

FlexPod Datacenter with VMware vSphere 7.0, Cisco UCS C125 M5, and NetApp ONTAP 9.7

Deployment Guide for FlexPod Datacenter with VMware vSphere 7.0, Cisco UCS C125 M5 Servers with 2nd Generation AMD EPYC CPUs, and NetApp ONTAP 9.7

Published: October 2020



In partnership with:



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 Inter-network 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, Giga-Drive, 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. 2021.

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)

© 2021 Cisco Systems, Inc. All rights reserved.

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 continuing to use 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 A400 storage, Cisco UCS Manager unified software release 4.1(2) with 2nd Generation AMD EPYC Processors and VMware vSphere 7.0. Cisco UCS Manager (UCSM) 4.1(2) 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(2), and VMware vSphere 7.0 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 primary FlexPod Datacenter with VMware vSphere 7.0 validated design introduced new hardware and software into the portfolio, enabling 10/25/40/100GbE along with native 32Gb FC via the Cisco MDS Fibre Channel switch or the Cisco Nexus 93180YC-FX switch. This primary design has been updated to include the latest Cisco and NetApp hardware and software. New pieces include:

- Support for the Cisco UCS 4.1(2) unified software release, Cisco UCS C125 servers with AMD EPYC 2nd Generation Processors, Cisco UCS B200-M5 and C220-M5 servers with 2nd Generation Intel Xeon Scalable Processors, and Cisco 1400 Series Virtual Interface Cards (VICs)
- Support for the latest Cisco UCS 6454 and 64108 (supported but not validated) Fabric Interconnects
- Support for the latest Cisco UCS 2408 Fabric Extender
- Support for Intel Optane Persistent Memory in Memory Mode with specific memory configurations and App Direct Mode
- Support for 32Gb FC SAN Switching in the Cisco Nexus 93180YC-FX switch
- Support for Cisco Data Center Network Manager (DCNM)-SAN Version 11.4(1)
- Support for Cisco Intersight Software as a Service (SaaS) Management
- Support for the NetApp AFF A400 and AFF A800 (supported but not validated) Storage Controller
- Support for the latest release of NetApp ONTAP® 9.7
- Support for NetApp Virtual Storage Console (VSC) 9.7P2
- Support for VVOL Datastores/VM over FC, iSCSI and NFS protocol.
- Support for NetApp SnapCenter and NetApp SnapCenter Plug-in for VMware vSphere Version 4.3.1

-
- Support for NetApp Active IQ Unified Manager 9.7P1
 - Support for NetApp Active IQ
 - Fibre channel, NFS, iSCSI (appendix) storage design
 - Validation of VMware vSphere 7.0
 - Unified Extensible Firmware Interface (UEFI) Secure Boot of VMware ESXi 7.0
 - 25 or 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.

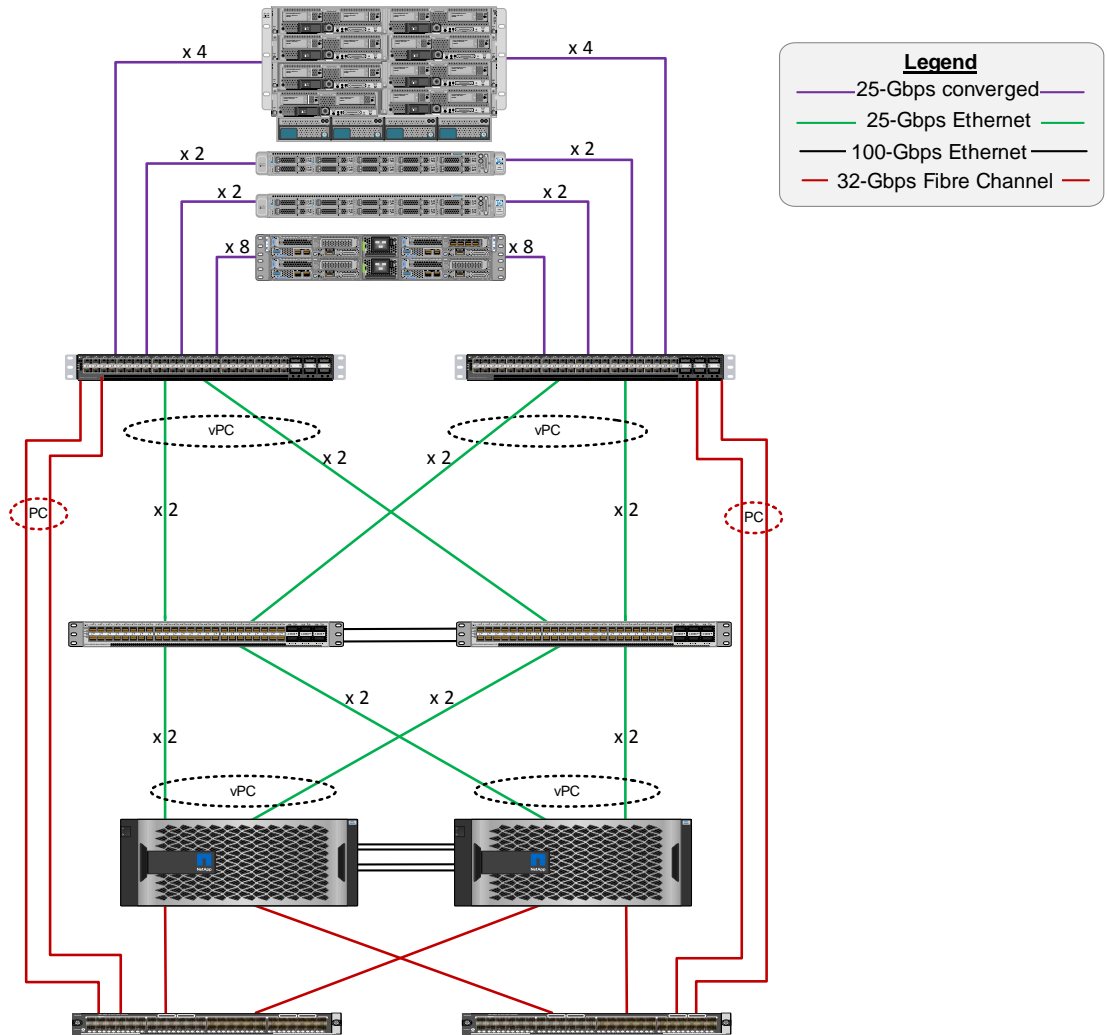
[Figure 1](#) shows the VMware vSphere built on FlexPod components and the network connections for a configuration with the Cisco UCS 6454 Fabric Interconnects. This design has port-channelled 25 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects via the Cisco UCS 2408 Fabric Extenders, port-channelled 25 Gb Ethernet connections between the C-Series rackmounts and the Cisco UCS Fabric Interconnects, and port-channel 25 Gb Ethernet connections between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A400 storage array. This infrastructure option expanded with Cisco MDS switches sitting between the Cisco UCS Fabric Interconnect and the NetApp AFF A400 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 A Series

Cisco Unified Computing System

Cisco UCS 6454 Fabric Interconnects, UCS 2408 Fabric Extenders, UCS B-Series Blade Servers with UCS VIC 1440, UCS C-Series Rack Servers with UCS VIC 1457, UCS C4200 Chassis, and UCS C125 Servers with UCS VIC 1455



Cisco Nexus 93180YC-FX

NetApp storage controllers AFF-A400 or AFF-A800

Cisco MDS 9132T or 9148T switch



Although this diagram includes the Cisco UCS 5108 chassis with Cisco UCS B-Series servers and Cisco UCS C-220 servers, this document describes the configuration of only the Cisco UCS C125 servers in the Cisco UCS C4200 chassis with AMD EPYC CPUs. For the configuration of Cisco UCS Intel-based servers, please see section [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7](#).

The reference 100Gb based hardware configuration includes:

- Two Cisco Nexus 93180YC-FX switches
- Two Cisco UCS 6454 fabric interconnects
- Two Cisco MDS 9132T multilayer fabric switches
- One NetApp AFF A400 or A800 (HA pair) running ONTAP 9.7 with NVMe SSD disks

Software Revisions

[Table 1](#) lists the software revisions for this solution.

Table 1. Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6454, Cisco UCS C125 Servers with 2nd Generation AMD EPYC Processors	4.1(2b)	Includes the Cisco UCS Manager and Cisco UCS VIC 1455
Network	Cisco Nexus 93180YC-FX NX-OS	9.3(5)	
	Cisco MDS 9132T	8.4(1a)	
Storage	NetApp AFF 400	ONTAP 9.7	
Software	Cisco UCS Manager	4.1(2b)	
	Cisco Data Center Network Manager (SAN)	11.4(1)	
	VMware vSphere	7.0	
	VMware ESXi nfnic FC Driver	4.0.0.56	
	VMware ESXi nenic Ethernet Driver	1.0.33.0	
	NetApp Virtual Storage Console (VSC) / VASA Provider Appliance	9.7.1	
	NetApp NFS Plug-in for VMware VAAI	1.1.2	
	NetApp SnapCenter for vSphere	4.3.1	Includes the vSphere plug-in for SnapCenter
	NetApp Active IQ Unified Manager	9.7P1	
Management	Cisco Intersight	N/A	
Compute	NetApp Active IQ	N/A	Includes the Cisco UCS Manager and Cisco UCS VIC 1455

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](#) lists the VLANs necessary for deployment as outlined in this guide.

Table 2. Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating this Document
Out-of-Band Mgmt	VLAN for out-of-band management interfaces	13
In-Band Mgmt	VLAN for in-band management interfaces	113
Native	VLAN to which untagged frames are assigned	2
Infra-NFS	VLAN for Infrastructure NFS traffic	3050
FCoE-A	VLAN for FCoE encapsulation of VSAN-A	103
FCoE-B	VLAN for FCoE encapsulation of VSAN-B	104
vMotion	VLAN for VMware vMotion	3000
VM-Traffic	VLAN for Production VM Interfaces	900

[Table 3](#) lists the VMs necessary for deployment as outlined in this document.

Table 3. Virtual Machines

Virtual Machine Description	Host Name	IP Address
vCenter Server		
NetApp VSC		
NetApp SnapCenter for vSphere		
Active IQ Unified Manager		
Cisco Data Center Network Manager (DCNM)		
Cisco Intersight Assist		

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.



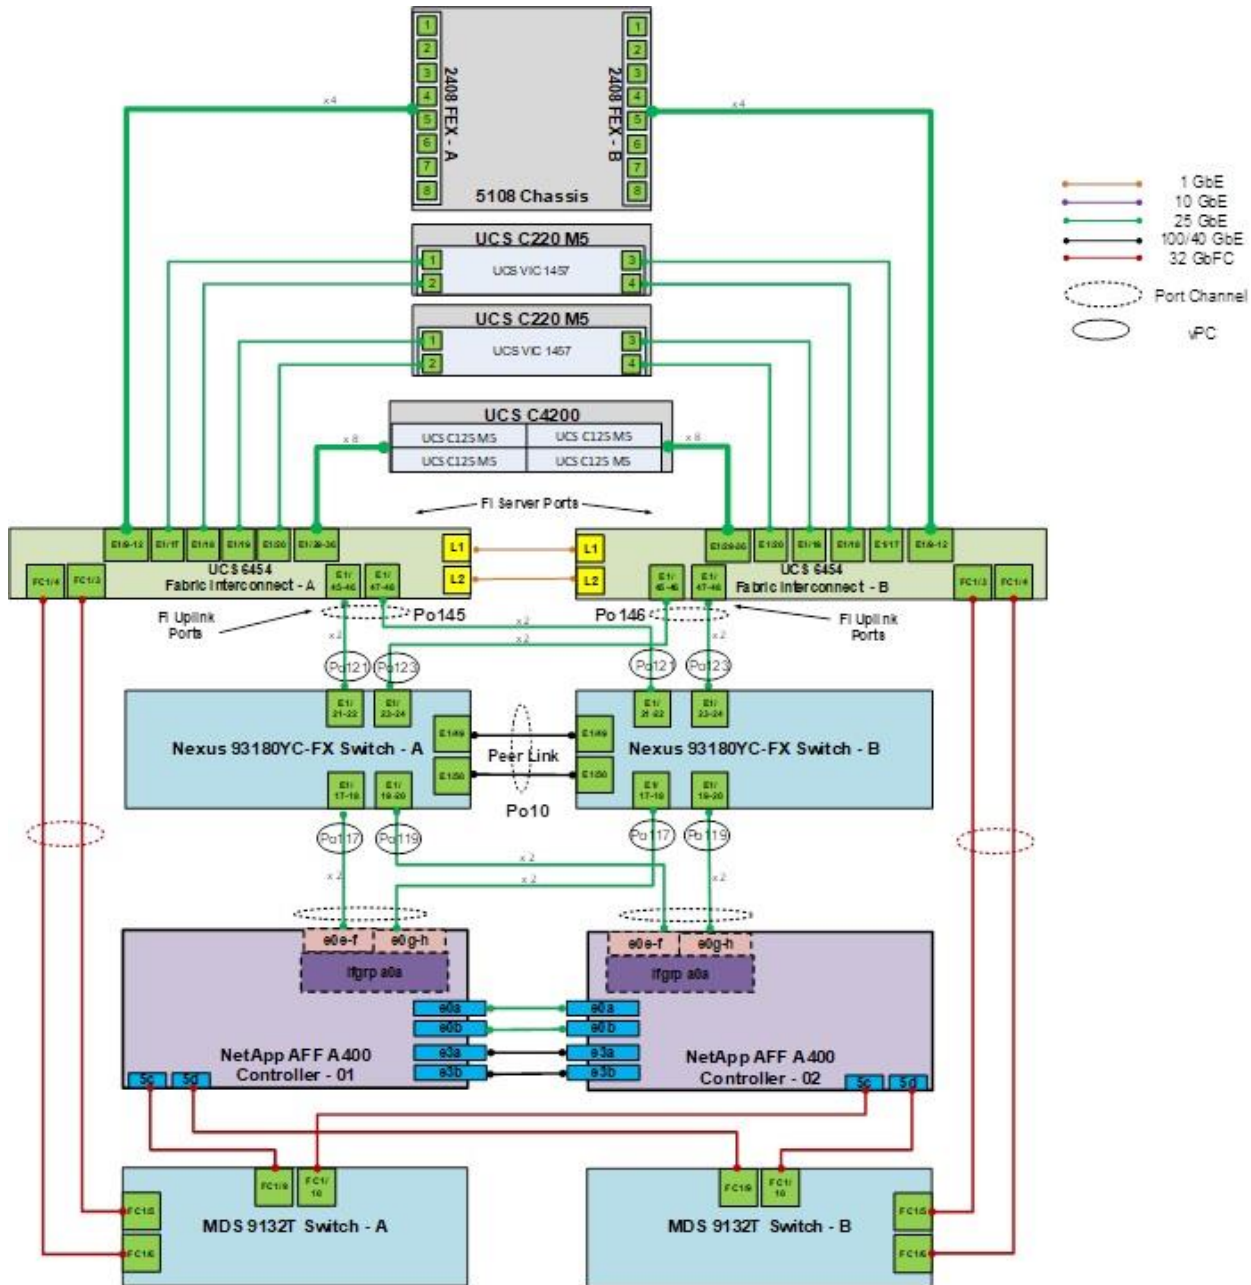
Be sure to use the cabling directions in this section as a guide.

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, 25Gb 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



Although this diagram includes the Cisco UCS 5108 chassis with Cisco UCS B-Series servers and Cisco UCS C-220 servers, this document describes configuration of only Cisco UCS C125 servers in the Cisco UCS C4200 chassis with AMD EPYC CPUs. For configuration of Cisco UCS Intel-based servers, please see the [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7](#) CVD.

Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 93180YC-FX switches for use in a FlexPod environment. The Nexus 93180YC-FX will be used for LAN switching in this solution and can also be used for SAN switching. 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 93180YC-FX switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 9.3(5), the Cisco suggested Nexus switch release at the time of this validation.



If using the Cisco Nexus 93180YC-FX switches for both LAN and SAN switching, please refer to section [FlexPod with Nexus 93180YC-FX SAN Switching Configuration - Part 1](#) in the Appendix.



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.

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]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: n
Configure CoPP system profile (strict/moderate/lenient/dense) [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]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense) [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 Features

Cisco Nexus A and Cisco Nexus B

SAN switching requires both the SAN_ENTERPRISE_PKG and FC_PORT_ACTIVATION_PKG licenses. Please ensure these licenses are installed on each Nexus 93180YC-FX switch. To enable the appropriate features on the Cisco Nexus switches, follow these steps:

1. Log in as admin.
2. Because basic FC configurations were entered in the setup script, feature-set fcoe has been automatically installed and enabled. Run the following commands:

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

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

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

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day> <end-month> <end-time> <offset-minutes>
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```



It is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summer time, please see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3\(x\)](#). Sample clock commands for the United States Eastern timezone are:

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

Create VLANs

Cisco Nexus A and Cisco Nexus B

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

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

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

Add NTP Distribution Interface

Cisco Nexus A

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

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

Cisco Nexus B

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

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

Add 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
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
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 Eth1/21
description <ucs-clustername>-a:1/45
udld enable
interface Eth1/22
description <ucs-clustername>-a:1/46
udld enable
interface Eth1/23
description <ucs-clustername>-b:1/45
udld enable
interface Eth1/24
description <ucs-clustername>-b:1/46
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. If you have fibre optic connections, do not enter the `udld enable` command.

```
interface Eth1/17
description <st-clustername>-01:e0e
interface Eth1/18
description <st-clustername>-01:e0f
interface Eth1/19
description <st-clustername>-02:e0e
interface Eth1/20
description <st-clustername>-02:e0f
interface Eth1/49
description <nexus-b-hostname>:1/49
interface Eth1/50
description <nexus-b-hostname>:1/50
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 Eth1/21
description <ucs-clustername>-a:1/47
udld enable
interface Eth1/22
description <ucs-clustername>-a:1/48
udld enable
interface Eth1/23
description <ucs-clustername>-b:1/47
udld enable
interface Eth1/24
description <ucs-clustername>-b:1/48
udld enable
```



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

```
interface Eth1/17
description <st-clustername>-01:e0g
interface Eth1/18
description <st-clustername>-01:e0h
interface Eth1/19
description <st-clustername>-02:e0g
interface Eth1/20
description <st-clustername>-02:e0h
interface Eth1/49
description <nexus-a-hostname>:1/49
interface Eth1/50
description <nexus-a-hostname>:1/50
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 Eth1/49-50
channel-group 10 mode active
no shutdown
interface Po117
description <st-clustername>-01
interface Eth1/17-18
channel-group 117 mode active
no shutdown
interface Po119
description <st-clustername>-02
interface Eth1/19-20
channel-group 119 mode active
no shutdown
interface Po121
description <ucs-clustername>-a
interface Eth1/21-22
channel-group 121 mode active
no shutdown
interface Po123
description <ucs-clustername>-b
interface Eth1/23-24
channel-group 123 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 Po117
inherit port-profile FP-ONTAP-Storage
interface Po119
inherit port-profile FP-ONTAP-Storage

interface Po121
inherit port-profile FP-UCS
interface Po123
inherit port-profile FP-UCS

exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

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

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

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po121
vpc 121
interface Po123
vpc 123
exit
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, follow this step:

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

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po121
vpc 121
interface Po123
vpc 123
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
show int status
```

Storage Configuration

NetApp All Flash FAS A400 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/index.jsp?topic=%2Fcom.netapp.nav.a400%2Fhome.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 16.

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 15.

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>
```

Enter the URL where the software can be found.



The web server must be pingable from node 02.

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

11. Press `Enter` for the user name, indicating no user name.

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

13. 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.

14. Press `Ctrl-C` when you see this message:

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

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

16. Enter `y` to zero disks, reset config, and install a new file system.

17. 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

To set up a node, follow these steps:

1. From a console port program attached to the storage controller A (node 01) console port, run the node set-up script. This script appears when ONTAP 9.7 boots on the node for the first time.
2. 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:
```

3. 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 Detail	Cluster Detail Value
Cluster management gateway	<clustermgmt-gateway>
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>
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>

Cluster Detail	Cluster Detail Value
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.

4. Complete the required information on the Initialize Storage System screen:

5. In the Cluster screen, follow these steps:

1. Enter the cluster name and administrator password.
2. Complete the Networking information for the cluster and each node.
3. Choose the box Use time services (NTP) and enter the IP addresses of the time servers in a comma separated list.

ONTAP System Manager [\(Return to classic version\)](#)

ONTAP 9.7 *Pro tips for initializing a storage system*

Health

2 healthy nodes were found.

AFF-A400

Initialize Storage System

STORAGE SYSTEM NAME

aa11-a400

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
722017000240	192.168.156.61
722017000239	192.168.156.62

Use Domain Name Service (DNS)



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.

6. Click Submit.

ONTAP System Manager (Return to classic version) ? <>

Networking

CLUSTER IP ADDRESS	SUBNET MASK	GATEWAY
192.168.156.60	24	192.168.156.1

NODE SERIAL NUMBERS	NODE IP ADDRESSES
722017000240	192.168.156.61
722017000239	192.168.156.62

Use Domain Name Service (DNS)

Others

Use time services (NTP)

NTP SERVERS
192.168.156.11
192.168.156.12
10.1.156.4
10.1.156.5

[+ Add](#)

[Submit](#)

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

The screenshot shows the ONTAP System Manager interface. The left sidebar has a menu with categories: DASHBOARD, STORAGE, NETWORK, EVENTS & JOBS, PROTECTION, and CLUSTER. Under CLUSTER, 'Overview' is selected and highlighted with a red box. The main content area is titled 'Overview' and shows the following details:

- NAME:** aa11-a400
- VERSION:** NetApp Release 9.7P5: Tue Jun 23 04:02:30 UTC 2020
- NTP SERVERS:** 10.1.156.4, 10.1.156.5, 192.168.156.11, 192.168.156.12
- DNS DOMAINS:** flexpod.cisco.com
- NAME SERVERS:** 10.1.156.250, 10.1.156.251
- MANAGEMENT INTERFACES:** 192.168.156.60
- DATE AND TIME:** July 20, 2020, 2:51 AM Etc/UTC

On the right side of the Overview card, there is a menu with options: ONTAP Update, History, Rename, and Edit. The 'Edit' option is highlighted with a red box. A red arrow points to the menu icon above these options.

Below the Overview card is a 'Nodes' section with a table:

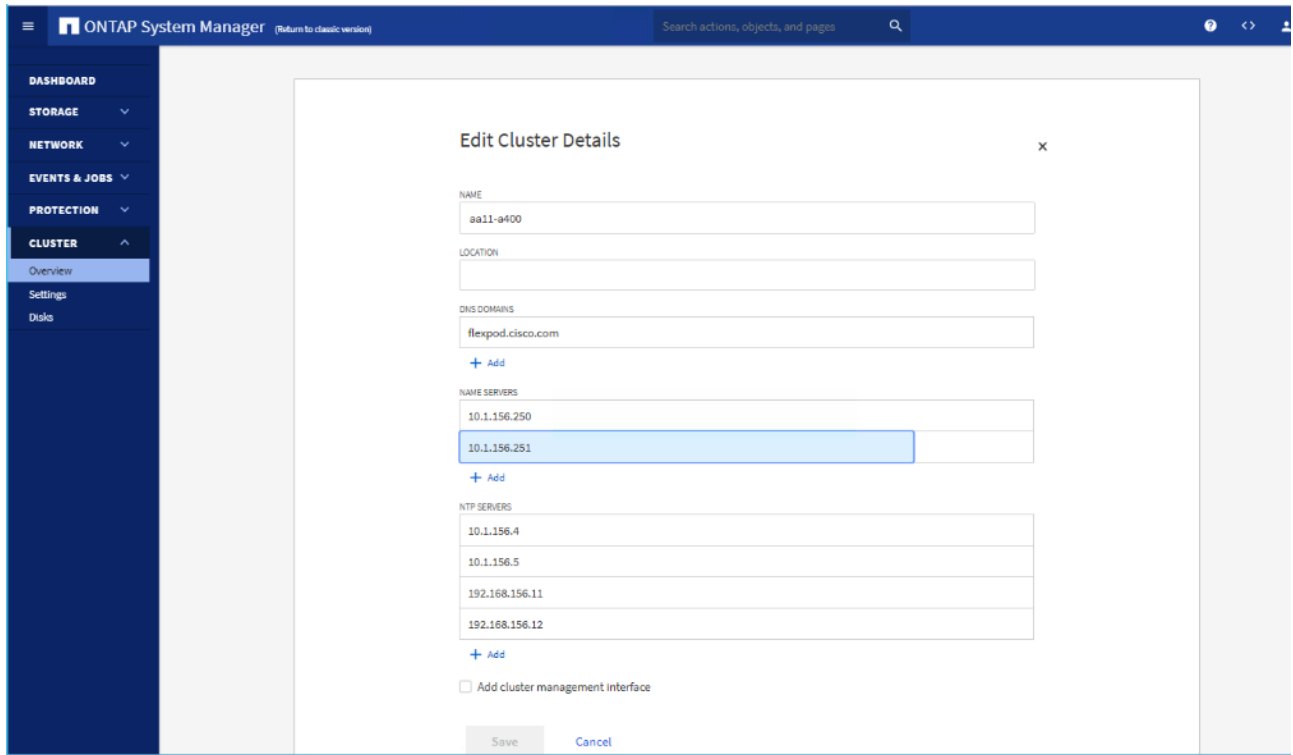
Nodes	Name	Up Time	Serial Num...	Management IP	Service Processor IP	System ID
aa11-a400-01 / aa11-a400-02						
	aa11-a400-01	00:18:12	722017000240	192.168.156.61	192.168.156.58 ::	0538141868
	aa11-a400-02	00:25:48	722017000239	192.168.156.62	192.168.156.59 ::	0538143135

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

1. Cluster location
2. DNS domain name
3. DNS server IP addresses



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



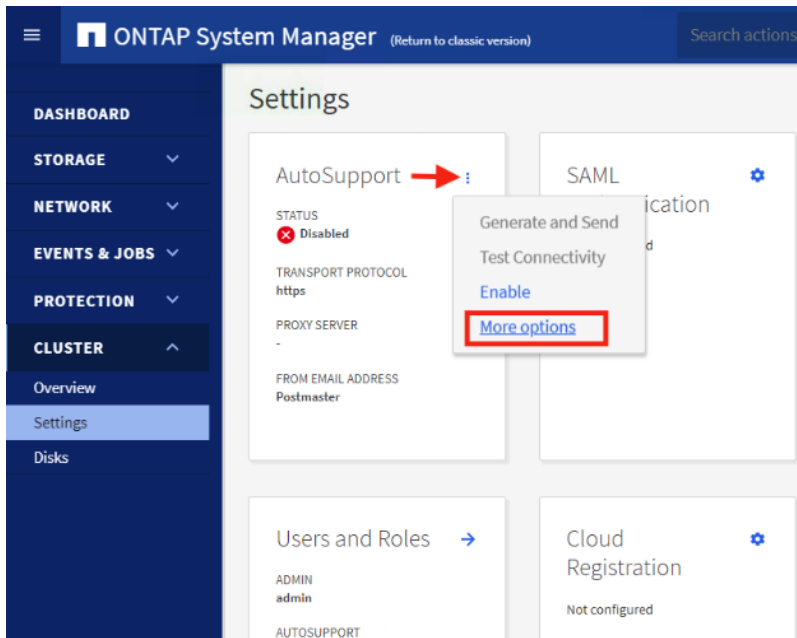
11. Click Save to make the changes persistent.



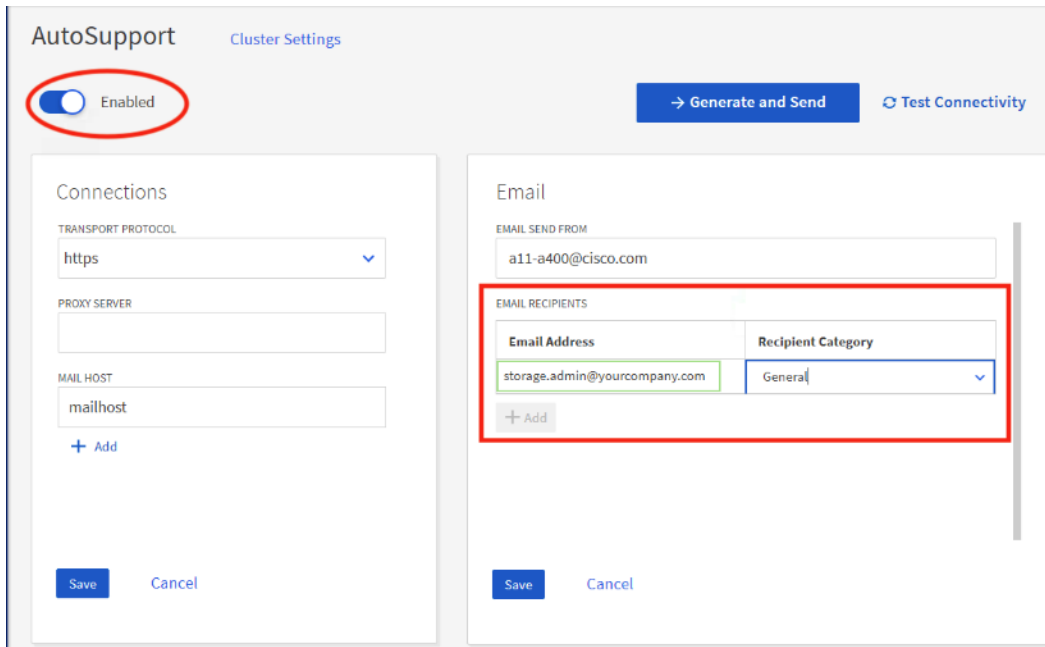
To configure AutoSupport, add licenses and create storage aggregates via the ONTAP CLI skip this section and configure the options in section.

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

13. Click Settings under the Cluster list.



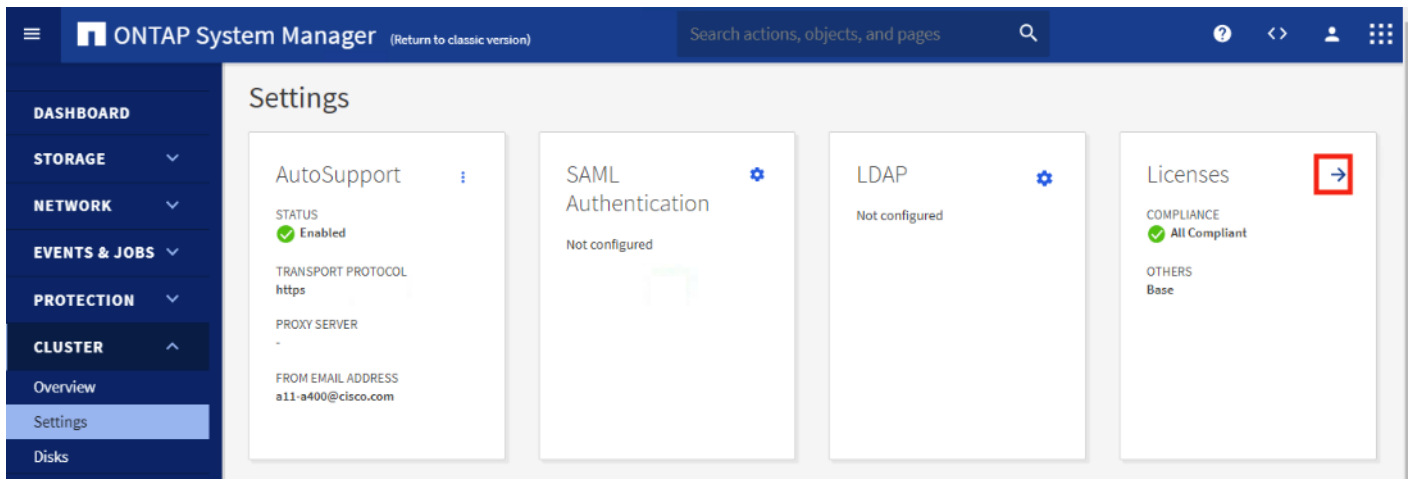
14. If AutoSupport was not configured during the initial setup, click the ellipsis in the AutoSupport tile and choose more options.
15. To enable AutoSupport click the slider button.
16. Click Edit to change the transport protocol, add a proxy server address and a mail host as needed.
17. Click Save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




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

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

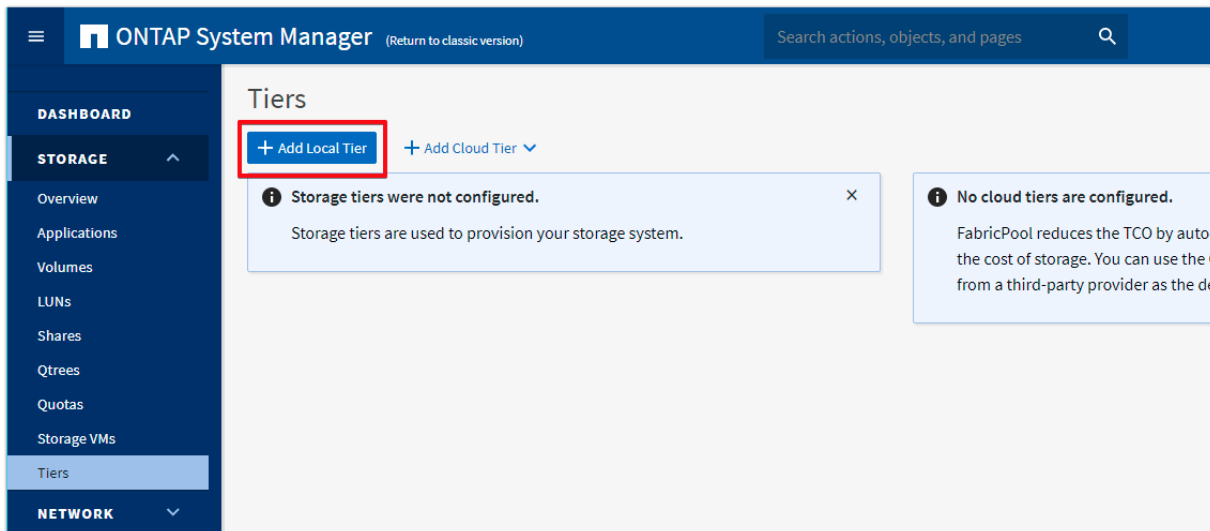
21. Add the desired licenses to the cluster by clicking Add and entering the license keys in a comma separated list.





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

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



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

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

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

27. Click Save to make the configuration persistent.

Add Local Tier x

Storage Recommendation

32.57 TB
USABLE

2 local tiers can be added on nodes "aa11-a400-01", "aa11-a400-02"

^ Recommendation details

LOCAL TIER DETAILS

Node Name	Local Tier	Usable Size	Type
aa11-a400-01	aa11_a400_01_NVME_...	16.29 TB	SSD
aa11-a400-02	aa11_a400_02_NVME_...	16.29 TB	SSD

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 69 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 72 if high availability is configured.
5. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

6. Verify that hardware assist is correctly configured.

```
storage failover hwassist show
```

Set Auto-Revert on Cluster Management

To set the auto-revert, follow these steps:

1. 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.

2. 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:

```
aa11-a400::*> 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-a400-01	1	16.29TB	0	24	0	1
aa11-a400-02	1	16.29TB	0	24	0	1
Total:	2	32.57TB	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-a400-01	aa11_a400_01_NVME_SSD_1	16.29TB	0	23
aa11-a400-02	aa11_a400_02_NVME_SSD_1	16.29TB	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_a400_01_NVME_SSD_1/plex0/rg0		NVMe-SSD	882.4GB	partition	21	2
/aa11_a400_02_NVME_SSD_1/plex0/rg0		NVMe-SSD	882.4GB	partition	21	2

```
Details about spare disks and partitions remaining after aggregate creation:
```

Node	Disk Type	Device Usable Size	Disk Or Partition	Remaining Spares
aa11-a400-01	NVMe-SSD	882.4GB	partition	1
aa11-a400-02	NVMe-SSD	882.4GB	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, `e0e`, `e0f`, 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 port broadcast-domain remove-ports -broadcast-domain Default -port <st-node01>:e0e,<st-
node02>:e0e,<st-node01>:e0f,<st-node02>:e0f,<st-node01>:e0g,<st-node02>:e0g,<st-node01>:e0h,<st-
node02>:e0h,<st-node01>:e4a,<st-node02>:e4a,<st-node01>:e4b,<st-node02>:e4b
```

```
network port broadcast-domain show
```

Disable Flow Control on 25/100GbE Ports

To disable flow control on 25 and 100GbE ports, follow these steps:

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

```
network port modify -node <st-node01> -port e4a,e4b -flowcontrol-admin none
modify -node <st-node01> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
```

network port

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

```
network port modify -node <st-node02> -port e4a,e4b -flowcontrol-admin none

network port modify -node <st-node02> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
aall-a400::*> net port show -node * -port e0e,e0f,e0g,e0h -fields speed-admin,duplex-admin,flowcontrol-admin
(network port show)
node          port duplex-admin speed-admin flowcontrol-admin
-----
aall-a400-01 e0e  full          25000     none
aall-a400-01 e0f  full          25000     none
aall-a400-01 e0g  full          25000     none
aall-a400-01 e0h  full          25000     none
aall-a400-02 e0e  full          25000     none
aall-a400-02 e0f  full          25000     none
aall-a400-02 e0g  full          25000     none
aall-a400-02 e0h  full          25000     none
8 entries were displayed.
aall-a400::*> net port show -node * -port e4a,e4b -fields speed-admin,duplex-admin,flowcontrol-admin
(network port show)
node          port duplex-admin speed-admin flowcontrol-admin
-----
aall-a400-01 e4a  full          100000    none
aall-a400-01 e4b  full          100000    none
aall-a400-02 e4a  full          100000    none
aall-a400-02 e4b  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. To disable auto-negotiate, follow these steps:

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

```
fcv adapter modify -node <st-node01> -adapter 5c -status-admin down
fcv adapter modify -node <st-node01> -adapter 5d -status-admin down
fcv adapter modify -node <st-node02> -adapter 5c -status-admin down
fcv adapter modify -node <st-node02> -adapter 5d -status-admin down
```

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

```
fcv adapter modify -node <st-node01> -adapter 5c -speed 32 -status-admin up
fcv adapter modify -node <st-node01> -adapter 5d -speed 32 -status-admin up
fcv adapter modify -node <st-node02> -adapter 5c -speed 32 -status-admin up
fcv adapter modify -node <st-node02> -adapter 5d -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 25GbE 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 e0e
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0f
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0h
network port ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0e
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0f
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0h

network port ifgrp show
```

Change MTU on Interface Groups

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

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

Create VLANs

To create VLANs, follow these steps:

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

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



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

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

network port vlan show
```

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

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

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

Configure Network Time Protocol

To configure time synchronization on the cluster, follow these steps:

1. Set the time zone for the cluster.

```
timezone -timezone <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 trapshot 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.

4. Generate and install self-signed certificates by running 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
```

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

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

7. 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.

8. 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:

1. 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 32GB -ostype vmware -space-reserve disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-02 -size 32GB -ostype vmware -space-reserve disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-03 -size 32GB -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-01a -role data -data-protocol fcp -home-node <st-node01> -home-port 5c -status-admin up

network interface create -vserver Infra-SVM -lif fcp-lif-01b -role data -data-protocol fcp -home-node <st-node01> -home-port 5d -status-admin up

network interface create -vserver Infra-SVM -lif fcp-lif-02a -role data -data-protocol fcp -home-node <st-node02> -home-port 5c -status-admin up
```

```
network interface create -vserver Infra-SVM -lif fcp-lif-02b -role data -data-protocol fcp -home-node <st-
node02> -home-port 5d -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-01 -role data -data-protocol nfs -home-node <st-
node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs-lif-01-ip> -netmask <node01-nfs-lif-01-mask>
-status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM -lif nfs-lif-02 -role data -data-protocol nfs -home-node <st-
node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs-lif-02-ip> -netmask <node02-nfs-lif-02-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, follow these steps:

1. run the following command:

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

2. 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 in ucsd managed mode, 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 management mode. (ucsm/intersight)? ucsm
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
You have chosen to setup a new Fabric interconnect in "ucsm" managed mode. 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(2b)

This document assumes the use of Cisco UCS 4.1(2b). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.1(2b), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Cisco Intersight can also be used to upgrade the Cisco UCS Infrastructure (Cisco UCS Manager, Cisco UCS Fabric Interconnects, and Cisco UCS Fabric Extenders) to version 4.1(2b). Before the upgrade can be done from Cisco Intersight, the UCS cluster will need to be claimed into Intersight. Please see the Cisco Intersight section in the FlexPod Management Tools section of this document. For the Cisco Intersight-based upgrade procedure, please see https://intersight.com/help/features#firmware_upgrade. This upgrade does require interacting with Cisco UCS Manager to reboot the Primary Fabric Interconnect when upgrading. Because the Cisco UCS servers are not yet connected to the Cisco UCS Infrastructure, the servers will not be upgraded using Cisco Intersight. However, the Cisco UCS B and C-Series 4.1(2b) bundles need to be manually downloaded to the Cisco UCS system.

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 :

OK

Cancel



We used the Cisco 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.

General

Events

Actions

Add NTP Server

Properties

Time Zone : America/New_York (Eastern ▼)

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) :

OK Cancel

-
6. Click OK and then click OK again.
 7. Repeat steps 1-6 for 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

OK

Cancel

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

Configure Unified Ports (FCP)

Fibre Channel port configurations differ between the Cisco UCS 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 Cisco UCS 6454 are from the first 16 ports starting from the first port and configured in increments of 4 ports from the left. For the Cisco UCS 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 Cisco UCS 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 re-

mainder of this section shows configuration of the Cisco UCS 6454. Modify as necessary for the Cisco UCS 6332-16UP or Cisco UCS 6248UP.

To enable the fibre channel ports, follow these steps for the Cisco UCS 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	

OK Cancel

- Click OK, then click Yes, then click OK to continue.
- Choose Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
- Choose Configure Unified Ports.
- 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.
- 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.
- Click OK, then click Yes, then OK to continue.
- 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 to 2 Link, the minimum recommended number of links for a FlexPod.

