

FlexPod Datacenter with Cisco ACI and VMware vSphere 6.0 U1

Deployment Guide for FlexPod Datacenter with Cisco ACI and VMware vSphere 6.0 U1

Last Updated: August 24, 2016



About Cisco Validated Designs

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

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

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

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

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

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE 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 NETAPP, ITS PARTNERS OR THEIR RESPECTIVE 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, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE 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 AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS,

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)

© 2016 Cisco Systems, Inc. All rights reserved.

Table of Contents

About Cisco Validated Designs	2
Executive Summary	9
Solution Overview.....	10
Introduction	10
Audience	10
Purpose of this Document.....	10
Solution Design.....	11
Architecture	11
Physical Topology.....	11
Deployment Hardware and Software	13
Software Revisions	13
Configuration Guidelines.....	13
Physical Infrastructure.....	15
FlexPod Cabling	15
Storage Configuration	23
Controller AFF80XX Series	23
NetApp Hardware Universe	23
Controllers.....	23
Disk Shelves	24
Clustered Data ONTAP 8.3.2	24
Complete the Configuration Worksheet	24
Configure ONTAP Nodes	24
Log in to Cluster	37
Zero All Spare Disks	37
Set Onboard UTA2 Ports Personality	37
Set Auto-Revert on Cluster Management	38
Set Up Management Broadcast Domain	38
Set Up Service Processor Network Interface	39
Create Aggregates	39
Verify Storage Failover.....	40
Disable Flow Control on 10GE Ports	41
Disable Unused FcoE Ports.....	41
Configure NTP	41

Configure SNMP	42
Configure AutoSupport	43
Enable Cisco Discovery Protocol	43
Create Broadcast Domains in ONTAP	43
Create Interface Groups	43
Create VLANs	43
Create Storage Virtual Machine	44
Create Load-Sharing Mirrors of SVM Root Volume	45
Configure HTTPS Access	45
Configure NFSv3	47
Create NetApp FlexVol Volumes	47
Create Boot LUNs	47
Adjust Storage Efficiency Settings	48
Create iSCSI LIFs	48
Create FCoE LIFs	48
Create NFS LIFs	49
Add Infrastructure SVM Administrator	49
Server Configuration	51
Cisco UCS Base Configuration	51
Perform Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments	51
Cisco UCS Setup	52
Log in to Cisco UCS Manager	52
Upgrade Cisco UCS Manager Software to Version 3.1(1h)	53
Configure Cisco UCS Call Home	53
Add Block of IP Addresses for KVM Access	54
Synchronize Cisco UCS to NTP	55
Change Fabric Interconnect FC Mode to Switching	55
Edit Chassis Discovery Policy	56
Enable Server, Uplink, and Storage Ports	56
Acknowledge Cisco UCS Chassis and FEX	58
Create Uplink Port Channels to Cisco Nexus Switches	58
Create MAC Address Pools	60
Create a WWNN Address Pool	62
Create a WWPN Address Pools	63
Create IQN Pools for iSCSI Boot	65

Create IP Pools for iSCSI Boot	66
Create UUID Suffix Pool	68
Create Server Pool	69
Create VLANs	69
Create VSANs and Configure FCoE Storage Ports	77
Modify Default Host Firmware Package	82
Set Jumbo Frames in Cisco UCS Fabric	83
Create Local Disk Configuration Policy (Optional)	84
Create Network Control Policy for Cisco Discovery Protocol and Link Layer Discovery Protocol	85
Create Power Control Policy	86
Create Server Pool Qualification Policy (Optional)	87
Create Server BIOS Policy	88
Update Default Maintenance Policy	89
Create vMedia Policy for VMware ESXi 6.0U1B Install Boot	90
Create vNIC Templates	91
Create LAN Connectivity Policies	103
Create vHBA Templates	115
Create Storage Connection Policies	117
Create a SAN Connectivity Policy	119
Create FCoE Boot Policy	125
Create iSCSI Boot Policy	128
Create Infrastructure FCoE Boot Service Profile Template	129
Create iSCSI Boot Service Profile Template	140
Create Service Profiles	151
Add More Servers to FlexPod Unit	151
Gather Necessary Information	152
Storage Configuration - SAN Boot	153
Clustered Data ONTAP iSCSI Boot Storage Setup	153
Create igroups	153
Clustered Data ONTAP FCoE Boot Storage Setup	153
Create igroups	153
Map Boot LUNs to igroups	154
Cisco ACI Fabric Configuration	155
Physical Connectivity	155
Cisco Application Policy Infrastructure Controller (APIC) Setup	155

