



## FlexPod Data Center with VMware vSphere 5.1 and Cisco Nexus 7000 using FCoE

Deployment Guide for FlexPod with VMware vSphere 5.1 and Cisco Nexus 7000 using FCoE and NetApp Clustered 8.1.2


Last Updated: November 22, 2013



Cisco  
Validated  
Design



Building Architectures to Solve Business Problems



## About the Authors

### **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.

### **Derek Huckaby, Technical Marketing Engineer, Unified Fabric Switching Services Product Group, Cisco Systems**

Derek Huckaby is a Technical Marketing Engineer for the Nexus 7000 Unified Fabric Switching products focusing on Nexus 7000 integration into FlexPod designs and Nexus 7000 services. Prior to joining the Nexus 7000 Product Marketing team, Derek led the team of Technical Marketing Engineers for the Data Center Application Services BU within Cisco. He began his work in network services at Cisco over 13 years ago specializing in application delivery and SSL termination solutions.

### **Haseeb Niazi, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems**

Haseeb has over 13 years of experience at Cisco dealing in Data Center, Security, and WAN Optimization related technologies. As a member of various solution teams and advanced services, Haseeb has helped many enterprise and service provider customers evaluate and deploy a wide range of Cisco solutions. Haseeb holds a master's degree in Computer Engineering from the University of Southern California.

---

**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.

**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 and 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>

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

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

© 2013 Cisco Systems, Inc. All rights reserved.

NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, FlexPod, MetroCluster, OnCommand, SnapLock, SnapMirror, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries.

# VMware vSphere 5.1 on FlexPod Clustered Data ONTAP with Nexus 7000 Using FCoE Deployment Guide

---

## Overview

Industry trends indicate a vast data center transformation toward shared infrastructures. By using virtualization, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements.

## 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 clustered Data ONTAP®.

## 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 non virtualized 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 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.

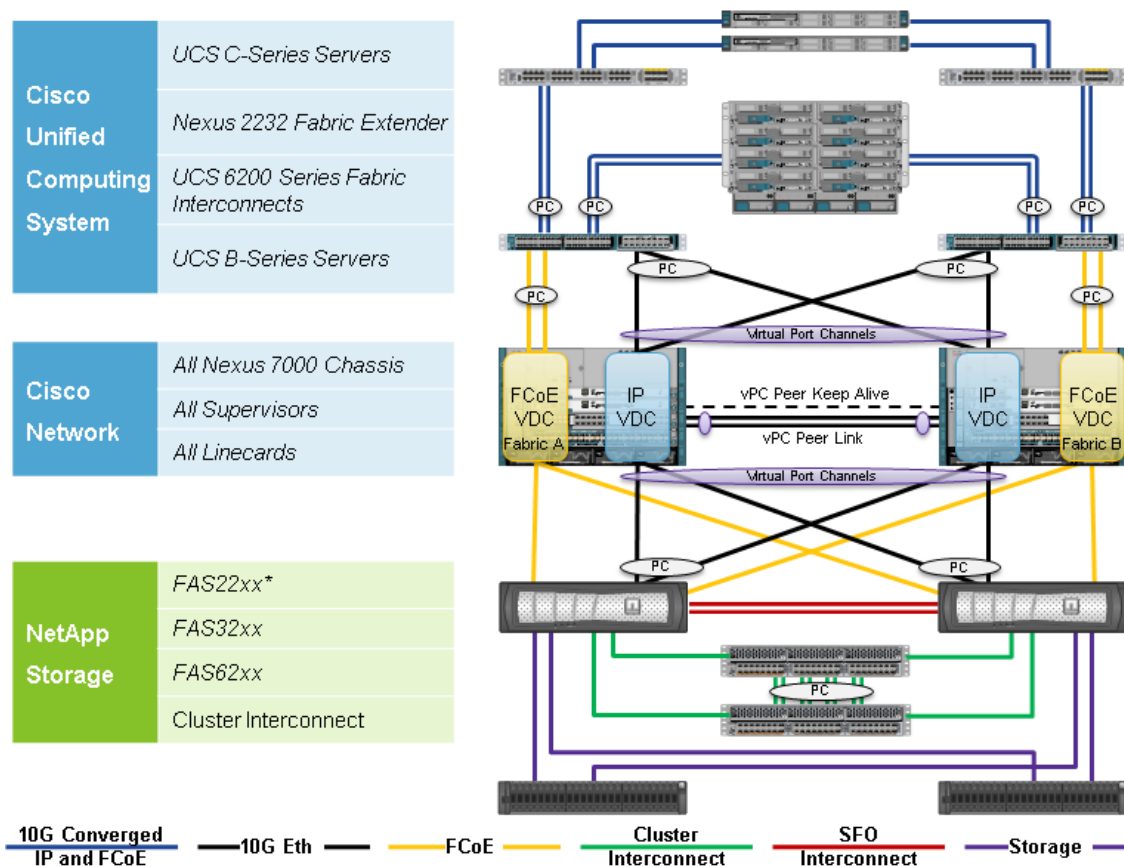
One benefit of the FlexPod architecture is the ability to customize or “flex” the environment to suit a customer’s requirements. This is why the reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an FCoE-based storage solution. A storage 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® 7000, 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

Port Channels (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.

Additionally, Figure 1 highlights the use of Cisco virtual device context (VDCs) as a means of logically segmenting the physical switch. By carving Cisco Nexus 7000 into the VDCs for IP and FCoE traffic, the architecture maximizes the hardware resource utilization while providing strong security and software fault isolation.

**Figure 1 FlexPod Distinct Uplink Design with Clustered Data ONTAP**



\*The FAS22xx fully supports IP-based storage, but does not support FCoE.

The reference configuration includes:

- Two Cisco Nexus 7000 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) running clustered Data ONTAP

**Note**

The FAS22xx is capable of supporting all IP-based storage protocols with the exception of FCoE.

Storage is provided by a NetApp FAS3250-AE (HA configuration in two chassis) operating in both clustered Data ONTAP and 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(1e)	Embedded management
	Cisco UCS C 200 M2 Server	2.1(1e)	Software bundle release
	Cisco UCS C 220 M3 Server	2.1(1e)	Software bundle release
	Cisco UCS B 200 M2 Server	2.1(1e)	Software bundle release
	Cisco UCS B 200 M3 Server	2.1(1e)	Software bundle release
	Cisco eNIC	2.1.2.38	Ethernet driver for Cisco VIC
	Cisco fNIC	1.5.0.20	FCoE driver for Cisco VIC
Network	Cisco Nexus 7000 Switch (F-series module required for FCoE support)	6.1(2)	Operating system version
Storage	NetApp FAS3250-AE	Clustered Data ONTAP 8.1.2	Operating system version
	Cisco Nexus 5596UP Cluster Interconnect	5.2(1)N1(1)	Operating System version

**Table 1**      **Software Revisions**

Layer	Compute	Version or Release	Details
Software	Cisco UCS Hosts	VMware vSphere ESXi™ 5.1	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™	5.1	VM (1 each): VMware vCenter
	NetApp OnCommand®	5.1	VM (1 each): OnCommand
	NetApp Virtual Storage Console (VSC)	4.1	Plug-in within VMware vCenter
	Cisco Nexus 1110-x	4.2(1)SV1(5.1a)	Virtual services appliance
	Cisco Nexus 1000v	4.2(1)SV2(1.1a)	Virtual services blade within the 1110-x
	NetApp NFS Plug-in for VMware vStorage APIs for Array Integration (VAAI)	1.0-018	Plug-in within VMware vCenter

## 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 steps 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:

Usage:

```
network port vlan create ?
[-node] <nodename>           Node
{ [-vlan-name] {<netport>|<ifgrp>} VLAN Name
|  -port {<netport>|<ifgrp>}    Associated Network Port
[-vlan-id] <integer> }         Network Switch VLAN Identifier
```

Example:

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

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. [Table 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 2**      **Necessary VLANs**

<b>VLAN Name</b>	<b>VLAN Purpose</b>	<b>ID Used in Validating This Document</b>
Mgmt in band	VLAN for in-band management interfaces	3175
Mgmt out of band	VLAN for out-of-band management interfaces	3171
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for NFS traffic	3170
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 (Cisco Nexus 1000v)	3176

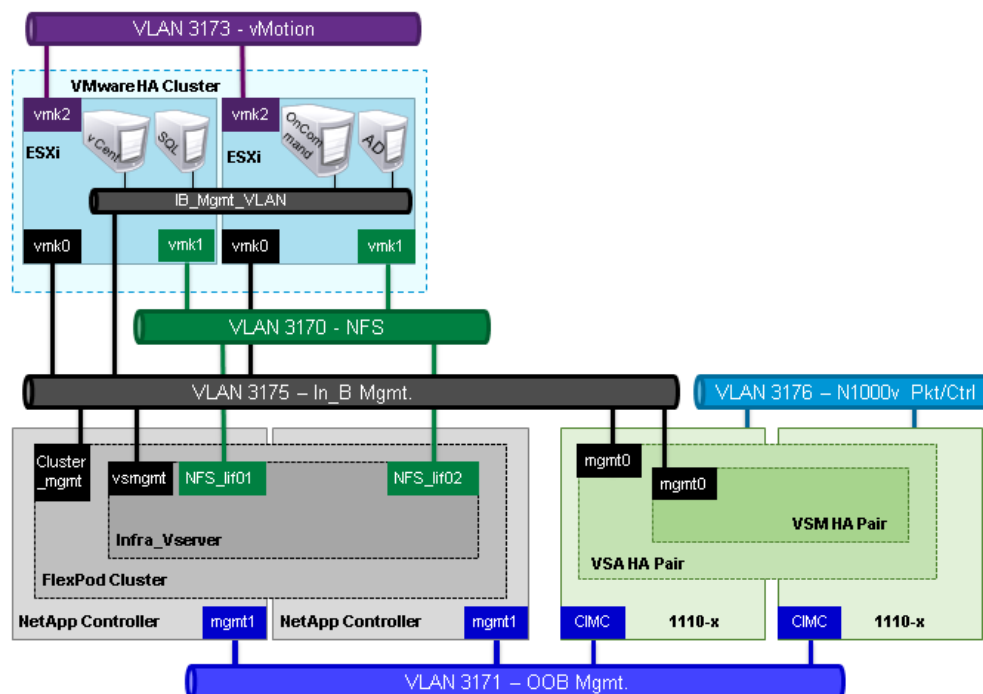
**Figure 2** Overview of LAN Usage

Table 3 lists the virtual storage area networks (VSANs) necessary for deployment as outlined in this guide.

**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 lists the virtual machines (VMs) necessary for deployment as outlined in this guide.

**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® core	

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 steps.

**Note**

In order for SNMP queries of the storage cluster to function properly, you should use separate in-band and out-of-band management VLANs. You must create a Layer 3 route between these VLANs.

**Table 5** Configuration Variables

Variable	Description	Customer Implementation Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01	
<<var_node01_mgmt_mask>>	Out-of-band management network netmask	
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_url_boot_software>>	Data ONTAP 8.1.2 URL; format: http://	
<<var_#_of_disks>>	Number of disks to assign to each storage controller	
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02	
<<var_node02_mgmt_mask>>	Out-of-band management network netmask	
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_clustername>>	Storage cluster host name	
<<var_cluster_base_license_key>>	Cluster base license key	
<<var_password>>	Global default administrative password	
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster	
<<var_clustermgmt_mask>>	In-band management network netmask	
<<var_clustermgmt_gateway>>	In-band management network default gateway	
<<var_dns_domain_name>>	DNS domain name	
<<var_nameserver_ip>>	DNS server IP(s)	
<<var_node_location>>	Node location string for each node	
<<var_node01>>	Cluster node 01 host name	
<<var_node02>>	Cluster node 02 host name	
<<var_raidsize>>	RAID group size for each node	



**Table 5**      **Configuration Variables**

<b>Variable</b>	<b>Description</b>	<b>Customer Implementation Value</b>
<<var_num_disks>>	Number of disks to assign to each storage data aggregate	
<<var_node01_sp_ip>>	Out-of-band cluster node 01 service processor management IP	
<<var_node01_sp_mask>>	Out-of-band management network netmask	
<<var_node01_sp_gateway>	Out-of-band management network default gateway	
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP	
<<var_node02_sp_mask>>	Out-of-band management network netmask	
<<var_node02_sp_gateway>	Out-of-band management network default gateway	
<<var_timezone>>	FlexPod time zone (for example, America/New_York)	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_snmp_contact>>	Administrator e-mail address	
<<var_snmp_location>>	Cluster location string	
<<var_oncommand_server_fqdn>>	VSC or OnCommand virtual machine fully qualified domain name (FQDN)	
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name	
<<var_mailhost>>	Mail server host name	
<<var_storage_admin_email>>	Administrator e-mail address	
<<var_security_cert_vserver_common_name>>	Infrastructure Vserver FQDN	
<<var_country_code>>	Two-letter country code	
<<var_state>>	State or province name	
<<var_city>>	City name	
<<var_org>>	Organization or company name	
<<var_unit>>	Organizational unit name	
<<var_security_cert_cluster_common_name>>	Storage cluster FQDN	
<<var_security_cert_node01_common_name>>	Cluster node 01 FQDN	
<<var_security_cert_node02_common_name>>	Cluster node 02 FQDN	

**Table 5 Configuration Variables**

Variable	Description	Customer Implementation Value
<<var_esxi_host1_nfs_ip>>	NFS VLAN IP address for each VMware ESXi host	
<<var_node01_nfs_lif_ip>>	Cluster node 01 NFS VLAN IP address	
<<var_node01_nfs_lif_mask>>	NFS VLAN netmask	
<<var_node02_nfs_lif_ip>>	Cluster node 02 NFS VLAN IP address	
<<var_node02_nfs_lif_mask>>	NFS VLAN netmask	
<<var_nexus_A_hostname>>	Cisco Nexus A host name	
<<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_nexus_A_ip_vdc>>	Name of the IP switching VDC on Cisco Nexus 7000 A	
<<var_nexus_A_ip_vdc_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address of the IP switching VDC on Cisco Nexus 7000 A	
<<var_nexus_A_fcoe_vdc>>	Name of the FCoE storage VDC on Cisco Nexus 7000 A	
<<var_nexus_A_fcoe_vdc_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address of the FCoE storage VDC on Cisco Nexus 7000 A	
<<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_nexus_B_ip_vdc>>	Name of the IP switching VDC on Cisco Nexus 7000 B	
<<var_nexus_B_ip_vdc_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address of the IP switching VDC on Cisco Nexus 7000 B	
<<var_nexus_B_fcoe_vdc>>	Name of the FCoE storage VDC on Cisco Nexus 7000 B	
<<var_nexus_B_fcoe_vdc_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address of the FCoE storage VDC on Cisco Nexus 7000 B	

**Table 5**      **Configuration Variables**

<b>Variable</b>	<b>Description</b>	<b>Customer Implementation Value</b>
<<var_ib-mgmt_vlan_id>>	In-band management network VLAN ID	
<<var_native_vlan_id>>	Native 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_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	
<<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_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_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<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	

**Table 5** Configuration Variables

Variable	Description	Customer Implementation Value
<<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_gateway>>	In-band management network default gateway	
<<var_vsm_hostname>>	Cisco Nexus 1000v VSM host name	
<<var_vcenter_server_ip>>	vCenter Server IP	
<<var_nodename>>	Name of node	
<<var_node01_rootaggrname>>	Root aggregate name of Node 01	
<<var_clustermgmt_port>>	Port for cluster management	
<<var_global_domain_name>>	Domain name	
<<var_dns_ip>>	IP address of the DNS server	
<<var_vsadmin_password>>	Password for VS admin account	
<<var_vserver_mgmt_ip>>	Management IP address for Vserver	
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver	
<<var_rule_index>>	Rule index number	
<<var_ftp_server>>	IP address for FTP server	
<<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_fcp_lif01a_wwpn>>	WWPN of FCP_LIF01a	
<<var_fcp_lif02a_wwpn>>	WWPN of FCP_LIF02a	
<<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_fcp_lif01b_wwpn>>	WWPN of FCP_LIF01b	
<<var_fcp_lif02b_wwpn>>	WWPN of FCP_LIF02b	
<<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	
<<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	

**Table 5**      **Configuration Variables**

Variable	Description	Customer Implementation Value
<<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_vmotion_vlan_id_ip_host-02>>	vMotion VLAN IP address for ESXi host 02	
<<var_vmotion_vlan_id_mask_host-02>>	vMotion VLAN netmask for ESXi host 02	

## Physical Infrastructure

### FlexPod Cabling on Clustered Data ONTAP

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

[Table 6](#) through [Table 22](#) in this section detail the prescribed and supported configuration of the NetApp FAS3250-AE running clustered Data ONTAP 8.1.2. This configuration uses a dual-port FCoE adapter, and external SAS disk shelves. For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool \(IMT\)](#).

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

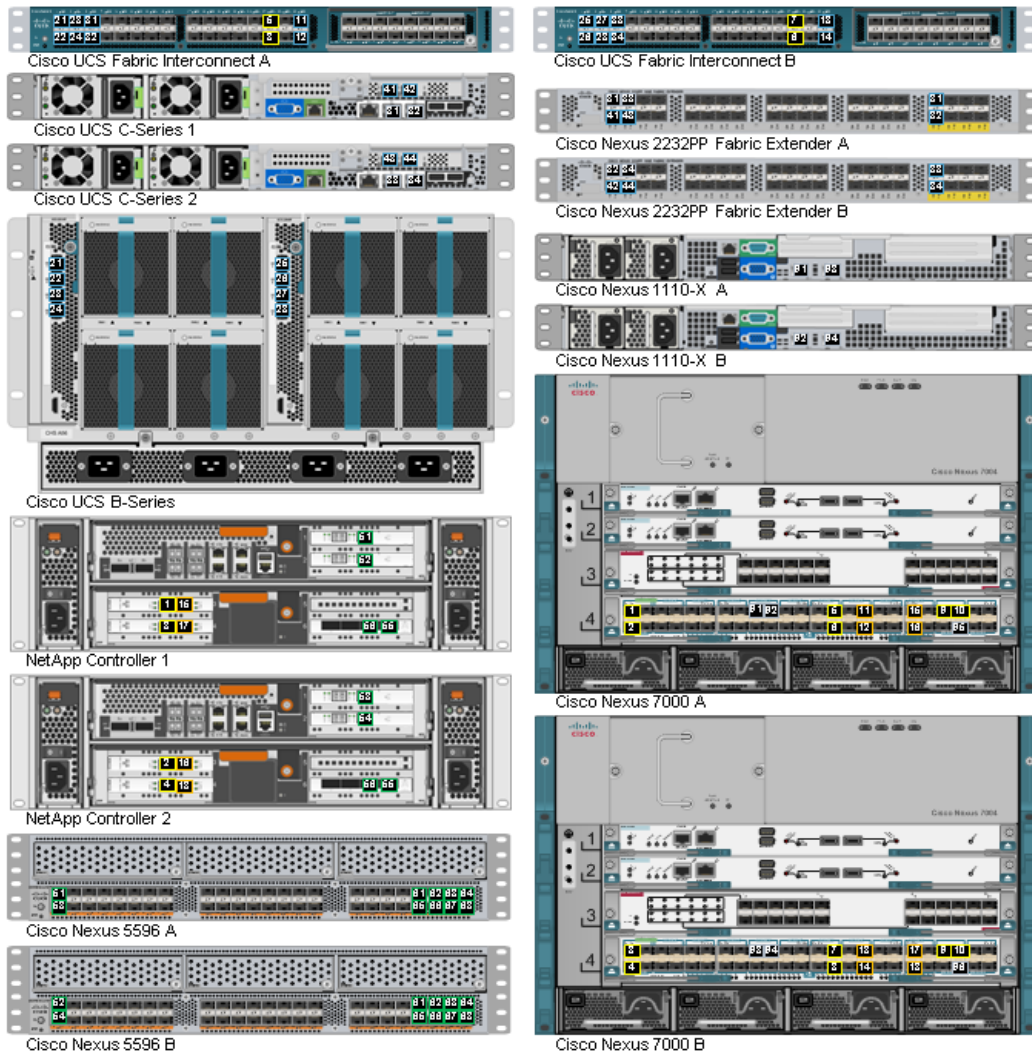
Be sure to follow the cabling directions in this section. Failure to do so will require necessary changes to the deployment procedures that follow because specific port locations are mentioned.

It is possible to order a FAS3250 system in a configuration different from what is prescribed in the tables in this section. Before starting, be sure that the configuration matches the descriptions in the tables and diagrams in this section.

[Figure 3](#) shows a cabling diagram for a FlexPod configuration using the Cisco Nexus 7000 and NetApp storage systems with clustered Data ONTAP. The labels indicate connections to endpoints rather than port numbers on the physical device. For example, SAS connections 27, 28, 29, and 30 as well as ACP connections 31 and 32 should be connected to the NetApp storage controller and disk shelves according to best practices for the specific storage controller and disk shelf quantity. For disk shelf cabling, see the Universal SAS and ACP Cabling Guide at:

[https://library.netapp.com/ecm/ecm\\_get\\_file/ECMM1280392](https://library.netapp.com/ecm/ecm_get_file/ECMM1280392)

**Figure 3** *FlexPod Cabling Diagram*



The information provided in [Table 6](#) through [Table 22](#) corresponds to each connection shown in [Figure 3](#).

**Table 6** *Cisco Nexus 7000 A Switching VDC Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 7000 A	Eth4/1	10GbE	NetApp controller 1	e3a
	Eth4/2	10GbE	NetApp controller 2	e3a
	Eth4/27	10GbE	Cisco UCS fabric interconnect A	Eth1/27
	Eth4/28	10GbE	Cisco UCS fabric interconnect B	Eth1/28
	Eth4/41	10GbE	Cisco Nexus 7000 B	Eth1/41
	Eth4/43	10GbE	Cisco Nexus 7000 B	Eth1/43
	Eth4/17	1GbE	Cisco Nexus 1110-X A	LOM A
	Eth4/19	1GbE	Cisco Nexus 1110-X B	LOM A
	Eth4/44	1GbE	GbE management switch	Any
	MGMT0	1GbE	GbE management switch	Any

**Note**

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

**Table 7** *Cisco Nexus 7000 B Switching VDC Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Ports
Cisco Nexus 7000 B	Eth4/1	10GbE	NetApp controller 1	e4a
	Eth4/2	10GbE	NetApp controller 2	e4a
	Eth4/28	10GbE	Cisco UCS fabric interconnect A	Eth1/28
	Eth4/27	10GbE	Cisco UCS fabric interconnect B	Eth1/27
	Eth4/41	10GbE	Cisco Nexus 7000 A	Eth1/41
	Eth4/43	10GbE	Cisco Nexus 7000 A	Eth1/43
	Eth4/17	1GbE	Cisco Nexus 1110-X A	LOM B
	Eth4/19	1GbE	Cisco Nexus 1110-X B	LOM B
	Eth4/44	1GbE	GbE management switch	Any
	MGMT0	1GbE	GbE management switch	Any

**Note**

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

**Table 8** *Cisco Nexus 7000 A Storage VDC Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 7000 A	Eth4/37	10GbE	NetApp controller 1	e3b
	Eth4/38	10GbE	NetApp controller 2	e3b
	Eth4/31	10GbE	Cisco Fabric Interconnect A	Eth1/31
	Eth4/32	10GbE	Cisco Fabric Interconnect A	Eth1/32
	MGMT0	1GbE	GbE management switch	Any

**Table 9** *Cisco Nexus 7000 B Storage VDC Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 7000 B	Eth4/37	10GbE	NetApp controller 1	e4b
	Eth4/38	10GbE	NetApp controller 2	e4b
	Eth4/31	10GbE	Cisco Fabric Interconnect B	Eth1/31
	Eth4/32	10GbE	Cisco Fabric Interconnect B	Eth1/32
	MGMT0	1GbE	GbE management switch	Any

**Table 10** *NetApp Controller 1 Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp Controller 1	e0M	100MbE	100MbE management switch	Any
	e0a	1GbE	GbE management switch	Any
	e0b	1GbE	GbE management switch	Any
	e0P	1 GbE	SAS shelves	ACP port
	c0a	10GbE	NetApp controller 2	c0a
	c0b	10GbE	NetApp controller 2	c0b
	e1a	10GbE	Cisco Nexus 5596 A	Eth1/1
	e2a	10GbE	Cisco Nexus 5596 B	Eth1/1
	e3a	10GbE	Cisco Nexus 7000 A (Switching)	Eth4/1
	e4a	10GbE	Cisco Nexus 7000 B (Switching)	Eth4/1
	e3b	10GbE	Cisco Nexus 7000 A (Storage)	Eth4/37
	e4b	10GbE	Cisco Nexus 7000 B (Storage)	Eth4/37



**Table 11** *NetApp Controller 2 Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp Controller 2	e0M	100MbE	100MbE management switch	Any
	e0a	1GbE	GbE management switch	Any
	e0b	1GbE	GbE management switch	Any
	e0P	1 GbE	SAS shelves	ACP port
	c0a	10GbE	NetApp controller 1	c0a
	c0b	10GbE	NetApp controller 1	c0b
	e1a	10GbE	Cisco Nexus 5596 A	Eth1/2
	e2a	10GbE	Cisco Nexus 5596 B	Eth1/2
	e3a	10GbE	Cisco Nexus 7000 A (Switching)	Eth4/2
	e4a	10GbE	Cisco Nexus 7000 B (Switching)	Eth4/2
	e3b	10GbE	Cisco Nexus 7000 A (Storage)	Eth4/38
	e4b	10GbE	Cisco Nexus 7000 B (Storage)	Eth4/38

**Table 12** *Cisco Nexus 5596 A Cluster Interconnect Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5596 A	Eth1/1	10GbE	NetApp controller 1	e1a
	Eth1/2	10GbE	NetApp controller 2	e1a
	Eth1/41	10GbE	Cisco Nexus 5596 B	Eth1/41
	Eth1/42	10GbE	Cisco Nexus 5596 B	Eth1/42
	Eth1/43	10GbE	Cisco Nexus 5596 B	Eth1/43
	Eth1/44	10GbE	Cisco Nexus 5596 B	Eth1/44
	Eth1/45	10GbE	Cisco Nexus 5596 B	Eth1/45
	Eth1/46	10GbE	Cisco Nexus 5596 B	Eth1/46
	Eth1/47	10GbE	Cisco Nexus 5596 B	Eth1/47
	Eth1/48	10GbE	Cisco Nexus 5596 B	Eth1/48
	MGMT0	1GbE	GbE management switch	Any

**Table 13** *Cisco Nexus 5596 B Cluster Interconnect Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5596 B	Eth1/1	10GbE	NetApp controller 1	e2a
	Eth1/2	10GbE	NetApp controller 2	e2a
	Eth1/41	10GbE	Cisco Nexus 5596 A	Eth1/41
	Eth1/42	10GbE	Cisco Nexus 5596 A	Eth1/42
	Eth1/43	10GbE	Cisco Nexus 5596 A	Eth1/43
	Eth1/44	10GbE	Cisco Nexus 5596 A	Eth1/44
	Eth1/45	10GbE	Cisco Nexus 5596 A	Eth1/45
	Eth1/46	10GbE	Cisco Nexus 5596 A	Eth1/46
	Eth1/47	10GbE	Cisco Nexus 5596 A	Eth1/47
	Eth1/48	10GbE	Cisco Nexus 5596 A	Eth1/48
	MGMT0	1GbE	GbE management switch	Any

**Note**

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

**Table 14**      **Cisco UCS Fabric Interconnect A Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/27	10GbE	Cisco Nexus 7000 A (Switching)	Eth4/27
	Eth1/28	10GbE	Cisco Nexus 7000 B (Switching)	Eth4/28
	Eth1/1	10GbE	Cisco UCS Chassis FEX A/Cisco Nexus 2232PP FEX A	IOM1/1
	Eth1/2	10GbE	Cisco UCS Chassis FEX A/Cisco Nexus 2232PP FEX A	IOM1/2
	Eth1/3	10GbE	Cisco UCS Chassis FEX A/Cisco Nexus 2232PP FEX A	IOM1/3
	Eth1/4	10GbE	Cisco UCS Chassis FEX A/Cisco Nexus 2232PP FEX A	IOM1/4
	Eth1/5	10GbE	Cisco UCS Chassis FEX A/Cisco Nexus 2232PP FEX A	Uplink 1
	Eth1/6	10GbE	Cisco UCS Chassis FEX A/Cisco Nexus 2232PP FEX A	Uplink 2
	Eth1/31	10GbE	Cisco Nexus 7000 A (Storage)	Eth4/31
	Eth1/32	10GbE	Cisco Nexus 7000 A (Storage)	Eth4/32
	MGMT0	1GbE	GbE management switch	Any
	L1	1GbE	Cisco UCS fabric interconnect B	L1
	L2	1GbE	Cisco UCS fabric interconnect B	L2

**Table 15** *Cisco UCS Fabric Interconnect B Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth1/28	10GbE	Cisco Nexus 7000 A (Switching)	Eth4/28
	Eth1/27	10GbE	Cisco Nexus 7000 B (Switching)	Eth4/27
	Eth1/1	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B	IOM2/1
	Eth1/2	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B	IOM2/2
	Eth1/3	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B	IOM2/3
	Eth1/4	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B	IOM2/4
	Eth1/5	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B	Uplink 1
	Eth1/6	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B	Uplink 2
	Eth1/31	10GbE	Cisco Nexus 7000 B (Storage)	Eth4/31
	Eth1/32	10GbE	Cisco Nexus 7000 B (Storage)	Eth4/32
	MGMT0	1GbE	GbE management switch	Any
	L1	1GbE	Cisco UCS fabric interconnect A	L1
	L2	1GbE	Cisco UCS fabric interconnect A	L2

**Table 16** *Cisco Nexus 2232PP FEX A*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 2232PP FEX A	Port 1	1GbE	Cisco UCS C-Series 1	M1
	Port 2	10GbE	Cisco UCS C-Series 1	Port 0
	Port 3	1GbE	Cisco UCS C-Series 2	M1
	Port 4	10GbE	Cisco UCS C-Series 2	Port 0
	Port 2/1	10GbE	Cisco UCS fabric interconnect A	Eth1/5
	Port 2/2	10GbE	Cisco UCS fabric interconnect A	Eth1/6

**Table 17** *Cisco Nexus 2232PP FEX B*

Local Device	Local Port	Connection	Remote Devices	Remote Port
Cisco Nexus 2232PP FEX B	Port 1	1GbE	Cisco UCS C-Series 1	M2
	Port 2	10GbE	Cisco UCS C-Series 1	Port 1
	Port 3	1GbE	Cisco UCS C-Series 2	M2
	Port 4	10GbE	Cisco UCS C-Series 2	Port 1
	Port 2/1	10GbE	Cisco UCS fabric interconnect B	Eth1/5
	Port 2/2	10GbE	Cisco UCS fabric interconnect B	Eth1/6

**Table 18** *Cisco UCS C-Series 1*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series Server 1	LOM1	1GbE	Cisco Nexus 2232PP FEX A	Port 1
	LOM2	1GbE	Cisco Nexus 2232PP FEX B	Port 1
	Port0	10GbE	Cisco Nexus 2232PP FEX A	Port 2
	Port1	10GbE	Cisco Nexus 2232PP FEX B	Port 2

**Table 19** *Cisco UCS C-Series 2*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series Server 2	LOM1	1GbE	Cisco Nexus 2232PP FEX A	Port 3
	LOM2	1GbE	Cisco Nexus 2232PP FEX B	Port 3
	Port0	10GbE	Cisco Nexus 2232PP FEX A	Port 4
	Port1	10GbE	Cisco Nexus 2232PP FEX B	Port 4

**Table 20** *Cisco Nexus 1110-XA*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series Server 3	LOM A	1GbE	Cisco Nexus 7000 A	Eth4/17
	LOM B	1GbE	Cisco Nexus 7000 B	Eth4/19

**Table 21** *Cisco Nexus 1110-XB*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series Server 4	LOM A	1GbE	Cisco Nexus 7000 A	Eth4/17
	LOM B	1GbE	Cisco Nexus 7000 B	Eth4/19

**Table 22** *NetApp FAS3250 Card Layout*

Slot	Part Number	Description
1	X1117A-R6	NIC 2-port 10GbE (ports e1a and e1b)
2	X1117A-R6	NIC 2-port 10GbE (ports e2a and e2b)
3	X1140A-R6	Unified target 2-port 10GbE (ports e3a and e3b)
4	X1140A-R6	Unified target 2-port 10GbE (ports e4a and e4b)
5	X1971A-R5	Flash Cache™ – 512GB
6	X2065A-R6	SAS, 4-port, 6Gb

# Storage Configuration

## Controller FAS32xx Series

**Table 23** *Controller FAS32XX Series Prerequisites*

Requirement	Reference	Comments
Physical site where storage system needs to be installed must be ready	Site Reference Guide: <a href="http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf">http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf</a>	Refer to the “Site Preparation” section
Storage system connectivity requirements	Site Reference Guide: <a href="http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf">http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf</a>	Refer to the “System Connectivity Requirements” section
Storage system general power requirements	Site Reference Guide: <a href="http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf">http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf</a>	Refer to the “Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements” section
Storage system model-specific requirements	Site Reference Guide: <a href="http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf">http://support.netapp.com/NOW/public/knowledge/docs/hardware/NetApp/site/pdf/site.pdf</a>	Refer to the “FAS32xx/V32xx Series Systems” section

## System Configuration Guides

System configuration guides provide supported hardware and software components for the specific Data ONTAP version. These online guides provide configuration information for all NetApp storage appliances currently supported by the Data ONTAP software. They also provide 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 System Configuration Guides at:  
<https://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/>
2. Click the appropriate NetApp storage appliance and then click the component you want to view. Alternatively, to compare components by storage appliance, click a component and then click the NetApp storage appliance you want to view.

## 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 DS2246 Series

### DS2246 Disk Shelves

To install a disk shelf for a new storage system, see:

<https://now.netapp.com/NOW/knowledge/docs/hardware/filer/210-04881+A0.pdf>

For information on cabling with the controller model, see SAS Disk Shelves Universal SAS and ACP Cabling Guide at:

[https://now.netapp.com/NOW/knowledge/docs/hardware/filer/215-05500\\_A0.pdf](https://now.netapp.com/NOW/knowledge/docs/hardware/filer/215-05500_A0.pdf)

The following information applies to DS2246 disk shelves:

- SAS disk drives use software-based disk ownership. Ownership of a disk drive is assigned to a specific storage system by writing software ownership information on the disk drive rather than by using the topography of the storage system's physical connections.
- Connectivity terms used: shelf-to-shelf (daisy-chain), controller-to-shelf (top connections), and shelf-to controller (bottom connections).
- Unique disk shelf IDs must be set per storage system (a number from 0 through 98).
- Disk shelf power must be turned on to change the digital display shelf ID. The digital display is on the front of the disk shelf.
- Disk shelves must be power-cycled after the shelf ID is changed for it to take effect.
- Changing the shelf ID on a disk shelf that is part of an existing storage system running Data ONTAP requires that you wait at least 30 seconds before turning the power back on so that Data ONTAP can properly delete the old disk shelf address and update the copy of the new disk shelf address.
- Changing the shelf ID on a disk shelf that is part of a new storage system installation (the disk shelf is not yet running Data ONTAP) requires no wait; you can immediately power-cycle the disk shelf.

## Cisco NX5596 Cluster Network Switch Configuration

**Table 24** *Cisco Nexus 5596 Cluster Network Switch Configuration Prerequisites*

Configuration Prerequisites
Rack and connect power to the new Cisco Nexus 5596 switches
Provide a terminal session that connects to the switch's serial console port (9600, 8, n, 1)
Connect the <b>mgmt0</b> port to the management network and be prepared to provide IP address information
Obtain password for admin
Determine switch name
Identify SSH key type (dsa, rsa, or rsa1)
Set up an e-mail server for Cisco Smart Call Home and IP connectivity between the switch and the e-mail server
Provide SNMP contact information for Cisco Smart Call Home (name, phone, street address)



**Table 24 Cisco Nexus 5596 Cluster Network Switch Configuration Prerequisites****Configuration Prerequisites**

Identify a CCO ID associated with an appropriate Cisco SMARTnet® Service contract for Cisco Smart Call Home

Enable Cisco SMARTnet Service for the device to be registered for Cisco Smart Call home

## Initial Setup of Cisco Nexus 5596 Cluster Interconnect

The first time a Cisco Nexus 5596 cluster interconnect is accessed, it runs a setup program that prompts the user to enter an IP address and other configuration information needed for the switch to communicate over the management Ethernet interface. This information is required to configure and manage the switch. If the configuration must be changed later, the setup wizard can be accessed again by running the setup command in EXEC mode.

To set up the Cisco Nexus 5596 cluster interconnect, follow these steps on both cluster interconnects.

1. Provide applicable responses to the setup prompts displayed on the Cisco Nexus 5596 cluster interconnect.

```
Do you want to enforce secure password standard (yes/no): yes
Enter the password for the "admin": <password>
Confirm the password for "admin": <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: <switchname>
Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <ic_mgmt0_ip>
Mgmt0 IPv4 netmask: <ic_mgmt0_netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <ic_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: <ntp_server_ip>
Enter basic FC configurations (yes/no) [n]: Enter
```

2. At the end of the setup, the configuration choices are displayed. Verify the information and save the configuration at this time.

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

## Download and Install NetApp Cluster Switch Software

When the Cisco Nexus 5596 is being used as a cluster network switch with Data ONTAP 8.1.2, it should be running NX-OS version 5.2(1)N1(1). The **show version** command from the switch command line interface will show the switch version currently running on the switch. If the currently running version is not 5.2(1)N1(1), go to the [NetApp Support site](#) and download and install NX-OS 5.2(1)N1(1) for the Cisco Nexus 5596 switch. Make sure both cluster interconnects are running NX-OS version 5.2(1)N1(1).

## Download and Merge of NetApp Cluster Switch Reference Configuration File

Cluster network and management network switches are shipped without the configuration files installed. These files must be downloaded to the switches during deployment. Configuration files must be downloaded when the cluster network and management network switches are first installed or after the Cisco switch software is updated or reinstalled.

After the initial setup is complete, the NetApp cluster network switch reference configuration must be transferred to the switch and merged with the existing configuration. Instructions for this task and the reference configuration files for the appropriate switches are available on the [NetApp Support site](#).

To download configuration files to a host and install them on a Cisco Nexus 5596 switch, follow these steps on both the cluster interconnects:

1. Obtain a console connection to the switch. Verify the existing configuration on the switch by running the **show run** command.
2. Log in to the switch. Make sure that the host recognizes the switch on the network (for example, use the ping utility).
3. Enter the following command:

```
copy <transfer protocol>: bootflash: vrf management
```

4. Verify that the configuration file is downloaded.
5. Merge the configuration file into the existing **running-config**. Run the following command, where **<config file name>** is the file name for the switch type. A series of warnings regarding PortFast is displayed as each port is configured.

```
copy <config file name> running-config
```

6. Verify the success of the configuration merge by running the show run command and comparing its output to the contents of the configuration file (**a .txt file**) that was downloaded.
  - a. The output for both installed-base switches and new switches should be identical to the contents of the configuration file for the following items:
    - **banner** (should match the expected version)
    - Switch port descriptions such as **description Cluster Node x**
    - The new ISL algorithm **port-channel load-balance Ethernet source-dest-port**
  - b. The output for new switches should be identical to the contents of the configuration file for the following items:
    - Port channel
    - Policy map
    - System QoS
    - Interface
    - Boot
  - c. The output for installed-base switches should have the flow control receive and send values on for the following items:
    - Interface port-channel 1 and 2
    - Ethernet interface 1/41 through Ethernet interface 1/48
7. Copy the **running-config to the startup-config**.

```
copy running-config startup-config
```

## Cisco Smart Call Home Setup

To configure Smart Call Home on a Cisco Nexus 5596 switch, follow these steps:

1. Enter the mandatory system contact using the **snmp-server contact** command in global configuration mode. Then run the **callhome** command to enter callhome configuration mode.
 

```
NX-5596#config t
NX-5596(config)#snmp-server contact <sys-contact>
NX-5596(config)#callhome
```
2. Configure the mandatory contact information (phone number, e-mail address, and street address).
 

```
NX-5596(config-callhome)#email-contact <email-address>
NX-5596(config-callhome)#phone-contact <+1-000-000-0000>
NX-5596(config-callhome)#streetaddress <a-street-address>
```
3. Configure the mandatory e-mail server information. The server address is an IPv4 address, IPv6 address, or the domain-name of a SMTP server to which Call Home will send e-mail messages. Optional port number (default=25) and VRF may be configured.
 

```
NX-5596(config-callhome)#transport email smtp-server <ip-address> port 25 use-vrf <vrf-name>
```
4. Set the destination profile CiscoTAC-1 e-mail address to callhome@cisco.com
 

```
NX-5596(config-callhome)#destination-profile CiscoTAC-1 email-addr callhome@cisco.com vrf management
```
5. Enable periodic inventory and set the interval.
 

```
NX-5596(config-callhome)#periodic-inventory notification
NX-5596(config-callhome)#periodic-inventory notification interval 30
```
6. Enable callhome, exit, and save the configuration.
 

```
NX-5596(config-callhome)#enable
NX-5596(config-callhome)#end
NX-5596#copy running-config startup-config
```
7. Send a callhome inventory message to start the registration process.
 

```
NX-5596#callhome test inventory
trying to send test callhome inventory message
successfully sent test callhome inventory message
```
8. Watch for an e-mail from Cisco regarding the registration of the switch. Follow the instructions in the e-mail to complete the registration for Smart Call Home.

## SNMP Monitoring Setup

Configure SNMP by using the following example as a guideline. This example configures a host receiver for SNMPv1 traps and enables all link up/down traps.

```
NX-5596(config)# snmp-server host <ip-address> traps { version 1 } <community>
[udp_port <number>]
NX-5596(config)# snmp-server enable traps link
```

## Clustered Data ONTAP 8.1.2

### Node 1

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 8.1.2, proceed to step 4 to load Data ONTAP 8.1.2 software. If Data ONTAP 8.1.2 is already loaded, proceed to step 16.

4. Allow the system to boot up.

```
boot_ontap
```

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




---

**Note** If Data ONTAP 8.1.2 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.1.2 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_node01_mgmt_ip>> <<var_node01_mgmt_mask>> <<var_node01_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 user name, 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.

**Note**

y

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the LOADER 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 not listed, the system is not set to boot in clustered Data ONTAP.

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

```
setenv bootarg.init.boot_clustered true
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.

```
y
```

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

```
y
```

**Note**

The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue to node 02 configuration while the disks for node 01 are zeroing.

## Node 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, enter:

```
printenv
```

3. If the last-OS-booted-ver parameter is not set to 8.1.2, proceed to step 4 to load Data ONTAP 8.1.2 software. If Data ONTAP 8.1.2 is already loaded, proceed to step 16.
4. Allow the system to boot up.

```
boot_ontap
```

5. Press Ctrl-C when **Press Ctrl-C for Boot Menu** is displayed.

```
Ctrl-C
```




---

**Note** If Data ONTAP 8.1.2 is not the version of software being booted, proceed with the following steps to install new software. If Data ONTAP 8.1.2 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_node02_mgmt_ip>> <<var_node02_mgmt_mask>> <<var_node02_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 user name, indicating no user name.

```
Enter
```

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

```
y
```

14. Select 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 not listed, the system is not set to boot in clustered Data ONTAP.

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

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

18. At the LOADER-A prompt, enter:

```
autoboot
```

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

```
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.**

```
y
```

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

```
y
```



**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.

## Cluster Create in Clustered Data ONTAP

**Table 25** *Creating Cluster in Clustered Data ONTAP Prerequisites*

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Clustered Data ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>

The first node in the cluster performs the cluster create operation. All other nodes perform a **cluster join** operation. The first node in the cluster is considered Node01.

