FlexPod Data Center with VMware vSphere 5.1Update1 with 7-Mode

Deployment Guide for FlexPod with VMware vSphere 5.1Update1 with Data ONTAP 8.2 operating in 7-Mode

Last Updated: May 15, 2014

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

John Kennedy, Technical Leader, Server Access Virtualization Business Unit, Cisco Systems

John Kennedy is a technical marketing engineer in the Server Access and Virtualization Technology group. Currently, John is focused on the validation of FlexPod architecture while contributing to future SAVTG products. John spent two years in the Systems Development unit at Cisco, researching methods of implementing long-distance vMotion for use in the Data Center Interconnect Cisco Validated Designs. Previously, John worked at VMware for eight and a half years as a senior systems engineer supporting channel partners outside the United States and serving on the HP Alliance team. He is a VMware Certified Professional on every version of VMware ESX and ESXi, vCenter, and Virtual Infrastructure, including vSphere 5. He has presented at various industry conferences.

Chris O'Brien, Technical Marketing Manager, Server Access Virtualization Business Unit, Cisco Systems

Chris O'Brien is currently focused on developing infrastructure best practices and solutions that are designed, tested, and documented to facilitate and improve customer deployments. Previously, O'Brien was an application developer and has worked in the IT industry for more than 15 years.

Arvind Ramakrishnan, Technical Marketing Engineer, Infrastructure and Cloud Engineering, NetApp Systems

Arvind Ramakrishnan is a Technical Marketing Engineer in the NetApp Infrastructure and Cloud Engineering team and is focused on developing, validating, and supporting converged infrastructure solutions that include NetApp products. Before his current role, he was a software engineer at EMC developing applications for cloud infrastructure management.

Karthick Radhakrishnan, Systems Architect, Infrastructure and Cloud Engineering, NetApp Systems

Karthick Radhakrishnan is a Systems Architect in the NetApp Infrastructure and Cloud Engineering team. He focuses on validating, supporting, and implementing cloud infrastructure solutions that include NetApp products. Prior to his current role, he was a networking tools developer at America Online supporting AOL transit data network. Karthick started his career in 2003, and he holds a master's degree in Computer Application.

Lindsey Street, Systems Architect, Infrastructure and Cloud Engineering, NetApp Systems

Lindsey Street is a systems 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 Master's of Science in Information Security from East Carolina University.

John George, Reference Architect, Infrastructure and Cloud Engineering, NetApp Systems

John George is a Reference Architect in the NetApp Infrastructure and Cloud Engineering team and is focused on developing, validating, and supporting cloud infrastructure solutions that include NetApp products. Before his current role, he supported and administered Nortel's worldwide training network and VPN infrastructure. John holds a Master's degree in computer engineering from Clemson University.

Chris Reno, Reference Architect, Infrastructure and Cloud Engineering, NetApp Systems

Chris Reno is a reference architect in the NetApp Infrastructure and Cloud Enablement group and is focused on creating, validating, supporting, and evangelizing solutions based on NetApp products. Before being employed in his current role, he worked with NetApp product engineers designing and developing innovative ways to perform Q&A for NetApp products, including enablement of a large grid infrastructure using physical and virtualized compute resources. In these roles, Chris gained expertise in stateless computing, netboot architectures, and virtualization.

About Cisco Validated Design (CVD) Program

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

FlexPod Data Center with VMware vSphere 5.1Update1 with 7-Mode

Overview

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with prevalidated 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

This document describes the architecture and deployment procedures of an infrastructure composed of Cisco®, NetApp®, and VMware® virtualization that uses FCoE-based storage serving NAS and SAN protocols. The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy the core FlexPod architecture with NetApp Data ONTAP® operating in 7-mode.

Architecture

The FlexPod architecture is highly modular or "podlike." Although each customer's FlexPod unit varies in its exact configuration, after a FlexPod unit is built, it 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).

Specifically, FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and nonvirtualized solutions. VMware vSphere® built on FlexPod includes NetApp storage, NetApp Data ONTAP, Cisco networking, the Cisco Unified Computing System™ (Cisco UCS®), and VMware vSphere® software in a single package. The design is flexible enough that the

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

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. This is why the reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an FCoE-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 FCoE-based storage. This design uses the Cisco Nexus® 5548UP, Cisco Nexus 2232PP FEX, and Cisco UCS C-Series and B-Series with the Cisco UCS virtual interface card (VIC) and the NetApp FAS family of storage controllers connected in a highly available design using Cisco Virtual PortChannels (vPCs). This infrastructure is deployed to provide FCoE-booted hosts with file- and block-level access to shared storage datastores. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture; be it FC, FCoE, or 10 Gigabit Ethernet, no recabling is required from the hosts to the Cisco UCS fabric interconnect.

Figure 1 VMware vSphere Built on FlexPod Components

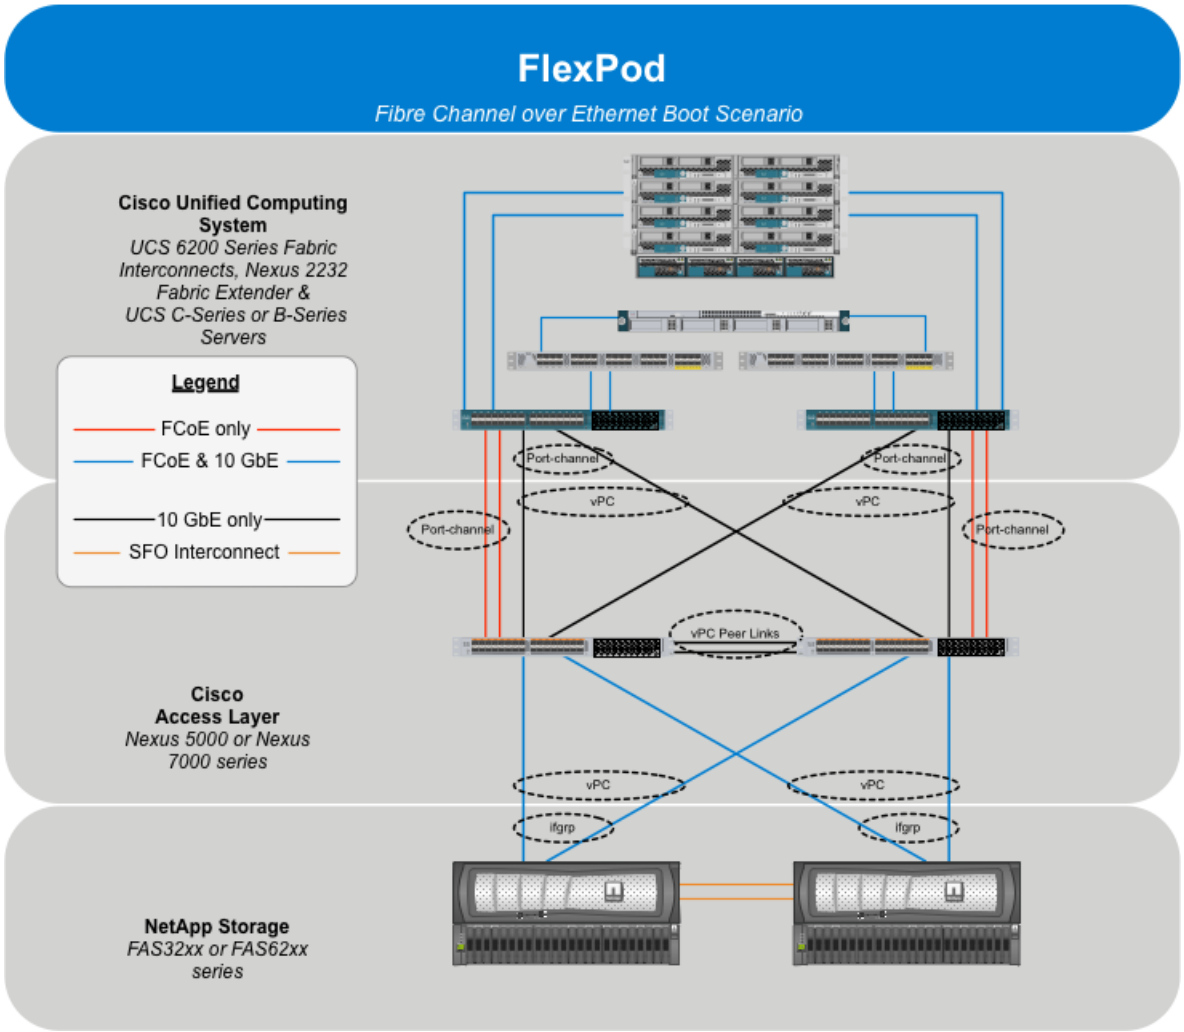

The reference configuration includes:

- Two Cisco Nexus 5548UP switches
- Two Cisco Nexus 2232PP fabric extenders
- Two Cisco UCS 6248UP fabric interconnects
- Support for 16 Cisco UCS C-Series servers without any additional networking components
- Support for 8 Cisco UCS B-Series servers without any additional blade server chassis
- Support for hundreds of Cisco UCS C-Series and B-Series servers by way of additional fabric extenders and blade server chassis
- One NetApp FAS3250-AE (HA pair) operating in 7-mode

Storage is provided by a NetApp FAS3250-AE (HA configuration in two chassis) operating 7-Mode. All system and network links feature redundancy, providing end-to-end high availability (HA). For server virtualization, the deployment includes VMware vSphere. Although this is the base design, each of the components can be scaled flexibly 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 capacity and throughput, and special hardware or software features can be added to introduce new capabilities.

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 compute and storage configuration to configuring virtualization with VMware vSphere.

Software Revisions

It is important to note the software versions used in this document. Table 1 details the software revisions used throughout this document.

Table 1 Software Revisions

Layer	Compute	Version or Release	Details
Compute	Cisco UCS Fabric Interconnect	2.1(3)	Embedded management
	Cisco UCS C 220 M3 Server	2.1(3)	Software bundle release
	Cisco UCS B 200 M3 Server	2.1(3)	Software bundle release
	Cisco eNIC	2.1.2.38	Ethernet driver for Cisco VIC
	Cisco fNIC	1.5.0.45	FCoE driver for Cisco VIC
Network	Cisco Nexus Switch	6.0(2)N1(2a)	Operating system version
Storage	NetApp FAS3250-AE	Data ONTAP 8.2P4 operating in 7-mode	Operating system version

Table 1 Software Revisions

Layer	Compute	Version or Release	Details
Software	Cisco UCS Hosts	VMware vSphere ESXi™ 5.1Update1	Operating system version
	Microsoft® .NET Framework	3.5.1	Feature enabled within Windows® operating system
	Microsoft SQL Server®	Microsoft SQL Server 2008 R2 SP1	VM (1 each): SQL Server DB
	VMware vCenter™	5.1Update1	VM (1 each): VMware vCenter
	NetApp OnCommand®	5.2	VM (1 each): OnCommand
	NetApp Virtual Storage Console (VSC)	4.2.1	Plug-in within VMware vCenter
	Cisco Nexus 1110-x	4.2(1)SP1(6.2)	Virtual Services Appliance (VSA)
	Cisco Nexus 1000v	4.2(1)SV2(2.1a) (Advanced Edition)	Virtual services blade within the VSA
	NetApp NFS Plug-in for VMware vStorage APIs for Array Integration (VAAI)	1.0.20	Plug-in within VMware vCenter
	Cisco UCS Central	1.1	Manager of multiple UCS domains
	Flash Accel™ for VMware Virtual Storage Console	1.2R1	Software that manages server flash storage

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. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the procedure for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, and so on. 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:

```
    controller A> vlan create

Usage:
        vlan create [-g {on|off}] <ifname> <vlanid_list>
        vlan add <ifname> <vlanid_list>
        vlan delete -q <ifname> [<vlanid_list>]
        vlan modify -g {on|off} <ifname>
        vlan stat <ifname> [<vlanid_list>]
Example:
```

```
controller A> vlan create vif0 <management VLAN ID>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 2 describes the VLANs necessary for deployment as outlined in this guide. The VM-Mgmt VLAN is used for management interfaces of the VMware vSphere hosts. Table 3 lists the VSANs necessary for deployment as outlined in this guide. Table 5 lists the configuration variables that are used throughout this document. Table 5 can be completed based on the specific site variables and used in implementing the document configuration procedures.

If you use separate in-band and out-of-band management VLANs, you must create a Layer 3 route between these VLANs. For this validation, a common management VLAN was used.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Mgmt in band	VLAN for in-band management interfaces	3175
Mgmt out of band	VLAN for out-of-band management interfaces	3170
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for NFS traffic	3172
FCoE - A	VLAN for FCoE traffic for fabric A	101
FCoE - B	VLAN for FCoE traffic for fabric B	102
vMotion	VLAN designated for the movement of VMs from one physical host to another	3173
VM Traffic	VLAN for VM application traffic	3174
Packet Control	VLAN for Packet Control traffic	3176

Table 3 Necessary VSANs

VSAN Name	VSAN Purpose	ID Used in Validating This Document
VSAN A	VSAN for fabric A traffic. ID matches FCoE-A VLAN	101
VSAN B	VSAN for fabric B traffic. ID matches FCoE-B VLAN	102

Table 4 Created VMware Virtual Machine

Virtual Machine Description	Host Name
vCenter SQL Server database	
vCenter Server	
NetApp Virtual Storage Console (VSC) and NetApp OnCommand® Unified Manager core	
Cisco UCS Central	
Active Directory (If not present)	

Table 5 Configuration Variables

Variable	Description	Customer Implementation Value
<<var_controller1>>	Storage Controller 1 Host Name	
<<var_controller1_e0m_ip>>	Out-of-band management IP for Storage Controller 1	
<<var_controller1_mask>>	Out-of-band management network netmask	
<<var_controller1_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_adminhost_ip>>>	Administration Host Server IP	
<<var_timezone>>	FlexPod time zone (for example, America/New_York)	
<<var_location>>	Node location string	
<<var_dns_domain_name>>	DNS domain name	
<<var_nameserver_ip>>	DNS server IP(s)	
<<var_controller1_sp_ip>>	Out-of-band service processor management IP for storage controller1	
<<var_controller1_sp_mask>>	Out-of-band management network netmask for controller1	
<<var_controller1_sp_gateway>>	Out-of-band management network default gateway for controller1	
<<var_controller2_sp_ip>>	Out-of-band service processor management IP for storage controller2	
<<var_controller2_sp_mask>>	Out-of-band management network netmask for controller2	
<<var_controller2_sp_gateway>>	Out-of-band management network default gateway for controller2	
<<var_mailhost>>	Mail server host name	

Table 5 Configuration Variables

Variable	Description	Customer Implementation Value
<<var_mailhost_ip>>	Mail server IP	
<<var_password>>	Global default administrative password	
<<var_controller2>>	Storage Controller 2 Host Name	
<<var_controller2_e0m_ip>>	Out-of-band management IP for Storage Controller 2	
<<var_controller2_mask>>	Out-of-band management network netmask	
<<var_controller2_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_#_of_disks>>	Number of disks to assign to each storage controller	
<<var_nfs_ip>>	NFS VLAN IP for each storage controller	
<<var_nfs_mask>>	NFS VLAN netmask	
<<var_nfs_license>>	Data ONTAP NFS License Code	
<<var_fc_license>>	Data ONTAP Fiber Channel Protocol License Code	
<<var_date>>	Current time in [[[[CC]yy]mm]dd]hhmm[.ss]] format	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_state>>	State or province name	
<<var_city>>	City name	
<<var_org>>	Organization or company name	
<<var_unit>>	Organizational unit name	
<<var_controller1_fqdn>>	Storage Controller 1 Fully Qualified Domain Name (FQDN)	
<<var_admin_email>>	Administrator e-mail address	
<<var_key_length>>	Number of bits in SSL/SSH Security Key	
<<var_controller2_fqdn>>	Storage Controller 2 FQDN	
<<var_snaprestore_license>>	Data ONTAP SnapRestore License Code	
<<var_flex_clone_license>>	Data ONTAP FlexClone License Code	
<<var_url_boot_software>>	Data ONTAP 8.2P4 URL; format: http://	
<<var_oncommand_server_fqdn>>	OnCommand/VSC Virtual Machine FQDN	
<<var_snmp_community>>	SNMP v1/v2 community name	
<<var_country_code>>	Two-letter country code	
<<var_nexus_A_hostname>>	Cisco Nexus A host name	

Table 5 Configuration Variables

Variable	Description	Customer Implementation Value
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_controller2_nfs_ip>>	NFS VLAN IP for storage controller 2	
<<var_nexus_B_hostname>>	Cisco Nexus B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_ib-mgmt_vlan_id>>	In-band management network VLAN ID	
<<var_native_vlan_id>>	Native VLAN ID	
<<var_oob-mgmt_vlan_id>>	Out of band management network VLAN ID	
<<var_nfs_vlan_id>>	NFS VLAN ID	
<<var_pkt-ctrl_vlan_id>>	Cisco Nexus 1000v packet control VLAN ID	
<<var_vmotion_vlan_id>>	VMware vMotion® VLAN ID	
<<var_vm-traffic_vlan_id>>	VM traffic VLAN ID	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_controller1_nfs_ip>	NFS VLAN IP for storage controller 1	
<<var_ucs_central_ip>>	UCS Central management IP address	
<<var_ucs_central_netmask>>	Out-of-band management network netmask	
<<var_ucs_central_gateway>>	Out-of-band management network default gateway	
<<var_ucs_central_hostname>>	UCS Central fully qualified domain name (FQDN)	
<<var_ftp_server>>	Accessible FTP Server IP	
<<var_nexus_vpc_domain_id>>	Unique Cisco Nexus switch VPC domain ID	
<<var_nexus_1110x-1>>	Cisco Nexus 1110X-1 host name	
<<var_nexus_1110x-2>>	Cisco Nexus 1110X-2 host name	
<<var_fabric _a_fcoe_vlan_id>>	Fabric A FCoE VLAN ID	

Table 5 Configuration Variables

Variable	Description	Customer Implementation Value
<<var_vsan_a_id>>	Fabric A VSAN ID	
<<var_fabric_b_fcoe_vlan_id>>	Fabric B FCoE VLAN ID	
<<var_vsan_b_id>>	Fabric B VSAN ID	
<<var_vsm_domain_id>>	Unique Cisco Nexus 1000v virtual supervisor module (VSM) domain ID	
<<var_vsm_mgmt_ip>>	Cisco Nexus 1000v VSM management IP address	
<<var_vsm_mgmt_mask>>	In-band management network netmask	
<<var_vsm_mgmt_gw>>	In-band management network default gateway	
<<var_vsm_hostname>>	Cisco Nexus 1000v VSM host name	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address	
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address	
<<var_cimc_ip>>	Out-of-band management IP for each Cisco Nexus 1110-X CIMC	
<<var_cimc_mask>>	Out-of-band management network netmask	
<<var_cimc_gateway>>	Out-of-band management network default gateway	
<<var_1110x_domain_id>>	Unique Cisco Nexus 110-X domain ID	
<<var_1110x_vsa>>	Virtual storage appliance (VSA) host name	
<<var_1110x_vsa_ip>>	In-band VSA management IP address	
<<var_1110x_vsa_mask>>	In-band management network netmask	
<<var_1110x_vsa_gateway>>	In-band management network default gateway	
<<var_vmhost_infra01_ip>>	VMware ESXi host 01 in-band management IP	
<<var_vmhost_infra02_ip>>	VMware ESXi host 02 in-band management IP	
<<var_nfs_vlan_id_ip_host-01>>	NFS VLAN IP address for ESXi host 01	

Table 5 Configuration Variables

Variable	Description	Customer Implementation Value
<<var_nfs_vlan_id_mask_host-01>>	NFS VLAN netmask for ESXi host 01	
<<var_vmotion_vlan_id_ip_host-01>>	vMotion VLAN IP address for ESXi host 01	
<<var_vmotion_vlan_id_mask_host-01>>	vMotion VLAN netmask for ESXi host 01	
<<var_nfs_vlan_id_ip_host-02>>	NFS VLAN IP address for ESXi host 02	
<<var_nfs_vlan_id_mask_host-02>>	NFS VLAN netmask for ESXi host 02	
<<var_vcenter_server_ip>>	vCenter Server IP	
<<var_vm_host_infra_01_A_wwpn>>	WWPN of VM-Host-Infra-01 vHBA-A	
<<var_vm_host_infra_02_A_wwpn>>	WWPN of VM-Host-Infra-02 vHBA-A	
<<var_controller1_1a_wwpn>>	WWPN of storage controller 1 port 1a	
<<var_controller2_1a_wwpn>>	WWPN of storage controller 2 port 1a	
<<var_vm_host_infra_01_B_wwpn>>	WWPN of VM-Host-Infra-01 vHBA-B	
<<var_vm_host_infra_02_B_wwpn>>	WWPN of VM-Host-Infra-02 vHBA-B	
<<var_controller1_2a_wwpn>>	WWPN of storage controller 1 port 2a	
<<var_controller2_2a_wwpn>>	WWPN of storage controller 2 port 2a	

Physical Infrastructure

FlexPod Cabling on Data ONTAP Operating in 7-Mode

Figure 2 shows the cabling diagram for a FlexPod configuration using Data ONTAP operating on 7-mode.

Figure 2 FlexPod Cabling Diagram in Data ONTAP 7-Mode

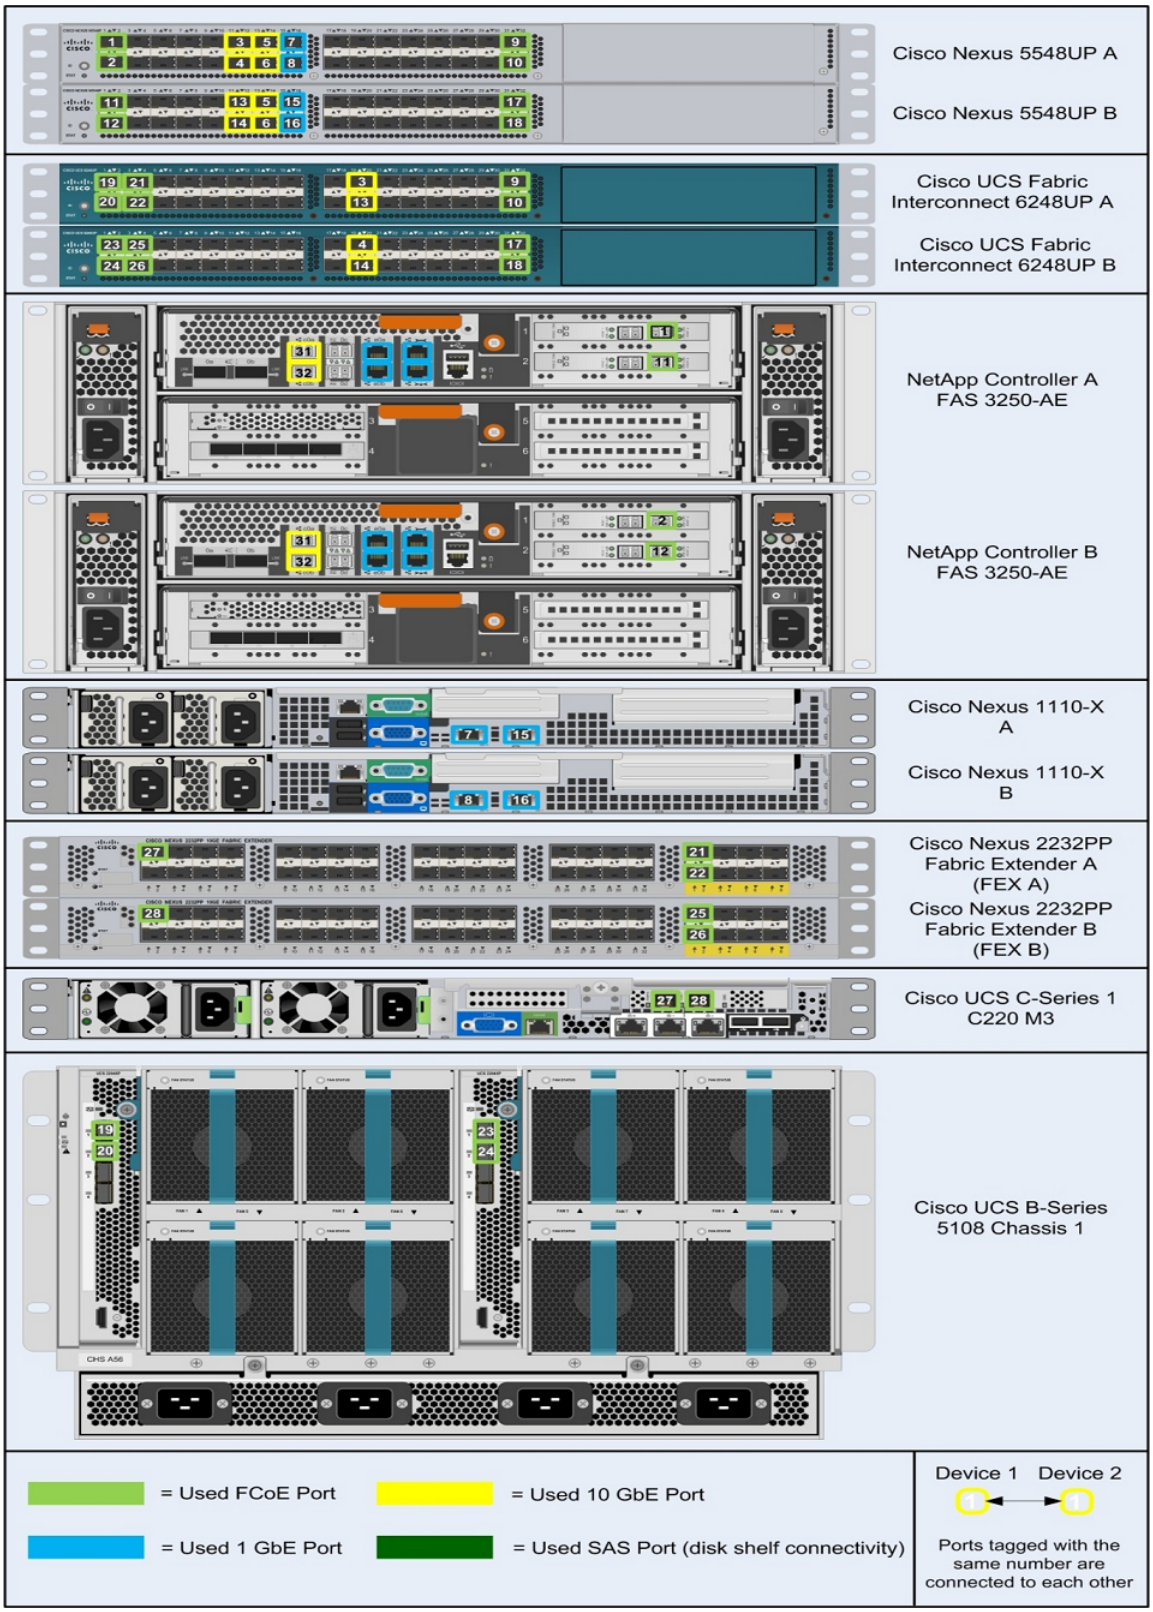

The information provided in Table 6 through Table 17 corresponds to each connection shown in Figure 2.

Table 6 Cisco Nexus 5548 A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5548 Switch A	Eth1/1	10GbE	NetApp controller A	e1a	1
	Eth1/2	10GbE	NetApp controller B	e1a	2
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/19	3
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/19	4
	Eth1/13	10GbE	Cisco Nexus 5548 B	Eth1/13	5
	Eth1/14	10GbE	Cisco Nexus 5548 B	Eth1/14	6
	Eth1/15	1GbE	Cisco Nexus 1110-X A	LOM A	7
	Eth1/16	1GbE	Cisco Nexus 1110-X B	LOM A	8
	Eth1/31	10GbE	Cisco UCS fabric interconnect A	Eth1/31	9
	Eth1/32	10GbE	Cisco UCS fabric interconnect A	Eth1/32	10
	MGMT0	1GbE	1GbE management switch	Any	

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

Table 7 Cisco Nexus 5548 B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Ports	Cabling Code
Cisco Nexus 5548 Switch B	Eth1/1	10GbE	NetApp controller 1	e2a	11
	Eth1/2	10GbE	NetApp controller 2	e2a	12
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/20	13
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/20	14
	Eth1/13	10GbE	Cisco Nexus 5548 A	Eth1/13	5
	Eth1/14	10GbE	Cisco Nexus 5548 A	Eth1/14	6
	Eth1/15	1GbE	Cisco Nexus 1110-X A	LOM B	15
	Eth1/16	1GbE	Cisco Nexus 1110-X B	LOM B	16
	Eth1/31	10GbE	Cisco UCS fabric interconnect B	Eth1/31	17
	Eth1/32	10GbE	Cisco UCS fabric interconnect B	Eth1/32	18
	MGMT0	1GbE	GbE management switch	Any	

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

Table 8 NetApp Controller 1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp Controller 1	e0M	100MbE	100MbE management switch	Any	
	e0P	1GbE	SAS shelves	ACP port	
	e1a	10GbE	Cisco Nexus 5548 A	Eth1/1	1
	e2a	10GbE	Cisco Nexus 5548 B	Eth1/1	11

Table 9 NetApp Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp Controller 2	e0M	100MbE	100MbE management switch	Any	
	e0P	1GbE	SAS shelves	ACP port	

Table 9 NetApp Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	e1a	10GbE	Cisco Nexus 5548 A	Eth1/2	2
	e2a	10GbE	Cisco Nexus 5548 B	Eth1/2	12

Table 10 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS Fabric Interconnect A	Eth1/19	10GbE	Cisco Nexus 5548 A	Eth1/11	3
	Eth1/20	10GbE	Cisco Nexus 5548 B	Eth1/11	13
	Eth1/1	10GbE	Cisco UCS Chassis 1 FEX A	Port1	19
	Eth1/2	10GbE	Cisco UCS Chassis 1 FEX A	Port2	20
	Eth1/3	10GbE	Cisco Nexus 2232PP FEX A	Port2/1	21
	Eth1/4	10GbE	Cisco Nexus 2232PP FEX A	Port2/2	22
	Eth1/31	10GbE	Cisco Nexus 5548 A	Eth1/31	9
	Eth1/32	10GbE	Cisco Nexus 5548 A	Eth1/32	10
	MGMT0	1GbE	1GbE management switch	Any	
	L1	1GbE	Cisco UCS fabric interconnect B	L1	
	L2	1GbE	Cisco UCS fabric interconnect B	L2	

Table 11 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS Fabric Interconnect B	Eth1/19	10GbE	Cisco Nexus 5548 A	Eth1/12	4
	Eth1/20	10GbE	Cisco Nexus 5548 B	Eth1/12	14
	Eth1/1	10GbE	Cisco UCS Chassis 1 FEX B	Port1	23
	Eth1/2	10GbE	Cisco UCS Chassis 1 FEX B	Port2	24
	Eth1/3	10GbE	Cisco Nexus 2232PP FEX B	Port2/1	25
	Eth1/4	10GbE	Cisco Nexus 2232PP FEX B	Port2/2	26
	Eth1/31	10GbE	Cisco Nexus 5548 B	Eth1/31	17
	Eth1/32	10GbE	Cisco Nexus 5548 B	Eth1/32	18
	MGMT0	1GbE	GbE management switch	Any	
	L1	1GbE	Cisco UCS fabric interconnect A	L1	
	L2	1GbE	Cisco UCS fabric interconnect A	L2	

Table 12 Cisco Nexus 2232PP FEX A

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 2232PP Fabric Extender (FEX A)	Port 1	10GbE	Cisco UCS C-Series 1	Port 0	27
	Port 2	10GbE	Cisco UCS C-Series 2	Port 0	
	Port 3	10GbE	Cisco UCS C-Series 3	Port 0	
	Port 4	10GbE	Cisco UCS C-Series 4	Port 0	
	Port 2/1	10GbE	Cisco UCS fabric interconnect A	Eth1/3	21
	Port 2/2	10GbE	Cisco UCS fabric interconnect A	Eth1/4	22

Table 13 Cisco Nexus 2232PP FEX B

Local Device	Local Port	Connection	Remote Devices	Remote Port	Cabling Code
Cisco Nexus 2232PP Fabric Extender (FEX B)	Port 1	10GbE	Cisco UCS C-Series 1	Port1	28
	Port 2	10GbE	Cisco UCS C-Series 2	Port1	
	Port 3	10GbE	Cisco UCS C-Series 3	Port1	
	Port 4	10GbE	Cisco UCS C-Series 4	Port1	
	Port 2/1	10GbE	Cisco UCS fabric interconnect B	Eth1/3	25
	Port 2/2	10GbE	Cisco UCS fabric interconnect B	Eth1/4	26

Table 14 Cisco UCS C-Series 1

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 1	Port0	10GbE	Cisco Nexus 2232PP Fabric Extender A (FEX A)	Port1	27
	Port1	10GbE	Cisco Nexus 2232PP Fabric Extender B (FEX B)	Port1	28

Table 15 Cisco UCS C-Series 2

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 2	Port0	10GbE	Cisco Nexus 2232PP Fabric Extender A (FEX A)	Port2	
	Port1	10GbE	Cisco Nexus 2232PP Fabric Extender B (FEX B)	Port2	

Table 16 Cisco UCS C-Series 3

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 3	Port0	10GbE	Cisco Nexus 2232PP Fabric Extender A (FEX A)	Port3	
	Port1	10GbE	Cisco Nexus 2232PP Fabric Extender B (FEX B)	Port3	

Table 17 Cisco UCS C-Series 4

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 4	Port0	10GbE	Cisco Nexus 2232PP Fabric Extender A (FEX A)	Port4	
	Port1	10GbE	Cisco Nexus 2232PP Fabric Extender B (FEX B)	Port4	

Table 18 NetApp FAS3250 Card Layout

Slot	Part Number	Description
1	X1140A-R6	Unified target 2-port 10GbE (ports e1a and e1b)
2	X1140A-R6	Unified target 2-port 10GbE (ports e2a and e2b)
3	X1971A-R5	Flash Cache™ – 512GB
4	X2065A-R6	SAS, 4-port, 6Gb

Table 19 Cisco C220M3 Card Layout for Single-wire Management

Slot	Part Number	Description
1	Cisco UCS VIC 1225	CNA 2-port 10GbE (ports 0 and 1)

Storage Configuration

Controller FAS32xx Series

Table 20 Controller FAS32xx Series Prerequisites

Requirements	Reference	Comments
Physical site where storage system needs to be installed must be ready	Site Reference Guide	Refer to the "Site Preparation" section.
Storage system connectivity requirements	Site Reference Guide	Refer to the "System Connectivity Requirements" section.
Storage system general power requirements	Site Reference Guide	Refer to the "Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pin-out Requirements" section.
Storage system model-specific requirements	Site Reference Guide	Refer to the "32xx Series Systems" section.

NetApp Hardware Universe

NetApp Hardware Universe provides supported hardware and software components for specific Data ONTAP versions. It provides configuration information for all NetApp storage appliances currently supported by the Data ONTAP software. It also provides a table of component compatibilities.

1. Make sure that the hardware and software components are supported with the version of Data ONTAP that you plan to install by checking the NetApp Hardware Universe at the NetApp Support site.
2. Access the Hardware Universe Application to view the System Configuration guides. Click the **Controllers** tab to view the compatibility between Data ONTAP software versions and NetApp storage appliances with the desired specifications.
3. Alternatively, to compare components by storage appliance, click **Compare Storage Systems**.

Controllers

Follow the physical installation procedures for the controllers in the FAS32xx documentation in NetApp Support site at:

https://now.netapp.com/NOW/knowledge/docs/hardware/filer/210-05224+A0.pdf

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives.

For a complete list of supported disk shelves, see:
http://support.netapp.com/documentation/productlibrary/index.html?productID=30147

When using SAS disk shelves with NetApp storage controllers, make sure that the cabling guidelines prescribed in the SAS Disk Shelves Universal SAS and ACP Cabling Guide is followed. The *SAS Disk Shelves Universal SAS and ACP Cabling Guide* is available at:

https://support.netapp.com/NOW/knowledge/docs/hardware/filer/215-05500_A0.pdf?isLegacy=true

Data ONTAP 8.2

Note The version of Data ONTAP 7-Mode used is 8.2P4, which is a patch release. Any reference to Data ONTAP 8.2 in this document refers to the patch release.

Complete the Configuration Worksheet

Before running the setup script, complete the configuration worksheet from the product manual.

For more information on accessing the Configuration Worksheet, see Configuration Guide at:

https://library.netapp.com/ecm/ecm_download_file/ECMP1155590

Note To access Configuration Worksheet, you may need to have access to NetApp Support site.

Assign Controller Disk Ownership and Initialize Storage

This section provides details for assigning disk ownership and disk initialization and verification.

Typical best practices should be followed when determining the number of disks to assign to each controller head. You may choose to assign a disproportionate number of disks to a given storage controller in an HA pair, depending on the intended workload.

In this reference architecture, half the total number of disks in the environment is assigned to one controller and the remainder to its partner.

Table 21 Controller Details

Detail	Detail Value
Controller 1 MGMT IP	<<var_controller1_e0m_ip>>
Controller 1 netmask	<<var_controller1_mask>>
Controller 1 gateway	<<var_controller1_mgmt_gateway>>
URL of the Data ONTAP boot software	<<var_url_boot_software>>
Controller 2 MGMT IP	<<var_controller2_e0m_ip>>
Controller 2 netmask	<<var_controller2_mask>>
Controller 2 gateway	<<var_controller2_mgmt_gateway>>

Controller1

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. From the Loader-A prompt:

```
printenv
```

3. If the last-OS-booted-ver parameter is not set to 8.2P4, proceed to step 4 to load Data ONTAP 8.2P4 software. If Data ONTAP 8.2P4 is already loaded, proceed to step 16.
4. Allow the system to boot up.

```
autoboot
```

5. Press Ctrl-C when the Press Ctrl-C for Boot Menu message appears.

Note If Data ONTAP 8.2P4 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.2P4 is the version being booted, then select option 8 and yes to reboot the node. Then proceed with step 15.

6. To install new software, select option 7.

```
7
```

7. Answer yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

9. Select yes to reboot now.

```
y
```

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

```
<<var_controller1_e0m_ip>>
<<var_controller1_mask>>
<<var_controller1_mgmt_gateway>>
```

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

Note This web server must be pingable.

```
<<var_url_boot_software>>
```

12. Press **Enter** for the username, indicating no user name.

```
Enter
```

13. Enter yes to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

14. Enter yes to reboot the node.

```
y
```

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

15. Press **Ctrl-C** to exit autoboot when you see this message:

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

16. From the LOADER-A prompt, enter:

```
printenv
```

Note If bootarg.init.boot_clustered true is listed, the system is not set to boot in Data ONTAP 7-Mode.

17. If the system is not set to boot in Data ONTAP 7-Mode, at the LOADER prompt, enter the following command to make sure the system boots in Data ONTAP 7-Mode.

```
unsetenv bootarg.init.boot_clustered
setenv bootarg.bsdportname e0M
```

18. At the LOADER-A prompt, enter:

```
autoboot
```

19. When you see Press Ctrl-C for Boot Menu:

```
Ctrl - C
```

20. Select option 4 for clean configuration and initialize all disks.

```
4
```

21. Answer yes to Zero disks, reset config and install a new file system.

```
yes
```

22. Enter yes to erase all the data on the disks.

```
yes
```

Note The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots. You can continue to controller 2 configuration while the disks for controller 1 are zeroing.

Controller 2

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. From the Loader-A prompt:

```
printenv
```

3. If the last-OS-booted-ver parameter is not set to 8.2P4, proceed to step 4 to load Data ONTAP 8.2P4 software. If Data ONTAP 8.2P4 is already loaded, proceed to step 16.
4. Allow the system to boot up.

```
autoboot
```

5. Press Ctrl-C when the Press Ctrl-C for Boot Menu message appears.

Note If Data ONTAP 8.2P4 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.2P4 is the version being booted, then select option 8 and yes to reboot the node. Then proceed with step 15.

6. To install new software, first select option 7.

```
7
```

7. Answer yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

9. Select yes to reboot now.

```
y
```

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

```
<<var_controller2_e0m_ip>>
<<var_controller2_mask>>
<<var_controller2_mgmt_gateway>>
```

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

Note This web server must be pingable.

```
<<var_url_boot_software>>
```

12. Press Enter for the username, indicating no user name.

```
Enter
```

13. Enter yes to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

14. Enter yes to reboot the node.

```
y
```

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

15. Press Ctrl-C to exit autoboot when you see this message:

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

16. From the LOADER-A prompt, enter:

```
printenv
```

Note If bootarg.init.boot_clustered true is listed, the system is not set to boot in Data ONTAP 7-Mode.

17. If the system is not set to boot in Data ONTAP 7-Mode, at the LOADER prompt, enter the following command to make sure the system boots in Data ONTAP 7-Mode.

```
unsetenv bootarg.init.boot_clustered
setenv bootarg.bsdportname e0M
```

18. At the LOADER-A prompt, enter:

```
autoboot
```

19. When you see Press Ctrl-C for Boot Menu:

```
Ctrl - C
```

20. Select option 4 for clean configuration and initialize all disks.

```
4
```

21. Answer yes to Zero disks, reset config and install a new file system

```
yes
```

22. Enter yes to erase all the data on the disks.

```
yes
```

Note The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Run Setup Process

When Data ONTAP is installed on a new storage system, the following files are not populated:

- /etc/rc
- /etc/exports
- /etc/hosts
- /etc/hosts.equiv

Controller 1

1. Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.
2. Enter the following information:

```
Please enter the new hostname []:<<var_controller1>>
Do you want to enable IPv6? [n]: Enter

Do you want to configure interface groups? [n]: Enter
Please enter the IP address for Network Interface e0a []: Enter
```

Note Press Enter to accept the blank IP address.

```
Please enter the IP address for Network Interface e0b []: Enter
Please enter the IP address for Network Interface e1a []: Enter
Please enter the IP address for Network Interface e1b []: Enter
Please enter the IP address for Network Interface e2a []: Enter
Please enter the IP address for Network Interface e2b []: Enter
No IP address specified. Please set an IP address.
        e0M is a Data ONTAP dedicated management port.

        NOTE: Dedicated management ports cannot be used for data
        protocols (NFS, CIFS, iSCSI, NDMP or Snap*),
        and if they are configured they should be on an isolated management LAN.
        The default route will use dedicated mgmt ports only as the last resort,
        since data protocol traffic will be blocked by default.
Please enter the IP address for Network Interface e0M: <<var_controller1_e0m_ip>>
Please enter the netmask for Network Interface e0M [255.255.255.0]:
<<var_controller1_mask>>
Please enter flow control for e0M {none, receive, send, full} [full]:
Please enter the name or IP address of the IPv4 default gateway:
<<var_controller1_mgmt_gateway>>
```

Note If additional interface cards are installed in your storage controller, you will have additional questions about the interfaces on those cards.

3. Enter the following information:

```
The administration host is given root access to the filer's /etc files for system
administration. To allow /etc root access to all NFS clients enter RETURN below.
```

```
Please enter the name or IP address of the administration host:
<<var_adminhost_ip>>

Please enter timezone [GMT]: <<var_timezone>>
```

Note Example time zone: America/New_York.

```
Where is the filer located? []: <<var_location>>
Enter the root directory for HTTP files [/home/http]: Enter
Do you want to run DNS resolver? [n]: y
Please enter DNS domain name []: <<var_dns_domain_name>>
You may enter up to 3 nameservers
Please enter the IP address for first nameserver []: <<var_nameserver_ip>>
Do you want another nameserver? [n]:
```

Note Optionally enter up to three name server IP addresses.

```
Do you want to run NIS client? [n]: Enter
Press the Return key to continue through AutoSupport message
Would you like to configure SP LAN interface [y]: Enter
Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_controller1_sp_ip>>
Please enter the netmask for the SP []: <<var_controller1_sp_mask>>
Please enter the IP address for the SP gateway: <<var_controller1_sp_gateway>>
Please enter the name or IP address of the mail host: <<var_mailhost>>
New password: <<var_password>>
Retype new password <<var_password>>
```

4. Enter the root password to log in to controller 1.
5. Reboot controller 1.

```
reboot
```

6. When you see Press Ctrl-C for Boot Menu, enter:

```
Ctrl - C
```

7. Select 5 to boot into maintenance mode.

```
5
```

8. At the question, Continue with boot? enter:

```
y
```

9. To verify the HA status of your environment, enter:

```
ha-config show
```

Note If either component is not in HA mode, use the ha-config modify command to put the components in HA mode.

10. Reboot the controller:

```
halt
```

11. At the LOADER-A prompt, enter:

```
autoboot
```

12. Log in into the controller.
13. Data ONTAP would assign disks to Storage Controllers automatically if the disk autoassign setting was turned on. Use the “options disk.auto_assign” to verify the setting.
14. If disk autoassign was turned on, proceed with controller 2 else continue with Step 15
15. Reboot the controller.

```
reboot
```

16. When you see Press Ctrl-C for Boot Menu, enter:

```
Ctrl - C
```

17. Select 5 to boot into maintenance mode.

```
5
```

18. When prompted Continue with boot?, enter y.
19. To see how many disks are unowned, enter:

```
disk show -a
```

Note No disks should be owned in this list.

20. Assign disks.

Note This reference architecture allocates half the disks to each controller. However, workload design could dictate different percentages.

```
disk assign -n <<var_#_of_disks>>
```

21. Reboot the controller.

```
halt
```

22. At the LOADER-A prompt, enter:

```
autoboot
```

Controller 2

1. Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.
2. Enter the following information:

```
Please enter the new hostname []:<<var_controller2>>
Do you want to enable IPv6? [n]: Enter

Do you want to configure interface groups? [n]: Enter
Please enter the IP address for Network Interface e0a []: Enter
```

Note Press Enter to accept the blank IP address.

```
Please enter the IP address for Network Interface e0b []: Enter
Please enter the IP address for Network Interface e1a []: Enter
Please enter the IP address for Network Interface e1b []: Enter
Please enter the IP address for Network Interface e2a []: Enter
Please enter the IP address for Network Interface e2b []: Enter
No IP address specified. Please set an IP address.
        e0M is a Data ONTAP dedicated management port.

        NOTE: Dedicated management ports cannot be used for data
        protocols (NFS, CIFS, iSCSI, NDMP or Snap*),
        and if they are configured they should be on an isolated management LAN.
        The default route will use dedicated mgmt ports only as the last resort,
        since data protocol traffic will be blocked by default.
Please enter the IP address for Network Interface e0M: <<var_controller2_e0m_ip>>
Please enter the netmask for Network Interface e0M [255.255.255.0]:
<<var_controller2_mask>>
Please enter flow control for e0M {none, receive, send, full} [full]:
```

```
Please enter the name or IP address of the IPv4 default gateway:
<<var_controller2_mgmt_gateway>>
```

Note If additional interface cards are installed in your storage controller, you will have additional questions about the interfaces on those cards.

3. Enter the following information:

```
The administration host is given root access to the filer's /etc files for system
administration.  To allow /etc root access to all NFS clients enter RETURN below.
Please enter the name or IP address of the administration host:
<<var_adminhost_ip>>

Please enter timezone [GTM]: <<var_timezone>>
```

Note Example time zone: America/New_York.

```
Where is the filer located? <<var_location>>
Enter the root directory for HTTP files [/home/http]: Enter
Do you want to run DNS resolver? [n]: y
Please enter DNS domain name []: <<var_dns_domain_name>>
You may enter up to 3 nameservers
Please enter the IP address for first nameserver []: <<var_nameserver_ip>>
Do you want another nameserver? [n]:
```

Note Optionally enter up to three name server IP addresses.

```
Do you want to run NIS client? [n]: Enter
Press the Return key to continue through AutoSupport message
Would you like to configure SP LAN interface [y]: Enter
Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_controller2_sp_ip>>
Please enter the netmask for the SP []: <<var_controller2_sp_mask>>
Please enter the IP address for the SP gateway: <<var_controller2_sp_gateway>>
Please enter the name or IP address of the mail host: <<var_mailhost>>
New password: <<var_password>>
Retype new password <<var_password>>
```

4. Enter the root password to log in to controller 2.
5. Reboot controller 2.

```
reboot
```

6. When you see Press Ctrl-C for Boot Menu, enter:

```
Ctrl - C
```

7. Select 5 to boot into maintenance mode.

```
5
```

8. At the question, Continue with boot? enter:

```
y
```

9. To verify the HA status of your environment, enter:

```
ha-config show
```

Note If either component is not in HA mode, use the ha-config modify command to put the components in HA mode.

10. Reboot the controller:

```
halt
```

11. At the LOADER-A prompt, enter:

```
autoboot
```

12. Log in into the controller.
13. Data ONTAP would assign disks to storage controllers automatically if the disk autoassign setting was turned on. Use the “options disk.auto_assign” to verify the setting.
14. If disk autoassign was turned on, proceed with controller 2 else continue with step 15.
15. Reboot the controller.

```
reboot
```

16. When you see Press Ctrl-C for Boot Menu, enter:

```
Ctrl - C
```

17. Select 5 to boot into maintenance mode.

```
5
```

18. When prompted Continue with boot?, enter y.
19. To see how many disks are unowned, enter:

```
disk show -a
```

Note No disks should be owned in this list.

20. Assign disks.

Note This reference architecture allocates half the disks to each controller. However, workload design could dictate different percentages.

```
disk assign -n <<var_#_of_disks>>
```

21. Reboot the controller.

```
halt
```

22. At the LOADER-A prompt, enter:

```
autoboot
```

Upgrade the Service Processor on Each Node to the Latest Release

With Data ONTAP 8.2, you must upgrade to the latest Service Processor (SP) firmware to take advantage of the latest updates available for the remote management device.

1. Using a web browser, connect to: http://support.netapp.com/NOW/cgi-bin/fw.
2. Navigate to the Service Process Image for installation from the Data ONTAP prompt page for your storage platform.
3. Proceed to the Download page for the latest release of the SP Firmware for your storage platform.
4. Using the instructions on this page, update the SPs on both controllers. You will need to download the .zip file to a web server that is reachable from the management interfaces of the controllers.

Aggregates in Data ONTAP 7-Mode

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

Controller 1

Execute the following command to create a new aggregate:

```
aggr create aggr1 <<var_#_of_disks>>
```

Note

- Leave 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 and you may add disks to an aggregate when there is a need for additional storage.
- The aggregate cannot be created until the disk zero completes. Use the aggr status command to display the aggregate creation status, Do not proceed until the aggregates are online

Controller 2

Execute the following command to create a new aggregate:

```
aggr create aggr1 <<var_#_of_disks>>
```

Note

- Leave 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 and you may add disks to an aggregate when there is a need for additional storage.
- The aggregate cannot be created until the disk zero completes. Use the aggr status command to display the aggregate creation status, Do not proceed until the aggregates are online

Flash Cache

Controller 1 and Controller 2

Execute the following commands to enable Flash Cache:

```
options flexscale.enable on
options flexscale.lopri_blocks off
options flexscale.normal_data_blocks on
```

Note For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, see TR-3832: Flash Cache and PAM Best Practices Guide at: http://media.netapp.com/documents/tr-3832.pdf. Before customizing the settings, determine whether the custom settings are required or whether the default settings are sufficient.

IFGRP LACP

Since this type of interface group requires two or more Ethernet interfaces and a switch that supports LACP, make sure that the switch is configured properly.

Controller 1 and Controller 2

Run the following command on the command line and also add it to the /etc/rc file, so it is activated upon boot:

```
ifgrp create lacp ifgrp0 -b port e1a e2a
wrfile -a /etc/rc "ifgrp create lacp ifgrp0 -b port e1a e2a"
```

Note All interfaces must be in down status before being added to an interface group.

VLAN

Controller 1 and Controller 2

Follow these steps to create a VLAN interface for NFS data traffic.

```
vlan create ifgrp0 <<var_nfs_vlan_id>>
wrfile -a /etc/rc "vlan create ifgrp0 <<var_nfs_vlan_id>>"
```

IP Config

Controller 1 and Controller 2

Run the following commands on the command line.

```
ifconfig ifgrp0-<<var_nfs_vlan_id>> <<var_nfs_ip>> netmask <<var_nfs_mask>> mtusize
9000 partner ifgrp0-<<var_nfs_vlan_id>>
wrfile -a /etc/rc " ifconfig ifgrp0-<<var_nfs_vlan_id>> <<var_nfs_ip>> netmask
<<var_nfs_mask>> mtusize 9000 partner ifgrp0-<<var_nfs_vlan_id>>"
```

Cisco Discovery Protocol

Follow these steps to enable Cisco Discovery Protocol (CDP) on controller 1 and controller 2.

Controller 1 and Controller 2

Enable CDP

```
options cdpd.enable on
```

Active-Active Controller Configuration

Controller 1 and Controller 2

Enable two storage controllers to an active-active configuration.

1. Enable high availability.

```
options cf.mode ha
```

2. Reboot each storage controller.

```
reboot
```

3. Log back in to both controllers.

Controller 1

Enable failover on Controller 1, if it is not enabled already.

```
cf enable
```

NFSv3

Controller 1 and Controller 2

1. Add a license for NFS.

```
license add <<var_nfs_license>>
```

2. Set the following recommended options that enable NFS version 3.

```
options nfs.tcp.enable on
options nfs.udp.enable off
options nfs.v3.enable on
```

3. Enable NFS.

```
nfs on
```

FCP

Controller 1 and Controller 2

1. License FCP.

```
license add <<var_fc_license>>
```

2. Start the FCP service.

```
options licensed_feature.fcp.enable on
fcp start
```

3. Record the WWPN or FC port name for later use.

```
fcp show adapters
```

NTP

The following commands configure and enable time synchronization on the storage controller. You must have either a publically available IP address or your company's standard NTP server name or IP address.

Controller 1 and Controller 2

1. Run the following commands to configure and enable the NTP server:

```
date <<var_date>>
```

2. Enter the current date in the format of [[[[CC]yy]mm]dd]hhmm[.ss]].

   For example: date 201208311436; would set the date to August 31st 2012 at 14:36.

```
options timed.servers <<var_global_ntp_server_ip>>
options timed.enable on
```

Data ONTAP SecureAdmin

Secure API access to the storage controller must be configured.

Controller 1

1. Issue the following as a one-time command to generate the certificates used by the Web services for the API.

```
secureadmin setup ssl
SSL Setup has already been done before. Do you want to proceed? [no] y
Country Name (2 letter code) [US]: <<var_country_code>>
State or Province Name (full name) [California]: <<var_state>>
Locality Name (city, town, etc.) [Santa Clara]: <<var_city>>
Organization Name (company) [Your Company]: <<var_org>>
Organization Unit Name (division): <<var_unit>>
Common Name (fully qualified domain name) [<<var_controller1_fqdn>>]: Enter
Administrator email: <<var_admin_email>>
Days until expires [5475] : Enter
Key length (bits) [512] : <<var_key_length>>
```

Note NetApp recommends that your key length be 1024.

   After the initialization, the CSR is available in the file /etc/keymgr/csr/secureadmin_tmp.pem.

2. Configure and enable SSL and HTTPS for API access using the following options.

```
options httpd.access none
options httpd.admin.enable off
options httpd.admin.ssl.enable on
options ssl.enable on
```

Controller 2

1. Issue the following as a one-time command to generate the certificates used by the Web services for the API.

```
secureadmin setup ssl
```

```
SSL Setup has already been done before. Do you want to proceed? [no] y
Country Name (2 letter code) [US]: <<var_country_code>>
State or Province Name (full name) [California]: <<var_state>>
Locality Name (city, town, etc.) [Santa Clara]: <<var_city>>
Organization Name (company) [Your Company]: <<var_org>>
Organization Unit Name (division): <<var_unit>>
Common Name (fully qualified domain name) [<<var_controller2_fqdn>>]: Enter
Administrator email: <<var_admin_email>>
Days until expires [5475] : Enter
Key length (bits) [512] : <<var_key_length>>
```

Note NetApp recommends that your key length be 1024.

After the initialization, the CSR is available in the file /etc/keymgr/csr/secureadmin_tmp.pem.

2. Configure and enable SSL and HTTPS for API access using the following options.

```
options httpd.access none
options httpd.admin.enable off
options httpd.admin.ssl.enable on
options ssl.enable on
```

Secure Shell

SSH must be configured and enabled.

Controller 1 and Controller 2

1. Use the following one-time command to generate host keys.

```
secureadmin disable ssh
secureadmin setup -f -q ssh 768 512 1024
```

2. Use the following options to configure and enable SSH.

```
options ssh.idle.timeout 60
options autologout.telnet.timeout 5
```

SNMP

Controller 1 and Controller 2

1. Run the following commands to configure SNMP basics, such as the local and contact information. When polled, this information displays as the sysLocation and sysContact variables in SNMP.

```
snmp contact "<<var_admin_email>>"
snmp location "<<var_location>>"
snmp init 1
options snmp.enable on
```

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

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

SNMPv1

Controller 1 and Controller 2

1. Set the shared secret plain-text password, which is called a community.

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

Note Use the delete all command with caution. If community strings are used for other monitoring products, the delete all command will remove them.

SNMPv3

SNMPv3 requires a user to be defined and configured for authentication.

Controller 1 and Controller 2

Create a user called snmpv3user.

```
useradmin role add snmp_requests -a login-snmp
useradmin group add snmp_managers -r snmp_requests
useradmin user add snmpv3user -g snmp_managers
New Password: <<var_password>>
Retype new password: <<var_password>>
```

AutoSupport HTTPS

AutoSupport™ sends support summary information to NetApp through HTTPS.

Controller 1 and Controller 2

Execute the following commands to configure AutoSupport:

```
options autosupport.noteto <<var_admin_email>>
```

Security Best Practices

Apply the following commands according to local security policies.

Controller 1 and Controller 2

Run the following commands to enhance security on the storage controller:

```
options rsh.access none
options webdav.enable off
options security.passwd.rules.maximum 14
options security.passwd.rules.minimum.symbol 1
options security.passwd.lockout.numtries 6
options autologout.console.timeout 5
```

Install Remaining Required Licenses and Enable MultiStore

Controller 1 and Controller 2

Install the following licenses to enable SnapRestore® and FlexClone®.

```
license add <<var_snaprestore_license>>
license add <<var_flex_clone_license>>
options licensed_feature.multistore.enable on
```

Enable NDMP

Run the following commands to enable NDMP.

Controller 1 and Controller 2

options ndmpd.enable on

Create FlexVol Volumes

Controller 1

Follow these steps to create two volumes on controller 1:

```
vol create esxi_boot -s none aggr1 100g
snap reserve esxi_boot 0
sis on /vol/esxi_boot
vol create infra_swap -s none aggr1 100g
snap reserve infra_swap 0
snap sched infra_swap 0 0 0
```

Controller 2

Follow these steps to create two volumes on controller 2:

```
vol create infra_datastore_1 -s none aggr1 500g
snap reserve infra_datastore_1 0
sis on /vol/infra_datastore_1
vol create OnCommandDB -s none aggr1 200g
snap reserve OnCommandDB 0
sis on /vol/OnCommandDB
```

NFS Exports

Follow these steps to create NFS exports on each controller.

Controller 1

```
exportfs -p
sec=sys,rw=<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,root=<<var_vm
_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,nosuid /vol/infra_swap

exportfs -p
sec=sys,ro,rw=<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs
_host_ip>>,root==<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_
nfs_host_ip>>,nosuid /vol/vol0
```

Controller 2

```
exportfs -p
sec=sys,rw=<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,root=<<var_vm
_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs_host_ip>>,nosuid /vol/infra_datastore_1

exportfs -p
sec=sys,ro,rw=<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_nfs
_host_ip>>,root==<<var_adminhost_ip>>:<<var_vm_infra01_nfs_host_ip>>:<<var_vm_infra02_
nfs_host_ip>>,nosuid /vol/vol0
```

LUN Creation

Follow these steps to create two LUNs on controller 1.

Controller 1

```
lun create -s 10g -t vmware -o noreserve /vol/esxi_boot/VM-Host-Infra-01
lun create -s 10g -t vmware -o noreserve /vol/esxi_boot/VM-Host-Infra-02
```

Server Configuration

FlexPod Cisco UCS Base

Perform Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlexPod environment. The following 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 6248UP Fabric Interconnect A

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

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect.

```
Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have choosen to setup a a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y
Enter the password for "admin": <<var_password>>
```

