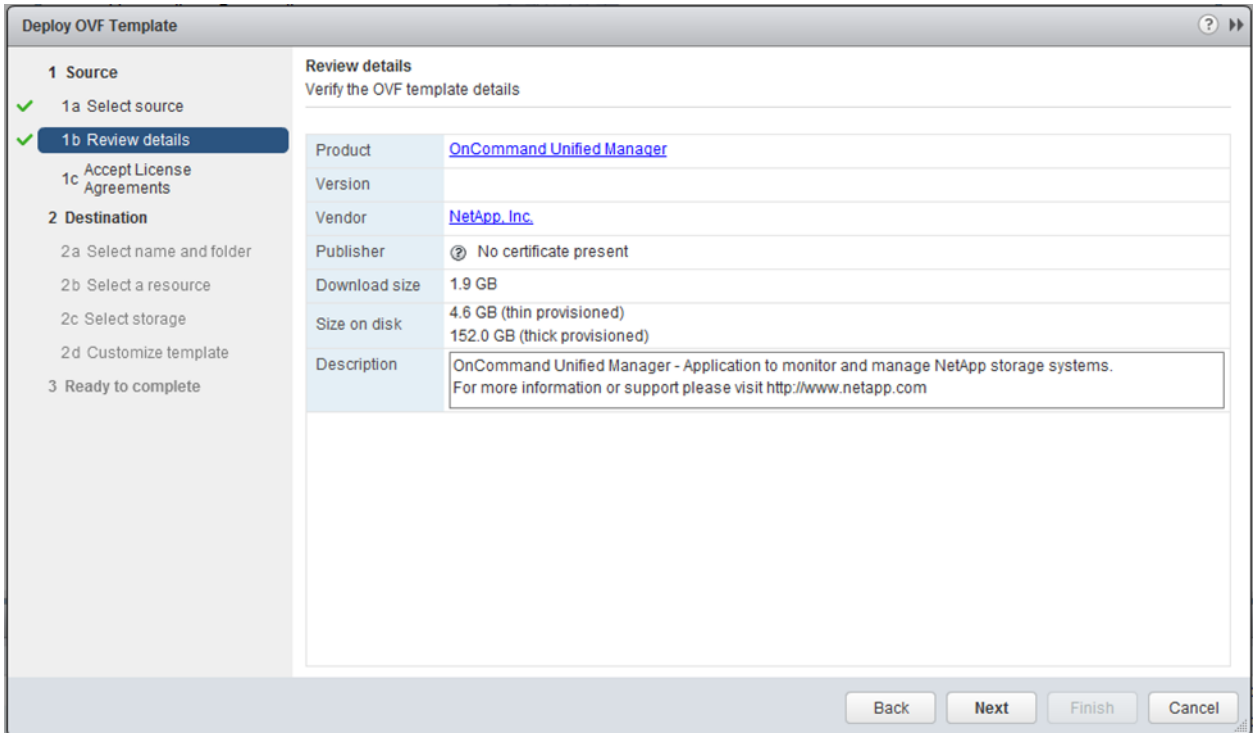

7. Read the end-user license agreement (EULA), and then click Accept to accept the agreement. Click Next.

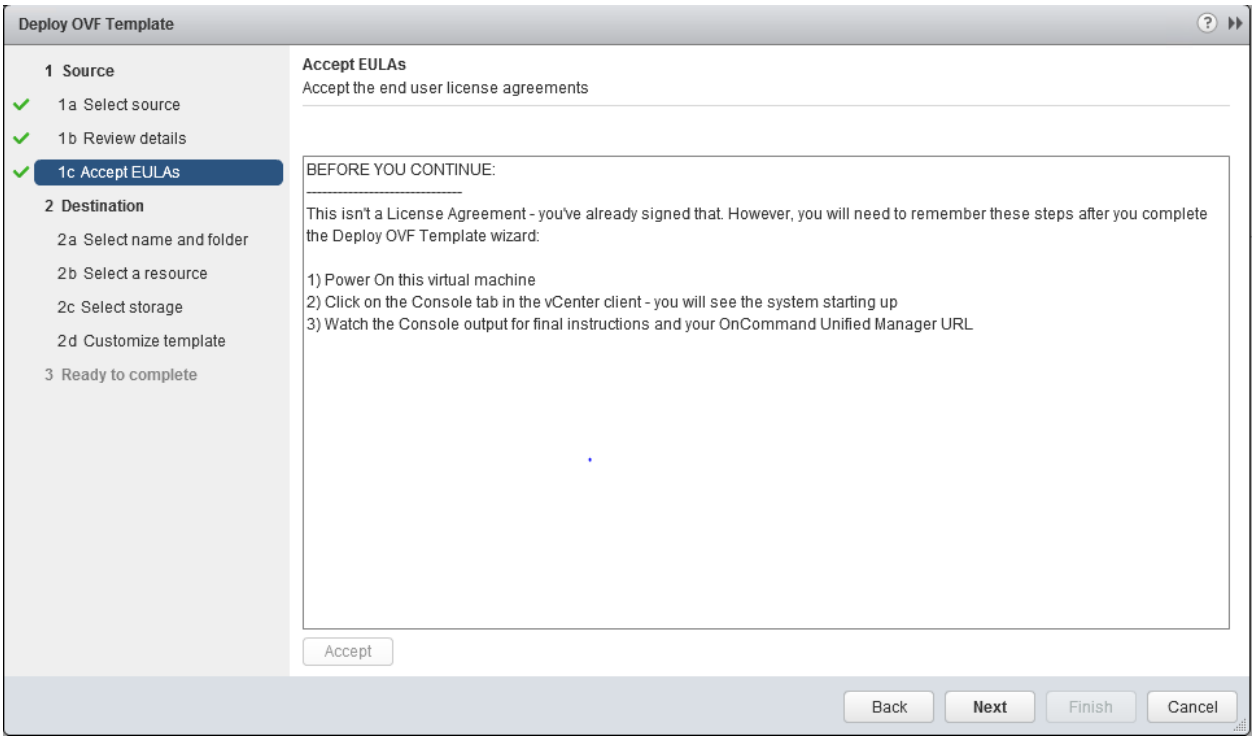

8. Enter the name of the VM and select the `FlexPod_DC` folder to hold it. Click Next to continue.

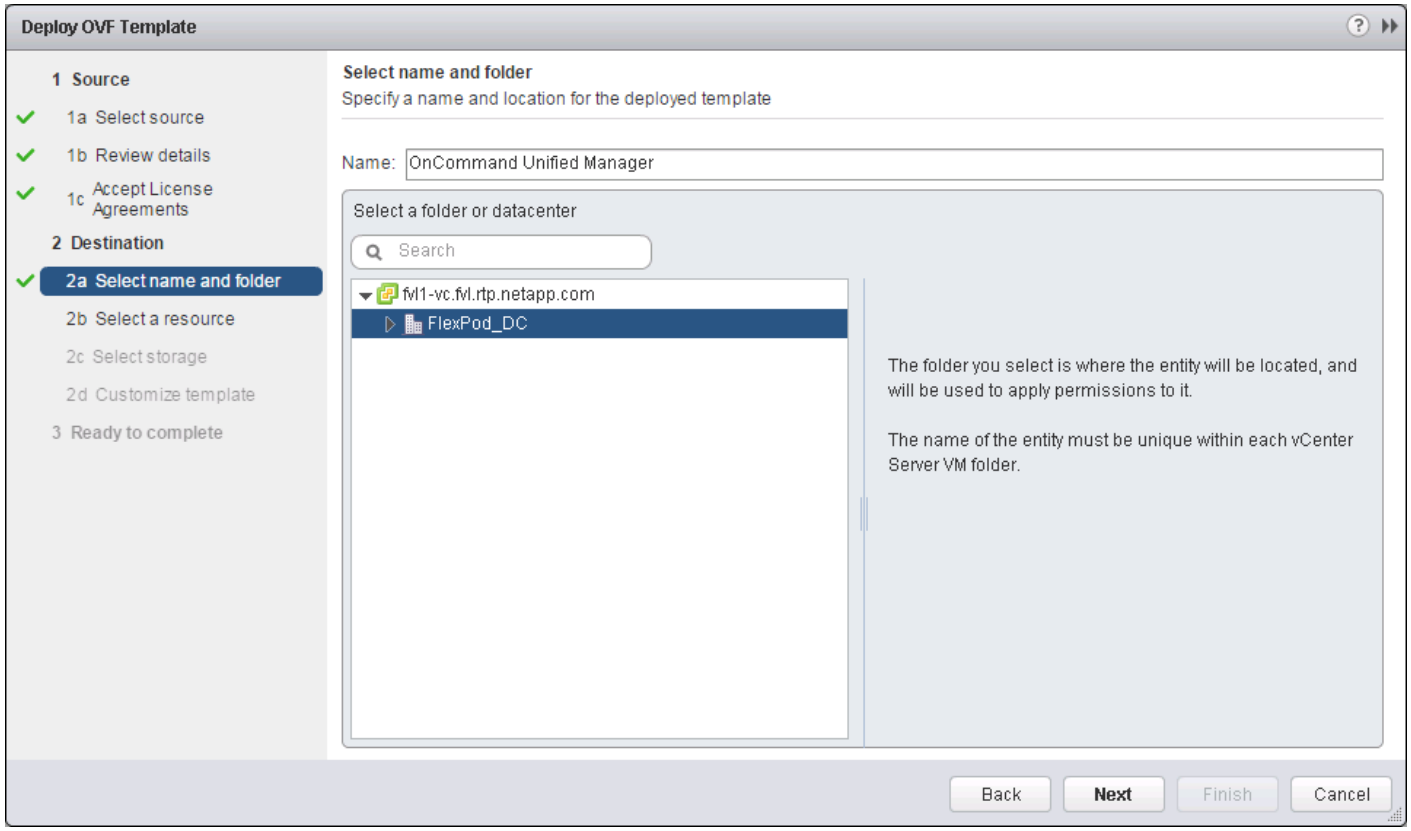

9. Select `FlexPod_Management` within the `FlexPod_DC` datacenter as the destination compute resource pool to host the VM. Click Next.

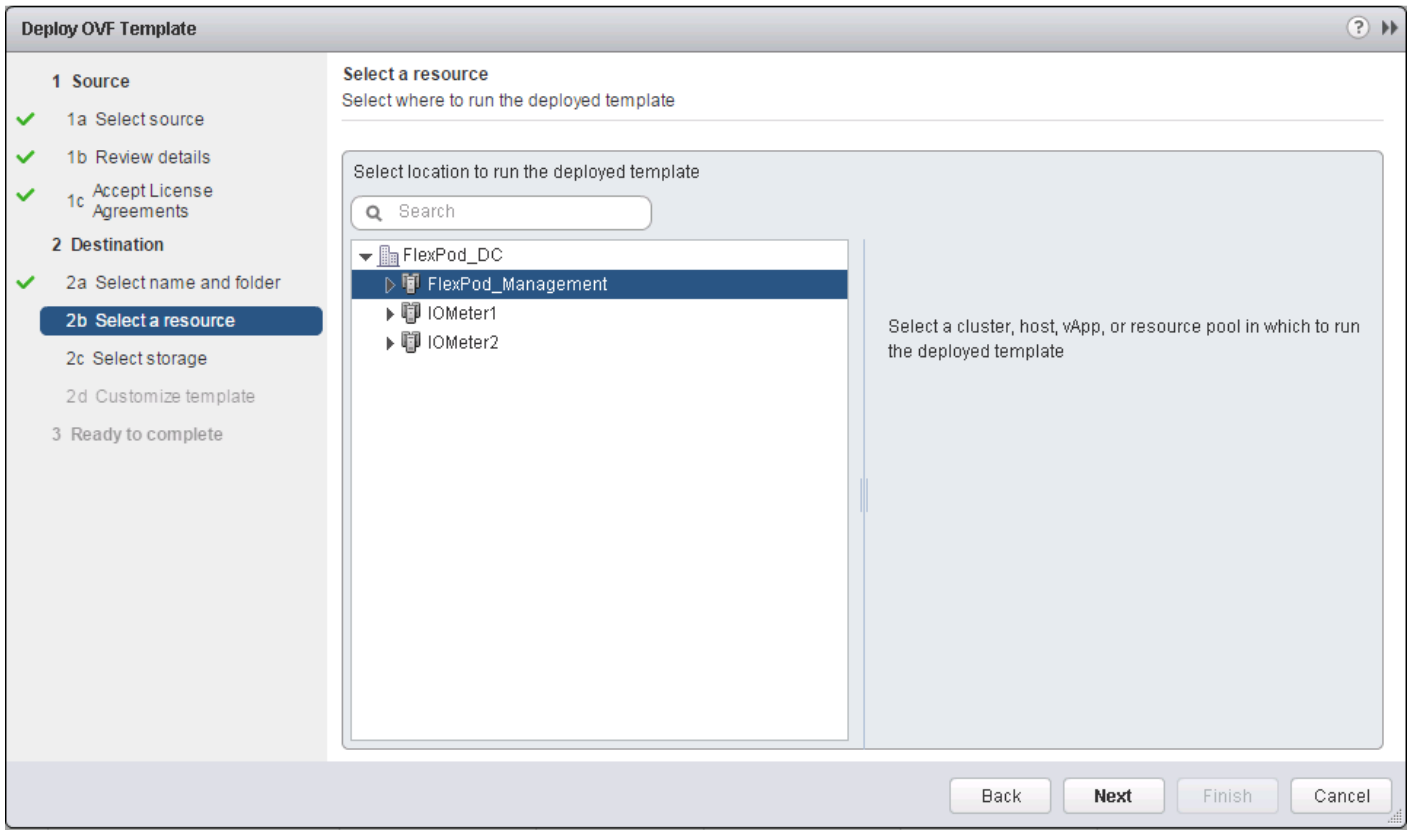

10. Select `infra_datastore_1` as the storage target for the VM and select Thin Provision as the virtual disk format. Click Next.

