

FlexPod Datacenter with VMware vSphere 6.5, NetApp AFF A-Series and IP-Based Storage

Deployment Guide for FlexPod Datacenter with IP-Based Storage, VMware vSphere 6.5, and ONTAP 9.1

Last Updated: June 29, 2017



About Cisco Validated Designs (CVDs)

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 Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, 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)

© 2017 Cisco Systems, Inc. All rights reserved.

Table of Contents

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.....	11
Deployment Hardware and Software	13
Software Revisions	13
Configuration Guidelines.....	13
Physical Infrastructure.....	14
FlexPod Cabling	14
Network Switch Configuration.....	18
Physical Connectivity	18
FlexPod Cisco Nexus Base	18
Set Up Initial Configuration	18
FlexPod Cisco Nexus Switch Configuration.....	20
Enable Licenses.....	20
Set Global Configurations	21
Create VLANs.....	21
Add NTP Distribution Interface.....	22
Add Individual Port Descriptions for Troubleshooting.....	22
Create Port Channels.....	23
Configure Port Channel Parameters.....	24
Configure Virtual Port Channels	25
Uplink into Existing Network Infrastructure	27
Storage Configuration	28
AFF A300 Controllers	28
NetApp Hardware Universe	28
Controllers.....	28
Disk Shelves	28

Clustered Data ONTAP 9.1	29
Complete Configuration Worksheet	29
Configure ONTAP Nodes	29
Log In to the Cluster	38
Zero All Spare Disks	38
Set Onboard Unified Target Adapter 2 Port Personality	38
Set Auto-Revert on Cluster Management	39
Set Up Management Broadcast Domain	39
Set Up Service Processor Network Interface	39
Create Aggregates	40
Verify Storage Failover.....	41
Disable Flow Control on 10GbE and 40GbE Ports	41
Disable Unused FCoE Capability on CNA Ports	42
Configure Network Time Protocol	42
Configure Simple Network Management Protocol.....	42
Configure AutoSupport	43
Enable Cisco Discovery Protocol	43
Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP	43
Create Interface Groups	43
Create VLANs.....	44
Create Storage Virtual Machine	44
Create Load-Sharing Mirrors of SVM Root Volume	44
Create Block Protocol (iSCSI) Service.....	45
Configure HTTPS Access	45
Configure NFSv3	46
Create FlexVol Volumes.....	46
Create Boot LUNs.....	47
Schedule Deduplication	47
Create iSCSI LIFs.....	47
Create NFS LIF	47
Add Infrastructure SVM Administrator.....	48
Server Configuration	49
Cisco UCS Base Configuration.....	49
Perform Initial Setup of Cisco UCS Fabric Interconnects for FlexPod Environments	49
Cisco UCS Setup.....	51

Log in to Cisco UCS Manager	51
Upgrade Cisco UCS Manager Software to Version 3.1(2f)	51
Anonymous Reporting	51
Configure Cisco UCS Call Home	52
Add Block of IP Addresses for KVM Access	52
Synchronize Cisco UCS to NTP.....	53
Edit Chassis Discovery Policy	54
Enable Server and Uplink Ports.....	55
Acknowledge Cisco UCS Chassis and FEX	56
Create Uplink Port Channels to Cisco Nexus Switches	57
Create MAC Address Pools	58
Create IQN Pools for iSCSI Boot	60
Create IP Pools for iSCSI Boot	61
Create UUID Suffix Pool.....	62
Create Server Pool	63
Create VLANs.....	64
Modify Default Host Firmware Package	67
Set Jumbo Frames in Cisco UCS Fabric.....	68
Create Local Disk Configuration Policy (Optional)	69
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	70
Create Power Control Policy.....	71
Create Server Pool Qualification Policy (Optional).....	72
Create Server BIOS Policy	73
Update the Default Maintenance Policy.....	76
Create vNIC Templates.....	77
Create LAN Connectivity Policy for iSCSI Boot.....	81
Create vMedia Policy for VMware ESXi 6.5a Install Boot	84
Create iSCSI Boot Policy	86
Create Service Profile Templates.....	88
Create vMedia-Enabled Service Profile Template	101
Create Service Profiles	102
Add More Servers to FlexPod Unit.....	102
Gather Necessary Information.....	102
Storage Configuration – Boot LUNs and Igroups	104
ONTAP Boot Storage Setup.....	104

Create igroups.....	104
Map Boot LUNs to igroups.....	104
VMware vSphere 6.5a Setup	105
VMware ESXi 6.5a	105
Download Cisco Custom Image for ESXi 6.5a.....	105
Log in to Cisco UCS 6300/6200 Fabric Interconnect	105
Set Up VMware ESXi Installation.....	106
Install ESXi.....	106
Set Up Management Networking for ESXi Hosts	107
Log in to VMware ESXi Hosts by Using VMware Host Client	109
Set Up VMkernel Ports and Virtual Switch.....	109
Setup iSCSI Multipathing	113
Mount Required Datastores	114
Configure NTP on ESXi Hosts	117
VMware vCenter 6.5a	118
Building the VMware vCenter Server Appliance	118
Setting Up VMware vCenter Server	129
ESXi Dump Collector setup for iSCSI-Booted Hosts.....	135
Cisco UCS Manager Plug-in for VMware vSphere Web Client	135
Cisco UCS Manager Plug-in Installation.....	136
Cisco UCS Domain Registration.....	138
Using the Cisco UCS vCenter Plugin.....	139
FlexPod VMware vSphere Distributed Switch (vDS).....	142
Configure the VMware vDS in vCenter	143
FlexPod Management Tools Setup.....	151
NetApp Virtual Storage Console 6.2.1 Deployment Procedure.....	151
Virtual Storage Console 6.2.1P1 Pre-installation Considerations	151
Install Virtual Storage Console 6.2.1P1	151
Register Virtual Storage Console with vCenter Server.....	153
Install NetApp NFS VAAI Plug-in.....	153
Discover and Add Storage Resources	154
Optimal Storage Settings for ESXi Hosts.....	154
Virtual Storage Console 6.2.1P1 Provisioning Datastores	156
Virtual Storage Console 6.2.1P1 Backup and Recovery	164
Sample Tenant Provisioning.....	170

Provisioning a Sample Application Tenant	170
Appendix	172
FlexPod Backups	172
Cisco UCS Backup	172
Cisco Nexus Backups	173
VMware VCSA Backup	174
Breakout Interface Configuration in the Cisco Nexus 9332PQ Switches	176
About the Authors.....	178
Acknowledgements	178



Executive Summary

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

This document describes the Cisco and NetApp® FlexPod Datacenter with Cisco UCS Manager unified software release 3.1(2f) and VMware vSphere 6.5a. Cisco UCS Manager (UCSM) 3.1 provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with Cisco UCS unified software release 3.1(2f), and VMware vSphere 6.5a is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches and NetApp AFF.

Solution Overview

Introduction

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

Audience

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

Purpose of this Document

This document provides a step by step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS Fabric Interconnects, NetApp AFF, and Cisco Nexus 9000 solution.

What's New?

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

- Support for the Cisco UCS 3.1(2f) unified software release, Cisco UCS B200-M4 servers, and Cisco UCS C220-M4 servers
- Support for the latest release of NetApp Data ONTAP® 9.1
- iSCSI, Fiber Channel and NFS Storage Design
- Validation of VMware vSphere 6.5a

Solution Design

Architecture

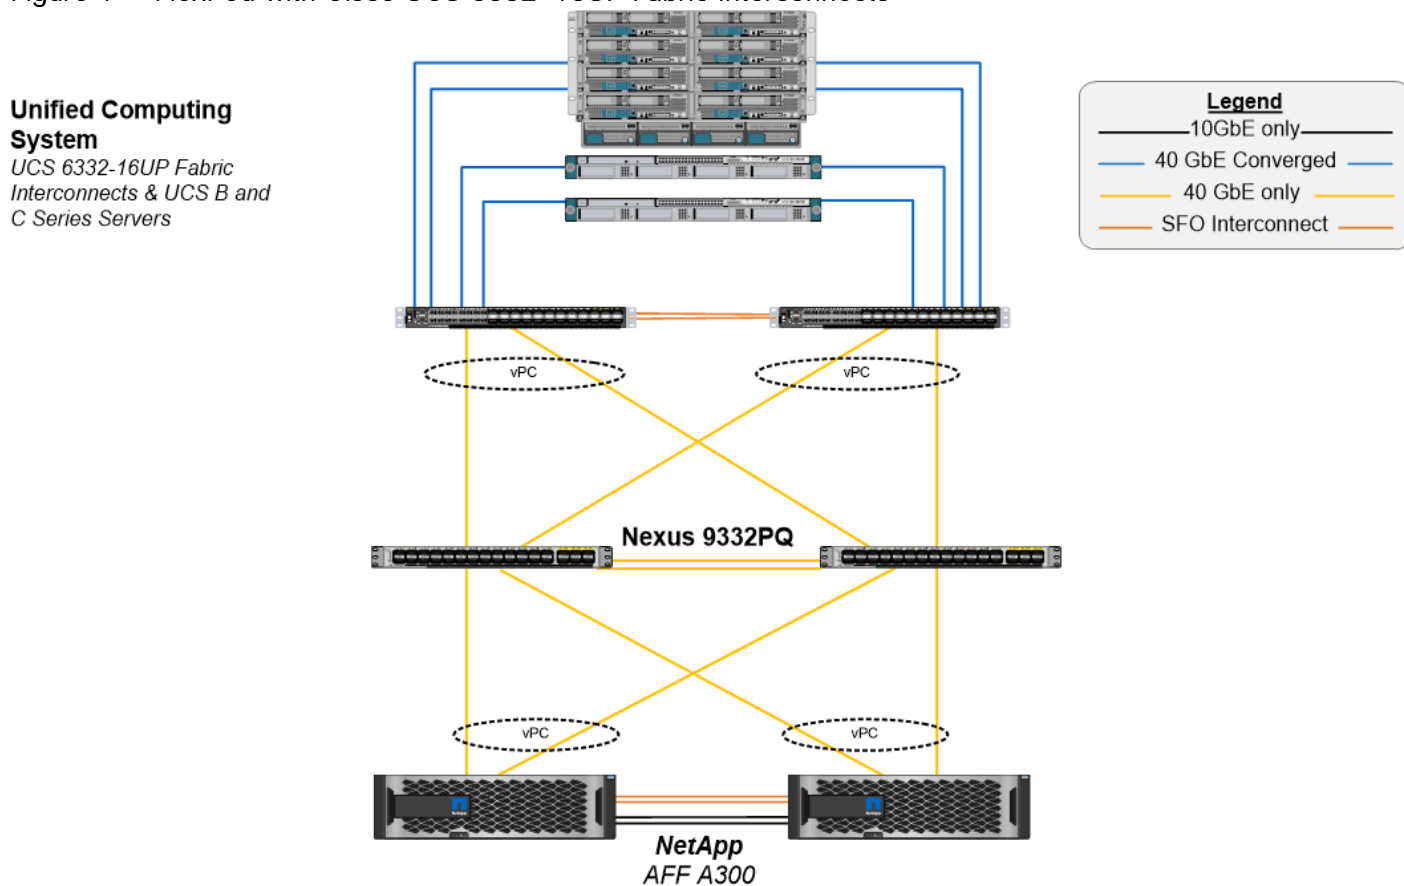
FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere® built on FlexPod includes NetApp All Flash FAS storage, Cisco Nexus® networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere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

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

Figure 1 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with the Cisco UCS 6332-16UP Fabric Interconnects. This design has end-to-end 40 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and C-Series rackmounts and the Cisco UCS Fabric Interconnect, between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A300. 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

Figure 1 FlexPod with Cisco UCS 6332-16UP Fabric Interconnects



The reference 40Gb based hardware configuration includes:


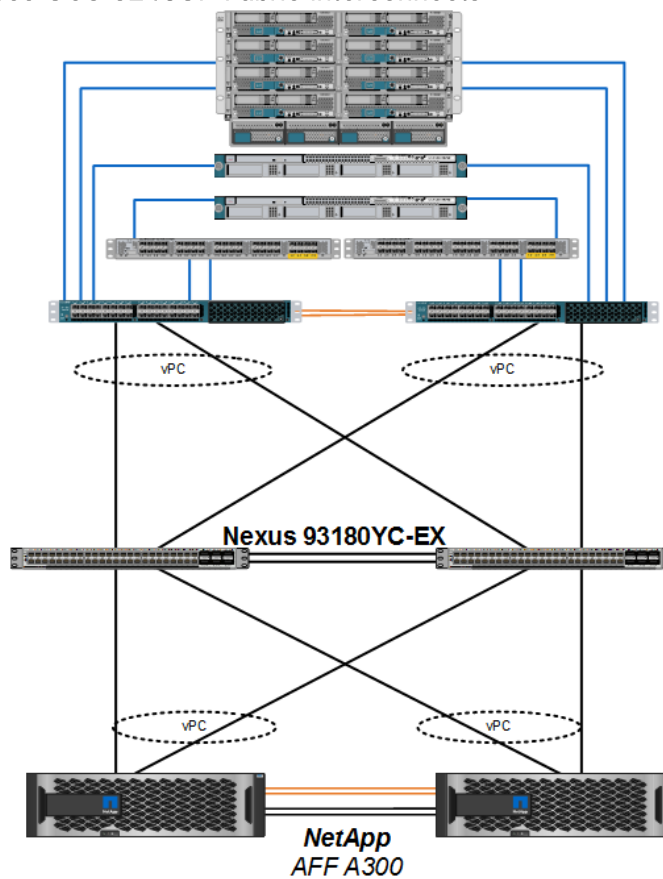
- Two Cisco Nexus 9332PQ switches
- Two Cisco UCS 6332-16UP fabric interconnects
- One NetApp AFF A300 (HA pair) running clustered Data ONTAP with Disk shelves and Solid State Drives (SSD)

Figure 2 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with the Cisco UCS 6248UP Fabric Interconnects. This design is identical to the 6332-16UP based topology, but has 10 Gb Ethernet connecting through a pair of Cisco Nexus 93180YC-EX switches to access iSCSI and NFS access to the AFF A300. Alternately, the Cisco Nexus 9332PQ switch can be used with the Cisco UCS 6248UP with a QSFP breakout cable and port configuration setting on the 9332PQ switch.

Figure 2 FlexPod with Cisco UCS 6248UP Fabric Interconnects

Unified Computing System

UCS 6248 Fabric Interconnects,
Nexus 2232 Fabric Extender
& UCS C and B Series Servers



The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-EX switches
- Two Cisco UCS 6248UP fabric interconnects
- One NetApp AFF A300 (HA pair) running clustered Data ONTAP with Disk shelves and Solid State Drives (SSD)

For server virtualization, the deployment includes VMware vSphere 6.5a. 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 storage controllers or disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps for deploying the base architecture, as shown in Figure 2. 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 and 6300 Series, UCS B-200 M4, UCS C-220 M4	3.1(2f)	Includes the Cisco UCS-IOM 2304 Cisco UCS Manager, Cisco UCS VIC 1340 and Cisco UCS VIC 1385
Network	Cisco Nexus 9000 NX-OS	7.0(3)I4(5)	
Storage	NetApp AFF A300	Data ONTAP 9.1	
Software	Cisco UCS Manager	3.1(2f)	
	Cisco UCS Manager Plugin for VMware vSphere Web Client	2.0.1	
	VMware vSphere ESXi	6.5a	
	VMware vCenter	6.5a	
	NetApp Virtual Storage Console (VSC)	6.2.1P1	

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-01, VM-Host-Infra-02 to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?
```

```
[-node] <nodename>
```

```
Node
```

```

{ [-vlan-name] {<netport>|<ifgrp>} VLAN Name
| -port {<netport>|<ifgrp>} Associated Network Port
[-vlan-id] <integer> } Network Switch VLAN Identifier
    
```

Example:

```

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

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

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out of Band Mgmt	VLAN for out-of-band management interfaces	13
In-Band Mgmt	VLAN for in-band management interfaces	113
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for Infrastructure NFS traffic	3050
iSCSI-A	VLAN for Infrastructure iSCSI A traffic	3010
iSCSI-B	VLAN for Infrastructure iSCSI B traffic	3020
vMotion	VLAN for VMware vMotion	3000
VM-Traffic	VLAN for Production VM Interfaces	900

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

Table 3 Virtual Machines

Virtual Machine Description	Host Name
Active Directory (AD)	
vCenter Server	
NetApp VSC	

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 details for the prescribed and supported configuration of the NetApp AFF A300 running NetApp ONTAP® 9.1.



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 use the cabling directions in this section as a guide.

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 3 details the cable connections used in the validation lab for the 40Gb end-to-end iSCSI topology based on the Cisco UCS 6332-16UP fabric interconnect. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has two connections to the out-of-band network switch.

Figure 3 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnect

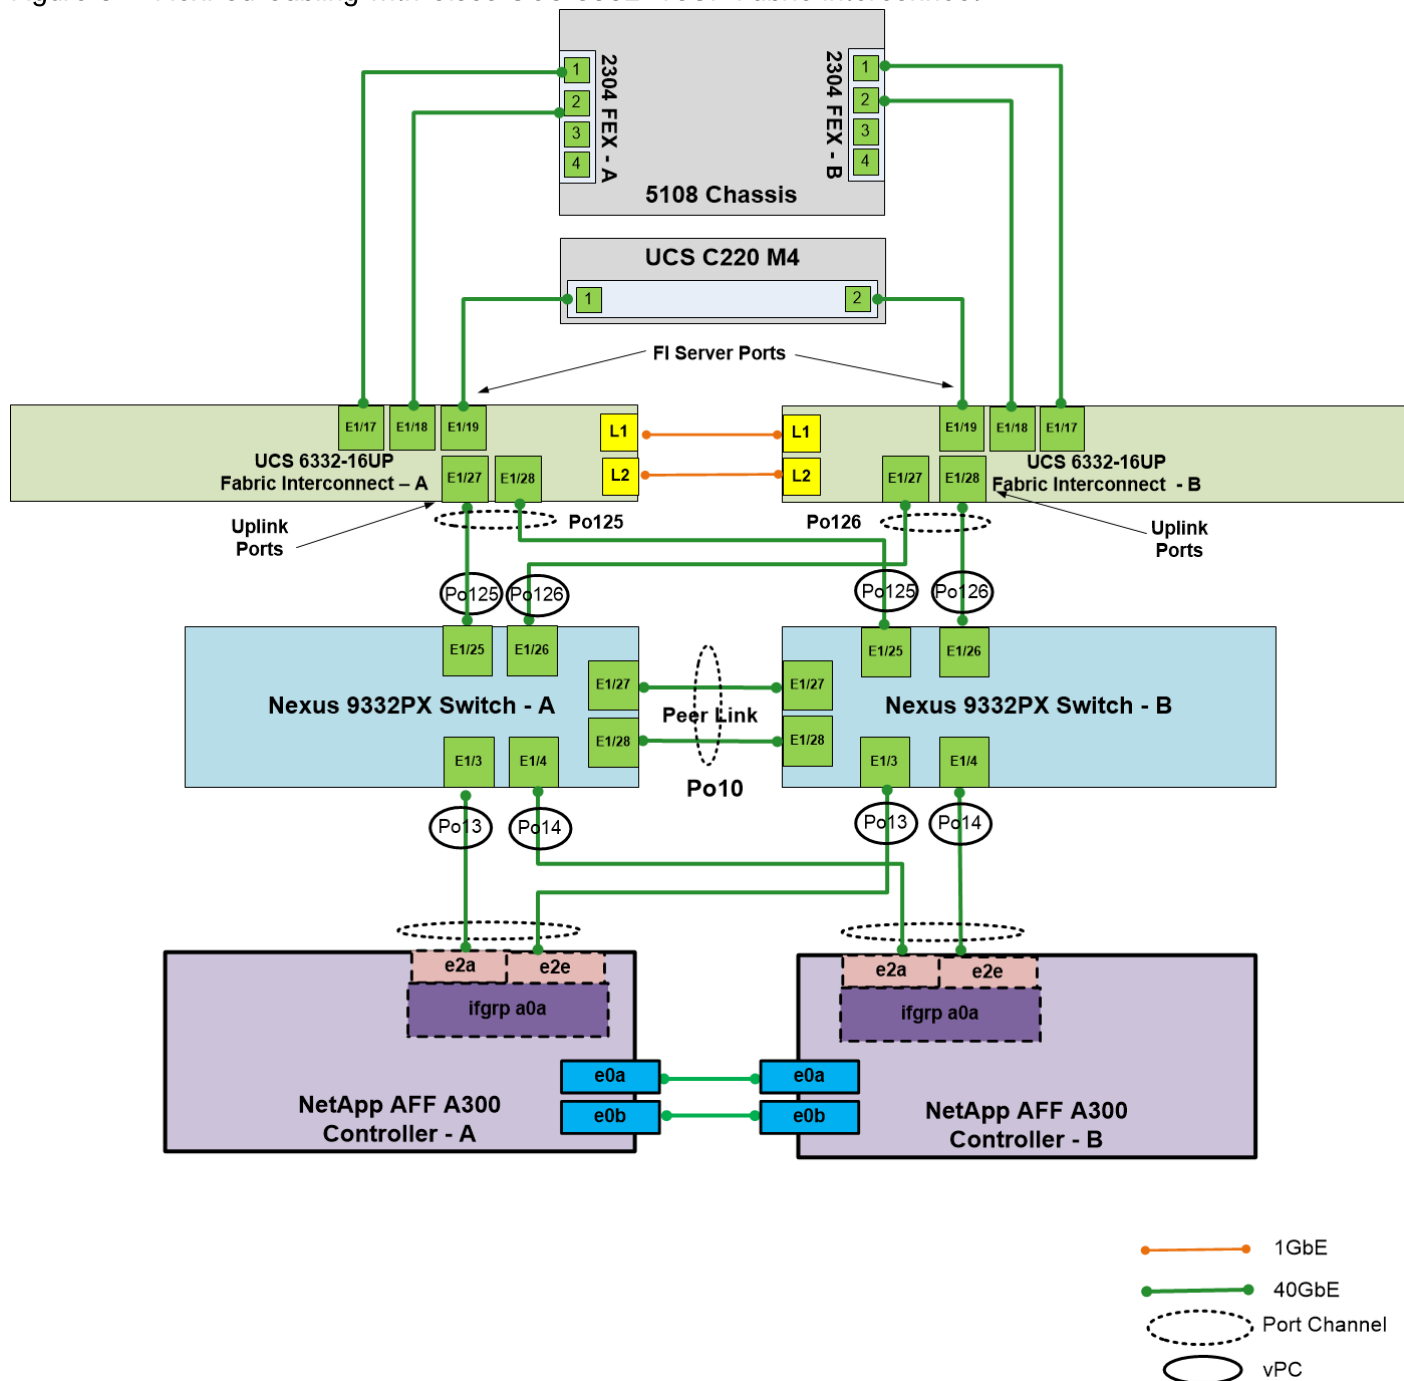
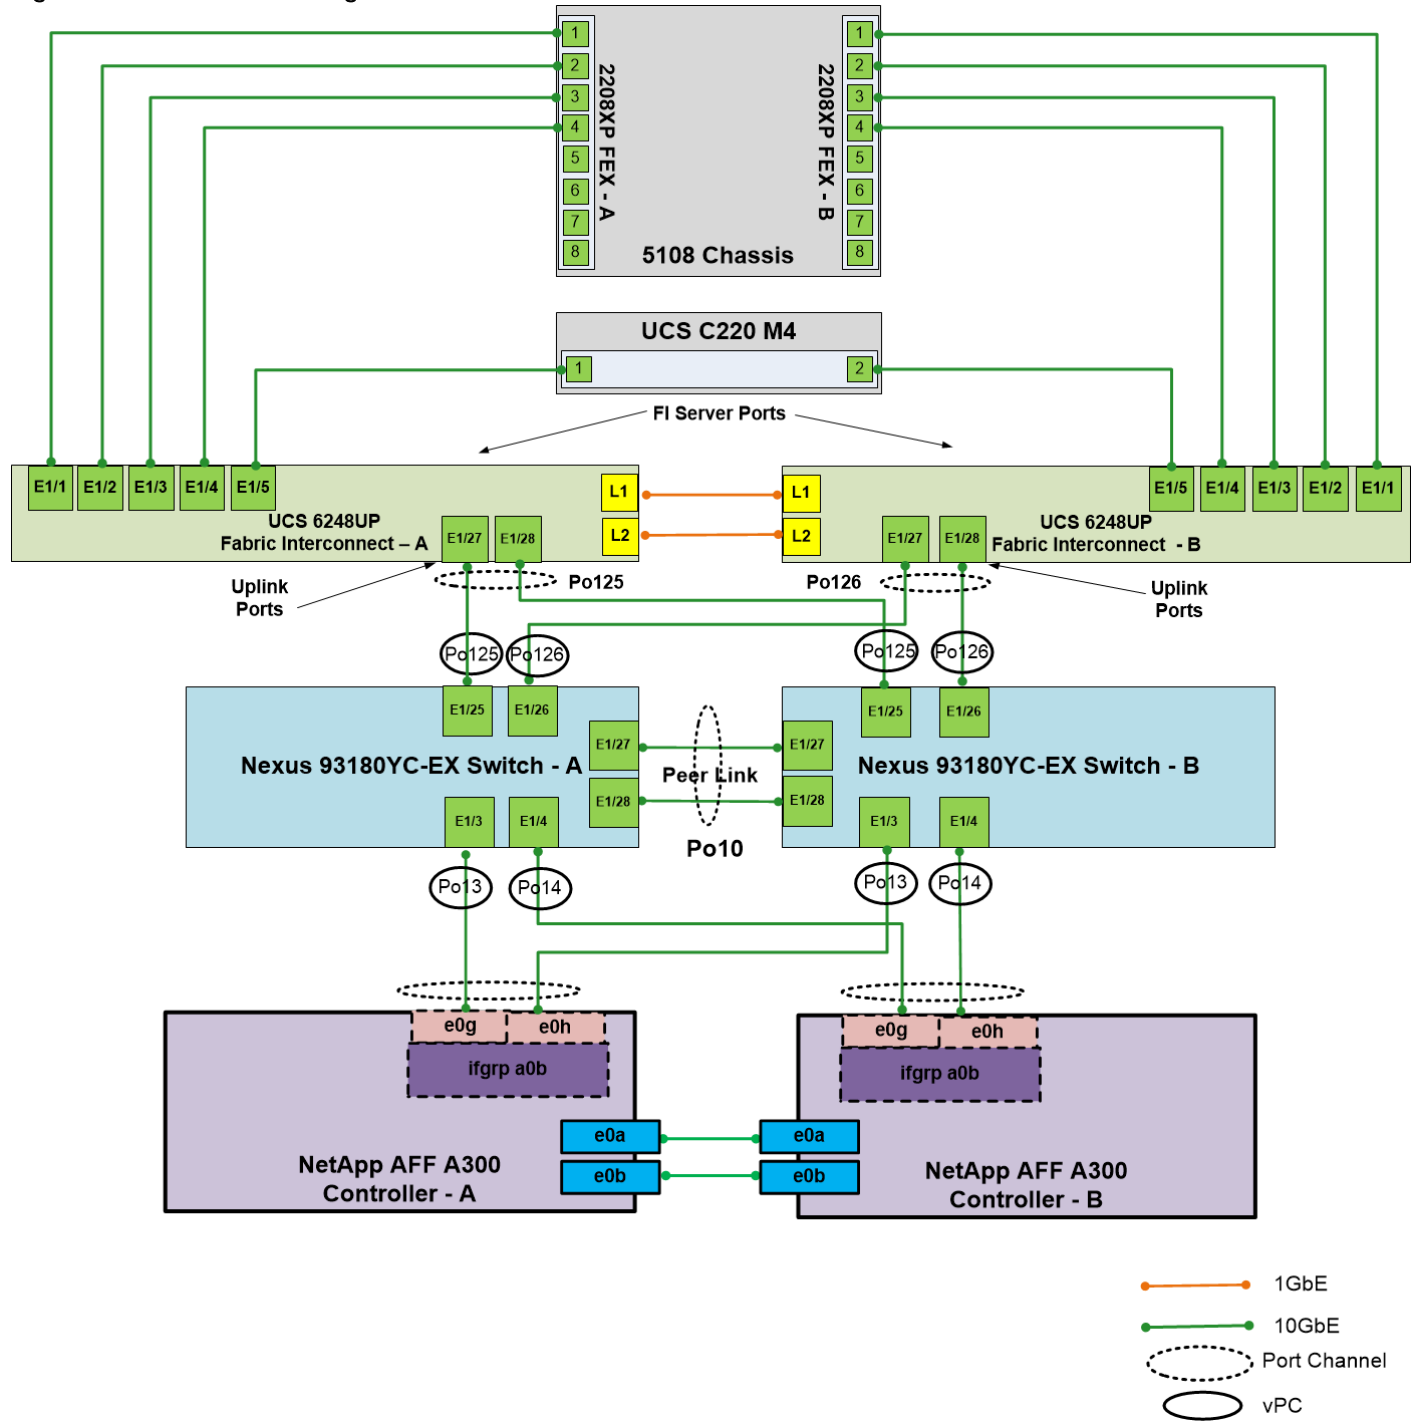


Figure 4 details the cabling connections used in the alternate 10Gb end-to-end iSCSI topology based on the Cisco UCS 6248UP fabric interconnect. As with the 40Gb topology, out-of-band connections will also be needed, with each Cisco UCS fabric interconnect and Cisco Nexus Switch will have a connection to the out-of-band network switch, and each AFF controller will have two connections to the out-of-band network switch.

Figure 4 FlexPod Cabling with Cisco UCS 6248UP Fabric Interconnect



Network Switch Configuration

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

Physical Connectivity

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

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 7.0(3)I4(5), and is valid for both the Cisco Nexus 9332PO switches deployed with the 40Gb end-to-end topology, and the Cisco Nexus 93180YC-EX switches used in the 10Gb based topology.



The following procedure includes the 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 that the default VRF is used to route the in-band management VLAN.

Set Up Initial Configuration

Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <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) [y]: Enter
```

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

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

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

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

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

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

```
Enter the switch name: <nexus-A-hostname>
```

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

```
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <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: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
2. Review the configuration summary before enabling the configuration.
Use this configuration and save it? (yes/no) [y]: Enter
```

Cisco Nexus B

To set up the initial configuration for the Cisco Nexus B switch on <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) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
```

```
Enter the switch name: <nexus-B-hostname>

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

Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>

Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>

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

IPv4 address of the default gateway: <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: <global-ntp-server-ip>

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

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

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

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 A and Cisco Nexus 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 A and Cisco Nexus B

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

Run the following commands to set global configurations:

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

Create VLANs

Cisco Nexus A and Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
vlan <infra-nfs-vlan-id>
name Infra-NFS-VLAN
vlan <infra-iSCSI-A-vlan-id>
name iSCSI-A-VLAN
vlan <infra-iSCSI-B-vlan-id>
name iSCSI-B-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus A

From the global configuration mode, run the following commands:

```
ntp source <switch-a-ntp-ip>

interface Vlan<ib-mgmt-vlan-id>

ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>

no shutdown

exit
```

Cisco Nexus B

From the global configuration mode, run the following commands:

```
ntp source <switch-b-ntp-ip>

interface Vlan<ib-mgmt-vlan-id>

ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>

no shutdown

exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus A

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



In this step and in the later sections, configure the AFF nodename <st-node> and Cisco UCS 6332-16UP or Cisco UCS 6248UP fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface Eth1/3

description <st-node>-1:e2a

interface Eth1/4

description <st-node>-2:e2a

interface Eth1/25

description <ucs-clustername>-a:1/27

interface Eth1/26

description <ucs-clustername>-b:1/27

interface Eth1/27
```

```
description <nexus-hostname>-b:1/27  
  
interface Eth1/28  
description <nexus-hostname>-b:1/28  
  
exit
```

Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
interface Eth1/3  
  
description <st-node>-1:e2e  
  
interface Eth1/4  
description <st-node>-2:e2e  
  
interface Eth1/25  
description <ucs-clustername>-a:1/28  
  
interface Eth1/26  
description <ucs-clustername>-b:1/28  
  
interface Eth1/27  
description <nexus-hostname>-a:1/27  
  
interface Eth1/28  
description <nexus-hostname>-a:1/28  
  
exit
```

Create Port Channels

Cisco Nexus A and Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
interface Po10  
  
description vPC peer-link  
  
interface Eth1/27-28  
channel-group 10 mode active  
  
no shutdown  
  
interface Po13  
  
description <st-node>-1
```

```

interface Eth1/3

channel-group 13 mode active

no shutdown

interface Po14

description <st-node>-2

interface Eth1/4

channel-group 14 mode active

no shutdown

interface Po125

description <ucs-clustername>-a

interface Eth1/25

channel-group 125 mode active

no shutdown

interface Po126

description <ucs-clustername>-b

interface Eth1/26

channel-group 126 mode active

no shutdown

exit

copy run start

```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

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

From the global configuration mode, run the following commands:

```

interface Po10

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>, <vm-traffic-
vlan-id>, <infra-iSCSI-A-vlan-id>, <infra-iSCSI-B-vlan-id>

spanning-tree port type network

interface Po13

switchport mode trunk

```



```
switchport trunk native vlan 2

switchport trunk allowed vlan <infra-nfs-vlan-id>, <infra-iSCSI-A-vlan-id>, <infra-iSCSI-B-vlan-id>

spanning-tree port type edge trunk

mtu 9216

interface Po14

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <infra-nfs-vlan-id>, <infra-iSCSI-A-vlan-id>, <infra-iSCSI-B-vlan-id>

spanning-tree port type edge trunk

mtu 9216

interface Po125

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>, <vm-traffic-
vlan-id>, <infra-iSCSI-A-vlan-id>, <infra-iSCSI-B-vlan-id>

spanning-tree port type edge trunk

mtu 9216

interface Po126

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>, <vm-traffic-
vlan-id>, <infra-iSCSI-A-vlan-id>, <infra-iSCSI-B-vlan-id>

spanning-tree port type edge trunk

mtu 9216

exit

copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

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

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>

role priority 10
```

```
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>

peer-switch

peer-gateway

auto-recovery

delay restore 150

interface Po10

vpc peer-link

interface Po13

vpc 13

interface Po14

vpc 14

interface Po125

vpc 125

interface Po126

vpc 126

exit

copy run start
```

Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>

role priority 20

peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>

peer-switch

peer-gateway

auto-recovery

delay restore 150

interface Po10

vpc peer-link

interface Po13

vpc 13

interface Po14

vpc 14
```

```
interface Po125
vpc 125
interface Po126
vpc 126
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, we recommend using vPCs to uplink the Cisco Nexus 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

AFF A300 Controllers

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

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

NetApp Hardware Universe

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



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

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

Controllers

Follow the physical installation procedures for the controllers found in the [AFF A300 Series product documentation](#) at the [NetApp Support](#) site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported by the AFF A300 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 9.1

Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the [ONTAP 9.1 Software Setup Guide](#). You must have access to the [NetApp Support](#) site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [ONTAP 9.1 Software Setup Guide](#) to learn about configuring ONTAP. Table 4 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 4 ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Data ONTAP 9.1 URL	<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 the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select `e0M` for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for `e0M`.