```
Confirm the password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)?
(yes/no) [n]: y
Enter the switch fabric (A/B)[]: A
Enter the system name: <<var_ucs_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 the DNS Server IPv4 address? (yes/no) [n]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? (yes/no) [n]: y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
```

2. Review the settings printed to the console. If they are correct, answer yes to apply and save the configuration.
3. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS 6248UP Fabric Interconnect B

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

1. Connect to the console port on the second Cisco UCS 6248 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.

FlexPod Cisco UCS FCoE vSphere on Data ONTAP 7-Mode

Log in to Cisco UCS Manager

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

1. Open a Web browser and navigate to the Cisco UCS 6248 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 2.1(3a)

This document assumes the use of Cisco UCS 2.1(3a). To upgrade the Cisco UCS Manager software and the UCS 6248 Fabric Interconnect software to version 2.1(3a), see Cisco UCS Manager Install and Upgrade Guides.

Add Block of IP Addresses for KVM Access

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

Note This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Choose **Pools** > **root** > **IP Pools** > **IP Pool ext-mgmt**.
3. In the Actions pane, choose **Create Block of IP Addresses**.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.
5. Click **OK** to create the IP block.
6. Click **OK** in the confirmation message window.

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, click the **Admin** tab in the navigation pane.
2. Choose **All** > **Timezone Management**.
3. In the Properties pane, choose the appropriate time zone in the Timezone menu.
4. Click **Save Changes**, and then click **OK**.
5. Click **Add NTP Server**.
6. Enter <<var_global_ntp_server_ip>> and click **OK**.
7. Click **OK**.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and of additional fabric extenders for further C-Series connectivity.

To modify the chassis discovery policy, follow these steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane and choose 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 2-link or set it 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 Server and Uplink Ports

To enable server and uplink ports, follow these steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
2. Choose **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand Ethernet Ports.
4. Choose the ports that are connected to the chassis or to the Cisco 2232 FEX (two per FEX), right-click them, and choose Configure as Server Port.
5. Click **Yes** to confirm server ports and click **OK**.
6. Verify that the ports connected to the chassis or to the Cisco 2232 FEX are now configured as server ports.

Figure 3 Configured Server Ports

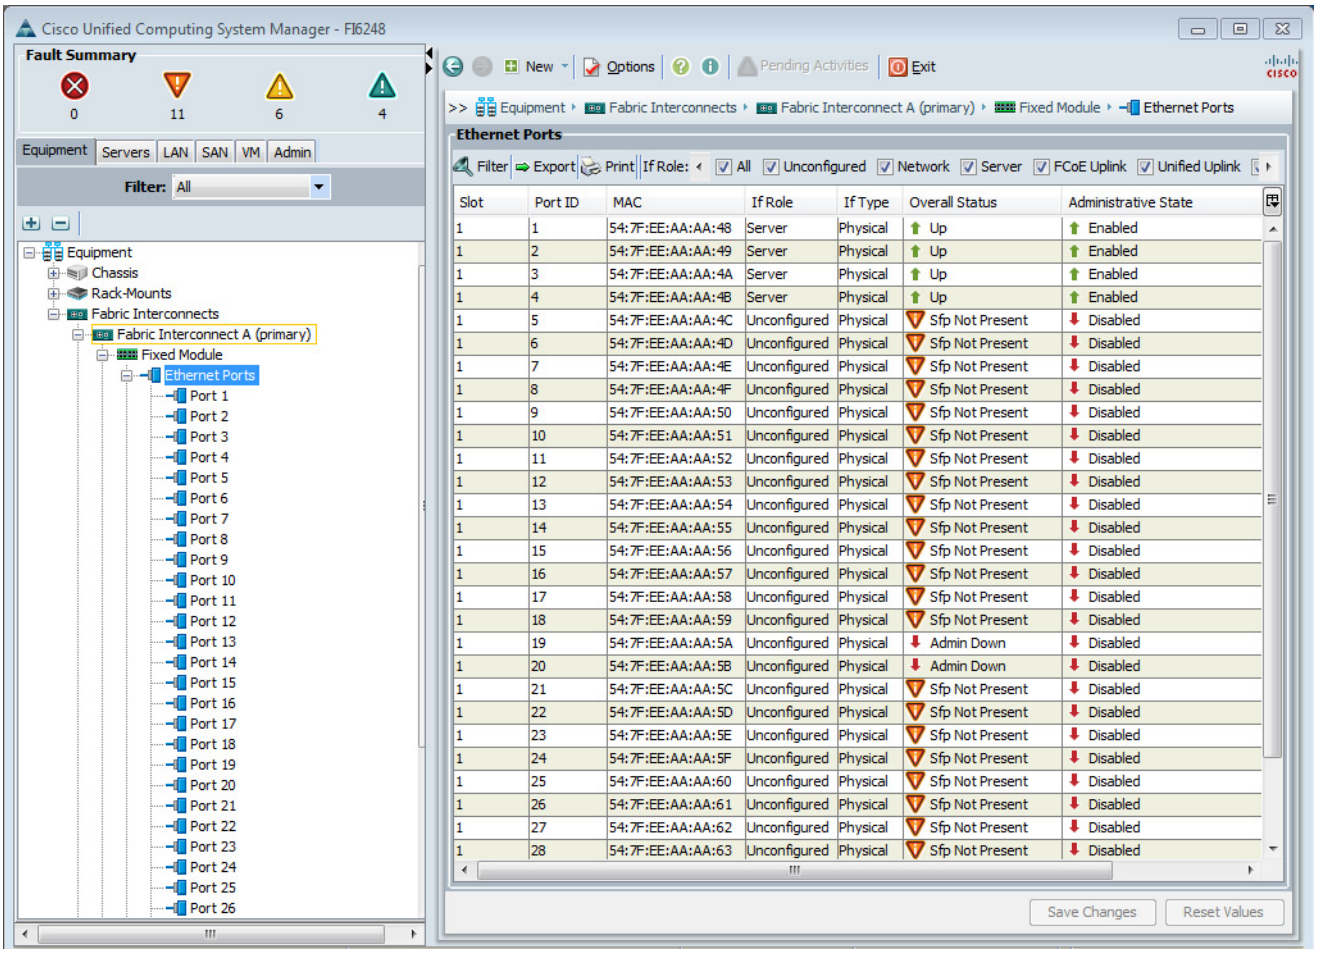

Slot	Port ID	MAC	If Role	If Type	Overall Status	Administrative State
1	1	54:7F:EE:AA:AA:48	Server	Physical	Up	Enabled
1	2	54:7F:EE:AA:AA:49	Server	Physical	Up	Enabled
1	3	54:7F:EE:AA:AA:4A	Server	Physical	Up	Enabled
1	4	54:7F:EE:AA:AA:4B	Server	Physical	Up	Enabled
1	5	54:7F:EE:AA:AA:4C	Unconfigured	Physical	Sfp Not Present	Disabled
1	6	54:7F:EE:AA:AA:4D	Unconfigured	Physical	Sfp Not Present	Disabled
1	7	54:7F:EE:AA:AA:4E	Unconfigured	Physical	Sfp Not Present	Disabled
1	8	54:7F:EE:AA:AA:4F	Unconfigured	Physical	Sfp Not Present	Disabled
1	9	54:7F:EE:AA:AA:50	Unconfigured	Physical	Sfp Not Present	Disabled
1	10	54:7F:EE:AA:AA:51	Unconfigured	Physical	Sfp Not Present	Disabled
1	11	54:7F:EE:AA:AA:52	Unconfigured	Physical	Sfp Not Present	Disabled
1	12	54:7F:EE:AA:AA:53	Unconfigured	Physical	Sfp Not Present	Disabled
1	13	54:7F:EE:AA:AA:54	Unconfigured	Physical	Sfp Not Present	Disabled
1	14	54:7F:EE:AA:AA:55	Unconfigured	Physical	Sfp Not Present	Disabled
1	15	54:7F:EE:AA:AA:56	Unconfigured	Physical	Sfp Not Present	Disabled
1	16	54:7F:EE:AA:AA:57	Unconfigured	Physical	Sfp Not Present	Disabled
1	17	54:7F:EE:AA:AA:58	Unconfigured	Physical	Sfp Not Present	Disabled
1	18	54:7F:EE:AA:AA:59	Unconfigured	Physical	Sfp Not Present	Disabled
1	19	54:7F:EE:AA:AA:5A	Unconfigured	Physical	Admin Down	Disabled
1	20	54:7F:EE:AA:AA:5B	Unconfigured	Physical	Admin Down	Disabled
1	21	54:7F:EE:AA:AA:5C	Unconfigured	Physical	Sfp Not Present	Disabled
1	22	54:7F:EE:AA:AA:5D	Unconfigured	Physical	Sfp Not Present	Disabled
1	23	54:7F:EE:AA:AA:5E	Unconfigured	Physical	Sfp Not Present	Disabled
1	24	54:7F:EE:AA:AA:5F	Unconfigured	Physical	Sfp Not Present	Disabled
1	25	54:7F:EE:AA:AA:60	Unconfigured	Physical	Sfp Not Present	Disabled
1	26	54:7F:EE:AA:AA:61	Unconfigured	Physical	Sfp Not Present	Disabled
1	27	54:7F:EE:AA:AA:62	Unconfigured	Physical	Sfp Not Present	Disabled
1	28	54:7F:EE:AA:AA:63	Unconfigured	Physical	Sfp Not Present	Disabled

7. Choose ports 19 and 20 that are connected to the Cisco Nexus 5548 switches, right-click them, and choose Configure as Uplink Port.
8. Click **Yes** to confirm uplink ports and click **OK**.
9. Choose ports 31 and 32, which will serve as FCoE uplinks to the Cisco Nexus 5548 switches; right-click them; and choose Configure as FCoE Uplink Port.

10. Click **Yes** to confirm FCoE uplink ports and click **OK**.
11. Choose **Equipment** > **Fabric Interconnects** > **Fabric Interconnect B (subordinate)** > **Fixed Module**.
12. Expand Ethernet Ports.
13. Choose the ports that are connected to the chassis or to the Cisco 2232 FEX (two per FEX), right-click them, and choose **Configure as Server Port**.
14. Click **Yes** to confirm server ports and click **OK**.
15. Choose ports 19 and 20 that are connected to the Cisco Nexus 5548 switches, right-click them, and choose **Configure as Uplink Port**.
16. Click **Yes** to confirm the uplink ports and click **OK**.
17. Choose ports 31 and 32 that will serve as FCoE uplinks to the Cisco Nexus 5548 switches, right-click them, and choose **Configure as FCoE Uplink Port**.
18. Click **Yes** to confirm FCoE uplink ports and click **OK**.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and external 2232 FEX modules, follow these steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
2. Expand Chassis.
3. Choose each chassis that is listed, right-click on each chassis and choose **Acknowledge Chassis**.

Figure 4 Acknowledging Cisco UCS Chassis

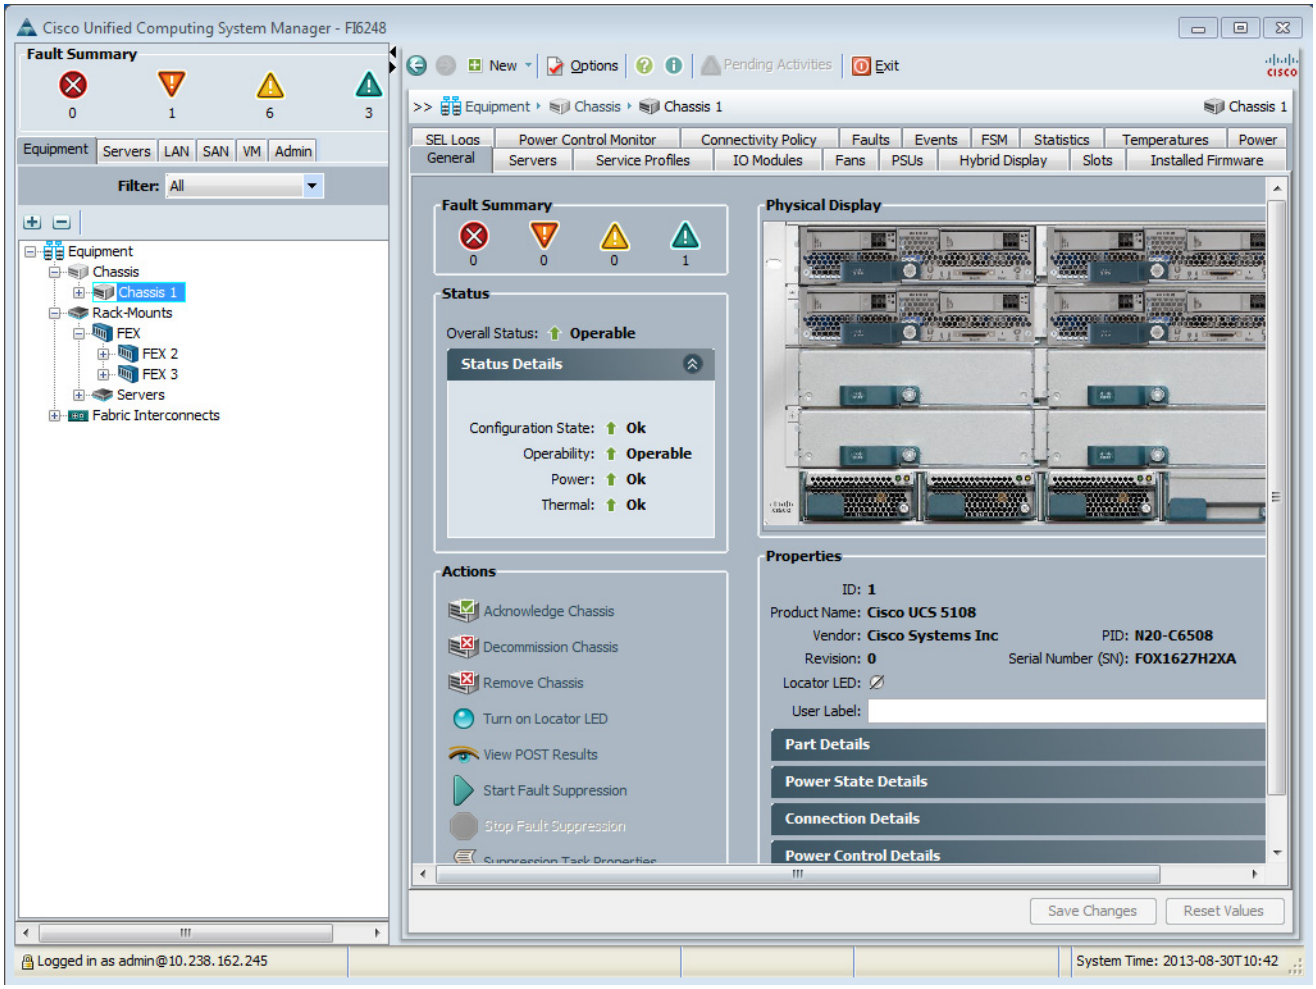

4. Click **Yes** and then click **OK** to complete acknowledging the chassis.
5. If C-Series servers are part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and choose **Acknowledge FEX**.

Figure 5 Acknowledging Cisco UCS Fabric Extenders

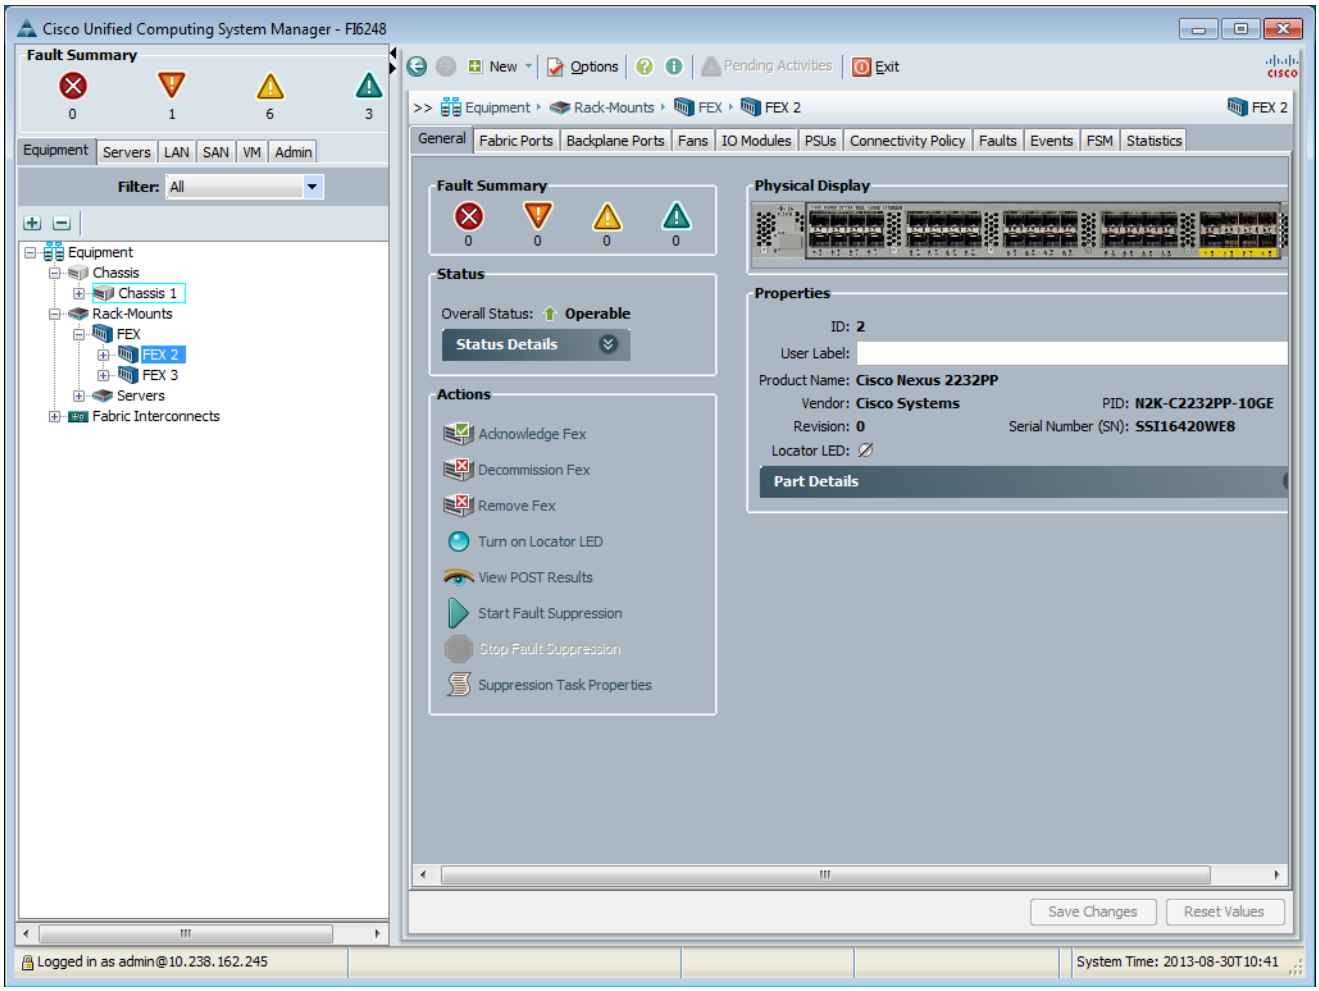

7. Click **Yes** and then click **OK** to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus 5548 Switches

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

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

Note In this procedure, two port channels are created: one from fabric A to both Cisco Nexus 5548 switches and one from fabric B to both Cisco Nexus 5548 switches.

2. Under **LAN > LAN Cloud**, expand the Fabric A tree.
3. Right-click Port Channels.
4. Choose Create Port Channel.
5. Enter 13 as the unique ID of the port channel.
6. Enter vPC-13-N5548 as the name of the port channel.

7. Click **Next**.

Figure 6 Creating Port Channels

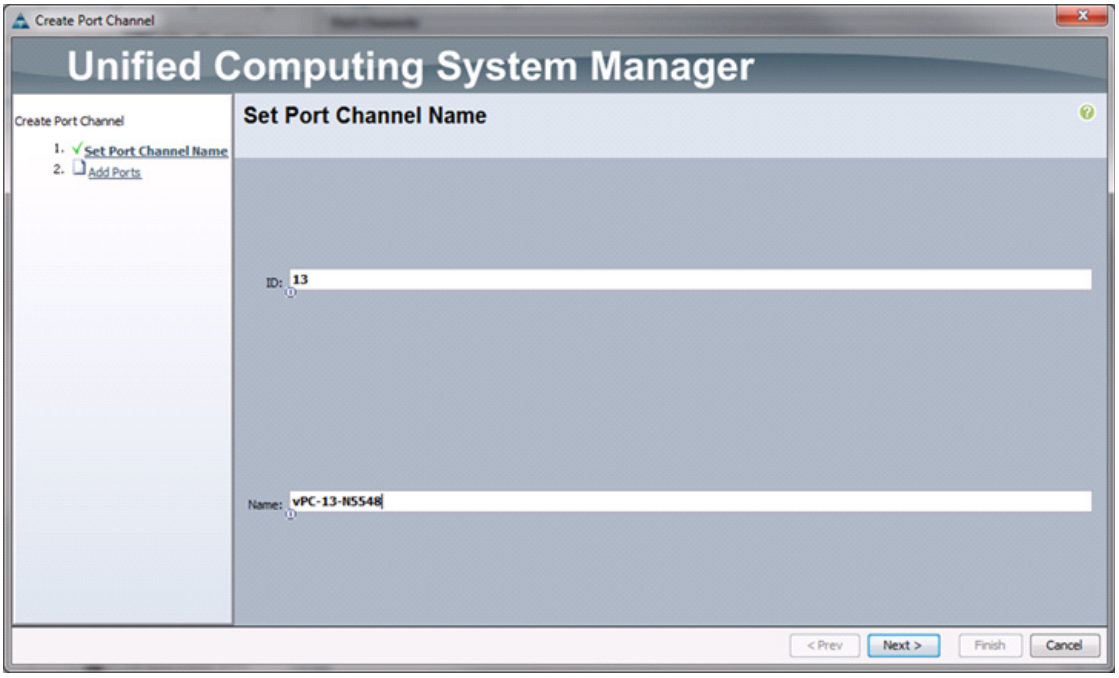

8. Choose the following ports to be added to the port channel:
   - Slot ID 1 and port 19
   - Slot ID 1 and port 20
9. Click **>>** to add the ports to the port channel.
10. Click **Finish** to create the port channel.
11. Click **OK**.
12. In the navigation pane, under **LAN** > **LAN Cloud**, expand the fabric B tree.
13. Right-click Port Channels.
14. Choose Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter vPC-14-N5548 as the name of the port channel.
17. Click **Next**.
18. Choose the following ports to be added to the port channel:
    - Slot ID 1 and port 19
    - Slot ID 1 and port 20
19. Click **>>** to add the ports to the port channel.
20. Click **Finish** to create the port channel.
21. Click **OK**.

Create an Organization

Organizations are used to organize resources and restrict access to various groups within the IT organization, thereby enabling multi-tenancy of the compute resources.

Note Although this document does not assume the use of organizations this procedure provides instructions for creating one.

To configure an organization in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, choose **Create Organization**.
2. Enter a name for the organization.
3. (Optional) Enter a description for the organization.
4. Click **OK**.
5. Click **OK** in the confirmation message window.

Create MAC Address Pools

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

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

Note In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Choose 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. Keep the Assignment Order at Default.
8. Click **Next**.
9. Click **Add**.
10. Specify a starting MAC address.

Note 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.

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

Figure 7 Creating MAC Address Pool for Fabric A

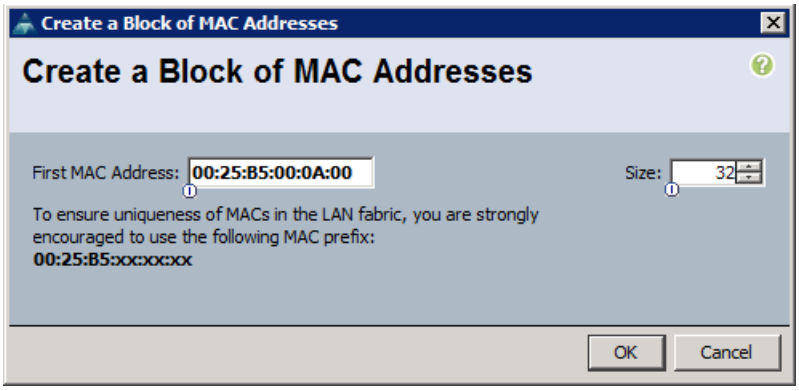

12. Click **OK**.
13. Click **Finish**.
14. In the confirmation message window, click **OK**.
15. Right-click MAC Pools under the root organization.
16. Choose Create MAC Pool to create the MAC address pool.
17. Enter MAC_Pool_B as the name of the MAC pool.
18. (Optional) Enter a description for the MAC pool.
19. Keep the Assignment Order at Default.
20. Click **Next**.
21. Click **Add**.
22. Specify a starting MAC address.

Note For the FlexPod solution, the recommendation is 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.

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

Figure 8 Creating MAC Address Pool for Fabric B

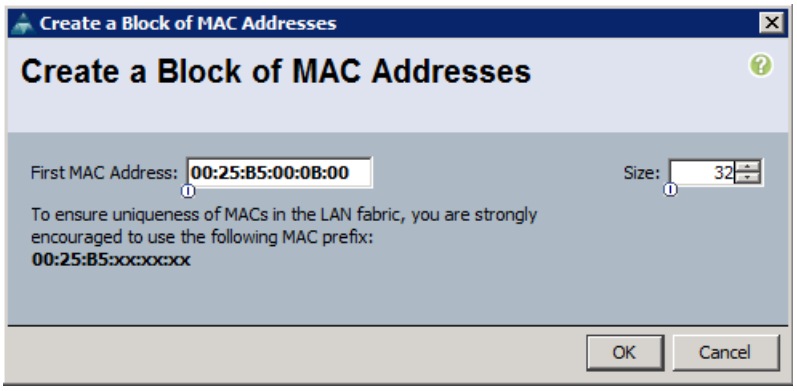

24. Click **OK**.

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

Create WWNN Pools

To configure the necessary World Wide Node Name (WWNN) pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Choose **Pools** > **root**.
3. Right-click WWNN Pools.
4. Choose Create WWNN Pool.
5. Enter WWNN_Pool as the name of the WWNN pool.
6. (Optional) Add a description for the WWNN pool.
7. Keep the Assignment Order at Default.
8. Click **Next**.
9. Click **Add** to add a block of WWNNs.
10. Keep the default block of WWNNs, or specify a base WWNN.
11. Specify a size for the WWNN block that is sufficient to support the available blade or server resources.

Figure 9 Creating WWNN Pool

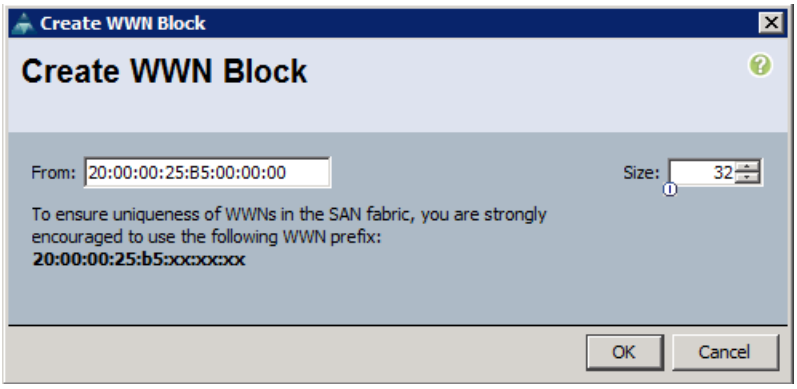

12. Click **OK**.
13. Click **Finish**.
14. Click **OK**.

Create WWPN Pools

To configure the necessary World Wide Port Name (WWPN) pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Choose **Pools** > **root**.

Note In this procedure, two WWPN pools are created: one for fabric A and one for fabric B.

3. Right-click WWPN Pools.
4. Choose Create WWPN Pool.
5. Enter WWPN_Pool_A as the name of the WWPN pool for fabric A.
6. (Optional) Enter a description for this WWPN pool.
7. Keep the Assignment Order at Default.
8. Click **Next**.
9. Click **Add** to add a block of WWPNs.
10. Specify the starting WWPN in the block for fabric A.

Note For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting WWPN to identify all the WWPNs in this pool as fabric A addresses.

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

Figure 10 Creating WWPN Pool

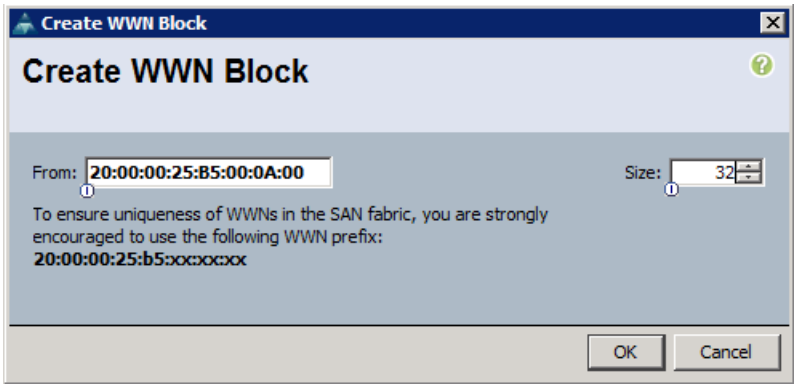

12. Click **OK**.
13. Click **Finish** to create the WWPN pool.
14. Click **OK**.
15. Right-click WWPN Pools.
16. Choose Create WWPN Pool.
17. Enter WWPN_Pool_B as the name for the WWPN pool for fabric B.
18. (Optional) Enter a description for this WWPN pool.
19. Keep the Assignment Order at Default.
20. Click **Next**.
21. Click **Add** to add a block of WWPNs.
22. Enter the starting WWPN address in the block for fabric B.

Note For the FlexPod solution, the recommendation is to place 0B in the next to last octet of the starting WWPN to identify all the WWPNs in this pool as fabric B addresses.

23. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.
24. Click **OK**.
25. Click **Finish**.
26. Click **OK**.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Pools** > **root**.
3. Right-click UUID Suffix Pools.
4. Choose 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. Keep the Assignment Order at Default.
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.

Figure 11 Creating UUID Suffix Pool

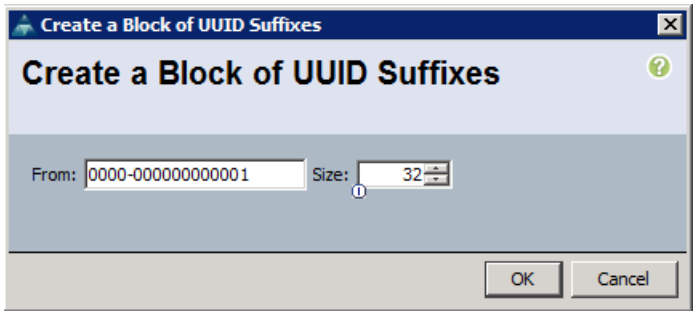

13. Click **OK**.
14. Click **Finish**.
15. Click **OK**.

Create Server Pool

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

Note 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. Choose **Pools** > **root**.
3. Right-click Server Pools.
4. Choose 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. Choose two servers to be used for the VMware management cluster and click **>>** to add them to the Infra_Pool server pool.
9. Click **Finish**.
10. Click **OK**.

Create VLANs

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

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

Note In this procedure, five VLANs are created.

2. Choose **LAN** > **LAN Cloud**.
3. Right-click VLANs.
4. Choose Create VLANs.
5. Enter IB-MGMT-VLAN as the name of the VLAN to be used for management traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter <<var_ib-mgmt_vlan_id>> as the ID of the management VLAN.
8. Keep the Sharing Type as None.
9. Click **OK**, and then click **OK** again.

Figure 12 Creating VLAN for Management Traffic

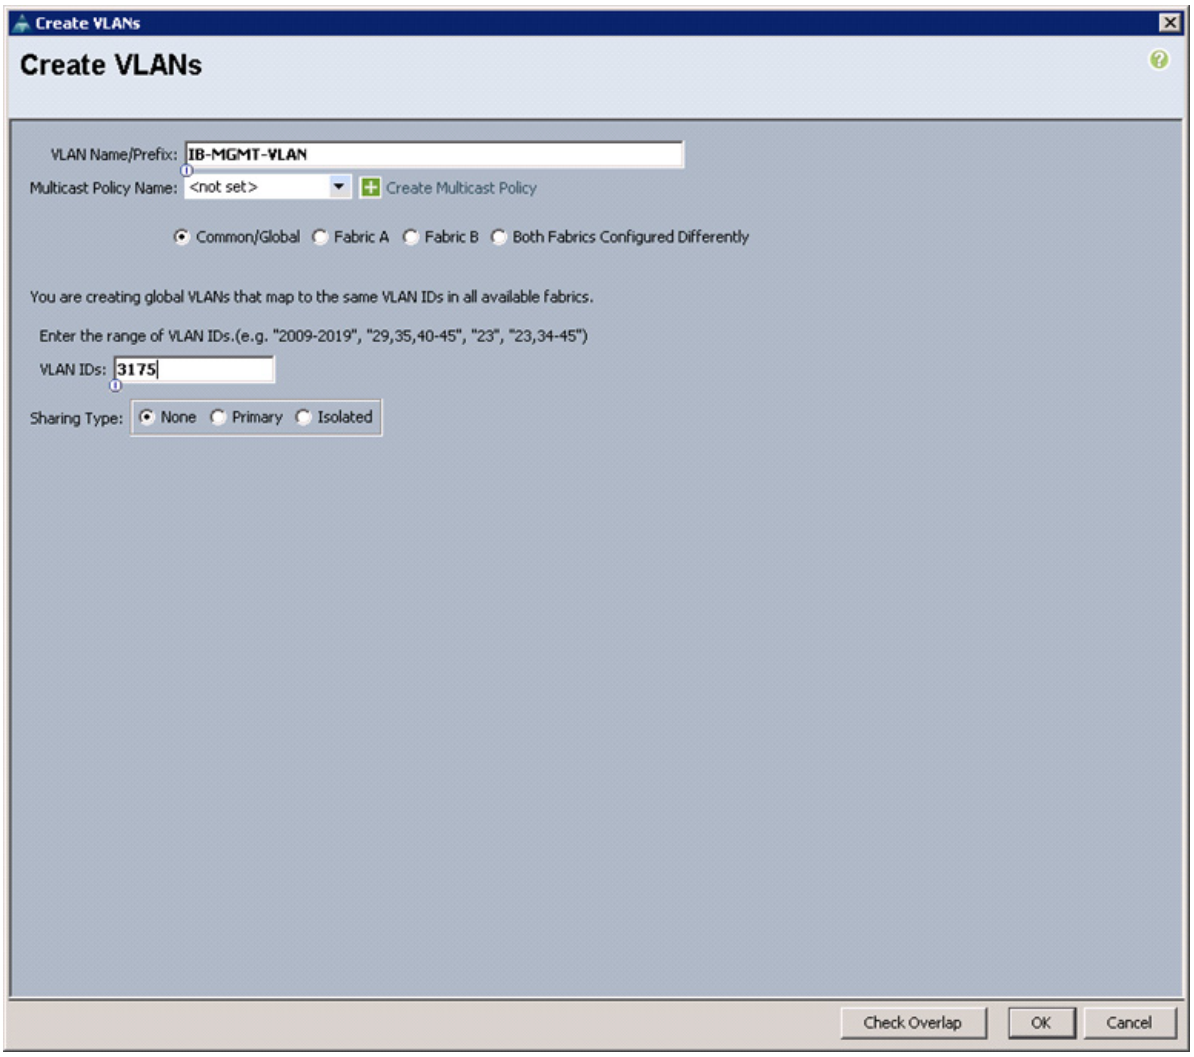

10. Right-click VLANs.
11. Choose Create VLANs.
12. Enter NFS-VLAN as the name of the VLAN to be used for NFS.
13. Keep the Common/Global option selected for the scope of the VLAN.
14. Enter the <<var_nfs_vlan_id>> for the NFS VLAN.
15. Keep the Sharing Type as None.
16. Click **OK**, and then click **OK** again.

Figure 13 Creating VLAN for NFS Traffic

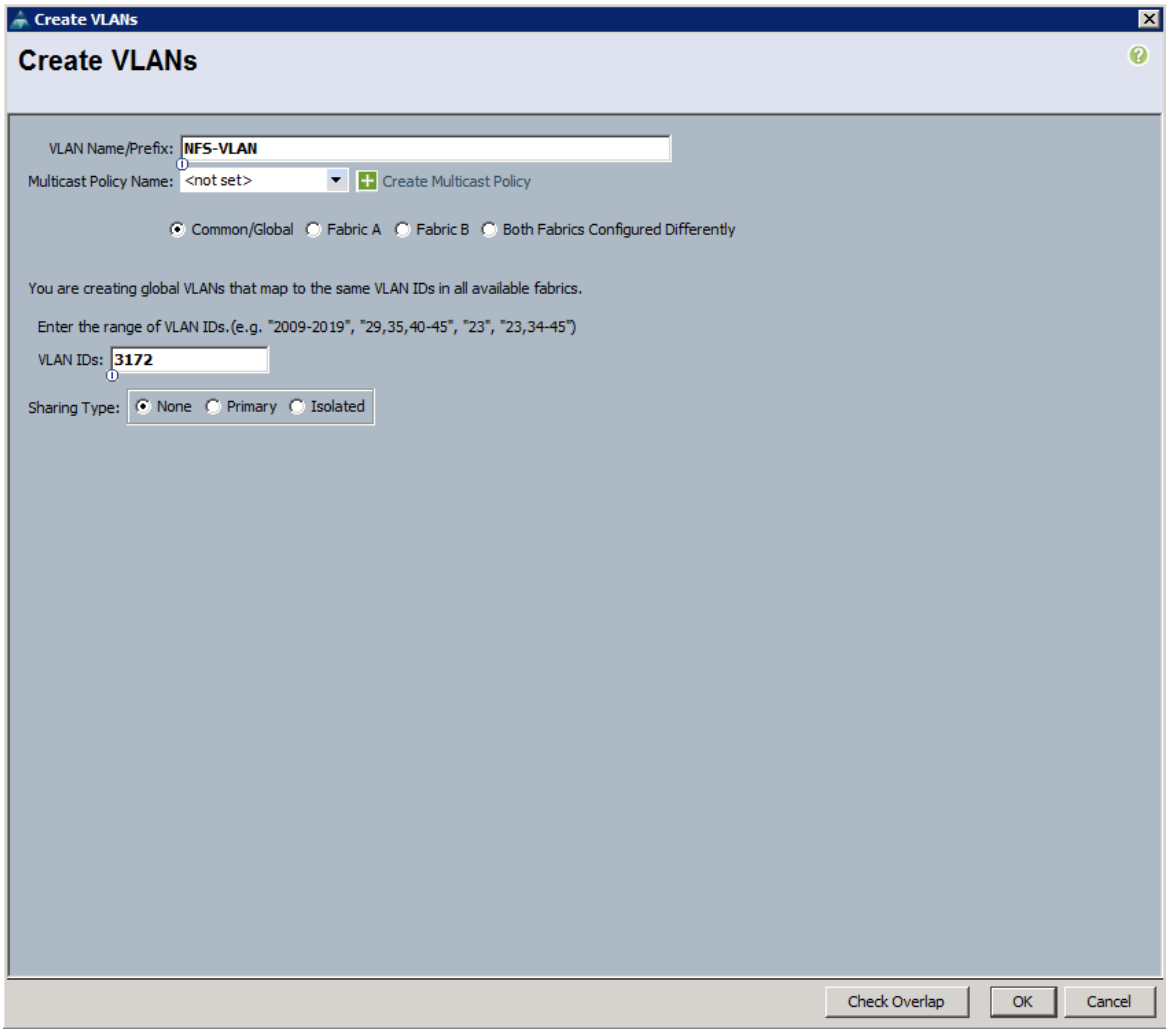

17. Right-click VLANs.
18. Choose Create VLANs.
19. Enter vMotion-VLAN as the name of the VLAN to be used for vMotion.
20. Keep the Common/Global option selected for the scope of the VLAN.
21. Enter the <<var_vmotion_vlan_id>> as the ID of the vMotion VLAN.
22. Keep the Sharing Type as None.
23. Click **OK**, and then click **OK** again.

Figure 14 Creating VLAN for vMotion

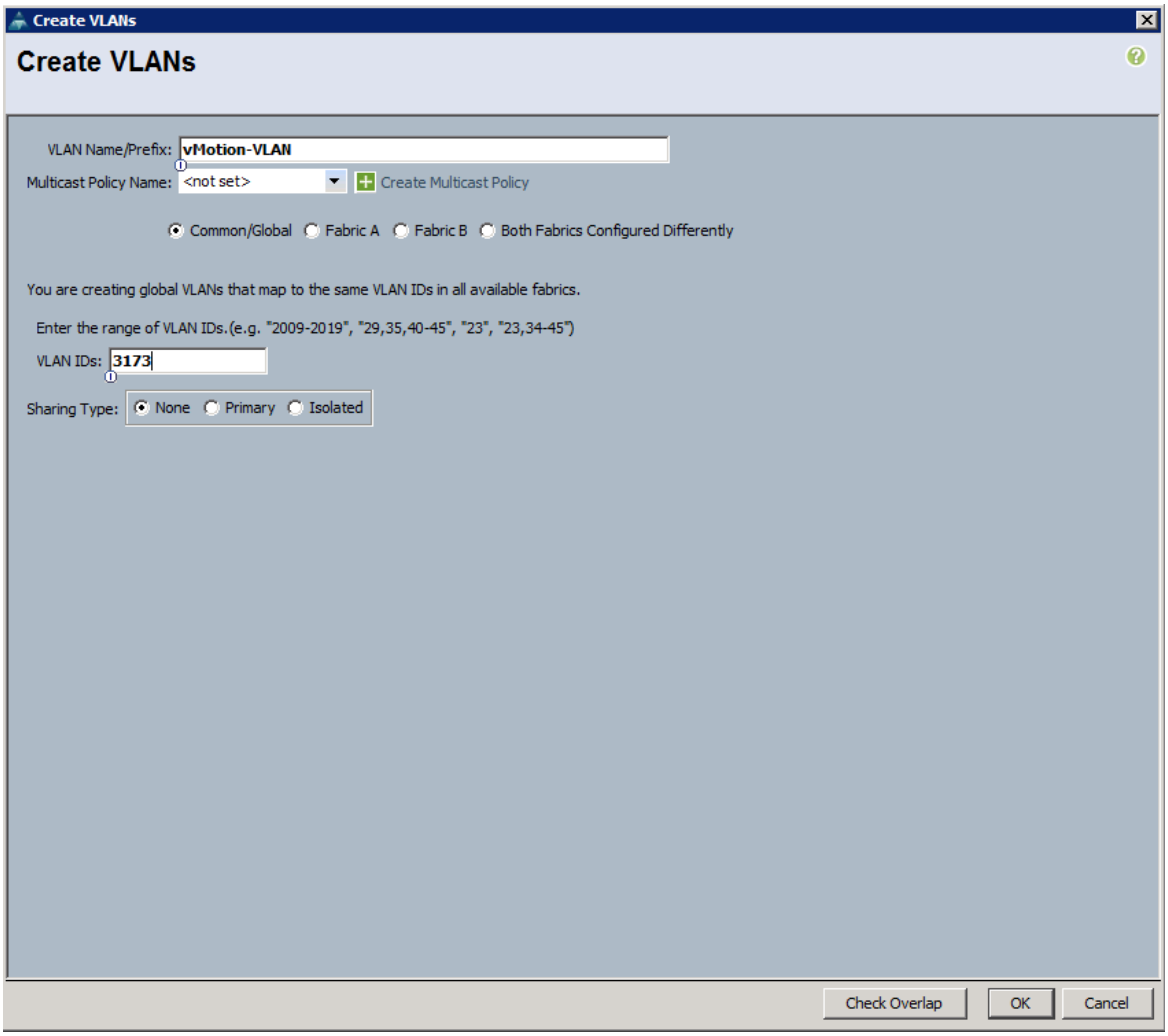

24. Right-click VLANs.
25. Choose Create VLANs.
26. Enter VM-Traffic-VLAN as the name of the VLAN to be used for the VM traffic.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the <<var_vm-traffic_vlan_id>> for the VM Traffic VLAN.
29. Keep the Sharing Type as None.
30. Click **OK**, and then click **OK** again.

Figure 15 Creating VLAN for VM Traffic

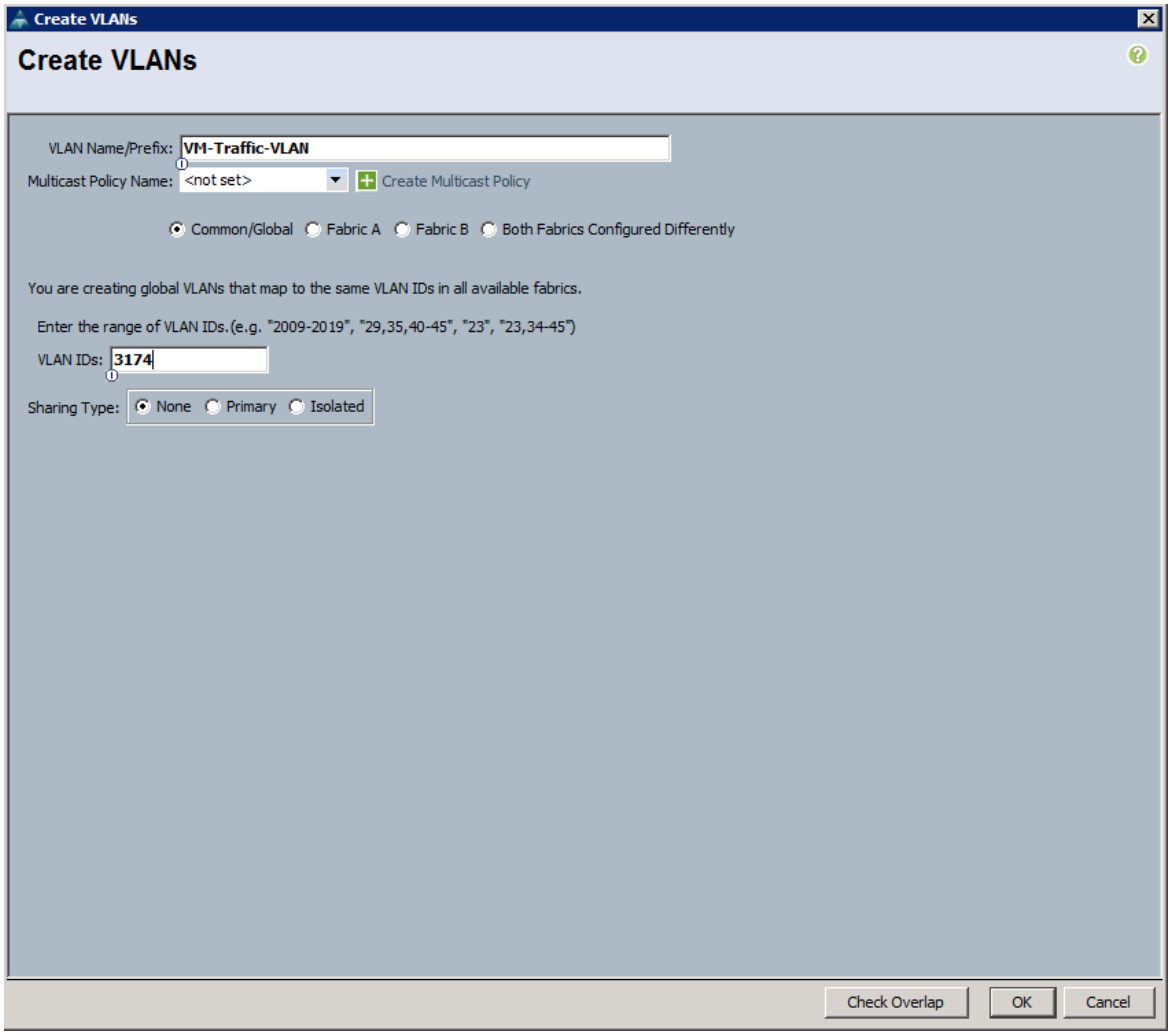

31. Right-click VLANs.
32. Choose Create VLANs.
33. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the <<var_native_vlan_id>> as the ID of the native VLAN.
36. Keep the Sharing Type as None.
37. Click **OK**, and then click **OK** again.

Figure 16 Creating Native VLAN

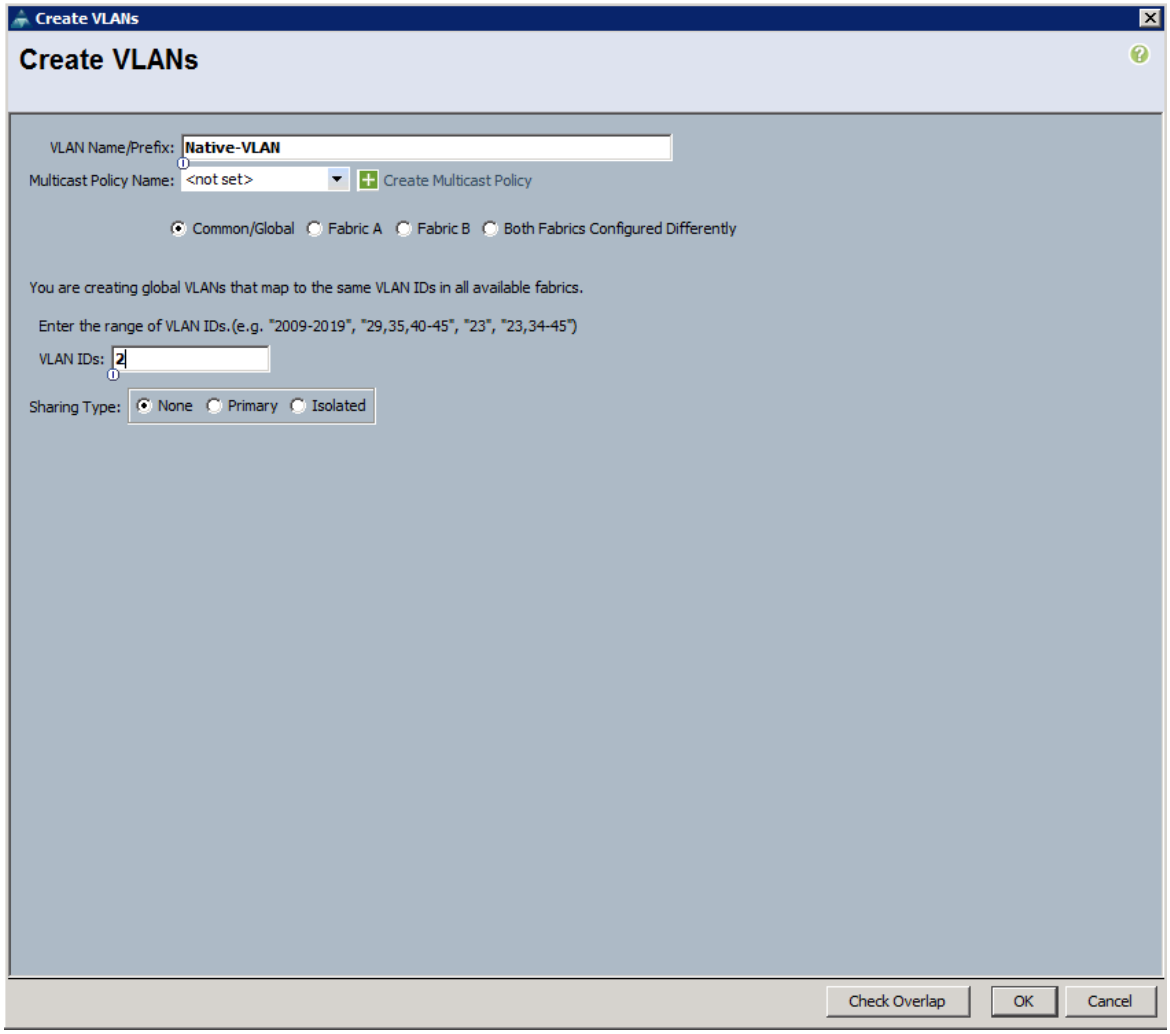

38. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and choose Set as Native VLAN.
39. Click **Yes**, and then click **OK**.

Create VSANs and FCoE Port Channels

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

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Expand the **SAN** > **SAN Cloud tree**.
3. Right-click VSANs.
4. Choose Create VSAN.
5. Enter VSAN_A as the name of the VSAN for fabric A.
6. Keep the Disabled option selected for FC Zoning.

7. Click the **Fabric A** radio button.
8. Enter <<var_vsan_a_id>> as the VSAN ID for fabric A.
9. Enter <<var_fabric_a_fcoe_vlan_id>> as the FCoE VLAN ID for fabric A.

Note For the FlexPod solution, it is recommended to use the same ID for the VSAN and the FCoE VLAN required for fabric A.

10. Click **OK**, and then click **OK** again to create the VSAN.

Figure 17 Creating VSAN for Fabric A

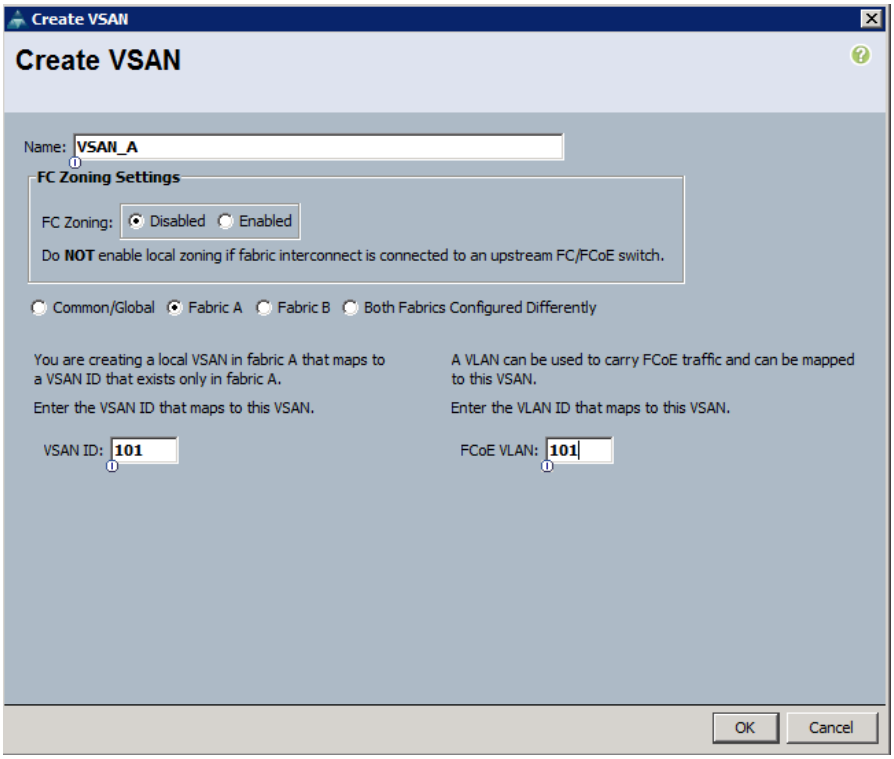

11. Right-click VSANs.
12. Choose Create VSAN.
13. Enter VSAN_B as the name of the VSAN for fabric B.
14. Keep the Disabled option selected for FC Zoning.
15. Click the **Fabric B** radio button.
16. Enter <<var_vsan_b_id>> as the VSAN ID for fabric B.
17. Enter <<var_fabric_b_fcoe_vlan_id>> as the FCoE VLAN ID for fabric B.

Note It is recommended to use the same ID for the VSAN and the FCoE VLAN required for fabric B.

18. Click **OK**, and then click **OK** again to create the VSAN.

Figure 18 Creating VSAN for Fabric B

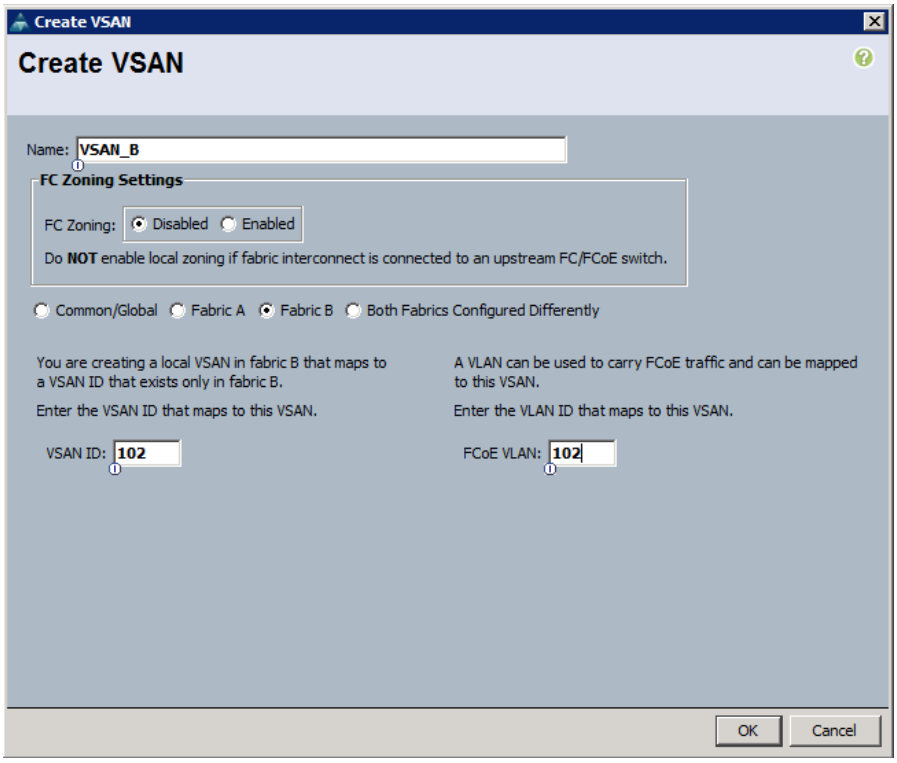

19. In the navigation pane, under **SAN** > **SAN Cloud**, expand the Fabric A tree.
20. Right-click FCoE Port Channels.
21. Choose Create FCoE Port Channel.
22. Enter 1 for the port channel ID and Po1 for the port channel name.
23. Click **Next**.
24. Choose ports 31 and 32 and click **>>** to add the ports to the port channel.
25. Click **Finish**.
26. check the check box for Show Navigator for FCoE Port-Channel 1 (Fabric A).
27. Click **OK** to create the port channel.
28. In the right pane, under Properties, choose VSAN VSAN_A for Fabric A in the VSAN list.
29. Click **Apply**, and then click **OK**.
30. Click **OK** to close the navigator.
31. In the navigation pane, under **SAN** > **SAN Cloud**, expand the fabric B tree.
32. Right-click FCoE Port Channels.
33. Choose Create FCoE Port Channel.
34. Enter 2 for the port channel ID and Po2 for the port channel name.
35. Click **Next**.
36. Choose ports 31 and 32 and click >> to add the ports to the port channel.
37. Click **Finish**.

38. Check the check box for Show Navigator for FCoE Port-Channel 2 (Fabric B).
39. Click **OK** to create the port channel.
40. In the right pane, under Properties, choose VSAN VSAN_B for Fabric B.
41. Click **Apply**, and then click **OK**.
42. Click **OK** to close the navigator.

Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click Host Firmware Packages.
4. Choose Create Host Firmware Package.
5. Enter VM-Host-Infra as the name of the host firmware package.
6. Keep the radio button Simple selected.
7. Choose the version 2.1(3a) for both the Blade and Rack Packages.
8. Click **OK** to create the host firmware package.
9. Click **OK**.

Figure 19 Creating Host Firmware Package

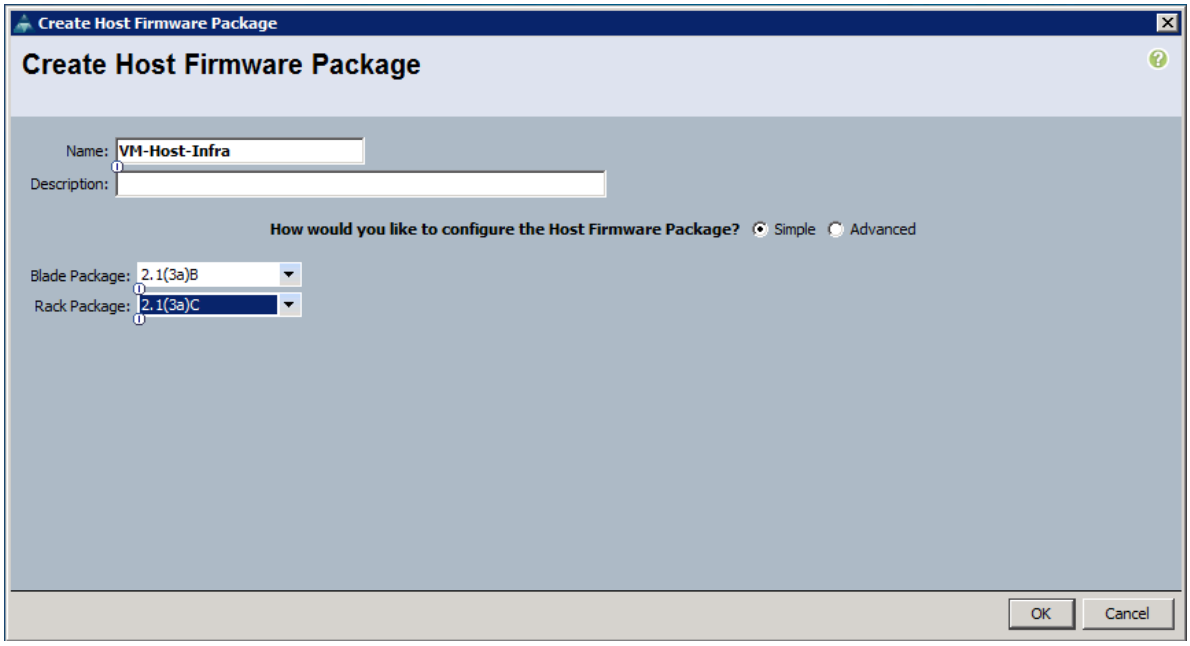

Set Jumbo Frames in Cisco UCS Fabric

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

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Choose **LAN > LAN Cloud > QoS System Class**.
3. In the right pane, click the **General** tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click **Save Changes**.
6. Click **OK**.

Figure 20 Setting Jumbo Frame

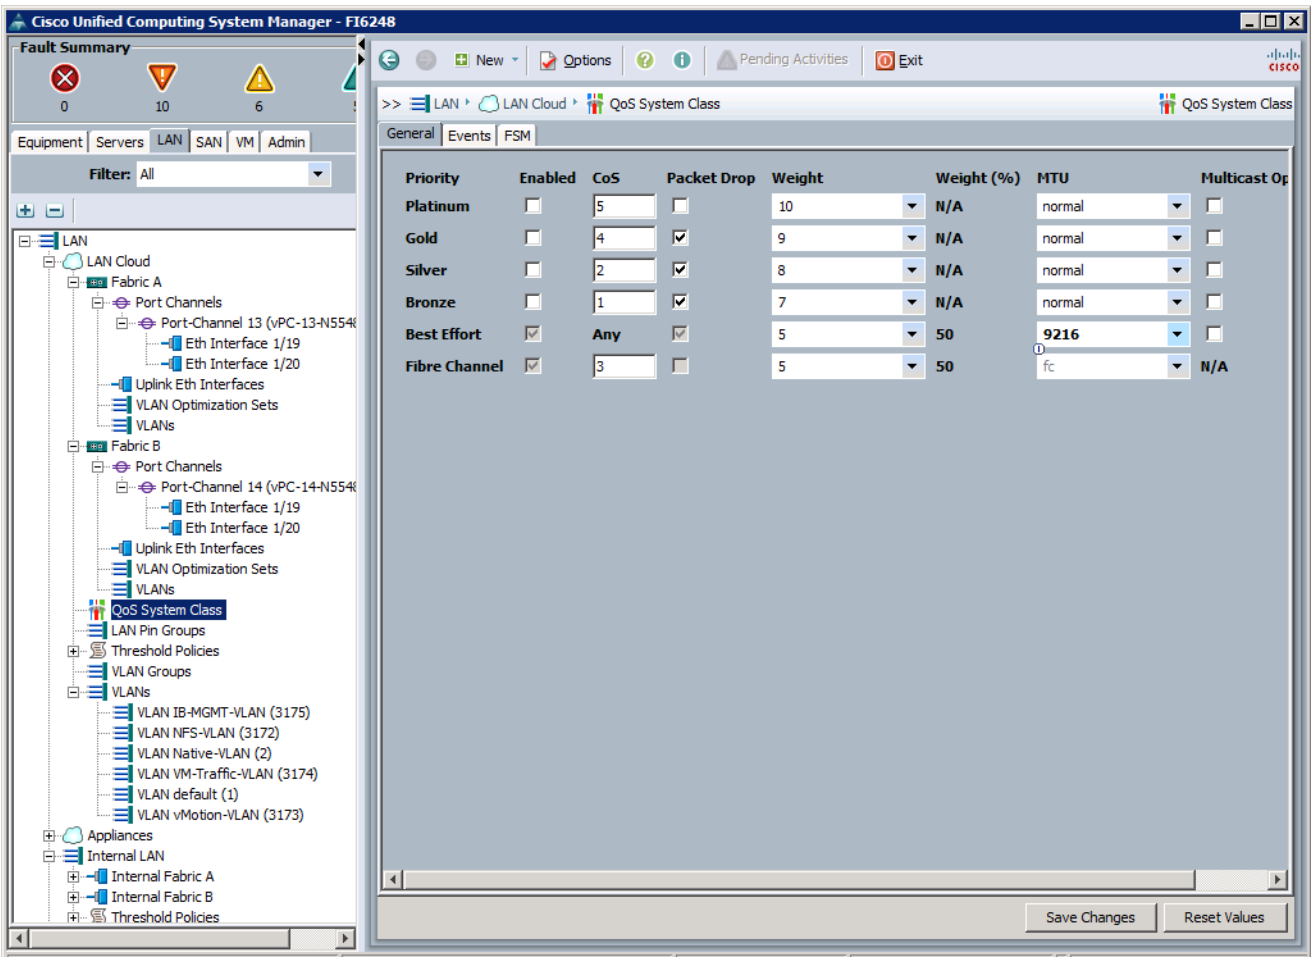

Create Local Disk Configuration Policy (Optional)

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

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

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click Local Disk Config Policies.
4. Choose 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. Keep the Flex Flash State at Disable.
8. Click **OK** to create the local disk configuration policy.

Figure 21 Creating Local Disk Configuration Policy

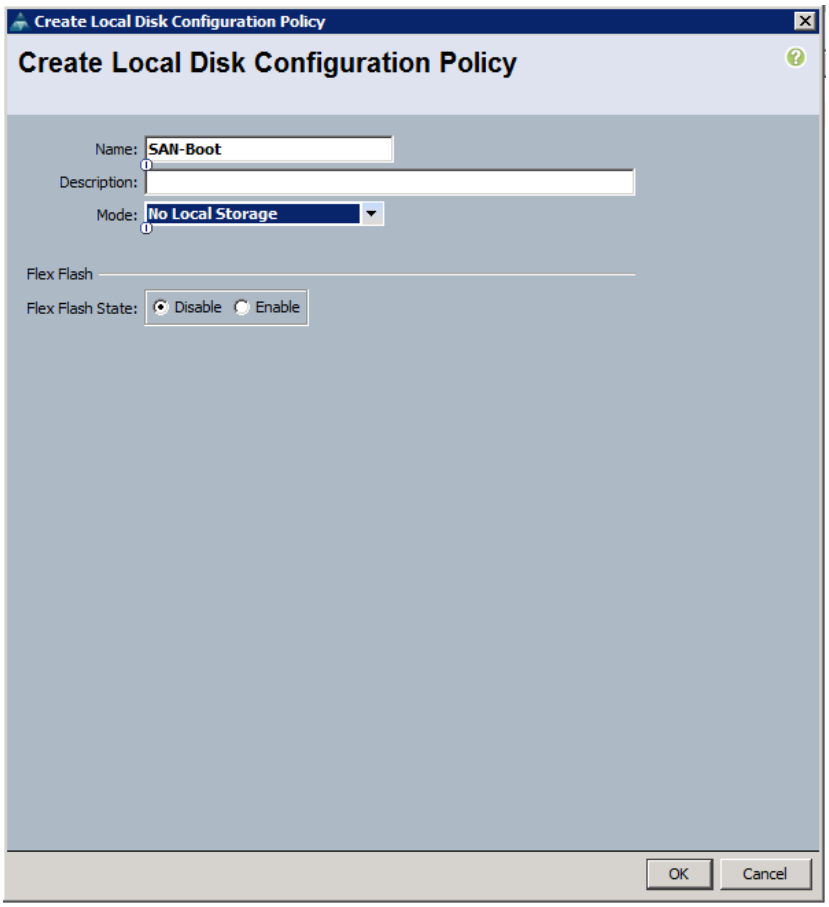

9. Click **OK**.

Create Network Control Policy for Cisco Discovery Protocol

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

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

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

Figure 22 Creating Network Control Policy

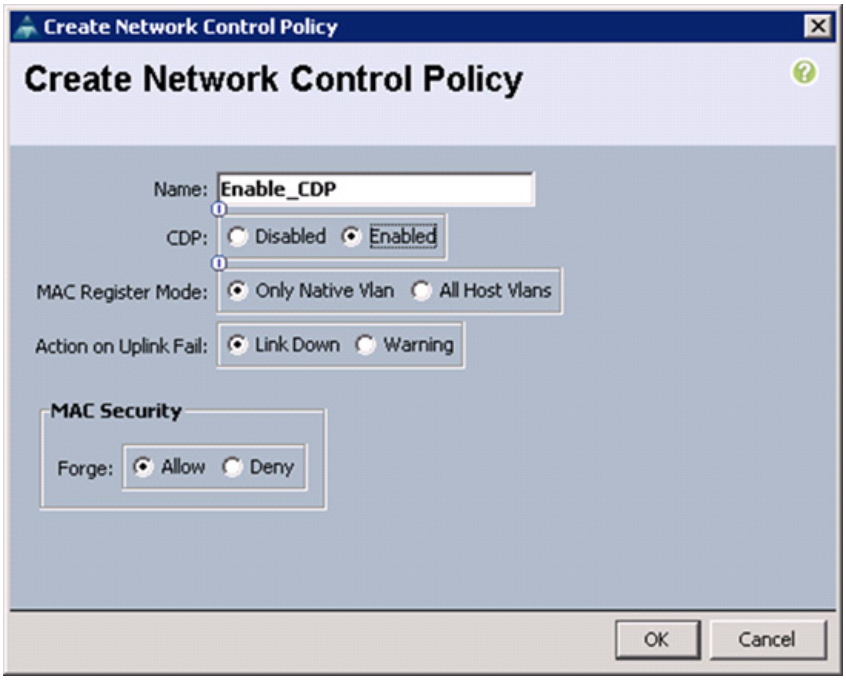

8. Click **OK**.

Create Power Control Policy

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click Power Control Policies.
4. Choose 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**.

Figure 23 Creating Power Control Policy

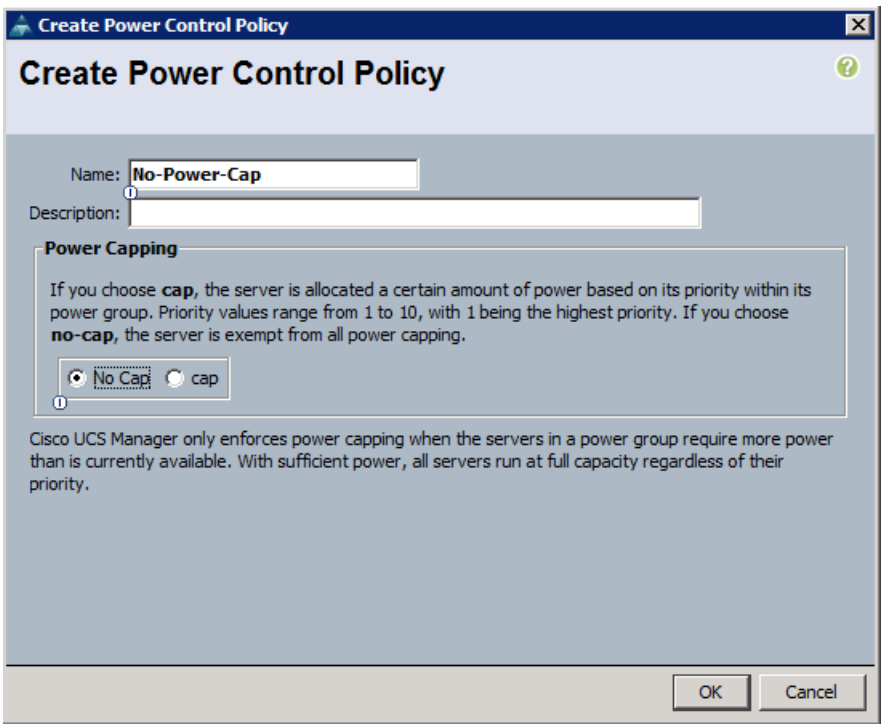

Create Server Pool Qualification Policy (Optional)

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

Note This example creates a policy for a B200-M3 server.

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click Server Pool Policy Qualifications.
4. Choose Create Server Pool Policy Qualification.
