

FlexPod Datacenter with VMware vSphere 6.7 U2, Cisco UCS 4th Generation Fabric and NetApp ONTAP 9.6

Deployment Guide for FlexPod Datacenter with NetApp ONTAP 9.6 on NetApp AFF A-Series, Cisco UCS M5 Servers with 2nd Generation Intel Xeon Scalable Processors, and VMware vSphere 6.7 U2

Published: February 2020



In partnership with: **NetApp**

About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

<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)

© 2020 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	9
Solution Overview	10
Introduction.....	10
Audience	10
Purpose of this Document	10
What's New in this Release?	10
Deployment Hardware and Software	12
Architecture	12
Topology	13
Software Revisions.....	13
Configuration Guidelines	14
Physical Infrastructure	16
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.....	20
Create VLANs	21
Add NTP Distribution Interface	21
Add Individual Port Descriptions for Troubleshooting and Enable UDLD for UCS Interfaces	21
Create Port Channels.....	23
Configure Port Channel Parameters	23
Configure Virtual Port Channels	24
Uplink into Existing Network Infrastructure	25
Switch Testing Commands	26
Storage Configuration	27
NetApp All Flash FAS A320 Controllers.....	27
NetApp Hardware Universe	27
Controllers	27
Disk Shelves	27
NetApp ONTAP 9.6	27
Complete Configuration Worksheet	27
Configure ONTAP Nodes	28
Set Up Node	31
Log into the Cluster.....	39
Verify Storage Failover	39
Set Auto-Revert on Cluster Management.....	39

Zero All Spare Disks.....	40
Set Up Service Processor Network Interface	40
Create Aggregates.....	40
Remove Ports from Default Broadcast Domain	41
Disable Flow Control on 10GbE and 100GbE ports	41
Disable Auto-Negotiate on 100GbE Ports	41
Disable Auto-Negotiate on Fibre Channel Ports	42
Enable Cisco Discovery Protocol	42
Enable Link-layer Discovery Protocol on all Ethernet Ports.....	42
Create Management Broadcast Domain	43
Create NFS Broadcast Domain	43
Create Interface Groups	43
Change MTU on Interface Groups	43
Create VLANs	43
Configure Network Time Protocol	44
Configure Simple Network Management Protocol	44
Configure SNMPv1 Access.....	45
Create SVM	45
Create Load-Sharing Mirrors of SVM Root Volume	45
Create Block Protocol (FC) Service	46
Configure HTTPS Access.....	46
Configure NFSv3.....	47
Create FlexVol Volumes.....	48
Create Boot LUNs.....	48
Modify Volume Efficiency	48
Create FC LIFs.....	48
Create NFS LIFs.....	49
Add Infrastructure SVM Administrator	49
Configure and Test AutoSupport.....	50
Cisco UCS Configuration	51
Cisco UCS Base Configuration.....	51
Perform Initial Setup of Cisco UCS 6454 Fabric Interconnects for FlexPod Environments.....	51
Cisco UCS Setup.....	52
Log into Cisco UCS Manager.....	52
Upgrade Cisco UCS Manager Software to Version 4.0(4g)	53
Anonymous Reporting.....	53
Configure Cisco UCS Call Home	53
Configure Unified Ports (FCP).....	54
Add Block of IP Addresses for KVM Access	56
Synchronize Cisco UCS to NTP	56
Add Additional DNS Server(s).....	57

Add an Additional Administrative User.....	58
Enable Port Auto-Discovery Policy	59
Enable Info Policy for Neighbor Discovery.....	60
Edit Chassis Discovery Policy.....	61
Enable Server and Uplink Ports.....	61
Acknowledge Cisco UCS Chassis and FEX	62
Create Uplink Port Channels to Cisco Nexus Switches	63
Add UDLD to Uplink Port Channels.....	66
Create an Organization	68
Create a WWNN Pool for FC Boot (FCP)	69
Create WWPN Pools (FCP)	71
Create VSANs (FCP).....	74
Enable FC Uplink VSAN Trunking (FCP).....	75
Create FC Uplink Port Channels (FCP)	76
Disable Unused FC Uplink Ports (FCP)	79
Create vHBA Templates (FCP)	79
Create SAN Connectivity Policy (FCP).....	81
Create MAC Address Pools	83
Create UUID Suffix Pool.....	85
Create Server Pool	86
Create VLANs	86
Modify Default Host Firmware Package	89
Set Jumbo Frames in Cisco UCS Fabric.....	90
Create Local Disk Configuration Policy (Optional)	91
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	93
Create Power Control Policy.....	93
Create Server BIOS Policy	95
Update the Default Maintenance Policy	100
Create vNIC Templates.....	101
Create High Traffic VMware Adapter Policy	105
Create LAN Connectivity Policy for FC Boot (FCP).....	107
Create vMedia Policy for VMware ESXi 6.7U2 ISO Install Boot.....	111
Create FC Boot Policy (FCP)	113
Create Service Profile Template (FCP)	118
Configure Storage Provisioning.....	119
Configure Networking	119
Configure SAN Connectivity.....	120
Configure Zoning	121
Configure vNIC/HBA Placement.....	121
Configure vMedia Policy.....	121
Configure Server Boot Order	122

Configure Maintenance Policy.....	122
Configure Server Assignment.....	123
Configure Operational Policies.....	124
Create vMedia-Enabled Service Profile Template.....	125
Create Service Profiles.....	126
Add More Servers to FlexPod Unit.....	126
Gather Necessary Information.....	127
SAN Switch Configuration.....	128
Physical Connectivity.....	128
FlexPod Cisco MDS Base.....	128
Cisco MDS 9132T A.....	128
FlexPod Cisco MDS Switch Configuration.....	131
Enable Licenses.....	131
Add Second NTP Server.....	131
Configure Individual Ports.....	132
Create VSANs.....	133
Create Device Aliases.....	134
Create Zones and Zoneset.....	135
Storage Configuration – Boot LUNs.....	136
ONTAP Boot Storage Setup.....	136
Create igroups.....	136
Map Boot LUNs to igroups.....	136
VMware vSphere 6.7U2 Setup.....	137
VMware ESXi 6.7U2.....	137
Download ESXi 6.7U2 from VMware.....	137
Log into Cisco UCS 6454 Fabric Interconnect.....	137
Set Up VMware ESXi Installation.....	138
Install ESXi.....	138
Set Up Management Networking for ESXi Hosts.....	139
Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional).....	141
Install VMware Drivers for the Cisco Virtual Interface Card (VIC) and ESXi Host.....	142
Log into the First VMware ESXi Host by Using VMware Host Client.....	142
Set Up VMkernel Ports and Virtual Switch.....	143
Mount Required Datastores.....	145
Configure NTP on First ESXi Host.....	147
Configure ESXi Host Swap.....	148
VMware vCenter 6.7U2.....	149
Build the VMware vCenter Server Appliance.....	149
Adjust vCenter CPU Settings.....	160
Setup VMware vCenter Server.....	162
Add AD User Authentication to vCenter (Optional).....	166

FlexPod VMware vSphere Distributed Switch (vDS).....	167
Configure the VMware vDS in vCenter	168
Add and Configure a VMware ESXi Host in vCenter.....	172
Add the ESXi Host to vCenter	172
Set Up VMkernel Ports and Virtual Switch	173
Mount Required Datastores.....	176
Configure NTP on ESXi Host.....	176
Configure ESXi Host Swap.....	178
Check ESXi Host Fibre Channel Pathing	178
Add the ESXi Host(s) to the VMware Virtual Distributed Switch	179
Add the vMotion VMkernel Port to the ESXi Host.....	180
VMware ESXi 6.7U2 TPM Attestation	181
FlexPod Management Tools Setup	183
Cisco UCS Manager Plug-in for VMware vSphere Web Client	183
Cisco UCS Manager Plug-in Installation.....	183
FlexPod UCS Domain Registration	185
Register the Proactive High Availability Provider for Cisco UCS (Optional)	186
Use the Cisco UCS vCenter Plugin.....	188
NetApp Virtual Storage Console 9.6 Deployment Procedure	191
Virtual Storage Console 9.6 Pre-installation Considerations.....	191
The requirements for deploying VSC are listed here.	191
Install Virtual Storage Console 9.6	191
Download the NetApp NFS VAAI Plug-In.....	194
Install the NetApp NFS VAAI Plug-In.....	194
Enable the VASA Provider	195
Discover and Add Storage Resources	196
Optimal Storage Settings for ESXi Hosts.....	198
Virtual Storage Console 9.6 Provisioning Datastores	199
NetApp SnapCenter 4.2	206
Deploy NetApp SnapCenter	206
SnapCenter Server Requirements	207
SnapCenter 4.2 License Requirements	208
SnapCenter Server	209
SnapCenter Plug-In for VMware vSphere	209
Support for Virtualized Databases and File Systems.....	209
SnapCenter Installation	209
SnapCenter Configuration.....	209
Install SnapCenter Plug-In for VMware.....	211
Host and Privilege Requirements for the NetApp Data Broker	211
Download and Deploy the NetApp Data Broker OVA (Open Virtual Appliance).....	212
Create a Virtual Appliance Credential for Migrating Backups	217

Register the SnapCenter Plug-in for VMware vSphere with SnapCenter Server	218
Register the SnapCenter Plug-in for VMware vSphere with VMware Vcenter.....	218
SnapCenter Plug-in for VMware vSphere in VCenter Server	219
Add Storage Systems (SVM)	220
View Virtual Machine Backups and Restore from vCenter by Using SnapCenter Plug-In	228
View Backups.....	228
Restore from vCenter by Using SnapCenter Plug-In.....	231
Provision Virtual Volumes (VVols)	235
Verify NDMP Vserver Scope Mode	235
OnCommand API Services for ONTAP VASA Provider	236
Create the Storage Capability Profile	239
Provision VVOL from Virtual Storage Console.....	240
Create Virtual Machine with Assigned VM Storage Policy.....	243
VSC VVOI Dashboard and Reports	244
Active IQ Unified Manager 9.6	245
Configure Active IQ Unified Manager	251
Add Storage to Active IQ Unified Manager.....	252
Sample Tenant Provisioning.....	255
Provision a Sample Application Tenant.....	255
Appendix.....	257
FlexPod iSCSI Addition	257
Cisco Nexus Switch Configuration	257
NetApp Storage Configuration – Part 1	257
Cisco UCS iSCSI Configuration.....	259
NetApp Storage Configuration – Part 2	281
VMware vSphere Configuration	282
ESXi Dump Collector Setup for iSCSI-Booted Hosts	297
Create a FlexPod ESXi Custom ISO using VMware vCenter	297
FlexPod Backups	302
Cisco UCS Backup.....	302
Cisco Nexus and MDS Backups.....	305
VMware VCSA Backup.....	306
About the Authors.....	309
Acknowledgements	309



Executive Summary

Cisco Validated Designs (CVDs) 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 NetApp ONTAP 9.6, Cisco UCS Manager unified software release 4.0(4) with 2nd Generation Intel Xeon Scalable Processors and VMware vSphere 6.7 U2. Cisco UCS Manager (UCSM) 4.0(4) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 6454, 2200/2300/2400 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with NetApp ONTAP 9.6, Cisco UCS unified software release 4.0(4), and VMware vSphere 6.7 U2 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays running ONTAP® 9.6 storage OS.

Solution Overview

Introduction

The current industry trend in datacenter 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 storage, Cisco MDS, and Cisco Nexus 9000 solution.

What's New in this Release?

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

- Support for the Cisco UCS 4.0(4) unified software release, Cisco UCS B200-M5 and C220-M5 servers with 2nd Generation Intel Xeon Scalable Processors, and Cisco 1400 Series Virtual Interface Cards (VICs)
- Support for the latest Cisco UCS 6454 Fabric Interconnect
- Support for the latest Cisco UCS 2408 Fabric Extender
- Support for the NetApp AFF A320 Storage Controller
- Support for the latest release of NetApp ONTAP® 9.6
- Support for NetApp Virtual Storage Console (VSC) 9.6
- Support for NetApp SnapCenter 4.2 and NetApp Data Broker 1.0
- Support for NetApp Active IQ Unified Manager 9.6
- Fibre channel, NFS, iSCSI (appendix) storage design
- Validation of VMware vSphere 6.7 U2
- Unified Extensible Firmware Interface (UEFI) Secure Boot of VMware ESXi 6.7 U2
- Trusted Platform Module (TPM) 2.0 Attestation of UEFI Secure Boot of VMware ESXi 6.7 U2
- 100 Gigabit per second Ethernet Connectivity

- 32 Gigabit per second Fibre Channel Connectivity

Deployment Hardware and Software

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, Cisco MDS storage 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 a Fibre Channel and 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 6454 Fabric Interconnects. This design has port-channelled 25 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects via the Cisco UCS 2408 Fabric Extenders, port-channelled 25 Gb Ethernet connections between the C-Series rackmounts and the Cisco UCS Fabric Interconnects, and 100 Gb Ethernet connections between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A320 storage array. This infrastructure option expanded with Cisco MDS switches sitting between the Cisco UCS Fabric Interconnect and the NetApp AFF A320 to provide FC-booted hosts with 32 Gb FC 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.

Topology

Figure 1 FlexPod with Cisco UCS 6454 Fabric Interconnects and NetApp AFF A320

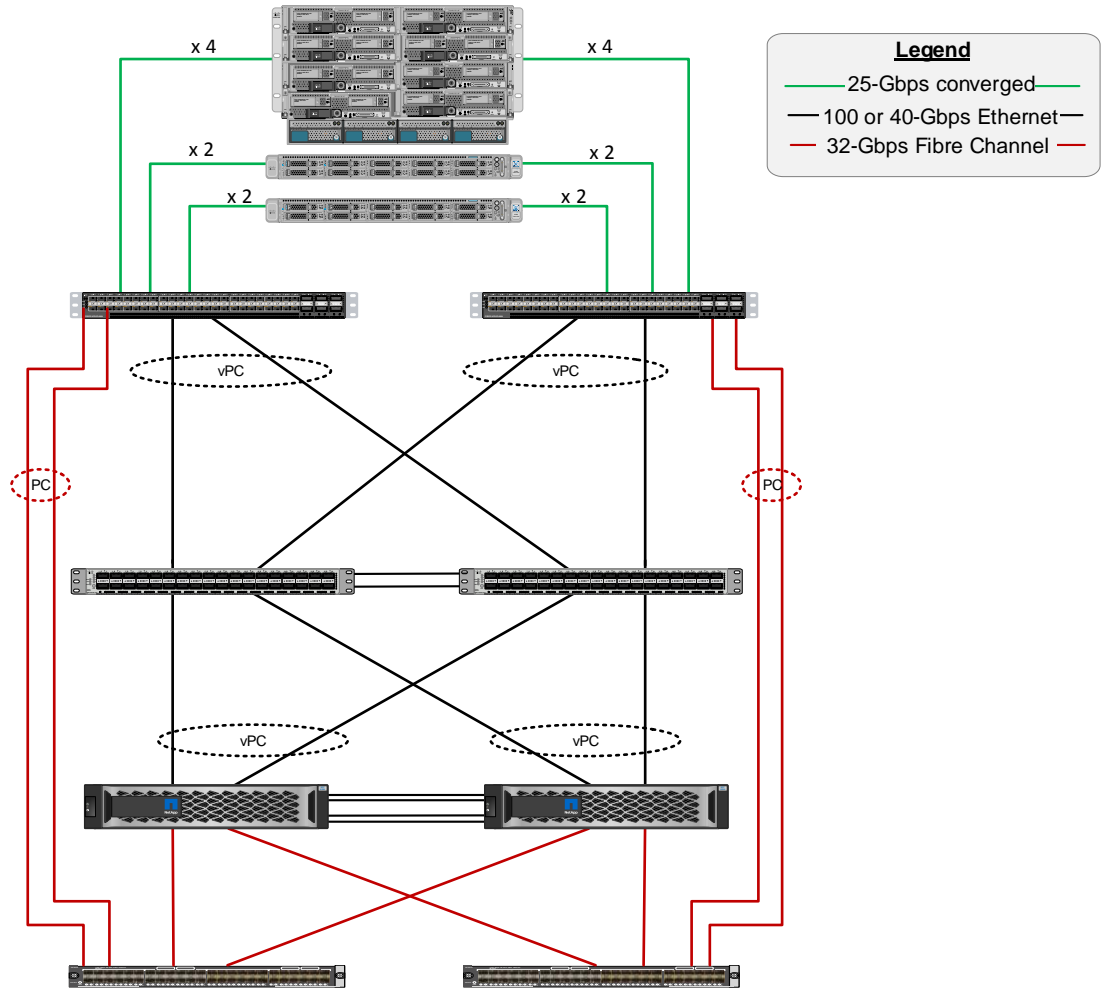
Cisco Unified Computing System

Cisco UCS 6454 Fabric Interconnects, UCS 2408 Fabric Extenders, UCS B-Series Blade Servers with UCS VIC 1440, and UCS C-Series Rack Servers with UCS VIC 1457

Cisco Nexus 9336C-FX2

NetApp storage controllers AFF-A320

Cisco MDS 9132T or 9148T switch



The reference 100Gb based hardware configuration includes:

- Two Cisco Nexus 9336c-FX2 switches
- Two Cisco UCS 6454 fabric interconnects
- Two Cisco MDS 9132T multilayer fabric switches
- One NetApp AFF 320 (HA pair) running ONTAP 9.6 with internal NVMe SSD disks

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
-------	--------	-------	----------

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6454, Cisco UCS B200 M5 and Cisco UCS C220 M5 with 2 nd Generation Intel Xeon Scalable Processors	4.0(4g)	Includes the Cisco UCS 2408 Fabric Extender, Cisco UCS Manager, Cisco UCS VIC 1440 and Cisco UCS VIC 1457
Network	Cisco Nexus 9336C-FX2 NX-OS	7.0(3)I7(7)	
	Cisco MDS 9132T	8.4(1)	
Storage	NetApp AFF 320	ONTAP 9.6	
Software	Cisco UCS Manager	4.0(4g)	
	Cisco UCS Manager Plugin for VMware vSphere Web Client	2.0.5	
	VMware vSphere	6.7U2	
	VMware ESXi nfnic FC Driver	4.0.0.38	
	VMware ESXi nenic Ethernet Driver	1.0.29.0	
	NetApp Virtual Storage Console (VSC) / VASA Provider Appliance	9.6	
	NetApp NFS Plug-in for VMware VAAI	1.1.2	
	NetApp SnapCenter Backup Management	4.2	Includes the vSphere plug-in for SnapCenter
	NetApp Data Broker	1.0	Required for SnapCenter
	NetApp Active IQ Unified Manager	9.6	

Configuration Guidelines

This document explains how to configure a fully redundant, highly available configuration for a FlexPod unit with 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 create -node <node01> -vlan-name a0a-<vlan id>
```

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

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
Infra-NFS	VLAN for Infrastructure NFS traffic	3050
FCoE-A	VLAN for FCoE encapsulation of VSAN-A	103
FCoE-B	VLAN for FCoE encapsulation of VSAN-B	104
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	IP Address
vCenter Server		
NetApp VSC		
NetApp SnapCenter		
NetApp Data Broker		
Active IQ Unified Manager		

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, a cabling diagram was used.

The cabling diagram in this section contains details for the prescribed and supported configuration of the NetApp AFF 320 running NetApp ONTAP® 9.6.



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.

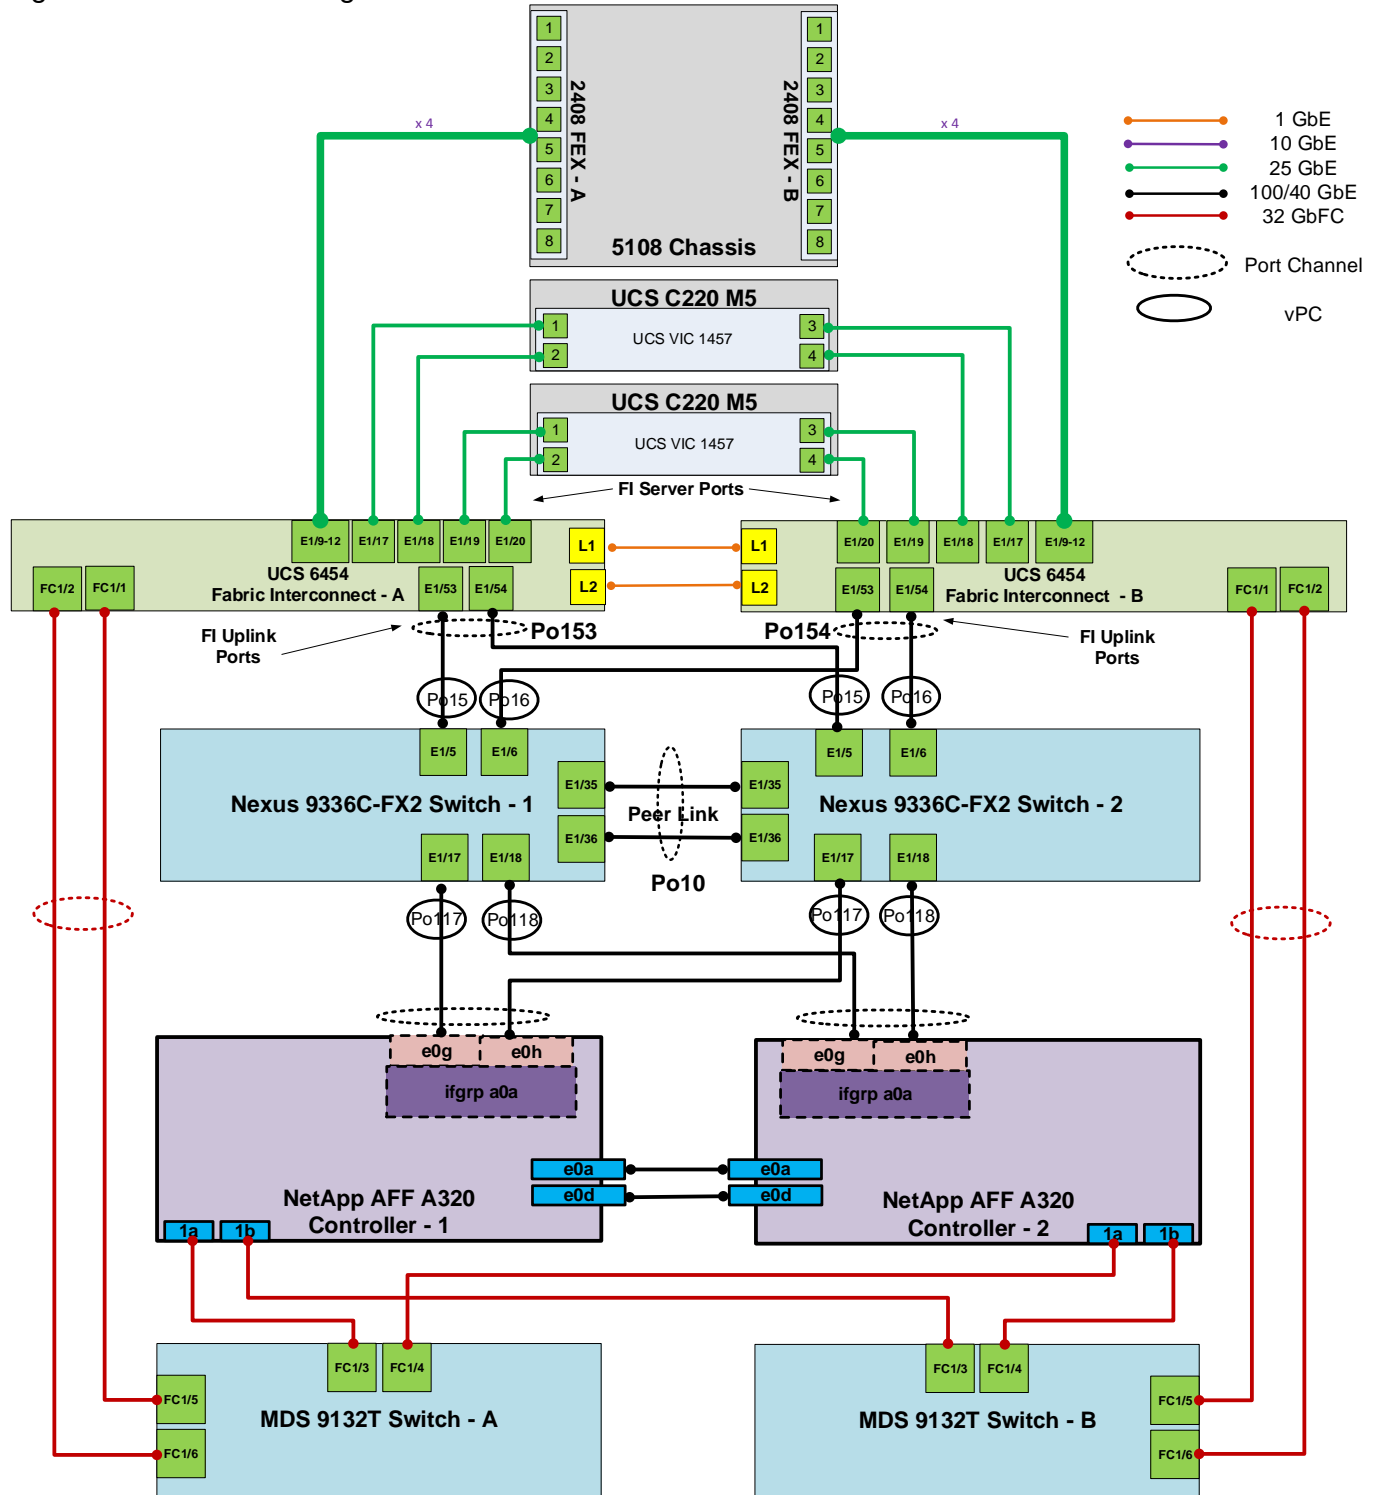


Be sure to use the cabling directions in this section as a guide.

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 [NetApp Support](#).

Figure 2 details the cable connections used in the validation lab for the FlexPod topology based on the Cisco UCS 6454 fabric interconnect. Two 32Gb uplinks connect as port-channels to each Cisco UCS Fabric Interconnect from the MDS switches, and a total of four 32Gb links connect the MDS switches to the NetApp AFF controllers. Also, 100Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the NetApp AFF controllers to the Cisco Nexus Switches. 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. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

Figure 2 FlexPod Cabling with Cisco UCS 6454 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 explained in [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)I7(7), the Cisco suggested Nexus switch release at the time of this validation.



The following procedure includes the setup of NTP distribution on both the mgmt0 port and 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.



With Cisco Nexus 9000 release 7.0(3)I7(7), autonegotiation (40G/100G) is not supported on ports 1-6 and 33-36 on the Cisco Nexus 9336C-FX2 switch. Since those ports are in use in this validation, port speed and duplex are hard set at both ends of the connection.

Set Up Initial Configuration

Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>, follow these 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 [yes - continue with normal setup, skip - bypass
password and basic configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

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

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) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus B

To set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>, follow these 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 [yes - continue with normal setup, skip - bypass
password and basic configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

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

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) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

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

FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus A and Cisco Nexus B

To license the Cisco Nexus switches, follow these steps:

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

```

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

```

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

To set global configurations, follow this step on both switches:

1. Run the following commands to set global configurations:

```

spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <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), follow this step on both switches:

1. 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
exit
```

Add NTP Distribution Interface

Cisco Nexus A

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

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-b-ntp-ip> use-vrf default
```

Cisco Nexus B

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

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-a-ntp-ip> use-vrf default
```

Add Individual Port Descriptions for Troubleshooting and Enable UDLD for UCS Interfaces

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, follow this step:



In this step and in the following sections, configure the AFF nodename <st-node> and Cisco UCS 6454 fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

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

```
interface Eth1/5
description <ucs-clustername>-a:1/53
udld enable
interface Eth1/6
description <ucs-clustername>-b:1/53
udld enable
```



For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected.

```
interface Eth1/17
description <st-clustername>-1:e0g
interface Eth1/18
description <st-clustername>-2:e0g
interface Eth1/35
description <nexus-b-hostname>:1/35
interface Eth1/36
description <nexus-b-hostname>:1/36
exit
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B and to enable aggressive UDLD on copper interfaces connected to Cisco UCS systems, follow this step:

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

```
interface Eth1/5
description <ucs-clustername>-a:1/54
udld enable
interface Eth1/6
description <ucs-clustername>-b:1/54
udld enable
```



For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected.

```
interface Eth1/17
description <st-clustername>-1:e0h
interface Eth1/18
description <st-clustername>-2:e0h
interface Eth1/35
description <nexus-a-hostname>:1/35
interface Eth1/36
description <nexus-a-hostname>:1/36
exit
```



For fibre optic connections to Cisco UCS systems (AOC or SFP-based), enter `udld enable` when configuring the interface.

Create Port Channels

Cisco Nexus A and Cisco Nexus B

To create the necessary port channels between devices, follow this step on both switches:

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

```
interface Po10
description vPC peer-link
interface Eth1/35-36
channel-group 10 mode active
no shutdown
interface Po117
description <st-clustername>-1
interface Eth1/17
channel-group 117 mode active
no shutdown
interface Po118
description <st-clustername>-2
interface Eth1/18
channel-group 118 mode active
no shutdown
interface Po15
description <ucs-clustername>-a
interface Eth1/5
channel-group 15 mode active
no shutdown
interface Po16
description <ucs-clustername>-b
interface Eth1/6
channel-group 16 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, follow this step on both switches:

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

```
interface Po10
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-
id>, <vm-traffic-vlan-id>
spanning-tree port type network
speed 100000
duplex full
no negotiate auto

interface Po117
switchport mode trunk
```

```

switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto

interface Po118
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto

interface Po15
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-
id>, <vm-traffic-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto

interface Po16
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-
id>, <vm-traffic-vlan-id>
spanning-tree port type edge trunk
mtu 9216
speed 100000
duplex full
no negotiate auto
exit
copy run start

```



Lab testing confirmed that the speed and duplex needed to be hard set on both ends for the NetApp storage interfaces even though these interfaces were not on ports 1-6 or 33-36.

Configure Virtual Port Channels

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, follow this step:

1. 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
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po118
vpc 118
interface Po15
vpc 15
interface Po16
vpc 16
exit
copy run start

```

Cisco Nexus B

To configure vPCs for switch B, follow this step:

1. 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
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po118
vpc 118
interface Po15
vpc 15
interface Po16
vpc 16
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.

Switch Testing Commands

The following commands can be used to check for correct switch configuration:



Some of these commands need to run after further configuration of the FlexPod components are complete to see complete results.

```
show run
show vpc
show port-channel summary
show ntp peer-status
show cdp neighbors
show lldp neighbors
show run int
show int
show udld neighbors
```

Storage Configuration

NetApp All Flash FAS A320 Controllers

See the following section ([NetApp Hardware Universe](#)) 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
- AFF Series Systems

NetApp Hardware Universe

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

To 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, follow these steps at the [NetApp Support](#) site.

1. Access the [HWU application](#) to view the System Configuration guides. Click the Platforms menu 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 A320 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 A320 is available at the [NetApp Support](#) site.

When using SAS disk shelves with NetApp storage controllers, refer to the [SAS cabling rules section in the AFF and FAS System Documentation Center](#) for proper cabling guidelines.

NetApp ONTAP 9.6

Complete Configuration Worksheet

Before running the setup script, complete the [Cluster setup worksheet](#) in the ONTAP 9 Documentation Center. 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 [Software setup section of the ONTAP 9 Documentation Center](#) 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>
ONTAP 9.6 URL	<url-boot-software>

Configure Node 01

To configure node 01, follow these 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 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 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, follow these 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 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 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 9.6 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 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 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>
NTP server B IP address	<switch-b-ntp-ip>



Cluster setup can also be performed using the CLI. This document describes the cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup on the Welcome screen.

Cluster Setup Workflow - X

Not Secure | <https://192.168.156.61/sysmgr/SysMgr.html>

NetApp OnCommand System Manager

Getting Started

Language English (English)

Welcome to the Guided Cluster Setup

Perform the following to set up a cluster:

- Create a cluster, add nodes and admin credentials
- Create management LIFs, configure Service Processor, DNS, and NTP servers
- Configure AutoSupport Messages and Event Notifications


[For information related to setting up the cluster, click here](#)

Template File

Browse to select a .csv file...

[To download the template, click file.csv or file.xlsx](#)

Important: You can download the template in ".csv" or ".xlsx" format. However, you can upload only those templates that are in ".csv" format.


Guided Setup

Click to set up the cluster

4. In the Cluster screen, follow these steps:
 - a. Enter the cluster and node names.
 - b. Select the cluster configuration.
 - c. Enter and confirm the password.
 - d. (Optional) Enter the cluster base and feature licenses.

ONTAP System Manager

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



Cluster Name

Nodes

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



Username

Password

Confirm Password

Not sure of your license number? [Click here](#) to contact technical support.

Feature Licenses (Optional)

Recommended: To provision an SVM in the following steps, respective feature licenses are mandatory for CIFS,NFS,iSCSI,NVMe-oF and/or FC/FCoE protocols.

[Submit and Continue](#)



The nodes should be discovered automatically; if they are not, click the Refresh link. By default, the cluster interfaces are 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 cluster creation.

5. Click Submit.

6. In the network page, complete the following sections:

- a. Cluster Management
 - Enter the IP address, netmask, gateway, and port details.
- b. Node Management
 - Enter the node management IP addresses and port details for all the nodes.
- c. Service Processor Management
 - Enter the IP addresses for all nodes.

- d. DNS Details
 - Enter the DNS domain names and server address.
- e. NTP Details
 - Enter the primary and alternate NTP server.

7. Click Submit.

ONTAP System Manager
Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Cluster (1) | **Network (2)** | Support (3) | Storage (4) | SVM (5) | Summary

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	192.168.156.60	255.255.255.0	192.168.156.1	e0M
Node Management	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
aa11-a...a320_1	192.168.156.61	e0M	Browse	
aa11-a...a320_2	192.168.156.62	e0M	Browse	
Service Processor Management	Default values have been detected for the Service Processor.			
	<input type="checkbox"/> Override the default values (Gateway is mandatory)			
	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
aa11-a...a320_1	192.168.156.58			
aa11-a...a320_2	192.168.156.59			

DNS Details

DNS Domain Names: flexpod.cisco.com

DNS Server IP Address: 10.1.156.250,10.1.156.251

NTP Details

Primary NTP Server: 192.168.156.11

Alternative NTP Server (Optional): 192.168.156.12


Previous | **Submit and Continue**

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

ONTAP System Manager

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



Event Notifications

Notify me through:

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

Previous

Submit and Continue

9. Click Submit.
10. In the Storage page, review the configuration details.
11. Skip this step if you want to configure it later. Otherwise, click Submit and Continue.

ONTAP System Manager


Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



48 nvme_ssd spares have been discovered. The recommended storage is shown below

Storage Recommendation



2 Data aggregates to be created
2 Spare disks will be available after the aggregates are created

[Hide Details](#)

Data Aggregate Details

■ Parity ■ Data □ Empty Slot

Node Name	Data Aggregate	Usable Size	Disks	Layout
aa11-affa320_1	aa11_affa320_1_NVME_SSD_1	16.42 TB	23 x 889.42 GB (NVME-SSD partition RAID-DP)	rg0
aa11-affa320_2	aa11_affa320_2_NVME_SSD_1	16.42 TB	23 x 889.42 GB (NVME-SSD partition RAID-DP)	rg0

Spare Disks

Node Name	Spare Disks	Disk Type
aa11-affa320_1	1 x 889.42 GB	NVME-SSD partition
aa11-affa320_2	1 x 889.42 GB	NVME-SSD partition

[Previous](#) [Submit and Continue](#) [Skip this Step](#)

12. In the SVM Tab, you can create your own SVM, or you can skip this step.

13. Click Skip this Step.

ONTAP System Manager

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



Create Storage Virtual Machine

SVM Name:

Data Protocols:

- CIFS *License required*
- NFS *License required*
- iSCSI *License required*
- FC/FCoE *License required*
- NVMe *License required*


Buttons: Previous, Submit and Continue, Skip this Step

14. In the Summary page, review the configuration details if needed and click Manage Your Cluster.

Click to go back, hold to see history

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 SVM and Storage Objects.

✔ Congratulations!
Cluster aa11-ffa320 ready for use.
There are a few advanced cluster features that you can configure.

- Cluster Peering**
Peers two clusters/SVMs to transfer and share data between them. [Proceed](#)

[Manage Your Cluster](#) [Export Configuration](#)



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 into the Cluster

To log into the cluster, follow these steps:

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

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 2 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes if it was not completed during the installation.

```
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 5 if high availability is not 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.

```
storage failover hwassist show
```

Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, follow this step:



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

Run the following command:

```
net interface modify -vserver aall-affa320 -lif cluster_mgmt_lif_1 -auto-revert true
```

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 AFF 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 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, run the following commands:

```
storage aggregate create -aggregate aggr1_node01 -node <st-node01> -diskcount <num-
disks>
storage aggregate 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 `storage aggregate show` command to display the aggregate creation status. Do not proceed until both `aggr1_node1` and `aggr1_node2` are online.

Remove Ports from Default Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, `e5a`, `e5b`, and so on) should be removed from the default broadcast domain, leaving just the management network port (`e0M`). To perform this task, run the following commands:

```
net port broadcast-domain remove-ports -broadcast-domain Default -ports <st-
node01>:e2a,<st-node01>:e2b,<st-node01>:e0g,<st-node01>:e0h,<st-node02>:e2a,<st-
node02>:e2b,<st-node02>:e0g,<st-node02>:e0h

network port broadcast-domain show
```

Disable Flow Control on 10GbE and 100GbE ports

NetApp recommends disabling flow control on all the 10/40/100GbE and UTA2 ports that are connected to external devices. To disable flow control, follow these steps:

1. Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port e0g,e0h,e2a,e2b -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 e0g,e0h,e2a,e2b -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 Auto-Negotiate on 100GbE Ports

To disable auto-negotiate on the 100GbE ports, follow these steps:

1. Run the following command to configure the ports on node 01:

```
network port modify -node <st-node01> -port e0g,e0h -autonegotiate-admin false -
speed-admin 100000 -duplex-admin full -flowcontrol-admin none
```

2. Run the following command to configure the ports on node 02:

```
network port modify -node <st-node02> -port e0g,e0h -autonegotiate-admin false -
speed-admin 100000 -duplex-admin full -flowcontrol-admin none

net port show -node * -port e0g,e0h -fields speed-admin ,duplex-admin,flowcontrol-
admin
(network port show)
node          port duplex-admin speed-admin flowcontrol-admin
-----
aa11-affa320-01 e0g full 100000 none
aa11-affa320-01 e0h full 100000 none
aa11-affa320-02 e0g full 100000 none
aa11-affa320-02 e0h full 100000 none
4 entries were displayed.
```

Disable Auto-Negotiate on Fibre Channel Ports

In accordance with the best practices for Fibre Channel host ports, disable auto-negotiate on each FCP adapter in each controller node.

1. Disable each Fibre Channel adapter in the controllers with the `fc adapter modify` command.

```
fc adapter modify -node <st-node01> -adapter 1a -status-admin down
fc adapter modify -node <st-node01> -adapter 1b -status-admin down
fc adapter modify -node <st-node02> -adapter 1a -status-admin down
fc adapter modify -node <st-node02> -adapter 1b -status-admin down
```

2. Set the desired speed on the adapter and return it to the online state.

```
fc adapter modify -node <st-node01> -adapter 1a -speed 32 -status-admin up
fc adapter modify -node <st-node01> -adapter 1b -speed 32 -status-admin up
fc adapter modify -node <st-node02> -adapter 1a -speed 32 -status-admin up
fc adapter modify -node <st-node02> -adapter 1b -speed 32 -status-admin up
```

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.

Enable Link-layer Discovery Protocol on all Ethernet Ports

Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches with the following step:

1. Enable LLDP on all ports of all nodes in the cluster.

```
node run * options lldp.enable on
```

Create Management Broadcast Domain

If the management interfaces are required to be on a separate VLAN, create a new broadcast domain for those interfaces.

```
network port broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500
network port broadcast-domain show
```

Create NFS Broadcast Domain

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

```
network port broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
network port broadcast-domain show
```

Create Interface Groups

To create the LACP interface groups for the 100GbE data interfaces, run the following commands:

```
network port ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0h

network port ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0h

network port ifgrp show
```

Change MTU on Interface Groups

Change the MTU size on the base interface-group ports before creating the VLAN ports.

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

Create VLANs

To create VLANs, follow these steps:

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

```
network port vlan create -node <st-node01> -vlan-name a0a-<ib-mgmt-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<ib-mgmt-vlan-id>
```

```
network port broadcast-domain add-ports -broadcast-domain IB-MGMT -ports <st-
node01>:a0a-<ib-mgmt-vlan-id>,<st-node02>:a0a-<ib-mgmt-vlan-id>

network port vlan show
```

2. Create the NFS VLAN ports and add them to the `Infra_NFS` broadcast domain.

```
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>

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

Configure Network Time Protocol

To configure time synchronization on the cluster, follow these 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, `201903271549.30`).

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

```
cluster time-service ntp server create -server <nexus-A-mgmt0-ip>
cluster time-service ntp server create -server <nexus-B-mgmt0-ip>
```

Configure Simple Network Management Protocol

To configure the Simple Network Management Protocol (SNMP), follow these 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>
```



To enable SNMPv3, refer to the [SNMP Configuration Express Guide](#) to configure SNMPv3 and setup one or more security users.

Create SVM

To create an infrastructure SVM, follow these steps:

1. Run the `vserver create` command.

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

2. Remove the unused data protocols from the SVM: CIFS, iSCSI, and NVMe.

```
vserver remove-protocols -vserver Infra-SVM -protocols iscsi,cifs
```

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.

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



If the NFS license was not installed during the cluster configuration, make sure to install the license before starting the NFS service.

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plug-in.

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

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, follow these 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 infra_svm_root_m01 -aggregate aggr1_node01
-size 1GB -type DP

volume create -vserver Infra-SVM -volume infra_svm_root_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:infra_svm_root -destination-path Infra-SVM:infra_svm_root_m01 -type LS -schedule 15min
```

```
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-SVM:infra_svm_root_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship.

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

Create Block Protocol (FC) Service

Run the following command to create the FCP service on each SVM. This command also starts the FCP service and sets the WWN for the SVM.

```
vserver fcp create -vserver Infra-SVM
vserver fcp show
```



If the FC license was not installed during the cluster configuration, make sure to install the license before creating the FC service.

Configure HTTPS Access

To configure secure access to the storage controller, follow these 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, the <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
-security 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.

- 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 verify that the system logs are available in a web browser.

```
set -privilege admin
```

```
https://<node01-mgmt-ip>/spi
```

```
https://<node02-mgmt-ip>/spi
```

Configure NFSv3

To configure NFSv3 on the SVM, follow these 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 infra_svm_root -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 -aggregate aggr1_node02 -
size 1TB -state online -policy default -junction-path /infra_datastore -space-
guarantee none -percent-snapshot-space 0

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

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

snapmirror update-ls-set -source-path Infra-SVM:infra_svm_root
```



If you are going to setup and use SnapCenter to backup the infra_datastore volume, add “-snapshot-policy none” to the end of the volume create command for the infra_datastore volume.

Create Boot LUNs

To create three boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -size 15GB -
ostype vmware -space-reserve disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-02 -size 15GB -
ostype vmware -space-reserve disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-03 -size 15GB -
ostype vmware -space-reserve disabled
```

Modify Volume Efficiency

On NetApp All Flash FAS systems, deduplication is enabled by default. To disable the efficiency policy on the infra_swap volume, follow this step:

```
volume efficiency off -vserver Infra-SVM -volume infra_swap
```

Create FC LIFs

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

```
network interface create -vserver Infra-SVM -lif fcp-lif011a -role data -data-
protocol fcp -home-node <st-node01> -home-port 1a -status-admin up
```

```

network interface create -vserver Infra-SVM -lif fcp-lif011b -role data -data-
protocol fcp -home-node <st-node01> -home-port 1b -status-admin up

network interface create -vserver Infra-SVM -lif fcp-lif021a -role data -data-
protocol fcp -home-node <st-node02> -home-port 1a -status-admin up

network interface create -vserver Infra-SVM -lif fcp-lif021b -role data -data-
protocol fcp -home-node <st-node02> -home-port 1b -status-admin up

network interface show

```

Create NFS LIFs

To create NFS LIFs, 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, follow these 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 a0a-<ib-mgmt-vlan-id> -address <svm-mgmt-ip>
-netmask <svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -
firewall-policy mgmt -auto-revert true

```

2. Create a default route that enables 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 them.

Configure and Test 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>
```

Test the AutoSupport configuration by sending a message from all nodes of the cluster:

```
autosupport invoke -node * -type all -message "FlexPod storage configuration  
completed"
```

Cisco UCS Configuration

Cisco UCS Base Configuration

This FlexPod deployment explains the configuration steps for the Cisco UCS 6454 Fabric Interconnects (FI) in a design that will support Fibre Channel SAN boot.



If setting up a system with iSCSI boot, the sections with (FCP) in the heading can be skipped and then complete the [Cisco UCS iSCSI Configuration](#) section in the Appendix.

Perform Initial Setup of Cisco UCS 6454 Fabric Interconnects for FlexPod Environments

This section provides the 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, follow these steps:

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

```

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

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

Enter the switch fabric (A/B) []: A

Enter the system name: <ucs-cluster-name>

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>

Cluster IPv4 address : <ucs-cluster-ip>

Configure the DNS Server IP address? (yes/no) [n]: y

DNS IP address : <dns-server-1-ip>

```

```

Configure the default domain name? (yes/no) [n]: y

Default domain name : <ad-dns-domain-name>

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

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

```

2. Wait for the login prompt for UCS Fabric Interconnect A before proceeding to the next section.

Cisco UCS Fabric Interconnect B

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

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

```

Enter the configuration method. (console/gui) ? console

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

Enter the admin password of the peer Fabric interconnect: <password>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucsa-mgmt-mask>
Cluster IPv4 address          : <ucs-cluster-ip>

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0
IPv4 Address

Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>

Local fabric interconnect model(UCS-FI-6454)
Peer fabric interconnect is compatible with the local fabric interconnect.
Continuing with the installer...

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

```

2. Wait for the login prompt for UCS Fabric Interconnect B before proceeding to the next section.

Cisco UCS Setup

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, follow these 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 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 4.0(4g)

This document assumes the use of Cisco UCS 4.0(4g). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.0(4g), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To enable anonymous reporting, follow this 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

SMTP Server

Host (IP Address or Hostname):

Port:

Don't show this message again.

OK

Cancel

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, follow these steps:

1. In Cisco UCS Manager, click Admin.

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.

Configure Unified Ports (FCP)

Fibre Channel port configurations differ between the 6454, 6332-16UP and the 6248UP Fabric Interconnects. All Fabric Interconnects have a slider mechanism within the Cisco UCS Manager GUI interface, but the fibre channel port selection options for the 6454 are from the first 8 ports starting from the first port and configured in increments of 4 ports from the left. For the 6332-16UP the port selection options are from the first 16 ports starting from the first port, and configured in increments of the first 6, 12, or all 16 of the unified ports. With the 6248UP, the port selection options will start from the right of the 32 fixed ports, or the right of the 16 ports of the expansion module, going down in contiguous increments of 2. The remainder of this section shows configuration of the 6454. Modify as necessary for the 6332-16UP or 6248UP.

To enable the fibre channel ports, follow these steps for the 6454:

1. In Cisco UCS Manager, click Equipment.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate).
3. Select Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
5. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 4 or 8 ports to be set as FC Uplinks.

Configure Unified Ports

**Instructions**

The position of the slider determines the type of the ports.

All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	

■ Up ■ Admin Down ■ Fail ■ Link Down

6. Click OK, then click Yes, then click OK to continue.
7. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
8. Select Configure Unified Ports.
9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
10. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 4 or 8 ports to be set as FC Uplinks.
11. Click OK, then click Yes, then OK to continue.
12. Wait for both Fabric Interconnects to reboot.
13. Log back into Cisco UCS Manager.

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, follow these steps:

1. In Cisco UCS Manager, click LAN.
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. Optionally, enter the Primary and Secondary DNS server addresses.

Create Block of IPv4 Addresses



From	:	<input type="text" value="192.168.156.240"/>	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="10.1.156.250"/>	Secondary DNS	:	<input type="text" value="10.1.156.251"/>

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, follow these steps:

1. In Cisco UCS Manager, click Admin.
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 <nexus-A-mgmt0-ip> and click OK. Click OK on the confirmation.

Add NTP Server



NTP Server :

OK

Cancel



We used the Nexus switch mgmt0 interface IP here because it is in the same L2 domain as the UCS mgmt0 IPs. We could also use the Nexus NTP IPs, but that traffic would then have to pass through an L3 router.

8. Click Add NTP Server.
9. Enter <nexus-B-mgmt0-ip> and click OK. Click OK.

All / Time Zone Management / Timezone

General

Events

Actions

Add NTP Server

Properties

Time Zone :

NTP Servers

Advanced Filter Export Print

Name

NTP Server 192.168.156.11

NTP Server 192.168.156.12

Add Additional DNS Server(s)

To add one or more additional DNS servers to the UCS environment, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand All > Communications Management.
3. Select DNS Management.
4. In the Properties pane, select Specify DNS Server.
5. Enter the IP address of the additional DNS server.

Specify DNS Server



DNS Server (IP Address) :



6. Click OK and then click OK again. Repeat this process for any additional DNS servers.

Add an Additional Administrative User

To add an additional locally authenticated Administrative user (flexadmin) to the Cisco UCS environment in case issues arise with the admin user, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand User Management > User Services > Locally Authenticated Users.
3. Right-click Locally Authenticated Users and select Create User.
4. In the Create User fields it is only necessary to fill in the Login ID, Password, and Confirm Password fields. Fill in the Create User fields according to your local security policy.
5. Leave the Account Status field set to Active.
6. Set Account Expires according to your local security policy.
7. Under Roles, select admin.
8. Leave Password Required selected for the SSH Type field.

Create User



Login ID :
 First Name :
 Last Name :
 Email :
 Phone :
 Password :
 Confirm Password :
 Account Status : Active Inactive
 Account Expires :

Roles

- aaa
- admin
- facility-manager
- network
- operations
- read-only
- server-compute
- server-equipment
- server-profile
- server-security
- storage

Locales

9. Click OK and then click OK again to complete adding the user.

Enable Port Auto-Discovery Policy

Setting the port auto-discovery policy enables automatic discovery of Cisco UCS B-Series chassis server ports. To modify the port auto-discovery policy, follow these steps:

1. In Cisco UCS Manager, click Equipment, select All > Equipment in the Navigation Pane, and select the Policies tab on the right.
2. Under Port Auto-Discovery Policy, set Auto Configure Server Port to Enabled.

Equipment

Main Topology View Fabric Interconnects Servers Thermal Decommissioned Firmware Management **Policies** Faults Diagnostics

Global Policies Autoconfig Policies Server Inheritance Policies Server Discovery Policies SEL Policy Power Groups **Port Auto-Discovery Policy** Security

Actions

Use Global

Properties

Owner : **Local**

Auto Configure Server Port : Disabled Enabled

[Save Changes](#) [Reset Values](#)

3. Click Save Changes and then OK.

Enable Info Policy for Neighbor Discovery

Enabling the info policy enables Fabric Interconnect neighbor information to be displayed. To modify the info policy, follow these steps:

1. In Cisco UCS Manager, click Equipment, select All > Equipment in the Navigation Pane, and select the Policies tab on the right.
2. Under Global Policies, scroll down to Info Policy and select Enabled for Action.

Info Policy

Action : Disabled Enabled

3. Click Save Changes and then OK.

- Under Equipment, select Fabric Interconnect A (primary). On the right, select the Neighbors tab. CDP information is shown under the LAN tab and LLDP information is shown under the LLDP tab.

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, follow these steps:

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



If varying numbers of links between chassis and the Fabric Interconnects will be used, leave Action set to 1 Link.

- On the 6454 Fabric Interconnects, the Link Grouping Preference is automatically set to Port Channel and is greyed out. On a 6300 Series or 6200 Series Fabric Interconnect, set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G.

Equipment

The screenshot shows the Cisco UCS Manager interface. The top navigation bar includes tabs for Main Topology View, Fabric Interconnects, Servers, Thermal, Decommissioned, Firmware Management, and Policies. The Policies tab is active. Below it, there are sub-tabs for Global Policies, Autoconfig Policies, Server Inheritance Policies, Server Discovery Policies, SEL Policy, and Power Groups. The Global Policies sub-tab is selected, and the 'Chassis/FEX Discovery Policy' is expanded. The configuration for this policy is shown below:

Chassis/FEX Discovery Policy

Action : 1 Link

Link Grouping Preference : None Port Channel

- If any changes have been made, click Save Changes and then click OK.

Enable Server and Uplink Ports

To enable and verify server and uplink ports, follow these steps:

- In Cisco UCS Manager, click Equipment.
- Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
- Expand and select Ethernet Ports.
- Verify that all ports connected to UCS chassis are configured as Server ports and have a status of Up.
- Select the ports that are connected to Cisco FEXes and direct connect Cisco UCS C-Series servers, right-click them, and select Configure as Server Port.
- Click Yes to confirm server ports and click OK.

7. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
8. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
9. Click Yes to confirm uplink ports and click OK.
10. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
11. Expand and select Ethernet Ports.
12. Verify that all ports connected to UCS chassis are configured as Server ports and have a status of Up.
13. Select the ports that are connected to Cisco FEXes and direct connect C-series servers, right-click them, and select Configure as Server Port.
14. Click Yes to confirm server ports and click OK.
15. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
16. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
17. 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 FEX modules, follow these steps:

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

Acknowledge Chassis



Are you sure you want to acknowledge Chassis 1 ?

This operation will rebuild the network connectivity between the Chassis and the Fabrics it is connected to. Currently there are 8 active links to Fabric A and there are 8 active links to Fabric B.

Yes

No

4. Click Yes and then click OK to complete acknowledging the chassis.
5. If Nexus FEXes 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, follow these steps:

1. In Cisco UCS Manager, click LAN.



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 under Fabric A.
4. Select Create Port Channel.
5. Enter 153 as the unique ID of the port channel.
6. Enter `Po153-Nexus` as the name of the port channel.
7. Click Next.
8. Select the uplink ports connected to the Nexus switches to be added to the port channel.
9. Click >> to add the ports to the port channel.

Create Port Channel ? X

1

Set Port Channel Name

2

Add Ports

Ports			
Slot ID	Aggr. Po...	Port	MAC
No data available			

>>
<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
1	0	53	00:3A:9...
1	0	54	00:3A:9...