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

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

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : None 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. If you plan to attach B-Series servers to this UCS Domain, enable this policy. 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 click 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 any Cisco UCS chassis and Cisco UCS C125 servers are configured as Server ports and have a status of Up.
5. If any Cisco UCS C125 rack mount ports are missing, choose the ports that are connected to Cisco FEXes and direct connect Cisco UCS C125 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 any chassis, Cisco UCS C125 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 Cisco UCS chassis and rack mounts are configured as Server ports and have a status of Up.
13. If any Cisco UCS C125 rack mount ports are missing, choose the ports that are connected to Cisco FEXes and direct connect Cisco UCS C125 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 any chassis, Cisco UCS C125 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 any 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.

Acknowledge Chassis



Are you sure you want to acknowledge Chassis 1 ?

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

Yes

No

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

Create an Organization

To this point in the Cisco UCS deployment, all items have been deployed at the root level in Cisco UCS Manager. To allow this Cisco UCS to be shared among different projects, Cisco 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 :

OK

Cancel

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 Cisco 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

Create WWNN Pool

Name : WWNN-Pool

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

8. Click Next.

9. Click Add.

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



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



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

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

1 Define Name and Description

2 Add WWN Blocks

Create WWPN Pool

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



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 global VSAN that maps to the same VSAN ID in all available fabrics.

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

- Click OK and then click OK again.
- Under SAN Cloud, right-click VSANs.
- Choose Create VSAN.
- Enter VSAN-B as the name of the VSAN to be used for Fabric B.
- Leave FC Zoning set at Disabled.
- Choose Fabric B.
- 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.
- 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 or Cisco Nexus 93180YC-FX 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 9132T A and use >> to add them to the port channel.

Create FC Port Channel

? ✕

1 Set FC Port Channel Name

2 Add Ports

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

Ports		
Port	Slot ID	WWPN
1	1	20:01:00:3A...
2	1	20:02:00:3A...

>>
<<

Ports in the port channel		
Port	Slot ID	WWPN
3	1	20:03:00:3A...
4	1	20:04:00:3A...

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.

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

13. Click Save Changes to assign the VSAN.
14. Click OK.
15. On the left under FC Port Channels, expand the newly created FC Port-Channel. Under the port-channel choose the first FC Interface. Enter a User Label to indicate the connectivity on the Nexus 93180YC-FX switch, such as <nexus-A-hostname>:fc1/5. Click Save Changes and OK. Repeat this process for the other FC Interface.
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 9132T 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 the newly created FC Port-Channel. Under the FC Port-Channel choose the first FC Interface. Enter a User Label to indicate the connectivity on the Nexus 93180YC-FX switch, such as <nexus-B-hostname>:fc1/5. Click Save Changes and OK. Repeat this process for the other FC Interface.

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

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.
4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information. Optionally, enter the Primary and Secondary DNS server addresses.

Create Block of IPv4 Addresses



From :	<input type="text" value="192.168.156.240"/>	Size :	<input type="text" value="12"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="192.168.156.1"/>
Primary DNS :	<input type="text" value="10.1.156.250"/>	Secondary DNS :	<input type="text" value="10.1.156.251"/>

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

Create Uplink Port Channels to Cisco Nexus Switches

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

1. In Cisco UCS Manager, click LAN.



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

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels under Fabric A.
4. Choose Create Port Channel.
5. Enter 145 as the unique ID of the port channel.

6. Enter Po145-Nexus as the name of the port channel.
7. Click Next.
8. 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
1	0	53	00:3A:9...
1	0	54	00:3A:9...

>>

<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
1	0	45	00:3A:9...
1	0	46	00:3A:9...
1	0	47	00:3A:9...
1	0	48	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 145. Ensure Auto is selected for the Admin Speed. After a few minutes, verify that the Overall Status is Up, and the Operational Speed is correct.

General	Ports	Faults	Events	Statistics
Status Overall Status : ↑ Up Additional Info : none		Properties ID : 145 Fabric ID : A Port Type : Aggregation Transport Type : Ether		
Actions Enable Port Channel Disable Port Channel Add Ports		Name : <input type="text" value="Po145-Nexus"/> Description : <input type="text"/> Flow Control Policy : <input type="text" value="default"/> LACP Policy : <input type="text" value="default"/> Note: Changing LACP policy may flap the port-channel if the suspend-individual value changes! Admin Speed : <input type="radio"/> 1 Gbps <input type="radio"/> 10 Gbps <input type="radio"/> 40 Gbps <input type="radio"/> 25 Gbps <input type="radio"/> 100 Gbps <input checked="" type="radio"/> Auto Operational Speed(Gbps) : 100		

13. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.
14. Right-click Port Channels under Fabric B.
15. Choose Create Port Channel.
16. Enter 146 as the unique ID of the port channel.
17. Enter Po146-Nexus as the name of the port channel.
18. Click Next.
19. Choose the ports connected to the Nexus switches to be added to the port channel:
20. Click >> to add the ports to the port channel.
21. Click Finish to create the port channel.
22. Click OK.
23. In the navigation pane, under LAN > LAN Cloud > Fabric B > Port Channels, choose Port-Channel 146. Ensure Auto is selected for the Admin Speed. After a few minutes, verify that the Overall Status is Up, and the Operational Speed is correct.
24. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 145. Under Port-Channel 145, choose Eth Interface 1/45. In the center pane under Properties, enter a User Label to indicate the port connectivity, such as <nexus-a-hostname>:Eth1/21. Click Save Changes and OK.

25. Repeat steps 1-25 for the remaining seven 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 :

- Click OK, then click OK again to complete creating the profile.
- In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 145. Choose the first Eth Interface under Port-Channel 145. 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 145 and for each Eth Interface under Port-Channel 146 on Fabric B.

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel ... / Eth Interface 1...

General | Faults | Events

Actions	Properties
Delete	ID : 45
Enable Interface	Slot ID : 1
Disable Interface	Fabric ID : A
	Transport Type : Ether
	Port : sys/switch-A/slot-1/switch-ether/port-45
	Membership : Up
	Link Profile : <input type="text" value="UDLD-Normal"/>
	User Label : <input type="text" value="aa11-93180-a:Eth1/21"/>

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 Cisco 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 Cisco UCS Manager version 4.1, the MTU for the Best Effort QoS System Class is again settable. To configure jumbo frames in the Cisco 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

General Events FSM

Actions Properties

Use Global Owner : Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	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
 - ▶ Fabric A
 - ▶ Fabric B
 - ▶ QoS System Class
 - ▶ LAN Pin Groups
 - ▼ Threshold Policies
 - ▶ thr-policy-default
 - ▶ VLAN Groups
 - ▼ VLANs
 - VLAN default (1)

[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
ℹ 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.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created. 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



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



For the FlexPod solution, it is recommended to place B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward our example of also embedding the cabinet number information giving us 00:25:B5:A1:3B:00 as our first MAC address.

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created. 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



CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

OK

Cancel

8. Click OK 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 : 9000

MAC Pool : MAC-Pool-A(238/256) ▼

QoS Policy : <not set> ▼

Network Control Policy : Enable-CDP-LLDP ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

Connection Policies

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

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 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	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Transmit Queues	:	<input type="text" value="1"/>	[1-1000]
Ring Size	:	<input type="text" value="256"/>	[64-4096]
Receive Queues	:	<input type="text" value="8"/>	[1-1000]
Ring Size	:	<input type="text" value="512"/>	[64-4096]
Completion Queues	:	<input type="text" value="9"/>	[1-2000]
Interrupts	:	<input type="text" value="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.

6. Expand Options and choose Enabled for Receive Side Scaling (RSS).

Create Ethernet Adapter Policy



Name : VMware-HighTrf

Description :

+ Resources

- Options

- | | | |
|--|---|--|
| Transmit Checksum Offload | : | <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled |
| Receive Checksum Offload | : | <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled |
| TCP Segmentation Offload | : | <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled |
| TCP Large Receive Offload | : | <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled |
| Receive Side Scaling (RSS) | : | <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled |
| Accelerated Receive Flow Steering | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |
| Network Virtualization using Generic Routing Encapsulation | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |
| Virtual Extensible LAN | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |
| GENEVE | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |
| AzureStack-Host QoS | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |
| Failback Timeout (Seconds) | : | <input type="text" value="5"/> [0-600] |
| Interrupt Mode | : | <input checked="" type="radio"/> MSI X <input type="radio"/> MSI <input type="radio"/> IN Tx |
| Interrupt Coalescing Type | : | <input checked="" type="radio"/> Min <input type="radio"/> Idle |
| Interrupt Timer (us) | : | <input type="text" value="125"/> [0-65535] |
| RoCE | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |
| Advance Filter | : | <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled |

OK

Cancel

7. 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 :

Peer Name :

vNIC Template :

[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.

- Expand Host Firmware Packages.
- Choose default.
- In the Actions pane, choose Modify Package Versions.
- Choose version 4.1(2b) for both the Blade and Rack Packages.

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

- 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 :

Description :

Mode :

FlexFlash

FlexFlash State : Disable Enable

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

FlexFlash RAID Reporting State : Disable Enable

FlexFlash Removable State : Yes No No Change

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

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.

OK

Cancel

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 C125 servers for a server pool.

9. In Cisco UCS Manager, click Servers.
10. Expand Policies > root.
11. Right-click Server Pool Policy Qualifications.
12. Choose Create Server Pool Policy Qualification.

13. Name the policy UCS-C125.

14. Choose Create Server PID Qualifications.

Choose UCSB-B200-M5 from the PID drop-down list.

Create Server PID Qualifications



PID :

OK

Cancel

15. Click OK

16. Optionally, choose additional qualifications to refine server selection parameters for the server pool.

17. 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.
4. Change the Reboot Policy to User Ack.
5. Choose "On Next Boot" to delegate maintenance windows to server administrators.

General Events

Actions

- Delete
- Show Policy Usage
- Use Global

Properties

Name : default

Description :

Owner : Local

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : Immediate User Ack

Reboot Policy : Immediate User Ack Timer Automatic

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

6. Click Save Changes.

7. Click OK to accept the changes.

Create vMedia Policy for VMware ESXi 7.0 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 [Cisco Custom ISO for UCS 4.1.2a](#) to the Cisco UCS server in order to boot the ESXi installation. To create this policy, follow these steps:



The Cisco Custom ISO for UCS 4.1.2a should also be used for Cisco UCS software release 4.1(2b) and VMware vSphere 7.0.

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-7.0-HTTP.
 6. Enter “Mounts Cisco Custom ISO ESXi7 for UCS 4.1(2a)” in the Description field.
 7. Click Add to add a vMedia Mount.
 8. Name the mount ESXi-7.0-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-7.0.0-16324942-Custom-Cisco-4.1.2a.iso as the Remote File name.



This VMware ESXi 7.0 Cisco Custom ISO can be downloaded from VMware Downloads.



If a working vCenter 7.0 installation is already in your environment, a FlexPod custom ISO for installing ESXi 7.0 with all the necessary drivers for this FlexPod deployment can be created. See section [Create a FlexPod ESXi Custom ISO using VMware vCenter](#) in the Appendix for a procedure to build 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-7.0-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-7.0.0-16324942-Custom-Cisco-4.1."/>
Remote Path	:	<input type="text" value="software/vSphere7"/>
Username	:	<input type="text"/>
Password	:	<input type="password"/>
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:

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 AMD-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:
 1. CDN Control -> Enabled
 2. Quiet Boot -> Disabled

Main Advanced Boot Options Server Management Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name : **AMD-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. Scroll down and set the following within the Processor tab:
 1. Core Performance Boost -> Auto
 2. Global C-state Control -> Disabled
 3. L1 Stream HW Prefetcher -> Enabled
 4. L2 Stream HW Prefetcher -> Enabled
 5. Determinism Slider -> Power
 6. IOMMU -> Enabled
 7. AMD Memory Interleaving -> Auto
 8. AMD Memory Interleaving Size -> Auto
 9. SMEE -> Enabled
 10. SMT Mode -> Auto
 11. SVM Mode -> Enabled

Main **Advanced** Boot Options Server Management Events

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

Advanced Filter Export Print

BIOS Setting	Value
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
L1PI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Core Performance Boost	Auto
Downcore control	Platform Default
Global C-state Control	Disabled
L1 Stream HW Prefetcher	Enabled
L2 Stream HW Prefetcher	Enabled
Determinism Slider	Power
IOMMU	Enabled
Bank Group Swap	Platform Default
Chipselect Interleaving	Platform Default
Configurable TDP Control	Platform Default
AMD Memory Interleaving	Auto
AMD Memory Interleaving Size	Auto
SMEE	Enabled
SMT Mode	Auto
SVM Mode	Enabled
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	Platform Default

Add Delete Info

Save Changes Reset Values

9. Click Save Changes.

10. Click OK.



For more information, see [Performance Tuning for Cisco UCS C125 Rack Server Nodes with AMD Processors](#).

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-lif01a and fcp-lif01b) and two Fibre Channel LIFs are on cluster node 2 (fcp-lif02a and fcp-lif02b). Also, it is assumed that the A LIFs are connected to switching Fabric A and the B LIFs are connected to switching Fabric B.



One boot policy is configured in this procedure. The policy configures the primary target to be fcp-lif01a.

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 7.0 with or without a TPM 2.0 module in the UCS server. TPM 2.0 Modules were not validated with Cisco UCS C125 servers in this validation.

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

OK

Cancel

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-01a.



To obtain this information, log in to the storage cluster and run the `network interface show -vserver In-fra-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-02a.
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-01b.
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-02b.
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](#)

38. 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 AMD-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.

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

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-NA-FlexPod**

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

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

UUID Assignment:

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

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

< Prev Next > **Finish** Cancel

8. Click Next.

Configure Storage Provisioning

To configure storage provisioning, follow these steps:

1. If you have servers with no physical disks in the UCS C4200 chassis, 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:

3. Choose the “Use Connectivity Policy” option to configure the LAN connectivity.
4. Choose FC-Boot from the LAN Connectivity Policy drop-down list.
5. Leave the Initiator Name Assignment at <not set>.

The screenshot shows the 'Create Service Profile Template' wizard in the Networking step. On the left is a vertical navigation menu with 11 steps: 1. Identify Service Profile Template, 2. Storage Provisioning, 3. Networking (highlighted), 4. SAN Connectivity, 5. Zoning, 6. vNIC/vHBA Placement, 7. vMedia Policy, 8. Server Boot Order, 9. Maintenance Policy, 10. Server Assignment, and 11. Operational Policies. The main content area is titled 'Create Service Profile Template' and includes a help icon. It contains the following elements: a header 'Optionally specify LAN configuration information.'; a 'Dynamic vNIC Connection Policy' dropdown menu with the text 'Select a Policy to use (no Dynamic vNIC Policy by default)'; a link 'Create Dynamic vNIC Connection Policy'; a section 'How would you like to configure LAN connectivity?' with radio buttons for 'Simple', 'Expert', 'No vNICs', and 'Use Connectivity Policy' (selected); a 'LAN Connectivity Policy' dropdown menu with 'FC-Boot' selected and a link 'Create LAN Connectivity Policy'; an 'Initiator Name' section with an 'Initiator Name Assignment' dropdown menu set to '<not set>' and a link 'Create IQN Suffix Pool'; a warning message: 'WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.'; and a bottom navigation bar with buttons for '< Prev', 'Next >', 'Finish' (highlighted), and 'Cancel'.

6. 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 ? ×

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 : [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.



Set no zoning options here since the fabric interconnects are in end host (NPV) mode and zoning is being done in the upstream SAN switch.

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								

[Create ISCSI vNIC](#) [Set ISCSI Boot Parameters](#) [Set Uefi Boot Parameters](#)

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

2. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.

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

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 “UCS-C125” for the Server Pool Qualification to choose only UCS C125 servers in the pool.
4. Expand Firmware Management and choose the default Host Firmware Package.

Create Service Profile Template ? X

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](#)

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

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

The screenshot displays the 'Create Service Profile Template' wizard. On the left, a vertical navigation bar lists 11 steps: 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, and 11. Operational Policies (highlighted). The main content area shows the 'Operational Policies' configuration page. It includes a header 'Create Service Profile Template' with a help icon and a close icon. Below the header is a note: 'Optionally specify information that affects how the system operates.' The configuration is organized into sections, each with a collapse/expand icon:

- BIOS Configuration:** A note states 'If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile'. The 'BIOS Policy' dropdown is set to 'AMD-VM-Host'.
- External IPMI/Redfish Management Configuration:** Collapsed.
- Management IP Address:** Collapsed.
- Monitoring Configuration (Thresholds):** Collapsed.
- Power Control Policy Configuration:** A note states 'Power control policy determines power allocation for a server in a given power group.' The 'Power Control Policy' dropdown is set to 'No-Power-Cap'. A link 'Create Power Control Policy' is visible.
- Scrub Policy:** Collapsed.
- KVM Management Policy:** Collapsed.
- Graphics Card Policy:** Collapsed.
- Persistent Memory Policy:** Collapsed.

 At the bottom right, there are four buttons: '< Prev', 'Next >', 'Finish' (highlighted in blue), and '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 AMD-VM-Host-Infra-FCP-A.
3. Right-click AMD-VM-Host-Infra-FCP-A and choose Create a Clone.
4. Name the clone AMD-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 AMD-VM-Host-Infra-FCP-A-vM and choose the vMedia Policy tab.
7. Click Modify vMedia Policy.

8. Choose the ESXi-7.0-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 > FlexPod > Service Template AMD-VM-Host-Infra-FCP-A-VM.
3. Right-click AMD-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 7.0 has been installed on the hosts, the host Service Profiles can be bound to the AMD-VM-Host-Infra-FCP-A Service Profile Template to remove the vMedia Mapping from the host.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All 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: WWPN
Infra-SVM	fcp-lif-01a	Fabric A	<fcp-lif-01a-wwpn>
	fcp-lif-01b	Fabric B	<fcp-lif-01b-wwpn>
	fcp-lif-02a	Fabric A	<fcp-lif-02a-wwpn>
	fcp-lif-02b	Fabric B	<fcp-lif-02b-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 WWPN
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 WWPN information in Cisco UCS Manager GUI, go to `Servers > Service Profiles > root > Sub-Organizations > Organization`. Expand each service profile and then expand vHBAs. Select each vHBA. The WWPN is shown under Properties on the right.

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 using the Cisco Nexus 93180YC-FX for both LAN and SAN switching, please refer to FlexPod with Nexus 93180YC-FX SAN Switching Configuration - Part 2 in the Appendix.



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(1a).

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 and Local Time Configuration

Cisco MDS 9132T A and Cisco MDS 9132T B

To configure the second NTP server and add local time configuration, follow this step:

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

```
ntp server <nexus-B-mgmt0-ip>
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day> <end-month> <end-time> <offset-minutes>
```



It is important to configure the local time so that logging time alignment, any backup schedules, and SAN Analytics forwarding are correct. For more information on configuring the timezone and daylight savings time or summer time, please see [Cisco MDS 9000 Series Fundamentals Configuration Guide, Release 8.x](#). Sample clock commands for the United States Eastern timezone are:

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

Configure Individual Ports

Cisco MDS 9132T A

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

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

```
interface fc1/9
switchport description <st-clustername>-1:2a
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/10
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 Cisco 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:

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

```
interface fc1/9
switchport description <st-clustername>-1:2b
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/10
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 Cisco 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 this step:

1. 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 fcl/9
vsan <vsan-a-id> interface fcl/10
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 this steps:

1. 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 fcl/9
vsan <vsan-b-id> interface fcl/10
vsan <vsan-b-id> interface port-channel15
exit
```

At this point, it may be necessary to go into Cisco 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 this step:

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

```
device-alias mode enhanced
device-alias database
device-alias name Infra-SVM-fcp-lif-01a pwnn <fcp-lif-01a-wwpn>
device-alias name Infra-SVM-fcp-lif-02a pwnn <fcp-lif-02a-wwpn>
```

```
device-alias name VM-Host-Infra-01-A pwnn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwnn <vm-host-infra-02-wwpna>
device-alias name VM-Host-Infra-03-A pwnn <vm-host-infra-03-wwpna>
device-alias commit
```

Cisco MDS 9132T B

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

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

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

Create Zones and Zoneset

Cisco MDS 9132T A

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

```
configure terminal
zone name Infra-SVM-Fabric-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-01-A init
member device-alias VM-Host-Infra-02-A init
member device-alias VM-Host-Infra-03-A init
member device-alias Infra-SVM-fcp-lif-01a target
member device-alias Infra-SVM-fcp-lif-02a 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-01b target
member device-alias Infra-SVM-fcp-lif-02b 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 7.0 Setup

VMware ESXi 7.0

This section provides detailed instructions for installing VMware ESXi 7.0 in a FlexPod environment. After the procedures are completed, three 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 7.0 from VMware

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

1. Click the following link: [Cisco Custom ISO for UCS 4.1.2a](#).



The Cisco Custom ISO for UCS 4.1.2a should also be used for Cisco UCS software release 4.1(2b) and VMware vSphere 7.0.

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 into 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.



If the ESXi installer fails to load because the software certificates cannot be validated, reset the server, and when prompted, press F2 to go into BIOS and set the system time and date to current. Then the ESXi installer should load properly.

-
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

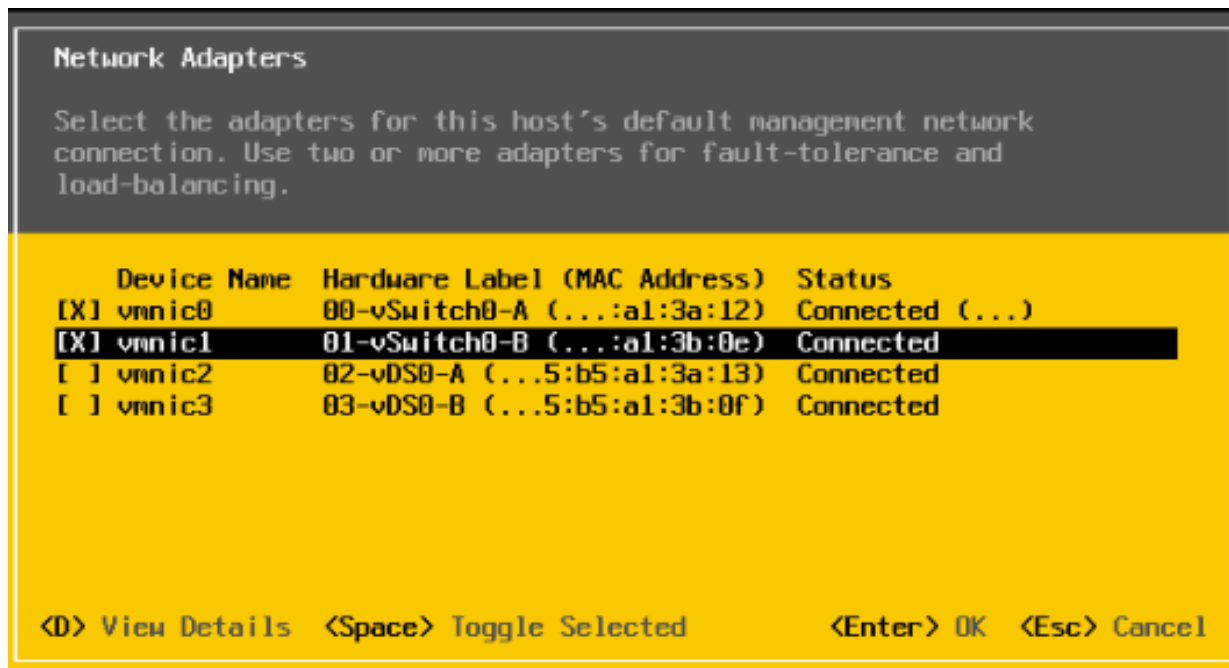
Adding a management network for each VMware host is necessary for managing the host. The following section details how to add a management network for the VMware hosts.

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.

- 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.
- 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.

- Press Enter.
- Choose the VLAN (Optional) option and press Enter.
- Enter the <ib-mgmt-vlan-id> and press Enter.
- Choose IPv4 Configuration and press Enter.
- Choose the “Set static IPv4 address and network configuration” option by using the arrow keys and space bar.
- Move to the IPv4 Address field and enter the IP address for managing the ESXi host.
- Move to the Subnet Mask field and enter the subnet mask for the ESXi host.
- Move to the Default Gateway field and enter the default gateway for the ESXi host.
- Press Enter to accept the changes to the IP configuration.
- 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.



Since 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. If this VMware ESXi host is iSCSI booted, the vmk1, iScsiBootPG-A interface's MAC address can also be reset to a random, VMware-assigned MAC address.
 1. Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk1. vmk1 should be a part of the "iScsiBootPG-A" port group and should have a MAC address from the UCS MAC Pool. Note the IP address and netmask of vmk1.
 2. To remove vmk1, type `esxcfg-vmknic -d "iScsiBootPG-A"`.
 3. To re-add vmk1 with a random MAC address, type `esxcfg-vmknic -a -i <vmk1-ip> -n <vmk1-netmask> -m 9000 "iScsiBootPG-A"`.
 4. Verify vmk1 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.
10. Type `exit` to log out of the command line interface.
11. Type `Ctrl-Alt-F2` to return to the ESXi console menu interface.

Install VMware and Cisco VIC Drivers for the ESXi Host

Download the offline bundle for the UCS Tools Component and the NetApp NFS Plug-in for VMware VAAI to the Management workstation:

[UCS Tools Component for ESXi 7.0 1.1.5](#) (ucs-tool-esxi_1.1.5-1OEM.zip)

[NetApp NFS Plug-in 1.1.2-3 for VMware VAAI](#) (ucs-tool-esxi_1.1.2-1OEM.zip)



This document is using the driver versions shown above along with Cisco VIC `nenic` version 1.0.33.0 and `nfnic` version 4.0.0.56 along with VMware vSphere version 7.0.0, Cisco UCS version 4.1(2b), and the latest patch of NetApp ONTAP 9.7. These were the versions validated and supported at the time this document was published. This document can be used as a guide for configuring future versions of software. Consult the [Cisco UCS Hardware Compatibility List](#) and the [NetApp Interoperability Matrix Tool](#) to determine supported combinations of firmware and software.

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 two offline bundles referenced above to the `/tmp` directory on each ESXi host.
2. Using a 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 component apply -d /tmp/ucs-tool-esxi_1.1.5-1OEM.zip
esxcli software vib install -d /tmp/NetAppNasPlugin.v23.zip
reboot
```

-
5. After reboot, log back into each host and run the following commands and ensure the correct version is installed:

