

# FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7

Deployment Guide for FlexPod Datacenter with Cisco UCS 4<sup>th</sup> Generation, NetApp ONTAP 9.7, NetApp Active IQ, and VMware vSphere 6.7 U3

Published: May 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 Port Profiles .....	21
Add Individual Port Descriptions for Troubleshooting and Enable UDLD for UCS Interfaces.....	22
Create Port Channels.....	23
Configure Port Channel Parameters .....	24
Configure Virtual Port Channels .....	24
Uplink into Existing Network Infrastructure .....	25
Switch Testing Commands .....	26
Storage Configuration .....	27
NetApp All Flash FAS A800 Controllers.....	27
NetApp Hardware Universe .....	27
Controllers .....	27
Disk Shelves .....	27
NetApp ONTAP 9.7 .....	28
Complete Configuration Worksheet .....	28
Configure ONTAP Nodes.....	28
Set Up Node.....	31
Log into the Cluster.....	41
Verify Storage Failover .....	41

Set Auto-Revert on Cluster Management.....	42
Zero All Spare Disks.....	42
Set Up Service Processor Network Interface .....	43
Create Auto-provisioned Aggregates .....	43
Create Manual Provisioned Aggregates (Optional) .....	44
Remove Ports from Default Broadcast Domain .....	45
Disable Flow Control on 100GbE ports.....	45
Disable Auto-Negotiate on 100GbE Ports .....	45
Disable Auto-Negotiate on Fibre Channel Ports .....	46
Enable Cisco Discovery Protocol .....	46
Enable Link-layer Discovery Protocol on all Ethernet Ports.....	46
Create Management Broadcast Domain .....	46
Create NFS Broadcast Domain .....	47
Create Interface Groups .....	47
Change MTU on Interface Groups .....	47
Create VLANs .....	47
Configure Network Time Protocol .....	48
Configure Simple Network Management Protocol .....	48
Configure SNMPv3 Access.....	48
Create SVM .....	49
Create Load-Sharing Mirrors of SVM Root Volume .....	49
Create Block Protocol (FC) Service .....	50
Configure HTTPS Access.....	50
Configure NFSv3.....	51
Create FlexVol Volumes.....	52
Create Boot LUNs.....	52
Modify Volume Efficiency .....	52
Create FC LIFs.....	53
Create NFS LIFs.....	53
Add Infrastructure SVM Administrator .....	53
Configure and Test AutoSupport.....	54
Cisco UCS Configuration .....	55
Cisco UCS Base Configuration .....	55
Perform Initial Setup of Cisco UCS 6454 Fabric Interconnects for FlexPod Environments.....	55
Cisco UCS Setup.....	56
Log into Cisco UCS Manager.....	56
Anonymous Reporting.....	57
Upgrade Cisco UCS Manager Software to Version 4.1(1c).....	57
Configure Cisco UCS Call Home .....	57
Synchronize Cisco UCS to NTP .....	58
Add Additional DNS Server(s).....	59



Add an Additional Administrative User.....	60
Configure Unified Ports (FCP).....	61
Edit Chassis Discovery Policy.....	63
Enable Port Auto-Discovery Policy .....	63
Enable Server and Uplink Ports.....	64
Enable Info Policy for Neighbor Discovery.....	65
Acknowledge Cisco UCS Chassis and FEX .....	65
Create an Organization .....	66
Create a WWNN Pool for FC Boot (FCP) .....	67
Create WWPN Pools (FCP) .....	69
Create VSANs (FCP).....	72
Enable FC Uplink VSAN Trunking (FCP).....	73
Create FC Uplink Port Channels (FCP) .....	74
Disable Unused FC Uplink Ports (FCP) .....	77
Create vHBA Templates (FCP) .....	77
Create SAN Connectivity Policy (FCP).....	79
Add Block of IP Addresses for KVM Access .....	81
Create Uplink Port Channels to Cisco Nexus Switches .....	82
Add UDLD to Uplink Port Channels.....	85
Set Jumbo Frames in Cisco UCS Fabric.....	87
Create VLANs .....	88
Create MAC Address Pools .....	91
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) .....	93
Create vNIC Templates.....	94
Create High Traffic VMware Adapter Policy .....	98
Create LAN Connectivity Policy for FC Boot (FCP).....	100
Create Server Pool .....	104
Create UUID Suffix Pool.....	105
Modify Default Host Firmware Package .....	105
Create Local Disk Configuration Policy (Optional) .....	106
Create Power Control Policy.....	108
Create Server Pool Qualification Policy (Optional).....	109
Update the Default Maintenance Policy .....	109
Create 100 Percent App Direct Mode Persistent Memory Policy .....	110
Create vMedia Policy for VMware ESXi 6.7U3 ISO Install Boot.....	111
Create Server BIOS Policy .....	113
Create FC Boot Policy (FCP) .....	116
Create Service Profile Template (FCP) .....	121
Configure Storage Provisioning.....	122
Configure Networking .....	122
Configure SAN Connectivity.....	123

Configure Zoning .....	124
Configure vNIC/HBA Placement.....	124
Configure vMedia Policy.....	124
Configure Server Boot Order .....	125
Configure Maintenance Policy.....	125
Configure Server Assignment .....	126
Configure Operational Policies .....	127
Create vMedia-Enabled Service Profile Template .....	128
Create Service Profile Template for Servers with Optane Memory.....	129
Create vMedia-Enabled Service Profile Template for Servers with Optane Memory.....	129
Create Service Profiles .....	129
Add More Servers to FlexPod Unit.....	130
Gather Necessary Information.....	130
SAN Switch Configuration.....	132
Physical Connectivity .....	132
FlexPod Cisco MDS Base .....	132
Cisco MDS 9132T A.....	132
FlexPod Cisco MDS Switch Configuration .....	135
Enable Licenses.....	135
Add Second NTP Server.....	135
Configure Individual Ports.....	136
Create VSANs.....	137
Create Device Aliases .....	138
Create Zones and Zoneset.....	139
Storage Configuration - Boot LUNs .....	140
ONTAP Boot Storage Setup .....	140
Create igroups .....	140
Map Boot LUNs to igroups.....	140
VMware vSphere 6.7U3 Setup.....	141
VMware ESXi 6.7U3 .....	141
Download ESXi 6.7U3 from VMware.....	141
Log into Cisco UCS 6454 Fabric Interconnect .....	141
Set Up VMware ESXi Installation .....	142
Install ESXi.....	142
Set Up Management Networking for ESXi Hosts .....	143
Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional).....	145
Install VMware Drivers for the Cisco Virtual Interface Card (VIC) and ESXi Host.....	146
Log into the First VMware ESXi Host by Using VMware Host Client .....	146
Set Up VMkernel Ports and Virtual Switch .....	147
Mount Required Datastores.....	149
Configure NTP on First ESXi Host .....	150

Configure ESXi Host Swap .....	151
VMware vCenter 6.7U3 .....	152
Build the VMware vCenter Server Appliance .....	152
Adjust vCenter CPU Settings .....	163
Setup VMware vCenter Server .....	164
Add AD User Authentication to vCenter (Optional) .....	170
FlexPod VMware vSphere Distributed Switch (vDS) .....	171
Configure the VMware vDS in vCenter .....	172
Add and Configure a VMware ESXi Host in vCenter .....	176
Add the ESXi Host to vCenter .....	176
Set Up VMkernel Ports and Virtual Switch .....	177
Mount Required Datastores .....	180
Configure NTP on ESXi Host .....	180
Configure ESXi Host Swap .....	182
Check ESXi Host Fibre Channel Pathing .....	182
Add the ESXi Host(s) to the VMware Virtual Distributed Switch .....	183
Add the vMotion VMkernel Port to the ESXi Host .....	184
VMware ESXi 6.7U3 TPM Attestation .....	186
FlexPod Management Tools Setup .....	187
NetApp Virtual Storage Console 9.7 Deployment Procedure .....	187
Virtual Storage Console 9.7 Pre-installation Considerations .....	187
The requirements for deploying VSC are listed here. ....	187
Install Virtual Storage Console 9.7 .....	187
Download the NetApp NFS VAAI Plug-In .....	192
Install the NetApp NFS VAAI Plug-In .....	192
Verify the VASA Provider .....	193
Discover and Add Storage Resources .....	194
Optimal Storage Settings for ESXi Hosts .....	196
Virtual Storage Console 9.7 Provisioning Datastores .....	197
Create the Storage Capability Profile .....	198
Create a VM Storage Policy .....	200
Create Virtual Machine with Assigned VM Storage Policy .....	206
NetApp SnapCenter 4.3 .....	207
NetApp SnapCenter Architecture .....	208
Install SnapCenter Plug-In for VMware vSphere 4.3 .....	208
Host and Privilege Requirements for the SnapCenter Plug-in for VMware vSphere .....	208
License requirements for SnapCenter Plug-in for VMware vSphere .....	209
Download and Deploy the SnapCenter Plug-in for VMware vSphere 4.3 .....	209
SnapCenter Plug-in for VMware vSphere in vCenter Server .....	215
Add Storage Systems (SVM) .....	216
View Virtual Machine Backups from vCenter by Using SnapCenter Plug-In .....	226

Create On-Demand Backup .....	228
Restore from vCenter by Using SnapCenter Plug-In .....	229
Active IQ Unified Manager 9.7 .....	233
Configure Active IQ Unified Manager .....	240
Add Local Users to Active IQ Unified Manager .....	244
Configure Remote Authentication .....	245
Add a Remote User to Active IQ Unified Manager .....	248
Add the vCenter Server to Active IQ Unified Manager .....	250
View Virtual Machine Inventory .....	252
Review Security Compliance with Active IQ Unified Manager .....	254
Remediate Security Compliance Findings .....	256
NetApp Active IQ .....	257
Add a Watchlist to the Discovery Dashboard .....	258
Create Active IQ Digital Advisor Dashboard .....	260
Cisco Data Center Network Manager (DCNM)-SAN .....	262
Prerequisites .....	262
Deploying the Cisco DCNM-SAN OVA .....	263
Configuring DCNM-SAN .....	271
Configure SAN Insights in DCNM SAN .....	274
Cisco Intersight .....	276
Sample Tenant Provisioning .....	285
Provision a Sample Application Tenant .....	285
Appendix .....	287
FlexPod iSCSI Addition .....	287
Cisco Nexus Switch Configuration .....	287
NetApp Storage Configuration – Part 1 .....	287
Cisco UCS iSCSI Configuration .....	289
Create Service Profile Template for Servers with Optane Memory .....	311
Create vMedia-Enabled Service Profile Template for Servers with Optane Memory .....	311
NetApp Storage Configuration – Part 2 .....	312
VMware vSphere Configuration .....	313
ESXi Dump Collector Setup for iSCSI-Booted Hosts .....	328
Create a FlexPod ESXi Custom ISO using VMware vCenter .....	328
FlexPod Backups .....	334
Cisco UCS Backup .....	334
Cisco Nexus and MDS Backups .....	336
VMware VCSA Backup .....	336
About the Authors .....	339
Acknowledgements .....	339



## 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.

Hybrid cloud adoption is accelerating and FlexPod is at the center of on premises infrastructure. As business applications move into the cloud, management applications must also follow suit where practical. In this updated design, FlexPod is introducing SaaS based management with Cisco Intersight and NetApp Active IQ. These platforms offer AI powered analytics for infrastructure management and operational intelligence.

This document describes the Cisco and NetApp® FlexPod Datacenter with NetApp ONTAP 9.7 on NetApp AFF A800 storage, Cisco UCS Manager unified software release 4.1(1) with 2<sup>nd</sup> Generation Intel Xeon Scalable Processors and VMware vSphere 6.7 U3. Cisco UCS Manager (UCSM) 4.1(1) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 6400, 2200/2300/2400 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. Also included are Cisco Intersight and NetApp Active IQ SaaS management platforms. FlexPod Datacenter with NetApp ONTAP 9.7, Cisco UCS unified software release 4.1(1), and VMware vSphere 6.7 U3 is a predesigned, best-practice datacenter 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® 9.7 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 4.1(1) unified software release, Cisco UCS B200-M5 and C220-M5 servers with 2<sup>nd</sup> Generation Intel Xeon Scalable Processors, and Cisco 1400 Series Virtual Interface Cards (VICs)
- Support for the latest Cisco UCS 6454 and 64108 (supported but not validated) Fabric Interconnects
- Support for the latest Cisco UCS 2408 Fabric Extender
- Addition of Cisco Data Center Network Manager (DCNM)-SAN Version 11.3(1)
- Addition of Cisco Intersight Software as a Service (SaaS) Management
- Support for the NetApp AFF A800 Storage Controller
- Support for the latest release of NetApp ONTAP® 9.7
- Support for NetApp Virtual Storage Console (VSC) 9.7
- Support for NetApp SnapCenter Plug-in for VMware vSphere 4.3
- Support for NetApp Active IQ Unified Manager 9.7
- Addition of NetApp Active IQ Data Optimization and Proactive care
- Fibre channel, NFS, iSCSI (appendix) storage design
- Validation of VMware vSphere 6.7 U3



- Unified Extensible Firmware Interface (UEFI) Secure Boot of VMware ESXi 6.7 U3
- Trusted Platform Module (TPM) 2.0 Attestation of UEFI Secure Boot of VMware ESXi 6.7 U3
- 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 A800 storage array. This infrastructure option expanded with Cisco MDS switches sitting between the Cisco UCS Fabric Interconnect and the NetApp AFF A800 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 A800

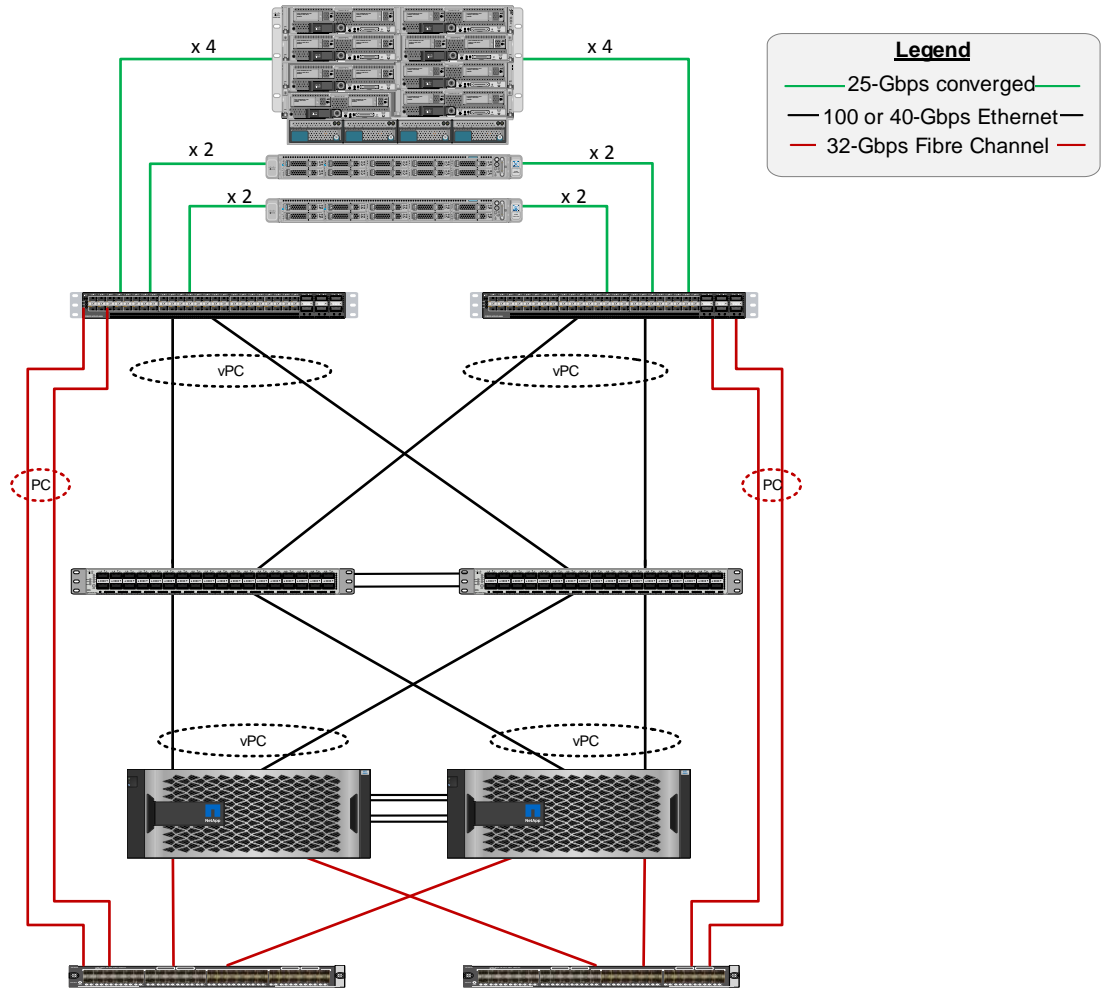
### Cisco Unified Computing System

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-A800

### 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 800 (HA pair) running ONTAP 9.7 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 <sup>nd</sup> Generation Intel Xeon Scalable Processors	4.1(1b)	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	9.3(4)	
	Cisco MDS 9132T	8.4(1a)	
Storage	NetApp AFF 800	ONTAP 9.7P1	
Software	Cisco UCS Manager	4.1(1c)	
	Cisco Data Center Network Manager (SAN)	11.3(1)	
	VMware vSphere	6.7U3	
	VMware ESXi nfnic FC Driver	4.0.0.52	
	VMware ESXi nenic Ethernet Driver	1.0.31.0	
	NetApp Virtual Storage Console (VSC) / VASA Provider Appliance	9.7	
	NetApp NFS Plug-in for VMware VAAI	1.1.2-3	
	NetApp SnapCenter for vSphere	4.3	Includes the vSphere plug-in for SnapCenter
	NetApp Active IQ Unified Manager	9.7	
Management	Cisco Intersight	N/A	
	NetApp Active IQ	N/A	

## 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	101
FCoE-B	VLAN for FCoE encapsulation of VSAN-B	102
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 for vSphere		
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 800 running NetApp ONTAP® 9.7.



**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.**

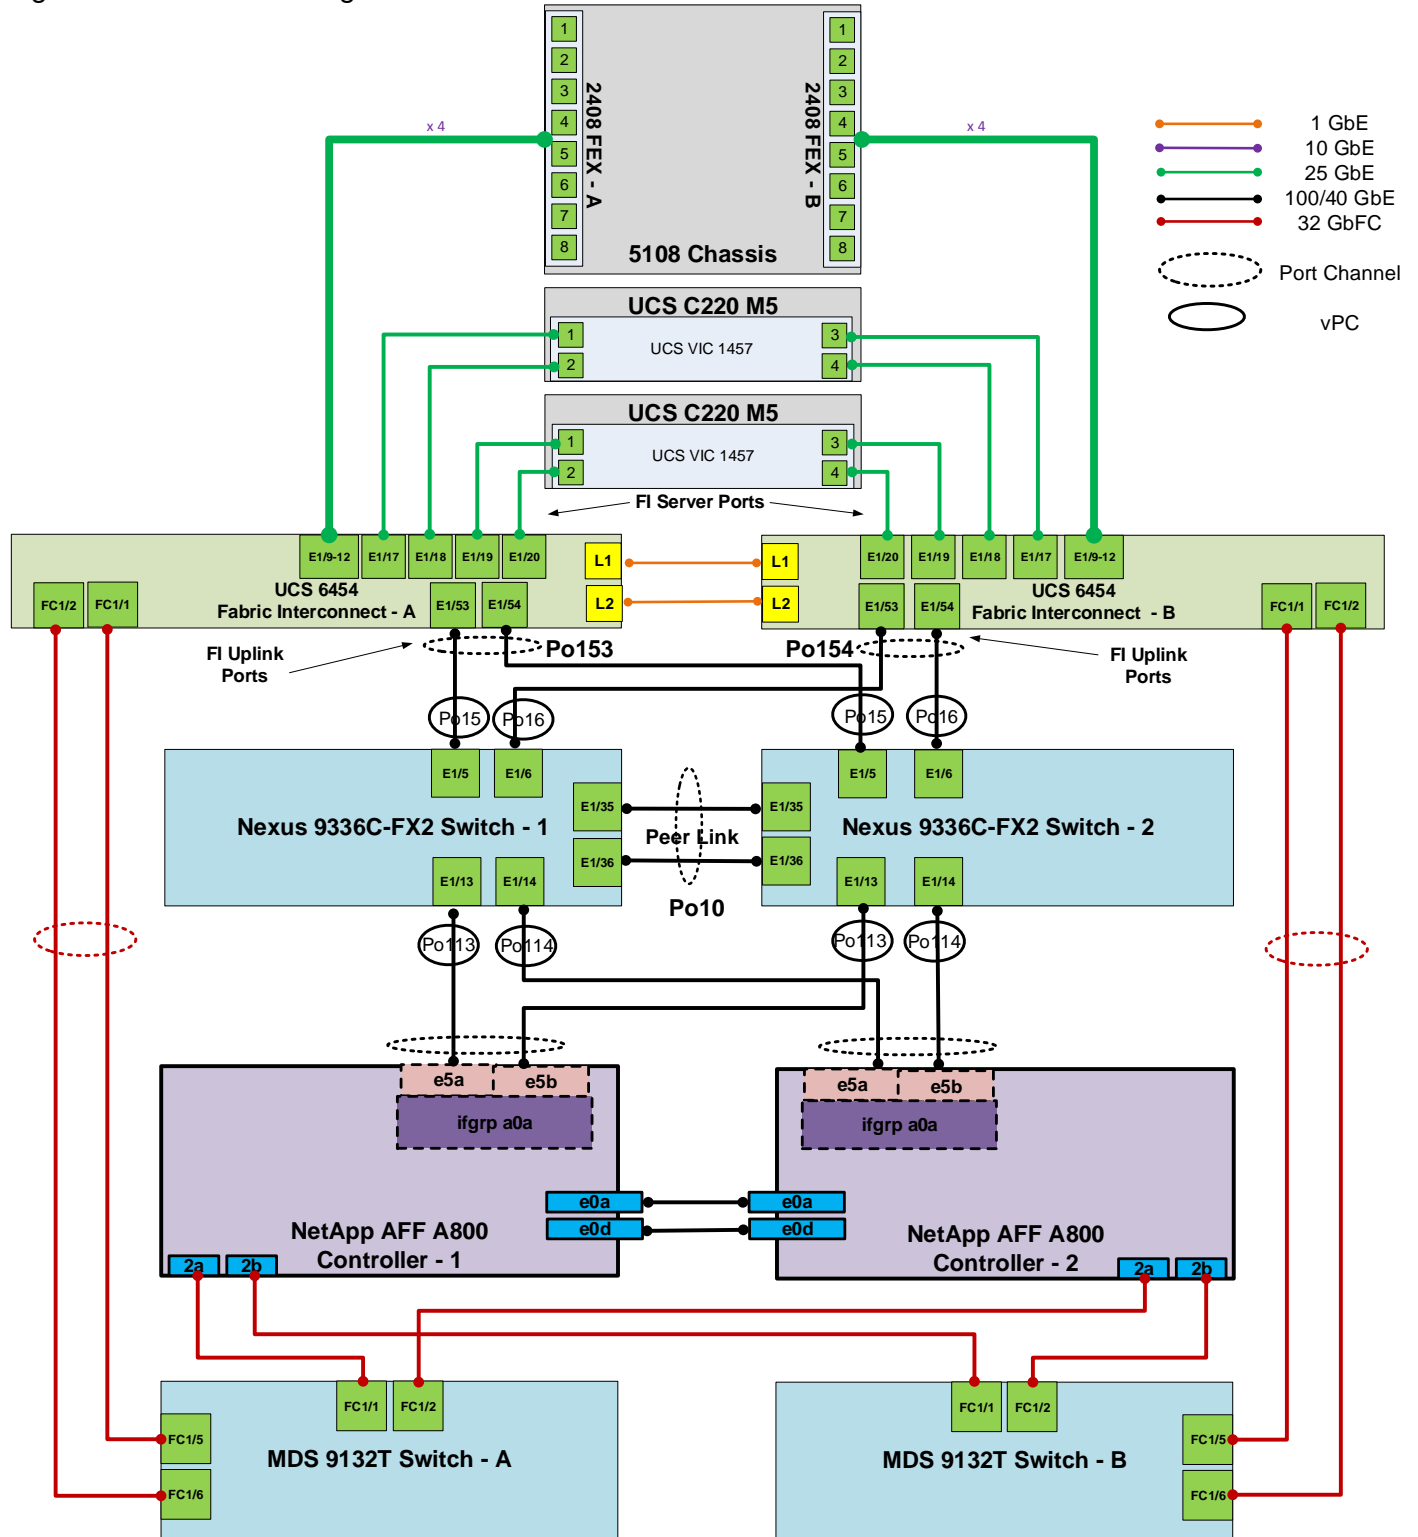
---

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 a connection 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 section [FlexPod Cabling](#).

### FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 9.3(4), 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.



In this validation, port speed and duplex are hard set at both ends of every 100GE 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 udld
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

2. 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 Port Profiles

This version of the FlexPod solution uses port profiles for virtual port channel (vPC) connections to NetApp Storage, Cisco UCS, and the vPC peer link.

### Cisco Nexus A and Cisco Nexus B

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

```
port-profile type port-channel FP-ONTAP-Storage
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
state enabled

port-profile type port-channel FP-UCS
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
state enabled

port-profile type port-channel vPC-Peer-Link
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
state enabled

```

## 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 these steps:



**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 Eth1/13
description <st-clustername>-1:e5a
interface Eth1/14
description <st-clustername>-2:e5a

```



```
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 Eth1/13
description <st-clustername>-1:e5b
interface Eth1/14
description <st-clustername>-2:e5b
interface Eth1/35
description <nexus-a-hostname>:1/35
interface Eth1/36
description <nexus-a-hostname>:1/36
exit
```

## 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 Po113
description <st-clustername>-1
interface Eth1/13
channel-group 113 mode active
no shutdown
interface Po114
description <st-clustername>-2
interface Eth1/14
```

```

channel-group 114 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
inherit port-profile vPC-Peer-Link

interface Po113
inherit port-profile FP-ONTAP-Storage
interface Po114
inherit port-profile FP-ONTAP-Storage

interface Po15
inherit port-profile FP-UCS

interface Po16
inherit port-profile FP-UCS

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 Po113
vpc 117
interface Po114
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 Po113
vpc 117
interface Po114
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 A800 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 here: <http://docs.netapp.com/platstor/topic/com.netapp.nav.a800/home.html> on 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 A800 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.

When using NVMe drive shelves with NetApp storage controllers, refer to the [NS224 Drive Shelves](#) documentation for installation and servicing guidelines.

## NetApp ONTAP 9.7

### 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 9.7 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 9.7 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.7 is the version being booted, choose option 8 and `y` to reboot the node. Then continue with step 14.**

4. To install new software, choose option 7 from the menu.



5. Enter `y` to continue the installation.
6. Choose `e0M` for the network port you want to use for the download.
7. Enter `n` to skip the reboot.
8. Choose option 7 from the menu: `Install new software first`
9. Enter `y` to continue the installation
10. Enter the IP address, netmask, and default gateway for `e0M`.

```
Enter the IP address for port e0M: <node01-mgmt-ip>
Enter the netmask for port e0M: <node01-mgmt-mask>
Enter the IP address of the default gateway: <node01-mgmt-gateway>
```

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



**The web server must be pingable from node 01.**

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

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

```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y
Please answer yes or no

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} yes
```



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.



During the ONTAP installation a prompt to reboot the node requests a Y/N response. The prompt requires the entire Yes or No response to reboot the node and continue the installation.

15. Press `Ctrl-C` when the following message displays:

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

16. Choose option 4 for Clean Configuration and Initialize All Disks.
17. Enter `y` to zero disks, reset config, and install a new file system.
18. Enter `yes` 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 configuration of node 02 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-B 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 9.7 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.7 is the version being booted, choose option 8 and `y` to reboot the node. Then continue with step 14.**

---

4. To install new software, choose option 7.
5. Enter `y` to continue the installation.
6. Choose `e0M` for the network port you want to use for the download.
7. Enter `n` to skip the reboot.
8. Choose option 7: `Install new software first`
9. Enter `y` to continue the installation
10. Enter the IP address, netmask, and default gateway for `e0M`.

```
Enter the IP address for port e0M: <node02-mgmt-ip>
Enter the netmask for port e0M: <node02-mgmt-mask>
Enter the IP address of the default gateway: <node02-mgmt-gateway>
```

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



The web server must be pingable from node 02.

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

12. Press **Enter** for the user name, indicating no user name.

13. Enter **y** to set the newly installed software as the default to be used for subsequent reboots.

14. Enter **yes** to reboot the node.

```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y

Please answer yes or no

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} yes
```



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.



During the ONTAP installation a prompt to reboot the node requests a Y/N response. The prompt requires the entire **Yes** or **No** response to reboot the node and continue the installation.

15. Press **Ctrl-C** when you see this message:

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

16. Choose option 4 for Clean Configuration and Initialize All Disks.

17. Enter **y** to zero disks, reset config, and install a new file system.

18. Enter **yes** 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 9.7 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 accessing <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>
Cluster Admin SVM	<cluster-adm-svm>
Infrastructure Data SVM	<infra-data-svm>
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 node 01 netmask	<node01-mgmt-mask>

Cluster Detail	Cluster Detail Value
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>
SNMPv3 User	<snmp-v3-usr>
SNMPv3 Authentication Protocol	<snmp-v3-auth-PROTO>
SNMPv3 Privacy Protocol	<snmpv3-priv-PROTO>



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

- Complete the required information on the Initialize Storage System screen:

The screenshot shows the ONTAP System Manager interface for initializing a storage system. The page is titled "ONTAP 9.7" and includes a "Health" section on the left indicating "2 healthy nodes were found" for a cluster named "AFF-A800". The main "Initialize Storage System" section contains the following fields:

- STORAGE SYSTEM NAME:** "Enter cluster name"
- ADMINISTRATIVE PASSWORD:** "Enter new password" and "Confirm password"
- Networking:**
  - CLUSTER IP ADDRESS:** "IP Address"
  - SUBNET MASK:** "Length"
  - GATEWAY:** "IP Address"
  - NODE SERIAL NUMBERS:** "86" and "87"
  - NODE IP ADDRESSES:** "192.168.156.61" and "IP Address"
- Use Domain Name Service (DNS)

4. In the Cluster screen, follow these steps:
  - a. Enter the cluster name and administrator password.
  - b. Complete the Networking information for the cluster and each node.
  - c. Choose the box Use time services (NTP) and enter the IP addresses of the time servers in a comma separated list.

The screenshot shows the ONTAP System Manager web interface. The top navigation bar includes the ONTAP logo, the text "ONTAP System Manager (Return to classic version)", and navigation icons. The main content area is titled "ONTAP 9.7 Pro tips for initializing a storage system".

**Health**

2 healthy nodes were found.

AFF-A800

**Initialize Storage System**

STORAGE SYSTEM NAME  
aa14-a800  
You will see this name when managing the storage system.

ADMINISTRATIVE PASSWORD  
\*\*\*\*\*  
\*\*\*\*\*

**Networking**

CLUSTER IP ADDRESS	SUBNET MASK	GATEWAY
192.168.156.60	24	192.168.156.1

NODE SERIAL NUMBERS	NODE IP ADDRESSES
86	192.168.156.61
87	192.168.156.62



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

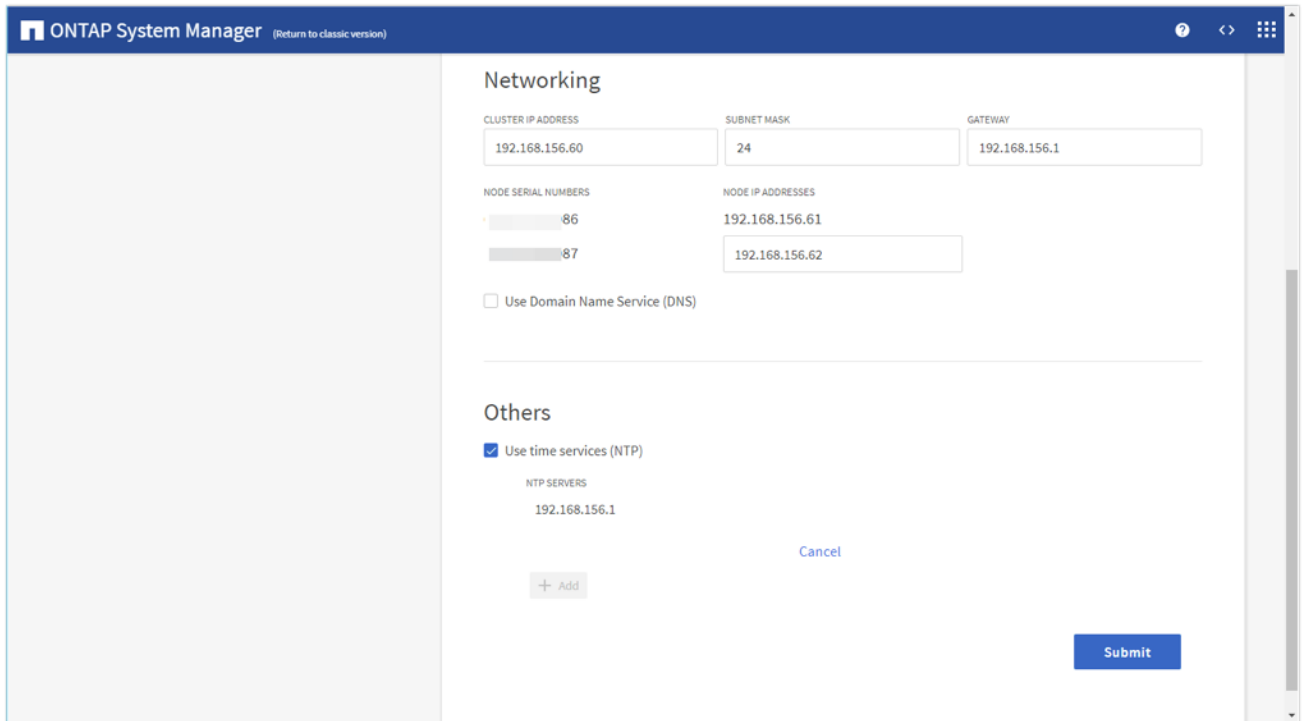


If all the nodes are not discovered, then configure the cluster using the command line.

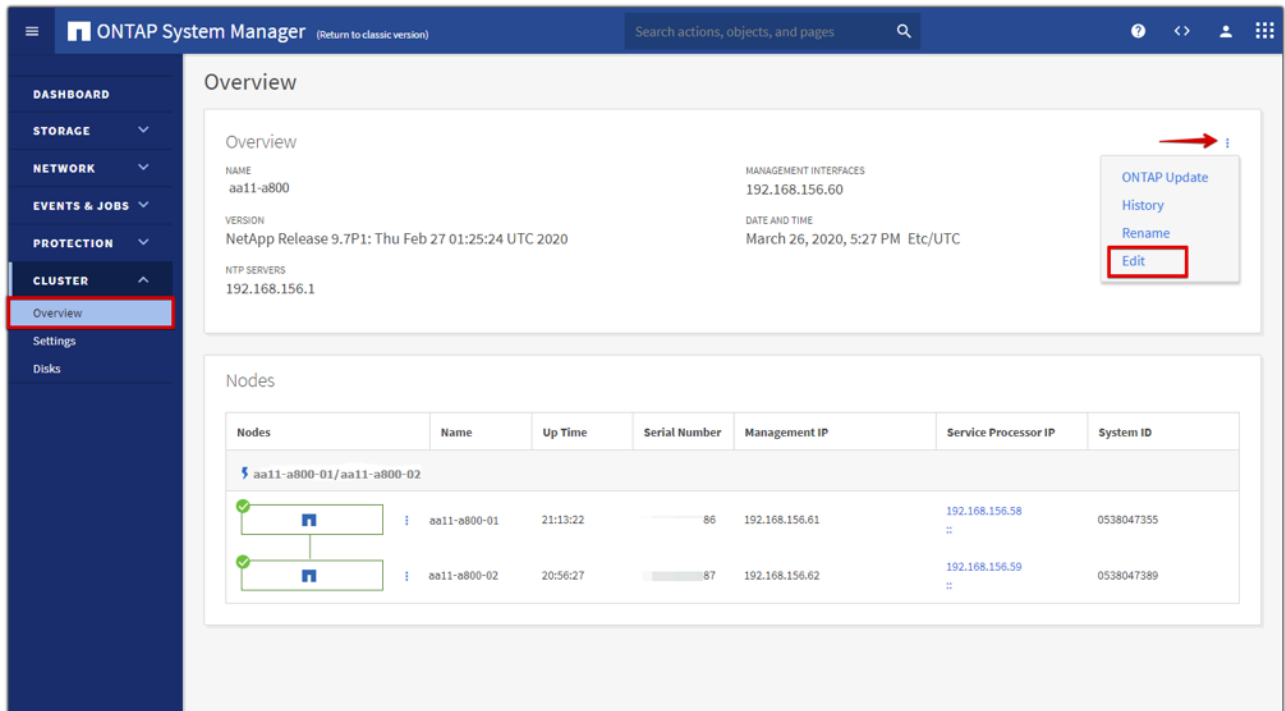


The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

5. Click Submit.



6. A few minutes will pass while the cluster is configured. When prompted, login to ONTAP System Manager to continue the cluster configuration.
7. From the Dashboard click the Cluster menu on the left and choose Overview.
8. Click the Details ellipsis button in the Overview pane at the top right of the screen and choose Edit.





9. Add additional cluster configuration details and click Save to make the changes persistent.

- a. Cluster location
- b. DNS domain name
- c. DNS server IP addresses



DNS server IP addresses can be added individually or with a comma separated list on a single line.

The screenshot shows the 'Edit Cluster Details' form in the ONTAP System Manager. The form is titled 'Edit Cluster Details' and has a close button (X) in the top right corner. The form contains several sections:

- NAME:** A text input field containing 'aa11-a800'.
- LOCATION:** A text input field containing 'RTP Lab'.
- DNS DOMAINS:** A text input field containing 'flexpod-ad.cisco.com'. Below it is a '+ Add' button.
- NAME SERVERS:** A list of text input fields. The first contains '10.1.156.250' and the second contains '10.1.156.251'. Below the list is a '+ Add' button.
- NTP SERVERS:** A text input field containing '192.168.156.1'. Below it is a '+ Add' button.
- Additional options:** A checkbox labeled 'Add cluster management interface' which is currently unchecked.
- Buttons:** At the bottom of the form are two buttons: 'Save' (in a blue box) and 'Cancel'.

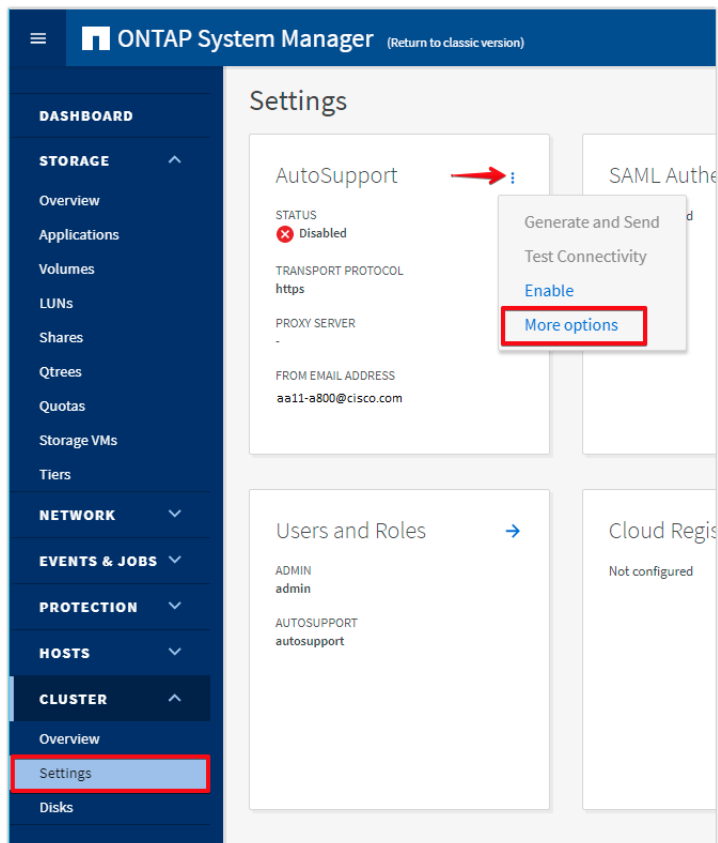
10. Click Save at the bottom of the page to make the changes persistent.



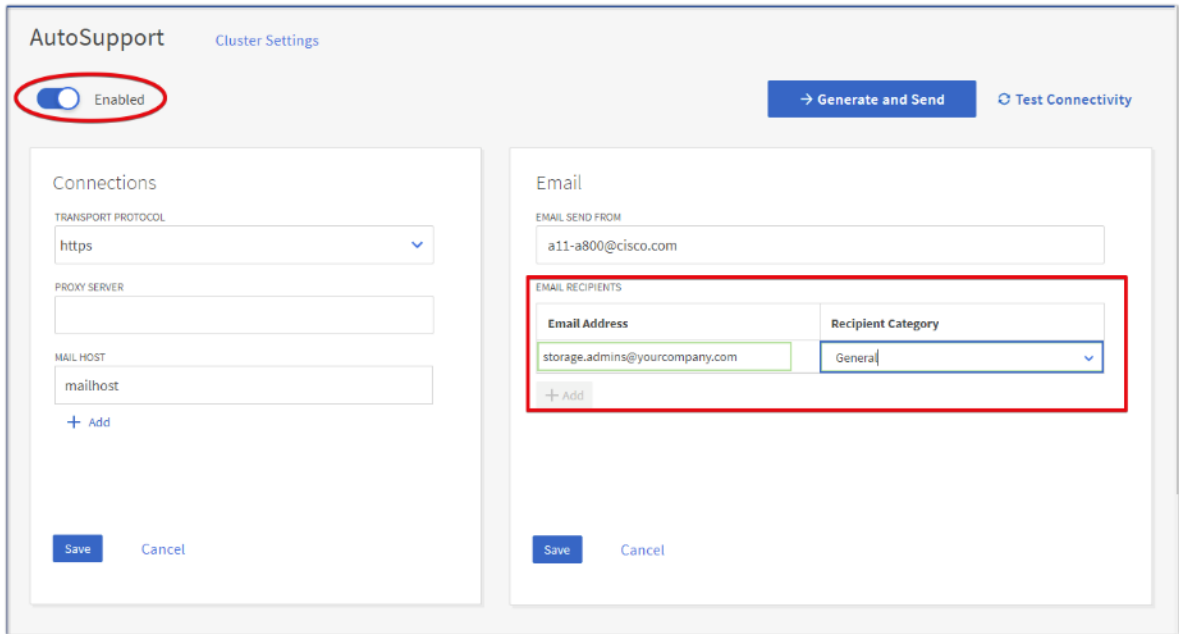
To configure AutoSupport, add licenses and create storage aggregates via the ONTAP CLI skip this section and configure the options in section.

11. Click the ellipsis in the top right of the AutoSupport tile and choose More options.

12. Choose the Settings menu under the Cluster menu.



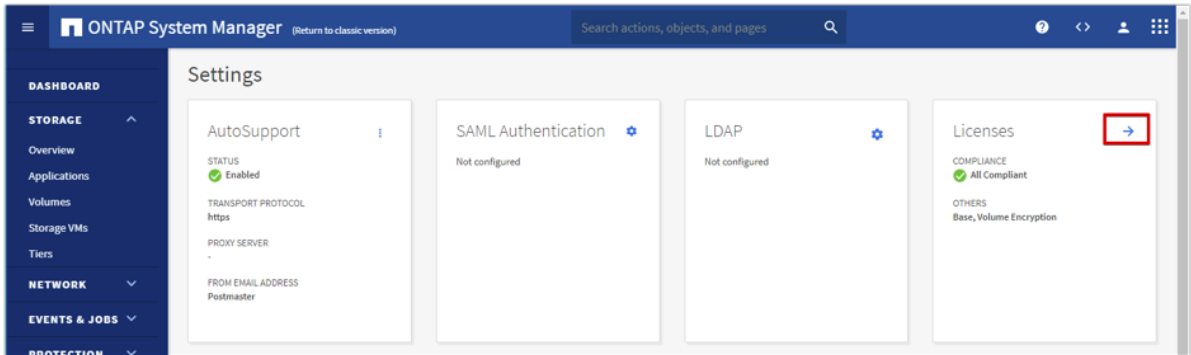
13. If AutoSupport was not configured during the initial setup, click the ellipsis in the AutoSupport tile and choose more options.
14. To enable AutoSupport click the slider button.
15. Click Edit to change the transport protocol, add a proxy server address and a mail host as needed.
16. Click Save to enable the changes.
17. In the Email tile to the right, click Edit and enter the desired email information:
  - a. Email send from address
  - b. Email recipient addresses
  - c. Recipient Category
  - d. Click Save when complete.

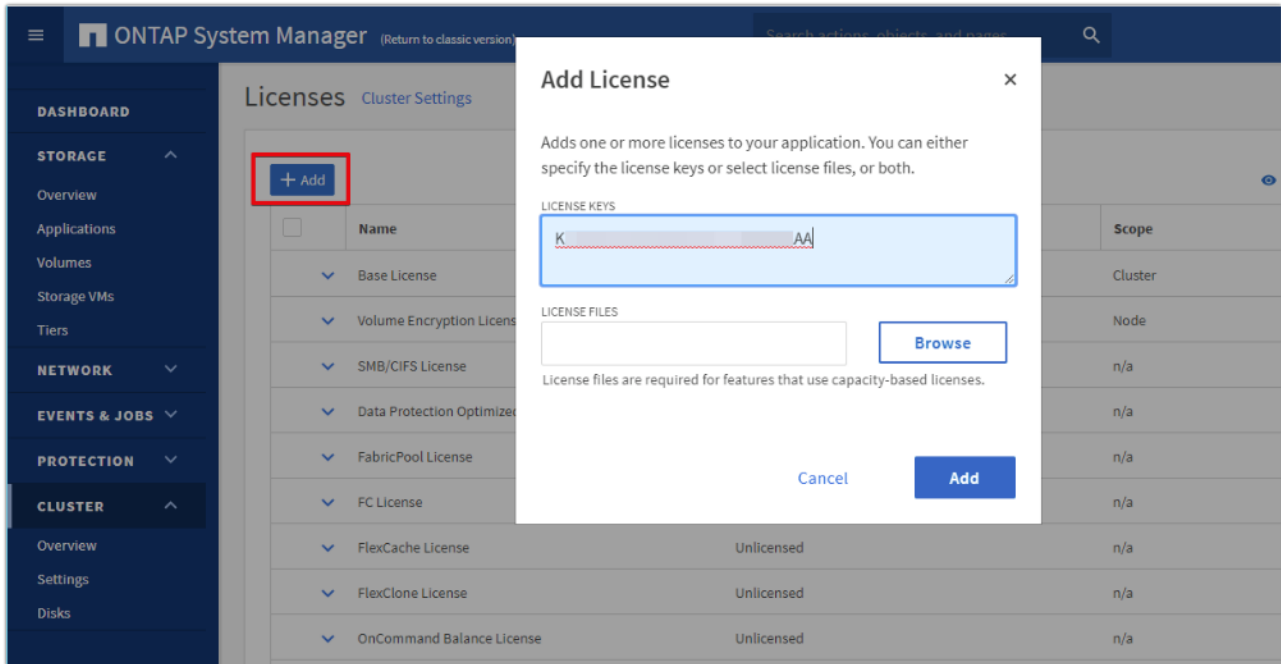


18. Choose Cluster Settings at the top left of the page to return to the cluster settings page.

19. Locate the Licenses tile on the right and click the detail arrow.

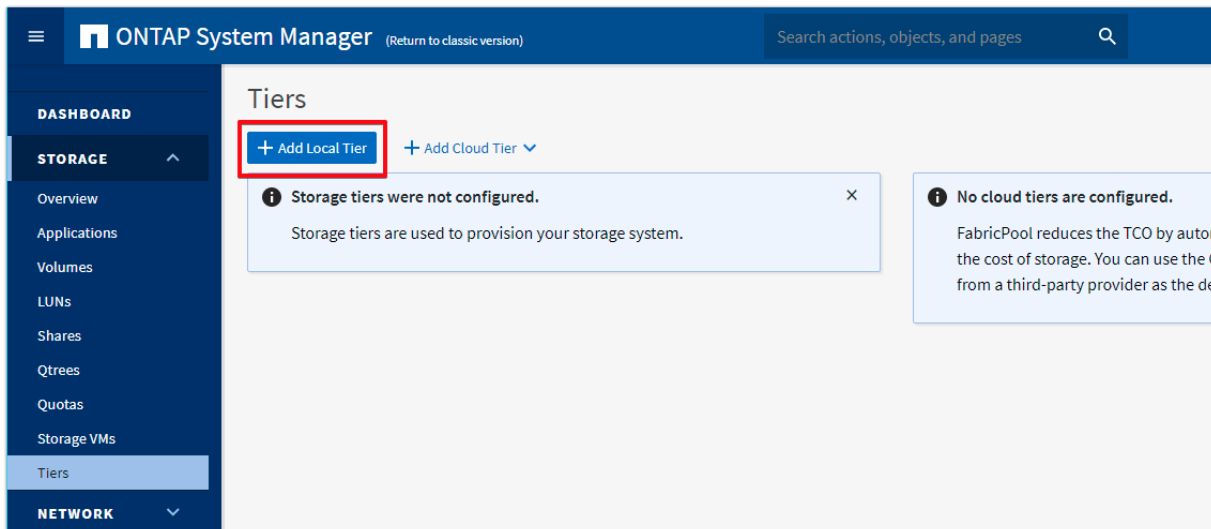
20. Add the desired licenses to the cluster by clicking Add and entering the license keys in a comma separated list.





21. Configure storage aggregates by selecting the Storage menu on the left and choosing Tiers.

22. Click Add Local Tier and allow ONTAP System Manager to recommend a storage aggregate configuration.



23. ONTAP will use best practices to recommend an aggregate layout. Click the Recommended details link to view the aggregate information.

24. Optionally, enable NetApp Aggregate Encryption (NAE) by selecting the Configure Onboard Key Manager for encryption check box.

25. Enter and confirm the passphrase and save it in a secure location for future use.

26. Click Save to make the configuration persistent.

### Add Local Tier ×

#### Storage Recommendation

**32.83 TB**  
USABLE

2 local tiers can be added on nodes "aa11-a800-01", "aa11-a800-02"

^ Recommendation details →

LOCAL TIER DETAILS

Node Name	Local Tier	Usable Size	Type
aa11-a800-01	aa11_a800_01_NVME_...	16.42 TB	SSD
aa11-a800-02	aa11_a800_02_NVME_...	16.42 TB	SSD

---

#### Encryption Considerations

Configure Onboard Key Manager for encryption →

..... X 👁

.....

i Save the passphrase for future use. You will need the passphrase if the system needs to be recovered.

Cancel Save



Careful consideration should be taken before enabling aggregate encryption. Aggregate encryption may not be supported for all deployments. Please review the [NetApp Encryption Power Guide](#) and the [Security Hardening Guide for NetApp ONTAP 9 \(TR-4569\)](#) to help determine if aggregate encryption is right for your environment.

## 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 into 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 <clustername> -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 a high availability

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 Auto-provisioned Aggregates

It is a best practice to allow ONTAP to create auto-provisioned aggregates. The auto-provisioning tool will create a storage layout including the appropriate number of spare disks according to ONTAP best practices. To create new storage aggregates with the auto-provisioning tool run the following commands, or skip to the manual aggregate creation steps below.