Cisco ACI Fabric Discovery.....	157
Initial ACI Fabric Setup.....	160
Fabric Access Policy Setup.....	166
Create Virtual Port Channels (vPCs).....	174
Create In-Band Management External Bridged Network and Core-Services EPG	185
Create Security Filters in Tenant common.....	198
Deploy Infrastructure (Foundation) Tenant	201
VMware vSphere 6.0 U1b Setup.....	224
VMware ESXi 6.0 U1b.....	224
Log in to Cisco UCS 6200 Fabric Interconnect.....	224
Install ESXi.....	225
Set Up Management Networking for ESXi Hosts	225
Download VMware vSphere Client.....	227
Log in to VMware ESXi Hosts by Using VMware vSphere Client.....	228
Set Up VMkernel Ports and Virtual Switch.....	228
Install VMware Drivers for the Cisco Virtual Interface Card (VIC).....	239
Mount Required Datastores	241
Configure NTP on ESXi Hosts	243
Move VM Swap File Location.....	244
VMware vCenter 6.0U1b.....	245
Install the Client Integration Plug-in	245
Building the VMware vCenter Server Appliance	246
Setting Up VMware vCenter Server	254
ESXi Dump Collector Setup for iSCSI-Booted Hosts	260
Add Active Directory (AD) Servers to Core-Services Network.....	261
Add AD User Authentication to vCenter (Optional)	261
Add vSphere Distributed Switch (vDS).....	263
Add vDS in APIC.....	263
Add VMware ESXi Host Servers to vDS	267
Create In-Band Management Port-Profile on vDS.....	269
Add Cisco Application Virtual Switch (AVS)	273
Install Cisco Virtual Switch Update Manager (VSUM) Virtual Appliance	273
Add Cisco AVS in APIC	275
Add VMware ESXi Host Servers to AVS.....	280
Add Second VXLAN Tunnel Endpoint (VTEP) to Each ESXi Host for Load Balancing	282

Create In-Band Management Port-Profile on AVS	283
FlexPod Management Tools Setup.....	290
NetApp Virtual Storage Console 6.2P2 Deployment Procedure.....	290
Virtual Storage Console 6.2 Pre-installation Considerations	290
Install Virtual Storage Console 6.2P2	290
Register Virtual Storage Console with vCenter Server.....	292
Install NetApp NFS VAAI Plug-in.....	292
Discover and Add Storage Resources	293
Optimal Storage Settings for ESXi Hosts.....	293
Virtual Storage Console 6.2P2 Backup and Recovery	295
OnCommand Performance Manager 2.1	299
OnCommand Performance Manager Open Virtualization Format (OVF) Deployment	299
OnCommand Performance Manager Basic Setup	303
OnCommand Unified Manager 6.4	305
OnCommand Unified Manager OVF Deployment.....	305
OnCommand Unified Manager Basic Setup	311
Link OnCommand Performance Manager and OnCommand Unified Manager.....	314
Sample Tenant Setup	318
Add Supernet Routes to Core-Services Devices.....	318
Adding the Supernet Route in a Windows VM.....	318
Adding the Supernet Route in the vCenter Server Appliance	318
Adding the Supernet Route in VMware ESXi	318
ACI Shared Layer 3 Out Setup	319
Configuring the Cisco Nexus 7000s for ACI Connectivity (Sample).....	320
Configuring ACI Shared Layer 3 Out	323
Lab Validation Tenant Configuration.....	339
Configure Tenant Storage.....	340
Create Tenant IPspace	340
Create Tenant Broadcast Domains in ONTAP	340
Create VLAN Interfaces	340
Create Tenant Storage Virtual Machine.....	341
Create Load-Sharing Mirrors of SVM Root Volume	342
Configure HTTPS Access	342
Configure NFSv3	343
Create FlexVol Volumes.....	344

Adjust Storage Efficiency Settings	344
Create iSCSI LIFs.....	344
Create FCoE LIFs.....	345
Create NFS LIF	345
Add Tenant SVM Administrator.....	345
Add Quality of Service (QoS) Policy to Monitor Application Workload.....	346
Configure Cisco UCS for the Tenant	346
Add Tenant iSCSI VLANs.....	347
Add Tenant iSCSI VLANs to iSCSI vNIC Templates.....	348
Add Tenant SAN Connectivity Policy	349
Create Application-Specific Service Profile Templates.....	356
Add New Application-Specific Server Pool.....	356
Create New Service Profiles for Application-Specific Servers	357
Gather Necessary Information	357
Configure Storage SAN Boot for the Tenant.....	358
Clustered VMware ESXi Boot LUNs in Infra-SVM.....	358
Clustered Data ONTAP iSCSI Boot Storage Setup	358
Clustered Data ONTAP FCoE Boot Storage Setup.....	358
Map Boot LUNs to igroups.....	359
Deploy ACI Application (App-A) Tenant.....	359
Install and Configure VMware ESXi on Tenant Hosts.....	391
Build a Second Tenant (Optional).....	391
Deploy L4-L7 VLAN Stitching in Sample Tenants.....	392
Deploy Sample Cisco ASA VPCs	392
APIC Advanced GUI.....	392
Create Tenant Firewall Outside Bridge Domain and Subnet	396
APIC Advanced GUI.....	396
Create Tenant L4-L7 Device and Service Graph	399
APIC Advanced GUI.....	399
About the Authors.....	406
Acknowledgements	406



Executive Summary

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

This document describes the Cisco and NetApp® FlexPod Datacenter with NetApp All Flash FAS (AFF), Cisco Application Centric Infrastructure (ACI), and VMware vSphere 6.0 Update 1b. FlexPod Datacenter with NetApp AFF and Cisco ACI is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, and NetApp AFF.