```
esxcli software component list | grep ucs  
esxcli software vib list | grep NetApp
```

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. Optionally, create two more vMotion VMkernel NICs to increase the speed of multiple simultaneous vMotion on this solution's 50GE vNICs.
 1. Click Add VMkernel NIC.

-
2. For New port group, enter VMkernel-vMotion1.
 3. For Virtual switch, choose vSwitch0.
 4. Enter <vmotion-vlan-id> for the VLAN ID.
 5. Change the MTU to 9000.
 6. Choose Static IPv4 settings and expand IPv4 settings.
 7. Enter the ESXi host's second vMotion IP address and netmask.
 8. Choose the vMotion stack for TCP/IP stack.
 9. Click Create.
 10. Click Add VMkernel NIC.
 11. For New port group, enter VMkernel-vMotion2.
 12. For Virtual switch, choose vSwitch0.
 13. Enter <vmotion-vlan-id> for the VLAN ID.
 14. Change the MTU to 9000.
 15. Choose Static IPv4 settings and expand IPv4 settings.
 16. Enter the ESXi host's third vMotion IP address and netmask.
 17. Choose the vMotion stack for TCP/IP stack.
 18. Click Create.
34. 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: 6
 Uplinks: 2

vSwitch Details

MTU	9000
Ports	11776 (11763 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	No
Allow MAC changes	No

Shaping policy

Enabled	No
---------	----

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

Port groups	Virtual switches	Physical NICs	VMkernel NICs	TCP/IP stacks	Firewall rules
Add VMkernel NIC Edit settings Refresh Actions Search					
Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.156.191	None
vmk1	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.191	None
vmk2	VMkernel-vMotion	vMotion stack	vMotion	192.168.100.191	None
vmk3	VMkernel-vMotion1	vMotion stack	vMotion	192.168.100.201	None
vmk4	VMkernel-vMotion2	vMotion stack	vMotion	192.168.100.211	None

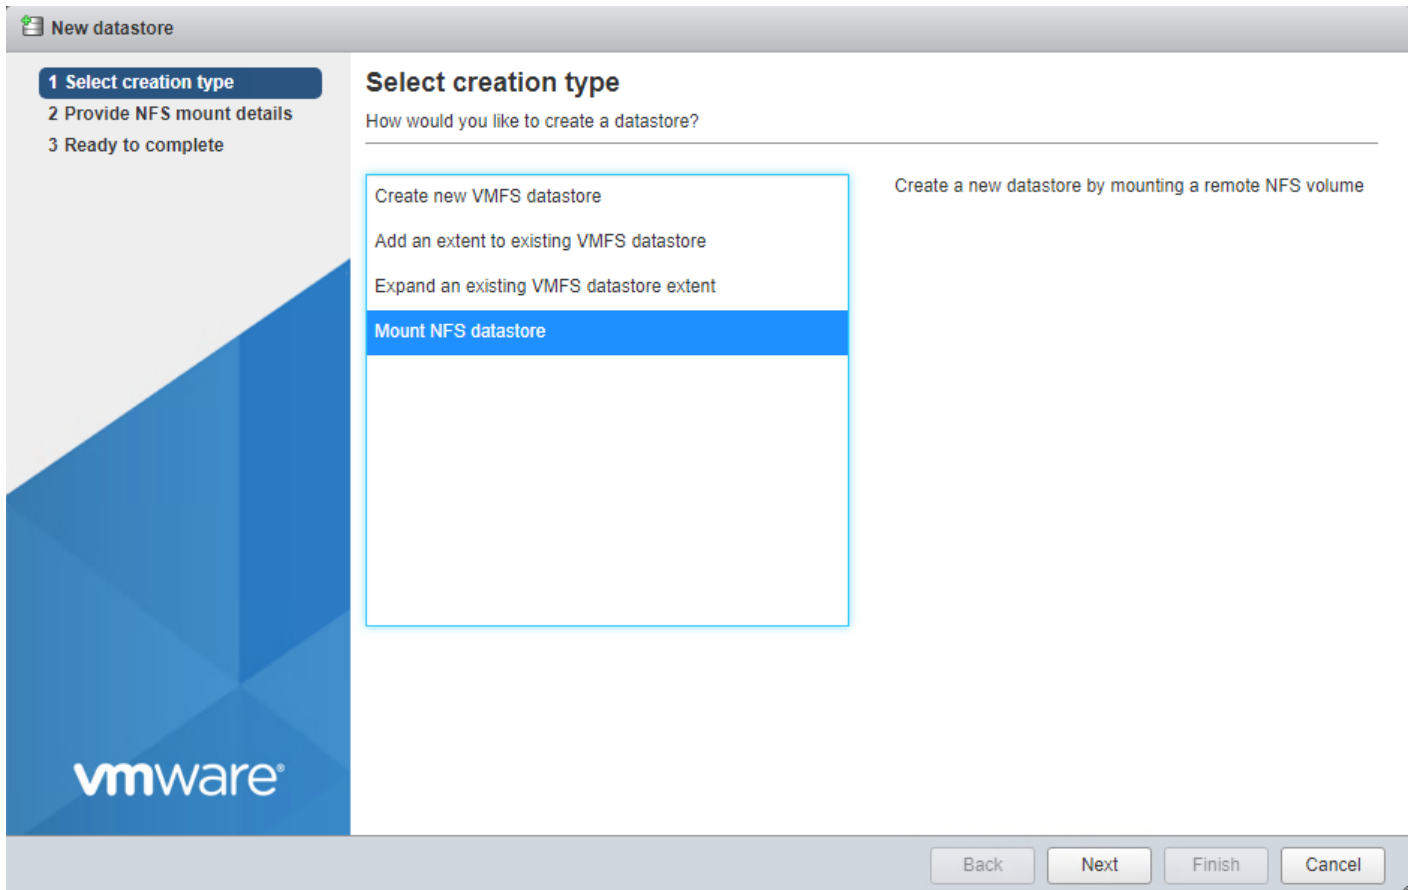
5 items

Mount Required Datastores

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, 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-02` LIF for the NFS server. Input `/infra_datastore` for the NFS share. Leave the NFS version set at NFS 3. Click Next.

New datastore - infra_datastore


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

Provide NFS mount details

Provide the details of the NFS share you wish to mount

Name	infra_datastore
NFS server	192.168.50.52
NFS share	/infra_datastore
NFS version	<input checked="" type="radio"/> NFS 3 <input type="radio"/> NFS 4

Back Next Finish Cancel



- Click Finish. The datastore should now appear in the datastore list.
- In the center pane, choose New Datastore to add a new datastore.
- In the New datastore popup, choose Mount NFS datastore and click Next.
- Input infra_swap for the datastore name. Input the IP address for the nfs-lif-01 LIF for the NFS server. Input /infra_swap for the NFS share. Leave the NFS version set at NFS 3. Click Next.
- Click Finish. The datastore should now appear in the datastore list.

Datstores Adapters Devices Persistent Memory

New datastore
 Increase capacity
 Register a VM
 Datastore browser
 Refresh
 Actions

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
infra_datastore	Unknown	1,024 GB	3.85 MB	1,024 GB	NFS	Supported	Single
infra_swap	Unknown	100 GB	364 KB	100 GB	NFS	Supported	Single

2 items

Configure NTP on First ESXi Host

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, choose Manage.
2. In the center pane, choose System > Time & date.
3. Click Edit NTP settings.
4. Make sure “Manually configure the date and time on this host and enter the approximate date and time.
5. Select Use Network Time Protocol (enable NTP client).
6. Use the drop-down list to choose Start and stop with host.
7. 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

07/22/2020 6:56 PM

Use Network Time Protocol (enable NTP client)

NTP service startup policy	Start and stop with host
NTP servers	10.1.156.11,10.1.156.12

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

Save Cancel

8. Click Save to save the configuration changes.



It currently is not possible to start NTP from the ESXi Host Client. NTP will be started from vCenter. The NTP server time may initially vary slightly from the host time.

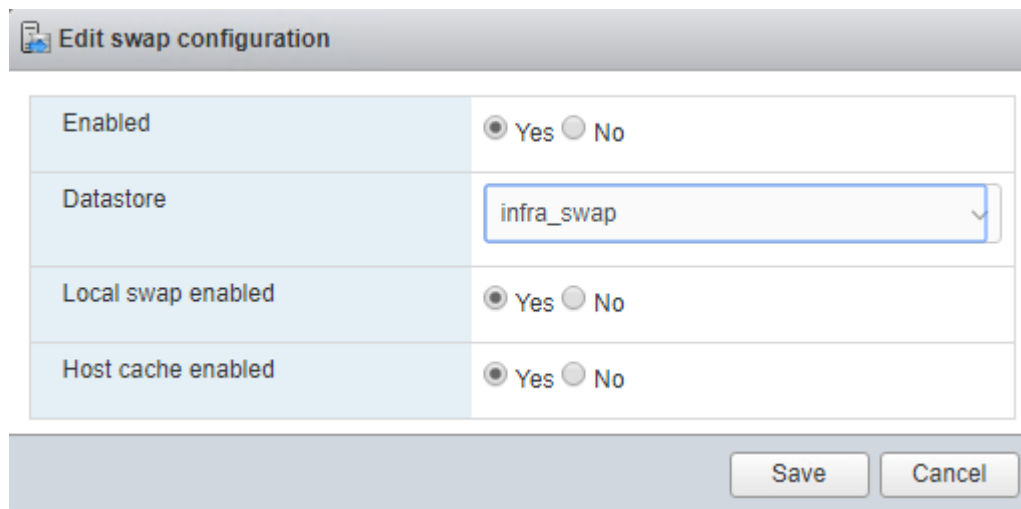
Configure ESXi Host Swap

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.

Configure Host Power Policy


ESXi Host VM-Host-Infra-01

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



Implementation of this policy is recommended in [Performance Tuning for Cisco UCS C125 Rack Server Nodes with AMD Processors](#) for maximum performance. If your organization has specific power policies, please set this policy accordingly.

1. From the Host Client, choose Manage.
2. In the center pane, choose Hardware > Power Management.
3. Choose Change policy.
4. Choose High performance and click OK.



If you are implementing iSCSI boot, execute the VMware ESXi setup scripts in the iSCSI Addition Appendix.

VMware vCenter 7.0D

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 7.0D Server Appliance in a FlexPod environment. After the procedures are completed, a VMware vCenter Server will be configured.

Build the VMware vCenter Server Appliance

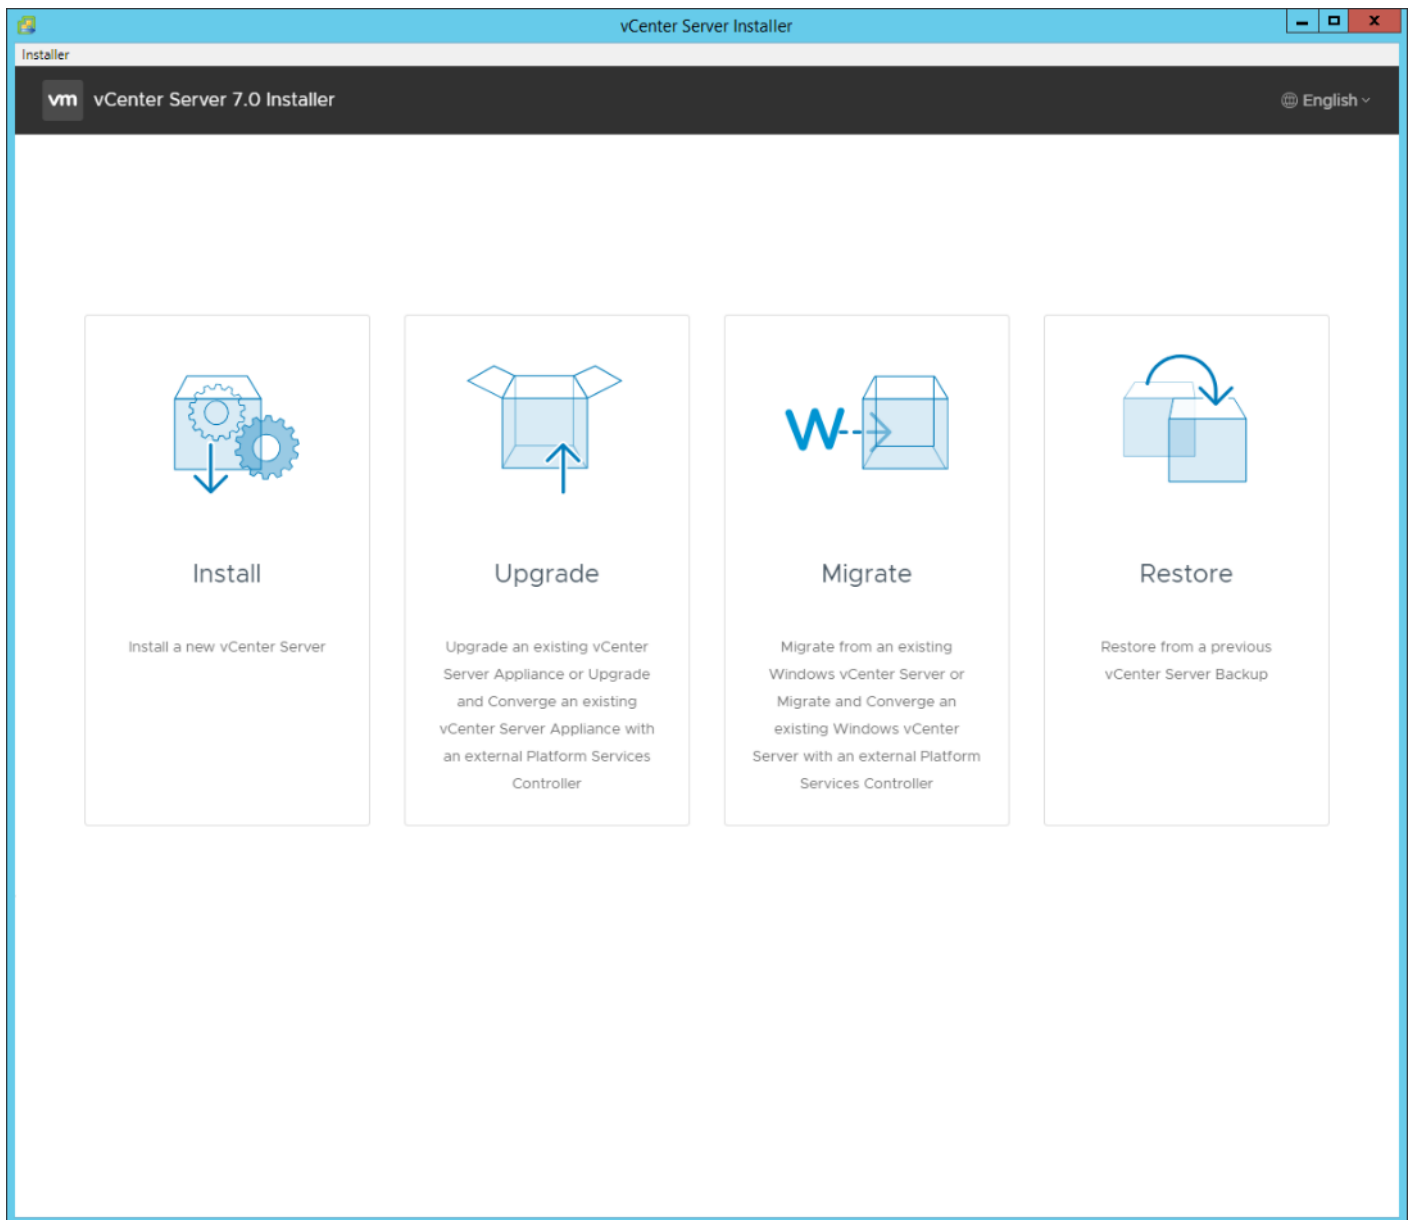
The VCSA deployment consists of 2 stages: install and configuration. To build the VMware vCenter virtual machine, follow these steps:

1. Locate and copy the VMware-VCSA-all-7.0.0-16749653.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 7.0 vCenter Server Appliance.

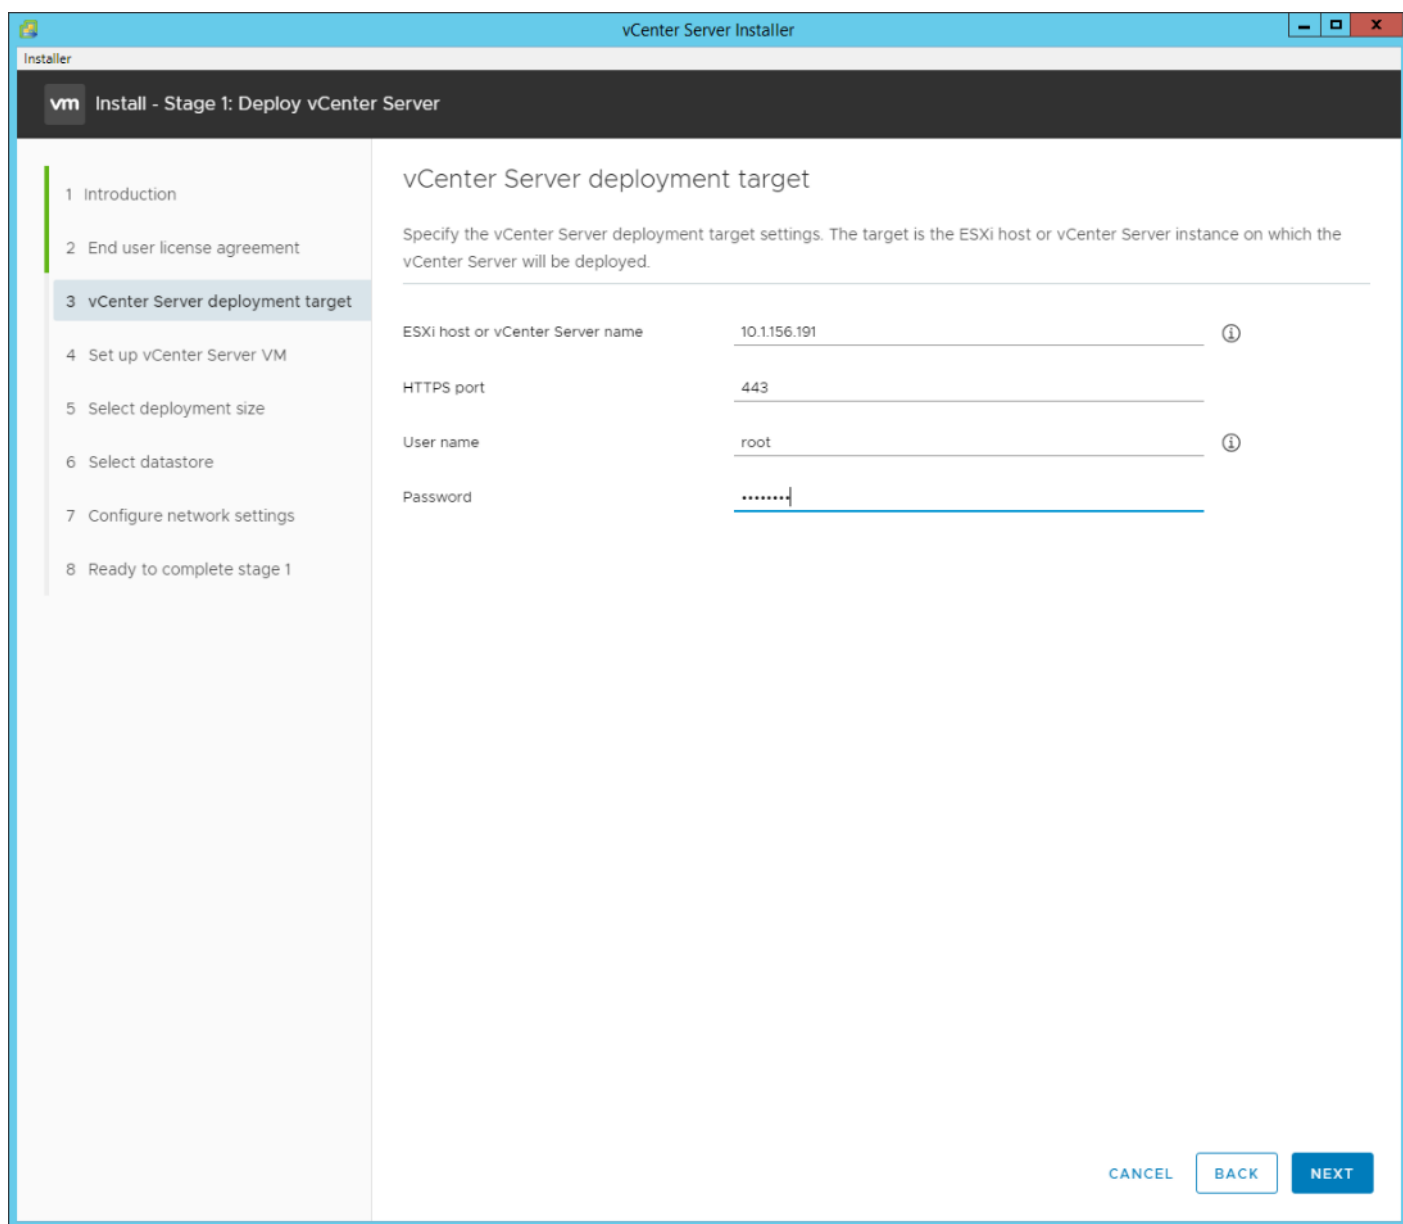


It is important to use at minimum VMware vCenter release 7.0B to ensure access to all needed features.

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 and above).
3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard 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 "vCenter Server deployment target" window, enter the host name or IP address of the first ESXi host, User name (root) and Password. Click NEXT.

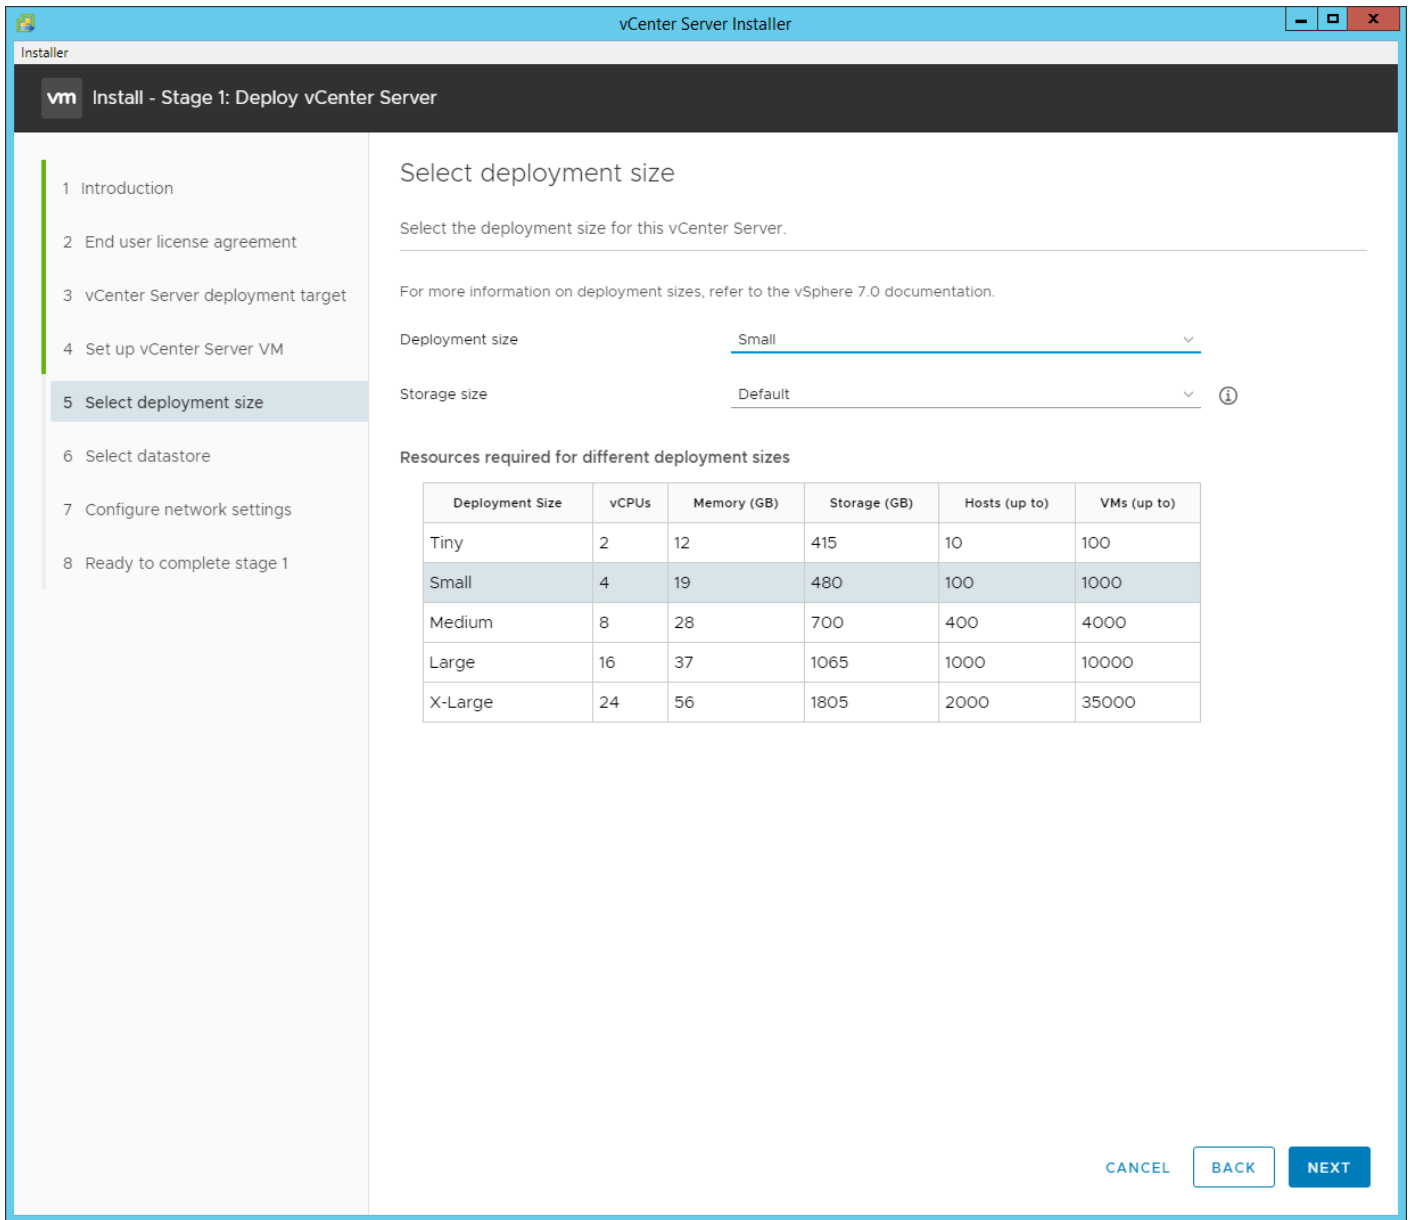


8. Click YES to accept the certificate.

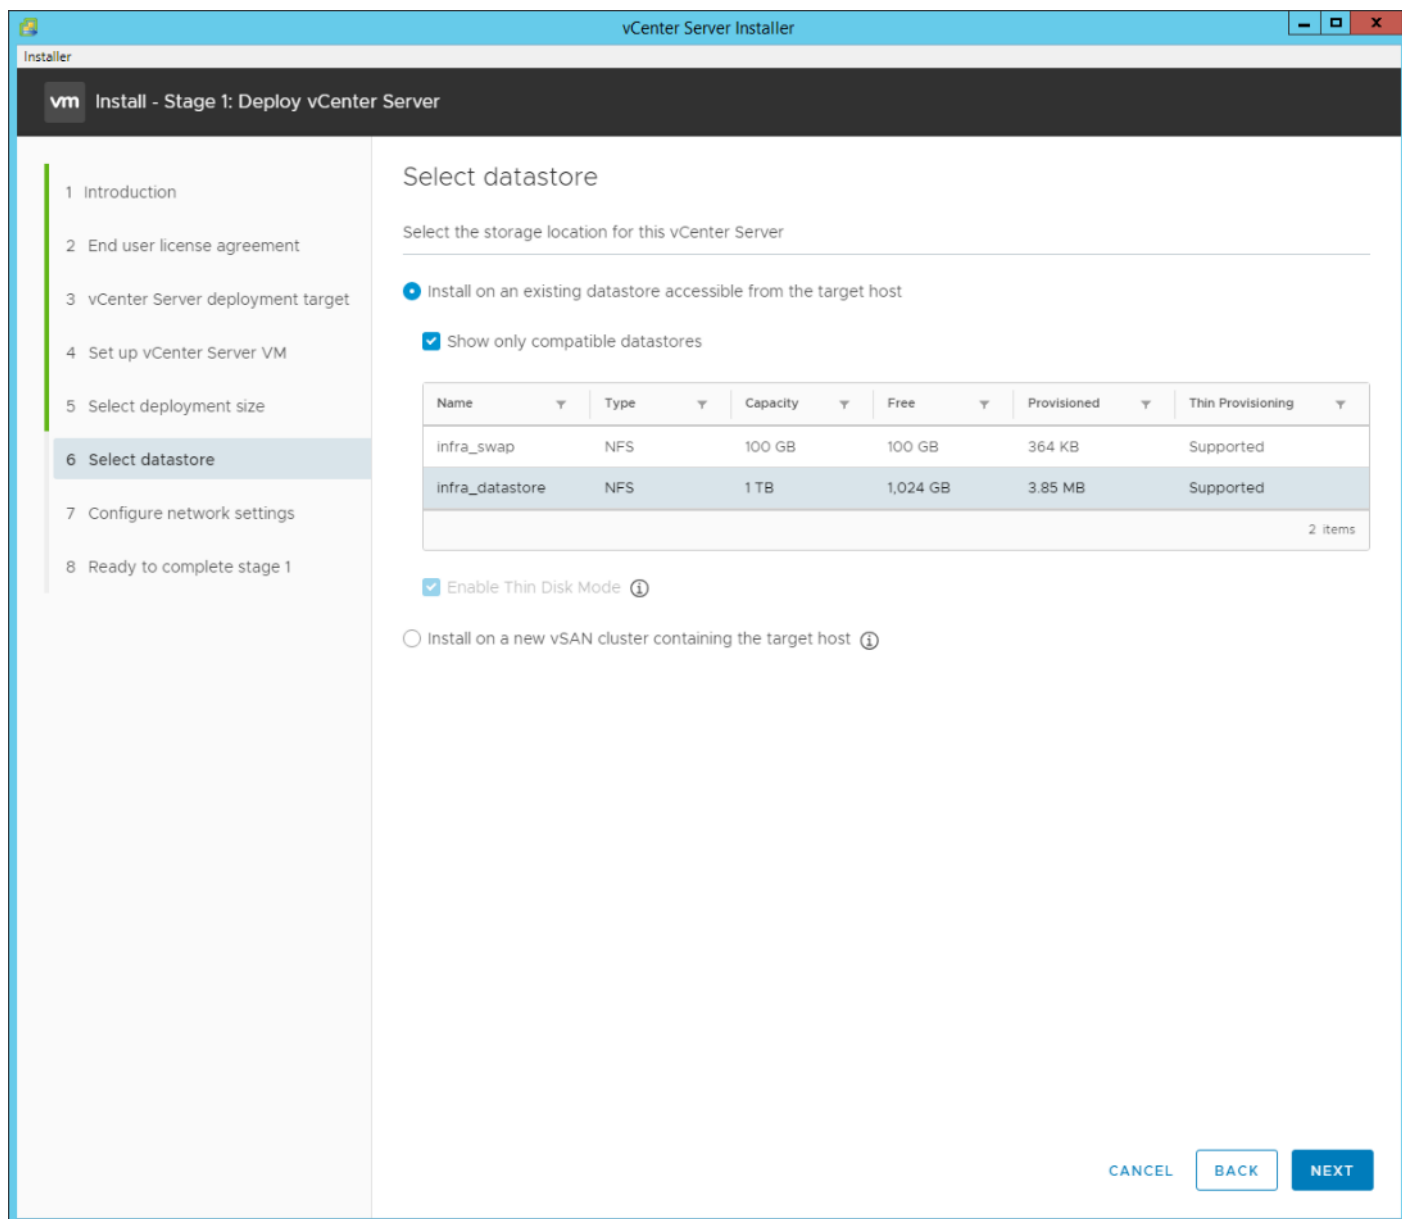
9. Enter the Appliance VM name and password details in the “Set up vCenter Server VM” section. Click NEXT.



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



11. Choose infra_datastore for storage. Click NEXT.



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

1. Choose a Network: IB-MGMT Network.



It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it isn't 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. Mov-

ing vCenter to a different host on vSwitch0 to be brought up always occurs correctly without requiring vCenter to already be up and running.

2. IP version: IPV4
3. IP assignment: static
4. FQDN: <vcenter-fqdn>
5. IP address: <vcenter-ip>
6. Subnet mask or prefix length: <vcenter-subnet-mask>
7. Default gateway: <vcenter-gateway>
8. DNS Servers: <dns-server1>,<dns-server2>

Configure network settings		
Configure network settings for this vCenter Server		
Network	IB-MGMT Network	ⓘ
IP version	IPv4	
IP assignment	static	
FQDN	na-vc.flexpod.cisco.com	ⓘ
IP address	10.1156.200	
Subnet mask or prefix length	255.255.255.0	ⓘ
Default gateway	10.1156.254	
DNS servers	10.1156.250,10.1156.251	
Common Ports		
HTTP	80	
HTTPS	443	


CANCEL BACK NEXT

13. Click NEXT.

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



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

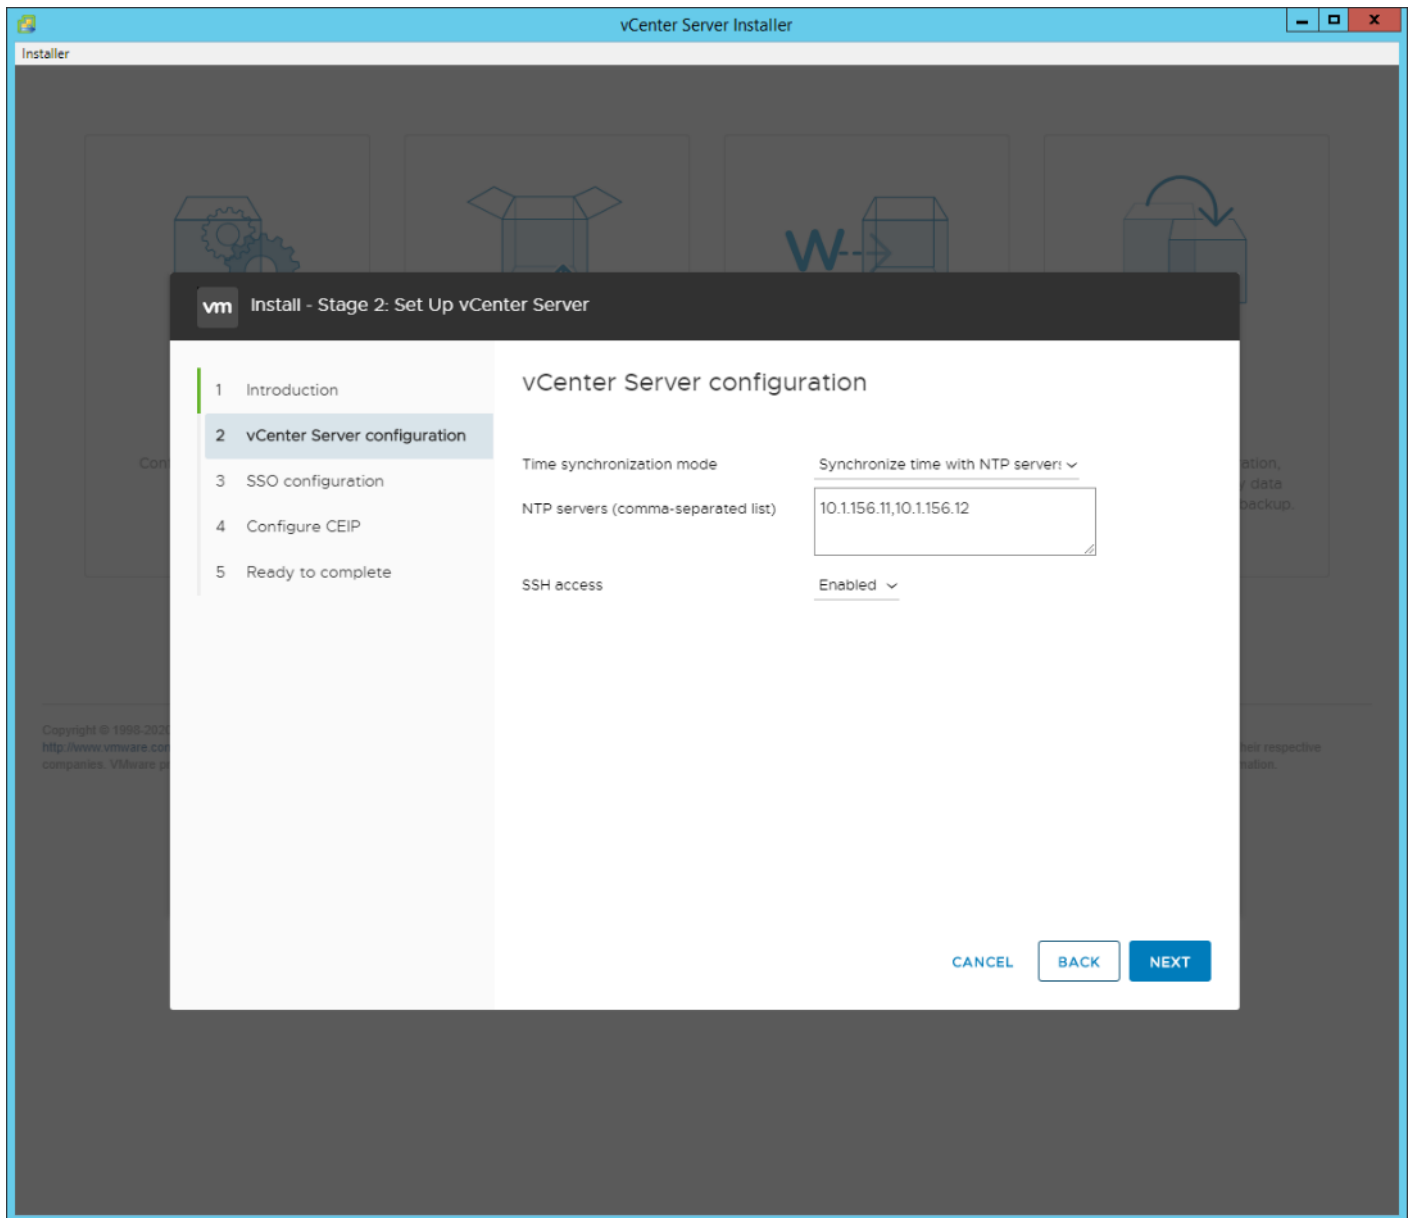


15. Click CONTINUE to proceed with stage 2 configuration.

16. Click NEXT.

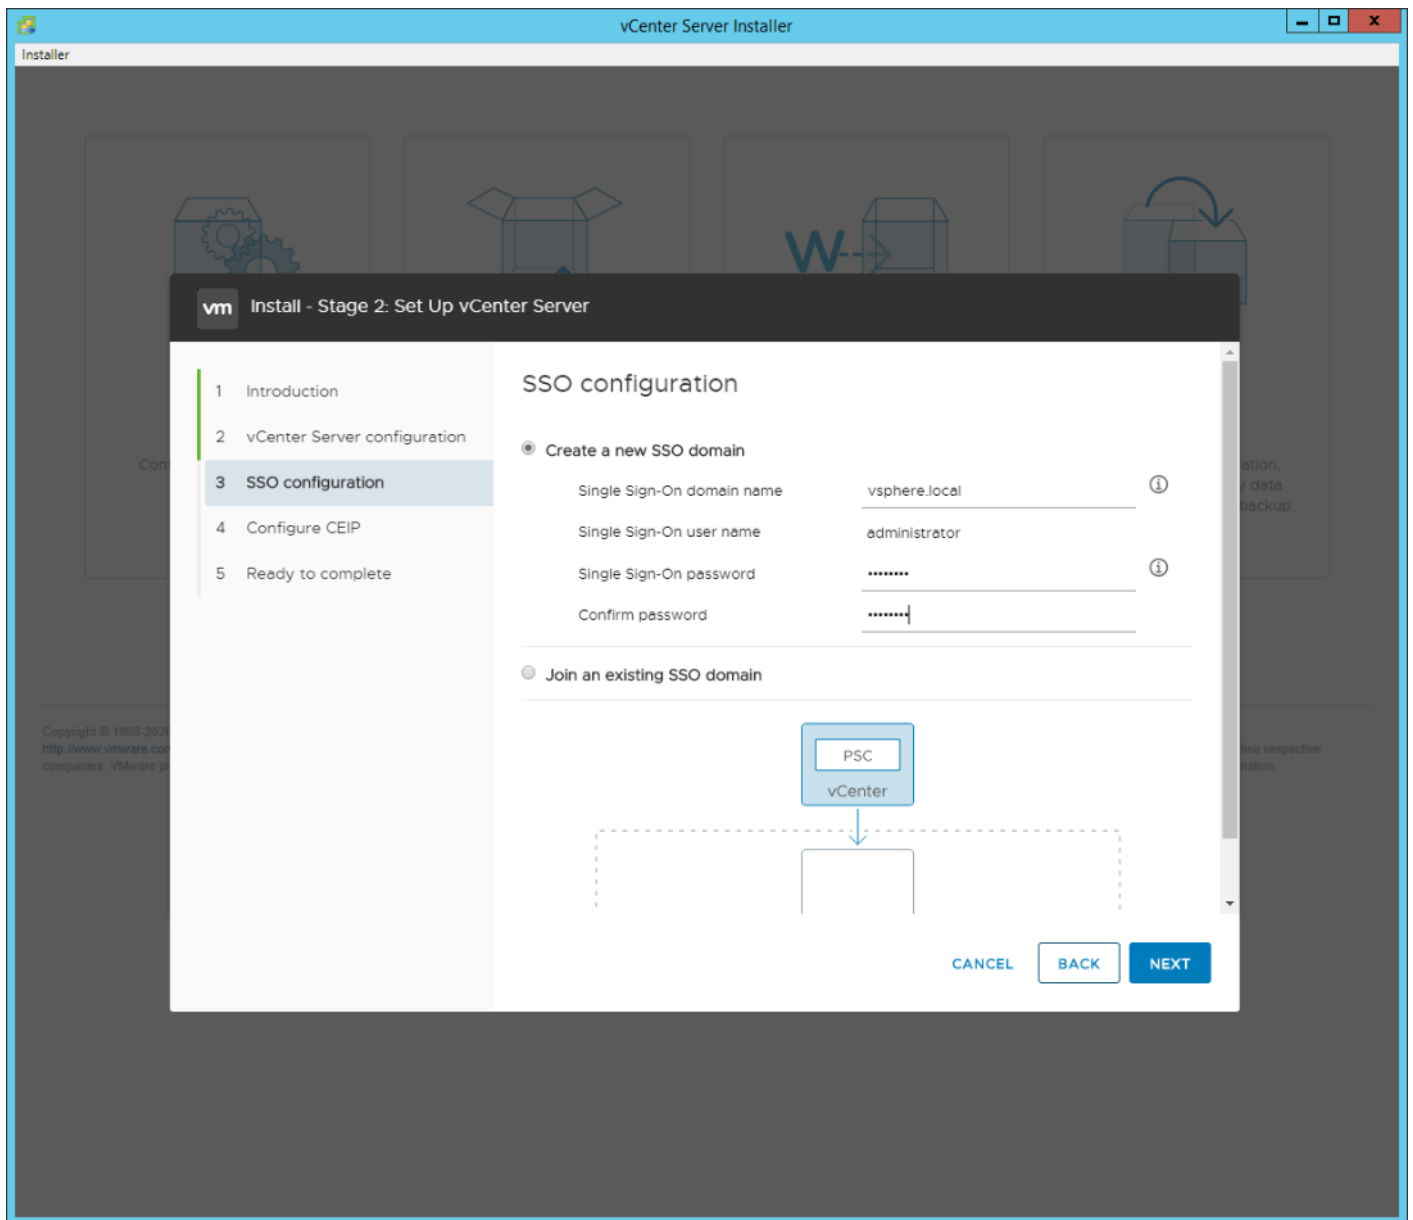
17. In the vCenter Server configuration window, configure these settings:

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



18. Click NEXT.

19. Complete the SSO configuration as shown below, or according to your organization's security policies:



20. Click NEXT.

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


22. Click NEXT.

23. Review the configuration and click FINISH.

24. Click OK.



vCenter Server setup will take a few minutes to complete.



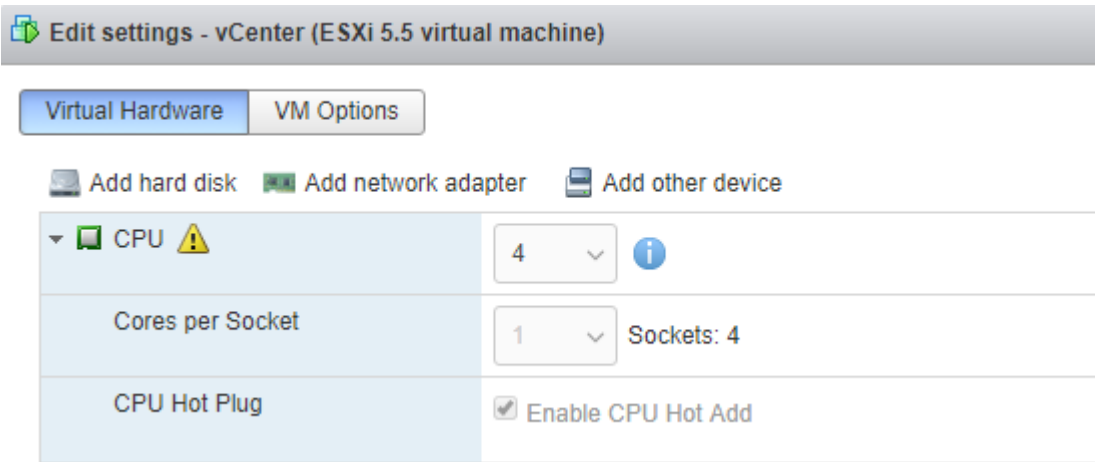
25. Click CLOSE. Eject or unmount the VCSA installer ISO.

Adjust vCenter CPU Settings

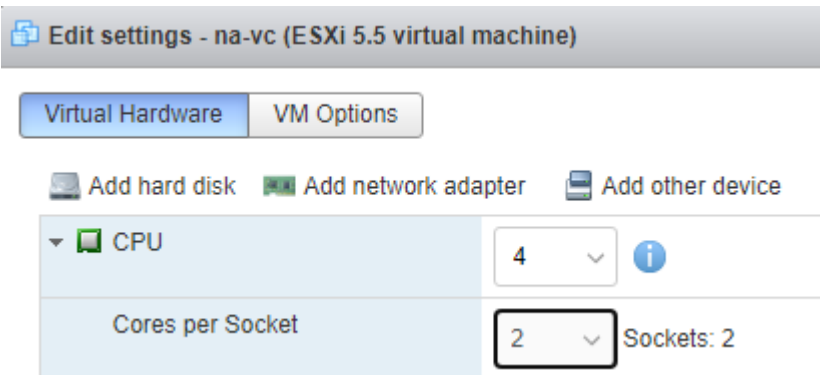
If a vCenter deployment size of Small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration. Cisco UCS B and C-Series servers are normally 2-socket servers. In this validation, the Small deployment size was 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. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.
5. On the left, choose Virtual Machines.
6. In the center pane, right-click the vCenter VM and choose Edit settings.
7. In the Edit settings window, expand CPU and check the value of Sockets.



8. If the number of Sockets does not match your server configuration, it will need to be adjusted. Click Cancel.
9. If the number of Sockets needs to be adjusted:
 1. Right-click the vCenter VM and choose Guest OS > Shut down. Click Yes on the confirmation.
 2. Once vCenter is shut down, right-click the vCenter VM and choose Edit settings.
 3. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to your server configuration (normally 2).



4. Click Save.

-
5. 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 VMware vCenter Server 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.
6. In the menu on the left choose Administration.
7. According to your Security Policy, adjust the settings for the root user and password.
8. In the menu on the left choose Update.
9. Follow the prompts to STAGE AND INSTALL any available vCenter updates. In this validation, vCenter version 7.0.0.10700 was installed.
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-fqdn>>. You will need to navigate security screens.



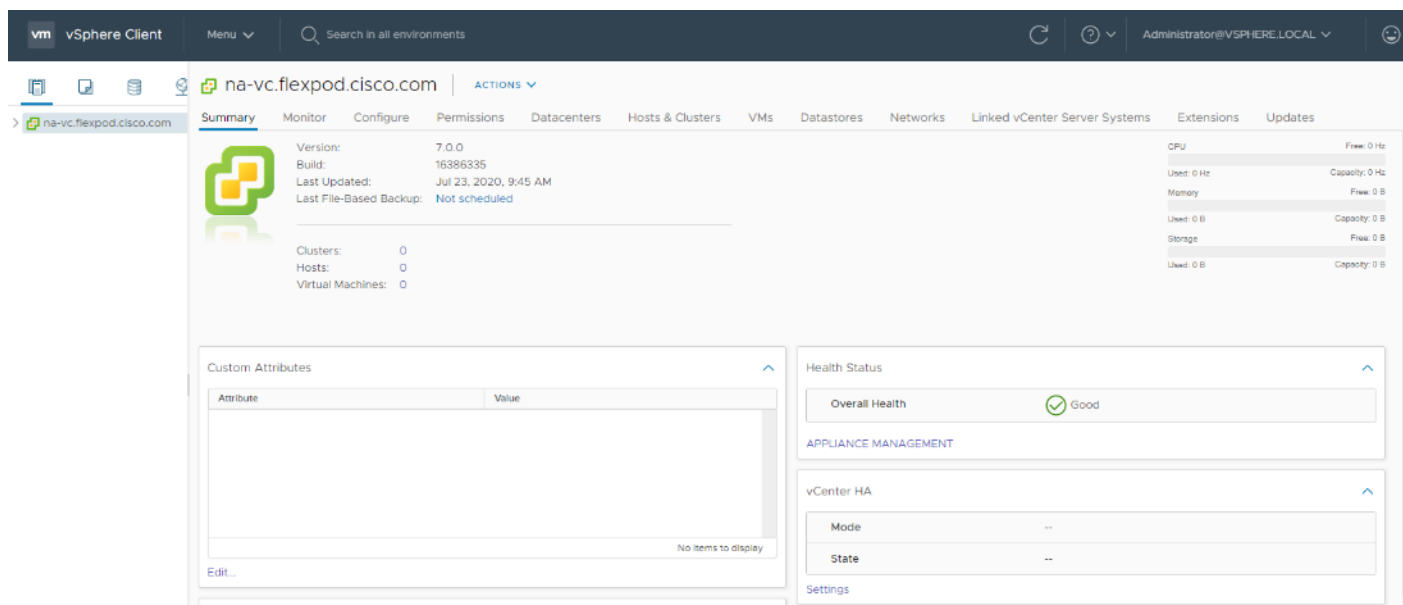
With VMware vCenter 7.0, the use of the vCenter FQDN is required.

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 the only option in vSphere 7 and will be used going forward.

13. Log in using the Single Sign-On username (administrator@vsphere.local) and password created during the vCenter installation. Dismiss the Licensing warning at this time.



14. In the center pane, choose ACTIONS > New Datacenter.

15. Type “FlexPod-DC” in the Datacenter name field.

New Datacenter



Name

FlexPod-DC

Location:

 na-vc.flexpod.cisco.com

CANCEL

OK

16. Click OK.

17. Expand the vCenter on the left.

18. Right-click the datacenter FlexPod-DC in the list in the left pane. Choose New Cluster.

19. Name the cluster FlexPod-Management.

20. Turn on DRS and vSphere HA. Do not turn on vSAN.

New Cluster

FlexPod-DC



Name	FlexPod-Management
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.

Manage all hosts in the cluster with a single image

CANCEL

OK

21. Click OK to create the new cluster.
22. Right-click "FlexPod-Management" and choose Settings.
23. Choose Configuration > General in the list located on the left and choose EDIT located on the right of General.
24. 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

25. Right-click “FlexPod-Management” and click Add Hosts.
26. In the IP address or FQDN field, enter either the IP address or the FQDN of the first VMware ESXi host. Enter root as the Username and the root password. Click NEXT.
27. In the Security Alert window, choose the host and click OK.
28. Verify the Host summary information and click NEXT.
29. Ignore warnings about the host being moved to Maintenance Mode and click FINISH to complete adding the host to the cluster.
30. The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.
31. In the list, right-click the added ESXi host and choose Settings.
32. In the center pane under Virtual Machines, choose Swap File location.
33. On the right, click EDIT.
34. Choose the infra_swap datastore and click OK.

Edit Swap File Location

na-esxi-1.flexpod.cisco.com




Select a location to store the swap files.

Virtual machine directory

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

Use a specific datastore

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

Name	Capacity	Provisioned	Free Space	Type	Thin Provisioned
Infra_datastore	1.00 TB	504.92 GB	1,011.55 GB	NFS	Supported
Infra_swap	100.00 GB	8.42 MB	99.99 GB	NFS	Supported

2 items

CANCEL

OK

35. In the list under System, choose Time Configuration.

36. Click EDIT to the right of Manual Time Configuration. Set the time and date to the correct local time and click OK.


37. Click EDIT to the right of Network Time Protocol.

38. In the Edit Network Time Protocol window, select Enable and then select Start NTP Service. Ensure the other fields are filled in correctly and click OK.

Edit Network Time Protocol

na-esxi-1.flexpod.cisco.com



Enable 

NTP Servers	<input type="text" value="10.1156.11,10.1156.12"/>
Separate servers with commas, e.g. 10.31.21.2, fe00::2800	
NTP Service Status:	Stopped <input checked="" type="checkbox"/> Start NTP Service
NTP Service Startup Policy:	<input type="text" value="Start and stop with host"/>

CANCEL

OK

39. In the list under Hardware, choose Overview. Scroll to the bottom and ensure the Power Management Active policy is High Performance. If the Power Management Active policy is not High Performance, to the right of Power Management, choose EDIT POWER POLICY. Choose High performance and click OK.
40. In the list under Storage, choose Storage Devices. Make sure the NETAPP Fibre Channel Disk LUN 0 or NETAPP iSCSI Disk LUN 0 is selected.
41. Choose the Paths tab.
42. Ensure that 4 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 Mark as Perennially Reserved

Name	LUN	Type	Capacity	Datastore	Operational St...	Hardware Accelerat...	Drive Ty...	Transport
Local ATA Disk (t10.ATA____Micron_5100_MTF...	0	disk	223.57 GB	Not Consu...	Attached	Not supported	Flash	Block Adapter
NETAPP Fibre Channel Disk (naa.600a09803831...	0	disk	32.00 GB	Not Consu...	Attached	Supported	Flash	Fibre Channel
Local ATA Disk (t10.ATA____Micron_5100_MTF...	0	disk	223.57 GB	Not Consu...	Attached	Not supported	Flash	Block Adapter

Copy All 3 items

Properties Paths Partition Details

Enable Disable

Runtime Name	Status	Target	Name	Preferred
vmhba0:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:16:6b:8b 20:01:d0:39:ea:16:6b:8b	vmhba0:C0:T1:L0	
vmhba1:C0:T2:L0	Active	20:00:d0:39:ea:16:6b:8b 20:04:d0:39:ea:16:6b:...	vmhba1:C0:T2:L0	
vmhba1:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:16:6b:8b 20:02:d0:39:ea:16:6b:...	vmhba1:C0:T1:L0	
vmhba0:C0:T2:L0	Active	20:00:d0:39:ea:16:6b:8b 20:03:d0:39:ea:16:6b:...	vmhba0:C0:T2:L0	

Add AD User Authentication to vCenter (Optional)

If an AD Infrastructure is set up in this FlexPod environment, you can set up 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 Identity Provider tab.
6. In the list under Type, select Active Directory Domain.
7. Choose JOIN AD.
8. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Do not fill in an Organizational unit. Click JOIN.
9. Click Acknowledge.
10. In the list on the left under Deployment, choose System Configuration. Choose the radio button to choose the vCenter, then choose REBOOT NODE.

-
11. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.
 12. Log back into the vCenter vSphere HTML5 Client as Administrator@vsphere.local.
 13. Under Menu, choose Administration. In the list on the left, under Single Sign On, choose Configuration.
 14. In the center pane, under Configuration, choose the Identity Provider tab. Under Type, select Identity Sources. Click ADD.
 15. Make sure your Active Directory (Integrated Windows Authentication) is selected, your Windows Domain name is listed, and Use machine account is selected. Click ADD.
 16. In the list select the Active Directory (Integrated Windows Authentication) Identity source type. If desired, select SET AS DEFAULT and click OK.
 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 Domain.
 20. On the User/Group 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(s) 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 for the 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 version 7.0.0 - ESXi 7.0 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="IGMP/MLD snooping"/>
Discovery protocol	
Type	<input type="text" value="Link Layer Discovery Protocol"/>
Operation	<input type="text" value="Both"/>
Administrator contact	
Name	<input type="text"/>
Other details	<input type="text"/>

CANCEL

OK

14. For the 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 Cisco 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, failback, and uplink failover order.

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

Failover order ⓘ

↑ ↓

Active uplinks
Uplink 2
Standby uplinks
Uplink 1
Unused uplinks

CANCEL

BACK

NEXT

20. Click NEXT.

21. Leave NetFlow disabled and click NEXT.

22. Leave Block all ports set as No and click NEXT.

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

24. 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
na-esxi-1.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	--	--

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. Repeat steps 1-30 to assign all vMotion VMkernel ports to the vMotion destination port group.



Do not migrate the other VMkernel ports.

vDS0 - Add and Manage Hosts

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

Manage VMkernel adapters

Manage and assign VMkernel network adapters to the distributed switch.

Assign port group Reset changes View settings

Host/VMkernel Network Adapters	In Use by Switch	Source Port Group	Destination Port Gr...
na-esxi-1.flexpod.cisco.com			
▲ On this switch			
vmk2 (Reassigned)	vSwitch0	VMkernel-vMotion	vMotion
vmk3 (Reassigned)	vSwitch0	VMkernel-vMotion1	vMotion
vmk4 (Reassigned)	vSwitch0	VMkernel-vMotion2	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(s) have a valid and correct Destination Port Group and click NEXT.



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

33. 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 7.0b installed, the management IP address set, and the Cisco UCS Tool 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 for 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, optionally selecting "Use the same credentials for all hosts". 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(s).

The screenshot shows the 'Add hosts' wizard in vCenter. The left sidebar has three steps: '1 Add hosts', '2 Host summary', and '3 Ready to complete', with the third step highlighted. The main area is titled 'Review and finish' and contains an information box stating: 'Hosts will enter maintenance mode before they are moved to the cluster. You might need to either power off or migrate powered on and suspended virtual machines.' Below this, it lists '2 new hosts will be connected to vCenter Server and moved to this cluster:' followed by the hostnames 'nx-esxi-2.flexpod.cisco.com' and 'nx-esxi-3.flexpod.cisco.com'. At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'FINISH'.

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 for the 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 Standard Switch: 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

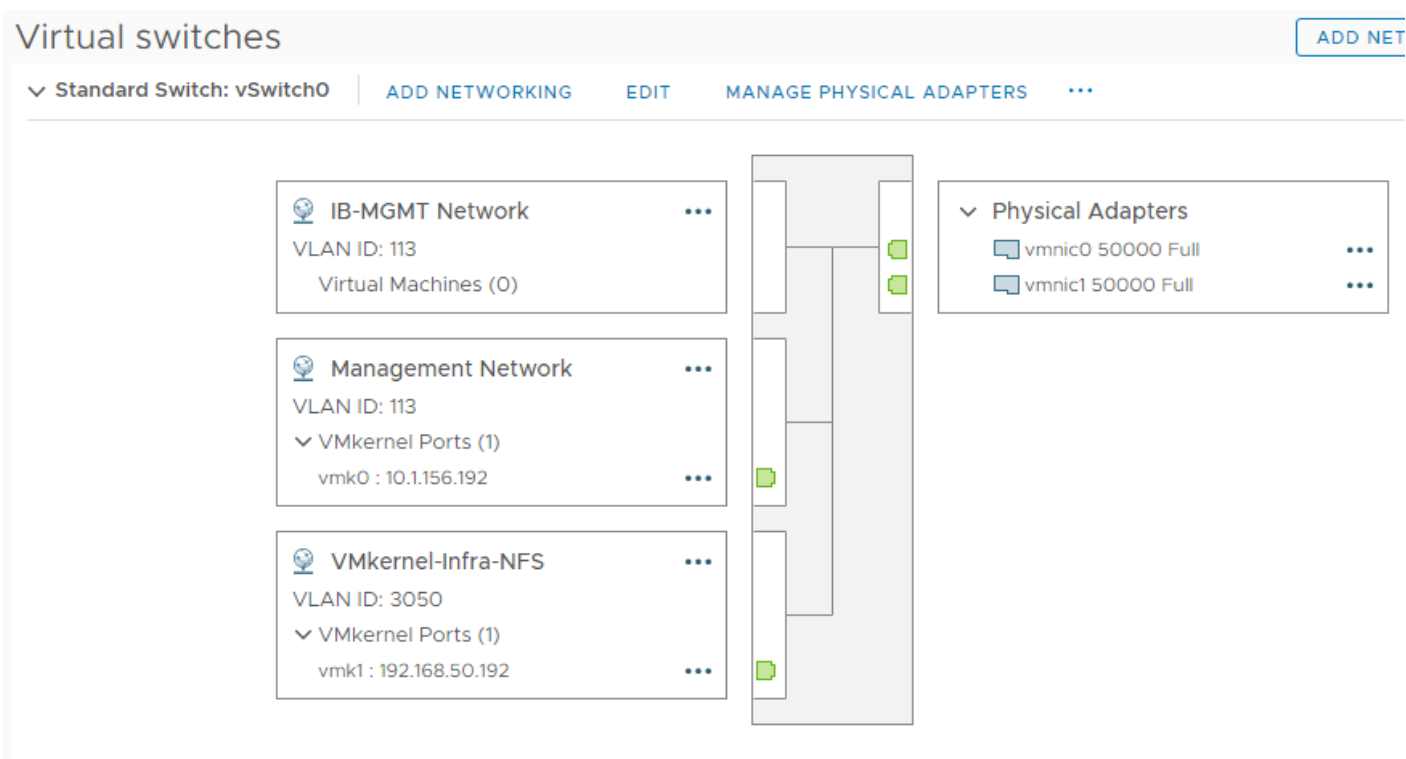
All		Properties	CDP	LLDP	RDMA
Adapter Name	Cisco Systems Inc Cisco VIC Ethernet NI				
Location	vmnic1				
Driver	PCI 0000:45:00.1				
	nenic				
Status					
Status	Connected				
Actual speed, Duplex	50 Gbit/s, Full Duplex				
Configured speed, Duplex	50 Gbit/s, Full Duplex				
Networks	0.0.0.1-255.255.255.254 (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:



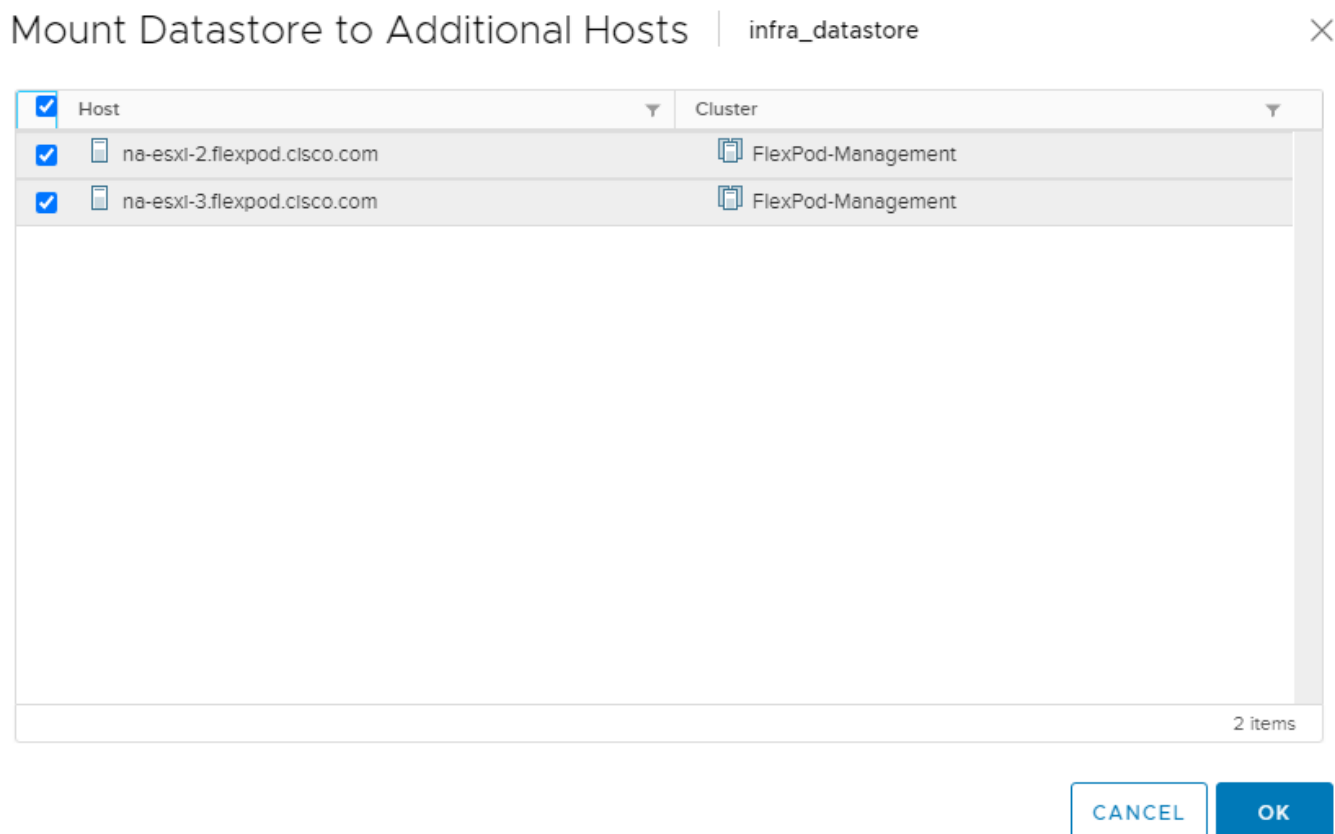
25. Repeat steps 1-24 for all hosts being added.

Mount Required Datastores for the 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. 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 for the 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. To the right of Manual Time Configuration, click EDIT.

5. Set the correct local time and click OK.
6. To the right of Network Time Protocol, click EDIT.
7. Choose the Enable checkbox.
8. Enter the two Nexus switch NTP IP addresses in the NTP servers box separated by a comma.
9. Click the Start NTP Service checkbox.
10. Use the drop-down list to choose Start and stop with host.

Edit Network Time Protocol | na-esxi-2.flexpod.cisco.com ✕

Enable ⓘ

NTP Servers	<input type="text" value="10.1156.11,10.1156.12"/>
Separate servers with commas, e.g. 10.31.21.2, fe00::2800	
NTP Service Status:	Stopped <input checked="" type="checkbox"/> Start NTP Service
NTP Service Startup Policy:	<input type="text" value="Start and stop with host"/> ▼

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

Configure ESXi Host Swap for the 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.

Edit System Swap Settings | na-esxi-2.flexpod.cisc... X

- Can use datastore:
- Can use host cache
- Can use datastore specified by host for swap files



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.

Change ESXi Power Management Policy for Cisco UCS C125 Hosts for the ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To change the ESXi power management policy for the AMD-base Cisco UCS C125 hosts, follow these steps:



Implementation of this policy is recommended in [Performance Tuning for Cisco UCS C125 Rack Server Nodes with AMD Processors](#) for maximum performance. If your organization has specific power policies, please set this policy accordingly.

1. In the list under Hardware, choose Overview. Scroll to the bottom and to the right of Power Management, choose EDIT POWER POLICY.
2. Choose High performance and click OK.

Edit Power Policy Settings

na-esxi-1.flexpod.cisco... X

High performance

Do not use any power management features

Balanced

Reduce energy consumption with minimal performance compromise

Low power

Reduce energy consumption at the risk of lower performance

Custom

User-defined power management policy

CANCEL

OK

Check ESXi Host Fibre Channel Pathing for the ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

For the 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

Mark as Perennially Reserved

Name	L...	Type	Capacity	Datasto...	Operational ...	Hardware Accelera...	Drive T...	Transp
Local ATA Disk (t10.ATA____Micron_5100_M...	0	disk	223.57 GB	Not Cons...	Attached	Not supported	Flash	Block
Local ATA Disk (t10.ATA____Micron_5100_M...	0	disk	223.57 GB	Not Cons...	Attached	Not supported	Flash	Block
NETAPP Fibre Channel Disk (naa.600a098038...	0	disk	32.00 GB	Not Cons...	Attached	Supported	Flash	Fibre

Copy All | 3 items

Properties | Paths | Partition Details

Enable | Disable

Runtime Name	Status	Target	Name	Preferred
vmhba0:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:16:6b:8b 20:01:d0:39:ea:1...	vmhba0:C0:T1:L0	
vmhba1:C0:T2:L0	Active	20:00:d0:39:ea:16:6b:8b 20:04:d0:39:ea:1...	vmhba1:C0:T2:L0	
vmhba1:C0:T1:L0	Active (I/O)	20:00:d0:39:ea:16:6b:8b 20:02:d0:39:ea:1...	vmhba1:C0:T1:L0	
vmhba0:C0:T2:L0	Active	20:00:d0:39:ea:16:6b:8b 20:03:d0:39:ea:1...	vmhba0:C0:T2:L0	

Add the ESXi Host(s) to the VMware Virtual Distributed Switch to the 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
na-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	--	--
na-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(s) to the ESXi Host to the 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.
 13. Optionally, repeat this process to add two more vMotion VMkernel ports.
 14. If NetApp VSC is installed, under Hosts and Clusters, right-click the host and click NetApp VSC > Set Recommended Values. Reboot the host.
 15. If this is an iSCSI-booted host, execute the instructions in the Appendix for an iSCSI-booted host being added in vCenter.
 16. Exit Maintenance Mode on each ESXi host in Maintenance Mode.

FlexPod Management Tools Setup

NetApp Virtual Storage Console 9.7.1 Deployment Procedure

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

Virtual Storage Console 9.7.1 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.7P2 or above:

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

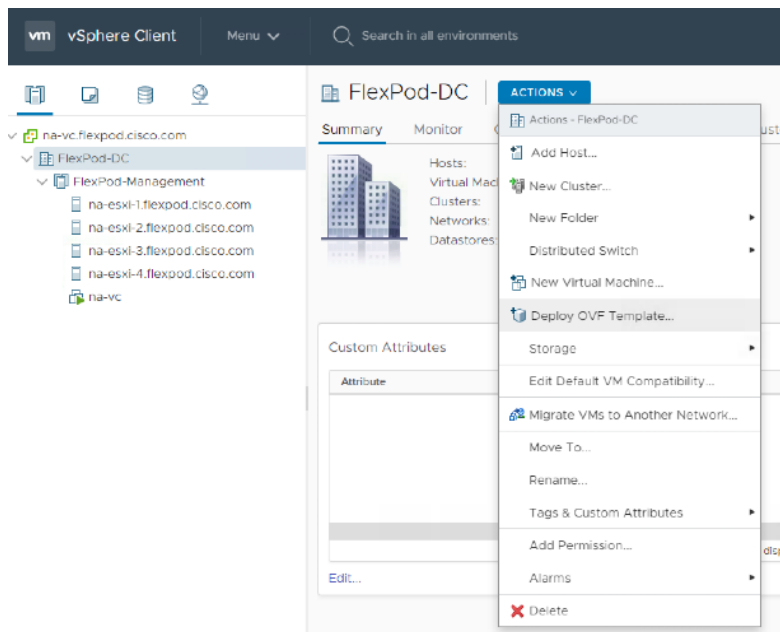
Port	Requirement
443 (HTTPS)	Secure communications between VMware vCenter Server and the storage systems
8143 (HTTPS)	VSC listens for secure communications
9083 (HTTPS)	VASA Provider uses this port to communicate with the vCenter Server and obtain TCP/IP settings

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

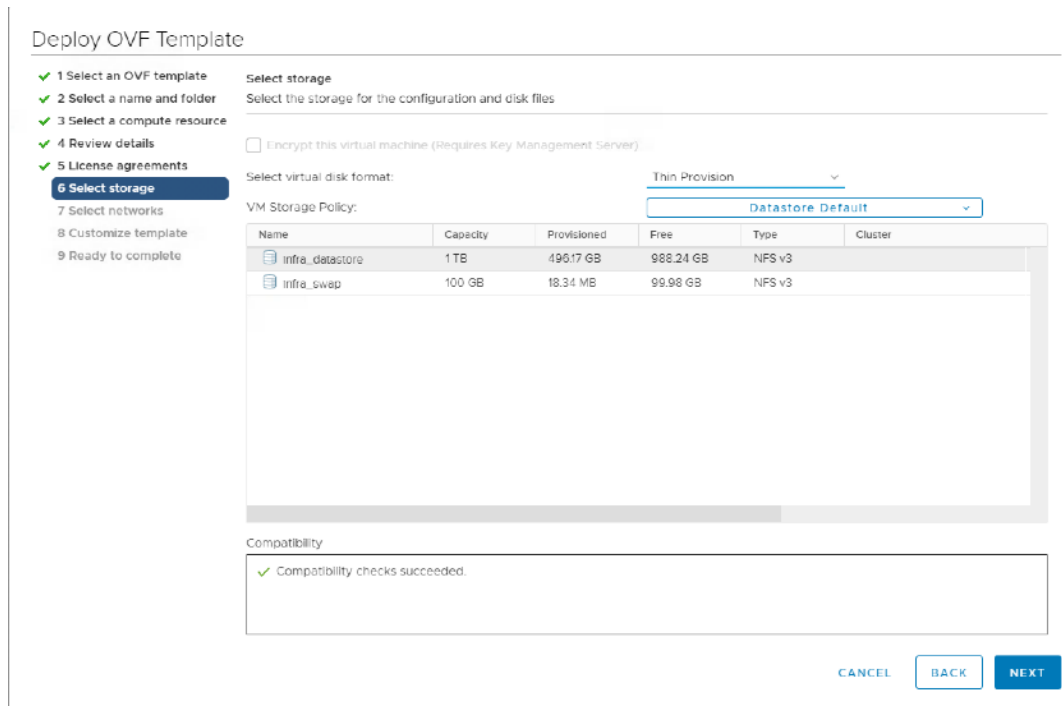
Install Virtual Storage Console 9.7.1

To install the VSC 9.7.1 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. Select ACTIONS for 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.



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

<ul style="list-style-type: none"> ✓ 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 <li style="background-color: #0070C0; color: white; padding: 2px;">8 Customize template 9 Ready to complete 	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr style="background-color: #D9E1F2;"> <th style="text-align: left; padding: 5px;">System Configuration</th> <th style="text-align: left; padding: 5px;">3 settings</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px;">Application User Password (*)</td> <td style="padding: 5px;"> Password to assign to the administrator account. Password <input style="width: 100%;" type="password" value="*****"/> Confirm Password <input style="width: 100%;" type="password" value="*****"/> </td> </tr> <tr> <td style="padding: 5px;">NTP Servers</td> <td style="padding: 5px;"> A comma-separated list of hostnames or IP addresses of NTP Servers. If left blank, VMware tools based time synchronization will be used. <input style="width: 100%;" type="text" value="192.168.156.1,10.1.156.4"/> </td> </tr> <tr> <td style="padding: 5px;">Maintenance User Password (*)</td> <td style="padding: 5px;"> Password to assign to maint user account. Password <input style="width: 100%;" type="password" value="*****"/> Confirm Password <input style="width: 100%;" type="password" value="*****"/> </td> </tr> <tr style="background-color: #D9E1F2;"> <th style="text-align: left; padding: 5px;">vCenter Registration Configuration</th> <th style="text-align: left; padding: 5px;">4 settings</th> </tr> <tr> <td style="padding: 5px;">vCenter Server Address (*)</td> <td style="padding: 5px;"> Specify the IP address/hostname of an existing vCenter to register to. <input style="width: 100%;" type="text" value="na-vc.flexpod.cisco.com"/> </td> </tr> <tr> <td style="padding: 5px;">Port (*)</td> <td style="padding: 5px;"> Specify the HTTPS port of an existing vCenter to register to. <input style="width: 100%;" type="text" value="443"/> </td> </tr> <tr> <td style="padding: 5px;">Username (*)</td> <td style="padding: 5px;"> Specify the username of an existing vCenter to register to. <input style="width: 100%;" type="text" value="administrator@vsphere.io"/> </td> </tr> </tbody> </table> <div style="text-align: right; margin-top: 10px;"> CANCEL BACK NEXT </div>	System Configuration	3 settings	Application User Password (*)	Password to assign to the administrator account. Password <input style="width: 100%;" type="password" value="*****"/> Confirm Password <input style="width: 100%;" type="password" value="*****"/>	NTP Servers	A comma-separated list of hostnames or IP addresses of NTP Servers. If left blank, VMware tools based time synchronization will be used. <input style="width: 100%;" type="text" value="192.168.156.1,10.1.156.4"/>	Maintenance User Password (*)	Password to assign to maint user account. Password <input style="width: 100%;" type="password" value="*****"/> Confirm Password <input style="width: 100%;" type="password" value="*****"/>	vCenter Registration Configuration	4 settings	vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to. <input style="width: 100%;" type="text" value="na-vc.flexpod.cisco.com"/>	Port (*)	Specify the HTTPS port of an existing vCenter to register to. <input style="width: 100%;" type="text" value="443"/>	Username (*)	Specify the username of an existing vCenter to register to. <input style="width: 100%;" type="text" value="administrator@vsphere.io"/>
System Configuration	3 settings																
Application User Password (*)	Password to assign to the administrator account. Password <input style="width: 100%;" type="password" value="*****"/> Confirm Password <input style="width: 100%;" type="password" value="*****"/>																
NTP Servers	A comma-separated list of hostnames or IP addresses of NTP Servers. If left blank, VMware tools based time synchronization will be used. <input style="width: 100%;" type="text" value="192.168.156.1,10.1.156.4"/>																
Maintenance User Password (*)	Password to assign to maint user account. Password <input style="width: 100%;" type="password" value="*****"/> Confirm Password <input style="width: 100%;" type="password" value="*****"/>																
vCenter Registration Configuration	4 settings																
vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to. <input style="width: 100%;" type="text" value="na-vc.flexpod.cisco.com"/>																
Port (*)	Specify the HTTPS port of an existing vCenter to register to. <input style="width: 100%;" type="text" value="443"/>																
Username (*)	Specify the username of an existing vCenter to register to. <input style="width: 100%;" type="text" value="administrator@vsphere.io"/>																

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

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.

Name	na-vsc
Template name	netapp-unified-virtual-appliance-for-vsc-vp-sra-9.7P2-5860-20200707_0700
Download size	1.6 GB
Size on disk	2.9 GB
Folder	FlexPod-DC
Resource	FlexPod-Management
Storage mapping	1
All disks	Datstore: infra_datstore, Format: Thin provision
Network mapping	1
net	IB-MGMT Network
IP allocation settings	
IP protocol	IPv4
IP allocation	Static - Manual

CANCEL

BACK

FINISH

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.

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

VMware Tools installation

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

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

OR

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

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

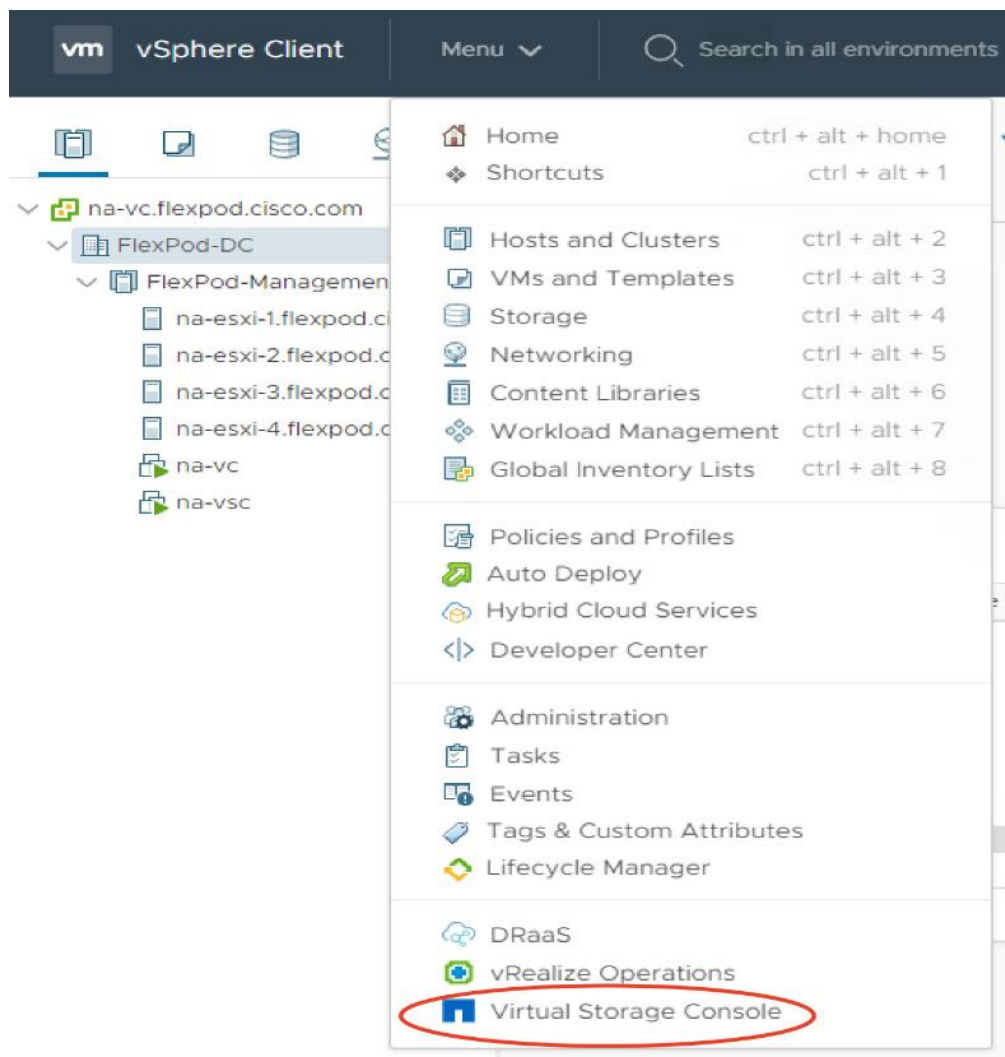
Once you click on mount, the installation process will automatically continue.
```