```
storage aggregate auto-provision -verbose
```

Per node summary of new aggregates to create, discovered spares, and also remaining spare disks and partitions after aggregate creation:

Node	New Aggrs	Total Usable Size	New -Discovered Disks	Spare- Partitions	-Remaining Disks	Spare- Partitions
aa11-a800-01	1	16.42TB	0	24	0	1
aa11-a800-02	1	16.42TB	0	24	0	1
Total:	2	32.83TB	0	48	0	2

New data aggregates to create with counts of disks and partitions to be used:

Node	New Data Aggregate	Usable Size	-Devices To Use- Disks	Partitions
aa11-a800-01	aa11_a800_01_NVME_SSD_1	16.42TB	0	23
aa11-a800-02	aa11_a800_02_NVME_SSD_1	16.42TB	0	23

RAID group layout showing how spare disks and partitions will be used in new data aggregates to be created:

RAID Group In New Data Aggregate To Be Created	Disk Type	Usable Size	Disk Or Partition	---Count--- Data Parity
/aa11_a800_01_NVME_SSD_1/plex0/rg0	NVMe-SSD	889.4GB	partition	21 2

```
/aall_a800_02_NVME_SSD_1/plex0/rg0      NVMe-SSD  889.4GB  partition  21      2
```

Details about spare disks and partitions remaining after aggregate creation:

Node	Disk Type	Device Usable Size	Disk Or Partition	Remaining Spares
aall-a800-01	NVMe-SSD	889.4GB	partition	1
aall-a800-02	NVMe-SSD	889.4GB	partition	1

Do you want to create recommended aggregates? {y|n}: y

Info: Aggregate auto provision has started. Use the "storage aggregate show-auto-provision-progress" command to track the progress.



Auto-provisioning is not supported for use with MetroCluster or third-party array LUNs. Refer to the *Aggregate creation workflow* within the Disk and Aggregate Management chapter of the [ONTAP 9 Cluster Administration guide](#) for more information.



When using aggregate auto-provisioning you cannot specify the aggregate names, however they can be changed via the ONTAP CLI or System Manager after the aggregates have been created.

## Create Manual Provisioned Aggregates (Optional)

An aggregate containing the root volume is created during the ONTAP setup process. To manually 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>
storage aggregate auto-provision -verbose
```



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.



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:

```
network broadcast-domain remove-ports -broadcast-domain Default -port <st-
node01>:e5a,<st-node01>:e5b,<st-node02>:e5a,<st-node02>:e5b

network port broadcast-domain show
```

## Disable Flow Control on 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 e0a,e1a,e5a,e5b -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption
in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <st-node02> -port e0a,e1a,e5a,e5b -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 e5a,e5b -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 e5a,e5b -autonegotiate-admin false -
speed-admin 100000 -duplex-admin full -flowcontrol-admin none

net port show -node * -port e5a,e5b -fields speed-admin,duplex-admin,flowcontrol-
admin
(network port show)
node          port duplex-admin speed-admin flowcontrol-admin
-----
aa11-a800-01 e5a   full           100000      none
aa11-a800-01 e5b   full           100000      none
```

```

aa11-a800-02 e5a full 100000 none
aa11-a800-02 e5b 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 2a -status-admin down
fc adapter modify -node <st-node01> -adapter 2b -status-admin down
fc adapter modify -node <st-node02> -adapter 2a -status-admin down
fc adapter modify -node <st-node02> -adapter 2b -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 2a -speed 32 -status-admin up
fc adapter modify -node <st-node01> -adapter 2b -speed 32 -status-admin up
fc adapter modify -node <st-node02> -adapter 2a -speed 32 -status-admin up
fc adapter modify -node <st-node02> -adapter 2b -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 e5a
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e5b

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 e5a
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e5b

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 an Active IQ Unified Manager server or another fault management system.

```
snmp traphost add <oncommand-um-server-fqdn>
```

## Configure SNMPv3 Access

SNMPv3 offers advanced security by using encryption and passphrases. The SNMPv3 user can run SNMP utilities from the traphost using the authentication and privacy settings that you specify.

```
security login create -user-or-group-name <<snmp-v3-usr>> -application snmp -
authentication-method usm
```

Enter the authoritative entity's EngineID [local EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256) [none]: <<snmp-v3-auth-PROTO>>

```

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128) [none]: <<snmpv3-
priv-PROTO>>

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

```



Refer to the [SNMP Configuration Express Guide](#) for additional details when configuring SNMPv3 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,nvme
```

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 -v3 enabled -vstorage enabled
```



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

5. Verify the NFS `vstorage` parameter for the NetApp NFS VAAI plug-in was 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 -status-admin up
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
-type server -serial <serial-number>
```



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

- 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 true
```

2. 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-lif-1a -role data -data-protocol fcp -home-node <st-node01> -home-port 1a -status-admin up

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

network interface create -vserver Infra-SVM -lif fcp-lif-2a -role data -data-protocol fcp -home-node <st-node02> -home-port 1a -status-admin up

network interface create -vserver Infra-SVM -lif fcp-lif-2b -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-lif-1 -role data -data-protocol nfs -home-node <st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs-lif-1-ip> -netmask <node01-nfs-lif-1-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM -lif nfs-lif-2 -role data -data-protocol nfs -home-node <st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs-lif-2-ip> -netmask <node02-nfs-lif-2-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 in-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. These steps have created a single data 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 open.**

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 into Cisco UCS Manager.

## Anonymous Reporting

To enable anonymous reporting, follow this step:

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

## Upgrade Cisco UCS Manager Software to Version 4.1(1c)

This document assumes the use of Cisco UCS 4.1(1c). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.1(1c), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

## 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. Choose 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.

## 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. Choose Timezone.
4. In the Properties pane, choose 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 :



**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, then click OK again.

## 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. Choose DNS Management.
4. In the Properties pane, choose 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 choose 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, choose 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.


## 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 16 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. Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate).
3. Choose 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 choose either 4, 8, 12, or 16 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	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Unconfigured	



6. Click OK, then click Yes, then click OK to continue.
7. Choose Equipment > Fabric Interconnects > Fabric Interconnect A (primary).

8. Choose 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 choose 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.

## 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:

1. In Cisco UCS Manager, click Equipment and choose the Policies tab.
2. 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, set Action set to 2 Link, the minimum recommended number of links for a FlexPod.**

3. 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 'Equipment' section of the Cisco UCS Manager interface. The 'Policies' tab is selected. Under 'Global Policies', the 'Chassis/FEX Discovery Policy' is configured. The 'Action' is set to '2 Link' and the 'Link Grouping Preference' is set to 'Port Channel'.

Chassis/FEX Discovery Policy	
Action	: 2 Link
Link Grouping Preference	: <input type="radio"/> None <input checked="" type="radio"/> Port Channel

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

## 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, choose All > Equipment in the Navigation Pane, and choose the Policies tab.
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 Server and Uplink Ports

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

1. In Cisco UCS Manager, click Equipment.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand and choose Ethernet Ports.
4. Verify that all ports connected to UCS chassis and rack mounts are configured as Server ports and have a status of Up.
5. If any rack mount ports are missing, choose the ports that are connected to Cisco FEXes and direct connect Cisco UCS C-Series servers, right-click them, and choose Configure as Server Port.
6. 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. Choose the ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.
9. Click Yes to confirm uplink ports and click OK.
10. Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
11. Expand and choose Ethernet Ports.
12. Verify that all ports connected to UCS chassis and rack mounts are configured as Server ports and have a status of Up.
13. If any rack mount ports are missing, choose the ports that are connected to Cisco FEXes and direct connect C-series servers, right-click them, and choose 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. Choose the ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.
17. Click Yes to confirm the uplink ports and click 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, choose All > Equipment in the Navigation Pane, and choose the Policies tab on the right.
2. Under Global Policies, scroll down to Info Policy and choose Enabled for Action.

### Info Policy

---

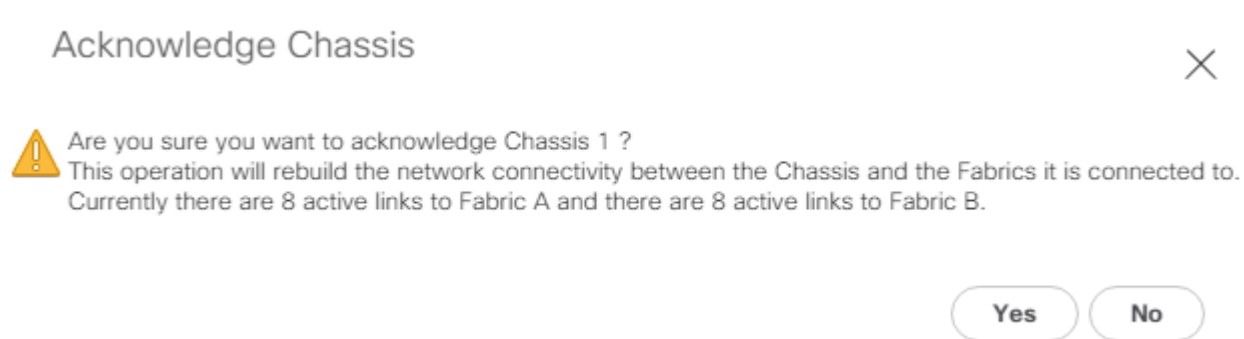
Action :  Disabled  Enabled

3. Click Save Changes and then click OK.
4. Under Equipment, choose Fabric Interconnect A or B. On the right, choose the Neighbors tab. CDP information is shown under the LAN tab and LLDP information is shown under the LLDP tab.

## 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 choose each chassis that is listed.
3. Right-click each chassis and choose Acknowledge Chassis.



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

## 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 is 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 choose Create Organization.
4. Provide a name for the Organization to indicate this FlexPod deployment and optionally provide a Description.

## Create Organization



Name :

Description :



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

### Create a WWNN Pool for FC Boot (FCP)

In this FlexPod implementation, a WWNN pool is created at the root organization level to avoid WWNN address pool overlaps. If your deployment plan calls for different WWNN ranges in different UCS organizations, place the WWNN pool at the organizational level. To configure the necessary WWNN pool for the Cisco UCS environment, follow these steps on Cisco UCS Manager.

1. Choose SAN.
2. Choose Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Choose 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. Choose Sequential for Assignment Order.

**1** Define Name and Description

**2** Add WWN Blocks

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. In this example, with the WWNN block modification, a maximum of 256 addresses are available.



## 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)

In this FlexPod implementation, WWPN address pools are created at the root organization level to avoid WWPN address pool overlaps. If your deployment plan calls for different WWPN address ranges in different UCS organizations, place the WWPN pools at the organizational level. To configure the necessary WWPN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Choose 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. Choose 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. Choose Sequential for Assignment Order.

**Create WWPN Pool** ? X

**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. In this example, with the WWPN block modification, a maximum of 256 addresses are available.

## 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**

OK

Cancel

12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.
15. Right-click WWPN Pools under the root organization.
16. Choose 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. Choose 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. In this example, with the WWPN block modification, a maximum of 256 addresses are available.
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. Choose SAN > SAN Cloud.

3. Right-click VSANs.

4. Choose 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. Choose 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. Choose 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. Choose 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. Choose Fabric A and in the Actions pane choose Enable FC Uplink Trunking.
4. Click Yes on the Confirmation and Warning.
5. Click OK.
6. Choose Fabric B and in the Actions pane choose 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. Choose SAN > SAN Cloud.
3. Expand Fabric A and choose FC Port Channels.
4. Right-click FC Port Channels and choose 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. Choose the appropriate Port Channel Admin Speed.
8. Choose 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 on the confirmation.
11. Under FC Port-Channels, choose the newly created port channel.
12. From the drop-down list to choose VSAN-A.

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

General	Ports	Faults	Events	Statistics
<b>Status</b> Overall Status : <span style="color: red;">▼ Failed</span> Additional Info : <b>No operational members</b>		<b>Properties</b> ID : <b>11</b> Fabric ID : <b>A</b> Port Type : <b>Aggregation</b> Transport Type : <b>Fc</b>		
<b>Actions</b> 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) : <b>0</b>		

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 choose 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 choose FC Port Channels.

17. Right-click FC Port Channels and choose 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. Choose 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, choose the newly created port channel.

24. In the right pane, use the drop-down to choose 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 choose 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 SAN.
2. In the Navigation Pane, expand SAN > SAN Cloud > Fabric A > Uplink FC Interfaces.
3. Right-click FC Interface 1/3 and choose Disable Interface.
4. Click Yes and OK to complete disabling FC Interface 1/3.
5. Repeat this process to disable FC Interface 1/4.
6. In the Navigation Pane, expand SAN > SAN Cloud > Fabric B > Uplink FC Interfaces.
7. Right-click FC Interface 1/3 and choose Disable Interface.
8. Click Yes and OK to complete disabling FC Interface 1/3.
9. Repeat step 1-8 to disable FC Interface 1/4.

### 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 > FlexPod.
3. Right-click vHBA Templates under the FlexPod Organization.
4. Choose 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. Choose VSAN-A.
9. Leave Initial Template as the Template Type.

10. Choose 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(250/256) ▼

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. Choose Create vHBA Template.
15. Enter FCP-vHBA-B as the vHBA template name.
16. Choose B as the Fabric ID.

17. Leave Redundancy Type set to No Redundancy.
18. Choose VSAN-B.
19. Leave Initial Template as the Template Type.
20. Choose 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. Choose SAN > Policies > root > Sub-Organizations > FlexPod.
3. Right-click SAN Connectivity Policies under the FlexPod Organization.
4. Choose Create SAN Connectivity Policy.
5. Enter FC-Boot as the name of the policy.
6. Choose 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. Choose the Use vHBA Template checkbox.
10. In the vHBA Template list, choose FCP-vHBA-A.
11. In the Adapter Policy list, choose 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. Choose the Use vHBA Template checkbox.
16. In the vHBA Template list, choose FCP-vHBA-B.
17. In the Adapter Policy list, choose 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.

## 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 choose Create Block of IPv4 Addresses.

- 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"/>



- Click OK to create the block.
- Click OK in the confirmation message.

## Create Uplink Port Channels to Cisco Nexus Switches

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

- 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.**

- Under LAN > LAN Cloud, expand the Fabric A tree.
- Right-click Port Channels under Fabric A.
- Choose Create Port Channel.
- Enter 153 as the unique ID of the port channel.
- Enter Po153-Nexus as the name of the port channel.
- Click Next.
- Choose 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.

**1** Set Port Channel Name

**2** Add Ports

### Create Port Channel

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, choose Port-Channel 153. Choose 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. Choose 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. Choose 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, choose Port-Channel 154. Choose 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, choose Eth Interface 1/53. In the center pane under Properties, enter a User Label to indicate the port connectivity, such as `<nexus-a-hostname>:Eth1/5`. 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 choose Create UDLD Link Policy.
4. Name the Policy UDLD-Normal and choose 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 click OK again to complete creating the policy.
6. Expand Policies > LAN Cloud > Link Profile.
7. Right-click Link Profile and choose Create Link Profile.
8. Name the Profile UDLD-Normal and choose 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. Choose the first Eth Interface under Port-Channel 153. From the drop-down list, choose 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
<b>Actions</b>		<b>Properties</b>	
Delete		ID	: <b>54</b>
Enable Interface		Slot ID	: <b>1</b>
Disable Interface		Fabric ID	: <b>B</b>
		Transport Type	: <b>Ether</b>
		Port	: sys/switch-B/slot-1/switch-ether/port-54
		Membership	: <b>Up</b>
		Link Profile	: UDLD-Normal ▼
		User Label	: <input type="text"/>

## 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 with UCS Manager version 4.0 software 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. In UCS Manager version 4.1, the MTU for the Best Effort QoS System Class is again settable. To configure jumbo frames in the UCS fabric, follow these steps:

1. 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

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

Configure Slow Drain Timers



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 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. Choose 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 choose Set as Native VLAN.
11. Click Yes and then click OK.
12. Right-click VLANs.
13. Choose 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. Choose 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. Choose 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. Choose 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.

LAN / LAN Cloud / VLANs

VLANs

Advanced Filter Export Print

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

+ Add Delete Info

## Create MAC Address Pools

In this FlexPod implementation, MAC address pools are created at the root organization level to avoid MAC address pool overlaps. If your deployment plan calls for different MAC address ranges in different UCS organizations, place the MAC pools at the organizational level. To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand 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. Choose Create MAC Pool to create the MAC address pool.
5. Enter MAC-Pool-A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Choose 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.**

- 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. In this example, with the MAC block modification, a maximum of 256 addresses are available.

## 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**



- Click OK.
- Click Finish.
- In the confirmation message, click OK.
- Right-click MAC Pools under the root organization.
- Choose Create MAC Pool to create the MAC address pool.
- Enter MAC-Pool-B as the name of the MAC pool.
- Optional: Enter a description for the MAC pool.
- Choose Sequential as the option for Assignment Order.
- Click Next.
- Click Add.
- 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.**

- 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. In this example, with the MAC block modification, a maximum of 256 addresses are available.



24. Click OK.
25. Click Finish.
26. In the confirmation message, 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. Choose Create Network Control Policy.
5. Enter Enable-CDP-LLDP as the policy name.
6. For CDP, choose the Enabled option.
7. For LLDP, scroll down and choose Enabled for both Transmit and Receive.

### Create Network Control Policy ? X

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 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 (vSwitch0-A and vSwitch0-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 (vDS0-A and vDS0-B) will be created for vNICs to connect to the VMware Virtual Distributed Switch (vDS0). The vDS will have port groups for the vMotion and VM-Traffic VLANs. The vMotion VLAN is being placed on both vSwitch0 and vDS0 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. Any tenant or application VLANs can be placed on the vDS in the future.

### 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 > FlexPod.
3. Under the FlexPod Organization, right-click vNIC Templates.
4. Choose 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. Choose 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. Choose Updating Template as the Template Type.
12. Under VLANs, choose the checkboxes for IB-MGMT, Infra-NFS, vMotion, and Native-VLAN VLANs.
13. Set Native-VLAN as the native VLAN.
14. Choose vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, choose MAC-Pool-A.
17. In the Network Control Policy list, choose Enable-CDP-LLDP.

## Create vNIC Template



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	<input type="radio"/>	
<input checked="" type="checkbox"/>	IB-MGMT	<input type="radio"/>	113
<input checked="" type="checkbox"/>	Infra-NFS	<input type="radio"/>	3050
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>	2
<input 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 :  ▼

**Connection Policies**

OK

Cancel

18. Click OK to create the vNIC template.
19. Click OK.
20. Under the FlexPod organization, right-click vNIC Templates.
21. Choose Create vNIC Template.
22. Enter vSwitch0-B as the vNIC template name.

23. Choose Fabric B.
24. Do not select the Enable Failover checkbox.
25. Set Redundancy Type to Secondary Template.
26. Choose vSwitch0-A for the Peer Redundancy Template.
27. In the MAC Pool list, choose 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. Choose 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.
35. Choose 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. Choose Updating Template as the Template Type.
39. Under VLANs, choose the checkboxes for vMotion, VM-Traffic, and Native-VLAN VLANs.
40. Set Native-VLAN as the native VLAN.
41. Choose vNIC Name for the CDN Source.
42. For MTU, enter 9000.
43. In the MAC Pool list, choose MAC-Pool-A.
44. In the Network Control Policy list, choose Enable-CDP-LLDP.

## Create vNIC Template



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	<input type="radio"/>	1
<input 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 :  ▼

**Connection Policies**

45. Click OK to create the vNIC template.

46. Click OK.

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

48. Choose Create vNIC Template

49. Enter vDS0-B as the vNIC template name.

50. Choose Fabric B.
51. Do not select the Enable Failover checkbox.
52. Set Redundancy Type to Secondary Template.
53. Choose vDS0-A for the Peer Redundancy Template.
54. In the MAC Pool list, choose 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 choose 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 : 1 [1-1000]

Ring Size : 256 [64-4096]

Receive Queues : 8 [1-1000]

Ring Size : 512 [64-4096]

Completion Queues : 9 [1-2000]

Interrupts : 11 [1-1024]

## Options

OK

Cancel



In this policy, Receive Queues can be set to 1-16. Completion Queues = Transmit Queues + Receive Queues. Interrupts = Completion Queues + 2. For more information, see [Cisco UCS Manager Network Management Guide, Release 4.1, Network-Related Policies](#).



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 choose Enabled for Receive Side Scaling (RSS).

## Create Ethernet Adapter Policy



Name : VMware-HighTrf

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 > FlexPod.
3. Under the FlexPod Organization, right-click LAN Connectivity Policies.
4. Choose Create LAN Connectivity Policy.
5. Enter FC-Boot as the name of the policy.
6. Click OK then OK again to add the policy.
7. In the menu on the left under LAN > Policies > root > Sub-Organizations > FlexPod > LAN Connectivity Policies, choose FC-Boot.
8. Click the Add button to add a vNIC.
9. In the Create vNIC dialog box, enter 00-vSwitch0-A as the name of the vNIC.
10. Choose the Use vNIC Template checkbox.
11. In the vNIC Template list, choose vSwitch0-A.
12. In the Adapter Policy list, choose 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](#)

13. Click OK to add this vNIC to the policy.
14. Click Save Changes and OK.
15. Click Add to add another vNIC to the policy.
16. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.
17. Check the box for the Use vNIC Template.
18. In the vNIC Template list, choose vSwitch0-B.
19. In the Adapter Policy list, choose VMWare.
20. Click OK to add the vNIC to the policy.
21. Click Save Changes and OK.

22. Click Add to add another vNIC to the policy.
23. In the Create vNIC dialog box, enter 02-vDS0-A as the name of the vNIC.
24. Choose the Use vNIC Template checkbox.
25. In the vNIC Template list, choose vDS0-A.
26. In the Adapter Policy list, choose VMWare-HighTrf.



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

---

27. Click OK to add this vNIC to the policy.
28. Click Save Changes and OK.
29. Click Add to add another vNIC to the policy.
30. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.
31. Choose the Use vNIC Template checkbox.
32. In the vNIC Template list, choose vDS0-B.
33. In the Adapter Policy list, choose VMWare-HighTrf.



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

---

34. Click OK to add this vNIC to the policy.
35. Click Save Changes and OK.

General
Events

**Actions**

---

[Delete](#)

[Show Policy Usage](#)

[Use Global](#)

Name : **FC-Boot**

Description :

Owner : **Local**

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

Name	MAC Address
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

## 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. Expand Pools > root > Sub-Organizations > FlexPod.
3. Right-click Server Pools under the FlexPod Organization.
4. Choose Create Server Pool.
5. Enter Infra-Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Choose 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 use cases three servers are recommended.**

9. Click Finish.
10. 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. Expand Pools > root.
3. Right-click UUID Suffix Pools.
4. Choose Create UUID Suffix Pool.
5. Enter UUID-Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Choose 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 and the number of Service Profiles that will be created.
13. Click OK.
14. Click Finish.
15. Click OK.

## Modify Default Host Firmware Package

Firmware management policies allow the administrator to choose 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. Choose default.
5. In the Actions pane, choose Modify Package Versions.

## Modify Package Versions



Blade Package :

Rack Package :

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

6. Choose version 4.1(1c) for both the Blade and Rack Packages.
7. Click OK, then click OK again to modify the host firmware package.

### Create Local Disk Configuration Policy (Optional)

A local disk configuration specifying no local disks for the Cisco UCS environment can be used to ensure that servers with no local disks are used for SAN Boot.



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. Choose Create Local Disk Configuration Policy.
5. Enter SAN-Boot as the local disk configuration policy name.
6. Change the mode to No Local Storage.

## Create Local Disk Configuration Policy



Name : SAN-Boot

Description :

Mode : No Local Storage ▼

### FlexFlash

FlexFlash State :  Disable  Enable

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

FlexFlash RAID Reporting State :  Disable  Enable

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 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. Choose Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.

### 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. Choose Create Server Pool Policy Qualification.
5. Name the policy UCS-B200M5.
6. Choose Create Server PID Qualifications.
7. Choose UCSB-B200-M5 from the PID drop-down list.

### Create Server PID Qualifications



PID :



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

## 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. Choose Maintenance Policies > default.

- Change the Reboot Policy to User Ack.
- Choose “On Next Boot” to delegate maintenance windows to server administrators.

Servers / Policies / root / Maintenance Policies / default

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

- Click Save Changes.
- Click OK to accept the changes.

## Create 100 Percent App Direct Mode Persistent Memory Policy

If any servers in your environment are equipped with Intel Optane DC Persistent Memory (PMEM), a Persistent Memory Policy should be used. Intel Optane DC PMEM can be used in App Direct Mode, Memory Mode, or Mixed Mode where the Intel Optane DC PMEM is divided between App Direct Mode and Memory Mode. In a Cisco UCS server that is equipped with Intel Optane DC PMEM, if a Persistent Memory Policy is not assigned, 100% of the Intel Optane DC PMEM will be used in Memory Mode and the standard DIMMs in the server will be used as cache and the DIMM capacity will not be visible. In VMware vSphere 6.7U3, usage of Intel Optane DC PMEM in Memory Mode in production environments requires a waiver from VMware. In order to avoid having to get this waiver, a Persistent Memory Policy that assigns 100% of the Intel Optane DC PMEM to App Direct Mode should be assigned to all servers that are equipped with Intel Optane DC PMEM.

- In Cisco UCS Manager, choose Servers.

2. Expand Policies > root.
3. Right-click Persistent Memory Policy.
4. Choose Create Persistent Memory Policy.
5. Name the policy App-Direct-Mode.
6. Under Goals, click Add.
7. Leave Memory Mode (%) set to zero and Persistent Memory Type set to App Direct.

## Create Goal



### Properties

Socket ID	:	<input checked="" type="radio"/> All Sockets
Memory Mode (%)	:	<input type="text" value="0"/>
Persistent Memory Type	:	<input checked="" type="radio"/> App Direct <input type="radio"/> App Direct Non Interleaved



8. Click OK to complete creating the Goal.
9. Click OK to complete creating the policy and click OK on the confirmation.

## Create vMedia Policy for VMware ESXi 6.7U3 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 6.7U3 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, choose Servers.
2. Expand Policies > root.
3. Right-click vMedia Policies.
4. Choose Create vMedia Policy.
5. Name the policy ESXi-6.7U3-HTTP.

6. Enter "Mounts Cisco Custom ISO for ESXi 6.7U3" in the Description field.
7. Click Add to add a vMedia Mount.
8. Name the mount ESXi-6.7U3-HTTP.
9. Choose the CDD Device Type.
10. Choose 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\_14320388\_Custom\_Cisco\_6.7.3.1.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-6.7U3-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_14320388_Custom_Cisco_6.7"/>
Remote Path	:	<input type="text" value="software/vSphere-6.7U3"/>
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 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 6.7U3, the TPM module must be enabled and Trusted Execution Technology (TXT) disabled in BIOS. According to the [Cisco UCS Server BIOS Tolerances, Release 4.1](#) 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 > FlexPod.
3. Right-click BIOS Policies under FlexPod Organization.
4. Choose 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 choose the newly created BIOS Policy. Set the following within the Main tab of the Policy:
  - a. CDN Control -> Enabled
  - b. Quiet Boot -> Disabled

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

Main Advanced Boot Options Server Management Events

## 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. Processor C State -> Disabled
  - b. Processor C1E -> Disabled
  - c. Processor C3 Report -> Disabled
  - d. Processor C6 Report -> Disabled
  - e. Processor C7 Report -> Disabled
  - f. Power Technology -> Custom

OS Setting	Value
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMCi	Platform Default
Power Technology	Custom
Energy Performance	Platform Default
ProcessorEpgProfile	Platform Default

9. Click the RAS Memory tab, and choose:

- a. Memory RAS configuration -> Maximum Performance

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

10. Click Save Changes.

11. Click OK.

## 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-lif-1a and fcp-lif-1b) and two Fibre Channel LIFs are on cluster node 2 (fcp-lif-2a and fcp-lif-2b). 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-lif-1a.**

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 > FlexPod.
3. Under the FlexPod Organization, right-click Boot Policies.
4. Choose 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. Choose the Uefi Boot Mode.
9. Choose 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 6.7U3 with or without a TPM 2.0 module in the UCS server.**

10. Expand Local Devices and choose Add Remote CD/DVD.
11. Expand vHBAs and choose Add SAN Boot.
12. Choose 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, choose Add SAN Boot Target.
- 16. Keep 0 as the value for Boot Target LUN.
- 17. Enter the WWPN for fcp-lif-1a.

---

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

---

- 18. Choose 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, choose Add SAN Boot Target.
21. Enter 0 as the value for Boot Target LUN.
22. Enter the WWPN for fcp-lif-2a.
23. Click OK to add the SAN boot target.
24. From vHBAs, choose 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, choose Add SAN Boot Target.
29. Keep 0 as the value for Boot Target LUN.
30. Enter the WWPN for fcp-lif-1b.
31. Choose Primary for the SAN boot target type.
32. Click OK to add the SAN boot target.
33. From vHBAs, choose Add SAN Boot Target.
34. Keep 0 as the value for Boot Target LUN.

35. Enter the WWPN for fcp-lif-2b.
36. Click OK to add the SAN boot target.
37. Expand CIMC Mounted Media and choose 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**

+ - Advanced Filter Export Print ⚙

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								

↑ Move Up 
 ↓ Move Down 
 🗑 Delete

Set Uefi Boot Parameters

OK
Cancel

38. Expand San > SAN Primary and choose SAN Target Primary. Choose 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 6.7U3 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 > FlexPod.
3. Right-click the FlexPod Organization.
4. Choose 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. Choose the Updating Template option.
7. Under UUID, choose 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 tab and choose the SAN-Boot Local Storage Policy. Otherwise, choose the default Local Storage Policy.
2. Click Next.

## Configure Networking

To configure networking, follow these steps:

1. Choose the “Use Connectivity Policy” option to configure the LAN connectivity.
2. Choose FC-Boot from the LAN Connectivity Policy drop-down list.
3. 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

4. Click Next.

## Configure SAN Connectivity

To configure SAN connectivity, follow these steps:

1. Choose the Use Connectivity Policy option for the “How would you like to configure SAN connectivity?” field.
2. Choose 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](#)

< Prev
Next >
Finish
Cancel

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. Choose 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, choose Infra-Pool.
2. Choose Down as the power state to be applied when the profile is associated with the server.
3. Optional: Choose “B200-M5” for the Server Pool Qualification to choose only B200 M5 servers in the pool.
4. Expand Firmware Management and choose 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 :  [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, choose VM-Host.
2. Expand Power Control Policy Configuration and choose 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. Choose Service Profile Templates > root > Sub-Organizations > FlexPod > Service Template VM-Host-Infra-FCP-A.
3. Right-click VM-Host-Infra-FCP-A and choose Create a Clone.
4. Name the clone VM-Host-Infra-FCP-A-vM.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly-created VM-Host-Infra-FCP-A-vM and choose the vMedia Policy tab.
7. Click Modify vMedia Policy.

8. Choose the ESXi-6.7U3-HTTP vMedia Policy and click OK.
9. Click OK to confirm.

## Create Service Profile Template for Servers with Optane Memory

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlexPod > Service Template VM-Host-Infra-FCP-A.
3. Right-click VM-Host-Infra-FCP-A and choose Create a Clone.
4. Name the clone Optane-Host-Infra-FCP-A.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly-created Optane-Host-Infra-FCP-A and choose the Policies tab.
7. Expand Persistent Memory Policy and choose the App-Direct-Mode policy.
8. Click Save Changes and then click OK to confirm.

## Create vMedia-Enabled Service Profile Template for Servers with Optane Memory

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlexPod > Service Template Optane-Host-Infra-FCP-A.
3. Right-click Optane-Host-Infra-FCP-A and choose Create a Clone.
4. Name the clone Optane-Host-Infra-FCP-A-vM.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly-created Optane-Host-Infra-FCP-A vM and choose the vMedia Policy tab.
7. Click Modify vMedia Policy.
8. Choose the ESXi-6.7U3-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. Choose 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 choose 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 ? ×

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 6.7U3 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.



**If the Service Profiles being built will be associated with servers equipped with Intel Optane DC PMEM, use the Optane-Host-Infra-FCP-A-vM service profile template instead.**

### 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 pools and policies created at the organizational level will need to be recreated within other 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.

**Table 6 WWPNS from NetApp Storage**

SVM	Adapter	MDS Switch	Target: WWPNS
Infra-SVM	fcplif-1a	Fabric A	<fcplif-1a-wwpn>
	fcplif-1b	Fabric B	<fcplif-1b-wwpn>
	fcplif-2a	Fabric A	<fcplif-2a-wwpn>
	fcplif-2b	Fabric B	<fcplif-2b-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 click the 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 : <nexus-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 : <nexus-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 fc1/1
switchport description <st-clustername>-1:2a
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description <st-clustername>-2:2a
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 fc1/1
switchport description <st-clustername>-1:2b
switchport speed 32000
switchport trunk mode off
```

```

no shutdown
exit

interface fc1/2
switchport description <st-clustername>-2:2b
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 fc1/1
vsan <vsan-a-id> interface fc1/2
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 fc1/1
vsan <vsan-b-id> interface fc1/2
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-lif-1a pwwn <fcp-lif-1a-wwpn>
device-alias name Infra-SVM-fcp-lif-2a pwwn <fcp-lif-2a-wwpn>
device-alias name VM-Host-Infra-01-A pwwn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwwn <vm-host-infra-02-wwpna>
device-alias name VM-Host-Infra-03-A pwwn <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-lif-1b pwwn <fcp-lif-1b-wwpn>
device-alias name Infra-SVM-fcp-lif-2b pwwn <fcp-lif-2b-wwpn>
device-alias name VM-Host-Infra-01-B pwwn <vm-host-infra-01-wwpnb>
device-alias name VM-Host-Infra-02-B pwwn <vm-host-infra-02-wwpnb>
device-alias name VM-Host-Infra-03-B pwwn <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-lif-1a target
member device-alias Infra-SVM-fcp-lif-2a 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-lif-1b target
member device-alias Infra-SVM-fcp-lif-2b 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 6.7U3 Setup

---

### VMware ESXi 6.7U3

This section provides detailed instructions for installing VMware ESXi 6.7U3 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 6.7U3 from VMware

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

1. Click the following link: [Cisco Custom ESXi 6.7U3 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. Choose Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-01.
8. In the Actions pane, click KVM Console.
9. Follow the prompts to launch the HTML5 KVM console.
10. Choose Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-02.

11. In the Actions pane, click KVM Console.
12. Follow the prompts to launch the HTML5 KVM console.
13. Choose Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-03.
14. In the Actions pane, click KVM Console.
15. Follow the prompts to launch the HTML5 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. Choose Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and choose 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. Choose the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

6. Choose 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.

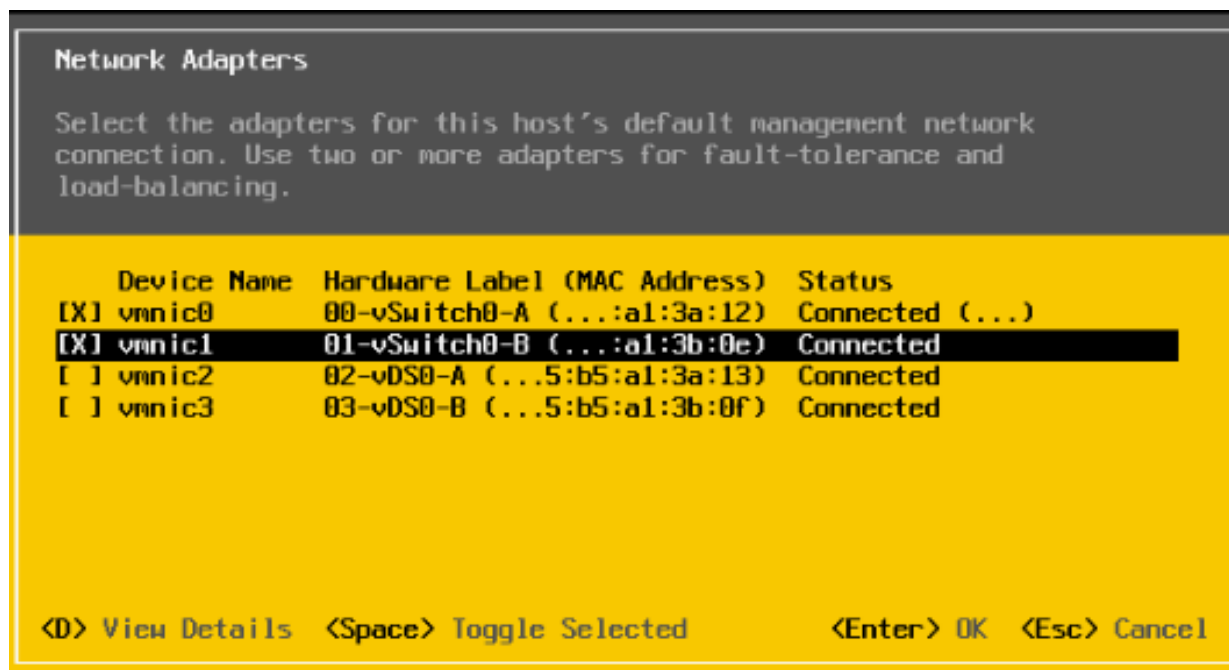
## Set Up Management Networking for ESXi Hosts

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, in the UCS KVM console, press F2 to customize VMware ESXi.
2. Log in as root, enter the corresponding password, and press Enter to log in.
3. Use the down arrow key to choose Troubleshooting Options and press Enter.
4. Choose Enable ESXi Shell and press Enter.
5. Choose Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Choose the Configure Management Network option and press Enter.
8. Choose Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field. If the numbers do not match, note the mapping of vmnic ports to vNIC ports for later use.
10. Using the spacebar, choose 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. Choose the VLAN (Optional) option and press Enter.
13. Enter the <ib-mgmt-vlan-id> and press Enter.
14. Choose IPv4 Configuration and press Enter.
15. Choose 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. Choose the IPv6 Configuration option and press Enter.
21. Using the spacebar, choose Disable IPv6 (restart required) and press Enter.
22. Choose 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, choose "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:

[nfnic Driver version 4.0.0.52](#)

[nenic Driver version 1.0.31.0](#)

[NetApp NFS Plug-in 1.1.2-3 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/Cisco-nfnic_4.0.0.52-10EM.670.0.0.8169922-offline_bundle-15920005.zip

esxcli software vib update -d /tmp/VMW-ESX-6.7.0-nenic-1.0.31.0-offline_bundle-15180549.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 Navigator, choose Networking.
2. In the center pane, choose the Virtual switches tab.
3. Highlight the vSwitch0 line.
4. Choose Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, choose vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.
9. Click Save.
10. Choose Networking, then choose the Port groups tab.
11. In the center pane, right-click VM Network and choose Edit settings.
12. Name the port group IB-MGMT Network and enter <ib-mgmt-vlan-id> in the VLAN ID field.
13. Click Save to finalize the edits for the IB-MGMT Network.
14. At the top, choose the VMkernel NICs tab.
15. Click Add VMkernel NIC.
16. For New port group, enter VMkernel-Infra-NFS.
17. For Virtual switch, choose vSwitch0.
18. Enter <infra-nfs-vlan-id> for the VLAN ID.
19. Change the MTU to 9000.
20. Choose Static IPv4 settings and expand IPv4 settings.
21. Enter the ESXi host Infrastructure NFS IP address and netmask.
22. Leave TCP/IP stack set at Default TCP/IP stack and do not choose any of the Services.

23. Click Create.
24. Click Add VMkernel NIC.
25. For New port group, enter VMkernel-vMotion.
26. For Virtual switch, choose vSwitch0.
27. Enter <vmotion-vlan-id> for the VLAN ID.
28. Change the MTU to 9000.
29. Choose Static IPv4 settings and expand IPv4 settings.
30. Enter the ESXi host vMotion IP address and netmask.
31. Choose the vMotion stack for TCP/IP stack.
32. Click Create.
33. Choose 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	11776 (11765 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 four port groups on the left connected to a central vSwitch, which is then connected to physical adapters on the right:

- IB-MGMT Network** (VLAN ID: 113) connects to vmnic1 and vmnic0.
- VMkernel-vMotion** (VLAN ID: 3000) connects to vmnic1 and vmnic0. It contains one VMkernel port: vmk2 (192.168.100.21).
- VMkernel-Infra-NFS** (VLAN ID: 3050) connects to vmnic1 and vmnic0. It contains one VMkernel port: vmk1 (192.168.50.21).
- Management Network** (VLAN ID: 113) connects to vmnic1 and vmnic0. It contains one VMkernel port: vmk0 (10.1.156.21).

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

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.156.21	None
vmk1	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.21	None
vmk2	VMkernel-vMotion	vMotion stack	vMotion	192.168.100.21	None

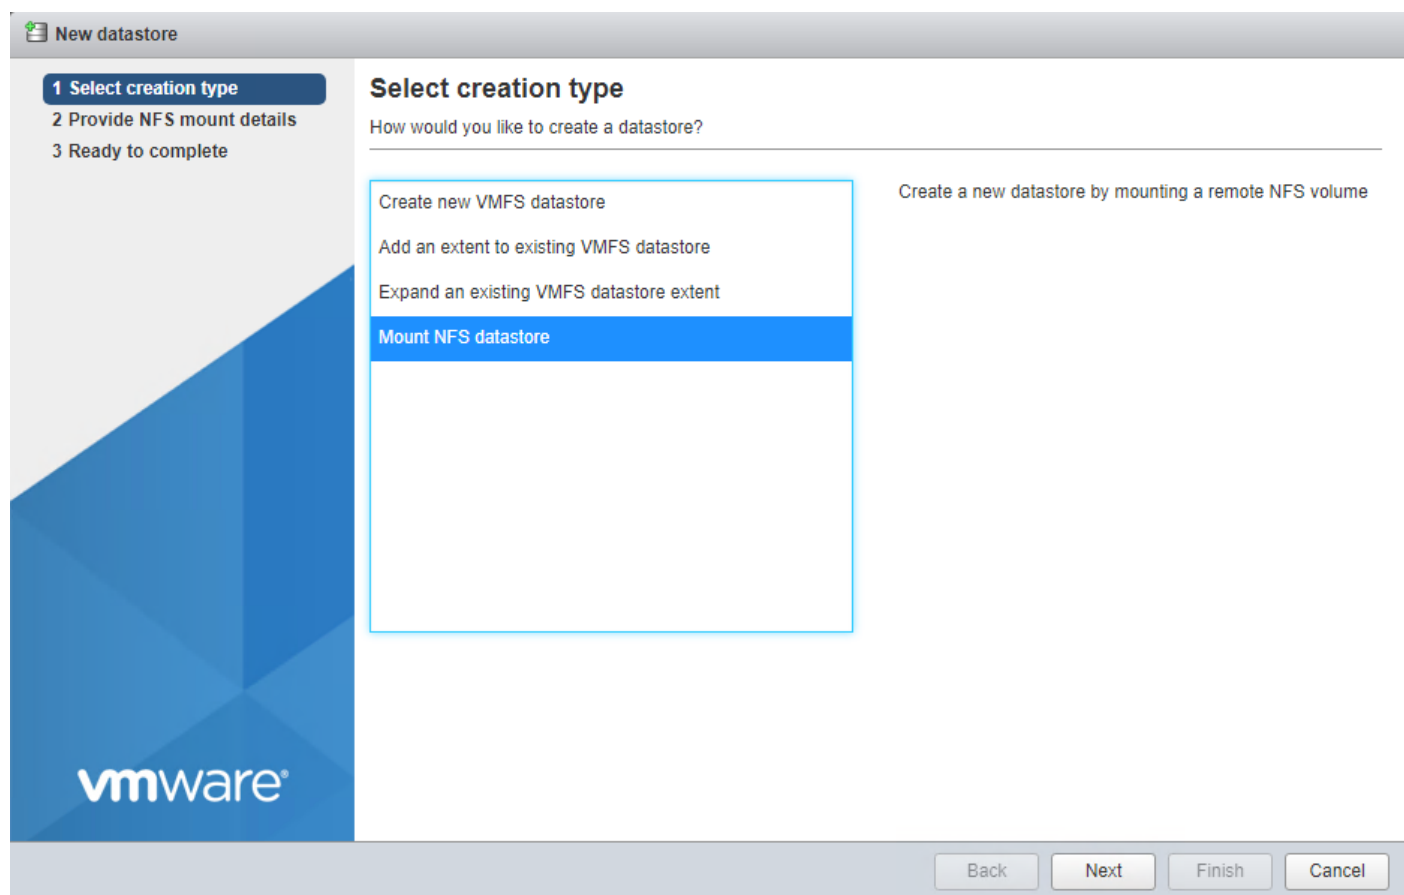


## Mount Required Datastores

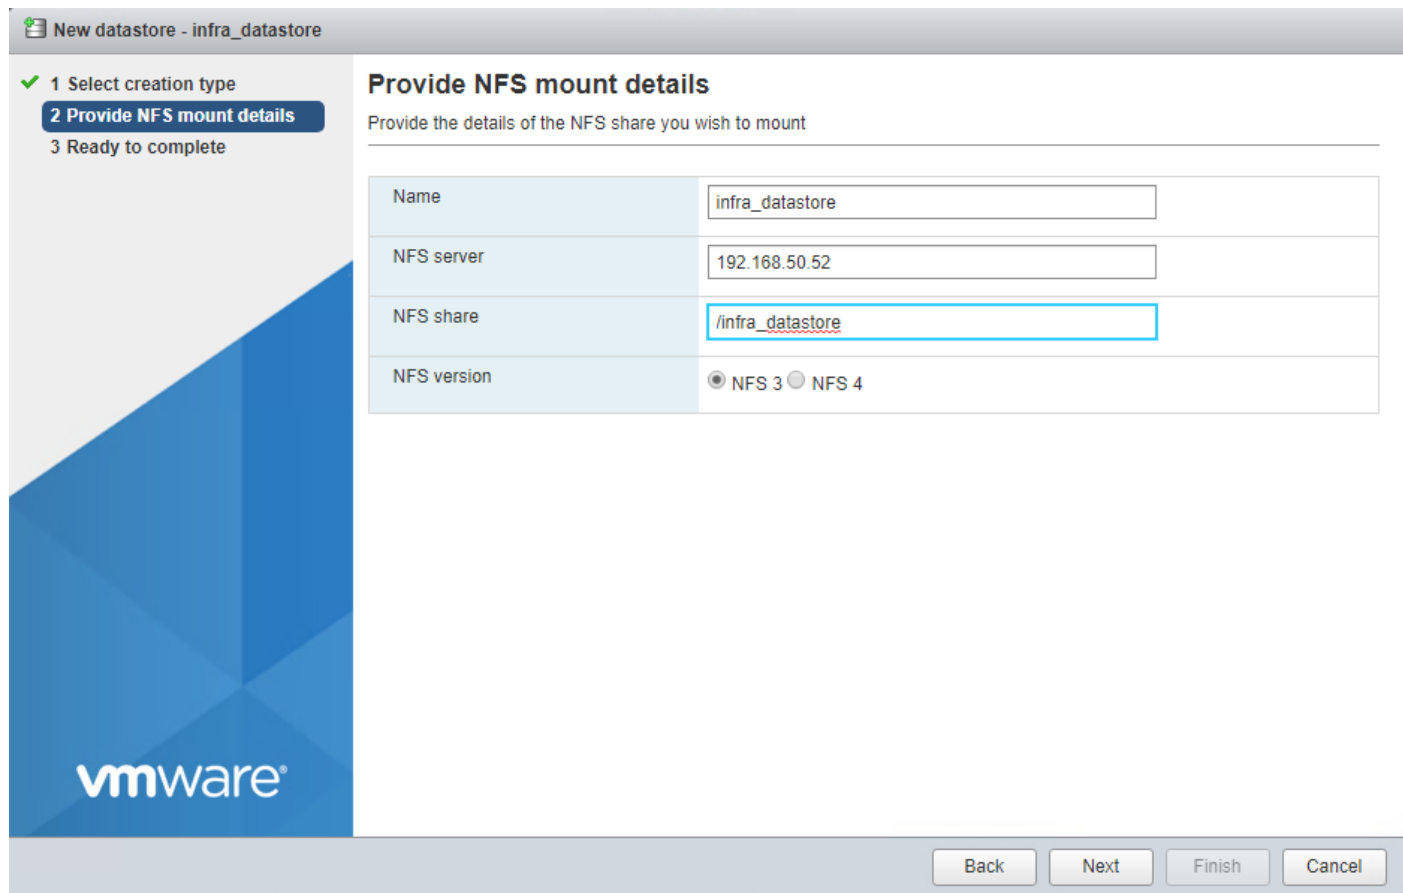
### ESXi Host VM-Host-Infra-01

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

1. From the Host Client, choose Storage.
2. In the center pane, choose the Datastores tab.
3. In the center pane, choose New Datastore to add a new datastore.
4. In the New datastore popup, choose Mount NFS datastore and click Next.



5. Input infra\_datastore for the datastore name. Input the IP address for the nfs-lif-2 LIF for the NFS server. Input /infra\_datastore for the NFS share. Leave the NFS version set at NFS 3. Click Next.



6. Click Finish. The datastore should now appear in the datastore list.
7. In the center pane, choose New Datastore to add a new datastore.
8. In the New datastore popup, choose Mount NFS datastore and click Next.
9. Input infra\_swap for the datastore name. Input the IP address for the nfs-lif-1 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 provisioning	Access		
datastore1	SSD	7.5 GB	1.41 GB	6.09 GB	VMFS6	Supported	Single		
infra_datastore	Unknown	1,024 GB	540 KB	1,024 GB	NFS	Supported	Single		
infra_swap	Unknown	100 GB	372 KB	100 GB	NFS	Supported	Single		

## 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, choose Manage.
2. In the center pane, choose 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 choose 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

04/03/2020 8:48 AM

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. Choose 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, choose Manage.
2. In the center pane, choose System > Swap.

3. Click Edit settings.
4. Use the drop-down list to choose `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.

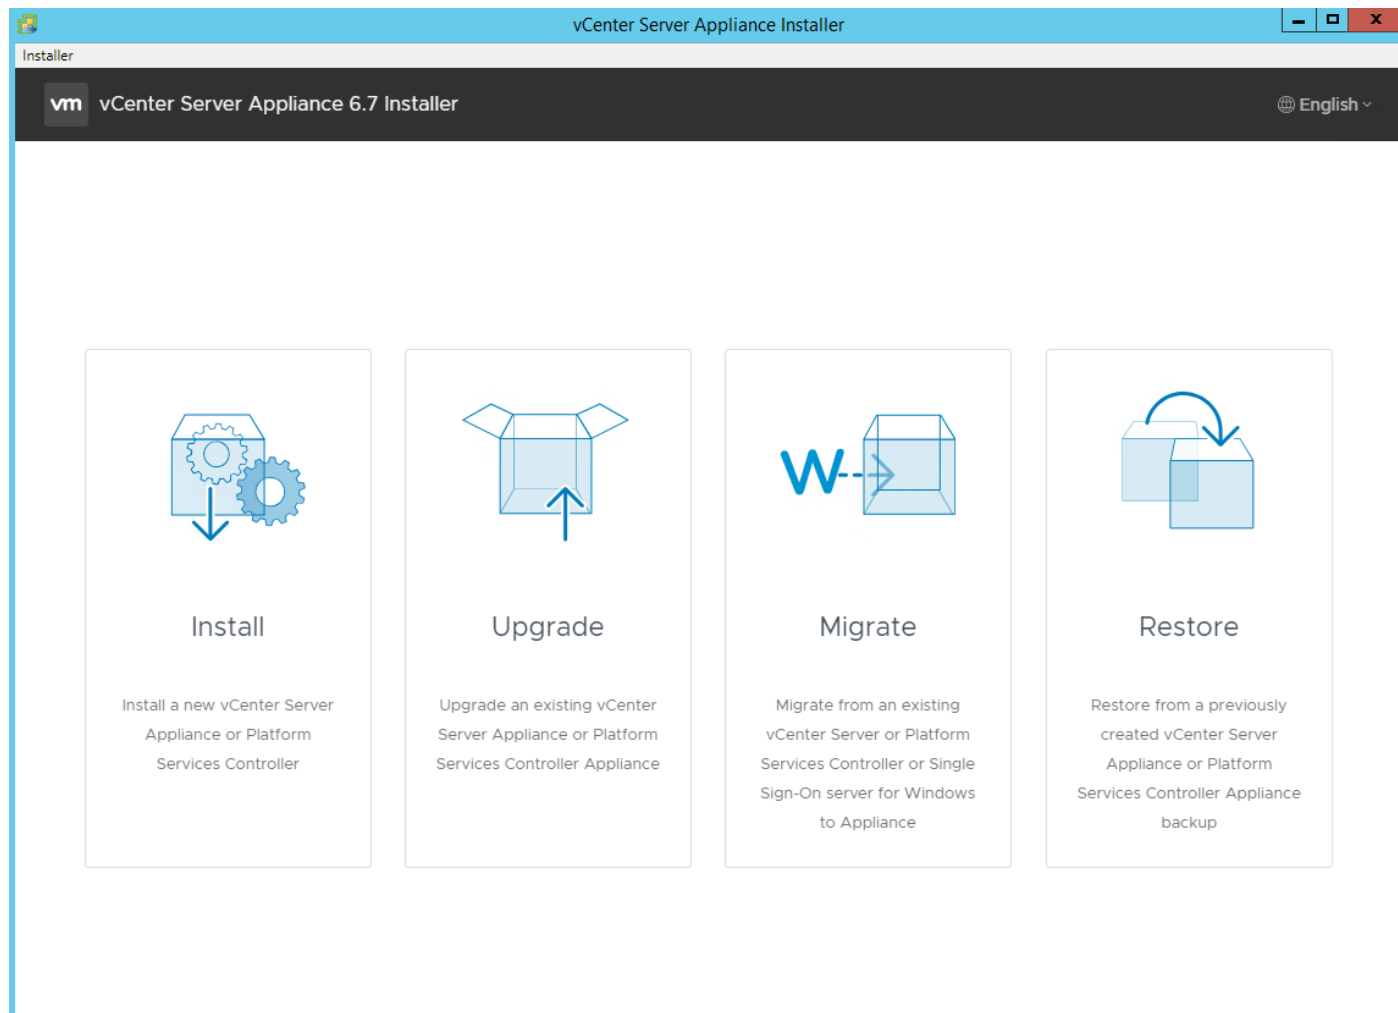
## VMware vCenter 6.7U3

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.

