

FlexPod Datacenter with VMware vSphere 6.0

Deployment Guide for FlexPod with VMware vSphere 6.0 and NetApp AFF 8000 Series and Cisco Nexus 9000 Series Switches for Top of Rack

Last Updated: November 11, 2015



About Cisco Validated Designs

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

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

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

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

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

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

© 2015 Cisco Systems, Inc. All rights reserved.

Table of Contents

About Cisco Validated Designs	2
Executive Summary	8
Solution Overview.....	9
Introduction	9
Audience.....	9
Purpose of this Document.....	9
What's New?	9
Solution Design.....	10
Architecture.....	10
Physical Topology.....	10
Deployment Hardware and Software	12
Software Revisions	12
Configuration Guidelines.....	12
Physical Infrastructure.....	17
FlexPod Cabling	17
Network Switch Configuration.....	23
Physical Connectivity	23
FlexPod Cisco Nexus Base	23
Set Up Initial Configuration	23
FlexPod Cisco Nexus Switch Configuration.....	26
Enable Licenses.....	26
Set Global Configurations	26
Create VLANs.....	26
Add NTP Distribution Interface.....	27
Add Individual Port Descriptions for Troubleshooting.....	28
Create Port Channels.....	30
Configure Port Channel Parameters.....	32
Configure Virtual Port Channels	34
Uplink into Existing Network Infrastructure	36
Storage Configuration.....	37
Controller AFF80XX Series	37
NetApp Hardware Universe	37
Controllers.....	37

Disk Shelves	37
Clustered Data ONTAP 8.3.1	38
Complete the Configuration Worksheet	38
Configure Clustered Data ONTAP Nodes.....	38
Log In to the Cluster	53
Zero All Spare Disks.....	53
Set Onboard UTA2 Ports Personality	53
Set Auto-Revert on Cluster Management	54
Set Up Management Broadcast Domain	54
Set Up Service Processor Network Interface	54
Create Aggregates	55
Verify Storage Failover.....	56
Disable Flow Control on UTA2 Ports.....	57
Disable Unused FcoE Ports.....	57
Configure NTP	57
Configure SNMP.....	58
Configure AutoSupport.....	59
Enable Cisco Discovery Protocol	59
Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP	59
Create Interface Groups	60
Create VLANs.....	60
Create Storage Virtual Machine	61
Create Load-Sharing Mirrors of SVM Root Volume	61
Create iSCSI Service	62
Configure HTTPS Access	62
Configure NFSv3	63
Create FlexVol Volumes.....	64
Create Boot LUNs.....	65
Schedule Deduplication.....	65
Create iSCSI LIFs.....	65
Create NFS LIF	66
Add Infrastructure SVM Administrator.....	67
Server Configuration.....	68
Cisco UCS Base Configuration.....	68
Perform Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments	68

Cisco UCS Setup	69
Log in to Cisco UCS Manager	69
Upgrade Cisco UCS Manager Software to Version 2.2(5b)	69
Anonymous Reporting	70
Add Block of IP Addresses for Inband KVM Access	70
Synchronize Cisco UCS to NTP	71
Edit Chassis Discovery Policy	71
Enable Server and Uplink Ports	71
Acknowledge Cisco UCS Chassis and FEX	72
Create Uplink Port Channels to Cisco Nexus Switches	74
Create MAC Address Pools	76
Create IQN Pools for iSCSI Boot	77
Create IP Pools for iSCSI Boot	78
Create UUID Suffix Pool	80
Create Server Pool	81
Create VLANs	81
Create VLAN Group and Assign Inband Profile	85
Create Host Firmware Package	87
Set Jumbo Frames in Cisco UCS Fabric	87
Create Local Disk Configuration Policy (Optional)	88
Create Network Control Policy for Cisco Discovery Protocol	89
Create Power Control Policy	90
Create Server Pool Qualification Policy (Optional)	91
Create Server BIOS Policy	92
Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts	93
Update the Default Maintenance Policy	94
Create vNIC Templates	95
Create Boot Policies	101
Create Service Profile Template	102
Create Service Profiles	118
Add More Servers to FlexPod Unit	119
Gather Necessary Information	119
Storage Configuration – iSCSI Boot	120
Clustered Data ONTAP iSCSI Boot Storage Setup	120
Create igroups	120

Map Boot LUNs to igroups.....	120
VMware vSphere 6.0 Setup	121
VMware ESXi 6.0	121
Download Cisco Custom Image for ESXi 6.0.....	121
Log in to Cisco UCS 6200 Fabric Interconnect.....	121
Set Up VMware ESXi Installation.....	122
Install ESXi.....	122
Set Up Management Networking for ESXi Hosts	123
Download VMware vSphere Client.....	125
Download VMware vSphere CLI 6.0	126
Log in to VMware ESXi Hosts by Using VMware vSphere Client.....	126
Set Up VMkernel Ports and Virtual Switch.....	127
Setup iSCSI Multipathing	135
Install VMware Drivers for the Cisco Virtual Interface Card (VIC).....	136
Mount Required Datastores	138
Configure NTP on ESXi Hosts	142
Move VM Swap File Location.....	143
VMware vCenter 6.0.....	144
Install the Client Integration Plug-in	144
Building the VMware vCenter Server Appliance	145
Setting Up VMware vCenter Server	154
ESXi Dump Collector Setup for iSCSI-Booted Hosts.....	164
Cisco UCS Virtual Media (vMedia) Policy for VMware ESXi Installation.....	165
Storage Controller Setup for vMedia Policy.....	165
FlexPod Cisco Nexus 1110-X and 1000V vSphere	170
Configure CIMC Interface on Both Cisco Nexus 1110-Xs	170
Configure Serial over LAN for Both Cisco Nexus 1110-Xs.....	171
Configure Cisco Nexus 1110-X Virtual Appliances	172
Set Up the Primary Cisco Nexus 1000V VSM.....	174
Set Up the Secondary Cisco Nexus 1000V VSM.....	175
Install Cisco Virtual Switch Update Manager	176
Register the Cisco Nexus 1000V in VMware vCenter.....	179
Perform Base Configuration of the Primary VSM.....	179
Add VMware ESXi Hosts to Cisco Nexus 1000V	183
Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V	185

Cisco Nexus 1000V vTracker.....	187
FlexPod Management Tools Setup.....	189
NetApp Virtual Storage Console (VSC) 6.1 Deployment Procedure.....	189
VSC 6.1 Pre-installation Considerations	189
Install VSC 6.1	189
Register VSC with vCenter Server	191
Discover and Add Storage Resources	192
Optimal Storage Settings for ESXi Hosts.....	193
VSC 6.1 Backup and Recovery	194
OnCommand Unified Manager 6.2P1.....	198
OnCommand Unified Manager OVF Deployment.....	198
OnCommand Unified Manager Basic Setup	204
OnCommand Performance Manager 2.0.....	210
OnCommand Performance Manager OVF Deployment.....	210
OnCommand Performance Manager Basic Setup	215
Link OnCommand Performance Manager to OnCommand Unified Manager.....	217
NetApp NFS Plug-In 1.1.0 for VMware VAAI.....	219
Enable VMware vStorage for NFS in Clustered Data ONTAP.....	219
Install NetApp NFS Plug-In for VMware VAAI.....	220
About the Authors.....	223
Acknowledgements	223



Executive Summary

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

This document describes the Cisco and NetApp® FlexPod Datacenter with VMware vSphere 6. FlexPod Datacenter with VMware vSphere 6 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, and NetApp AFF.

Solution Overview

Introduction

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

Audience

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

Purpose of this Document

This document provides a step by step configuration and implementation guide for the FlexPod Datacenter with NetApp AFF and Cisco Nexus 9000 solution. For the design decisions and technology discussion of the solution, please refer to FlexPod Datacenter with NetApp All Flash FAS, Cisco Nexus 9000 and VMware vSphere 6 Design Guide:

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

What's New?

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

- Validation of Cisco Nexus 9000 with a NetApp All-Flash FAS storage array
- Support for the Cisco UCS 2.2(5) release and Cisco UCS B200-M4 servers
- Support for the latest release of NetApp Data ONTAP® 8.3.1
- An IP-based storage design supporting both NAS datastores and iSCSI based SAN LUNs
- Cisco Nexus 1000v vTracker technology
- Cisco UCS Inband KVM Access
- Cisco UCS vMedia client for vSphere Installation
- Cisco UCS Firmware Auto Sync Server policy

Solution Design

Architecture

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

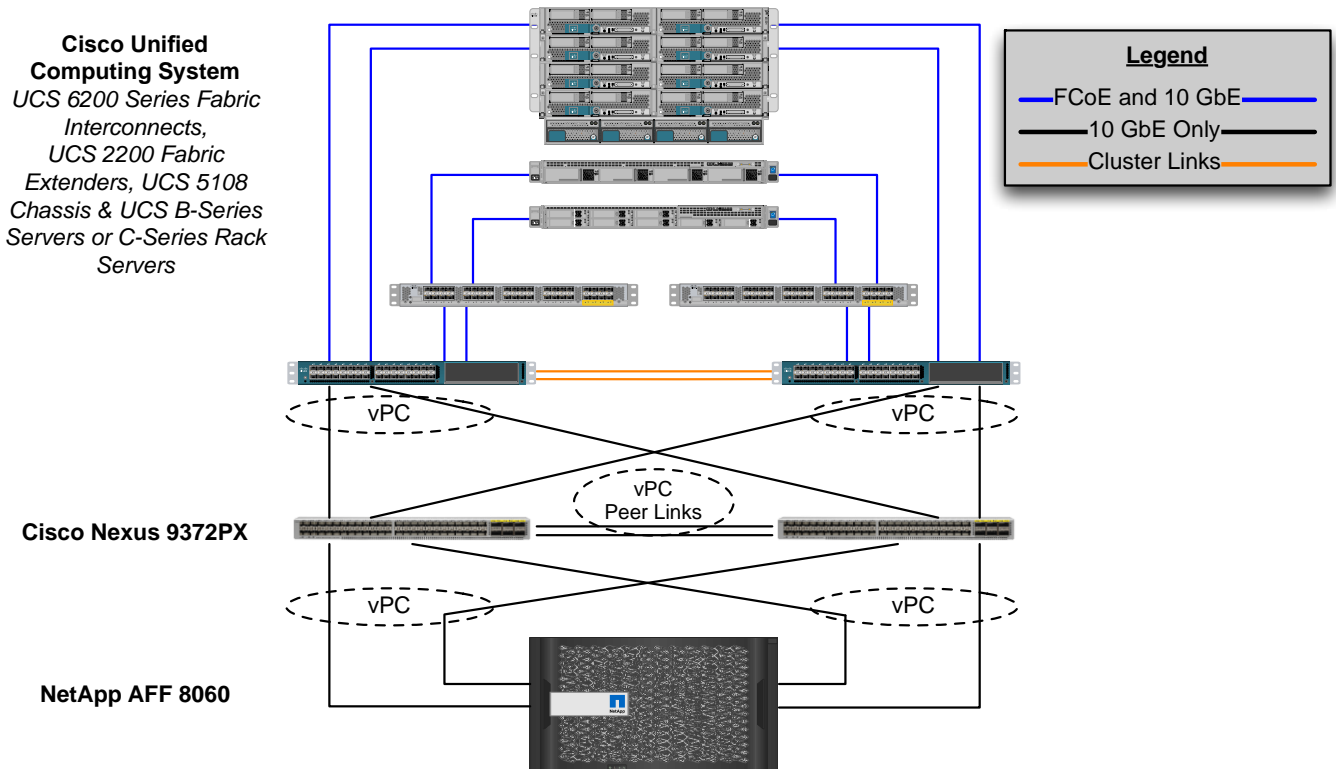
One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The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

0 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with IP-based storage. This design uses the Cisco Nexus 9000, Cisco Nexus 2232PP FEX, and Cisco UCS C-Series and B-Series servers and the NetApp AFF family of storage controllers connected in a highly available modular design. This infrastructure is deployed to provide iSCSI-booted hosts with file-level and block-level access to shared storage. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Physical Topology

0 illustrates the physical architecture.

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



The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches
- Two Cisco UCS 6248UP fabric interconnects
- One NetApp AFF8060 (HA pair) running clustered Data ONTAP with Disk shelves and Solid State Drives (SSD)

For server virtualization, the deployment includes VMware vSphere 6. Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps for deploying the base architecture, as shown in 0. These procedures cover everything from physical cabling to network, compute and storage device configurations.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6200 Series, UCS B-200 M4, UCS C-220 M4	2.2(5b)	Includes the Cisco UCS-IOM 2208XP, Cisco UCS Manager, UCS VIC 1240 and UCS VIC 1340
	Cisco eNIC	2.1.2.71	
	Cisco fNIC	1.6.0.17a	
Network	Cisco Nexus 9000 NX-OS	7.0(3)I1(1a)	
	Cisco Nexus 1000V	5.2(1)SV3(1.5a)	
	Cisco Nexus 1110-X	5.2(1)SP1(7.3)	
Storage	NetApp AFF 8060	Data ONTAP 8.3.1	
Software	VMware vSphere ESXi	6.0	
	VMware vCenter	6.0	
	NetApp Virtual Storage Console (VSC)	6.1	
	OnCommand Performance Manager	2.0	

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these 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 3 lists the virtual machines (VMs) necessary for deployment as outlined in this guide. Table 2 describe the VLANs necessary for deployment as outlined in this guide.

Table 2 Necessary VLANs

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

Table 3 lists the virtual machines (VMs) necessary for deployment as outlined in this document.

Table 3 Virtual Machines

Virtual Machine Description	Host Name
Active Directory	
vCenter Server	
NetApp Virtual Storage Console (VSC)	
NetApp OnCommand Unified Manager	
OnCommand Performance Manager	

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

Table 4 Configuration Variables

Variable	Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01
<<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 8.3.1 URL; format: http://
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02
<<var_node02_mgmt_mask>>	Out-of-band management network netmask
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway
<<var_clustername>>	Storage cluster host name
<<var_cluster_base_license_key>>	Cluster base license key
<<var_nfs_license>>	NFS license key
<<var_iscsi_license>>	iSCSI license key
<<var_password>>	Global default administrative password
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster
<<var_clustermgmt_mask>>	Out-of-band management network netmask
<<var_clustermgmt_gateway>>	Out-of-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_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_node01>>	Cluster node 01 hostname

Variable	Value
<<var_node02>>	Cluster node 02 hostname
<<var_num_disks>>	Number of disks to assign to each storage controller
<<var_nfs_vlan_id>>	Infrastructure NFS VLAN ID for LIF
<<var_iscsi_vlan_A_id>>	Infrastructure iSCSI-A VLAN ID for LIF
<<var_iscsi_vlan_B_id>>	Infrastructure iSCSI-B VLAN ID for LIF
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID
<<var_oob_mgmt_vlan_id>>	Out-of-band management network VLAN ID
<<var_timezone>>	FlexPod time zone (for example, America/New_York)
<<var_global_ntp_server_ip>>	NTP server IP address for out-of-band mgmt
<<var_switch_a_ntp_ip>>	NTP server IP address for Nexus 9372 Switch A
<<var_switch_b_ntp_ip>>	NTP server IP address for Nexus 9372 Switch B
<<var_ib-mgmt_vlan_netmask_length>>	Length of IB-MGMT-VLAN Netmask
<<var_snmp_contact>>	Administrator e-mail address
<<var_snmp_location>>	Cluster location string
<<var_oncommand_server_fqdn>>	VSC or OnCommand virtual machine fully qualified domain name (FQDN)
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name
<<var_mailhost>>	Mail server host name
<<var_storage_admin_email>>	Administrator e-mail address
<<var_esxi_host1_nfs_ip>>	NFS VLAN IP address for VMware ESXi host 1
<<var_esxi_host2_nfs_ip>>	NFS VLAN IP address for VMware ESXi host 2
<<var_node01_nfs_lif_infra_swap_ip>>	IP address of Infra Swap
<<var_node01_nfs_lif_infra_swap_mask>>	Subnet Mask of Infra Swap
<<var_node02_nfs_lif_infra_datastore_1_ip>>	IP address of Datastore 1
<<var_node02_nfs_lif_infra_datastore_1_mask>>	Subnet mask of Datastore 1
<<var_vserver_mgmt_ip>>	Management IP address for Vserver
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver
<<var_vserver_mgmt_gateway>>	Default Gateway for Vserver
<<var_vsadmin_password>>	Password for VS admin account

Variable	Value
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address
<<var_vm_host_infra_01_iqn>>	IQN of Infra 01
<<var_vm_host_infra_02_iqn>>	IQN of Infra 02
<<var_vm_host_infra_01_ip>>	VMware ESXi host 01 out-of-band management IP
<<var_vm_host_infra_02_ip>>	VMware ESXi host 02 out-of-band management IP
<<var_nfs_vlan_ip_host_01>>	ESXi host 1, NFS VLAN IP
<<var_nfs_vlan_ip_mask_host_01>>	ESXi host1, NFS VLAN subnet mask
<<var_nfs_vlan_ip_host_02>>	ESXi host 2, NFS VLAN IP
<<var_nfs_vlan_ip_mask_host_02>>	ESXi host2, NFS VLAN subnet mask
<<var_vcenter_server_ip>>	IP address of the vCenter Server
<<var_svm_mgmt_vlan_id>>	Infrastructure Vserver management VLAN ID
<<var_node01_iscsi_lif01a_ip>>	iSCSI LIF 01a IP address
<<var_node01_iscsi_lif01a_mask>>	iSCSI LIF 01a subnet mask
<<var_node01_iscsi_lif01b_ip>>	iSCSI LIF 01b IP address
<<var_node01_iscsi_lif01b_mask>>	iSCSI LIF 01b subnet mask
<<var_node01_iscsi_lif02a_ip>>	iSCSI LIF 02a IP address
<<var_node01_iscsi_lif02a_mask>>	iSCSI LIF 02a subnet mask
<<var_node01_iscsi_lif02b_ip>>	iSCSI LIF 02b IP address
<<var_node01_iscsi_lif02b_mask>>	iSCSI LIF 02b subnet mask
<<var_vserver_mgmt_ip>>	Management IP address for Infrastructure Vserver
<<var_vserver_mgmt_mask>>	Management subnet mask for Infrastructure Vserver
<<var_oncommand_server_ip>>	IP address of the OnCommand Unified Manager
<<var_rule_index>>	Rule index number

Variable	Value
<<var_server_nfs_vlan_id>>	NFS VLAN ID
<<var_nfs_lif02_ip>>	NFS LIF 02 IP Address
<<var_nfs_lif01_ip>>	NFS LIF 01 IP Address

Physical Infrastructure

FlexPod Cabling

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

The tables in this section contain the details for the prescribed and supported configuration of the NetApp AFF8060 running clustered Data ONTAP 8.3.1.



Data ONTAP 8.3.1. is the minimum supported version for the AFF8060. For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool](#) (IMT).

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

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

Figure 2 shows a cabling diagram for a FlexPod configuration using the Cisco Nexus 9000 and NetApp storage systems with clustered Data ONTAP. The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 2 FlexPod Cabling Diagram

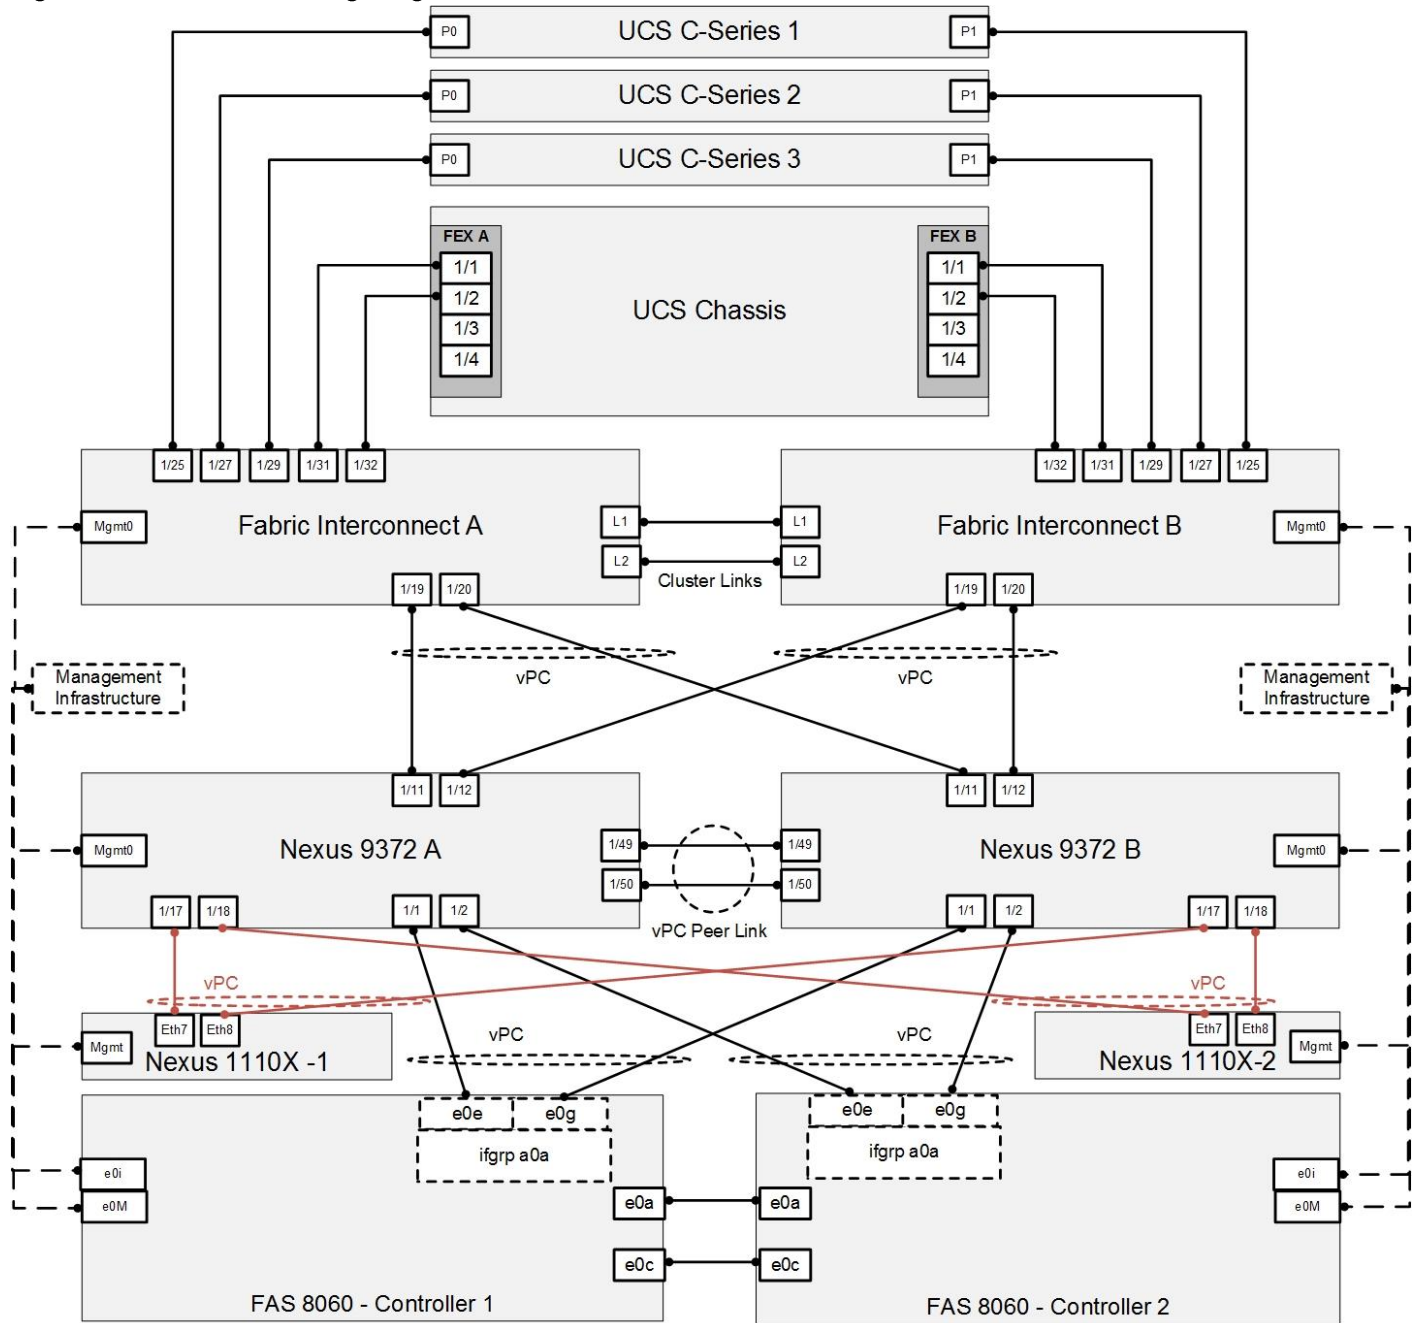


Table 5 through Table 13 provide the details of all the connections in use.

Table 5 Cisco Nexus 9372-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 A	Eth1/1	10GbE	NetApp Controller 1	e0e
	Eth1/2	10GbE	NetApp Controller 2	e0e

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/19
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/19
	Eth1/17	10GbE	Nexus 1110-X 1	Eth7
	Eth1/18	10GbE	Nexus 1110-X 2	Eth7
	Eth1/49	40GbE	Cisco Nexus 9372 B	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 B	Eth1/50
	MGMT0	GbE	GbE management switch	Any



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

Table 6 Cisco Nexus 9372-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 B	Eth1/1	10GbE	NetApp Controller 1	e0g
	Eth1/2	10GbE	NetApp Controller 2	e0g
	Eth1/11	10GbE	Cisco UCS fabric interconnect A	Eth1/20
	Eth1/12	10GbE	Cisco UCS fabric interconnect B	Eth1/20
	Eth1/17	10GbE	Nexus 1110-X 1	Eth8
	Eth1/18	10GbE	Nexus 1110-X 2	Eth8
	Eth1/49	40GbE	Cisco Nexus 9372 A	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 A	Eth1/50
	MGMT0	GbE	GbE management switch	Any

Table 7 NetApp Controller-1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port

Local Device	Local Port	Connection	Remote Device	Remote Port
	e0a	10GbE	NetApp Controller 2	e0a
	e0b	10GbE	NetApp Controller 2	e0b
	e0c	10GbE	NetApp Controller 2	e0c
	e0d	10GbE	NetApp Controller 2	e0d
	e0e	10GbE	Cisco Nexus 9372 A	Eth1/1
	e0g	10GbE	Cisco Nexus 9372 B	Eth1/1



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

Table 8 NetApp Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 1	e0a
	e0b	10GbE	NetApp Controller 1	e0b
	e0c	10GbE	NetApp Controller 1	e0c
	e0d	10GbE	NetApp Controller 1	e0d
	e0e	10GbE	Cisco Nexus 9372 A	Eth1/2
	e0g	10GbE	Cisco Nexus 9372 B	Eth1/2

Table 9 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/19	10GbE	Cisco Nexus 9372 A	Eth1/11
	Eth1/20	10GbE	Cisco Nexus 9372 B	Eth1/11
	Eth1/25	10GbE	Cisco UCS C-Series 1	Port 0

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/27	10GbE	Cisco UCS C-Series 2	Port 0
	Eth1/29	10GbE	Cisco UCS C-Series 3	Port 0
	Eth1/31	10GbE	Cisco UCS Chassis FEX A	IOM 1/1
	Eth1/32	10GbE	Cisco UCS Chassis FEX A	IOM 1/2
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 10 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/19	10GbE	Cisco Nexus 9372 A	Eth1/12
	Eth1/20	10GbE	Cisco Nexus 9372 B	Eth1/12
	Eth1/25	10GbE	Cisco UCS C-Series 1	Port 1
	Eth1/27	10GbE	Cisco UCS C-Series 2	Port 1
	Eth1/29	10GbE	Cisco UCS C-Series 3	Port 1
	Eth1/31	10GbE	Cisco UCS Chassis FEX B	IOM 2/1
	Eth1/32	10GbE	Cisco UCS Chassis FEX B	IOM 2/2
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 11 Cisco UCS C-Series 1

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 1	Port 0	10GbE	Cisco UCS fabric interconnect A	Eth1/25
	Port 1	10GbE	Cisco UCS fabric interconnect B	Eth1/25

Table 12 Cisco UCS C-Series 2

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 2	Port 0	10GbE	Cisco UCS fabric interconnect A	Eth1/27
	Port 1	10GbE	Cisco UCS fabric interconnect B	Eth1/27

Table 13 Cisco UCS C-Series 3

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 3	Port 0	10GbE	Cisco UCS fabric interconnect A	Eth1/29
	Port 1	10GbE	Cisco UCS fabric interconnect B	Eth1/29

Network Switch Configuration

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

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as covered in the section **Error! Reference source not found.** FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Nexus 9000 7.0(3)1(1a).



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

Set Up Initial Configuration

Cisco Nexus 9372PX A

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

1. Configure the switch.



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

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

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

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

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

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

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

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

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

```
Enter the switch name: <<var_nexus_A_hostname>>
Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_A_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_A_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <<var_nexus_A_mgmt0_gw>>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <<var_global_ntp_server_ip>>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut)
[noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip)
[strict]: Enter

Would you like to edit the configuration? (yes/no) [n]: Enter
2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus 9372PX B

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

1. Configure the switch.



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

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

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

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

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

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

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

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

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

Enter the switch name: <<var_nexus_B_hostname>>

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

Mgmt0 IPv4 address: <<var_nexus_B_mgmt0_ip>>

Mgmt0 IPv4 netmask: <<var_nexus_B_mgmt0_netmask>>

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

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

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

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

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

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

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

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

NTP server IPv4 address: <<var_global_ntp_server_ip>>

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

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

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

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

2. Review the configuration summary before enabling the configuration.

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

FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

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

```
config t
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

Set Global Configurations

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <<var_global_ntp_server_ip>> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <<var_ib-mgmt-vlan_gateway>>
copy run start
```

Create VLANs

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

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

```
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
exit
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
exit
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
exit
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
exit
vlan <<var_iscsi-a_vlan_id>>
name iSCSI-A-VLAN
exit
vlan <<var_iscsi-b_vlan_id>>
name iSCSI-B-VLAN
exit
vlan <<var_packet-ctrl_vlan_id>>
name Packet-Ctrl-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus 9372PX A

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

```
ntp source <<var_switch_a_ntp_ip>>
interface Vlan<<var_ib-mgmt_vlan_id>>
```

```
ip address <<var_switch_a_ntp_ip>>/<<var_ib-  
mgmt_vlan_netmask_length>>  
  
no shutdown  
  
exit
```

Cisco Nexus 9372PX B

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

```
ntp source <<var_switch_b_ntp_ip>>  
  
interface Vlan<<var_ib-mgmt_vlan_id>>  
  
ip address <<var_switch_b_ntp_ip>>/<<var_ib-  
mgmt_vlan_netmask_length>>  
  
no shutdown  
  
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 9372PX A

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

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

```
interface Eth1/1  
description <<var_node01>>:e0e  
exit  
  
interface Eth1/2  
description <<var_node02>>:e0e  
exit  
  
interface Eth1/11  
description <<var_ucs_clustername>>-a:1/19  
exit  
  
interface Eth1/12  
description <<var_ucs_clustername>>-b:1/19  
exit  
  
interface Eth1/17
```



```
description <<var_n1110-x>>-1:eth 7
exit
interface Eth1/18
description <<var_n1110-x>>-2:eth7
exit
interface Eth1/49
description <<var_nexus_B_hostname>>:1/49
exit
interface Eth1/50
description <<var_nexus_B_hostname>>:1/50
exit
```

Cisco Nexus 9372PX B

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

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

```
interface Eth1/1
description <<var_node01>>:e0g
exit
interface Eth1/2
description <<var_node02>>:e0g
exit
interface Eth1/11
description <<var_ucs_clustername>>-a:1/20
exit
interface Eth1/12
description <<var_ucs_clustername>>-b:1/20
exit
interface Eth1/17
description <<var_n1110-x>>-1:eth 8
exit
```

```
interface Eth1/18
description <<var_n1110-x>>-2:eth 8
exit
interface Eth1/49
description <<var_nexus_A_hostname>>:1/49
exit
interface Eth1/50
description <<var_nexus_A_hostname>>:1/50
exit
```

Create Port Channels

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

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

```
interface Po10
description vPC peer-link
exit
interface Eth1/49-50
channel-group 10 mode active
no shutdown
exit
interface Po11
description <<var_node01>>
exit
interface Eth1/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
```

```
exit
interface Eth1/2
channel-group 12 mode active
no shutdown
exit
interface Pol11
description <<var_ucs_clustername>>-a
exit
interface Eth1/11
channel-group 111 mode active
no shutdown
exit
interface Pol12
description <<var_ucs_clustername>>-b
exit
interface Eth1/12
channel-group 112 mode active
no shutdown
exit
interface Pol17
description <<var_n1110-x>>-1
exit
interface Eth1/17
channel-group 117 mode active
no shutdown
exit
interface Pol18
description <<var_n1110-x>>-2
exit
```

```
interface Eth1/18
channel-group 118 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

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

```
interface Po10
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>,
<<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
traffic_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-
b_vlan_id>>, <<var_packet-ctrl_vlan_id>>
spanning-tree port type network
exit
interface Po11
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>,
<<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-
b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po12
switchport mode trunk
switchport trunk native vlan 2
```

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

spanning-tree port type edge trunk

mtu 9216

exit

interface Po111

switchport mode trunk

switchport trunk native vlan 2

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

spanning-tree port type edge trunk

mtu 9216

exit

interface Po112

switchport mode trunk

switchport trunk native vlan 2

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

spanning-tree port type edge trunk

mtu 9216

exit

interface Po117

switchport mode trunk

switchport trunk native vlan 2

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

spanning-tree port type edge trunk

exit

interface Po118
```

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>,
<<var_packet-ctrl_vlan_id>>
spanning-tree port type edge trunk
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus 9372PX A

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

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

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 10
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source
<<var_nexus_A_mgmt0_ip>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po111
```

```
vpc 111
exit
interface Po112
vpc 112
exit
interface Po117
vpc 117
exit
interface Po118
vpc 118
exit
copy run start
```

Cisco Nexus 9372PX B

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

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

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 20
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source
<<var_nexus_B_mgmt0_ip>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
```

```
interface Po12
vpc 12
exit
interface Po111
vpc 111
exit
interface Po112
vpc 112
exit
interface Po117
vpc 117
exit
interface Po118
vpc 118
exit
copy run start
```

Uplink into Existing Network Infrastructure

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

Storage Configuration

Controller AFF80XX Series

Refer to the [Site Requirements Guide](#) for planning the physical location of the storage systems. From the downloaded guide, refer the following sections:

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

NetApp Hardware Universe

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

1. Confirm that the hardware and software components are supported with the version of Data ONTAP that you plan to install by using the [NetApp Hardware Universe \(HWU\) application](#) at the [NetApp Support](#) site.
2. Access the [HWU](#) 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. These procedures can be found in the [AFF8000 Series product documentation](#) at the [NetApp Support](#) site.

Disk Shelves

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

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

Clustered Data ONTAP 8.3.1

Complete the Configuration Worksheet

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

Configure Clustered Data ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [Clustered Data ONTAP 8.3 Software Setup Guide](#) to learn about the information required to configure clustered Data ONTAP. Table 14 lists the information that you will need to configure two clustered Data ONTAP nodes. You should customize the cluster detail values with the information that is applicable to your deployment.

Table 14 Clustered Data ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.3.1 URL	<<var_url_boot_software>>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If Data ONTAP 8.3.1 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.1 is the version being booted, select option 8 and y (Yes) to reboot the node, then continue with step 14.

4. To install new software, select option 7.

7

5. Enter y (Yes) to perform an upgrade.

y

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

e0M

7. Enter y (Yes) to reboot now.

y

8. After reboot, 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>>

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



This web server must be pingable.

<<var_url_boot_software>>

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

Enter

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

y

12. Enter y (Yes) to reboot the node.

y



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

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

Press Ctrl-C for Boot Menu

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

4

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

`y`

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

`y`



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

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press `Ctrl-C` when prompted.

```
Ctrl-C
```



If Data ONTAP 8.3.1 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.1 is the version being booted, select option 8 and `yes` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.

7

5. Enter `y` (Yes) to perform a non-disruptive upgrade.

`y`

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

`e0M`

7. Enter `y` (Yes) to reboot now.

y

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

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

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



This web server must be pingable.

```
<<var_url_boot_software>>
```

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

Enter

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

y

12. Enter y (Yes) to reboot the node.

y



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

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

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

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

4

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

y

16. Enter y (Yes) to erase all the data on the disks.

y



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

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when Data ONTAP 8.3.1 boots on the node for the first time.

1. Follow the prompts to set up node 01.

```
Welcome to node setup.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
```

```
"back" - if you want to change previously answered questions, and
```

```
"exit" or "quit" - if you want to quit the setup wizard.
```

```
Any changes you made before quitting will be saved.
```

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

```
This system will send event messages and weekly reports to NetApp  
Technical
```

```
Support.
```

```
To disable this feature, enter "autosupport modify -support  
disable" within 24
```

```
hours.
```

```
Enabling AutoSupport can significantly speed problem determination  
and
```

```
resolution should a problem occur on your system.
```

```
For further information on AutoSupport, see:
```

```
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}: yes
```

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

```
Enter the node management interface IP address:
```

```
<<var_node01_mgmt_ip>>
```

Enter the node management interface netmask:

<<var_node01_mgmt_mask>>

Enter the node management interface default gateway:

<<var_node01_mgmt_gateway>>

A node management interface on port e0M with IP address

<<var_node01_mgmt_ip>> has been created

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

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

For System Setup, this node's management address is:

<<var_node01_mgmt_ip>>.

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

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

3. At the node command prompt, enter the following commands:

```
::> storage failover modify -mode ha
```

Mode set to HA. Reboot node to activate HA.

```
::> system node reboot
```

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

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

Welcome to node setup.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,

"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.

Any changes you made before quitting will be saved.

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

Enter the node management interface port [e0M]: Enter

Enter the node management interface IP address
[<<var_node01_mgmt_ip>>]: Enter

Enter the node management interface netmask
[<<var_node01_mgmt_mask>>]: Enter

Enter the node management interface default gateway
[<<var_node01_mgmt_gateway>>]: Enter

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

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

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

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

5. Log in to the node as the admin user and no password.
6. Repeat this procedure for storage cluster node 02.

Create Cluster on Node 01

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

Table 15 Cluster **create** in Clustered Data ONTAP Prerequisites

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

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

```
cluster setup
```

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,

"back" - if you want to change previously answered questions, and

"exit" or "quit" - if you want to quit the cluster setup wizard.

Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".

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

Do you want to create a new cluster or join an existing cluster?
{create, join}:



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

To create a new cluster, complete the following steps:

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

```
create
```

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

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

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

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

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

Existing cluster interface configuration found:

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

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

System Defaults:

Private cluster network ports `[e0a,e0c]`.

Cluster port MTU values will be set to 9000.

Cluster interface IP addresses will be automatically generated.

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



If 4 ports are being used for the switchless cluster interconnect, enter e0a, e0b, e0c, e0d for the private cluster network ports above.

5. The steps to create a cluster are displayed.

Enter the cluster administrators (username "admin") password:

`<<var_password>>`

Retype the password: `<<var_password>>`

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

Step 1 of 5: Create a Cluster

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

Enter the cluster name: <<var_clustername>>

Enter the cluster base license key:

<<var_cluster_base_license_key>>

Creating cluster <<var_clustername>>

Enter an additional license key []:<<var_iscsi_license>>



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



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

Enter the cluster management interface port [e0e]: e0i

Enter the cluster management interface IP address:

<<var_clustermgmt_ip>>

Enter the cluster management interface netmask:

<<var_clustermgmt_mask>>

Enter the cluster management interface default gateway:

<<var_clustermgmt_gateway>>

6. Enter the DNS domain name.

Enter the DNS domain names:<<var_dns_domain_name>>

Enter the name server IP addresses:<<var_nameserver_ip>>



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

7. Set up the node.

Where is the controller located []:<<var_node_location>>

Enter the node management interface port [e0M]: e0M

Enter the node management interface IP address

[<<var_node01_mgmt_ip>>]: Enter

Enter the node management interface netmask

[<<var_node01_mgmt_mask>>]: Enter

Enter the node management interface default gateway

[<<var_node01_mgmt_gateway>>]: Enter

The node management interface has been modified to use port e0M with IP address <<var_node01_mgmt_ip>>.

This system will send event messages and weekly reports to NetApp Technical Support.

To disable this feature, enter "autosupport modify -support disable" within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.

For further information on AutoSupport, please see:
<http://support.netapp.com/autosupport/>

Press enter to continue: Enter
Cluster "<<var_clustername>>" has been created.

To complete cluster setup, you must join each additional node to the cluster

by running "cluster setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP

Software Setup Guide for information about additional system configuration

tasks. You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager

or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster

management IP address (<<var_clustermgmt_ip>>).

To access the command-line interface, connect to the cluster management

IP address (for example, ssh admin@<<var_clustermgmt_ip>>).

```
<<var_clustername>>::>
```



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

Join Node 02 to Cluster

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

Table 16 Cluster `join` in Clustered Data ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<code><<var_clustername>></code>
Cluster management IP address	<code><<var_clustermgmt_ip>></code>
Cluster node02 IP address	<code><<var_node02_mgmt_ip>></code>
Cluster node02 netmask	<code><<var_node02_mgmt_mask>></code>
Cluster node02 gateway	<code><<var_node02_mgmt_gateway>></code>

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

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

```
admin
```

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

```
cluster setup
```

This node's storage failover partner is already a member of a cluster.

Storage failover partners must be members of the same cluster.

The cluster setup wizard will default to the cluster join dialog.

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,

"back" - if you want to change previously answered questions, and

"exit" or "quit" - if you want to quit the cluster setup wizard.

Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".

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

Do you want to create a new cluster or join an existing cluster?

{join}:



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

3. Run the following command to join a cluster:

```
join
```

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

Existing cluster interface configuration found:

Port	MTU	IP	Netmask
e0a	9000	169.254.1.79	255.255.0.0
e0c	9000	169.254.100.157	255.255.0.0

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

System Defaults:

Private cluster network ports [e0a,e0c].

Cluster port MTU values will be set to 9000.

Cluster interface IP addresses will be automatically generated.



If 4 ports are being used for the switchless cluster interconnect, enter e0a, e0b, e0c, e0d for the private cluster network ports above.

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

5. The steps to join a cluster are displayed.

Step 1 of 3: Join an Existing Cluster

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

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

Starting cluster support services ..
```

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

Step 2 of 3: Configure Storage Failover (SFO)

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

SFO is enabled.

Step 3 of 3: Set Up the Node

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

Notice: HA is configured in management.



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

6. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address
[<<var_node02_mgmt_ip>>]: Enter
Enter the node management interface netmask
[<<var_node02_netmask>>]: Enter
Enter the node management interface default gateway
[<<var_node02_gw>>]: Enter
The node management interface has been modified to use port e0M
with IP address <<var_node02_mgmt_ip>>.
```

This system will send event messages and weekly reports to NetApp Technical Support.

To disable this feature, enter "autosupport modify -support disable" within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.

For further information on AutoSupport, please see:
<http://support.netapp.com/autosupport/>

Press enter to continue: Enter

This node has been joined to cluster "<<var_clustername>>".

To complete cluster setup, you must join each additional node to the cluster

by running "cluster setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP

Software Setup Guide for information about additional system configuration

tasks. You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager

or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster

management IP address (<<var_clustermgmt_ip>>).

To access the command-line interface, connect to the cluster management

IP address (for example, ssh admin@<<var_clustermgmt_ip>>).



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

Log In to the Cluster

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

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

Zero All Spare Disks

To zero all spare disks in the cluster, complete the following step:

1. Run the following command:

```
disk zerospares
```



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

Set Onboard UTA2 Ports Personality

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

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

```
ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
<<var_node01>>	0e	cna	target	-	-	online
<<var_node01>>	0f	cna	target	-	-	online
<<var_node01>>	0g	cna	target	-	-	online
<<var_node01>>	0h	cna	target	-	-	online
<<var_node02>>	0e	cna	target	-	-	online
<<var_node02>>	0f	cna	target	-	-	online

```

<<var_node02>>
           0g      cna      target      -      -      online
<<var_node02>>
           0h      cna      target      -      -      online
8 entries were displayed.

```

2. Verify that the Current Mode of all the ports in use is `cna` and the Current Type is set to `target`. If not, change the port personality by running the following command:

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



The ports must be offline to run this command. To take an adapter offline, run the `fcport adapter modify -node <home node of the port> -adapter <port name> -state down` command. Ports must be converted in pairs, for example, 0c and 0d, after which, a reboot is required, and the ports must be brought back to the up state.

Set Auto-Revert on Cluster Management

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



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

1. Run the following command:

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

Set Up Management Broadcast Domain

To set up the default broadcast domain for management network interfaces, complete the following step:

1. Run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_node01>>:e0e,<<var_node01>>:e0f,<<var_node01>>:e0g,<<var_node
01>>:e0h,<<var_node01>>:e0j,<<var_node01>>:e0k,<<var_node01>>:e0l,<
<var_node02>>:e0e,<<var_node02>>:e0f,<<var_node02>>:e0g,<<var_node0
2>>:e0h,<<var_node02>>:e0j,<<var_node02>>:e0k,<<var_node02>>:e0l
broadcast-domain show
```

Set Up Service Processor Network Interface

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

1. Run the following commands:

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

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



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

Create Aggregates

An aggregate containing the root volume is created during the 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.

To create new aggregates, complete the following steps:

1. Run the following commands:

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

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



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



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



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

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

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

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

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

```
node run <<var_node01>> snap delete -A -a -f aggr1_node01
```

```
node run <<var_node02>> snap delete -A -a -f aggr1_node02
```

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

```
aggr show
```

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

Verify Storage Failover

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

1. Verify the status of storage failover.

```
storage failover show
```



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

2. Enable failover on one of the two nodes.

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



Enabling failover on one node enables it for both nodes.

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



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

```
cluster ha show
```

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

5. Enable HA mode only for the two-node cluster.



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>>
```

Disable Flow Control on UTA2 Ports

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

1. Run the following commands to configure node 01:

```
network port modify -node <<var_node01>> -port
e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0h -flowcontrol-admin none
```

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

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

2. Run the following commands to configure node 02:

```
network port modify -node <<var_node02>> -port
e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0h -flowcontrol-admin none
```

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

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

```
network port show -fields flowcontrol-admin
```

Disable Unused FcoE Ports

Unused data FCoE ports on active interfaces should be disabled. To disable these ports, complete the following step:

1. Run the following commands:

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

Configure NTP

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

1. To set the time zone for the cluster, run the following command:

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



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

2. To set the date for the cluster, run the following command:

```
date <ccyyymmddhhmm.ss>
```



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

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

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

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

Configure SNMP

To configure SNMP, complete the following steps:

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>>
```

Configure SNMPv1 Access

To configure SNMPv1 access, complete the following step:

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

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

Create SNMPv3 User

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

1. Create a user called `snmpv3user`.

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

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

Configure AutoSupport

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

1. Run the following command:

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

Enable Cisco Discovery Protocol

To enable Cisco Discovery Protocol (CDP) on the NetApp storage controllers, complete the following step:



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

2. Run the following command to enable CDP on Data ONTAP:

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

Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

To create a data broadcast domain with an MTU of 9000, complete the following step:

1. Run the following commands to create a broadcast domain on Data ONTAP:

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

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, complete the following step:

1. Run the following commands to create the LACP interface groups:

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode
multimode_lacp
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0e
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0g

ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode
multimode_lacp
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0e
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0g
```

```
ifgrp show
```

Create VLANs

To create VLANs, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```
network port modify -node <<var_node01>> -port a0a -mtu 9000
network port modify -node <<var_node02>> -port a0a -mtu 9000

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

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

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

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

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



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

Create Storage Virtual Machine

To create an infrastructure storage virtual machine (SVM, formerly known as Vserver), complete the following steps:



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

1. Run the `vserver create` command.

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

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

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

3. Add the two data aggregates to the Infra-SVM aggregate list for NetApp Virtual Storage Console (VSC).

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

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

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

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

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

```
vserver nfs show
```

Create Load-Sharing Mirrors of SVM Root Volume

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

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

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

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

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

3. Create the mirroring relationships.

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

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

4. Initialize the mirroring relationship.

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

Create iSCSI Service

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM.

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

Configure HTTPS Access

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

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

```
set -privilege diag
```

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

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a Certificate Authority (CA) To delete the default certificates, run the following commands:



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

```
security certificate delete [TAB] ...
```

```
Example: security certificate delete -vserver Infra-SVM -common-
name Infra-SVM -ca Infra-SVM -type server -serial 552429A6
```

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

```
security certificate create [TAB] ...
```

```
Example: security certificate create -common-name infra-
svm.ciscorobo.com -type server -size 2048 -country US -state
"California" -locality "San Jose" -organization "Cisco" -unit "UCS"
-email-addr "abc@cisco.com" -expire-days 365 -protocol SSL -hash-
function SHA256 -vserver Infra-SVM
```

5. To obtain the values for the parameters that would be required in step 6, run the `security certificate show` command.
6. Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify [TAB] ...
```

```
Example: security ssl modify -vserver clus -server-enabled true -
client-enabled false -ca clus.ciscorobo.com -serial 55243646 -
common-name clus.ciscorobo.com
```

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

```
system services web modify -external true -ssl3-enabled true
```

Warning: Modifying the cluster configuration will cause pending web service requests to be

interrupted as the web servers are restarted.

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

```
system services firewall policy delete -policy mgmt -service http -
vserver <<var_clustername>>
```

8. It is normal for some of these commands to return an error message stating that the entry does not exist.
9. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
```

```
vserver services web modify -name spi|ontapi|compat -vserver * -
enabled true
```

Configure NFSv3

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

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

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

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

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

Create FlexVol Volumes

The following information is required to create a FlexVol volume:

- Volume name
- Volume size
- Aggregate on which the volume exists

To create a NetApp FlexVol® volume, complete the following step:

1. Run the following commands:

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

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

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

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

Create Boot LUNs

To create two boot LUNs, complete the following step:

1. Run the following commands:

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

Schedule Deduplication

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

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

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

2. Create the Always_On_Deduplication efficiency policy:

```
cron create -name lmin -minute
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,3
0,31,32,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47,48,48,50,51,52,53,54,55,56,5
7,58,59
```

```
efficiency policy create -vserver Infra-SVM -policy Always_On_Deduplication -type
scheduled -schedule lmin -qos-policy background -enabled true
```

3. Optionally, assign the Always_On_Deduplication policy to infra_datastore_1:

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

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

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

Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), complete the following step:

1. Run the following commands:

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

```
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role
```

```
data -data-protocol iscsi -home-node <<var_node01>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_node01_iscsi_lif01b_ip>> -
netmask <<var_node01_iscsi_lif01b_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false
```

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

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

```
network interface show
```

Create NFS LIF

To create an NFS LIF, complete the following step:

1. Run the following commands:

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

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

```
network interface show
```



NetApp recommends creating a new LIF for each datastore.

Add Infrastructure SVM Administrator

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

1. Run the following commands:

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



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

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

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

```
network route show
```

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

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

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

Server Configuration

Cisco UCS Base Configuration

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

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

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

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

```
Enter the configuration method: console
```

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

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

```
Enforce strong passwords? (y/n) [y]: y
```

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

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

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

```
Which switch fabric (A|B): A
```

```
Enter the system name: <<var_ucs_clustername>>
```

```
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
```

```
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
```

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

```
Cluster IPv4 address: <<var_ucs_cluster_ip>>
```

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

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

```
Configure the default domain name? y
```

```
Default domain name: <<var_dns_domain_name>>
```



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

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

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

Cisco UCS 6248 B

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

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

```
Enter the configuration method: console
```

```
Installer has detected the presence of a peer Fabric
interconnect. This Fabric interconnect will be added to the
cluster. Continue (y|n)? y
```

```
Enter the admin password for the peer fabric interconnect:
<<var_password>>
```

```
Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>
```

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

```
(yes/no): y
```

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

Cisco UCS Setup

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6248 fabric interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

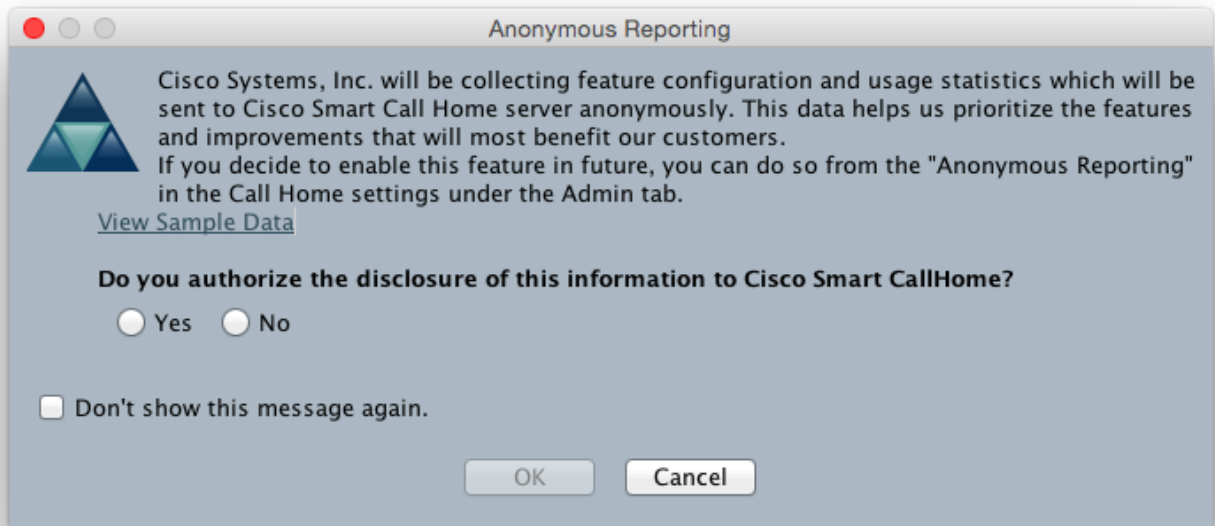
Upgrade Cisco UCS Manager Software to Version 2.2(5b)

This document assumes the use of Cisco UCS 2.2(5b). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to version 2.2(5b), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To create anonymous reporting, complete the following steps:

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



Add Block of IP Addresses for Inband KVM Access

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > IP Pools.
3. Right-click IP Pools and select Create IP Pool.
4. Name the pool in-band-mgmt.
5. Click Next.
6. Click Add.
7. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.
8. Click Next.
9. Click Finish to create the IP block.

10. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

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

Edit Chassis Discovery Policy

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

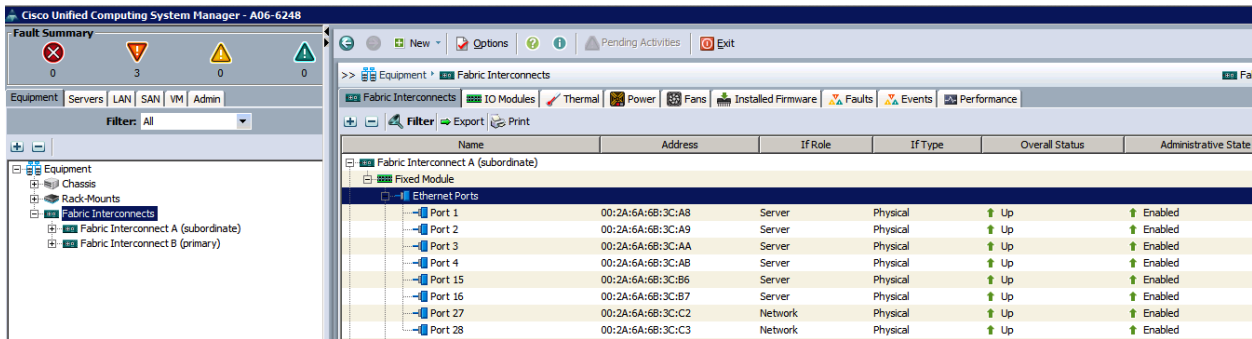
1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and select Equipment in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the number of uplink ports that are cabled between the chassis or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

Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.

3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis, Cisco 2232 FEX (two per FEX), and direct connect UCS C-Series servers, right-click them, and select “Configure as Server Port.”
5. Click Yes to confirm server ports and click OK.
6. Verify that the ports connected to the chassis, C-series servers and to the Cisco 2232 FEX are now configured as server ports.
7. Select ports 19 and 20 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

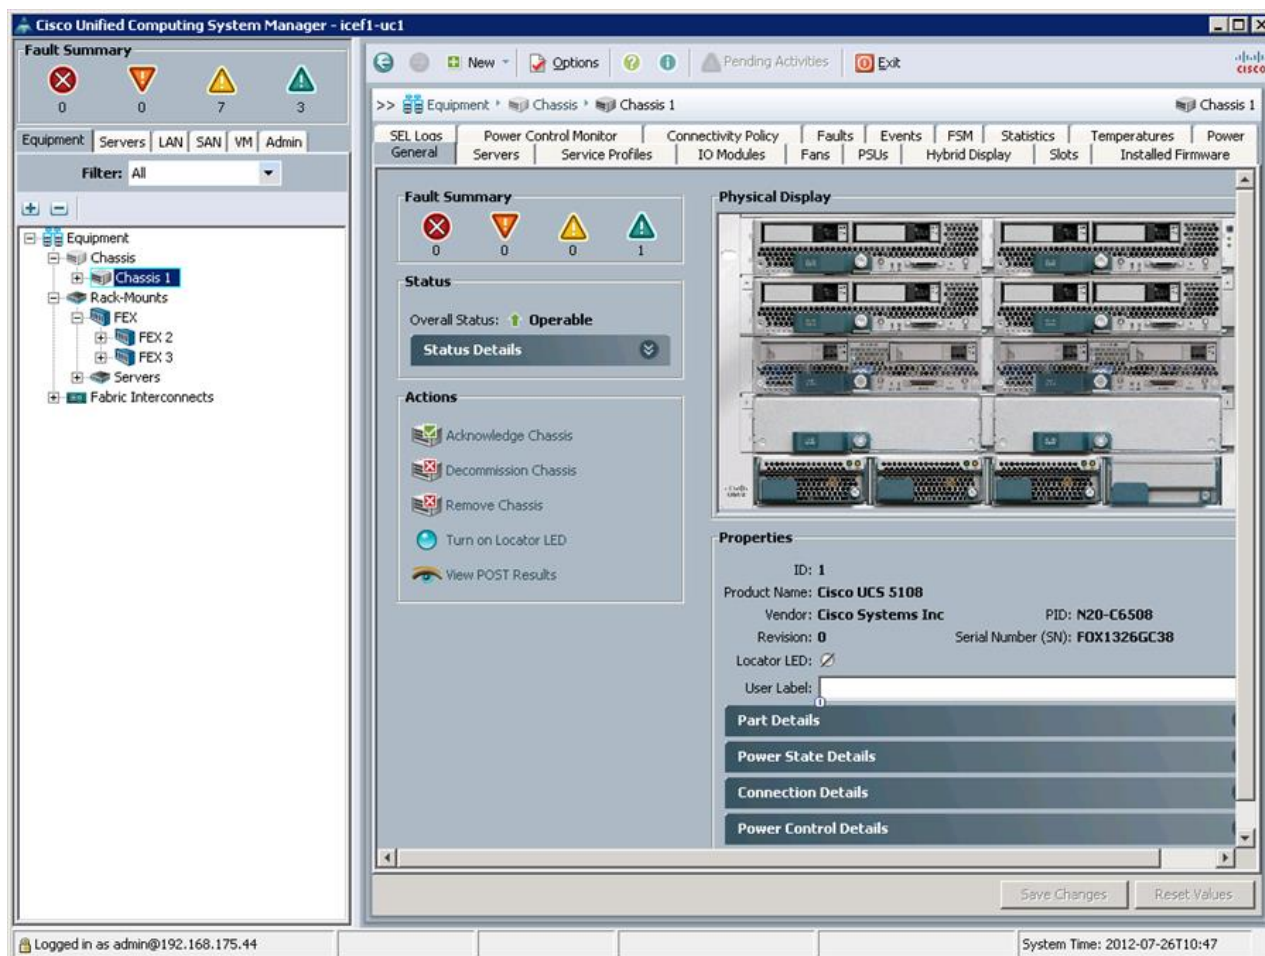


8. Click Yes to confirm uplink ports and click OK.
9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
10. Expand Ethernet Ports.
11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them,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

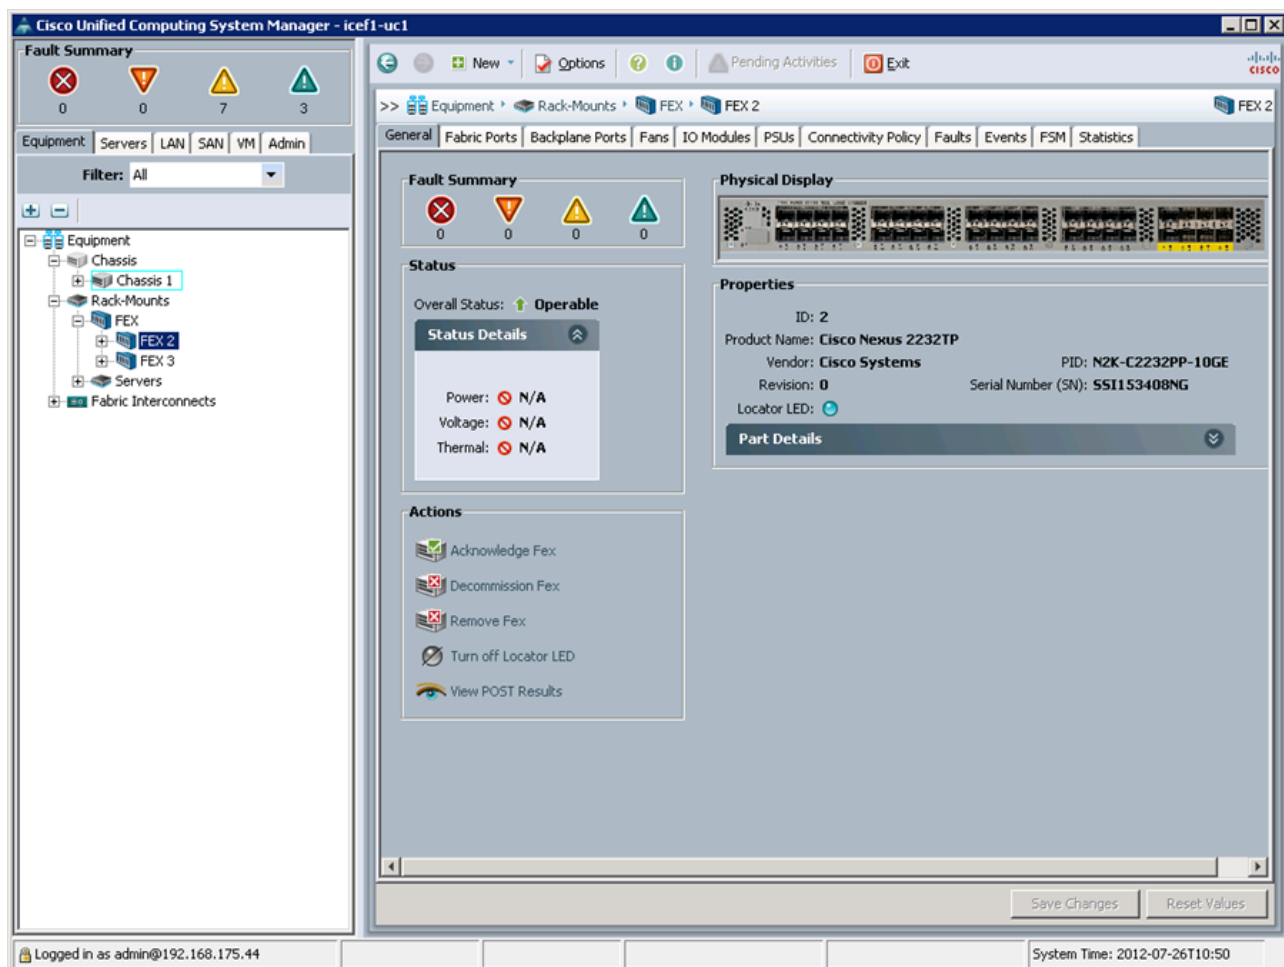
Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and external 2232 FEX modules, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.



4. Click Yes and then click OK to complete acknowledging the chassis.
5. If the Nexus 2232 FEX is part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and select Acknowledge FEX.



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

Create Uplink Port Channels to Cisco Nexus Switches

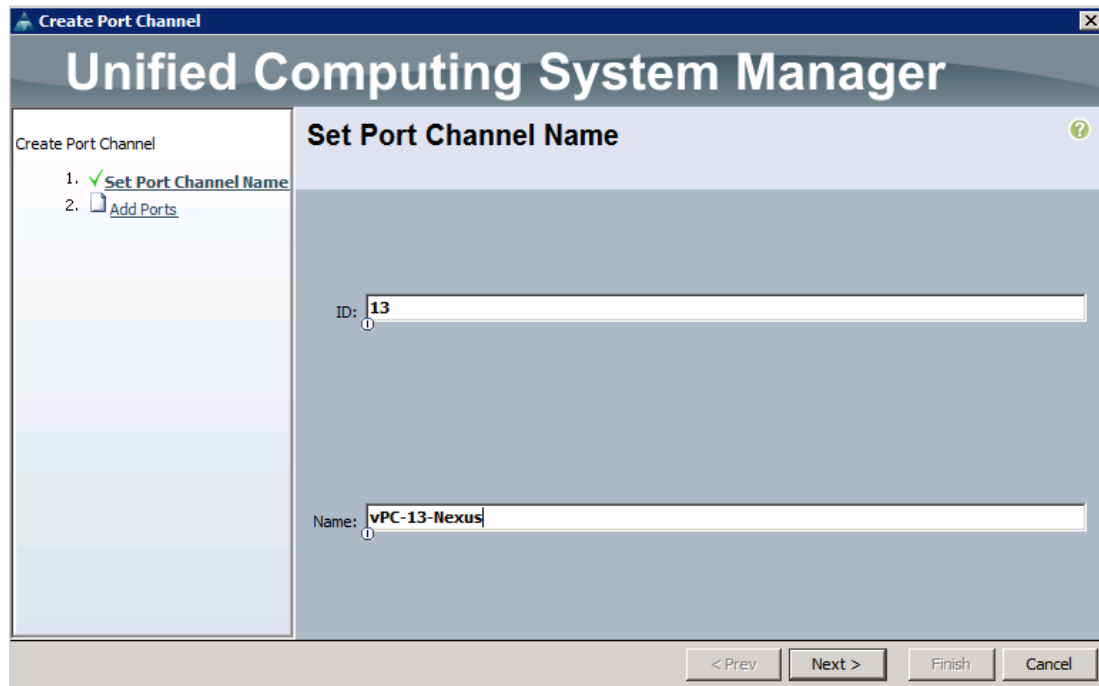
To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

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



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

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



8. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 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. Select Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter vPC-14-Nexus as the name of the port channel.
17. Click Next.
18. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 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 MAC Address Pools

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

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



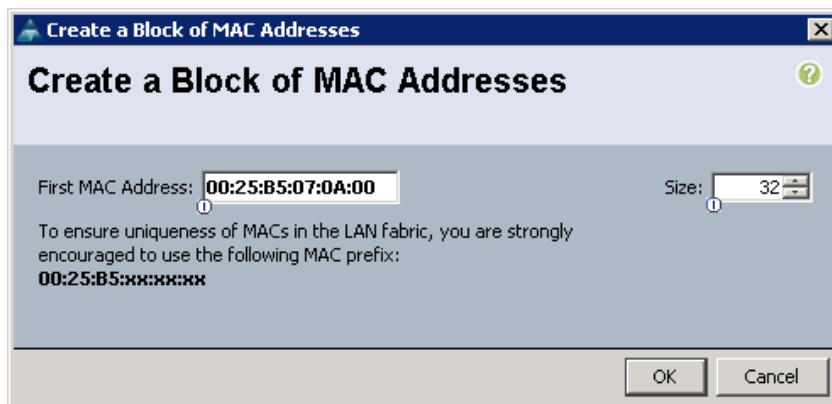
In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter `MAC_Pool_A` as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Click Next.
8. Click Add.
9. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses.

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



Create a Block of MAC Addresses

First MAC Address: Size:

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

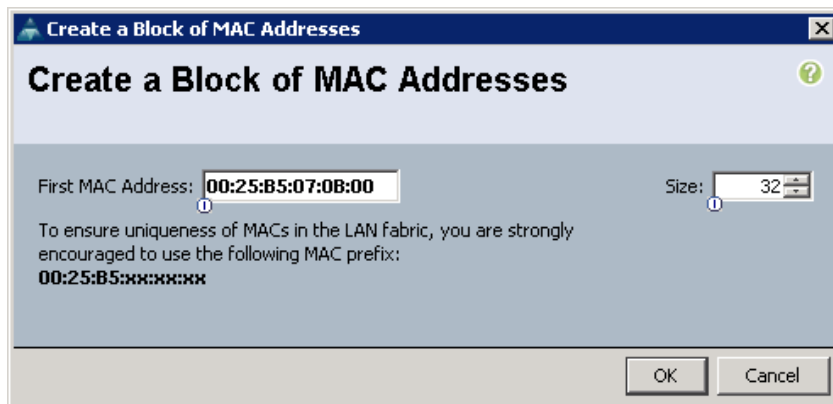
OK Cancel

11. Click OK.
12. Click Finish.
13. In the confirmation message, click OK.
14. Right-click MAC Pools under the root organization.
15. Select 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.



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses.

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



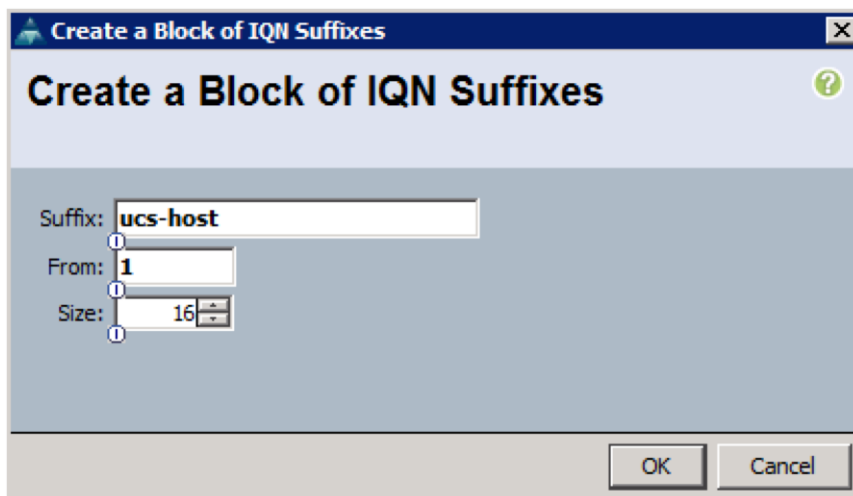
22. Click OK.
23. Click Finish.
24. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot

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

1. In the UCS Manager, select the SAN tab on the left.

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



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

Create IP Pools for iSCSI Boot

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

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

2. Select Pools > root.



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

3. Right-click IP Pools under the root organization.
4. Select Create IP Pool to create the IP pool.
5. Enter `iSCSI_IP_Pool_A` for the name of the IP pool.
6. Optional: Enter a description of the IP pool.
7. Select Sequential for Assignment Order.
8. Click Next.
9. Click Add.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Set the size to enough addresses to accommodate the servers.
12. Click OK.
13. Click Finish.
14. Right-click IP Pools under the root organization.
15. Select Create IP Pool to create the IP pool.
16. Enter `iSCSI_IP_Pool_B` for the name of the IP pool.
17. Optional: Enter a description of the IP pool.
18. Select Sequential for Assignment Order.
19. Click Next.
20. Click Add.
21. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
22. Set the size to enough addresses to accommodate the servers.
23. Click OK.
24. Click Finish.

Create Block of IP Addresses

From: Size:

Subnet Mask: Default Gateway:

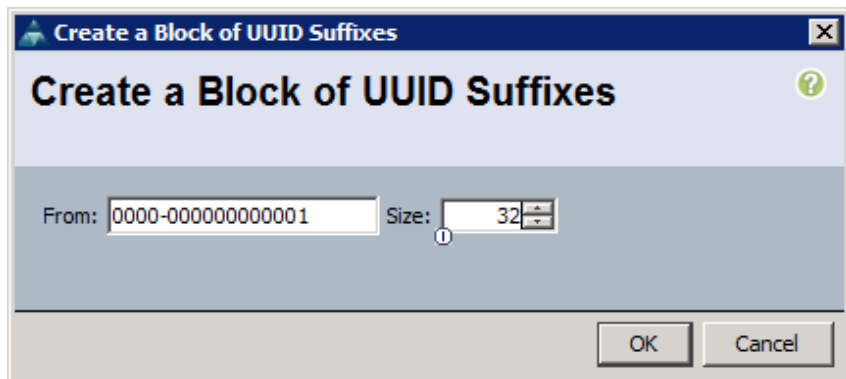
Primary DNS: Secondary DNS:

OK Cancel

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID_Pool` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. 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.



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

Create Server Pool

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



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

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

Create VLANs

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

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



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

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `Native-VLAN` as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name:

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.

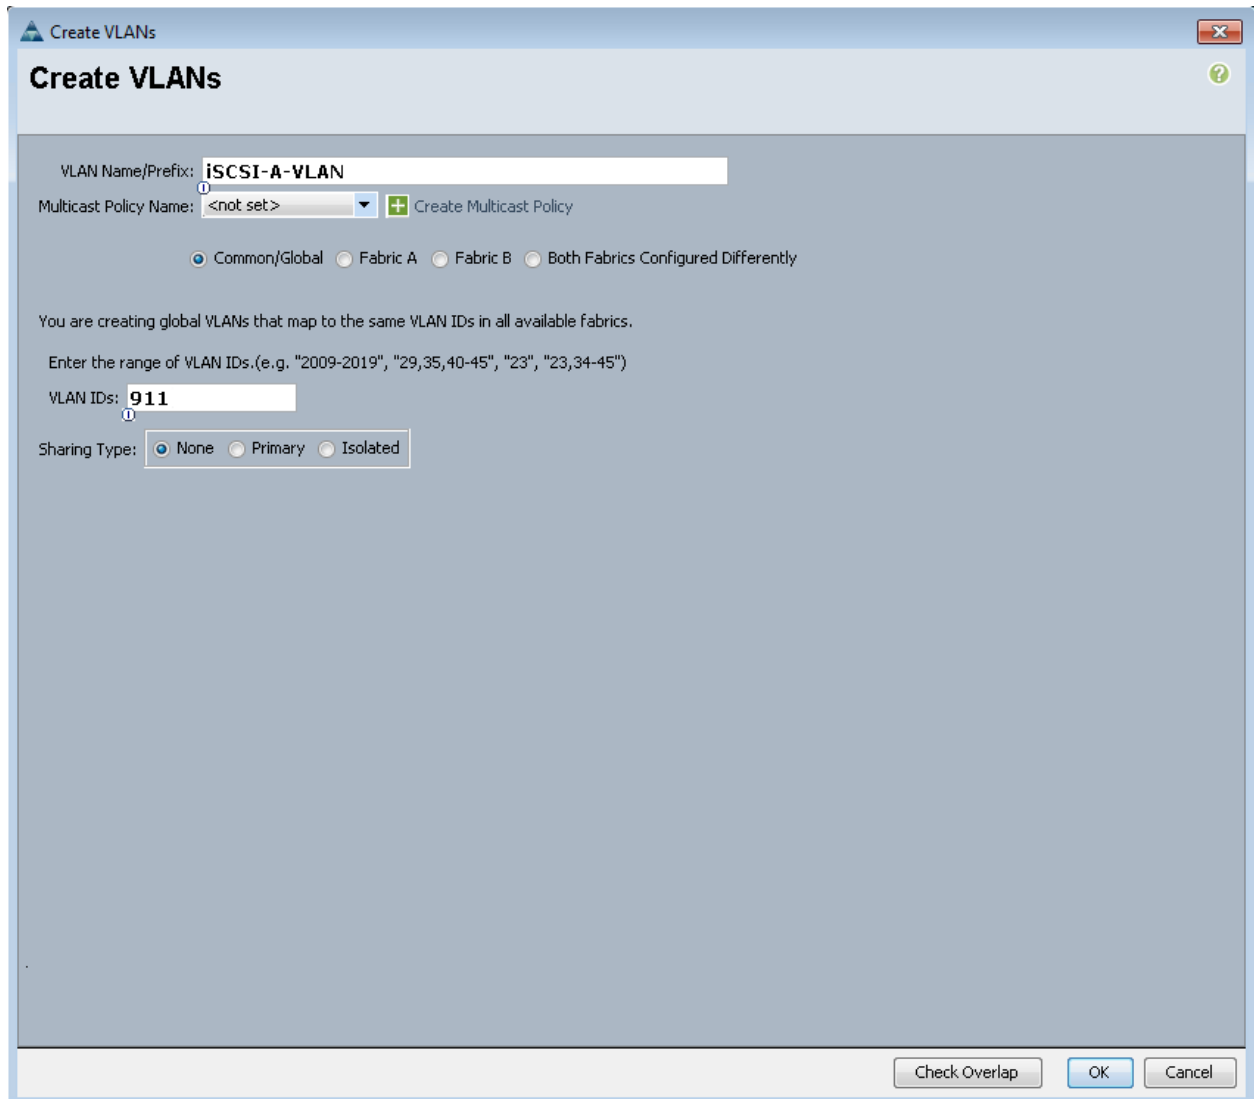
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type: None Primary Isolated

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs.

14. Enter `iSCSI-A-VLAN` as the name of the VLAN to be used for the first iSCSI VLAN.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the VLAN ID for the first iSCSI VLAN.
17. Click OK, then OK.



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

23. Click OK, then OK.

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name:

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. Select Create VLANs

26. Enter `IB-Mgmt` as the name of the VLAN to be used for management traffic.

27. Keep the Common/Global option selected for the scope of the VLAN.

28. Enter the In-Band management VLAN ID.

29. Keep the Sharing Type as None.

30. Click OK, and then click OK again.

31. Right-click VLANs.

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

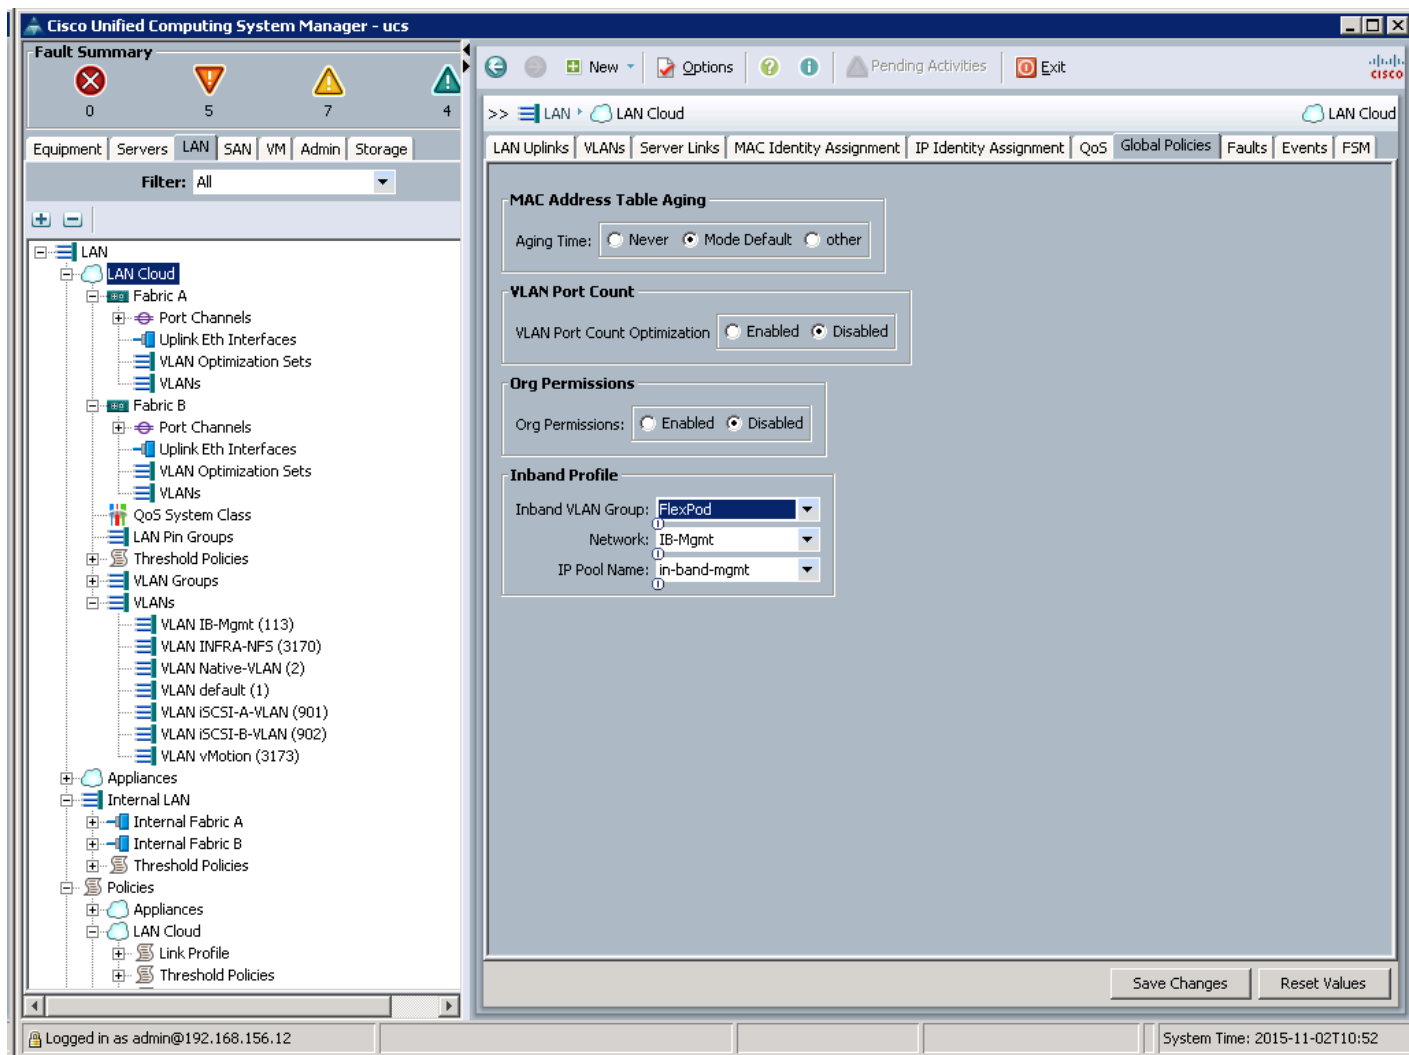
Create VLAN Group and Assign Inband Profile

A VLAN group is required in order to set up Inband KVM Access..

To create a VLAN group, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLAN Groups and select Create VLAN Group.

4. Name the VLAN group FlexPod and select all VLANs.
5. Select the radio button next to the Native-VLAN and click Next.
6. Click Next.
7. Select the two uplink port channels and use the >> button to add them to the VLAN group.
8. Click Finish.
9. Click OK.
10. Select LAN > LAN Cloud. Then select the Global Policies tab.
11. In the Inband Profile box, select the FlexPod VLAN Group, the IB-MGMT Network, and the in-band-mgmt IP Pool Name.
12. Select Save Changes and OK.

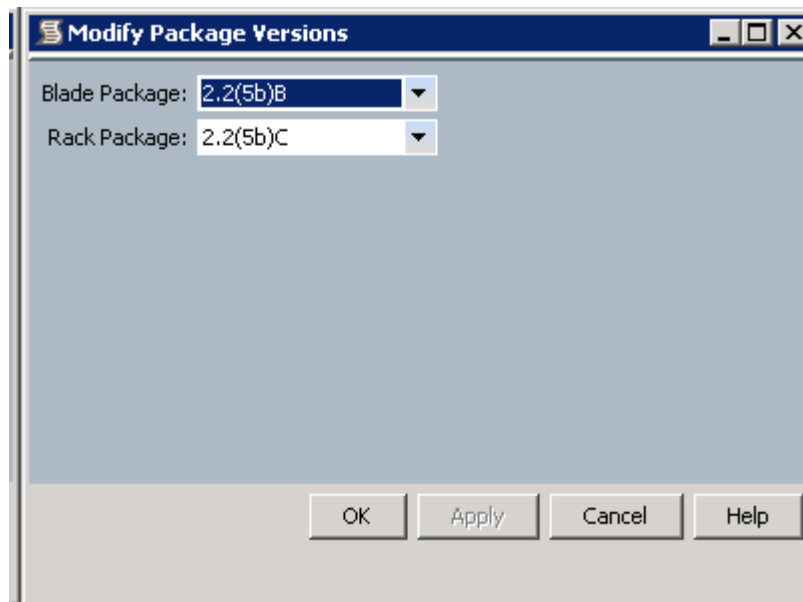


Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 2.2(5b) for both the Blade and Rack Packages.
7. Click OK to modify the host firmware package.



Set Jumbo Frames in Cisco UCS Fabric

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.

4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes in the bottom of the window.
6. Click OK

The screenshot shows the Cisco Unified Computing System Manager interface. The left pane displays a tree view of the network configuration, including LAN Cloud, Fabric A, Fabric B, and QoS System Class. The right pane shows the configuration for the QoS System Class, with the 'Best Effort' row selected. The MTU for the 'Best Effort' row is set to 9216.

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

Create Local Disk Configuration Policy (Optional)

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

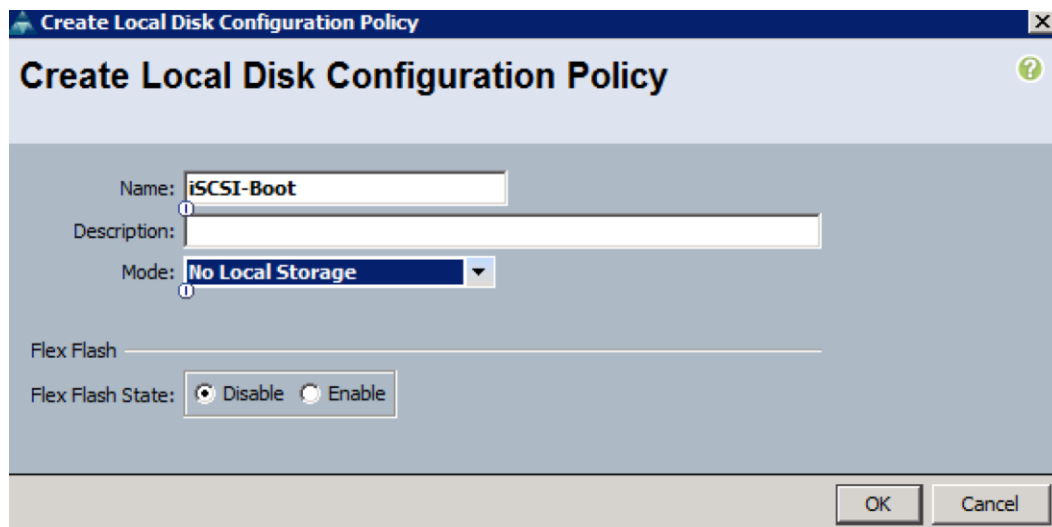


This policy should not be used on servers that contain local disks.

To create a local disk configuration policy, complete the following steps:

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

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



8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

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

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

Create Network Control Policy

Name:

Description:

CDP: Disabled Enabled

MAC Register Mode: Only Native Vlan All Host Vlans

Action on Uplink Fail: Link Down Warning

MAC Security

Forge: Allow Deny

LLDP

Transmit: Disabled Enabled

Receive: Disabled Enabled

OK Cancel

8. Click OK.

Create Power Control Policy

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

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

Create 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

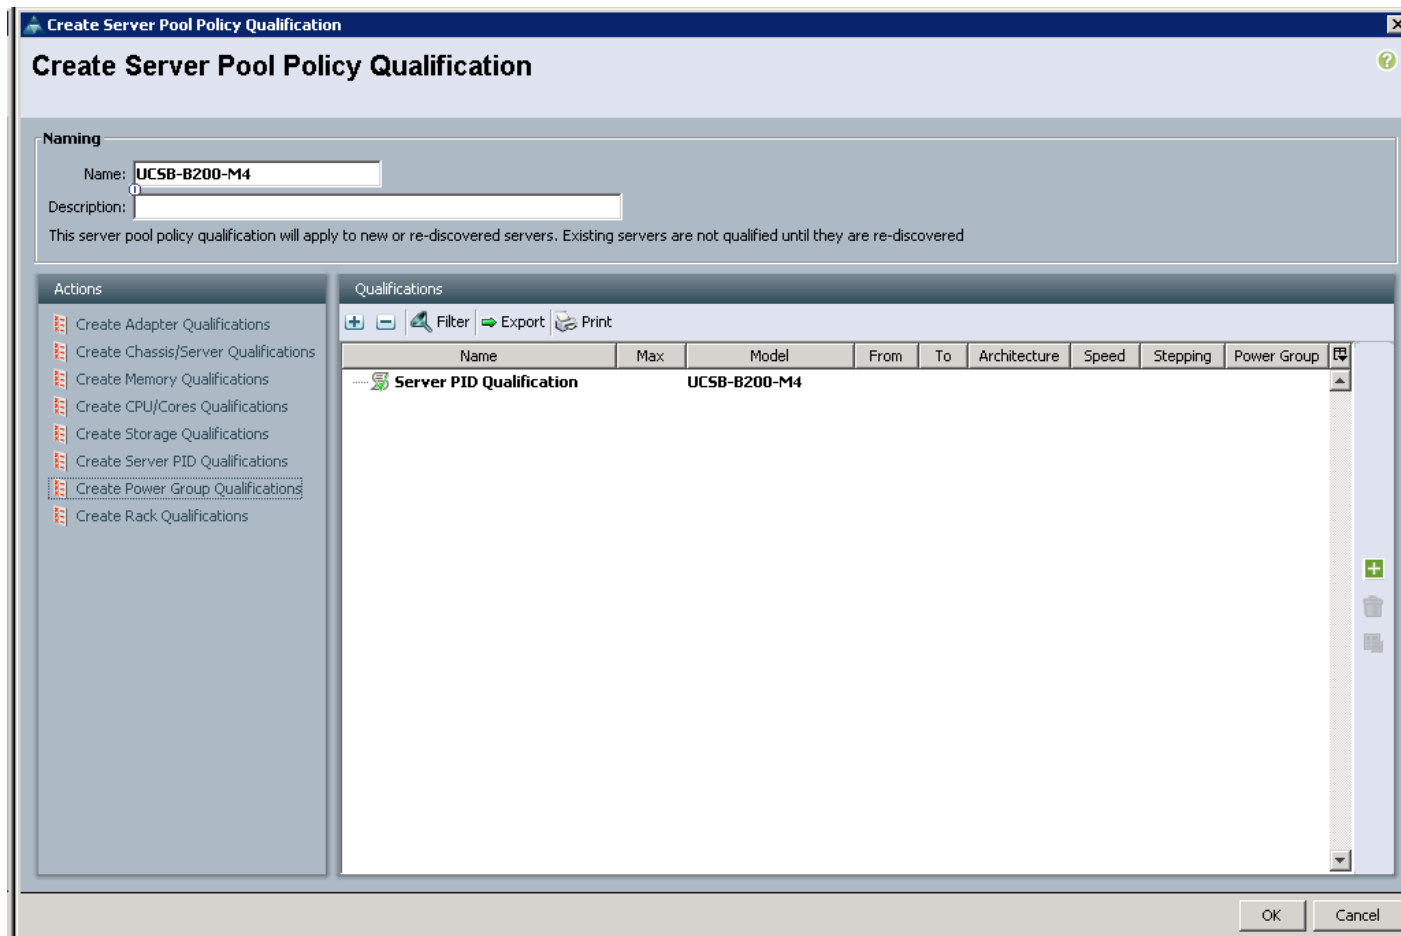
Create Server Pool Qualification Policy (Optional)

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



This example creates a policy for a Cisco UCS B200-M4 server.

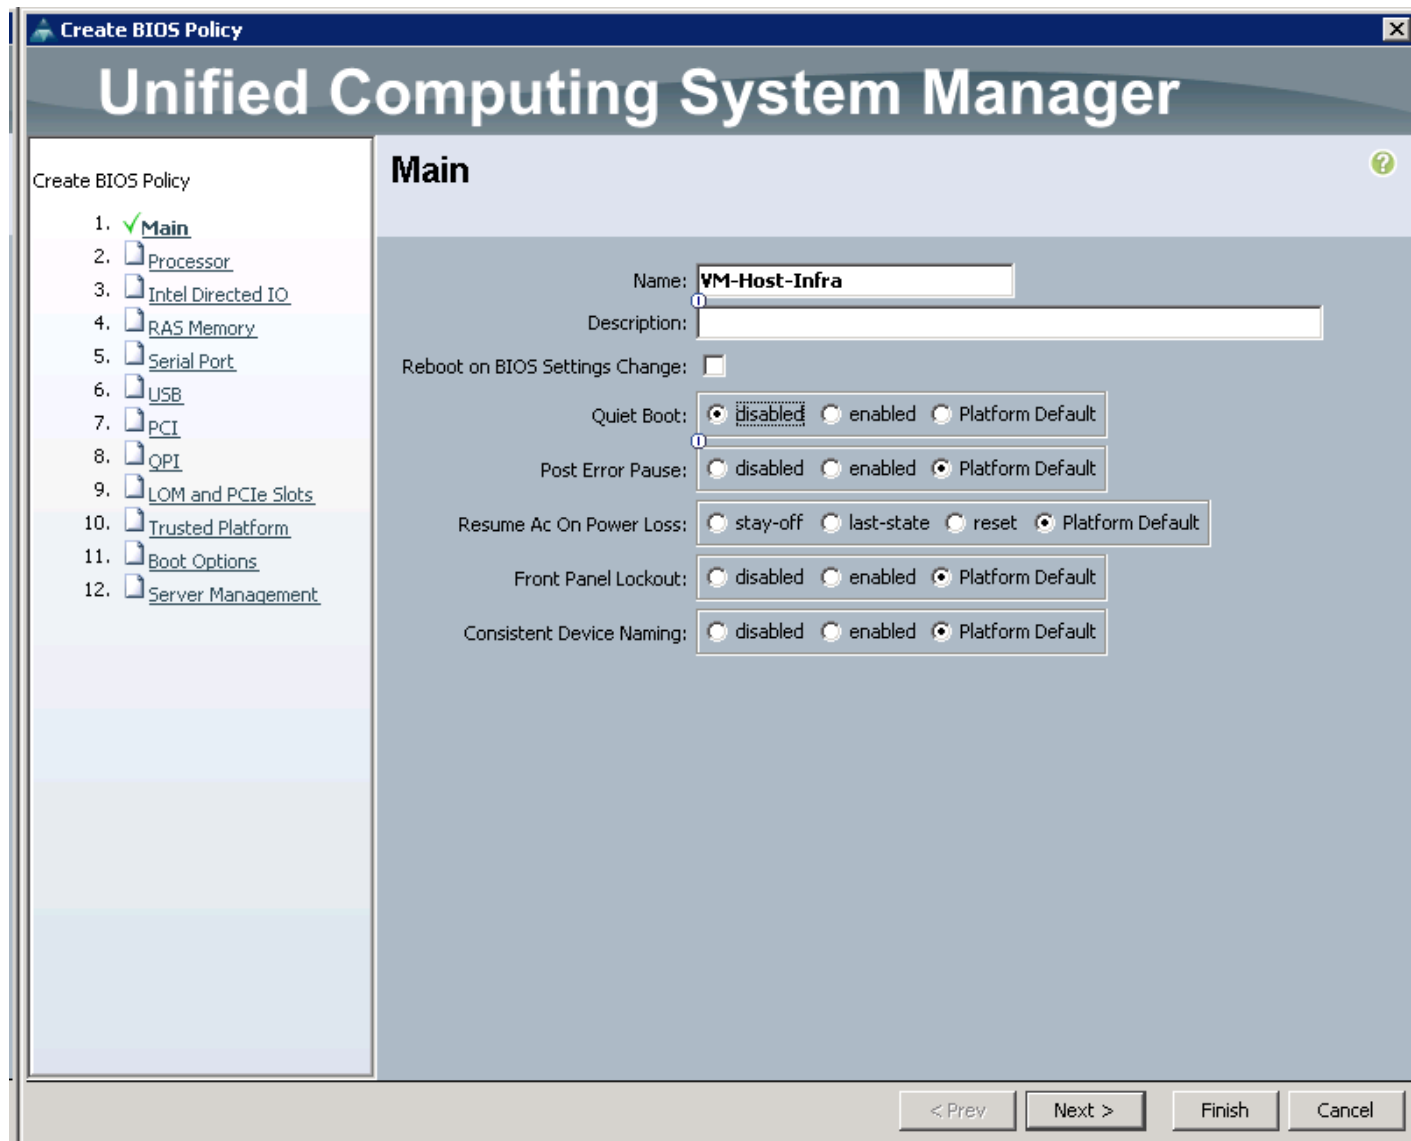
1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCSB-B200-M4.
6. Select Create Server PID Qualifications.
7. Select UCSB-B200-M4 as the name.
8. Click OK to create the server PID qualification.
9. Click OK to create the policy then OK for the confirmation.



Create Server BIOS Policy

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

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



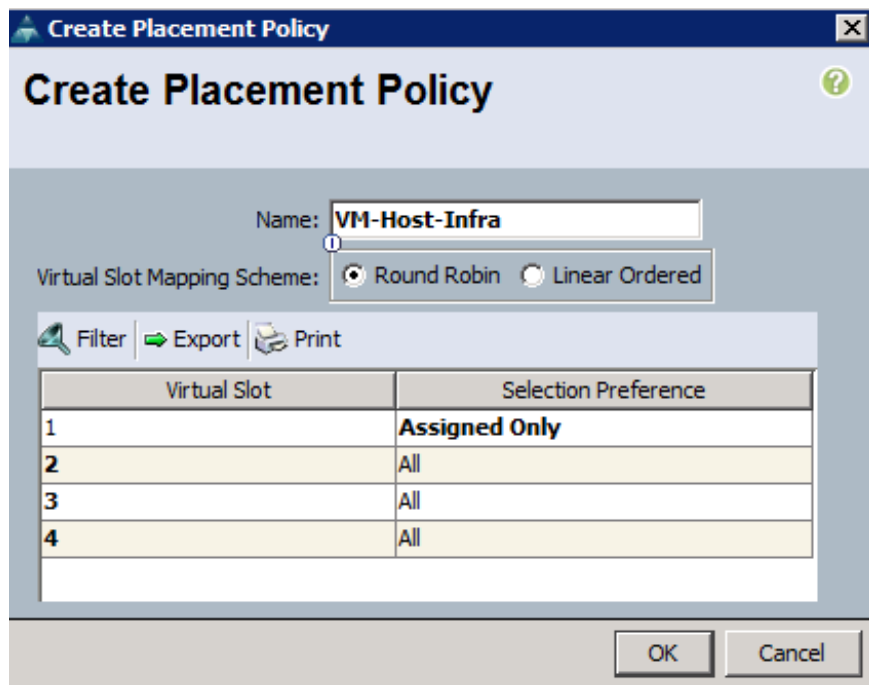
8. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

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

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

6. Click 1 and select Assigned Only.
7. Click OK, and then click OK again.



Create Placement Policy

Name:

Virtual Slot Mapping Scheme: Round Robin Linear Ordered

Filter | Export | Print

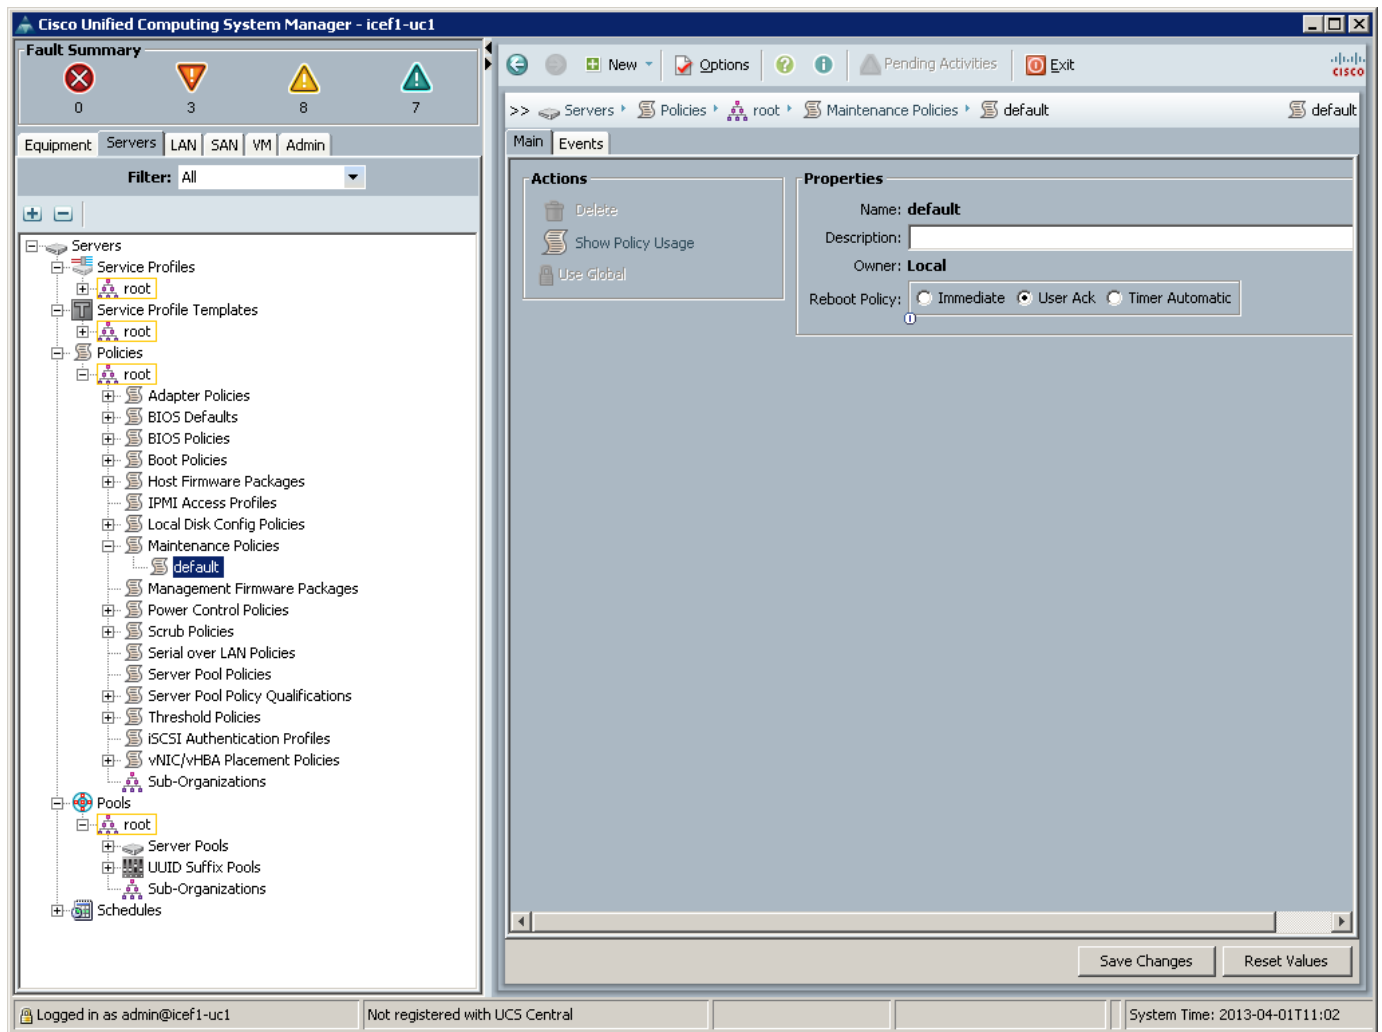
Virtual Slot	Selection Preference
1	Assigned Only
2	All
3	All
4	All

OK Cancel

Update the Default Maintenance Policy

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

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



Create vNIC Templates

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

Create Data vNICs

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.

8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.
10. Under VLANs, select the checkboxes for IB-MGMT, INFRA-NFS, Native-VLAN, VM-Traffic, and vMotion VLANs.
11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_A.
14. In the Network Control Policy list, select Enable_CDP.
15. Click OK to create the vNIC template.
16. Click OK.

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

Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	INFRA-NFS	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	VM-Traffic	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A-VLAN	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

+ Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy:

OK Cancel

17. In the navigation pane, select the LAN tab.

18. Select Policies > root.
19. Right-click vNIC Templates.
20. Select Create vNIC Template
21. Enter `vNIC_Template_B` as the vNIC template name.
22. Select Fabric B.
23. Do not select the Enable Failover checkbox.
24. Under Target, make sure the VM checkbox is not selected.
25. Select Updating Template as the template type.
26. Under VLANs, select the checkboxes for `IB-MGMT`, `INFRA-NFS`, `Native-VLAN`, and `vMotion VLANs`.
27. Set `default` as the native VLAN.
28. For MTU, enter `9000`.
29. In the MAC Pool list, select `MAC_Pool_B`.
30. In the Network Control Policy list, select `Enable_CDP`.
31. Click OK to create the vNIC template.
32. Click OK.

Create iSCSI vNICs

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `iSCSI_Template_A` as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.
9. Under VLANs, select `iSCSI-A-VLAN`.
10. Set `iSCSI-A-VLAN` as the native VLAN.

11. Under MTU, enter 9000.
12. From the MAC Pool list, select `MAC_Pool_A`.
13. From the Network Control Policy list, select `Enable_CDP`.
14. Click OK to complete creating the vNIC template.
15. Click OK.

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

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-Mgmt	<input type="radio"/>
<input type="checkbox"/>	INFRA-NFS	<input type="radio"/>
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>
<input type="checkbox"/>	VM-Traffic	<input type="radio"/>
<input checked="" type="checkbox"/>	iSCSI-A-VLAN	<input checked="" type="radio"/>

+ Create VLAN

MTU:

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy:

OK Cancel

16. Select the LAN tab on the left.

17. Select Policies > root.
18. Right-click vNIC Templates.
19. Select Create vNIC Template.
20. Enter `iscsi_Template_B` as the vNIC template name.
21. Select Fabric B. Do not select the Enable Failover checkbox.
22. Under Target, make sure that the VM checkbox is not selected.
23. Select Updating Template for Template Type.
24. Under VLANs, select `iscsi-B-VLAN`.
25. Set `iscsi-B-VLAN` as the native VLAN.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select `MAC_Pool_B`.
28. From the Network Control Policy list, select `Enable_CDP`.
29. Click OK to complete creating the vNIC template.
30. Click OK.

Create Boot Policies

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (`iscsi_lif01a` and `iscsi_lif01b`) and two iSCSI LIFs are on cluster node 2 (`iscsi_lif02a` and `iscsi_lif02b`). One boot policy is configured in this procedure. This policy configures the primary target to be `iscsi_lif01a`.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-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 option cleared.
8. Expand the Local Devices drop-down menu and select `Add Remote CD/DVD`.

9. Expand the iSCSI vNICs section and select Add iSCSI Boot.
10. In the Add iSCSI Boot dialog box, enter iSCSI-A-vNIC.
11. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
14. Click OK.
15. Expand CIMC Mounted vMedia.
16. Select Add CIMC Mounted CD/DVD.
17. Click OK.
18. Click OK to save the boot policy. Click OK to close the Boot Policy window.

The screenshot displays the configuration interface for a service profile template. The main area shows the **Properties** for a boot policy named **Boot-Fabric-A**. The **Owner** is set to **Local**. The **Reboot on Boot Order Change** checkbox is unchecked. The **Enforce vNIC/vHBA/iSCSI Name** checkbox is checked. The **Boot Mode** is set to **Legacy**.