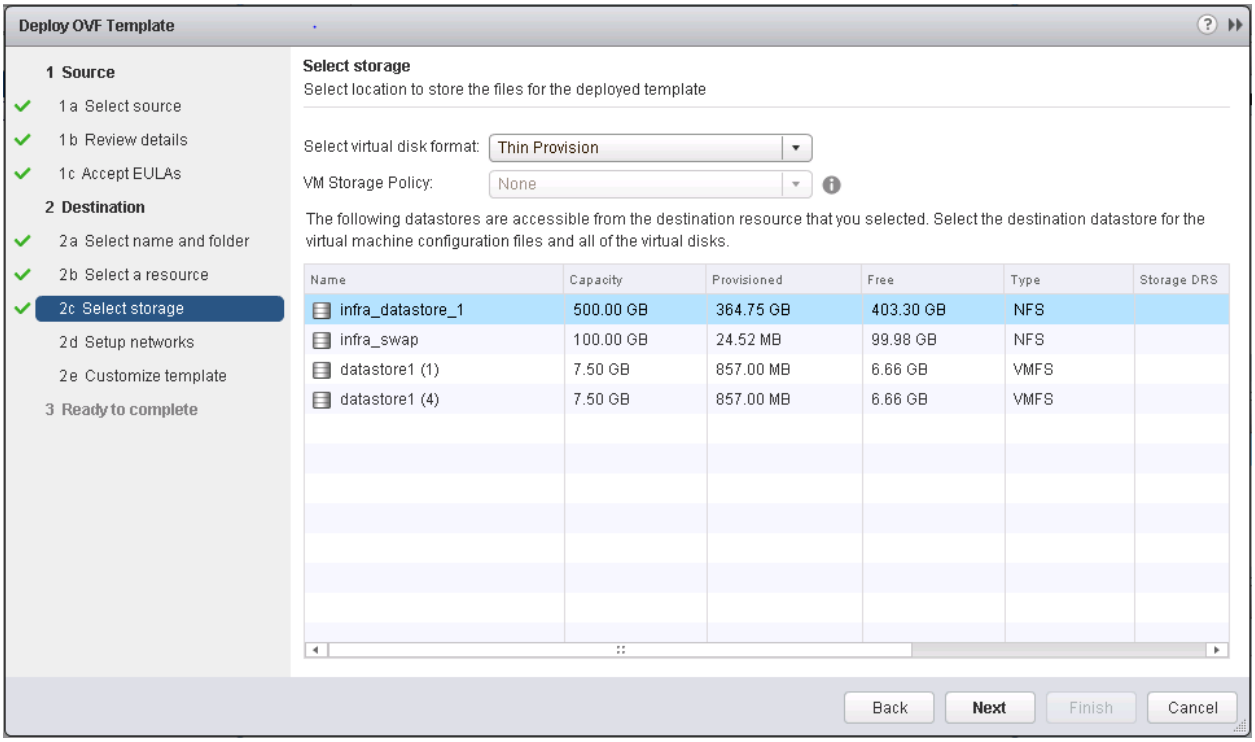

11. Select `IB-MGMT-VLAN` as the destination network for the `nat` source network. Click Next.

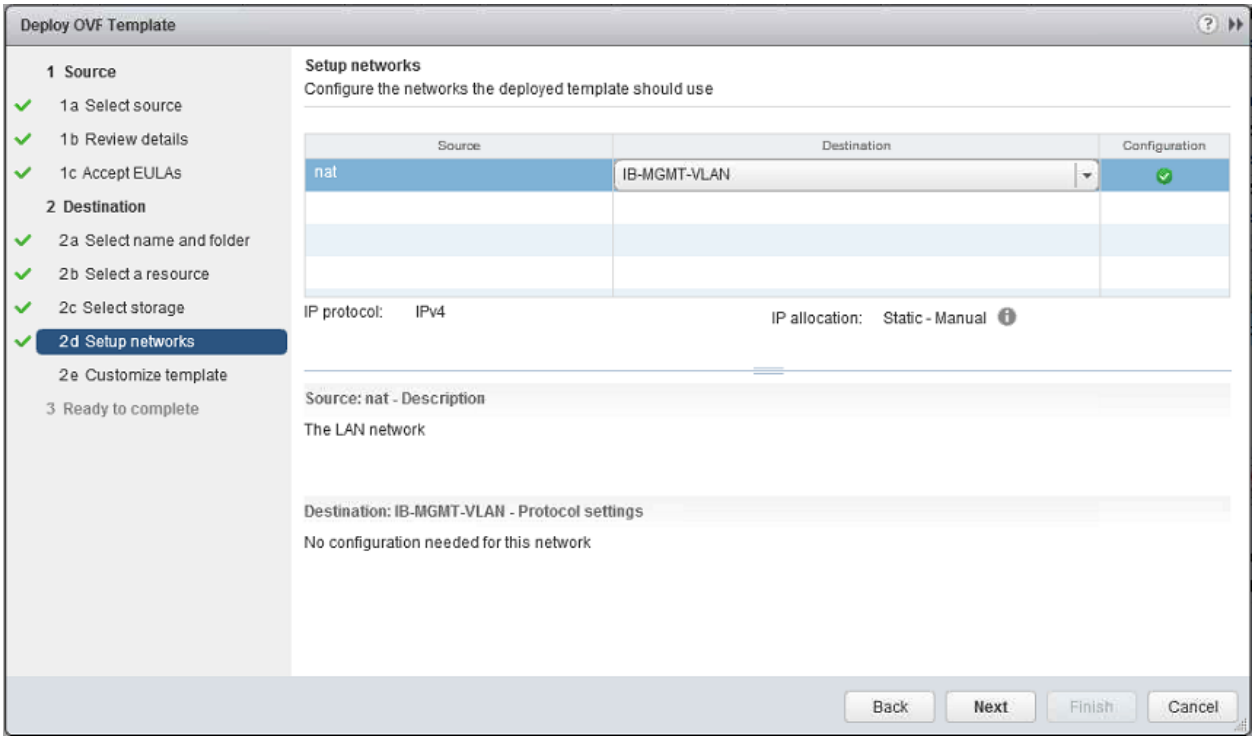

12. Complete the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS fields. Click Next.

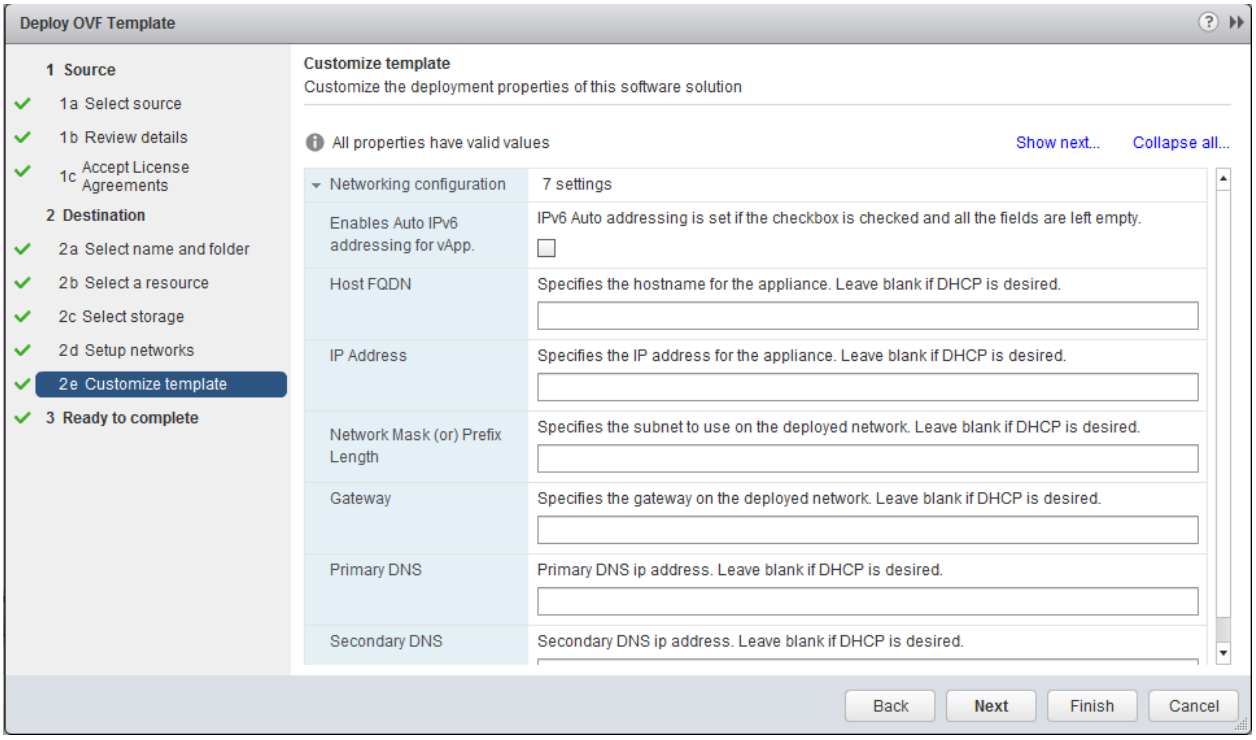

13. Clear the Power On After Deployment checkbox.

14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration information.

15. In the left pane, select vCenter -> Virtual Machines. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.

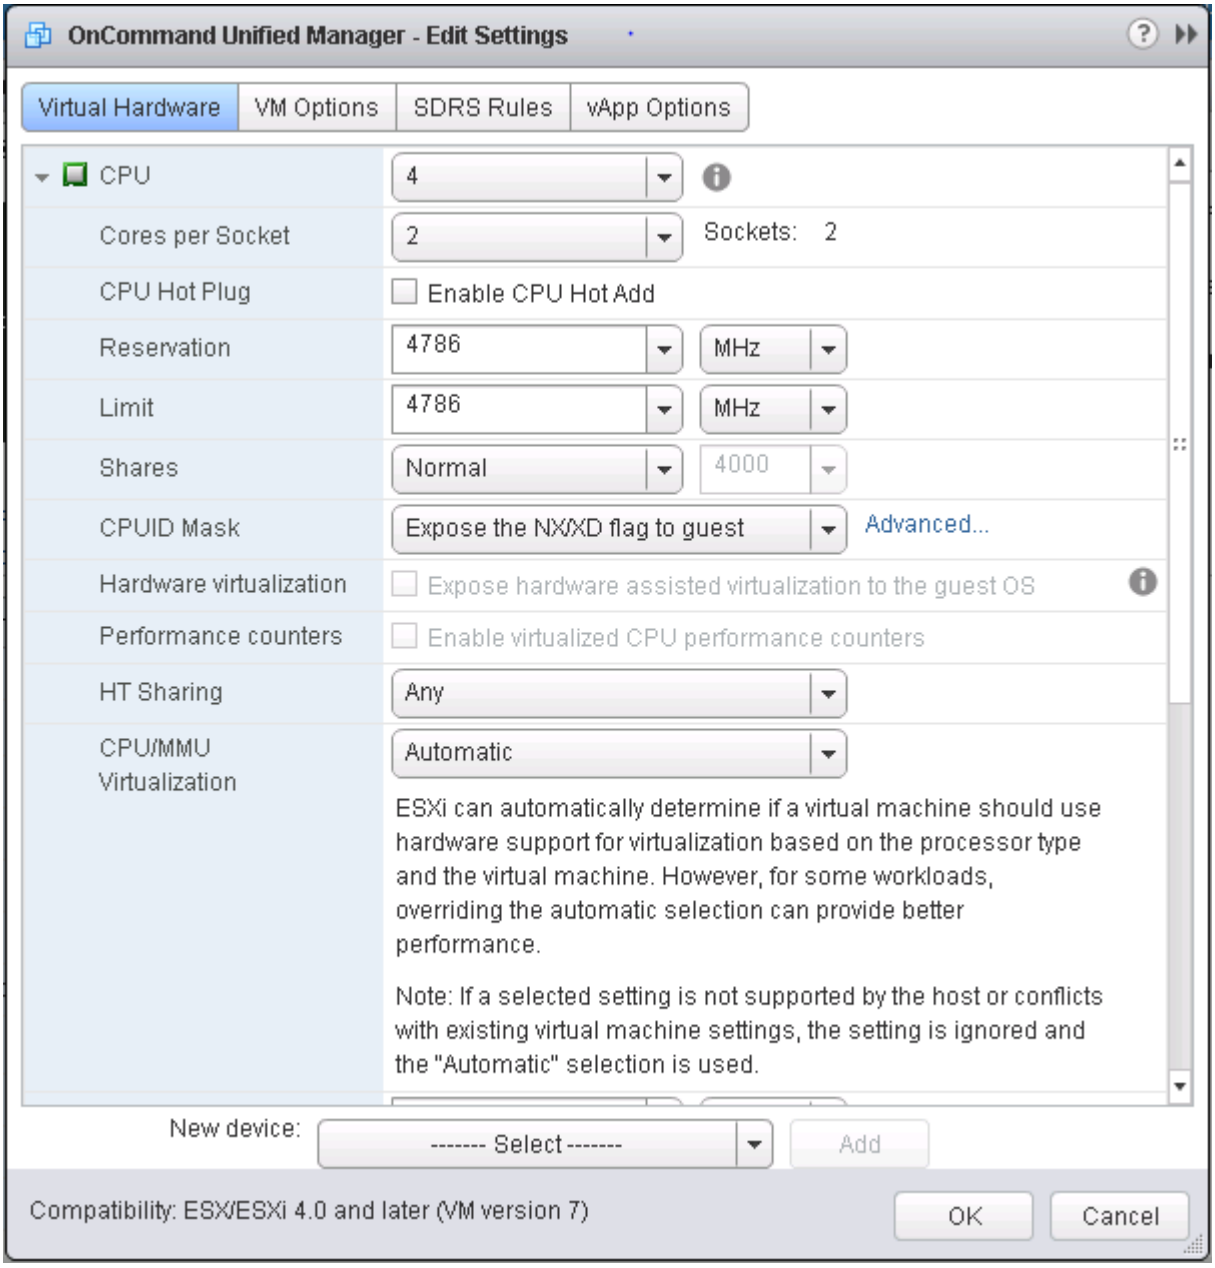

16. Size the VM's CPU and memory parameters according to the OnCommand Unified Manager 6.3 Installation and Setup Guide.

17. Click OK to accept the changes.

18. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Unified Manager Basic Setup

To setup the OnCommand Unified Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.

