




# FlexPod Data Center with VMware vSphere 5.1Update 1

Deployment Guide for FlexPod with VMware vSphere 5.1Update1

Last Updated: February 3, 2014



Building Architectures to Solve Business Problems



---

## About the Authors

---

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

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

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

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

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

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

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

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

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

Lindsey Street is a systems architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innova-

---

tive vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Master's of Science in Information Security from East Carolina University.

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

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

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

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



# About Cisco Validated Design (CVD) Program

---

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit:

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

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.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned 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. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2014 Cisco Systems, Inc. All rights reserved.



# FlexPod Data Center with VMware vSphere 5.1 Update 1

---

## Overview

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

## Audience

This document describes the architecture and deployment procedures of an infrastructure composed of Cisco®, NetApp®, and VMware® virtualization that uses FCoE-based storage serving NAS and SAN protocols. The intended audience for this document include, 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 nonvirtualized solutions. VMware vSphere® built on FlexPod includes NetApp storage, NetApp Data ONTAP, Cisco networking, the Cisco Unified Computing System™ (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the



---

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

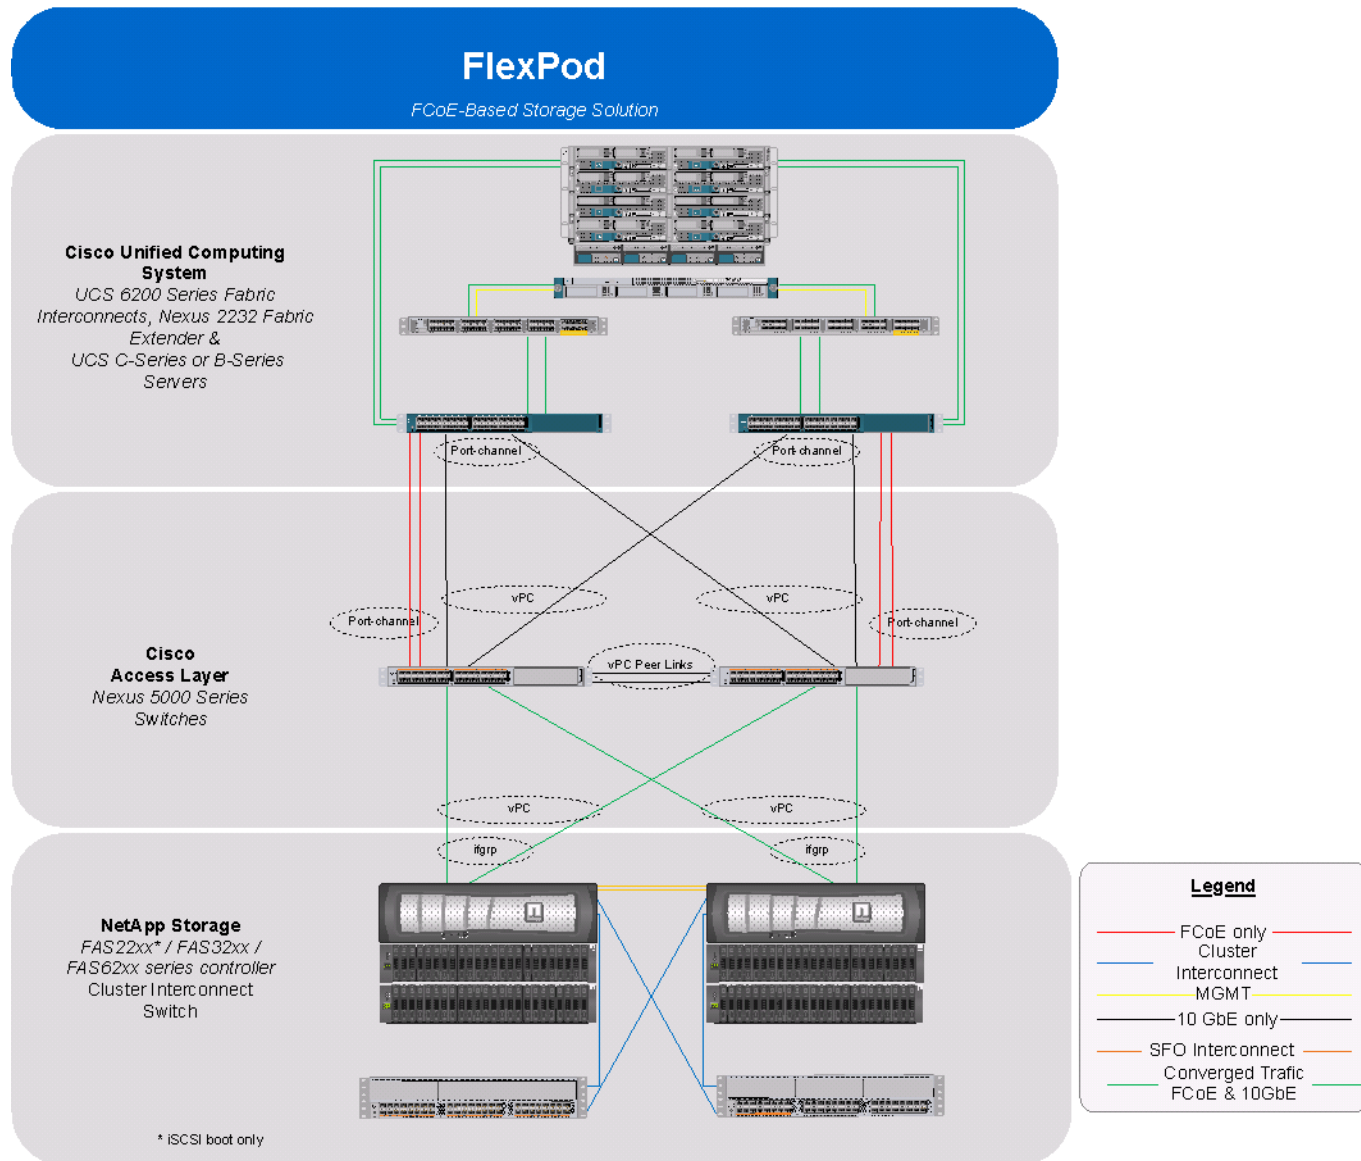
Copyright © 2014 Cisco Systems, Inc. All rights reserved.

networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or “flex” the environment to suit a customer's requirements. This is why the reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an FCoE-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

[Figure 1](#) shows the VMware vSphere built on FlexPod components and the network connections for a configuration with FCoE-based storage. This design uses the Cisco Nexus® 5548UP, Cisco Nexus 2232PP FEX, and Cisco UCS C-Series and B-Series with the Cisco UCS virtual interface card (VIC) and the NetApp FAS family of storage controllers connected in a highly available design using Cisco Virtual PortChannels (vPCs). This infrastructure is deployed to provide FCoE-booted hosts with file- and block-level access to shared storage datastores. The reference architecture reinforces the “wire-once” strategy, because as additional storage is added to the architecture; be it FC, FCoE, or 10 Gigabit Ethernet, no recabling is required from the hosts to the Cisco UCS fabric interconnect.

**Figure 1** VMware vSphere Built on Flexpod Components



The reference configuration includes:

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

Storage is provided by a NetApp FAS3250-AE (HA configuration in two chassis) operating in clustered Data ONTAP. All system and network links feature redundancy, providing end-to-end high availability (HA). For server virtualization, the deployment includes VMware vSphere. Although this is the base design, each of the components can be scaled flexibly to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capacity and throughput, and special hardware or software features can be added to introduce new capabilities.

This document guides you through the low-level steps for deploying the base architecture, as shown in [Figure 1](#). These procedures cover everything from physical cabling to compute and storage configuration to configuring virtualization with VMware vSphere.

## Software Revisions

It is important to note the software versions used in this document. [Table 1](#) details the software revisions used throughout this document.

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

Layer	Compute	Version or Release	Details
Compute	Cisco UCS Fabric Interconnect	2.1(3)	Embedded management
	Cisco UCS Rack Server C 220 M3	2.1(3)	Software bundle release
	Cisco UCS Blade Server B 200 M3	2.1(3)	Software bundle release
	Cisco eNIC	2.1.2.38	Ethernet driver for Cisco VIC
	Cisco fNIC	1.5.0.45	FCoE driver for Cisco VIC
Network	Cisco Nexus Fabric Switch	6.0(2)N2(2)	Operating system version
Storage	NetApp FAS3250-AE	Clustered Data ONTAP 8.2P4	Operating system version

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

Layer	Compute	Version or Release	Details
Software	Cisco UCS Hosts	VMware vSphere ESXi™ 5.1Update1	Operating system version
	Microsoft® .NET Framework	3.5.1	Feature enabled within Windows® operating system
	Microsoft SQL Server®	Microsoft SQL Server 2008 R2 SP1	VM (1 each): SQL Server DB
	VMware vCenter™	5.1Update1	VM (1 each): VMware vCenter
	NetApp OnCommand®	6.0	VM (1 each): OnCommand
	NetApp Virtual Storage Console (VSC)	4.2.1	Plug-in within VMware vCenter
	Cisco Nexus 1110-x	4.2(1)SP1(6.2)	Virtual services appliance
	Cisco Nexus 1000v	4.2(1)SV2(2.1a) (Advanced Edition)	Virtual services blade within the 1110-x
	NetApp NFS Plug-in for VMware vStorage APIs for Array Integration (VAAI)	1.0.20	Plug-in within VMware vCenter
	Cisco UCS Central	1.1	Manager of multiple Cisco UCS domains
	Flash Accel™ for VMware Virtual Storage Console	1.2R1	Software that manages server flash storage

## Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the 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 3](#) lists the virtual storage area networks (VSANs) necessary for deployment as outlined in this guide.

[Table 4](#) lists the configuration variables that are used throughout this document. [Table 4](#) can be completed based on the specific site variables and used in implementing the document configuration steps.

**Note**

The Cluster management and Node management interfaces will be on the Out-of-band management VLAN. Ensure that there is a Layer 3 route between the Out-of band and In-band management VLANs.

**Table 2**                      **Necessary VLANs**

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Mgmt in band	VLAN for in-band management interfaces	3175
Mgmt out of band	VLAN for out-of-band management interfaces	3170
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for NFS traffic	3172
FCoE - A	VLAN for FCoE traffic for fabric A	101
FCoE - B	VLAN for FCoE traffic for fabric B	102
vMotion	VLAN designated for the movement of VMs from one physical host to another	3173
VM Traffic	VLAN for VM application traffic	3174
Packet Control	VLAN for Packet Control traffic (Cisco Nexus 1000v)	3176

**Table 3**      **Necessary VSANs**

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

**Table 4**      **Created VMware Virtual Machine**

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

**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 8.2 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	



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

<b>Variable</b>	<b>Description</b>	<b>Customer Implementation Value</b>
<<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_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_oncommand-server-ip>>	OnCommand virtual machine management IP Address	
<<var_oncommand_server_netmask>>	Out-of-band management network netmask	
<<var_oncommand_server_gateway>>	Out-of-band management network default gateway	

**Table 5 Configuration Variables**

Variable	Description	Customer Implementation Value
<<var_ucs_central_ip>>	UCS Central management IP	
<<var_ucs_central_netmask>>	Out-of-band management network netmask	
<<var_ucs_central_gateway>>	Out-of-band management network default gateway	
<<var_ucs_central_hostname>>	UCS Central fully qualified domain name (FQDN)	
<<var_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_security_cert_vserver_authority>>	Infrastructure Vserver Security Certificate Authority	
<<var_security_cert_vserver_serial_no>>	Infrastructure Vserver security certificate serial number	
<<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_cluster_authority>>	Storage cluster security certificate authority	
<<var_security_cert_cluster_serial_no>>	Storage cluster security certificate serial number	
<<var_security_cert_node01_common_name>>	Cluster node 01 FQDN	
<<var_security_cert_node01_authority>>	Cluster node 01 security certificate authority	
<<var_security_cert_node01_serial_no>>	Cluster node 01 security certificate serial number	
<<var_security_cert_node02_common_name>>	Cluster node 02 FQDN	
<<var_security_cert_node02_authority>>	Cluster node 02 security certificate authority	

**Table 5** Configuration Variables

Variable	Description	Customer Implementation Value
<<var_security_cert_node02_serial_no>	Cluster node 02 security certificate serial number	
<<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_B_hostname>>	Cisco Nexus B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_ib-mgmt_vlan_id>>	In-band management network VLAN ID	
<<var_native_vlan_id>>	Native VLAN ID	
<<var_oob-mgmt_vlan_id>>	Out-of-band management network VLAN ID	
<<var_nfs_vlan_id>>	NFS VLAN ID	
<<var_pkt-ctrl_vlan_id>>	Cisco Nexus 1000v packet control VLAN ID	
<<var_vmotion_vlan_id>>	VMware vMotion® VLAN ID	
<<var_vm-traffic_vlan_id>>	VM traffic VLAN ID	
<<var_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	

**Table 5** Configuration Variables

Variable	Description	Customer Implementation Value
<<var_vsan_b_id>>	Fabric B VSAN ID	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucs_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A, out-of-band management IP address	
<<var_ucs_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucs_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	
<<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	

**Table 5** Configuration Variables

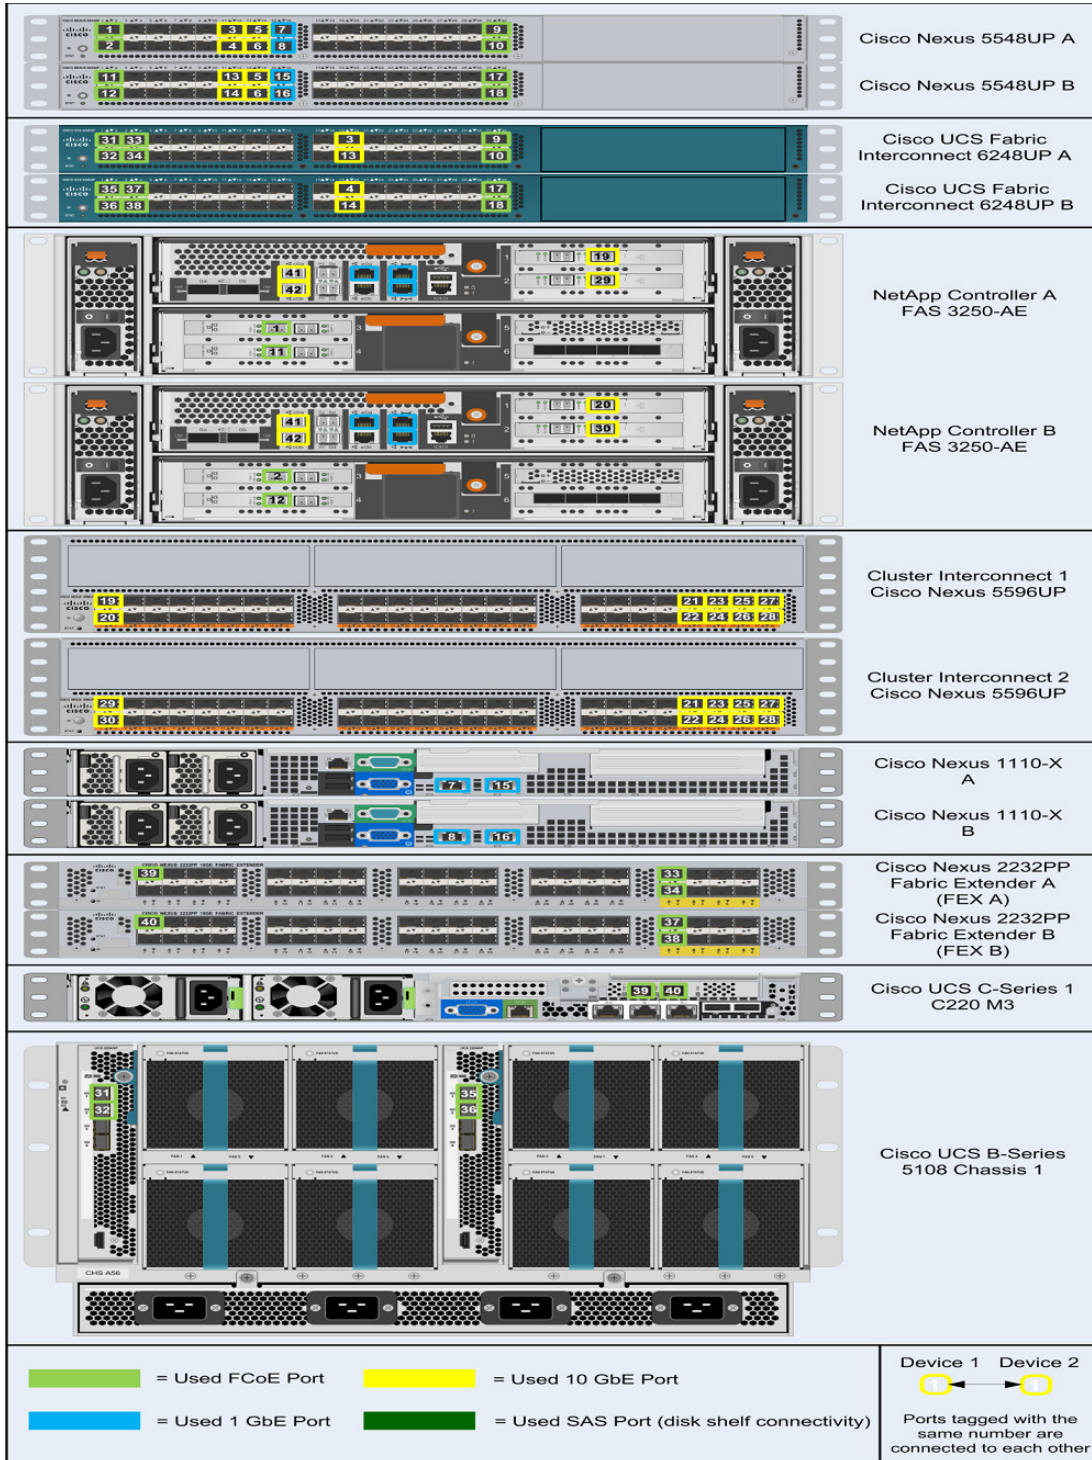
Variable	Description	Customer Implementation Value
<<var_vsadmin_password>>	Password for VS admin account	
<<var_vserver_mgmt_ip>>	Management IP address for Vserver	
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver	
<<var_rule_index>>	Rule index number	
<<var_ftp_server>>	IP address for FTP server	
<<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	
<<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

Figure 2 shows the cabling diagram for a FlexPod configuration using clustered Data ONTAP.

**Figure 2** *Flexpod Cabling Diagram in Clustered Data ONTAP*



The information provided in [Table 6](#) through [Table 20](#) corresponds to each connection shown in [Figure 2](#).

**Table 6** *Cisco Nexus 5548 A Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5548 Switch A	Eth1/1	10GbE	NetApp controller 1	e3a	1
	Eth1/2	10GbE	NetApp controller 2	e3a	2
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/19	3
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/19	4
	Eth1/13	10GbE	Cisco Nexus 5548 B	Eth1/13	5
	Eth1/14	10GbE	Cisco Nexus 5548 B	Eth1/14	6
	Eth1/15	1GbE	Cisco Nexus 1110-X A	LOM A	7
	Eth1/16	1GbE	Cisco Nexus 1110-X B	LOM A	8
	Eth1/31	10GbE	Cisco UCS fabric interconnect A	Eth1/31	9
	Eth1/32	10GbE	Cisco UCS fabric interconnect A	Eth1/32	10
	MGMT0	1GbE	GbE management switch	Any	

**Note**

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

**Table 7** *Cisco Nexus 5548 B Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Ports	Cabling Code
Cisco Nexus 5548 Switch B	Eth1/1	10GbE	NetApp controller 1	e4a	<b>11</b>
	Eth1/2	10GbE	NetApp controller 2	e4a	<b>12</b>
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/20	<b>13</b>
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/20	<b>14</b>
	Eth1/13	10GbE	Cisco Nexus 5548 A	Eth1/13	<b>5</b>
	Eth1/14	10GbE	Cisco Nexus 5548 A	Eth1/14	<b>6</b>
	Eth1/15	1GbE	Cisco Nexus 1110-X A	LOM B	<b>15</b>
	Eth1/16	1GbE	Cisco Nexus 1110-X B	LOM B	<b>16</b>
	Eth1/31	10GbE	Cisco UCS fabric interconnect B	Eth1/31	<b>17</b>
	Eth1/32	10GbE	Cisco UCS fabric interconnect B	Eth1/32	<b>18</b>
	MGMT0	1GbE	GbE management switch	Any	

**Note**

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

**Table 8** *Cisco Nexus 5596 A Cluster Interconnect Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5596 Switch A	Eth1/1	10GbE	NetApp controller 1	e1a	<b>19</b>
	Eth1/2	10GbE	NetApp controller 2	e1a	<b>20</b>
	Eth1/41	10GbE	Cisco Nexus 5596 B	Eth1/41	<b>21</b>
	Eth1/42	10GbE	Cisco Nexus 5596 B	Eth1/42	<b>22</b>
	Eth1/43	10GbE	Cisco Nexus 5596 B	Eth1/43	<b>23</b>
	Eth1/44	10GbE	Cisco Nexus 5596 B	Eth1/44	<b>24</b>
	Eth1/45	10GbE	Cisco Nexus 5596 B	Eth1/45	<b>25</b>
	Eth1/46	10GbE	Cisco Nexus 5596 B	Eth1/46	<b>26</b>
	Eth1/47	10GbE	Cisco Nexus 5596 B	Eth1/47	<b>27</b>
	Eth1/48	10GbE	Cisco Nexus 5596 B	Eth1/48	<b>28</b>
	MGMT0	1GbE	GbE management switch	Any	



**Table 9** *Cisco Nexus 5596 B Cluster Interconnect Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5596 Switch B	Eth1/1	10GbE	NetApp controller 1	e2a	<b>29</b>
	Eth1/2	10GbE	NetApp controller 2	e2a	<b>30</b>
	Eth1/41	10GbE	Cisco Nexus 5596 A	Eth1/41	<b>21</b>
	Eth1/42	10GbE	Cisco Nexus 5596 A	Eth1/42	<b>22</b>
	Eth1/43	10GbE	Cisco Nexus 5596 A	Eth1/43	<b>23</b>
	Eth1/44	10GbE	Cisco Nexus 5596 A	Eth1/44	<b>24</b>
	Eth1/45	10GbE	Cisco Nexus 5596 A	Eth1/45	<b>25</b>
	Eth1/46	10GbE	Cisco Nexus 5596 A	Eth1/46	<b>26</b>
	Eth1/47	10GbE	Cisco Nexus 5596 A	Eth1/47	<b>27</b>
	Eth1/48	10GbE	Cisco Nexus 5596 A	Eth1/48	<b>28</b>
	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 10** *NetApp Controller 1 Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp Controller 1	e0M	100MbE	100MbE management switch	Any	
	e0a	1GbE	GbE management switch	Any	
	e0P	1 GbE	SAS shelves	ACP port	
	c0a	10GbE	NetApp controller 2	c0a	<b>41</b>
	c0b	10GbE	NetApp controller 2	c0b	<b>42</b>
	e1a	10GbE	Cisco Nexus 5596 A	Eth1/1	<b>19</b>
	e2a	10GbE	Cisco Nexus 5596 B	Eth1/1	<b>29</b>
	e3a	10GbE	Cisco Nexus 5548 A	Eth1/1	<b>1</b>
	e4a	10GbE	Cisco Nexus 5548 B	Eth1/1	<b>11</b>

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

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp Controller 2	e0M	100MbE	100MbE management switch	Any	
	e0a	1GbE	GbE management switch	Any	
	e0P	1 GbE	SAS shelves	ACP port	
	c0a	10GbE	NetApp controller 1	c0a	41
	c0b	10GbE	NetApp controller 1	c0b	42
	e1a	10GbE	Cisco Nexus 5596 A	Eth1/2	20
	e2a	10GbE	Cisco Nexus 5596 B	Eth1/2	30
	e3a	10GbE	Cisco Nexus 5548 A	Eth1/2	2
	e4a	10GbE	Cisco Nexus 5548 B	Eth1/2	12

**Table 12** *Cisco UCS Fabric Interconnect A Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS Fabric Interconnect A	Eth1/19	10GbE	Cisco Nexus 5548 A	Eth1/11	3
	Eth1/20	10GbE	Cisco Nexus 5548 B	Eth1/11	13
	Eth1/1	10GbE	Cisco UCS Chassis 1 FEX A	Port1	31
	Eth1/2	10GbE	Cisco UCS Chassis 1 FEX A	Port2	32
	Eth1/3	10GbE	Cisco Nexus 2232PP FEX A	Port2/1	33
	Eth1/4	10GbE	Cisco Nexus 2232PP FEX A	Port2/2	34
	Eth1/31	10GbE	Cisco Nexus 5548 A	Eth1/31	9
	Eth1/32	10GbE	Cisco Nexus 5548 A	Eth1/32	10
	MGMT0	1GbE	GbE management switch	Any	
	L1	1GbE	Cisco UCS fabric interconnect B	L1	
	L2	1GbE	Cisco UCS fabric interconnect B	L2	

**Table 13** *Cisco UCS Fabric Interconnect B Cabling Information*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS Fabric Interconnect B	Eth1/19	10GbE	Cisco Nexus 5548 A	Eth1/12	4
	Eth1/20	10GbE	Cisco Nexus 5548 B	Eth1/12	14
	Eth1/1	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B		35
	Eth1/2	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B		36
	Eth1/3	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B		37
	Eth1/4	10GbE	Cisco UCS Chassis FEX B/Cisco Nexus 2232PP FEX B		38
	Eth1/31	10GbE	Cisco Nexus 5548 B	Eth1/31	17
	Eth1/32	10GbE	Cisco Nexus 5548 B	Eth1/32	18
	MGMT0	1GbE	GbE management switch	Any	
	L1	1GbE	Cisco UCS fabric interconnect A	L1	
	L2	1GbE	Cisco UCS fabric interconnect A	L2	

**Table 14** *Cisco Nexus 2232PP FEX A*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 2232PP FEX A	Port 1	10GbE	Cisco UCS C-Series 1	Port 0	39
	Port 2	10GbE	Cisco UCS C-Series 1	Port 0	
	Port 3	10GbE	Cisco UCS C-Series 2	Port 0	
	Port 4	10GbE	Cisco UCS C-Series 2	Port 0	
	Port 2/1	10GbE	Cisco UCS fabric interconnect A	Eth1/1	33
	Port 2/2	10GbE	Cisco UCS fabric interconnect A	Eth1/2	34

**Table 15** *Cisco Nexus 2232PP FEX B*

Local Device	Local Port	Connection	Remote Devices	Remote Port	Cabling Code
Cisco Nexus 2232PP FEX B	Port 1	10GbE	Cisco UCS C-Series 1	Port 1	40
	Port 2	10GbE	Cisco UCS C-Series 1	Port 1	
	Port 3	10GbE	Cisco UCS C-Series 2	Port 1	
	Port 4	10GbE	Cisco UCS C-Series 2	Port 1	
	Port 2/1	10GbE	Cisco UCS fabric interconnect B	Eth1/1	37
	Port 2/2	10GbE	Cisco UCS fabric interconnect B	Eth1/2	38

**Table 16** *Cisco UCS C-Series 1*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 1	Port 0	10GbE	Cisco Nexus 2232PP FEX A	Port 1	39
	Port 1	10GbE	Cisco Nexus 2232PP FEX B	Port 1	40

**Table 17** *Cisco UCS C-Series 2*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 2	Port 0	10GbE	Cisco Nexus 2232PP FEX A	Port 2	
	Port 1	10GbE	Cisco Nexus 2232PP FEX B	Port 2	

**Table 18** *Cisco UCS C-Series 3*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 3	Port 0	10GbE	Cisco Nexus 2232PP FEX A	Port 3	
	Port 1	10GbE	Cisco Nexus 2232PP FEX B	Port 3	

**Table 19** *Cisco UCS C-Series 4*

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS C-Series Server 4	Port 0	10GbE	Cisco Nexus 2232PP FEX A	Port 4	
	Port 1	10GbE	Cisco Nexus 2232PP FEX B	Port 4	

**Table 20** *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

**Table 21** *Cisco C220M3 Card Layout for Single-wire Management*

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

## Storage Configuration

### Controller FAS32xx Series

**Table 22** *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

**Table 22**      **Controller FAS32XX Series Prerequisites**

Requirement	Reference	Comments
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

The NetApp Hardware Universe provides supported hardware and software components for the specific Data ONTAP version. It provides 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. Access the [Hardware Universe Application](#) to view the System Configuration guides. Click the **Controllers** tab to view the compatibility between Data ONTAP software versions and NetApp storage appliances with the desired specifications.
3. Alternatively, to compare components by storage appliance, click **Compare Storage Systems**.

## Controllers

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

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

## Disk Shelves

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)

## Cisco NX5596 Cluster Network Switch Configuration

**Table 23** *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)
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.

```

Abort Power On Auto Provisioning and continue with normal setup ?(yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for 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
```

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.

```
***** Transfer of file Completed Successfully *****
Copy complete, now saving to disk (please wait)...
```

5. Enter the following command to view the saved configuration file.

```
dir bootflash:
```

6. 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
```

7. 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
- 8. 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 8.2



### Note

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

## Complete the Configuration Worksheet

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

**Table 24**      *Controller FAS32XX Series Prerequisites*

Requirement	Reference	Comments
Configuration Worksheet	<a href="https://library.netapp.com/ecm/ecm_download_file/ECMP1196796">https://library.netapp.com/ecm/ecm_download_file/ECMP1196796</a>	Requires access to the <a href="#">NetApp Support site</a> .

**Table 25**      *Clustered Data ONTAP Software Installation Prerequisites*

Cluster Details	Cluster Detail Values
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>

**Table 25**      **Clustered Data ONTAP Software Installation Prerequisites**

Cluster Details	Cluster Detail Values
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.2P4 URL	<<var_url_boot_software>>

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

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

```
7
```

7. Answer yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

9. Select yes to reboot now.

```
y
```

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

```
<<var_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.

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:

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

See [Table 26](#) for prerequisites to create clustered Data ONTAP.

**Table 26** *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>>

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

Cluster Detail	Cluster Detail Value
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. Follow these instructions to activate HA and set storage failover:

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

```
Non-HA mode, Reboot node to activate HA Do you want to reboot now to set storage
failover (SFO) to HA mode? {yes, no} [yes]: Enter
```

4. After the reboot, proceed with steps for creating cluster.

5. 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.
The cluster will be connected using network switches.
```

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

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



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

7. 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>>
```

**8. 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.

**9. 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.

**10. Press **Enter** to accept the AutoSupport™ message.****11. Reboot node 01.**

```
system node reboot -node <<var_node01>>
Warning: Are you sure you want to reboot the node? {y|n}: y
```

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

```
Ctrl - C
```

**13. Select 5 to boot into maintenance mode.**

```
5
```

**14. When prompted **Continue with boot?**, enter **y**.****15. 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.

**16. Reboot the controller.**

```
halt
```



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

```
autoboot
```

18. Log in to the cluster.

19. Data ONTAP automatically assigns disks to storage controllers if the disk autoassign setting is turned on. Use the storage disk option **show -fields autoassign** command to verify this setting.

20. If disk autoassign was turned on, skip to “[Cluster Join in Clustered Data ONTAP](#)” section on [page 37](#). Else, continue to step 21.

21. Reboot node 01.

```
system node reboot -node <<var_node01>>
Warning: Are you sure you want to reboot the node? {y|n}: y
```

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

```
Ctrl - C
```

23. Select 5 to boot into maintenance mode.

```
5
```

24. When prompted Continue with boot?, enter y.

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

```
disk show -a
```



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

26. 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>>
```

27. Reboot the controller.

```
halt
```

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

```
autoboot
```

## Cluster Join in Clustered Data ONTAP

**Table 27** *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>>

**Table 27**      *Joining Cluster in Clustered Data ONTAP Prerequisites*

Cluster Detail	Cluster Detail Value
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. Follow these instructions to activate HA and set storage failover:

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

Non-HA mode, Reboot node to activate HA Do you want to reboot now to set storage
failover (SFO) to HA mode? {yes, no} [yes]: Enter
```

4. After the reboot, proceed with steps for creating cluster.
5. Data ONTAP detects that its Storage Failover Partner is part of a cluster, agree to join the same cluster.

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

```
Existing cluster interface configuration found:
```

```
Port      MTU      IP              Netmask
e1a       9000     169.254.251.110 255.255.0.0
e2a       9000     169.254.56.206  255.255.0.0
Do you want to use this configuration? {yes, no} [yes]: Enter
```



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

6. 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.

7. Set up the node.

```
Enter the node management interface port [e0M]: e0a
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
```

8. 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.

9. Press **Enter** to accept the AutoSupport message.

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

11. Reboot node 02.

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

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

```
Ctrl - C
```

13. Select 5 to boot into maintenance mode.

```
5
```

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

```
y
```

15. 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
```

16. Reboot the controller.

```
halt
```

17. At the LOADER-A prompt, enter:

```
Autoboot
```

18. Log in to the cluster.

19. Data ONTAP automatically assigns disks to storage controllers if the disk autoassign setting was turned on. Use the storage disk option **show -fields autoassign** command to verify the setting.

20. If disk autoassign was turned on, go to [“Cluster Join in Clustered Data ONTAP” section on page 37](#). Else, continue to step 21.

21. Reboot node 02.

```
system node reboot -node <<var_node02>>
Warning: Are you sure you want to reboot the node? {y|n}: y
```

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

```
Ctrl - C
```

23. Select 5 to boot into maintenance mode.

```
5
```

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

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

```
disk show -a
```

26. 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>>
```

27. Reboot the controller:

```
halt
```

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

```
autoboot
```

## 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 e0a
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 e0a
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 8.2 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.

## Aggregates in Clustered Data ONTAP

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

1. Execute the following command to create new aggregates:

```
aggr create -aggregate aggr01 -nodes <<var_node01>> -diskcount <<var_num_disks>>
aggr create -aggregate aggr02 -nodes <<var_node02>> -diskcount <<var_num_disks>>
```



### Note

- Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.
  - Start with five disks initially; you can add disks to an aggregate when there is a need for additional storage.
  - 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 8.2, you must upgrade to the latest service processor (SP) firmware to take advantage of the latest updates available for the remote management device.

1. Use the command:

```
system node service-processor show
```

2. Get the version of the Service Processor firmware that is currently running on your storage system.
3. Using a Web browser, connect to <http://support.netapp.com/NOW/cgi-bin/fw>.
4. Navigate to the Service Process Image for installation from the Data ONTAP prompt page for your storage platform.
5. Check the latest firmware version that is available for your storage platform. If your storage system is not running the latest version, proceed to the download page for the latest release of the SP firmware for your storage platform.
6. 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)**.
7. Execute step 4 on each node, if SP automatic update is not enabled.
8. View the status of the SP upgrade using steps 5-6.

## 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. Verify the status of storage failover.

```
storage failover show
```

2. Both the nodes <<var\_node01>> and <<var\_node02>> must be capable of performing a takeover.
3. Proceed to step 4 if the nodes are capable of performing a takeover.
4. 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.

5. 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
```

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

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

## 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 “a<number><letter>”, where <number> is an integer in the range [0-999] without leading zeros and <letter> is a lowercase letter.



## 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 serveral 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.ss>
```



**Note** The format for the date is <[Century][Year][Month][Day][Hour][Minute]>; for example, 201309081735.17.

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>>
```

## 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 ro <<var_snmp_community>>
```



### Note

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

## SNMPv3 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. Enter the authoritative entity's EngineID and select md5 as the authentication protocol.
3. Use the command **security snmpusers** to view the EngineID.
4. Enter an eight-character minimum-length password for the authentication protocol, when prompted.
5. Select **des** as the privacy protocol.
6. 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, execute the following steps:



### Note

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

Enable CDP on Data ONTAP:

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

## Vserver (Storage Virtual Machine)

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, ndmp}: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 1GB -type DP
```

```
volume create -vserver Infra_Vserver -volume root_vol_m02 -aggregate aggr02 -size 1GB -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. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
snapmirror modify -source-path //Infra_Vserver/rootvol -destination-path * -schedule 15min
```

## 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 -sslsv3-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

```

5. Use the security certificate show command to obtain the values for the parameters that would be needed in the step below.