A **Warning** message is displayed below the properties:

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 are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

The **Boot Order** table is visible at the bottom of the interface:

Name	Order	vNIC/vHBA/iSCSI v...	Type	Lun ID	WWN	Slot Number	Lun ID/NAME	Boot Name	Boot Path	Description
Remote CD/DVD	1									
iSCSI	2									
iSCSI		iSCSI-A-vNIC	Primary							
iSCSI		iSCSI-B-vNIC	Secondary							
CIMC Mounted CD/DV3										

Create Service Profile Template

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.
8. Click Next.

The screenshot shows the 'Create Service Profile Template' wizard in Cisco Unified Computing System Manager. The window title is 'Create Service Profile Template'. The main heading is 'Unified Computing System Manager'. The left sidebar shows a list of steps: 1. Identify Service Profile Template (checked), 2. Storage Provisioning, 3. Networking, 4. SAN Connectivity, 5. Zoning, 6. vNIC/vHBA Placement, 7. vMedia Policy, 8. Server Boot Order, 9. Maintenance Policy, 10. Server Assignment, and 11. Operational Policies.

The main content area is titled 'Identify Service Profile Template'. It contains the following text and form elements:

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.

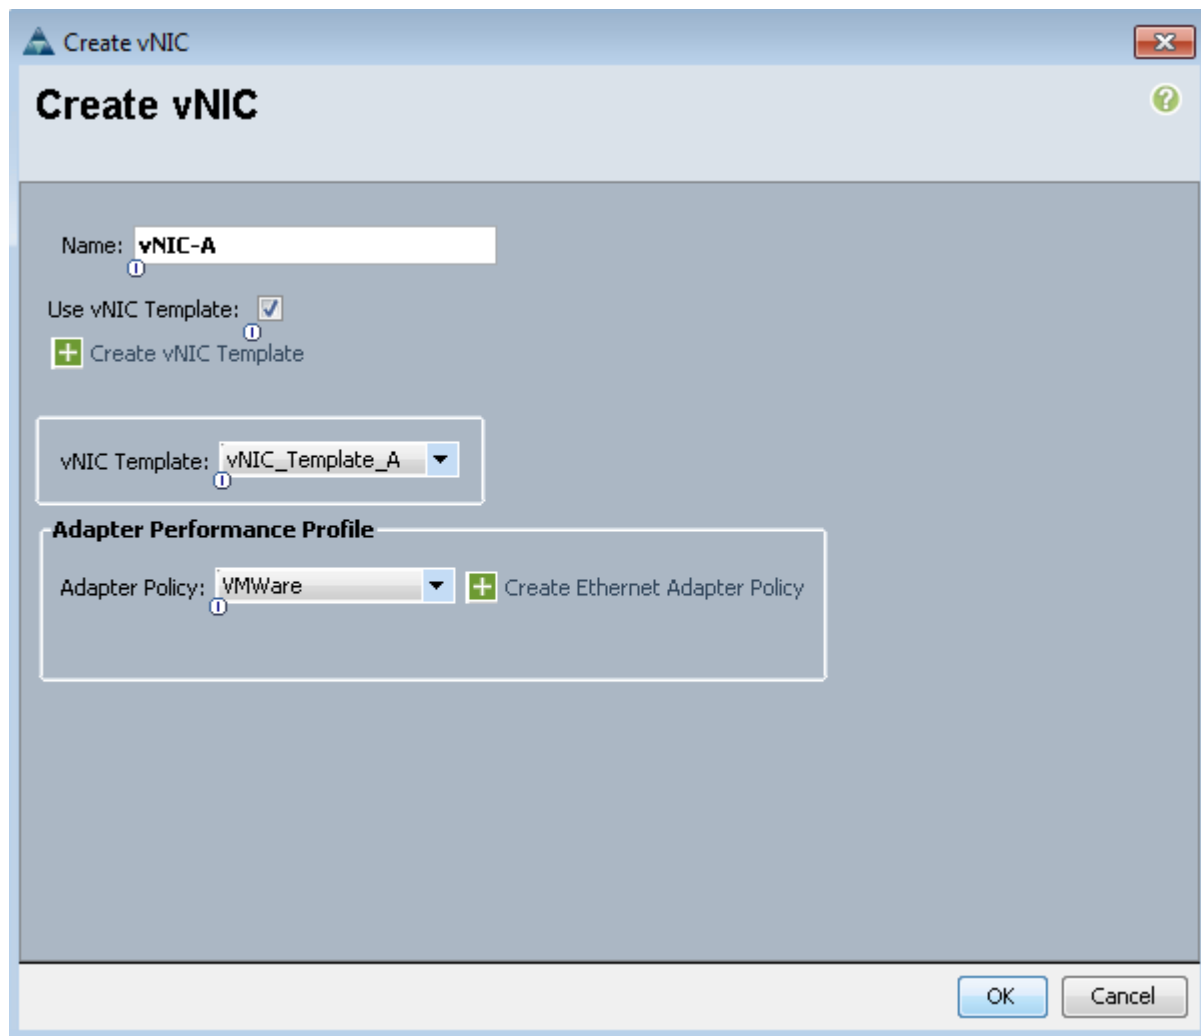
At the bottom right, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

Configure Storage Provisioning

1. If you have servers with no physical disks, select the iSCSI-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

Configure Networking Options

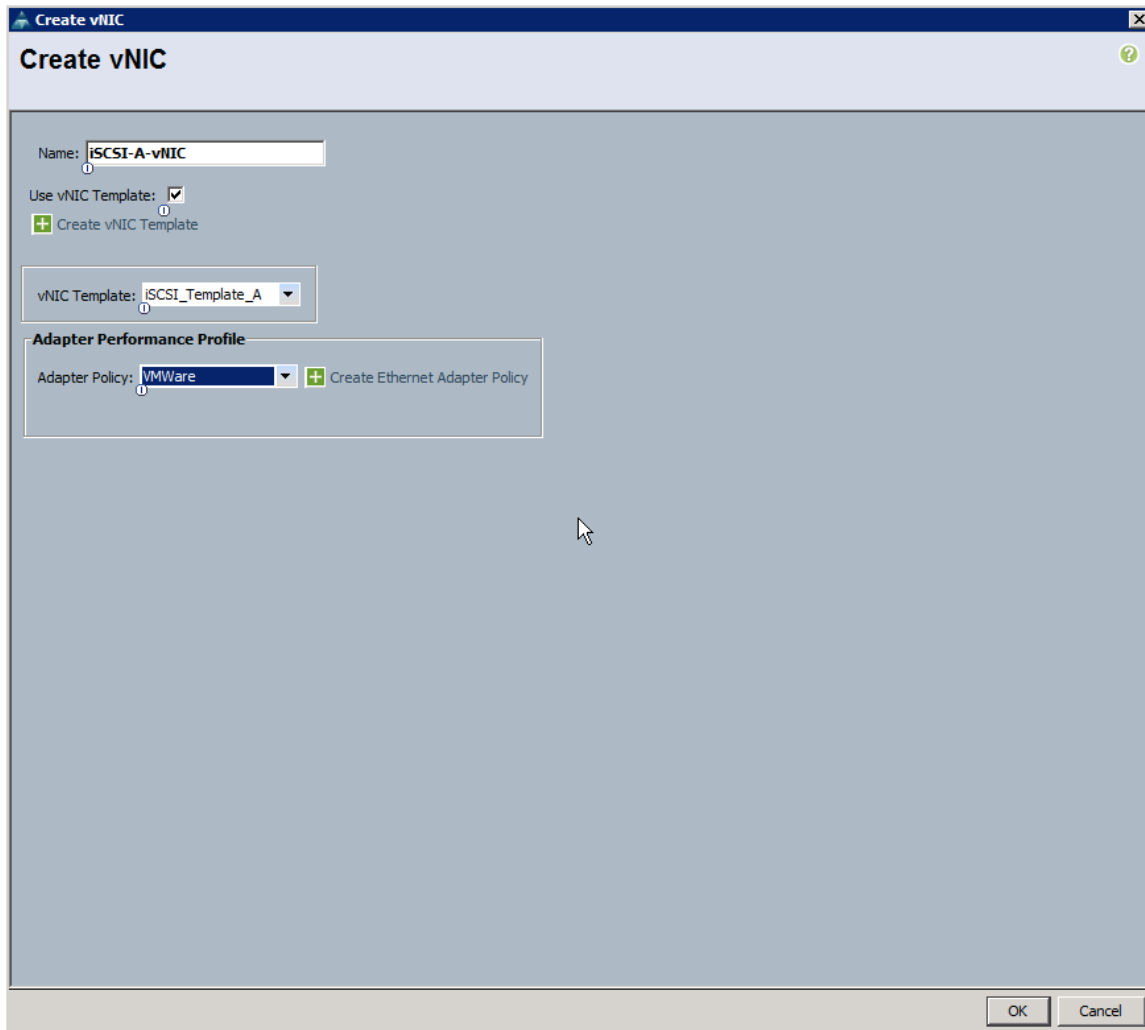
1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Expert” option to configure the LAN connectivity.
3. Click the upper Add button to add a vNIC to the template.
4. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
5. Select the Use vNIC Template checkbox.
6. In the vNIC Template list, select vNIC_Template_A.
7. In the Adapter Policy list, select VMWare.
8. Click OK to add this vNIC to the template.



9. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.
10. In the Create vNIC box, enter vNIC-B as the name of the vNIC.
11. Select the Use vNIC Template checkbox.
12. In the vNIC Template list, select vNIC_Template_B.
13. In the Adapter Policy list, select VMWare.
14. Click OK to add the vNIC to the template.
15. Click the upper Add button to add a vNIC to the template.
16. In the Create vNIC dialog box, enter iSCSI-A-vNIC as the name of the vNIC.
17. Select the Use vNIC Template checkbox.
18. In the vNIC Template list, select iSCSI_Template_A.

19. In the Adapter Policy list, select VMWare.

20. Click OK to add this vNIC to the template.



21. Click the upper Add button to add a vNIC to the template.

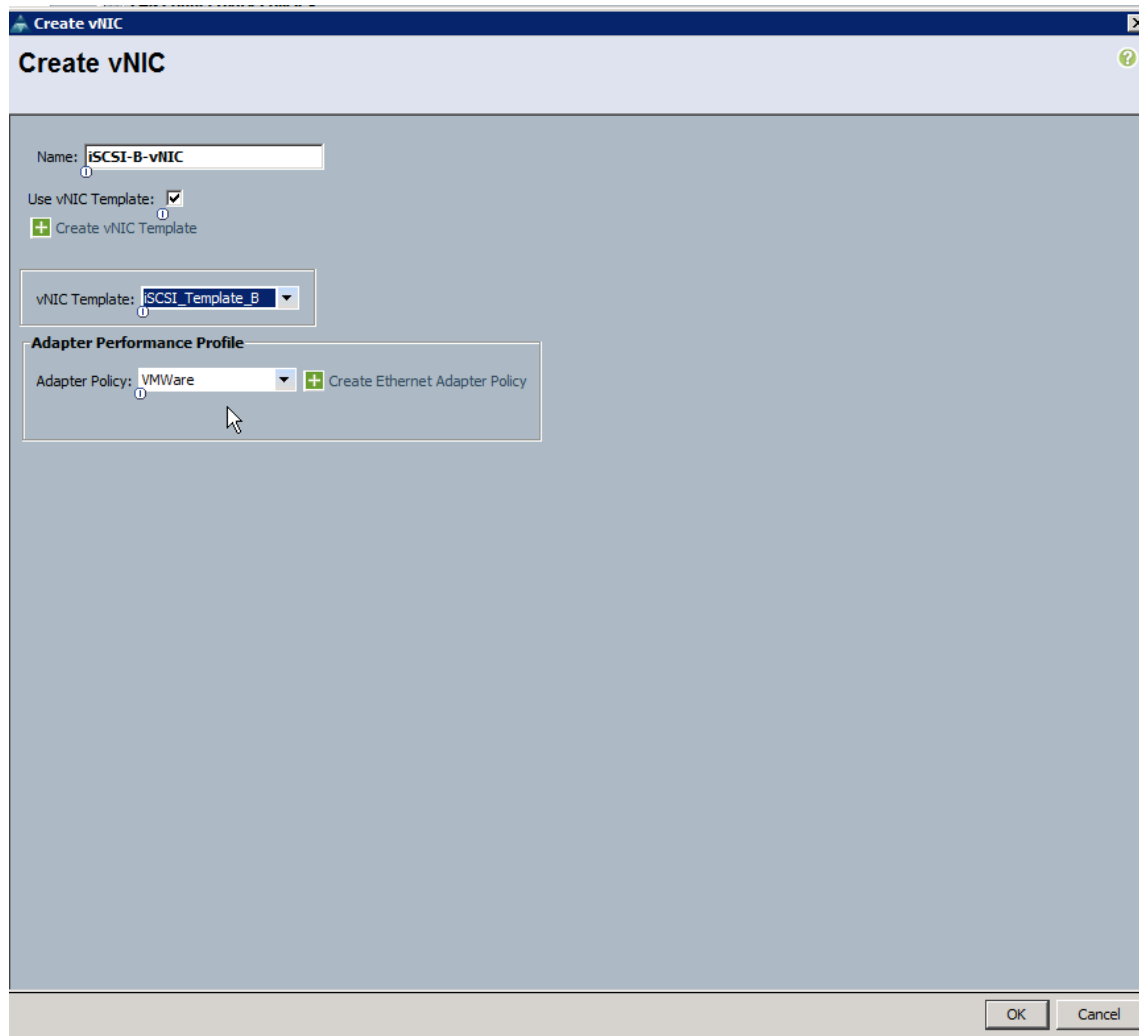
22. In the Create vNIC dialog box, enter iSCSI-B-vNIC as the name of the vNIC.

23. Select the Use vNIC Template checkbox.

24. In the vNIC Template list, select iSCSI_Template_B.

25. In the Adapter Policy list, select VMWare.

26. Click OK to add this vNIC to the template.



27. Expand the iSCSI vNICs section (if not already expanded).
28. Select `IQN-Pool` under “Initiator Name Assignment.”
29. Click the **lower** Add button in the iSCSI vNIC section to define a vNIC.
30. Enter `iSCSI-A-vNIC` as the name of the vNIC.
31. Select `iSCSI-A-vNIC` for Overlay vNIC.
32. Set the iSCSI Adapter Policy to `default`.
33. Set the VLAN to `iSCSI-A-VLAN`.
34. Leave the MAC Address set to `None`.
35. Click OK.

Create iSCSI vNIC

Name:

Overlay vNIC:

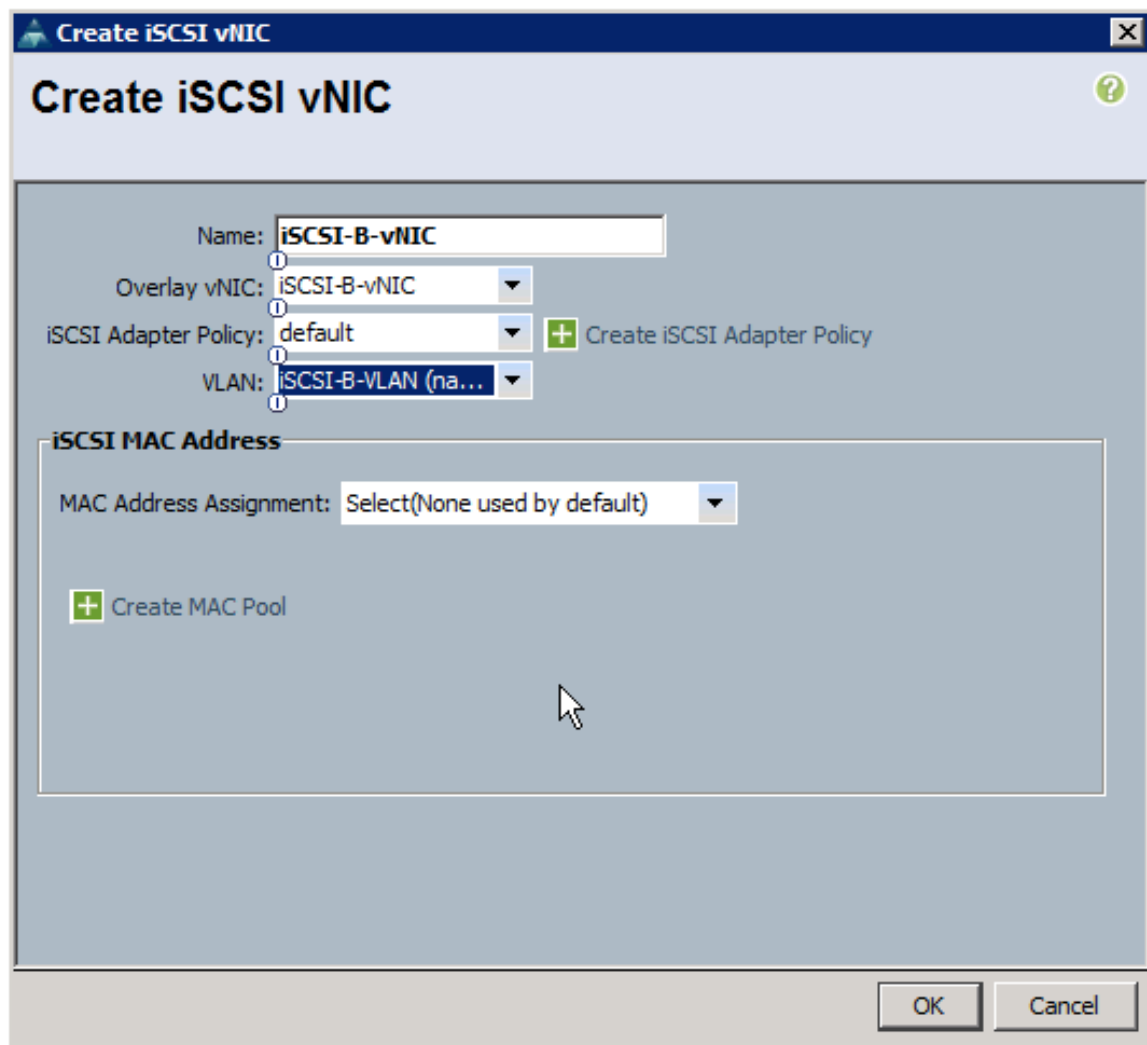
iSCSI Adapter Policy:

VLAN:

iSCSI MAC Address

MAC Address Assignment:

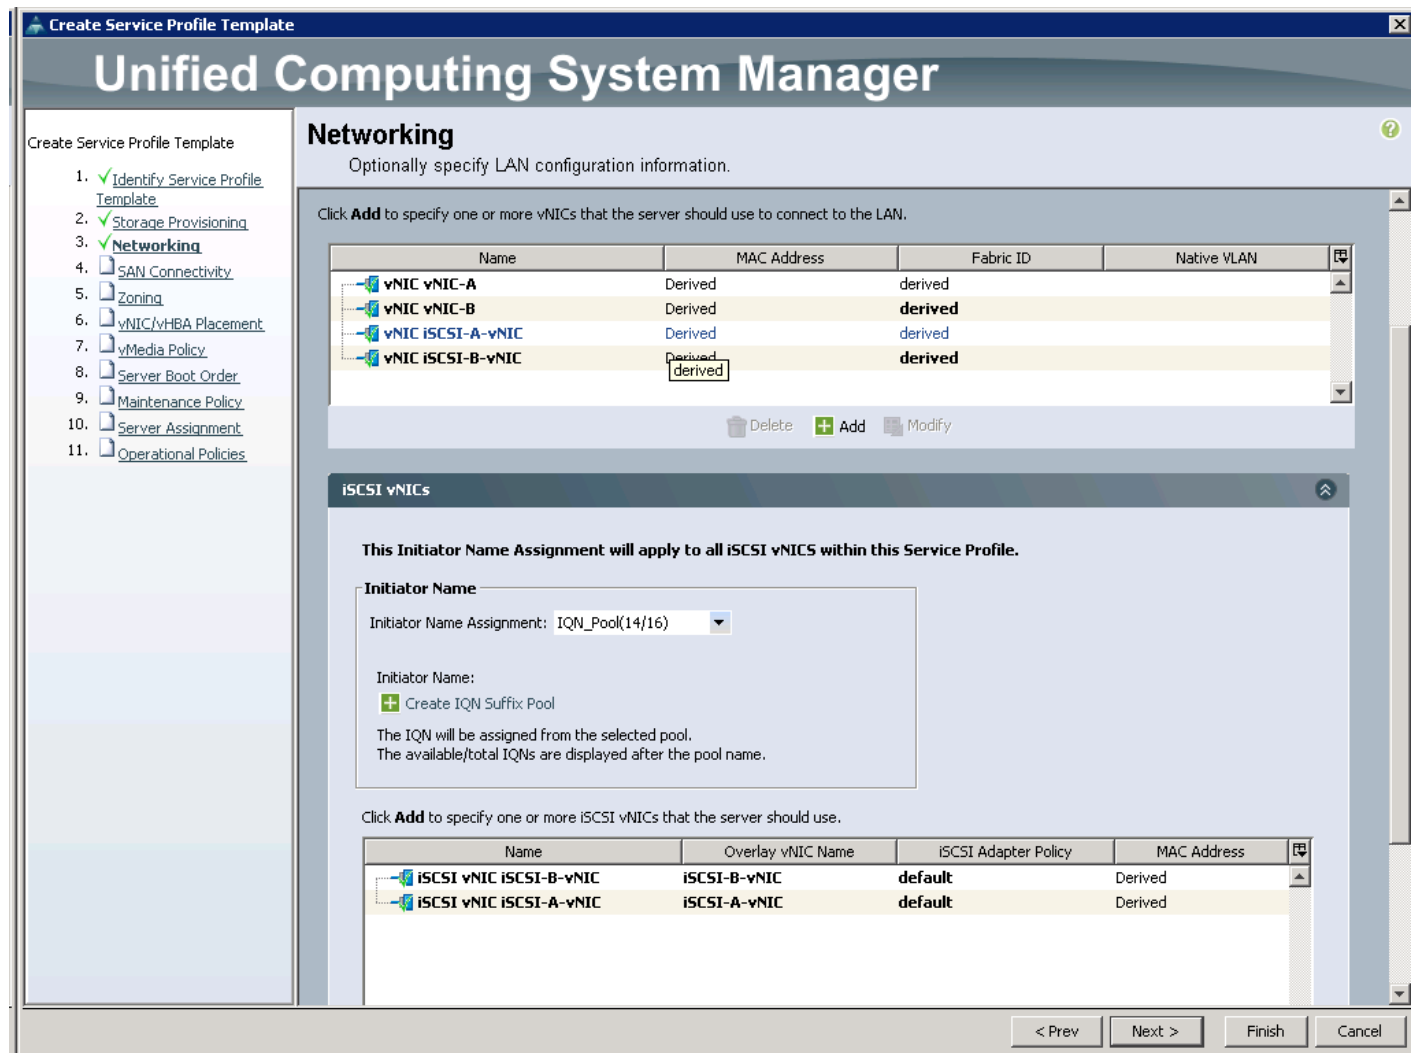
36. Click the lower Add button in the iSCSI vNIC section to define a vNIC.
37. Enter `iSCSI-B-vNIC` as the name of the vNIC.
38. Set the Overlay vNIC to `iSCSI-B-vNIC`
39. Set the iSCSI Adapter Policy to `default`.
40. Set the VLAN to `iSCSI-B-VLAN`
41. Leave the MAC Address set to `None`.
42. Click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



43. Click OK.

44. Review the table in the Networking page to make sure that all vNICs were created.

45. Click Next.



Configure Storage Options

- Select the **No vHBAs** option for the “How would you like to configure SAN connectivity?” field.
- Click **Next**.

Configure Zoning Options

- Set no Zoning options and click **Next**.

Configure vNIC/HBA Placement

- In the “Select Placement” list, select the **VM-Host-Infra** placement policy.
- Select **vCon1** and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - vNIC-A
 - vNIC-B
 - iSCSI-vNIC-A


d. iSCSI-vNIC-B

3. Review the table to verify that all vNICs were assigned to the policy in the appropriate order.
4. Click Next.

Configure vMedia Policy

1. Do not configure a vMedia Policy at this time.
2. Click Next.

Configure Server Boot Order

1. Select `Boot-Fabric-A` for Boot Policy.
2. In the Boot Order pane, select `iSCSI-A-vNIC`.
3. Click the “Set iSCSI Boot Parameters” button.
4. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps
5. Set `iSCSI_IP_Pool_A` as the “Initiator IP address Policy”.
6. Keep the “iSCSI Static Target Interface” button selected and click the  button at the bottom right.
7. Log in to the storage cluster management interface and run the following command:

```
iscsi show
```
8. Note or copy the iSCSI target name for `Infra-SVM`.
9. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from `Infra-SVM`.
10. Enter the IP address of `iSCSI_lif02a` for the IPv4 Address field.

Create iSCSI Static Target

iSCSI Target Name:


Priority:

Port:

Authentication Profile: + Create iSCSI Authentication Profile

IPv4 Address:

LUN ID:

11. Click OK to add the iSCSI static target.
12. Keep the iSCSI Static Target Interface option selected and click the  button.
13. In the Create iSCSI Static Target window, paste the iSCSI target node name from `Infra-SVM` into the iSCSI Target Name field.
14. Enter the IP address of `iscsi_lif01a` in the IPv4 Address field.
15. Click OK.

Set iSCSI Boot Parameters
✕

Set iSCSI Boot Parameters ?

Name: **iSCSI-A-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set>

+ Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_A(14/16)

IPv4 Address: **0.0.0.0**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**

+ Create IP Pool

The IP address will be automatically assigned from the selected pool.


iSCSI Static Target Interface
 iSCSI Auto Target Interface

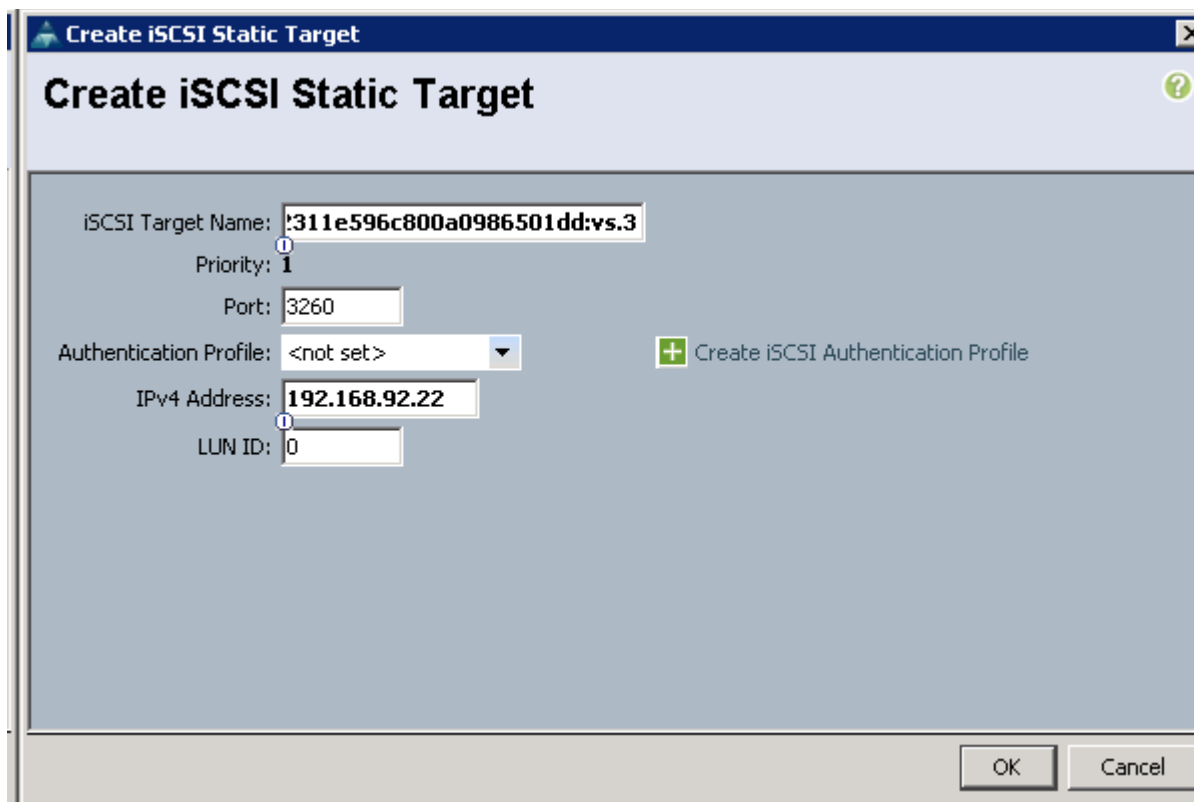
Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	IPv4 Address	LUN Id	⌵
iqn.1992-08.c...	2	3260		192.168.91.21	0	▲
iqn.1992-08.c...	1	3260		192.168.91.22	0	▲

OK
Cancel

16. Click OK.

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




Create iSCSI Static Target

iSCSI Target Name: `311e596c800a0986501dd:vs.3`

Priority: `1`


Port: `3260`

Authentication Profile: `<not set>`  Create iSCSI Authentication Profile

IPv4 Address: `192.168.92.22`

LUN ID: `0`

OK Cancel

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

Set iSCSI Boot Parameters
✕

Set iSCSI Boot Parameters ?

Name: **iSCSI-B-vNIC**

Authentication Profile: <not set> + Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set> ▼

+ Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_B(14/16) ▼

IPv4 Address: **0.0.0.0**

Subnet Mask: **255.255.255.0**

Default Gateway: **0.0.0.0**

Primary DNS: **0.0.0.0**

Secondary DNS: **0.0.0.0**

+ Create IP Pool

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface
 iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	IPv4 Address	LUN Id	
iqn.1992-08.c...	2	3260		192.168.92.21	0	⬆
iqn.1992-08.c...	1	3260		192.168.92.22	0	⬆

+
-
✖
📄
⬇

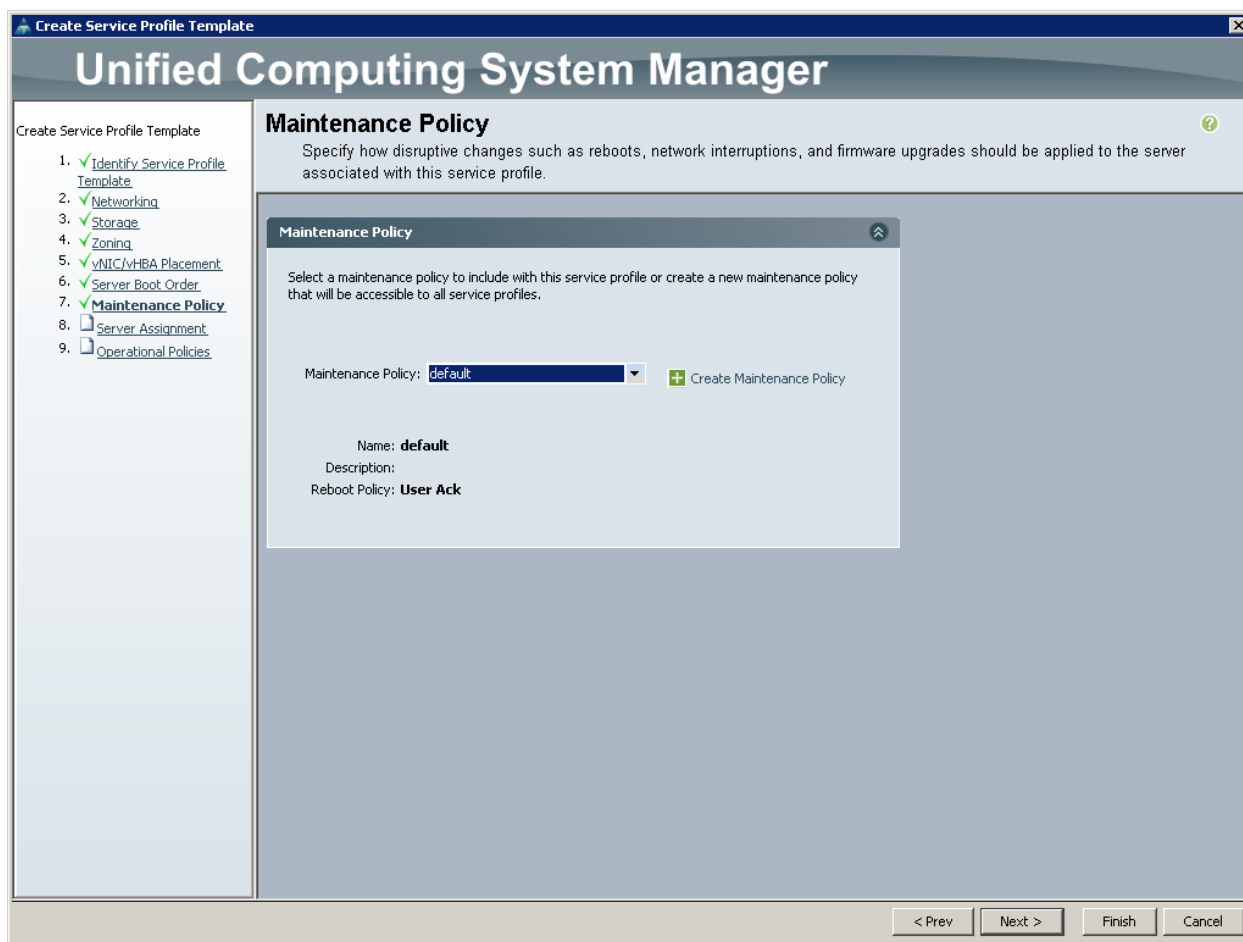
OK
Cancel

29. Click OK.

30. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
31. Click Next to continue to the next section.

Configure Maintenance Policy

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

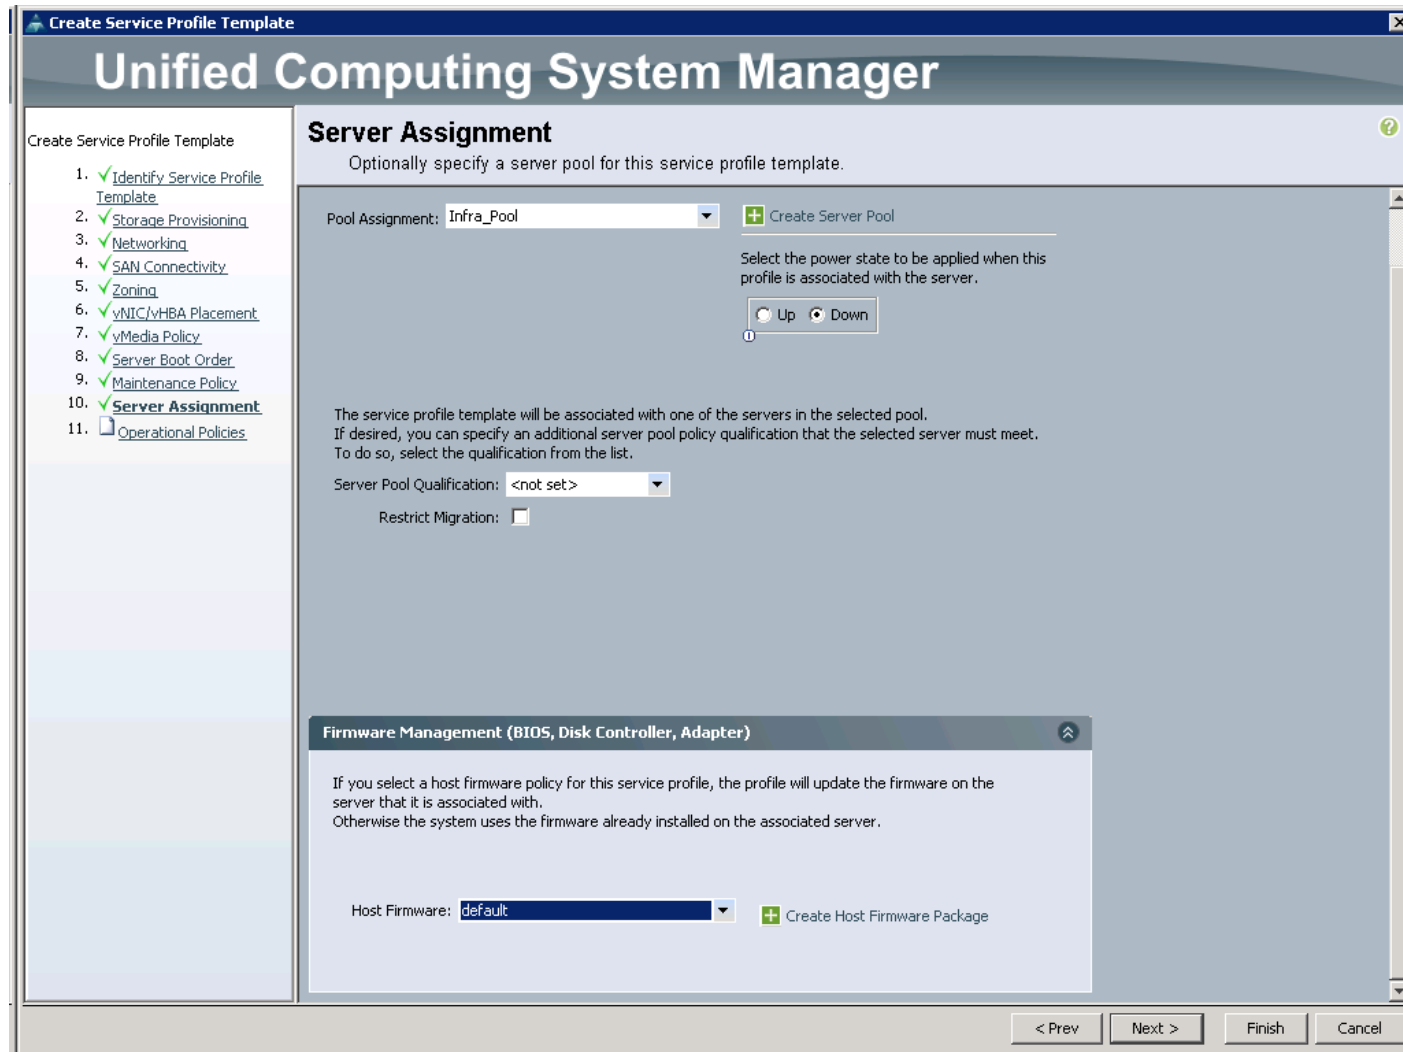


Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select `Infra_Pool`.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management at the bottom of the page and select `default` from the Host Firmware list.

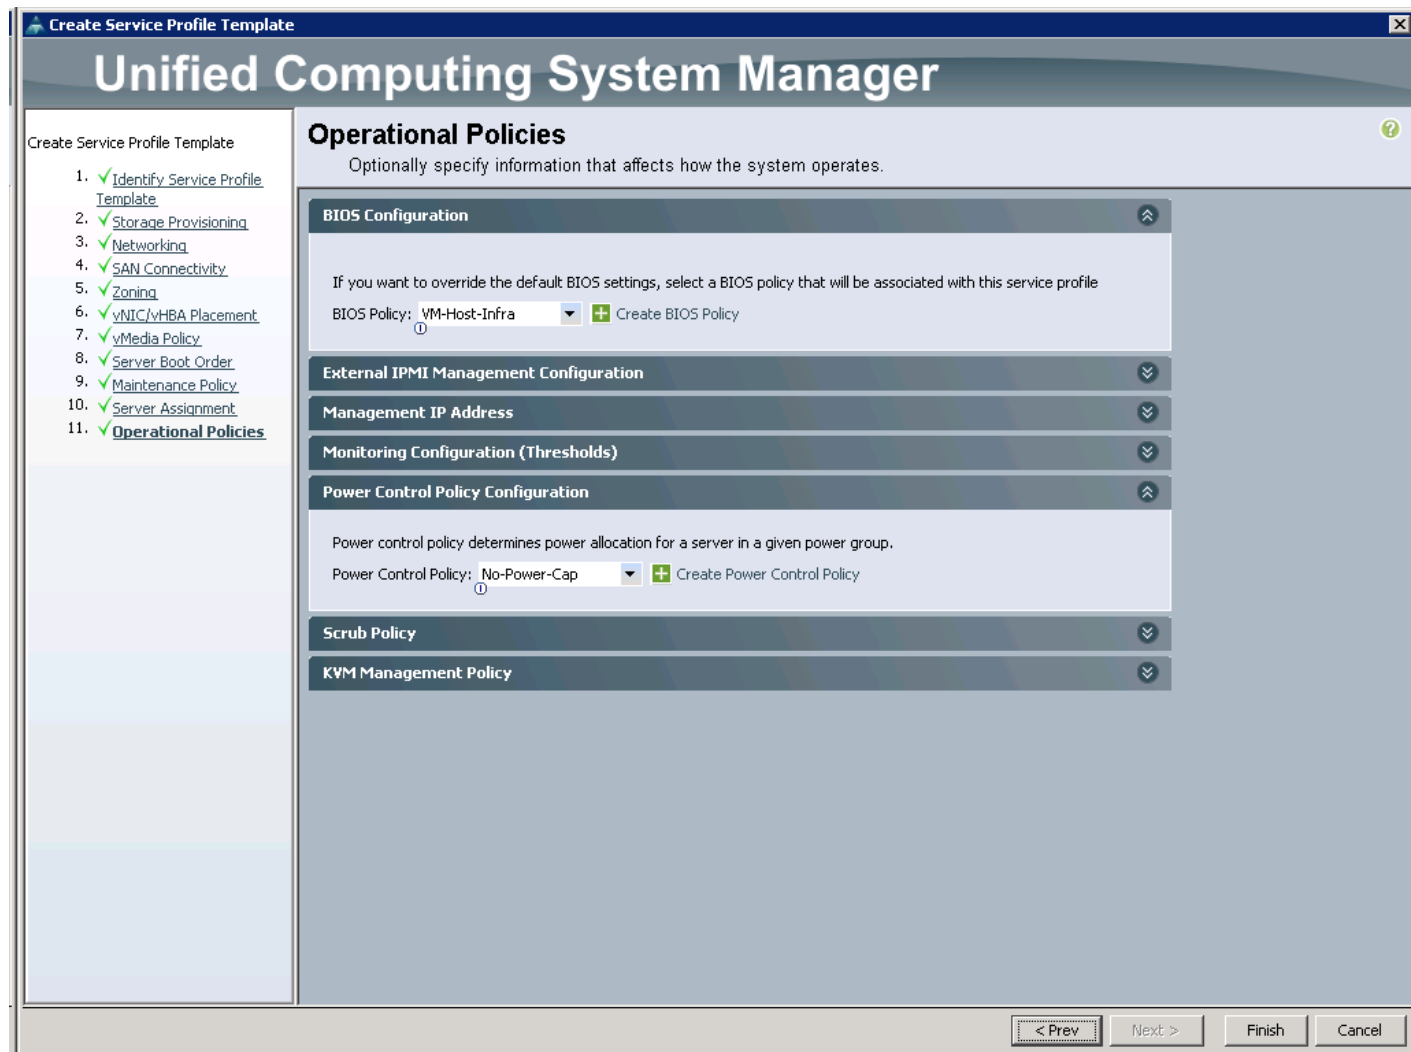
5. Click Next.



Configure Operational Policies

To configure the operational policies, complete the following steps:

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



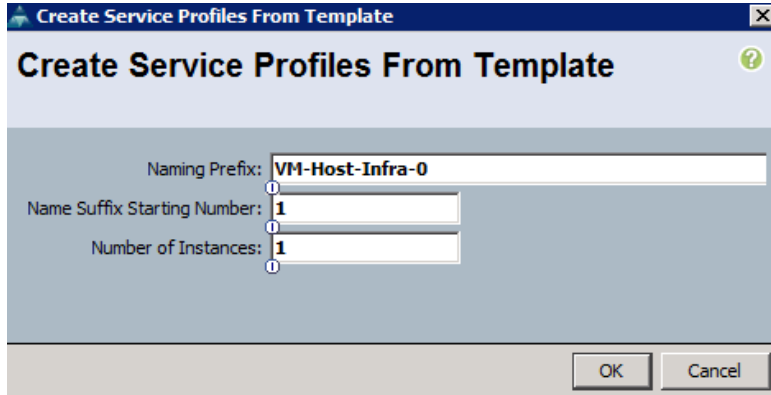
3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Service Profiles

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

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

- Click OK to create the service profile.



- Click OK in the confirmation message.

Add More Servers to FlexPod Unit

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

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into Table 17 and Table 18 .

Table 17 iSCSI LIFs for iSCSI IQN

Vserver	iSCSI Target IQN
Infra-SVM	



To gather the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface. For 7-Mode storage, run the `iscsi nodename` command on each storage controller.

Table 18 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN	Variables
VM-Host-Infra-01		<< var_vm_host_infra_01_iqn >>
VM-Host-Infra-02		<< var_vm_host_infra_02_iqn >>



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

Storage Configuration – iSCSI Boot

Clustered Data ONTAP iSCSI Boot Storage Setup

Create igroups

To create igroups, complete the following steps:

1. From the cluster management node SSH connection, enter the following:

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol
iscsi -ostype vmware -initiator <<var_vm_host_infra_01_iqn>>

igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol
iscsi -ostype vmware -initiator <<var_vm_host_infra_02_iqn>>

igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_01_iqn>>,
<<var_vm_host_infra_02_iqn>>
```



Use the values listed in Table 17 and Table 18 for the IQN information.



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

Map Boot LUNs to igroups

1. From the storage cluster management SSH connection, enter the following:

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

VMware vSphere 6.0 Setup

VMware ESXi 6.0

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

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

Download Cisco Custom Image for ESXi 6.0

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



This ESXi 6.0 Cisco custom image includes updates for the fnic and eNIC drivers. The versions that are part of this image are: eNIC: 2.1.2.59; fnic: 1.6.0.12

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, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. To download the Cisco UCS Manager software, click the Launch UCS Manager link.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click the Servers tab.

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

Set Up VMware ESXi Installation

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



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

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

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.
8. Boot the server by selecting Boot Server and clicking OK. Then click OK again.

Install ESXi

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

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

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

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



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

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

Set Up Management Networking for ESXi Hosts

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

ESXi Host VM-Host-Infra-01

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

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the `<<var_ib_mgmt_vlan_id>>` and press Enter.
6. Select Network Adapters option and select `vmnic04` and press Enter.
7. From the Configure Management Network menu, select IP Configuration and press Enter.
8. Select the Set Static IP Address and Network Configuration option by using the space bar.
9. Enter the IP address for managing the first ESXi host: `<<var_vm_host_infra_01_ip>>`.
10. Enter the subnet mask for the first ESXi host.
11. Enter the default gateway for the first ESXi host.
12. Press Enter to accept the changes to the IP configuration.
13. Select the IPv6 Configuration option and press Enter.

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



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

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

ESXi Host VM-Host-Infra-02

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

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root` and enter the corresponding password.
3. Select the `Configure the Management Network` option and press Enter.
4. Select the `VLAN (Optional)` option and press Enter.
5. Enter the `<<var_ib-mgmt_vlan_id>>` and press Enter.
6. Select `Network Adapters` option and select `vmnic4` (defined earlier as OOB vNIC) and press Enter.
7. From the `Configure Management Network` menu, select `IP Configuration` and press Enter.
8. Select the `Set Static IP Address and Network Configuration` option by using the space bar.
9. Enter the IP address for managing the second ESXi host: `<<var_vm_host_infra_02_ip>>`.

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



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

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

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the `VM-Host-Infra-01` management IP address.
2. Download and install the vSphere Client.



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

Download VMware vSphere CLI 6.0

1. Click the following link [VMware vSphere CLI 6.0](#)
2. Select your OS and Click **Download**.
3. Save it to your destination folder.
4. Run the VMware-vSphere-CLI-6.0.0-2503617.exe
5. Click Next.
6. Accept the terms for the license and click **Next**.
7. Click **Next** on the Destination Folder screen.
8. Click Instal.
9. Click Finish.



Note: Install VMware vSphere CLI 6.0 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, complete the following steps:

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

ESXi Host VM-Host-Infra-02

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

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-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.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

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

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of vSwitch0, click Properties.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to 9000.
7. Click OK.
8. Select the Management Network configuration and click Edit.
9. Change the network label to VMkernel-MGMT and select the Management Traffic checkbox.
10. Click OK to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to MGMT Network and enter <<var_ib-mgmt_vlan_id>> in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.
14. Click Close.
15. On the right side of iScsiBootvSwitch, click Properties.
16. Select the vSwitch configuration and click Edit.
17. Change the MTU to 9000.
18. Click OK.
19. Select iScsiBootPG and click Edit.
20. Change the Network Label to VMkernel-iSCSI-A.
21. Change the MTU to 9000.
22. Click OK.

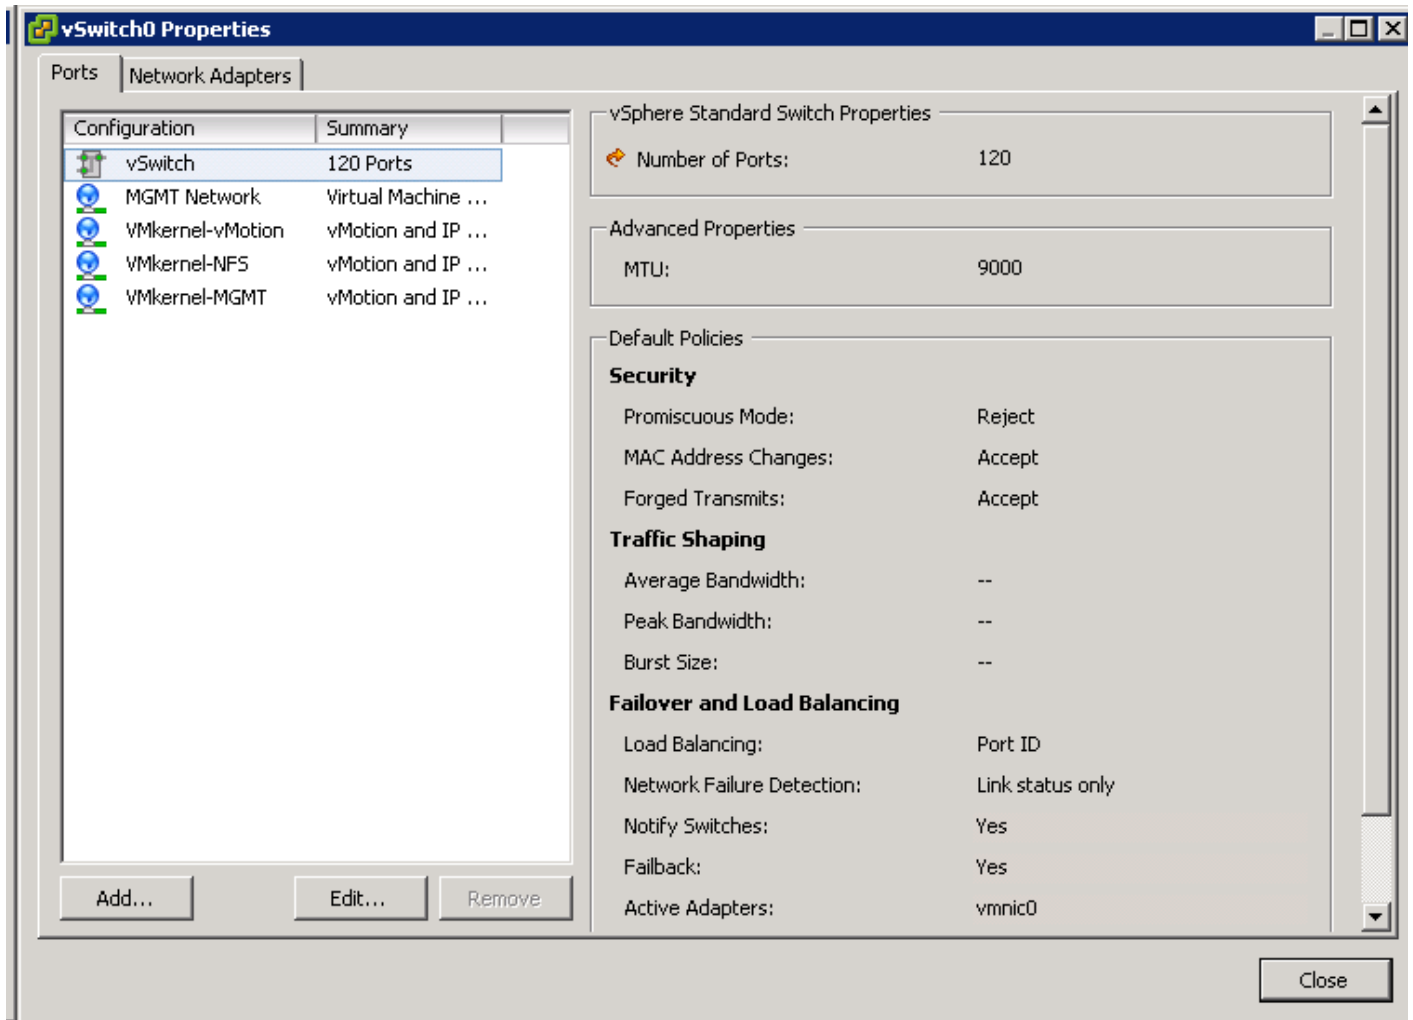
23. Click Close.
24. In the vSphere Standard Switch view, click Add Networking.
25. Select VMkernel and click Next.
26. Select Create a vSphere standard switch to create a new vSphere standard switch.
27. Select the check boxes for the network adapter vmnic3.
28. Click Next.
29. Change the network label to `VMkernel-iSCSI-B`.
30. Click Next.
31. Enter the IP address and the subnet mask for the iSCSI VLAN B interface for `VM-Host-Infra-01`.



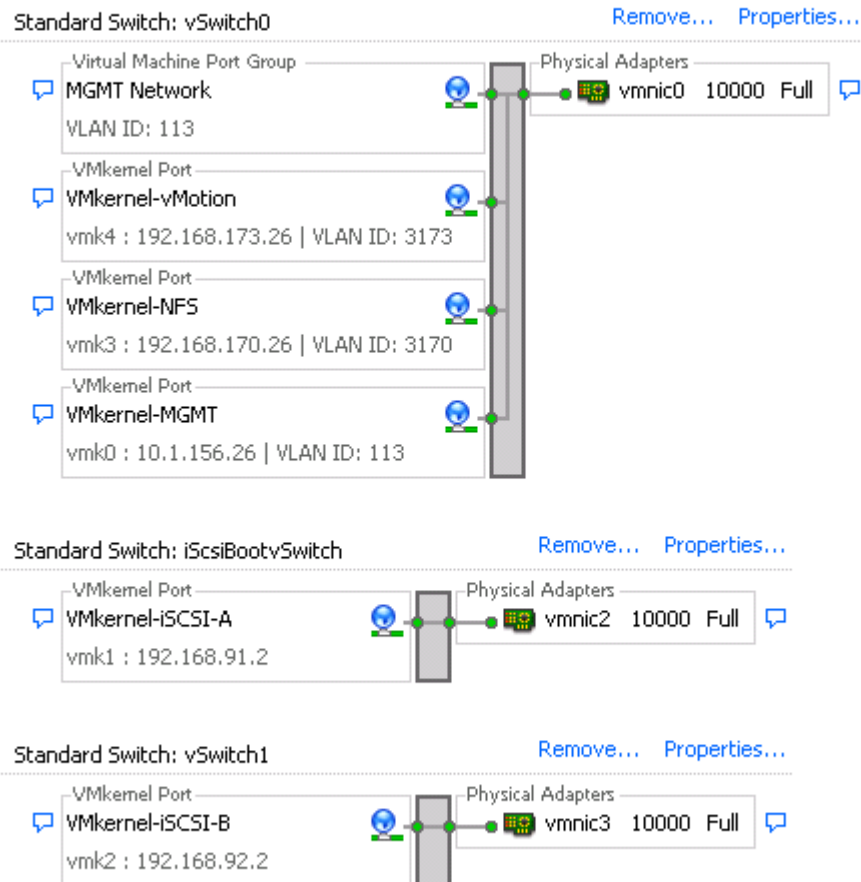
To obtain the iSCSI IP address information; login to the Cisco UCS Manager, in the servers tab select the corresponding service profiles. In the right pane, click the boot order and select the iSCSI-B-vNIC; click set iSCSI boot parameters; the IP address should appear as the initiator IP address.

32. Click Next.
33. Click Finish.
34. On the right side of `vSwitch1`, click Properties.
35. Select the vSwitch configuration and click Edit.
36. Change the MTU to 9000.
37. Click OK.
38. Select `VMkernel-iSCSI-B` and click Edit.
39. Change the MTU to 9000.
40. Click OK.
41. Click Close.
42. On the right side of `vSwitch0`, click Properties.
43. Click Add.
44. Change the network label to `VMkernel-NFS` and enter `<<var_nfs_vlan_id>>` in the VLAN ID (Optional) field.
45. Click Next.

46. Enter the IP address <<var_nfs_vlan_ip_host_01>> and the subnet mask <<var_nfs_vlan_ip_mask_host_01>> for the NFS VLAN interface for VM-Host-Infra-01.
47. To continue with the NFS VMkernel creation, click Next.
48. To finalize the creation of the NFS VMkernel interface, click Finish.
49. Select the VMkernel-NFS configuration and click Edit.
50. Change the MTU to 9000.
51. Click OK to finalize the edits for the VMkernel-NFS network.
52. Click Add.
53. Change the network label to VMkernel-vMotion and enter <<var_vmotion_vlan_id>> in the VLAN ID (Optional) field.
54. Click Next.
55. Enter the IP address <<var_vmotion_vlan_ip_host_01>> and the subnet mask <<var_vmotion_vlan_ip_mask_host_01>> for the vMotion VLAN interface for VM-Host-Infra-01.
56. To continue with the vMotion VMkernel creation, click Next.
57. To finalize the creation of the vMotion VMkernel interface, click Finish.
58. Select the VMkernel-vMotion configuration and click Edit.
59. Change the MTU to 9000.
60. Click OK to finalize the edits for the VMkernel-vMotion network.
61. The properties for vSwitch0 should be similar to the following example:



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

View: vSphere Standard Switch vSphere Distributed Switch**Networking**

ESXi Host VM-Host-Infra-02

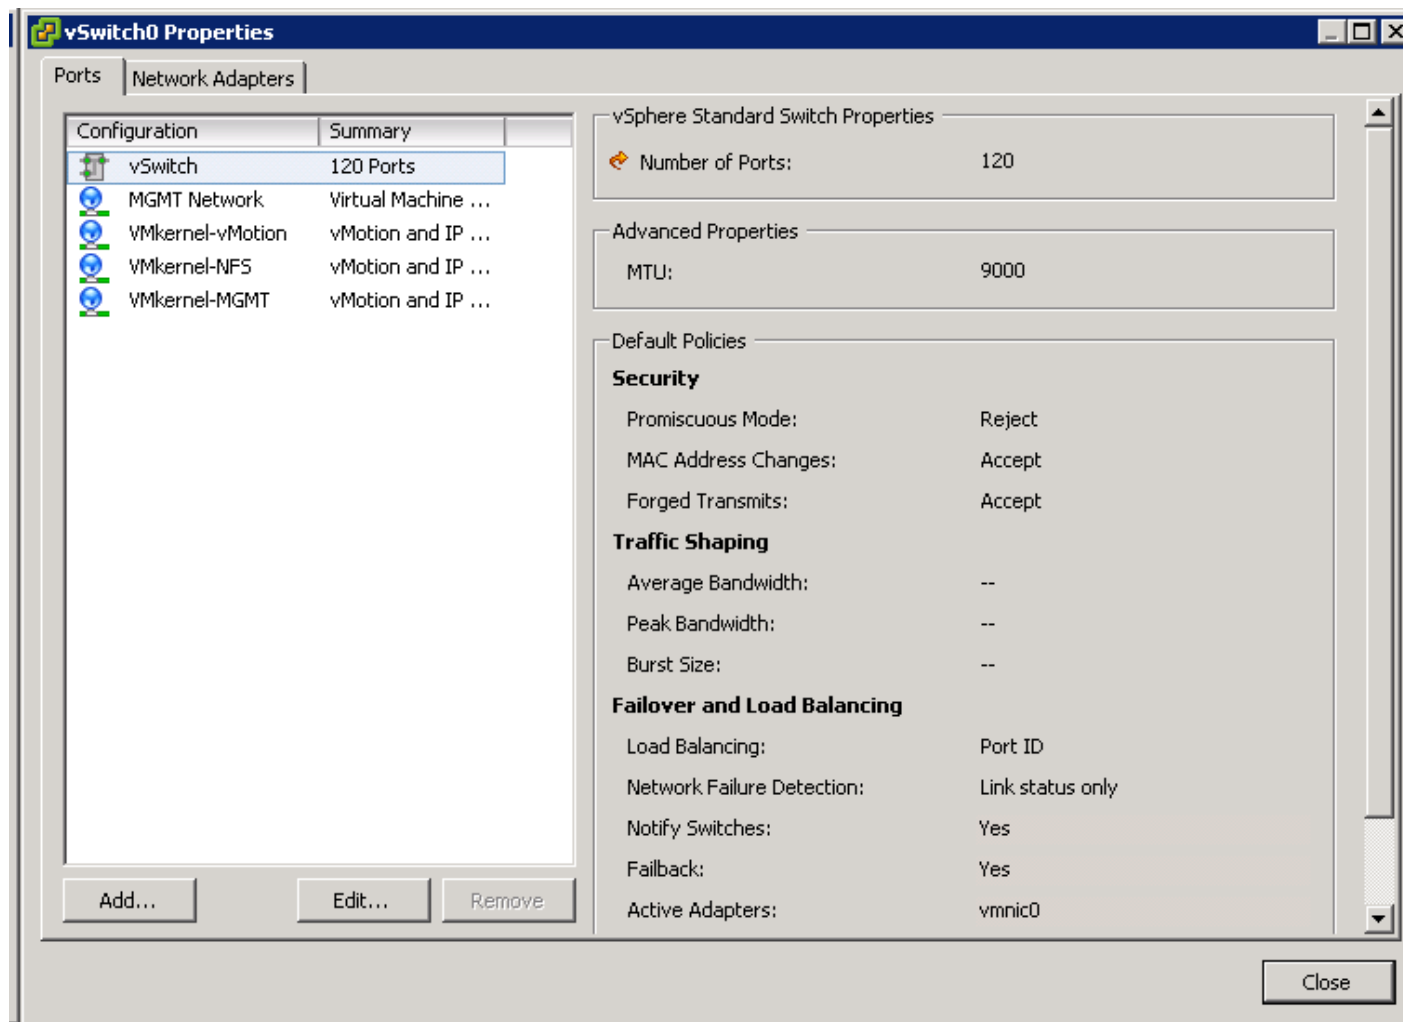
To set up the VMkernel ports and the virtual switches on the VM-Host-Infra-02. ESXi host, complete the following steps:

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

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

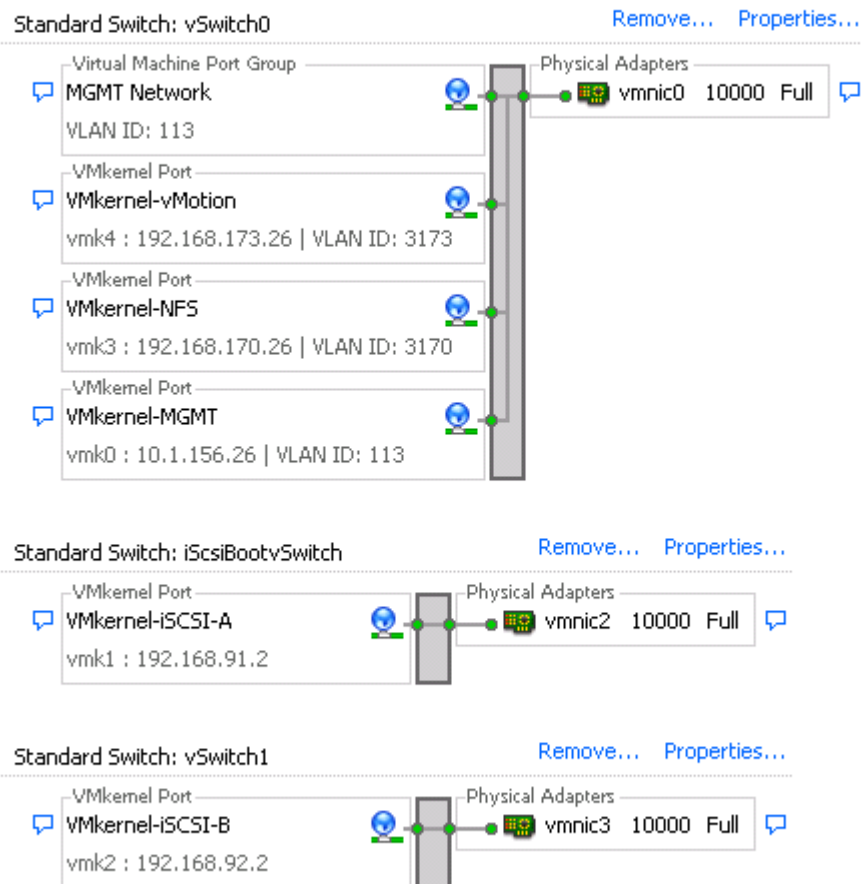
32. Click Next.
33. Click Finish.
34. On the right side of vSwitch1, click Properties.
35. Select the vSwitch configuration and click Edit.
36. Change the MTU to 9000.
37. Click OK.
38. Select VMkernel-iSCSI-B and click Edit.
39. Change the MTU to 9000.
40. Click OK.
41. Click Close.
42. On the right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

56. Enter the IP address <<var_vmotion_vlan_ip_host_02>> and the subnet mask <<var_vmotion_vlan_ip_mask_host_02>> for the vMotion VLAN interface for VM-Host-Infra-02.
57. To continue with the vMotion VMkernel creation, click Next.
58. To finalize the creation of the vMotion VMkernel interface, click Finish.
59. Select the VMkernel-vMotion configuration and click Edit.
60. Change the MTU to 9000.
61. Click OK to finalize the edits for the VMkernel-vMotion network.



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

View: vSphere Standard Switch vSphere Distributed Switch

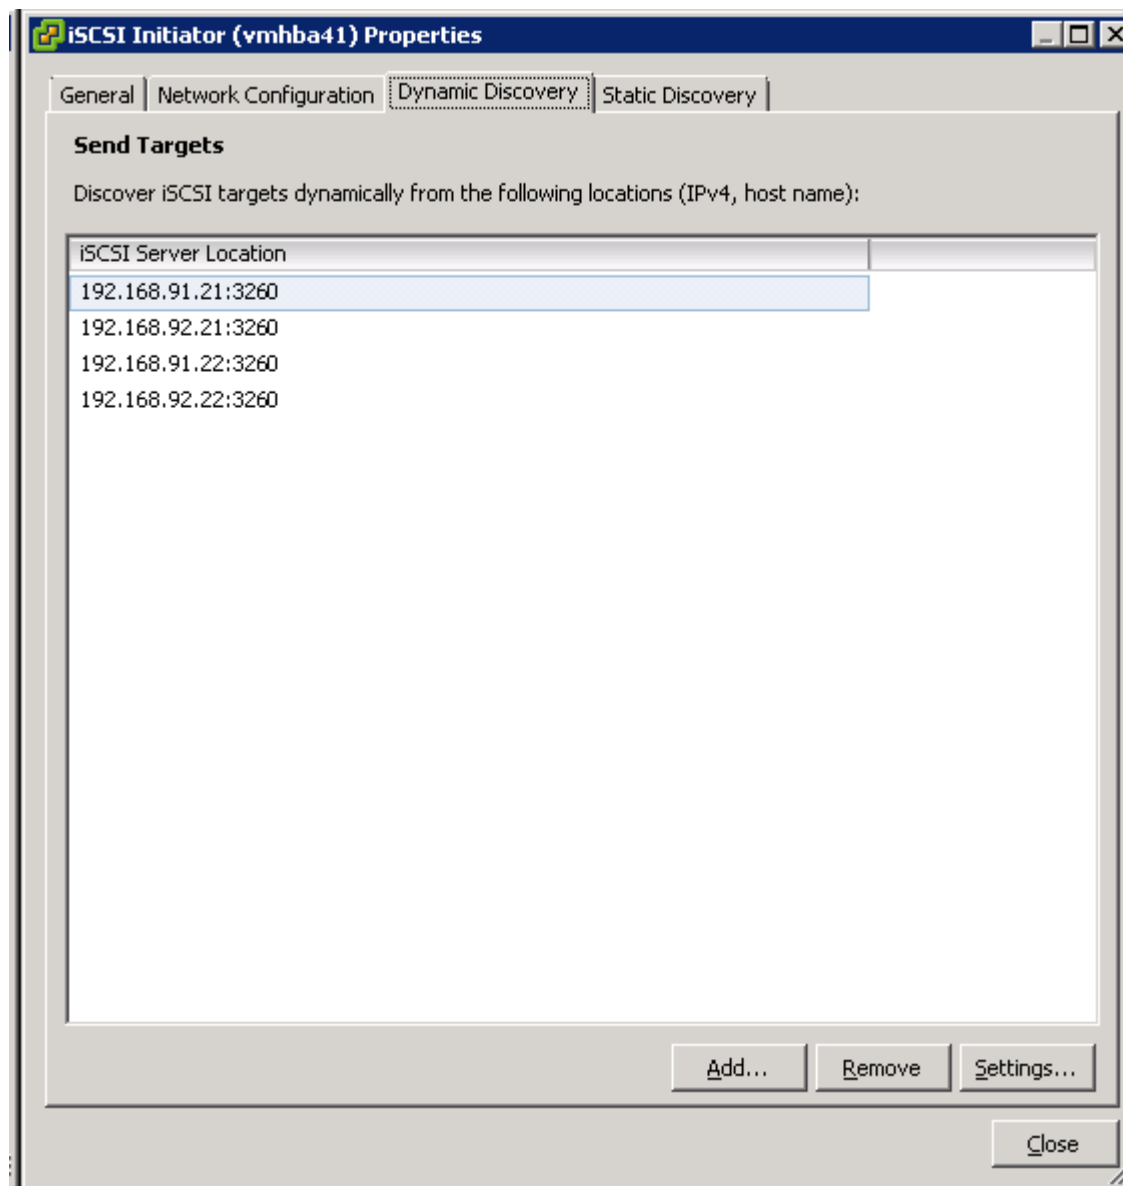
Networking

Setup iSCSI Multipathing

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

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

1. From the vSphere Client, click Storage Adapters in the Hardware pane.
2. Select the iSCSI Software Adapter and click Properties.
3. Select the Dynamic Discovery tab and click Add.
4. Enter the IP address of iscsi_lif01a.
5. Click OK.
6. Repeat putting in the IP addresses of iscsi_lif01b, iscsi_lif02a and iscsi_lif02b.



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

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

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

- [fnic Driver version 1.6.0.17a](#)
- [enic Driver version 2.2.2.71](#)



VMware vSphere 5.5 drivers are supported to work with vSphere 6.0.

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

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. From Resources > Storage, right-click datastore1 and select Browse Datastore.
4. Click the fourth button and select Upload File.
5. Navigate to the saved location for the downloaded VIC drivers and select `fnic_driver_1.6.0.17a_ESX55-offline_bundle-2774889.zip`.
6. Click Open and Yes to upload the file to datastore1.
7. Click the fourth button and select Upload File.
8. Navigate to the saved location for the downloaded VIC drivers and select `enic-2.1.2.71_esx55-offline_bundle-2739120.zip`.
9. Click Open and Yes to upload the file to datastore1.
10. Make sure the files have been uploaded to both ESXi hosts.
11. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
12. At the command prompt, run the following commands to account for each host

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d
/vmfs/volumes/datastore1/fnic_driver_1.6.0.17a_ESX55-offline_bundle-2774889.zip
```



To get the host thumbprint, type the command without the `--thumbprint` option, then copy and paste the thumbprint into the command.

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d
/vmfs/volumes/datastore1/fnic_driver_1.6.0.17a_ESX55-offline_bundle-2774889.zip
```

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d /vmfs/volumes/datastore1/enic-
2.1.2.71_esx55-offline_bundle-2739120.zip
```