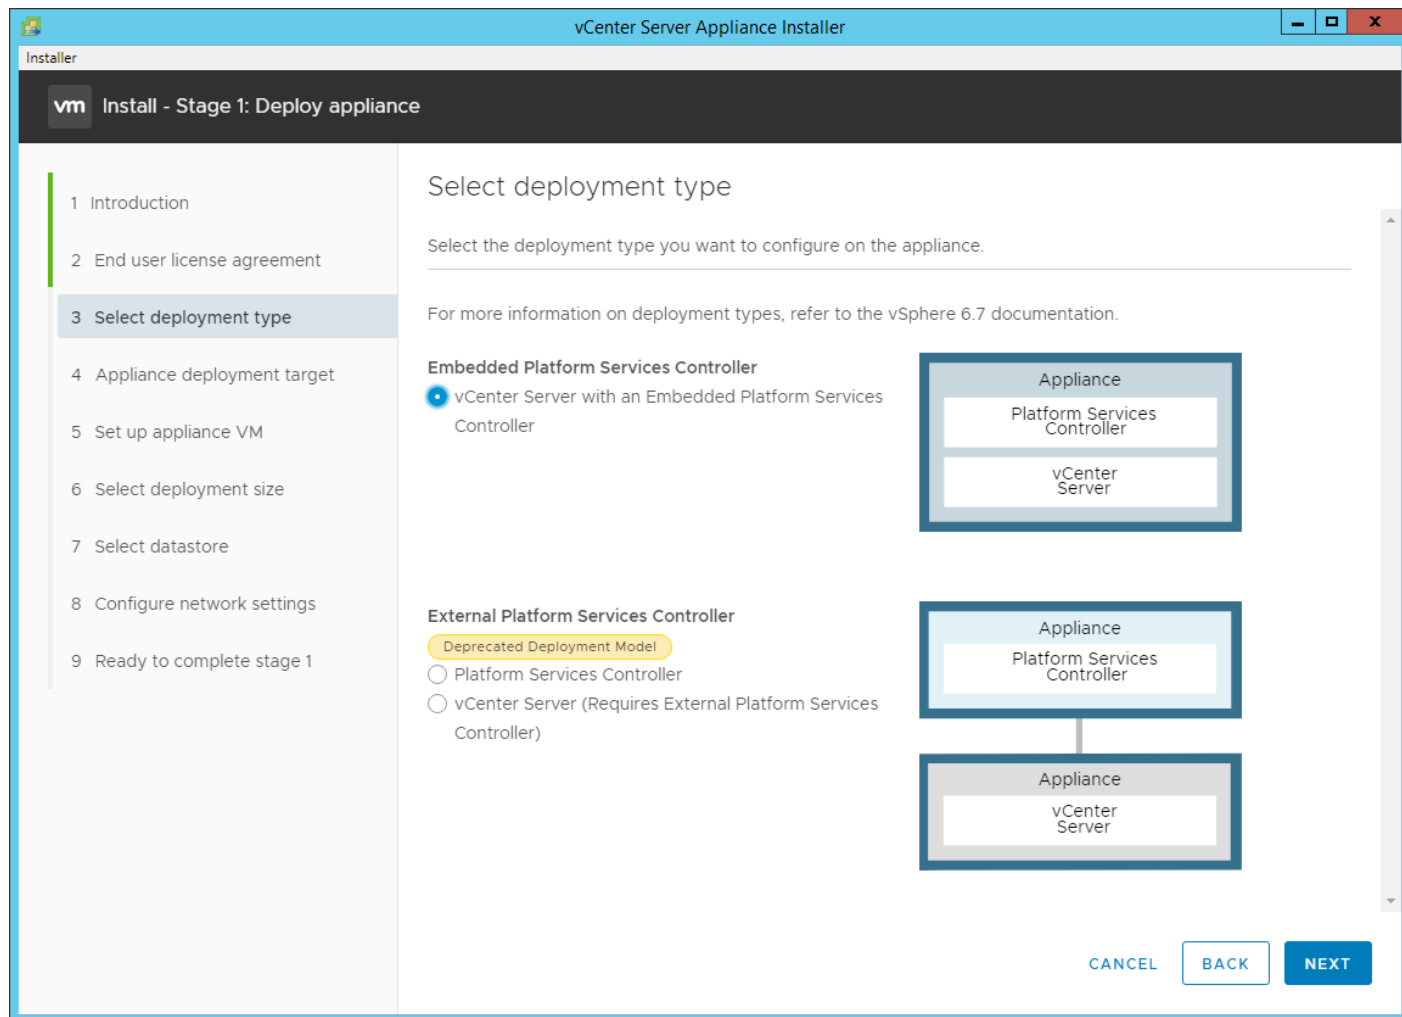
### Build the VMware vCenter Server Appliance

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

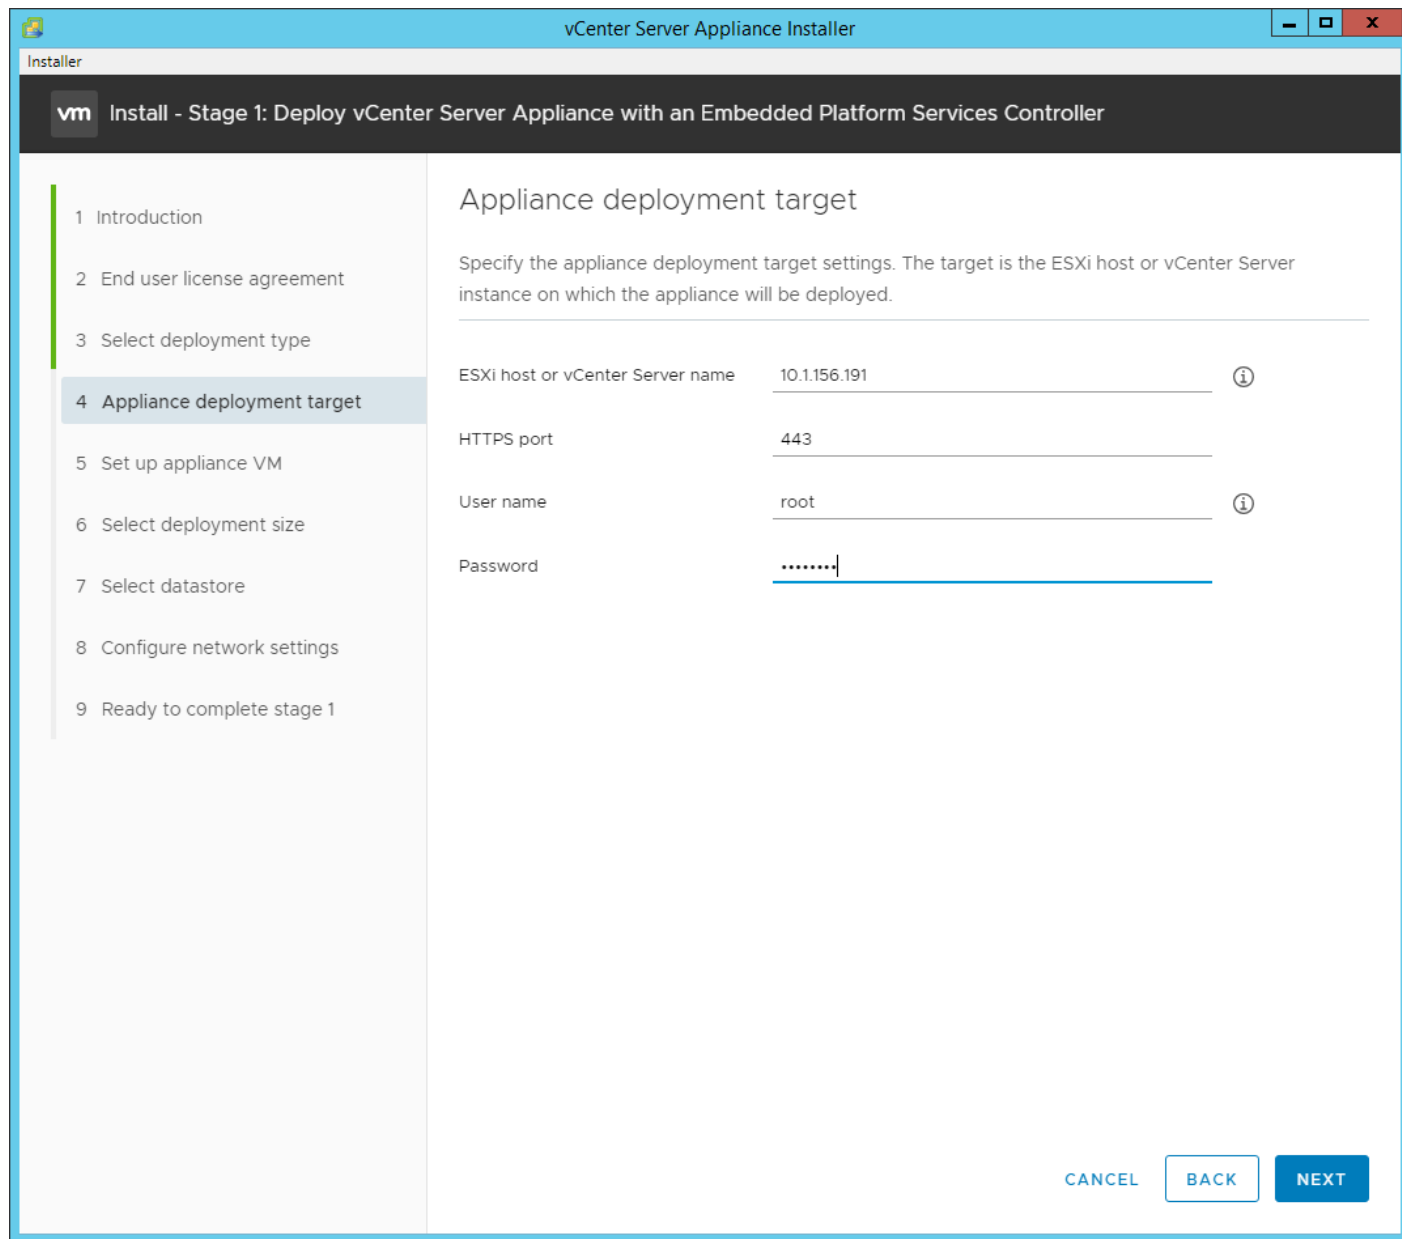
1. Locate and copy the `VMware-VCSA-all-6.7.0-15132721.iso` file to the desktop of the management workstation. This ISO is for the VMware vSphere 6.7 U3 vCenter Server Appliance.
2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the `Mount` command in Windows Server 2012).
3. In the mounted disk directory, navigate to the `vcsa-ui-installer > win32` directory and double-click `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, choose Embedded Platform Services Controller and click NEXT.

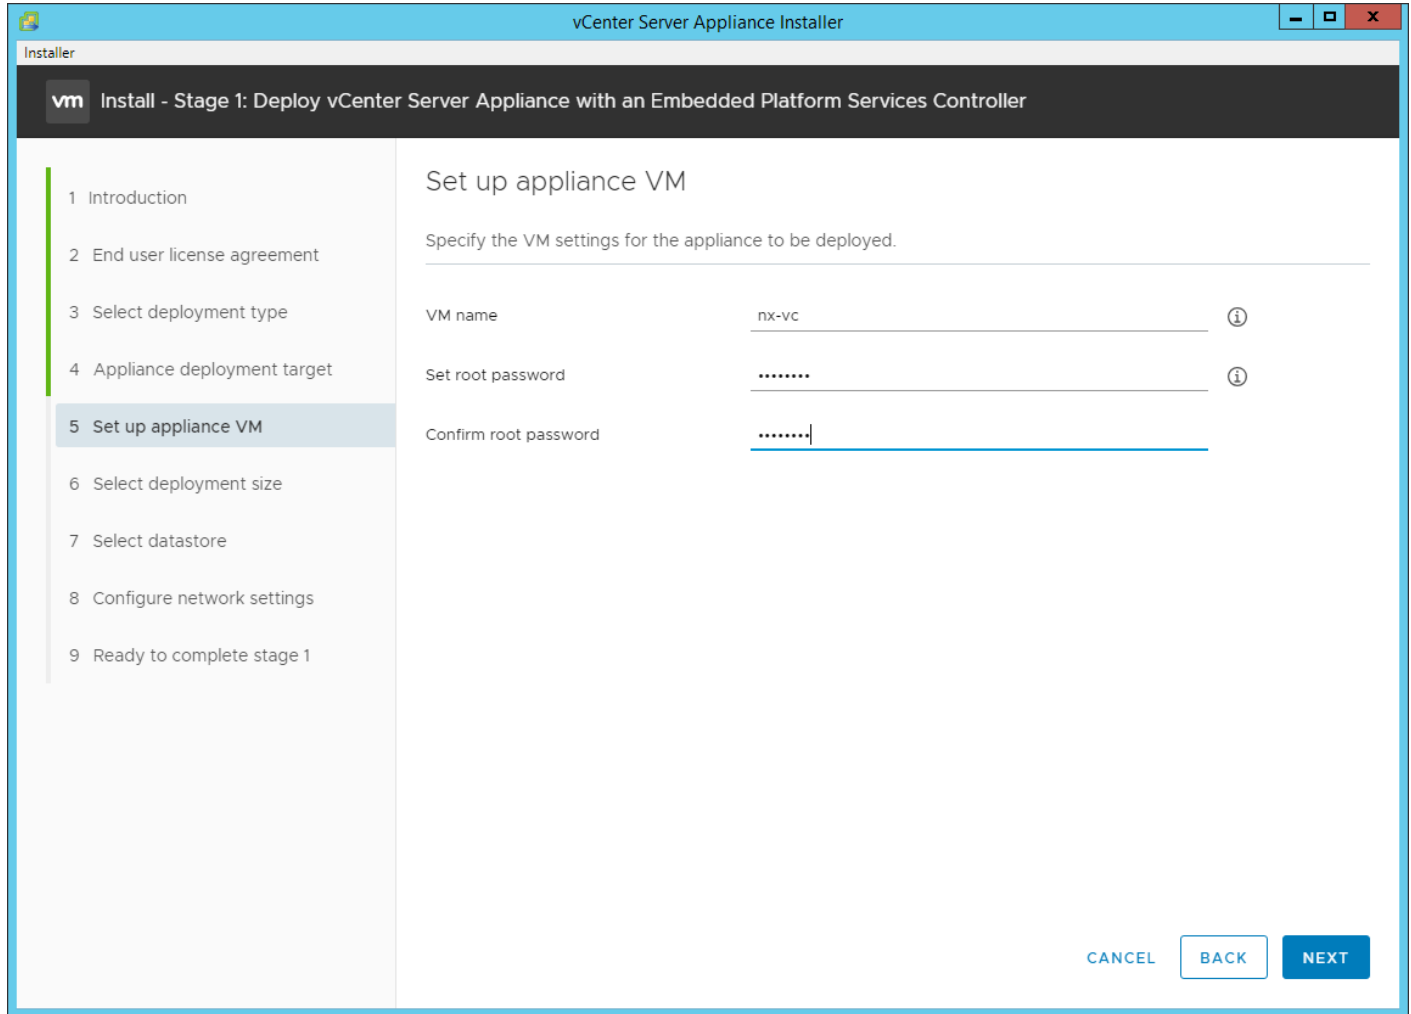


8. In the "Appliance deployment target", enter the host name or IP address of the first ESXi host, User name (root) and Password.



9. Click Yes to accept the certificate.

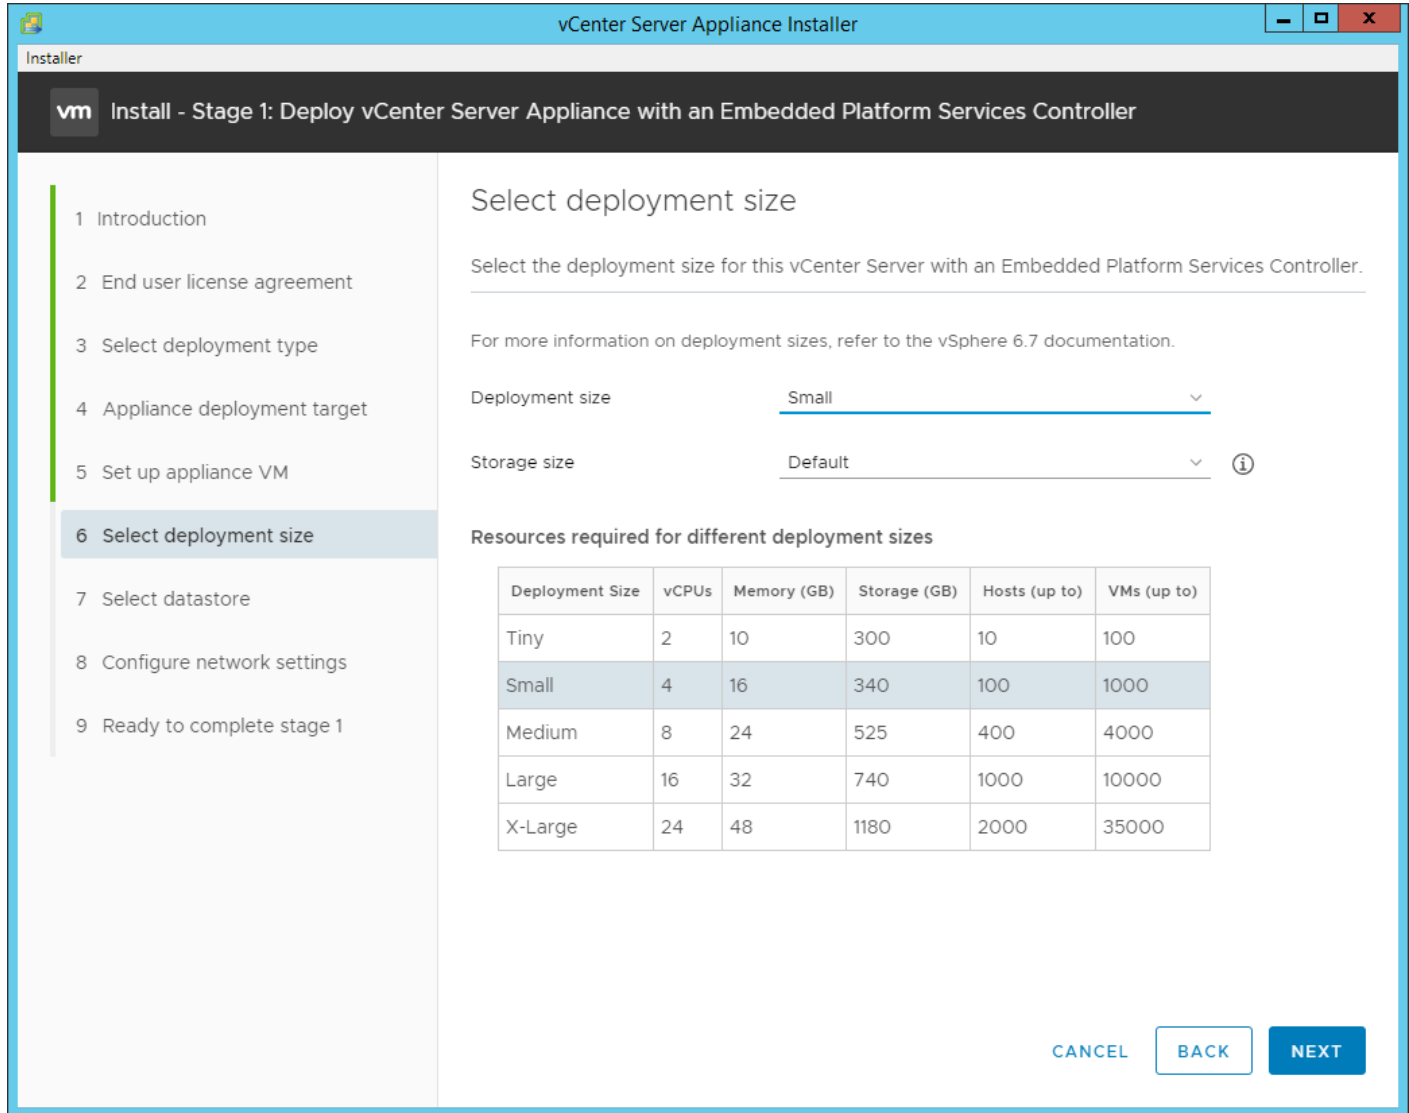
10. Enter the Appliance VM name and password details in the “Set up appliance VM” section. Click NEXT.



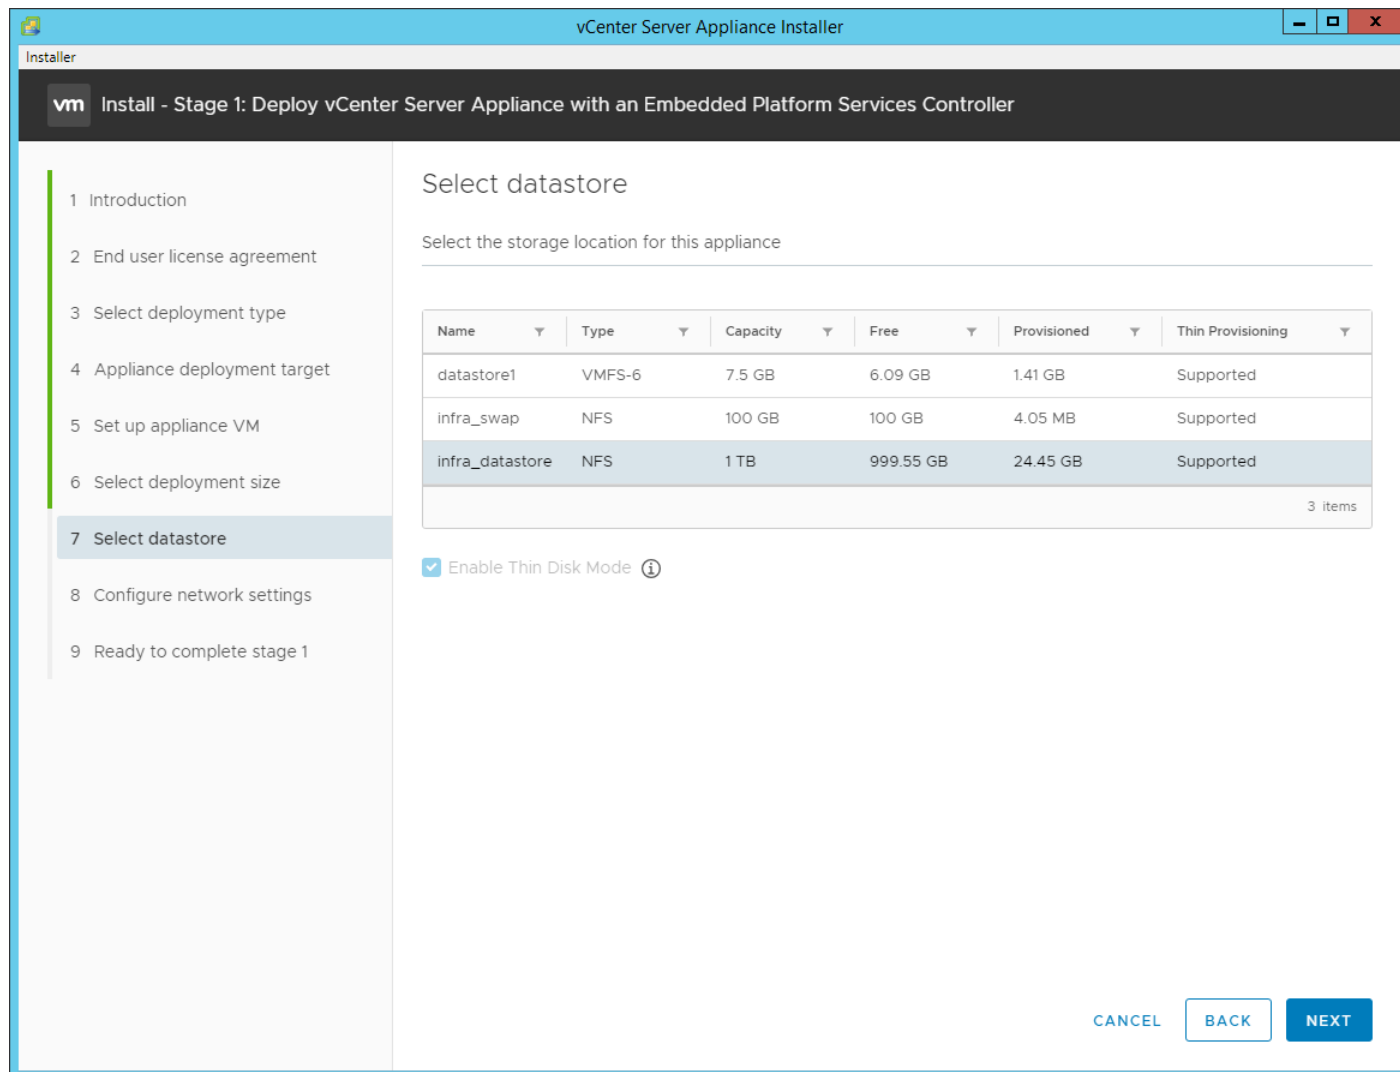
11. In the “Select deployment size” section, choose the Deployment size and Storage size. For example, choose “Small” and “Default”.

12. Click NEXT.





13. Choose infra\_datastore for storage. Click NEXT.



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

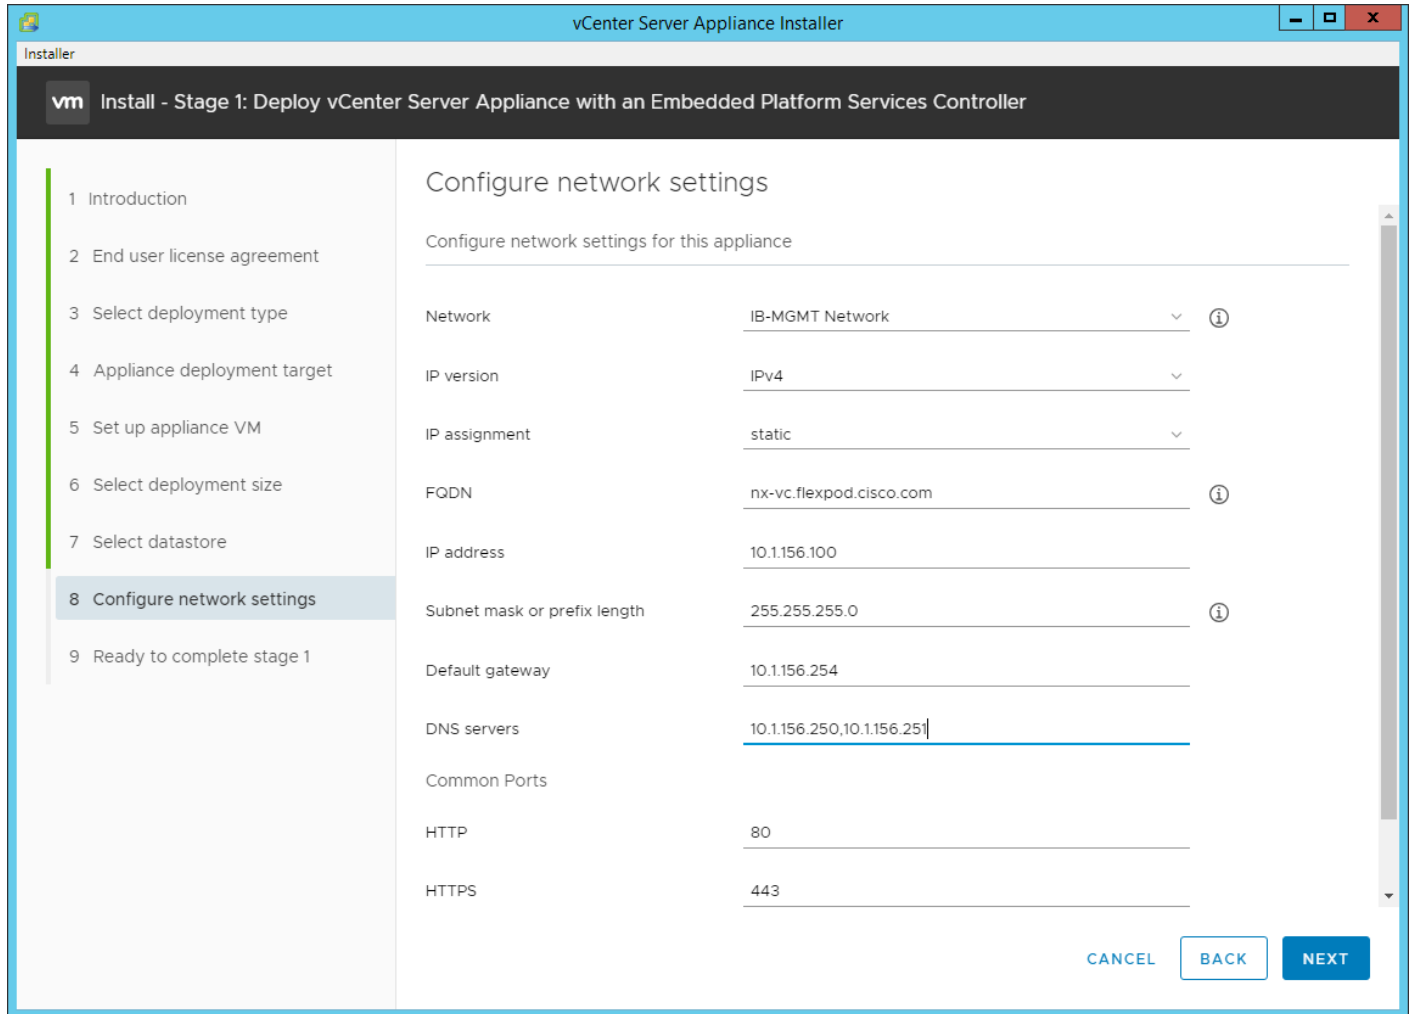
- a. Choose a Network: IB-MGMT Network.



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. FQDN: <vcenter-fqdn>

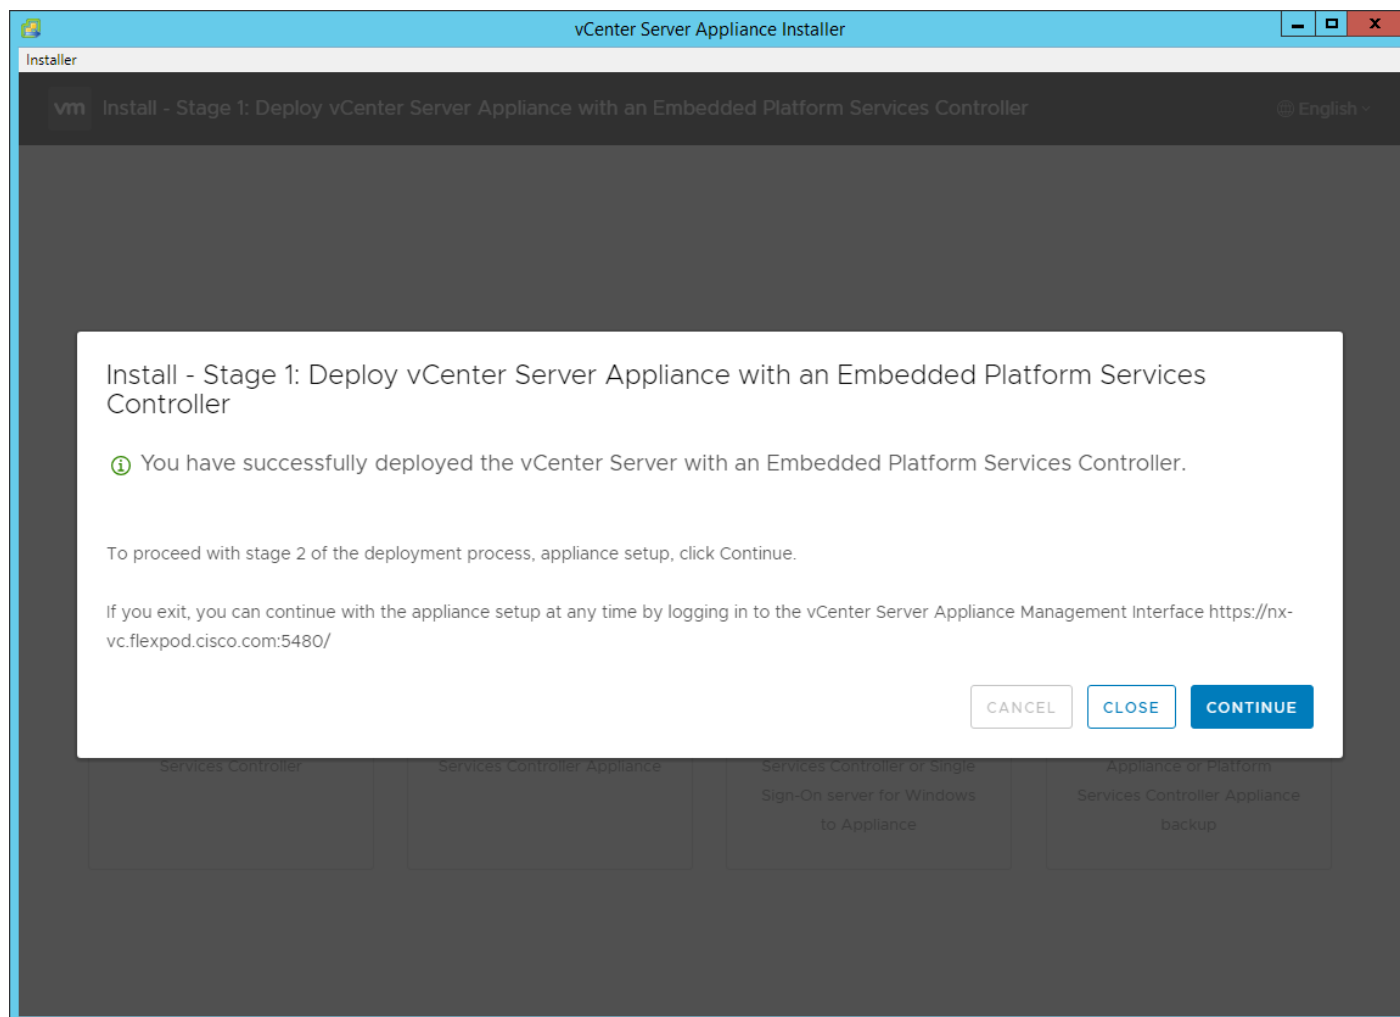
- e. IP address: <vcenter-ip>
- f. Subnet mask or prefix length: <vcenter-subnet-mask>
- g. Default gateway: <vcenter-gateway>
- h. DNS Servers: <dns-server1>, <dns-server2>



15. Click NEXT.

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

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

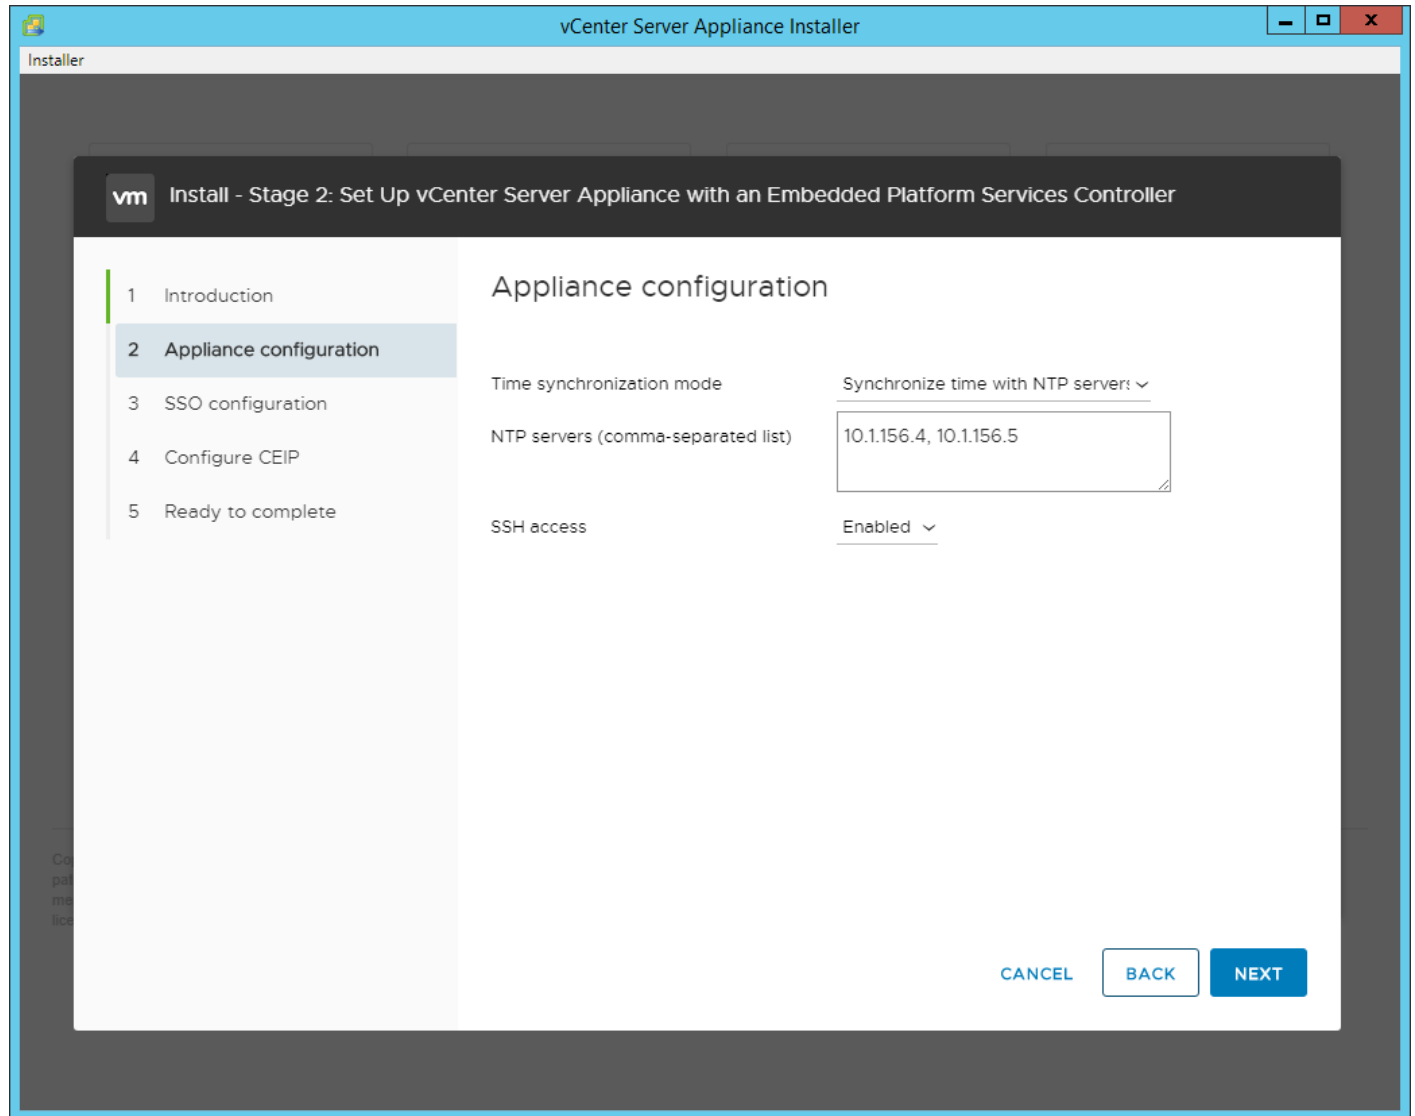


17. Click CONTINUE to proceed with stage 2 configuration.

18. Click NEXT.

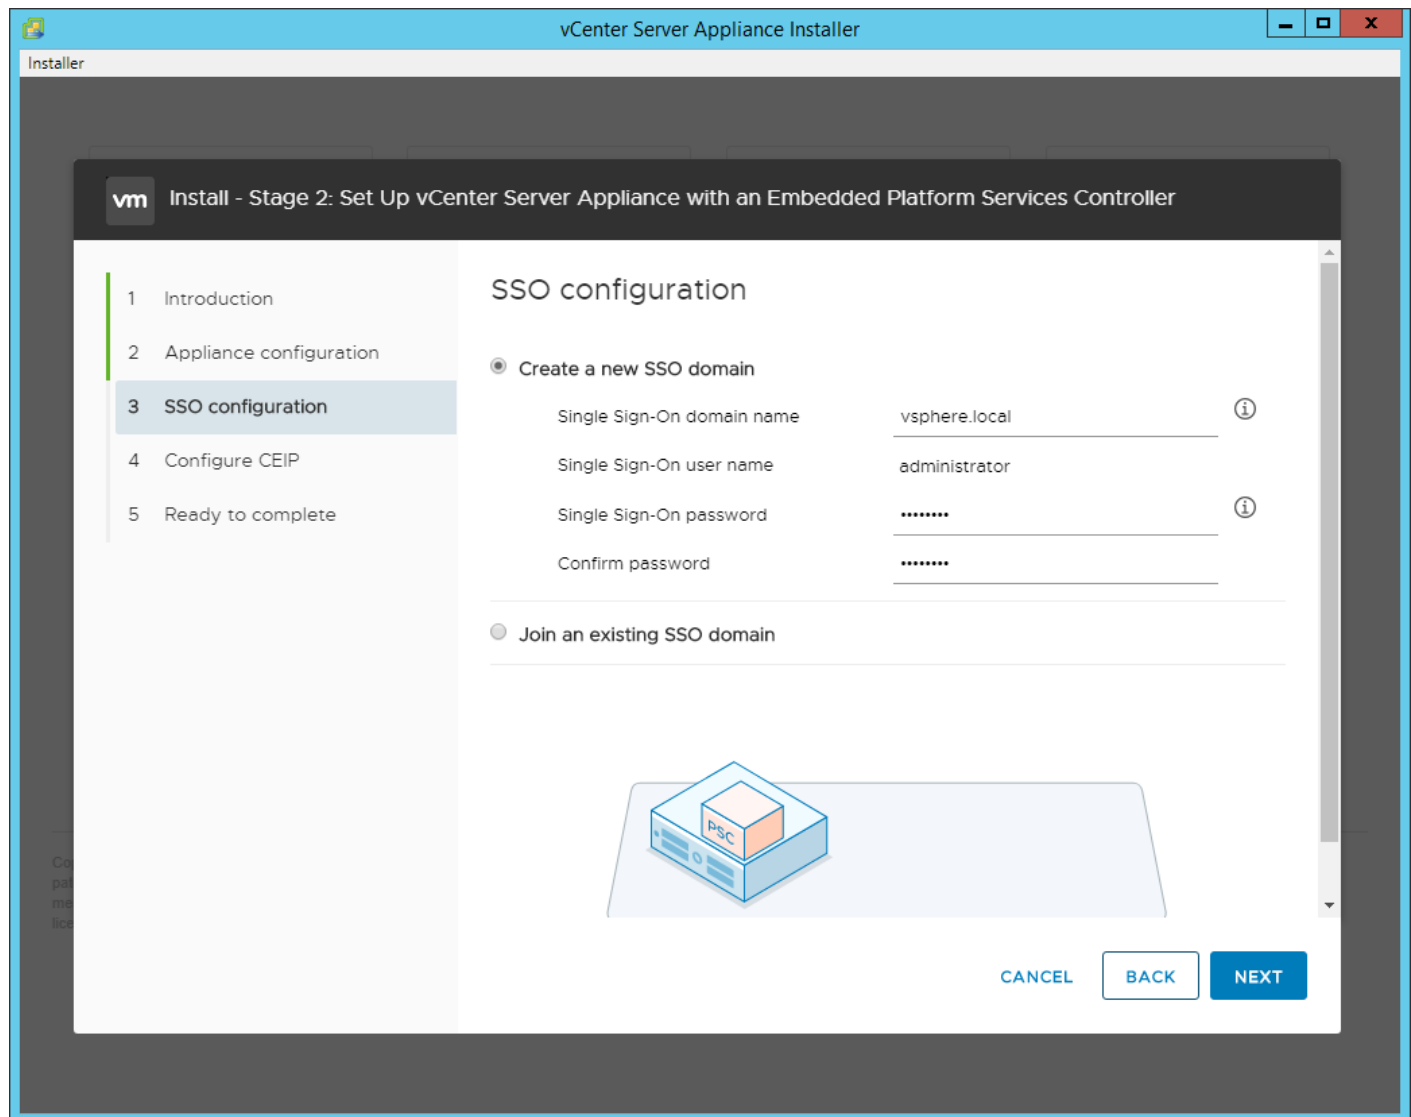
19. In the Appliance Configuration, configure these settings:

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



20. Click NEXT.

21. Complete the SSO configuration as shown below:



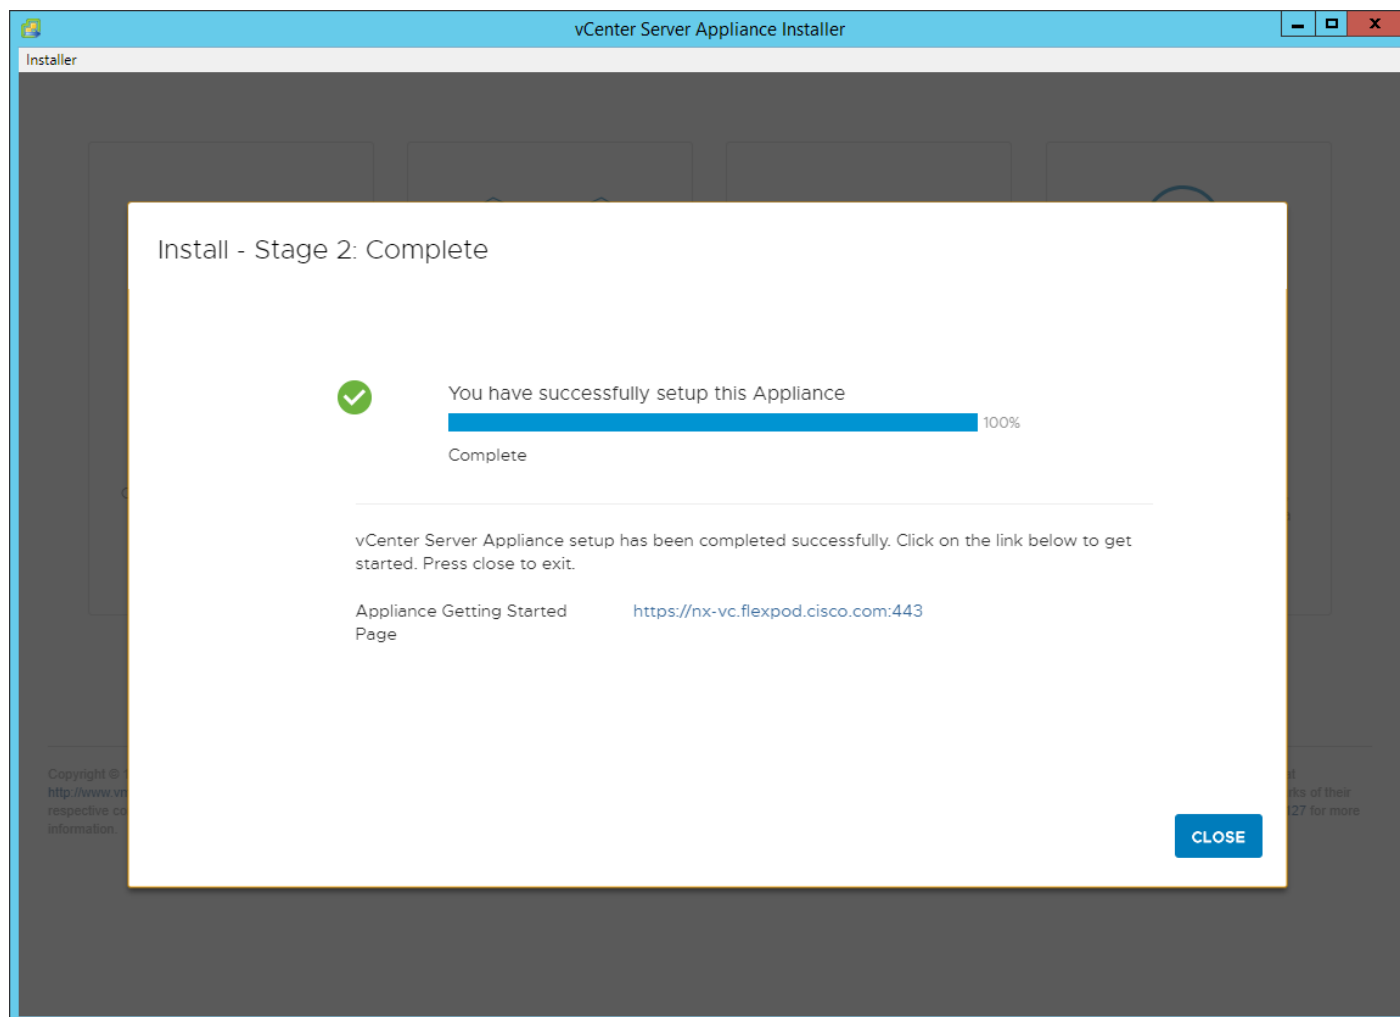
22. Click NEXT.

23. Decide whether to join VMware's Customer Experience Improvement Program (CEIP).

24. Click NEXT.

25. Review the configuration and click FINISH.

26. Click OK.



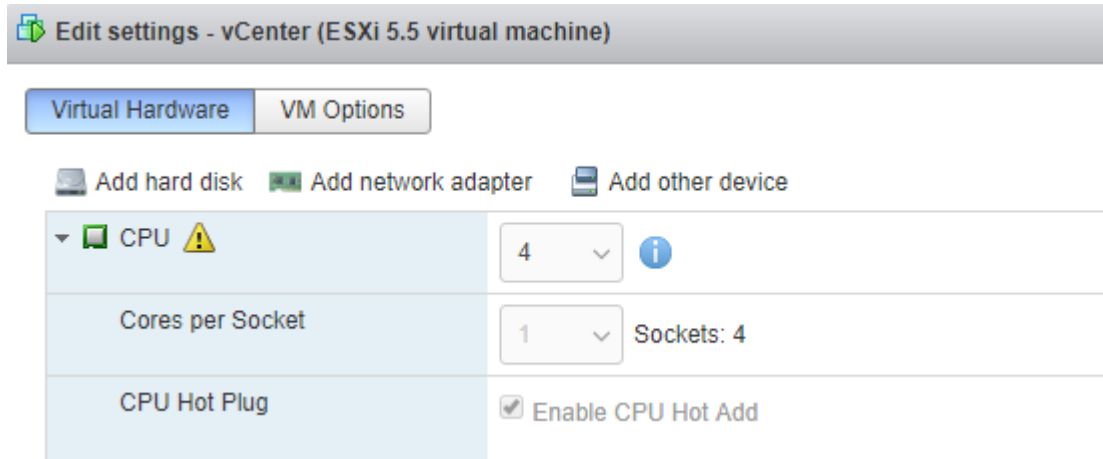
27. Click CLOSE. Eject or unmount the VCSA installer ISO.