Solution Overview

Introduction

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. Business agility requires application agility, so IT teams must provision applications in hours instead of months. Resources must scale up (or down) in minutes, not hours.

To simplify the evolution to a shared cloud infrastructure based on an application-driven policy model, Cisco and NetApp have developed a solution called FlexPod Datacenter with NetApp AFF and Cisco ACI. Cisco ACI provides a holistic architecture with centralized automation and policy-driven application profiles that delivers software flexibility with hardware performance. NetApp AFF addresses enterprise storage requirements with high performance, superior flexibility, and best-in-class data management.

Audience

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

Purpose of this Document

This document provides a step-by-step configuration and implementation guide for the FlexPod Datacenter with NetApp AFF and Cisco ACI solution. For the design decisions and technology discussion of the solution, see the [FlexPod Datacenter with Cisco ACI and VMware vSphere 6.0U1 Design Guide](#).

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

- Validation of the latest version of Cisco ACI with the latest version of the NetApp AFF storage array
- Validation of the Cisco ACI 1.3 release on Cisco Nexus 9000 Series switches
- Support for the Cisco UCS 3.1 release and Cisco UCS B200-M4 and C220-M4 servers with Intel E5-2600 v4 Series processors
- Support for NetApp Data ONTAP® 8.3.2
- A storage design supporting both NAS datastores and iSCSI and Fibre Channel over Ethernet (FCoE) SAN LUNs

Solution Design

Architecture

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and nonvirtualized solutions. VMware vSphere® built on FlexPod includes NetApp storage, NetApp ONTAP, NetApp AFF, Cisco Nexus® networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that networking, computing, and storage can fit in one data center rack or be deployed according to your data center design. The 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 your requirements. A FlexPod system can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an integrated infrastructure 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.