```
<node01-mgmt-ip> <node01-mgmt-mask> <node01-mgmt-gateway>
```

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



This web server must be pingable.

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.



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 the following message displays:

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

14. Select option 4 for **Clean Configuration and Initialize All Disks**.
15. Enter `y` to **zero disks, reset config, and install a new file system**.
16. Enter `y` to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. 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 the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M.

```
<node02-mgmt-ip> <node02-mgmt-mask> <node02-mgmt-gateway>
```

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



This web server must be pingable.

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.



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.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

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 ONTAP 9.1 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.

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.

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: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created

Use your web browser to complete cluster setup by accesing https://<node01-mgmt-ip>

Otherwise press Enter to complete cluster setup using the command line interface:
```

2. To complete the cluster setup, open a web browser and navigate to <https://<node01-mgmt-ip>>.

Table 5 Cluster create in ONTAP prerequisites

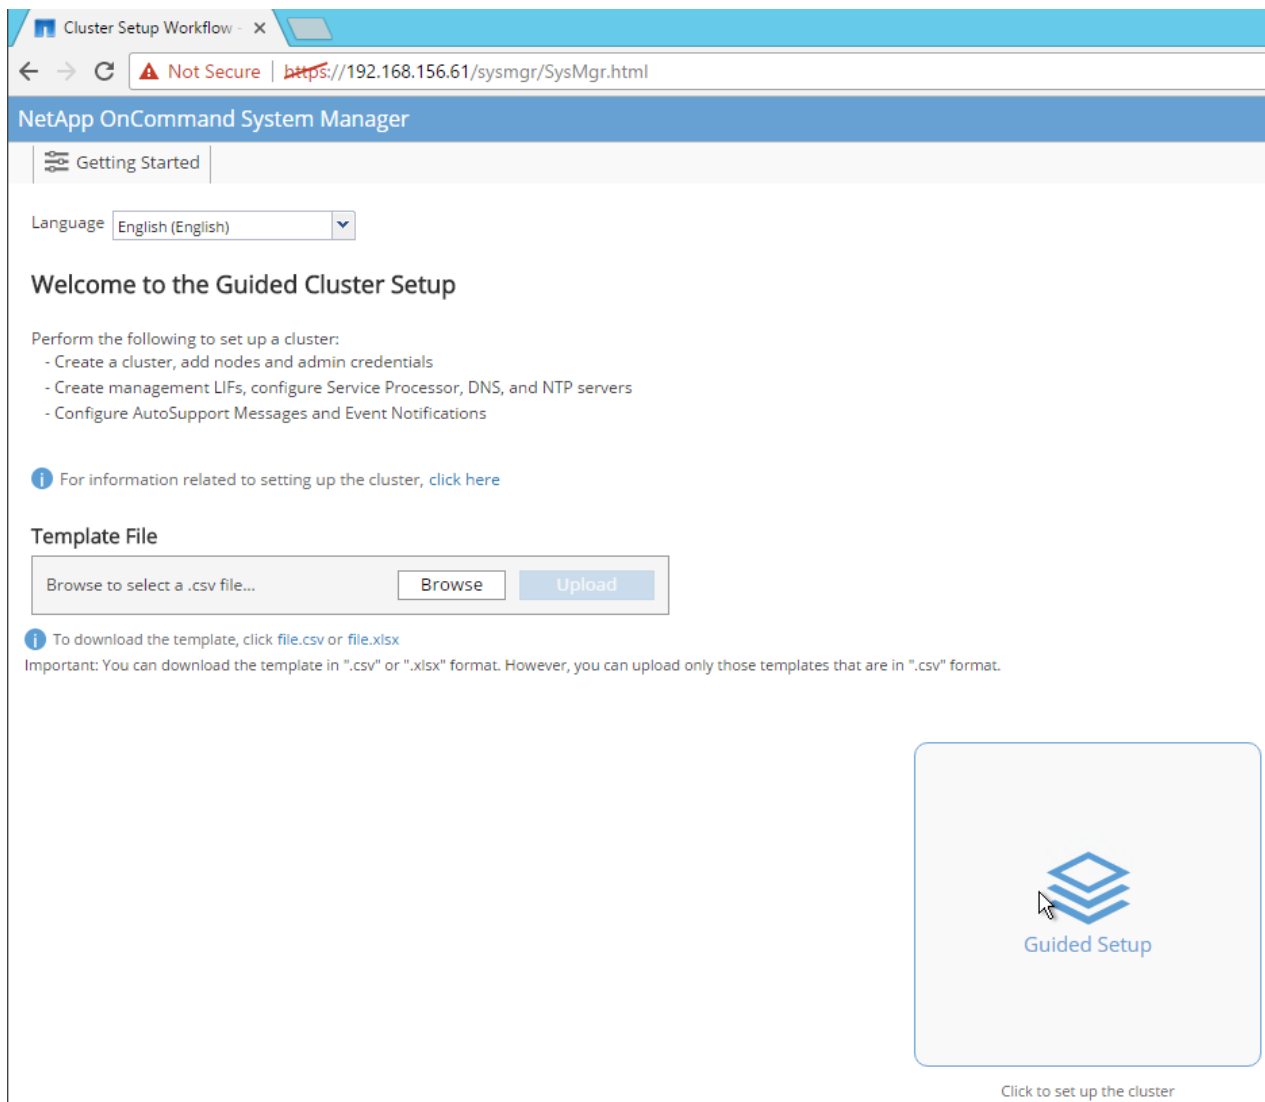
Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>

Cluster Detail	Cluster Detail Value
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-SP-ip>
Node 02 service processor IP address	<node02-SP-ip>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server IP address	<ntp-ip>




Cluster setup can also be done using command line interface. This document describes the cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup on the Welcome screen.



4. In the Cluster screen, do as follows:


- Enter the cluster and node names.
- Select the cluster configuration.
- Enter and confirm the password.
- (Optional) Enter the cluster base and feature licenses.



Cluster Name

Nodes

Not sure all nodes have been discovered? [Refresh](#)



Cluster Configuration: Switched Cluster Switchless Cluster

Ensure that the hardware connectivity is set up for the two-node switchless cluster.

Username

Password

Confirm Password

Cluster Base License (Optional)

For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional)

Cluster Base License is mandatory to add Feature Licenses.



The nodes are discovered automatically; if they are not, click the Refresh link. By default, the cluster interfaces will be created on all the new factory shipping storage controllers.



If all the nodes are not discovered, then configure the cluster using the command line.




Cluster license and feature licenses can also be installed after completing the cluster creation.

5. Click Submit.
6. In the network page, complete the following sections:
 - Cluster Management
 - Enter the IP address, netmask, gateway and port details.
 - Node Management
 - Enter the node management IP addresses and port details for all the nodes.
 - Service Processor Management
 - Enter the IP addresses for all the nodes.
 - DNS Details
 - Enter the DNS domain names and server address.
 - NTP Details
 - Enter the primary and alternate NTP server.
7. Click Submit.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



Network (Management)

IP Addresses (IPv4) required Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

IP Address Range
 You must enter the default network details manually.

	IP Address	Netmask	Gateway (Optional)	Port
Cluster Management	<input type="text" value="192.168.156.60"/>	<input type="text" value="255.255.255.0"/>	<input type="text" value="192.168.156.1"/>	<input type="text" value="e0c"/>
<p>Node Management <input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.</p>				
bb04-aff300-1	<input type="text" value="192.168.156.61"/>	<input type="text" value="e0M"/>		<input type="text"/>
bb04-aff300-2	<input type="text" value="192.168.156.62"/>	<input type="text" value="e0M"/>		<input type="text"/>
<p>Service Processor Management Default values have been detected for the Service Processor.</p> <p><input type="checkbox"/> Override the default values (Gateway is mandatory)</p> <p><input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.</p>				
bb04-aff300-1	<input type="text" value="192.168.156.58"/>			
bb04-aff300-2	<input type="text" value="192.168.156.59"/>			

DNS Details

DNS Domain Names

DNS Server IP Address

NTP Details

Primary NTP Server

Alternative NTP Server (Optional)

- In the Support page, configure the AutoSupport and Event Notifications sections.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? AutoSupport

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

? Event Notifications

Notify me through:

<input checked="" type="checkbox"/>	Email	SMTP Mail Host <input type="text" value="testvikings.smtp.cisco.com"/>	Email Addresses <input type="text" value="adminvikings@cisco.com"/>
<input type="checkbox"/>	SNMP	SNMP Trap Host <input type="text"/>	
<input type="checkbox"/>	Syslog	Syslog Server <input type="text"/>	

Submit

- Click Submit.
- In the Summary page, review the configuration details if needed.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



[Click here to view the summary](#)

The next step will be to configure your aggregates, SVM and Storage Objects.
Click the button below to start provisioning your storage.

[Manage your cluster](#)



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, we assume that it is on the same subnet.

Log In to the Cluster

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

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

Zero All Spare Disks

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

```
disk zerospares
```



Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an All Flash FAS configuration. Disk autoassign should have assigned one data partition to each node in an 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 partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Onboard Unified Target Adapter 2 Port Personality

To set the personality of the onboard unified target adapter 2 (UTA2), complete the following steps:

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

```
ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
<st-node01>	0e	cna	target	-	-	online
<st-node01>	0f	cna	target	-	-	online
<st-node01>	0g	cna	target	-	-	online
<st-node01>	0h	cna	target	-	-	online
<st-node02>	0e	cna	target	-	-	online
<st-node02>	0f	cna	target	-	-	online
<st-node02>	0g	cna	target	-	-	online
<st-node02>	0h	cna	target	-	-	online

8 entries were displayed.

- Verify that the Current Mode and Current Type properties for all ports are set properly. Set the ports used for iSCSI connectivity to mode cna. The port type for all protocols should be set to target. Change the port personality by running the following command:

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



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

Set Auto-Revert on Cluster Management

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



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

Run the following command:

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

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, e0d, e2a, and e2e) should be removed from the default broadcast domain, leaving just the management network ports (e0c and e0M). To perform this task, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports <st-node01>:e0d,<st-node01>:e0e,<st-node01>:e0f,<st-node01>:e0g,<st-node01>:e0h,<st-node01>:e2a, <st-node01>:e2e,<st-node02>:e0d,<st-node02>:e0e,<st-node02>:e0f,<st-node02>:e0g,<st-node02>:e0h <st-node02>:e2a,<st-node02>:e2e
broadcast-domain show
```

Set Up Service Processor Network Interface

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

```
system service-processor network modify -node <st-node01> -address-family IPv4 -enable true -dhcp
none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>

system service-processor network modify -node <st-node02> -address-family IPv4 -enable true -dhcp
none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```



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

Create Aggregates

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

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <st-node01> -diskcount <num-disks>
aggr create -aggregate aggr1_node02 -node <st-node02> -diskcount <num-disks>
```



You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate.



For all flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.



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



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

2. (Optional) Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02. The aggregate will be automatically renamed if System guided setup is used.

```
aggr show
aggr rename -aggregate aggr0 -newname <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 the storage failover.

```
storage failover show
```



Both <st-node01> and <st-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 <st-node01> -enabled true
```



Enabling failover on one node enables it for both nodes.

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



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

```
cluster ha show
```

4. Continue with step 6 if high availability is configured.
5. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

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

```
storage failover hwassist show
storage failover modify -hwassist-partner-ip <node02-mgmt-ip> -node <st-node01>
storage failover modify -hwassist-partner-ip <node01-mgmt-ip> -node <st-node02>
```

Disable Flow Control on 10GbE and 40GbE Ports

NetApp recommends disabling flow control on all 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 <st-node01> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -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 <st-node02> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -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 Capability on CNA Ports

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

```
fc adapter modify -node <st-node01> -adapter 0e -status-admin down
fc adapter modify -node <st-node01> -adapter 0f -status-admin down
fc adapter modify -node <st-node01> -adapter 0g -status-admin down
fc adapter modify -node <st-node01> -adapter 0h -status-admin down
fc adapter modify -node <st-node02> -adapter 0e -status-admin down
fc adapter modify -node <st-node02> -adapter 0f -status-admin down
fc adapter modify -node <st-node02> -adapter 0g -status-admin down
fc adapter modify -node <st-node02> -adapter 0h -status-admin down
fc adapter show -fields -status-admin
```

Configure Network Time Protocol

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

1. Set the time zone for the cluster.

```
timezone <timezone>
```



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

2. Set the date for the cluster.

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



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

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

```
cluster time-service ntp server create -server <switch-a-ntp-ip>
cluster time-service ntp server create -server <switch-b-ntp-ip>
```

Configure Simple Network Management Protocol

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

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

```
snmp contact <snmp-contact>
snmp location "<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 <oncommand-um-server-fqdn>
```

Configure SNMPv1 Access

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

```
snmp community add ro <snmp-community>
```

Configure AutoSupport

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

```
system node autosupport modify -node * -state enable -mail-hosts <mailhost> -transport https -support enable -noteto <storage-admin-email>
```

Enable Cisco Discovery Protocol

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

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



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

Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

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

```
ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2e

ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2e

ifgrp show
```

Create VLANs

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

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <st-node01>:a0a-<infra-nfs-vlan-id>,
<st-node02>:a0a-<infra-nfs-vlan-id>
```

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

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-b-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-b-vlan-id>

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <st-node01>:a0a-<infra-iscsi-a-
vlan-id>, <st-node02>:a0a-<infra-iscsi-a-vlan-id>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <st-node01>:a0a-<infra-iscsi-b-
vlan-id>, <st-node02>:a0a-<infra-iscsi-b-vlan-id>
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

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

2. Select the SVM data protocols to configure, keeping iSCSI and NFS.

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

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

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

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

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

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

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

Create Load-Sharing Mirrors of SVM Root Volume

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

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

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

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

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

3. Create the mirroring relationships.

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

4. Initialize the mirroring relationship.

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

Create Block Protocol (iSCSI) Service

Run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

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

Configure HTTPS Access

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

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

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

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

```
security certificate show
```

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

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



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

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

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

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

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

- Disable HTTP cluster management access.

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



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

- 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:

- Create a new rule for the infrastructure NFS subnet in the default export policy.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -clientmatch <infra-nfs-subnet-cidr> -rorule sys -rwrule sys -superuser sys -allow-suid false
```

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

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

Create FlexVol Volumes

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

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

To create a FlexVol volume, run the following commands:

```

volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node01 -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_datastore_2 -aggregate aggr1_node02 -size 500GB -state
online -policy default -junction-path /infra_datastore_2 -space-guarantee none -percent-snapshot-
space 0
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_node01 -size 100GB -state online
-policy default -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path Infra-SVM:rootvol

```

Create Boot LUNs

To create two boot LUNs, run the following commands:

```

lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 15GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-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 deduplication schedule to `esxi_boot` and `infra_datastore_1`:

```

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

```

Create iSCSI LIFs

Run the following commands to create four iSCSI LIFs (two on each node):

```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<infra-iscsi-a-vlan-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 <st-node01> -home-port a0a-<infra-iscsi-b-vlan-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 <st-node02> -home-port a0a-<infra-iscsi-a-vlan-id> -address <var_node02_iscsi_lif02a_ip> -
netmask <var_node02_iscsi_lif02a_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 <st-node02> -home-port a0a-<infra-iscsi-b-vlan-id> -address <var_node02_iscsi_lif02b_ip> -
netmask <var_node02_iscsi_lif02b_mask> -status-admin up -failover-policy disabled -firewall-policy
data -auto-revert false

network interface show

```

Create NFS LIF

To create an NFS LIF, run the following commands:

```

network interface create -vserver Infra-SVM -lif nfs_lif01 -role data -data-protocol nfs -home-node
<st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs_lif01-ip> -netmask <node01-
nfs_lif01-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface create -vserver Infra-SVM -lif nfs_lif02 -role data -data-protocol nfs -home-node
<st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs_lif02-ip> -netmask <node02-
nfs_lif02-mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface show

```

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 svm-mgmt -role data -data-protocol none -home-node
<st-node02> -home-port e0c -address <svm-mgmt-ip> -netmask <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 <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: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver Infra-SVM

```



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

Server Configuration

Cisco UCS Base Configuration

This FlexPod deployment will show configuration steps for the Cisco UCS Fabric Interconnects (FI) in a design that will support iSCSI to the NetApp AFF through the Cisco Nexus.

Perform Initial Setup of Cisco UCS Fabric Interconnects for FlexPod Environments

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

Cisco UCS Fabric Interconnect 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 fabric interconnect.

```
Enter the configuration method: gui
```

```
Physical switch Mgmt0 IP address: <ucsa-mgmt-ip>
```

```
Physical switch Mgmt0 IPv4 netmask: <ucsa-mgmt-mask>
```

```
IPv4 address of the default gateway: <ucsa-mgmt-gateway>
```

2. Using a supported web browser, connect to <https://<ucsa-mgmt-ip>>, accept the security prompts, and click the **'Express Setup'** link under HTML.
3. Select Initial Setup and click Submit.
4. Select Enable clustering, Fabric A, and IPv4.
5. Fill in the Virtual IP Address with the UCS cluster IP.
6. Completely fill in the System setup section. For system name, use the overall UCS system name. For the Mgmt IP Address, use <ucsa-mgmt-ip>.

Basic Settings

Cluster and Fabric setup

Enable clustering
 Standalone mode
 Synchronize

Fabric Setup: Fabric A Fabric B

IPv4
 IPv6

Virtual IP Address: . . .

System setup

Enforce strong password?: Yes No

System name:

Admin Password: Confirm Admin password:

Mgmt IP Address: . . . Mgmt IP Netmask: . . .

Default Gateway: . . .

DNS Server IP: . . . Domain Name :

UCS Central managed environment

UCS Central IP: . . . Shared Secret:

7. Click Submit.

Cisco UCS Fabric Interconnect 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 fabric interconnect.

Enter the configuration method: `gui`

Physical switch Mgmt0 IP address: `<ucsb-mgmt-ip>`

Physical switch Mgmt0 IPv4 netmask: `<ucsb-mgmt-mask>`

IPv4 address of the default gateway: `<ucsb-mgmt-gateway>`

2. Using a supported web browser, connect to <https://<ucsb-mgmt-ip>>, accept the security prompts, and click the **'Express Setup'** link under HTML.
3. Under System setup, enter the Admin Password entered above and click Submit.
4. Enter <ucsb-mgmt-ip> for the Mgmt IP Address and click Submit.

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 fabric interconnect cluster address.



You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to come up.

2. Click the Launch UCS Manager link under HTML to launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 3.1(2f)

This document assumes the use of Cisco UCS 3.1(2f). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.1(2f), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To create anonymous reporting, complete the following step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products. If you select Yes, enter the IP address of your SMTP Server. Click OK.

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.

If you decide to enable this feature in future, you can do so from the " Anonymous Reporting" in the Call Home settings under the Admin tab.
[View Sample Data](#)

Do you authorize the disclosure of this information to Cisco Smart CallHome?

Yes No

Don't show this message again.

Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in Cisco UCS Manager. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

Add Block of IP Addresses for KVM Access

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

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

Create Block of IPv4 Addresses

From :	<input type="text" value="192.168.156.101"/>	Size :	<input type="text" value="12"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="192.168.156.1"/>
Primary DNS :	<input type="text" value="0.0.0.0"/>	Secondary DNS :	<input type="text" value="0.0.0.0"/>

OK **Cancel**

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

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Expand All > Time Zone Management.
3. Select Timezone.
4. In the Properties pane, select the appropriate time zone in the Timezone menu.
5. Click Save Changes, and then click OK.
6. Click Add NTP Server.
7. Enter <switch-a-ntp-ip> and click OK. Click OK.

8. Click Add NTP Server.

9. Enter <switch-b-ntp-ip> and click OK. Click OK on the confirmation.

All /

Edit Chassis Discovery Policy

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

1. In Cisco UCS Manager, click Equipment on the left and select Equipment in the second list.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.

- Set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.

Equipment



Chassis/FEX Discovery Policy

Action :

Link Grouping Preference : None Port Channel

Backplane Speed Preference : 40G 4x10G

- Click Save Changes.
- Click OK.

Enable Server and Uplink Ports

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

- In Cisco UCS Manager, click Equipment on the left.
- Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
- Expand Ethernet Ports.
- Select the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers, right-click them, and select “Configure as Server Port.”
- Click Yes to confirm server ports and click OK.
- Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
- Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.



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

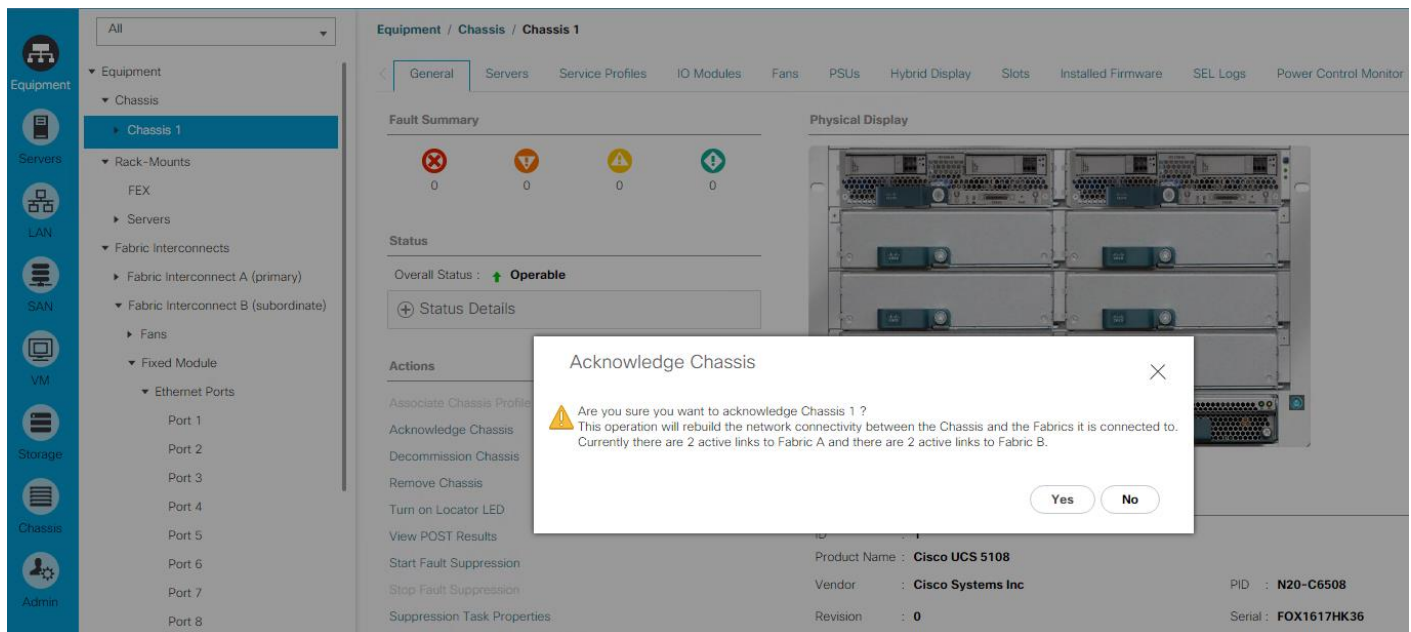
- Click Yes to confirm uplink ports and click OK.
- Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.

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 the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

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

1. In Cisco UCS Manager, click Equipment on the left.
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 Nexus 2232 FEX are 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 LAN on the left.



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 125 as the unique ID of the port channel.
6. Enter vPC-125-Nexus as the name of the port channel.
7. Click Next.
8. Select the ports connected to the Nexus switches to be added to the port channel:
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 126 as the unique ID of the port channel.
16. Enter vPC-126-Nexus as the name of the port channel.
17. Click Next.
18. Select the ports connected to the Nexus switches to be added to the port channel:
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 LAN on the left.
2. Select Pools > root.




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

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



For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the example of also embedding the Cisco UCS domain number information giving us `00:25:B5:32:0A:00` as our first MAC address.

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



12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.
15. Right-click MAC Pools under the root organization.
16. Select Create MAC Pool to create the MAC address pool.
17. Enter `MAC-Pool1-B` as the name of the MAC pool.
18. Optional: Enter a description for the MAC pool.
19. Select Sequential as the option for Assignment Order.
20. Click Next.
21. Click Add.

22. Specify a starting MAC address.



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0B:00 as our first MAC address.

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

24. Click OK.

25. Click Finish.

26. 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 Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. Right click IQN Pools.
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.

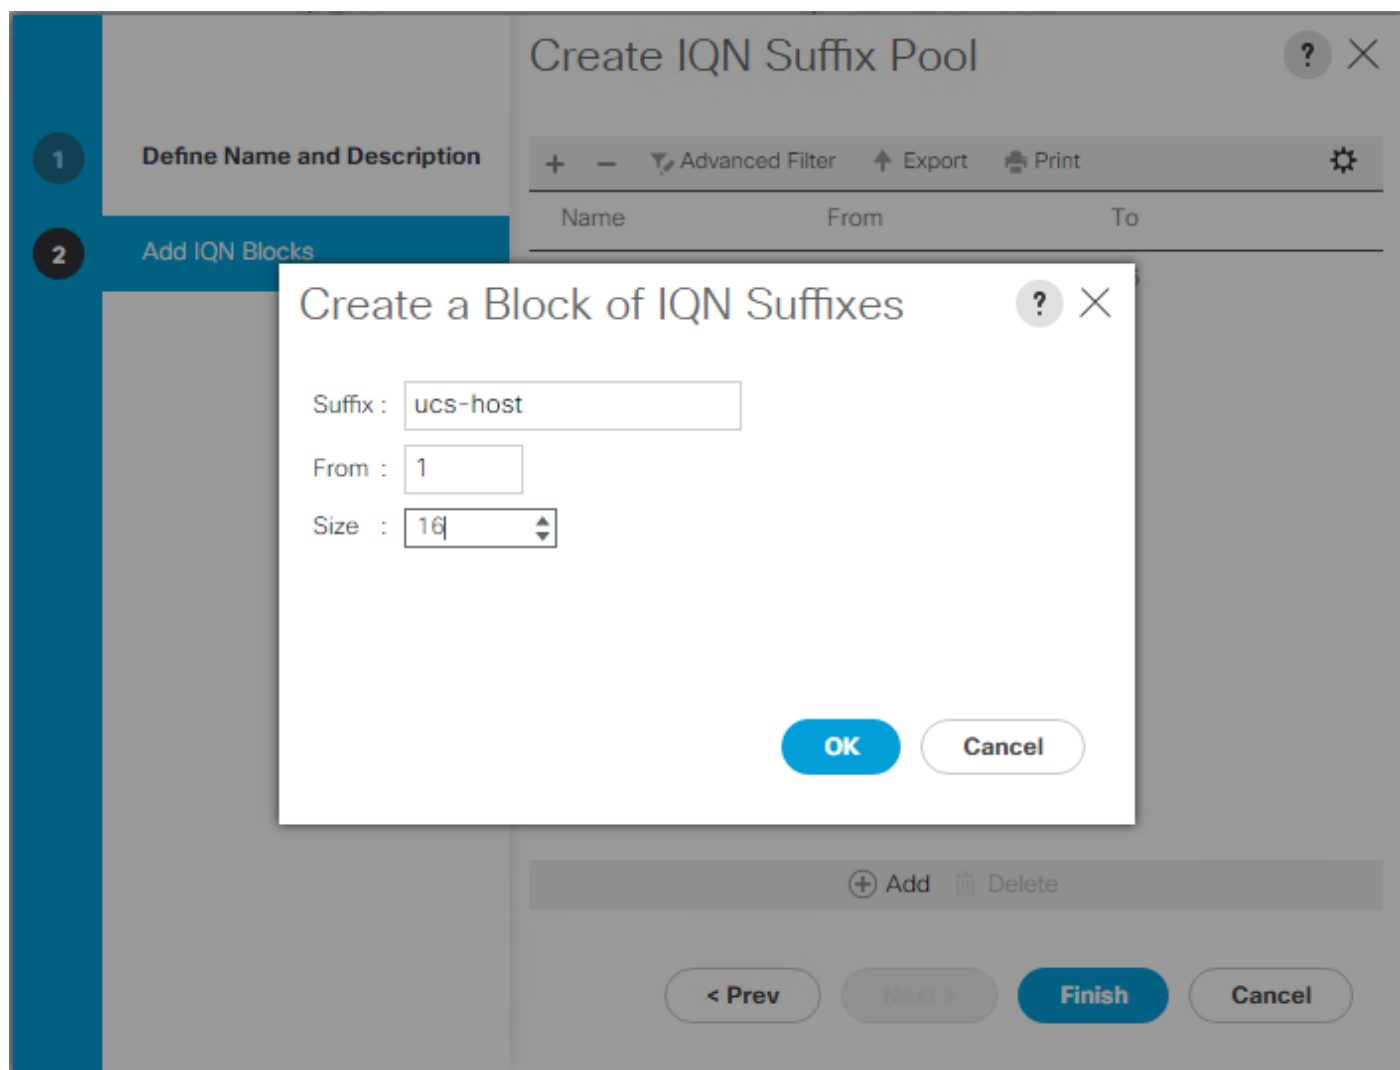


If multiple Cisco UCS domains are being used, a more specific IQN suffix may need to be used.

12. Enter 1 in the From field.

13. Specify the size of the IQN block sufficient to support the available server resources.

14. Click OK.



15. Click Finish.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.
3. Right-click IP Pools.
4. Select Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.

7. Select Sequential for the assignment order.
8. Click Next.
9. Click Add to add a block of IP address.
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 Next.
14. Click Finish.
15. Right-click IP Pools.
16. Select Create IP Pool.
17. Enter iSCSI-IP-Pool-B as the name of IP pool.
18. Optional: Enter a description for the IP pool.
19. Select Sequential for the assignment order.
20. Click Next.
21. Click Add to add a block of IP address.
22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses
23. Set the size to enough addresses to accommodate the servers
24. Click OK.
25. Click Next.
26. Click Finish.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.

4. Select Create UUID Suffix Pool.
5. Enter `UUID-Pool` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool

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



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

1. In Cisco UCS Manager, click Servers on the left.
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 LAN on the left.



In this procedure, five 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 ? X

VLAN Name/Prefix :

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

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
 Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs :

Sharing Type : None Primary Isolated Community

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 `IB-MGMT` as the name of the VLAN to be used for management traffic.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the In-Band management VLAN ID.
17. Keep the Sharing Type as None.
18. Click OK, and then click OK again.
19. Right-click VLANs.
20. Select Create VLANs.