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. Refresh the Home Screen and confirm that the NetApp VSC is installed.



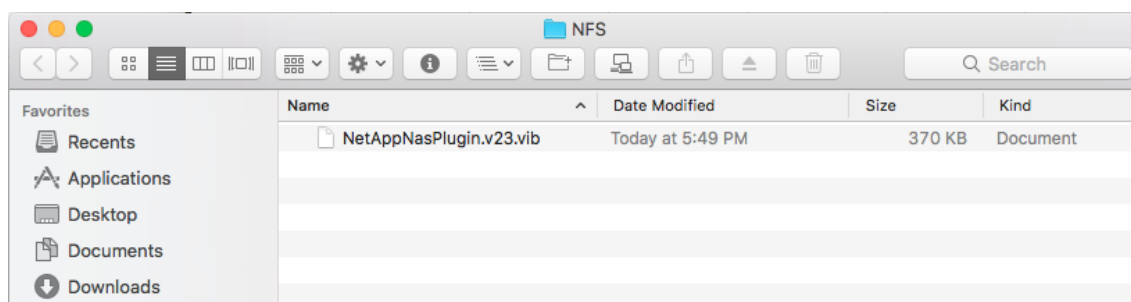
The NetApp VSC 9.7.1 vCenter plug-in is only available in the vSphere HTML5 Client and is not available in the vSphere Web Client.



Download the NetApp NFS Plug-in for VAAI

To download the NetApp NFS Plug-in for VAAI, 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 Plug-in for VAAI



The NFS Plug-in for VAAI 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 Plug-in for VAAI, 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 Plug-in for VAAI, 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) 9.7.1. To verify the VASA provider was enabled, follow these steps:

1. From the vSphere Client, click Menu > Virtual Storage Console.
2. Click Settings.
3. Click Manage Capabilities in the Administrative Settings tab.
4. In the Manage Capabilities dialog box if not enabled, click Enable VASA Provider slider.
5. Enter the IP address of the virtual appliance for VSC, VASA Provider, and VMware Storage Replication Adapter (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 interfaces (APIs) that enables vSphere vCenter to recognize the capabilities of storage arrays.



Enable Storage Replication Adapter (SRA)

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

Enter authentication details for VASA Provider and SRA server:

IP address or hostname: 10.1.156.201
Username: Administrator
Password: _____

CANCEL

APPLY

Discover and Add Storage Resources

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

1. Using the vSphere Web Client, log in to the vCenter Server as the FlexPod admin user. If the vSphere Web Client was previously opened, close the tab, and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.

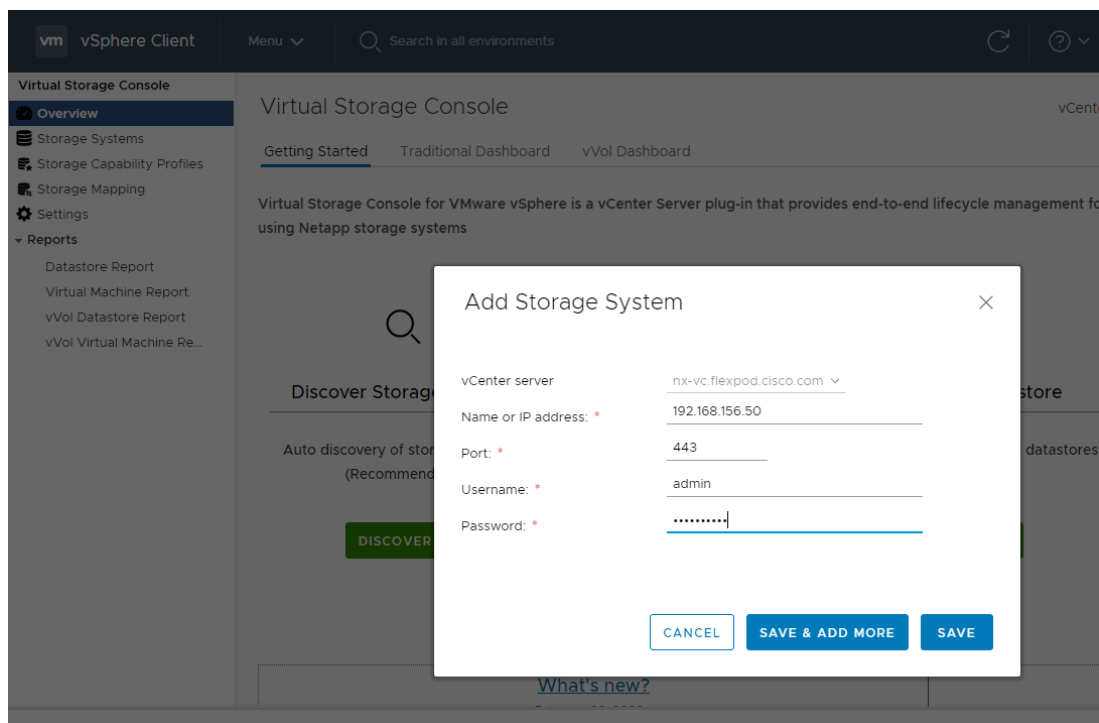


When using the cluster admin account, add storage from the cluster level.



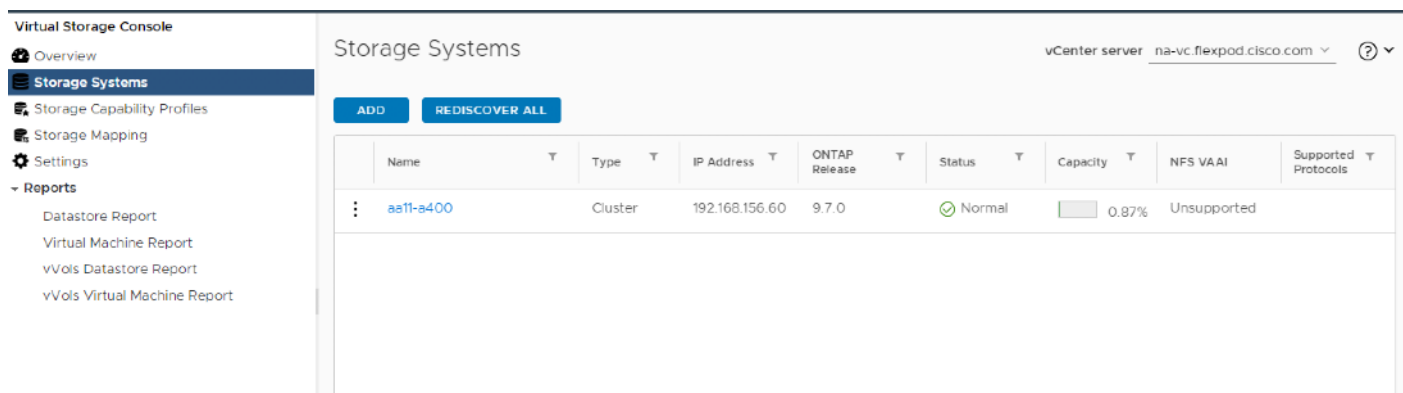
You can modify the storage credentials with the vsadmin account or another SVM level account with role-based access control (RBAC) privileges. Refer to the [ONTAP 9 Administrator Authentication and RBAC Power Guide](#) for additional information.

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.

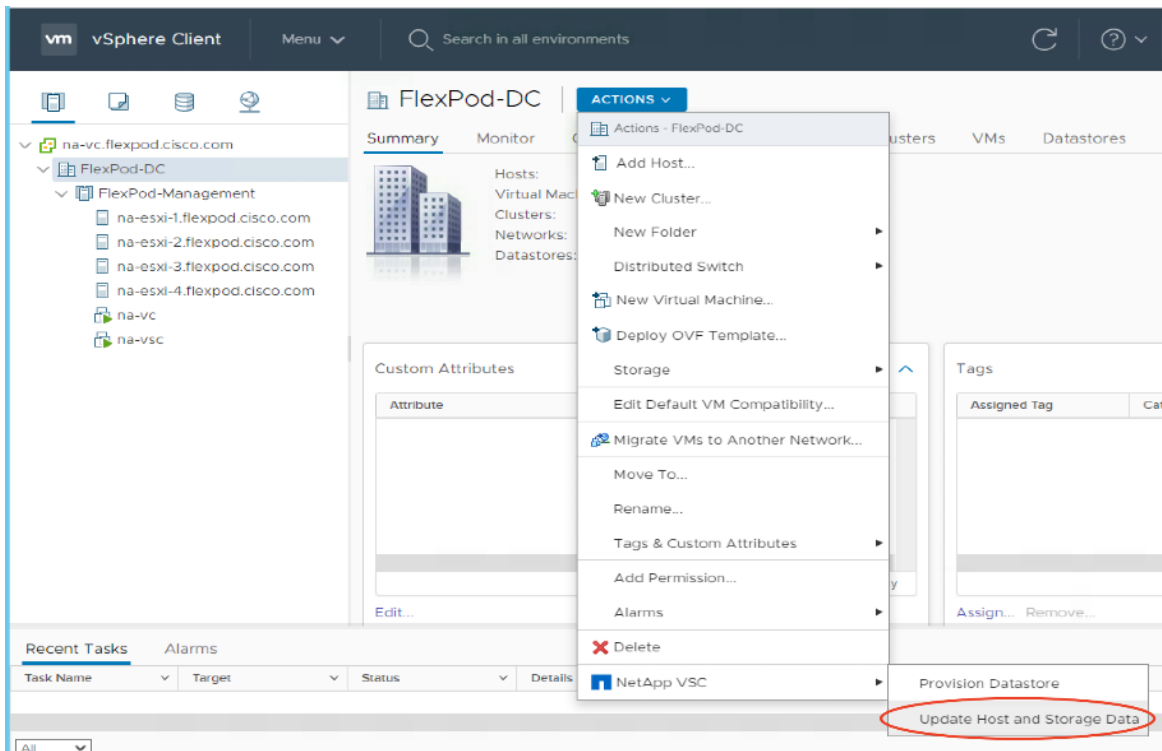


10. Wait for the Storage Systems to update. You might need to click Refresh to complete this update.

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, then 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 9.7.1 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 9.7.1. 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 9.7.1 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.

The screenshot displays the NetApp Virtual Storage Console interface. On the left, a navigation sidebar lists various sections: Overview, Storage Systems, Storage Capability Profiles (which is currently selected and highlighted in blue), Storage Mapping, Settings, and Reports. Under Reports, there are sub-items: Datastore Report, Virtual Machine Report, vVols Datastore Report, and vVols Virtual Machine Report. The main content area is titled 'Storage Capability Profiles' and includes a blue 'CREATE' button. Below this, a table of profiles is visible, with a context menu open over one of the entries. The context menu contains three options: 'Update', 'Clone' (which is circled in red), and 'Delete'.

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

Clone Storage Capability Profile

1 General

2 Platform

3 Performance

4 Storage attributes

5 Summary

General

Specify the name and description of storage capability profile. ?

Name: * AFF_No_Encrypt

Description: Cloned profile for Platinum level Service for AFF platform

CANCEL NEXT

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

Clone Storage Capability Profile

1 General

2 Platform

3 Performance

4 Storage attributes

5 Summary

Platform

Platform: All Flash FAS(AFF)

CANCEL BACK NEXT

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

Clone Storage Capability Profile

1 General

2 Platform

3 Performance

4 Storage attributes

5 Summary

Storage attributes

Deduplication: Yes

Compression: Yes

Space reserve: Thin

Encryption: No

Tiering policy (FabricPool): Any

CANCEL BACK NEXT

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

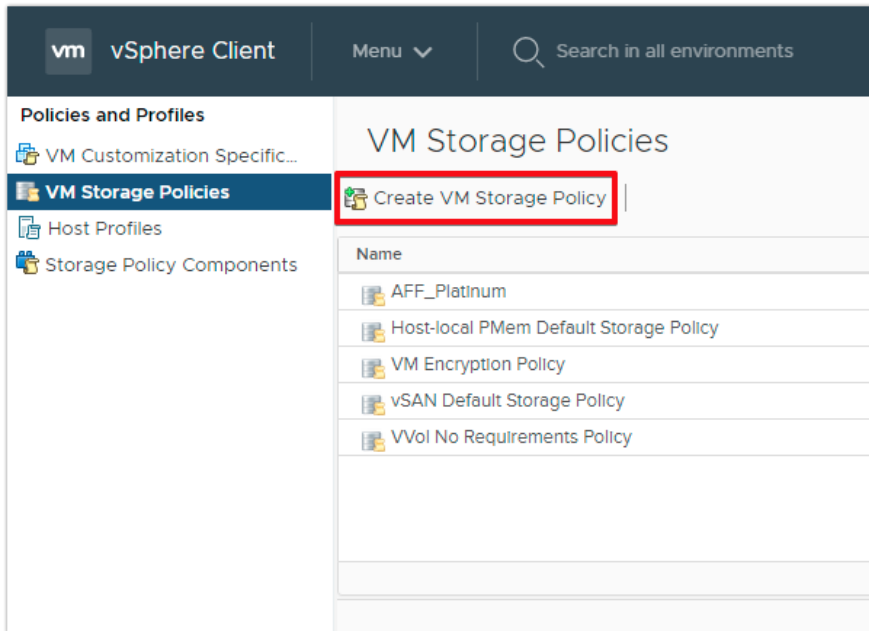


It is recommended to Clone the Storage Capability Profile if you wish to make any changes to the default profiles rather than editing the built-in profile.

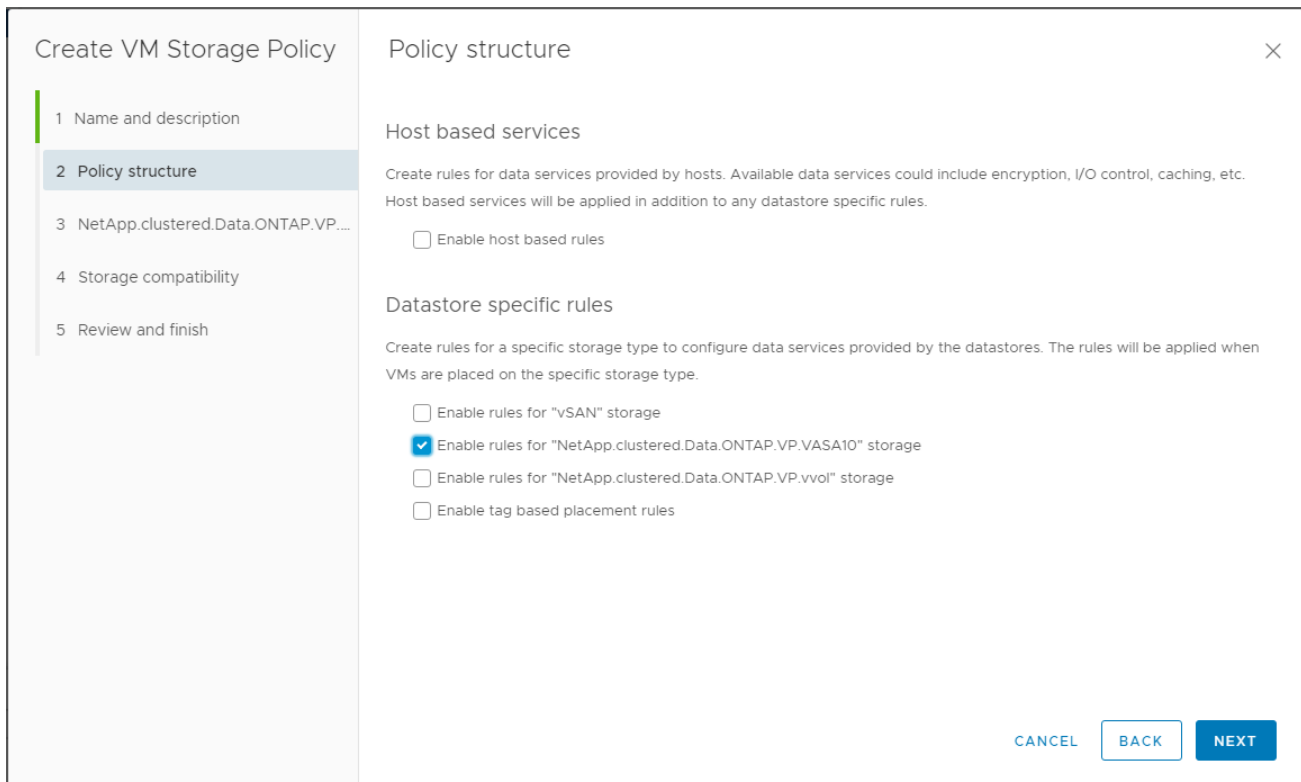
Create a VM Storage Policy

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.

The screenshot shows a wizard window titled "Create VM Storage Policy" with a sub-header "NetApp.clustered.Data.ONTAP.VP.VASA10 rules". On the left, a vertical list of steps is shown: 1 Name and description, 2 Policy structure, 3 NetApp.clustered.Data.ONTAP.VP... (highlighted), 4 Storage compatibility, and 5 Review and finish. The main area has two tabs: "Placement" (selected) and "Tags". Under "Placement", there is a field "SystemLabel.label" with a help icon and a dropdown menu showing "AFF_No_Encrypt". At the bottom right, there are three buttons: "CANCEL", "BACK", and "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.