2. In the VMRC window, select VMRC > Manage > Install VMware Tools. VMware Tools will install in the VM.

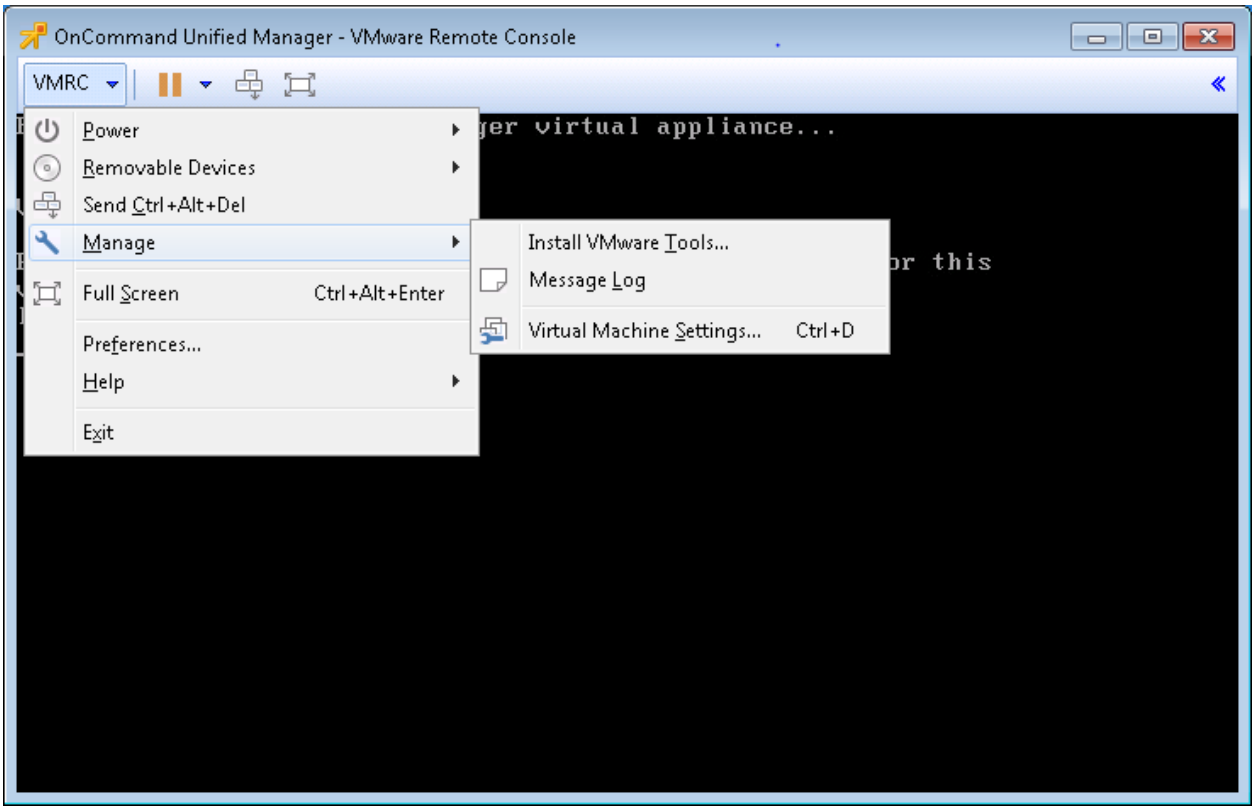

3. Set up OnCommand Unified Manager by answering the following questions in the console window:

```
Geographic area: <<Enter the number corresponding to your time zone>>
```

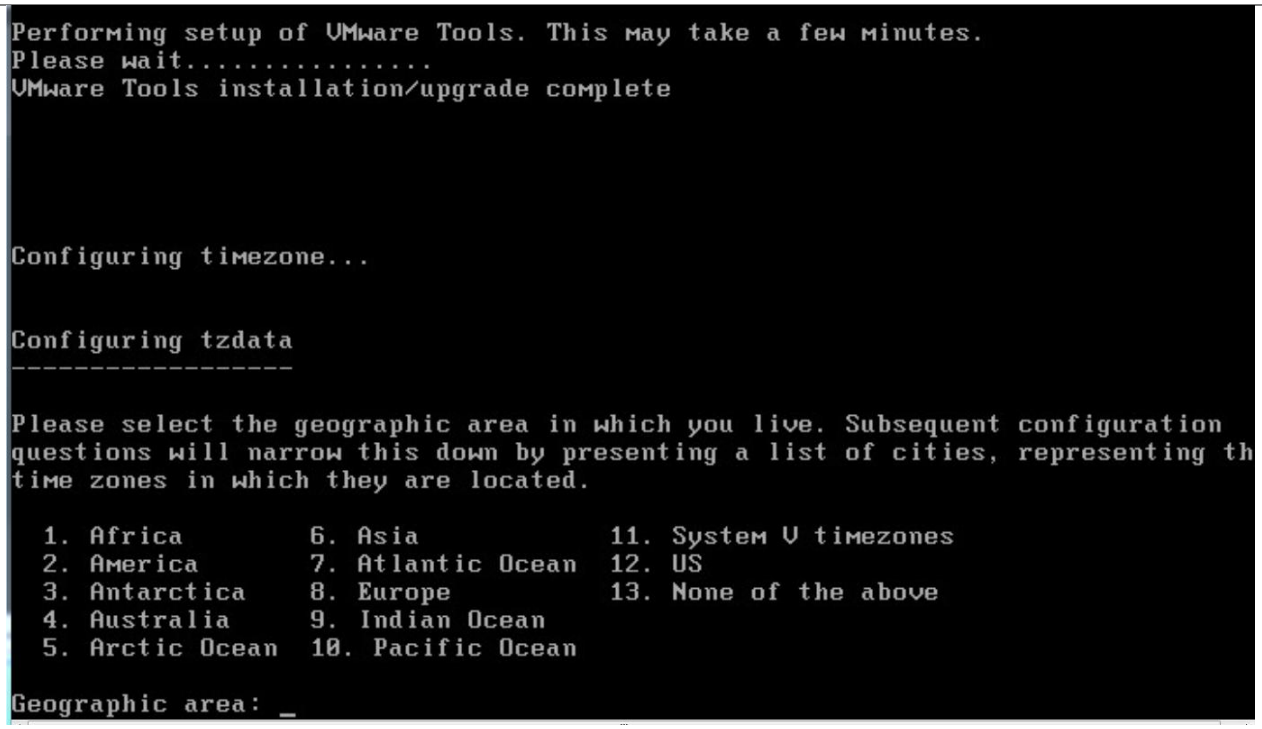

These commands complete the network configuration checks, generate SSL certificates for HTTPS and start the OnCommand Unified Manager services.

4. To Create a Maintenance User account, run the following commands:

The maintenance user manages and maintains the settings on the OnCommand Unified Manager virtual appliance.

```
Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>
```