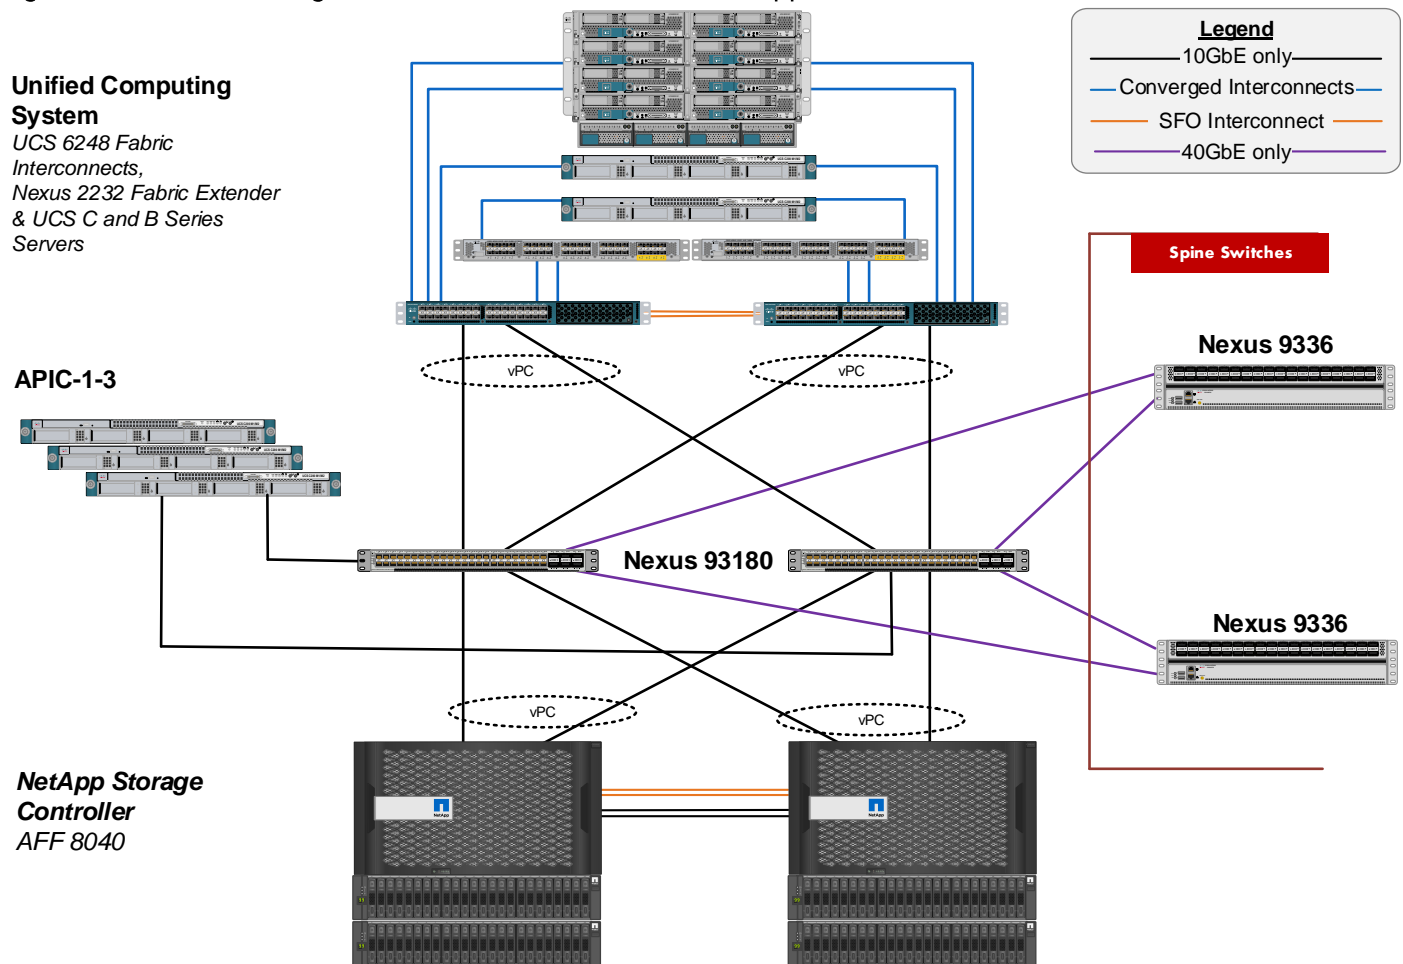
Figure 1 shows the FlexPod with Cisco ACI components and network connections for a configuration with IP-based storage. This design uses the Cisco Nexus 9000, the Cisco Application Policy Infrastructure Controller (APIC), Cisco UCS C-Series and B-Series servers, and the NetApp AFF family of storage controllers connected in a highly available modular design. This infrastructure can include FCoE-based storage and is deployed to provide iSCSI or FCoE-booted hosts with file-level and block-level access to shared storage. The reference architecture reinforces the wire-once strategy, because, as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

The ACI switching architecture is laid out in a leaf-and-spine topology where every leaf connects to every spine using 40G Ethernet interface(s). The software controller (the APIC) is delivered as an appliance, and three or more of these appliances form a cluster for high availability and enhanced performance.

Physical Topology

0 illustrates the physical architecture.

Figure 1 FlexPod Design with Cisco Nexus 9000 and NetApp ONTAP



The reference hardware configuration includes the following components:

- Two Cisco Nexus 93180YC-EX, 9372PX, or 9396 leaf switches
- Two Cisco Nexus 9336PQ spine switches
- Three Cisco APIC-M1s
- Two Cisco UCS 6248UP fabric interconnects
- One NetApp AFF8040 (an HA pair) running ONTAP with disk shelves and solid state drives (SSD)

For server virtualization, the deployment includes VMware vSphere 6.0 Update 1b. Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps needed to deploy the base architecture, as shown in 0. These procedures cover everything from physical cabling to network, compute, and storage-device configurations.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6200 Series, UCS B-200 M4, UCS C-220 M4	3.1(1h)	Includes the Cisco UCS-IOM 2208XP, Cisco UCS Manager, UCS VIC 1240 and UCS VIC 1340
	Cisco eNIC	2.3.0.7	
	Cisco fNIC	1.6.0.25	
Network	Cisco APIC	1.3(2f)	
	Cisco Nexus 9000 iNX-OS	11.3(2f)	
	Cisco Virtual Switch Update Manager (VSUM)	2.0	
	Cisco Application Virtual Switch (AVS)	5.2(1)SV3(1.25)	
Storage	NetApp AFF 8040	Data ONTAP 8.3.2	
Software	VMware vSphere ESXi	6.0u1b	
	VMware vCenter	6.0u1b	
	NetApp OnCommand® Unified Manager for Clustered Data ONTAP	6.4	
	NetApp Virtual Storage Console (VSC)	6.2	
	OnCommand Performance Manager	2.1	

Configuration Guidelines

This document provides details for configuring 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 or B). For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, and so on. Finally, to indicate that

you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the `network port vlan create` command:

Usage:

```
network port vlan create ?
[-node] <nodename>                Node
{ [-vlan-name] {<netport>|<ifgrp>} VLAN Name
| -port {<netport>|<ifgrp>}        Associated Network Port
[-vlan-id] <integer> }            Network Switch VLAN Identifier
```

Example:

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

This document enables you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 3 lists the virtual machines (VMs) necessary for deployment as outlined in this guide. [Error! Reference source not found.](#) describe the ACI End Point Groups (EPGs), VLANs, subnets, and ACI bridge domains that were deployed for the FlexPod infrastructure and two sample tenants as outlined in this guide.