## Adjust vCenter CPU Settings

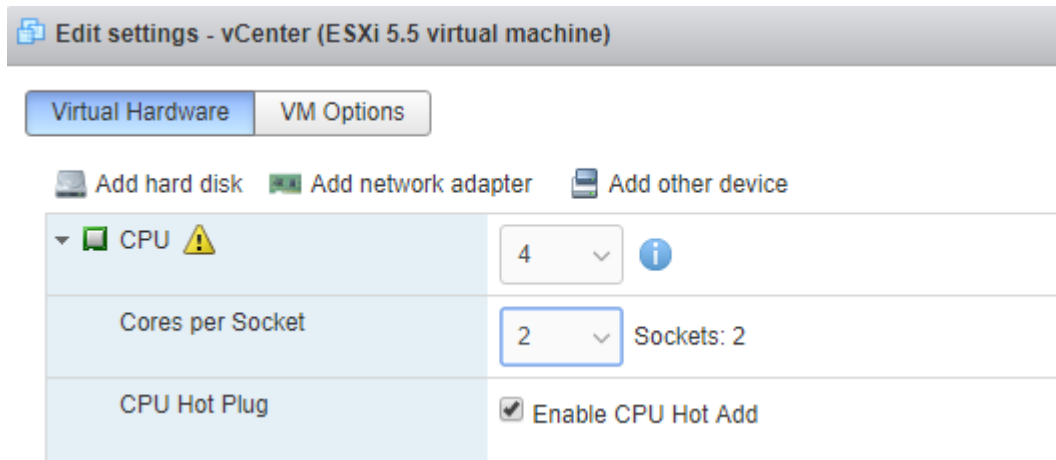
If a vCenter deployment size of Small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration. Cisco UCS B200 and C220 servers are 2-socket servers. In this validation, the Small deployment size was selected 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, choose Virtual Machines.
7. In the center pane, right-click the vCenter VM and choose 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 choose Guest OS > Shut down. Click Yes on the confirmation.
  - b. Once vCenter is shut down, right-click the vCenter VM and choose 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 choose 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:



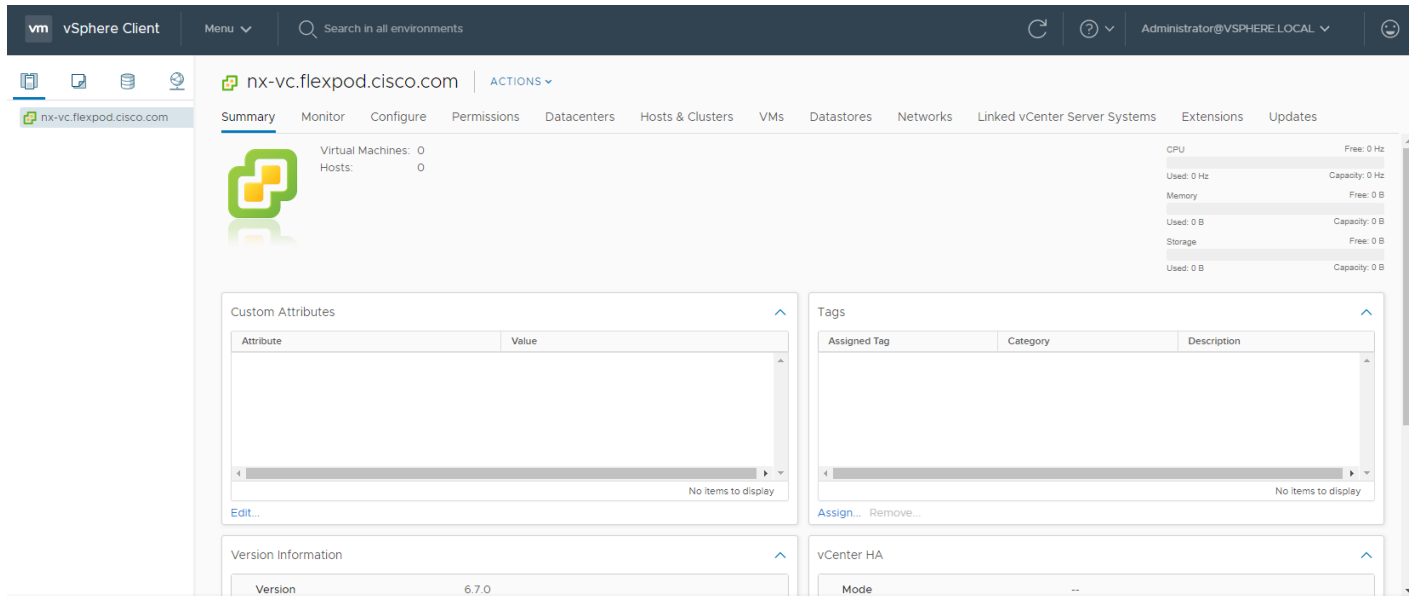
1. Using a web browser, navigate to <https://<vcenter-ip-address>:5480>. You will need to navigate security screens.
2. Log into the Appliance Management interface as root with the root password set in the vCenter installation.
3. In the menu on the left, choose Time.
4. Choose EDIT to the right of Time zone.
5. Choose the appropriate Time zone and click SAVE. The correct time should appear under Time synchronization.
6. In the menu on the left, choose Update.
7. Choose the latest vCenter update and choose STAGE AND INSTALL and follow the prompts to get to the latest version of vCenter 6.7 U3. You will need to click refresh to get back to the Appliance Management interface.
8. In the menu on the left choose Administration.
9. According to your Security Policy, adjust the settings for the root user and password.
10. In the upper right-hand corner of the screen, choose root > Logout to logout of the Appliance Management interface.
11. Using a web browser, navigate to <https://<vcenter-ip-address>>. You will need to navigate security screens.
12. Choose LAUNCH VSPHERE CLIENT (HTML5).



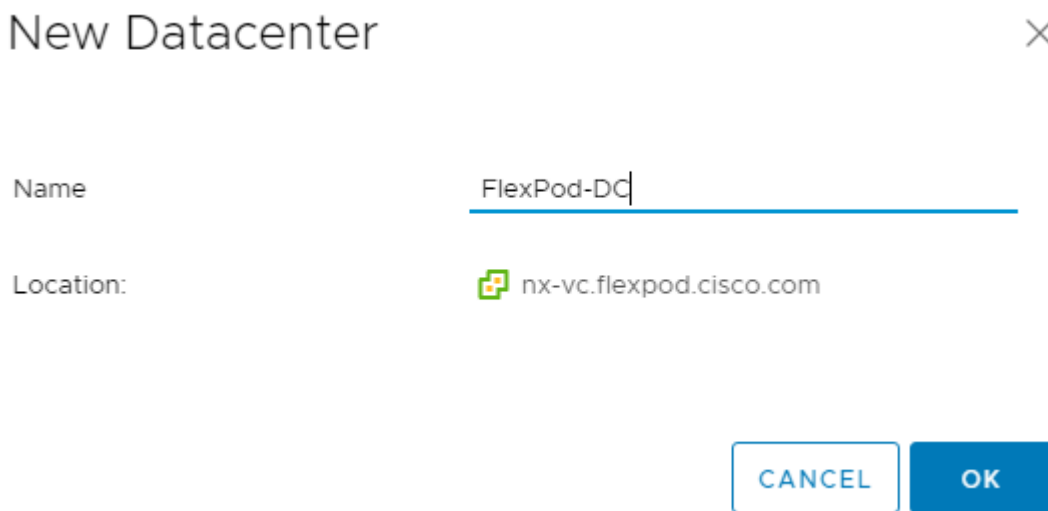
**Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is fully featured in vSphere 6.7U3 and will be used going forward.**

---

13. Log in using the Single Sign-On username ([administrator@vsphere.local](mailto:administrator@vsphere.local)) and password created during the vCenter installation.






14. In the center pane, choose ACTIONS > New Datacenter.
15. Type "FlexPod-DC" in the Datacenter name field.



16. Click OK.
17. Right-click the datacenter FlexPod-DC in the list in the left pane. Choose New Cluster.
18. Name the cluster FlexPod-Management.
19. Turn on DRS and vSphere HA. Do not turn on vSAN.

## New Cluster | FlexPod-DC ×

Name	<u>FlexPod-Management</u>
Location	 FlexPod-DC
 vSphere 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.

20. Click OK to create the new cluster.

21. Right-click "FlexPod-Management" and choose Settings.

22. Choose Configuration > General in the list located on the left and choose EDIT located on the right of General.

23. Choose Datastore specified by host and click OK.


## Edit Cluster Settings | FlexPod-Management ×

Virtual machine directory

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

Datastore specified by host

Store the swap files in the datastore specified by the host to be used for swap files. 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.

CANCEL

OK

24. Right-click "FlexPod-Management" and click Add Hosts.
25. 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.
26. In the Security Alert window, choose the host and click OK.
27. Verify the Host summary information and click NEXT.
28. Ignore warnings about the host being moved to Maintenance Mode and click FINISH to complete adding the host to the cluster.
29. The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.
30. In the list, right-click the added ESXi host and choose Settings.
31. In the center pane under Virtual Machines, choose Swap file location.
32. On the right, click EDIT.
33. Choose the infra\_swap datastore and click OK.

## Edit Swap File Location | nx-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
datastore1	7.50 GB	1.41 GB	6.09 GB	VMFS	Supported
Infra_datastore	1.00 TB	316.77 GB	1,024.00 GB	NFS41	Supported
Infra_swap	100.00 GB	3.62 MB	100.00 GB	NFS41	Supported

3 items

CANCEL

OK

34. In the list under Storage, choose Storage Devices. Make sure the NETAPP Fibre Channel Disk 0 is selected.
35. Choose the Paths tab.
36. 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 HDD Disk | Mark as Local

Name	LUN	Type	Capacity	Datastore	Operational St...	Hardware Accelerati...	Drive Ty...	Transport
NETAPP Fibre Channel Disk (naa.600a098038313...	0	disk	15.00 GB	datast...	Attached	Supported	Flash	Fibre Channel

Copy All | 1 items

Properties | **Paths** | Partition Details

Enable | Disable

Runtime Name	Status	Target	Name	Preferred
vmhba2:C0:T3:L0	Active	20:02:00:a0:98:e2:17:ca 20:04:00:a0:98:e2:17:ca	vmhba2:C0:T3:L0	
vmhba2:C0:T2:L0	Active (I/O)	20:02:00:a0:98:e2:17:ca 20:1d:00:a0:98:e2:17:ca	vmhba2:C0:T2:L0	
vmhba0:C0:T5:L0	Active	20:02:00:a0:98:e2:17:ca 20:03:00:a0:98:e2:17:ca	vmhba0:C0:T5:L0	
vmhba0:C0:T3:L0	Active (I/O)	20:02:00:a0:98:e2:17:ca 20:1c:00:a0:98:e2:17:ca	vmhba0:C0:T3: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 choose 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, choose Administration. In the list on the left, under Single Sign On, choose Configuration.
5. In the center pane, under Configuration, choose the Active Directory Domain tab.
6. Choose 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, choose System Configuration. Choose the radio button to choose the vCenter, then choose 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 as Administrator@vsphere.local.
11. Under Menu, choose Administration. In the list on the left, under Single Sign On, choose Configuration.
12. In the center pane, under Configuration, choose the Active Directory Domain tab.

13. Make sure your Active Directory Domain is listed.
14. Choose the Identity Sources tab.
15. Choose ADD IDENTITY SOURCE.
16. Choose 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, choose Global Permissions.
18. In the center pane, click the + sign to add a Global Permission.
19. In the Add Permission window, choose 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. Choose 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 in vCenter and on the first 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.

A 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 In-Band management interfaces including 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 should be 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, choose Networking under Menu.
2. Right-click the FlexPod-DC datacenter and choose 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 – ESXi 6.7 and later 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. Choose the newly created vDS.
8. Right-click the VM-Traffic port group and choose Edit Settings.
9. Choose VLAN on the left.
10. Choose VLAN for VLAN type and enter the VM-Traffic VLAN ID. Click OK.
11. Right-click the vDS and choose Settings > Edit Settings.
12. In the Edit Settings window, choose Advanced on the left.
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"/> ▾
<b>Discovery protocol</b>	
Type	<input type="text" value="Link Layer Discovery Protocol"/> ▾
Operation	<input type="text" value="Both"/> ▾
<b>Administrator contact</b>	
Name	<input type="text"/>
Other details	<input type="text"/>

14. For the vMotion port group, right-click the vDS, choose Distributed Port Group, and choose 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. Choose 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 ⓘ

↑ ↓

<b>Active uplinks</b>
Uplink 2
<b>Standby uplinks</b>
Uplink 1
<b>Unused uplinks</b>

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. Right-click the vDS and choose Add and Manage Hosts.
25. Make sure Add hosts is selected and click NEXT.
26. Click the green + sign to add New hosts. Choose the one configured FlexPod Management host and click OK. Click NEXT.
27. Choose vmnic2 and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 and click Assign uplink. Choose 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.

## vDSO - 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
nx-esxi-1.flexpod.cisco.com <ul style="list-style-type: none"> <li>On this switch               <ul style="list-style-type: none"> <li>  vmnic2 (Assigned)                   <span style="margin-left: 20px;">--</span> <span style="margin-left: 20px;">Uplink 1</span> <span style="margin-left: 20px;">vDSO-DVUplinks-...</span> </li> <li style="background-color: #e0e0e0;">  vmnic3 (Assigned)                   <span style="margin-left: 20px;">--</span> <span style="margin-left: 20px;">Uplink 2</span> <span style="margin-left: 20px;">vDSO-DVUplinks-...</span> </li> </ul> </li> <li>On other switches/unclaimed               <ul style="list-style-type: none"> <li>  vmnic0                   <span style="margin-left: 20px;">vSwitch0</span> <span style="margin-left: 20px;">--</span> <span style="margin-left: 20px;">--</span> </li> <li>  vmnic1                   <span style="margin-left: 20px;">vSwitch0</span> <span style="margin-left: 20px;">--</span> <span style="margin-left: 20px;">--</span> </li> </ul> </li> </ul>			

CANCEL

BACK

NEXT

28. Click NEXT.

29. Choose vmk2 (VMkernel vMotion) and click Assign port group.

30. Choose the vMotion destination port group and click OK.

31. Do not migrate the other VMkernel ports.

## vDSO - 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...
nx-esxi-1.flexpod.cisco.com			
On this switch			
vmk2 (Reassigned)	vSwitch0	VMkernel-vMotion	vMotion
On other switches/unclaimed			
vmk0	vSwitch0	Management Net...	Do not migrate
vmk1	vSwitch0	VMkernel-Infra-NFS	Do not migrate

CANCEL

BACK

NEXT

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

33. Do not migrate any virtual machine networking ports. Click NEXT.

34. 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 6.7U3 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, choose 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. Choose 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 be placed in Maintenance Mode and 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 choose the ESXi host.
2. In the center pane, choose the Configure tab.
3. In the list, choose Virtual switches under Networking.
4. Expand vSwitch0.

5. Choose EDIT to Edit settings.
6. Change the MTU to 9000.
7. Choose 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

**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: vmnic0, vmnic1

Standby adapters:

Unused adapters:

Adapter Name: Cisco Systems Inc Cisco VIC Ethernet NI  
 Location: PCI 0000:62:00.1  
 Driver: nenic

**Status**

Status: Connected  
 Actual speed, Duplex: 20 Gbit/s, Full Duplex  
 Configured speed, Duplex: 20 Gbit/s, Full Duplex  
 Networks: 10.1.156.1-10.1.156.254 ( VLAN113 )

**SR-IOV**


Status: Not supported

Select active and standby adapters. During a failover, standby adapters activate in the order specified above.

CANCEL OK

9. Click OK.
10. In the center pane, to the right of VM Network 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, choose VMkernel adapters.
14. In the center pane, click Add Networking.
15. Make sure VMkernel Network Adapter is selected and click NEXT.
16. Choose an existing standard switch and click BROWSE. Choose 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. Choose Custom for MTU and make sure 9000 is entered.
20. Leave the Default TCP/IP stack selected and do not choose any of the Enabled services. Click NEXT.
21. Choose 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, choose Virtual switches. Then expand vSwitch0. The properties for vSwitch0 should be similar to the following example:



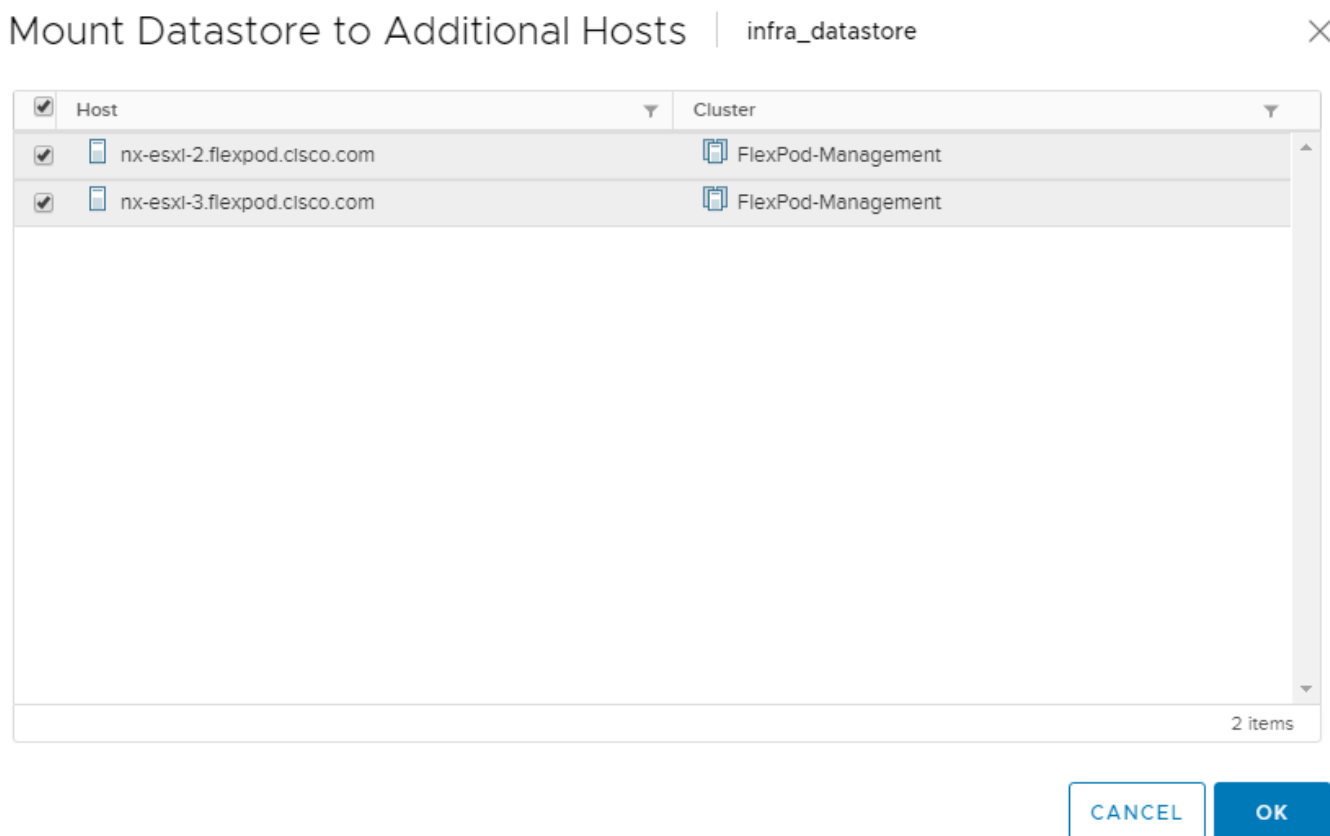
25. Repeat this procedure for all hosts being added.

## 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, choose Menu > Storage.
2. Located on the left, expand FlexPod-DC.
3. Located on the left, right-click infra\_datastore and choose Mount Datastore to Additional Hosts.
4. Choose the ESXi host(s) and click OK.



5. Repeat steps 1-4 to mount the infra\_swap datastore to the ESXi host(s).
6. Choose infra\_datastore. In the center pane, choose 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 choose the ESXi host.
2. In the center pane, choose the Configure tab.
3. In the list under System, choose Time Configuration.
4. Located on the right, click EDIT.
5. Choose 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 choose 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

Start NTP Service

**NTP Service Startup Policy:**

Start and stop with host ▼

CANCEL
OK

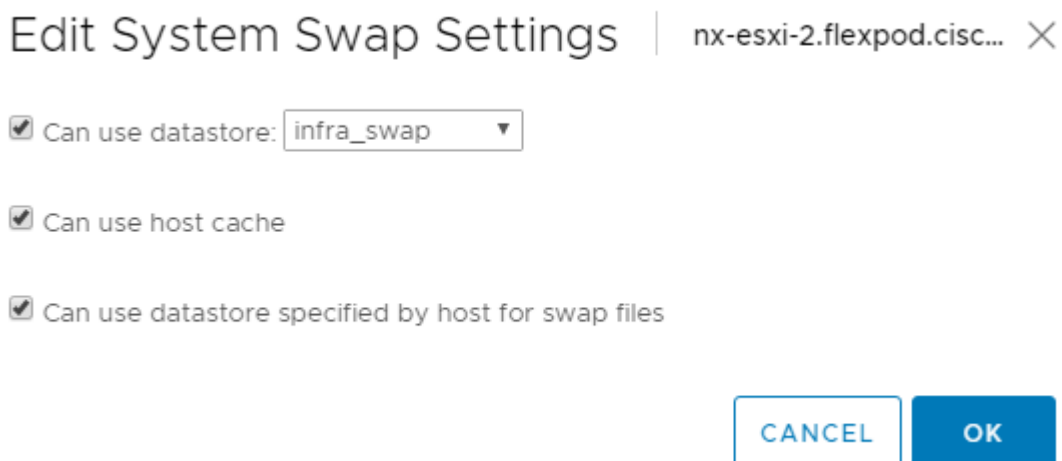
9. Click OK to save the configuration changes.
10. Verify that NTP service is now enabled and running and the clock is now set to approximately the correct time.

## 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 choose the ESXi host.
2. In the center pane, choose the Configure tab.
3. In the list under System, choose System Swap.
4. Located on the right, click EDIT.
5. Choose Can use datastore and use the drop-down list to choose infra\_swap. Leave all other settings unchanged.



6. Click OK to save the configuration changes.
7. In the list under Virtual Machines, choose Swap File Location.
8. Located on the right, click Edit.
9. Choose infra\_swap and click OK.

## Check ESXi Host Fibre Channel Pathing

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

For fibre channel SAN-booted ESXi hosts, to ensure that the host(s) boot disk contains all required fibre channel paths, follow these steps:

1. In the list under Storage, choose Storage Devices. Make sure the NETAPP Fibre Channel Disk is selected.
2. Choose 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 HDD Disk | Mark as Local

Name	LUN	Type	Capacity	Datastore	Operational St...	Hardware Accelerati...	Drive Ty...	Transport
NETAPP Fibre Channel Disk (naa.600a098038313...	0	disk	15.00 GB	datast...	Attached	Supported	Flash	Fibre Channel

Copy All | 1 items

Properties | **Paths** | Partition Details

Enable | Disable

Runtime Name	Status	Target	Name	Preferred
vmhba2:C0:T3:L0	Active (I/O)	20:02:00:a0:98:e2:17:ca 20:04:00:a0:98:e2:17:ca	vmhba2:C0:T3:L0	
vmhba2:C0:T2:L0	Active	20:02:00:a0:98:e2:17:ca 20:1d:00:a0:98:e2:17:ca	vmhba2:C0:T2:L0	
vmhba0:C0:T5:L0	Active (I/O)	20:02:00:a0:98:e2:17:ca 20:03:00:a0:98:e2:17:ca	vmhba0:C0:T5:L0	
vmhba0:C0:T3:L0	Active	20:02:00:a0:98:e2:17:ca 20:1c:00:a0:98:e2:17:ca	vmhba0:C0:T3:L0	

### 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, choose Networking under Menu.
2. Right-click the vDS (vDS0) and click Add and Manage Hosts.
3. Make sure Add hosts is selected and click NEXT.
4. Click the green + sign to add New hosts. Choose the configured FlexPod Management host(s) and click OK. Click NEXT.
5. Choose vmnic2 on each host and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 on each host and click Assign uplink. Choose 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.**







## vDSO - 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
nx-esxi-2.flexpod.cisco.com			
On this switch			
 vmnic2 (Assigned)	--	Uplink 1	vDSO-DVUplinks...
 vmnic3 (Assigned)	--	Uplink 2	vDSO-DVUplinks...
On other switches/unclaimed			
 vmnic0	vSwitch0	--	--
 vmnic1	vSwitch0	--	--
nx-esxi-3.flexpod.cisco.com			
On this switch			
 vmnic2 (Assigned)	--	Uplink 1	vDSO-DVUplinks...
 vmnic3 (Assigned)	--	Uplink 2	vDSO-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 choose the ESXi host.
2. In the center pane, click the Configure tab.
3. In the list under Networking, choose VMkernel adapters.
4. Choose Add Networking to Add host networking.
5. Make sure VMkernel Network Adapter is selected and click NEXT.

6. Choose BROWSE to the right of Select an existing network.
7. Choose vMotion on the vDS and click OK.
8. Click NEXT.
9. Make sure the Network label is vMotion with the vDS in parenthesis. From the drop-down list, select Custom for MTU and make sure the MTU is set to 9000. Choose the vMotion TCP/IP stack and click NEXT.
10. Choose 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.

## nx-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	vDS0
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	192.168.100.22 (static)
Subnet mask	255.255.255.0

CANCEL

BACK

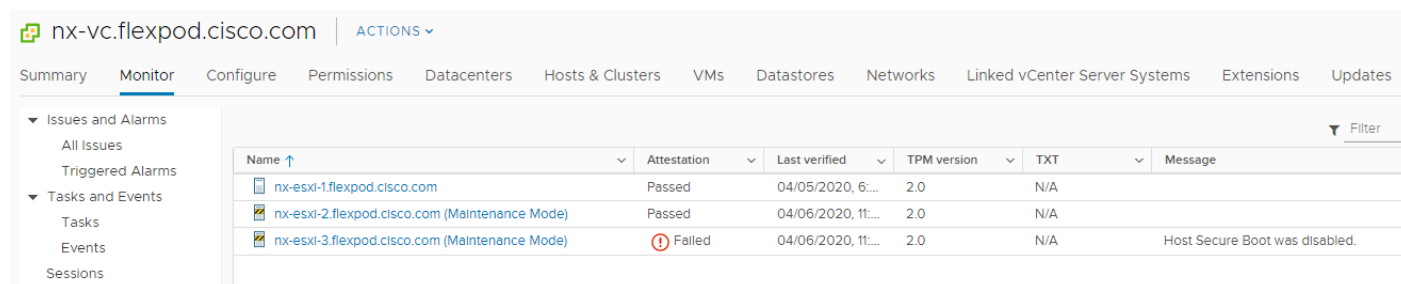
FINISH

13. If NetApp VSC is installed, under Hosts and Clusters, right-click the host and click NetApp VSC > Set Recommended Values. Reboot the host.
14. If this is an iSCSI-booted hosts, execute the instructions in the Appendix for an iSCSI-booted host being added in vCenter.
15. Exit Maintenance Mode on each ESXi host in Maintenance Mode.

## VMware ESXi 6.7U3 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, click vCenter. In the center pane, click Monitor > Security. The Attestation status will appear as shown below:



Name	Attestation	Last verified	TPM version	TXT	Message
nx-esxi-1.flexpod.cisco.com	Passed	04/05/2020, 6:...	2.0	N/A	
nx-esxi-2.flexpod.cisco.com (Maintenance Mode)	Passed	04/06/2020, 11:...	2.0	N/A	
nx-esxi-3.flexpod.cisco.com (Maintenance Mode)	Failed	04/06/2020, 11:...	2.0	N/A	Host Secure Boot was disabled.



It may be necessary to disconnect and reconnect a host from vCenter to get it to pass attestation the first time. Also, in this example, the third host was iSCSI booted and did not use UEFI Secure Boot.

## FlexPod Management Tools Setup

### NetApp Virtual Storage Console 9.7 Deployment Procedure

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



During testing and validation for VMware Virtual Volumes with Virtual Storage Console 9.7, an issue was encountered with a FlexGroup volume configured on the storage. While FlexGroups are not currently supported by VSC 9.7, the presence of a FlexGroup volume on a different SVM should not have interfered with creating Virtual Volumes on a second SVM. This issue has been documented as NetApp Bug #1317261 and will be fixed in Virtual Storage Console 9.7P2. Due to this issue, the configuration steps for virtual volumes have been removed from this guide until the issue has been fixed in a publicly available release of VSC 9.7.

### Virtual Storage Console 9.7 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.7:

- Protocol licenses (NFS)
- NetApp FlexClone® (for provisioning and cloning and vVol)
- NetApp SnapRestore® (for backup and recovery)
- The NetApp SnapManager® Suite
- SnapMirror or SnapVault



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.

**Table 8 Port Requirements for VSC**

Port	Requirement
443 (HTTPS)	Secure communications between VMware vCenter Server and the storage systems
8143 (HTTPS)	VSC listens for secure communications
9083 (HTTPS)	VASA Provider uses this port to communicate with the vCenter Server and obtain TCP/IP settings

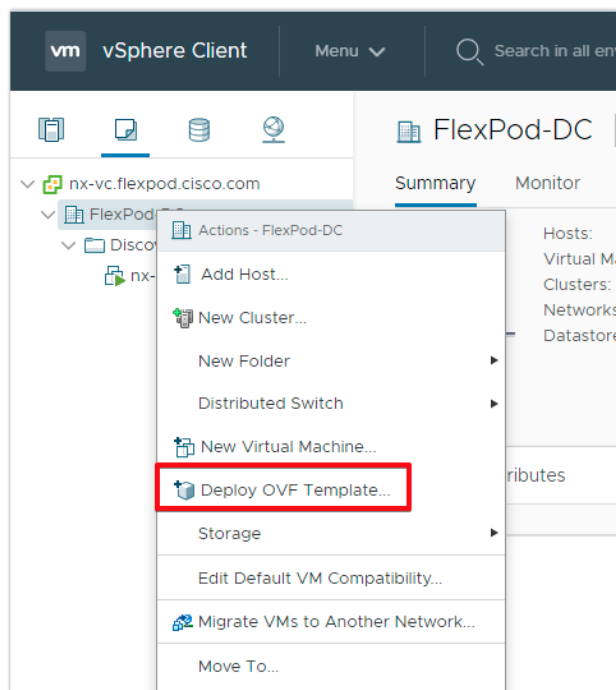
The requirements for deploying VSC are listed [here](#).

### Install Virtual Storage Console 9.7

To install the VSC 9.7 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

1. Launch the vSphere Web Client and navigate to Hosts and Clusters.

2. Right-click the FlexPod-DC datacenter and choose Deploy OVF Template.



3. Browse to the VSC OVA file downloaded from the NetApp Support site.
4. Enter the VM name and choose a datacenter or folder in which to deploy and click NEXT.
5. Choose a host cluster resource in which to deploy OVA and click NEXT.
6. Review the details and accept the license agreement.
7. Choose the infra\_datastore volume and choose the Thin Provision option for the virtual disk format.



### 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: Thin Provision

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type
datastore1	7.5 GB	1.41 GB	6.09 GB	VM
datastore1 (1)	7.5 GB	1.41 GB	6.09 GB	VM
datastore1 (2)	7.5 GB	1.41 GB	6.09 GB	VM
Infra_datastore	1 TB	331.21 GB	1,000.13 GB	NF
Infra_swap	100 GB	9.16 MB	99.99 GB	NF

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

8. From Select Networks, choose a destination network (IB-MGMT) and click NEXT.
9. From Customize Template, enter the VSC administrator password, vCenter name or IP address and other configuration details and click NEXT.

### Deploy OVF Template

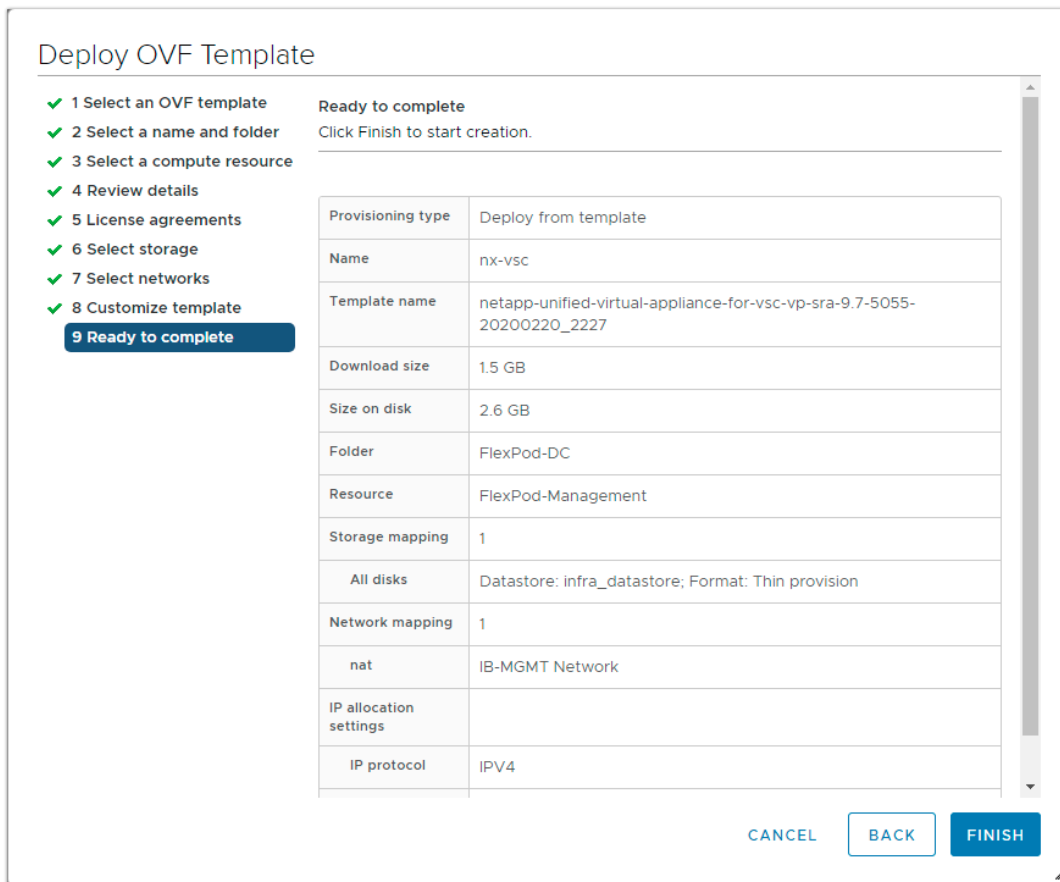
- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

Password		.....
Confirm Password		.....
vCenter Registration Configuration 4 settings		
vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to.	<u>nx-vc.flexpod.cisco.com</u>
Port (*)	Specify the HTTPS port of an existing vCenter to register to.	<u>443</u>
Username (*)	Specify the username of an existing vCenter to register to.	<u>administrator@vsphere.local</u>
Password (*)	Specify the password to register to.	<input type="password"/>
	Confirm Password	<input type="password"/>
Network Properties 8 settings		

**Enter a password to enable authentication.**

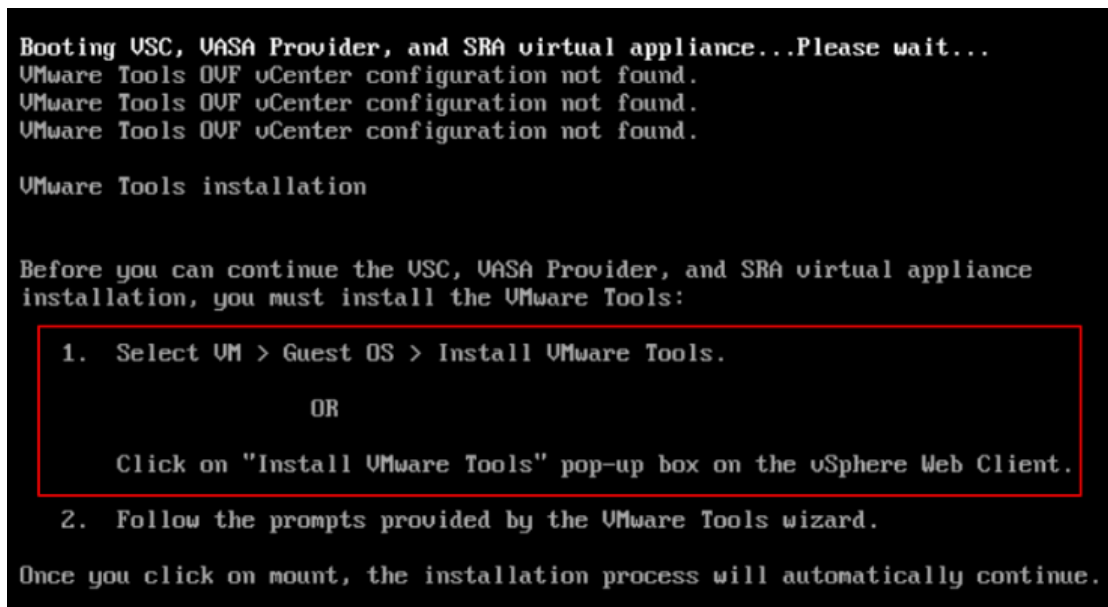
[CANCEL](#) [BACK](#) [NEXT](#)

10. Review the configuration details entered and click FINISH to complete the deployment of NetApp-VSC VM.



11. Power on the NetApp-VSC VM and open the VM console.

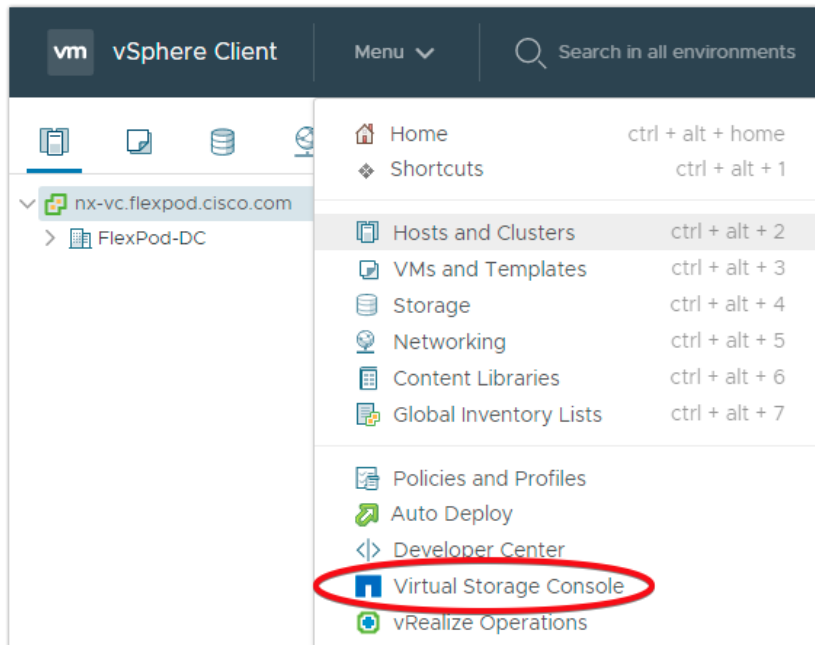
12. During the NetApp-VSC VM boot process, you see a prompt to install VMware Tools. From vCenter, right-click the NetApp-VSC VM > Guest OS > Install VMware Tools.



13. Networking configuration and vCenter registration information was provided during the OVF template customization, therefore after the VM is running, VSC and vSphere API for Storage Awareness (VASA) is registered with vCenter.
14. From the Home screen confirm that the NetApp VSC is installed.



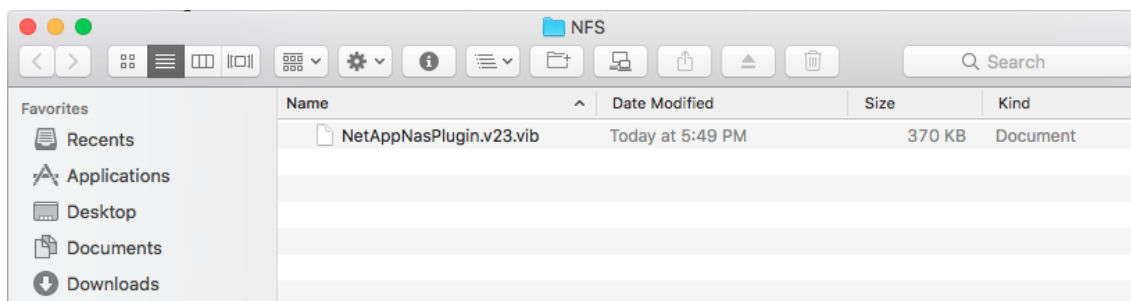
**The NetApp VSC 9.7 vCenter Plugin is only available in the vSphere HTML5 Client and is not available in the vSphere Web Client.**



## 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 Getting Started page.
3. Click NFS VAAI Tools tab.
4. Click Change in the Existing version section.
5. Browse and choose the renamed .vib file, and then click Upload to upload the file to the virtual appliance.
6. In the Install on ESXi Hosts section, choose 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.

### Verify the VASA Provider

The VASA provider for ONTAP is enabled by default during the installation of the NetApp Virtual Storage Console (VSC) 9.7. To verify the VASA provider was enabled, perform the following 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 if not enabled, click Enable VASA Provider slider.
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:

Username:

Password: \*

## 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 the tab and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.



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



**VSC cannot manage the storage system from the cluster admin level if a network interface with the fc-nvme data protocol is enabled. This issue is documented in [NetApp Bug ID 1302377](#) and is fixed in VSC 9.7P1.**

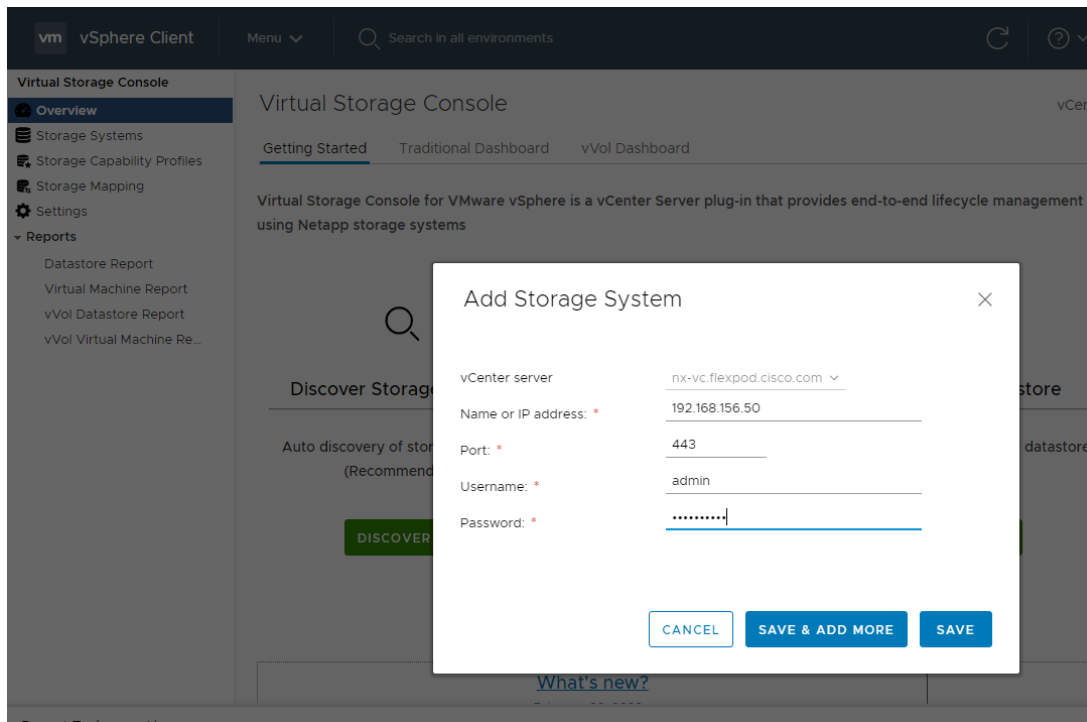


**You can modify the storage credentials with the vsadmin account or another SVM level account with RBAC privileges. Refer to the [ONTAP 9 Administrator Authentication and RBAC Power Guide](#) for additional information.**

---

3. Choose Storage Systems >Add
4. Click Overview > Getting Started, and then click ADD button under Add Storage System.

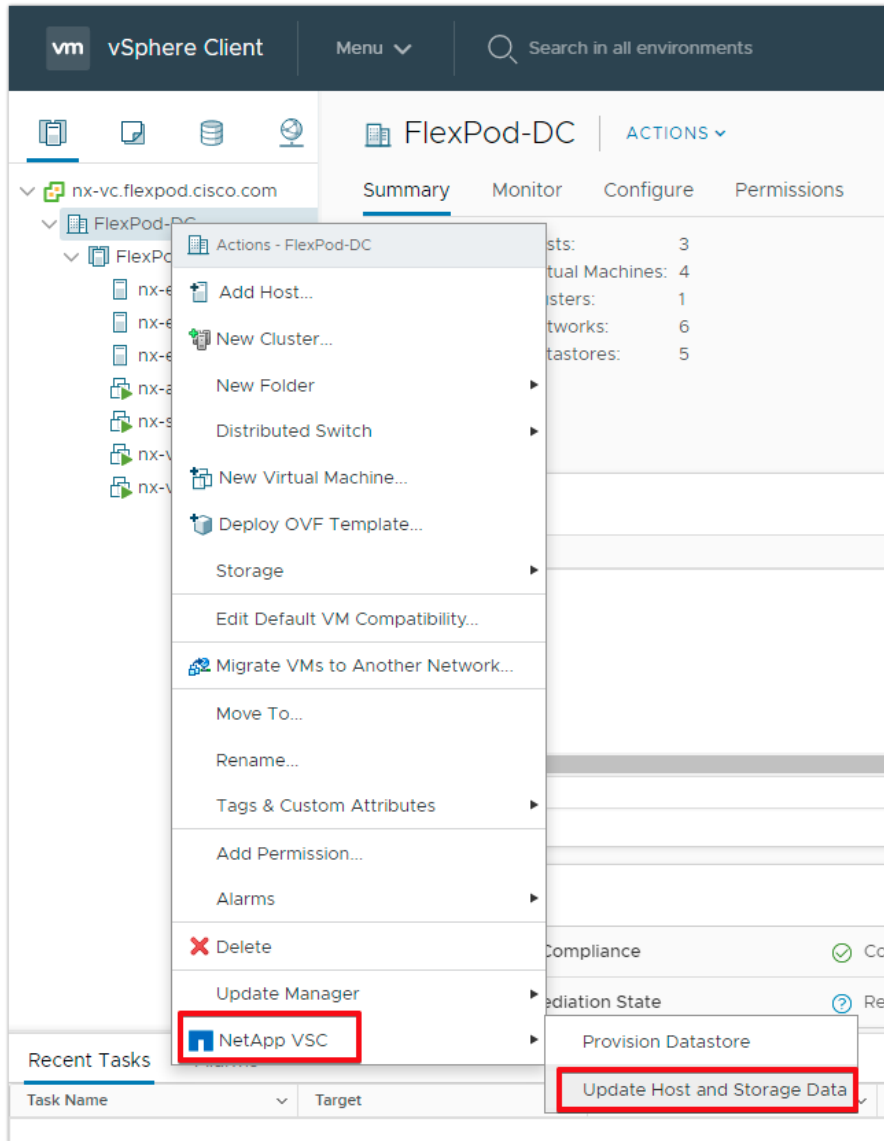
5. Specify the vCenter Server instance where the storage will be located.
6. In the IP Address/Hostname field, enter the storage cluster management IP.
7. Confirm Port 443 to Connect to this storage system.
8. Enter admin for the user name and the admin password for the cluster.
9. Click Save to add 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, click Hosts and Clusters.
2. Right-click the FlexPod-DC datacenter, click NetApp VSC > Update Host and Storage Data.



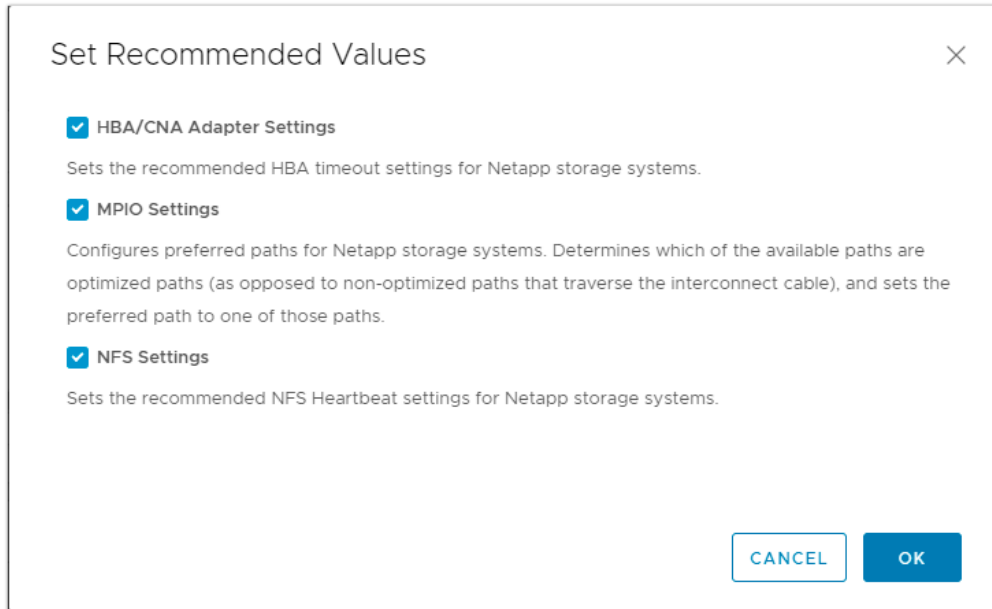
3. VSC displays a Confirm dialog box that informs you that this operation might take a few minutes.
4. Click OK.

## 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. Choose a host and then click Actions > NetApp VSC > Set Recommended Values.
3. In the NetApp Recommended Settings dialog box, choose 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). A vSphere host reboot may be required after applying the settings.

4. Click OK.



## Virtual Storage Console 9.7 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.

## Create the Storage Capability Profile

In order to leverage the automation features of VASA 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 Machine. 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 two pre-configured Storage Capability Profiles- Platinum and Bronze.