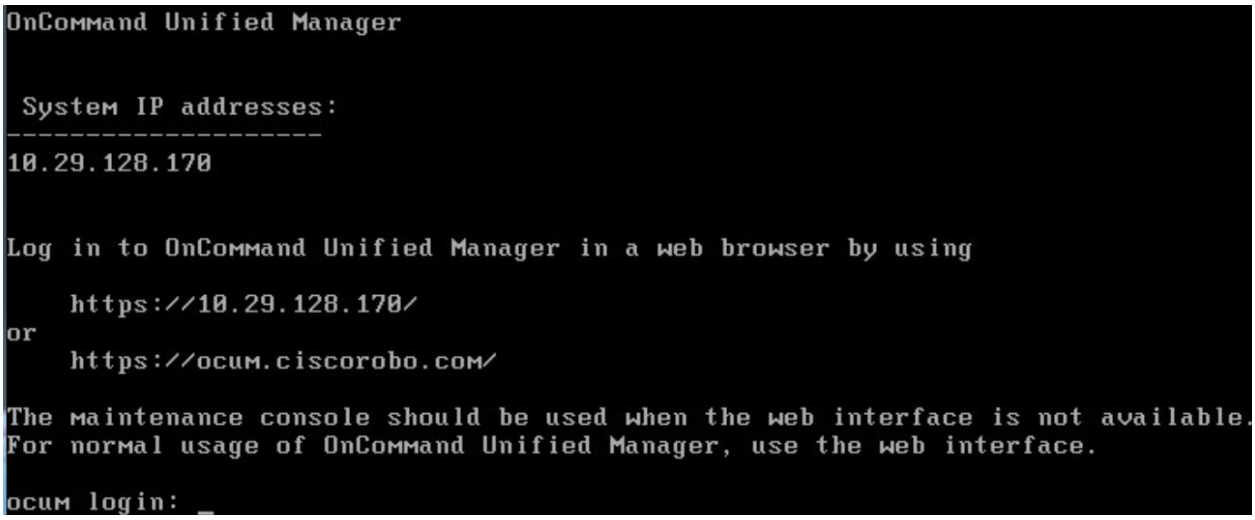

5. With a web browser, navigate to the OnCommand Unified Manager using URL: `https://<<var_oncommand_server_ip>>`.

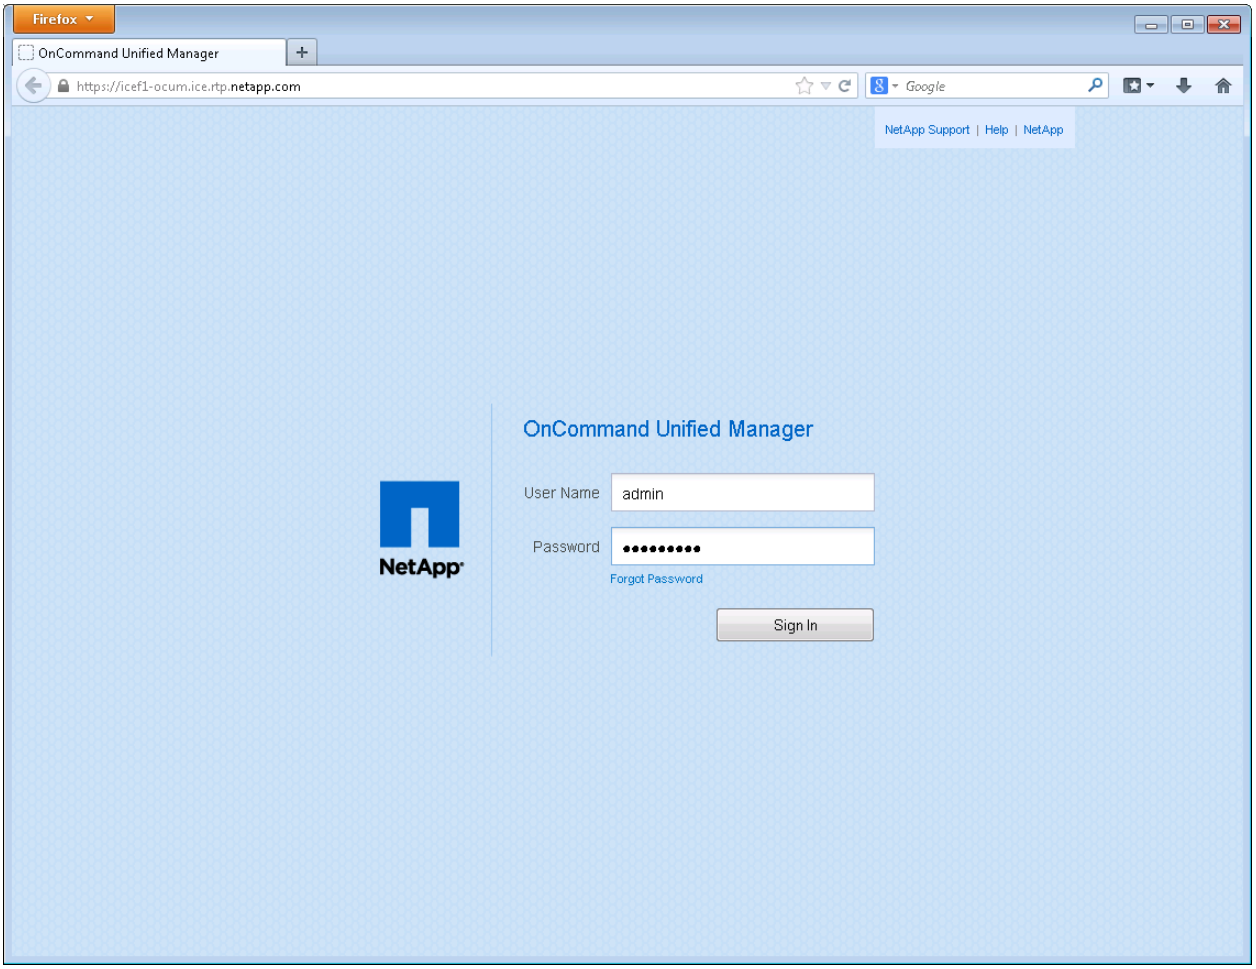

6. Log in using the maintenance user account credentials.

7. Select `Yes` to enable AutoSupport capabilities.

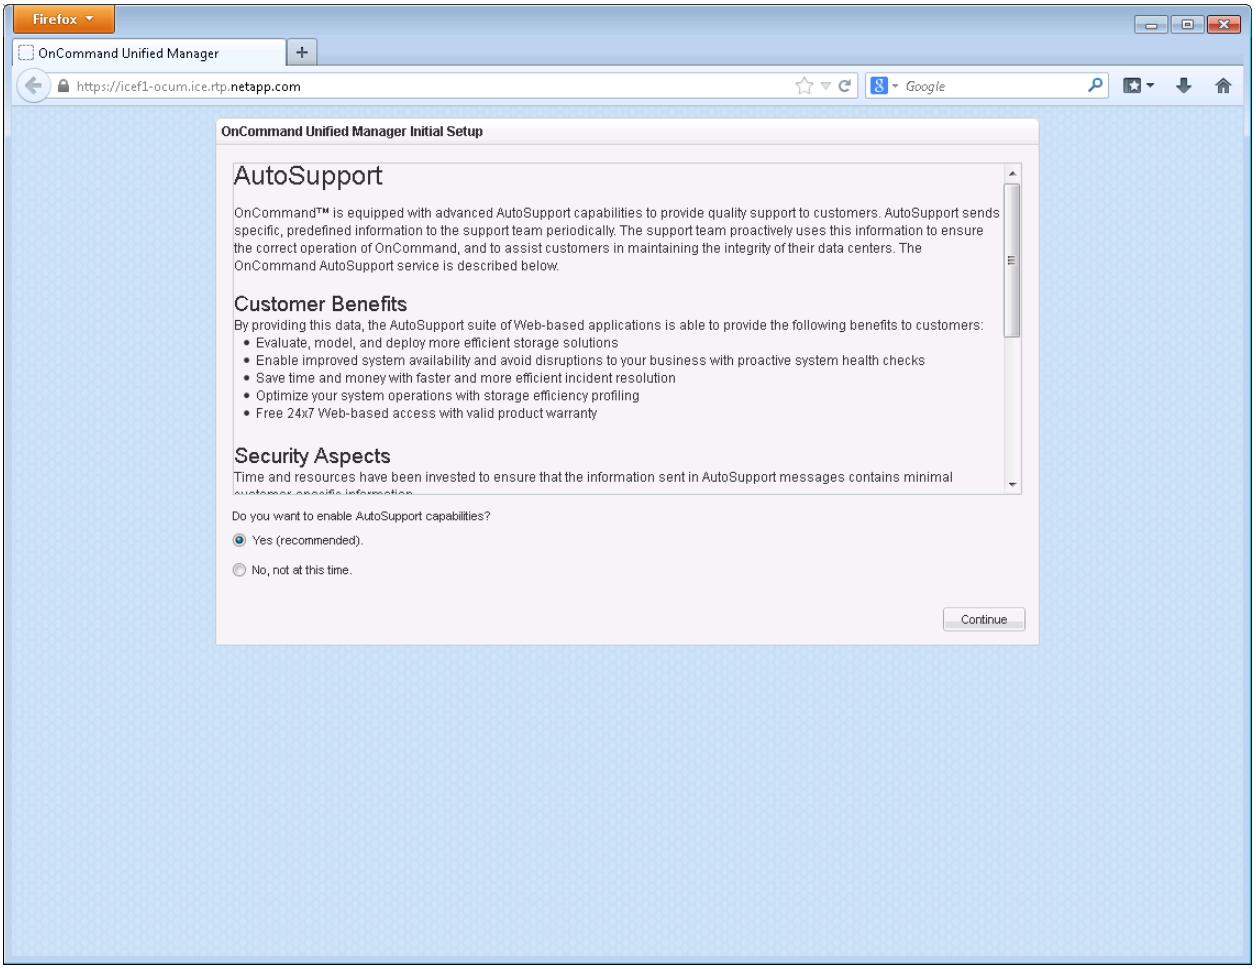

8. Click Continue.
9. Enter the NTP server IP address `<<var_switch_a_ntp_ip>>`.
10. Enter the storage admin e-mail address `<<var_storage_admin_email>>`.
11. Enter the SMTP server host name.

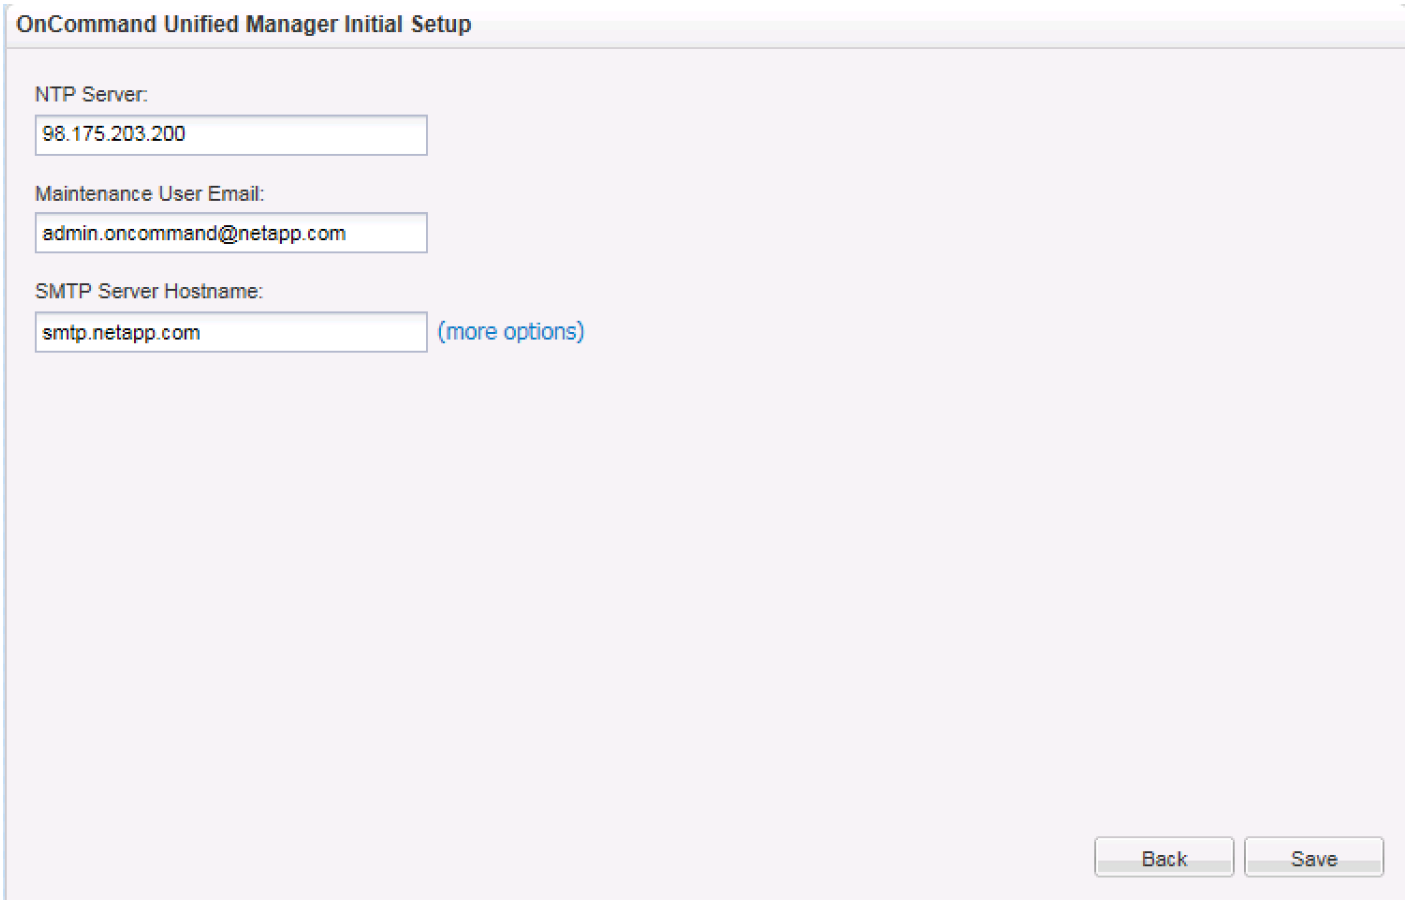

12. Click Save.

13. Click Add Cluster.

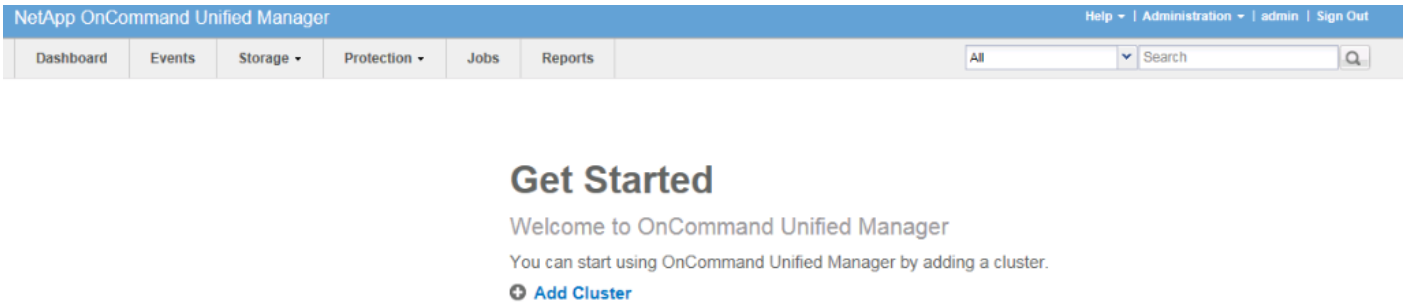

14. Provide the cluster management IP address, user name, password, protocol, and port.

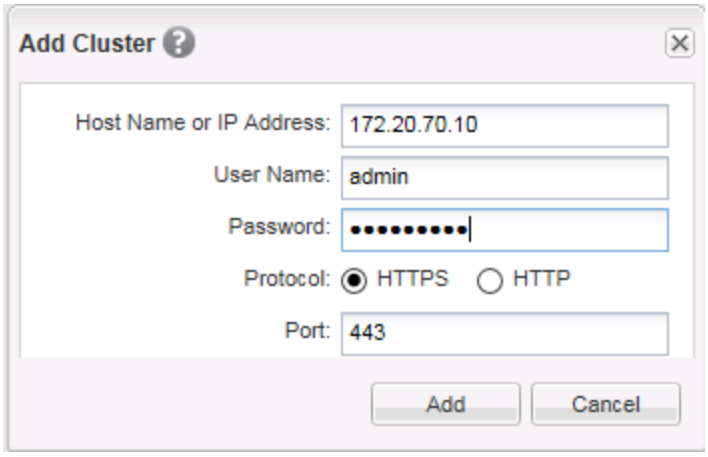

15. Click Add.

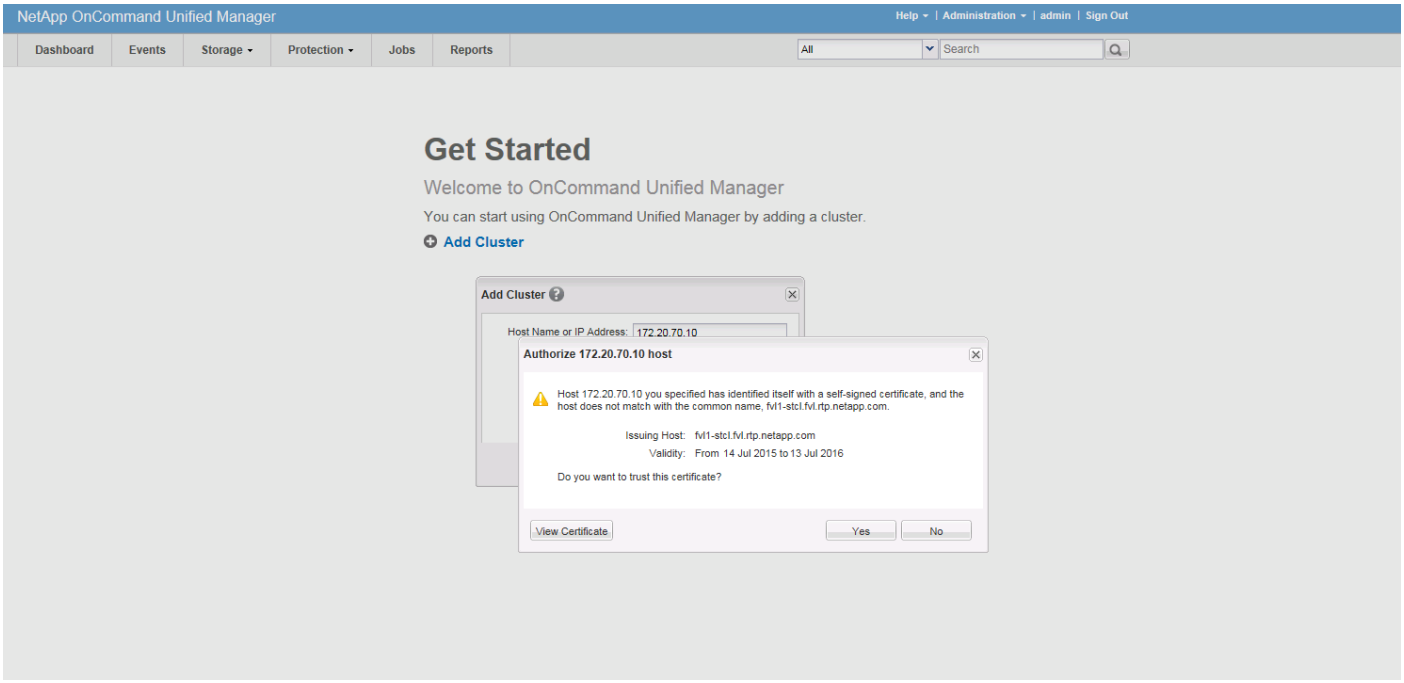

16. Click Yes to trust the certificate from the controller.

The Cluster Add operation might take a couple of minutes.