Table 2 Lab Validation Infrastructure (Foundation) Tenant Configuration

EPG	Storage VLAN	UCS VLAN	External VLAN	Subnet / Gateway	Bridge Domain
IB-MGMT	N/A	DVS	163	172.26.163.0/24 - L2	BD-common-Internal
Core-Services	N/A	363	163	172.26.163.10/24	BD-common-Internal
SVM-MGMT	263	N/A	163	172.26.163.0/24 - L2	BD-common-Internal
iSCSI-A	3010	3110	N/A	192.168.110.0/24 - L2	BD-iSCSI-A
iSCSI-B	3020	3120	N/A	192.168.120.0/24 - L2	BD-iSCSI-B
NFS-LIF	3050	N/A	N/A	192.168.150.0/24 - L2	BD-NFS
NFS-VMK	N/A	3150	N/A	192.168.150.0/24 - L2	BD-NFS
vMotion	N/A	3000	N/A	192.168.100.0/24 - L2	BD-Internal
VMware vDS Pool	N/A	1101-1120	N/A	Varies	Varies
ACI System VLAN for AVS	N/A	4093	N/A	Varies	Varies

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

Table 3 Lab Validation Tenant App-A Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI-A	3011	3111	192.168.111.0/24 - L2	BD-iSCSI-A
iSCSI-B	3021	3121	192.168.121.0/24 - L2	BD-iSCSI-B
NFS-LIF	3051	N/A	192.168.151.0/24 - L2	BD-NFS
NFS-VMK	N/A	DVS	192.168.151.0/24 - L2	BD-NFS
SVM-MGMT	264	N/A	172.16.254.6/29	BD-Internal
Web	N/A	DVS	172.16.0.254/24	BD-Internal
App	N/A	DVS	172.16.1.254/24	BD-Internal

DB	N/A	DVS	172.16.2.254/24	BD-Internal
----	-----	-----	-----------------	-------------

Table 4 Lab Validation Tenant App-B Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI-A	3012	3112	192.168.111.0/24 - L2	BD-iSCSI-A
iSCSI-B	3022	3122	192.168.121.0/24 - L2	BD-iSCSI-B
NFS-LIF	3052	N/A	192.168.151.0/24 - L2	BD-NFS
NFS-VMK	N/A	DVS	192.168.151.0/24 - L2	BD-NFS
SVM-MGMT	265	N/A	172.16.254.14/29	BD-Internal
Web	N/A	DVS	172.16.3.254/24	BD-Internal
App	N/A	DVS	172.16.4.254/24	BD-Internal
DB	N/A	DVS	172.16.5.254/24	BD-Internal

When planning this FlexPod deployment, you must make several decisions. You must first determine which block-based storage protocols you wish to deploy. Both FCoE and iSCSI can be deployed for both application LUN access and SAN-boot. This document provides optional steps for both protocols. Perform the steps necessary for the storage protocols that you intend to install. The second decision to make is which Distributed Virtual Switch (DVS) to implement. The first of two choices in this document is the VMware vSphere Distributed Switch (vDS) that is included with VMware Enterprise Plus licensing and uses a pool of dynamic VLANs in the Cisco APIC. The second choice is the Cisco AVS in VXLAN local switching mode, which is free from Cisco with VMware Enterprise Plus licensing and includes VXLANs automatically assigned by the APIC. This document again has optional steps for both DVSs.

You should set up an HTTP server that is accessible from the out-of-band management network. At a minimum, the Data ONTAP 8.3.2 software file and the VMware ESXi 6.0U1b .iso file should be placed on this server.

Physical Infrastructure

FlexPod Cabling

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

The tables in this section contain the details for the prescribed and supported configuration of a NetApp AFF8040 running clustered Data ONTAP 8.3.2. For any modifications to 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 are used in various configuration steps.

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

Figure 2 shows a cabling diagram for a FlexPod configuration using the Cisco Nexus 9000 and NetApp storage systems with ONTAP. The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk-shelf cabling, refer to the [Universal SAS and ACP Cabling Guide](#).