21. Enter `Infra-NFS` as the name of the VLAN to be used for NFS.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the Infrastructure NFS VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK, and then click OK again.
26. Right-click VLANs.
27. Select Create VLANs.
28. Enter `vMotion` as the name of the VLAN to be used for vMotion.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the vMotion VLAN ID.
31. Keep the Sharing Type as None.
32. Click OK, and then click OK again.
33. Select Create VLANs.
16. Enter `VM-Traffic` as the name of the VLAN to be used for VM Traffic.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the VM-Traffic 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 `iSCSI-A-VLAN` as the name of the VLAN to be used for iSCSI-A.
41. Keep the Common/Global option selected for the scope of the VLAN.
42. Enter the iSCSI-A 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 `iSCSI-B-VLAN` as the name of the VLAN to be used for iSCSI-B.
48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the iSCSI-B VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK, and then click OK again.

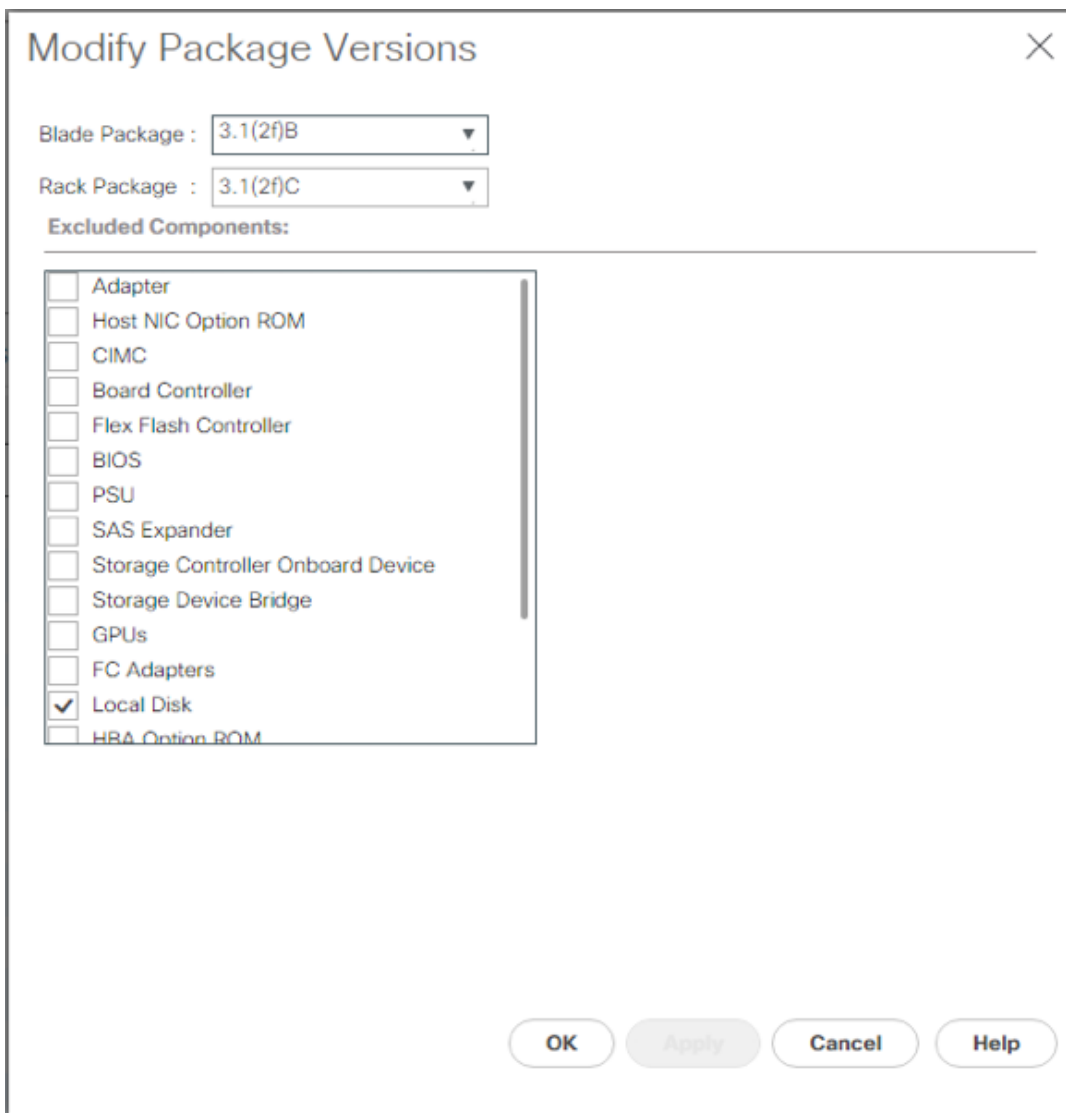
Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN IB-MGMT (113)	113	Lan	Ether	No	None
VLAN infra-NFS (3050)	3050	Lan	Ether	No	None
VLAN iSCSI-A-VLAN (3010)	3010	Lan	Ether	No	None
VLAN iSCSI-B-VLAN (3020)	3020	Lan	Ether	No	None
VLAN Native-VLAN (2)	2	Lan	Ether	No	None

Modify Default 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 Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.1(2f) for both the Blade and Rack Packages.



7. Click OK then OK again 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 LAN on the left.
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

LAN / LAN Cloud / QoS System Class

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 Servers on the left.
2. Select Policies > root.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter SAN-Boot as the local disk configuration policy name.
6. Change the mode to No Local Storage.
7. Click OK to create the local disk configuration policy.

Create Local Disk Configuration Policy ? ×

Name : SAN-Boot

Description :

Mode : No Local Storage ▾

FlexFlash

FlexFlash State : Disable Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately.
Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : Disable Enable

8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)

To create a network control policy that enables CDP and LLDP on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.

4. Select Create Network Control Policy.
5. Enter `Enable-CDP-LLDP` as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

Create Network Control Policy

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**

9. 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 Servers tab on the left.
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 ? X

Name :

Description :

Fan Speed Policy :

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.

Create Server Pool Qualification Policy (Optional)

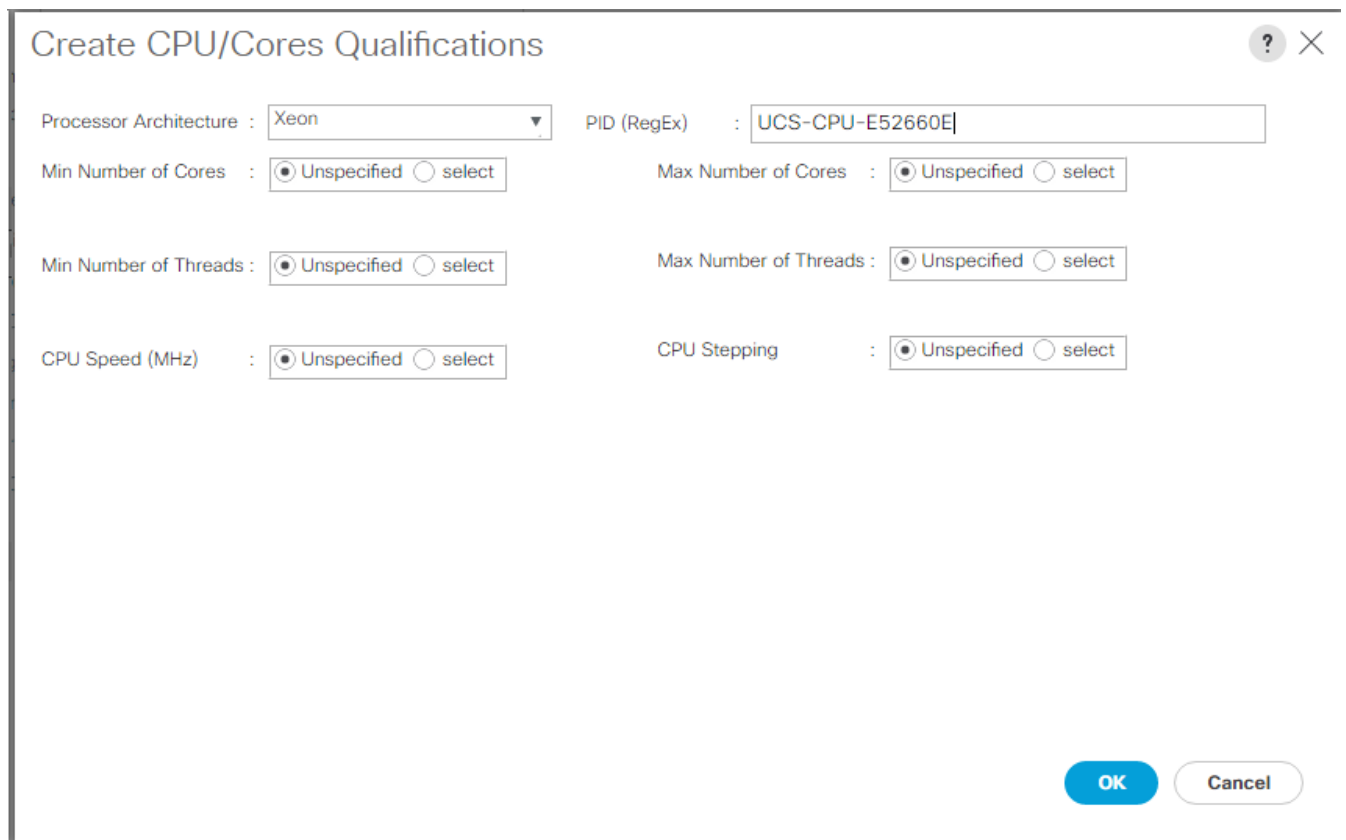
To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:



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

1. In Cisco UCS Manager, click Servers on the left.

2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCS-Broadwell.
6. Select Create CPU/Cores Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Enter UCS-CPU-E52660E as the PID.
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy then OK for the confirmation.



Create CPU/Cores Qualifications ? X

Processor Architecture : Xeon ▼ PID (RegEx) : UCS-CPU-E52660E

Min Number of Cores : Unspecified select Max Number of Cores : Unspecified select

Min Number of Threads : Unspecified select Max Number of Threads : Unspecified select

CPU Speed (MHz) : Unspecified select CPU Stepping : Unspecified select

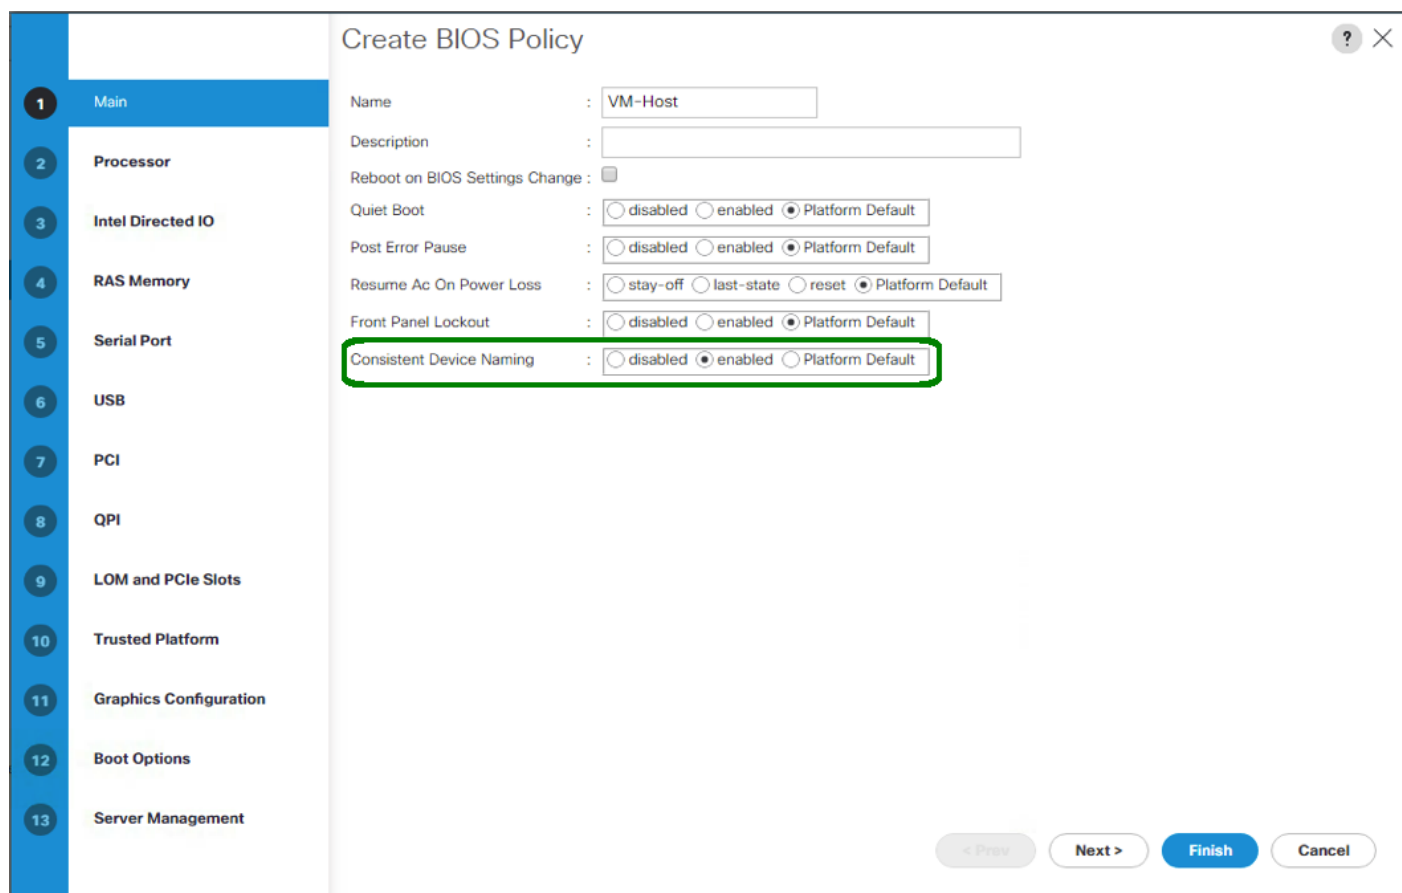
OK Cancel

Create Server BIOS Policy

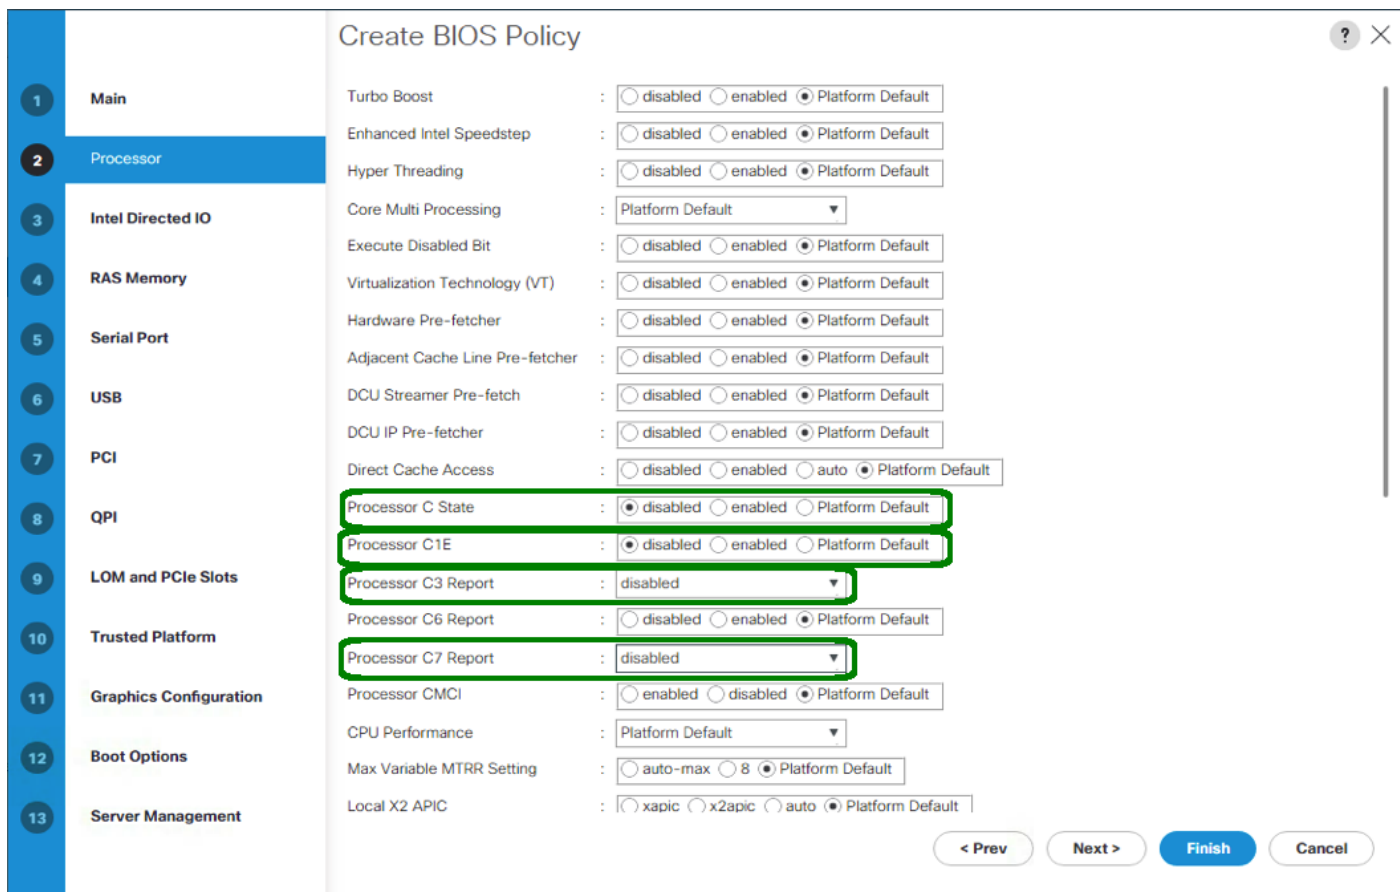
To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.

3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter VM-Host as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.



8. Click the Processor tab on the left.
9. Set the following within the Processor tab
 10. Processor C State -> disabled
 11. Processor C1E -> disabled
 12. Processor C3 Report -> disabled
 13. Processor C7 Report -> disabled

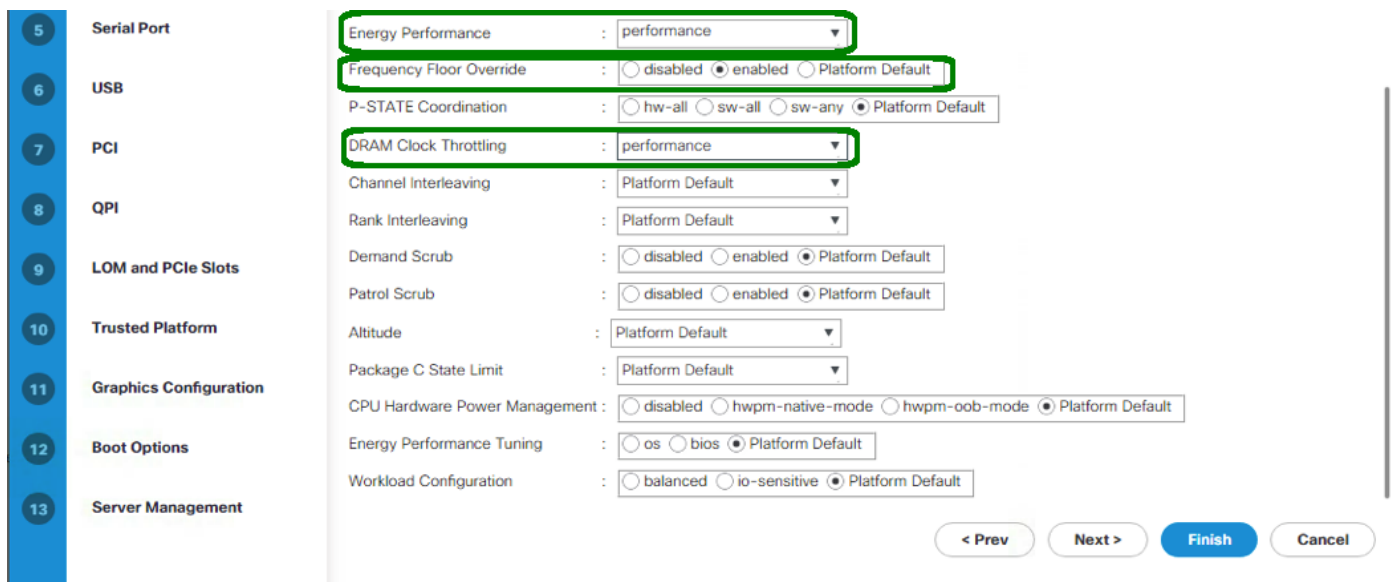


14. Scroll down to the remaining Processor options and select:

15. Energy Performance -> performance

16. Frequency Floor Override -> enabled

17. DRAM Clock Throttling -> performance



18. Click the RAS Memory option, and select:

19. LV DDR Mode -> performance-mode

The screenshot shows the 'Create BIOS Policy' configuration window. The left sidebar lists 13 categories, with 'RAS Memory' selected. The main configuration area shows the following settings:

- Memory RAS Config : Platform Default
- NUMA : disabled enabled Platform Default
- LV DDR Mode** : power-saving-mode performance-mode auto Platform Default
- DRAM Refresh Rate : Platform Default
- DDR3 Voltage Selection : ddr3-1500mv ddr3-1350mv Platform Default

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

20. Click Finish to create the BIOS policy.

21. Click OK.

Update the Default Maintenance Policy

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select **"On Next Boot"** to delegate maintenance windows to server administrators.

General	Events
Actions Delete Show Policy Usage Use Global	Properties Name : default Description : <input type="text"/> Owner : Local Soft Shutdown Timer : <input type="text" value="150 Secs"/> Reboot Policy : <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack <input type="radio"/> Timer Automatic <input checked="" type="checkbox"/> On Next Boot (Apply pending changes at next reboot.)

- Click Save Changes.
- 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 Infrastructure vNICs


- In Cisco UCS Manager, click LAN on the left.
- Select Policies > root.
- Right-click vNIC Templates.
- Select Create vNIC Template.
- Enter `Infra-A` as the vNIC template name.
- Keep Fabric A selected.
- Select the Enable Failover checkbox.



Selecting Failover is a critical step to improve link failover time by handling it at the hardware level, and to guard against any potential for NIC failure not being detected by the virtual switch.

- Select Primary Template for Redundancy Type.
- Leave the Peer Redundancy Template set to <not set>.
- Under Target, make sure that only the Adapter checkbox is selected.

11. Select Updating Template as the Template Type.
12. Under VLANs, select the checkboxes for IB-MGMT, Infra-NFS, vMotion VM-Traffic, and Native-VLAN VLANs.
13. Set Native-VLAN as the native VLAN.
14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-Pool-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.



18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template Infra-B:

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

4. Select Create vNIC Template
5. Enter `Infra-B` as the vNIC template name.
6. Select Fabric B.
7. Select the Enable Failover checkbox.
8. Set Redundancy Type to Secondary Template.
9. Select `Infra-A` for the Peer Redundancy Template.
10. In the MAC Pool list, select `MAC-Pool-B`. The MAC Pool is all that needs to be selected for the Secondary Template.
11. Click OK to create the vNIC template.
12. Click OK.

Create iSCSI vNICs

1. Select LAN 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. Select Fabric A. Do not select the Enable Failover checkbox.
7. Leave Redundancy Type set at No Redundancy.
8. Under Target, make sure that only the Adapter checkbox is selected.
9. Select Updating Template for Template Type.
10. Under VLANs, select only `iSCSI-A-VLAN`.
11. Select `iSCSI-A-VLAN` as the native VLAN.
12. Leave vNIC Name set for the CDN Source.
13. Under MTU, enter 9000.
14. From the MAC Pool list, select `MAC-Pool-A`.
15. From the Network Control Policy list, select `Enable-CDP-LLDP`.
16. Click OK to complete creating the vNIC template.

17. Click OK.
18. Select LAN on the left.
19. Select Policies > root.
20. Right-click vNIC Templates.
21. Select Create vNIC Template.
22. Enter `iSCSI-Template-B` as the vNIC template name.
23. Select Fabric B. Do not select the Enable Failover checkbox.
24. Leave Redundancy Type set at No Redundancy.
25. Under Target, make sure that only the Adapter checkbox is selected.
26. Select Updating Template for Template Type.
27. Under VLANs, select only `iSCSI-B-VLAN`.
28. Select `iSCSI-B-VLAN` as the native VLAN.
29. Leave vNIC Name set for the CDN Source.
30. Under MTU, enter 9000.
31. From the MAC Pool list, select `MAC-Pool-B`.
32. From the Network Control Policy list, select `Enable-CDP-LLDP`.
33. Click OK to complete creating the vNIC template.
34. Click OK.

Create LAN Connectivity Policy for iSCSI Boot

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

1. In Cisco UCS Manager, click LAN on the left.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter `iSCSI-Boot` as the name of the policy.
6. Click the upper Add button to add a vNIC.

7. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Infra-A.
10. In the Adapter Policy list, select VMWare.
11. Click OK to add this vNIC to the policy.

Create vNIC ? ×

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

Create Ethernet Adapter Policy

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

15. In the vNIC Template list, select Infra-B.
16. In the Adapter Policy list, select VMWare.
17. Click OK to add the vNIC to the policy.
18. Click the upper Add button to add a vNIC.
19. In the Create vNIC dialog box, enter 02-iSCSI-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select iSCSI-Template-A.
22. In the Adapter Policy list, select VMWare.
23. Click OK to add this vNIC to the policy.
24. Click the upper Add button to add a vNIC to the policy.
25. In the Create vNIC dialog box, enter 03-iSCSI-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select iSCSI-Template-B.
28. In the Adapter Policy list, select VMWare.
29. Click OK to add this vNIC to the policy.
30. Expand the Add iSCSI vNICs.
31. Select Add in the Add iSCSI vNICs section.
32. Set the name to iSCSI-A-vNIC.
33. Select the 02-iSCSI-A as Overlay vNIC.
34. Set the VLAN to iSCSI-A-VLAN (native).
35. Set the iSCSI Adapter Policy to default
36. Leave the MAC Address set to None.
37. Click OK.
38. Select Add in the Add iSCSI vNICs section.
39. Set the name to iSCSI-B-vNIC.
40. Select the 03-iSCSI-A as Overlay vNIC.

41. Set the VLAN to iSCSI-B-VLAN.
42. Set the iSCSI Adapter Policy to default.
43. Leave the MAC Address set to None.

Create LAN Connectivity Policy ? X

Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

🗑️ Delete ➕ Add ⓘ Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-B-vNIC	03-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-A-vNIC	02-iSCSI-A	default	Derived

➕ Add 🗑️ Delete ⓘ Modify

OK
Cancel

44. Click OK, then click OK again to create the LAN Connectivity Policy.

Create vMedia Policy for VMware ESXi 6.5a Install Boot

In the NetApp Data ONTAP setup steps an HTTP web server is required, which will be used for hosting NetApp Data ONTAP as well as VMware software. The vMedia Policy created here will map the VMware ESXi 6.5a ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, complete the following steps:

1. In Cisco UCS Manager, select Servers on the left.

2. Select Policies > root.
3. Right-click vMedia Policies.
4. Select Create vMedia Policy.
5. Name the policy ESXi-6.5a-HTTP.
6. **Enter “Mounts ISO for ESXi 6.5a” in the Description field.**
7. Click Add.
8. Name the mount ESXi-6.5a-HTTP.
9. Select the CDD Device Type.
10. Select the HTTP Protocol.
11. Enter the IP Address of the web server.



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

12. Enter VMware-VMvisor-Installer-201701001-4887370.x86_64.iso as the Remote File name.



This VMware ESXi 6.5a ISO can be downloaded from [VMware Downloads](#).

13. Enter the web server path to the ISO file in the Remote Path field.

Create vMedia Mount ? X

Name :

Description :

Device Type : CDD HDD

Protocol : NFS CIFS HTTP HTTPS

Hostname/IP Address :

Image Name Variable : None Service Profile Name

Remote File :

Remote Path :

Username :

Password :

14. Click OK to create the vMedia Mount.

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



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

Create iSCSI Boot Policy

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). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).



One boot policy is configured in this procedure. The policy configures the primary target to be iscsi_lif01a.

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

1. In Cisco UCS Manager, click Servers on the left.

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.



Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Add Remote CD/DVD`.
9. Expand the iSCSI vNICs drop-down menu 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 Media and select `Add CIMC Mounted CD/DVD`.
16. Click OK to create the policy.

Properties for: Boot Policy Boot-Fabric-A

General
Events
✕

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **Boot-Fabric-A**

Description :

Owner : **Local**

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Warning

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.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

+ - ⌵ Advanced Filter ⬆ Export 🖨 Print ⚙

Name	Or...	vNIC/vHBA/iS...	Type	WWN	LUN ...	Slot ...	Boot ...	Boot ...	Desc...
Remote CD/DVD	1								
▼ iSCSI	2								
iSCSI		iSCSI-A-vNIC	Primary						
iSCSI		iSCSI-B-vNIC	Secondary						
CIMC Mounted CD/DVD	3								

⬆ Move Up ⬇ Move Down 🗑 Delete

OK
Apply
Cancel
Help

Create Service Profile Templates

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 Servers on the left.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.

- Under UUID, select UUID_Pool as the UUID pool.

Create Service Profile Template

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

Name :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root**

The template will be created in the following organization. Its name must be unique within this organization.
Type : Initial Template Updating Template

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

UUID Assignment:

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

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

< Prev Next > **Finish** Cancel

- Click Next.

Configure Storage Provisioning

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

Configure Networking Options

- Keep the default setting for Dynamic vNIC Connection Policy.
- Select the **“Use Connectivity Policy”** option to configure the LAN connectivity.
- Select iSCSI-Boot from the LAN Connectivity Policy pull-down.
- Select IQN_Pool in Initiator Name Assignment.

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: ▼

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

LAN Connectivity Policy: ▼ [Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment: ▼

Initiator Name:

[Create IQN Suffix Pool](#)

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

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

5. Click Next.

Configure Storage Options

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

Configure Zoning Options

1. Click Next.

Configure vNIC/HBA Placement

1. In the “Select Placement” list, leave the placement policy as “Let System Perform Placement”.
2. Click Next.

Configure vMedia Policy

1. Do not select a vMedia Policy.
2. Click Next.

Configure Server Boot Order

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

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: `Boot-Fabric-A` [Create Boot Policy](#)

Name : **Boot-Fabric-A**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **No**
 Boot Mode : **Legacy**

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

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	WWN	LUN Na...	Slot Nu...	Boot N...	Boot Pa...	Descrip...
CIMC Mounted CD/D...	3								
▼ iSCSI	2								
iSCSI		iSCSI-A-vNIC	Primary						
iSCSI		iSCSI-B-vNIC	Second...						
Remote CD/DVD	1								

[Create iSCSI vNIC](#) [Set iSCSI Boot Parameters](#) [Set UEFI Boot Parameters](#)

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

2. In the Boot order, select `iSCSI-A-vNIC`.
3. Click `Set iSCSI Boot Parameters` button.
4. In the `Set iSCSI Boot Parameters` pop-up, leave `Authentication Profile` to `<not set>` unless you have independently created one appropriate to your environment.
5. Leave the `"Initiator Name Assignment"` dialog box `<not set>` to use the single `Service Profile Initiator Name` defined in the previous steps.
6. Set `iSCSI_IP_Pool_A` as the `"Initiator IP address Policy"`.
7. Select `iSCSI Static Target Interface` option.
8. Click `Add`.
9. Enter the `iSCSI Target Name`. To get the `iSCSI target name` of `Infra-SVM`, login into `storage cluster management interface` and run `"iscsi show" command`.

```
bb04-aff300::> iscsi show
```

Vserver	Target Name	Target Alias	Status Admin
Infra-SVM	iqn.1992-08.com.netapp:sn.b5acab9ef1c811e68d9d00a098a9fec2:vs.3	Infra-SVM	up

10. Enter the IP address of iscsi_lif_02a 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. Click Add.

13. Enter the iSCSI Target Name.

14. Enter the IP address of iscsi_lif_01a for the IPv4 Address field.

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority : **2**

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

15. Click OK to add the iSCSI static target.

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(12/16) ▼

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](#)

[Reset Initiator Address](#)

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

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pro..	iSCSI IPV4 Address	LUN Id
iqn.1992-08.c...	1	3260		192.168.10.62	0
iqn.1992-08.c...	2	3260		192.168.10.61	0

OK [Cancel](#)



The target IPs were put in with the storage Node 02 IP first and the storage Node 01 IP second. This is assuming the boot LUN is on Node 01. The host will boot using the path to Node 01 if the order in this procedure is used.