< Prev
Next >
Finish
Cancel

10. Click Finish to create the port channel.

11. Click OK.

12. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, select Port-Channel 153. Select 100 Gbps for the Admin Speed.

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 153 Po153...

General Ports Faults Events Statistics

Status

Overall Status : ▼ **Failed**

Additional Info : **port-channel-members-down**

Actions

Enable Port Channel

Disable Port Channel

Add Ports

Properties

ID : **153**

Fabric ID : **A**

Port Type : **Aggregation**

Transport Type : **Ether**

Name :

Description :

Flow Control Policy :

LACP Policy :

Note: Changing LACP policy may flap the port-channel if the suspend-individual value changes!

Admin Speed : 1 Gbps 10 Gbps 40 Gbps 25 Gbps 100 Gbps Auto

Operational Speed(Gbps) : **0**

13. Click Save Changes and OK. After a few minutes, verify that the Overall Status is Up and the Operational Speed is correct.

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 153 Po153...

General Ports Faults Events Statistics

Status

Overall Status : ▲ **Up**

Additional Info : **none**

Actions

Enable Port Channel

Disable Port Channel

Add Ports

Properties

ID : **153**

Fabric ID : **A**

Port Type : **Aggregation**

Transport Type : **Ether**

Name :

Description :

Flow Control Policy :

LACP Policy :

Note: Changing LACP policy may flap the port-channel if the suspend-individual value changes!

Admin Speed : 1 Gbps 10 Gbps 40 Gbps 25 Gbps 100 Gbps Auto

Operational Speed(Gbps) : **200**

14. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.

15. Right-click Port Channels under Fabric B.

16. Select Create Port Channel.

17. Enter 154 as the unique ID of the port channel.

18. Enter `Po154-Nexus` as the name of the port channel.
19. Click Next.
20. Select the ports connected to the Nexus switches to be added to the port channel:
21. Click >> to add the ports to the port channel.
22. Click Finish to create the port channel.
23. Click OK.
24. In the navigation pane, under LAN > LAN Cloud > Fabric B > Port Channels, select Port-Channel 154. Select 100 Gbps for the Admin Speed.
25. Click Save Changes and OK. After a few minutes, verify that the Overall Status is Up and the Operational Speed is correct.
26. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 153. Under Port-Channel 153, select Eth Interface 1/53. In the center pane under Properties, enter a User Label to indicate the port connectivity, such as `<nexus-a-hostname>:1/5`. Click Save Changes and OK. Repeat this process for the remaining three uplink ports.

Add UDLD to Uplink Port Channels

To configure the unidirectional link detection (UDLD) on the Uplink Port Channels to the Nexus switches for fibre optic connections, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > LAN Cloud > UDLD Link Policy.
3. Right-click UDLD Link Policy and select Create UDLD Link Policy.
4. Name the Policy UDLD-Normal and select Enabled for the Admin State and Normal for the Mode.

Create UDLD Link Policy



Name :

Admin State : Enabled Disabled

Mode : Normal Aggressive

5. Click OK, then OK again to complete creating the policy.
6. Expand Policies > LAN Cloud > Link Profile.
7. Right-click Link Profile and select Create Link Profile.
8. Name the Profile UDLD-Normal and select the UDLD-Normal Link Policy created above.

Create Link Profile



Name :

UDLD Link Policy :

9. Click OK, then click OK again to complete creating the profile.
10. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 153. Select the first Eth Interface under Port-Channel 153. From the drop-down list, select the UDLD-Normal Link Profile created above, click Save Changes and OK. Repeat this process for each Eth Interface under Port-Channel 153 and for each Eth Interface under Port-Channel 154 on Fabric B.

LAN / LAN Cloud / Fabric B / Port Channels / Port-Channel 15... / Eth Interface 1/54

General		Faults	Events
Actions		Properties	
Delete		ID	: 54
Enable Interface		Slot ID	: 1
Disable Interface		Fabric ID	: B
		Transport Type	: Ether
		Port	: sys/switch-B/slot-1/switch-ether/port-54
		Membership	: Up
		Link Profile	: UDLD-Normal ▼
		User Label	: <input type="text"/>

Create an Organization

To this point in the UCS deployment, all items have been deployed at the root level in Cisco UCS Manager. To allow this UCS to be shared among different projects, UCS Organizations can be created. In this validation, the organization for this FlexPod deployment in NA-FlexPod. To create an organization for this FlexPod deployment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. In the Navigation Pane, expand Servers > Service Profiles.
3. Right-click root under Service Profiles and select Create Organization.
4. Provide a name for the Organization to indicate this FlexPod deployment and optionally provide a Description.

Create Organization



Name :

Description :

OK

Cancel

5. Click OK, then click OK again to complete creating the organization.

Create a WWNN Pool for FC Boot (FCP)

To configure the necessary WWNN pool for the Cisco UCS environment, follow these steps on Cisco UCS Manager.

1. Select SAN.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Select Create WWNN Pool to create the WWNN pool.
5. Enter WWNN-Pool for the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.
7. Select Sequential for Assignment Order.

1 Define Name and Description

2 Add WWN Blocks

Create WWNN Pool

Name : WWNN-Pool

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

8. Click Next.

9. Click Add.

10. Modify the From field as necessary for the UCS Environment



Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the Cisco UCS domain. Within the From field in our example, the sixth and seventh octets were changed from 00:00 to A1:30 to represent these WWNNs being in the A13 cabinet.



When there are multiple UCS domains sitting in adjacency, it is important that these blocks; the WWNN, WWPN, and MAC, hold differing values between each set.

11. Specify a size of the WWNN block sufficient to support the available server resources.

Create WWN Block



From : Size :

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx



12. Click OK.

13. Click Finish and click OK to complete creating the WWNN pool.

Create WWPN Pools (FCP)

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

1. In Cisco UCS Manager, click SAN.
2. Select Pools > root.



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

3. Right-click WWPN Pools under the root organization.
4. Select Create WWPN Pool to create the WWPN pool.
5. Enter WWPN-Pool-A as the name of the WWPN pool.
6. Optional: Enter a description for the WWPN pool.
7. Select Sequential for Assignment Order.

1 Define Name and Description

2 Add WWN Blocks

Name : WWPN-Pool-A

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

8. Click Next.

9. Click Add.

10. Specify a starting WWPN.



For the FlexPod solution, the recommendation is to place `A` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:A1:3A:00`

11. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources remembering that servers could have multiple vHBAs and unassociated service profiles could be created.

Create WWN Block



From : Size :

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx



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



For the FlexPod solution, the recommendation is to place B in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric B addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with 20:00:00:25:B5:A1:3B:00.

23. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources remembering that servers could have multiple vHBAs and unassociated service profiles could be created.
24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Create VSANs (FCP)

To configure the necessary virtual storage area networks (VSANs) for the FlexPod Organization in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.



In this procedure, two VSANs are created, one for each SAN switching fabric.

2. Select SAN > SAN Cloud.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter VSAN-A as the name of the VSAN to be used for Fabric A.
6. Leave FC Zoning set at Disabled.
7. Select Fabric A.
8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

Create VSAN



Name :

FC Zoning Settings

FC Zoning : Disabled Enabled

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

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

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

Enter the VSAN ID that maps to this VSAN.

VSAN ID :

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN :

OK

Cancel

9. Click OK and then click OK again.
10. Under SAN Cloud, right-click VSANs.
11. Select Create VSAN.
12. Enter VSAN-B as the name of the VSAN to be used for Fabric B.
13. Leave FC Zoning set at Disabled.
14. Select Fabric B.
15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric B. It is recommended use the same ID for both parameters and to use something other than 1.
16. Click OK and then click OK again.

Enable FC Uplink VSAN Trunking (FCP)

To enable VSAN trunking on the FC Uplinks in the Cisco UCS environment, follow these steps:



Enabling VSAN trunking is optional. It is important that the Cisco MDS 9132T VSAN trunking configuration match the configuration set in Cisco UCS Manager.

1. In Cisco UCS Manager, click SAN.
2. Expand SAN > SAN Cloud.
3. Select Fabric A and in the Actions pane select Enable FC Uplink Trunking.
4. Click Yes on the Confirmation and Warning.
5. Click OK.
6. Select Fabric B and in the Actions pane select Enable FC Uplink Trunking.
7. Click Yes on the Confirmation and Warning.
8. Click OK.

Create FC Uplink Port Channels (FCP)

To create the FC Uplink Port Channels and assign the appropriate VSANs to them for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Select SAN > SAN Cloud.
3. Expand Fabric A and select FC Port Channels.
4. Right-click FC Port Channels and select Create FC Port Channel.
5. Set a unique ID for the port channel and provide a unique name for the port channel.
6. Click Next.
7. Select the appropriate Port Channel Admin Speed.
8. Select the ports connected to Cisco MDS A and use >> to add them to the port channel.

1 Set FC Port Channel Name

2 Add Ports

Create FC Port Channel ? X

Port Channel Admin Speed : 4 Gbps 8 Gbps 16gbps 32gbps

Ports		
Port	Slot ID	WWPN
3	1	20:03:00:3A...
4	1	20:04:00:3A...

>>

<<

Ports in the port channel		
Port	Slot ID	WWPN
1	1	20:01:00:3A...
2	1	20:02:00:3A...

Slot ID:

WWPN:

Slot ID:

WWPN:

< Prev
Next >
Finish
Cancel

9. Click Finish to complete creating the port channel.
10. Click OK the confirmation.
11. Under FC Port-Channels, select the newly created port channel.
12. In the right pane, view the drop-down list to select NA-VSAN-A.

SAN / SAN Cloud / Fabric A / FC Port Channels / FC Port-Channel 11 S...

General	Ports	Faults	Events	Statistics
Status Overall Status : ▼ Failed Additional Info : No operational members		Properties ID : 11 Fabric ID : A Port Type : Aggregation Transport Type : Fc		
Actions Enable Port Channel Disable Port Channel Add Ports		Name : <input type="text" value="SPo11"/> Description : <input type="text"/> VSAN : <input type="text" value="Fabric A/vsan NA-VSAN"/> ▾ Port Channel Admin Speed : <input type="radio"/> 4 Gbps <input type="radio"/> 8 Gbps <input type="radio"/> 16gbps <input checked="" type="radio"/> 32gbps Operational Speed(Gbps) : 0		

13. Click Save Changes to assign the VSAN.

14. Click OK.



At this point in the deployment, since the Cisco MDS has not yet been configured, the SAN port-channel will not come up.

15. On the left under FC Port Channels, expand FC Port-Channel 11. Under FC Port-Channel 11 select FC Interface 1/1. Enter a User Label to indicate the connectivity on the MDS 9132T switch, such as <mds-A-hostname>.fc1/5. Click Save Changes and OK. Repeat this process for FC Interface 1/2.

16. Expand Fabric B and select FC Port Channels.

17. Right-click FC Port Channels and select Create FC Port Channel.

18. Set a unique ID for the port channel and provide a unique name for the port channel.

19. Click Next.

20. Select the ports connected to Cisco MDS B and use >> to add them to the port channel.

21. Click Finish to complete creating the port channel.

22. Click OK on the confirmation.

23. Under FC Port-Channels, select the newly created port channel.

24. In the right pane, use the drop-down to select VSAN-B.

25. Click Save Changes to assign the VSAN.

26. Click OK.
27. On the left under FC Port Channels, expand FC Port-Channel 12. Under FC Port-Channel 12 select FC Interface 1/1. Enter a User Label to indicate the connectivity on the MDS 9132T switch, such as <mds-B-hostname>:fc1/5. Click Save Changes and OK. Repeat this process for FC Interface 1/2.

Disable Unused FC Uplink Ports (FCP)

When Unified Ports were configured earlier in this procedure, on the Cisco UCS 6454 FI and the Cisco UCS 6332-16UP FI, FC ports were configured in groups. Because of this group configuration, some FC ports are unused and need to be disabled to prevent alerts. To disable the unused FC ports 3 and 4 on the Cisco UCS 6454 FIs, follow these steps:

1. In Cisco UCS Manager, click Equipment.
2. In the Navigation Pane, expand Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > FC Ports.
3. Select FC Port 3 and FC Port 4. Right-click and select Disable.
4. Click Yes and OK to complete disabling the unused FC ports.
5. In the Navigation Pane, expand Equipment > Fabric Interconnects > Fabric Interconnect B > Fixed Module > FC Ports.
6. Select FC Port 3 and FC Port 4. Right-click and select Disable.
7. Click Yes and OK to complete disabling the unused FC ports.

Create vHBA Templates (FCP)

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

1. In Cisco UCS Manager, click SAN.
2. Expand Policies > root > Sub-Organizations > NA-FlexPod.
3. Right-click vHBA Templates under the FlexPod Organization.
4. Select Create vHBA Template.
5. Enter FCP-vHBA-A as the vHBA template name.
6. Keep Fabric A selected.
7. Leave Redundancy Type set to No Redundancy.
8. Select VSAN-A.
9. Leave Initial Template as the Template Type.
10. Select WWPN-Pool-A as the WWPN Pool.

Create vHBA Template



Name : FCP-vHBA-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : VSAN-A [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-Pool-A(49/64) ▼

QoS Policy : <not set> ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

OK

Cancel

11. Click OK to create the vHBA template.
12. Click OK.
13. Right-click vHBA Templates under the FlexPod Organization.
14. Select Create vHBA Template.
15. Enter FCP-vHBA-B as the vHBA template name.
16. Select Fabric B as the Fabric ID.
17. Leave Redundancy Type set to No Redundancy.

18. Select VSAN-B.
19. Leave Initial Template as the Template Type.
20. Select WWPN-Pool-B as the WWPN Pool.
21. Click OK to create the vHBA template.
22. Click OK.

Create SAN Connectivity Policy (FCP)

To configure the necessary Infrastructure SAN Connectivity Policy within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Select SAN > Policies > root > Sub-Organizations > NA-FlexPod.
3. Right-click SAN Connectivity Policies under the FlexPod Organization.
4. Select Create SAN Connectivity Policy.
5. Enter FC-Boot as the name of the policy.
6. Select the previously created WWNN-Pool for the WWNN Assignment.
7. Click the Add button at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter FCP-Fabric-A as the name of the vHBA.
9. Select the Use vHBA Template checkbox.
10. In the vHBA Template list, select FCP-vHBA-A.
11. In the Adapter Policy list, select VMWare.

Create vHBA



Name : FCP-Fabric-A

Use vHBA Template : Redundancy Pair : Peer Name :

vHBA Template : FCP-vHBA-A ▼

[Create vHBA Template](#)**Adapter Performance Profile**

Adapter Policy : VMWare ▼

[Create Fibre Channel Adapter Policy](#)

OK

Cancel

12. Click OK.
13. Click the Add button at the bottom to add a second vHBA.
14. In the Create vHBA dialog box, enter FCP-Fabric-B as the name of the vHBA.
15. Select the Use vHBA Template checkbox.
16. In the vHBA Template list, select FCP-vHBA-B.
17. In the Adapter Policy list, select VMWare.
18. Click OK.

Create SAN Connectivity Policy



Name :

Description :

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

World Wide Node Name

WWNN Assignment:

[Create WWNN Pool](#)

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

Name	WWPN
▶ vHBA FCP-Fabric-B	Derived
▶ vHBA FCP-Fabric-A	Derived

Delete Add Modify

OK

Cancel

19. Click OK to create the SAN Connectivity Policy.

20. Click OK to confirm creation.

Create MAC Address Pools

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

1. In Cisco UCS Manager, click LAN.
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 A 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 cabinet number information giving us 00:25:B5:A1:3A: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 remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created.

Create a Block of MAC Addresses



First MAC Address : Size :

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

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-Pool-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 B 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 our example of also embedding the cabinet number information giving us 00:25:B5:A1:3B: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 remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created.
24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click Servers.
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 in the FlexPod Organization, follow these steps:



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

1. In Cisco UCS Manager, click Servers.
2. Select and expand Pools > root > Sub-Organizations > NA-FlexPod.
3. Right-click Server Pools under the FlexPod Organization.
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 three (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra-Pool server pool.



Although the VMware minimum host cluster size is two, in most cases three servers are recommended.

9. Click Finish.
10. Click OK.

Create VLANs

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

1. In Cisco UCS Manager, click LAN.



In this procedure, five unique VLANs are created. See [Table 2](#)

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

Create VLANs



VLAN Name/Prefix :

Multicast Policy Name : [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

Check Overlap

OK

Cancel

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.



Modify these VLAN names as necessary for your environment.

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.
34. Enter VM-Traffic as the name of the VLAN to be used for VM Traffic.
35. Keep the Common/Global option selected for the scope of the VLAN.
36. Enter the VM-Traffic VLAN ID.

37. Keep the Sharing Type as None.

38. Click OK and then click OK again.

The screenshot shows the Cisco UCS Manager interface for configuring VLANs. On the left, a navigation pane shows a tree structure under 'LAN' > 'LAN Cloud' > 'VLANs', with 'VLAN default (1)' selected. The main content area is titled 'LAN / LAN Cloud / VLANs' and contains a table of VLANs. The table has columns for Name, ID, Type, Transport, Native, and VLAN Sharing. The first row, 'VLAN default (1)', is highlighted in blue. Below the table are buttons for '+ Add', 'Delete', and 'Info'.

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	No	None
VLAN IB-MGMT (...)	113	Lan	Ether	No	None
VLAN Infra-NFS (...)	3050	Lan	Ether	No	None
VLAN Native-VLA...	2	Lan	Ether	Yes	None
VLAN VM-Traffic ...	900	Lan	Ether	No	None
VLAN vMotion (30...	3000	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 modify the default firmware management policy in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select version 4.0(4g) for both the Blade and Rack Packages.

Modify Package Versions



Blade Package : 4.0(4g)B

Rack Package : 4.0(4g)C

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

<input type="checkbox"/>	Adapter
<input type="checkbox"/>	BIOS
<input type="checkbox"/>	Board Controller
<input type="checkbox"/>	CIMC
<input type="checkbox"/>	FC Adapters
<input type="checkbox"/>	Flex Flash Controller
<input type="checkbox"/>	GPUs
<input type="checkbox"/>	HBA Option ROM
<input type="checkbox"/>	Host NIC
<input type="checkbox"/>	Host NIC Option ROM
<input checked="" type="checkbox"/>	Local Disk
<input type="checkbox"/>	NVME Mswitch Firmware
<input type="checkbox"/>	PSU
<input type="checkbox"/>	Port Switch Firmware

OK

Apply

Cancel

Help

- Click OK, then click OK again to modify the host firmware package.

Set Jumbo Frames in Cisco UCS Fabric

Jumbo Frames are used in FlexPod for the NFS and iSCSI storage protocols. The normal best practice in FlexPod has been to set the MTU of the Best Effort QoS System Class in Cisco UCS Manager to 9216 for Jumbo Frames. In the Cisco UCS 6454 Fabric Interconnect the MTU for the Best Effort QoS System Class is fixed at normal and cannot be changed. With this setting of normal in the 6454, Jumbo Frames can pass through the Cisco UCS fabric without being dropped. The screenshot below is from Cisco UCS Manager on a 6454 Fabric Interconnect, where the MTU for the Best Effort class is not settable. To configure jumbo frames in the Cisco UCS fabric in a 6300 or 6200 series Fabric Interconnect, follow these steps:

- In Cisco UCS Manager, click LAN.

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

LAN / LAN Cloud / QoS System Class

Actions		Properties						
Use Global		Owner : Local						
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	normal	<input type="checkbox"/>	
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A	

[Configure Slow Drain Timers](#)



The only the Fibre Channel and Best Effort QoS System Classes are enabled in this FlexPod implementation. The Cisco UCS and Nexus switches are intentionally configured this way so that all IP traffic within the FlexPod will be treated as Best Effort. Enabling the other QoS System Classes without having a comprehensive, end-to-end QoS setup in place can cause difficult to troubleshoot issues. For example, NetApp storage controllers by default mark IP-based, VLAN-tagged packets with a CoS value of 4. With the default configuration on the Nexus switches in this implementation, storage packets will pass through the switches and into the Cisco UCS Fabric Interconnects with CoS 4 set in the packet header. If the Gold QoS System Class in the Cisco UCS is enabled and the corresponding CoS value left at 4, these storage packets will be treated according to that class and if Jumbo Frames is being used for the storage protocols, but the MTU of the Gold QoS System Class is not set to Jumbo (9216), packet drops will occur. Note also that if the Platinum class is enabled, the MTU must be set to 9216 to use Jumbo Frames in that class.

Create Local Disk Configuration Policy (Optional)

A local disk configuration specifying no local disks 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, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand 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.

Create Local Disk Configuration Policy



Name :

Description :

Mode :

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

FlexFlash Removable State : Yes No No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily.
Please ensure SD cards are not in use before changing the FlexFlash Removable State.

OK

Cancel

7. Click OK to create the local disk configuration policy.
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 server virtual network controller (vNIC) ports, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand 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.

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

8. Click OK to create the network control policy.
9. Click OK.

Create Power Control Policy

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

1. In Cisco UCS Manager, click Servers.

2. Expand 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.

Create Power Control Policy



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.

7. Click OK to create the power control policy.
8. Click OK.

Create Server Pool Qualification Policy (Optional)

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



This example creates a policy for Cisco UCS B200 M5 servers for a server pool.

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCS-B200M5.
6. Select Create Server PID Qualifications.
7. Select UCSB-B200-M5 from the PID drop-down list.

Create Server PID Qualifications



PID :



8. Click OK
9. Optionally select additional qualifications to refine server selection parameters for the server pool.
10. Click OK to create the policy then OK for the confirmation.

Create Server BIOS Policy

To create a server BIOS policy for VMware ESXi hosts within the FlexPod organization, follow these steps:



In this lab validation, some Cisco UCS B200 M5 and Cisco UCS C220 M5 servers had TPM2.0 modules installed. To utilize TPM2.0 functionality with VMware vSphere 6.7U2, the TPM module must be enabled and Trusted Execution Technology(TXT) disabled in BIOS. According to the [Cisco UCS Server BIOS Tolerances, Release 4.0](#) document, these settings are the default or Platform Default settings for all M5 servers. Because of this, these settings do not have to be added to this BIOS policy.

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Sub-Organizations > NA-FlexPod.
3. Right-click BIOS Policies under FlexPod Organization.
4. Select Create BIOS Policy.

5. Enter VM-Host as the BIOS policy name.

Create BIOS Policy



Name :

Description :

Reboot on BIOS Settings Change :



6. Click OK, then click OK again to create the BIOS Policy.
7. Under the FlexPod Organization, expand BIOS Policies and select the newly created BIOS Policy. Set the following within the Main tab of the Policy:
 - a. CDN Control -> Enabled
 - b. Quiet Boot -> Disabled

Actions

- Delete
- Show Policy Usage
- Use Global

Properties

Name : **VM-Host**
 Description :
 Owner : **Local**
 Reboot on BIOS Settings Change :

Advanced Filter Export Print

BIOS Setting	Value
CDN Control	Enabled
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Disabled
Resume on AC power loss	Platform Default

Add Delete Info

Save Changes Reset Values

8. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab. Set the following within the Processor tab:
 - a. DRAM Clock Throttling -> Performance
 - b. Frequency Floor Override -> Enabled

Servers / Policies / root / Sub-Organizations / NX-FlexPod / BIOS Policies / VM-Host

Main **Advanced** Boot Options Server Management Events**Processor** Intel Directed IO RAS Memory Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform Graphics Configuration

Advanced Filter Export Print

BIOS Setting	Value
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
Boot Performance Mode	Platform Default
CPU Performance	Platform Default
Core Multi Processing	Platform Default
DRAM Clock Throttling	Performance
Direct Cache Access	Platform Default
Energy Performance Tuning	Platform Default
Enhanced Intel SpeedStep Tech	Platform Default
Execute Disable Bit	Platform Default
Frequency Floor Override	Enabled
Intel HyperThreading Tech	Platform Default
Energy Efficient Turbo	Platform Default
Intel Turbo Boost Tech	Platform Default
Intel Virtualization Technology	Platform Default
Channel Interleaving	Platform Default
IMC Inteleave	Platform Default
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default

Add Delete Info

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

- a. Processor C State -> Disabled
- b. Processor C1E -> Disabled
- c. Processor C3 Report -> Disabled
- d. Processor C7 Report -> Disabled
- e. Energy Performance -> Performance

Servers / Policies / root / Sub-Organizations / NX-FlexPod / BIOS Policies / VM-Host

Main **Advanced** Boot Options Server Management Events

Processor Intel Directed IO RAS Memory Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform Graphics Configuration

Advanced Filter Export Print

BIOS Setting	Value
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Platform Default
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Platform Default
Energy Performance	Performance
ProcessorEppProfile	Platform Default
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Core Performance Boost	Platform Default
Downcore control	Platform Default
Global C-state Control	Platform Default

Add Delete Info

10. Click the RAS Memory tab, and select:

- a. LV DDR Mode -> Performance Mode

Servers / Policies / root / Sub-Organizations / NX-FlexPod / BIOS Policies / VM-Host

Main **Advanced** Boot Options Server Management Events

Processor Intel Directed IO **RAS Memory** Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform Graphics Configuration

Advanced Filter Export Print

BIOS Setting	Value
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

11. Click Save Changes.

12. Click OK.

Update the Default Maintenance Policy

To update the default Maintenance Policy to either require user acknowledgement before server boot when service profiles change or to make the changes on the next server reboot, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand 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.

Servers / Policies / root / Maintenance Policies / default

General

Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **default**

Description :

Owner : **Local**

Soft Shutdown Timer :

Storage Config. Deployment Policy : Immediate User Ack

Reboot Policy : Immediate User Ack Timer Automatic

On Next Boot (Apply pending changes at next reboot.)

Save Changes

Reset Values

6. Click Save Changes.
7. Click OK to accept the changes.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates within the FlexPod organization, follow these steps. A total of 4 vNIC Templates will be created. Two of the vNIC templates (Infra-A and Infra-B) will be created for vNICs to connect to VMware ESXi vSwitch0. vSwitch0 will have port groups for the IB-MGMT, Infra-NFS, vMotion, and VM-Traffic VLANs. The third and fourth vNIC templates (Infra-vDS-A and Infra-vDS-B) will be created for vNICs to connect to the VMware Virtual Distributed Switch (vDS). The vDS will have port groups for the IB-MGMT, vMotion, and VM-Traffic VLANs. The vMotion VLAN is being placed on both vSwitch0 and the vDS so that the vMotion VMkernel port can initially be created on vSwitch0 then migrated to the vDS to allow QoS marking of vMotion packets to occur within the vDS if QoS policies need to be applied to vMotion in the future. The IB-MGMT VLAN is being placed on both vSwitch0 and the vDS so that management virtual machines other than vCenter can be placed on the vDS. The VM-Traffic VLAN is being placed on both vSwitch0 and the vDS so that it will be allowed on vSwitch0 in case the customer chooses not to implement the vDS on the management ESXi hosts.

Create Infrastructure vNIC Templates

To create the infrastructure vNIC templates, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > root > Sub-Organizations > NA-FlexPod.
3. Under the FlexPod Organization, right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vSwitch0-A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. 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.

Create vNIC Template



Warning

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

Template Type : Initial Template Updating Template

VLANs

VLAN Groups

Advanced Filter
Export
Print
⚙️

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	default	○	1
<input checked="" type="checkbox"/>	IB-MGMT	○	113
<input checked="" type="checkbox"/>	Infra-NFS	○	3050
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>	2
<input checked="" type="checkbox"/>	VM-Traffic	○	900
<input checked="" type="checkbox"/>	vMotion	○	3000

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

18. Click OK to create the vNIC template.

19. Click OK.

20. Under the FlexPod organization, right-click vNIC Templates.

21. Select Create vNIC Template.

22. Enter vSwitch0-B as the vNIC template name.
23. Select Fabric B.
24. Do not select the Enable Failover checkbox.
25. Set Redundancy Type to Secondary Template.
26. Select vSwitch0-A for the Peer Redundancy Template.
27. In the MAC Pool list, select MAC-Pool-B.



The MAC Pool is all that needs to be selected for the Secondary Template, all other values will either be propagated from the Primary Template or set at default values.

28. Click OK to create the vNIC template.
29. Click OK.
30. Under the FlexPod Organization, right-click vNIC Templates.
31. Select Create vNIC Template.
32. Enter vDS0-A as the vNIC template name.
33. Keep Fabric A selected.
34. Do not select the Enable Failover checkbox.



If only Cisco UCS B-Series servers are being used in your implementation, do not enable Failover.

35. Select Primary Template for Redundancy Type.
36. Leave the Peer Redundancy Template set to <not set>.
37. Under Target, make sure that only the Adapter checkbox is selected.
38. Select Updating Template as the Template Type.
39. Under VLANs, select the checkboxes for IB-MGMT, vMotion, VM-Traffic, and Native-VLAN VLANs.
40. Set Native-VLAN as the native VLAN.
41. Select vNIC Name for the CDN Source.
42. For MTU, enter 9000.
43. In the MAC Pool list, select MAC-Pool-A.
44. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



Warning

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

Template Type : Initial Template Updating Template

VLAN Groups
Advanced Filter ↑ Export ↑ Print
⚙️

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	default	<input type="radio"/>	1
<input checked="" type="checkbox"/>	IB-MGMT	<input type="radio"/>	113
<input type="checkbox"/>	Infra-NFS	<input type="radio"/>	3050
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>	2
<input checked="" type="checkbox"/>	VM-Traffic	<input type="radio"/>	900
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>	3000

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

45. Click OK to create the vNIC template.

46. Click OK.

47. Under the FlexPod organization, right-click vNIC Templates.

48. Select Create vNIC Template

49. Enter vDS0-B as the vNIC template name.
50. Select Fabric B.
51. Do not select the Enable Failover checkbox.
52. Set Redundancy Type to Secondary Template.
53. Select vDS0-A for the Peer Redundancy Template.
54. In the MAC Pool list, select MAC-Pool-B.



The MAC Pool is all that needs to be selected for the Secondary Template, all other values will either be propagated from the Primary Template or set at default values.

55. Click OK to create the vNIC template.
56. Click OK.

Create High Traffic VMware Adapter Policy

To create the optional VMware-High-Traffic Ethernet Adapter policy to provide higher vNIC performance, follow these steps:



This Ethernet Adapter policy can be attached to vNICs when creating the LAN Connectivity policy for vNICs that have large amounts of traffic on multiple flows or TCP sessions. This policy provides more hardware transmit and receive queues handled by multiple CPUs to the vNIC.

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Right-click Adapter Policies and select Create Ethernet Adapter Policy.
4. Name the policy VMware-HighTrf.
5. Expand Resources and set the values as shown below.

Create Ethernet Adapter Policy



Name : VMware-HighTrf

Description :

Resources

Pooled : Disabled Enabled

Transmit Queues : 8 [1-1000]

Ring Size : 256 [64-4096]

Receive Queues : 8 [1-1000]

Ring Size : 512 [64-4096]

Completion Queues : 16 [1-2000]

Interrupts : 18 [1-1024]

Options

OK

Cancel



Although previous versions of this document set the Ring Sizes for the Transmit and Receive Queues to 4096, [Tuning Guidelines for Cisco UCS Virtual Interface Cards](#) states that the sizes should be increased only if packet drops are observed on the vNIC interfaces.

- Expand Options and select Enabled for Receive Side Scaling (RSS).

Create Ethernet Adapter Policy



Name :

Description :

+ Resources

- Options

Transmit Checksum Offload : Disabled Enabled

Receive Checksum Offload : Disabled Enabled

TCP Segmentation Offload : Disabled Enabled

TCP Large Receive Offload : Disabled Enabled

Receive Side Scaling (RSS) : Disabled Enabled

Accelerated Receive Flow Steering : Disabled Enabled

Network Virtualization using Generic Routing Encapsulation : Disabled Enabled

Virtual Extensible LAN : Disabled Enabled

Failback Timeout (Seconds) : **[0-600]**

Interrupt Mode : MSI X MSI IN Tx

Interrupt Coalescing Type : Min Idle

Interrupt Timer (us) : **[0-65535]**

RoCE : Disabled Enabled

Advance Filter : Disabled Enabled

Interrupt Scaling : Disabled Enabled

OK

Cancel

- Click OK, then click OK again to complete creating the Ethernet Adapter Policy.

Create LAN Connectivity Policy for FC Boot (FCP)

To configure the necessary Infrastructure LAN Connectivity Policy within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > Policies > root > Sub-Organizations > NA-FlexPod.
3. Under the FlexPod Organization, right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter FC-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-vSwitch0-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select vSwitch0-A.
10. In the Adapter Policy list, select VMWare.

Create vNIC

Name : Use vNIC Template : Redundancy Pair : vNIC Template : Peer Name : [Create vNIC Template](#)**Adapter Performance Profile**Adapter Policy : [Create Ethernet Adapter Policy](#)**OK****Cancel**

11. Click OK to add this vNIC to the policy.
12. Click Add to add another vNIC to the policy.
13. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select vSwitch0-B.
16. In the Adapter Policy list, select VMWare.
17. Click OK to add the vNIC to the policy.
18. Click Add to add another vNIC to the policy.
19. In the Create vNIC dialog box, enter 02-vDS0-A as the name of the vNIC.

20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select vDS0-A.
22. In the Adapter Policy list, select VMWare-HighTrf.



The VMware Adapter Policy can also be selected for this vNIC.

23. Click OK to add this vNIC to the policy.
24. Click Add to add another vNIC to the policy.
25. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select vDS0-B.
28. In the Adapter Policy list, select VMWare-HighTrf.



Select the same Adapter Policy that was selected for 02-Infra-vDS-A.

29. Click OK to add this vNIC to the policy.

Create LAN Connectivity Policy



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-vDS0-B	Derived	
vNIC 02-vDS0-A	Derived	
vNIC 01-vSwitch0-B	Derived	
vNIC 00-vSwitch0-A	Derived	

Delete Add Modify

Add iSCSI vNICs

OK

Cancel

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

Create vMedia Policy for VMware ESXi 6.7U2 ISO Install Boot

In the NetApp ONTAP setup steps, an HTTP web server is required, which is used for hosting ONTAP as well as VMware software. The vMedia Policy created will map the VMware ESXi 6.7U2 Cisco Custom ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, follow these steps:

1. In Cisco UCS Manager, select Servers.
2. Select Policies > root.
3. Right-click vMedia Policies.
4. Select Create vMedia Policy.

5. Name the policy ESXi-6.7U2-HTTP.
6. Enter "Mounts Cisco Custom ISO for ESXi 6.7U2" in the Description field.
7. Click Add.
8. Name the mount ESXi-6.7U2-HTTP.
9. Select the CDD Device Type.
10. Select the HTTP Protocol.
11. Enter the IP Address of the web server.



To avoid any DNS lookup issues, enter the IP of the web server instead of the hostname.

12. Enter VMware_ESXi_6.7.0_13006603_Custom_Cisco_6.7.2.1.iso as the Remote File name.



This VMware ESXi 6.7U2 Cisco Custom ISO can be downloaded from VMware Downloads.



If a working vCenter 6.7U2 installation is already in your environment, a FlexPod custom ISO for installing ESXi 6.7U2 with all necessary drivers for this FlexPod deployment can be created. Please see the [Appendix](#) for a procedure for building this custom ISO.

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

Create vMedia Mount



Name	:	<input type="text" value="ESXi-6.7U2-HTTP"/>
Description	:	<input type="text"/>
Device Type	:	<input checked="" type="radio"/> CDD <input type="radio"/> HDD
Protocol	:	<input type="radio"/> NFS <input type="radio"/> CIFS <input checked="" type="radio"/> HTTP <input type="radio"/> HTTPS
Hostname/IP Address	:	<input type="text" value="10.1.156.150"/>
Image Name Variable	:	<input checked="" type="radio"/> None <input type="radio"/> Service Profile Name
Remote File	:	<input type="text" value="VMware_ESXi_6.7.0_13006603_Custom_Cisco_6.7"/>
Remote Path	:	<input type="text" value="software/vSphere-6.7U2"/>
Username	:	<input type="text"/>
Password	:	<input type="text"/>
Remap on Eject	:	<input type="checkbox"/>

14. Click OK to create the vMedia Mount.

15. Click OK then click 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 FC Boot Policy (FCP)

This procedure applies to a Cisco UCS environment in which two Fibre Channel logical interfaces (LIFs) are on cluster node 1 (fcp-lif011a and fcp-lif011b) and two Fibre Channel LIFs are on cluster node 2 (fcp-lif021a and fcp-lif021b). 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 fcp-lif011a.

To create a boot policy for the within the FlexPod organization, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root > Sub-Organizations > NA-FlexPod.
3. Under the FlexPod Organization, right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter Boot-FCP-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Select the Uefi Boot Mode.
9. Select the Boot Security checkbox.

Create Boot Policy

Name	:	<input type="text" value="Boot-FCP-A"/>
Description	:	<input type="text"/>
Reboot on Boot Order Change	:	<input type="checkbox"/>
Enforce vNIC/vHBA/iSCSI Name	:	<input checked="" type="checkbox"/>
Boot Mode	:	<input type="radio"/> Legacy <input checked="" type="radio"/> Uefi
Boot Security	:	<input checked="" type="checkbox"/>



UEFI Secure Boot can be used to boot VMware ESXi 6.7U2 with or without a TPM 2.0 module in the UCS server.

10. Expand Local Devices and select Add Remote CD/DVD.
11. Expand vHBAs and select Add SAN Boot.
12. Select Primary for the type field.
13. Enter FCP-Fabric-A in the vHBA field.


Add SAN Boot



vHBA :

Type : Primary Secondary Any

- 14. Click OK.
- 15. From vHBAs, select Add SAN Boot Target.
- 16. Keep 0 as the value for Boot Target LUN.
- 17. Enter the WWPN for fcp-lif011a.

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

- 18. Select Primary for the SAN boot target type.

Add SAN Boot Target



Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary



19. Click OK to add the SAN boot target.
20. From vHBAs, select Add SAN Boot Target.
21. Enter 0 as the value for Boot Target LUN.
22. Enter the WWPN for fcp-lif021a.
23. Click OK to add the SAN boot target.
24. From vHBAs, select Add SAN Boot.
25. In the Add SAN Boot dialog box, enter FCP-Fabric-B in the vHBA box.
26. The SAN boot type should automatically be set to Secondary.
27. Click OK.
28. From vHBAs, select Add SAN Boot Target.
29. Keep 0 as the value for Boot Target LUN.
30. Enter the WWPN for fcp-lif011b.
31. Select Primary for the SAN boot target type.
32. Click OK to add the SAN boot target.
33. From vHBAs, select Add SAN Boot Target.
34. Keep 0 as the value for Boot Target LUN.

35. Enter the WWPN for fcp-lif021b.
36. Click OK to add the SAN boot target.
37. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

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.

- Local Devices
- CIMC Mounted vMedia
 - Add CIMC Mounted CD/DVD
 - Add CIMC Mounted HDD
- vNICs
- vHBAs
- iSCSI vNICs

Boot Order

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descript...
Rem...	1								
San	2								
▶ S...		FCP-Fa...	Primary						
▶ S...		FCP-Fa...	Second...						
CIM...	3								

38. Expand San > SAN Primary and select SAN Target Primary. Select Set Uefi Boot Parameters.
39. Fill in the Set Uefi Boot Parameters exactly as shown in the following screenshot:

Set Uefi Boot Parameters



Uefi Boot Parameters

Boot Loader Name	:	<input type="text" value="BOOTX64.EFI"/>
Boot Loader Path	:	<input "="" type="text" value="\EFI\BOOT\"/>
Boot Loader Description	:	<input type="text"/>



40. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target.

41. Repeat this process to set Uefi Boot Parameters for each of the 4 SAN Boot Targets.



For Cisco UCS B200 M5 and Cisco UCS C220 M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for M4 and earlier servers, VMware ESXi 6.7U2 will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

42. Click OK, then click OK again to create the boot policy.

Create Service Profile Template (FCP)

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot within the FlexPod organization. To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > NA-FlexPod.
3. Right-click the FlexPod Organization.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-FCP-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the "Updating Template" option.
7. Under UUID, select UUID_Pool as the UUID pool.

Create Service Profile Template ? X

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/org-NX-FlexPod**

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.

8. Click Next.

Configure Storage Provisioning

To configure storage provisioning, follow these steps:

1. 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.
2. Click Next.

Configure Networking

To configure networking, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select FC-Boot from the LAN Connectivity Policy drop-down list.
4. Leave Initiator Name Assignment at <not set>.

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:

[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.

< Prev Next > **Finish** Cancel

5. Click Next.

Configure SAN Connectivity

To configure SAN connectivity, follow these steps:

1. Select the Use Connectivity Policy option for the “How would you like to configure SAN connectivity?” field.
2. Select the FC-Boot option from the SAN Connectivity Policy drop-down list.

Create Service Profile Template ? X

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

Simple
 Expert
 No vHBAs
 Use Connectivity Policy

SAN Connectivity Policy : FC-Boot ▼ [Create SAN Connectivity Policy](#)

3. Click Next.

Configure Zoning

To configure zoning, follow this step:

1. Set no zoning options and click Next.

Configure vNIC/HBA Placement

To configure vNIC/HBA placement, follow these steps:

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

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

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

Configure Server Boot Order

To configure the server boot order, follow these steps:

1. Select Boot-FCP-A for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **Boot-FCP-A**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Uefi**
 Boot Security : **Yes**

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

Boot Order

Name	Order	vNIC/vHB...	Type	LUN Name	WWN	Slot Num...	Boot Name	Boot Path	Description
Remot...	1								
▶ San	2								
CIMC ...	3								

2. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

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**
 Storage Config. Deployment Policy : **User Ack**
 Reboot Policy : **User Ack**

< Prev Next > **Finish** Cancel

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these 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 “B200-M5” for the Server Pool Qualification to select only B200 M5 servers in the pool.
4. Expand Firmware Management and select the default Host Firmware Package.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: [Create Server Pool](#)

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

Up Down

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

Server Pool Qualification :

Restrict Migration :

Firmware Management (BIOS, Disk Controller, Adapter)

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

Host Firmware Package:

[Create Host Firmware Package](#)

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

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these 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.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

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 :

+ External IPMI/Redfish 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 : [Create Power Control Policy](#)

+ Scrub Policy

+ KVM Management Policy

+ Graphics Card 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, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organizations > NA-FlexPod > Service Template VM-Host-Infra-FCP-A.
3. Right-click VM-Host-Infra-FCP-A and select Create a Clone.
4. Name the clone VM-Host-Infra-FC-A-vM.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Select the newly-created VM-Host-Infra-FCP-A-vM and select the vMedia Policy tab.
7. Click Modify vMedia Policy.

8. Select the ESXi-6.7U2-HTTP vMedia Policy and click OK.
9. Click OK to confirm.

Create Service Profiles

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

1. Connect to UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organizations > NA-FlexPod > Service Template VM-Host-Infra-FCP-A-vm.
3. Right-click VM-Host-Infra-FCP-A-vm 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 3 as the “Number of Instances.”

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :



7. Click OK to create the service profiles.
8. Click OK in the confirmation message.
9. When VMware ESXi 6.7U2 has been installed on the hosts, the host Service Profiles can be bound to the VM-Host-Infra-FCP-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 WWPNS from NetApp Storage

SVM	Adapter	MDS Switch	Target: WWPNS
Infra-SVM	fcp-lif011a	Fabric A	<fcp-lif011a-wwpn>
	fcp-lif011b	Fabric B	<fcp-lif011b-wwpn>
	fcp-lif021a	Fabric A	<fcp-lif021a-wwpn>
	fcp-lif021b	Fabric B	<fcp-lif021b-wwpn>



To obtain the FC WWPNS, run the `network interface show` command on the storage cluster management interface.

Table 7 WWPNS for Cisco UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPNS
VM-Host-Infra-01	Fabric A	<vm-host-infra-01-wwpna>
	Fabric B	<vm-host-infra-01-wwpnb>
VM-Host-Infra-02	Fabric A	<vm-host-infra-02-wwpna>
	Fabric B	<vm-host-infra-02-wwpnb>
VM-Host-Infra-03	Fabric A	<vm-host-infra-03-wwpna>
	Fabric B	<vm-host-infra-03-wwpnb>



To obtain the FC vHBA WWPNS information in Cisco UCS Manager GUI, go to `Servers > Service Profiles > root > Sub-Organizations > Organization`. Click each service profile and then click the “Storage” tab, then “vHBAs” tab on the right. The WWPNS are displayed in the table at the bottom of the page.

SAN Switch Configuration

This section explains how to configure the Cisco MDS 9000s for use in a FlexPod environment. Follow the steps precisely because failure to do so could result in an improper configuration.

If directly connecting storage to the Cisco UCS fabric interconnects, skip this section.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in the section [FlexPod Cabling](#).

FlexPod Cisco MDS Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes you are using the Cisco MDS 9132T with NX-OS 8.4(1).

Cisco MDS 9132T A

To set up the initial configuration for the Cisco MDS A switch, <mds-A-hostname>, follow these steps:



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. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

```

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

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 : <mds-A-hostname>

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

Mgmt0 IPv4 address : <mds-A-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-A-mgmt0-netmask>

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

```

```

IPv4 address of the default gateway : <mds-A-mgmt0-gw>
Configure advanced IP options? (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
Enable the telnet service? (yes/no) [n]: Enter
Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter
Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter
Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter
Enable the http-server? (yes/no) [y]: Enter
Configure clock? (yes/no) [n]: Enter
Configure timezone? (yes/no) [n]: Enter
Configure summertime? (yes/no) [n]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <mds-A-mgmt0-ip>
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: Enter
Configure default zone mode (basic/enhanced) [basic]: Enter

```

2. Review the configuration.

```

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

```

Cisco MDS 9132T B

To set up the initial configuration for the Cisco MDS B switch, <mds-B-hostname>, follow these steps:



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. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

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

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 : <mds-B-hostname>

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

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

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

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

IPv4 address of the default gateway : <mds-B-mgmt0-gw>

Configure advanced IP options? (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

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

Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter
```

```

Configure timezone? (yes/no) [n]: Enter
Configure summertime? (yes/no) [n]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <mds-A-mgmt0-ip>
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: Enter
Configure default zone mode (basic/enhanced) [basic]: Enter

```

2. Review the configuration.

```

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

```

FlexPod Cisco MDS Switch Configuration

Enable Licenses

Cisco MDS 9132T A and Cisco MDS 9132T B

To enable the correct features on the Cisco MDS switches, follow these steps:

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

```

configure terminal
feature npiv
feature fport-channel-trunk

```

Add Second NTP Server

Cisco MDS 9132T A and Cisco MDS 9132T B

To configure the second NTP server, follow this step:

From the global configuration mode, run the following command:

```

ntp server <nexus-B-mgmt0-ip>

```

Configure Individual Ports

Cisco MDS 9132T A

To configure individual ports and port-channels for switch A, follow this step:

From the global configuration mode, run the following commands:

```
interface fc1/3
switchport description <st-clustername>-1:1a
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/4
switchport description <st-clustername>-2:1a
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description <ucs-clustername>-a:1/1
channel-group 15
no shutdown
exit

interface fc1/6
switchport description <ucs-clustername>-a:1/2
channel-group 15
no shutdown
exit

interface port-channel15
channel mode active
switchport trunk allowed vsan <vsan-a-id>
switchport description <ucs-clustername>-a
switchport speed 32000
no shutdown
exit
```



If VSAN trunking is not being used between the UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-a-id>” for interface port-channel15. Note also that the default setting of switchport trunk mode auto is being used for the port channel.

Cisco MDS 9132T B

To configure individual ports and port-channels for switch B, follow this step:

From the global configuration mode, run the following commands:

```
interface fc1/3
switchport description <st-clustername>-1:1b
switchport speed 32000
switchport trunk mode off
```



```

no shutdown
exit

interface fc1/4
switchport description <st-clustername>-2:1b
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description <ucs-clustername>-b:1/1
channel-group 15
no shutdown
exit

interface fc1/6
switchport description <ucs-clustername>-b:1/2
channel-group 15
no shutdown
exit

interface port-channel15
channel mode active
switchport trunk allowed vsan <vsan-b-id>
switchport description <ucs-clustername>-b
switchport speed 32000
no shutdown
exit

```



If VSAN trunking is not being used between the UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-b-id>” for interface port-channel15. Note also that the default setting of switchport trunk mode auto is being used for the port channel.

Create VSANs

Cisco MDS 9132T A

To create the necessary VSANs for fabric A and add ports to them, follow these steps:

From the global configuration mode, run the following commands:

```

vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fc1/3
vsan <vsan-a-id> interface fc1/4
vsan <vsan-a-id> interface port-channel15
exit

```

Cisco MDS 9132T B

To create the necessary VSANs for fabric B and add ports to them, follow these steps:

From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fc1/3
vsan <vsan-b-id> interface fc1/4
vsan <vsan-b-id> interface port-channel15
exit
```



At this point, it may be necessary to go into UCS Manager and disable and enable the FC port-channel interfaces to get the port-channels to come up.

Create Device Aliases

Cisco MDS 9132T A

To create device aliases for Fabric A that will be used to create zones, follow these steps:

From the global configuration mode, run the following commands:

```
device-alias mode enhanced
device-alias database
device-alias name Infra-SVM-fcp-lif011a pwnn <fcp-lif011a-wwpn>
device-alias name Infra-SVM-fcp-lif021a pwnn <fcp-lif021a-wwpn>
device-alias name VM-Host-Infra-01-A pwnn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwnn <vm-host-infra-02-wwpna>
device-alias name VM-Host-Infra-03-A pwnn <vm-host-infra-03-wwpna>
device-alias commit
```

Cisco MDS 9132T B

To create device aliases for Fabric B that will be used to create zones, follow these steps:

From the global configuration mode, run the following commands:

```
device-alias mode enhanced
device-alias database
device-alias name Infra-SVM-fcp-lif011b pwnn <fcp-lif011b-wwpn>
device-alias name Infra-SVM-fcp-lif021b pwnn <fcp-lif021b-wwpn>
device-alias name VM-Host-Infra-01-B pwnn <vm-host-infra-01-wwpnb>
device-alias name VM-Host-Infra-02-B pwnn <vm-host-infra-02-wwpnb>
device-alias name VM-Host-Infra-03-B pwnn <vm-host-infra-03-wwpnb>
device-alias commit
```

Create Zones and Zoneset

Cisco MDS 9132T A

To create the required zones and zoneset on Fabric A, run the following commands:

```
configure terminal
zone name Infra-SVM-Fabric-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-01-A init
member device-alias VM-Host-Infra-02-A init
member device-alias VM-Host-Infra-03-A init
member device-alias Infra-SVM-fcp-lif011a target
member device-alias Infra-SVM-fcp-lif021a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member Infra-SVM-Fabric-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
show zoneset active
copy r s
```



Since Smart Zoning is enabled, a single zone is created with all host boot initiators and boot targets for the Infra-SVM instead of creating a separate zone for each host with the host initiator and boot targets. If a new host is added, its boot initiator can simply be added to the single zone in each MDS switch and then the zoneset reactivated. If another SVM is added to the FlexPod with FC targets, a new zone can be added for that SVM.

Cisco MDS 9132T B

To create the required zones and zoneset on Fabric B, run the following commands:

```
configure terminal
zone name Infra-SVM-Fabric-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-01-B init
member device-alias VM-Host-Infra-02-B init
member device-alias VM-Host-Infra-03-B init
member device-alias Infra-SVM-fcp-lif011b target
member device-alias Infra-SVM-fcp-lif021b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member Infra-SVM-Fabric-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active
copy r s
```

Storage Configuration – Boot LUNs

ONTAP Boot Storage Setup

Create igroups

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

```
lun igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol fcp -ostype
vmware -initiator <vm-host-infra-01-wwpna>, <vm-host-infra-01-wwpnb>

lun igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol fcp -ostype
vmware -initiator <vm-host-infra-02-wwpna>, <vm-host-infra-02-wwpnb>

lun igroup create -vserver Infra-SVM -igroup VM-Host-Infra-03 -protocol fcp -ostype
vmware -initiator <vm-host-infra-03-wwpna>, <vm-host-infra-03-wwpnb>

lun igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol fcp -ostype vmware
-initiator <vm-host-infra-01-wwpna>, <vm-host-infra-01-wwpnb>, <vm-host-infra-02-
wwpna>, <vm-host-infra-02-wwpnb>, <vm-host-infra-03-wwpna>, <vm-host-infra-03-wwpnb>
```



Use the values listed in [Table 6](#) and [Table 7](#) for the WWPN information.

To view the three igroups just created, use the command `lun igroup show`.

```
lun igroup show -protocol fcp
```

Map Boot LUNs to igroups

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

```
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -igroup
VM-Host-Infra-01 -lun-id 0

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-02 -igroup
VM-Host-Infra-02 -lun-id 0

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-03 -igroup
VM-Host-Infra-03 -lun-id 0
```

VMware vSphere 6.7U2 Setup

VMware ESXi 6.7U2

This section provides detailed instructions for installing VMware ESXi 6.7U2 in a FlexPod 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 ESXi 6.7U2 from VMware

If the VMware ESXi ISO has not been downloaded, follow these steps:

1. Click the following link: [Cisco Custom ESXi 6.7U2 ISO](#).
2. You will need a user id and password on vmware.com to download this software.
3. Download the .iso file.

Log into Cisco UCS 6454 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 into the Cisco UCS environment, follow these steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link to 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.
7. Select Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-01.
8. In the Actions pane, click >> to the right of KVM Console.
9. Select the "Launch Java KVM Console" checkbox and click OK to launch the console.
10. Follow the prompts to launch the Java-based KVM console.

11. Select Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-02.
12. In the Actions pane, click >> to the right of KVM Console.
13. Select the “Launch Java KVM Console” checkbox and click OK to launch the console.
14. Follow the prompts to launch the Java-based KVM console.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-01, VM-Host-Infra-02, and VM-Host-Infra-03



Skip this section if you’re using vMedia policies; the ISO file will already be connected to KVM.

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

1. In the KVM window, click Virtual Media.
2. Select 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 Console tab to monitor the server boot.

Install ESXi

ESXi Hosts VM-Host-Infra-01, VM-Host-Infra-02, and VM-Host-Infra-03

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

1. Boot the server by selecting Boot Server in the KVM and click OK, then click OK again.
2. On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.
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.



It may be necessary to map function keys as User Defined Macros under the Macros menu in the UCS KVM console.

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, press Enter to reboot the server.



The ESXi installation image will be automatically unmapped in the KVM when Enter is pressed.