17. After the cluster is added, it can be accessed by clicking on the Storage tab and selecting Clusters.

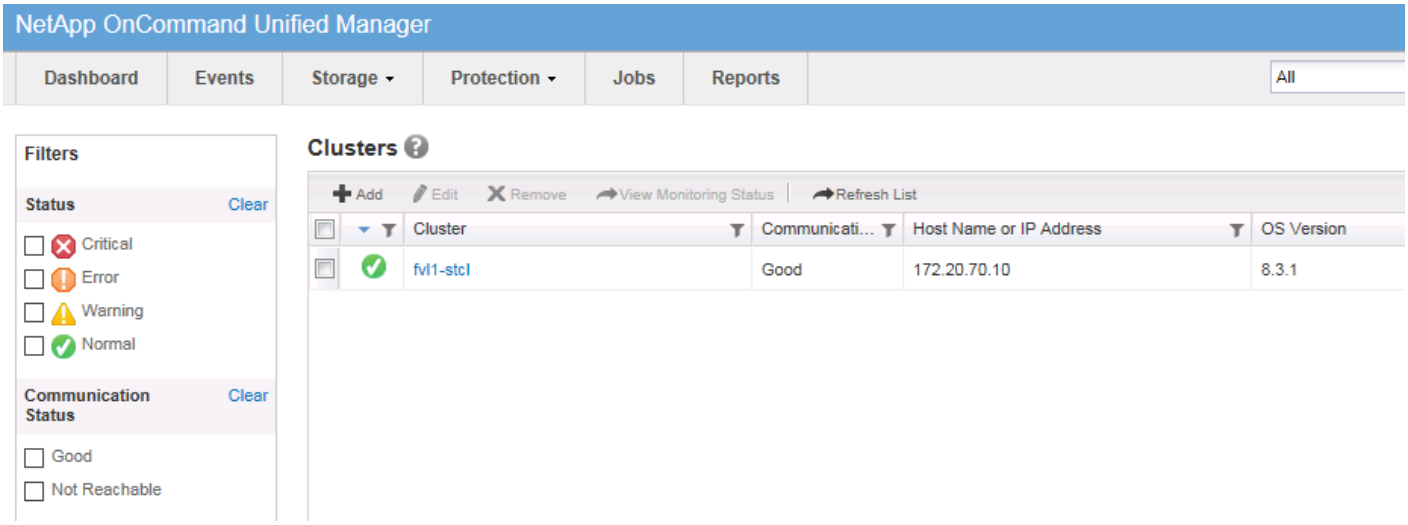

OnCommand Performance Manager 2.0

OnCommand Performance Manager OVF Deployment

To install OnCommand Performance Manager, complete the following steps:

1. Download and review the OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances.
2. Download OnCommand Performance Manager version 2.0 (`OnCommandPerformanceManager-netapp-2.0.0.ova`).
3. Log in to the vSphere Web Client. Select Home > VMs and Templates.
4. At the top of the center pane, click Actions > Deploy OVF Template.

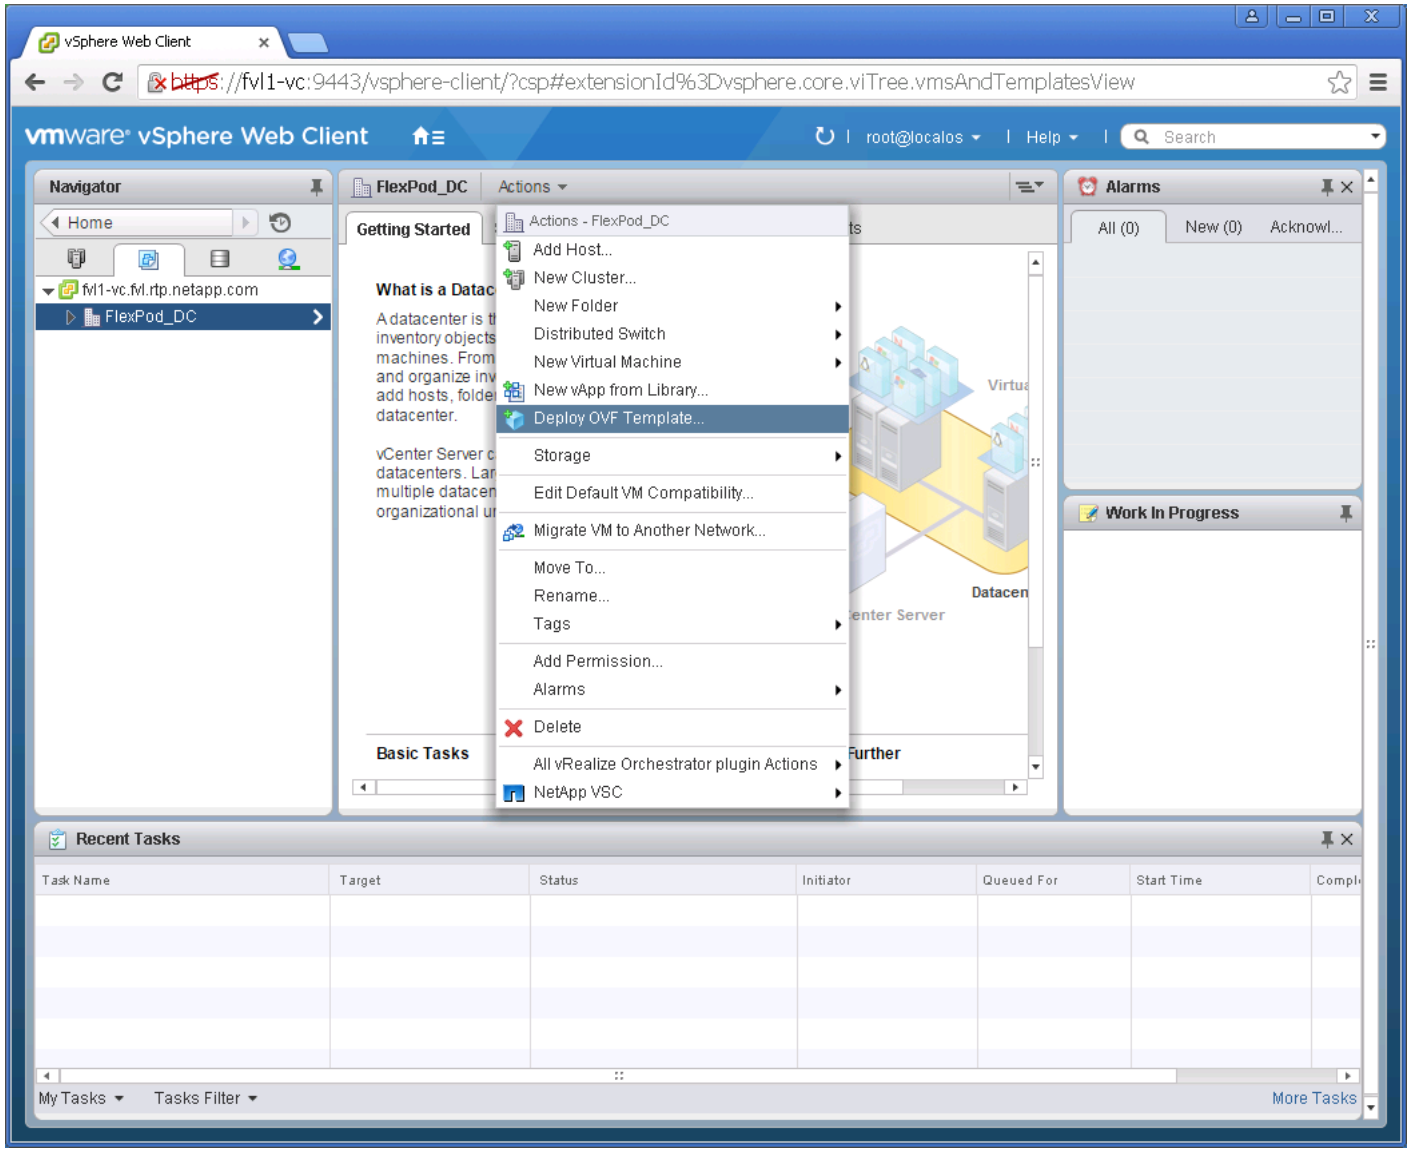

5. Browse to the `OnCommandPerformanceManager-netapp-2.0.0.ova` file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.

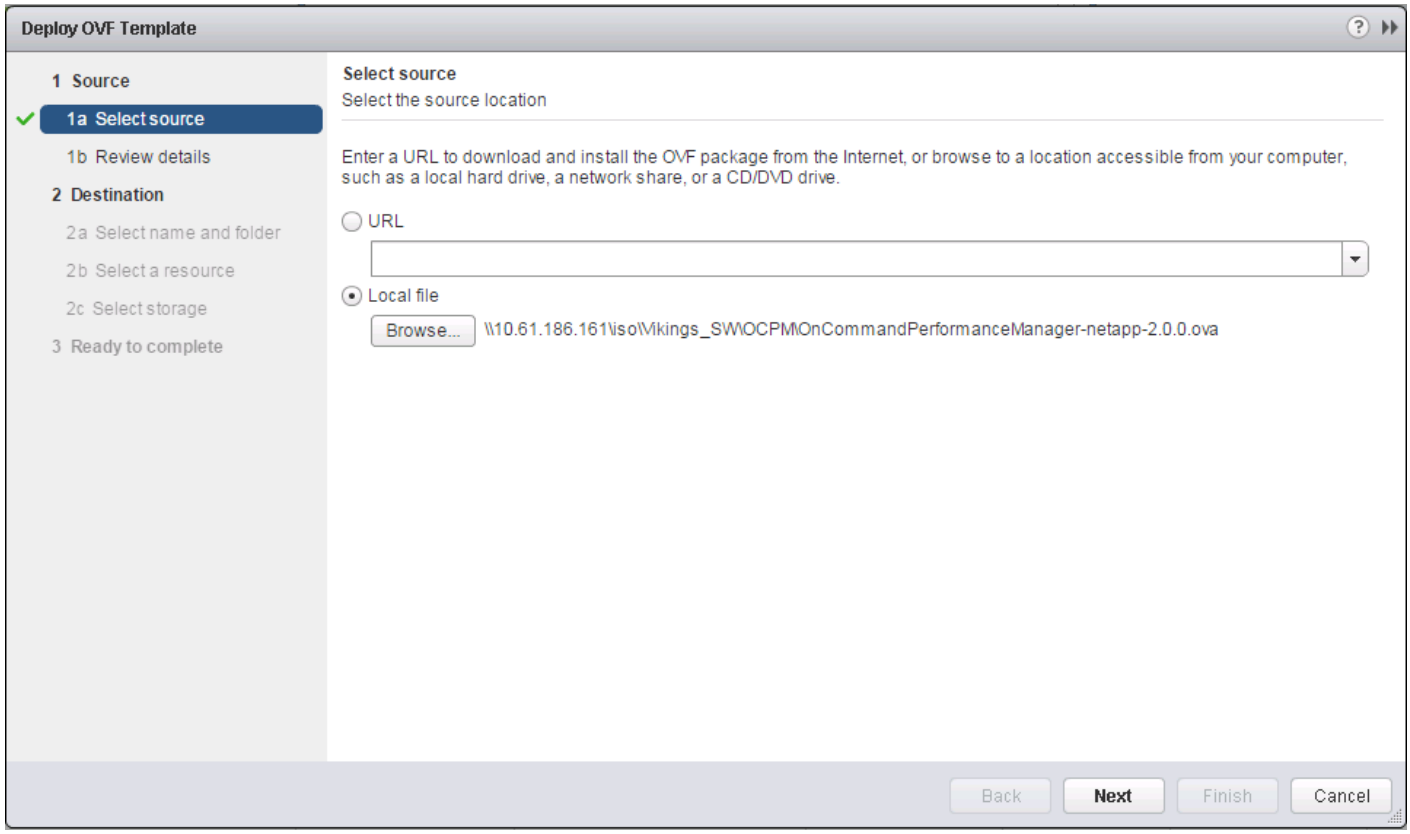

6. Review the details and click Next.

7. Read the EULA and click the Accept button to accept the agreement. Click Next.

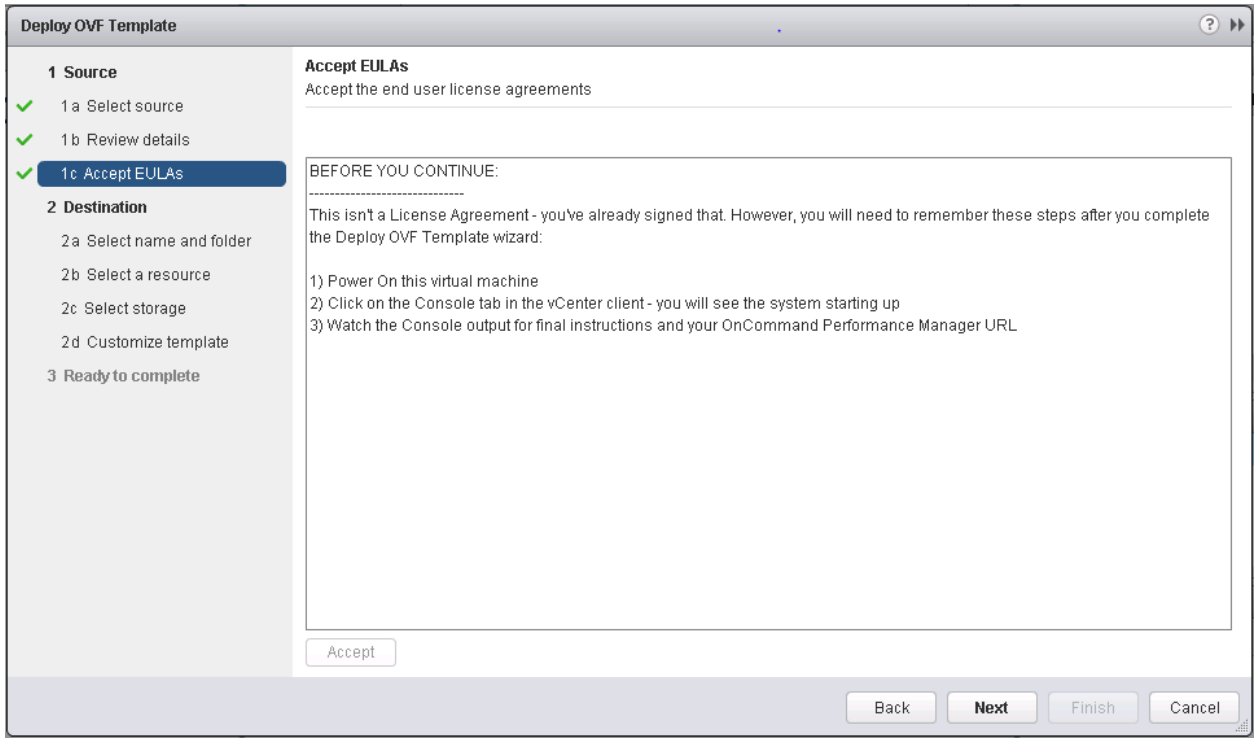

8. Enter the name of the VM and select the `FlexPod_DC` folder to hold the VM. Click Next.

9. Select `FlexPod_Management` within the `FlexPod_DC` datacenter as the destination compute resource pool to host the VM. Click Next.