5. Enter UCSB-B200-M3 as the name for the policy.
6. Choose Create Server PID Qualifications.
7. Enter UCSB-B200-M3 as the PID.
8. Click **OK** to create the server pool qualification policy.
9. Click **OK**, and then click **OK** again.

Figure 24 Creating Server PID Qualifications

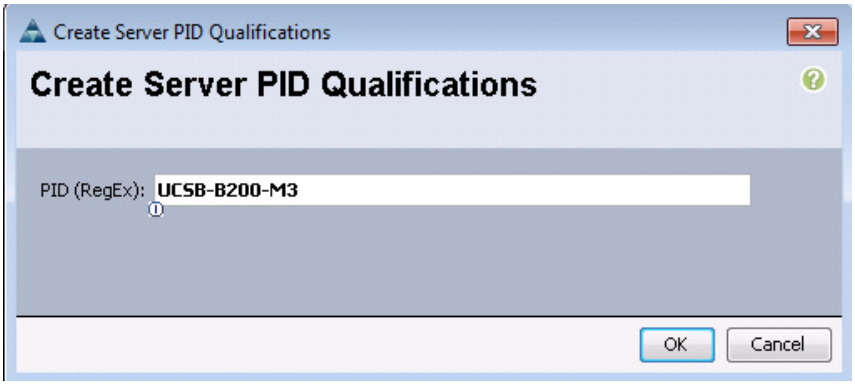

Create Server BIOS Policy

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click BIOS Policies.
4. Choose Create BIOS Policy.
5. Enter VM-Host-Infra as the BIOS policy name.
6. Change the Quiet Boot setting to Disabled.
7. Click **Finish** to create the BIOS policy.

Figure 25 Creating BIOS Policy

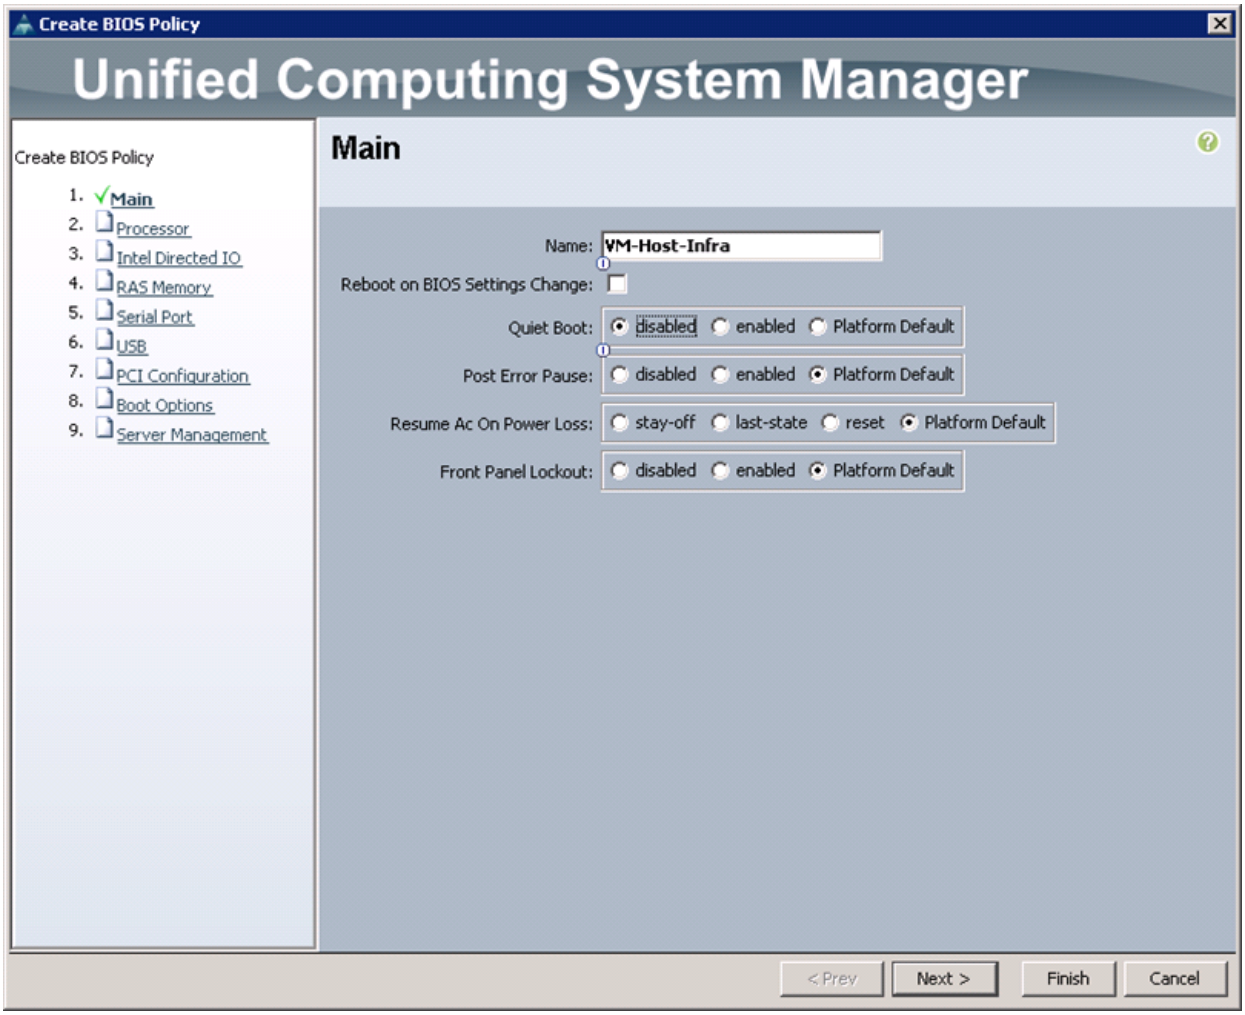

8. Click **OK**.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

To create a vNIC/vHBA placement policy for the infrastructure hosts, follow these steps:

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

Figure 26 Creating Placement Policy

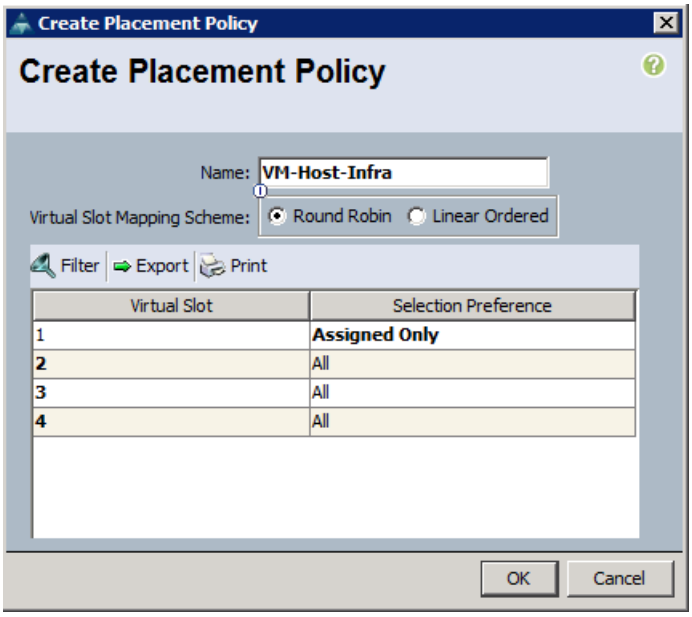

Update default Maintenance Policy

To update the default Maintenance Policy, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Choose **Maintenance Policies** > **default**.
4. Change the Reboot Policy to User Ack.
5. Click **Save Changes**.
6. Click **OK** to accept the change.

Figure 27 Updating Maintenance Policy

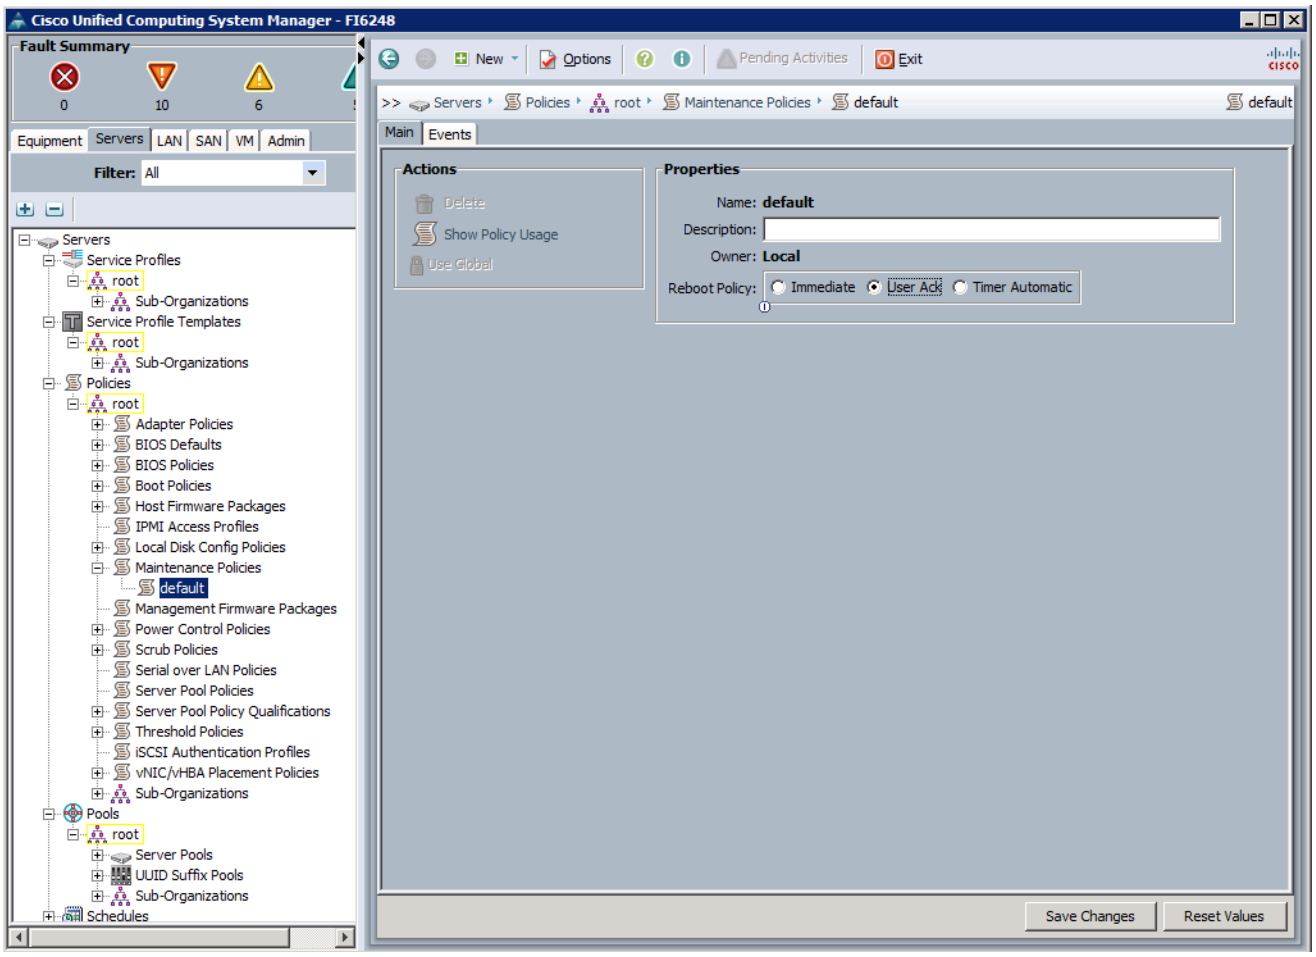

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Choose **Policies > root**.
3. Right-click vNIC Templates.
4. Choose **Create vNIC Template**.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep the radio button **Fabric A** selected.
7. Do not check the Enable Failover check box.
8. Under Target, make sure that the VM check box is not checked.
9. Click the **Updating Template** radio button as the Template Type.
10. Under VLANs, check the check boxes for IB-MGMT-VLAN, NFS-VLAN, Native-VLAN, VM-Traffic-VLAN, and vMotion-VLAN.

11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, Choose MAC_Pool_A.
14. In the Network Control Policy list, Choose Enable_CDP.
15. Click **OK** to create the vNIC template.
16. Click **OK**.

Figure 28 Creating vNIC Template for Fabric A

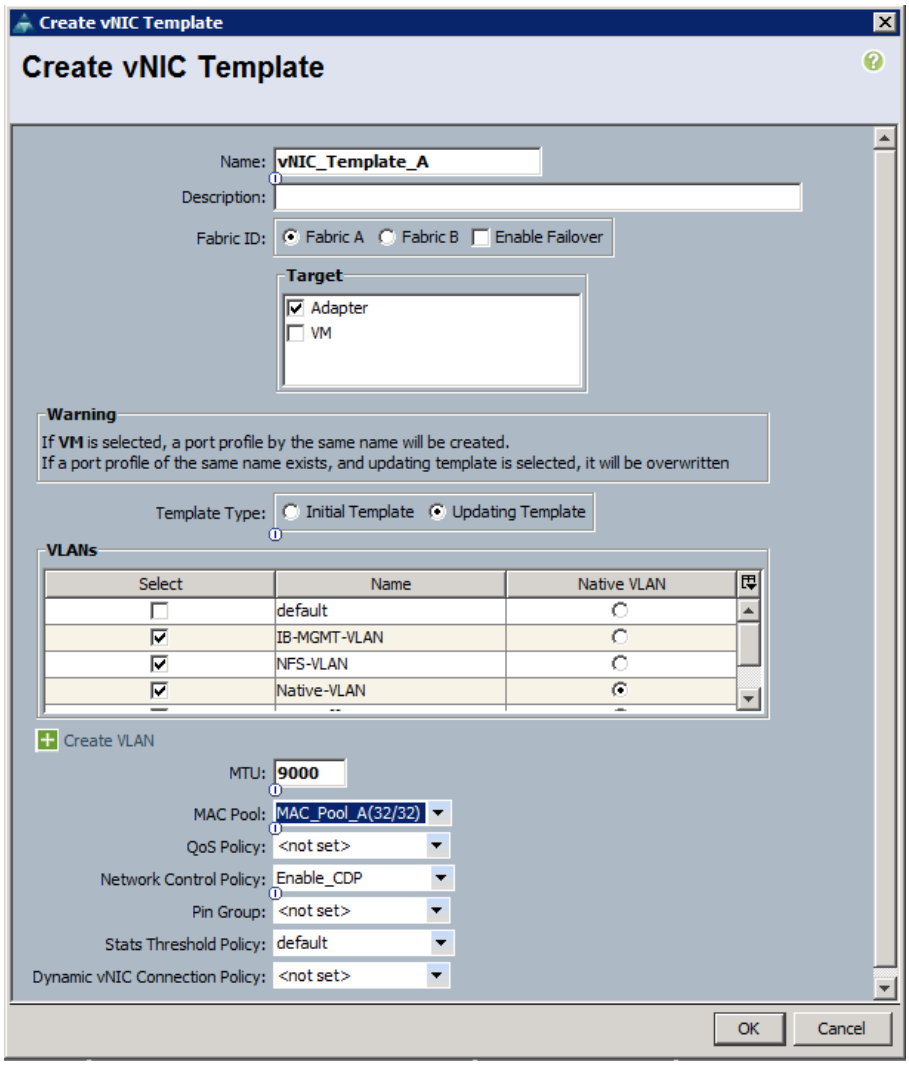

17. In the navigation pane, click the **LAN** tab.
18. Choose **Policies** > **root**.
19. Right-click vNIC Templates.
20. Choose **Create vNIC Template**.
21. Enter vNIC_Template_B as the vNIC template name.

22. Click the radio button **Fabric B**.
23. Do not check the Enable Failover check box.
24. Under Target, make sure the VM check box is not checked.
25. Click the **Updating Template** radio button as the template type.
26. Under VLANs, check the check boxes for IB-MGMT-VLAN, NFS-VLAN, Native-VLAN, VM-Traffic-VLAN, and vMotion-VLAN.
27. Set Native-VLAN as the native VLAN.
28. For MTU, enter 9000.
29. In the MAC Pool list, Choose MAC_Pool_B.
30. In the Network Control Policy list, Choose Enable_CDP.
31. Click **OK** to create the vNIC template.
32. Click **OK**.

Figure 29 Creating vNIC Template for Fabric B

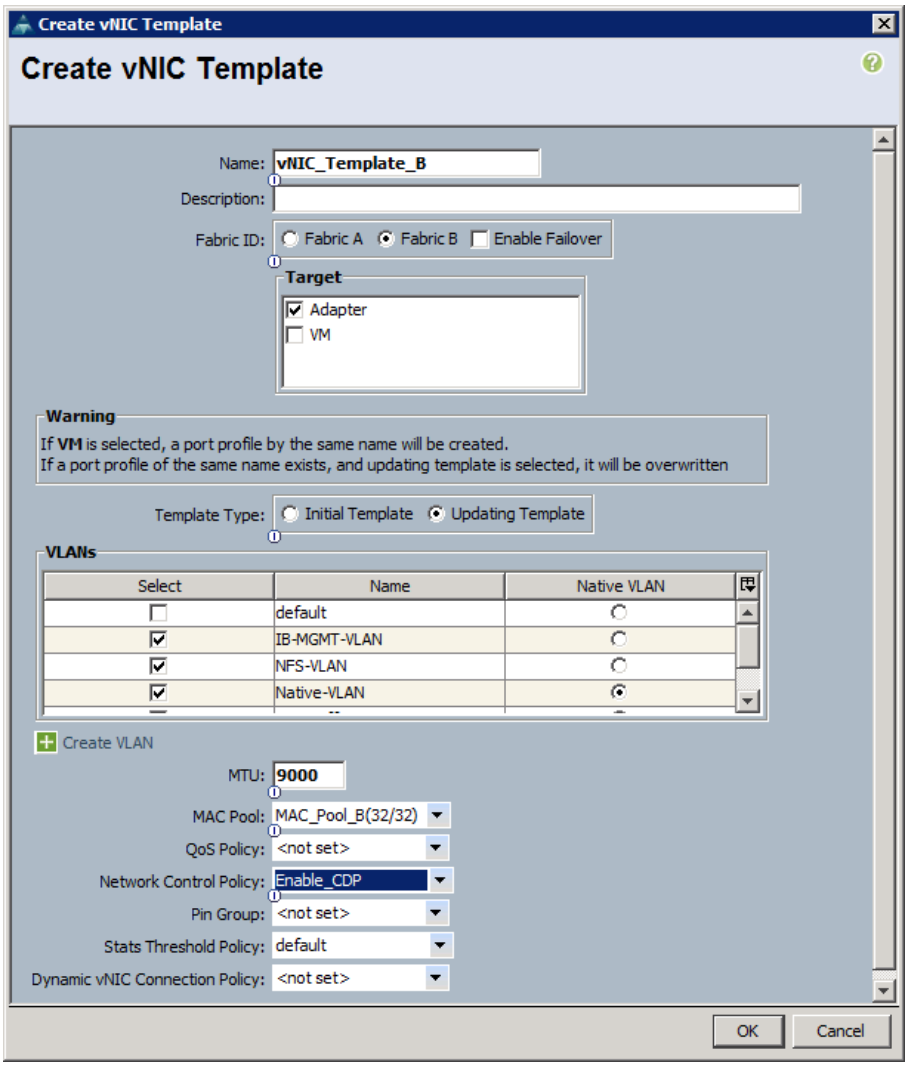

Create vHBA Templates for Fabric A and Fabric B

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

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click vHBA Templates.
4. Choose **Create vHBA Template**.
5. Enter vHBA_Template_A as the vHBA template name.
6. Click the radio button **Fabric A**.
7. In the Select VSAN list, Choose VSAN_A.
8. In the WWPN Pool list, Choose WWPN_Pool_A.
9. Click **OK** to create the vHBA template.
10. Click **OK**.

Figure 30 Creating vHBA Template for Fabric A

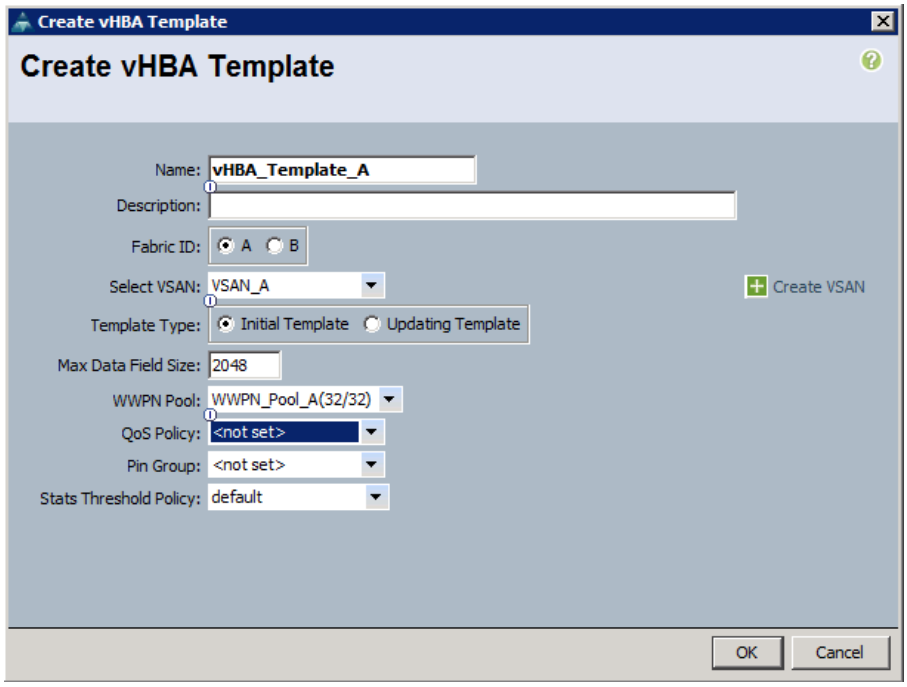

11. In the navigation pane, click the **SAN** tab.
12. Choose **Policies** > **root**.
13. Right-click vHBA Templates.
14. Choose **Create vHBA Template**.
15. Enter vHBA_Template_B as the vHBA template name.
16. Click the radio button **Fabric B**.
17. In the Select VSAN list, Choose VSAN_B.

18. In the WWPN Pool, Choose WWPN_Pool_B.
19. Click **OK** to create the vHBA template.
20. Click **OK**.

Figure 31 Creating vHBA Template for Fabric B

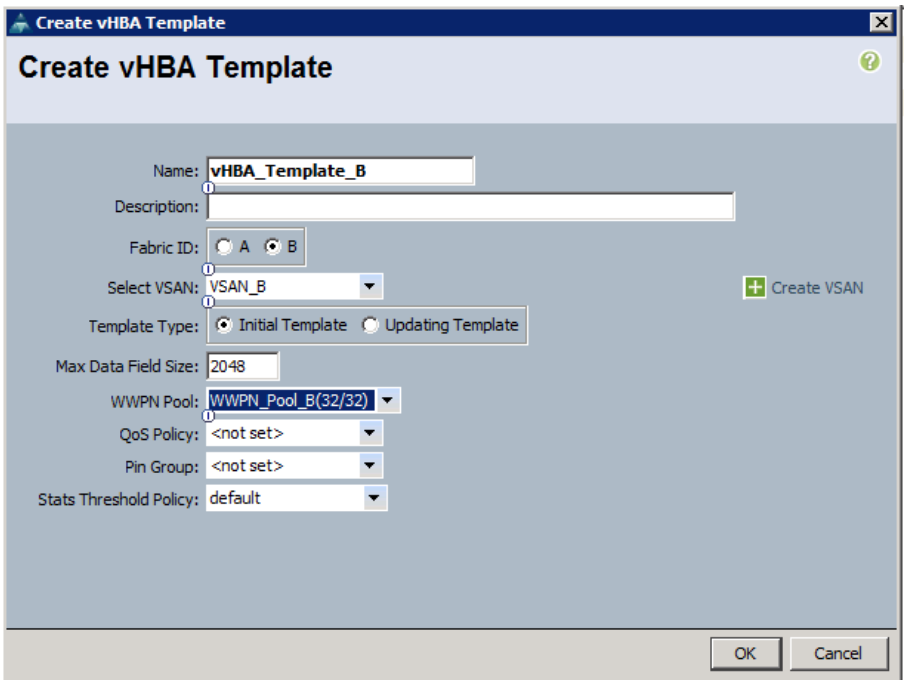

Create Boot Policies

This procedure applies to a Cisco UCS environment in which the storage FCoE ports are configured in the following ways:

- The FCoE ports 1a on storage controllers 1 and 2 are connected to the Cisco Nexus 5548 switch A.
- The FCoE ports 2a on storage controllers 1 and 2 are connected to the Cisco Nexus 5548 switch B.

Two boot policies are configured in this procedure:

- The first configures the primary target to be FCoE port 1a on storage controller 1.
- The second configures the primary target to be FCoE port 2a on storage controller 1.

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click Boot Policies.
4. Choose **Create Boot Policy**.
5. Enter Boot-Fabric-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 check box unchecked.

8. Expand the Local Devices drop-down menu and Choose Add CD-ROM.
9. Expand the vHBAs drop-down menu and Choose Add SAN Boot.
10. In the Add SAN Boot dialog box, enter Fabric-A in the vHBA field.
11. Make sure that the Primary radio button is selected as the SAN boot type.
12. Click **OK** to add the SAN boot initiator.

Figure 32 Adding SAN Boot Initiator for Fabric A

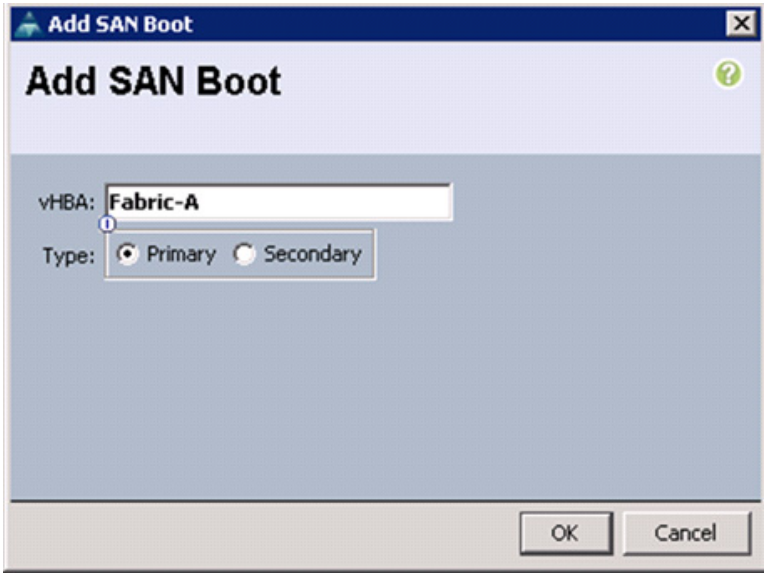

13. From the vHBA drop-down menu, choose Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for FCoE port 1a on storage controller 1.

Note To obtain this information, log in to storage controller 1 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

16. Keep the **Primary** radio button selected as the SAN boot target type.
17. Click **OK** to add the SAN boot target.

Figure 33 Adding SAN Boot Target for Fabric A

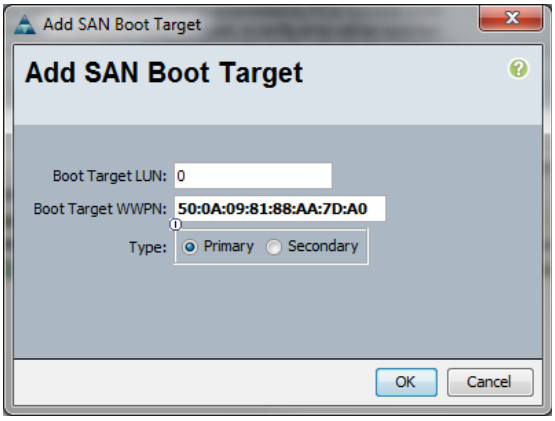

18. From the vHBA drop-down menu, choose Add SAN Boot Target.
19. Keep 0 as the value for Boot Target LUN.
20. Enter the WWPN for FCoE port 1a on storage controller 2.

Note To obtain this information, log in to storage controller 2 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

21. Click **OK** to add the SAN boot target.

Figure 34 Adding Secondary SAN Boot Target for Fabric A

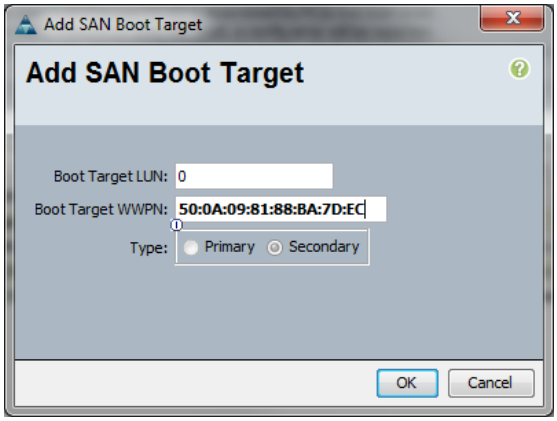

22. From the vHBA drop-down menu, choose Add SAN Boot.
23. In the Add SAN Boot dialog box, enter Fabric-B in the vHBA box.
24. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.
25. Click **OK** to add the SAN boot initiator.

Figure 35 Adding SAN Boot Initiator for Fabric B

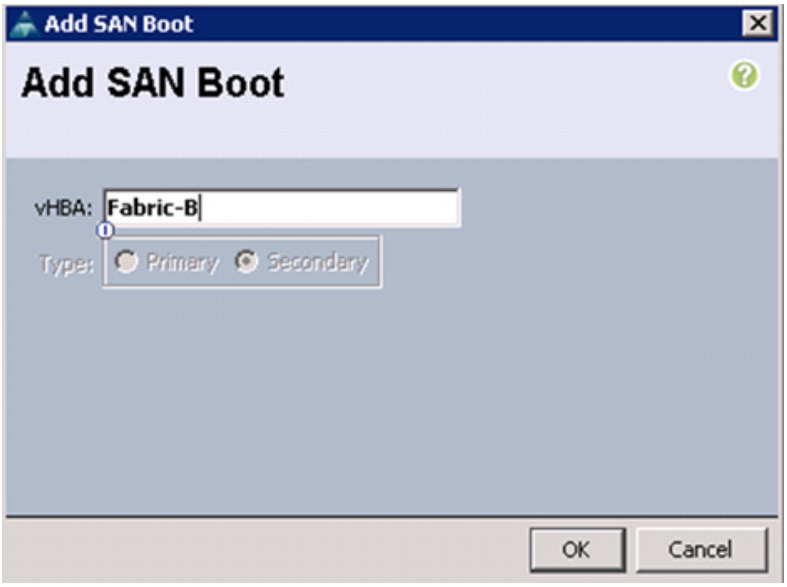

26. From the vHBA drop-down menu, choose Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN for FCoE port 2a on storage controller 1.

Note To obtain this information, log in to storage controller 1 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

29. Keep Primary as the SAN boot target type.
30. Click **OK** to add the SAN boot target.

Figure 36 Adding Primary SAN Boot Target for Fabric B

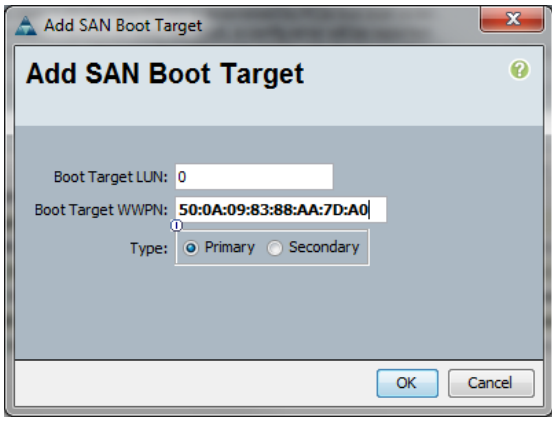

31. From the vHBA drop-down menu, choose Add SAN Boot Target.
32. Keep 0 as the value for Boot Target LUN.

33. Enter the WWPN for FCoE port 2a on storage controller 2.

Note To obtain this information, log in to storage controller 2 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

34. Click **OK** to add the SAN boot target.

Figure 37 Adding Secondary SAN Boot Target

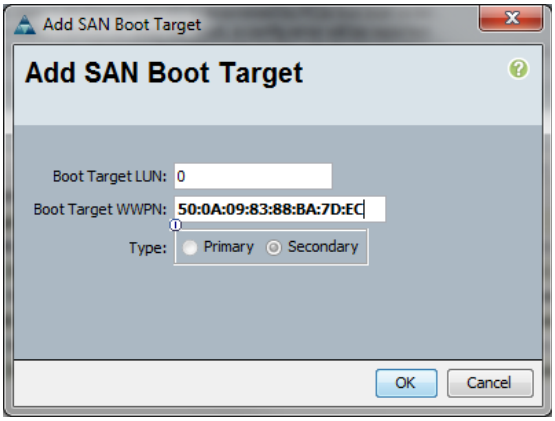

35. Click **OK**, and then **OK** again to create the boot policy.
36. Right-click Boot Policies again.
37. Choose **Create Boot Policy**.
38. Enter Boot-Fabric-B as the name of the boot policy.
39. (Optional) Enter a description of the boot policy.
40. Keep the Reboot on Boot Order Change check box unchecked.
41. From the Local Devices drop-down menu choose Add CD-ROM.
42. From the vHBA drop-down menu choose Add SAN Boot.
43. In the Add SAN Boot dialog box, enter Fabric-B in the vHBA box.
44. Make sure that the Primary radiobutton is selected as the SAN boot type.
45. Click **OK** to add the SAN boot initiator.

Figure 38 Adding SAN Boot Initiator for Fabric B

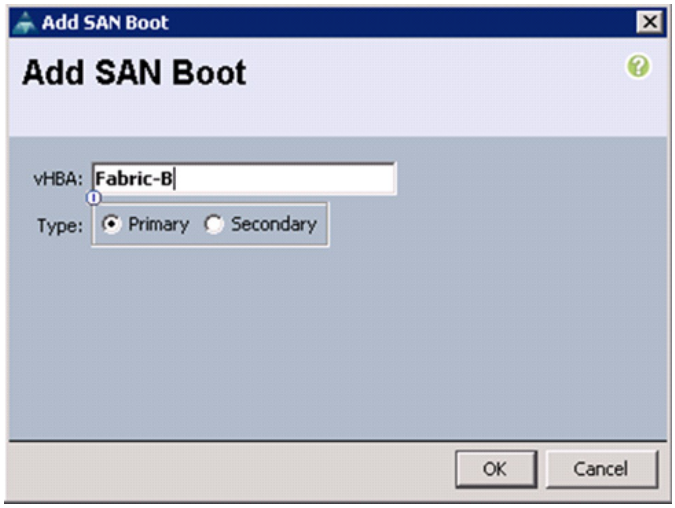

46. From the vHBA drop-down menu, choose Add SAN Boot Target.
47. Keep 0 as the value for Boot Target LUN.
48. Enter the WWPN for FCoE port 2a on storage controller 1.

Note To obtain this information, log in to storage controller 1 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

49. Keep Primary as the SAN boot target type.
50. Click **OK** to add the SAN boot target.

Figure 39 Adding Primary SAN Boot Target for Fabric B

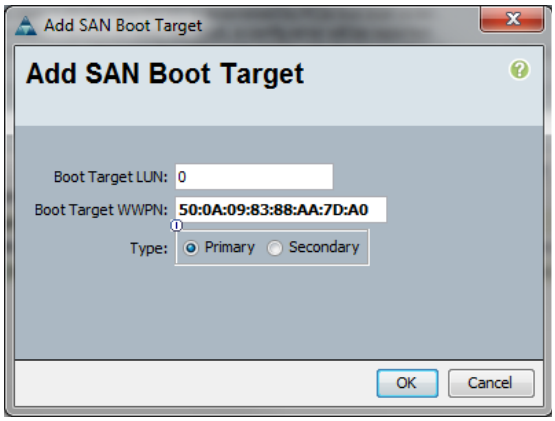

51. From the vHBA drop-down menu, choose Add SAN Boot Target.
52. Keep 0 as the value for Boot Target LUN.
53. Enter the WWPN for FCoE port 2a on storage controller 2.

Note To obtain this information, log in to storage controller 2 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

54. Click **OK** to add the SAN boot target.

Figure 40 Adding Secondary SAN Boot Target for Fabric B

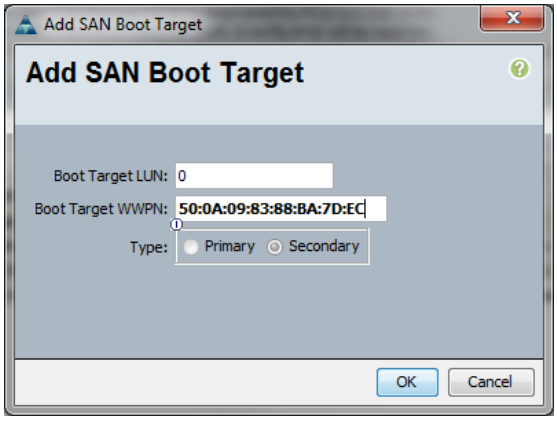

55. From the vHBA menu, choose Add SAN Boot.
56. In the Add SAN Boot dialog box, enter Fabric-A in the vHBA box.
57. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.
58. Click **OK** to add the SAN boot initiator.

Figure 41 Adding SAN Boot for Fabric A

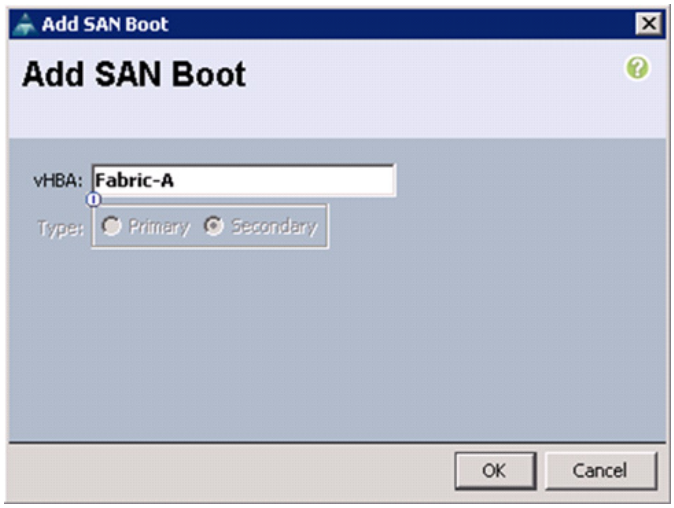

59. From the vHBA menu, choose Add SAN Boot Target.
60. Keep 0 as the value for Boot Target LUN.
61. Enter the WWPN for FCoE port 1a on storage controller 1.

Note To obtain this information, log in to storage controller 1 and run the fcp show adapters command. Make sure you enter the port name and not the node name.

62. Keep Primary as the SAN boot target type.
63. Click **OK** to add the SAN boot target.

Figure 42 Adding Primary SAN Boot Target for Fabric A

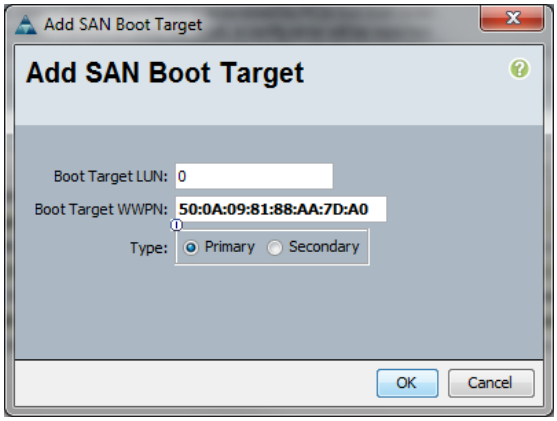

64. From the vHBA drop-down menu, choose Add SAN Boot Target.
65. Keep 0 as the value for Boot Target LUN.
66. Enter the WWPN for FCoE port 1a on storage controller 2.

Note To obtain this information, log in to storage controller 2 and run the **fcp show adapters** command. Make sure you enter the port name and not the node name.

67. Click **OK** to add the SAN boot target.

Figure 43 Adding Secondary SAN Boot Target for Fabric A

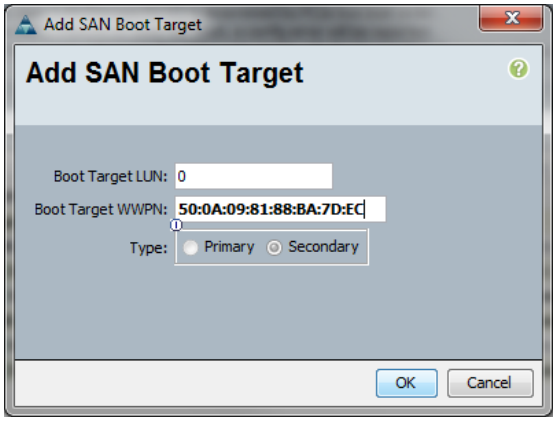

68. Click **OK**, and then click **OK** again to create the boot policy.

Create Service Profile Templates

In this procedure, two service profile templates are created: one for fabric A boot and one for fabric B boot. The first profile is created and then cloned and modified for the second host.

To create service profile templates, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Service Profile Templates** > **root**.
3. Right-click root.
4. Choose **Create Service Profile Template** to open the Create Service Profile Template wizard.
5. Identify the Service Profile Template:
   a. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from node 1 on fabric A.
   b. Click the **Updating Template** radio button.
   c. Under UUID, choose UUID_Pool as the UUID pool.
   d. Click **Next**.

Figure 44 Details for Creating Service Profile Template

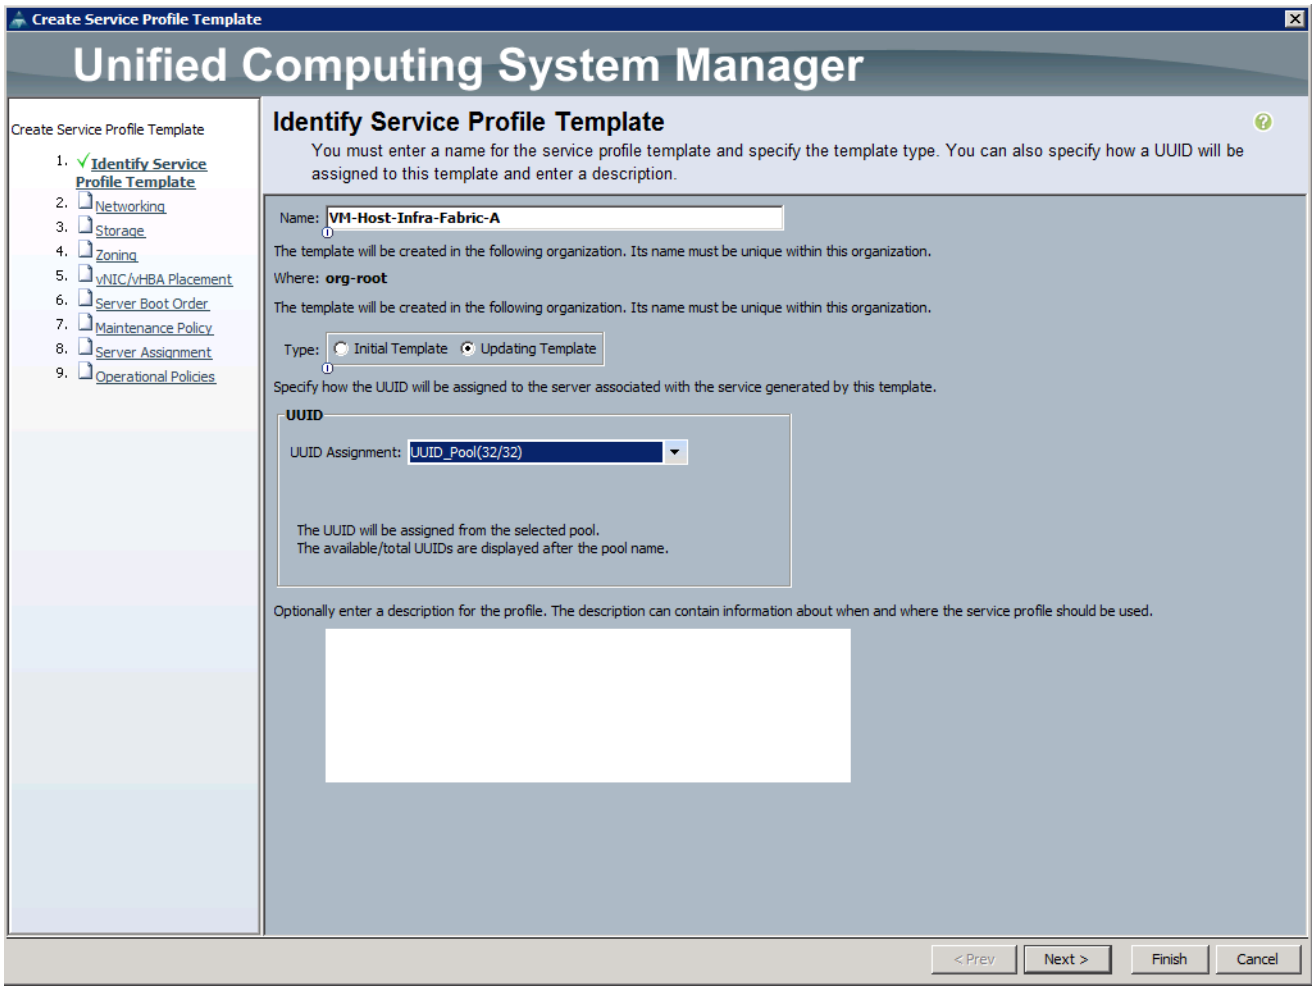

6. Configure the Networking options:
   a. Keep the default setting for Dynamic vNIC Connection Policy.
   b. Click the Expert radio button to configure the LAN connectivity.
   c. Click **Add** to add a vNIC to the template.
   d. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
   e. Check the Use vNIC Template check box.
   f. In the vNIC Template list, choose vNIC_Template_A.
   g. In the Adapter Policy list, choose VMWare.
   h. Click **OK** to add this vNIC to the template.

Figure 45 Creating vNIC Using vNIC Template

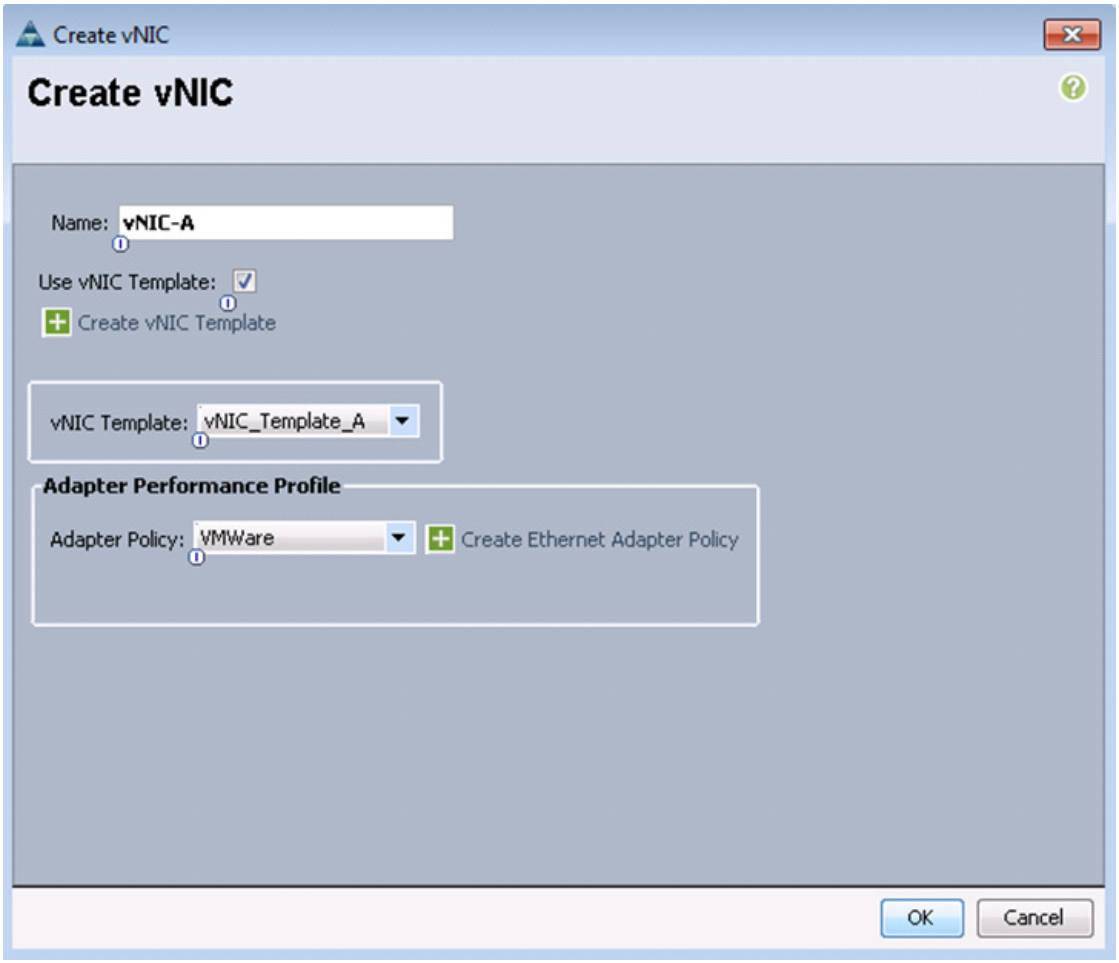

i. On the Networking page of the wizard, click **Add** to add another vNIC to the template.

j. In the Create vNIC box, enter vNIC-B as the name of the vNIC.

k. Check the Use vNIC Template check box.

l. In the vNIC Template list, choose vNIC_Template_B.

m. In the Adapter Policy list, choose VMWare.

n. Click **OK** to add the vNIC to the template.

o. Review the table in the Networking page to make sure that both vNICs were created.

p. Click **Next**.

Figure 46 LAN Configuration Details

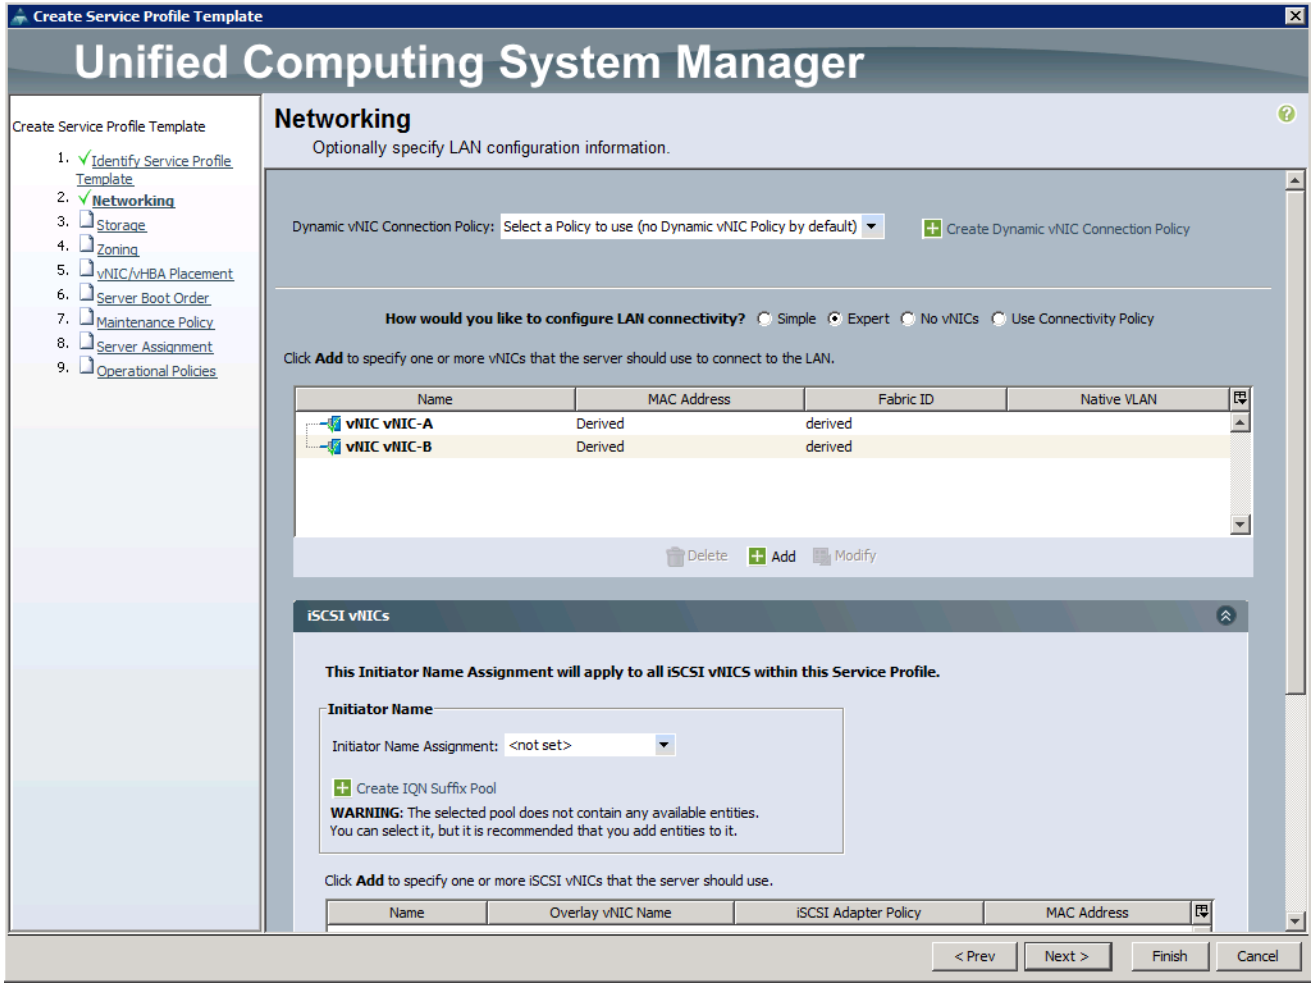

7. Configure the Storage options:
   a. Choose a local disk configuration policy:
   – If the server in question has local disks, choose default in the Local Storage list.
   – If the server in question does not have local disks, choose SAN-Boot.
   b. Click the **Expert** radio button to configure the SAN connectivity.
   c. In the WWNN Assignment list, choose WWNN_Pool.
   d. Click **Add** at the bottom of the page to add a vHBA to the template.
   e. In the Create vHBA dialog box, enter Fabric-A as the name of the vHBA.
   f. Check the Use vHBA Template check box.
   g. In the vHBA Template list, choose vHBA_Template_A.
   h. In the Adapter Policy list, choose VMware.
   i. Click **OK** to add this vHBA to the template.

Figure 47 Creating vHBA Using vHBA Template

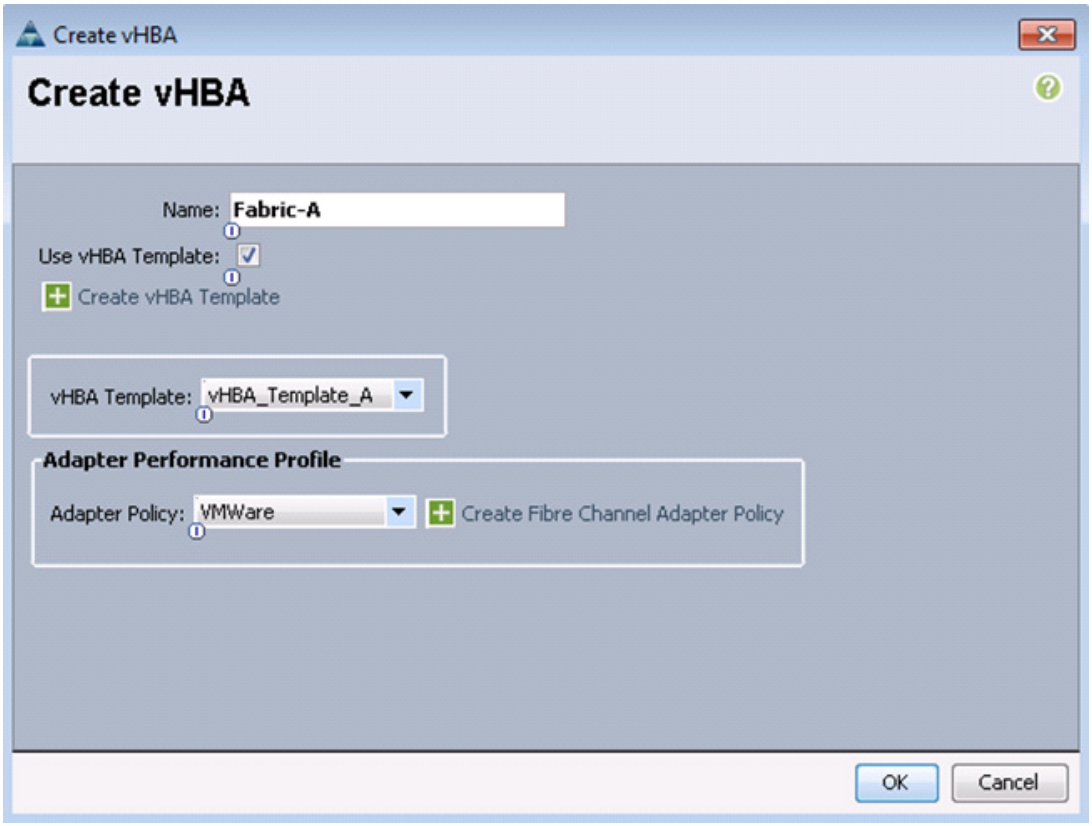

j. On the Storage page of the wizard, click **Add** at the bottom of the page to add another vHBA to the template.

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

l. Check the check box for Use HBA Template.

m. In the vHBA Template list, choose vHBA_Template_B.

n. In the Adapter Policy list, choose VMware.

o. Click **OK** to add the vHBA to the template.

p. Review the table in the Storage page to verify that both vHBAs were created.

q. Click **Next**.

Figure 48 Storage Window Showing Created vHBAs

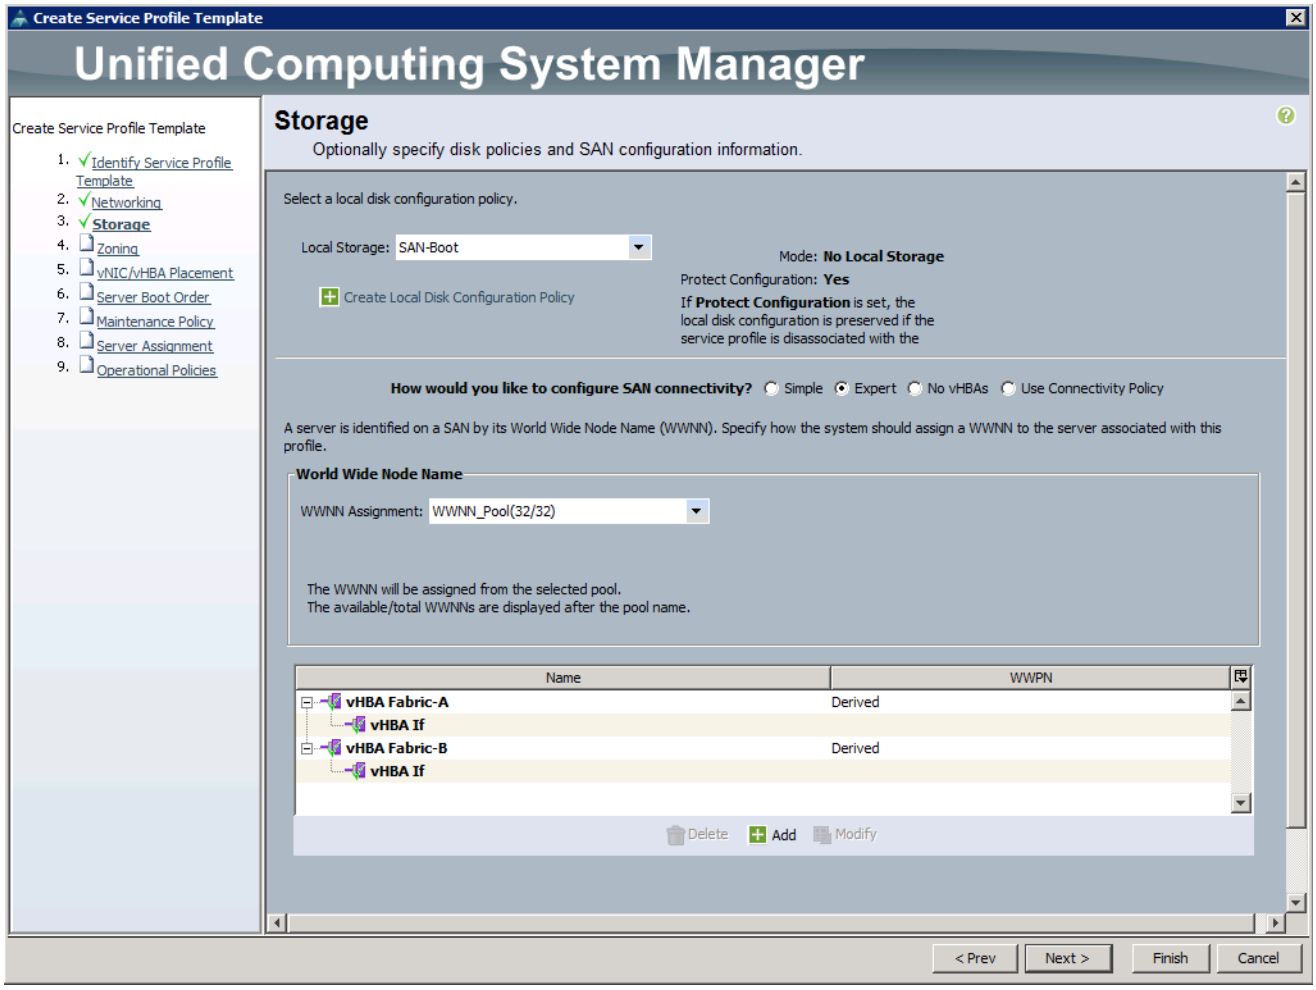

8. Set no Zoning options and click **Next**.
9. Set the vNIC/vHBA placement options.
   a. In the Select Placement list, choose the VM-Host-Infra placement policy.
   b. Choose vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
      - vHBA Fabric-A
      - vHBA Fabric-B
      - vNIC-A
      - vNIC-B
   c. Review the table to verify that all vNICs and vHBAs were assigned to the policy in the appropriate order.
   d. Click **Next**.

Figure 49 Placing vNIC and vHBA on Physical Adapters

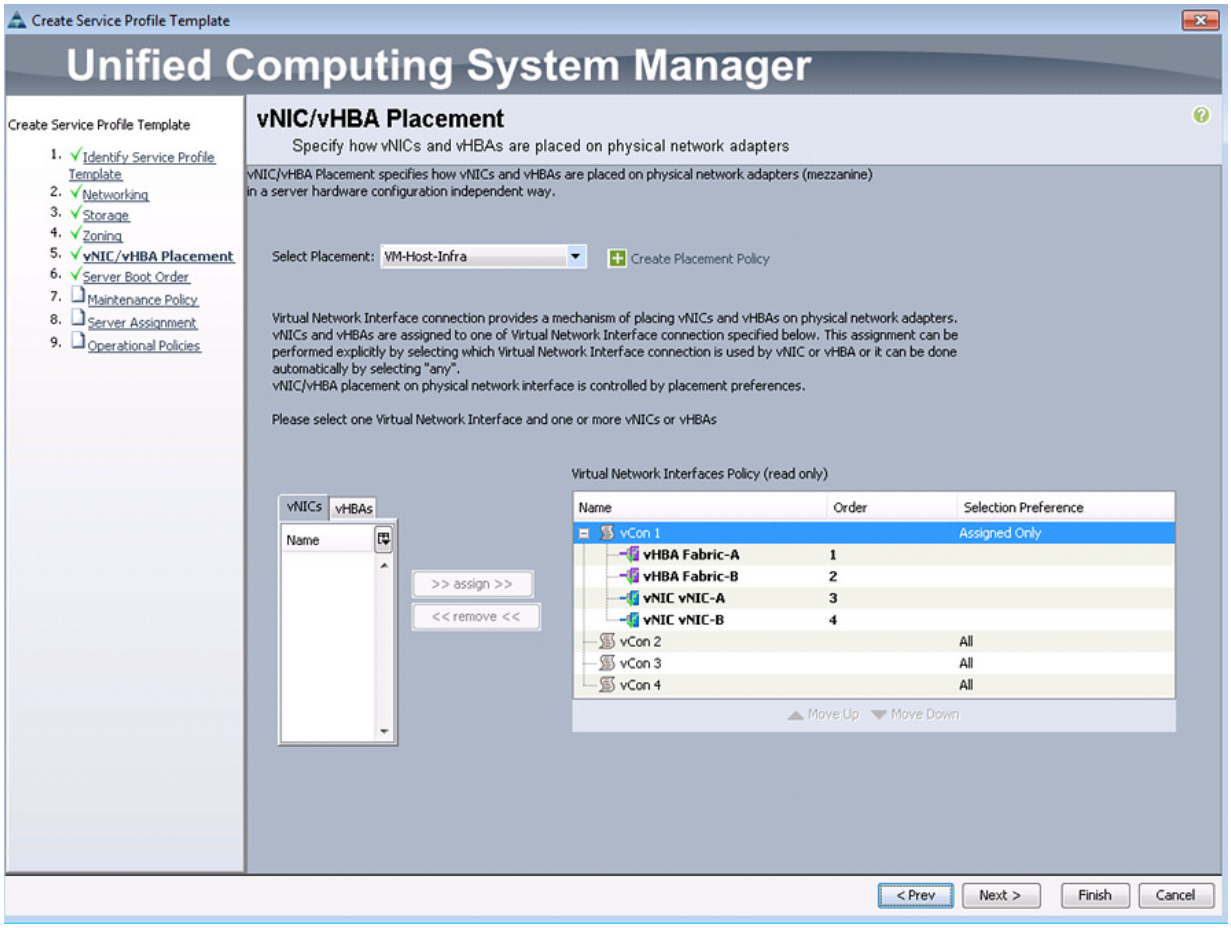

10. Set the Server Boot Order:
    a. In the Boot Policy list, choose Boot-Fabric-A.
    b. Review the table to verify that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
    c. Click **Next**.

Figure 50 Setting Boot Order for the Service Profile Template

11. Add a Maintenance Policy:
    a. Choose the Default Maintenance Policy.
    b. Click **Next**.
12. Specify the Server Assignment:
    a. In the Pool Assignment list, choose Infra_Pool.
    b. (Optional) Choose a Server Pool Qualification policy.
    c. Choose Down as the power state to be applied when the profile is associated with the server.
    d. Expand Firmware Management at the bottom of the page and choose VM-Host-Infra from the Host Firmware list.
    e. Click **Next**.

Figure 51 Assigning a Server Pool to the Service Profile Template

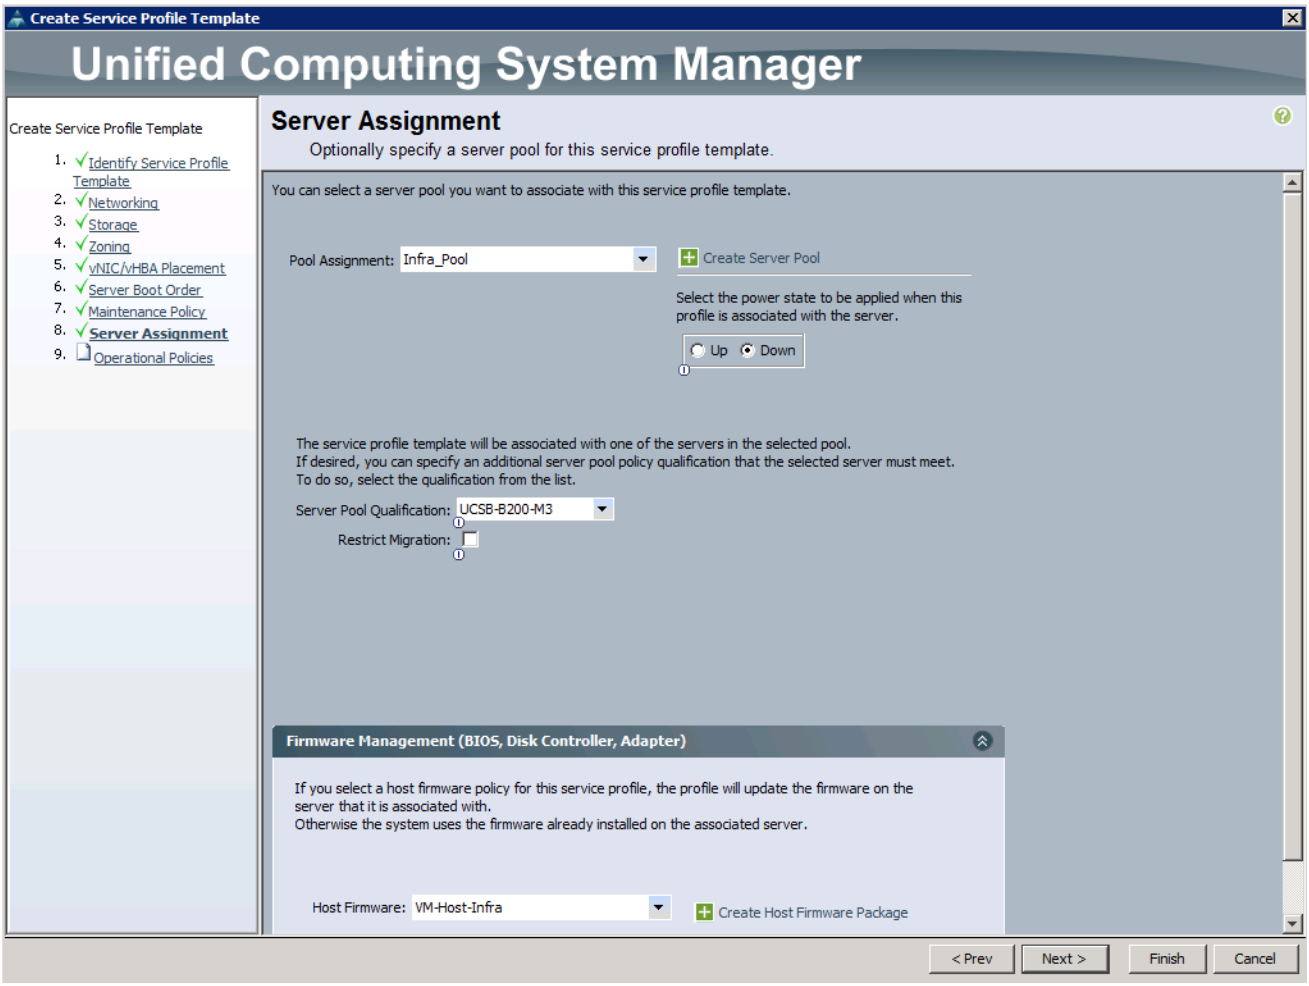

13. Add Operational Policies:
    a. In the BIOS Policy list, choose VM-Host-Infra.
    b. Expand Power Control Policy Configuration and choose No-Power-Cap in the Power Control Policy list.

Figure 52 Setting Operational Policy

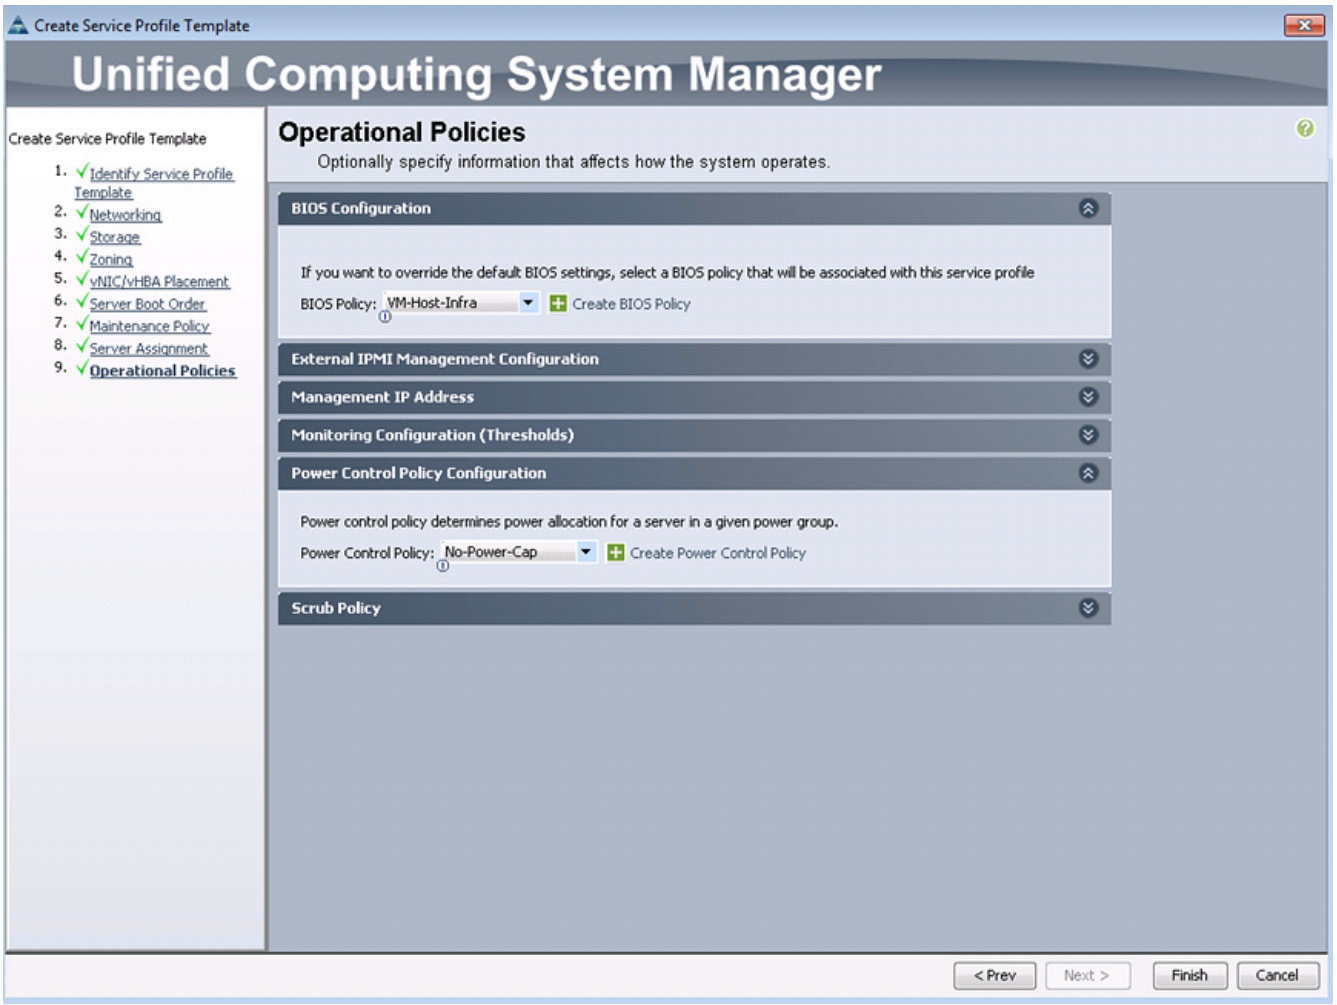

14. Click **Finish** to create the service profile template.
15. Click **OK** in the confirmation message.
16. Click the **Servers** tab in the navigation pane.
17. Choose **Service Profile Templates > root**.
18. Right-click the previously created VM-Host-Infra-Fabric-A template.
19. Choose **Create a Clone**.
20. In the dialog box, enter VM-Host-Infra-Fabric-B as the name of the clone, choose the root Org, and click **OK**.

Figure 53 Cloning a Service Profile Template

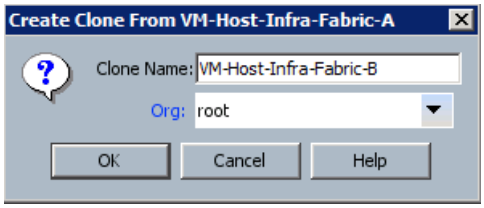

21. Click **OK**.
22. Choose the newly cloned service profile template and click the **Boot Order** tab.
23. Click **Modify Boot Policy**.
24. In the Boot Policy list, choose Boot-Fabric-B.

Figure 54 Modifying Boot Policy

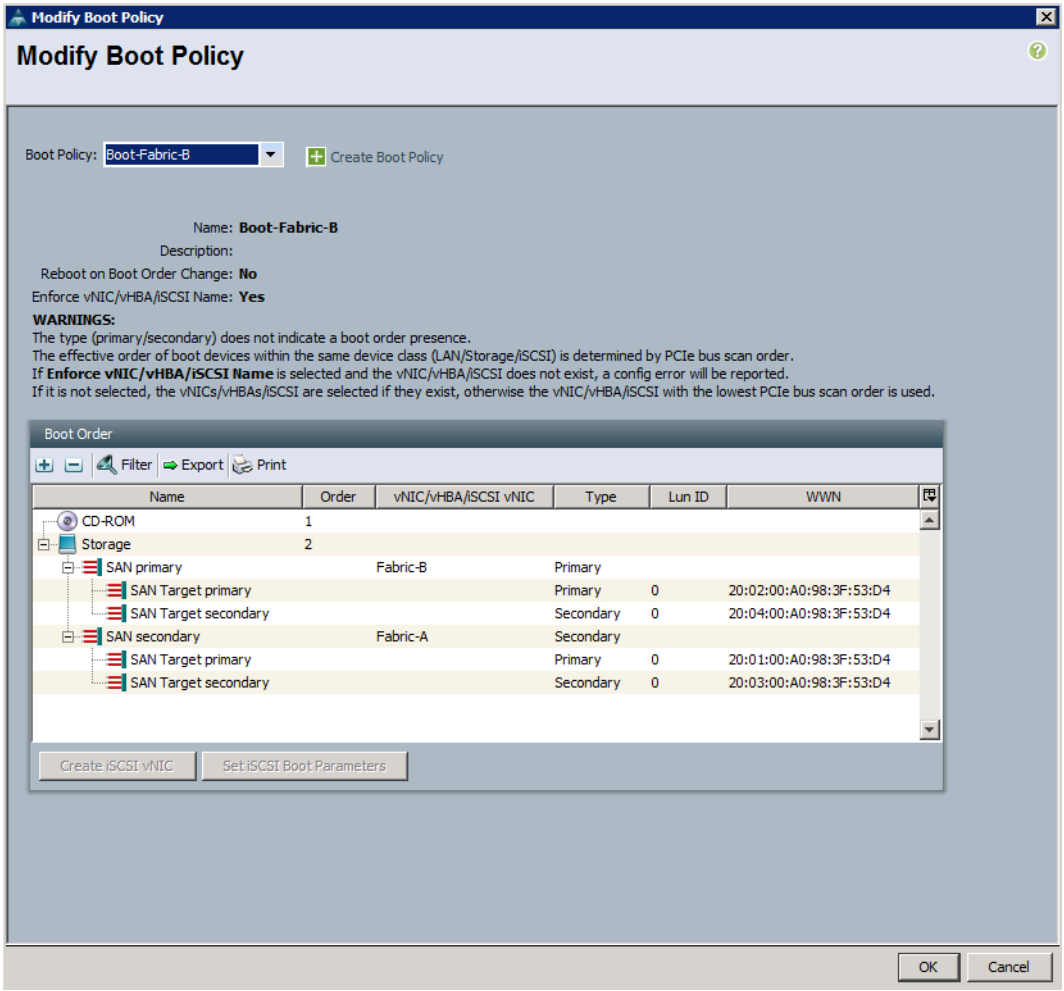

25. Click **OK**, and then click **OK** again.
26. In the right pane, click the **Network** tab and then click **Modify vNIC/HBA Placement**.
27. Expand vCon 1 and move vHBA Fabric-B ahead of vHBA Fabric-A in the placement order.

Figure 55 Modifying Placement Policy

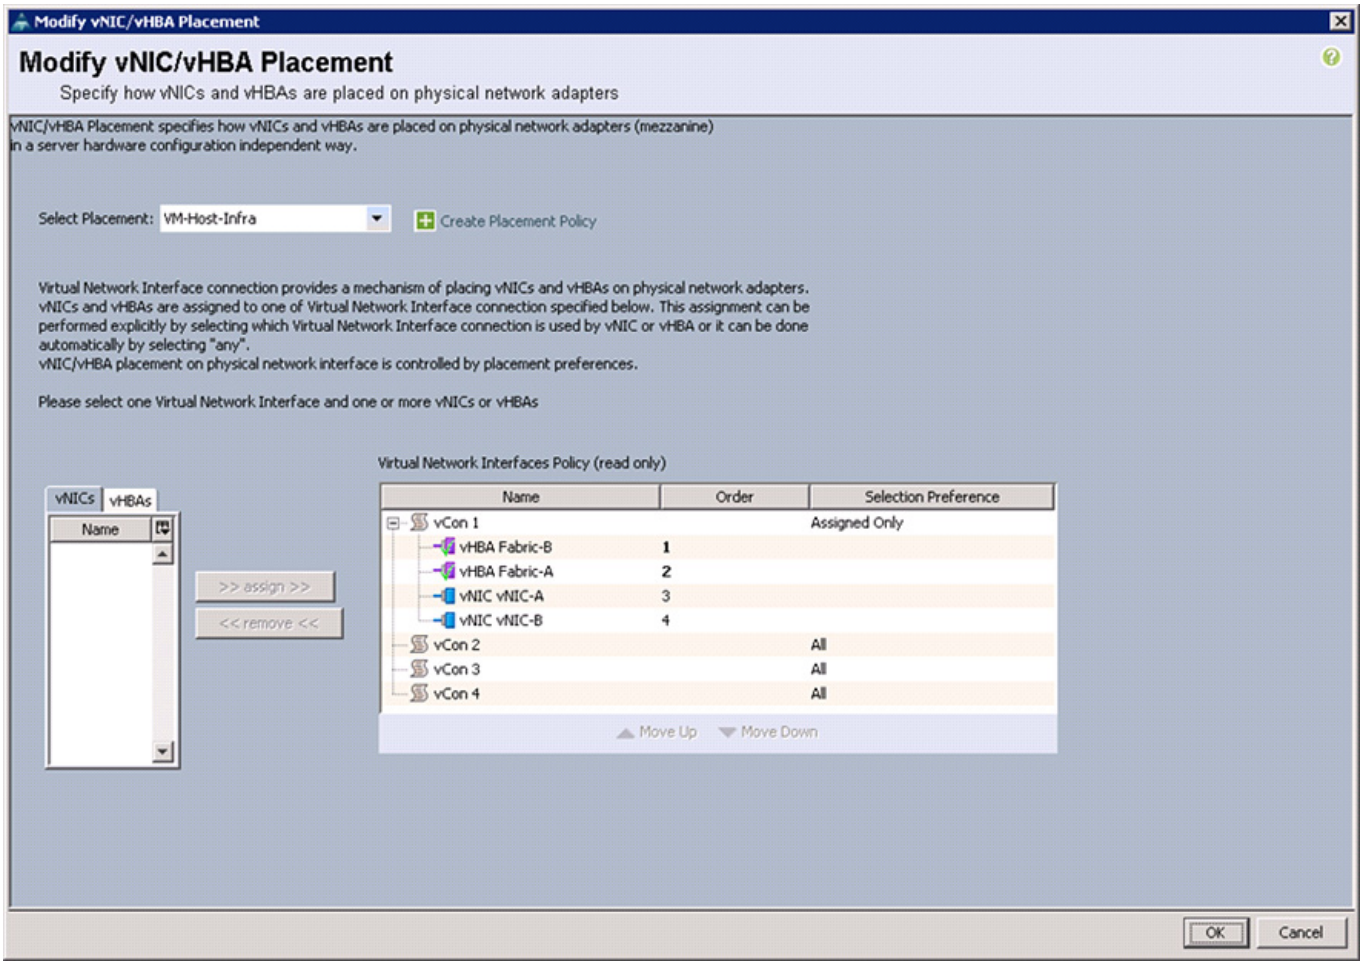

28. Click **OK**, and then click **OK** again.

Create Service Profiles

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Service Profile Templates > root > Service Template VM-Host-Infra-Fabric-A**.
3. Right-click VM-Host-Infra-Fabric-A and choose **Create Service Profiles from Template**.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Keep 1 as the Suffix Starting Number.
6. Enter 1 as the number of service profiles to create.
7. Click **OK** to create the service profile.

Figure 56 Creating Service Profile from a Service Profile Template

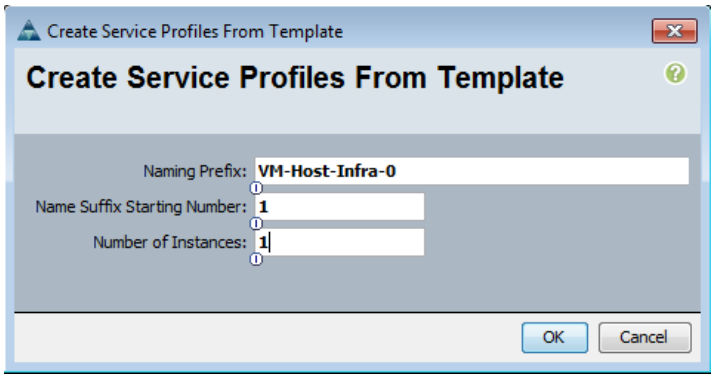

8. Click **OK** in the confirmation message.
9. Choose **Service Profile Templates** > **root** > **Service Template VM-Host-Infra-Fabric-B**.
10. Right-click VM-Host-Infra-Fabric-B and choose **Create Service Profiles from Template**.
11. Enter VM-Host-Infra-0 as the service profile prefix.
12. Keep 2 as the Suffix Starting Number.
13. Enter 1 as the number of service profiles to create.
14. Click **OK** to create the service profile.

Figure 57 Creating Service Profile from a Service Profile Template

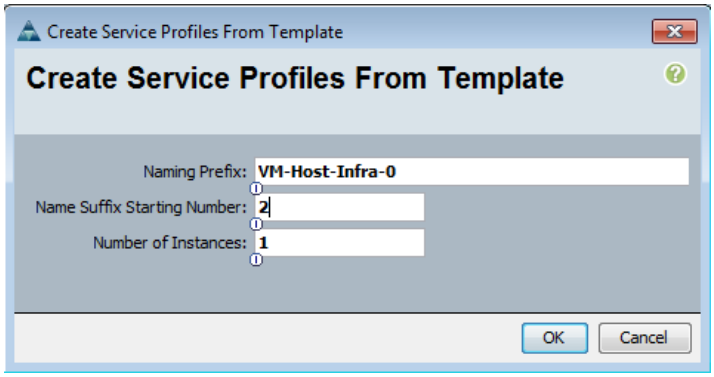

15. Click **OK** in the confirmation message.

    Verify that the service profiles VM-Host-Infra-01 and VM-Host-Infra-02 have been created. The service profiles are automatically associated with the servers in their assigned server pools.
16. (Optional) Choose each newly created service profile and enter the server host name or the FQDN in the User Label field in the General tab. Click Save Changes to map the server host name to the service profile name.

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 22 and Table 23.

Table 22 FC Port Names for Storage Controllers 1 and 2

Storage Controller	FCoE Port	FC Port Name
1	1a	
1	2a	
2	1a	
2	2a	

Note To gather the FC port name information, run the **fcp show adapters** command on the storage controller.

Table 23 vHBA WWPNs for Fabric A and Fabric B

Cisco UCS Service Profile Name	Fabric A vHBA WWPN	Fabric B vHBA WWPN
VM-Host-infra-01		
VM-Host-infra-02		

Note To gather the vHBA WWPN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the **Servers** tab. Expand **Servers** > **Service Profiles** > **root**. Click each service profile and then click the **Storage** tab in the right pane. In Table 23, record the WWPN information that is displayed in the right pane for both the Fabric A vHBA and the Fabric B vHBA for each service profile.

Storage Networking

FlexPod Cisco Nexus Base

Table 24 FlexPod Cisco Nexus Base Prerequisite

Description
The Cisco Nexus switch must be running Cisco Nexus NX-OS 6.0(2)N1(2a) or late

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. Follow these steps precisely; failure to do so might result in an improper configuration.

Set up Initial Configuration

Cisco Nexus 5548UP A

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

1. Configure the switch.

Note 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) [no]:
yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for the "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>>
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
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Enter basic FC configurations (yes/no) [n]: 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 5548UP B

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

1. Configure the switch.

Note 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) [no]:
yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for the "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>>
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
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Enter basic FC configurations (yes/no) [n]: 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 FCoE Storage vSphere on Data ONTAP 7-Mode

Enable Licenses

Cisco Nexus 5548UP A

To license the Cisco Nexus A switch on <<var_nexus_A_hostname>>, follow these steps:

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