10. 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, follow these steps on each ESXi host:

ESXi Host VM-Host-Infra-01, VM-Host-Infra-02, and VM-Host-Infra-03

To configure each ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, press F2 to customize VMware ESXi.
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.
10. Using the spacebar, select vmnic1.



In lab testing, examples have been seen 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.

11. Press Enter.
12. Select the VLAN (Optional) option and press Enter.
13. Enter the <ib-mgmt-vlan-id> and press Enter.
14. Select IPv4 Configuration and press Enter.
15. Select the "Set static IPv4 address and network configuration" option by using the arrow keys and space bar.
16. Move to the IPv4 Address field and enter the IP address for managing the ESXi host.
17. Move to the Subnet Mask field and enter the subnet mask for the ESXi host.
18. Move to the Default Gateway field and enter the default gateway for the ESXi host.
19. Press Enter to accept the changes to the IP configuration.
20. Select the IPv6 Configuration option and press Enter.
21. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
22. Select the DNS Configuration option and press Enter.



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

23. Using the spacebar, select "Use the following DNS server addresses and hostname:"
24. Move to the Primary DNS Server field and enter the IP address of the primary DNS server.
25. Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.
26. Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.
27. Press Enter to accept the changes to the DNS configuration.
28. Press Esc to exit the Configure Management Network submenu.
29. Press Y to confirm the changes and reboot the ESXi host.

Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional)

ESXi Host VM-Host-Infra-01, VM-Host-Infra-02, and VM-Host-Infra-03

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port it is placed on. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will exist because vmk0 will retain the assigned MAC address unless the ESXi System Configuration is reset. To reset the MAC address of vmk0 to a random VMware-assigned MAC address, follow these steps:

1. From the ESXi console menu main screen, type Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of Static Macros.
2. Log in as root.
3. Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk0. vmk0 should be a part of the "Management Network" port group. Note the IP address and netmask of vmk0.
4. To remove vmk0, type `esxcfg-vmknic -d "Management Network"`.
5. To re-add vmk0 with a random MAC address, type `esxcfg-vmknic -a -i <vmk0-ip> -n <vmk0-netmask> "Management Network"`.
6. Verify vmk0 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.
7. Tag vmk0 as the management interface by typing `esxcli network ip interface tag add -i vmk0 -t Management`.
8. When vmk0 was re-added, if a message popped up saying vmk1 was marked as the management interface, type `esxcli network ip interface tag remove -i vmk1 -t Management`.
9. Type `exit` to log out of the command line interface.
10. Type Ctrl-Alt-F2 to return to the ESXi console menu interface.

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

Download and extract the offline bundle for the following VMware VIC Drivers and the NetApp NFS Plug-in for VMware VAAI to the Management workstation:



The VMware ESXi 6.7U2 Cisco Custom ISO contains the nenic driver version 1.0.29.0. It is not necessary to download or update the nenic driver but the commands are left here to be used for future updates.

[nfnic Driver version 4.0.0.38](#)

[nenic Driver version 1.0.29.0](#)

[NetApp NFS Plug-in 1.1.2 for VMware VAAI](#) (This is the offline bundle and does not need to be extracted.)

ESXi Hosts VM-Host-Infra-01, VM-Host-Infra-02, and VM-Host-Infra-03

To install VMware VIC Drivers and the NetApp NFS Plug-in for VMware VAAI on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, follow these steps:

1. Using an SCP program such as WinSCP, copy the three offline bundles referenced above to the /tmp directory on each ESXi host.
2. Using an ssh tool such as PuTTY, ssh to each VMware ESXi host. Log in as root with the root password.
3. Type `cd /tmp`.
4. Run the following commands on each host:

```
esxcli software vib update -d /tmp/VMW-ESX-6.7.0-nenic-1.0.29.0-offline_bundle-12897497.zip

esxcli software vib update -d /tmp/VMW-ESX-6.7.0-nfnic-4.0.0.38-offline_bundle-14003014.zip

esxcli software vib install -d /tmp/NetAppNasPlugin.v23.zip

reboot
```

Log into the First VMware ESXi Host by Using VMware Host Client

ESXi Host VM-Host-Infra-01

To log into the VM-Host-Infra-01 ESXi host by using the VMware Host Client, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

5. Decide whether to join the VMware Customer Experience Improvement Program and click OK.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

To set up the VMkernel ports and the virtual switches on the first ESXi host, follow these steps:



In this procedure, you're only setting up the first ESXi host. The second and third hosts will be added to vCenter and setup from the vCenter HTML5 Interface.

1. From the Host Client, select Networking.
2. In the center pane, select the Virtual switches tab.
3. Highlight the vSwitch0 line.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, select vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.
9. Click Save.
10. Select Networking, then select the Port groups tab.
11. In the center pane, right-click VM Network and select Remove.
12. Click Remove to complete removing the port group.
13. In the center pane, select Add port group.
14. 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.
15. Click Add to finalize the edits for the IB-MGMT Network.
16. At the top, select the VMkernel NICs tab.
17. Click Add VMkernel NIC.
18. For New port group, enter VMkernel-vMotion.
19. For Virtual switch, select vSwitch0.
20. Enter <vmotion-vlan-id> for the VLAN ID.

21. Change the MTU to 9000.
22. Select Static IPv4 settings and expand IPv4 settings.
23. Enter the ESXi host vMotion IP address and netmask.
24. Select the vMotion stack for TCP/IP stack.
25. Click Create.
26. Click Add VMkernel NIC.
27. For New port group, enter VMkernel-Infra-NFS
28. For Virtual switch, select vSwitch0.
29. Enter <infra-nfs-vlan-id> for the VLAN ID
30. Change the MTU to 9000.
31. Select Static IPv4 settings and expand IPv4 settings.
32. Enter the ESXi host Infrastructure NFS IP address and netmask.
33. Do not select any of the Services.
34. Click Create.
35. 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: 2

vSwitch Details	
MTU	9000
Ports	7936 (7925 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

NIC teaming policy	
Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Failback	Yes

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

Shaping policy	
Enabled	No

vSwitch topology

The diagram shows a central vSwitch with four network types connected to it:

- IB-MGMT Network** (VLAN ID: 113)
- VMkernel-Infra-NFS** (VLAN ID: 3050) with VMkernel ports: vmk2 (192.168.50.191)
- VMkernel-vMotion** (VLAN ID: 3000) with VMkernel ports: vmk1 (192.168.100.191)
- Management Network** (VLAN ID: 113) with VMkernel ports: vmk0 (10.1.156.191)

On the right, physical adapters are shown:

- vmnic1, 20000 Mbps, Full
- vmnic0, 20000 Mbps, Full

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

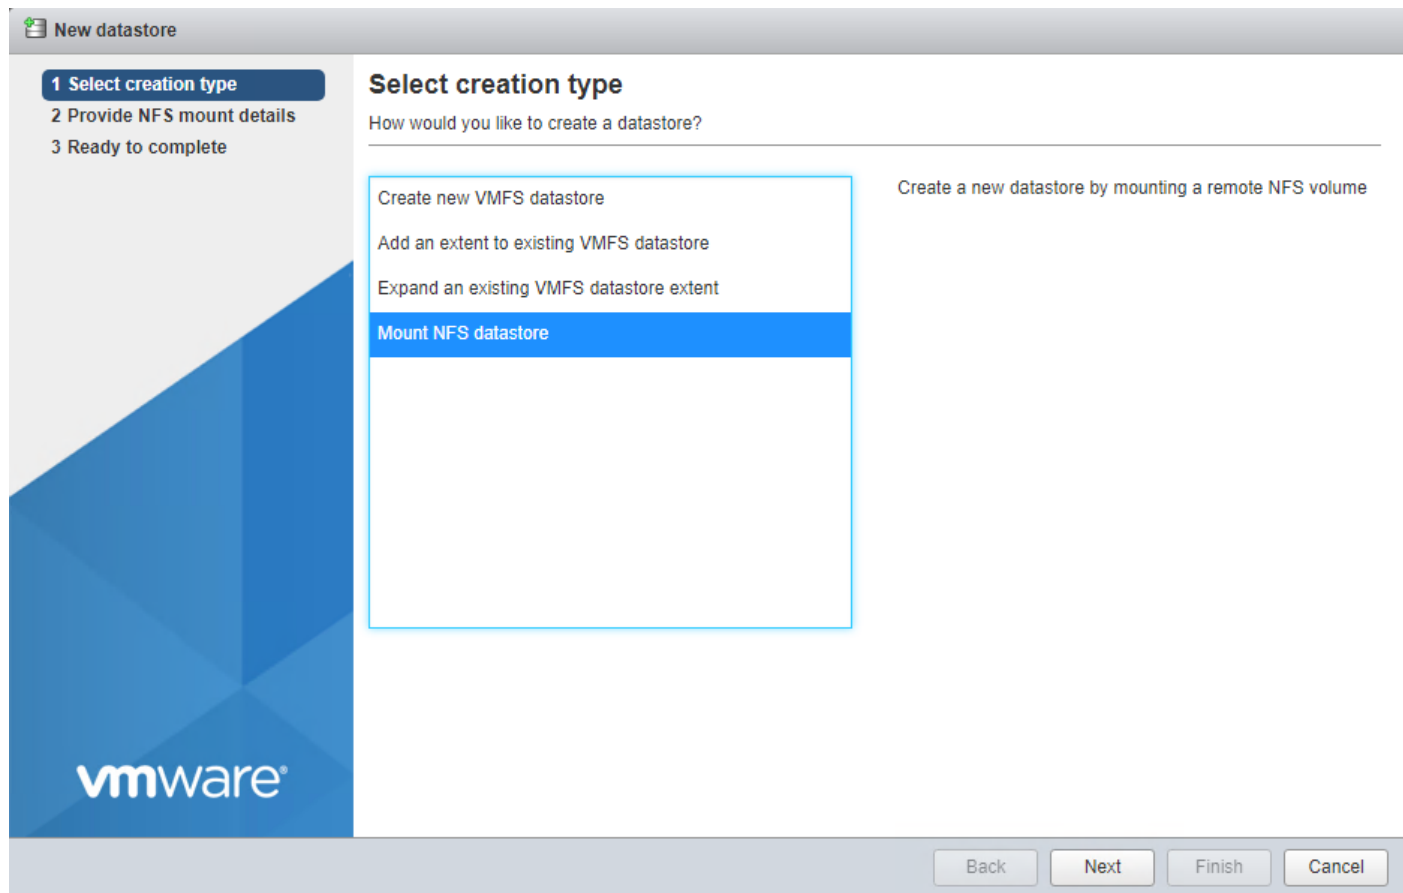
Port groups	Virtual switches	Physical NICs	VMkernel NICs	TCP/IP stacks	Firewall rules
Add VMkernel NIC Edit settings Refresh Actions Search					
Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.156.191	None
vmk2	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.191	None
vmk1	VMkernel-vMotion	vMotion stack	vMotion	192.168.100.191	None
3 items					

Mount Required Datastores

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, select Storage.
2. In the center pane, select the Datastores tab.
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` for the datastore name. Input the IP address for the `nfs-lif02` LIF for the NFS server. Input `/infra_datastore` for the NFS share. Leave the NFS version set at NFS 3. Click Next.

New datastore - infra_datastore

1 Select creation type
2 Provide NFS mount details
 3 Ready to complete

Provide NFS mount details

Provide the details of the NFS share you wish to mount

Name	infra_datastore
NFS server	192.168.50.62
NFS share	/infra_datastore
NFS version	<input checked="" type="radio"/> NFS 3 <input type="radio"/> NFS 4

Back Next Finish Cancel

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_swap for the datastore name. Input the IP address for the nfs-lif01 LIF for the NFS server. Input /infra_swap 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.

Datastores									
Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provision...	Access		
datastore1	Non-SSD	7.5 GB	4.37 GB	3.13 GB	VMFS6	Supported	Single		
infra_datastore	Unknown	500 GB	16.96 MB	499.98 GB	NFS	Supported	Single		
infra_swap	Unknown	100 GB	23.07 MB	99.98 GB	NFS	Supported	Single		

3 items

Configure NTP on First ESXi Host

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, select Manage.
2. In the center pane, select System > Time & Date.
3. Click Edit settings.
4. Make sure Use Network Time Protocol (enable NTP client) is selected.
5. Use the drop-down list 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

03/01/2019 5:38 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 time-zone will be UTC.



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

Configure ESXi Host Swap

ESXi Host VM-Host-Infra-01

To configure host swap on the first ESXi host, follow these steps on the host:

1. From the Host Client, select Manage.
2. In the center pane, select System > Swap.

3. Click Edit settings.
4. Use the drop-down list to select `infra_swap`. Leave all other settings unchanged.

Edit swap configuration	
Enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No
Datastore	infra_swap
Local swap enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No
Host cache enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No

Save Cancel

5. Click Save to save the configuration changes.
6. If you are implementing iSCSI boot, execute the VMware ESXi setup scripts in the iSCSI Addition Appendix.

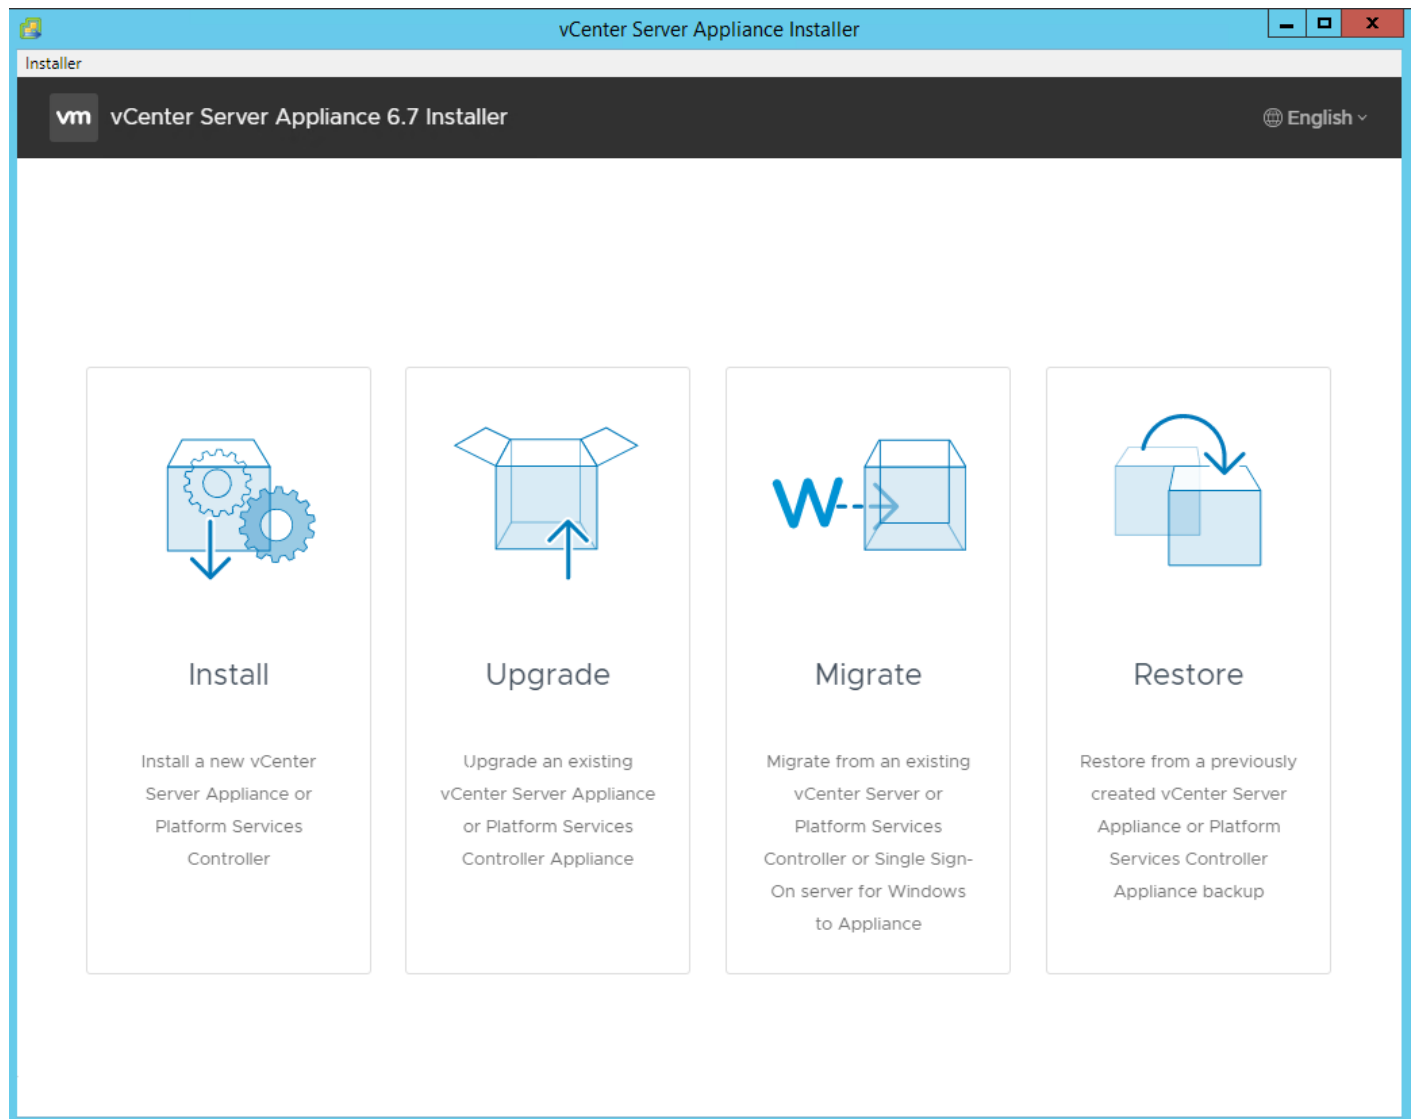
VMware vCenter 6.7U2

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.7U2 Server Appliance in a FlexPod environment. After the procedures are completed, a VMware vCenter Server will be configured.

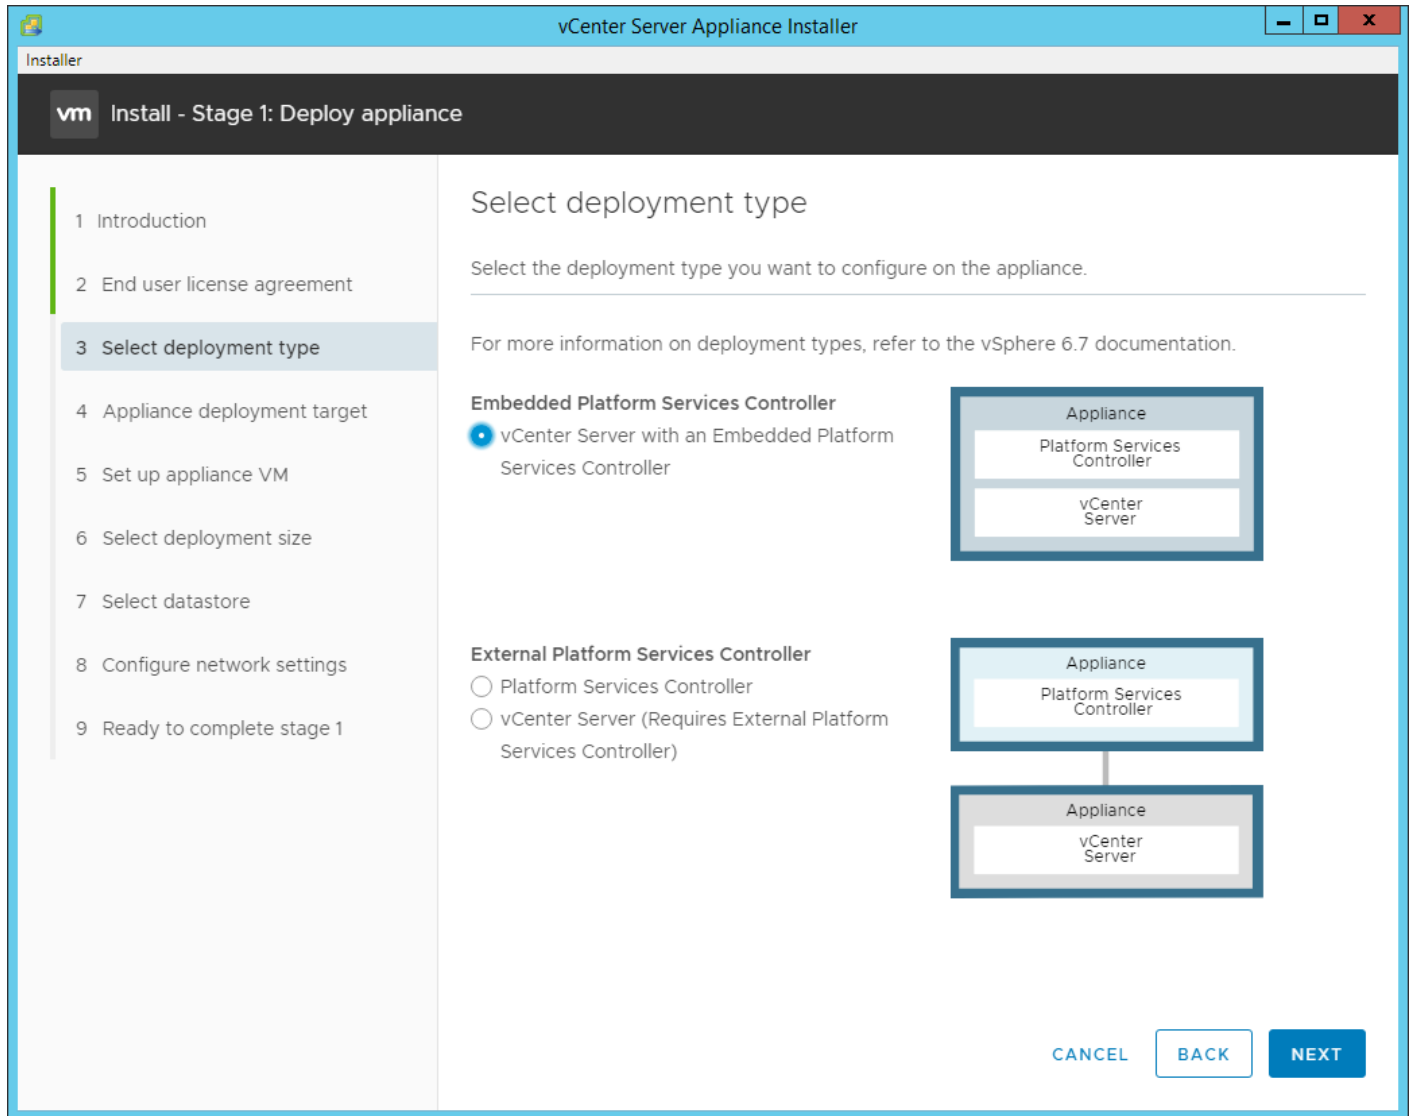
Build the VMware vCenter Server Appliance

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

1. Locate and copy the `VMware-VCSA-all-6.7.0-13643870.iso` file to the desktop of the management workstation. This ISO is for the VMware vSphere 6.5 U2 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 `vcasa-ui-installer > win32` directory and double-click `install-er.exe`. The vCenter Server Appliance Installer wizard appears.



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 and click NEXT.

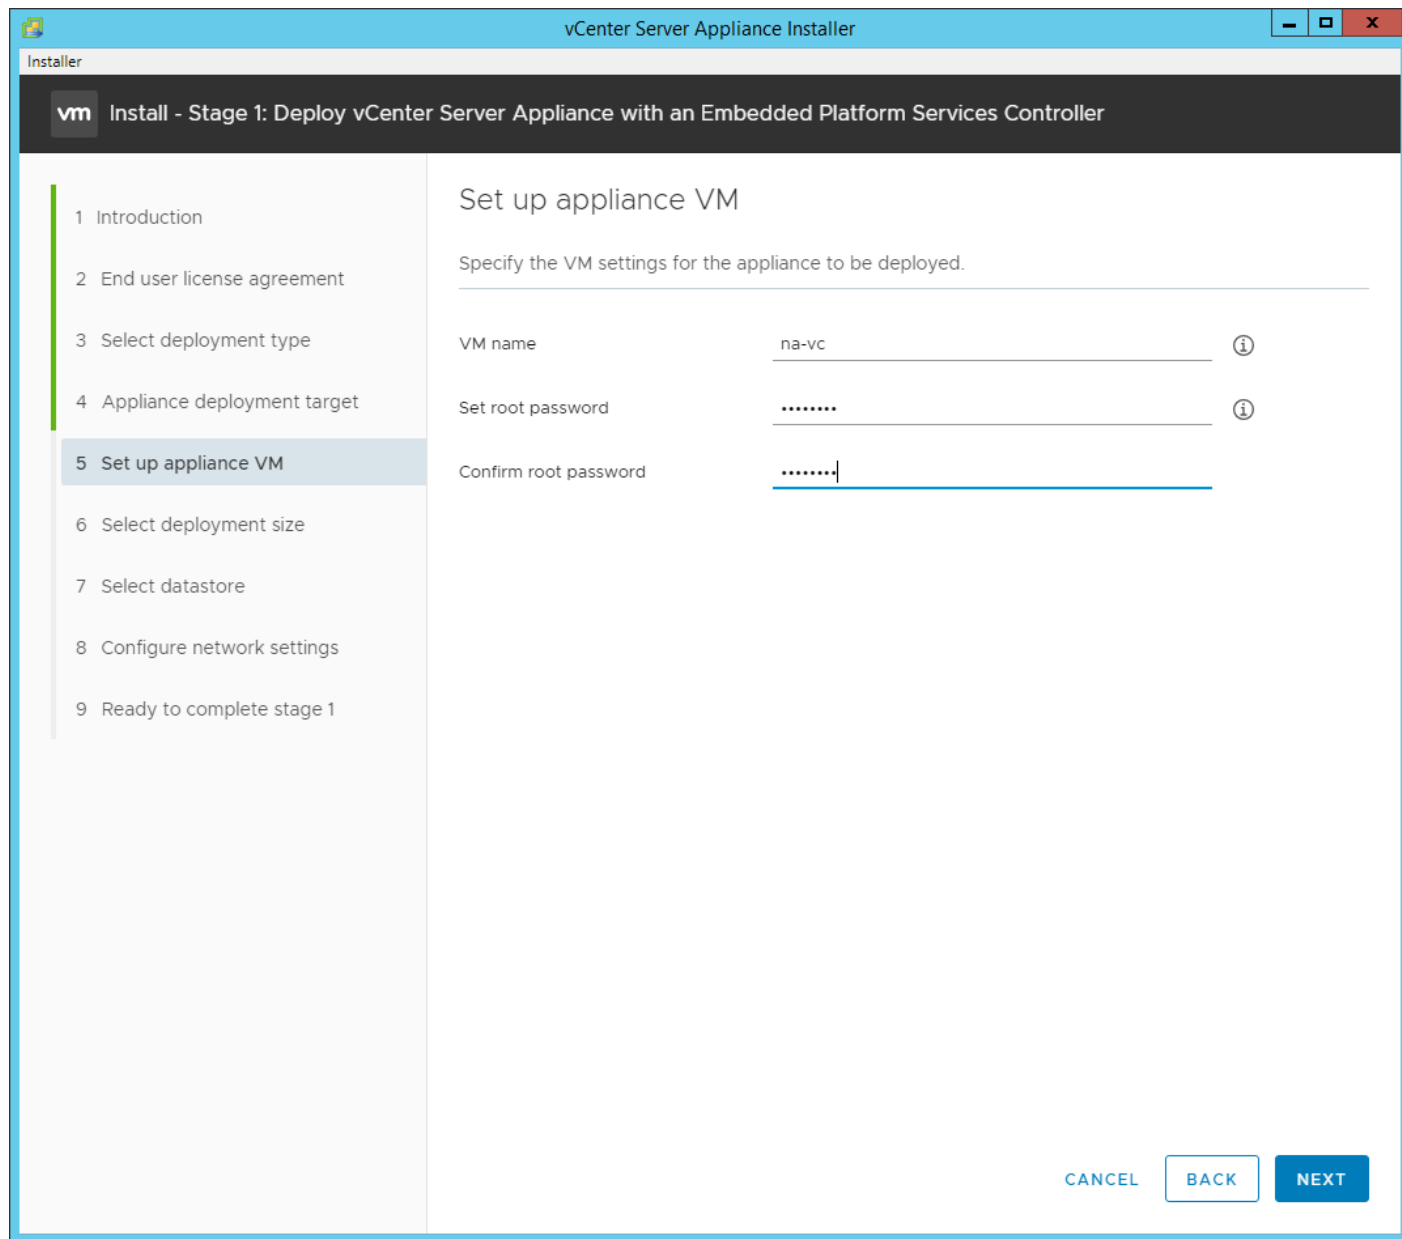


8. In the “Appliance deployment target”, enter the host name or IP address of the first ESXi host, User name (root) and Password.

The screenshot shows the 'vCenter Server Appliance Installer' window. The title bar reads 'vCenter Server Appliance Installer'. Below the title bar, the text 'Installer' is visible. The main header area contains the VMware logo and the text 'Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller'. On the left side, there is a vertical list of steps: 1 Introduction, 2 End user license agreement, 3 Select deployment type, 4 Appliance deployment target (highlighted), 5 Set up appliance VM, 6 Select deployment size, 7 Select datastore, 8 Configure network settings, and 9 Ready to complete stage 1. The main content area is titled 'Appliance deployment target' and contains the following text: 'Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.' Below this text are four input fields: 'ESXi host or vCenter Server name' with the value '10.1.156.191', 'HTTPS port' with the value '443', 'User name' with the value 'root', and 'Password' with a masked value '.....'. Each field has an information icon (i) to its right. At the bottom right of the window, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.


9. Click Yes to accept the certificate.

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

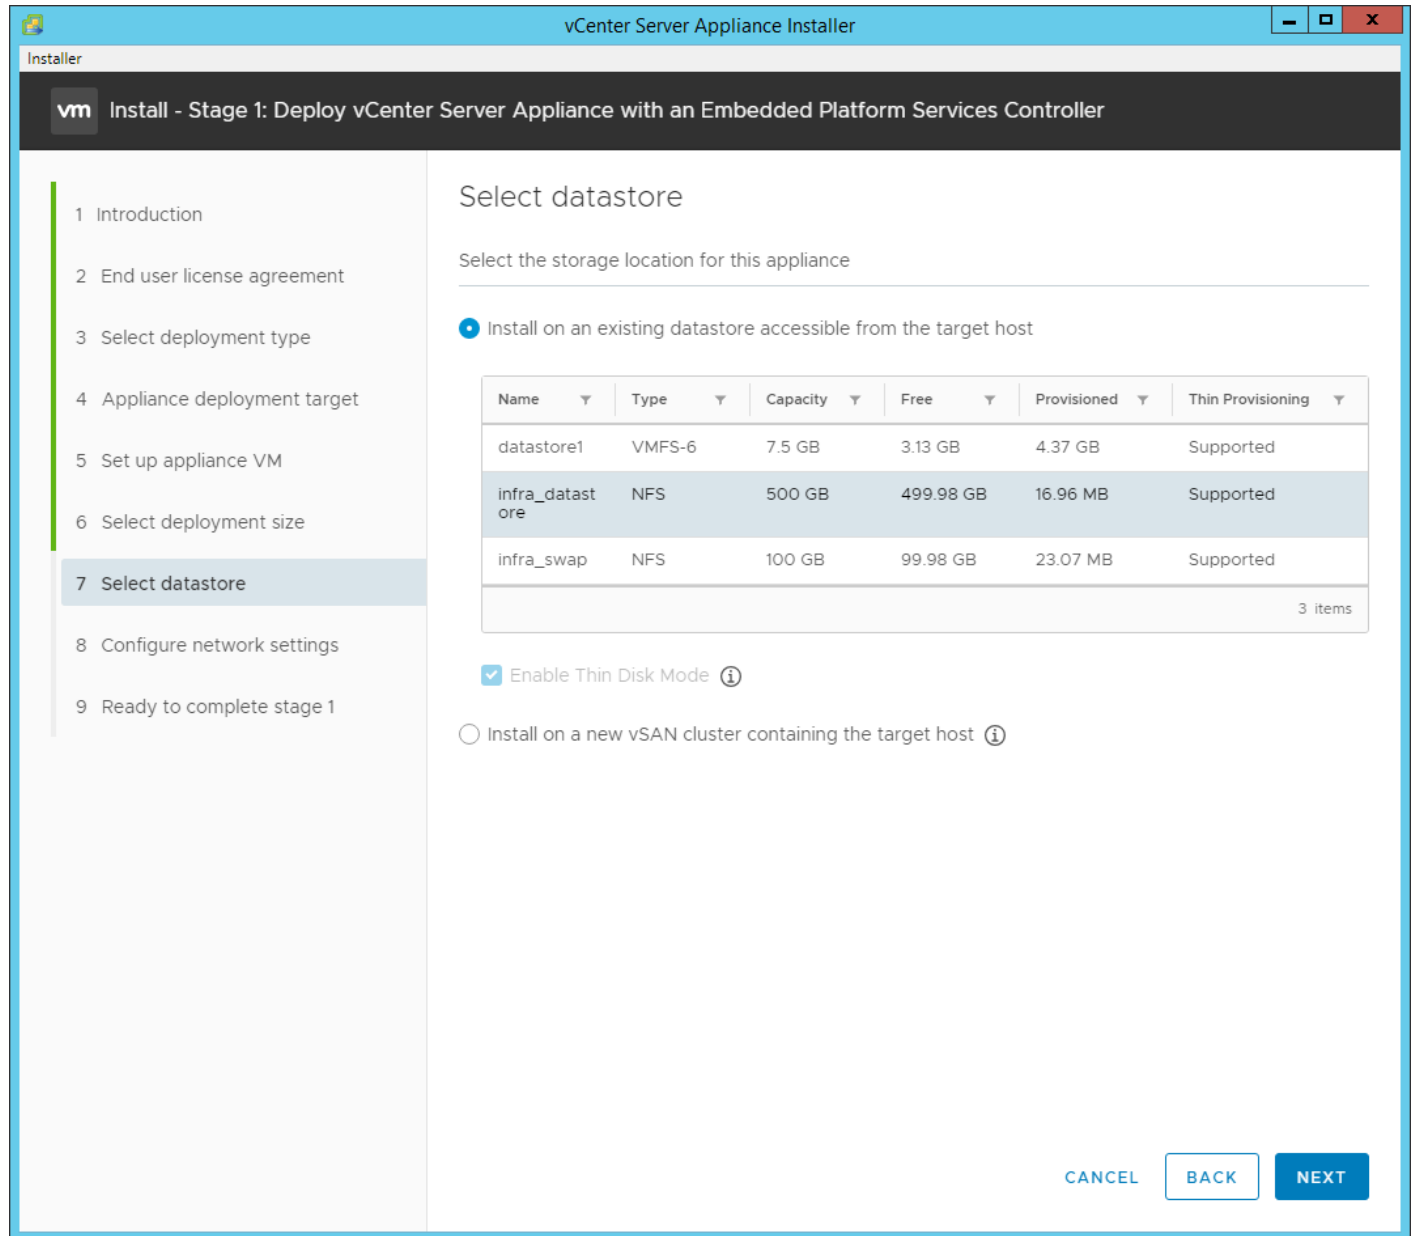


11. In the “Select deployment size” section, Select the deployment size and Storage size. For example, “Small.”

12. Click NEXT.



13. Select infra_datastore. Click NEXT.



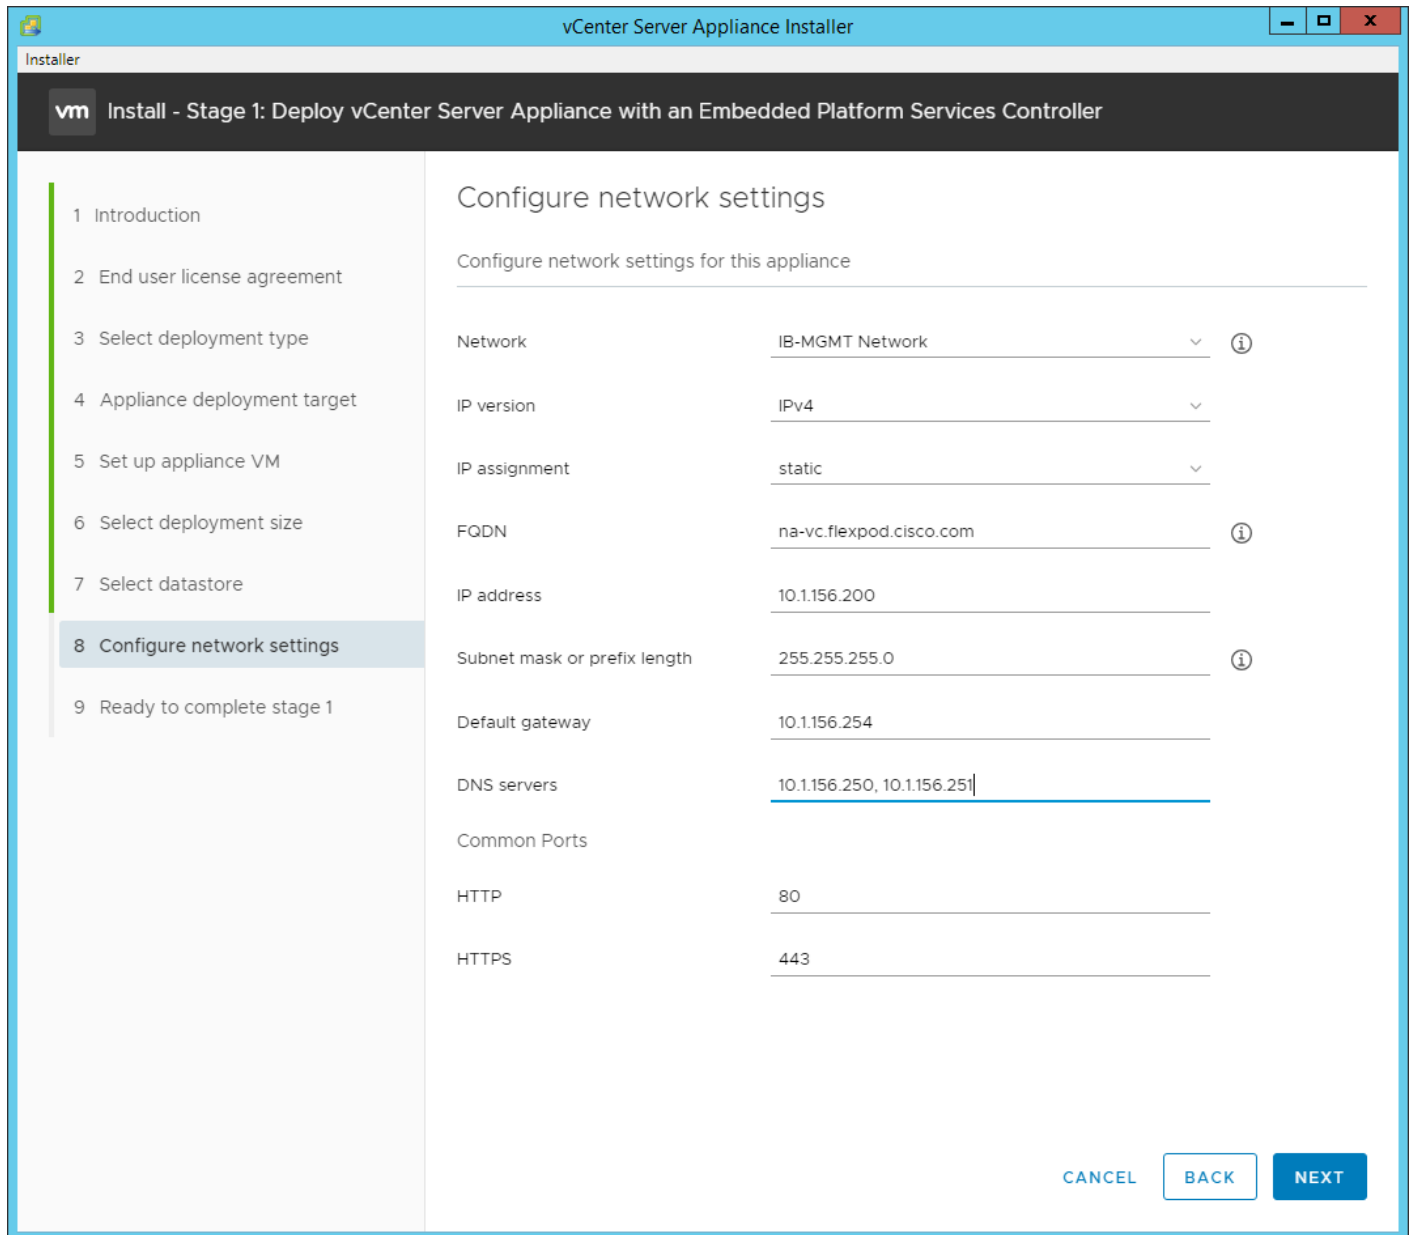
14. In the "Network Settings" section, configure the below settings:

- a. Choose a Network: IB-MGMT Network.



It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it not get moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, and it is attempted to bring up vCenter on a different host than the one it was running on before the shutdown, vCenter will not have a functional network connection. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. Moving vCenter to vSwitch0 to be brought up on a different ESXi host always occurs correctly without requiring vCenter to be up and running.

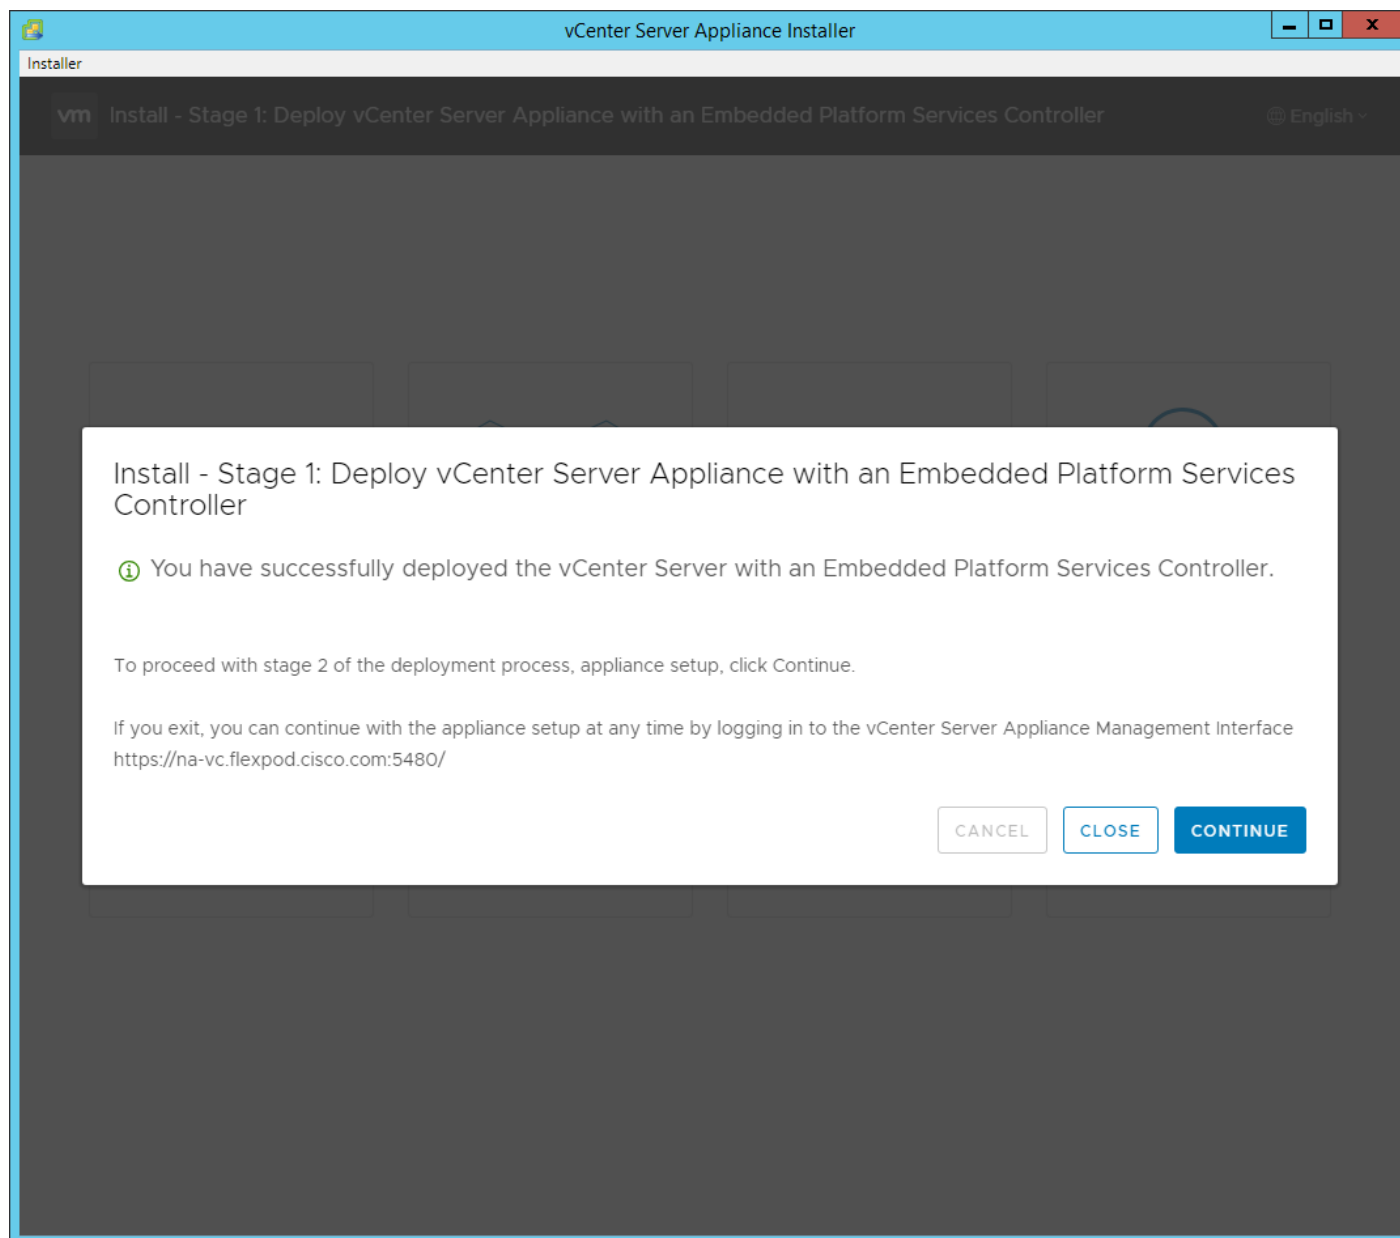
- b. IP version: IPV4
- c. IP assignment: static
- d. System name: <vcenter-fqdn>
- e. IP address: <vcenter-ip>
- f. Subnet mask or prefix length: <vcenter-subnet-mask>
- g. Default gateway: <vcenter-gateway>
- h. DNS Servers: <dns-server1>, <dns-server2>



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.