6. The following is an example output in which the relevant parameters are highlighted.

```

FAS3250-Cluster::> security certificate show
Vserver      Serial Number  Common Name                                     Type
-----
FAS3250-Cluster
          52173806      FAS3250-Cluster.cert                         server
Certificate Authority: FAS3250-Cluster.cert
Expiration Date: Sat Aug 23 15:53:03 2014

```

7. The required parameters are:

- -common-name : FAS3250-Cluster.cert
- -ca : Certificate Authority: FAS3250-Cluster.cert
- -serial : 52173806

8. Use these values in the command as shown:

```

security ssl modify -vserver Infra_Vserver -common-name
<<var_security_cert_vserver_common_name>>
-server-enabled true -client-enabled false -ca
<<var_security_certificate_vserver_authority>> -serial
<<var_security_certificate_vserver_serial_no>>

security ssl modify -vserver <<var_clustername>> -common-name
<<var_security_cert_cluster_common_name>>
-server-enabled true -client-enabled false -ca
<<var_security_certificate_cluster_authority>> -serial
<<var_security_certificate_cluster_serial_no>>

security ssl modify -vserver <<var_node01>> -common-name
<<var_security_cert_node01_common_name>>

```

```

-server-enabled true -client-enabled false -ca
<<var_security_certificate_node01_authority>> -serial
<<var_security_certificate_node01_serial_no>>

security ssl modify -vserver <<var_node02>>-common-name
<<var_security_cert_node02_common_name>>
-server-enabled true -client-enabled false -ca
<<var_security_certificate_node02_authority>> -serial
<<var_security_certificate_node02_serial_no>>

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**

Vserver services web access create –name compat –role admin –vserver <<var\_clustername>>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 none
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 -vsriver 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 -vsriver Infra_Vserver -volume esxi_boot -aggregate aggr01 -size
100g -state online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vsriver Infra_Vserver -volume OnCommandDB -aggregate aggr02 -size
200g -state online -policy Flexpod -junction-path /OnCommandDB -space-guarantee
none -percent-snapshot-space 0

snapmirror update-ls-set -source-path //Infra_Vserver/rootvol

```

## LUN in Clustered Data ONTAP

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

```

lun create -vsriver Infra_Vserver -volume esxi_boot -lun VM-Host-Infra-01 -size
10g -ostype vmware -space-reserve disabled
lun create -vsriver 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 -vsriver Infra_Vserver -volume infra_datastore_1
volume efficiency on -vsriver Infra_Vserver -volume esxi_boot
volume efficiency on -vsriver 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 -vsriver 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 -vsriver 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 3a
network interface create -vserver Infra_Vserver -lif fcp_lif01b -role data
-data-protocol fcp -home-node <<var_node01>> -home-port 4a
network interface create -vserver Infra_Vserver -lif fcp_lif02a -role data
-data-protocol fcp -home-node <<var_node02>> -home-port 3a
network interface create -vserver Infra_Vserver -lif fcp_lif02b -role data
-data-protocol fcp -home-node <<var_node02>> -home-port 4a
```

## 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 -failover-group
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
Enter a new password: <<var_vsadmin_password>>
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/gui) ? console
Enter the setup mode; setup newly or restore from backup. (setup/restore)? setup
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
Enforce strong password? (y/n) [y]: y
```



```

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

```

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

## Cisco UCS 6248UP Fabric Interconnect B

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

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

```

Enter the configuration method. (console/gui) ? console
Installer has detected the presence of a peer Fabric interconnect. This Fabric
interconnect will be added to the cluster. Continue (y/n) ? y
Enter the admin password of the peer Fabric interconnect: <<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(3a). To upgrade the Cisco UCS Manager software and the UCS 6248 Fabric Interconnect software to version 2.1(3a), see [Cisco UCS Manager Install and Upgrade Guides](#).

## Add Block of IP Addresses for KVM Access

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



### Note

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

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

## Synchronize Cisco UCS to NTP

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

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

## Edit Chassis Discovery Policy

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

To modify the chassis discovery policy, follow these steps:

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

## Enable Server and Uplink Ports

To enable server and uplink ports, follow these steps:

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

**Figure 3** Configured Server Ports

The screenshot displays the Cisco Unified Computing System Manager interface. The left navigation pane shows the hierarchy: Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module > Ethernet Ports. The main pane shows a table of Ethernet Ports with columns: Slot, Port ID, MAC, If Role, If Type, Overall Status, and Administrative State. Ports 1 through 18 are configured as Server ports (Physical, Up, Enabled). Ports 19 and 20 are configured as Uplink ports (Physical, Admin Down, Disabled). Ports 21 through 29 are unconfigured (Physical, Sfp Not Pres..., Disabled).

Slot	Port ID	MAC	If Role	If Type	Overall Status	Administrative State
1	1	54:7F:EE:23:52...	Server	Physical	Up	Enabled
1	2	54:7F:EE:23:52...	Server	Physical	Up	Enabled
1	3	54:7F:EE:23:52...	Server	Physical	Up	Enabled
1	4	54:7F:EE:23:52...	Server	Physical	Up	Enabled
1	5	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	6	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	7	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	8	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	9	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	10	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	11	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	12	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	13	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	14	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	15	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	16	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	17	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	18	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	19	54:7F:EE:23:52...	Unconfigured	Physical	Admin Down	Disabled
1	20	54:7F:EE:23:52...	Unconfigured	Physical	Admin Down	Disabled
1	21	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	22	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	23	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	24	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	25	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	26	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	27	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	28	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled
1	29	54:7F:EE:23:52...	Unconfigured	Physical	Sfp Not Pres...	Disabled

7. Choose ports 19 and 20 that are connected to the Cisco Nexus 5548 switches, right-click them, and choose **Configure as Uplink Port**.
8. Click **Yes** to confirm uplink ports and click **OK**.

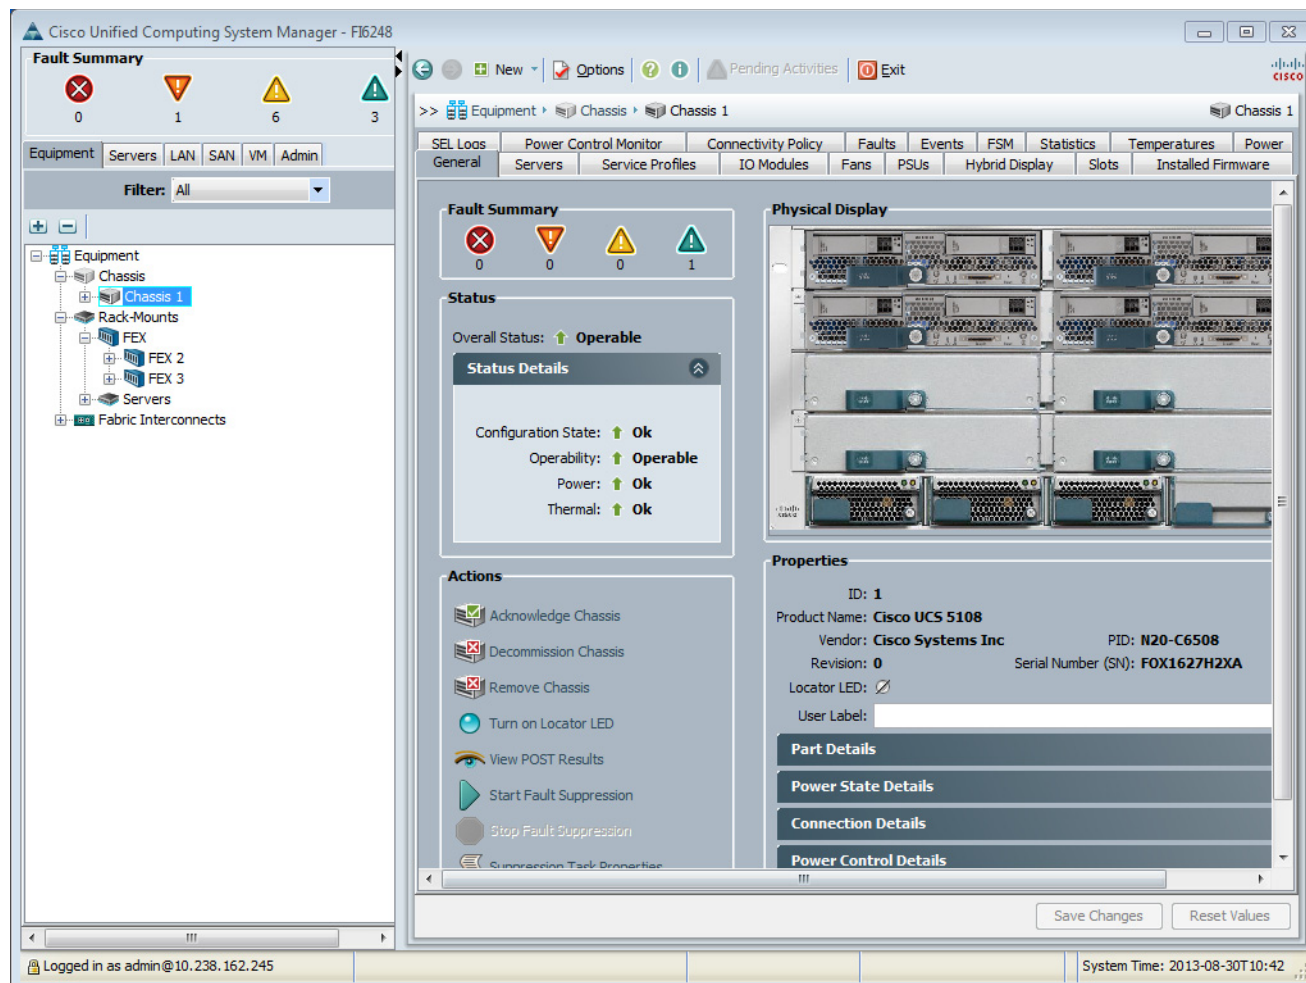
9. Choose ports 31 and 32, which will serve as FCoE uplinks to the Cisco Nexus 5548 switches; right-click them; and choose **Configure as FCoE Uplink Port**.
10. Click **Yes** to confirm FCoE uplink ports and click **OK**.
11. Choose **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate)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

## Acknowledge Cisco UCS Chassis and FEX

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

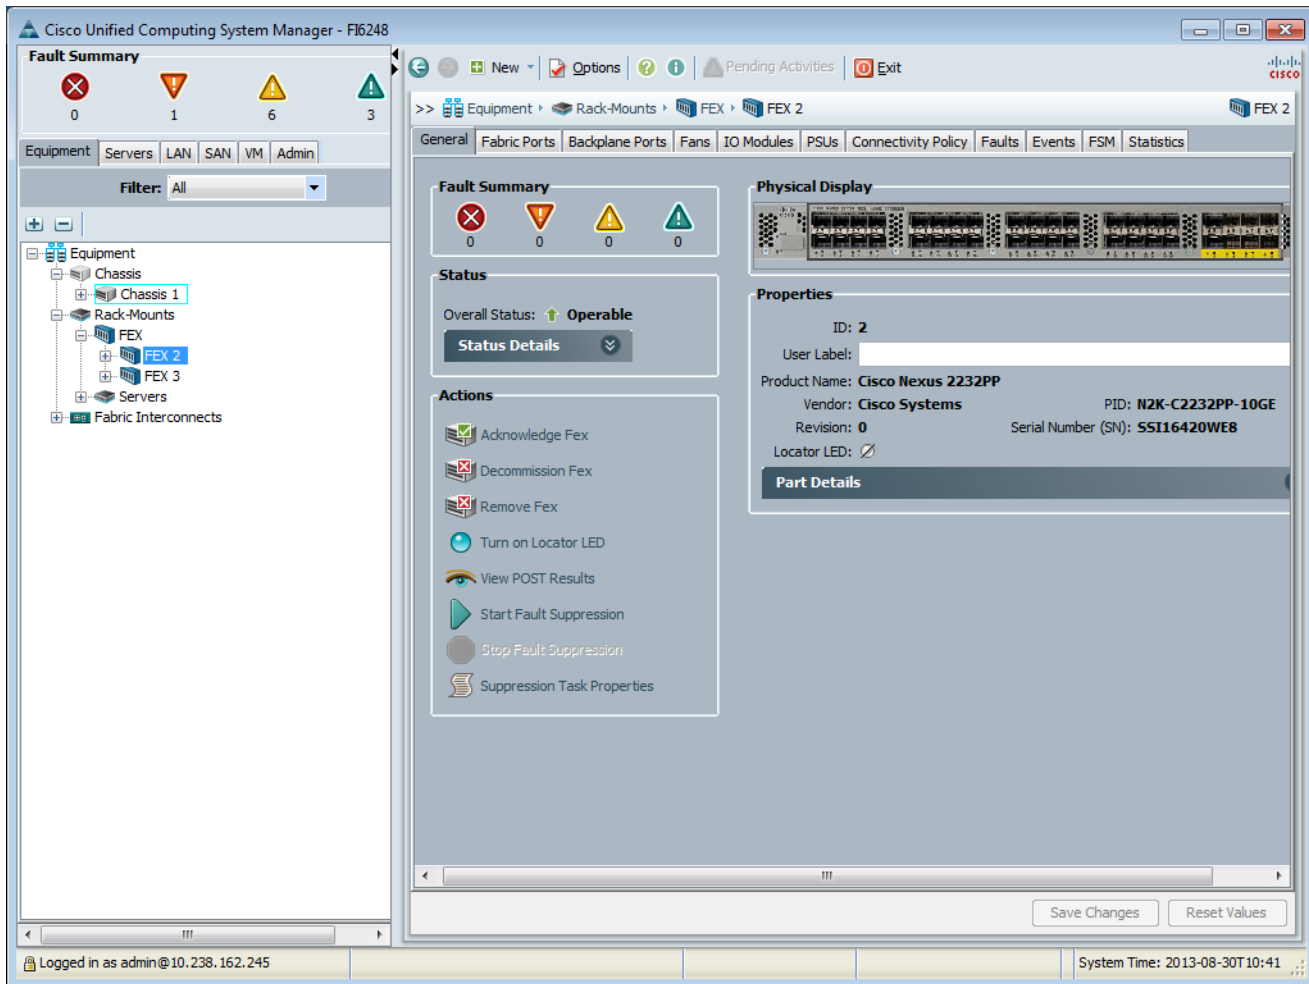
1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
2. Expand Chassis.
3. Choose each chassis that is listed, right-click on each chassis and choose **Acknowledge Chassis**.

**Figure 4 Acknowledging Cisco UCS Chassis**



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

**Figure 5** Acknowledging Cisco UCS Fabric Extenders



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

## Create Uplink Port Channels to Cisco Nexus 5548 Switches

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

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



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

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

7. Click **Next**.

**Figure 6**      *Creating Port Channels*

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



21. Click **OK**.

## Create an Organization

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



### Note

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

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

1. In Cisco UCS Manager, from the **New** drop down 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.



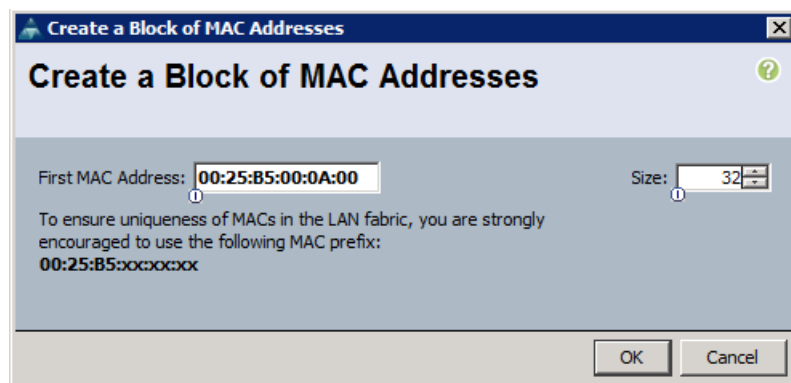
### Note

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 7** *Creating MAC Address Pool for Fabric A*



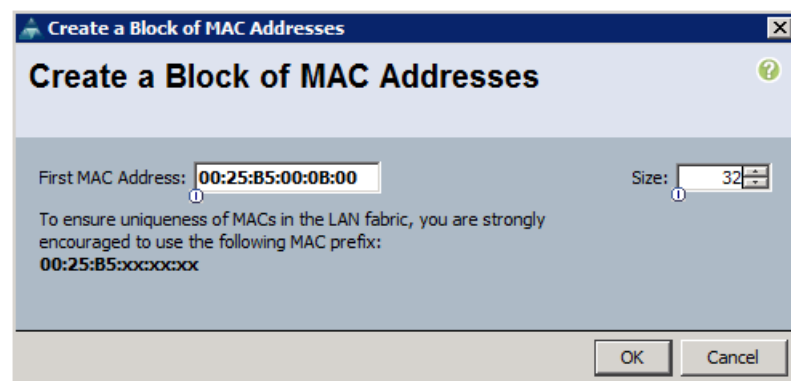
11. Click **OK**.
12. Click **Finish**.
13. In the confirmation message window, click **OK**.
14. Right-click MAC Pools under the root organization.
15. Choose Create MAC Pool to create the MAC address pool.
16. Enter MAC\_Pool\_B as the name of the MAC pool.
17. (Optional) Enter a description for the 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 8** *Creating MAC Address Pool for Fabric B*



22. Click **OK**.
23. Click **Finish**.

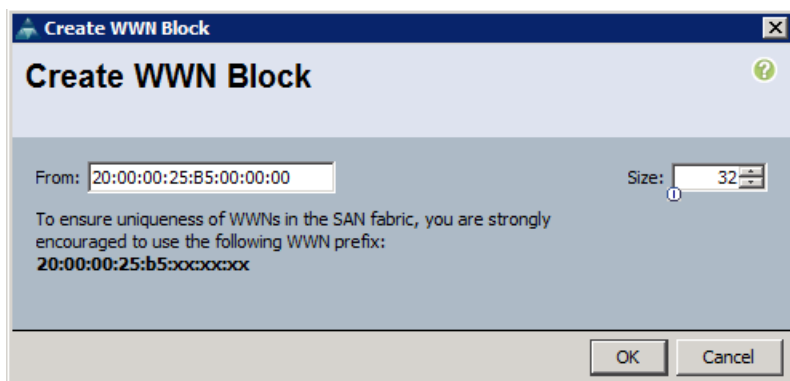
24. In the confirmation message window, click **OK**.

## 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 9**      *Creating 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 10**      **Creating 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.

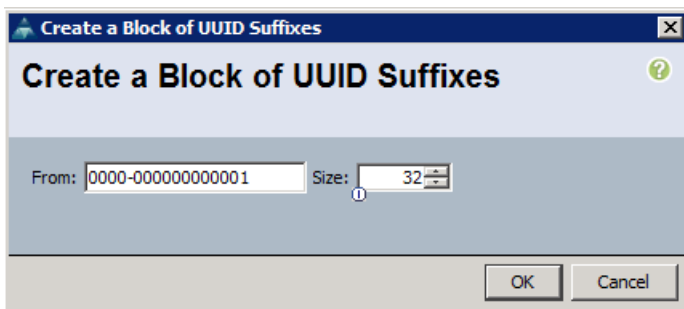
22. Click **OK**.
23. Click **Finish**.
24. Click **OK**.

## 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 11**      *Creating UUID Suffix Pool*



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 12**      **Creating VLAN for Management Traffic**

**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 13**      **Creating VLAN for NFS Traffic**

**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 14**      **Creating VLAN for vMotion**

**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 15**      **Creating VLAN for VM Traffic**

**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

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

**Figure 16**      **Creating Native VLAN**

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

## Create VSANs and FCoE Port Channels

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

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

7. Click the **Fabric A** radio button.
8. Enter <<var\_vsan\_a\_id>> as the VSAN ID for fabric A.
9. Enter <<var\_fabric\_a\_fcoe\_vlan\_id>> as the FCoE VLAN ID for fabric A.



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

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

**Figure 17**      *Creating VSAN for Fabric A*

11. Right-click VSANs.
12. Choose Create VSAN.
13. Enter VSAN\_B as the name of the VSAN for fabric B.
14. Keep the Disabled option selected for FC Zoning.
15. Click the **Fabric B** radio button.
16. Enter <<var\_vsan\_b\_id>> as the VSAN ID for fabric B.
17. Enter <<var\_fabric\_b\_fcoe\_vlan\_id>> as the FCoE VLAN ID for fabric B.



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

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

**Figure 18**      **Creating VSAN for Fabric B**

**Create VSAN**

Name:

**FC Zoning Settings**

FC Zoning: ☒ Disabled ☐ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

☐ 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:

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

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

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

## Create Host Firmware Package

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

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

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

**Figure 19**      *Creating Host Firmware Package*

**Create Host Firmware Package**

Name:

Description:

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

Blade Package:

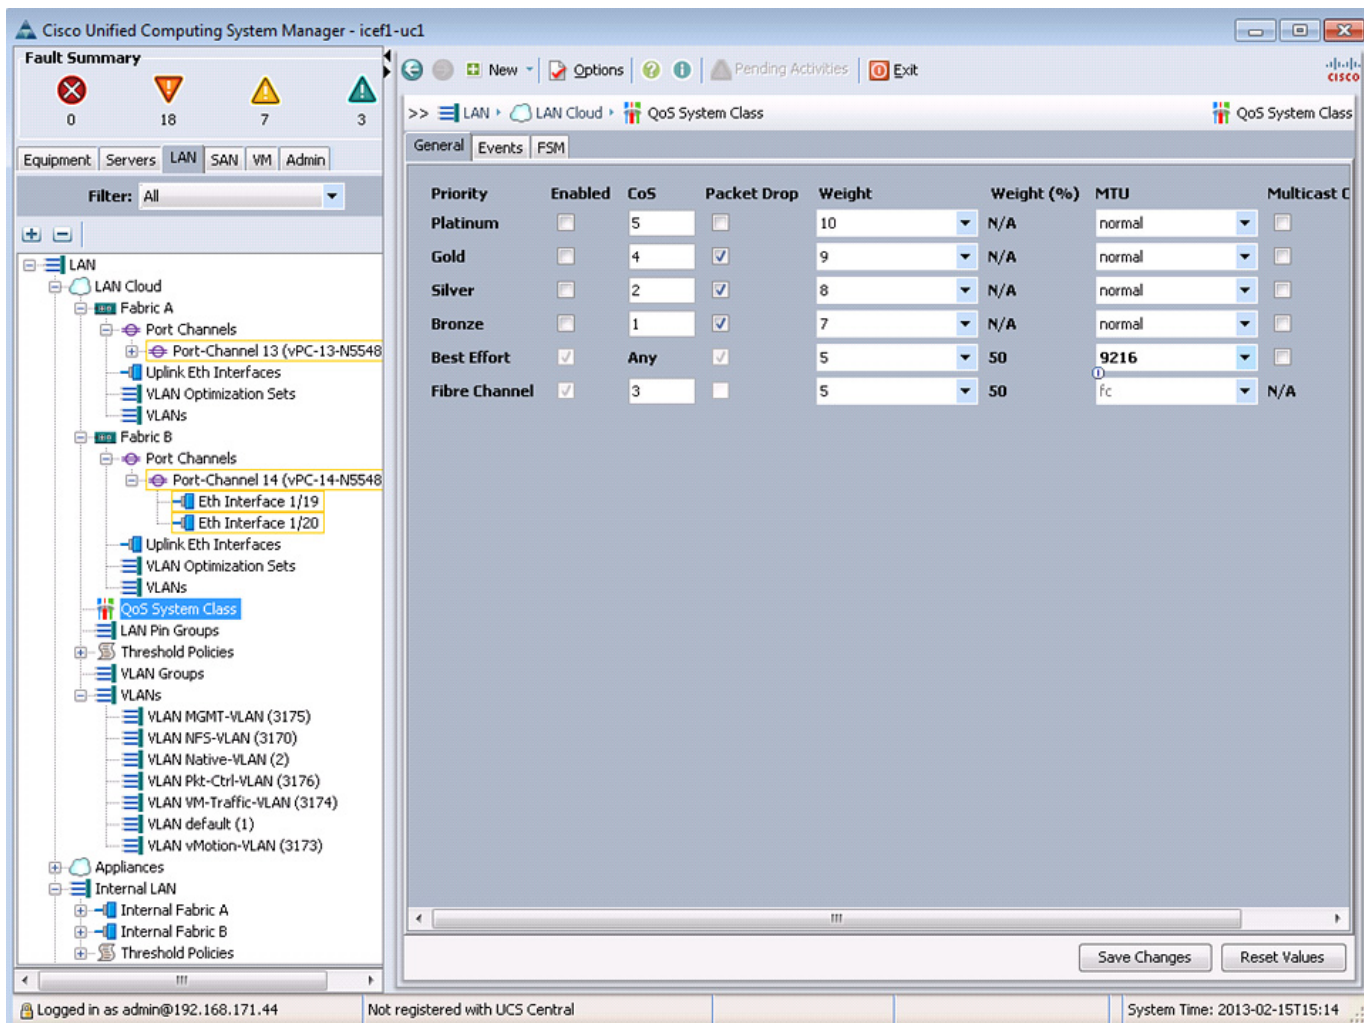
Rack Package:

## Set Jumbo Frames in Cisco UCS Fabric

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

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

**Figure 20**      **Setting Jumbo Frame**



## Create Local Disk Configuration Policy (Optional)

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

**Note**

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

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

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose **Policies > root**.
3. Right-click Local Disk Config Policies.
4. Choose Create Local Disk Configuration Policy.
5. Enter SAN-Boot as the local disk configuration policy name.
6. Change the mode to No Local Storage.
7. Click **OK** to create the local disk configuration policy.

**Figure 21**      *Creating Local Disk Configuration Policy*

**Create Local Disk Configuration Policy**

Name:

Description:

Mode:

Flex Flash

Flex Flash State: ☒ Disable ☐ Enable

OK Cancel

8. Click **OK**.

## Create Network Control Policy for Cisco Discovery Protocol

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

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

**Figure 22**      *Creating Network Control Policy*

**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

8. Click **OK**.

## Create Power Control Policy

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

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



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

**Figure 23**      **Creating Power Control Policy**

**Create Power Control Policy**

Name:

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:



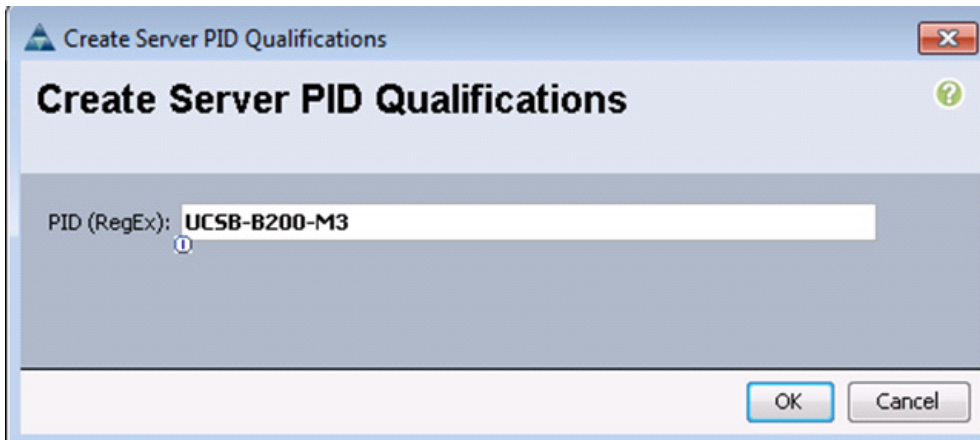
### Note

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 24**      *Creating Server PID Qualifications*

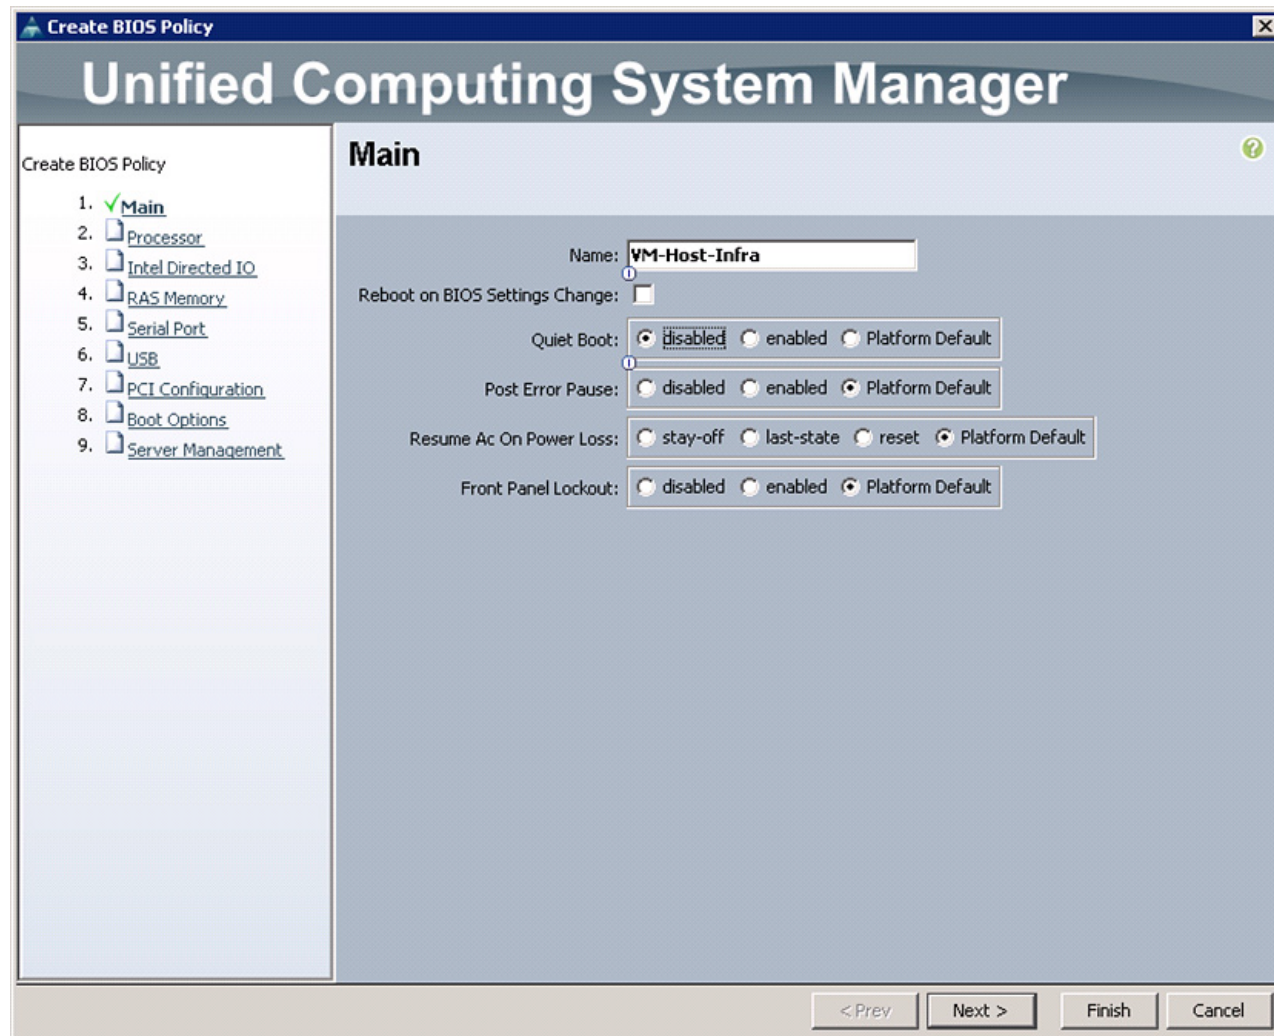


## Create Server BIOS Policy

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

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

**Figure 25**      **Creating BIOS Policy**



8. Click **OK**.

## Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

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

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

**Figure 26**      **Creating Placement Policy**

**Create Placement Policy**

Name:

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 27**      **Updating Maintenance Policy**



## Create vNIC Templates

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

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

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

**Figure 28**      *Creating vNIC Template for Fabric A*

**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

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 29**      **Creating 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 30**      *Creating 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 31**      *Creating 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 5548 A) and the B LIFs are connected to fabric B (Cisco Nexus 5548 B).

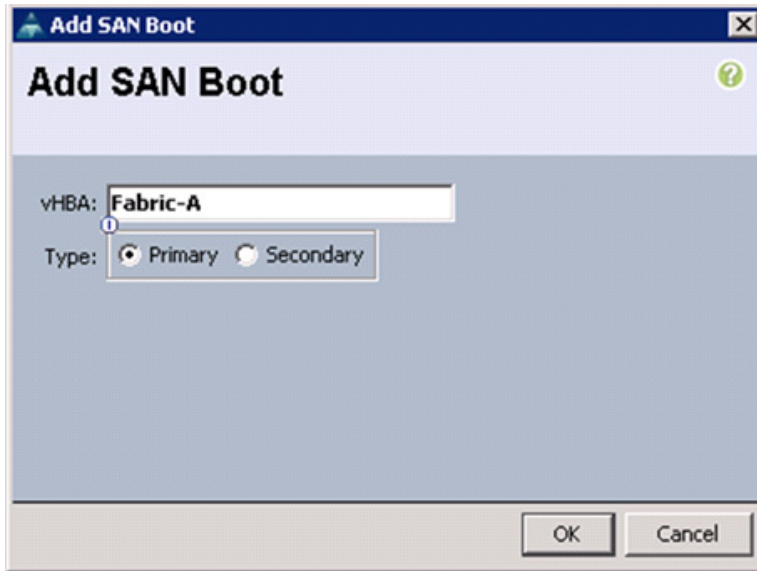
Two boot policies are configured in this procedure. The first policy configures the primary target to be fcp\_lif01a and the second boot policy configures the primary target to be fcp\_lif01b.

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

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

10. In the Add SAN Boot dialog box, enter Fabric-A in the vHBA field.
11. Make sure that the Primary radio button is selected as the SAN boot type.
12. Click **OK** to add the SAN boot initiator.

**Figure 32**      *Adding SAN Boot Initiator for Fabric A*



13. From the vHBA drop-down menu, choose Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for 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 33**      **Adding SAN Boot Target for Fabric A**

**Add SAN Boot Target**

Boot Target LUN: 0

Boot Target WWPN: 20:01:00:A0:98:3F:53:D4

Type: ☒ Primary ☐ Secondary

OK 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 34**      **Adding Secondary SAN Boot Target for Fabric A**

**Add SAN Boot Target**

Boot Target LUN: 0

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

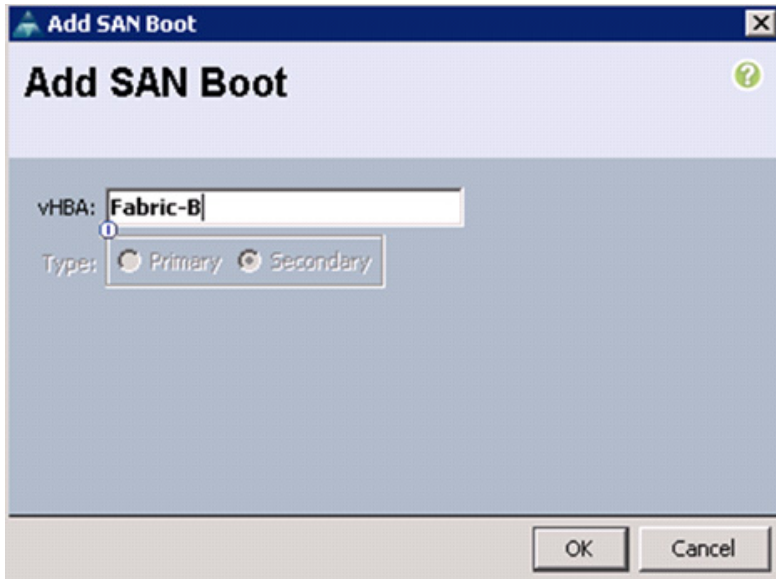
Type: ☐ Primary ☒ Secondary

OK Cancel

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

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

**Figure 35**      *Adding SAN Boot Initiator for Fabric B*



26. From the vHBA drop-down menu, choose Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN 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 36** Adding Primary SAN Boot Target for Fabric B

**Add SAN Boot Target**

Boot Target LUN: 0

Boot Target WWPN: 20:02:00:A0:98:3F:53:D4

Type: ☒ Primary ☐ Secondary

OK 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 37** Adding Secondary SAN Boot Target

**Add SAN Boot Target**

Boot Target LUN: 0

Boot Target WWPN: 20:04:00:A0:98:3F:53:D4

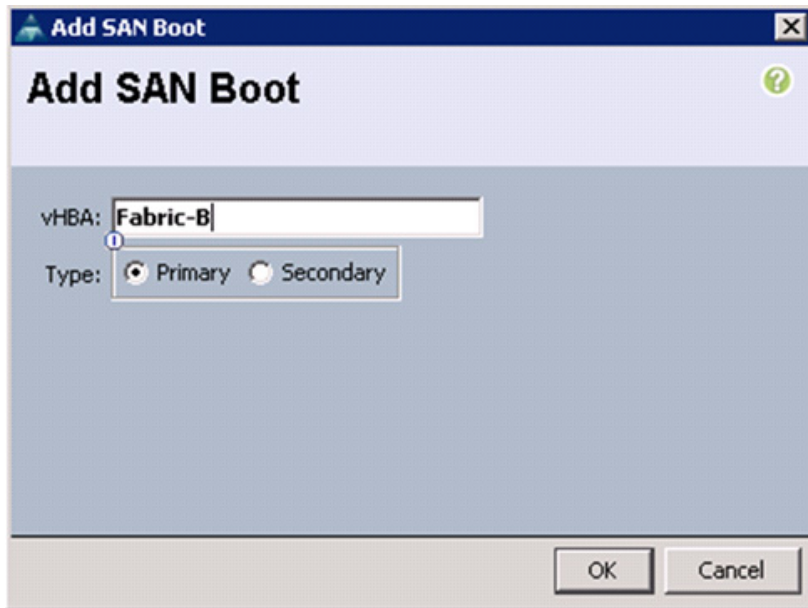
Type: ☐ Primary ☒ Secondary

OK Cancel

35. Click **OK**, and then **OK** again to create the boot policy.
36. Right-click Boot Policies again.
37. Choose **Create Boot Policy**.
38. Enter Boot-Fabric-B as the name of the boot policy.
39. (Optional) Enter a description of the boot policy.
40. Keep the Reboot on Boot Order Change check box unchecked.

41. From the Local Devices drop-down menu choose Add CD-ROM.
42. From the vHBA drop-down menu choose Add SAN Boot.
43. In the Add SAN Boot dialog box, enter Fabric-B in the vHBA box.
44. Make sure that the **Primary** radio button is selected as the SAN boot type.
45. Click **OK** to add the SAN boot initiator.