New Datastore

General

Specify the details of the datastore to provision

Provisioning destination: FlexPod-DC BROWSE

Type: NFS VMFS vVol

Name: NX_NFS_DS_01

Size: 500 GB

Protocol: NFS 3 NFS 4.1

Use storage capability profile for provisioning

CANCEL 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

Storage system

Specify the storage capability profiles and the storage system you want to use.

Storage capability profile: AFF_No_Encrypt

Platform: All Flash FAS(AFF) Performance: None
 Compression: Yes Deduplication: Yes Tiering policy (FabricPool): Any
 Space reserve: Thin Encryption: No

Storage system: aa11-a400 (192.168.156.60)

Storage VM: Infra-SVM

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: aa11_a400_01_NVME_SSD_1 - (16667.25 GB Free)

Volumes: Automatically creates a new volume.

ⓘ Advance options are pre-selected for optimum results.

Advanced options ▾

Space reserve: Thin

CANCEL
BACK
NEXT

10. Review the Summary and click FINISH.

New Datastore

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

Summary

vCenter server:	na-vc.flexpod.cisco.com
Provisioning destination:	FlexPod-DC
Datastore name:	NX_NFS_DS_0
Datastore size:	500 GB
Datastore type:	NFS
Protocol:	NFS 3
Datastore cluster:	None
Storage capability profile:	AFF_No_Encrypt

Storage system details

Storage system:	aa11-a400
SVM:	Infra-SVM

Storage attributes

Aggregate:	aa11_a400_01_NVME_SSD_1
Space reserve:	Thin

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. Also, VSC Home page > Reports > Datastore Report should be listing the newly created datastore.

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.
3. Click Browse to choose the destination to provision the datastore.
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.

New Datastore

1 General

2 Storage system

3 Storage attributes

4 Summary

General

Provisioning destination: FlexPod-DC [BROWSE](#)

Type: NFS VMFS vVols

Name: NX_FC_DS_01

Size: 500 GB

Protocol: iSCSI FC / FCoE

Use storage capability profile for provisioning

Advanced options

File system: VMFS6

Datastore cluster: None

[CANCEL](#) [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.

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: AFF_No_Encrypt ▼

Platform: All Flash FAS(AFF)	Performance: None	
Compression: Yes	Deduplication: Yes	Tiering policy (FabricPool): Any
Space reserve: Thin	Encryption: No	

Storage system: aa11-a400 (192.168.156.60) ▼

Storage VM: Infra-SVM ▼

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: aa11_a400_01_NVME_SSD_1 - (16666.39 GB Free) ▼

Volumes: Automatically creates a new volume.

i Advance options are pre-selected for optimum results

Advanced options ▼

Space reserve: Thin ▼

CANCEL
BACK
NEXT

10. Review the Summary and click FINISH.

New Datastore

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

Summary

vCenter server: na-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 capability profile: AFF_No_Encrypt

Storage system details

Storage system: aa11-a400
SVM: infra-SVM

Storage attributes

Aggregate: aa11_a400_01_NVME_SSD_1
Space reserve: Thin

CANCEL BACK FINISH

15. 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. Also, VSC Home page > Reports > Datastore Report should be listing the newly created datastore.

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.

New Virtual Machine

✓ 1 Select a creation type
 ✓ 2 Select a name and folder
 ✓ 3 Select a compute resource
4 Select storage
 5 Select compatibility
 6 Select a guest OS
 7 Customize hardware
 8 Ready to complete

Select storage
Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
infra_datastore	1 TB	588.85 GB	975.41 GB	NFS v3	
infra_swap	100 GB	7.58 MB	99.99 GB	NFS v3	
NX_FC_DS_01	500.25 GB	1.41 GB	498.84 GB	VMFS 6	
NX_NFS_DS_0	500 GB	300 KB	500 GB	NFS v3	

Compatibility

✓ Compatibility checks succeeded.

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.

Virtual Volumes(vVols)

NetApp VASA Provider enables you to create and manage VMware virtual volumes (vVols). A vVols datastore consists of one or more FlexVol volumes within a storage container (also called "backing storage"). A virtual machine can be spread across one vVols datastore or multiple vVols datastores. All of the FlexVol volumes within the storage container must use the same protocol (NFS, iSCSI, or FCP) and the same SVMs.



Lab testing has shown that if a virtual machine (VM) has one or more disks in vVol datastores and the VM is migrated to another host, just at the end of the migration the VM can be stunned or frozen for 45 or more seconds.

Verify NDMP Vserver Scope Mode

To verify the NDMP Vserver scope mode, follow these steps:

1. View NDMP scope mode with the following command:

```
system services ndmp node-scope-mode status
NDMP node-scope-mode is enabled.
```

2. Disable NDMP node-scoped mode.

```
system services ndmp node-scope-mode off
NDMP node-scope-mode is disabled.
```

3. Enable NDMP services on the vserver.

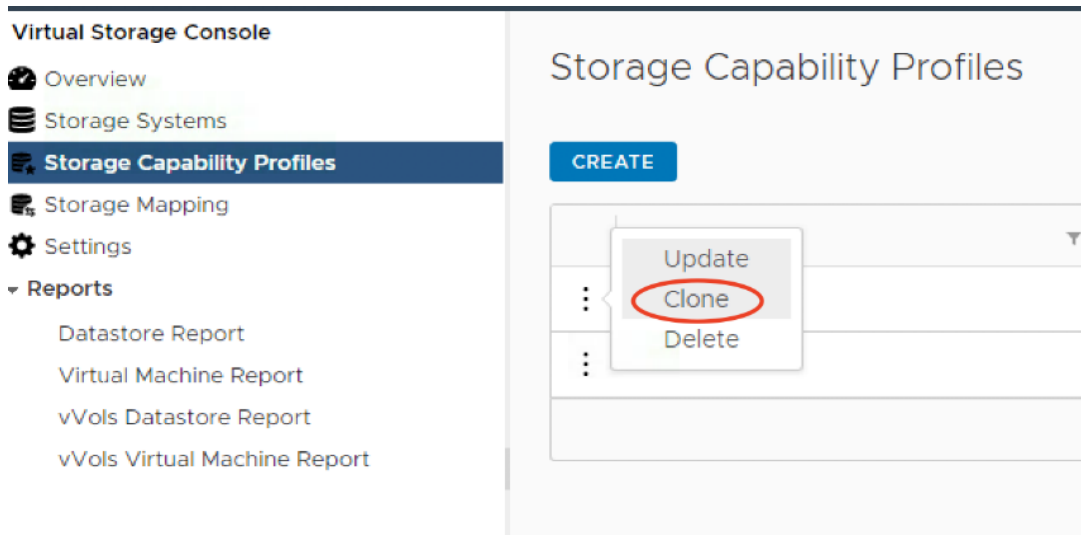
```
vserver services ndmp on -vserver Infra-SVM
```

Create the Storage Capability Profile

You can select one or more VASA Provider storage capability profiles for a vVols datastore. You can also specify a default storage capability profile for any vVols datastores that are automatically created in that storage container.

To create storage capability profile for the vVol datastore, 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.



The screenshot shows a configuration page for a 'Clone Storage Capability Profile'. On the left is a navigation pane with five items: '1 General' (selected), '2 Platform', '3 Performance', '4 Storage attributes', and '5 Summary'. The main area is titled 'General' and contains the following fields:

- A header instruction: 'Specify a name and description for the storage capability profile. ?'
- 'Name:' field with the value 'AFF_Cloned_Gold_No_encrypt'.
- 'Description:' field with the value 'Clone profile for Gold level Service for AFF platform'.
- 'CANCEL' and 'NEXT' buttons at the bottom right.

1. Choose All Flash FAS(AFF) for the storage platform and click Next.
2. Choose None to allow unlimited performance or set a the desired minimum and maximum IOPS for the QoS policy group. You can set the value for Max IOPS, which enables you to use the QoS functionality.

When applied for a virtual datastore, a QoS policy with " MAX IOPS" value is created for each data vVols.

When you select ONTAP Service Level, then the existing adaptive QoS policies of ONTAP are applied to a data vVols. You can select one of three service levels: Extreme, Performance, or Value. The ONTAP service level is applicable only to vVols datastores.

1. On the Storage attributes page, change the Encryption and Tiering policy to the desired settings and click NEXT.

Clone Storage Capability Profile

1 General

2 Platform

3 Performance

4 Storage attributes

5 Summary

Storage attributes

Deduplication: Yes

Compression: Yes

Space reserve: Thin

Encryption: No

Tiering policy (FabricPool): Any

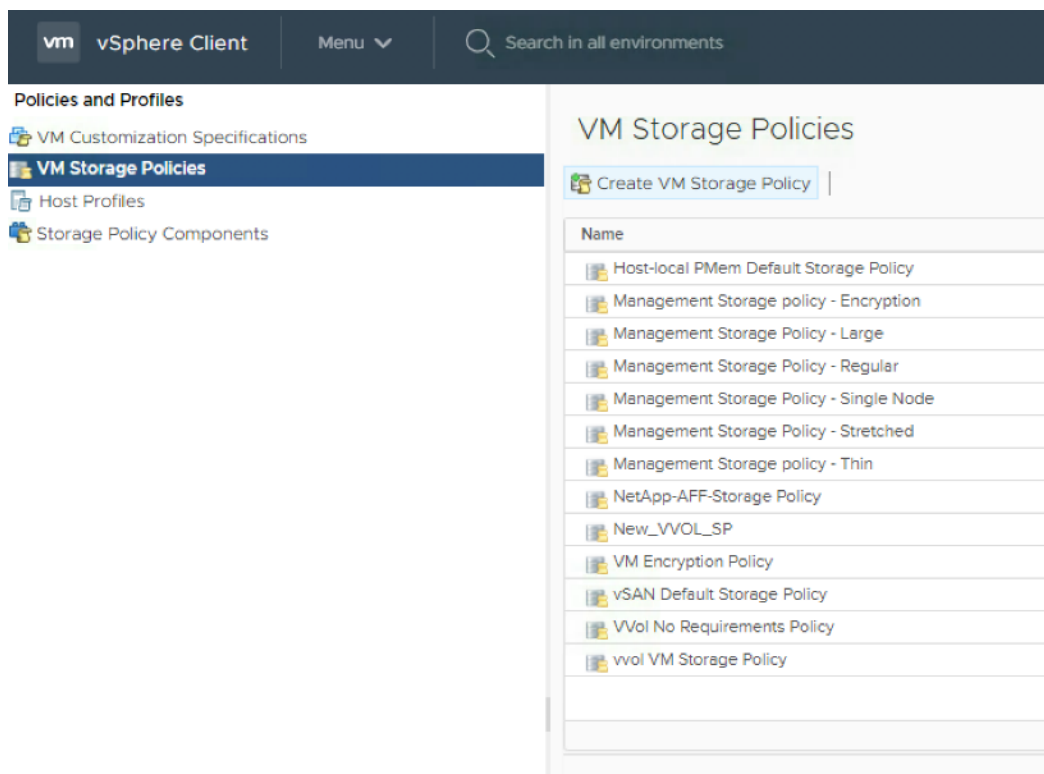
CANCEL BACK NEXT

2. Review the summary page and choose FINISH to create the storage capability 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. Click Create VM Storage Policy.
3. Create a new name for the VM storage Policy and enter a description and click NEXT.
4. Choose Enable rules for NetApp.clustered.Data.ONTAP.VP.VASA.10 storage and NetApp.clustered.Data.ONTAP.VP.vvol storage, located under the Datastore specific rules section and click NEXT.

Create VM Storage Policy

- 1 Name and description
- 2 Policy structure
- 3 NetApp.clustered.Data.ONTAP.VP...
- 4 NetApp.clustered.Data.ONTAP.VP...
- 5 Storage compatibility
- 6 Review and finish

Policy structure

Host based services

Create rules for data services provided by hosts. Available data services could include encryption, I/O control, caching, etc. Host based services will be applied in addition to any datastore specific rules.

Enable host based rules

Datastore specific rules

Create rules for a specific storage type to configure data services provided by the datastores. The rules will be applied when VMs are placed on the specific storage type.

Enable rules for "vSAN" storage

Enable rules for "NetApp.clustered.Data.ONTAP.VP.VASA10" storage

Enable rules for "NetApp.clustered.Data.ONTAP.VP.vvol" storage

Enable tag based placement rules

CANCEL
BACK
NEXT

5. On the Placement tab for VP.VASA and VP.vvol storage rules select the SCP created in the previous step.

Create VM Storage Policy

- 1 Name and description
- 2 Policy structure
- 3 NetApp.clustered.Data.ONTAP.VP...
- 4 NetApp.clustered.Data.ONTAP.VP...
- 5 Storage compatibility
- 6 Review and finish

NetApp.clustered.Data.ONTAP.VP.VASA10 rules

Placement	Tags
SystemLabelLabel ⓘ	AFF_Cloned_Gold_No_encrypt

CANCEL
BACK
NEXT

Create VM Storage Policy

NetApp.clustered.Data.ONTAP.VP.vvol rules

- 1 Name and description
- 2 Policy structure
- 3 NetApp.clustered.Data.ONTAP.VP...
- 4 NetApp.clustered.Data.ONTAP.VP...
- 5 Storage compatibility
- 6 Review and finish

Placement Tags

ProfileName ⓘ AFF_Cloned_Gold_No_encrypt ▼

CANCEL BACK NEXT

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

7. Review the Policy Summary and click Finish.

Provision a vVols Datastore

To provision the vVols datastore over NFS protocol, 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 vVols and Enter the datastore name.
5. Select NFS for protocol and click Next.

New Datastore

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

General

Specify the details of the datastore to provision. ?

Provisioning destination: FlexPod-Management BROWSE

Type: NFS VMFS vVols

Name: vvol_DS1

Description:

Provision vVOI Datastore for NFS

Protocol: NFS iSCSI FC / FCoE

6. Select the Storage capability profile created earlier for vVols.
7. Select the NFS storage server and the NetApp Storage SVM where the vVols needs to be created and 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 profiles:

- Bronze
- Custom profiles**
- AFF_No_Encrypt
- Gold
- AFF_Cloned_Gold_No_encrypt

Storage system: aa11-a400 (192.168.156.60) ▼

Storage VM: Infra-SVM ▼

8. Create new vVols or select existing vVols.



You can create multiple vVols for a datastore.

New Datastore


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

Storage attributes

Specify the storage details for provisioning the datastore.

Volumes: Create new volumes Select volumes

Create new volumes

Name	Size	Storage Capability Profile	Aggregate	
 FlexVol volumes are not added.				
Name	Size(GB)	Storage capability profile	Aggregates	Space reserve
<u>vvol1_DS1_01</u>	<u>10</u>	<u>AFF_Cloned_Gold_No_er</u> ▾	<u>aa11_a400_01_NVME_SSI</u> ▾	<u>Thin</u>

Auto Grow

New Datastore

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

Storage attributes

FlexVol volumes are not added.

Name	Size(GB)	Storage capability profile	Aggregates	Space reserve
<u>vvol1_DS1_01</u>	<u>10</u>	<u>AFF_Cloned_Gold_No_er</u> ▾	<u>aa11_a400_01_NVME_SSI</u> ▾	<u>Thin</u>

Auto Grow

Grow Grow/Shrink

Maximum Size(GB): 12

→ **ADD**

9. Check the storage capability profile and click NEXT.

New Datastore

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

Storage attributes

Name	Size(GB)	Storage capability profile	Aggregates	Space reserve
vvol1_DS1_01	10 GB	AFF_Cloned_Gold_No_encrypt	aa11_a400_01_NVME_SSD_1	Thin

1 - 1 of 1 item

Auto Grow

Grow
 Grow/Shrink

Maximum Size(GB): _____

ADD

Default storage capability profile: AFF_Cloned_Gold_No_encrypt ▼

CANCEL BACK **NEXT**

10. Review all the fields on the summary page and click Finish.

New Datastore

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

Summary

Datastore name: vvol_DS1
Datastore type: vVols
Protocol: NFS
Storage capability profile: AFF_Cloned_Gold_No_encrypt

Storage system details

Storage system: aa11-a400
SVM: Infra-SVM

Storage attributes

New FlexVol Name	New FlexVol Size	Aggregate	Storage Capability Profile
vvol1_DS1_01	10 GB	aa11_a400_01_NVME_SSD_1	AFF_Cloned_Gold_No_encrypt

Click 'Finish' to provision this datastore.

CANCEL BACK **FINISH**

11. Verify in the vVols Datastore report the vVols is mounted correctly, go to VSC->Reports-> vVols Datastore Report.

vVols Datastore Report

EXPORT TO CSV

Name	Total Space	Free Space	Used Space	Space Utilized (%)	Available Space (%)	IOPS	Latency
exp_vvol	100.00 GB	99.99 GB	40.00 MB	0%	99%	0	0 ms
FC_VVOL_DS	40.00 GB	39.99 GB	8.00 MB	0%	99%	0	0 ms
vvol_mk_new_sp	100.00 GB	100.00 GB	2.00 MB	0%	99%	0	0 ms
vvol_iscsi_datastore	10.00 GB	10.00 GB	5.00 MB	0%	99%	0	0 ms
vvol_ds1	10.00 GB	10.00 GB	0 B	0%	100%	0	0 ms
sree_iscsi_ds	30.00 GB	29.99 GB	8.00 MB	0%	99%	0	0 ms
NFS_VVOL	50.00 GB	50.00 GB	1.00 MB	0%	99%	0	0 ms
sree_vvol	25.00 GB	25.00 GB	1.00 MB	0%	99%	0	0 ms

Items per page: 10 | 1 - 8 of 8 items



To provision vVols for FC or ISCSI protocol, select it in the General tab and provide protocol-specific storage attributes in the Storage Attributes Inputs to create vVols successfully.

Update a vVols Datastore

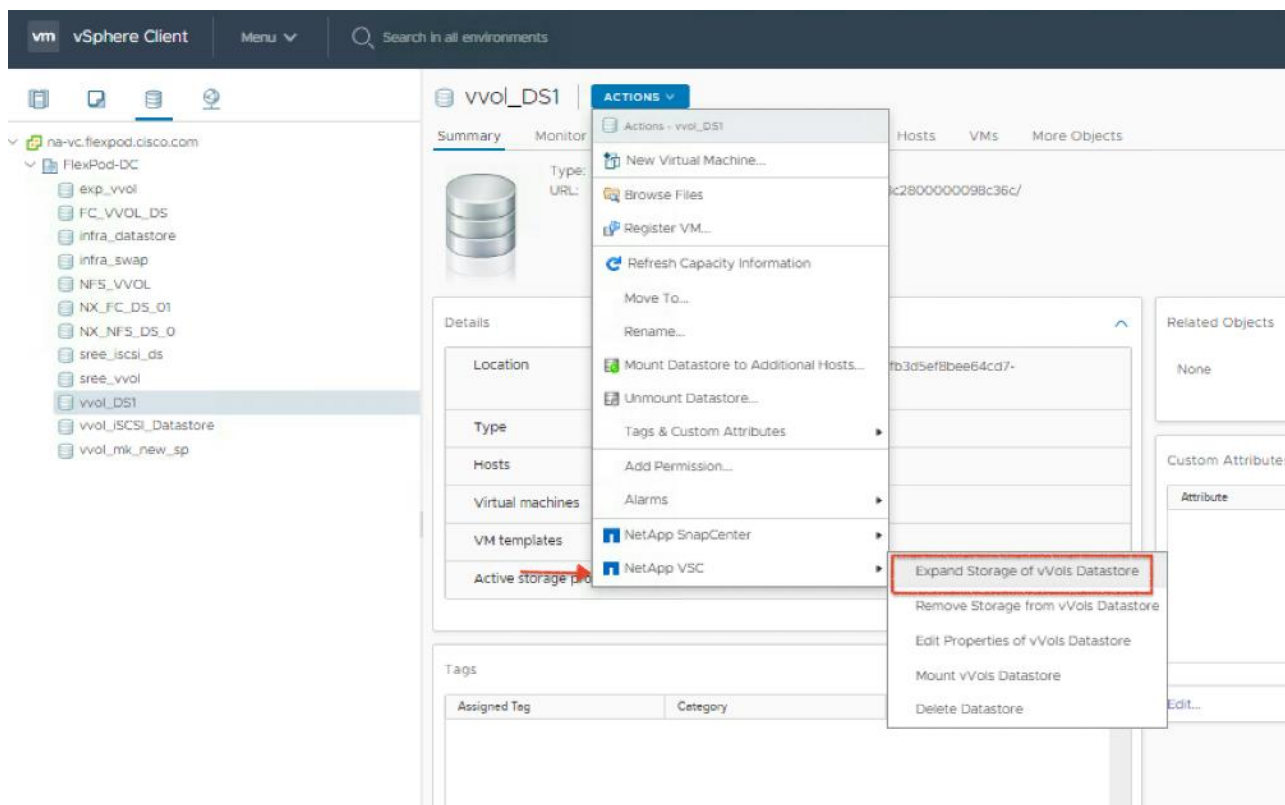
The following actions can be performed on a vVols Datastore

- Expand Storage on a vVols Datastore
- Remove Storage on a vVols Datastore
- Edit Properties of vVols Datastore
- Mount vVols Datastore
- Delete vVols Datastore

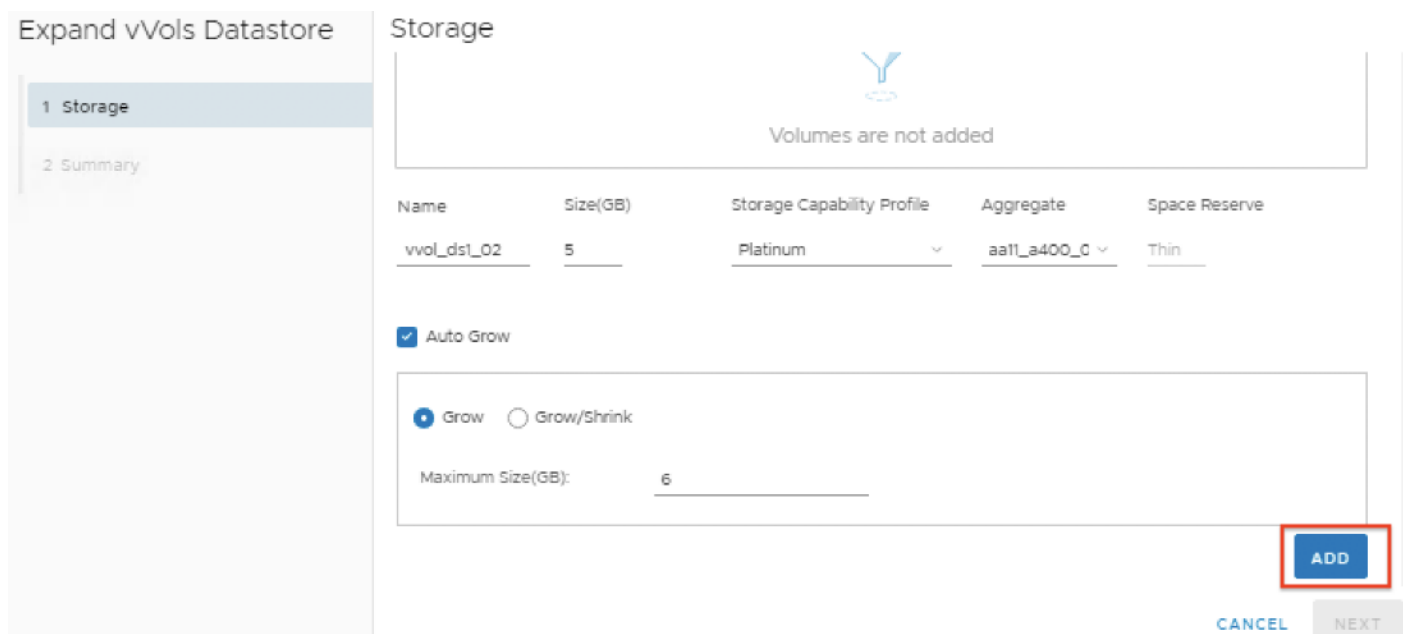
Expand Storage on a vVol Datastore

To expand storage on a vVol datastore, follow these steps:

1. In the Storage tab, click the vVols datastore to expand -> Actions-> NetApp VSC -> Expand Storage on a vVols Datastore.



2. Provide the storage attributes. Create new vVols or select existing vVols. Click Add.



3. Review the details in the Summary page and click Finish.

Expand vVols Datastore

- 1 Storage
- 2 Summary

Summary

General

vCenter server: na-vc.flexpod.cisco.com
 Protocol: nfs
 Storage system: aa11-a400
 SVM: infra-SVM

Created new volumes:

Name	Size	Storage Capability Profile	Aggregate
vvol_ds1_02	5 GB	Platinum	aa11_a400_01_NVME_SS...

CANCEL BACK FINISH

4. Verify that the size of the vVols datastore has been expanded successfully.

Virtual Storage Console

- Overview
- Storage Systems
- Storage Capability Profiles
- Storage Mapping
- Settings
- Reports
 - Datastore Report
 - Virtual Machine Report
 - vVols Datastore Report**
 - vVols Virtual Machine Report

vVols Datastore Report

vCenter server: na-vc.flexpod.cisco.com
 Last refreshad: 09/02/2020 03:25:12

EXPORT TO CSV

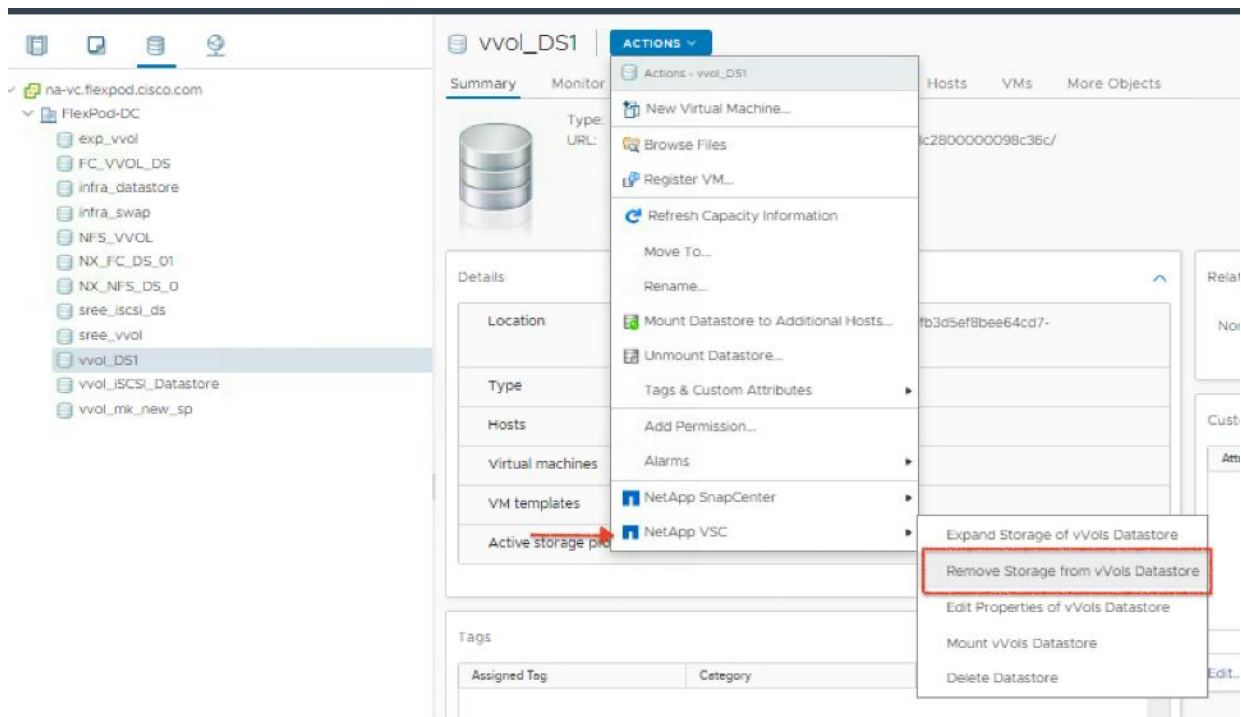
Name	Total Space	Free Space	Used Space	Space Utilized (%)	Available Space (%)	IOPS	Latency
exp_vvol	100.00 GB	99.96 GB	40.00 MB	0%	99%	0	0 ms
PC_VVOL_DS	40.00 GB	39.99 GB	8.00 MB	0%	99%	0	0 ms
vvol_mk_new_sp	100.00 GB	100.00 GB	2.00 MB	0%	99%	0	0 ms
vvol_SCSI_Datastore	10.00 GB	10.00 GB	5.00 MB	0%	99%	0	0 ms
vvol_DS1	15.00 GB	15.00 GB	1.00 MB	0%	99%	0	0 ms
stree_scsi_ds	30.00 GB	29.99 GB	8.00 MB	0%	99%	0	0 ms
NFS_VVOL	50.00 GB	50.00 GB	1.00 MB	0%	99%	0	0 ms
stree_vvol	25.00 GB	25.00 GB	1.00 MB	0%	99%	0	0 ms

Items per page: 10 1 - 8 of 8 items

Remove Storage from a vVol Datastore

To remove storage from a vVol datastore, follow these steps:

1. In the Storage tab, click the vVol datastore to remove storage -> Actions-> NetApp VSC -> Remove Storage from vVols Datastore.



2. Select the vVols within the datastore to remove and click Remove.

Remove Storage from vvol_DS1

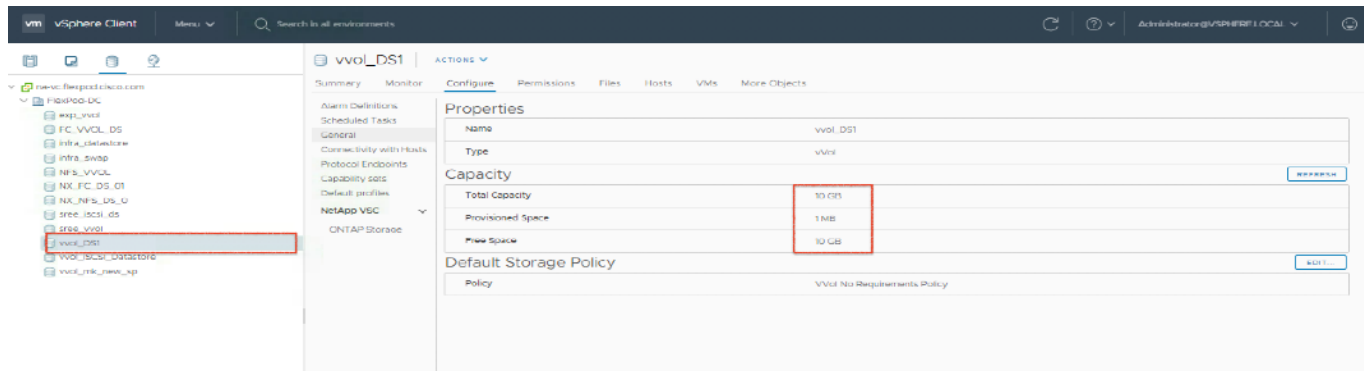
vCenter server: na-vc.flexpod.cisco.com

Select the FlexVol volumes that you want to remove from the vVols datastore.

<input type="checkbox"/>	FlexVol Name	FlexVol Size	Storage Capability Profile	Aggregate Name	No. of vVols
<input type="checkbox"/>	vvol_ds1_02	5.00 GB	Platinum	aa11_a400_01_NVME_SSD_1	0
<input checked="" type="checkbox"/>	vvol_ds1_01	10.00 GB	AFF_Gold_no_encrypt	aa11_a400_01_NVME_SSD_1	1

1 - 2 of 2 items

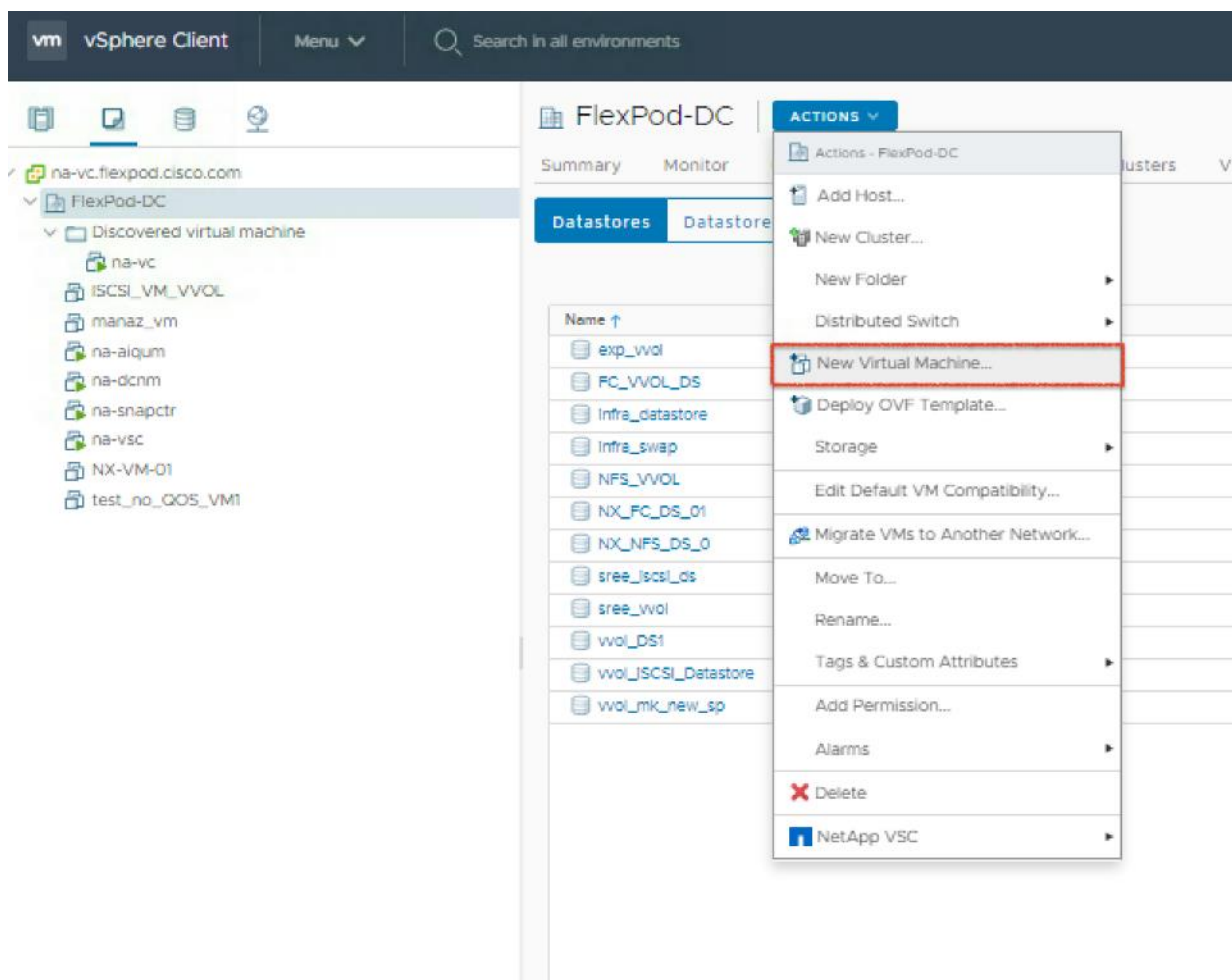
3. Verify the size of the datastore to validate the successful removal of storage from the vVols Datastore.



Create a Virtual Machine on a vVols Datastore with Assigned Virtual Machine Storage Policy

To provision a virtual machine on a vVols datastore, follow these steps:

1. Navigate to vSphere Client-> VMs and Templates-> Actions_> New Virtual Machine.



2. Enter the name for the VM and click the datacenter.

New Virtual Machine

✓ 1 Select a creation type

2 Select a name and folder

3 Select a compute resource

4 Select storage

5 Select compatibility

6 Select a guest OS

7 Customize hardware

8 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name: IO_VM1

Select a location for the virtual machine.

▼ na-vc.flexpod.cisco.com

> FlexPod-DC

CANCEL

BACK

NEXT

3. Choose the FlexPod-Management Data compute Resource.

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Select storage**
- 5 Select compatibility
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select storage

Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy:

Datastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
exp_vvol	100 GB	41 MB	99.96 GB	VVol	
FC_VVOL_DS	40 GB	8 MB	39.99 GB	VVol	
Infra_datastore	1 TB	1.35 TB	625.24 GB	NFS v3	
Infra_swap	200 GB	32.85 MB	199.97 GB	NFS v3	
NFS_VVOL	50 GB	1 MB	50 GB	VVol	
NX_FC_DS_01	500.25 GB	7.31 GB	492.94 GB	VMFS 6	
NX_NFS_DS_0	500 GB	19.6 GB	500 GB	NFS v3	
sree_iscsi_ds	30 GB	3.6 GB	29.99 GB	VVol	
sree_vvol	25 GB	1 MB	25 GB	VVol	
vvol_DS1	10 GB	0 B	10 GB	VVol	
vvol_ISCSI_Datastore	10 GB	5 MB	10 GB	VVol	
vvol_mk_new_so	100 GB	3.6 GB	100 GB	VVol	

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- 5 Select compatibility**
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select compatibility

Select compatibility for this virtual machine depending on the hosts in your environment

The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.

Compatible with: ESXi 7.0 and later ⓘ

This virtual machine uses hardware version 17, which provides the best performance and latest features available in ESXi 7.0

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- ✓ 6 Select a guest OS
- 7 Customize hardware**
- 8 Ready to complete

Customize hardware

Configure the virtual machine hardware

Virtual Hardware VM Options

> CPU	1	
> Memory	1	GB
> New Hard disk *	16	GB
> New SCSI controller *	VMware Paravirtual	
> New Network *	IB-MGMT Network	<input checked="" type="checkbox"/> Connect...
> New CD/DVD Drive *	Client Device	<input type="checkbox"/> Connect...
> Video card *	Specify custom settings	
> Security Devices	Not Configured	
VMCI device		
> New SATA Controller	New SATA Controller	
> Other	Additional Hardware	

Compatibility: ESXi 7.0 and later (VM version 17)

CANCEL BACK NEXT

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- ✓ 6 Select a guest OS
- ✓ 7 Customize hardware
- 8 Ready to complete**

Ready to complete

Click Finish to start creation.

Virtual machine name	IO_VMI
Folder	FlexPod-DC
Cluster	FlexPod-Management
Datastore	NFS_VVOL
Guest OS name	SUSE Linux Enterprise 15 (64-bit)
Virtualization Based Security	Disabled
CPUs	1
Memory	1 GB
NICs	1
NIC 1 network	IB-MGMT Network
NIC 1 type	VMXNET 3
SCSI controller 1	VMware Paravirtual
Create hard disk 1	New virtual disk
Capacity	16 GB
Datastore	NFS_VVOL

CANCEL BACK FINISH

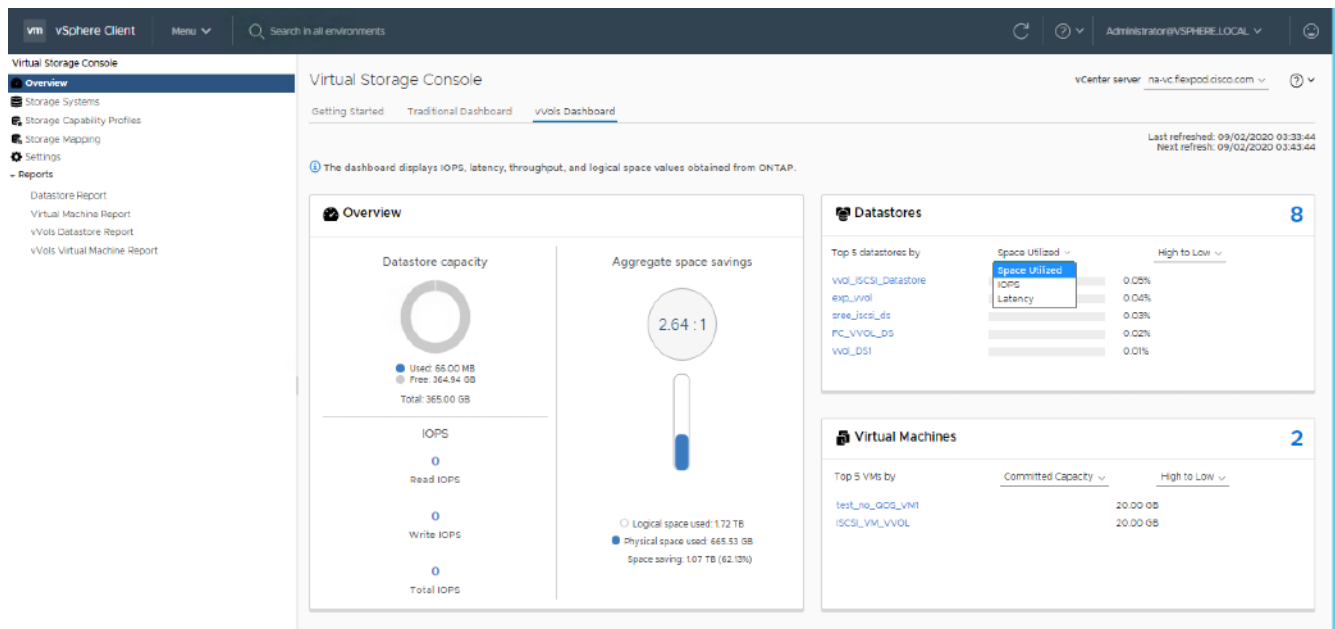
4. Validate that the VM is created successfully.

Monitor a vVols Datastore

To monitor a vVols datastore, follow these steps:

1. The vVols dashboard in the VSC Overview tab provides:

- Overview
- Top 5 datastores by Space Utilization, IOPs and Latency in ascending and descending order.
- A list of VMs on the vVols Datastore.