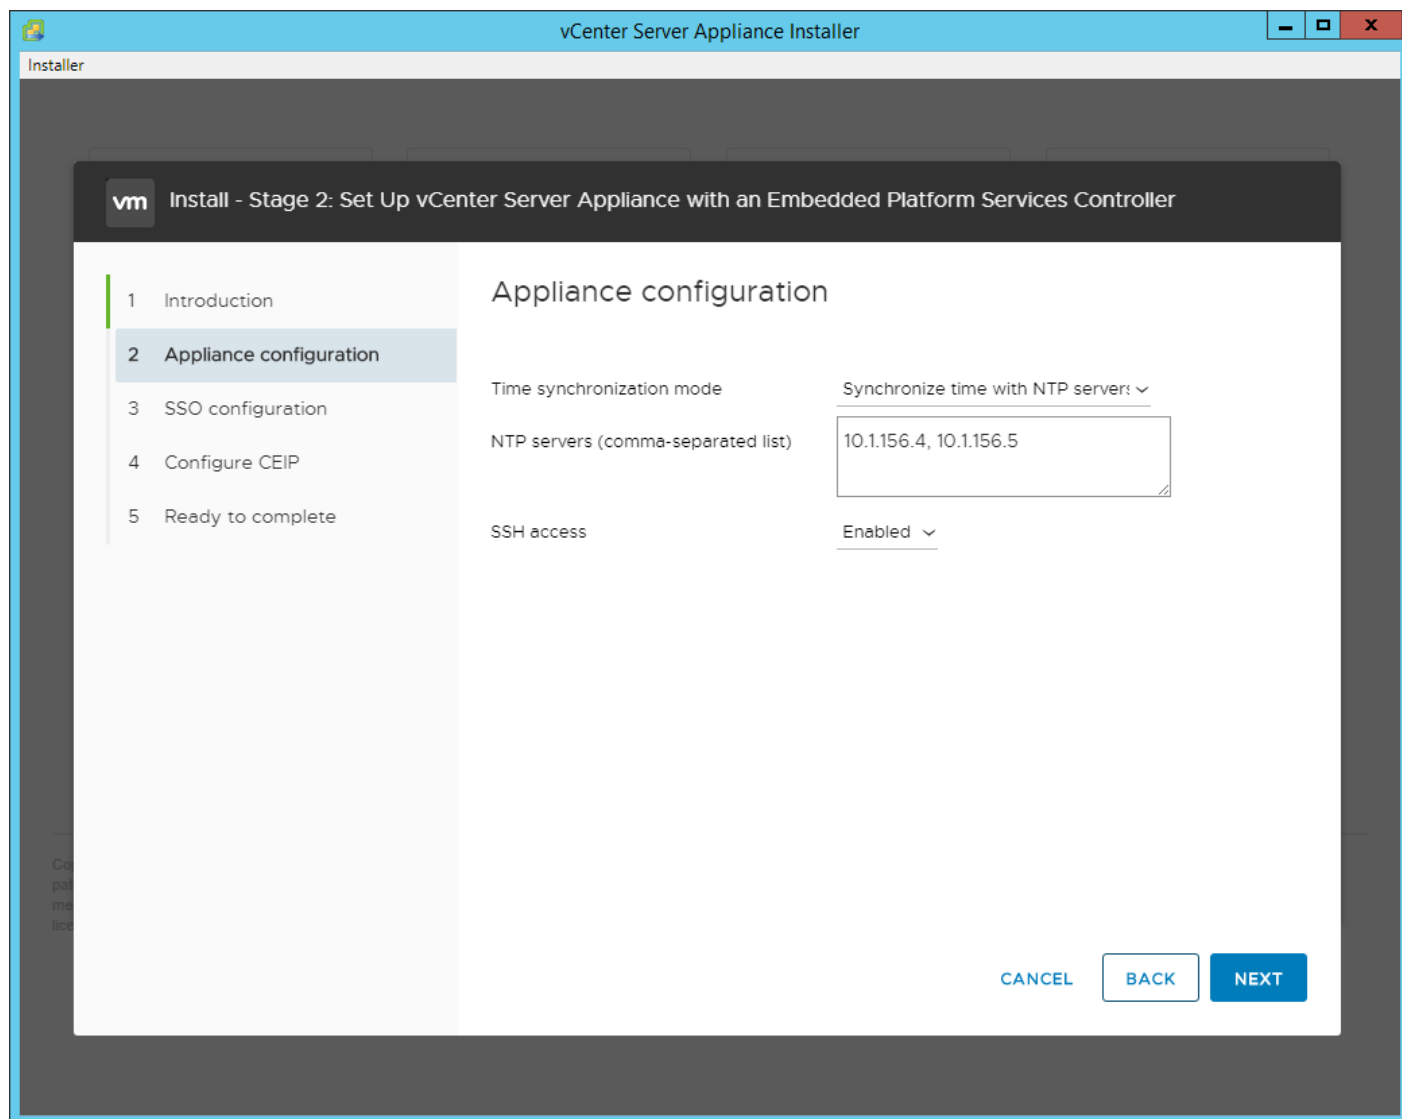


18. Click CONTINUE to proceed with stage 2 configuration.

19. Click NEXT.

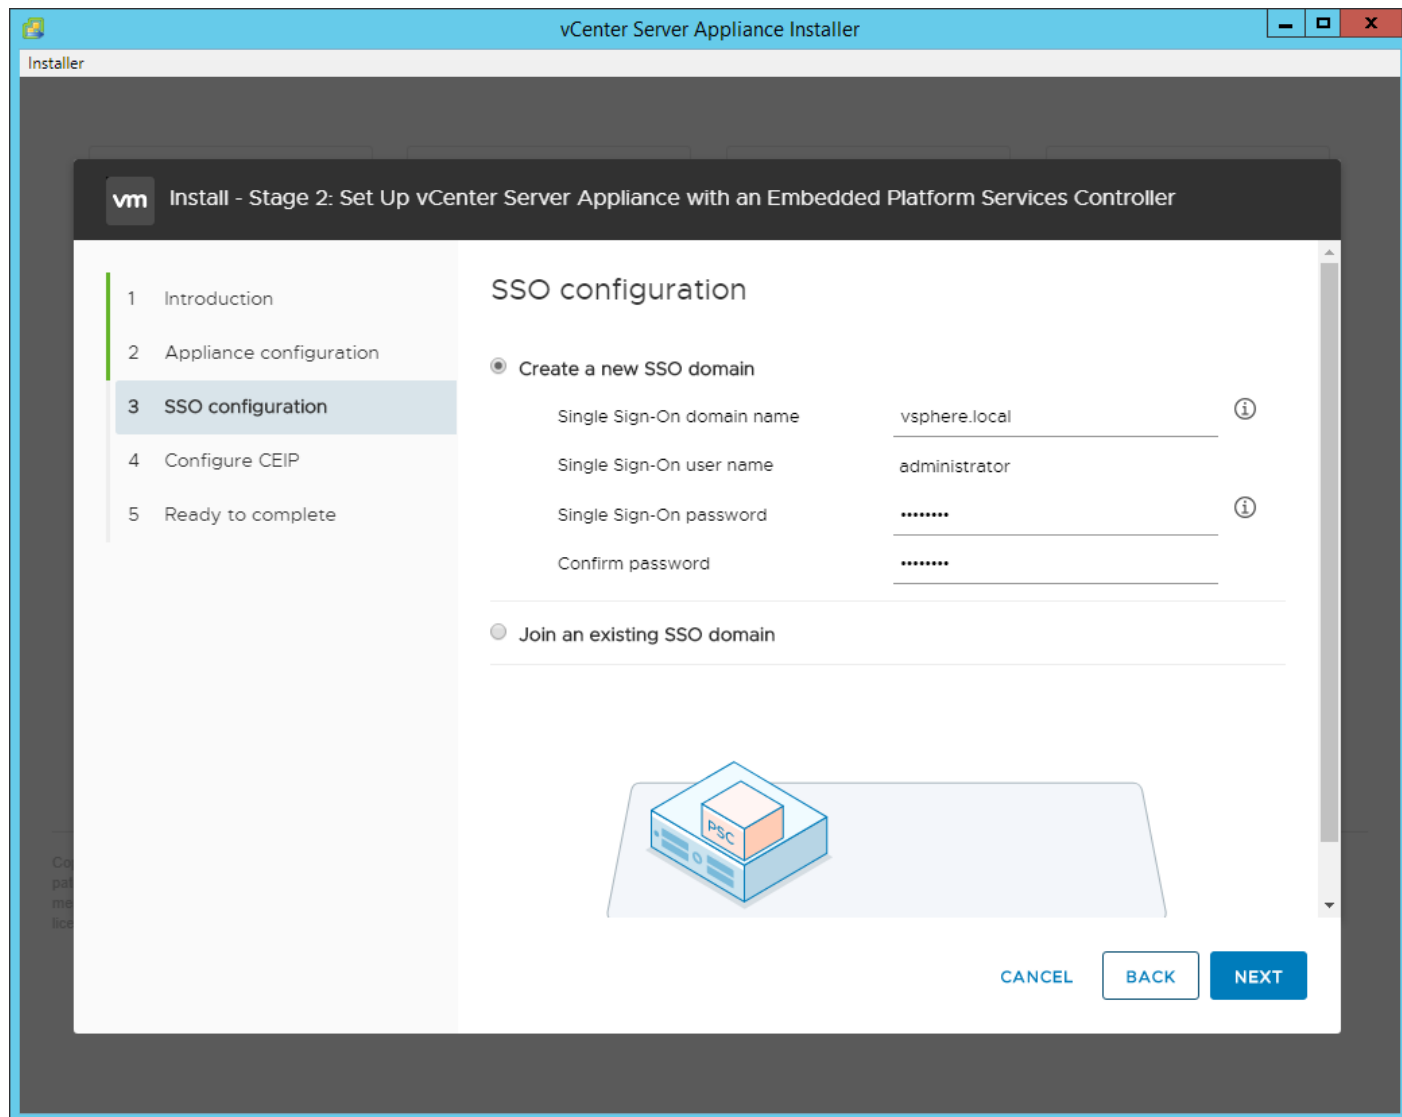
20. In the Appliance Configuration, configure these settings:

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



21. Click NEXT.

22. Complete the SSO configuration as shown below:



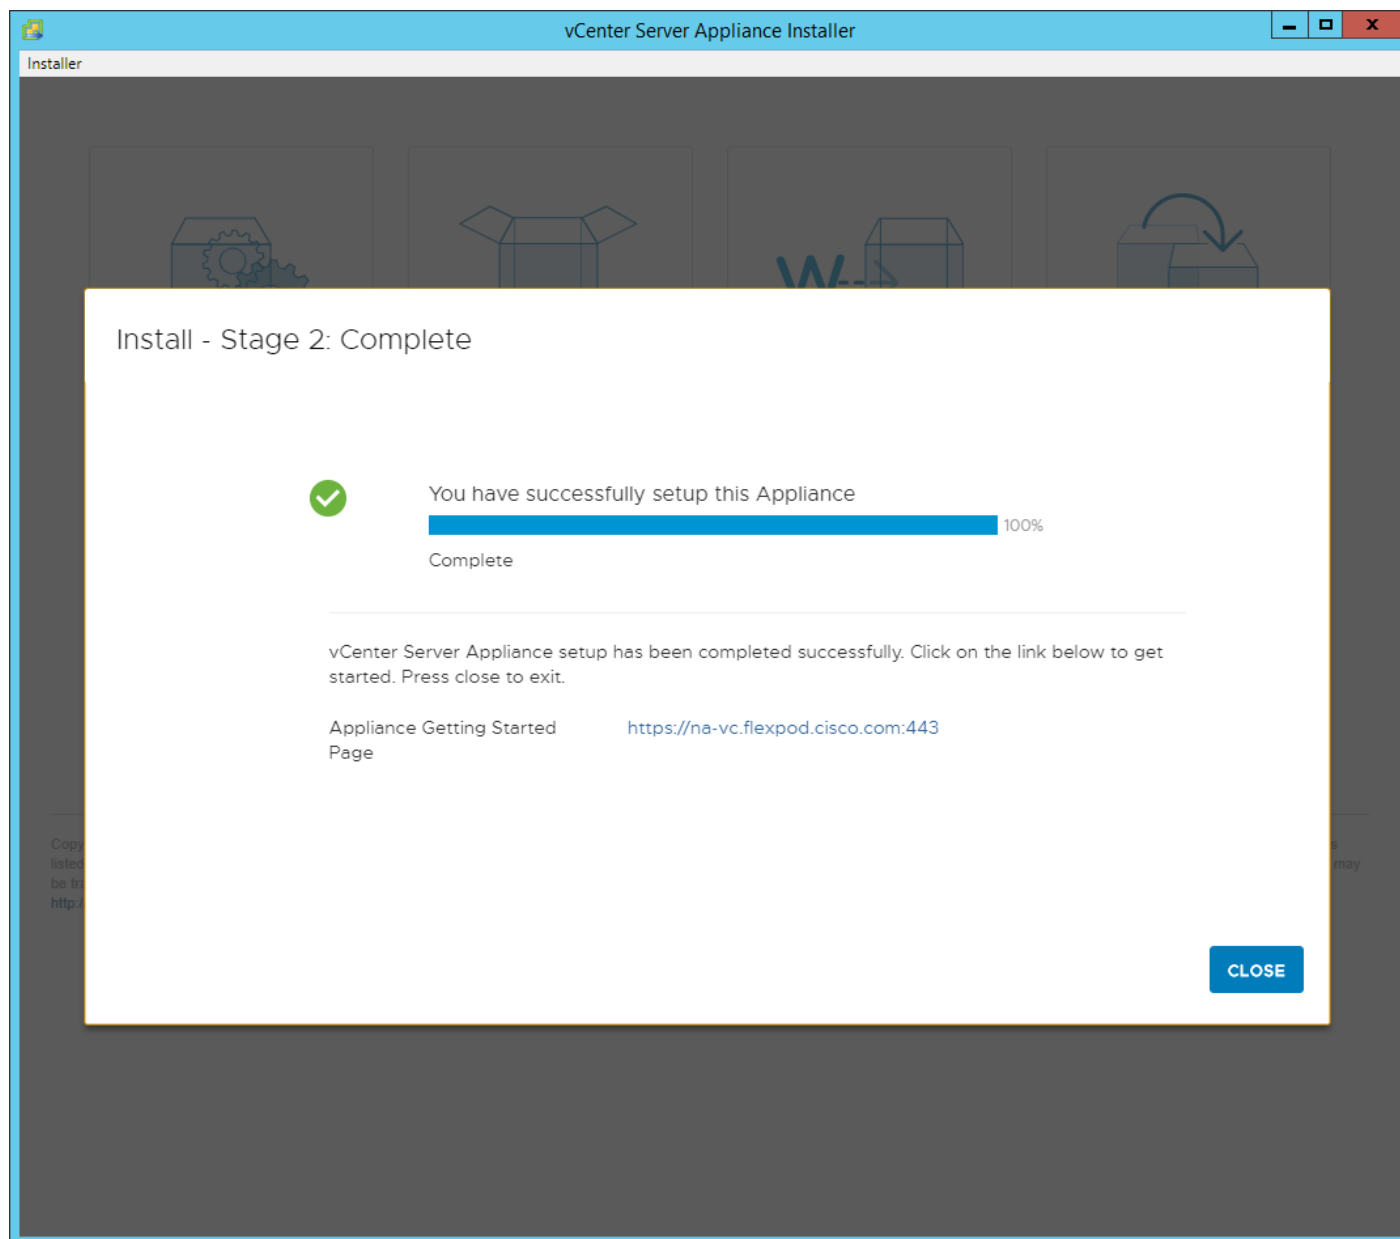
23. Click NEXT.

24. Decide whether to join VMware’s Customer Experience Improvement Program (CEIP).

25. Click NEXT.

26. Review the configuration and click FINISH.

27. Click OK.



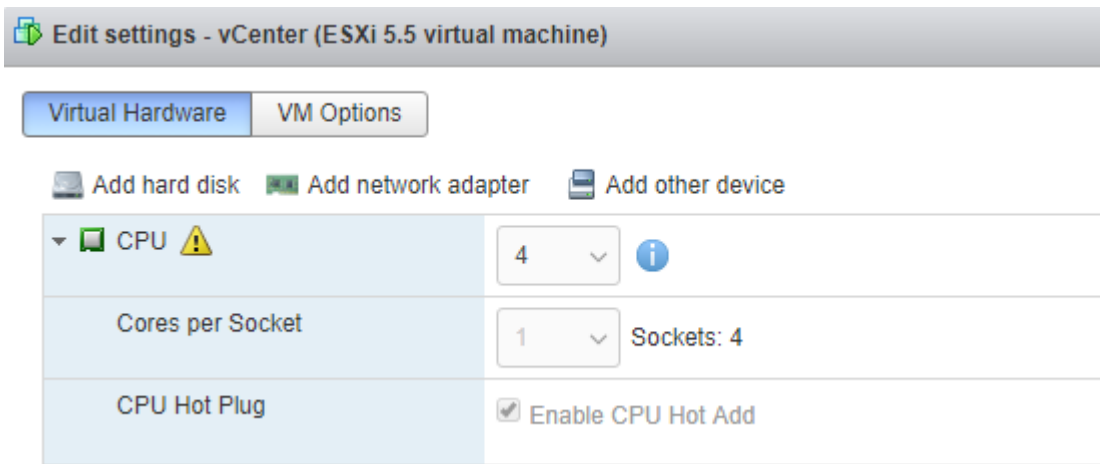
28. Click CLOSE. Eject or unmount the VCSA installer ISO.

Adjust vCenter CPU Settings

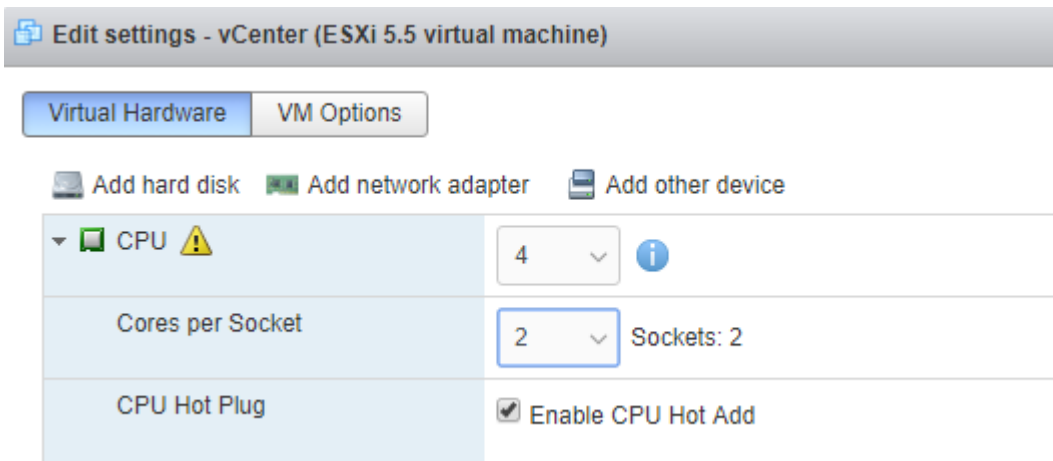
If a vCenter deployment size of Small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the UCS server CPU hardware configuration. Cisco UCS B200 and C220 servers are 2-socket servers. In this validation, the Small deployment size was selected and vCenter was setup for a 4-socket server. This setup will cause issues in the VMware ESXi cluster Admission Control. To resolve the Admission Control issue, follow these 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. On the left, select Virtual Machines.
7. In the center pane, right-click the vCenter VM and select Edit settings.
8. In the Edit settings window, expand CPU and check the value of Sockets.



9. If the number of Sockets is greater than 2, it will need to be adjusted. Click Cancel.
10. If the number of Sockets needs to be adjusted:
 - a. Right-click the vCenter VM and select Guest OS > Shut down. Click Yes on the confirmation.
 - b. Once vCenter is shut down, right-click the vCenter VM and select Edit settings.
 - c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value 2.



- d. Click Save.
- e. Right-click the vCenter VM and select Power > Power on. Wait approximately 10 minutes for vCenter to come up.

Setup VMware vCenter Server

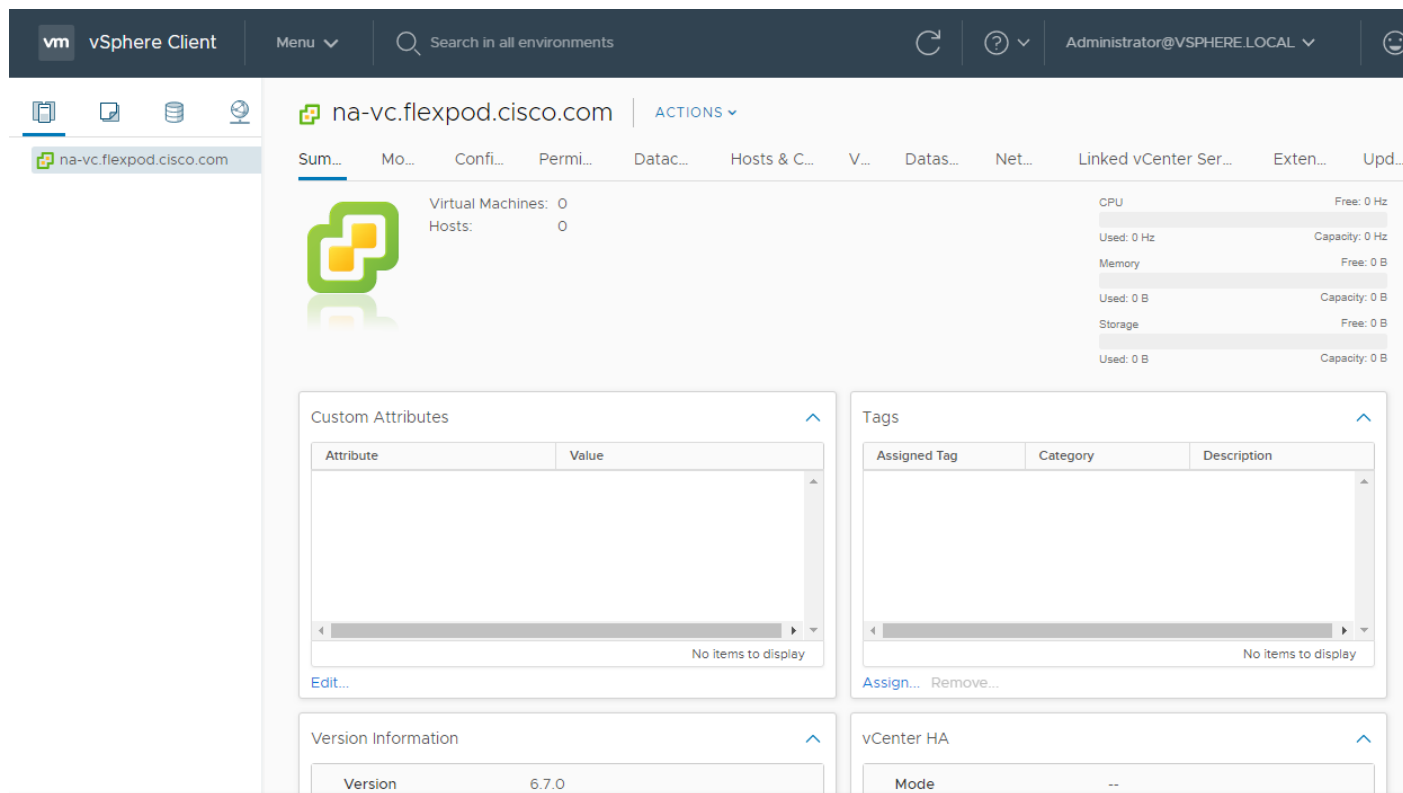
To setup the VMware vCenter Server, follow these steps:

1. Using a web browser, navigate to <https://<vcenter-ip-address>>. You will need to navigate security screens.
2. Select LAUNCH VSPHERE CLIENT (HTML5).



Although the previous versions of this document used FLEX vSphere Web Client, the VMware vSphere HTML5 Client is fully featured in vSphere 6.7U2 and will be used going forward.

3. Log in using the Single Sign-On username (administrator@vsphere.local) and password created during the vCenter installation.



4. In the center pane, select ACTIONS > New Datacenter.
5. Type “FlexPod-DC” in the Datacenter name field.

New Datacenter



Name

FlexPod-DC

Location:

 na-vc.flexpod.cisco.com

CANCEL




OK

6. Click OK.
7. Right-click the datacenter FlexPod-DC in the list in the left pane. Select New Cluster.
8. Name the cluster FlexPod-Management.
9. Turn on DRS and vSphere HA. Do not turn on vSAN.

New Cluster

FlexPod-DC



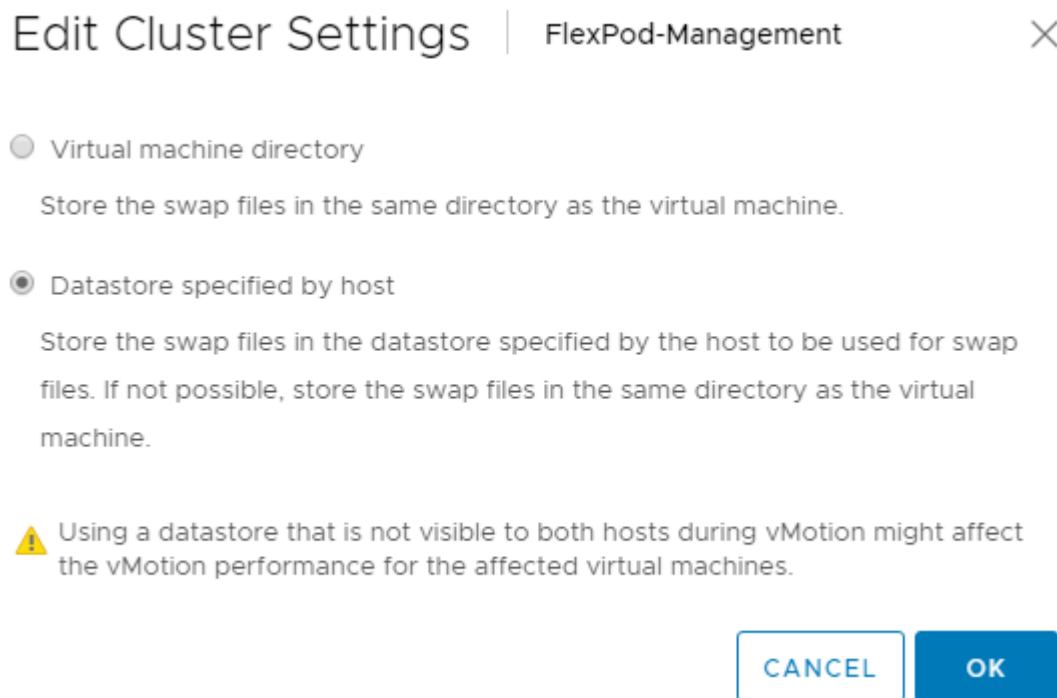
Name	FlexPod-Management
Location	 FlexPod-DC
 DRS	<input checked="" type="checkbox"/>
 vSphere HA	<input checked="" type="checkbox"/>
vSAN	<input type="checkbox"/>

These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

CANCEL

OK

10. Click OK to create the new cluster.
11. Expand “FlexPod-DC”.
12. Right-click “FlexPod-Management” and select Settings.
13. Select Configuration > General in the list located on the left and select EDIT located on the right of General.
14. Select Datastore specified by host and click OK.



15. Right-click “FlexPod-Management” and click Add Hosts.
16. In the IP address or FQDN field, enter either the IP address or the FQDN name of the first VMware ESXi host. Enter root as the Username and the root password. Click NEXT.
17. Verify the Host summary information and click NEXT.
18. Ignore warnings about the host being moved to Maintenance Mode and click FINISH to complete adding the host to the cluster.
19. The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.
20. In the list, right-click the added ESXi host and select Settings.
21. In the center pane under Virtual Machines, select Swap file location.
22. On the right, click EDIT.

23. Select the infra_swap datastore and click OK.

Edit Swap File Location | na-esxi-1.flexpod.cisco.com ✕

Select a location to store the swap files.

Virtual machine directory
Store the swap files in the same directory as the virtual machine.

Use a specific datastore
⚠ Store the swap files in the specified datastore. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

Name	Capacity	Provisioned	Free Space	Type	Thin Provisioned
infra_datastore	3.00 TB	972.06 GB	2.32 TB	NFS	Supported
datastore1	7.50 GB	4.37 GB	3.13 GB	VMFS	Supported
infra_swap	200.00 GB	95.52 MB	199.91 GB	NFS	Supported

3 items

24. In the list under Storage, select Storage Devices. Make sure the NETAPP Fibre Channel Disk 0 is selected.

25. Select the Paths tab.

26. Ensure that 4 fibre channel paths appear, two of which should have the status Active (I/O).

Storage Devices

Refresh | Attach | Detach | Rename... | Turn On LED | Turn Off LED | Erase Partitions... | Mark as Flash Disk | Mark as Local

Name	LUN	Type	Capacity	Datastore	Operational St...	Hardware Acceleration	Drive Type	Transport
NETAPP Fibre Channel Disk (naa.600a0980383137...	0	disk	15.00 GB	datast...	Attached	Supported	HDD	Fibre Channel

Copy All | 1 items

Properties | **Paths** | Partition Details

Enable | Disable

Runtime Name	Status	Target	Name	Preferred
vmhba0:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:02:a9:15 20:04:d0:39:ea:02:a9:15	vmhba0:C0:T1:L0	
vmhba2:C0:T2:L0	Active	20:00:d0:39:ea:02:a9:15 20:01:d0:39:ea:02:a9:15	vmhba2:C0:T2:L0	
vmhba2:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:02:a9:15 20:03:d0:39:ea:02:a9:15	vmhba2:C0:T1:L0	
vmhba0:C0:T2:L0	Active	20:00:d0:39:ea:02:a9:15 20:02:d0:39:ea:02:a9:15	vmhba0:C0:T2:L0	

Copy All | 4 items

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.

To add an AD user authentication to the vCenter, follow these steps:

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 LAUNCH VSPHERE CLIENT (HTML5).
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. Under Menu, select Administration. In the list on the left, under Single Sign On, select Configuration.
5. In the center pane, under Configuration, select the Active Directory Domain tab.
6. Select JOIN AD.
7. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Do not fill in an Organizational unit. Click JOIN.
8. In the list on the left under Deployment, select System Configuration. Select the radio button to select the vCenter, then select REBOOT NODE.
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 vSphere HTML5 Client.
11. Under Menu, select Administration. In the list on the left, under Single Sign On, select Configuration.
12. In the center pane, under Configuration, select the Active Directory Domain tab.

13. Make sure your Active Directory Domain is listed.
14. Select the Identity Sources tab.
15. Select ADD IDENTITY SOURCE.
16. Select the Active Directory (Integrated Windows Authentication) Identity source type. Your AD domain name should be filled in. Leave Use machine account selected and click ADD. Your AD domain should now appear in the Identity Sources list.
17. On the left under Access Control, select Global Permissions.
18. In the center pane, click the + sign to add a Global Permission.
19. In the Add Permission window, select your AD domain for the User.
20. On the search line, enter either the FlexPod Admin username or the Domain Admins group. Leave the Role set to Administrator. Select the Propagate to children checkbox.



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.

21. Click OK to add the selected User or Group. The user or group should now appear in the Global Permissions list with the Administrator role.
22. Log out and log back into the vCenter HTML5 Client as the FlexPod Admin user. You will need to add the domain name to the user, for example, flexadmin@domain.

FlexPod VMware vSphere Distributed Switch (vDS)

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

In the Cisco UCS setup section of this document two sets of vNICs were setup. The vmnic ports associated with the vDS0-A and B vNICs will be placed on the VMware vDS in this procedure. The vMotion VMkernel port will be placed on the vDS.

An IB-MGMT, vMotion, and a VM-Traffic 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 vDS0-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC peer-link interfaces on the switches.

In this document, the infrastructure ESXi management VMkernel ports, the vCenter management interface, and the infrastructure NFS VMkernel ports are left on vSwitch0 to facilitate bringing the virtual environment back up in the event it needs to be completely shut down. The vMotion VMkernel ports are moved to the vDS to allow QoS marking of vMotion to be done at the VLAN level in the vDS if vMotion needs to have QoS policies applied in the future. The vMotion port group is also pinned to Cisco UCS fabric B. Pinning is better done in a vDS to ensure consistency across all ESXi hosts.

Configure the VMware vDS in vCenter

VMware vSphere Web Client

To configure the vDS, follow these steps:

1. After logging into the VMware vSphere HTML5 Client, select Networking under Menu.
2. Right-click the FlexPod-DC datacenter and select Distributed Switch > New Distributed Switch.
3. Give the Distributed Switch a descriptive name (vDS0) and click NEXT.
4. Make sure Distributed switch: 6.6.0 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. Expand the FlexPod-DC datacenter and the newly created vDS. Select the newly created vDS.
8. Right-click the VM-Traffic port group and select Edit Settings.
9. Select VLAN.
10. Select VLAN for VLAN type and enter the VM-Traffic VLAN ID. Click OK.
11. Right-click the vDS and select Settings > Edit Settings.
12. in the Edit Settings window, select Advanced.
13. Change the MTU to 9000. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.

vDS0 - Edit Settings

General

Advanced

MTU (Bytes)	<input type="text" value="9000"/>
Multicast filtering mode	<input type="text" value="Basic"/> ▾
Discovery protocol	
Type	<input type="text" value="Link Layer Discovery Protocol"/> ▾
Operation	<input type="text" value="Both"/> ▾
Administrator contact	
Name	<input type="text"/>
Other details	<input type="text"/>

CANCEL

OK

14. For the first port group used for vMotion, right-click the vDS, select Distributed Port Group, and select New Distributed Port Group.
15. Enter vMotion as the name and click NEXT.
16. Set the VLAN type to VLAN, enter the VLAN ID used for vMotion, click the Customize default policies configuration check box, and click NEXT.
17. Leave the Security options set to Reject and click NEXT.
18. Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.
19. 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. This will pin all vMotion traffic to UCS Fabric Interconnect B except when a failure occurs.

New Distributed Port Group

✓ 1 Name and location

✓ 2 Configure settings

✓ 3 Security

✓ 4 Traffic shaping

5 Teaming and failover

6 Monitoring

7 Miscellaneous

8 Ready to complete

Teaming and failover

Controls load balancing, network failure detection, switches notification, fallback, and uplink failover order.

Load balancing	Route based on originating virtual port
Network failure detection	Link status only
Notify switches	Yes
Failback	Yes

Failover order ⓘ

↑ ↓

Active uplinks
Uplink 2
Standby uplinks
Uplink 1
Unused uplinks

CANCEL

BACK

NEXT

20. Click NEXT.

21. Leave NetFlow disabled and click NEXT.

22. Leave Block all ports set as No 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, right-click the vDS, select Distributed Port Group, and select New Distributed Port Group.

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

26. Set the VLAN type to VLAN, enter the VLAN ID used for In-Band Management, do not select the Customize default policies configuration check box, and click NEXT.

27. Confirm the options and click FINISH to create the port group.

28. Right-click the vDS and select Add and Manage Hosts.

29. Make sure Add hosts is selected and click NEXT.
30. Click the green + sign to add New hosts. Select the one configured FlexPod Management host and click OK. Click NEXT.
31. Select vmnic2 and click Assign uplink. Select Uplink 1 and click OK. Select vmnic3 and click Assign uplink. Select Uplink 2 and click OK.



It is important to assign the uplinks as shown below. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.

vDS0 - Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- 3 Manage physical adapters**
- 4 Manage VMkernel adapt...
- 5 Migrate VM networking
- 6 Ready to complete

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Assign uplink
 Unassign adapter
 View settings

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
▲ na-esxi-1.flexpod.cisco.com			
▲ On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
▲ On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--

32. Click NEXT.
33. Select vmk1 (VMkernel vMotion) and click Assign port group.
34. Select the vMotion destination port group and click OK.
35. Do not migrate the other VMkernel ports.

vDS0 - Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- ✓ 3 Manage physical adapters
- 4 Manage VMkernel adapt...**
- 5 Migrate VM networking
- 6 Ready to complete

Manage VMkernel adapters

Manage and assign VMkernel network adapters to the distributed switch.

Assign port group Reset changes View settings

Host/VMkernel Network Adapters	In Use by Switch	Source Port Group	Destination Port Gr...
na-esxi-1.flexpod.cisco.com			
On this switch			
vmk1 (Reassigned)	vSwitch0	VMkernel-vMotion	vMotion
On other switches/unclaimed			
vmk0	vSwitch0	Management Net...	Do not migrate
vmk2	vSwitch0	VMkernel-Infra-NFS	Do not migrate

CANCEL

BACK

NEXT

36. Confirm the vMotion VMkernel adapter has a valid and correct Destination Port Group and click NEXT.

37. Click NEXT.

38. Click FINISH to complete adding the ESXi host to the vDS.

Add and Configure a VMware ESXi Host in vCenter

This section details the steps to add and configure an ESXi host in vCenter. This section assumes the host has had VMware ESXi 6.7U2 installed, the management IP address set, and the VIC drivers and NetApp NFS Plug-in for VMware VAAI installed. This procedure is initially being run on the second and third ESXi management hosts but can be run on any added ESXi host.

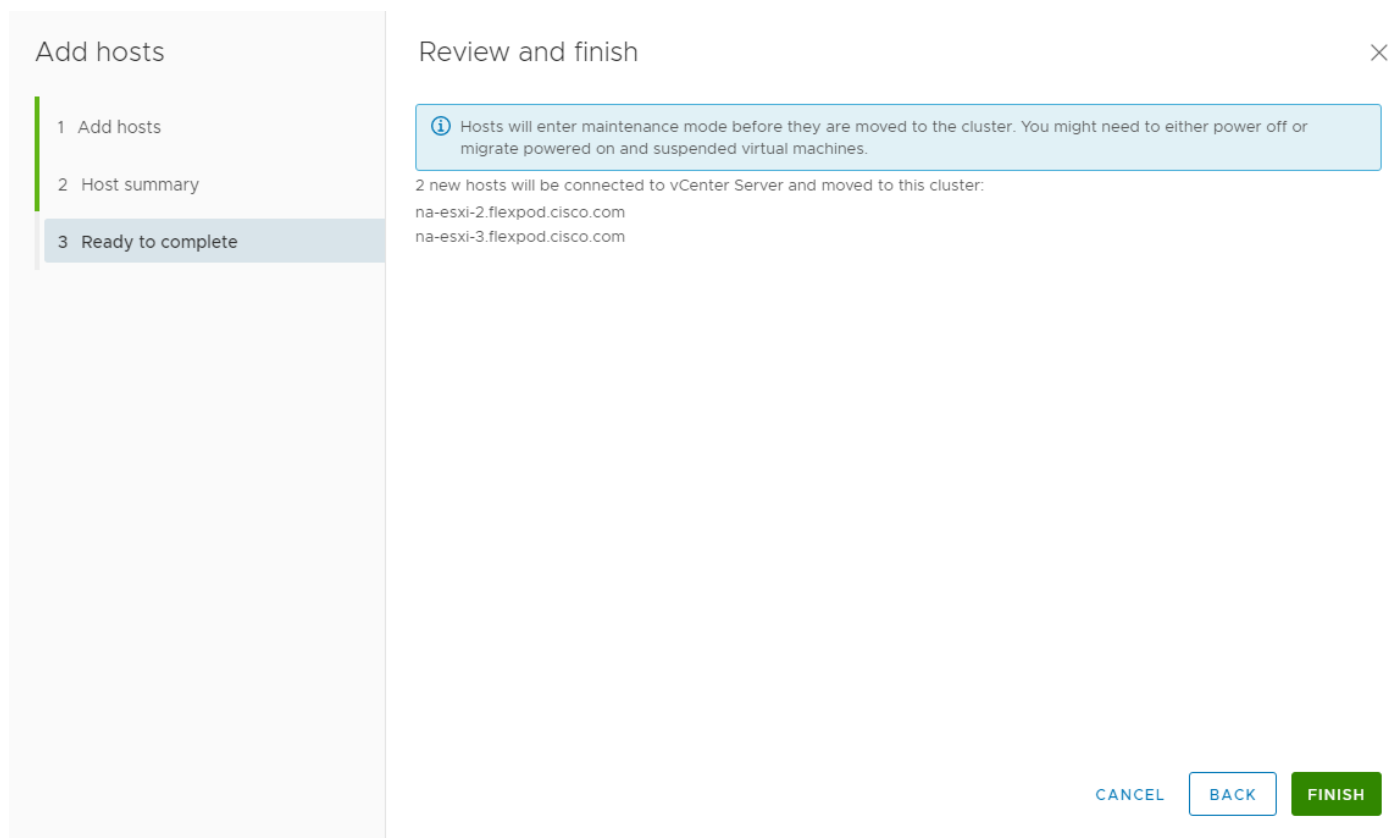
Add the ESXi Host to vCenter

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To add the ESXi host(s) to vCenter, follow these steps:

1. From the Home screen in the VMware vCenter HTML5 Interface, select Menu > Hosts and Clusters.
2. Right-click the "FlexPod-Management" cluster and click Add Hosts.

3. In the IP address or FQDN field, enter either the IP address or the FQDN name of the configured VMware ESXi host. Also enter the user id (root) and associated password. If more than one host is being added, add the corresponding host information. Click NEXT.
4. Select all hosts being added and click OK to accept the certificate(s).
5. Review the host details and click NEXT to continue.
6. Review the configuration parameters and click FINISH to add the host.



7. The added ESXi host(s) will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To set up the VMkernel ports and the virtual switches on the ESXi host, follow these steps:

1. In the vCenter HTML5 Interface, under Hosts and Clusters select the ESXi host.
2. In the center pane, select the Configure tab.
3. In the list, select Virtual switches under Networking.
4. Select vSwitch0.

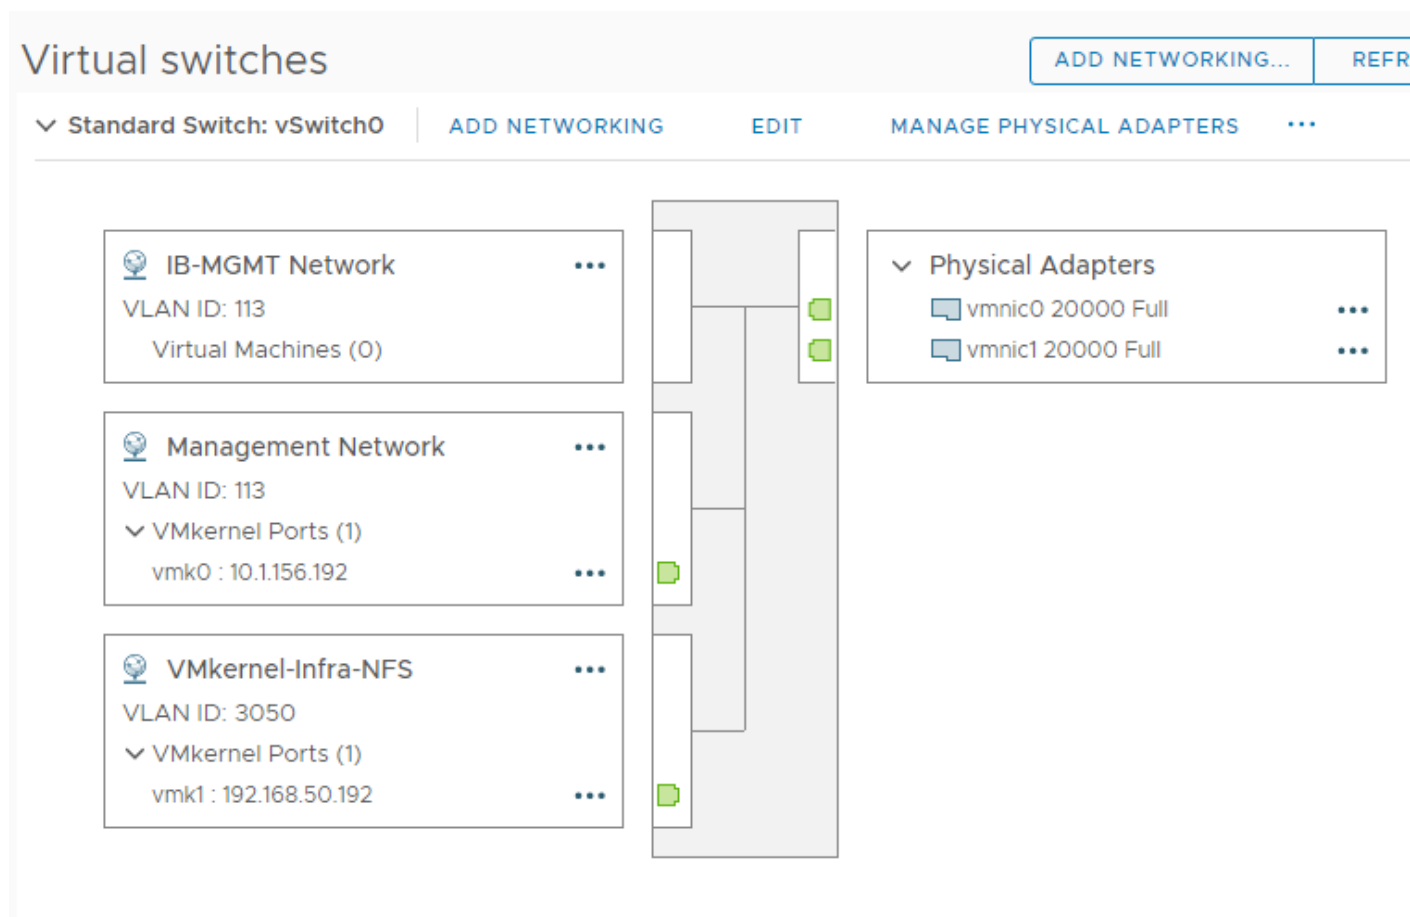
5. Select EDIT to Edit settings.
6. Change the MTU to 9000.
7. Select Teaming and failover located on the left.
8. In the Failover order section, use the arrow icons to move the vmnics until both are Active adapters.

vSwitch0 - Edit Settings

The screenshot shows the 'vSwitch0 - Edit Settings' dialog box. On the left, a sidebar contains tabs for 'Properties', 'Security', 'Traffic shaping', and 'Teaming and failover', with 'Teaming and failover' being the active tab. The main area is divided into two sections: 'Load balancing' and 'Network failure detection'. Under 'Load balancing', the setting is 'Route based on originating virtual port'. Under 'Network failure detection', the settings are 'Link status only', 'Notify switches' (Yes), and 'Failback' (Yes). Below this is the 'Failover order' section, which includes a list of adapters categorized into 'Active adapters', 'Standby adapters', and 'Unused adapters'. In the 'Active adapters' list, 'vmnic0' and 'vmnic1' are shown. To the right of this list is a detailed view of the selected adapter, 'vmnic1', showing its properties: Adapter Name (vmnic1), Location (PCI 0000:62:00.1), Driver (nenic), Status (Connected), Actual speed, Duplex (20000 Mb, Full Duplex), Configured speed, Duplex (20000 Mb, Full Duplex), Networks (10.1156.250-10.1156.251 (VLAN113)), and SR-IOV Status (Not supported). At the bottom right of the dialog are 'CANCEL' and 'OK' buttons.

9. Click OK.
10. In the center pane, select VM Network and click ... > Edit Settings to edit settings.
11. Rename the port group IB-MGMT Network and enter <ib-mgmt-vlan-id> in the VLAN ID field.
12. Click OK to finalize the edits for the IB-MGMT Network.
13. Located on the left under Networking, select VMkernel adapters.
14. In the center pane, click Add Networking.
15. Make sure VMkernel Network Adapter is selected and click NEXT.
16. Select an existing standard switch and click BROWSE. Select vSwitch0 and click OK. Click NEXT.

17. For Network label, enter VMkernel-Infra-NFS.
18. Enter <infra-nfs-vlan-id> for the VLAN ID.
19. Select Custom for MTU and make sure 9000 is entered.
20. Leave the Default TCP/IP stack selected and do not select any of the Enabled services. Click NEXT.
21. Select Use static IPv4 settings and enter the IPv4 address and subnet mask for the Infra-NFS VMkernel port for this ESXi host.
22. Click NEXT.
23. Review the settings and click FINISH to create the VMkernel port.
24. On the left under Networking, select Virtual switches. Then select vSwitch0. The properties for vSwitch0 should be similar to the following example:



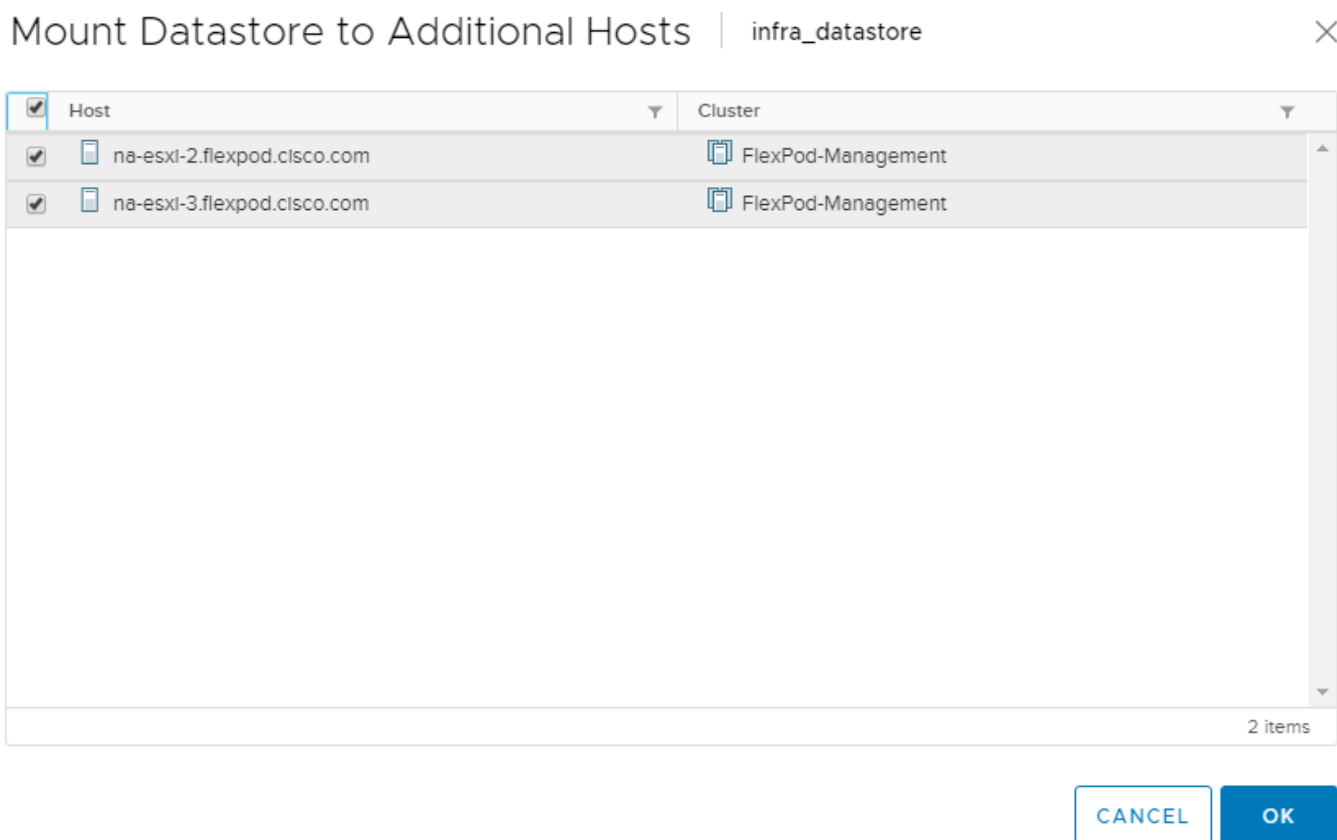
25. Repeat this procedure for all hosts being added.

Mount Required Datastores

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To mount the required datastores, follow these steps on the ESXi host(s):

1. From the vCenter Home screen, select Menu > Storage.
2. Located on the left, expand FlexPod-DC.
3. Located on the left, right-click infra_datastore and select Mount Datastore to Additional Hosts.
4. Select the ESXi host(s) and click OK.



5. Repeat this process to mount the infra_swap datastore to the ESXi host(s).
6. On the left, select infra_datastore. In the center pane, select Hosts. Verify the ESXi host(s) now has the datastore mounted. Repeat this process to also verify that infra_swap is also mounted.

Configure NTP on ESXi Host

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

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

1. In the vCenter HTML5 Interface, under Hosts and Clusters select the ESXi host.
2. In the center pane, select the Configure tab.
3. In the list under System, select Time Configuration.
4. Located on the right, click EDIT.
5. Select Use Network Time Protocol (enable NTP client).
6. Enter the two Nexus switch NTP IP addresses in the NTP servers box separated by a comma.
7. Click the Start NTP Service checkbox.
8. Use the drop-down list to select Start and stop with host.

Edit Time Configuration na-esxi-2.flexpod.cisco.com ✕

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

Manually configure the date and time on this host

01/02/2020
18:29:52

(time is in ISO 8601 format)

Use Network Time Protocol (Enable NTP client)

NTP Servers

10.1.156.4,10.1.156.5

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

NTP Service Status:	Stopped
	<input checked="" type="checkbox"/> Start NTP Service
NTP Service Startup Policy:	Start and stop with host ▼

CANCEL
OK

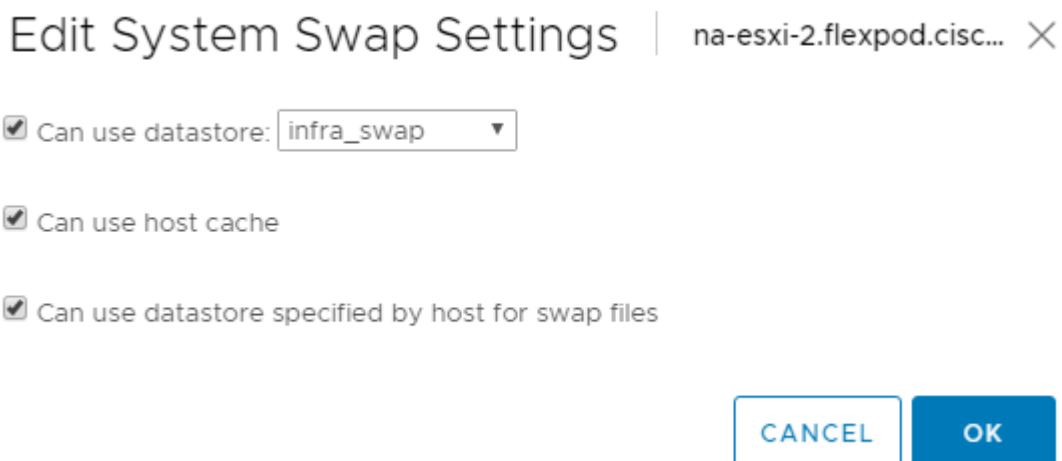
9. Click OK to save the configuration changes.
10. Verify that NTP service is now enabled and running and the clock is now set to approximately the correct time.

Configure ESXi Host Swap

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To configure host swap on the ESXi host(s), follow these steps on the host:

1. In the vCenter HTML5 Interface, under Hosts and Clusters select the ESXi host.
2. In the center pane, select the Configure tab.
3. In the list under System, select System Swap.
4. Located on the right, click EDIT.
5. Select Can use datastore and use the drop-down list to select infra_swap. Leave all other settings unchanged.



6. Click OK to save the configuration changes.
7. In the list under Virtual Machines, select Swap File Location.
8. Located on the right, select Edit.
9. Select infra_swap and click OK.

Check ESXi Host Fibre Channel Pathing

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To ensure that the ESXi host(s) boot disk contains all required fibre channel paths, follow these steps:

1. In the list under Storage, select Storage Devices. Make sure the NETAPP Fibre Channel Disk is selected.
2. Select the Paths tab.
3. Ensure that 4 fibre channel paths appear, two of which should have the status Active (I/O).

Storage Devices

Refresh | Attach | Detach | Rename... | Turn On LED | Turn Off LED | Erase Partitions... | Mark as Flash Disk | Mark as Local

Name	LUN	Type	Capacity	Datastore	Operational St...	Hardware Acceleration	Drive Type	Transport
NETAPP Fibre Channel Disk (naa.600a0980383137...)	0	disk	15.00 GB	datast...	Attached	Supported	HDD	Fibre Channel

Copy All | 1 items

Properties | **Paths** | Partition Details

Enable | Disable

Runtime Name	Status	Target	Name	Preferred
vmhba0:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:02:a9:15 20:04:d0:39:ea:02:a9:15	vmhba0:C0:T1:L0	
vmhba2:C0:T2:L0	Active	20:00:d0:39:ea:02:a9:15 20:01:d0:39:ea:02:a9:15	vmhba2:C0:T2:L0	
vmhba2:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:02:a9:15 20:03:d0:39:ea:02:a9:15	vmhba2:C0:T1:L0	
vmhba0:C0:T2:L0	Active	20:00:d0:39:ea:02:a9:15 20:02:d0:39:ea:02:a9:15	vmhba0:C0:T2:L0	

Copy All | 4 items

Add the ESXi Host(s) to the VMware Virtual Distributed Switch

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To add the ESXi host(s) to the VMware vDS, follow these steps on the host:

1. After logging into the VMware vSphere HTML5 Client, select Networking under Menu.
2. Right-click the vDS (vDS0) and select Add and Manage Hosts.
3. Make sure Add hosts is selected and click NEXT.
4. Click the green + sign to add New hosts. Select the configured FlexPod Management host(s) and click OK. Click NEXT.
5. Select vmnic2 on each host and click Assign uplink. Select Uplink 1 and click OK. Select vmnic3 on each host and click Assign uplink. Select Uplink 2 and click OK. If more than one host is being connected to the vDS, use the Apply this uplink assignment to the rest of the hosts checkbox.




It is important to assign the uplinks as shown below. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.










vDS0 - Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- 3 Manage physical adapters**
- 4 Manage VMkernel adapt...
- 5 Migrate VM networking
- 6 Ready to complete

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

 Assign uplink
  Unassign adapter
  View settings

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
 vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
 vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
▲ On other switches/unclaimed			
 vmnic0	vSwitch0	--	--
 vmnic1	vSwitch0	--	--
 vmnic4	--	--	--
 vmnic5	--	--	--
▲  na-esxi-3.flexpod.cisco.com			
▲ On this switch			
 vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
 vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
▲ On other switches/unclaimed			

CANCEL

BACK

NEXT

6. Click NEXT.
7. Do not migrate any VMkernel ports and click NEXT.
8. Do not migrate any VM ports and click NEXT.
9. Click FINISH to complete adding the ESXi host(s) to the vDS.

Add the vMotion VMkernel Port to the ESXi Host

ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To add the vMotion VMkernel Port to the ESXi host(s) on the VMware vDS, follow these steps on the host:

1. In the vCenter HTML5 Interface, under Hosts and Clusters select the ESXi host.
2. In the center pane, select the Configure tab.
3. In the list under Networking, select VMkernel adapters.
4. Select Add Networking to Add host networking.
5. Make sure VMkernel Network Adapter is selected and click NEXT.

6. Select BROWSE to the right of Select an existing network.
7. Select vMotion on the vDS and click OK.
8. Click NEXT.
9. Make sure the Network label is vMotion with the vDS in parenthesis. Using the pulldown, select Custom for MTU and make sure the MTU is set to 9000. Select the vMotion TCP/IP stack and click NEXT.
10. Select Use static IPv4 settings and input the host's vMotion IPv4 address and Subnet mask.
11. Click NEXT.
12. Review the parameters and click FINISH to add the vMotion VMkernel port.

na-esxi-2.flexpod.cisco.com - Add Networking

- ✓ 1 Select connection type
- ✓ 2 Select target device
- ✓ 3 Port properties
- ✓ 4 IPv4 settings
- 5 Ready to complete**

Ready to complete

Review your settings selections before finishing the wizard.

Distributed port group	vMotion
Distributed switch	vDSO
vMotion	Enabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled

NIC settings

MTU	9000
TCP/IP stack	vMotion

IPv4 settings

IPv4 address	192.168.100.192 (static)
Subnet mask	255.255.255.0

CANCEL

BACK

FINISH

13. If NetApp VSC is installed, under Hosts and Clusters, right-click the host and select NetApp VSC > Set Recommended Values. Reboot the host.

VMware ESXi 6.7U2 TPM Attestation

If your Cisco UCS servers have Trusted Platform Module (TPM) 2.0 modules installed, the TPM can provide assurance that ESXi has booted with UEFI Secure Boot enabled and using only digitally signed code. In the Cisco

UCS section of this document, UEFI secure boot was enabled in the boot policy. A server can boot with UEFI Secure Boot with or without a TPM 2.0 module. If it has a TPM, VMware vCenter can attest that the server booted with UEFI Secure Boot. Follow these steps:

1. If your Cisco UCS servers have TPM 2.0 modules installed, TPM Attestation can be verified in the vSphere HTML5 Client. To get to the HTML5 client from the Web Client, click "Launch vSphere Client (HTML5)" in the upper center portion of the Web Client window.
2. From the Hosts and Clusters window in the vSphere Client, select the vCenter. In the center pane, select Monitor > Security. The Attestation status will appear as shown below:

Name ↑	Attestation	Last verified	TPM version	TXT	Message
na-esxi-1.flexpod.cisco.com	Passed	01/01/2020, 4:4...	2.0	N/A	
na-esxi-2.flexpod.cisco.com (Maintenance Mode)	Passed	01/02/2020, 5:...	2.0	N/A	
na-esxi-3.flexpod.cisco.com (Maintenance Mode)	Passed	01/02/2020, 5:...	2.0	N/A	



It may be necessary to disconnect and reconnect a host from vCenter to get it to pass attestation the first time.