**Figure 38**      *Adding SAN Boot Initiator for Fabric B*



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



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

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

**Figure 39** Adding Primary SAN Boot Target for Fabric B

**Add SAN Boot Target**

Boot Target LUN: 0

Boot Target WWPN: 20:02:00:A0:98:3F:53:D4

Type: ☒ Primary ☐ Secondary

OK Cancel

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



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

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

**Figure 40** Adding Secondary SAN Boot Target for Fabric B

**Add SAN Boot Target**

Boot Target LUN: 0

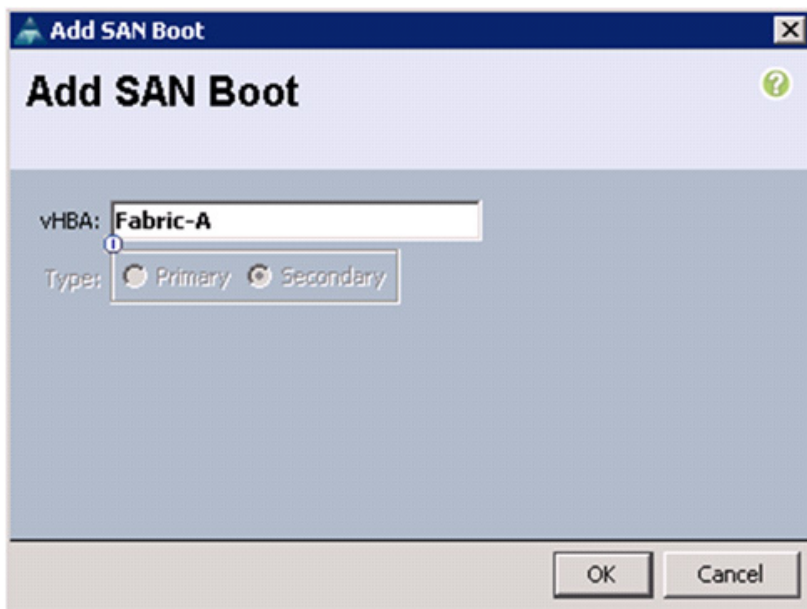
Boot Target WWPN: 20:04:00:A0:98:3F:53:D4

Type: ☐ Primary ☒ Secondary

OK Cancel

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

**Figure 41**      **Adding SAN Boot for Fabric A**



59. From the vHBA menu, choose Add SAN Boot Target.

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

61. Enter the WWPN for fcp\_lif01a.

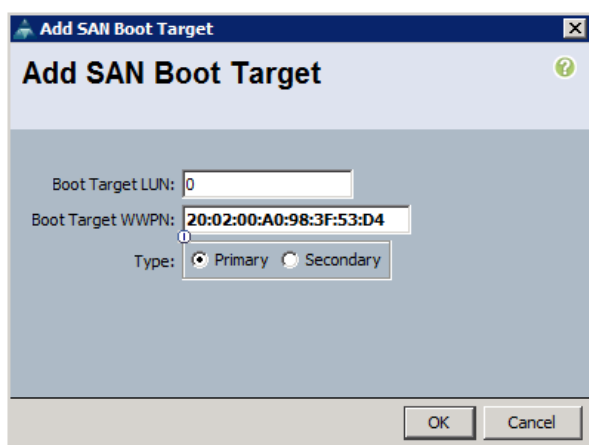


**Note** To obtain this information, log in to the storage cluster and run the **network interface show** command.

62. Keep Primary as the SAN boot target type.

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

**Figure 42**      **Adding Primary SAN Boot Target for Fabric A**



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



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

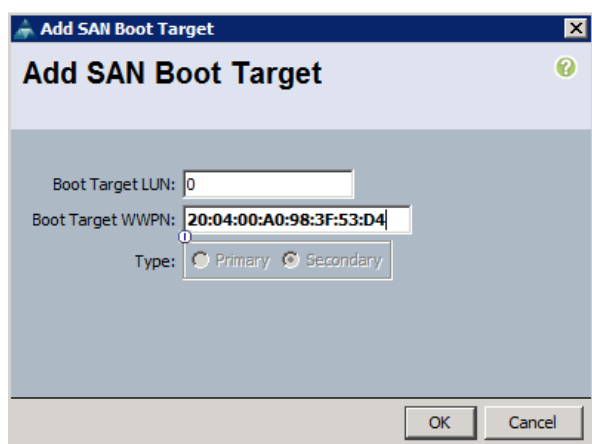
66. Enter the WWPN for fcp\_lif02a.



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

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

**Figure 43** Adding Secondary SAN Boot Target for Fabric A



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

## Create Service Profile Templates

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

To create service profile templates, follow these steps:

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

**Figure 44**      *Details for Creating Service Profile Template*

**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 45**      **Creating vNIC Using vNIC Template**

**Create vNIC**

Name:

Use vNIC Template: ☒

+ Create vNIC Template

vNIC Template:

**Adapter Performance Profile**

Adapter Policy:  + Create Ethernet Adapter Policy

OK Cancel

- i. On the Networking page of the wizard, click **Add** to add another vNIC to the template.
- j. In the Create vNIC box, enter vNIC-B as the name of the vNIC.
- k. Check the Use vNIC Template check box.
- l. In the vNIC Template list, choose vNIC\_Template\_B.
- m. In the Adapter Policy list, choose VMWare.
- n. Click **OK** to add the vNIC to the template.
- o. Review the table in the Networking page to make sure that both vNICs were created.
- p. Click **Next**.

Figure 46 LAN Configuration Details

**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

**Networking**

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:  [+ Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity? ☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
vNIC vNIC-A	Derived	derived	
vNIC vNIC-B	Derived	derived	

[Delete](#) [+ Add](#) [Modify](#)

**iSCSI vNICs**

This Initiator Name Assignment will apply to all iSCSI vNICs within this Service Profile.

**Initiator Name**

Initiator Name Assignment:

[+ Create IQN Suffix Pool](#)

**WARNING:** The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

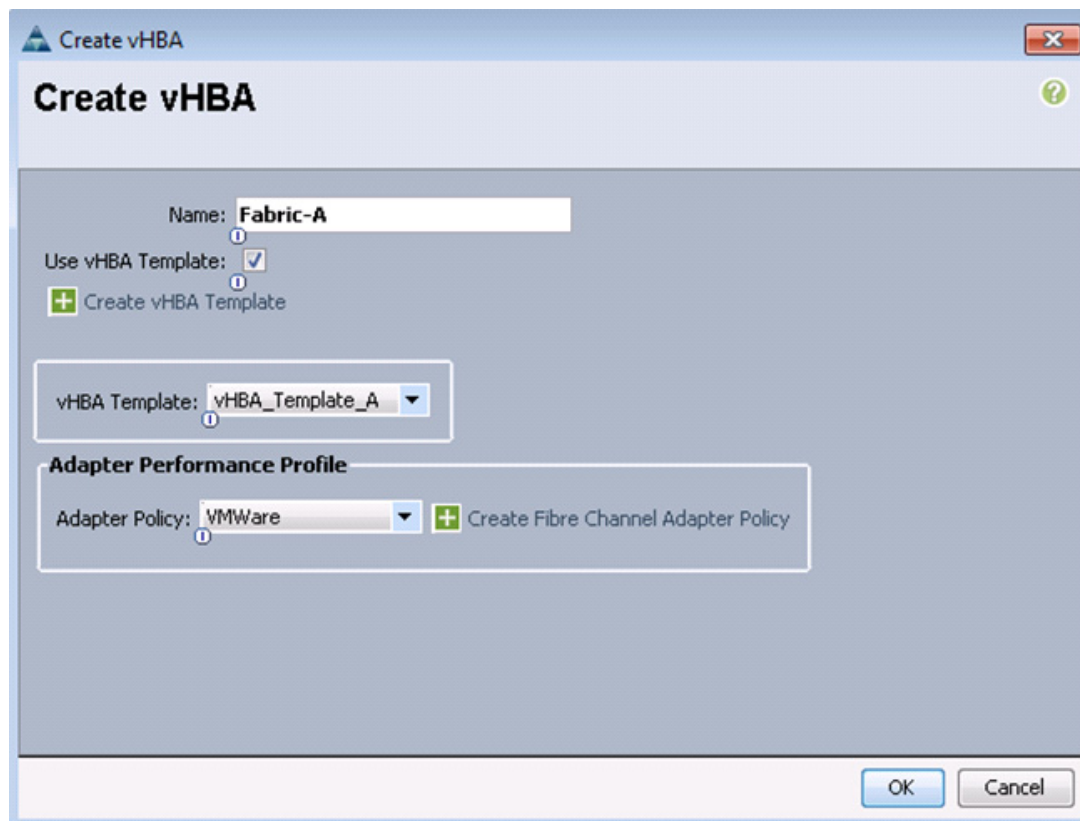
Click **Add** to specify one or more iSCSI vNICs that the server should use.

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
------	-------------------	----------------------	-------------

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

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

**Figure 47**      **Creating vHBA Using vHBA Template**



- j. On the Storage page of the wizard, click **Add** at the bottom of the page to add another vHBA to the template.
- k. In the Create vHBA dialog box, enter Fabric-B as the name of the vHBA.
- l. Check the check box for Use HBA Template.
- m. In the vHBA Template list, choose vHBA\_Template\_B.
- n. In the Adapter Policy list, choose VMware.
- o. Click **OK** to add the vHBA to the template.
- p. Review the table in the Storage page to verify that both vHBAs were created.
- q. Click **Next**.

**Figure 48** Storage Window Showing Created vHBAs

**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**

Protect Configuration: **Yes**  
If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the

**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 49** *Placing vNIC and vHBA on Physical Adapters*



**10. Set the Server Boot Order:**

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

**Figure 50**      **Setting Boot Order for the Service Profile Template**

**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

**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/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**

+ Add - Remove Filter Export Print

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		Fabric-A	Primary		
SAN Target primary			Primary	0	20:01:00:A0:98:3F:53:D4
SAN Target secondary			Secondary	0	20:03:00:A0:98:3F:53:D4
SAN secondary		Fabric-B	Secondary		
SAN Target primary			Primary	0	20:02:00:A0:98:3F:53:D4
SAN Target secondary			Secondary	0	20:04:00:A0:98:3F:53:D4

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



**Figure 51**      **Assigning a Server Pool to the Service Profile Template**

**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

**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:  [+ 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:

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:  [+ Create Host Firmware Package](#)

< Prev   Next >   Finish   Cancel

**13. Add Operational Policies:**


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

**Figure 52**      **Setting Operational Policy**

The screenshot shows the 'Create Service Profile Template' window in the Unified Computing System Manager. The window has a title bar with a close button. Below the title bar is a header area with the text 'Unified Computing System Manager'. On the left side, there is a navigation pane titled 'Create Service Profile Template' with a list of steps: 1. Identify Service Profile Template (checked), 2. Networking (checked), 3. Storage (checked), 4. Zoning (checked), 5. vNIC/vHBA Placement (checked), 6. Server Boot Order (checked), 7. Maintenance Policy (checked), 8. Server Assignment (checked), and 9. Operational Policies (checked). The main area is titled 'Operational Policies' and contains the text 'Optionally specify information that affects how the system operates.' Below this, there are several expandable sections: 'BIOS Configuration' (expanded), 'External IPMI Management Configuration' (collapsed), 'Management IP Address' (collapsed), 'Monitoring Configuration (Thresholds)' (collapsed), 'Power Control Policy Configuration' (expanded), and 'Scrub Policy' (collapsed). The 'BIOS Configuration' section contains the text 'If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile' and a dropdown menu for 'BIOS Policy' with the value 'VM-Host-Infra' and a '+ Create BIOS Policy' button. The 'Power Control Policy Configuration' section contains the text 'Power control policy determines power allocation for a server in a given power group.' and a dropdown menu for 'Power Control Policy' with the value 'No-Power-Cap' and a '+ Create Power Control Policy' button. At the bottom of the window, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

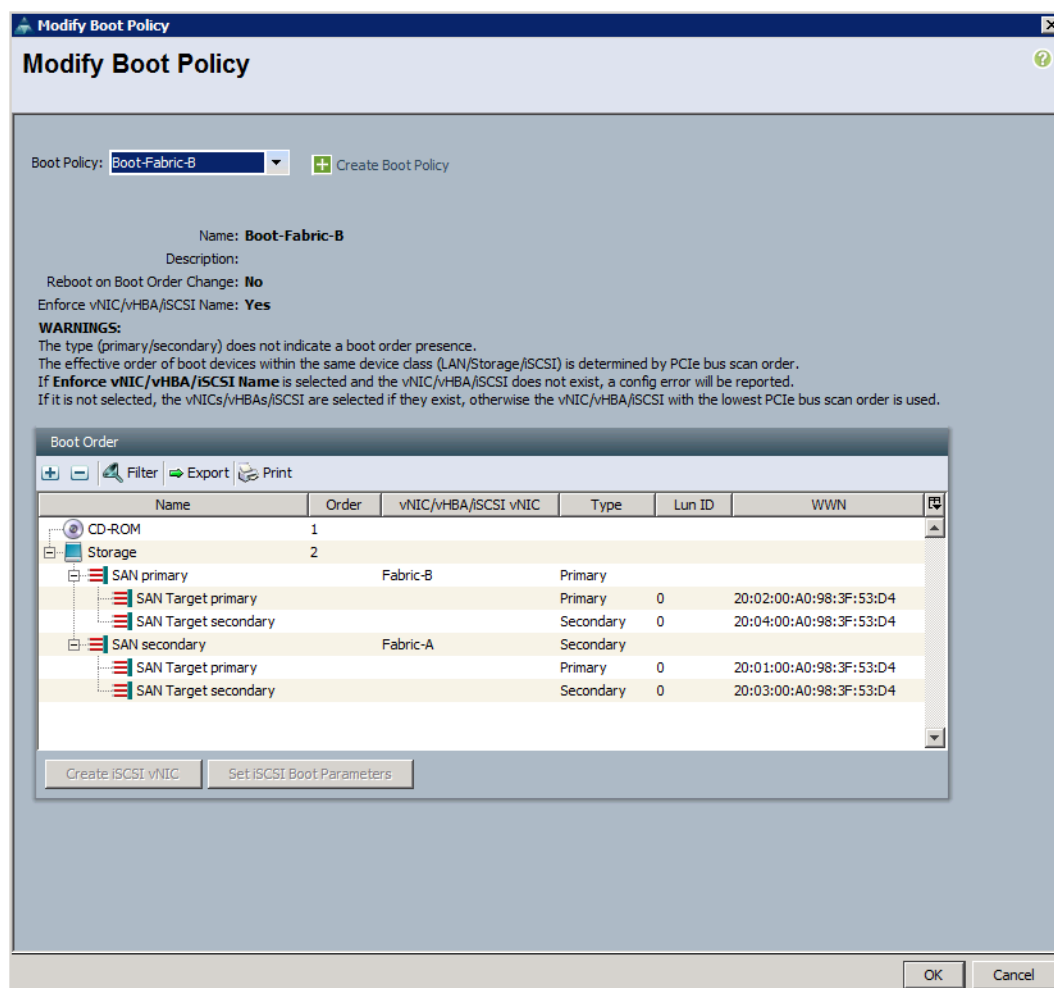
14. Click **Finish** to create the service profile template.
15. Click **OK** in the confirmation message.
16. Click the **Servers** tab in the navigation pane.
17. Choose **Service Profile Templates > root**.
18. Right-click the previously created VM-Host-Infra-Fabric-A template.
19. Choose **Create 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 53**      **Cloning a Service Profile Template**



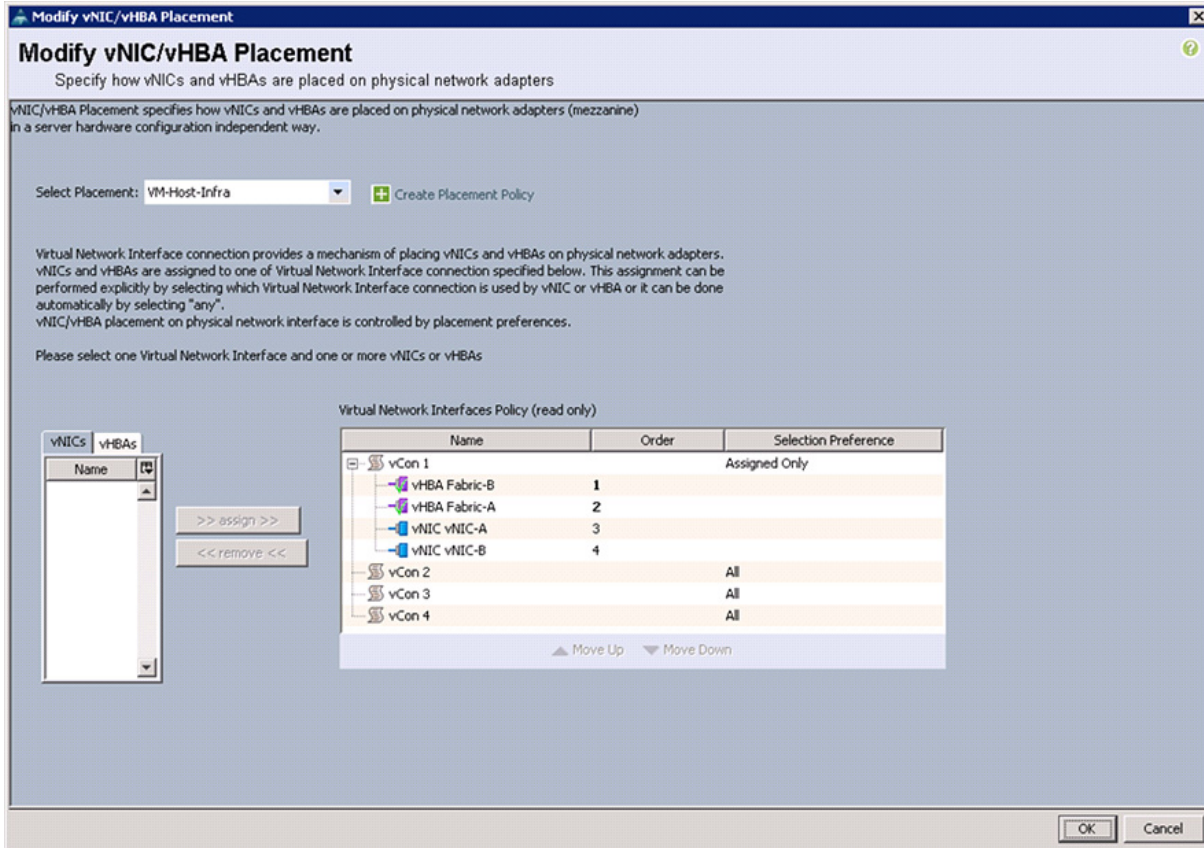
21. Click **OK**.
22. Choose the newly cloned service profile template and click the **Boot Order** tab.
23. Click **Modify Boot Policy**.
24. In the Boot Policy list, choose Boot-Fabric-B.

**Figure 54**      **Modifying Boot Policy**



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

**Figure 55**      **Modifying Placement Policy**



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

## Create Service Profiles

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

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

**Figure 56** *Creating Service Profile from a Service Profile Template*

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

**Figure 57** *Creating Service Profile from a Service Profile Template*

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

## Add More Servers to FlexPod Unit

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

## Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into [Table 28](#) and [Table 29](#).

**Table 28** *FCP LIFs for FC WWPNS*

FCP LIFS	FC WWPNS
fcp_lif01a	
fcp_lif01b	
fcp_lif02a	
fcp_lif02b	



### Note

To gather the FC WWPNS, log in to the storage cluster and run the **network interface show** command.

**Table 29** *vHBA WWPNS for Fabric A and Fabric B*

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



### Note

To gather the vHBA WWPNS 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 29](#), record the WWPNS information that is displayed in the right pane for both the Fabric A vHBA and the Fabric B vHBA for each service profile.

# Storage Networking

## FlexPod Cisco Nexus Base

**Table 30** *Flexpod Cisco Nexus Base Prerequisite*

Description
The Cisco Nexus switch must be running Cisco Nexus NX-OS 6.0(2)N2(2) or later

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

## Set Up Initial Configuration

### Cisco Nexus A

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

1. Configure the switch.


**Note**

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

```

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]:
yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for the "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <<var_nexus_A_hostname>>
Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_A_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_A_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <<var_nexus_A_mgmt0_gw>>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa): rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Enter basic FC configurations (yes/no) [n]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

```

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

```

### Cisco Nexus B

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

1. Configure the switch.


**Note**

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

```

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter

```

```

Enter the switch name: <<var_nexus_B_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_B_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_B_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <<var_nexus_B_mgmt0_gw>>
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa): rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Enter basic FC configurations (yes/no) [n]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

```

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

```

## FlexPod Cisco Nexus FCoE Storage vSphere on Clustered Data ONTAP

### Enable Licenses

#### Cisco Nexus A

To license the Cisco Nexus A switch on <<var\_nexus\_A\_hostname>>, follow these steps:

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

```

config t
feature fcoe
feature npiv
feature lacp
feature vpc

```

#### Cisco Nexus B

To license the Cisco Nexus B switch on <<var\_nexus\_B\_hostname>>, follow these steps:

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

```

config t
feature fcoe
feature npiv
feature lacp
feature vpc

```

### Set Global Configurations

#### Cisco Nexus 5548 A and Cisco Nexus 5548 B

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

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

```

spanning-tree port type network default

```



```

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

```

## Create VLANs

### Cisco Nexus 5548 A and Cisco Nexus 5548 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
exit
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
exit
vlan <<var_pkt-ctrl_vlan_id>>
name Packet-Control-VLAN
exit
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
exit
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
exit

```

## Add Individual Port Descriptions for Troubleshooting

### Cisco Nexus 5548 A

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

From the global configuration mode, run the following commands:

```

interface Eth1/1
description <<var_node01>>:e3a
exit
interface Eth1/2
description <<var_node02>>:e3a
exit
interface Eth1/11
description <<var_ucs_clustername>>-A:1/19
exit

```

```

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

```

## Cisco Nexus 5548 B

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

From the global configuration mode, run the following commands:

```

interface Eth1/1
description <<var_node01>>:e4a
exit
interface Eth1/2
description <<var_node02>>:e4a
exit
interface Eth1/11
description <<var_ucs_clustername>>-A:1/20
exit
interface Eth1/12
description <<var_ucs_clustername>>-B:1/20
exit
interface Eth1/13
description <<var_nexus_A_hostname>>:1/13
exit
interface Eth1/14
description <<var_nexus_A_hostname>>:1/14
exit
interface eth1/31
description <<var_ucs_clustername>>-B:1/31
exit
interface eth1/32
description <<var_ucs_clustername>>-B:1/32
exit

```

## Create Port Channels

### Cisco Nexus 5548 A and Cisco Nexus 5548 B

To create the necessary port channels between devices, follow these steps on both switches:

From the global configuration mode, run the following commands:

```

interface Po10
description vPC peer-link
exit
interface Eth1/13-14
channel-group 10 mode active
no shutdown

```

```

exit
interface Po11
description <<var_node01>>
exit
interface Eth1/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
exit
interface Eth1/2
channel-group 12 mode active
no shutdown
exit
interface Po13
description <<var_ucs_clustername>>-A
exit
interface Eth1/11
channel-group 13 mode active
no shutdown
exit
interface Po14
description <<var_ucs_clustername>>-B
exit
interface Eth1/12
channel-group 14 mode active
no shutdown
exit
copy run start

```

## Configure Port Channels

### Cisco Nexus 5548 A and Cisco Nexus 5548 B

To configure the port channels, follow these steps on both switches:

From the global configuration mode, run the following commands:

```

interface Po10
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_pkt-ctrl_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
spanning-tree port type network
no shutdown
exit
interface Po11
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface Po12
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_nfs_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface Po13

```

```

switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface Po14
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan
<<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>,
<<var_vm-traffic_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
copy run start

```

## Configure Virtual Port Channels

### Cisco Nexus 5548 A

To configure virtual port channels (vPCs) for switch A, follow these steps:

From the global configuration mode, run the following commands:

```

vpc domain <<var_nexus_vpc_domain_id>>
role priority 10
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source
<<var_nexus_A_mgmt0_ip>>
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start

```

### Cisco Nexus 5548 B

To configure vPCs for switch B, follow these steps:

From the global configuration mode, run the following commands.

```

vpc domain <<var_nexus_vpc_domain_id>>
role priority 20
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source
<<var_nexus_B_mgmt0_ip>>
auto-recovery
exit
interface Po10

```

```

vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start

```

## Configure Ports for Cisco Nexus 1110-X Virtual Appliances

### Cisco Nexus 5548 A

To configure the ports in switch A that are connected to the Cisco Nexus 1110-X, follow these steps:

From the global configuration mode, run the following commands:

```

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

```

### Cisco Nexus 5548 B

To configure the ports in switch B that are connected to the Cisco Nexus 1110-X, follow these steps:

From the global configuration mode, run the following commands:

```

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

```

```
no shutdown
exit
copy run start
```

## Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 5548 switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run the **copy run start** command to save the configuration on each switch after configuration is completed.

## Create VSANs, Assign and Enable Virtual Fibre Channel Ports

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

### Cisco Nexus 5548 A

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

From the global configuration mode, run the following commands:

```
vlan <<var_fabric_a_fcoe_vlan_id>>
name FCoE_Fabric_A
fcoe vsan <<var_vsan_a_id>>
exit
interface po11
switchport trunk allowed vlan add <<var_fabric_a_fcoe_vlan_id>>
exit
interface vfc11
switchport description <<var_node01>>:3a
bind interface Eth1/1
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
exit
interface po12
switchport trunk allowed vlan add <<var_fabric_a_fcoe_vlan_id>>
exit
interface vfc12
switchport description <<var_node02>>:3a
bind interface Eth1/2
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
exit
interface po15
description <<var_ucs_clustername>>-A:FCoE
exit
interface Eth1/31-32
channel-group 15 mode active
exit
interface po15
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_fabric_a_fcoe_vlan_id>>
spanning-tree port type edge trunk
no shutdown
```

```

exit
interface vfc15
switchport description <<var_ucs_clustername>>-A:FCoE
bind interface po15
switchport trunk allowed vsan <<var_vsan_a_id>>
no shutdown
vsan database
vsan <<var_vsan_a_id>> name Fabric_A
vsan <<var_vsan_a_id>> interface vfc11
vsan <<var_vsan_a_id>> interface vfc12
vsan <<var_vsan_a_id>> interface vfc15
exit

```

## Cisco Nexus 5548 B

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

From the global configuration mode, run the following commands:

```

vlan <<var_fabric_b_fcoe_vlan_id>>
name FCoE_Fabric_B
fcoe vsan <<var_vsan_b_id>>
exit
interface po11
switchport trunk allowed vlan add <<var_fabric_b_fcoe_vlan_id>>
exit
interface vfc11
switchport description <<var_node01>>:4a
bind interface Eth1/1
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
exit
interface po12
switchport trunk allowed vlan add <<var_fabric_b_fcoe_vlan_id>>
exit
interface vfc12
switchport description <<var_node02>>:4a
bind interface Eth1/2
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
exit
interface po15
description <<var_ucs_clustername>>-B:FCoE
exit
interface Eth1/31-32
channel-group 15 mode active
exit
interface po15
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_fabric_b_fcoe_vlan_id>>
spanning-tree port type edge trunk
no shutdown
exit
interface vfc15
switchport description <<var_ucs_clustername>>-B:FCoE
bind interface po15
switchport trunk allowed vsan <<var_vsan_b_id>>
no shutdown
vsan database
vsan <<var_vsan_b_id>> name Fabric_B
vsan <<var_vsan_b_id>> interface vfc11
vsan <<var_vsan_b_id>> interface vfc12
vsan <<var_vsan_b_id>> interface vfc15

```

```
exit
```

## Create Device Aliases

### Cisco Nexus 5548 A

To configure device aliases and zones for the primary boot paths of switch A on <<var\_nexus\_A\_hostname>>, follow these steps:

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 pwwn <<var_fcp_lif01a_wwpn>>
device-alias name fcp_lif02a pwwn <<var_fcp_lif02a_wwpn>>
exit
device-alias commit
```

### Cisco Nexus 5548 B

To configure device aliases and zones for the boot paths of switch B on <<var\_nexus\_B\_hostname>>, follow these steps:

From the global configuration mode, run the following commands:

```
device-alias database
device-alias name VM-Host-Infra-01_B pwwn <<var_vm_host_infra_01_B_wwpn>>
device-alias name VM-Host-Infra-02_B pwwn <<var_vm_host_infra_02_B_wwpn>>
device-alias name fcp_lif01b pwwn <<var_fcp_lif01b_wwpn>>
device-alias name fcp_lif02b pwwn <<var_fcp_lif02b_wwpn>>
exit
device-alias commit
```

## Create Zones

### Cisco Nexus 5548 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 5548 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

## 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 5.1Update1 Setup

## FlexPod VMware ESXi 5.1Update1 FCoE on Clustered Data ONTAP

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



### Note

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

## Log in to Cisco UCS 6200 Fabric Interconnect

### Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, follow these steps:

1. Open a Web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the **Launch UCS Manager** link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click **Login** to log in to the Cisco UCS Manager.
6. From the main menu, click the **Servers** tab.
7. Choose **Servers > Service Profiles > root > VM-Host-Infra-01**.
8. Right-click VM-Host-Infra-01 and choose KVM Console.
9. If prompted to accept an Unencrypted KVM Session, accept as necessary.
10. Choose **Servers > Service Profiles > root > VM-Host-Infra-02**.
11. Right-click VM-Host-Infra-02 and choose KVM Console Actions > KVM Console.
12. If prompted to accept an Unencrypted KVM Session, accept as necessary.

## Set Up VMware ESXi Installation

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

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

1. In the KVM window, click the **Virtual Media** tab.
2. 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:
  - [VMware ESXi 5.x Driver for Cisco enic](#)
  - [VMware ESXi 5.x Driver for Cisco fnic](#)
2. Log in and download the eNIC and fNIC drivers.
3. Extract the vib files from the downloaded zip files:
  - a. Navigate to enic\_driver\_2.1.2.38-1023014.zip >  
 enic\_driver\_2.1.2.38-offline\_bundle-1023014.zip > vib20 > net-enic.  
 Network: Cisco\_bootbank\_net-enic\_2.1.2.38-1OEM.500.0.0.472560.vib
  - b. Navigate to fnic\_driver\_1.5.0.45-1233300.zip >  
 fnic\_driver\_1.5.0.45-offline\_bundle-1233300.zip > vib20 > scsi-fnic.  
 Storage: Cisco\_bootbank\_scsi-fnic\_1.5.0.45-1OEM.500.0.0.472560.vib
4. Document the saved location.

## Load Updated Cisco VIC enic and fnic Drivers

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

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

1. From each vSphere Client, choose the host in the inventory.
2. 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 net-enic-2.1.2.38-1OEM.500.0.0.472560.x86\_64.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 scsi-fnic-1.5.0.20-1OEM.500.0.0.472560.x86\_64.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 -v
/vmfs/volumes/datastore1/Cisco_bootbank_net-enic_2.1.2.38-1OEM.500.0.0.472560.vib
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -v
/vmfs/volumes/datastore1/Cisco_bootbank_net-enic_2.1.2.38-1OEM.500.0.0.472560.vib
```
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 -v
/vmfs/volumes/datastore1/Cisco_bootbank_scsi-fnic_1.5.0.45-1OEM.500.0.0.472560.vib
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
install --no-sig-check -v
/vmfs/volumes/datastore1/Cisco_bootbank_scsi-fnic_1.5.0.45-1OEM.500.0.0.472560.vib
```
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

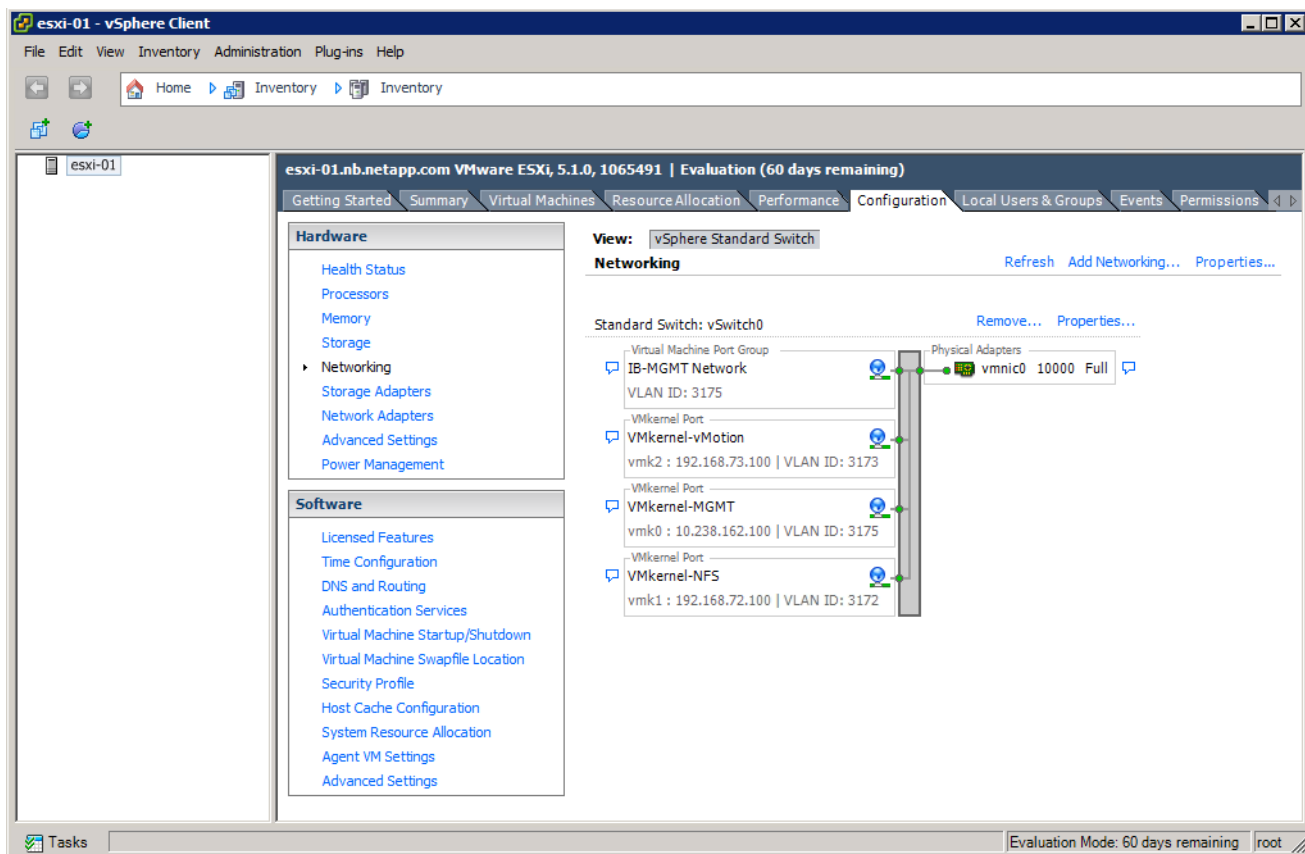
To set up the VMkernel ports and the virtual switches on the VM-Host-Infra-01 ESXi host, follow these steps:

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 58](#).

**Figure 58** vSphere Client Showing VMKernel Ports and Virtual Switch



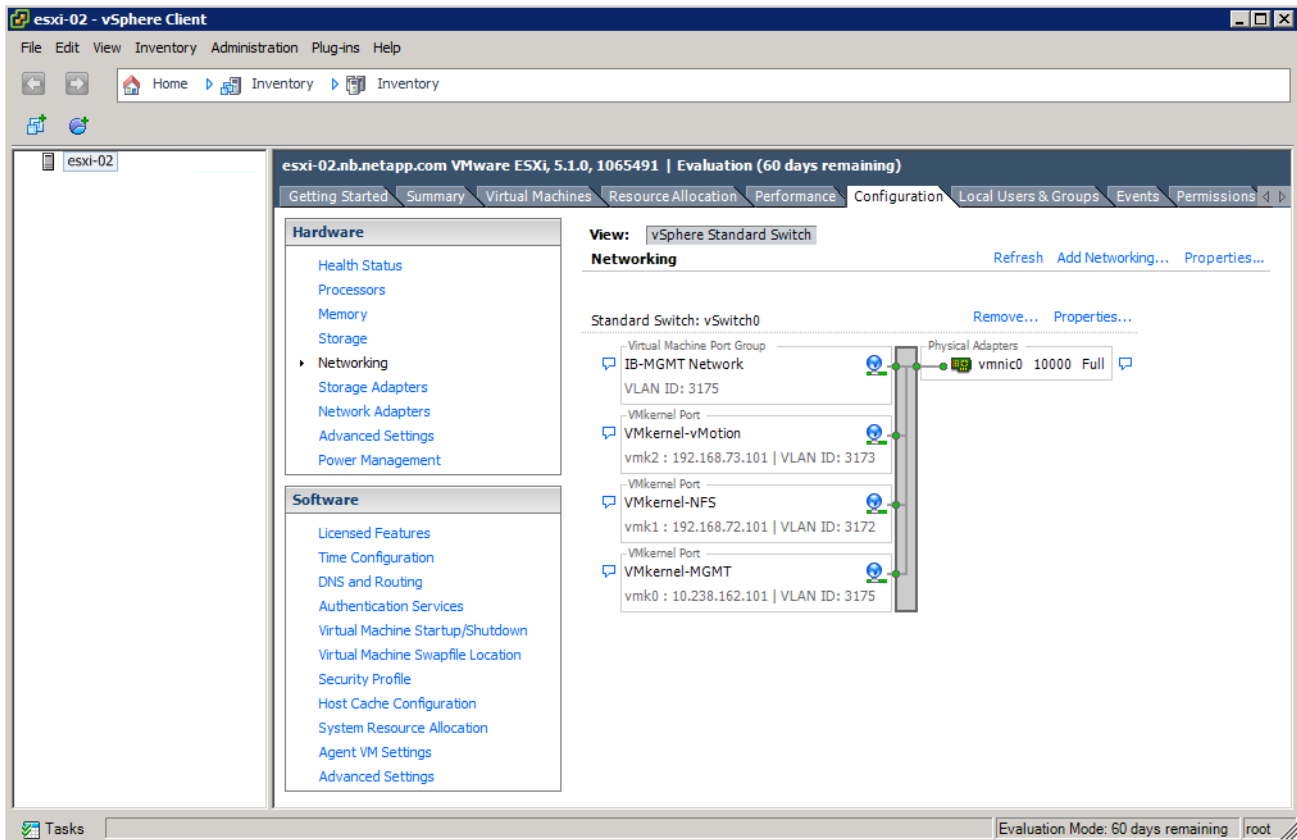
## ESXi Host VM-Host-Infra-02

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

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Configuration** tab.
3. Click **Networking** in the Hardware pane.
4. Click **Properties** on the right side of vSwitch0.
5. Choose the **vSwitch configuration** and click **Edit**.
6. From the General tab, change the MTU to 9000.