Figure 2 FlexPod Cabling Diagram

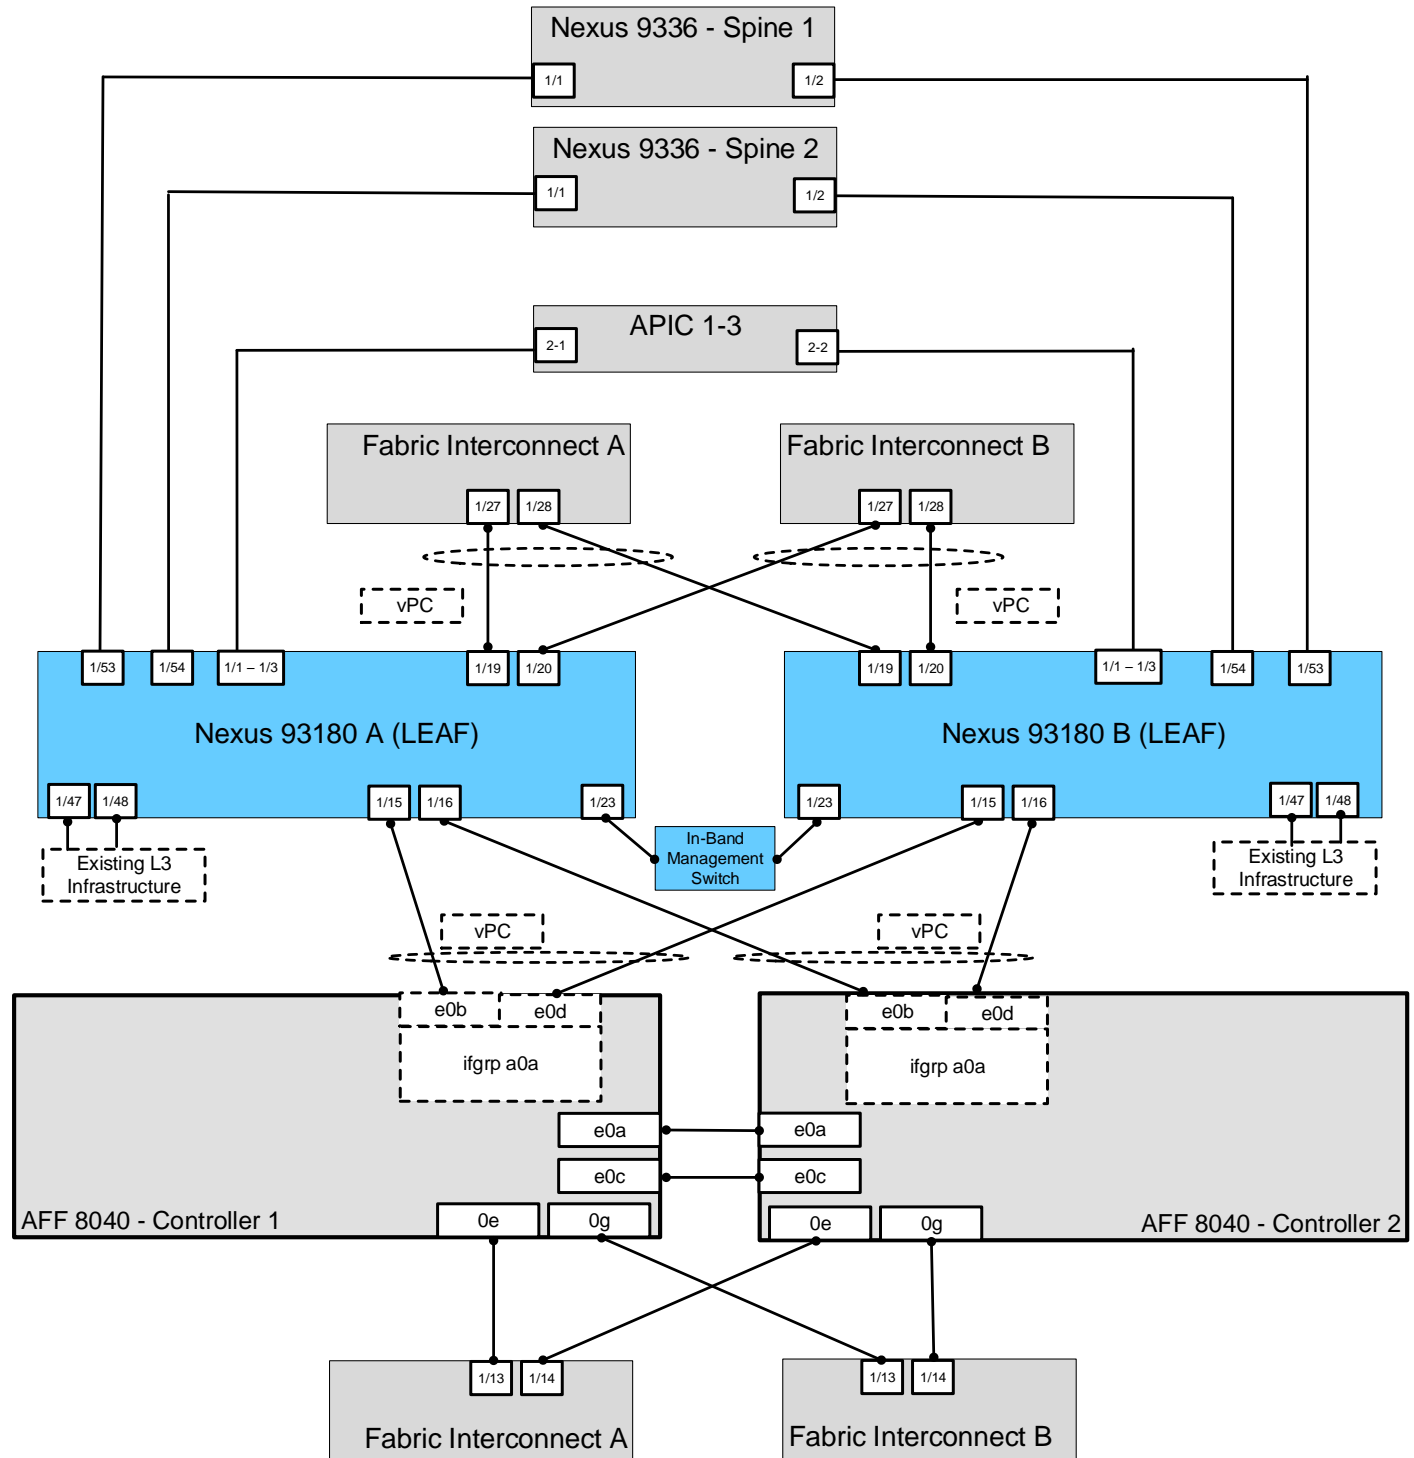


Table 5 through Table 14 describes all of the connections in use.

Table 5 Cisco Nexus 93180-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
--------------	------------	------------	---------------	-------------

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 A	Eth1/1	10GbE	APIC 1	Eth2/1
	Eth1/2	10GbE	APIC 2	Eth2/1
	Eth1/3	10GbE	APIC 3	Eth2/1
	Eth1/15	10GbE	NetApp controller 01	e0b
	Eth1/16	10GbE	NetApp controller 02	e0b
	Eth1/19	10GbE	Cisco UCS fabric interconnect A	Eth1/27
	Eth1/20	10GbE	Cisco UCS fabric interconnect B	Eth1/27
	Eth1/23	1GbE	In-band management switch	Any
	Eth1/47	10GbE	Nexus 7K A	Eth4/21
	Eth1/48	10GbE	Nexus 7K B	Eth4/21
	Eth1/53	40GbE	Cisco Nexus 9336 A	Eth1/1
	Eth1/54	40GbE	Cisco Nexus 9336 B	Eth1/1
	MGMT0	GbE	GbE management switch	Any



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

Table 6 Cisco Nexus 93180-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 B	Eth1/1	10GbE	APIC 1	Eth2/2
	Eth1/2	10GbE	APIC 2	Eth2/2
	Eth1/3	10GbE	APIC 3	Eth2/2
	Eth1/15	10GbE	NetApp controller 01	e0d
	Eth1/16	10GbE	NetApp controller 02	e0d
	Eth1/19	10GbE	Cisco UCS fabric interconnect A	Eth1/28
	Eth1/20	10GbE	Cisco UCS fabric interconnect B	Eth1/28