FlexPod Management Tools Setup

Cisco UCS Manager Plug-in for VMware vSphere Web Client

The Cisco UCS Manager Plug-in for VMware vSphere Web Client allows administration of UCS domains through the VMware vCenter administrative interface. The 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 Cisco UCS B and C-Series Servers
- View VIF paths for servers
- Launch the Cisco UCS Manager GUI
- Launch the KVM consoles of Cisco UCS Servers
- Switch the existing state of the locator LEDs

The 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.

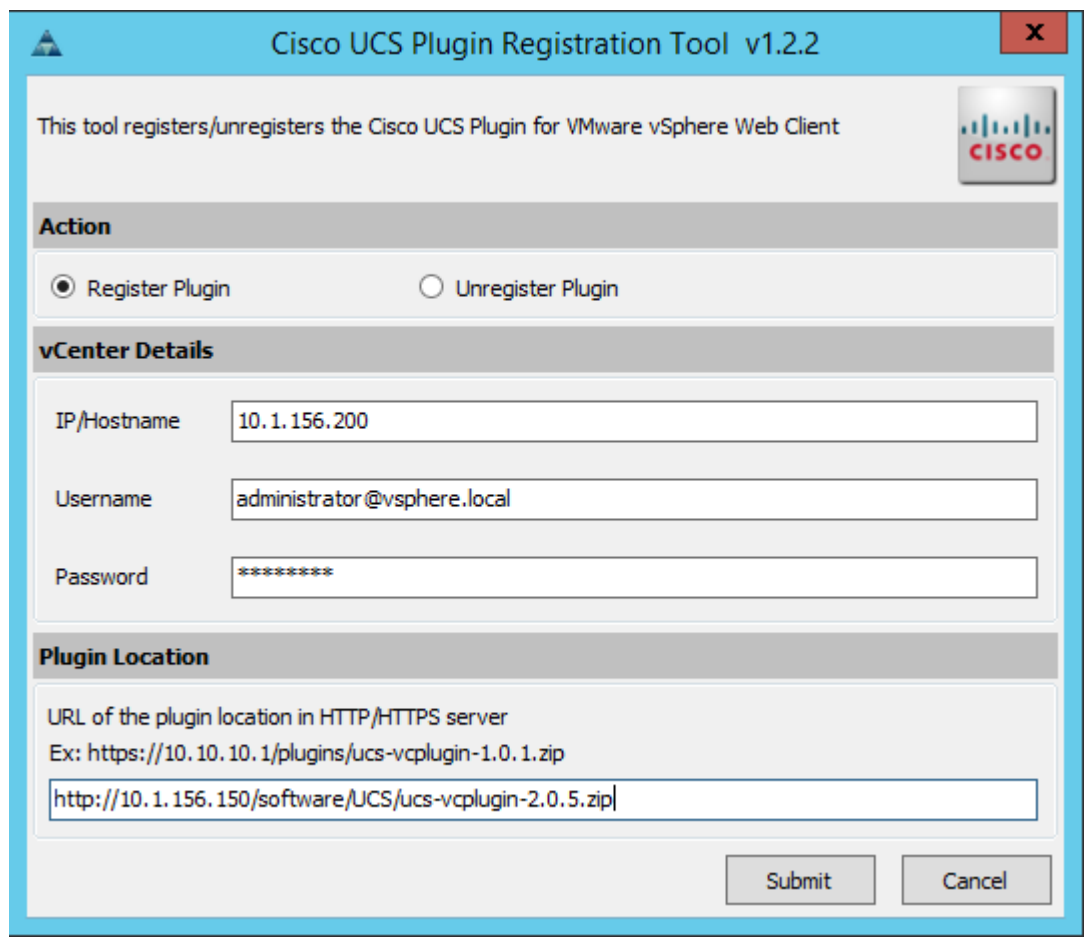


This plug-in is currently only available for the VMware vSphere Web Client and not for the HTML5 client.


Cisco UCS Manager Plug-in Installation

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

1. Download the plugin and registration tool from:
<https://software.cisco.com/download/home/286282669/type/286282010/release/2.0.5?catid=282558030>.
2. Place the downloaded ucs-vcplugin-2.0.5.zip file onto the web server used for hosting the ONTAP software and VMware ESXi ISO.
3. Unzip the Cisco_UCS_Plugin_Registration_Tool_1.2.2.zip and open the executable file within it.
4. Leave Register Plugin selected for the Action, fill in the following vCenter information and click Submit:
 - a. IP/Hostname
 - b. Username
 - c. Password
 - d. URL that plugin has been uploaded to



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. This issue can be taken care of after the plugin has been registered, click OK to close the Information dialogue box.
7. Click OK to confirm that the Cisco UCS Plugin registered successfully.
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, open the BASH shell, and type:

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



This will add "allowHttp=true" to the end of the webclient.properties file. Make sure 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.

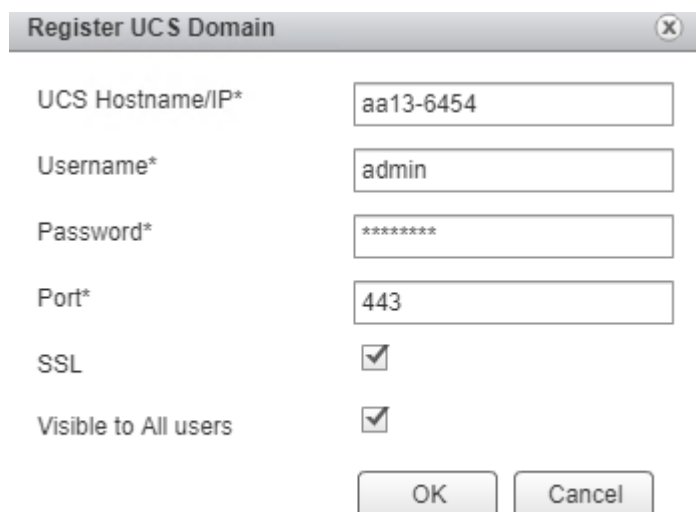
FlexPod UCS Domain Registration

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 UCS Domain, follow these steps:

1. Open the vSphere Web Client.
2. Select the Home from the Navigator or drop-down list and select 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:
 - a. UCS Hostname/IP
 - b. Username
 - c. Password
 - d. Port (if different than 443)
 - e. Leave SSL selected and click the Visible to All users option



UCS Hostname/IP*	aa13-6454
Username*	admin
Password*	*****
Port*	443
SSL	<input checked="" type="checkbox"/>
Visible to All users	<input checked="" type="checkbox"/>

OK Cancel

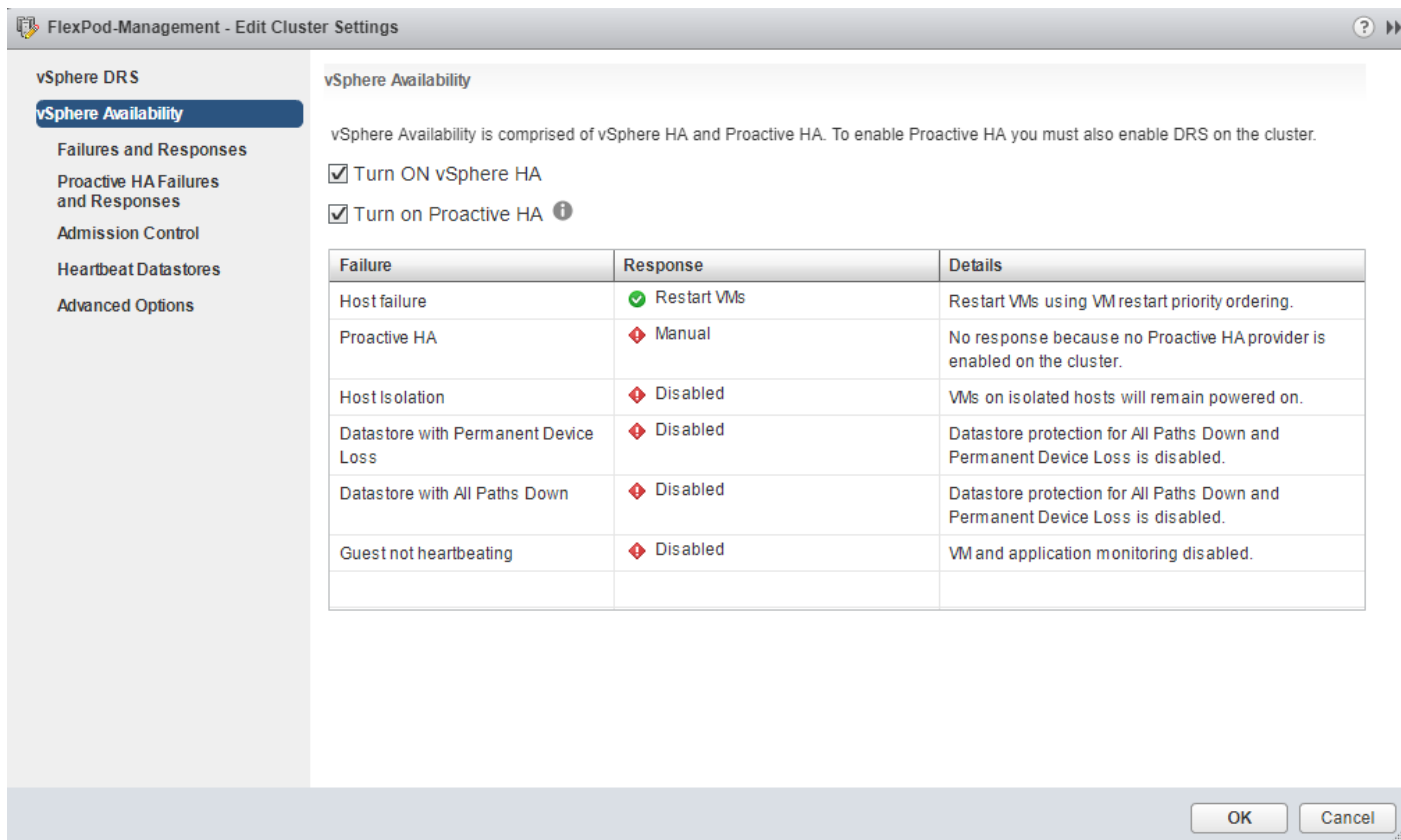
4. Click OK then click OK again to register the UCS Domain.

Register the Proactive High Availability Provider for Cisco UCS (Optional)

Cisco UCS Provider for Proactive high availability feature allows the system to assess the health of the server running the ESXi host. It assesses if the server is healthy, moderately or severely degraded. Any fault which occurs from the Cisco approved predefined list of faults with critical or major severity is reported to the vCenter.

To register the proactive HA provider for Cisco UCS, follow these steps:

1. From the Cisco UCS Management Center page in the vSphere Web Client, select the Proactive HA Registration tab.
2. Enter administrator@vsphere.local and the associated password. Click Register.
3. Click OK to confirm registration.
4. Enable the Proactive HA Provider for Cisco UCS by selecting Hosts and Clusters > FlexPod-Management > Configure > vSphere Availability in the vSphere Web Client. Select Turn on Proactive HA.



- In the list on the left, select Proactive HA Failures and Responses.
- Select the Automated Automation Level, Mixed Mode Remediation and the checkbox next to Cisco UCS Provider. Click OK.

The screenshot shows the 'FlexPod-Management - Edit Cluster Settings' window. The left sidebar contains a navigation menu with the following items: vSphere DRS, vSphere Availability, Failures and Responses, Proactive HA Failures and Responses (highlighted), Admission Control, Heartbeat Datastores, and Advanced Options. The main content area is titled 'Proactive HA Failures and Responses' and includes the following text: 'You can configure how Proactive HA responds when a provider has notified its health degradation to vCenter, indicating a partial failure of that host. In the event of a partial failure, vCenter Server can proactively migrate the host's running VMs to a healthier host.'

Below the text are two configuration options:

- Automation Level: Automated
- Remediation: Mixed mode

Below these options is a paragraph: 'Select the check boxes to enable Proactive HA providers for this cluster. Providers appear below when their corresponding vSphere Web Client plugin has been installed and the providers support every host in the cluster. Click on the edit link to view/edit the failure conditions supported by the provider.'

	Proactive HA provider	Failure conditions blocked	Action
<input checked="" type="checkbox"/>	Cisco UCS Provider	No	edit

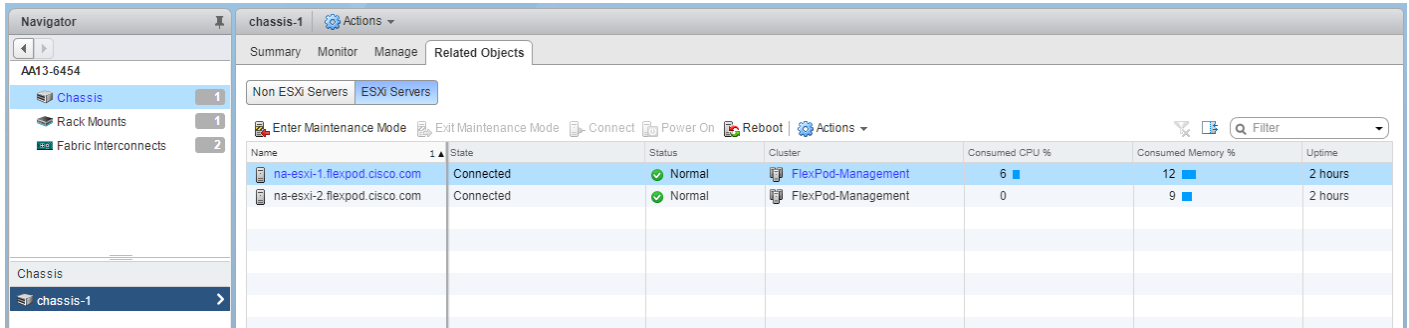
At the bottom right of the window are 'OK' and 'Cancel' buttons.



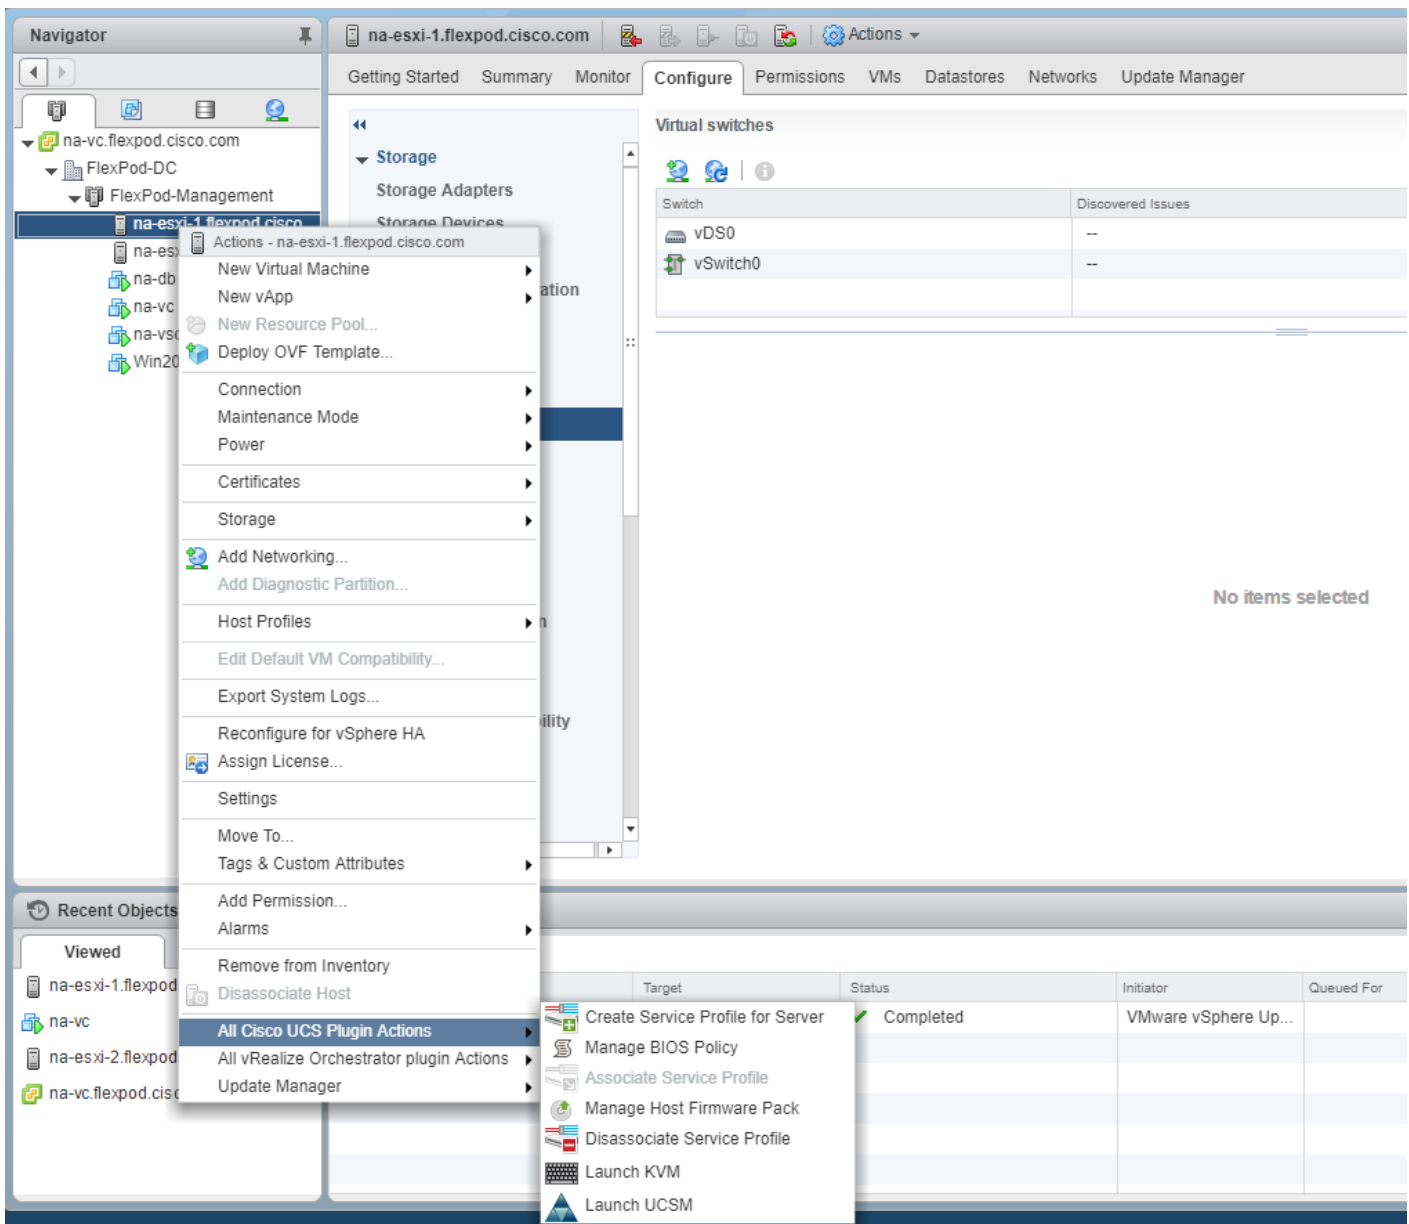
When adding new ESXi hosts to vCenter, it will be necessary to disable Proactive HA in order to add the host. Once the host is added, simply re-enable Proactive HA.

Use the Cisco UCS vCenter Plugin

The plugin can now enable the functions described at the start of this section by double-clicking the registered UCS Domain under the Home tab in the Cisco UCS Management Center:



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



For detailed installation instructions and usage information, please refer to the Cisco UCS Manager Plug-in for VMware vSphere Web Client User Guide:

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.

NetApp Virtual Storage Console 9.6 Deployment Procedure

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

Virtual Storage Console 9.6 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.6:

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



The Backup and Recovery capability has been integrated with SnapCenter and requires additional licenses for SnapCenter to perform backup and recovery of virtual machines and applications.

The requirements for deploying VSC are listed [here](#).

Install Virtual Storage Console 9.6

To install the VSC 9.6 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

1. Go to vSphere HTML5 Client > Host Cluster > Deploy OVF Template.
2. Browse to the VSC OVA file downloaded from the NetApp Support site.
3. Enter the VM name and select a datacenter or folder in which to deploy. Click NEXT.
4. Select a host cluster resource in which to deploy OVA. Click NEXT.
5. Review the details and accept the license agreement.
6. Select Storage (infra_datastore – Thin Provision).
7. From Select Networks, choose a destination network (IB-MGMT – vDS0) and click NEXT.
8. From Customize Template, enter the VSC administrator password, vCenter name or IP address and other configuration details and click NEXT.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- ✓ 8 Customize template
- 9 Ready to complete**

Click Finish to start creation.

Provisioning type	Deploy from template
Name	NetApp-VSC
Template name	unified-virtual-appliance-for-vsc-vp-sra-9.6-4209-20190812_0930
Download size	1.0 GB
Size on disk	2.1 GB
Folder	FlexPod-DC
Resource	FlexPod-Management
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thin provision
Network mapping	1
nat	IB-MGMT
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual

CANCEL

BACK

FINISH

9. Review the configuration details entered and click FINISH to complete the deployment of NetApp-VSC VM.
10. Power on the NetApp-VSC VM and open the VM console.
11. During the NetApp-VSC VM boot process, you see a prompt to install VMware Tools. From vCenter, select NetApp-VSC VM > Guest OS > Install VMware Tools.

```
Booting VSC, VASA Provider, and SRA virtual appliance...Please wait...  
VMware Tools OVF vCenter configuration not found.  
VMware Tools OVF vCenter configuration not found.  
VMware Tools OVF vCenter configuration not found.
```

VMware Tools installation

Before you can continue the VSC, VASA Provider, and SRA virtual appliance installation, you must install the VMware Tools:

1. Select VM > Guest OS > Install VMware Tools.

OR

Click on "Install VMware Tools" pop-up box on the vSphere Web Client.

2. Follow the prompts provided by the VMware Tools wizard.

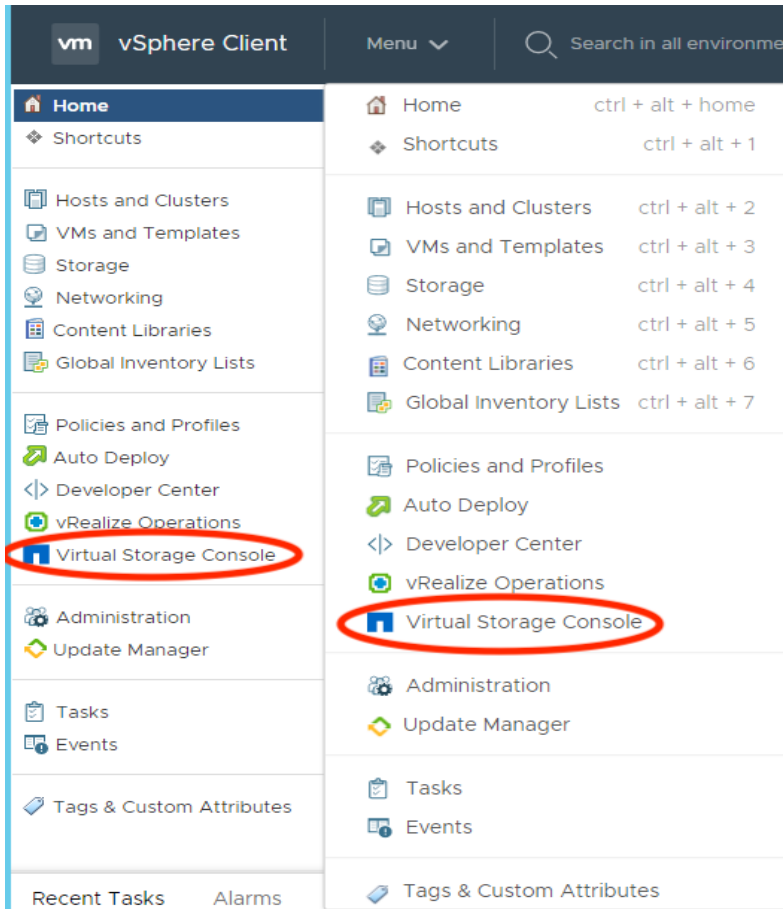
Once you click on mount, the installation process will automatically continue.

12. Networking configuration and vCenter registration information was provided during the OVF template customization, therefore after NetApp-VSC VM is running, VSC, vSphere API for Storage Awareness (VASA), and VMware Storage Replication Adapter (SRA) are registered with vCenter.

13. From the Home screen, confirm that the NetApp VSC is installed.



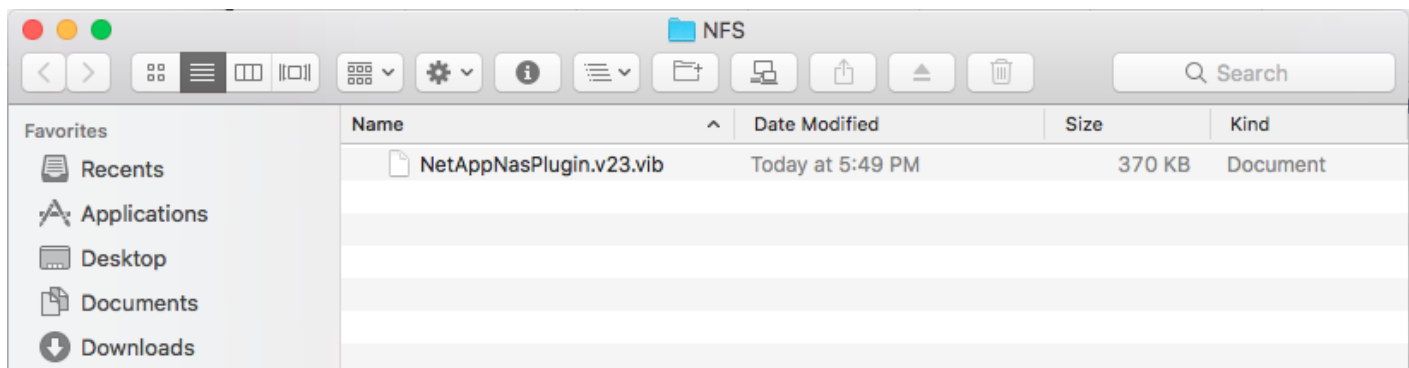
The NetApp VSC 9.6 vCenter Plugin is only available in the vSphere HTML5 Client and is not available in the vSphere Web Client.



Download the NetApp NFS VAAI Plug-In

To download the NetApp NFS VAAI Plug-In, follow this step:

1. Download the NetApp NFS Plug-In 1.1.2 for VMware .vib file from the [NFS Plugin Download](#) page and save it to your local machine or admin host.



Install the NetApp NFS VAAI Plug-In



The NFS VAAI plug-in was already installed on the ESXi hosts along with the Cisco UCS VIC drivers. It is not necessary to re-install it here.

To install the NetApp NFS VAAI plug-in, follow these steps:

1. Rename the .vib file that you downloaded from the NetApp Support Site to NetAppNasPlugin.vib to match the predefined name that VSC uses.
2. Click Settings in the VSC home page.
3. Click NFS VAAI Tools tab.
4. Click Change in the Existing version section.
5. Browse and select the renamed .vib file, and then click Upload to upload the file to the virtual appliance.
6. In the Install on ESXi Hosts section, select the ESXi host on which you want to install the NFS VAAI plug-in, and then click Install.
7. Reboot the ESXi host after the installation finishes.

Enable the VASA Provider

To enable the VASA provider, follow these steps:

1. From the vSphere Client, click Menu > Virtual Storage Console.
2. Click Settings.
3. Click Manage Capabilities in the Administrative Settings tab.
4. In the Manage Capabilities dialog box, select the VASA Provider extension to enable.
5. Enter the IP address of the virtual appliance for VSC, VASA Provider, and SRA and the administrator password, and then click Apply.

Manage Capabilities ×

Enable VASA Provider

vStorage APIs for Storage Awareness (VASA) is a set of application program interface (APIs) that enables vSphere vCenter to recognize the capabilities of storage arrays.

Enable Storage Replication Adapter (SRA)

Storage Replication Adapter (SRA) allows VMware Site Recovery Manager (SRM) to integrate with third party storage array technology.

Enter authentication details for VASA provider and SRA server:

IP address or hostname: 10.1.156.201

Username: Administrator

Password: *

CANCEL
APPLY

Discover and Add Storage Resources

To Add storage resources for the Monitoring and Host Configuration capability and the Provisioning and Cloning capability, follow these 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.



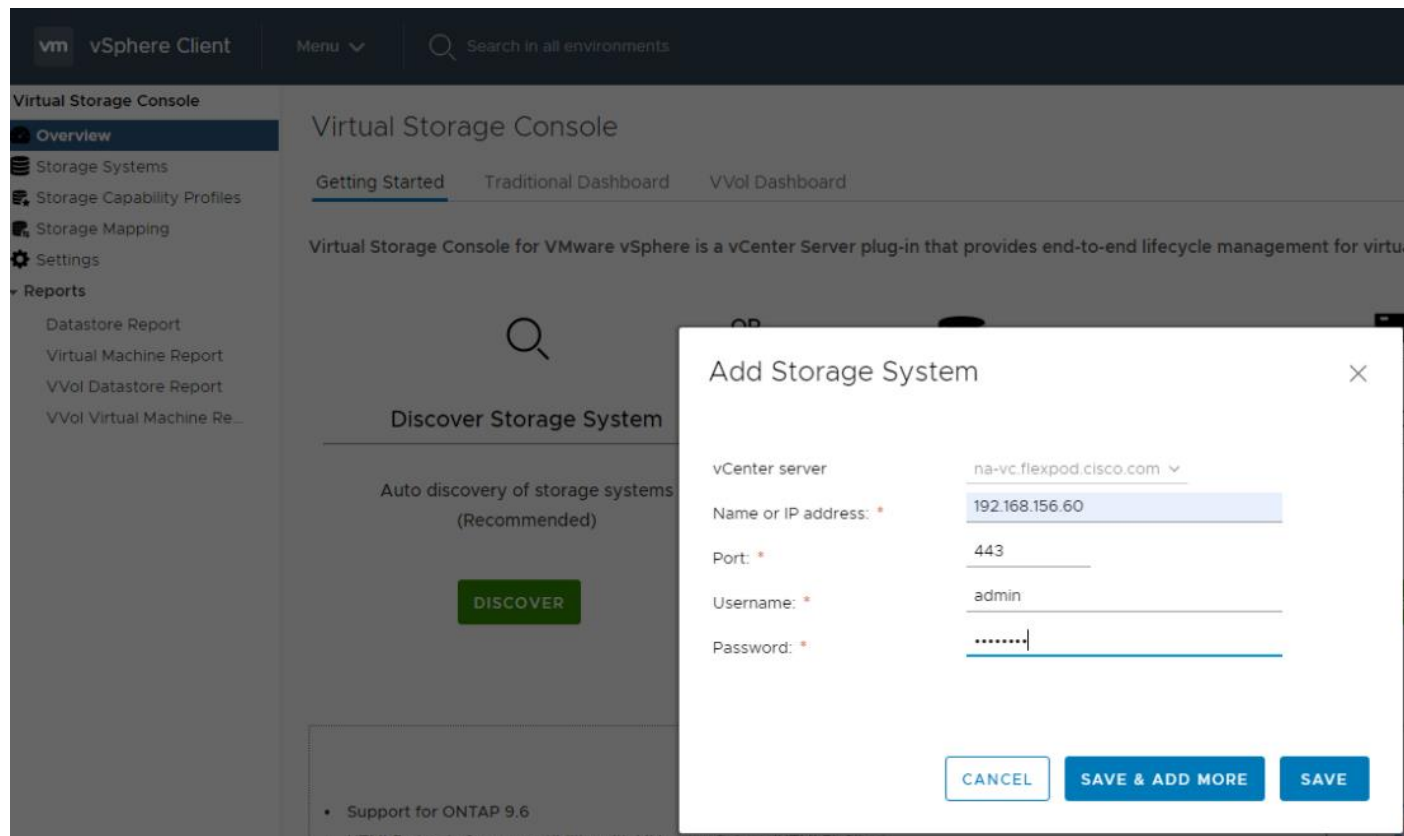
You can modify the storage credentials with the vsadmin account or another SVM level account with RBAC privileges.



When using the cluster admin account, add storage from the cluster level.

3. Select Storage Systems >Add
4. Click Overview > Getting Started, and then click ADD button under Add Storage System.
5. Specify the vCenter Server instance where the storage will be located.
6. In the IP Address/Hostname field, enter the storage cluster management IP.
7. Confirm Port 443 to Connect to this storage system.

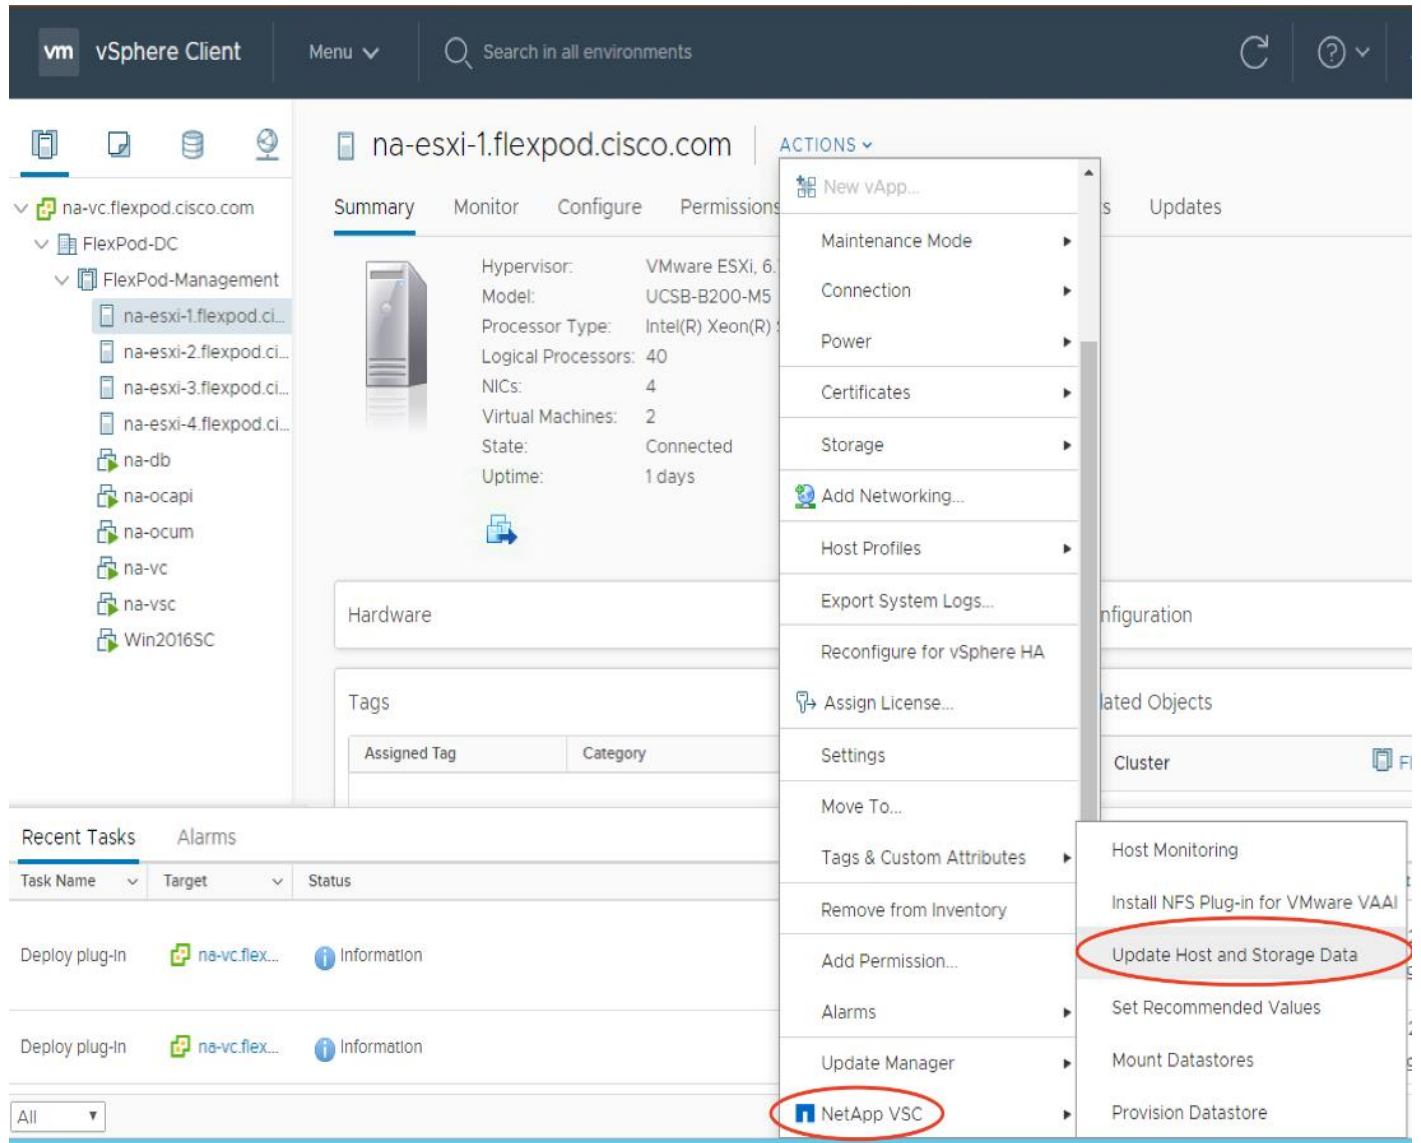
8. Enter admin for the user name and the admin password for password.
9. Click Save to add the storage configuration to VSC.



10. Wait for the Storage Systems to update. You might need to click Refresh to complete this update.

To Discover the cluster and SVMs with the cluster admin account, follow these steps:

1. From the vSphere Client Home page, select Hosts and Clusters.
2. From the Actions menu of required datacenter, select NetApp VSC > Update Host and Storage Data.

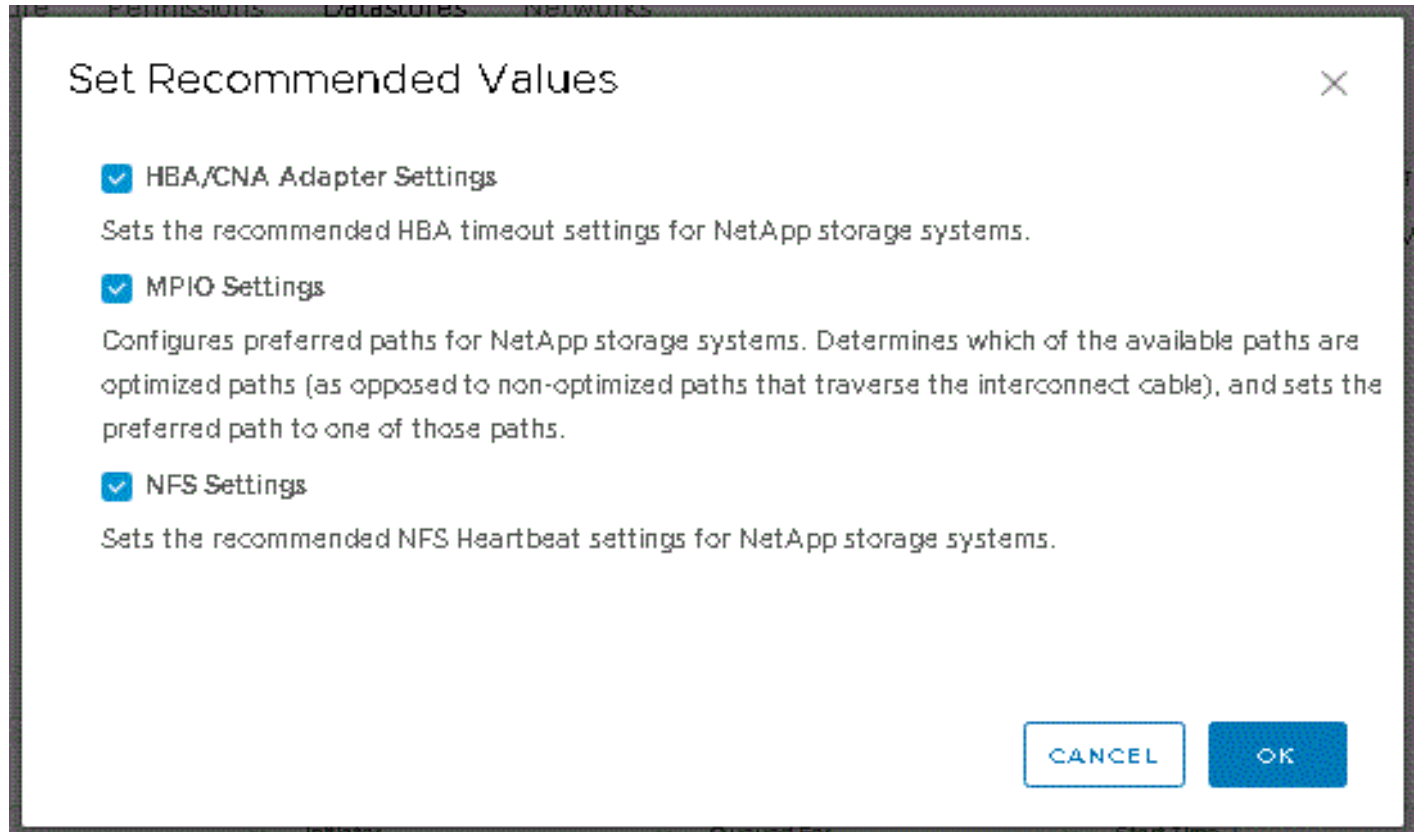


3. VSC displays a Confirm dialog box that informs you that this operation might take a few minutes.
4. Click OK.

Optimal Storage Settings for ESXi Hosts


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

1. From the VMware vSphere Web Client Home page, click vCenter > Hosts.
2. Select a host, and then select Actions > NetApp VSC > Set Recommended Values.
3. In the NetApp Recommended Settings dialog box, select all the values for your system.



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).

4. Click OK.



Virtual Storage Console 9.6 Provisioning Datastores

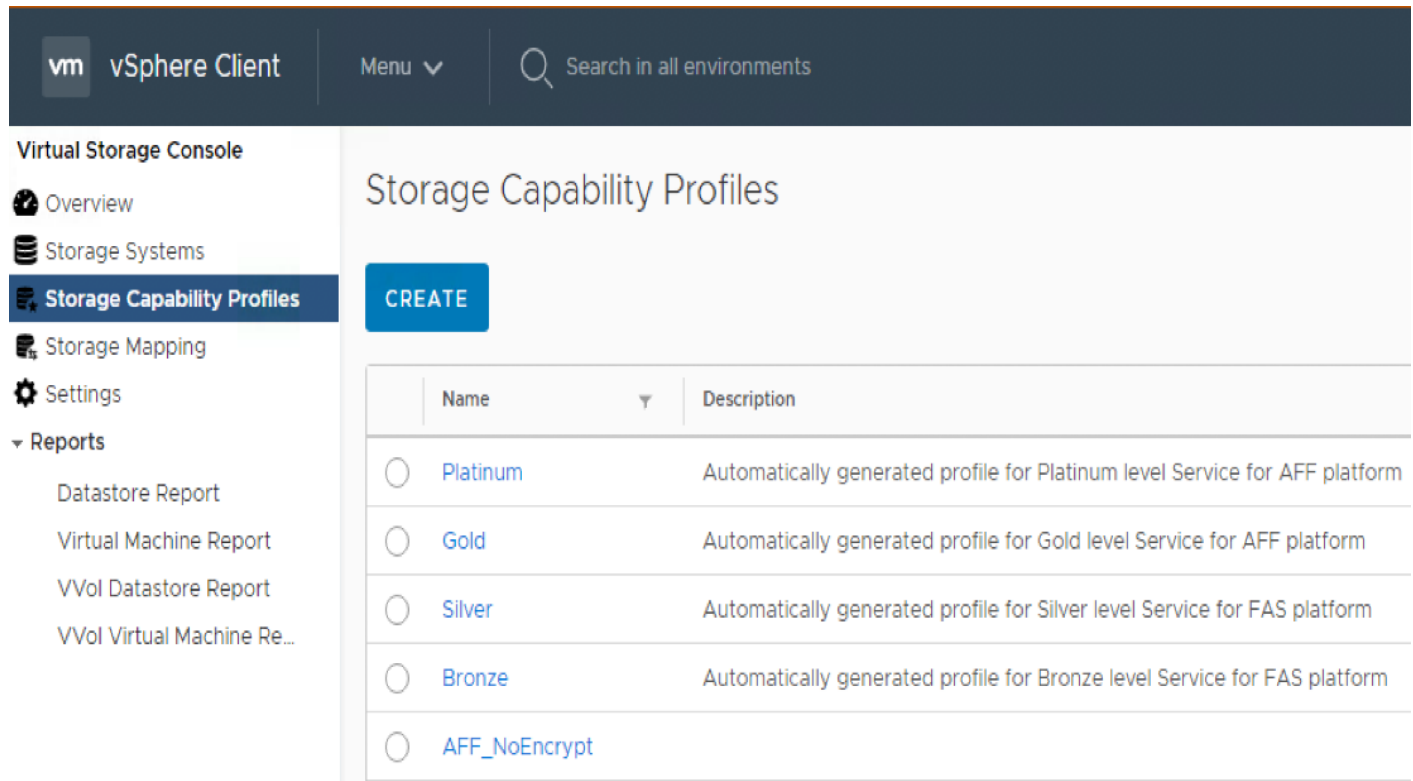
Using VSC, the administrator can provision an NFS, FC or iSCSI datastore and attach it to a single host or multiple hosts in the cluster. The following steps describe provisioning a datastore and attaching it to the cluster.



It is a NetApp best practice to use Virtual Storage Console (VSC) to provision datastores for the FlexPod infrastructure. When using VSC to create vSphere datastores, all NetApp storage best practices are implemented during volume creation and no additional configuration is needed to optimize performance of the datastore volumes.

Storage Capabilities

A storage capability is a set of storage system attributes that identifies a specific level of storage performance (storage service level), storage efficiency, and other capabilities such as encryption for the storage object that is associated with the storage capability. Use the default storage capabilities.



The screenshot shows the vSphere Client interface. The top navigation bar includes the 'vm vSphere Client' logo, a 'Menu' dropdown, and a search bar labeled 'Search in all environments'. The left sidebar is titled 'Virtual Storage Console' and contains the following menu items: Overview, Storage Systems, Storage Capability Profiles (highlighted), Storage Mapping, Settings, and Reports. Under Reports, there are sub-items: Datastore Report, Virtual Machine Report, VVol Datastore Report, and VVol Virtual Machine Re... The main content area is titled 'Storage Capability Profiles' and features a blue 'CREATE' button. Below the button is a table with the following data:

	Name	Description
<input type="radio"/>	Platinum	Automatically generated profile for Platinum level Service for AFF platform
<input type="radio"/>	Gold	Automatically generated profile for Gold level Service for AFF platform
<input type="radio"/>	Silver	Automatically generated profile for Silver level Service for FAS platform
<input type="radio"/>	Bronze	Automatically generated profile for Bronze level Service for FAS platform
<input type="radio"/>	AFF_NoEncrypt	

Provision NFS Datastore

To provision the NFS datastore, follow these steps:

1. From the Virtual Storage Console Home page, click Overview
2. In the Getting Started Tab, Click Provision Button.
3. Click Browse to select the destination to provision the datastore as per the next step.
4. Select the type as NFS and Enter the datastore name.
5. Provide the size of the datastore and the NFS Protocol.
6. Check the storage capability profile and click Next

New datastore

✕

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

General

Specify the details of the datastore to provision. ?

Provisioning destination: * BROWSE

Type: * NFS VMFS VVol

Name: *

Size: * GB v

Protocol: * NFS 3 NFS 4.1

Use storage capability profile for provisioning

CANCEL
NEXT

7. Select the Storage Capability Profile, cluster name and the desired SVM to create the datastore. In this example, the Infra-SVM is selected.
8. Click Next.

New datastore

✕

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage system

Specify the storage capability profiles and the storage system you want to use.

Storage capability profile: * v

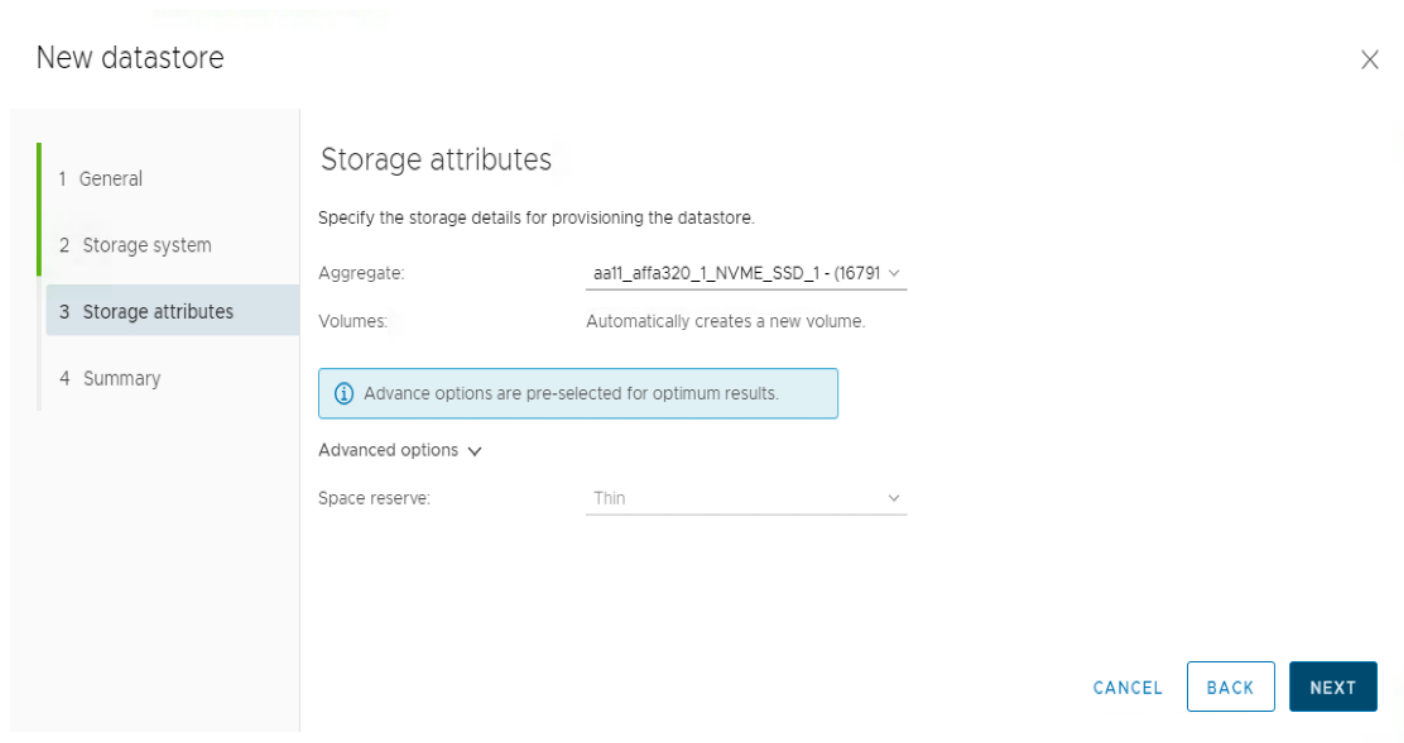
Platform: All Flash FAS(AFF)	Performance: None	Encryption: No
Compression: Yes	Deduplication: Yes	Encryption: No
Space reserve: Thin	Encryption: No	Tiering policy (FabricPool): Any

Storage system: * v

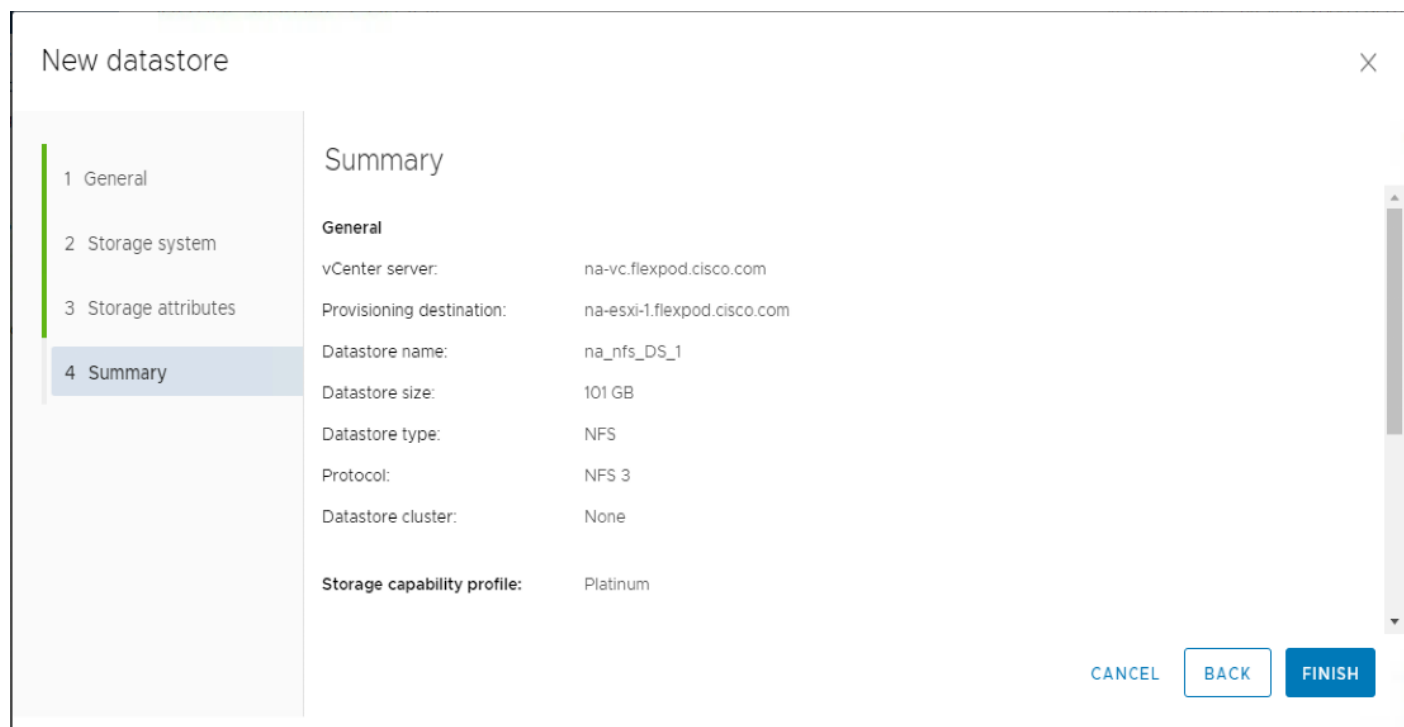
Storage VM: * v

CANCEL
BACK
NEXT

9. Select the aggregate name and Click Next.



10. Review the Summary and click Finish.





The datastore is created and mounted on the hosts in the cluster. Click Refresh from the vSphere web client to see the newly created datastore or it is also listed in the VSC home page > Traditional Dashboard > Datastores view.