1. During the first node boot, the Cluster Setup wizard starts running on the console.

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



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

2. Enter the following command to create a new cluster:

```
create
```

3. The system defaults are displayed.

```
System Defaults:
Private cluster network ports [e1a,e2a].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
Do you want to use these defaults? {yes, no} [yes]:
```

4. NetApp recommends accepting the system defaults. To accept the system defaults, press Enter.



**Note** Cluster is created; this can take a minute or two.

5. The steps to create a cluster are displayed.

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



**Note** For this validated architecture we recommend you install license keys for SnapRestore®, NFS, FCP, FlexClone®, and SnapManager® Suite. After you finish entering the license keys, press Enter.

```
Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
Enter the cluster management interface port [e0a]: e0a
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway:
<<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

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





**Note** If you have more than one name server IP address, separate them with a comma.

7. Set up the node.

```
Where is the controller located []:<<var_node_location>>
Enter the node management interface port [e0M]: e0b
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
enter the node management interface netmask:<<var_node01_mgmt_mask>>
Enter the node management interface default gateway:<<var_node01_mgmt_gateway>>
```



**Note** The node management interface should be in a different subnet than the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

8. Press Enter to accept the AutoSupport™ message.

9. Reboot node 01.

```
system node reboot <<var_node01>>
y
```

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

```
Ctrl - C
```

11. Select 5 to boot into maintenance mode.

```
5
```

12. When prompted **Continue with boot?**, enter **y**.

13. To verify the HA status of your environment, run the following command:

```
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.

14. To see how many disks are unowned, enter:

```
disk show -a
```



**Note** No disks should be owned in this list.

15. 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>>
```

16. Reboot the controller.

```
halt
```

17. At the **LOADER-A** prompt, enter:

```
autoboot
```

## Cluster Join in Clustered Data ONTAP

**Table 26** *Joining Cluster in Clustered Data ONTAP Prerequisites*

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>

The first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered Node01, and the node joining the cluster in this example is Node02.

1. During the node boot, the Cluster Setup wizard starts running on the console.

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



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

2. Enter the following command to join a cluster:

```
join
```

3. The system defaults are displayed.

```
System Defaults:
Private cluster network ports [e1a,e2a].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
Do you want to use these defaults? {yes, no} [yes]:
```

4. NetApp recommends accepting the system defaults. To accept the system defaults, press Enter.



**Note** The cluster creation can take a minute or two.

5. The steps to create a cluster are displayed.

```
Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter
```



**Note** The node should find the cluster name.

6. Set up the node.

```
Enter the node management interface port [e0M]: e0b
Enter the node management interface IP address: <<var_node02_mgmt_ip>>
Enter the node management interface netmask: Enter
Enter the node management interface default gateway: Enter
```

7. The node management interface should be in a subnet different from the cluster management interface. The node management interfaces can reside on the out-of-band management network, and the cluster management interface can be on the in-band management network.

8. Press Enter to accept the AutoSupport message.

9. Log in to the Cluster Interface with the admin user id and <<var\_password>>.

10. Reboot node 02.

```
system node reboot <<var_node02>>
y
```

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

```
Ctrl - C
```

12. Select 5 to boot into maintenance mode.

```
5
```

13. At the question, **Continue with boot?** enter:

```
y
```

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



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

```
ha-config show
```

15. To see how many disks are unowned, enter:

```
disk show -a
```

16. Assign disks.



**Note** This reference architecture allocates half the disks to each controller. Workload design could dictate different percentages, however. Assign all remaining disks to node 02.

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

17. Reboot the controller:

```
halt
```

18. At the **LOADER-A** prompt, enter:

```
autoboot
```

19. Press **Ctrl-C** for boot menu when prompted.

```
Ctrl-C
```

## Log in to the Cluster

Open an SSH connection to cluster IP or host name and log in to the admin user with the password you provided earlier.

## Zero All Spare Disks

Zero all spare disks in the cluster.

```
disk zerospares
```

## Set Auto-Revert on Cluster Management

To set the auto-revert parameter on the cluster management interface, enter:

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

## Failover Groups Management in Clustered Data ONTAP

Create a management port failover group.

```
network interface failover-groups create -failover-group fg-cluster-mgmt -node <<var_node01>> -port e0a
network interface failover-groups create -failover-group fg-cluster-mgmt -node <<var_node02>> -port e0a
```

## Assign Management Failover Group to Cluster Management LIF

Assign the management port failover group to the cluster management LIF.

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt -failover-group fg-cluster-mgmt
```

## Failover Groups Node Management in Clustered Data ONTAP

Create a management port failover group.

```
network interface failover-groups create -failover-group fg-node-mgmt-01 -node <<var_node01>> -port e0b
network interface failover-groups create -failover-group fg-node-mgmt-01 -node <<var_node01>> -port e0M
network interface failover-groups create -failover-group fg-node-mgmt-02 -node <<var_node02>> -port e0b
network interface failover-groups create -failover-group fg-node-mgmt-02 -node <<var_node02>> -port e0M
```

## Assign Node Management Failover Groups to Node Management LIFs

Assign the management port failover group to the cluster management LIF.

```
network interface modify -vserver <<var_node01>> -lif mgmt1 -auto-revert true
-use-failover-group enabled -failover-group fg-node-mgmt-01
network interface modify -vserver <<var_node02>> -lif mgmt1 -auto-revert true
-use-failover-group enabled -failover-group fg-node-mgmt-02
```

## Flash Cache in Clustered Data ONTAP

Follow these steps to enable Flash Cache on each node:

Run the following commands from the cluster management interface:

```
system node run -node <<var_node01>> options flexscale.enable on
system node run -node <<var_node01>> options flexscale.lopri_blocks off
system node run -node <<var_node01>> options flexscale.normal_data_blocks on
system node run -node <<var_node02>> options flexscale.enable on
system node run -node <<var_node02>> options flexscale.lopri_blocks off
system node run -node <<var_node02>> options flexscale.normal_data_blocks on
```



### Note

- Data ONTAP 8.1 and later does not require a separate license for Flash Cache.
- For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, see [TR-3832: Flash Cache Best Practices Guide](#). Before customizing the settings, determine whether the custom settings are required or if the default settings are sufficient.

## 64-Bit Aggregates in Clustered Data ONTAP

A 64-bit aggregate containing the root volume is created during the Data ONTAP setup process. To create additional 64-bit aggregates, determine the aggregate name, the node on which to create it, and the number of disks it will contain.

1. Execute the following command to create new aggregates:

```
aggr create -aggregate aggr01 -nodes <<var_node01>> -B 64 -diskcount 3
aggr create -aggregate aggr02 -nodes <<var_node02>> -B 64 -diskcount 3
```



### Note

- In this configuration an aggregate with a minimum size of three disks is created for the FlexPod management infrastructure. This provides flexibility to either add to this aggregate or create new aggregates for production workloads. If the disks in this implementation are smaller than 600GB, consider adding a 4-disk aggregate.
- Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.
- The default RAID group size for the aggregate can be specified by adding the “-s <raidsize>” parameter. Calculate the RAID group size to allow for roughly balanced (same size) RAID groups of from 12 through 20 disks (for SAS disks) within the aggregate. For example, if 52 disks were being assigned to the aggregate, select a RAID group size of 18. A RAID group size of 18 would yield two 18-disk RAID groups and one 16-disk RAID group. Keep in mind that the default RAID group size is 16 disks, and that the larger the RAID group size, the longer the disk rebuild time in case of a failure.

- The aggregate cannot be created until disk zeroing completes. Use the **aggr show** command to display aggregate creation status. Do not proceed until both aggr01 and aggr02 are online.

2. Disable Snapshot copies for the two data aggregates just created.

```
node run <<var_node01>> aggr options aggr01 nosnap on
node run <<var_node02>> aggr options aggr02 nosnap on
```

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

```
node run <<var_node01>> snap delete -A -a -f aggr01
node run <<var_node02>> snap delete -A -a -f aggr02
```

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

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

## Service Processor

Gather information about the network and the AutoSupport settings before configuring the Service Processor (SP).

Configure the SP using DHCP or static addressing. If the SP uses a static IP address, verify that the following SP prerequisites have been met:

- An available static IP address
- The network netmask
- The network gateway IP
- AutoSupport information

A best practice is to configure the AutoSupport recipients and mail host before configuring the SP. Data ONTAP automatically sends AutoSupport configuration to the SP, allowing the SP to send alerts and notifications through an AutoSupport message to the system administrative recipients specified in AutoSupport. When configuring the SP, enter the name or the IP address of the AutoSupport mail host, when prompted.

A service processor needs to be set up on each node.

## Upgrade the Service Processor on Each Node to the Latest Release

With Data ONTAP 8.1.2, 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 nodes in your cluster. You will need to download the .zip file to a Web server that is reachable from the cluster management interface. In step 1a of the instructions substitute the following command: **system image get -node \* -package [http://web\\_server\\_name/path/SP\\_FW.zip](http://web_server_name/path/SP_FW.zip)**.

Also, instead of **run local**, use **system node run <<var\_nodename>>**, then execute steps 2–6 on each node.

## Configure the Service Processor on Node 01

1. From the cluster shell, enter the following command:

```
system node run <<var_node01>> sp setup
```

2. Enter the following to set up the SP:

```
Would you like to configure the SP? Y
Would you like to enable DHCP on the SP LAN interface? no
Please enter the IP address of the SP[]: <<var_node01_sp_ip>>
Please enter the netmask of the SP[]: <<var_node01_sp_mask>>
Please enter the IP address for the SP gateway[]: <<var_node01_sp_gateway>>
```

## Configure the Service Processor on Node 02

1. From the cluster shell, enter the following command:

```
system node run <<var_node02>> sp setup
```

2. Enter the following to set up the SP:

```
Would you like to configure the SP? Y
Would you like to enable DHCP on the SP LAN interface? no
Please enter the IP address of the SP[]: <<var_node02_sp_ip>>
Please enter the netmask of the SP[]: <<var_node02_sp_mask>>
Please enter the IP address for the SP gateway[]: <<var_node02_sp_gateway>>
```

## Storage Failover in Clustered Data ONTAP

Run the following commands in a failover pair to enable storage failover:

1. Enable failover on one of the two nodes.

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



**Note** Enabling failover on one node enables it for both nodes.

2. Enable HA mode for two-node clusters only.



**Note** Do not run this command for clusters with more than two nodes because it will cause problems with failover.

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

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

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

## IFGRP LACP in Clustered Data ONTAP

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

1. Run the following commands on the command line to create interface groups (ifgrps).

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e3a
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e4a
ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e3a
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e4a
```



### Note

- All interfaces must be in the down status before being added to an interface group.
- The interface group name must follow the standard naming convention of a0x.

## VLAN in Clustered Data ONTAP

Create NFS VLANs.

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>
```

## Jumbo Frames in Clustered Data ONTAP

To configure a clustered Data ONTAP network port to use jumbo frames (which usually have an MTU of 9,000 bytes), run the following command from the cluster shell:

```
network port modify -node <<var_node01>> -port a0a-<<var_nfs_vlan_id>> -mtu 9000
```

Warning: Changing the network port settings will cause a several second interruption in carrier.

```
Do you want to continue? {y|n}: y
```

```
network port modify -node <<var_node02>> -port a0a-<<var_nfs_vlan_id>> -mtu 9000
```

Warning: Changing the network port settings will cause a several second interruption in carrier.

```
Do you want to continue? {y|n}: y
```

## NTP in Clustered Data ONTAP

To configure time synchronization on the cluster, follow these steps:

1. Set the time zone for the cluster.

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



### Note

For example, in the Eastern United States, the time zone is America/New\_York.

2. Set the date for the cluster.



```
date <ccyyymmddhhmm>
```

**Note**

The format for the date is <[Century][Year][Month][Day][Hour][Minute]>; for example, 201208081240.

3. Configure the Network Time Protocol (NTP) for each node in the cluster.

```
system services ntp server create -node <<var_node01>> -server
<<var_global_ntp_server_ip>>
system services ntp server create -node <<var_node02>> -server
<<var_global_ntp_server_ip>>
```

4. Enable the NTP for the cluster.

```
system services ntp config modify -enabled true
```

## SNMP in Clustered Data ONTAP

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

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

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

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

## SNMPv1 in Clustered Data ONTAP

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

```
snmp community delete all
snmp community add row <<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 in Clustered Data ONTAP

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

1. Create a user called snmpv3user.

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

2. Select all of the default authoritative entities and select md5 as the authentication protocol.
3. Enter an eight-character minimum-length password for the authentication protocol, when prompted.
4. Select des as the privacy protocol.
5. Enter an eight-character minimum-length password for the privacy protocol, when prompted.

## AutoSupport HTTPS in Clustered Data ONTAP

AutoSupport sends support summary information to NetApp through HTTPS.

Execute the following commands to configure AutoSupport:

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

## Cisco Discovery Protocol in Clustered Data ONTAP

To enable Cisco Discovery Protocol (CDP) on the NetApp storage controllers, follow these steps:



### Note

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

To enable CDP on the NetApp storage controllers, execute the following commands:

Enable CDP on Data ONTAP:

```
node run -node <<var_node01>> options cdpd.enable on
node run -node <<var_node02>> options cdpd.enable on
```

## Vserver

To create an infrastructure Vserver, follow these steps:

1. Run the Vserver setup wizard.

```
vserver setup
```

Welcome to the Vserver Setup Wizard, which will lead you through the steps to create a virtual storage server that serves data to clients.

You can enter the following commands at any time:

"help" or "?" if you want to have a question clarified,  
 "back" if you want to change your answers to previous questions, and  
 "exit" if you want to quit the Vserver Setup Wizard. Any changes you made before typing "exit" will be applied.

You can restart the Vserver Setup Wizard by typing "vserver setup". To accept a default or omit a question, do not enter a value.

Step 1. Create a Vserver.

You can type "back", "exit", or "help" at any question.

2. Enter the Vserver name.

```
Enter the Vserver name:Infra_Vserver
```

3. Select the Vserver data protocols to configure.

```
Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi}:nfs,
fcp
```

4. Select the Vserver client services to configure.

```
Choose the Vserver client services to configure {ldap, nis, dns}:Enter
```

5. Enter the Vserver's root volume aggregate:

```
Enter the Vserver's root volume aggregate {aggr01, aggr02} [aggr01]:aggr01
```

6. Enter the Vserver language setting. English is the default [C].

```
Enter the Vserver language setting, or "help" to see all languages [C]:Enter
```

7. Enter the Vserver's security style:

```
Enter the Vservers root volume's security style {unix, ntfs, mixed} [unix]: Enter
```

8. Answer no to Do you want to create a data volume?

```
Do you want to create a data volume? {yes, no} [Yes]: no
```

9. Answer no to Do you want to create a logical interface?

```
Do you want to create a logical interface? {yes, no} [Yes]: no
```

10. Answer no to Do you want to Configure FCP? {yes, no} [yes]: no.

```
Do you want to Configure FCP? {yes, no} [yes]: no
```

11. Add the two data aggregates to the Infra\_Vserver aggregate list for NetApp Virtual Console.

```
vserver modify -vserver Infra_Vserver -aggr-list aggr01, aggr02
```

## Create Load Sharing Mirror of Vserver Root Volume in Clustered Data ONTAP

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

```
volume create -vserver Infra_Vserver -volume root_vol_m01 -aggregate aggr01 -size 20MB -type DP
volume create -vserver Infra_Vserver -volume root_vol_m02 -aggregate aggr02 -size 20MB -type DP
```

2. Create the mirroring relationships.

```
snapmirror create -source-path //Infra_Vserver/root_vol -destination-path //Infra_Vserver/root_vol_m01 -type LS
snapmirror create -source-path //Infra_Vserver/root_vol -destination-path //Infra_Vserver/root_vol_m02 -type LS
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path //Infra_Vserver/root_vol
```

4. Set an hourly (at 5 minutes past the hour) update schedule on each mirroring relationship.

```
snapmirror modify -source-path //Infra_Vserver/root_vol -destination-path * -schedule hourly
```

## FC Service in Clustered Data ONTAP

Create the FC service on each Vserver. This command also starts the FC service and sets the FC alias to the name of the Vserver.

```
fcv create -vserver Infra_Vserver
```

## HTTPS Access in Clustered Data ONTAP

Secure access to the storage controller must be configured.

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

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

2. Generally, a self-signed certificate is already in place. Check it with the following command:

```
security certificate show
```

3. Run the following commands as one-time commands to generate and install self-signed certificates:



**Note** You can also use the security certificate delete command to delete expired certificates

```
security certificate create -vserver Infra_Vserver -common-name
<<var_security_cert_vserver_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_clustername>> -common-name
<<var_security_cert_cluster_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_node01>> -common-name
<<var_security_cert_node01_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
security certificate create -vserver <<var_node02>> -common-name
<<var_security_cert_node02_common_name>> -size 2048 -country <<var_country_code>>
-state <<var_state>> -locality <<var_city>> -organization <<var_org>> -unit
<<var_unit>> -email <<var_storage_admin_email>>
```

4. Configure and enable SSL and HTTPS access and disable Telnet access.

```
system services web modify -external true -ssl3-enabled true
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -action allow
system services firewall policy create -policy mgmt -service http -action deny
-ip-list 0.0.0.0/0
system services firewall policy delete -policy mgmt -service telnet -action allow
system services firewall policy create -policy mgmt -service telnet -action deny
-ip-list 0.0.0.0/0
security ssl modify -vserver Infra_Vserver -certificate
<<var_security_cert_vserver_common_name>> -enabled true
y
security ssl modify -vserver <<var_clustername>> -certificate
<<var_security_cert_cluster_common_name>> -enabled true
y
security ssl modify -vserver <<var_node01>> -certificate
<<var_security_cert_node01_common_name>> -enabled true
y
security ssl modify -vserver <<var_node02>> -certificate
<<var_security_cert_node02_common_name>> -enabled true
y
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
vserver services web access create -name spi -role admin -vserver
<<var_clustername>>
vserver services web access create -name ontapi -role admin -vserver
<<var_clustername>>
```

**Note**

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

## NFSv3 in Clustered Data ONTAP

Run all commands to configure NFS on the Vserver.

1. Secure the default rule for the default export policy and create the FlexPod export policy.

```
vserver export-policy rule modify -vserver Infra_Vserver -policyname default
-ruleindex 1 -rorule never -rwrule never -superuser never
vserver export-policy create -vserver Infra_Vserver FlexPod
```

2. Create a new rule for the FlexPod export policy.

**Note**

For each ESXi host being created, create a rule. Each host will have its own rule index. Your first ESXi host will have rule index 1, your second ESXi host will have rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra_Vserver -policyname FlexPod
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys
-rwrule sys -superuser sys -allow-suid false
```

3. Assign the FlexPod export policy to the infrastructure Vserver root volume.

```
volume modify -vserver Infra_Vserver -volume root_vol -policy FlexPod
```

## FlexVol in Clustered Data ONTAP

The following information is required to create a FlexVol® volume: the volume's name and size, and the aggregate on which it will exist. Create two VMware datastore volumes, a server boot volume, and a volume to hold the OnCommand database LUN. Also, update the Vserver root volume load sharing mirrors to make the NFS mounts accessible.

```
volume create -vserver Infra_Vserver -volume infra_datastore_1 -aggregate aggr02
-size 500g -state online -policy FlexPod -junction-path /infra_datastore_1
-space-guarantee none -percent-snapshot-space 0
```

```
volume create -vserver Infra_Vserver -volume infra_swap -aggregate aggr01 -size
100g -state online -policy FlexPod -junction-path /infra_swap -space-guarantee
none -percent-snapshot-space 0 -snapshot-policy none
```

```
volume create -vserver Infra_Vserver -volume esxi_boot -aggregate aggr01 -size
100g -state online -policy default -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra_Vserver -volume OnCommandDB -aggregate aggr02 -size
200g -state online -policy default -space-guarantee none -percent-snapshot-space 0
```

```
snapmirror update-ls-set -source-path //Infra_Vserver/root_vol
```

## LUN in Clustered Data ONTAP

1. Create two boot LUNS: VM-Host-Infra-01 and VM-Host-Infra-02.

```
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-01 -size
10g -ostype vmware -space-reserve disabled
lun create -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-02 -size
10g -ostype vmware -space-reserve disabled
```

## Deduplication in Clustered Data ONTAP

Enable deduplication on appropriate volumes.

```
volume efficiency on -vserver Infra_Vserver -volume infra_datastore_1
volume efficiency on -vserver Infra_Vserver -volume esxi_boot
volume efficiency on -vserver Infra_Vserver -volume OnCommandDB
```

## Failover Groups NAS in Clustered Data ONTAP

Create an NFS port failover group.

```
network interface failover-groups create -failover-group
fg-nfs-<<var_nfs_vlan_id>> -node <<var_node01>> -port a0a-<<var_nfs_vlan_id>>
network interface failover-groups create -failover-group
fg-nfs-<<var_nfs_vlan_id>> -node <<var_node02>> -port a0a-<<var_nfs_vlan_id>>
```

## NFS LIF in Clustered Data ONTAP

Create an NFS logical interface (LIF).

```
network interface create -vserver Infra_Vserver -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>>
-address <<var_node01_nfs_lif_ip>> -netmask <<var_node01_nfs_lif_mask>>
-status-admin up -failover-policy nextavail -firewall-policy data -auto-revert
true -use-failover-group enabled -failover-group fg-nfs-<<var_nfs_vlan_id>>

network interface create -vserver Infra_Vserver -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>>
-address <<var_node02_nfs_lif_ip>> -netmask <<var_node02_nfs_lif_mask>>
-status-admin up -failover-policy nextavail -firewall-policy data -auto-revert
true -use-failover-group enabled -failover-group fg-nfs-<<var_nfs_vlan_id>>
```

## FCP LIF in Clustered Data ONTAP

Create four FCoE LIFs, two on each node.

```
network interface create -vserver Infra_Vserver -lif fcp_lif01a -role data
-data-protocol fcp -home-node <<var_node01>> -home-port 3b
network interface create -vserver Infra_Vserver -lif fcp_lif01b -role data
-data-protocol fcp -home-node <<var_node01>> -home-port 4b
network interface create -vserver Infra_Vserver -lif fcp_lif02a -role data
-data-protocol fcp -home-node <<var_node02>> -home-port 3b
network interface create -vserver Infra_Vserver -lif fcp_lif02b -role data
-data-protocol fcp -home-node <<var_node02>> -home-port 4b
```

## Add Infrastructure Vserver Administrator

Add the infrastructure Vserver administrator and Vserver administration logical interface in the out-of-band management network with the following commands:

```
network interface create -vserver Infra_Vserver -lif vsmgmt -role data
-data-protocol none -home-node <<var_node02>> -home-port e0a -address
<<var_vserver_mgmt_ip>> -netmask <<var_vserver_mgmt_mask>> -status-admin up
-failover-policy nextavail -firewall-policy mgmt -auto-revert true
-use-failover-group enabled -failover-group fg-cluster-mgmt

network routing-groups route create -vserver Infra_Vserver -routing-group
d<<var_clustermgmt_ip>> -destination 0.0.0.0/0 -gateway
<<var_clustermgmt_gateway>>
security login password -username vsadmin -vserver Infra_Vserver
Please enter a new password: <<var_vsadmin_password>>
Please enter it again: <<var_vsadmin_password>>

security login unlock -username vsadmin -vserver Infra_Vserver
```