```
config t
feature fcoe
feature npiv
feature lacp
feature vpc
```

Cisco Nexus 5548UP B

To license the Cisco Nexus B switch on <<var_nexus_B_hostname>>, follow these steps:

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

```
config t
feature fcoe
feature npiv
feature lacp
feature vpc
```

Set Global Configurations

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To set global configurations, follow these steps on both switches:

Run the following commands to set global configurations and jumbo frames in QoS:

```
spanning-tree port type network default
```

```
spanning-tree port type edge bpduguard default
port-channel load-balance ethernet source-dest-port
policy-map type network-qos jumbo
class type network-qos class-default
mtu 9216
exit
class type network-qos class-fcoe
pause no-drop
mtu 2158
exit
exit
system qos
service-policy type network-qos jumbo
exit
copy run start
```

Create VLANs

Cisco Nexus 5548UP A and Cisco Nexus 5548UP 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_nfs_vlan_id>>
name NFS-VLAN
exit
vlan <<var_pkt-ctrl_vlan_id>>
name Packet-Control-VLAN
exit
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
exit
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 5548UP A

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

From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_controller1>>:e1a
exit
interface Eth1/2
description <<var_controller2>>:e1a
exit
interface Eth1/11
description <<var_ucs_clustername>>-A:1/19
```

```
exit
interface Eth1/12
description <<var_ucs_clustername>>-B:1/19
exit
interface Eth1/13
description <<var_nexus_B_hostname>>:1/13
exit
interface Eth1/14
description <<var_nexus_B_hostname>>:1/14
exit
interface eth1/31
description <<var_ucs_clustername>>-A:1/31
exit
interface eth1/32
description <<var_ucs_clustername>>-A:1/32
exit
```

Cisco Nexus 5548UP B

To add individual port descriptions for troubleshooting activity and verification for switch B, follow these steps:

From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_controller1>>:e2a
exit
interface Eth1/2
description <<var_controller2>>:e2a
exit
interface Eth1/11
description <<var_ucs_clustername>>-A:1/20
exit
interface Eth1/12
description <<var_ucs_clustername>>-B:1/20
exit
interface Eth1/13
description <<var_nexus_A_hostname>>:1/13
exit
interface Eth1/14
description <<var_nexus_A_hostname>>:1/14
exit
interface eth1/31
description <<var_ucs_clustername>>-B:1/31
exit
interface eth1/32
description <<var_ucs_clustername>>-B:1/32
exit
```

Create Port Channels

Cisco Nexus 5548UP A and Cisco Nexus 5548UP 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 Eth1/13-14
channel-group 10 mode active
```

```
no shutdown
exit
interface Po11
description <<var_controller1>>
exit
interface Eth1/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_controller2>>
exit
interface Eth1/2
channel-group 12 mode active
no shutdown
exit
interface Po13
description <<var_ucs_clustername>>-A
exit
interface Eth1/11
channel-group 13 mode active
no shutdown
exit
interface Po14
description <<var_ucs_clustername>>-B
exit
interface Eth1/12
channel-group 14 mode active
no shutdown
exit
copy run start
```

Configure Port Channels

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To configure the port channels, 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 <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_pkt-ctrl_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
spanning-tree port type network
no shutdown
exit
interface Po11
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface Po12
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
```

```
interface Po13
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>>
spanning-tree port type edge trunk
no shutdown
exit
interface Po14
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>>
spanning-tree port type edge trunk
no shutdown
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus 5548UP 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>>
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Cisco Nexus 5548UP 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>>
auto-recovery
exit
interface Po10
vpc peer-link
```

```
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Configure Ports for Cisco Nexus 1110-X Virtual Appliances

Cisco Nexus 5548 A

To configure the ports in switch A that are connected to the Cisco Nexus 1110-X, complete the following step:

From the global configuration mode, run the following commands:

```
interface Eth1/15
description <<var_nexus_1110x-1>>:Eth1
switchport mode trunk
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_pkt-ctrl_vlan_id>>
speed 1000
spanning-tree port type edge trunk
vpc orphan-port suspend
no shutdown
exit
interface Eth1/16
description <<var_nexus_1110x-2>>:Eth1
switchport mode trunk
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_pkt-ctrl_vlan_id>>
speed 1000
spanning-tree port type edge trunk
vpc orphan-port suspend
no shutdown
exit
copy run start
```

Cisco Nexus 5548 B

To configure the ports in switch B that are connected to the Cisco Nexus 1110-X, complete the following step:

From the global configuration mode, run the following commands:

```
interface Eth1/15
description <<var_nexus_1110x-1>>:Eth2
switchport mode trunk
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_pkt-ctrl_vlan_id>>
speed 1000
spanning-tree port type edge trunk
vpc orphan-port suspend
no shutdown
exit
interface Eth1/16
description <<var_nexus_1110x-2>>:Eth2
```

```
switchport mode trunk
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_pkt-ctrl_vlan_id>>
speed 1000
spanning-tree port type edge trunk
vpc orphan-port suspend
no shutdown
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 5548UP 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 configuration is completed.

Create VSANs, Assign and Enable Virtual Fibre Channel Ports

This procedure sets up Fibre Channel over Ethernet (FCoE) connections between the Cisco Nexus 5548UP switches, the Cisco UCS Fabric Interconnects, and the NetApp storage systems.

Cisco Nexus 5548UP A

To configure virtual storage area networks (VSANs), assign virtual Fibre Channel (vFC) ports, and enable vFC ports on switch A, follow these steps:

From the global configuration mode, run the following commands:

```
vlan <<var_fabric_a_fcoe_vlan_id>>
name FCoE_Fabric_A
fcoe vsan <<var_vsan_a_id>>
exit
interface po11
switchport trunk allowed vlan add <<var_fabric_a_fcoe_vlan_id>>
exit
interface vfc11
switchport description <<var_controller1>>:1a
bind interface Eth1/1
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
exit
interface po12
switchport trunk allowed vlan add <<var_fabric_a_fcoe_vlan_id>>
exit
interface vfc12
switchport description <<var_controller2>>:1a
bind interface Eth1/2
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
exit
interface po15
description <<var_ucs_clustername>>-A:FCoE
exit
interface Eth1/31-32
channel-group 15 mode active
exit
interface po15
```

```
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_fabric_a_fcoe_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface vfc15
switchport description <<var_ucs_clustername>>-A:FCoE
bind interface po15
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
vsan database
vsan <<var_vsan_a_id>> name Fabric_A
vsan <<var_vsan_a_id>> interface vfc11
vsan <<var_vsan_a_id>> interface vfc12
vsan <<var_vsan_a_id>> interface vfc15
exit
```

Cisco Nexus 5548UP B

To configure VSANs, assign vFC ports, and enable vFC ports on switch B, follow these steps:

From the global configuration mode, run the following commands:

```
vlan <<var_fabric_b_fcoe_vlan_id>>
name FCoE_Fabric_B
fcoe vsan <<var_vsan_b_id>>
exit
interface po11
switchport trunk allowed vlan add <<var_fabric_b_fcoe_vlan_id>>
exit
interface vfc11
switchport description <<var_controller1>>:2a
bind interface Eth1/1
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
exit
interface po12
switchport trunk allowed vlan add <<var_fabric_b_fcoe_vlan_id>>
exit
interface vfc12
switchport description <<var_controller2>>:2a
bind interface Eth1/2
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
exit
interface po15
description <<var_ucs_clustername>>-B:FCoE
exit
interface Eth1/31-32
channel-group 15 mode active
exit
interface po15
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_fabric_b_fcoe_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface vfc15
switchport description <<var_ucs_clustername>>-B:FCoE
bind interface po15
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
```

```
vsan database
vsan <<var_vsan_b_id>> name Fabric_B
vsan <<var_vsan_b_id>> interface vfc11
vsan <<var_vsan_b_id>> interface vfc12
vsan <<var_vsan_b_id>> interface vfc15
exit
```

Create Device Aliases

Cisco Nexus 5548UP A

To configure device aliases and zones for the primary boot paths of switch A on <<var_nexus_A_hostname>>, complete the following step:

From the global configuration mode, run the following commands:

```
device-alias database
device-alias name VM-Host-Infra-01_A pwwn <<var_vm_host_infra_01_A_wwpn>>
device-alias name VM-Host-Infra-02_A pwwn <<var_vm_host_infra_02_A_wwpn>>
device-alias name <<var_controller1>>_1a pwwn <<var_controller1_1a_wwpn>>
device-alias name <<var_controller2>>_1a pwwn <<var_controller2_1a_wwpn>>
exit
device-alias commit
```

Cisco Nexus 5548UP B

To configure device aliases and zones for the boot paths of switch B on <<var_nexus_B_hostname>>, complete the following step:

From the global configuration mode, run the following commands:

```
device-alias database
device-alias name VM-Host-Infra-01_B pwwn <<var_vm_host_infra_01_B_wwpn>>
device-alias name VM-Host-Infra-02_B pwwn <<var_vm_host_infra_02_B_wwpn>>
device-alias name <<var_controller1>>_2a pwwn <<var_controller1_2a_wwpn>>
device-alias name <<var_controller2>>_2a pwwn <<var_controller2_2a_wwpn>>
exit
device-alias commit
```

Create Zones

Cisco Nexus 5548UP A

To create zones for the service profiles on switch A, follow these steps:

1. Create a zone for each service profile.

```
zone name VM-Host-Infra-01_A vsan <<var_vsan_a_id>>
member device-alias VM-Host-Infra-01_A
member device-alias <<var_controller1>>_1a
member device-alias <<var_controller2>>_1a
exit
zone name VM-Host-Infra-02_A vsan <<var_ vsan_a_id>>
member device-alias VM-Host-Infra-02_A
member device-alias <<var_controller1>>_1a
member device-alias <<var_controller2>>_1a
exit
```

2. After the zone for the Cisco UCS service profiles has been created, create the zone set and add the necessary members.

```
zoneset name FlexPod vsan <<var_ vsan_a_id>>
member VM-Host-Infra-01_A
member VM-Host-Infra-02_A
exit
```

3. Activate the zone set.

```
zoneset activate name FlexPod vsan <<var_ vsan_a_id>>
exit
copy run start
```

Cisco Nexus 5548UP B

To create zones for the service profiles on switch B, follow these steps:

1. Create a zone for each service profile.

```
zone name VM-Host-Infra-01_B vsan <<var_ vsan_b_id>>
member device-alias VM-Host-Infra-01_B
member device-alias <<var_controller1>>_2a
member device-alias <<var_controller2>>_2a
exit
zone name VM-Host-Infra-02_B vsan <<var_ vsan_b_id>>
member device-alias VM-Host-Infra-02_B
member device-alias <<var_controller1>>_2a
member device-alias <<var_controller2>>_2a
exit
```

2. After all of the zones for the Cisco UCS service profiles have been created, create the zone set and add the necessary members.

```
zoneset name FlexPod vsan <<var_ vsan_b_id>>
member VM-Host-Infra-01_B
member VM-Host-Infra-02_B
exit
```

3. Activate the zone set.

```
zoneset activate name FlexPod vsan <<var_ vsan_b_id>>
exit
copy run start
```

Storage Part 2

Data ONTAP 7-Mode SAN Boot Storage Setup

The following subsections create initiator groups (igroups) on storage controller 1 and map the SAN boot LUNs to these igroups so that VMware ESXi can be installed on the LUNs for the two management hosts created.

Create Igroups

Enter the following commands to create two igroups:

Controller 1 Command Line Interface

```
igroup create -f -t vmware VM-Host-Infra-01 <<var_vm_host_infra_01_A_wwpn >>
<<var_vm_host_infra_01_B_wwpn >>
igroup create -f -t vmware VM-Host-Infra-02 <<var_vm_host_infra_02_A__wwpn >>
<<var_vm_host_infra_02_B_wwpn >>
```

Note To view the two igroups just created, type **igroup show**.

Controller 2 Command Line Interface

```
igroup create -f -t vmware MGMT-Hosts <<var_vm_host_infra_01_A_wwpn >>
<<var_vm_host_infra_01_B_wwpn >> <<var_vm_host_infra_02_A_wwpn >>
<<var_vm_host_infra_02_B_wwpn >>
```

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

Map Boot LUNs to Igroups

Enter the following commands to map the boot LUNS to igroups.

Controller 1 Command Line Interface

```
lun map /vol/esxi_boot/VM-Host-Infra-01 VM-Host-Infra-01 0
lun map /vol/esxi_boot/VM-Host-Infra-02 VM-Host-Infra-02 0
```

VMware vSphere 5.1Update1 Setup

FlexPod VMware ESXi 5.1Update1 FCoE on 7-Mode

This section provides detailed instructions for installing VMware ESXi 5.1Update1 in a FlexPod environment. After the procedures are completed, two FCP-booted ESXi hosts will be provisioned. These deployment procedures are customized to include the environment variables.

Note 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 Fibre Channel Protocol (FCP) boot Logical Unit Numbers (LUNs).

Log in to Cisco UCS 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, follow these steps:

1. Open a Web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the **Launch UCS Manager** link to 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.
6. From the main menu, click the **Servers** tab.
7. Click **Servers** > **Service Profiles** > **root** > **VM-Host-Infra-01**.
8. Right-click VM-Host-Infra-01 and click **KVM Console**.
9. If prompted to accept an unencrypted KVM session, accept as necessary.
10. Click **Servers** > **Service Profiles** > **root** > **VM-Host-Infra-02**.
11. Right-click VM-Host-Infra-02 and click **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-Infra-02

To prepare the server for the OS installation, follow these steps on each ESXi host:

1. In the KVM window, click the **Virtual Media** tab.
2. If prompted to accept an unencrypted KVM session, accept as necessary.
3. Click **Add Image**.
4. Browse to the ESXi installer ISO image file and click **Open**.
5. Check the Mapped check box to map the newly added image.
6. Click the **KVM** tab to monitor the server boot.
7. Boot the server by selecting Boot Server and click **OK**. Then click **OK** again.

Install ESXi

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

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

1. On reboot, the machine detects the presence of the ESXi installation media. Choose the ESXi installer from the 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. Choose the NetApp LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Choose the appropriate keyboard layout and press **Enter**.
6. Enter and confirm the root password and press **Enter**.

7. The installer issues a warning that existing partitions will be removed from the volume. Press **F11** to continue with the installation.
8. After the installation is complete, uncheck the **Mapped** check box (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.

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

9. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, simply click **Yes** to unmap the image.
10. In 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, follow these 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, follow these 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. Choose the **Configure the Management Network** option and press **Enter**.
4. (Optional) Choose the **VLAN** option and press **Enter**.
5. Enter the <<var_ib-mgmt_vlan_id>> and press Enter.
6. From the Configure Management Network menu, choose **IP Configuration** and press **Enter**.
7. Choose the **Set Static IP Address** and **Network Configuration** option by using the space bar.
8. Enter the IP address for managing the first ESXi host: <<var_vm_host_infra_01_ip>>.
9. Enter the subnet mask for the first ESXi host.
10. Enter the default gateway for the first ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Choose the **IPv6 Configuration** option and press **Enter**.
13. Using the spacebar, deselect Enable IPv6 (restart required) and press Enter.
14. Choose the **DNS Configuration** option and press **Enter**.

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

15. Enter the IP address of the primary DNS server.
16. (Optional) Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.

18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Choose **Test Management Network** to verify that the management network is set up correctly and press **Enter**.
23. Press **Enter** to run the test.
24. Press **Enter** to exit the window.
25. Press **Esc** to log out of the VMware console.

ESXi Host VM-Host-Infra-02

To configure the VM-Host-Infra-02 ESXi host with access to the management network, follow these 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. Choose the **Configure the Management Network** option and press **Enter**.
4. Choose the **VLAN** (Optional) option and press **Enter**.
5. Enter the <<var_ib-mgmt_vlan_id>> and press Enter.
6. From the Configure Management Network menu, choose **IP Configuration** and press **Enter**.
7. Choose the **Set Static IP Address** and **Network Configuration** option by using the space bar.
8. Enter the IP address for managing the second ESXi host: <<var_vm_host_infra_02_ip>>.
9. Enter the subnet mask for the second ESXi host.
10. Enter the default gateway for the second ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Choose the **IPv6 Configuration** option and press **Enter**.
13. Using the spacebar, deselect Enable IPv6 (restart required) and press Enter.
14. Choose the **DNS Configuration** option and press **Enter**.

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

15. Enter the IP address of the primary DNS server.
16. (Optional) Enter the IP address of the secondary DNS server.
17. Enter the FQDN for the second ESXi host.
18. Press **Enter** to accept the changes to the DNS configuration.
19. Press **Esc** to exit the Configure Management Network submenu.
20. Press **Y** to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.

22. Choose **Test Management Network** to verify that the management network is set up correctly and press **Enter**.
23. Press **Enter** to run the test.
24. Press **Enter** to exit the window.
25. Press **Esc** to log out of the VMware console.

Download VMware vSphere Client and vSphere Remote CLI

To download the VMware vSphere Client and install Remote CLI, follow these 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 both the vSphere Client and the Windows version of vSphere Remote Command Line.

Note These applications are downloaded from the VMware Web site 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 by using the VMware vSphere Client, follow these 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-Infra-02

To log in to the VM-Host-Infra-02 ESXi host by using the VMware vSphere Client, follow these steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-Infra-02 as the host you are trying to connect to: <<var_vm_host_infra_02_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click **Login** to connect.

Download Updated Cisco VIC enic and fnic Drivers

To download the Cisco virtual interface card (VIC) eNIC and fNIC drivers, complete the following steps:

Note The eNIC version used in this configuration is 2.1.2.38, and the fNIC version is 1.5.0.45.

1. Open a web browser on the management workstation and navigate to:
   - VMware ESXi 5.x Driver for Cisco enic
   - VMware ESXi 5.x Driver for Cisco fnic
2. Log in and download the eNIC and fNIC drivers.
3. Extract the vib files from the downloaded zip files:
   a. Navigate to enic_driver_2.1.2.38-1023014.zip > enic_driver_2.1.2.38-offline_bundle-1023014.zip > vib20 > net-enic.

   Network: Cisco_bootbank_net-enic_2.1.2.38-1OEM.500.0.0.472560.vib

   b. Navigate to fnic_driver_1.5.0.45-1233300.zip > fnic_driver_1.5.0.45-offline_bundle-1233300.zip > vib20 > scsi-fnic.

   Storage: Cisco_bootbank_scsi-fnic_1.5.0.45-1OEM.500.0.0.472560.vib
4. Document the saved location.

Load Updated Cisco VIC enic and fnic Drivers

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

To load the updated versions of the enic and fnic drivers for the Cisco VIC, follow these steps for the hosts on each vSphere Client:

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Summary** tab to view the environment summary.
3. From **Resources** > **Storage**, right-click datastore1 and choose Browse Datastore.
4. Click the fourth button and choose Upload File.
5. Navigate to the saved location for the downloaded enic driver version and choose Cisco_bootbank_net-enic_2.1.2.38-1OEM.500.0.0.472560.vib.
6. Click **Open** to open the file.
7. Click the fourth button and choose Upload File.
8. Navigate to the saved location for the downloaded fnic driver version and choose Cisco_bootbank_scsi-fnic_1.5.0.45-1OEM.500.0.0.472560.vib.
9. Click **Open** to open the file.
10. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
11. At the command prompt, run the following commands to account for each host (enic):

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -v
/vmfs/volumes/datastore1/Cisco_bootbank_net-enic_2.1.2.38-1OEM.500.0.0.472560.vib
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -v
/vmfs/volumes/datastore1/Cisco_bootbank_net-enic_2.1.2.38-1OEM.500.0.0.472560.vib
```

12. At the command prompt, run the following commands to account for each host (fnic):

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -v
/vmfs/volumes/datastore1/Cisco_bootbank_scsi-fnic_1.5.0.45-1OEM.500.0.0.472560.vib
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -v
/vmfs/volumes/datastore1/Cisco_bootbank_scsi-fnic_1.5.0.45-1OEM.500.0.0.472560.vib
```

13. From the vSphere Client, right-click each host in the inventory and choose **Reboot**.
14. Choose **Yes** to continue.
15. Enter a reason for the reboot and click **OK**.
16. After the reboot is complete, log back in to both hosts using the vSphere Client.

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, follow these steps:

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Configuration** tab.
3. Click **Networking** in the Hardware pane.
4. Click **Properties** on the right side of vSwitch0.
5. Choose the **vSwitch configuration** and click **Edit**.
6. From the General tab, change the MTU to 9000.
7. Click **OK** to close the properties for vSwitch0.
8. Choose the **Management Network configuration** and click **Edit**.
9. Change the network label to VMkernel-MGMT and check the Management Traffic check box.
10. Click **OK** to finalize the edits for Management Network.
11. Choose the **VM Network configuration** and click **Edit**.
12. Change the network label to IB-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 **Add** to add a network element.
15. Choose **VMkernel** and click **Next**.
16. Change the network label to VMkernel-NFS and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.
17. Click **Next** to continue with the NFS VMkernel creation.
18. Enter the IP address <<var_nfs_vlan_id_ip_host-01>> and the subnet mask <<var_nfs_vlan_id_mask_host01>> for the NFS VLAN interface for VM-Host-Infra-01.
19. Click **Next** to continue with the NFS VMkernel creation.
20. Click **Finish** to finalize the creation of the NFS VMkernel interface.
21. Choose the **VMkernel-NFS configuration** and click **Edit**.
22. Change the MTU to 9000.

23. Click **OK** to finalize the edits for the VMkernel-NFS network.
24. Click **Add** to add a network element.
25. Choose **VMkernel** and click **Next**.
26. Change the network label to VMkernel-vMotion and enter **<<var_vmotion_vlan_id>>** in the VLAN ID (Optional) field.
27. Check the Use This Port Group for vMotion check box.
28. Click **Next** to continue with the vMotion VMkernel creation.
29. Enter the IP address <<var_vmotion_vlan_id_ip_host-01>> and the subnet mask <<var_vmotion_vlan_id_mask_host-01>> for the vMotion VLAN interface for VM-Host-Infra-01.
30. Click **Next** to continue with the vMotion VMkernel creation.
31. Click **Finish** to finalize the creation of the vMotion VMkernel interface.
32. Choose the **VMkernel-vMotion configuration** and click **Edit**.
33. Change the MTU to 9000.
34. Click **OK** to finalize the edits for the VMkernel-vMotion network.
35. Close the dialog box to finalize the ESXi host networking setup. The networking for the ESXi host should be similar to Figure 58.

Figure 58 vSphere Client Showing VMKernel Ports and Virtual Switch

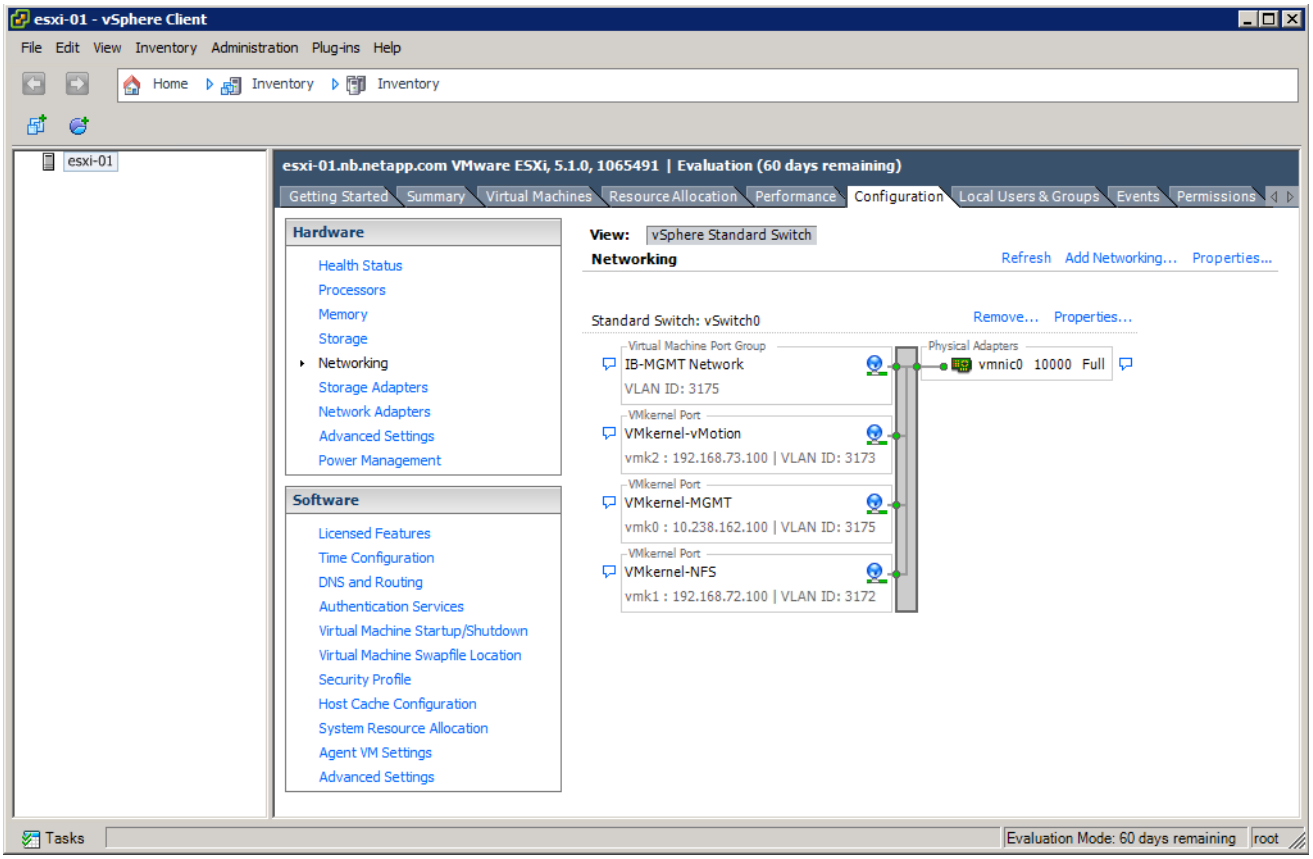

ESXi Host VM-Host-Infra-02

To set up the VMkernel ports and the virtual switches on the VM-Host-Infra-02 ESXi host, follow these steps:

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Configuration** tab.
3. Click **Networking** in the Hardware pane.
4. Click **Properties** on the right side of vSwitch0.
5. Choose the **vSwitch configuration** and click **Edit**.
6. From the General tab, change the MTU to 9000.
7. Click **OK** to close the properties for vSwitch0.
8. Choose the **Management Network configuration** and click **Edit**.
9. Change the network label to VMkernel-MGMT and check the Management Traffic check box.
10. Click **OK** to finalize the edits for Management Network.
11. Choose the **VM Network configuration** and click **Edit**.
12. Change the network label to IB-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 **Add** to add a network element.
15. Choose **VMkernel** and click **Next**.
16. Change the network label to VMkernel-NFS and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.
17. Click **Next** to continue with the NFS VMkernel creation.
18. Enter the IP address <<var_nfs_vlan_id_ip_host-02>> and the subnet mask <<var_nfs_vlan_id_mask_host02>> for the NFS VLAN interface for VM-Host-Infra-02.
19. Click **Next** to continue with the NFS VMkernel creation.
20. Click **Finish** to finalize the creation of the NFS VMkernel interface.
21. Choose the **VMkernel-NFS configuration** and click **Edit**.
22. Change the MTU to 9000.
23. Click **OK** to finalize the edits for the VMkernel-NFS network.
24. Click **Add** to add a network element.
25. Choose **VMkernel** and click **Next**.
26. Change the network label to VMkernel-vMotion and enter <<var_vmotion_vlan_id>> in the VLAN ID (Optional) field.
27. Check the Use This Port Group for vMotion check box.
28. Click **Next** to continue with the vMotion VMkernel creation.
29. Enter the IP address <<var_vmotion_vlan_id_ip_host-02>> and the subnet mask <<var_vmotion_vlan_id_mask_host-02>> for the vMotion VLAN interface for VM-Host-Infra-02.
30. Click **Next** to continue with the vMotion VMkernel creation.
31. Click **Finish** to finalize the creation of the vMotion VMkernel interface.

32. Choose the **VMkernel-vMotion configuration** and click **Edit**.
33. Change the MTU to 9000.
34. Click **OK** to finalize the edits for the VMkernel-vMotion network.
35. Close the dialog box to finalize the ESXi host networking setup. The networking for the ESXi host should be similar to Figure 59.

Figure 59 vSphere Client Showing VMKernel Ports and Virtual Switch

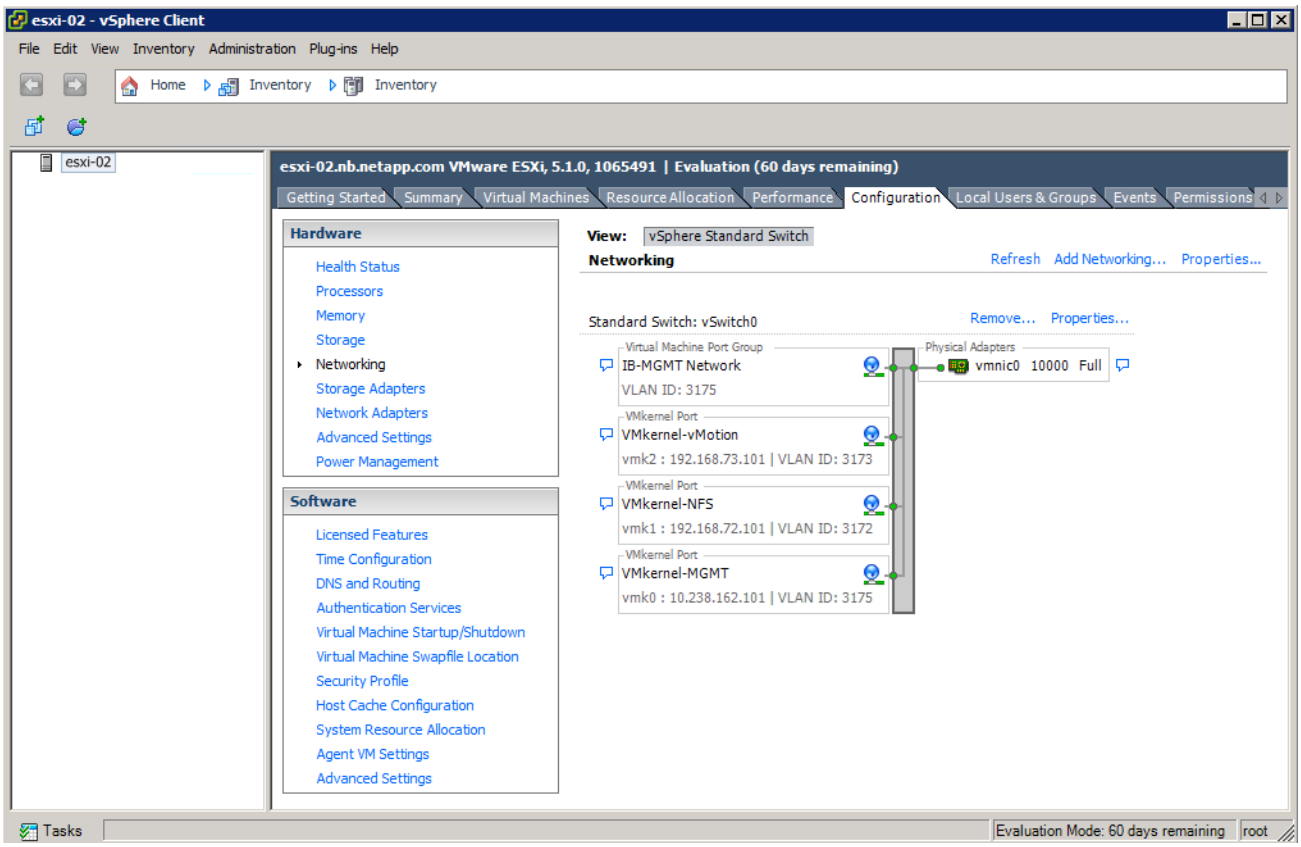

Mount Required Datastores

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

To mount the required datastores, follow these steps on each ESXi host:

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Configuration** tab to enable configurations.
3. Click **Storage** in the Hardware pane.
4. From the Datastore area, click **Add Storage** to open the Add Storage wizard.
5. Choose **Network File System** and click **Next**.
6. The wizard prompts for the location of the NFS export. Enter <<var_controller2_nfs_ip>> as the IP address for NFSIP Address for Controller 2.
7. Enter **/vol/infra_datastore_1** as the path for the NFS export.

8. Make sure that the Mount NFS read only check box is unchecked.
9. Enter **infra_datastore_1** as the datastore name.
10. Click **Next** to continue with the NFS datastore creation.
11. Click **Finish** to finalize the creation of the NFS datastore.
12. From the Datastore area, click **Add Storage** to open the Add Storage wizard.
13. Choose **Network File System** and click **Next**.
14. The wizard prompts for the location of the NFS export. Enter <<var_controller1_nfs_ip>> as the IP address for NFSIP Address for Controller 1.
15. Enter **/vol/infra_swap** as the path for the NFS export.
16. Make sure that the Mount NFS read only check box is unchecked.
17. Enter **infra_swap** as the datastore name.
18. Click **Next** to continue with the NFS datastore creation.
19. Click **Finish** to finalize the creation of the NFS datastore.

Configure NTP on ESXi Hosts

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

To configure Network Time Protocol (NTP) on the ESXi hosts, follow these steps on each host:

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Configuration** tab to enable configurations.
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 Options dialog box, follow these steps:
   a. Click General in the left pane and choose 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_global_ntp_server_ip>> as the IP address of the NTP server and click **OK**.
8. In the NTP Daemon Options dialog box, check the Restart NTP Service to Apply Changes check box and click **OK**.
9. In the Time Configuration dialog box, follow these steps:
   a. Check the NTP Client Enabled check box and click **OK**.
   b. Verify that the clock is now set to approximately the correct time.

Note The NTP server time may vary slightly from the host time.

Move VM Swap File Location

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

To move the VM swap file location, follow these steps on each ESXi host:

1. From each vSphere Client, choose the host in the inventory.
2. Click the Configuration tab to enable configurations.
3. Click **Virtual Machine Swapfile Location** in the Software pane.
4. Click **Edit** at the upper-right side of the window.
5. Choose Store the swapfile in a swapfile datastore selected below.
6. Select infra_swap as the datastore in which to house the swap files.
7. Click **OK** to finalize moving the swap file location.

FlexPod VMware vCenter 5.1Update1

The procedures in the following subsections provide detailed instructions for installing VMware vCenter 5.1Update1 in a FlexPod environment. After the procedures are completed, a VMware vCenter Server will be configured along with a Microsoft SQL Server database to provide database support to vCenter. These deployment procedures are customized to include the environment variables.

Note This procedure focuses on the installation and configuration of an external Microsoft SQL Server 2008 R2 database, but other types of external databases are also supported by vCenter. For information about how to configure the database and integrate it into vCenter, see the VMware vSphere5.1 documentation at: http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html

To install VMware vCenter 5.1Update1, an accessible Windows Active Directory® (AD) Domain is necessary. If an existing AD Domain is not available, an AD virtual machine, or AD pair, can be set up in this FlexPod environment. See "Appendix" section on page 169.

Build Microsoft SQL Server VM

ESXi Host VM-Host-Infra-01

To build a SQL Server virtual machine (VM) for the VM-Host-Infra-01 ESXi host, follow these steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the vSphere Client, choose the host in the inventory pane.
3. Right-click the host and choose **New Virtual Machine**.
4. Click **Custom** and then click **Next**.
5. Enter a name for the VM. Click **Next**.
6. Choose infra_datastore_1. Click **Next**.
7. Choose Virtual Machine Version: 8. Click **Next**.
8. Verify that the Windows option and the Microsoft Windows Server 2008 R2 (64-bit) version are selected. Click **Next**.

9. Choose two virtual sockets and one core per virtual socket. Click **Next**.
10. Choose 4GB of memory. Click **Next**.
11. Choose one network interface card (NIC).
12. For NIC 1, choose the IB-MGMT Network option and the VMXNET 3 adapter. Click **Next**.
13. Keep the LSI Logic SAS option for the SCSI controller selected. Click **Next**.
14. Keep the Create a New Virtual Disk option selected. Click **Next**.
15. Make the disk size at least 60GB. Click **Next**.
16. Click **Next**.
17. Check the Edit the Virtual Machine Settings Before Completion check box. Click **Continue**.
18. Choose the **Options** tab.
19. Choose **Boot Options**.
20. Check the **Force BIOS Setup** check box.
21. Click **Finish**.
22. From the left pane, expand the host field by clicking the plus sign (+).
23. Right-click the newly created SQL Server VM and click **Open Console**.
24. Click the third button (green right arrow) to power on the VM.
25. Click the ninth button (CD with a wrench) to map the Windows Server 2008 R2 SP1 ISO, and then choose **Connect to ISO Image on Local Disk**.
26. Navigate to the Windows Server 2008 R2 SP1 ISO, select it, and click **Open**.
27. In the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to select CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.
28. The Windows Installer boots. Choose the appropriate language, time and currency format, and keyboard. Click **Next**.
29. Click **Install Now**.
30. Make sure that the Windows Server 2008 R2 Standard (Full Installation) option is selected. Click **Next**.
31. Read and accept the license terms and click **Next**.
32. Choose Custom (Advanced). Make sure that Disk 0 Unallocated Space is selected. Click **Next** to allow the Windows installation to complete.
33. After the Windows installation is complete and the VM has rebooted, click **OK** to set the Administrator password.
34. Enter and confirm the Administrator password and choose the blue arrow to log in. Click **OK** to confirm the password change.
35. After logging in to the VM desktop, from the VM console window, choose the VM menu. Under Guest, choose Install/Upgrade VMware Tools. Click **OK**.
36. If prompted to eject the Windows installation media before running the setup for the VMware tools, click **OK**, then click **OK**.
37. In the dialog box, choose Run setup64.exe.
38. In the VMware Tools installer window, click **Next**.

39. Make sure that Typical is selected and click **Next**.
40. Click **Install**.
41. If prompted to Trust Software from VMware, Inc, check the checkbox to always trust, and click **Install**.
42. Click **Finish**.
43. Click **Yes** to restart the VM.
44. After the reboot is complete, choose the VM menu. Under Guest, choose Send Ctrl+Alt+Del and then enter the password to log in to the VM.
45. Set the time zone for the VM, IP address, gateway, and host name. Add the VM to the Windows AD domain.

Note A reboot is required.

46. If necessary, activate Windows.
47. Log back in to the VM and download and install all required Windows updates.

Note This process requires several reboots.

Install Microsoft SQL Server 2008 R2

vCenter SQL Server VM

To install SQL Server on the vCenter SQL Server VM, follow these steps:

1. Connect to an AD domain controller in the FlexPod Windows domain and add an admin user for the FlexPod using the Active Directory Users and Computers tool. This user should be a member of the Domain Administrators security group.
2. Log in to the vCenter SQL Server VM as the FlexPod admin user and open Server Manager.
3. Expand Features and click **Add Features**.
4. Expand .NET Framework 3.5.1 Features and choose only .NET Framework 3.5.1.

Figure 60 Adding Features for Installing SQL Server

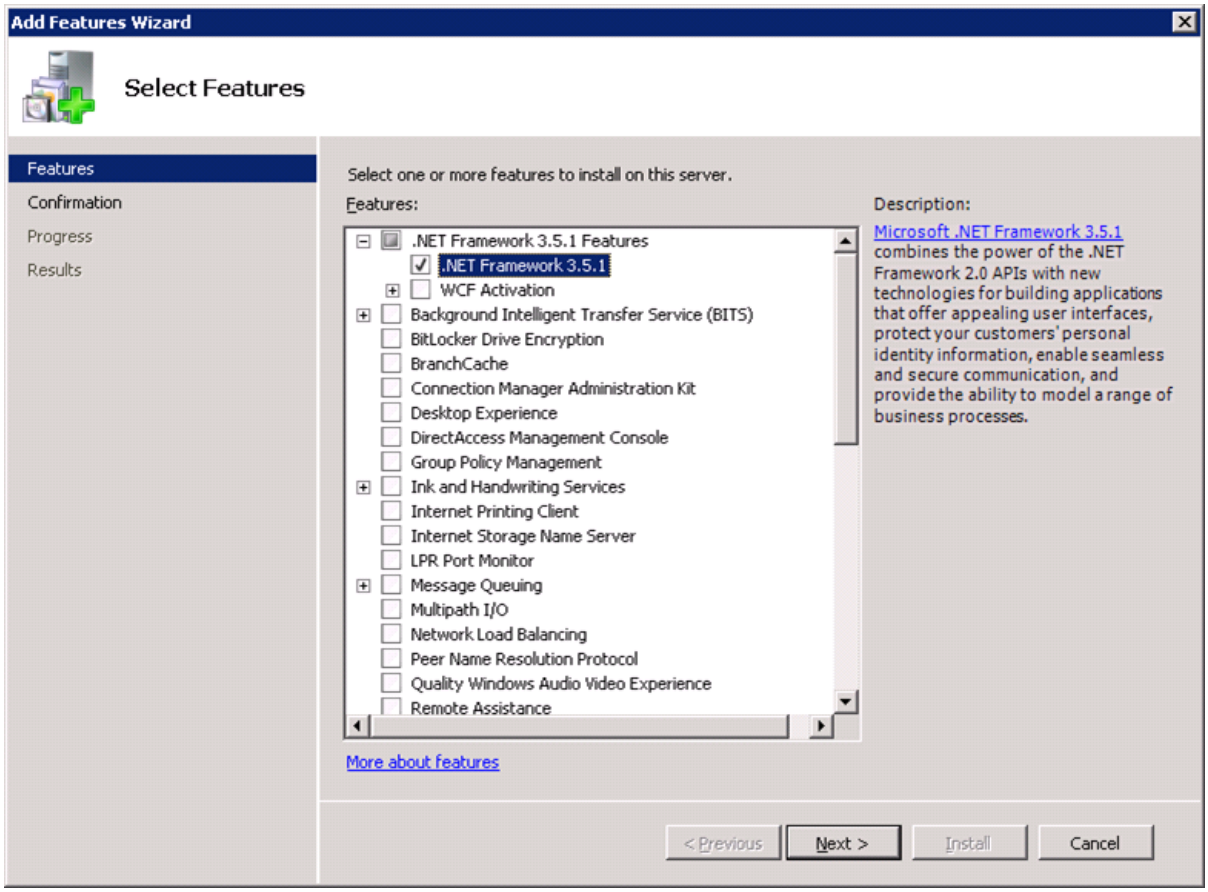

5. Click **Next**.
6. Click **Install**.
7. Click **Close**.
8. Open Windows Firewall with Advanced Security by navigating to **Start** > **Administrative Tools** > **Windows Firewall with Advanced Security**.
9. Choose Inbound Rules and click **New Rule**.
10. Choose Port and click **Next**.
11. Choose TCP and enter the specific local port 1433. Click **Next**.
12. Choose Allow the Connection. Click **Next**, and then click **Next** again.
13. Name the rule SQL Server and click **Finish**.
14. Close Windows Firewall with Advanced Security.
15. In the vCenter SQL Server VMware console, click the ninth button (CD with a wrench) to map the Microsoft SQL Server 2008 R2 ISO. Choose Connect to ISO Image on Local Disk.
16. Navigate to the SQL Server 2008 R2 ISO, select it, and click **Open**.
17. In the dialog box, click Run setup.exe.
18. In the SQL Server Installation Center window, click **Installation** on the left.
19. Choose New Installation or Add Features to an Existing Installation.

20. Click **OK**.
21. Choose Enter the Product Key. Enter a product key and click **Next**.
22. Read and accept the license terms and choose whether to check the second check box. Click **Next**.
23. Click **Install** to install the setup support files.
24. Address any warnings except for the Windows firewall warning. Click **Next**.

Note The Windows firewall issue was addressed in Step 13.

25. Choose SQL Server Feature Installation and click **Next**.
26. Under Instance Features, choose only Database Engine Services.
27. Under Shared Features, choose **Management Tools > Basic and Management Tools > Complete**. Click **Next**.

Figure 61 Selecting SQL Server Features

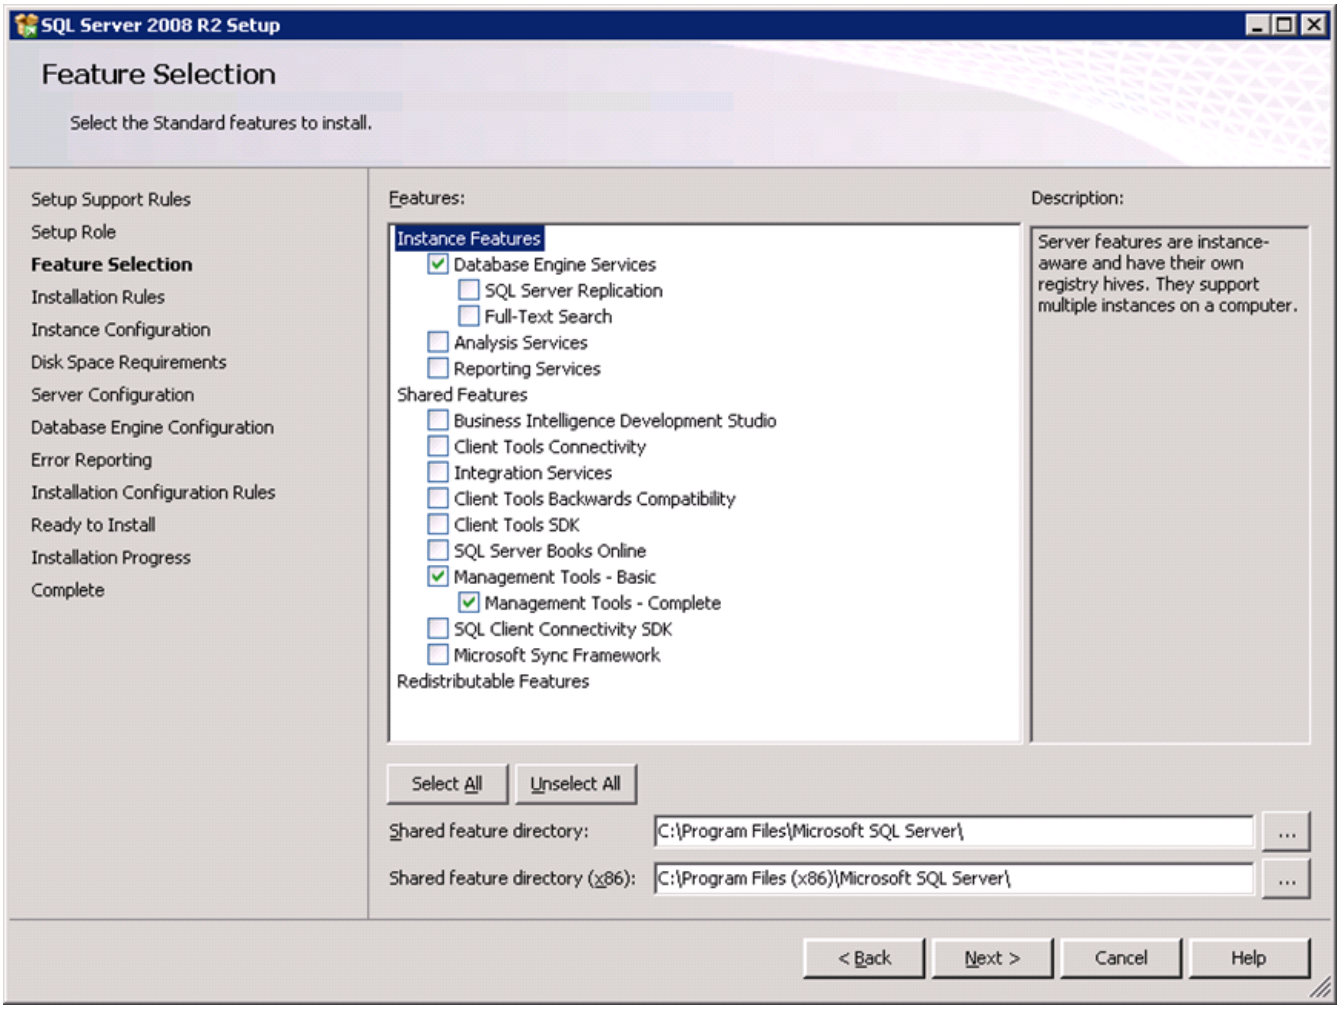

28. Click **Next**.

29. Keep Default Instance selected. Click **Next**.

Figure 62 SQL Server Instance Configuration

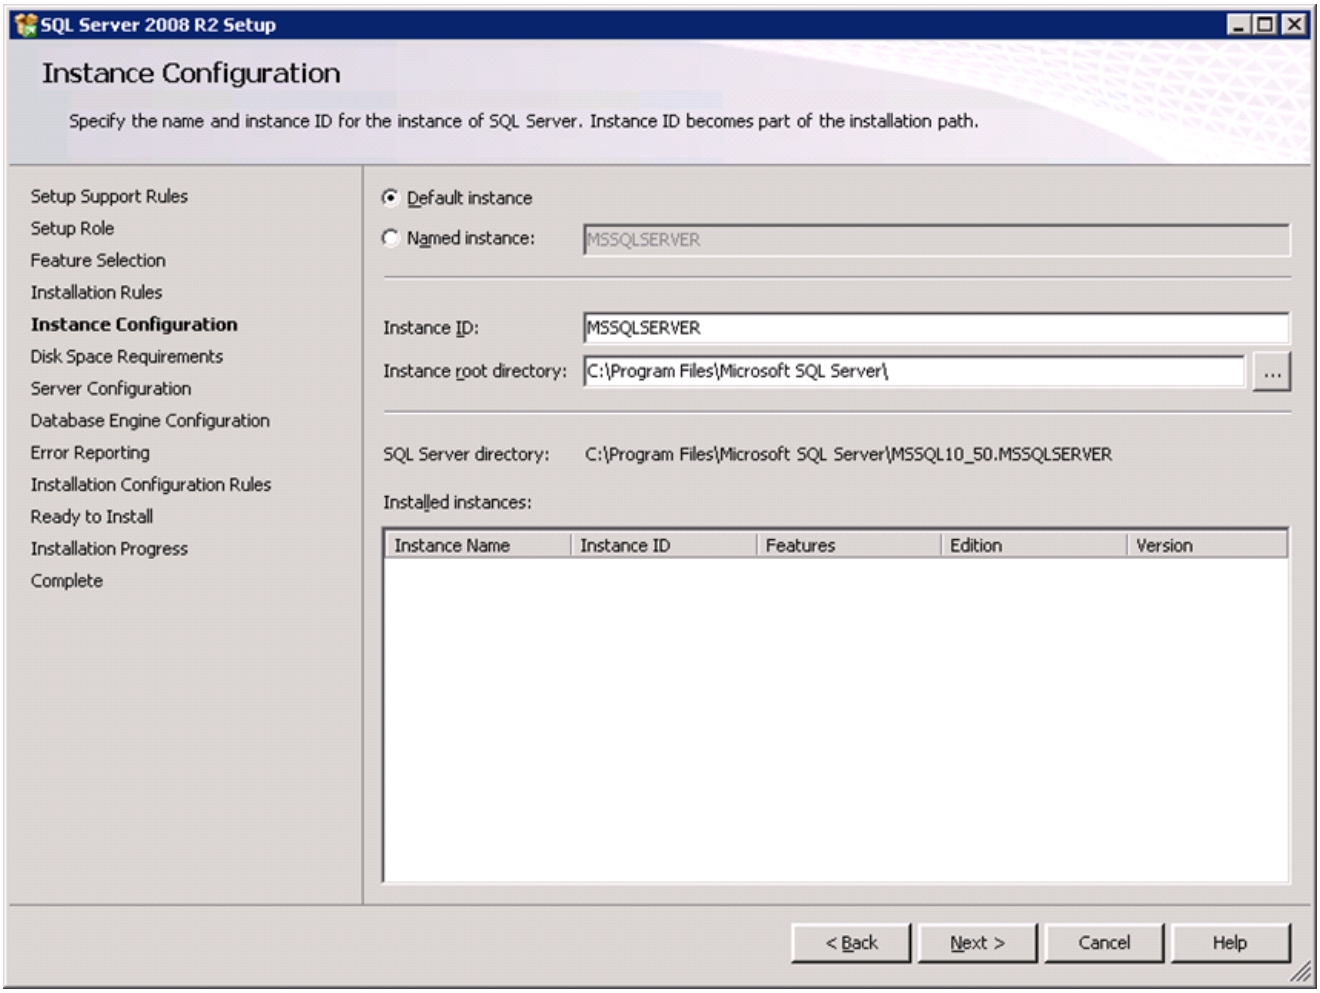

30. Click **Next** for Disk Space Requirements.
31. For the SQL Server Agent service, click in the first cell in the Account Name column and then click **<<Browse…>>**.
32. Enter the local machine administrator name (for example, systemname\Administrator), click **Check Names**, and click **OK**.
33. Enter the administrator password in the first cell under Password.
34. Change the startup type for SQL Server Agent to Automatic.
35. For the SQL Server Database Engine service, choose Administrator in the Account Name column and enter the administrator password again. Click **Next**.

Figure 63 Configuring SQL Server

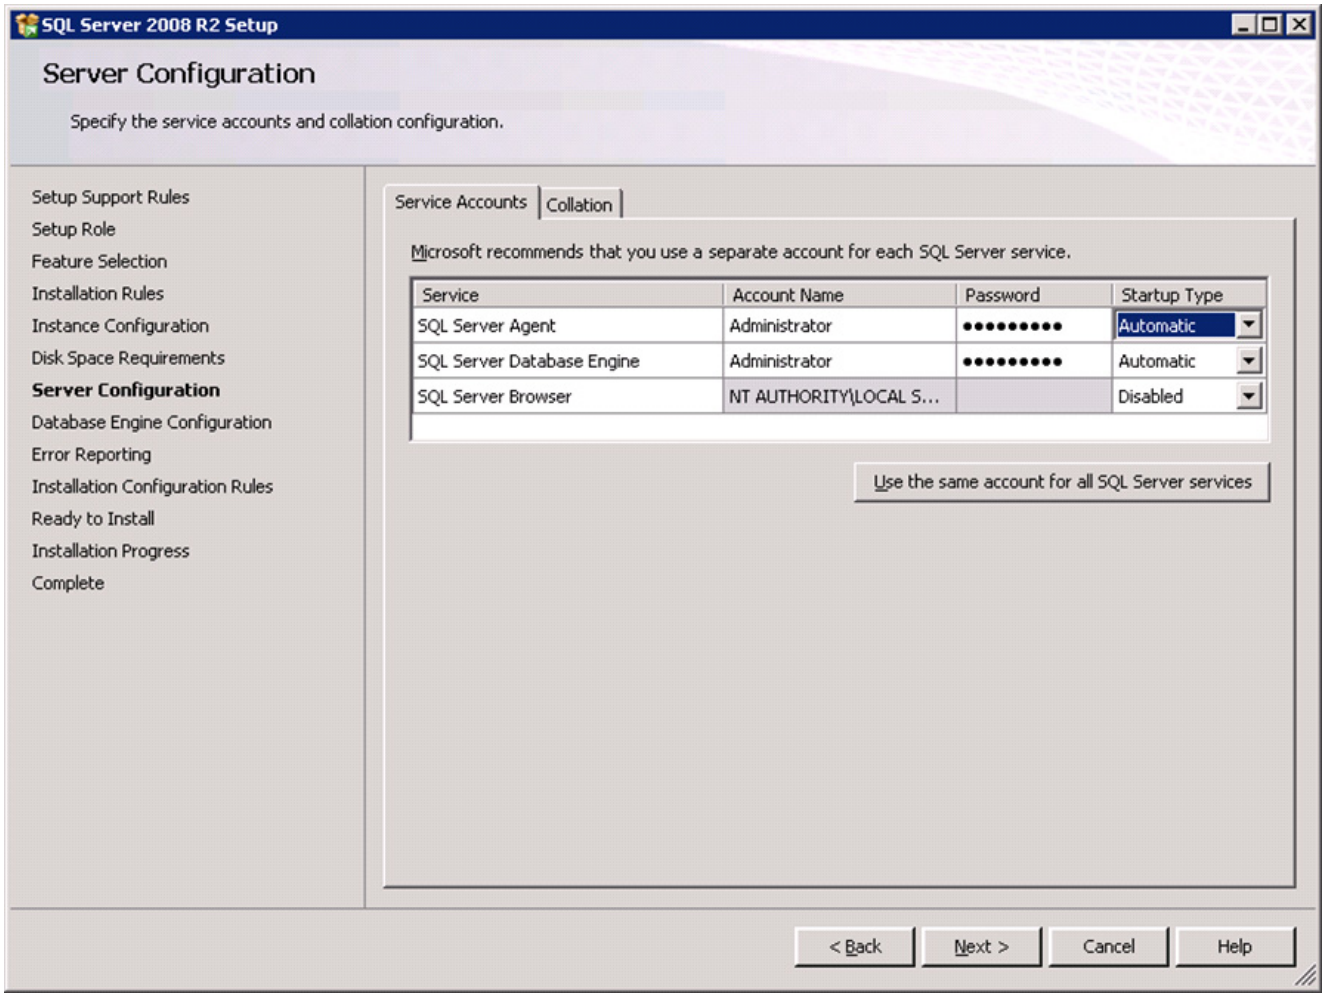

36. Choose Mixed Mode (SQL Server Authentication and Windows Authentication). Enter and confirm the password for the SQL Server system administrator (sa) account, click Add Current User, and Click **Next**.

Figure 64 Configuring Database Engine

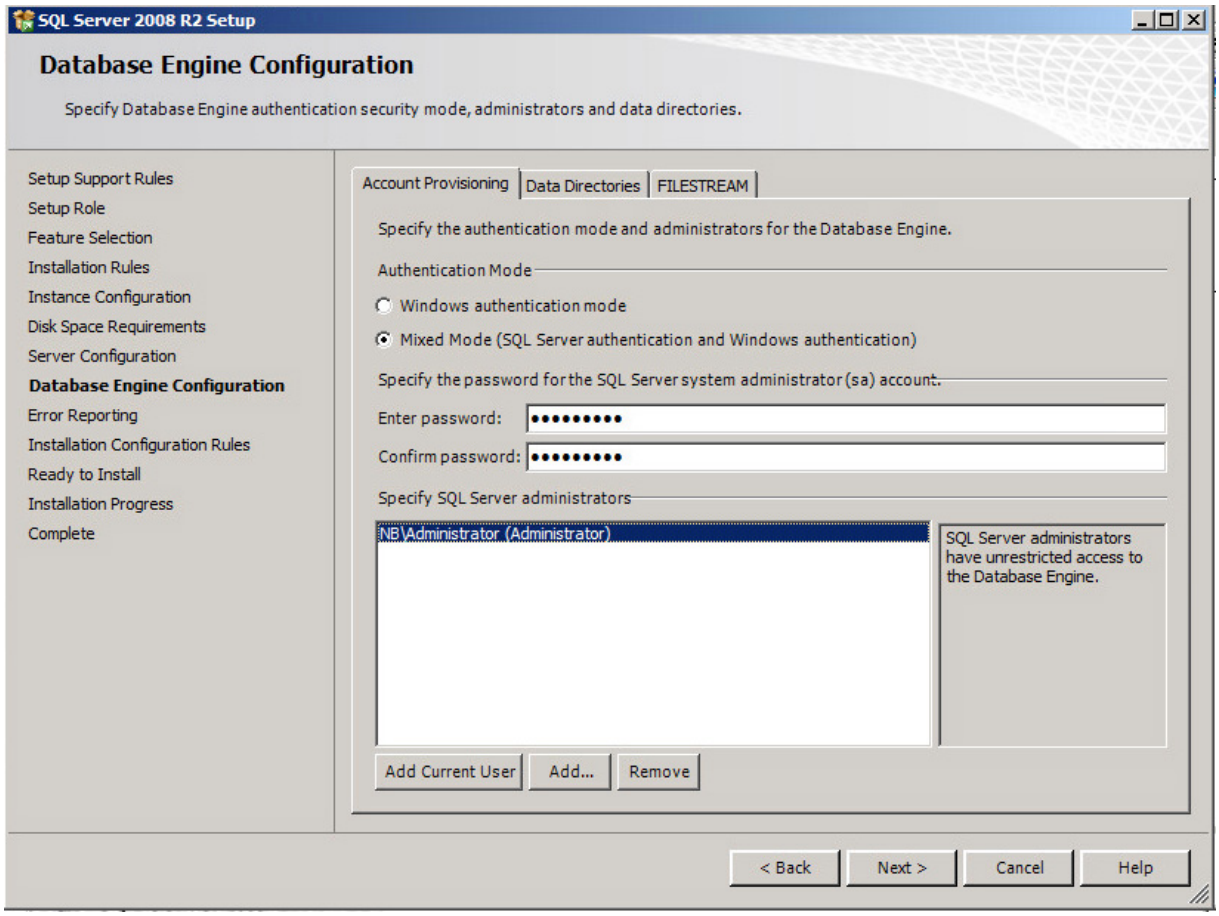

37. Choose whether to send error reports to Microsoft. Click **Next**.
38. Click **Next**.
39. Click **Install**.
40. After the installation is complete, click **Close** to close the SQL Server installer.
41. Close the SQL Server Installation Center.
42. Install all available Microsoft Windows updates by navigating to **Start > All Programs > Windows Update**.
43. Choose **Start > All Programs > Microsoft SQL Server 2008 R2 > SQL Server Management Studio** to open the SQL Server Management Studio.
44. Under Server Name, choose the local machine name. Under Authentication, choose SQL Server Authentication. Enter sa in the Login field and enter the sa password. Click **Connect**.
45. Click **New Query**.
46. Run the following script, substituting the **vpxuser** password for **<Password>**:

```
use [master]
go
CREATE DATABASE [VCDB] ON PRIMARY
(NAME = N'vcdb', FILENAME = N'C:\VCDB.mdf', SIZE = 2000KB, FILEGROWTH = 10% )
LOG ON
(NAME = N'vcdb_log', FILENAME = N'C:\VCDB.ldf', SIZE = 1000KB, FILEGROWTH = 10%)
```