	Eth1/23	1GbE	In-band management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/47	10GbE	Cisco Nexus 7K A	Eth4/22
	Eth1/48	10GbE	Cisco Nexus 7K B	Eth4/22
	Eth1/53	40GbE	Cisco Nexus 9336 A	Eth1/2
	Eth1/54	40GbE	Cisco Nexus 9336 B	Eth1/2
	MGMT0	GbE	GbE management switch	Any

Table 7 Cisco Nexus 9336-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336 A	Eth1/1	40GbE	Cisco Nexus 93180 A	Eth1/53
	Eth1/2	40GbE	Cisco Nexus 93180 B	Eth1/53
	MGMT0	GbE	GbE management switch	Any

Table 8 Cisco Nexus 9336-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336 B	Eth1/1	40GbE	Cisco Nexus 93180 A	Eth1/54
	Eth1/2	40GbE	Cisco Nexus 93180 B	Eth1/54
	MGMT0	GbE	GbE management switch	Any

Table 9 NetApp Controller-01 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 01	e0M	GbE	GbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp controller 02	e0a
	e0b	10GbE	Cisco Nexus 93180 A	Eth1/15
	e0c	10GbE	NetApp controller 02	e0c
	e0d	10GbE	Cisco Nexus 93180 A	Eth1/15

Local Device	Local Port	Connection	Remote Device	Remote Port
	0e	10GbE	Cisco UCS fabric interconnect A	Eth1/13
	0g	10GbE	Cisco UCS fabric interconnect B	Eth1/13



Note: The term e0M refers to the physical Ethernet port labeled with a wrench icon on the rear of the chassis.