## 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. These 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 chosen 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>>
Enter the same password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)?
(yes/no) [n]: y
Which switch fabric (A|B): A
Enter the system name: <<var_ucs_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_cluster_ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
```

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. Do you want to 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 Clustered Data ONTAP

### 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 **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(1b)

This document assumes the use of Cisco UCS 2.1(1b). To upgrade the Cisco UCS Manager software and the UCS 6248 Fabric Interconnect software to version 2.1(1b), see Cisco UCS Manager Install and Upgrade Guides at:

[http://www.cisco.com/en/US/docs/unified\\_computing/ucs/sw/upgrading/from2.0/to2.1/b\\_UpgradingCiscoUCSFrom2.0To2.1.html](http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/upgrading/from2.0/to2.1/b_UpgradingCiscoUCSFrom2.0To2.1.html)

### 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**.

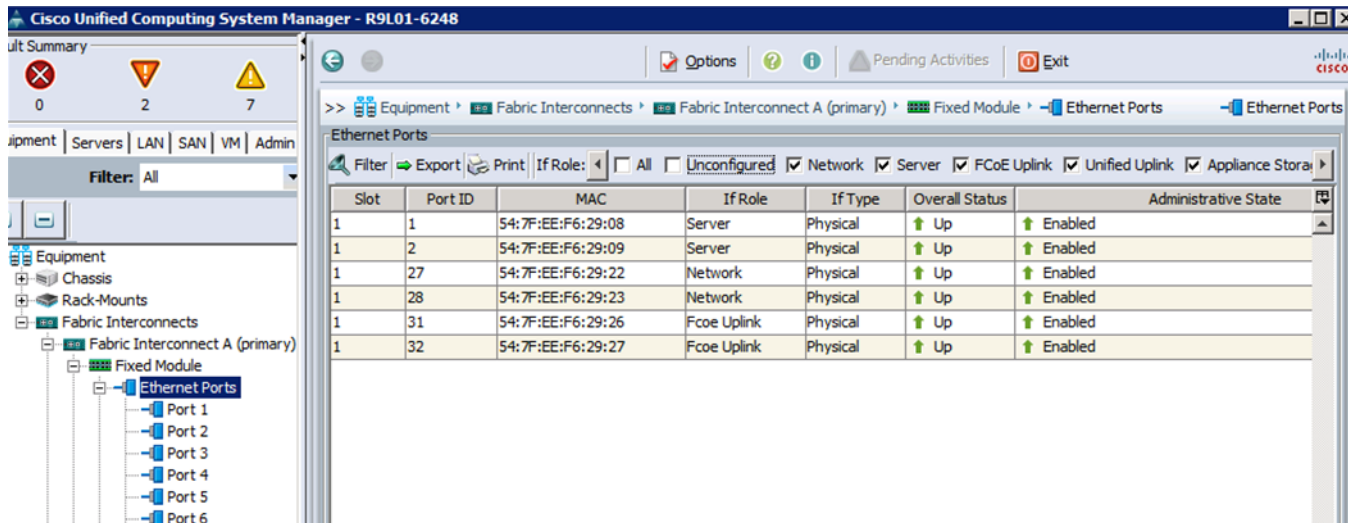
## Enable Server and Uplink Ports

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**.

- Verify that the ports connected to the chassis or to the Cisco 2232 FEX are now configured as server ports.

**Figure 4 UCS - Port Configuration Example**



- Choose ports 27 and 28 that are connected to the Cisco Nexus switches, right-click them, and choose **Configure as Uplink Port**.



**Note** The UCS ports that are connected to the Cisco Nexus switches for Ethernet traffic should be connected to switch ports that are allocated to the Cisco Nexus 7000 switch VDC.

- Click **Yes** to confirm uplink ports and click **OK**.
- Choose ports 31 and 32, which will serve as FCoE uplinks to the Cisco Nexus switches; right-click them; and choose **Configure as FCoE Uplink Port**.



**Note** The UCS ports that are connected to the Cisco Nexus switches for FCoE traffic should be connected to switch ports that are allocated to the Cisco Nexus 7000 storage VDC.

- Click **Yes** to confirm FCoE uplink ports and click **OK**.
- Choose **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
- Expand Ethernet Ports.
- 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**.
- Click **Yes** to confirm server ports and click **OK**.
- Choose ports 27 and 28 that are connected to the Cisco Nexus switches, right-click them, and choose **Configure as Uplink Port**.

**Note**

The UCS ports that are connected to the Cisco Nexus switches for Ethernet traffic should be connected to switch ports that are allocated to the Cisco Nexus 7000 switch VDC.

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 switches, right-click them, and choose **Configure as FCoE Uplink Port**.

**Note**

The UCS ports that are connected to the Cisco Nexus switches for FCoE traffic should be connected to switch ports that are allocated to the Cisco Nexus 7000 storage VDC.

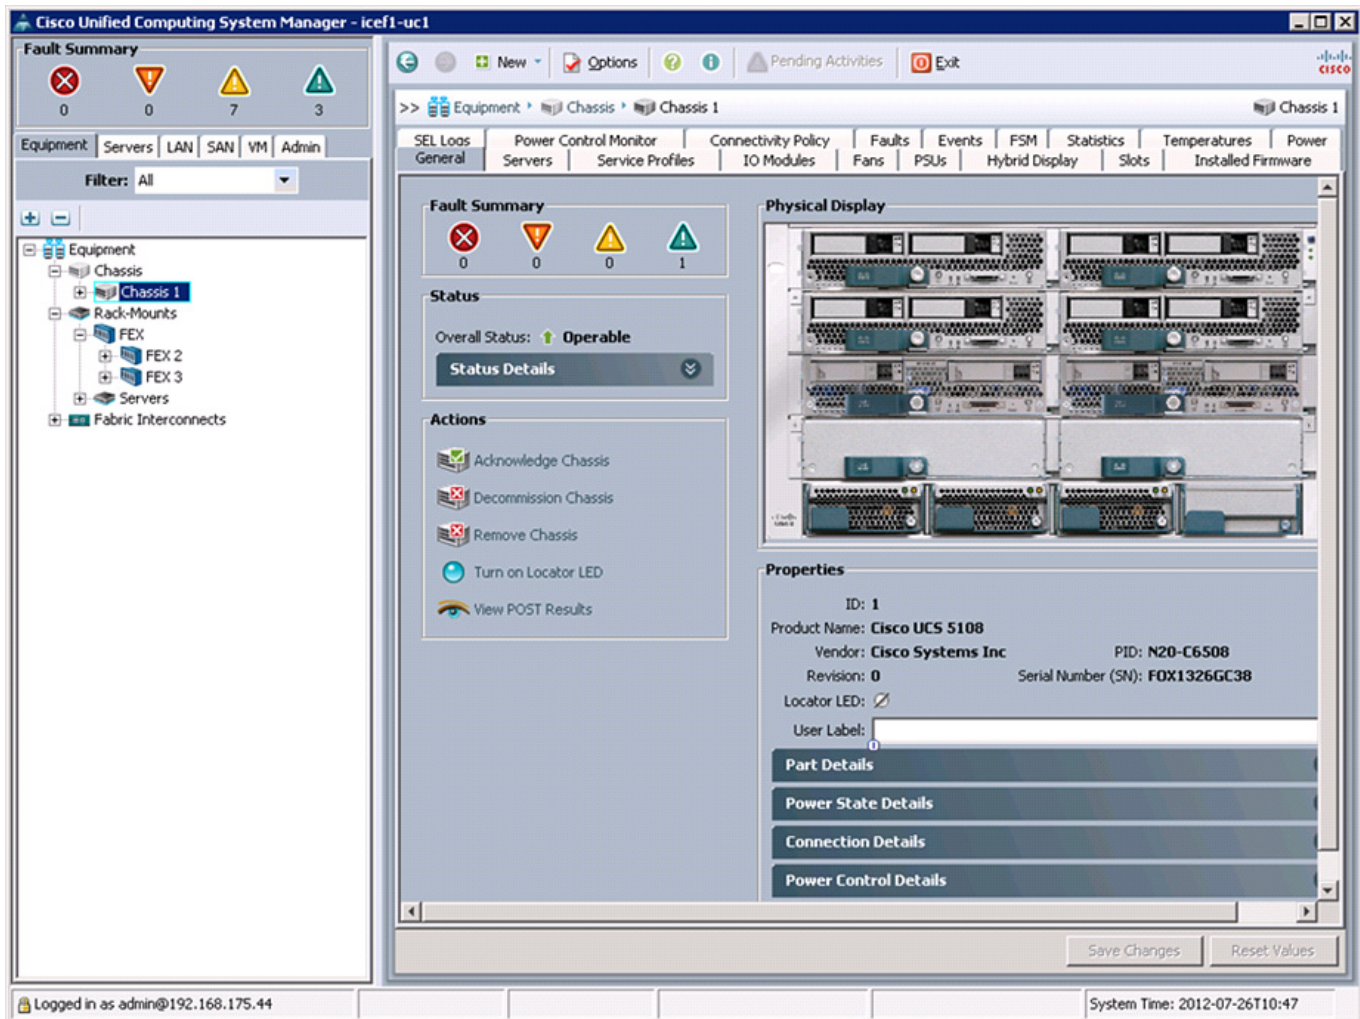
18. Click **Yes** to confirm FCoE uplink ports and click **OK**.

## Acknowledge Cisco UCS Chassis and FEX

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

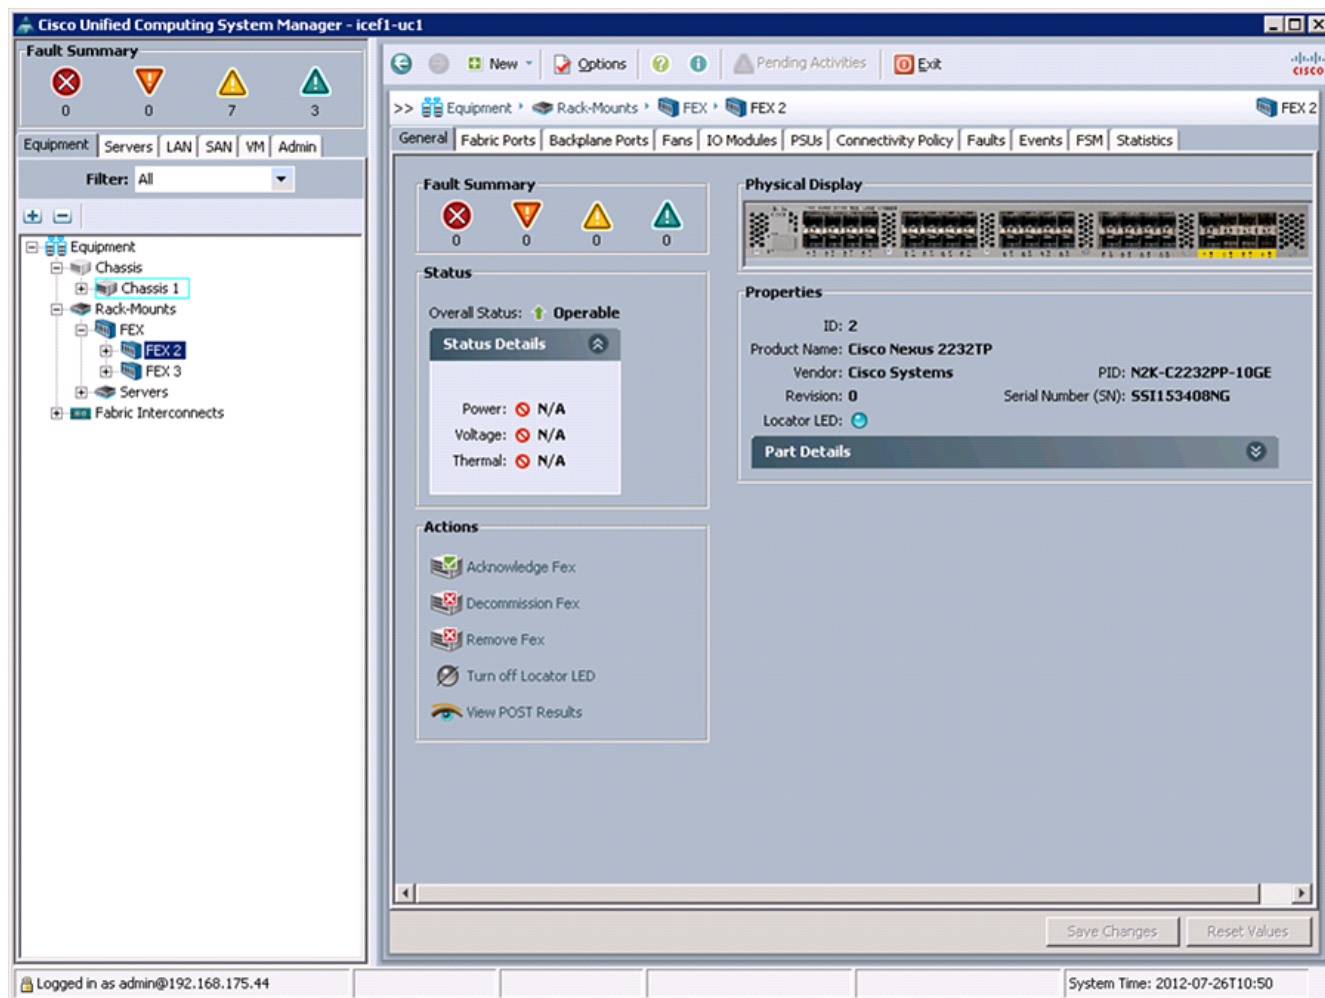
1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
2. Expand Chassis.
3. Choose each chassis that is listed, right-click on each chassis and choose **Acknowledge Chassis**.

**Figure 5 UCS - Chassis Overview**



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 6 UCS - FEX Management Using UCSM**



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

## Create Uplink Port Channels to Cisco Nexus Switches

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

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

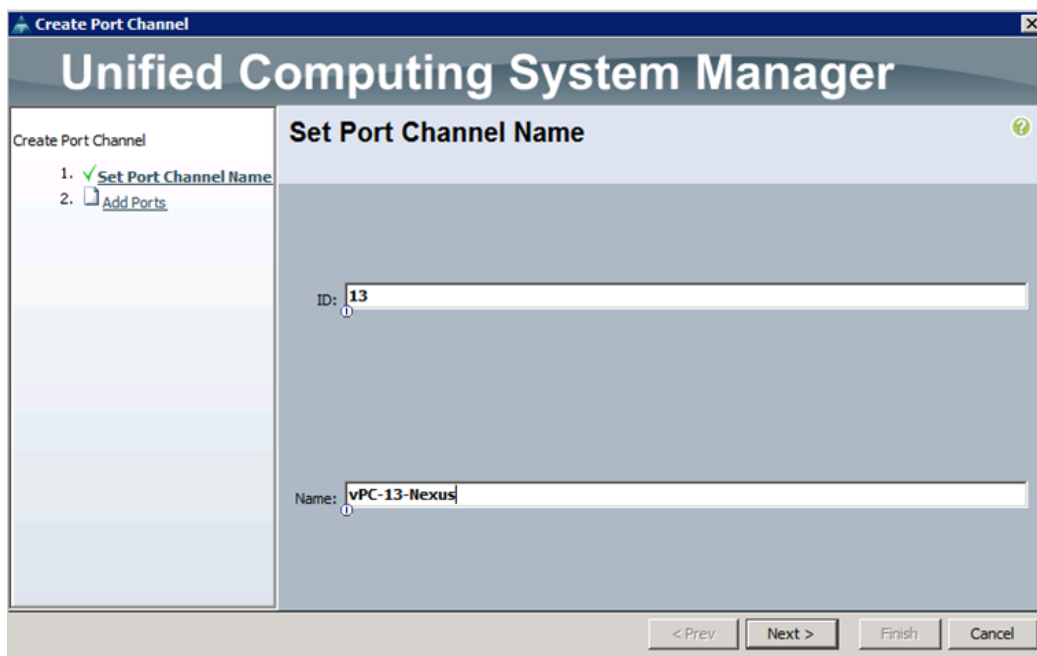


**Note** In this procedure, two port channels are created: one from fabric A to both Cisco Nexus switches and one from fabric B to both Cisco Nexus switches.

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

7. Click **Next**.

**Figure 7 UCS - Port Channel Wizard**



8. Choose the following ports to be added to the port channel:
  - Slot ID 1 and port 27
  - Slot ID 1 and port 28
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-Nexus 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 27
  - Slot ID 1 and port 28
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. Click **Next**.
8. Click **Add**.
9. Specify a starting MAC address.



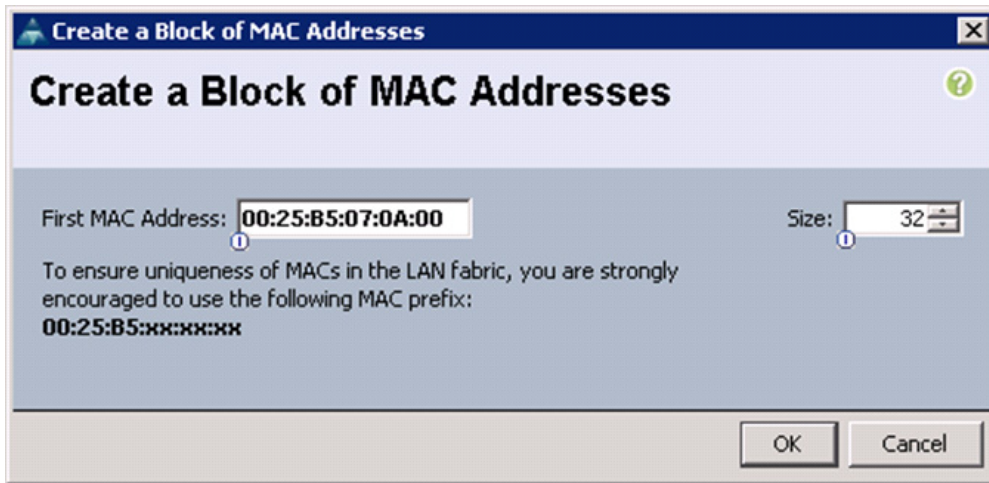
### Note

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.

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



**Figure 8** UCS - Create MAC Address Pool for Fabric A



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



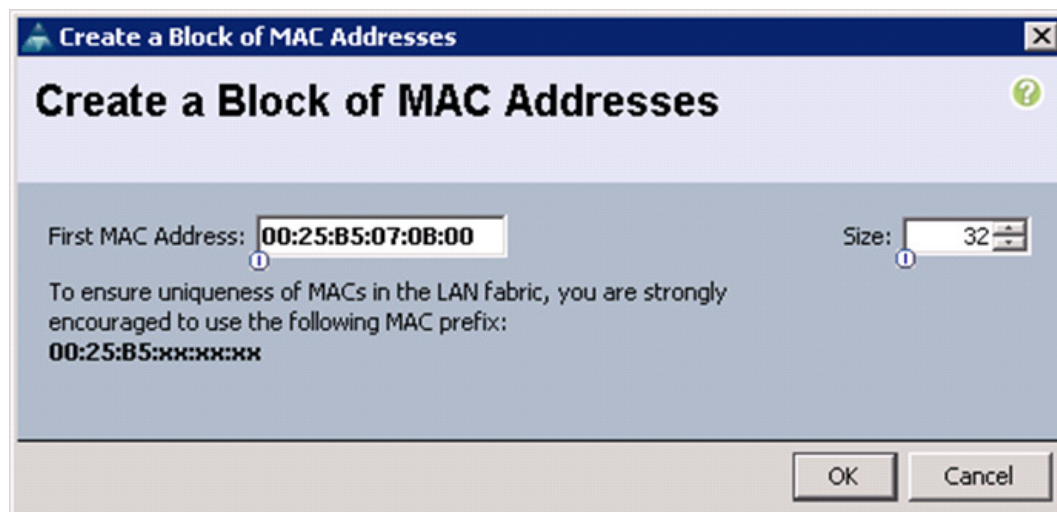
**Note**

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.

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



**Figure 9 UCS - Create MAC Address Pool for Fabric B**



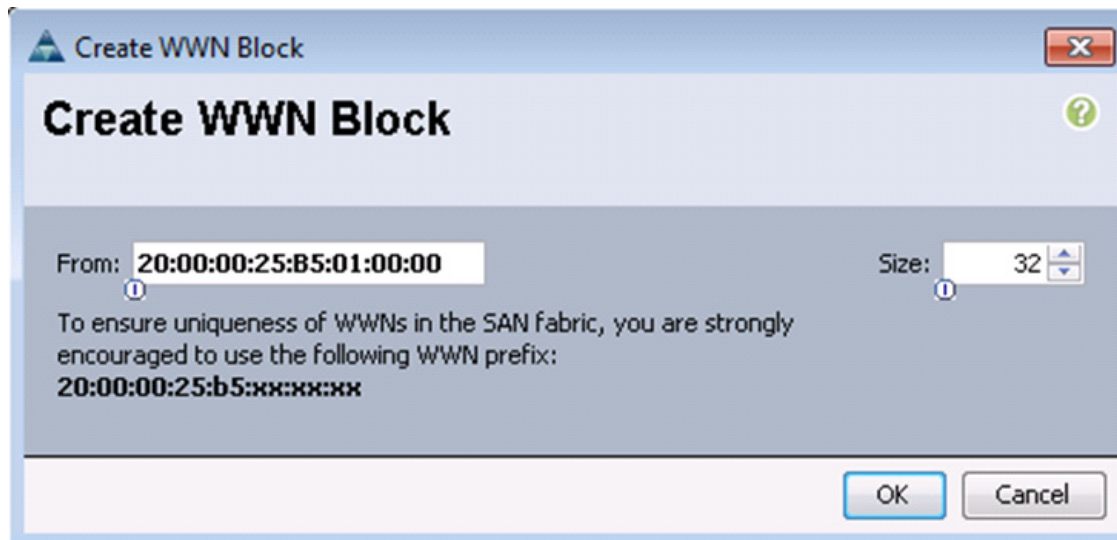
22. Click **OK**.
23. Click **Finish**.
24. In the confirmation message window, click **OK**.

## 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. Click **Next**.
8. Click **Add** to add a block of WWNNs.
9. Keep the default block of WWNNs, or specify a base WWNN.
10. Specify a size for the WWNN block that is sufficient to support the available blade or server resources.

Figure 10 UCS - Create WWNN Pool



11. Click **OK**.
12. Click **Finish**.
13. 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. Click **Next**.
8. Click **Add** to add a block of WWPNs.
9. 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.

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

**Figure 11 UCS - Create WWPN Pool**



11. Click **OK**.
12. Click **Finish** to create the WWPN pool.
13. Click **OK**.
14. Right-click WWPN Pools.
15. Choose Create WWPN Pool.
16. Enter WWPN\_Pool\_B as the name for the WWPN pool for fabric B.
17. (Optional) Enter a description for this WWPN pool.
18. Click **Next**.
19. Click **Add** to add a block of WWPNs.
20. 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.

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

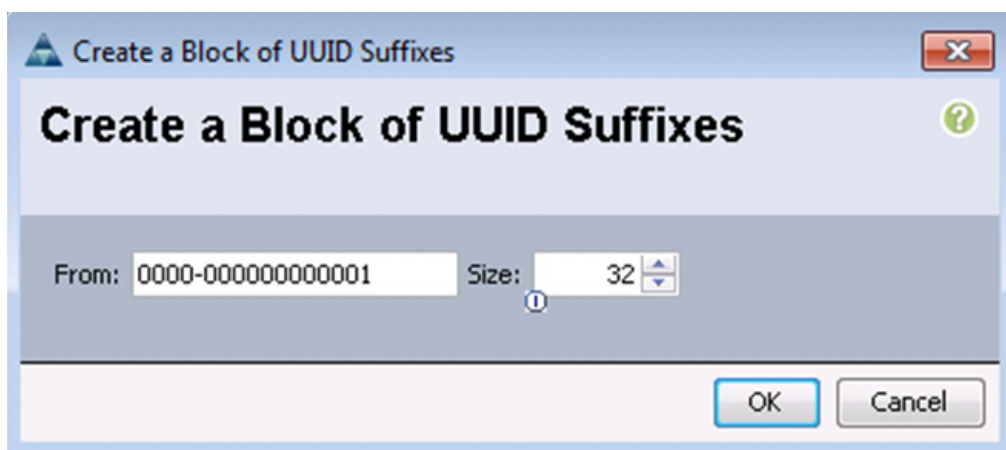
## Create UUID Suffix Pool

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. Click **Next**.
9. Click **Add** to add a block of UUIDs.
10. Keep the From field at the default setting.
11. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

**Figure 12 UCS - Create UUID Block**



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

## Create Server Pool

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



### 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 13** UCS - Create In-Band Management VLAN

**Create VLANs**

VLAN Name/Prefix:

Multicast Policy Name:  + Create Multicast Policy

☒ Common/Global
 ☐ Fabric A
 ☐ Fabric B
 ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type:
 ☒ None
 ☐ Primary
 ☐ Isolated

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 14 UCS - Create NFS VLAN

**Create VLANs**

VLAN Name/Prefix:

Multicast Policy Name:  + Create Multicast Policy

☒ Common/Global
 ☐ Fabric A
 ☐ Fabric B
 ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type:
 ☒ None
 ☐ Primary
 ☐ Isolated

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 15 UCS - Create vMotion VLAN**

**Create VLANs**

VLAN Name/Prefix:

Multicast Policy Name:  [+ Create Multicast Policy](#)

☒ Common/Global
 ☐ Fabric A
 ☐ Fabric B
 ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs. (e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type:
 ☒ None
 ☐ Primary
 ☐ Isolated

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 16 UCS - Create VM Traffic VLAN

**Create VLANs**

VLAN Name/Prefix: **VM-Traffic-VLAN**

Multicast Policy Name: **<not set>** + Create Multicast Policy

☒ Common/Global
 ☐ Fabric A
 ☐ Fabric B
 ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: **3174**

Sharing Type:
 ☒ None
 ☐ Primary
 ☐ Isolated

Check Overlap
OK
Cancel

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 17** UCS - Create Native VLAN

**Create VLANs**

VLAN Name/Prefix:

Multicast Policy Name:  [Create Multicast Policy](#)

☒ Common/Global
 ☐ Fabric A
 ☐ Fabric B
 ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type: ☒ None ☐ Primary ☐ Isolated

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 18 UCS - Create VSAN for Fabric A

**Create VSAN**

Name:

**FC Zoning Settings**

FC Zoning: ☒ Disabled ☐ Enabled

Do **NOT** enable zoning for this VSAN if the fabric interconnect is connected to an upstream switch that has zoning enabled on the same VSAN.

☐ Common/Global ☒ Fabric A ☐ Fabric B ☐ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.  
Enter the VSAN ID that maps to this VSAN.

VSAN ID:

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.  
Enter the VLAN ID that maps to this VSAN.

FCoE VLAN:

OK Cancel

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 19 UCS - Create VSAN for Fabric B**

**Create VSAN**

Name: **VSAN\_B**

**FC Zoning Settings**

FC Zoning: ☒ Disabled ☐ Enabled

Do **NOT** enable zoning for this VSAN if the fabric interconnect is connected to an upstream switch that has zoning enabled on the same VSAN.

☐ Common/Global ☐ Fabric A ☒ Fabric B ☐ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.  
Enter the VSAN ID that maps to this VSAN.

VSAN ID: **102**

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.  
Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: **102**

OK Cancel

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(1e) for both the Blade and Rack Packages.
8. Click **OK** to create the host firmware package.
9. Click **OK**.

**Figure 20** UCS - Create Host Firmware Package

**Create Host Firmware Package**

Name:

Description:

How would you like to configure the Host Firmware Package? ☒ Simple ☐ Advanced

Blade Package:

Rack Package:

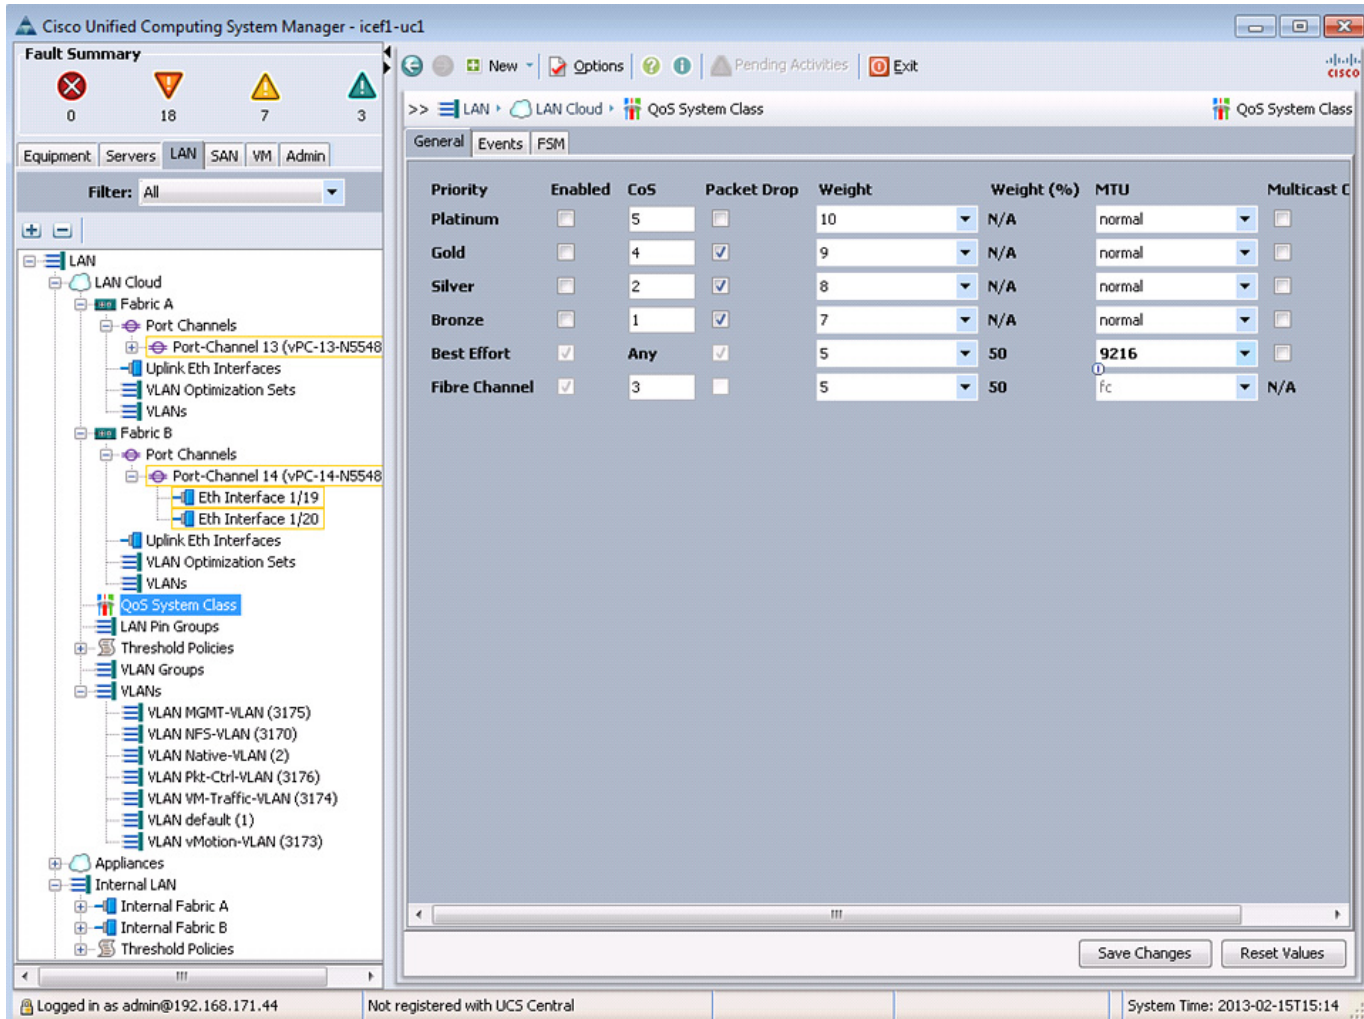
OK Cancel

## 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 21 UCS - Setting Jumbo Frames



## 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.



- Click **OK** to create the local disk configuration policy.

**Figure 22** UCS - Create Local Disk Policy

**Create Local Disk Configuration Policy**

Name:

Description:

Mode:

OK Cancel

- 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:

- In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
- Choose **Policies > root**.
- Right-click Network Control Policies.
- 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.
8. Click **OK**.

**Figure 23** UCS - Create Network Control Policy

**Create Network Control Policy**

Name:

CDP: ☐ Disabled ☒ Enabled

MAC Register Mode: ☒ Only Native Vlan ☐ All Host Vlans

Action on Uplink Fail: ☒ Link Down ☐ Warning

**MAC Security**

Forge: ☒ Allow ☐ Deny

OK Cancel

## 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 24 UCS - Create Power Control Policy**

**Create Power Control Policy**

Name: **No-Power-Cap**

Description:

**Power Capping**

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

☒ No Cap ☐ cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK Cancel

## Create Server Pool Qualification Policy (Optional)

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

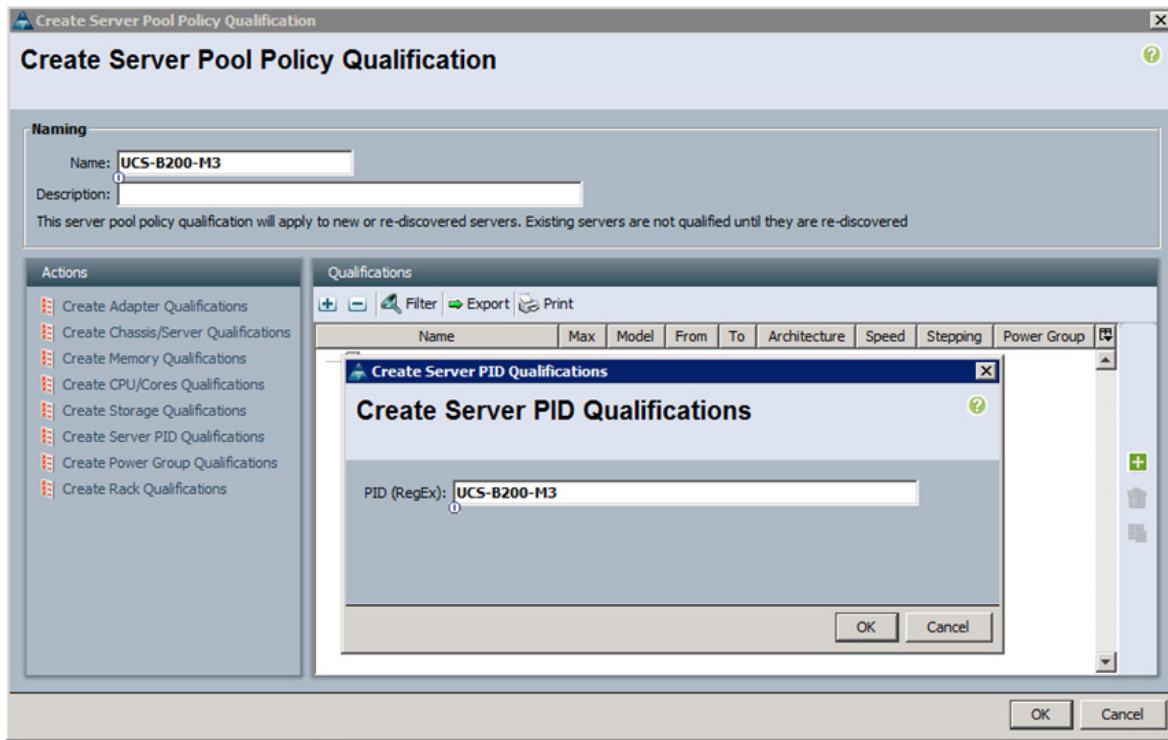


### Note

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 25 UCS - Create 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.
8. Click **OK**.

Figure 26 UCS - Create BIOS Policy

**Create BIOS Policy**

**Unified Computing System Manager**

Create BIOS Policy

1. **Main**
2. Processor
3. Intel Directed IO
4. RAS Memory
5. Serial Port
6. USB
7. PCI Configuration
8. Boot Options
9. Server Management

**Main**

Name:

Reboot on BIOS Settings Change: ☐

Quiet Boot: ☒ disabled ☐ enabled ☐ Platform Default

Post Error Pause: ☐ disabled ☐ enabled ☒ Platform Default

Resume Ac On Power Loss: ☐ stay-off ☐ last-state ☐ reset ☒ Platform Default

Front Panel Lockout: ☐ disabled ☐ enabled ☒ Platform Default

< Prev   Next >   Finish   Cancel

## 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 choose Assigned Only.
7. Click **OK**, and then click **OK** again.

Figure 27 UCS - Create vNIC/vHBA Placement Policy

**Create Placement Policy**

Name: **VM-Host-Infra**

Virtual Slot Mapping Scheme: ☒ Round Robin ☐ Linear Ordered

Filter Export Print

Virtual Slot	Selection Preference
1	<b>Assigned Only</b>
2	All
3	All
4	All

OK Cancel


## 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 28** UCS - Update Default Server Reboot 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 29** UCS - Create vNIC Template for Fabric A

**Create vNIC Template**

Name:

Description:

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

**Target**

☒ Adapter  
☐ VM

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

Template Type: ☐ Initial Template ☒ Updating Template

**VLANs**

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	IB-MGMT-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>

**Create VLAN**

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

**OK** **Cancel**

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 30** UCS - Create vNIC Template for Fabric B

**Create vNIC Template**

Name:

Description:

Fabric ID: ☐ Fabric A ☒ Fabric B ☐ Enable Failover

**Target**

☒ Adapter  
☐ VM

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

Template Type: ☐ Initial Template ☒ Updating Template

**VLANs**

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	IB-MGMT-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>

[+ Create VLAN](#)

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

OK Cancel

## 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 31 UCS - Create vHBA Template for Fabric A**

**Create vHBA Template**

Name:

Description:

Fabric ID: ☒ A ☐ B

Select VSAN:  + Create VSAN

Template Type: ☐ Initial Template ☒ Updating Template

Max Data Field Size:

WWPN Pool:

QoS Policy:

Pin Group:

Stats Threshold Policy:

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 32 UCS - Create vHBA Template for Fabric B

**Create vHBA Template**

Name:

Description:

Fabric ID: ☐ A ☒ B

Select VSAN:  + Create VSAN

Template Type: ☐ Initial Template ☒ Updating Template

Max Data Field Size:

WWPN Pool:

QoS Policy:

Pin Group:

Stats Threshold Policy:

## Create Boot Policies

This procedure applies to a Cisco UCS environment in which two FCoE logical interfaces (LIFs) are on cluster node 1 (fcp\_lif01a and fcp\_lif01b) and two FCoE LIFs are on cluster node 2 (fcp\_lif02a and fcp\_lif02b). Also, it is assumed that the A LIFs are connected to fabric A (Cisco Nexus Switch A) and the B LIFs are connected to fabric B (Cisco Nexus Switch B).

Two boot policies are configured in this procedure. The first policy configures the primary target to be fcp\_lif01a and the second boot policy configures the primary target to be fcp\_lif01b.

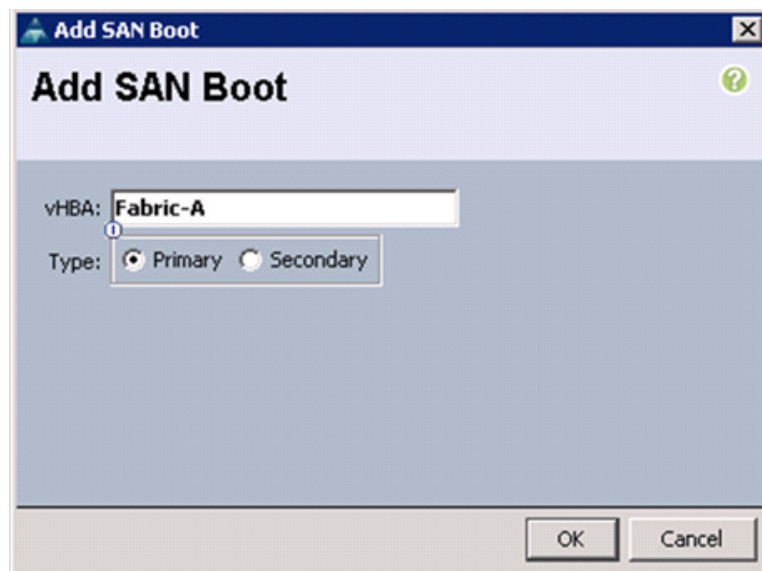
### Creating First Boot Policy

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 33 UCS - Setting Fabric-A as Primary in the 1st Boot Policy**



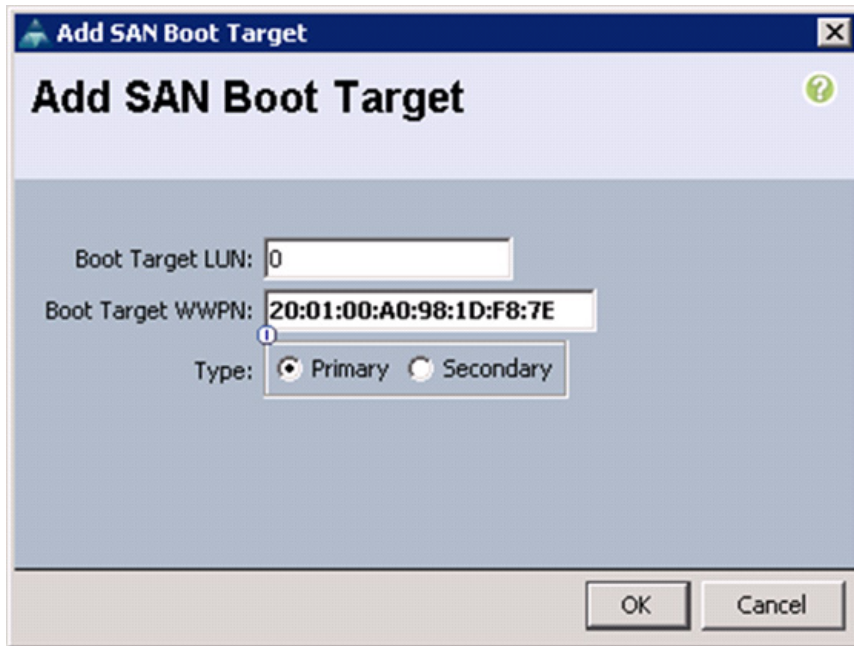
13. From the vHBA drop-down menu, choose Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for fcp\_lif01a.



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

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

Figure 34 UCS - Adding FCP\_LIF01a WWPN as Primary SAN Boot Target



The image shows a Windows-style dialog box titled "Add SAN Boot Target". It has a blue header bar with a small icon on the left and a close button (X) on the right. Below the header, the title "Add SAN Boot Target" is displayed in a large, bold, black font. To the right of the title is a green question mark icon. The main area of the dialog is light blue and contains three input fields. The first field is labeled "Boot Target LUN:" and contains the value "0". The second field is labeled "Boot Target WWPN:" and contains the value "20:01:00:A0:98:1D:F8:7E". The third field is labeled "Type:" and contains two radio buttons: "Primary" (which is selected) and "Secondary". At the bottom of the dialog, there are two buttons: "OK" and "Cancel".

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 fcp\_lif02a.

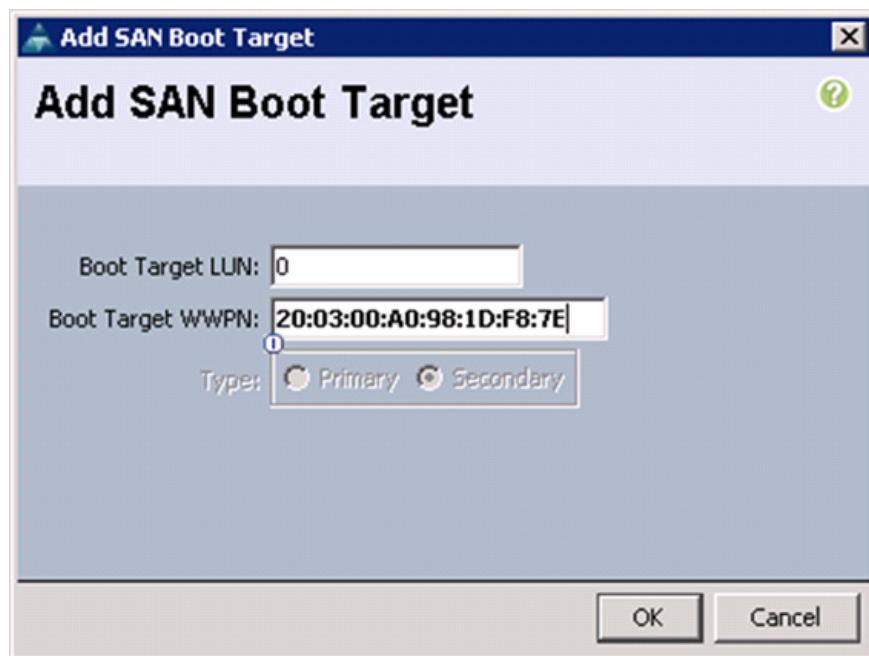


**Note**

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

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

**Figure 35 UCS - Adding FCP\_LIF02a WWPN as Secondary SAN Boot Target**



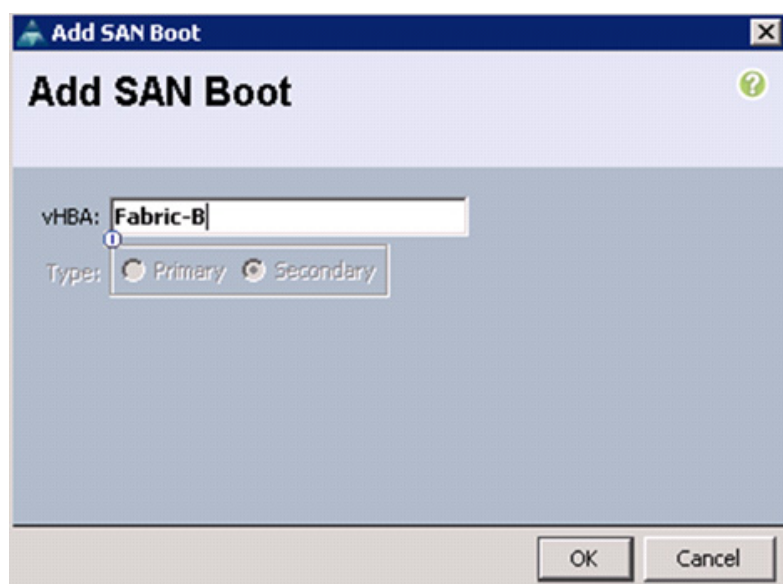
The dialog box titled "Add SAN Boot Target" contains the following fields and options:

- Boot Target LUN:** A text box containing the value "0".
- Boot Target WWPN:** A text box containing the value "20:03:00:A0:98:1D:F8:7E".
- Type:** A group box containing two radio buttons: "Primary" (unselected) and "Secondary" (selected).

At the bottom of the dialog are "OK" and "Cancel" buttons.

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 36 UCS - Setting Fabric-B as Secondary in the 1st Boot Policy**



The dialog box titled "Add SAN Boot" contains the following fields and options:

- vHBA:** A text box containing the value "Fabric-B".
- Type:** A group box containing two radio buttons: "Primary" (unselected) and "Secondary" (selected).

At the bottom of the dialog are "OK" and "Cancel" buttons.

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 fcp\_lif01b.



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

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

**Figure 37** UCS - Adding FCP\_LIF01b WWPN as Primary SAN Boot Target

The screenshot shows a Windows-style dialog box titled "Add SAN Boot Target". Inside the dialog, there are three input fields: "Boot Target LUN:" containing the number "0", "Boot Target WWPN:" containing the hexadecimal string "20:02:00:A0:98:1D:F8:7E", and "Type:" with two radio button options, "Primary" (which is selected) and "Secondary". At the bottom right of the dialog are two buttons, "OK" and "Cancel".

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 fcp\_lif02b.

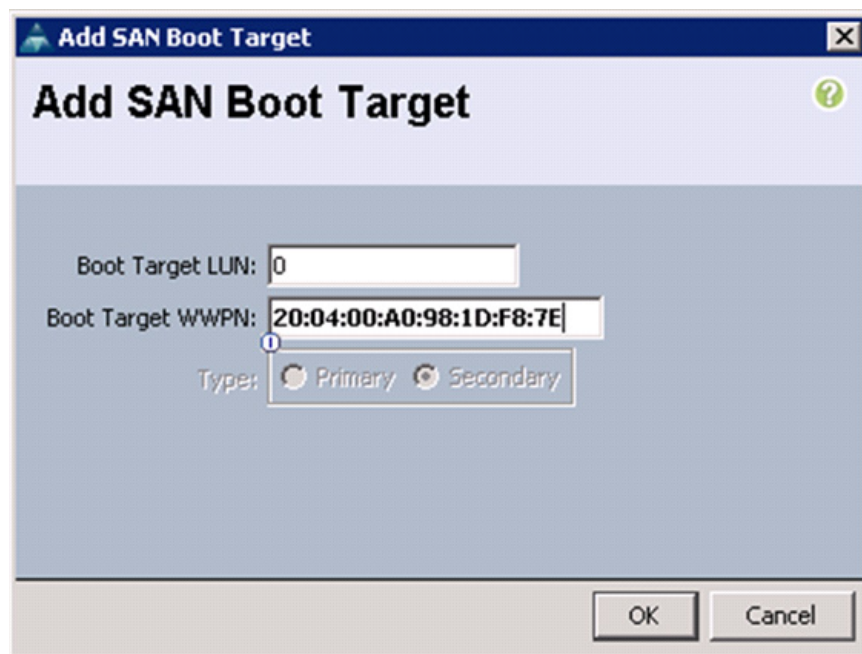


**Note** To obtain this information, log in to the storage cluster and run the **network interface show** command.

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



**Figure 38 UCS - Adding FCP\_LIF02b WWPN as Secondary SAN Boot Target**




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

### Creating Second Boot Policy

1. Right-click Boot Policies again.
2. Choose **Create Boot Policy**.
3. Enter Boot-Fabric-B as the name of the boot policy.
4. (Optional) Enter a description of the boot policy.
5. Keep the Reboot on Boot Order Change check box unchecked.
6. From the Local Devices drop-down menu choose Add CD-ROM.
7. From the vHBA drop-down menu choose Add SAN Boot.
8. In the Add SAN Boot dialog box, enter Fabric-B in the vHBA box.
9. Make sure that the Primary radio button is selected as the SAN boot type.
10. Click **OK** to add the SAN boot initiator.

**Figure 39** UCS - Setting Fabric-B as Primary in the 2st Boot Policy



11. From the vHBA drop-down menu, choose Add SAN Boot Target.
12. Keep 0 as the value for Boot Target LUN.
13. Enter the WWPN fcp\_lif01b.



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

14. Keep Primary as the SAN boot target type.
15. Click **OK** to add the SAN boot target.

**Figure 40** UCS - Adding FCP\_LIF01b WWPN as Primary SAN Boot Target

**Add SAN Boot Target**

Boot Target LUN: 0

Boot Target WWPN: 20:02:00:A0:98:1D:F8:7E

Type: ☒ Primary ☐ Secondary

OK Cancel

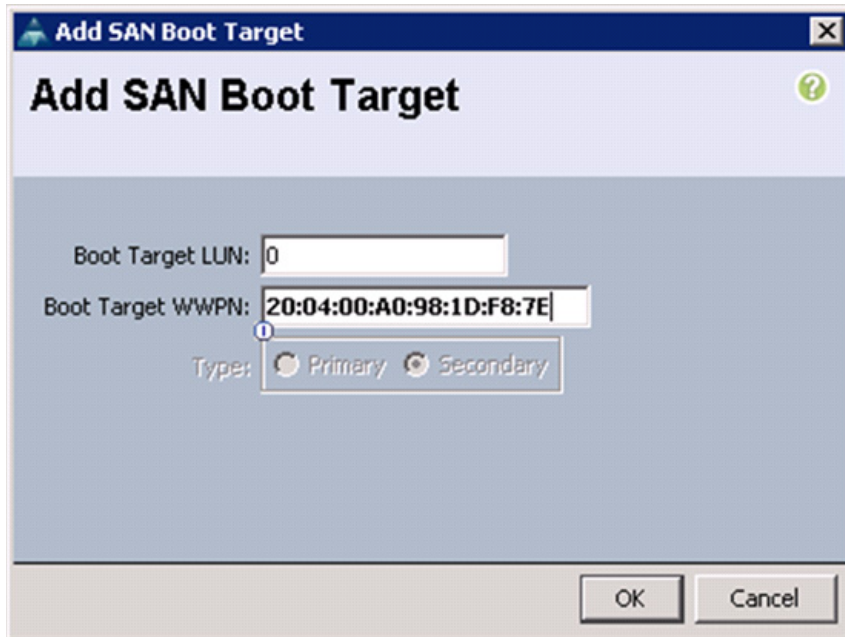
16. From the vHBA drop-down menu, choose Add SAN Boot Target.
17. Keep 0 as the value for Boot Target LUN.
18. Enter the WWPN for fcp\_lif02b.



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

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

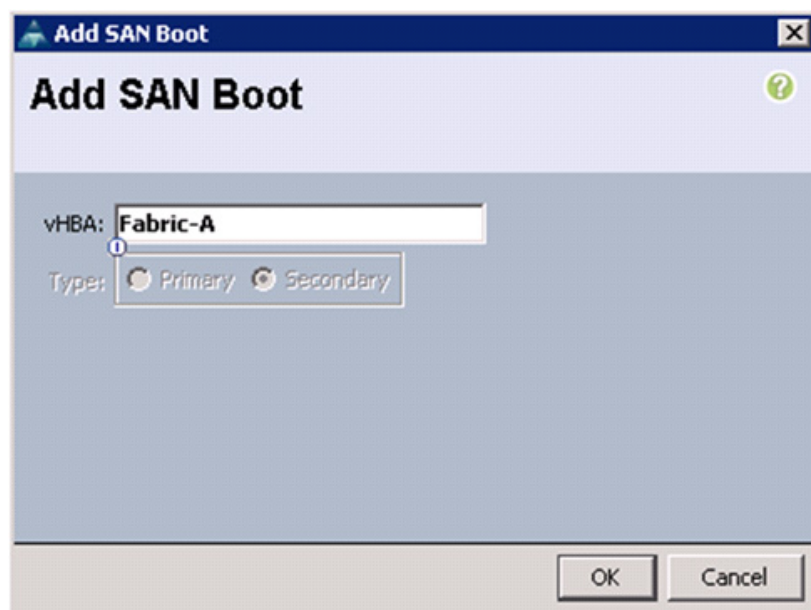
**Figure 41** UCS - Adding FCP\_LIF02b WWPN as Secondary SAN Boot Target



The dialog box titled "Add SAN Boot Target" has a close button (X) in the top right corner. Below the title bar, there is a question mark icon. The main area contains three input fields: "Boot Target LUN:" with the value "0", "Boot Target WWPN:" with the value "20:04:00:A0:98:1D:F8:7E", and "Type:" with two radio buttons, "Primary" and "Secondary". The "Secondary" radio button is selected. At the bottom right, there are "OK" and "Cancel" buttons.

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

**Figure 42** UCS - Setting Fabric-A as Secondary in the 2st Boot Policy



The dialog box titled "Add SAN Boot" has a close button (X) in the top right corner. Below the title bar, there is a question mark icon. The main area contains two input fields: "vHBA:" with the value "Fabric-A", and "Type:" with two radio buttons, "Primary" and "Secondary". The "Secondary" radio button is selected. At the bottom right, there are "OK" and "Cancel" buttons.

24. From the vHBA menu, choose Add SAN Boot Target.
25. Keep 0 as the value for Boot Target LUN.
26. Enter the WWPN for fcp\_lif01a.



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

27. Keep Primary as the SAN boot target type.
28. Click **OK** to add the SAN boot target.

**Figure 43** UCS - Adding FCP\_LIF01a WWPN as Primary SAN Boot Target

29. From the vHBA drop-down menu, choose Add SAN Boot Target.
30. Keep 0 as the value for Boot Target LUN.
31. Enter the WWPN for fcp\_lif02a.



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

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

Figure 44 UCS - Adding FCP\_LIF02a WWPN as Secondary SAN Boot Target

**Add SAN Boot Target**

Boot Target LUN: 0

Boot Target WWPN: 20:03:00:A0:98:1D:F8:7E

Type: ☐ Primary ☒ Secondary

OK Cancel

33. 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 45 UCS - Create Service Profile Template**

**Create Service Profile Template**

**Unified Computing System Manager**

Create Service Profile Template

1. **Identify Service Profile Template**
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

**Identify Service Profile Template**

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name:

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

Where: **org-root**

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

Type: ☐ Initial Template ☒ Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

**UUID**

UUID Assignment:

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

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev   Next >   Finish   Cancel

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 46 UCS - Create vNIC Using vNIC Template

**Create vNIC**

Name:

Use vNIC Template: ☒

vNIC Template:

**Adapter Performance Profile**

Adapter Policy:

- 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 47 UCS - Validation of vNIC Creation

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The 'Networking' step is active, showing options for LAN configuration. The left sidebar lists the steps: 1. Identify Service Profile Template, 2. Networking, 3. Storage, 4. Zoning, 5. vNIC/vHBA Placement, 6. Server Boot Order, 7. Maintenance Policy, 8. Server Assignment, and 9. Operational Policies. The main area is titled 'Networking' and includes a sub-header 'Optionally specify LAN configuration information.' Below this, there is a dropdown for 'Dynamic vNIC Connection Policy' set to 'Select a Policy to use (no Dynamic vNIC Policy by default)' and a '+ Create Dynamic vNIC Connection Policy' button. A section titled 'How would you like to configure LAN connectivity?' has four radio buttons: Simple, Expert (selected), No vNICs, and Use Connectivity Policy. Below this, a note says 'Click Add to specify one or more vNICs that the server should use to connect to the LAN.' A table lists two vNICs: vNIC vNIC-A and vNIC vNIC-B, both with 'Derived' MAC and Fabric IDs. Below the table are 'Delete', '+ Add', and 'Modify' buttons. Another section says 'Click Add to specify one or more iSCSI vNICs that the server should use.' Below this is an empty table with columns: Name, Overlay vNIC Name, iSCSI Adapter Policy, and MAC Address, with '+ Add', 'Delete', and 'Modify' buttons at the bottom. At the very bottom are '< Prev', 'Next >', 'Finish', and 'Cancel' buttons.

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 48 UCS - Create vHBA Using vHBA Template**

**Create vHBA**

Name:

Use vHBA Template: ☒

+ Create vHBA Template

vHBA Template:

**Adapter Performance Profile**

Adapter Policy:  + Create Fibre Channel Adapter Policy

OK Cancel

- 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 49 UCS - Validation of the vHBA Creation

**Create Service Profile Template**

**Unified Computing System Manager**

Create Service Profile Template

1. ☒ Identify Service Profile Template
2. ☒ Networking
3. ☒ Storage
4. ☐ Zoning
5. ☐ vNIC/vHBA Placement
6. ☐ Server Boot Order
7. ☐ Maintenance Policy
8. ☐ Server Assignment
9. ☐ Operational Policies

**Storage**

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage:  Mode: **No Local Storage**

☒ Create Local Disk Configuration Policy

Protect Configuration: **Yes**

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

**How would you like to configure SAN connectivity?** ☐ Simple ☒ Expert ☐ No vHBAs ☐ Use Connectivity Policy

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.

**World Wide Node Name**

WWNN Assignment:

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

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

< Prev Next > Finish Cancel

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 50 UCS - vNIC and vHBA Placement

**Unified Computing System Manager**

Create Service Profile Template

1. ☒ Identify Service Profile Template  
 2. ☒ Networking  
 3. ☒ Storage  
 4. ☒ Zoning  
 5. ☒ **vNIC/vHBA Placement**  
 6. ☒ Server Boot Order  
 7. ☐ Maintenance Policy  
 8. ☐ Server Assignment  
 9. ☐ Operational Policies

### vNIC/vHBA Placement

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement:

Virtual Network Interface connection provides a mechanism of placing vNICs and vHBAs on physical network adapters. vNICs and vHBAs are assigned to one of Virtual Network Interface connection specified below. This assignment can be performed explicitly by selecting which Virtual Network Interface connection is used by vNIC or vHBA or it can be done automatically by selecting "any".  
 vNIC/vHBA placement on physical network interface is controlled by placement preferences.

Please select one Virtual Network Interface and one or more vNICs or vHBAs

vNICs vHBAs

Name

>> assign >>

<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Selection Preference
vCon 1		Assigned Only
vHBA Fabric-A	1	
vHBA Fabric-B	2	
vNIC vNIC-A	3	
vNIC vNIC-B	4	
vCon 2		All
vCon 3		All
vCon 4		All

▲ Move Up ▼ Move Down

< Prev Next > Finish Cancel

10. Set the Server Boot Order:

- In the Boot Policy list, choose Boot-Fabric-A.
- Review the table to verify that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
- Click **Next**.

Figure 51 UCS - Setting Server Boot Order

**Unified Computing System Manager**

Create Service Profile Template

1. ☒ Identify Service Profile Template  
 2. ☒ Networking  
 3. ☒ Storage  
 4. ☒ Zoning  
 5. ☒ vNIC/vHBA Placement  
 6. ☒ **Server Boot Order**  
 7. ☐ Maintenance Policy  
 8. ☐ Server Assignment  
 9. ☐ Operational Policies

### Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **Boot-Fabric-A** + Create Boot Policy

Name: **Boot-Fabric-A**

Description:

Reboot on Boot Order Change: **No**

Enforce vNIC/vHBA/SCSI Name: **Yes**

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/SCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/SCSI with the lowest PCIe bus scan order is used.

Name	Order	vNIC/vHBA/SCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		Fabric-A	Primary		
SAN Target primary			Primary	0	20:06:00:A0:98:37:78:08
SAN Target secondary			Secondary	0	20:08:00:A0:98:37:78:08
SAN secondary		Fabric-B	Secondary		
SAN Target primary			Primary	0	20:07:00:A0:98:37:78:08
SAN Target secondary			Secondary	0	20:09:00:A0:98:37:78:08

Create iSCSI vNIC    Set iSCSI Boot Parameters

< Prev    Next >    Finish    Cancel

11. Add a Maintenance Policy:
  - a. Choose the Default Maintenance Policy.
  - b. Click **Next**.



**Figure 52 UCS - Choosing a Maintenance Policy**

**Create Service Profile Template**

## Unified Computing System Manager

Create Service Profile Template

1. ☒ Identify Service Profile Template
2. ☒ Networking
3. ☒ Storage
4. ☒ Zoning
5. ☒ vNIC/vHBA Placement
6. ☒ Server Boot Order
7. ☒ **Maintenance Policy**
8. ☐ Server Assignment
9. ☐ Operational Policies

### Maintenance Policy

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

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: **default** + Create Maintenance Policy

Name: **default**  
 Description:  
 Reboot Policy: **User Ack**

< Prev   Next >   Finish   Cancel

**12. Specify the Server Assignment:**

- In the Pool Assignment list, choose **Infra\_Pool**.
- (Optional) Choose a Server Pool Qualification policy.
- Choose **Down** as the power state to be applied when the profile is associated with the server.
- Expand **Firmware Management** at the bottom of the page and choose **VM-Host-Infra** from the Host Firmware list.
- Click **Next**.

**Figure 53 UCS - Assigning a Server Pool and Setting Host Firmware Management Policy**

**Unified Computing System Manager**

Create Service Profile Template

1. ✓ Identify Service Profile Template  
 2. ✓ Networking  
 3. ✓ Storage  
 4. ✓ Zoning  
 5. ✓ vNIC/vHBA Placement  
 6. ✓ Server Boot Order  
 7. ✓ Maintenance Policy  
 8. ✓ **Server Assignment**  
 9. Operational Policies

### Server Assignment

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: **Infra\_Pool** + Create Server Pool

Select the power state to be applied when this profile is associated with the server.

☐ Up ☒ Down

The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification: **<not set>**

Restrict Migration: ☐

**Firmware Management (BIOS, Disk Controller, Adapter)**

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware: **VM-Host-Infra** + Create Host Firmware Package

< Prev   Next >   Finish   Cancel

**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 54** UCS - Adding BIOS Setting and Power Control Policies

**Create Service Profile Template**

## Unified Computing System Manager

Create Service Profile Template

1. ☒ Identify Service Profile Template
2. ☒ Networking
3. ☒ Storage
4. ☒ Zoning
5. ☒ vNIC/vHBA Placement
6. ☒ Server Boot Order
7. ☒ Maintenance Policy
8. ☒ Server Assignment
9. ☒ **Operational Policies**

### Operational Policies

Optionally specify information that affects how the system operates.

**BIOS Configuration**

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy:

**External IPMI Management Configuration**

**Management IP Address**

**Monitoring Configuration (Thresholds)**

**Power Control Policy Configuration**

Power control policy determines power allocation for a server in a given power group.

Power Control Policy:


**Scrub Policy**

< Prev Next > Finish Cancel

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 55** UCS - Dialog box for Service Profile Cloning Options



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 56 UCS - Modify Boot Policy in Cloned Service Profile Template**

**Modify Boot Policy**

Boot Policy: **Boot-Fabric-B** + Create Boot Policy

Name: **Boot-Fabric-B**

Description:

Reboot on Boot Order Change: **No**

Enforce vNIC/vHBA/iSCSI Name: **Yes**

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Boot Order**

+ - Filter Export Print

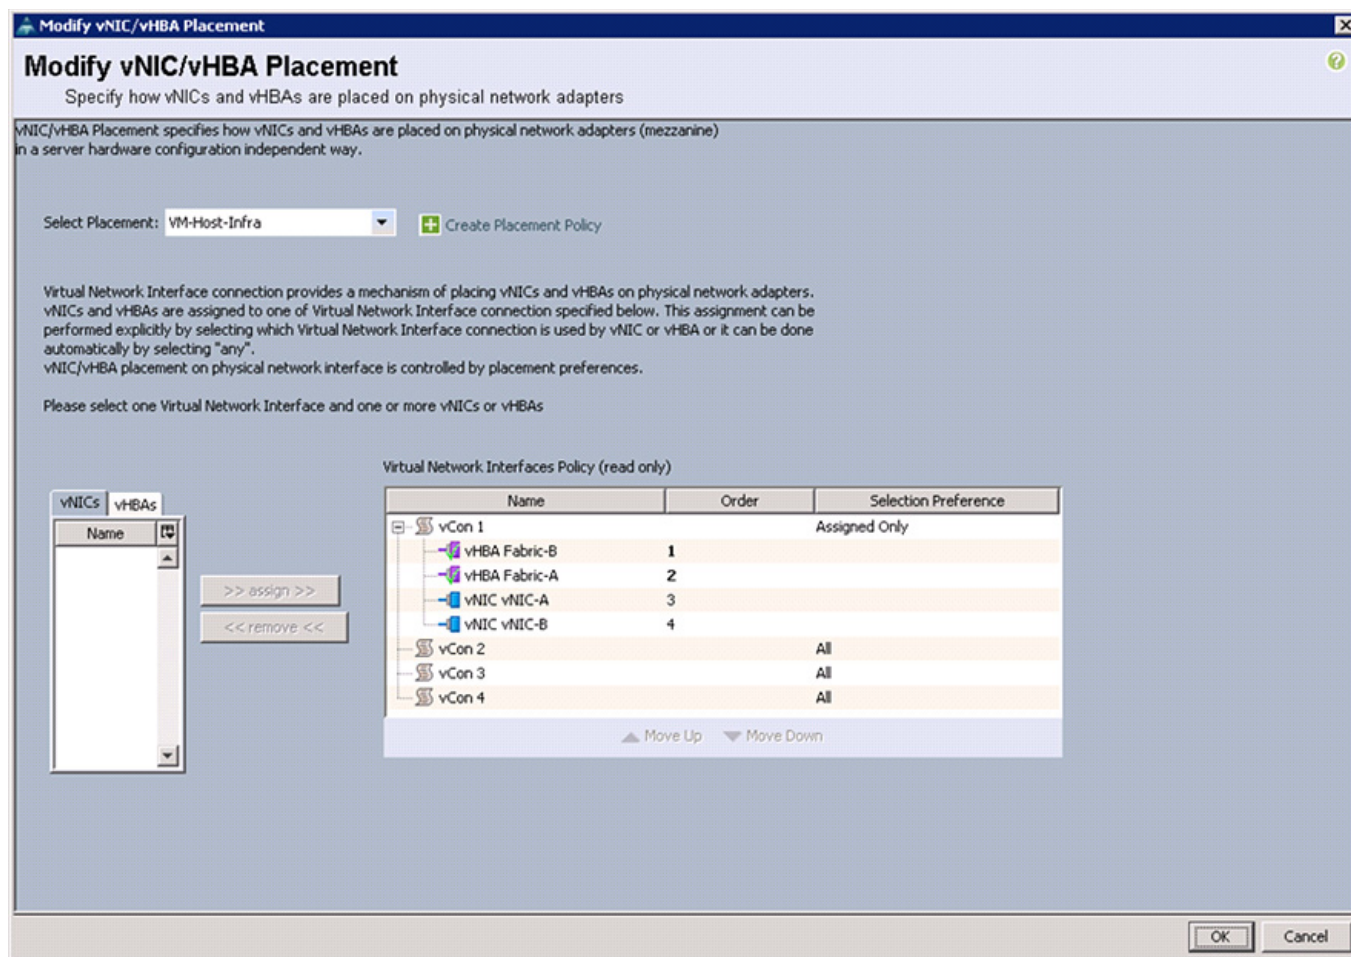
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		Fabric-B	Primary		
SAN Target primary			Primary	0	20:07:00:A0:98:37:78:08
SAN Target secondary			Secondary	0	20:09:00:A0:98:37:78:08
SAN secondary		Fabric-A	Secondary		
SAN Target primary			Primary	0	20:06:00:A0:98:37:78:08
SAN Target secondary			Secondary	0	20:08:00:A0:98:37:78:08

Create iSCSI vNIC    Set iSCSI Boot Parameters

OK    Cancel

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 57 UCS - Modify vNIC/vHBA Placement in Cloned Service Profile Template**




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. Enter 1 as the number of service profiles to create.
6. Click **OK** to create the service profile.

**Figure 58 UCS - Create Service Profile from Template VM-Host-Infra-Fabric-A**



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

**Figure 59 UCS - Create Service Profile from Template VM-Host-Infra-Fabric-B**



13. Click **OK** in the confirmation message.
14. 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.
15. (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 27](#) and [Table 28](#).

**Table 27** *FCP LIFs for FC WWPNs*

FCP LIFS	FC WWPN
fcp_lif01a	
fcp_lif01b	
fcp_lif02a	
fcp_lif02b	



**Note**

To gather the FC WWPN, run the network interface show command on the storage controller.

**Table 28** *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 28](#), 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.

## Network Configuration

The following section provides a detailed procedure for configuring the Cisco Nexus 7000 Switches for use in a FlexPod environment. Follow these steps precisely because failure to do so might result in an improper configuration.

**Note**

The configuration steps detailed in this section provides guidance for configuring the Nexus 7000 running release 6.1(2) within a multi-VDC environment

## Cisco Nexus 7000 Network Initial Configuration Setup

This section provides initial configuration details for setting up Cisco Nexus 7000 Switch pair.

### Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <<var\_nexus\_A\_hostname>>, follow these steps:

**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) [n]: yes
```