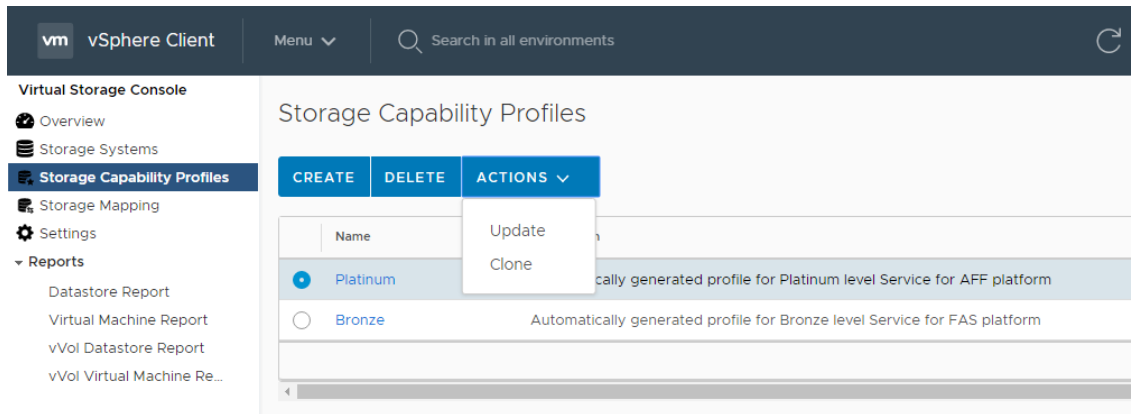
---

 **Adaptive QoS policies are not currently supported with VSC 9.7. Storage Capability Profiles (SCP) can still be created with Max IOPS and Min IOPS defined.**

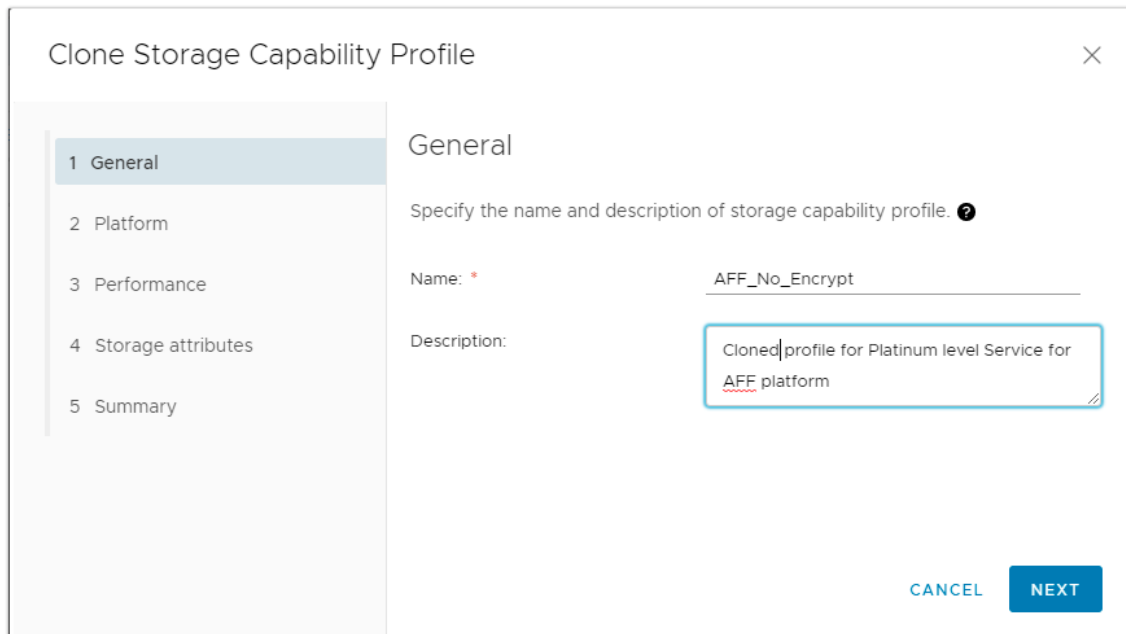
---

To review or edit one of the built-in profiles pre-configured with VSC 9.7 follow these steps:

1. In the NetApp Virtual Storage Console click Storage Capability Profiles.
2. Choose the Platinum Storage Capability Profile and choose Clone from the toolbar.



3. Enter a name for the cloned SCP and add a description if desired.



- Choose All Flash FAS(AFF) for the storage platform and click Next.

The screenshot shows the 'Clone Storage Capability Profile' dialog box with the 'Platform' tab selected. The left sidebar contains a list of tabs: 1 General, 2 Platform (selected), 3 Performance, 4 Storage attributes, and 5 Summary. The main area is titled 'Platform' and contains a 'Platform:' label followed by a dropdown menu showing 'All Flash FAS(AFF)'. At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

- Choose None to allow unlimited performance or set a the desired minimum and maximum IOPS for the QoS policy group.
- On the Storage attributes page, Change the Encryption and Tiering policy to the desired settings and click NEXT.

The screenshot shows the 'Clone Storage Capability Profile' dialog box with the 'Storage attributes' tab selected. The left sidebar contains a list of tabs: 1 General, 2 Platform, 3 Performance, 4 Storage attributes (selected), and 5 Summary. The main area is titled 'Storage attributes' and contains five dropdown menus: 'Deduplication:' (Yes), 'Compression:' (Yes), 'Space reserve:' (Thin), 'Encryption:' (No), and 'Tiering policy (FabricPool):' (Any). At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

- Review the summary page and choose FINISH to create the storage capability profile.

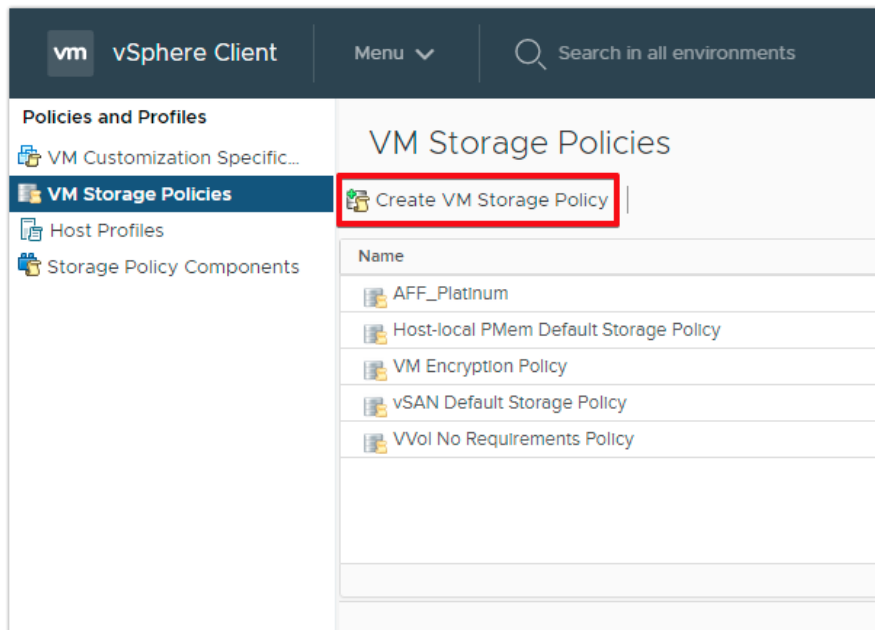


It is recommended to Clone the Storage Capability Profile if you wish to make any changes to the default profiles rather than editing the built-in profile.

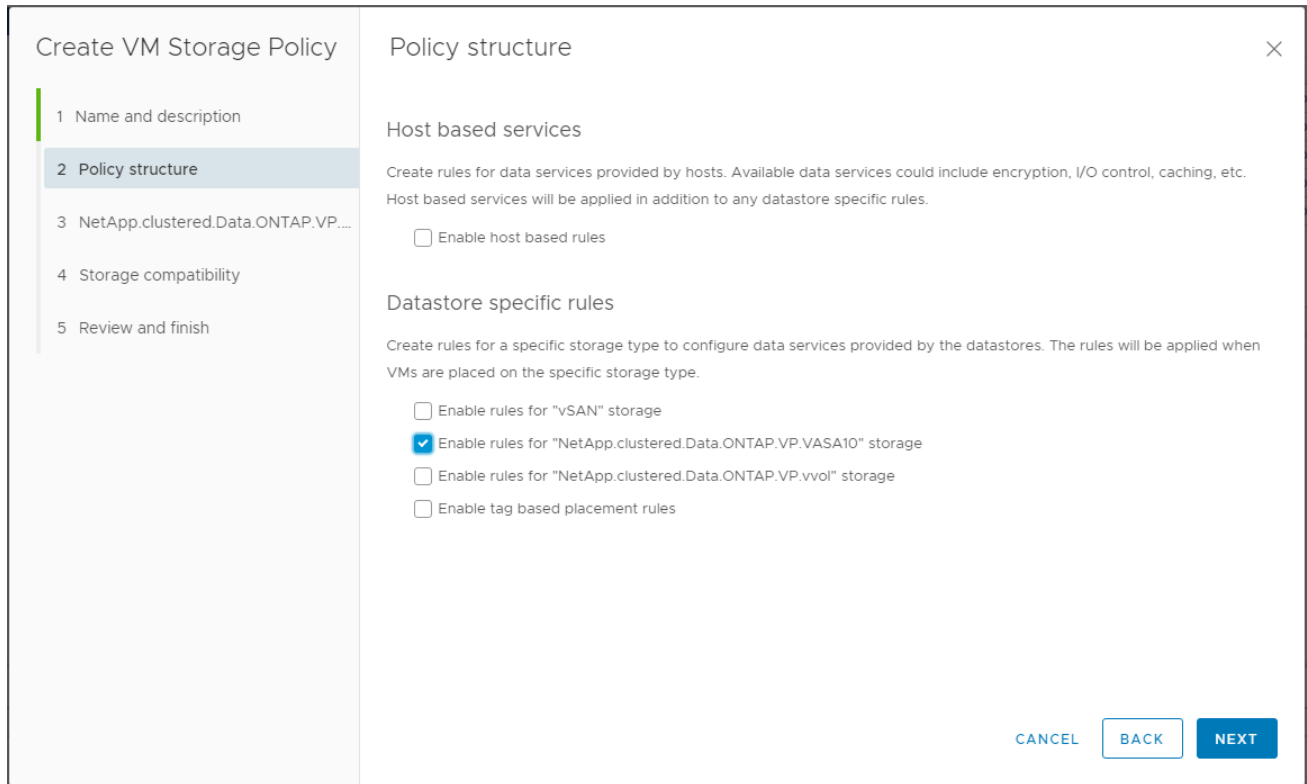
## Create a VM Storage Policy

Create a VM storage policy and associate a storage capability profile (SCP) to the datastore that meets the requirements defined in the SCP. To create a new VM Storage policy, follow these steps:

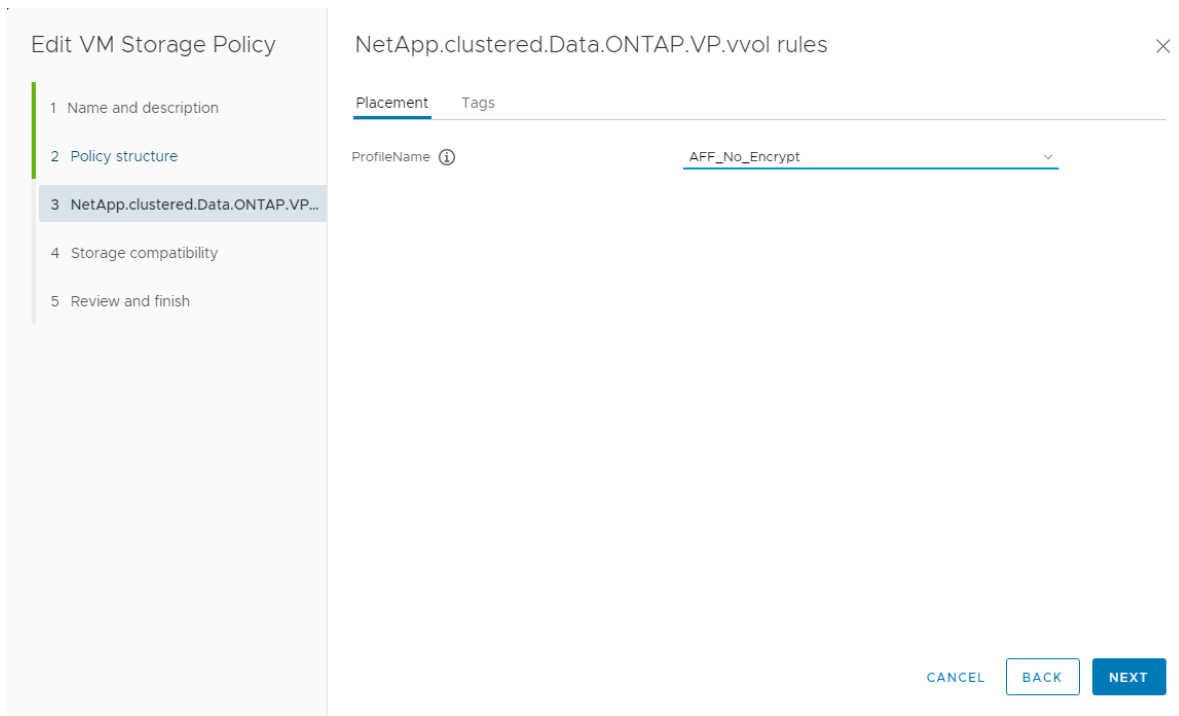
1. Navigate to Policies and Profiles from the vSphere Client menu.



2. Choose VM Storage Policies and click Create VM Storage Policy.
3. Create a name for the VM storage policy and enter a description and click NEXT.
4. Choose Enable rules for NetApp.clustered.Data.ONTAP.VP.VASA10 storage located under the Datastore specific rules section and click NEXT.



5. On the Placement tab select the SCP created in the previous step and click NEXT.



6. The datastores with matching capabilities are displayed, click NEXT.

7. Review the policy summary and click FINISH.

### 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.
3. Click Browse to choose the destination to provision the datastore as per the next step.
4. Choose 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.



7. Choose the desired 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: *	<u>AFF_No_Encrypt</u> ▾	
Platform: All Flash FAS(AFF)	Performance: None	
Compression: Yes	Deduplication: Yes	Tiering policy (FabricPool): Any
Space reserve: Thin	Encryption: No	
Storage system: *	<u>aa14-a800 (192.168.156.50)</u> ▾	
Storage VM: *	<u>MtWhitney-Infra</u> ▾	

[CANCEL](#) [BACK](#) [NEXT](#)

9. Choose 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:	<u>aggr1_node01 - (15520.38 GB Free)</u> ▾
Volumes:	Automatically creates a new volume.

Advanced options >

[CANCEL](#) [BACK](#) [NEXT](#)

10. Review the Summary and click FINISH.

### New Datastore

✕

- 1 General
- 2 Storage system
- 3 Storage attributes
- 4 Summary

#### Summary

**General**

vCenter server: nx-vc.flexpod.cisco.com

Provisioning destination: FlexPod-DC

Datastore name: NX\_NFS\_DS\_01

Datastore size: 500 GB

Datastore type: NFS

Protocol: NFS 3

Datastore cluster: None

**Storage system details**

Storage system: MtWhitney-Infra

CANCEL BACK 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:

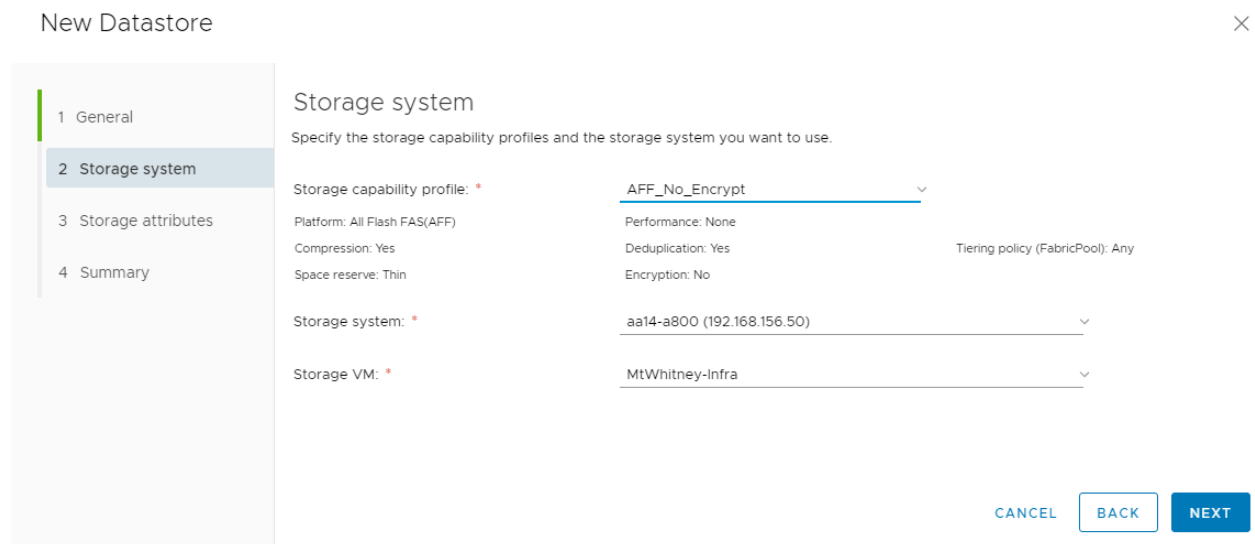
1. From the Virtual Storage Console Home page, click Overview
2. In the Getting Started Tab, Click Provision Button.
3. Click Browse to choose the destination to provision the datastore as per the next step.
4. Choose the type as VMFS and Enter the datastore name.
5. Provide the size of the datastore and the FC Protocol.
6. Check the storage capability profile and click NEXT.





7. Choose 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.



9. Choose 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: aggr1\_node01 - (15520.38 GB Free) ▾

Volumes: Automatically creates a new volume.

Advanced options >

CANCEL BACK NEXT

10. Review the Summary and click FINISH.

**New Datastore**

1 General  
2 Storage system  
3 Storage attributes  
4 **Summary**

**Summary**

**General**

vCenter server:	nx-vc.flexpod.cisco.com
Provisioning destination:	FlexPod-DC
Datastore name:	NX_FC_DS_01
Datastore size:	500 GB
Datastore type:	VMFS
Protocol:	FCP
File system:	VMFS6
Datastore cluster:	None

**Storage system details**

CANCEL BACK FINISH

11. Click OK.



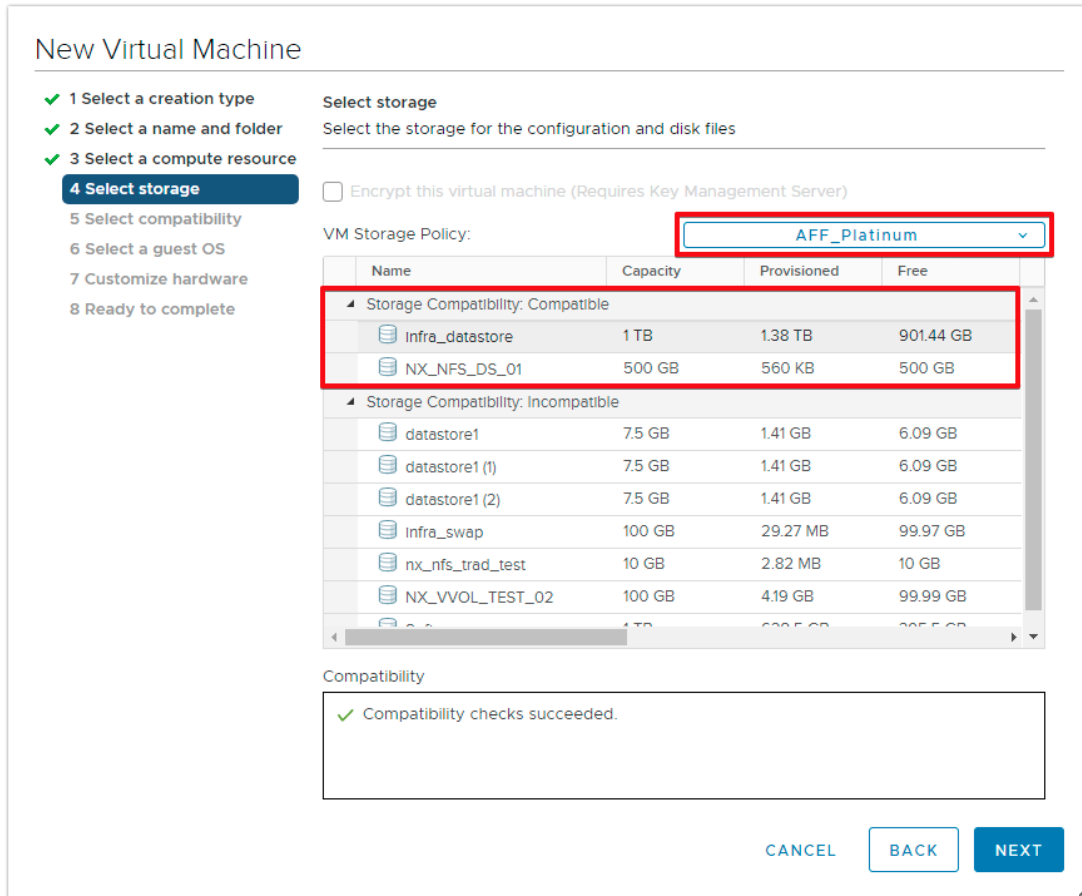
The datastore is created and mounted on all the hosts in the cluster. Click Refresh from the vSphere Web Client to see the newly created datastore or it is also listed in the VSC home page > Traditional Dashboard > Datastores view.

## Create Virtual Machine with Assigned VM Storage Policy

To create a virtual machine assigned to a VM storage policy, follow these steps:

1. Navigate to the VMs and Templates tab and click the FlexPod-DC datacenter.
2. Click Actions and click New Virtual Machine.

3. Choose Create a new virtual machine and choose NEXT.
4. Enter a name for the VM and click the FlexPod-DC datacenter.
5. Choose the FlexPod-Management Data compute Resource.
6. Choose the VM storage policy from the selections and choose a compatible datastore and click NEXT.



7. Choose Compatibility and click NEXT.
8. Choose the Guest OS and click NEXT.
9. Customize the hardware for the VM and click NEXT.
10. Review the details and click FINISH.

### NetApp SnapCenter 4.3

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.

## NetApp SnapCenter Architecture

The SnapCenter platform is based on a multitier architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter host agent. The host agent that performs virtual machine and datastore backups for VMware vSphere is the SnapCenter Plug-in for VMware vSphere. It is packaged as a Linux appliance (Debian-based Open Virtual Appliance format) and is no longer part of the SnapCenter Plug-ins Package for Windows. Additional information on deploying SnapCenter server for application backups can be found in the documentation listed below.

This guide focuses on deploying and configuring the SnapCenter plug-in for VMware vSphere to protect virtual machines and VM datastores.

You must install SnapCenter Server and the necessary plug-ins to support application-consistent backups for Microsoft SQL, Microsoft Exchange, Oracle databases and SAP HANA. Application level protection is beyond the scope of this deployment guide. Refer to the SnapCenter documentation for more information or the application specific CVD's and technical reports for detailed information on how to deploy SnapCenter for a specific application configuration.

- [SnapCenter 4.3 Documentation Center](#)
- [SAP HANA Backup and Recovery with SnapCenter](#)
- [FlexPod Datacenter with Microsoft SQL Server 2017 on Linux VM Running on VMware and Hyper-V](#)
- [SnapCenter Plug-in for VMware vSphere Documentation](#)

## Install SnapCenter Plug-In for VMware vSphere 4.3

NetApp SnapCenter Plug-in for VMware vSphere is a Linux-based virtual appliance which enables the SnapCenter Plug-in for VMware vSphere to protect virtual machines and VMware datastores.

### Host and Privilege Requirements for the SnapCenter Plug-in for VMware vSphere

Review the following requirements before you install the SnapCenter Plug-in for VMware vSphere virtual appliance:

- You must deploy the SnapCenter Plug-in for VMware vSphere 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 9 Port Requirements**

Port	Requirement
8080(HTTPS) bidirectional	This port is used to manage the virtual appliance
8144(HTTP) bidirectional	Communication between SnapCenter Plug-in for VMware vSphere and vCenter

Port	Requirement
443 (HTTPS)	Communication between SnapCenter Plug-in for VMware vSphere and vCenter

## License requirements for SnapCenter Plug-in for VMware vSphere

The following licenses are required to be installed on the ONTAP storage system to backup and restore VM's in the virtual infrastructure:

**Table 10 SnapCenter Plug-in for VMware vSphere License Requirements**

Product	License Requirements
ONTAP	SnapManager Suite: Used for backup operations  One of these: SnapMirror or SnapVault (for secondary data protection regardless of the type of relationship)
ONTAP Primary Destinations	To perform protection of VMware VMs and datastores the following licenses should be installed:  SnapRestore: used for restore operations  FlexClone: used for mount and attach operations
ONTAP Secondary Destinations	To perform protection of VMware VMs and datastores only:  FlexClone: used for mount and attach operations
VMware	vSphere Standard, Enterprise, or Enterprise Plus  A vSphere license is required to perform restore operations, which use Storage vMotion. vSphere Essentials or Essentials Plus licenses do not include Storage vMotion.



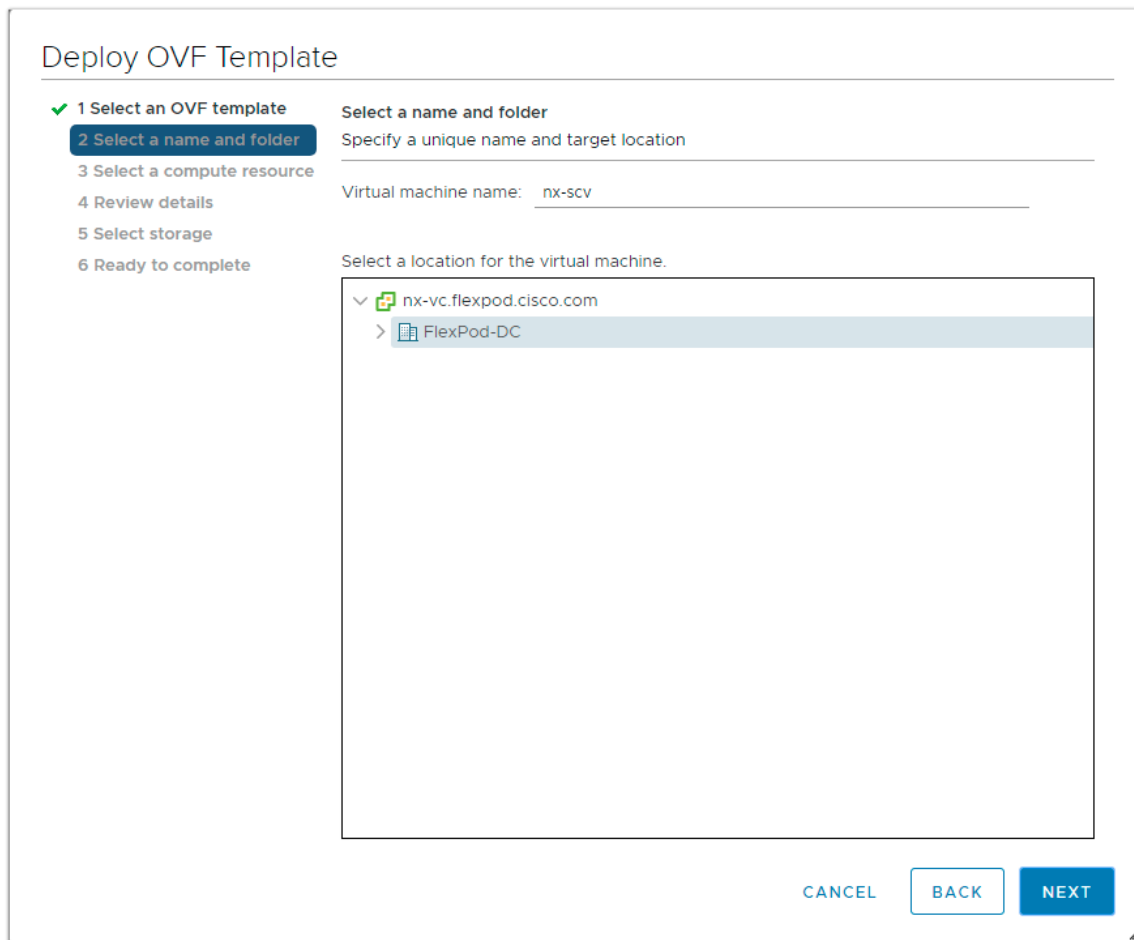
**It is recommended but not required, that you add SnapCenter Standard licenses to secondary destinations. If SnapCenter Standard licenses are not enabled on secondary systems, you cannot use SnapCenter after performing a failover operation. A FlexClone license on secondary storage is required to perform mount and attach operations. A SnapRestore license is required to perform restore operations.**

## Download and Deploy the SnapCenter Plug-in for VMware vSphere 4.3

To download and deploy the SnapCenter Plug-in for VMware vSphere appliance, follow these steps:

1. Download SnapCenter Plug-in for VMware vSphere OVA file from NetApp support site (<https://mysupport.netapp.com>).
2. From VMware vCenter, navigate to the VMs and Templates tab, right-click FlexPod-DC and choose 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 and location for the VM and click NEXT to continue.



5. On the Select a compute resource page, choose 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.
7. On the License agreements page, check the box I accept all license agreements.
8. On the Select storage page, change the datastore virtual disk format to Thin Provision 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

**Select storage**  
Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: **Thin Provision** ▾

VM Storage Policy: **Datastore Default** ▾

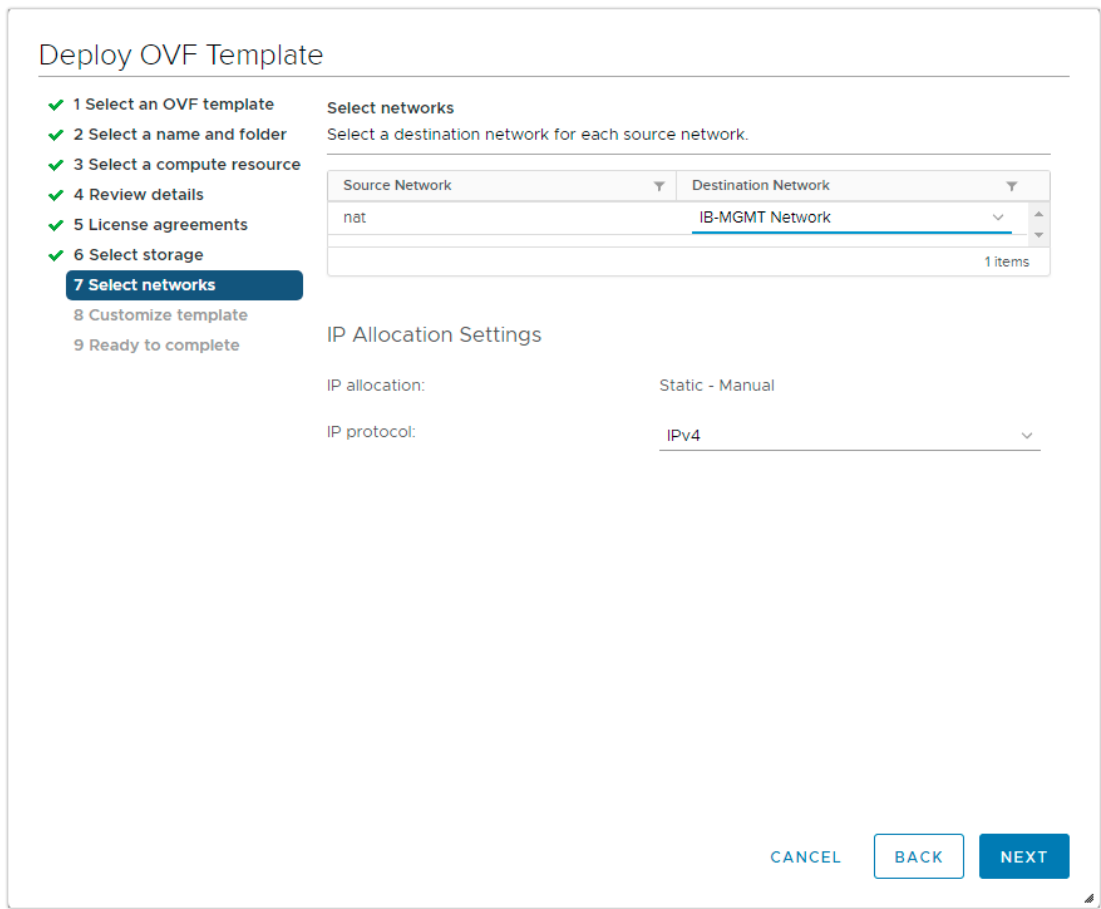
Name	Capacity	Provisioned	Free	Type
datastore1	7.5 GB	1.41 GB	6.09 GB	VM
datastore1 (1)	7.5 GB	1.41 GB	6.09 GB	VM
datastore1 (2)	7.5 GB	1.41 GB	6.09 GB	VM
infra_datastore	1 TB	331.21 GB	1,000.13 GB	NF
infra_swap	100 GB	9.16 MB	99.99 GB	NF

Compatibility

✓ Compatibility checks succeeded.

[CANCEL](#) [BACK](#) [NEXT](#)

9. On the Select networks page, choose a source network and map it to a destination network, and then click NEXT.



10. On the Customize template page, do the following:
  - a. In Register to existing vCenter, enter the vCenter credentials.
  - b. In Create SnapCenter Plug-in for VMware vSphere credentials, enter the SnapCenter Plug-in for VMware vSphere credentials.
  - c. In Create SCV credentials, create a username and password for the SCV maintenance user.
  - d. In Setup Network Properties, enter the network information.
  - e. In Setup Date and Time, choose the time zone where the vCenter is located.



### 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

2. Create SCV Credentials 2 settings	
2.1 Username	flexadmin
2.2 Password	
Password	.....
Confirm Password	.....
3. Setup Network Properties 1 settings	
3.1 Host Name	
Hostname for the appliance	nx-scv
3.2 Setup IPv4 Network Properties 6 settings	
3.2.1 IPv4 Address	
IP address for the appliance	10.1156.104
3.2.2 IPv4 Netmask	
Subnet to use on the deployed network	255.255.255.0
3.2.3 IPv4 Gateway	
Gateway on the deployed network	10.1156.254

CANCEL BACK NEXT

11. On the Ready to complete page, review the page and click FINISH.

### Deploy OVF Template

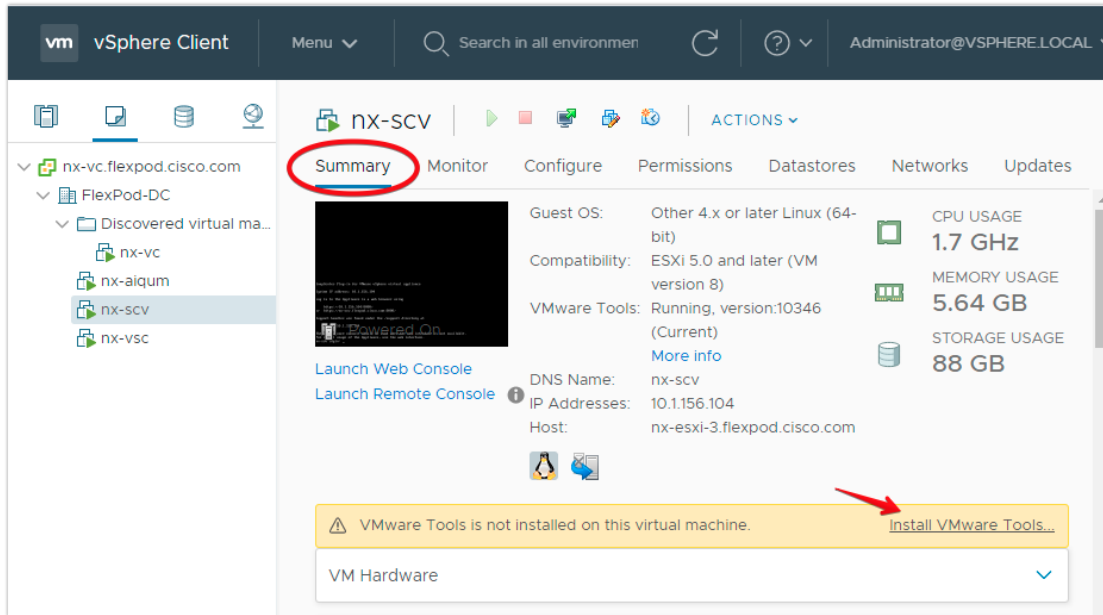
- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- ✓ 8 Customize template
- 9 Ready to complete**

**Ready to complete**  
Click Finish to start creation.

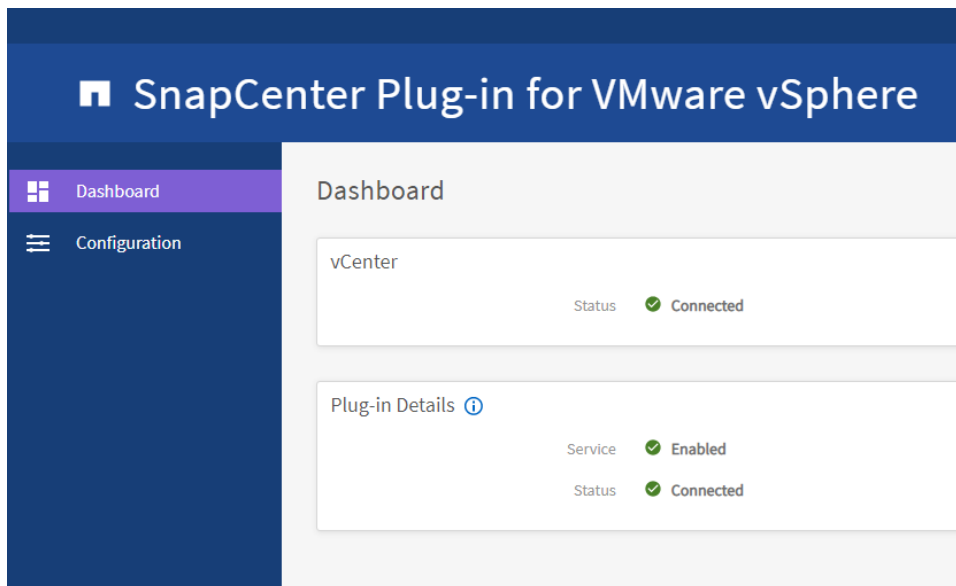
Provisioning type	Deploy from template
Name	nx-svc
Template name	scv-4.3.0.5595301-200319_1051
Download size	3.4 GB
Size on disk	5.3 GB
Folder	FlexPod-DC
Resource	FlexPod-Management
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thin provision
Network mapping	1
nat	IB-MGMT Network
IP 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, click Install VMware tools in the Yellow banner displayed in the summary tab of the appliance.



14. Log into SnapCenter Plug-in for VMware vSphere using the IP address displayed on the appliance console screen with the credentials that you provided in the deployment wizard. 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.



### SnapCenter Plug-in for VMware vSphere in vCenter Server

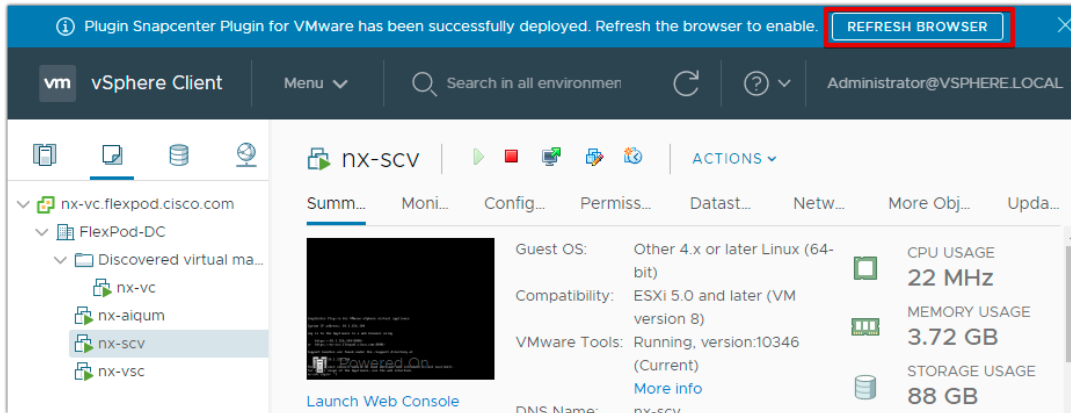
After you have successfully installed the Plug-in for VMware vSphere, to configure SnapCenter and make it ready to backup virtual machines, follow these steps:

1. In your browser, navigate to VMware vSphere Web Client URL <https://<vCenter Server>/ui>.

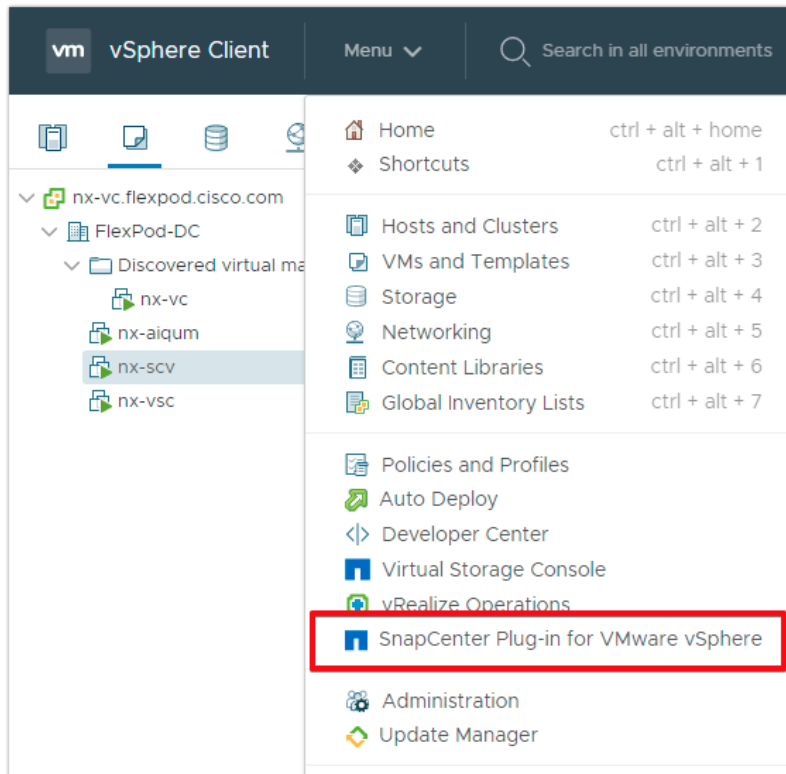


If currently logged into vCenter, logoff, close the open tab and sign-on again to access the SnapCenter Plug-In for VMware vSphere.

- After logging on to the vSphere Web Client you will see a blue banner indicating the SnapCenter plugin was successfully deployed. Click the refresh button to activate the plugin.



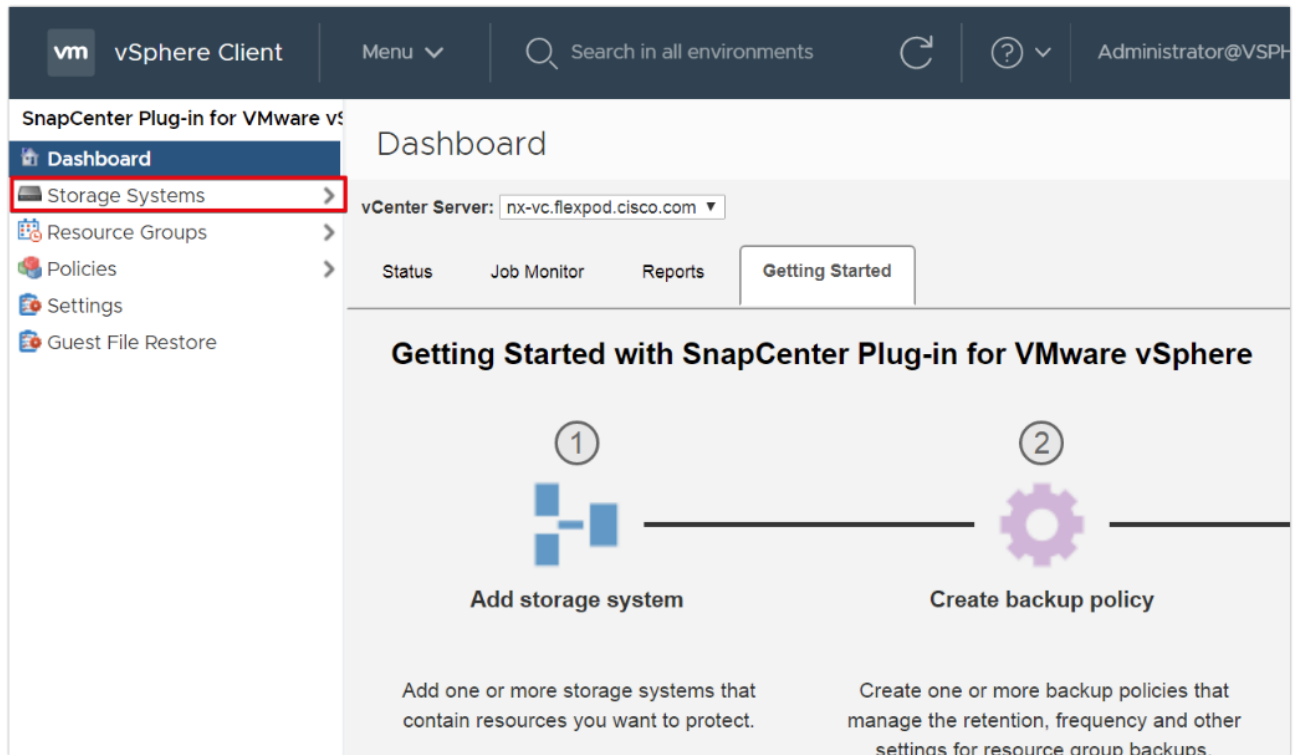
- On the VMware vSphere Web Client page, click the menu and click SnapCenter Plug-In for VMware vSphere to launch the SnapCenter Plug-In for VMware GUI.



## 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 a cluster or SVM.
3. Enter vCenter, Storage System, user credentials, and other required information in following dialog box.
4. Check the box for Log SnapCenter server events to syslog and Send AutoSupport Notification for failed operation to storage system.

**+ Add Storage System** ×

vCenter Server:

Storage System:

Platform:

Username:

Password:

Protocol:

Port:

Timeout:

Preferred IP:

**Event Management System(EMS) & AutoSupport Setting**

Log Snapcenter server events to syslog

Send AutoSupport Notification for failed operation to storage system

### 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, choose the backup frequency (hourly, daily, weekly, monthly, and on-demand only).
  - d. Expand the Advanced options and select VM Consistency and Include datastore with independent disks.
  - e. Click Add.

**+ New Backup Policy**
✕

**Name**

**Description**

**vCenter Server**

**Retention**   ⓘ

**Frequency**

**Replication**

Update SnapMirror after backup ⓘ

Update SnapVault after backup

Snapshot label

**Advanced** ▾

VM consistency

Include datastores with independent disks

**Scripts** ⓘ

4. Create multiple policies as required for different sets of VMs or datastores.

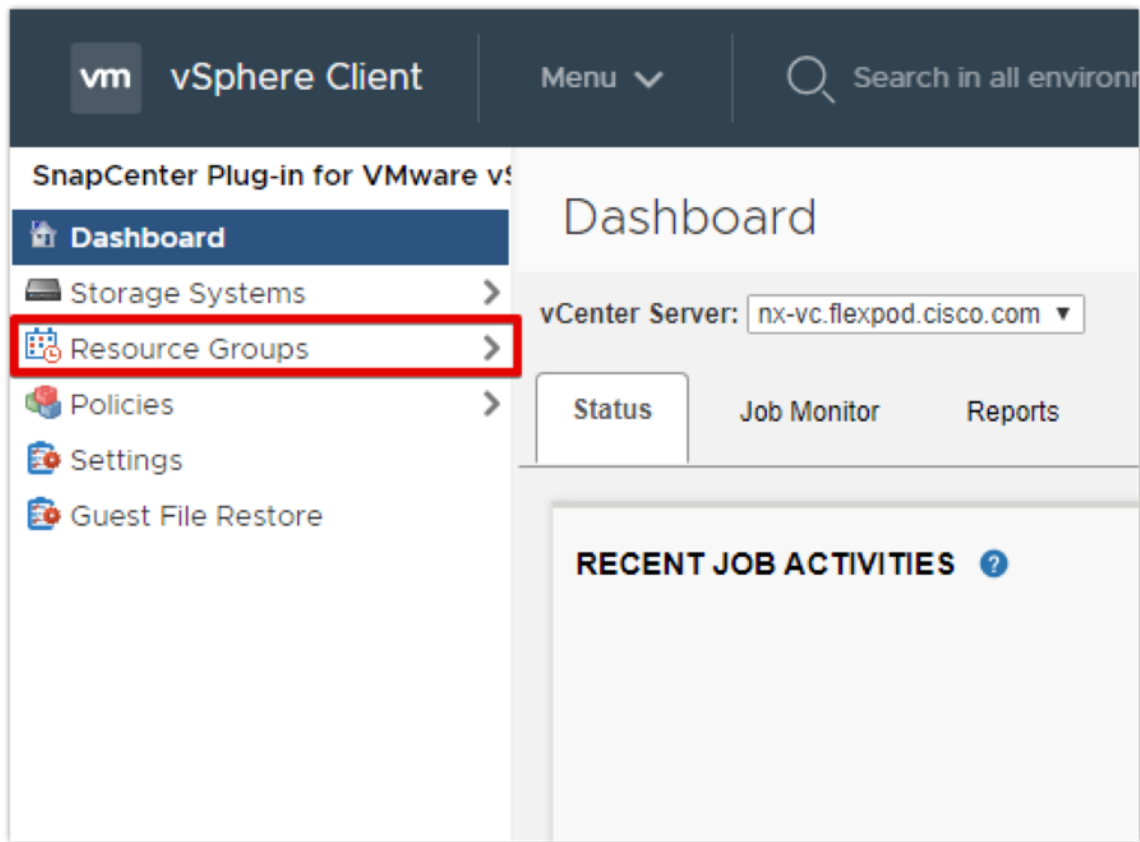
### Create Resource Groups

Resource groups are groups of virtual machines 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 SnapCenter Plugin for VMware vSphere, 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, choose NetApp SnapCenter from the drop-down list, and then choose Create Resource Group from the secondary drop-down list.



- b. To create a resource group for one datastore, click Storage, right-click a datastore, choose NetApp SnapCenter from the drop-down list, and then choose 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.



Simplify the task of locating virtual machine and datastore snapshots by selecting the **Custom snapshot format** option and choose the desired label such as \$ResourceGroup to have the resource group name appended to the snapshot name during snapshot operation.

- Choose a datastore as the parent entity to create a resource group of virtual machines, and then choose 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:

**Available entities**

**Selected entities**

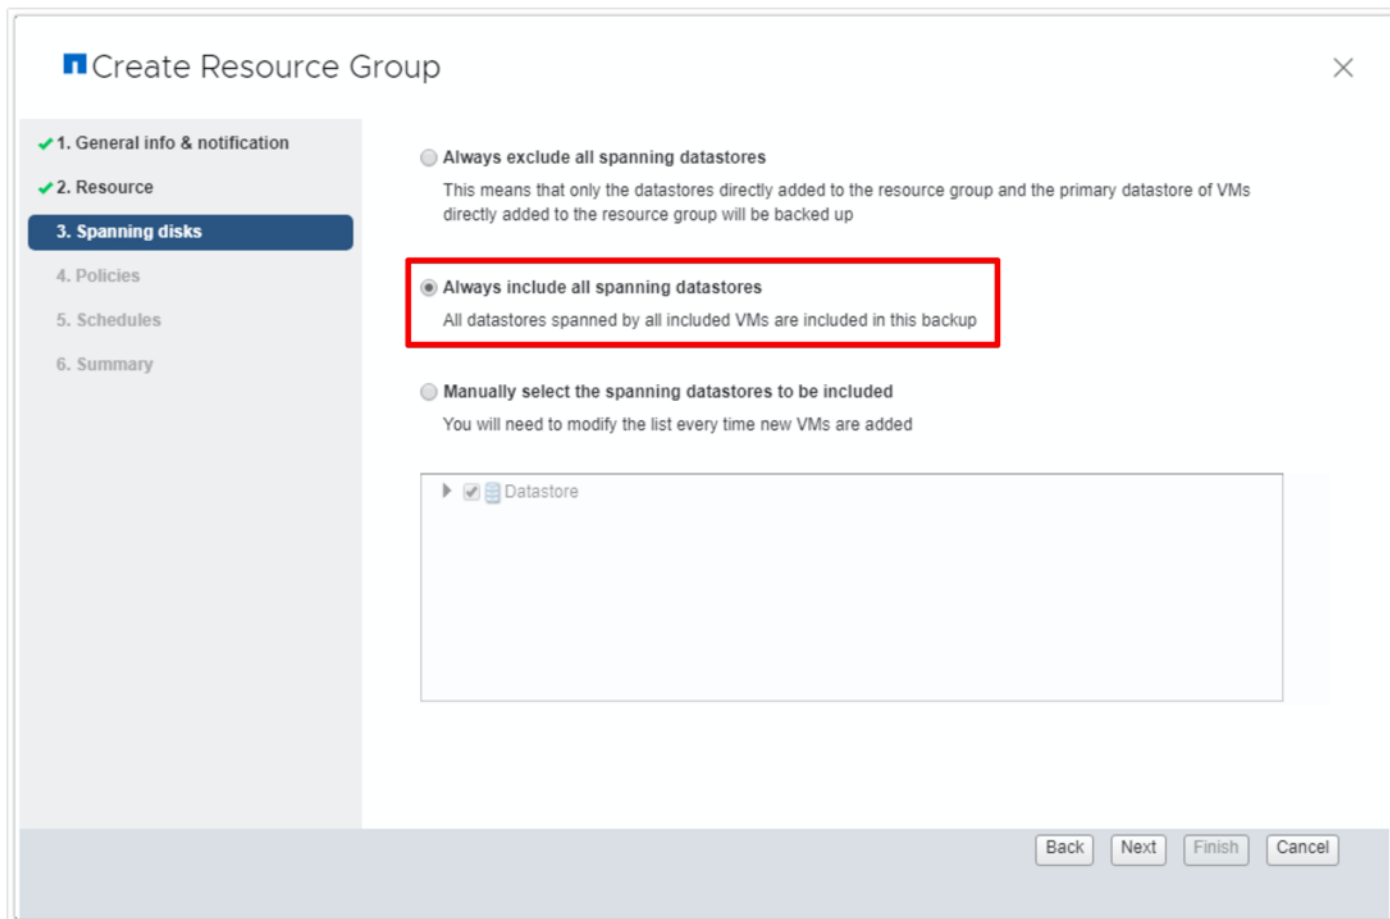
- nx-aiqum
- nx-dcnm
- nx-scv
- nx-vc
- nx-vsc

Back Next Finish Cancel



Entire datastores can be backed up by selecting FlexPod-DC in the parent entity list box and selecting the datastore.