Provision FC Datastore

To provision the FC datastore, follow these steps:

1. From the Virtual Storage Console Home page, click Overview
2. In the Getting Started Tab, Click Provision Button.
3. Click Browse to select the destination to provision the datastore as per the next step.
4. Select the type as VMFS and Enter the datastore name.
5. Provide the size of the datastore and the FC Protocol.
6. Check the storage capability profile and click Next

New datastore X

1 General

2 Storage system

3 Storage attributes

4 Summary

General

Specify the details of the datastore to provision. ?

Provisioning destination: * BROWSE

Type: * NFS VMFS VVol

Name: *

Size: *

Protocol: * iSCSI FC / FCoE

Use storage capability profile for provisioning

CANCEL
NEXT

7. Select the Storage Capability Profile, cluster name and the desired SVM to create the datastore. In this example, the Infra-SVM is selected.
8. Click Next.

New datastore ✕

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage system

Specify the storage capability profiles and the storage system you want to use.

Storage capability profile: *	Platinum	▼
Platform: All Flash FAS(AFF)	Performance: None	
Compression: Yes	Deduplication: Yes	Encryption: No
Space reserve: Thin	Encryption: No	Tiering policy (FabricPool): Any
Storage system: *	aa11-affa320 (192.168.156.60)	▼
Storage VM: *	Infra-SVM	▼

CANCEL BACK NEXT

9. Select the aggregate name and click Next.

New datastore ✕

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Storage attributes

Specify the storage details for provisioning the datastore.

Aggregate: aa11_afa320_1_NVME_SSD_1 - (16791 ▼

Volumes: Automatically creates a new volume.

[Advanced options >](#)

CANCEL BACK NEXT

10. Review the Summary and click Finish.

New datastore X

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

Summary

General

vCenter server: na-vc.flexpod.cisco.com

Provisioning destination: na-esxi-2.flexpod.cisco.com

Datastore name: na_FC_DS_1

Datastore size: 100 GB

Datastore type: VMFS

Protocol: FCP

File system: VMFS6

Datastore cluster: None

Storage capability profile: Platinum

CANCEL
BACK
FINISH

11. Click Ok.



The datastore is created and mounted on all the hosts in the cluster. Click Refresh from the vSphere web client to see the newly created datastore or it is also listed in the VSC home page > Traditional Dashboard > Datastores view.

NetApp SnapCenter 4.2

SnapCenter Software is a simple, centralized, scalable platform that provides application-consistent data protection for applications, databases, host file systems, and VMs running on ONTAP systems anywhere in the Hybrid Cloud.

Deploy NetApp SnapCenter

The SnapCenter platform is based on a multitier architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter host agent. Snap Center Plug-in for VMware vSphere is packaged as a Linux Appliance (Debian-based Open Virtual Appliance format) and will no longer be part of the SnapCenter Plug-ins Package for Windows. The new virtual appliance is called NetApp Data Broker.

You need to install SnapCenter Server and the necessary plug-ins to support application-consistent backups. The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has more information.

- [Deployment Guide for SnapCenter Plug-in for VMware vSphere](#)
- [Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere](#)

SnapCenter Server Requirements

[Table 8](#) lists the minimum requirements for installing SnapCenter Server and the SnapCenter plug-in on a Windows server. For the latest version compatibility and other plug-in information, refer to the [NetApp Interoperability Matrix Tool](#).

See the [SnapCenter 4.2 Installation and Setup Guide](#) for the complete documentation.

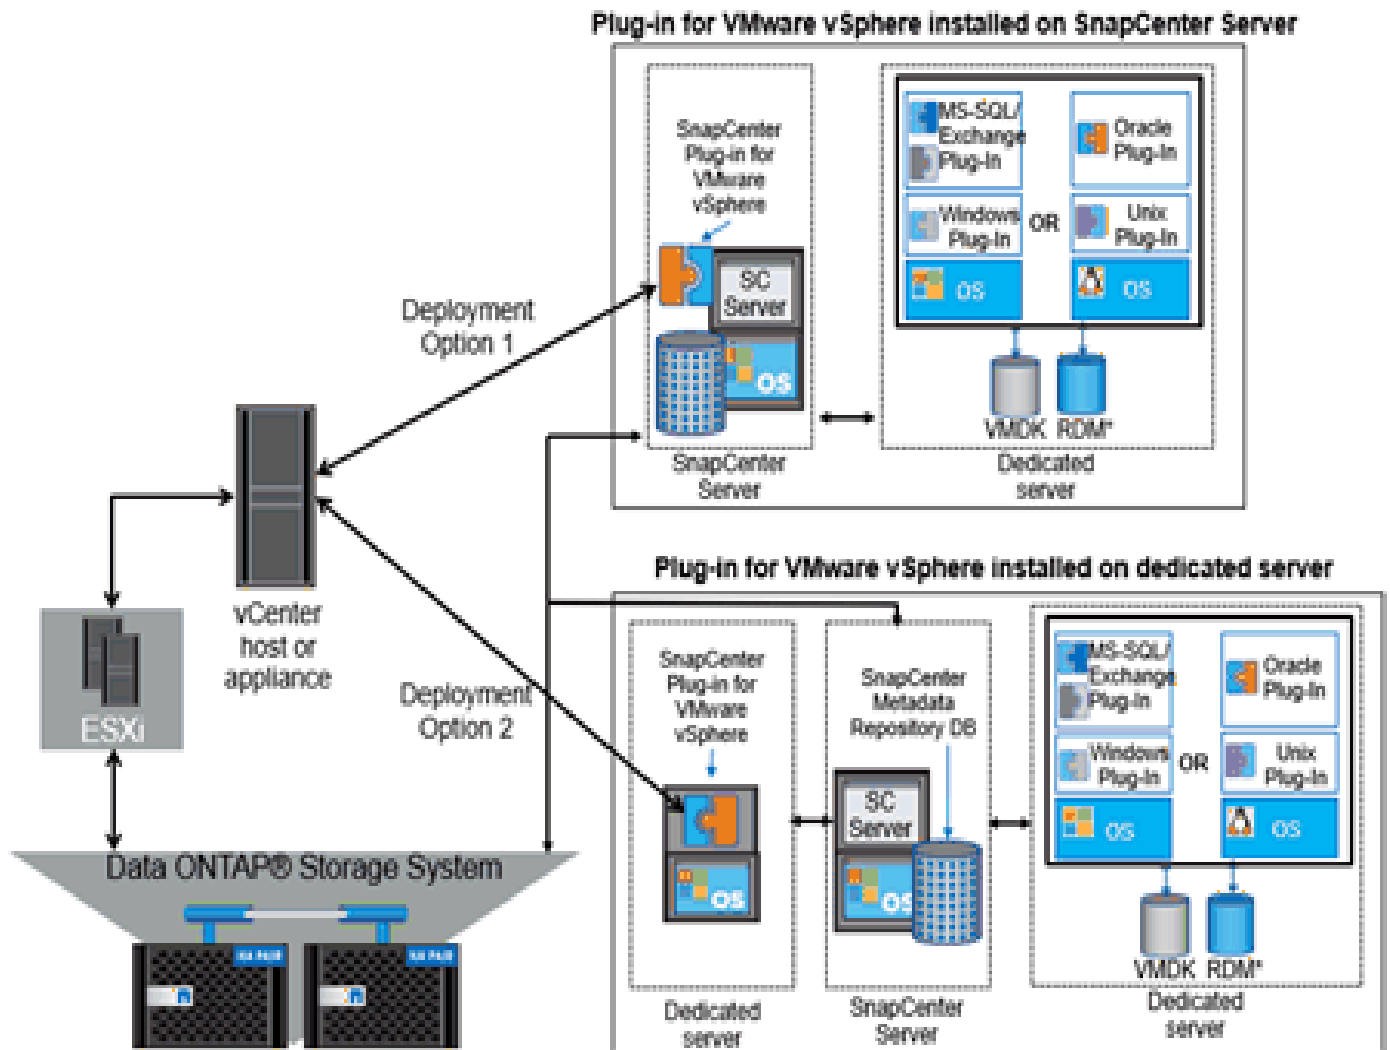
Table 8 SnapCenter Server Requirements

Component	Requirements
Minimum CPU count	4 cores/vCPUs
Memory	Minimum: 8GB Recommended: 32GB
Storage space	Minimum space for installation: 10GB Minimum space for repository: 20GB
Supported operating systems	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016
Software packages	.NET 4.5.2 or later Windows Management Framework 4.0 or later PowerShell 4.0 or later Java 1.8 (64-bit)
Active Directory domain membership	Windows Host must be joined to an Active Directory domain.
Database for SnapCenter repository	MySQL Server 5.7.22 (installed as part of the SnapCenter installation)

Table 9 Port Requirements

Port	Requirement
443	vCenter Server to SnapCenter Server API Access over HTTPS
8144	SnapCenter GUI to SnapCenter Plug-in for VMWare
8145	SnapCenter Server Core API
8146	For SnapCenter server REST API
3306	MySQL

Figure 3 NetApp Snapshot Deployment



SnapCenter 4.2 License Requirements

The following licenses are required for SnapCenter on storage systems that run ONTAP 9.6. These licenses do not need to be added into the SnapCenter GUI:

- SnapCenter Standard controller-based licenses
- Protocol licenses (NFS, FCP, or iSCSI as per protocol selection for deployment)
- NetApp FlexClone (for provisioning and cloning only)
- NetApp SnapRestore (for backup and recovery)
- The NetApp SnapManager Suite

SnapCenter Server

The SnapCenter Server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, APIs, and the SnapCenter repository. SnapCenter enables load balancing, high availability, and horizontal scaling across multiple SnapCenter Servers within a single user interface. You can accomplish high availability by using network load balancing (NLB) and Application Request Routing (ARR) with SnapCenter. For larger environments with thousands of hosts, adding multiple SnapCenter Servers can help balance the load.

SnapCenter Plug-In for VMware vSphere

For SnapCenter 4.2 and later, the SnapCenter Plug-in for VMware vSphere is deployed as part of the NetApp Data Broker virtual appliance Linux-based VM. It provides a vSphere web client GUI on vCenter to protect VMware virtual machines (VMs) and datastores and supports SnapCenter application-specific plug-ins in protecting virtualized databases on primary and secondary storage.

Support for Virtualized Databases and File Systems

The Plug-in for VMware vSphere provides native backup, recovery, and cloning of virtualized applications (virtualized SQL and Oracle databases and Windows file systems) when using the SnapCenter GUI.

SnapCenter natively leverages the Plug-in for VMware vSphere for all SQL, Oracle, and Windows file system data protection operations on VM disks (VMDKs), raw device mappings (RDMs), and NFS datastores.

SnapCenter Installation

To install SnapCenter, follow these steps. However, before you install SnapCenter, the SnapCenter Server host system must be up-to-date with Windows updates and no pending system restarts.

1. Download the SnapCenter Server installation package from NetApp Support Site.

[NetApp Downloads: SnapCenter Software](#)

2. Initiate the SnapCenter Server installation by double-clicking the downloaded .exe file.

After you initiate the installation, all the prechecks are performed and if the minimum requirements are not met appropriate error or warning messages are displayed.

You can ignore the warning messages and proceed with installation; however, errors should be fixed.

3. Review the pre-populated values required for the SnapCenter Server installation and modify if required.

You do not have to specify the password for MySQL Server repository database. During SnapCenter Server installation the password is auto generated.

4. Click Install Now.

SnapCenter Configuration

The default GUI URL is a secure connection to the default port 8146 on the server where the SnapCenter Server is installed (<https://server:8146>). If you provided a different server port during the SnapCenter installation, that port is used instead.

For NLB deployment, you must access SnapCenter using the NLB cluster IP (https://NLB_Cluster_IP:8146). If you do not see the SnapCenter UI when you navigate to https://NLB_Cluster_IP:8146 in IE, you must add the NLB IP

address as a trusted site in IE on each plug-in host, or you must disable IE Enhanced Security on each plug-in host.

Add a User or Group to a Role

To add a user or group to a role, follow these steps:

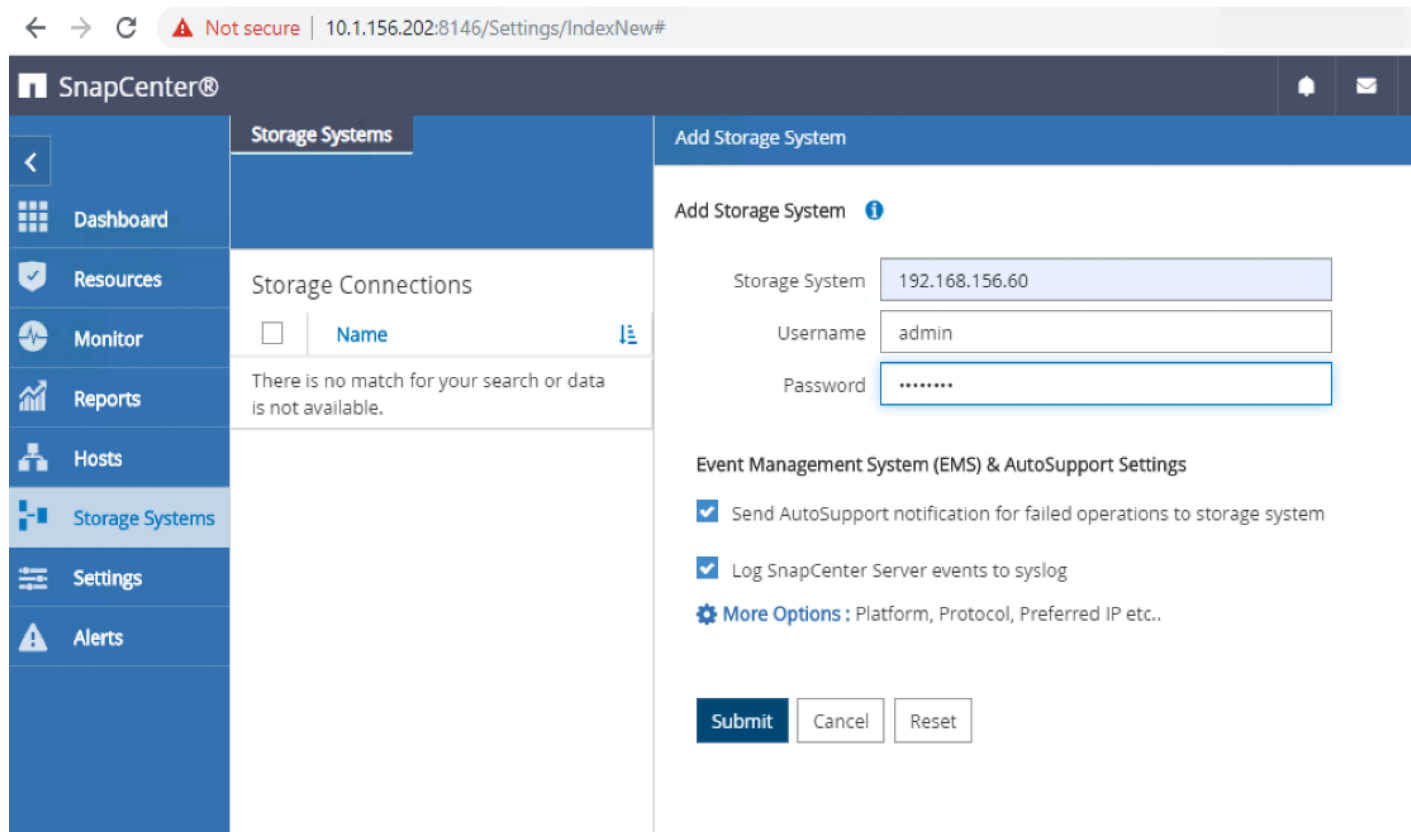
1. In the left navigation pane, click Settings.
2. In the Settings page, click <Users and Access >.
3. Select Domain or Workgroup as the Access type. For Domain authentication type, you should specify the domain name of the user or group to which you want to add the user to a role.
4. Select User/Group for Type and add the User Name. Repeat this step to add additional users or groups to the selected role.
5. Select the role to which add the user.
6. Click Assign and add assets to the user.
7. Click Submit.

Configure Storage System Connections

To perform data protection and provisioning operations with SnapCenter, you must first set up the storage system connections that give SnapCenter access to ONTAP storage. Storage systems can also be added thru SnapCenter Plugin for vCenter. Both methods require the same set of information.

To add storage systems in SnapCenter GUI, follow these steps:

1. Select Storage Systems -> Add
2. Provide Storage system IP/FQDN and credentials.
3. Check all event management options.
4. Click Submit.



Install SnapCenter Plug-In for VMware

NetApp Data Broker is a linux-based virtual appliance (Open Virtual Appliance format) which enables the SnapCenter Plug-in for VMware vSphere.

Host and Privilege Requirements for the NetApp Data Broker

Review the following requirements before you install the NetApp Data Broker virtual appliance for SnapCenter Plug-in for VMware vSphere:

- You must deploy the NetApp Data Broker virtual appliance as a Linux VM.
- You should deploy the virtual appliance on the vCenter Server.
- You must not deploy the virtual appliance in a folder that has a name with special characters.
- You must deploy and register a separate, unique instance of the virtual appliance for each vCenter Server.

Table 10 Port Requirements

Port	Requirement
8144(HTTPS)	SnapCenter Plug-in for VMware vSphere port
443 (HTTPS)	VMware vSphere vCenter Server port

Download and Deploy the NetApp Data Broker OVA (Open Virtual Appliance)

To download and deploy the NetApp Data Broker OVA, follow these steps:

1. Download NetApp Data Broker for VMware vSphere .ova file from Netapp support site (<https://mysupport.netapp.com>).
2. From the VMware VCenter, VMs and Templates Tab, In Actions> Deploy OVF Template.
3. Specify the location of the OVF Template and Click Next.
4. On the Select a name and folder page, enter a unique name for the VM or vApp, and select a deployment location, and then click Next.

Deploy OVF Template

✓ 1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name: na-db

Select a location for the virtual machine.

- na-vc.flexpod.cisco.com
 - FlexPod-DC

CANCEL BACK NEXT

5. On the Select a compute resource page, select a resource where you want to run the deployed VM template, and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- 3 Select a compute resource**
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select a compute resource

Select the destination compute resource for this operation

- ✓ FlexPod-DC
 - ✓ FlexPod-Management
 - na-esxi-1.flexpod.cisco.com
 - na-esxi-2.flexpod.cisco.com

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

6. On the Review details page, verify the .ova template details and click Next.
7. On the License agreements page, select the checkbox for I accept all license agreements.
8. On the Select storage page, define where and how to store the files for the deployed OVF template:
 - a. Select the disk format for the VMDKs.
 - b. Select a VM Storage Policy.
 - c. Select a datastore to store the deployed OVA template.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- 6 Select storage**
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed ▾

VM Storage Policy: Datastore Default ▾

Name	Capacity	Provisioned	Free	Type
datastore1	7.5 GB	4.37 GB	3.13 GB	VM
infra_datastore	500 GB	405.96 GB	393.8 GB	NF
infra_nfs_datastore1	10 GB	424 KB	10 GB	NF
infra_nfs_datastore2	20 GB	4.56 MB	20 GB	NF
infra_swap	100 GB	94.58 MB	99.91 GB	NF

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

9. On the Select networks page, select a source network and map it to a destination network, and then click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- 7 Select networks**
- 8 Customize template
- 9 Ready to complete

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	IB-MGMT Network

1 items

IP Allocation Settings

IP allocation: Static - Manual
IP protocol: IPv4

CANCEL BACK **NEXT**

10. On the Customize template page, do the following:
 - a. In Register to existing vCenter, enter the vCenter credentials.
 - b. In Create NetApp Data Broker credentials, enter the NetApp Data Broker credentials.
 - c. In Setup Network Properties, enter the network information.
 - d. In Setup Date and Time, select the time zone where the vCenter is located.
11. On the Ready to complete page, review the page and click Finish.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- ✓ 8 Customize template
- 9 Ready to complete**

Ready to complete
Click Finish to start creation.

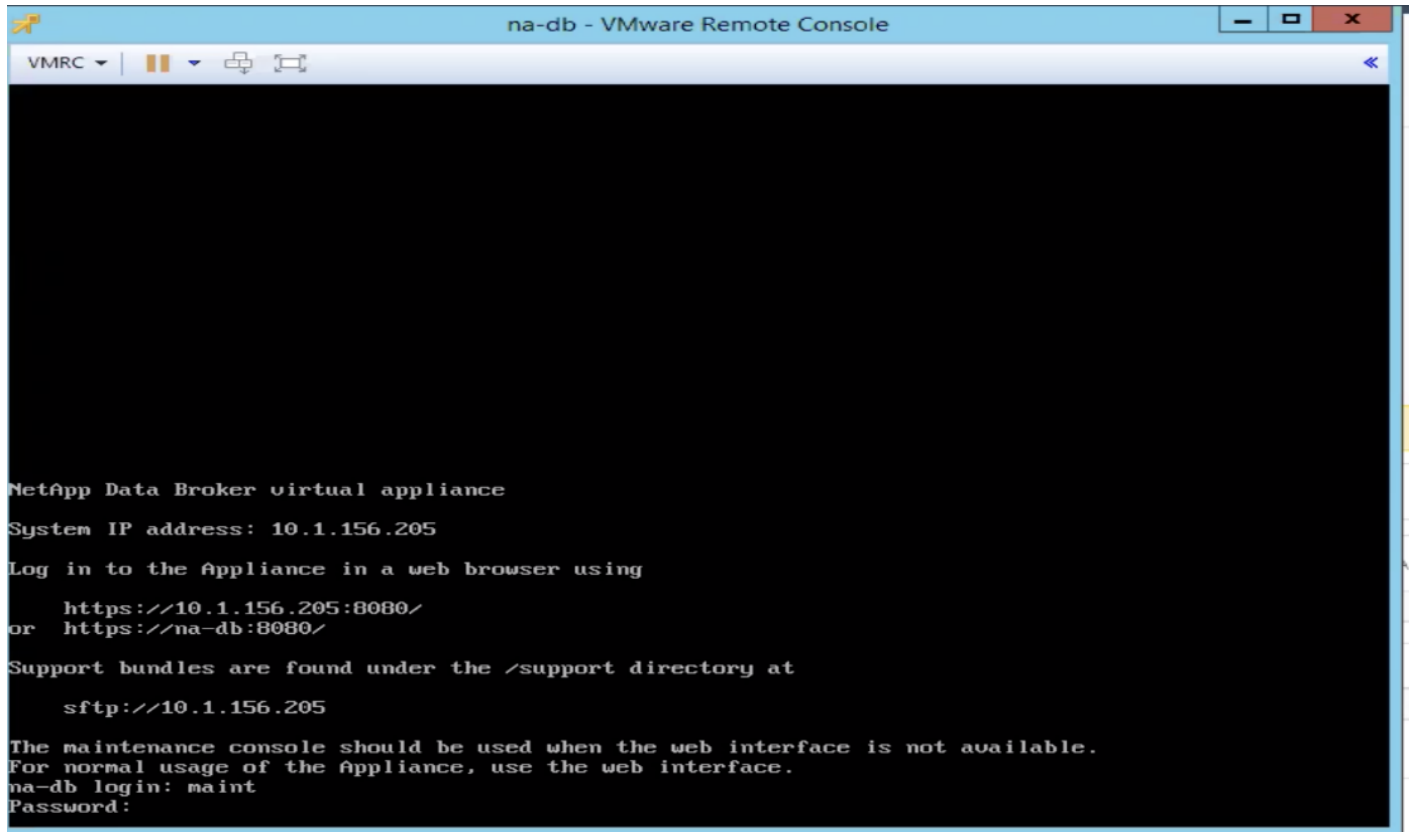
Provisioning type	Deploy from template
Name	na-db
Template name	netapp-data-broker-1.0.0.5368440-190730_0042
Download size	3.5 GB
Size on disk	88.0 GB
Folder	FlexPod-DC
Resource	na-esxi-1.flexpod.cisco.com
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thick provision lazy zeroed
Network mapping	1
nat	IB-MGMT Network
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual

CANCEL

BACK

FINISH

12. Navigate to the VM where the virtual appliance was deployed, then click the Summary tab, and then click the Power On box to start the virtual appliance.
13. While the virtual appliance is powering on, right-click the deployed virtual appliance and then click Install VMware tools.
14. Log in to NetApp Data Broker using the IP address displayed on the deployment screen and the credentials that you provided in the deployment wizard, then verify on the Dashboard that the virtual appliance is successfully connected to vCenter and the SnapCenter Plug-in for VMware vSphere is successfully enabled and connected.
15. The maintenance console user name is set to "maint" and the password is set to "admin123" by default.



Create a Virtual Appliance Credential for Migrating Backups

For VM consistent or VM crash-consistent backups, or application-consistent backups of virtualized data, you must migrate those backups to the virtual appliance for SnapCenter Plug-in for VMware vSphere. Before migrating, you must add the virtual appliance credentials to SnapCenter Server.

To create a virtual appliance credential for migrating backups, follow these steps:

1. In the left navigation pane of the SnapCenter GUI, click Settings.
2. In the Settings page, click Credentials, and then click New to start the wizard.
3. Enter the credential information:

Field Name	Details
Credential name	Enter a name for the credentials.
Username	Enter the username specified when the virtual appliance was deployed.
Password	Enter the password specified when the virtual appliance was deployed.
Authentication	Select Linux.

Register the SnapCenter Plug-in for VMware vSphere with SnapCenter Server

If you want to perform application-over-VMDK workflows in SnapCenter (application-based protection workflows for virtualized databases and file systems), you must register the SnapCenter Plug-in for VMware vSphere virtual appliance with the SnapCenter Server.

To register the SnapCenter plug-in, follow these steps:

1. In the SnapCenter GUI left navigation pane, click Hosts.
2. Verify that the Managed Hosts tab is selected at the top.
3. Click Add to start the wizard.

Field Name	Details
Host Type	Select "vSphere" as the type of host.
Host name	Enter the IP address of the NetApp Data Broker virtual appliance.
Credential	Enter the username and password for the NetApp Data Broker virtual appliance that was provided during the appliance deployment.

4. On the Add Hosts dialog box, specify the host you want to add to the SnapCenter Server:
5. Click Submit.
When the VM host is successfully added, it is displayed on the Managed Hosts tab.
6. Provide the credential information that was created above.

Register the SnapCenter Plug-in for VMware vSphere with VMware Vcenter

To register the SnapCenter plug-in, follow these steps:

1. In your browser, navigate to <http://<Data Broker IP>:8144/register.html>
2. Provide the VCenter Server Information and Click on Register.

- You should see the message Registration successful.

← → ↻ **Not secure** | 10.1.156.205:8144/register.html

vSphere Plugin Registration

The SnapCenter Plugin for VMware is registered as specified below. If you need to change the registration settings update the fields below and then click "Register".

If you specify a new vCenter Server IP address the SnapCenter Plugin will unregister with the previously specified vCenter Server and then register with the newly specified vCenter Server.

vCenter Server information

Host name or IP address:	<input type="text" value="10.1.156.205"/>
Port:	<input type="text" value="443"/>
User name:	<input type="text" value="administrator@vsphere.local"/>
User password:	<input type="password" value="*****"/>

SnapCenter Plug-in for VMware vSphere in VCenter Server

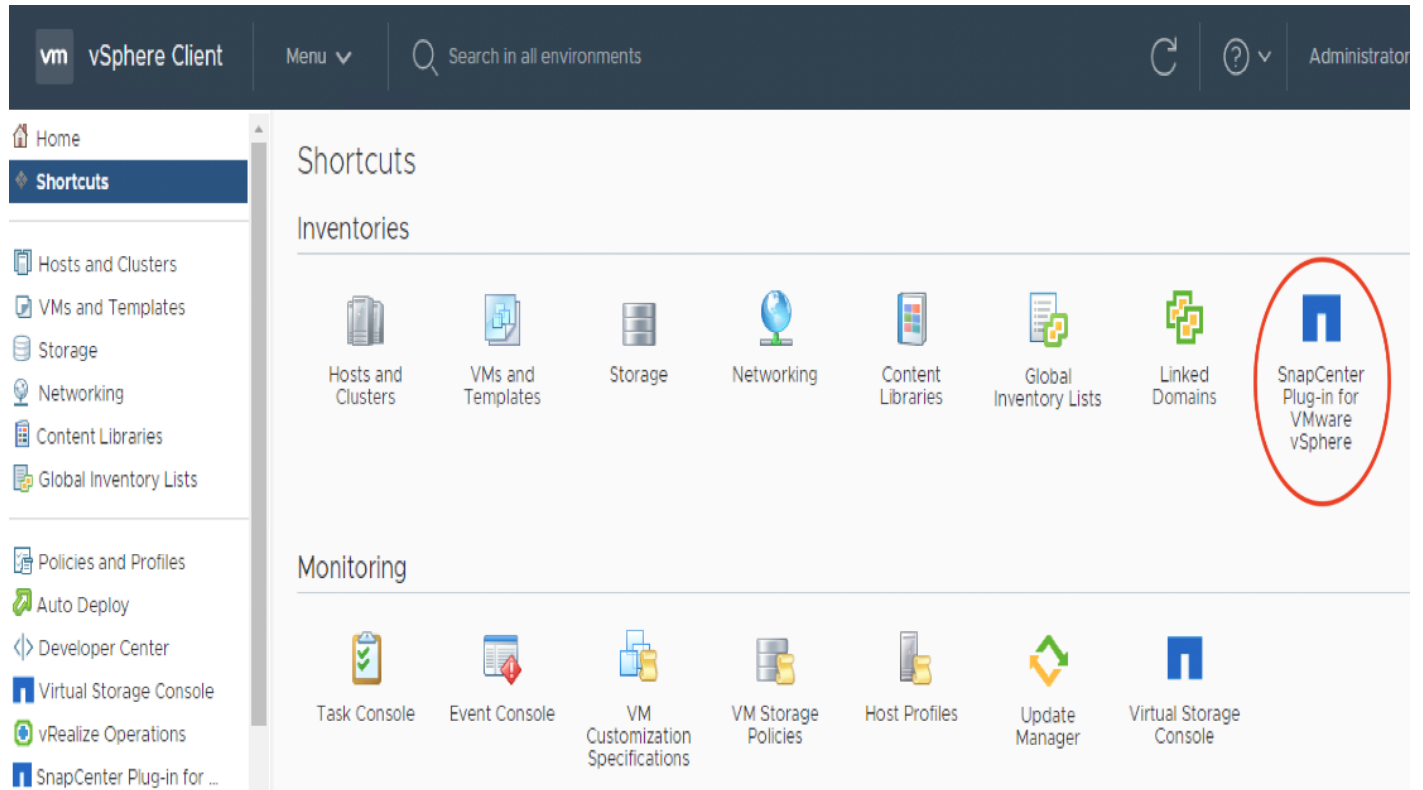
After you have successfully installed the Plug-in for VMware vSphere, to configure SnapCenter and make it ready to backup virtual machines, follow these steps:

- In your browser, navigate to VMware vSphere web client URL <https://<vCenter Server>/vsphere-client/?csp>.



If currently logged into vCenter, logoff and sign-on again to access the SnapCenter Plug-In for VMware vSphere.

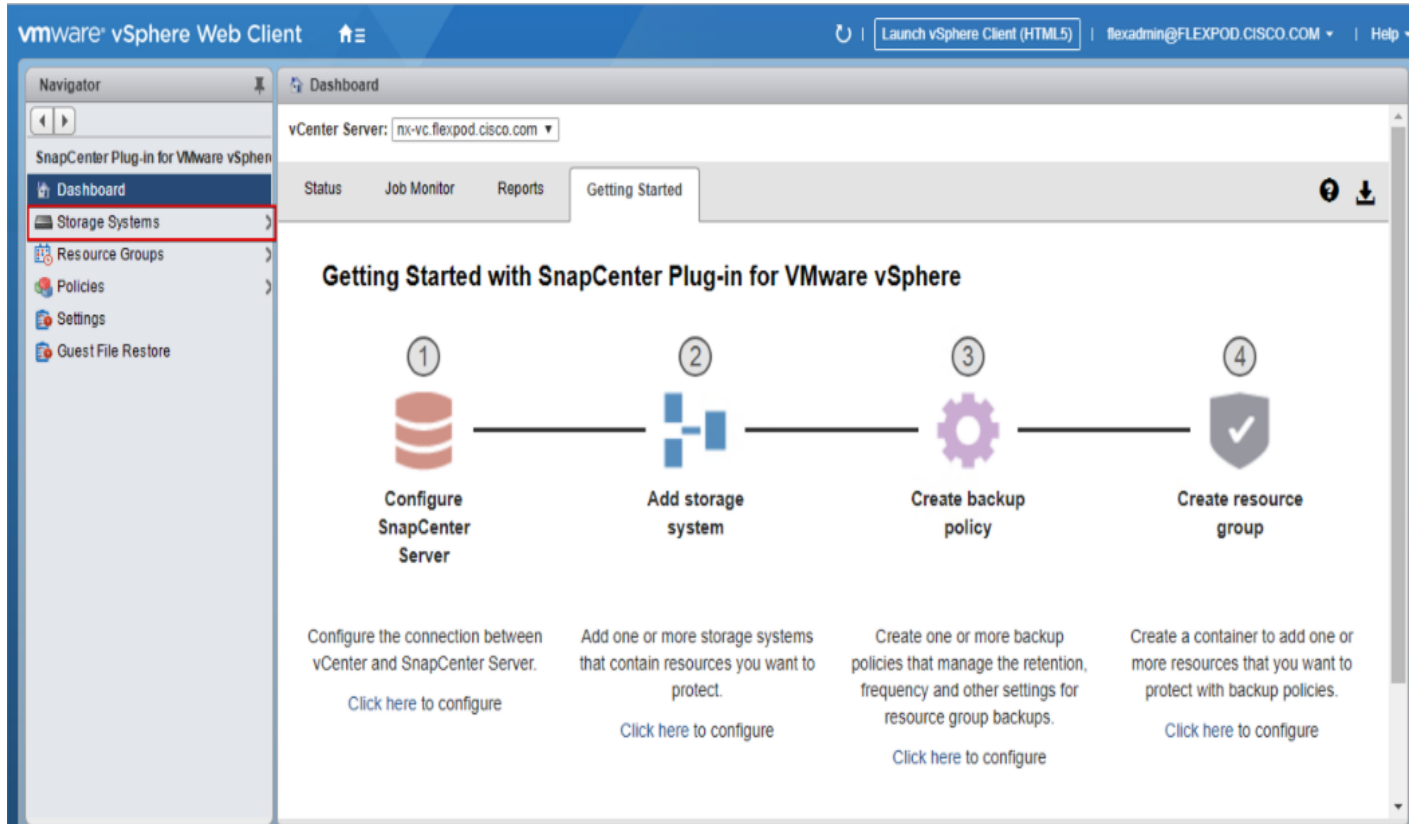
- On the VMware vCenter Single Sign-On page, login.
- On the VMware vSphere web client page, click Home and select SnapCenter Plug-In for VMware vSphere to bring up the SnapCenter Plug-In for VMware GUI. It will take you to the SnapCenter Plug-In for VMware dashboard.



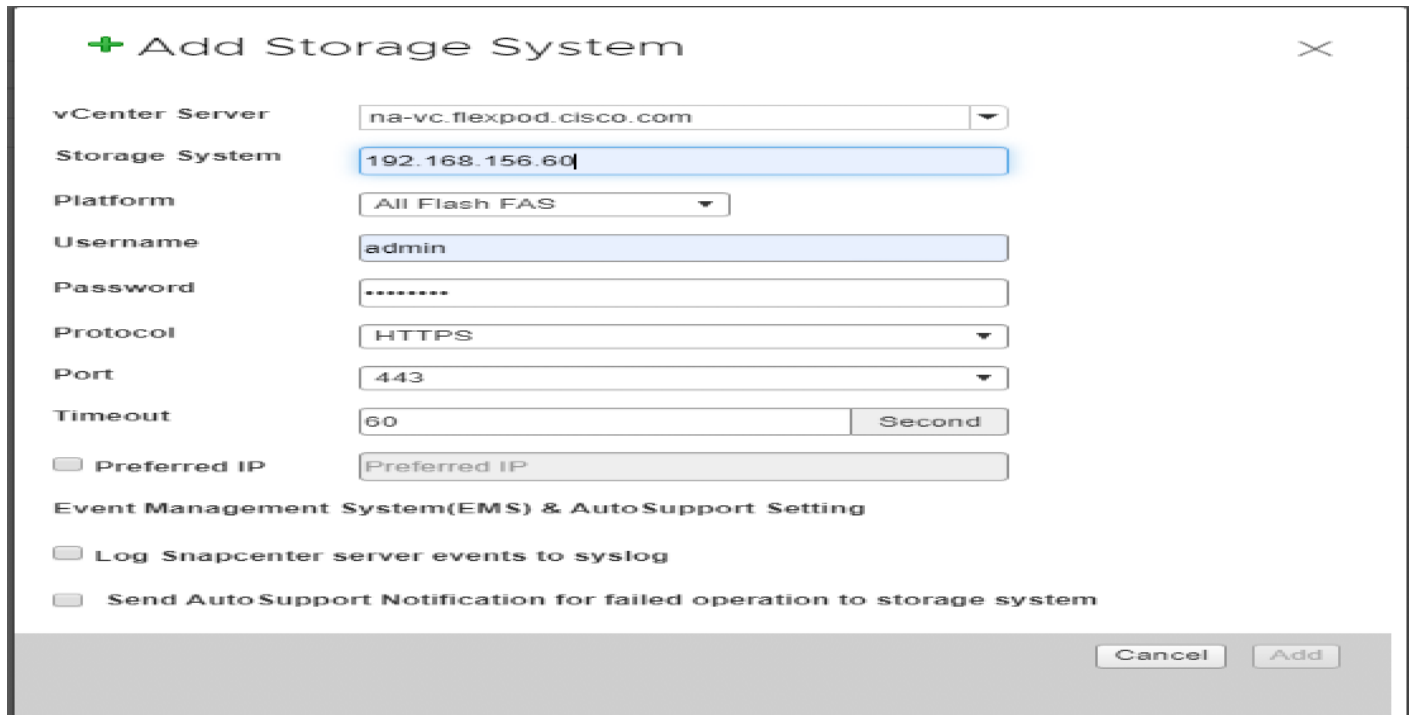
Add Storage Systems (SVM)

To add storage systems, follow these steps:

1. Go to the Storage Systems tab.



2. Click + Add Storage System to add an SVM.
3. Enter vCenter, Storage System, user credentials, and other required information in following dialog box.





You must login to the SVM with the vsadmin or another user with vserver admin privileges not the cluster admin user.

Create Backup Policies for Virtual Machines and Datastores

To create backup policies for VMs and datastores, follow these steps:

1. In the left Navigator pane of the VMware vSphere web client, click Policies.
2. On the Policies page, click + New Policy in the toolbar.
3. On the New Backup Policy page, follow these steps:
 - a. Enter the policy name and a description.
 - b. Enter the backups to keep.
 - c. From the Frequency drop-down list, select the backup frequency (hourly, daily, weekly, monthly, and on-demand only).
 - d. Click Add.

The screenshot shows the 'New Backup Policy' configuration window. The fields are as follows:

- Name:** infra_vm_backups
- Description:** description
- vCenter Server:** na-vc.flexpod.cisco.com
- Retention:** Backup(s) to keep, 7
- Frequency:** Daily
- Replication:**
 - Update SnapMirror after backup
 - Update SnapVault after backup
- Snapshot label:** (empty field)
- Advanced:**
 - VM consistency
 - Include datastores with independent disks
- Scripts:** Enter script path

4. Create multiple policies as required for different sets of VMs or datastores.

Create Resource Groups

Resource groups are groups of virtual machine or datastores that are backed up together. A backup policy is associated with the resource group to back up the virtual machines and retain a certain number of backups as defined in the policy.

To create resource groups, follow these steps:

1. In the left Navigator pane of the VMware vSphere web client, click Resource Groups and then click Create Resource Group. This is the easiest way to create a resource group. However, you can also create a resource group with one resource by performing one of the following steps:
 - a. To create a resource group for one virtual machine, click VMs and Templates, right-click a virtual machine, select NetApp SnapCenter from the drop-down list, and then select Create Resource Group from the secondary drop-down list.
 - b. To create a resource group for one datastore, click Storage, right-click a datastore, select NetApp SnapCenter from the drop-down list, and then select Create Resource Group from the secondary drop-down list.
2. In the General Info & Notification page, enter the resource group name and complete the notification settings. Click Next.

Create Resource Group

✕

1. General info & notification

2. Resource

3. Spanning disks

4. Policies

5. Schedules

6. Summary

vCenter Server:

Name:

Description:

Notification:

Email send from:

Email send to:

Email subject:

Custom snapshot format: Use custom name format for Snapshot copy

Note that the Plug-in for VMware vSphere cannot do the following:

- 1) Backup RDMs attached to VMs
- 2) Perform application-consistent Snapshot copies of applications or databases running inside the VMs.

Action: Use a corresponding SnapCenter plug-in for the supported application or database and access it using the SnapCenter GUI.

3. Select a datastore as a parent entity to create a resource group of virtual machines, and then select the virtual machines from the available list. Click Next.

Create Resource Group



1. General info & notification

2. Resource

3. Spanning disks

4. Policies

5. Schedules

6. Summary

Parent entity: FlexPod-DC

Q Enter entity name

Available entities

- datastore1
- datastore1 (1)
- datastore1 (2)
- datastore1 (3)
- infra_nfs_datastore1
- infra_nfs_datastore2
- infra_swap
- jam1
- na_FC_DS_1

Selected entities

- infra_datastore

»

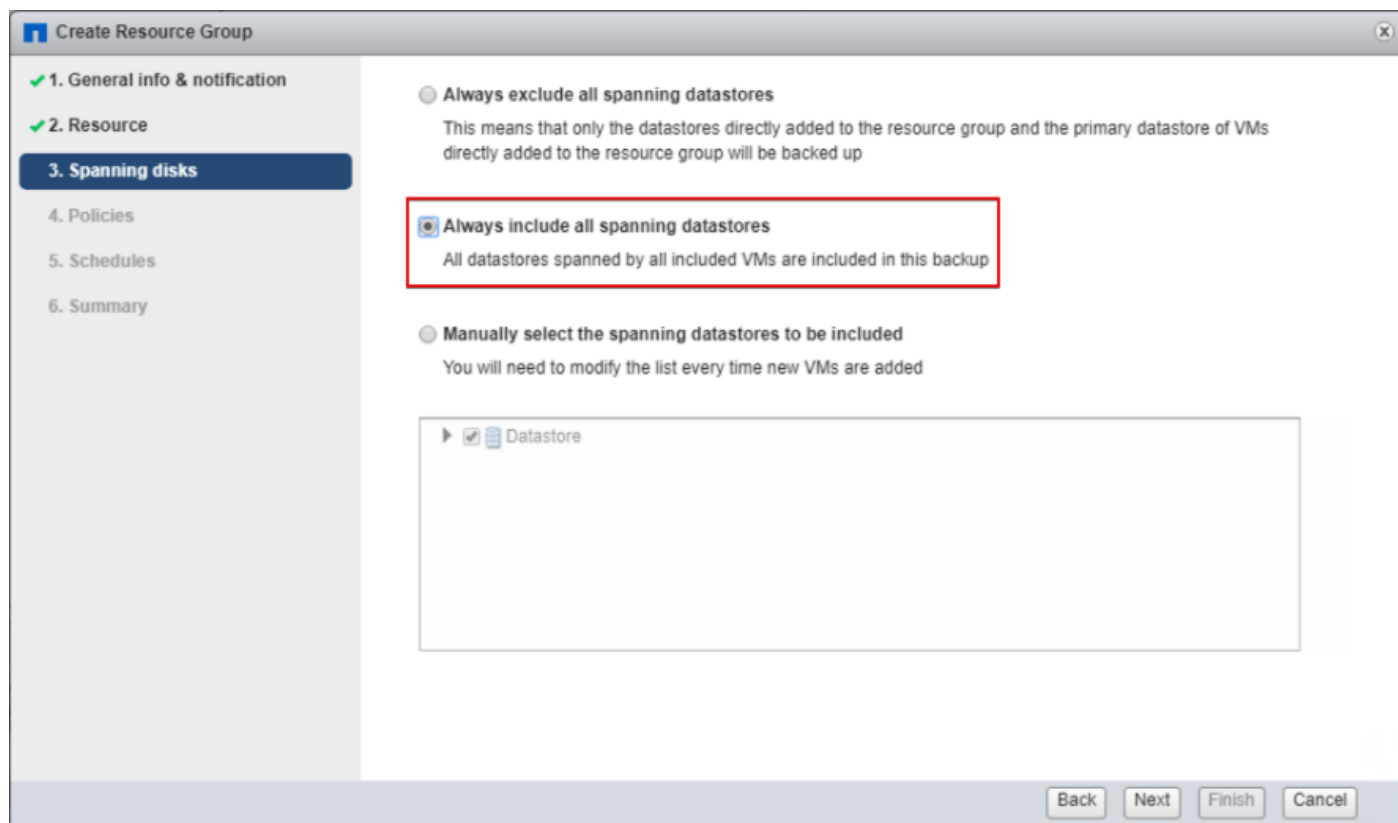
>

<

«

Back Next Finish Cancel

4. From the Spanning Disks options, select the Always Include all Spanning Datastores option.



5. From the Policies options, select the policies (from those already created earlier) that you want to associate with the resource group. Click Next.


Create Resource Group ✕



- ✓ 1. General info & notification
- ✓ 2. Resource
- ✓ 3. Spanning disks
- ✓ 4. Policies
- 5. Schedules**
- 6. Summary

Infra_vm_bac... ▼

Type Daily

Every Day(s)

Starting 

At AM  

7. Review the summary and click Finish to complete the creation of the resource group.

Create Resource Group X

- ✓ 1. General info & notification
- ✓ 2. Resource
- ✓ 3. Spanning disks
- ✓ 4. Policies
- ✓ 5. Schedules
- ✓ 6. Summary

Name	Infra_VMs
Description	
Send email	Errors
Email send from	snapcenter@flexpod.cisco.com
Email send to	flexadmin@flexpod.cisco.com
Email subject	Backup errors: Infra-VMs
Custom snapshot format	None ⓘ
Entities	infra_datastore
Spanning	True
Policies	Infra_vm_back... : Daily

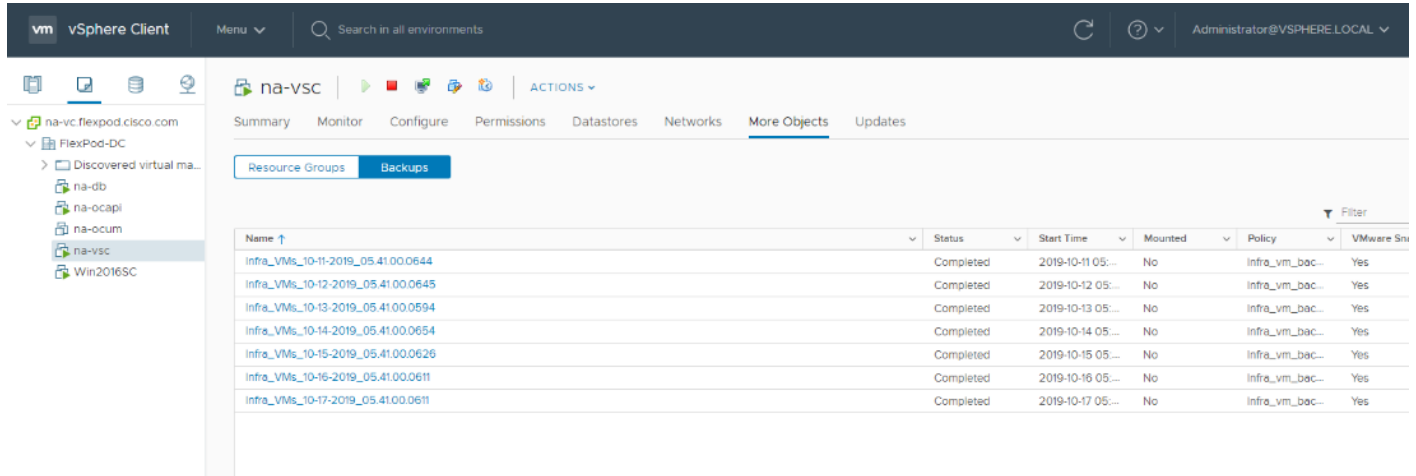
Back
Next
Finish
Cancel

View Virtual Machine Backups and Restore from vCenter by Using SnapCenter Plug-In

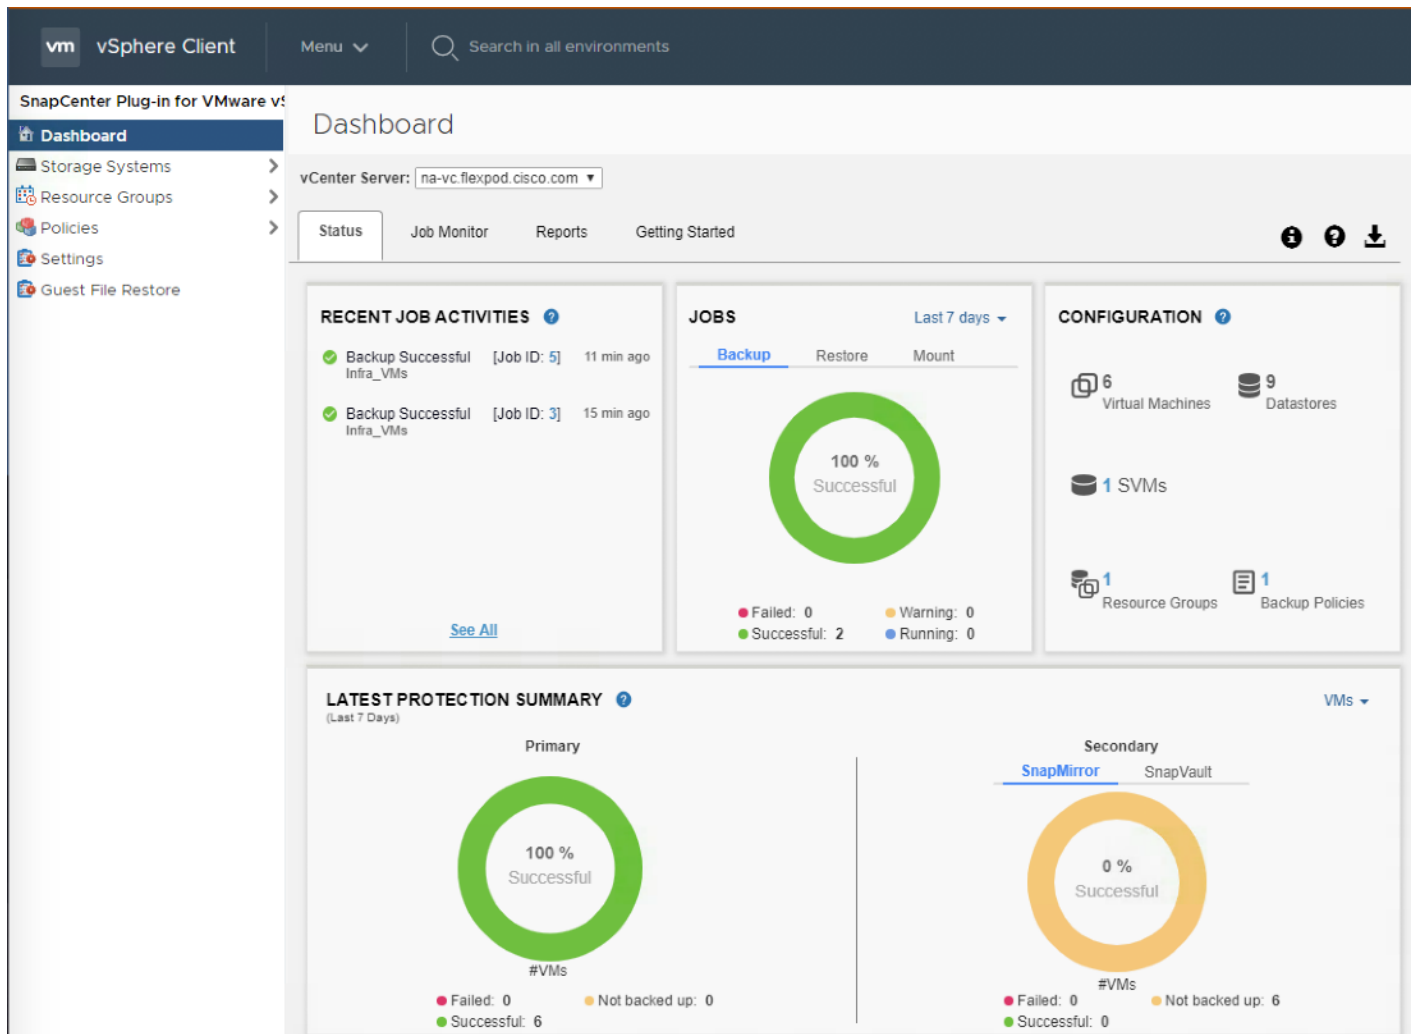
View Backups

Backups of the resources included in the resource group occurs according to the schedule of the policies associated with the resource group. To view the backups associated with each schedule, follow these steps:

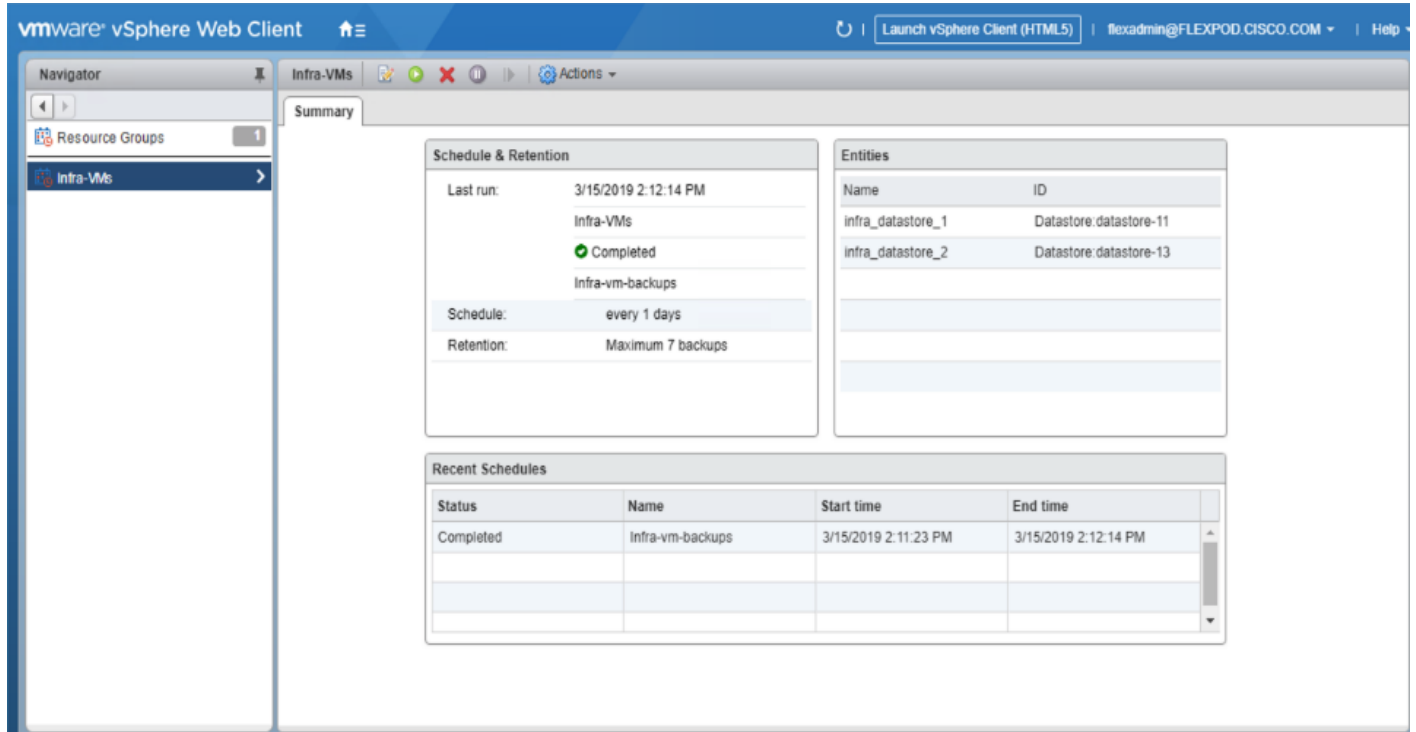
1. Go to any virtual machine that has been added to Resource Group > More Objects > Backups. It shows all the backups for that resource.