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

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

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

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

```
Do you want to enable admin vdc (yes/no) [n]: y
```

```
---- Basic System Configuration Dialog VDC: 2 ----
```

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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

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

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

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

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

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

```
Enable license grace period? (yes/no) [n]: y
```

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

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

```

Mgmt0 IPv4 netmask : <<var_nexus_A_mgmt0_netmask>>

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

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

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

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

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

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

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

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

NTP server IPv4 address : <<var_global_ntp_server_ip>>

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

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

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

The following configuration will be applied:
password strength-check
switchname <<var_nexus_A_hostname>>
license grace-period
vrf context management
ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
ntp server <<var_global_ntp_server_ip>>
system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address <<var_nexus_A_mgmt0_ip>> <<var_nexus_A_mgmt0_netmask>>
no shutdown

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

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

[#####] 100%
Copy complete.

```

## Cisco Nexus B

To set up the initial configuration for the Cisco Nexus B switch on <<var\_nexus\_B\_hostname>>, follow these steps:



### 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) [n]: yes

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

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

Enter the password for "admin": <<var\_password>>

Confirm the password for "admin": <<var\_password>>

Do you want to enable admin vdc (yes/no) [n]: y

---- Basic System Configuration Dialog VDC: 2 ----

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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

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

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

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

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

Enter the switch name : <<var\_nexus\_B\_hostname>>

Enable license grace period? (yes/no) [n]: y

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

Mgmt0 IPv4 address : <<var\_nexus\_B\_mgmt0\_ip>>

Mgmt0 IPv4 netmask : <<var\_nexus\_B\_mgmt0\_netmask>>

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

IPv4 address of the default gateway : <<var\_nexus\_B\_mgmt0\_gw>>

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

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

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

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

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

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

NTP server IPv4 address : <<var\_global\_ntp\_server\_ip>>

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

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



```

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

The following configuration will be applied:
  password strength-check
  switchname <<var_nexus_B_hostname>>
  license grace-period
vrf context management
ip route 0.0.0.0/0 <<var_nexus_B_mgmt0_gw>>
exit
  no feature telnet
  ssh key rsa 1024 force
  feature ssh
  ntp server <<var_global_ntp_server_ip>>
  system default switchport
  system default switchport shutdown
  copp profile strict
interface mgmt0
ip address <<var_nexus_B_mgmt0_ip>> <<var_nexus_B_mgmt0_netmask>>
no shutdown

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

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

[#####] 100%
Copy complete.

```

## Set Global Configurations and Create VDCs

This section provides information on setting global configurations and creating VDCs for Cisco Nexus 7000 Switches.

### Cisco Nexus 7000 A and Cisco Nexus 7000 B

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

1. Install the feature set for FCoE:

```
install feature-set fcoe
```

2. Enable apply the FCoE license on the F2-series module to be used for FCoE storage traffic:

```
license fcoe module 4
```

3. Enable system wide FCoE QoS and define the QoS policy to be used:

```
system qos
service-policy type network-qos default-nq-7e-policy
```

### Create Ethernet switching and FCoE Storage VDCs on Nexus 7000 A

1. Create IP Switching VDC. This process will take a moment:

```
vdc <<var_nexus_A_ip_vdc>>
```

2. Make the VDC compliant for F-series module(s) being used.

**Note**

f2 and f2e are the possible types:

```
limit-resource module-type f2
This will cause all ports of unallowed types to be removed from this vdc. Continue
(y/n)? [yes] Enter
```

**3. Allocate Interfaces used for IP Networking:**

```
allocate int eth 4/1-2, eth 4/41, eth 4/43, eth 4/17, eth 4/19, eth 4/27-28
Entire port-group is not present in the command. Missing ports will be included
automatically
Moving ports will cause all config associated to them in source vdc to be removed.
Are you sure you want to move the ports (y/n)? [yes] Enter
```

**4. Create Storage VDC for Fabric A. As earlier, this process will take a moment:**

```
vdc <<var_nexus_A_fcoe_vdc>> type storage
```

**5. Allocate Interfaces used for FCoE Networking:**

```
allocate interface ethernet 4/31-32, ethernet 4/37-38
Entire port-group is not present in the command. Missing ports will be included
automatically
Moving ports will cause all config associated to them in source vdc to be removed.
Are you sure you want to move the ports (y/n)? [yes] Enter
```

**6. Allocate FCoE VLAN used for Fabric A:**

```
allocate fcoe-vlan-range 101
```

**7. Save Configuration:**

```
copy run start
[#####] 100%
Copy complete.
```

## Create Ethernet switching and FCoE Storage VDCs on Nexus 7000 B

**1. Create IP Switching VDC. This process will take a moment:**

```
vdc <<var_nexus_B_ip_vdc>>
```

**2. Make the VDC compliant for F-series module(s) being used.****Note**

f2 and f2e are the possible types:

```
This will cause all ports of unallowed types to be removed from this vdc. Continue
(y/n)? [yes] Enter
```

**3. Allocate Interfaces used for IP Networking:**

```
allocate int eth 4/1-2, eth 4/41, eth 4/43, eth 4/17, eth 4/19, eth 4/27-28
Entire port-group is not present in the command. Missing ports will be included
automatically
Moving ports will cause all config associated to them in source vdc to be removed.
Are you sure you want to move the ports (y/n)? [yes] Enter
```

**4. Create Storage VDC for Fabric B. As earlier, this process will take a moment:**

```
vdc <<var_nexus_B_fcoe_vdc>> type storage
```

**5. Allocate Interfaces used for IP Networking:**

```
allocate interface ethernet 4/31-32, ethernet 4/37-38
Entire port-group is not present in the command. Missing ports will be included
automatically
Moving ports will cause all config associated to them in source vdc to be removed.
Are you sure you want to move the ports (y/n)? [yes] Enter
```

**6. Allocate FCoE VLAN used for Fabric B:**

```
allocate fcoe-vlan-range 102
```

**7. Save Configuration:**

```
copy run start
[#####] 100%
Copy complete.
```

## Initial Configuration Setup for the IP switching VDC

This section provides information on setting up IP switching VDC individually for both the Cisco Nexus 7000 switches.

### Setup the IP Switching VDC for Cisco Nexus A

From the Admin VDC on Nexus A, switch to the IP switching VDC:

```
switchto vdc <<var_nexus_A_ip_vdc>>
```

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

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

```
Enter the password for "admin":
```

```
Confirm the password for "admin":
```

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

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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

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

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

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

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

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

```
Enable license grace period? (yes/no) [n]: Enter
```

```

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

Mgmt0 IPv4 address : <<var_nexus_A_IP_VDC_mgmt0_ip>>

Mgmt0 IPv4 netmask : <<var_nexus_A_mgmt0_netmask>>

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

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

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

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

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

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

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

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

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

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

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

The following configuration will be applied:
password strength-check
switchname <<var_nexus_A_ip_vdc>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
system default switchport shutdown
interface mgmt0
ip address <<var_nexus_A_IP_VDC_mgmt0_ip>> <<var_nexus_A_mgmt0_netmask>>
no shutdown

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

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

[#####] 100%
Copy complete.

```

## Setup the IP Switching VDC for Cisco Nexus B

From the Admin VDC on Nexus B, switch to the IP switching VDC:

```

switchto vdc <<var_nexus_B_ip_vdc>>

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

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

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

```

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

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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

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

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

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

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

Enter the switch name : <<var\_nexus\_A\_ip\_vdc>>

Enable license grace period? (yes/no) [n]: Enter

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

Mgmt0 IPv4 address : <<var\_nexus\_B\_IP\_VDC\_mgmt0\_ip>>

Mgmt0 IPv4 netmask : <<var\_nexus\_B\_mgmt0\_netmask>>

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

IPv4 address of the default gateway : <<var\_nexus\_B\_mgmt0\_gw>>

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

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

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

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

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

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

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

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

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

The following configuration will be applied:

```
password strength-check
switchname <<var_nexus_B_ip_vdc>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_B_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
```

```

    system default switchport shutdown
interface mgmt0
ip address <<var_nexus_B_ip_vdc_mgmt0_ip>> <<var_nexus_B_mgmt0_netmask>>
no shutdown

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

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

[#####] 100%
Copy complete.

```

## Enable Appropriate Cisco Nexus Features and Settings for IP Switching

This section provides information on enabling the required Cisco Nexus features and IP switching settings.

### Cisco Nexus 7000 A and Cisco Nexus 7000 B

The following commands enable IP switching feature and set default spanning tree behaviors for the VDC:

1. From the IP switching VDC of each Nexus enter configuration mode:

```
config terminal
```

2. Use the following commands to enable the necessary features for this VDC:

```
feature udd
feature lacp
feature vpc
```

3. Configure spanning tree defaults for the IP switching VDC:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
```

4. Enable jumbo frame support:

```
system jumbomtu 9216
```

5. Save the running configuration to start-up:

```
copy run start
```

## Create VLANs for FlexPod Ethernet Traffic

This section provides details on creating VLANs for Cisco Nexus 7000 switches for handling FlexPod Ethernet traffic.

### Cisco Nexus 7000 A and Cisco Nexus 7000 B

To create the necessary virtual local area networks (VLANs), follow these steps on both switches:

From the global configuration mode, run the following commands:

```
vlan <<var_ib-mgmt_vlan_id>>
```

```

name IB-MGMT-VLAN

vlan <<var_native_vlan_id>>
name Native-VLAN

vlan <<var_nfs_vlan_id>>
name NFS-VLAN

vlan <<var_pkt_ctrl_vlan_id>>
name Packet-Control-VLAN

vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN

vlan <<var_vm_traffic_vlan_id>>
name VM-Traffic-VLAN

```

## Configure Virtual Port Channel Domain

This section provides details on configuring port channel domain for Cisco Nexus 7000 switches.

### Cisco Nexus 7000 A

To configure virtual port channels (vPCs) for switch A and its VDC for IP traffic, follow these steps:

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

```
vpc domain <<var_nexus_vpc_domain_id>>
```

2. Make Nexus 7000A the primary vPC peer by defining a small priority value:

```
role priority 10
```

3. Use the management interfaces on the supervisors of the Cisco Nexus 7000s to establish a keepalive link in the event of a complete vPC peer link failure:

```
peer-keepalive destination <<var_nexus_B_IP_VDC_mgmt0_ip>> source
<<var_nexus_A_IP_VDC_mgmt0_ip>>
```



#### Note

The management interfaces on the supervisors should be connected to an external management switch. They should never be cross connected between Nexus 7000 chassis.

4. Enable auto-recovery for this vPC domain:

```
auto-recovery
```

### Cisco Nexus 7000 B

To configure vPCs for switch B and its VDC for IP traffic, follow these steps:

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

```
vpc domain <<var_nexus_vpc_domain_id>>
```

2. Make Nexus 7000 B the secondary vPC peer by defining a larger priority value than that of the Nexus 7000 A:

```
role priority 20
```

3. Use the management interfaces on the supervisors of the Cisco Nexus 7000s to establish a keepalive link in the event of a complete vPC peer link failure:

```
peer-keepalive destination <<var_nexus_A_IP_VDC_mgmt0_ip>> source
<<var_nexus_B_IP_VDC_mgmt0_ip>>
```

4. Enable auto-recovery for this vPC domain:

```
auto-recovery
```

## Configure Network Interfaces for the VPC Peer Links

This section provides details on configuring network interfaces for the vPC peer links between the Cisco Nexus 7000 switches.

### Cisco Nexus 7000 A

1. Define a port description for the interfaces connecting to VPC Peer <<var\_nexus\_B\_hostname>>

```
interface Eth4/41
description VPC Peer <<var_nexus_B_hostname>>:4/41
```

```
interface Eth4/43
description VPC Peer <<var_nexus_B_hostname>>:4/43
```

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

```
interface Eth4/41,Eth4/43
channel-group 10 mode active
no shutdown
```

3. Define a description for the port channel connecting to <<var\_nexus\_B\_hostname>>

```
interface Po10
description vpc peer-link
```

4. Make the port-channel a switchport, and configure a trunk to allow InBand management, NFS, and VM traffic, Packet Control VLANs and the native VLAN

```
switchport
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>>, <<var_pkt_ctrl_vlan_id>>
```

5. Make this port channel the VPC Peer Link and bring it up

```
vpc peer-link
no shutdown
```

### Cisco Nexus 7000 B

1. Define a port description for the interfaces connecting to VPC Peer <<var\_nexus\_A\_hostname>>

```
interface Eth4/41
description VPC Peer <<var_nexus_A_hostname>>:4/41
```

```
interface Eth4/43
description VPC Peer <<var_nexus_A_hostname>>:4/43
```



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

```
interface Eth4/41,Eth4/43
channel-group 10 mode active
no shutdown
```

3. Define a description for the port channel connecting to <<var\_nexus\_A\_hostname>>

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

4. Make the port-channel a switchport, and configure a trunk to allow InBand management, NFS, and VM traffic, Packet Control VLANs and the native VLAN

```
switchport
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>>, <<var_pkt_ctrl_vlan_id>>
```

5. Make this port channel the VPC Peer Link and bring it up

```
vpc peer-link
no shutdown
```

## Configure Network Interfaces to NetApp Storage for Data Traffic

This section provides details on configuring network interfaces to NetApp storage device.

### Cisco Nexus 7000 A

1. Define a port description for the interface connecting to <<var\_node01>>

```
interface Eth4/1
description <<var_node01>>:e3a
```

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

```
channel-group 11 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var\_node01>>

```
interface Po11
description <<var_node01>>
```

4. Make the port-channel a switchport, and configure a trunk to allow NFS traffic vlan and the native vlan

```
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
```

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

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

6. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 11
no shutdown
```

8. Define a port description for the interface connecting to <<var\_node02>>

```
interface Eth4/2
description <<var_node02>>:e3a
```

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

```
channel-group 12 mode active
no shutdown
```

10. Define a description for the port channel connecting to <<var\_node02>>

```
interface Po12
description <<var_node02>>
```

11. Make the port channel a switchport, and configure a trunk to allow NFS traffic vlan and the native vlan

```
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
```

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

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

13. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

14. Make this a VPC port channel and bring it up

```
vpc 12
no shutdown
```

## Cisco Nexus 7000 B

1. Define a port description for the interface connecting to <<var\_node01>>

```
interface Eth4/1
description <<var_node01>>:e4a
```

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

```
channel-group 11 mode active
no shutdown
```

3. Define a description for the port channel connecting to <<var\_node01>>

```
interface Po11
description <<var_node01>>
```

4. Make the port channel a switchport, and configure a trunk to allow NFS traffic vlan and the native vlan

```
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
```

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

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

6. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 11
no shutdown
```

8. Define a port description for the interface connecting to <<var\_node02>>

```
interface Eth4/2
description <<var_node02>>:e4a
```

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

```
channel-group 12 mode active
no shutdown
```

10. Define a description for the port-channel connecting to <<var\_node02>>

```
interface Po12
description <<var_node02>>
```

11. Make the port-channel a switchport, and configure a trunk to allow NFS traffic vlan and the native vlan

```
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
```

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

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

13. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 12
no shutdown
```

## Configure Network Interfaces to UCS Fabric Interconnect

This section provides details on configuring network interfaces to cisco UCS fabric Interconnect.

### Cisco Nexus 7000 A

1. Define a port description for the interface connecting to <<var\_ucs\_clustername>>-A

```
interface Eth4/27
description <<var_ucs_clustername>>-A:1/27
```

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

```
channel-group 13 mode active
```

```
no shutdown
```

3. Define a description for the port-channel connecting to <<var\_ucs\_clustername>>-A

```
interface Po13
description <<var_ucs_clustername>>-A
```

4. Make the port-channel a switchport, and configure a trunk to allow InBand management, NFS, and VM traffic VLANs and the native VLAN

```
switchport
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>>,
```

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

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

6. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 13
no shutdown
```

8. Define a port description for the interface connecting to <<var\_ucs\_clustername>>-B

```
interface Eth4/28
description <<var_ucs_clustername>>-B:1/28
```

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

```
channel-group 14 mode active
no shutdown
```

10. Define a description for the port-channel connecting to <<var\_ucs\_clustername>>-B

```
interface Po14
description <<var_ucs_clustername>>-B
```

11. Make the port-channel a switchport, and configure a trunk to allow InBand management, NFS, and VM traffic VLANs and the native VLAN

```
switchport
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>>,
```

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

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

13. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 14
no shutdown
```

## Cisco Nexus 7000 B

1. Define a port description for the interface connecting to <<var\_ucs\_clustername>>-B

```
interface Eth4/27
description <<var_ucs_clustername>>-B:1/27
```

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

```
channel-group 14 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var\_ucs\_clustername>>-B

```
interface Po14
description <<var_ucs_clustername>>-B
```

4. Make the port-channel a switchport, and configure a trunk to allow InBand management, NFS, and VM traffic VLANs and the native VLAN

```
switchport
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>>,
```

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

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

6. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 14
no shutdown
```

8. Define a port description for the interface connecting to <<var\_ucs\_clustername>>-A

```
interface Eth4/28
description <<var_ucs_clustername>>-A:1/28
```

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

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

10. Define a description for the port-channel connecting to <<var\_ucs\_clustername>>-A

```
interface Po13
description <<var_ucs_clustername>>-A
```

11. Make the port-channel a switchport, and configure a trunk to allow InBand management, NFS, and VM traffic VLANs and the native VLAN

```
switchport
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>>,
```

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

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

13. Set the MTU to be 9216 to support jumbo frames:

```
mtu 9216
```

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

```
vpc 13
no shutdown
```

## Configure Ports for Cisco Nexus 1110-X Virtual Appliances

This section provides details on configuring ports for Cisco Nexus 1110-X virtual appliances.

### Cisco Nexus 7000 A

To configure the ports in switch A that are connected to the Cisco Nexus 1110-X, follow these steps:

1. Define a port description for the interface connecting to Cisco Nexus 1110-X-1

```
interface Eth4/17
description <<var_nexus_1110x-1>>:Eth1
```

2. Define a port description for the interface connecting to Cisco Nexus 1110-X-2

```
interface Eth4/19
description <<var_nexus_1110x-2>>:Eth1
```

3. Configure both Nexus 1110-X ports to be trunks carrying the InBand management and Packet Control VLANs

```
interface Eth4/17, Eth4/19
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_pkt-ctrl_vlan_id>>
```

4. Make the interfaces spanning tree edge ports

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

5. Bring it up the interfaces

```
no shutdown
```

### Cisco Nexus 7000 B

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

1. Define a port description for the interface connecting to Cisco Nexus 1110-X-1

```
interface Eth4/17
description <<var_nexus_1110x-1>>:Eth2
```

2. Configure both Nexus 1110-X ports to be trunks carrying the InBand management and Packet Control VLANs

```
interface Eth4/17, Eth4/19
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
```

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

3. Make the interfaces spanning tree edge ports

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

4. Bring it up the interfaces

```
no shutdown
```

## Interconnect Nexus Data Plane to Management Plane

This section provides details on the interconnection of Cisco Nexus data plane with the management pane.

### Cisco Nexus 7000 A and Cisco Nexus 7000 B

The Nexus data plane needs access to the management plane to enable management access across the IP switching environment.

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

```
interface Eth4/44
description IB-Mgmt:<<mgmt_uplink_port>>
```

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

```
channel-group 9 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var\_ucs\_clustername>>-A

```
interface Po9
description IB-Mgmt
```

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

```
switchport
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
```

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

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

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

```
vpc 9
no shutdown
```



#### Note

It may be desired to create a dedicated Switch Virtual Interface (SVI) on the Nexus data plane to test and troubleshoot the management plane. If a L3 interface is deployed, be sure it is deployed on both Cisco Nexus 7000 IP Switching VDCs to ensure Type-2 VPC consistency.

7. Save the running configuration to start-up in both Nexus 7000 IP switching VDCs:

```
copy run start
```

8. Exit to the Admin VDC: (multiple exit commands may be required)

```
exit
```

## Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment or use the Nexus 7000 of the FlexPod as your distribution block. If an existing Cisco Nexus environment is present, Cisco recommends using vPCs to uplink the Cisco Nexus 7000 switches included in the FlexPod environment into the existing 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.

## Configure FCoE Storage VDCs

This section provides details on setting up the FCoE storage VDCs for Cisco Nexus switches.

### Setup the FCoE Storage VDC for Cisco Nexus A

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

Switch to the switching VDC.

```
switchto vdc <<var_nexus_A_fcoe_vdc>>
```

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

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

```
Enter the password for "admin":
```

```
Confirm the password for "admin":
```

```
---- Basic System Configuration Dialog VDC: 3 ----
```

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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

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

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

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

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

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

```
Enable license grace period? (yes/no) [n]: Enter
```

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

```
Mgmt0 IPv4 address : <<var_nexus_A_fcoe_vdc_mgmt0_ip>>
```

```
Mgmt0 IPv4 netmask : <<var_necus_A_mgmt0_netmask>>
```



```

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

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

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

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

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

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

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

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

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

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

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

The following configuration will be applied:
password strength-check
switchname <<var_nexus_A_fcoe_vdc>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
system default switchport shutdown
interface mgmt0
ip address <<var_nexus_A_fcoe_vdc_mgmt0_ip>> <<var_nexus_A_mgmt0_netmask>>
no shutdown

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

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

Disabling ssh: as its enabled right now:
generating rsa key(1024 bits).....
.
generated rsa key
Enabling ssh: as it has been disabled
% All 0s mask is invalid
2. Review the configuration summary before enabling the configuration.
Would you like to save the running-config to startup-config? (yes/no) [n]: y

[#####] 100%
Copy complete.

```

## Setup the FCoE Storage VDC for Cisco Nexus B

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

1. Switch to the switching VDC.

```

switchto vdc <<var_nexus_B_fcoe_vdc>>

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

```

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

Enter the password for "admin":

Confirm the password for "admin":

---- Basic System Configuration Dialog VDC: 3 ----

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

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

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

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

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

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

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

Enter the switch name : <<var\_nexus\_B\_hostname>>

Enable license grace period? (yes/no) [n]: Enter

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

Mgmt0 IPv4 address : <<var\_nexus\_B\_fcoe\_vdc\_mgmt0\_ip>>

Mgmt0 IPv4 netmask : <<var\_nexus\_B\_mgmt0\_netmask>>

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

IPv4 address of the default gateway : <<var\_nexus\_B\_mgmt0\_gw>>

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

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

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

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

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

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

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

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

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

The following configuration will be applied:

password strength-check

switchname <<var\_nexus\_B\_fcoe\_vdc>>

```

vrf context management
ip route 0.0.0.0/0 <<var_nexus_B_mgmt0_gw>>
exit
  no feature telnet
  ssh key rsa 1024 force
  feature ssh
  system default switchport
  system default switchport shutdown
interface mgmt0
ip address <<var_nexus_B_fcoe_vdc_mgmt0_ip>> <<var_nexus_B_mgmt0_netmask>>
no shutdown

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

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

Disabling ssh: as its enabled right now:
generating rsa key(1024 bits).....
generated rsa key
Enabling ssh: as it has been disabled
% All 0s mask is invalid

```

2. Review the configuration summary before enabling the configuration.

```

Would you like to save the running-config to startup-config? (yes/no) [n]: y

[#####] 100%
Copy complete.

```

## Enable FCoE Features

This section provides information on enabling the required FCoE features in the storage VDCs of Cisco Nexus 7000 switches.

### Cisco Nexus A – Storage A VDC and Cisco Nexus B – Storage B VDC

Use switchto from the Admin VDC to change to the FCoE storage VDC and use the following commands enable FCoE features.

1. From the FCoE storage switching VDC of each Nexus 7000 enter into configuration mode:

```
config terminal
```

2. Use the following commands to enable the necessary features for this VDC:

```

feature-set fcoe
feature npiv
feature lacp
feature lldp

```

## Create VSANs, Assign and Enable Virtual Fibre Channel Ports

This procedure sets up Fibre Channel over Ethernet (FCoE) connections between the Cisco Nexus 7000 switches, the Cisco UCS Fabric Interconnects, and the NetApp storage systems.

## Cisco Nexus 7000 A

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

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

```
vsan database
vsan <<var_vsan_a_id>> name FCoE_Fabric_A
```

2. Create a VLAN for FCoE

```
vlan <<var_fabric_a_fcoe_vlan_id>>
name FCoE_Fabric_A
```

3. Apply the FCoE VSAN

```
fcoe vsan <<var_vsan_a_id>>
```

### Create FCoE port channel to the UCS FI

1. Define a port description for the FCoE interfaces connecting to <<var\_ucs\_clusternam>>-A

```
interface Eth4/31
description <<var_ucs_clusternam>>-A:1/31

interface Eth4/32
description <<var_ucs_clusternam>>-A:1/32
```

2. Apply it to a port channel and bring up the interfaces.

```
interface Eth4/31, Eth4/32
channel-group 2 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var\_ucs\_clusternam>>-A.

```
interface Po2
description <<var_ucs_clusternam>> Fabric A
```

4. Make the port-channel a switchport, and configure a trunk to allow FCoE VLAN.

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_fabric_a_fcoe_vlan_id>>
```

5. Bring up this FCoE Port-channel

```
no shutdown
```

### Create vFC port channel to the UCS FI

1. Define a description for the vFC port-channel

```
interface vfc-port-channel 2
switchport description <<var_ucs_clusternam>>-A:FCoE
```

2. Bind the vFC to the Port-Channel

```
bind interface po2
```

3. Allow the FCoE VLAN on the vFC

```
switchport trunk allowed vsan <<var_fabric_a_fcoe_vlan_id>>
```

4. Bring up this vFC

```
no shutdown
```

5. Add the vFC to the FCoE VSAN Database

```
vsan database
vsan <<var_fabric_a_fcoe_vlan_id>> interface vfc-po2
```

### Create FCoE port to the NetApp Storage <<var\_node01>>

1. Define a port description for the FCoE interfaces connecting to <<var\_node01>>

```
interface Eth4/37
description <<var_node01>>:3b
```

2. Make the interface a switchport and configure a trunk to allow the FCoE VLAN

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_fabric_a_fcoe_vlan_id>>
```

3. Bring up this FCoE interface

```
no shutdown
```

### Create vFC port to the NetApp Storage <<var\_node01>>

1. Define a description for the vFC port

```
interface vfc 437
switchport description <<var_node01>>:FCoE
```

2. Bind the vFC to the Port

```
bind interface Eth4/37
```

3. Allow the FCoE VSAN on the vFC

```
switchport trunk allowed vsan <<var_fabric_a_fcoe_vlan_id>>
```

4. Bring up this vFC

```
no shutdown
```

5. Add the vFC to the FCoE VSAN Database

```
vsan database
vsan <<var_fabric_a_fcoe_vlan_id>> interface vfc 437
```

### Create FCoE port to the NetApp Storage <<var\_node02>>

1. Define a port description for the FCoE interfaces connecting to <<var\_node02>>

```
interface Eth4/38
description <<var_node02>>:4b
```

2. Make the interface a switchport and configure a trunk to allow the FCoE VLAN

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_fabric_a_fcoe_vlan_id>>
```

3. Bring up this FCoE interface

```
no shutdown
```

## Create vFC port to the NetApp Storage <<var\_node02>>

1. Define a description for the vFC port.

```
interface vfc 438
switchport description <<var_node02>>:FCoE
```

2. Bind the vFC to the Port.

```
bind interface Eth4/38
```

3. Allow the FCoE VSAN on the vFC

```
switchport trunk allowed vsan <<var_fabric_a_fcoe_vlan_id>>
```

4. Bring up this vFC

```
no shutdown
```

5. Add the vFC to the FCoE VSAN Database

```
vsan database
vsan <<var_fabric_a_fcoe_vlan_id>> interface vfc 438
```

## Cisco Nexus 7000 B

To configure virtual storage area networks (VSANs), assign virtual Fibre Channel (vFC) ports, and enable vFC ports on switch B, complete the following steps:

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

```
vsan database
vsan <<var_vsan_b_id>> name FCoE_Fabric_B
```

2. Create a VLAN for FCoE

```
vlan <<var_fabric_b_fcoe_vlan_id>>
name FCoE_Fabric_B
```

3. Apply the FCoE VSAN

```
fcoe vsan <<var_vsan_b_id>>
```

## Create FCoE port channel to the UCS FI

1. Define a port description for the FCoE interfaces connecting to <<var\_ucs\_clustername>>-B

```
interface Eth4/31
description <<var_ucs_clustername>>-B:1/31

interface Eth4/32
description <<var_ucs_clustername>>-B:1/32
```

2. Apply it to a port channel and bring up the interfaces

```
interface Eth4/31, Eth4/32
channel-group 2 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var\_ucs\_clustername>>-B

```
interface Po2
description <<var_ucs_clustername>> Fabric B
```

4. Make the port-channel a switchport, and configure a trunk to allow FCoE VLAN

```

switchport
switchport mode trunk
switchport trunk allowed vlan <<var_fabric_b_fcoe_vlan_id>>

```

#### 5. Bring up this FCoE Port-channel

```
no shutdown
```

### Create vFC port channel to the UCS FI

#### 1. Define a description for the vFC port-channel

```

interface vfc-port-channel 2
switchport description <<var_ucs_clustername>>-B:FCoE

```

#### 2. Bind the vFC to the Port-Channel

```
bind interface po2
```

#### 3. Allow the FCoE VLAN on the vFC

```
switchport trunk allowed vsan <<var_fabric_b_fcoe_vlan_id>>
```

#### 4. Bring up this vFC

```
no shutdown
```

#### 5. Add the vFC to the FCoE VSAN Database

```

vsan database
vsan <<var_fabric_b_fcoe_vlan_id>> interface vfc-po2

```

### Create FCoE port to the NetApp Storage <<var\_node01>>

#### 1. Define a port description for the FCoE interfaces connecting to <<var\_node01>>

```

interface Eth4/37
description <<var_node02>>:3b

```

#### 2. Make the interface a switchport and configure a trunk to allow the FCoE VLAN

```

switchport
switchport mode trunk
switchport trunk allowed vlan <<var_fabric_b_fcoe_vlan_id>>

```

#### 3. Bring up this FCoE interface

```
no shutdown
```

### Create vFC port to the NetApp Storage <<var\_node01>>

#### 1. Define a description for the vFC port

```

interface vfc 437
switchport description <<var_node01>>:FCoE

```

#### 2. Bind the vFC to the Port

```
bind interface Eth4/37
```

#### 3. Allow the FCoE VSAN on the vFC

```
switchport trunk allowed vsan <<var_fabric_b_fcoe_vlan_id>>
```

#### 4. Bring up this vFC

```
no shutdown
```

5. Add the vFC to the FCoE VSAN Database

```
vsan database
vsan <<var_fabric_b_fcoe_vlan_id>> interface vfc 437
```

## Create FCoE port to the NetApp Storage <<var\_node02>>

1. Define a port description for the FCoE interfaces connecting to <<var\_node02>>

```
interface Eth4/38
description <<var_node01>>:4b
```

2. Make the interface a switchport and configure a trunk to allow the FCoE VLAN

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_fabric_b_fcoe_vlan_id>>
```

3. Bring up this FCoE interface

```
no shutdown
```

## Create vFC port to the NetApp Storage <<var\_node02>>

1. Define a description for the vFC port

```
interface vfc 438
switchport description <<var_node02>>:FCoE
```

2. Bind the vFC to the Port

```
bind interface Eth4/38
```

3. Allow the FCoE VSAN on the vFC

```
switchport trunk allowed vsan <<var_fabric_b_fcoe_vlan_id>>
```

4. Bring up this vFC

```
no shutdown
```

5. Add the vFC to the FCoE VSAN Database

```
vsan database
vsan <<var_fabric_b_fcoe_vlan_id>> interface vfc 438
```

## Create Device Aliases

This section provides details on configuring device aliases and zones on the Cisco Nexus 7000 switches.

### Cisco Nexus 7000 A

To configure device aliases and zones for the primary boot paths of switch A on <<var\_nexus\_A\_hostname>>, follow this 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 fcp_lif01a pwnn <<var_fcp_lif01a_wwpn>>
device-alias name fcp_lif02a pwnn <<var_fcp_lif02a_wwpn>>
exit
device-alias commit

```

## Cisco Nexus 7000 B

To configure device aliases and zones for the boot paths of switch B on <<var\_nexus\_B\_hostname>>, follow this step:

From the global configuration mode, run the following commands:

```

device-alias database
device-alias name VM-Host-Infra-01_B pwnn <<var_vm_host_infra_01_B_wwpn>>
device-alias name VM-Host-Infra-02_B pwnn <<var_vm_host_infra_02_B_wwpn>>
device-alias name fcp_lif01b pwnn <<var_fcp_lif01b_wwpn>>
device-alias name fcp_lif02b pwnn <<var_fcp_lif02b_wwpn>>
exit
device-alias commit

```

## Create Zones

This section provides details on creating zones on Cisco Nexus 7000 switches.

## Cisco Nexus 7000 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 fcp_lif01a
member device-alias fcp_lif02a
exit
zone name VM-Host-Infra-02_A vsan <<var_vsan_a_id>>
member device-alias VM-Host-Infra-02_A
member device-alias fcp_lif01a
member device-alias fcp_lif02a
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 7000 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 fcp_lif01b
member device-alias fcp_lif02b
exit
zone name VM-Host-Infra-02_B vsan <<var_vsan_b_id>>
member device-alias VM-Host-Infra-02_B
member device-alias fcp_lif01b
member device-alias fcp_lif02b
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 - SAN Boot

### Clustered Data ONTAP SAN Boot Storage Setup

#### Create Igroups

From the cluster management node SSH connection, enter the following:

```
igroup create -vserver Infra_Vserver -igroup VM-Host-Infra-01 -protocol fcp -ostype
vmware -initiator <<var_vm_host_infra_01_A_wwpn>>, <<var_vm_host_infra_01_B_wwpn>>
igroup create -vserver Infra_Vserver -igroup VM-Host-Infra-02 -protocol fcp -ostype
vmware -initiator <<var_vm_host_infra_02_A_wwpn>>, <<var_vm_host_infra_02_B_wwpn>>
igroup create -vserver Infra_Vserver -igroup MGMT-Hosts -protocol fcp -ostype vmware
-initiator
<<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

From the cluster management SSH connection, enter the following:

```
lun map -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-01 -igroup
VM-Host-Infra-01 -lun-id 0
```

```
lun map -vserver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-02 -igroup
VM-Host-Infra-02 -lun-id 0
```

# VMware vSphere 5.1 Setup

## FlexPod VMware ESXi 5.1 FCoE on Clustered Data ONTAP

This section provides detailed instructions for installing VMware ESXi 5.1 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. Log in to Cisco UCS Manager by using the admin user name and password.
3. From the main menu, click the **Servers** tab.
4. Choose **Servers > Service Profiles > root > VM-Host-Infra-01**.
5. Right-click VM-Host-Infra-01 and choose KVM Console.
6. Choose **Servers > Service Profiles > root > VM-Host-Infra-02**.
7. Right-click VM-Host-Infra-02 and choose KVM Console Actions > KVM Console.

## 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. Click **Add Image**.
3. Browse to the ESXi installer ISO image file and click **Open**.
4. Check the Mapped check box to map the newly added image.
5. Click the **KVM** tab to monitor the server boot.

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

## Set Up Management Networking for ESXi Hosts

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



### Note

The enic version used in this configuration is 2.1.2.38, and the fnic version is 1.5.0.20.

1. Open a web browser on the management workstation and navigate to:  
<http://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI5X-CISCO-ENIC-21238&productId=285>
2. Download the enic\_driver\_2.1.2.38-1023014.zip driver bundle.
3. Open a web browser on the management workstation and navigate to:  
<http://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI5X-CISCO-FNIC-15020&productId=285>
4. Download the fnic\_driver\_1.5.0.20-1021375.zip driver bundle.
5. Open both the enic and fnic driver bundles. These bundles include the VMware driver bundles that will be uploaded to vCenter.
  - Network: enic\_driver\_2.1.2.38-offline\_bundle-1023014.zip
  - Storage: fnic\_driver\_1.5.0.20-offline\_bundle-1021375.zip
6. Save the location of these driver bundles for uploading to vCenter in the next section.



### Note

If the links above have changed, go to [www.cisco.com](http://www.cisco.com) for the latest ISO image of UCS-related drivers. This ISO will either have the drivers included or may have an HTML file with the location of the latest network and FCoE storage drivers.

- The network driver link can be found in the README.html file in this directory of the ISO: VMware/Network/Cisco/1280/ESXi\_5.1
- The storage driver link can be found in the README.html file in this directory of the ISO: VMware/Storage/Cisco/1280/ESXi\_5.1

## 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. Choose 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 `enic_driver_2.1.2.38-offline_bundle-1023014.zip`.
6. Click **Open** to open the file.
7. Click **Yes** to upload the .zip file to datastore1.
8. Click the fourth button and choose Upload File.
9. Navigate to the saved location for the downloaded fnic driver version and choose `fnic_driver_1.5.0.20-offline_bundle-1021375.zip`.
10. Click **Open** to open the file.
11. Click **Yes** to upload the .zip file to datastore1.
12. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
13. 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 -d
/vmfs/volumes/datastore1/enic_driver_2.1.2.38-offline_bundle-1023014.zip
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -d
/vmfs/volumes/datastore1/enic_driver_2.1.2.38-offline_bundle-1023014.zip
```
14. 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 -d
/vmfs/volumes/datastore1/fnic_driver_1.5.0.20-offline_bundle-1021375.zip
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -d
/vmfs/volumes/datastore1/fnic_driver_1.5.0.20-offline_bundle-1021375.zip
```
15. From the vSphere Client, right-click each host in the inventory and choose Reboot.
16. Click **Yes** to continue.
17. Enter a reason for the reboot and click **OK**.
18. 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 60](#).