7. Click **OK** to close the properties for vSwitch0.
8. Choose the **Management Network configuration** and click **Edit**.
9. Change the network label to **VMkernel-MGMT** and check the **Management Traffic** check box.
10. Click **OK** to finalize the edits for Management Network.
11. Choose the **VM Network configuration** and click **Edit**.
12. Change the network label to IB-MGMT Network and enter <<var\_ib-mgmt\_vlan\_id>> in the VLAN ID (Optional) field.
13. Click **OK** to finalize the edits for VM Network.
14. Click **Add** to add a network element.
15. Choose **VMkernel** and click **Next**.
16. Change the network label to VMkernel-NFS and enter <<var\_nfs\_vlan\_id>> in the VLAN ID (Optional) field.
17. Click **Next** to continue with the NFS VMkernel creation.
18. Enter the IP address <<var\_nfs\_vlan\_id\_ip\_host-02>> and the subnet mask <<var\_nfs\_vlan\_id\_mask\_host02>> for the NFS VLAN interface for VM-Host-Infra-02.
19. Click **Next** to continue with the NFS VMkernel creation.
20. Click **Finish** to finalize the creation of the NFS VMkernel interface.
21. Choose the **VMkernel-NFS configuration** and click **Edit**.
22. Change the MTU to 9000.
23. Click **OK** to finalize the edits for the VMkernel-NFS network.
24. Click **Add** to add a network element.
25. Choose **VMkernel** and click **Next**.
26. Change the network label to VMkernel-vMotion and enter <<var\_vmotion\_vlan\_id>> in the VLAN ID (Optional) field.
27. Check the Use This Port Group for vMotion check box.
28. Click **Next** to continue with the vMotion VMkernel creation.
29. Enter the IP address <<var\_vmotion\_vlan\_id\_ip\_host-02>> and the subnet mask <<var\_vmotion\_vlan\_id\_mask\_host-02>> for the vMotion VLAN interface for VM-Host-Infra-02.
30. Click **Next** to continue with the vMotion VMkernel creation.
31. Click **Finish** to finalize the creation of the vMotion VMkernel interface.
32. Choose the **VMkernel-vMotion configuration** and click **Edit**.
33. Change the MTU to 9000.
34. Click **OK** to finalize the edits for the VMkernel-vMotion network.
35. Close the dialog box to finalize the ESXi host networking setup. The networking for the ESXi host should be similar to [Figure 59](#).

**Figure 59 vSphere Client Showing VMKernel Ports and Virtual Switch**



## Mount Required Datastores

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

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

1. From each vSphere Client, choose the host in the inventory.
2. Click the **Configuration** tab to enable configurations.
3. Click **Storage** in the Hardware pane.
4. From the Datastore area, click **Add Storage** to open the Add Storage wizard.
5. 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 (ntpd) 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 5.1Update1

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 179 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.

---

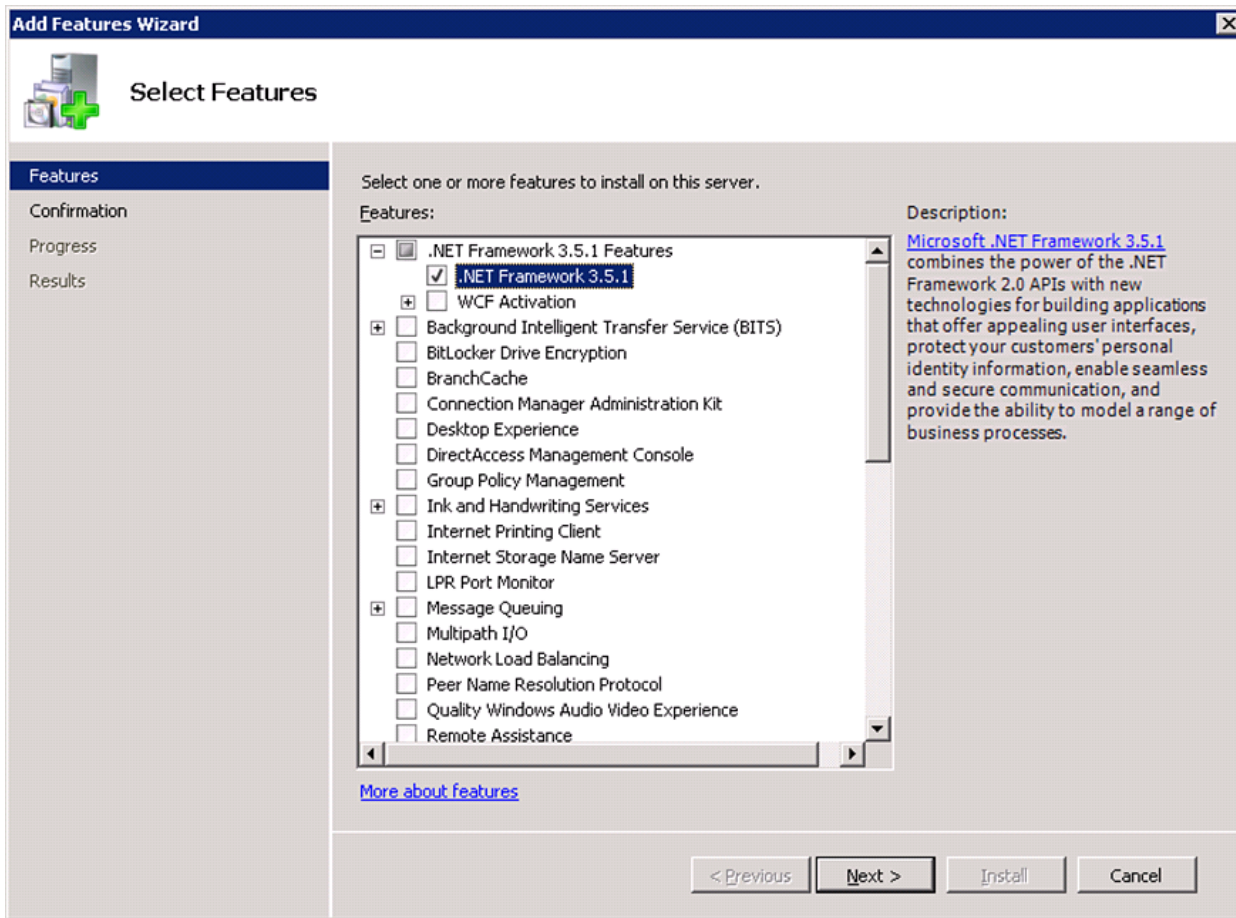
## Install Microsoft SQL Server 2008 R2

### 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 60**      **Adding Features for Installing SQL Server**



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



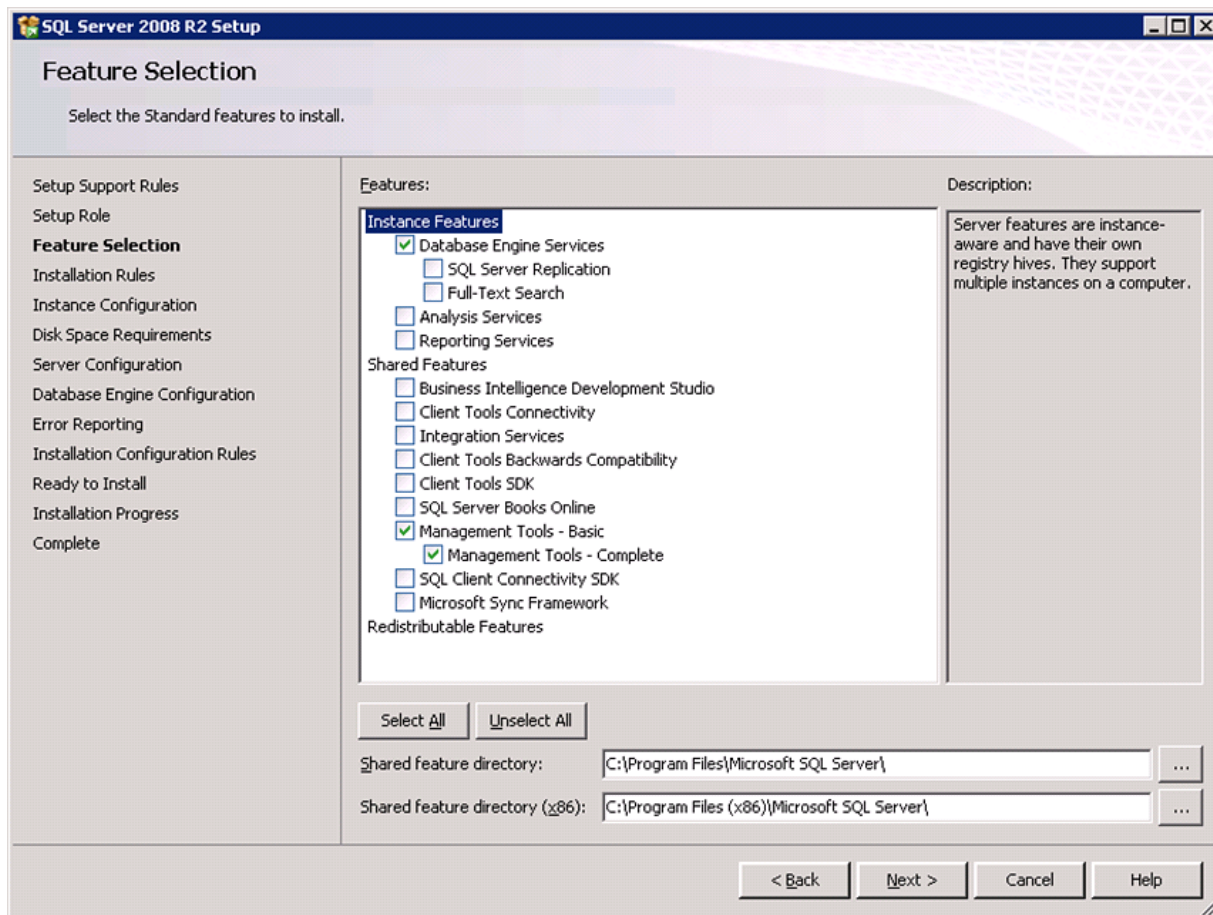
19. Choose New Installation or Add Features to an Existing Installation.
20. Click **OK**.
21. Choose Enter the Product Key. Enter a product key and click **Next**.
22. Read and accept the license terms and choose whether to check the second check box. Click **Next**.
23. Click **Install** to install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



**Note** The Windows firewall issue was addressed in Step 13.

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

**Figure 61**      *Selecting SQL Server Features*



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

**Figure 62**      **SQL Server Instance Configuration**

**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.

☒ **Default instance**  
☐ **Named instance:**

**Instance ID:**   
**Instance root directory:**  ...

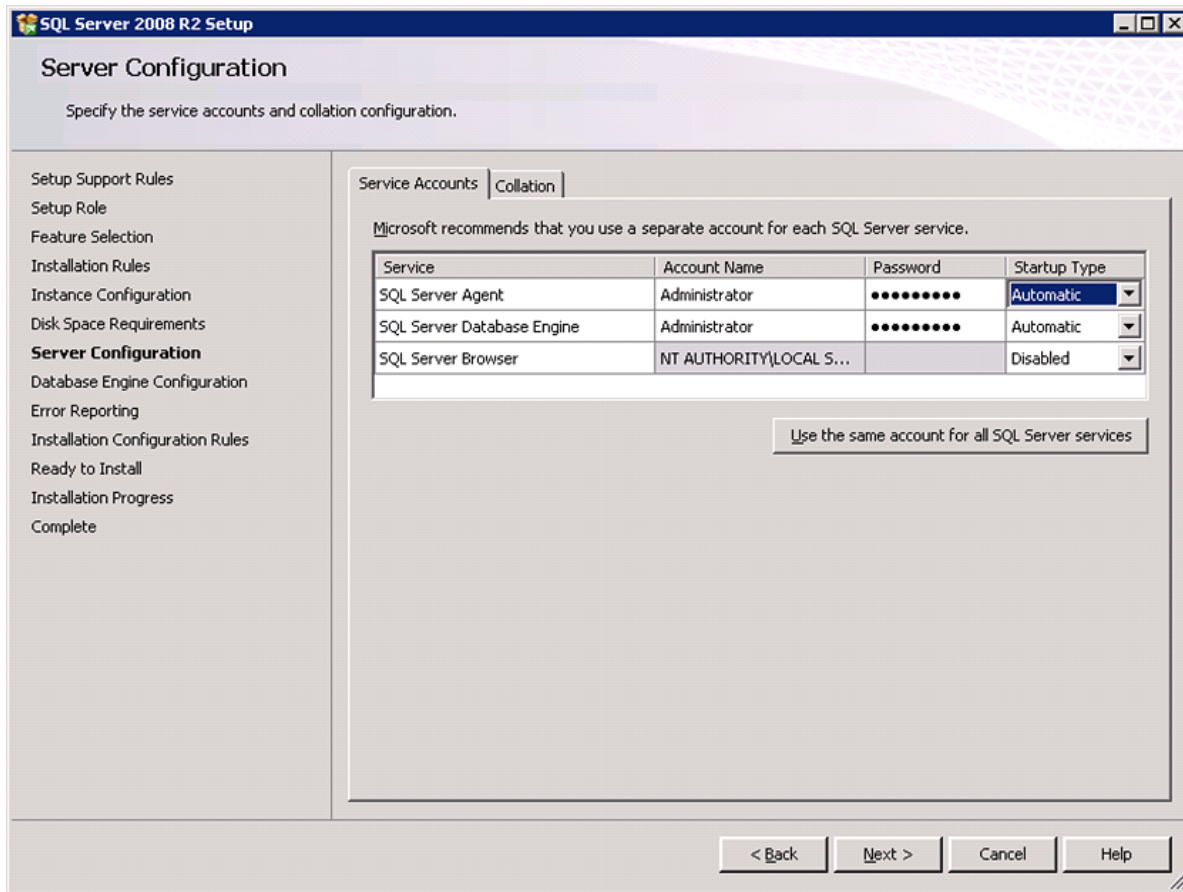
**SQL Server directory:**

**Installed instances:**

Instance Name	Instance ID	Features	Edition	Version

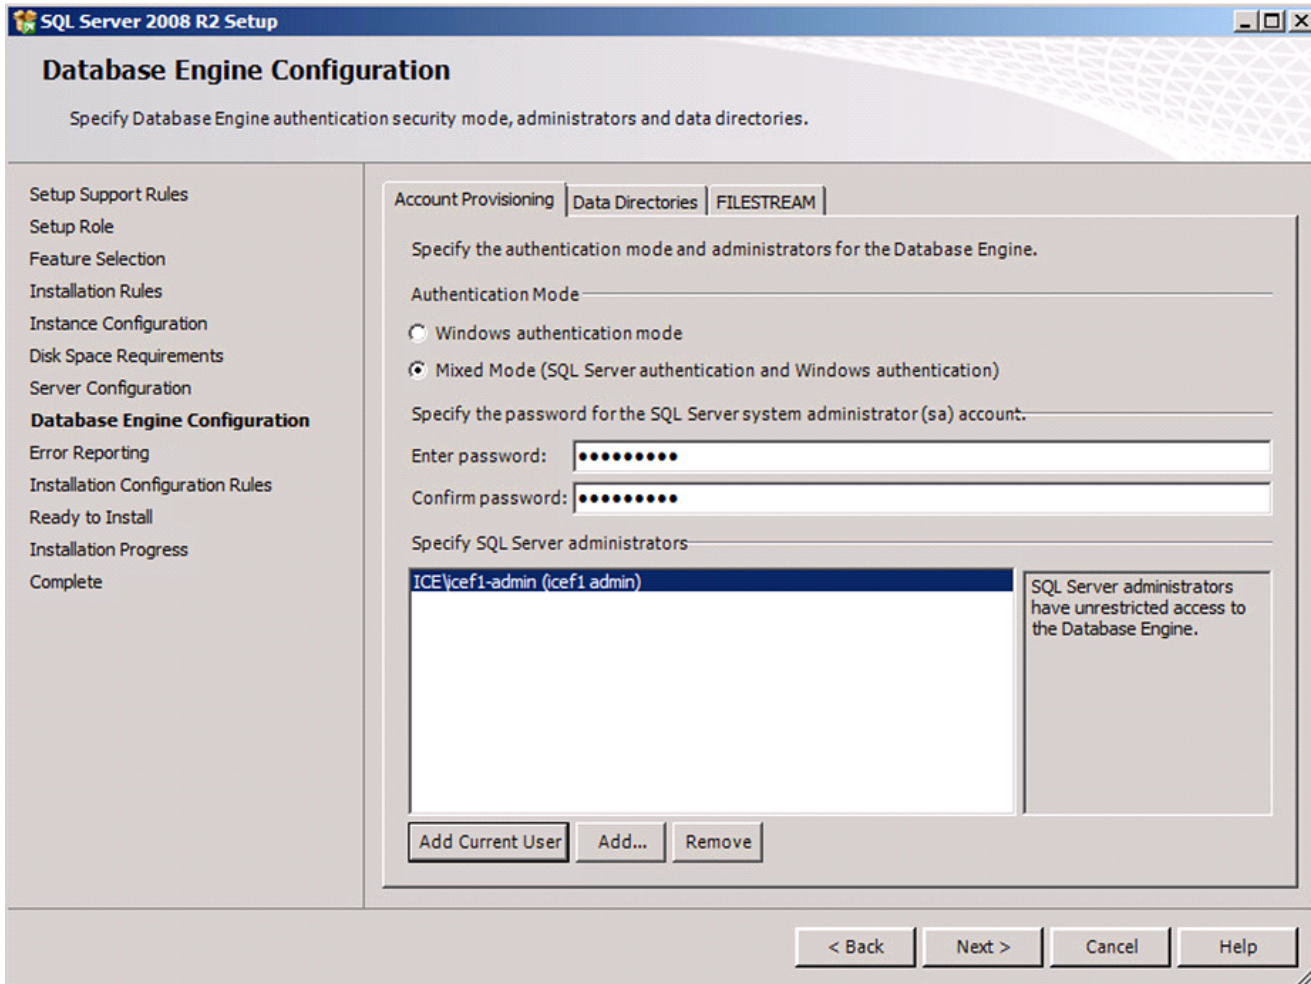
30. Click **Next** for Disk Space Requirements.
31. For the SQL Server Agent service, choose the first cell in the Account Name column and then click **<<Browse...>>**.
32. Enter the local machine administrator name (for example, systemname\Administrator), click **Check Names**, and click **OK**.
33.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

**Figure 63**      **Configuring SQL Server**



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

**Figure 64**      **Configuring Database Engine**



37. Choose whether to send error reports to Microsoft. Click **Next**.
38. Click **Next**.
39. Click **Install**.
40. After the installation is complete, click **Close** to close the SQL Server installer.
41. Close the SQL Server Installation Center.
42. Install all available Microsoft Windows updates by navigating to **Start > All Programs > Windows Update**.
43. 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 **vpxuser** password for **<Password>**:

```
use [master]
go
CREATE DATABASE [VCDB] ON PRIMARY
```

```

(NAME = N'vcdb', FILENAME = N'C:\VCDB.mdf', SIZE = 2000KB, FILEGROWTH = 10% )
LOG ON
(NAME = N'vcdb_log', FILENAME = N'C:\VCDB.ldf', SIZE = 1000KB, FILEGROWTH = 10%)
COLLATE SQL_Latin1_General_CP1_CI_AS
go
use VCDB
go
sp_addlogin @loginame=[vpxuser], @passwd=N'<Password>', @defdb='VCDB',
@deflanguage='us_english'
go
ALTER LOGIN [vpxuser] WITH CHECK_POLICY = OFF
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
use MSDB
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
use VCDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use MSDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go

```

**Note**

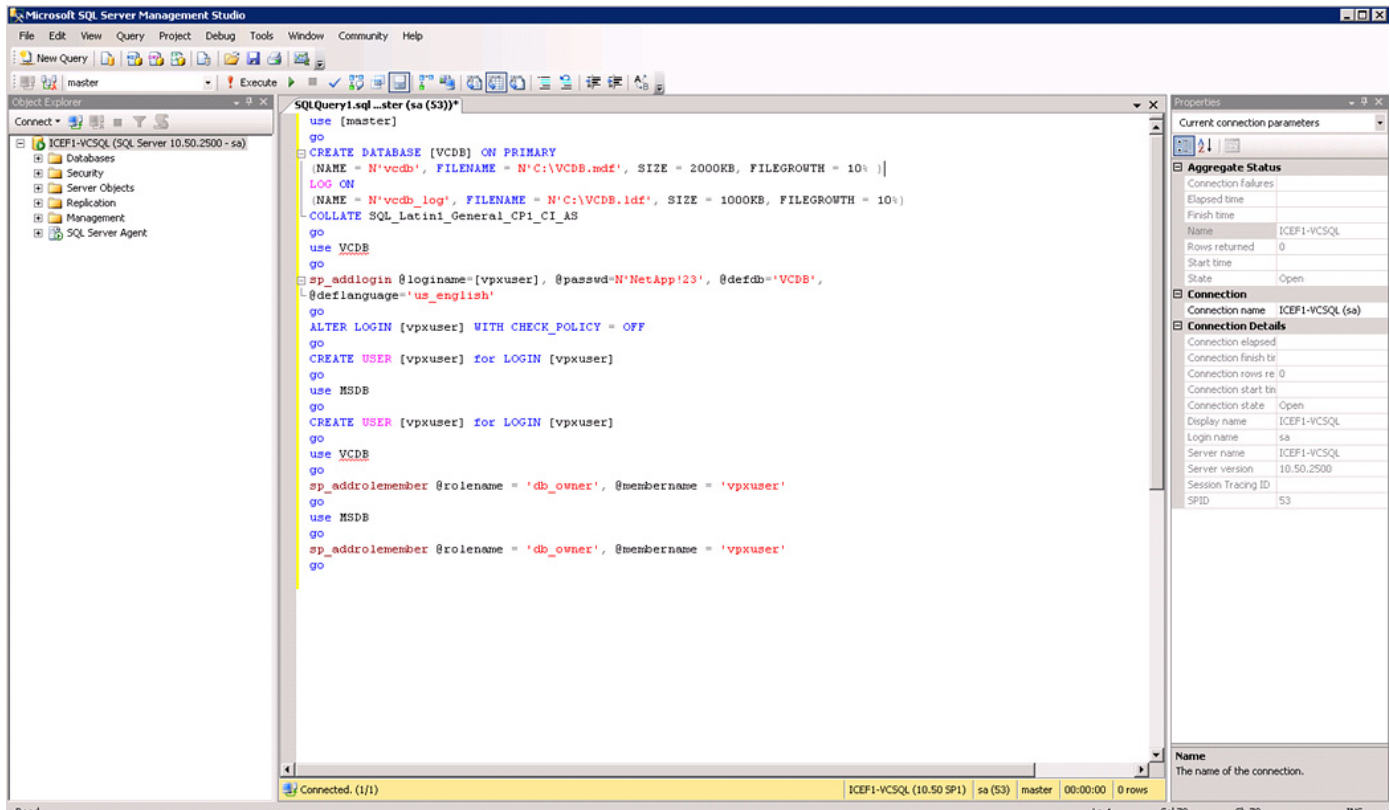

---

This example illustrates the script.

---



Figure 65 Executing the Query



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

## Build and Set Up VMware vCenter VM

### Build VMware vCenter VM

To build the VMware vCenter VM, follow these steps:

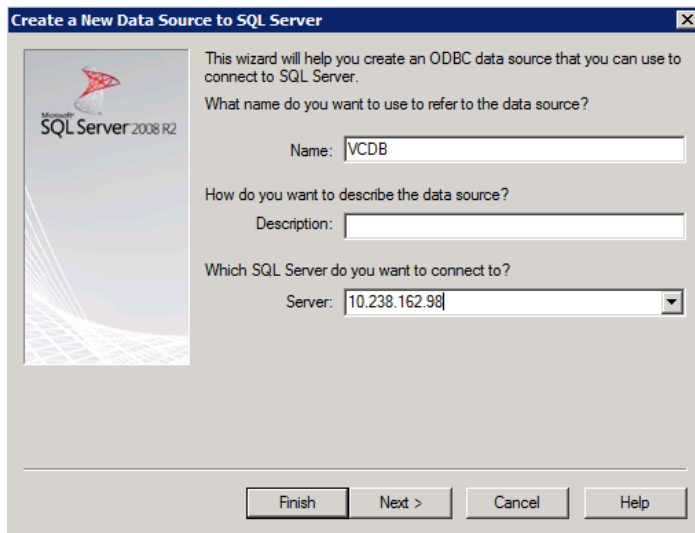
1. Using the instructions for building a SQL Server VM provided in the section “Build Microsoft SQL Server VM,” build a VMware vCenter VM with the following configuration in the <<var\_ib-mgmt\_vlan\_id>> VLAN:
  - 4GB RAM
  - Two CPUs
  - One virtual network interface
2. Start the VM, install VMware 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 66**      **Creating New Data Source**



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

**Figure 67**      *Login Credentials for SQL Server*

**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. Check the **Change the default database to** check box and choose VCDB from the drop down list. Click **Next**.

**Figure 68**      *Setting Up the Database*

**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 69**      *Verifying ODBC Data Source*



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



**Note** A restart might be required.

## Install VMware vCenter Server

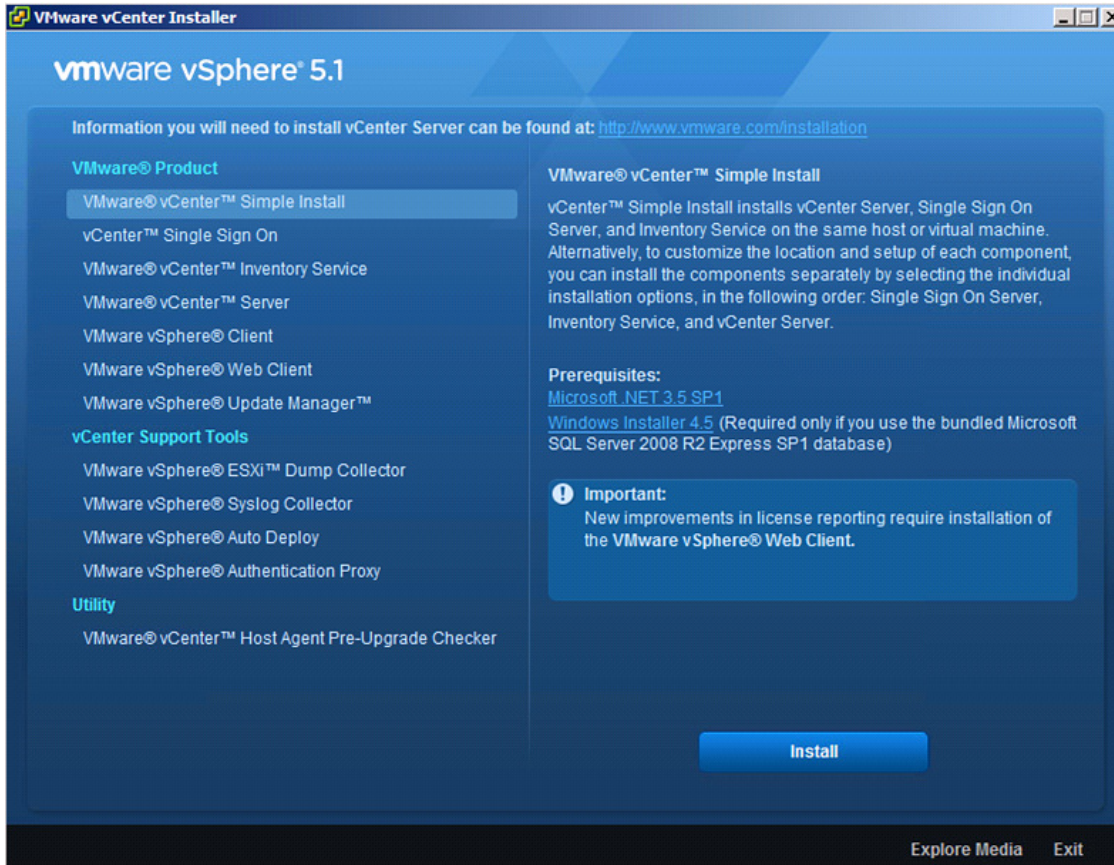
### vCenter Server VM

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

1. In the vCenter Server VMware console, click the ninth button (CD with a wrench) to map the VMware vCenter ISO and choose **Connect to ISO Image on Local Disk**.
2. Navigate to the VMware vCenter 5.1 (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 70**      *Installing VMware vSphere*

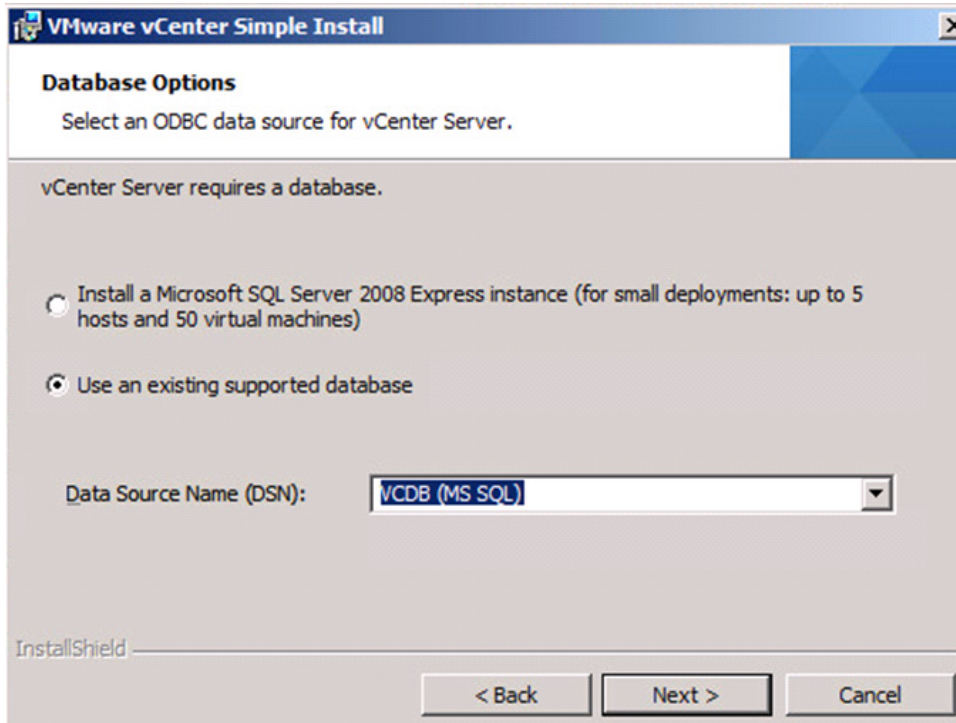


5. Click **Yes** at the User Account Control warning.
6. Click **Next** to install vCenter Single Sign On.
7. Click **Next**.
8. Accept the terms of the license agreement and click **Next**.
9. Enter and confirm <<var\_password>> for admin@System-Domain. Click **Next**.
10. Keep the radio button checked to install a local Microsoft SQL Server 2008 R2 Express instance and click **Next**.

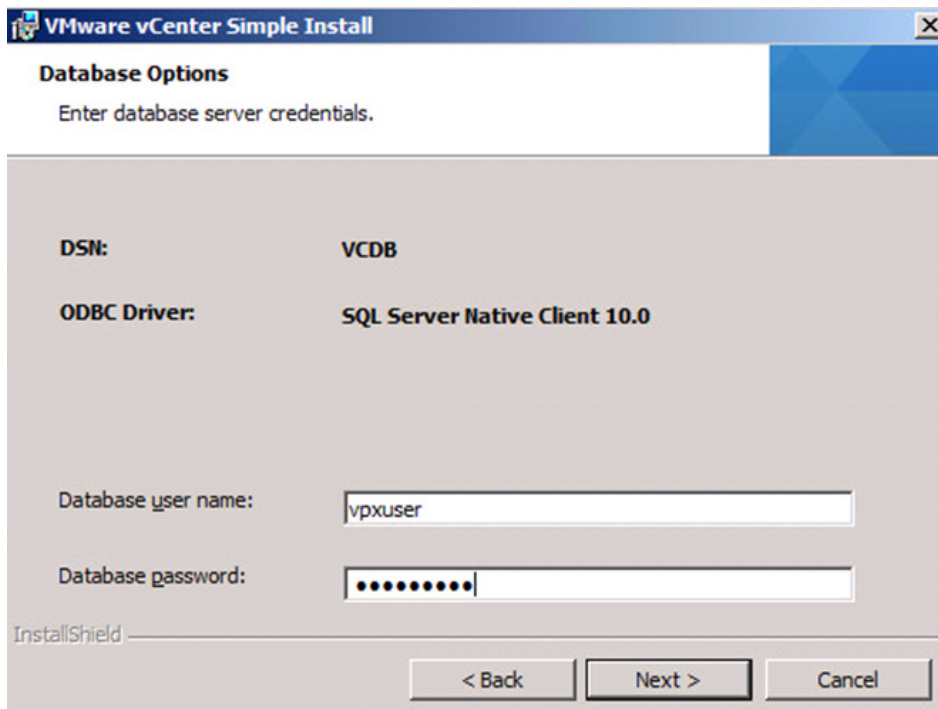
**Figure 71**      **Setting Database User Information**

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 72**      **Selecting Database for VMware vSphere**



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

**Figure 73**      **Entering Database Login 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 as **FlexPod\_Management** and check the check boxes for **Turn On vSphere HA** and **Turn on vSphere DRS**. Click **Next**.

**Figure 74**      **Selecting Cluster Features**

**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. Check the Assign a New License Key to the host check box. Press **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 5.1Update1 Enterprise Plus licensing is installed.

### Configure CIMC Interface on Both Cisco Nexus 1110-Xs

#### Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure the Cisco Integrated Management Controller (CIMC) interface on the Cisco Nexus 1110-X VSAs, follow these steps:

1. Using the supplied dongle, connect a monitor and USB keyboard to the KVM console port on the front of the Cisco Nexus 1110-X virtual appliance.
2. Reboot the virtual appliance.
3. Press **F8** when prompted to configure the CIMC interface.
4. Using the spacebar, set the NIC mode to Dedicated.
5. Uncheck the check box **DHCP enabled**.
6. Set the CIMC IP address (<<var\_cimc\_ip>>) in the out-of-band management VLAN.

7. Set the CIMC subnet mask (<<var\_cimc\_mask>>).
8. Set the CIMC gateway (<<var\_cimc\_gateway>>).
9. Set the NIC redundancy to None.
10. Set and reenter the CIMC default password (<<var\_password>>).
11. Press **F10** to save the configuration.
12. Continue pressing **F5** until Network settings configured is shown.
13. Press **Esc** to reboot the virtual appliance.

## Configure Serial over LAN for Both Cisco Nexus 1110-Xs


### Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure serial over LAN on the Cisco Nexus 1110-X VSAs, follow these steps:


1. Use a Web browser to open the URL at **http://<var\_cimc\_ip>**.
2. Log in to the CIMC with the admin user id and the CIMC default password (<<var\_password>>).
3. In the left column, click **Remote Presence**.
4. Click the **option Serial over LAN** tab.
5. Check the **Enabled** check box for Serial over LAN Properties.
6. From the Baud Rate drop-down menu, choose **9600 bps**.
7. Click **Save Changes**.



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



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

**Figure 76**      *Execute Command for Connecting the Host*

## Configure Cisco Nexus 1110-X Virtual Appliances

### Cisco Nexus 1110-X A

To configure Cisco Nexus 1110-X A, follow these steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.

```

Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_password>>
Enter HA role[primary/secondary]: primary
Enter network-uplink type <1-5>: 1
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>>
Confirm 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. Ports 1-2 carry all management, control and data vlans
  2. Ports 1-2 management and control, ports 3-6 data
  3. Ports 1-2 management, ports 3-6 control and data
  4. Ports 1-2 management, ports 3-4 control, ports 5-6 data
  5. Flexible
Choose 1
Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
Enter the domain id<1-4095>: <<var_1110x_domain_id>>
```



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

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

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.SV2.2.1a.iso) is not present on the Cisco Nexus 1110-X, run the copy command to copy it to the Cisco Nexus 1110-X disk. You must obtain the file from [www.cisco.com](http://www.cisco.com) and place the file either on an FTP server or on a UNIX® or Linux® machine (using scp) that is accessible from the Cisco Nexus 1110-X management interface. An example copy command from an FTP server is copy  
ftp://<<var\_ftp\_server>>/nexus-1000v.4.2.1.SV2.2.1a.iso /repository/..

```

virtual-service-blade-type new nexus-1000v.4.2.1.SV2.1.1a.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 in both the subsections:

### Cisco Nexus 1110-X A

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

### Cisco Nexus 1110-X B

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

```

config t
virtual-service-blade VSM-1
enable secondary
Enter vsb image: [nexus-1000v.4.2.1.SV2.2.1a.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_mgmt_ip>>
Enter Management subnet mask: <<var_vsm_mgmt_mask>>
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>
Enter HostName: <<var_vsm_hostname>>
Enter the password for 'admin': <<var_password>>

```

2. Type show virtual-service-blade summary. Continue periodically entering this command until both the primary and secondary VSM-1s have a state of VSB POWERED ON.