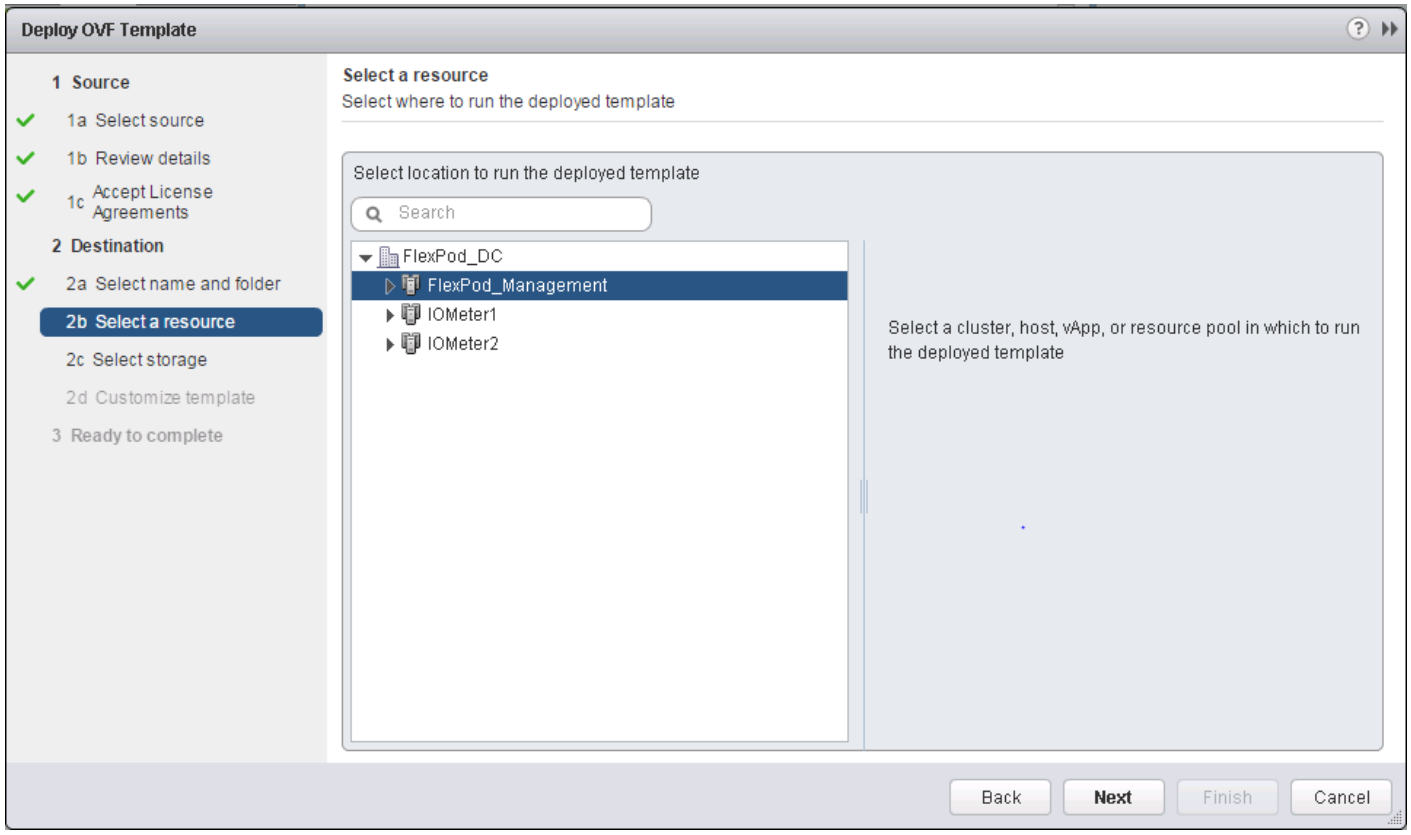

10. Select `infra_datastore_1` as the storage target for the VM, and select Thin Provision as the virtual disk format. Click Next.

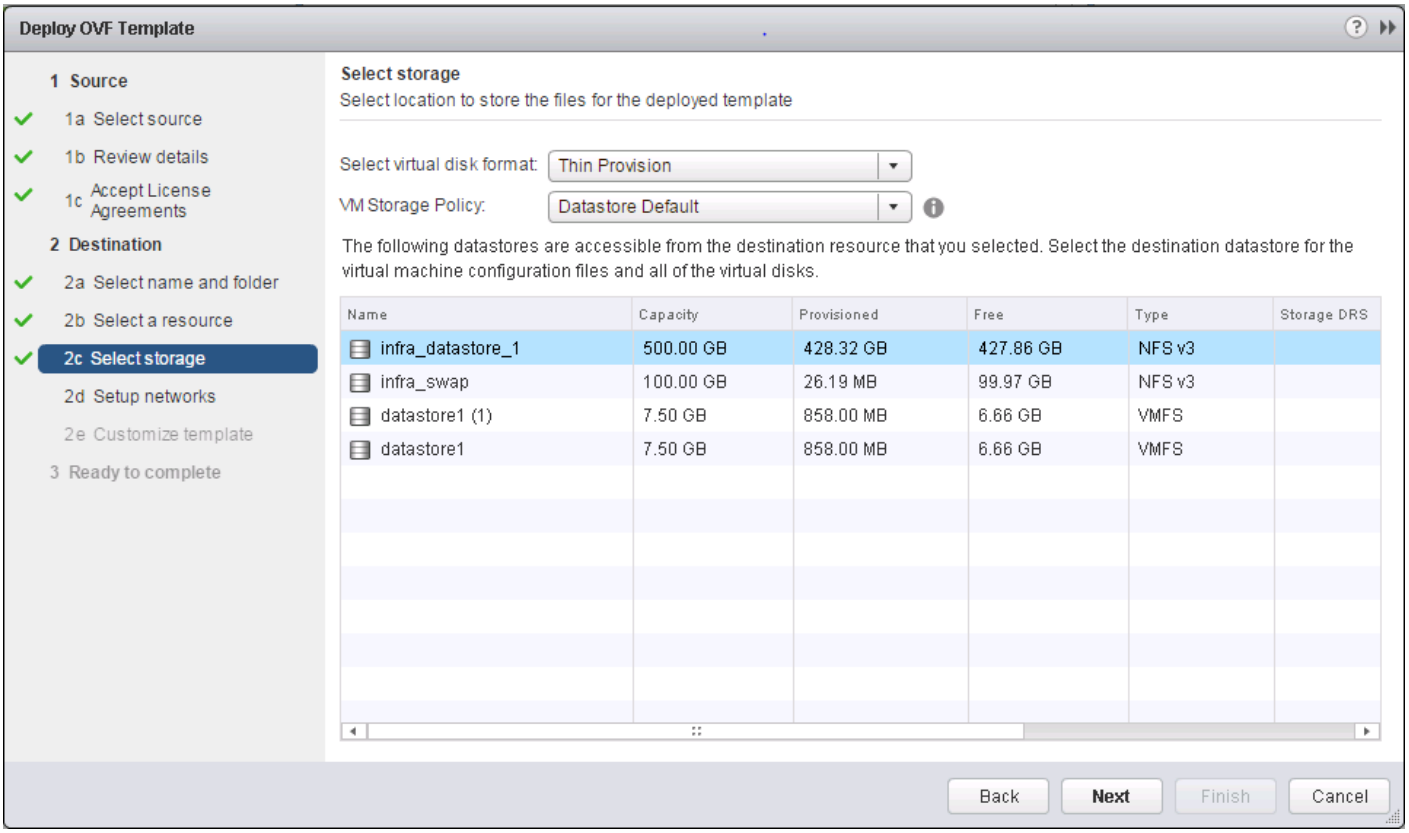

11. Select `IB-MGMT-VLAN` as the destination network for the `nat` source network. Click Next.

12. Enter the host name, IP address, network mask, gateway, primary DNS, and secondary DNS. Click Next.

13. Deselect Power On After Deployment.

14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.

15. In the left pane, navigate to Home -> Hosts and Clusters. Expand the `FlexPod_Management` cluster and select the newly created OnCommand Performance Manager VM. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.

16. Expand the CPU options.

    a. The minimum required CPU reservation is 9572MHz. Determine the CPU frequency of the host.

    b. Set the required number of CPUs (9572 / CPU frequency of the host).

    c. Set the number of cores per socket where the socket number on the right matches the number of CPU sockets in the host. For example, if a host has two CPUs operating at a speed of 1999MHz, then the VM requires six virtual CPUs (9572 / 1999 = 4.79 – rounded to 6 virtual CPUs). If the host has two physical CPU sockets, then set three cores per socket.

For detailed information, see OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances.

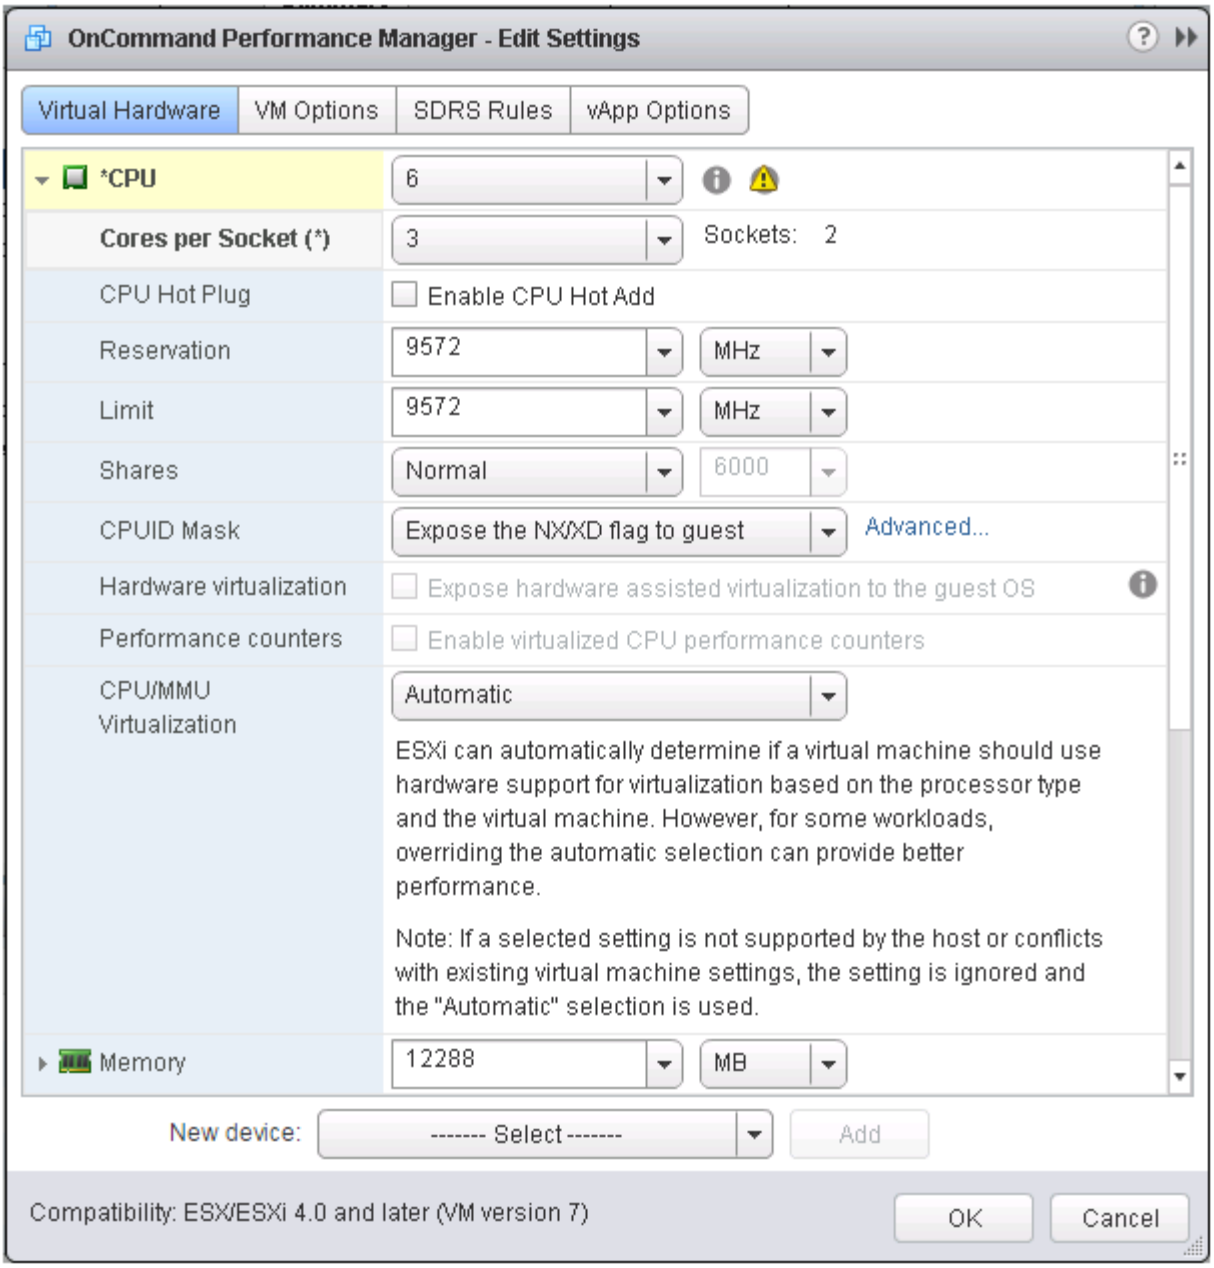

17. Click OK to accept the changes.

18. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Performance Manager Basic Setup

To setup the OnCommand Performance Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.