**Figure 60** ESXi - VM-Host-Infra-01 vSphere Standard Switch Network Configuration



## 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 61](#).

**Figure 61** *ESXi - VM-Host-Infra-02 vSphere Standard Switch Network Configuration*

The screenshot displays the vSphere Client interface for the host 'icef1-h12'. The left sidebar shows the 'Hardware' section expanded, with 'Networking' selected. The main pane shows the 'Configuration' tab for the 'vSphere Standard Switch'. The 'View' dropdown is set to 'vSphere Standard Switch'. The 'Networking' section shows the configuration for 'Standard Switch: vSwitch0'. The 'Virtual Machine Port Group' list includes 'IB-MGMT Network' (VLAN ID: 3175), 'VMkernel-vMotion' (VLAN ID: 3173), 'VMkernel-NFS' (VLAN ID: 3170), and 'VMkernel-MGMT' (VLAN ID: 3175). The 'Physical Adapters' section shows 'vmnic0' with a speed of 10000 and a status of 'Full'. The 'Recent Tasks' section at the bottom shows two completed tasks: 'Reconfigure port group' and 'Update virtual NIC'.

Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completion Time
Reconfigure port group	icef1-h12	Completed		root	2/20/2013 4:51:45 PM	2/20/2013 4:51:45 PM	2/20/2013 4:51:45 PM
Update virtual NIC	icef1-h12	Completed		root	2/20/2013 4:51:45 PM	2/20/2013 4:51:45 PM	2/20/2013 4:51:45 PM

## 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. Select Network File System and click **Next**.
6. The wizard prompts for the location of the NFS export. Enter <<var\_nfs\_lif02\_ip>> as the IP address for nfs\_lif02.
7. Enter **/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\_nfs\_lif01\_ip>> as the IP address for nfs\_lif01.
15. Enter **/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. Choose 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 5.1

The procedures in the following subsections provide detailed instructions for installing VMware vCenter 5.1 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 vSphere 5.1 documentation at: <http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html>

---

To install VMware vCenter 5.1, 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 217](#) for this setup.

## 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. 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 Edit the Virtual Machine Settings Before Completion check box. 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 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 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 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. 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 and then enter the password to log in to the VM.
44. 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.

---

45. If necessary, activate Windows.
46. 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 62 SQL Server - .NET Framework Installation**



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 63** SQL Server - Feature Selection



28. Click **Next**.

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

**Figure 64** *SQL Server - Instance Configuration*

**SQL Server 2008 R2 Setup**

**Instance Configuration**

Specify the name and instance ID for the instance of SQL Server. Instance ID becomes part of the installation path.

Setup Support Rules  
Setup Role  
Feature Selection  
Installation Rules  
**Instance Configuration**  
Disk Space Requirements  
Server Configuration  
Database Engine Configuration  
Error Reporting  
Installation Configuration Rules  
Ready to Install  
Installation Progress  
Complete

☒ **Default instance**  
☐ **Named instance:**

Instance ID:   
 Instance root directory:  ...

SQL Server directory: C:\Program Files\Microsoft SQL Server\MSSQL10\_50.MSSQLSERVER

Installed instances:

Instance Name	Instance ID	Features	Edition	Version
---------------	-------------	----------	---------	---------

< Back   Next >   Cancel   Help

30. Click **Next** for Disk Space Requirements.
31. For the SQL Server Agent service, choose 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 Password field.
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 65**      **SQL Server - Service Account Configuration**

SQL Server 2008 R2 Setup

Server Configuration

Specify the service accounts and collation configuration.

Setup Support Rules  
Setup Role  
Feature Selection  
Installation Rules  
Instance Configuration  
Disk Space Requirements  
**Server Configuration**  
Database Engine Configuration  
Error Reporting  
Installation Configuration Rules  
Ready to Install  
Installation Progress  
Complete

Service Accounts | Collation

Microsoft recommends that you use a separate account for each SQL Server service.

Service	Account Name	Password	Startup Type
SQL Server Agent	Administrator	••••••••	Automatic
SQL Server Database Engine	Administrator	••••••••	Automatic
SQL Server Browser	NT AUTHORITY\LOCAL S...		Disabled

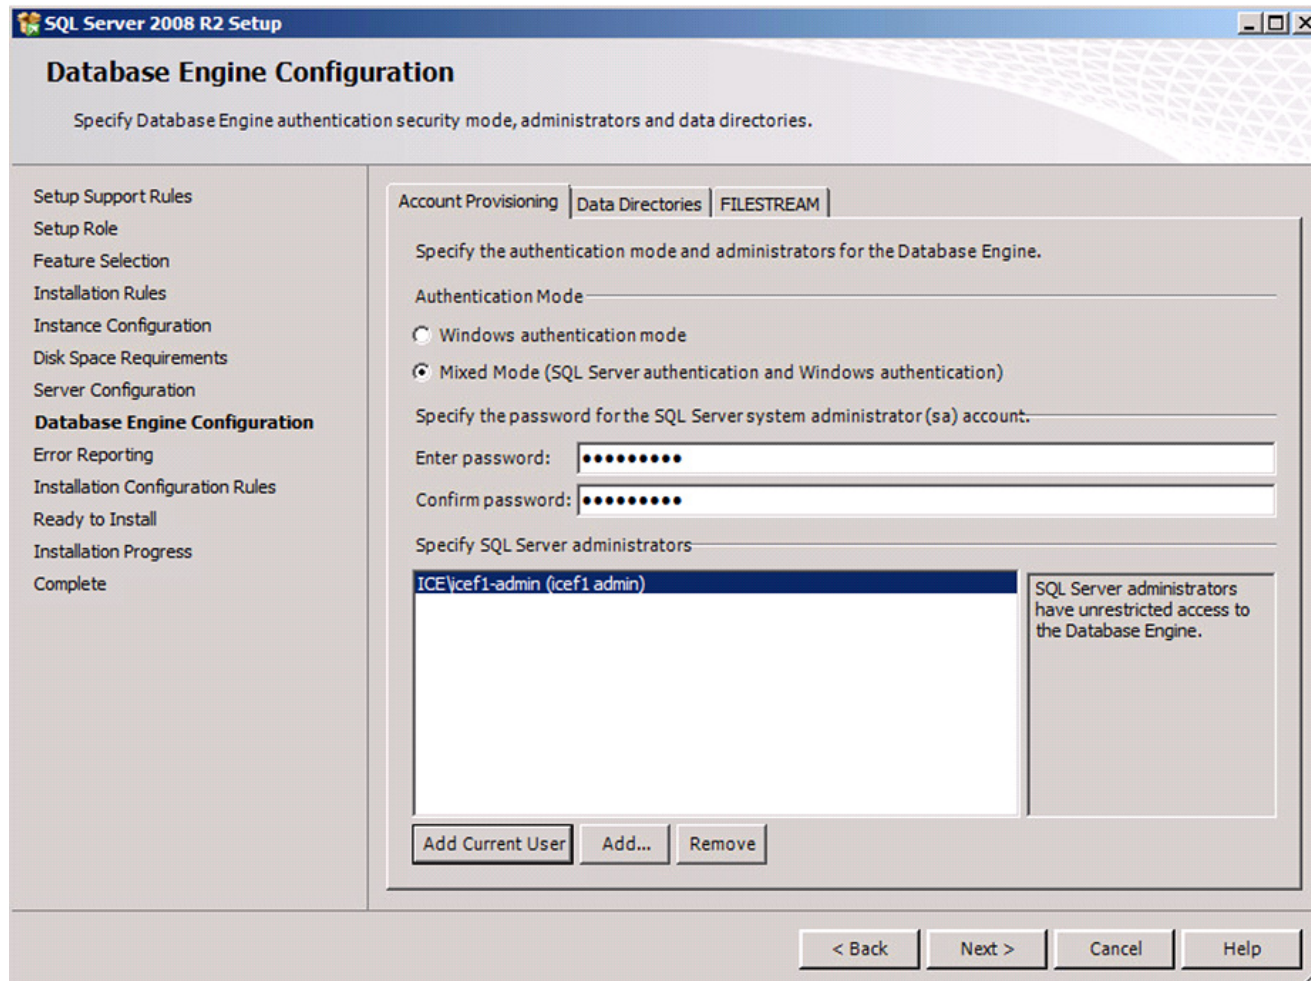
Use the same account for all SQL Server services

< Back   Next >   Cancel   Help

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 66 SQL Server - Database Engine Configuration**



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. Open the SQL Server Management Studio by selecting **Start > All Programs > Microsoft SQL Server 2008 R2 > 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 **vpuser** 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 67 illustrates the execution of the script.

**Figure 67**      **SQL Server - Configuration Script**

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:

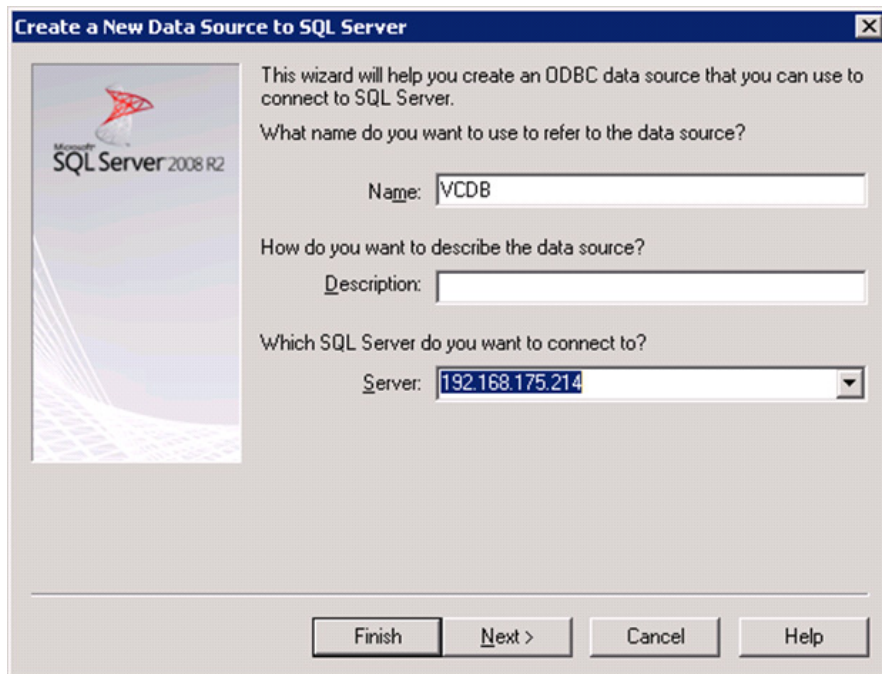
1. Using the instructions for building a SQL Server VM provided in the section “Build Microsoft SQL Server VM,” build a VMware vCenter VM with the following configuration in the <<var\_ib-mgmt\_vlan\_id>> VLAN:
  - 4GB RAM
  - Two CPUs
  - One virtual network interface
2. Start the VM, install VMware 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). Open Data Sources (ODBC) by selecting **Start > Administrative Tools > Data Sources (ODBC)**.
10. Click 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 68**      **SQL Server – Creating an ODBC 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 69** *SQL Server – Setting up SQL Server Authentication*

**Create a New Data Source to SQL Server**

How should SQL Server verify the authenticity of the login ID?

☐ With Integrated Windows authentication.

SPN (Optional):

☒ With SQL Server authentication using a login ID and password entered by the user.

Login ID: vpxuser

Password: .....

☒ Connect to SQL Server to obtain default settings for the additional configuration options.

< Back   Next >   Cancel   Help

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

**Figure 70** *SQL Server – Default Database Selection*

**Create a New Data Source to SQL Server**

☒ Change the default database to:

VCDB

Mirror server:

SPN for mirror server (Optional):

☐ Attach database filename:

☒ Use ANSI quoted identifiers.

☒ Use ANSI nulls, paddings and warnings.

< Back   Next >   Cancel   Help

16. Click **Finish**.

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

**Figure 71**      *SQL Server – ODBC Data Source Test*



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

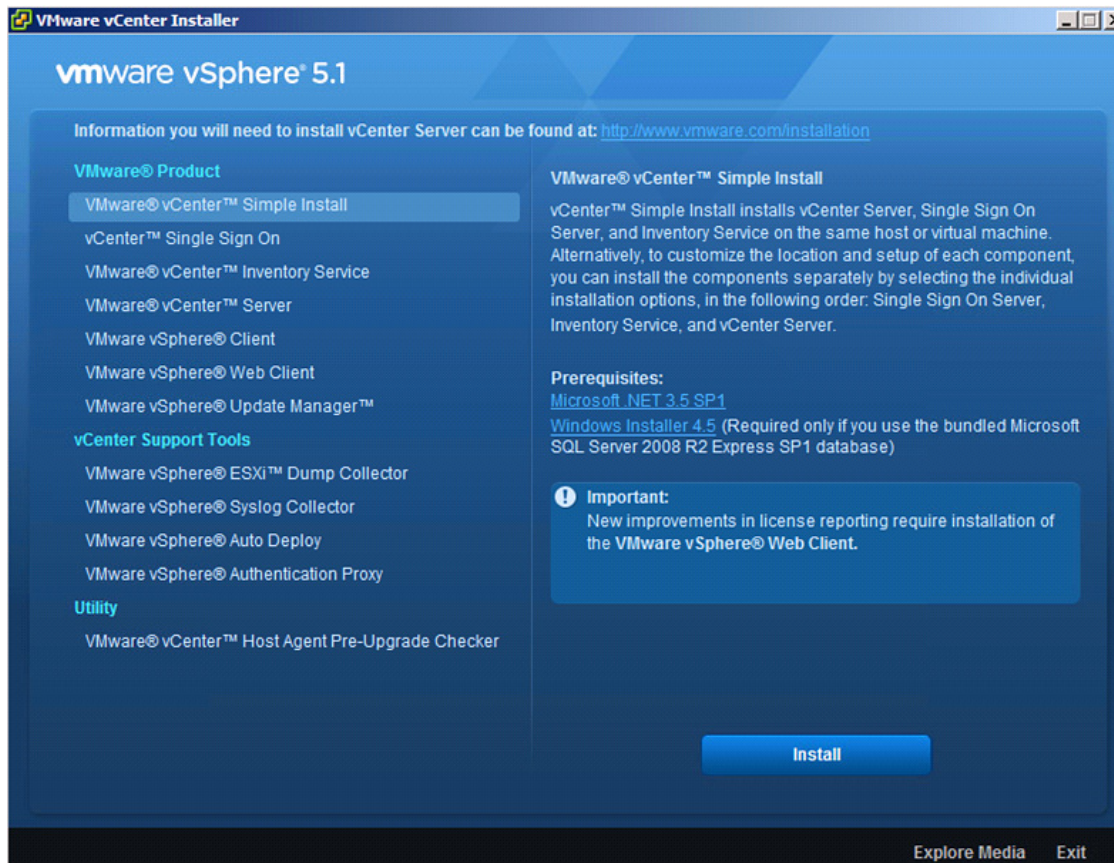
### vCenter Server VM

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 5.1 (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 72** VMware – vSphere 5.1 Installation



5. Click **Yes** at the User Account Control warning.
6. Click **Next** to install vCenter Single Sign On.
7. Click **Next**.
8. Accept the terms of the license agreement and click **Next**.
9. Enter and confirm <<var\_password>> for admin@System-Domain. 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**.
12. Verify the vCenter VM FQDN and click **Next**.
13. Leave Use network service account selected and click **Next**.
14. Click **Next** to choose the default destination folder.
15. Click **Next** to choose the default HTTPS port.
16. Click **Install** to install vCenter Single Sign On.

17. Click **Yes** at the User Account Control warning.
18. Click **Yes** at the User Account Control warning.
19. Enter the vCenter 5.1 license key and click **Next**.
20. Choose Use an Existing Supported Database. Choose VCDB from the Data Source Name list and click **Next**.

**Figure 73** VMware - Selecting Existing Database for vSphere



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

**Figure 74** VMware - Entering Database Credentials

22. Review the warning and click **OK**.
23. Click **Next** to use the SYSTEM Account.
24. Click **Next** to accept the default ports.
25. Choose the appropriate inventory size. Click **Next**.
26. Click **Install**.
27. Click **Finish**.
28. Click **OK** to confirm the installation.
29. Click **Exit** in the VMware vCenter Installer window.
30. Disconnect the VMware vCenter ISO from the vCenter VM.
31. 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 check the check boxes for Turn On vSphere HA and Turn on vSphere DRS. Click **Next**.

**Figure 75** *VMware - Setting up the Cluster*

**New Cluster Wizard**

**Cluster Features**  
What features do you want to enable for this cluster?

**Cluster Features**

- vSphere DRS
  - Power Management
- vSphere HA
  - Virtual Machine Options
  - VM Monitoring
- VMware EVC
- VM Swapfile Location
- Ready to Complete

**Name**

FlexPod\_Management

**Cluster Features**

Select the features you would like to use with this cluster.

☒ **Turn On vSphere HA**

vSphere HA detects failures and provides rapid recovery for the virtual machines running within a cluster. Core functionality includes host and virtual machine monitoring to minimize downtime when heartbeats cannot be detected.

vSphere HA must be turned on to use Fault Tolerance.

☒ **Turn On vSphere DRS**

vSphere DRS enables vCenter Server to manage hosts as an aggregate pool of resources. Cluster resources can be divided into smaller resource pools for users, groups, and virtual machines.

vSphere DRS also enables vCenter Server to manage the assignment of virtual machines to hosts automatically, suggesting placement when virtual machines are powered on, and migrating running virtual machines to balance load and enforce resource allocation policies.

vSphere DRS and VMware EVC should be enabled in the cluster in order to permit placing and migrating VMs with Fault Tolerance turned on, during load balancing.

Help    < Back    Next >    Cancel

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 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](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. Hit 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(1.1a) has been downloaded from [www.cisco.com](http://www.cisco.com) and expanded. This procedure also assumes that VMware vSphere 5.1 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>>](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 76**      **Enabling Serial over LAN for Cisco Nexus 1110-X**



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.

**Figure 77 Cisco Nexus 1110-x Base Configuration**



## 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
Enter control VLAN <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
Enter the domain<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 system configuration dialogue (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the VSA name : <<var_1110x_vsa>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IP address type V4/V6? (V4): 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 (das/rsa) [rsa]: Enter
Number of rsa key bits <768-2048> [1024]: Enter
Enable the http server? (yes/no) [y]: Enter
Configure the ntp server? (yes/no) [n]: y
  
```

NTP server IPv4 address: <<var\_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>>



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

2. Enter the admin password again to confirm: <<var\_password>>.

Enter HA role[primary/secondary]: secondary  
Enter network-uplink type <1-5>: 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>>

3. The Cisco Nexus 1110-X saves the configuration and reboots.

## Set Up the Primary Cisco Nexus 1000V VSM

### Cisco Nexus 1110-X A

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, 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.SV2.1.1a.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.SV2.1.1a.iso /repository/.

```
virtual-service-blade-type new nexus-1000v.4.2.1.SV2.1.1a.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.SV2.1.1a.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
```

**Note**

This domain ID should be different than the VSA domain ID.

```
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_mgmt_ip>>
Enter Management subnet mask: <<var_vsm_mgmt_mask>>
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>
Enter HostName: <<var_vsm_hostname>>
Enter the password for 'admin': <<var_password>>
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:

### Cisco Nexus 1110-X A

1. Enable the secondary VSM.

```
enable secondary
Enter vsb image: [nexus-1000v.4.2.1.SV2.1.1a.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>>
```

2. Enter the admin password <<var\_password>>.
3. 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
```

## 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>>](http://<<var_vsm_mgmt_ip>>).
2. Right-click the [cross\\_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib](#) hyperlink and choose Save target as.
3. Save the file as [cross\\_cisco-vem-v152-4.2.1.2.1.1a.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, click 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, select the `VEM` folder.
10. Click the fourth button at the top and choose **Upload File**.
11. Navigate to the `cross_cisco-vem-v152-4.2.1.2.1.1a.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-v152-4.2.1.2.1.1a.0-3.1.1.vib
```

**Figure 78** Installing VEM on the ESXi Servers

```

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.62 -u root -p NetApp!23 software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  UIBs Installed: Cisco_bootbank_cisco-vem-v152-esx_4.2.1.2.1.1a.0-3.1.1
  UIBs Removed:
  UIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.175.101 -u root -p NetApp!23 software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v152-4.2.1.2.1.1a.0-3.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  UIBs Installed: Cisco_bootbank_cisco-vem-v152-esx_4.2.1.2.1.1a.0-3.1.1
  UIBs Removed:
  UIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>_

```

## 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

```

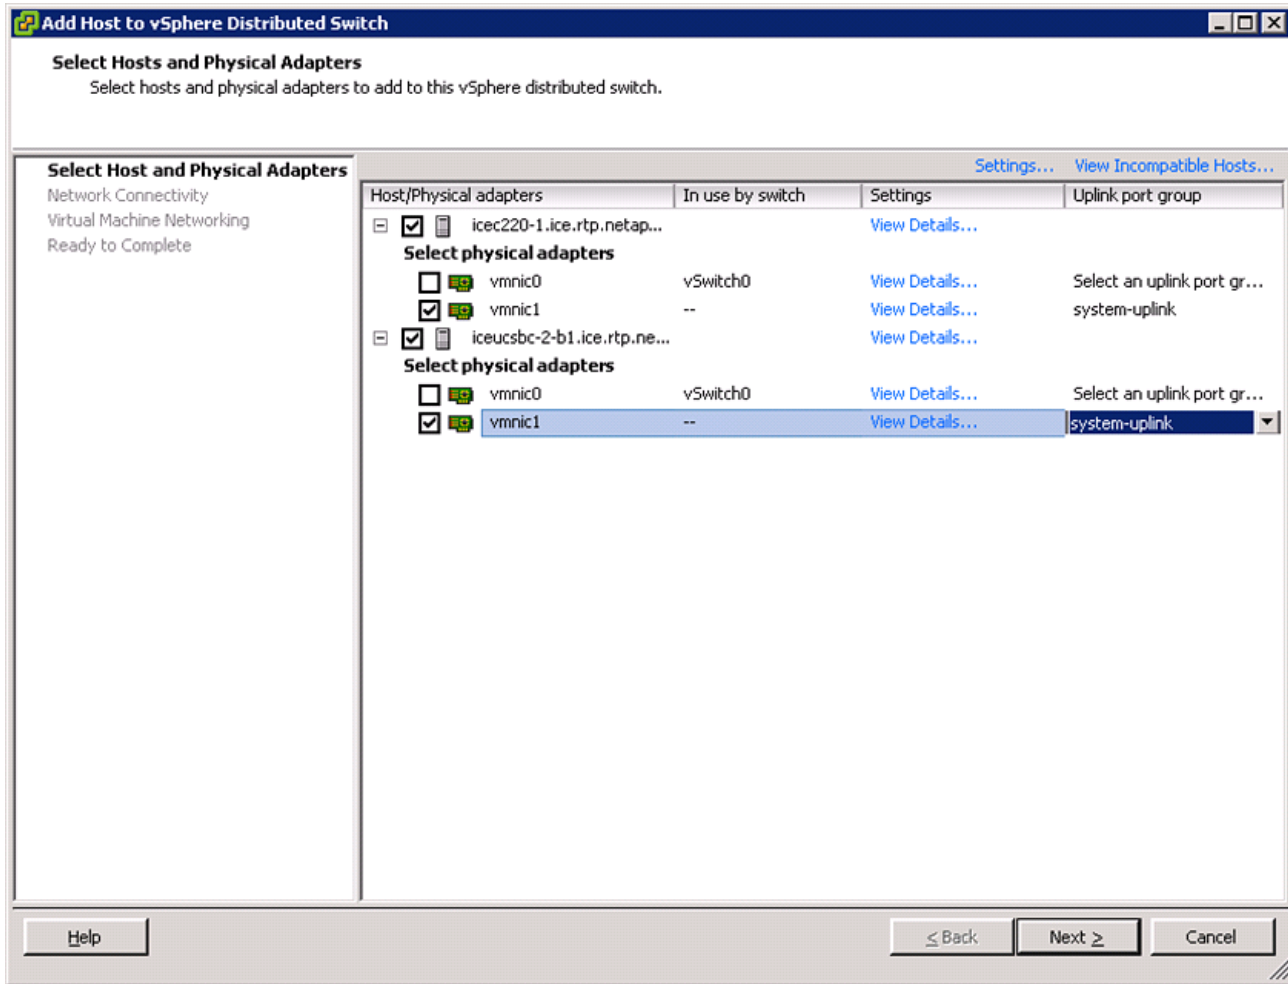
## Migrate Networking Components for ESXi Hosts to Cisco Nexus 1000V

### 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**.

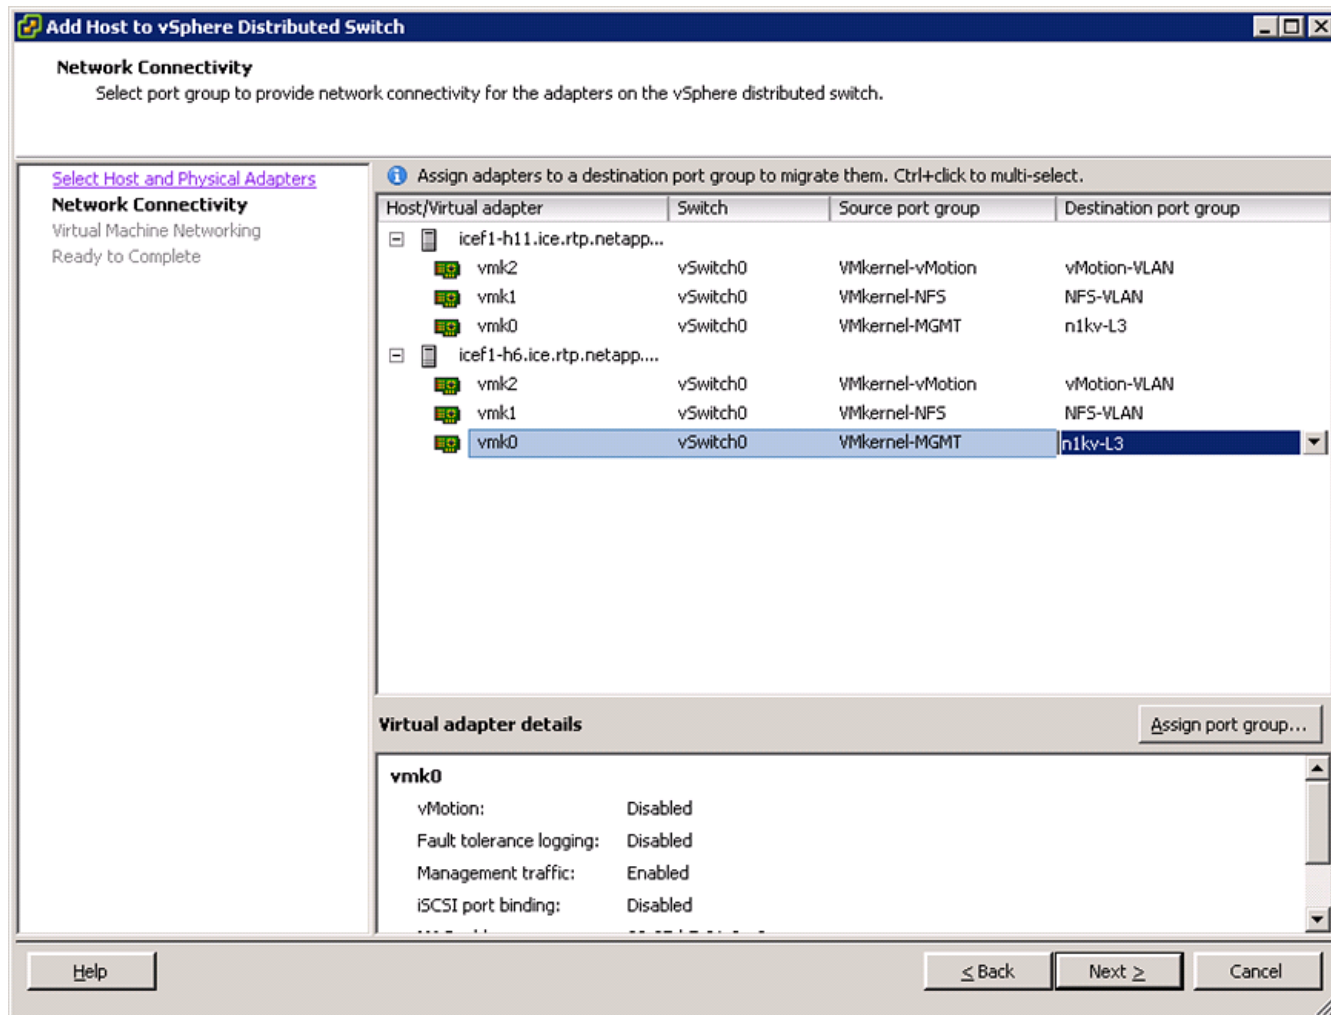
**Figure 79** VMware – Selecting vmnic1 as system-uplink for N1Kv



- 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**.

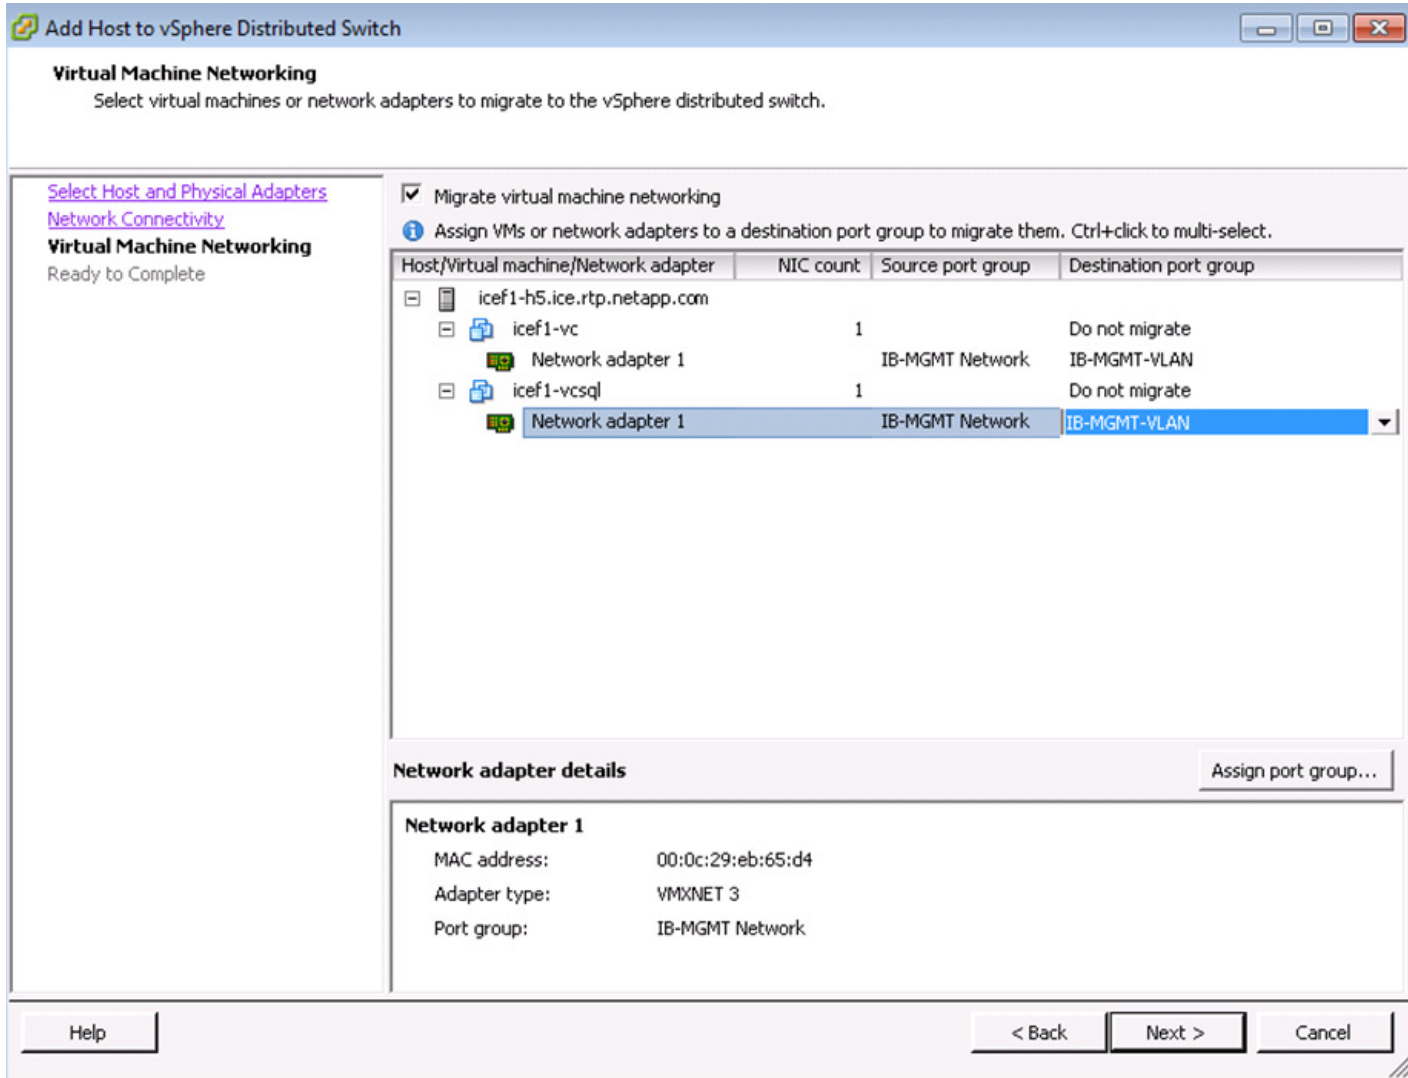


**Figure 80** VMware – Selecting n1kv-L3 as the Management Port-Group



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

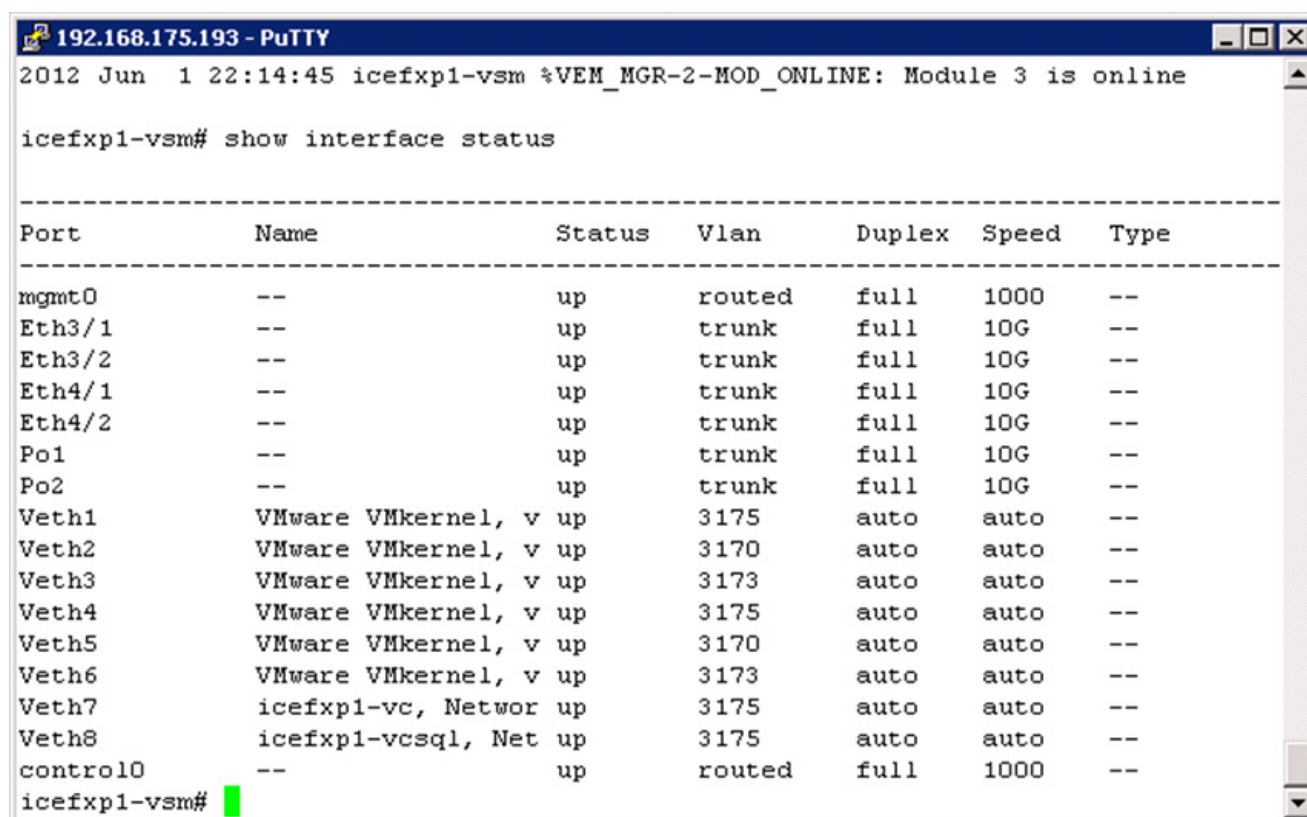
**Figure 81** VMware – VM Network Migration to N1Kv



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, click **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 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 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 82** VSM - "Show Interface Status"



```

2012 Jun  1 22:14:45 icefxp1-vsm %VEM_MGR-2-MOD_ONLINE: Module 3 is online

icefxp1-vsm# 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	3170	auto	auto	--
Veth3	VMware VMkernel, v	up	3173	auto	auto	--
Veth4	VMware VMkernel, v	up	3175	auto	auto	--
Veth5	VMware VMkernel, v	up	3170	auto	auto	--
Veth6	VMware VMkernel, v	up	3173	auto	auto	--
Veth7	icefxp1-vc, Networ	up	3175	auto	auto	--
Veth8	icefxp1-vcsql, Net	up	3175	auto	auto	--
control0	--	up	routed	full	1000	--

```

icefxp1-vsm#