2. The vVol Datastore Report can be retrieved from VSC-> Reports-> vVol Datastore Report.

The screenshot shows the vSphere Client interface with the Virtual Storage Console (VSC) vVols Datastore Report. The report is a table with the following columns: Name, Total Space, Free Space, Used Space, Space Utilized (%), Available Space (%), IOPS, and Latency. The data is as follows:

Name	Total Space	Free Space	Used Space	Space Utilized (%)	Available Space (%)	IOPS	Latency
exp_vvol	100.00 GB	99.99 GB	41.00 MB	0%	99%	0	0 ms
FC_VVOL_DS	40.00 GB	39.99 GB	8.00 MB	0%	99%	0	0 ms
vVol_mk_new_sp	100.00 GB	100.00 GB	2.00 MB	0%	99%	0	0 ms
vVol_ISCSI_Datastore	10.00 GB	10.00 GB	5.00 MB	0%	99%	0	0 ms
vVol_DS1	10.00 GB	10.00 GB	1.00 MB	0%	99%	0	0 ms
tree_iscsi_ds	30.00 GB	29.99 GB	8.00 MB	0%	99%	0	0 ms
NFS_VVOL	50.00 GB	50.00 GB	1.00 MB	0%	99%	0	0 ms
tree_vvol	25.00 GB	25.00 GB	1.00 MB	0%	99%	0	0 ms

3. The vVol Virtual Machine Report can be retrieved from VSC-> Reports-> vVol Virtual Machine Report.

Virtual Storage Console

- Overview
- Storage Systems
- Storage Capability Profiles
- Storage Mapping
- Settings
- Reports
 - Datastore Report
 - Virtual Machine Report
 - vVols Datastore Report
 - vVols Virtual Machine Report**

vVols Virtual Machine Report

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

Last refreshed: 09/02/2020 04:02:35

EXPORT TO CSV

Name	Committed Capacity	Latency	Uptime	Throughput	Logical Base	Host	Power State
test_na_005_VMI	20.00 GB	0 ms	N/A	0 Bytes/s	0 B	na-esxi-2.flexpod.cisco.com	Powered Off
ISCS_VM_VVOL	20.00 GB	0 ms	N/A	0 Bytes/s	0 B	na-esxi-2.flexpod.cisco.com	Powered Off
IO_VMI	20.00 GB	0 ms	N/A	0 Bytes/s	0 B	na-esxi-1.flexpod.cisco.com	Powered Off

Items per page: 10 | 1 - 3 of 3 items

NetApp SnapCenter 4.3.1

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 4.3.1

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
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

Product	License Requirements
	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 4.3.1

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.

Deploy OVF Template

✓ 1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name: na-snapstr

Select a location for the virtual machine.

na-vc.flexpod.cisco.com

- > FlexPod-DC

CANCEL

BACK

NEXT

5. On the Select a compute resource page, 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:

Dalastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
infra_datastore	1 TB	805.75 GB	932.12 GB	NFS v3	
infra_swap	100 GB	12.51 MB	99.99 GB	NFS v3	
NX_FC_DG_01	500.25 GB	1.41 GB	498.84 GB	VMFS 6	
NX_NFS_VK_01	500 GB	372 KB	500 GB	NFS v3	

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.

Deploy OVF Template

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

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	IB-MGMT Network

1 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL BACK NEXT

10. On the Customize template page, do the following:

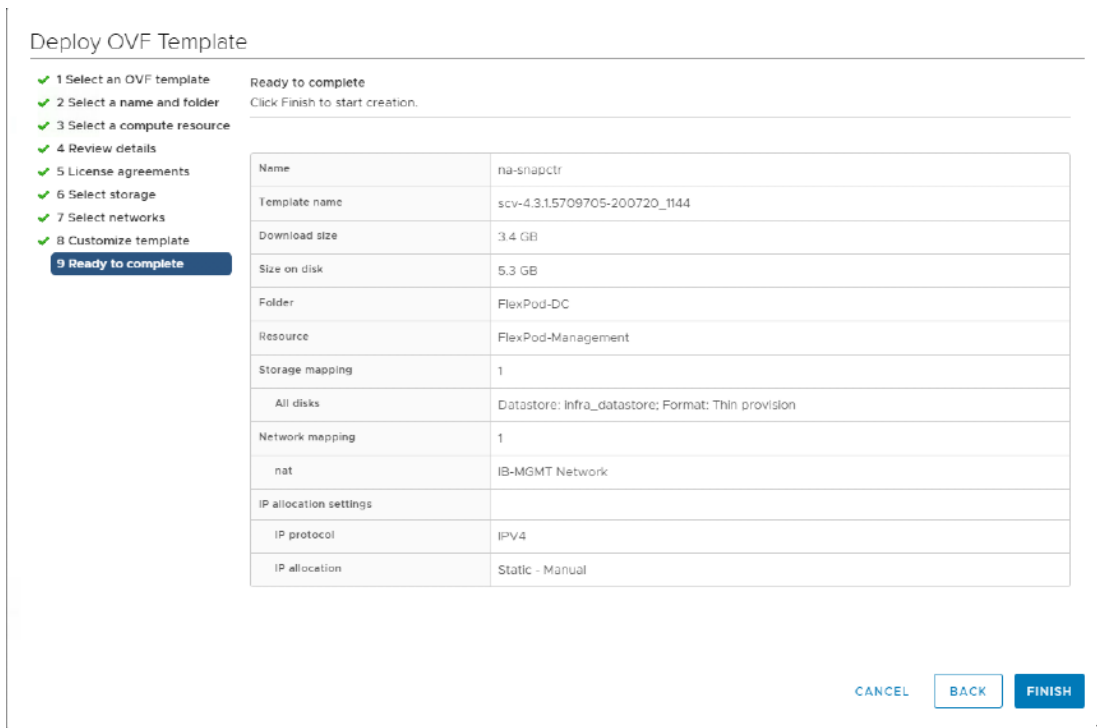
1. In Register to existing vCenter, enter the vCenter credentials.
2. In Create SnapCenter Plug-in for VMware vSphere credentials, enter the SnapCenter Plug-in for VMware vSphere credentials.
3. In Create SCV credentials, create a username and password for the SCV maintenance user.
4. In Setup Network Properties, enter the network information.
5. In Setup Date and Time, choose the time zone where the vCenter is located.

Deploy OVF Template

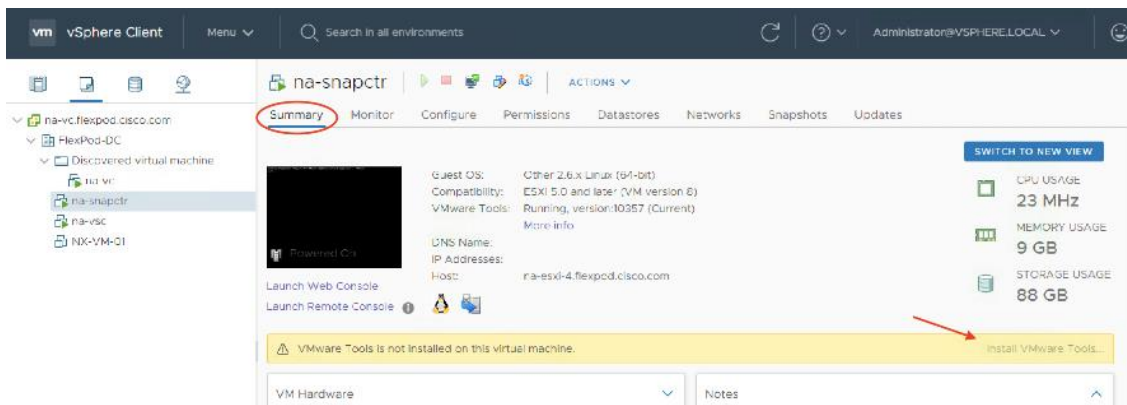
✓ 1 Select an OVF template	2. Create SCV Credentials	2 settings
✓ 2 Select a name and folder	2.1 Username	admin
✓ 3 Select a compute resource	2.2 Password	Password
✓ 4 Review details		Confirm Password
✓ 5 License agreements		
✓ 6 Select storage		
✓ 7 Select networks	3. Setup Network Properties	1 settings
✓ 8 Customize template	3.1 Host Name	Hostname for the appliance
9 Ready to complete		
	3.2 Setup IPv4 Network Properties	6 settings
	3.2.1 IPv4 Address	IP address for the appliance
		10.1156.202
	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
		10.1156.1
	3.2.4 IPv4 Primary DNS	Primary DNS server's IP address
		10.1156.250
	3.2.5 IPv4 Secondary DNS (optional)	Secondary DNS server's IP address

CANCEL BACK NEXT

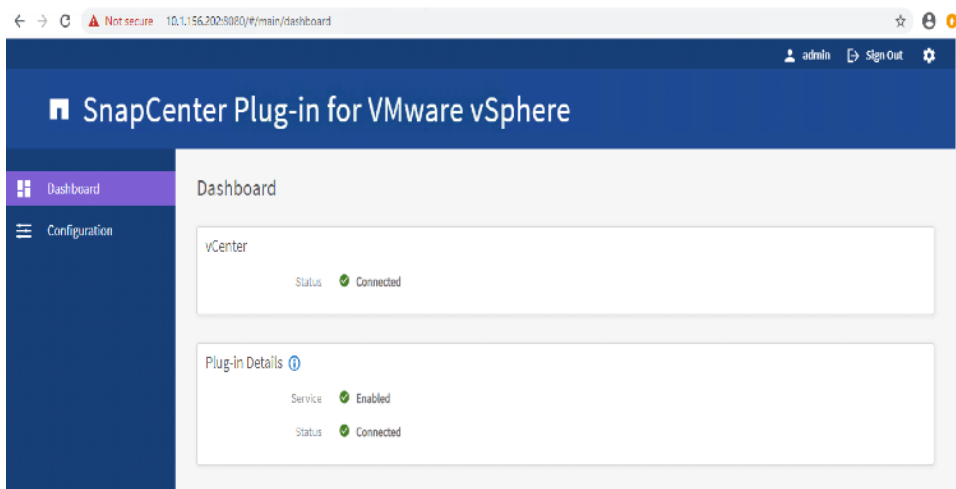
11. On the Ready to complete page, review the page and click FINISH.



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

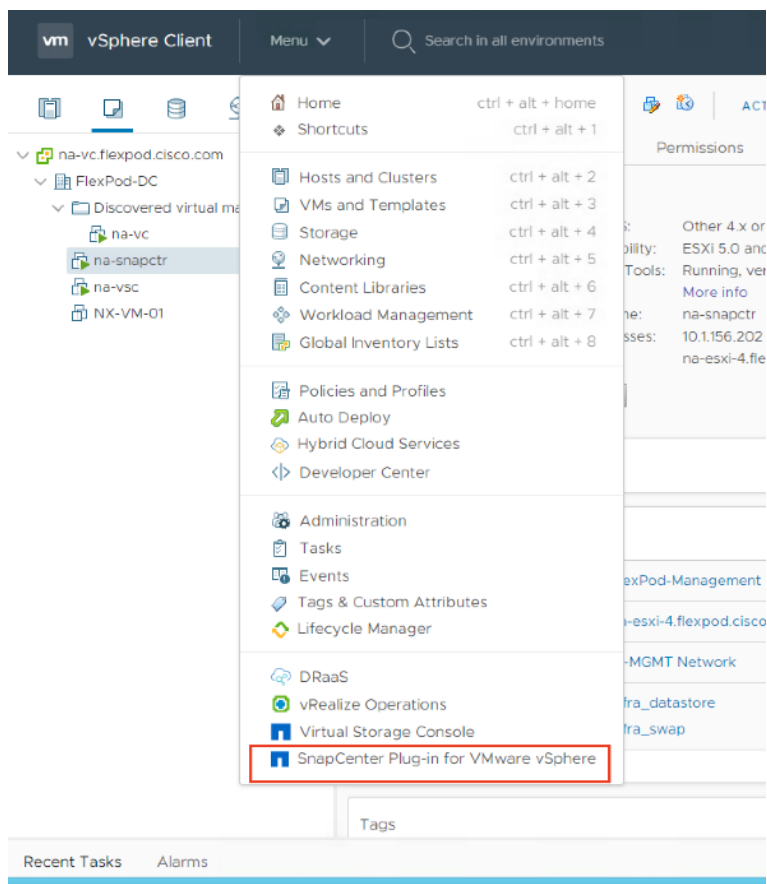
After you have successfully installed the Plug-in for VMware vSphere, to configure SnapCenter and make it ready to backup virtual machines, follow these steps:

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.

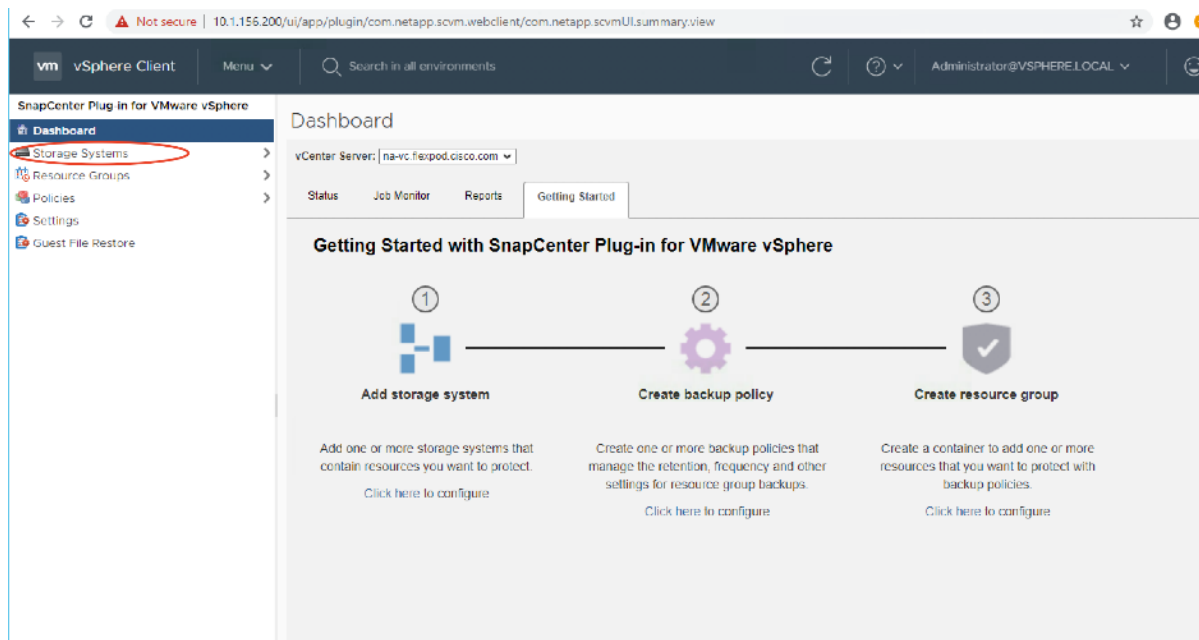
2. After logging on to the vSphere Web Client you will see a blue banner indicating the SnapCenter plug-in was successfully deployed. Click the refresh button to activate the plug-in.
3. On the VMware vSphere Web Client page, click the menu and click SnapCenter Plug-in for VMware vSphere to launch the 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:
 1. Enter the policy name and a description.
 2. Enter the backups to keep.
 3. From the Frequency drop-down list, choose the backup frequency (hourly, daily, weekly, monthly, and on-demand only).
 4. Expand the Advanced options and select VM Consistency and Include datastore with independent disks.
4. Click Add.

+ New Backup Policy ×

Name:

Description:

vCenter Server:

Retention: Days to keep ⓘ

Frequency:

Replication: Update SnapMirror after backup ⓘ
 Update SnapVault after backup

Snapshot label:

Advanced ▾: VM consistency
 Include datastores with independent disks

Scripts ⓘ:

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

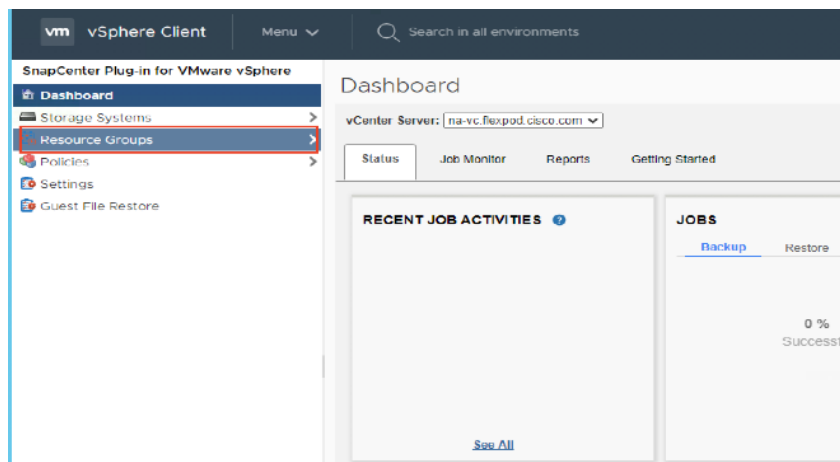
Create Resource Groups

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 Plug-in 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:

1. 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.



2. 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.
3. 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.

4. 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

- na-snapctr
- na-vc
- na-vsc
- NX-VM-01

>>

>

<

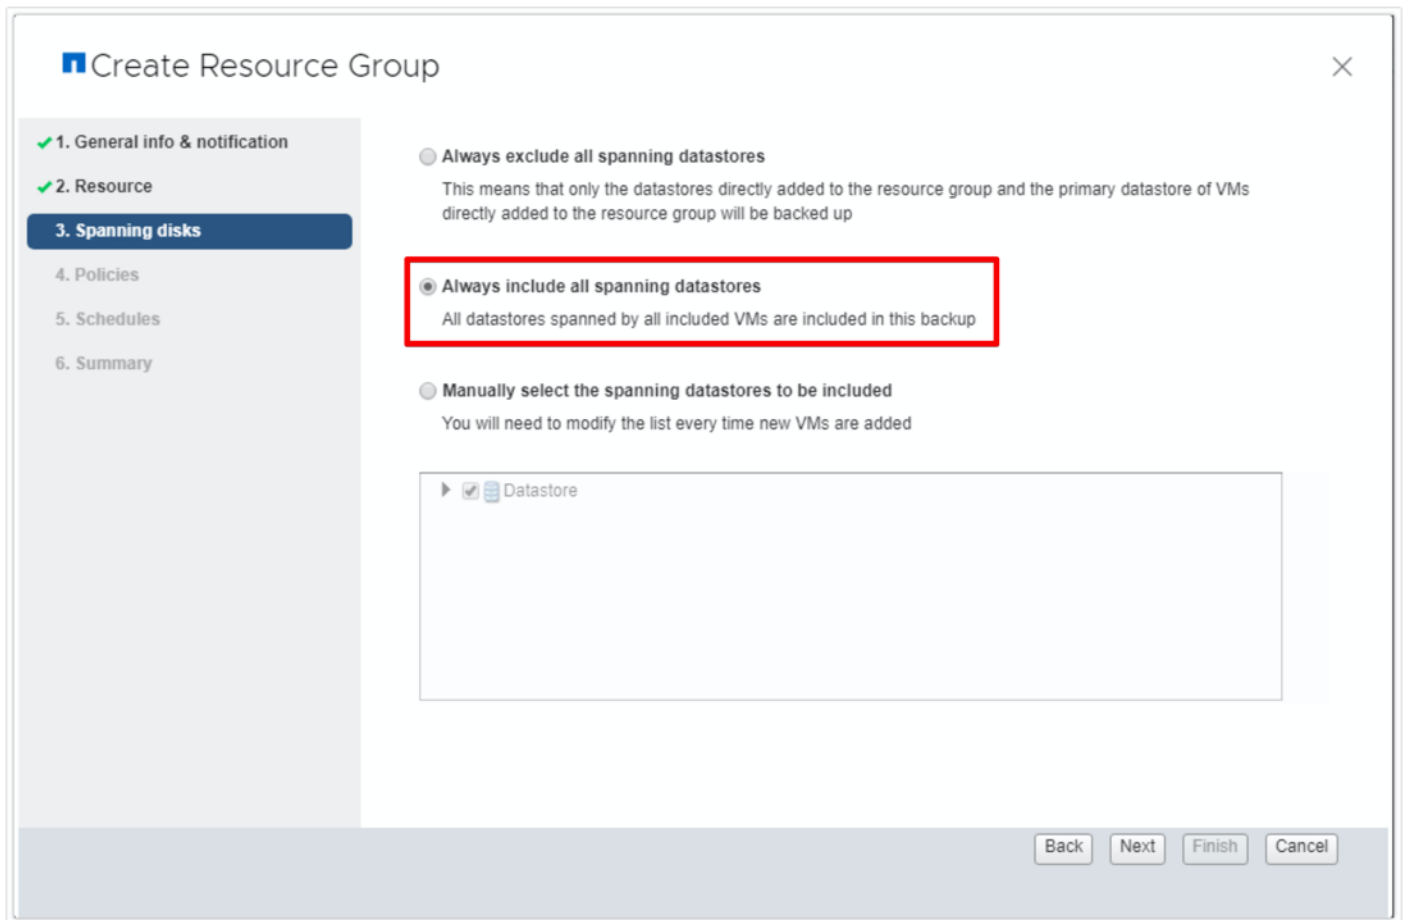
<<

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.



- 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

Infra_vm_bac... ▼

Type Daily

Every Day(s)

Starting 

At

8. 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

Name	Mgmt_VMs
Description	Infrastructure Management
Send email	Errors
Email send from	flexadmin@flexpod.cisco.com
Email send to	flexadmin@flexpod.cisco.com
Email subject	Mgmt VM Backup Errors
Custom snapshot format	<ResourceGroup>_<TimeStamp>
Entities	na-snapctr, na-vc, na-vsc, NX-VM-01
Spanning	True
Policies	infra_vm_back... : Hourly

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.

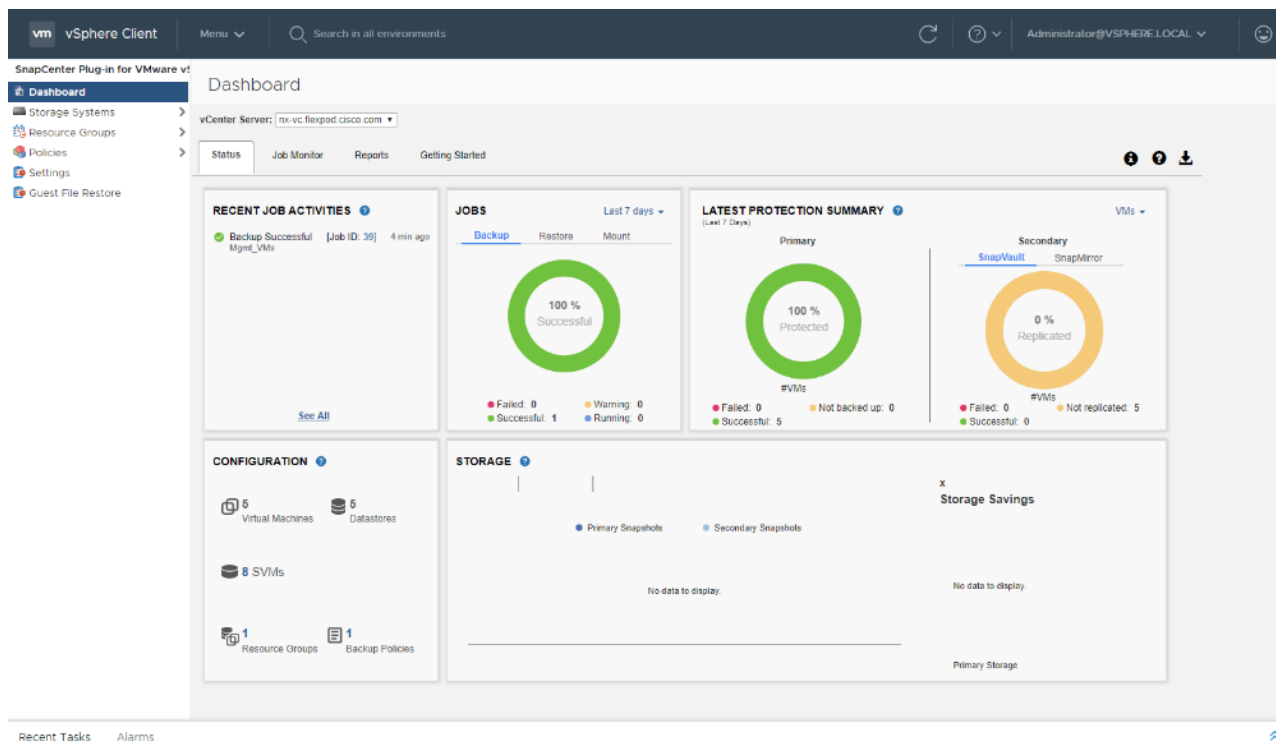
The screenshot shows the vSphere Client interface. The top navigation bar includes the 'vm' logo, 'vSphere Client', a 'Menu' dropdown, and a search bar labeled 'Search in all environments'. Below the navigation bar, there are icons for various objects. The left pane shows a tree view with the following structure:

- na-vc.flexpod.cisco.com
 - FlexPod-DC
 - Discovered virtual machine
 - na-vc (selected)
 - na-snapctr
 - na-vsc
 - NX-VM-01

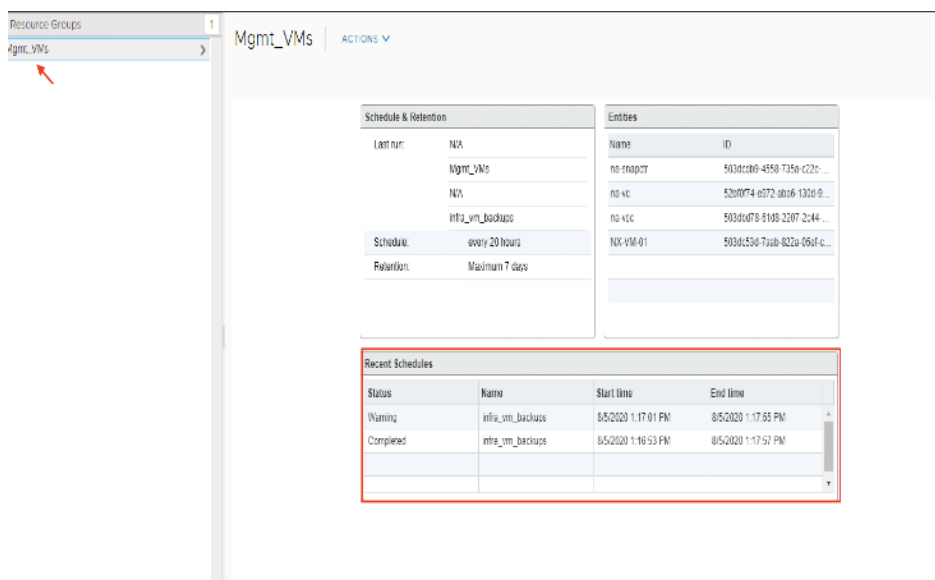
The right pane shows the 'na-vc' object selected. The 'Summary' tab is active, and the 'Backups' sub-tab is selected. Below the sub-tabs, there are two buttons: 'Resource Groups' and 'Backups'. The 'Backups' button is highlighted. Below the buttons, there is a table with the following data:

Name ↑
Mgmt_VMs_08-05-2020_13.16.54.0905
Mgmt_VMs_08-05-2020_13.17.01.0318

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



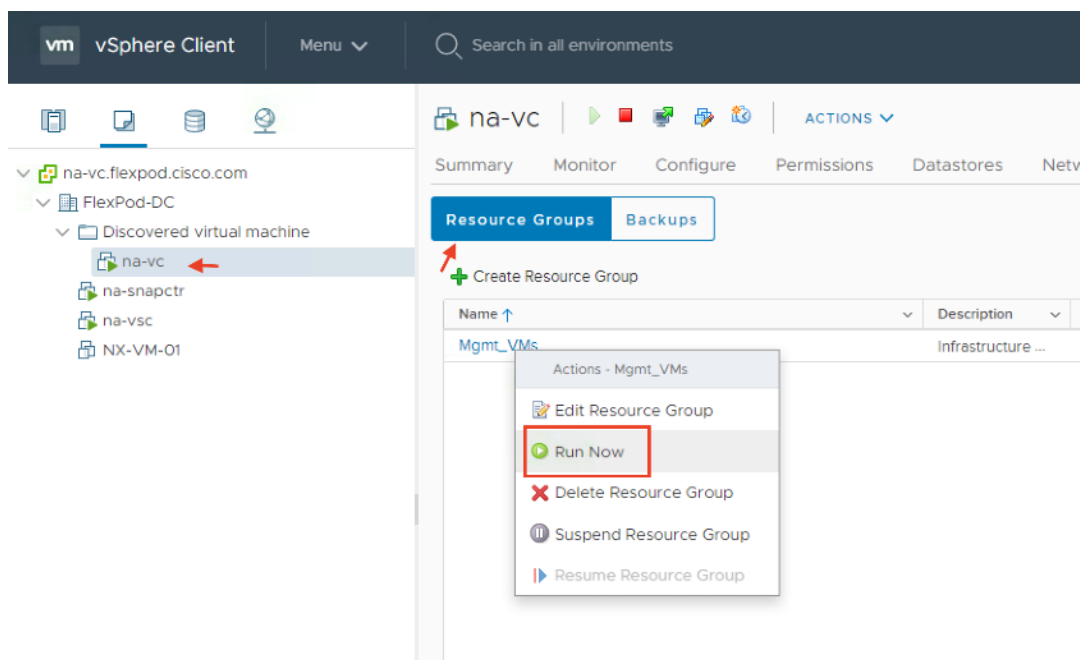
- In the SnapCenter Plug-in for VMware vSphere, click Resource Groups and choose any resource group. In the right pane, the completed backups are displayed.



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.



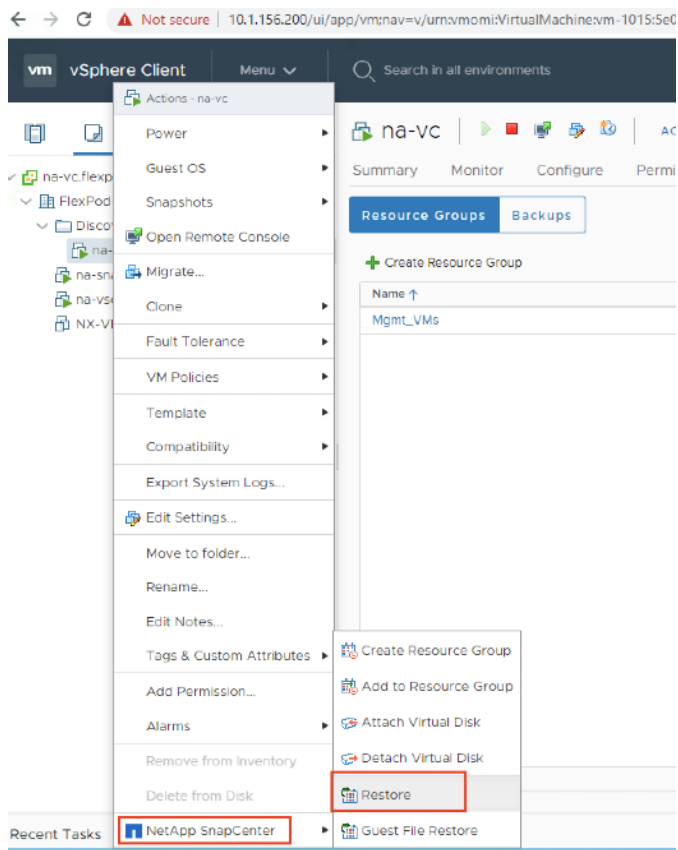
Restore from vCenter by Using SnapCenter Plug-In

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

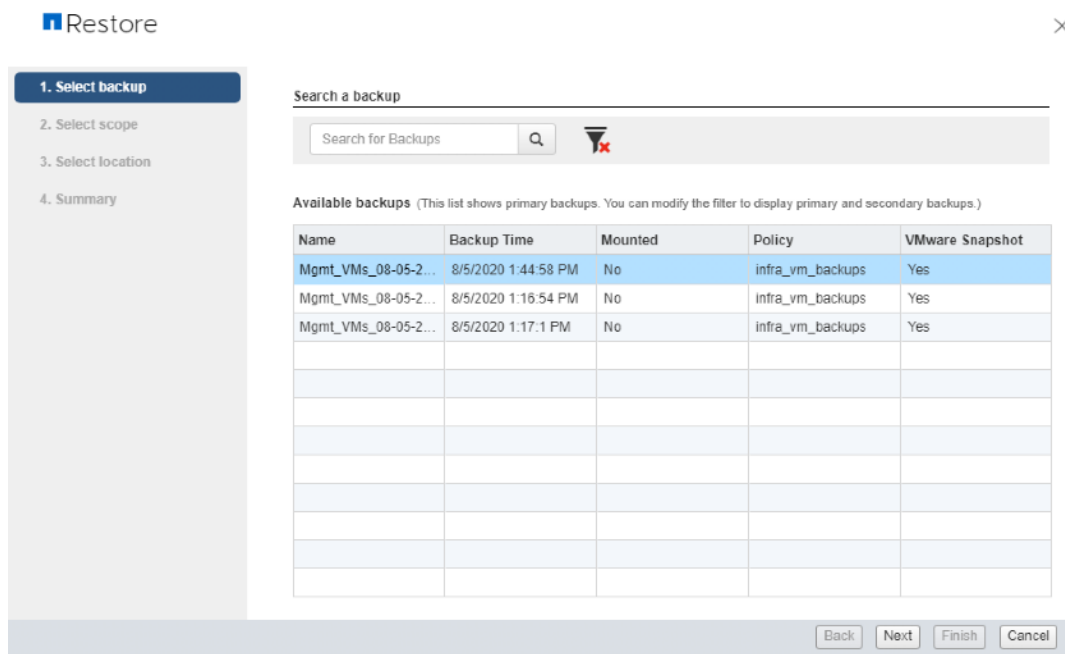


The Plug-in for VMware vSphere provides native backup, recovery, and cloning of virtualized applications.

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

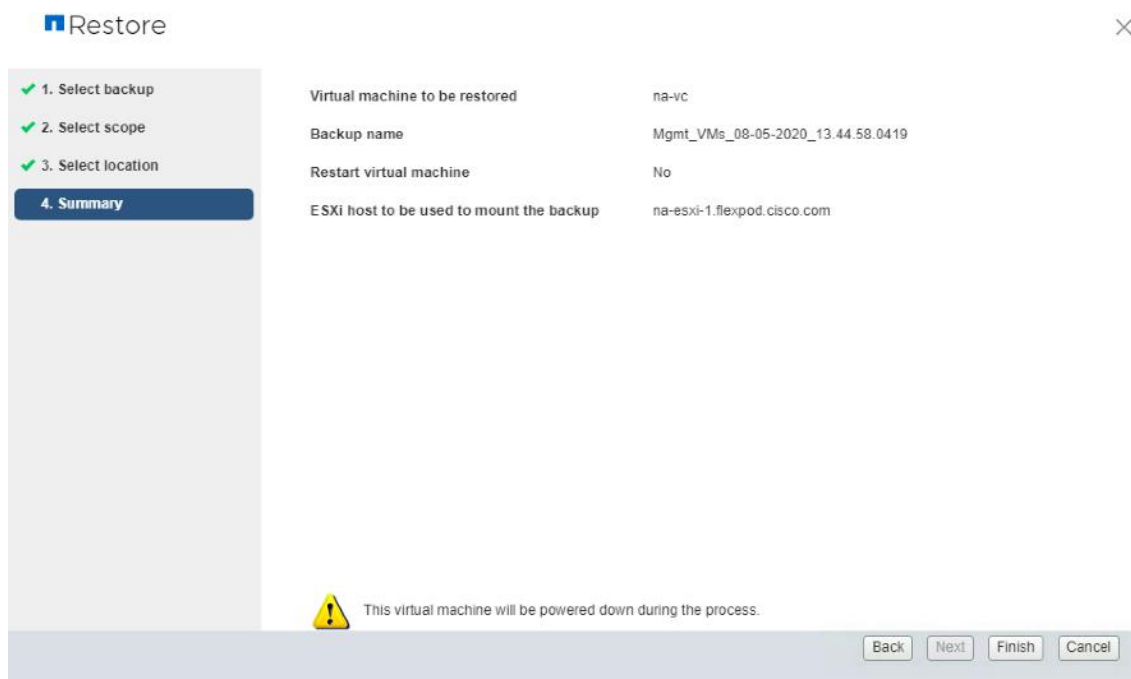


2. Choose a backup from which to restore. Click Next.



3. From the Restore Scope drop-down list:

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



Active IQ Unified Manager 9.7P1

Active IQ Unified Manager 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 9.7P1 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"> • 5 GB - Thin provisioned • 152 GB - Thick provisioned



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 Active IQ Unified Manager. See the [Unified Manager Best Practices Guide](#) (TR-4621) for more details.

To install Active IQ Unified Manager 9.7P1, 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.

- na-vc.flexood.cisco.com
 - FlexPod-OC

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 http://www.netapp.com
Download size	2.5 GB
Size on disk	4.4 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:
 1. Choose the disk format for the VMDKs.
9. Choose a VM Storage Policy.
10. 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: Thick Provision Lazy Zeroed

VM Storage Policy: **Datstore Default**

Name	Capacity	Provisioned	Free	Type	Cluster
infra_datstore	1 TB	693.28 GB	846.35 GB	NFS v3	
infra_swap	100 GB	14.16 MB	99.99 GB	NFS v3	
NX_FC_DS_01	500.25 GB	1.41 GB	498.84 GB	VMFS 6	
NX_NFS_DS_0	500 GB	372 KB	500 GB	NFS v3	

Compatibility

✓ Compatibility checks succeeded

CANCEL BACK NEXT

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

12. 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

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.
ne-eiurq.flexpod.cisco.co

IP Address
Specifies the IP address for the appliance. Leave blank if DHCP is desired.
10.1.156.203

Network Mask (or) Prefix Length
Specifies the subnet to use on the deployed network. Leave blank if DHCP is desired.
255.255.255.0

Gateway
Specifies the gateway on the deployed network. Leave blank if DHCP is desired.
10.1.156.1

Primary DNS
Primary DNS ip address. Leave blank if DHCP is desired.
10.1.156.250

Secondary DNS
Secondary DNS ip address. Leave blank if DHCP is desired.
10.1.156.251

CANCEL BACK NEXT



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

13. 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.

Name	na-aiqum
Template name	ActiveIQUnifiedManager-9.7P1
Download size	2.6 GB
Size on disk	152.0 GB
Folder	FlexPod-DC
Resource	FlexPod-Management
Storage mapping	1
All disks	Datastore: Intra_datastore; Format: Thick provision lazy zeroed
Network mapping	1
net	IB-MGMT Network
IP allocation settings	
IP protocol	IPv4
IP allocation	Static - Manual

[CANCEL](#)
[BACK](#)
[FINISH](#)

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

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

vm vSphere Client | Search in all environments | Administrator@VSPHERE LOCAL

na-aiqum | ACTIONS

Summary | Monitor | Configure | Permissions | Datastores | Networks | Snapshots | More Objects | Updates

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

Guest OS | ACTIONS

- Power Status: Powered On
- Guest OS: Other (32-bit)
- Encryption: Not encrypted
- VMware Tools: Not running, not installed
- DNS Name: --
- IP Addresses: --

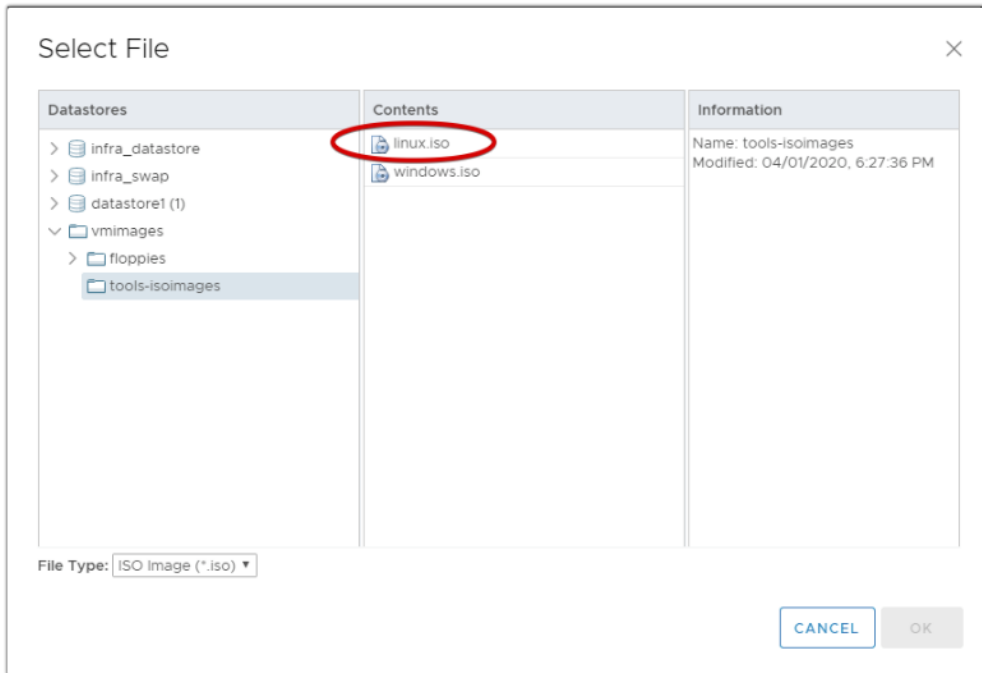
[LAUNCH REMOTE CONSOLE](#) [LAUNCH WEB CONSOLE](#)

Capacity and Usage

- CPU: 1.32 GHz used / 4 CPU(s) allocated
- Memory: 9 GB used / 12 GB allocated
- Storage: 152 GB used / 152 GB allocated

[VIEW STATS](#) Created: 8/5/20, 3:08 PM

16. 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.