2. In the VMware Remote Console window, select VMRC > Manage > Install VMware Tools. VMware Tools installs in the VM.

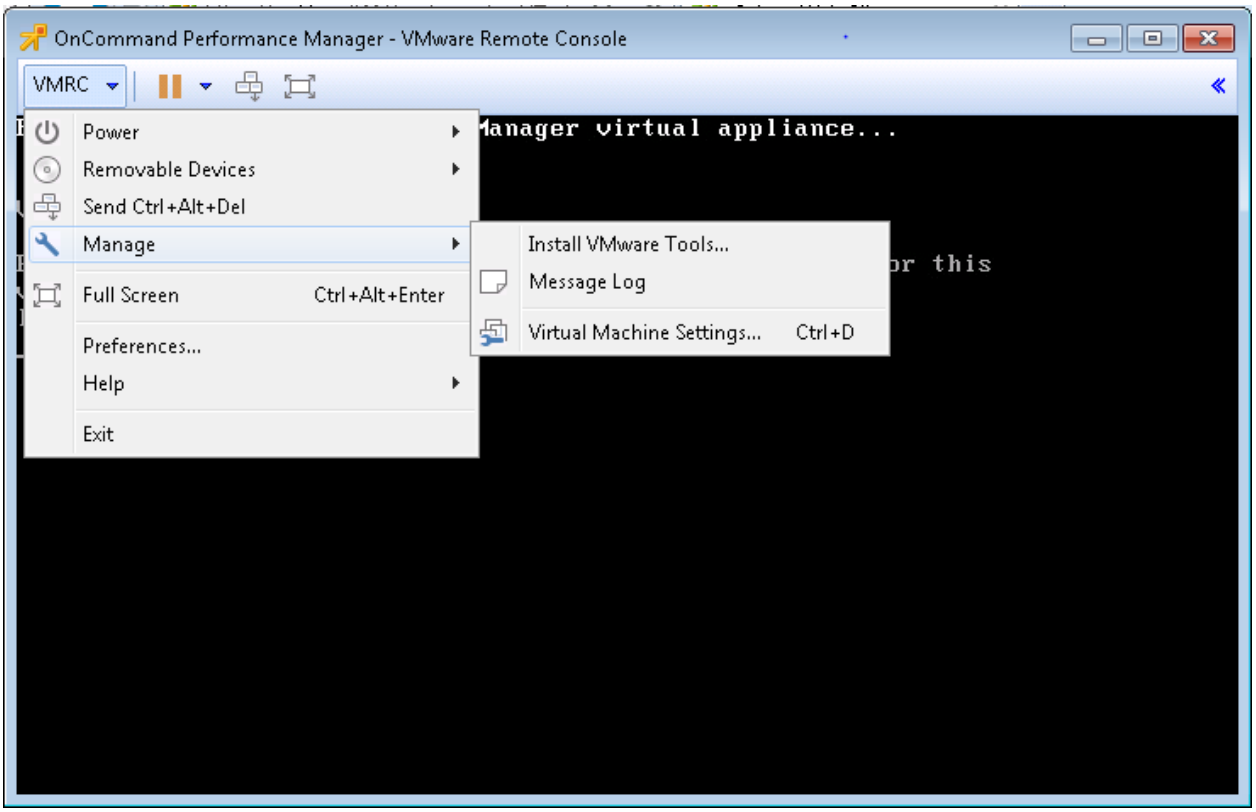

3. Set up OnCommand Performance Manager by answering the following questions in the console window:

```
Geographic area: <<Enter your geographic location>>
Time zone: <<Select the city or region corresponding to your time zone>>
```

These commands complete the network configuration checks, generate SSL certificates, and start OnCommand Performance Manager services.

4. To create a maintenance user account, run the following commands:

The maintenance user manages and maintains the settings on the OnCommand Performance Manager virtual appliance.

```
Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>
```

5. Using a web browser, navigate to OnCommand Performance Manager using the URL `https://<<var_oncommand_pm_ip>>`.

6. Log in using the maintenance user account (admin) credentials.

7. Enter a maintenance user e-mail address, SMTP mail server information, and the NTP server IP address. Click Save.

8. Select the `Yes` option to enable AutoSupport capabilities. Click Save.

9. Click Save and go to next step to not change the admin password.

10. Enter the storage cluster host name or IP address, the storage cluster admin user name, and the storage cluster admin password. Click Add Cluster, and then click Save and Complete Configuration. It can take up to 15 minutes for the cluster to be visible in OnCommand Performance Manager.

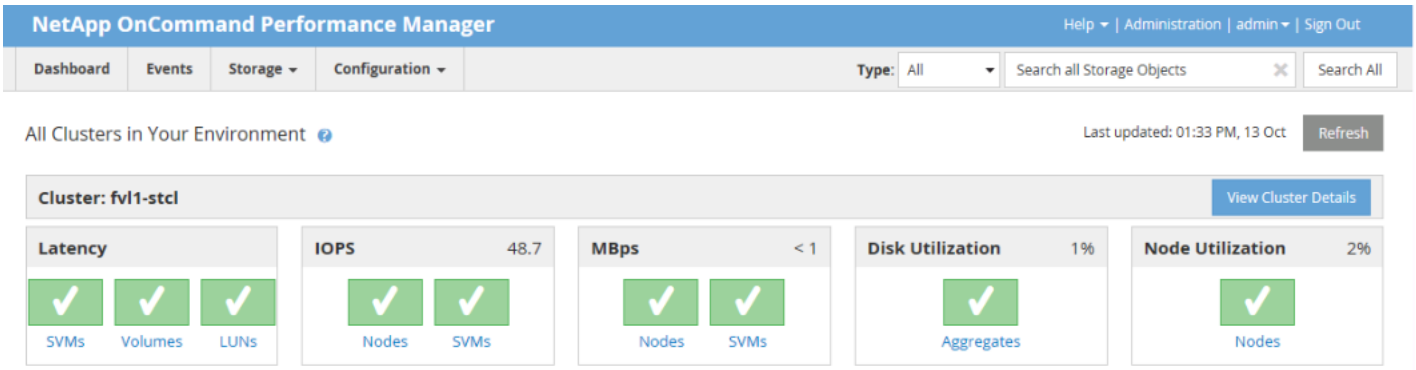

11. After the cluster is added, it can be accessed by clicking on Administration > Manage Data Sources.

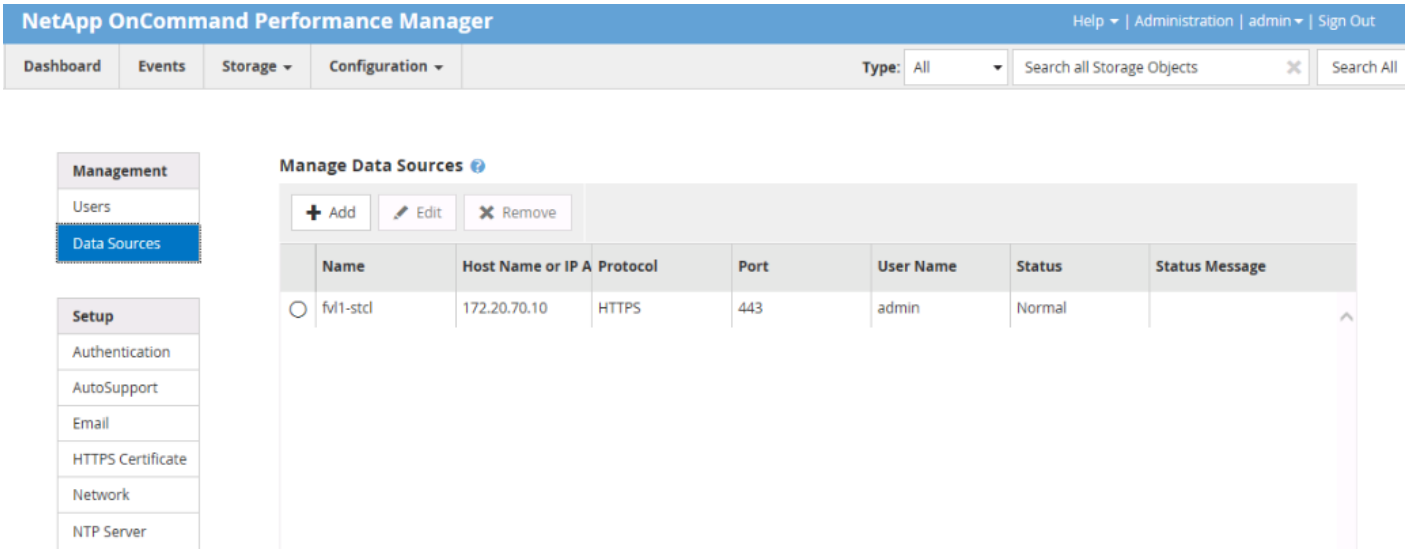

Link OnCommand Performance Manager to OnCommand Unified Manager

To link OnCommand Performance Manager to the OnCommand Unified Manager, complete the following steps:

1. Using a web browser, navigate to OnCommand Unified Manager using the URL `https://<<var_oncommand_server_ip>>`. Log in with the maintenance user ID and password set up earlier.

2. In the OnCommand Unified Manager web interface, select Administration > Manage Users to set up an Event Publication user.

3. Click Add to add a user.

4. Leave the Type set to Local User. Use eventpub as the name and enter and confirm a password. Enter an e-mail address for this user and set the Role to Event Publisher. Click Add.

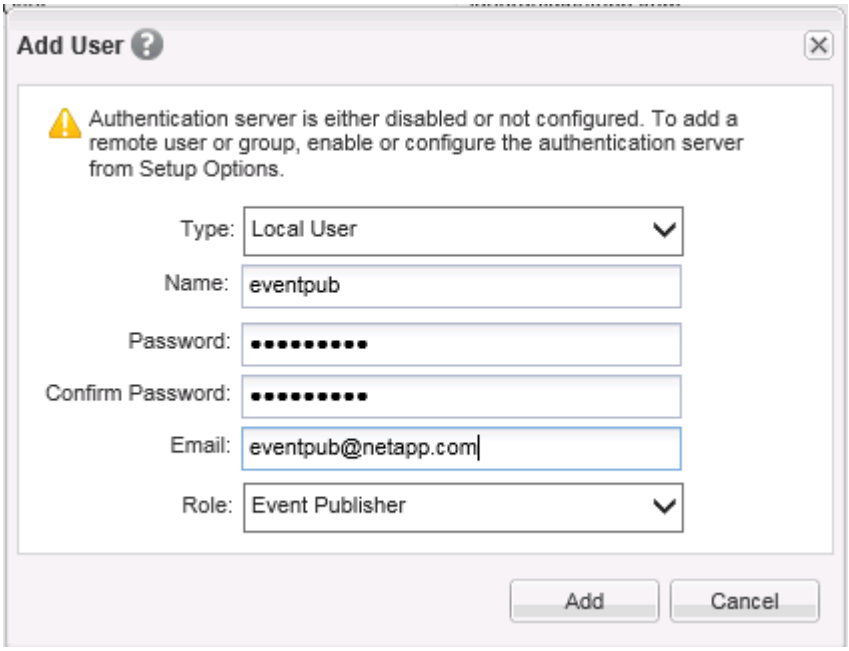

5. In the OnCommand Performance Manager console window, log into the Command Line Interface with the maintenance user (admin) defined earlier.