```
copy run start
```

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

## Install Virtual Ethernet Module on Each ESXi Host

### vCenter Server Virtual Machine

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

1. Launch a Web browser to `http://<<var_vsm_mgmt_ip>>`.
2. Right-click the `cross_cisco-vem-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-v162-4.2.1.2.2.1a.0-3.1.1.vib` file and click **Open**.
12. Click **Yes**. The VEM file should now appear in the VEM folder in the datastore.
13. Open the VMware vSphere CLI command prompt.
14. For each ESXi host in the VMware vSphere CLI, run the following command:

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

**Figure 77** Executing Command for Each ESXi Hosts

```
Administrator: Command Prompt
C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 10.238.162.100 -u root -p nbice123! software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VIBs Installed: Cisco_bootbank_cisco-vem-v162-esx_4.2.1.2.2.1a.0-3.1.1
  VIBs Removed:
  VIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 10.238.162.101 -u root -p nbice123! software vib install -v /vmfs/volumes/infra_datastore_1/VEM/cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VIBs Installed: Cisco_bootbank_cisco-vem-v162-esx_4.2.1.2.2.1a.0-3.1.1
  VIBs Removed:
  VIBs 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

```

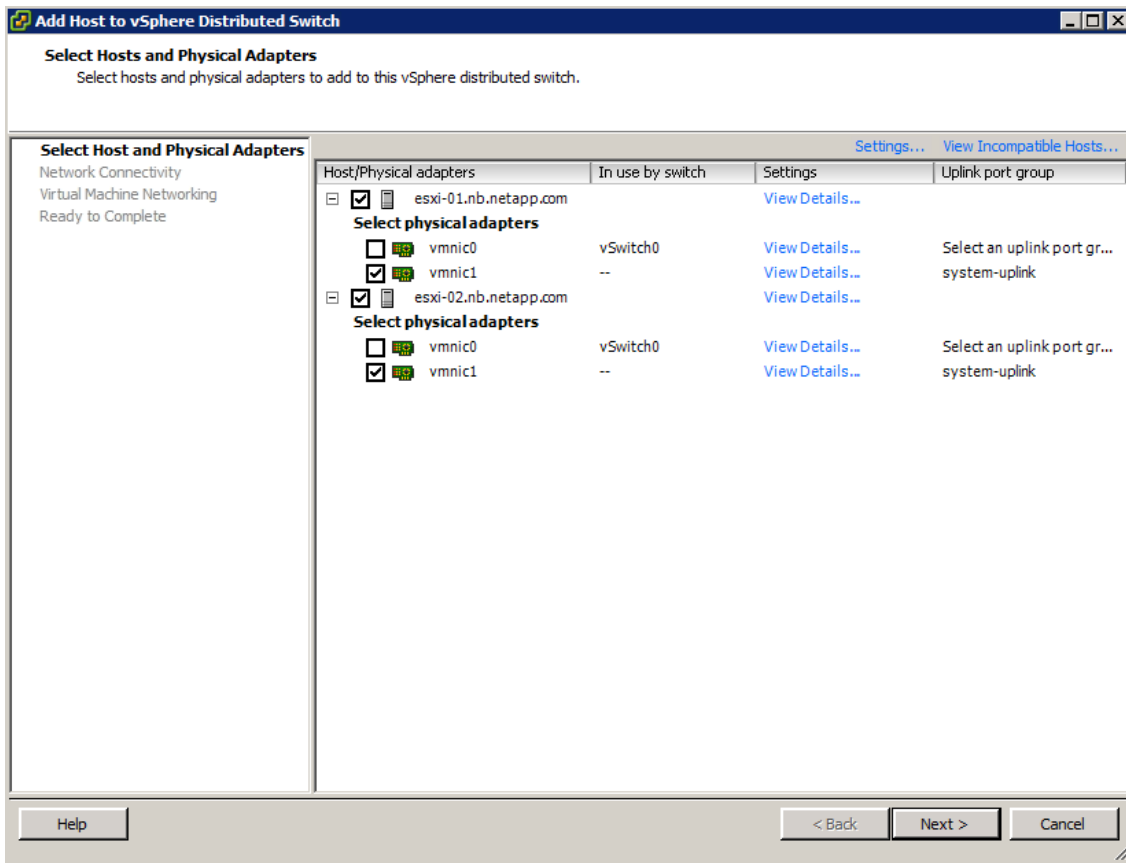
## Migrate Networking Components for ESXi Hosts to Cisco Nexus 1000V

### 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**.

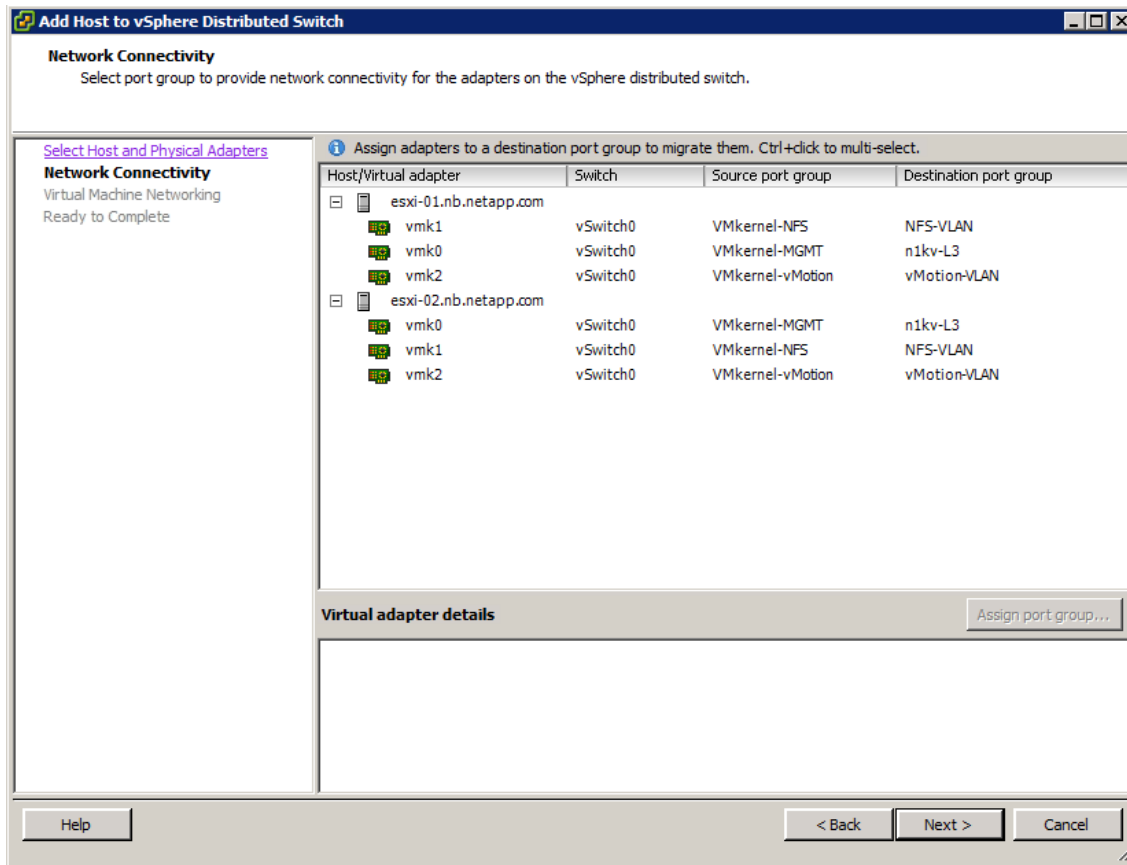
**Figure 78**      *Adding Host to the vSphere Distributed Switch*



- 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**.

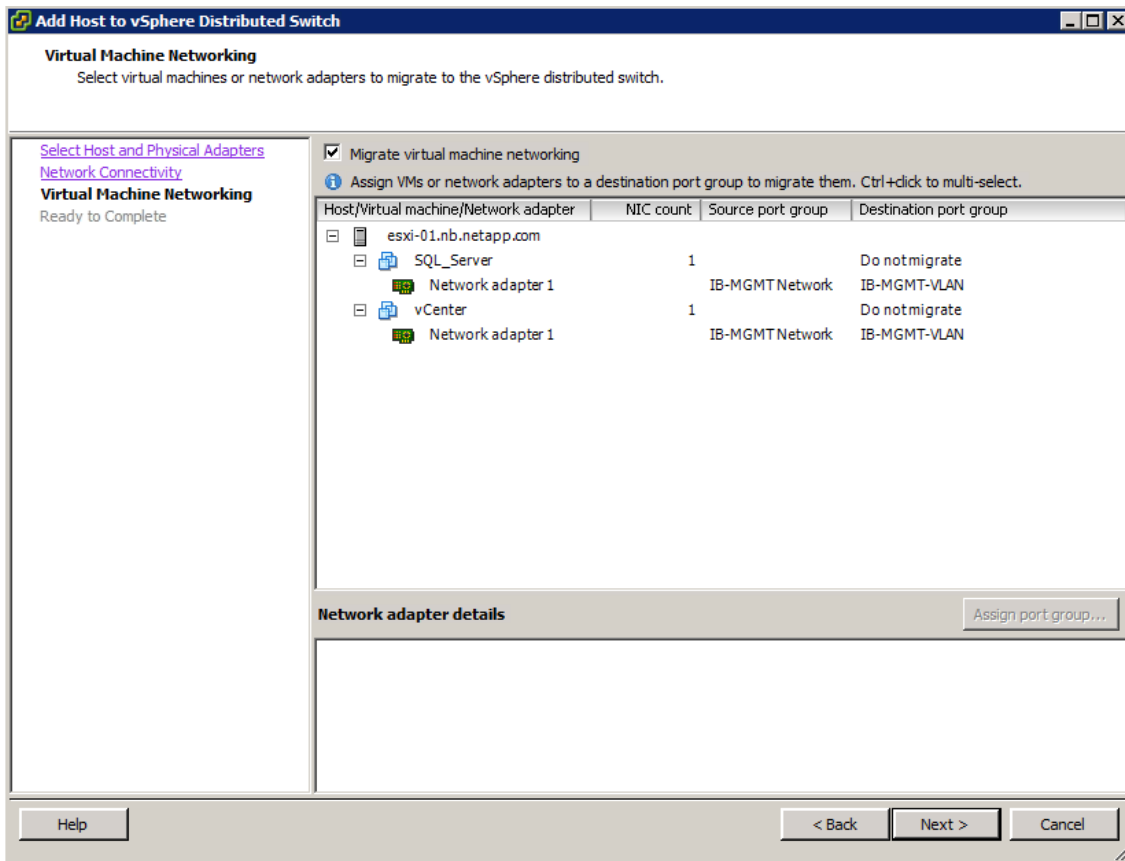


**Figure 79**      **Assigning Port Groups**



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

**Figure 80**      **Choosing Port Groups for Each VMs**



7. Click **Finish**. Wait for the migration process to complete.
8. In the vSphere Client window, choose **Home > Hosts and Clusters**.
9. Select 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 side of vSwitch0, click **Remove**.
11. Click **Yes**.
12. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View.
13. Click **Manage Physical Adapters**.
14. Scroll down to the system-uplink box and click **Add NIC**.
15. Choose vmnic0 and click **OK**.
16. Click **OK** to close the Manage Physical Adapters window. Two system uplinks should now be present.
17. Choose the second ESXi host and Click the **Configuration** tab. In the Hardware field, click **Networking**.
18. Make sure vSphere Standard Switch is selected at the top next to View. vSwitch0 should have no active VMkernel or VM Network ports on it. On the upper right side of vSwitch0, click **Remove**.

19. Click **Yes**.
20. After vSwitch0 has disappeared from the screen, click **vSphere Distributed Switch** at the top next to View.
21. Click **Manage Physical Adapters**.
22. Scroll down to the system-uplink box and click **Add NIC**.
23. Choose vmnic0 and click **OK**.
24. Click **OK** to close the Manage Physical Adapters window. Two system-uplinks should now be present.
25. From the SSH client that is connected to the Cisco Nexus 1000V, run **show interface status** to verify that all interfaces and port channels have been correctly configured.

**Figure 81**      *Verifying Interfaces and Port Channels*

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

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

Figure 82 Verifying the ESXi Hosts

```

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

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

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

```

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

## FlexPod Management Tool Setup

### NetApp Virtual Storage Console (VSC) 4.2.1 Deployment Procedure

#### VSC 4.2.1 Preinstallation Considerations

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

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

#### Install VSC 4.2.1

To install the VSC 4.1 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: [Virtual Storage Console 4.2.1](#) from the [NetApp Support site](#).
4. Right-click the file downloaded in step 3 and choose Run As Administrator.
5. On the Installation wizard Welcome page, click **Next**.
6. Check the check box to accept the installation license agreement.
7. Check the backup and recovery check box. Click **Next**.




**Note** The backup and recovery capability requires an additional license.


**Figure 83**      *Selecting the Desired Capabilities of VSC 4.1*



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

**Figure 84**      **Installation Location**

9. Click **Install**.

**Figure 85**      **Ready to Install VSC 4.1**

10. Click **Finish**.

## Register VSC with vCenter Server

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

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

3. In the Plug-in Service Information section, choose the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.
4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user), and user password for the vCenter Server. Click **Register** to complete the registration.

**Figure 86**      **Registering VSC with vCenter Server**

vSphere Plugin Registration - Windows Internet Explorer

https://localhost:8143/ Certificate Error Bing

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: 10.238.162.97

vCenter Server information

Host name or IP Address: 10.238.162.99

Port: 443

User name: NB\Administrator

User password: .....

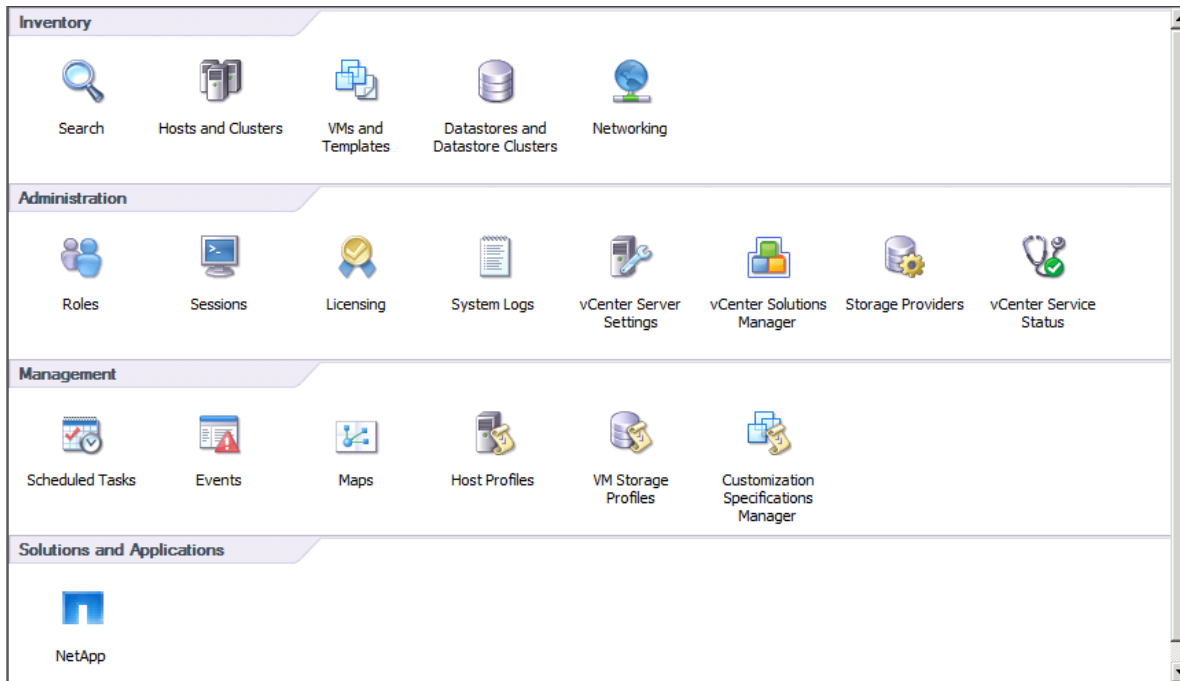
Register

## Discover and Add Storage Resources

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

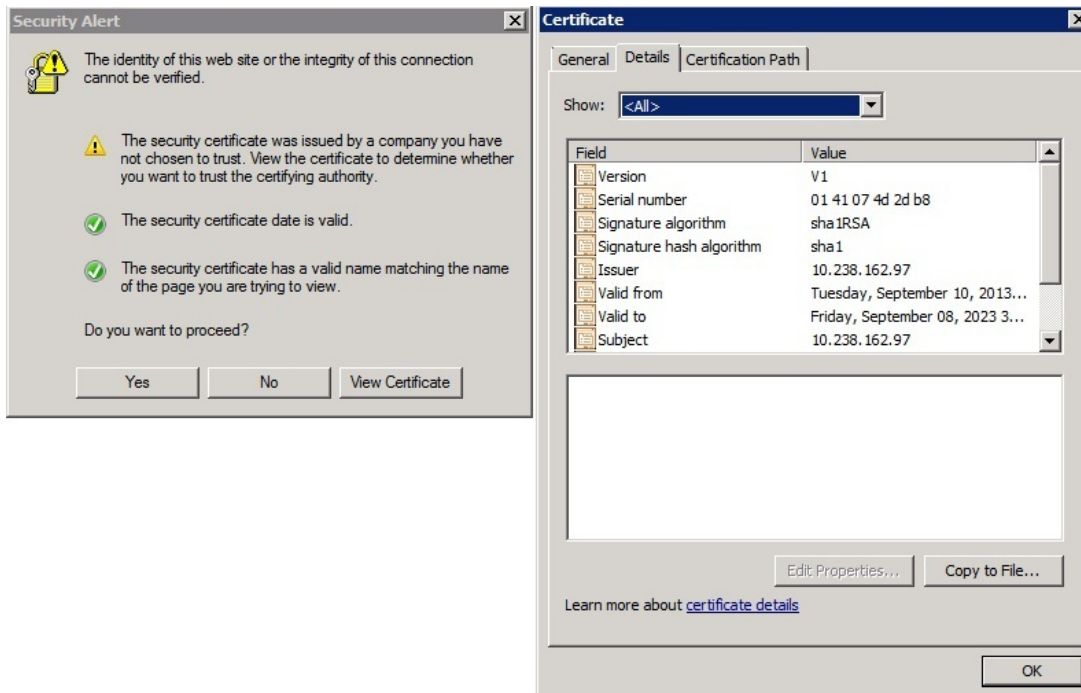
1. Using the vSphere Client, log in to the vCenter Server as FlexPod admin user. If the vSphere Client was previously opened, close it and then reopen it.
2. If an SSL certificate warning from the VSC is displayed, check the checkbox to Install the certificate and then click **Ignore**.
3. Click the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

**Figure 87**      **Adding Storage Resource**



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

**Figure 88**      **Certificate Details**





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

**Figure 89** vSphere Client Showing Storage Controllers

The screenshot shows the vSphere Client interface. On the left is a navigation pane with the following options: Overview, Storage Details - SAN, Storage Details - NAS, Data Collection, Tools, and Discovery Status. The main area is titled 'Monitoring and Host Configuration' and contains 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-	10.238.162.238		Authenti...	0.00B (0%)	Unknown	Unknown
Controller: -unknown- (192.168.72.241)			Unknown	0.00B (0%)	Unknown	Unknown
Controller: -unknown- (192.168.72.242)			Unknown	0.00B (0%)	Unknown	Unknown

**ESX Hosts**

Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
esxi-01.nb.netapp.com		5.1.0	Alert	Alert	Alert	Alert
esxi-02.nb.netapp.com		5.1.0	Alert	Alert	Alert	Alert

Last update: Fri Nov 29 18:52:02 GMT+530 2013

- In the list of storage controllers, right-click the first controller listed and choose **Modify Credentials**.
- 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**.
- Click **OK** to accept the controller privileges.

## Optimal Storage Settings for ESXi Hosts

VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, follow these steps:

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

**Figure 90** Setting Recommended Values for the Hosts

The screenshot shows the 'ESX Hosts' section of the vSphere Client. A context menu is open over the first host, 'esxi-01.nb.netapp.com', with the following options: 'Set Recommended Values...', 'Show Details...', and 'Skip Host...'. The table below shows the hosts and their status.

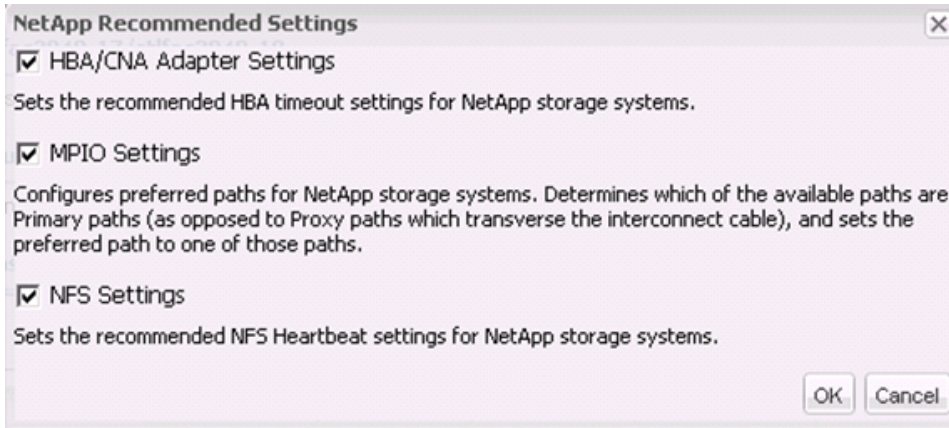
Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
esxi-01.nb.netapp.com		5.1.0			Alert	Alert
esxi-02.nb.netapp.com		5.1.0			Alert	Alert

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



**Note** 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 91** *Recommended Settings for NetApp Storage System*



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

**Figure 92** *ESXi Servers with Changed Network Related Parameter*

ESX Hosts						
Hostname	IP Address	Version	Status	Adapter Settings	MPIO Settings	NFS Settings
esxi-01.nb.netapp.com		5.1.0	Pending Reboot	Normal	Normal	Normal
esxi-02.nb.netapp.com		5.1.0	Pending Reboot	Normal	Normal	Normal

## VSC 4.2.1 Backup and Recovery

### Prerequisites to use Backup and Recovery Capability

You can use the Backup and Recovery capability to schedule backups and restore your datastores, virtual machines, or virtual disk files. Before you begin using this capability you must add storage systems that contain the datastores and virtual machines for which backups need to be created.



**Note** The Backup and Recovery capability does not use the user credentials from the Monitoring and Host Configuration capability.


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

## Backup and Recovery Configuration

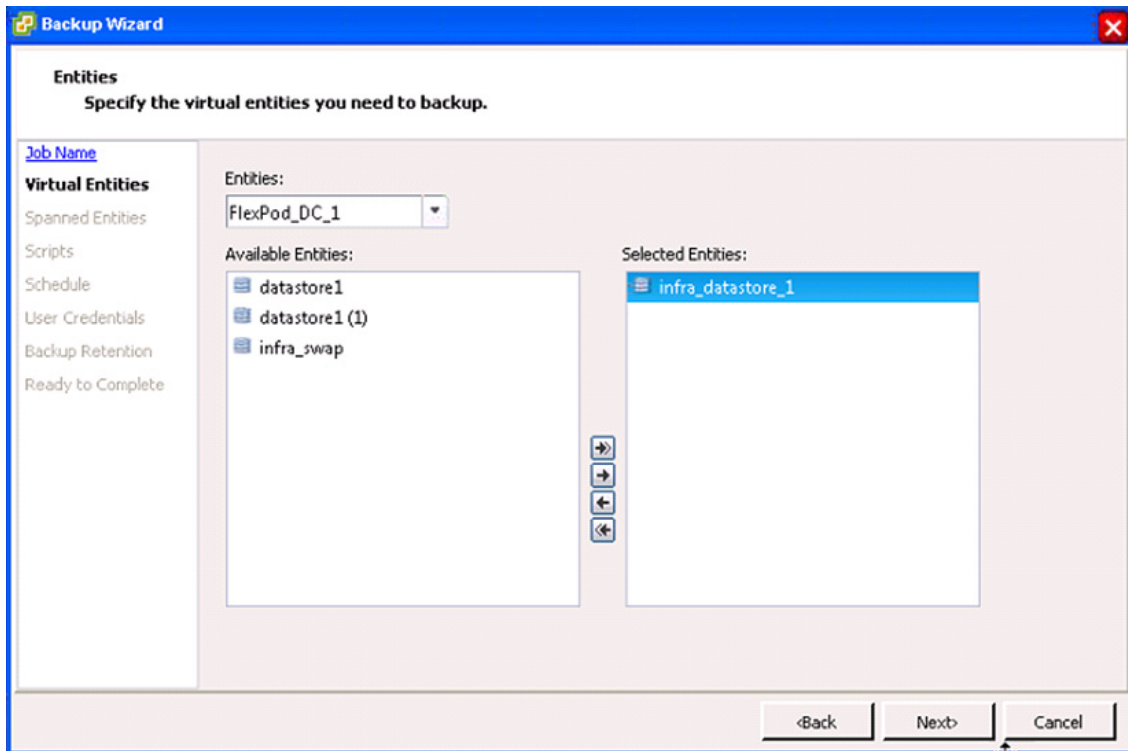
The following steps detail the procedure to configure a backup job for a datastore:

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

**Figure 93**      *Configuring Backup*

3. Type a backup job name and description.
4. If you want to create a VMware snapshot for each backup, choose Perform VMware consistency snapshot in the options pane.
5. Click **Next**.
6. Select infra\_datastore\_1 and then click  to move it to the selected entities. Click **Next**.

**Figure 94**      *Selecting Entities to Backup*



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 95**      **Setting Schedule for Backup**

**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 96** Specifying Backup Retention Period

**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

- 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 97 Summary of Backup Settings**

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 6.0

### OnCommand Unified Manager OVF Deployment

To install the OnCommand Unified Manager, complete the following steps:

Download and review the [OnCommand Unified Manager for Clustered Data ONTAP 6.0 Installation and Setup Guide](#).



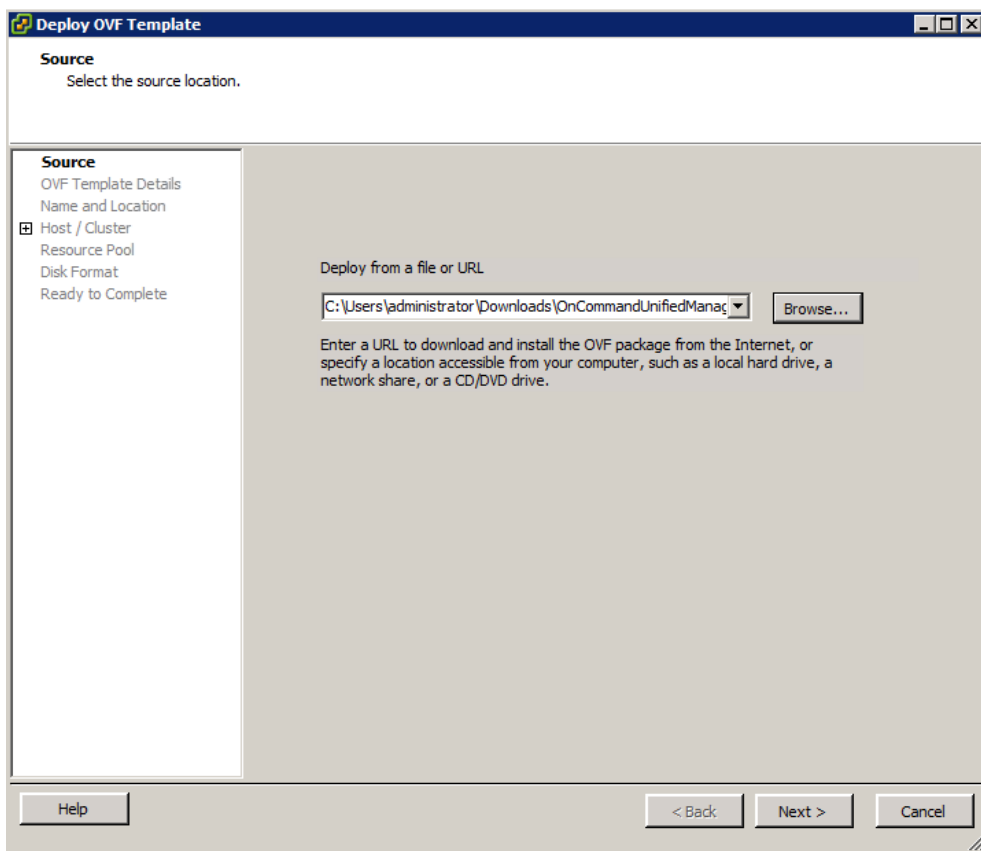
**Note**

VMware High Availability for the Unified Manager virtual appliance is not supported. The virtual appliance can be deployed on a VMware server that is a member of a VMware high availability environment, but utilizing the VMware High Availability functionality is not supported.

If deployment fails when using an high-availability-enabled environment due to insufficient resources, modify the following default VMware settings

- Decrease the VM resources CPU and memory settings.
  - Decrease the vSphere HA Admission Control Policy to use less than the default percentage of CPU and memory.
  - Modify the Cluster Features Virtual Machine Options by disabling the VM Restart Priority and leaving the Host Isolation Response powered on.
1. Download the OnCommand Unified Manager (OnCommandUnifiedManager-6.0.ova), from [http://support.netapp.com/NOW/download/software/oncommand\\_cdot/6.0/](http://support.netapp.com/NOW/download/software/oncommand_cdot/6.0/)
  2. Log in to the vSphere Client. From the Menu, choose **File > Deploy OVF Template**.
  3. Browse the .ova file that was downloaded locally. Click **Open**.
  4. Click **Next**.

**Figure 98**      *Deploying OVF Template: Select Source Locator*



5. Click **Next**.
6. Accept the EULA by clicking **Accept** and then **Next**.
7. Enter the name of the VM and select the FlexPod\_DC\_1 data center to host the VM. Click **Next**.
8. Select OnCommandDB as the storage, click **Next**.
9. Click **Next**.
10. Choose **IB-MGMT Network** as the destination network to the nat source network. Click **Next**.



**Figure 99**      **Deploying OVF Template: Network Mapping**

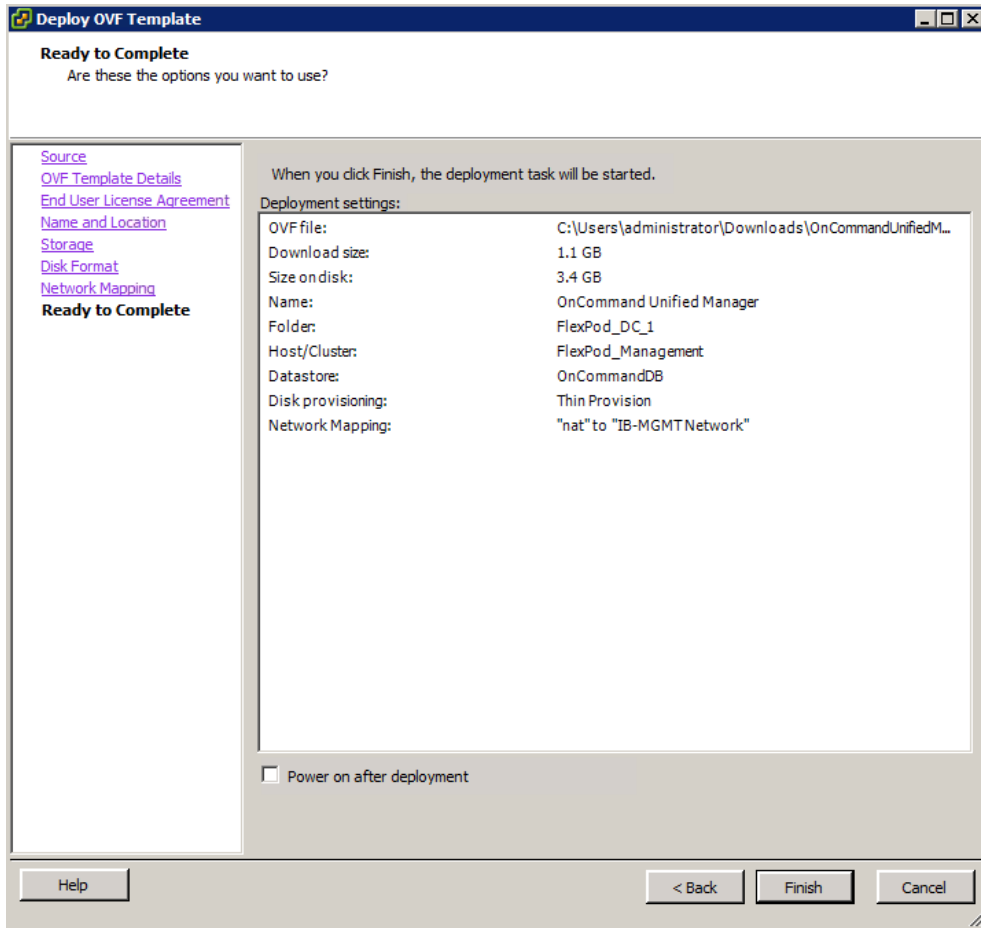
The screenshot shows the 'Deploy OVF Template' wizard at the 'Network Mapping' step. The title bar reads 'Deploy OVF Template'. Below the title bar, the step is labeled 'Network Mapping' with the question 'What networks should the deployed template use?'. On the left, a navigation pane lists steps: 'Source', 'OVF Template Details', 'End User License Agreement', 'Name and Location', 'Storage', 'Disk Format', and 'Network Mapping' (which is bolded and followed by 'Ready to Complete'). The main area is titled 'Map the networks used in this OVF template to networks in your inventory'. It contains two columns: 'Source Networks' and 'Destination Networks'. The 'Source Networks' column has one entry, 'nat', which is selected. The 'Destination Networks' column has one entry, 'IB-MGMTNetwork', which is also selected. Below these columns is a 'Description:' label and a text box containing 'The LAN network'. At the bottom of the window are three buttons: 'Help', '< Back', and 'Next >', and a 'Cancel' button.

Source Networks	Destination Networks
nat	IB-MGMTNetwork

Description:  
The LAN network

11. Review the settings to be configured. Click **Finish**.

**Figure 100**      *Deploying OVF Template: Review Deployment Settings*



12. Click **Close**.
13. Power on the virtual machine and open a console window.

## OnCommand Unified Manager Basic Setup

1. Set up OnCommand Unified Manager by answering the following questions in the console window:

```

Geographic area: <<Enter your geographic location>>
Please select the city or region corresponding to your time-zone: <<Enter your
city or region>>
Time zone: <<Choose the preferred time zone>>
Please enter your choice (1/2): 2
Enter the host FQDN: <<var_oncommand_server_fqdn>>
Enter the host IP address: <<var_oncommand_server_ip>>
Enter the network mask: <<var_oncommand_server_netmask>>
Enter the IP address of the default gateway: <<var_oncommand_server_gateway>>
Enter the primary DNS address: <<var_nameserver_ip>>
Enter the secondary DNS address: <<optional>>
Enter additional search domains separated by spaces (Optional field): <<optional>>
Are these values correct? (y/n):y
  
```

The network configuration and SSL certificate generation for HTTPS are completed. The OnCommand Unified Manager services are started.

2. Create a Maintenance User account.

**Note**

The maintenance user manages and maintains the settings on the OnCommand Unified Manager virtual appliance.