17. 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 U timezones
  4. Australia       8. Europe          12. US

Geographic area: 2_

```

18. 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: _
```

19. Log into NetApp Active IQ Unified Manager 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://10.1.156.203/
```

```
or
```

```
https://na-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.
```

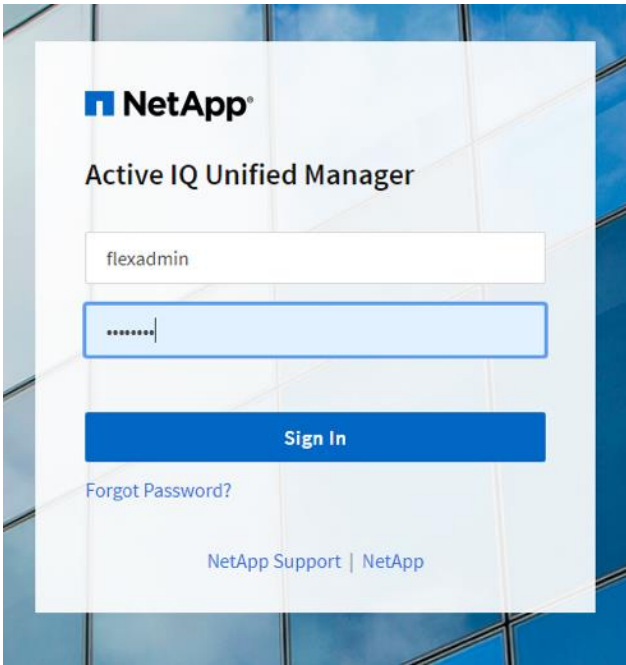
```
Hint: Num Lock on
```

```
na-aiqum login: _
```

Configure Active IQ Unified Manager

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

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



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.

Notifications

Configure your email server to allow Active IQ Unified Manager to assist in the event of a forgotten password.

Maintenance User Email

Email

SMTP Server

Host Name or IP Address

Port

User Name

Password

 Use START / TLS

Use SSL

Network Time Protocol (NTP) server

NTP server

[Continue](#)

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

Active IQ Unified Manager

Getting Started

Progress: 1. Email (Completed), 2. AutoSupport (Current), 3. API Gateway, 4. Add ONTAP Clusters, 5. Finish

Set up AutoSupport

AutoSupport is a service that sends periodic data to NetApp® Active IQ. Active IQ is a data-driven service that uses artificial intelligence, machine learning, and community wisdom to provide predictive analytics, actionable insights, and proactive support that help maximize availability and optimize performance in your NetApp data management environment. For more information see: <https://www.netapp.com/us/products/data-infrastructure-management/active-iq-predictive-technology.aspx>

AutoSupport will be enabled on this system. You can disable AutoSupport at a later time from Settings -> AutoSupport

Agree and Continue

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

Active IQ Unified Manager

Getting Started

Progress: 1. Email (Completed), 2. AutoSupport (Completed), 3. API Gateway (Current), 4. Add ONTAP Clusters, 5. Finish

Set up API Gateway

The API Gateway for Active IQ Unified Manager REST APIs enables you to control multiple ONTAP clusters by leveraging the cluster authentication and cluster management capabilities of Active IQ Unified Manager. This capability enables you to use Unified Manager as the single entry point for using ONTAP REST APIs without the need to log in to individual clusters.

Enable API Gateway

Continue

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

Active IQ Unified Manager

Getting Started

Email
 AutoSupport
 API Gateway
 4 Add ONTAP Clusters
 5 Finish

Add ONTAP Clusters

HOST NAME OR IP ADDRESS

CLUSTER USERNAME

CLUSTER PASSWORD

PORT

Recently added clusters (0)

Host name/IP Address	Data Acquisition Status
No clusters listed.	

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

Authorize Cluster Certificate

Host aa11-a400.flexpod.cisco.com you specified has identified itself with a self signed certificate for and the host does not match with the name (CN or DN): aa11-a400.

[View Certificate](#)

Do you want to trust this 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:
 1. Choose Local User for the Type.
3. Enter a username and password.
4. Add the user's email address.
5. Choose the appropriate role for the new user.
6. Click Save to add the new user to Active IQ Unified Manager.

Users: Add

TYPE

Local User 

 Authentication server is either disabled or not configured. To add a remote user or group, enable or configure the authentication server from Setup Options.

NAME

sree.lanka

PASSWORD

.....

CONFIRM PASSWORD

.....

EMAIL

.....@netapp.com

ROLE

Application Administrator 

Save

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 Active Directory 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 Active Directory users reside.
6. If Active Directory LDAP communications are protected via SSL enable the Use Secure Connection option.
7. Add one or more Active Directory 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 Active Directory 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:

i Disable Nested Group Lookup:

i 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

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

Port

389

389

Test User ✕

Enter the username to find the user in the authentication server.
Enter the username and password to authenticate the user.

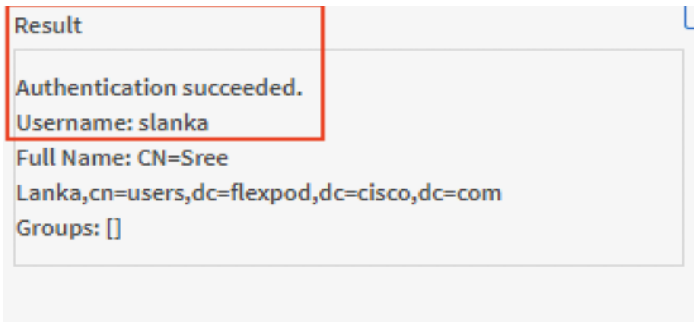
Username:

Password:

Start Cancel

Test Authentication

A result message displays 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.

A screenshot of the 'Users: Add' form. The form has a title 'Users: Add' with a help icon. It contains four fields: 'TYPE' with a dropdown menu set to 'Remote User', 'NAME' with a text input containing 'slanka', 'EMAIL' with a text input containing a redacted name followed by '@flexpod.cisco.com', and 'ROLE' with a dropdown menu set to 'Application Administrator'. At the bottom, there are two buttons: 'Save' (blue) and 'Cancel' (light blue).

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

Users ⓘ

+ Add Edit Delete

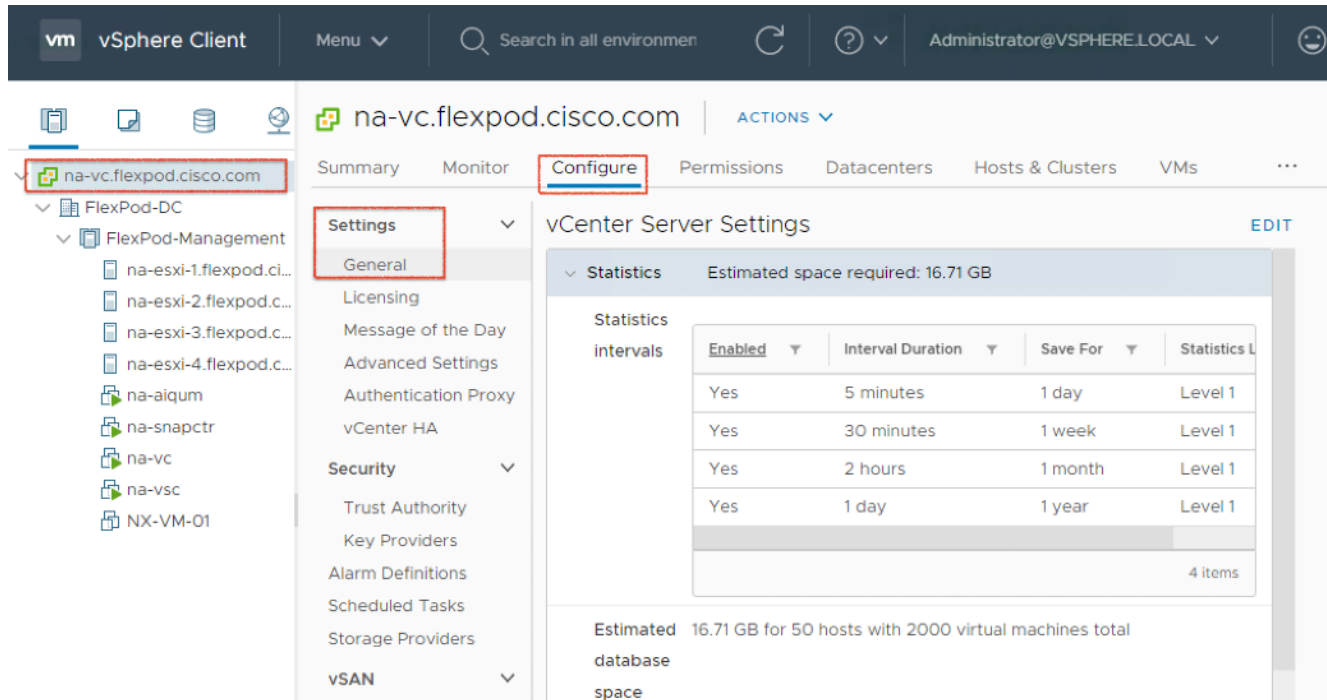
	Name	Type
<input type="checkbox"/>	flexadmin	Maintenance User
<input type="checkbox"/>	slanka	Remote User

Add the vCenter Server to Active IQ Unified Manager

Active IQ Unified Manager 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 Active IQ Unified Manager 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	1 day	Level 3
<input checked="" type="checkbox"/>	30 minutes	1 week	Level 1
<input checked="" type="checkbox"/>	2 hours	1 month	Level 1
<input checked="" type="checkbox"/>	1 day	1 year	Level 1

Database size

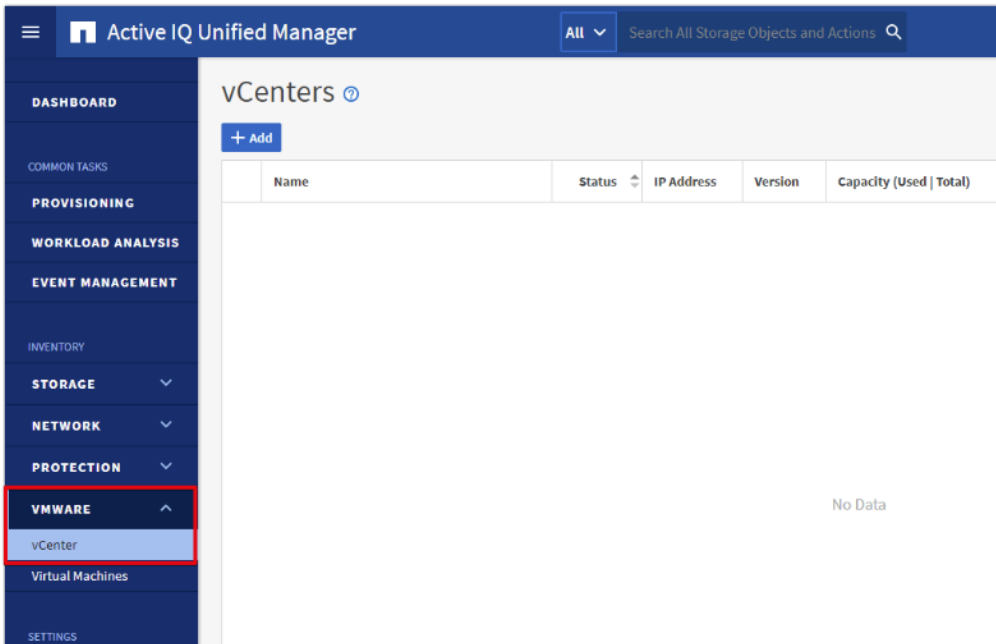
Based on the current vCenter Server inventory size, the vCenter Server database can be estimated. Enter the expected number of hosts and virtual machines in the inventory to calculate an estimate.

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.



- Expand the section and choose vCenter and click Add.
- Enter the VMware vCenter server details and click Save.

Add VMware vCenter Server

VCENTER SERVER IP ADDRESS OR HOST NAME

na-vc.flexpod.cisco.com

USERNAME

administrator@vsphere.local

PASSWORD

PORT

443

Save

Cancel

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

Authorize Certificate

Host na-vc.flexpod.cisco.com you specified has identified itself with a ca signed certificate for Active IQ Unified Manager.

[View Certificate](#)

Do you want to trust this certificate?

Yes

No



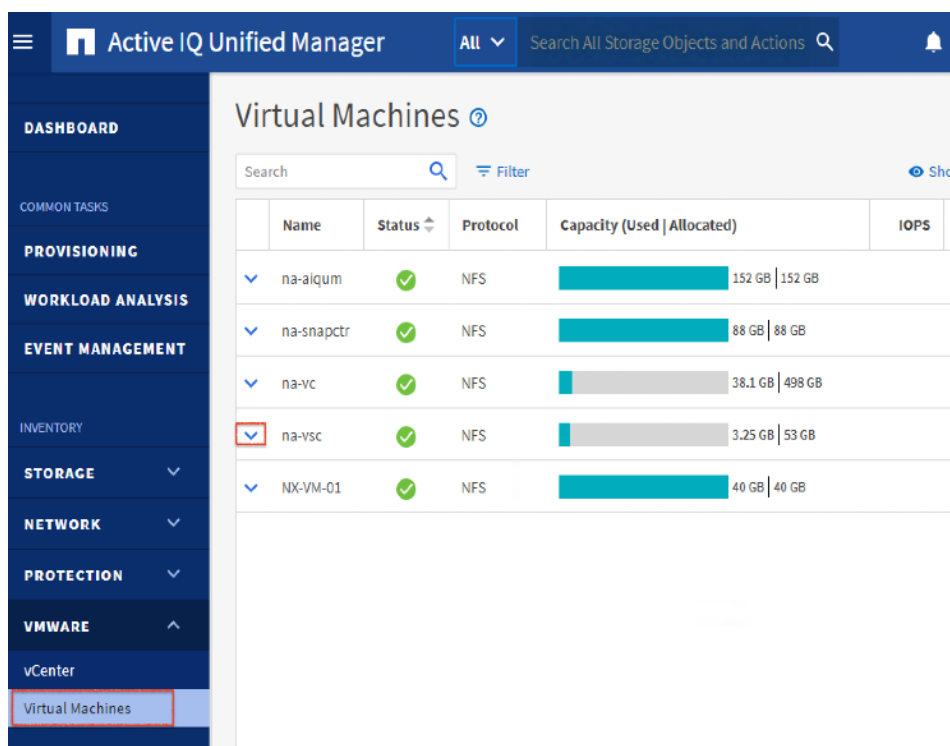
It may take up to 15 minutes to discover the vCenter server. Performance data can take up to an hour after discovery to become available.

View Virtual Machine Inventory

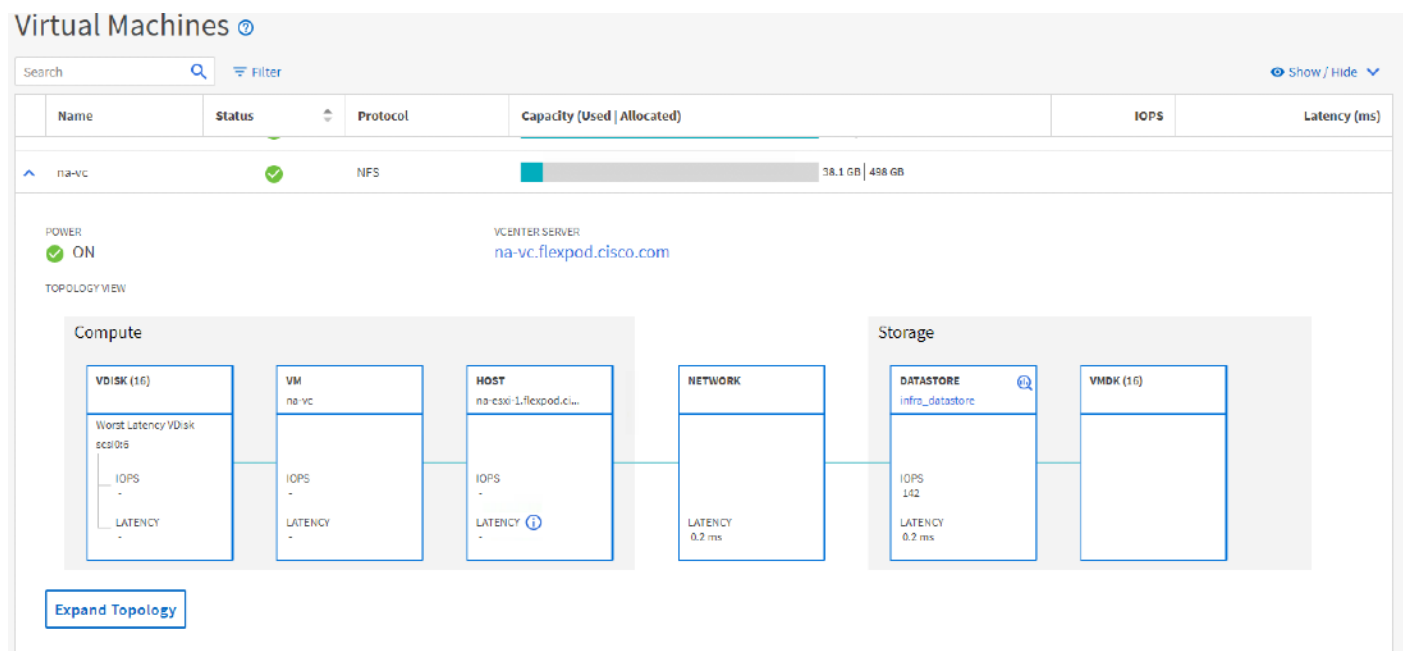
The virtual machine inventory is automatically added to Active IQ Unified Manager during discovery of the vCenter server. Virtual machines can be viewed in a hierarchical display detailing storage capacity, IOPS and latency for each component in the virtual infrastructure to troubleshoot the source of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

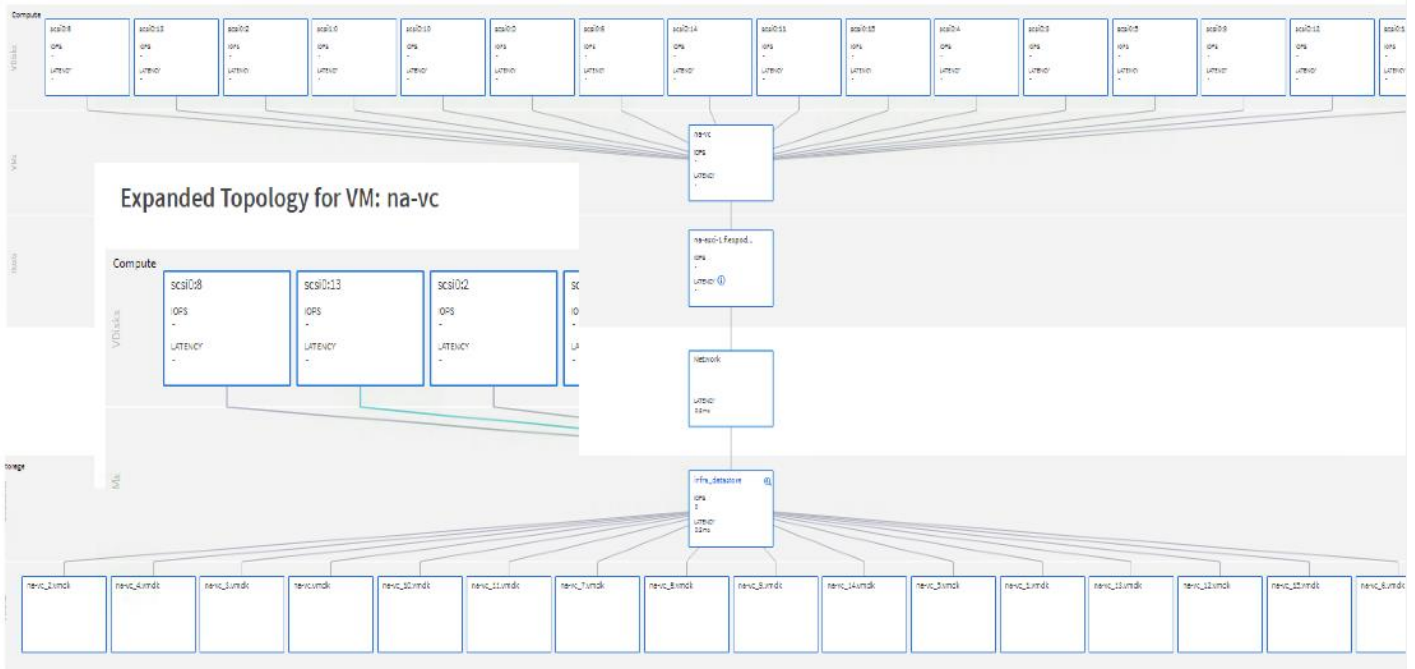


- Choose a VM and click the blue caret to expose the topology view. Review the compute, network, and storage components and their associated IOPS and latency statistics.



- 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.




Expanded Topology for VM: na-vc



Review Security Compliance with Active IQ Unified Manager

Active IQ Unified Manager identifies issues and makes recommendations to improve the security posture of ON-TAP. Active IQ Unified Manager 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:

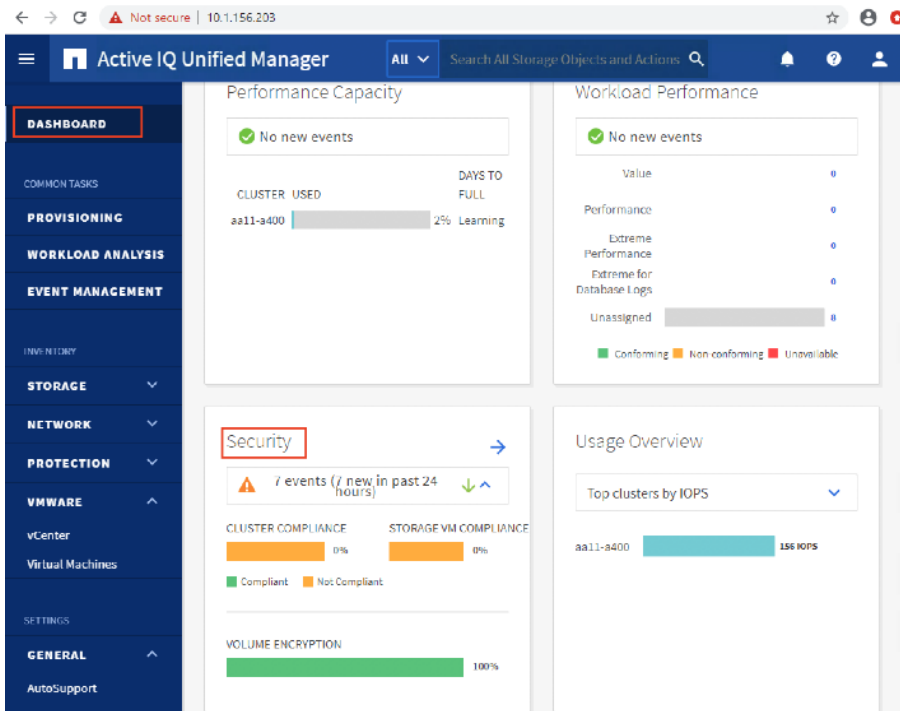
-  - The parameter is configured as recommended.
-  - The parameter is not configured as recommended.
-  - Either the functionality is not enabled on the cluster, or the parameter is not configured as recommended, but this parameter does not contribute to the compliance of the object.

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:

1. Navigate to the URL of the Active IQ Unified Manager installation and login.
2. Choose the Dashboard from the left menu bar in Active IQ Unified Manager.
3. 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.

Individual Cluster

aa11-a400

Cluster Compliance

Pro tips for Cluster compliance

SELECTED CLUSTER AND ALL STORAGE VM EVENTS

7 events (7 new in past 24 hours)

- SSH is using insecure ciphers (Infra-SVM)
- Audit Log Disabled (Infra-SVM)
- Login Banner Disabled (Infra-SVM)
- FIPS Mode Disabled (aa11-a400)

View All

Storage VM Compliance

Pro tips for Storage VM compliance

All Storage VMs

0% COMPLI

Compliant Storage VMs: 0
Not Compliant Storage VMs: 1

Individual Storage VM

- Infra-SVM
 - General Settings
 - iSCSI Settings
 - NFS Settings

10 ENCR

0% 20% 40% 60% 80%

Software Encrypted: 0
Hardware Encrypted: 8

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

Event Management

Last updated: Aug 5, 2020, 5:34

VIEW [dropdown] Search Events [input] Filter [icon]

Assign To [dropdown] Acknowledge [checkbox] Mark as Resolved [checkbox] Add Alert [icon] Show / Hide [dropdown]

<input type="checkbox"/>	Triggered Time	State	Severity	Impact Level	Impact Area	Name
<input type="checkbox"/>	Aug 5, 2020, 3:51 PM	New	Warning	Risk	Security	SSH is using insecure ciphers
<input type="checkbox"/>	Aug 5, 2020, 3:51 PM	New	Warning	Risk	Security	Audit Log Disabled
<input type="checkbox"/>	Aug 5, 2020, 3:51 PM	New	Warning	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.

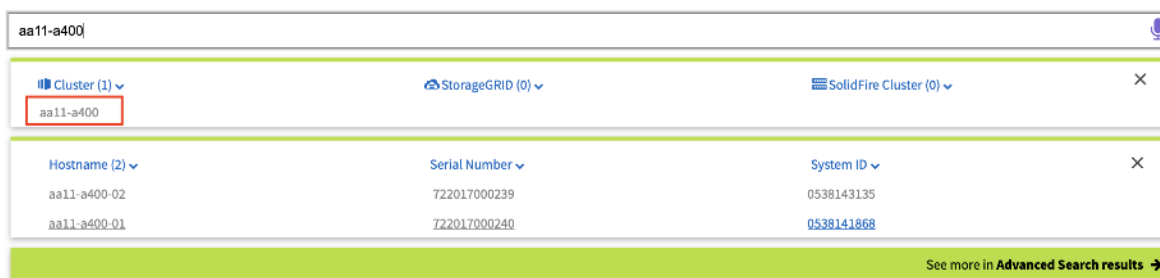
Additional Active IQ documentation is available on the [Active IQ Documentation Resources](#) web page.

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.

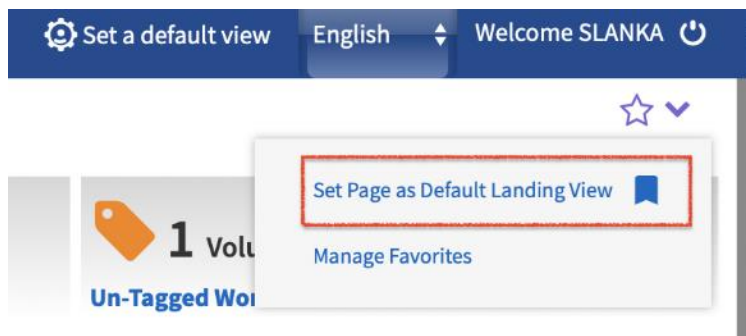


The screenshot shows the search results for the cluster 'aa11-a400'. The search bar at the top contains 'aa11-a400'. Below the search bar, there are three filter sections: 'Cluster (1)' with 'aa11-a400' selected, 'StorageGRID (0)', and 'SolidFire Cluster (0)'. Below these filters is a table with three columns: 'Hostname (2)', 'Serial Number', and 'System ID'. The table contains two rows of data.

Hostname (2) ▾	Serial Number ▾	System ID ▾
aa11-a400-02	722017000239	0538143135
aa11-a400-01	722017000240	0538141868

At the bottom right of the table, there is a link: 'See more in Advanced Search results →'.

5. Choose the cluster name to launch the main dashboard.
6. Click the drop-down list and select to make it the default dashboard.



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:

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

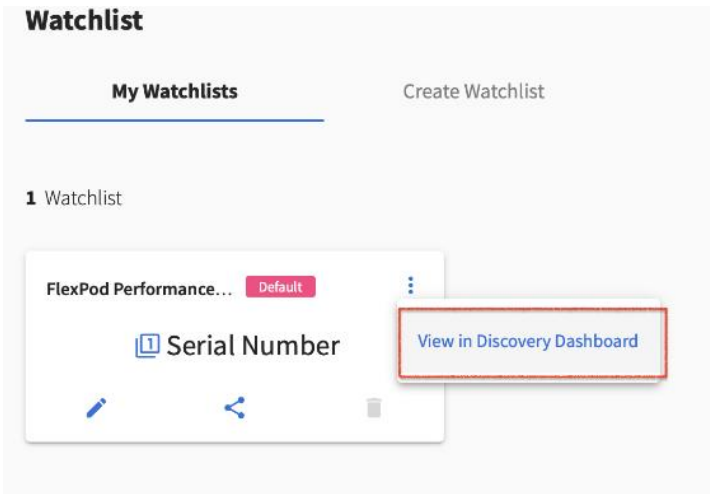
The screenshot shows the Active IQ Discovery Dashboard for Cisco Systems Inc. The breadcrumb trail is: Home > Cisco Systems Inc. > CISCO SYSTEMS - RTP - BUILDING 9 > aa11-a400. The dashboard displays a risk level of 0 High and an upgrade recommendation for 9.7P6 with 2 upgrades available. Below this, there is a table of nodes with the following data:

Hostname	Serial Number	System ID	OS Version	Model	FlexPod
aa11-a400-01	722017000240	0538141868	9.7P5	AFF-A400	No
aa11-a400-02	722017000239	0538143135	9.7P5	AFF-A400	No

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

The screenshot shows the Watchlist creation form. The form is titled "Watchlist" and has a sub-section "Create Watchlist". The "Name the Watchlist" field contains "FlexPod Performance Insights". Under "Add Systems by", the "Serial Number" radio button is selected. The "Choose Category" dropdown is set to "Serial Number". The "Paste Serial Numbers (Maximum Limit: 500)" field contains two input boxes with the values "39" and "40". The "Make this my default watchlist" checkbox is checked. A "Create Watchlist" button is at the bottom.

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



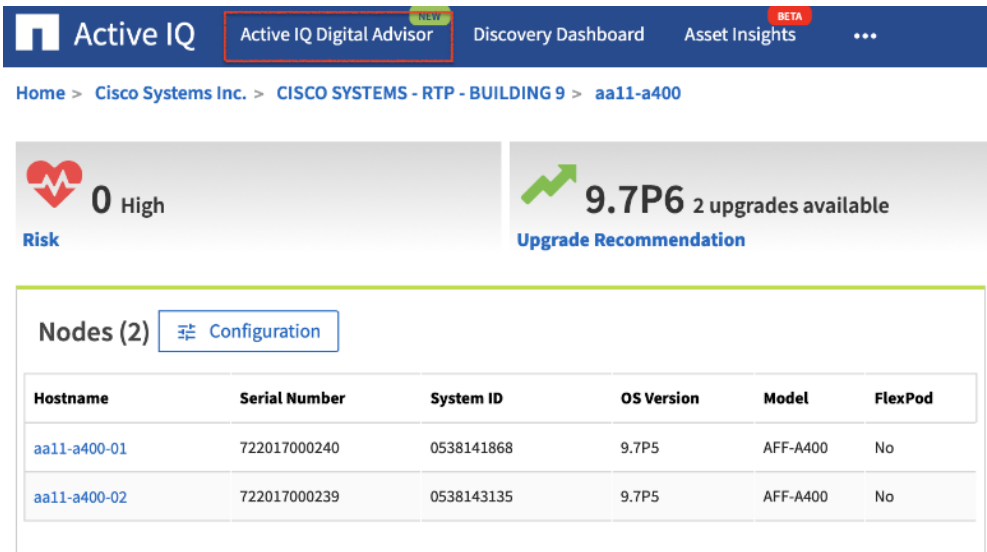
6. View the health and risk overview for the cluster.

Create Active IQ Digital Advisor Dashboard

The 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.

1 Select or Create Watchlist

2 Create Dashboard

Select Watchlist

1 Watchlist found

FlexPod Performance Insights

Watchlist Details

* Mandatory fields

Name the Watchlist *

FlexPod Performance Insights

Add Systems by

Category Serial Number

Choose Category

Serial Number

Paste Serial Numbers (Maximum Limit 500) *

39 240

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 FlexPod Performance Insights

* Mandatory fields

Dashboard name (Ex. Joey) *

FlexPod Performance Insights

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.

FlexPod Performance Insights

Wellness | Actions | Risks | View All

- Performance & Efficiency: No Pending Acti...
- Availability & Protection: No Pending Acti...
- Capacity: No Pending Acti...
- Configuration: 1 Action
- Security: No Pending Acti...
- Renewal: No Pending Acti...

Inventory | View All

Overview

● ONTAP

2 Systems, 1 Cluster, 1 Site

Planning

Capacity Addition, Renewal

Upgrades

Upgrades | Current Interoperability

1 Action

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.

Wellness | Actions | Risks

Performance & Efficiency: No Pending Acti...

Availability & Protection: No Pending Acti...

FlexPod Performance Insights > Wellness

Wellness | Update FAS / AFF Firmware

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

Actions (1) | **Unique Risks (3)** | Affected Systems

View Acknowledged Risks

Filter by: SW Config Change | Search by Risk Name

Fix It	Risk Name ↑	Mitigation ↑	Corrective Action	Systems	Impact ↑	Acknowledge	Public	Internal Info
	This system requires an updated Disk Qualificatio...	Potentially Non-disruptive	Disk Qualification Package KB ID: 1363	2	Medium	Ack	Yes	Signature: 2648
	Verify that DNS is configured for admin and data S...	Potentially Non-disruptive	Configuring DNS services for the SVM Creating SVM - Includes SAN SVM Information	2	Best Practice	Ack	No	Signature: 3676
	The node is not configured to save configuration ...	Potentially Non-disruptive	Backing up and restoring cluster configurations Commands for managing configuration backup schedules	2	Best Practice	Ack	No	Signature: 3191

7. Click the link in the Corrective Action column to read the bug information or knowledge base article about 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 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:

1. 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.



If using the Cisco Nexus 93180YC-FX for SAN switching, it does not support SAN Analytics.

2. Passwords. Cisco DCNM-SAN passwords should adhere to the following password requirements:
 1. It must be at least eight characters long and contain at least one alphabet and one numeral.
 2. It can contain a combination of alphabets, numerals, and special characters.
 3. Do not use any of these special characters in the DCNM password for all platforms: <SPACE> " & \$ % ' ^ = < > ; : ` \ | / , . *
3. 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>
4. 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.

5. 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	:	SHA
Use AES-128	:	Yes
Password	:	<input type="password" value="*****"/>
Confirm Password	:	<input type="password" value="*****"/>
Privacy Password	:	<input type="password" value="*****"/>
Confirm Privacy Password	:	<input type="password" value="*****"/>

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

Deploy the Cisco DCNM-SAN OVA

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

1. Download the Cisco DCNM 11.4.1 Open Virtual Appliance for VMware from [https://software.cisco.com/download/home/281722751/type/282088134/release/11.4\(1\)](https://software.cisco.com/download/home/281722751/type/282088134/release/11.4(1)). Extract dcnm-va.11.4.1.ova from the ZIP file.
2. In the VMware vCenter HTML5 interface, click Menu > Hosts and Clusters.
3. Right-click the FlexPod-Management cluster and select Deploy OVF Template.
4. Choose Local file then click UPLOAD FILES. Navigate to choose dcnm-va.11.4.1.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

Local file

dcnm-va.11.4.1.ova

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)	Description 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	
<input type="radio"/> ComputeHuge	
5 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	Cluster
Infra_datastore	500 GB	683.79 GB	392.85 GB	NFS v3	
Infra_swap	100 GB	33.77 MB	99.97 GB	NFS v3	

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.

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**

Ready to complete

Click Finish to start creation.

Name	na-dcnm
Template name	dcnm
Download size	4.8 GB
Size on disk	9.1 GB
Folder	FlexPod-DC
Resource	FlexPod-MGMT
Storage mapping	1
All disks	Datastore: infra_datastore; Format: Thin provision
Network mapping	3
dcnm-mgmt	IB-MGMT Network
enhanced-fabric-mgmt	IB-MGMT Network
enhanced-fabric-inband	IB-MGMT Network
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual

CANCEL

BACK

FINISH


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 set to match the number of CPUs in the UCS servers used in this deployment. The following example shows 2 sockets.

Edit Settings | na-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	▼ MHz ▼
Limit	Unlimited	▼ MHz ▼
Shares	Normal ▼	32000
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. Choose SAN only for the Installation mode and leave Cisco Systems, Inc. for the OEM vendor and click Next.
20. Enter and repeat the administrator and database passwords and click Next.
21. Enter the DCNM FQDN, a comma-separated list of DNS servers, a comma-separated list of NTP servers, and select the appropriate time zone. Click Next.

Cisco DCNM Installer

[Install Mode](#)[Administration](#)[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. Digit-only host names are not allowed.

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)

Timezone*[Previous](#)[Next](#)

- 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.
- Leave Internal Application Services Network set at the default setting and click Next.
- Review the Summary details and click 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.

Configure DCNM-SAN

To configure the DCNM-SAN, follow these steps:



When the DCNM installation is complete, the browser should redirect to the DCNM management URL.

- Log in as admin with the password entered above.

2. On the message that appears, choose Do not show this message again and click No.



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.

3. In the menu on the left, click Inventory > Discovery > LAN Switches.



4. 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:

5. 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.
6. After a few minutes (click 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.

	<input type="checkbox"/>	Switch	IP Address	Serial No	Managed	SNMP Status	Role	Last Updated Time	Group	User	Auth/Priv...
1	<input type="checkbox"/>	AA13-6454-A	192.168.156.18	FDO22191DZ5	true	SSH: There w...	leaf	2020-08-11 09:52:57	Default_LAN	snmpadmin	SHA_AES
2	<input type="checkbox"/>	AA13-6454-B	192.168.156.19	FDO22191DNN	true	SSH: There w...	leaf	2020-08-11 09:53:02	Default_LAN	snmpadmin	SHA_AES
3	<input type="checkbox"/>	aa11-93180-a	192.168.156.21	FDO24170XVH	true	ok	leaf	2020-08-11 09:53:02	Default_LAN	snmpadmin	SHA_AES
4	<input type="checkbox"/>	aa11-93180-b	192.168.156.22	FDO24170Y56	true	ok	leaf	2020-08-11 09:53:02	Default_LAN	snmpadmin	SHA_AES

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



8. Click  to add a switching fabric.

9. Enter either the IP address or 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.



If the Cisco Nexus 93180YC-FX switches are being used for SAN switching, substitute them here for MDS 9132Ts. They will need to be added again under SAN switches since LAN and SAN switching are handled separately in DCNM.

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:

10. Repeat steps 1-9 to add the second Cisco MDS 9132T and Fabric Interconnect.

The two SAN fabrics should now appear in the Inventory.

Inventory / Discovery / SAN Switches

<input type="checkbox"/>	Name	Seed Switch	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

11. Choose Inventory > Discovery > Virtual Machine Manager.



12. Click  to add the vCenter.

13. In the Add VCenter window, enter the IP address of the vCenter VCSA. Enter the administrator@vsphere.local user name and password. Click Add.


The vCenter should now appear in the inventory.



14. Choose Administration > Performance Setup > LAN Collections.

15. 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

-  aa11-93180-a
-  aa11-93180-b

16. Choose Administration > Performance Setup > SAN Collections.

17. 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"/>

18. 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.

19. 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. See [Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.4\(1\)](#).

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.
- Only Cisco MDS switches support SAN Analytics. Nexus 93180YC-FX switches do not support SAN Analytics.
- For more information on SAN Insights, see the SAN Insights sections: [Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.4\(1\)](#).
- 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.

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

1. 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: 10:06:10.494 EDT Tuesday August 11 2020

Selected 1 / Total 1

Switch		Model	Release	Licensed	Switch Time	Subscriptions	Install Query	Receiver
<input checked="" type="checkbox"/>	aa13-9132t-a	DS-C9132T-K9	8.4(1a)	Yes	10:06:12.790 EDT Tue Aug 11 2020	SCSI	Storage	10.1.156.210

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 on the ports connected to the NetApp AFF A400. 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)	6 interface(s)		Storage			
	▼ DS-C9132T-K9-S...	6 interface(s)					
		fc1/1	NX-Infra-SVM-fc-lif-1a	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	
		fc1/2	NX-Infra-SVM-fc-lif-2a	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	
		fc1/7	NetApp_50:0a:09:83:8...	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	
		fc1/8	NetApp_50:0a:09:83:8...	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	
		fc1/9	NetApp_50:0a:09:83:8...	Storage	disabled	<input checked="" type="checkbox"/> <input type="checkbox"/>	pending enable
		fc1/10	NetApp_50:0a:09:83:8...	Storage	disabled	<input checked="" type="checkbox"/> <input 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 steps 1-7 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.

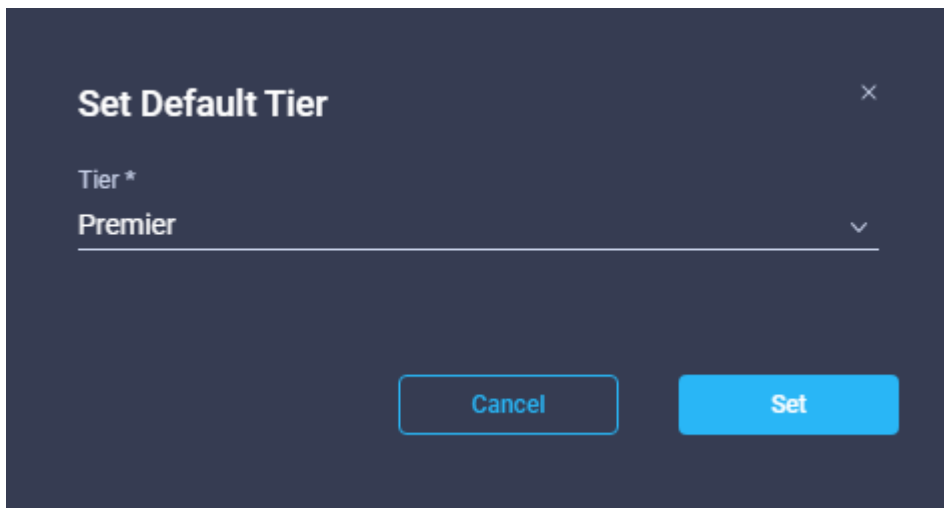
Cisco Intersight



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>. If you have an existing Intersight account, connect to <https://intersight.com> and sign in with your Cisco ID, select the appropriate account, and skip to step 6.
2. Click Create an account.
3. Sign in with your Cisco ID.
4. Read, scroll through, and accept the End User License Agreement and click Next.
5. Enter an Account Name and click Create.
6. Choose ADMIN > Targets. Click Claim a New Target. Select Cisco UCS Domain (UCSM Managed) and click Start. Fill in the Device ID and Claim Code and click Claim. 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.
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.
9. From the Licensing Window, click Set Default Tier. From the drop-down list choose Premier for Tier and click Set.



10. 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.



11. Click Refresh to refresh the Intersight window with Premier, Advantage, and Essentials features added.
12. Click  in the Intersight window and click Take a Site Tour. Follow the prompts for a tour of Cisco Intersight.
13. 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 page in Cisco Intersight. It is divided into three main sections, each with a 'Validated' status:

- Server Hardware Compliance (1):**
 - Server Model: UCSC-C125
 - CPU: AMD EPYC 7552 48-Core Processor
 - Server Firmware Version: 4.1(2b)C
- Server Software Compliance (2):**
 - OS Vendor: VMware ESXi
 - OS Version: 7.0.0
- Adapter Compliance (3):**

Model	Hardware Status	Software Status	Firmware Version	Driver Protocol	Driver Version
UCSC-PCIE-C25Q-04	Validated	Validated	5.1(2e)	nfnic	4.0.0.56-10EM.670.0.0.81
UCSC-PCIE-C25Q-04	Validated	Validated	5.1(2e)	nenic	1.0.33.0-10EM.670.0.0.81

14. Using the Intersight Assist personality of the Cisco Intersight Virtual Appliance VMware vCenter, it 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 latest release of the OVA from <https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-148>.
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 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 browse to the intersight-virtual-appliance-1.0.9-148.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

Local file

intersight-virtual-appliance-1.0.9-148.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

▼ CPU	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 ▼ Advanced...		
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	▼	i
> 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>.




It may take a few minutes for <https://intersight-assist-fqdn> to respond.

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

What would you like to Install ?

Intersight Connected Virtual Appliance 

Intersight Private Virtual Appliance 

Intersight Assist 

 Recover from backup

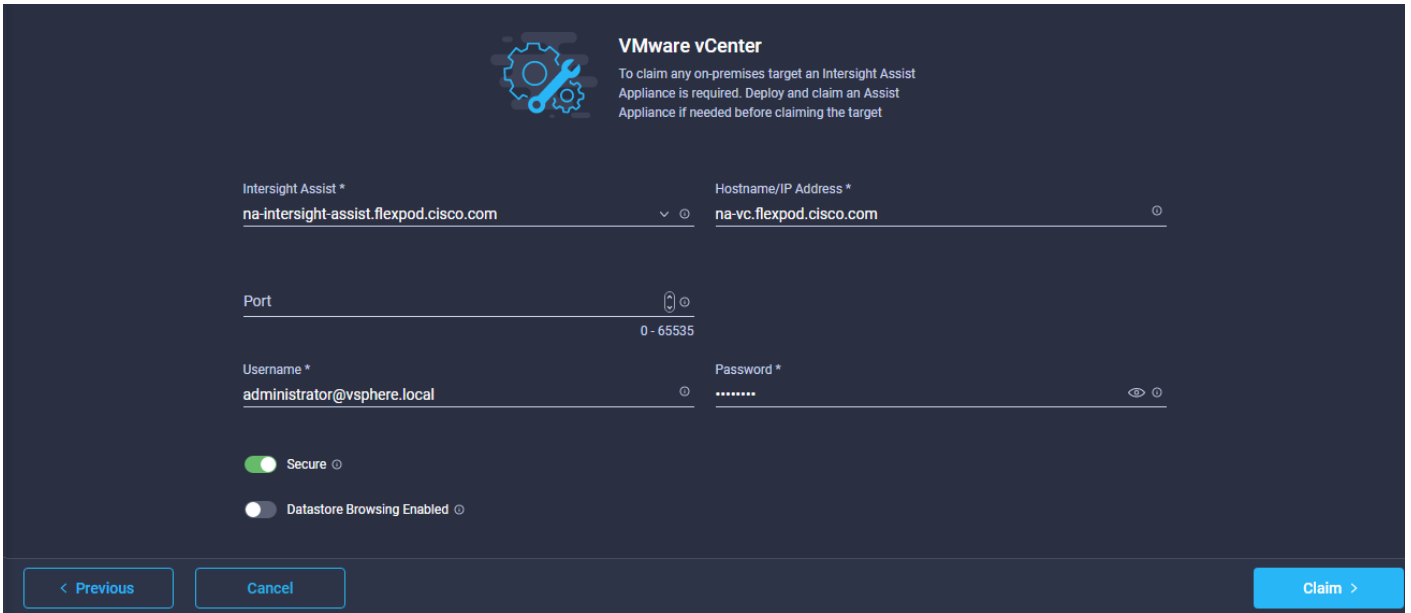
Proceed

32. From Cisco Intersight, click ADMIN > 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 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.



The Intersight Assist VM will reboot during the software download process. It will be necessary to refresh the Web Browser after the reboot is complete to follow the status of the download process.

35. When the software download is complete, navigate the security prompts and an Intersight Assist login screen will appear. Log into Intersight Assist with the admin 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 ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select VMware vCenter under Hypervisor and click Start. In the VMware vCenter window, make sure the Intersight Assist is correctly selected, fill in the vCenter information, and click Claim.



VMware vCenter
To claim any on-premises target an Intersight Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target

Intersight Assist *
na-intersight-assist.flexpod.cisco.com

Hostname/IP Address *
na-vc.flexpod.cisco.com

Port
0 - 65535

Username *
administrator@vsphere.local

Password *

Secure

Datastore Browsing Enabled

[< Previous](#) [Cancel](#) [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.

OPERATE > Virtualization > Datacenters

Datacenters Clusters Hosts Virtual Machines Datastores

Search

Export 1 items found 10 per page 1 of 1

Name	Datastores	Networks	Clusters	Hosts	Virtual Machines	Hypervisor Manager
FlexPod-DC		43	17	3	11	10.1.156.200

OPERATE > Virtualization > Clusters > FlexPod-Management

General Hosts Virtual Machines

Search

Export 3 items found 10 per page 1 of 1

Name	Datacenter	Cluster	CPU Capacity	CPU Utilization	Memory Capacity	Memory Utilization	CPUs
na-esxi-1.flexpod.cisco.com	FlexPod-DC	FlexPod-Management	73.60 GHz	2.9%	382.65 GiB	17.0%	32
na-esxi-2.flexpod.cisco.com	FlexPod-DC	FlexPod-Management	73.60 GHz	4.4%	1007.66 GiB	12.7%	32
na-esxi-3.flexpod.cisco.com	FlexPod-DC	FlexPod-Management	100.00 GHz	0.2%	1007.66 GiB	0.7%	40

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 has 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:
 1. 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.
 2. 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.
 3. Create the tenant SVM and follow all procedures in that section.
 4. Create Load-Sharing Mirrors for the tenant SVM.
 5. Create the FC or iSCSI service for the tenant SVM if fibre channel or iSCSI is being deployed in this tenant.
 6. Optionally, create a self-signed security certificate for the tenant SVM.
 7. Configure NFSv3 for the tenant SVM.
 8. Create a VM datastore volume in the tenant SVM.
 9. 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.
 10. If iSCSI is being deployed in this tenant, configure four iSCSI LIFs in the tenant SVM on the iSCSI VLAN interfaces.
 11. Create an NFS LIF in the tenant SVM on each storage node.
 12. Create a boot LUN in the esxi_boot volume in the Infra-SVM for each tenant VMware ESXi host.

-
13. 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.
 1. Create a Server Pool for the tenant ESXi host servers.
 2. Create all tenant VLANs in the LAN Cloud.
 3. Add the tenant VLANs to the vDS vNIC templates.
 4. 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, Cisco Nexus 93180YC-FX switches, or Cisco DCNM:
 1. Create device aliases for the tenant ESXi host vHBAs and the FC LIFs in the tenant storage SVM.
 2. Add the tenant host initiators to the Infra-SVM zone.
 3. Create a zone for the tenant SVM with fibre channel targets from the tenant SVM.
 4. Add these zones to the Fabric zoneset and activate the zoneset.
 6. In the storage cluster:
 1. 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.
 2. 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 with Cisco Nexus 93180YC-FX SAN Switching Configuration - Part 1

FlexPod Cisco Nexus Switch Base Configuration

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment using the Nexus 93180YC-FX switches for both LAN and SAN switching. This procedure assumes the use of Cisco Nexus 9000 9.3(5), 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.

Set Up Initial Configuration in Cisco Nexus 93180YC-FX A

To set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>, follow this step:

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
Enter basic FC configurations (yes/no) [n]: y
Configure default physical FC switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: y
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
```