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d /vmfs/volumes/datastore1/enic-
2.1.2.71_esx55-offline_bundle-2739120.zip
```

13. Back in the vSphere Client for each host, right click the host and select Reboot.

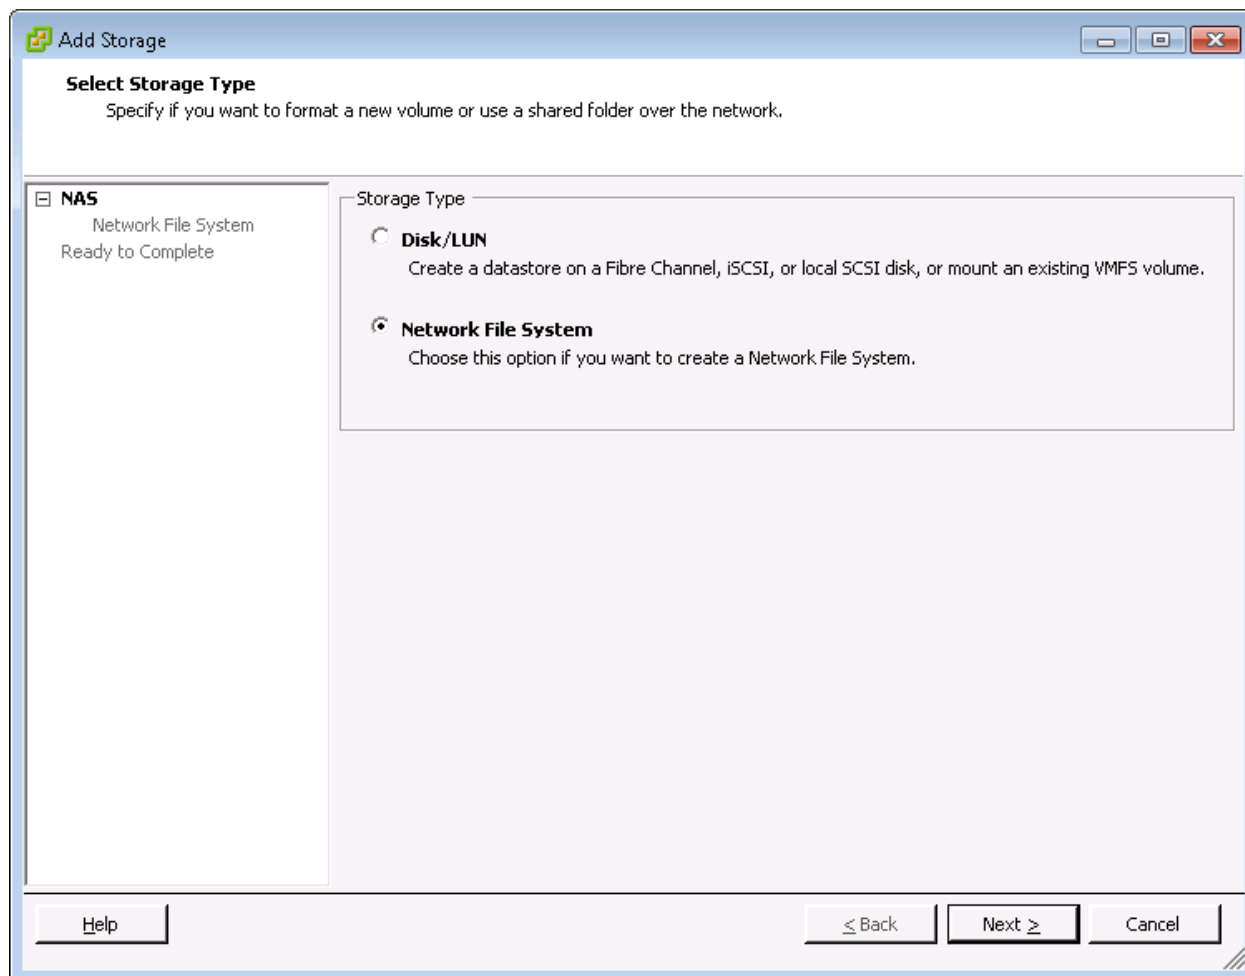
14. Click Yes and OK to reboot the host.
15. Log back into each host with vSphere Client.

Mount Required Datastores

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

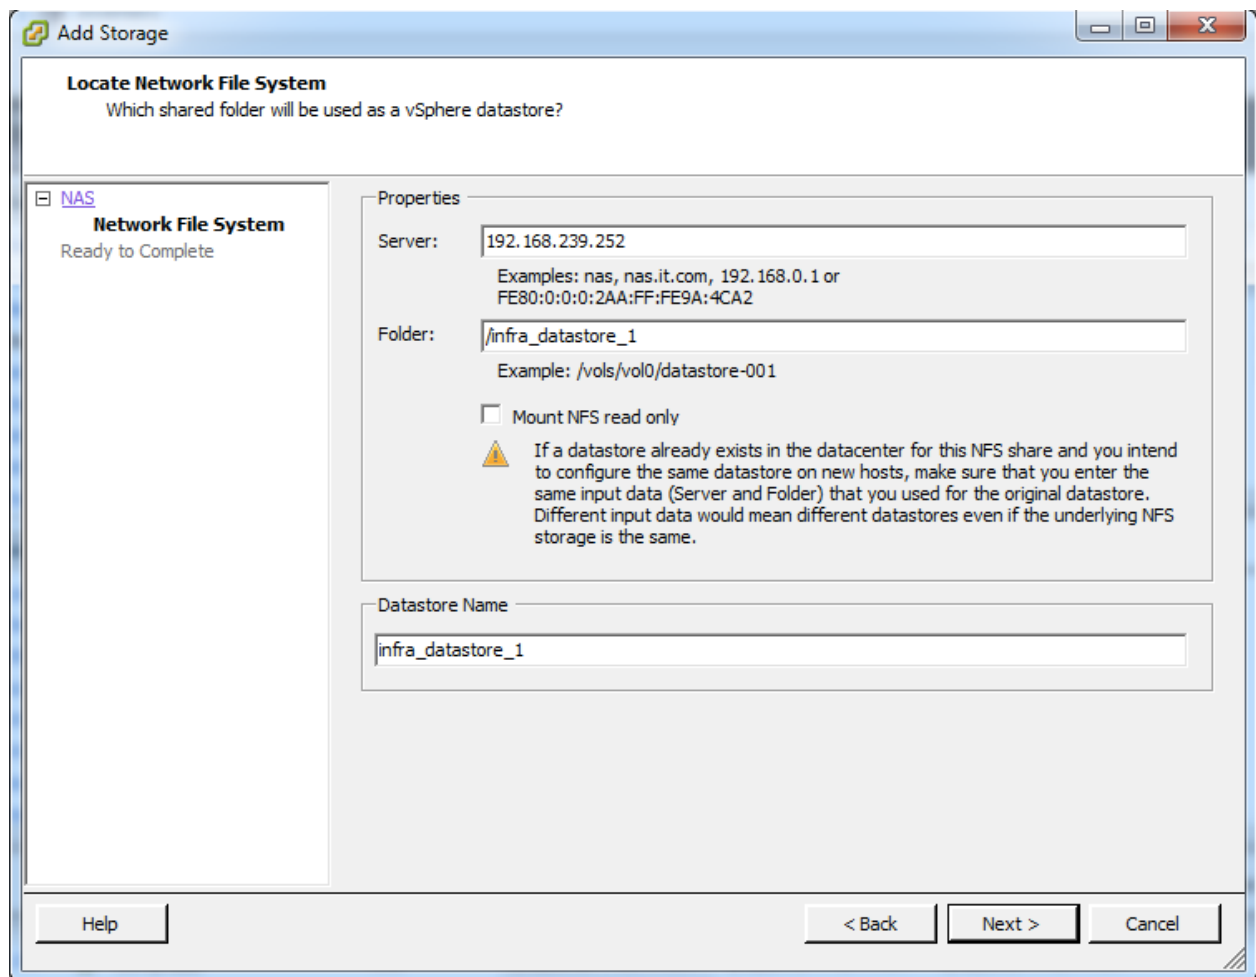
To mount the required datastores, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the Configuration tab.
3. Click Storage in the Hardware pane.
4. From the Datastores area, click Add Storage to open the Add Storage wizard.

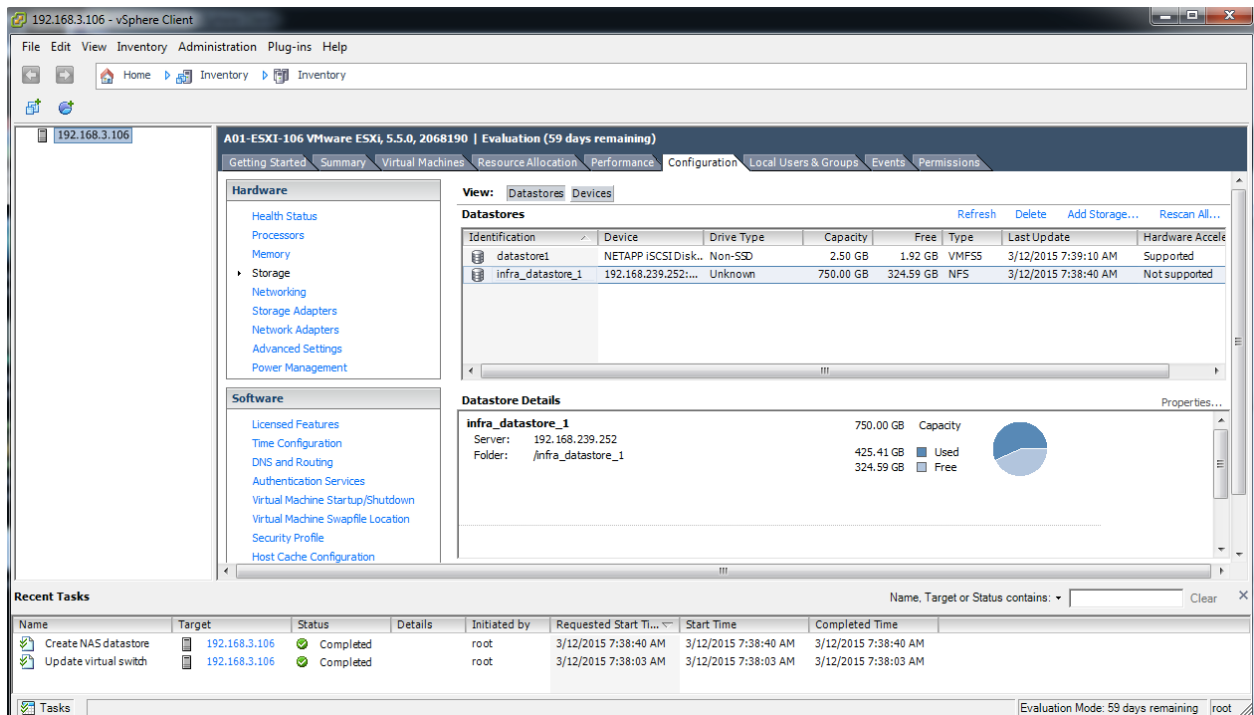


5. Select Network File System and click Next.

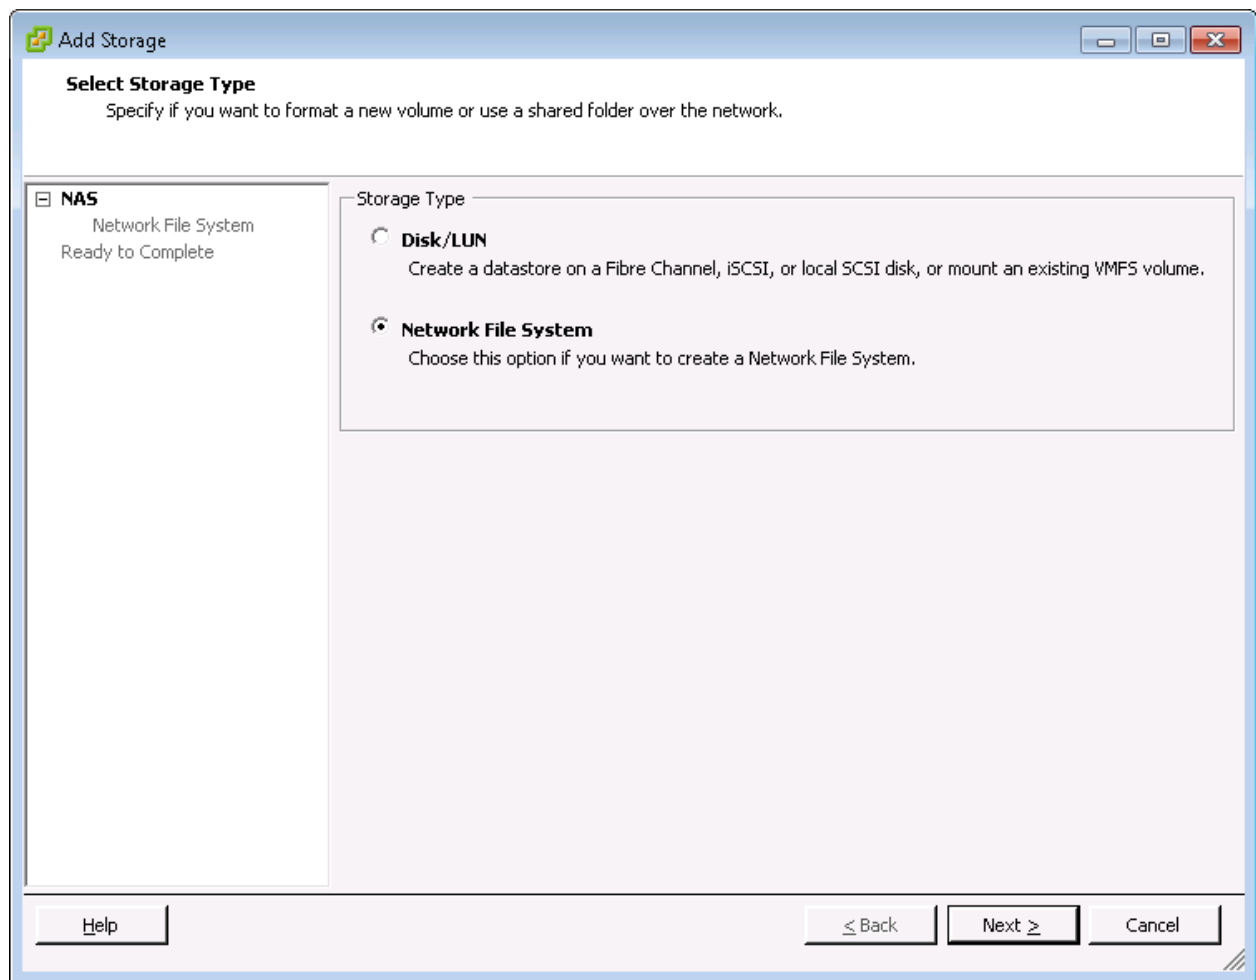
- The wizard prompts for the location of the NFS export. Enter <<var_node02_nfs_lif_infra_datastore_1_ip>> as the IP address for nfs_lif_infra_datastore_1.
- Enter /infra_datastore_1 as the path for the NFS export.
- Confirm that the Mount NFS read only checkbox is not selected.
- Enter infra_datastore_1 as the datastore name.



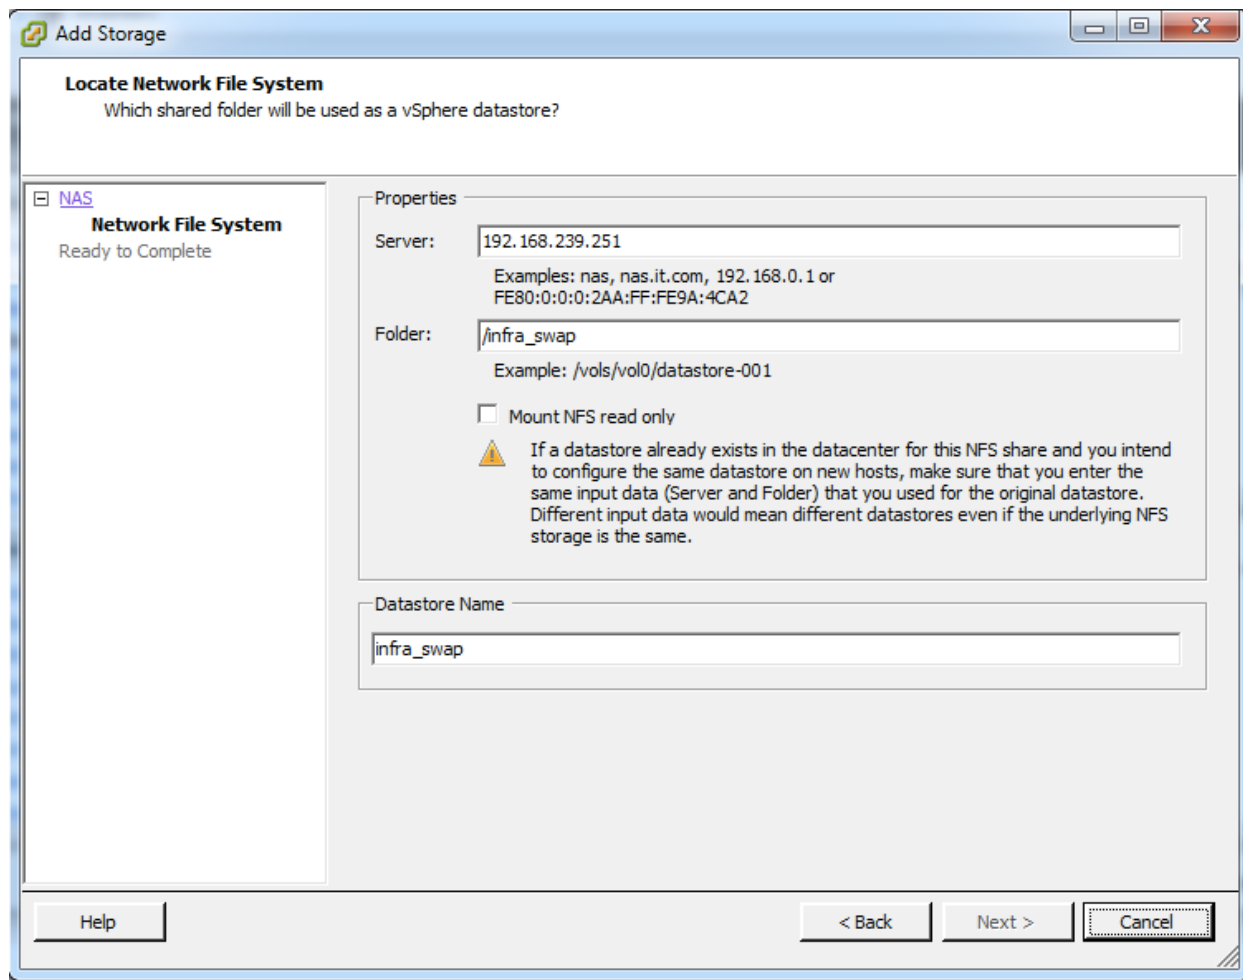
- To continue with the NFS datastore creation, click Next.
- To finalize the creation of the NFS datastore, click Finish.



12. From the Datastores area, click Add Storage to open the Add Storage wizard.



13. Select Network File System and click Next.
14. The wizard prompts for the location of the NFS export. Enter `<<var_node01_nfs_lif_infra_swap_ip>>` as the IP address for `nfs_lif_infra_swap`.
15. Enter `/infra_swap` as the path for the NFS export.
16. Confirm that the Mount NFS read only checkbox is not selected.
17. Enter `infra_swap` as the datastore name.



18. To continue with the NFS datastore creation, click Next.

19. To finalize the creation of the NFS datastore, click Finish.

Configure NTP on ESXi Hosts

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

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

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click **Time Configuration** in the Software pane.
4. Click **Properties** at the upper-right side of the window.
5. At the bottom of the Time Configuration dialog box, click **Options**.
6. In the NTP Daemon (ntpd) Options dialog box, complete the following steps:

- a. Click **General** in the left pane and select Start and stop with host.
 - b. Click **NTP Settings** in the left pane and click **Add**.
7. In the Add NTP Server dialog box, enter <<var_switch_a_ntp_ip>> as the IP address of the NTP server and click **OK**.
8. Click **Add**.
9. In the Add NTP Server dialog box, enter <<var_switch_b_ntp_ip>> as the IP address of the NTP server and click **OK**.
10. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes checkbox and click **OK**.
11. In the Time Configuration dialog box, complete the following steps:
 - a. Select the NTP Client Enabled checkbox and click **OK**.
 - b. Verify that the clock is now set to approximately the correct time.




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

Move VM Swap File Location

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

To move the VM swap file location, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the **Configuration** tab.
3. Click **Virtual Machine Swapfile** Location in the Software pane.
4. Click **Edit** at the upper-right side of the window.
5. Select “Store the swapfile in a swapfile datastore selected below.”
6. Select the <infra_swap> datastore in which to house the swap files.



7. Click OK to finalize moving the swap file location.

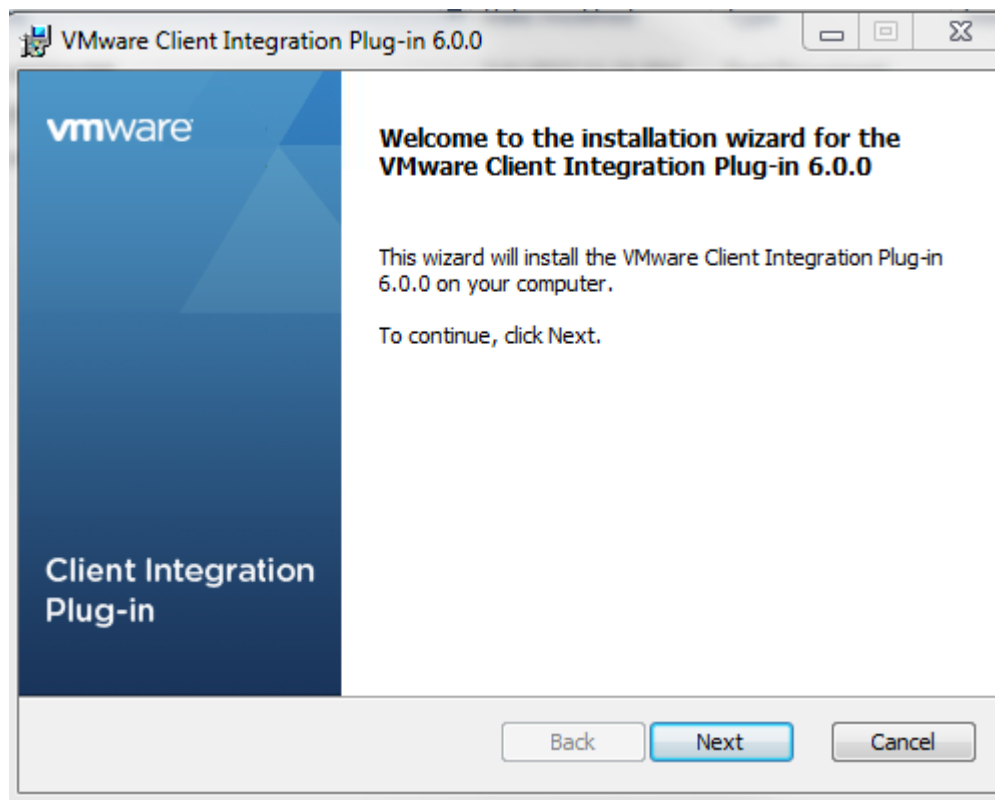
VMware vCenter 6.0

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.0 Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Install the Client Integration Plug-in

To install the client integration plug-in, complete the following steps:

1. Download the .iso installer for the vCenter Server Appliance and Client Integration Plug-in.
2. Mount the ISO image to the Windows virtual machine or physical server on which you want to install the Client Integration Plug-In to deploy the vCenter Server Appliance.
3. In the software installer directory, navigate to the vcsa directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe. The Client Integration Plug-in installation wizard appears.

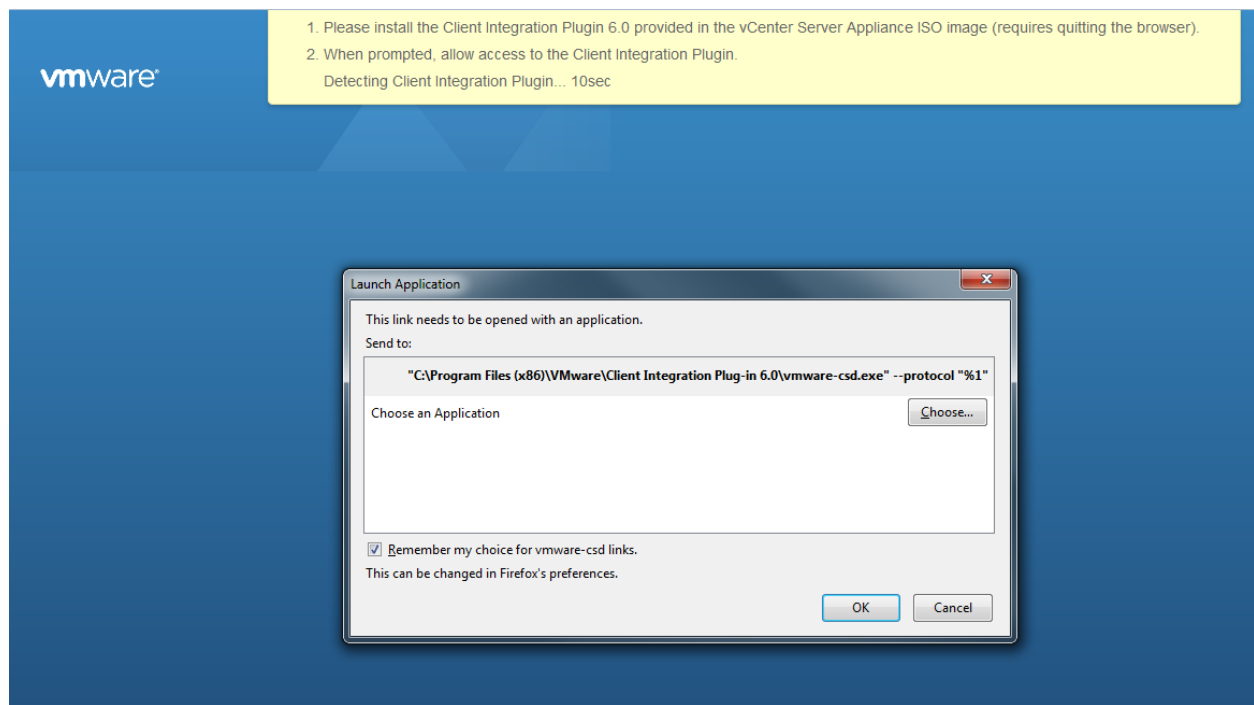


4. On the Welcome page, click Next.
5. Read and accept the terms in the End-User License Agreement and click Next.
6. Click Next.
7. Click Install.


Building the VMware vCenter Server Appliance

To build the VMware vCenter virtual machine, complete the following steps:

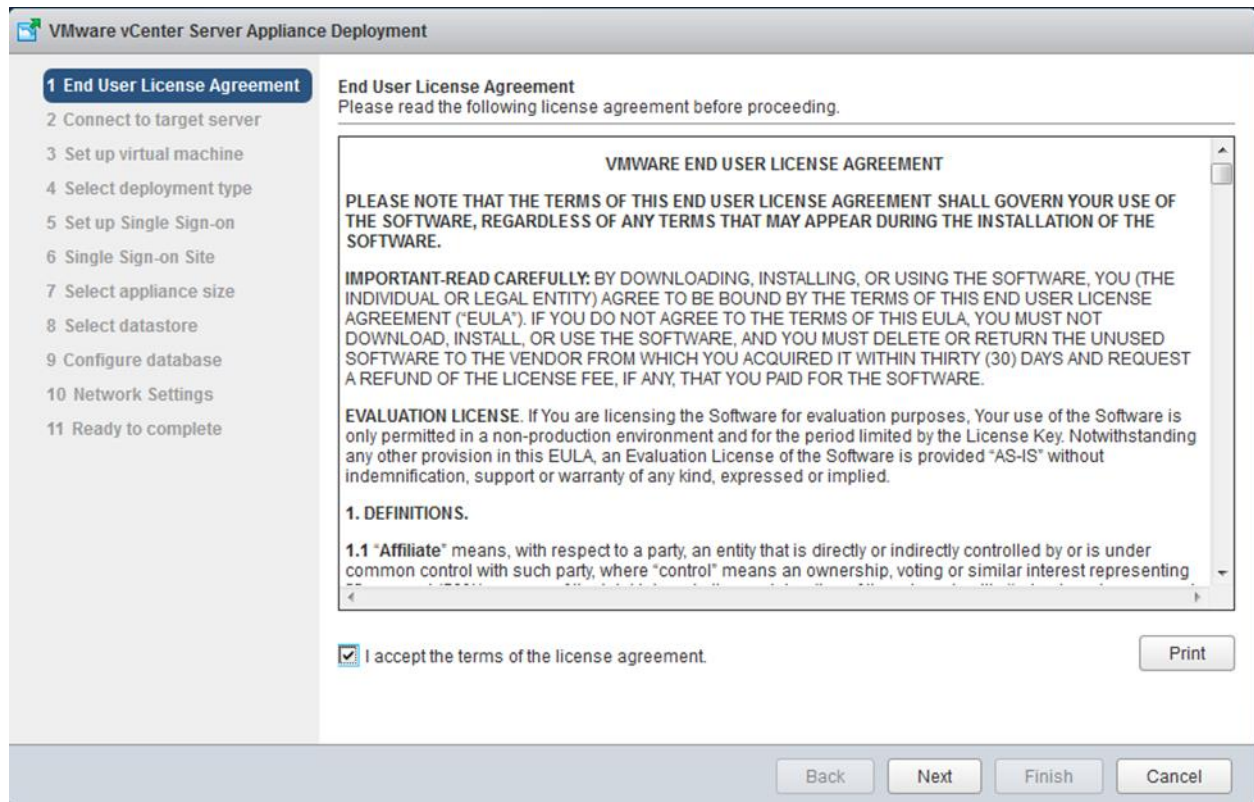
1. In the software installer directory, double-click `vcsa-setup.html`.
2. Allow the plug-in to run on the browser when prompted.



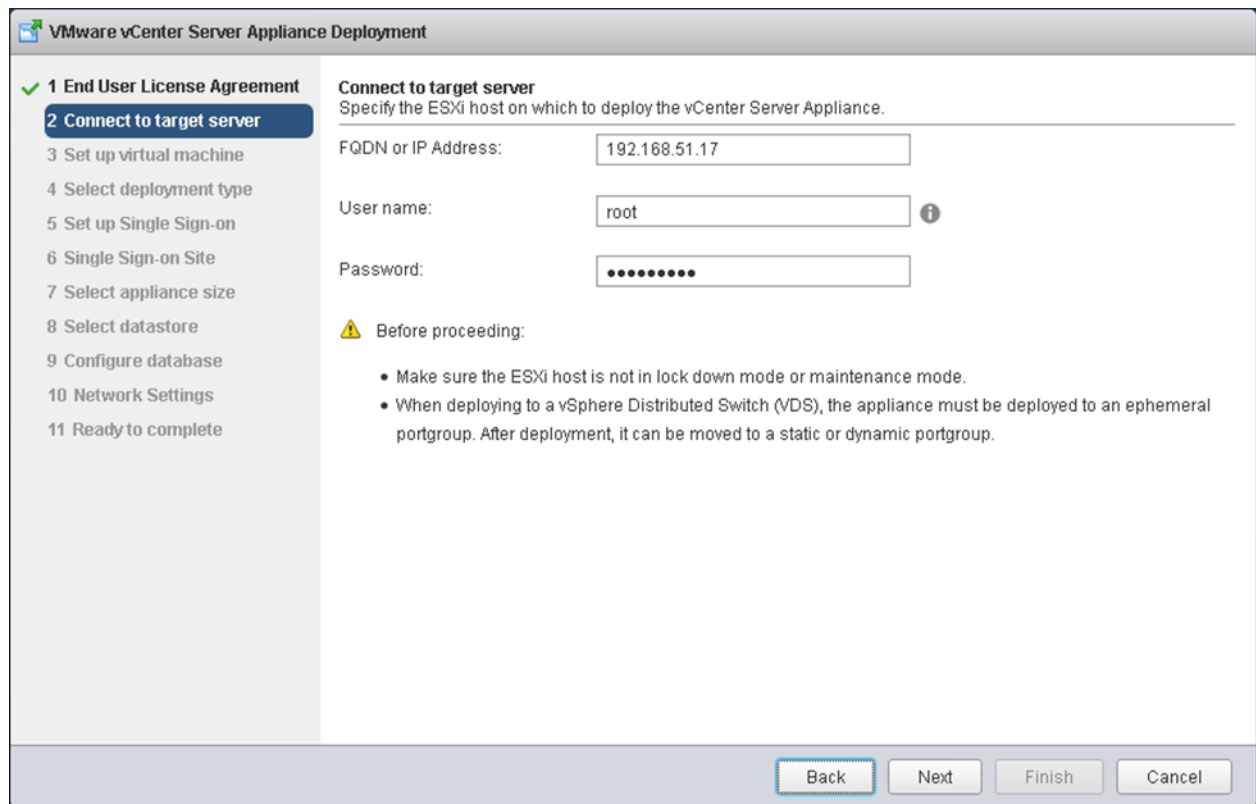
3. On the Home page, click Install to start the vCenter Server Appliance deployment wizard.



4. Read and accept the license agreement, and click Next.



5. In the "Connect to target server" page, enter the ESXi host name, User name and Password.



6. Click Yes to accept the certificate.

7. Enter the Appliance name and password details in the “Set up virtual machine” page.

The screenshot shows the 'Set up virtual machine' step of the VMware vCenter Server Appliance Deployment wizard. The left sidebar contains a list of steps: 1 End User License Agreement, 2 Connect to target server, 3 Set up virtual machine (highlighted), 4 Select deployment type, 5 Set up Single Sign-on, 6 Single Sign-on Site, 7 Select appliance size, 8 Select datastore, 9 Configure database, 10 Network Settings, and 11 Ready to complete. The main area is titled 'Set up virtual machine' and contains the following fields: 'Appliance name' with the value 'icee1-vcenter', 'OS user name' with the value 'root', 'OS password' (masked with dots), and 'Confirm OS password' (masked with dots). Information icons are present next to the 'Appliance name' and 'OS password' fields. At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

8. In the “Select deployment type” page, choose “Install vCenter Server with an embedded Platform Services Controller”.

VMware vCenter Server Appliance Deployment

✓ 1 End User License Agreement
✓ 2 Connect to target server
✓ 3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Single Sign-on Site
7 Select appliance size
8 Select datastore
9 Configure database
10 Network Settings
11 Ready to complete

Select deployment type
Select the services to deploy onto this appliance.

vCenter Server 6.0 requires a Platform Services Controller, which contains shared services such as Single Sign-On, Licensing, and Certificate Management. An embedded Platform Services Controller is deployed on the same Appliance VM as vCenter Server. An external Platform Services Controller is deployed in a separate Appliance VM. For smaller installations, consider vCenter Server with an embedded Platform Services Controller. For larger installations with multiple vCenter Servers, consider one or more external Platform Services Controllers. Refer to the vCenter Server documentation for more information.

Note: Once you install vCenter Server, you can only change from an embedded to an external Platform Services Controller with a fresh install.

Embedded Platform Services Controller

Install vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

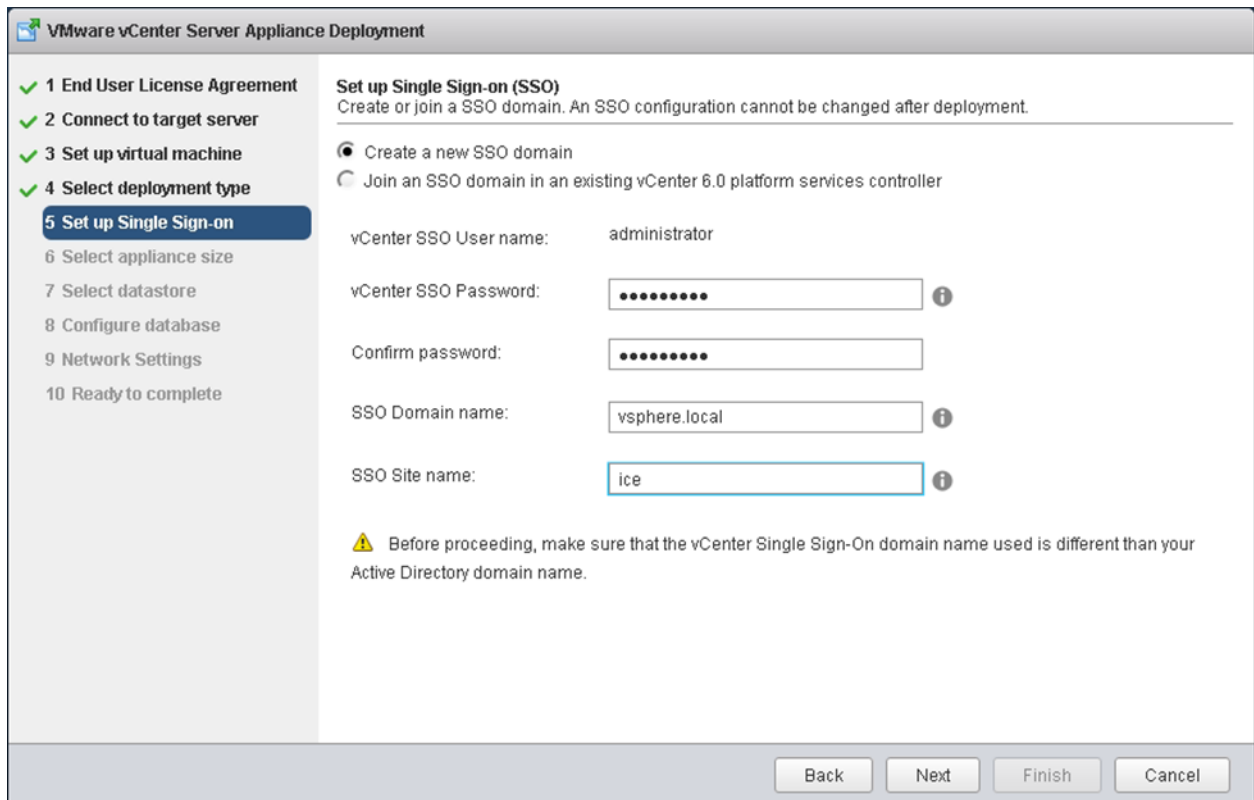
Install Platform Services Controller
 Install vCenter Server (Requires External Platform Services Controller)

Back Next Finish Cancel

9. Click Next.

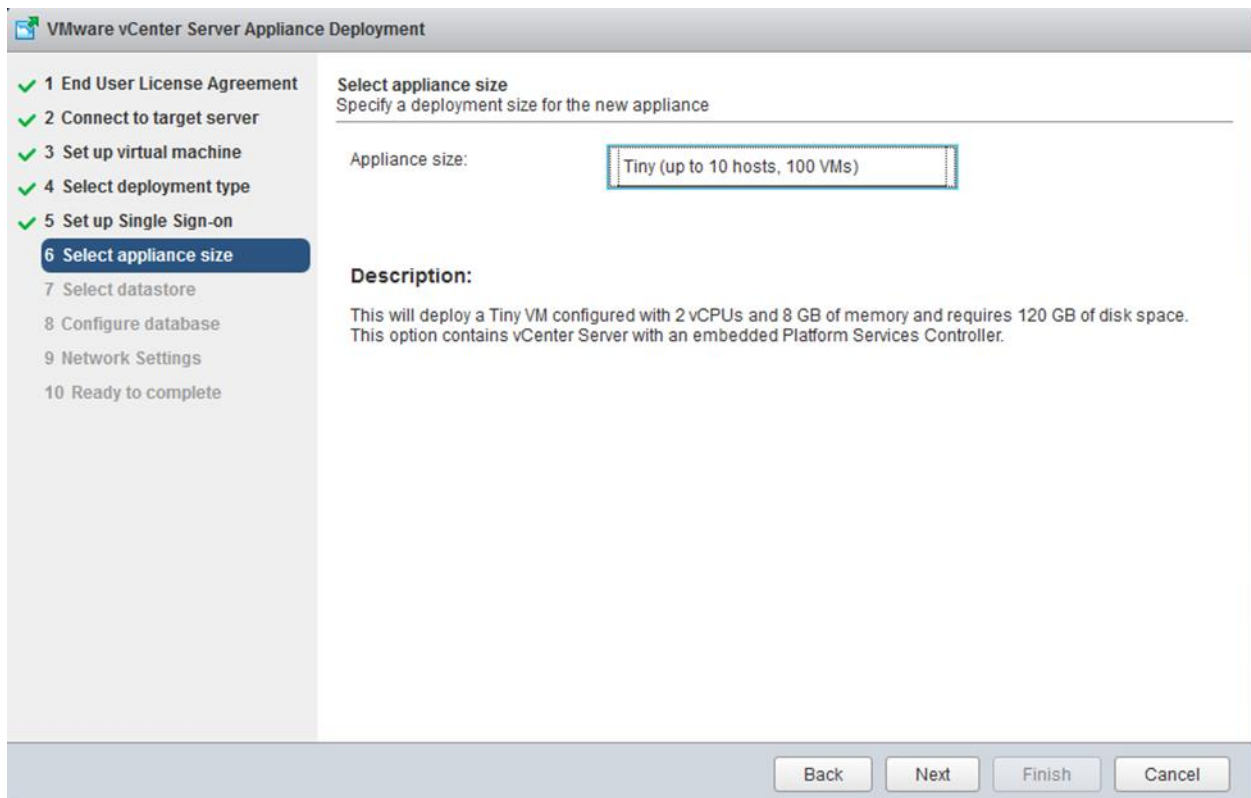
10. In the “Set up Single Sign-On” page, select “Create a new SSO domain.”

11. Enter the SSO password, Domain name and Site name.



12. Click Next.

13. Select the appliance size. For example, "Tiny (up to 10 hosts, 100 VMs)."



14. Click Next.

15. In the “Select datastore” page, choose infra_datastore_1.

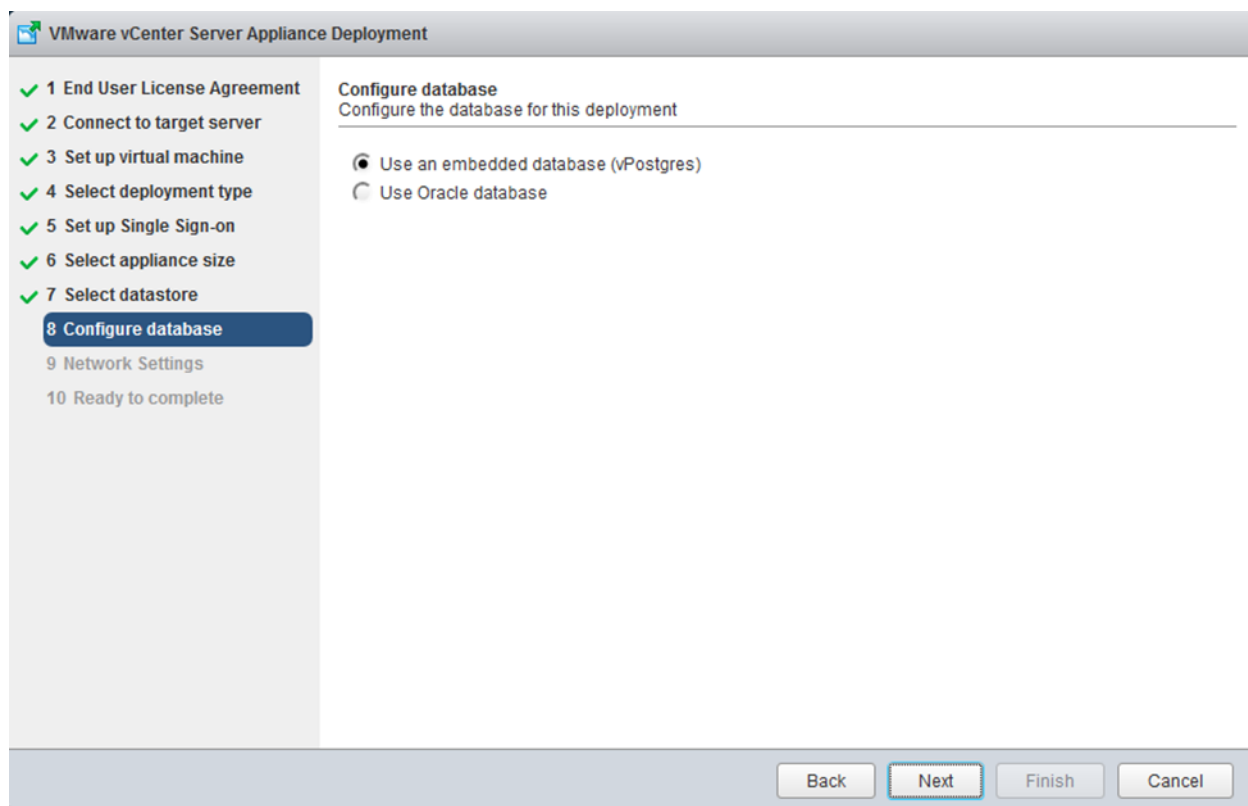
The screenshot shows the 'VMware vCenter Server Appliance Deployment' wizard. The left sidebar contains a list of steps: 1 End User License Agreement, 2 Connect to target server, 3 Set up virtual machine, 4 Select deployment type, 5 Set up Single Sign-on, 6 Select appliance size, 7 Select datastore (highlighted), 8 Configure database, 9 Network Settings, and 10 Ready to complete. The main area is titled 'Select datastore' and includes the instruction 'Select the storage location for this deployment'. Below this, a text block states: 'The following datastores are accessible. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.' A table lists the available datastores:

Name	Type	Capacity	Free	Provisioned	Thin Provisioning
datastore1	VMFS	12.5 GB	11.63 GB	0.87 GB	true
infra_datastore_1	NFS	500 GB	499.99 GB	0.01 GB	true
infra_swap	NFS	100 GB	99.99 GB	0.01 GB	true

Below the table is a horizontal scrollbar. At the bottom of the main area, there is a checkbox labeled 'Enable Thin Disk Mode' with an information icon. The bottom of the wizard features four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

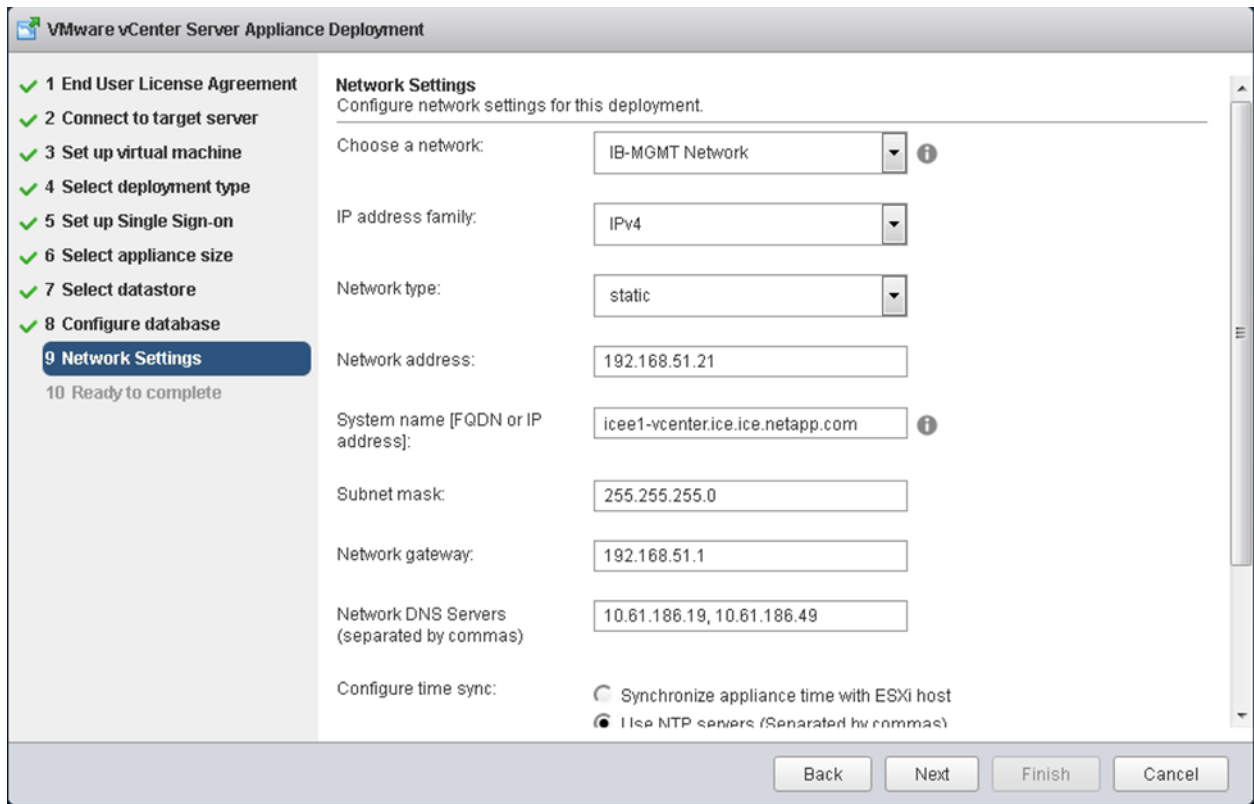
16. Click Next.

17. Select embedded database in the “Configure database” page. Click Next.

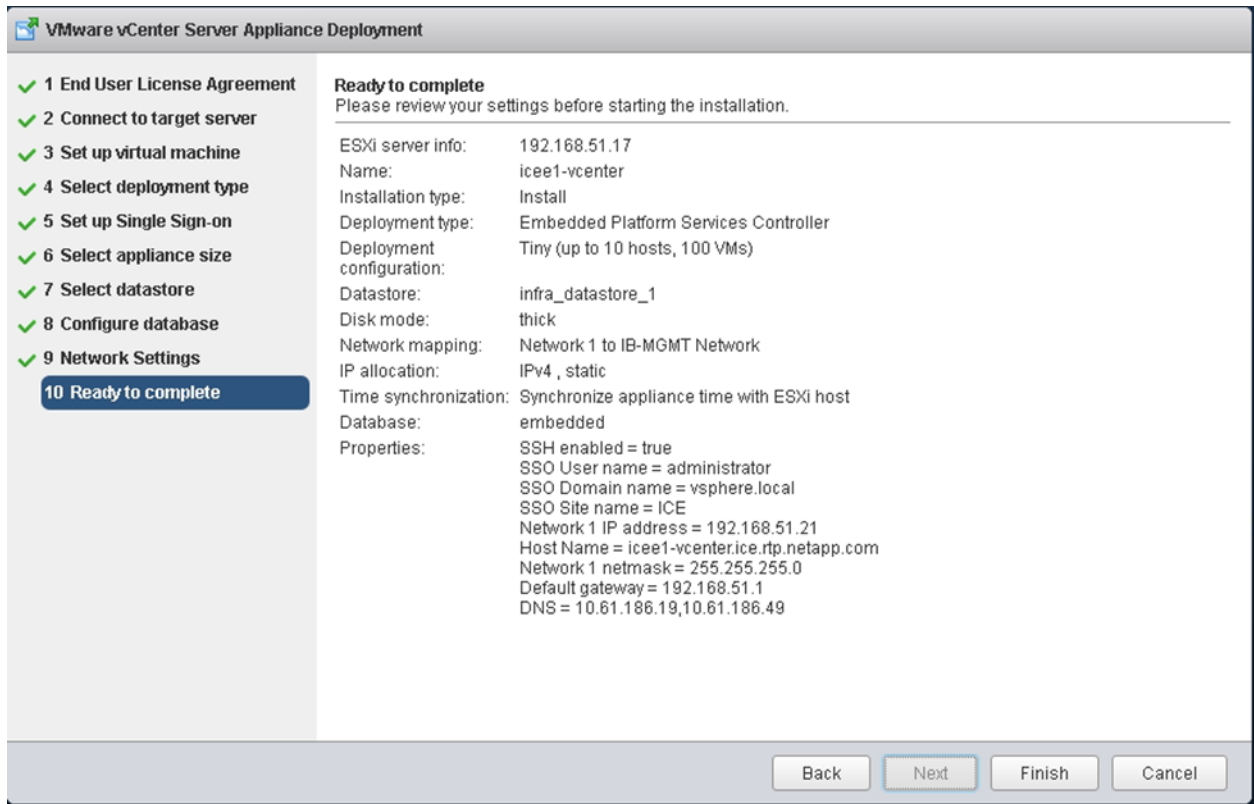


18. In the “Network Settings” page, configure the below settings:

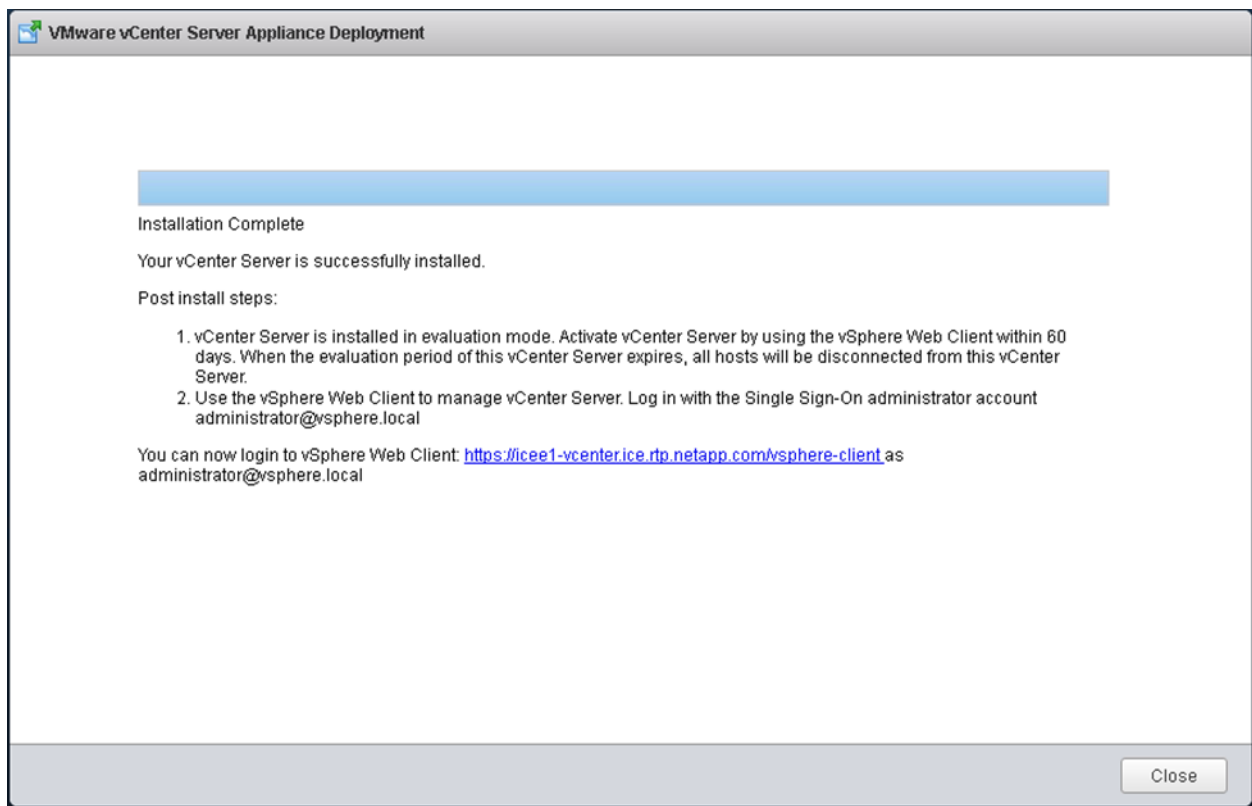
- a. Choose a Network: MGMT-Network
- b. IP address family: IPV4
- c. Network type: static
- d. Network address: <<var_vcenter_ip>>
- e. System name: <<var_vcenter_fqdn>>
- f. Subnet mask: <<var_vcenter_subnet_mask>>
- g. Network gateway: <<var_vcenter_gateway>>
- h. Network DNS Servers: <<var_dns_server>>
- i. Configure time sync: Use NTP servers
- j. (Optional). Enable SSH



19. Review the configuration and click Finish.



20. The vCenter appliance installation will take few minutes to complete.



Setting Up VMware vCenter Server

1. Using a web browser, navigate to `https://<var_vcenter_ip>`.

The screenshot shows the VMware vCenter Server web interface. The VMware logo is in the top left corner. The main content area is divided into two columns. The left column is titled "Getting Started" and contains the following text:

To access vSphere remotely, use the vSphere Web Client.

[Log in to vSphere Web Client](#)

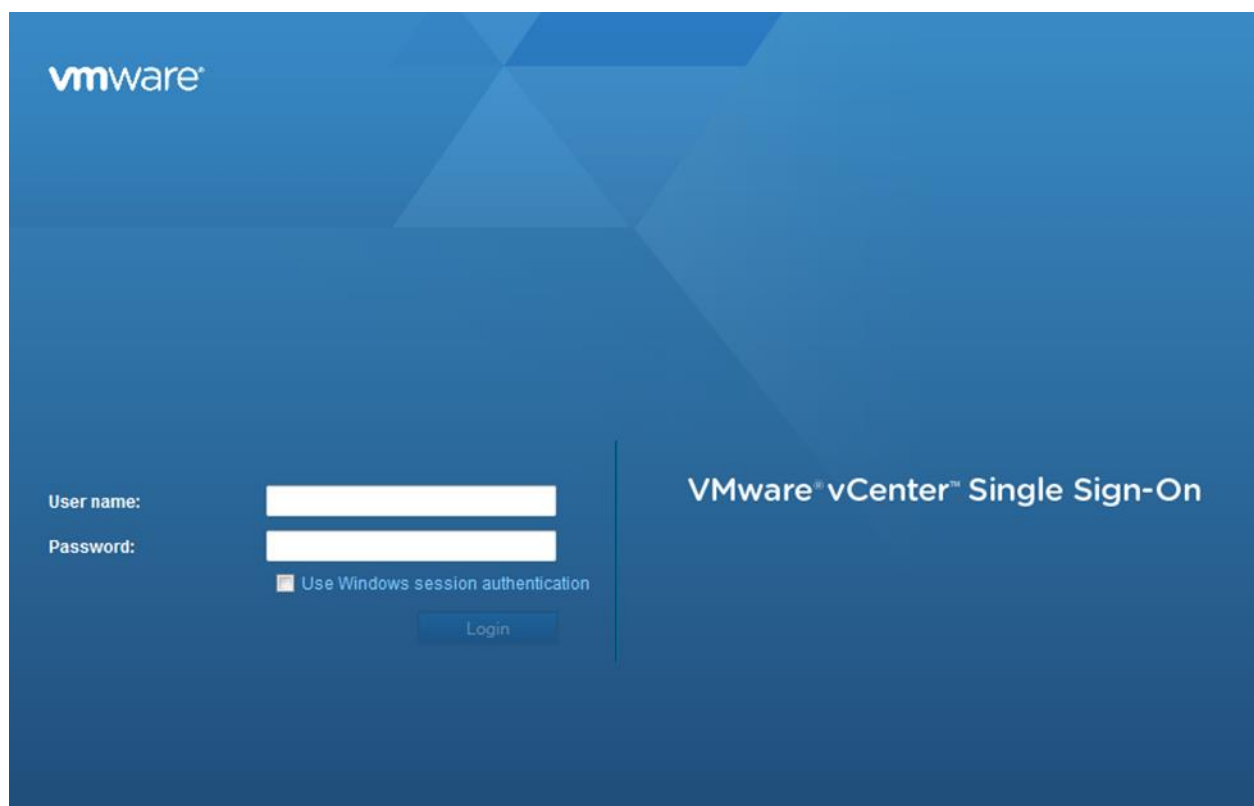
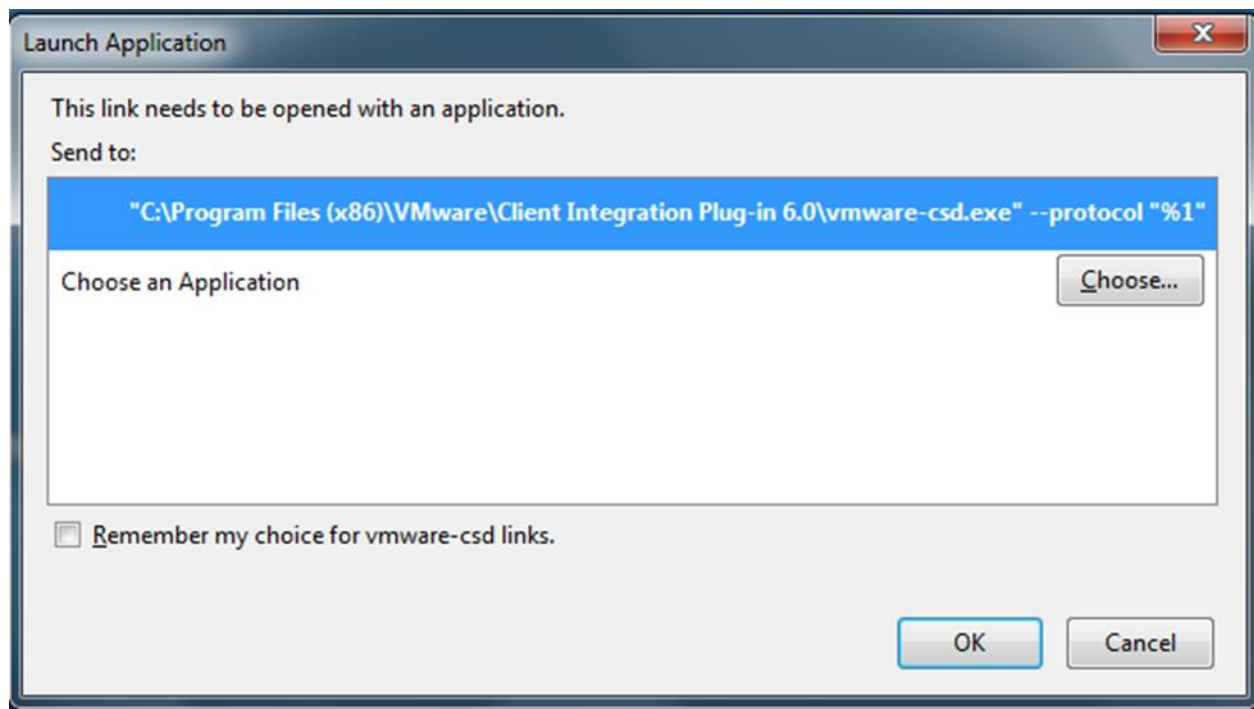
For help, see [vSphere Documentation](#)

In the center, there is an illustration of vCenter Servers represented as server racks.

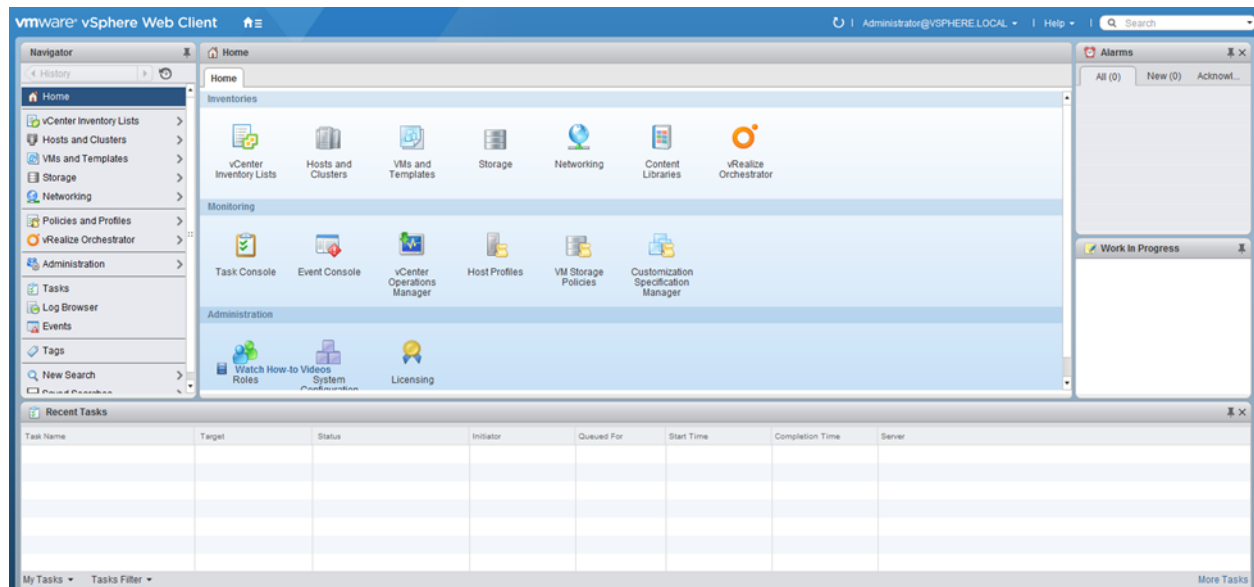
The right column is titled "For Administrators" and "For Developers". Under "For Administrators", there is a section for "Web-Based Datastore Browser" with the text: "Use your web browser to find and download files (for example, virtual machine and virtual disk files). [Browse datastores in the vSphere inventory](#)". Under "For Developers", there is a section for "vSphere Web Services SDK" with the text: "Learn about our latest SDKs, Toolkits, and APIs for managing VMware ESX, ESXi, and VMware vCenter. Get sample code, reference documentation, participate in our Forum Discussions, and view our latest Sessions and Webinars." Below this, there are links for "Learn more about the Web Services SDK", "Browse objects managed by vSphere", and "Download trusted root CA certificates".

Copyright © 1998-2015 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware products may contain individual open source software components, each of which

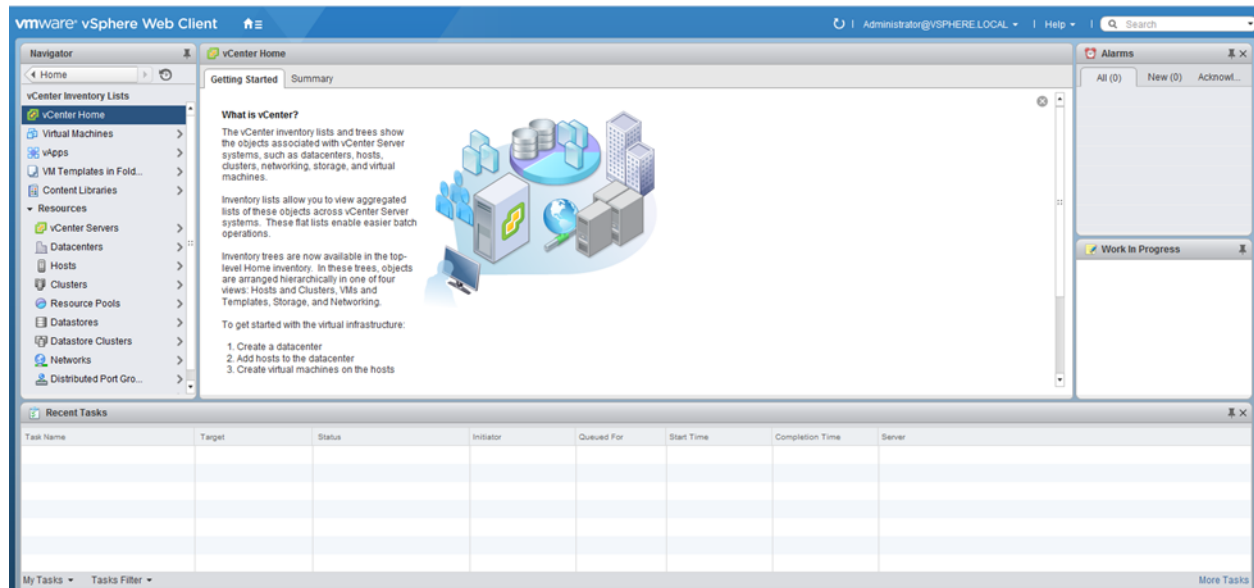
2. Click Log in to vSphere Web Client.
3. Click OK if "Launch Application" window appears.



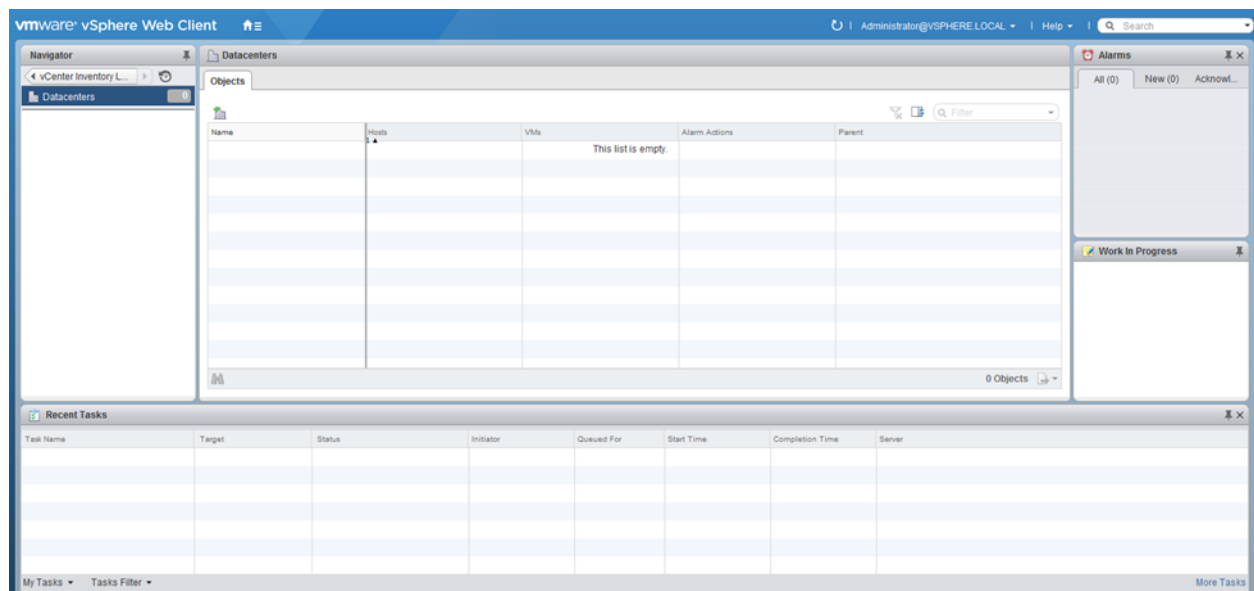
4. Log in using Single Sign-On username and password created during the vCenter installation.




5. Navigate to vCenter Inventory Lists on the left pane.



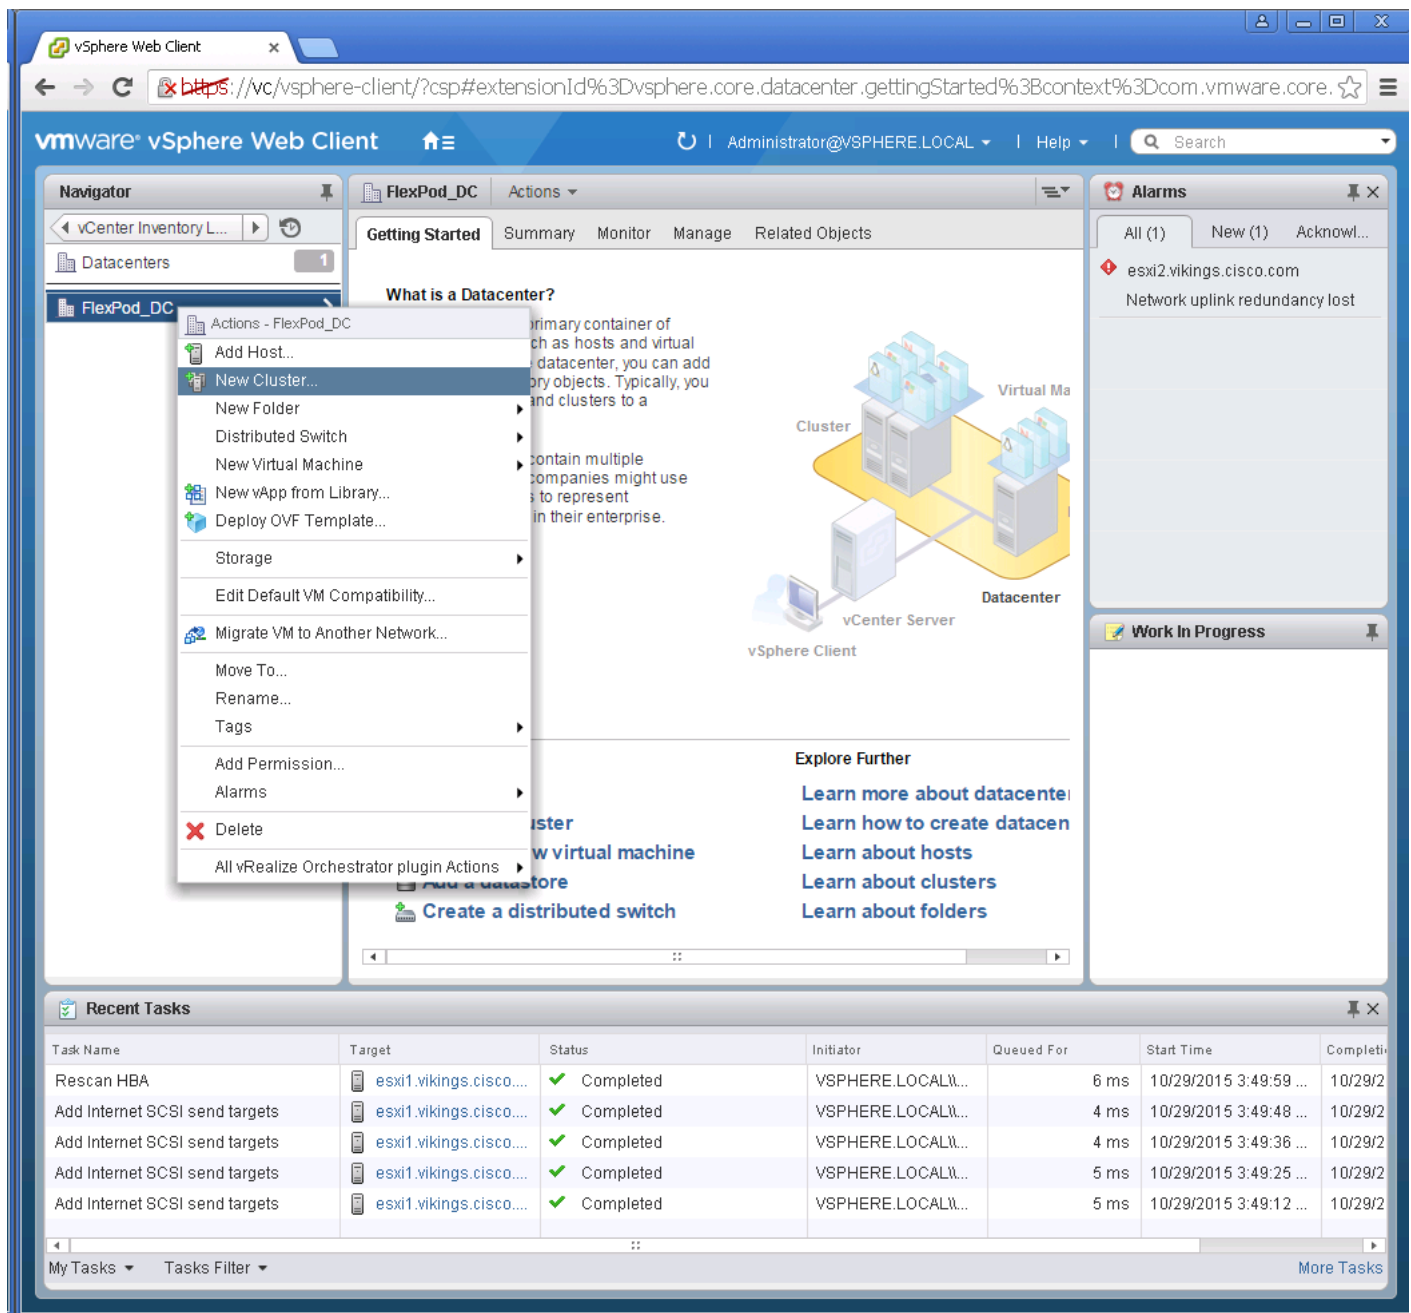
6. Under Resources, click Datacenters in the left pane.



7. To create a Data center, click the leftmost icon in the center pane that has a green plus symbol above it.
8. Type "FlexPod_DC" in the Datacenter name field.
9. Select the vCenter Name/IP option.
10. Click OK.



11. Right-click the data center FlexPod_DC in the list in the center pane. Click New Cluster.






12. Name the cluster FlexPod_Management.

13. Check the box beside DRS. Leave the default values.

14. Check the box beside vSphere HA. Leave the default values.

New Cluster ? >>

Name	FlexPod_Management
Location	 FlexPod_DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated ▾
Migration Threshold	Conservative  Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	
Admission Control Status	Admission control will prevent powering on VMs that violate availability constraints <input checked="" type="checkbox"/> Enable admission control
Policy	Specify the type of the policy that admission control should enforce. <input checked="" type="radio"/> Host failures cluster tolerates: <input type="text" value="1"/> ▾ <input type="radio"/> Percentage of cluster resources reserved as failover spare capacity: Reserved failover CPU capacity: <input type="text" value="25"/> ▾ % CPU Reserved failover Memory capacity: <input type="text" value="25"/> ▾ % Memory
VM Monitoring	
VM Monitoring Status	Disabled ▾ Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low  High
EVC	Disable ▾
Virtual SAN	<input type="checkbox"/> Turn ON

15. Click OK to create the new cluster.

16. On the left pane, double click the “FlexPod_DC”.

17. Click Clusters.

The screenshot displays the VMware vSphere Web Client interface. The main pane shows the 'FlexPod_DC' cluster with the 'Clusters' tab selected. The 'FlexPod_Management' cluster is highlighted in the list. The 'Related Objects' table shows the following data:

Name	Available CPU (GHz)	Available Memory (GB)
FlexPod_Management	96.62 GHz	237.05 GB

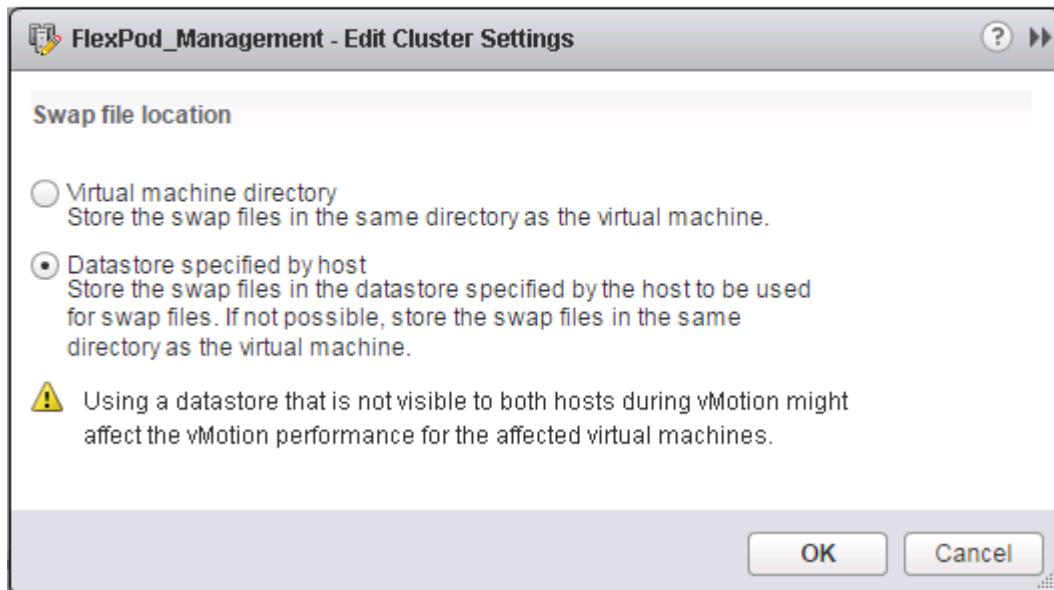
The 'Recent Tasks' pane at the bottom shows a list of completed tasks:

Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time
Rescan HBA	esxi1.vikings.cisco...	Completed	VSPHERE.LOCAL\...	6 ms	10/29/2015 3:49:59 ...	10/29/2015 3:49:59 ...
Add Internet SCSI send targets	esxi1.vikings.cisco...	Completed	VSPHERE.LOCAL\...	4 ms	10/29/2015 3:49:48 ...	10/29/2015 3:49:48 ...
Add Internet SCSI send targets	esxi1.vikings.cisco...	Completed	VSPHERE.LOCAL\...	4 ms	10/29/2015 3:49:36 ...	10/29/2015 3:49:36 ...
Add Internet SCSI send targets	esxi1.vikings.cisco...	Completed	VSPHERE.LOCAL\...	5 ms	10/29/2015 3:49:25 ...	10/29/2015 3:49:25 ...
Add Internet SCSI send targets	esxi1.vikings.cisco...	Completed	VSPHERE.LOCAL\...	5 ms	10/29/2015 3:49:12 ...	10/29/2015 3:49:12 ...

18. Under the Clusters pane, right click the “FlexPod_Management” and select Settings.

19. Select Configuration > General in the list on the left and select Edit to the right of General.

20. Select Datastore specified by host and click OK.



21. Under the Clusters pane, right click the "FlexPod_Management" and click Add Host.

The screenshot shows the VMware vSphere Web Client interface. The main view is the 'FlexPod_DC' cluster configuration page, specifically the 'Related Objects' tab. A table lists the objects in the cluster:

Name	Available CPU (GHz)	Available Memory (GB)
FlexPod_Management	96.62 GHz	237.05 GB

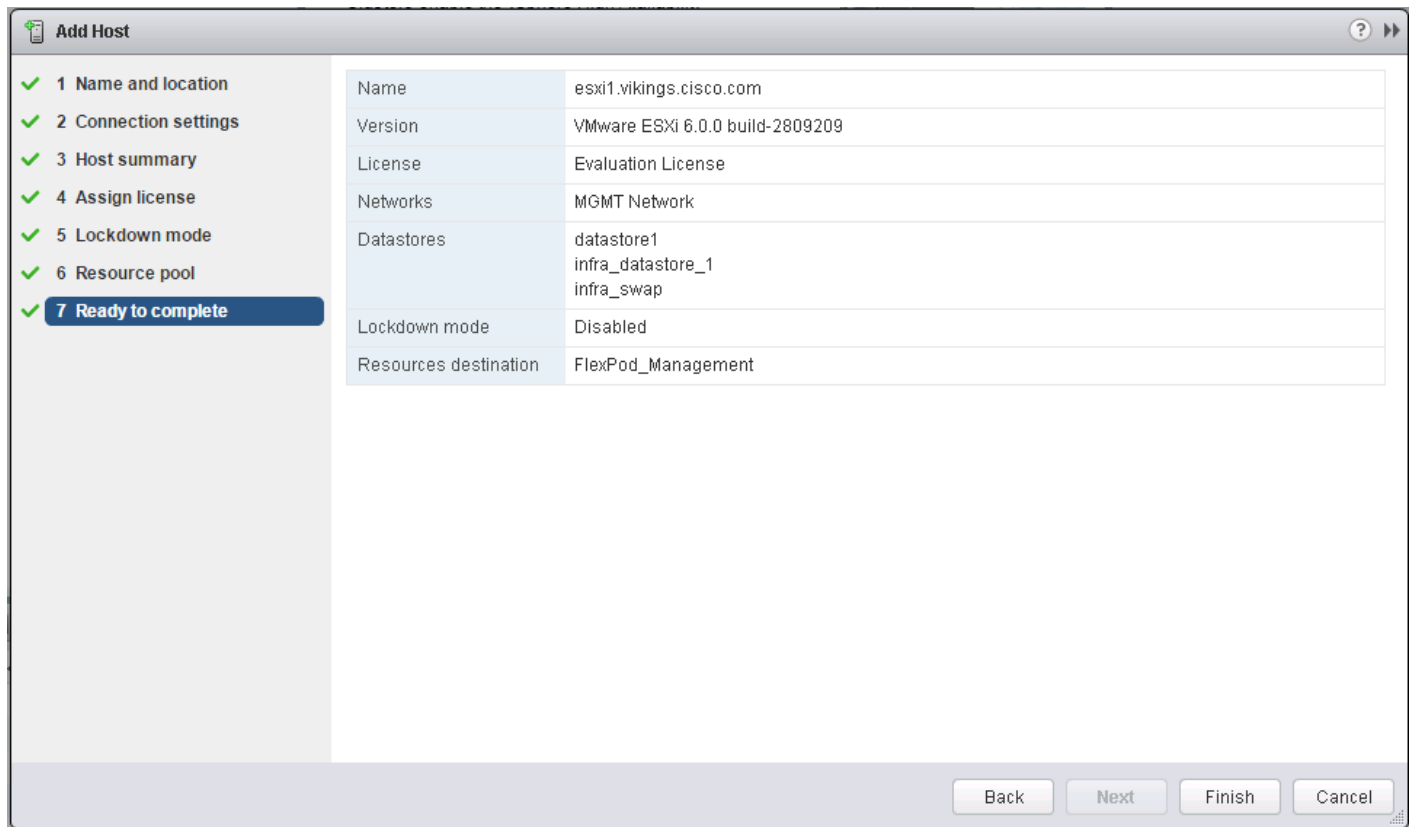
A context menu is open over the 'FlexPod_Management' object, showing various actions. The 'Assign License...' option is highlighted. The 'Alarms' panel on the right shows a warning for 'esxi2.vikings.cisco.com' regarding 'Network uplink redundancy lost'. The 'Work In Progress' panel is also visible.

22. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.
23. Type root as the user name and the root password. Click Next to continue.
24. Click Yes to accept the certificate.
25. Review the host details and click Next to continue.
26. Assign a license and click Next to continue.

27. Click Next to continue.

28. Click Next to continue.

29. Review the configuration parameters. Then click Finish to add the host.



30. Repeat the steps 18 to 27 to add the remaining VMware ESXi hosts to the cluster.



Two VMware ESXi hosts will be added to the cluster.

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. To setup the ESXi Dump Collector, complete the following steps:

1. In the vSphere web client, select Home.
2. In the center pane, click System Configuration.
3. In the left hand pane, click VMware vSphere ESXi Dump Collector.
4. In the Actions menu, choose Start.
5. In the Actions menu, click Edit Startup Type.

6. Select Automatic.
7. Click OK.
8. On the Management Workstation, open the VMware vSphere CLI command prompt.
9. Set each iSCSI-booted ESXi Host to coredump to the ESXi Dump Collector by running the following commands:

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> -
-thumbprint <host_thumbprint> system coredump network set --
interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip> --
server-port 6500
```



To get the host thumbprint, type the command without the --thumbprint option, then copy and paste the thumbprint into the command.