- From the Spanning Disks options, choose the Always include all spanning datastores option.



5. From the Policies tab, choose one of the previously created policies that you want to associate with the resource group and click Next.

### Create Resource Group

- 1. General info & notification
- 2. Resource
- 3. Spanning disks
- 4. Policies**
- 5. Schedules
- 6. Summary

**+ Create Policy**

<input type="checkbox"/>	Name	VM Consistent	Include independent dis...	Schedule
<input checked="" type="checkbox"/>	Infra_vm_backups	Yes	Yes	Daily
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				

Back Next Finish Cancel

6. From the Schedules option, choose the schedule for each selected policy and 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

7. Review the summary and click Finish to complete the creation of the resource group.

**Create Resource Group**
✕

- ✓ 1. General info & notification
- ✓ 2. Resource
- ✓ 3. Spanning disks
- ✓ 4. Policies
- ✓ 5. Schedules
- ✓ 6. Summary

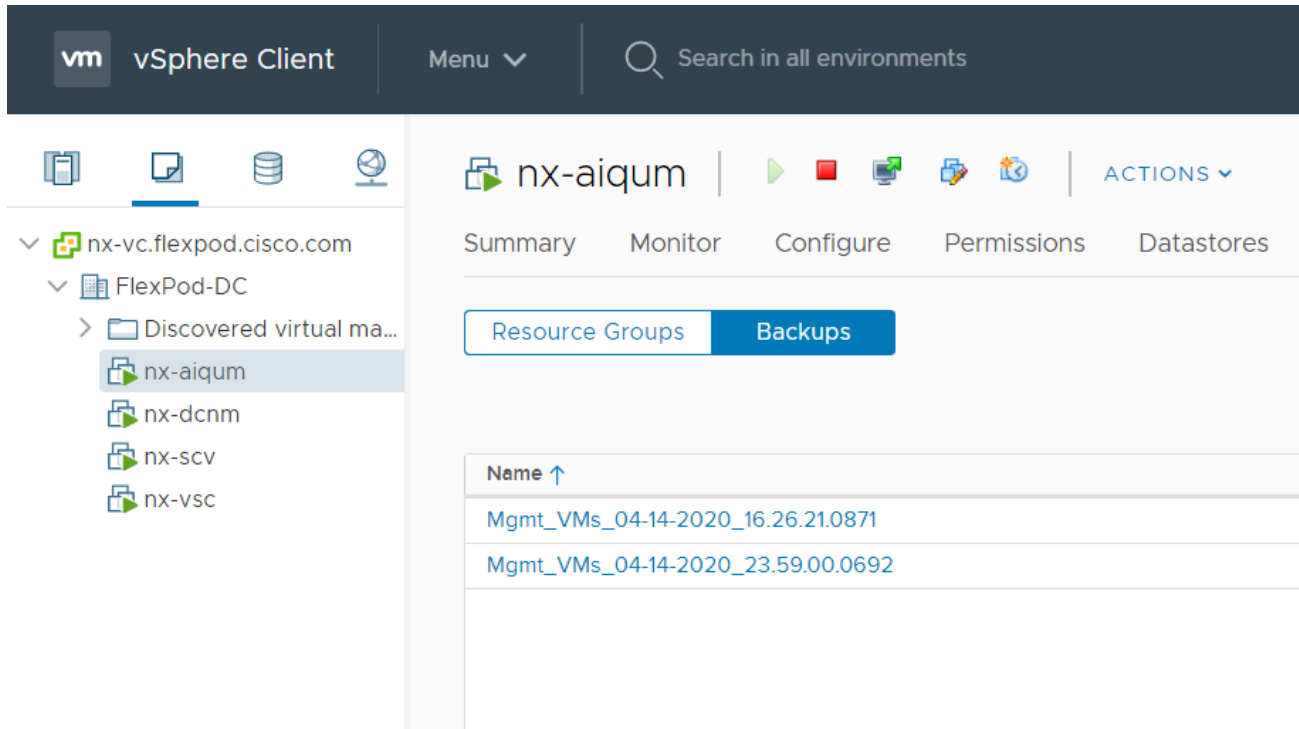
<b>Name</b>	Mgmt_VMs
<b>Description</b>	Infrastructure management
<b>Send email</b>	Errors
<b>Email send from</b>	flexadmin@flexpod.cisco.com
<b>Email send to</b>	flexadmin@flexpod.cisco.com
<b>Email subject</b>	Mgmt VM Backup Errors
<b>Custom snapshot format</b>	<ResourceGroup>_<TimeStamp>
<b>Entities</b>	nx-aicum, nx-dcnm, nx-scv, nx-vc, nx-vsc
<b>Spanning</b>	True
<b>Policies</b>	Infra_vm_back... : Daily

Back
Next
Finish
Cancel

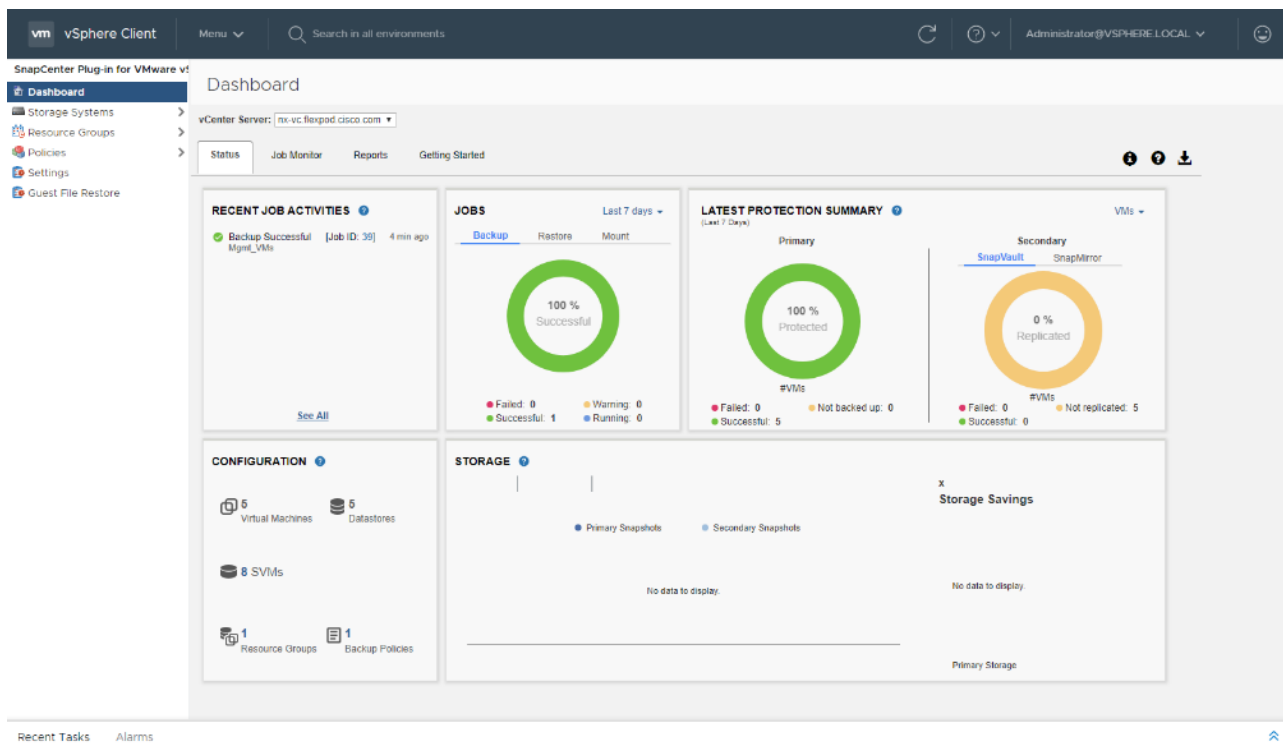
## View Virtual Machine Backups from vCenter by Using SnapCenter Plug-In

Backups of the virtual machines 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:

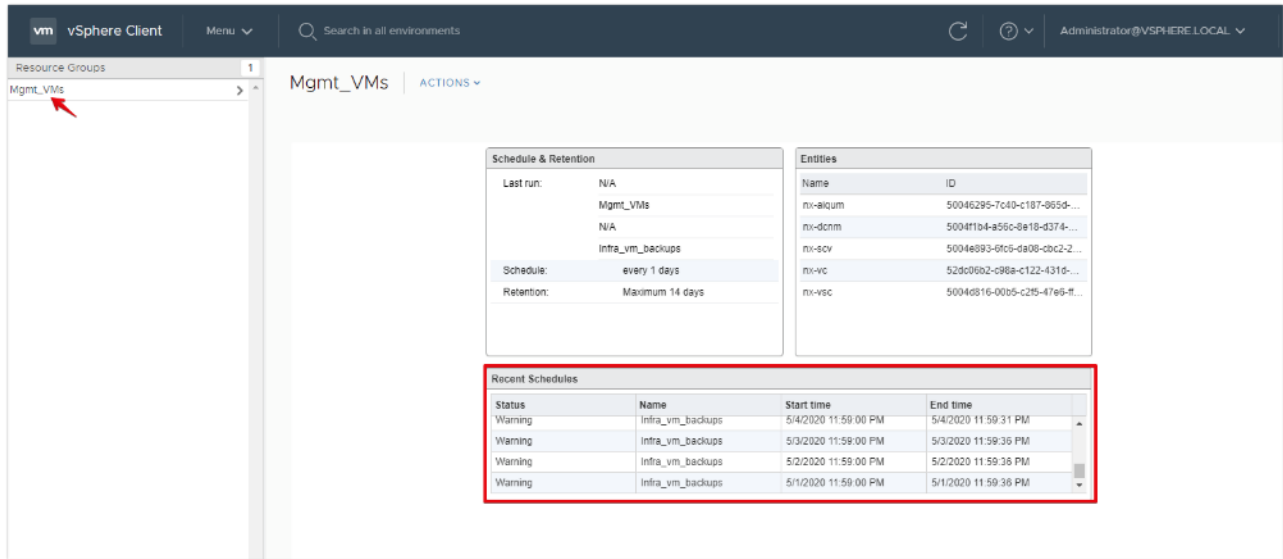
1. Navigate to the VMs and Templates tab.
2. Go to any virtual machine that is a member of a Resource Group and click the More Objects tab. Choose the Backups tab to view all the backups available for the virtual machine.



3. Navigate to the SnapCenter Plugin for VMware vSphere and choose the Dashboard tab to view recent job activity, backup jobs and configuration details.



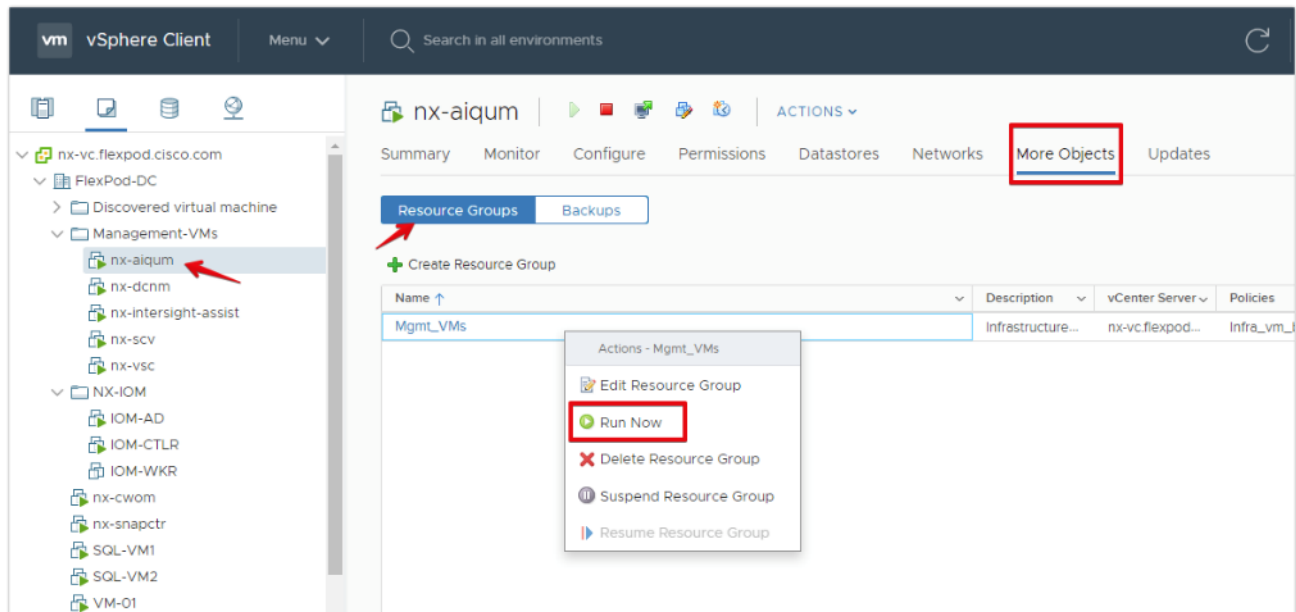
4. In the SnapCenter Plug-In for VMware vSphere, click Resource Groups and choose any resource group. In the right pane, the completed backups are displayed.



## Create On-Demand Backup

To create an on-demand backup for any resource group, follow these steps:

1. From the VMs and Templates tab, choose a virtual machine contained in the resource group where you want to create an on-demand backup.
2. Click the More Objects tab and choose the Resource Groups tab from the toolbar to display the list of resource groups.
3. Right-click the resource group and click Run Now to run the backup immediately.






## Restore from vCenter by Using SnapCenter Plug-In

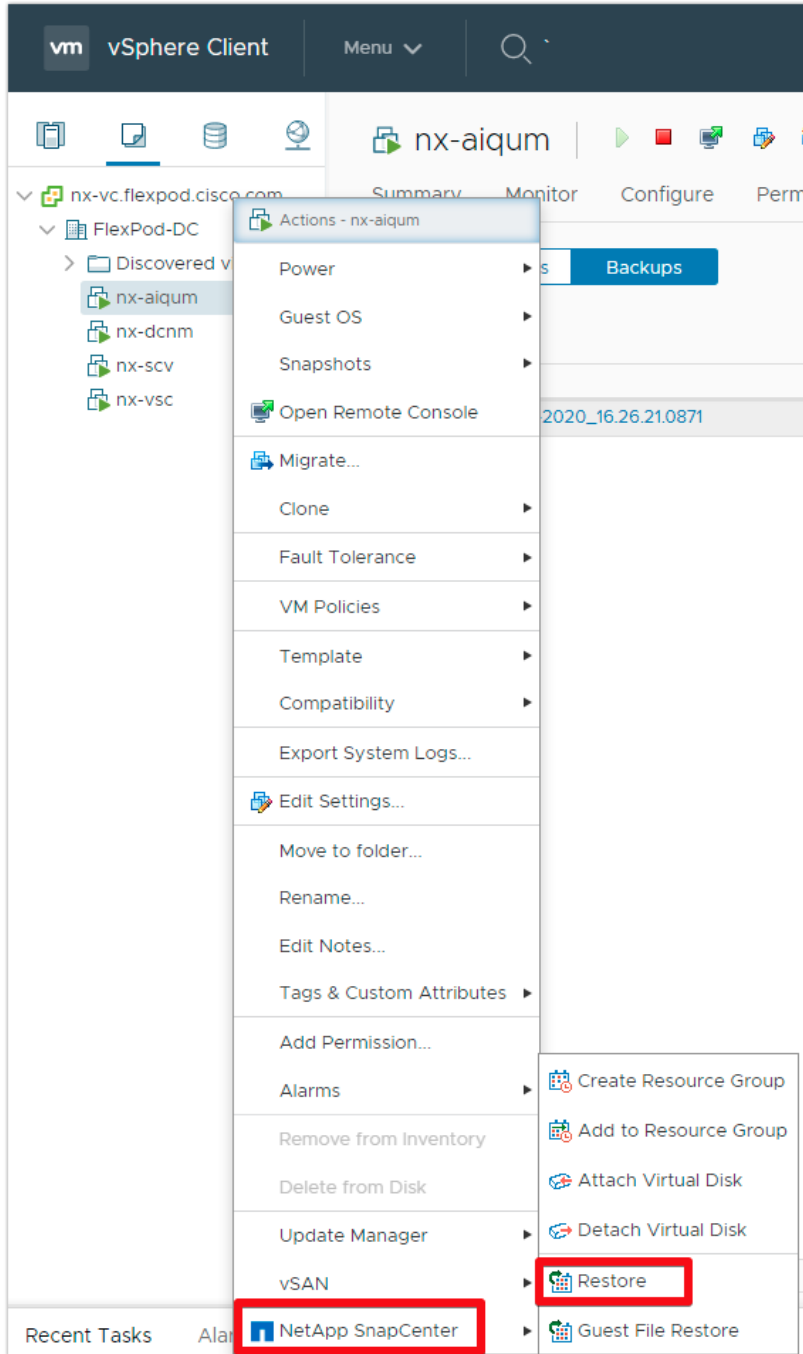
To restore from vCenter by using SnapCenter Plug-In, follow these steps:

---

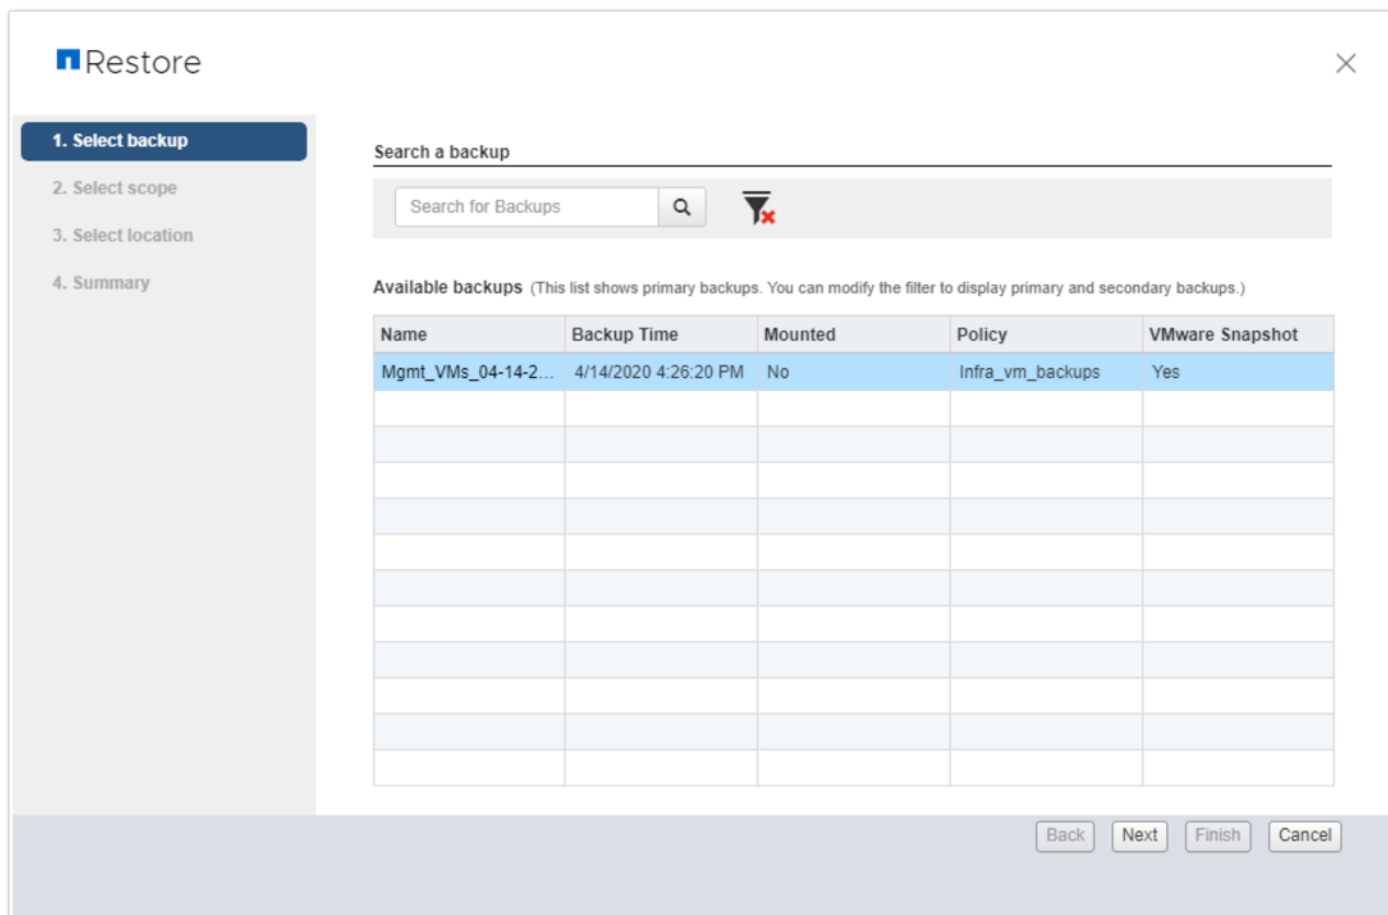
 **The Plug-In for VMware vSphere provides native backup, recovery, and cloning of virtualized applications.**

---

1. Navigate to VMs and Templates, choose a VM and right-click to access the context menu. Choose NetApp SnapCenter > Restore.



2. Choose a backup from which to restore. Click Next.



3. From the Restore Scope drop-down list:
- Choose either " Entire virtual machine" to restore the virtual machine with all VMDKs, or choose " Particular Virtual Disk" to restore the VMDK without affecting the virtual machine configuration and other VMDKs.
  - Choose the ESXi host that the VM should be restored to and check the box if you wish to restart the VM upon being restored. Click Next.

**Restore**

- 1. Select backup
- 2. Select scope**
- 3. Select location
- 4. Summary

Restore scope: Entire virtual machine

Restored VM name: Entire virtual machine

ESXi host name: nx-esxi-2.flexpod.cisco.com

Restart VM:

Back Next Finish Cancel

4. Choose the destination datastore and click Next.

**Restore** ✕

✓ 1. Select backup

✓ 2. Select scope

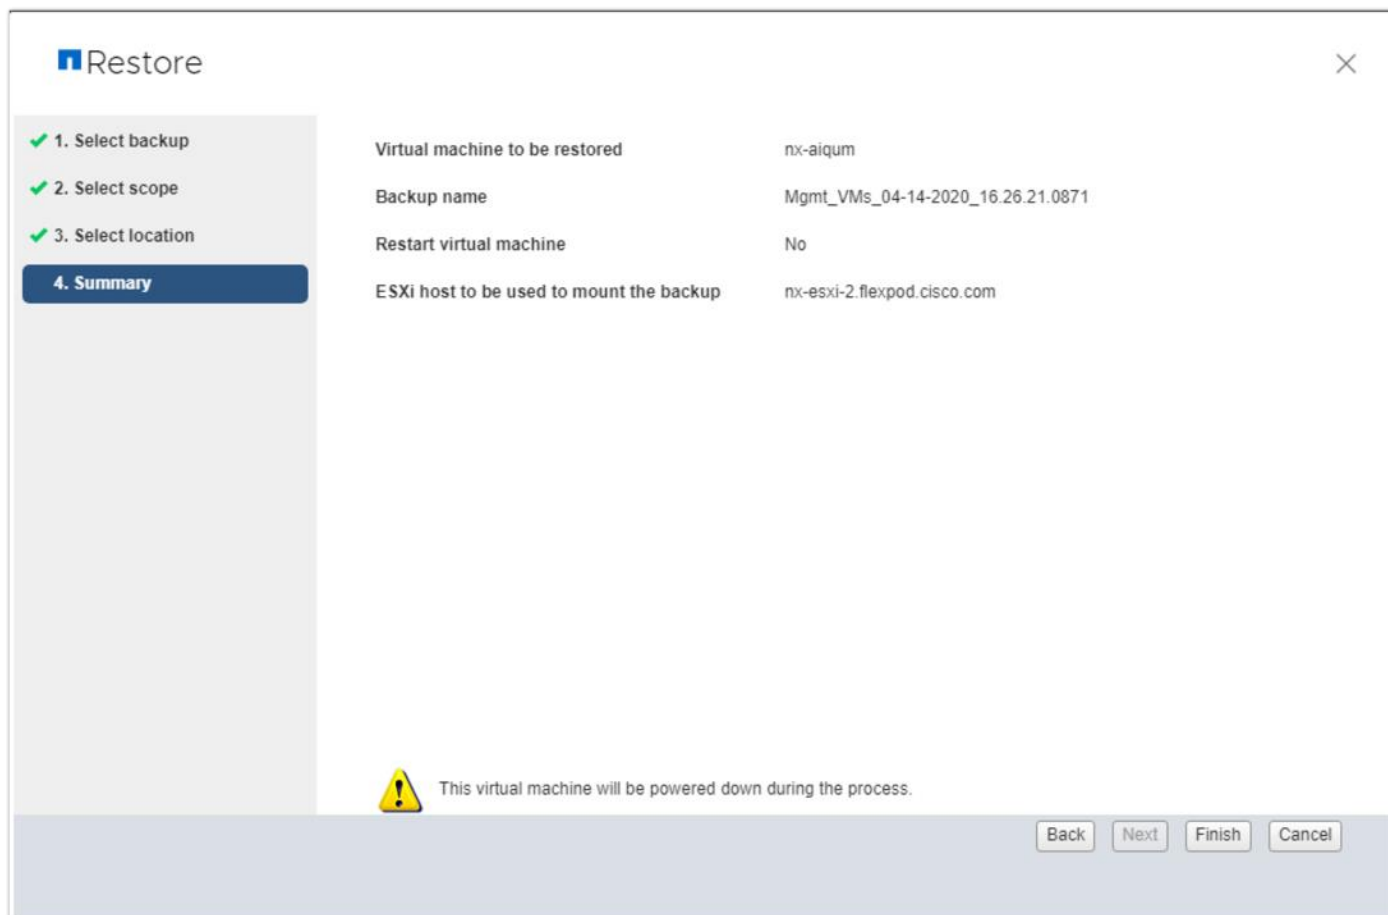
**3. Select location**

4. Summary

Destination datastore	Locations
infra_datastore	(Primary) 10.1.156.10:MTW_Infra_DS ▾

Back Next Finish Cancel

5. Review the Summary and click Finish to complete the restore process.



## Active IQ Unified Manager 9.7

Active IQ Unified Manager (AIQ UM) enables you to monitor and manage the health and performance of your ONTAP storage systems and virtual infrastructure 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 9.7 as a virtual appliance. 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
Free Disk Space/virtual disk size	<ul style="list-style-type: none"> <li>• 5 GB - Thin provisioned</li> <li>• 152 GB - Thick provisioned</li> </ul>



There is a limit to the number of nodes that a single instance of Active IQ Unified Manager can monitor before you need to install a second instance of AIQ UM. See the [Unified Manager Best Practices Guide \(TR-4621\)](#) for more details.

To install Active IQ Unified Manager 9.7, follow these steps:

1. Download NetApp Active IQ Unified Manager for VMware vSphere OVA file from [NetApp support site](#).
2. From the VMware vCenter, click the VMs and Templates tab, then click 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, and choose 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:

Select a location for the virtual machine.

- nx-vc.flexpod.cisco.com
  - FlexPod-DC

CANCEL BACK NEXT

5. On the Select a compute resource page, choose 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 <a href="http://www.netapp.com">http://www.netapp.com</a>
Download size	2.2 GB
Size on disk	3.9 GB (thin provisioned)
	152.0 GB (thick provisioned)
Extra configuration	keyboard.typematicMinDelay = 2000000

CANCEL
BACK
NEXT

7. On the License agreements page, check the box 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. Choose the disk format for the VMDKs.
  - b. Choose a VM Storage Policy.
  - c. Choose 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: Thin Provision ▼

VM Storage Policy: Datastore Default ▼

Name	Capacity	Provisioned	Free	Type
datastore1	7.5 GB	1.41 GB	6.09 GB	VM <span style="font-size: 0.8em;">▲</span>
datastore1 (1)	7.5 GB	1.41 GB	6.09 GB	VM
datastore1 (2)	7.5 GB	1.41 GB	6.09 GB	VM
<b>infra_datastore</b>	<b>1 TB</b>	<b>93774 GB</b>	<b>899.64 GB</b>	<b>NF</b>
infra_swap	100 GB	15.12 MB	99.99 GB	NF
nx_nfs_trad_test	10 GB	620 KB	10 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.
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.

Host FQDN

Specifies the hostname for the appliance. Leave blank if DHCP is desired.

nx-aiqum.flexpod.cisco.co

IP Address

Specifies the IP address for the appliance. Leave blank if DHCP is desired.

10.1156.106

Network Mask (or) Prefix Length

CANCEL BACK NEXT



Scroll through the customization template to ensure all required values are entered.

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	nx-aiqum
Template name	ActiveIQUnifiedManager-9.7
Download size	2.2 GB
Size on disk	3.9 GB
Folder	FlexPod-DC
Resource	FlexPod-Management
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thin provision
Network mapping	1
nat	IB-MGMT Network
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual

CANCEL BACK FINISH

12. Choose the newly created Active IQ VM, right-click it and choose Power > Power On to start the virtual machine.

13. While the virtual Machine is powering on, click the prompt in the yellow banner to Install VMware tools.

nx-aiqum | ACTIONS

Summary | Monitor | Configure | Permissions | Datastores | Networks | Updates

Powered On

Guest OS: Other 4.x or later Linux (64-bit)  
 Compatibility: ESXi 5.5 and later (VM version 10)  
 VMware Tools: Running, version:10346 (Current)  
[More info](#)

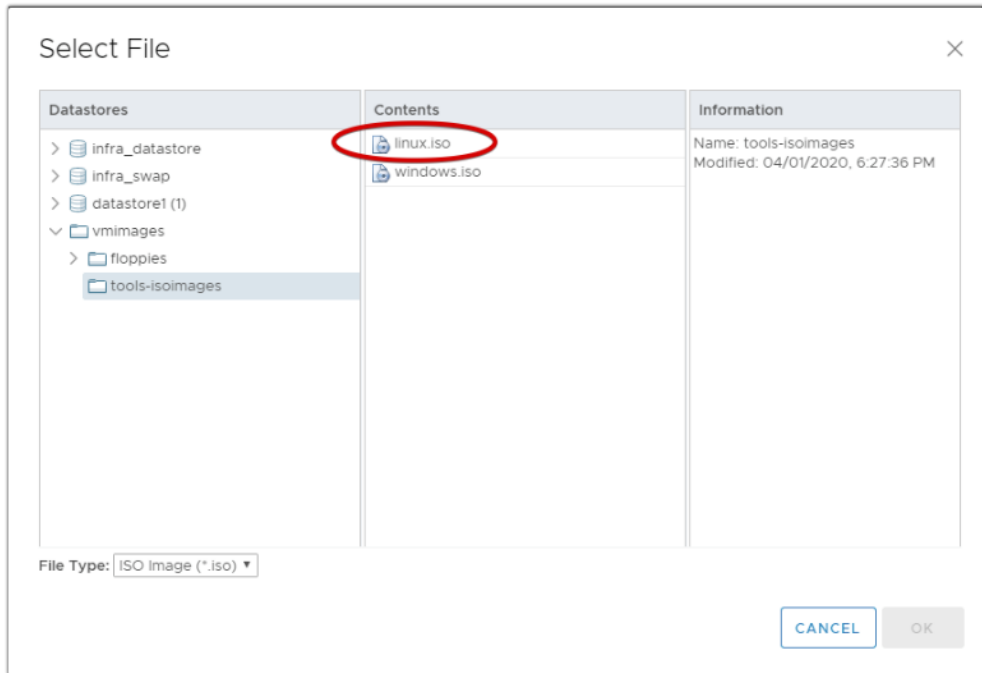
DNS Name: UnifiedManager  
 IP Addresses:  
 Host: nx-esxi-3.flexpod.cisco.com

CPU USAGE: 275 MHz  
 MEMORY USAGE: 368 MB  
 STORAGE USAGE: 152 GB

VMware Tools is not installed on this virtual machine. [Install VMware Tools...](#)

VM Hardware | Notes

- Click Mount in the Install VMware Tools dialog box and browse to the vmimages > tools-isoimages folder and choose linux.iso and click OK to proceed with installing VMware tools.



- Open a console session to the Active IQ Unified Manager appliance and configure the time zone information when displayed.

```

Configuring timezone...

Configuring tzdata
-----
Please select the geographic area in which you live. Subsequent configuration questions will narrow
this down by presenting a list of cities, representing the time zones in which they are located.

  1. Africa          5. Arctic Ocean    9. Indian Ocean    13. None of the above
  2. America         6. Asia            10. Pacific Ocean
  3. Antarctica     7. Atlantic Ocean  11. System V timezones
  4. Australia      8. Europe          12. US

Geographic area: 2_
    
```

- Create the maintenance user account when prompted by specifying a user account name and password.



Store the maintenance user account and password in a secure location. It is required for the initial GUI login and to make any configuration changes to the appliance settings that may be needed in the future.

```

Create the maintenance user.

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

For example, the maintenance user can do the following:

- Change network settings
- Upgrade to a newer version of Active IQ Unified Manager or apply patches
- Create and manage other users and their permissions using the web interface

At the prompt, specify the username and password for the new maintenance user.

The maintenance user name should start with any letter between a-z,
followed by any combination of -, a-z or 0-9.

Username: flexadmin
Enter new UNIX password:
Retype new UNIX password: _

```

17. Log into NetApp Active IQ UM using the IP address or URL displayed on the deployment screen and the maintenance user credentials you created in the previous step.

```

Active IQ Unified Manager

Log in to Active IQ Unified Manager in a web browser by using

  https://10.1.156.106/

or

  https://nx-aiqum.flexpod.cisco.com/

The maintenance console should be used when the web interface is not available.
For normal usage of Active IQ Unified Manager, use the web interface.

Hint: Num Lock on

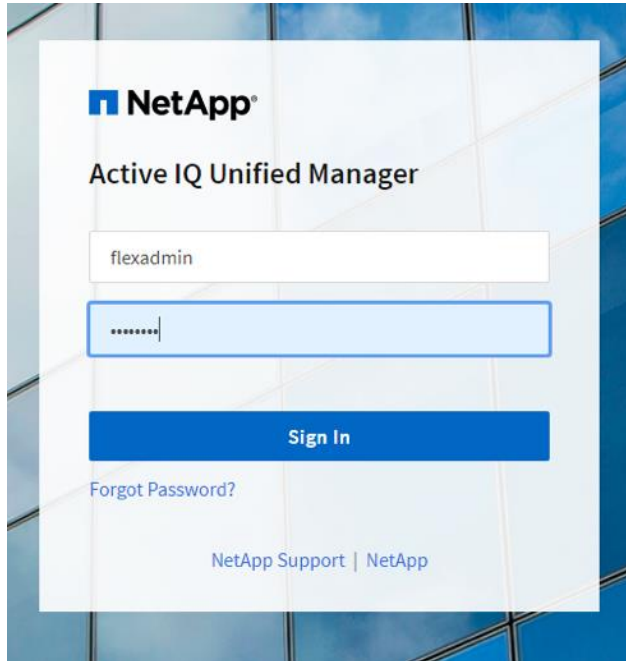
nx-aiqum login: _

```

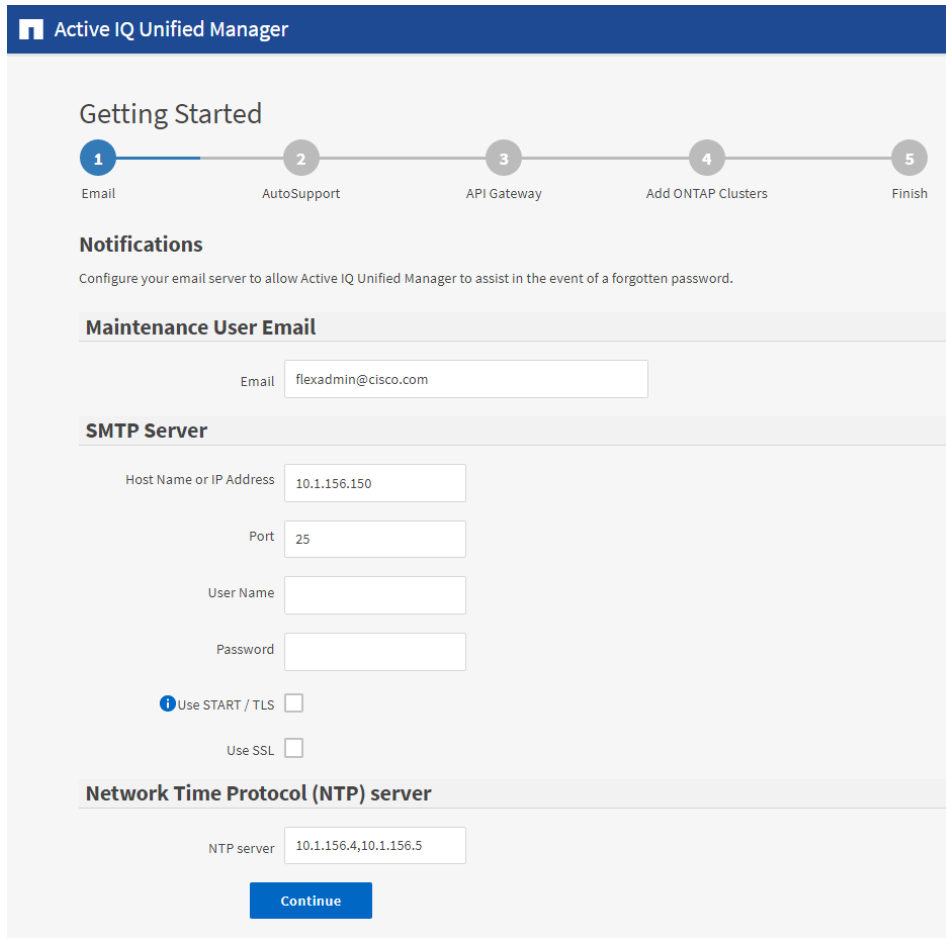
## Configure Active IQ Unified Manager

To configure Active IQ Unified Manager (AIQ UM) and add a storage system for monitoring, follow these steps:

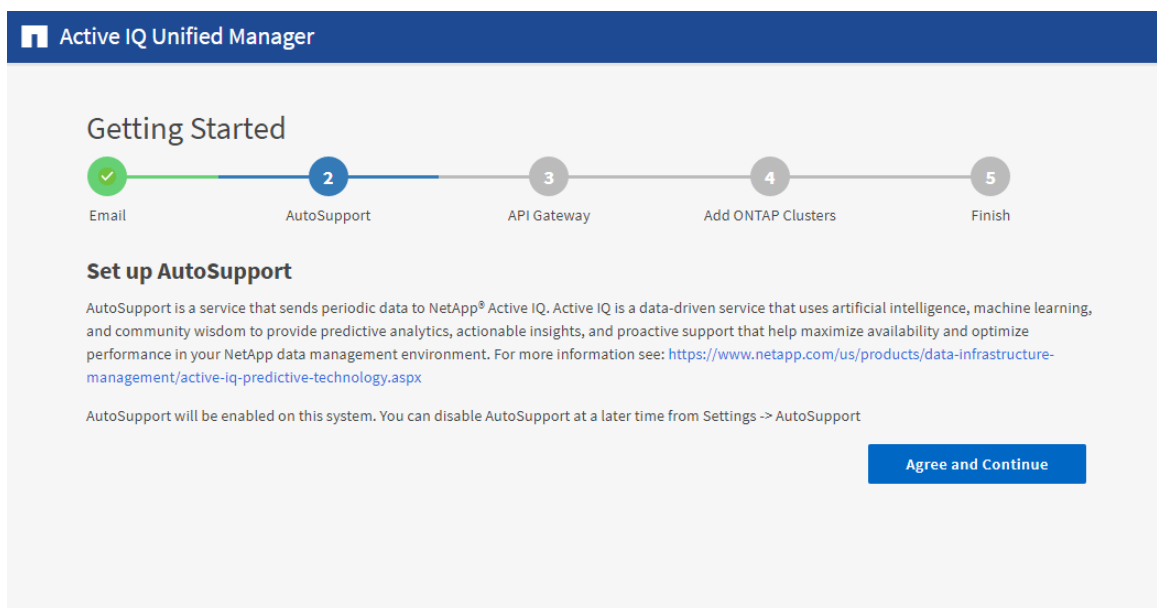
1. Launch a web browser and log into Active IQ Unified Manger.



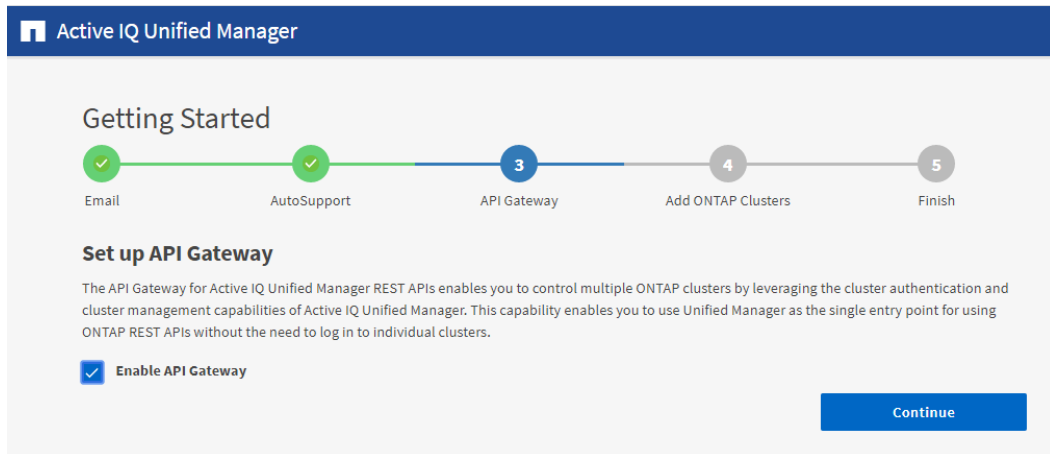
2. Enter the email address that Unified Manager will use to send alerts, enter the mail server configuration, and the IP address or hostname of the NTP server. Choose Continue and complete the AutoSupport configuration.



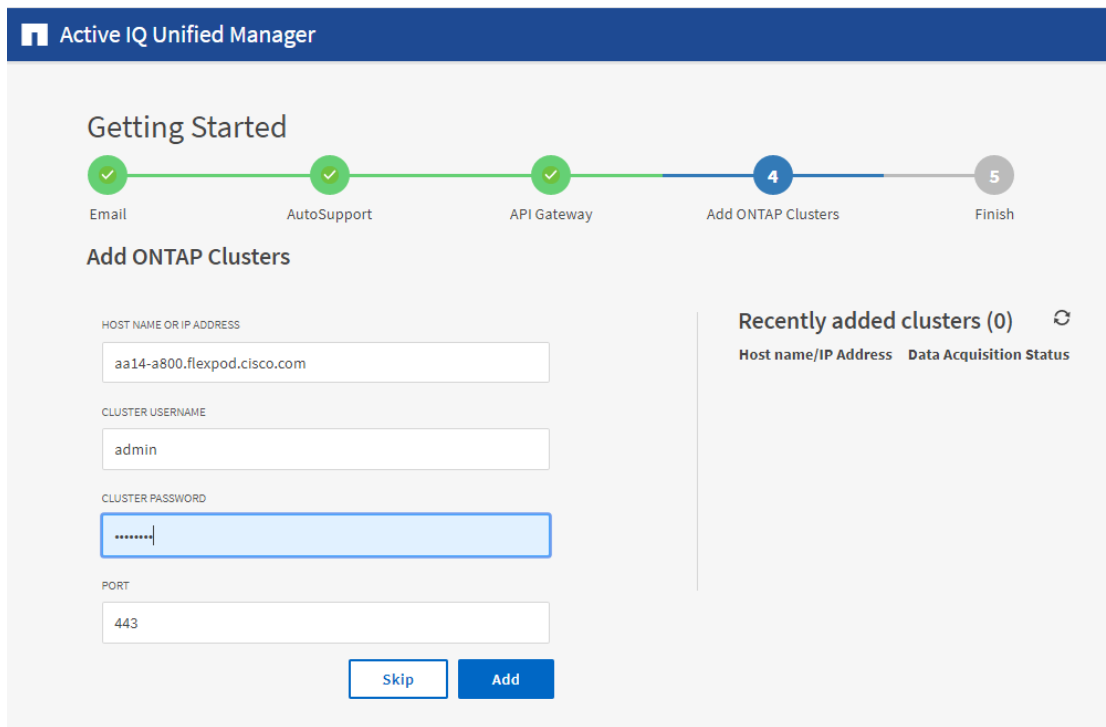
3. Configure AutoSupport for Unified Manager by clicking Agree and Continue.



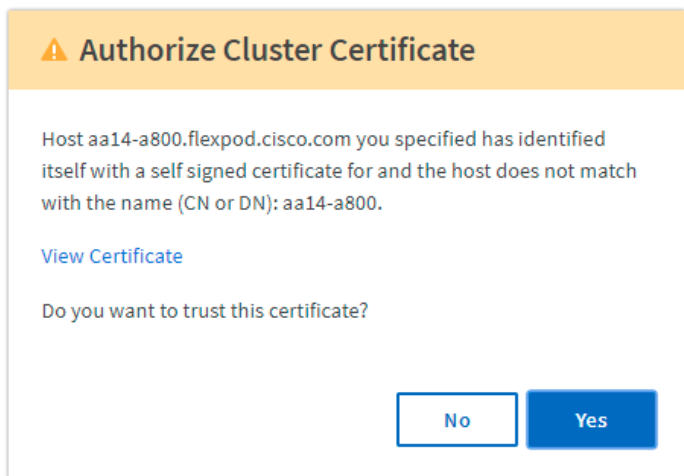
4. Choose the Enable API Gateway checkbox and click Continue to setup the API gateway for Active IQ Unified Manager.



5. Enter the ONTAP cluster hostname or IP address and the admin login credentials then click Add.



6. A security prompt will be displayed to authorize the cluster certificate. Choose Yes to trust the certificate.



7. When prompted to trust the self-signed certificate from Active IQ Unified Manager, click Yes to finish and add the storage system.

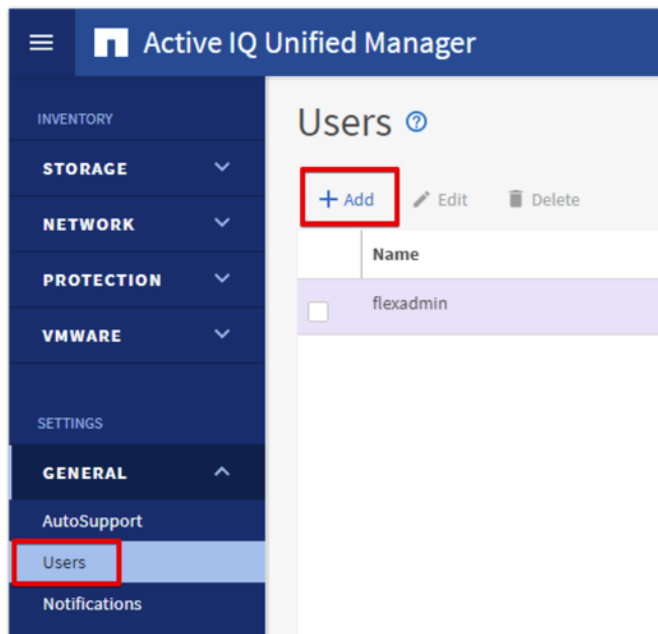


The initial discovery process can take up to 15 minutes to complete.

## Add Local Users to Active IQ Unified Manager

To add a local user to Active IQ Unified Manager, follow these steps:

1. Navigate to the General section and click Users.



2. Click Add and complete the requested information:
3. Choose Local User for the Type.
4. Enter a username and password.



5. Add the user's email address.
6. Choose the appropriate role for the new user.
7. Click Save to add the new user to AIQ UM.

The screenshot shows a web form titled "Users: Add" with a help icon. The form contains the following fields:

- TYPE:** A dropdown menu with "Local User" selected.
- NAME:** A text input field containing "scott.kovacs".
- PASSWORD:** A text input field with masked characters "\*\*\*\*\*".
- CONFIRM PASSWORD:** A text input field with masked characters "\*\*\*\*\*".
- EMAIL:** A text input field containing "@netapp.com".
- ROLE:** A dropdown menu with "Application Administrator" selected.

At the bottom of the form, there are two buttons: "Save" (in blue) and "Cancel".

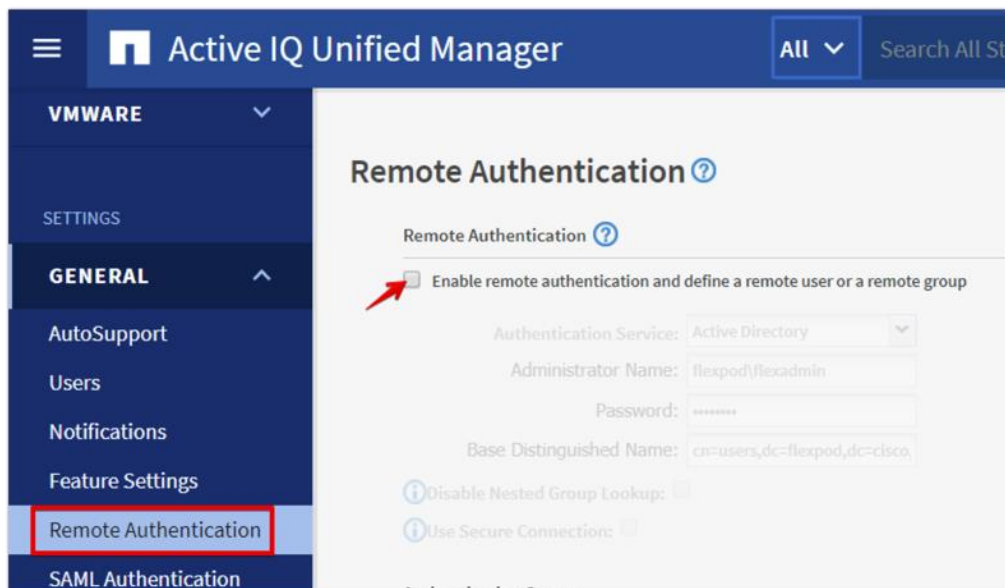
## Configure Remote Authentication

Simplify user management and authentication for Active IQ Unified Manager by integrating it with Microsoft Active Directory. To connect Active IQ Unified Manager to Active Directory and perform user authentication with the AD domain, follow these steps:



**You must be logged on as the maintenance user created during the installation or another user with Application Administrator privileges to configure remote authentication.**

1. Navigate to the General section and choose Remote Authentication.
2. Choose the option to Enable remote authentication and define a remote user or remote group.



3. Choose Active Directory from the authentication service list.
4. Enter the Active Directory service account name and password. The account name can be in the format of domain\user or user@domain.
5. Enter the base DN where your AD users reside.
6. If AD LDAP communications are protected via SSL enable the Use Secure Connection option.
7. Add one or more AD domain controllers by clicking Add and entering the IP or FQDN of the domain controller.
8. Click Save to enable the configuration.



**If you don't know the base DN to your AD user organizational unit, contact the Active Directory administrator at your organization to provide this information.**

## Remote Authentication ?

Remote Authentication ?