```

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

Figure 83 VSM - "Show Module"

```
icefl-vsm
icefl-vsm(config)# show module
Mod  Ports  Module-Type                Model                Status
---  ---
1     0       Virtual Supervisor Module  Nexus1000V           ha-standby
2     0       Virtual Supervisor Module  Nexus1000V           active *
3    248     Virtual Ethernet Module    NA                    ok
4    248     Virtual Ethernet Module    NA                    ok

Mod  Sw                Hw
---  ---
1     4.2 (1)SV2 (1.1a)  0.0
2     4.2 (1)SV2 (1.1a)  0.0
3     4.2 (1)SV2 (1.1a)  VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
4     4.2 (1)SV2 (1.1a)  VMware ESXi 5.1.0 Releasebuild-838463 (3.1)

Mod  MAC-Address(es)                Serial-Num
---  ---
1     00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2     00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
3     02-00-0c-00-03-00 to 02-00-0c-00-03-80  NA
4     02-00-0c-00-04-00 to 02-00-0c-00-04-80  NA

Mod  Server-IP          Server-UUID          Server-Name
--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) 4.1 Deployment Procedure

#### VSC 4.1 Pre Installation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.1.2:

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

## Install VSC 4.1


To install the VSC 4.1 software, follow these steps:

1. Using the instructions in section “Build Microsoft SQL Server VM,” 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 4.1 at: <http://support.netapp.com/NOW/cgi-bin/software/?product=Virtual+Storage+Console&platform=VMware+vSphere> from the [NetApp Support site](#).
4. Right-click the file downloaded in step 3 and choose Run As Administrator.
5. Click **Yes** at the User Access Control warning.
6. On the Installation wizard Welcome page, click **Next**.
7. Choose the backup and recovery capability. Click **Next**.




**Note** The backup and recovery capability requires an additional license.

**Figure 84** NetApp VSC - Backup and Recovery Capability




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

**Figure 85**      *NetApp VSC – Default Installation Location*



9. Click **Install**.
10. Click **Finish**.

**Figure 86**      **NetApp VSC – Start Installation**



## 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 87**      **NetApp VSC – VMware Plug-in Registration**

vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.

Plugin service information

Host name or IP Address: 192.168.175.191

vCenter Server information

Host name or IP Address: 192.168.175.188

Port: 443

User name: ice\ycef1-admin

User password: ••••••••

Register

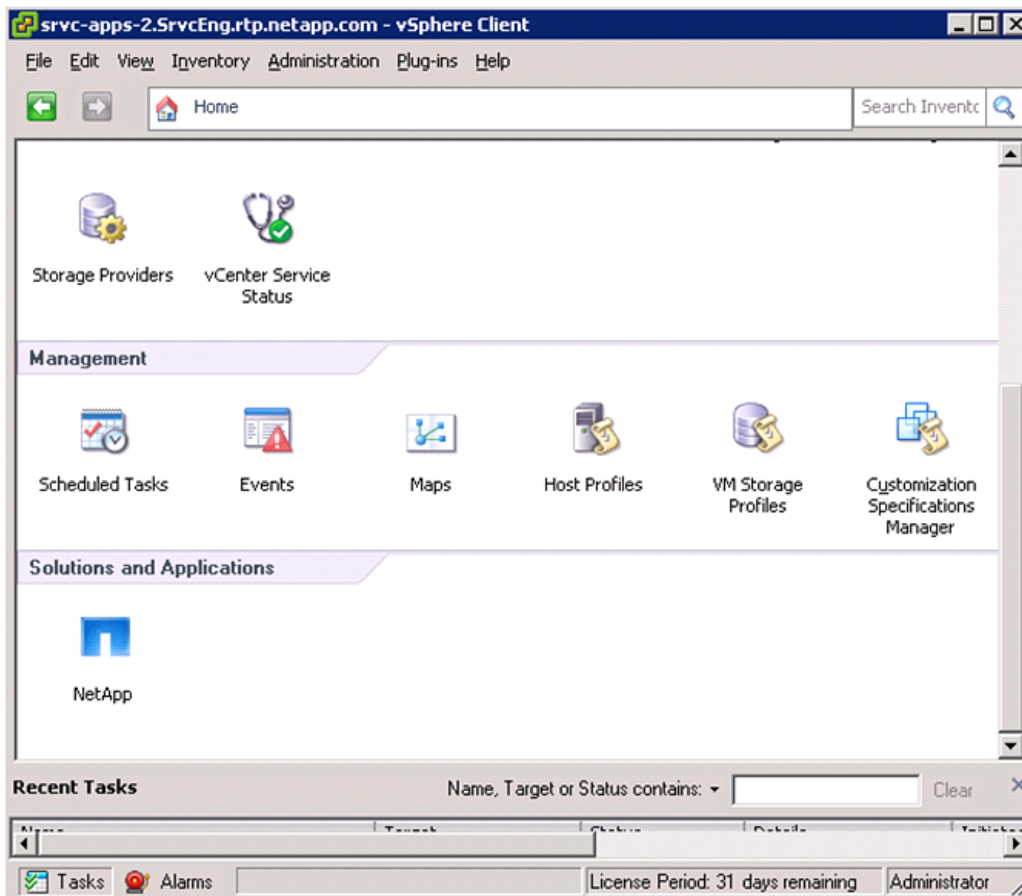
## Discover and Add Storage Resources

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. Click the **Home** tab in the left side of the vSphere Client window.
3. Under Solutions and Applications, click the **NetApp** icon.

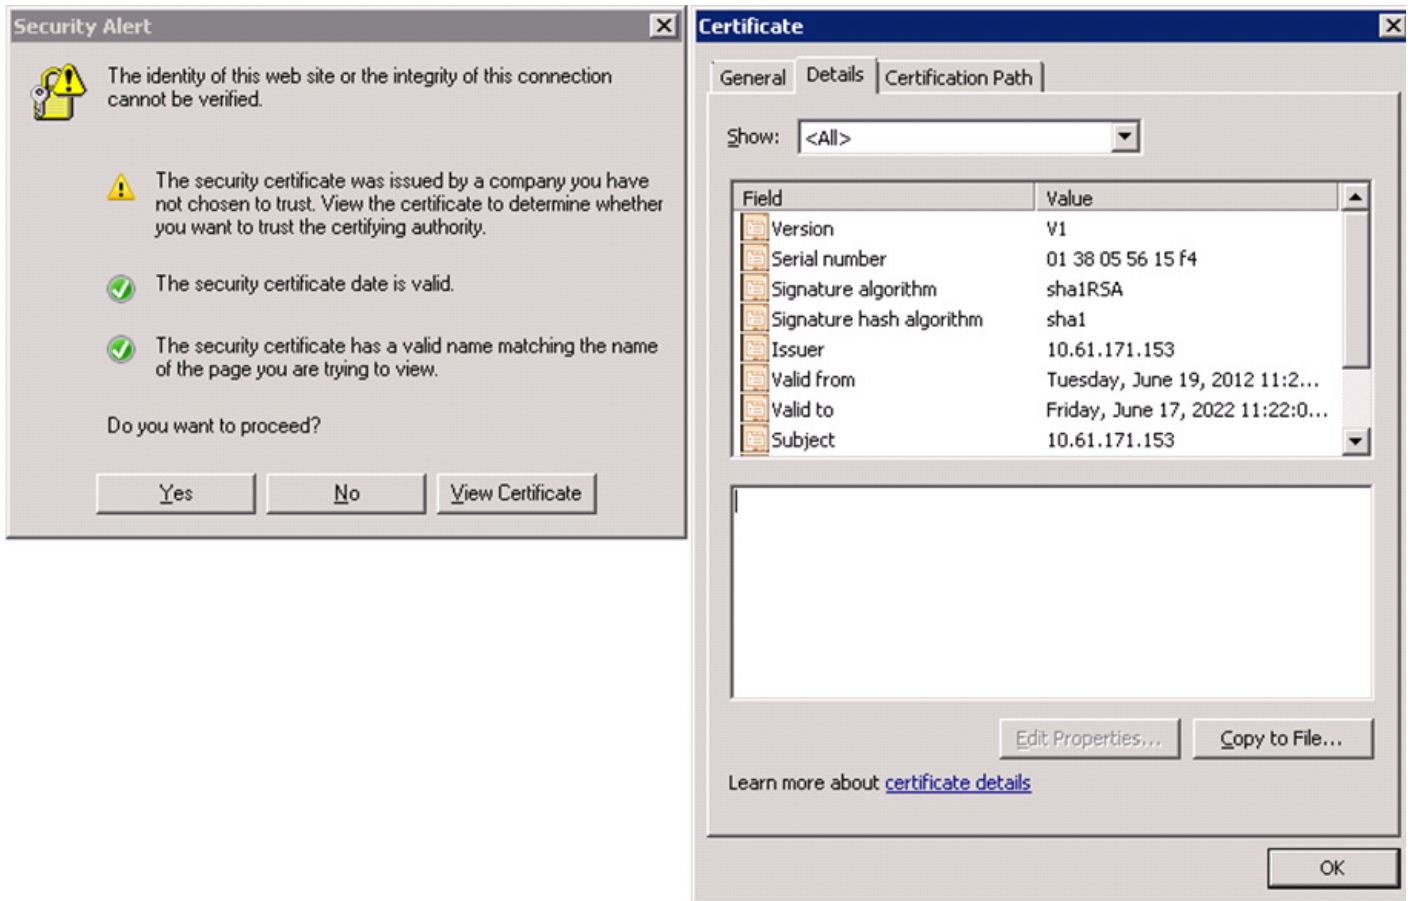


**Figure 88**      **NetApp VSC - Configuration in vSphere Client**



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

**Figure 89**      **NetApp VSC - Client Certificate Warning**



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

**Figure 90 NetApp VSC - Select Monitoring and Host Configuration**

The screenshot shows the NetApp VSC interface within a vSphere Client window. The left sidebar contains navigation options: Overview, Storage Details - SAN, Storage Details - NAS, Data Collection, Tools, and Discovery Status. The main area is divided into two sections: Storage Controllers and ESX Hosts.

**Storage Controllers**

Controller	IP Address	Version	Status	Free Capacity	VAAI Capable	Supported Protocols
Unknown (3 Unknown)						
Controller: -unknown-n-	192.168.175.160		Authenti...	0.00B (0%)	Unknown	Unknown
Controller: -unknown-n- (192.168.170.151)			Unknown	0.00B (0%)	Unknown	Unknown
Controller: -unknown-n- (192.168.170.152)			Unknown	0.00B (0%)	Unknown	Unknown

**ESX Hosts**

Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
icef1-h12.ice.rtp.netapp.com	192.168.175.101	5.1.0	Alert	Normal	Normal	Alert
icef1-h5.ice.rtp.netapp.com	192.168.175.62	5.1.0	Alert	Normal	Normal	Alert

Last update: Fri Feb 22 12:12:00 GMT-500 2013

**Recent Tasks**

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
------	--------	--------	---------	--------------	----------------	----------------------	------------	----------------

6. In the list of storage controllers, right-click the first controller listed and choose Modify Credentials.
7. Enter the storage cluster management IP address in the Management IP address field. Enter admin for the User name, and the admin password for the Password. Make sure that Use SSL is selected. Click **OK**.
8. Click **OK** to accept the controller privileges.

Figure 91 NetApp VSC - Setting up Storage Controller

The screenshot displays the NetApp VSC interface within the vSphere Client. The left sidebar shows the 'Monitoring and Host Configuration' menu with options like Overview, Storage Details - SAN, Storage Details - NAS, Data Collection, Tools, and Discovery Status. The main area is divided into two sections: 'Storage Controllers' and 'ESX Hosts'.

**Storage Controllers**

Controller	IP Address	Version	Status	Free Capacity	VAAI Capable	Supported Protocols
Cluster: icef1-stcl (1 Vserver)						
Vserver: Infra_Vserver		8.1.2 Clu...	Normal	850.61GB (9...	Supported	NFS, FC/FCoE
Cluster: icef1-stcl	192.168.175.150	8.1.2 Clu...	Normal	19.76TB (95%)	Supported	NFS, FC/FCoE

**ESX Hosts**

Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
icef1-h12.ice.rtp.netapp.com	192.168.175.101	5.1.0	Alert	Normal	Normal	Alert
icef1-h5.ice.rtp.netapp.com	192.168.175.62	5.1.0	Alert	Normal	Normal	Alert

Last update: Fri Feb 22 12:22:45 GMT-500 2013

**Recent Tasks**

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
NetApp Storage Discov...	FlexPod_DC_1	Completed	[icef1-stcl-...	ICE icef1-ad...	ICEF1-VC.ice.r...	2/22/2013 12:22:39 ...	2/22/2013 12:22:39 ...	2/22/2013 12:22:41 ...
NetApp Storage Discov...	FlexPod_DC_1	Completed	[icef1-stcl-...	ICE icef1-ad...	ICEF1-VC.ice.r...	2/22/2013 12:22:39 ...	2/22/2013 12:22:39 ...	2/22/2013 12:22:41 ...
NetApp Storage Discov...	FlexPod_DC_1	Completed	[Infra_Vser...	ICE icef1-ad...	ICEF1-VC.ice.r...	2/22/2013 12:22:39 ...	2/22/2013 12:22:39 ...	2/22/2013 12:22:42 ...

Tasks: Alarms License Period: 195 days remaining ICE|icef1-admin

## 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 92**      **NetApp VSC - Selecting Optimal Storage Settings**

The screenshot displays the NetApp VSC vSphere Client interface. The left sidebar shows the 'Monitoring and Host Configuration' menu with options like Overview, Storage Details - SAN, Storage Details - NAS, Data Collection, Tools, and Discovery Status. The main content area is divided into two sections: 'Storage Controllers' and 'ESX Hosts'.

**Storage Controllers**

Controller	IP Address	Version	Status	Free Capacity	VAAI Capable	Supported Protocols
<b>Clusteracef1-stcl (1 Vserver)</b>						
Vserver: infra_Vserver		8.1.2 Clu...	Normal	850.61GB (9...	Supported	NFS, FC/FCoE
Cluster: icef1-stcl	192.168.175.150	8.1.2 Clu...	Normal	19.76TB (95%)	Supported	NFS, FC/FCoE

**ESX Hosts**

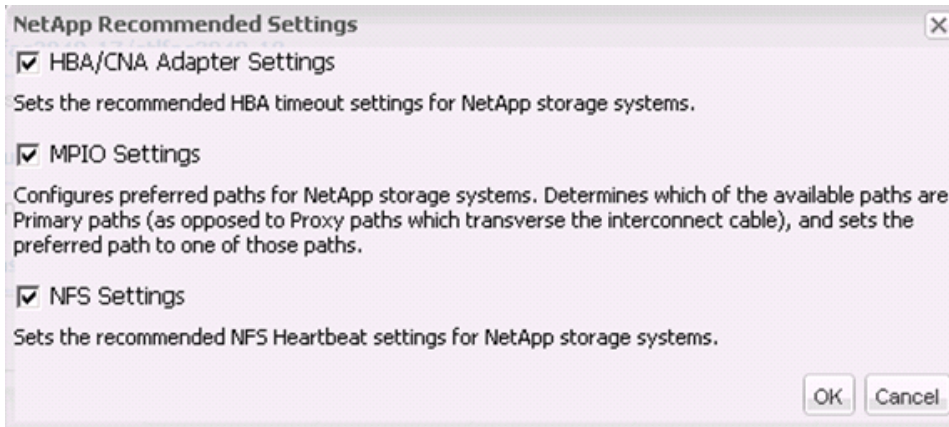
Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
icef1-h12.ice.rtp.netapp.com	192.168.175.101	5.1.0	Alert	Normal	Normal	Alert
icef1-h5.ice.rtp.netapp.com	192.168.175.62	5.1.0	Alert	Normal	Normal	Alert

A context menu is open over the second host, showing options: 'Set Recommended Values...', 'Show Details...', and 'Skip Host...'. The status bar at the bottom indicates 'Last update: Fri Feb 22 12:26:48 GMT-500 2013'.

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 93**      **NetApp VSC – Selecting Individual Recommended Settings**



**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.

## VSC 4.1 Backup and Recovery

### Adding Storage Systems to the 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 add the storage systems that contain the datastores and virtual machines for which you are creating backups.



**Note**

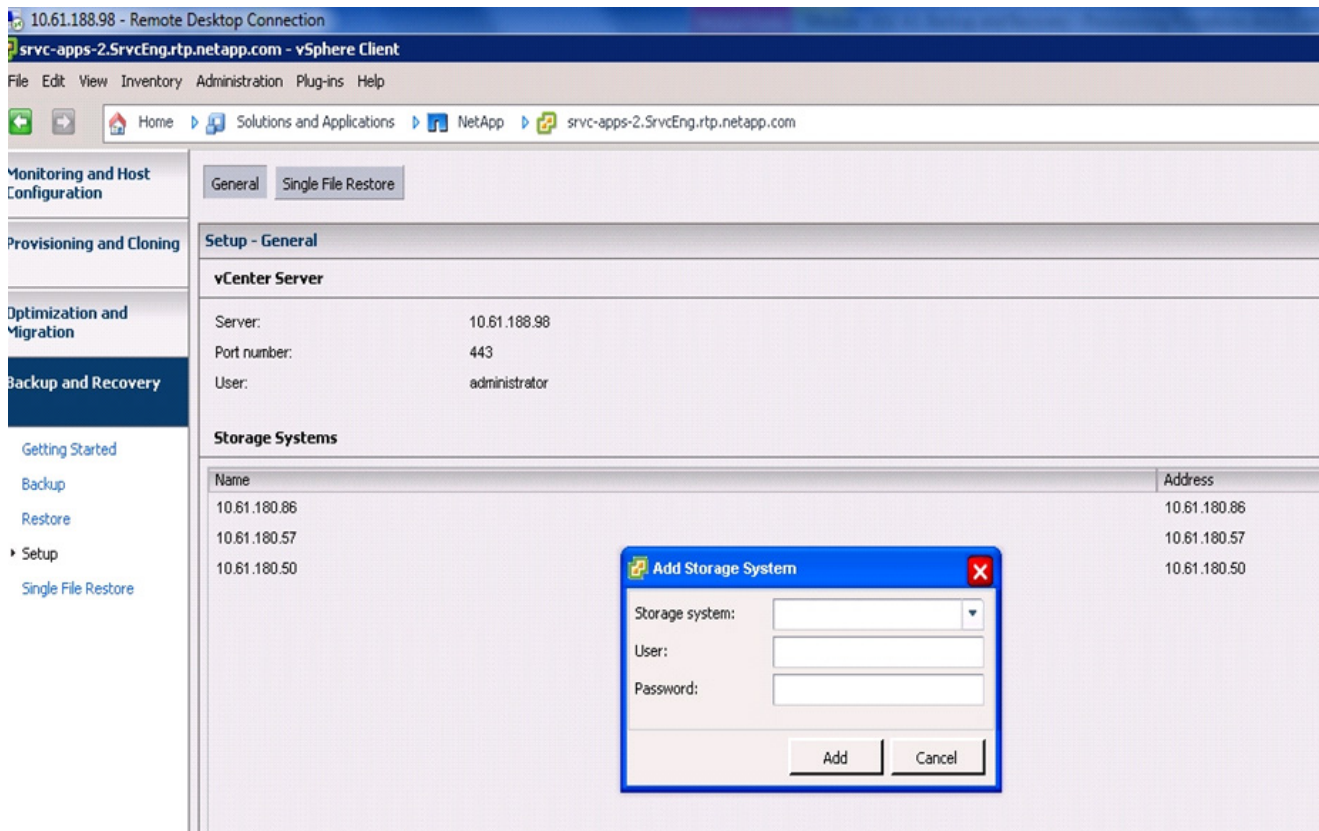
The Backup and Recovery capability does not use the user credentials from the Monitoring and Host Configuration capability.

Follow these steps to add the storage systems to the Backup and Recovery capability:

1. Click **Backup and Recovery** and then click **Setup**.
2. Click **Add**. The Add Storage System dialog box appears.
3. Enter the DNS name or IP address and the user credentials of the storage controller 1.
4. Click **Add** to add the storage cluster.
5. Repeat this process for storage controller 2.



**Figure 94**      **NetApp VSC - Adding Storage System to Backup and Recovery**



## Backup and Recovery Configuration

The following steps detail the procedure to configure a backup job for a datastore.

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

Figure 95 NetApp VSC - Configuring Backup Job

**Backup Wizard**

**Job**  
Specify a name for the backup job that you want to create.

**Job Name**

Virtual Entities  
Spanned Entities  
Scripts  
Schedule  
User Credentials  
Backup Retention  
Ready to Complete


Name: VSC\_backup

Description: VM backup

Options

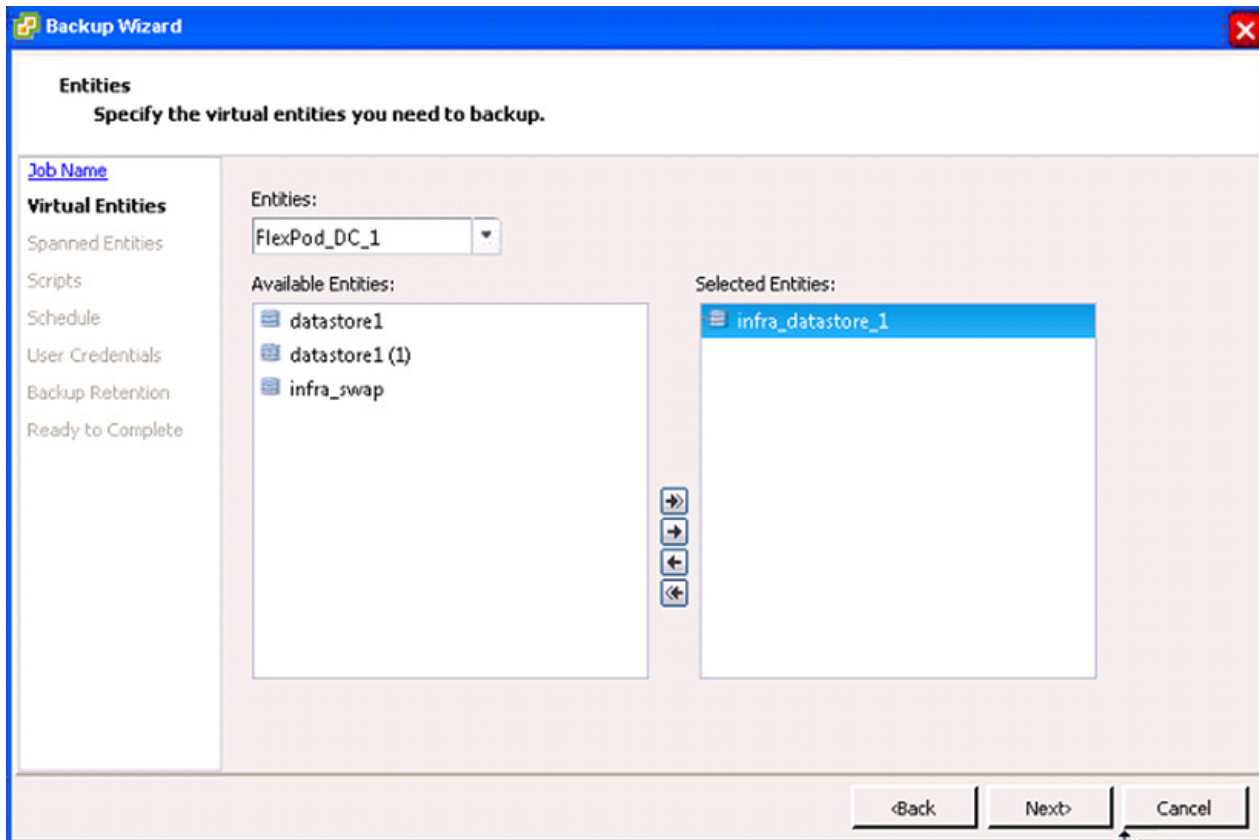
- ☐ Initiate SnapMirror update
- ☒ Perform VMware consistency snapshot
- ☐ Include datastores with independent disks

<Back   Next>   Cancel

3. Enter 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**      **NetApp VSC - Selecting Backup Datastore**



7. choose one or more backup scripts if available and click **Next**.
8. choose the hourly, daily, weekly, or monthly schedule that you want for this backup job and click **Next**.

**Figure 97**      **NetApp VSC - Setting a Backup Schedule**

**Backup Wizard**

**Schedule**  
You can specify an hourly, daily, weekly, monthly or no schedule at all for the backup job.

[Job Name](#)  
[Virtual Entities](#)  
Spanned Entities  
[Scripts](#)  
**Schedule**  
User Credentials  
Backup Retention  
Ready to Complete

Perform this backup

☒ Hourly  
☐ Daily  
☐ Weekly  
☐ Monthly  
☐ One time only

Hourly schedule details

Backups will be performed

Every: 1 hour  
At: 11: 49 AM  
Starting: 07/08/2012

<Back    Next>    Cancel

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 98 NetApp VSC - Backup Retention**

**Backup Wizard**

**Retention and Alerts**  
You can specify backup retention based on maximum days, maximum no of backups or backup indefinitely.

[Job Name](#)  
[Virtual Entities](#)  
[Spanned Entities](#)  
[Scripts](#)  
[Schedule](#)  
[User Credentials](#)  
**Backup Retention**  
Ready to Complete

**Retention**

☒ A maximum of days: 1

☐ A maximum of backups: 1

☐ Never expires

**Email alerts**

Source email address: test1@example.com

Destination email address (s): test2@example.com

SMTP host: smtp.example.com

Notify on: Always

[Send test email](#)

<Back    Next>    Cancel

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 99 NetApp VSC - Backup Summary

12. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by typing the command:

```
volume modify -volume infra_datastore_1 -snapshot-policy none
```

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

```
volume snapshot show -volume infra_datastore_1
volume snapshot delete -volume infra_datastore_1 <snapshot name>
```

## OnCommand Unified Manager 5.1

### 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. Make sure the Infra\_Vserver is selected in Vserver drop-down menu and click **Next**.
4. Choose VMFS as the Datastore type and click **Next**.

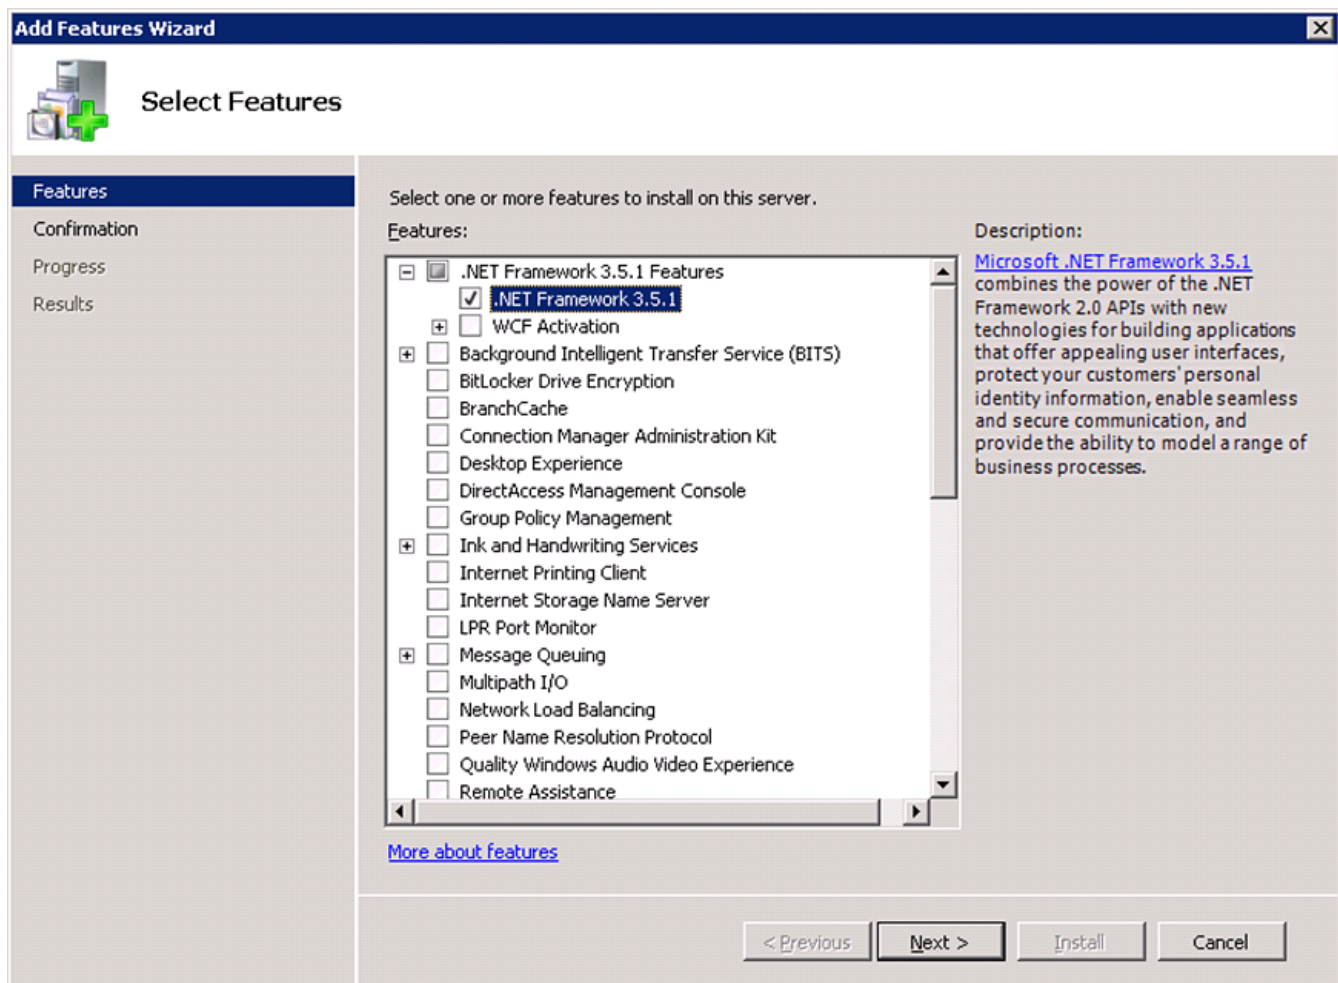
5. Choose FCP as the Protocol type, set the Size to 100, enter the datastore name as RDM\_Map, check the check box to create new volume container, choose aggr01 as the Aggregate, check the Thin Provision check box, and click **Next**.
6. Verify settings and click **Apply**.

## Install .NET Framework 3.5.1 Feature

From the Virtual Storage Console (VSC) and OnCommand VM:

1. Log in to the VSC and OnCommand VM as the FlexPod admin and open Server Manager.
2. Click **Features** and click **Add Features**.
3. Expand .NET Framework 3.5.1 Features and choose only .NET Framework 3.5.1.

**Figure 100** OnCommand VM .NET Setup



4. Click **Next**.
5. Click **Install**.
6. Click **Close**.

## 7. Close **Server Manager**.

## Install SnapDrive 6.4.2

Follow these steps to install SnapDrive® 6.4.2:

1. Download SnapDrive 6.4.2 from the [NetApp Support Site](#).
2. 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.
3. Click **Next** in the Welcome page of the SnapDrive installation wizard.
4. If this is a new SnapDrive installation, read and accept the license agreement. Click **Next**.
5. If this is a SnapDrive upgrade, choose Modify/Upgrade in the Program Maintenance page. Click **Next**.
6. Choose “Per Storage System” as the license type. Click **Next**.



### Note

- In the case of upgrading SnapDrive, the license information will already be populated.
  - In the case of selecting storage system licensing, SnapDrive can be installed without entering a license key. SnapDrive operations can be executed only on storage systems that have a SnapDrive or SnapManager license installed.
  - In the case of clustered Data ONTAP 8.1–based systems, the storage system licensing for SnapDrive is bundled with the other SnapManager product licenses. They are now a single license called the SnapManager\_suite license.
7. In the Customer Information page, type the user name and organization name. Click **Next**.
  8. The Destination Folder page prompts for a directory in which to install SnapDrive on the host. For new installations, by default this directory is C:\Program Files\NetApp\SnapDrive\. To accept the default, click **Next**.
  9. Check the Enable VirtualCenter or ESX Server Settings check box. Enter the IP address, user name, and password for the vCenter Server and click **Next**.



### Note

Selecting Enable VirtualCenter or ESX Server Settings enables SnapDrive to use RDM pass-through LUNs. Choose this option to use RDM pass-through disks. By default, this option is not selected.

**Figure 101 SnapDrive - Installation Wizard**

10. Check the Enable SnapManager for Virtual Infrastructure Configuration Details check box. Enter the IP address of the VSC and OnCommand Server, and accept the default port. Click **Next**.

**Figure 102 SnapDrive - Enable SnapManager for Virtual Infrastructure**

11. Enter your credentials or follow these steps to select a user account:



- a. In the Enter object name to select box, enter the local machine administrator in Domain name\user name format. Click **Add**.
- b. Click **Check Names**.
- c. Click **OK**.
- d. Enter the Administrator password.
- e. Click **Next**.
- f. Click **OK**.

**Figure 103**      **SnapDrive - Service Credentials**

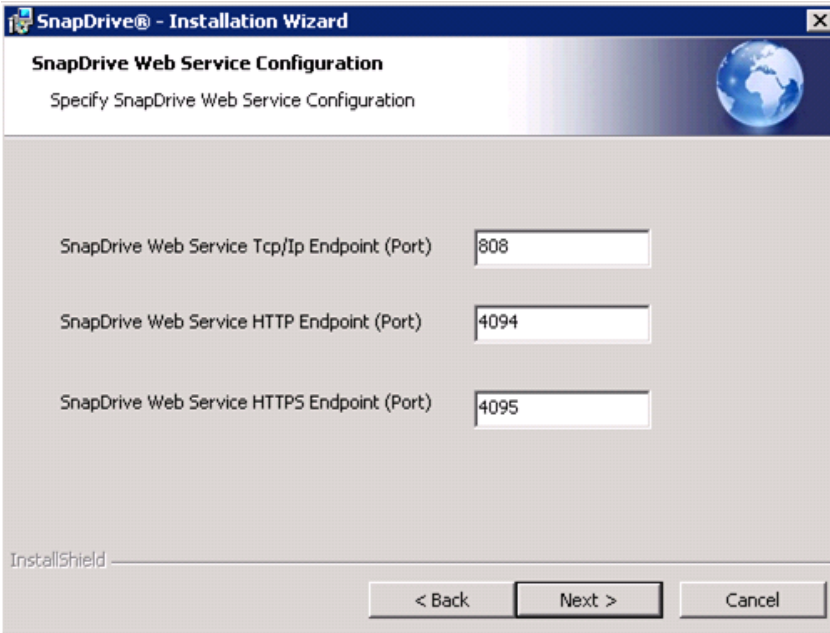


**Note**      The specified account must be a member of the local administrators' group of this system.

12. In the SnapDrive Web Service Configuration page, keep the default ports unless any of them are already being used exclusively by another service. Click **Next**.



**Figure 104**      **SnapDrive - Web Service Configuration**

The image shows a screenshot of the 'SnapDrive® - Installation Wizard' window, specifically the 'SnapDrive Web Service Configuration' step. The window has a title bar with the SnapDrive logo and a close button. Below the title bar, the text 'SnapDrive Web Service Configuration' is displayed, followed by the instruction 'Specify SnapDrive Web Service Configuration'. To the right of this text is a small globe icon. The main area of the window contains three input fields, each with a label and a text box. The first field is labeled 'SnapDrive Web Service Tcp/Ip Endpoint (Port)' and contains the value '808'. The second field is labeled 'SnapDrive Web Service HTTP Endpoint (Port)' and contains the value '4094'. The third field is labeled 'SnapDrive Web Service HTTPS Endpoint (Port)' and contains the value '4095'. At the bottom left of the window, the 'InstallShield' logo is visible. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a black border.

13. In the Transport Protocol Default Setting screen:
  - a. choose Enable Transport Protocol Settings.
  - b. choose HTTPS as the transport protocol.
  - c. Enter the user name (vsadmin) and password for the Infra\_Vserver vserver.
  - d. Verify that port ID is set to 443 and click **Next**.

**Figure 105**      **SnapDrive - Transport Protocol Default Setting**

**SnapDrive - InstallShield Wizard**

**Transport Protocol Default Setting**  
Specify Default Transport Setting for Storage System(s)

☒ Enable Transport Protocol Settings

☐ RPC  
☐ HTTP  
☒ HTTPS

Specify the user name and password for the HTTP/HTTPS Protocol selection.

User Name:

Password:

Port ID:

InstallShield

< Back    Next >    Exit

14. Click **Next > Next > Install > Finish**.
15. From the Start menu, open SnapDrive.
16. In the left pane, expand the local machine and choose Disks.
17. In the right pane, choose Create Disk.
18. In the create disk Wizard Window, click **Next**.
19. In the storage system name field, enter the Infra\_Vserver management IP address, and click **Add**.
20. In the list that appears, choose OnCommandDB.
21. Enter OnCommandDB for the LUN Name and click **Next**.

**Figure 106 SnapDrive - Create Disk Wizard**

**Create Disk Wizard**

**Provide a Storage System Name, LUN Path and Name**  
Enter the Storage System Name and LUN path information below.

Select an existing Storage System, or enter a new Storage System name and press "Add".

Storage System Name: 192.168.175.160 Add

Name	Type
infra_datastore_1	volume
infra_swap	volume
OnCommandDB	volume
RDM_Map	volume

LUN Path: /vol/OnCommandDB

LUN Name: OnCommandDB

LUN Description:

< Back Next > Cancel

22. Make sure the LUN type is set to Dedicated and click **Next**.

23. Assign drive letter O and set LUN size to 100GB. Click **Next**.

**Figure 107 SnapDrive - LUN Properties**

**Create Disk Wizard**

**Select LUN Properties**  
Provide the drive letter and the size of the LUN to create

**Drive Parameters**

☒ Assign a Drive Letter: O

☐ Use a Volume Mount Point:

☐ Do not assign a Drive letter or Volume Mount Point

**Snapshot Copies**

Do you want to limit the maximum disk size to accommodate at least one snapshot on the volume?

☒ Limit ☐ Do not limit

**LUN Size**

Maximum: 16349.9 GB

Minimum: 64 MB

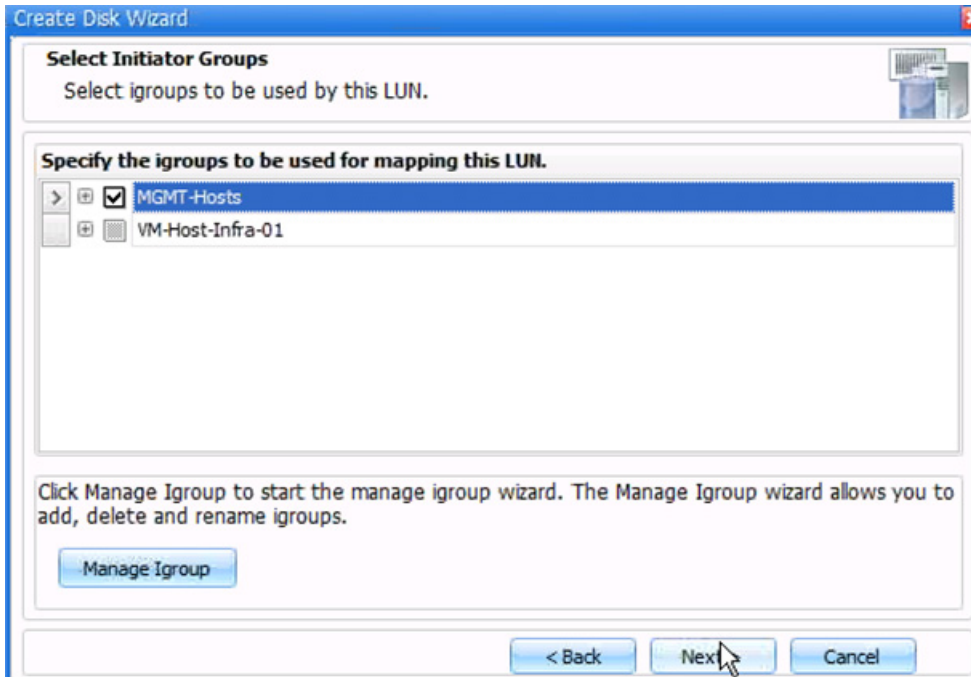
LUN Size: 100.0 GB

< Back Next > Cancel

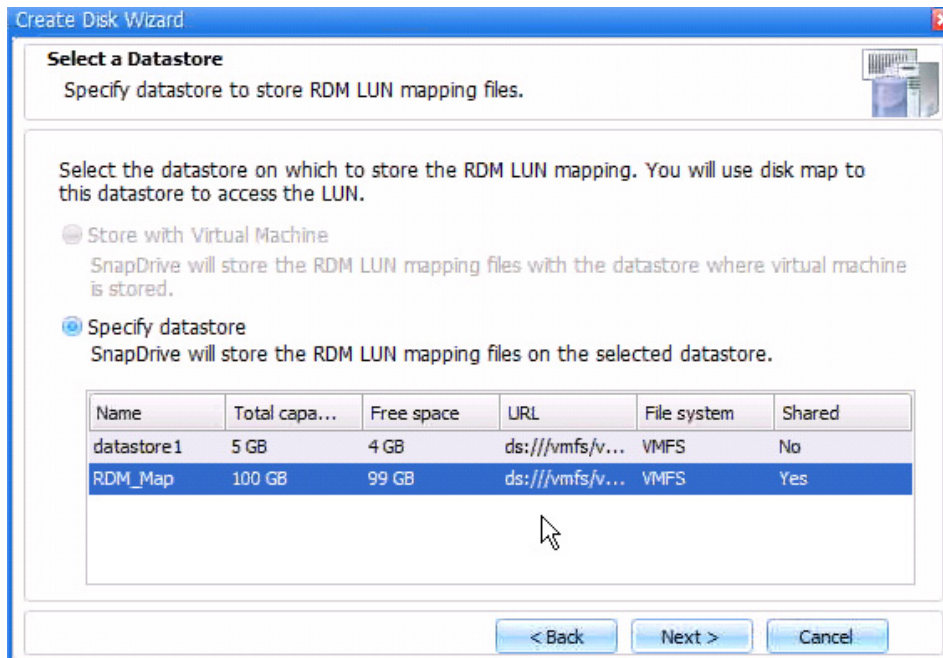
24. Choose all initiators on the Initiator List, and click **Next**.

25. Choose manual as the Initiator group management, and click **Next**.
26. Choose the MGMT-Hosts igroup, and click **Next**.

**Figure 108**      *SnapDrive - Select Initiator Groups*



27. Choose the RDM\_Map Datastore in the Select Datastore section. Click **Next**.

**Figure 109 Snapdrive - Storing RDM LUN Mapping**

28. Click **Finish** to create the disk.

29. 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, click [here](#).
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.
4. The DataFabric Manager Server license key
5. Navigate to the path or directory containing the downloaded file and launch the file.
6. In the Security Warning message, click **Yes** to start the installation.
7. In the Welcome screen, click **Next**.
8. Accept the AutoSupport notice and click **Next**.
9. Identify whether the OnCommand Unified Manager instance should manage systems with clustered Data ONTAP or 7-Mode and click **Next**.



### Note

- For a 7-Mode environment, either the Express edition or the Standard edition of the software is available.
- For a clustered Data ONTAP environment, only the Standard edition of the software is available.
- If the infrastructure has both 7-Mode and clustered Data ONTAP systems, two OnCommand instances are needed to manage the respective 7-Mode or clustered Data ONTAP systems.

10. Enter the 14-character license key when prompted and click **Next**.
11. 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.

12. Follow the remaining setup prompts to complete the installation.

From an MS-DOS command prompt, perform the following steps as an administrator:

1. 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
```

2. 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:\
```

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

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

4. 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 multi-word entries must be encased in double quotation marks, for example, "North Carolina."

5. Turn off automatic discovery.

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

6. 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.

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

```
dfm service start http
```

8. 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 -X <<var_password>> default
```

9. 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
```

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

```
dfm host add <<var_clustername>>
```

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

```
dfm host set <<var_clustername>> hostlogin=admin
dfm host set <<var_clustername>> hostPassword=<<var_password>>
```

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

```
dfm host list
dfm host get <<var_clustername>>
```

13. 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_clustername>>
```

14. (optional) Configure an SNMP trap host.

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

15. 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
```

16. Create a manual backup.

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

17. 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
```

18. To open Windows Firewall with Advanced Security, click Start > Administrative Tools > Windows Firewall with Advanced Security.

19. Choose Inbound Rules.
20. Click New Rule.
21. Choose Port and click Next.
22. Leave TCP selected and enter 8443 in the Specific local ports text box. Click **Next**.
23. Click **Next**.
24. Click **Next**.
25. Name the rule OnCommand Console External Access and click **Finish**.
26. Click **New Rule**.
27. Choose Port and click **Next**.
28. Choose UDP and enter 162 in the Specific local ports text box. Click **Next**.
29. Click **Next**.
30. Click **Next**.
31. Name the rule OnCommand SNMP Trap and click **Finish**.
32. Close Windows Firewall with Advanced Security.

## NetApp NFS Plug-In 1.0 for VMware VAAI

### Enable VMware vStorage for NFS in Clustered Data ONTAP

To enable VMware vStorage for NFS in clustered Data ONTAP, complete the following steps:

1. From an SSH session to the storage cluster management address, log in with the admin user name and password.
2. Enable vStorage on the Vserver.

```
vserver nfs modify -vserver Infra_Vserver -vstorage enabled
```

3. Verify that the export policy rules are set up correctly.

```
vserver export-policy rule show -vserver Infra_Vserver
```

4. The access protocol for the FlexPod policy name should be NFS. If the access protocol is not NFS for a given rule index, run the following command to set NFS as the access protocol:

```
vserver export-policy rule modify -vserver Infra_Vserver -policyname FlexPod  
-ruleindex <<var_rule_index>> -protocol nfs
```

## Install NetApp NFS Plug-In for VMware VAAI

To install the NetApp NFS plug-in for VMware vStorage APIs for Array Integration (VAAI), 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 4.1 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\. However, VSC 4.1 for VMware vSphere lets you change this directory. For example, if you are using the default installation directory, the path to the NetAppNasPlugin.vib file is the following: C:\Program Files\Virtual Storage Console\etc\vsc\web\NetAppNasPlugin.vib.

---

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 NFS Plug-in for VMware VAAI

The screenshot shows the vSphere Client interface for the ICEF1-VC.ice.rtp.netapp.com. The left sidebar contains navigation links: Overview, Storage Details - SAN, Storage Details - NAS, Data Collection, Tools, and Discovery Status. The main content area is titled 'NFS Plug-in for VMware VAAI' and includes the following sections:

- MBR Tools:** This ESX console-based tool tests and aligns guest file systems on a VMDK for VMFS and NFS datastores. It includes two download buttons: 'Download (For ESX 4.x)' and 'Download (For ESXi 4.x and ESXi 5.x)'.
- Guest OS Tools:** Guest OS timeout scripts are used to set the SCSI I/O timeout values for supported guest operating systems. It includes links for Linux OS, Windows OS, and Solaris OS.
- NFS Plug-in for VMware VAAI:** The NFS Plug-in for VMware VAAI is a software library that integrates with VMware's Virtual Disk Libraries. It includes a note about checking the Release Notes for more information and an 'Install on Host' button.

The bottom of the interface shows a 'Recent Tasks' table with columns: Name, Target, Status, Details, Initiated by, vCenter Server, and Requested Start Time.

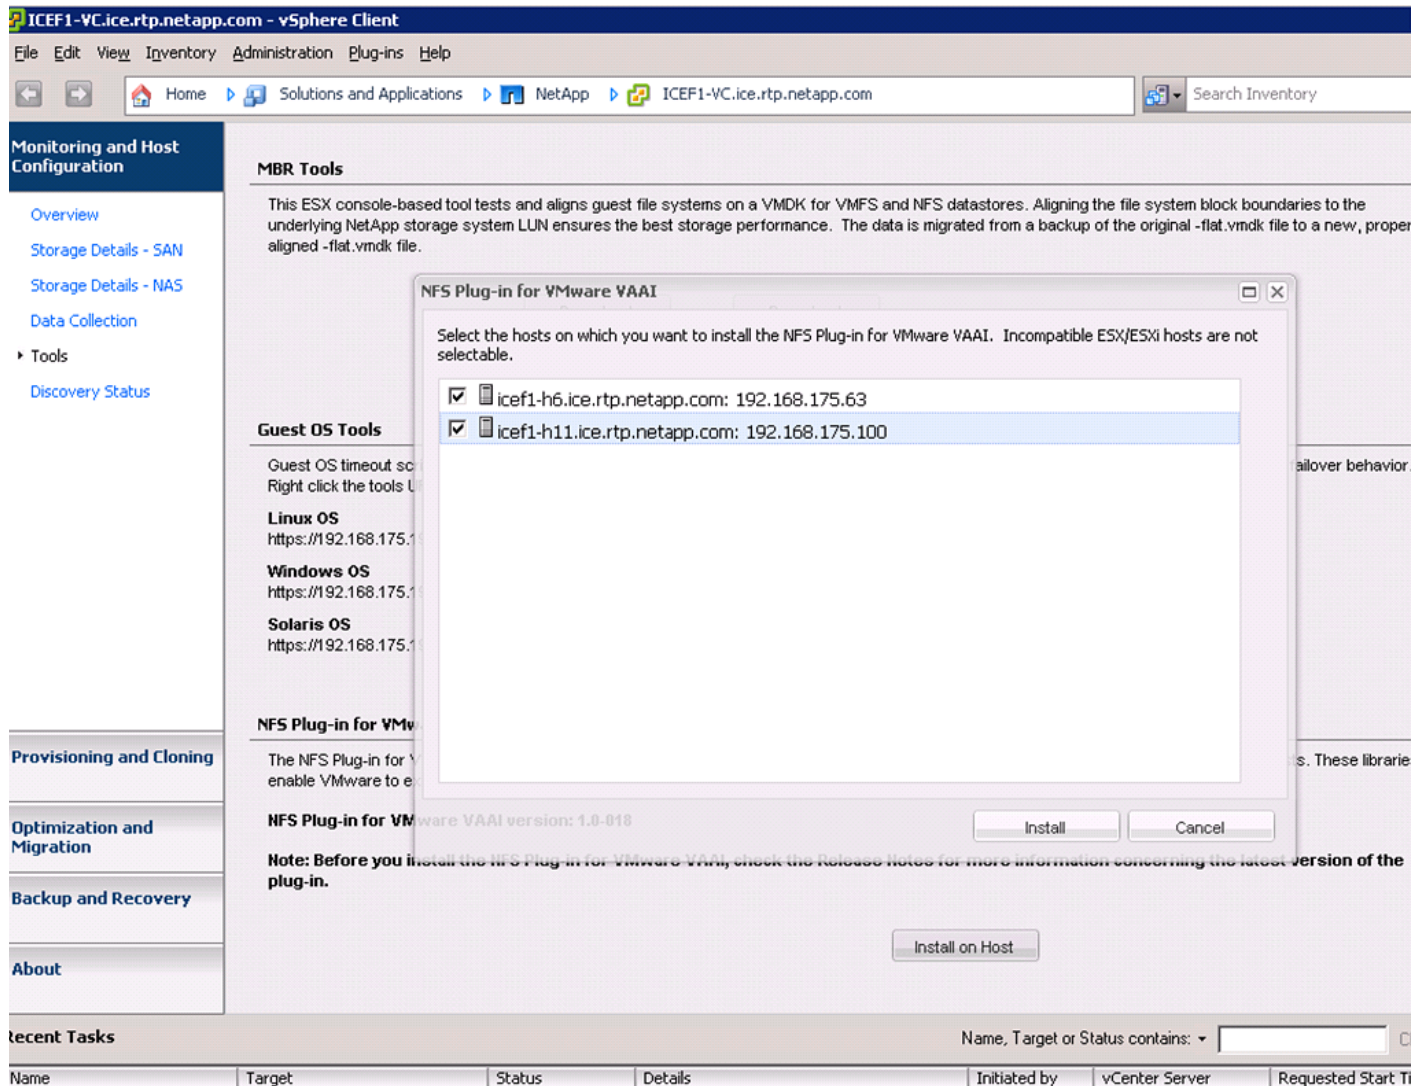
- Choose all ESXi hosts and click **Install**, and then click **Yes**.