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> -
-thumbprint <host_thumbprint> system coredump network set --
interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip> --
server-port 6500
```

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> -
-thumbprint <host_thumbprint> system coredump network set --enable
true
```

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> -
-thumbprint <host_thumbprint> system coredump network set --enable
true
```

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> -
-thumbprint <host_thumbprint> system coredump network check
```

```
esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> -
-thumbprint <host_thumbprint> system coredump network check
```

Cisco UCS Virtual Media (vMedia) Policy for VMware ESXi Installation

Storage Controller Setup for vMedia Policy

NetApp Storage Cluster Setup

To setup the NetApp storage cluster, complete the following steps:

1. From an SSH session connected to the NetApp Storage cluster.
2. Allow the in-band mgmt vlan to access the infrastructure datastores.
3. Enter the following command:

```
vserver export-policy rule create -policyname default -
<<var_inband_mgmt_subnet_cidr>> -rorule sys -rwrule sys -allow-suid
false -vserver Infra-SVM -ruleindex 3 -protocol nfs -superuser sys
```

4. Create 2 ports and then add those ports to the newly created broadcast domain, enter the following commands:

```
network port vlan create -node clus-01 -vlan-name a0a-113
```

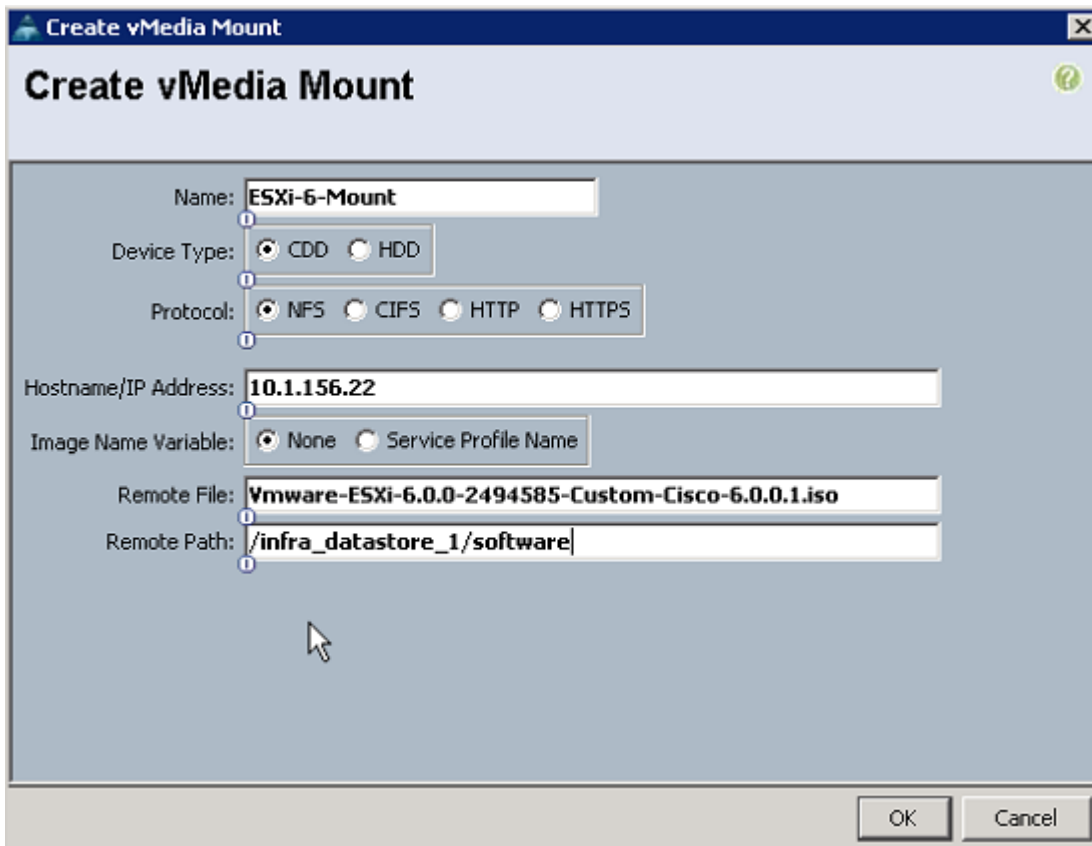
```
network port vlan create -node clus-01 -vlan-name a0a-113
```

```
broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500 -ports
clus-01:a0a-113, clus-02:a0a-113
```

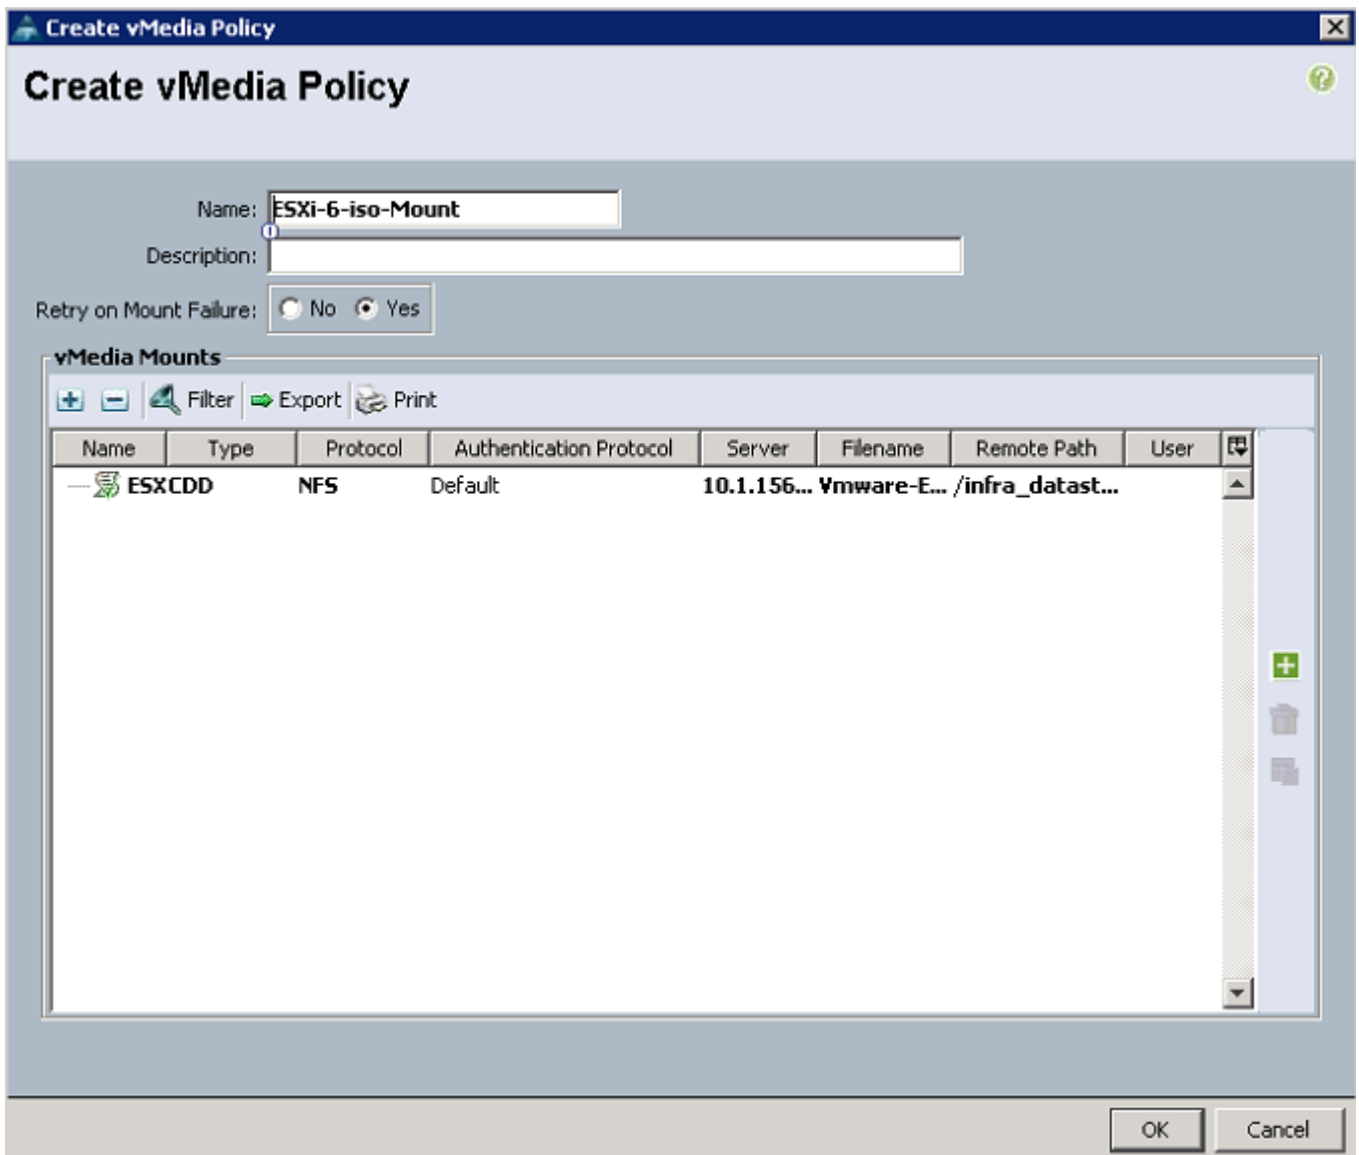
5. To Create an additional network interface for the infrastructure datastore to be accessed from the Inband mgmt network, enter the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_IB-MGMT -role
data -data-protocol nfs -home-node clus-02 -home-port a0a-113 -
address <<var_inband_nfs_ip>> -netmask
<<var_inband_mgmt_vlan_mask>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true -
failover-group IB-MGMT
```

6. From the vSphere interface; Select the home tab and select Storage.
7. Expand FlexPod_DC and Infra_datastore 1.
8. Right-click and select Browse Files.
9. Click the third icon to create a new folder; name this folder software and click Create.
10. Select the software folder in the list on the left; click the first icon to upload a file to the datastore.
11. Browse to the VMware Cisco custom ISO and upload it to the datastore.
12. Log in to the Cisco UCS Manager.
13. Select the servers Tab , Policies → root → vMedia Policies.
14. Right-click vMedia Policies and select create vMedia Policy.
15. Fill in the Policy name and then click the green + sign.
16. Enter the vMedia Mount name, Ip Address, Remote file and remote path.



17. Select OK in the create vMedia mount window.




18. Click OK in the Create vMedia Policy window, select OK to pop up window.

19. Select Service Profile Templates → root.

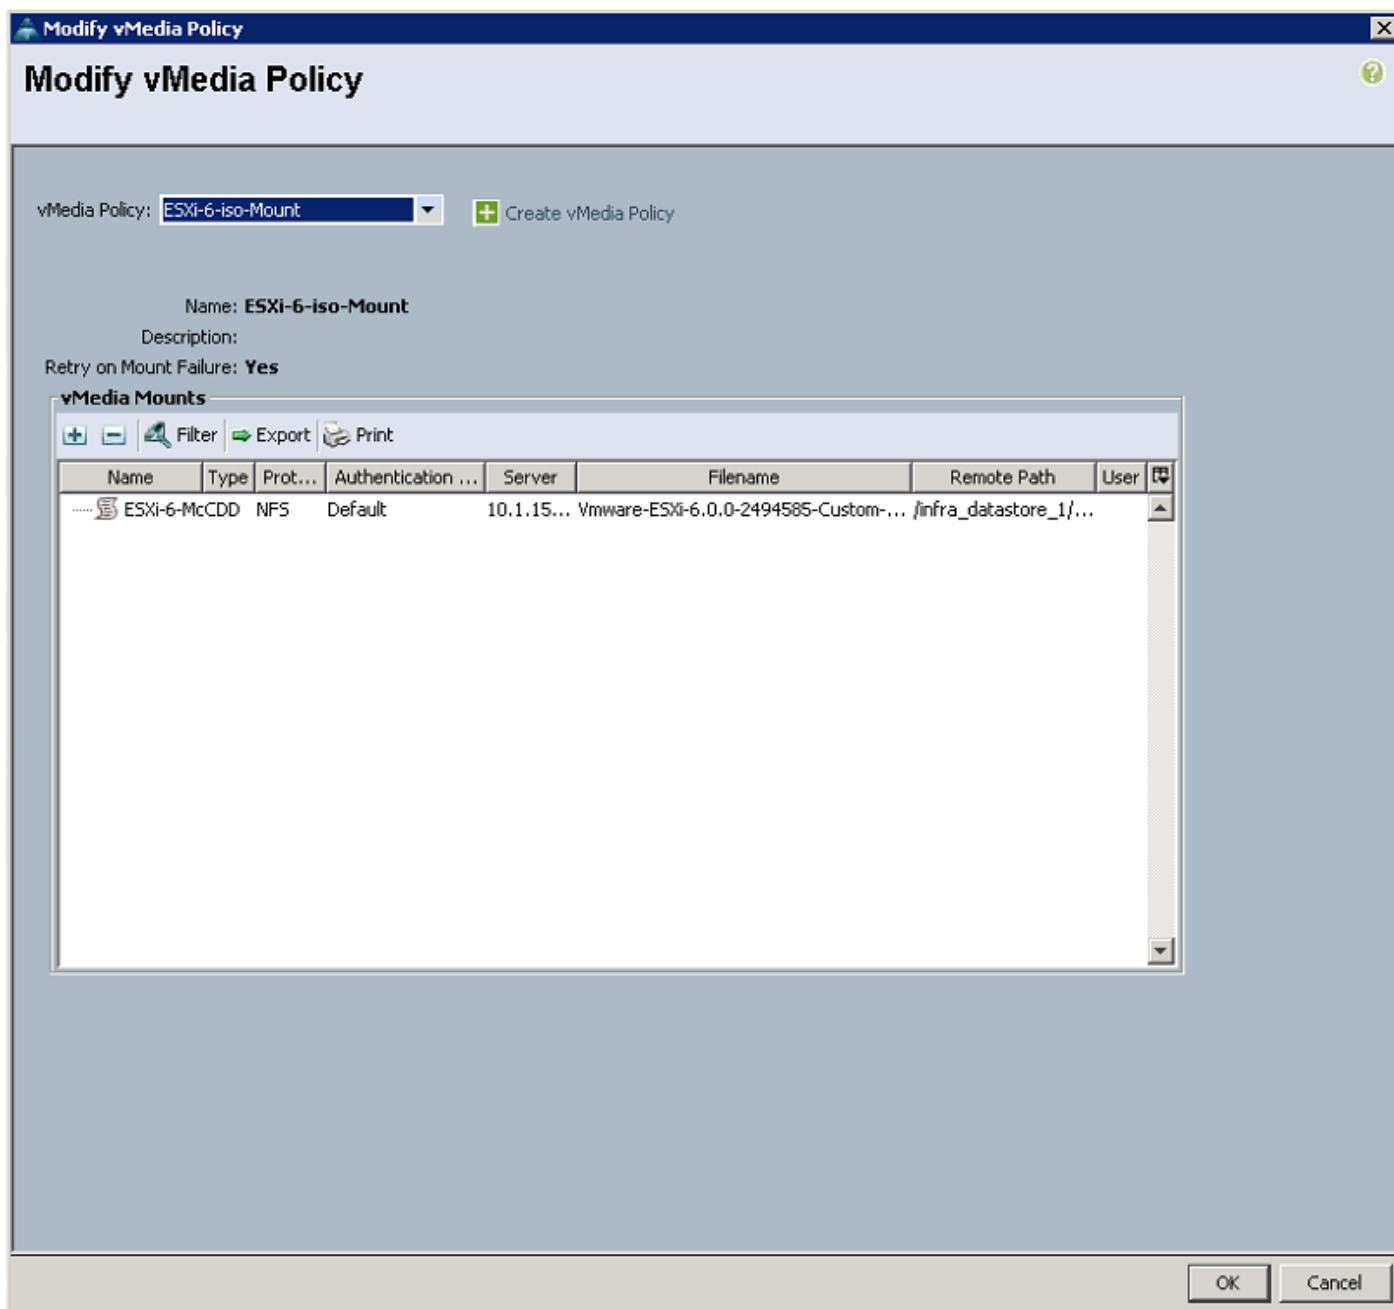
20. Right click Service Profile Template and select Create a Clone.

21. Name the Clone, and select the root Org.



22. Select OK to close the create clone window.

23. Select the Template that was just created and select the vMedia Policy tab.
24. In the Actions Pane, select Modify vMedia Policy.
25. Select the ESXi-6-iso-Mount policy from the drop-down menu.



26. Click OK to close the Modify vMedia Policy window.
27. Click OK in the pop up window.



For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to create the service profile. On first boot the host will boot into the ESXi installer. After ESXi is installed,

you can unbind from the vMedia Service Profile Template and bind to the original Service Profile Template and the ESXi installer will not be automatically mounted.

FlexPod Cisco Nexus 1110-X and 1000V vSphere

This section provides detailed procedures for installing a pair of high-availability (HA) Cisco Nexus 1110-X Virtual Services Appliances (VSAs) in a FlexPod configuration. This validation effort used a preexisting management infrastructure to support the VSA devices and therefore does not document the cabling configuration.

Primary and standby Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) are installed on the 1110-Xs and Cisco Nexus 1000V distributed virtual switch (DVS) will be provisioned. This procedure assumes that the Cisco Nexus 1000V software version 5.2(1)SV3(1.5b) has been downloaded from [Cisco Nexus 1000V Download Link](#) and expanded. It is recommended to install software version 5.2(1)SP1(7.3) on the Nexus 1110-Xs using [Cisco Nexus Cloud Services Platform Software Installation and Upgrade Guide](#). Additionally, this procedure assumes that Cisco Virtual Switch Update Manager (VSUM) version 1.5.3 has been downloaded from [Cisco VSUM Download Link](#) and expanded. This procedure also assumes that VMware vSphere 6.0 Enterprise Plus licensing is installed.

Configure CIMC Interface on Both Cisco Nexus 1110-Xs

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

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

1. Using the supplied dongle, connect a monitor and USB keyboard to the KVM console port on the front of the Cisco Nexus 1110-X virtual appliance.
2. Reboot the virtual appliance.
3. Press F8 when prompted to configure the CIMC interface.
4. Using the spacebar, set the NIC mode to Dedicated.
5. Clear the checkbox for DHCP enabled.
6. Set the CIMC IP address (<<var_cimc_ip>>) in the out-of-band management VLAN.
7. Set the CIMC subnet mask (<<var_cimc_mask>>).
8. Set the CIMC gateway (<<var_cimc_gateway>>).
9. Set the NIC redundancy to None.
10. Set and reenter the CIMC default password (<<var_password>>).
11. Press F10 to save the configuration.
12. Continue pressing F5 until Network settings configured is shown.

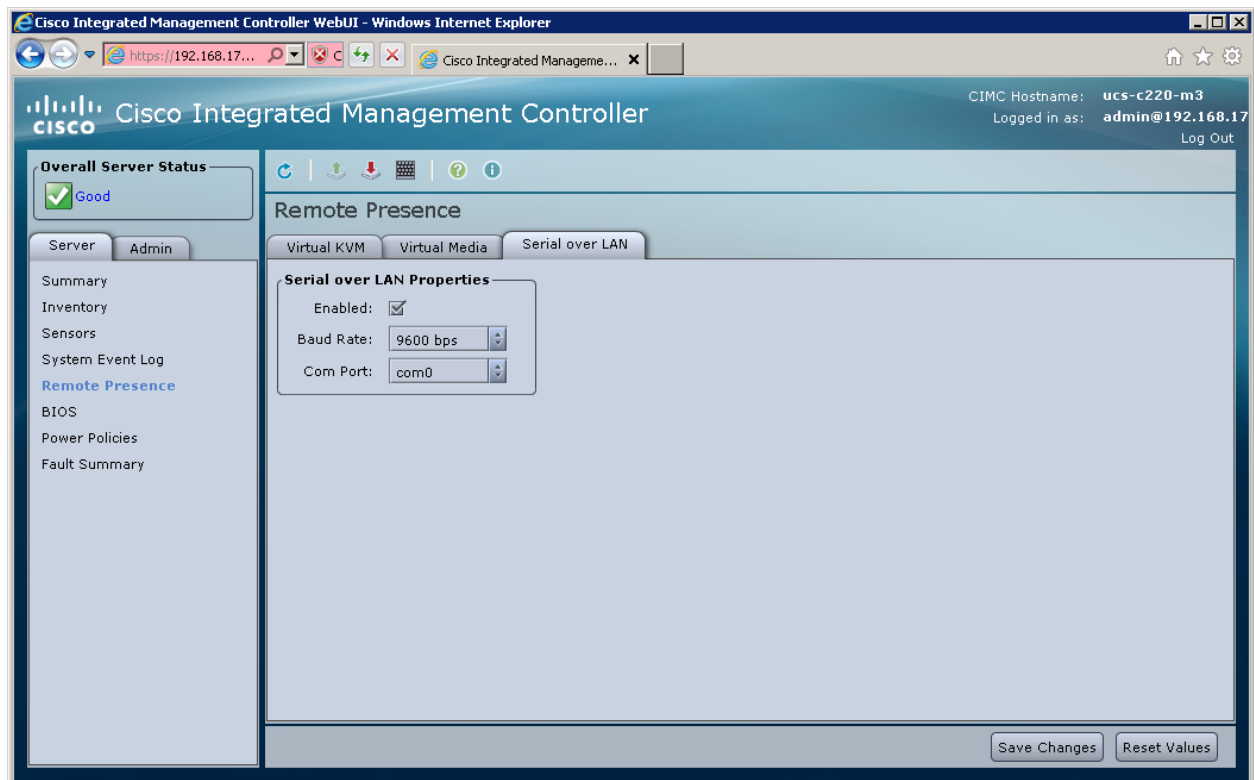
13. Press Esc to reboot the virtual appliance.

Configure Serial over LAN for Both Cisco Nexus 1110-Xs

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

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

1. Use a Web browser to open the URL at `http://<<var_cimc_ip>>`.
2. Log in to the CIMC with the admin user id and the CIMC default password (`<<var_password>>`).
3. In the left column, click Remote Presence.
4. Click the Serial over LAN tab.
5. Select the Enabled checkbox for Serial over LAN Properties.
6. From the Baud Rate drop-down menu, select 9600 bps.
7. Click Save Changes.



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

```

192.168.171.127 - PuTTY
login as: admin
admin@192.168.171.127's password:
ucs-c220-m3# connect host
CISCO Serial Over LAN:
Close Network Connection to Exit

Invalid admin password. Please try again.

Enter the password for "admin": █

```

Configure Cisco Nexus 1110-X Virtual Appliances

Cisco Nexus 1110-X A

To configure Cisco Nexus 1110-X A, complete the following steps:

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

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

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

Enter HA role[primary/secondary]: primary

Enter the domain id<1-4095>: <<var_vsa_domain_id>>

Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>

Control Channel Setup.

Choose Uplink: < Gig:1,2 10Gig:7,8 NewPortChannel:0 >[0]: Enter

Choose type of portchannel <ha/lacp>[ha]: lacp

PortChannell - Choose uplinks < Gig:1,2 10Gig:7,8 >[1,2]: 7,8

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

Management Channel setup

Choose Uplink: < Gig:1,2 Po1:9 NewPortChannel:0 >[9]: Enter


```

Would you like to enter the basic system configuration dialogue (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the VSA name : <<var_1110x_vsa>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IP address type V4/V6? (V4): V4
Mgmt0 IPv4 address : <<var_1110x_vsa_ip>>
Mgmt0 IPv4 netmask : <<var_1110x_vsa_mask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <<var_1110x_vsa_gateway>>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (das/rsa) [rsa]: Enter
Number of rsa key bits <768-2048> [1024]: Enter
Enable the http server? (yes/no) [y]: Enter
Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address: <<var_switch_a_ntp_ip>>

```

2. Review the configuration summary. If everything is correct, enter no to skip editing the configuration.

```

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

```

3. The Cisco Nexus 1110-X saves the configuration and reboots. After reboot, log back in as admin.

Cisco Nexus 1110-X B

To configure the Cisco Nexus 1110-X B, complete the following steps:

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

```

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

```



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

2. Enter the admin password again to confirm: <<var_password>>.

```

Enter HA role[primary/secondary]: secondary
Enter the domain id<1-4095>: <<var_vsa_domain_id>>

```



This is the same domain id that you entered on the primary Cisco Nexus 1110-X.

```

Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
  Control Channel Setup.
Choose Uplink: < Gig:1,2 10Gig:7,8 NewPortChannel:0 >[0]: Enter
  Choose type of portchannel <ha/lacp>[ha]: lacp
PortChannel1 - Choose uplinks < Gig:1,2 10Gig:7,8 >[1,2]: 7,8
Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>
  Management Channel setup
  Choose Uplink: < Gig:1,2 Po1:9 NewPortChannel:0 >[9]: Enter

```

3. The Cisco Nexus 1110-X saves the configuration and reboots.

Set Up the Primary Cisco Nexus 1000V VSM

Cisco Nexus 1110-X A

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, complete the following steps:



These steps are completed from the primary Nexus 1110-X A

1. Continue periodically running the following command until module 2 (Cisco Nexus 1110-X B) has a status of ha-standby.

```
show module
```

2. Enter the global configuration mode and create a virtual service blade.

```
config t
```

```
virtual-service-blade VSM-1
```

```
dir /repository
```

3. If the desired Cisco Nexus 1000V ISO file (n1000v-dk9.5.2.1.SV3.1.5b.iso) is not present on the Cisco Nexus 1110-X, run the copy command to copy it to the Cisco Nexus 1110-X disk. You must place the file either on an FTP server or on a UNIX or Linux® machine (using scp) that is accessible from the Cisco Nexus 1110-X management interface. An example copy command from an FTP server is copy ftp://<<var_ftp_server>>/n1000v-dk9.5.2.1.SV3.1.5b.iso /repository/.

```
virtual-service-blade-type new n1000v-dk9.5.2.1.SV3.1.5b.iso
```

```
interface control vlan <<var_pkt-ctrl_vlan_id>>
```

```
interface packet vlan <<var_pkt-ctrl_vlan_id>>
```

```
enable primary
```

```
Enter vsb image:[n1000v-dk9.5.2.1.SV3.1.5b.iso] Enter
```

Enter domain id[1-4095]: <<var_vsm_domain_id>>



This domain ID should be different than the VSA domain ID.

Enter SVS Control mode (L2 / L3): [L3] Enter

Management IP version [V4/V6]: [V4] Enter

Enter Management IP address: <<var_vsm_mgmt_ip>>

Enter Management subnet mask: <<var_vsm_mgmt_mask>>

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

Enter HostName: <<var_vsm_hostname>>

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



This password must be entered with only uppercase and lowercase letters. No special characters can be used in this password.

Do you want to continue with installation with entered details (Y/N)? [Y] Enter
copy run start

4. Run show virtual-service-blade summary. Continue periodically entering this command until the primary VSM-1 has a state of VSB POWERED ON.
5. Modify the management, control and packet interface and set PortChannel 1 as the uplink interface (if needed):

```
virtual-service-blade VSM-1
  interface control uplink PortChannel1
  interface management uplink PortChannel1
  interface packet uplink PortChannel1
```

Set Up the Secondary Cisco Nexus 1000V VSM

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

Cisco Nexus 1110-X A

```
enable secondary
Enter vsb image: [n1000v-dk9.5.2.1.SV3.1.5b.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_mgmt_ip>>
```

Enter Management subnet mask: <<var_vsm_mgmt_mask>>

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

Enter HostName: <<var_vsm_hostname>>

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

This password must be entered with only uppercase and lowercase letters. No special characters can be used in this password. Do you want to continue installation with entered details (Y/N)? [Y]

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

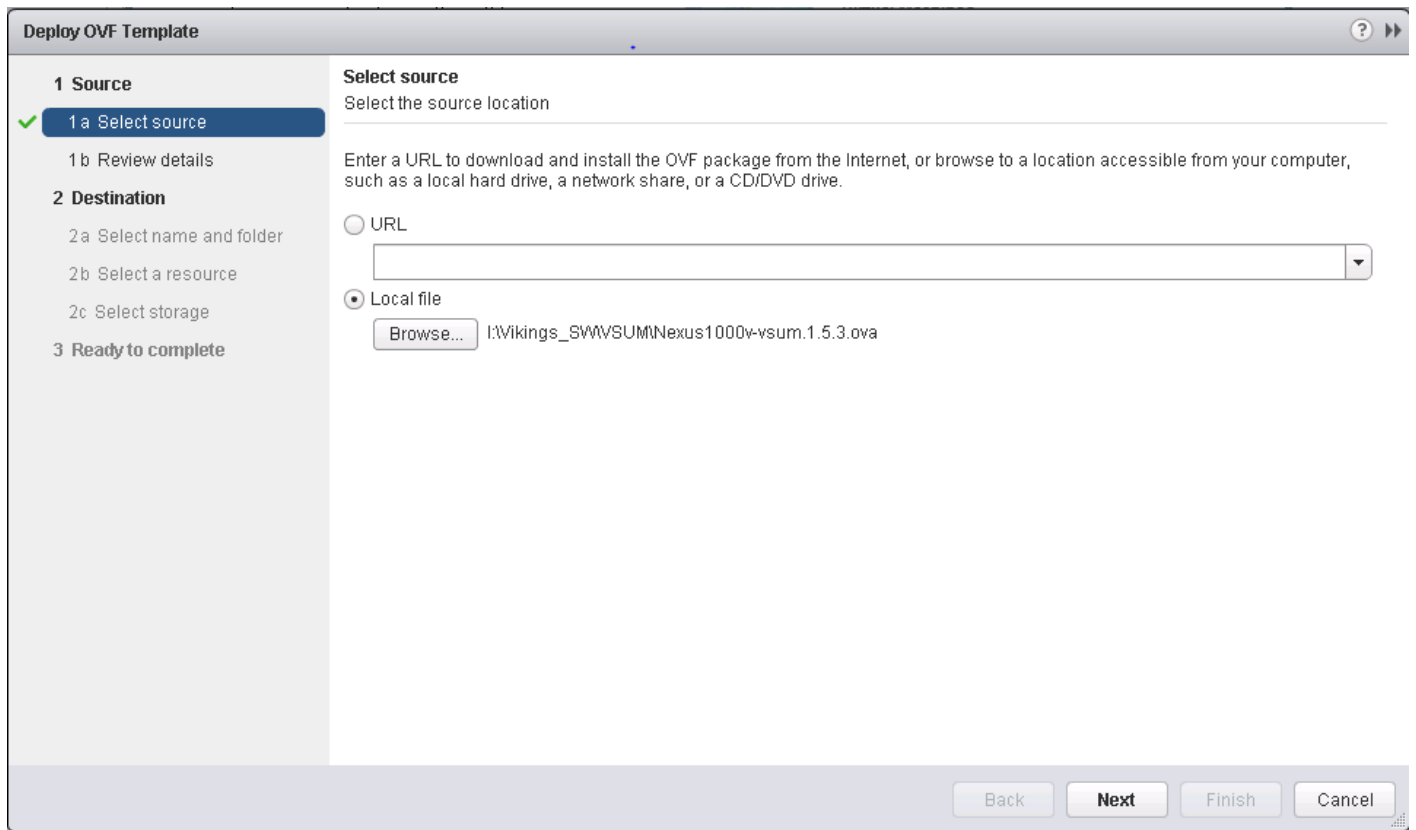
`copy run start`

Install Cisco Virtual Switch Update Manager

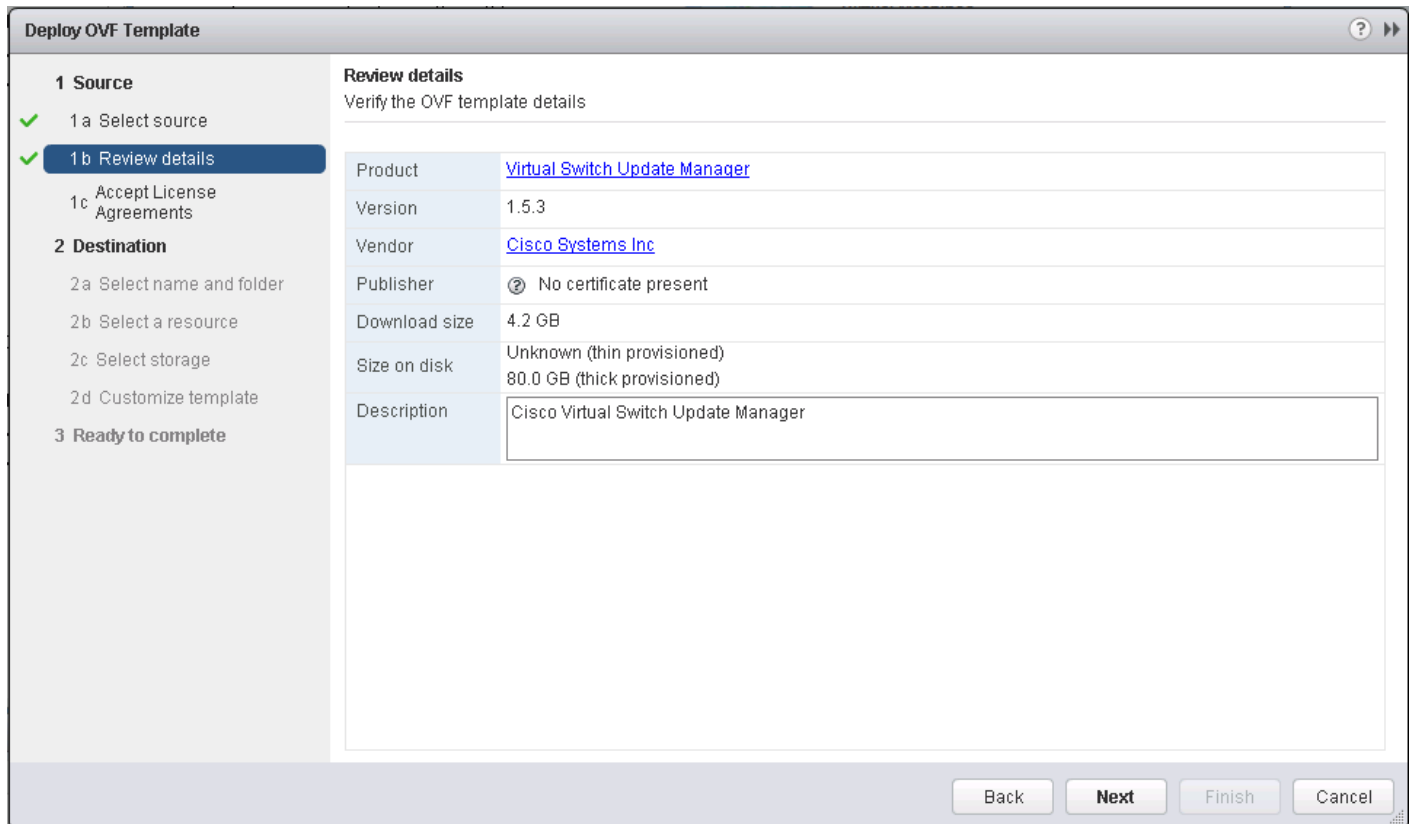
VMware vSphere Web Client

To install the Cisco Virtual Switch Upgrade Manager from OVA in the VMware virtual environment, complete the following steps:

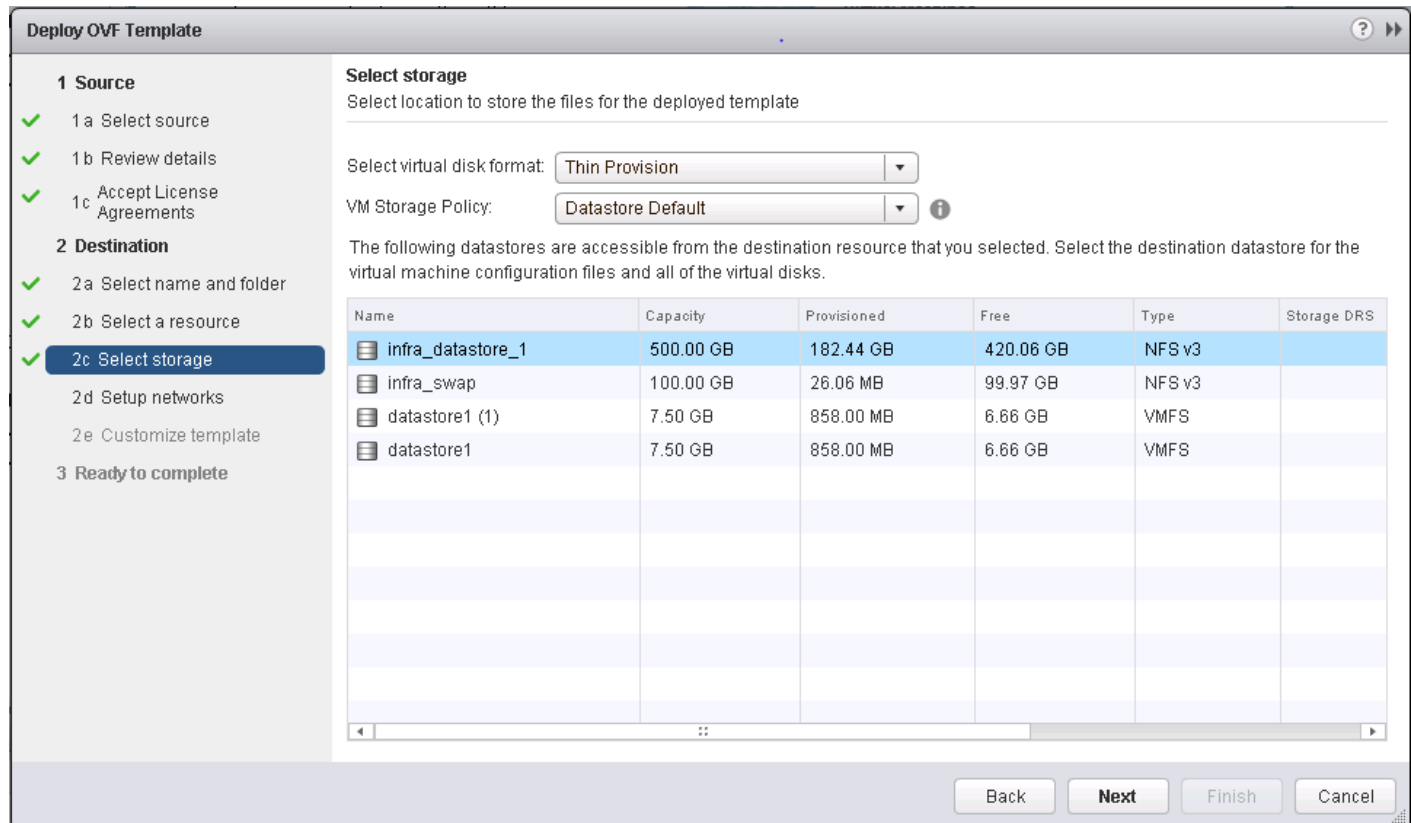
1. Log into the VMware vSphere Web Client.
2. In the pane on the right, click VMs and Templates.
3. In the center pane, select Actions > Deploy OVF Template.
4. Select Browse and browse to and select the Nexus1000v-vsum.1.5.3.ova file.
5. Click Open.
6. Click Next.



7. Review the details and click Next.



8. Click Accept to accept the License Agreement and click Next.
9. Name the Virtual Machine, select the FlexPod_DC datacenter and click Next.
10. Select the FlexPod_Management cluster and click Next.
11. Select infra_datastore_1 and the Thin Provision virtual disk format and click Next.



12. Select the MGMT Network and click Next.
13. Fill in the Networking Properties.
14. Expand the vCenter Properties and fill those in.
15. Click Next.
16. Review all settings and click Finish.
17. Wait for the Deploy OVF template task to complete.
18. Select the Home button in VMware vSphere Web Client and select Hosts and Clusters.
19. Expand the FlexPod_Management cluster and select the Virtual Switch Update Manager VM.
20. In the center pane, select Launch Remote Console. If a security warning pops up, click Allow.

21. If a security certificate warning pops up, click Connect Anyway.
22. Power on the Virtual Switch Update Manager VM.
23. Once the VM has completely booted up, log out and log back into the VMware vSphere Web Client.

Register the Cisco Nexus 1000V in VMware vCenter

VMware vSphere Web Client

To register the Cisco Nexus 1000V, complete the following steps:

1. After logging back into the VMware vSphere Web Client, Cisco Virtual Switch Update Manager should now appear under the Home tab. Select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Click Install.
4. In the pane on the right, select FlexPod_DC.
5. Under Nexus1000v Switch Deployment Process, select I already have a control plane (VSM) deployed.
6. Enter the IP Address of the VSM and the admin password.
7. Click Finish.
8. Click the Home button.
9. Select Cisco Virtual Switch Update manager.
10. Under Basic tasks, select Nexus 1000v.
11. Click Configure.
12. In the pane on the right, select FlexPod_DC.
13. The Nexus 1000v Switch should appear under the *Choose an associated Distributed Virtual Switch* section.

Perform Base Configuration of the Primary VSM

SSH Connection to Primary VSM

To perform the base configuration of the primary VSM, complete the following steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands.

```
config t
```

```

ntp server <<var_switch_a_ntp_ip>> use-vrf management
ntp server <<var_switch_b_ntp_ip>> use-vrf management
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
vlan <<var_native_vlan_id>>
name Native-VLAN
vlan <<var_iscsi_a_vlan_id>>
name iSCSI-A-VLAN
vlan <<var_iscsi_b_vlan_id>>
name iSCSI-B-VLAN
vlan <<var_pkt-ctrl_vlan_id>>
name Pkt-Ctrl-VLAN
exit
port-profile type ethernet system-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>

```



Any VLAN that has a VMKernel port should be assigned as a system vlan on both the uplink and the vEthernet ports of the virtual switch.

```

system mtu 9000
state enabled
port-profile type ethernet iscsi-a-uplink

```



```
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_a_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_a_vlan_id>>
no shutdown
system vlan <<var_iscsi_a_vlan_id>>
system mtu 9000
state enabled
port-profile type ethernet iscsi-b-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_b_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
system mtu 9000
state enabled
port-profile type vethernet IB-MGMT-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet NFS-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_nfs_vlan_id>>
no shutdown
system vlan <<var_nfs_vlan_id>>
state enabled
port-profile type vethernet vMotion-VLAN
vmware port-group
```

```
switchport mode access
switchport access vlan <<var_vmotion_vlan_id>>
no shutdown
system vlan <<var_vmotion_vlan_id>>
state enabled
port-profile type vethernet VM-Traffic-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vm-traffic_vlan_id>>
no shutdown
system vlan <<var_vm-traffic_vlan_id>>
state enabled
port-profile type vethernet nlkv-L3
capability l3control
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet iSCSI-A-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_a_vlan_id>>
no shutdown
system vlan <<var_iscsi_a_vlan_id>>
state enabled
port-profile type vethernet iSCSI-B-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
```

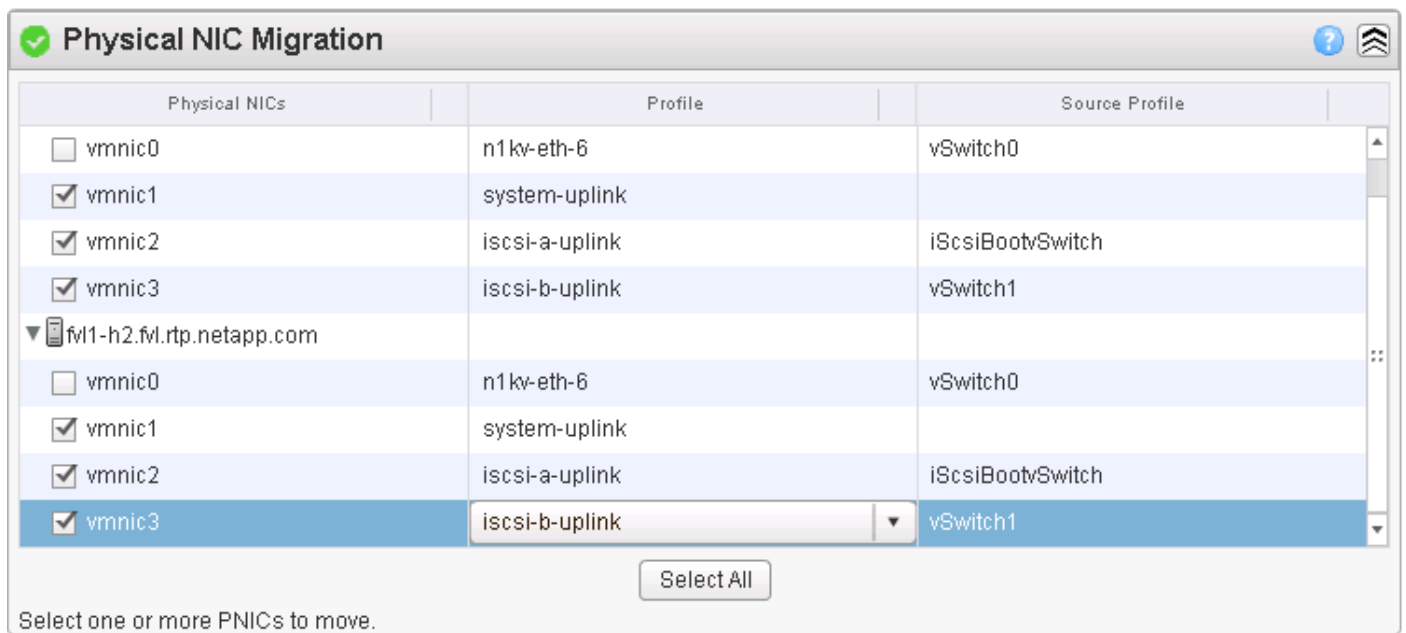
```
state enabled  
exit  
copy run start
```

Add VMware ESXi Hosts to Cisco Nexus 1000V

VMware vSphere Web Client

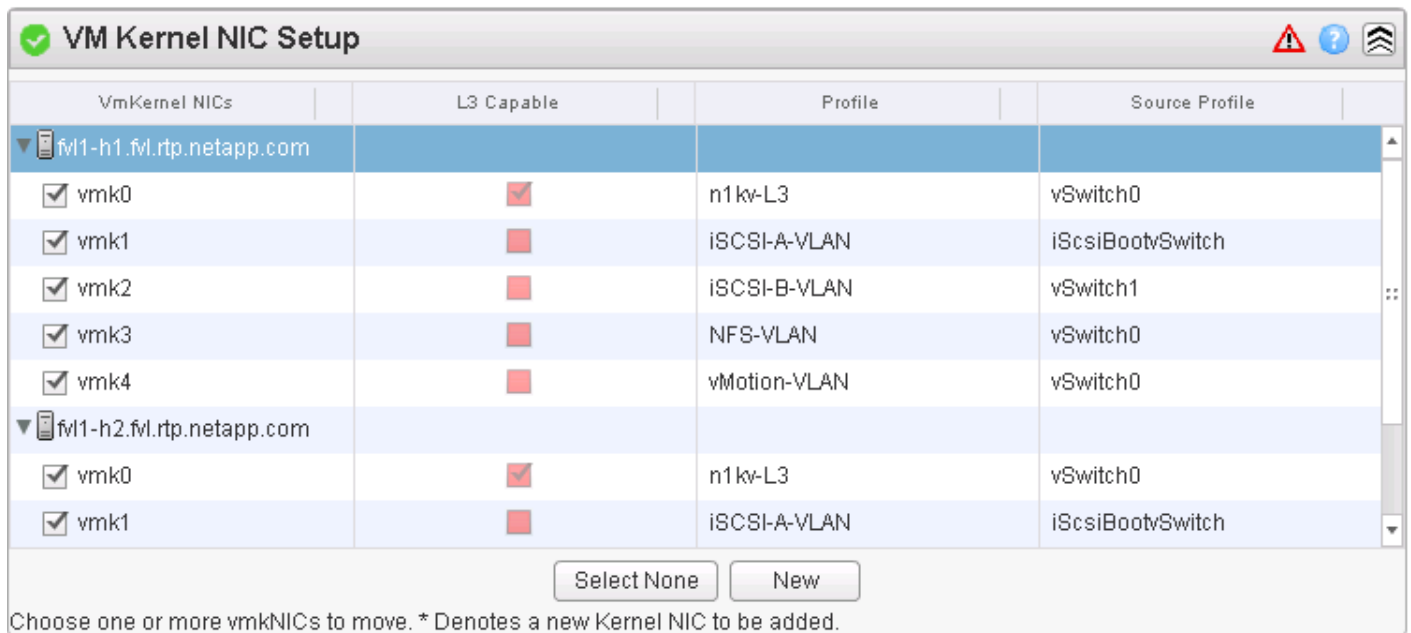
To add VMware ESXi hosts, complete the following steps:

1. Back in the VMware vSphere Web Client, from the Home tab, select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Select Configure.
4. Select the FlexPod_DC datacenter on the right.
5. Select the VSM on the lower right.
6. Click Manage.
7. In the center pane, select the Add Host tab.
8. Expand the FlexPod_Management ESXi Cluster and select both FlexPod Management Hosts.
9. Click Suggest.
10. Scroll down to Physical NIC Migration and expand each ESXi host.
11. On both hosts, unselect vmnic0, and select vmnic1. For vmnic1, select the system-uplink Profile. Select vmnic2 and select the iscsi-a-uplink Profile. Select vmnic3 and select the iscsi-b-uplink Profile.



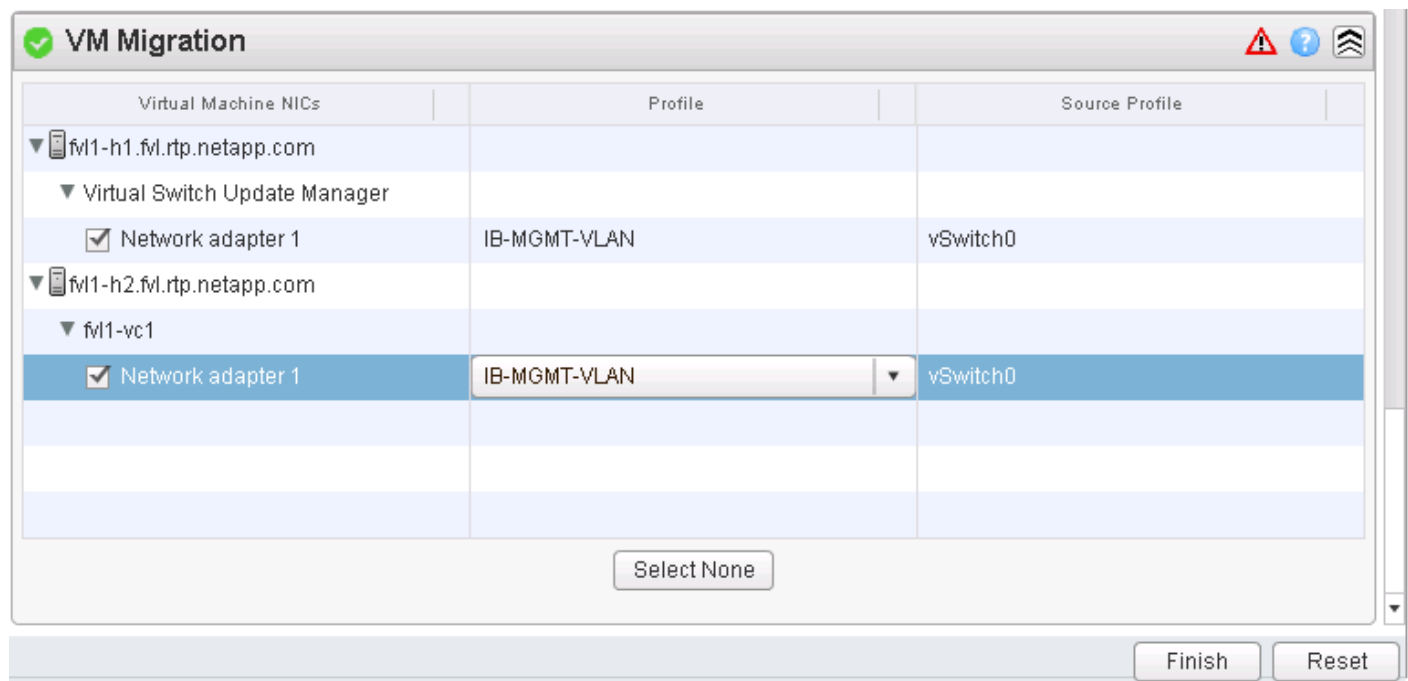
12. Scroll down to VM Kernel NIC Setup and expand both ESXi hosts.

13. All VMkernel ports should already have the appropriate checkboxes selected.



14. Scroll down to VM Migration and expand both ESXi hosts.

15. Select the IB-MGMT-VLAN profile for the VSUM and vCenter Virtual Machines.



16. Click Finish.



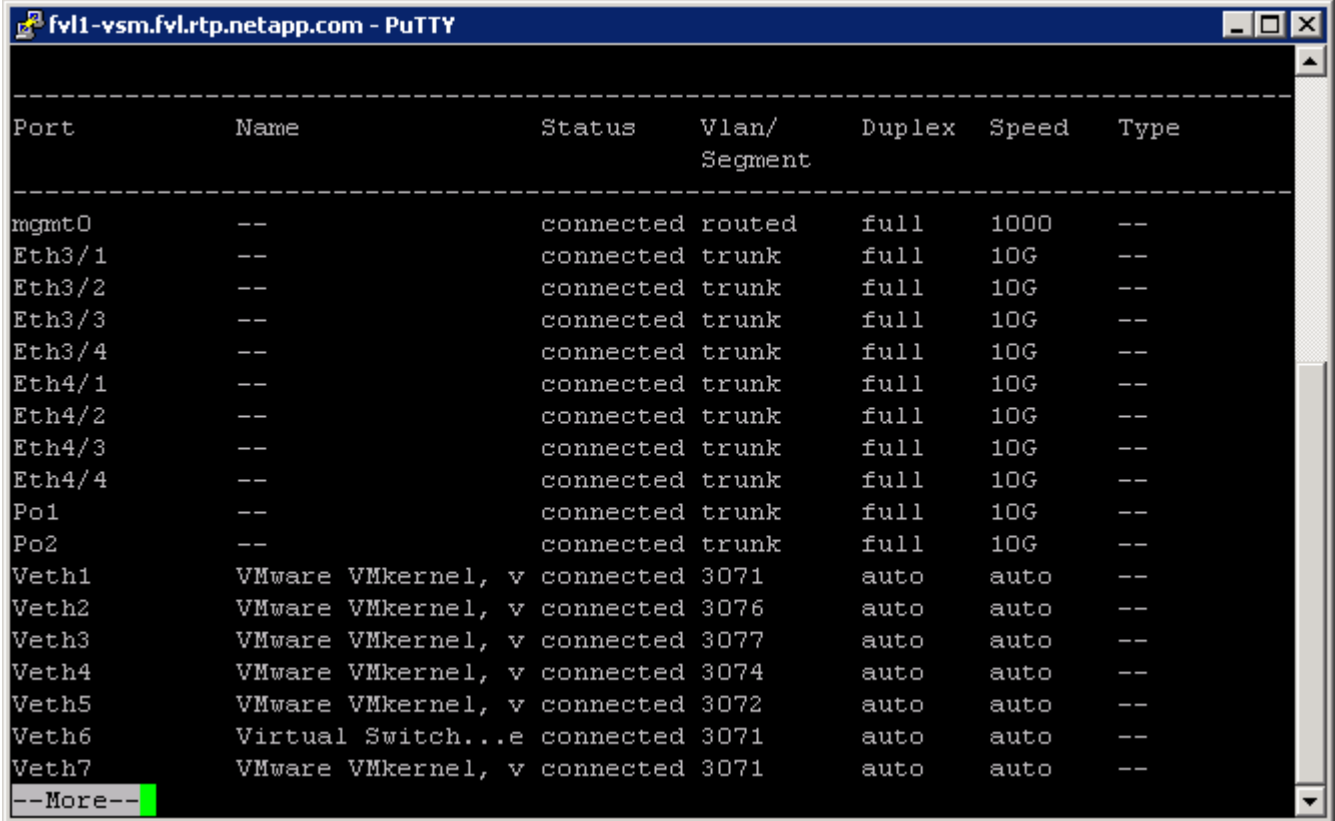
The progress of the virtual switch installation can be monitored from the c# interface.

Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V

To migrate the ESXi host redundant network ports, complete the following steps:

1. In the VMware vSphere Web Client window, select Home > Hosts and Clusters.
2. On the left expand the Datacenter and cluster, and select the first VMware ESXi host.
3. In the center pane, select the Manage tab, then select Networking.
4. Select vSwitch0. All of the port groups on vSwitch0 should be empty. Click the red X under Virtual switches to delete vSwitch0.
5. Click Yes to remove vSwitch0. It may be necessary to refresh the Web Client to see the deletion.
6. Delete iScsiBootvSwitch and vSwitch1.
7. The Nexus 1000V VSM should now be the only virtual switch. Select it and select the third icon above it under Virtual switches (Manage the physical network adapters connected to the selected switch).
8. Click the green plus sign to add an adapter.
9. For UpLink03, select the system-uplink port group and make sure vmnic0 is the Network adapter. Click OK.

10. Click OK to complete adding the Uplink. It may be necessary to refresh the Web Client to see the addition.
11. Repeat this procedure for the second ESXi host.
12. From the SSH client that is connected to the Cisco Nexus 1000V, run show interface status to verify that all interfaces and port channels have been correctly configured.