Enable remote authentication and define a remote user or a remote group

Authentication Service:

Administrator Name:

Password:

Base Distinguished Name:

? Disable Nested Group Lookup:

? Use Secure Connection:

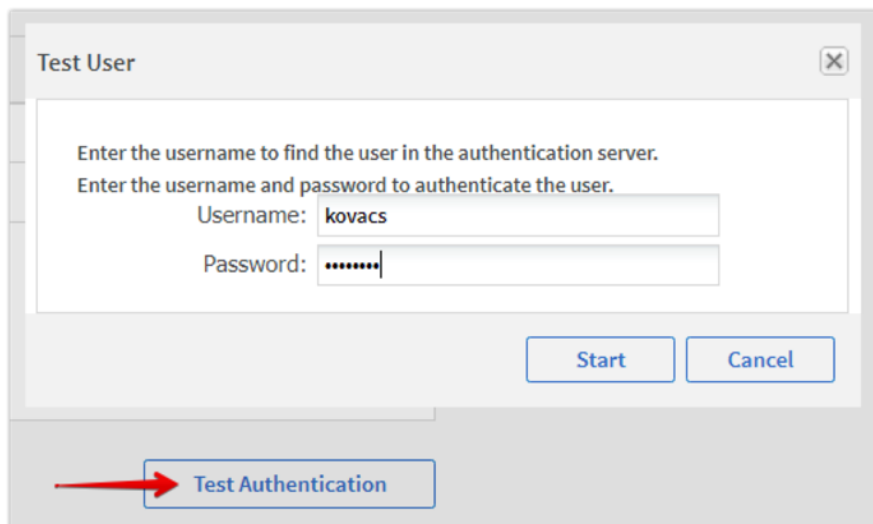
### Authentication Servers

Add Edit Delete

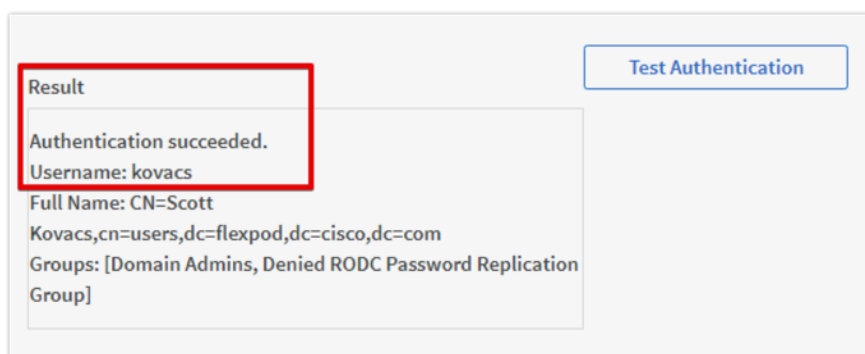
Name or IP Address	Port
10.1.156.251	389
10.1.156.250	389

Save Test Authentication

9. Click Test Authentication and enter an AD username and password to test authentication with the Active Directory authentication servers.



10. A result message should be returned indicating authentication was successful.



### Add a Remote User to Active IQ Unified Manager

To add remote users that need to access Active IQ Unified Manager and authenticate with the Active Directory servers, follow these steps:

1. Navigate to the General section and choose Users.
2. Click Add and choose Remote User from the Type list box.

### Users: Add ?

TYPE

Remote User ▼

NAME

kovacs

EMAIL

\_\_\_\_\_@netapp.com

ROLE

Application Administrator ▼

**Save** Cancel

3. Enter the following information into the form:
  - a. The user name of the Active Directory user.
  - b. Email address of the user.
  - c. Choose the appropriate role for the user
4. Click Save when finished to add the remote user to Active IQ Unified Manager.

### Users ?

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

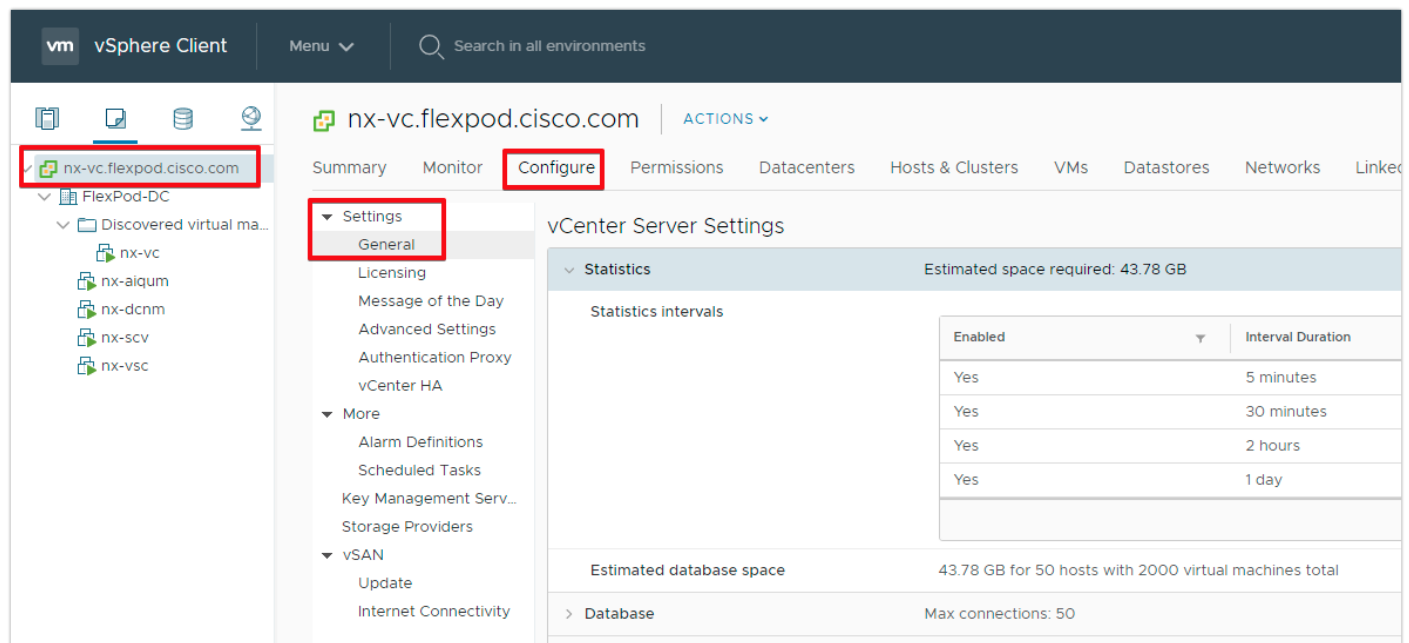
	Name	Type
<input type="checkbox"/>	flexadmin	Maintenance User
<input type="checkbox"/>	kovacs	Remote User

## Add the vCenter Server to Active IQ Unified Manager

Active IQ Unified Manager (AIQ UM) provides visibility into vCenter and the virtual machines running inside the datastores backed by ONTAP storage. Virtual machines and storage are monitored to enable fast identification of performance issues within the various components of the virtual infrastructure stack.

Before adding vCenter into AIQ UM the log level of the vCenter server must be changed by following these steps:

1. In the vSphere client navigate to VMs and Templates and choose the vCenter instance from the top of the object tree.
2. Click the Configure tab, expand the Settings and choose General.



3. Click EDIT.
4. In the pop-up window under Statistics, locate the 5 minutes Interval Duration row and change the setting to Level 3 under the Statistics Level column. Click SAVE.

### Edit vCenter general settings

- Statistics
- Database
- Runtime settings
- User directory
- Mail
- SNMP receivers
- Ports
- Timeout settings
- Logging settings
- SSL settings

#### Statistics

Enter settings for collecting vCenter Server statistics.

Enabled	Interval Duration	Save For	Statistics Level
<input checked="" type="checkbox"/>	5 minutes <span style="font-size: 0.8em;">▼</span>	1 day <span style="font-size: 0.8em;">▼</span>	Level 3 <span style="font-size: 0.8em;">▼</span>
<input checked="" type="checkbox"/>	30 minutes <span style="font-size: 0.8em;">▼</span>	1 week <span style="font-size: 0.8em;">▼</span>	Level 1 <span style="font-size: 0.8em;">▼</span>
<input checked="" type="checkbox"/>	2 hours <span style="font-size: 0.8em;">▼</span>	1 month <span style="font-size: 0.8em;">▼</span>	Level 1 <span style="font-size: 0.8em;">▼</span>
<input checked="" type="checkbox"/>	1 day <span style="font-size: 0.8em;">▼</span>	1 year <span style="font-size: 0.8em;">▼</span>	Level 1 <span style="font-size: 0.8em;">▼</span>

**Database size**  
Based on the current vCenter Server inventory size, the vCenter Server database can be estimated. Enter the expected number of hosts and virtual machines in the inventory to calculate an estimate.

Physical hosts	50	Estimated space required:	43.78 GB
Virtual machines	2000		

Monitor vCenter database consumption and disk partition in Appliance Management UI

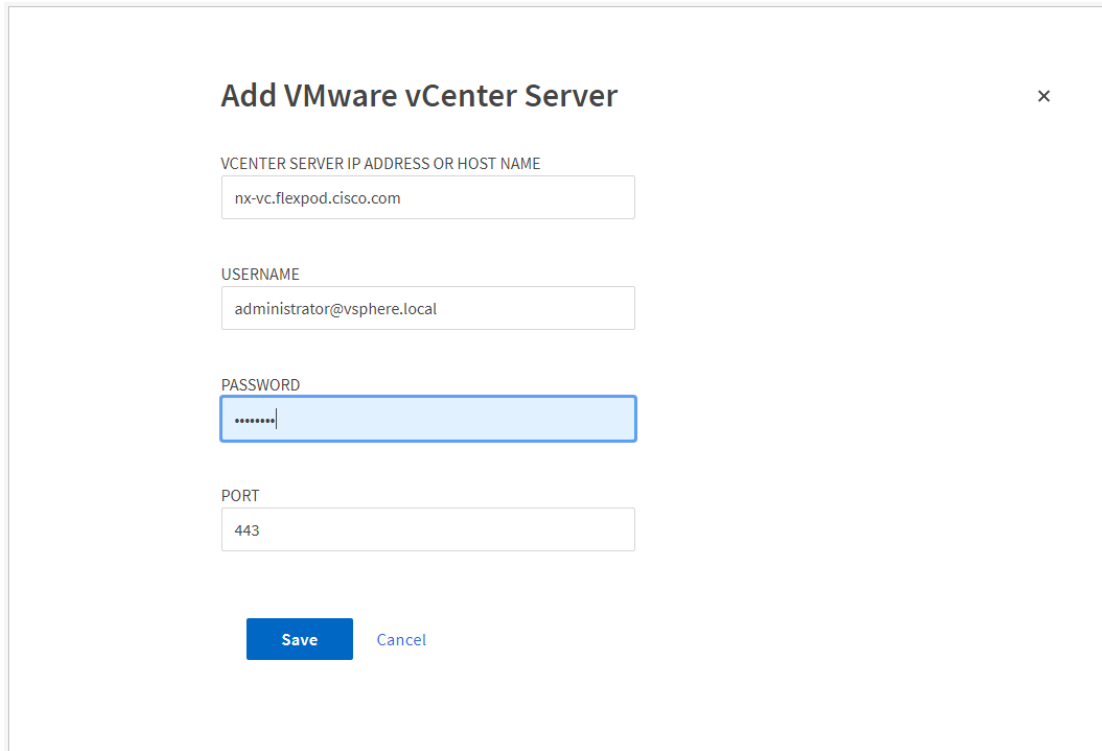
CANCEL
SAVE

5. Return to Active IQ Unified Manager and navigate to the VMware section located under Inventory.

The screenshot shows the Active IQ Unified Manager interface. The left-hand navigation menu is expanded to show the 'VMWARE' section, which is highlighted with a red box. Under 'VMWARE', the 'vCenter' option is also highlighted with a red box. The main content area shows a table for 'vCenters' with columns for Name, Status, IP Address, Version, and Capacity (Used | Total). The table is currently empty, displaying 'No Data'.

6. Expand the section and choose vCenter and click Add.

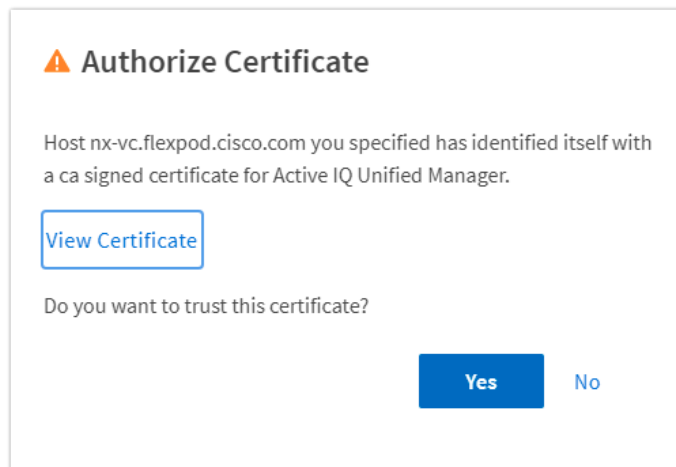
7. Enter the VMware vCenter server details and click Save.



The dialog box is titled "Add VMware vCenter Server" and contains the following fields and buttons:

- VCENTER SERVER IP ADDRESS OR HOST NAME:** nx-vc.flexpod.cisco.com
- USERNAME:** administrator@vsphere.local
- PASSWORD:** [Redacted with dots]
- PORT:** 443
- Buttons:** Save (blue), Cancel (grey)

8. A dialog box will appear asking to authorize the certificate. Click Yes to trust the certificate and add the vCenter server.



The dialog box is titled "Authorize Certificate" and contains the following text and buttons:

- Warning Icon:** A small orange triangle with an exclamation mark.
- Text:** Host nx-vc.flexpod.cisco.com you specified has identified itself with a ca signed certificate for Active IQ Unified Manager.
- Button:** View Certificate (blue)
- Text:** Do you want to trust this certificate?
- Buttons:** Yes (blue), No (grey)



It may take up to 15 minutes to discover the vCenter server. Performance data can take up to an hour after discovery to become available.

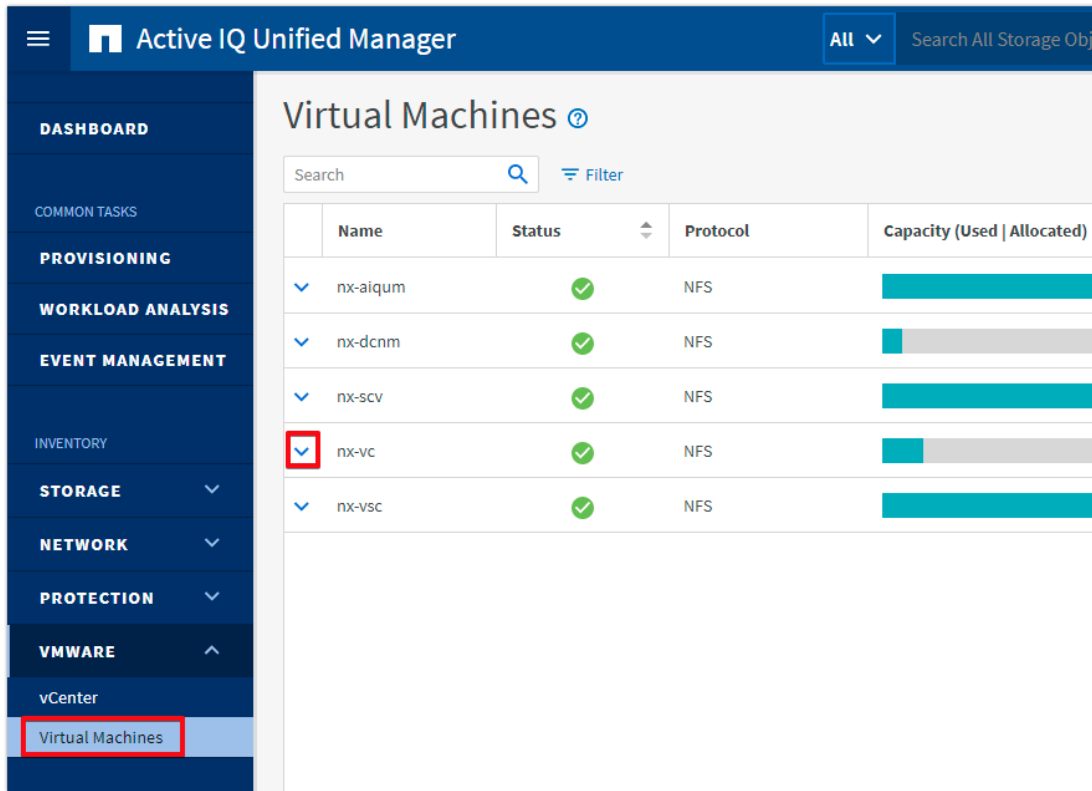
## View Virtual Machine Inventory

The virtual machine inventory is automatically added to Active IQ Unified Manager during discovery of the vCenter server. Virtual machines can be viewed in a hierarchical display detailing storage capacity, IOPS and latency for each component in the virtual infrastructure to troubleshoot the source of any performance related issues.

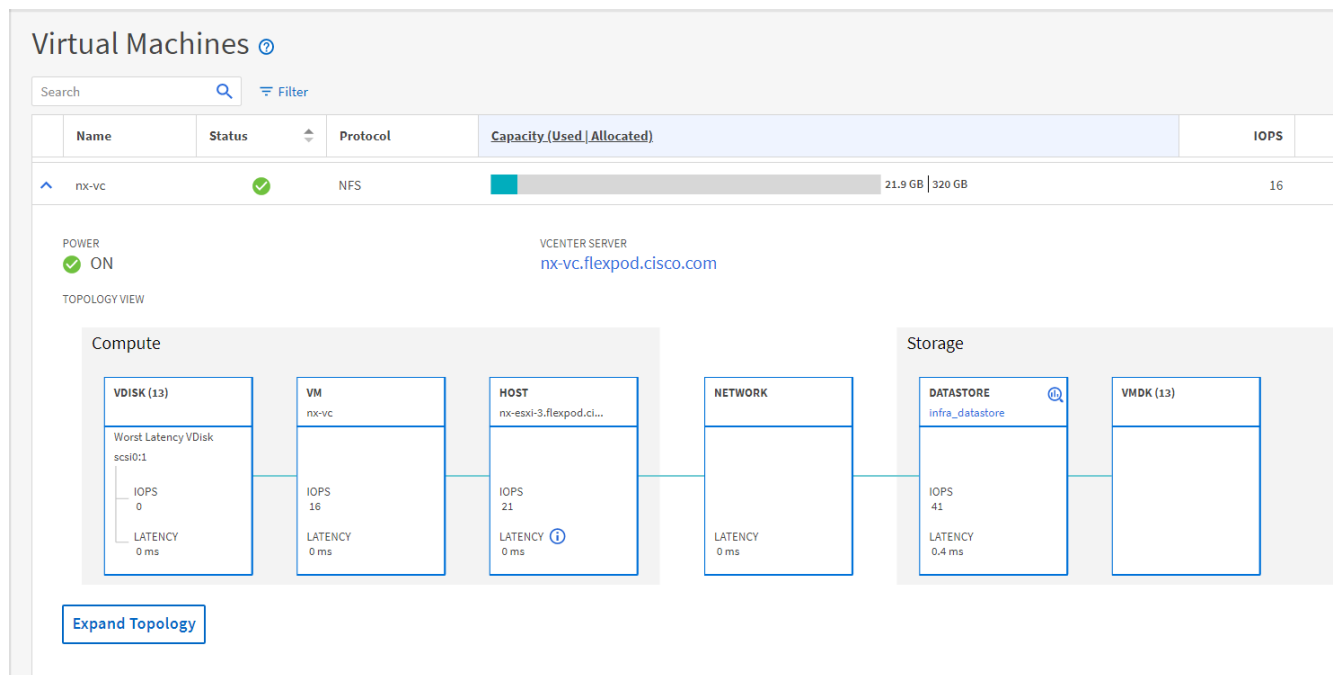


Review the virtual machine topology and statics by following these steps:

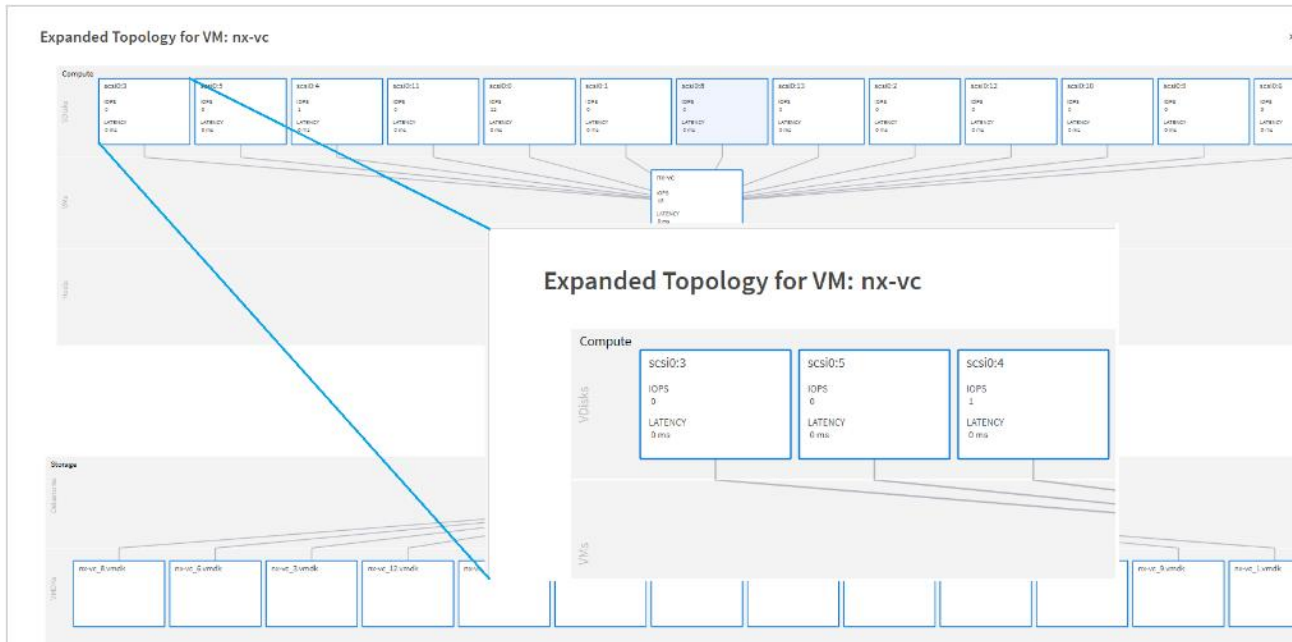
1. Navigate to the VMware section located under Inventory, expand the section and click Virtual Machines.



2. Choose a VM and click the blue caret to expose the topology view. Review the compute, network and storage components and their associated IOPS and latency statistics.






- Click Expand Topology to see the entire hierarchy of the virtual machine and its virtual disks as it is connected through the virtual infrastructure stack. The VM components are mapped from vSphere and compute through the network to the storage.



### Review Security Compliance with Active IQ Unified Manager

Active IQ Unified Manager (AIQ UM) identifies issues and makes recommendations to improve the security posture of ONTAP. AIQ UM evaluates ONTAP storage based on recommendations made in the Security Hardening Guide for ONTAP 9. Items are identified according to their level of compliance with the recommendations. All events identified do not inherently apply to all environments, for example, FIPS compliance. Review the [Security Hardening Guide for NetApp ONTAP 9](#) (TR-4569) for additional information and recommendations for securing ONTAP 9.

The status icons in the security cards have the following meanings in relation to their compliance:

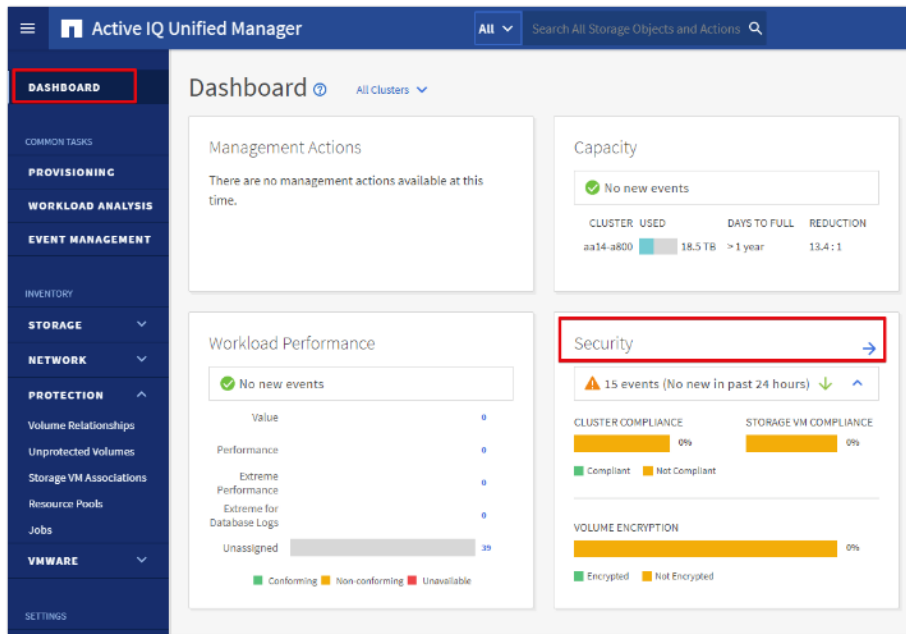
-  - The parameter is configured as recommended.
-  - The parameter is not configured as recommended.
-  - Either the functionality is not enabled on the cluster, or the parameter is not configured as recommended, but this parameter does not contribute to the compliance of the object.

Note that volume encryption status does not contribute to whether the cluster or SVM are considered compliant.

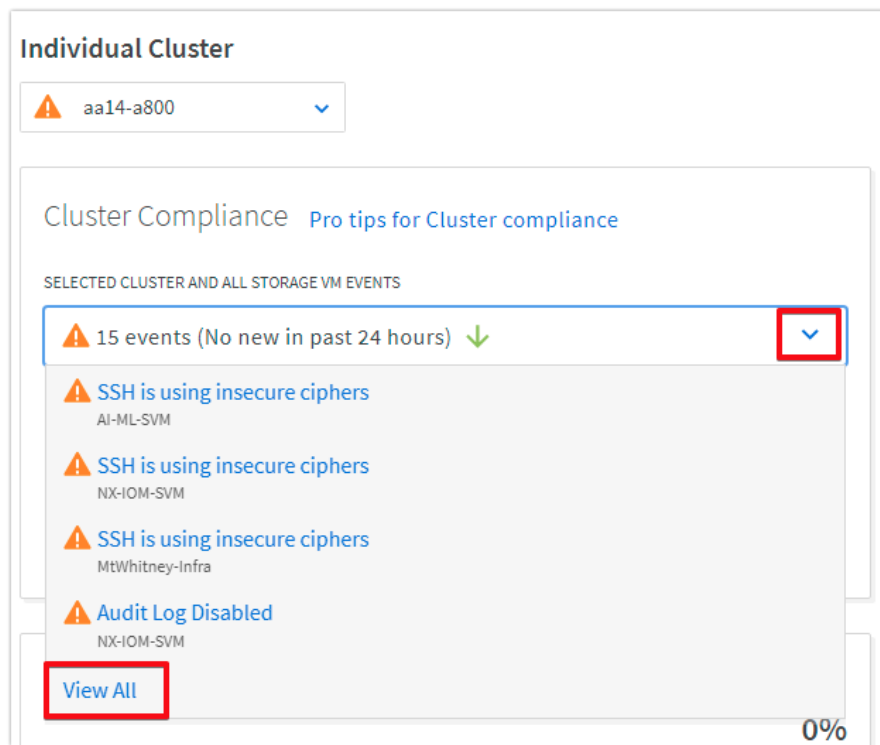


To identify security events in Active IQ Unified Manager, follow these steps:

- Navigate to the URL of the Active IQ Unified Manager installation and login.
- Choose the Dashboard from the left menu bar in Active IQ Unified Manager.
- Locate the Security card and note the compliance level of the cluster and SVM. Click the blue arrow to expand the findings.



4. Locate Individual Cluster section and the Cluster Compliance card. From the drop-down list choose View All.



5. Choose an event from the list and click the name of the event to view the remediation steps.

**Event Management** ⓘ

VIEW  Search Events  Filter 1

<input type="checkbox"/>	Triggered Time	State	Severity	Impact Level	Impact Area	Name
<input type="checkbox"/>	Apr 8, 2020, 10:56 AM	New	⚠	Risk	Security	SSH is using insecure ciphers
<input type="checkbox"/>	Apr 8, 2020, 10:56 AM	New	⚠	Risk	Security	Audit Log Disabled
<input type="checkbox"/>	Apr 8, 2020, 10:56 AM	New	⚠	Risk	Security	Login Banner Disabled

6. Remediate the risk if desired and perform the suggested actions to fix the issue.

**Event: SSH is using insecure ciphers** ⓘ

SSH is using insecure ciphers.

Suggested Actions to Fix The Issue ⓘ

- Ciphers with the suffix CBC are considered insecure.
- To remove the CBC ciphers, run the ONTAP command  

```
security ssh remove -vserver <vserver name> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

## Remediate Security Compliance Findings

Active IQ identifies several security compliance risks after installation that can be immediately corrected to improve the security posture of ONTAP.

### Correct Cluster Risks

To correct cluster risks, follow these steps:

1. Remove insecure SSH ciphers from the cluster administrative SVM:

```
security ssh remove -vserver <clus-adm-svm> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

2. Enable the login banner on the cluster:

```
security login banner modify -vserver <clustername> -message "Access restricted to authorized users"
```

### Correct Infrastructure Storage VM Risks

To correct infrastructure storage VM risks, follow these steps:

1. Remove insecure ciphers from the data SVM:

```
security ssh remove -vserver <infra-data-svm> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

2. Enable the login banner on the SVM:

```
security login banner modify -vserver <infra-data-svm> -message "Access restricted to authorized users"
```

## NetApp Active IQ

NetApp Active IQ is a data-driven service that leverages artificial intelligence and machine learning to provide analytics and actionable intelligence for ONTAP storage systems. Active IQ uses AutoSupport data to deliver proactive guidance and best practices recommendations to optimize storage performance and minimize risk.

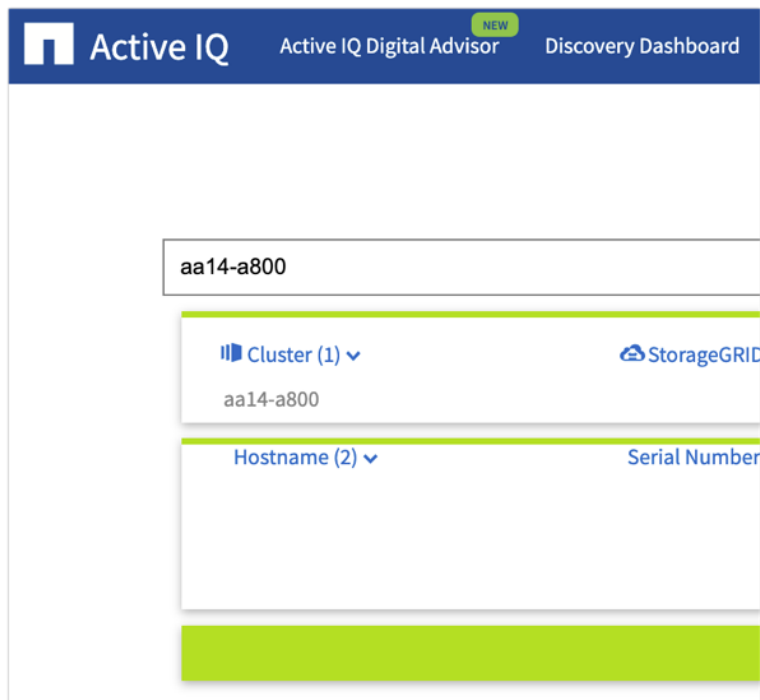
Additional Active IQ documentation is available on the [Active IQ Documentation Resources](#) web page.

Active IQ is automatically enabled when you configure AutoSupport on the ONTAP storage controllers. To get started with Active IQ follow these steps:

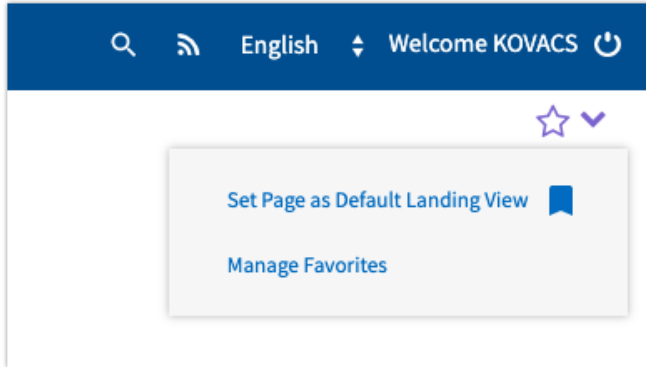
1. Obtain the controller serial numbers from your ONTAP system with the following command:

```
system node show -fields serialnumber
```

2. Navigate to the Active IQ portal at <https://activeiq.netapp.com/>
3. Login with you NetApp support account ID
4. At the welcome screen enter the cluster name or one of controller serial numbers in the search box. Active IQ will automatically begin searching for the cluster and display results below.



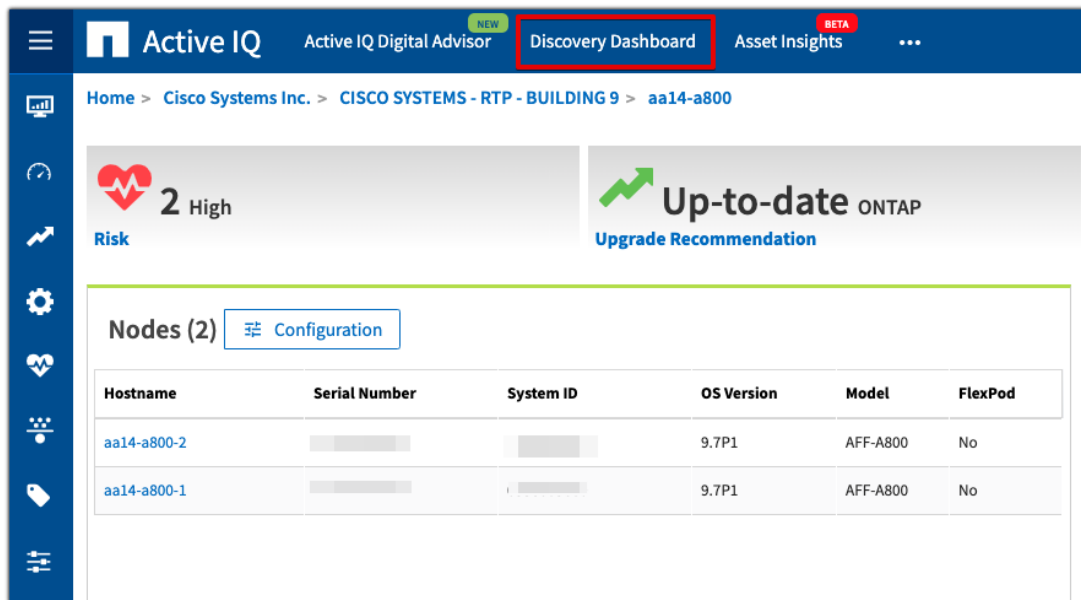
- Choose the cluster name to launch the main dashboard.
- Click the star in the upper right corner to add the cluster to your favorites for easy location in the future. Click the drop-down list to make the cluster dashboard your default view.



### Add a Watchlist to the Discovery Dashboard

The system level dashboard is the default view for systems in Active IQ. To create a watchlist for the quick access cluster to cluster health and risk information, follow these steps:

- Click Discovery Dashboard in the toolbar at the top of the Active IQ screen.



- Click Create Watchlist and enter a name for the watchlist.
- Choose the radio button to add systems by serial number and enter the cluster serial numbers to the watchlist.
- Check the box for Make this my default watchlist if desired and click Create Watchlist.

**Watchlist**

My Watchlists **Create Watchlist**

Name the Watchlist \*  
aa14-a800

Add Systems by ⓘ  
 Category  Serial Number

Choose Category  
Serial Number

Paste Serial Numbers (Maximum Limit 500) \*  
: 83 59

Make this my default watchlist

**Create Watchlist** Cancel

5. Click the ellipsis on the cluster watchlist card you created and click View in Discovery Dashboard.

**Watchlist**

My Watchlists Create Watchlist

1 Watchlist

aa14-a800 Default

Serial Number

View in Discovery Dashboard

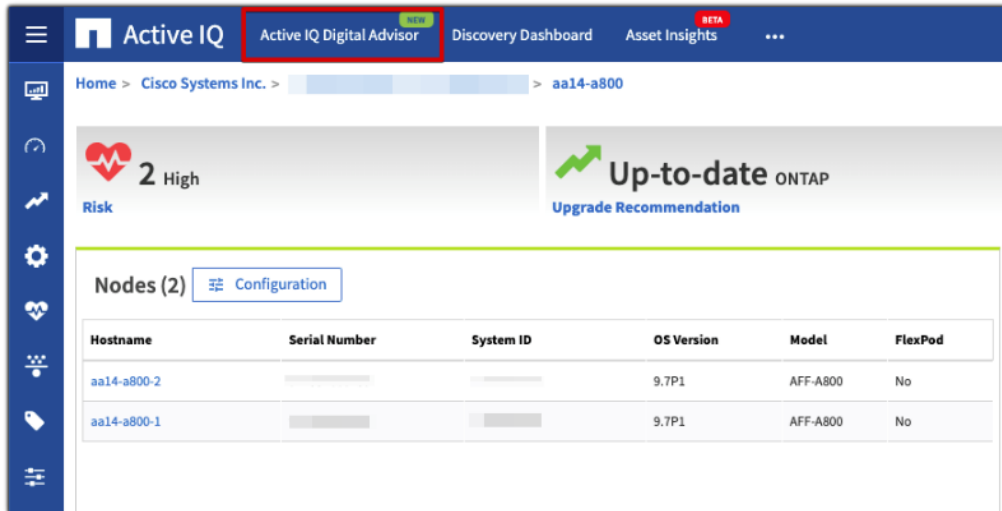
6. View the health and risk overview for the cluster.

## Create Active IQ Digital Advisor Dashboard

The April 2020 release of Active IQ Digital advisor provides a summary dashboard and system wellness score based on the health and risks that Active IQ have identified. The dashboard provides a quick way to identify and get proactive recommendations on how to mitigate risks in the storage environment including links to technical reports and mitigation plans.

To create an Active IQ Digital Advisor dashboard, follow these steps:

1. At the cluster dashboard, click Active IQ Digital Advisor from the top menu.



2. Choose the watchlist created in the previous step and click Next.

**Select Watchlist** +

1 Watchlist found

aa14-a800

**Watchlist Details** \* Mandatory fields

Name the Watchlist \*

aa14-a800

Add Systems by

Category  Serial Number

Choose Category

Serial Number

Paste Serial Numbers (Maximum Limit 500) \*

59 83

Next

3. Accept the dashboard default name and choose all the available widgets.
4. Check the box Make this the default dashboard and click Create.



1 Select or Create Watchlist 2 Create Dashboard

### Create Dashboard using watchlist aa14-a800 \* Mandatory fields

Dashboard name (Ex. Joey) \*  
aa14-a800

**Add widgets**

Inventory  Upgrades  Planning

Make this my default dashboard

[Previous](#) [Create](#)

5. Review the enhanced dashboard including the Wellness Score and any recommended actions or risks.

NetApp Active IQ

Dashboard

Default aa14-a800

+ Add New Dashboard

## aa14-a800

Wellness Score **Good** Actions Risks View All →

- Performance & Efficiency: No Pending Actions
- Availability & Protection: 2 Actions
- Capacity: No Pending Actions
- Configuration: 2 Actions
- Security: No Pending Actions
- Renewal: No Pending Actions

**Inventory** View All →

Overview

ONTAP

2 Systems 1 Cluster 1 Site

**Planning**

- Capacity Addition
- Renewal

**Upgrades**

Upgrades Current Interoperability

No Pending Actions

6. Switch between the Actions and Risks tabs to view the risks broken down by category or a list of all risks with their impact and links to corrective actions.

NetApp Active IQ

Dashboard

Default aa14-a800

+ Add New Dashboard

## aa14-a800

Wellness Score **Good** Actions Risks

- Performance & Efficiency: No Pending Actions
- Availability & Protection: 2 Actions

**Inventory** View All →

aa14-a800 > Wellness

Wellness

All Performance & Efficiency **Availability & Protection** Capacity Configuration Security

Actions **Unique Risks (2)**

View Acknowledged Risks

Filter by All Search by Risk Name

Risk Name	Mitigation	Corrective Action	System	Impact	Acknowledge
On AFF A800 systems an erroneous 'Critical High' sensor reading can result in...	Potentially Non-disruptive	<a href="#">Bug ID: 1279964</a>	2	High	Ack
Digital certificate for a SVM (Storage Virtual Machine) has expired.	Disruptive	<a href="#">KB ID: 27617</a>	2	Medium	Ack

- Click the link in the Corrective Action column to read the bug information or knowledge base article how to remediate the risk.



Additional tutorials and video walk-throughs of Active IQ features can be viewed on the [Active IQ documentation](#) web page.

## Cisco Data Center Network Manager (DCNM)-SAN

Cisco DCNM-SAN can be used to monitor, configure, and analyze Cisco 32Gbps fibre channel fabrics. Cisco DCNM-SAN is deployed as a virtual appliance from an OVA and is managed through a web browser. SAN Analytics can be added to provide insights into your fabric by allowing you to monitor, analyze, identify, and troubleshoot performance issues.

### Prerequisites

The following prerequisites need to be configured:

- Licensing. Cisco DCNM-SAN includes a 60-day server-based trial license that can be used to monitor and configure Cisco MDS Fibre Channel switches and monitor Cisco Nexus switches. Both DCNM server-based and switch-based licenses can be purchased. Additionally, SAN Insights and SAN Analytics requires an additional switch-based license on each switch. Cisco MDS 32Gbps Fibre Channel switches provide a 120-day grace period to trial SAN Analytics.
- Passwords. Cisco DCNM-SAN passwords should adhere to the following password requirements:
  - It must be at least eight characters long and contain at least one alphabet and one numeral.
  - It can contain a combination of alphabets, numerals, and special characters.
  - Do not use any of these special characters in the DCNM password for all platforms: <SPACE> " & \$ % ' ^ = < > ; : ` \ | / , . \*
- DCNM SNMPv3 user on switches. Each switch (both Cisco MDS and Nexus) needs an SNMPv3 user added for DCNM to use to query and configure the switch. On each switch, enter the following command in configure terminal mode (in the example, the userid is snmpuser):

```
snmp-server user snmpadmin network-admin auth sha <password> priv aes-128 <privacy-password>
```

- On Cisco MDS switches, type `show run`. If `snmpadmin passphrase lifetime 0` is present, enter `username snmpadmin passphrase lifetime 99999 warntime 14 gracetime 3`



**It is important to use auth type sha and privacy auth aes-128 for both the switch and UCS snmpadmin users.**

- DCNM SNMPv3 user in UCSM. A SNMPv3 user needs to be added to UCSM to allow DCNM to query the LAN side of the fabric interconnects. In Cisco UCS Manager, click Admin. Navigate to All > Communication Management > Communication Services. Under SNMP, click Enabled, click Save Changes, and then click OK. Under SNMP Users, click Add. Enter the user name and enter and confirm the Password and Privacy Password.

## Create SNMP User



Name	:	<input type="text" value="snmpadmin"/>
Auth Type	:	<b>SHA</b>
Use AES-128	:	<b>Yes</b>
Password	:	<input type="password" value="*****"/>
Confirm Password	:	<input type="password" value="*****"/>
Privacy Password	:	<input type="password" value="*****"/>
Confirm Privacy Password	:	<input type="password" value="*****"/>

**OK** **Cancel**

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

## Deploying the Cisco DCNM-SAN OVA

To deploy the Cisco DCNM-SAN OVA, follow these steps:

- Download the Cisco DCNM 11.3.1 Open Virtual Appliance for VMware from [https://software.cisco.com/download/home/281722751/type/282088134/release/11.3\(1\)](https://software.cisco.com/download/home/281722751/type/282088134/release/11.3(1)). Extract `dcnm-va.11.3.1.ova` from the ZIP file.
- In the VMware vCenter HTML5 interface, click Menu > Hosts and Clusters.
- Right-click the FlexPod-Management cluster and select Deploy OVF Template.

4. Choose Local file then click Choose Files. Navigate to choose dcnm-va.11.3.1.ova and click Open. 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 an OVF template**

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

http | https://remoteserver-address/filetodeploy.ovf | .ova

Local file

Choose Files dcnm-va.11.3.1.ova

CANCEL

BACK

NEXT

5. Name the virtual machine and choose the FlexPod-DC datacenter. Click NEXT.
6. Choose the FlexPod-Management cluster and click NEXT.
7. Review the details and click NEXT.
8. Scroll through and accept the license agreements. Click NEXT.
9. Choose the appropriate deployment configuration size and click NEXT.



**If using the SAN Insights and SAN Analytics feature, it is recommended to use the Huge size.**

## 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 Configuration**
- 7 Select storage
- 8 Select networks
- 9 Customize template
- 10 Ready to complete

### Configuration

Select a deployment configuration

<input type="radio"/> Large (Production)	<b>Description</b> Use this deployment option to configure a huge version of appliance with 32vCPUs and 128GB RAM. This is recommended when using SAN Insights feature.
<input type="radio"/> Small (Lab/PoC)	
<input checked="" type="radio"/> Huge	
<input type="radio"/> Compute	
4 Items	

CANCEL

BACK

NEXT

10. Choose infra\_datastore and the Thin Provision virtual disk format. 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 Configuration
- 7 Select storage**
- 8 Select networks
- 9 Customize template
- 10 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: Thin Provision

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type
 datastore1	7.5 GB	1.41 GB	6.09 GB	VM
 datastore1 (1)	7.5 GB	1.41 GB	6.09 GB	VM
 datastore1 (2)	7.5 GB	1.41 GB	6.09 GB	VM
 Infra_datastore	1 TB	331.23 GB	1,000.1 GB	NF
 Infra_swap	100 GB	13.95 MB	99.99 GB	NF

### Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

11. Choose IB-MGMT Network for all three Source Networks. 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 Configuration
- ✓ 7 Select storage
- 8 Select networks**
- 9 Customize template
- 10 Ready to complete

### Select networks

Select a destination network for each source network.

Source Network	Destination Network
dcnm-mgmt	IB-MGMT Network
enhanced-fabric-mgmt	IB-MGMT Network
enhanced-fabric-inband	IB-MGMT Network

3 items

### IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL

BACK

NEXT

12. Fill in the management IP address, subnet mask, and gateway. Set the Extra Disk Size according to how many Cisco MDS switches you will be monitoring with this DCNM. If you are only monitoring the two Cisco MDS switches in this FlexPod deployment, set this field to 32. Click NEXT.
13. Review the settings and click FINISH to deploy the OVA.
14. After deployment is complete, right-click the newly deployed DCNM VM and click Edit Settings. Expand CPU and adjust the Cores per Socket setting until the number of Sockets is 2 to match the 2-socket UCS servers used in this deployment.

# Edit Settings | nx-dcnm




Virtual Hardware

VM Options

ADD NEW DEVICE

▼ CPU	32 ▼	
Cores per Socket	16 ▼ Sockets: 2	
CPU Hot Plug	<input type="checkbox"/> Enable CPU Hot Add	
Reservation	0 <input type="text"/> MHz ▼	
Limit	Unlimited <input type="text"/> MHz ▼	
Shares	Normal ▼ 32000 <input type="text"/>	
CPUID Mask	Expose the NX/XD flag to guest ▼ <a href="#">Advanced...</a>	
Hardware virtualization	<input type="checkbox"/> Expose hardware assisted virtualization to the guest OS	
Performance Counters	<input type="checkbox"/> Enable virtualized CPU performance counters	
CPU/MMU Virtualization	Automatic ▼	

15. Click OK to complete the change.
16. Right-click the newly deployed DCNM VM and click Open Remote Console. Once the console is up, click to power on the VM. Once the VM has powered up, point a web browser to the URL displayed on the console. 
17. Navigate the security prompts and click Get started.
18. Make sure Fresh installation – Standalone is selected and click Continue.
19. Enter and repeat the admin password and click Next.
20. Choose SAN only for the Installation mode and leave Cisco Systems, Inc. for the OEM vendor and click Next.
21. Enter the DCNM FQDN, a comma-separated list of DNS servers, and a comma-separated list of NTP servers. Click Next.



# Cisco DCNM Installer

Administration   Install Mode   **System Settings**   Network Settings   Applications   HA Settings   Summary

Please enter the following system settings

**Fully Qualified Host Name \***

Fully Qualified Host Name as per RFC1123, section 2.1, for example:  
myhost.mydomain.com

**DNS Server Address List \***

Comma-separated list of DNS Server addresses (IPv4 or IPv6)

**NTP Server Address List \***

Comma-separated list of NTP Server addresses (RFC1123-compliant name, IPv4 or IPv6)

[Previous](#)

[Next](#)

- 22. The Management Network settings should be filled in. For Out-of-Band Network, a different IP address in the same subnet as the management address should be used. Only input the IPV4 address with prefix. Do not put in the Gateway IPv4 Address. Scroll down and click Next.
- 23. Leave Internal Application Services Network set at the default setting and click Next.
- 24. Review the Summary details and click Start installation.

# Cisco DCNM Installer

Administration Install Mode System Settings Network Settings Applications HA Settings **Summary**

Please review the configuration details

<b>Installation Mode</b>	SAN only
<b>OEM Vendor</b>	Cisco Systems, Inc.
<b>Fully Qualified Host Name</b>	nx-dcnm.flexpod.cisco.com
<b>DNS Server Address List</b>	10.1.156.250 10.1.156.251
<b>NTP Server Address List</b>	10.1.156.4 10.1.156.5
<b>Management Network IP Address</b>	10.1.156.102/24
<b>Management Network Default Gateway</b>	10.1.156.254
<b>Management Network IPv6 Address</b>	
<b>Management Network Default IPv6 Gateway</b>	
<b>Out-of-Band Network IP Address</b>	10.1.156.103/24
<b>Out-of-Band Gateway IP Address</b>	
<b>Out-of-Band Network IPv6 Address</b>	
<b>Out-of-Band Gateway IPv6 Address</b>	
<b>In-Band Network IP Address</b>	
<b>In-Band Gateway IP Address</b>	
<b>In-Band Network IPv6 Address</b>	
<b>In-Band Gateway IPv6 Address</b>	
<b>Administration Password</b>	*****
<b>Internal App Services IP Subnet</b>	172.17.0.0/20
<b>Enable Clustered Mode?</b>	No

Previous

Start installation

- When the Installation status is complete, click Continue.
- In the vCenter HTML5 client under Hosts and Clusters, choose the DCNM VM and click the Summary Tab. If an alert is present that states "A newer version of VMware Tools is available for this virtual machine.", click Upgrade VMware Tools. Choose Automatic Upgrade and click UPGRADE. Wait for the VMware Tools upgrade to complete.
- Right-click the DCNM VM and choose Open Remote Console.
- Log in as root with the administrator password entered in the DCNM installation.
- If it is desired to use the local timezone for DCNM and switch logs, set the local timezone. The following example sets the America/New\_York timezone in the Eastern United States:

```
timedatectl list-timezones | grep New
timedatectl set-timezone America/New_York
timedatectl
```

30. If you set the local timezone in DCNM, it is necessary to also set the local timezone in the Cisco Nexus and MDS switches. The following example commands will set the EST/EDT timezone for the Eastern United States in both Cisco Nexus and MDS switches:

```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```




For more information on configuring timezones, see

[https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/fundamentals/configuration/guide/b\\_Cisco\\_Nexus\\_9000\\_Series\\_NX-OS\\_Fundamentals\\_Configuration\\_Guide/b\\_Cisco\\_Nexus\\_9000\\_Series\\_NX-OS\\_Fundamentals\\_Configuration\\_Guide\\_chapter\\_0110.html#task\\_1231769](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/fundamentals/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Fundamentals_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Fundamentals_Configuration_Guide_chapter_0110.html#task_1231769).