16. In the Boot order, select iSCSI-B-vNIC.
17. Click Set iSCSI Boot Parameters button.
18. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment
19. **Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps**
20. **Set iSCSI_IP_Pool_B as the “Initiator IP address Policy”.**
21. Select iSCSI Static Target Interface option.
22. Click Add.
23. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, login into storage cluster management interface and run “iscsi show” command”.

```
bb04-aff300::> iscsi show
```

Vserver	Target Name	Target Alias	Status Admin
Infra-SVM	iqn.1992-08.com.netapp:sn.b5acab9ef1c811e68d9d00a098a9fec2:vs.3	Infra-SVM	up

24. Enter the IP address of iscsi_lif_02b for the IPv4 Address field.

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority : **1**

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

25. Click OK to add the iSCSI static target.

26. Click Add.

27. Enter the iSCSI Target Name.

28. Enter the IP address of `iscsi_lif_01b` for the IPv4 Address field.

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority : **2**

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

29. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters ? X

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(12/16) ▼

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

Reset Initiator Address

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

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pro..	iSCSI IPV4 Address	LUN Id
iqn.1992-08.c...	1	3260		192.168.20.62	0
iqn.1992-08.c...	2	3260		192.168.20.61	0

+ Add
🗑 Delete
ℹ Info

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

OK
Cancel

30. Click Next.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

Create Service Profile Template

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

⊖ Maintenance Policy

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

Maintenance Policy: [Create Maintenance Policy](#)

Name : **default**
 Description :
 Soft Shutdown Timer : **150 Secs**
 Reboot Policy : **User Ack**


< Prev Next > **Finish** Cancel

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select Infra-Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select **“UCS-Broadwell”** for the **Server Pool Qualification**.
4. Expand Firmware Management at the bottom of the page and select the default policy.



5. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

Create Service Profile Template

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

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

BIOS Policy : VM-Host ▼

⊕ External IPMI Management Configuration

⊕ Management IP Address

⊕ Monitoring Configuration (Thresholds)

⊖ Power Control Policy Configuration

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

Power Control Policy : No-Power-Cap ▼ [Create Power Control Policy](#)

⊕ Scrub Policy

⊕ KVM Management Policy

< Prev Next > **Finish** Cancel

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

Create vMedia-Enabled Service Profile Template

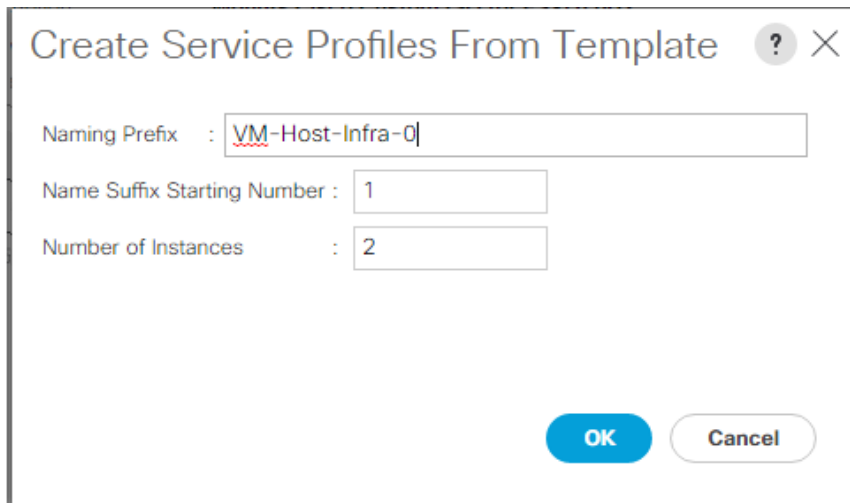
To create a service profile template with vMedia enabled, complete the following steps:

1. Connect to UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and select Create a Clone.
4. Name the clone VM-Host-Infra-iSCSI-A-vm.
5. Select the newly-created VM-Host-Infra-iSCSI-A-vm and select the vMedia Policy tab on the right.
6. Click Modify vMedia Policy.
7. Select the ESXi-6.5a-HTTP vMedia Policy and click OK.
8. Click OK to confirm.

Create Service Profiles

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

1. Connect to Cisco UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-iSCSI-A-VM.
3. Right-click VM-Host-Infra-iSCSI-A and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. **Enter 1 as “Name Suffix Starting Number.”**
6. **Enter 2 as the “Number of Instances.”**
7. Click OK to create the service profiles.



The screenshot shows a dialog box titled "Create Service Profiles From Template". It has a title bar with a question mark icon and a close button (X). The dialog contains three input fields:

- Naming Prefix :**
- Name Suffix Starting Number :**
- Number of Instances :**

At the bottom right of the dialog, there are two buttons: a blue "OK" button and a grey "Cancel" button.

8. Click OK in the confirmation message.
9. Once VMware ESXi 6.5a has been installed on the hosts, the host Service Profiles can be bound to the VM-Host-Infra-iSCSI-A Service Profile Template to remove the vMedia Mapping from the host.

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 server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 6 and Table 7 .

Table 6 iSCSI LIFs for iSCSI IQN

SVM	Target: IQN
Infra-SVM	



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

Table 7 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN	Variables
VM-Host-Infra-01		<vm-host-infra-01-iqn>
VM-Host-Infra-02		<vm-host-infra-02-iqn>



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

Storage Configuration – Boot LUNs and Igroups

ONTAP Boot Storage Setup

Create igroups

Create igroups by entering the following commands from the cluster management node SSH connection:

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol iscsi -ostype vmware -initiator
<vm-host-infra-01-iqn>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol iscsi -ostype vmware -initiator
<vm-host-infra-02-iqn>
igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol iscsi -ostype vmware -initiator <vm-
host-infra-01-iqn>, <vm-host-infra-02-iqn>
```



Use the values listed in **Error! Reference source not found.** and **Error! Reference source not found.** for the IQN information.

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

Map Boot LUNs to igroups

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

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


VMware vSphere 6.5a Setup

VMware ESXi 6.5a

This section provides detailed instructions for installing VMware ESXi 6.5a 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 6.5a

If the VMware ESXi custom image has not been downloaded, complete the following steps to complete the download.

1. Click the following link: [VMware vSphere Hypervisor \(ESXi\) 6.5a](#).
2. You will need a user id and password on vmware.com to download this software.
3. Download the .iso file.

Log in to Cisco UCS 6300/6200 Fabric Interconnect

Cisco UCS Manager

The Cisco UCS 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 Cisco 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. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
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 Servers on the left.
7. Select Servers > Service Profiles > root > `VM-Host-Infra-01`.
8. Right-click `VM-Host-Infra-01` and select KVM Console.
9. Follow the prompts to launch the Java-based KVM console.

10. Select Servers > Service Profiles > root > VM-Host-Infra-02.
11. Right-click VM-Host-Infra-02. and select KVM Console.
12. Follow the prompts to launch the Java-based KVM console.

Set Up VMware ESXi Installation

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



Skip this section if you are 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.

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. Boot the server by selecting Boot Server and clicking OK. Then click OK again.
2. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
5. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.

8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. 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.

10. After the installation is complete, press Enter to reboot the server.
11. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.


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 and VM-Host-Infra-02

To configure each 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 Troubleshooting Options and press Enter.
4. Select Enable ESXi Shell and press Enter.
5. Select Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Select the Configure Management Network option and press Enter.
8. Select Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field.



In lab testing, examples have been seen with the Cisco UCS C220M4 server and VIC 1385/1387 where the vmnic and device ordering do not match. If this is the case, use the Consistent Device Naming (CDN) to note which vmnics are mapped to which vNICs and adjust the upcoming procedure accordingly.

10. Press Enter.
11. Select the VLAN (Optional) option and press Enter.
12. Enter the <ib-mgmt-vlan-id> and press Enter.
13. Select IPv4 Configuration and press Enter.
14. Select the Set static IPv4 address and network configuration option by using the space bar.
15. Enter the IP address for managing the first ESXi host.
16. Enter the subnet mask for the first ESXi host.
17. Enter the default gateway for the first ESXi host.
18. Press Enter to accept the changes to the IP configuration.
19. Select the DNS Configuration option and press Enter.



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

20. Enter the IP address of the primary DNS server.
21. Optional: Enter the IP address of the secondary DNS server.
22. Enter the fully qualified domain name (FQDN) for the first ESXi host.
23. Press Enter to accept the changes to the DNS configuration.
24. Press Esc to exit the Configure Management Network menu.
25. Select Test Management Network to verify that the management network is set up correctly and press Enter.
26. Press Enter to run the test, press Enter again once the test has completed, review environment if there is a failure.
27. Re-select the Configure Management Network and press Enter.
28. Select the IPv6 Configuration option and press Enter.
29. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
30. Press Esc to exit the Configure Management Network submenu.
31. Press Y to confirm the changes and reboot the ESXi host.

Log in to VMware ESXi Hosts by Using VMware Host Client

ESXi Host VM-Host-Infra-01

To log in to the `VM-Host-Infra-01` ESXi host by using the VMware Host 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. Click Open the VMware Host Client.
3. Enter `root` for the user name.
4. Enter the root password.
5. Click Login to connect.
6. Repeat this process to log into `VM-Host-Infra-02` in a separate browser tab or window.

Set Up VMkernel Ports and Virtual Switch

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

To set up the VMkernel ports and the virtual switches on the ESXi hosts, complete the following steps:

1. From the Host Client, select Networking on the left.
2. In the center pane, select the Virtual switches tab.
3. Select vSwitch0.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Click Save.
7. Select Networking on the left.
8. In the center pane, select the Virtual switches tab.
9. Select iScsiBootvSwitch.
10. Select Edit settings.
11. Change the MTU to 9000.
12. Click Save.
13. Select the VMkernel NICs tab.
14. Select vmk1 iScsiBootPG.
15. Select Edit settings.
16. Change the MTU to 9000.
17. Expand IPv4 settings and change the IP address to an address outside of the UCS iSCSI-IP-Pool-A.



To avoid IP address conflicts if the Cisco UCS iSCSI IP Pool addresses should get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

18. Click Save.
19. Select the Virtual switches tab.
20. Select the Add standard virtual switch.
21. Provide a name of `iScsciBootvSwitch-B` for the vSwitch Name.
22. Set the MTU to 9000.
23. Select vmnic3 from the Uplink 1 pulldown options.
24. Click Add.

25. In the center pane, select the VMkernel NICs tab.
26. Select Add VMkernel NIC
27. Specify a New port group name of `iScsiBootPG-B`.
28. Select `iScsiBootvSwitch-B` for Virtual switch.
29. Set the MTU to 9000. Do not enter a VLAN ID.
30. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask within the Configuration.



To avoid IP address conflicts, if the Cisco UCS iSCSI IP Pool addresses should get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

31. Click Create.
32. On the left, select Networking, then select the Port groups tab.
33. In the center pane, right-click VM Network and select Remove.
34. Click Remove to complete removing the port group.
35. In the center pane, select Add port group.
36. Name the port group `IB-MGMT Network` and enter `<ib-mgmt-vlan-id>` in the VLAN ID field, and make sure Virtual switch `vSwitch0` is selected.
37. Click Add to finalize the edits for the IB-MGMT Network.
38. At the top, select the VMkernel NICs tab.
39. Click Add VMkernel NIC.
40. For New port group, enter `VMkernel-vMotion`.
41. For Virtual switch, select `vSwitch0` selected.
42. Enter `<vmotion-vlan-id>` for the VLAN ID.
43. Change the MTU to 9000.
44. Select Static IPv4 settings and expand IPv4 settings.
45. Enter the ESXi host vMotion IP address and netmask.
46. Select the vMotion stack TCP/IP stack.
47. Select vMotion under Services.

48. Click Create.
49. Click Add VMkernel NIC.
50. For New port group, enter VMkernel-Infra-NFS
51. For Virtual switch, select vSwitch0 selected.
52. Enter <infra-nfs-vlan-id> for the VLAN ID
53. Change the MTU to 9000.
54. Select Static IPv4 settings and expand IPv4 settings.
55. Enter the ESXi host Infrastructure NFS IP address and netmask.
56. Do not select any of the Services.
57. Click Create.
58. Select the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:

vSwitch0
 Type: Standard vSwitch
 Port groups: 4
 Uplinks: 1

⚠ This virtual switch has no uplink redundancy. You should add another uplink adapter. ⚙ Actions

vSwitch Details	
MTU	9000
Ports	6912 (6894 available)
Link discovery	Unknown
Attached VMs	0 (0 active)
Beacon interval	1

NIC teaming policy	
Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Rolling order	No

Security policy	
Allow promiscuous mode	No
Allow forged transmits	Yes
Allow MAC changes	Yes

vSwitch topology

- IB-MGMT Network (VLAN ID: 113)
- VMkernel-vMotion (VLAN ID: 3000)
 - VMkernel ports (1)
 - vmk4: 192.168.100.25
- VMkernel-Infra-NFS (VLAN ID: 3050)
 - VMkernel ports (1)
 - vmk3: 192.168.50.25
- Management Network (VLAN ID: 113)
 - VMkernel ports (1)
 - vmk0: 10.1.156.25

Physical adapters: vmnic0, 20000 Mbps, Full

59. Select the VMkernel NICs tab to confirm configured virtual adapters. The adapters listed should be similar to the following example:

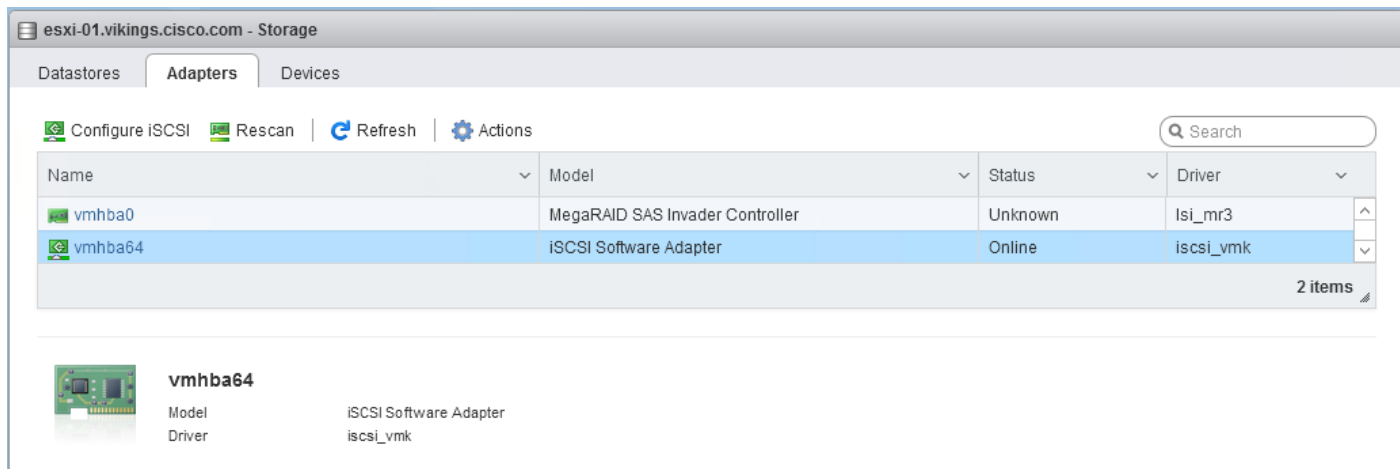


Setup iSCSI Multipathing

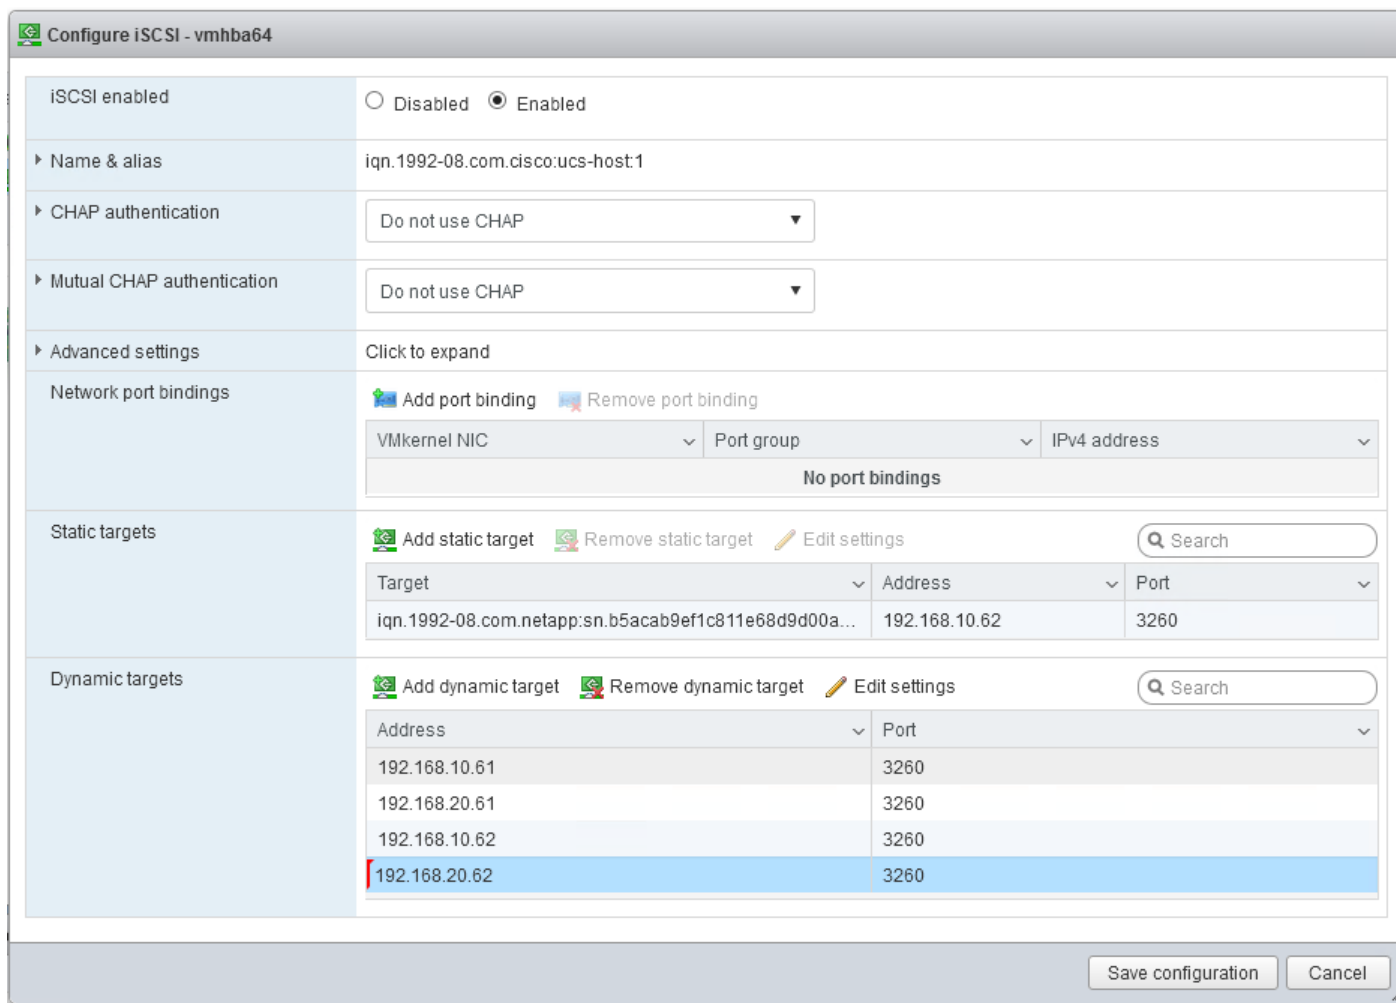
ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02


To setup the iSCSI multipathing on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each Host Client, select Storage on the left.
2. In the center pane, click Adapters.
3. Select the iSCSI software adapter and click Configure iSCSI.



4. Under Dynamic targets, click Add dynamic target.
5. Enter the IP Address of iSCSI_lif01a.
6. Repeat putting the ip address of iscsi_lif01b, iscsi_lif02a, iscsi_lif02b.
7. Click Save configuration.



 To get all the iscsi_lif IP address, login to NetApp storage cluster management interface and type “network interface show” command.

 The host will automatically rescan the storage adapter and the targets will be added to Static targets.

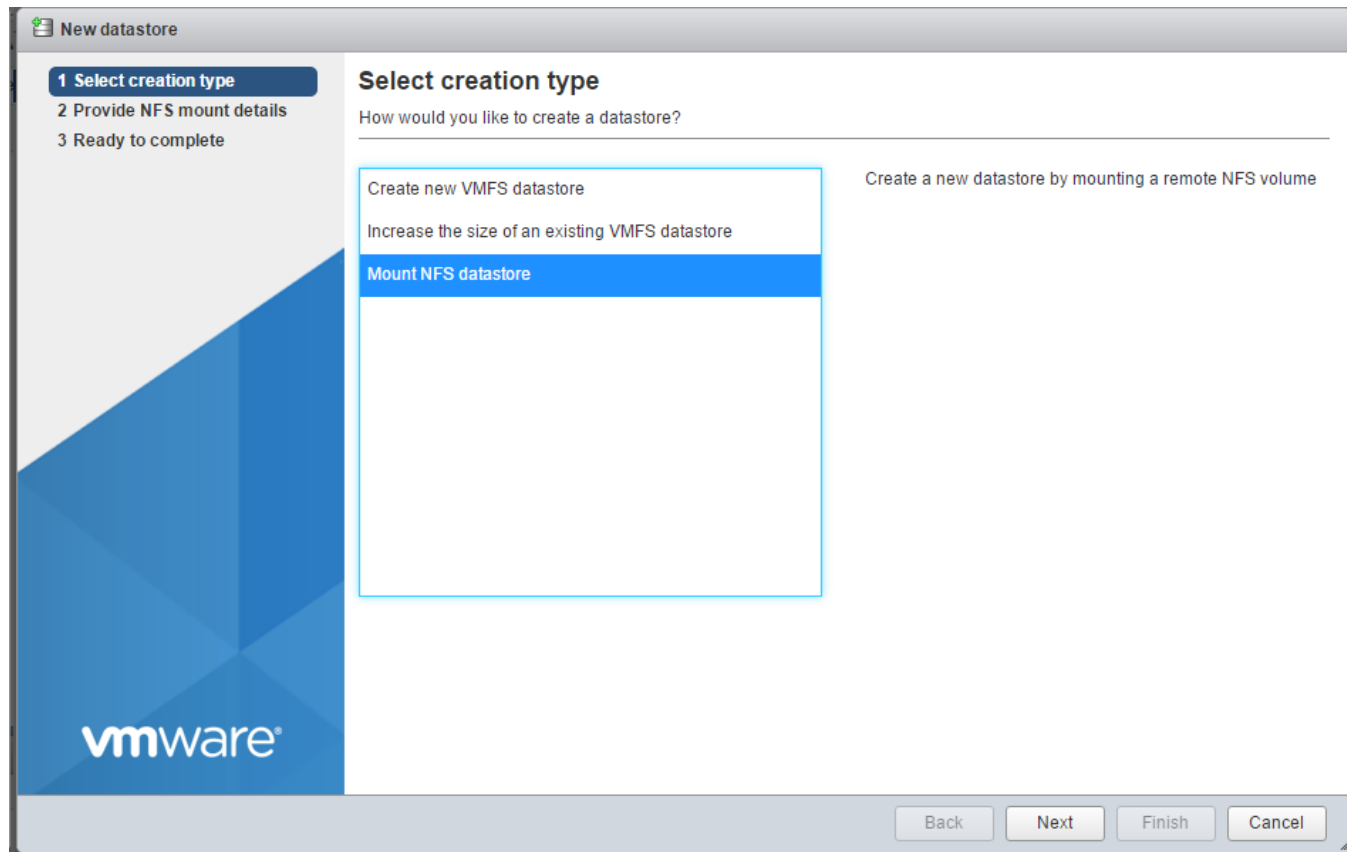
Mount Required Datastores

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

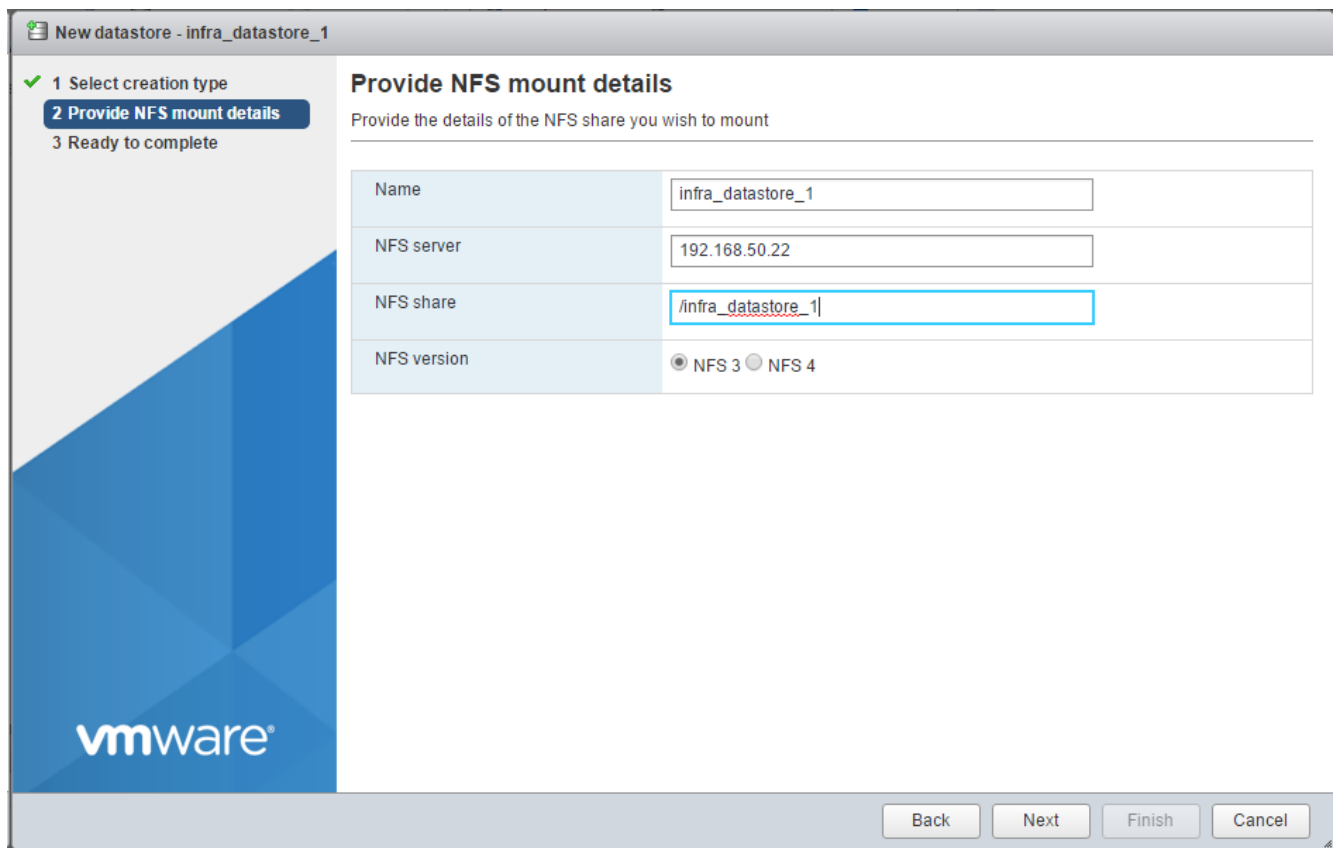
To mount the required datastores, complete the following steps on each ESXi host:

1. From the Host Client, select Storage on the left.
2. In the center pane, select Datastores.
3. In the center pane, select New Datastore to add a new datastore.

4. In the New datastore popup, select Mount NFS datastore and click Next.



5. Input `infra_datastore_1` for the datastore name. Input the IP address for the `nfs_lif01` LIF for the NFS server. Input `/infra_datastore_1` for the NFS share. Leave the NFS version set at NFS 3. Click Next.



6. Click Finish. The datastore should now appear in the datastore list.
7. In the center pane, select New Datastore to add a new datastore.
8. In the New datastore popup, select Mount NFS datastore and click Next.
9. Input infra_datastore_2 for the datastore name. Input the IP address for the `nfs_lif02` LIF for the NFS server. Input /infra_datastore_2 for the NFS share. Leave the NFS version set at NFS 3. Click Next.
10. Click Finish. The datastore should now appear in the datastore list.

The screenshot shows the 'Storage' view in vSphere for host 'esxi-01.vikings.cisco.com'. The 'Datastores' tab is selected, and a table lists the available datastores:

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provision...	Access
datastore1	Non-SSD	7.5 GB	3.95 GB	3.55 GB	VMFS6	Supported	Single
infra_datastore_1	Unknown	500 GB	37.19 GB	462.81 GB	NFS	Supported	Single
infra_datastore_2	Unknown	500 GB	60.79 GB	439.21 GB	NFS	Supported	Single

11. Mount both datastores on both ESXi hosts.

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 Host Client, select Manage on the left.
2. In the center pane, select the Time & Date tab.
3. Click Edit settings.
4. Make sure Use Network Time Protocol (enable NTP client) is selected.
5. Use the pulldown to select Start and stop with host.
6. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.

Edit time configuration

Specify how the date and time of this host should be set.

Manually configure the date and time on this host

10/13/2016 4:09 PM

Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 10.1.156.4,10.1.156.5

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

Save Cancel

7. Click Save to save the configuration changes.
8. Select Actions > NTP service > Start.
9. Verify that NTP service is now running and the clock is now set to approximately the correct time.



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

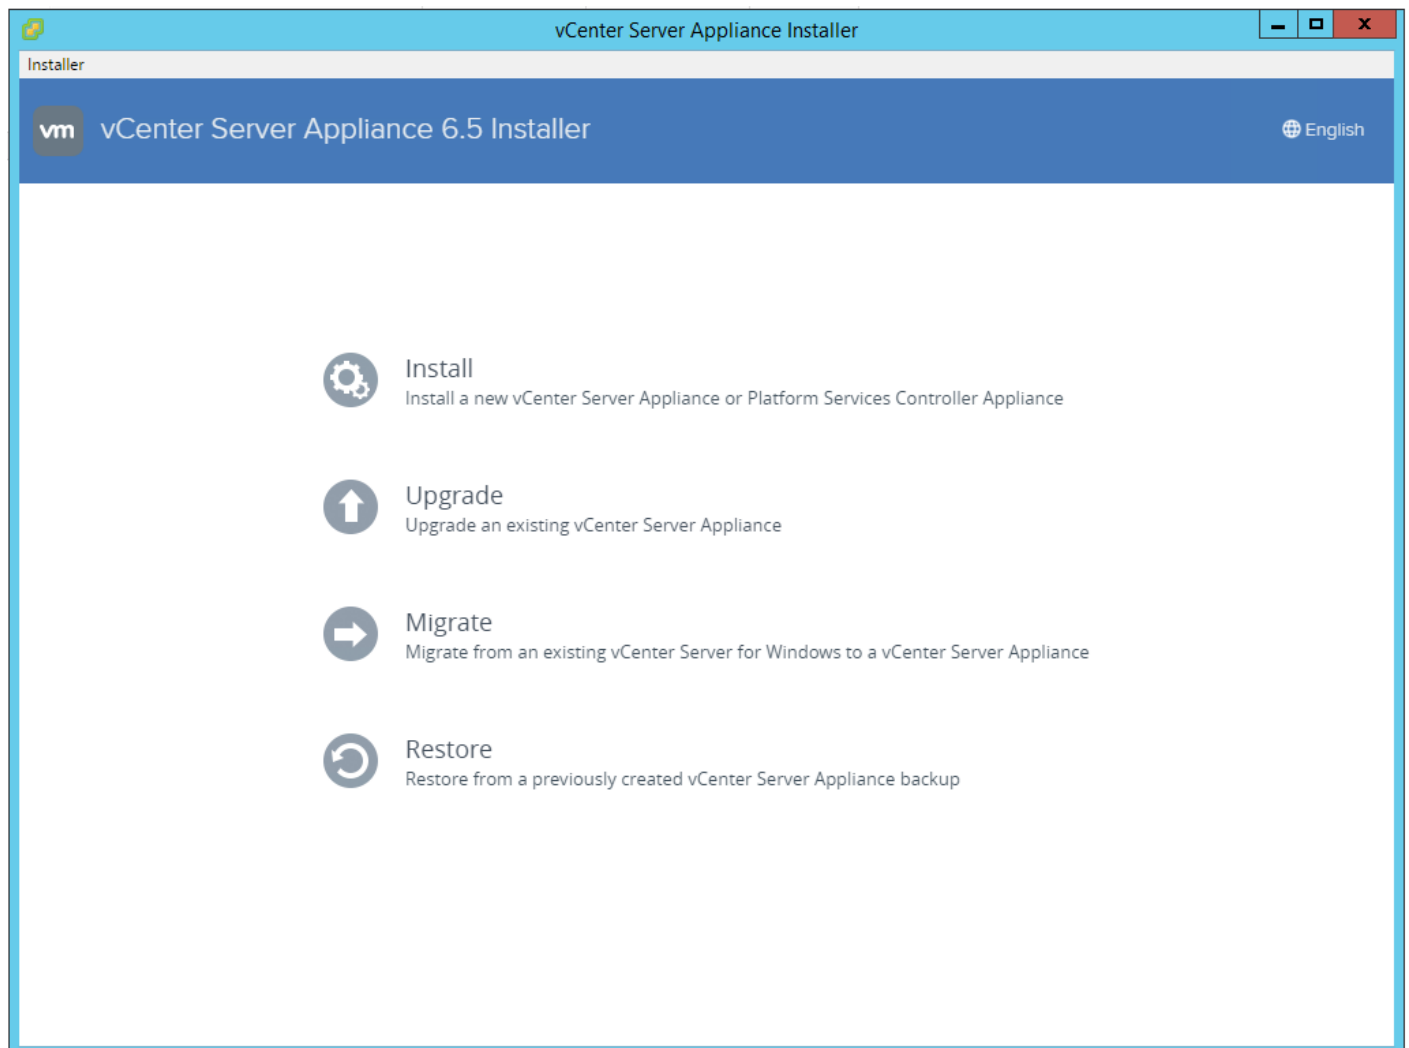
VMware vCenter 6.5a

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.5a Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

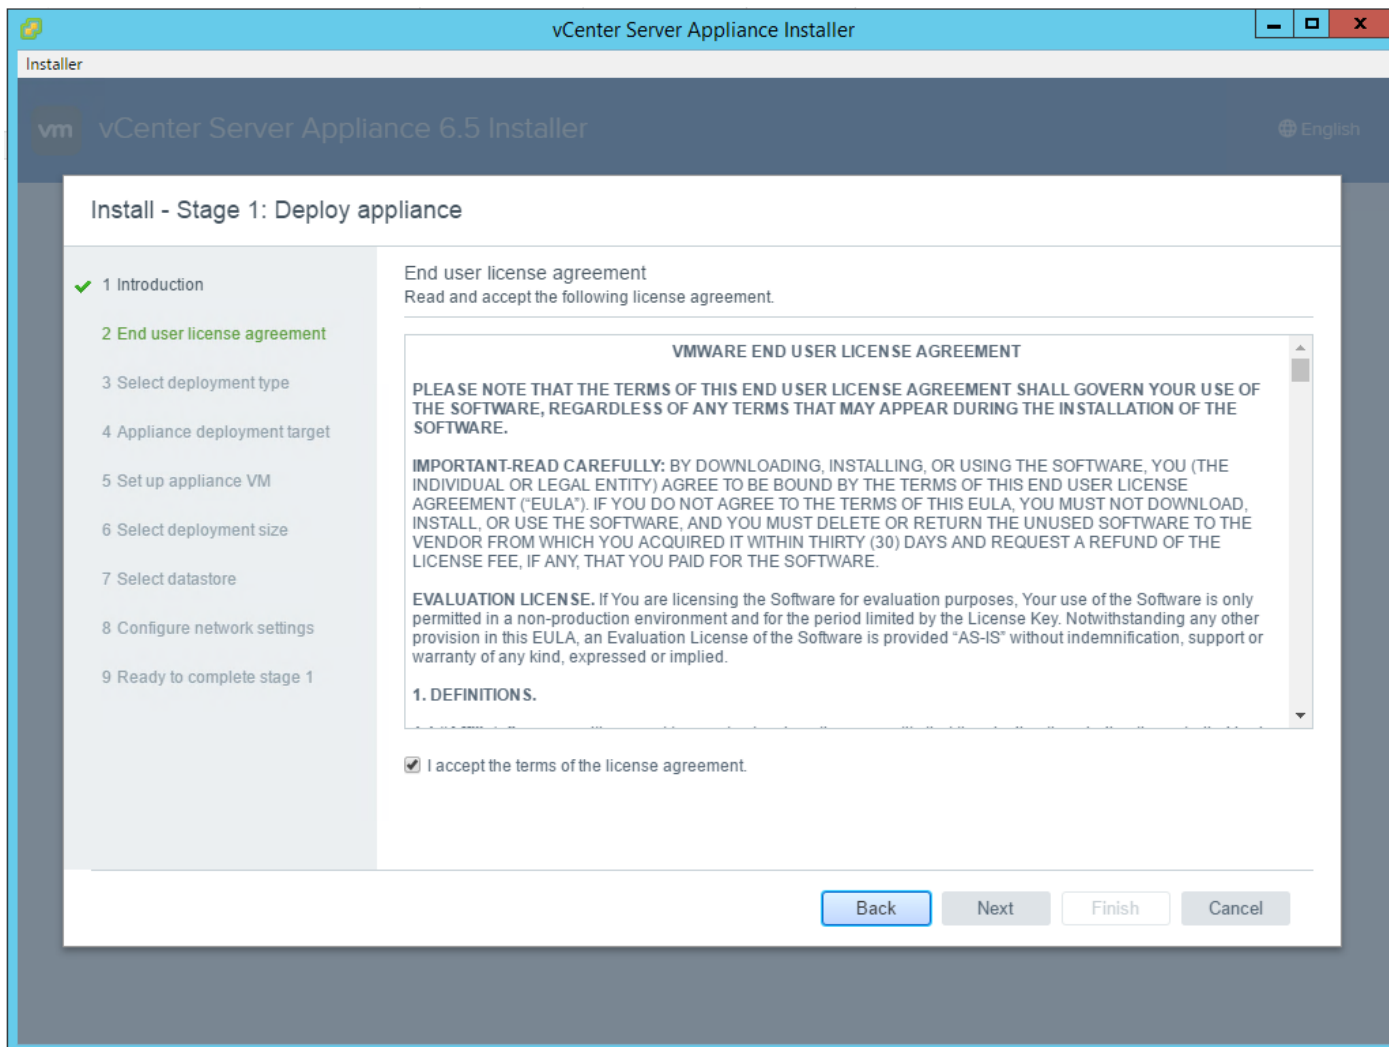
Building the VMware vCenter Server Appliance

The VCSA deployment consists of 2 stages: install and configuration. To build the VMware vCenter virtual machine, complete the following steps:

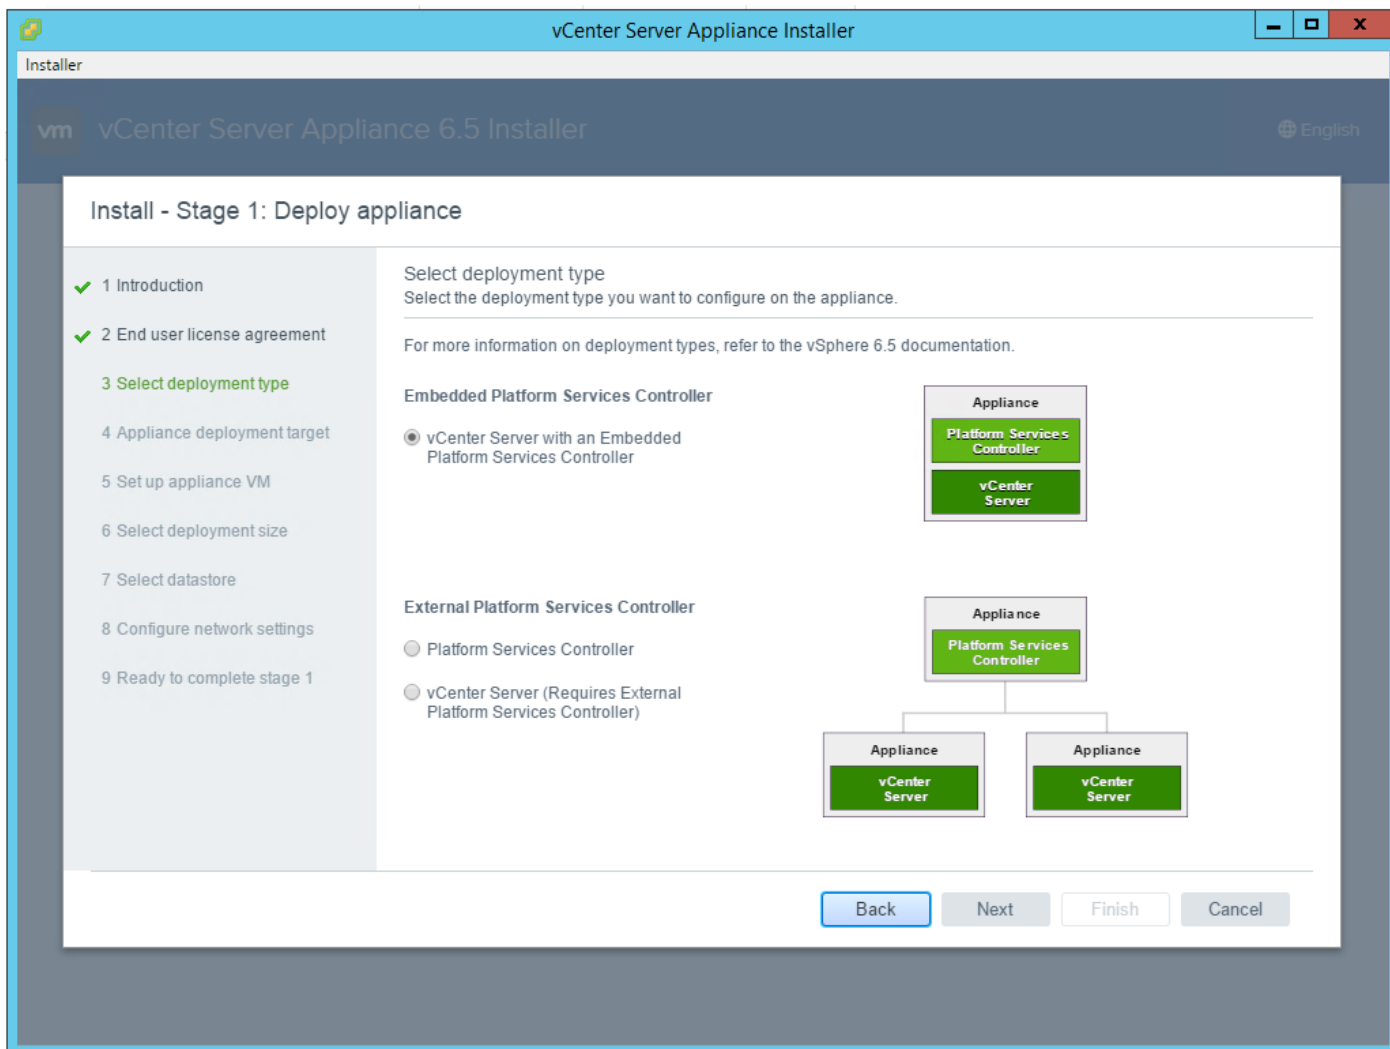
1. Locate and copy the VMware-VCSA-all-6.5.0-4944578.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 6.5 vCenter Server Appliance.
2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012).
3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appear.



4. Click Install to start the vCenter Server Appliance deployment wizard.
5. Click Next in the Introduction section.
6. Read and accept the license agreement and click Next.



7. In the “Select deployment type” section, select Embedded Platform Services Controller.



8. In the “Appliance deployment target”, enter the ESXi host name or IP address, User name and Password.

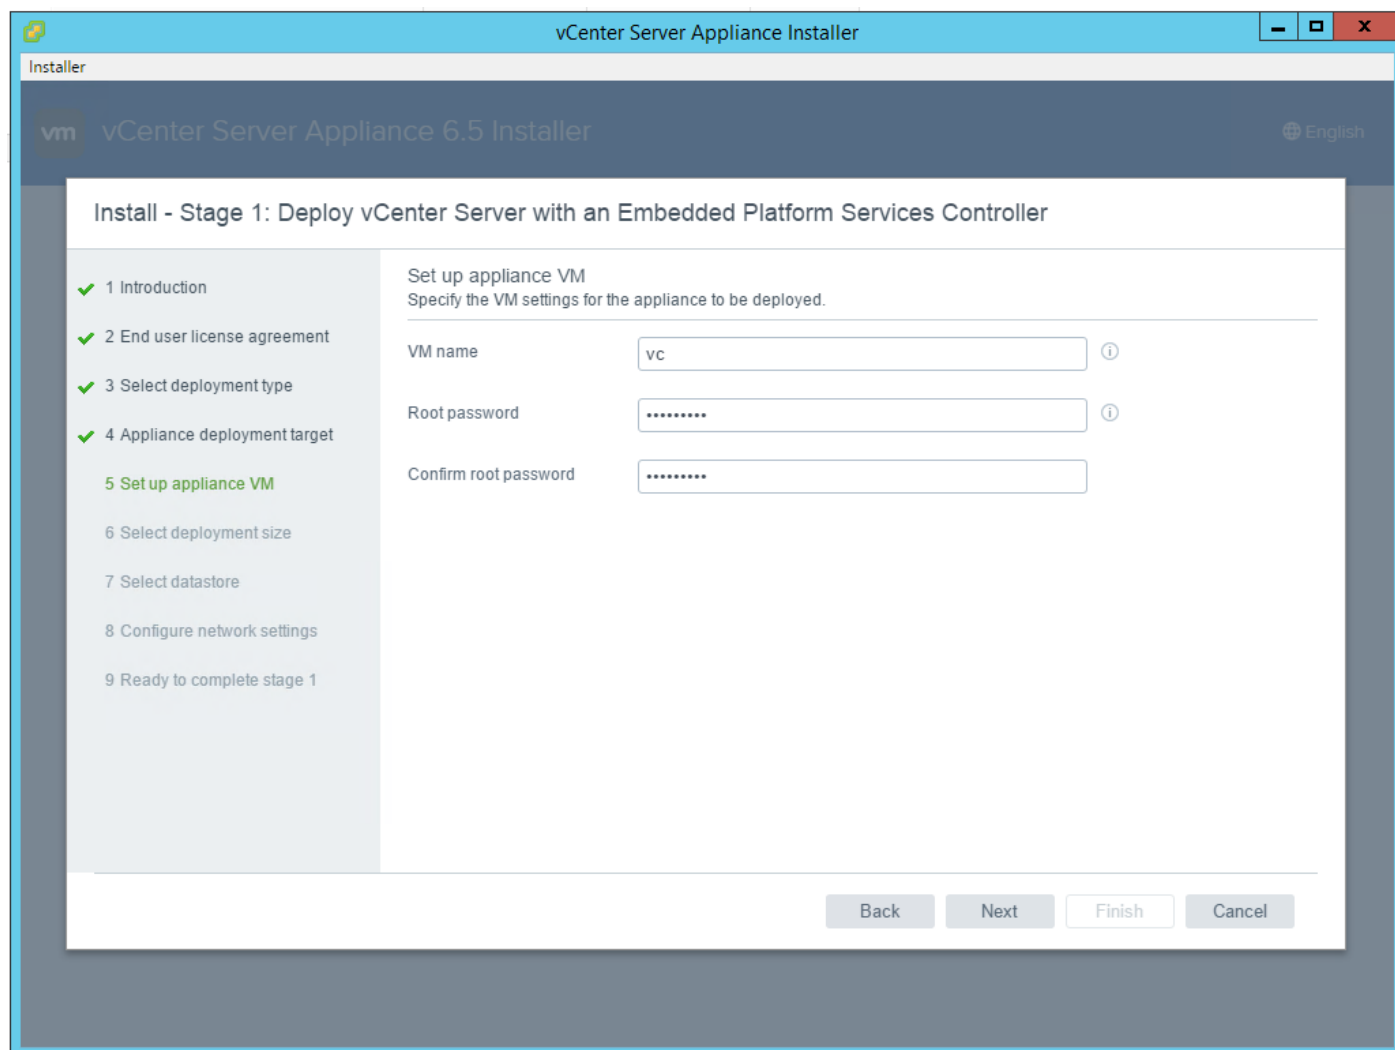
The screenshot shows the 'vCenter Server Appliance Installer' window. The title bar reads 'vCenter Server Appliance Installer'. The main window has a dark blue header with the VMware logo and 'vCenter Server Appliance 6.5 Installer' on the left, and 'English' on the right. The main content area is titled 'Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller'. On the left, a progress list shows steps 1 through 9, with step 4 'Appliance deployment target' highlighted in green. The main area contains the following configuration fields:

- Appliance deployment target**: Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.
- ESXi host or vCenter Server name**: ⓘ
- HTTPS port**:
- User name**: ⓘ
- Password**:

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

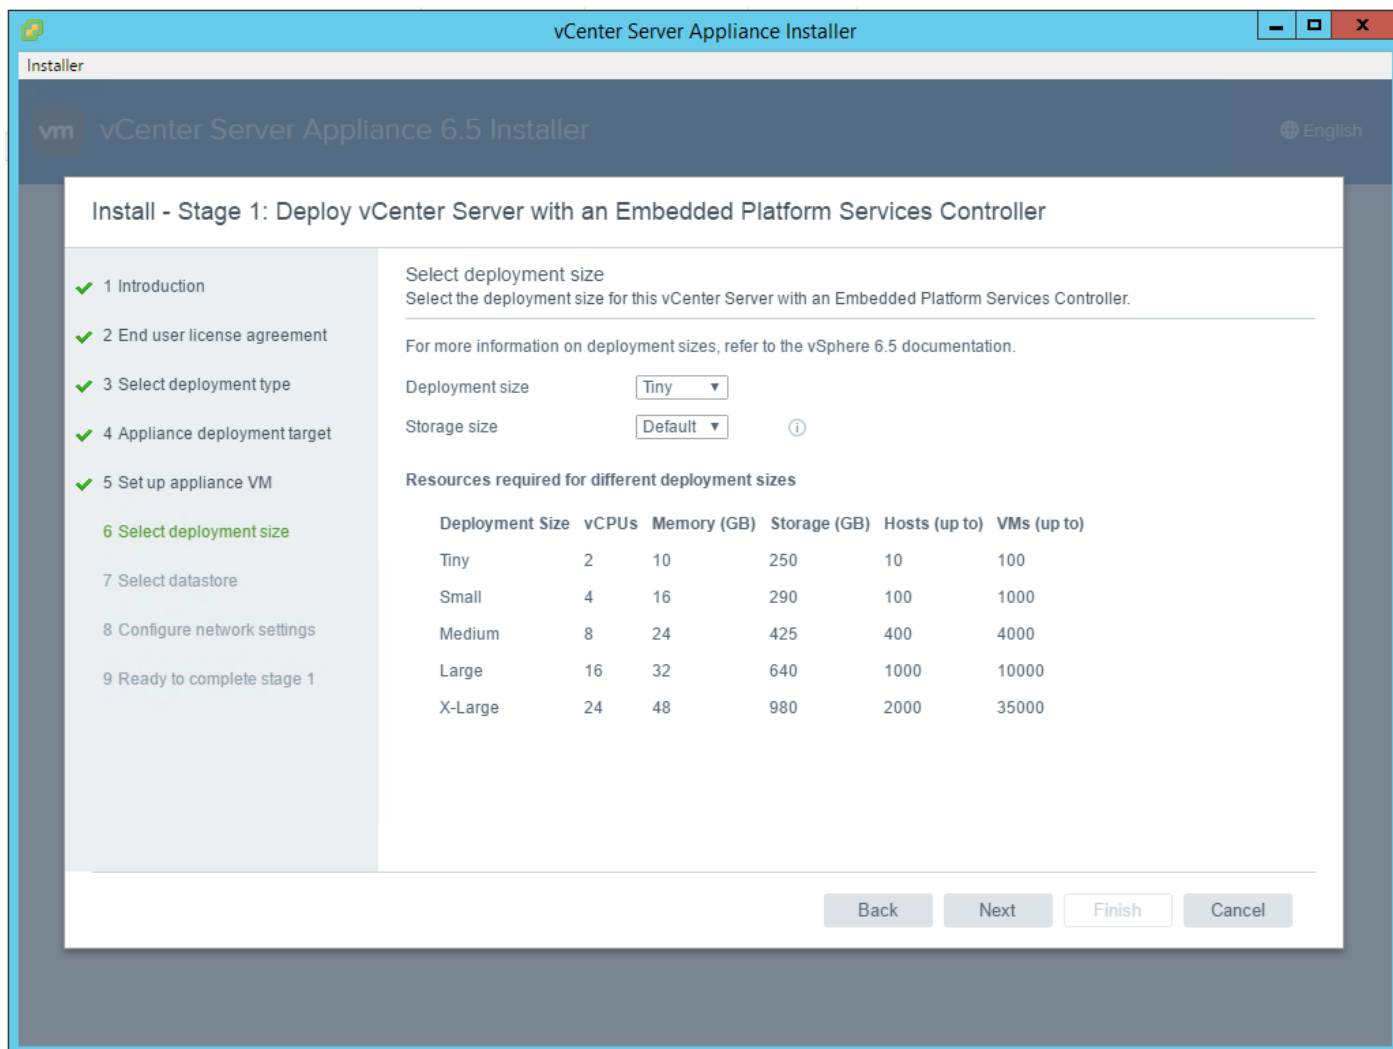
9. Click Yes to accept the certificate.

10. Enter the Appliance name and password details in the “Set up appliance VM” section. Click Next.




11. In the “Select deployment size” section, Select the deployment size and Storage size. For example, “Tiny.”

12. Click Next.



13. Select the infra_datastore_1. Click Next.



14. In the “Network Settings” section, configure the following settings:

- Choose a Network: IB-MGMT Network
- IP version: IPV4
- IP assignment: static
- System name: <vcenter-fqdn>
- IP address: <vcenter-ip>
- Subnet mask or prefix length: <vcenter-subnet-mask>
- Default gateway: <vcenter-gateway>
- DNS Servers: <dns-server>

The screenshot shows the 'vCenter Server Appliance Installer' window. The title bar reads 'vCenter Server Appliance Installer'. The main window has a dark blue header with 'vm vCenter Server Appliance 6.5 Installer' and 'English' on the right. The main content area is titled 'Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller'. On the left, a progress list shows steps 1 through 9, with step 8 'Configure network settings' highlighted in green. The main area contains the following configuration fields:

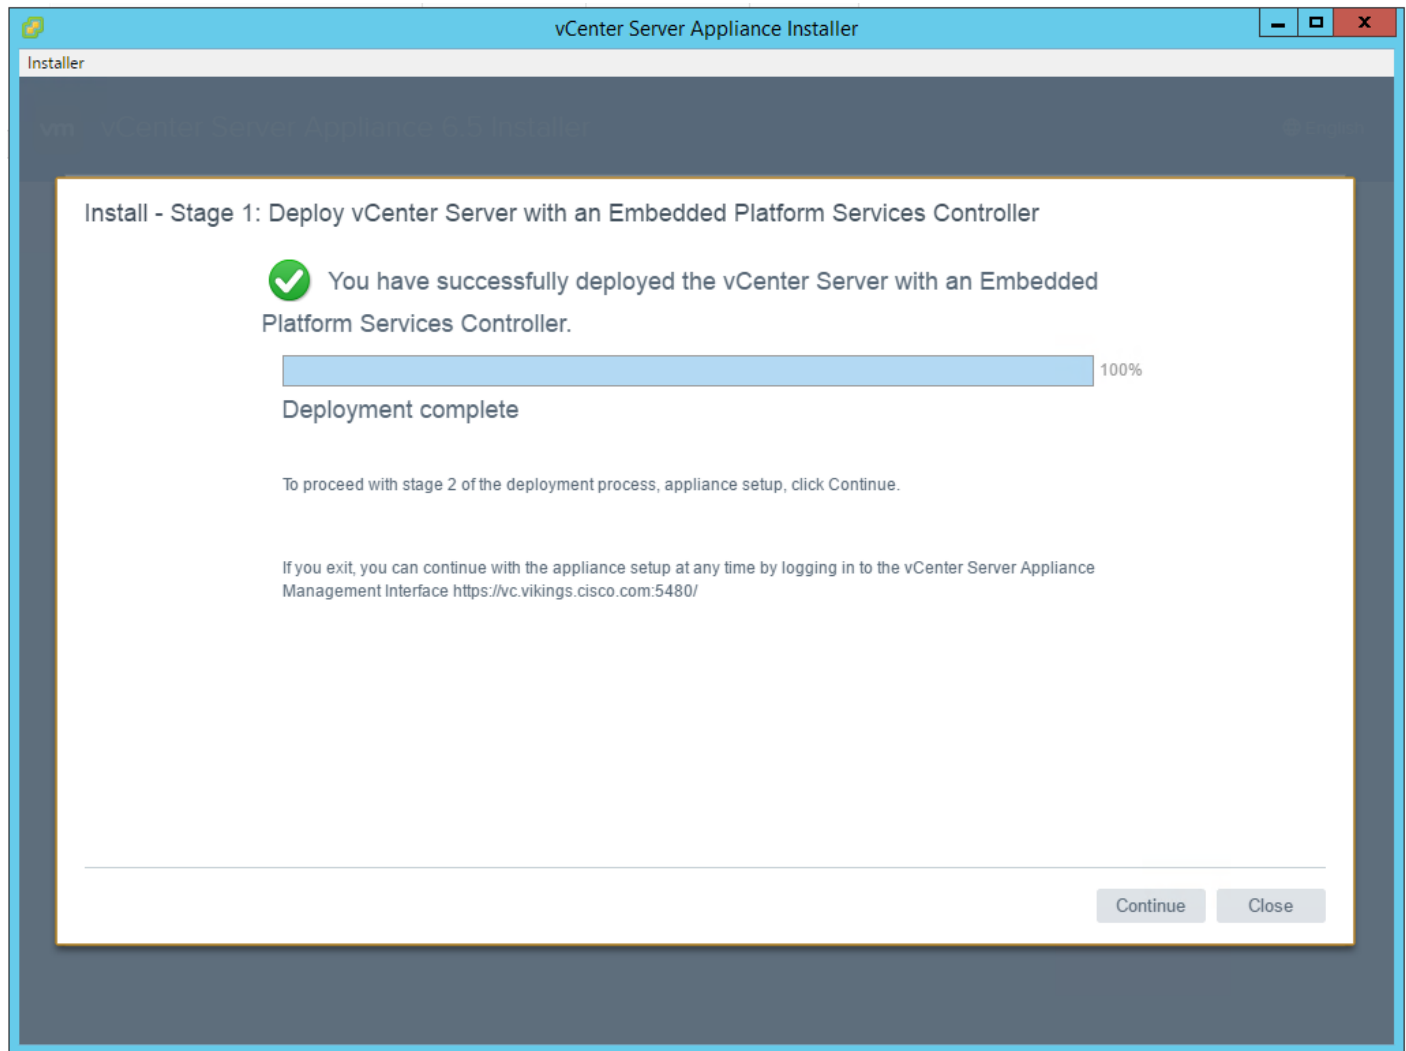
Field	Value
Network	IB-MGMT-Network
IP version	IPv4
IP assignment	static
System name	vc.vikings.cisco.com
IP address	10.1.156.100
Subnet mask or prefix length	255.255.255.0
Default gateway	10.1.156.1
DNS servers	192.168.156.9

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

15. Click Next.

16. Review all values and click Finish to complete the installation.

17. The vCenter appliance installation will take a few minutes to complete.




18. Click Continue to proceed with stage 2 configuration.

19. Click Next.


20. In the Appliance Configuration, configure the below settings:

- a. Time Synchronization Mode: Synchronize time with NTP servers.
- b. NTP Servers: <ntp_server_ip>
- c. SSH access: Enabled.



21. Click Next.

22. Complete the SSO configuration as show below.



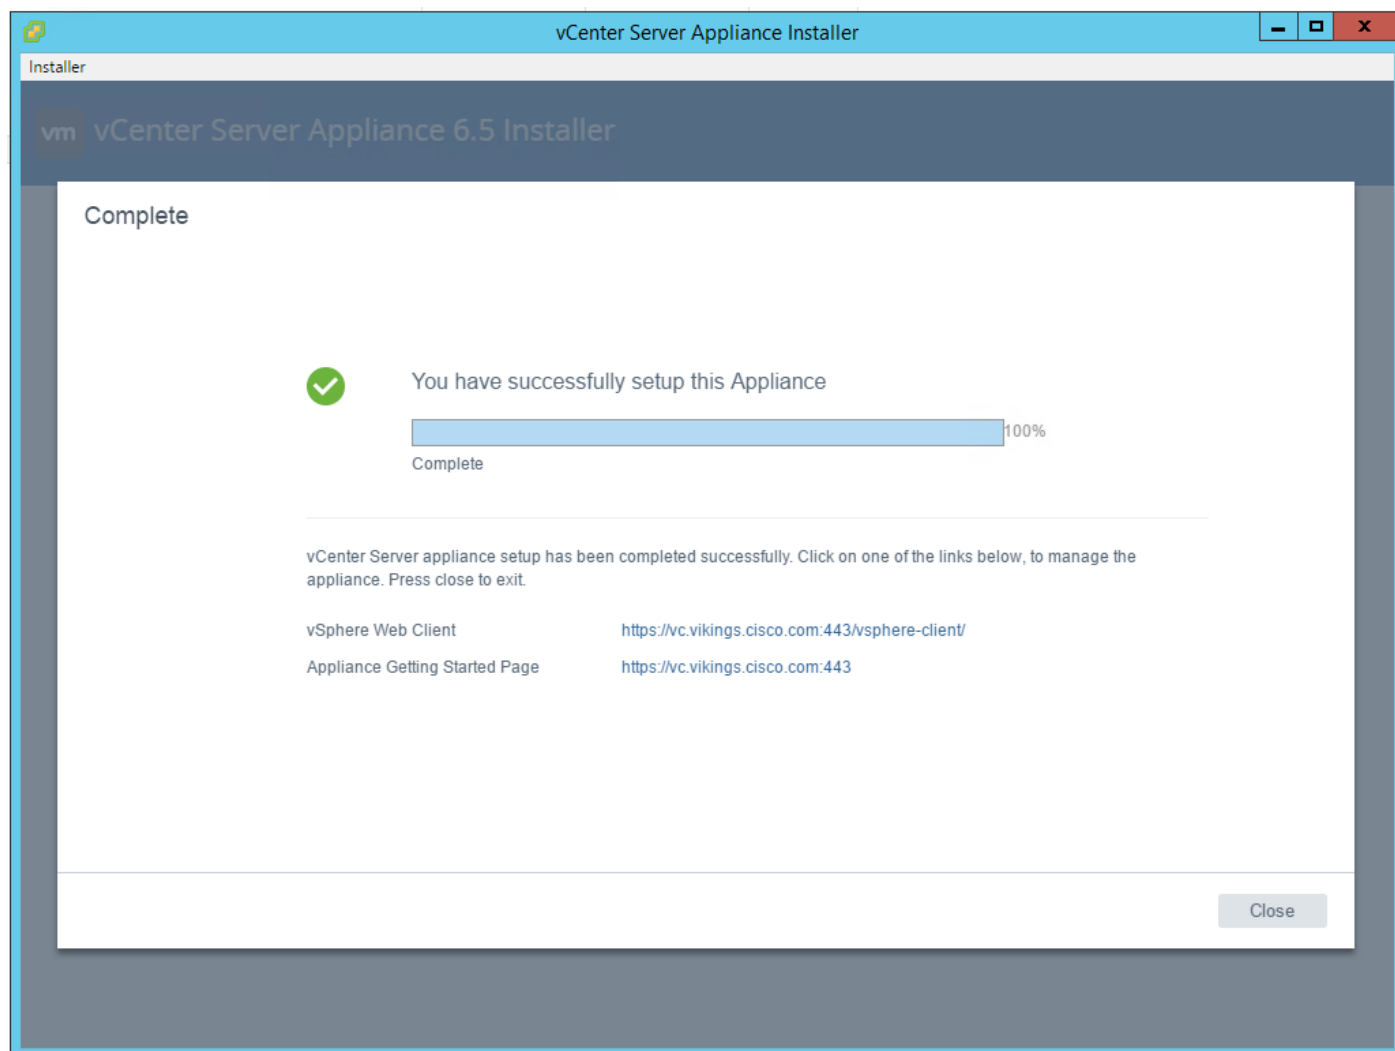
23. Click Next.

24. If needed, select Join the VMware's Customer Experience Improvement Program (CEIP).

25. Click Next.

26. Review the configuration and click Finish.

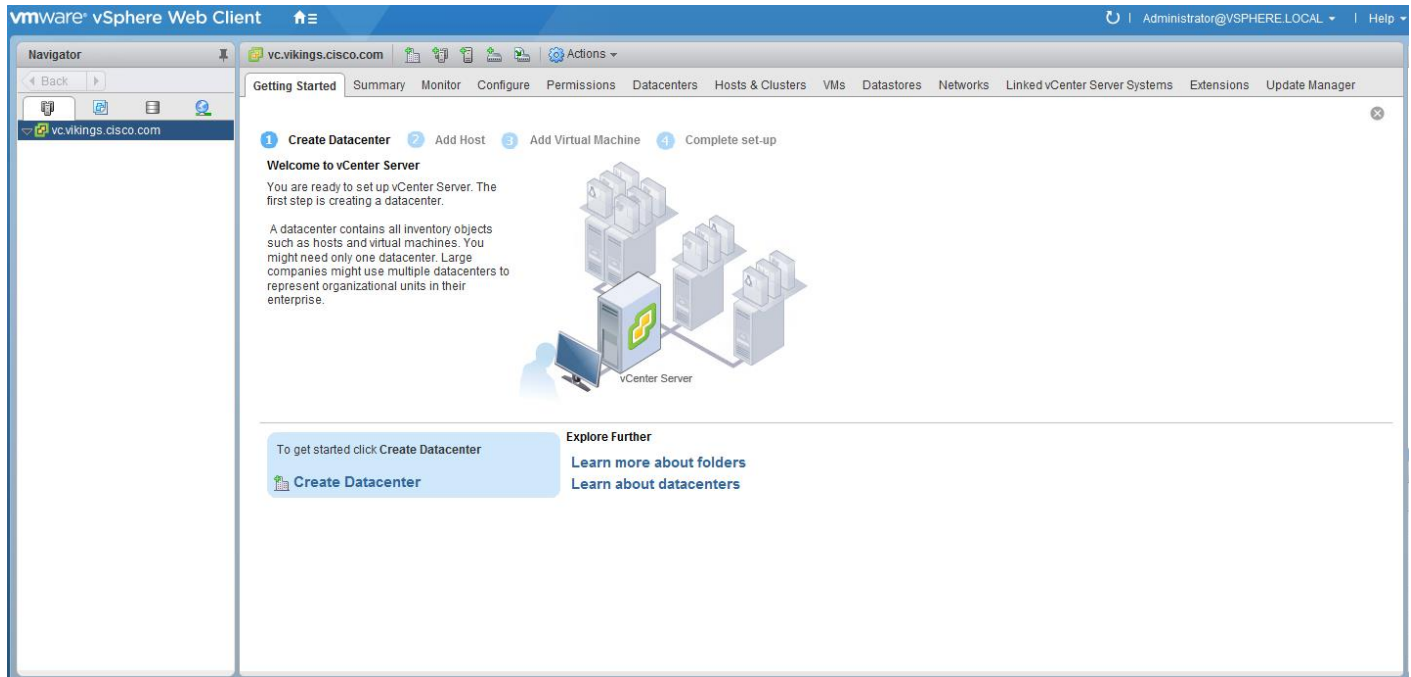
27. Click OK.



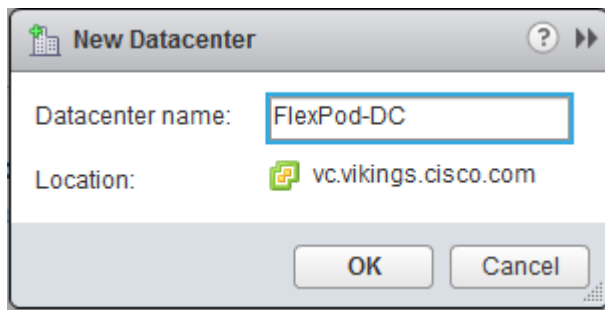
Setting Up VMware vCenter Server

To set up the VMware vCenter Server, complete the following steps:

1. Using a web browser, navigate to <https://<vcenter-ip>/vsphere-client>
2. Click Download Enhanced Authentication Plugin. Install the same by double clicking the downloaded file.
3. Log in using the Single Sign-On username and password created during the vCenter installation.



4. Click "Create Datacenter" in the center pane.
5. Type "FlexPod-DC" in the Datacenter name field.
6. Click OK.

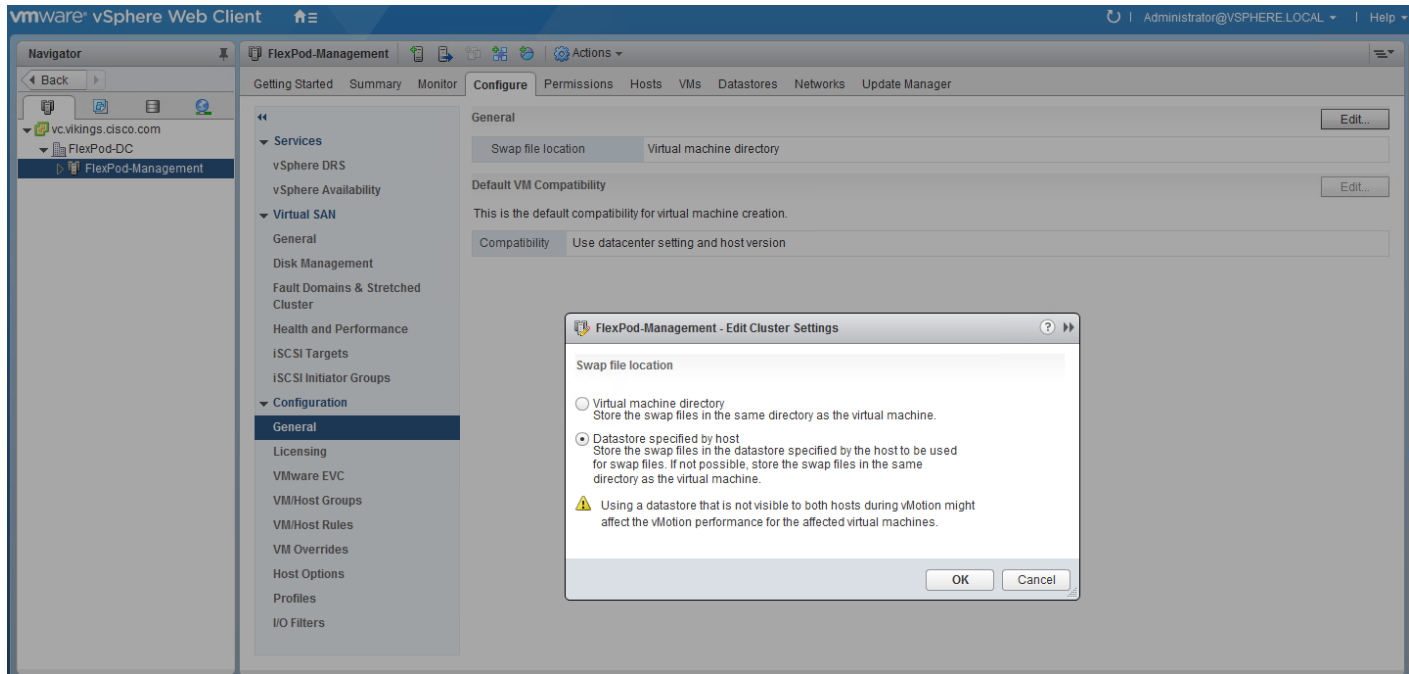


7. Right-click the data center FlexPod-DC in the list in the center pane. Click New Cluster.
8. Name the cluster FlexPod-Management.
9. Check the box to turn on DRS. Leave the default values.
10. Check the box to turn on vSphere HA. Leave the default values.

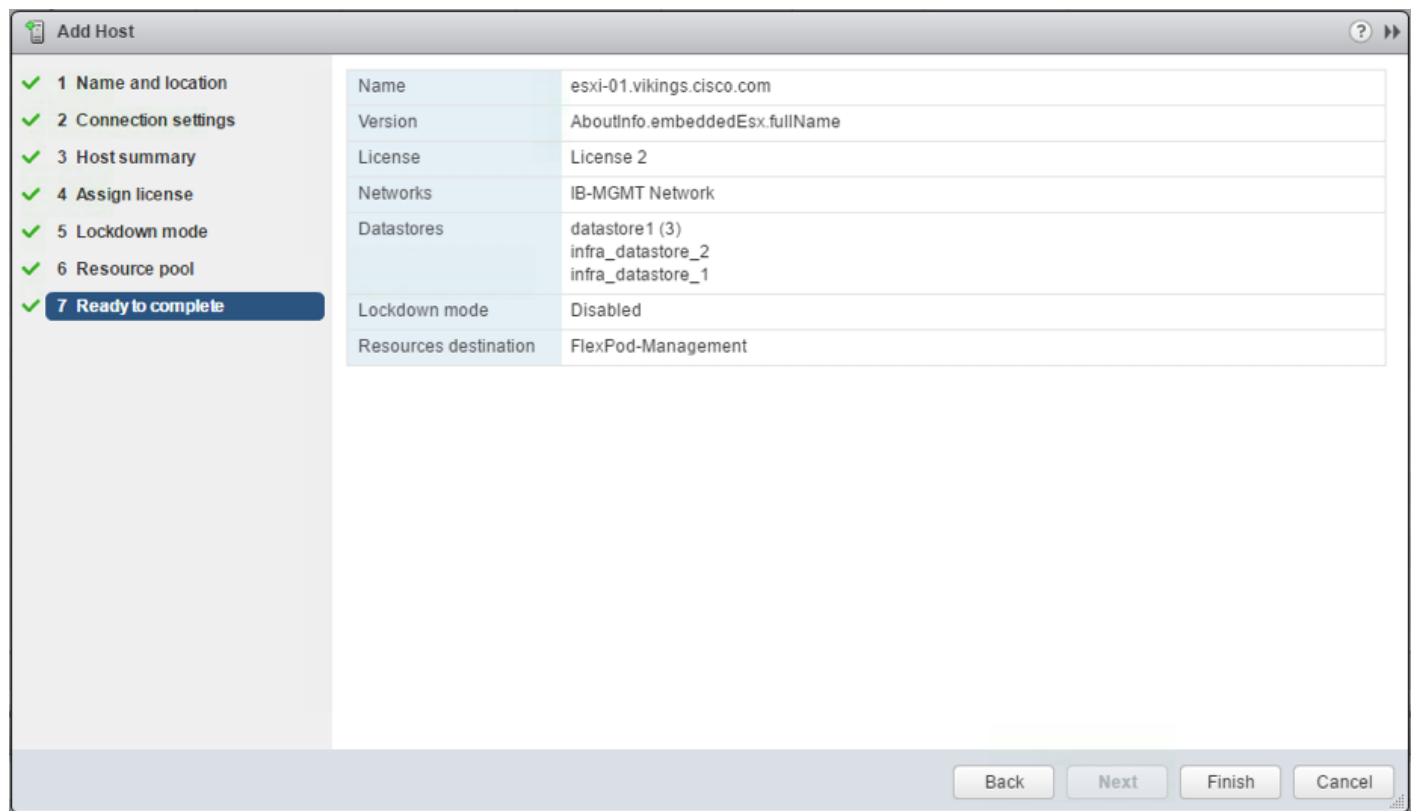
Name	FlexPod-Management
Location	FlexPod-DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	<input checked="" type="checkbox"/> Enable admission control
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

OK Cancel

11. Click OK to create the new cluster.
12. On the left pane, double click the “FlexPod-DC”.
13. Click Clusters.
14. Under the Clusters pane, right click FlexPod-Management and select Settings.
15. Select Configuration > General in the list on the left and select Edit to the right of General.
16. Select Datastore specified by host and click OK.



17. On the left, right-click FlexPod-Management and click Add Host.
18. In the Host field, enter either the IP address or the FQDN name of one of the VMware ESXi hosts. Click Next.
19. Type root as the user name and the root password. Click Next to continue.
20. Click Yes to accept the certificate.
21. Review the host details and click Next to continue.
22. Assign a license or leave in evaluation mode and click Next to continue.
23. Click Next to continue.
24. Click Next to continue.
25. Review the configuration parameters. Click Finish to add the host.



26. Repeat the steps 17 to 25 to add the remaining VMware ESXi hosts to the cluster.



Two VMware ESXi hosts will be added to the cluster.

Add AD User Authentication to vCenter (Optional)

If an AD Infrastructure is set up in this FlexPod environment, you can setup in AD and authenticate from vCenter.

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).
2. Connect to <https://<vcenter-ip>>, and select Log in to vSphere Web Client.
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. Navigate to Home. In the center pane, select System Configuration under Administration.
5. On the left, select Nodes and under Nodes select the vCenter.
6. In the center pane, select the manage tab, and within the Settings select Active Directory and click Join.
7. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Click OK.

8. On the left, right-click the vCenter and select Reboot.
9. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.
10. Log back into the vCenter Web Client.
11. In the center pane, select System Configuration under Administration.
12. On the left, select Nodes and under Nodes select the vCenter.
13. In the center pane under the Manage tab, select Active Directory. Make sure your Active Directory Domain is listed.
14. Navigate back to the vCenter Home.
15. In the center pane under Administration, select Roles.
16. On the left under Single Sign-On, select Configuration.
17. In the center pane, select the Identity Sources tab.
18. Click the green + sign to add an Identity Source.
19. Select the Active Directory (Integrated Windows Authentication) Identity source type.
20. Your AD domain name should be filled in. Leave Use machine account selected and click OK.
21. Your AD domain should now appear in the Identity Sources list.
22. On the left, under Single Sign-On, select Users and Groups.
23. In the center pane, select your AD domain for the Domain.
24. Make sure the FlexPod Admin user setup in step 1 appears in the list.
25. On the left under Administration, select Global Permissions.
26. Select the Manage tab, and click the green + sign to add a User or Group.
27. In the Global Permission Root - Add Permission window, click Add.
28. In the Select Users/Groups window, select your AD Domain.
29. Under Users and Groups, select either the FlexPod Admin user or the Domain Admins group.



The FlexPod Admin user was created in the Domain Admins group. The selection here depends on whether the FlexPod Admin user will be the only user used in this FlexPod or you would like to add other users later. By selecting the Domain Admins group, any user placed in that group in the AD domain will be able to login to vCenter as an Administrator.

30. Click Add. Click Check names to verify correctness of the names. Click OK to acknowledge the correctness of the names.
31. Click OK to add the selected User or Group.
32. Verify the added User or Group is listed under Users and Groups and the Administrator role is assigned.
33. Click OK.
34. Log out and log back into the vCenter Web Client as the FlexPod Admin user. You will need to add the domain name to the user, i.e. flexadmin@domain.

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 pane, select Services.
4. Under services, click VMware vSphere ESXi Dump Collector.
5. In the center pane, click the green start icon to start the service.
6. In the Actions menu, click Edit Startup Type.
7. Select Automatic.
8. Click OK.
9. Connect to each ESXi host via ssh as root
10. Run the following commands:

```
esxcli system coredump network set -v vmk0 -j <vcenter-ip>
esxcli system coredump network set -e true
esxcli system coredump network check
```

Cisco UCS Manager Plug-in for VMware vSphere Web Client

The Cisco UCS Manager Plug-in for VMware vSphere Web Client allows administration of Cisco UCS domains through the VMware's vCenter administrative interface. Capabilities of the plug-in include:

- View Cisco UCS physical hierarchy
- View inventory, installed firmware, faults, power and temperature statistics

- Map the ESXi host to the physical server
- Manage firmware for B and C series servers
- View VIF paths for servers
- Launch the Cisco UCS Manager GUI
- Launch the KVM consoles of UCS servers
- Switch the existing state of the locator LEDs

Installation is only valid for VMware vCenter 5.5 or higher, and will require revisions of .NET Framework 4.5 and VMware PowerCLI 5.1 or greater.

Cisco UCS Manager Plug-in Installation

To begin the plug-in installation on a Windows system that meets the previously stated requirements:

1. Download the plugin and registration tool from:
<https://software.cisco.com/download/release.html?mdfid=286282669&reltype=latest&relind=AVAILABLE&dwld=true&softwareid=286282010&catid=282558030&rellifecycle=&atcFlag=N&release=2.0.1&dwldImageGuid=5963DCD457E74C173F61512F00BB6FE78C5D8B72>
2. Place the downloaded ucs-vcplugin-2.0.1.zip file onto the web server used for hosting the ONTAP software and VMware ESXi ISO.
3. Unzip the Cisco_UCS_Plugin_Registration_Tool_1_1_3.zip and open the executable file within it.
4. Leave Register Plugin selected for the Action and fill in:
 - IP/Hostname
 - Username
 - Password
 - URL that plugin has been uploaded

Cisco UCS Plugin Registration Tool v1.1.3

This tool registers/unregisters the Cisco UCS Plugin for VMware vSphere Web Client

Action

Register Plugin Unregister Plugin

vCenter Details

IP/Hostname: 10.1.156.100

Username: administrator@vsphere.local

Password: *****


Plugin Location

URL of the plugin location in HTTP/HTTPS server
Ex: https://10.10.10.1/plugins/ucs-vcplugin-1.0.1.zip

http://10.1.156.150/bears/UCS/ucs-vcplugin-2.0.1.zip

Submit Cancel

5. A pop-up will appear explaining that 'allowHttp=true' will need to be added to the webclient.properties file on the VCSA in the /etc/vmware/vsphere-client directory.



6. Take care of this issue after the plugin has been registered, click OK to close the Information dialog box.
7. Click Submit to register the plugin with the vCenter Server Appliance.
8. To resolve the change needed for the HTTP download of the vSphere Web Client launch, connect to the VCSA with ssh using the root account and type:

```
echo "allowHttp=true" >> /etc/vmware/vsphere-client/webclient.properties
```



This will add "allowHttp=true" to the end of the webclient.properties file. Take care with this command to use two greater than symbols ">>" to append to the end of the configuration file, a single greater than symbol will replace the entire pre-existing file with what has been sent with the echo command.

9. Reboot the VCSA.

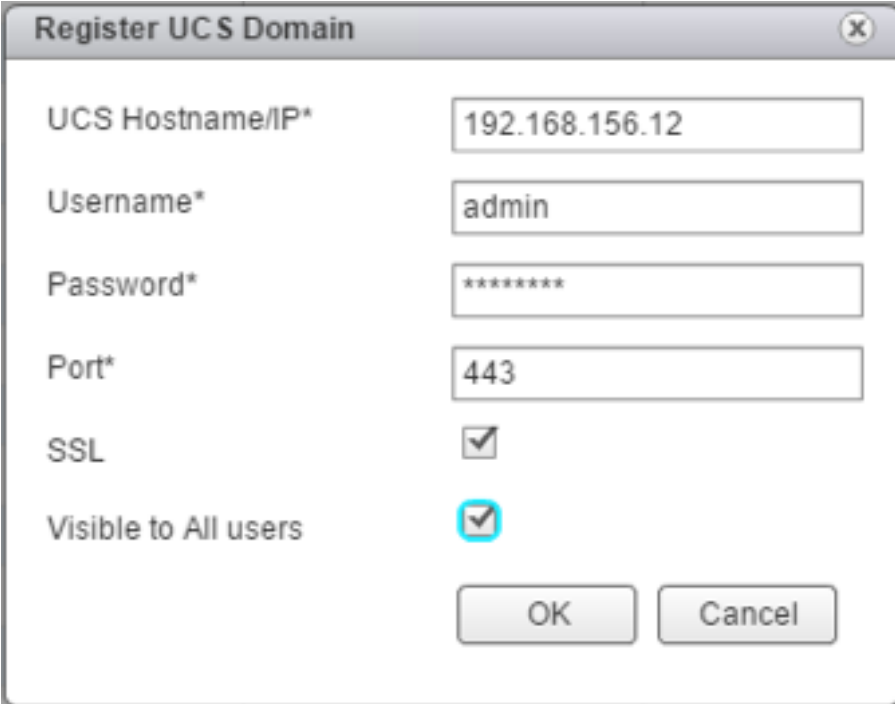
Cisco UCS Domain Registration

Registration of the FlexPod UCS Domain can now be performed. The account used will correlate to the permissions allowed to the plugin, admin will be used in our example, but a read only account could be used with the plugin if that was appropriate for the environment.

To register the Cisco UCS Domain, complete the following steps::

1. Opening up the vSphere Web Client.

2. Select the Home from the Navigator or pull-down options, and double click the Cisco UCS icon appearing in the Administration section.
3. Click the Register button and provide the following options in the Register UCS Domain dialogue box that appears:
 - UCS Hostname/IP
 - Username
 - Password
 - Port (if different than 443)
 - Leave SSL selected and click on the Visible to All users option



Register UCS Domain

UCS Hostname/IP* 192.168.156.12

Username* admin

Password* *****

Port* 443

SSL

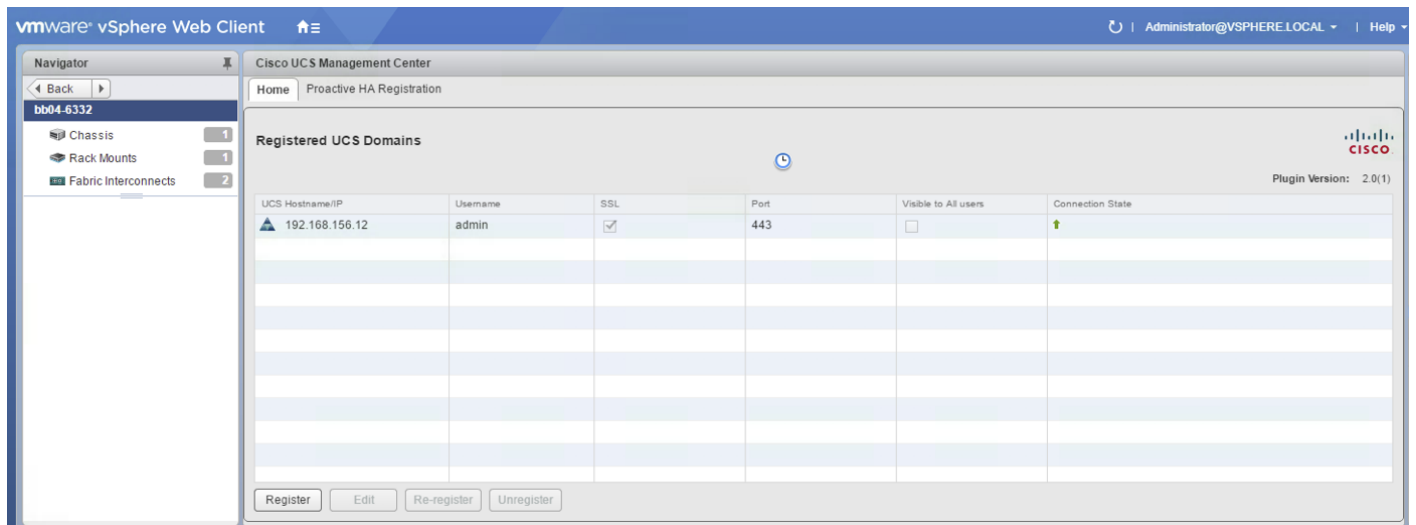
Visible to All users

OK Cancel


4. Click OK to register the UCS Domain.

Using the Cisco UCS vCenter Plugin

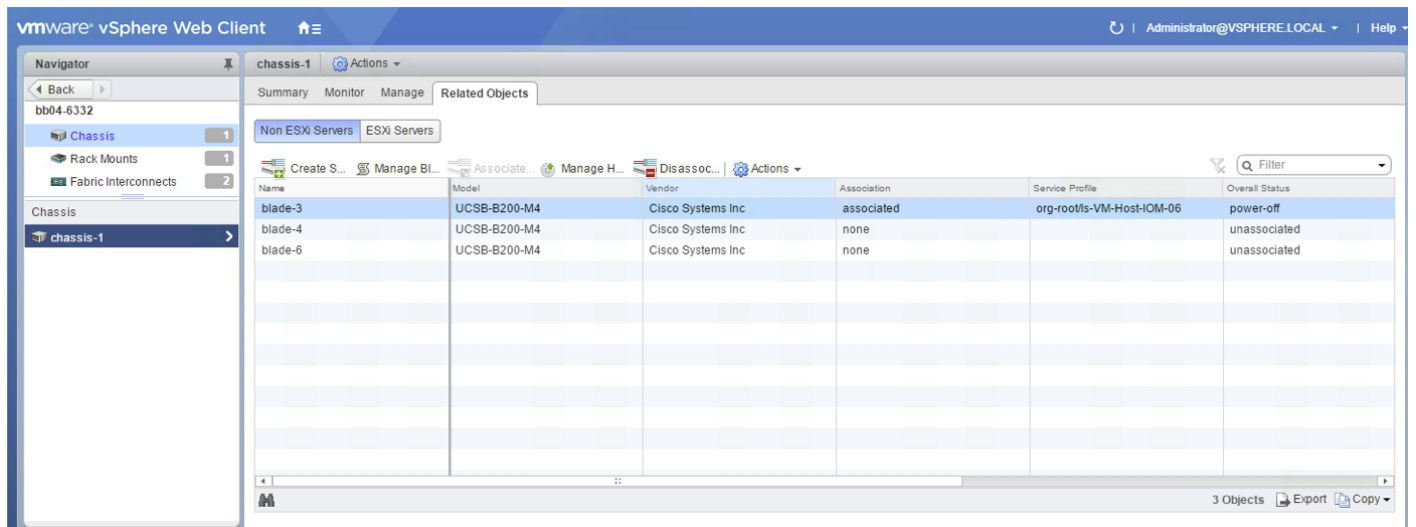
The plugin can now enable the functions described at the start of this section by double-clicking the registered Cisco UCS Domain:



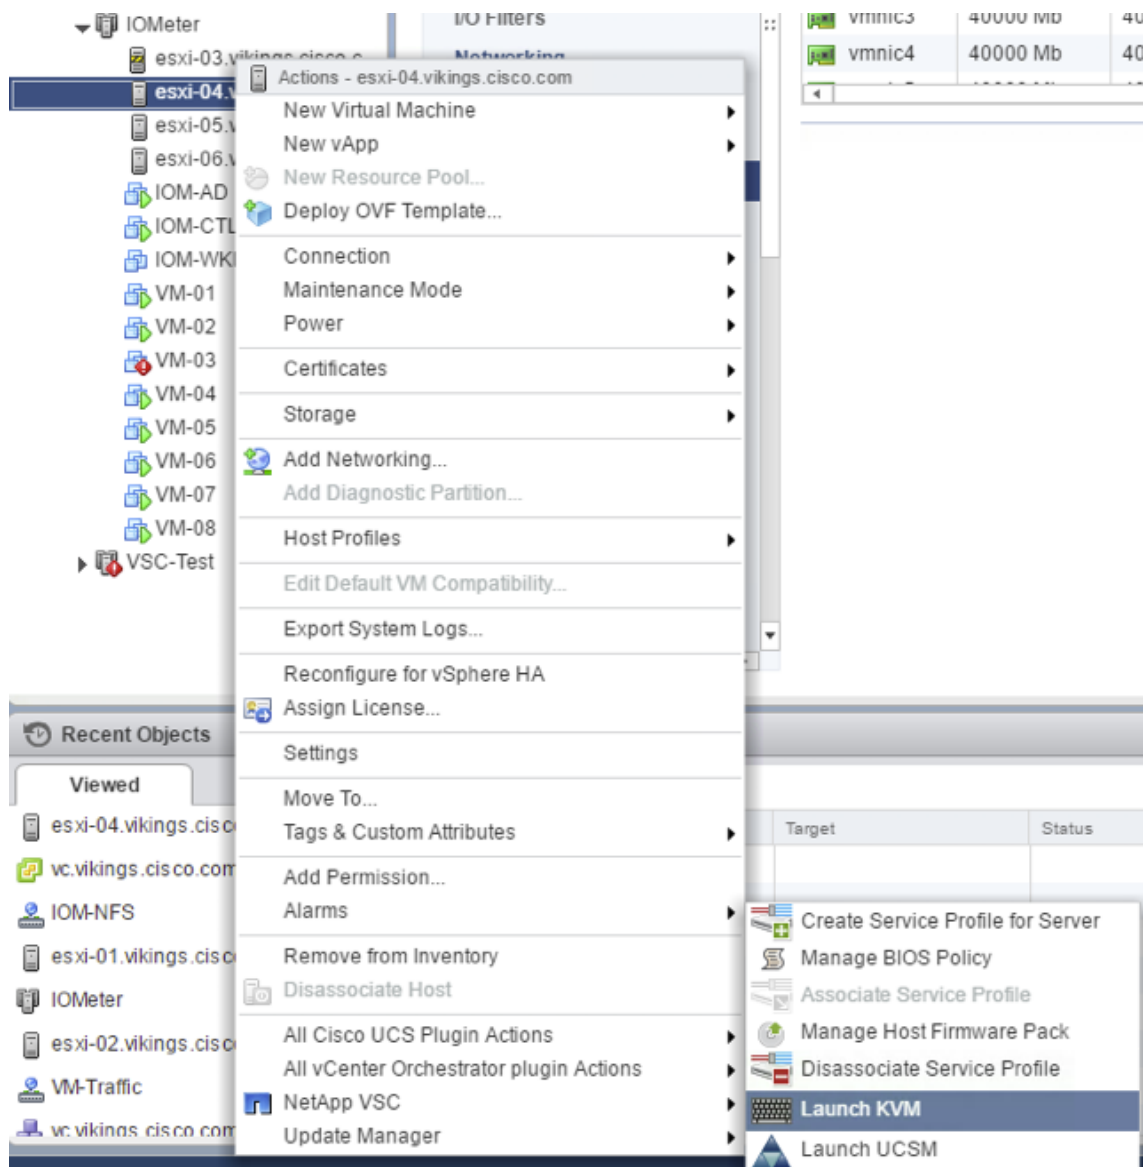
This will pull up a view of the components associated to the domain:



Selecting within the chassis or rack mounts will provide you with a list of ESXi or non-ESXi servers to perform operations:



In addition to viewing and working within objects shown in the Cisco **UCS Plugin's view of the** Cisco UCS Domain, direct access of Cisco UCS functions provided by the plugin can be selected within the pulldown options of hosts registered to vCenter:



For the installation instructions and usage information, please refer to the Cisco UCS Manager Plug-in for VMware vSphere Web Client User Guide at:
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vmware_tools/vCenter/vCenter_Plugin_User_Guide/2x/b_vCenter_2x.html

FlexPod VMware vSphere Distributed Switch (vDS)

This section provides detailed procedures for installing the VMware vDS on the FlexPod ESXi Management Hosts.

In the Cisco UCS setup section of this document two sets of vNICs (Infra-A and B, and iSCSI-A and B) were setup. The vmnic ports associated with the Infra-A and B vNICs will be migrated to VMware vDS in this procedure. The critical infrastructure VLAN interfaces and vMotion interfaces will be placed on the vDS.

An IB-Mgmt VLAN and a VM-Traffic VLAN port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would need to have the corresponding VLANs added to the Cisco UCS LAN


cloud, to the Cisco UCS Infra-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC and peer-link interfaces on the switches.

Configure the VMware vDS in vCenter


VMware vSphere Web Client

To configure the vDS, complete the following steps:

1. After logging into the VMware vSphere Web Client, select Networking under the Home tab.
2. Right-click the FlexPod-DC datacenter and select Distributed Switch > New Distributed Switch.
3. Give the Distributed Switch a descriptive name and click Next.
4. Make sure Distributed switch: 6.5.0 is selected and click Next.
5. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter VM-Traffic for the Port group name. Click Next.
6. Review the information and click Finish to complete creating the vDS.
7. On the left, expand the FlexPod-DC datacenter and the newly created vDS. Select the newly created vDS.
8. Select the VM-Traffic port group. In the center pane, select the Edit distributed port group settings icon. The Edit button can be used to change the number of ports in the port group to a number larger than the default of 8. All of the other properties of the port group can also be changed under Edit.
9. Select the vDS on the left. Click the Edit distributed switch settings icon on the right.
10. On the left in the Edit Settings window, select Advanced.
11. Change the MTU to 9000. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.



12. Three port groups will be created for infrastructure use, first being for the management vmkernel. On the left, right-click the vDS, select Distributed Port Group, and select **New Distributed Port Group... within the pull-down options of Distributed Port Group.**
13. For the first port group used for vMotion on the left, right-click the vDS, select Distributed Port Group, and select **New Distributed Port Group... within the pull-down options of Distributed Port Group.**
14. Enter vMotion as the name and click Next.
15. Set the VLAN type to VLAN, enter the VLAN used for vMotion, click the Customize default policies configuration check box, and click Next.
16. Leave the Security options set to Reject and click Next.
17. Leave the Ingress and Egress traffic shaping options as Disabled, and click Next.
18. Select Uplink 1 from the list of Active uplinks, and click the down arrow icon twice to place Uplink 1 in the list of Standby uplinks.



19. Click Next.

20. Leave Netflow Disabled and click Next.

21. Leave Block all ports set as No and click Next.

22. Leave the additional settings dialogue options as they are shown and click Next.

23. Confirm the options and click Finish to create the port group.

24. For the second port group used for Infrastructure In-Band Management on the left, right-click the vDS, select Distributed Port Group, and select New Distributed Port Group... within the pull-down options of Distributed Port Group.

25. Enter IB-MGMT as the name and click Next.

26. Set the VLAN type to VLAN, enter the VLAN used for In-Band Management, click the Customize default policies configuration check box, and click Next.

27. Leave the Security options set to Reject and click Next.

28. Leave the Ingress and Egress traffic shaping options as Disabled, and click Next.

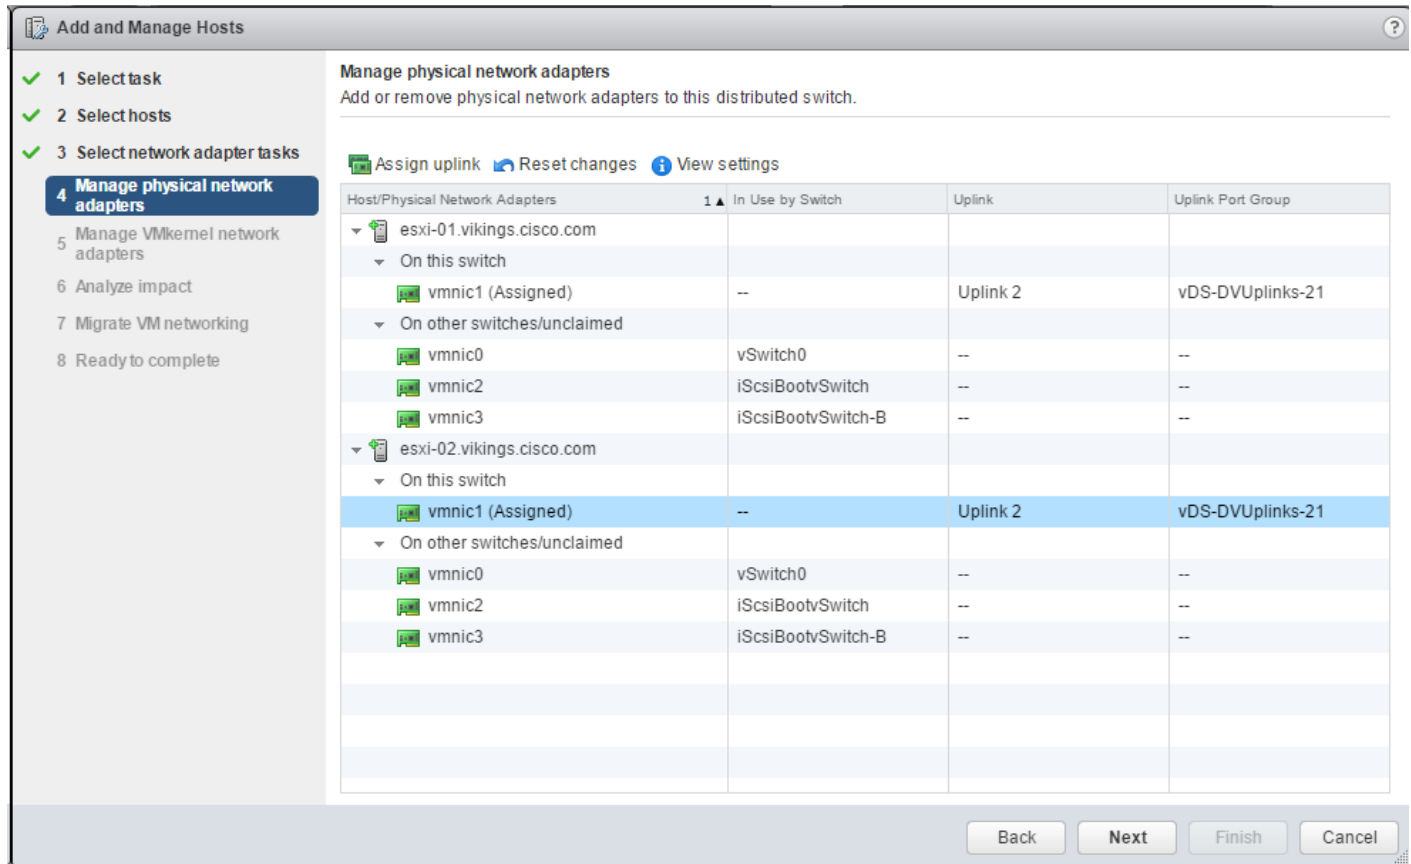
29. Select Uplink 2 from the list of Active uplinks, and click the down arrow icon twice to place Uplink 2 in the list of Standby uplinks. This step is pinning In-Band Management traffic to UCS Fabric A.

30. Click Next.

31. Leave Netflow Disabled and click Next.
32. Leave Block all ports set as No and click Next.
33. Leave the additional settings dialogue options as they are shown and click Next.
34. Confirm the options and click Finish to create the port group.
35. Repeating this same procedure, a second time for the Infrastructure NFS vmkernel. On the left, right-click the vDS, select **Distributed Port Group**, and select **New Distributed Port Group...** within the pull-down options of Distributed Port Group.
36. Enter Infra-NFS as the name and click Next.
37. Set the VLAN type to VLAN, enter the VLAN used for NFS, and click Next.
38. Click Finish to create the port group.
39. On the left, right-click the vDS and select Add and Manage Hosts.
40. Make sure Add hosts is selected and click Next.
41. Click the green + sign to add hosts. Select the two FlexPod Management hosts and click OK. Click Next.
42. Leave Manage physical adapters and Manage VMkernel adapters selected. Select Migrate virtual machine networking and click Next.
43. Select vmnic1 on the first host and click Assign uplink. Select Uplink 2 and click OK. Repeat this process to assign vmnic1(Uplink2) from both hosts to the vDS.



It is important to assign vmnic1 (UCS Fabric B) to Uplink2. This allows the port groups to be pinned to the appropriate fabric.



44. Click Next.

45. Select vmk0 on the first host and click Assign port group.

46. Select the IB-MGMT destination port group and click OK.


47. Select vmk3 on the first host and click Assign port group.

48. Select the vMotion destination port group and click OK.

49. Select vmk4 on the first host and click Assign port group.

50. Select the Infra-NFS destination port group and click OK.

51. Repeat this process for the second ESXi host.

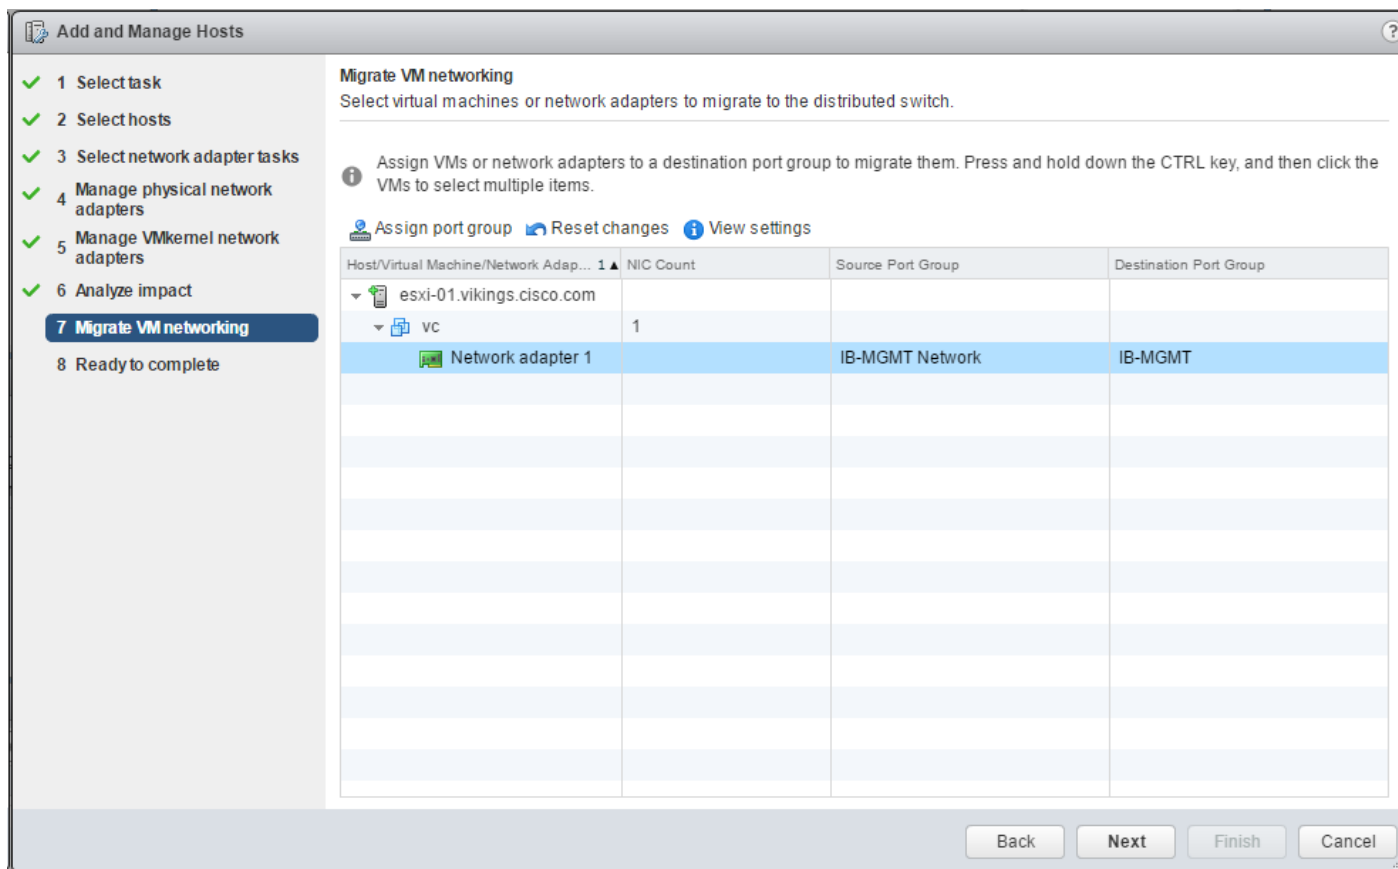


52. Confirm the 3 VMkernel adapters on each host have valid Destination Port Groups and Click Next.

53. Click Next after confirming there is no impact detected in the Analyze impact screen.

54. In the Migrate VM networking window, expand the vCenter VM and select Network adapter 1.

55. Click Assign port group, select the IB-MGMT port group and click OK.



56. Click Next.

57. Click Finish to complete adding the two ESXi hosts to the vDS.

58. Select Hosts and Clusters and select ESXi Host 1.

59. Under the Configure tab, in the pane on the left, select Virtual switches.

60. In the center pane under Virtual switches, select vSwitch0.


61. Under Virtual switches, select the red X icon to delete vSwitch0. Click Yes to confirm.

62. Under Virtual switches, select the vDS.

63. Under Virtual switches, select the third icon to Manage the physical network adapters connected to the virtual switch.

64. Click the green + icon to add an uplink.

65. Make sure vmnic0 and Uplink 1 are selected and click OK.



66. Click OK to complete adding the adapter to the vDS on the host.

67. Repeat this process to add vmnic0 to the vDS on ESXi Host 2.

FlexPod Management Tools Setup

NetApp Virtual Storage Console 6.2.1 Deployment Procedure

This section describes the deployment procedures for the NetApp Virtual Storage Console (VSC).

Virtual Storage Console 6.2.1P1 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.1:

- Protocol licenses (NFS and iSCSI)
- NetApp FlexClone® (for provisioning and cloning only)
- NetApp SnapRestore® (for backup and recovery)
- The NetApp SnapManager® Suite

Install Virtual Storage Console 6.2.1P1


To install the VSC 6.2.1P1 software, complete the following steps:

1. Build a VSC VM with Windows Server 2012 R2, 4GB of RAM, two CPUs, and one virtual network interface in the `IB-MGMT Network` port group. The virtual network interface should be a VMXNET 3 adapter.
2. Bring up the VM, install VMware Tools, assign the IP address and gateway in the IB-MGMT subnet, 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 the FlexPod Admin user using the VMware console.
6. From the VMware console on the VSC VM, download the x64 version of [Virtual Storage Console 6.2.1P1](#) from the [NetApp Support](#) site.
7. Right-click the `VSC-6.2.1P1-win64.exe` file downloaded in step 6 and select Run as Administrator.
8. Select the appropriate language and click OK.
9. On the Installation wizard Welcome page, click Next.
10. Select the checkbox to accept the message and 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 Virtual Storage Console 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 of the VSC VM.
4. In the vCenter Server Information section, enter the host name or IP address, the user name (FlexPod admin user or root), and the user password for the vCenter Server. Click Register to complete the registration.

vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.

Plugin service information	
Host name or IP Address:	<input type="text" value="vsc.vikings.cisco.com"/>
vCenter Server information	
Host name or IP Address:	<input type="text" value="vc.vikings.cisco.com"/>
Port:	<input type="text" value="443"/>
User name:	<input type="text" value="administrator@vsphere.local"/>
User password:	<input type="password" value="••••••••"/>

Register


5. With a successful registration, the storage controller discovery automatically begins.

Install NetApp NFS VAAI Plug-in

To install the NetApp NFS VAAI Plug-in, complete the following steps:

1. Download the NetApp NFS Plug-in 1.1.2 for VMware .vib file from the [NFS Plugin Download](#) on the VSC VM.
2. Rename the downloaded file `NetAppNasPlugin.vib`.

3. Move the file to the C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web folder.



Discover and Add Storage Resources

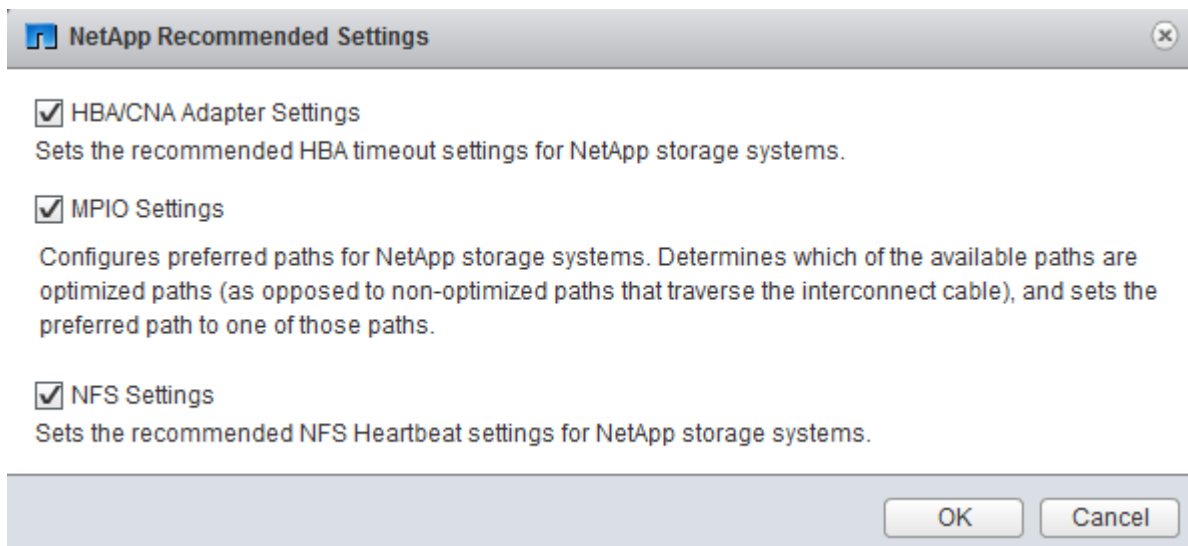
To discover storage resources for the Monitoring and Host Configuration capability and the Provisioning and Cloning capability, complete the following steps:

1. Using the vSphere web client, log in to the vCenter Server as the FlexPod admin user. 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 Use TLS to Connect to This Storage System is selected. Click OK.
5. Click OK to accept the controller privileges.
6. Wait for the Storage Systems to update. You may need to click Refresh to complete this update.

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 on vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for these hosts.

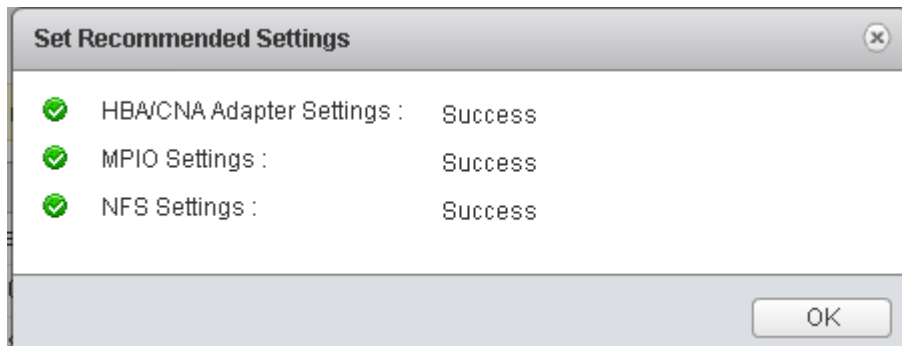


2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings.



This functionality sets values for HBAs and converged network adapters (CNAs), sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).

3. Click OK.



4. From the Home screen in the vSphere Web Client, select Virtual Storage Console.
5. On the left under Virtual Storage Console, select NFS VAAI Tools.
6. Make sure that NFS Plug-in for VMware VAAI Version 1.1.2-3 is shown.
7. Click Install on Host.
8. Select both ESXi hosts and click Install.
9. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.



In testing a conflict has been identified between NetApp VSC, vSphere 6.5, and UCS servers containing the LSI Megaraid SAS Invader local disk controller. If NetApp VSC cannot discover and set optimized settings for a server, the server most likely has the Megaraid SAS Invader controller installed. This controller utilizes the lsi-mr3 ESXi vib. A workaround for this problem is to disable the disk controller by running “esxcli system module set --enabled=false --module=lsi_mr3” from an ESXi console or ssh prompt and then rebooting the host. When the host comes back up, the NetApp VSC functions should then work.


Virtual Storage Console 6.2.1P1 Provisioning Datastores

Using VSC, the administrator can provision NFS, FC or iSCSI datastore and attach it to single host or multiple hosts in the cluster. Following steps illustrates provisioning a datastore and attach it to the cluster.

Provisioning NFS Datastore

To provision the NFS datastore, complete the following steps:

1. From the Home screen of the vSphere Web Client, right-click the FlexPod-Management cluster and **select “NetApp VSC > Provision Datastore”**.



2. Enter the datastore name and select the type as NFS.
3. Click Next.

NetApp Datastore Provisioning Wizard

1 Name and type
2 Storage system
3 Details
4 Ready to complete

Specify the name and type of datastore you want to provision.

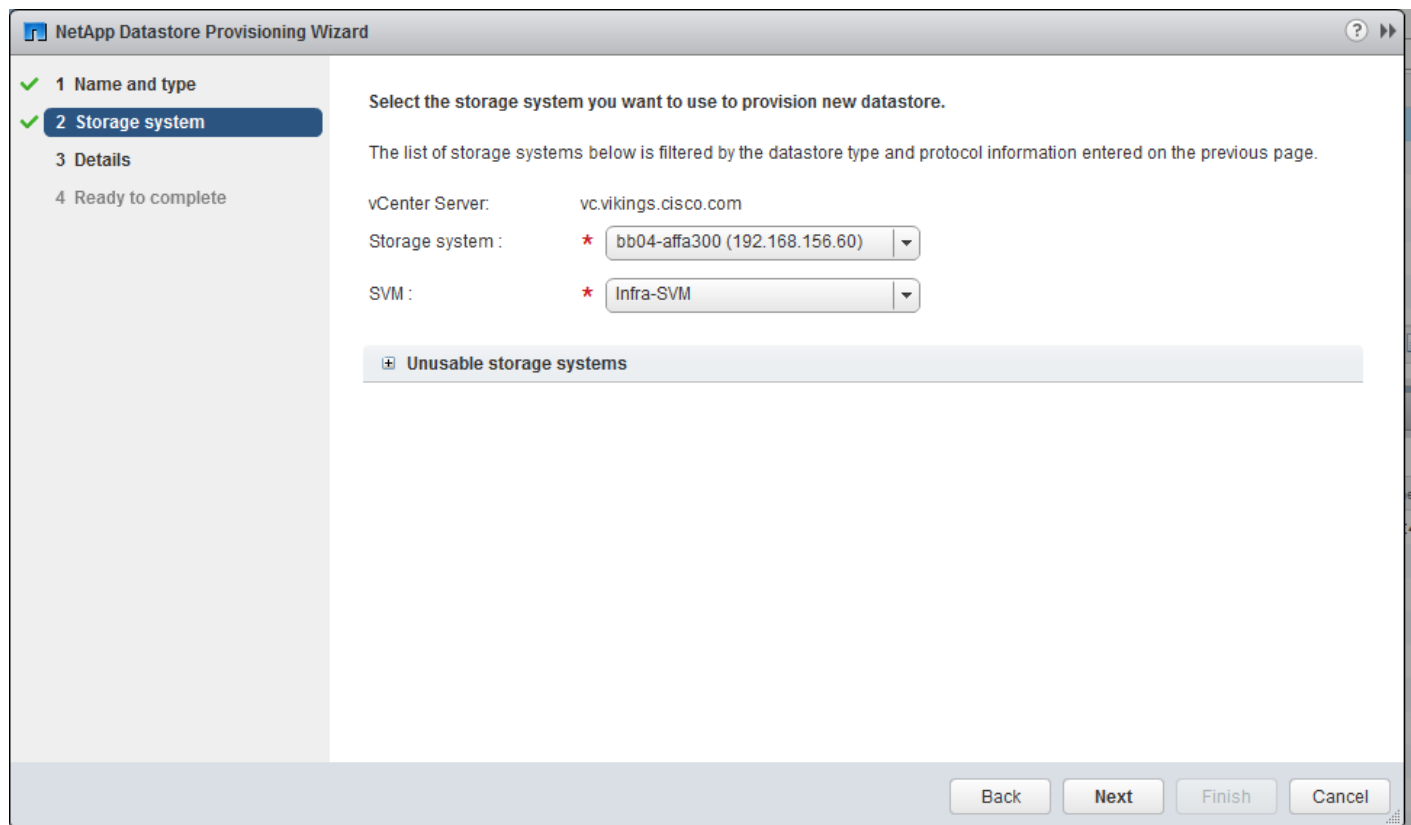
You will be able to select the storage system for your datastore in the next page of this wizard.

Name : * infra_NFS_datastore_1

Type : * NFS VMFS

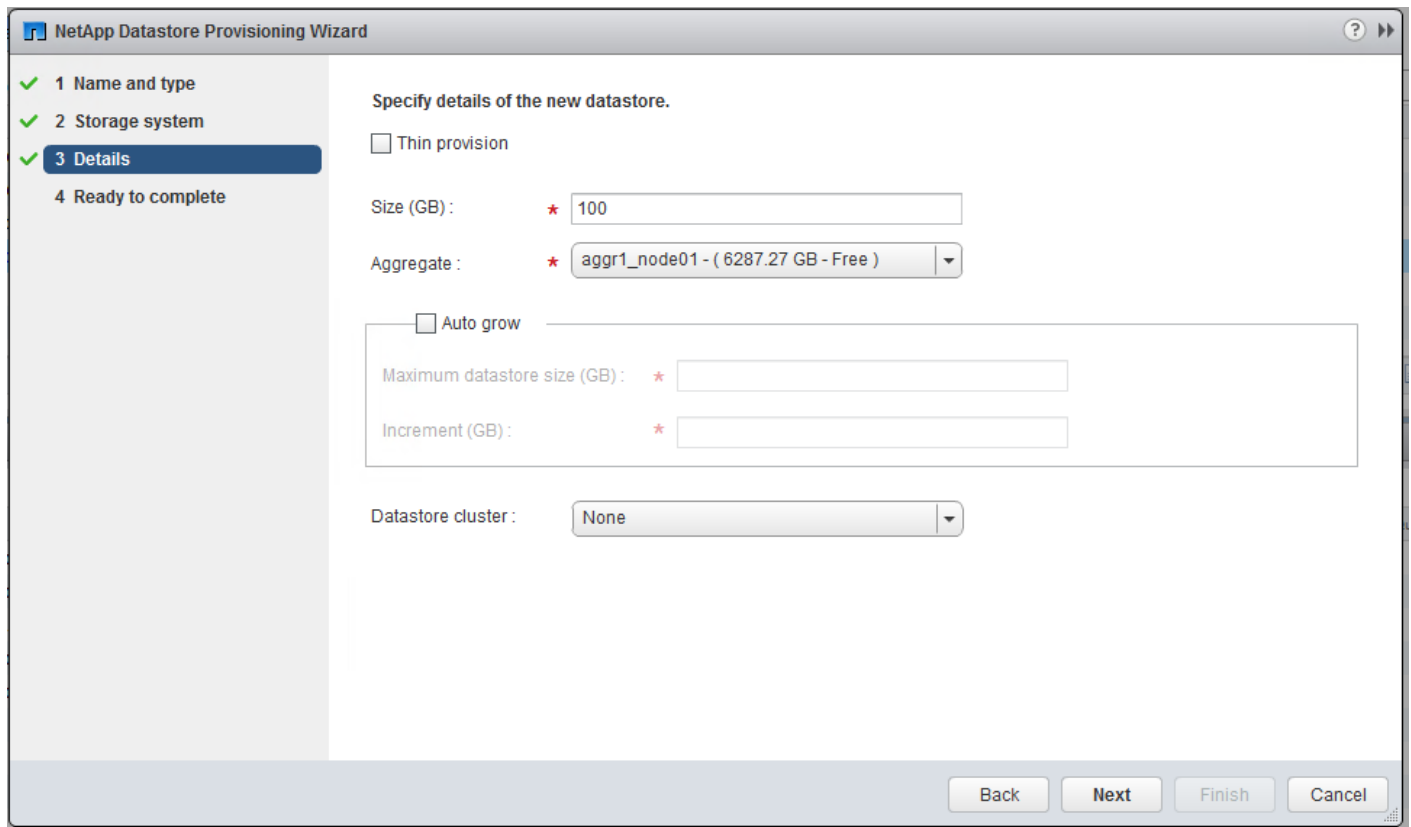
Back Next Finish Cancel

4. Select the cluster name in the Storage system and desired SVM to create the datastore. In this example, Infra-SVM is selected.
5. Click Next.

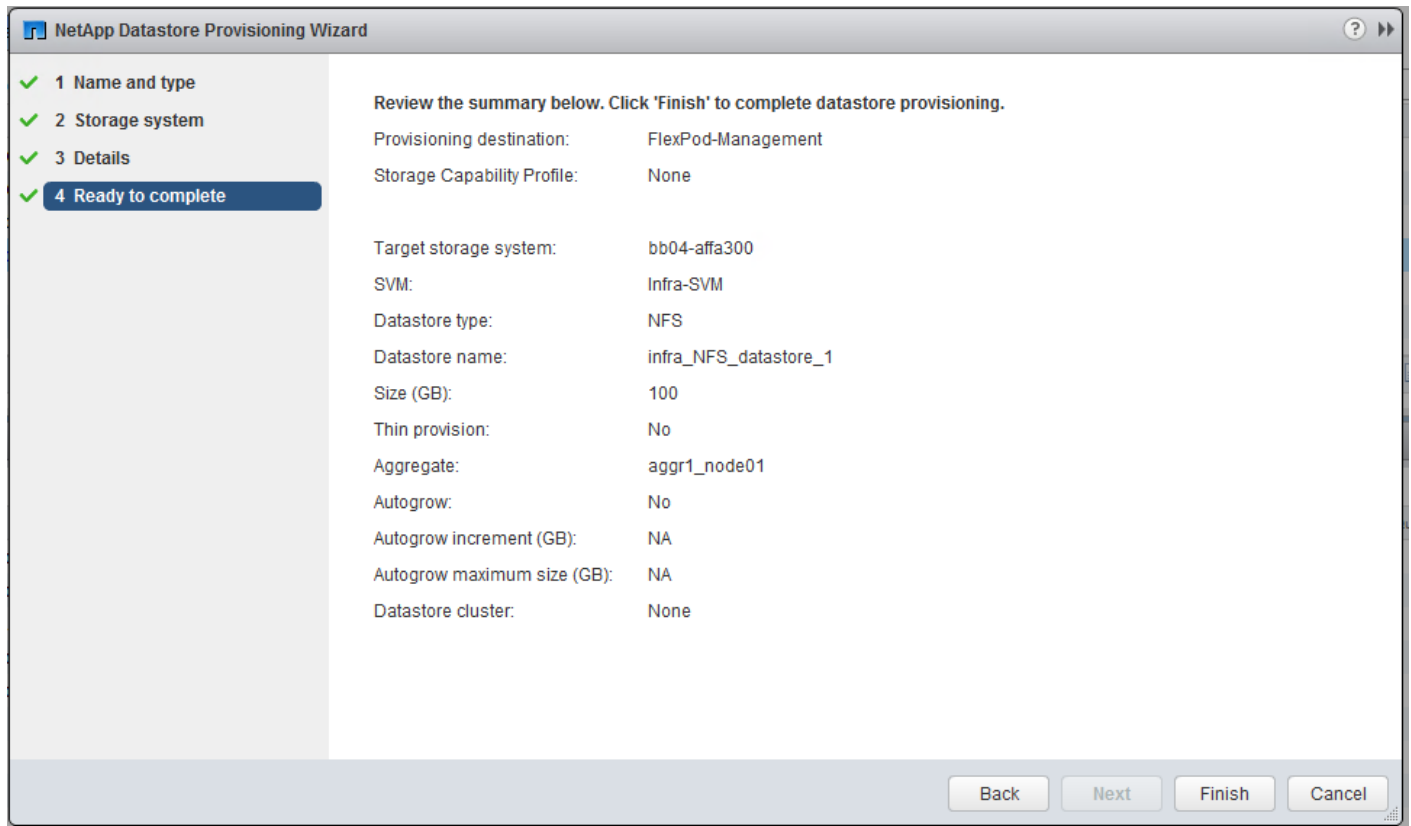


The screenshot shows the 'NetApp Datastore Provisioning Wizard' window. The left sidebar contains a progress indicator with four steps: '1 Name and type' (checked), '2 Storage system' (highlighted), '3 Details', and '4 Ready to complete'. The main area is titled 'Select the storage system you want to use to provision new datastore.' and includes a note: 'The list of storage systems below is filtered by the datastore type and protocol information entered on the previous page.' Below this, there are three fields: 'vCenter Server:' with the value 'vc.vikings.cisco.com', 'Storage system :' with a dropdown menu showing 'bb04-affa300 (192.168.156.60)', and 'SVM :' with a dropdown menu showing 'Infra-SVM'. A section titled 'Unusable storage systems' is visible but empty. At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

6. Enter the size of the datastore and select the aggregate name.
7. Click Next.



8. Review the details and click Finish.



9. Click Ok.

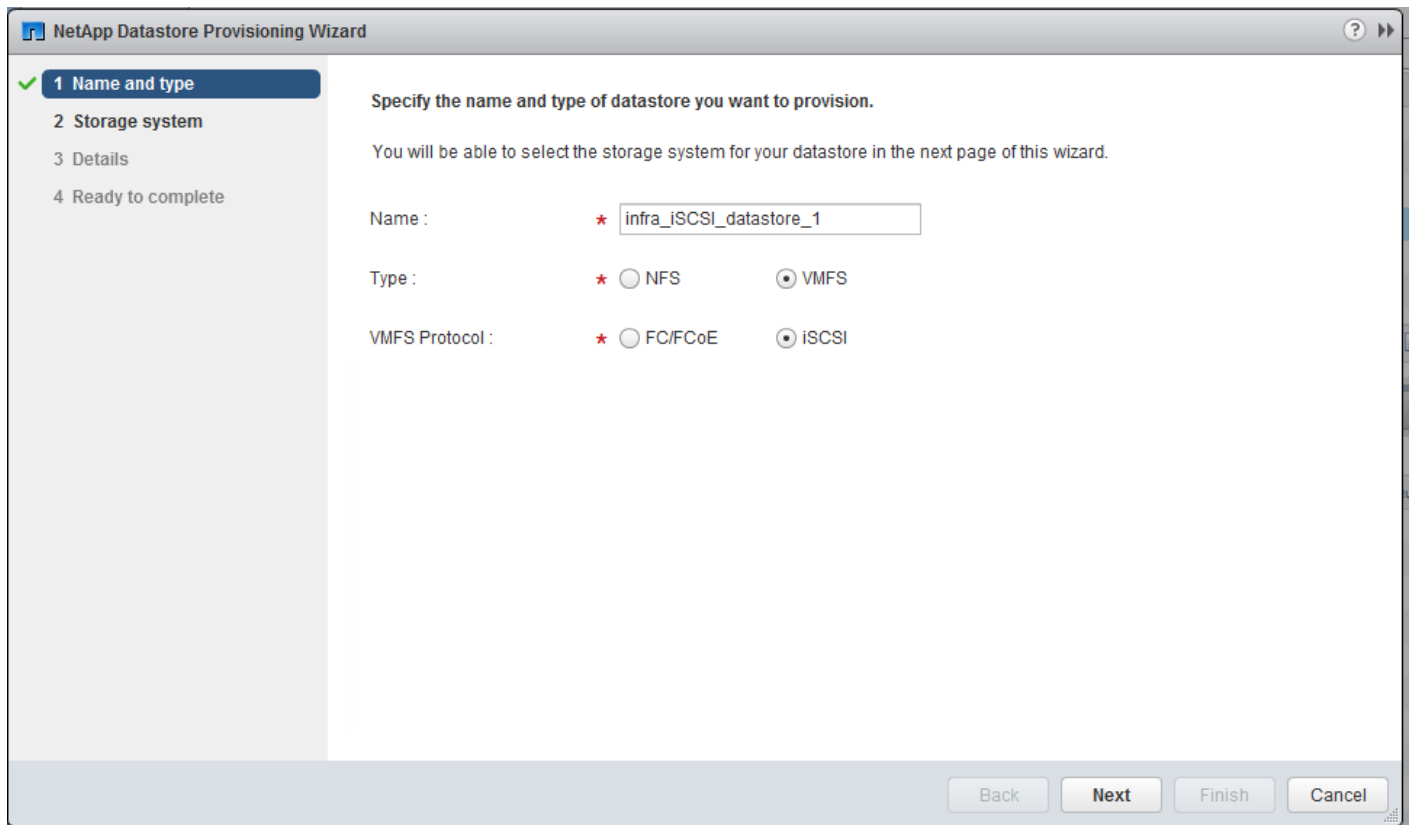


The datastore will be created and mounted on all the hosts in the cluster. Click Refresh screen from the vSphere web client to see the newly created datastore.

Provisioning iSCSI Datastore

To provision the iSCSI datastore, complete the following steps:


1. From the Home screen of the vSphere Web Client, right click the FlexPod-Management cluster and select **“NetApp VSC > Provision Datastore”**.
2. Enter the datastore name and select the type as VMFS. For VMFS protocol, select iSCSI.
3. Click Next.




4. Select the cluster name in the Storage system and desired SVM to create the datastore. In this example, Infra-SVM is selected.
5. Click Next.

The screenshot shows the 'NetApp Datastore Provisioning Wizard' window. The left sidebar contains a progress indicator with four steps: '1 Name and type' (checked), '2 Storage system' (highlighted), '3 Details', and '4 Ready to complete'. The main content area is titled 'Select the storage system you want to use to provision new datastore.' and includes a note: 'The list of storage systems below is filtered by the datastore type and protocol information entered on the previous page.' Below this, there are three fields: 'vCenter Server:' with the value 'vc.vikings.cisco.com', 'Storage system :' with a dropdown menu showing 'bb04-affa300 (192.168.156.60)', and 'SVM :' with a dropdown menu showing 'Infra-SVM'. A section titled 'Unusable storage systems' is visible but empty. At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

6. Enter the size of the datastore. Select **“Create new volume”** check box and select the aggregate name.
7. Click Next.



8. Review the details and click Finish.



9. Click Ok.



The datastore will be created and mounted on all the hosts in the cluster. Click Refresh screen from vSphere web client to see the newly created datastore

Virtual Storage Console 6.2.1P1 Backup and Recovery

Prerequisites for Use of Backup and Recovery Capability

Before you begin using the Backup and Recovery capability to schedule backups and restores of your datastores, VMs, or virtual disk files, you must confirm that the storage systems that contain the datastores and VMs for which you are creating backups have valid storage credentials.

If you plan to leverage the SnapMirror update option, add all of the destination storage systems with valid storage credentials.

Backup and Recovery Configuration

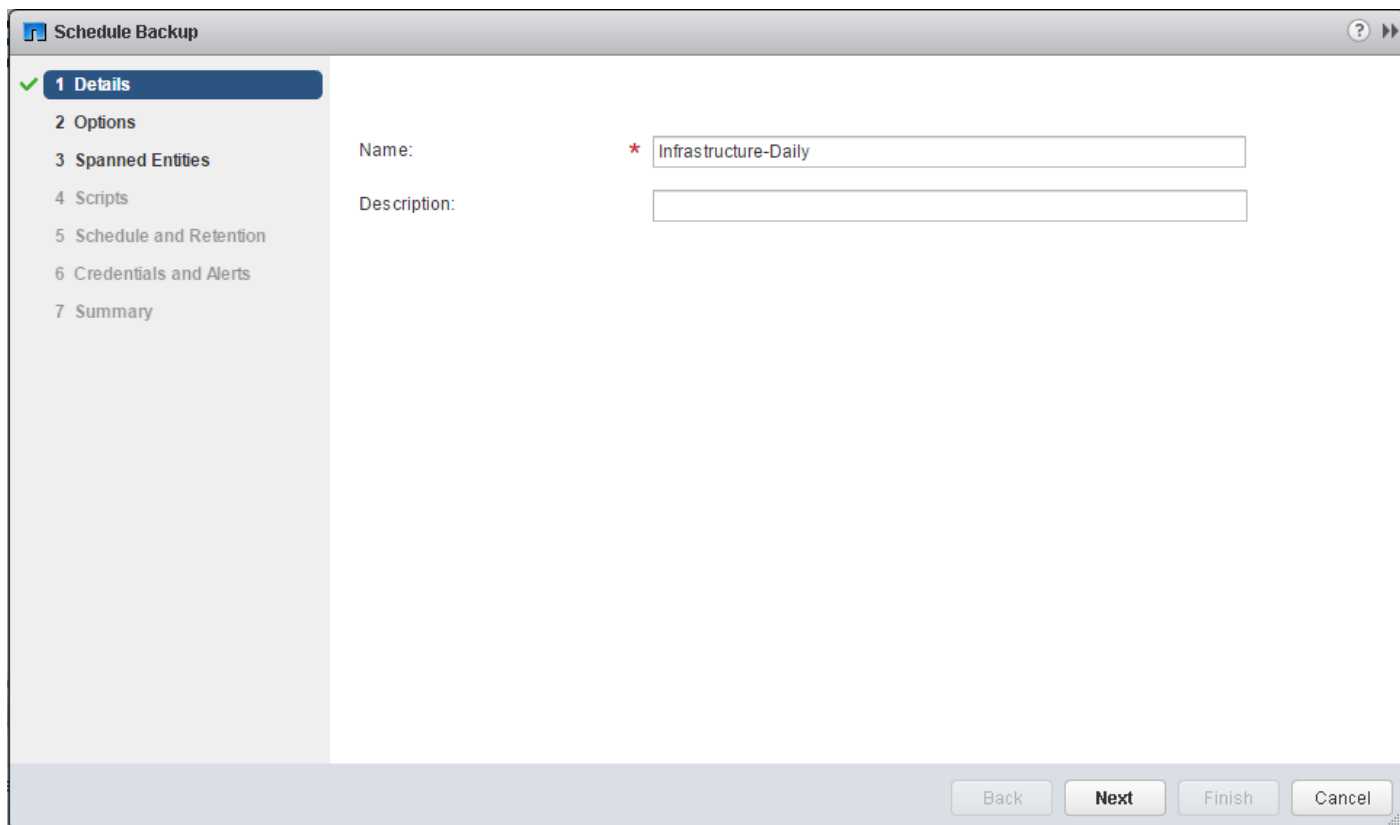
To configure a backup job for a datastore, complete the following steps

1. From the Home screen of the vSphere Web Client, select the Home tab and click Storage.
2. On the left, expand the datacenter.
3. Right-click the datastore that you need to backup. Select NetApp VSC > Schedule Backup.

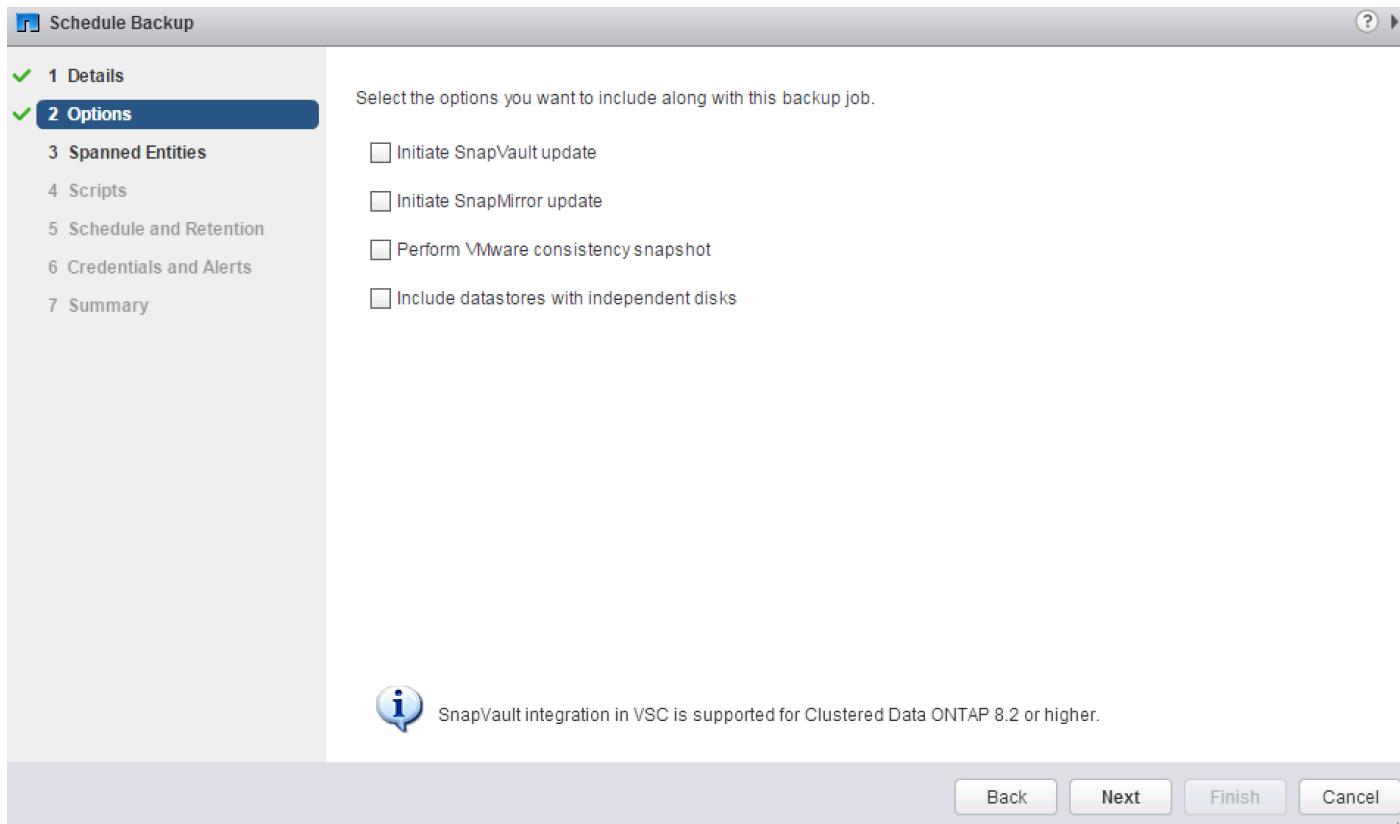


If you prefer a one-time backup, choose Backup Now instead of Schedule Backup.

4. Type a backup job name and description. Click Next.



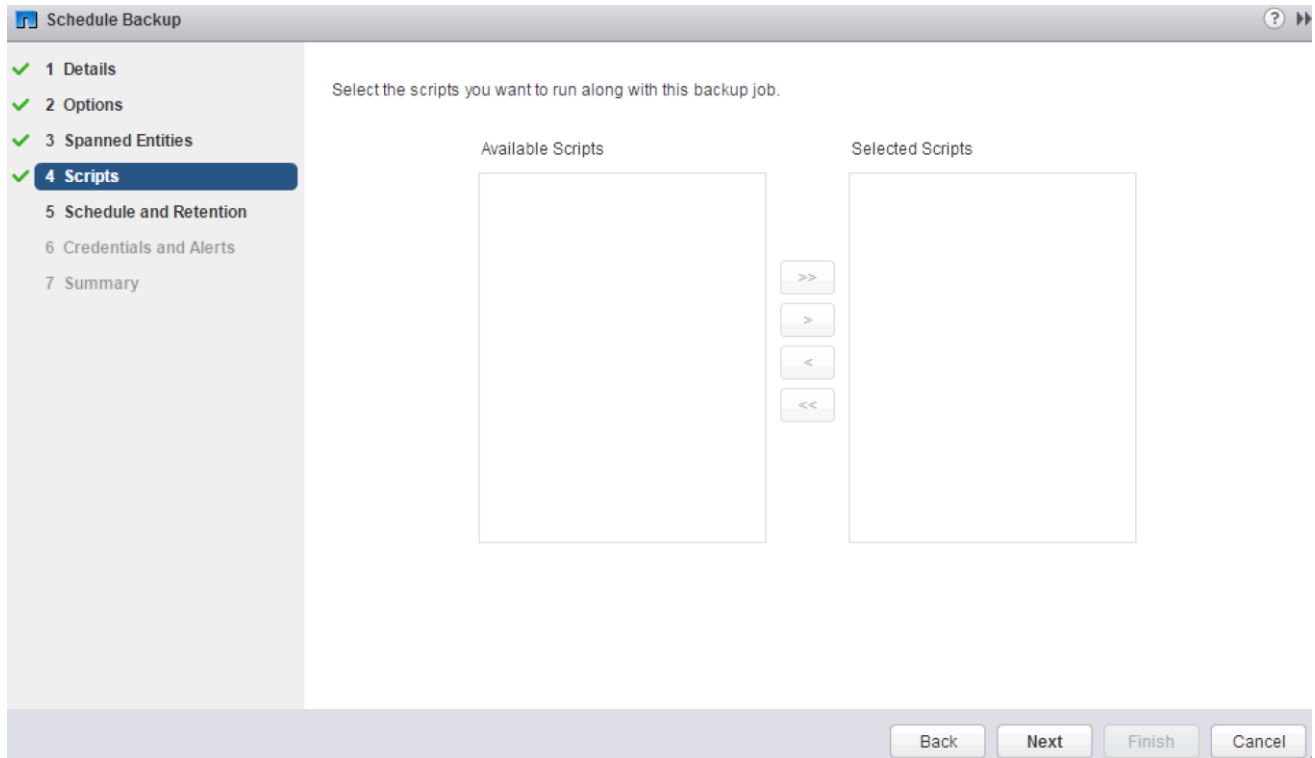
5. Select the options necessary for the backup.





For consistent VM snapshots, select Perform VMware Consistency Snapshot to make a VMware snapshot of each VM just before the NetApp Snapshot copy is made. The VMware snapshot is then deleted after the NetApp Snapshot copy is made.

6. Click Next on the Options screen.
7. Click Next on the Spanned Entities screen.
8. Select one or more backup scripts if available, and click Next in the Scripts screen.



9. Select the hourly, daily, weekly, or monthly schedule and retention policy that you want for this backup job. Click Next.

Schedule Backup

Configure the schedule and retention settings for this job.

Schedule

Hourly

Daily

Weekly

Monthly

On demand only

Daily schedule details

Backup will be performed daily

Starting: 02/27/2017 01:00 AM

Retention

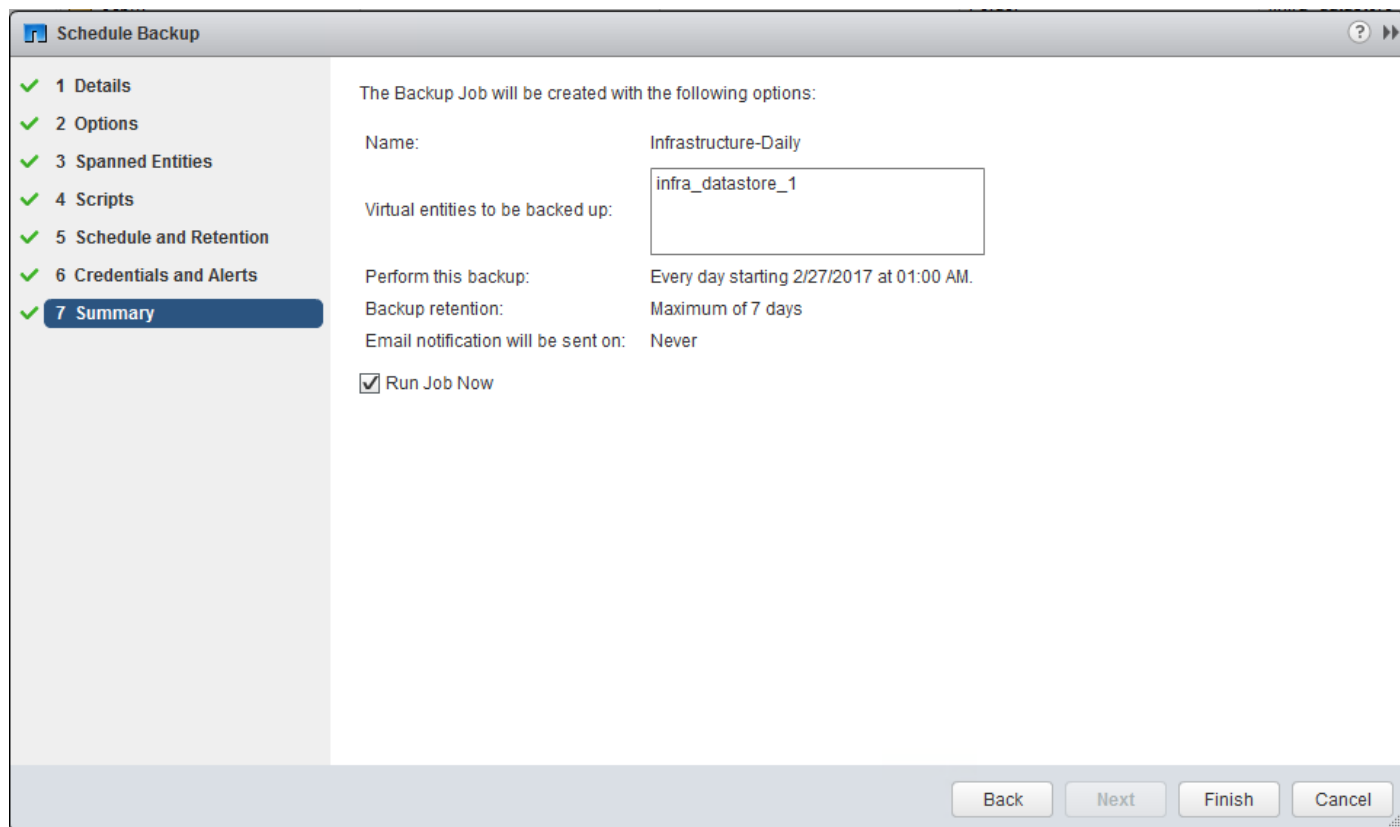
Maximum Days: 7

Maximum Backups: 1

Backups Never Expire


Back Next Finish Cancel

10. Use the default vCenter credentials or enter the user name and password for the vCenter server. Click Next.
11. Specify any needed backup notification details. Enter an e-mail address and mail server address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate them. Click Next.



12. 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.

13. Click OK.



14. You can also create other backup jobs with overlapping schedules. For example, you can create weekly or monthly backups that overlay daily backups.

15. On the storage cluster interface, automatic Snapshot copies of the volume can now be disabled because NetApp VSC is now handling scheduled backups. To do so, enter the following command:

```
volume modify -vserver Infra-SVM -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 -vserver Infra-SVM -volume infra_datastore_1
volume snapshot delete -vserver Infra-SVM -volume infra_datastore_1 -snapshot <snapshot-name>
```




The wildcard character * can be used in Snapshot names in the previous command.

Sample Tenant Provisioning

Provisioning a Sample Application Tenant

This section describes a sample procedure for provisioning an application tenant. The procedure here refers back to previous sections of this document and can be used as a guide and modified as needed when provisioning an application tenant.

1. Plan your application tenant and determine what storage protocols will be provided in the tenant. In the architecture covered in this document, NFS, iSCSI, and CIFS/SMB can be provided to the tenant. Also, plan what network VLANs the tenant will use. It is recommended to have a VLAN for virtual machine management traffic. One or two VLANs (iSCSI needs two if VMware RDM LUNs or iSCSI datastores will be provisioned) are also needed for each storage protocol used.
2. In the Nexus switches, declare all added VLANs and configure the VM VLAN as an allowed VLAN on the Cisco UCS port channels and the vPC peer link. Also, Layer 3 with HSRP or VRRP can be configured in the Nexus switches to provide this VLAN access to the outside. Layer 3 setup is not covered in this document, but is covered in the Cisco Nexus 9000 documentation. Configure the storage VLANs on the Cisco UCS and storage port channels, and on the vPC peer link. The VM VLAN can also be added to the storage port channels in order to configure the tenant SVM management interface on this VLAN.
3. In the storage cluster:
 - a. Create a broadcast domain with MTU 1500 for the tenant SVM management interface. Create a broadcast domain with MTU 9000 for each tenant storage protocol except fiber channel.
 - b. Create VLAN interface ports on the node interface group on each node for tenant SVM management (VM VLAN) and for the VLAN for each storage protocol. Add these VLAN ports to the appropriate broadcast domains.
 - c. Create the tenant SVM and follow all procedures in that section.
 - d. Create Load-Sharing Mirrors for the tenant SVM.
 - e. Create the iSCSI service for the tenant SVM if iSCSI is being deployed in this tenant.
 - f. Optionally, create a self-signed security certificate for the tenant SVM.
 - g. Configure NFSv3 for the tenant SVM.
 - h. Create a VM datastore volume in the tenant SVM.
 - i. Create a once-a-day deduplication schedule on the VM datastore volume.
 - j. If iSCSI is being deployed in this tenant, configure four iSCSI LIFs in the tenant SVM on the appropriate VLAN interface ports.
 - k. Create one NFS LIF in the tenant SVM on each storage node.
 - l. Create a boot LUN in the esxi_boot volume in the Infra-SVM for each tenant VMware ESXi host.

- m. Add the tenant SVM Administrator, SVM management LIF on the SVM management VLAN port, and default route for the SVM.
4. In Cisco UCS, one method of tenant setup is to dedicate a VMware ESXi cluster and set of Cisco UCS servers to each tenant. Service profiles will be generated for at least two tenant ESXi hosts. These hosts can boot from LUNs from the esxi_boot volume in the Infra-SVM, but will also have access to FC storage in the tenant SVM.
 - a. Create a Server Pool for the tenant ESXi host servers.
 - b. Create all tenant VLANs in the LAN Cloud.
 - c. Add the tenant VLANs to the Infra vNIC templates.
 - d. Generate service profiles from the service profile template with the vMedia policy for the tenant ESXi hosts. Remember to bind these service profiles to the service profile template without the vMedia policy after VMware ESXi installation.
5. In the storage cluster:
 - a. Create igroups for the tenant ESXi hosts in both the Infra-SVM and tenant SVM. Also, create an igroup in the tenant SVM that includes the IQNs for all tenant ESXi hosts to support shared storage from the tenant SVM.
 - b. In Infra-SVM, map the boot LUNs created earlier to the tenant ESXi hosts. Tenant iSCSI storage can be created later using either NetApp VSC or NetApp SnapDrive.
6. Install and configure VMware ESXi on all tenant host servers. It is not necessary to map infra_datastore_1.
7. In VMware vCenter, create a cluster for the tenant ESXi hosts. Add the hosts to the cluster.
8. Using the vCenter Web Client, add the tenant hosts to the VMware vDS. In the VMware vDS, add port-groups for the tenant VLANs. If iSCSI port-groups are created, pin them to the appropriate up-link for the UCS Fabric they are connected to.
9. Back in vCenter, add in any necessary VMkernel ports for storage interfaces remembering to set the MTU correctly on these interfaces. Mount the tenant NFS datastore on the tenant cluster if one was created.
10. Using the NetApp VSC plugin to the vCenter Web Client, set recommended values for all tenant ESXi hosts. Install the NetApp NFS Plug-in for VMware VAAI for all tenant hosts and reboot each host. Optionally, create a VSC backup job for the tenant NFS datastore.
11. You can now begin provisioning virtual machines on the tenant cluster. The NetApp VSC plugin can be used to provision both fiber channel and NFS datastores.

Appendix

FlexPod Backups

Cisco UCS Backup

Automated backup of the Cisco UCS domain is important for recovery of the Cisco UCS Domain from issues ranging from catastrophic failure to human error. There is a native backup solution within Cisco UCS that allows local or remote backup using FTP/TFTP/SCP/SFTP as options, and is detailed below.

Backups created can be a binary file containing the Full State, which can be used for a restore to the original or a replacement pair of fabric interconnects. Alternately this XML configuration file consisting of All configurations, just System configurations, or just Logical configurations of the Cisco UCS Domain. For scheduled backups, options will be Full State or All Configuration, backup of just the System or Logical configurations can be manually initiated.

To schedule the backup, complete the following steps within the Cisco UCS Manager GUI:

1. Select Admin within the Navigation pane and select All.
2. Click the Policy Backup & Export tab within All.
3. For a Full State Backup, All Configuration Backup, or both, specify the following:
 - a. Hostname : <IP or FQDN of host that will receive the backup>
 - b. Protocol: [FTP/TFTP/SCP/SFTP]
 - c. User: <account on host to authenticate>
 - d. Password: <password for account on host>
 - e. Remote File: <full path and filename prefix for backup file>
 - f. Admin State: <select Enable to activate the schedule on save, Disable to disable schedule on save>
 - g. Schedule: [Daily/Weekly/Bi Weekly]

4. Click Save Changes to create the Policy.

Cisco Nexus Backups

The configuration of the Cisco Nexus 9000 switches can be backed up manually at any time with the copy command, but automated backups can be put in place with the NX-OS feature scheduler. An example of setting up automated configuration backups of one of the FlexPod 9332PQ switches is shown below:

```
bb04-9332-a# conf t
Enter configuration commands, one per line. End with CNTL/Z.
bb04-9332-a(config)# feature scheduler
bb04-9332-a(config)# scheduler logfile size 1024
bb04-9332-a(config)# scheduler job name backup-cfg
bb04-9332-a(config-job)# copy running-config
tftp://192.168.156.155/9332/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
bb04-9332-a(config-job)# exit
bb04-9332-a(config)# scheduler schedule name daily
bb04-9332-a(config-schedule)# job name backup-cfg
bb04-9332-a(config-schedule)# time daily 2:00
```

```
bb04-9332-a (config-schedule) # end
```

Show the job that has been setup:

```
bb04-9332-a# sh scheduler job
Job Name: backup-cfg
-----
copy running-config tftp://192.168.156.155/9332/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf
management
=====
```

```
bb04-9332-a# show scheduler schedule
Schedule Name      : daily
-----
User Name          : admin
Schedule Type      : Run every day at 2 Hrs 0 Mins
Last Execution Time : Sun Apr  9 02:00:00 2017
Last Completion Time: Sun Apr  9 02:00:01 2017
Execution count    : 3
-----
```

Job Name	Last Execution Status
backup-cfg	Success (0)

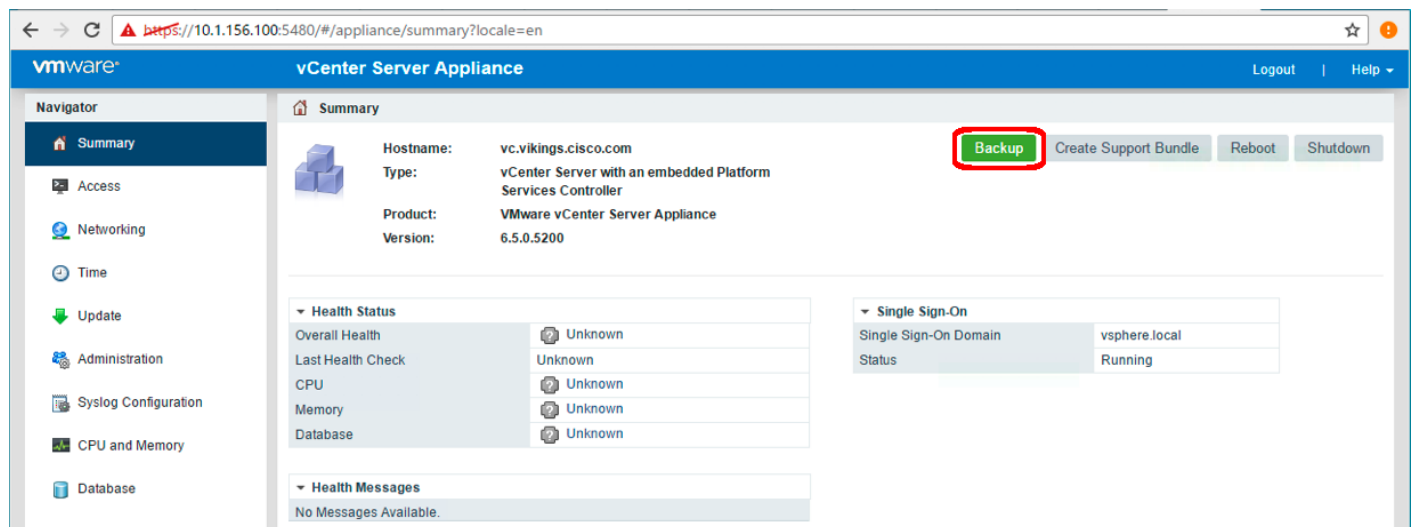
For detailed information about the scheduler, refer to:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_01010.html

VMware VCSA Backup

Basic backup of the vCenter Server Appliance is also available within the native capabilities of the VCSA, though within the default solution this is manually initiated for each backup operation. To create a backup, complete the following steps:

1. Connect to the VCSA Console at <https://<VCSA IP>:5480>



2. Click Backup within the Summary section to open up the Backup Appliance Dialogue.
3. Specify:
 - a. the Protocol to use [HTTPS/HTTP/SCP/FTPS/FTP]
 - b. location of an empty directory to be used for the backup
 - c. the User name and password

The screenshot shows the 'Backup Appliance' dialog box. The 'Enter backup details' section is active, showing the following fields:

- Protocol:** SCP
- Location:** 10.1.156.150/var/www/html/bears/configs/ucs
- Port:** 22
- User name:** root
- Password:** (masked with dots)

There is an unchecked checkbox for **Encrypt Backup Data**.

Navigation buttons at the bottom are: Back, Next, Finish, and Cancel.

4. Click Next.

- De-select some parts if they should be excluded from the backup.

Backup Appliance

✓ 1 Enter backup details

2 Select parts to backup

3 Ready to complete

Select parts to backup

Select files you want to backup and optionally provide a description for your backup.

A minimum set of data needed to restore the appliance will be backed up by default. This includes data such as OS, VC services and Inventory. In addition to this, you can also choose to backup additional parts below.

Parts

<input checked="" type="checkbox"/>	common	Inventory and configuration.	532 MB
<input checked="" type="checkbox"/>	Stats, Events, Alarms, and Tasks	Historical data (Statistics, Events and Tasks) in vCenter Server database.	34 MB

Description:

Back Next Finish Cancel

- Click Next.

- Review the options selected and click Finish to begin the backup.



Restoration can be initiated with the backed-up files using the Restore function of the VCSA 6.5 Installer.

Breakout Interface Configuration in the Cisco Nexus 9332PQ Switches

The 40Gb end to end FlexPod design in this document uses a pair of Nexus 9332PQ which is built with all ports being of the 40 Gbps Quad Small Form Factor Pluggable Plus (QSFP+) type. If there is a need to directly support a 10Gb Small Form Pluggable Plus (SFP+), this can be configured within the switch, and connected to the 10Gb SFP+ device using a supported QSFP+ Breakout Cable.

Configuring the QSFP+ ports uses the interface breakout command as shown in this example to turn the 40G interface Ethernet 1/1 into 4x10G interfaces:

```
bb04-9332-a(config)# show running-config interface Ethernet1/1

interface Ethernet1/1
  no switchport

bb04-9332-a(config)# interface breakout module 1 port 5 map 10g-4x

bb04-9332-a(config)# show running-config interface Ethernet1/1/1-4
```



```
interface Ethernet1/1/1
interface Ethernet1/1/2
interface Ethernet1/1/3
interface Ethernet1/1/4
```

Breakout configurations that are no longer needed can be reverted with the no interface breakout command:

```
bb04-9332-a(config)# no interface breakout module 1 port 1 map 10g-4x
bb04-9332-a(config)#
```

About the Authors

John George, Technical Marketing Engineer, Cisco UCS Data Center Solutions Engineering, Cisco Systems, Inc.

John George recently moved to Cisco from NetApp and is focused on designing, developing, validating, and supporting the FlexPod Converged Infrastructure since its inception. Before his role with FlexPod, he supported and administered a large worldwide training network and VPN infrastructure. John holds a Master's degree in Computer Engineering from Clemson University.

Karthick Radhakrishnan, Systems Architect, Converged Infrastructure Engineering, NetApp

Karthick Radhakrishnan is a Systems Architect in the NetApp Infrastructure and Cloud Engineering team. He focuses on the validating, supporting, implementing cloud infrastructure solutions that include NetApp products. Prior to his current role, he was a networking tools developer at America Online supporting AOL transit data network. Karthick worked in the IT industry for more than 14 years and he holds Master's degree in Computer Application.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Ramesh Isaac, Cisco Systems, Inc.
- Aaron Kirk, NetApp