```
fv11-vsm.fvl.rtp.netapp.com - PuTTY
-----
Port          Name          Status      Vlan/Segment Duplex  Speed  Type
-----
mgmt0         --           connected  routed        full   1000   --
Eth3/1        --           connected  trunk         full   10G    --
Eth3/2        --           connected  trunk         full   10G    --
Eth3/3        --           connected  trunk         full   10G    --
Eth3/4        --           connected  trunk         full   10G    --
Eth4/1        --           connected  trunk         full   10G    --
Eth4/2        --           connected  trunk         full   10G    --
Eth4/3        --           connected  trunk         full   10G    --
Eth4/4        --           connected  trunk         full   10G    --
Po1           --           connected  trunk         full   10G    --
Po2           --           connected  trunk         full   10G    --
Veth1         VMware VMkernel, v connected 3071   auto   auto   --
Veth2         VMware VMkernel, v connected 3076   auto   auto   --
Veth3         VMware VMkernel, v connected 3077   auto   auto   --
Veth4         VMware VMkernel, v connected 3074   auto   auto   --
Veth5         VMware VMkernel, v connected 3072   auto   auto   --
Veth6         Virtual Switch...e connected 3071   auto   auto   --
Veth7         VMware VMkernel, v connected 3071   auto   auto   --
--More--
```

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

```

fv11-vsm.fvl.rtp.netapp.com - PuTTY
fv11-vsm# sho module
Mod  Ports  Module-Type          Model          Status
-----
1    0      Virtual Supervisor   Nexus1000V    active *
2    0      Virtual Supervisor   Nexus1000V    ha-standby
3    1022   Virtual Ethernet    NA            ok
4    1022   Virtual Ethernet    NA            ok

Mod  Sw          Hw
-----
1    5.2(1)SV3(1.5b)  0.0
2    5.2(1)SV3(1.5b)  0.0
3    5.2(1)SV3(1.5b)  VMware ESXi 6.0.0 Releasebuild-2494585 (6.0)
4    5.2(1)SV3(1.5b)  VMware ESXi 6.0.0 Releasebuild-2494585 (6.0)

Mod  Server-IP      Server-UUID          Server-Name
-----
1    172.20.71.44   NA                   NA
2    172.20.71.44   NA                   NA
3    172.20.71.26   722201d0-e027-e511-0000-000000110001  fv11-h1
4    172.20.71.27   722201d0-e027-e511-0000-000000110002  fv11-h2

* this terminal session
fv11-vsm# █

```

14. Run copy run start.

Cisco Nexus 1000V vTracker

SSH Connection to Primary VSM

The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network environment. To connect SSH to the primary VSM, complete the following steps:

1. From an ssh interface connected to the Cisco Nexus 1000V VSM, enter the following:

```

config t
feature vtracker
copy run start
show vtracker upstream-view
show vtracker vm-view vnic
show vtracker vm-view info
show vtracker module-view pnic
show vtracker vlan-view

```

```
fv11-vsm.fvl.rtp.netapp.com - PuTTY
fv11-vsm(config)#
fv11-vsm(config)# show vtracker vlan-view

* R = Regular Vlan, P = Primary Vlan, C = Community Vlan
  I = Isolated Vlan, U = Invalid

-----
VLAN    Type  VethPort  VM Name                Adapter Name          Mod
-----
1       R     -         -                       -                     -
2       R     -         -                       -                     -
3071    R     Veth1     Module 3                vmk0                   3
        Veth6     Virtual Switch...e ManagerNet Adapter 1  3
        Veth7     Module 4                vmk0                   4
        Veth12    fv11-vc1                Net Adapter 1          3
        Veth13    fv11-vsc                Net Adapter 1          3
3072    R     Veth5     Module 3                vmk4                   3
        Veth11    Module 4                vmk4                   4
3073    R     -         -                       -                     -
3074    R     Veth4     Module 3                vmk3                   3
        Veth10    Module 4                vmk3                   4
3075    R     -         -                       -                     -
3076    R     Veth2     Module 3                vmk1                   3
        Veth8     Module 4                vmk1                   4
3077    R     Veth3     Module 3                vmk2                   3
        Veth9     Module 4                vmk2                   4
-----

fv11-vsm(config)#
```


FlexPod Management Tools Setup

NetApp Virtual Storage Console (VSC) 6.1 Deployment Procedure

This section describes the deployment procedures for the NetApp Virtual Storage Console (VSC).

VSC 6.1 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.3.1:

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

Install VSC 6.1


To install the VSC 6.1 software, complete the following steps:

1. Build a VSC virtual machine with Windows Server 2012 R2, 4GB RAM, two CPUs, and one virtual network interface in the <<var_ib_mgmt_vlan_id>> VLAN. The virtual network interface should be a VMXNET 3 adapter.
2. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain.
3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the User Interfaces and Infrastructure Feature on the VM.
4. Install all Windows updates on the VM.
5. Log in to the VSC VM as FlexPod admin user.
6. Download the x64 version of the [Virtual Storage Console 6.1](#) from the [NetApp Support](#) site.
7. From the VMware Console, right-click the VSC-6.1-win64.exe file downloaded in step 3 and select Run as administrator.
8. Select the appropriate language and click OK.
9. On the Installation wizard Welcome page, click Next.
10. Select the checkbox to accept the message, click Next.

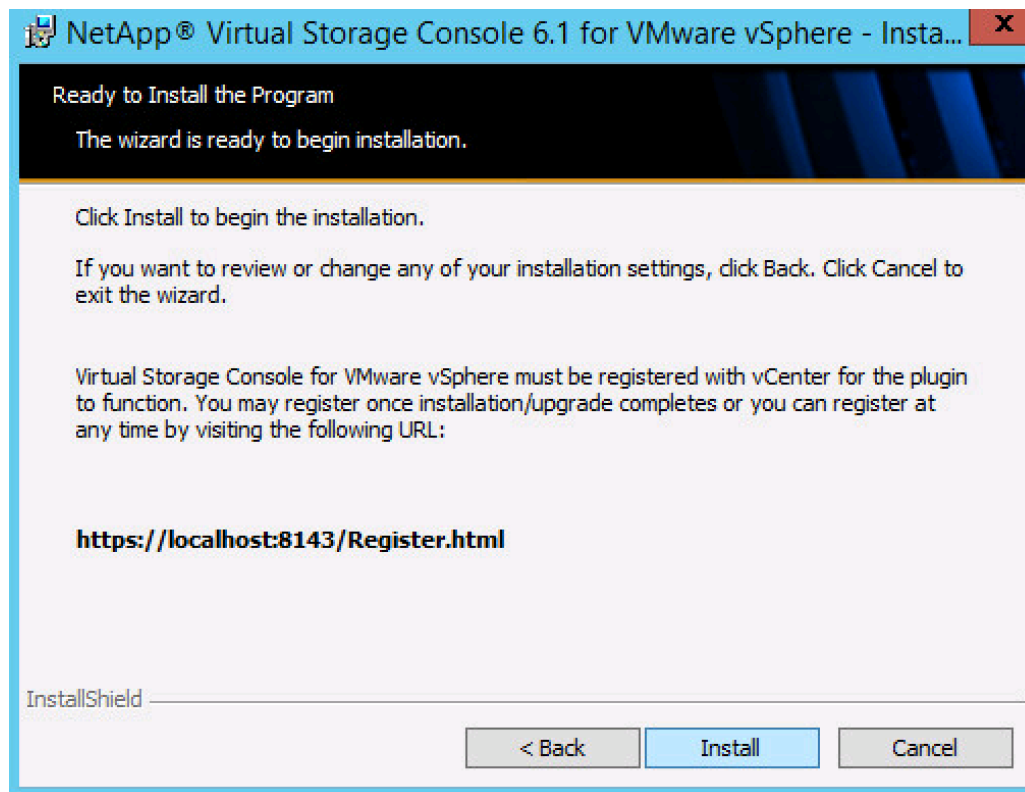


The Backup and Recovery capability requires an additional license.

11. Click Next to accept the default installation location.



12. Click Install.

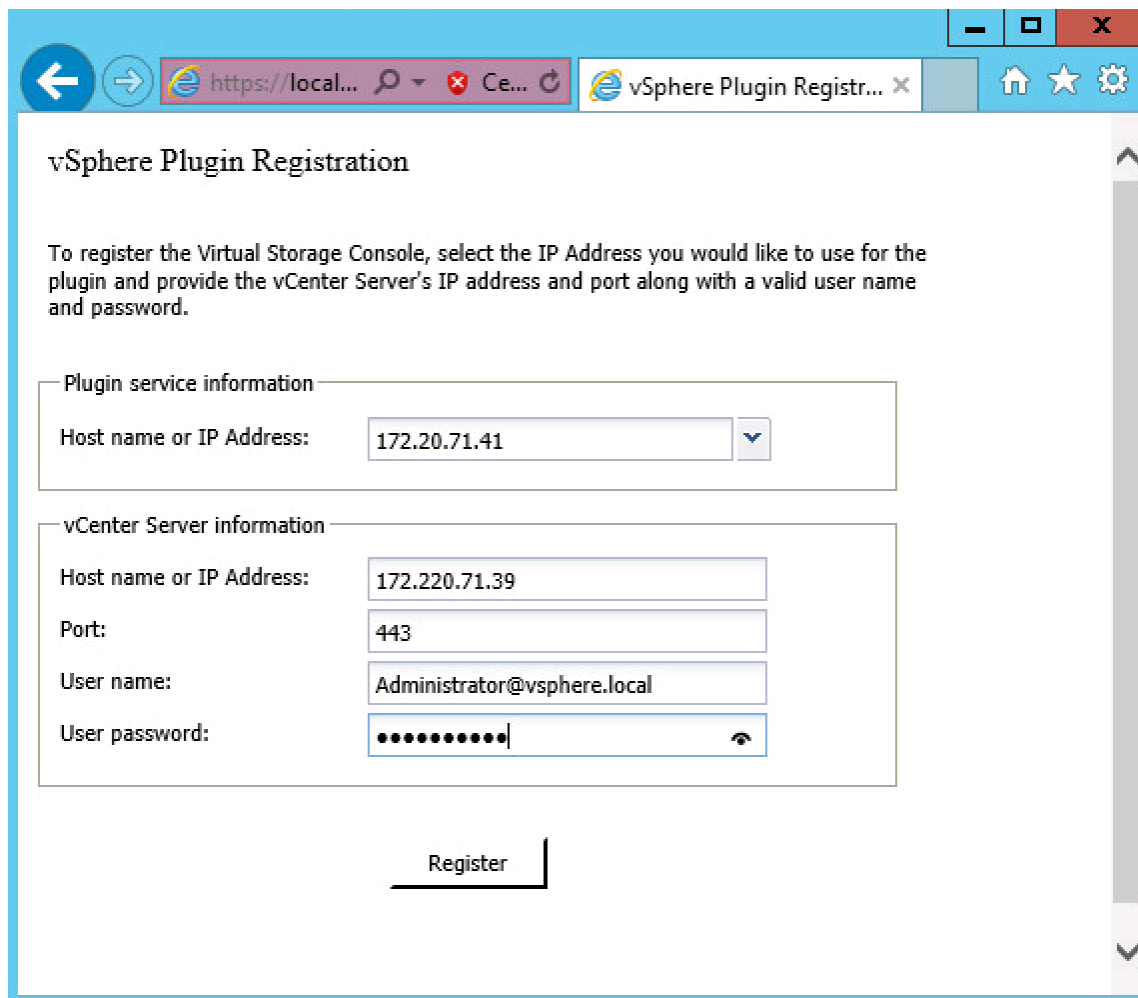


13. Click Finish.

Register VSC with vCenter Server

To register the VSC with the vCenter Server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open <https://localhost:8143/Register.html> in Internet Explorer.
2. Click Continue to this website (not recommended).
3. In the Plug-in Service Information section, select the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.
4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user or root), and user password for the vCenter Server. Click Register to complete the registration.



The screenshot shows a web browser window titled "vSphere Plugin Registration". The page contains the following text and form fields:

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:

vCenter Server information

Host name or IP Address:

Port:

User name:

User password:

5. After a successful registration, the storage controllers are discovered automatically.

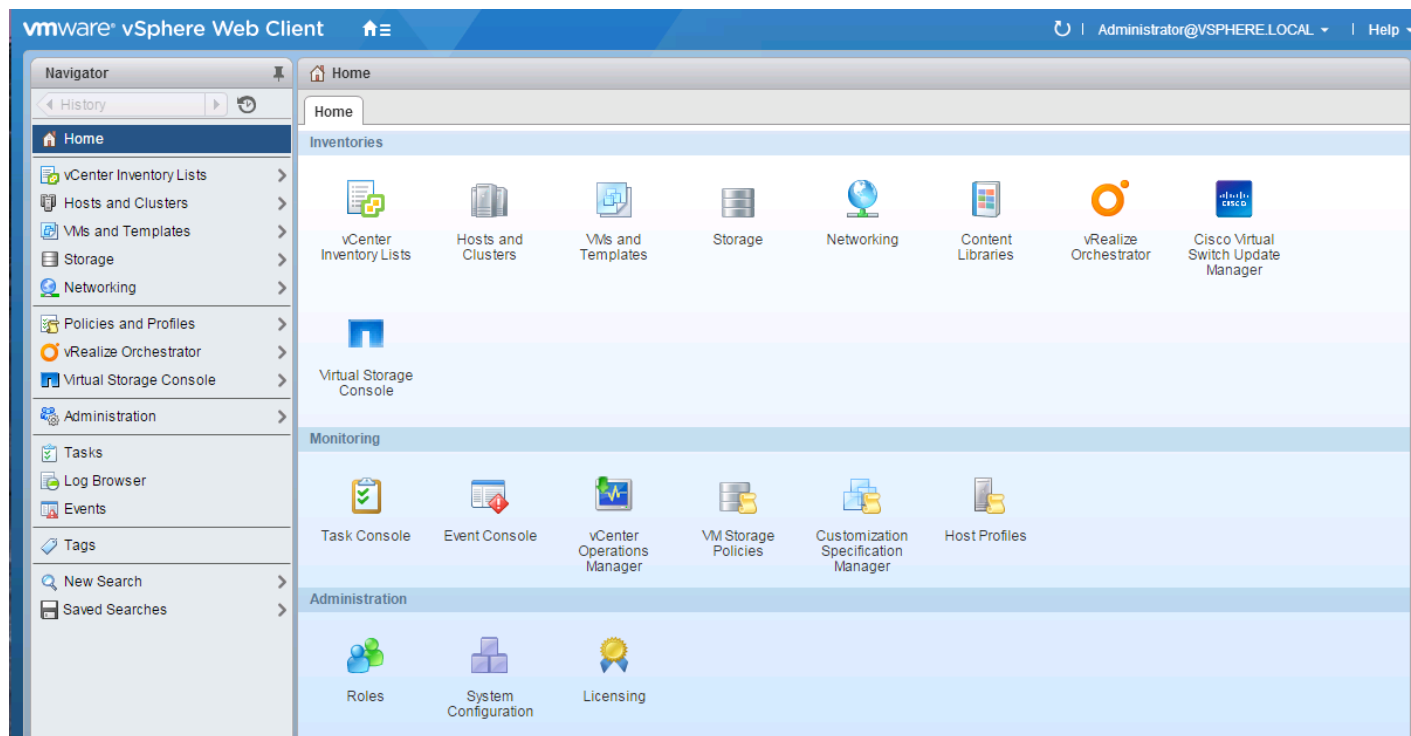


Storage discovery process will take some time.

Discover and Add Storage Resources

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

1. Using the vSphere web client, log in to the vCenter Server as FlexPod admin user or root. If the vSphere web client was previously opened, close it and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.

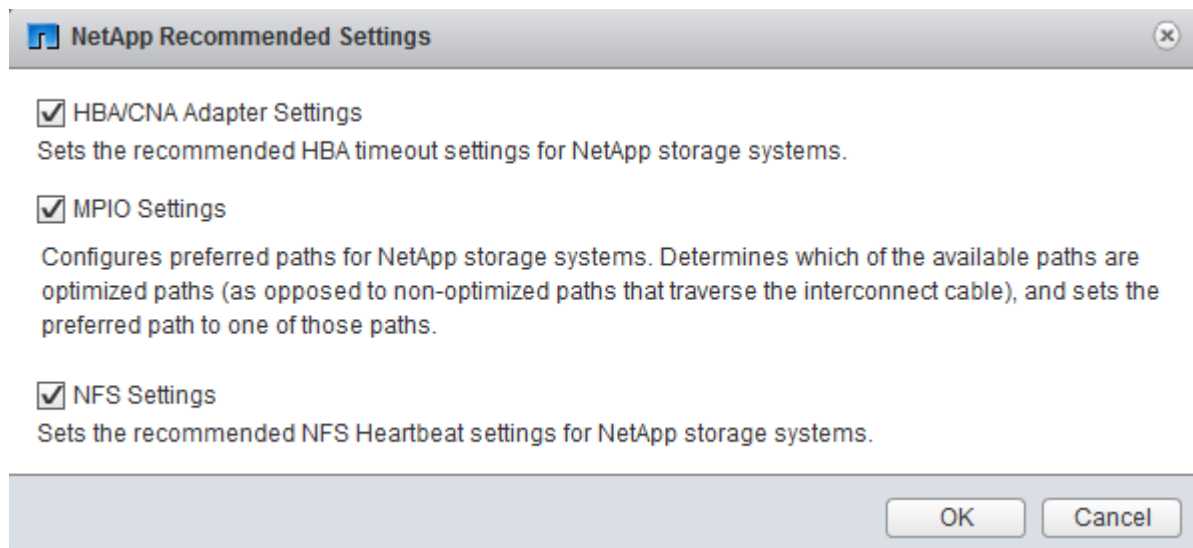


3. Select Storage Systems. Under the Objects tab, click Actions > Modify.
4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name, and the admin password for password. Confirm that Use SSL to connect to this storage system is selected. Click OK.
5. Click OK to accept the controller privileges.

Optimal Storage Settings for ESXi Hosts

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

1. From the Home screen, click vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for these hosts.




2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings.



This functionality sets values for HBAs and CNAs, sets appropriate paths, and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).

3. Click OK.



For each host where settings were adjusted in the previous step, place the host in Maintenance Mode, reboot the host, and exit Maintenance Mode.

VSC 6.1 Backup and Recovery

Prerequisites to Use Backup and Recovery Capability

Before you begin using the Backup and Recovery capability to schedule backups and restore your datastores, virtual machines, or virtual disk files, you must confirm that the storage systems that contain the datastores and virtual machines for which you are creating backups have valid storage credentials.



If you plan to leverage the SnapMirror update option, add all the destination storage systems with valid storage credentials.

Backup and Recovery Configuration

To configure a backup job for a datastore, complete the following steps:

1. From Home screen, select the Home tab and click Storage.
2. On the left, expand the Datacenter and select Datastores.
3. Right-click the datastore which you need to backup. Select NetApp VSC > Backup > Schedule Backup Job.



If you prefer a one-time backup, then choose Backup Now instead of Schedule Backup.

4. Type a backup job name and description.



If you want to create a VMware snapshot for each backup, select Perform VMware consistency snapshot in the options pane.

5. Click Next.
6. Select any options to include in the backup.

Schedule Backup

1 Details
 2 Options
 3 Spanned Entities
 4 Scripts
 5 Schedule and Retention
 6 Credentials and Alerts
 7 Summary

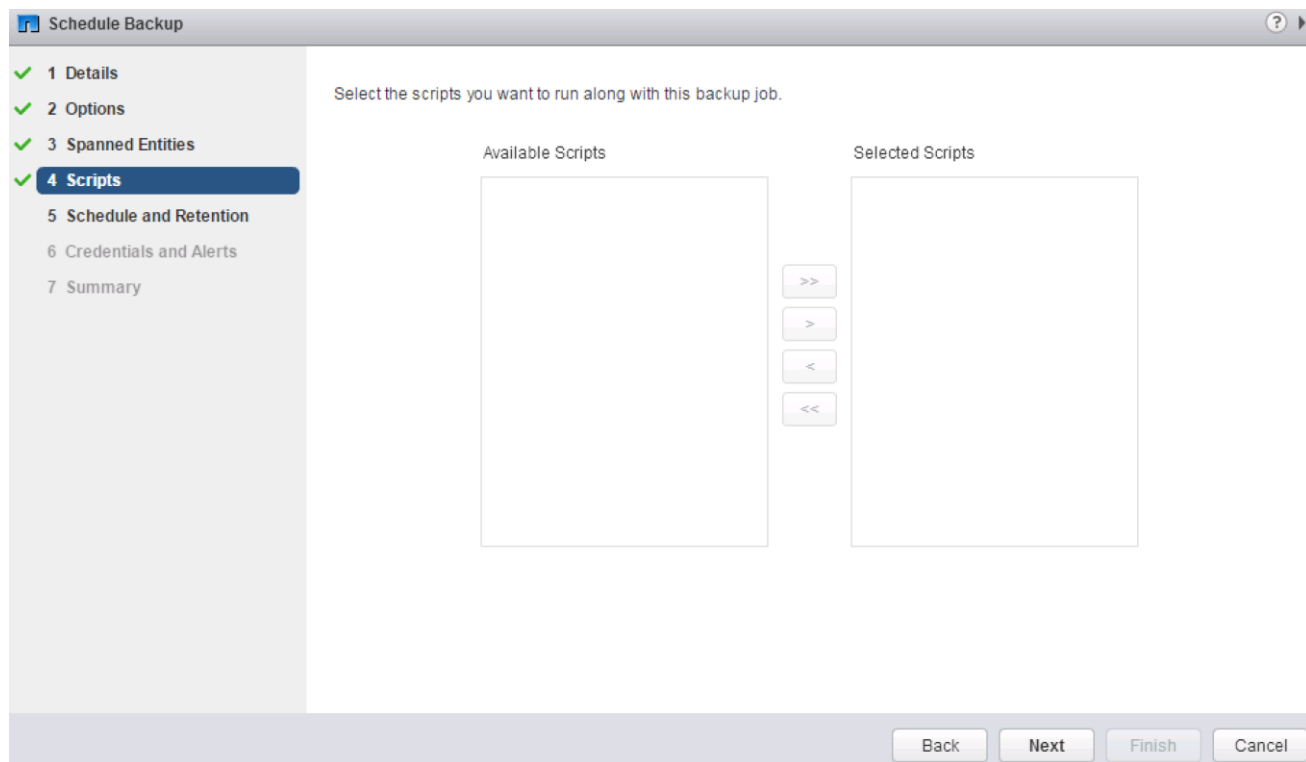
Select the options you want to include along with this backup job.

Initiate SnapVault update
 Initiate SnapMirror update
 Perform VMware consistency snapshot
 Include datastores with independent disks

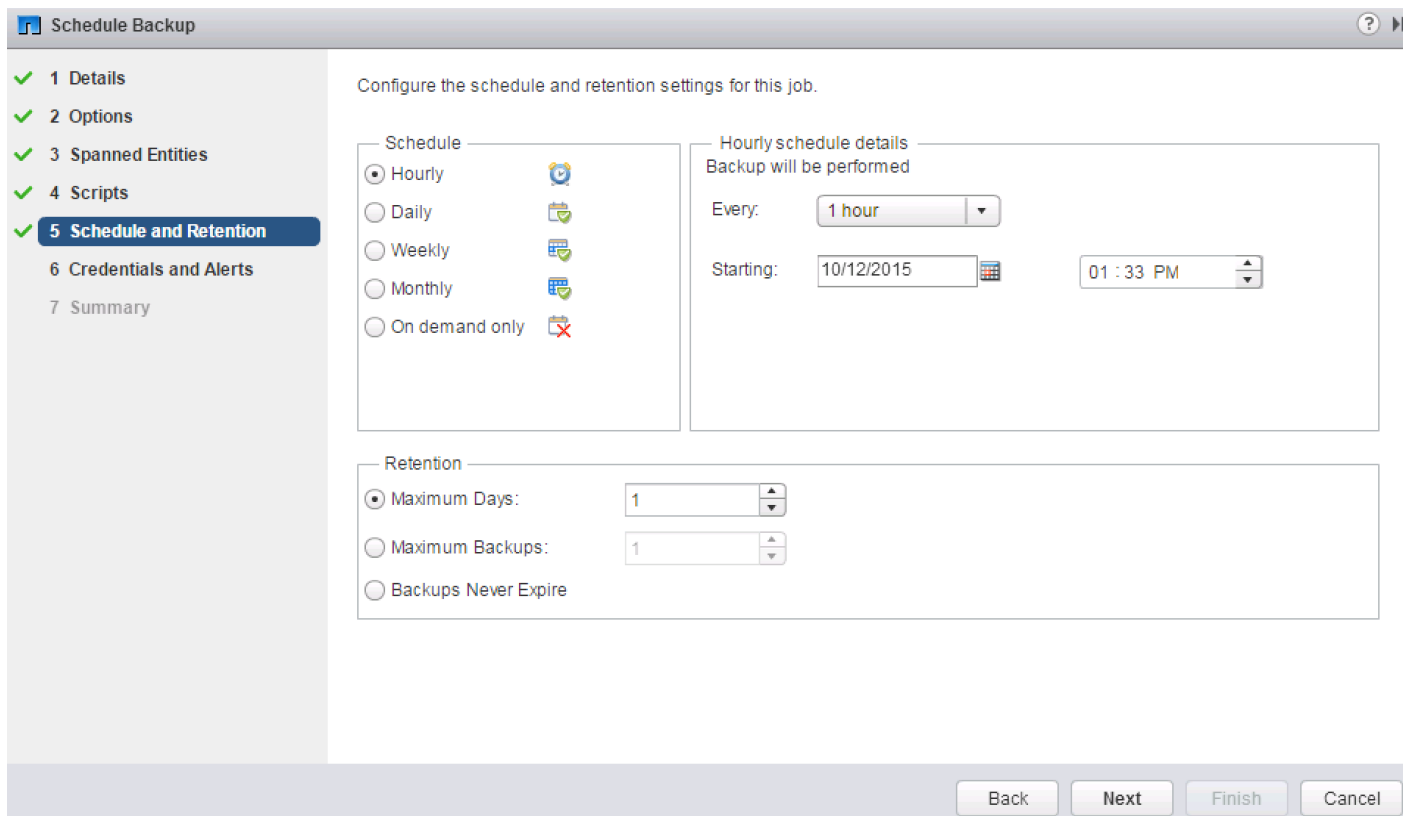
SnapVault integration in VSC is supported for Clustered Data ONTAP 8.2 or higher.

Back Next Finish Cancel

7. Click Next on the Options screen.
8. Click Next on the Spanned Entities screen.
9. Select one or more backup scripts if available and click Next in the Scripts screen.



10. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click Next.



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

12. Specify backup retention details as per requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate e-mail addresses. Click Next.

Schedule Backup

✓ 1 Details
 ✓ 2 Options
 ✓ 3 Spanned Entities
 ✓ 4 Scripts
 ✓ 5 Schedule and Retention
 ✓ 6 Credentials and Alerts
 ✓ 7 Summary


The Backup Job will be created with the following options:

Name: VSC_backup
 Description: VM backup
 Virtual entities to be backed up: infra_datastore_1
 Perform this backup: Every 1 hour starting 01:33 PM at 10/12/2015.
 Backup retention: Maximum of 1 day
 Email notification will be sent on: Always
 Email notification will be sent from: vsc_backup@example.com
 Email notification will be sent to: admin@example.com
 Email notification SMTP host: smtp.example.com
 Run Job Now

Back Next Finish Cancel

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

14. Click OK.



15. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by entering the following command:

```
volume modify -volume infra_datastore_1 -snapshot-policy none
```

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

```
volume snapshot show -volume infra_datastore_1
```

```
volume snapshot delete -volume infra_datastore_1 -vserver Infra-SVM -snapshot
<snapshot name>
```



The wildcard character, *, can be used in snapshot names in the previous command.

OnCommand Unified Manager 6.2P1

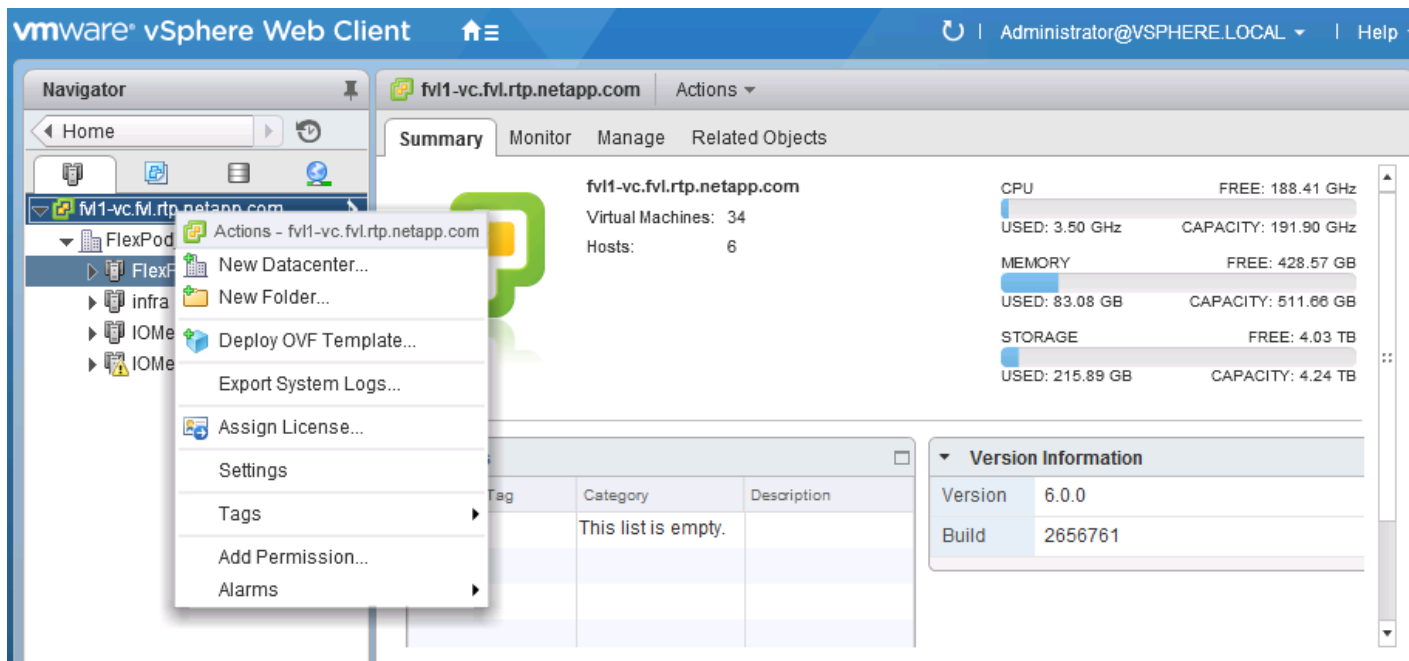
OnCommand Unified Manager OVF Deployment

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

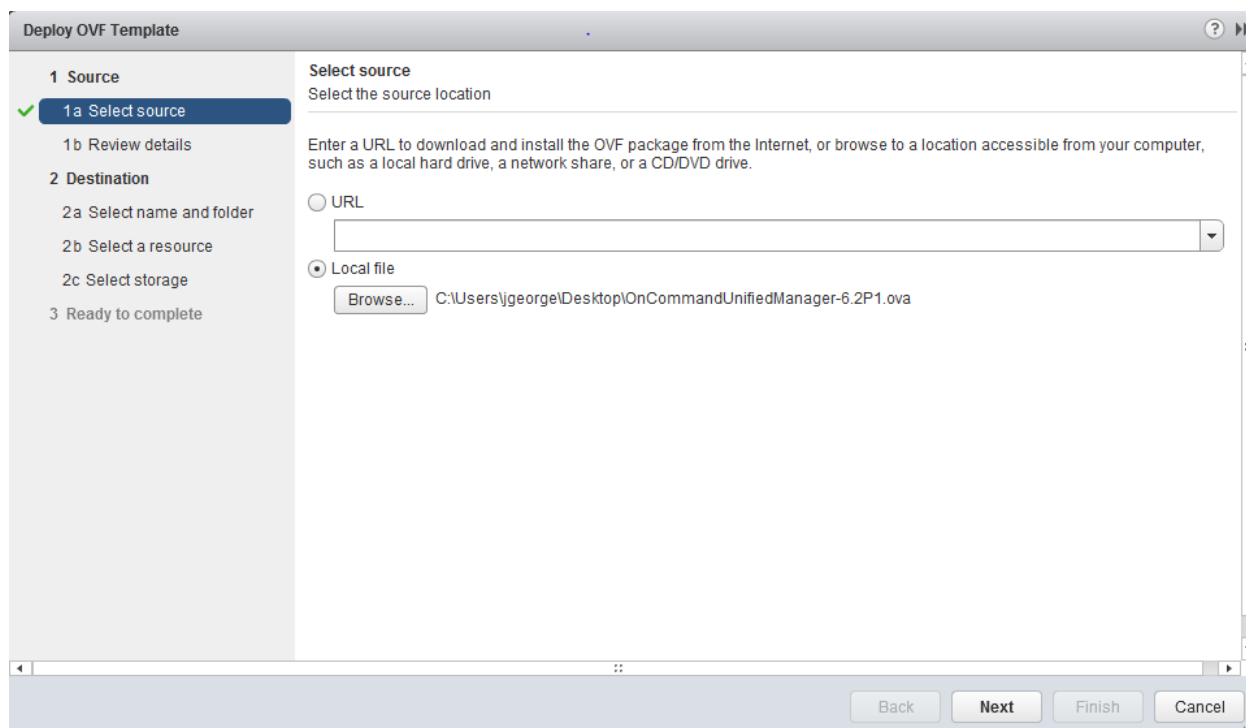


Download and review the [OnCommand Unified Manager Installation and Setup Guide](#).

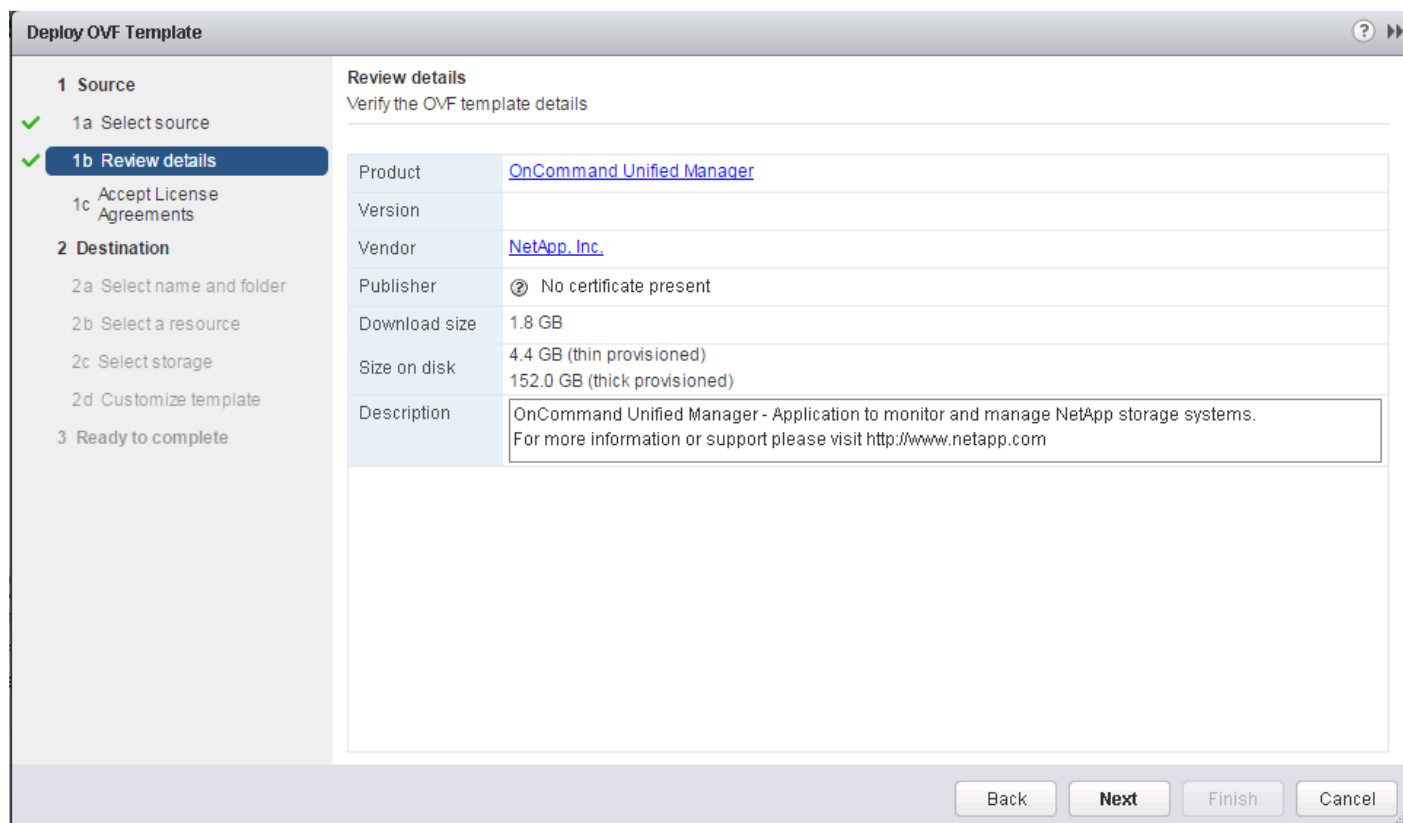
1. Download the OnCommand Unified Manager version 6.2P1 (OnCommandUnifiedManager-6.2P1.ova), from http://mysupport.netapp.com/NOW/download/software/oncommand_cdot/6.2P1/
2. Log in to the vSphere Web Client. Go to vCenter > VMs and Templates.
3. At the top of the center pane, click Actions > Deploy OVF Template.



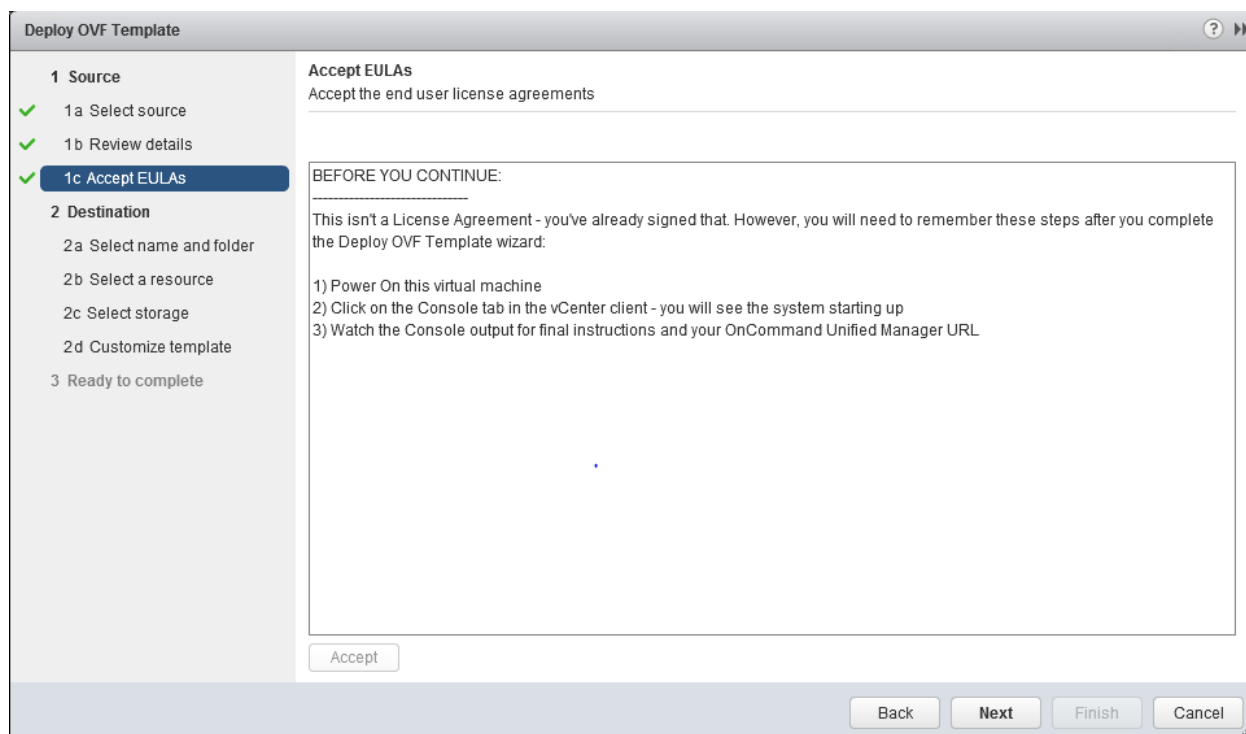
4. Browse the .ova file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.



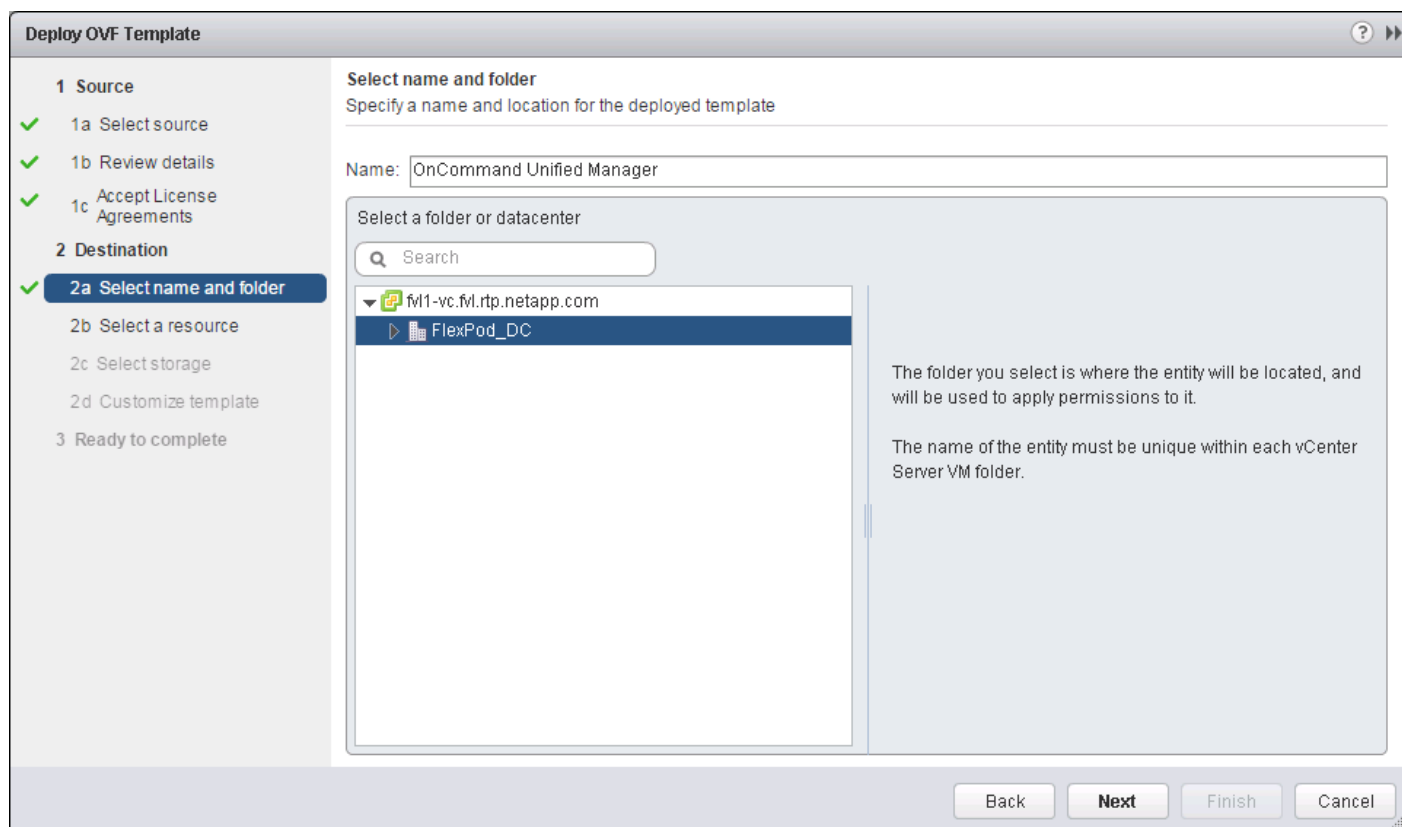
5. Click Next.



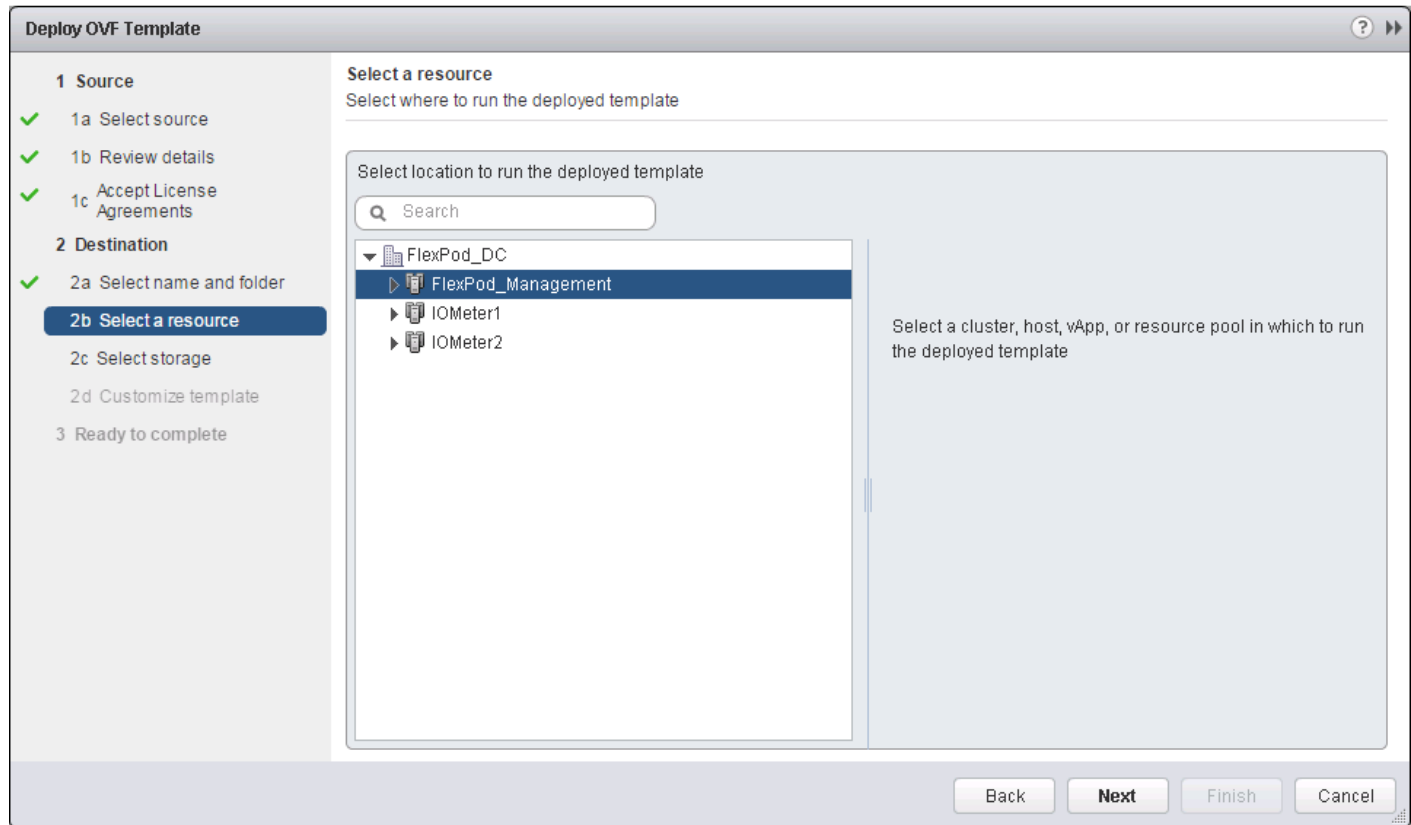
6. Read the EULA, then click the Accept button to accept the agreement. Click Next to continue.



7. Enter the name of the VM and select the FlexPod_DC folder to hold the VM. Click Next to continue.



8. Select FlexPod_Management within the FlexPod_DC datacenter as the destination compute resource pool to host the VM. Click Next to continue.



9. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the Virtual disk format. Click Next to continue.

Deploy OVF Template

1 Source

- ✓ 1 a Select source
- ✓ 1 b Review details
- ✓ 1 c Accept EULAs

2 Destination

- ✓ 2 a Select name and folder
- ✓ 2 b Select a resource
- ✓ **2 c Select storage**
- 2 d Setup networks
- 2 e Customize template

3 Ready to complete

Select storage
Select location to store the files for the deployed template

Select virtual disk format:

VM Storage Policy:

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
infra_datastore_1	500.00 GB	364.75 GB	403.30 GB	NFS	
infra_swap	100.00 GB	24.52 MB	99.98 GB	NFS	
datastore1 (1)	7.50 GB	857.00 MB	6.66 GB	VMFS	
datastore1 (4)	7.50 GB	857.00 MB	6.66 GB	VMFS	

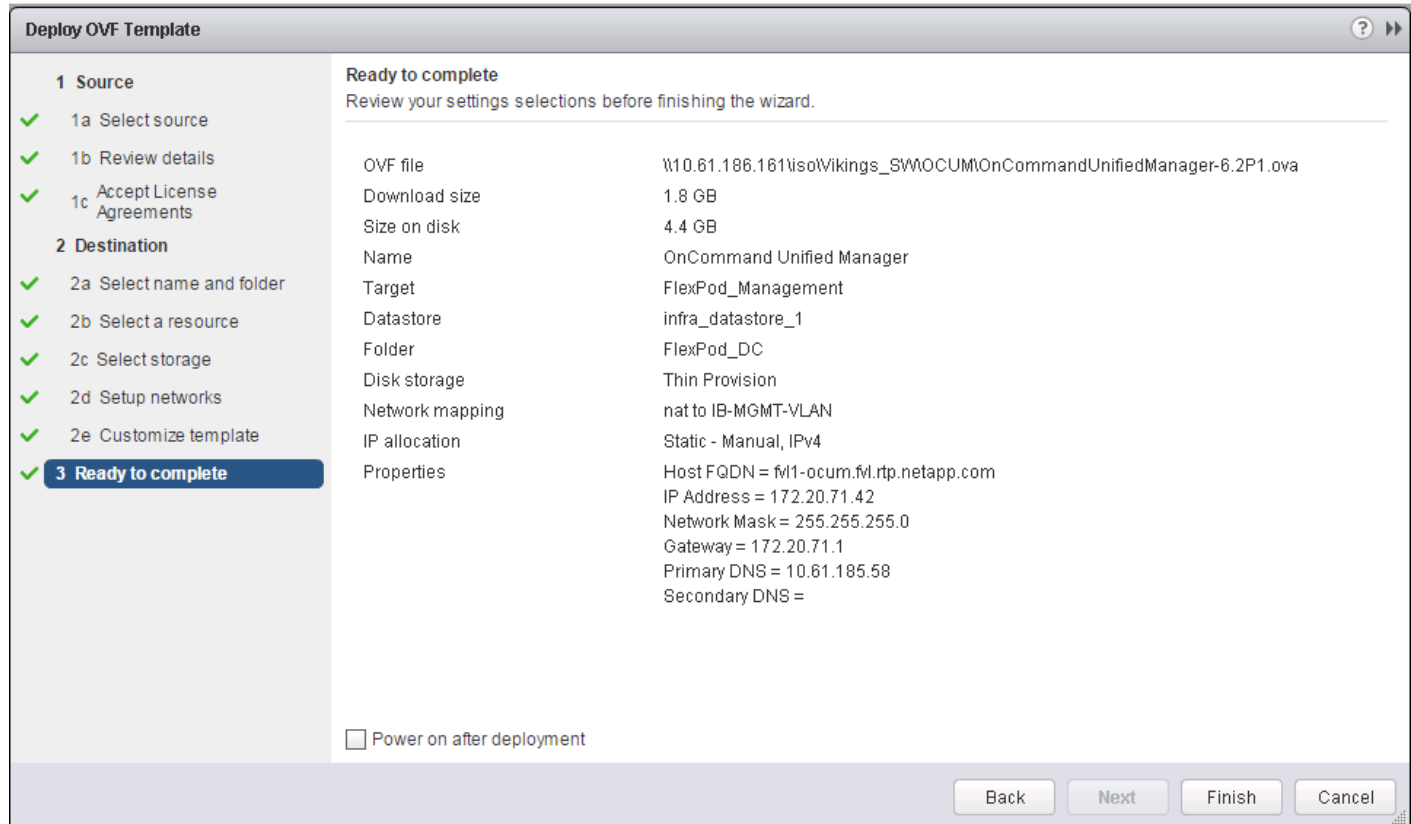
Back Next Finish Cancel

10. Select `IB-MGMT-VLAN` as the destination network to the nat source network. Click Next.

11. Fill out the details for the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS. Click Next to continue.

12. Clear the Power on after deployment checkbox.

13. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.



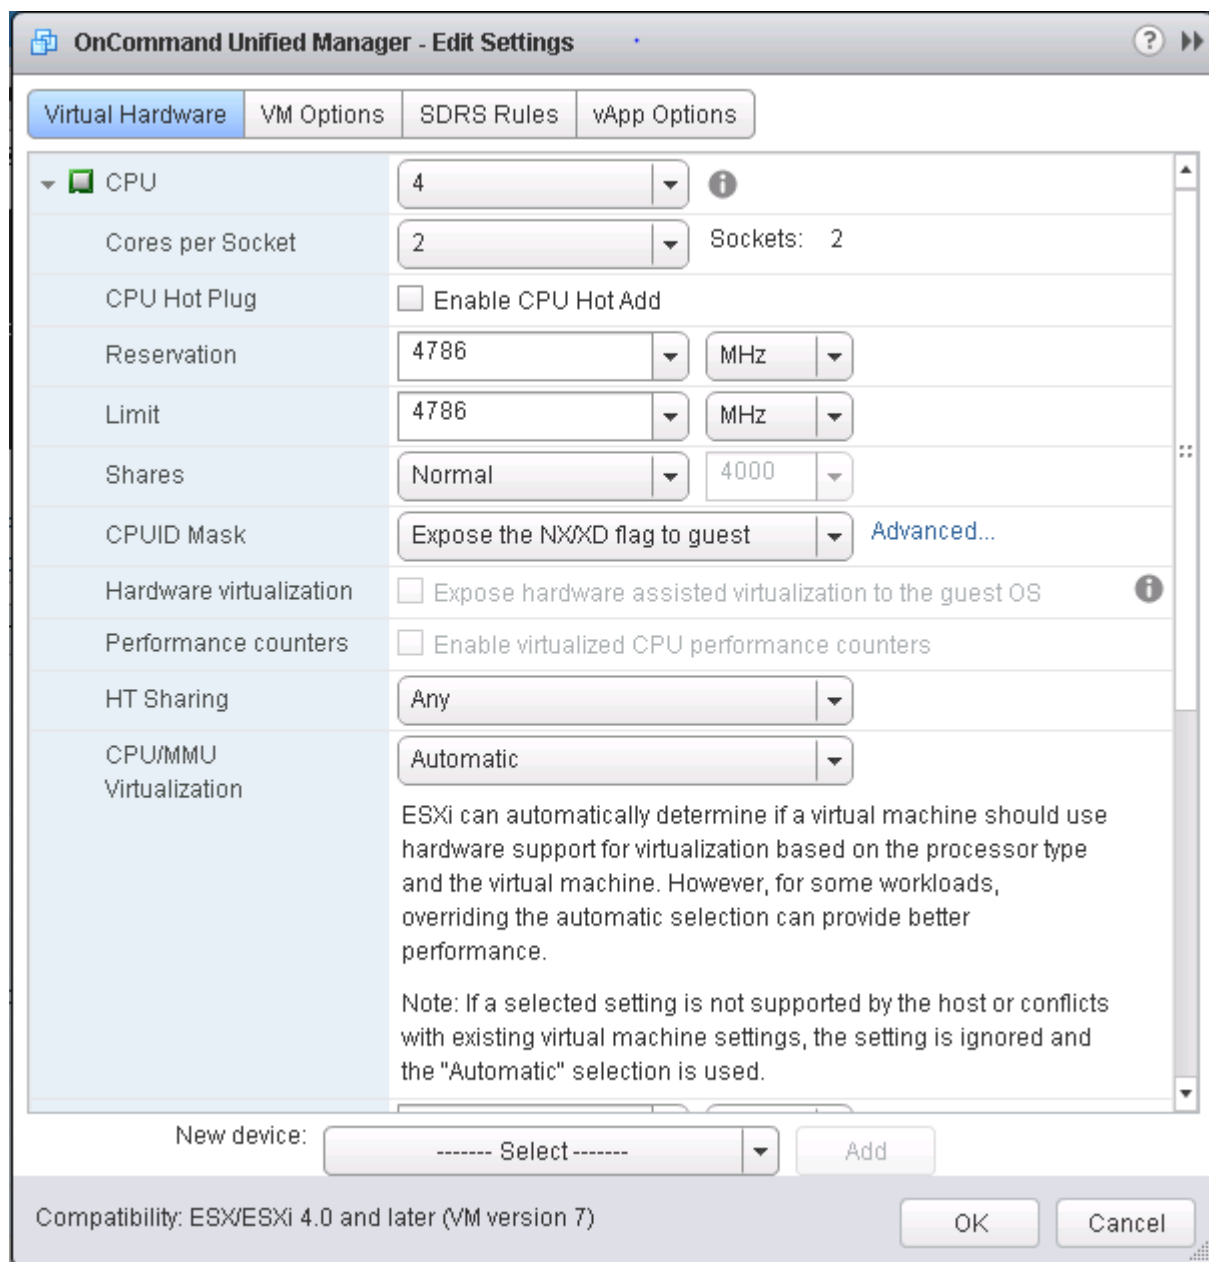
14. On the left pane, navigate to vCenter -> Virtual machines. After OVF deployment is complete, right-click the newly created virtual machine and select Edit Settings.

15. Click the CPU tab to expand the CPU options:

- a. The minimum required CPU Reservation is 4786 MHz. Determine the CPU frequency of the host.
- b. Set the number of CPUs to the number of CPUs required ($4786 / \text{CPU Frequency of host}$).
- c. Set the number of Cores per Socket where the Sockets number on the right matches the number of CPU sockets in the host. For example, if a host has 2 CPUs operating at a speed of 1999MHz, then the VM would need 4 virtual CPUs ($4786 / 1999 = 2.39$ - rounded to 4 virtual CPUs). If the host has 2 physical CPU sockets, 2 Cores per Socket, set the CPU Reservation and Limit to 4786 MHz
- d. The amount of memory can be set to 8 GB.



For detailed information, refer to the [OnCommand Unified Manager Installation and Setup Guide](#).



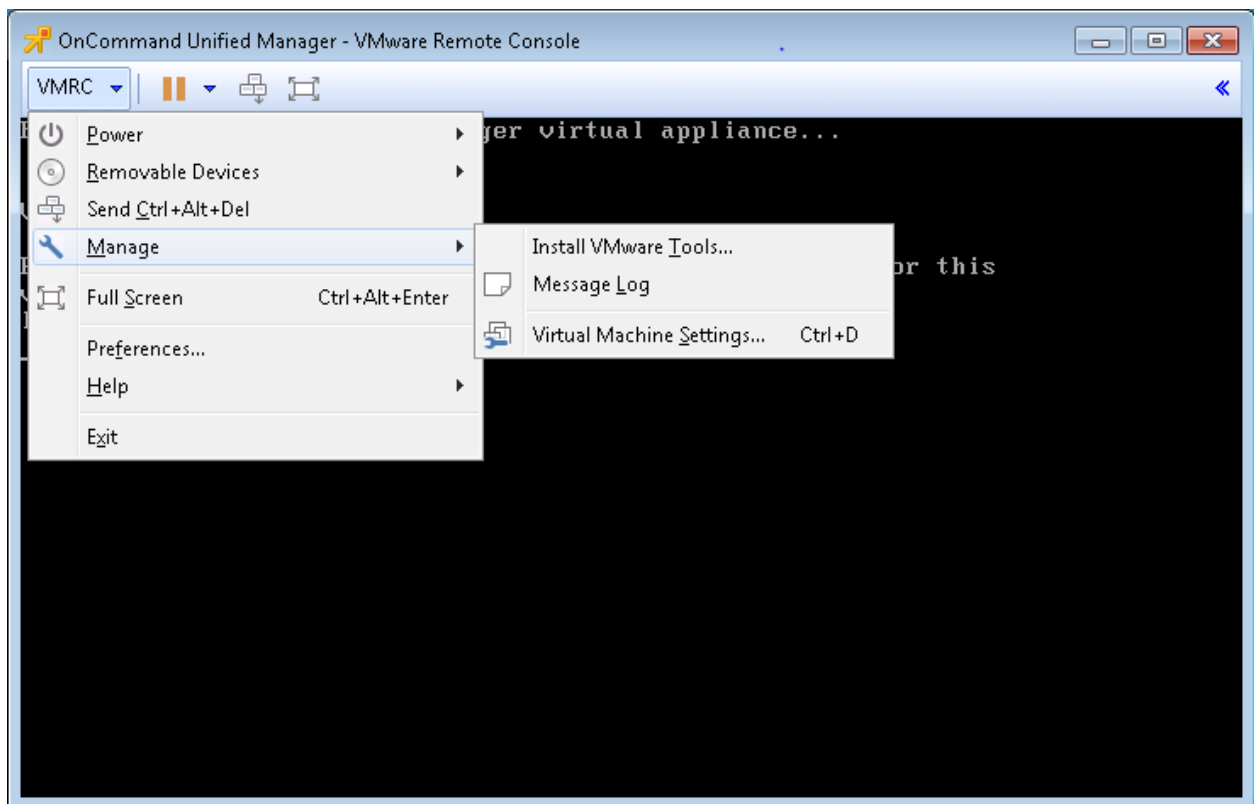
16. Click OK to accept the changes.

17. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Unified Manager Basic Setup

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

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.
2. In the VMRC window, select VMRC > Manage > Install VMware Tools. VMware Tools will install in the VM.



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

Geographic area: <<Enter your geographic location>>

Time zone: <<Select the city or region corresponding to your time zone>>

These commands complete the network configuration checks, generate SSL certificates for HTTPS and start the OnCommand Unified Manager services.

4. To Create a Maintenance User account, run the following commands:




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