## Configuring DCNM-SAN

To configure the DCNM-SAN, follow these steps:

1. When the DCNM installation is complete, the browser should redirect to the DCNM management URL.
2. Log in as admin with the password entered above.
3. On the message that appears, choose Do not show this message again and click No.
4. If you have purchased DCNM server-based or switch-based licenses, follow the instructions that came with the licenses to install them. A new DCNM installation also has a 60-day trial license.
5. In the menu on the left, click Inventory > Discovery > LAN Switches.
6. Click  to add LAN switches. In the Add LAN Devices window, enter the mgmt0 IP address of Nexus switch A in the Seed Switch box. Enter the snmpadmin user name and password set up in the Prerequisites section above. Set Auth-Privacy to SHA\_AES. Click Next.

## Add LAN Devices

**Discovery Type:**  Hops from seed switch  Switch list

**Seed Switch:**

**Max Hops from Seed:**

**User Name:**

**Password:**

**Auth-Privacy:**

**Add Switches To Group:**

**Scan Time:**

- LAN switch discovery will take a few minutes. In the LAN Discovery list that appears, the two Nexus switches and two Fabric Interconnects that are part of this FlexPod should appear with a status of “manageable”. Using the checkboxes on the left, choose the two Nexus switches and two Fabric Interconnects that are part of this FlexPod. Click Add.
- After a few minutes (hit the Refresh icon in the upper right-hand corner), the two Nexus switches and two Fabric Interconnects that are part of this FlexPod will appear with detailed information. The SSH warning under SNMP Status can be ignored since only SNMP can be used to monitor Fabric Interconnects.

### [Inventory / Discovery / LAN Switches](#)

Selected 0 / Total 4

Switch	IP Address	Serial No	Managed	SNMP Status	Role	Last Updated Time	Group	User	Auth/Priv...
AA13-6454-A	192.168.156.18	FDO22191DZ5	true	SSH: There w...	leaf	2020-04-14 06:42:51	Default_LAN	snmpadmin	SHA_AES
AA13-6454-B	192.168.156.19	FDO22191DNN	true	SSH: There w...	leaf	2020-04-14 06:42:51	Default_LAN	snmpadmin	SHA_AES
AA14-9336C-1	192.168.156.11	FDO22272BB0	true	ok	spine	2020-04-14 06:42:54	Default_LAN	snmpadmin	SHA_AES
AA14-9336C-2	192.168.156.12	FDO22272HLE	true	ok	spine	2020-04-14 06:42:54	Default_LAN	snmpadmin	SHA_AES

- In the menu on the left, click Inventory > Discovery > SAN Switches.



- Click  to add a switching fabric.

- Enter either the IP address of hostname of the first Cisco MDS 9132T switch. Leave Use SNMPv3/SSH selected. Set Auth-Privacy to SHA\_AES. Enter the snmpadmin user name and password set up in the Prerequisites section above. Click Options>>. Enter the UCS admin user name and password. Click Add.

## Add Fabric

**Fabric Seed Switch:**   
**SNMP:**  Use SNMPv3/SSH  
**Auth-Privacy:**  ▼  
**User Name:**   
**Password:**   
 Limit Discovery by VSAN  
 Enable NPV Discovery in All Fabrics  
**UCS User Name:**   
**UCS Password:**

12. Repeat steps 1-11 to add the second Cisco MDS 9132T and Fabric Interconnect.

13. The two SAN fabrics should now appear in the Inventory.

🏠 | [Inventory / Discovery / SAN Switches](#)

<input type="button" value="+"/> <input type="button" value="X"/> <input type="button" value="Move"/> <input type="button" value="Rediscover"/> <input type="button" value="Purge"/>						
<input type="checkbox"/>	Name	SeedSwitch	Status	SNMPv3/SSH	User/Cmnty	Auth/P...
<input type="checkbox"/>	Fabric_aa13-9132t-a	192.168.156.13	managedContinuously	true	snmpadmin	SHA_AES
<input type="checkbox"/>	Fabric_aa13-9132t-b	192.168.156.14	managedContinuously	true	snmpadmin	SHA_AES

14. Choose Inventory > Discovery > Virtual Machine Manager.



15. Click  to add the vCenter.

16. In the Add VCenter window, enter the IP address of the vCenter VCSA. Enter the [administrator@vsphere.local](mailto:administrator@vsphere.local) user name and password. Click Add.

The vCenter should now appear in the inventory.

17. Choose Administration > Performance Setup > LAN Collections.

18. Choose the Default\_LAN group and all information you would like to collect. Click Apply. Click Yes to restart the Performance Collector.

## Administration / Performance Setup / LAN Collections

For all selected licensed LAN Switches collect:  Trunks  Access  Errors & Discards  Temperature Sensor

Apply

Default\_LAN

AA14-9336C-1

AA14-9336C-2

19. Choose Administration > Performance Setup > SAN Collections.

20. Choose both fabrics. Choose all information you would like to collect and click Apply. Click Yes to restart the Performance Collector.

## Administration / Performance Setup / SAN Collections

Apply							
		Name	ISL/NPV Links	Hosts	Storage	FC Flows	FC Ethernet
1	<input checked="" type="checkbox"/>	Fabric_aa13-9132t-a	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
2	<input checked="" type="checkbox"/>	Fabric_aa13-9132t-b	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

21. Choose Configure > SAN > Device Alias. Since device-alias mode enhanced was configured in the Cisco MDS 9132T switches, Device Aliases can be created and deleted from DCNM and pushed to the MDS switches.

22. Choose Configure > SAN > Zoning. Just as Device Aliases can be created and deleted from DCNM, zones can be created, deleted, and modified in DCNM and pushed to the MDS switches. Remember to enable Smart Zoning and to Zone by Device Alias.

You can now explore all of the different options and information provided by DCNM SAN. Please see [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/11\\_3\\_1/config\\_guide/sanovaiso/b\\_dcnm\\_san\\_o\\_va-iso.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/11_3_1/config_guide/sanovaiso/b_dcnm_san_o_va-iso.html).

## Configure SAN Insights in DCNM SAN

The SAN Insights feature enables you to configure, monitor, and view the flow analytics in fabrics. Cisco DCNM enables you to visually see health-related indicators in the interface so that you can quickly identify issues in fabrics. Also, the health indicators enable you to understand the problems in fabrics. The SAN Insights feature also provides more comprehensive end-to-end flow-based data from host to LUN.

- Ensure that the time configurations set above, including daylight savings settings are consistent across the MDS switches and Cisco DCNM.
- SAN Insights requires installation of a switch-based SAN Analytics license on each switch. To trial the feature, each switch includes a one-time 120-day grace period for SAN Analytics from the time the feature is first enabled.
- SAN Insights supports current Fibre Channel Protocol (SCSI) and NVMe over Fibre Channel (NVMe).

- SAN Insights works by enabling SAN Analytics and Telemetry Streaming on each switch. The switches then stream the SAN Analytics data to DCNM, which collects, correlates, and displays statistics. All configuration can be done from DCNM.
- For more information on SAN Insights, see the SAN Insights sections of [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/11\\_3\\_1/config\\_guide/san/b\\_dcnm\\_san.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/11_3_1/config_guide/san/b_dcnm_san.html).
- For more information on SAN Analytics, see [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8\\_x/config/san\\_analytics/cisco-mds9000-san-analytics-telemetry-streaming-config-guide-8x.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/san_analytics/cisco-mds9000-san-analytics-telemetry-streaming-config-guide-8x.html).

To configure SAN Insights in DCNM SAN, follow these steps:

1. In the menu on the left, click Configure > SAN > SAN Insights. Click Continue.
2. Choose Fabric A. Click Continue.
3. Choose the Fabric A Cisco MDS switch. Under Install Query click None and from the drop-down list click Storage. Under Subscriptions, choose SCSI. Optionally, under Receiver, choose the second IP address in the In-Band Management subnet configured for DCNM. Click Save, then click Continue.

## 2. Select Switches

Choose the switch(es) on which SAN Insights is to be configured in Fabric\_aa13-9132t-a

DCNM server time: 13:29:20.273 PDT Wednesday April 15 2020

Selected 1 / Total 1

Disable Analytics...									
Show Quick Filter									
<input type="checkbox"/>	Switch	Model	Release	Licensed	Switch Time	Subscriptions	Install Query	Receiver	
<input checked="" type="checkbox"/>	aa13-9132t-a	DS-C9132T-K9	8.4(1a)	Yes	16:29:22.387 EDT Wed Apr 15 2020	SCSI	Storage	10.1.156.102	

4. Review the information and click Continue.
5. Expand the switch and then the module. Under Enable / Disable SCSI Telemetry, click the left icon to enable telemetry. Click Continue.

## 4. Select Interfaces

Choose the switch interfaces that will generate analytics data within Fabric\_aa13-9132t-a

Total Top Level Rows 1

Switch	Module	Interface	Connected To	Type	Analytics Status	Enable / Disable SCSI Telemetry	Enable / Disable NVMe Telemetry
▼ aa13-9132t-a	1 module(s)	4 interface(s)		Storage			
	DS-C9132T-K9-S...	4 interface(s)					
		fc1/1	Infra-SVM-fc-lif-1a	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	pending enable
		fc1/2	20:31:00:a0:98:e2:17:ca	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	pending enable

6. Review the information and click Commit to push the configuration to the Cisco MDS switch.


7. Ensure that the two operations were successful and click Close.
8. Repeat this process to install SAN Analytics and Telemetry on the Fabric B switch.
9. After approximately two hours, you can view SAN Analytics data under the Dashboard and Monitor.

## Cisco Intersight

Cisco Intersight™ is a management platform delivered as a service with embedded analytics for your Cisco and 3rd party IT infrastructure. This platform offers an intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in more advanced ways than the prior generations of tools. Cisco Intersight provides an integrated and intuitive management experience for resources in the traditional data center and at the edge. With flexible deployment options to address complex security needs, getting started with Intersight is quick and easy.

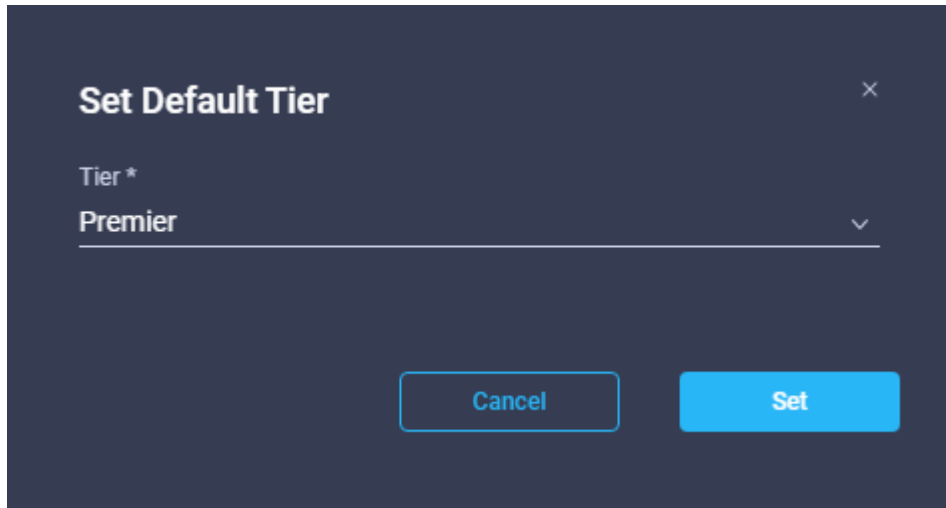
Cisco Intersight offers flexible deployment either as Software as a Service (SaaS) on Intersight.com or running on your premises as Cisco Intersight Virtual Appliance. The virtual appliance provides the benefits of Cisco Intersight while allowing more flexibility for those with additional data locality and security requirements. The remainder of this section details Intersight deployment as SaaS on Intersight.com. To learn more about the virtual appliance, see the [Cisco Intersight Virtual Appliance Getting Started Guide](#).



To configure Cisco Intersight, follow these steps:


1. If you do not already have a Cisco Intersight account, to claim your Cisco UCS system into a new account on Cisco Intersight, connect to <https://Intersight.com>.
2. Click Create an account.
3. Click Continue. Complete the Sign in process with your Cisco ID.
4. Read the Offer Description carefully and accept it. Click Next.
5. Enter an Account Name, Device ID, and Claim Code. The Device ID and Claim Code can be obtained by connecting to Cisco UCS Manager and selecting Admin > All > Device Connector. The Device ID and Claim Code are on the right. Click Claim.
6. Click Create. After the account has been created, click Log me in to log into Cisco Intersight.
7. To claim your Cisco UCS system into an existing Intersight account, log into the account at <https://Intersight.com>. Choose Administration > Devices. Click Claim a New Device. Under Direct Claim, fill in the Device ID and Claim Code. The Device ID and Claim Code can be obtained by connecting to Cisco UCS Manager and selecting Admin > All > Device Connector. The Device ID and Claim Code are on the right.
8. From the Cisco Intersight window, click  and then click Licensing. If this is a new account, all servers connected to the UCS Domain will appear under the Base license Tier. If you have purchased Cisco Intersight licenses and have them in your Cisco Smart Account, click Register and follow the prompts to register this Cisco Intersight account to your Cisco Smart Account. Cisco Intersight also offers a one-time 90-day trial of Premier licensing for new accounts. Click Start Trial and then Start to begin this evaluation. The remainder of this section will assume Premier licensing.

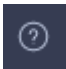


- From the Licensing Window, click Set Default Tier. From the drop-down list choose Premier for Tier and click Set.



- To set all Cisco UCS Servers to Premier licensing, click Servers. Click  to the left of the Name heading to choose all servers. Click  above the headings and click Set License Tier. From the drop-down list choose Premier for the Tier and click Set License Tier.



- Click Refresh to refresh the Intersight window with Premier, Advantage, and Essentials features added.
- Click  in the Intersight window and click Take a Site Tour. Follow the prompts for a tour of Cisco Intersight.
- The Essentials tier of Cisco Intersight includes a Cisco driver check against the Cisco Hardware Compatibility List (HCL). In the Servers list, choose one of the servers in your VMware FlexPod-Management cluster by clicking the server name. Review the detailed General and Inventory information for the server. Click the HCL tab. Review the server information, the version of VMware ESXi, and the Cisco VIC driver versions.

The screenshot displays the HCL Validation section of a management console. It features three numbered steps, each with a 'Validated' status:

- 1 Server Hardware Compliance (Validated):**
  - Server Model: UCSB-B200-M5
  - CPU: Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
  - Server Firmware Version: 4.1(1b)B
- 2 Server Software Compliance (Validated):**
  - OS Vendor: VMware ESXi
  - OS Version: 6.7.0.3
- 3 Adapter Compliance (Validated):**

Below the validation steps is a table with 2 items found. The table has columns for Model, Hardware Status, Software Status, Firmware Version, Driver Protocol, and Driver Version.

Model	Hardware Status	Software Status	Firmware Version	Driver Protocol	Driver Version
UCSB-MLOM-40G-04	Validated	Validated	5.1(1e)	nenic	1.0.31.0-10EM.670.0.0.81
UCSB-MLOM-40G-04	Validated	Validated	5.1(1e)	nfnic	4.0.0.48-10EM.670.0.0.81

14. Using the Intersight Assist personality of the Cisco Intersight Virtual Appliance VMware vCenter currently can be monitored (Advantage Licensing Tier) and configured (Premier Licensing Tier). To install Intersight Assist from an Open Virtual Appliance (OVA) in your VMware FlexPod-Management Cluster, first download the OVA from <https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-113>.
15. Refer to [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/Intersight/cisco-intersight-assist-getting-started-guide/m-installing-cisco-intersight-assist.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/cisco-intersight-assist-getting-started-guide/m-installing-cisco-intersight-assist.html) and set up the DNS entries for the Intersight Assist hostname as specified under Before you begin.
16. From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the FlexPod-Management cluster and click Deploy OVF Template.
17. Specify a URL or either browse to the intersight-virtual-appliance-1.0.9-113.ova file. 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 an OVF template**

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

http | https://remoteserver-address/filetodeploy.ovf | .ova

Local file

intersight-virtua...ce-1.0.9-113.ova

CANCEL

BACK

NEXT

- 18. Name the Intersight Assist VM and choose the location. Click NEXT.
- 19. Choose the FlexPod-Management cluster and click NEXT.
- 20. Review details and click NEXT.
- 21. Choose a deployment configuration (Tiny recommended) 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 Configuration**
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

### Configuration

Select a deployment configuration

	Description
<input type="radio"/> Small(16 vCPU, 32 Gi RAM)	Deployment size supports Intersight Assist only.
<input type="radio"/> Medium(24 vCPU, 64 Gi RAM)	
<input checked="" type="radio"/> Tiny(8 vCPU, 16 Gi RAM)	
3 Items	

CANCEL

BACK

NEXT

22. Choose infra\_datastore for storage and choose the Thin Provision virtual disk format. Click NEXT.

23. Choose IB-MGMT Network for the VM Network. Click NEXT.

24. Fill in all values to customize the template. Click NEXT.

25. Review the deployment information and click FINISH to deploy the appliance.

26. Once the OVA deployment is complete, right-click the Intersight Assist VM and click Edit Settings.

27. Expand CPU and adjust the Cores per Socket so that 2 Sockets are shown. Click OK.

Edit Settings | nx-intersight-assist
✕

Virtual Hardware
VM Options

ADD NEW DEVICE

<b>▼ CPU</b>	8	▼		i
Cores per Socket	4	▼	Sockets: 2	
CPU Hot Plug	<input checked="" type="checkbox"/> Enable CPU Hot Add			
Reservation	0	▼	MHz	▼
Limit	Unlimited	▼	MHz	▼
Shares	Normal	▼	8000	
CPUID Mask	Expose the NX/XD flag to guest ▼ <span style="color: #0070c0;">Advanced...</span>			
Hardware virtualization	<input type="checkbox"/> Expose hardware assisted virtualization to the guest OS			
Performance Counters	<input type="checkbox"/> Enable virtualized CPU performance counters			
CPU/MMU Virtualization	Automatic			▼
> Memory	16	▼	GB	▼
> Hard disks	8 total   500 GB			
> SCSI controller 0	LSI Logic SAS			

CANCEL
OK

28. Right-click the Intersight Assist VM and choose Open Remote Console.

29. Click  to power on the VM.

30. When you see the login prompt, close the Remote Console and connect to <https://intersight-assist-fqdn>.

31. Navigate the security prompts and select Intersight Assist. Click Proceed.

### What would you like to Install ?

Intersight Connected Virtual Appliance ?

Intersight Assist ?


 Recover from backup

**Proceed**

32. From Cisco Intersight, click Administration > Devices. Click Claim a New Device. Copy and paste the Device ID and Claim Code shown in the Intersight Assist web interface to the Cisco Intersight Device Claim Direct Claim window. In Cisco Intersight, click Claim.
33. In the Intersight Assist web interface, click Connect Intersight Virtual Appliance. This will open a new tab showing the Intersight login. Close this tab and return to the Intersight Assist web interface tab. The Device Connector should now show Claimed. Click Continue.
34. The Intersight Assist software will now be downloaded and installed into the Intersight Assist VM. This can take up to an hour to complete.
35. When the software download is complete, an Intersight Assist login screen will appear. Log into Intersight Assist with the admin@local user and the password supplied in the OVA installation. Check the Intersight Assist status and log out of Intersight Assist.
36. To claim the vCenter, from Cisco Intersight, click Administration > Devices. Click Claim a New Device. In the Device Claim window, choose Claim Through Intersight Assist. Fill in the vCenter information and click Claim.

- 37. After a few minutes, the VMware vCenter will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.
- 38. Detailed information obtained from the vCenter can now be viewed by clicking Virtualization from the menu.
- 39. Click Orchestration. VMware workflows can be created by clicking Create New Workflow. Click the Tasks tab and scroll down to the Virtualization tasks. Virtualization tasks can be stitched together into workflows to automate VMware functions.

---

 **The Cisco Intersight Orchestration feature is currently a preview feature for testing and feedback purposes only. Please do not use this on production systems. Also, only the Virtualization tasks can be used to build workflows.**

---

▲ This is a preview feature for testing and feedback purposes only. Please do not use this on production systems.

Tools ☰ 🗑️ 👁️ Mapping 🖋️ Properties { } JSON View 🕒 History ⚡ Execute

Tasks Workflows

🔍 Search

**Virtualization**

- Confirm OVA or OVF Installation
- Confirm Virtual Machine Provisioning
- Deploy Virtual Machine from OVA or OVF
- Expand Hypervisor Datastore
- Find Hypervisor Storage
- Get Hypervisor Datastore
- Invoke Guest Customization for Linux Virtual Machine
- Invoke Guest Customization for Windows Virtual Machine
- New Hypervisor Cluster
- New Hypervisor Datacenter
- New Hypervisor Host
- New VMFS Datastore
- New Virtual Machine from Template
- Remove Hypervisor Cluster
- Remove Hypervisor Datacenter
- Remove Hypervisor Host
- Remove VMFS Datastore
- Rename Hypervisor Datacenter

Start

Completed Failed

Workflow Properties

General Inputs Outputs

Organization \*

Workflow Name \*

Description

Version 1

Retryable

Add Tag

Close Designer ● Save the workflow to validate Save Workflow



## 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. If the infrastructure NFS VLAN will be used in the tenant, consider migrating the infrastructure NFS VMkernel port on each host to the vDS to take advantage of Ethernet adapter policy queuing. 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. Add the tenant host initiators to the Infra-SVM zone.
  - c. Create a zone for the tenant SVM with fibre channel targets from the tenant SVM.
  - d. 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 HTML5 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 HTML5 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

To create the block protocol iSCSI service, follow these steps:



**If the FCP protocol is not being used in the environment it should be removed from the vserver 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 vserver.

```
vserver remove-protocols -vserver <infra-data-svm> -protocols fcp
```

2. Enable the iSCSI protocol on the vservers.

```
vserver add-protocols -vserver <infra-data-svm> -protocols iscsi
```

3. Create the iSCSI block service.

```
vserver iscsi create -vserver <infra-data-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-data-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-data-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-data-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-data-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. Choose 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. Choose 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. Choose 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. Choose 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. Choose 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. Choose 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. Choose 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. Choose 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. Choose LAN.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click vNIC Templates under the FlexPod Organization.
4. Choose Create vNIC Template.
5. Enter iSCSI-Template-A as the vNIC template name.
6. Choose Fabric A. Do not choose 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. Choose Updating Template for Template Type.
10. Under VLANs, choose only Infra-iSCSI-A.
11. Choose 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, choose MAC-Pool-A.
15. From the Network Control Policy list, choose 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	<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 : 9000

MAC Pool : MAC-Pool-A(256/256) ▼

QoS Policy : <not set> ▼

Network Control Policy : Enable-CDP-LLDP ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

OK

Cancel

- Click OK to complete creating the vNIC template.
- Click OK.
- Right-click vNIC Templates.
- Choose Create vNIC Template.
- Enter iSCSI-Template-B as the vNIC template name.

21. Choose Fabric B. Do not choose 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. Choose Updating Template for Template Type.
25. Under VLANs, choose only Infra-iSCSI-B.
26. Choose 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, choose MAC-Pool-B.
30. From the Network Control Policy list, choose 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. Choose Create LAN Connectivity Policy.
5. Enter iSCSI-Boot as the name of the policy.
6. Click OK then OK again to create the policy.
7. On the left under LAN > Policies > root > Sub-Organizations > FlexPod Organization > LAN Connectivity Policies, choose iSCSI-Boot.
8. Click the Add button to add a vNIC.
9. In the Create vNIC dialog box, enter 00-vSwitch0-A as the name of the vNIC.
10. Choose the Use vNIC Template checkbox.
11. In the vNIC Template list, choose vSwitch0-A.
12. In the Adapter Policy list, choose VMWare.

13. Click OK to add this vNIC to the policy.
14. Click Save Changes and OK.
15. Click the Add button to add another vNIC to the policy.
16. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.
17. Choose the Use vNIC Template checkbox.
18. In the vNIC Template list, choose vSwitch0-B.
19. In the Adapter Policy list, choose VMWare.
20. Click OK to add the vNIC to the policy.
21. Click Save Changes and OK.
22. Click the Add button to add a vNIC.
23. In the Create vNIC dialog box, enter 02-vDS0-A as the name of the vNIC.
24. Choose the Use vNIC Template checkbox.
25. In the vNIC Template list, choose vDS0-A.
26. In the Adapter Policy list, choose VMWare-HighTrf.
27. Click OK to add this vNIC to the policy.
28. Click Save Changes and OK.
29. Click the Add button to add another vNIC to the policy.
30. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.
31. Choose the Use vNIC Template checkbox.
32. In the vNIC Template list, choose vDS0-B.
33. In the Adapter Policy list, choose VMWare-HighTrf.
34. Click OK to add the vNIC to the policy.
35. Click Save Changes and OK.
36. Click the Add button to add a vNIC.
37. In the Create vNIC dialog box, enter 04-iSCSI-A as the name of the vNIC.
38. Choose the Use vNIC Template checkbox.

39. In the vNIC Template list, choose iSCSI-Template-A.
40. In the Adapter Policy list, choose VMWare.
41. Click OK to add this vNIC to the policy.
42. Click Save Changes and OK.
43. Click Add to add a vNIC to the policy.
44. In the Create vNIC dialog box, enter 05-iSCSI-B as the name of the vNIC.
45. Choose the Use vNIC Template checkbox.
46. In the vNIC Template list, choose iSCSI-Template-B.
47. In the Adapter Policy list, choose VMWare.
48. Click OK to add this vNIC to the policy.
49. Click Save Changes and OK.
50. Expand Add iSCSI vNICs.
51. Choose Add in the Add iSCSI vNICs section.
52. Set the name to iSCSI-Boot-A.
53. Choose 04-iSCSI-A as the Overlay vNIC.
54. Set the iSCSI Adapter Policy to default.
55. Leave the VLAN set to Infra-iSCSI-A (native).
56. Leave the MAC Address set to None.
57. Click OK.
58. Click Save Changes and OK.
59. Choose Add in the Add iSCSI vNICs section.
60. Set the name to iSCSI-Boot-B.
61. Choose 05-iSCSI-B as the Overlay vNIC.
62. Set the iSCSI Adapter Policy to default.
63. Leave the VLAN set to Infra-iSCSI-B (native).
64. Leave the MAC Address set to None.

65. Click OK.

66. Click Save Changes and OK.

General Events

---

**Actions**

Delete

Show Policy Usage

Use Global

Name : **iSCSI-Boot**

Description :

Owner : **Local**

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

## 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-lif-1a and iscsi-lif-1b) and two iSCSI LIFs are on cluster node 2 (iscsi-lif-2a and iscsi-lif-2b). 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-lif-1a.**



**It is not recommended to use UEFI Secure Boot for iSCSI Boot at this time. The Legacy boot mode will be used. If a TPM 2.0 module is installed in a UCS server using Legacy boot, whether enabled or disabled in the BIOS policy, a TPM Attestation Failed alert will appear in VMware vCenter on every reboot of the server.**

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. Choose 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 choose the Reboot on Boot Order Change checkbox.
8. Choose the Legacy Boot Mode.
9. Expand the Local Devices drop-down menu and click Add Remote CD/DVD.
10. Expand the iSCSI vNICs drop-down menu and click Add iSCSI Boot.
11. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.
12. Click OK.
13. Choose 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...
<b>Remote CD/DVD</b>	1								
<b>iSCSI</b>	2								
iSCSI		iSCSI-Boot-A	Pri...						
iSCSI		iSCSI-Boot-B	Sec...						
<b>CIMC Mounted CD/DVD</b>	3								

↑ Move Up ↓ Move Down Delete

Set Uefi Boot Parameters

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. Choose 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. Choose the Updating Template option.
7. Under UUID Assignment, choose UUID\_Pool.

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 choose the SAN-Boot Local Storage Policy. Otherwise, choose the default Local Storage Policy.
2. Click Next.

### Configure Networking Options

To configure the network options, follow these steps:

1. Choose the “Use Connectivity Policy” option to configure the LAN connectivity.
2. Choose iSCSI-Boot from the LAN Connectivity Policy drop-down list.

- Choose 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](#)

- Click Next.

### Configure Storage Options

To configure the storage options, follow these steps:

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

### Configure Zoning Options

To configure the zoning options, follow this step:

- Make no changes and click Next.

### Configure vNIC/HBA Placement

To configure the vNIC/HBA placement, follow these steps:

- In the “Select Placement” list, leave the placement policy as “Let System Perform Placement”.
- 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. Choose 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 Numb...	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 choose 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. Choose 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="&lt;not set&gt;"/>	<a href="#">Create iSCSI Authentication Profile</a>
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
<b>iqn.1992-08....</b>	1	3260		192.168.10.62	0
<b>iqn.1992-08....</b>	2	3260		192.168.10.61	0

**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, choose 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. Choose 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
<b>iqn.1992-08....</b>	1	3260		192.168.20.62	0
<b>iqn.1992-08....</b>	2	3260		192.168.20.61	0

[+](#) Add [-](#) Delete [i](#) Info

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

OK

Cancel

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	: <b>default</b>
Description	:
Soft Shutdown Timer	: <b>150 Secs</b>
Storage Config. Deployment Policy	: <b>User Ack</b>
Reboot Policy	: <b>User Ack</b>

< Prev    Next >    **Finish**    Cancel

2. Click Next.

## Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, choose Infra-Pool.
2. Choose Down as the power state to be applied when the profile is associated with the server.
3. Optional: choose “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 choose 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 :

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

**Create Service Profile Template**

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

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

BIOS Policy :

⊕ 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. Choose Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and click 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. Choose the newly-created VM-Host-Infra-iSCSI-A-vM and choose the vMedia Policy tab.
6. Click Modify vMedia Policy.
7. Choose the ESXi-6.7U3-HTTP vMedia Policy and click OK.
8. Click OK to confirm.

## Create Service Profile Template for Servers with Optane Memory

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlexPod > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and click Create a Clone.
4. Name the clone Optane-Host-Infra-iSCSI-A.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly-created Optane-Host-Infra-iSCSI-A and choose the Policies tab.
7. Expand Persistent Memory Policy and choose the App-Direct-Mode policy.
8. Click Save Changes and OK to confirm.

## Create vMedia-Enabled Service Profile Template for Servers with Optane Memory

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlexPod > Service Template Optane-Host-Infra-iSCSI-A.
3. Right-click Optane-Host-Infra-iSCSI-A and click Create a Clone.
4. Name the clone Optane-Host-Infra-iSCSI-A-vM.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly-created Optane-Host-Infra-iSCSI-A vM and select the vMedia Policy tab.
7. Click Modify vMedia Policy.
8. Choose the ESXi-6.7U3-HTTP vMedia Policy and click OK.
9. 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. Choose 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 choose Create Service Profiles from Template.

4. For Naming Prefix, enter VM-Host-Infra-0.
5. For Name Suffix Starting Number, enter 1.
6. For Number of Instances, enter 3.

## 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.



**If the Service Profiles being built will be associated with servers equipped with Intel Optane DC PMEM, use the Optane-Host-Infra-FCP-A-VM service profile template instead.**

When VMware ESXi 6.7 U3 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 three 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-data-svm> -igroup vm-host-infra-01 -protocol iscsi
-ostype vmware -initiator <vm-host-infra-01-iqn>

lun igroup create -vserver <infra-data-svm> -igroup vm-host-infra-02 -protocol iscsi
-ostype vmware -initiator <vm-host-infra-02-iqn>

lun igroup create -vserver <infra-data-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-data-svm> -path /vol/esxi_boot/vm-host-infra-01 -
igroup vm-host-infra-01 -lun-id 0

lun mapping create -vserver <infra-data-svm> -path /vol/esxi_boot/vm-host-infra-02 -
igroup vm-host-infra-02 -lun-id 0

lun mapping create -vserver <infra-data-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 Navigator, click Networking.
2. In the center pane, choose the Virtual switches tab.
3. Highlight the iScsiBootvSwitch line.
4. Choose Edit settings.

5. Change the MTU to 9000.
6. Choose Add uplink to add an uplink to iScsiBootvSwitch.
7. From the drop-down list select vmnic5 for Uplink 2.
8. Expand NIC teaming, choose vmnic5, and choose Mark standby.

✎ **Edit standard virtual switch - iScsiBootvSwitch**

✎ Add uplink

MTU	<input style="width: 90%;" type="text" value="9000"/>									
Uplink 1	<div style="border: 1px solid #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span>vmnic4 - Up, 20000 mbps</span> <span>⌵</span> <span style="color: #e67e22;">✕</span> </div>									
Uplink 2	<div style="border: 1px solid #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span>vmnic5 - Up, 20000 mbps</span> <span>⌵</span> <span style="color: #e67e22;">✕</span> </div>									
▶ Link discovery	Click to expand									
▶ Security	Click to expand									
▼ NIC teaming										
Load balancing	<div style="border: 1px solid #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span>Route based on originating port ID</span> <span>⌵</span> </div>									
Network failover detection	<div style="border: 1px solid #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span>Link status only</span> <span>⌵</span> </div>									
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: 5px;"> <div style="border: 1px solid #e67e22; padding: 2px; display: flex; align-items: center; gap: 5px;"> <span style="color: #e67e22;">✎</span> Mark active           </div> <div style="display: flex; align-items: center; gap: 5px;"> <span>⇅↑</span> Move up           <span>⇅↓</span> Move down         </div> </div> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr> <th style="width: 40%;">Name</th> <th style="width: 30%;">Speed</th> <th style="width: 30%;">Status</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px;"><span style="color: #e67e22;">✎</span> vmnic4</td> <td style="padding: 2px;">20000 Mbps, full duplex</td> <td style="padding: 2px;">Active</td> </tr> <tr style="background-color: #e6f2e6;"> <td style="padding: 2px;"><span style="color: #e67e22;">✎</span> vmnic5</td> <td style="padding: 2px;">20000 Mbps, full duplex</td> <td style="padding: 2px;">Standby</td> </tr> </tbody> </table>	Name	Speed	Status	<span style="color: #e67e22;">✎</span> vmnic4	20000 Mbps, full duplex	Active	<span style="color: #e67e22;">✎</span> vmnic5	20000 Mbps, full duplex	Standby
Name	Speed	Status								
<span style="color: #e67e22;">✎</span> vmnic4	20000 Mbps, full duplex	Active								
<span style="color: #e67e22;">✎</span> vmnic5	20000 Mbps, full duplex	Standby								
▶ Traffic shaping	Click to expand									

Save

Cancel

9. Click Save.
10. Choose the VMkernel NICs tab.
11. Choose the vmk1 iScsiBootPG line. Choose 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 in Cisco UCS.**

---

14. Click Save to save the changes to the VMkernel port.
15. Choose the Port groups tab.
16. Choose the iScsiBootPG line. Choose Edit Settings to edit the properties of this port group.
17. Expand NIC teaming and click Yes to the right of Override failover order.
18. Choose vmnic5 and click 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;"> <span> Mark active</span> <span> Mark unused</span> <span> Move up</span> <span> Move down</span> </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 style="background-color: #e0f0ff;"> <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. Choose the iScsiBootvSwitch line and click Edit settings.

22. Expand NIC teaming and click vmnic5. Choose 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, choose the VMkernel NICs tab.

25. Click Add VMkernel NIC.

26. For New port group, enter iScsiBootPG-B

27. For Virtual switch, choose iScsiBootvSwitch.

28. Leave the VLAN ID set at 0.
29. Change the MTU to 9000.
30. Choose 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 in Cisco UCS.**











---

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
<input type="button" value="Create"/> <input type="button" value="Cancel"/>	

33. Click Create.
34. Choose the Port groups tab.
35. Choose the iScsiBootPG-B line. Choose Edit Settings to edit the properties of this port group.
36. Expand NIC teaming and click Yes to the right of Override failover order.
37. To the right of Failover order, select vmnic4 and click 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;"> <span> Mark active</span> <span> Mark unused</span> <span> Move up</span> <span> Move down</span> </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 choose Storage, then in the center pane choose the Adapters tab.

40. Click Software iSCSI to configure software iSCSI for the host.

41. In the Configure iSCSI window, under Dynamic targets, click Add dynamic target.

42. Choose to add address and enter the IP address of iscsi-lif-1a from storage SVM Infra-SVM. Press Return.

43. Repeat steps 1-42 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**

<b>iSCSI enabled</b>	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled															
<b>Name &amp; alias</b>	iqn.2010-11.com.flexpod.ucs-host.4															
<b>CHAP authentication</b>	Do not use CHAP															
<b>Mutual CHAP authentication</b>	Do not use CHAP															
<b>Advanced settings</b>	Click to expand															
<b>Network port bindings</b>	No port bindings															
<b>Static targets</b>	<div style="display: flex; justify-content: space-between; align-items: center;"> <span>  Add static target            Remove static target            Edit settings         </span> <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.9c7aefb072b611eab6b000a098e...</td> <td>192.168.10.51</td> <td>3260</td> </tr> <tr> <td>iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...</td> <td>192.168.20.52</td> <td>3260</td> </tr> <tr> <td>iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...</td> <td>192.168.10.52</td> <td>3260</td> </tr> <tr> <td>iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...</td> <td>192.168.20.51</td> <td>3260</td> </tr> </tbody> </table>	Target	Address	Port	iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.10.51	3260	iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.20.52	3260	iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.10.52	3260	iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.20.51	3260
Target	Address	Port														
iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.10.51	3260														
iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.20.52	3260														
iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.10.52	3260														
iqn.1992-08.com.netapp:sn.9c7aefb072b611eab6b000a098e...	192.168.20.51	3260														
<b>Dynamic targets</b>	<div style="display: flex; justify-content: space-between; align-items: center;"> <span>  Add dynamic target            Remove dynamic target            Edit settings         </span> <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>192.168.10.51</td> <td>3260</td> </tr> <tr> <td>192.168.10.52</td> <td>3260</td> </tr> <tr> <td>192.168.20.51</td> <td>3260</td> </tr> <tr> <td>192.168.20.52</td> <td>3260</td> </tr> </tbody> </table>	Address	Port	192.168.10.51	3260	192.168.10.52	3260	192.168.20.51	3260	192.168.20.52	3260					
Address	Port															
192.168.10.51	3260															
192.168.10.52	3260															
192.168.20.51	3260															
192.168.20.52	3260															

47. Click Cancel to close the window.



**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, choose the ESXi host.

2. In the center pane, click 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. Choose vmnic5.

## Add Physical Adapters to the Switch ×

### Network Adapters

vmnic5

All
Properties
CDP
LLDP

<b>Adapter Name</b>	Cisco Systems Inc Cisco VIC Ethernet NIC vmnic5
<b>Location</b>	PCI 0000:1c:00.5
<b>Driver</b>	nenic
<b>Status</b>	
Status	Connected
Actual speed, Duplex	50 Gbit/s, Full Duplex
Configured speed, Duplex	50 Gbit/s, Full Duplex
Networks	No networks
<b>Network I/O Control</b>	
Status	Allowed
<b>SR-IOV</b>	
Status	Not supported
<b>Cisco Discovery Protocol</b>	
Version	2
Timeout	60
Time to live	163
Samples	4294
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
<b>Adapter Name</b>		Cisco Systems Inc Cisco VIC Ethernet NIC	
<b>Location</b>		vmnic5	
<b>Driver</b>		PCI 0000:1c:00.5	
<b>Driver</b>		nenic	
<b>Status</b>			
<b>Status</b>		Connected	
<b>Actual speed, Duplex</b>		50 Gbit/s, Full Duplex	
<b>Configured speed, Duplex</b>		50 Gbit/s, Full Duplex	
<b>Networks</b>		No networks	
<b>Network I/O Control</b>			
<b>Status</b>		Allowed	
<b>SR-IOV</b>			
<b>Status</b>		Not supported	
<b>Cisco Discovery Protocol</b>			

CANCEL OK

9. Click OK to complete adding the vmnic to the vSwitch.
10. Choose VMkernel adapters.
11. Choose the iScsiBootPG VMkernel adapter line. Choose 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. Choose Virtual switches. In the center pane, expand iScsiBootvSwitch.
16. Choose ... to the right of iScsiBootPG in the center of the window. Choose Edit Settings.
17. In the Edit Settings window, choose Teaming and failover. In the center of the window under Failover order, choose Override.

## iScsiBootPG - Edit Settings

<b>Properties</b>	Load balancing	<input type="checkbox"/> Override	Route based on originating virtual port	▼
<b>Security</b>	Network failure detection	<input type="checkbox"/> Override	Link status only	▼
<b>Traffic shaping</b>	Notify switches	<input type="checkbox"/> Override	Yes	▼
<b>Teaming and failover</b>	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, choose Virtual switches.
20. In the center pane, expand iScsiBootvSwitch. Click EDIT to edit settings for the vSwitch.
21. Choose Teaming and failover. Under Failover order, choose 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 0000:1c:00.5  
Driver: nenic

**Status**

Status: Connected  
Actual speed, Duplex: 50 Gbit/s, Full Duplex  
Configured speed, Duplex: 50 Gbit/s, 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, choose VMkernel adapters.
24. In the center pane, choose Add Networking to add a VMkernel adapter.
25. Make sure VMkernel Network Adapter is selected and click NEXT.
26. Choose an existing standard switch and click BROWSE.
27. Choose iScsiBootvSwitch and click OK.

## nx-esxi-3.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

iScsiBootvSwitch BROWSE ...

New standard switch

MTU (Bytes)

CANCEL BACK NEXT

28. Click NEXT.

29. Enter iScsiBootPG-B for the Network label, leave VLAN ID set to None (0), choose Custom - 9000 for MTU, and click NEXT.

30. Choose 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, choose Virtual switches. Expand iScsiBootvSwitch.

34. In the center of the window, choose ... to the right of iScsiBootPG-B. Choose Edit Settings.

35. Choose Teaming and failover. In the center of the window under Failover order, choose Override.

36. Choose 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
vmnic5

Standby adapters

Unused adapters
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 0000:1c:00.4		
<b>Status</b>			
Status	Connected		
Actual speed, Duplex	50 Gbit/s, Full Duplex		
Configured speed, Duplex	50 Gbit/s, Full Duplex		
Networks	No networks		

**CANCEL** **OK**

37. Click OK to complete the changes to the port group.

38. In the list under Storage, choose Storage Adapters.

39. Choose the iSCSI Software Adapter and below, choose 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 now be listed.

45. If NetApp VSC is installed, 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](mailto:administrator@vsphere.local) and choose 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, choose 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. Choose 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
```

11. Exit Maintenance Mode on each ESXi host in Maintenance Mode.



**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.**

## Create a FlexPod ESXi Custom ISO using VMware vCenter

In this validation document, the [Cisco Custom Image for ESXi 6.7 U3 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 U3 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

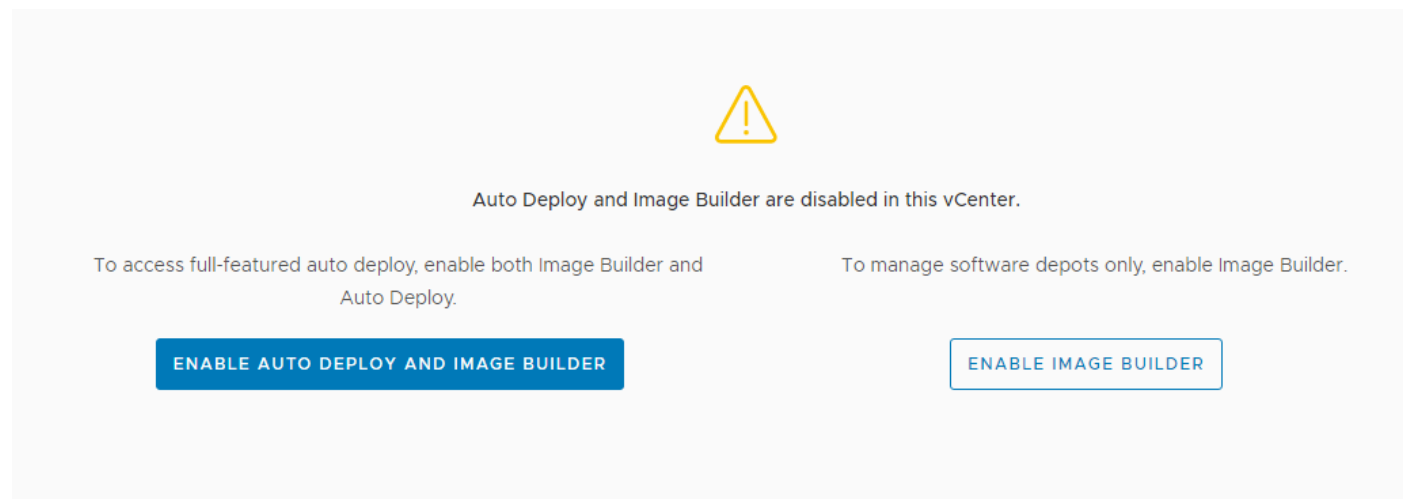
can be used to install VMware ESXi 6.7 U3 without having to do any additional driver updates. This ISO can be produced by following these steps:

1. Download the [Cisco Custom Image for ESXi 6.7 U3 Offline Bundle](#). This file (VMware\_ESXi\_6.7.0\_14320388\_Custom\_Cisco\_6.7.3.1\_Bundle.zip) can be used to produce the Cisco Custom Image for ESXi 6.7 U3 GA Install CD ISO.
2. Download the following driver images and extract the listed .zip file:
  - [VMware ESXi 6.7 nfnic 4.0.0.52 FC Driver for Cisco nfnic](#) - Cisco-nfnic\_4.0.0.52-1OEM.670.0.0.8169922-15920005.zip
  - [VMware ESXi 6.7 nenic 1.0.31.0 NIC Driver for Cisco nenic](#) - VMW-ESX-6.7.0-nenic-1.0.31.0-offline\_bundle-15180549.zip
  - [NetApp NFS Plug-in for VMware VAAI 1.1.2](#) - NetAppNasPlugin.v23.zip



**It is not necessary to extract the NetAppNasPlugin.v23.zip file, that file is directly downloaded.**

3. Log into the VMware vCenter HTML5 Client as administrator@vsphere.local.
4. Under Menu, choose Auto Deploy.
5. If you see the following, choose ENABLE IMAGE BUILDER.



6. In the center pane, choose Software Depots.
7. Click IMPORT to upload a software depot.
8. Name the depot ESXi-6.7U3-Custom-Cisco. Click BROWSE. Browse to the local location of the VMware\_ESXi\_6.7.0\_14320388\_Custom\_Cisco\_6.7.3.1\_Bundle.zip file downloaded above, highlight it, and click Open.

## Import Software Depot



Name \*

File \*

9. Click **UPLOAD** to upload the software depot.
10. Repeat steps 6-9 to add software depots for nfnic-4.0.0.52, nenic-1.0.31.0, and NetAppNasPlugin-v23.
11. Click **NEW** to add a custom software depot.
12. Choose Custom depot and name the custom depot FlexPod-ESXi-6.7U3.

## Add Software Depot



Online depot

Name:

URL:

Custom depot

Name: \*

CANCEL

ADD

13. Click ADD to add the custom software depot.
14. From the drop-down list, choose the ESXi-6.7U3-Custom-Cisco (ZIP) software depot. Make sure the Image Profiles tab is selected and then click the radio button to select the VMware-ESXi-6.7.0-14320388-Custom-Cisco-6.7.3.1 image profile. Click CLONE to clone the image profile.
15. Name the clone FlexPod-ESXi-6.7U3. For Vendor, enter Cisco-NetApp. For Description, enter "Cisco Custom ISO for ESXi 6.7U3 GA with nfnic-4.0.0.52, nenic-1.0.31.0, and NetAppNasPlugin-v23". Choose FlexPod-ESXi-6.7U3 for Software depot.

### Clone Image Profile

- 1 Name and details
- 2 Select software packages
- 3 Ready to complete

### Name and details

Name \* FlexPod-ESXi-6.7U3

Vendor \* Cisco-NetApp

Description

Cisco Custom ISO for ESXi 6.7U3 GA with nfnc-4.0.0.52, nenic-1.0.31.0, and NetAppNasPlugin-v23

Software depot \* FlexPod-ESXi-6.7U3 ⓘ

CANCEL NEXT

16. Click NEXT.

17. Under Available software packages, uncheck `nenic-1.0.29.0`, check `nenic-1.0.31.0`, check `NetAppNasPlugin 1.1.2-3`, check `nfnc 4.0.0.52`, and uncheck `nfnc 4.0.0.40`. 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"/>	misc-drivers	6.7.0-2.48.13006603	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	mtip32xx-native	3.9.8-1vmw.670.1.28.103...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	native-misc-drive...	6.7.0-2.48.13006603	VMware certified	VMware	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	ne1000	0.8.4-2vmw.670.2.48.13...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input type="checkbox"/>	nenic	1.0.29.0-1vmw.670.3.73...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	nenic	1.0.31.0-1OEM.670.0.0....	VMware certified	Cisco	nenic-1.0.31.0
<input checked="" type="checkbox"/>	net-bnx2	2.2.4f.v60.10-2vmw.67...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	net-bnx2x	1.78.80.v60.12-2vmw.67...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	net-cdc-ether	1.0-3vmw.670.0.0.8169...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	net-cnic	1.78.76.v60.13-2vmw.67...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco

148 selected of 150 items

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-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	net-usbnet	1.0-3vmw.670.0.0.8169...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	net-vmxnet3	1.1.3.0-3vmw.670.2.48.1...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	NetAppNasPlugin	1.1.2-3	VMware accepted	NetApp	NetAppNasPlugin-v23
<input checked="" type="checkbox"/>	nfnic	4.0.0.52-1OEM.670.0.0....	VMware certified	Cisco	nfnic-4.0.0.52
<input type="checkbox"/>	nfnic	4.0.0.40-1OEM.670.0.0....	VMware certified	Cisco	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	nhpsa	2.0.22-3vmw.670.1.28.1...	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx4-core	3.17.13.1-1vmw.670.2.48....	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx4-en	3.17.13.1-1vmw.670.2.48....	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco
<input checked="" type="checkbox"/>	nmlx4-rdma	3.17.13.1-1vmw.670.2.48....	VMware certified	VMW	ESXI-6.7U3-Custom-Cisco

148 selected of 150 items

CANCEL BACK NEXT

18. Click NEXT.

19. Click FINISH.
20. Using the Software Depot pulldown, choose the FlexPod-ESXi-6.7U3 (Custom) software depot. Under Image Profiles choose the FlexPod-ESXi-6.7U3 image profile. Click ... > Export to export an image profile. ISO should be selected. Click OK to generate a bootable ESXi installable image.
21. Once the Image profile export completes, click DOWNLOAD to download the ISO.
22. 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 create the 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. Choose Admin within the Navigation pane and choose 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: <choose 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:

```
show 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 : Yet to be executed
-----
      Job Name          Last Execution Status
-----
backup-cfg              -NA-
=====
```

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](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 scheduled backup of the vCenter Server Appliance is available within the native capabilities of the VCSA. To create a scheduled backup, follow these steps:

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

2. Click Backup in the list to open up the Backup Appliance Dialogue.
3. To the right of Backup Schedule, click CONFIGURE.
4. 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 ⓘ	<u>scp://nx-ftp.flexpod.cisco.com/var/www/html/software/Configs/nx-vc/</u>	
Backup server credentials	User name	<u>admin</u>
	Password	<u>*****</u>
Schedule ⓘ	<u>Daily</u> ▼	<u>02</u> : <u>15</u> <u>A.M.</u> <u>America/New_York</u>
Encrypt backup (optional)	Encryption Password	<u></u>
	Confirm Password	<u></u>
Number of backups to retain	<input type="radio"/> Retain all backups	
	<input checked="" type="radio"/> Retain last <u>7</u> <input type="text" value="↑"/> <input type="text" value="↓"/> backups	
Data	<input checked="" type="checkbox"/> Inventory and configuration	1158 MB
	<input checked="" type="checkbox"/> Stats, Events, and Tasks	131 MB
	Total size (compressed) 1289 MB	

CANCEL
CREATE

5. 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/nx-vc/
Backup data	<ul style="list-style-type: none"><li>• Inventory and configuration</li><li>• Stats, Events, and Tasks</li></ul>
Number of backups to retain	7

Activity [BACKUP NOW](#)

6. The Backup Schedule should now show a Status of Enabled.
7. 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.

Scott Kovacs, Technical Marketing Engineer, Converged Infrastructure Engineering, NetApp, Inc.

Scott is a Technical Marketing Engineer on the Converged Infrastructure Engineering team at NetApp. He has been with NetApp since 2007, serving in a variety of Technical Support, Professional Services and Engineering roles. Scott has over 22 years of experience in the IT industry specializing in data management, Fibre Channel networking, 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.
- Sree Lakshmi Lanka, Solutions Engineer, NetApp, Inc.