2. On the SnapCenter Plugin for VMware vSphere check-out the Dashboard.

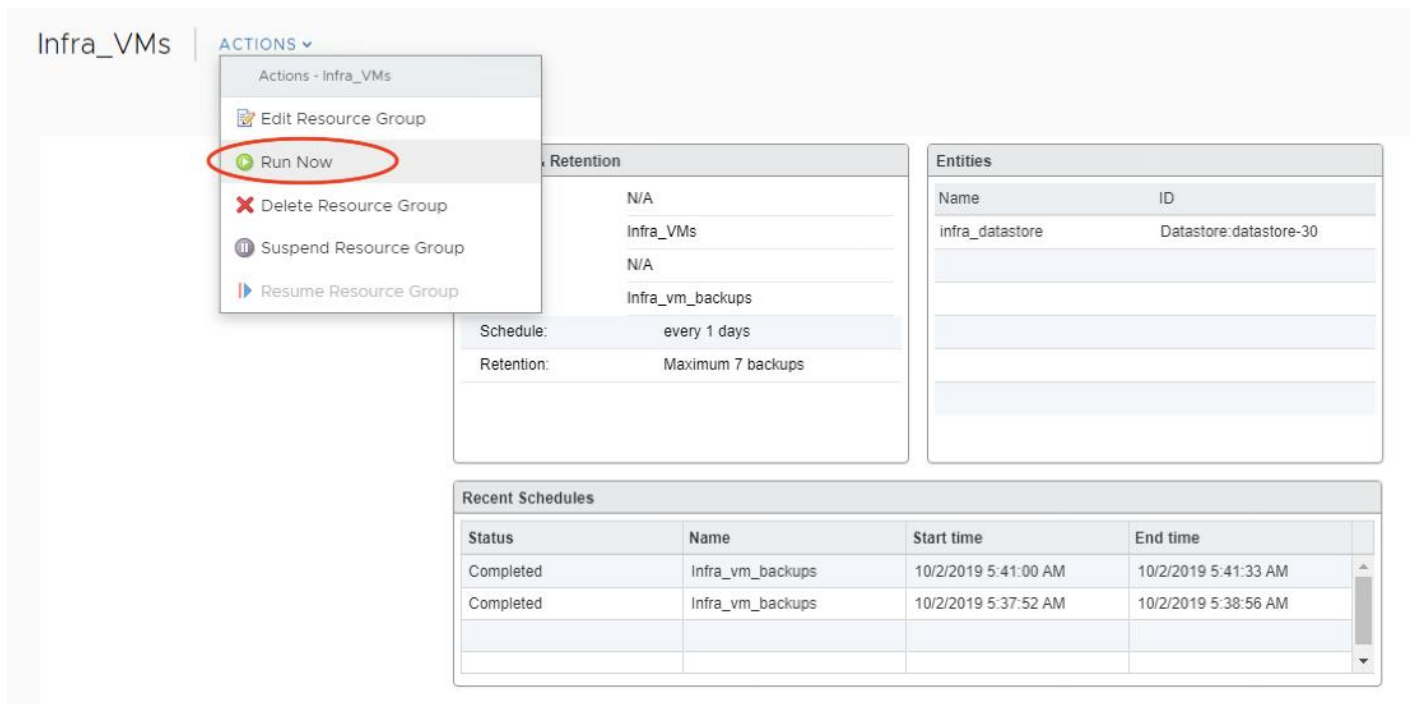


3. In the SnapCenter Plug-In for VMware vSphere, Click Resource Groups and select any resource group. In the right pane, the completed backups are displayed.



4. An on-demand backup can also be created for any resource group by following these steps:

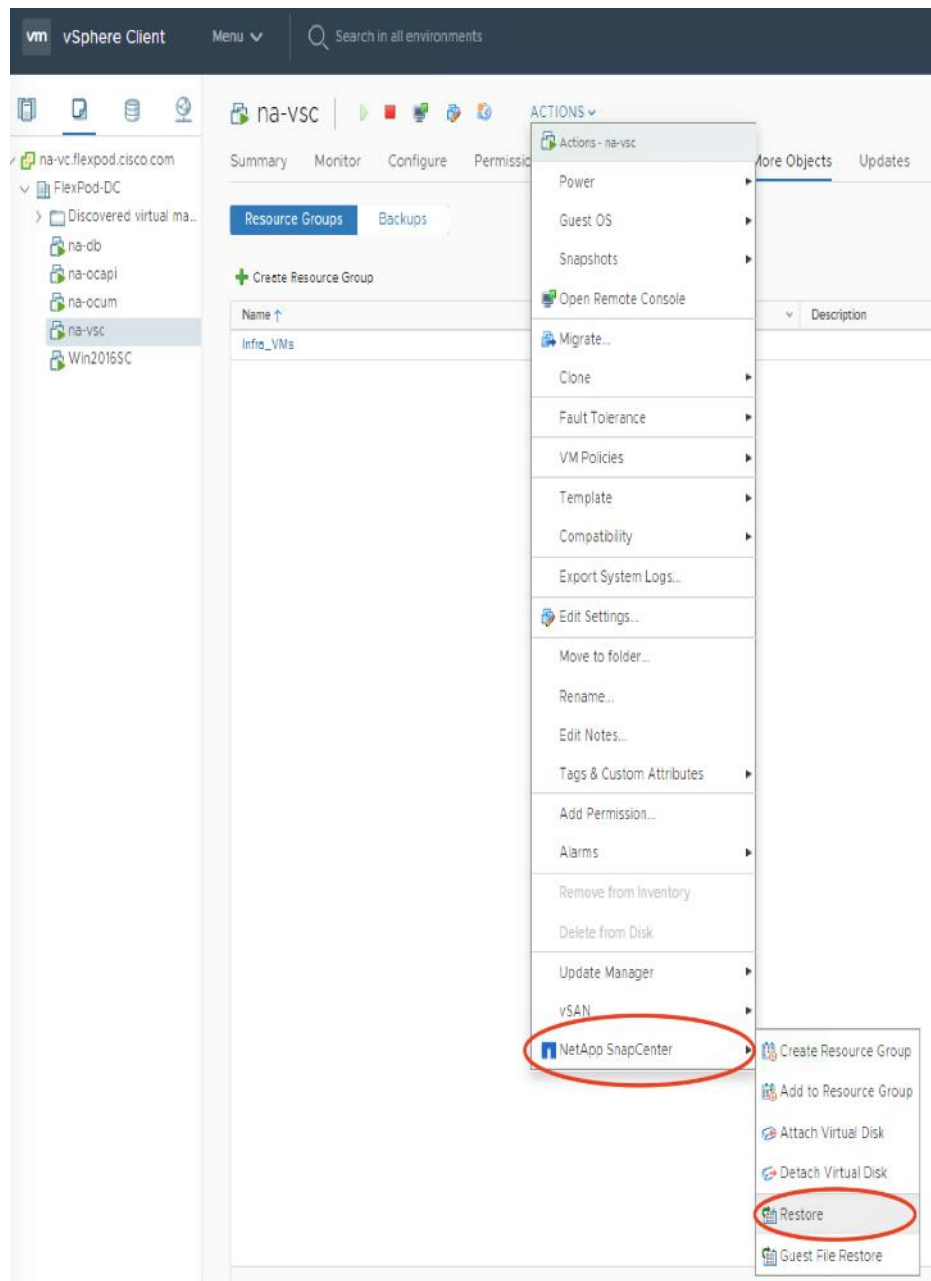
- a. Click the virtual machine or datastore resource > More Objects > Resource Groups. Select the Resource Group and click Actions. Select Run Now. This process triggers the backup of the resource group containing a particular resource.



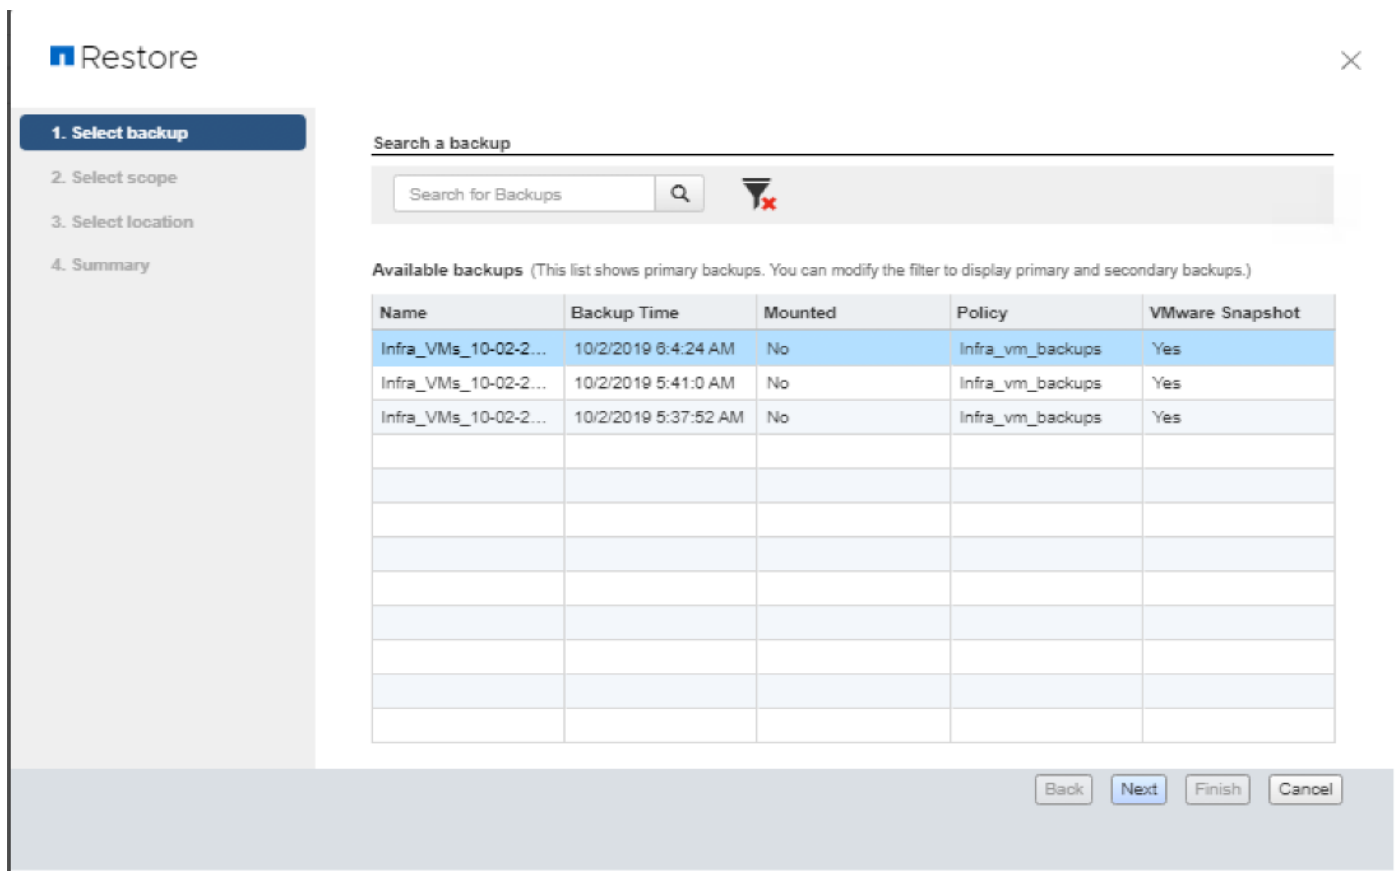
Restore from vCenter by Using SnapCenter Plug-In

To restore from vCenter by using SnapCenter Plug-In, follow these steps:

1. The Plug-In for VMware vSphere provides native backup, recovery, and cloning of virtualized applications. Go to vCenter Web Client GUI > Select a VM > right-click and select NetApp SnapCenter.



2. Select a backup from which to restore. Click Next.




- From the Restore Scope drop-down list, select either Entire virtual machine to restore the virtual machine with all VMDKs, or select Particular Virtual Disk to restore the VMDK without affecting the virtual machine configuration and other VMDKs.

Restore ×

- ✓ 1. Select backup
- ✓ 2. Select scope**
- 3. Select location
- 4. Summary


Restore scope	Entire virtual machine
Restored VM name	na-vmc
ESXi host name	na-esxi-1.flexpod.cisco.com
Restart VM	<input type="checkbox"/>

4. Select the destination datastore. Click Next.

 Restore
✕

- ✓ 1. Select backup
- ✓ 2. Select scope
- ✓ 3. Select location
- 4. Summary

Virtual machine to be restored	na-vsc
Backup name	Infra_VMs_10-02-2019_08.04.24.0309
Restart virtual machine	No
ESXi host to be used to mount the backup	na-esxi-1.flexpod.cisco.com


This virtual machine will be powered down during the process.

Provision Virtual Volumes (VVols)

The NetApp vSphere APIs for Storage Awareness or VASA enables the vSphere Administrator to take advantage of the new capabilities for storage-based policy management and automation. The NetApp VASA provider is used to provide automation for DevOps and private cloud deployments and enable the vSphere Administrator to provision storage for virtual machines based on pre-configured policies created by the Storage Administrator. Specific storage types, performance, and storage efficiency characteristics for each storage tier in the environment can be expressed in a set of policies to ensure the appropriate type of storage is provisioned to the virtual machine every time.

Verify NDMP Vserver Scope Mode

To verify the NDMP Vserver scope mode, follow these steps:

1. View NDMP scope mode with the following command:

```
system services ndmp node-scope-mode status
NDMP node-scope-mode is enabled.
```

2. Disable NDMP node-scoped mode.

```
system services ndmp node-scope-mode off
NDMP node-scope-mode is disabled.
```

3. Enable NDMP services on the vserver.

```
vserver services ndmp on -vserver Infra-SVM
```

OnCommand API Services for ONTAP VASA Provider

OnCommand API Services enables you to add storage systems running ONTAP, Cloud Volumes ONTAP, ONTAP Select to retrieve storage-related information, and to provision storage resources.

To configure API services to enable the VSC rich dashboard for Virtual Volumes, follow these steps:

Download the following packages before beginning the installation and upload them to the /tmp directory on a RHEL 7.5 or CentOS 7.5 host you will use as the API services host:

- MySQL YUM Repository: <https://dev.mysql.com/downloads/repo/yum>
- mysql80-community-release-el7-3.noarch.rpm
- [OnCommand API Services 2.2.0 binary](#)

1. Install a separate CentOS 7.5 VM.
2. Open the required firewall ports with the following command:

```
firewall-cmd --zone=public --add-port=443/tcp --permanent;firewall-cmd --reload
```

3. Download the Extra Package for Enterprise Linux (EPEL).

```
cd /tmp
wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

4. Install the EPEL repository.

```
yum install epel-release-latest-7.noarch.rpm
```

5. Install the MySQL community edition repository.

```
yum install mysql80-community-release-el7-3.noarch.rpm
```

6. Remove any mariadb components.

```
yum remove mariadb*
```

7. Set the executable bit on the OnCommand API Services binary.

```
chmod +x OnCommand-API-services-2.2.0.bin
```

8. Edit the /etc/yum.repos.d/mysql-community.repo. Locate the subrepository block and set enable=0 on mysql80-community, and enable=1 on mysql57-community.

```
# Enable to use MySQL 5.7
[mysql57-community]
name=MySQL 5.7 Community Server
baseurl=http://repo.mysql.com/yum/mysql-5.7-community/el/7/$basearch/
```



```

enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

# Disabled
[mysql80-community]
name=MySQL 8.0 Community Server
baseurl=http://repo.mysql.com/yum/mysql-8.0-community/el/7/$basearch/
enabled=0
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

```

9. Install the MySQL 5.7 community server and additional libraries.

```
yum install mysql-community-server mysql-community-libs-compat*
```

10. Install OnCommand API Services 2.2.

```
./OnCommand-API-services-2.2.0.bin
```

11. When OnCommand API services installation is complete record the URL for use to connect with the VSC.

```

Complete!

OnCommand API Services installed successfully.

Use a web browser and one of the following URLs to configure and access the OnCommand API Services GUI.

    https://10.1.156.206/admin/
    https://na-ocapi/admin/
    https://na-ocapi.flexpod.cisco.com/admin/

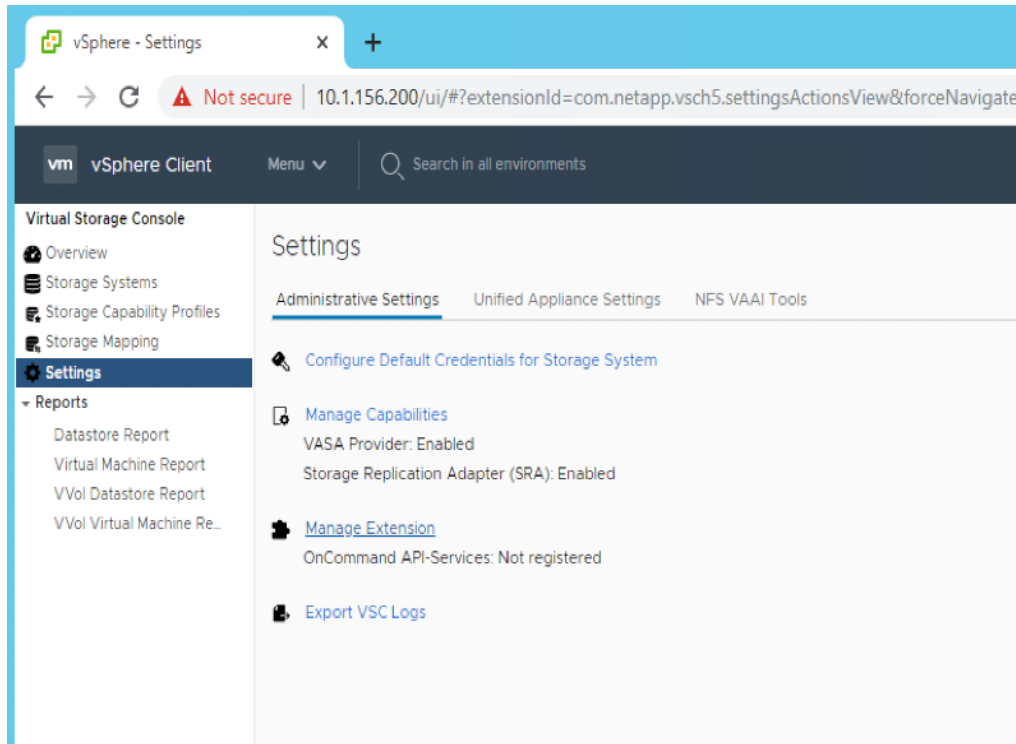
Use a web browser and the following URL to access OnCommand API Services documentation.

    https://10.1.156.206/
    https://na-ocapi/
    https://na-ocapi.flexpod.cisco.com/

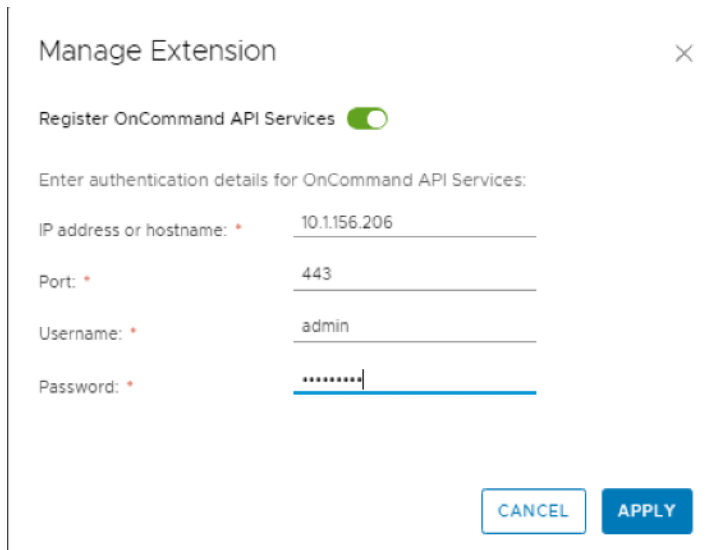
[root@na-ocapi /]# █

```

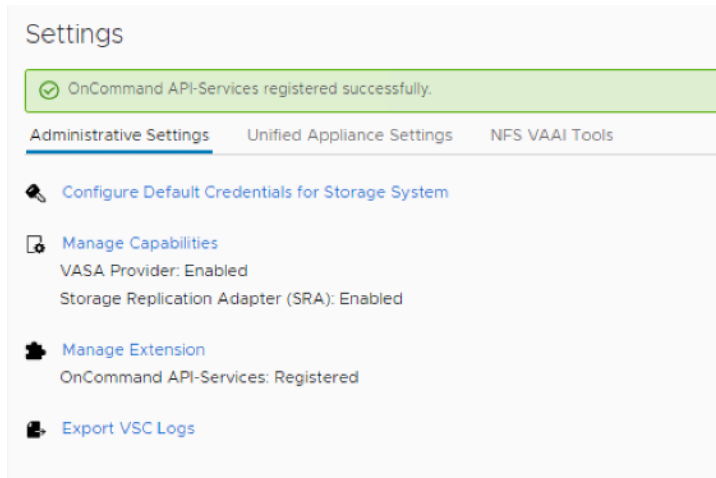
12. Navigate to the NetApp Virtual Storage Console from Home.



13. Select Manage VASA Provider Extensions and enter the configuration information recorded from step 11



14. Click Apply and notice that the API-Services registered successfully.



15. Navigate back to the Summary option in the Navigator pane to display the WVOL dashboard

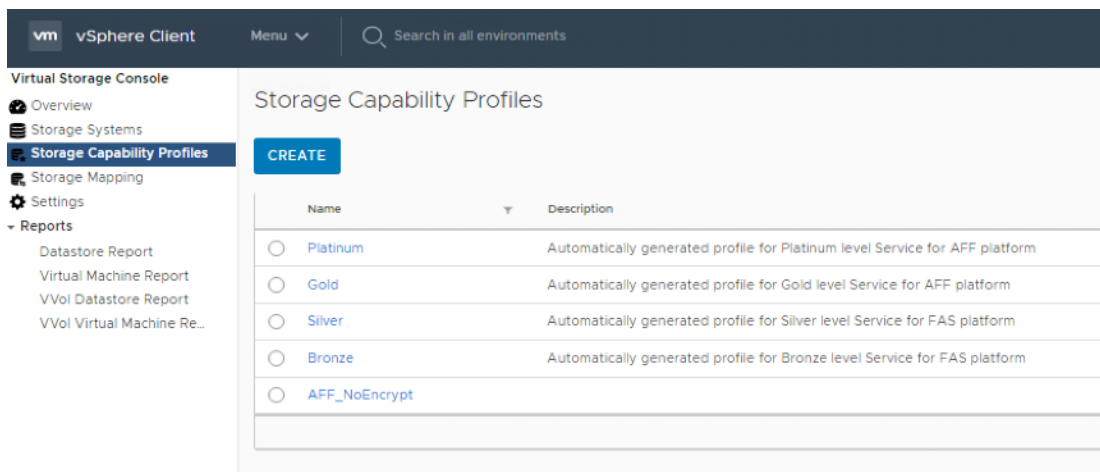
Create the Storage Capability Profile

In order to leverage the automation features of VASA and Virtual Volumes two primary components must first be configured. The Storage Capability Profile (SCP) and the VM Storage Policy. The Storage Capability Profile expresses a specific set of storage characteristics into one or more profiles used to when provisioning a Virtual Volume. The SCP is specified as part of VM Storage Policy which is specified when you deploy a virtual machine. NetApp Virtual Storage Console comes with Five pre-configured Storage Capability Profiles mapped to different storages types from FAS to AFF. These storage profiles match the adaptive QoS policies in ONTAP 9.6.

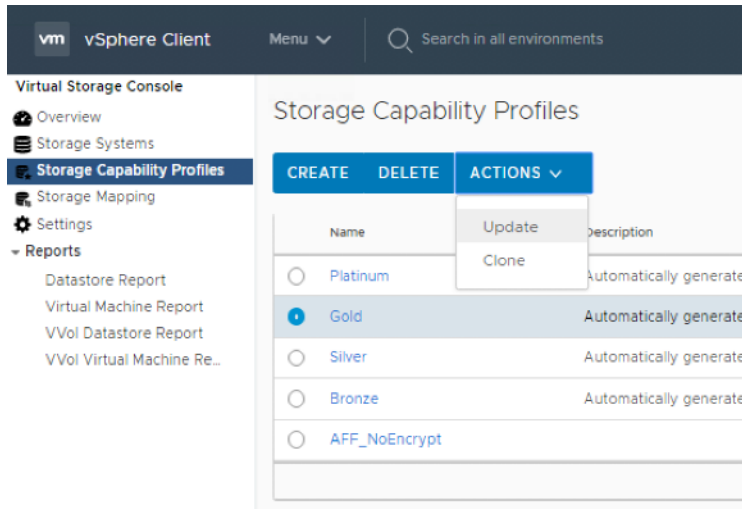
Default Policy Group	Expected IOPS / TB GB	Peak IOPS / TB GB	Absolute Minimum IOPS / TB GB
Extreme	6144	12,288	1000
Performance	2148	4096	500
Value (FAS only)	128	512	75

To review or edit one of the built-in profiles pre-configured with VSC 7.2.1, follow these steps:

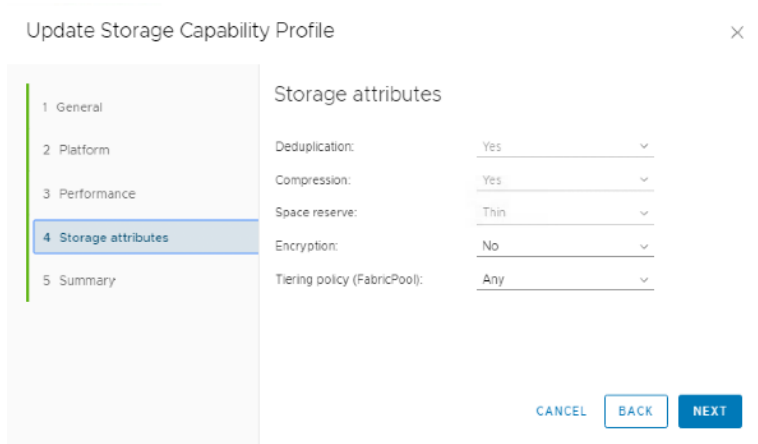
1. In the NetApp Virtual Storage Console select Storage Capability Profiles.




2. Select the Gold Storage Capability Profile and choose Update from the toolbar.



3. Click number 4. Storage Attributes to edit the storage capabilities.



4. Review the Storage Attributes and make any desired changes.

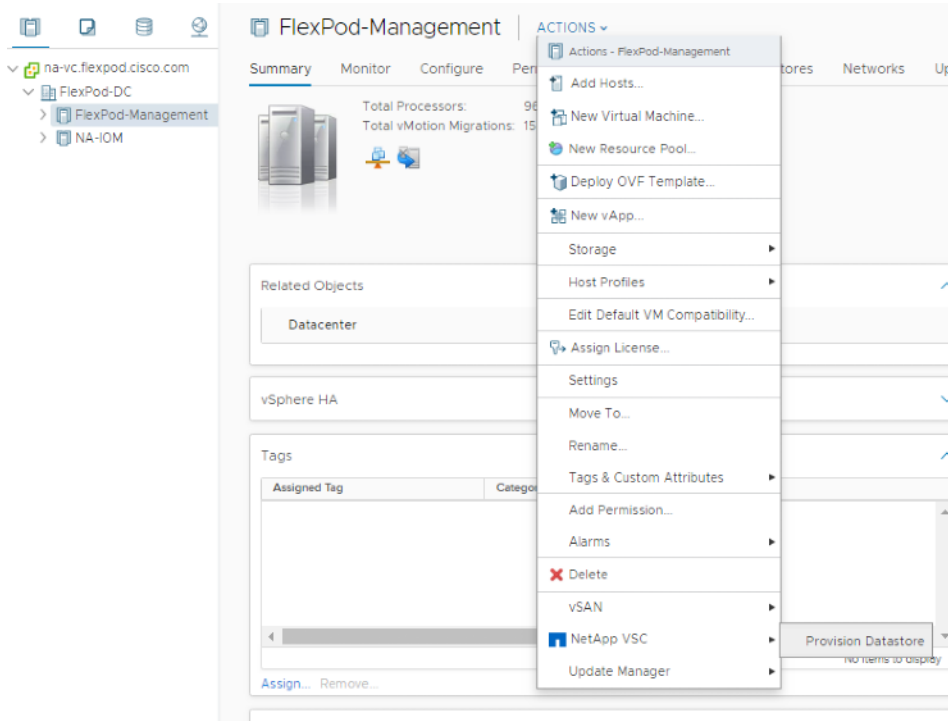
 **It is recommended to Clone the Storage Capability Profile if you wish to make any changes to the default profiles rather than editing the built-in profile.**

5. Click Cancel or Finish if any changes were made.

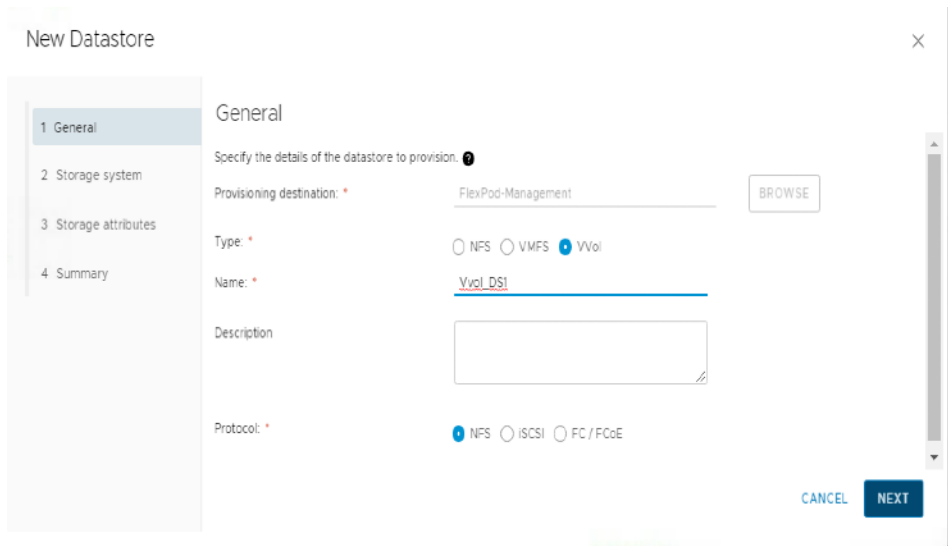
Provision VVOL from Virtual Storage Console

To provision VVOL from Virtual Storage Console, follow these steps:

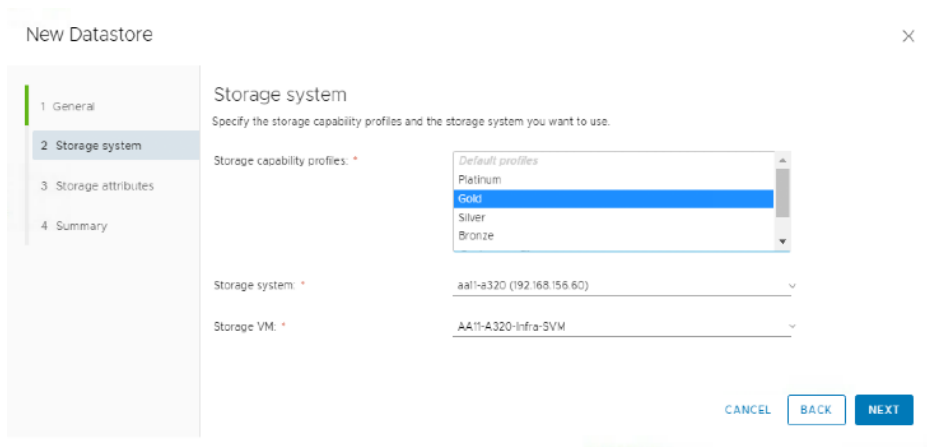
1. Starting from Home > Hosts and Clusters.
2. Select FlexPod-Management datacenter > Actions > Provision Datastore.



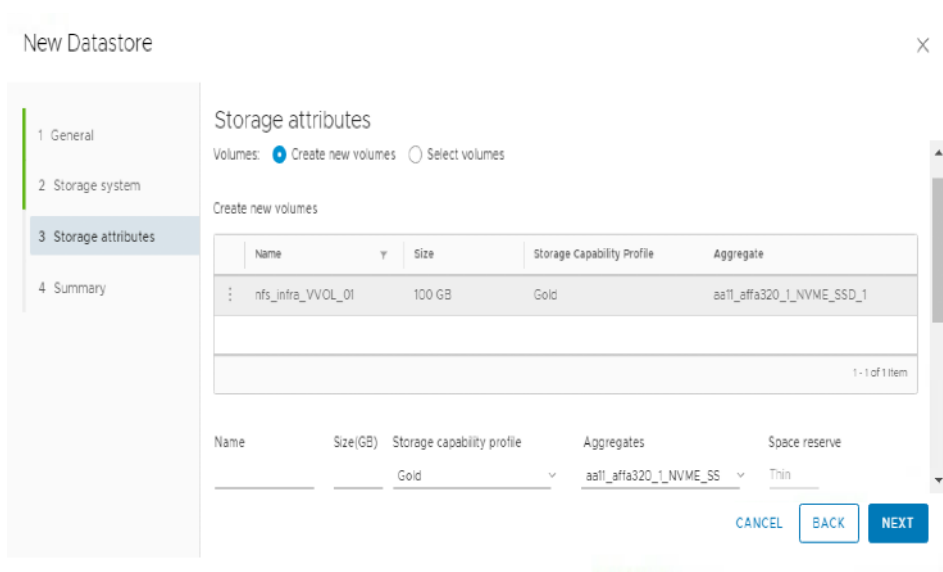
3. Provide a name, a description for the volume, and choose NFS for the protocol. Select Next.



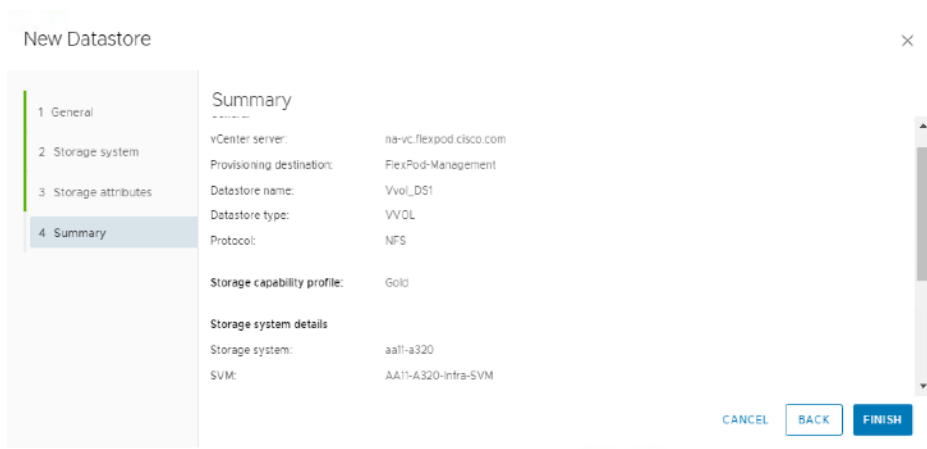
4. In the Storage System section Select the Gold Storage Capability Profile or another appropriate SCP for the target workload. Provide the Storage system and the Storage VM detail. Choose Next.



5. In the Storage Attributes Section, select Create New Volume and provide the details.



6. Review the details and click Finish.



Create Virtual Machine with Assigned VM Storage Policy

To create a virtual machine and assigned a VM storage policy, follow these steps:

1. Navigate to the VMs and Templates tab and select the FlexPod-DC datacenter.
2. Click Actions and select New Virtual Machine.
3. Select Create a new virtual machine and choose Next.
4. Enter a name for the VM and select the FlexPod-DC datacenter.
5. Choose the FlexPod-Management Data compute Resource
6. Select the VVOI Datastore for the storage and click Next.

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Select storage
- 5 Select compatibility
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select storage
Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy: Datastore Default ▾

Name	Capacity	Provisioned	Free	Ty
datastore1	7.5 GB	4.37 GB	3.13 GB	VI
datastore1 (f)	7.5 GB	1.41 GB	6.09 GB	VI
datastore1 (6)	7.5 GB	1.41 GB	6.09 GB	VI
infra_datastore	3 TB	2.61 TB	2.6 TB	NI
infra_datastore_2	3 TB	228.39 GB	2.89 TB	NI
infra_swap	200 GB	88.21 MB	199.91 GB	NI
iom_datastore	500 GB	938.99 GB	418.56 GB	NI
na_FC_DS_1	100 GB	6.17 GB	93.83 GB	VI
VVOI_DS1	100 GB	14 MB	99.99 GB	VI

Compatibility

✓ Compatibility checks succeeded.

CANCEL
BACK
NEXT

7. Select Compatibility and click Next.
8. Select the Guest OS and click Next.
9. Customize the hardware for the VM and click Next.
10. Review the details and click Finish.

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- ✓ 6 Select a guest OS
- ✓ 7 Customize hardware
- 8 Ready to complete

Ready to complete
Click Finish to start creation.

Provisioning type	Create a new virtual machine
Virtual machine name	VVOL-VM-01
Folder	FlexPod-DC
Cluster	FlexPod-Management
Datastore	Vvol_DS1
Guest OS name	Microsoft Windows Server 2012 (64-bit)
Virtualization Based Security	Disabled
CPUs	2
Memory	4 GB
NICs	1
NIC 1 network	IB-MGMT Network
NIC 1 type	E1000E
SCSI controller 1	LSI Logic SAS
Create hard disk 1	New virtual disk
Capacity	40 GB

CANCEL
BACK
FINISH

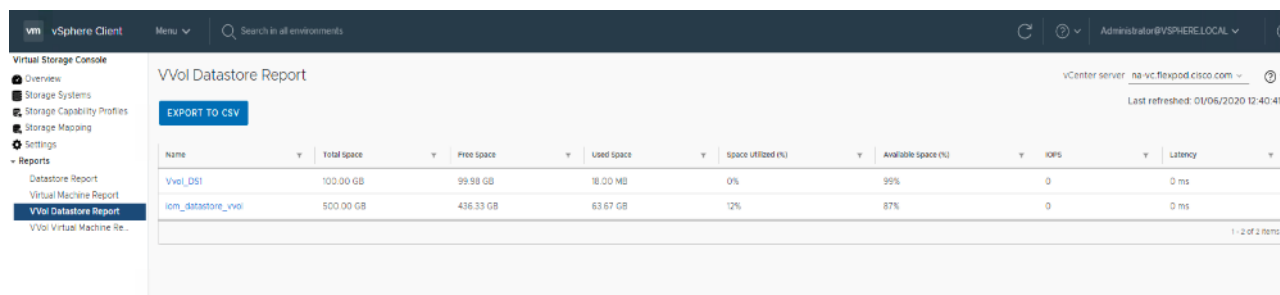
VSC VVOL Dashboard and Reports

To configure the VSC VVOL dashboard and reports, follow these steps:

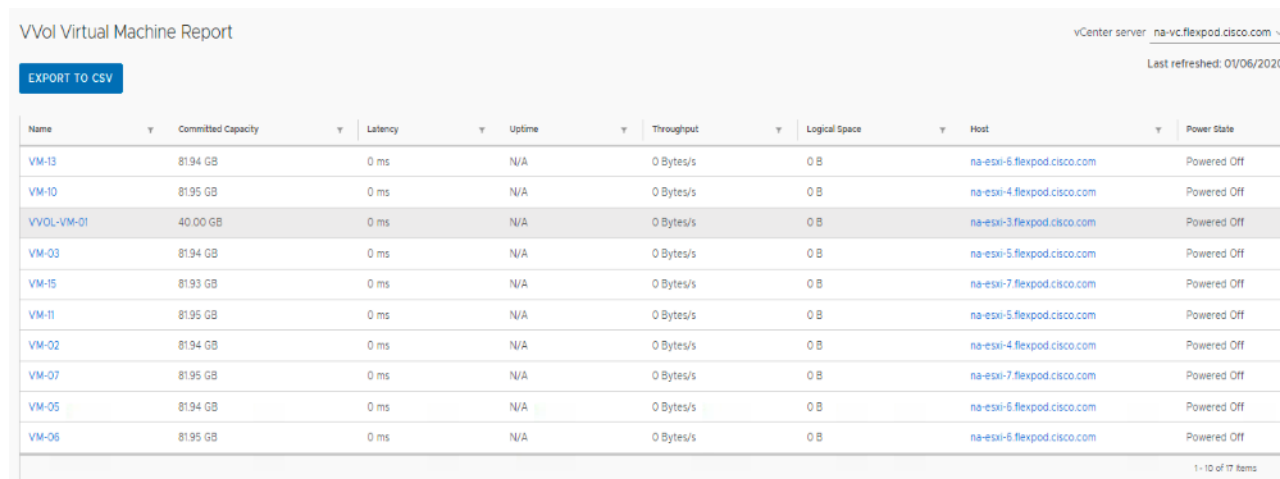
1. In the Virtual Storage Console-> Overview-> Click the VVol dashboard.

The screenshot displays the VVol Dashboard in the Virtual Storage Console. The interface includes a navigation bar with 'Getting Started', 'Traditional Dashboard', and 'VVol Dashboard'. A status bar at the top right shows 'vCenter server na-vc.flexpod.cisco.com' and refresh information. The main content area is divided into three sections: Overview, Datasstores, and Virtual Machines. The Overview section features a donut chart for Datastore capacity (63.68 GB used, 536.32 GB free, total 600.00 GB) and IOPS metrics (Read, Write, Total). The Datasstores section shows a table of the top 5 datasstores by space utilized, with 'lvm_datastore_vvol' at 12.73% and 'Vvol_DS1' at 0.01%. The Virtual Machines section shows a table of the top 5 VMs by committed capacity, all at 81.95 GB.

2. In the VSC -> Reports -> click the VVOL Datastore reports.



3. In the VSC-> Reports -> Click the VVol VM reports.



Active IQ Unified Manager 9.6

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface. Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems.

This section describes the steps to deploy NetApp Active IQ Unified Manager 9.6 on a CentOS Linux 7.6 virtual machine. The following table lists the recommended configuration for the virtual machine to install and run Active IQ Unified Manager to ensure acceptable performance.

Table 11 Virtual Machine Configuration

Hardware Configuration	Recommended Settings
RAM	12 GB
Processors	4 CPUs/ vCPUs
CPU Cycle Capacity	9572 MHz total

Hardware Configuration	Recommended Settings
Free Disk Space	Red Hat or CentOS <ul style="list-style-type: none"> • 50 GB allotted to the root partition • 100 GB of free disk space allotted to the /opt/netapp/data directory, which is mounted on an LVM drive on a separate local disk attached to the target system



The /tmp directory should have at least 10 GB of free space and the /var/log directory should have at least 16GB of free space.

To install Unified Manager 9.6, follow these steps:

1. Download NetApp Active IQ Unified Manager for VMware vSphere .ova file from Netapp support site (<https://mysupport.netapp.com>).
2. From the VMware VCenter, click the VMs and Templates Tab, then click In Actions> Deploy OVF Template.
3. Specify the location of the OVF Template and click Next.
4. On the Select a name and folder page, enter a unique name for the VM or vApp, and select a deployment location, and then click Next.

Deploy OVF Template

5. On the Select a compute resource page, select a resource where you want to run the deployed VM template, and click Next.
6. On the Review details page, verify the .ova template details and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Review details**
- 5 License agreements
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Review details
Verify the template details.

⚠ The OVF package contains advanced configuration options, which might pose a security risk. Review the advanced configuration options below. Click next to accept the advanced configuration options.

Publisher	No certificate present
Product	Active IQ Unified Manager
Vendor	NetApp, Inc.
Description	Active IQ Unified Manager - Application to monitor and manage NetApp storage systems. For more information or support please visit http://www.netapp.com
Download size	1.7 GB
Size on disk	3.5 GB (thin provisioned)
	152.0 GB (thick provisioned)
Extra configuration	keyboard.typematicMinDelay = 2000000

CANCEL BACK NEXT

7. On the License agreements page, select the checkbox for I accept all license agreements.
8. On the Select storage page, define where and how to store the files for the deployed OVF template.
 - a. Select the disk format for the VMDKs.
 - b. Select a VM Storage Policy.
 - c. Select a datastore to store the deployed OVA template.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- 6 Select storage**
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Select storage
Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type
datastore1	7.5 GB	4.37 GB	3.13 GB	Virtual
datastore1 (1)	7.5 GB	4.37 GB	3.13 GB	Virtual
datastore1 (2)	7.5 GB	4.33 GB	3.17 GB	Virtual
datastore1 (3)	7.5 GB	4.33 GB	3.17 GB	Virtual
infra_datastore	500 GB	891.86 GB	52.79 GB	NFS
infra_nfs_datastore1	10 GB	512 KB	10 GB	NFS
infra_nfs_datastore2	20 GB	3.86 MB	20 GB	NFS
infra_swap	100 GB	144.55 MB	99.86 GB	NFS
jem1	100 GB	0 B	100 GB	Virtual
...

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

9. On the Select networks page, select a source network and map it to a destination network, and then click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- 7 Select networks**
- 8 Customize template
- 9 Ready to complete

Select networks
Select a destination network for each source network.

Source Network	Destination Network
nat	IB-MGMT Network

1 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL BACK NEXT

10. On the Customize template page, provide network details.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

Customize template

Customize the deployment properties of this software solution.

✓ All properties have valid values
✕

Networking configuration	7 settings
Enables Auto IPv6 addressing for vApp.	IPv6 Auto addressing is set if the checkbox is checked and all the fields are left empty. <input type="checkbox"/>
Host FQDN	Specifies the hostname for the appliance. Leave blank if DHCP is desired. <input type="text" value="na-aiqum"/>
IP Address	Specifies the IP address for the appliance. Leave blank if DHCP is desired. <input type="text" value="10.156.203"/>
Network Mask (or) Prefix Length	Specifies the subnet to use on the deployed network. Leave blank if DHCP is desired. <input type="text" value="255.255.255.0"/>
Gateway	Specifies the gateway on the deployed network. Leave blank if DHCP is desired. <input type="text" value="10.156.254"/>
Primary DNS	Primary DNS ip address. Leave blank if DHCP is desired. <input type="text" value="10.156.250"/>
Secondary DNS	Secondary DNS ip address. Leave blank if DHCP is desired. <input type="text" value="10.156.251"/>

CANCEL

BACK

NEXT

11. On the Ready to complete page, review the page and click Finish.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- ✓ 8 Customize template
- 9 Ready to complete**

Ready to complete

Click Finish to start creation.

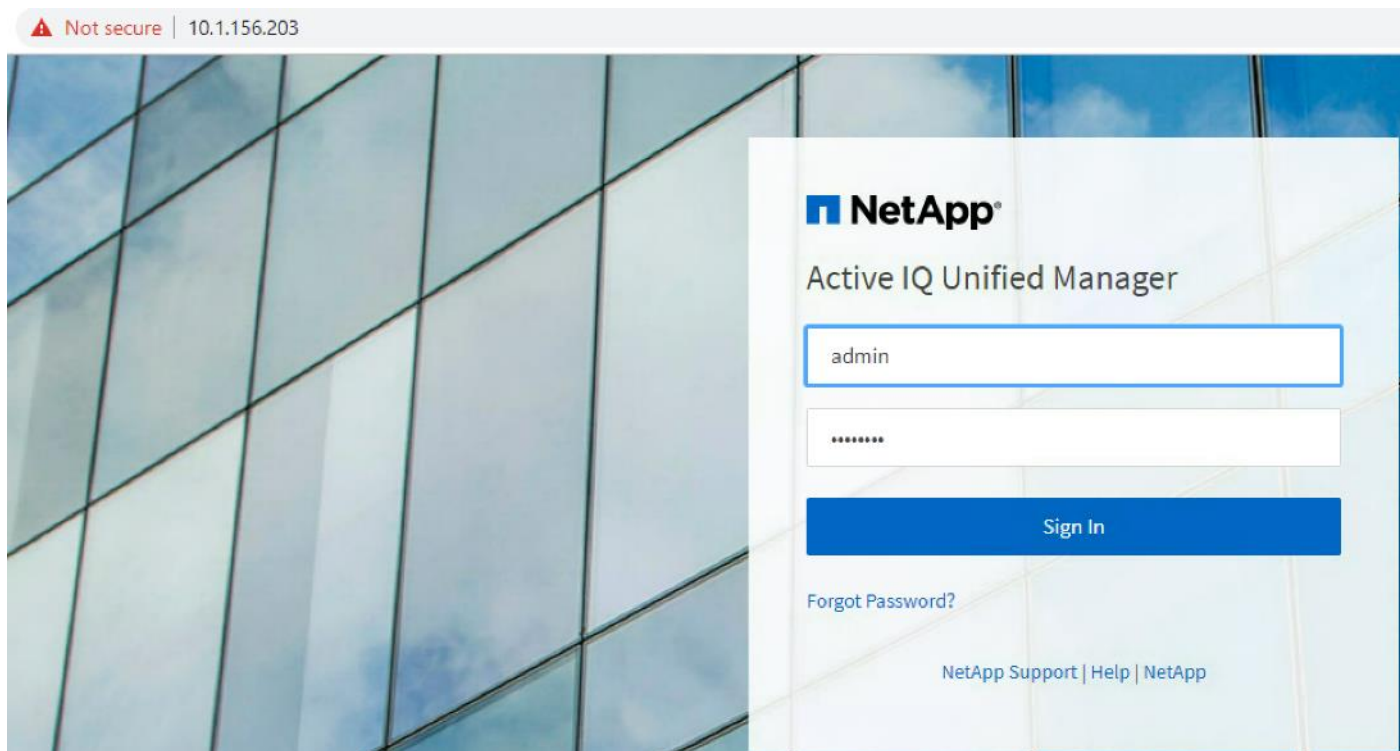
Provisioning type	Deploy from template
Name	na-aiqum
Template name	ActiveIQUnifiedManager-9.6
Download size	1.7 GB
Size on disk	152.0 GB
Folder	Discovered virtual machine
Resource	FlexPod-Management
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thick provision lazy zeroed
Network mapping	1
nat	IB-MGMT Network
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual

CANCEL

BACK

FINISH

12. Navigate to the VM in the VMs and Templates Tab, then click the Summary tab, and then click the Power On box to start the virtual machine.
13. While the virtual Machine is powering on, right-click the deployed virtual machine and then click Install VMware tools.
14. Log into NetApp Active IQ UM using the IP address displayed on the deployment screen. Please use the credentials for the admin user or the maintenance user.



Configure Active IQ Unified Manager

To configure Active IQ Unified Manager and add a storage system for monitoring, follow these steps:

1. Log into the web browser of Active IQ Unified Manger.
2. Enter the email address that Unified Manager will use to send alerts, enter the mail server configuration, and the IP address or hostname of the NTP server. Select Next and complete the AutoSupport configuration.

Active IQ Unified Manager

1 Email 2 AutoSupport 3 Finish

Setup Email Settings

Maintenance User Email

Email:

SMTP Server

Hostname:

Port:

Username:

Password:

Use START / TLS:

Use SSL:

NTP Server

Host Name or IP Address:

Next

3. Save and complete the configuration and continue to the next step.

Active IQ Unified Manager

1 Email 2 AutoSupport 3 Finish

Finish and Start Adding Clusters

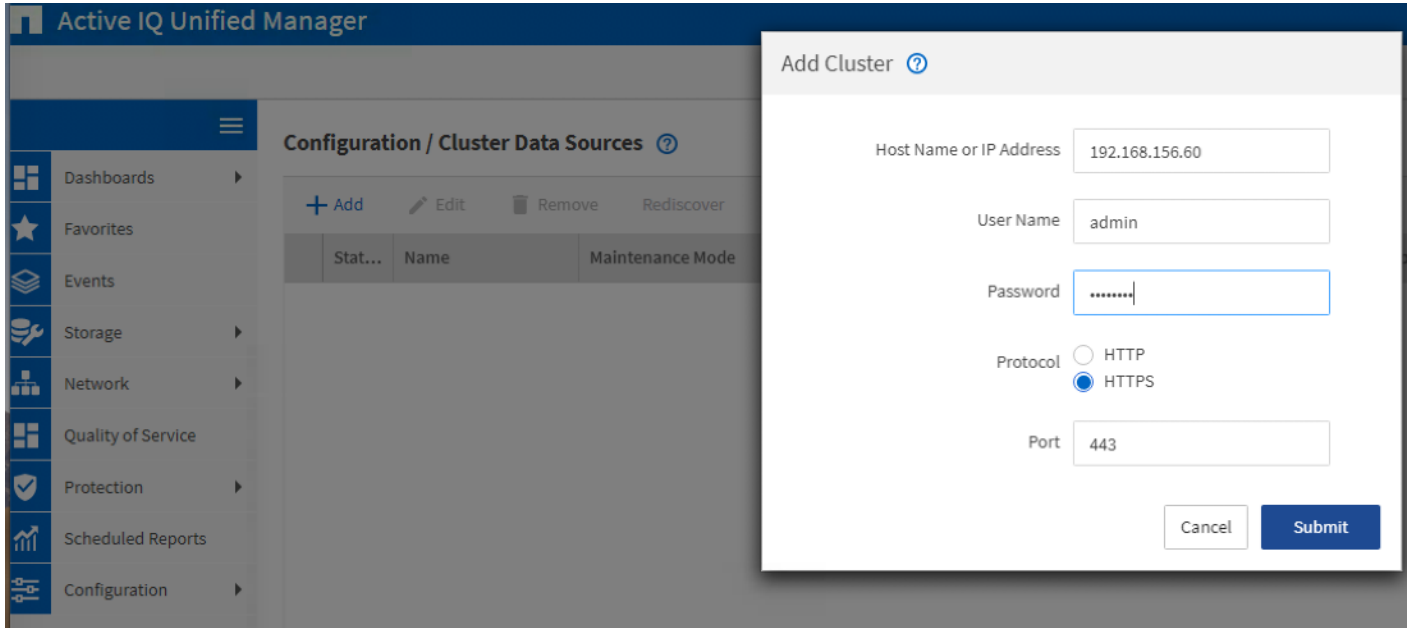
Click 'Save and Continue' to finish and proceed adding clusters to be monitored.

Save and Continue

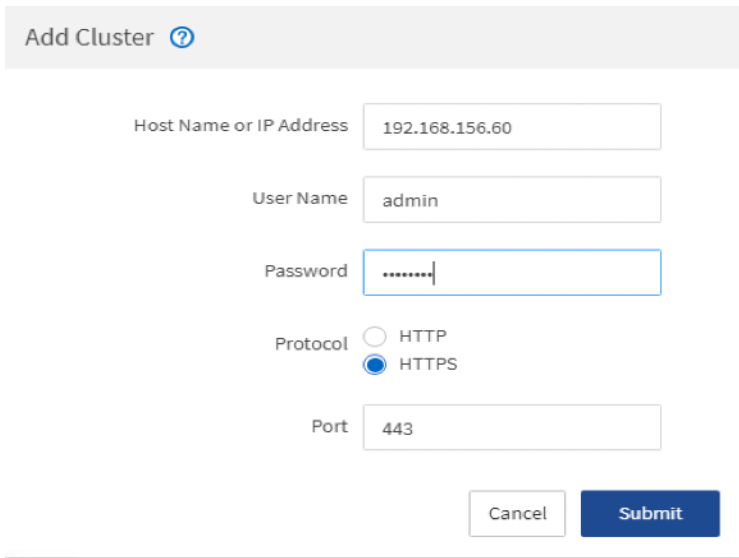
Add Storage to Active IQ Unified Manager

To add storage to the Active IQ Unified Manager, follow these steps:


1. Click Add on the Configuration/Cluster Data Sources screen to add the storage cluster.



2. Enter the hostname and login credentials to the storage cluster.



3. When prompted to trust the self-signed certificate from Active IQ Unified Manager, select Yes to finish and add the storage system. The initial discovery process can take up to 15 minutes to complete.

 **Authorize Cluster Certificate**

Host aa14-a800 you specified has identified itself with a self signed certificate for OnCommand Unified Manager.

[View Certificate](#)

Do you want to trust this certificate?

Sample Tenant Provisioning

Provision a Sample Application Tenant

This section describes a sample procedure for provisioning an application tenant. The procedure refers 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 explained in this document, fibre channel, NFS, iSCSI, and CIFS/SMB (CIFS/SMB have not been discussed in this document) 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 except fibre channel. Fibre channel will have new storage LIFs defined with the same VSANs configured for the FlexPod Infrastructure.
2. In the Nexus switches, declare all added VLANs and configure the VM VLAN as an allowed VLAN on the 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 explained in this document but is explained in the Nexus 9000 documentation. Configure the storage VLANs on the 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 fibre 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 except fibre channel. 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 FC or iSCSI service for the tenant SVM if fibre channel or 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. If fibre channel is being deployed in this tenant, configure four FCP LIFs in the tenant SVM on the same fibre channel ports as in the Infrastructure SVM.
 - j. If iSCSI is being deployed in this tenant, configure four iSCSI LIFs in the tenant SVM on the iSCSI VLAN interfaces.
 - k. Create an 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 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 vDS 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 Cisco MDS 9132T switches:
 - a. Create device aliases for the tenant ESXi host vHBAs and the FC LIFs in the tenant storage SVM.
 - b. Create zones for the tenant ESXi hosts with fibre channel targets from both the storage Infra-SVM and the tenant SVM.
 - c. Add these zones to the Fabric zoneset and activate the zoneset.
6. 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 WWPNs 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 FC or iSCSI storage can be created later using NetApp VSC.
7. Install and configure VMware ESXi on all tenant host servers. It is not necessary to map infra_datastore unless you want the tenant ESXi hosts to have access to VMs or VM templates in these datastores.
8. In VMware vCenter, create a cluster for the tenant ESXi hosts. Add the hosts to the cluster.
9. Using the vCenter Web Client, add the tenant hosts to vDS0 or create a tenant vDS and add the hosts to it. In vDS0, add port-profiles for the tenant VLANs. When migrating the hosts to the vDS, leave only the ESXi management interfaces on vSwitch0.
10. 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. Tenant iSCSI VMkernel ports can be created on the vDS with the port groups pinned to the appropriate fabric.
11. Using the NetApp VSC plugin to the vCenter Web Client, set recommended values for all tenant ESXi hosts. Ensure the NetApp NFS Plug-in for VMware VAAI is installed on all tenant hosts and reboot each host.
12. You can now begin provisioning virtual machines on the tenant cluster. The NetApp VSC plugin can be used to provision fibre channel, iSCSI, and NFS datastores.
13. Optionally, use NetApp SnapCenter to provision backups of tenant virtual machines.

Appendix

FlexPod iSCSI Addition

Cisco Nexus Switch Configuration

This section is a delta section for adding infrastructure iSCSI to the Nexus switches. This section should be executed after the Cisco Nexus Switch Configuration section in the main document is completed.

Create Infrastructure iSCSI VLANs on Cisco Nexus A and Cisco Nexus B

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

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

```
vlan <infra-iscsi-a-vlan-id>
name Infra-iSCSI-A-VLAN
vlan <infra-iscsi-b-vlan-id>
name Infra-iSCSI-B-VLAN
exit
```

Add Infrastructure iSCSI VLANs to Port-Channels on Cisco Nexus A and Cisco Nexus B

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

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

```
interface Po10,Po117,Po118,Po15,Po16
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
copy run start
```

NetApp Storage Configuration – Part 1

When using the iSCSI protocol for SAN boot connectivity, use the Storage Configuration steps outlined in the body of this document. Where appropriate, replace the Fibre Channel configuration steps with the steps listed here.



If the iSCSI license was not installed during the cluster configuration, make sure to install the license before creating the iSCSI service.

Create Block Protocol (iSCSI) Service



If the FCP protocol is not being used in the environment it should be removed from the vservers configured in a previous step. If FCP will be used in addition to iSCSI or in the future, step 1 can be omitted.

1. Remove FCP protocol from the vservers.