```
Username : admin
```

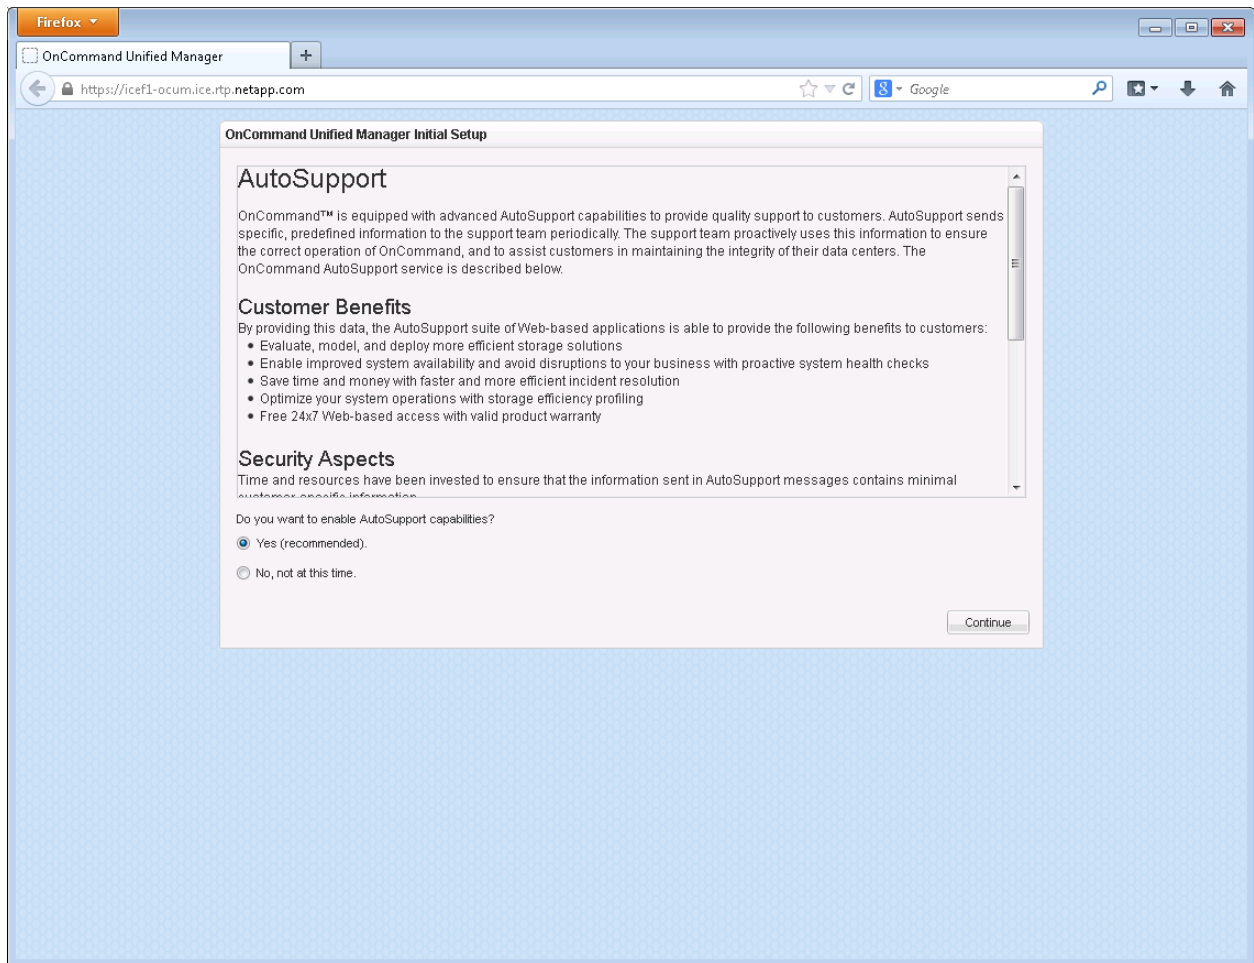
```
Enter new UNIX password: <<var_password>>
```

```
Retype new UNIX password: <<var_password>>
```

5. Using a web browser navigate to the OnCommand Unified Manager using URL:
https://<<var_oncommand_server_ip>>.



6. Log in using the Maintenance User account credentials.
7. Select `Yes` option to enable AutoSupport capabilities.



8. Click Continue.
9. Provide the NTP Server IP address <<var_switch_a_ntp_ip>>.
10. Provide the Maintenance User Email <<var_storage_admin_email>>.
11. Provide the SMTP Server Hostname.

OnCommand Unified Manager Initial Setup


NTP Server:

Maintenance User Email:

SMTP Server Hostname:
 [\(more options\)](#)

12. Click Save.

13. Click Add Cluster



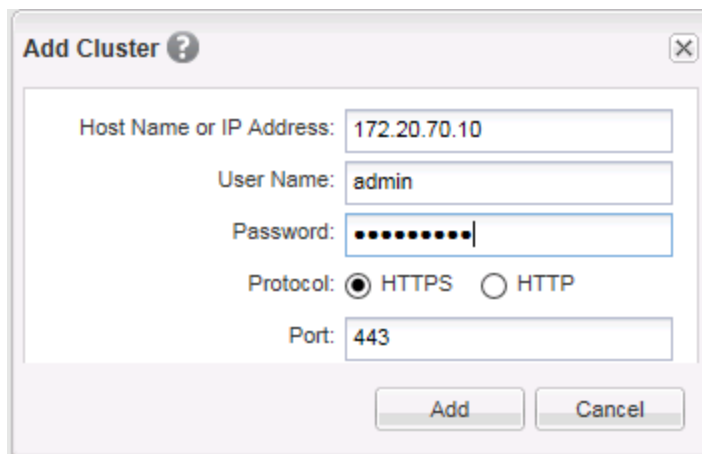
Get Started

Welcome to OnCommand Unified Manager

You can start using OnCommand Unified Manager by adding a cluster.

[➕ Add Cluster](#)

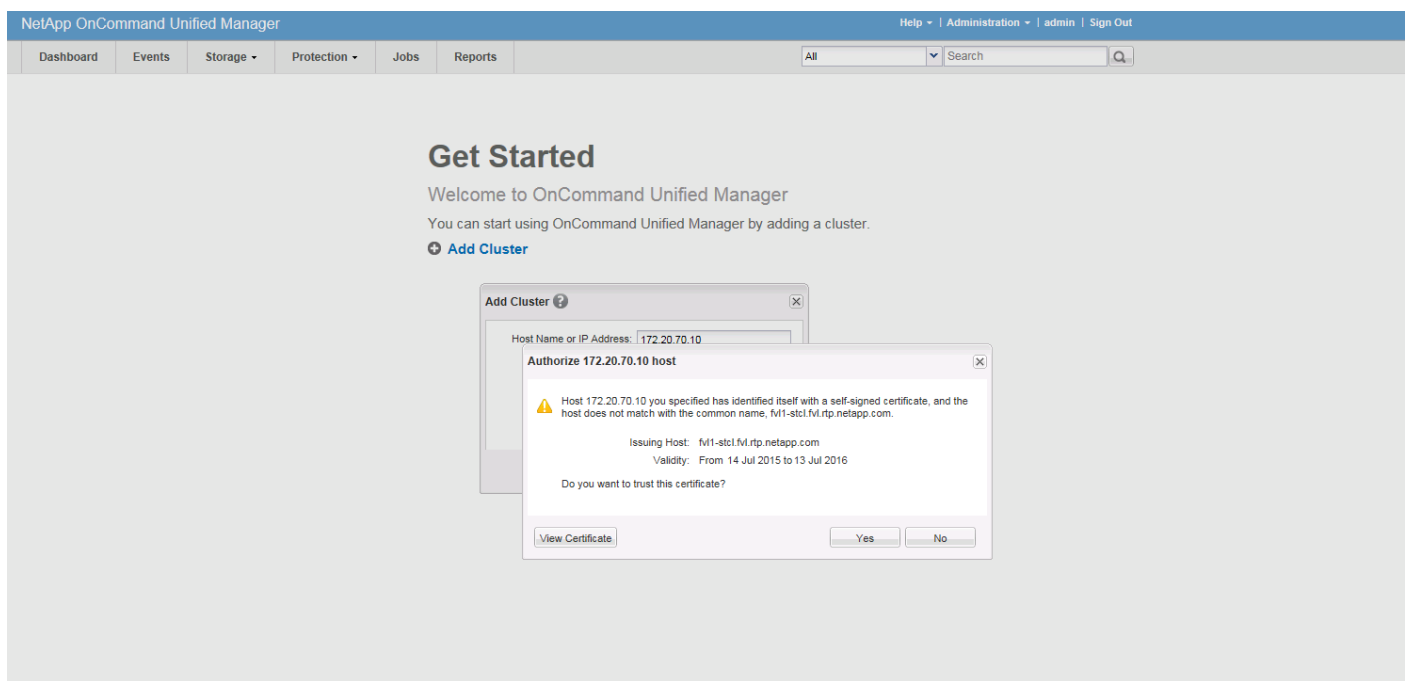
14. Provide the Cluster Management IP address, User Name, Password, Protocol, and Port.



The 'Add Cluster' dialog box contains the following fields and options:

- Host Name or IP Address: 172.20.70.10
- User Name: admin
- Password: [masked]
- Protocol: HTTPS HTTP
- Port: 443
- Buttons: Add, Cancel

15. Click Add.



The screenshot shows the NetApp OnCommand Unified Manager interface. The main content area displays a 'Get Started' message with a link to 'Add Cluster'. An 'Add Cluster' dialog box is open, and a secondary 'Authorize 172.20.70.10 host' dialog box is displayed over it. The authorization dialog contains the following information:

- Warning: Host 172.20.70.10 you specified has identified itself with a self-signed certificate, and the host does not match with the common name, fv1-stcl.fv1.rtp.netapp.com.
- Issuing Host: fv1-stcl.fv1.rtp.netapp.com
- Validity: From 14 Jul 2015 to 13 Jul 2016
- Question: Do you want to trust this certificate?
- Buttons: View Certificate, Yes, No

16. Click Yes to trust the certificate from the controller.



The Cluster Add operation might take a couple of minutes.

17. After the cluster is added it can be accessed by clicking on the Storage tab and selecting Clusters.

NetApp OnCommand Unified Manager

Dashboard Events Storage Protection Jobs Reports All

Filters

Status [Clear](#)

Critical

Error

Warning

Normal


Communication Status [Clear](#)

Good

Not Reachable

Clusters ?

+ Add Edit Remove View Monitoring Status Refresh List

<input type="checkbox"/>	Cluster	Communication...	Host Name or IP Address	OS Version
<input type="checkbox"/>	 fv11-stcl	Good	172.20.70.10	8.3.1

OnCommand Performance Manager 2.0

OnCommand Performance Manager OVF Deployment

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




Download and review the [OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances](#).

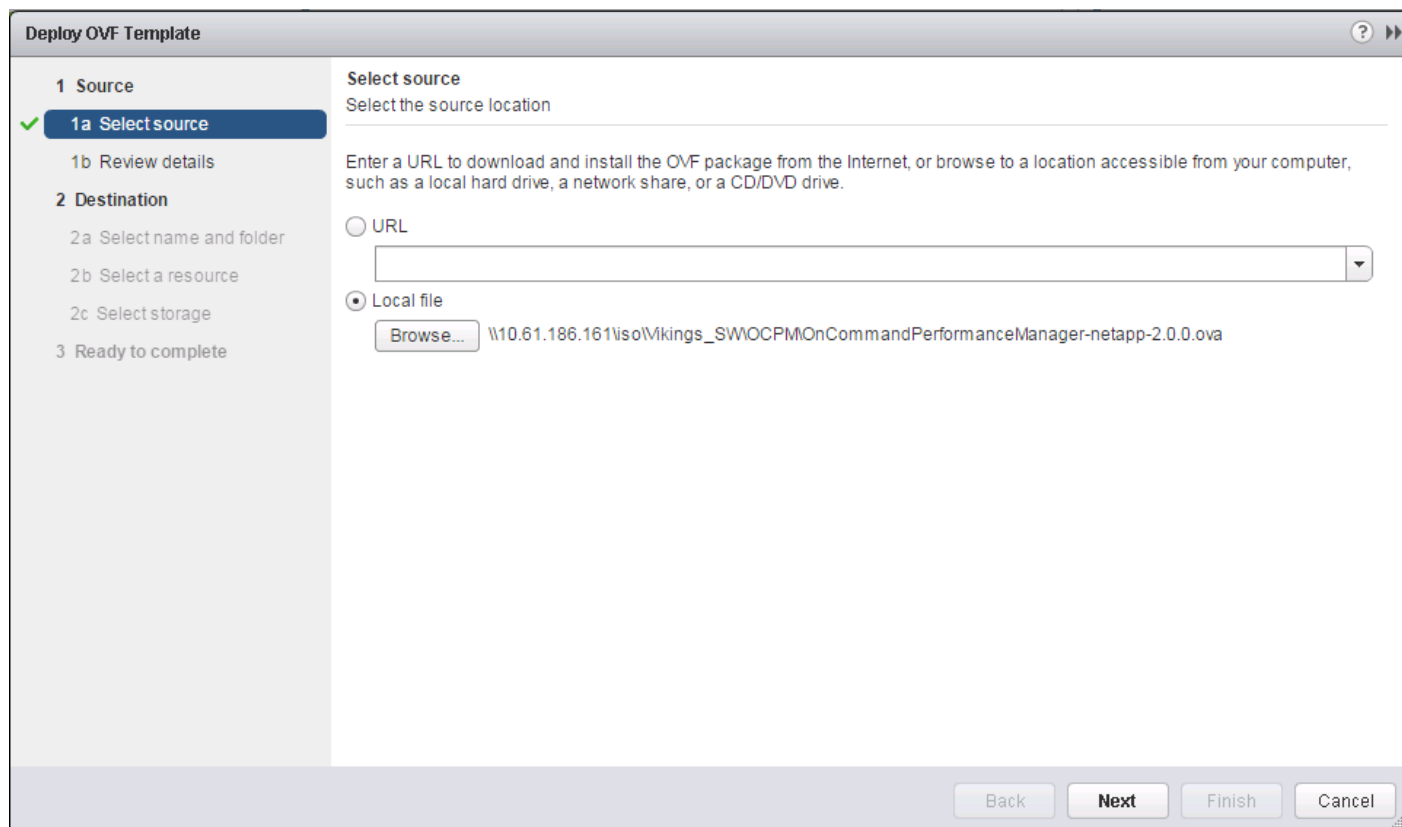


Download the OnCommand Performance Manager version 2.0 (OnCommandPerformanceManager-netapp-2.0.0.ova), from http://mysupport.netapp.com/NOW/download/software/oncommand_pm/2.0/.

1. Log in to the vSphere Web Client. Go to Home > VMs and Templates.
2. At the top of the center pane, click Actions > Deploy OVF Template.

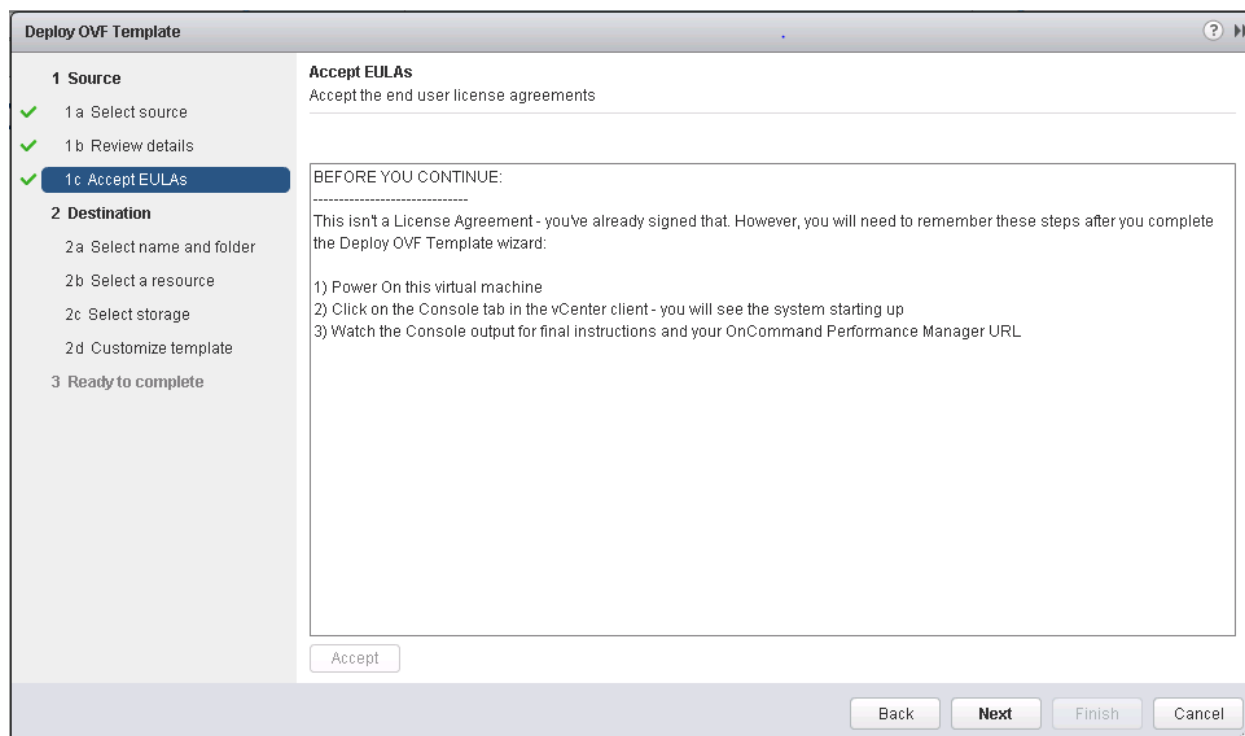


3. Browse to the `OnCommandPerformanceManager-netapp-2.0.0.ova` file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.

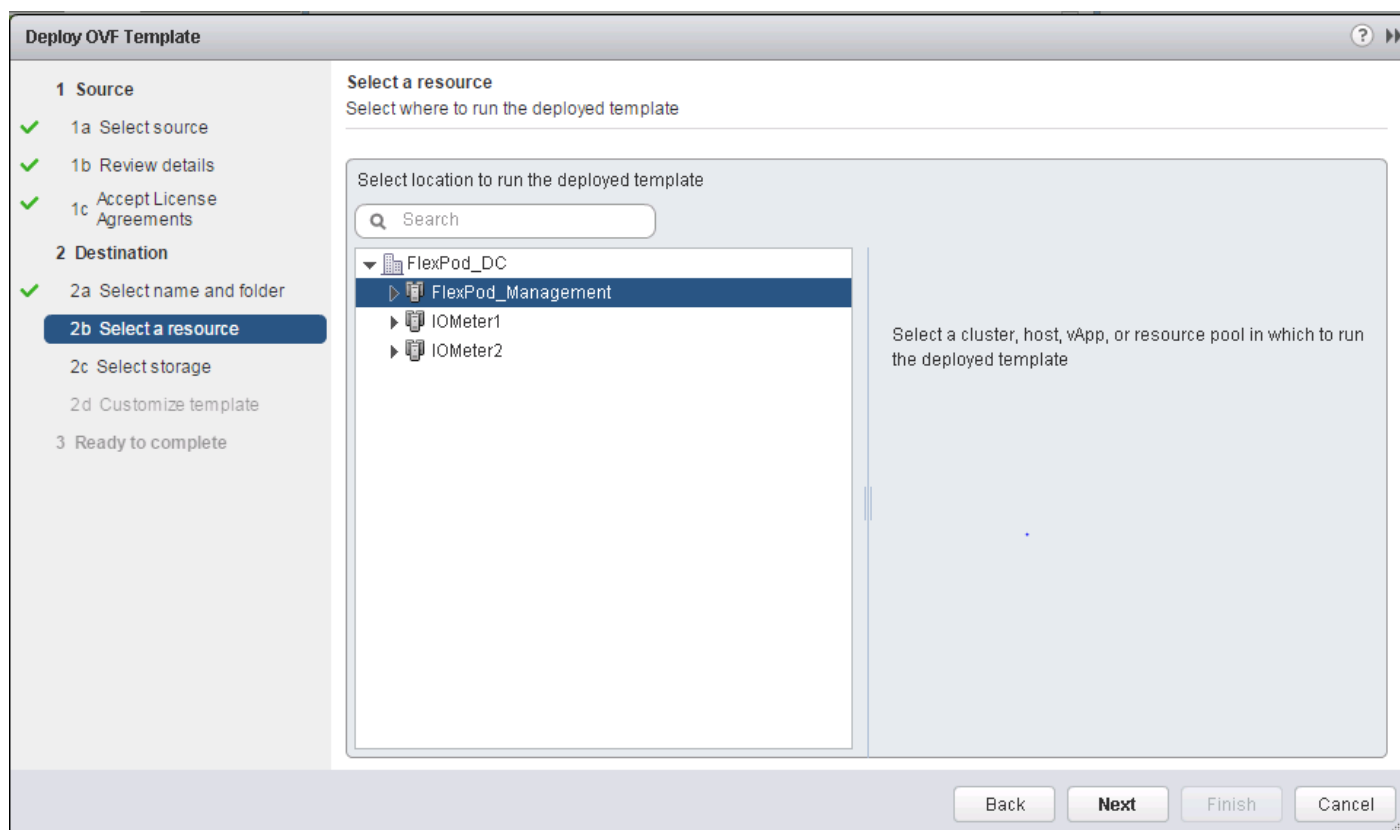


4. Review the details and click Next.


5. Read the EULA, and then click the Accept button to accept the agreement. Click Next to continue.



6. Enter the name of the VM and select the FlexPod_DC folder to hold the VM. Click Next to continue.
7. Select FlexPod_Management within the FlexPod_DC datacenter as the destination compute resource pool to host the VM. Click Next to continue.



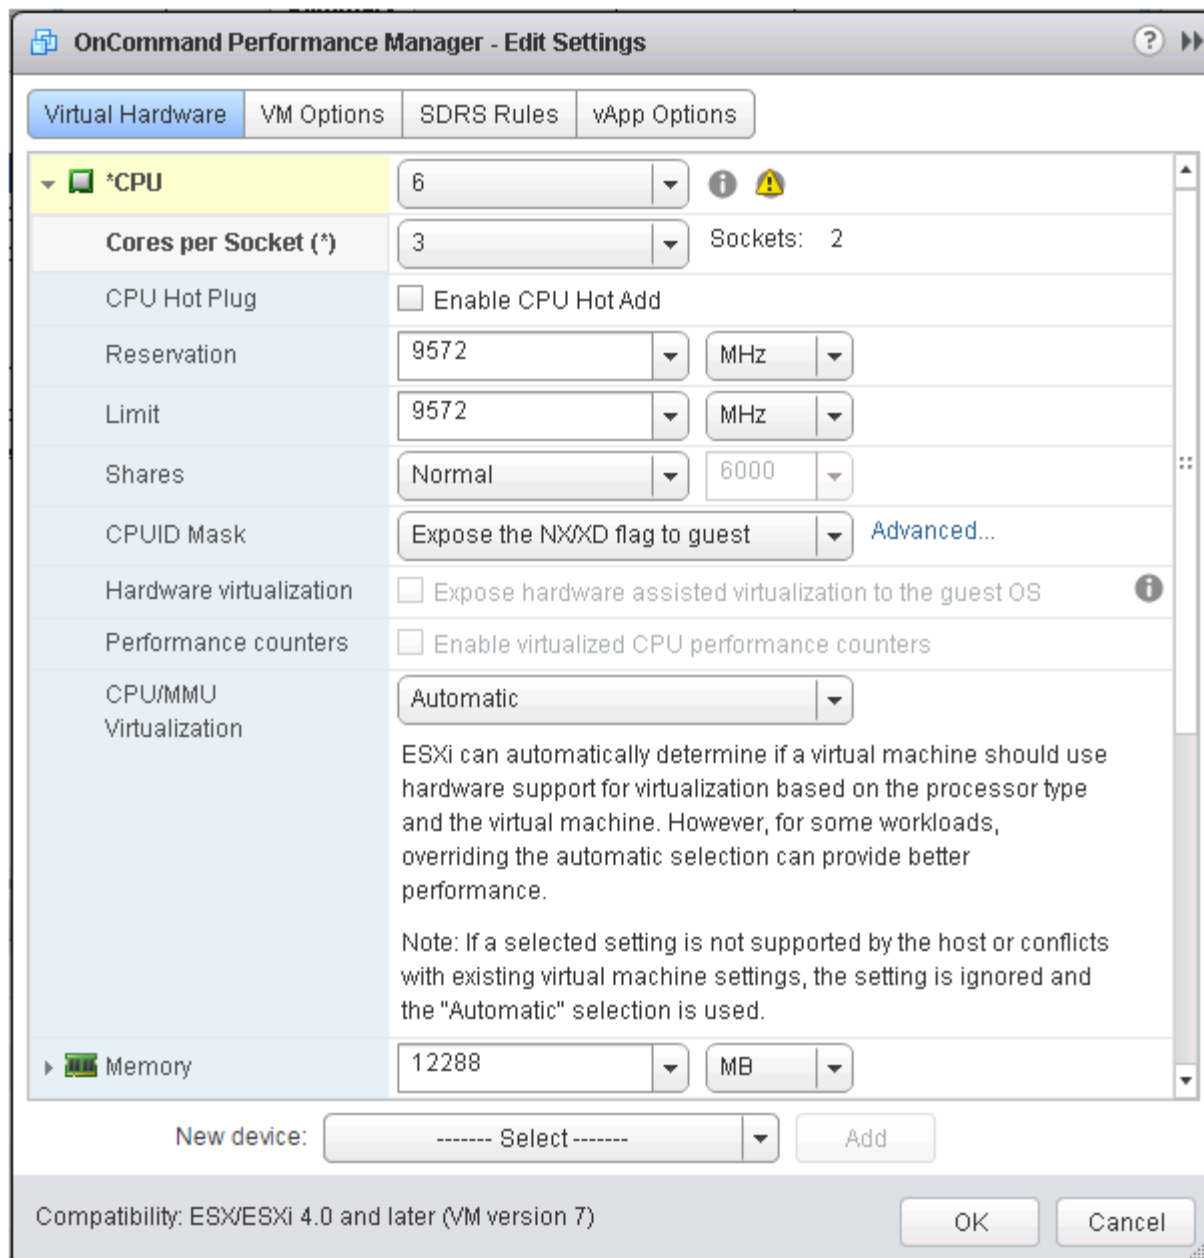
8. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the Virtual disk format. Click Next to continue.



9. Select `IB-MGMT-VLAN` as the destination network to the nat source network. Click Next.
10. Fill out the details for the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS. Click Next to continue.
11. Clear the Power on after deployment checkbox.
12. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.
13. On the left pane, navigate to Home -> Hosts and Clusters. Expand the FlexPod_Management cluster and select the newly create OnCommand Performance Manager VM. After OVF deployment is complete, right-click the newly created virtual machine and select Edit Settings.
14. Expand the CPU options.
 - a. The minimum required CPU Reservation is 9572 MHz. Determine the CPU frequency of the host.
 - b. Set the number of CPUs to the number of CPUs required ($9572 / \text{CPU Frequency of host}$).
 - c. Set the number of Cores per Socket where the Sockets number on the right matches the number of CPU sockets in the host. For example, if a host has 2 CPUs operating at a speed of 1999MHz, then the VM would need 6 virtual CPUs ($9572 / 1999 = 4.79$ - rounded to 6 virtual CPUs). If the host has 2 physical CPU sockets, 3 Cores per Socket.



For detailed information, refer to the [OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances](#).



15. Click OK to accept the changes.


16. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Performance Manager Basic Setup

To setup the OnCommand Performance Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.

2. In the VMware Remote Console window, select VMRC > Manage > Install VMware Tools. VMware Tools will install in the VM.



3. Set up OnCommand Performance Manager by answering the following questions in the console window:

Geographic area: <<Enter your geographic location>>

Time zone: <<Select the city or region corresponding to your time zone>>

These commands complete the network configuration checks, generates SSL certificates and starts the OnCommand Performance Manager services.

1. To Create a Maintenance User account, run the following commands:



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

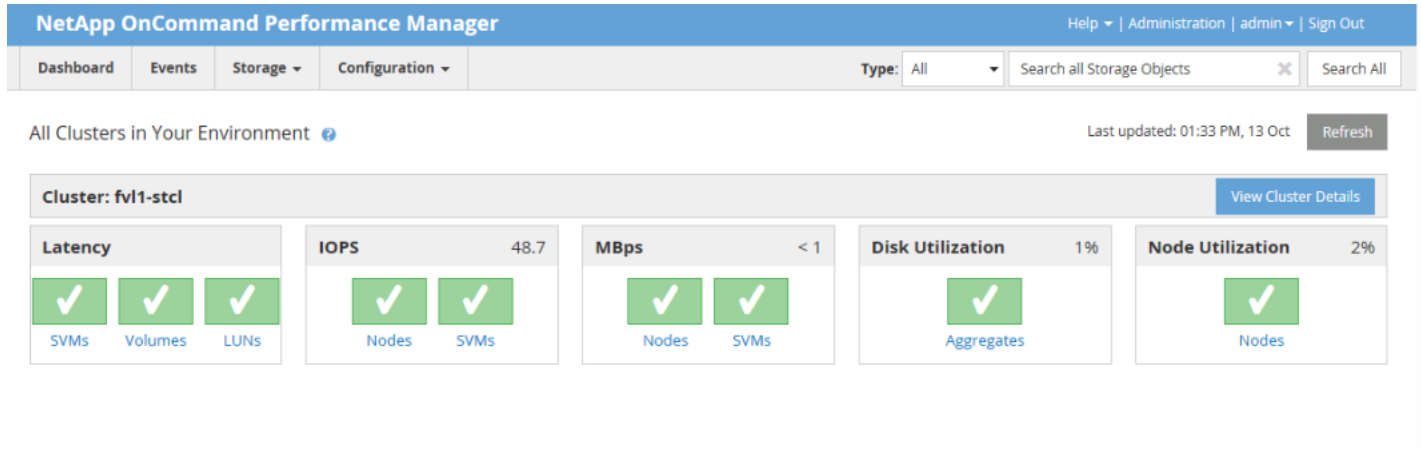
```
Username : admin
```

```
Enter new UNIX password: <<var_password>>
```

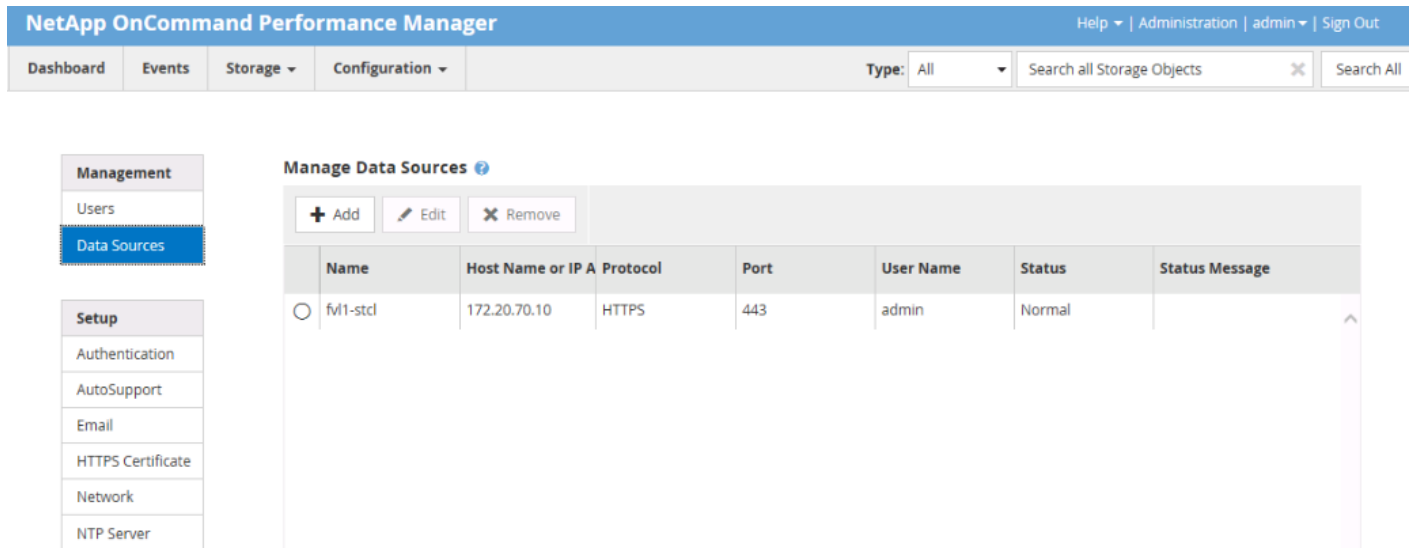
```
Retype new UNIX password: <<var_password>>
```

2. Using a web browser navigate to the OnCommand Performance Manager using URL: `https://<<var_oncommand_pm_ip>>`.
3. Log in using the Maintenance User account (admin) credentials.

4. Enter a Maintenance User Email Address, SMTP Mail Server information, and the NTP server IP address. Click Save and go to next step.
5. Select Yes option to enable AutoSupport capabilities. Click Save and go to next step.
6. Click Save and go to next step to not change the admin password.
7. Enter the storage cluster host name or IP address, the storage cluster admin user name, and the storage cluster admin password. Click Add Cluster, then click Save and Complete Configuration. It may take up to 15 minutes for the cluster to be visible in OnCommand Performance Manager.



8. After the cluster is added it can be accessed by clicking on Administration > Manage Data Sources.



Link OnCommand Performance Manager to OnCommand Unified Manager

To link OnCommand Performance Manager to the OnCommand Unified Manager, complete the following steps:

1. Using a web browser navigate to the **OnCommand Unified Manager** using URL: `https://<<var_oncommand_server_ip>>`. Log in with the Maintenance user id and password setup earlier.

2. In the OnCommand Unified Manager web interface, select Administration > Manage Users to set up an Event Publication user.
3. Click Add to add a user.
4. Leave the Type set to Local User. Use eventpub as the Name and enter and confirm a password. Enter an email address for this user and set the Role to Event Publisher. Click Add.

The screenshot shows a dialog box titled "Add User" with a close button (X) in the top right corner. Inside the dialog, there is a warning icon (yellow triangle with an exclamation mark) and a message: "Authentication server is either disabled or not configured. To add a remote user or group, enable or configure the authentication server from Setup Options." Below the message are several input fields: "Type" is a dropdown menu set to "Local User"; "Name" is a text box containing "eventpub"; "Password" and "Confirm Password" are text boxes filled with dots; "Email" is a text box containing "eventpub@netapp.com"; and "Role" is a dropdown menu set to "Event Publisher". At the bottom of the dialog are two buttons: "Add" and "Cancel".

5. At the OnCommand Performance Manager console window, log into the Command Line Interface with the Maintenance User (admin) defined earlier.
6. Enter 5 to select Unified Manager Connection.

```

OnCommand Performance Manager - VMware Remote Console
VMRC | [Icons]
For regular system operation and usage, use the UI.
ocpm login: admin
Password:
Linux OnCommand 3.2.0-4-amd64 #1 SMP Debian 3.2.68-1+deb7u1 x86_64

OnCommand Performance Manager Maintenance Console

Version      : 2.0.0
System ID    : 8cde1672-8af5-446b-8e08-45deef876778
Status       : Running

Main Menu
-----
1 ) Upgrade
2 ) Network Configuration
3 ) System Configuration
4 ) Support/Diagnostics
5 ) Unified Manager Connection
6 ) External Data Provider
7 ) Backup/Restore

x ) Exit

Enter your choice: 5_

```

7. Enter 2 to Add / Modify Unified Manager Server Connection.
8. Enter y to continue.
9. Enter the OnCommand Unified Manager FQDN or IP Address.
10. Hit Enter to accept the default port 443.
11. Enter y to accept the Unified Manager Security Certificate.
12. Enter eventpub for the Event Publisher User Name.
13. Enter the eventpub password.
14. Enter y to accept the entered settings.
15. Press any key to continue.
16. Exit the OnCommand Performance Manager console. OnCommand Performance Manager events will now appear in the OnCommand Unified Manager Dashboard.

NetApp NFS Plug-In 1.1.0 for VMware VAAI

Enable VMware vStorage for NFS in Clustered Data ONTAP

To enable VMware vStorage for NFS in 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-SVM -vstorage enabled
```

3. Verify that the export policy rules are set up correctly.

```
vserver export-policy rule show -vserver Infra-SVM
```

Sample output:

```
NetApp::> vserver export-policy rule show -vserver Infra-SVM
```

	Policy	Rule	Access	Client	RO
Vserver	Name	Index	Protocol	Match	Rule
Infra-SVM	default	1	nfs	192.168.170.61	sys
Infra-SVM	default	2	nfs	192.168.170.60	sys
Infra-SVM	default	3	nfs	192.168.170.58	sys
Infra-SVM	default	4	nfs	192.168.170.59	sys
Infra-SVM	default	5	nfs	192.168.170.62	sys
Infra-SVM	default	6	nfs	192.168.170.63	sys

6 entries were displayed.

4. The access protocol for the FlexPod policy name should be NFS. If the access protocol is not “nfs” for a given rule index, run the following command to set NFS as the access protocol:

```
vserver export-policy rule modify -vserver Infra-SVM -policyname default -  
ruleindex <<var_rule_index>> -protocol nfs
```

Install NetApp NFS Plug-In for VMware VAAI

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

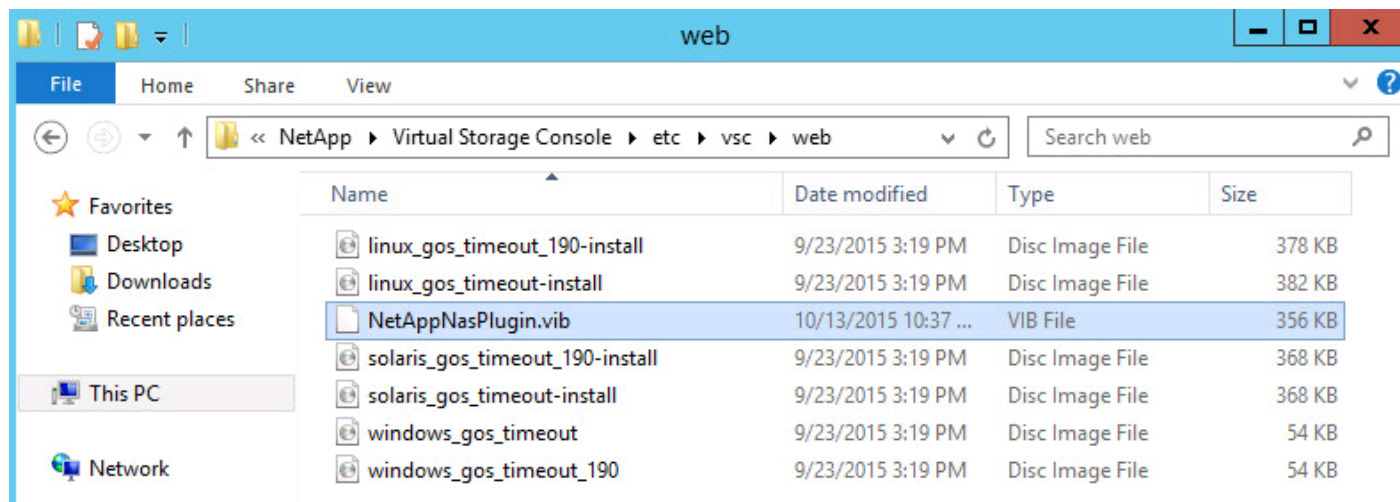
1. From a console interface on the NetApp VSC VM, go to the [Software Downloads](#) page in the [NetApp Support](#) site.
2. Scroll down to locate the NetApp NFS Plug-in for VMware VAAI, select the ESXi6.0 platform, and click Go.
3. Click View & Download.
4. Click CONTINUE.
5. Click Accept.

- Download the .vib file of the most recent plug-in version to the VSC VM Desktop as NetAppNasPlugin.vib.

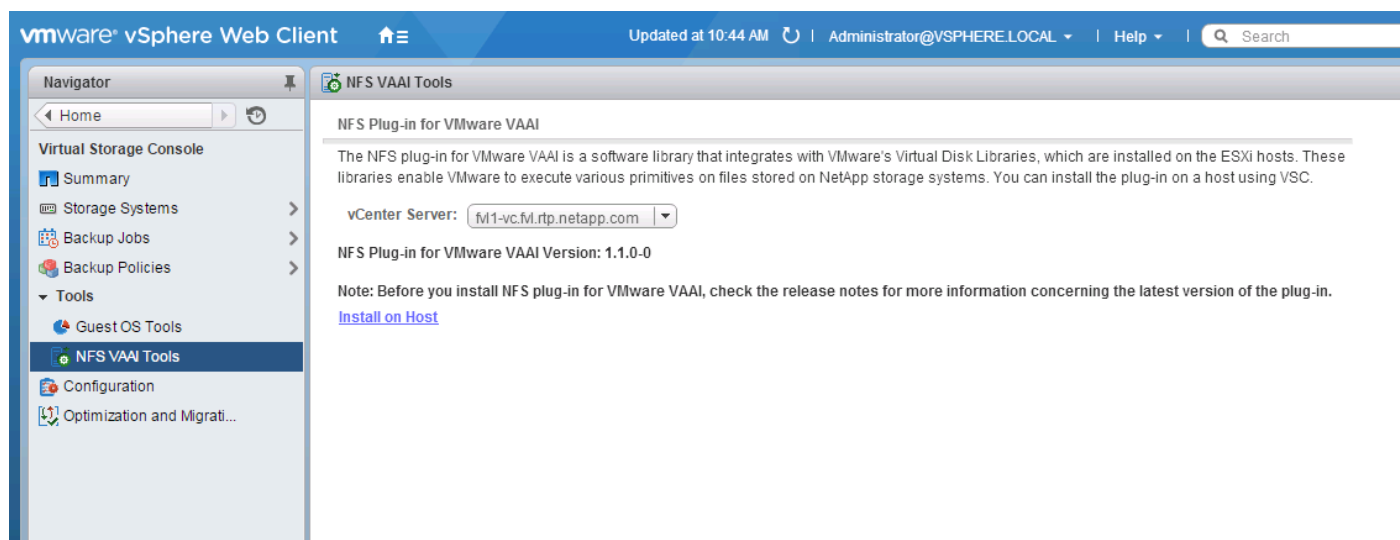


It is important that the file be saved as NetAppNasPlugin.vib.

- On the VSC VM Desktop, move the NetAppNasPlugin.vib file to the C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web folder.



- Go to the VMware vSphere Web Client and select VSC. Click on NFS VAAI Tools. Make sure NFS Plug-in for VMware VAAI Version: 1.1.0-0 is shown.



- Click Install on Host. Select all Hosts on which you want to install the plug-in.

NFS Plug-in for VMware VAAI

Select the hosts on which you want to install the NFS Plug-in for VMware VAAI. Incompatible ESX/ESXi hosts are not selectable.

<input checked="" type="checkbox"/>	fm1-h1.fvl.rtp.netapp.com
<input checked="" type="checkbox"/>	fm1-h2.fvl.rtp.netapp.com

10. Click Install then click OK.

11. One at a time, put each ESXi host into Maintenance Mode, reboot the host, then Exit Maintenance Mode. It may be necessary to manually migrate VMs to the other host to allow the host to enter Maintenance Mode.

12. When the reboots have completed, in the vSphere Web Client from the Home page, click Storage, then select the infra_datastore_1 datastore. Select Settings under the Manage tab in the center pane. Hardware Acceleration should now show Supported on all hosts as shown below. All NFS datastores should now support Hardware Acceleration.

The screenshot shows the VMware vSphere Web Client interface. The left sidebar displays a tree view of the storage infrastructure, with 'infra_datastore_1' selected. The main pane shows the 'Manage' tab for 'infra_datastore_1', with the 'Settings' sub-tab active. The 'Datastore Capabilities' section is expanded, showing 'Storage I/O Control' as 'Disabled' and 'Hardware Acceleration' as 'Supported on all hosts'.

Properties	
Name	infra_datastore_1
Type	NFS 3
Maximum file size	15.97 TB
Maximum virtual disk size	15.81 TB

Capacity	
Capacity	432.41 GB free out of 500.00 GB

Datastore Capabilities	
Storage I/O Control	Disabled
Hardware Acceleration	Supported on all hosts

About the Authors

Haseeb Niazi, Technical Marketing Engineer, Cisco UCS Data Center Solutions Engineering, Cisco Systems, Inc.

Haseeb Niazi has over 16 years of experience at Cisco focused on Data Center, Security, WAN Optimization, and related technologies. As a member of various solution teams and advanced services, Haseeb has helped many enterprise and service provider customers evaluate and deploy a wide range of Cisco solutions. Haseeb holds a master's degree in Computer Engineering from the University of Southern California.

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

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Masters of Science in Information Security from East Carolina University.

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

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.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Chris O' Brien, Cisco Systems, Inc.
- Melissa Palmer, NetApp