6. Enter 5 to select Unified Manager Connection.

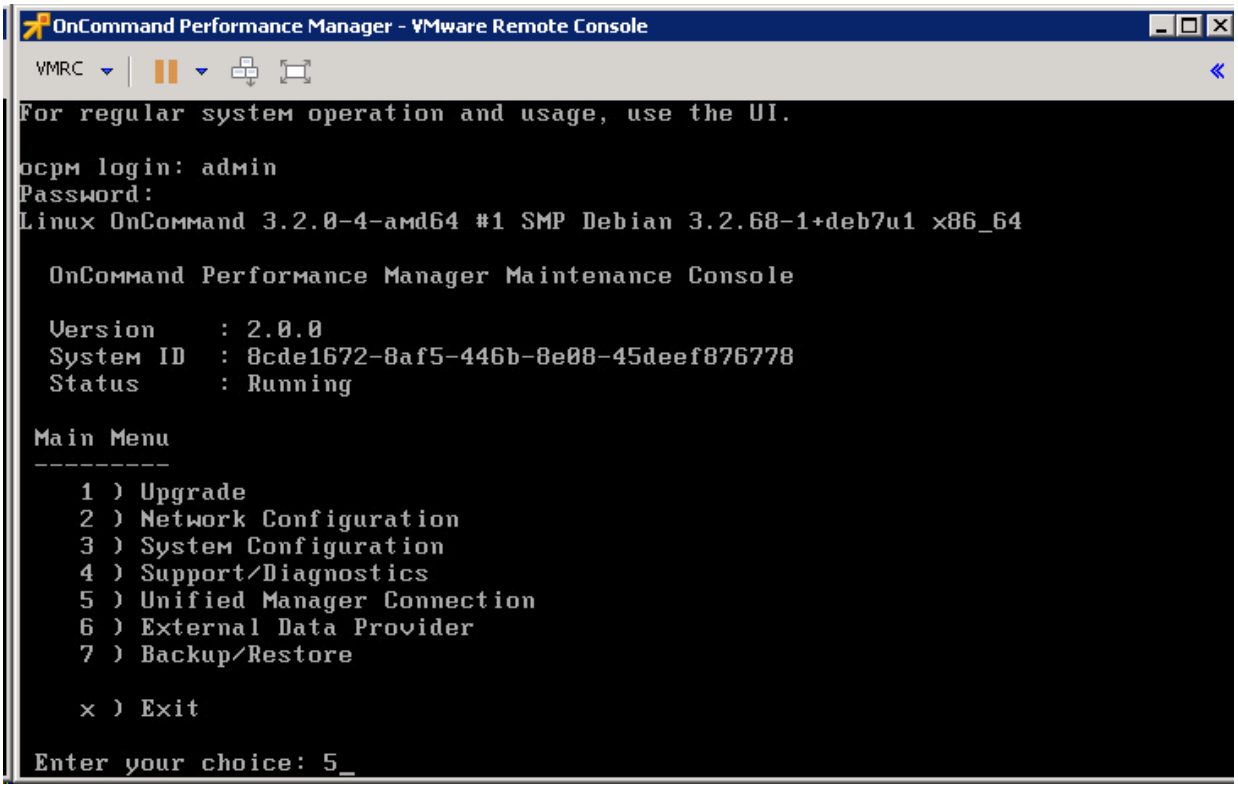

7. Enter `2` to Add / Modify Unified Manager Server Connection.
8. Enter `y` to continue.
9. Enter the OnCommand Unified Manager FQDN or IP address.
10. Click Enter to accept the default port 443.
11. Enter `y` to accept the Unified Manager security certificate.
12. Enter `eventpub` for the Event Publisher User Name.
13. Enter the eventpub password.
14. Enter `y` to accept the entered settings.
15. Press any key to continue.
16. Exit the OnCommand Performance Manager console. OnCommand Performance Manager events now appear in the OnCommand Unified Manager Dashboard.

NetApp NFS Plug-In 1.1.0 for VMware VAAI

Enable VMware vStorage for NFS in Clustered Data ONTAP

To enable VMware vStorage for NFS in ONTAP, complete the following steps:

1. From an SSH session to the storage cluster management address, log in with the admin user name and password.

2. Enable vStorage on the SVM.

```
vserver nfs modify –vserver Infra-SVM –vstorage enabled
```

3. Verify that the export policy rules are set up correctly.

```
vserver export-policy rule show –vserver Infra-SVM
```

The following text provides a sample output:

```
NetApp::> vserver export-policy rule show –vserver Infra-SVM
             Policy          Rule    Access   Client                RO
Vserver      Name            Index   Protocol Match                 Rule
------------ --------------- ------  -------- --------------------- ---------
Infra-SVM    default         1       nfs      192.168.170.61       sys
Infra-SVM    default         2       nfs      192.168.170.60        sys
Infra-SVM    default         3       nfs      192.168.170.58        sys
Infra-SVM    default         4       nfs      192.168.170.59        sys
Infra-SVM    default         5       nfs      192.168.170.62        sys
Infra-SVM    default         6       nfs      192.168.170.63        sys
6 entries were displayed.
```

4. The access protocol for the FlexPod policy name should be `nfs`. If the access protocol is not `nfs` for a given rule index, run the following command to set `nfs` as the access protocol:

```
vserver export-policy rule modify –vserver Infra-SVM –policyname default –ruleindex <<var_rule_index>> -
protocol nfs
```

Install NetApp NFS Plug-In for VMware VAAI

To install the NetApp NFS plug-in for VMware vStorage APIs for Array Integration (VAAI), complete the following steps:

1. From a console interface on the NetApp VSC VM, go to the Software Downloads page in the NetApp Support site.

2. Scroll down to locate the NetApp NFS Plug-in for VMware VAAI, select the ESXi6.0 platform, and click Go.

3. Click View & Download.

4. Click CONTINUE.

5. Click Accept.

6. Download the `.vib` file of the most recent plug-in version to the VSC VM Desktop as `NetAppNasPlugin.vib`.

It is important that the file be saved as `NetAppNasPlugin.vib`.

7. On the VSC VM desktop, move the `NetAppNasPlugin.vib` file to the `C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web` folder.

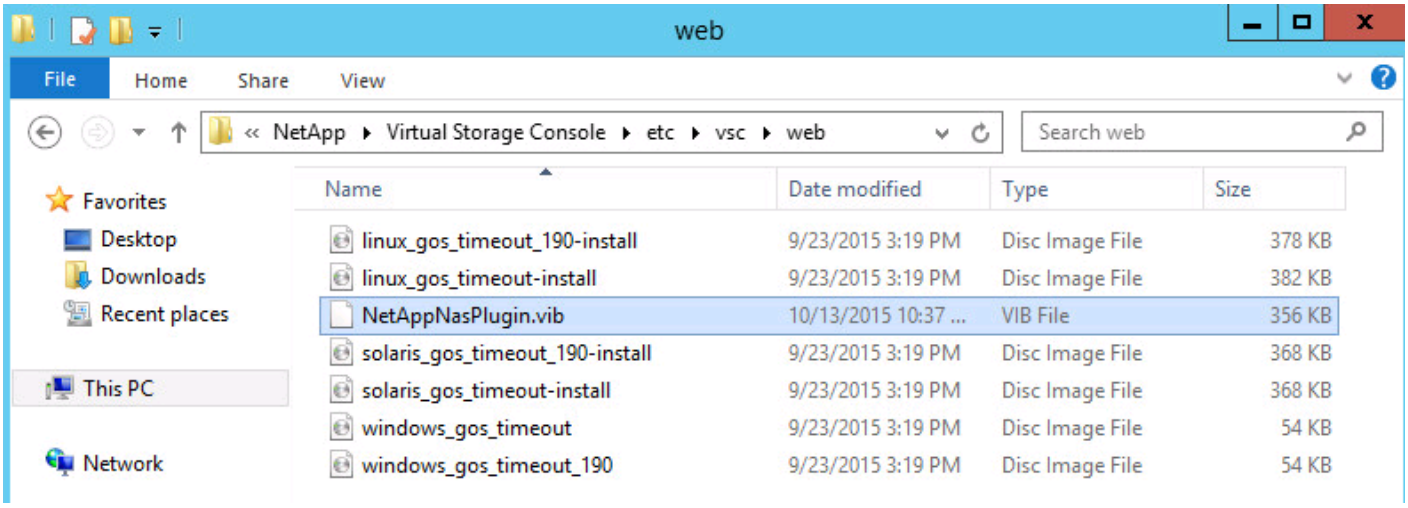

8. Go to the VMware vSphere Web Client and select VSC. Click NFS VAAI Tools. Make sure NFS Plug-in for VMware VAAI Version: 1.1.0-0 is shown.

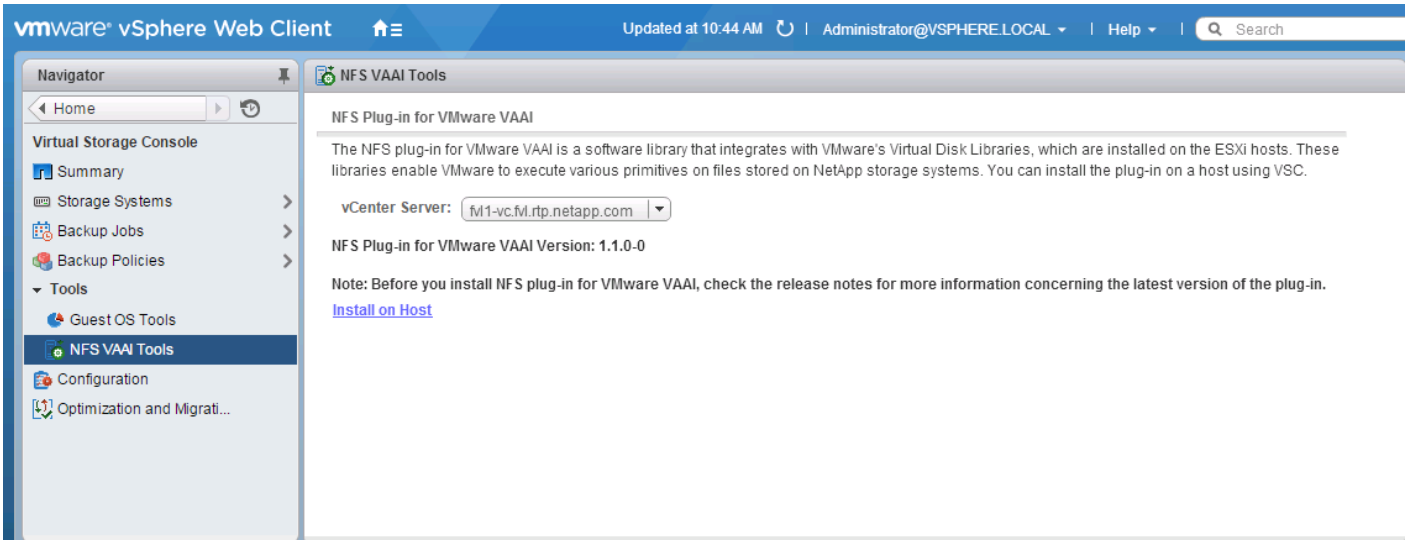

9. Click Install on Host. Select all of the hosts on which you want to install the plug-in.

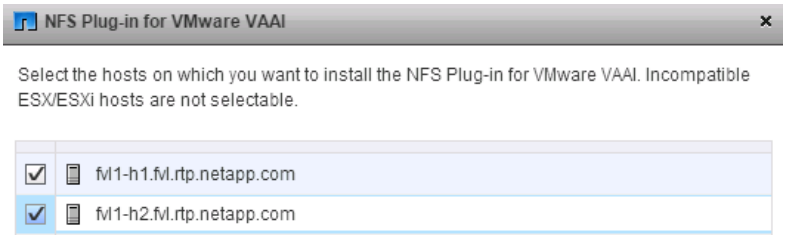

10. Click Install and then click OK.

11. One at a time, put each ESXi host into maintenance mode, reboot the host, and then exit maintenance mode. It might be necessary to manually migrate VMs to the other host to allow the host to enter maintenance mode.

12. When the reboots have completed, click Storage in the vSphere web client from the Home page. Then select the `infra_datastore_1` datastore. Select Settings under the Manage tab in the center pane. Hardware Acceleration should now read Supported on All Hosts, as is shown in the following screenshot. All NFS datastores should now support hardware acceleration.

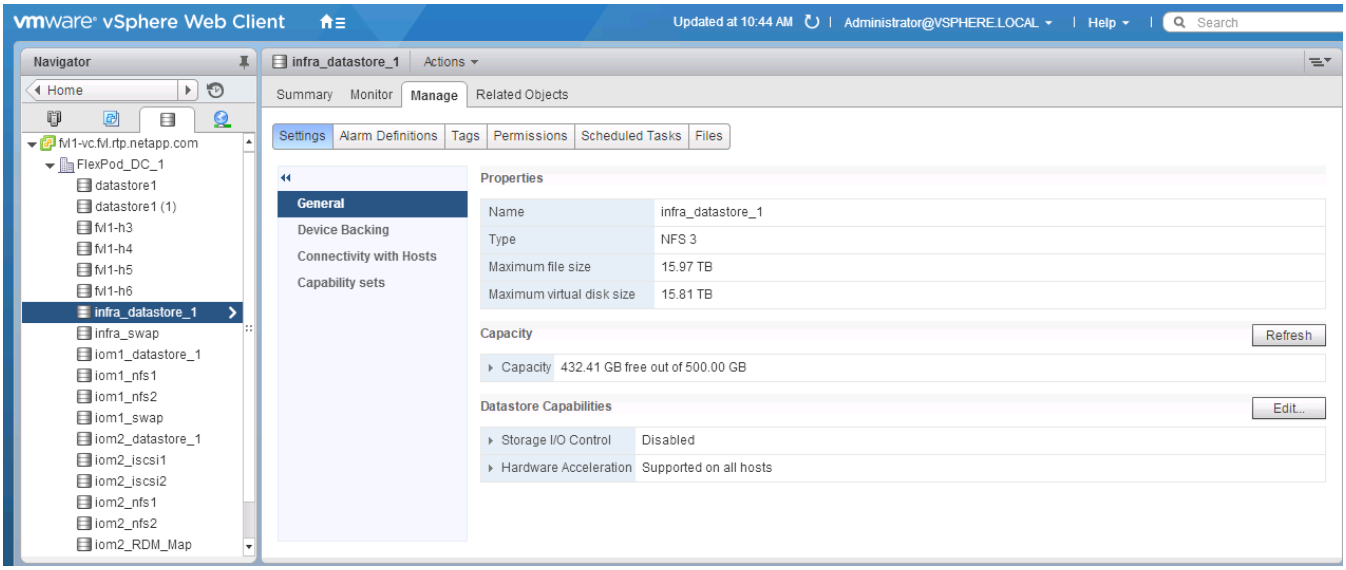

About the Authors

Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.

Ramesh Isaac is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Group. Ramesh has worked in data center and mixed-use lab settings since 1995. He started in information technology supporting UNIX environments and focused on designing and implementing multi-tenant virtualization solutions in Cisco labs over the last couple of years. Ramesh holds certifications from Cisco, VMware, and Red Hat.

Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Masters of Science in Information Security from East Carolina University.

Dave Derry, Technical Marketing Engineer, Infrastructure and Cloud Engineering, NetApp

Dave Derry is a Technical Marketing Engineer in the Converged Infrastructure Engineering team at NetApp. He has been with NetApp since 2012, serving in a variety of engineering roles. Prior to that, he was an engineer at Cisco Systems for over ten years, in a variety of development and solution test roles.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- John George, Cisco Systems, Inc.
- Haseeb Niazi, Cisco Systems, Inc.
- Melissa Palmer, NetApp