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

Review the configuration summary before enabling the configuration.

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

Set Up Initial Configuration in Cisco Nexus 93180YC-FX B

To set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>, follow this step:

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
Enter basic FC configurations (yes/no) [n]: y
Configure default physical FC switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: y
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

Review the configuration summary before enabling the configuration.

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

Enable Features in Cisco Nexus 93180YC-FX A and B

SAN switching requires both the SAN_ENTERPRISE_PKG and FC_PORT_ACTIVATION_PKG licenses. Please ensure these licenses are installed on each Nexus 93180YC-FX switch. To enable the appropriate features on the Cisco Nexus switches, follow these steps:

1. Log in as admin.
2. Since basic FC configurations were entered in the setup script, feature-set fcoe has been automatically installed and enabled. Run the following commands:

```
config t
feature npiv
feature fport-channel-trunk
feature udd
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

Perform TCAM Carving and Configure Unified Ports in Cisco Nexus 93180YC-FX A and B

SAN switching requires TCAM carving for lossless fibre channel no-drop support. Also, unified ports need to be converted to fc ports. To perform TCAM carving on the Cisco Nexus switches and to convert ports 1-16 to fc, follow these steps:

1. Run the following commands:

```
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region ing-ifacl 256
hardware access-list tcam region ing-redirect 256
slot 1
port 1-16 type fc
copy running-config startup-config
reload
This command will reboot the system. (y/n)? [n] y
```

2. After the switch reboots, log back in as admin. Run the following commands:

```
show hardware access-list tcam region |i i ing-racl
show hardware access-list tcam region |i i ifacl
show hardware access-list tcam region |i i ing-redirect
show int status
```

Set System-Wide QoS Configurations in Cisco Nexus 93180YC-FX A and B

To set system-wide QoS configurations for FCoE for no-drop traffic support, follow this step on both switches:

1. Run the following commands to set global configurations:

```
config t
system qos
service-policy type queuing input default-fcoe-in-que-policy
service-policy type queuing output default-fcoe-8q-out-policy
service-policy type network-qos default-fcoe-8q-nq-policy
copy run start
```

Set Global Configurations in Cisco Nexus 93180YC-FX A and B

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

1. Run the following commands to set global configurations:

```
system default switchport mode F
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
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day> <end-month> <end-time> <offset-minutes>
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```



It is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summer time, please see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3\(x\)](#). Sample clock commands for the United States Eastern timezone are:

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

Create VLANs in Cisco Nexus 93180YC-FX A and 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 in Cisco Nexus 93180YC-FX 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
```

Add NTP Distribution Interface in Cisco Nexus 93180YC-FX B

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

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

Add Port Profiles in Cisco Nexus 93180YC-FX A and B

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. To add port profiles, run the following:

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
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
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 in Cisco Nexus 93180YC-FX A

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



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

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

```
interface Eth1/21
description <ucs-clustername>-a:1/45
udld enable
interface Eth1/22
description <ucs-clustername>-a:1/46
udld enable
interface Eth1/23
description <ucs-clustername>-b:1/45
udld enable
interface Eth1/24
description <ucs-clustername>-b:1/46
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. If you have fibre optic connections, do not enter the `udld enable` command.

```
interface Eth1/17
description <st-clustername>-01:e0e
interface Eth1/18
description <st-clustername>-01:e0f
interface Eth1/19
description <st-clustername>-02:e0e
interface Eth1/20
description <st-clustername>-02:e0f
interface Eth1/49
description <nexus-b-hostname>:1/49
interface Eth1/50
description <nexus-b-hostname>:1/50
exit
```

Add Individual Port Descriptions for Troubleshooting and Enable UDLD for UCS Interfaces in Cisco Nexus 93180YC-FX 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 Eth1/21
description <ucs-clustername>-a:1/47
udld enable
interface Eth1/22
description <ucs-clustername>-a:1/48
udld enable
interface Eth1/23
description <ucs-clustername>-b:1/47
udld enable
interface Eth1/24
description <ucs-clustername>-b:1/48
udld enable
```



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

```
interface Eth1/17
description <st-clustername>-01:e0g
interface Eth1/18
description <st-clustername>-01:e0h
interface Eth1/19
description <st-clustername>-02:e0g
interface Eth1/20
description <st-clustername>-02:e0h
interface Eth1/49
description <nexus-a-hostname>:1/49
interface Eth1/50
description <nexus-a-hostname>:1/50
exit
```


Create Port Channels in Cisco Nexus 93180YC-FX A and 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 Eth1/49-50
channel-group 10 mode active
no shutdown
interface Po117
description <st-clustername>-01
interface Eth1/17-18
channel-group 117 mode active
no shutdown
interface Po119
description <st-clustername>-02
interface Eth1/19-20
channel-group 119 mode active
no shutdown
interface Po121
description <ucs-clustername>-a
interface Eth1/21-22
channel-group 121 mode active
no shutdown
interface Po123
description <ucs-clustername>-b
interface Eth1/23-24
channel-group 123 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters in Cisco Nexus 93180YC-FX A and 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 Po117
inherit port-profile FP-ONTAP-Storage
interface Po119
inherit port-profile FP-ONTAP-Storage

interface Po121
inherit port-profile FP-UCS
interface Po123
inherit port-profile FP-UCS

exit
copy run start
```

Configure Virtual Port Channels in Cisco Nexus 93180YC-FX A

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

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

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po121
vpc 121
interface Po123
vpc 123
exit
copy run start
```

Configure Virtual Port Channels in Cisco Nexus 93180YC-FX B

To configure vPCs for switch B, follow this step:

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

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po121
vpc 121
interface Po123
vpc 123
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.

FlexPod with Nexus 93180YC-FX SAN Switching Configuration - Part 2

If the Nexus 93180YC-FX switch is being used for SAN Switching, this section should be completed in place of the Cisco MDS section of this document.

Configure Fibre Channel Ports in Cisco Nexus 93180YC-FX A

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

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

```
interface fc1/1
switchport description <st-clustername>-01:5a
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description <st-clustername>-02:5a
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description <ucs-clustername>-a:1/3
port-license acquire
channel-group 15
no shutdown
exit

interface fc1/6
switchport description <ucs-clustername>-a:1/4
port-license acquire
channel-group 15
no shutdown
exit

interface san-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.

Configure Fibre Channel Ports in Cisco Nexus 93180YC-FX B

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

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

```
interface fc1/1
switchport description <st-clustername>-01:5b
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
```

```

switchport description <st-clustername>-02:5b
port-license acquire
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description <ucs-clustername>-b:1/3
port-license acquire
channel-group 15
no shutdown
exit

interface fc1/6
switchport description <ucs-clustername>-b:1/4
port-license acquire
channel-group 15
no shutdown
exit

interface san-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 Nexus 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 in Cisco Nexus 93180YC-FX A

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

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

```

vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
vsan <vsan-a-id> interface fc1/1
Traffic on fc1/1 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-a-id> interface fc1/2
Traffic on fc1/2 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-a-id> interface san-port-channel15
exit
copy run start

```

Create VSANs in Cisco Nexus 93180YC-FX B

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

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

```

vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
vsan <vsan-b-id> interface fc1/1
Traffic on fc1/1 may be impacted. Do you want to continue? (y/n) [n] y

```

```
vsan <vsan-b-id> interface fcl/2
Traffic on fcl/2 may be impacted. Do you want to continue? (y/n) [n] y
vsan <vsan-b-id> interface san-port-channel15
exit
copy run start
```

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 uddl neighbors
show int status
show int brief
show flogi database
```

Cisco Nexus 93180YC-FX Zoning

This section explains how to configure Device Aliases, Zoning, and Zonesets in the Cisco Nexus 93180YC-FX switches for use in a FlexPod environment. Follow the steps precisely because failure to do so could result in an improper configuration.

Create Device Aliases in Cisco Nexus 93180YC-FX A

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

1. Login as admin and from the global configuration mode, run the following commands:

```
config t
device-alias mode enhanced
device-alias database
device-alias name Infra-SVM-fcp-lif01a pwn <fcp-lif01a-wwpn>
device-alias name Infra-SVM-fcp-lif02a pwn <fcp-lif02a-wwpn>
device-alias name VM-Host-Infra-01-A pwn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwn <vm-host-infra-02-wwpna>
device-alias name VM-Host-Infra-03-A pwn <vm-host-infra-03-wwpna>
device-alias commit
show device-alias database
```

Create Device Aliases in Cisco Nexus 93180YC-FX B

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

1. Login as admin and from the global configuration mode, run the following commands:

```
config t
device-alias mode enhanced
device-alias database
```

```
device-alias name Infra-SVM-fcp-lif01b pwnn <fcp-lif01b-wwpn>
device-alias name Infra-SVM-fcp-lif02b pwnn <fcp-lif02b-wwpn>
device-alias name VM-Host-Infra-01-B pwnn <vm-host-infra-01-wwpnb>
device-alias name VM-Host-Infra-02-B pwnn <vm-host-infra-02-wwpnb>
device-alias name VM-Host-Infra-03-B pwnn <vm-host-infra-03-wwpnb>
device-alias commit
show device-alias database
```

Create Zones and Zoneset in Cisco Nexus 93180YC-FX A

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

```
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-01a target
member device-alias Infra-SVM-fcp-lif-02a 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
```

Create Zones and Zoneset in Cisco Nexus 93180YC-FX B

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

```
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-01b target
member device-alias Infra-SVM-fcp-lif-02b 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
```

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:

```
config t
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:

```
port-profile type port-channel vPC-Peer-Link
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
port-profile type port-channel FP-ONTAP-Storage
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
port-profile type port-channel FP-UCS
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 vservers configured in a previous step. If FCP will be used in addition to iSCSI or in the future, step 1 can be omitted.

1. Remove FCP protocol from the vservers.

```
vserver remove-protocols -vserver <infra-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-01a -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-01b -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-02a -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-02b -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



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.

- Keep the Common/Global option selected for the scope of the VLAN.
- Enter the Infra-iSCSI-A VLAN ID.
- 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

- Click OK and then click OK again.
- Right-click VLANs.
- Choose Create VLANs.
- Enter Infra-iSCSI-B as the name of the VLAN to be used for iSCSI-B.
- Keep the Common/Global option selected for the scope of the VLAN.
- Enter the iSCSI-B VLAN ID.
- 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-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

16. Click OK to complete creating the vNIC template.
17. Click OK.
18. Right-click vNIC Templates.
19. Choose Create vNIC Template.

-
20. Enter iSCSI-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-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-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.

Create iSCSI Boot Policy

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi-lif01a and iscsi-lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi-lif02a and iscsi-lif02b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).



One boot policy is configured in this procedure. The policy configures the primary target to be iscsi-lif01a.

To create a boot policy for the Cisco UCS environment 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 Uefi Boot Mode.
 9. Check the checkbox for Boot Security.
 10. Expand the Local Devices drop-down menu and click Add Remote CD/DVD.
 11. Expand the iSCSI vNICs drop-down menu and click Add iSCSI Boot.
 12. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.
 13. Click OK.
 14. Choose Add iSCSI Boot.
 15. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-B.
 16. Click OK.
 17. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:

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

Boot Order

Name	O...	vNIC/vHBA/iSCSI...	Type	LUN...	WWN	Slot ...	Boot...	Boot...	Des...
Remote CD/DVD									
	1								
iSCSI									
	2								
iSCSI		iSCSI-Boot-A	Pri...						
iSCSI		iSCSI-Boot-B	Sec...						
CIMC Mounted CD/DVD									
	3								

18. 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 AMD-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.

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-NA-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 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.
3. 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](#)

4. Click Next.

Configure Storage Options

To configure the storage options, follow these steps:

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

Configure Zoning Options

To configure the zoning options, follow this step:

1. Make no changes and click Next.

Configure vNIC/HBA Placement

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

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

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

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

Configure Server Boot Order

To configure the server boot orders, follow these steps:

1. 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: **Boot-iSCSI-A** [Create Boot Policy](#)

Name : **Boot-iSCSI-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 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-01a for the IPv4 Address field.

Create iSCSI Static Target



iSCSI Target Name	:	<input type="text" value="iqn.1992-08.com.netapp::"/>	
Priority	:	<input type="text" value="1"/>	
Port	:	<input type="text" value="3260"/>	
Authentication Profile	:	<input type="text" value="<not set> ▼"/>	Create iSCSI Authentication Profile
IPv4 Address	:	<input type="text" value="192.168.10.61"/>	
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-02a for the IPv4 Address field.

15. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters



You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy:

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

[Create IP Pool](#)

[Reset Initiator Address](#)

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

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPv4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.10.61	0
iqn.1992-08....	2	3260		192.168.10.62	0

[+](#) Add [🗑](#) Delete [i](#) Info

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

OK

Cancel

16. Click OK to complete setting the iSCSI Boot Parameters.

-
17. In the Boot order, 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-01b 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-02b for the IPv4 Address field.
 30. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters



You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy:

IPv4 Address : **0.0.0.0**

Subnet Mask : **255.255.255.0**

Default Gateway : **0.0.0.0**

Primary DNS : **0.0.0.0**

Secondary DNS : **0.0.0.0**

[Create IP Pool](#)

[Reset Initiator Address](#)

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

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.20.61	0
iqn.1992-08....	2	3260		192.168.20.62	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 : **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 “UCS-C125” for the Server Pool Qualification to select only UCS C125 servers in the pool.
4. Expand Firmware Management at the bottom of the page and choose the default policy.

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](#)

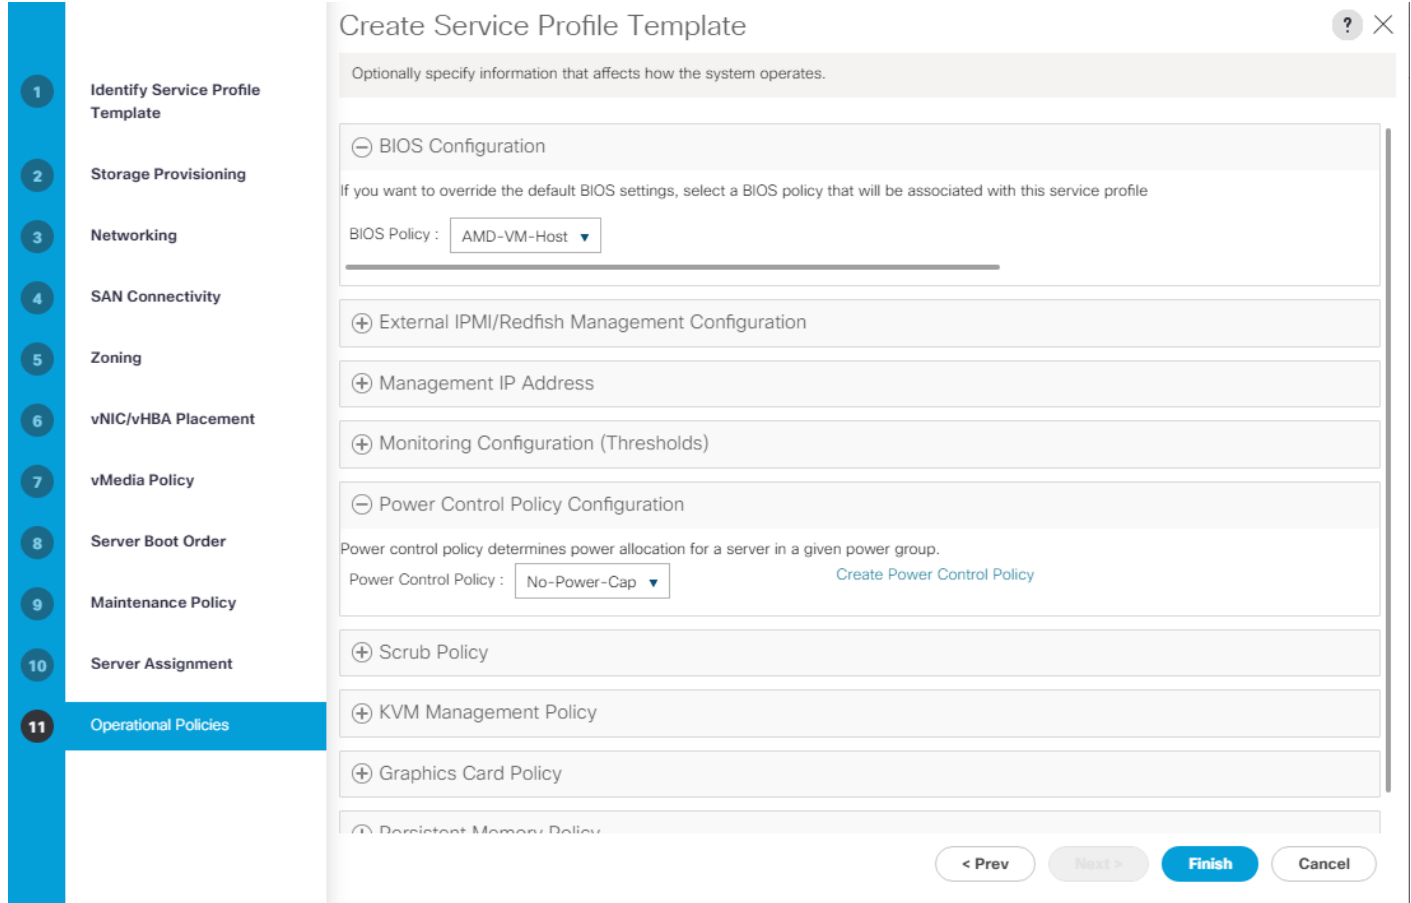
< Prev Next > **Finish** Cancel

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

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



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

Create 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 AMD-VM-Host-Infra-iSCSI-A.
3. Right-click AMD-VM-Host-Infra-iSCSI-A and click Create a Clone.
4. Name the clone AMD-VM-Host-Infra-iSCSI-A-vM and click OK then click OK again to create the clone.
5. Choose the newly created AMD-VM-Host-Infra-iSCSI-A-vM and choose the vMedia Policy tab.
6. Click Modify vMedia Policy.
7. Choose the ESXi-7.0-HTTP vMedia Policy and click OK.

8. Click OK to confirm.

Create Service Profiles

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

1. Connect to Cisco UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template AMD-VM-Host-Infra-iSCSI-A-VM.
3. Right-click AMD-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 ? ×

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :



7. Click OK to create the service profiles.
8. Click OK in the confirmation message.



When VMware ESXi 7.0 has been installed on the hosts, the host Service Profiles can be bound to the AMD-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. To create igroups, follow these steps:

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>
VM-Host-Infra-02		<vm-host-infra-02-iqn>
VM-Host-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.”

1. 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.

2. 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.


Add uplink	
MTU	<input type="text" value="9000"/>
Uplink 1	<input type="text" value="vmnic4 - Up, 50000 mbps"/>
▶ Link discovery	Click to expand
▶ Security	Click to expand
▶ NIC teaming	Click to expand
▶ Traffic shaping	Click to expand


Save Cancel

6. Click Save to save the changes to iScsiBootvSwitch.
7. Choose Add standard virtual switch.
8. Name the switch vSwitch1.
9. Change the MTU to 9000.

10. From the drop-down list select vmnic5 for Uplink 1.

 Add standard virtual switch - vSwitch1

 Add uplink

vSwitch Name	<input type="text" value="vSwitch1"/>
MTU	<input type="text" value="9000"/>
Uplink 1	<input type="text" value="vmnic5 - Up, 50000 mbps"/> 
▶ Link discovery	Click to expand
▶ Security	Click to expand

11. Choose Add to add vSwitch1.

12. In the center pane, choose the VMkernel NICs tab.

13. Highlight the iScsiBootPG line.


14. Choose Edit settings.

15. Change the MTU to 9000.

16. 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.

 Edit settings - vmk1

Port group	iScsiBootPG
MTU	9000
IP version	IPv4 only
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	192.168.10.193
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

17. Click Save to save the changes to iScsiBootPG VMkernel NIC.
18. Choose Add VMkernel NIC.
19. For New port group, enter iScsiBootPG-B.
20. For Virtual switch, use the pull-down to choose vSwitch1.
21. Change the MTU to 9000.
22. For IPv4 settings, choose Static.
23. Expand IPv4 Settings and enter a unique IP address in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B.

 Add VMkernel NIC

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	vSwitch1
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

24. Click Create to complete creating the VMkernel NIC.

25. In the center pane, choose the Port groups tab.

26. Highlight the iScsiBootPG line.

27. Choose Edit settings.

28. Change the Name to iScsiBootPG-A.

29. Click Save to complete editing the port group name.
30. Click Storage, then in the center pane choose the Adapters tab.
31. Click Software iSCSI to configure software iSCSI for the host.
32. In the Configure iSCSI window, under Dynamic targets, click Add dynamic target.
33. Choose to add address and enter the IP address of iscsi-lif-01a from storage SVM Infra-SVM. Press Return.
34. Repeat steps 32-33 to add the IP addresses for iscsi-lif-02a, iscsi-lif-01b, and iscsi-lif-02b.
35. Click Save configuration.
36. Click Software iSCSI to configure software iSCSI for the host.
37. Verify that four static targets and four dynamic targets are listed for the host.

Configure iSCSI - vmhba64

iSCSI enabled Disabled Enabled

▶ Name & alias: iqn.2010-11.com.flexpod:ucs-host:1

▶ CHAP authentication: Do not use CHAP

▶ Mutual CHAP authentication: Do not use CHAP

▶ Advanced settings: Click to expand

Network port bindings

VMkernel NIC	Port group	IPv4 address
No port bindings		

Static targets

Target	Address	Port
iqn.1992-08.com.netapp:sn.3333eff6ca4711eaa866d039ea1...	192.168.10.61	3260
iqn.1992-08.com.netapp:sn.3333eff6ca4711eaa866d039ea1...	192.168.20.62	3260
iqn.1992-08.com.netapp:sn.3333eff6ca4711eaa866d039ea1...	192.168.10.62	3260
iqn.1992-08.com.netapp:sn.3333eff6ca4711eaa866d039ea1...	192.168.20.61	3260

Dynamic targets

Address	Port
192.168.10.61	3260
192.168.10.62	3260
192.168.20.61	3260
192.168.20.62	3260

38. 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. Choose ... > Edit Settings to the right of iScsiBootPG. Change the Network label to iScsiBootPG-A and click OK.
6. Choose ... > Edit Settings to the right of the VMkernel Port IP address. Change the MTU to 9000.
7. 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.

8. Click OK.
9. In the upper right-hand corner, choose ADD NETWORKING to add another vSwitch.
10. Make sure VMkernel Network Adapter is selected and click NEXT.
11. Choose New standard switch and change the MTU to 9000. Click NEXT.
12. Choose  to add an adapter. Make sure vmnic5 is highlighted and click OK. vmnic5 should now be under Active adapters. Click NEXT.
13. Enter iScsiBootPG-B for the Network label, leave VLAN ID set to None (0), choose Custom - 9000 for MTU, and click NEXT.
14. 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. Click NEXT.

-
15. Click FINISH to complete creating the vSwitch and the VMkernel port.
 16. In the list under Storage, choose Storage Adapters.
 17. Choose the iSCSI Software Adapter and below, choose the Dynamic Discovery tab.
 18. Click Add.
 19. Enter the IP address of the storage controller's Infra-SVM LIF iscsi-lif-01a and click OK.
 20. Repeat this process to add the IPs for iscsi-lif-02a, iscsi-lif-01b, and iscsi-lif-02b.
 21. Under Storage Adapters, click Rescan Adapter to rescan the iSCSI Software Adapter.
 22. Under Static Discovery, four static targets should now be listed.
 23. Under Paths, four paths should now be listed with two of the paths having the "Active (I/O)" Status.

Create a FlexPod ESXi Custom ISO using VMware vCenter

In this validation document, the [Cisco Custom ISO for UCS 4.1.2a](#) was used to install VMware ESXi. After this installation the Cisco UCS Tools and the NetApp NFS Plug-in for VMware VAAI had to be installed during the FlexPod deployment. vCenter 7.0b or later can be used to produce a FlexPod custom ISO containing the updated UCS Tools and the NetApp NFS Plug-in for VMware VAAI. This ISO can be used to install VMware ESXi 7.0 without having to do any additional driver updates. To create the FlexPod ESXi custom ISO, follow these steps:



The Cisco Custom ISO for UCS 4.1.2a should also be used for Cisco UCS software release 4.1(2b) and VMware vSphere 7.0.

1. Download the [Cisco Custom Offline Bundle](#) for UCS 4.1.2a. This file (VMware-ESXi-7.0.0-16324942-Custom-Cisco-4.1.2a-Bundle.zip) can be used to produce the FlexPod ESXi 7.0 CD ISO.
2. Download the following listed .zip files:
 - [UCS Tools Component for ESXi 7.0 1.1.5](#) (ucs-tool-esxi_1.1.5-1OEM.zip)
 - [NetApp NFS Plug-in for VMware VAAI 1.1.2](#) - NetAppNasPlugin.v23.zip
3. Log into the VMware vCenter HTML5 Client as administrator@vsphere.local.
4. Under Menu, choose Auto Deploy.
5. If you receive the message "Auto Deploy and Image Builder are disabled in this vCenter", click ENABLE IMAGE BUILDER.



Auto Deploy and Image Builder are disabled in this vCenter.

To access full-featured auto deploy, enable both Image Builder and Auto Deploy.

To manage software depots only, enable Image Builder.

ENABLE AUTO DEPLOY AND IMAGE BUILDER

ENABLE IMAGE BUILDER

6. Click IMPORT to upload a software depot.
7. Name the depot Cisco Custom ESXi 7.0 for UCS 4.1(2a). Click BROWSE. Browse to the local location of the VMware-ESXi-7.0.0-16324942-Custom-Cisco-4.1.2a-Bundle.zip file downloaded above, highlight it, and click Open.

Import Software Depot



Name *

File * [BROWSE](#)

CANCEL

UPLOAD

8. Click UPLOAD to upload the software depot.
9. Repeat steps 6-8 to add software depots for ucs-tool-esxi_1.1.5-1OEM, and NetAppNasPlugin-v23.
10. Click NEW to add a custom software depot.
11. Choose Custom depot and name the custom depot FlexPod-ESXi-7.0.

Add Software Depot



Online depot

Name:

URL:

Custom depot

Name: *

CANCEL

ADD

12. Click ADD to add the custom software depot.

13. From the drop-down list, choose the Cisco Custom ESXi-7.0 for UCS 4.1(2a) (ZIP) software depot. Make sure the Image Profiles tab is selected and then click the radio button to select the Cisco-UCS-Custom-ESXi-7-1632492_4.1.2-a image profile. Click CLONE to clone the image profile.

14. Name the clone FlexPod-ESXi-7.0. For Vendor, enter Cisco-NetApp. For Description, enter "Cisco Custom ISO ESXi 7.0 for UCS 4.1(2a) with UCS Tool-1.1.5 and NetAppNasPlugin-v23". Choose FlexPod-ESXi-7.0 for Software depot.

Clone Image Profile

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

Name and details

Name * FlexPod-ESXi-7.0

Vendor * Cisco-NetApp

Description

Cisco Custom ISO [ESXi 7.0](#) for [UCS 4.1\(2a\)](#) with [UCS Tool-1.1.5](#) and [NetAppNasPlugin-v23](#)

Software depot * FlexPod-ESXi-7.0 ⓘ

CANCEL NEXT

15. Click NEXT.

16. Under Available software packages, check NetAppNasPlugin 1.1.2-3 and ucs-tool-esxi 1.1.5-1OEM. Uncheck ucs-tool-esxi 1.1.2-1OEM. 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"/>	lsuv2-smartpqi2...	1.0.0-3vmw.700.1.0.158...	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	mtip32xx-native	3.9.8-1vmw.700.1.0.158...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	native-misc-drive...	7.0.0-1.25.16324942	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	ne1000	0.8.4-10vmw.700.1.0.15...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nenic	1.0.33.0-1OEM.670.0.0...	VMware certified	Cisco	Cisco Custom ESXi 7.0 for
<input type="checkbox"/>	nenic	1.0.29.0-1vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	NetAppNasPlugin	1.1.2-3	VMware accepted	NetApp	NetAppNasPlugin-v23
<input checked="" type="checkbox"/>	nfnic	4.0.0.56-1OEM.670.0.0...	VMware certified	Cisco	Cisco Custom ESXi 7.0 for
<input type="checkbox"/>	nfnic	4.0.0.44-1vmw.700.1.0...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nhpsa	2.0.50-1vmw.700.1.0.15...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nmlx4-core	3.19.16.7-1vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nmlx4-en	3.19.16.7-1vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nmlx4-rdma	3.19.16.7-1vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nmlx5-core	4.19.16.7-1vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nmlx5-rdma	4.19.16.7-1vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	ntg3	4.1.4.1-1vmw.700.1.0.158...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nvme-pcie	1.2.2.14-1vmw.700.1.25.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	nvmerdma	1.0.0.0-1vmw.700.1.0.15...	VMware certified	VMW	Cisco Custom ESXi 7.0 for

76 selected of 87 items

CANCEL

BACK

NEXT

Clone Image Profile

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

Select software packages X

Acceptance level Partner supported v

<input type="checkbox"/>	Name	Version	Acceptance Level	Vendor	Depot
<input checked="" type="checkbox"/>	qfle3i	2.1.2.0-1OEM.700.1.0.15...	VMware certified	QLC	Cisco Custom ESXi 7.0 for
<input type="checkbox"/>	qfle3i	1.0.15.0-6vmw.700.1.0.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	qflge	1.1.0.11-1vmw.700.1.0.15...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	qlnativefc	4.0.1.0-3vmw.700.1.0.1...	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	rste	2.0.2.0088-7vmw.700.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	sfvmk	2.0.0.1004-3vmw.700.1...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	smartpqi	1.0.4.3011-1vmw.700.1.0...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	tools-light	11.1.0.16036546-16321839	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	ucs-tool-esxi	1.1.5-1OEM	Partner supported	CIS	ucs-tool-esxi_1.1.5-1OEM
<input type="checkbox"/>	ucs-tool-esxi	1.1.2-1OEM	Partner supported	CIS	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vdfs	7.0.0-1.25.16324942	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vmkata	0.1-1vmw.700.1.0.15843...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vmkfcoe	1.0.0.2-1vmw.700.1.0.15...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vmkusb	0.1-1vmw.700.1.25.1632...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vmw-ahci	1.3.9-1vmw.700.1.0.1584...	VMware certified	VMW	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vmware-esx-esx...	1.2.0.37-1vmw.700.1.0.1...	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vsan	7.0.0-1.25.16324942	VMware certified	VMware	Cisco Custom ESXi 7.0 for
<input checked="" type="checkbox"/>	vsanhealth	7.0.0-1.25.16324942	VMware certified	VMware	Cisco Custom ESXi 7.0 for

76 selected of 87 items

CANCEL
BACK
NEXT

17. Click NEXT.

18. Click FINISH.

19. From the Software Depot drop-down list, choose the FlexPod-ESXi-7.0 (Custom) software depot. Under Image Profiles choose the FlexPod-ESXi-7.0 image profile. Click EXPORT to export an image profile. The ISO should be highlighted. Click OK to generate a bootable ESXi installable image.

20. Once the Image profile export completes, click DOWNLOAD to download the ISO.

21. Once downloaded, you can rename the ISO to a more descriptive name.

22. Optionally, generate the ZIP archive to generate an offline bundle for the FlexPod image using ... -> Export.

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:
 1. Hostname: <IP or FQDN of host that will receive the backup>
 2. Protocol: [FTP/TFTP/SCP/SFTP]
 3. User: <account on host to authenticate>
 4. Password: <password for account on host>
 5. Remote File: <full path and filename prefix for backup file>



Admin State must be Enabled to fill in the Remote File field.

6. Admin State: <choose Enable to activate the schedule on save, Disable to disable schedule on Save>
7. 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

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 the following:

1. The Backup location with the protocol to use [HTTPS/HTTP/SCP/FTPS/FTP]
2. The User name and password.
3. The Number of backups to retain.

Create Backup Schedule

Backup location ⓘ	<input type="text" value="scp://nx-ftp.flexpod.cisco.com/var/www/html/software/Configs/nx-vc/"/>		
Backup server credentials	User name	<input type="text" value="admin"/>	
	Password	<input type="password" value="*****"/>	
Schedule ⓘ	Daily ▾	02 : 15	A.M. America/New_York
Encrypt backup (optional)	Encryption Password	<input type="password"/>	
	Confirm Password	<input type="password"/>	
Number of backups to retain	<input type="radio"/> Retain all backups		
	<input checked="" type="radio"/> Retain last <input type="text" value="7"/> 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

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">• Inventory and configuration• Stats, Events, and Tasks
Number of backups to retain	7

Activity [BACKUP NOW](#)

The Backup Schedule should now show a Status of Enabled.

Restoration can be initiated with the backed-up files using the Restore function of the VCSA 7.0 Installer.

About the Authors

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

John has been involved in designing, developing, validating, and supporting the FlexPod Converged Infrastructure since it was developed almost nine years ago. Before his roles with FlexPod, he supported and administered a large worldwide training network and VPN infrastructure. John holds a master's degree in Computer Engineering from Clemson University.

Sree Lakshmi Lanka, Solutions Engineer, Infrastructure and Cloud Engineering, NetApp

Sree is a Senior Solutions Architect at NetApp. She has more than 12 years of experience in data center infrastructure solutions, both in traditional and in hybrid/public cloud space. She collaborates with marketing, product management and engineering teams to develop and deliver technical marketing product material, which includes reference architectures, technical reports, presentations, blogs, demo videos and white papers. This material is aimed at educating customers, partners, or sales team. She has a bachelor's degree in Computer Science and an artist in the field of Kuchipudi dance, an Indian classical dance form from Andhra Pradesh. She enjoys organic gardening, hiking, and running.

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.

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](https://cs.co/en-cvds) at <https://cs.co/en-cvds>.

Americas Headquarters

Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters

Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)