#### Note

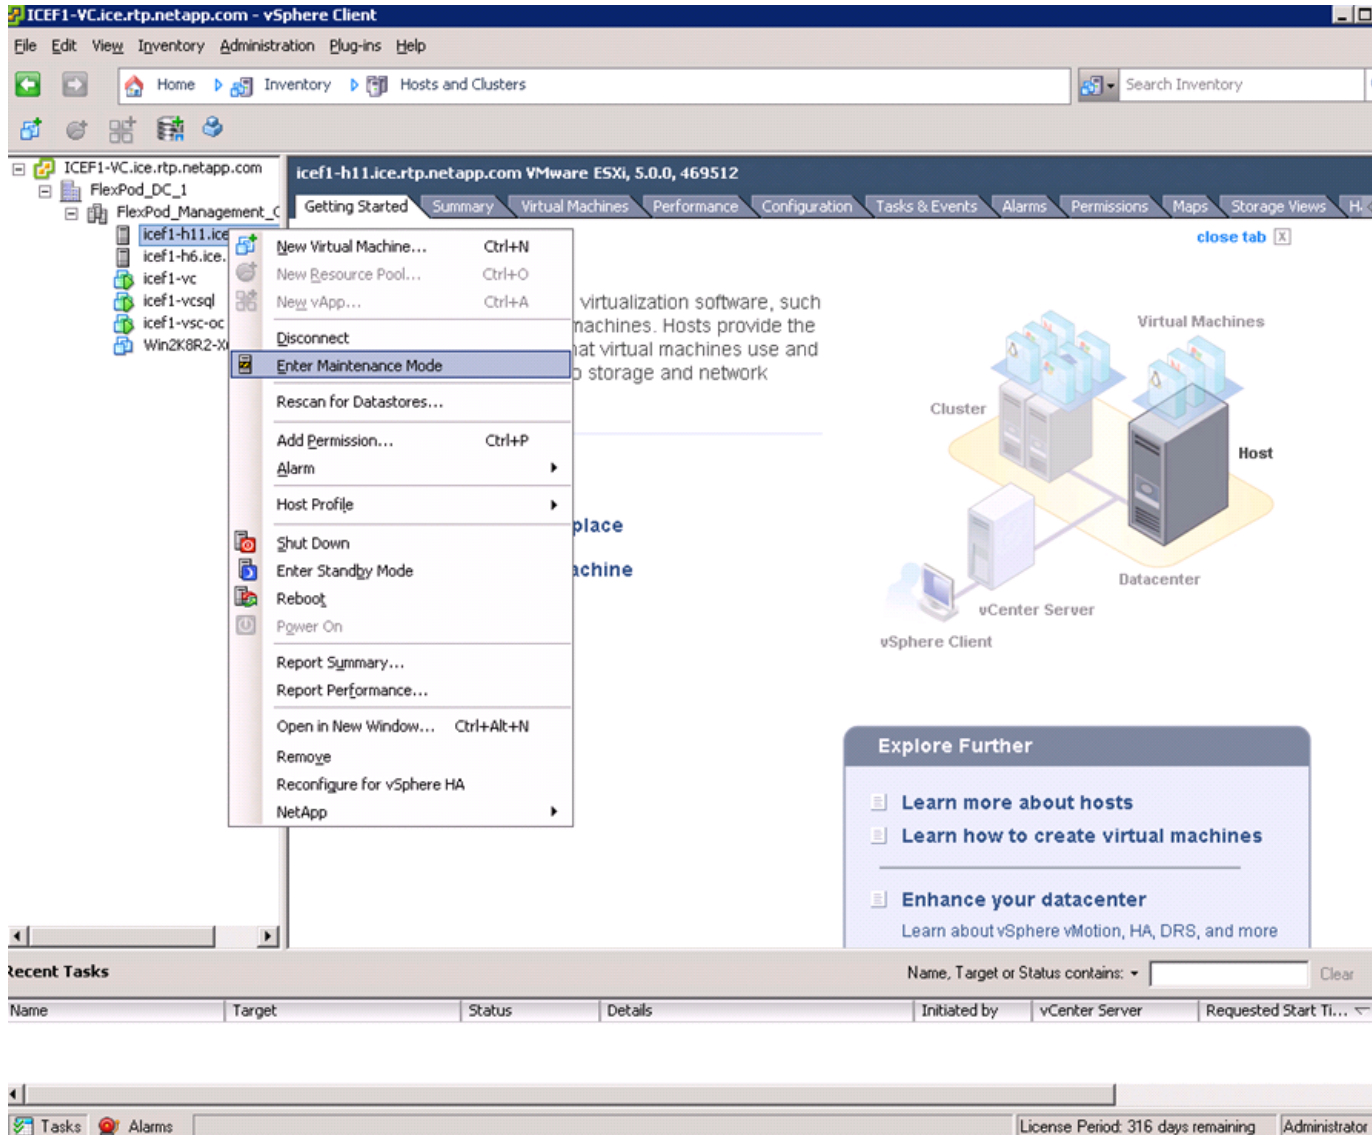
The Monitoring and Host Configuration capability automatically installs the plug-in on the hosts selected.

**Figure 111 NFS Plug-in Host Selection**



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

Figure 112 ESXi - Entering 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.

# Appendix

## Build Windows Active Directory Server VM(s)

### ESXi Host VM-Host-Infra-01

To build an Active Directory 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. 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.



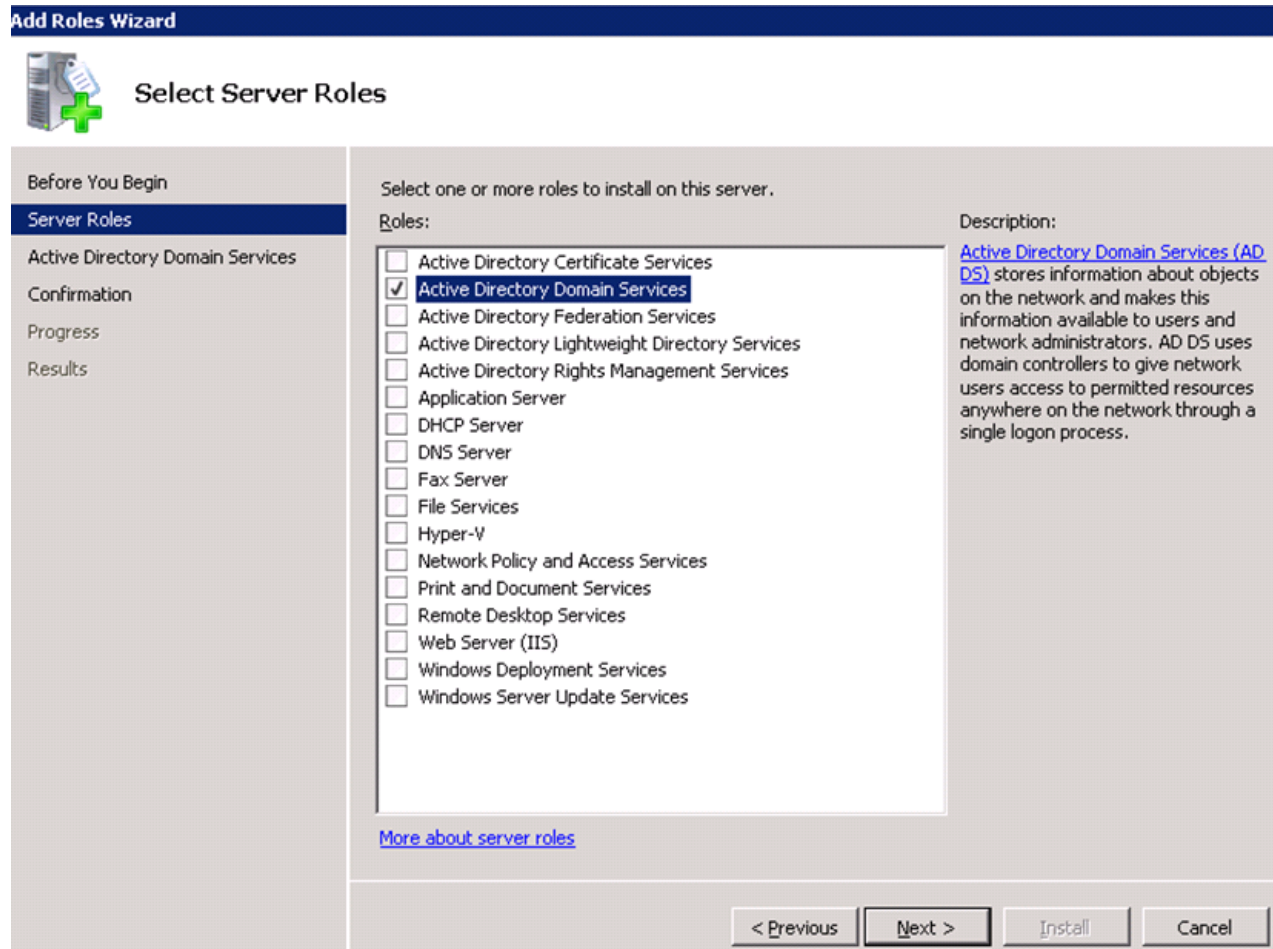

---

**Note** This process requires several reboots.

---


47. Open Server Manager.
48. On the left, click Roles, then choose Add Roles on the right.
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 113**      **Installing Active Directory Domain Services**



52. Click **Next**.

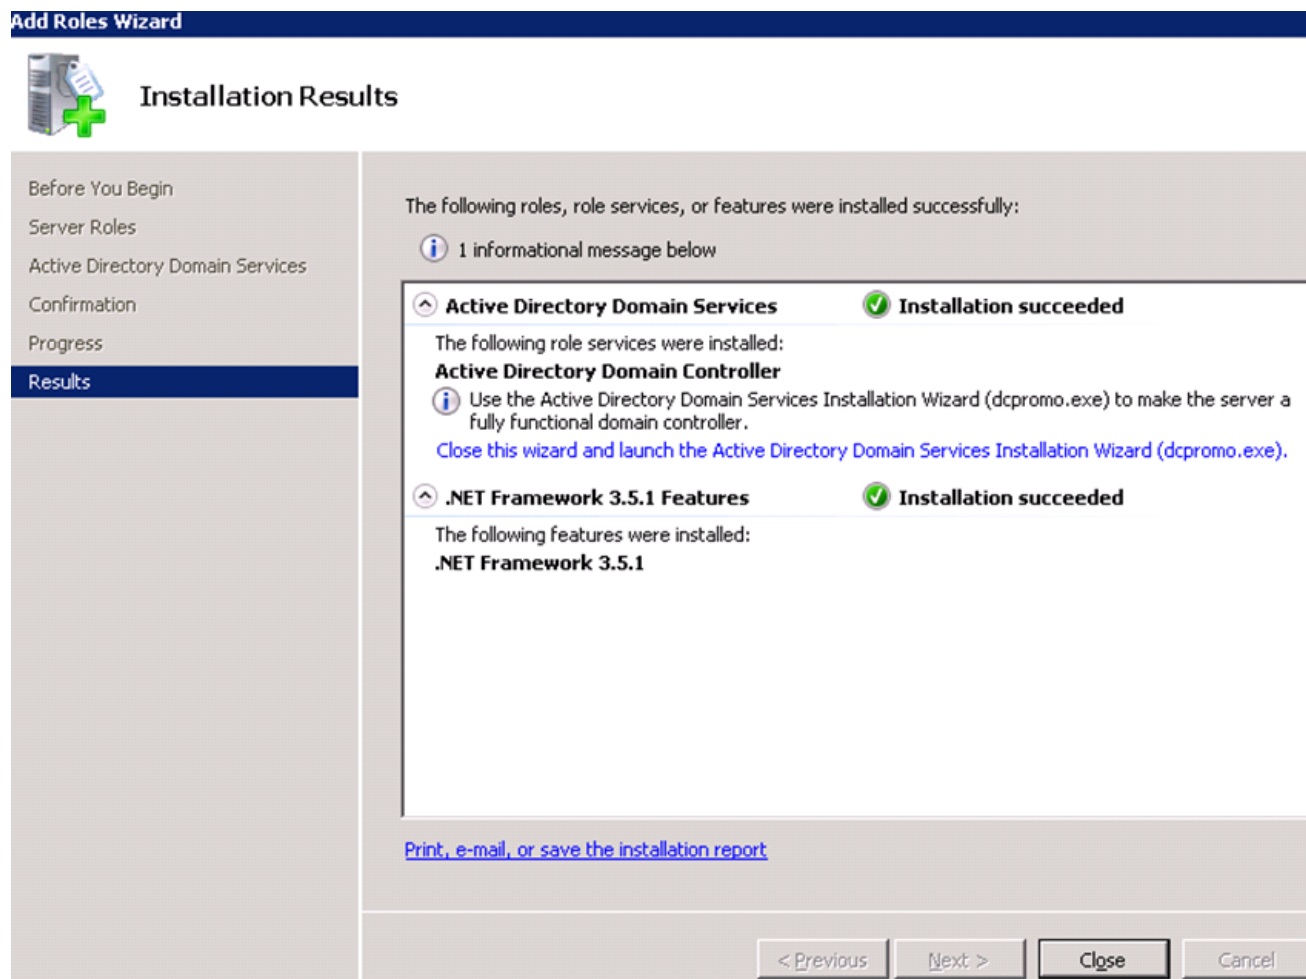
53. Click **Next**.

**Figure 114**      **Active Directory - Confirm Installation Options**

54. Click **Install**.

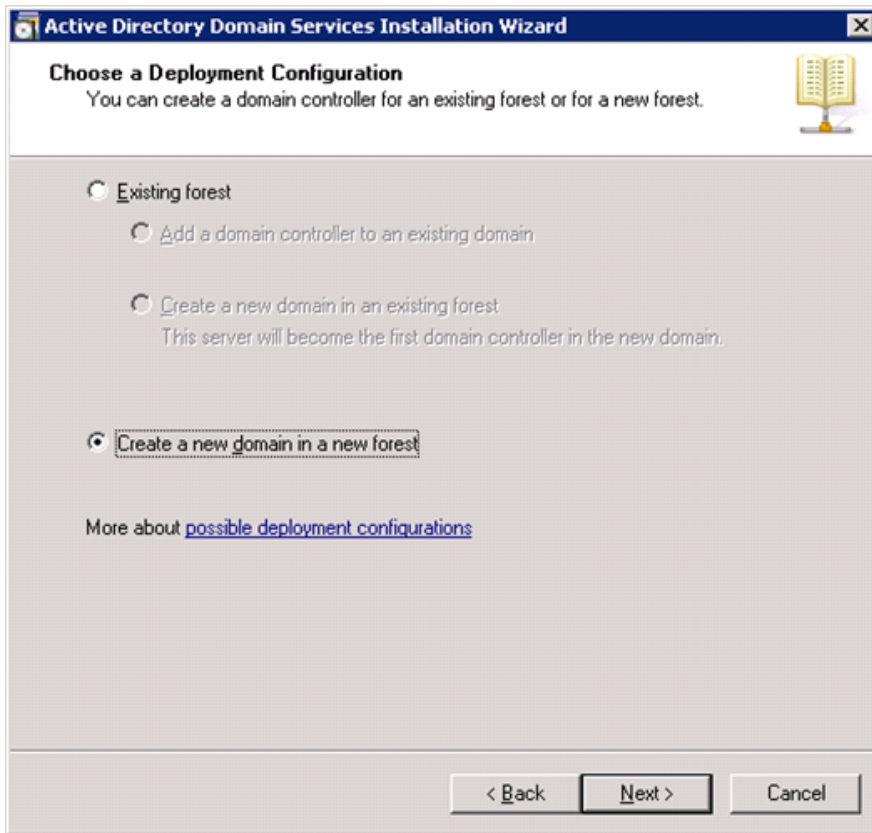


**Figure 115**      **Active Directory - Successful Installation**



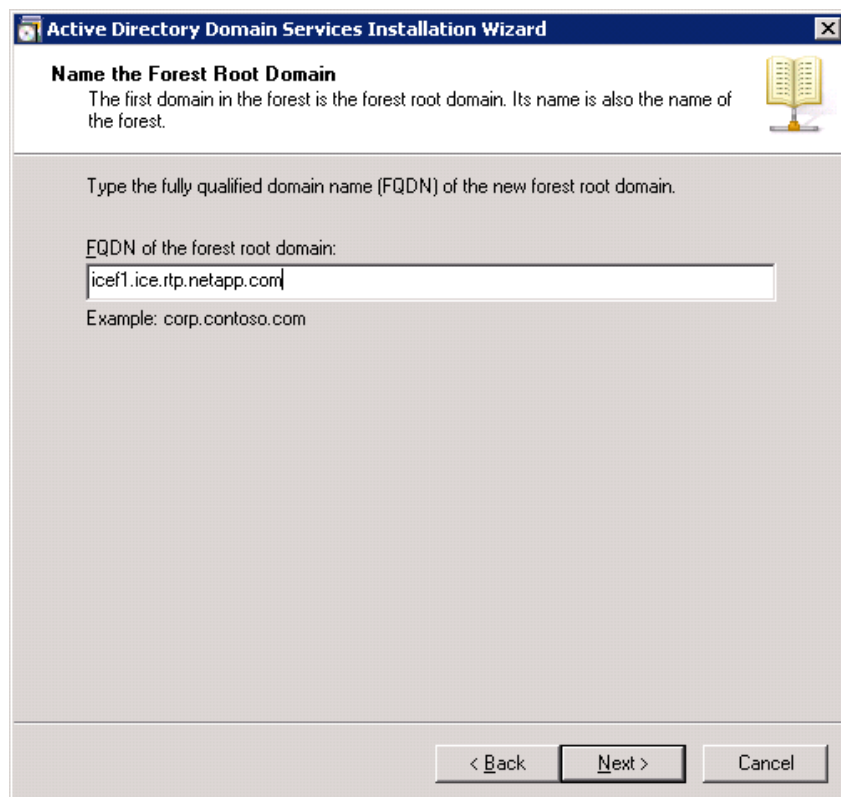
55. In the middle of the window, click **Close** this wizard and launch the Active Directory Domain Services Installation Wizard (dcpromo.exe).
56. In the Active Directory Domain Services Installation Wizard, click **Next**.
57. Click **Next**.
58. Choose Create a new domain in a new forest and click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

**Figure 116**      **Active Directory - Creating a New Domain**



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

**Figure 117**      **Active Directory - Naming Root Domain**



The screenshot shows the 'Active Directory Domain Services Installation Wizard' window. The title bar reads 'Active Directory Domain Services Installation Wizard'. The main heading is 'Name the Forest Root Domain'. Below the heading, a text box explains: 'The first domain in the forest is the forest root domain. Its name is also the name of the forest.' To the right of this text is an icon of an open book. Below the explanation, a prompt says 'Type the fully qualified domain name (FQDN) of the new forest root domain.' followed by 'FQDN of the forest root domain:'. A text input field contains the value 'icef1.ice.rtp.netapp.com'. Below the input field, an example is provided: 'Example: corp.contoso.com'. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

**Active Directory Domain Services Installation Wizard**

**Name the Forest Root Domain**

The first domain in the forest is the forest root domain. Its name is also the name of the forest.

Type the fully qualified domain name (FQDN) of the new forest root domain.

FQDN of the forest root domain:

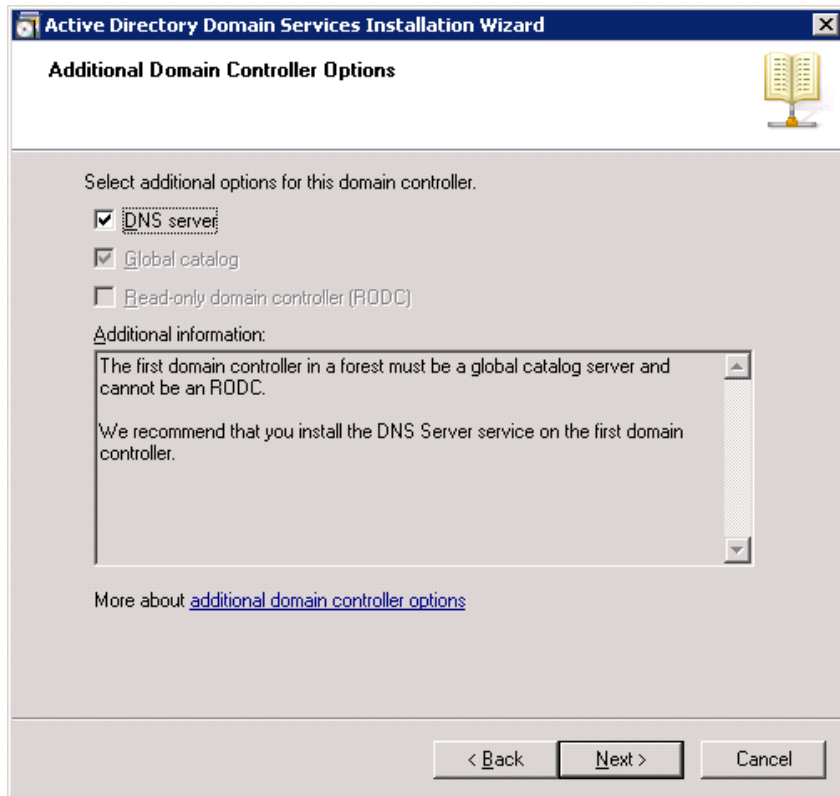
icef1.ice.rtp.netapp.com

Example: corp.contoso.com

< Back   Next >   Cancel

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

**Figure 118 Active Directory - Domain Controller Options**



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](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 sub organization 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. Refer to appendix 13.3, “VM-FEX Virtual Interfaces.”

**Figure 119 UCS - Dynamic vNIC Connection Policy**

The screenshot shows the 'Create Dynamic vNIC Connection Policy' wizard. The fields are as follows:

- Name:** FEX
- Description:** vNIC Connection Policy for FEX
- Number of Dynamic vNICs:** 64
- Adapter Policy:** VMWare
- Protection:** Protected (selected), Protected Pref A, Protected Pref B

Buttons for 'OK' and 'Cancel' are at the bottom right.

## Create a Server BIOS Policy

To define the BIOS policy for a service profile that supports VM-FEX on ESXi, follow these 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 sub organization 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 120 UCS - Create BIOS Policy

Create BIOS Policy

# Unified Computing System Manager

Create BIOS Policy

1. **Main**
2. Processor
3. Intel Directed IO
4. RAS Memory
5. Serial Port
6. USB
7. PCI Configuration
8. Boot Options
9. Server Management

## Main

Name:

Reboot on BIOS Settings Change: ☐

Quiet Boot: ☐ disabled ☐ enabled ☒ Platform Default

Post Error Pause: ☐ disabled ☐ enabled ☒ Platform Default

Resume Ac On Power Loss: ☐ stay-off ☐ last-state ☐ reset ☒ Platform Default

Front Panel Lockout: ☐ disabled ☐ enabled ☒ Platform Default

ACPI10 Support: ☐ disabled ☐ enabled ☒ Platform Default

< Prev Next > Finish Cancel

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



**Figure 121 UCS - Enabling VT and Direct Cache Access Policies**

**Create BIOS Policy**

# Unified Computing System Manager

Create BIOS Policy

1. ☒ Main
2. ☒ **Processor**
3. ☐ Intel Directed IO
4. ☐ RAS Memory
5. ☐ Serial Port
6. ☐ USB
7. ☐ PCI Configuration
8. ☐ Boot Options
9. ☐ Server Management

## Processor

Turbo Boost: ☐ disabled ☐ enabled ☒ Platform Default

Enhanced Intel Speedstep: ☐ disabled ☐ enabled ☒ Platform Default

Hyper Threading: ☐ disabled ☐ enabled ☒ Platform Default

Core Multi Processing: Platform Default

Execute Disabled Bit: ☐ disabled ☐ enabled ☒ Platform Default

Virtualization Technology (VT): ☐ disabled ☒ enabled ☐ Platform Default

Direct Cache Access: ☐ disabled ☒ enabled ☐ Platform Default

Processor C State: ☐ disabled ☐ enabled ☒ Platform Default

Processor C1E: ☐ disabled ☐ enabled ☒ Platform Default

Processor C3 Report: ☐ disabled ☐ acpi-c2 ☐ acpi-c3 ☒ Platform Default

Processor C6 Report: ☐ disabled ☐ enabled ☒ Platform Default

Processor C7 Report: ☐ disabled ☐ enabled ☒ Platform Default

CPU Performance: ☐ enterprise ☐ high-throughput ☐ hpc ☒ Platform Default

Max Variable MTRR Setting: ☐ auto-max ☐ 8 ☒ Platform Default

< Prev Next > Finish Cancel

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



Figure 122 UCS - Enable Intel Directed IO




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 completing the following 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.

**Figure 123 UCS - Cloning a Service Profile Template**



4. Click **OK** when notified that the service profile clone was successfully created. The Service Template navigation window appears.
5. 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.
6. Choose Use a Dynamic vNIC Connection Policy from the drop-down menu and the previously created Dynamic vNIC policy. Click **OK**.

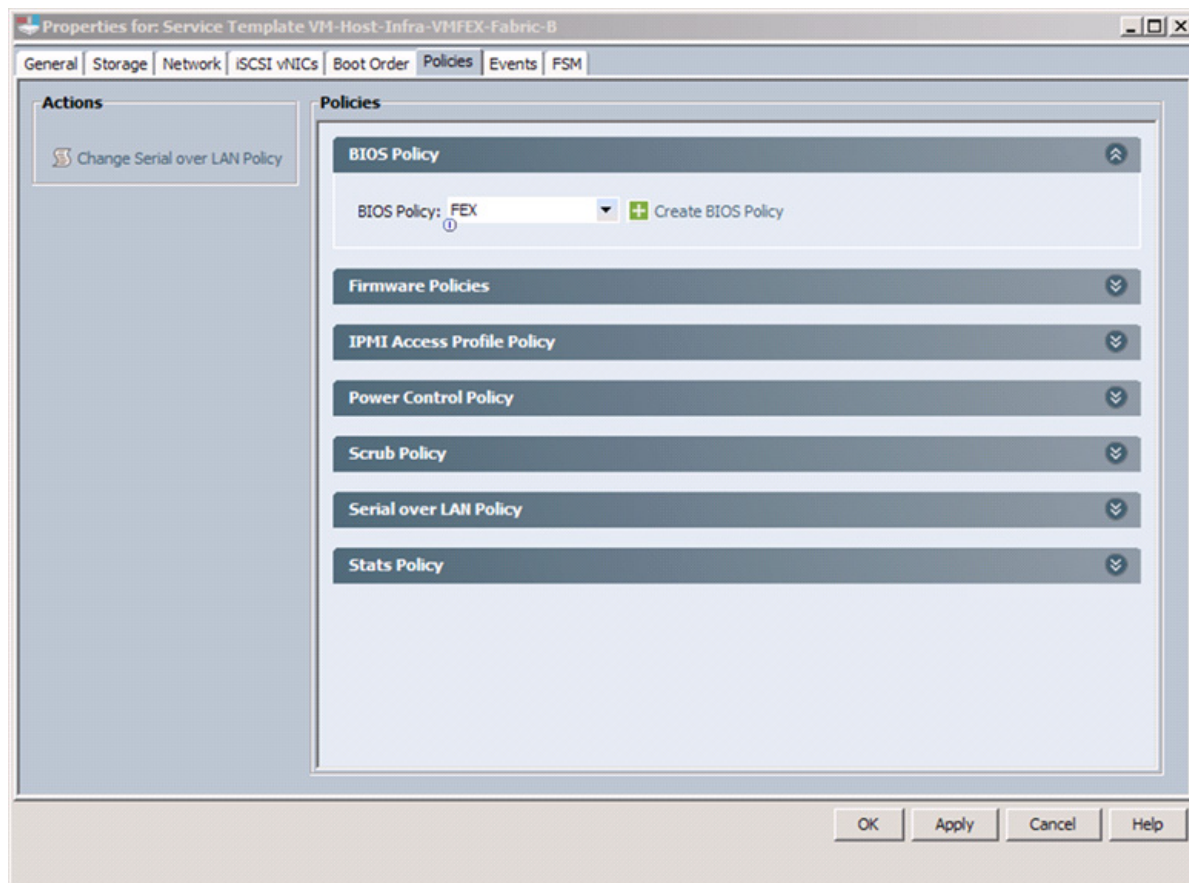
**Figure 124 UCS - Create Dynamic vNIC Connection Policy**



7. Click **OK** when notified that the vNIC connection policy was successfully modified.
8. From the Service Template properties window, click the **Policies** tab.

9. Expand the BIOS Policies in the Policies section of the working pane.
10. Choose the previously defined FEX BIOS policy and click **OK**.

**Figure 125** UCS - Choosing a BIOS Policy in the Service Profile Template

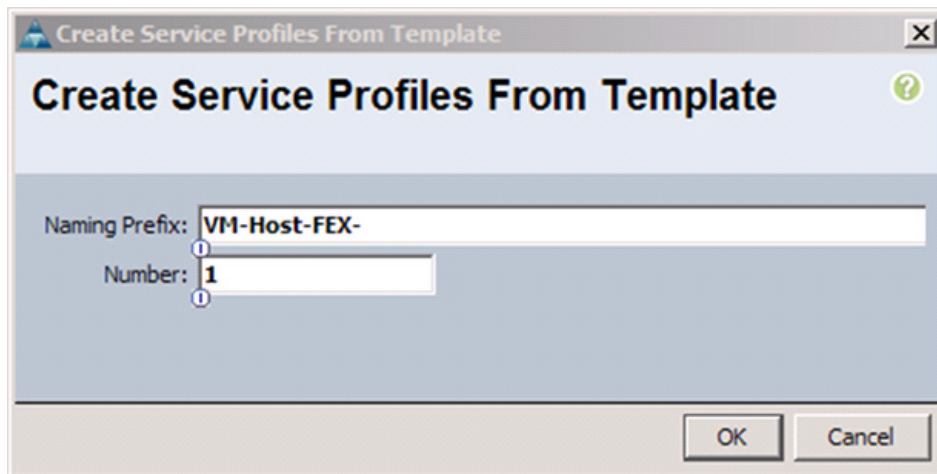


## Create VM-FEX Service Profile

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 number of service profiles to create.
6. Click **OK** to create the service profile.

Figure 126 UCS - Create Service Profile



7. Click **OK** in the confirmation message.
8. Verify that the service profile VM-Host-FEX-1 has been created. The service profile is automatically associated with the servers in their assigned server pools.

## Install and Set Up VMware ESXi

Refer to section 11.1 to install and completely set up VMware ESXi version 5.1 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(1d)
  - 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:

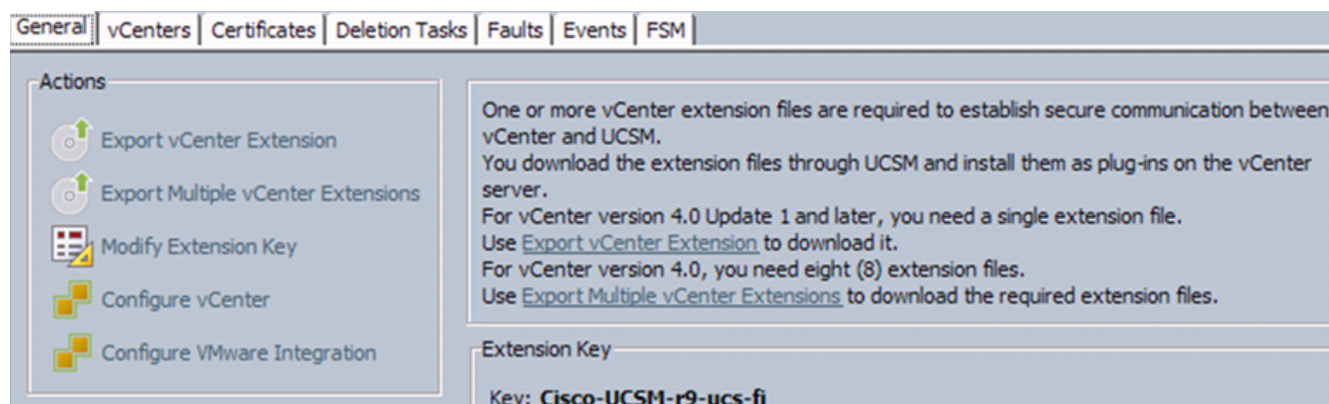
```
esxcli -s <host_ip> -u root -p <host_password> software vib install -v
/vmfs/volumes/infra_datastore_1/VM-FEX/cisco-vem-v151-5.1-1.1.1.1.vib
```

## 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.

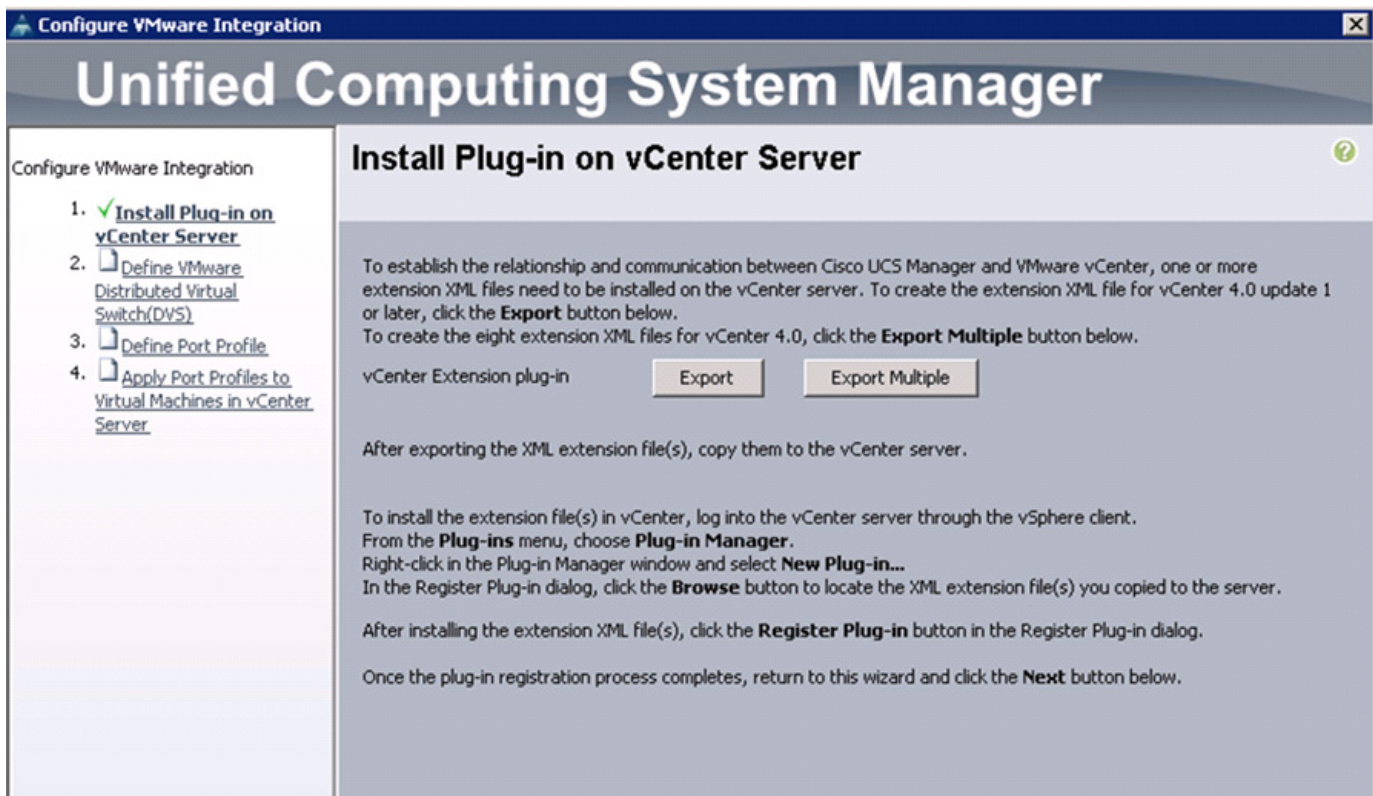
**Figure 127** UCS - VM Tab



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 128 UCS - Configuring VMware Integration



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 129 UCS - Define VMware Distributed Virtual Switch**

**Unified Computing System Manager**

Configure VMware Integration

1. ☒ Install Plug-in on vCenter Server
2. ☒ **Define VMware Distributed Virtual Switch(DVS)**
3. ☐ Define Port Profile
4. ☐ Apply Port Profiles to Virtual Machines in vCenter Server

### Define VMware Distributed Virtual Switch(DVS)

**vCenter Server**

vCenter Server Name: <<var vcenter Server Name>>  
 Description:   
 vCenter Server Hostname or IP Address: <<var\_vcenter\_server\_ip>>

**Datacenter**

vCenter Datacenter Name: FlexPod\_DC\_1  
 Description:

**DVS Folder**

Folder Name: DVS-FEX  
 Description:

**DVS**

DVS Name: DVS-FEX  
 Description:   
 DVS: ☐ Disable ☒ Enable

< Prev   Next >   Finish   Cancel

8. Create the FEX-MGMT port profile, choose the MGMT-VLAN, and indicate it is the native VLAN.

Figure 130 UCS - Define Port Profile

**Configure VMware Integration**

**Unified Computing System Manager**

**Define Port Profile**

**Configure VMware Integration**

1. ☒ [Install Plug-in on vCenter Server](#)
2. ☒ [Define VMware Distributed Virtual Switch\(DVS\)](#)
3. ☒ [Define Port Profile](#)
4. ☐ [Apply Port Profiles to Virtual Machines in vCenter Server](#)

**Port Profile**

Name:

QoS Policy:

Network Control Policy:

Max Ports:

Pin Group:

**VLANs**

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	FooBar1_public	<input type="radio"/>
<input checked="" type="checkbox"/>	MGMT-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	Packet-Control-VLAN	<input type="radio"/>
<input type="checkbox"/>	Service-HA	<input type="radio"/>
<input type="checkbox"/>	ServiceNodeServices	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

**Profile Client**

Name:

Description:

Datacenter:

Folder:

Distributed Virtual Switch:

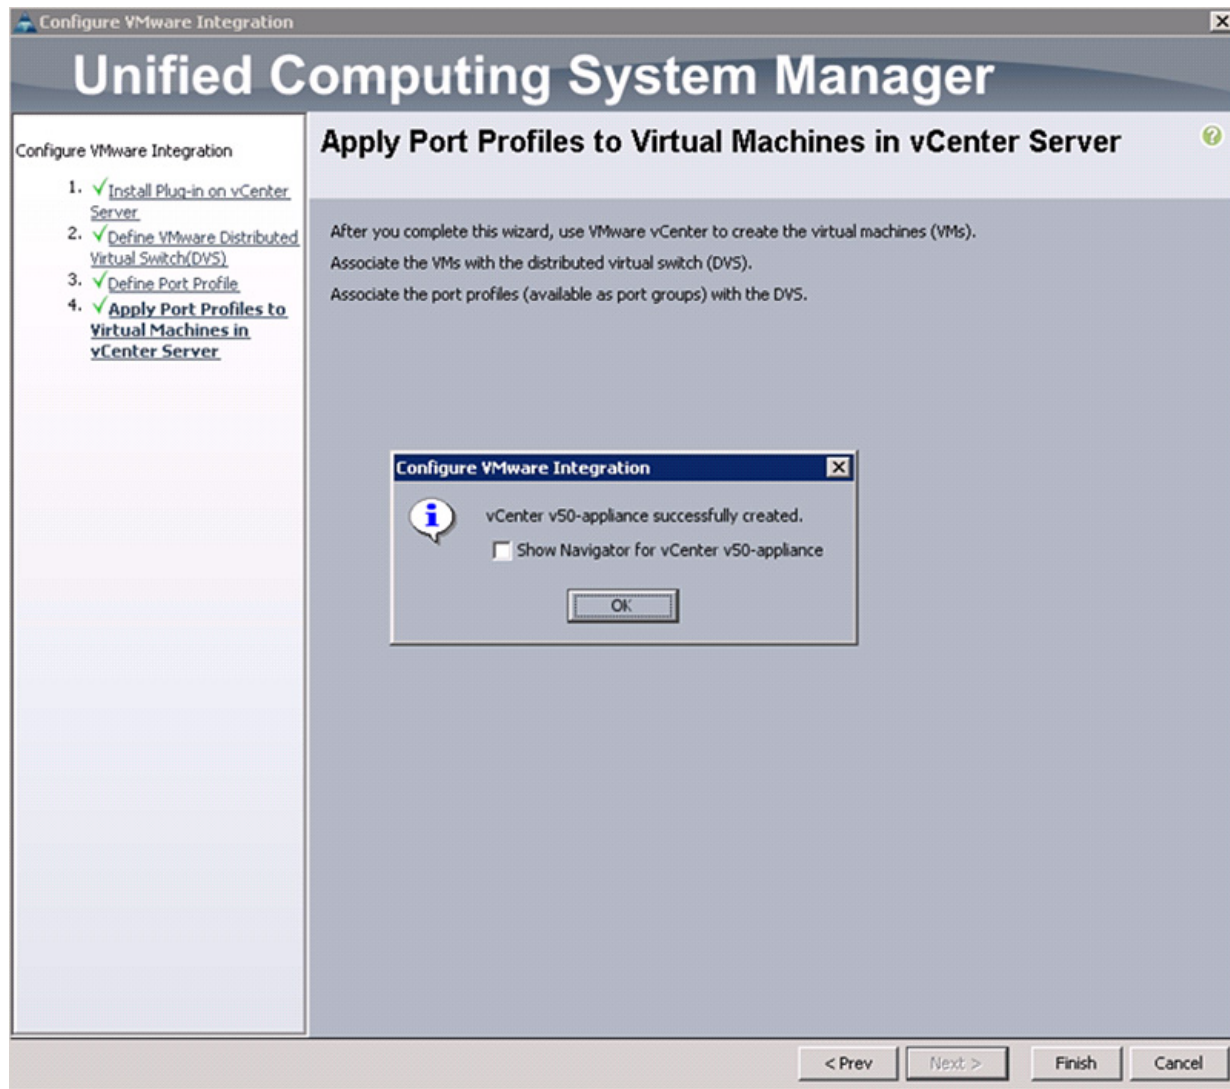
< Prev    Next >    Finish    Cancel

9. Click **Next**.

10. When finishing the wizard, the Cisco UCS Manager connects to vCenter and adds the plug-in.



Figure 131 UCS - Apply Port profile to VMs

**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 132** *VMware - UCS Plug-in Installation Verification*



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

**Figure 133** *VMware - Addition of FEX*



## 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 134 UCS - Create Port Profile

**Create Port Profile**

Name:

Description:

QoS Policy:

Network Control Policy:

Max Ports:

Host Network IO Performance: ☐ None ☒ High Performance

Pin Group:

**VLANs**

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	FooBar1_public	<input type="radio"/>
<input type="checkbox"/>	MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	Packet-Control-VLAN	<input type="radio"/>
<input type="checkbox"/>	Service-HA	<input type="radio"/>
<input type="checkbox"/>	ServiceNodeServices	<input type="radio"/>
<input checked="" type="checkbox"/>	VM-Traffic-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

OK Cancel

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 135 UCS - Create Profile Client**

**Create Profile Client**

Name: FEX-Traffic-VLAN

Description:

Datacenter: r9-dc-1

Folder: DVS-FEX

Distributed Virtual Switch: DVS-FEX

OK Cancel

The client profile created will appear in your distributed virtual switch DVS-FEX in vCenter as a port group.