```
COLLATE SQL_Latin1_General_CP1_CI_AS
go
use VCDB
go
sp_addlogin @loginame=[vpxuser], @passwd=N'<Password>', @defdb='VCDB',
@deflanguage='us_english'
go
ALTER LOGIN [vpxuser] WITH CHECK_POLICY = OFF
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
use MSDB
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
use VCDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use MSDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
```

Note Figure 65 illustrates the script.

Figure 65 Executing the Query

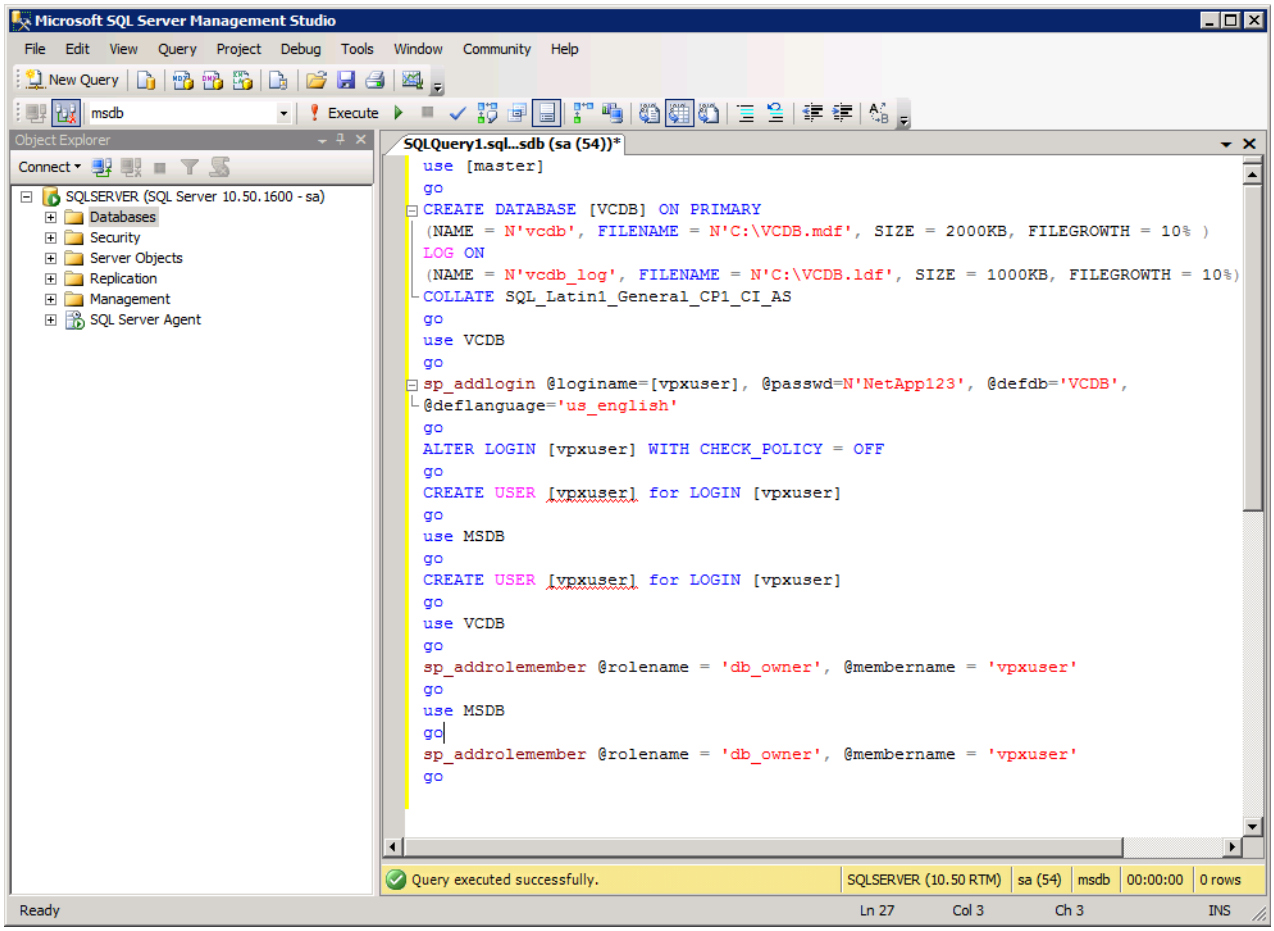

47. **Click** Execute and verify that the query executes successfully.
48. Close Microsoft SQL Server Management Studio.
49. Disconnect the Microsoft SQL Server 2008 R2 ISO from the SQL Server VM.

Build and Set up VMware vCenter VM

Build VMware vCenter VM

To build the VMware vCenter VM, follow these steps:

1. Using the instructions for building a SQL Server VM provided in "Build and Set up VMware vCenter VM" section on page 126, build a VMware vCenter VM with the following configuration in the <<var_ib-mgmt_vlan_id>> VLAN:
   - 4GB RAM
   - Two CPUs
   - One virtual network interface
2. Start the VM, install VMware Tools, and assign an IP address and host name to it in the Active Directory domain.

Set up VMware vCenter VM

To set up the newly built VMware vCenter VM, follow these steps:

1. Log in to the vCenter VM as the FlexPod admin user and open Server Manager.
2. Expand Features and click **Add Features**.
3. Expand .NET Framework 3.5.1 Features and choose only .NET Framework 3.5.1.
4. Click **Next**.
5. Click **Install**.
6. Click **Close** to close the Add Features wizard.
7. Close **Server Manager**.
8. Download and install the client components of the Microsoft SQL Server 2008 R2 Native Client from the Microsoft Download Center.
9. Create the vCenter database data source name (DSN). Choose **Start > Administrative Tools > Data Sources (ODBC)** to open Data Sources (ODBC).
10. Choose the System DSN tab.
11. Click **Add**.
12. Choose SQL Server Native Client 10.0 and click **Finish**.
13. Name the data source VCDB. In the Server field, enter the IP address of the vCenter SQL server. Click **Next**.

Figure 66 Creating New Data Source

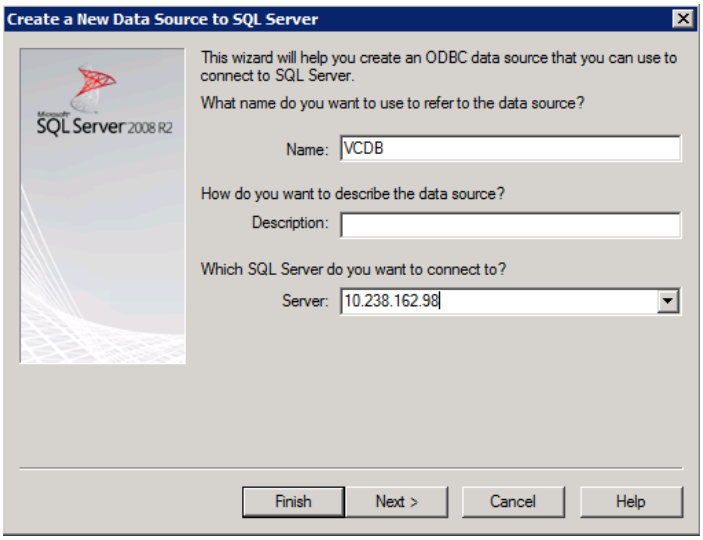

14. Choose With SQL Server authentication using a login ID and password entered by the user. Enter vpxuser as the login ID and the vpxuser password. Click **Next**.

Figure 67 Login Credentials for SQL Server

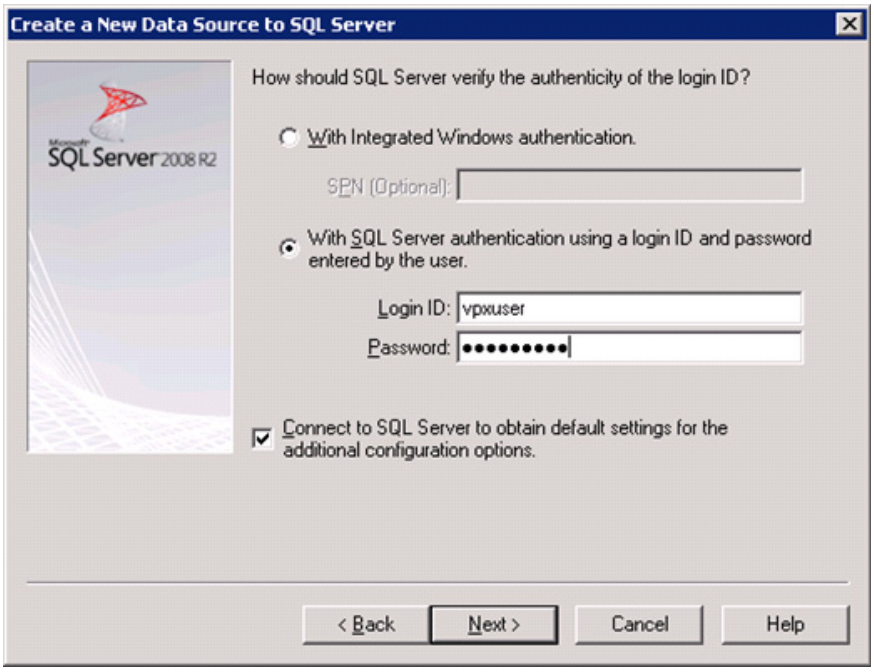

15. Choose Change the Default Database To and choose VCDB from the list. Click **Next**.

Figure 68 Setting Up the Database

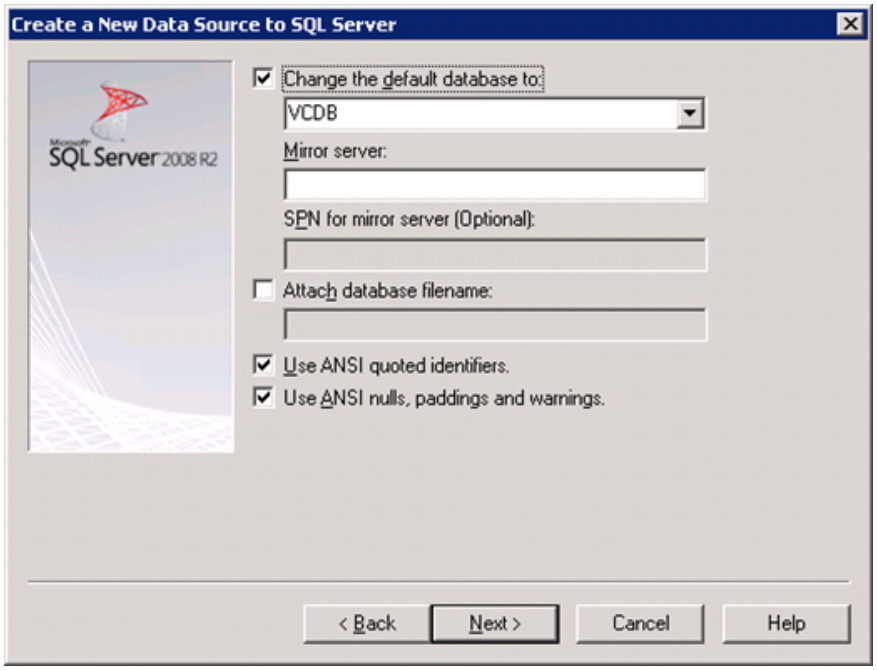

16. Click **Finish**.
17. Click **Test Data Source**. Verify that the test completes successfully.

Figure 69 Verifying ODBC Data Source

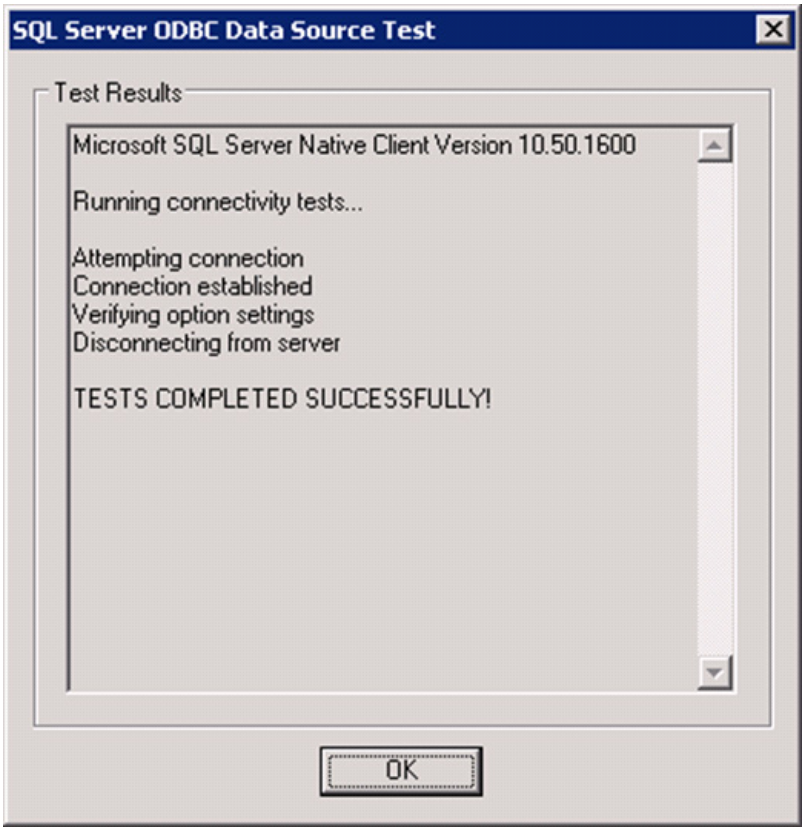

18. Click **OK** and then click **OK** again.
19. Click **OK** to close the ODBC Data Source Administrator window.
20. Install all available Microsoft Windows updates by navigating to **Start > All Programs > Windows Update**.

Note A restart might be required.

Install VMware vCenter Server

vCenter Server VM

To install vCenter Server on the vCenter Server VM, follow these steps:

1. In the vCenter Server VMware console, click the ninth button (CD with a wrench) to map the VMware vCenter ISO and choose Connect to ISO Image on Local Disk.
2. Navigate to the VMware vCenter 5.1Update1 (VIMSetup) ISO, select it, and click **Open**.
3. In the dialog box, click **Run autorun.exe**.
4. In the VMware vCenter Installer window, make sure that VMware vCenter Simple Install is selected and click **Install**.

Figure 70 Installing VMware vSphere

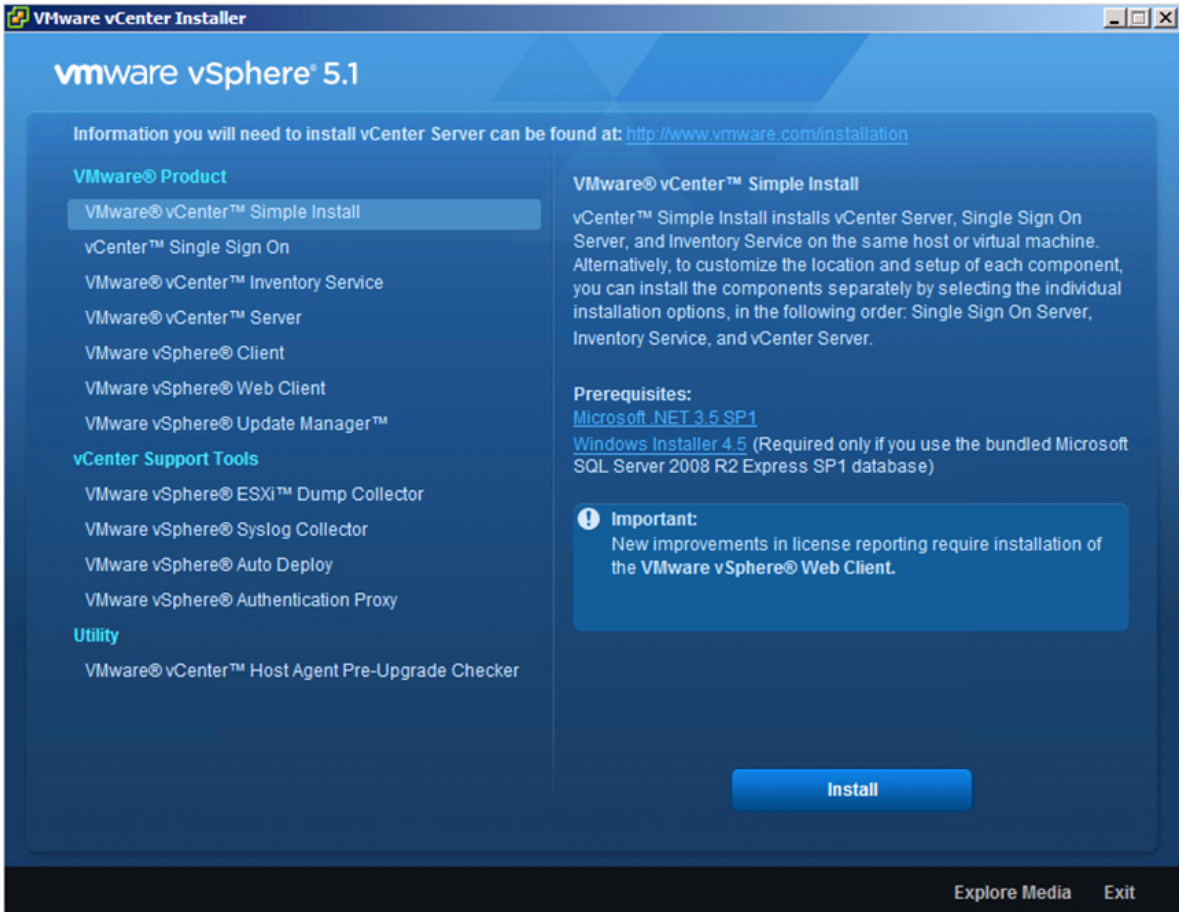

5. Click **Next** to install vCenter Single Sign On.
6. Click **Next**.
7. Accept the terms of the license agreement and click **Next**.
8. Enter and confirm <<var_password>> for admin@System-Domain.
9. Click **Next**.
10. Keep the radio button checked to install a local Microsoft SQL Server 2008 R2 Express instance and click **Next**.
11. Enter and confirm <<var_password>> for both user names. Click **Next**.

Figure 71 Setting Database User Information

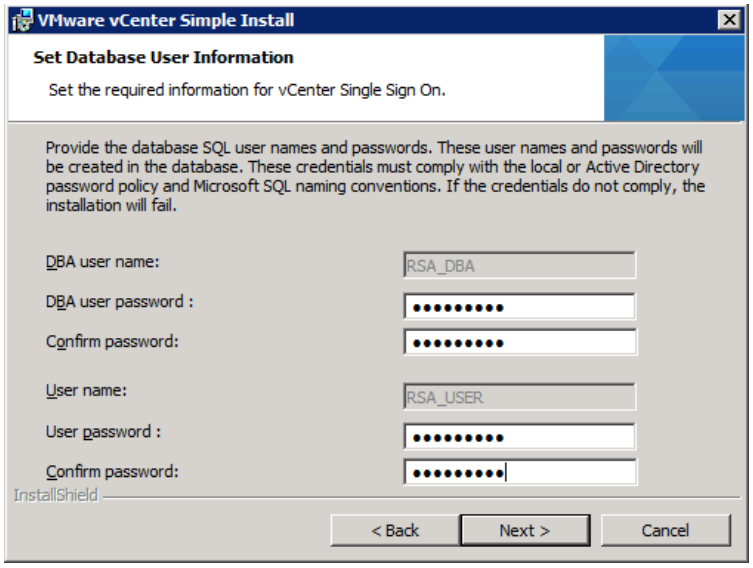

12. Verify the vCenter VM FQDN and click **Next**.
13. Leave Use network service account selected and click **Next**.
14. Click **Next** to select the default destination folder.
15. Click **Next** to select the default HTTPS port.
16. Click **Install** to install vCenter Single Sign On.
17. Enter the vCenter 5.1Update1 license key and click **Next**.
18. Choose Use an Existing Supported Database. Choose VCDB from the Data Source Name list and click **Next**.

Figure 72 Selecting Database for VMware vSphere

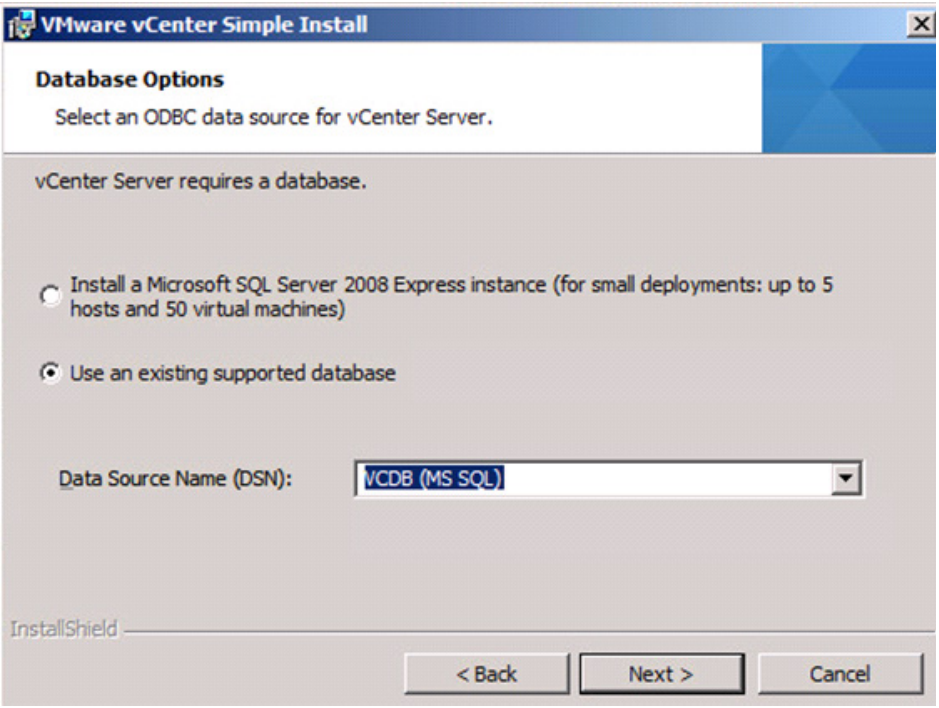

19. Enter the vpxuser password and click **Next**.

Figure 73 Entering Database Login Credentials

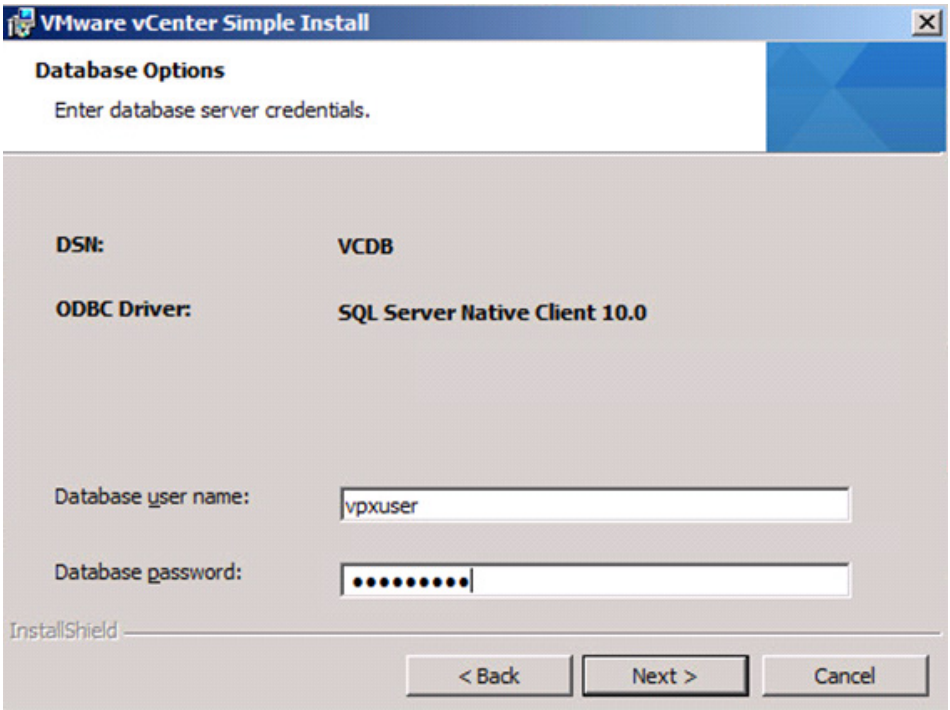

20. Review the warning and click **OK**.
21. Click **Next** to use the SYSTEM Account.
22. Click **Next** to accept the default ports.
23. Choose the appropriate inventory size. Click **Next**.
24. Click **Install**.
25. Click **Finish**.
26. Click **OK** to confirm the installation.
27. Click **Exit** in the VMware vCenter Installer window.
28. Disconnect the VMware vCenter ISO from the vCenter VM.
29. Install all available Microsoft Windows updates by navigating to **Start > All Programs > Windows Updates**.

Note A restart might be required.

Set up vCenter Server

vCenter Server VM

To set up vCenter Server on the vCenter Server VM, follow these steps:

1. Using the vSphere Client, log in to the newly created vCenter Server as the FlexPod admin user.
2. Click **Create a data center**.
3. Enter FlexPod_DC_1 as the data center name.
4. Right-click the newly created FlexPod_DC_1 data center and choose **New Cluster**.
5. Name the cluster FlexPod_Management and click the check boxes for Turn On vSphere HA and Turn on vSphere DRS. Click **Next**.

Figure 74 Selecting Cluster Features

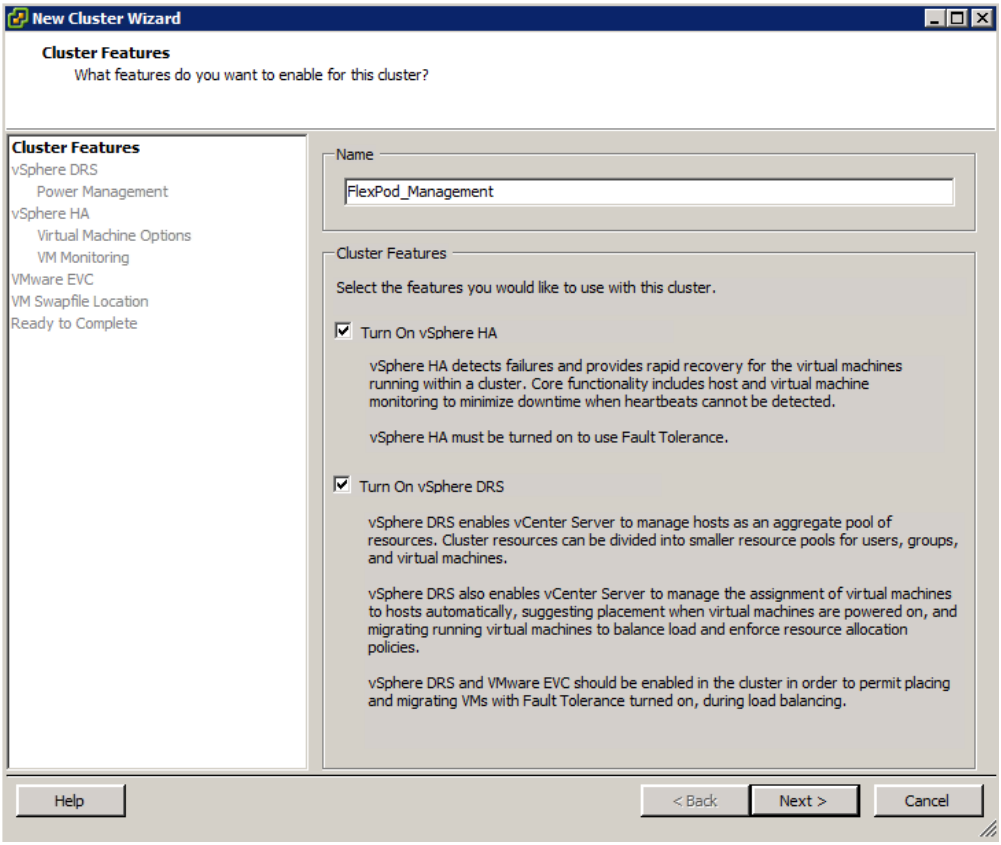

6. Accept the defaults for vSphere DRS. Click **Next**.
7. Accept the defaults for Power Management. Click **Next**.
8. Accept the defaults for vSphere HA. Click **Next**.
9. Accept the defaults for Virtual Machine Options. Click **Next**.
10. Accept the defaults for VM Monitoring. Click **Next**.
11. Accept the defaults for VMware EVC. Click **Next**.

Note If we are mixing UCS B or C-Series M2 and M3 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, see Enhanced vMotion Compatibility (EVC) Processor Support at: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003212

12. Choose Store the swapfile in the datastore specified by the host. Click **Next**.
13. Click **Finish**.
14. Right-click the newly created FlexPod_Management cluster and choose **Add Host**.
15. In the Host field, enter either the IP address or the host name of the VM-Host-Infra_01 host. Enter root as the user name and the root password for this host. Click **Next**.
16. Click **Yes**.

17. Click **Next**.
18. Choose Assign a New License Key to the Host. Click Enter Key and enter a vSphere license key. Click **OK**, and then click **Next**.
19. Click **Next**.
20. Click **Next**.
21. Click **Finish**. VM-Host-Infra-01 is added to the cluster.
22. Repeat this procedure to add VM-Host-Infra-02 to the cluster.

FlexPod Cisco Nexus 1110-X and 1000V vSphere

The following sections provide detailed procedures for installing a pair of high-availability (HA) Cisco Nexus 1110-X Virtual Services Appliances (VSAs) in a FlexPod configuration. Primary and standby Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) are installed on the 1110-Xs. By the end of this section, a Cisco Nexus 1000V distributed virtual switch (DVS) will be provisioned. This procedure assumes that the Cisco Nexus 1000V software version 4.2(1)SV2(2.1a) has been downloaded from www.cisco.com and expanded. This procedure also assumes that VMware vSphere 5.1Update1 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, follow these 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 check box 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, follow these 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. Check the Enabled check box for Serial over LAN Properties.
6. From the Baud Rate drop-down menu, choose 9600 bps.
7. Click **Save Changes**.

Figure 75 Configuring Serial Over LAN on Cisco Nexus 1110-Xs

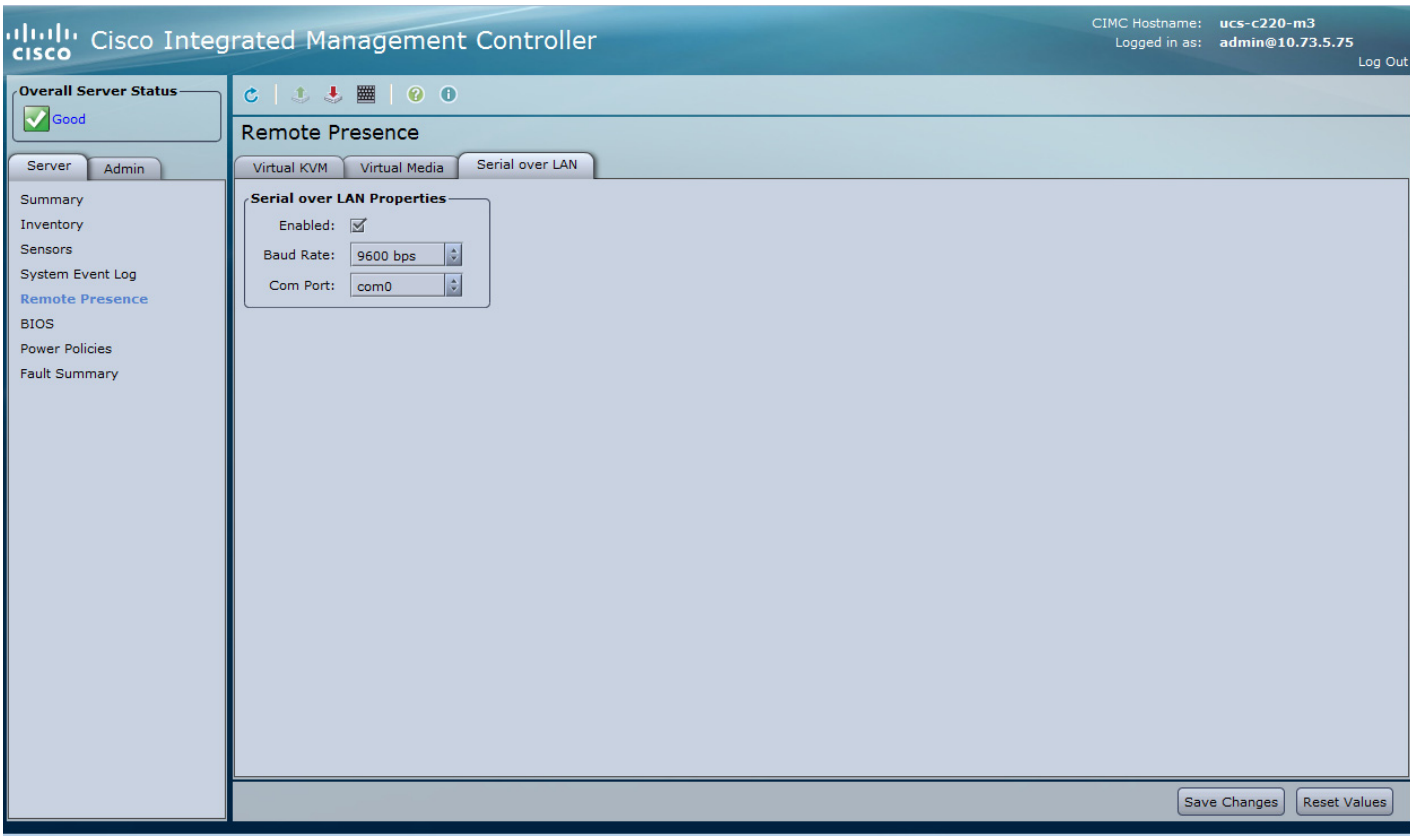

8. Log out of the CIMC Web interface.
9. Use an SSH client to connect to <<var_cimc_ip>> with the default CIMC user name and password.
10. Run connect host.

Configure Cisco Nexus 1110-X Virtual Appliances

Cisco Nexus 1110-X A

To configure Cisco Nexus 1110-X A, follow these 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 network-uplink type <1-5>:
  1. Ports 1-2 carry all management, control and data vlans
  2. Ports 1-2  management and control, ports 3-6 data
  3. Ports 1-2 management, ports 3-6 control and data
  4. Ports 1-2 management, ports 3-4 control, ports 5-6 data
  5. Flexible
Choose 1
Enter control VLAN <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
Enter the domain id<1-4095>: <<var_1110x_domain_id>>
Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>
Would you like to enter the basic 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): Enter
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 (dsa/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_global_ntp_server_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, follow these 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>>
```

Note This is the same password that you entered on the primary Cisco Nexus 1110-X.

```
Enter HA role[primary/secondary]: secondary
Enter network-uplink type <1-5>:
  1. Ports 1-2 carry all management, control and data vlans
  2. Ports 1-2  management and control, ports 3-6 data
```

```
  3. Ports 1-2 management, ports 3-6 control and data
  4. Ports 1-2 management, ports 3-4 control, ports 5-6 data
  5. Flexible
Choose 1
Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
Enter the domain id<1-4095>: <<var_1110x_domain_id>>
```

Note This is the same unique Cisco Nexus 1110 domain ID entered on Cisco Nexus 1110-X A.

```
Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>
```

2. 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, follow these steps:

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 (nexus-1000v.4.2.1.SV2.2.1a.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>>/nexus-1000v.4.2.1.SV2.2.1a.iso /repository/.

```
virtual-service-blade-type new nexus-1000v.4.2.1.SV2.2.1a.iso
interface control vlan <<var_pkt-ctrl_vlan_id>>
interface packet vlan <<var_pkt-ctrl_vlan_id>>
enable primary
Enter vsb image:[nexus-1000v.4.2.1.SV2.2.1a.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
```

Note This domain ID should be different from 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>>
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.

Set up the Secondary Cisco Nexus 1000V VSM

To set up the secondary Cisco Nexus 1000V VSM on Cisco Nexus 1110-X B, follow these steps in two subsections:

Cisco Nexus 1110-X A

Run system switchover to activate Cisco Nexus 1110-X B.

Cisco Nexus 1110-X B

1. Log in to Cisco Nexus 1110-X B as the admin user.

```
config t
virtual-service-blade VSM-1
enable secondary
Enter vsb image: [nexus-1000v.4.2.1.SV2.2.1a.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>>
```

2. 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.

```
copy run start
```

3. Run system switchover on Cisco Nexus 1110-X B to activate Cisco Nexus 1110-X A. This causes Cisco Nexus 1110-X B to reboot.

Install Virtual Ethernet Module on Each ESXi Host

vCenter Server VM

To install the Virtual Ethernet Module (VEM) on the ESXi hosts, follow these steps:

1. Launch a Web browser to http://<<var_vsm_mgmt_ip>>.
2. Right-click the cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib hyperlink and choose Save target as.
3. Save the file as cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib, type All Files, on the Desktop of the management workstation.
4. From the main window in the vSphere Client connected to vCenter, choose the first server in the list under the FlexPod Management cluster.
5. Click the **Summary** tab.
6. Under Storage on the right, right-click infra_datastore_1 and choose Browse Datastore.
7. Choose the root folder (/) and click the third button at the top to add a folder.
8. Name the folder VEM and click **OK**.
9. On the left, choose the VEM folder.
10. Click the fourth button at the top and choose Upload File.
11. Navigate to the cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib file and click **Open**.
12. Click **Yes**. The VEM file should now appear in the VEM folder in the datastore.

13. Open the VMware vSphere CLI command prompt.
14. For each ESXi host in the VMware vSphere CLI, run the following command:

```
esxcli -s <Host Server IP> -u root -p <Root Password> software vib install -v
/vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib
```

Figure 76 Executing Command for Each ESXi Hosts

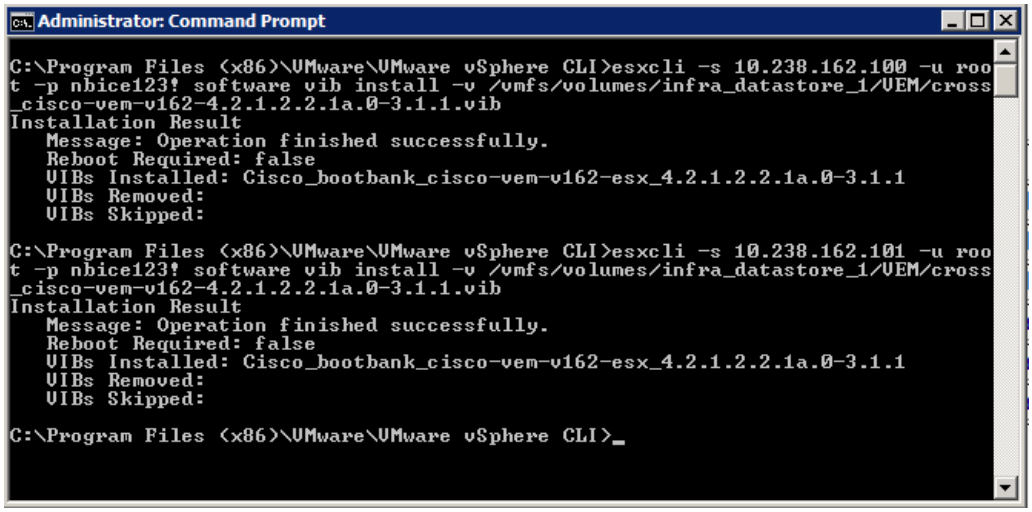

Register Cisco Nexus 1000V as a vCenter Plug-in

To register the Cisco Nexus 1000V as a vCenter plug-in, follow these steps:

1. Using a Web browser, navigate to the <<var_vsm_mgmt_ip>> using http://<<var_vsm_mgmt_ip>>.
2. Right-click the cisco_nexus_1000v_extension.xml hyperlink and choose Save target as.
3. Save the XML file to the local desktop.
4. In the vSphere Client connected to vCenter, choose **Plug-ins** > **Manage Plug-ins**.
5. Right-click the white space in the window and choose **New Plug-in**.
6. Browse to the desktop and choose the cisco_nexus_1000v_extension.xml document that was previously saved. Click **Open**.
7. Click **Register Plug-in**.
8. Click **Ignore**.
9. Click **OK**.
10. The Cisco_Nexus_1000V should now appear in the list of available plug-ins.
11. Click **Close** to close the Plug-in Manager.

Perform Base Configuration of the Primary VSM

To perform the base configuration of the primary VSM, follow these 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
```

```
svs connection vCenter
protocol vmware-vim
remote ip address <<var_vcenter_server_ip>> port 80
vmware dvs datacenter-name FlexPod_DC_1
connect
exit
ntp server <<var_global_ntp_server_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-VLAN
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
vlan <<var_native_vlan_id>>
name Native-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_mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_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
exit
copy run start
```

Migrate Networking Components for ESXi Hosts to Cisco Nexus 1000V

vSphere Client Connect to vCenter

To migrate the networking components for the ESXi hosts to the Cisco Nexus 1000V, follow these steps:

1. In the VMware vSphere Client connected to vCenter, choose **Home** > **Networking**.
2. Expand the vCenter, DataCenter, and Cisco Nexus 1000V folders. choose the Cisco Nexus 1000V switch.
3. Under Basic Tasks for the vSphere distributed switch, choose Add a Host.
4. For both hosts, choose vmnic1 and choose the system-uplink Uplink port group. Click **Next**.

Figure 77 Adding Host to the vSphere Distributed Switch

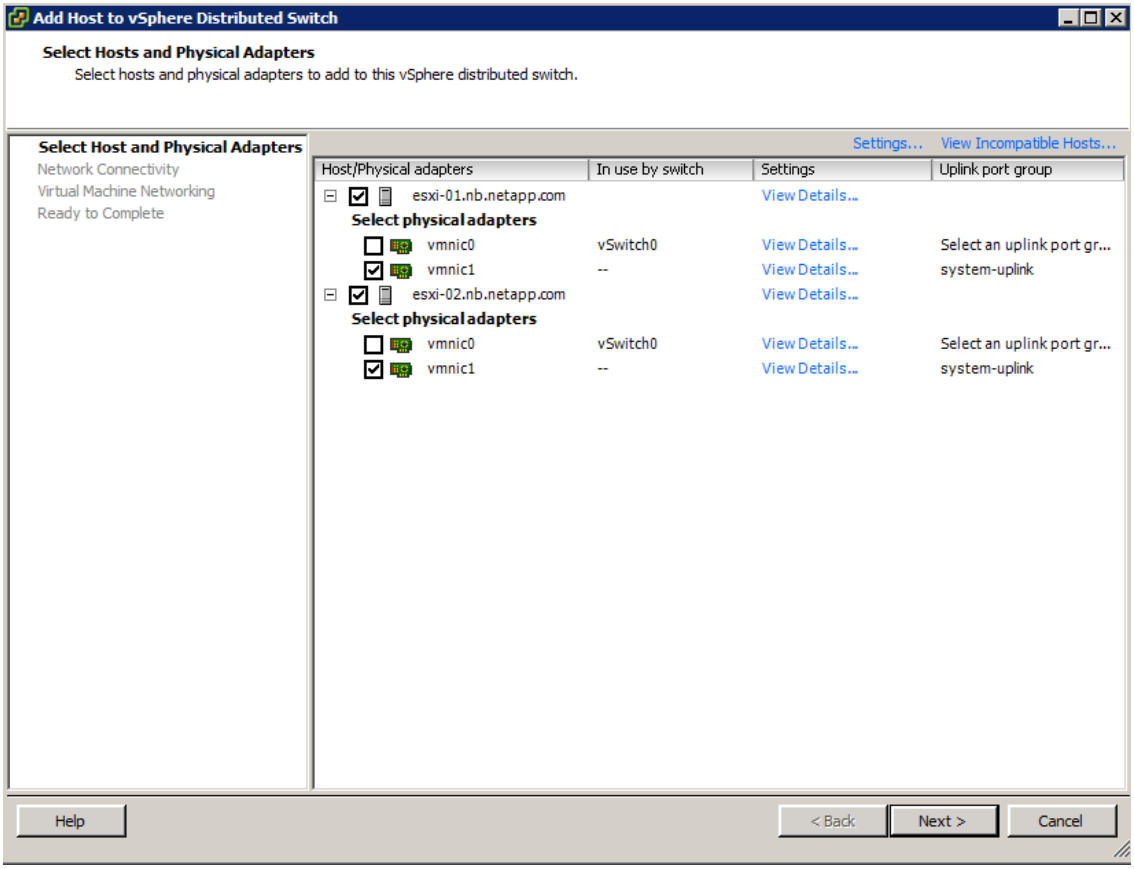

5. For all VMkernel ports, choose the appropriate Destination Port Group from the Cisco Nexus1000V, making sure to choose the “n1kv-L3” destination port group for the MGMT VMkernel ports. Click **Next**.

Figure 78 Assigning Port Groups

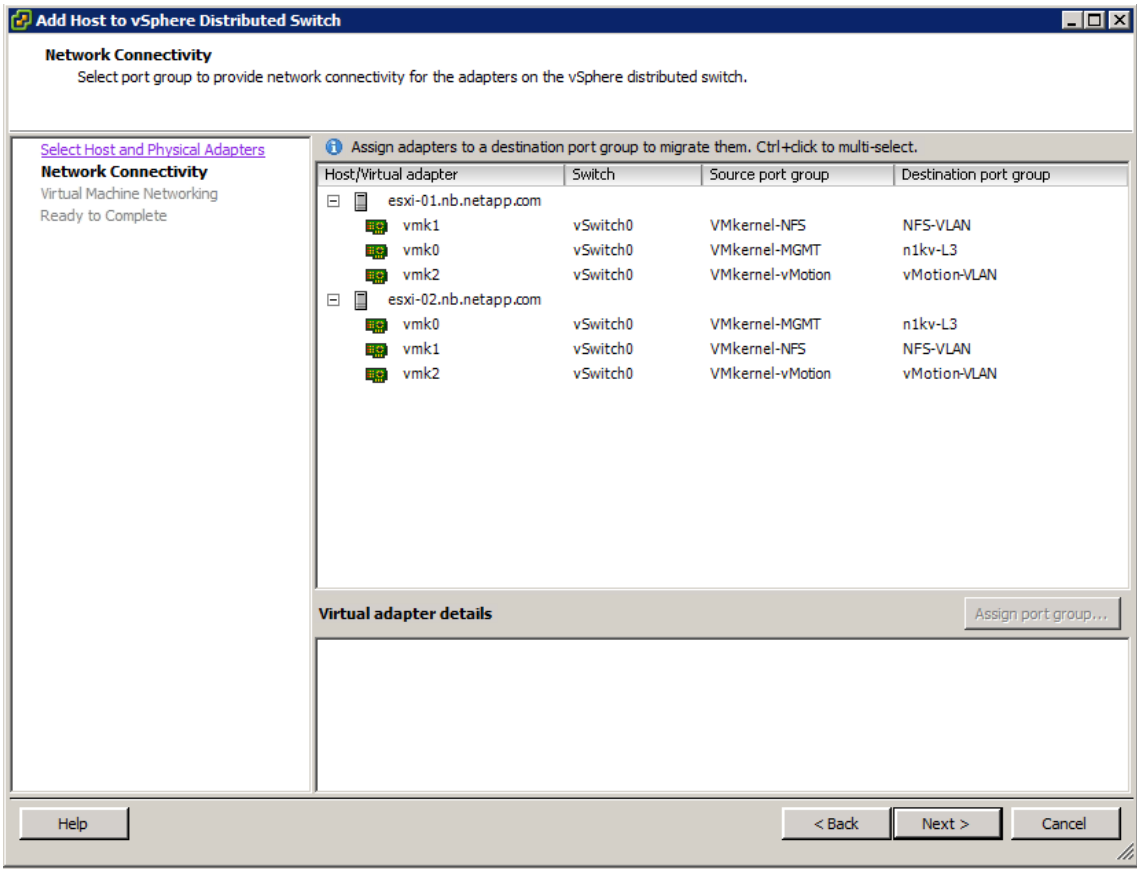

6. Choose the Migrate Virtual Machine Networking check box. Expand each VM and choose the port groups for migration individually. Click **Next**.

Figure 79 Choosing Port Groups for Each VMs

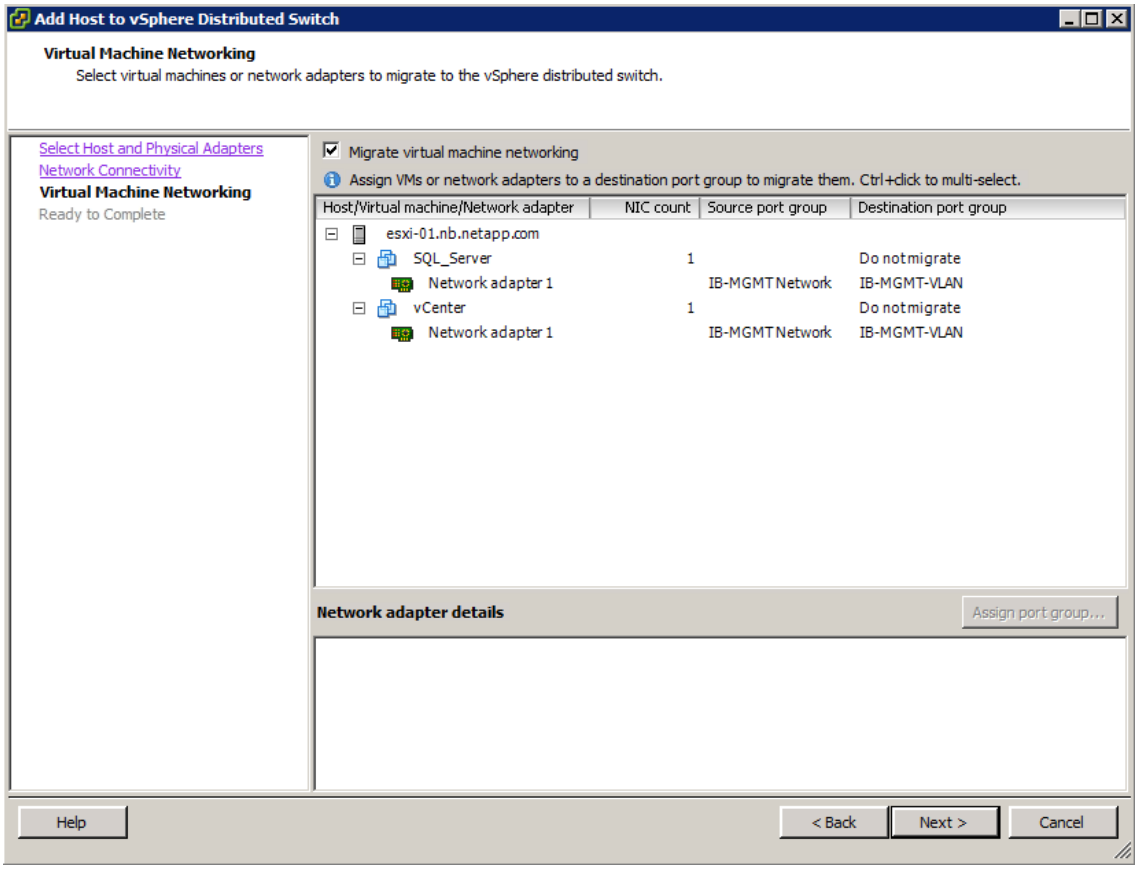

7. Click **Finish**. Wait for the migration process to complete.
8. In the vSphere Client window, choose **Home > Hosts and Clusters**.
9. Choose the first ESXi host and click the **Configuration** tab. In the Hardware box, choose Networking.
10. Make sure that vSphere Standard Switch is selected at the top next to View. vSwitch0 should not have any active VMkernel or VM Network ports on it. On the upper-right side of vSwitch0, click **Remove**.
11. Click **Yes**.
12. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View.
13. Click **Manage Physical Adapters**.
14. Scroll down to the system-uplink box and click **Add NIC**.
15. choose vmnic0 and click **OK**.
16. Click **OK** to close the Manage Physical Adapters window. Two system uplinks should now be present.
17. choose the second ESXi host and click the **Configuration** tab. In the Hardware field, click **Networking**.

18. Make sure vSphere Standard Switch is selected at the top next to View. vSwitch0 should have no active VMkernel or VM Network ports on it. On the upper-right side of vSwitch0, click **Remove**.
19. Click **Yes**.
20. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View.
21. Click **Manage Physical Adapters**.
22. Scroll down to the system-uplink box and click **Add NIC**.
23. choose vmnic0 and click **OK**.
24. Click **OK** to close the Manage Physical Adapters window. Two system-uplinks should now be present.
25. 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.

Figure 80 Verifying Interfaces and Port Channels

```
10.238.162.203 - PuTTY
Cisco-Nexus-VSM-B# show interface status

--------------------------------------------------------------------------------
Port          Name               Status    Vlan      Duplex  Speed   Type
--------------------------------------------------------------------------------
mgmt0         --                 up        routed    full    1000    --
Eth3/1        --                 up        trunk     full    10G     --
Eth3/2        --                 up        trunk     full    10G     --
Eth4/1        --                 up        trunk     full    10G     --
Eth4/2        --                 up        trunk     full    10G     --
Po1           --                 up        trunk     full    10G     --
Po2           --                 up        trunk     full    10G     --
Veth1         VMware VMkernel, v up        3175      auto    auto    --
Veth2         VMware VMkernel, v up        3172      auto    auto    --
Veth3         VMware VMkernel, v up        3173      auto    auto    --
Veth4         VMware VMkernel, v up        3172      auto    auto    --
Veth5         VMware VMkernel, v up        3175      auto    auto    --
Veth6         VMware VMkernel, v up        3173      auto    auto    --
Veth7         SQL_Server, Networ up        3175      auto    auto    --
Veth8         vCenter, Network A up        3175      auto    auto    --
control0      --                 up        routed    full    1000    --
Cisco-Nexus-VSM-B#
```

26. Run show module and verify that the two ESXi hosts are present as modules.

Figure 81 Verifying the ESXi Hosts are Shown as Modules

```
10.238.162.203 - PuTTY
Cisco-Nexus-VSM-B# show module
Mod  Ports  Module-Type                        Model               Status
---  -----  --------------------------------   ------------------  ------------
1    0      Virtual Supervisor Module          Nexus1000V          active *
2    0      Virtual Supervisor Module          Nexus1000V          ha-standby
3    332    Virtual Ethernet Module            NA                  ok
4    332    Virtual Ethernet Module            NA                  ok

Mod  Sw                  Hw
---  ------------------  ------------------------------------------------
1    4.2(1)SV2(2.1a)     0.0
2    4.2(1)SV2(2.1a)     0.0
3    4.2(1)SV2(2.1a)     VMware ESXi 5.1.0 Releasebuild-1065491 (3.1)
4    4.2(1)SV2(2.1a)     VMware ESXi 5.1.0 Releasebuild-1065491 (3.1)

Mod  Server-IP        Server-UUID                           Server-Name
---  ---------------  ------------------------------------  --------------------
1    10.238.162.203   NA                                    NA
2    10.238.162.203   NA                                    NA
3    10.238.162.101   ae007b7c-970f-e311-0000-00000000001f  esxi-02.nb.netapp.co
m
--More--
```

27. Run **copy run start.**
28. Type **exit** two times to log out of the Cisco Nexus 1000v.

FlexPod Management Tool Setup

NetApp Virtual Storage Console (VSC) 4.2.1 Deployment Procedure

VSC 4.2.1 Preinstallation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.2P4:

- Protocol licenses (NFS and FCP)
- FlexClone (for provisioning and cloning only)
- SnapRestore (for backup and recovery)
- SnapManager suite

Install VSC 4.2.1

To install the VSC 4.2.1 software, follow these steps:

1. Using the instructions given in "Build and Set up VMware vCenter VM" section on page 126, build a VSC and an OnCommand virtual machine with 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. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain. Install the current version of Adobe Flash Player on the VM. Install all Windows updates on the VM.

2. Log in to the VSC and OnCommand VM as the FlexPod admin user.
3. Download the x64 version of the Virtual Storage Console 4.2.1 at: Virtual Storage Console 4.2.1 from the NetApp Support site.
4. Right-click the file downloaded in step 3 and choose Run As Administrator.
5. On the Installation wizard Welcome page, click **Next**.
6. Check the checkbox to accept the message, click **Next**.
7. Choose the backup and recovery capability. Click **Next**.

Note The backup and recovery capability requires an additional license.

Figure 82 Selecting the Desired Capabilities of VSC 4.2.1

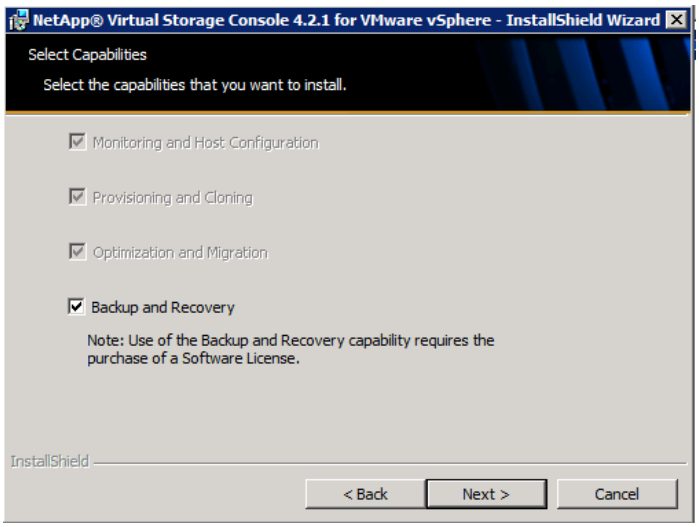

8. Click **Next** to accept the default installation location.

Figure 83 Installation Location

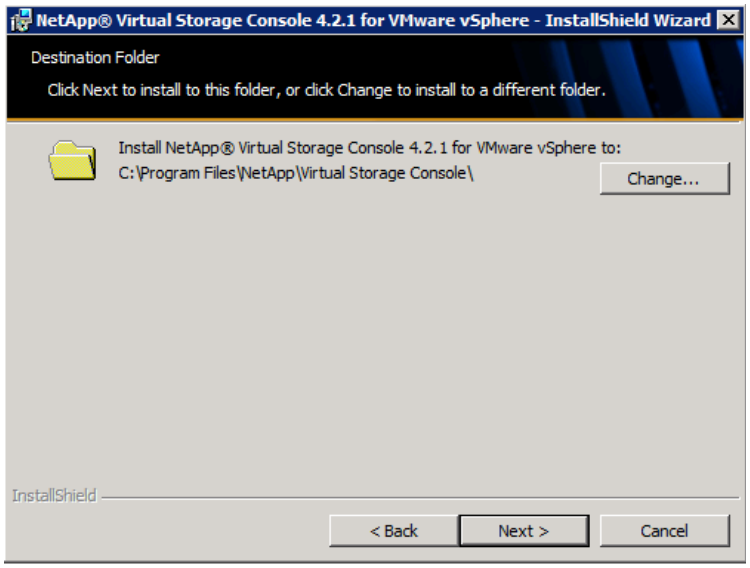

9. Click **Install**.

Figure 84 Ready to Install VSC 4.2.1

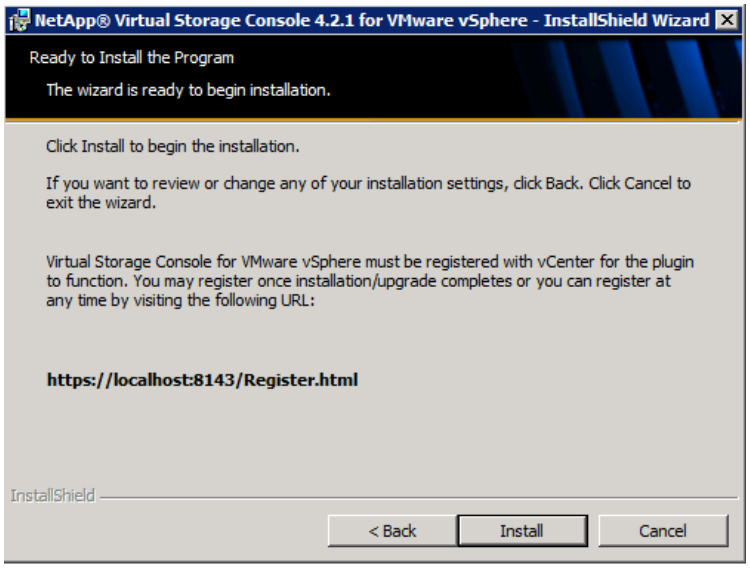

10. Click **Finish**.

Register VSC with vCenter Server

To register the VSC with the vCenter Server, follow these steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete.
2. Click **Continue** to this website (not recommended).

3. In the Plug-in Service Information section, choose 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), and user password for the vCenter Server. Click **Register** to complete the registration.

Figure 85 Registering VSC with vCenter Server

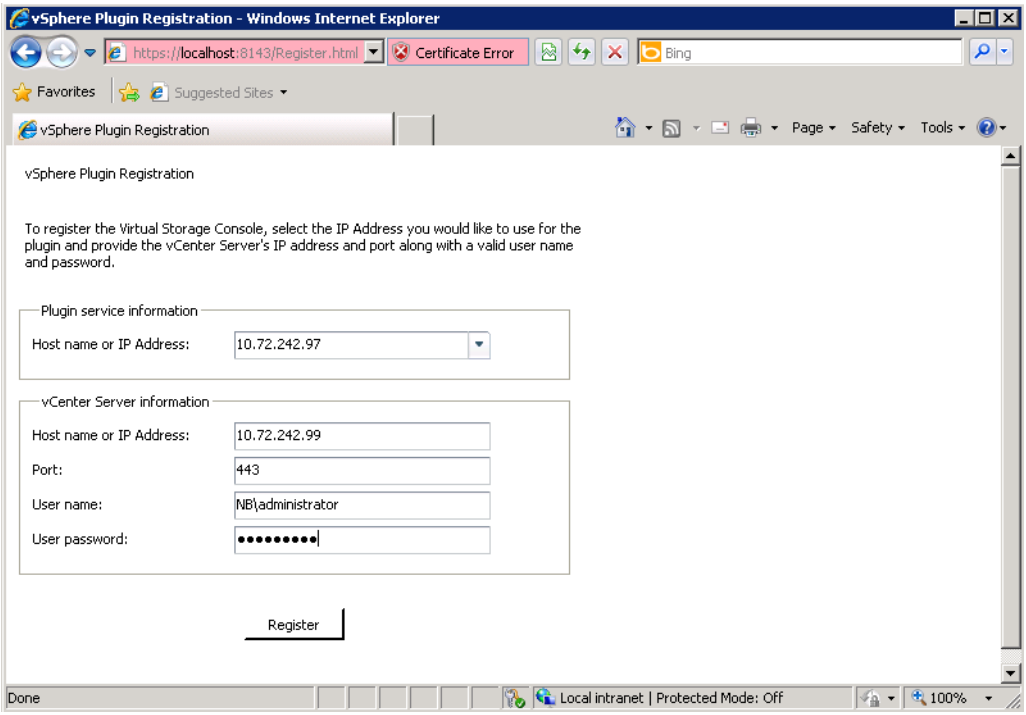

Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, follow these steps:

1. Using the vSphere Client, log in to the vCenter Server as FlexPod admin user. If the vSphere Client was previously opened, close it and then reopen it.
2. If an SSL Certificate warning from the VSC is displayed, select the checkbox to Install the certificate and Click **Ignore**.
3. Click the **Home** tab in the left side of the vSphere Client window.
4. Under Solutions and Applications, click **NetApp**.

Figure 86 Adding Storage Resource

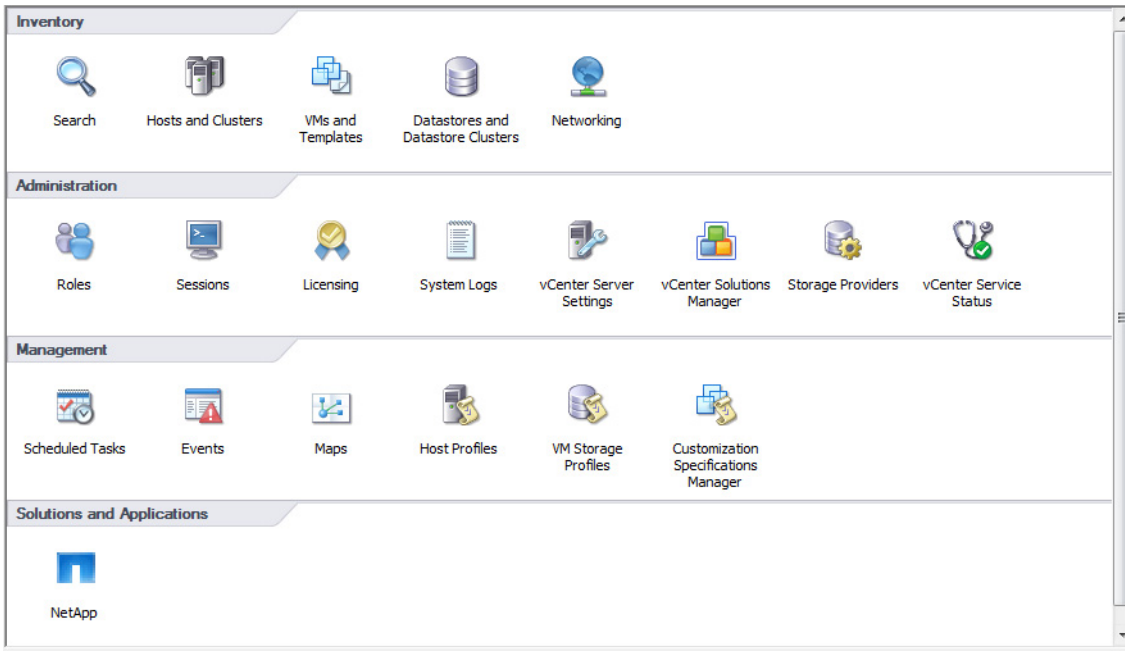

5. Click **Yes** when the security certificate warning appears. To view the certificate, click **View Certificate**.

Figure 87 Certificate Details

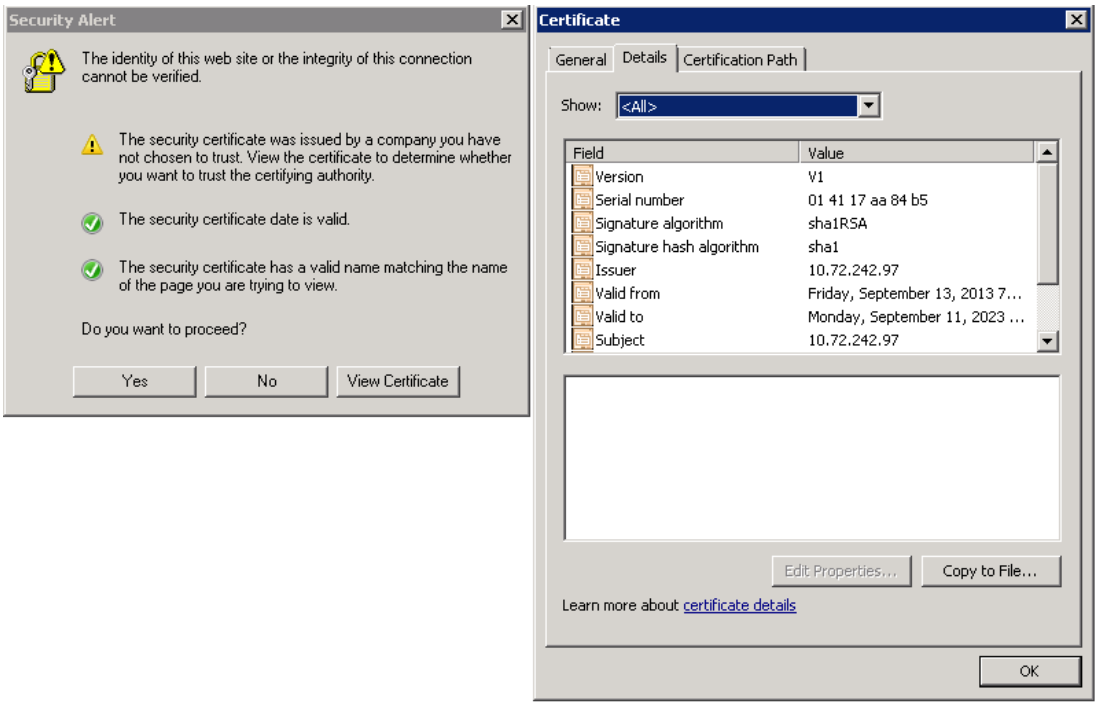

6. In the navigation pane, choose Monitoring and Host Configuration if it is not selected by default

7. In the list of storage controllers, right-click the first controller listed and choose Modify Credentials.

Figure 88 vSphere Client Showing Storage Controllers

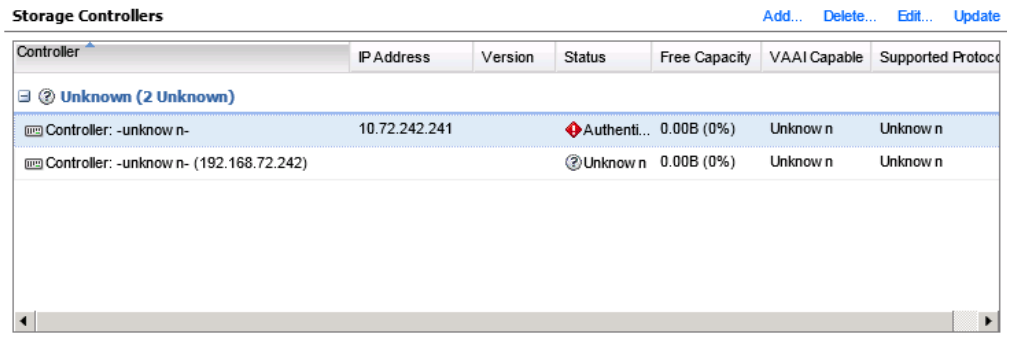

8. Enter the storage management IP address in the Management IP Address field. Enter root for the user name, and the admin password for the Password. Make sure that Use SSL is selected. Click **OK**.
9. Click **OK** to accept the controller privileges.
10. Add any additional storage controllers that you may want to manage using VSC by using the **Add** option in the upper-right side of the window.

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, follow these steps:

1. Choose individual or multiple ESXi hosts.
2. Right-click and choose **Set Recommended Values** for these hosts.

Figure 89 Setting Recommended Values for the Hosts

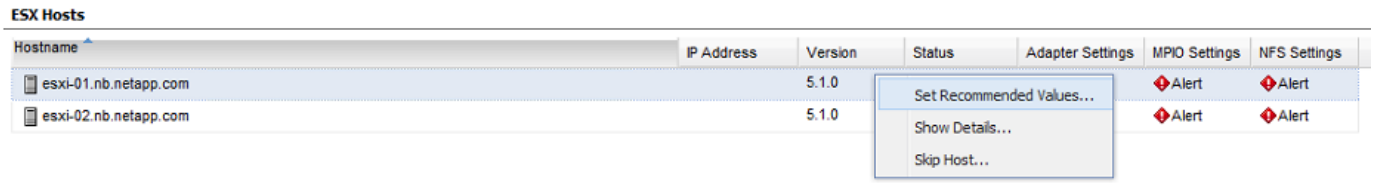

3. Check the settings to apply to selected vSphere hosts. Click **OK** to apply the settings.

   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).

Figure 90 Recommended Settings for NetApp Storage System

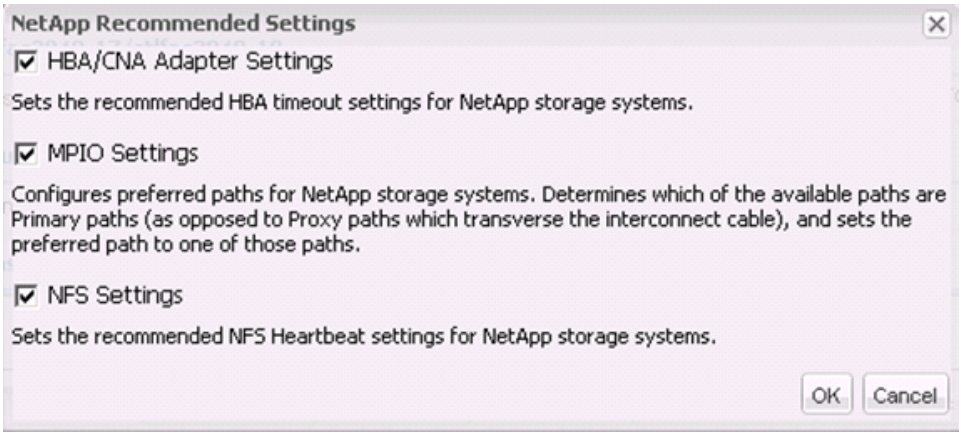

Note Depending on what changes have been made, the servers might require a restart for network-related parameter changes to take effect. If no reboot is required, the Status value is set to Normal. If a reboot is required, the Status value is set to Pending Reboot. If a reboot is required, the ESX or ESXi servers should be placed into Maintenance Mode, evacuate (if necessary), and be restarted before proceeding.

Figure 91 ESXi Servers with Changed Network Related Parameter

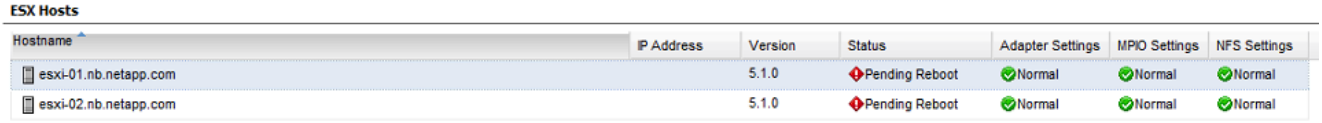

4. After the recommended values have been set and the ESXi servers are rebooted, the status of the ESXi servers must show normal.

Figure 92 ESXi Servers with Changed Network Related Parameter - After Reboot

VSC 4.2.1 Provisioning and Cloning Setup

Provisioning and cloning in VSC 4.2.1 helps administrators to provision both VMFS and NFS datastores at the data center, datastore cluster, or host level in VMware environments.

1. In a vSphere Client connected to vCenter, choose **Home** > **Solutions and Applications** > **NetApp** and click the **Provisioning and Cloning** tab on the left. Choose Storage controllers.
2. In the main part of the window, right-click <<var_controller1>> and choose Resources.
3. In the <<var_controller1>> resources window, use the arrows to move volumes ifgrp0-<<var_nfs_vlan_id>>, esxi_boot and aggr1 to the right. Also choose the Prevent further changes check box as shown in Figure 93.

Figure 93 List of Parameters of Storage Controller 1 Ready for Provisioning and Cloning

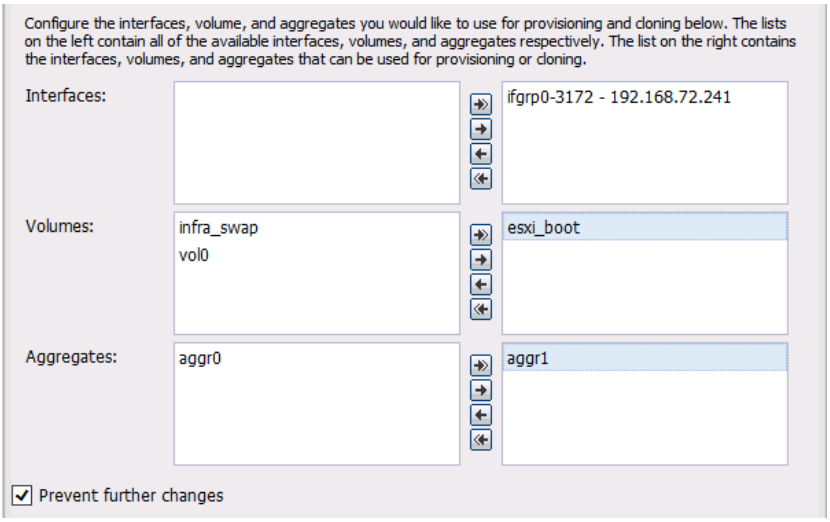

4. Click **Save**.
5. In the main part of the window, right-click <<var_controller2>> and choose Resources.
6. In the <<var_controller2>> resources window, use the arrows to move volumes ifgrp0-<<var_nfs_vlan_id>>, infra_datastore_1 and aggr1 to the right. choose the Prevent Further changes check box as shown in Figure 94.

Figure 94 List of Parameters of Storage Controller 2 Ready for Provisioning and Cloning

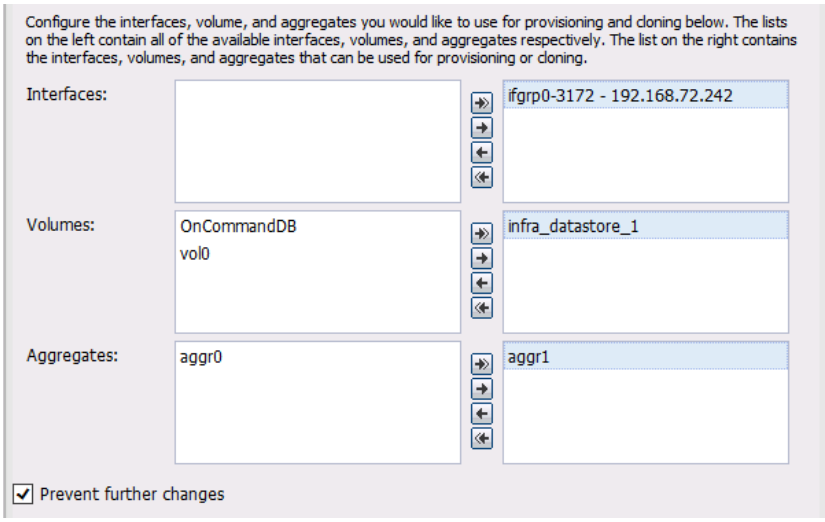

7. Click **Save**.

VSC 4.2.1 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, virtual machines, or virtual disk files, you must make sure that the storage systems that contain the datastores and virtual machines for which you are creating backups have valid storage credentials in the Monitoring and Host Configuration section.

If you are planning to leverage the SnapMirror update option, add all the destination storage systems with valid storage credentials to the Monitoring and Host Configuration section.

Backup and Recovery Configuration

To configure a backup job for a datastore, follow these steps:

1. Click **Backup and Recovery**, then choose Backup.
2. Click **Add**. The Backup wizard appears.

Figure 95 Configuring Backup

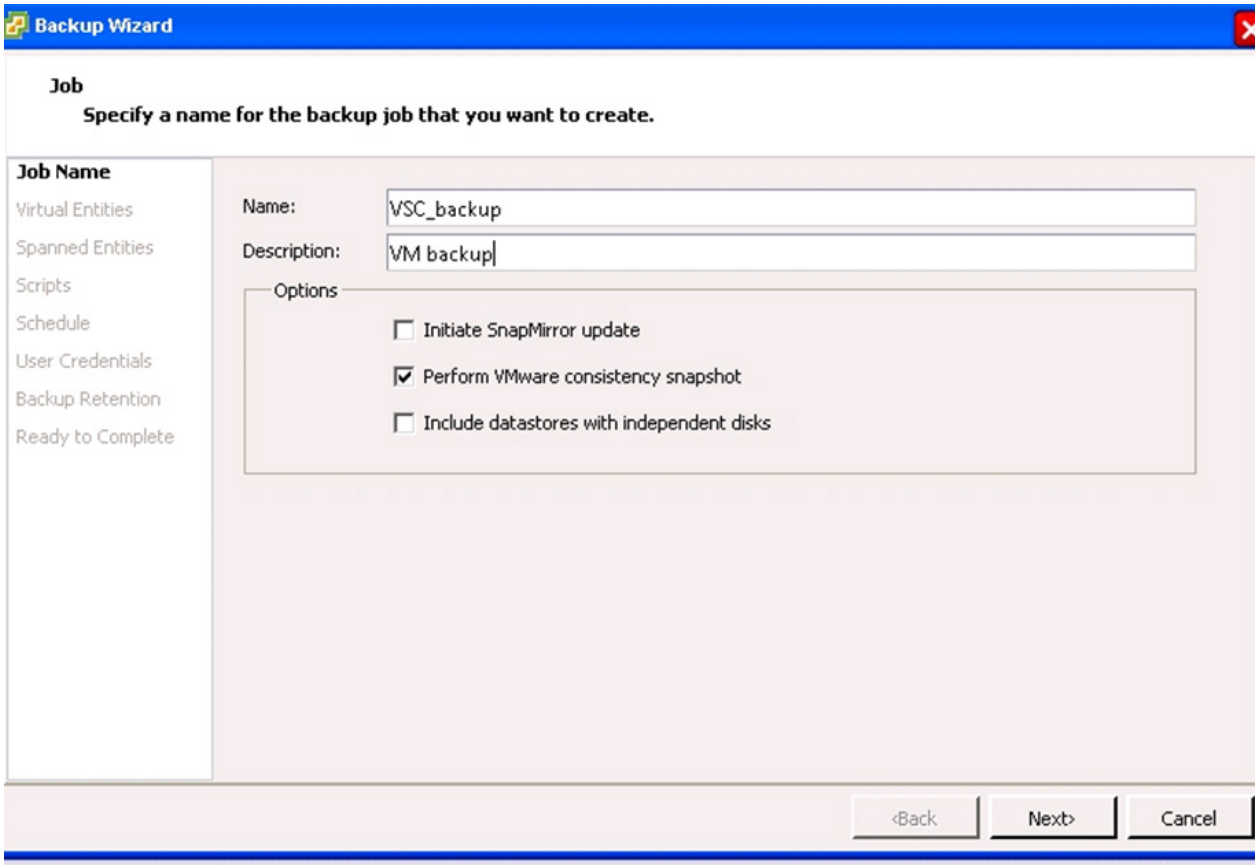

3. Type a backup job name and description.
4. If you want to create a VMware snapshot for each backup, choose Perform VMware consistency snapshot in the options pane.
5. Click **Next**.

6. choose infra_datastore_1 and then click to move it to the selected entities. Click **Next**.

Figure 96 Selecting Entities to Backup

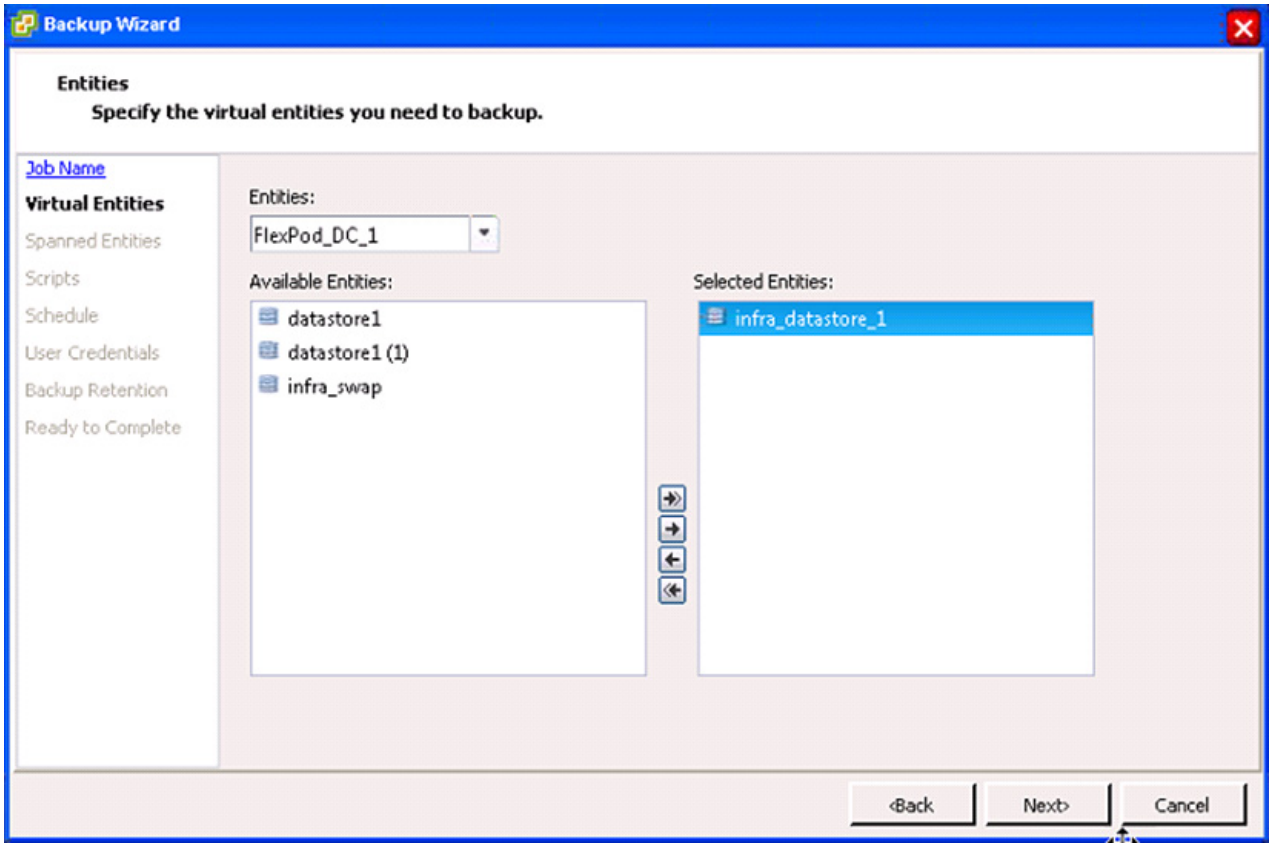

7. choose one or more backup scripts if available and click **Next**.

Figure 97 Select Scripts to Run Along the Backup

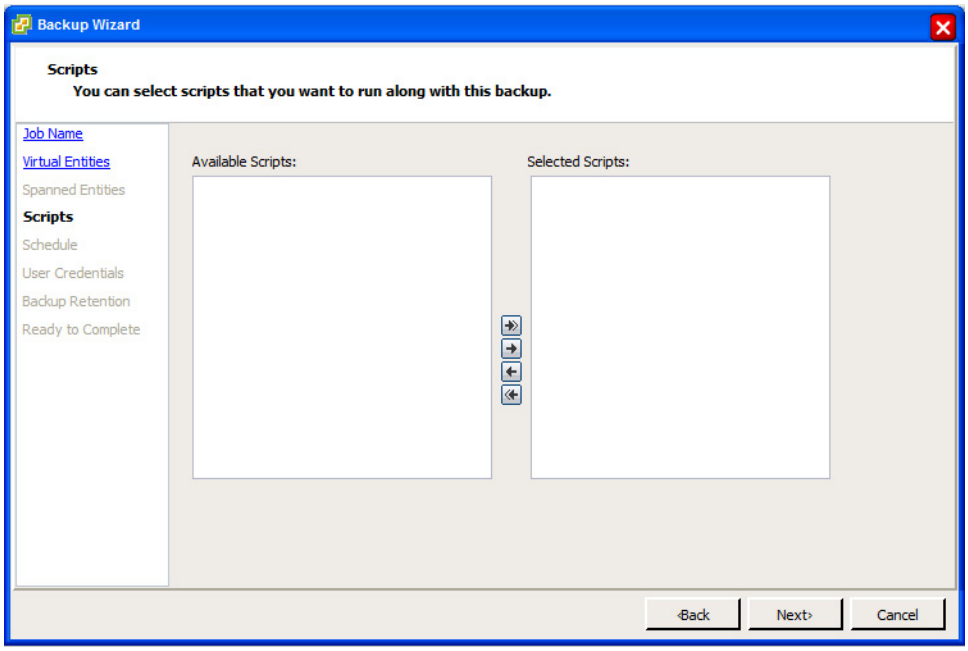

8. choose the hourly, daily, weekly, or monthly schedule that you want for this backup job and click **Next**.

Figure 98 Setting Schedule for Backup

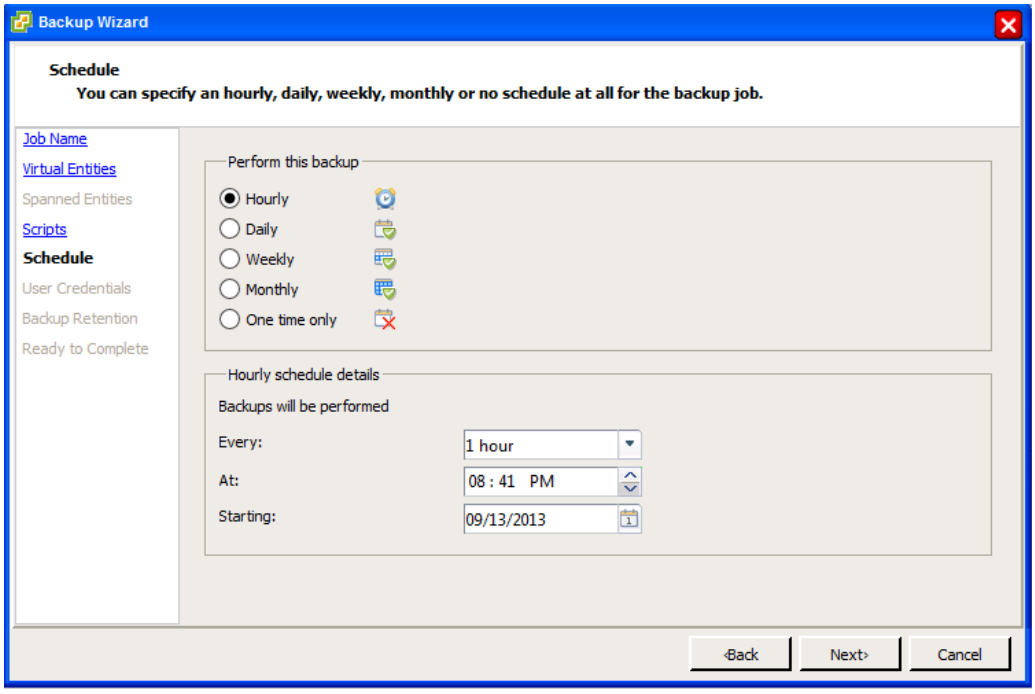

9. Use the default vCenter credentials or type the user name and password for the vCenter Server and click **Next**.

10. 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 e-mail addresses. Click **Next**.

Figure 99 Specifying Backup Retention Period

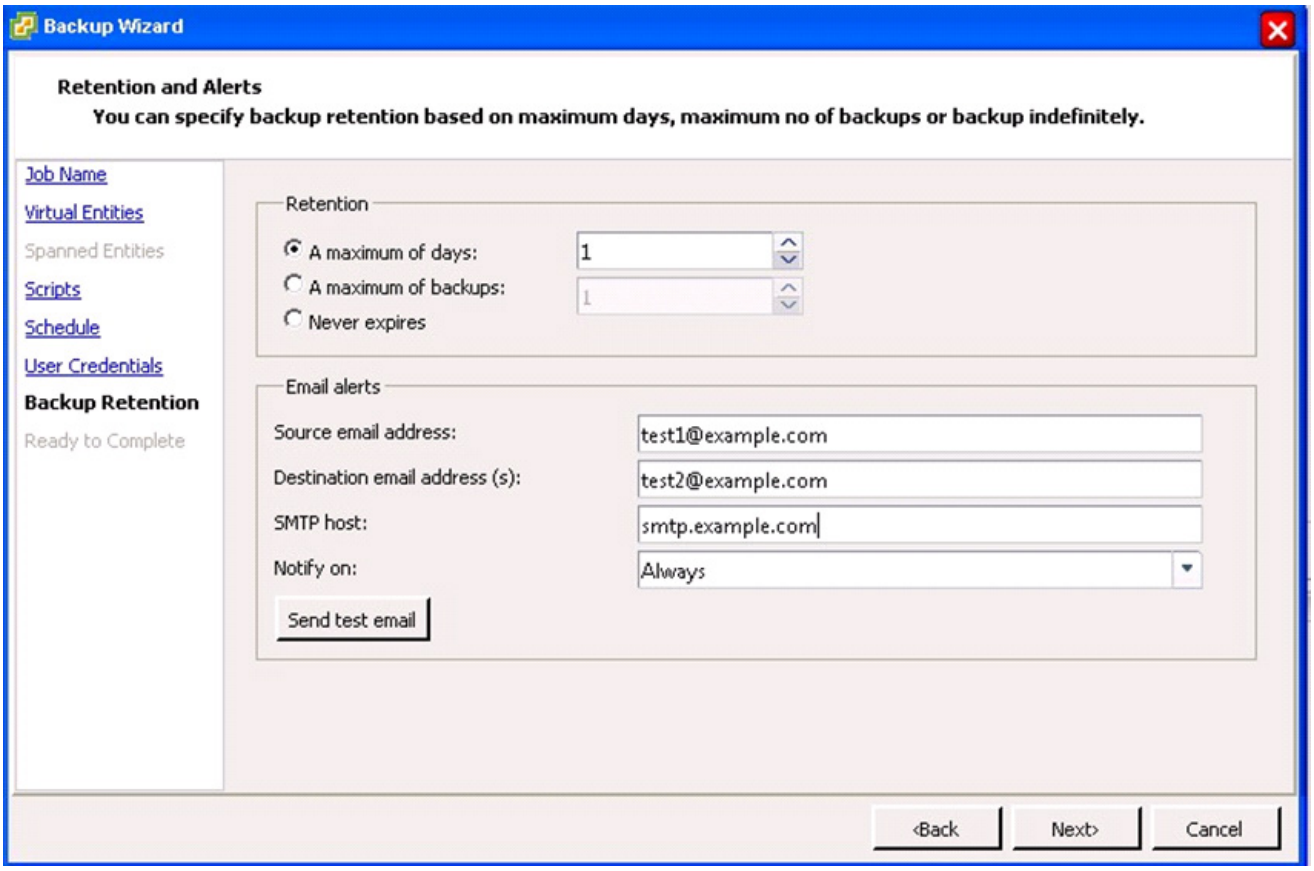

11. Review the summary page and click **Finish**. If you want to run the job immediately, choose the Run Job Now option and then click **Finish**.

Figure 100 Summary of Backup Settings

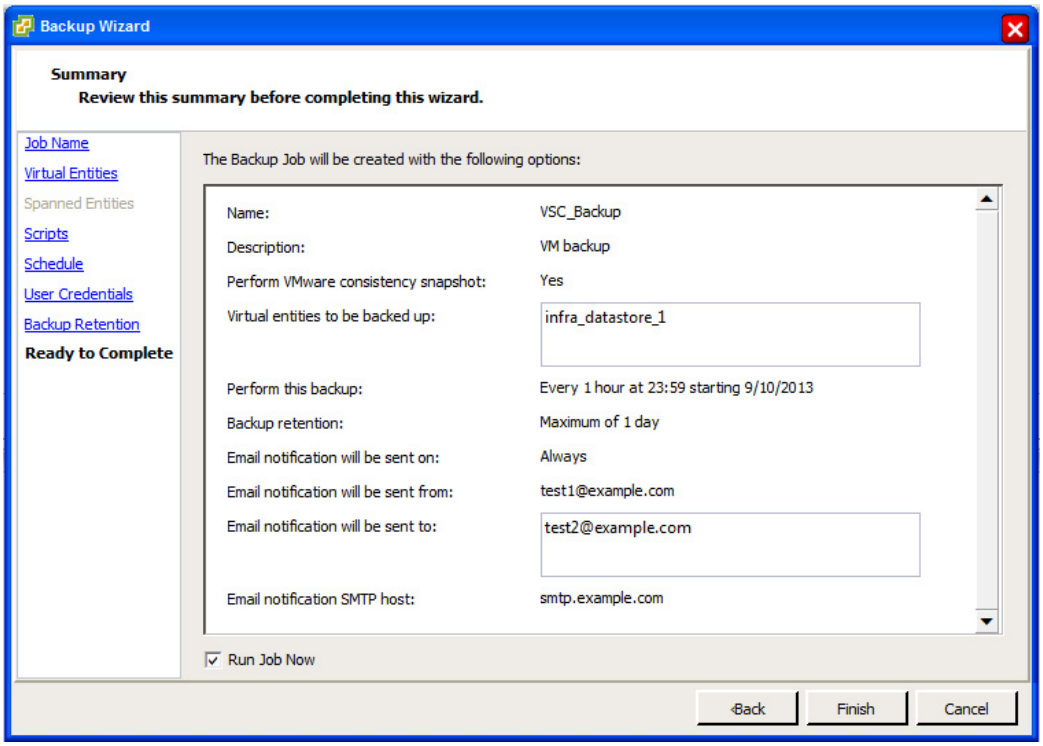

12. On the management interface of storage controller 2, automatic Snapshot copies of the infrastructure datastore volume can be disabled by typing the command:

```
snap sched infra_datastore_1 0 0 0
```

13. Also, to delete any existing automatic Snapshot copies that have been created on the volume type the following command:

```
snap list infra_datastore_1
snap delete infra_datastore_1 <snapshot name>
```

OnCommand Unified Manager 5.2

Create Raw Device Mapping (RDM) Datastore

From the VMware vCenter Client, do as follows:

1. In the VMware vCenter Client, from **Home** > **Inventory** > **Hosts and Clusters**, right-click the FlexPod_Management cluster.
2. choose **NetApp** > **Provisioning and Cloning** > **Provision Datastore**.
3. Select storage controller 2 as the Target Storage Controller and click **Next**.
4. choose VMFS as the Datastore type and click **Next**.
5. Select FCP as the Protocol type, set the Size to 100, enter the datastore name as RDM_Map, select the checkbox to create new volume container, select aggr1 as the Aggregate, check the Thin Provision checkbox, and click **Next**.
6. Verify settings and click **Apply**.

Install .NET Framework 4.0 Feature

From the Virtual Storage Console (VSC) and OnCommand VM, follow these steps:

1. Download the .NET Framework 4.0 executable file from the Microsoft Download Center.
2. Launch the Installer.
3. Read and accept the license terms by selecting the checkbox.
4. Click **Install**.
5. Click **Finish**.

Figure 101 .Net Framework 4 - Installation Complete

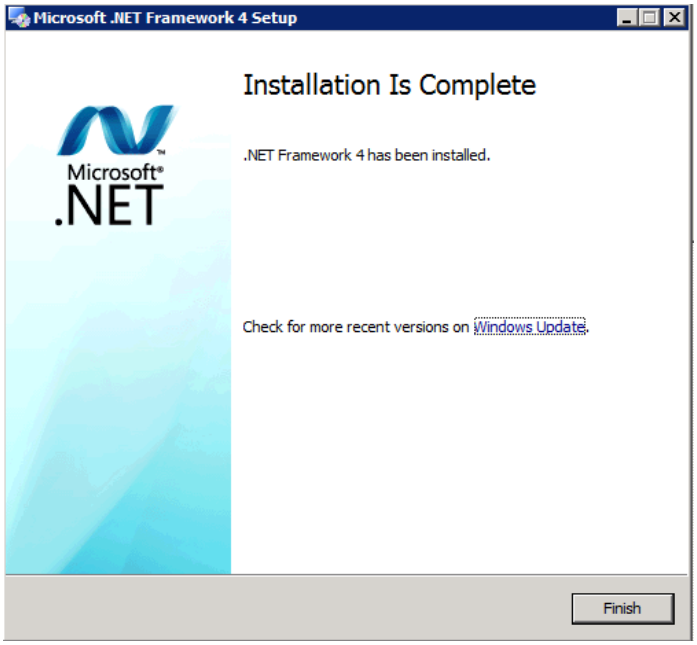

6. Click **Restart Now**, to restart the VSC and OnCommand VM.

Install SnapDrive 7.0

Follow these steps to install SnapDrive® 7.0.

1. Log in to the VSC and OnCommand VM using the FlexPod admin credentials.
2. Install all Windows updates on the virtual machine
3. Download SnapDrive 7.0 from the NetApp Support Site.
4. Browse to the location of the SnapDrive installation package and double-click the executable file. This launches the SnapDrive installation wizard and opens the Welcome page.
5. Click **Next** in the Welcome page of the SnapDrive installation wizard.
6. If this is a new SnapDrive installation, read and accept the license agreement. Click **Next**.
7. Select the appropriate license type, provide a license if necessary. Click **Next**.
8. Enter the user name and organization. Click **Next**.
9. To accept the default installation folder, click **Next**.

10. Check the checkbox to allow SnapDrive to communicate via the firewall. Click **Next**.
11. Check the checkbox to enable VMware vCenter integration.

Figure 102 Specifying Account Information for Installing SnapDrive 7.0

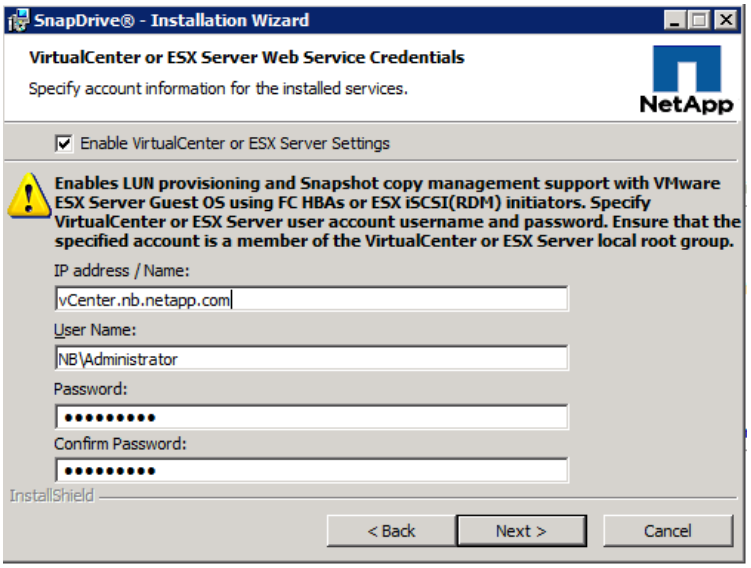

12. Provide the IP address/host name, user name, and password for the vCenter Server. Click Next.
13. If this is a SnapDrive upgrade, select Modify/Upgrade in the Program Maintenance page. Click Next.
14. Select the checkbox to enable VSC integration.

Figure 103 Specifying Virtual Infrastructure Details

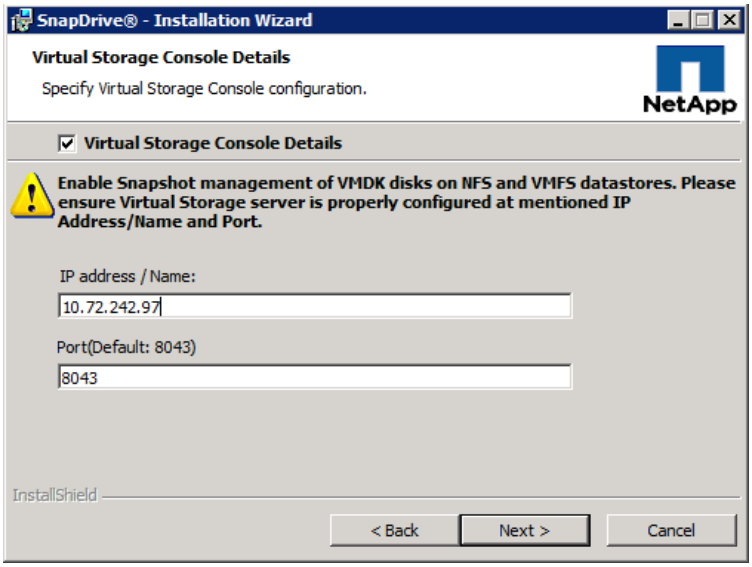

15. Provide the IP address/host name of the VSC virtual machine, accept the default port.
16. Click Next.

17. Provide the SnapDrive Service credentials. Click Next.
18. Accept the default ports for Web Service Configuration. Click Next.
19. (Optional) Provide a preferred IP address for management traffic. Click Next.
20. Select the checkbox to Enable Transport Protocol Settings:
    a. Select HTTPS as the protocol.
    b. Provide the user name as root.
    c. Provide the root password for the storage controllers.
    d. Verify that the port ID is 443.

Figure 104 Transport Protocol Settings

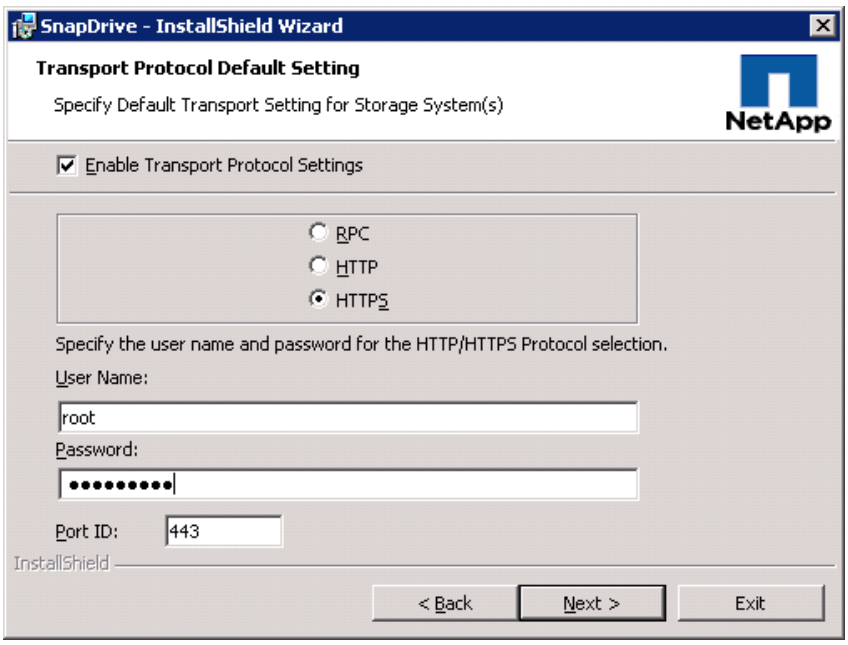

21. Click **Next** > **Next** > **Install** > **Finish**.
22. From the Start menu, open SnapDrive.
23. In the left pane, expand the local machine and choose Disks.
24. In the right pane, choose Create Disk.
25. In the create disk Wizard Window, click **Next**.
26. In the storage system name field, enter the management IP address of storage controller 2, and click **Add**.
27. In the list that appears, choose OnCommandDB.
28. Enter OnCommandDB for the LUN Name and click **Next**.

Figure 105 Adding LUN Information

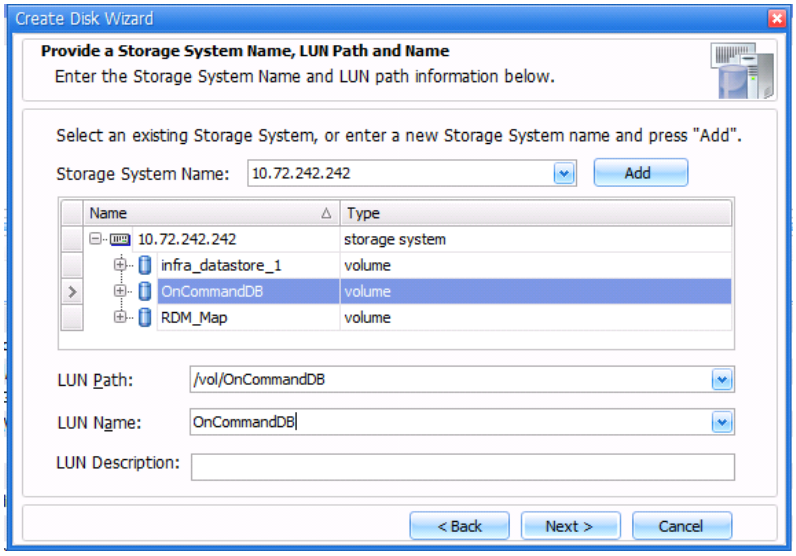

29. Make sure the LUN type is set to Dedicated and click **Next**.
30. Assign drive letter O and set LUN size to 100GB. Click **Next**.
31. Click **Next** to accept the volume properties changes.

Figure 106 Entering LUN Properties

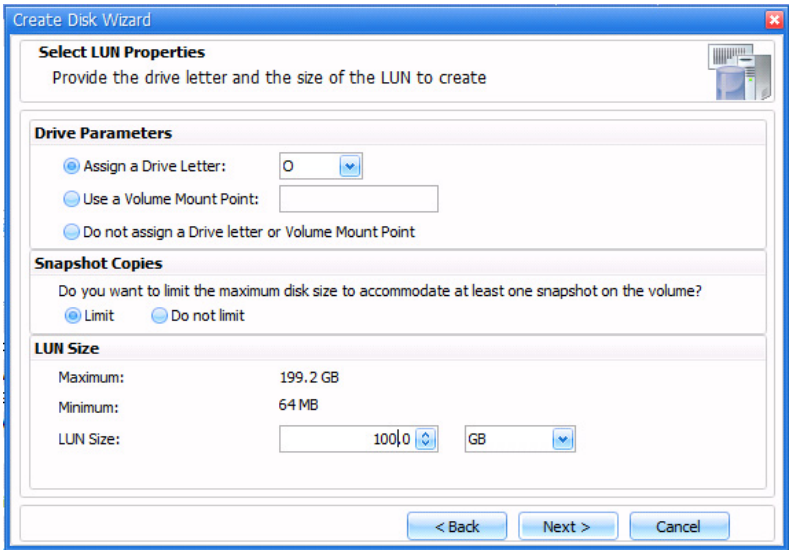

32. Choose all initiators on the Initiator List, and click **Next**.
33. Choose manual as the Initiator group management, and click **Next**.
34. Choose the MGMT-Hosts igroup, and click **Next**.

Figure 107 Specifying igroups for the LUN

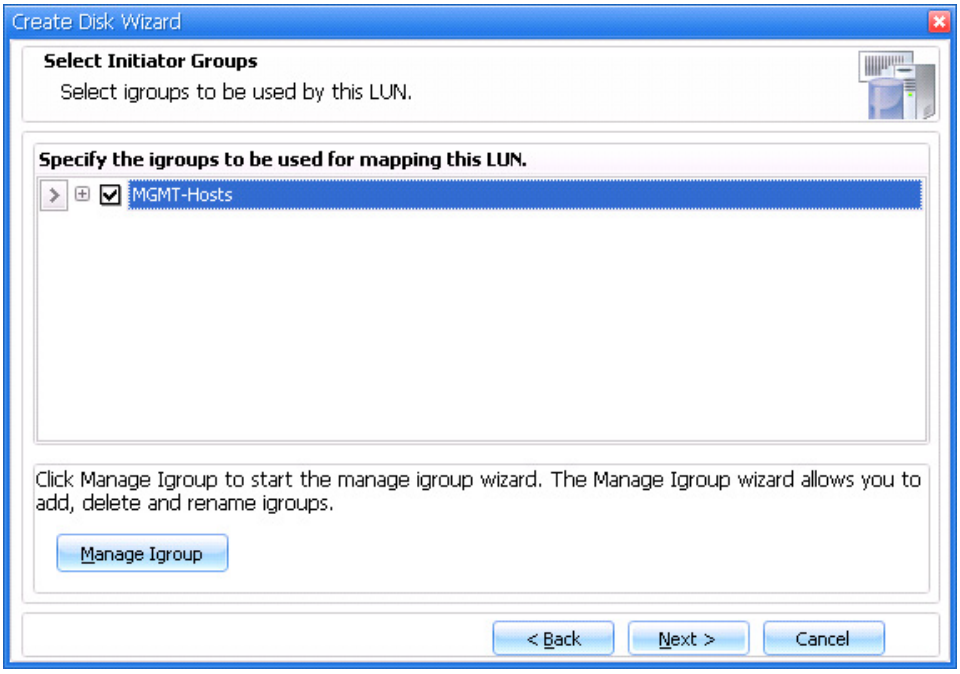

35. Choose the RDM_Map Datastore in the Select Datastore section. Click **Next**.

Figure 108 Selecting a Datastore to Store LUN Mapping

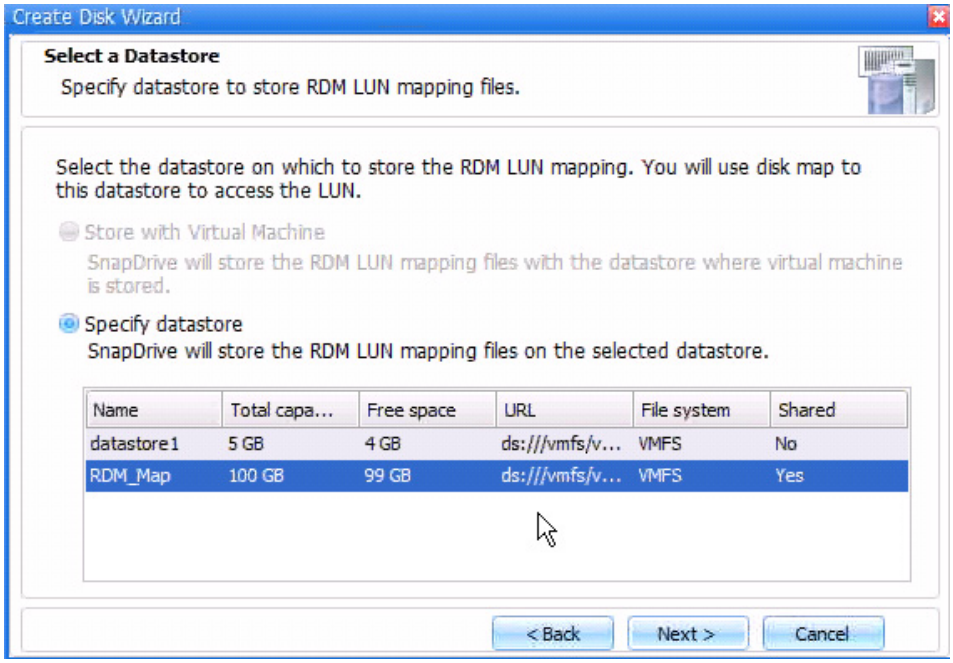

36. Click **Finish** to create the disk.
37. Close **SnapDrive**.

Install NetApp OnCommand Core Package

To install the OnCommand Unified Manager Core Package, follow these steps:

1. To download the OnCommand Unified Manager Core Package for Windows, go to http://support.netapp.com/NOW/download/software/occore_win/5.2/.
2. Using the FlexPod admin credentials, log in to the VSC and OnCommand VM.
3. Identify the DataFabric® Manager Server license key before starting the installation.

   The DataFabric Manager Server license key
4. Navigate to the path or directory containing the downloaded file and launch the file.
5. In the Welcome screen, click **Next**.
6. Accept the AutoSupport notice and click **Next**.
7. Choose 7-Mode as the Mode of Operation. Click **Next**.

Figure 109 Selecting the Operation Mode

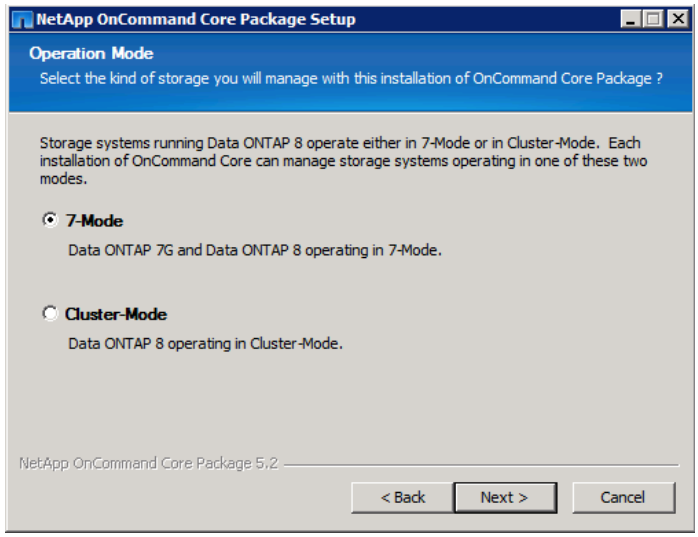

8. Click the 7-Mode radio button for the Operation Mode. Click **Next**.
9. Enter the 14-character license key when prompted and click **Next**.
10. choose the installation location, if different from the default.

Note Do not change the default location of the local Temp Folder directory, or the installation will fail. The installer automatically extracts the installation files to the %TEMP% location.

11. Click **Install**.
12. Click **Next**.
13. Click **Finish**.
14. Follow the remaining setup prompts to complete the installation.

    From an MS-DOS command prompt, follow these steps as an administrator:

15. In preparation for the database movement to the previously created LUN from local storage, stop all OnCommand Unified Manager services and verify that the services have stopped.

```
dfm service stop
dfm service list
```

16. Move the data to the previously created LUN.

Note The **dfm datastore setup help** command provides switch options available with the command.

```
dfm datastore setup O:\
```

17. Start OnCommand Unified Manager and then verify that all services have started.

```
dfm service start
dfm service list
```

18. Generate an SSL key.

```
dfm ssl server setup
Key Size (minimum = 512..1024..2048..) [default=512]: 1024
Certificate Duration (days) [default=365]: Enter
Country Name (e.g., 2 letter code): <<var_country_code>>
State or Province Name (full name): <<var_state>>
Locality Name (city): <<var_city>>
Organization Name (e.g., company): <<var_org>>
Organizational Unit Name (e.g., section): <<var_ unit>>
Common Name (fully-qualified hostname): <<var_oncommand_server_fqdn>>
Email Address:  <<var_admin_email>>
```

Note The SSL key command fails if certain command line option inputs do not follow specified character lengths (for example, a two-letter country code), and any multiword entries must be encased in double quotation marks, for example, “North Carolina.”

19. Turn off automatic discovery.

```
dfm option set discoverEnabled=no
```

20. Set the protocol security options for communication with various devices.

```
dfm service stop http
dfm option set httpsEnabled=yes
dfm option set httpEnabled=no
dfm option set httpsPort=8443
dfm option set hostLoginProtocol=ssh
dfm option set hostAdminTransport=https
```

Note The HTTPS and SSH protocols must be enabled on the storage controllers that are monitored by OnCommand Unified Manager.

21. Restart the DataFabric Manager HTTP services to make sure that the security options take effect.

```
dfm service start http
```

22. Configure OnCommand Unified Manager to use SNMPv3 to poll configuration information from the storage devices. Use the user name and password generated for SNMPv3.

```
dfm snmp modify -v 3 -c <<var_snmp_community>> -U snmpv3user -P <<var_password>>
-A MD5 default
```

23. Set up OnCommand Unified Manager to send AutoSupport through HTTPS to NetApp.

```
dfm option set SMTPServerName=<<var_mailhost>>
dfm option set autosupportAdminContact=<<var_storage_admin_email>>
dfm option set autosupportContent=complete
dfm option set autosupportProtocol=https
```

24. Manually add the storage cluster to the OnCommand server.

```
dfm host add <<var_controller1>>
dfm host add <<var_controller2>>
```

25. Set the array login and password credentials in OnCommand Unified Manager. This is the root or administrator account.

```
dfm host set <<var_controller1>> hostlogin=root
dfm host set <<var_controller1>> hostPassword=<<var_password>>
dfm host set <<var_controller2>> hostlogin=root
dfm host set <<var_controller2>> hostPassword=<<var_password>>
```

26. List the storage systems discovered by OnCommand Unified Manager and their properties.

```
dfm host list
dfm host get <<var_controller1>>
dfm host get <<var_controller2>>
```

27. Test the network configuration and connectivity between the OnCommand server and the named host. This test helps identify misconfigurations that prevent the OnCommand server from monitoring or managing a particular appliance. The test should be the first command used if a problem using the OnCommand server occurs with only some of the appliances.

```
dfm host diag <<var_controller1>>
dfm host diag <<var_controller2>>
```

28. (optional) Configure an SNMP trap host.

```
dfm alarm create -T <<var_oncommand_server_fqdn>>
```

29. Configure OnCommand Unified Manager to generate and send e-mails for every event whose importance ranks as critical or higher.

```
dfm alarm create -E <<var_admin_email>> -v Critical
```

30. Create a manual backup.

```
dfm backup create -t snapshot
```

31. Schedule backups to a virtual backup directory on the 100GB FC LUN.

```
dfm option set backupRetentionCount=20
dfm backup schedule set -t snapshot -D 21:00
```

32. To open Windows Firewall with Advanced Security, click **Start > Administrative Tools > Windows Firewall with Advanced Security**.
33. Choose Inbound Rules.
34. Click **New Rule**.
35. Choose Port and click Next.
36. Leave TCP selected and enter 8443 in the Specific local ports text box. Click **Next**.
37. Click **Next**.
38. Click **Next**.

39. Name the rule OnCommand Console External Access and click **Finish**.
40. Click **New Rule**.
41. Choose Port and click **Next**.
42. Choose UDP and enter 162 in the Specific local ports text box. Click **Next**.
43. Click **Next**.
44. Click **Next**.
45. Name the rule OnCommand SNMP Trap and click **Finish**.
46. Close Windows Firewall with Advanced Security.

NetApp NFS Plug-In 1.0.20 for VMware VAAI

Enable VMware vStorage for NFS in Data ONTAP 7-Mode

To enable VMware vStorage for NFS when Data ONTAP is operating in 7-mode, follow these steps:

1. From a Secure Shell (SSH) session on each storage controller, log in with the root user name and password.
2. Enable vStorage on the storage system.

```
options nfs.vstorage.enable on
```

Install NetApp NFS Plug-In for VMware VAAI

To install the NetApp NFS plug-in for VMware vStorage APIs for Array Integration (VAAI), follow these steps:

1. From the vSphere console of the VSC and OnCommand virtual machine (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, choose the ESXi platform, and click **Go**.
3. Download the .vib file of the most recent plug-in version.
4. Verify that the file name of the .vib file matches the predefined name that VSC 4.2.1 for VMware vSphere uses: NetAppNasPlugin.vib.

Note If the .vib file name does not match the predefined name, rename the .vib file. Neither the VSC client nor the NetApp vSphere Plug-in Framework (NVPF) service needs to be restarted after the .vib file is renamed.

5. Copy the plug-in .vib file (NetAppNasPlugin.vib) to C:\Program Files\Virtual Storage Console\etc\vsc\web.

Note The default directory path is C:\Program Files\NetApp\Virtual Storage Console\.

6. In the VMware vSphere Client connected to the vCenter Server, choose **Home > Solutions and Applications > NetApp**.
7. In the Monitoring and Host Configuration capability navigation pane, choose Tools.

8. Under NFS Plug-in for VMware VAAI, click **Install on Host**.

Figure 110 Installing NFS Plug-in for VMware VAAI on Host

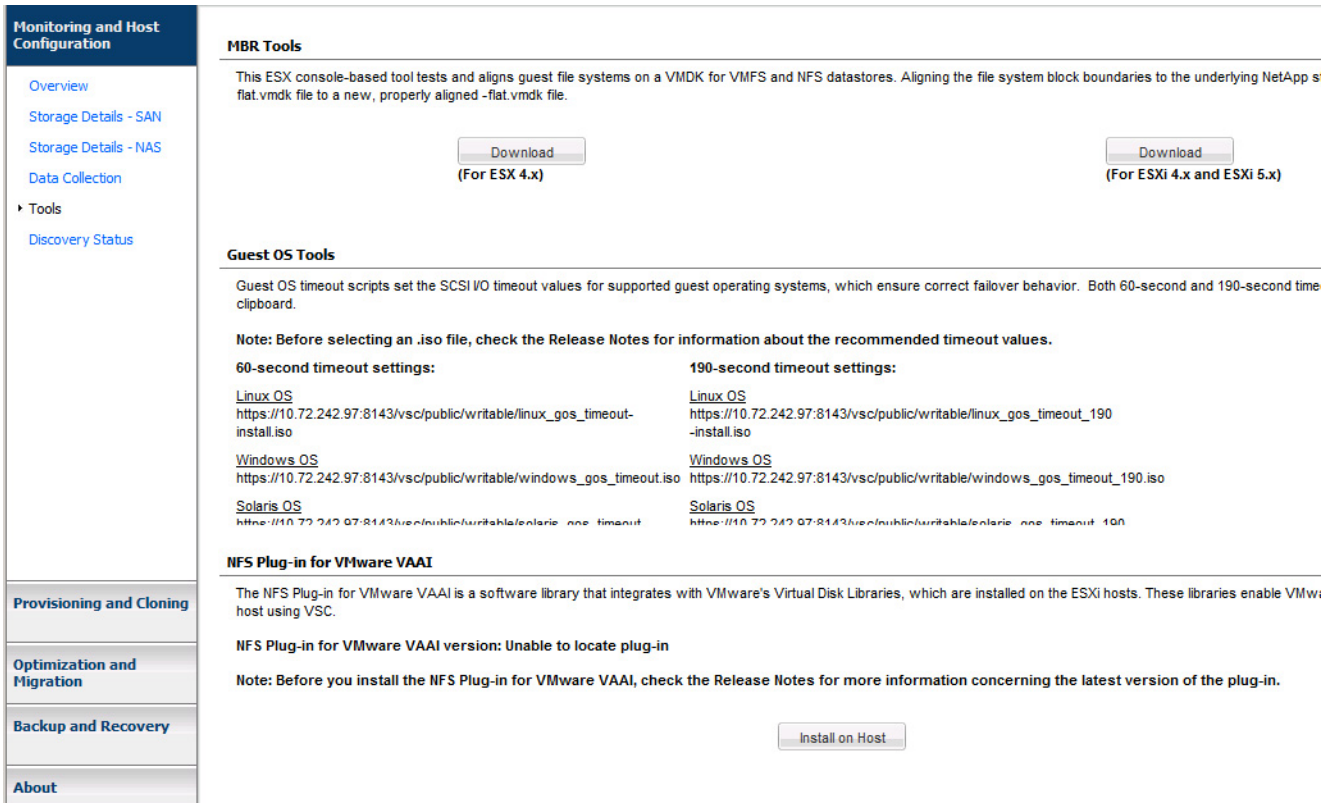

9. Choose all ESXi hosts and click **Install**, and then click **Yes**.

Note The Monitoring and Host Configuration capability automatically installs the plug-in on the hosts selected.

10. Choose **Home** > **Inventory** > **Host and Clusters**.
11. For each host (one at a time), right-click the host and choose Enter Maintenance Mode.
12. Click **Yes**, click **Yes** again, and then click **OK**.

Note It might be necessary to migrate all VMs away from the host.

13. After the host is in maintenance mode, right-click the host and choose Reboot.
14. Enter a reason for the reboot and click **OK**.
15. After the host reconnects to the vCenter Server, right-click the host and choose Exit Maintenance Mode.
16. Make sure that all ESXi hosts get rebooted.
17. Click **OK** to close the VASA Configuration.

Appendix

Build Windows Active Directory Server VM(s)

ESXi Host VM-Host-Infra-01

To build an Active Directory Server virtual machine (VM) for VM-Host-Infra-01 ESXi host, follow these steps:

1. Log in to the host using VMware vSphere Client.
2. In the vSphere Client, choose the host in the inventory pane.
3. Right-click the host and choose New Virtual Machine.
4. Choose Custom and click **Next**.
5. Enter a name for the VM. Click **Next**.
6. Choose infra_datastore_1. Click **Next**.
7. Choose Virtual Machine Version: 8. Click **Next**.
8. Verify that the Windows option and the Microsoft Windows Server 2008 R2 (64-bit) version are selected. Click **Next**.
9. Choose two virtual sockets and one core per virtual socket. Click **Next**.
10. Choose 4GB of memory. Click **Next**.
11. Choose one network interface card (NIC).
12. For NIC 1, choose the IB-MGMT Network option and the VMXNET 3 adapter. Click **Next**.
13. Keep the LSI Logic SAS option for the SCSI controller selected. Click **Next**.
14. Keep the Create a New Virtual Disk option selected. Click **Next**.
15. Make the disk size at least 60GB. Click **Next**.
16. Click **Next**.
17. Check the check box for Edit the Virtual Machine Settings Before Completion. Click **Continue**.
18. Click the **Options** tab.
19. choose Boot Options.
20. Check the Force BIOS Setup check box.
21. Click **Finish**.
22. From the left pane, expand the host field by clicking the plus sign (+).
23. Right-click the newly created AD Server VM and click **Open** Console.
24. Click the third button (green right arrow) to power on the VM.
25. Click the ninth button (CD with a wrench) to map the Windows Server 2008 R2 SP1 ISO, and then choose Connect to ISO Image on Local Disk.
26. Navigate to the Windows Server 2008 R2 SP1 ISO, select it, and click **Open**.
27. In the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to choose CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.

28. The Windows Installer boots. Choose the appropriate language, time and currency format, and keyboard. Click **Next**.
29. Click Install now.
30. Make sure that the Windows Server 2008 R2 Standard (Full Installation) option is selected. Click **Next**.
31. Read and accept the license terms and click **Next**.
32. Choose Custom (Advanced). Make sure that Disk 0 Unallocated Space is selected. Click **Next** to allow the Windows installation to complete.
33. After the Windows installation is complete and the VM has rebooted, click **OK** to set the Administrator password.
34. Enter and confirm the Administrator password and click the blue arrow to log in. Click **OK** to confirm the password change.
35. After logging in to the VM desktop, from the VM console window, choose the VM menu. Under Guest, choose Install/Upgrade VMware Tools. Click **OK**.
36. If prompted to eject the Windows installation media before running the setup for the VMware tools, click **OK**, then click **OK**.
37. In the dialog box, choose Run setup64.exe.
38. In the VMware Tools installer window, click **Next**.
39. Make sure that Typical is selected and click **Next**.
40. Click **Install**.
41. Click **Finish**.
42. Click **Yes** to restart the VM.
43. After the reboot is complete, choose the VM menu. Under Guest, choose Send Ctrl+Alt+Del. Then enter the password to log in to the VM.
44. Set the time zone for the VM, IP address, gateway, and host name.

Note A reboot is required.

45. If necessary, activate Windows.
46. Download and install all required Windows updates.

Note This process requires several reboots.

47. Open Server Manager.
48. On the left, click **Roles** > **Add Roles**.
49. Click **Next**.
50. In the list, check the check box next to Active Directory Domain Services.
51. In the popup, click **Add** Required Features to add .NET Framework 3.5.1.

Figure 111 Selecting Roles to Install on the Server

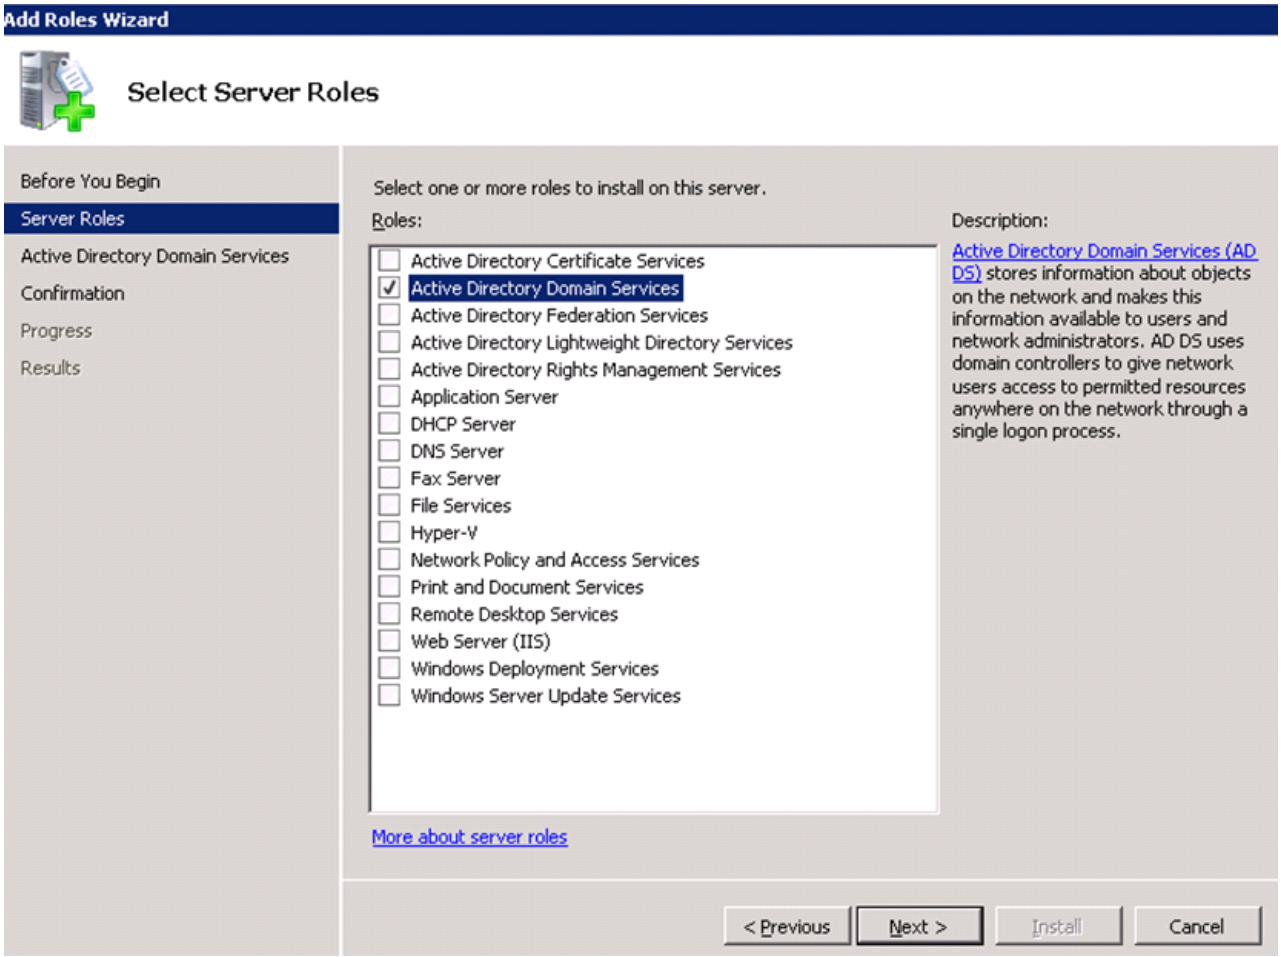

52. Click **Next**.
53. Click **Next**.

Figure 112 Confirming Installation Selections

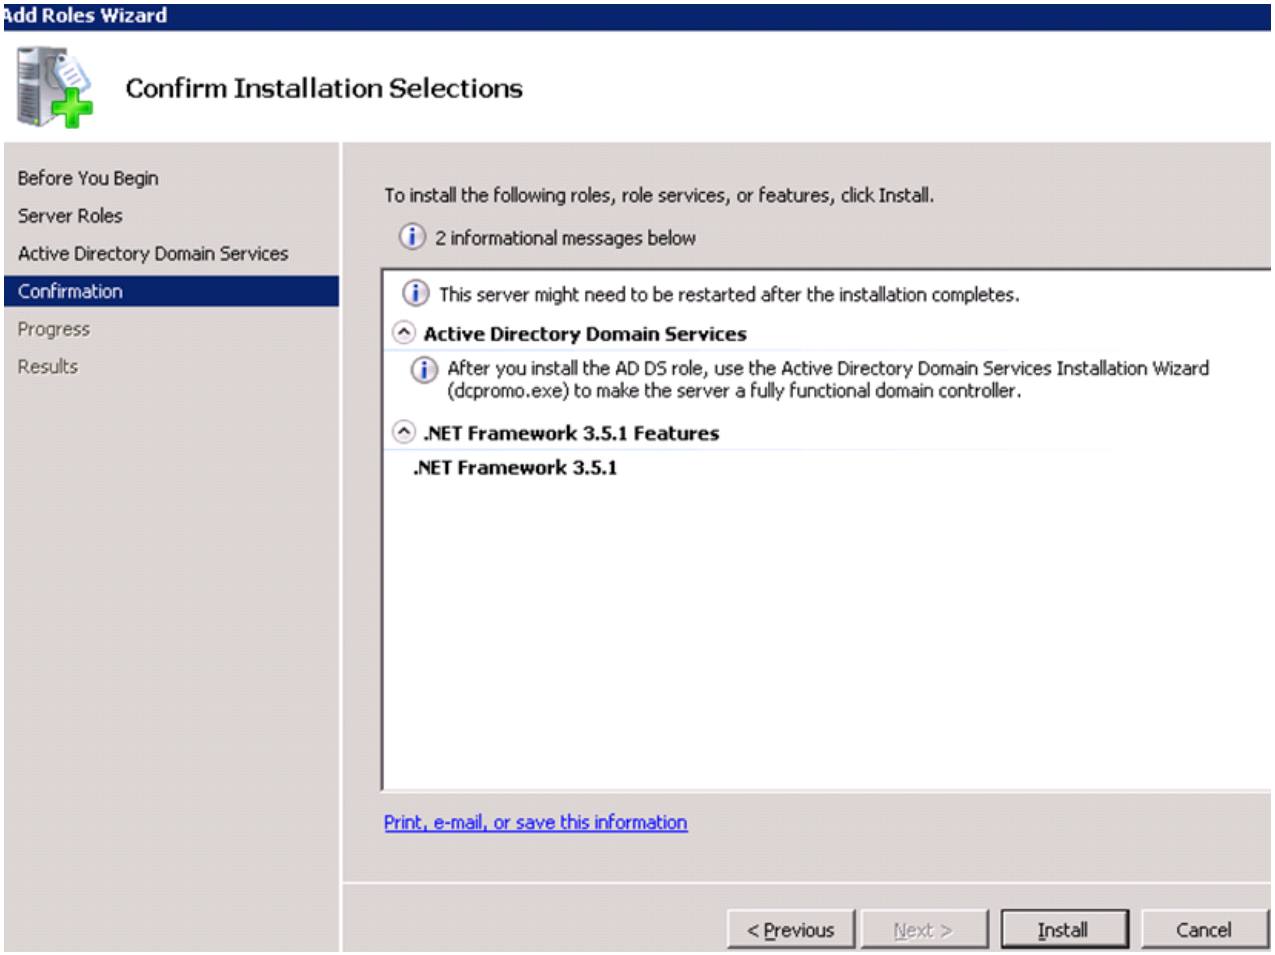

54. Click **Install**.

Figure 113 Installation Completion Summary

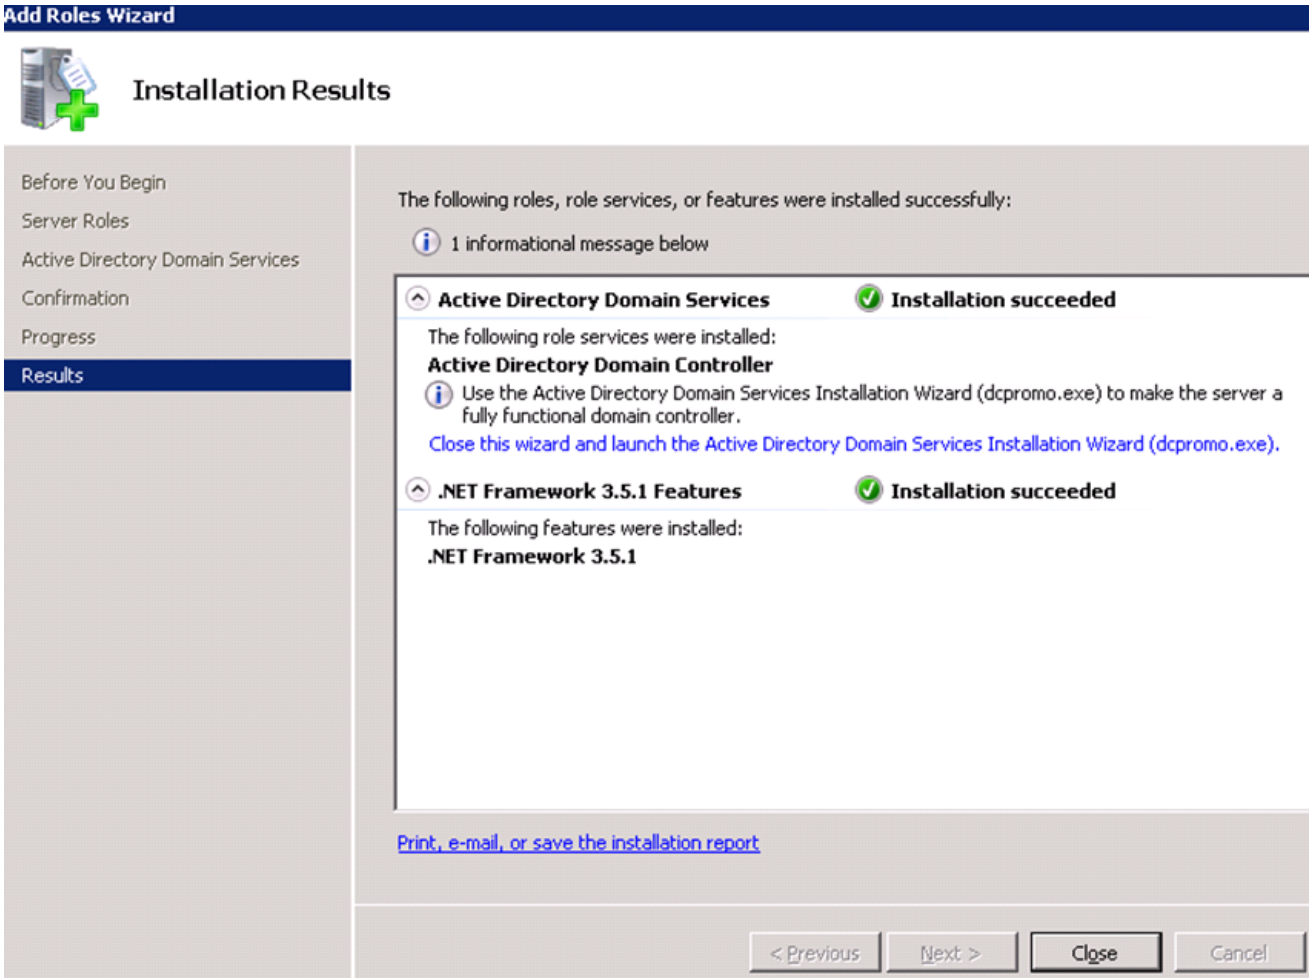

55. In the middle of the window, click **Close** this wizard and launch the Active Directory Domain Services Installation Wizard (dcpromo.exe).
56. In the Active Directory Domain Services Installation Wizard, click **Next**.
57. Click **Next**.
58. Choose Create a new domain in a new forest and click **Next**.

Figure 114 Creating a Domain Controller for a New Forest

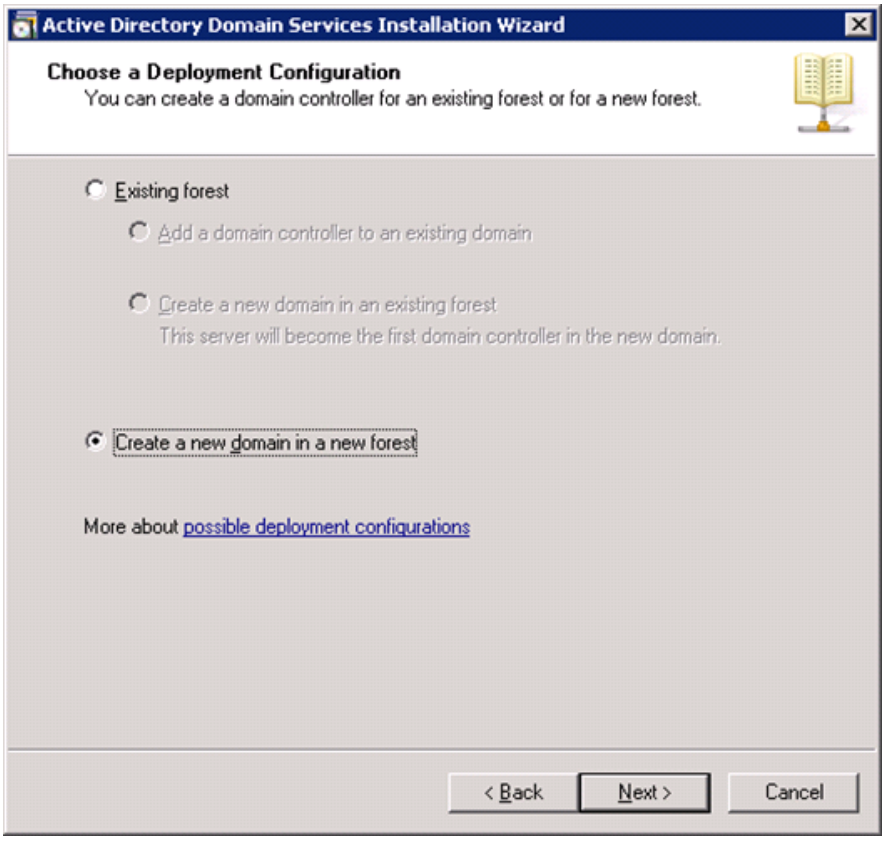

59. Type the FQDN of the Windows domain for this FlexPod and click **Next**.

Figure 115 Naming the Forest Root Domain

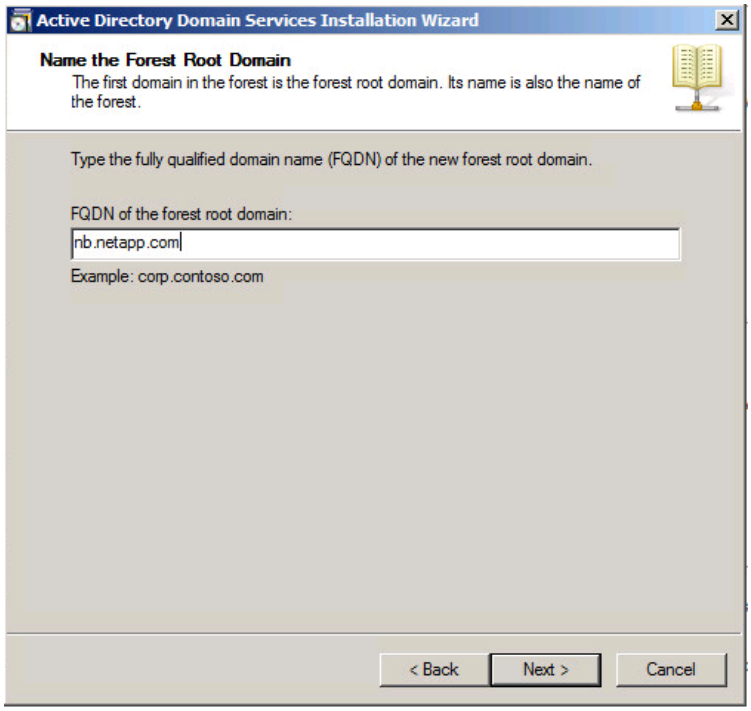

60. Choose the appropriate forest functional level and click **Next**.
61. Keep DNS server selected and click **Next**.

Figure 116 Selecting Additional Options for the Domain Controller

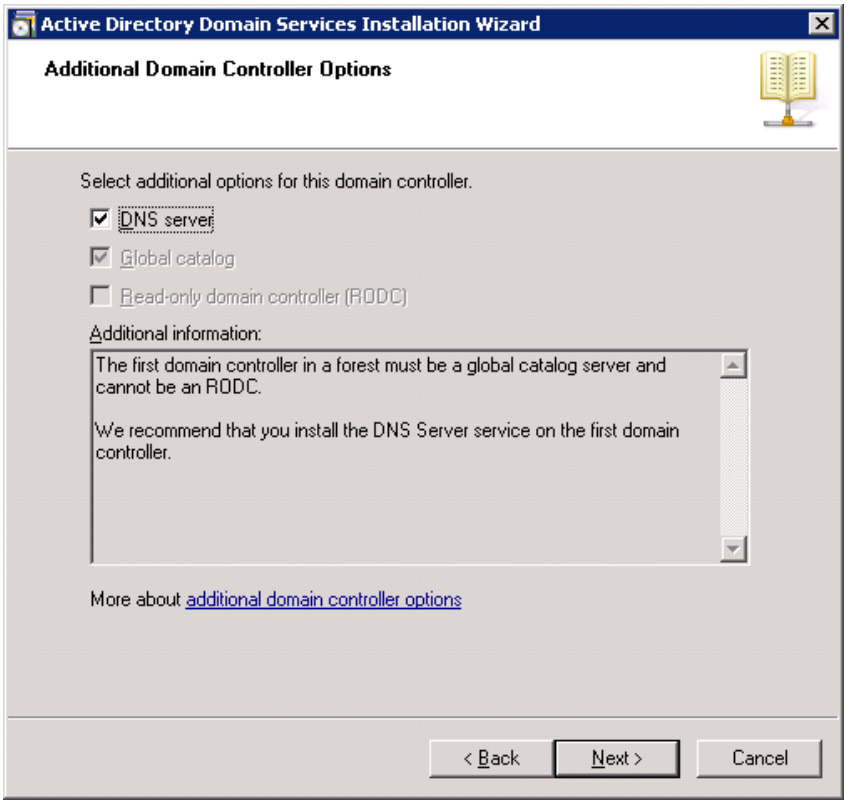

62. If one or more DNS servers exist that this domain can resolve from, Click **Yes** to create a DNS delegation. If this is AD server is being created on an isolated network, click **No**, to not create a DNS delegation. The remaining steps in this procedure assume a DNS delegation is not created. Click **Next**.
63. Click **Next** to accept the default locations for database and log files.
64. Enter and confirm <<var_password>> for the Directory Services Restore Mode Administrator Password. Click **Next**.
65. Review the Summary information and click **Next**. Active Directory Domain Services will install.
66. Click **Finish**.
67. Click **Restart Now** to restart the AD Server.
68. After the machine has rebooted, log in as the domain Administrator.
69. Open the DNS Manager by clicking **Start** > **Administrative Tools** > **DNS**.
70. (Optional) Add Reverse Lookup Zones for your IP address ranges.
71. Expand the Server and Forward Lookup Zones. Choose the zone for the domain. Right-click and choose New Host (A or AAAA). Populate the DNS Server with Host Records for all components in the FlexPod.
72. (Optional) Build a second AD server VM. Add this server to the newly created Windows Domain and activate Windows. Install Active Directory Domain Services on this machine. Launch dcpromo.exe at the end of this installation. Choose to add a domain controller to a domain in an

existing forest. Add this domain controller to the domain created earlier. Complete the installation of this second domain controller. After vCenter Server is installed, affinity rules can be created to keep the two AD servers running on different hosts.

Configuring Cisco VM-FEX with the UCS Manager

Background

FlexPod for VMware utilizes distributed virtual switching to manage the virtual access layer from a central point. While previous versions of FlexPod have only described the use of the Cisco Nexus 1000V, there exists an option to use the built-in virtual switching functionality delivered through hardware on the Cisco UCS known as VM-FEX. This has several advantages:

- There is no need for extra HW such as Cisco Nexus 1110-X.
- Cisco UCS provides a central configuration environment with which the administrator is already familiar.
- Compared to using the Cisco Nexus 1000v as virtual appliances within vCenter itself, this setup avoids an SPOF and common restart issues when running the distributed switches in an environment in which they are required for the network functionality of the ESX servers on which they are running. This is a common problem that needs to be addressed in the solution design.

In other words, it dramatically simplifies the hardware setup and operation by optimally utilizing the new hardware features.

Process Overview

This section provides a detailed overview of VM-FEX setup, configuration, and operation using Cisco UCS Manager.

This section describes:

- Initial setup and configuration
- Operation, that is, adding networks for additional tenants

For configuration details, see Configuration Guide:
http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/2.1/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide_2_1.html

Initial Setup

For initial setup, follow these steps:

1. Create a vNIC connection policy in Cisco UCS Manager.
2. Create a server BIOS policy.
3. Clone an existing service profile.
4. Install the VEM software on the ESX server.
5. Install the plug-in into vCenter.

Create a Dynamic vNIC Connection Policy

To define the dynamic vNIC connection policy that vNICs created from a vNIC template should use, follow these steps in Cisco UCS Manager:

1. Log in to Cisco UCS Manager.
2. Click the **LAN** tab in the left navigation pane and click **LAN** > **Policies** > **root** > **Sub-organizations** (name of the suborganization if applicable) > **Dynamic vNIC Connection Profile**.
3. Right-click and choose **Create Dynamic vNIC Connection Policy** to start the wizard.
4. Type a name and description for the vNIC connection policy. Choose VMWare from the Adapter Policy drop-down menu. Choose the Protected option. Click **OK**.

Note

- The Protected option allows the vNIC to use both fabric A and fabric B.
- With Cisco UCS C-Series servers, the number of dynamic vNICs that can be used depends on the hardware in use. For more information, see "VM-FEX Virtual Interfaces" section on page 198.

Figure 117 Creating Dynamic vNIC Policy

Create a Server BIOS Policy

To define the BIOS policy for a service profile that supports VM-FEX on ESXi, complete the following steps in Cisco UCS Manager:

1. Click the **Server** tab in the left navigation pane, and choose **Server** > **Policies** > **root** > **Sub-organizations** (name of the suborganization if applicable) > **BIOS Policies**.
2. Right-click and choose **Create BIOS Policy** to start the wizard.
3. Type a name for the policy and retain the platform defaults.

Figure 118 Creating BIOS Policy

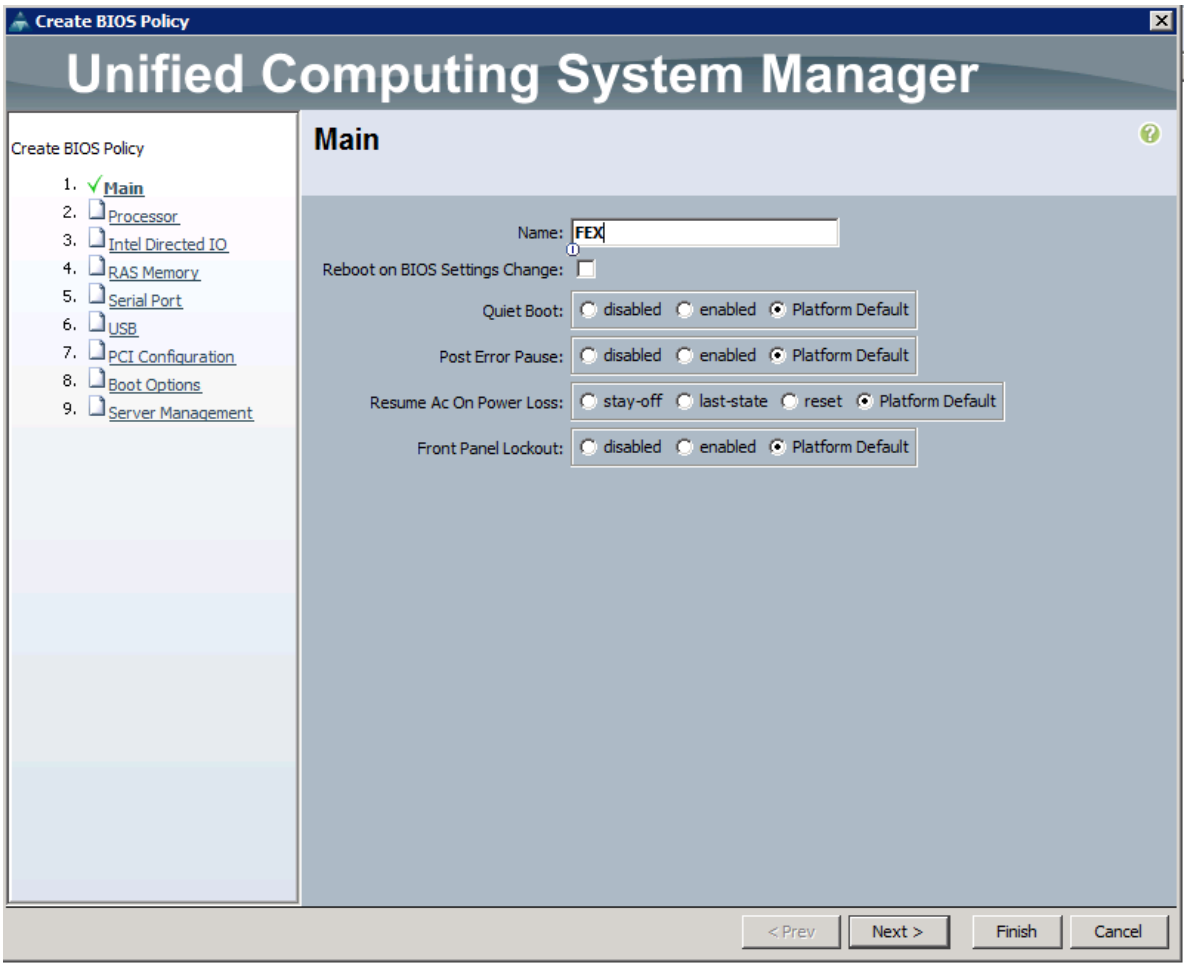

4. For Virtualization Technology (VT) and Direct Cache Access, choose enabled.

Figure 119 Processor Settings for BIOS Policy

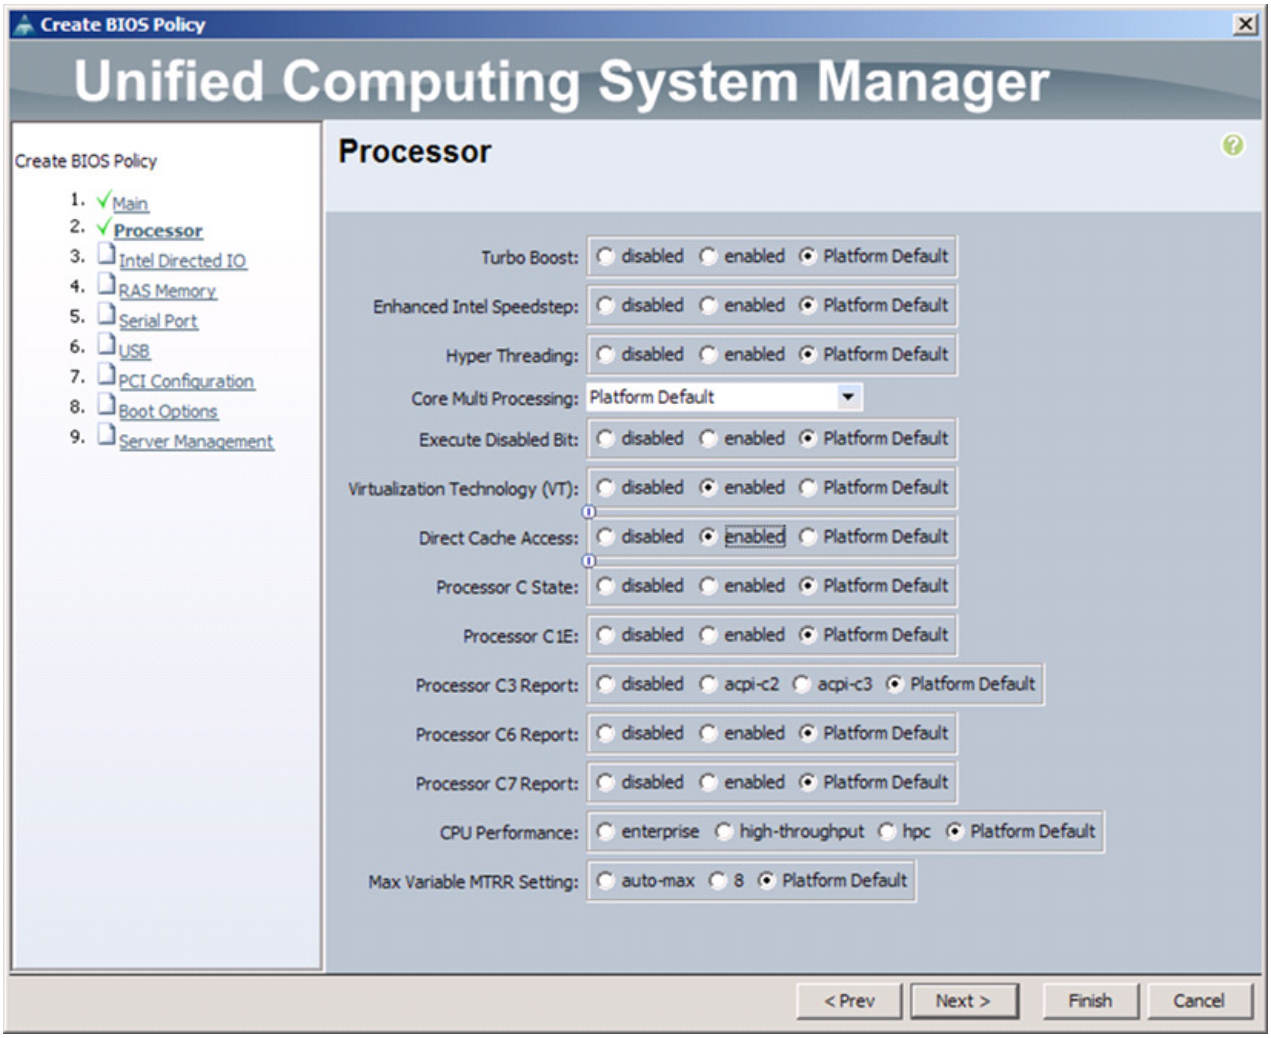

5. Click **Next**.
6. For VT For Directed IO, choose enabled.

Figure 120 Intel Directed IO Settings for BIOS Policy

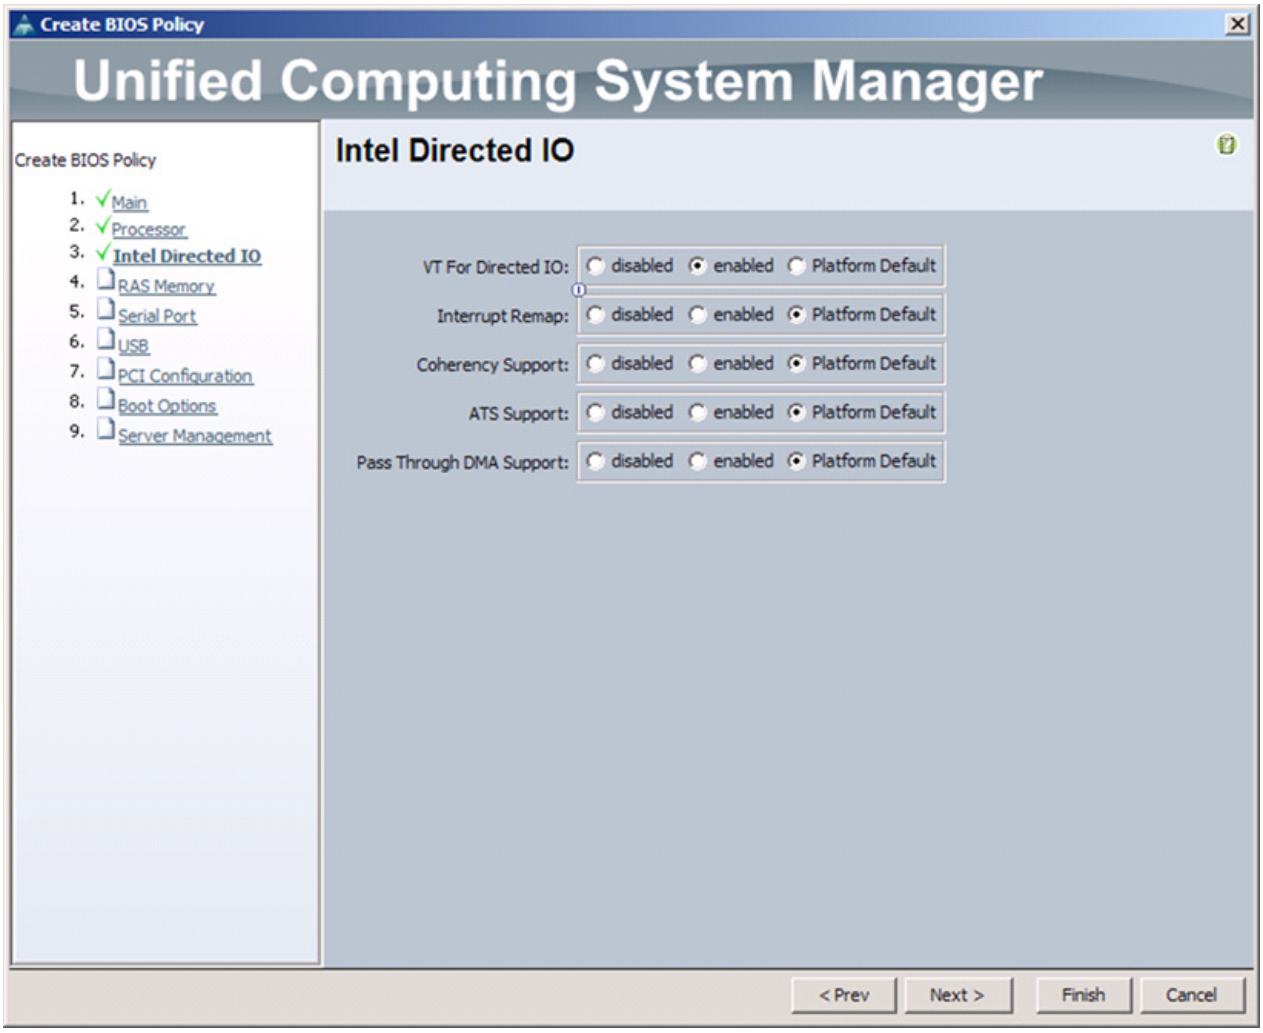

7. Click **Next**.
8. The remaining sections of the Create BIOS Policy wizard (RAS Memory, Serial Port, USB, PCI Configuration, Boot Options, and Server Management) can retain the Platform Default option. Click **Next** on each of these windows and then click **Finish** to complete the wizard.

Create a VM-FEX Enabled Service Profile Template

To create a Cisco UCS service profile using VM-FEX, clone a previously defined Cisco UCS service profile and apply the dynamic vNIC and BIOS policies by following these steps in the Cisco UCS Manager:

1. Click the **Servers** tab in the left navigation pane and expand the Service Profile Templates.
2. Right-click VM-Host-Infra-Fabric-A and choose **Create a Clone**.
3. Type a clone name and choose an organizational owner for the new service profile template.
4. Click **OK**.

Figure 121 Cloning Service Profile Template

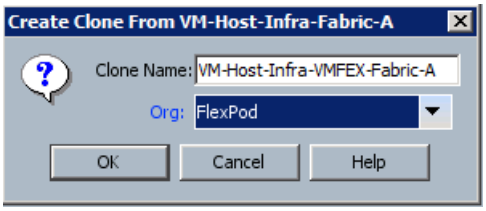

5. Click **OK.**
6. When notified that the service profile clone was successfully created, check the checkbox to navigate to the Service Profile Template. Click **OK**.
7. Click the **Network** tab and choose **Change Dynamic vNIC Connection Policy** under the Actions section of the working pane. The Change Dynamic vNIC Connection Policy form appears.
8. Choose Use a Dynamic vNIC Connection Policy from the drop-down menu and the previously created Dynamic vNIC policy. Click **OK**.

Figure 122 Changing the Dynamic vNIC Connection Policy

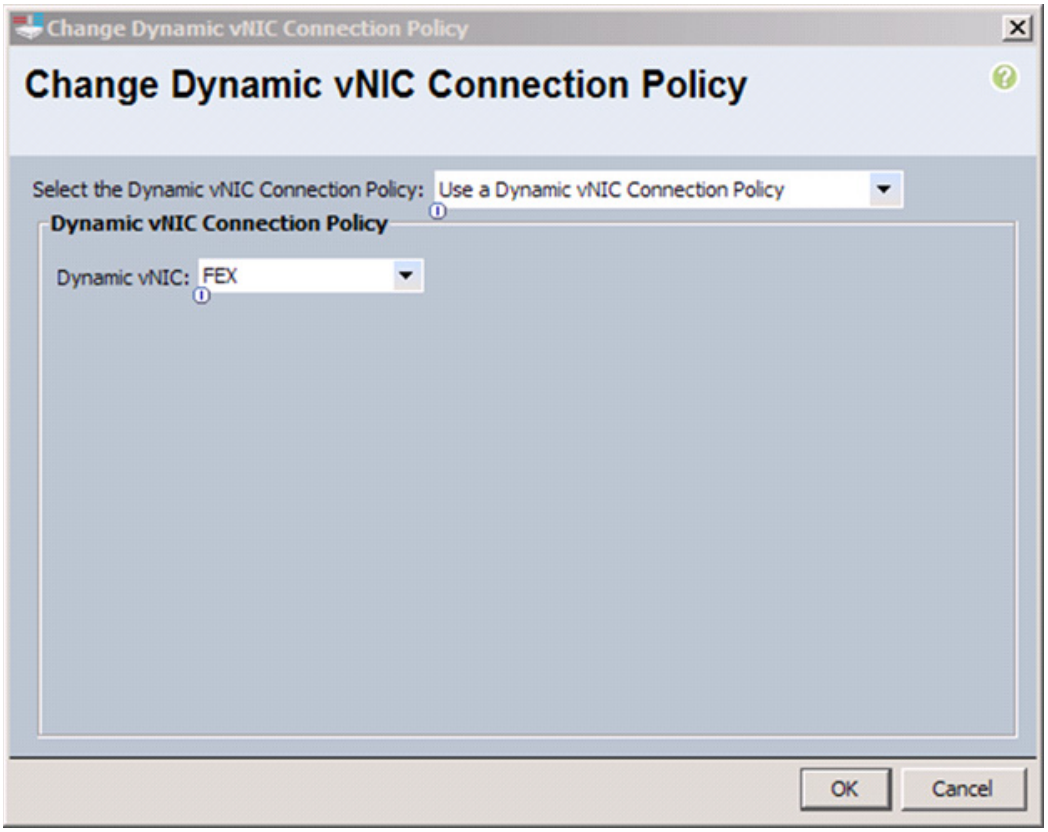

9. Click **OK** when notified that the vNIC connection policy was successfully modified.
10. From the Service Template properties window, click the **Policies** tab.
11. Expand the BIOS Policies in the Policies section of the working pane.
12. Choose the previously defined FEX BIOS policy and click **Save Changes**.

13. Click **OK** when notified that the Service Profile Template was successfully modified.

Figure 123 Choosing a BIOS Policy

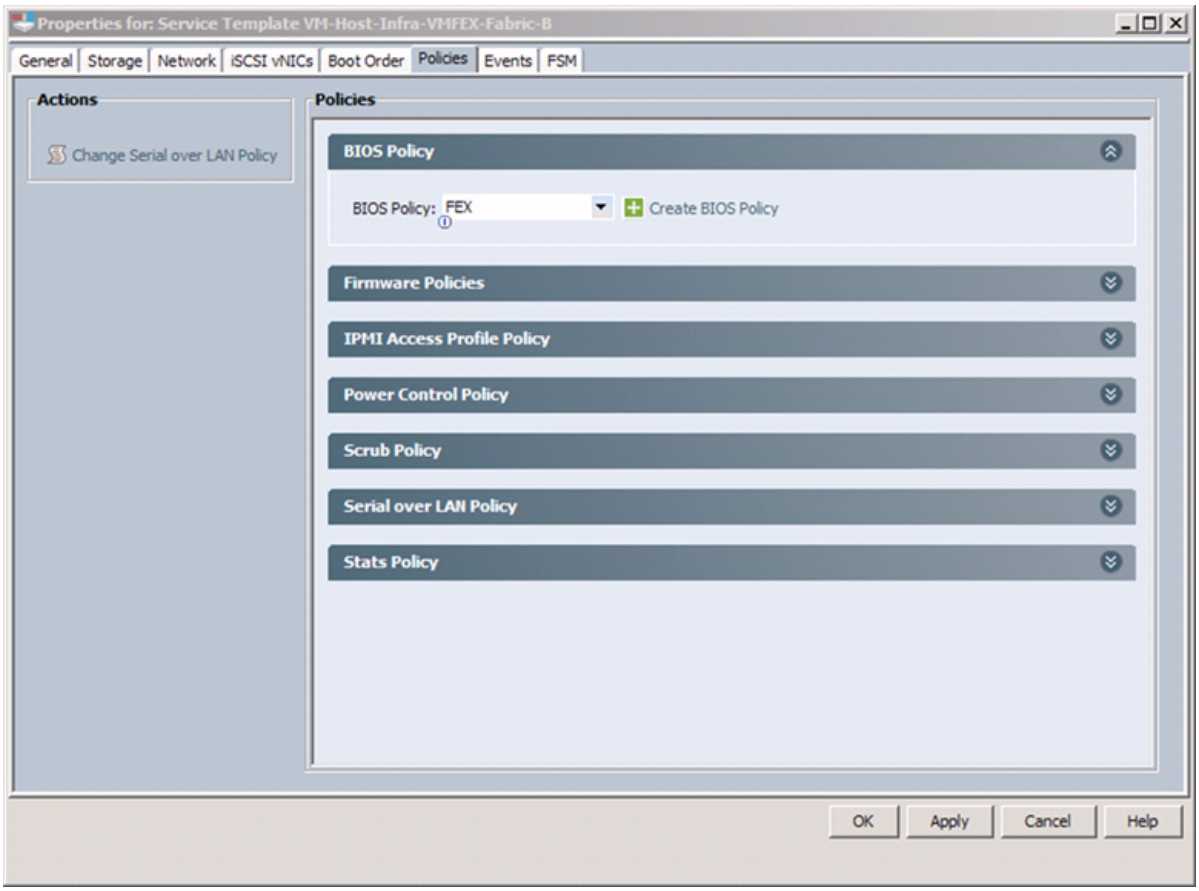

Create VM-FEX Service Profile

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Service Profile Templates** > **Service Template VM-Host-Infra-VMFEX-Fabric-A**.
3. Right-click VM-Host-Infra-FEX-Fabric-A and choose **Create Service Profiles** from Template.
4. Enter VM-Host-FEX-0 as the service profile prefix.
5. Enter 1 as the suffix starting number.
6. Enter 1 for the number of instances.
7. Click **OK** to create the service profile.

Figure 124 Creating Service Profile from Template

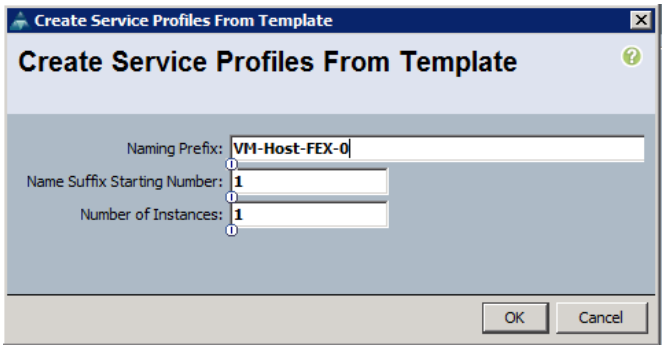

8. Click **OK** in the confirmation message.
9. Verify that the service profile VM-Host-FEX-01 has been created. The service profile is automatically associated with the servers in their assigned server pools.

Install and Set up VMware ESXi

See "FlexPod VMware ESXi 5.1Update1 FCoE on 7-Mode" section on page 106 to install and completely set up VMware ESXi version 5.1Update1 on the two ESXi hosts. After ESXi setup is complete, add the two new hosts to VMware vCenter.

Download Cisco VEM Software Bundle

To download the Cisco UCS B-Series or C-Series server drivers, follow these steps:

Note The following bundle was used during validation cisco-vem-v151-5.1-1.1.1.1.vib.

1. Open a Web browser on the management workstation and navigate to the following Cisco Download Software pages:
   a. Downloads Home > Products > Servers - Unified Computing > Cisco UCS B-Series Blade Server Software > Unified Computing System (UCS) Drivers-2.1(3)
   b. Downloads Home > Products > Servers - Unified Computing > Cisco UCS C-Series Rack-Mount UCS-Managed Server Software > Unified Computing System (UCS) Drivers-1.4(5b)
2. Follow the steps necessary to download the software bundles located on the ISO image.
3. Mount the ISO image and copy the appropriate vib file from the **VMware** > **VM-FEX** > **Cisco** directory to the local machine.
4. From the vCenter vSphere Client, choose the infra_datastore_1 in the **Inventory** > **Datastores and Datastore Clusters** navigation menu.
5. Under the Basic Tasks choose Browse this Datastore
6. Choose the root folder (/) and click the third button at the top to add a folder.
7. Name the folder VM-FEX and click **OK**.
8. On the left, choose the VM-FEX folder.
9. Click the fourth button at the top and choose Upload File.

10. Navigate to the cisco-vem-v151-5.1-1.1.1.1.vib file and click **Open**.
11. Click Yes to upload the .vib file to infra_datastore_1.

    The VM-FEX file should now appear in the VM-FEX folder in the datastore.

Install the FEX Virtual Ethernet Module on Each ESXi Host

To install the Virtual Ethernet Module (VEM) on the ESXi hosts, follow these steps:

1. Open the VMware vSphere CLI command prompt.
2. For each ESXi host in the VMware vSphere CLI, run the following command:

```
esxcli -s <host_ip> -u root -p <host_password> software vib install -v
/vmfs/volumes/infra_datastore_1/VM-FEX/cisco-vem-v151-5.1-1.1.1.1.vib
```

Integrate Cisco UCS with vCenter

To integrate Cisco UCS Manager and vCenter, follow these steps:

1. Log in to the Cisco UCS Manager.
2. In the navigation pane, click the **VM** tab, and in the VM tab, expand the All folder. Choose the VMware node, and in the Working Area, click the **General** tab.

Figure 125 Configuring VMware Integration

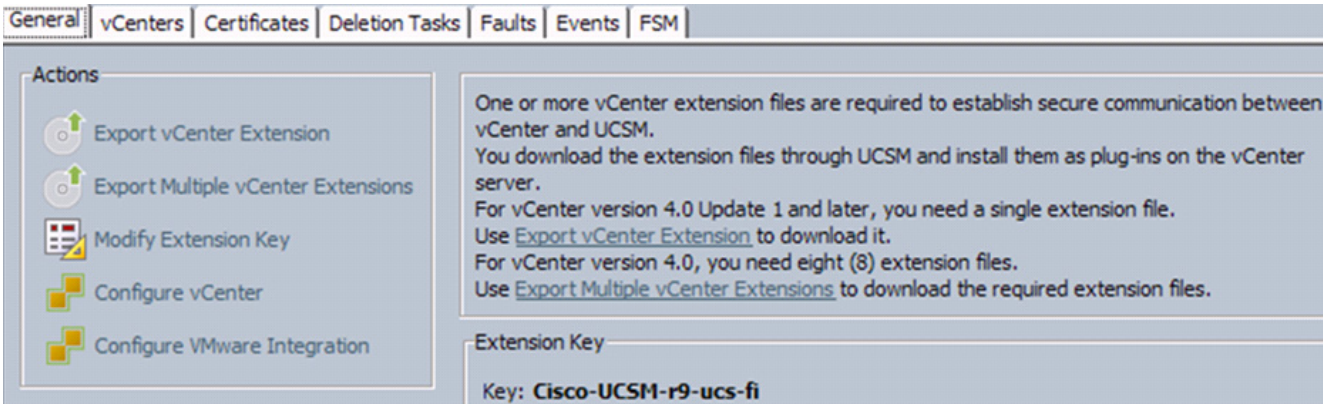

3. Choose **Configure VMware Integration** in the Actions area to start the Configuration wizard.
4. Follow the instructions and click **Export** and complete the steps to install the UCSM extension file in vCenter.

Figure 126 Exporting vCenter Extension Plug-in

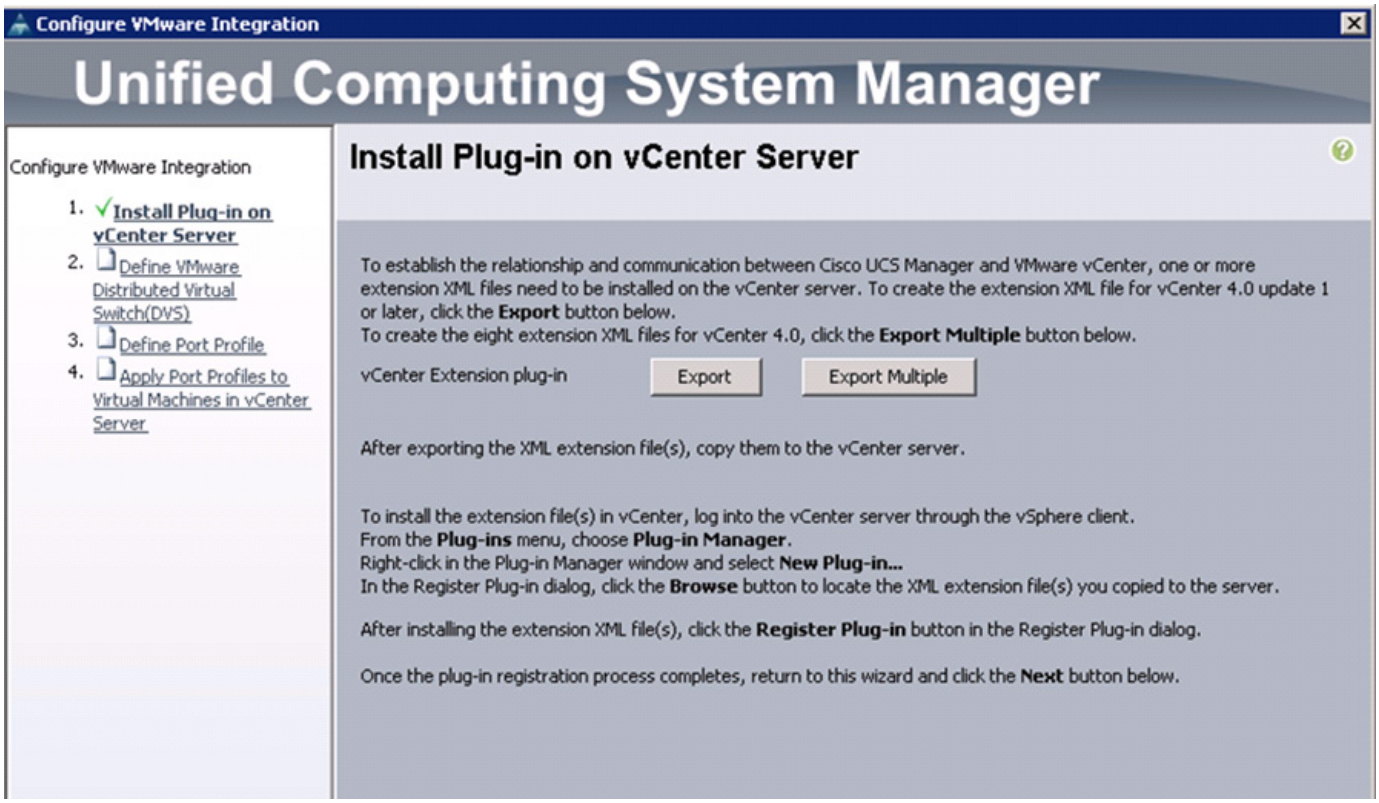

5. Click **Next**.
6. Enter the VMware vCenter Server name, vCenter Server host name or IP address, vCenter data center name, DVS folder, and DVS name.
7. Click **Next**.

Figure 127 Defining VMware DVS

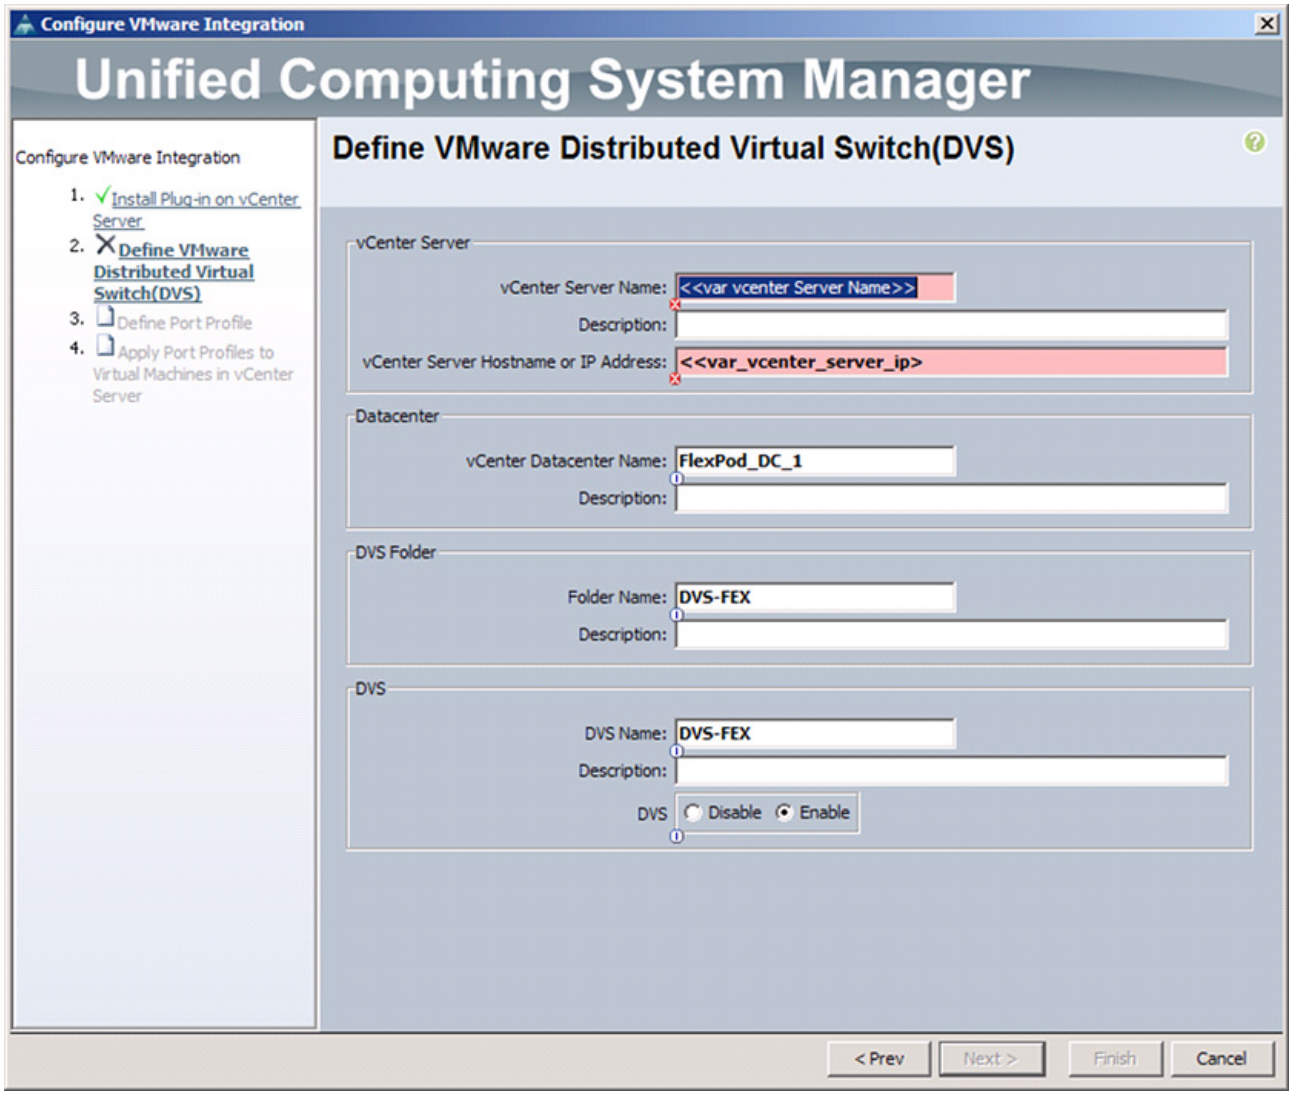

8. Create the FEX-MGMT port profile, select the IB-MGMT-VLAN, and mark it as the native VLAN.

Figure 128 Defining Port Profile for VMware Integration

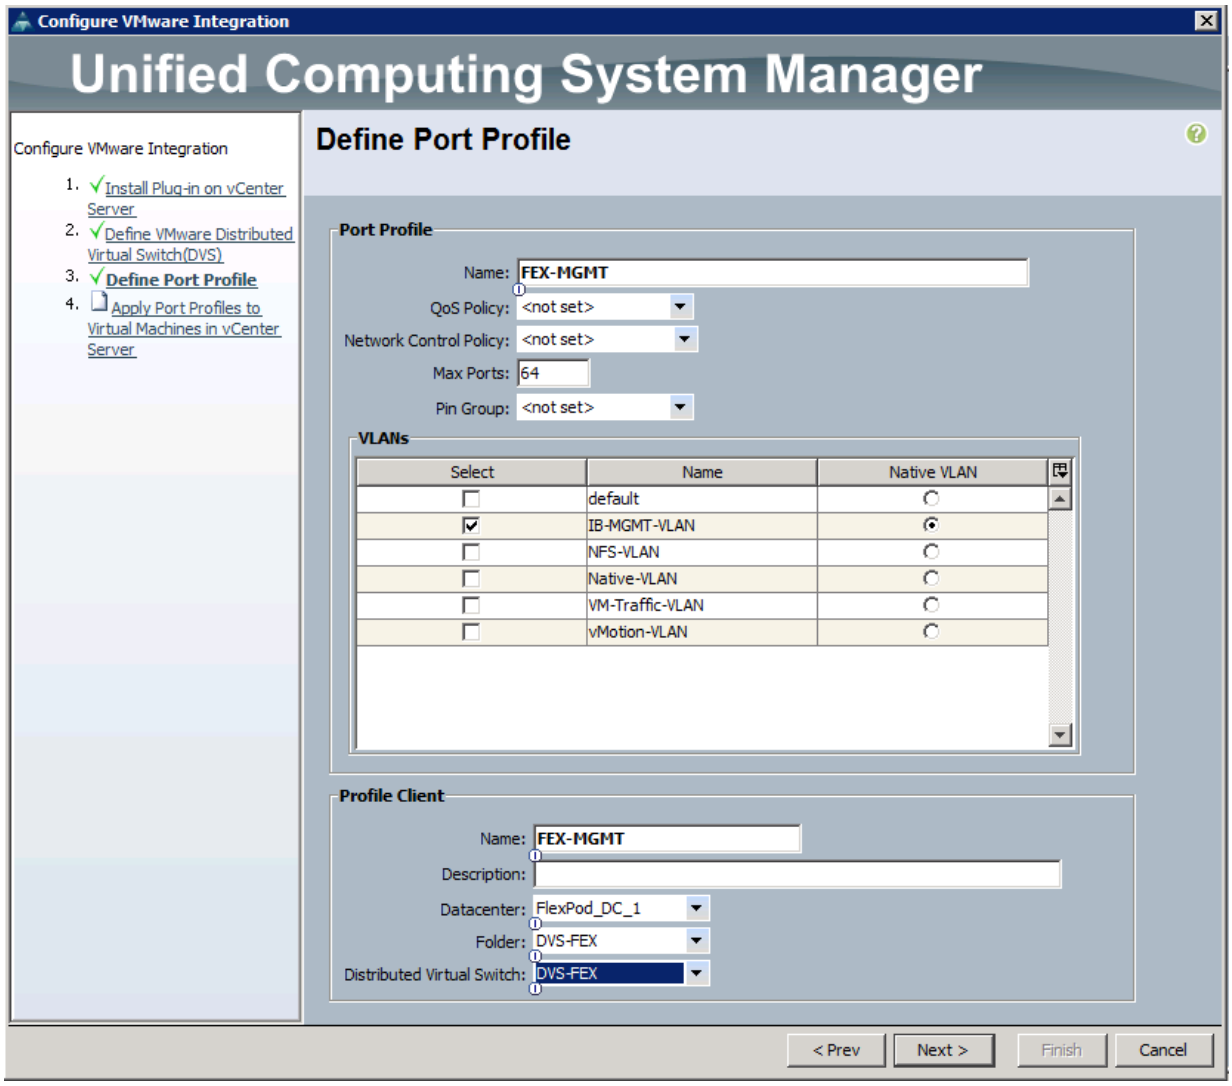

9. Click **Next**.
10. Click **Finish**.
11. When finishing the wizard, the Cisco UCS Manager connects to vCenter and adds the plug-in.

Figure 129 VMware Integration is Successful

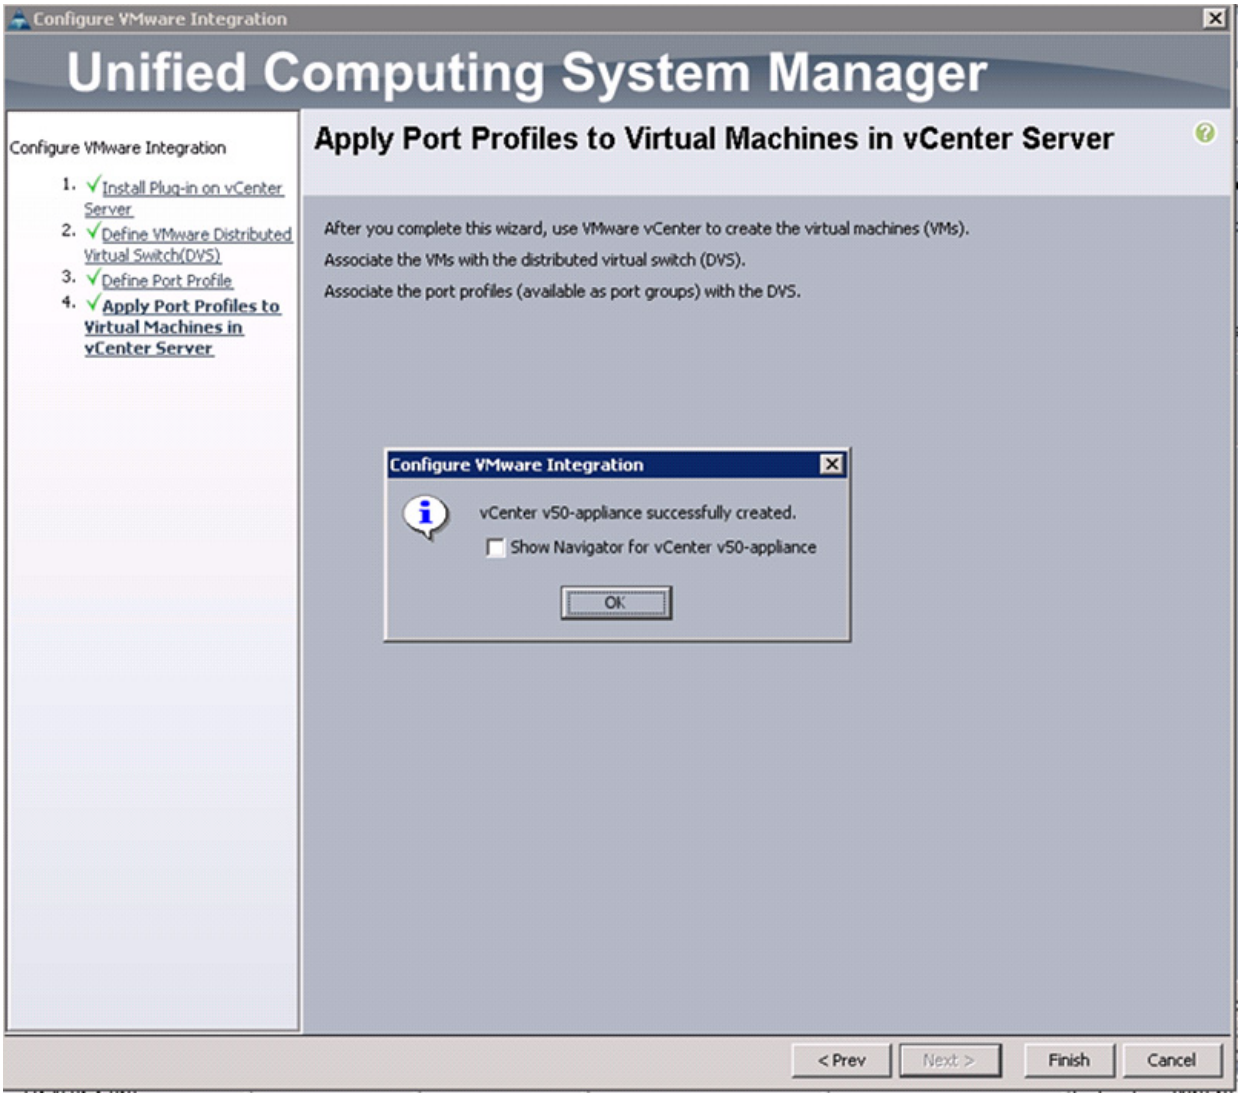

Note The ESXi host will require additional hypervisor vNICs to support VMware vMotion, and NFS traffic uses the generic port-profile creation steps documented in section "Standard Operations" to establish a FEX-vMotion and FEX-NFS Port Profile.

Validate Setting in VMware vCenter

To validate the successful installation of the Cisco UCS Manager plug-in, follow these steps:

1. Log in to the vCenter Server.
2. In the Main menu, choose **Plug-ins** > **Manage Plug-ins**.

   The popup windows shows that the Cisco UCS Manager is already integrated in vCenter.

Figure 130 Cisco UCS Manager Integrated in vCenter

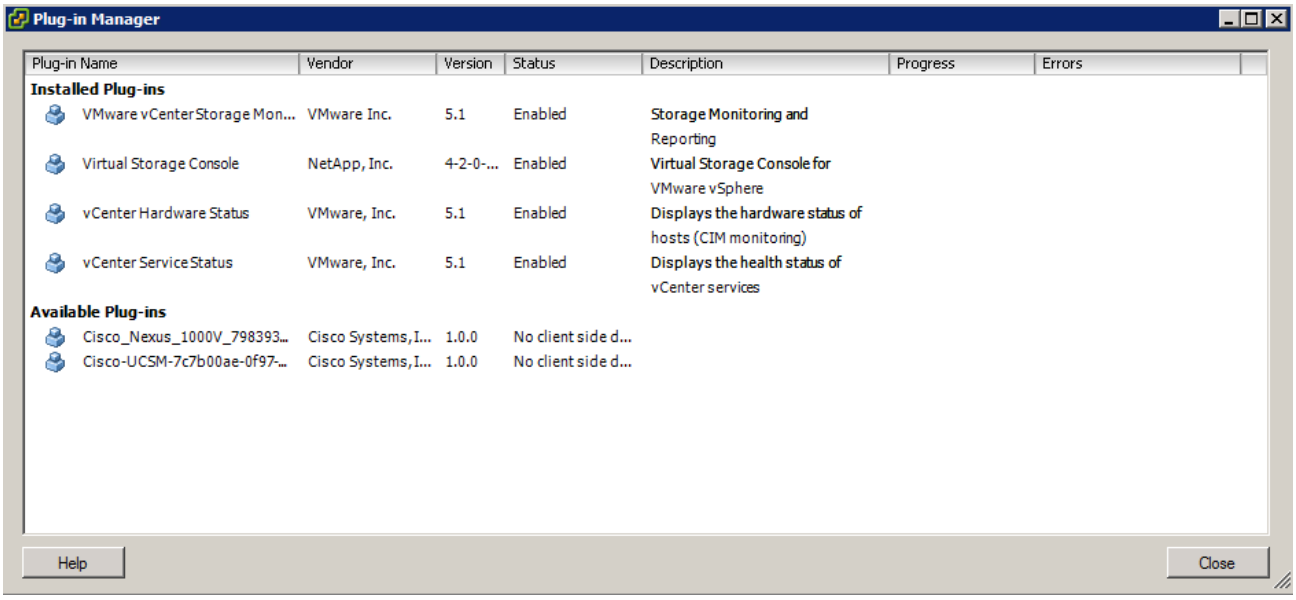

3. Click **Inventory** > **Networking** to see FEX added to distributed switch from Cisco UCS Manager.

Figure 131 Verifying FEX-DVS in Cisco UCS Manager

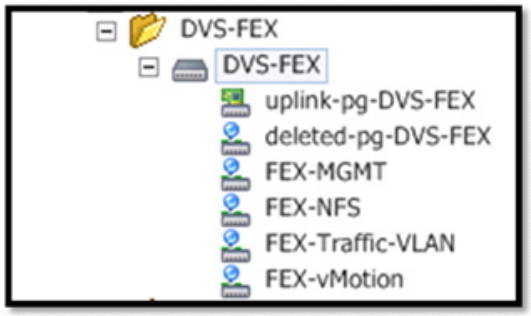

Standard Operations

The VM-FEX environment supports the addition of port profiles to the distributed switch. The following section describes how to add these distributed port groups.

Add Distributed Port Group to the VDS (vSphere Distributed Switch)

Port Profiles

Port profiles contain the properties and settings that you can use to configure virtual interfaces in Cisco UCS for VM-FEX. The port profiles are created and administered in Cisco UCS Manager. After a port profile is created, assigned to, and actively used by one or more distributed virtual switches (DVSs), any changes made to the networking properties of the port profile in Cisco UCS Manager are immediately applied to those DVSs.

In VMware vCenter, a port profile is represented as a port group. Cisco UCS Manager pushes the port profile names to VMware vCenter, which displays the names as port groups. None of the specific networking properties or settings in the port profile is visible in VMware vCenter. You must configure at least one port profile client for a port profile if you want Cisco UCS Manager to push the port profile to VMware vCenter.

Port Profile Client

The port profile client determines the DVSs to which a port profile is applied. By default, the port profile client specifies that the associated port profile applies to all DVSs in VMware vCenter. However, you can configure the client to apply the port profile to all DVSs in a specific data center or data center folder or to only one DVS.

Create a VM-FEX Port Profile

Follow these steps to create VM-FEX port profiles for use on the Cisco UCS distributed virtual switch.

1. Log in to Cisco UCS Manager.
2. Click the **VM** tab.
3. Right-click **Port Profile** > **Create Port Profile**.
4. Enter the name of the Port Profile.
5. (Optional) Enter a description.
6. (Optional) Choose a QoS policy.
7. (Optional) Choose a network control policy.
8. Enter the maximum number of ports that can be associated with this port profile. The default is 64 ports.

Note The maximum number of ports that can be associated with a single DVS is 4096. If the DVS has only one associated port profile, that port profile can be configured with up to 4096 ports. However, if the DVS has more than one associated port profile, the total number of ports associated with all of those port profiles combined cannot exceed 4096.

9. (Optional) Choose High Performance.

Note Select None—Traffic to and from a virtual machine passes through the DVS.

Select High Performance— Traffic to and from a virtual machine bypasses the DVS and hypervisor and travels directly between the virtual machines and a virtual interface card (VIC) adapter.

10. Choose the VLAN.
11. Choose Native-VLAN.
12. Click **OK**.

Figure 132 Creating Port Profile

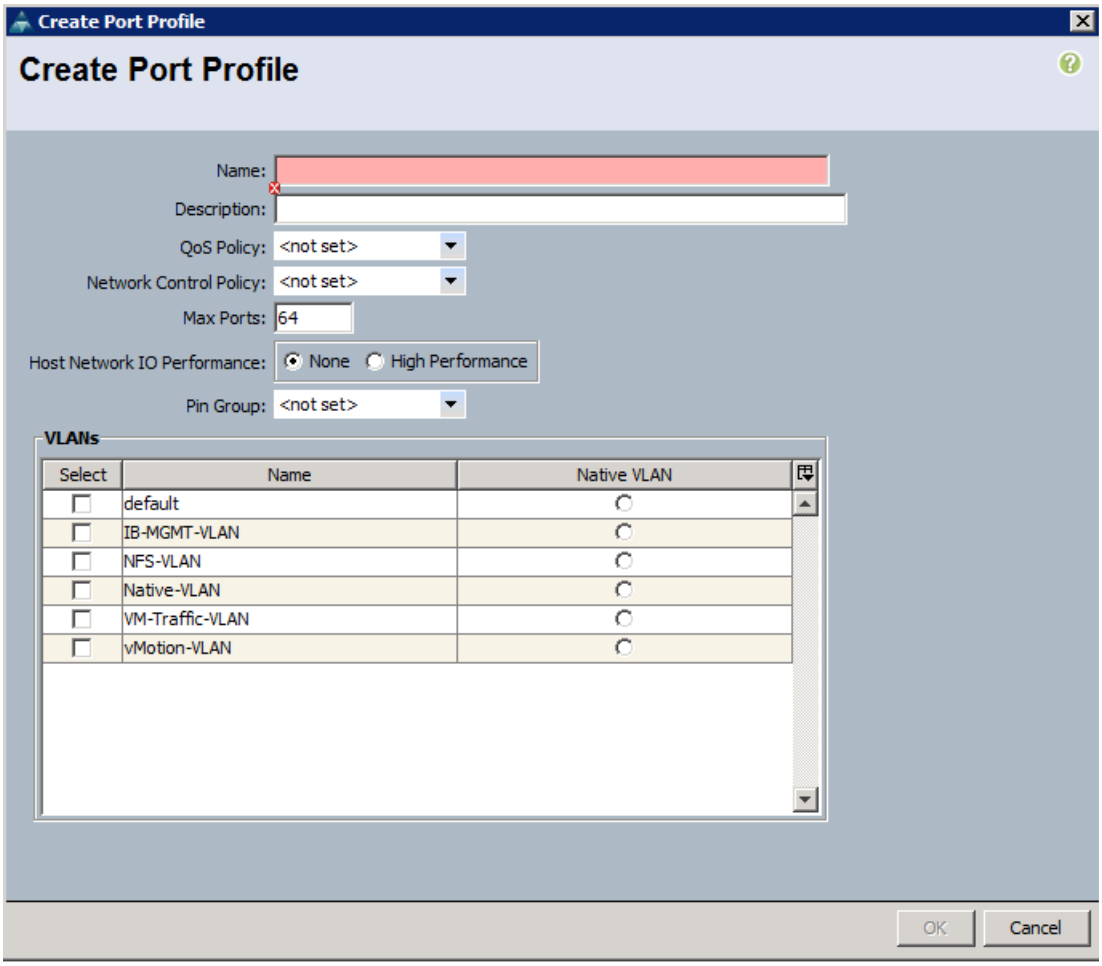

Or

Figure 133 Creating Port Profile with High Performance

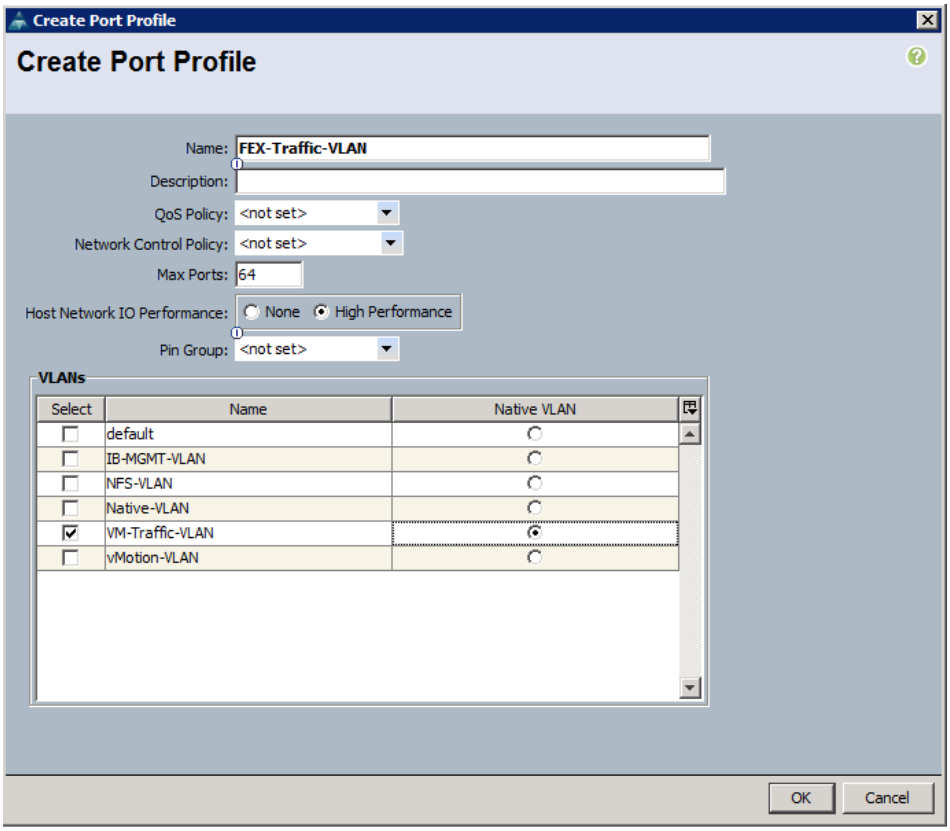

The port profile created will appear in the working pane.

Create the Port Profile Client

To create the client profile for use in the Cisco UCS virtual distributed switch, follow these steps:

1. In the navigation pane under the VM tab, expand **All** > **Port Profiles**. Right-click the Port Profile and click **Create Profile Client**.
2. Choose the data center created in your vCenter Server, folder, and distributed virtual switch created in section "Integrate Cisco UCS with vCenter."
3. Click **OK**.

Figure 134 Creating Profile Client

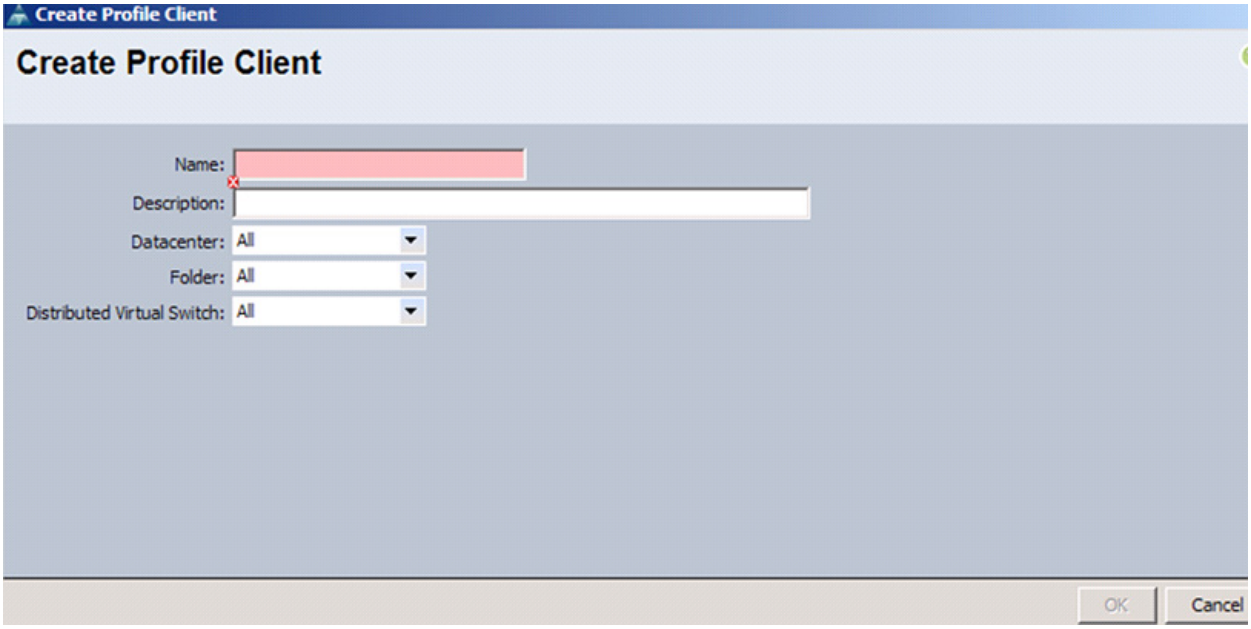

Or

Figure 135 Creating Profile Client for DVS-FEX

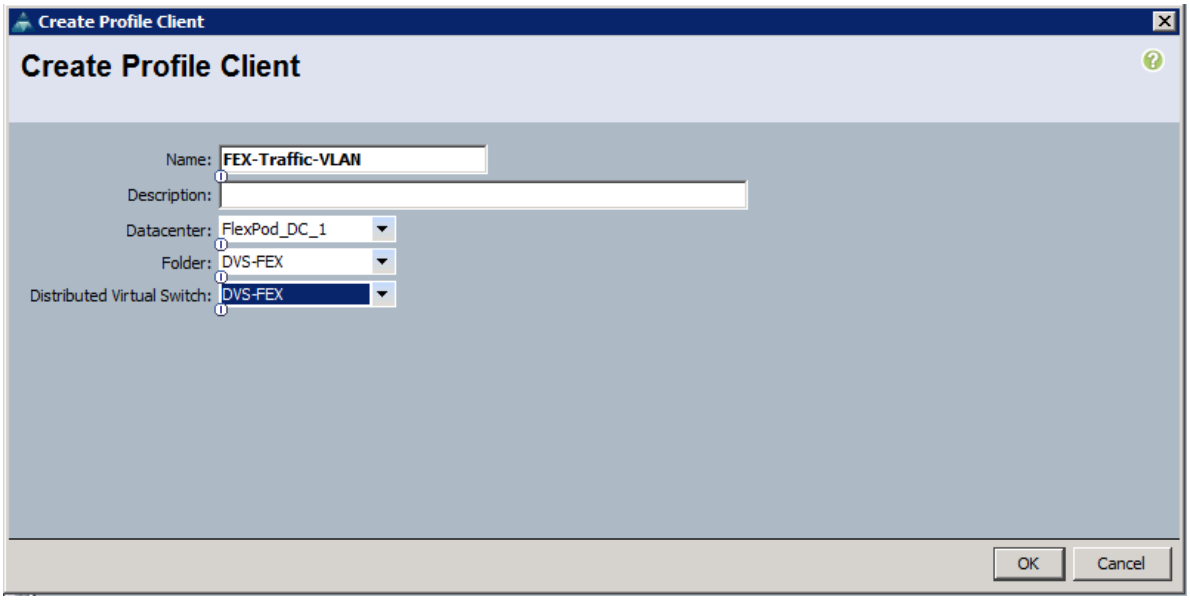

4. The client profile created will appear in your distributed virtual switch DVS-FEX in vCenter as a port group.
5. Repeat these steps as necessary for the workloads in the environment.

Migrate Networking Components for ESXi Hosts to Cisco DVS-FEX

vCenter Server VM

To migrate the networking components for the ESXi hosts to the Cisco FEX-DVS, follow these steps:

1. In the VMware vSphere client connected to vCenter, choose **Home** > **Networking**.
2. Expand the vCenter, DataCenter, and DVS-FEX folders. choose the DVS-FEX switch.
3. Under Basic Tasks for the vSphere distributed switch, choose Add a Host.
4. For both hosts, choose vmnic1 and choose the uplink-pg-DVS-FEX Uplink port group. Click Next.

Figure 136 Selecting Physical Adapters to Add to the Vsphere Distributed Switch

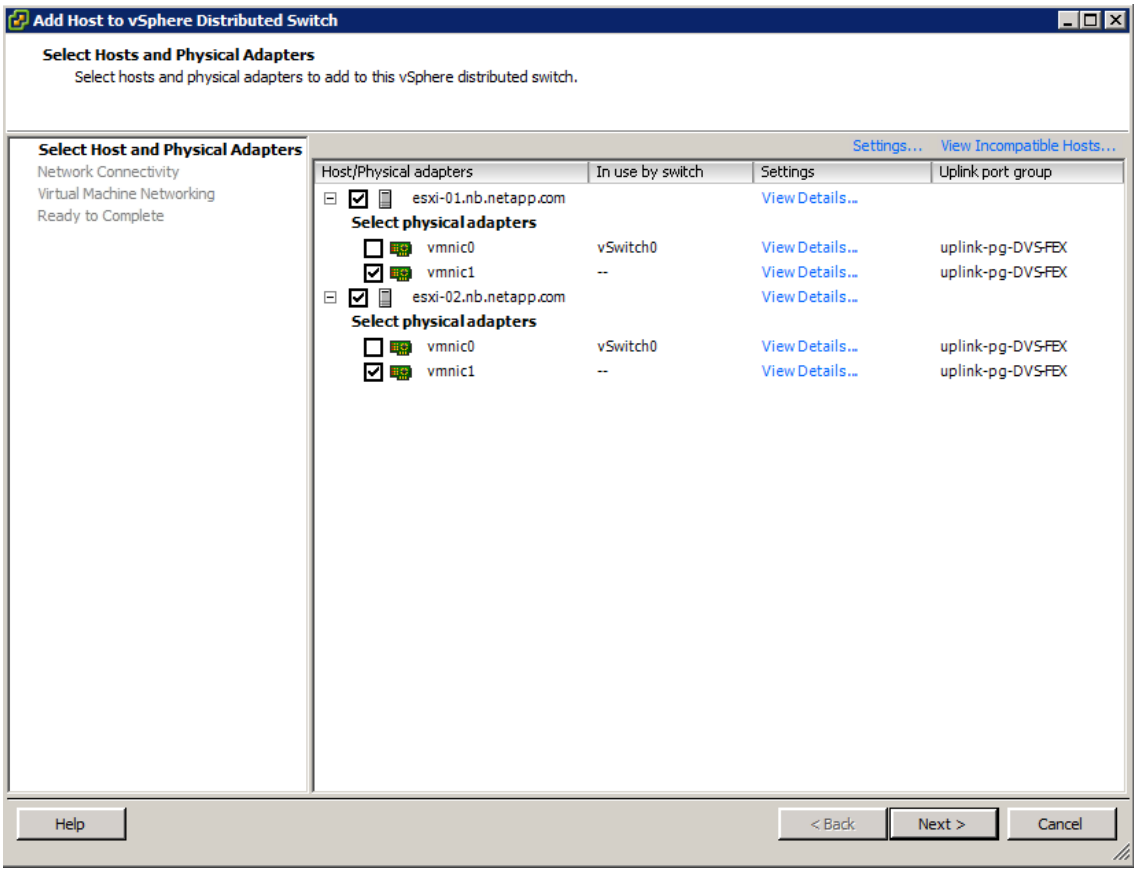

5. For all VMkernel ports, choose the appropriate destination Port Group from the Cisco DVS-FEX. Click **Next**.

Figure 137 Selecting Destination Port Groups

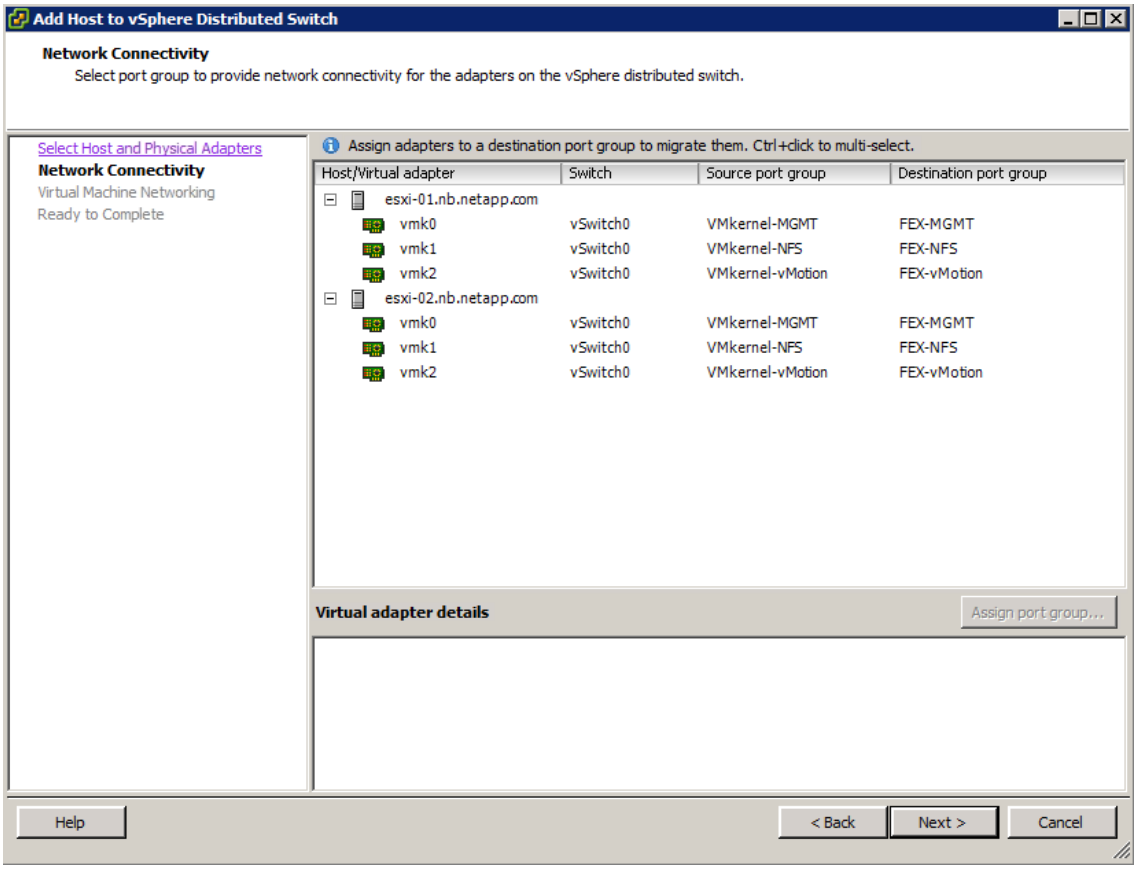

6. Check the **Migrate Virtual Machine Networking** check box. Expand each VM and choose the port groups for migration individually. Click **Next**.

Figure 138 Selecting Port Groups on Each VM to Migrate Individually

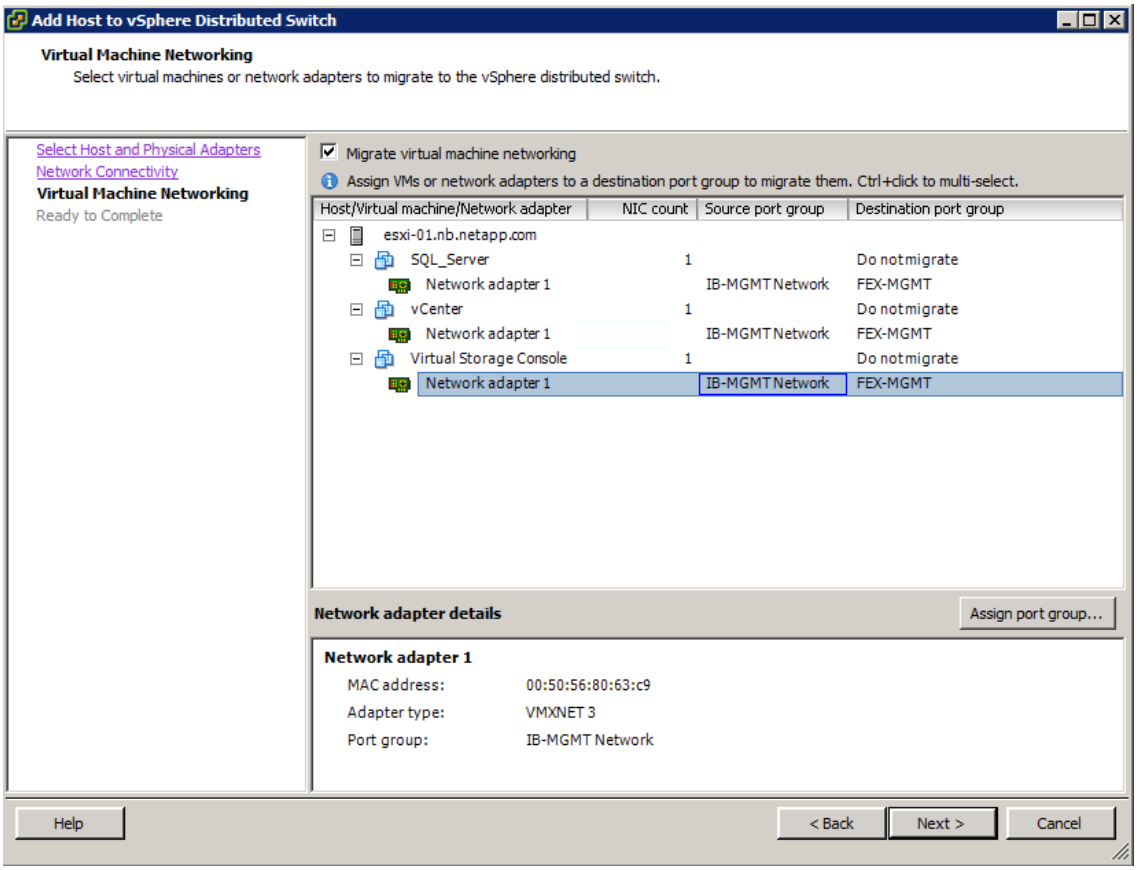

7. Click **Finish**. Wait for the migration process to complete.
8. In the vSphere Client window, choose **Home > Hosts and Clusters**.
9. Choose the first ESXi host and click the **Configuration** tab. In the Hardware field, choose Networking.
10. Make sure that vSphere Standard Switch is selected at the top next to View. vSwitch0 should not have any active VMkernel or VM Network ports on it. On the upper-right side of vSwitch0, click **Remove**.
11. Click **Yes**.
12. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View.
13. Click **Manage Physical Adapters**.
14. In the uplink-pg-DVS-FEX field click **Add NIC**.
15. Choose vmnic0 and click **OK**.
16. Click **OK** to close the Manage Physical Adapters window. Two uplinks should now be present.
17. Choose the second ESXi host and click the Configuration tab. In the Hardware field, choose Networking.
18. Make sure vSphere Standard Switch is selected at the top next to View. vSwitch0 should have no active VMkernel or VM Network ports on it. On the upper-right side of vSwitch0, click **Remove**.

19. Click **Yes**.
20. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch**.
21. Click **Manage Physical Adapters.**
22. In the uplink-pg-DVS-FEX field click **Add NIC**.
23. Choose vmnic0 and click **OK**.
24. Click **OK** to close the Manage Physical Adapters window. Two uplinks should now be present.

VM-FEX Virtual Interfaces

In a blade server environment, the number of vNICs and vHBAs configurable for a service profile is determined by adapter capability and the amount of virtual interface (VIF) namespace available in the adapter. In Cisco UCS, portions of VIF namespace are allotted in chunks called VIFs. Depending on your hardware, the maximum number of VIFs is allocated on a predefined, per-port basis.

The maximum number of VIFs varies based on hardware capability and port connectivity. For each configured vNIC or vHBA, one or two VIFs are allocated. Standalone vNICs and vHBAs use one VIF, and failover vNICs and vHBAs use two.

The following variables affect the number of VIFs available to a blade server, and therefore, the number of vNICs and vHBAs you can configure for a service profile.

- The maximum number of VIFs supported on your fabric interconnect
- How the fabric interconnects are cabled
- If the fabric interconnect and IOM are configured in fabric port channel mode

For more information about the maximum number of VIFs supported by your hardware configuration, refer to the Cisco UCS 6100 and 6200 Series Configuration Limits for Cisco UCS Manager for your software release. Table 25 and Table 26 reference these limits.

Table 25 VM-FEX Environment Configuration Limits

Feature	Cisco UCS 6200 Series Fabric Interconnect
Host per DVS	52
DVSs per Cisco UCS Domain	1
vCenter Server units per Cisco UCS Domain	4
Port profiles per Cisco UCS Domain	512
Dynamic ports per port profile	4096
Dynamic ports per DVS	4096

Table 26 Cisco UCS Fabric Interconnect and Cisco UCS C-Series Server VIF Support

Acknowledge Link Between FEX and FI	Maximum VIFs (vNICs+vHBAs) per VIC Adapter in Single-Wire Management	Maximum VIFs (vNICs+vHBAs) per VIC Adapter in Dual-Wire Management
1	12	13
2	27	28

Table 26 Cisco UCS Fabric Interconnect and Cisco UCS C-Series Server VIF Support

Acknowledge Link Between FEX and FI	Maximum VIFs (vNICs+vHBAs) per VIC Adapter in Single-Wire Management	Maximum VIFs (vNICs+vHBAs) per VIC Adapter in Dual-Wire Management
4	57	58
8	117	118

Note

- For a non-VIC adapter the maximum number of vNICs is two and the maximum number of vHBAs is two.
- If the server in single-wire mode has two VIC adapters, the maximum number of VIFs (vNICs + vHBAs) available for the second adapter would be same as for an adapter in a dual-wire mode server.
- For more information on Cisco UCS C-Series Server integration into UCSM, see: http://www.cisco.com/en/US/docs/unified_computing/ucs/c-series_integration/ucsm2.1/b_UCSM2-1_C-Integration.pdf

Server-Side Flash—NetApp Flash Accel with Fusion-IO

This section provides detailed instructions for installing NetApp Flash Accel software on NetApp Virtual Storage Console to manage server-side Flash. For the Flash Accel plug-in to install properly on the NetApp Virtual Storage Console, some network reconfiguration of the environment is necessary.

Create Out-of-Band Management VLAN on Cisco Nexus 5548UP Switches

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

To create out-of-band management VLAN on both the Cisco Nexus switches, follow these steps:

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