```
vserver remove-protocols -vserver Infra-SVM -protocols fcp
```

2. Enable the iSCSI protocol on the vservers.

```
vserver add-protocols -vserver Infra-SVM -protocols iscsi
```

3. Create the iSCSI block service.

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

Create iSCSI Broadcast Domains

To create the broadcast domains for each of the iSCSI VLANs, run the following commands:

```
network port broadcast-domain create -broadcast-domain Infra-iSCSI-A -mtu 9000
network port broadcast-domain create -broadcast-domain Infra-iSCSI-B -mtu 9000
```

Create iSCSI VLANs

To create iSCSI VLANs, follow these steps:

1. Modify the MTU size on the parent interface group hosting the iSCSI traffic using the following commands:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

2. Create VLAN ports for the iSCSI LIFs on each storage controller.

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

Add VLANs to iSCSI Broadcast Domains

To add each of the iSCSI VLAN ports to the corresponding broadcast domain, run the following commands:

```
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-
node01>:a0a-<infra-iscsi-a-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-
node01>:a0a-<infra-iscsi-b-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-
node02>:a0a-<infra-iscsi-a-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-
node02>:a0a-<infra-iscsi-b-vlan-id>

network port broadcast-domain show
```

Create iSCSI LIFs

To create four iSCSI LIFs, run the following commands (two on each node):

```

network interface create -vserver Infra-SVM -lif iscsi-lif-1a -role data -data-
protocol iscsi -home-node <st-node01> -home-port a0a-<infra-iscsi-a-vlan-id> -
address <st-node01-infra-iscsi-a-ip> -netmask <infra-iscsi-a-mask> -status-admin up

network interface create -vserver Infra-SVM -lif iscsi-lif-1b -role data -data-
protocol iscsi -home-node <st-node01> -home-port a0a-<infra-iscsi-b-vlan-id> -
address <st-node01-infra-iscsi-b-ip> -netmask <infra-iscsi-b-mask> -status-admin up

network interface create -vserver Infra-SVM -lif iscsi-lif-2a -role data -data-
protocol iscsi -home-node <st-node02> -home-port a0a-<infra-iscsi-a-vlan-id> -
address <st-node02-infra-iscsi-a-ip> -netmask <infra-iscsi-a-mask> -status-admin up

network interface create -vserver Infra-SVM -lif iscsi-lif-2b -role data -data-
protocol iscsi -home-node <st-node02> -home-port a0a-<infra-iscsi-b-vlan-id> -
address <st-node02-infra-iscsi-b-ip> -netmask <infra-iscsi-b-mask> -status-admin up

network interface show

```

Cisco UCS iSCSI Configuration

The following subsections can be completed to add infrastructure iSCSI to the Cisco UCS. These subsections can be completed in place of the subsections in the Cisco UCS Configuration section of this document labeled (FCP), or they can be completed in addition to the FCP sections to have the option of FCP or iSCSI boot.

Create IQN Pools for iSCSI Boot

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

1. In Cisco UCS Manager, click SAN.
2. Expand 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.2010-11.com.flexpod 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.

Create a Block of IQN Suffixes ? X

Suffix :

From :

Size :



14. Click OK.

15. Click Finish and OK to complete creating the IQN pool.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools for iSCSI boot for the Cisco UCS environment, follow these steps:



The IP Pools for iSCSI Boot are created here in the root organization, assuming that all UCS servers will be booted from the NetApp Infrastructure SVM. If servers will be booted from tenant SVMs with UCS tenant organizations, consider creating the IP Pools for iSCSI Boot in the tenant organization.

1. In Cisco UCS Manager, click LAN.
2. Expand 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 addresses.
10. In the From field, enter the beginning of the range to assign as iSCSI boot IP addresses on Fabric A.

11. Set the size to enough addresses to accommodate the servers.
12. Enter the appropriate Subnet Mask.
13. Click OK.
14. Click Next.
15. Click Finish and OK to complete creating the Fabric A iSCSI IP Pool.
16. Right-click IP Pools.
17. Select Create IP Pool.
18. Enter iSCSI-IP-Pool-B as the name of IP pool.
19. Optional: Enter a description for the IP pool.
20. Select Sequential for the assignment order.
21. Click Next.
22. Click Add to add a block of IP addresses.
23. In the From field, enter the beginning of the range to assign as iSCSI IP addresses on Fabric B.
24. Set the size to enough addresses to accommodate the servers.
25. Enter the appropriate Subnet Mask.
26. Click OK.
27. Click Next.
28. Click Finish and OK to complete creating the Fabric B iSCSI IP Pool.

Create iSCSI VLANs

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

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter Infra-iSCSI-A as the name of the VLAN to be used for iSCSI-A.
6. Keep the Common/Global option selected for the scope of the VLAN.

7. Enter the Infra-iSCSI-A VLAN ID.
8. Keep the Sharing Type as None.

Create VLANs



VLAN Name/Prefix :

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

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

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

VLAN IDs :

Sharing Type : None Primary Isolated Community

Check Overlap

OK

Cancel

9. Click OK and then click OK again.
10. Right-click VLANs.
11. Select Create VLANs.
12. Enter Infra-iSCSI-B as the name of the VLAN to be used for iSCSI-B.
13. Keep the Common/Global option selected for the scope of the VLAN.
14. Enter the iSCSI-B VLAN ID.
15. Keep the Sharing Type as None.
16. Click OK and then click OK again.

Create iSCSI vNIC Templates

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

1. Select LAN.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click vNIC Templates under the FlexPod Organization.
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 Infra-iSCSI-A.
11. Select Infra-iSCSI-A 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.

Create vNIC Template



VLANS | VLAN Groups

Advanced Filter | Export | Print | Settings

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	default	<input type="radio"/>	1
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>	113
<input checked="" type="checkbox"/>	Infra-iSCSI-A	<input checked="" type="radio"/>	3010
<input type="checkbox"/>	Infra-iSCSI-B	<input type="radio"/>	3020
<input type="checkbox"/>	Infra-NFS	<input type="radio"/>	3050
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>	2

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Dynamic vNIC usNIC VMQ

usNIC Connection Policy : ▼

OK **Cancel**

16. Click OK to complete creating the vNIC template.
17. Click OK.
18. Right-click vNIC Templates.
19. Select Create vNIC Template.
20. Enter iSCSI-Template-B as the vNIC template name.

21. Select Fabric B. Do not select the Enable Failover checkbox.
22. Leave Redundancy Type set at No Redundancy.
23. Under Target, make sure that only the Adapter checkbox is selected.
24. Select Updating Template for Template Type.
25. Under VLANs, select only Infra-iSCSI-B.
26. Select Infra-iSCSI-B as the native VLAN.
27. Leave vNIC Name set for the CDN Source.
28. Under MTU, enter 9000.
29. From the MAC Pool list, select MAC-Pool-B.
30. From the Network Control Policy list, select Enable-CDP-LLDP.
31. Click OK to complete creating the vNIC template.
32. Click OK.

Create LAN Connectivity Policy for iSCSI Boot

To configure the necessary Infrastructure LAN Connectivity Policy within the FlexPod Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click LAN Connectivity Policies under the FlexPod Organization.
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-vSwitch0-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select vSwitch0-A.
10. In the Adapter Policy list, select VMWare.
11. Click OK to add this vNIC to the policy.
12. Click the upper Add button to add another vNIC to the policy.

13. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select vSwitch0-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-vDS0-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select vDS0-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 another vNIC to the policy.
25. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select vDS0-B.
28. In the Adapter Policy list, select VMWare.
29. Click OK to add the vNIC to the policy.
30. Click the upper Add button to add a vNIC.
31. In the Create vNIC dialog box, enter 04-iSCSI-A as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select iSCSI-Template-A.
34. In the Adapter Policy list, select VMWare.
35. Click OK to add this vNIC to the policy.
36. Click Add to add a vNIC to the policy.
37. In the Create vNIC dialog box, enter 05-iSCSI-B as the name of the vNIC.
38. Select the Use vNIC Template checkbox.

39. In the vNIC Template list, select iSCSI-Template-B.
40. In the Adapter Policy list, select VMWare.
41. Click OK to add this vNIC to the policy.
42. Expand Add iSCSI vNICs.
43. Select Add in the Add iSCSI vNICs section.
44. Set the name to iSCSI-Boot-A.
45. Select 04-iSCSI-A as the Overlay vNIC.
46. Set the iSCSI Adapter Policy to default.
47. Leave the VLAN set to Infra-iSCSI-A (native).
48. Leave the MAC Address set to None.
49. Click OK.
50. Select Add in the Add iSCSI vNICs section.
51. Set the name to iSCSI-Boot-B.
52. Select 05-iSCSI-B as the Overlay vNIC.
53. Set the iSCSI Adapter Policy to default.
54. Leave the VLAN set to Infra-iSCSI-B (native).
55. Leave the MAC Address set to None.

Create LAN Connectivity Policy



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 05-iSCSI-B	Derived	
vNIC 04-iSCSI-A	Derived	
vNIC 03-vDS0-B	Derived	
vNIC 02-vDS0-A	Derived	
vNIC 01-vSwitch0-B	Derived	
vNIC 00-vSwitch0-A	Derived	

Delete Add Modify

Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-B	05-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-Boot-A	04-iSCSI-A	default	Derived

Add Delete Modify

OK

Cancel

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

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.



It is not recommended to use UEFI Secure Boot for iSCSI Boot at this time. The Legacy boot mode will be used.

To create a boot policy for the Cisco UCS environment within the FlexPod Organization, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click Boot Policies under the FlexPod Organization.
4. Select Create Boot Policy.
5. Enter Boot-iSCSI-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Select the Legacy Boot Mode.
9. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
10. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
11. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.
12. Click OK.
13. Select Add iSCSI Boot.
14. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-B.
15. Click OK.
16. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

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.

Local Devices

CIMC Mounted vMedia

[Add CIMC Mounted CD/DVD](#)

[Add CIMC Mounted HDD](#)

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	O...	vNIC/vHBA/iSCSI vN...	Type	LU...	WWN	Slot...	Boo...	Boo...	Des...
Remote CD/DVD	1								
iSCSI	2								
iSCSI		iSCSI-Boot-A	Pri...						
iSCSI		iSCSI-Boot-B	Sec...						
CIMC Mounted CD/DVD	3								

OK

Cancel

17. Click OK then click OK again to create the policy.

Create iSCSI Boot Service Profile Template

In this procedure, one service profile template for Infrastructure ESXi hosts within the FlexPod Organization is created for Fabric A boot.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > FlexPod Organization.
3. Right-click the FlexPod Organization.
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.
7. Under UUID Assignment, select UUID_Pool.

Create Service Profile Template ? X

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/org-NX-VM-FlexPod**

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

8. Click Next.

Configure Storage Provisioning

To configure the storage provisioning, follow these steps:

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

Configure Networking Options

To configure the network options, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.

3. Select iSCSI-Boot from the LAN Connectivity Policy drop-down list.
4. 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

To configure the storage options, follow these steps:

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

Configure Zoning Options

To configure the zoning options, follow this step:

1. Make no changes and click Next.

Configure vNIC/HBA Placement

To configure the vNIC/HBA placement, follow these steps:

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

2. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

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

Configure Server Boot Order

To configure the server boot orders, follow these steps:

1. Select Boot-iSCSI-A for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **Boot-iSCSI-A**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 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/vHB...	Type	LUN Name	WWN	Slot Num...	Boot Name	Boot Path	Description
Remot...	1								
iSCSI	2								
CIMC ...	3								

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

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

2. In the Boot order, expand iSCSI and select iSCSI-Boot-A.
3. Click Set iSCSI Boot Parameters.
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, log into the storage cluster management interface and run the “iscsi show” command”.
10. Enter the IP address of iscsi-lif-2a for the IPv4 Address field.



The LIFs are being entered in reverse order because the second one will be used when booting the server.

Create iSCSI Static Target



iSCSI Target Name	:	<input type="text" value="iqn.1992-08.com.netapp:"/>	
Priority	:	<input type="text" value="1"/>	
Port	:	<input type="text" value="3260"/>	
Authentication Profile	:	<input type="text" value="<not set>"/>	Create iSCSI Authentication Profile
IPv4 Address	:	<input type="text" value="192.168.10.62"/>	
LUN ID	:	<input type="text" value="0"/>	

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-1a for the IPv4 Address field.
15. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters



[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(30/32) ▼

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 Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.10.62	0
iqn.1992-08....	2	3260		192.168.10.61	0

[+](#) Add [-](#) Delete [i](#) Info

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

OK

Cancel

16. Click OK to complete setting the iSCSI Boot Parameters.

17. In the Boot order, select iSCSI-Boot-B.

18. Click Set iSCSI Boot Parameters.
19. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.
20. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
21. Set iSCSI-IP-Pool-B as the “Initiator IP address Policy”.
22. Select the iSCSI Static Target Interface option.
23. Click Add.
24. 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”.
25. Enter the IP address of iscsi-lif-2b for the IPv4 Address field.
26. Click OK to add the iSCSI static target.
27. Click Add.
28. Enter the iSCSI Target Name.
29. Enter the IP address of iscsi-lif-1b for the IPv4 Address field.
30. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters



[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(30/32) ▼

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 Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.20.62	0
iqn.1992-08....	2	3260		192.168.20.61	0

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

OK

Cancel

31. Click OK to complete setting the iSCSI Boot Parameters.

32. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

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
Storage Config. Deployment Policy	: User Ack
Reboot Policy	: User Ack

< Prev Next > **Finish** Cancel

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these 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-B200-M5” for the Server Pool Qualification to select only B200 M5 servers in the pool.
4. Expand Firmware Management at the bottom of the page and select the default policy.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: [Create Server Pool](#)

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

Up Down

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

Server Pool Qualification : [Create Host Firmware Package](#)

Restrict Migration :

Firmware Management (BIOS, Disk Controller, Adapter)

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

Host Firmware Package: [Create Host Firmware Package](#)

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

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these 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.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

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 :

⊕ External IPMI/Redfish 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 : [Create Power Control Policy](#)

⊕ Scrub Policy

⊕ KVM Management Policy

⊕ Graphics Card 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, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organizations > FlexPod Organization > 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 and click OK then click OK again to create the clone.
5. Select the newly-created VM-Host-Infra-iSCSI-A-vM and select the vMedia Policy tab.
6. Click Modify vMedia Policy.
7. Select the ESXi-6.7U2-HTTP vMedia Policy and click OK.
8. Click OK to confirm.

Create Service Profiles

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

1. Connect to Cisco UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template VM-Host-Infra-iSCSI-A-VM.
3. Right-click VM-Host-Infra-iSCSI-A-VM 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.”

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :



7. Click OK to create the service profiles.
8. Click OK in the confirmation message.

When VMware ESXi 6.7 U2 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.

NetApp Storage Configuration – Part 2

Create igroups

It is assumed that boot LUNs have already been created for the two ESXi management hosts.

Table 12 iSCSI IQN for SVM

SVM Name	SVM Target IQN
Infra-SVM	

Table 13 iSCSI vNIC IQN Configuration

Cisco UCS Service Profile Name	iSCSI IQN	Variable
vm-host-infra-01		<vm-host-infra-01-iqn>
vmhost-infra-02		<vm-host-infra-02-iqn>
vmhost-infra-03		<vm-host-infra-03-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 top right. The “Initiator Name” is displayed at the top of the page under the “Service Profile Initiator Name.”

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

```
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-01 -protocol iscsi -
ostype vmware -initiator <vm-host-infra-01-iqn>

lun igroup create -vserver Infra-SVM -igroup vm-host-infra-02 -protocol iscsi -
ostype vmware -initiator <vm-host-infra-02-iqn>

lun igroup create -vserver Infra-SVM -igroup vm-host-infra-03 -protocol iscsi -
ostype vmware -initiator <vm-host-infra-03-iqn>
```



Use the values listed in Table 12 and Table 13 for the IQN information.

To view the two igroups just created, use the command `lun igroup show`.

```
lun igroup show -protocol iscsi
```

Map Boot LUNs to igroups

From the storage cluster management LIF SSH connection, run the following commands:

```
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/vm-host-infra-01 -igroup
vm-host-infra-01 -lun-id 0

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/vm-host-infra-02 -igroup
vm-host-infra-02 -lun-id 0

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/vm-host-infra-03 -igroup
vm-host-infra-03 -lun-id 0
```


VMware vSphere Configuration


Set Up VMkernel Ports and Virtual Switch on ESXi Host VM-Host-Infra-01





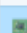

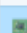

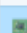
To add the iSCSI networking configuration on the first ESXi host, follow the steps at the end of section [Set Up VMkernel Ports and Virtual Switch](#). In this section, a single iSCSI Boot vSwitch is configured with two uplinks, one to UCS fabric A and the other to fabric B. The first VMkernel port will be mapped only to the fabric A uplink and the second one will be mapped to the fabric B uplink.

To setup Vmkernel ports and virtual switches on ESXi hosts on VM-Host-Infra-01, follow these steps:

1. From the Host Client, select Networking.
2. In the center pane, select the Virtual switches tab.
3. Highlight the iScsiBootvSwitch line.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Select Add uplink to add an uplink to iScsiBootvSwitch.
7. Use the pulldown to select vmnic5 for Uplink 2.
8. Expand NIC teaming, select vmnic5, and select Mark standby.

 **Edit standard virtual switch - iScsiBootvSwitch**

 Add uplink


MTU	<input type="text" value="9000"/>									
Uplink 1	vmnic4 - Up, 20000 mbps ✕									
Uplink 2	vmnic5 - Up, 20000 mbps ✕									
▶ Link discovery	Click to expand									
▶ Security	Click to expand									
▼ NIC teaming										
Load balancing	Route based on originating port ID ▼									
Network failover detection	Link status only ▼									
Notify switches	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failback	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<div style="display: flex; align-items: center; gap: 10px;">  Mark active  Move up  Move down </div> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr> <th style="width: 30%;">Name</th> <th style="width: 40%;">Speed</th> <th style="width: 30%;">Status</th> </tr> </thead> <tbody> <tr> <td> vmnic4</td> <td>20000 Mbps, full duplex</td> <td>Active</td> </tr> <tr style="background-color: #d9ead3;"> <td> vmnic5</td> <td>20000 Mbps, full duplex</td> <td>Standby</td> </tr> </tbody> </table>	Name	Speed	Status	 vmnic4	20000 Mbps, full duplex	Active	 vmnic5	20000 Mbps, full duplex	Standby
Name	Speed	Status								
 vmnic4	20000 Mbps, full duplex	Active								
 vmnic5	20000 Mbps, full duplex	Standby								
▶ Traffic shaping	Click to expand									











9. Click Save.
10. Select the VMkernel NICs tab.
11. Select the vmk1 iScsiBootPG line. Select Edit Settings to edit the properties of this VMkernel port.
12. Change the MTU to 9000.
13. Expand IPv4 Settings and enter a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.



It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

14. Click Save to save the changes to the VMkernel port.
15. Select the Port groups tab.
16. Select the iScsiBootPG line. Select Edit Settings to edit the properties of this port group.
17. Expand NIC teaming and select Yes to the right of Override failover order.
18. Select vmnic5 and select Mark unused.

 Edit port group - iScsiBootPG


Name	<input type="text" value="iScsiBootPG"/>									
VLAN ID	<input type="text" value="0"/>									
Virtual switch	<input type="text" value="iScsiBootvSwitch"/>									
▶ Security	Click to expand									
▼ NIC teaming										
Load balancing	<input type="text" value="Inherit from vSwitch"/>									
Network failover detection	<input type="text" value="Inherit from vSwitch"/>									
Notify switches	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Failback	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Override failover order	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<div style="display: flex; gap: 10px;">  Mark active  Mark unused  Move up  Move down </div> <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td> vmnic4</td> <td>20000 Mbps, full duplex</td> <td>Active</td> </tr> <tr> <td> vmnic5</td> <td>20000 Mbps, full duplex</td> <td>Unused</td> </tr> </tbody> </table>	Name	Speed	Status	 vmnic4	20000 Mbps, full duplex	Active	 vmnic5	20000 Mbps, full duplex	Unused
Name	Speed	Status								
 vmnic4	20000 Mbps, full duplex	Active								
 vmnic5	20000 Mbps, full duplex	Unused								
▶ Traffic shaping	Click to expand									


19. Click Save to complete the changes to the iScsiBootPG.











20. At the top, select the Virtual switches tab.

21. Select the iScsiBootvSwitch line and click Edit settings.

22. Expand NIC teaming and select vmnic5. Select Mark active to make vmnic5 active within the vSwitch.

 Edit standard virtual switch - iScsiBootvSwitch

 Add uplink

MTU	<input type="text" value="9000"/>									
Uplink 1	<input type="text" value="vmnic4 - Up, 20000 mbps"/> 									
Uplink 2	<input type="text" value="vmnic5 - Up, 20000 mbps"/> 									
▶ Link discovery	Click to expand									
▶ Security	Click to expand									
▼ NIC teaming										
Load balancing	<input type="text" value="Route based on originating port ID"/>									
Network failover detection	<input type="text" value="Link status only"/>									
Notify switches	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failback	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<input checked="" type="checkbox"/> Mark standby <input type="checkbox"/>  Move up <input type="checkbox"/>  Move down <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td> vmnic4</td> <td>20000 Mbps, full duplex</td> <td>Active</td> </tr> <tr> <td> vmnic5</td> <td>20000 Mbps, full duplex</td> <td>Active</td> </tr> </tbody> </table>	Name	Speed	Status	 vmnic4	20000 Mbps, full duplex	Active	 vmnic5	20000 Mbps, full duplex	Active
Name	Speed	Status								
 vmnic4	20000 Mbps, full duplex	Active								
 vmnic5	20000 Mbps, full duplex	Active								
▶ Traffic shaping	Click to expand									

23. Click Save to save the changes to the vSwitch.

24. At the top, select the VMkernel NICs tab.

25. Click Add VMkernel NIC.

26. For New port group, enter iScsiBootPG-B

27. For Virtual switch, select iScsiBootvSwitch.

28. Leave the VLAN ID set at 0.
29. Change the MTU to 9000.
30. Select Static IPv4 settings and expand IPv4 settings.
31. Enter a unique IP address and netmask in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B.



It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

32. Do not select any of the Services.

Add VMkernel NIC

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	iScsiBootvSwitch
VLAN ID	0
MTU	9000
IP version	IPv4 only
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	192.168.20.193
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication


33. Click Create.











34. Select the Port groups tab.

35. Select the iScsiBootPG-B line. Select Edit Settings to edit the properties of this port group.

36. Expand NIC teaming and select Yes to the right of Override failover order.

37. To the right of Failover order, select vmnic4 and select Mark unused.

 Edit port group - iScsiBootPG-B

Name	<input type="text" value="iScsiBootPG-B"/>									
VLAN ID	<input type="text" value="0"/>									
Virtual switch	<input type="text" value="iScsiBootvSwitch"/>									
▶ Security	Click to expand									
▼ NIC teaming										
Load balancing	<input type="text" value="Inherit from vSwitch"/>									
Network failover detection	<input type="text" value="Inherit from vSwitch"/>									
Notify switches	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Failback	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Inherit from vSwitch									
Override failover order	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<div style="display: flex; gap: 10px;">  Mark active  Mark unused  Move up  Move down </div> <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td> vmnic4</td> <td>20000 Mbps, full duplex</td> <td>Unused</td> </tr> <tr> <td> vmnic5</td> <td>20000 Mbps, full duplex</td> <td>Active</td> </tr> </tbody> </table>	Name	Speed	Status	 vmnic4	20000 Mbps, full duplex	Unused	 vmnic5	20000 Mbps, full duplex	Active
Name	Speed	Status								
 vmnic4	20000 Mbps, full duplex	Unused								
 vmnic5	20000 Mbps, full duplex	Active								
▶ Traffic shaping	Click to expand									

38. Click Save to save the changes to the port group.
39. On the left select Storage, then in the center pane select the Adapters tab.
40. Click Software iSCSI to configure software iSCSI for the host.
41. In the Configure iSCSI window, under Dynamic targets, select Add dynamic target.
42. Select to add address and enter the IP address of iscsi-lif-1a from storage SVM Infra-SVM. Press Return.
43. Repeat this process to add the IP addresses for iscsi-lif-2a, iscsi-lif-1b, and iscsi-lif-2b.

44. Click Save configuration.
45. Click Software iSCSI to configure software iSCSI for the host.
46. Verify that four static targets and four dynamic targets are listed for the host.

Configure iSCSI - vmhba64

iSCSI enabled	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled															
Name & alias	iqn.2010-11.com.flexpod:ucs-host:5															
CHAP authentication	Do not use CHAP															
Mutual CHAP authentication	Do not use CHAP															
Advanced settings	Click to expand															
Network port bindings	No port bindings															
Static targets	<div style="display: flex; justify-content: space-between; align-items: center;"> Add static target Remove static target Edit settings <input type="text" value="Search"/> </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Target</th> <th>Address</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td>iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...</td> <td>192.168.10.62</td> <td>3260</td> </tr> <tr> <td>iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...</td> <td>192.168.10.61</td> <td>3260</td> </tr> <tr> <td>iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...</td> <td>192.168.20.62</td> <td>3260</td> </tr> <tr> <td>iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...</td> <td>192.168.20.61</td> <td>3260</td> </tr> </tbody> </table>	Target	Address	Port	iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.10.62	3260	iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.10.61	3260	iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.20.62	3260	iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.20.61	3260
Target	Address	Port														
iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.10.62	3260														
iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.10.61	3260														
iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.20.62	3260														
iqn.1992-08.com.netapp:sn.13e45747cb5711e99f1ed039ea0...	192.168.20.61	3260														
Dynamic targets	<div style="display: flex; justify-content: space-between; align-items: center;"> Add dynamic target Remove dynamic target Edit settings <input type="text" value="Search"/> </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Address</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td>192.168.10.61</td> <td>3260</td> </tr> <tr> <td>192.168.10.62</td> <td>3260</td> </tr> <tr> <td>192.168.20.61</td> <td>3260</td> </tr> <tr> <td>192.168.20.62</td> <td>3260</td> </tr> </tbody> </table>	Address	Port	192.168.10.61	3260	192.168.10.62	3260	192.168.20.61	3260	192.168.20.62	3260					
Address	Port															
192.168.10.61	3260															
192.168.10.62	3260															
192.168.20.61	3260															
192.168.20.62	3260															

47. Click Cancel to close the window.
48. Add the host to vCenter, set the swap datastore on the host, and add the host to the VMware vDS, including migrating the vMotion VMkernel port to the vDS.
49. Using NetApp VSC, set recommended values for the host.



If the host shows an alarm stating that connectivity with the boot disk was lost, place the host in Maintenance Mode and reboot the host.

Add iSCSI Configuration to a VMware ESXi Host Added in vCenter

This section details the steps to add iSCSI configuration to an ESXi host added and configured in vCenter. This section assumes the host has been added to vCenter and the basic networking completed, NFS datastores set up, and the time configuration and swap files added.

To add an iSCSI configuration to an ESXi host, follow these steps:

1. In the vSphere HTML5 Client, under Hosts and Clusters, select the ESXi host.
2. In the center pane, select Configure. In the list under Networking, select Virtual switches.
3. In the center pane, expand iScsiBootvSwitch. Click EDIT to edit settings for the vSwitch.
4. Change the MTU to 9000 and click OK.
5. Click MANAGE PHYSICAL ADAPTERS to add a network adapter to the vSwitch.
6. Click the green plus sign to add an adapter.
7. Select vmnic5.

Add Physical Adapters to the Switch ✕

Network Adapters

- vmnic2
- vmnic3
- vmnic5

All	Properties	CDP	LLDP
Adapter Name		Cisco Systems Inc Cisco VIC Ethernet NIC vmnic5	
Location		PCI 0000:62:00.5	
Driver		nenic	
Status			
Status		Connected	
Actual speed, Duplex		20000 Mb, Full Duplex	
Configured speed, Duplex		20000 Mb, Full Duplex	
Networks		No networks	
Network I/O Control			
Status		Allowed	
SR-IOV			
Status		Not supported	
Cisco Discovery Protocol			
Version		2	
Timeout		60	
Time to live		124	
Samples		17	
Device ID		AA13-6454-B.flexpod.cisco.com(FDO22191DNN)	

CANCEL
OK

8. Click OK. Use the blue down arrow to move vmnic5 under Unused adapters.

Manage Physical Network Adapters | iScsiBootvSwitch

Assigned adapters

+ | × | ↑ | ↓

Active adapters

- vmnic4

Standby adapters

Unused adapters

- vmnic5

All	Properties	CDP	LLDP
Adapter Name		Cisco Systems Inc Cisco VIC Ethernet NIC	
Location		vmnic5	
Driver		PCI 0000:62:00.5	
Driver		nenic	
Status			
Status		Connected	
Actual speed, Duplex		20000 Mb, Full Duplex	
Configured speed, Duplex		20000 Mb, Full Duplex	
Networks		No networks	
Network I/O Control			
Status		Allowed	
SR-IOV			
Status		Not supported	
Cisco Discovery Protocol			

9. Click OK to complete adding the vmnic to the vSwitch.
10. Select VMkernel adapters.
11. Select the iScsiBootPG VMkernel adapter. Select EDIT to edit adapter settings.
12. Change the MTU to 9000.
13. Click IPv4 settings on the left. Change the IP address to a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.



It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

14. Click OK.
15. Select Virtual switches. In the center pane, expand iScsiBootvSwitch.
16. Select ... to the right of iScsiBootPG in the center of the window. Select Edit Settings.
17. In the Edit Settings window, select Teaming and failover. In the center of the window under Failover order, select Override.

iScsiBootPG - Edit Settings

Properties	Load balancing	<input type="checkbox"/> Override	Route based on originating virtual port	▼
Security	Network failure detection	<input type="checkbox"/> Override	Link status only	▼
Traffic shaping	Notify switches	<input type="checkbox"/> Override	Yes	▼
Teaming and failover	Failback	<input type="checkbox"/> Override	Yes	▼

Failover order

Override

↑ ↓

Active adapters
vmnic4
Standby adapters
Unused adapters
vmnic5

Select a physical network adapter from the list to view its details.

Select active and standby adapters. During a failover, standby adapters activate in the order specified above.

18. Click OK to complete editing the port group.

19. In the list located on the left under Networking, select Virtual switches.

20. In the center pane, expand iScsiBootvSwitch. Click EDIT to edit settings for the vSwitch.

21. Select Teaming and failover. Under Failover order, select vmnic5. Use the blue up arrow icon to move vmnic5 to the Active adapters list.

iScsiBootvSwitch - Edit Settings

Properties

Security

Traffic shaping

Teaming and failover

Load balancing: Route based on originating virtual port

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active adapters: vmnic4, vmnic5

Standby adapters:

Unused adapters:

Adapter Name: Cisco Systems Inc Cisco VIC Ethernet NI
vmnic5
Location: PCI 0000:62:00.5
Driver: nenic

Status

Status: Connected
Actual speed, Duplex: 20000 Mb, Full Duplex
Configured speed, Duplex: 20000 Mb, Full Duplex
Networks: No networks

SR-IOV

Status: Not supported

Select active and standby adapters. During a failover, standby adapters activate in the order specified above.

CANCEL OK

22. Click OK to complete the changes to the vSwitch.
23. In the list on the left under Networking, select VMkernel adapters.
24. In the center pane, select Add Networking to add a VMkernel adapter.
25. Make sure VMkernel Network Adapter is selected and click NEXT.
26. Select an existing standard switch and click BROWSE.
27. Select iScsiBootvSwitch and click OK.

na-esxi-1.flexpod.cisco.com - Add Networking

1 Select connection type
 2 Select target device
 3 Port properties
 4 IPv4 settings
 5 Ready to complete

Select target device
Select a target device for the new connection.

Select an existing network
 BROWSE ...

Select an existing standard switch
 BROWSE ...


New standard switch
 MTU (Bytes)

CANCEL

28. Click NEXT.

29. Enter iScsiBootPG-B for the Network label, leave VLAN ID set to None (0), select Custom - 9000 for MTU, and click NEXT.

30. Select Use static IPv4 settings. Enter a unique IP address and netmask in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B.

 **It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.**

31. Click NEXT.

32. Review the settings and click FINISH to complete creating the VMkernel port.

33. In the list under Networking, select Virtual switches. Expand iScsiBootvSwitch.

34. In the center of the window, select ... to the right of iScsiBootPG-B. Select Edit Settings.

35. Select Teaming and failover. In the center of the window under Failover order, select Override.

36. Select vmnic4. Use the blue down arrow to move vmnic4 to Unused adapters.

iScsiBootPG-B - Edit Settings

Properties

Security

Traffic shaping

Teaming and failover

Load balancing Override Route based on originating virtual port ▾

Network failure detection Override Link status only ▾

Notify switches Override Yes ▾

Failback Override Yes ▾

Failover order

Override

↑ ↓

Active adapters	Standby adapters	Unused adapters
vmnic5		vmnic4

Select active and standby adapters. During a failover, standby adapters activate in the order specified above.

All	Properties	CDP	LLDP
Adapter			
Name	Cisco Systems Inc Cisco VIC Ethernet NI		
Location	vmnic4		
Driver	PCI 0000:62:00.4		
Status			
Status	Connected		
Actual speed, Duplex	20000 Mb, Full Duplex		
Configured speed, Duplex	20000 Mb, Full Duplex		
Networks	No networks		

CANCEL **OK**

37. Click OK to complete the changes to the port group.

38. In the list under Storage, select Storage Adapters.

39. Select the iSCSI Software Adapter and below, select the Dynamic Discovery tab.

40. Click Add.

41. Enter the IP address of the storage controller's Infra-SVM LIF iscsi-lif-1a and click OK.

42. Repeat this process to add the IPs for iscsi-lif-2a, iscsi-lif-1b, and iscsi-lif-2b.

43. Under Storage Adapters, click Rescan Adapter to rescan the iSCSI Software Adapter.

44. Under Static Discovery. Four static targets should be listed.

45. Using NetApp VSC, set recommended values for the host.


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. This procedure should be run at the end of the vCenter deployment section. To setup the ESXi Dump Collector, follow these steps:

1. Log into the vSphere Web Client as administrator@vsphere.local and select Home.



This procedure cannot be done from the vSphere HTML5 Client, it must be done from the Flash-based Web Client.

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
esxcli system coredump network get
```

Create a FlexPod ESXi Custom ISO using VMware vCenter

In this validation document, the [Cisco Custom Image for ESXi 6.7 U2 GA Install CD](#) ISO was used to install VMware ESXi. After this installation the Cisco VIC nfnic and nenic drivers had to be updated and the NetApp NFS Plug-in for VMware VAAI had to be installed during the FlexPod deployment. vCenter 6.7 U2 can be used to produce a FlexPod custom ISO containing the updated VIC drivers and the NetApp NFS Plug-in for VMware VAAI. This ISO can be used to install VMware ESXi 6.7 U2 without having to do any additional driver updates. This ISO can be produced by following these steps:

1. Download the [Cisco Custom Image for ESXi 6.7 U2 Offline Bundle](#). This file (VMware_ESXi_6.7.0_13006603_Custom_Cisco_6.7.2.1_Bundle.zip) can be used to produce the Cisco Custom Image for ESXi 6.7 U2 GA Install CD ISO.
2. Download the following driver images and extract the listed .zip file:

- [VMware ESXi 6.7 nfnic 4.0.0.38 FC Driver for Cisco nfnic](#) - VMW-ESX-6.7.0-nfnic-4.0.0.38-offline_bundle-14003014.zip
- [VMware ESXi 6.7 nenic 1.0.29.0 NIC Driver for Cisco nenic](#) - VMW-ESX-6.7.0-nenic-1.0.29.0-offline_bundle-12897497.zip



It is not necessary to download the nenic 1.0.29.0 driver because it is included in the Cisco Custom Image for ESXi 6.7 U2 Offline Bundle. It is being left here for future updates.

- [NetApp NFS Plug-in for VMware VAAI 1.1.2](#) - NetAppNasPlugin.v23.zip
3. Connect to the VMware Appliance Management interface (<https://<vcenter-IP>:5480>) and login as root.
 4. In the menu on the left, select Services.
 5. In the center pane, select ImageBuilder Service and at the top, click START.
 6. Verify that the ImageBuilder Service now is Healthy and Started. Log out of the VMware Appliance Management interface.
 7. Log into the VMware vCenter HTML5 Client as administrator@vsphere.local.
 8. Under Menu, select Auto Deploy.
 9. In the center pane, select Software Depots.
 10. Click the IMPORT to upload a software depot.
 11. Name the depot ESXi-6.7U2-Custom-Cisco. Click BROWSE. Browse to the local location of the zip file, highlight it, and click Open.

Import Software Depot



Name *

File *

BROWSE

CANCEL

UPLOAD

12. Click UPLOAD to upload the software depot.
13. Repeat steps 9–12 to add software depots for nfnic-4.0.0.38 and NetAppNasPlugin-v23.
14. Click NEW to add a custom software depot.
15. Select Custom depot and name the custom depot FlexPod-ESXi-6.7U2.

Add Software Depot



Online depot

Name:

URL:

Custom depot

Name: *

CANCEL

ADD

16. Click ADD to add the custom software depot.
17. From the drop-down list, select the ESXi-6.7U2- Custom-Cisco software depot. Make sure the Image Profiles tab is selected and then use the radio button to select the VMware-ESXi-6.7.0-13006603-Custom-Cisco-6.7.2.1 image profile. Click CLONE to clone the image profile.
18. Name the clone FlexPod-ESXi-6.7U2. For Vendor, enter Cisco-NetApp. For Description, enter “Cisco Custom ISO for ESXi 6.7U2 GA with nfnic-4.0.0.38, nenic-1.0.29.0, and NetAppNasPlugin-v23”. Select FlexPod-ESXi-6.7U2 for Software depot.

Clone Image Profile

- 1 Name and details
- 2 Select software packages
- 3 Ready to complete

Name and details

Name * FlexPod-ESXI-6.7U2

Vendor * Cisco-NetApp

Description

Cisco Custom ISO for ESXi 6.7U2 GA with nfnc-4.0.0.38, nenic-1.0.29.0, and NetAppNasPlugin-v23

Software depot * FlexPod-ESXI-6.7U2 ⓘ

CANCEL NEXT

19. Click NEXT.

20. Under Available software packages, check NetAppNasPlugin 1.1.2-3, make sure nenic 1.0.29.0 is checked, uncheck nfnc 4.0.0.35, and check nfnc 4.0.0.38. Leave the remaining selections unchanged.

Clone Image Profile

- 1 Name and details
- 2 Select software packages
- 3 Ready to complete

Select software packages ✕

Acceptance level Partner supported ▼

<input type="checkbox"/>	Name ▼	Version ▼	Acceptance Level ▼	Vendor ▼	Depot ▼
<input checked="" type="checkbox"/>	nenic	1.0.29.0-1OEM.670.0.0...	VMware certified	Cisco	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-bnx2	2.2.4f.v60.10-2vmw.67...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-bnx2x	1.78.80.v60.12-2vmw.67...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-cdc-ether	1.0-3vmw.670.0.0.8169...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-cnic	1.78.76.v60.13-2vmw.67...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-e1000	8.0.3.1-5vmw.670.0.0.8...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-e1000e	3.2.2.1-2vmw.670.0.0.81...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-fcoe	1.0.29.9.3-7vmw.670.0.0...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-forcedeth	0.61-2vmw.670.0.0.816...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-i40e	2.0.7-1OEM.600.0.0.24...	VMware certified	INT	ESXI-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-igb	5.0.5.1.1-5vmw.670.0.0...	VMware certified	VMW	ESXI-6.7U2-Custom-Cisco

147 selected of 148 items 1 2 >

CANCEL

BACK

NEXT

Clone Image Profile

- 1 Name and details
- 2 Select software packages
- 3 Ready to complete

Select software packages ✕

Acceptance level Partner supported ▼

<input type="checkbox"/>	Name	Version	Acceptance Level	Vendor	Depot
<input checked="" type="checkbox"/>	net-tg3	3.131d.v60.4-2vmw.670...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-usbnet	1.0-3vmw.670.0.0.8169...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	net-vmxnet3	1.1.3.0-3vmw.670.2.48.1...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	NetAppNasPlugin	1.1.2-3	VMware accepted	NetApp	NetAppNasPlugin-v23
<input checked="" type="checkbox"/>	nfnic	4.0.0.38-1OEM.670.0.0...	VMware certified	Cisco	nfnic-4.0.0.38
<input type="checkbox"/>	nfnic	4.0.0.35-1OEM.670.0.0...	VMware certified	Cisco	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	nhpsa	2.0.22-3vmw.670.1.28.1...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx4-core	3.17.13.1-1vmw.670.2.48...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx4-en	3.17.13.1-1vmw.670.2.48...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx4-rdma	3.17.13.1-1vmw.670.2.48...	VMware certified	VMW	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx5-core	4.17.14.2-1OEM.670.0.0...	VMware certified	MEL	ESXi-6.7U2-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx5-rdma	4.17.14.2-1OEM.670.0.0...	VMware certified	MEL	ESXi-6.7U2-Custom-Cisco

147 selected of 148 items 1 2 >

CANCEL
BACK
NEXT

21. Click NEXT.

22. Click FINISH.

23. Using the Software Depot pulldown, select the FlexPod-ESXi-6.7U2 custom software depot. Under Image Profiles select the FlexPod-ESXi-6.7U2 image profile. Click ... > Export to export an image profile. ISO should be selected. Click OK to generate a bootable ESXi installable image.

24. Once the Image profile export completes, click DOWNLOAD to download the ISO.

25. Optionally, generate the ZIP archive to generate an offline bundle for the FlexPod image.

FlexPod Backups

Cisco UCS Backup

Automated backup of the UCS domain is important for recovery of the UCS Domain from issues ranging 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.

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 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 configure the backup, using the Cisco UCS Manager GUI, follow these steps:

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>



Admin State must be Enabled to fill in the Remote File field.

- f. Admin State: <select Enable to activate the schedule on save, Disable to disable schedule on Save>
- g. Schedule: [Daily/Weekly/Bi Weekly]

All

General Policy Backup & Export

Protocol : FTP TFTP SCP SFTP

User :

Password :

Remote File :

Admin State : Disable Enable

Schedule : Daily Weekly Bi Weekly

Max Files : **0**

Description :

All Configuration Backup Policy

Hostname :

Protocol : FTP TFTP SCP SFTP

User :

Password :

Remote File :

Admin State : Disable Enable

Schedule : Daily Weekly Bi Weekly

Max Files : **0**

Description :

Backup/Export Config Reminder

Admin State : Disable Enable

Remind me after(Days) :

4. Click Save Changes to create the Policy.

Cisco Nexus and MDS Backups

The configuration of the Cisco Nexus 9000 and Cisco MDS 9132T 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 9336C-FX2 switches is shown below:

```
conf t
feature scheduler
scheduler logfile size 1024
scheduler job name backup-cfg
copy running-config tftp://<server-ip>/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
exit
scheduler schedule name daily
job name backup-cfg
time daily 2:00
end
```



On the Cisco MDS 9132T, remove “vrf management” from the copy command.

Show the job that has been setup:

```
sh scheduler job
Job Name: backup-cfg
-----
copy running-config tftp://10.1.156.150/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf
management

=====

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)
=====
```

The documentation for the feature scheduler can be found here:

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 provide a scheduled solution. To create a backup, follow these steps:

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

The screenshot shows the VMware Appliance Management console. The top navigation bar includes 'vm Appliance Management', the date and time 'Fri 10-04-2019 07:29 PM UTC', and user information 'English', 'Help', 'Actions', and 'root'. A left sidebar contains a menu with items: Summary, Monitor, Access, Networking, Firewall, Time, Services, Update, Administration, Syslog, and Backup. The main content area displays system details: Hostname (na-vc.flexpod.cisco.com), Type (vCenter Server with an embedded Platform Services Controller), Product (VMware vCenter Server Appliance), Version (6.7.0.31000), and Build number (13643870). Below this is a 'Health Status' section with a table showing 'Overall Health' as 'Good (Last checked Oct 4, 2019, 03:28:53 PM)' and individual components (CPU, Memory, Database, Storage, Swap) all as 'Good'. To the right is a 'Single Sign-On' section with a table showing 'Domain' as 'vsphere.local' and 'Status' as 'Running'.

2. Click Time to set the vCenter time zone.
3. In the center pane, to the right of Timezone, click EDIT.
4. From the drop-down list select the current time zone and click SAVE.

Edit Time Zone

Time zone America/New_York ▼

5. Click Backup in the list to open up the Backup Appliance Dialogue.
6. To the right of Backup Schedule, click CONFIGURE.
7. Specify:
 - a. The Backup location with the protocol to use [HTTPS/HTTP/SCP/FTPS/FTP]
 - b. The User name and password.

c. The Number of backups to retain.

Create Backup Schedule

Backup location ⓘ scp://nx-ftp.flexpod.cisco.com/var/www/html/software/Configs/na-vc/

Backup server credentials

User name admin

Password

Schedule ⓘ Daily 02 : 15 A.M. America/New_York

Encrypt backup (optional)

Encryption Password _____

Confirm Password _____

Number of backups to retain

Retain all backups

Retain last 7 backups

Data

Stats, Events, and Tasks 61 MB

Inventory and configuration 778 MB

Total size (compressed) 839 MB

[CANCEL](#) [CREATE](#)

8. Click Create.

Backup Schedule		EDIT DISABLE DELETE
▼ Status	Enabled	
Schedule	Daily , 2:15 A.M. America/New_York	
Backup Location	scp://nx-ftp.flexpod.cisco.com/var/www/html/software/Configs/na-vc/	
Backup data	<ul style="list-style-type: none"> Stats, Events, and Tasks Inventory and configuration 	
Number of backups to retain	7	

Activity [BACKUP NOW](#)

9. The Backup Schedule should now show a Status of Enabled.

10. Restoration can be initiated with the backed-up files using the Restore function of the VCSA 6.7 Installer.

About the Authors

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

John has been involved in designing, developing, validating, and supporting the FlexPod Converged Infrastructure since it was developed almost nine years ago. Before his roles 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.

Sree Lakshmi Lanka, Solutions Engineer, Infrastructure and Cloud Engineering, NetApp

Sree is a Solutions Engineer in the Converged Infrastructure Engineering team at NetApp. She has been with NetApp since 2007, serving in a variety of Member Technical Staff- SAN Interoperability Engineering roles. Sree has over 15 years of experience in the IT industry specializing in data management, SAN and security.

Acknowledgements

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

- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.
- Scott Kovacs, Technical Marketing Engineer, NetApp