Repeat these steps as necessary for the workloads in the environment.

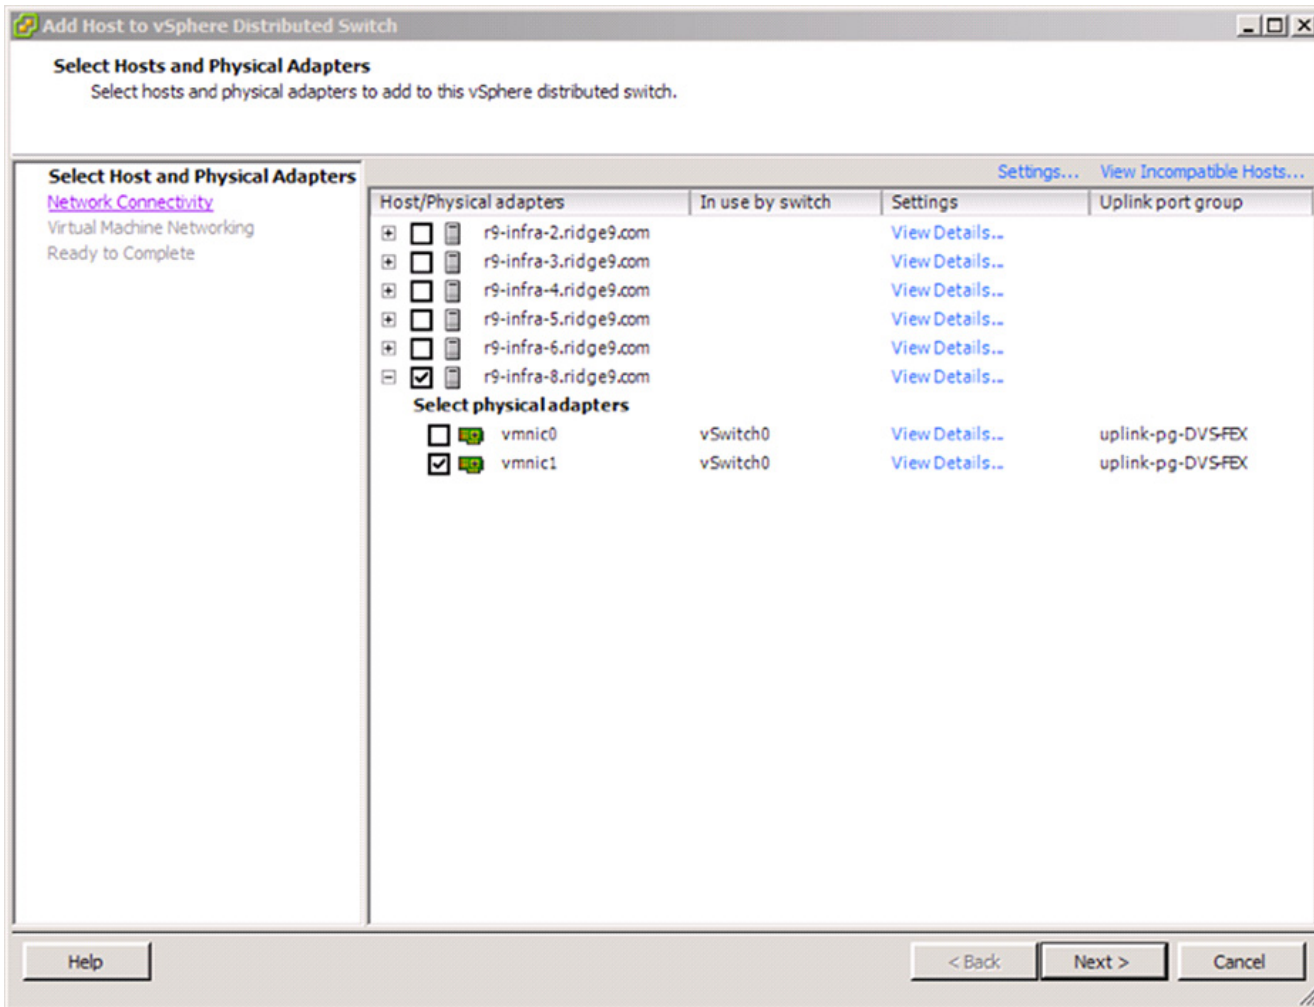
## Migrate Networking Components for ESXi Hosts to Cisco DVS-FEX

### 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** VMware – Select vmnic1 for uplink-pg-DVS-FEX



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

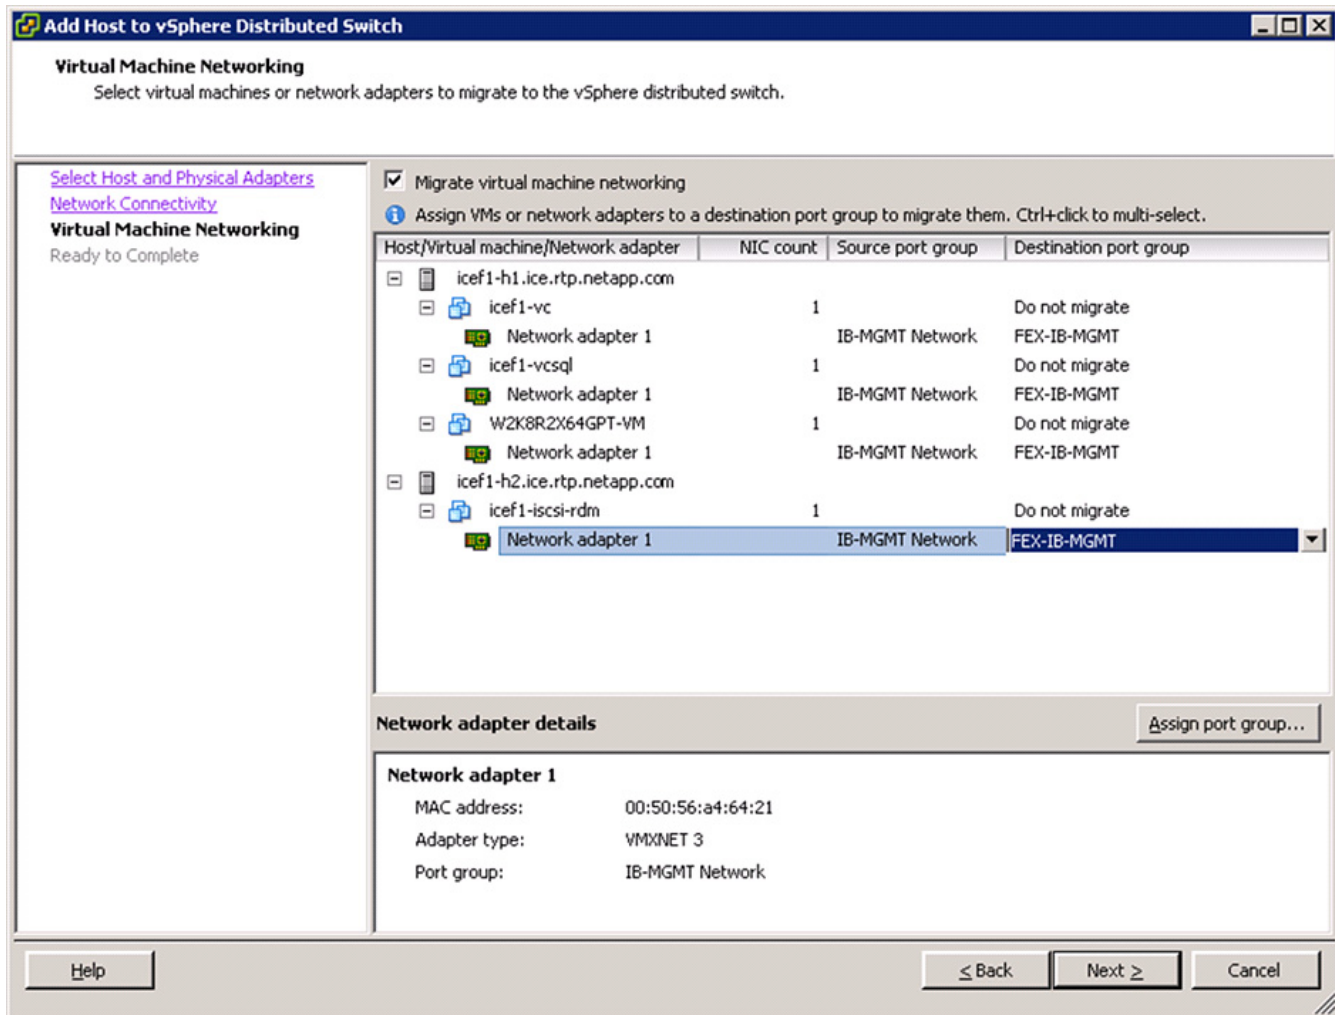
**Figure 137 VMware – FEX Port Group Selection**



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



Figure 138 VMware – Migrate the 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 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 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 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 29](#) and [Table 30](#) reference these limits.

**Table 29** *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 30** *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
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](http://www.cisco.com/en/US/docs/unified_computing/ucs/c-series_integration/ucsm2.1/b_UCSM2-1_C-Integration.pdf)

## Expand Two-Node Cluster to Four-Node Cluster FlexPod

1. Log in to the cluster interface and disable cluster HA by typing:
 

```
cluster ha modify -configured false
```
2. Build the two new storage cluster nodes using the procedure described in “[Storage Configuration](#)” section on page 27.
  - a. In “[Clustered Data ONTAP 8.1.2](#)” section on page 32, make sure Data ONTAP 8.1.2 is installed, initialize the disks, and assign disks for the two new controllers.
  - b. In “[Cluster Create in Clustered Data ONTAP](#)” section on page 35, use the Node 2 instructions to join Nodes 3 and 4 to the cluster.
  - c. For logging into the cluster, see “[Log in to the Cluster](#)” section on page 40.
  - d. In “[Zero All Spare Disks](#)” section on page 40, zero all spare disks on Nodes 3 and 4.
  - e. In “[Set Auto-Revert on Cluster Management](#)” section on page 40, no action is necessary.
  - f. In “[Failover Groups Management in Clustered Data ONTAP](#)” section on page 40, add Node 3 and 4 e0a ports to the mgmt failover group.
  - g. In “[Assign Management Failover Group to Cluster Management LIF](#)” section on page 40, no action is necessary.
  - h. In “[Failover Groups Node Management in Clustered Data ONTAP](#)” section on page 40, create failover groups node-mgmt03 and node-mgmt04.
  - i. In “[Assign Node Management Failover Groups to Node Management LIFs](#)” section on page 41, complete the assignments for Nodes 3 and 4.
  - j. In “[Flash Cache in Clustered Data ONTAP](#)” section on page 41, set up Flash Cache on Nodes 3 and 4.

- k. In [“64-Bit Aggregates in Clustered Data ONTAP” section on page 41](#), create aggr03 on Node3 and aggr04 on Node 4, disable Snapshot copies on these aggregates, and delete any existing Snapshot copies on these aggregates. Rename aggr0 on Node 3.
- l. In [“Service Processor” section on page 42](#) upgrade and configure the service processors on Nodes 3 and 4.
- m. In [“Storage Failover in Clustered Data ONTAP” section on page 43](#), execute steps 1 and 3 for Nodes 3 and 4.
- n. In [“IFGRP LACP in Clustered Data ONTAP” section on page 44](#), create ifgrp i0a on Nodes 3 and 4.
- o. In [“VLAN in Clustered Data ONTAP” section on page 44](#), add VLAN interfaces for the NFS VLAN on Nodes 3 and 4.
- p. In [“Jumbo Frames in Clustered Data ONTAP” section on page 44](#), modify the newly added VLAN interfaces for jumbo frames.
- q. In [“NTP in Clustered Data ONTAP” section on page 44](#), only create the NTP server services for Nodes 3 and 4.
- r. No action is necessary under the following sections:
  - [SNMPv1 in Clustered Data ONTAP, page 45](#)
  - [SNMPv3 in Clustered Data ONTAP, page 45](#)
- s. In [“AutoSupport HTTPS in Clustered Data ONTAP” section on page 46](#), execute the one step listed.
- t. In [“Cisco Discovery Protocol in Clustered Data ONTAP” section on page 46](#), enable CDP on Nodes 3 and 4.
- u. In [“Vserver” section on page 46](#), only execute the last step to add aggr03 and aggr04 to the aggregate list for Infra\_Vserver:
 

```
vserver modify -vserver Infra_Vserver -aggr-list aggr01, aggr02, aggr03, aggr04
```
- v. In [“Create Load Sharing Mirror of Vserver Root Volume in Clustered Data ONTAP” section on page 47](#), create root\_vol\_m03 on aggr03 and root\_vol\_m04 on aggr04. Create the two new SnapMirror relationships. Use the following commands to initialize the two new SnapMirror relationships.
 

```
snapmirror initialize -source-path //Infra_Vserver/root_vol -destination-path //Infra_Vserver/root_vol_m03
snapmirror initialize -source-path //Infra_Vserver/root_vol -destination-path //Infra_Vserver/root_vol_m04
```

 Finally, execute step 4 to set the SnapMirror relationships to an hourly schedule.
- w. In [“FC Service in Clustered Data ONTAP” section on page 47](#), no action is necessary.
- x. In [“HTTPS Access in Clustered Data ONTAP” section on page 48](#), generate certificates for the Node 3 and Node 4 Management Interfaces, and delete the preconfigured certificates for these interfaces. Using the security ssl modify command, assign these newly created certificates to the Node Management interfaces.
- y. No action is necessary under the following sections:
  - [NFSv3 in Clustered Data ONTAP, page 49](#)
  - [FlexVol in Clustered Data ONTAP, page 49](#)
  - [LUN in Clustered Data ONTAP, page 49](#)
  - [Deduplication in Clustered Data ONTAP, page 50](#)

- z. In “Failover Groups NAS in Clustered Data ONTAP” section on page 50, add Node 3 and 4 NFS ports to the NFS failover group.
- aa. In “NFS LIF in Clustered Data ONTAP” section on page 50, create LIF nfs\_lif03 on Node 3 and nfs\_lif04 on Node 4.
- ab. In “FCP LIF in Clustered Data ONTAP” section on page 50, create fcp\_lif03a and fcp\_lif03b on Node 3 and fcp\_lif-4a and fcp\_lif04b on Node 4.
- ac. No action is necessary for “Add Infrastructure Vserver Administrator” section on page 51.
- 3. Using the procedures described in “Network Configuration” section on page 113 provision the Ethernet Ports, Port Channels, and VPCs for the ports connected from Nodes 3 and 4 to the switches. Then, add device aliases for the new FCP LIFs, add the FCoE VLAN to the storage port channels on each switch, and configure the new vFC interfaces and add them to the VASN database on each switch.
- 4. You can now add datastores on the new nodes or migrate volumes and NAS LIFs to the two nodes in your cluster.

## Migrate from 7-Mode FlexPod to Clustered Data ONTAP FlexPod

This procedure describes one method of migrating the FlexPod VMware Management Cluster (two ESXi hosts) from existing 7-Mode storage in a FlexPod unit to added clustered Data ONTAP storage. For FlexPod workload migration, engage NetApp Professional Services to properly migrate application data LUNs to clustered Data ONTAP. This procedure assumes setting up two new ESXi hosts on the clustered Data ONTAP storage and migrating all management VMs to these two new servers instead of migrating the host boot LUNs to clustered Data ONTAP. To migrate the boot LUNs to clustered Data ONTAP, it is necessary to engage NetApp Professional Services.

1. Cable the two new clustered Data ONTAP nodes by referring to section 6, “Physical Infrastructure.”
2. Build the storage cluster according to section 7, “Storage Configuration.” Assume that two new servers will be added. Assign NFS IPs to these two new servers and use them to create FlexPod export policy rules.
3. On the 7-Mode storage systems, add the two new servers’ NFS IPs to the exports for infra\_datastore\_1.
4. In the Cisco UCS Manager, create clustered Data ONTAP boot policies, service profile templates, and two Service Profiles. Refer to section 8, “Server Configuration.”
5. In the Cisco Nexus 7000s, add the cluster node ports, vPCs, and vFCs. Add the new device aliases for the cluster FCP LIFs and the two new server HBAs. Add zones for the two new servers, put them in the FlexPod zoneset, and activate it. Refer to section 9, “Storage Networking.”
6. Create igroups in the cluster and map the two new boot LUNs to the igroups using section 10 as a guide.
7. Install and configure ESXi on the two new servers. Refer to section 11, “VMware vSphere 5.1 Setup.” Mount the infra\_datastore\_1 and infra\_swap datastores with different names on the two new servers, that is, infra\_cl\_datastore\_1 and infr\_cl\_swap.
8. Add the two new servers to the FlexPod\_Management cluster in vCenter.
9. Add the two new servers to the Cisco Nexus 1000v, including installing the VEM on each server.
10. Using VSC, add the storage cluster to VSC.
11. Using VSC set up the best practice parameters on the two new servers.

12. Install the NetApp VAAI NFS plug-in on the two new servers, including enabling vStorage on the infrastructure Vserver.
13. In the vSphere Client connected to vSphere, under Home > Inventory > Hosts and Clusters, right-click each of the two new ESXi hosts and using NetApp submenu, mount the 7-Mode infra\_datastore\_1 to the two new servers that are booted from the clustered storage.
14. If the 7-Mode storage will not be retained in the FlexPod unit, do the following:
  - a. Go in to the VSC-OnCommand VM and uninstall OnCommand Core. Using SnapDrive, delete and the OnCommandDB LUN and disk.
  - b. If no other VMs are using RDM mapped disks, using VSC, destroy the RDM\_Map datastore on the 7-Mode storage.
  - c. Shut down and remove the VASA VM.
  - d. Use vMotion to migrate the VC, VCSQL, and VSC-OC VMs to the two new servers in which the 7-Mode datastore is mounted.
  - e. Use Storage vMotion to migrate the VC, VCSQL, and VSC-OC VMs to the clustered Data ONTAP datastore.
  - f. Unmount the 7-Mode datastore from the two new servers.
  - g. Shut down the two old Management ESXi Servers that were booted from 7-Mode storage.
  - h. Remove these servers from vCenter and from the Cisco Nexus 1000v.
  - i. Halt and remove the 7-Mode storage controllers from the FlexPod unit.
  - j. Remove zones and any network port data for the 7-Mode storage controllers in the Cisco Nexus switches.
  - k. In VSC, remove the 7-Mode storage controllers from the configuration.
  - l. In VSC Backup and Recovery, remove the 7-Mode storage controllers and all associated backup jobs.
15. If the 7-Mode storage will be retained in the FlexPod unit, do the following:
  - a. Use vMotion to migrate the VC, VCSQL, VASA, and VSC-OC VMs to the two new servers where the 7-Mode datastore is mounted.
  - b. Use Storage vMotion to migrate the VC and VCSQL to the clustered Data ONTAP datastore.
  - c. Shut down the two old Management ESXi Servers that were booted from 7-Mode storage.
  - d. Remove these servers from vCenter and from the Cisco Nexus 1000v.
  - e. Remove the boot LUNs for these servers from the 7-Mode storage controllers.
  - f. Remove zones for the 7-Mode storage controllers in the Cisco Nexus switches.
  - g. The new servers now are booted from the clustered Data ONTAP storage but have the NFS datastores mounted from both types of storage.
  - h. Build a new VM on the clustered Data ONTAP datastore for OnCommand, and install SnapDrive on it. Refer to section 12.2, "OnCommand Unified Manager 5.1."
16. Using VSC, thin provision a new 100GB RDM\_Map\_CL datastore on aggr02 on the clustered Data ONTAP storage on the FlexPod\_Management cluster in vCenter.
17. Add the storage cluster to the VSC Backup and Recovery module, and optionally create a recurring backup job for the datastore now containing the management virtual machines. Refer to the section, "VSC 4.1 Backup and Recovery."

18. Log in to the VSC-OnCommand VM or the newly built clustered Data ONTAP OnCommand VM. Set the SnapDrive default storage system Transport Protocol settings to the login for the Infra\_Vserver credentials. Use SnapDrive to create the OnCommandDB LUN on the cluster. Install and configure OnCommand Core.
19. Contact NetApp Professional services to migrate your workload to the cluster.

## Cisco Nexus 7000 Example Configurations

### Cisco Nexus 7000 A

```
L01-7004-1# sh run vdc-all
!Running config for default vdc: L01-7004-1

!Command: show running-config
!Time: Mon May 6 15:46:46 2013

version 6.1(2)
license grace-period
license fcoe module 4

switchname L01-7004-1
system admin-vdc
system qos
  service-policy type network-qos default-nq-7e-policy
install feature-set fcoe
vdc L01-7004-1 id 1
  cpu-share 5
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc flexpod id 2
  limit-resource module-type f2
  cpu-share 5
  allocate interface Ethernet4/1-4,Ethernet4/17-20,Ethernet4/25-28,Ethernet4/41-44
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 8 maximum 8
  limit-resource u6route-mem minimum 4 maximum 4
  limit-resource m4route-mem minimum 8 maximum 8
  limit-resource m6route-mem minimum 5 maximum 5
  limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc storage-A id 3 type storage
  limit-resource module-type f1 f2
  allow feature-set fcoe
  cpu-share 5
  allocate interface Ethernet4/29-32,Ethernet4/37-40
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
```

```

limit-resource monitor-session minimum 0 maximum 2
limit-resource monitor-session-erspan-dst minimum 0 maximum 23
limit-resource vrf minimum 2 maximum 4096
limit-resource port-channel minimum 0 maximum 768
limit-resource u4route-mem minimum 8 maximum 8
limit-resource u6route-mem minimum 4 maximum 4
limit-resource m4route-mem minimum 8 maximum 8
limit-resource m6route-mem minimum 5 maximum 5
limit-resource monitor-session-inband-src minimum 0 maximum 1

username admin password 5 $1$GcNFEZqS$zwohuUxQ2ACIgcI.MAIXt0 role network-admin
ip domain-lookup
system default switchport
copp profile strict
snmp-server user admin network-admin auth md5 0x158f668b2f06cb1adfe2257633eb1782 priv
0x158f668b2f06cb1adfe2257633eb1782 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 172.26.164.254

vrf context management
ip route 0.0.0.0/0 172.26.164.1
vlan 1

vdc storage-A id 3
allocate fcoe-vlan-range 101

interface mgmt0
ip address 172.26.164.81/24
line console
line vty
boot kickstart bootflash:/n7000-s2-kickstart.6.1.2.bin sup-1
boot system bootflash:/n7000-s2-dk9.6.1.2.bin sup-1

!Running config for vdc: flexpod

switchto vdc flexpod

!Command: show running-config
!Time: Mon May 6 15:46:47 2013

version 6.1(2)
switchname flexpod

cfs eth distribute
feature udld
feature interface-vlan
feature lacp
feature vpc

username admin password 5 $1$apHdZJKq$J4vHctcowxiRCKsZH0Axo/ role vdc-admin
ip domain-lookup
system default switchport
snmp-server user admin vdc-admin auth md5 0x3ecb54326ab8d6fe97553a3127a5bed4 priv
0x3ecb54326ab8d6fe97553a3127a5bed4 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

```

```

vrf context management
  ip route 0.0.0.0/0 172.26.164.1
vlan 1-2,3170,3173-3176
vlan 2
  name Native-VLAN
vlan 3170
  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 edge bpdufilter default
spanning-tree port type network default
vpc domain 1
  role priority 10
  peer-keepalive destination 172.26.164.86 source 172.26.164.85
  auto-recovery

interface Vlan1

interface Vlan3175
  no shutdown
  no ip redirects
  ip address 172.26.164.231/24
  no ipv6 redirects

interface port-channel9
  description IB-Mgmt
  switchport access vlan 3175
  spanning-tree port type normal
  vpc 9

interface port-channel10
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3176
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description flexpodcl-01
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel12
  description flexpodcl-02
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

```



```

interface port-channel13
  description ucs_cluster-A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3175
  spanning-tree port type edge trunk
  mtu 9216
  vpc 13

interface port-channel14
  description ucs_cluster-B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3175
  spanning-tree port type edge trunk
  mtu 9216
  vpc 14

interface Ethernet4/1
  description node01:e3a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170
  mtu 9216
  channel-group 11 mode active
  no shutdown

interface Ethernet4/2
  description node02:e3a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170
  mtu 9216
  channel-group 12 mode active
  no shutdown

interface Ethernet4/3

interface Ethernet4/4

interface Ethernet4/17
  description Nexus-1110-X-1:Eth1
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge trunk
  no shutdown

interface Ethernet4/18

interface Ethernet4/19
  description Nexus-1110-X-2:Eth1
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge trunk
  no shutdown

interface Ethernet4/20

interface Ethernet4/25

interface Ethernet4/26

```

```

interface Ethernet4/27
  description ucs_cluster-A:1/27
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3175
  mtu 9216
  channel-group 13 mode active
  no shutdown

interface Ethernet4/28
  description ucs_cluster-B:1/28
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3175
  mtu 9216
  channel-group 14 mode active
  no shutdown

interface Ethernet4/41
  description VPC Peer nexus_B:4/41
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3176
  channel-group 10 mode active
  no shutdown

interface Ethernet4/42

interface Ethernet4/43
  description VPC Peer nexus_B:4/43
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3176
  channel-group 10 mode active
  no shutdown

interface Ethernet4/44
  description IB-Mgmt:1/6
  switchport access vlan 3175
  channel-group 9 mode active
  no shutdown

interface mgmt0
  ip address 172.26.164.85/24
  line vty

switchback
!Running config for vdc: storage-A

switchto vdc storage-A

!Command: show running-config
!Time: Mon May  6 15:46:47 2013

version 6.1(2)
feature-set fcoe
switchname storage-A
feature npiv
feature lacp
feature lldp
username admin password 5 $1$S4mnhDT9$M4z5JEM1XAo2C1A/ROJ6Y1 role vdc-admin
ip domain-lookup
system default switchport

```

```

snmp-server user admin vdc-admin auth md5 0xb2d55a5c14a50f69c3d2446d77efcf43 priv
0xb2d55a5c14a50f69c3d2446d77efcf43 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ip route 0.0.0.0/0 172.26.164.1
vlan 1,101
vlan 101
    fcoe vsan 101
    name FCoE_Fabric_A
vsan database
    vsan 101 name "FCoE_Fabric_A"
device-alias database
    device-alias name fcp_lif01a pwwn 20:06:00:a0:98:37:78:08
    device-alias name fcp_lif02a pwwn 20:08:00:a0:98:37:78:08
    device-alias name VM-Host-Infra-01 pwwn 20:00:00:25:b5:01:0a:00
    device-alias name VM-Host-Infra-02 pwwn 20:00:00:25:b5:01:0a:01

device-alias commit

fcdomain fcid database
    vsan 101 wwn 22:c4:54:7f:ee:f6:29:3f fcid 0x840000 dynamic
    vsan 101 wwn 20:00:00:25:b5:01:0a:00 fcid 0x840001 dynamic
    !
    [VM-Host-Infra-01]
    vsan 101 wwn 20:00:00:25:b5:01:0a:01 fcid 0x840002 dynamic
    !
    [VM-Host-Infra-02]
    vsan 101 wwn 50:0a:09:82:88:32:87:03 fcid 0x840020 dynamic
    vsan 101 wwn 20:06:00:a0:98:37:78:08 fcid 0x840021 dynamic
    !
    [fcp_lif01a]
    vsan 101 wwn 50:0a:09:82:88:12:87:64 fcid 0x840040 dynamic
    vsan 101 wwn 20:08:00:a0:98:37:78:08 fcid 0x840041 dynamic
    !
    [fcp_lif02a]

interface port-channel2
    description ucs_cluster Fabric A
    switchport mode trunk
    switchport trunk allowed vlan 101

interface vfc437
    bind interface Ethernet4/37
    switchport trunk allowed vsan 101
    switchport description node01:FCoE
    no shutdown

interface vfc438
    bind interface Ethernet4/38
    switchport trunk allowed vsan 101
    switchport description node02:FCoE
    no shutdown

interface vfc-po2
    bind interface port-channel2
    switchport description ucs_cluster-A:FCoE
    no shutdown
vsan database
    vsan 101 interface vfc-po2
    vsan 101 interface vfc437
    vsan 101 interface vfc438

interface Ethernet4/29

interface Ethernet4/30

```

```

interface Ethernet4/31
  description ucs_cluster-A:1/31
  switchport mode trunk
  switchport trunk allowed vlan 101
  channel-group 2 mode active
  no shutdown

interface Ethernet4/32
  description ucs_cluster-A:1/32
  switchport mode trunk
  switchport trunk allowed vlan 101
  channel-group 2 mode active
  no shutdown

interface Ethernet4/37
  description var_node01:3b
  switchport mode trunk
  switchport trunk allowed vlan 101
  no shutdown

interface Ethernet4/38
  description var_node02:4b
  switchport mode trunk
  switchport trunk allowed vlan 101
  no shutdown

interface Ethernet4/39

interface Ethernet4/40

interface mgmt0
  ip address 172.26.164.87/24
line vty
!Active Zone Database Section for vsan 101
zone name VM-Host-Infra-01_A vsan 101
  member pwnn 20:06:00:a0:98:37:78:08
!    [fcp_lif01a]
  member pwnn 20:08:00:a0:98:37:78:08
!    [fcp_lif02a]
  member pwnn 20:00:00:25:b5:01:0a:00
!    [VM-Host-Infra-01]

zone name VM-Host-Infra-02_A vsan 101
  member pwnn 20:00:00:25:b5:01:0a:01
!    [VM-Host-Infra-02]
  member pwnn 20:06:00:a0:98:37:78:08
!    [fcp_lif01a]
  member pwnn 20:08:00:a0:98:37:78:08
!    [fcp_lif02a]

zoneset name FlexPod vsan 101
  member VM-Host-Infra-01_A
  member VM-Host-Infra-02_A

zoneset activate name FlexPod vsan 101
do clear zone database vsan 101
!Full Zone Database Section for vsan 101
zone name VM-Host-Infra-01_A vsan 101
  member pwnn 20:06:00:a0:98:37:78:08
!    [fcp_lif01a]
  member pwnn 20:08:00:a0:98:37:78:08
!    [fcp_lif02a]
  member pwnn 20:00:00:25:b5:01:0a:00

```

```

!                               [VM-Host-Infra-01]

zone name VM-Host-Infra-02_A vsan 101
    member pwwn 20:00:00:25:b5:01:0a:01
!                               [VM-Host-Infra-02]
    member pwwn 20:06:00:a0:98:37:78:08
!                               [fcp_lif01a]
    member pwwn 20:08:00:a0:98:37:78:08
!                               [fcp_lif02a]

zoneset name FlexPod vsan 101
    member VM-Host-Infra-01_A
    member VM-Host-Infra-02_A

```

## Cisco Nexus 7000 B

```

L01-7004-2# sh run vdc-all
!Running config for default vdc: L01-7004-2

!Command: show running-config
!Time: Mon May  6 11:46:58 2013

version 6.1(2)
license grace-period
license fcoe module 4

switchname L01-7004-2
system admin-vdc
system qos
    service-policy type network-qos default-nq-7e-policy
install feature-set fcoe
vdc L01-7004-2 id 1
    cpu-share 5
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    limit-resource monitor-session-erspan-dst minimum 0 maximum 23
    limit-resource vrf minimum 2 maximum 4096
    limit-resource port-channel minimum 0 maximum 768
    limit-resource u4route-mem minimum 96 maximum 96
    limit-resource u6route-mem minimum 24 maximum 24
    limit-resource m4route-mem minimum 58 maximum 58
    limit-resource m6route-mem minimum 8 maximum 8
    limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc flexpod id 2
    limit-resource module-type f2
    cpu-share 5
    allocate interface Ethernet4/1-4,Ethernet4/17-20,Ethernet4/25-28,Ethernet4/41-44
    boot-order 1
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    limit-resource monitor-session-erspan-dst minimum 0 maximum 23
    limit-resource vrf minimum 2 maximum 4096
    limit-resource port-channel minimum 0 maximum 768
    limit-resource u4route-mem minimum 8 maximum 8
    limit-resource u6route-mem minimum 4 maximum 4
    limit-resource m4route-mem minimum 8 maximum 8
    limit-resource m6route-mem minimum 5 maximum 5
    limit-resource monitor-session-inband-src minimum 0 maximum 1
vdc storage-B id 3 type storage
    limit-resource module-type f1 f2
    allow feature-set fcoe

```

```

cpu-share 5
allocate interface Ethernet4/29-32,Ethernet4/37-40
boot-order 1
limit-resource vlan minimum 16 maximum 4094
limit-resource monitor-session minimum 0 maximum 2
limit-resource monitor-session-erspan-dst minimum 0 maximum 23
limit-resource vrf minimum 2 maximum 4096
limit-resource port-channel minimum 0 maximum 768
limit-resource u4route-mem minimum 8 maximum 8
limit-resource u6route-mem minimum 4 maximum 4
limit-resource m4route-mem minimum 8 maximum 8
limit-resource m6route-mem minimum 5 maximum 5
limit-resource monitor-session-inband-src minimum 0 maximum 1

feature telnet

username admin password 5 $1$EhCdf1Sj$V9LF7KntU3XmtV.AAxAoj/ role network-admin
ip domain-lookup
ip domain-name ridgepod.org
ip name-server 64.102.6.247
snmp-server user admin network-admin auth md5 0x67fe818704f4adf0a66d1a820695c51f priv
0x67fe818704f4adf0a66d1a820695c51f localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community RO group network-operator
snmp-server community RW group network-admin
ntp server 172.26.164.254 use-vrf management

vrf context management
  ip route 0.0.0.0/0 172.26.164.1
vlan 1

vdc storage-B id 3
  allocate fcoe-vlan-range 102

interface mgmt0
  ip address 172.26.164.82/24
clock timezone EST -5 0
clock summer-time EDT 2 sun march 02:00 1 sun nov 02:00 60
cli alias name flexpod switchto vdc flexpod
cli alias name wr copy run start
cli alias name fp switchto vdc flexpod
cli alias name storage switchto vdc storage-B
line console
line vty
boot kickstart bootflash:/n7000-s2-kickstart.6.1.2.bin sup-1
boot system bootflash:/n7000-s2-dk9.6.1.2.bin sup-1

!Running config for vdc: flexpod

switchto vdc flexpod

!Command: show running-config
!Time: Mon May  6 11:46:58 2013

version 6.1(2)
switchname flexpod

cfs eth distribute
feature udd
feature interface-vlan

```

```

feature lACP
feature vPC

username admin password 5 $1$6YvWm.3L$L4pjA8bu3mh./ZXzKInpw0 role vdc-admin
ip domain-lookup
ip domain-name ridgepod.org
ip name-server 64.102.6.247
snmp-server user admin vdc-admin auth md5 0x9a15f8c611020d45cc96832e9e228e42 priv
0x9a15f8c611020d45cc96832e9e228e42 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community RW group vdc-admin
snmp-server community RO group vdc-operator

vrf context management
  ip route 0.0.0.0/0 172.26.164.1
vlan 1-2,3170,3173-3176
vlan 2
  name Native-VLAN
vlan 3170
  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 edge bpdufilter default
spanning-tree port type network default
vpc domain 1
  role priority 20
  peer-keepalive destination 172.26.164.85 source 172.26.164.86
  peer-gateway
  auto-recovery

interface Vlan1
  no ip redirects
  no ipv6 redirects

interface Vlan3175
  no shutdown
  no ip redirects
  ip address 172.26.164.232/24
  no ipv6 redirects

interface port-channel1

interface port-channel9
  description IB-Mgmt
  switchport
  switchport access vlan 3175
  spanning-tree port type normal
  vpc 9

interface port-channel10
  description vPC peer-link
  switchport

```

```

switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3176
spanning-tree port type network
vpc peer-link

interface port-channel11
description flexpodcl-01
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170
spanning-tree port type edge trunk
mtu 9216
vpc 11

interface port-channel12
description flexpodcl-02
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170
spanning-tree port type edge trunk
mtu 9216
vpc 12

interface port-channel13
description ucs_cluster-B
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3175
spanning-tree port type edge trunk
mtu 9216
vpc 13

interface port-channel14
description ucs_cluster-A
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3173-3175
spanning-tree port type edge trunk
mtu 9216
vpc 14

interface Ethernet4/1
description node01:e4a
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170
mtu 9216
channel-group 11 mode active
no shutdown

interface Ethernet4/2
description node02:e4a
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170
mtu 9216
channel-group 12 mode active

```



```

no shutdown

interface Ethernet4/3

interface Ethernet4/4

interface Ethernet4/17
  description Nexus-1110-X-1:Eth2
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge
  no shutdown

interface Ethernet4/18

interface Ethernet4/19
  description Nexus-1110-X-2:Eth2
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3175-3176
  spanning-tree port type edge
  no shutdown

interface Ethernet4/20

interface Ethernet4/25

interface Ethernet4/26

interface Ethernet4/27
  description ucs_cluster-B:1/27
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3175
  mtu 9216
  channel-group 14 mode active
  no shutdown

interface Ethernet4/28
  description ucs_cluster-A:1/28
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3175
  mtu 9216
  channel-group 13 mode active
  no shutdown

interface Ethernet4/41
  description VPC Peer nexus_A:4/41
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3176
  spanning-tree port type network
  channel-group 10 mode active
  no shutdown

interface Ethernet4/42

```

```

interface Ethernet4/43
  description VPC Peer nexus_A:4/43
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3173-3176
  spanning-tree port type network
  channel-group 10 mode active
  no shutdown

interface Ethernet4/44
  description IB-Mgmt:1/5
  switchport
  switchport access vlan 3175
  channel-group 9 mode active
  no shutdown

interface mgmt0
  ip address 172.26.164.86/24
cli alias name wr copy run start
line vty

switchback
!Running config for vdc: storage-B

switchto vdc storage-B

!Command: show running-config
!Time: Mon May 6 11:46:59 2013

version 6.1(2)
feature-set fcoe
switchname storage-B
feature npiv
feature lacp
feature lldp
username admin password 5 $1$8LIpCPDC$koy/Jfm305MsS/yZIuQnT. role vdc-admin
ip domain-lookup
system default switchport
snmp-server user admin vdc-admin auth md5 0x9aaedc30bf8294fe1f1f175fb1e9a5dd priv
0x9aaedc30bf8294fe1f1f175fb1e9a5dd localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ip route 0.0.0.0/0 172.26.164.1
vlan 1,102
vlan 102
  fcoe vsan 102
  name FCoE_Fabric_B
vsan database
  vsan 102 name "FCoE_Fabric_B"
device-alias database
  device-alias name fcp_lif01b pwwn 20:07:00:a0:98:37:78:08
  device-alias name fcp_lif02b pwwn 20:09:00:a0:98:37:78:08
  device-alias name VM-Host-Infra-01 pwwn 20:00:00:25:b5:01:0b:00
  device-alias name VM-Host-Infra-02 pwwn 20:00:00:25:b5:01:0b:01

device-alias commit

fcdomain fcid database
  vsan 102 wwn 50:0a:09:82:88:12:87:64 fcid 0x870000 dynamic
  vsan 102 wwn 50:0a:09:84:88:32:87:03 fcid 0x870020 dynamic

```

```

vsan 102 wwn 20:07:00:a0:98:37:78:08 fcid 0x870021 dynamic
! [fcp_lif01b]
vsan 102 wwn 20:08:00:a0:98:37:78:08 fcid 0x870001 dynamic
vsan 102 wwn 50:0a:09:84:88:12:87:64 fcid 0x870002 dynamic
vsan 102 wwn 20:09:00:a0:98:37:78:08 fcid 0x870003 dynamic
! [fcp_lif02b]
vsan 102 wwn 22:c5:54:7f:ee:f6:34:ff fcid 0x870040 dynamic
vsan 102 wwn 20:00:00:25:b5:01:0b:00 fcid 0x870041 dynamic
! [VM-Host-Infra-01]
vsan 102 wwn 20:00:00:25:b5:01:0b:01 fcid 0x870042 dynamic
! [VM-Host-Infra-02]

interface port-channel2
  description ucs_cluster Fabric B
  switchport mode trunk
  switchport trunk allowed vlan 102

interface vfc437
  bind interface Ethernet4/37
  switchport trunk allowed vsan 102
  switchport description node02:FCoE
  no shutdown

interface vfc438
  bind interface Ethernet4/38
  switchport trunk allowed vsan 102
  switchport description node01:FCoE
  no shutdown

interface vfc-po2
  bind interface port-channel2
  switchport description ucs_cluster-B:FCoE
  no shutdown
vsan database
  vsan 102 interface vfc-po2
  vsan 102 interface vfc437
  vsan 102 interface vfc438

interface Ethernet4/29

interface Ethernet4/30

interface Ethernet4/31
  description ucs_cluster-B:1/31
  switchport mode trunk
  switchport trunk allowed vlan 102
  channel-group 2 mode active
  no shutdown

interface Ethernet4/32
  description ucs_cluster-B:1/32
  switchport mode trunk
  switchport trunk allowed vlan 102
  channel-group 2 mode active
  no shutdown

interface Ethernet4/37
  description var_node02:3b
  switchport mode trunk
  switchport trunk allowed vlan 102
  no shutdown

interface Ethernet4/38
  description var_node01:4b

```

```

switchport mode trunk
switchport trunk allowed vlan 102
no shutdown

interface Ethernet4/39

interface Ethernet4/40

interface mgmt0
 ip address 172.26.164.88/24
cli alias name wr copy run start
line vty
!Active Zone Database Section for vsan 102
zone name VM-Host-Infra-01_B vsan 102
    member pwnn 20:00:00:25:b5:01:0b:00
!        [VM-Host-Infra-01]
    member pwnn 20:07:00:a0:98:37:78:08
!        [fcp_lif01b]
    member pwnn 20:09:00:a0:98:37:78:08
!        [fcp_lif02b]

zone name VM-Host-Infra-02_B vsan 102
    member pwnn 20:07:00:a0:98:37:78:08
!        [fcp_lif01b]
    member pwnn 20:09:00:a0:98:37:78:08
!        [fcp_lif02b]
    member pwnn 20:00:00:25:b5:01:0b:01
!        [VM-Host-Infra-02]

zoneset name FlexPod vsan 102
    member VM-Host-Infra-01_B
    member VM-Host-Infra-02_B

zoneset activate name FlexPod vsan 102
do clear zone database vsan 102
!Full Zone Database Section for vsan 102
zone name VM-Host-Infra-01_B vsan 102
    member pwnn 20:00:00:25:b5:01:0b:00
!        [VM-Host-Infra-01]
    member pwnn 20:07:00:a0:98:37:78:08
!        [fcp_lif01b]
    member pwnn 20:09:00:a0:98:37:78:08
!        [fcp_lif02b]

zone name VM-Host-Infra-02_B vsan 102
    member pwnn 20:07:00:a0:98:37:78:08
!        [fcp_lif01b]
    member pwnn 20:09:00:a0:98:37:78:08
!        [fcp_lif02b]
    member pwnn 20:00:00:25:b5:01:0b:01
!        [VM-Host-Infra-02]

zoneset name FlexPod vsan 102
    member VM-Host-Infra-01_B
    member VM-Host-Infra-02_B

```



















































---

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

Copyright © 2013 Cisco Systems, Inc. All rights reserved