```
config t
vlan <<var_oob-mgmt_vlan_id>>
name OOB-MGMT-VLAN
exit
copy run start
```

Configure Port Channels with the Out-of Band Management VLAN on Cisco Nexus 5548UP Switches

Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

The out-of-band management VLAN is added to the vPC peer-link and the port channels between the Cisco Nexus 5548UP switches and the Cisco UCS Fabric Interconnects.

From the global configuration mode, run the following commands:

```
interface Po10
switchport trunk allowed vlan add 3170
exit
interface Po13
```

```
switchport trunk allowed vlan add 3170
exit
interface Po14
switchport trunk allowed vlan add 3170
exit
copy run start
```

Note Add the out-of-band management VLAN to the switch uplink that has been configured for the FlexPod environment.

Create Out-of-Band Management VLAN on the Cisco UCS

Cisco UCS Manager

1. Using a web browser, log in to the Cisco UCS Manager as admin.
2. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
3. Select **LAN** > **LAN Cloud**.
4. Right-click VLANs.
5. Choose **Create VLANs**.
6. Enter **OOB-MGMT-VLAN** as the name of the VLAN to be used for out-of-band management traffic.
7. Make sure the **Common/Global** radio button is selected for the scope of the VLAN.
8. Enter **<<var_oob-mgmt_vlan_id>>** as the ID of the out-of-band management VLAN.
9. Click **None** radio button for Sharing Type.
10. Click **OK** and then click **OK** again.

Figure 139 Creating Management VLAN

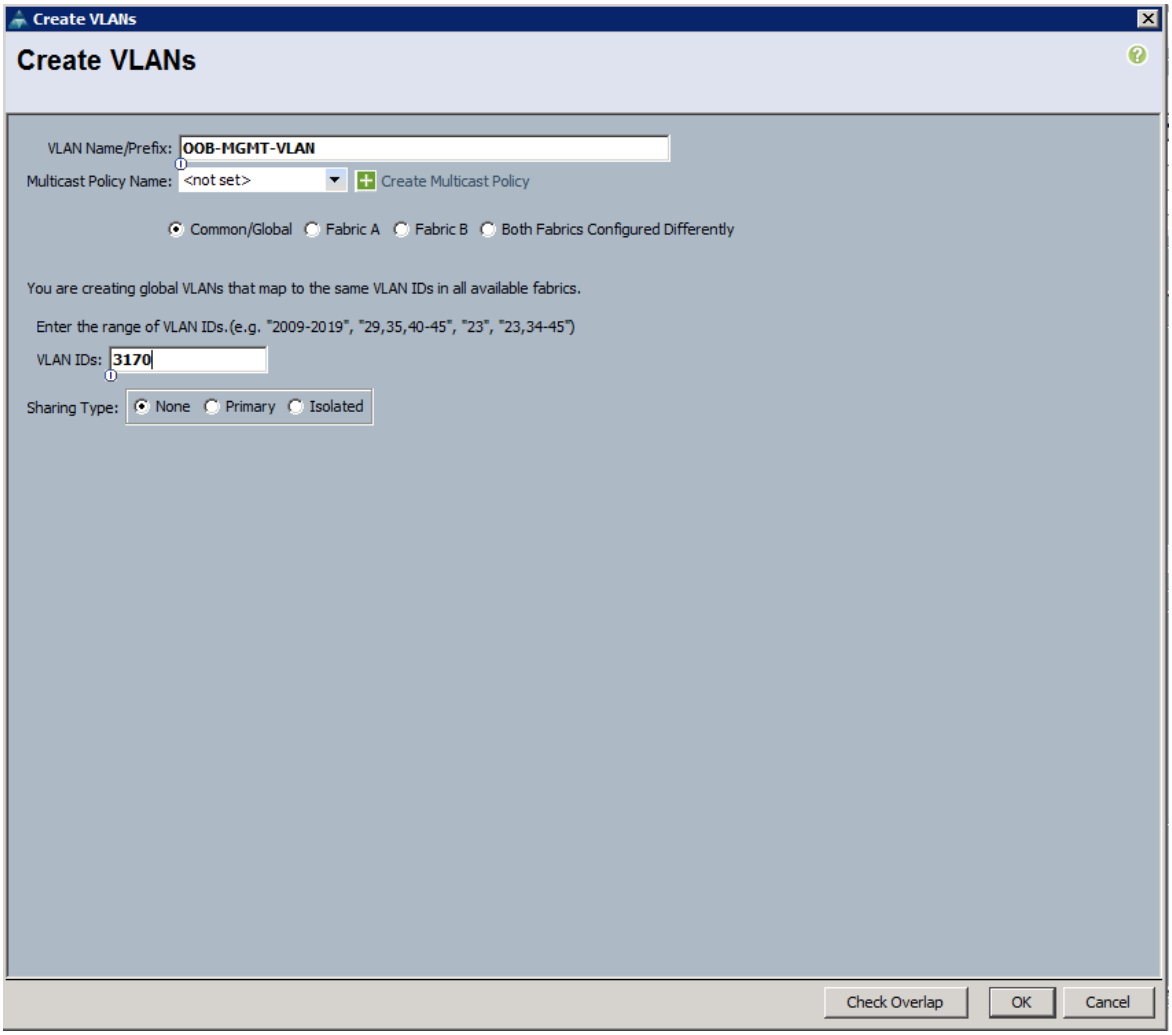

Create vNIC Templates on the Cisco UCS

Cisco UCS Manager

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Choose **Policies** > **root**.
3. Right-click vNIC Templates.
4. Choose **Create vNIC Template**.
5. Enter **vNIC_OOB_MGMT_A** as the vNIC template name.
6. Click the **Fabric A** radio button.
7. Keep the **Enable Failover** checkbox unckecked.

8. Under Target, make sure that the **VM** checkbox is unchecked.
9. Click the **Updating Template** radio button for the Template Type.
10. Under VLANs, check the checkbox for **OOB-MGMT-VLAN**.
11. Keep the MTU at 1500.
12. For MAC Pool , choose **MAC_Pool_A** from the drop-down list.
13. For Network Control Policy, choose **Enable_CDP** from the drop-down list.
14. Click **OK** to create the vNIC template.
15. Click **OK**.

Figure 140 Creating vNIC Template for Fabric A

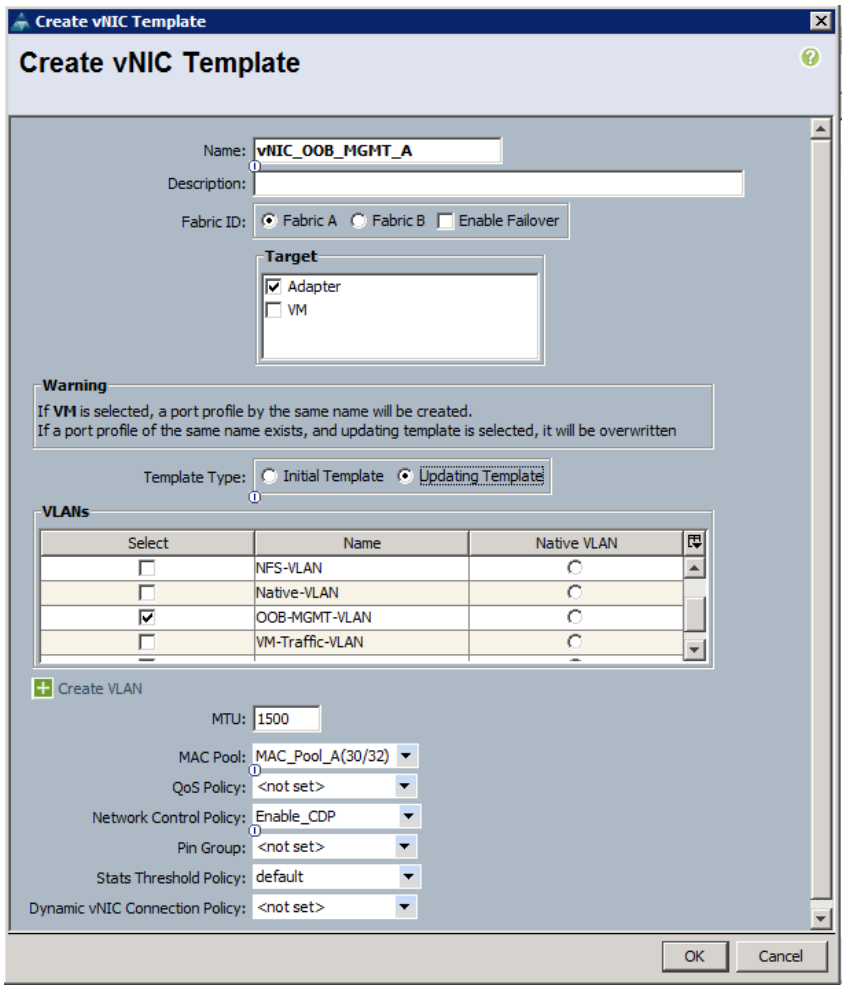

16. In the navigation pane, click the **LAN** tab.
17. Choose **Policies** > **root**.
18. Right-click vNIC Templates.
19. Choose **Create vNIC Template**.
20. Enter **vNIC_OOB_MGMT_B** as the vNIC template name.

21. Click the **Fabric B** radio button.
22. Uncheck the **Enable Failover** checkbox.
23. Under Target, make sure that the **VM** checkbox is unchecked.
24. Choose the **Updating Template** radio button for the Template Type.
25. Under VLANs, check the checkbox for OOB-MGMT-VLAN.
26. Keep the MTU at 1500.
27. For MAC Pool, choose **MAC_Pool_B** from the drop-down list.
28. For Network Control Policy, choose **Enable_CDP** from the drop-down list.
29. Click **OK** to create the vNIC template.
30. Click **OK**.

Figure 141 Creating vNIC Template for Fabric B

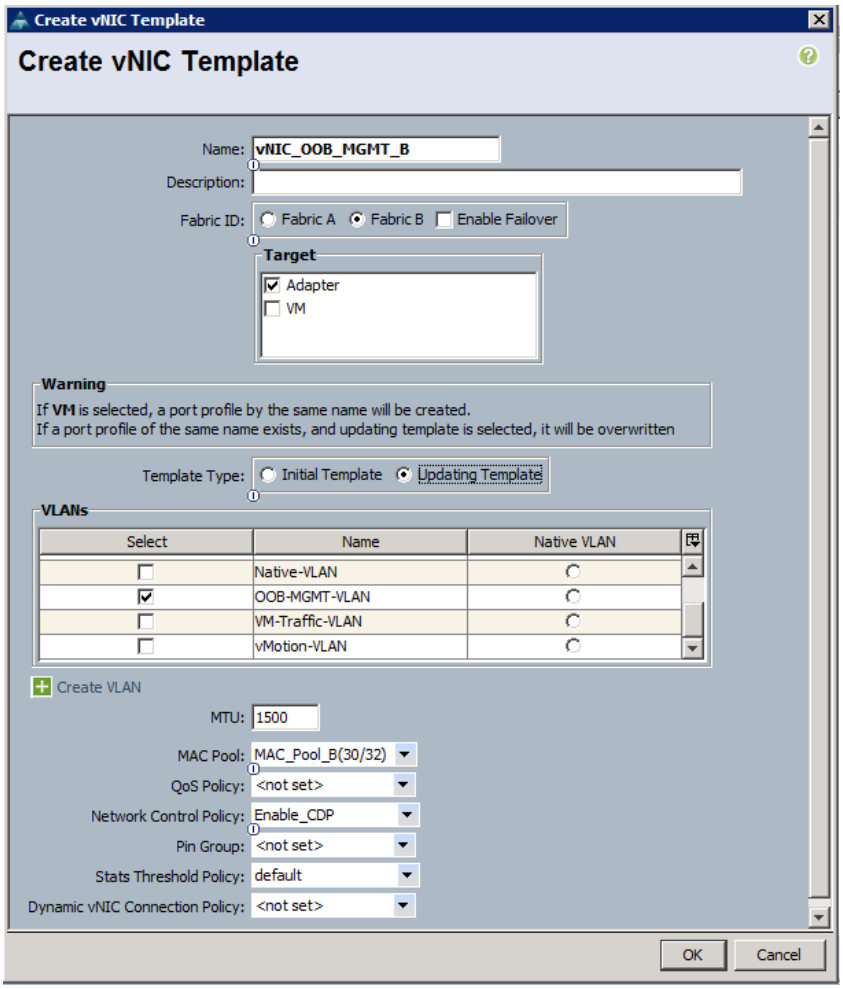

Configure ESXi Hosts

The procedures in the following two subsections will have to be repeated for each ESXi host on which Flash Accel and Fusion-io need to be configured.

Shut Down the ESXi Host

1. Log in to VMware vCenter Server using vSphere Client as FlexPod admin user.
2. Choose **Hosts and Clusters** from the Inventory pane.
3. Right-click the ESXi server on which you want to install Flash Accel and choose the option **Enter Maintenance Mode**.

Note This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

4. Right-click the ESXi server in the Maintenance Mode and choose **Shut Down**.
5. Click **OK**.

Add Out-of-Band Management vNICs to ESXi Host's Service Profile

This procedure will add Cisco Virtual Network Interfaces to the ESXi hosts and will require a server reboot. In this procedure, the service profile is unbounded from its template and modified. Alternatively, the Service Profile Template could be modified, but that would require all servers bound to the template to be rebooted.

Note If the Service Profile Template is modified, all servers bound to the template would inherit the configuration changes, whether or not Flash Accel is installed on them.

1. Using a web browser, log in to the Cisco UCS Manager as admin.
2. In the Cisco UCS Manager, click the **Servers** tab in the navigation pane.
3. Under **Servers** > **Service Profiles** > **root**, select the Service Profile for the server that was shut down.
4. Under the Actions pane, choose **Unbind from the Template**.
5. Click **Yes** to unbind the Service Profile from the template.
6. Click **OK**.
7. Choose the **Network** tab on the right pane.
8. At the bottom of the screen, click **Add**.
9. In the Create vNIC dialog box, enter **vNIC-OOB-MGMT-A** as the name of the vNIC.
10. Check the **Use vNIC Template** checkbox.
11. For vNIC Template, choose **vNIC_OOB_MGMT_A** from the drop-down list.
12. For Adapter Policy, choose **VMware** from the drop-down list.
13. Click **OK** to add the vNIC to the Service Profile.
14. At the bottom of the screen, click **Add**.

Figure 142 Creating a vNIC Using vNIC Template

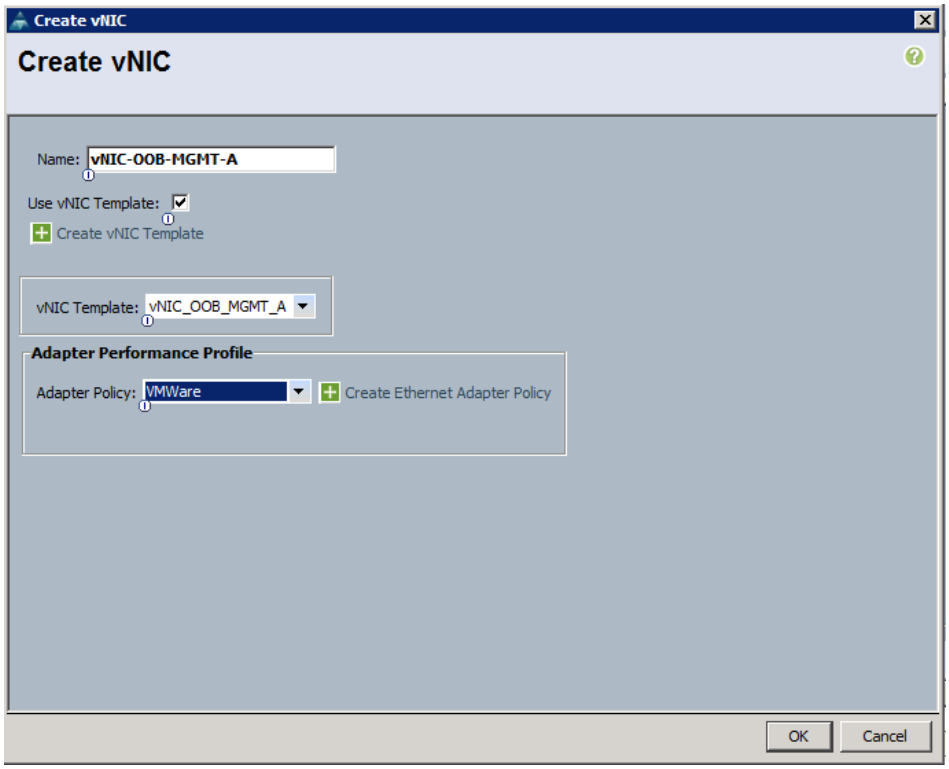

15. In the Create vNIC dialog box, enter **vNIC-OOB-MGMT-B** as the name of the vNIC.
16. Check the **Use vNIC Template** checkbox.
17. For vNIC Template, choose **vNIC_OOB_MGMT_B** from the drop-down list.
18. For Adapter Policy, choose VMware from the drop-down list.
19. Click **OK** to add the vNIC to the Service Profile.

Figure 143 Creating a vNIC Using vNIC Template

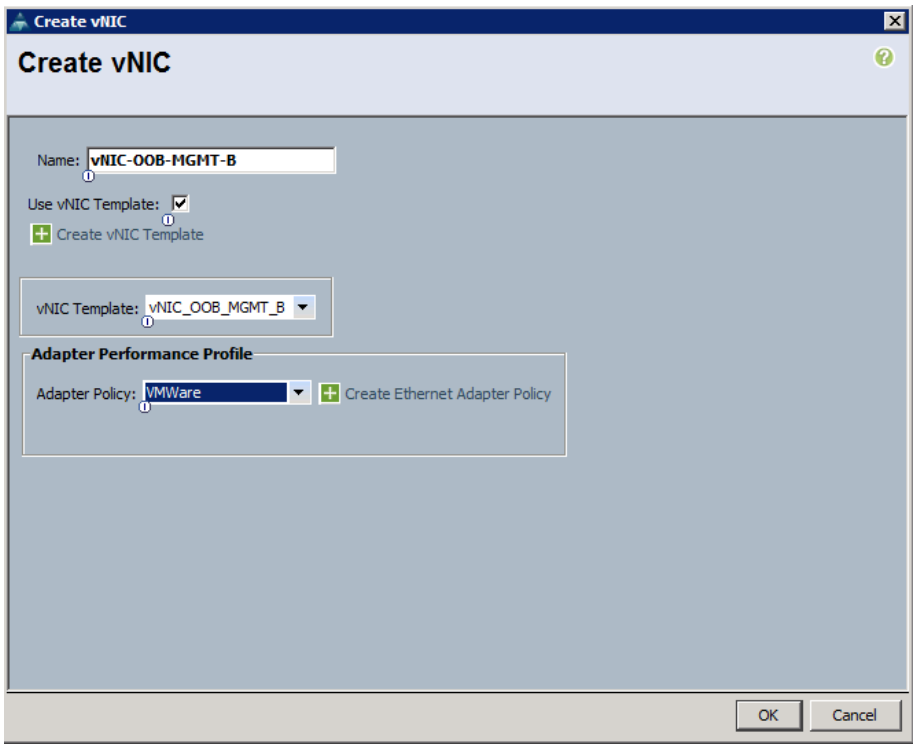

20. In the lower right-hand corner of the screen, click **Save Changes**.
21. Click **Yes**.
22. Click **OK**.
23. Under Actions, click **Modify vNIC/HBA Placement**.
24. Make sure that **VM-Host-Infra** is selected as the placement policy.
25. Select vCon1 and assign vNIC-OOB-MGMT-A and vNIC-OOB-MGMT-B to the virtual network interfaces policy and place them below vNIC-B.

Figure 144 Modifying vNIC/vHBA Placement

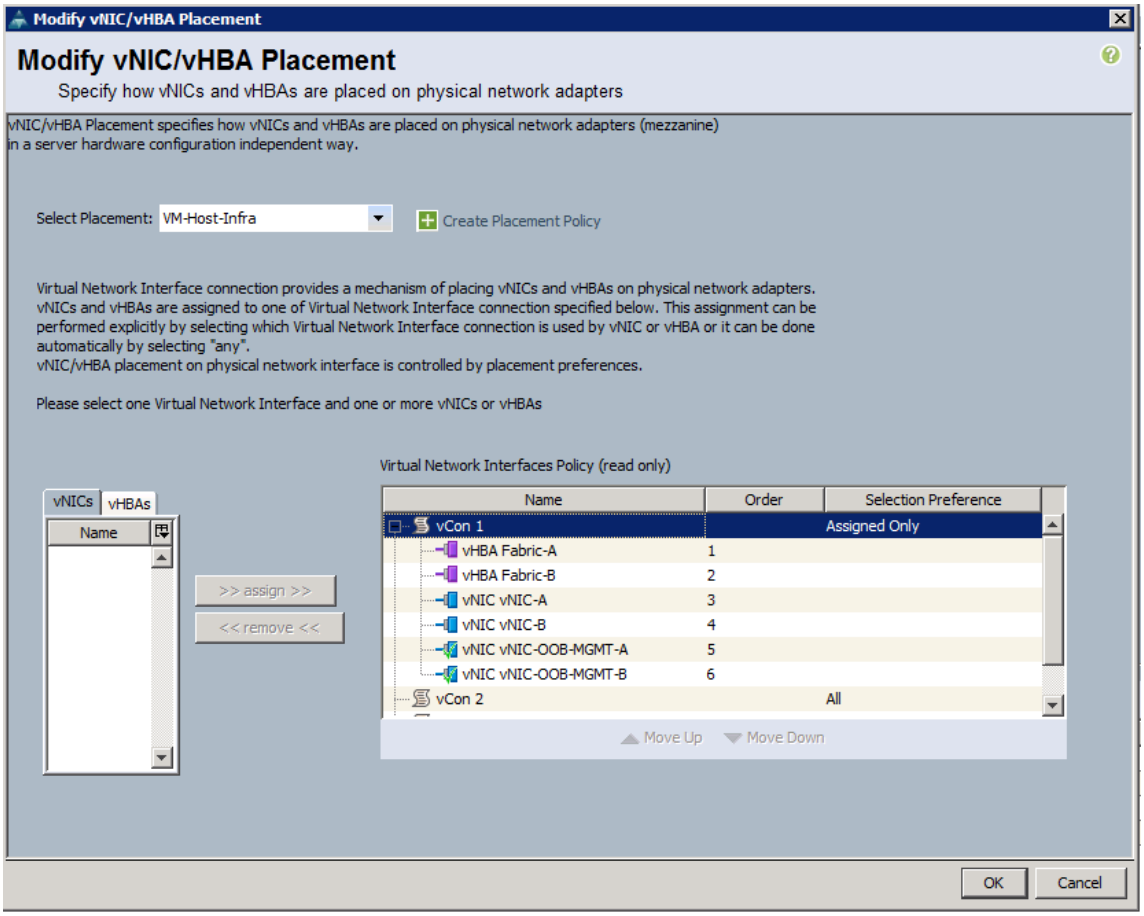

26. Click **OK** and click **Yes**.
27. Click **OK**.
28. At the top of the screen, choose Pending Activities.
29. Check the **Reboot Now** checkbox and click **OK**.
30. Check the **General** tab on the right pane.
31. In the Actions pane, select KVM Console to monitor the server until VMware ESXi has rebooted.
32. Repeat the preceding two sections for each ESXi host on which Flash Accel and Fusion-io needs to be configured.

Add Management Port Group to ESXi Host

The procedures in the following subsection will have to be repeated for each ESXi host in the VMware Cluster irrespective of whether or not Flash Accel and Fusion-io need to be configured on them.

1. Log in to VMware vCenter Server using vSphere Client as FlexPod admin user.
2. Select Hosts and Clusters from the Inventory pane.
3. Right-click on the ESXi server on which you want to install Flash Accel and select Enter Maintenance Mode.

Note This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

4. With the ESXi Host selected, choose the **Configuration** tab.
5. In the Hardware Pane, choose **Networking**.
6. On the right pane, click **Add Networking**.
7. Choose VMkernel, click **Next**.
8. Make sure **vmnic2** Physical Adapter is selected and click **Next**.
9. Change the Network Label to Management Network.

Note This is a mandatory requirement.

10. Enter the **<<var_oob-mgmt_vlan_id>>** for the VLAN ID.
11. Check the **Use this port group for management traffic** check box and click **Next**.

Figure 145 VMKernel Connection Settings

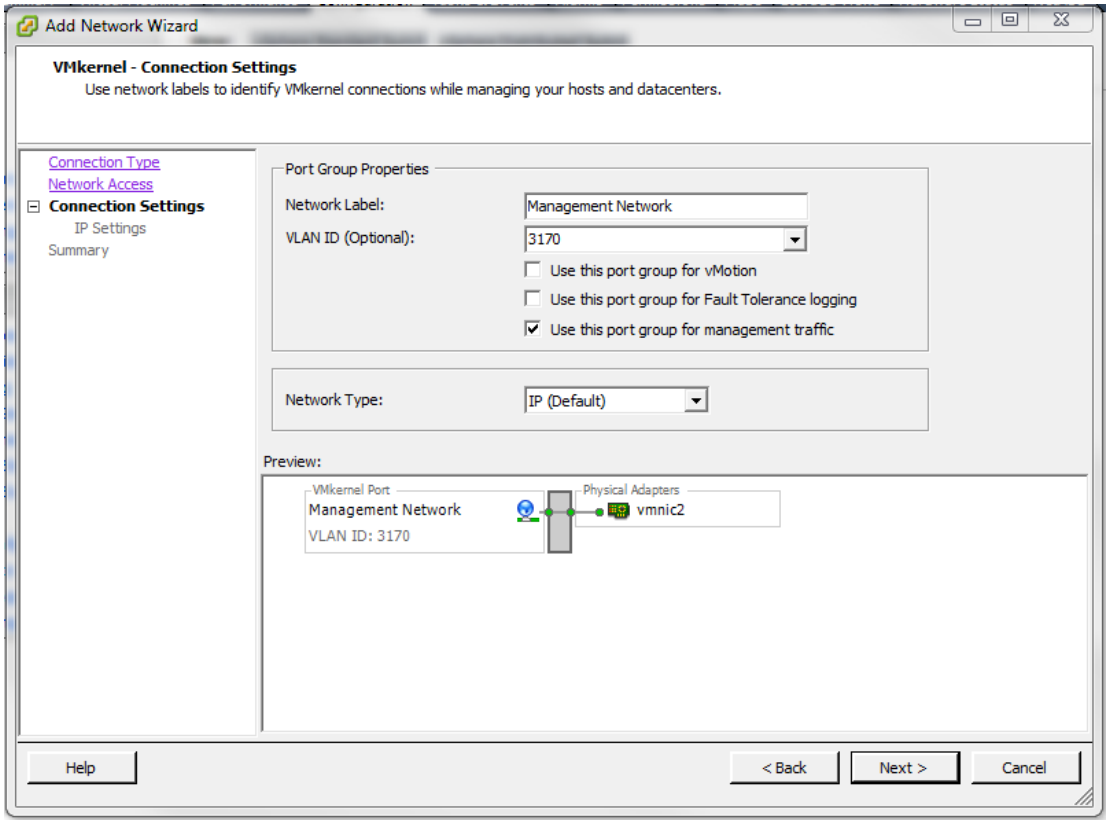

12. Enter the appropriate IP address and subnet mask and click **Next**.
13. Click **Finish** to create vSwitch0.
14. Choose Properties for vSwitch0.

15. Choose the **Network Adapters** tab.
16. Click **Add**.
17. Check the checkbox for **vmnic3** and click **Next**.
18. Click **Next**.
19. Click **Finish**.
20. Choose the **Ports** tab.
21. With vSwitch selected, click **Edit**.
22. Choose the **NIC Teaming** tab.
23. Under Load Balancing, choose Route based on source MAC hash.
24. Click **OK**.
25. Click **Close** to close the vSwitch0 properties.

Figure 146 vSwitch0 Properties

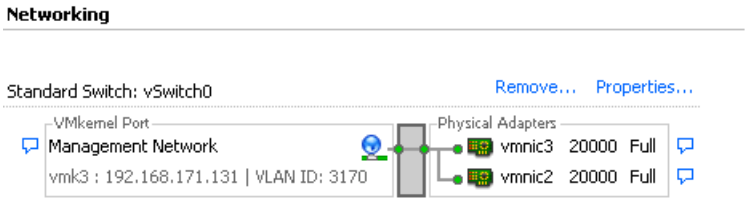

26. In the Software pane, choose **DNS and Routing**.
27. Choose Properties on the right corner.
28. Choose the **Routing** tab.
29. Change the Default gateway to the gateway for the OOB-MGMT-VLAN.
30. Click **OK**.
31. If host name resolution is being used for identifying the ESXi Hosts in vCenter, the DNS server host record for this ESXi host should be changed to use the OOB-MGMT-VLAN IP address.
32. Right-click the ESXi host and click **Exit Maintenance Mode**.

Download Flash Accel Software

1. From the management workstation, navigate to the NetApp Support site.
2. Choose **Flash Accel** and choose the **VMware – Virtual Storage Console** option.
3. Download the 1.2R1 version of Flash Accel.

Download Fusion-io Driver

1. From the management workstation, navigate to the Fusion-io driver download section in the VMware support site.

Note This procedure requires valid login credentials and Internet access is required on the management workstation.

2. Click **Download**, log in, and click **Download Now** to download the driver.
3. Extract the offline bundle scsi-iomemory-vsl-5X-3.2.4.1108-offline_bundle-1176140.zip from the downloaded driver file.

Install Flash Accel

Before proceeding with the following procedure, NetApp highly recommends verifying that the servers, flash devices, VMware software, and virtual machines meet the minimum resource requirements. These prerequisites can be found in the NetApp Flash Accel for VMware with Virtual Storage Console document.

The Flash Accel Base package will need to be installed on the host on which VMware Virtual Storage Console (VSC) is running.

1. Launch the downloaded Flash Accel binary.
2. Click **Next** in the Welcome screen.
3. Click **Next** to accept the default installation folder.
4. Click **Install**.
5. Accept to restart the VSC service, click **Finish**.

Figure 147 NetApp Flash Accel Base for VSC for VMware vSphere Setup Wizard

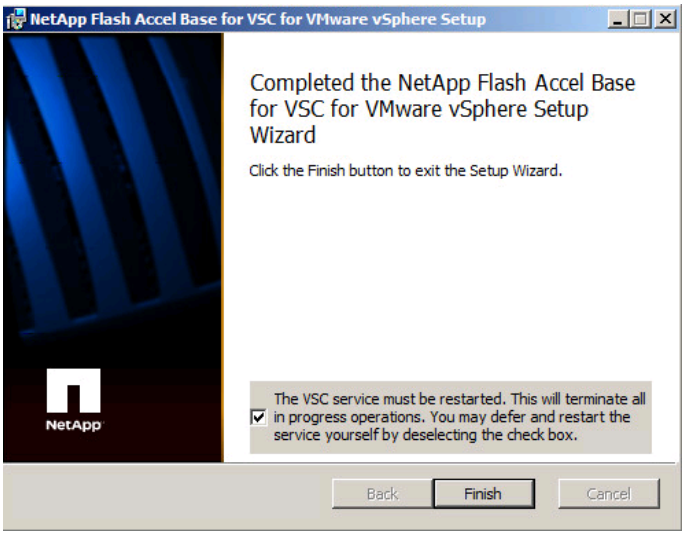

The Flash Accel base package is installed on the VSC host. The VSC service is restarted automatically.

Install Fusion-io Driver on the ESXi Servers

1. Log in to the vSphere Client.
2. Choose **Datastores** and **Datastore Clusters** in the Inventory pane.

3. Choose the **infra_datastore_1** from the drop-down list of datastores.
4. Right-click infra_datastore_1 and choose **Browse Datastore**.
5. Click the third button at the top to create a new folder.
6. Name the folder Fusion-io and click **OK**.
7. On the left, select the Fusion-io folder.
8. Click the fourth button and choose **Upload File**.
9. Navigate to the scsi-iomemory-vsl-5X-3.2.4.1108-offline_bundle-1176140.zip file and click **Open**.
10. The file should now appear in the Fusion-io folder in the datastore.
11. Open the VMware vSphere CLI command prompt.
12. For each ESXi host in the VMware vSphere CLI, run the following command:

```
esxcli -s <Host Server ip> -u root -p <Root Password> software vib install -d
/vmfs/volumes/infra_datastore_1/Fusion-io/scsi-iomemory-vsl-5X-3.2.4.1108-offline_
bundle-1176140.zip
```

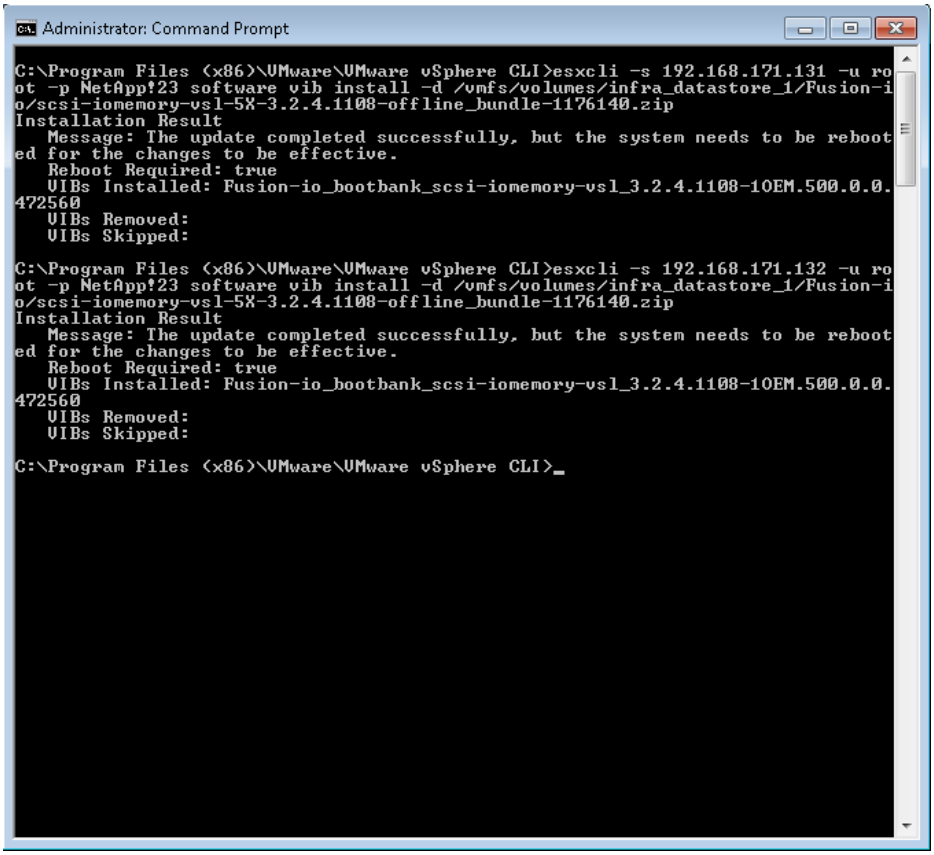

13. From the vSphere Client, click **Hosts and Clusters** from the Inventory pane.
14. Right-click an ESXi server and choose **Enter Maintenance Mode**.

Note This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

15. Right-click the ESXi server in Maintenance Mode and choose **Reboot**.

16. After reboot is completed, right-click the ESXi server and choose **Exit Maintenance Mode**.

Repeat steps 14 through 16 for the other ESXi servers.

Install Host Agent

1. From the vSphere Client, select Inventory > Hosts and Clusters.
2. Right-click an ESXi server and select Enter Maintenance Mode.

Note This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

3. Navigate to **Home** > **NetApp** (under Solutions and Applications).
4. Under the Monitoring and Host Configuration section, click **Flash Accel Cache**.

Figure 148 Flash Accel Cache

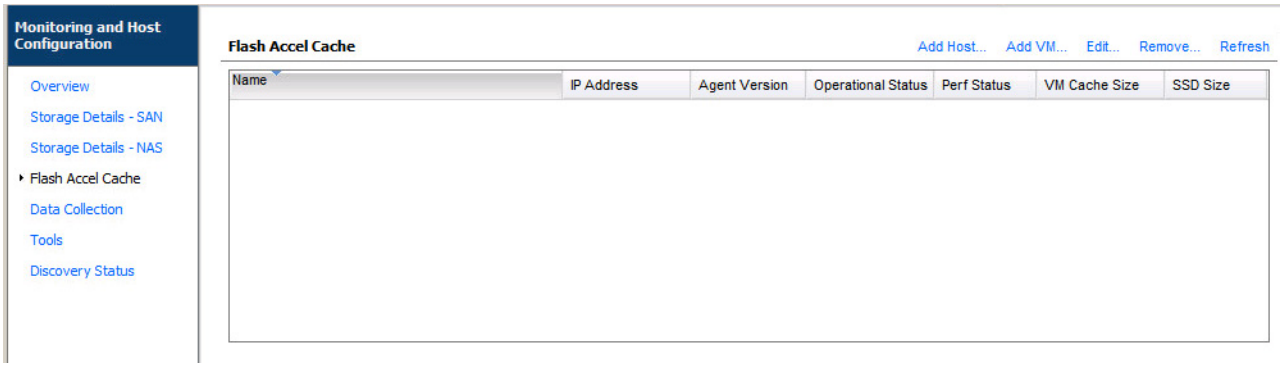

5. Click **Add Host** in the Flash Accel Cache window.
6. Select the ESXi server that is in Maintenance Mode and its underlying Fusion-io card. Click **Install**.

Figure 149 Selecting Hosts

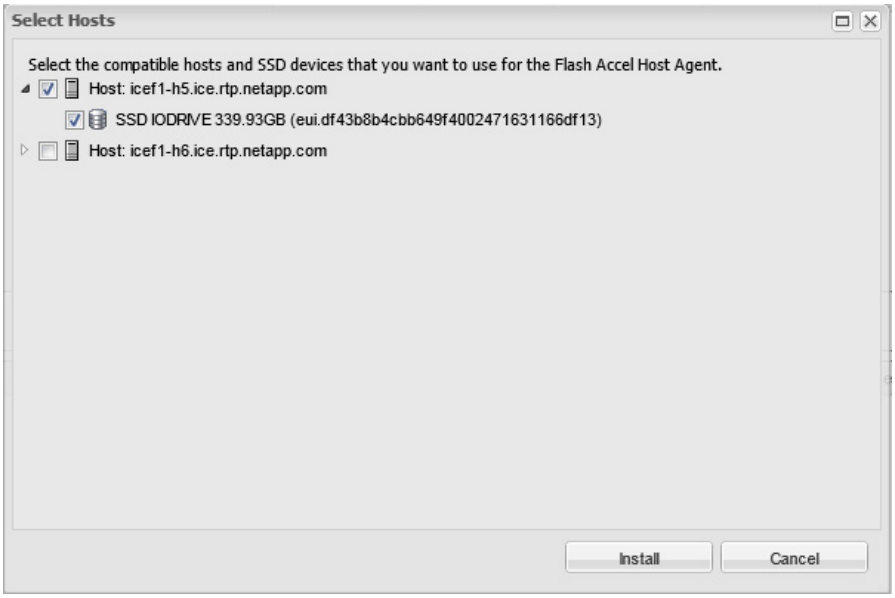

7. Click **Yes** in the confirmation window to install the Host Agent.
8. Upon successful installation, the ESXi server will be listed with the configured Flash Accel Cache details.

Figure 150 Configured Flash Accel Cache

9. In the vSphere Client, navigate to **Home > Inventory > Hosts and Clusters**.
10. Right-click the ESXi server and choose **Exit Maintenance Mode**.
11. Repeat steps 5 through 10 for the remaining ESXi servers.

Figure 151 Configured Flash Accel Cache Showing Added ESXi Servers

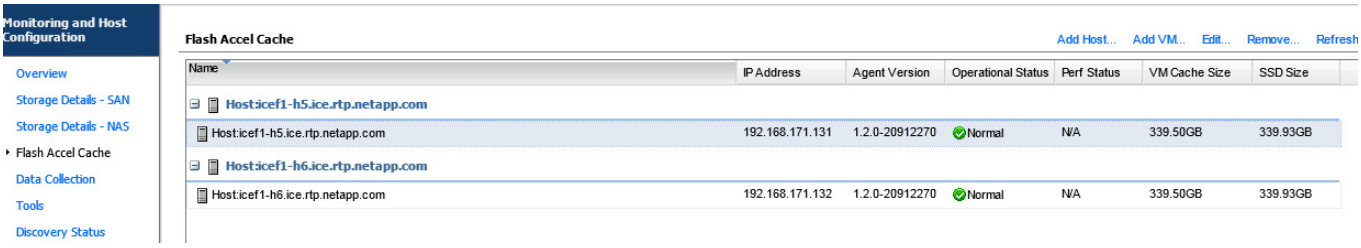

Install Guest OS Agent and Allocate Cache

To install the guest OS agent and allocate cache on VMs, complete the following steps.

1. From the Flash Accel Cache window, select the ESXi server on which the VM is running.
2. Click **Add VM** in the Flash Accel Cache window.

Figure 152 Adding VM

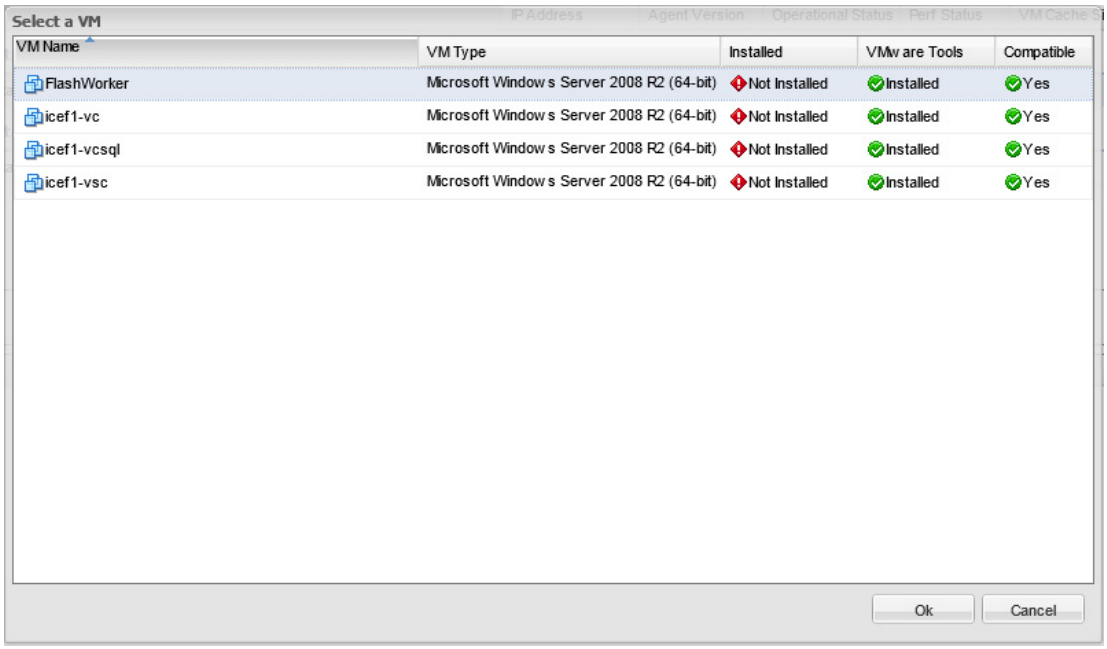

Note All the VMs that would use Flash Accel Cache should have the up-to-date versions of the VMware tools.

3. Select a compatible VM on which you want to install the guest OS agent. Click **OK**.
4. In the Agent Connection Information section, accept the default values for IP address and management port. Provide the administrator credentials and click the **Verify** to validate.
5. In the Cache Space Settings section, specify the cache size for the VM and the select the datastore where the Flash Accel mapping file should be saved. Check the **Enable VMotion** checkbox and make sure to select a shared datastore to which each of the potential target hosts has access. Click **OK**.

Figure 153 Adding Flash Accel Guest OS Agent

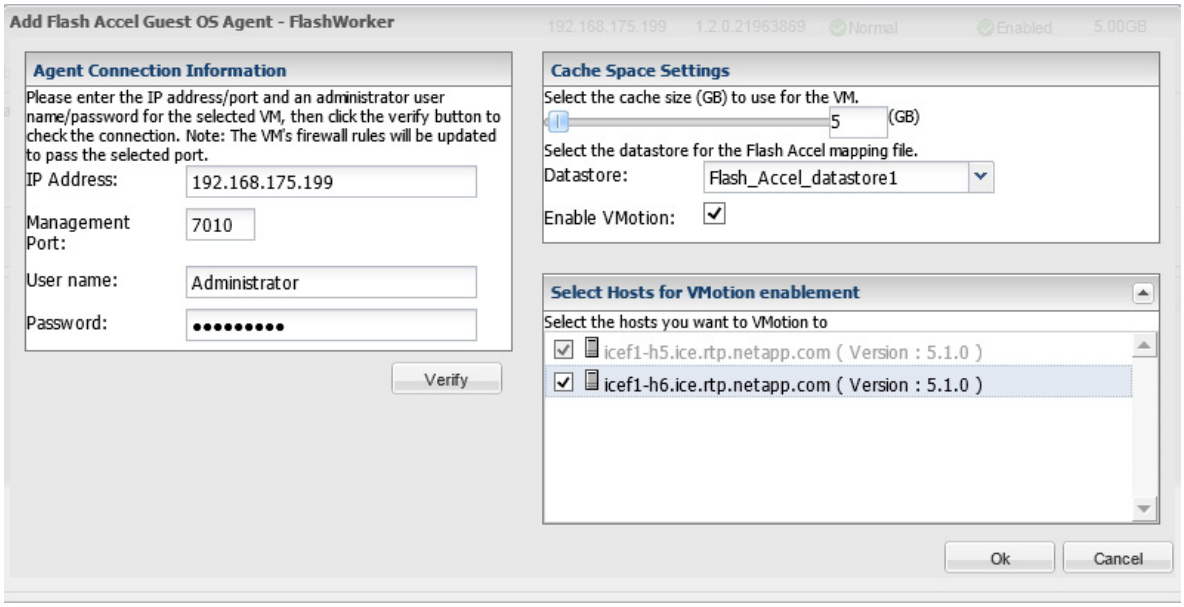

6. Click **Yes** to confirm the installation and reboot the VM.

Figure 154 Flash Accel Cache Showing Added VM

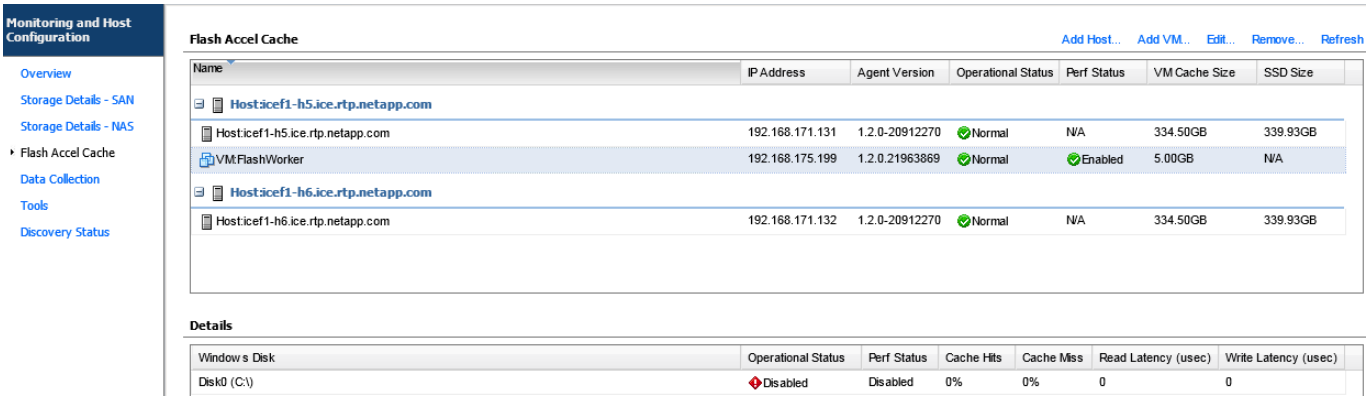

7. Select the VM on which the Guest OS Agent was installed and cache was allocated. Click **Edit** in the Flash Accel Cache Window.
8. Under the VM Cached Devices section, right-click Windows Disk and choose **Enabled** to activate caching.

Figure 155 Editing Flash Accel Guest OS Agent

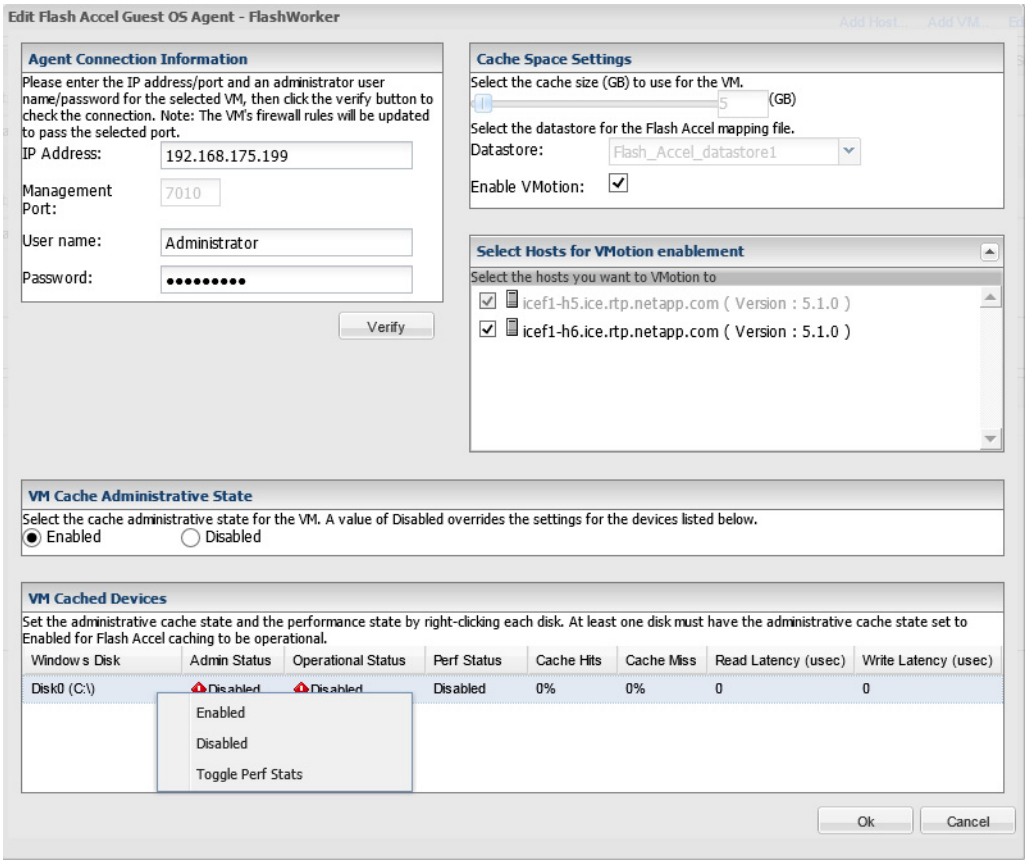

9. (Optional) Right-click the same Windows Disk and choose **Toggle Perf Stats** to collect performance statistics.

Figure 156 Flash Accel Cache Details

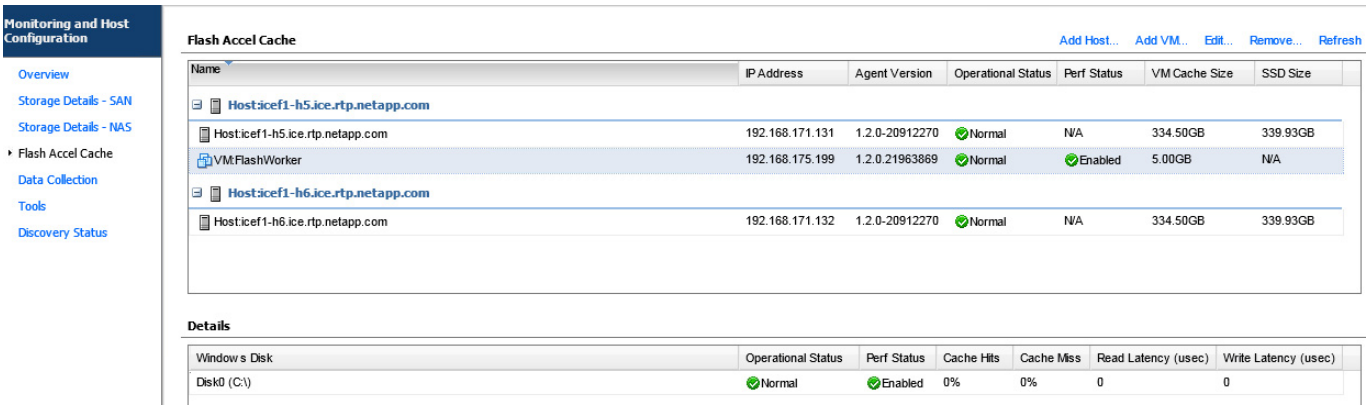

10. Click **OK**.

Cisco UCS Central—Multi Domain Management

Cisco UCS Central software manages multiple, globally distributed Cisco UCS domains with thousands of servers from a single pane.

This section provides a detailed overview of UCS Central setup in a standalone mode.

The installation and upgrade guide is available at http://www.cisco.com/en/US/docs/unified_computing/ucs/ucs-central/install-upgrade/1.1/b_UCSC_Installation_and_Upgrade_Guide_11.html

Obtain the UCS Central Software

1. Navigate to the Cisco UCS Central Download page.
2. Download the OVA file ucs-central.1.1.1a.ova.

Install the UCS Central Software

1. Using the vSphere Client, log in to the vCenter Server as FlexPod admin user.
2. In the Main menu, click **File > Deploy OVF Template**.
3. Browse to the OVA file that was downloaded. Click **Next**.
4. Click **Next**.
5. Modify the default name if desired and select the Inventory Location. Click Next.
6. Select a cluster/server on which you want to host the UCS Central virtual machine.
7. Select the datastore in which the virtual machine files will be stored. Click **Next**.
8. Click **Next**.
9. Check the checkbox to power on the VM after deployment.
10. Click **Finish**.

Note Do not proceed until the virtual machine has finished booting.

11. Open a console window to the UCS Central virtual machine.
12. Answer the following questions in the console window:

```
Setup new configuration or restore full-state configuration from
backup[setup/restore] – setup
Enter the UCS Central VM eth0 IPv4 Address : <<var_ucs_central_ip>>
Enter the UCS Central VM eth0 IPv4 Netmask : <<var_ucs_central_netmask>>
Enter the VM IPv4 Default Gateway : <<var_ucs_central_gateway>>
Is this VM part of a cluster (select 'no' for standalone) (yes/no)? no
Enter the UCS Central VM Hostname : <<var_ucs_central_hostname>>
Enter the DNS Server IPv4 Address : <<var_nameserver_ip>>
Enter the Default Domain Name : <<var_dns_domain_name>>
Use a Shared Storage Device for Database (yes/no)? no
Enforce Strong Password (yes/no)? yes
Enter the admin Password : <<var_password>>
Confirm admin Password : <<var_password>>
Enter the Shared Secret :  <<enter the shared secret (or password) that you want
to use to register one or more Cisco UCS domains with Cisco UCS Central>>
Confirm Shared Secret : <<re-enter the Shared Secret>>
```

```
Do you want Statistics collection [yes / no]? yes
Enter the Statistics DB Type [D=Default (internal Pstgres db) / P=Postgres /
O=Oracle] : D
Proceed with this configuration? Please confirm [yes/no] - yes
```

Note If you wish to modify/answer the questions again, enter no in the above prompt.

After you confirm that you want to proceed with the configuration, the network interface reinitializes with your settings and Cisco UCS Central becomes accessible via the IP address.

Accessing UCS Central GUI

1. Using a Web browser, navigate to the <<var_ucs_central_hostname>> using https://<<var_ucs_central_ip>>.
2. Log in with the user name as admin and the admin password.
3. Click the Operations Management tab, expand **Domain Groups** > **Domain Group root**.
4. Choose **Operational Policies**.
5. Select Time Zone in the right pane, choose the desired Time Zone.
6. Click **Add NTP Server.**
7. Provide the NTP Server IP Address, click **OK**.
8. Click **Save**.

Adding UCS Managers to UCS Central

UCS Managers are be added into the UCS Central by logging into the UCSM and registering the UCSM with UCS Central.

Follow these steps to add an UCS Manager to UCS Central:

1. Login to the Cisco UCS Manager.
2. In the navigation pane, click the **Admin** tab.
3. In the Admin tab expand the All folder, choose **Communication Management** > **UCS Central**.
4. In the **UCS Central** tab, click **Register With UCS Central** in the **Actions** pane.
5. Enter the host name or IP address of the UCS Central.
6. Enter the Shared Secret. Click **OK**.
7. Click **Accept** to terminate any open GUI sessions to the UCS Manager.
8. Select the checkbox to view the Navigator for the UCS Central. Click **OK**.
9. Verify the Registration Status.

Cisco Nexus 5548UP Example Configurations

Cisco Nexus 5548UP A

```
!Command: show running-config
!Time: Fri Sep 13 16:45:33 2013
```

```
version 6.0(2)N1(2a)
feature fcoe
switchname NX5548UP-A
feature npiv
no feature telnet
cfs eth distribute
feature lacp
feature vpc
feature lldp
username admin password 5 $1$kqCzt6lW$HzN4Xo02Z8.utBe9urANk1  role network-admin

banner motd #Nexus 5000 Switch
#

ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-default
    mtu 9216
    multicast-optimize
system qos
  service-policy type qos input fcoe-default-in-policy
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type network-qos jumbo
snmp-server user admin network-admin auth md5 0xd5d13e9b8ffb25541b2abe28b229104e priv
0xd5d13e9b8ffb25541b2abe28b229104e localizedkey
ntp server 10.72.242.70 use-vrf management
vrf context management
  ip route 0.0.0.0/0 10.72.242.1
vlan 1
vlan 2
  name Native-VLAN
vlan 101
  fcoe vsan 101
  name FCoE_Fabric_A
vlan 3170
  name OOB-MGMT-VLAN
vlan 3172
  name NFS-VLAN
vlan 3173
  name vMotion-VLAN
vlan 3174
  name VM-Traffic-VLAN
vlan 3175
  name IB-MGMT-VLAN
vlan 3176
  name Packet-Control-VLAN
```

```
spanning-tree port type edge bpduguard default
spanning-tree port type network default
port-channel load-balance ethernet source-dest-port
vpc domain 7
  role priority 10
  peer-keepalive destination 10.72.242.249 source 10.72.242.248
  auto-recovery
port-profile default max-ports 512
vsan database
  vsan 101 name "Fabric_A"
device-alias database
  device-alias name FAS3250-01_1a pwwn 50:0a:09:81:88:aa:7d:a0
  device-alias name FAS3250-02_1a pwwn 50:0a:09:81:88:ba:7d:ec
  device-alias name VM-Host-Infra-01_A pwwn 20:00:00:25:b5:00:0a:0f
  device-alias name VM-Host-Infra-02_A pwwn 20:00:00:25:b5:00:0a:1f

device-alias commit

fcdomain fcid database
  vsan 101 wwn 50:0a:09:81:88:aa:7d:a0 fcid 0x820001 dynamic
!              [FAS3250-01_1a]
  vsan 101 wwn 50:0a:09:81:88:ba:7d:ec fcid 0x820021 dynamic
!              [FAS3250-02_1a]
  vsan 101 wwn 22:c8:54:7f:ee:34:b8:bf fcid 0x820040 dynamic
  vsan 101 wwn 20:00:00:25:b5:00:0a:0f fcid 0x820041 dynamic
!              [VM-Host-Infra-01_A]
  vsan 101 wwn 20:00:00:25:b5:00:0a:1f fcid 0x820042 dynamic
!              [VM-Host-Infra-02_A]


interface port-channel1
  description "GLC-T connected to UPLINK INFRA"
  switchport mode trunk
  switchport trunk allowed vlan 3170,3175
  spanning-tree port type normal
  vpc 1

interface port-channel10
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FAS3250-01
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  spanning-tree port type edge trunk
  vpc 11

interface port-channel12
  description FAS3250-02
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  spanning-tree port type edge trunk
  vpc 12

interface port-channel13
  description FI6248-A
  switchport mode trunk
```

```
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  spanning-tree port type edge trunk
  vpc 13

interface port-channel14
  description FI6248-B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  spanning-tree port type edge trunk
  vpc 14

interface port-channel15
  description FI6248-A:FCoE
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101
  spanning-tree port type edge trunk

interface vfc11
  bind interface Ethernet1/1
  switchport trunk allowed vsan 101
  switchport description FAS3250-01:1a
  no shutdown

interface vfc12
  bind interface Ethernet1/2
  switchport trunk allowed vsan 101
  switchport description FAS3250-02:1a
  no shutdown

interface vfc15
  bind interface port-channel15
  switchport trunk allowed vsan 101
  switchport description FI6248-A:FCoE
  no shutdown
vsan database
  vsan 101 interface vfc11
  vsan 101 interface vfc12
  vsan 101 interface vfc15

interface Ethernet1/1
  description FAS3250-01:e1a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  channel-group 11 mode active

interface Ethernet1/2
  description FAS3250-02:e1a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  channel-group 12 mode active

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6
```

```
interface Ethernet1/7

interface Ethernet1/8
  switchport mode trunk
  switchport trunk allowed vlan 3170,3175
  speed 1000
  channel-group 1 mode active

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11
  description FI6248-A:1/19
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  channel-group 13 mode active

interface Ethernet1/12
  description FI6248-B:1/19
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  channel-group 14 mode active

interface Ethernet1/13
  description NX5548UP-B:1/13
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  channel-group 10 mode active

interface Ethernet1/14
  description NX5548UP-B:1/14
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  channel-group 10 mode active

interface Ethernet1/15
  description Nexus1110-X-A:Eth1
  switchport mode trunk
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge trunk
  speed 1000
  vpc orphan-port suspend

interface Ethernet1/16
  description Nexus1110-X-B:Eth1
  switchport mode trunk
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge trunk
  speed 1000
  vpc orphan-port suspend

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20
```

```
interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31
  description FI6248-A:1/31
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101
  channel-group 15 mode active

interface Ethernet1/32
  description FI6248-A:1/32
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101
  channel-group 15 mode active

interface mgmt0
  ip address 10.72.242.248/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.6.0.2.N1.2a.bin
boot system bootflash:/n5000-uk9.6.0.2.N1.2a.bin
!Full Zone Database Section for vsan 101
zone name VM-Host-Infra-01_A vsan 101
    member pwwn 20:00:00:25:b5:00:0a:0f
!               [VM-Host-Infra-01_A]
    member pwwn 50:0a:09:81:88:aa:7d:a0
!               [FAS3250-01_1a]
    member pwwn 50:0a:09:81:88:ba:7d:ec
!               [FAS3250-02_1a]

zone name VM-Host-Infra-02_A vsan 101
    member pwwn 20:00:00:25:b5:00:0a:1f
!               [VM-Host-Infra-02_A]
    member pwwn 50:0a:09:81:88:aa:7d:a0
!               [FAS3250-01_1a]
    member pwwn 50:0a:09:81:88:ba:7d:ec
!               [FAS3250-02_1a]

zoneset name FlexPod vsan 101
    member VM-Host-Infra-01_A
    member VM-Host-Infra-02_A

zoneset activate name FlexPod vsan 101
```

Cisco Nexus 5548UP B

```
!Command: show running-config
!Time: Fri Sep 13 17:10:52 2013

version 6.0(2)N1(2a)
feature fcoe
switchname NX5548UP-B
feature npiv
no feature telnet
cfs eth distribute
feature lacp
feature vpc
feature lldp
username admin password 5 $1$zIJsN.tA$UvbASyjxIARCstojspT93/  role network-admin

banner motd #Nexus 5000 Switch
#

ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-default
    mtu 9216
    multicast-optimize
system qos
  service-policy type qos input fcoe-default-in-policy
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type network-qos jumbo
snmp-server user admin network-admin auth md5 0xa241b7674af438b7cab3af0f7c77d2e4 priv
0xa241b7674af438b7cab3af0f7c77d2e4 localizedkey
ntp server 10.72.242.70 use-vrf management
vrf context management
  ip route 0.0.0.0/0 10.72.242.1
vlan 1
vlan 2
  name Native-VLAN
vlan 102
  fcoe vsan 102
  name FCoE_Fabric_B
vlan 3170
  name OOB-MGMT-VLAN
vlan 3172
  name NFS-VLAN
vlan 3173
  name vMotion-VLAN
vlan 3174
```

```
  name VM-Traffic-VLAN
vlan 3175
  name IB-MGMT-VLAN
vlan 3176
  name Packet-Control-VLAN
spanning-tree port type edge bpduguard default
spanning-tree port type network default
port-channel load-balance ethernet source-dest-port
vpc domain 7
  role priority 20
  peer-keepalive destination 10.72.242.248 source 10.72.242.249
  auto-recovery
port-profile default max-ports 512
vsan database
  vsan 102 name "Fabric_B"
device-alias database
  device-alias name FAS3250-01_2a pwwn 50:0a:09:83:88:aa:7d:a0
  device-alias name FAS3250-02_2a pwwn 50:0a:09:83:88:ba:7d:ec
  device-alias name VM-Host-Infra-01_B pwwn 20:00:00:25:b5:00:0b:0f
  device-alias name VM-Host-Infra-02_B pwwn 20:00:00:25:b5:00:0b:1f

device-alias commit

fcdomain fcid database
  vsan 102 wwn 50:0a:09:83:88:aa:7d:a0 fcid 0xbd0001 dynamic
!              [FAS3250-01_2a]
  vsan 102 wwn 50:0a:09:83:88:ba:7d:ec fcid 0xbd0021 dynamic
!              [FAS3250-02_2a]
  vsan 102 wwn 22:c8:54:7f:ee:34:a8:bf fcid 0xbd0040 dynamic
  vsan 102 wwn 20:00:00:25:b5:00:0b:0f fcid 0xbd0041 dynamic
!              [VM-Host-Infra-01_B]
  vsan 102 wwn 20:00:00:25:b5:00:0b:1f fcid 0xbd0042 dynamic
!              [VM-Host-Infra-02_B]


interface port-channel1
  description "GLC-T connected to UPLINK INFRA"
  switchport mode trunk
  switchport trunk allowed vlan 3170,3175
  spanning-tree port type normal
  vpc 1

interface port-channel10
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FAS3250-01
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102,3172
  spanning-tree port type edge trunk
  vpc 11

interface port-channel12
  description FAS3250-02
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102,3172
  spanning-tree port type edge trunk
```

```
  vpc 12

interface port-channel13
  description FI6248-A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  spanning-tree port type edge trunk
  vpc 13

interface port-channel14
  description FI6248-B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  spanning-tree port type edge trunk
  vpc 14

interface port-channel15
  description FI6248-B:FCoE
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102
  spanning-tree port type edge trunk

interface vfc11
  bind interface Ethernet1/1
  switchport trunk allowed vsan 102
  switchport description FAS3250-01:2a
  no shutdown

interface vfc12
  bind interface Ethernet1/2
  switchport trunk allowed vsan 102
  switchport description FAS3250-02:2a
  no shutdown

interface vfc15
  bind interface port-channel15
  switchport trunk allowed vsan 102
  switchport description NX5548UP-B:FCoE
  no shutdown
vsan database
  vsan 102 interface vfc11
  vsan 102 interface vfc12
  vsan 102 interface vfc15

interface Ethernet1/1
  description FAS3250-01:e1a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102,3172
  channel-group 11 mode active

interface Ethernet1/2
  description FAS3250-02:e1a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102,3172
  channel-group 12 mode active

interface Ethernet1/3

interface Ethernet1/4
```

```
interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8
  switchport mode trunk
  switchport trunk allowed vlan 3170,3175
  speed 1000
  channel-group 1 mode active

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11
  description FI6248-A:1/20
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  channel-group 13 mode active

interface Ethernet1/12
  description FI6248-B:1/20
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  channel-group 14 mode active

interface Ethernet1/13
  description NX5548UP-A:1/13
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  channel-group 10 mode active

interface Ethernet1/14
  description NX5548UP-A:1/14
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  channel-group 10 mode active

interface Ethernet1/15
  description Nexus1110-X-A:Eth2
  switchport mode trunk
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge trunk
  speed 1000
  vpc orphan-port suspend

interface Ethernet1/16
  description Nexus1110-X-B:Eth2
  switchport mode trunk
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge trunk
  speed 1000
  vpc orphan-port suspend

interface Ethernet1/17

interface Ethernet1/18
```

```
interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31
  description FI6248-B:1/31
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102
  channel-group 15 mode active

interface Ethernet1/32
  description FI6248-B:1/32
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102
  channel-group 15 mode active

interface mgmt0
  ip address 10.72.242.249/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.6.0.2.N1.2a.bin
boot system bootflash:/n5000-uk9.6.0.2.N1.2a.bin
!Full Zone Database Section for vsan 102
zone name VM-Host-Infra-01_B vsan 102
    member pwwn 20:00:00:25:b5:00:0b:0f
!               [VM-Host-Infra-01_B]
    member pwwn 50:0a:09:83:88:aa:7d:a0
!               [FAS3250-01_2a]
    member pwwn 50:0a:09:83:88:ba:7d:ec
!               [FAS3250-02_2a]

zone name VM-Host-Infra-02_B vsan 102
    member pwwn 50:0a:09:83:88:aa:7d:a0
!               [FAS3250-01_2a]
    member pwwn 50:0a:09:83:88:ba:7d:ec
!               [FAS3250-02_2a]
    member pwwn 20:00:00:25:b5:00:0b:1f
!               [VM-Host-Infra-02_B]

zoneset name FlexPod vsan 102
    member VM-Host-Infra-01_B
```

```
    member VM-Host-Infra-02_B

zoneset activate name FlexPod vsan 102d vsan 102
```