```
Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>
```

## OnCommand Unified Manager Initial Setup

1. Using a web browser navigate to the OnCommand Unified Manager using URL: `https://<<var_oncommand_server_ip>>`
2. Login using the Maintenance User account credentials.
3. Click the radio button to enable AutoSupport capabilities.
4. Click **Continue**.

**Figure 101**      **OnCommand Unified Manager: Initial Setup - Part1**

**OnCommand Unified Manager Initial Setup**

### AutoSupport

As a part of NetApp®'s effort to provide quality support to its customers, OnCommand™ Unified Manager is equipped with advanced NetApp AutoSupport capabilities. The AutoSupport service sends specific, predefined information to NetApp Support periodically. NetApp Support proactively uses this information to ensure the correct operation of OnCommand Unified Manager, and to assist customers in maintaining the integrity of their data centers. The OnCommand Unified Manager AutoSupport service is described below.

### Customer Benefits

By providing this data, the AutoSupport suite of Web-based applications is able to provide the following benefits to NetApp customers:

- Evaluate, model, and deploy more efficient storage solutions
- Enable improved system availability and avoid disruptions to your business with proactive system health checks
- Save time and money with faster and more efficient incident resolution
- Optimize your system operations with storage efficiency profiling
- Free 24x7 Web-based access with valid product warranty

### Security Aspects

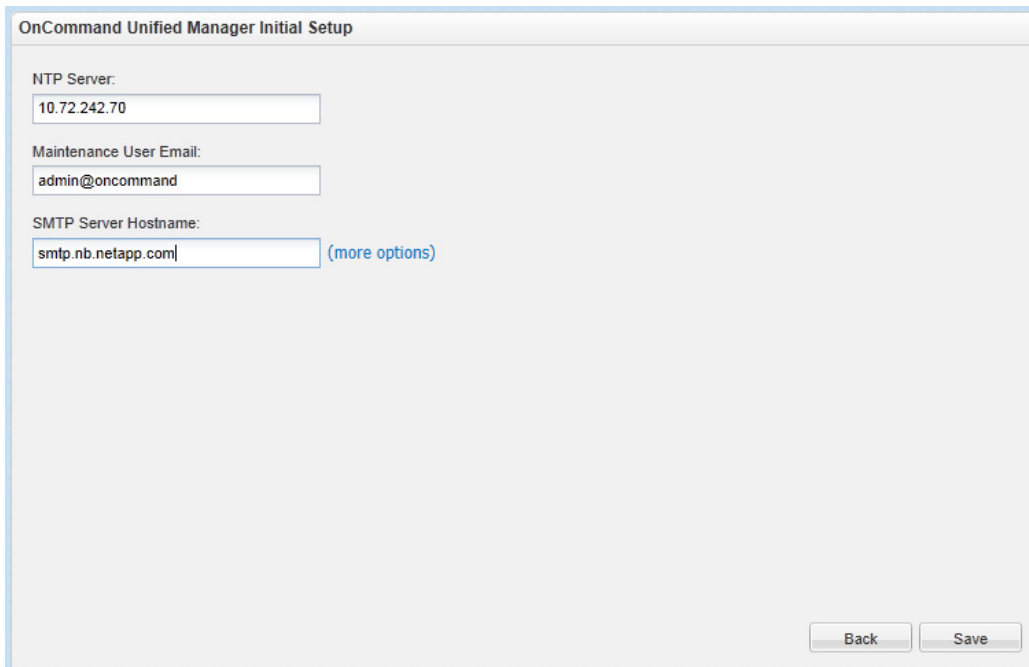
Do you want to enable AutoSupport capabilities?

☒ Yes (recommended).

☐ No, not at this time.

**Continue**

5. Provide the NTP Server IP address `<<var_global_ntp_server_ip>>`.
6. Provide the Maintenance User Email `<<var_storage_admin_email>>`.
7. Provide the SMTP Server Hostname.

**Figure 102**      **OnCommand Unified Manager: Initial Setup - Part2**


OnCommand Unified Manager Initial Setup

NTP Server:  
10.72.242.70

Maintenance User Email:  
admin@oncommand

SMTP Server Hostname:  
smtp.nb.netapp.com| (more options)

Back Save

8. Click **Save**.
9. Click **Add Cluster**.
10. Provide the Cluster Management IP address, Username, Password, Protocol and Port.

**Figure 103**      **Adding Cluster**


Add Cluster ?

Host Name or IP Address: 10.238.162.235

User Name: admin

Password: .....

Protocol: ☒ HTTPS ☐ HTTP

Port: 443

Add Cancel

11. Click **Add**.



**Note** The Cluster Add operation might take a couple of minutes.

12. Once the cluster is added it can be accessed by clicking on the **Storage** tab and choosing **Clusters**.

**Figure 104** Window Showing the Added Cluster

NetApp OnCommand Unified Manager

Help | Administration | admin | S

Dashboard Events Storage Jobs All Search

**Filters**

**Status** Clear

☐ Critical

☐ Error

☐ Warning

☒ Normal

**Communication Status** Clear

☐ Good

☐ Not Reachable

**System Health** Clear

☐ OK

☐ OK with suppressed

☐ Degraded

☐ Components not reachable

**Clusters** ?

+ Add Edit X Remove

Cluster	Communication	System Health	Host Name or IP Address	OS Version	Node Count
FAS3250-Cluster	Good	OK	10.238.162.235	8.2P4 Cluster-Mode	2

## NetApp NFS Plug-In 1.0.20 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.

- 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\.

- In the VMware vSphere Client connected to the vCenter Server, choose **Home > Solutions and Applications > NetApp**.
- In the Monitoring and Host Configuration capability navigation pane, choose **Tools**.
- Under NFS Plug-in for VMware VAAI, click **Install on Host**.

**Figure 105**      *Installing NFS Plug-in for VMware VAAI on Host*

**Monitoring and Host Configuration**

- Overview
- Storage Details - SAN
- Storage Details - NAS
- Data Collection
- Tools
- Discovery Status

**MBR Tools**

This ESX console-based tool tests and aligns guest file systems on a VMKFS for VMFS and NFS datastores. Aligning the file system block boundaries to the underlying NetApp storage system LUN ensures the best storage performance. The data is migrated from a backup of the original -flat.vmdk file to a new, properly aligned -flat.vmdk file.

Download (For ESX 4.x)

Download (For ESX 4.x and ESX 6.x)

**Guest OS Tools**

Guest OS timeout scripts set the SCSI I/O timeout values for supported guest operating systems, which ensure correct failover behavior. Both 60-second and 190-second timeout values are supported. Right-click the URL for the .iso file containing the script you need and copy it to the clipboard.

**Note:** Before selecting an .iso file, check the Release Notes for information about the recommended timeout values.

**60-second timeout settings:**

- Linux OS**  
https://10.238.162.97:8143/vsc/public/writable/linux\_gos\_timeout-install.iso
- Windows OS**  
https://10.238.162.97:8143/vsc/public/writable/windows\_gos\_timeout.iso
- Solaris OS**  
https://10.238.162.97:8143/vsc/public/writable/solaris\_gos\_timeout.iso

**190-second timeout settings:**

- Linux OS**  
https://10.238.162.97:8143/vsc/public/writable/linux\_gos\_timeout\_190-install.iso
- Windows OS**  
https://10.238.162.97:8143/vsc/public/writable/windows\_gos\_timeout\_190.iso
- Solaris OS**  
https://10.238.162.97:8143/vsc/public/writable/solaris\_gos\_timeout\_190.iso

**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, which are installed on the ESXi hosts. These libraries enable VMware to execute various primitives on files stored on NetApp storage systems. You can install the plug-in on a host using VSC.

NFS Plug-in for VMware VAAI version: 1.0-020

**Note:** Before you install the NFS Plug-in for VMware VAAI, check the Release Notes for more information concerning the latest version of the plug-in.

Install on Host

- 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.

- Choose **Home > Inventory > Host and Clusters**.
- For each host (one at a time), right-click the host and choose Enter Maintenance Mode.
- 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.



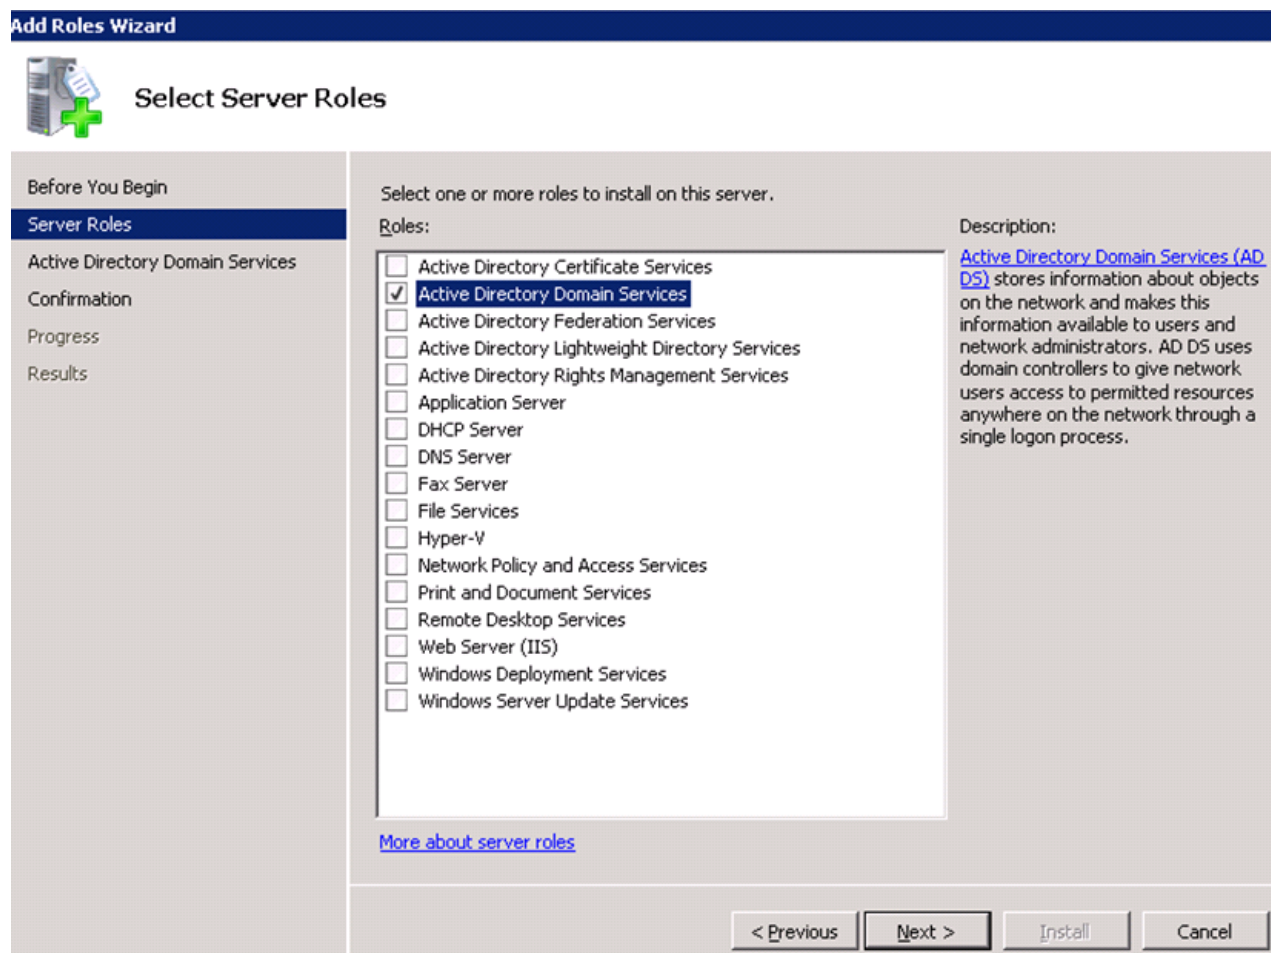
**Note** This process requires several reboots.

47. Open Server Manager.
48. On the left, click **Roles**, then choose **Add Roles** on the right.



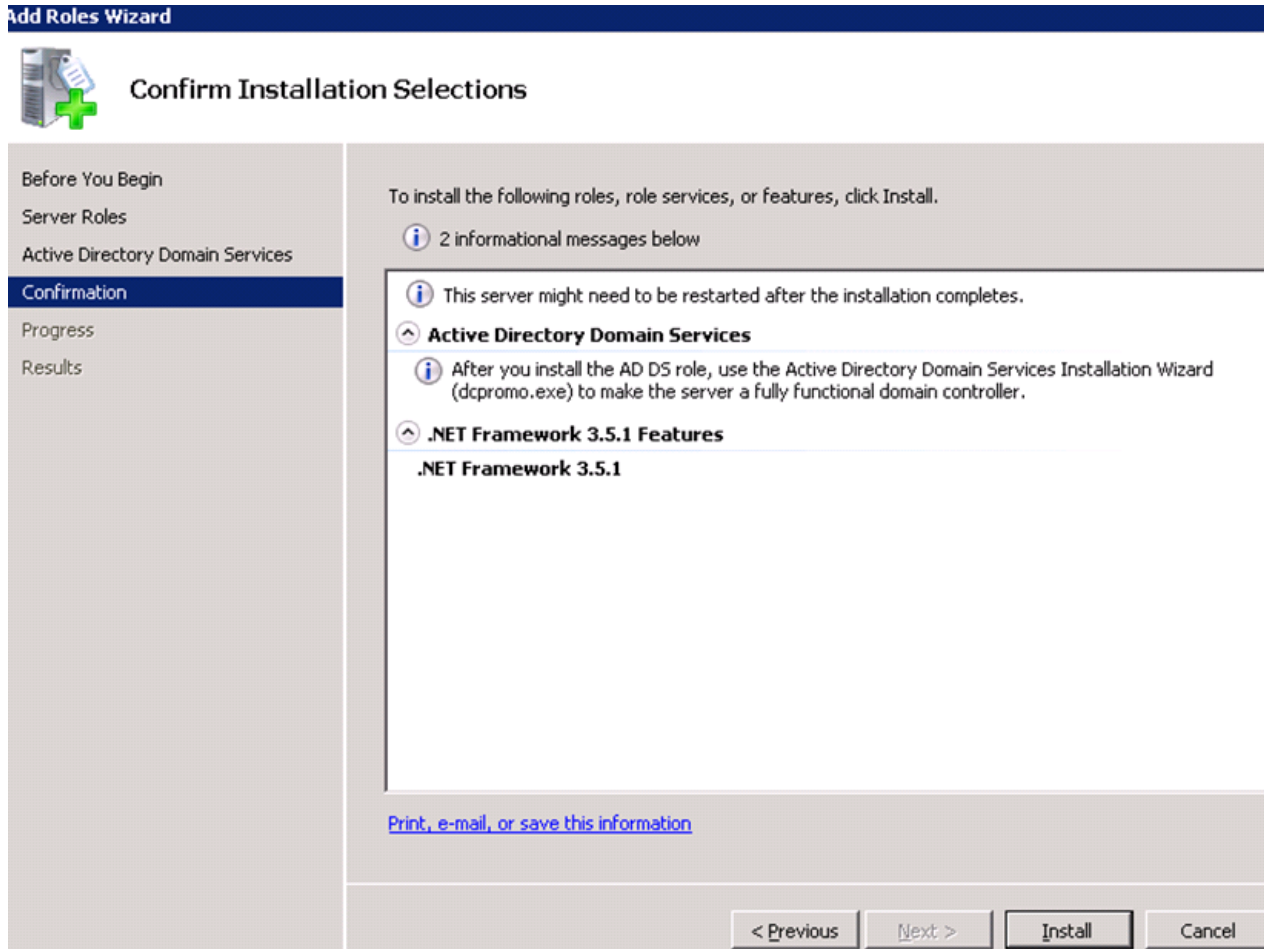
49. Click **Next**.
50. In the list, check the check box next to Active Directory Domain Services.
51. In the popup, click **Add Required Features** to add .NET Framework 3.5.1.

**Figure 106**      *Selecting Roles to Install on the Server*



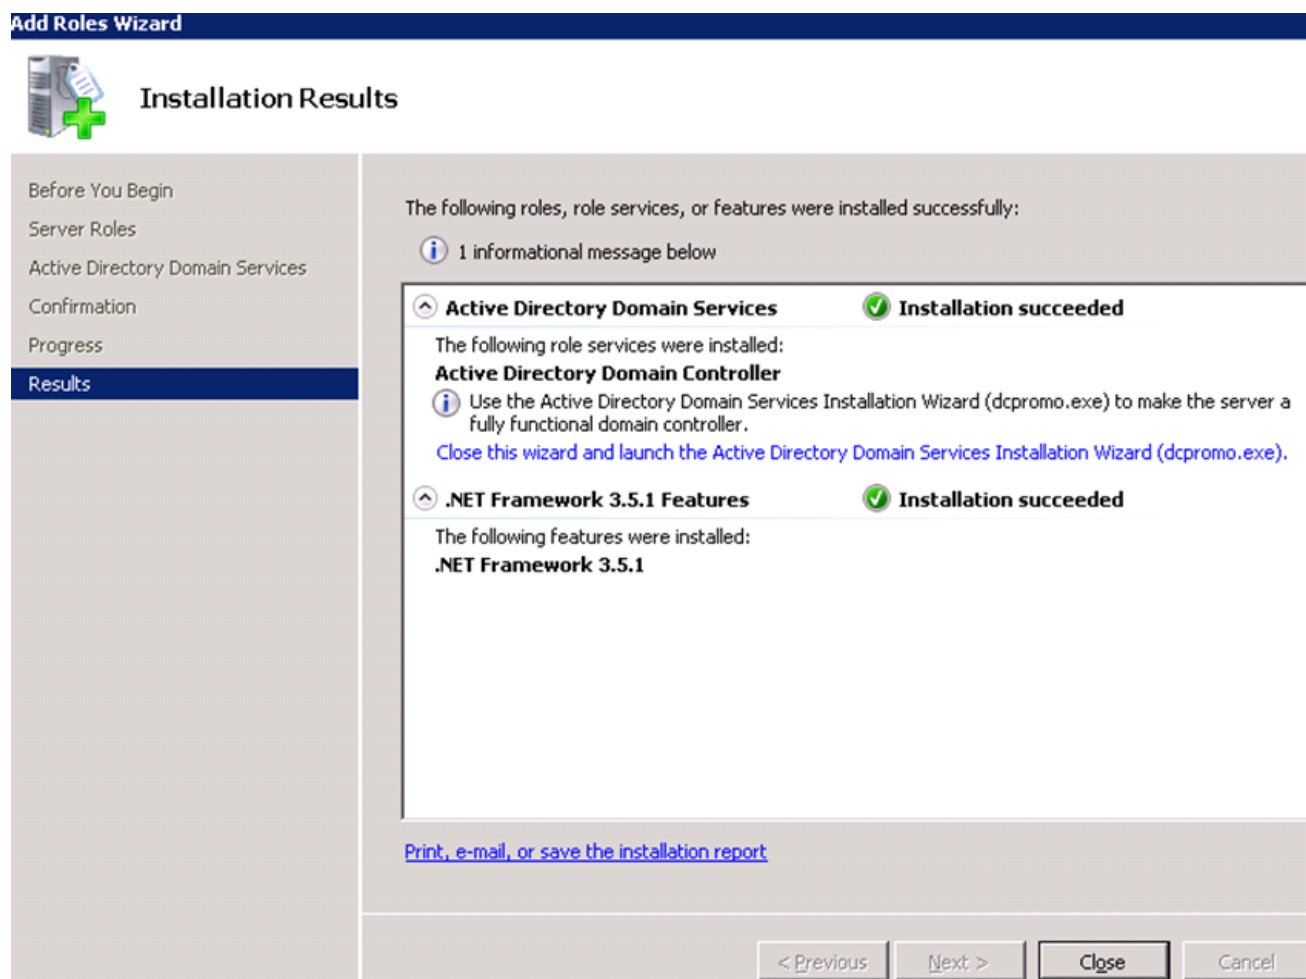
52. Click **Next**.
53. Click **Next**.

**Figure 107**      *Confirming Installation Selections*




54. Click **Install**.

**Figure 108**      **Installation Completion Summary**



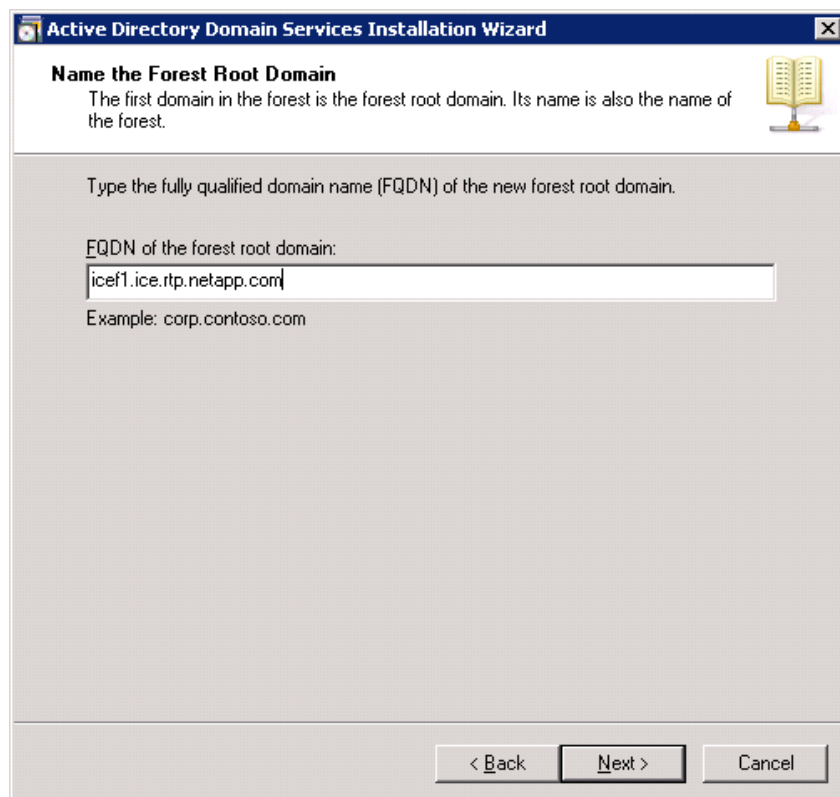
55. In the middle of the window, click **Close** this wizard and launch the Active Directory Domain Services Installation Wizard (dcpromo.exe).
56. In the Active Directory Domain Services Installation Wizard, click **Next**.
57. Click **Next**.
58. Choose Create a new domain in a new forest and click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

**Figure 109**      *Creating a Domain Controller for a New Forest*



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

**Figure 110**      **Naming the Forest Root Domain**



**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:

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 111**      **Selecting Additional Options for the Domain Controller**



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

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

## Configuring Cisco VM-FEX with the UCS Manager

### Background

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

- There is no need for extra hardware 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. See [“VM-FEX Virtual Interfaces” section on page 208](#), under [Appendix, page 179](#).

**Figure 112**      *Creating Dynamic vNIC 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.

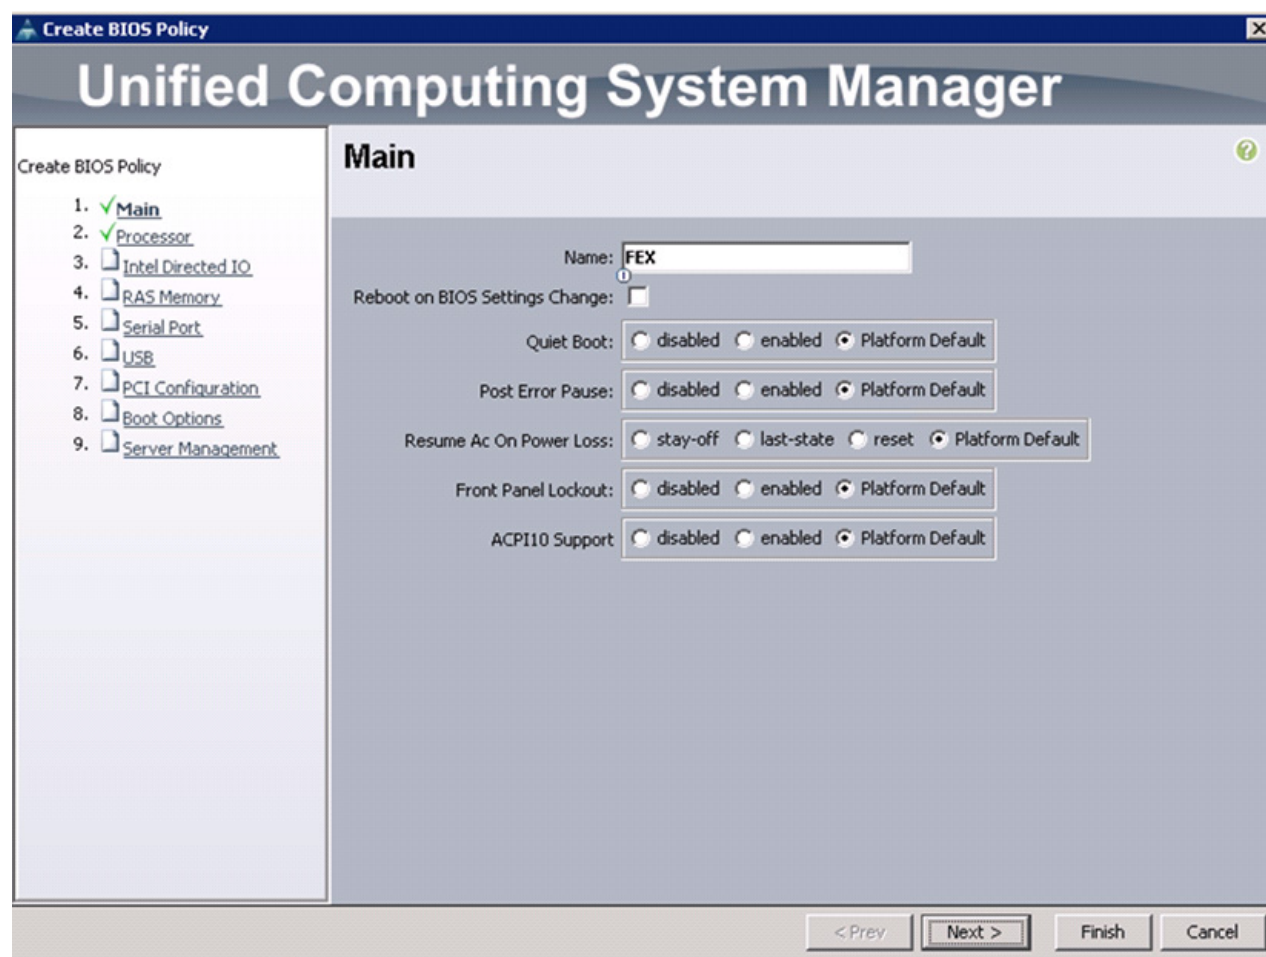
## Create a Server BIOS Policy

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 113 Creating BIOS Policy



4. For Virtualization Technology (VT) and Direct Cache Access, click the **enabled** radio button.

Figure 114 Processor Settings for 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

## 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, click the **enabled** radio button.

**Figure 115 Intel Directed IO Settings for BIOS Policy**




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

## Create a VM-FEX Enabled Service Profile Template

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

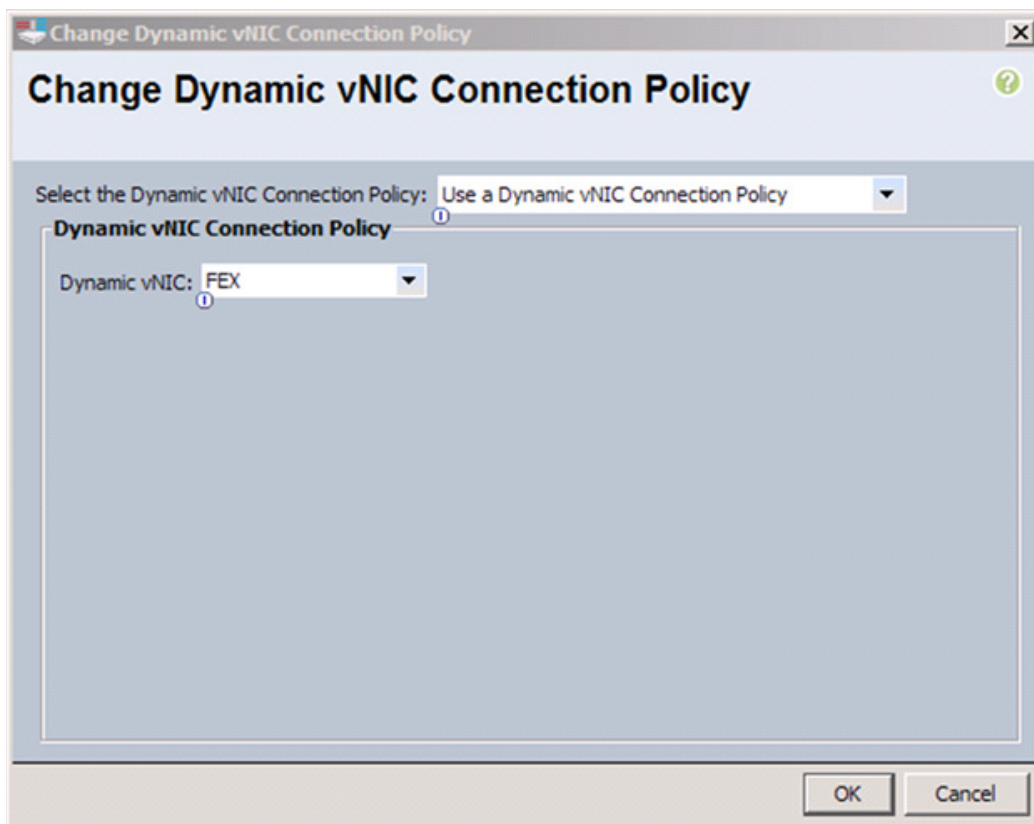
1. Click the **Servers** tab in the left navigation pane and expand the Service Profile Templates.
2. Right-click VM-Host-Infra-Fabric-A and choose **Create a Clone**.
3. Type a clone name and choose an organizational owner for the new service profile template.

**Figure 116**      **Cloning Service Profile Template**



4. Click **OK**.
5. When notified that the service profile clone was successfully created, check the checkbox to navigate to the Service Profile Template. Click **OK**. The Service Template navigation window appears.
6. 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.
7. Choose Use a Dynamic vNIC Connection Policy from the drop-down menu and the previously created Dynamic vNIC policy. Click **OK**.


**Figure 117**      **Changing the Dynamic vNIC Connection Policy**



8. Click **OK** when notified that the vNIC connection policy was successfully modified.
9. From the Service Template properties window, click the **Policies** tab.
10. Expand the BIOS Policies in the Policies section of the working pane.

11. Choose the previously defined FEX BIOS policy and click **Save Changes**.
12. Click **OK** when notified that the Service Profile Template was successfully modified.

**Figure 118**      *Choosing a BIOS Policy*



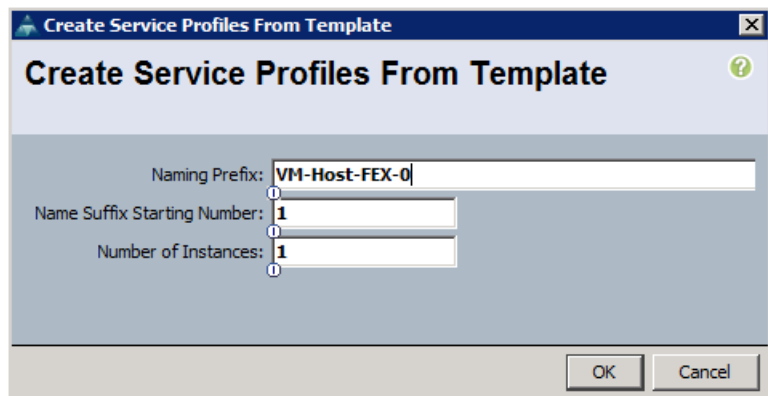
## Create VM-FEX Service Profile

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

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



**Figure 119**      *Creating Service Profile from Template*



8. Click **OK** in the confirmation message.
9. 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\(3\)](#)
  - b. [Downloads Home](#) > [Products](#) > [Servers - Unified Computing](#) > [Cisco UCS C-Series Rack-Mount UCS-Managed Server Software](#) > [Unified Computing System \(UCS\) Drivers-1.4\(5b\)](#)
2. Follow the steps necessary to download the software bundles located on the ISO image.
3. Mount the ISO image and copy the appropriate vib file from the VMware > VM-FEX > Cisco directory to the local machine.
4. From the vCenter vSphere Client, choose the infra\_datastore\_1 from the **Inventory > Datastores and Datastore Clusters** navigation menu.
5. Under Basic Tasks, click **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:

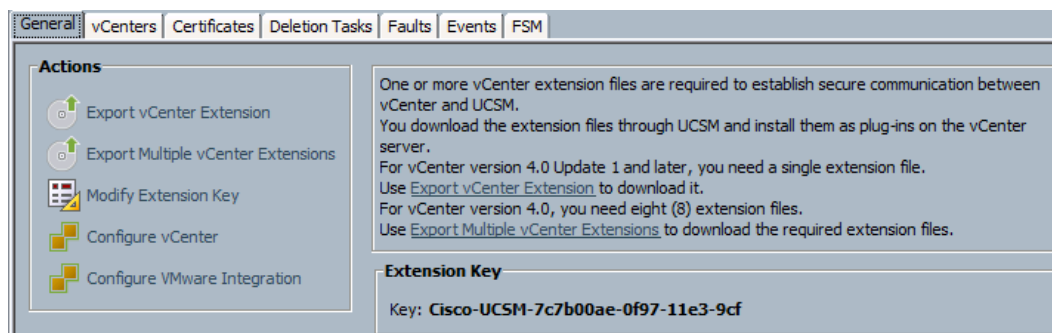
```
esxcli -s <host_ip> -u root -p <host_password> software vib install -v
/vmfs/volumes/infra_datastore_1/VM-FEX/cisco-vem-v151-5.1-1.1.1.1.vib
```

## 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.
3. Choose **Configure VMware Integration** in the Actions area to start the Configuration wizard.

**Figure 120**      *Configuring VMware Integration*



4. Follow the instructions and click **Export** and complete the steps to install the UCS Manager extension file in vCenter.

Figure 121 Exporting vCenter Extension Plug-in



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



Figure 122 Defining VMware DVS

**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 IB-MGMT-VLAN, and indicate it is the native VLAN.

**Figure 123**     *Defining Port Profile for VMware Integration*

**Configure VMware Integration**

# 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 Port Profile

**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 checked="" type="checkbox"/>	IB-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"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

**Profile Client**

Name:

Description:

Datacenter:

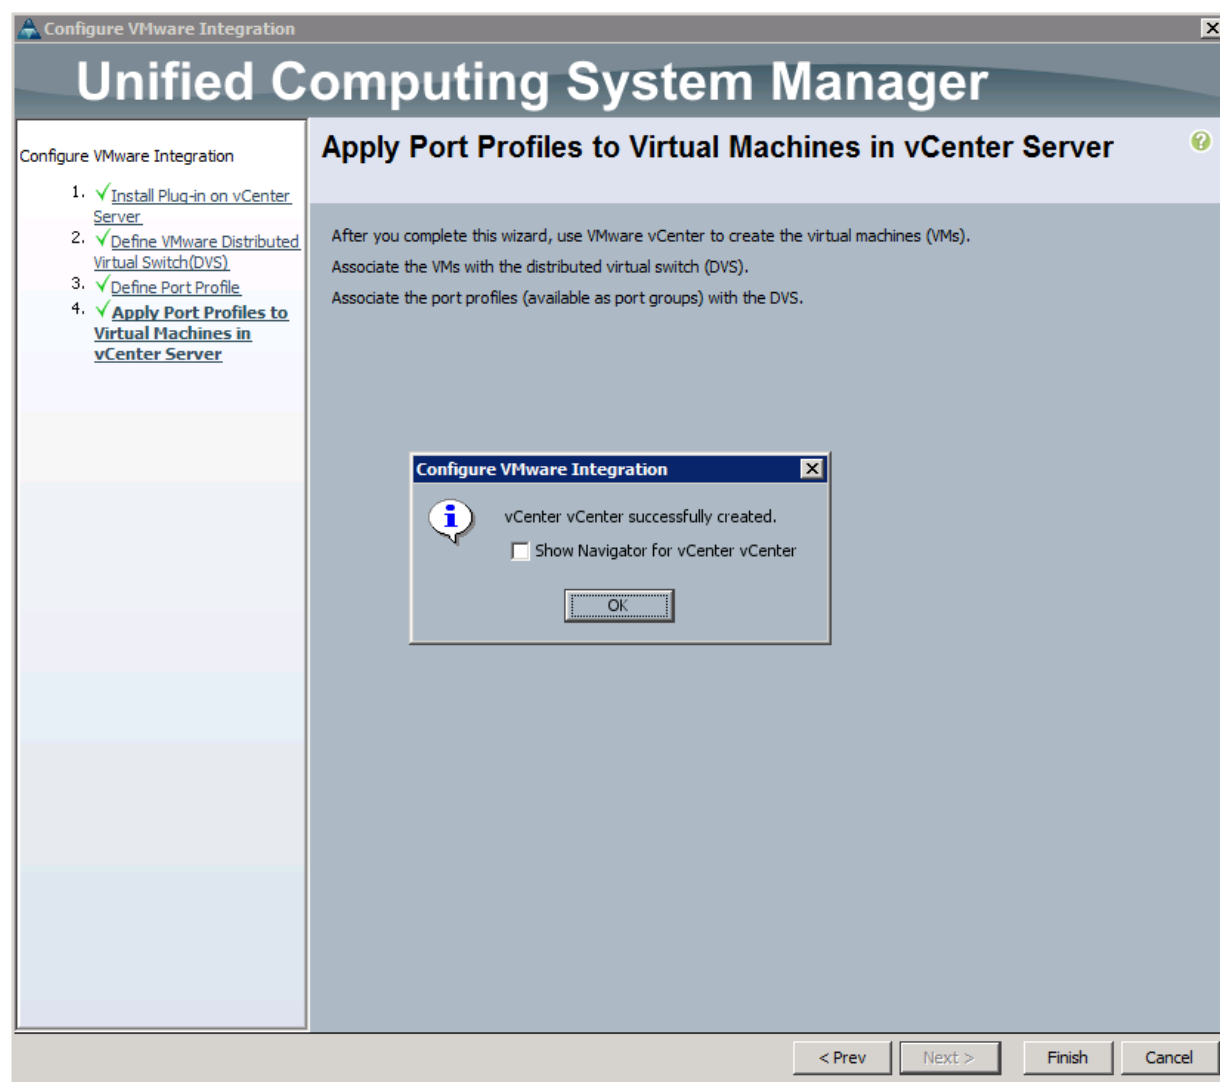
Folder:

Distributed Virtual Switch:

< Prev    Next >    Finish    Cancel

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

**Figure 124** VMware Integration is succesful



**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.

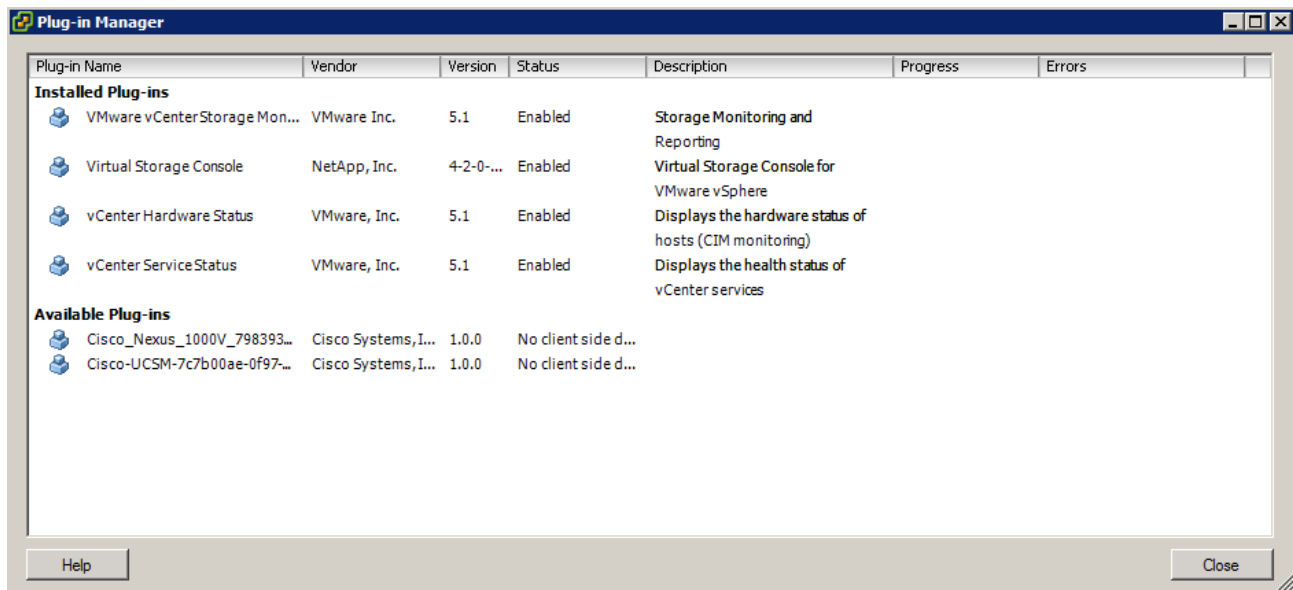
## Validate Setting in VMware vCenter

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**.

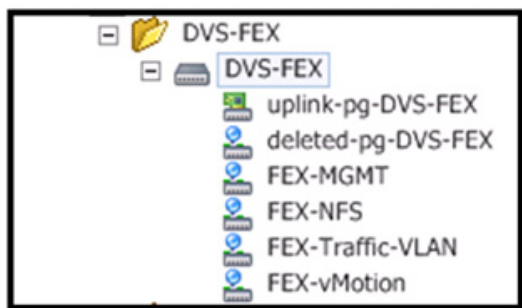
The popup windows shows that the Cisco UCS Manager is already integrated in vCenter.

**Figure 125** Cisco UCS Manager Integrated in vCenter



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

**Figure 126** Verifying FEX-DVS in Cisco UCS Manager



## Standard Operations

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

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

### Port Profiles

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

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

## Port Profile Client

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

## Create a VM-FEX Port Profile

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

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



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

9. (Optional) Choose High Performance.



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

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

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

**Figure 127**      **Creating 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"/>	IB-MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

OK Cancel

Or

**Figure 128**      **Creating Port Profile with High Performance**

**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"/>	IB-MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<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 129**      *Creating Profile Client*

Or

**Figure 130**      *Creating Profile Client for DVS-FEX*

4. The client profile created will appear in your distributed virtual switch DVS-FEX in vCenter as a port group.
5. Repeat these steps 1 through 4, as necessary for the workloads in the environment.



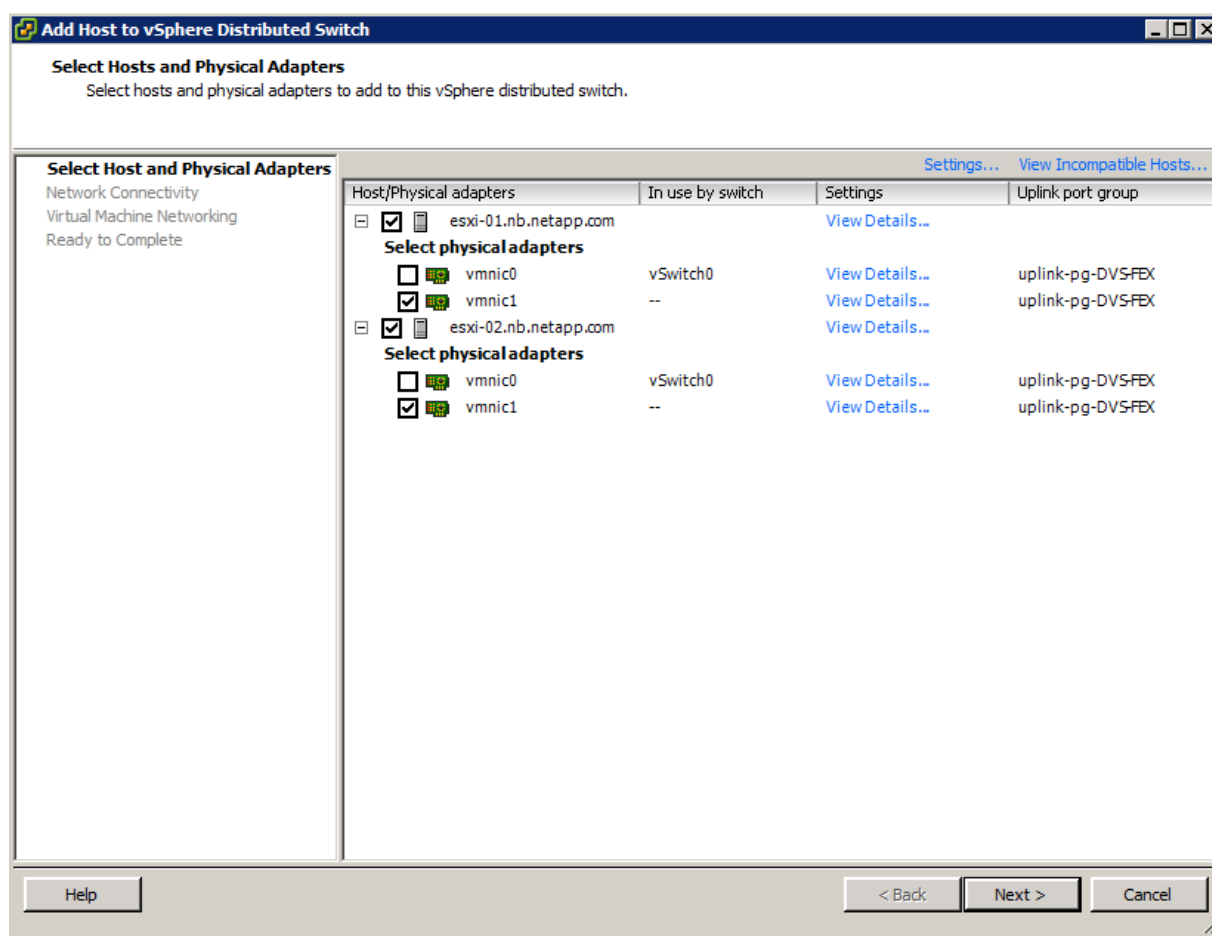
## Migrate Networking Components for ESXi Hosts to Cisco DVS-FEX

### vCenter Server VM

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

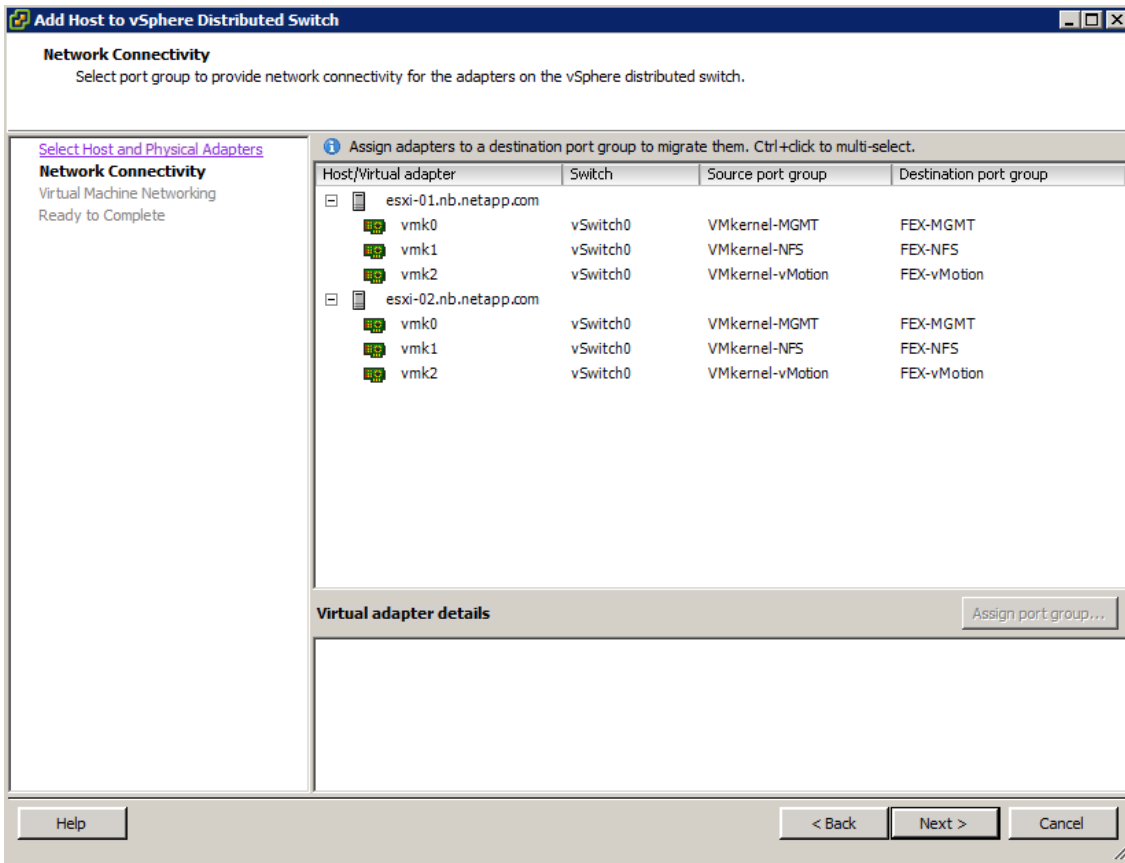
1. In the VMware vSphere client connected to vCenter, choose **Home > Networking**.
2. Expand the vCenter, DataCenter, and DVS-FEX folders. choose the DVS-FEX switch.
3. Under Basic Tasks for the vSphere distributed switch, choose **Add a Host**.
4. For both hosts, choose **vmnic1** and choose the **uplink-pg-DVS-FEX**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

**Figure 131**      *Selecting Physical Adapters to Add to the Vsphere Distributed Switch*



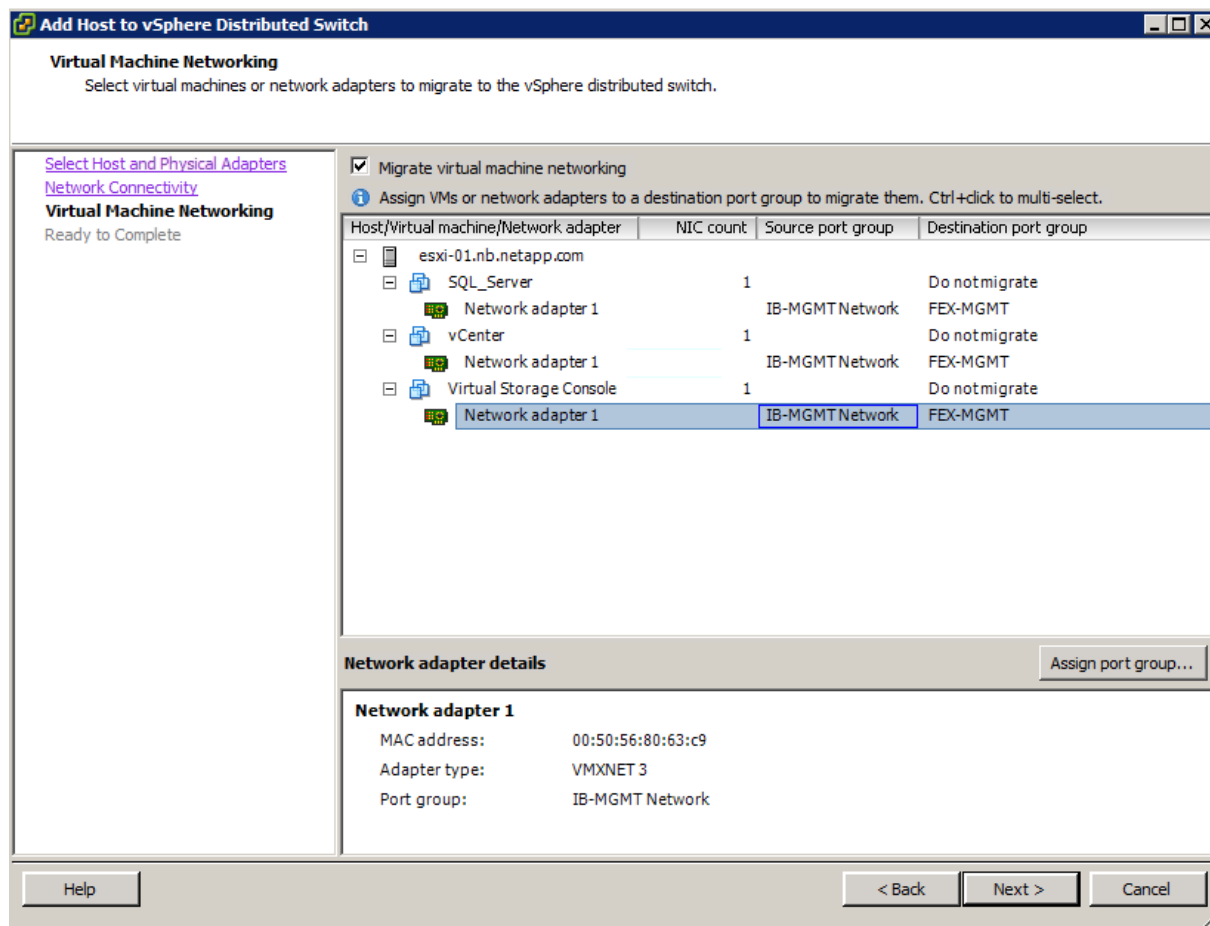
5. For all VMkernel ports, choose the appropriate destination Port Group from the Cisco DVS-FEX. Click **Next**.

**Figure 132**      **Selecting Destination Port Groups**



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

**Figure 133**      **Selecting Port Groups on Each VM to Migrate Individually**



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

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

## VM-FEX Virtual Interfaces

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

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

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

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

For more information about the maximum number of VIFs supported by your hardware configuration, see the [Cisco UCS 6100 and 6200 Series Configuration Limits for Cisco UCS Manager](#) for your software release. [Table 31](#) and [Table 32](#) reference these limits.

**Table 31** *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 32** *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)

## Server-Side Flash—NetApp Flash Accel with Fusion-I/O

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

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

#### Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

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

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

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

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

#### Cisco Nexus 5548UP A and Cisco Nexus 5548UP B

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

From the global configuration mode, run the following commands:

```
interface Po10
switchport trunk allowed vlan add 3170
exit
interface Po13
switchport trunk allowed vlan add 3170
exit
interface Po14
switchport trunk allowed vlan add 3170
exit
copy run start
```


**Note**

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

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

### Cisco UCS Manager

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

**Figure 134**      **Creating 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

## Create vNIC Templates on the Cisco UCS

### Cisco UCS Manager

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

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Choose **Policies > root**.
3. Right-click vNIC Templates.
4. Choose **Create vNIC Template**.
5. Enter **vNIC\_OOB\_MGMT\_A** as the vNIC template name.
6. Click the **Fabric A** radio button.
7. Keep the **Enable Failover** checkbox unchecked.

8. Under Target, make sure that the **VM** checkbox is unchecked.
9. Click the **Updating Template** radio button for the Template Type.
10. Under VLANs, check the checkbox for **OOB-MGMT-VLAN**.
11. Keep the MTU at 1500.
12. For MAC Pool , choose **MAC\_Pool\_A** from the drop-down list.
13. For Network Control Policy, choose **Enable\_CDP** from the drop-down list.
14. Click **OK** to create the vNIC template.
15. Click **OK**.

**Figure 135**      *Creating 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

Select	Name	Native VLAN
<input type="checkbox"/>	NFS-VLAN	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	OOB-MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>

**+** Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

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



21. Click the **Fabric B** radio button.
22. Uncheck the **Enable Failover** checkbox.
23. Under Target, make sure that the **VM** checkbox is unchecked.
24. Choose the **Updating Template** radio button for the Template Type.
25. Under VLANs, check the checkbox for OOB-MGMT-VLAN.
26. Keep the MTU at 1500.
27. For MAC Pool, choose **MAC\_Pool\_B** from the drop-down list.
28. For Network Control Policy, choose **Enable\_CDP** from the drop-down list.
29. Click **OK** to create the vNIC template.
30. Click **OK**.

**Figure 136**      *Creating 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"/>	Native-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	OOB-MGMT-VLAN	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>
<input type="checkbox"/>	vMotion-VLAN	<input type="radio"/>

**Create VLAN**

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Dynamic vNIC Connection Policy:

OK Cancel

## Configure ESXi Hosts

The procedures in the following two subsections will have to be repeated for each ESXi host on which Flash Accel and Fusion-io need to be configured.

### Shut Down the ESXi Host

1. Log in to VMware vCenter Server using vSphere Client as FlexPod admin user.
2. Choose **Hosts and Clusters** from the Inventory pane.
3. Right-click the ESXi server on which you want to install Flash Accel and choose the option **Enter Maintenance Mode**.




---

**Note** This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

---

4. Right-click the ESXi server in the Maintenance Mode and choose **Shut Down**.
5. Click **OK**.

### Add Out-of-Band Management vNICs to ESXi Host's Service Profile

This procedure will add Cisco Virtual Network Interfaces to the ESXi hosts and will require a server reboot. In this procedure, the service profile is unbound from its template and modified. Alternatively, the Service Profile Template could be modified, but that would require all servers bound to the template to be rebooted.



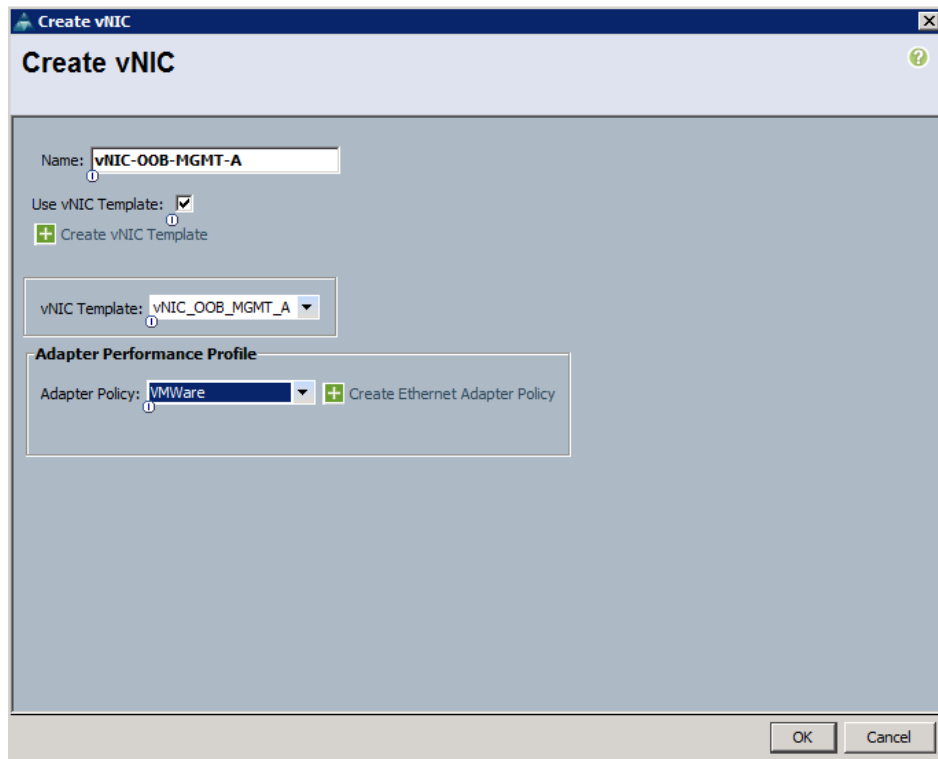

---

**Note** If the Service Profile Template is modified, all servers bound to the template would inherit the configuration changes, whether or not Flash Accel is installed on them.

---

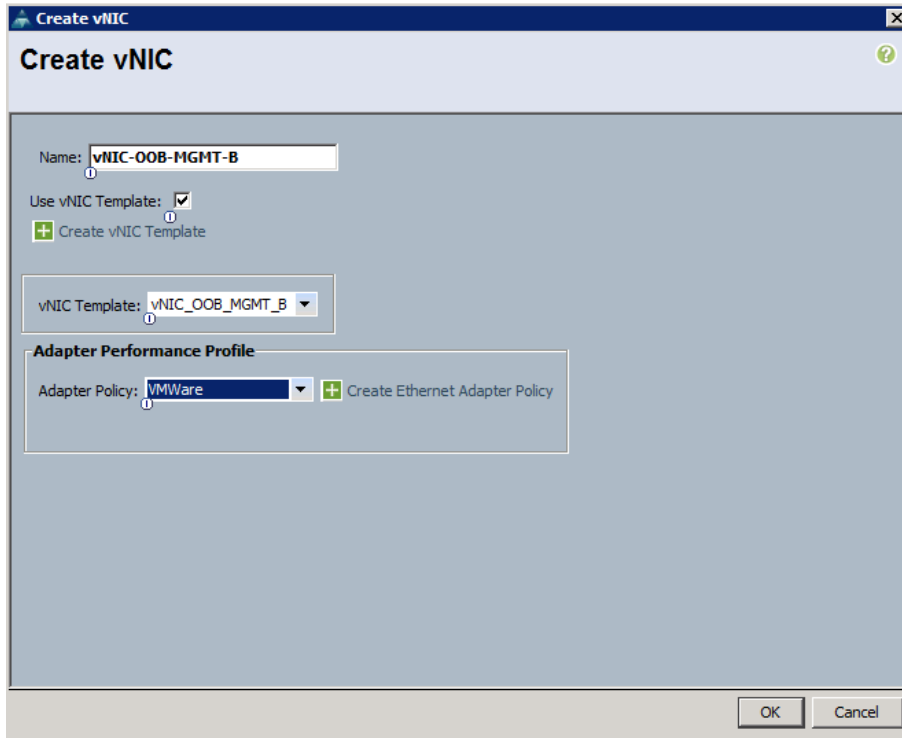
1. Using a web browser, log in to the Cisco UCS Manager as admin.
2. In the Cisco UCS Manager, click the **Servers** tab in the navigation pane.
3. Under **Servers > Service Profiles > root**, select the Service Profile for the server that was shut down.
4. Under the Actions pane, choose **Unbind from the Template**.
5. Click **Yes** to unbind the Service Profile from the template.
6. Click **OK**.
7. Choose the **Network** tab on the right pane.
8. At the bottom of the screen, click **Add**.
9. In the Create vNIC dialog box, enter **vNIC-OOB-MGMT-A** as the name of the vNIC.
10. Check the **Use vNIC Template** checkbox.
11. For vNIC Template, choose **vNIC\_OOB\_MGMT\_A** from the drop-down list.
12. For Adapter Policy, choose **VMware** from the drop-down list.
13. Click **OK** to add the vNIC to the Service Profile.

**Figure 137**      *Creating a vNIC Using vNIC Template*



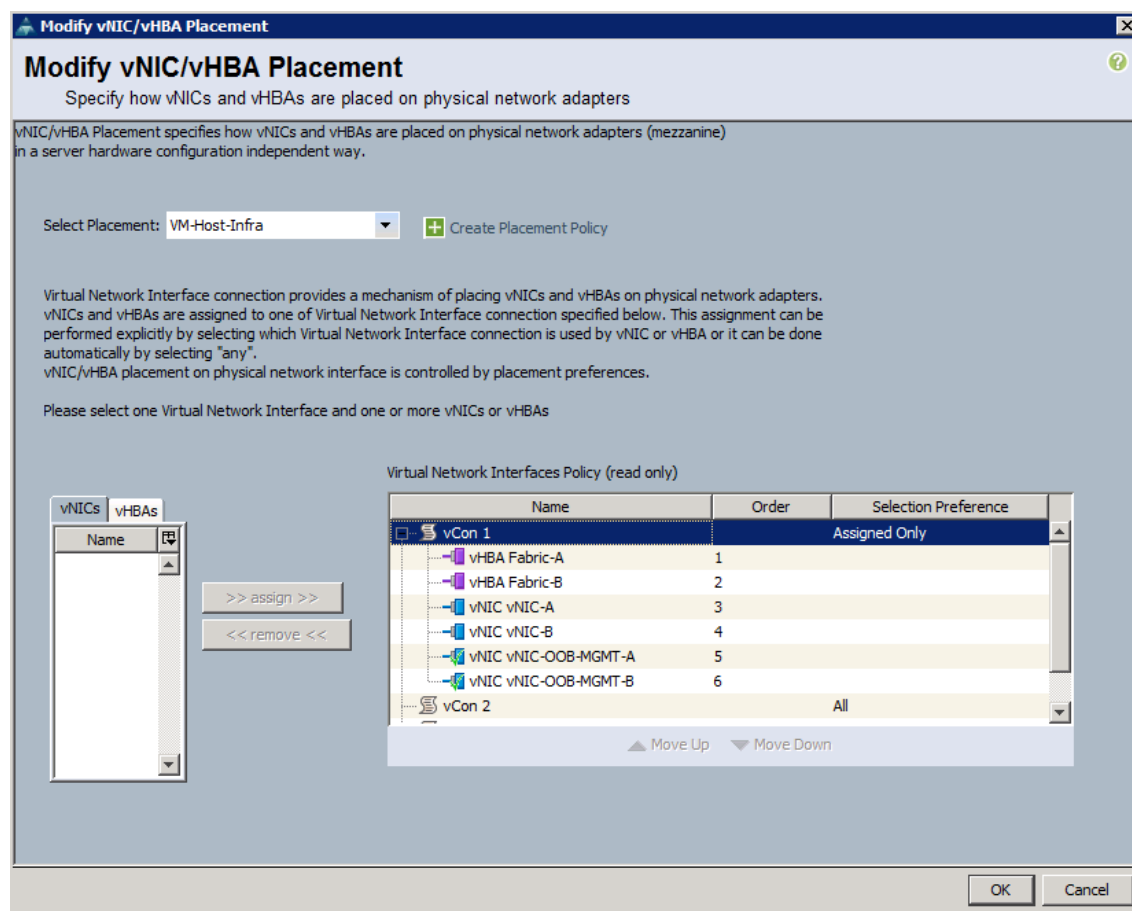
14. In the Create vNIC dialog box, enter **vNIC-OOB-MGMT-B** as the name of the vNIC.
15. Check the **Use vNIC Template** checkbox.
16. For vNIC Template, choose **vNIC\_OOB\_MGMT\_B** from the drop-down list.
17. For Adapter Policy, choose VMware from the drop-down list.
18. Click **OK** to add the vNIC to the Service Profile.

**Figure 138**      *Creating a vNIC Using vNIC Template*



19. In the lower right-hand corner of the screen, click **Save Changes**.
20. Click **Yes**.
21. Click **OK**.
22. Under Actions, click **Modify vNIC/HBA Placement**.
23. Make sure that **VM-Host-Infra** is selected as the placement policy.
24. Select vCon1 and assign vNIC-OOB-MGMT-A and vNIC-OOB-MGMT-B to the virtual network interfaces policy and place them below vNIC-B.

**Figure 139**      **Modifying vNIC/vHBA Placement**



25. Click **OK** and click **Yes**.
26. Click **OK**.
27. At the top of the screen, choose Pending Activities.
28. Check the **Reboot Now** checkbox and click **OK**.
29. Check the **General** tab on the right pane.
30. In the Actions pane, select KVM Console to monitor the server until VMware ESXi has rebooted.
31. Repeat the preceding two sections for each ESXi host on which Flash Accel and Fusion-io needs to be configured.

## Add Management Port Group to ESXi Host

The procedures in the following subsection will have to be repeated for each ESXi host in the VMware Cluster irrespective of whether or not Flash Accel and Fusion-io need to be configured on them.

1. Log in to VMware vCenter Server using vSphere Client as FlexPod admin user.
2. Select Hosts and Clusters from the Inventory pane.
3. Right-click on the ESXi server on which you want to install Flash Accel and select Enter Maintenance Mode.



**Note** This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

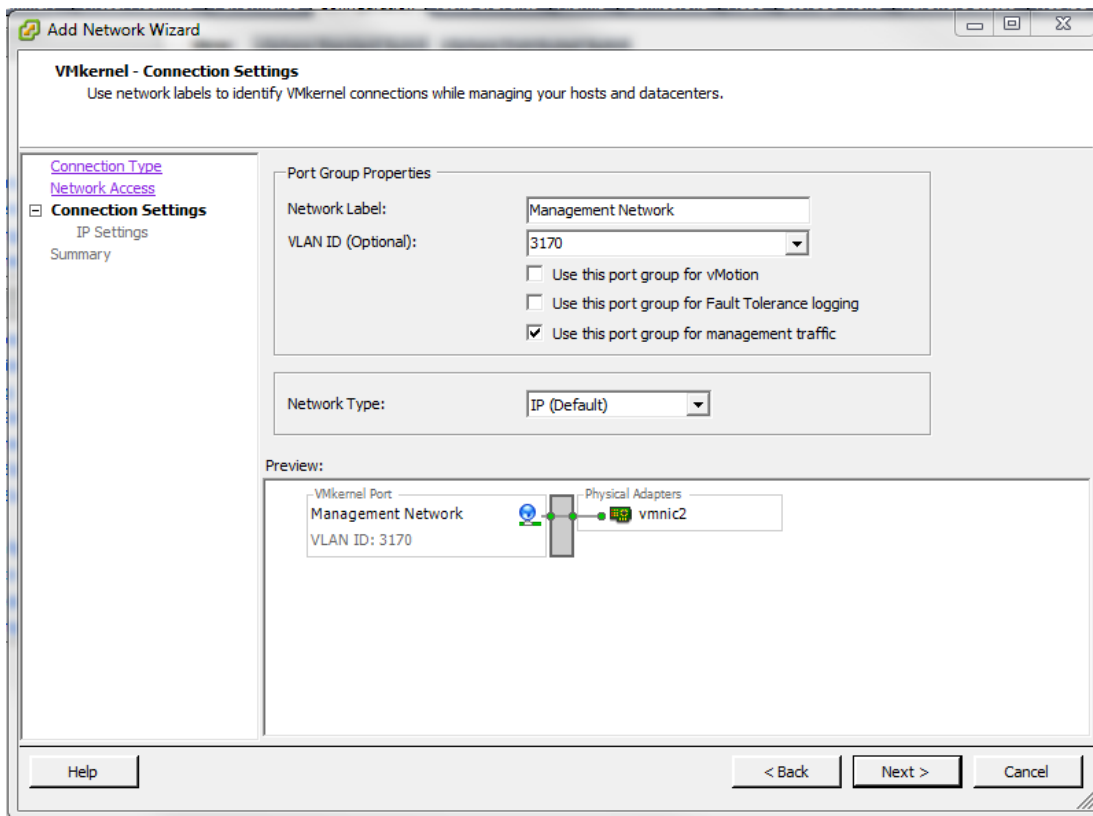
4. With the ESXi Host selected, choose the **Configuration** tab.
5. In the Hardware Pane, choose **Networking**.
6. On the right pane, click **Add Networking**.
7. Choose VMkernel, click **Next**.
8. Make sure **vmnic2** Physical Adapter is selected and click **Next**.
9. Change the Network Label to Management Network.



**Note** This is a mandatory requirement.

10. Enter the <<var\_oob-mgmt\_vlan\_id>> for the VLAN ID.
11. Check the **Use this port group for management traffic** check box and click **Next**.

**Figure 140** VMKernel Connection Settings



12. Enter the appropriate IP address and subnet mask and click **Next**.
13. Click **Finish** to create vSwitch0.
14. Choose Properties for vSwitch0.

15. Choose the **Network Adapters** tab.
16. Click **Add**.
17. Check the checkbox for **vmnic3** and click **Next**.
18. Click **Next**.
19. Click **Finish**.
20. Choose the **Ports** tab.
21. With vSwitch selected, click **Edit**.
22. Choose the **NIC Teaming** tab.
23. Under Load Balancing, choose Route based on source MAC hash.
24. Click **OK**.
25. Click **Close** to close the vSwitch0 properties.

**Figure 141** **vSwitch0 Properties**

#### Networking



26. In the Software pane, choose **DNS and Routing**.
27. Choose Properties on the right corner.
28. Choose the **Routing** tab.
29. Change the Default gateway to the gateway for the OOB-MGMT-VLAN.
30. Click **OK**.
31. If host name resolution is being used for identifying the ESXi Hosts in vCenter, the DNS server host record for this ESXi host should be changed to use the OOB-MGMT-VLAN IP address.
32. Right-click the ESXi host and click **Exit Maintenance Mode**.

## Download Flash Accel Software

1. From the management workstation, navigate to the [NetApp Support site](#).
2. Choose **Flash Accel** and choose the **VMware – Virtual Storage Console** option.
3. Download the 1.2R1 version of Flash Accel.

## Download Fusion-io Driver

1. From the management workstation, navigate to the Fusion-io driver download section in the [VMware support site](#).

**Note**

This procedure requires valid login credentials and Internet access is required on the management workstation.

2. Click **Download**, log in, and click **Download Now** to download the driver.
3. Extract the offline bundle `scsi-iomemory-vsl-5X-3.2.4.1108-offline_bundle-1176140.zip` from the downloaded driver file.


## Install Flash Accel

Before proceeding with the following procedure, NetApp highly recommends verifying that the servers, flash devices, VMware software, and virtual machines meet the minimum resource requirements. These prerequisites can be found in the [NetApp Flash Accel for VMware with Virtual Storage Console document](#).

The Flash Accel Base package will need to be installed on the host on which VMware Virtual Storage Console (VSC) is running.

1. Launch the downloaded Flash Accel binary.
2. Click **Next** in the Welcome screen.
3. Click **Next** to accept the default installation folder.
4. Click **Install**.
5. Accept to restart the VSC service, click **Finish**.

**Figure 142** *NetApp Flash Accel Base for VSC for VMware vSphere Setup Wizard*



The Flash Accel base package is installed on the VSC host. The VSC service is restarted automatically.

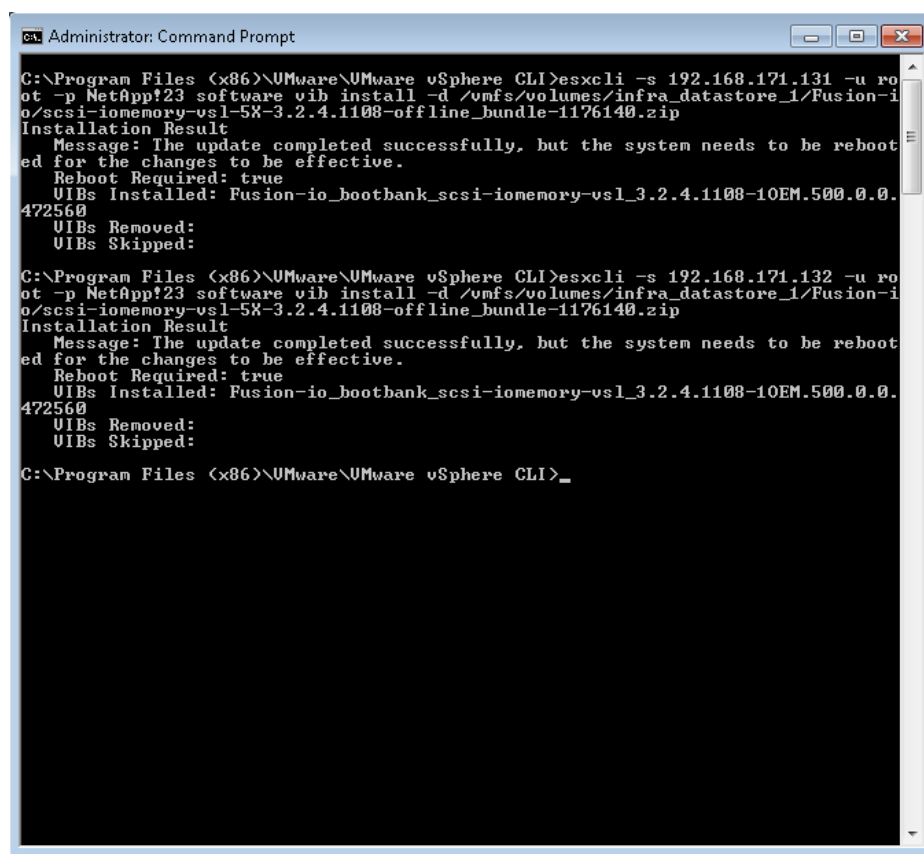
## Install Fusion-io Driver on the ESXi Servers

1. Log in to the vSphere Client.
2. Choose **Datastores** and **Datastore Clusters** in the Inventory pane.



3. Choose the **infra\_datastore\_1** from the drop-down list of datastores.
4. Right-click **infra\_datastore\_1** and choose **Browse Datastore**.
5. Click the third button at the top to create a new folder.
6. Name the folder **Fusion-io** and click **OK**.
7. On the left, select the **Fusion-io** folder.
8. Click the fourth button and choose **Upload File**.
9. Navigate to the **scsi-iomemory-vsl-5X-3.2.4.1108-offline\_bundle-1176140.zip** file and click **Open**.
10. The file should now appear in the **Fusion-io** folder in the datastore.
11. Open the VMware vSphere CLI command prompt.
12. For each ESXi host in the VMware vSphere CLI, run the following command:

```
esxcli -s <Host Server ip> -u root -p <Root Password> software vib install -d
/vmfs/volumes/infra_datastore_1/Fusion-io/scsi-iomemory-vsl-5X-3.2.4.1108-offline_
bundle-1176140.zip
```



```
C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.171.131 -u root -p NetApp!23 software vib install -d /vmfs/volumes/infra_datastore_1/Fusion-io/scsi-iomemory-vsl-5X-3.2.4.1108-offline_bundle-1176140.zip
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  UIBs Installed: Fusion-io_bootbank_scsi-iomemory-vsl_3.2.4.1108-10EM.500.0.0.472560
  UIBs Removed:
  UIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>esxcli -s 192.168.171.132 -u root -p NetApp!23 software vib install -d /vmfs/volumes/infra_datastore_1/Fusion-io/scsi-iomemory-vsl-5X-3.2.4.1108-offline_bundle-1176140.zip
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  UIBs Installed: Fusion-io_bootbank_scsi-iomemory-vsl_3.2.4.1108-10EM.500.0.0.472560
  UIBs Removed:
  UIBs Skipped:

C:\Program Files (x86)\VMware\VMware vSphere CLI>_
```

13. From the vSphere Client, click **Hosts and Clusters** from the Inventory pane.
14. Right-click an ESXi server and choose **Enter Maintenance Mode**.



**Note** This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

15. Right-click the ESXi server in Maintenance Mode and choose **Reboot**.

16. After reboot is completed, right-click the ESXi server and choose **Exit Maintenance Mode**.
17. Repeat steps 14 through 16 for the other ESXi servers.

## Install Host Agent

1. From the vSphere Client, select Inventory > Hosts and Clusters.
2. Right-click an ESXi server and select Enter Maintenance Mode.




**Note** This procedure would require the VMs running on the ESXi server to be migrated to another ESXi server.

3. Navigate to **Home > NetApp** (under Solutions and Applications).
4. Under the Monitoring and Host Configuration section, click **Flash Accel Cache**.

**Figure 143**      *Flash Accel Cache*



5. Click **Add Host** in the Flash Accel Cache window.
6. Select the ESXi server that is in Maintenance Mode and its underlying Fusion-io card. Click **Install**.

**Figure 144**      **Selecting Hosts**

7. Click **Yes** in the confirmation window to install the Host Agent.
8. Upon successful installation, the ESXi server will be listed with the configured Flash Accel Cache details.

**Figure 145**      **Configured Flash Accel Cache**

<b>Monitoring and Host Configuration</b> <a href="#">Overview</a> <a href="#">Storage Details - SAN</a> <a href="#">Storage Details - NAS</a> <b>Flash Accel Cache</b> <a href="#">Data Collection</a> <a href="#">Tools</a> <a href="#">Discovery Status</a>	<b>Flash Accel Cache</b>							<a href="#">Add Host...</a> <a href="#">Add VM...</a> <a href="#">Edit...</a> <a href="#">Remove...</a> <a href="#">Refresh</a>
	Name	IP Address	Agent Version	Operational Status	Perf Status	VM Cache Size	SSD Size	
	Host:icef1-h5.ice.rtp.netapp.com							
	Host:icef1-h5.ice.rtp.netapp.com	192.168.171.131	1.2.0-20912270	Normal	N/A	339.50GB	339.93GB	

9. In the vSphere Client, navigate to **Home > Inventory > Hosts and Clusters**.
10. Right-click the ESXi server and choose **Exit Maintenance Mode**.
11. Repeat steps 5 through 10 for the remaining ESXi servers.

**Figure 146**      **Configured Flash Accel Cache Showing Added ESXi Servers**

<b>Monitoring and Host Configuration</b> <a href="#">Overview</a> <a href="#">Storage Details - SAN</a> <a href="#">Storage Details - NAS</a> <b>Flash Accel Cache</b> <a href="#">Data Collection</a> <a href="#">Tools</a> <a href="#">Discovery Status</a>	<b>Flash Accel Cache</b>							<a href="#">Add Host...</a> <a href="#">Add VM...</a> <a href="#">Edit...</a> <a href="#">Remove...</a> <a href="#">Refresh</a>
	Name	IP Address	Agent Version	Operational Status	Perf Status	VM Cache Size	SSD Size	
	Host:icef1-h5.ice.rtp.netapp.com							
	Host:icef1-h5.ice.rtp.netapp.com	192.168.171.131	1.2.0-20912270	Normal	N/A	339.50GB	339.93GB	

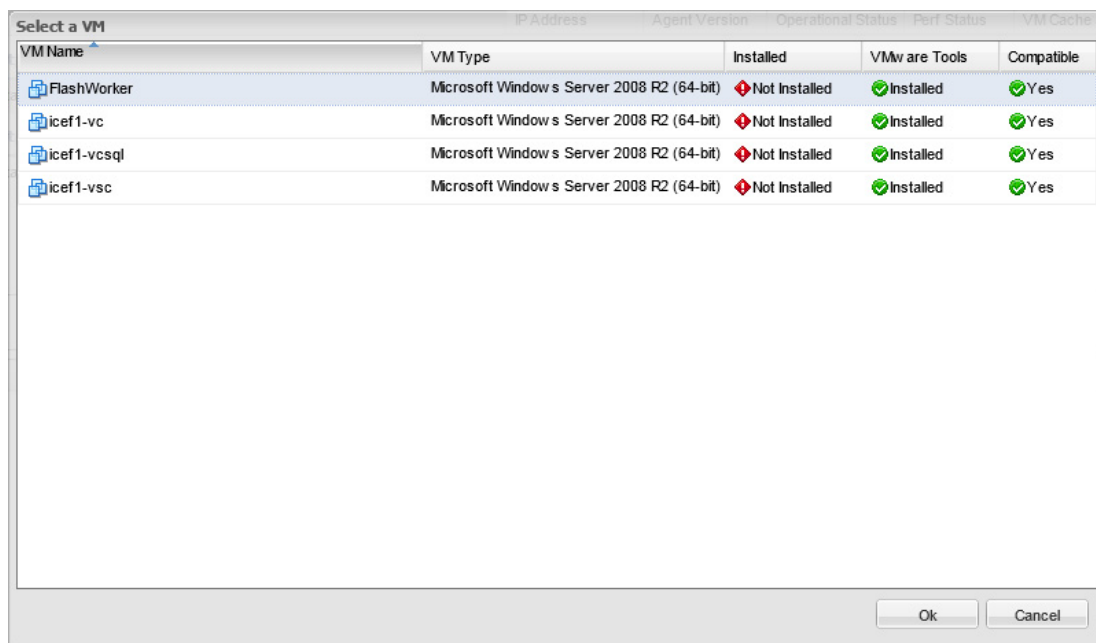
	Host:icef1-h6.ice.rtp.netapp.com							
	Host:icef1-h6.ice.rtp.netapp.com	192.168.171.132	1.2.0-20912270	Normal	N/A	339.50GB	339.93GB	

## Install Guest OS Agent and Allocate Cache

To install the guest OS agent and allocate cache on VMs, complete the following steps.

1. From the Flash Accel Cache window, select the ESXi server on which the VM is running.
2. Click **Add VM** in the Flash Accel Cache window.

**Figure 147**      **Adding VM**



VM Name	VM Type	Installed	VMware Tools	Compatible
FlashWorker	Microsoft Windows Server 2008 R2 (64-bit)	Not Installed	Installed	Yes
icef1-vc	Microsoft Windows Server 2008 R2 (64-bit)	Not Installed	Installed	Yes
icef1-vcsql	Microsoft Windows Server 2008 R2 (64-bit)	Not Installed	Installed	Yes
icef1-vsc	Microsoft Windows Server 2008 R2 (64-bit)	Not Installed	Installed	Yes



**Note** All the VMs that would use Flash Accel Cache should have the up-to-date versions of the VMware tools.

3. Select a compatible VM on which you want to install the guest OS agent. Click **OK**.
4. In the Agent Connection Information section, accept the default values for IP address and management port. Provide the administrator credentials and click the **Verify** to validate.
5. In the Cache Space Settings section, specify the cache size for the VM and the select the datastore where the Flash Accel mapping file should be saved. Check the **Enable VMotion** checkbox and make sure to select a shared datastore to which each of the potential target hosts has access. Click **OK**.

**Figure 148 Adding Flash Accel Guest OS Agent**

**Add Flash Accel Guest OS Agent - FlashWorker**

192.168.175.199 1.2.0.21963869 Normal Enabled 5.00GB

**Agent Connection Information**

Please enter the IP address/port and an administrator user name/password for the selected VM, then click the verify button to check the connection. Note: The VM's firewall rules will be updated to pass the selected port.

IP Address: 192.168.175.199

Management Port: 7010

User name: Administrator

Password: .....

Verify

**Cache Space Settings**

Select the cache size (GB) to use for the VM.

5 (GB)

Select the datastore for the Flash Accel mapping file.

Datastore: Flash\_Accel\_datastore1

Enable VMotion: ☒

**Select Hosts for VMotion enablement**

Select the hosts you want to VMotion to

☒ icef1-h5.ice.rtp.netapp.com ( Version : 5.1.0 )

☒ icef1-h6.ice.rtp.netapp.com ( Version : 5.1.0 )

Ok Cancel

- Click **Yes** to confirm the installation and reboot the VM.

**Figure 149 Flash Accel Cache Showing Added VM**

Monitoring and Host Configuration		Flash Accel Cache						
Overview		Add Host... Add VM... Edit... Remove... Refresh						
Storage Details - SAN								
Storage Details - NAS								
Flash Accel Cache								
Data Collection								
Tools								
Discovery Status								
Flash Accel Cache		Name	IP Address	Agent Version	Operational Status	Perf Status	VM Cache Size	SSD Size
		Hosticef1-h5.ice.rtp.netapp.com						
		Hosticef1-h5.ice.rtp.netapp.com	192.168.171.131	1.2.0-20912270	Normal	N/A	334.50GB	339.93GB
		VM:FlashWorker	192.168.175.199	1.2.0.21963869	Normal	Enabled	5.00GB	N/A
		Hosticef1-h6.ice.rtp.netapp.com						
		Hosticef1-h6.ice.rtp.netapp.com	192.168.171.132	1.2.0-20912270	Normal	N/A	334.50GB	339.93GB
Details								
Windows Disk		Operational Status	Perf Status	Cache Hits	Cache Miss	Read Latency (usec)	Write Latency (usec)	
Disk0 (C:)		Disabled	Disabled	0%	0%	0	0	

- Select the VM on which the Guest OS Agent was installed and cache was allocated. Click **Edit** in the Flash Accel Cache Window.
- Under the VM Cached Devices section, right-click Windows Disk and choose **Enabled** to activate caching.

**Figure 150**      **Editing Flash Accel Guest OS Agent**

**Edit Flash Accel Guest OS Agent - FlashWorker**

**Agent Connection Information**  
Please enter the IP address/port and an administrator user name/password for the selected VM, then click the verify button to check the connection. Note: The VM's firewall rules will be updated to pass the selected port.  
IP Address: 192.168.175.199  
Management Port: 7010  
User name: Administrator  
Password: .....  
Verify

**Cache Space Settings**  
Select the cache size (GB) to use for the VM.  
5 (GB)  
Select the datastore for the Flash Accel mapping file.  
Datastore: Flash\_Accel\_datastore1  
Enable VMotion: ☒

**Select Hosts for VMotion enablement**  
Select the hosts you want to VMotion to  
☒ icef1-h5.ice.rtp.netapp.com ( Version : 5.1.0 )  
☒ icef1-h6.ice.rtp.netapp.com ( Version : 5.1.0 )

**VM Cache Administrative State**  
Select the cache administrative state for the VM. A value of Disabled overrides the settings for the devices listed below.  
☒ Enabled ☐ Disabled

**VM Cached Devices**  
Set the administrative cache state and the performance state by right-clicking each disk. At least one disk must have the administrative cache state set to Enabled for Flash Accel caching to be operational.

Windows Disk	Admin Status	Operational Status	Perf Status	Cache Hits	Cache Miss	Read Latency (usec)	Write Latency (usec)
Disk0 (C:\)	Disabled	Disabled	Disabled	0%	0%	0	0

Right-click context menu for Disk0 (C:\):  
Enabled  
Disabled  
Toggle Perf Stats

Ok Cancel

- (Optional) Right-click the same Windows Disk and choose **Toggle Perf Stats** to collect performance statistics.

**Figure 151**      **Flash Accel Cache Details**

**Monitoring and Host Configuration**

**Flash Accel Cache**

Add Host... Add VM... Edit... Remove... Refresh

Name	IP Address	Agent Version	Operational Status	Perf Status	VM Cache Size	SSD Size
Hosticef1-h5.ice.rtp.netapp.com						
Hosticef1-h5.ice.rtp.netapp.com	192.168.171.131	1.2.0-20912270	Normal	N/A	334.50GB	339.93GB
VMFlashWorker	192.168.175.199	1.2.0.21963869	Normal	Enabled	5.00GB	N/A
Hosticef1-h6.ice.rtp.netapp.com						
Hosticef1-h6.ice.rtp.netapp.com	192.168.171.132	1.2.0-20912270	Normal	N/A	334.50GB	339.93GB

**Details**

Windows Disk	Operational Status	Perf Status	Cache Hits	Cache Miss	Read Latency (usec)	Write Latency (usec)
Disk0 (C:\)	Normal	Enabled	0%	0%	0	0

- Click **OK**.

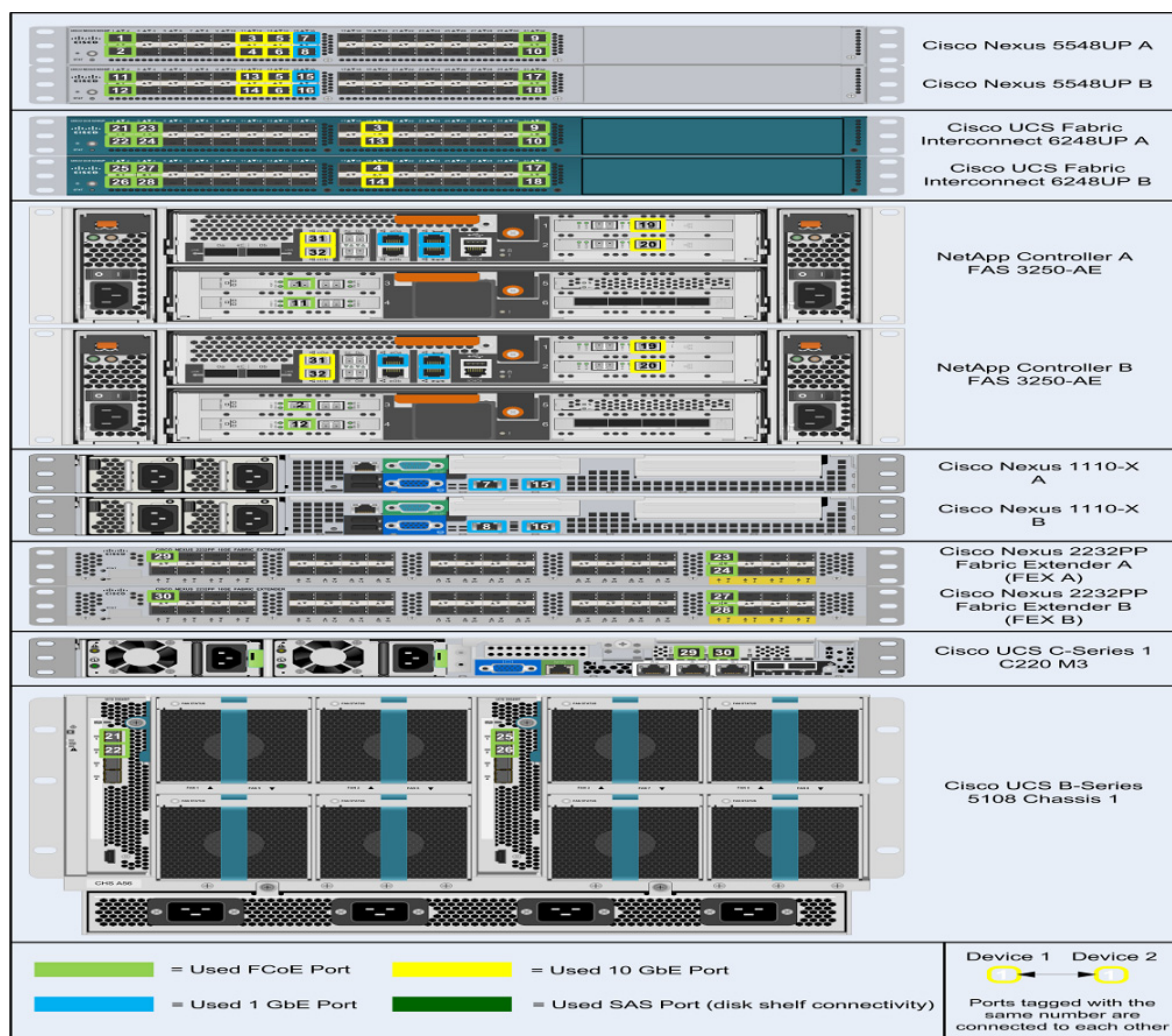
## Clustered Data ONTAP— Switchless Cluster Configuration

Starting with Data ONTAP 8.2, you can optionally configure two-node clusters without the cluster network switches. The cluster network switches would be replaced by direct back-to-back connections between the nodes.

### Storage Cabling for Switchless Cluster Configuration

Figure 3 shows the cabling diagram for a switchless cluster configuration using the FAS3250-AE storage controllers.

**Figure 152** FlexPod Cabling Diagram in Clustered Data ONTAP— Switchless Cluster



The cluster ports on the storage system are directly connected to its partner node storage system.



**Table 33** *NetApp Clustered Data ONTAP Switches Cluster Cabling*

Cluster Node1	Cluster Node2
c0a	c0a
c0b	c0b
e1a	e1a
e2a	e2a

**Note**

The rest of the NetApp storage systems cabling and the FlexPod cabling remain unchanged.

## Switchless Cluster Configuration

To set up a switchless cluster configuration, complete the following steps:

1. Create the cluster on the first node as described in “[Cluster Create in Clustered Data ONTAP](#)” section on page 34, perform steps 1 through 9.
2. Enter the following commands at the cluster shell prompt:

```
set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y
network options switchless-cluster show
```
3. If the output of the above command is “false,” enter the following commands:

```
network options switchless-cluster modify true
network options switchless-cluster show
```
4. The output of the previous command should be “true.” Revert to the admin privileges.

```
set -privilege admin
```
5. Perform steps 10 through 27 in “[Cluster Create in Clustered Data ONTAP](#)” section on page 34.
6. Join the second node to the cluster created as described in “[Cluster Join in Clustered Data ONTAP](#)” section on page 37.
7. Log in to the cluster as described in “[Log in to the Cluster](#)” section on page 40.
8. Configure the NetApp storage systems as described from “[Zero All Spare Disks](#)” section on page 40 onwards.

## Cisco UCS Central – Multi Domain Management

Cisco UCS Central software manages multiple, globally distributed Cisco UCS domains with thousands of servers from a single pane.

This section provides a detailed overview of UCS Central setup in standalone mode.

The installation and upgrade guide is available at:

[http://www.cisco.com/en/US/docs/unified\\_computing/ucs/ucs-central/install-upgrade/1.1/b\\_UCSC\\_Installation\\_and\\_Upgrade\\_Guide\\_11.html](http://www.cisco.com/en/US/docs/unified_computing/ucs/ucs-central/install-upgrade/1.1/b_UCSC_Installation_and_Upgrade_Guide_11.html)



## Obtain the UCS Central Software

1. Navigate to the [Cisco UCS Central Download](#) page.
2. Download the OVA file ucs-central.1.1.1a.ova.

## Install the UCS Central Software

1. Using the vSphere Client, log in to the vCenter Server as FlexPod admin user.
2. In the Main menu, click **File > Deploy OVF Template**.
3. Browse to the OVA file that was downloaded. Click **Next**.
4. Click **Next**.
5. Modify the default name if desired and select the Inventory Location. Click **Next**.
6. Select a cluster/server on which you want to host the UCS Central virtual machine.
7. Select the datastore in which the virtual machine files will be stored. Click **Next**.
8. Click **Next**.
9. Check the checkbox to power on the VM after deployment.
10. Click **Finish**.



**Note** Do not proceed until the virtual machine has finished booting.

11. Open a console window to the UCS Central virtual machine.
12. Answer the following prompts in the console window.

```
Setup new configuration or restore full-state configuration from
backup[setup/restore] - setup
Enter the UCS Central VM eth0 IPv4 Address : <<var_ucs_central_ip>>
Enter the UCS Central VM eth0 IPv4 Netmask : <<var_ucs_central_netmask>>
Enter the VM IPv4 Default Gateway : <<var_ucs_central_gateway>>
Is this VM part of a cluster (select 'no' for standalone) (yes/no)? no
Enter the UCS Central VM Hostname : <<var_ucs_central_hostname>>
Enter the DNS Server IPv4 Address : <<var_nameserver_ip>>
Enter the Default Domain Name : <<var_dns_domain_name>>
Use a Shared Storage Device for Database (yes/no)? no
Enforce Strong Password (yes/no)? yes
Enter the admin Password : <<var_password>>
Confirm admin Password : <<var_password>>
Enter the Shared Secret : enter the shared secret (or password) that you want
to use to register one or more Cisco UCS domains with Cisco UCS Central
Confirm Shared Secret : re-enter the Shared Secret
Do you want Statistics collection [yes / no]? yes
Enter the Statistics DB Type [D=Default (internal Pstgres db) / P=Postgres /
O=Oracle] : D
Proceed with this configuration? Please confirm [yes/no] - yes
```



**Note** If you wish to modify/answer the prompts again, enter no in the above prompt.

13. After confirming that you want to proceed with the configuration, the network interface reinitializes with your settings and Cisco UCS Central can be accessed using the IP address.

## Access UCS Central GUI

1. Using a web browser, navigate to the <<var\_ucs\_central\_hostname>> using `https://<var_ucs_central_ip>>`.
2. Log in with the user name as admin and the admin password.
3. Click the **Operations Management** tab, expand **Domain Groups > Domain Group root**.
4. Choose **Operational Policies**.
5. Click **Time Zone** in the right pane, and select the desired time zone.
6. Click **Add NTP Server**.
7. Provide the NTP Server IP Address <<var\_global\_ntp\_server\_ip>> and click **OK**.
8. Click **Save**.

## Add UCS Managers to UCS Central

UCS Managers are added into the UCS Central by logging into the UCS Manager and registering the UCS Manager with UCS Central.

To add UCS Manager to UCS Central, complete the following steps:

1. Log in to the Cisco UCS Manager.
2. In the navigation pane, click the **Admin** tab.
3. In the Admin tab expand the All folder, choose **Communication Management > UCS Central**.
4. In the **UCS Central** tab, click **Register With UCS Central** in the Actions section.
5. Enter the host name or IP address of the UCS Central.
6. Enter the Shared Secret. Click **OK**.
7. Click **Accept** to terminate any open GUI sessions to the UCS Manager.
8. Check the checkbox to view the Navigator for the UCS Central. Click **OK**.
9. Verify the Registration Status.

## Expand Two-Node Cluster to Four-Node Cluster FlexPod

1. Cable the two new storage cluster nodes in a similar manner as described in [Table 10](#) and [Table 11](#).
2. Log in to the cluster interface and disable cluster HA by typing:
 

```
cluster ha modify -configured false
```
3. Build the two new storage cluster nodes using the procedure described in [“Storage Configuration” section on page 25](#).
  - a. In [“Clustered Data ONTAP 8.2” section on page 30](#) make sure Data ONTAP 8.2 is installed, initialize the disks, and assign disks for the two new controllers.
  - b. In [“Cluster Create in Clustered Data ONTAP” section on page 34](#) use the Node 2 instructions to join Nodes 3 and 4 to the cluster.
  - c. Log in to 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 41 no action is necessary.
- h. In “Failover Groups Node Management in Clustered Data ONTAP” section on page 41 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 “Aggregates in Clustered Data ONTAP” section on page 42 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 44 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 45 add VLAN interfaces for the NFS VLAN on Nodes 3 and 4.
- p. In “Jumbo Frames in Clustered Data ONTAP” section on page 45 modify the newly added VLAN interfaces for jumbo frames.
- q. In “NTP in Clustered Data ONTAP” section on page 45 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” section on page 46
  - “SNMPv3 in Clustered Data ONTAP” section on page 46
- 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 (Storage Virtual Machine)” section on page 47 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 48 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 48](#) 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” section on page 50](#)
  - [“FlexVol in Clustered Data ONTAP” section on page 50](#)
  - [“LUN in Clustered Data ONTAP” section on page 51](#)
  - [“Deduplication in Clustered Data ONTAP” section on page 51](#)
- z. In [“Failover Groups NAS in Clustered Data ONTAP” section on page 51](#) add Node 3 and 4 NFS ports to the NFS failover group.
- aa. In [“NFS LIF in Clustered Data ONTAP” section on page 51](#) create LIF nfs\_lif03 on Node 3 and nfs\_lif04 on Node 4.
- ab. In [“FCP LIF in Clustered Data ONTAP” section on page 52](#) 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 52](#).
- 4. Using the procedures in described in [“Storage Networking” section on page 106](#) 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.
- 5. 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 [“Physical Infrastructure” section on page 17](#).
2. Build the storage cluster according to [“Storage Configuration” section on page 25](#) 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. See [“Server Configuration” section on page 52](#).

5. In the Cisco Nexus 5548s, 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. See [“Storage Networking” section on page 106](#)
6. Create igroups in the cluster and map the two new boot LUNs to the igroups.
7. Install and configure ESXi on the two new servers. See [“VMware vSphere 5.1Update1 Setup” section on page 118](#) 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. See [“OnCommand Unified Manager 6.0” section on page 171](#).
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 5548 Example Configurations

### Cisco Nexus 5548 A

```

!Command: show running-config
!Time: Wed Sep 11 14:32:27 2013

version 6.0(2)N2(2)
feature fcoe
switchname NX5548UP-A
feature npiv
no feature telnet
cfs eth distribute
feature lacp
feature vpc
feature lldp
username admin password 5 $1$Atf.qEX5$sLcEk1MaswuwjD2DsQwtq1 role network-admin

banner motd #Nexus 5000 Switch
#

ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-fcoe
    pause no-drop

```

```

        mtu 2158
    class type network-qos class-default
        mtu 9216
        multicast-optimize
system qos
    service-policy type qos input fcoe-default-in-policy
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type network-qos jumbo
snmp-server user admin network-admin auth md5 0x83cfe580474af369c86705270411d7eb priv
0x83cfe580474af369c86705270411d7eb localiz
edkey
ntp server 10.72.242.70 use-vrf management
vrf context management
    ip route 0.0.0.0/0 10.238.162.1
vlan 1
vlan 2
    name Native-VLAN
vlan 101
    fcoe vsan 101
    name FCoE_Fabric_A
vlan 3170
    name OOB-MGMT-VLAN
vlan 3172
    name NFS-VLAN
vlan 3173
    name vMotion-VLAN
vlan 3174
    name VM-Traffic-VLAN
vlan 3175
    name IB-MGMT-VLAN
vlan 3176
    name Packet-Control-VLAN
spanning-tree port type edge bpduguard default
spanning-tree port type network default
port-channel load-balance ethernet source-dest-port
vpc domain 7
    role priority 10
    peer-keepalive destination 10.238.162.249 source 10.238.162.248
    auto-recovery
port-profile default max-ports 512
vsan database
    vsan 101 name "Fabric_A"
device-alias database
    device-alias name fcp_lif01a pwwn 20:01:00:a0:98:3f:53:d4
    device-alias name fcp_lif02a pwwn 20:03:00:a0:98:3f:53:d4
    device-alias name VM-Host-Infra-01_A pwwn 20:00:00:25:b5:00:0a:0f
    device-alias name VM-Host-Infra-02_A pwwn 20:00:00:25:b5:00:0a:1f

device-alias commit

fcdomain fcid database
    vsan 101 wwn 50:0a:09:81:88:ad:35:09 fcid 0x820000 dynamic
    vsan 101 wwn 20:01:00:a0:98:3f:53:d4 fcid 0x820001 dynamic
    !
    [fcp_lif01a]
    vsan 101 wwn 50:0a:09:81:88:6c:0a:8c fcid 0x820020 dynamic
    vsan 101 wwn 20:03:00:a0:98:3f:53:d4 fcid 0x820021 dynamic
    !
    [fcp_lif02a]
    vsan 101 wwn 22:c8:54:7f:ee:aa:aa:7f fcid 0x820040 dynamic
    vsan 101 wwn 20:00:00:25:b5:00:0a:0f fcid 0x820041 dynamic
    !
    [VM-Host-Infra-01_A]
    vsan 101 wwn 20:00:00:25:b5:00:0a:1f fcid 0x820042 dynamic
    !
    [VM-Host-Infra-02_A]

```

```

interface port-channel1
  description "GLC-T connected to UPLINK INFRA"
  switchport mode trunk
  switchport trunk allowed vlan 3170,3175
  spanning-tree port type normal
  vpc 1

interface port-channel10
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3176
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FAS3250-Cluster-01
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  spanning-tree port type edge trunk
  vpc 11

interface port-channel12
  description FAS3250-Cluster-02
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  spanning-tree port type edge trunk
  vpc 12

interface port-channel13
  description FI6248-A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  spanning-tree port type edge trunk
  vpc 13

interface port-channel14
  description FI6248-B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  spanning-tree port type edge trunk
  vpc 14

interface port-channel15
  description FI6248-A:FCoE
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101
  spanning-tree port type edge trunk

interface vfc11
  bind interface Ethernet1/1
  switchport trunk allowed vsan 101
  switchport description FAS3250-Cluster-01:3a
  no shutdown

interface vfc12
  bind interface Ethernet1/2
  switchport trunk allowed vsan 101

```



```

switchport description FAS3250-Cluster-02:3a
no shutdown

interface vfc15
  bind interface port-channel15
  switchport trunk allowed vsan 101
  switchport description FI6248-A:FCoE
  no shutdown
vsan database
  vsan 101 interface vfc11
  vsan 101 interface vfc12
  vsan 101 interface vfc15

interface Ethernet1/1
  description FAS3250-Cluster-01:e3a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  channel-group 11 mode active

interface Ethernet1/2
  description FAS3250-Cluster-02:e3a
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 101,3172
  channel-group 12 mode active

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8
  switchport mode trunk
  switchport trunk allowed vlan 3170,3175
  speed 1000
  channel-group 1 mode active

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11
  description FI6248-A:1/19
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  channel-group 13 mode active

interface Ethernet1/12
  description FI6248-B:1/19
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3170,3172-3175
  channel-group 14 mode active

interface Ethernet1/13
  description NX5548UP-B:1/13
  switchport mode trunk

```

```
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3176
channel-group 10 mode active

interface Ethernet1/14
description NX5548UP-B:1/14
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3176
channel-group 10 mode active

interface Ethernet1/15
description Nexus1110-X-A:Eth1
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000
vpc orphan-port suspend

interface Ethernet1/16
description Nexus1110-X-B:Eth1
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000
vpc orphan-port suspend

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31
description FI6248-A:1/31
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 101
channel-group 15 mode active

interface Ethernet1/32
description FI6248-A:1/32
```

```

switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 101
channel-group 15 mode active

interface mgmt0
 ip address 10.238.162.248/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.6.0.2.N1.2a.bin
boot system bootflash:/n5000-uk9.6.0.2.N1.2a.bin
!Full Zone Database Section for vsan 101
zone name VM-Host-Infra-01_A vsan 101
 member pwnn 20:00:00:25:b5:00:0a:0f
!           [VM-Host-Infra-01_A]
 member pwnn 20:01:00:a0:98:3f:53:d4
!           [fcp_lif01a]
 member pwnn 20:03:00:a0:98:3f:53:d4
!           [fcp_lif02a]

zone name VM-Host-Infra-02_A vsan 101
 member pwnn 20:00:00:25:b5:00:0a:1f
!           [VM-Host-Infra-02_A]
 member pwnn 20:01:00:a0:98:3f:53:d4
!           [fcp_lif01a]
 member pwnn 20:03:00:a0:98:3f:53:d4
!           [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

```

## Cisco Nexus 5548 B

```

!Command: show running-config
!Time: Wed Sep 11 14:44:09 2013

version 6.0(2)N2(2)
feature fcoe
switchname NX5548UP-B
feature npiv
no feature telnet
cfs eth distribute
feature lacp
feature vpc
feature lldp
username admin password 5 $1$aVldNDJH$igwxQD3t9h9ojPB0VisZ8. role network-admin

banner motd #Nexus 5000 Switch
#

ip domain-lookup
class-map type qos class-fcoe
class-map type queuing class-fcoe
 match qos-group 1
class-map type queuing class-all-flood
 match qos-group 2
class-map type queuing class-ip-multicast
 match qos-group 2
class-map type network-qos class-fcoe

```

```

    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
policy-map type network-qos jumbo
    class type network-qos class-fcoe
        pause no-drop
        mtu 2158
    class type network-qos class-default
        mtu 9216
        multicast-optimize
system qos
    service-policy type qos input fcoe-default-in-policy
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type network-qos jumbo
snmp-server user admin network-admin auth md5 0xd1b0309fc2a032aef4bc1d2566bf7c23 priv
0xd1b0309fc2a032aef4bc1d2566bf7c23 localiz
edkey
ntp server 10.72.242.70 use-vrf management
vrf context management
    ip route 0.0.0.0/0 10.238.162.1
vlan 1
vlan 2
    name Native-VLAN
vlan 102
    fcoe vsan 102
    name FCoE_Fabric_B
vlan 3170
    name OOB-MGMT-VLAN
vlan 3172
    name NFS-VLAN

vlan 3173
    name vMotion-VLAN
vlan 3174
    name VM-Traffic-VLAN
vlan 3175
    name IB-MGMT-VLAN
vlan 3176
    name Packet-Control-VLAN
spanning-tree port type edge bpduguard default
spanning-tree port type network default
port-channel load-balance ethernet source-dest-port
vpc domain 7
    role priority 20
    peer-keepalive destination 10.238.162.248 source 10.238.162.249
    auto-recovery
port-profile default max-ports 512
vsan database
    vsan 102 name "Fabric_B"
device-alias database
    device-alias name fcp_lif01b pwwn 20:02:00:a0:98:3f:53:d4
    device-alias name fcp_lif02b pwwn 20:04:00:a0:98:3f:53:d4
    device-alias name VM-Host-Infra-01_B pwwn 20:00:00:25:b5:00:0b:0f
    device-alias name VM-Host-Infra-02_B pwwn 20:00:00:25:b5:00:0b:1f

device-alias commit

fcdomain fcid database
    vsan 102 wwn 50:0a:09:83:88:ad:35:09 fcid 0xbd0000 dynamic
    vsan 102 wwn 20:02:00:a0:98:3f:53:d4 fcid 0xbd0001 dynamic
!
```

```

vsan 102 wwn 50:0a:09:83:88:6c:0a:8c fcid 0xbd0020 dynamic
vsan 102 wwn 20:04:00:a0:98:3f:53:d4 fcid 0xbd0021 dynamic
!
[fcplif02b]
vsan 102 wwn 22:c9:54:7f:ee:aa:ac:bf fcid 0xbd0040 dynamic
vsan 102 wwn 20:00:00:25:b5:00:0b:0f fcid 0xbd0041 dynamic
!
[VM-Host-Infra-01_B]
vsan 102 wwn 20:00:00:25:b5:00:0b:1f fcid 0xbd0042 dynamic
!
[VM-Host-Infra-02_B]

interface port-channel1
description "GLC-T connected to UPLINK INFRA"
switchport mode trunk
switchport trunk allowed vlan 3170,3175
spanning-tree port type normal
vpc 1

interface port-channel10
description vPC peer-link
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3176
spanning-tree port type network
vpc peer-link

interface port-channel11
description FAS3250-Cluster-01
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 102,3172
spanning-tree port type edge trunk
vpc 11

interface port-channel12
description FAS3250-Cluster-02
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 102,3172
spanning-tree port type edge trunk
vpc 12

interface port-channel13
description FI6248-A
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3175
spanning-tree port type edge trunk
vpc 13

interface port-channel14
description FI6248-B
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3175
spanning-tree port type edge trunk
vpc 14

interface port-channel15
description FI6248-B:FCoE
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 102
spanning-tree port type edge trunk

```

```

interface vfc11
    bind interface Ethernet1/1
    switchport trunk allowed vsan 102
    switchport description FAS3250-Cluster-01:4a
    no shutdown

interface vfc12
    bind interface Ethernet1/2
    switchport trunk allowed vsan 102
    switchport description FAS3250-Cluster-02:4a
    no shutdown

interface vfc15
    bind interface port-channel15
    switchport trunk allowed vsan 102
    switchport description NX5548UP-B:FCoE
    no shutdown
vsan database
    vsan 102 interface vfc11
    vsan 102 interface vfc12
    vsan 102 interface vfc15

interface Ethernet1/1
    description FAS3250-Cluster-01:e4a
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 102,3172
    channel-group 11 mode active

interface Ethernet1/2
    description FAS3250-Cluster-02:e4a
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 102,3172
    channel-group 12 mode active

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8
    switchport mode trunk
    switchport trunk allowed vlan 3170,3175
    speed 1000
    channel-group 1 mode active

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11
    description FI6248-A:1/20
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 3170,3172-3175
    channel-group 13 mode active

interface Ethernet1/12

```

```

description FI6248-B:1/20
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3175
channel-group 14 mode active

interface Ethernet1/13
description NX5548UP-A:1/13
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3176
channel-group 10 mode active

interface Ethernet1/14
description NX5548UP-A:1/14
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 3170,3172-3176
channel-group 10 mode active

interface Ethernet1/15
description Nexus1110-X-A:Eth2
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000
vpc orphan-port suspend

interface Ethernet1/16
description Nexus1110-X-B:Eth2
switchport mode trunk
switchport trunk allowed vlan 3175-3176
spanning-tree port type edge trunk
speed 1000
vpc orphan-port suspend

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

```

```

interface Ethernet1/31
  description FI6248-B:1/31
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102
  channel-group 15 mode active

interface Ethernet1/32
  description FI6248-B:1/32
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 102
  channel-group 15 mode active

interface mgmt0
  ip address 10.238.162.249/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.6.0.2.N1.2a.bin
boot system bootflash:/n5000-uk9.6.0.2.N1.2a.bin
!Full Zone Database Section for vsan 102
zone name VM-Host-Infra-01_B vsan 102
  member pwnn 20:00:00:25:b5:00:0b:0f
!      [VM-Host-Infra-01_B]
  member pwnn 20:02:00:a0:98:3f:53:d4
!      [fcp_lif01b]
  member pwnn 20:04:00:a0:98:3f:53:d4
!      [fcp_lif02b]

zone name VM-Host-Infra-02_B vsan 102
  member pwnn 20:02:00:a0:98:3f:53:d4
!      [fcp_lif01b]
  member pwnn 20:04:00:a0:98:3f:53:d4
!      [fcp_lif02b]
  member pwnn 20:00:00:25:b5:00:0b:1f
!      [VM-Host-Infra-02_B]

zoneset name FlexPod vsan 102
  member VM-Host-Infra-01_B
  member VM-Host-Infra-02_B

zoneset activate name FlexPod vsan 102

```