Table 10 NetApp Controller 02 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 02	e0M	GbE	GbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp controller 01	e0a
	e0b	10GbE	Cisco Nexus 93180 A	Eth1/16
	e0c	10GbE	NetApp controller 01	e0c
	e0d	10GbE	Cisco Nexus 93180 B	Eth1/16
	0e	10GbE	Cisco UCS fabric interconnect A	Eth1/14
	0g	10GbE	Cisco UCS fabric interconnect B	Eth1/14

Table 11 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/1	10GbE	Cisco UCS Chassis FEX A	IOM1/1
	Eth1/2	10GbE	Cisco UCS Chassis FEX A	IOM1/2
	Eth1/5	10GbE	Cisco UCS C-Series 1	Port 0
	Eth1/13	10GbE	NetApp controller 01	0e
	Eth1/14	10GbE	NetApp controller 02	0e
	Eth1/27	10GbE	Cisco Nexus 93180 A	Eth1/19
	Eth1/28	10GbE	Cisco Nexus 93180 B	Eth1/19
	Eth1/31	10GbE	Cisco Nexus 2232 FEX A	Eth2/1

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/32	10GbE	Cisco Nexus 2232 FEX A	Eth2/2
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 12 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/1	10GbE	Cisco UCS Chassis FEX B	IOM1/1
	Eth1/2	10GbE	Cisco UCS Chassis FEX B	IOM1/2
	Eth1/5	10GbE	Cisco UCS C-Series 1	Port 1
	Eth1/13	10GbE	NetApp controller 01	0g
	Eth1/14	10GbE	NetApp controller 02	0g
	Eth1/27	10GbE	Cisco Nexus 93180 A	Eth1/20
	Eth1/28	10GbE	Cisco Nexus 93180 B	Eth1/20
	Eth1/31	10GbE	Cisco Nexus 2232 FEX B	Eth2/1
	Eth1/32	10GbE	Cisco Nexus 2232 FEX B	Eth2/2
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 13 Cisco UCS C-Series 1

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 1	Port 0	10GbE	Cisco UCS fabric interconnect A	Eth1/5
	Port 1	10GbE	Cisco UCS fabric interconnect B	Eth1/5

Table 14 Cisco UCS C-Series 2

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 2	Port 0	10GbE	Cisco Nexus 2232 FEX A	Eth1/1

Local Device	Local Port	Connection	Remote Device	Remote Port
	Port 1	10GbE	Cisco Nexus 2232 FEX B	Eth1/1

Table 15 Cisco UCS C-Series 3

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 3	Port 0	10GbE	Cisco Nexus 2232 FEX A	Eth1/2
	Port 1	10GbE	Cisco Nexus 2232 FEX B	Eth1/2

Storage Configuration

The following section details the initial NetApp AFF setup for a FlexPod with ACI. The following table shows the EPG VLANs, subnets, and bridge domains used in this lab validation.

Table 16 Lab Validation Infrastructure (Foundation) Tenant Configuration

EPG	Storage VLAN	UCS VLAN	External VLAN	Subnet / Gateway	Bridge Domain
IB-MGMT	N/A	DVS	163	172.26.163.0/24 - L2	BD-common-Internal
Core-Services	N/A	363	163	172.26.163.10/24	BD-common-Internal
SVM-MGMT	263	N/A	163	172.26.163.0/24 - L2	BD-common-Internal
iSCSI-A	3010	3110	N/A	192.168.110.0/24 - L2	BD-iSCSI-A
iSCSI-B	3020	3120	N/A	192.168.120.0/24 - L2	BD-iSCSI-B
NFS-LIF	3050	N/A	N/A	192.168.150.0/24 - L2	BD-NFS
NFS-VMK	N/A	3150	N/A	192.168.150.0/24 - L2	BD-NFS
vMotion	N/A	3000	N/A	192.168.100.0/24 - L2	BD-Internal
VMware vDS Pool	N/A	1101-1120	N/A	Varies	Varies
ACI System VLAN for AVS	N/A	4093	N/A	Varies	Varies

Controller AFF80XX Series

NetApp Hardware Universe

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

1. Confirm that the hardware and software components are supported with the version of ONTAP that you plan to install by using the [NetApp Hardware Universe \(HWU\)](#) site.
2. Access the [HWU](#) application to view the System Configuration guides. Click the Controllers tab to view the compatibility between ONTAP software versions and NetApp storage appliances with the desired specifications.
3. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers. These procedures can be found in the [AFF8000 Series product documentation](#) at the [NetApp Support](#) site.

Disk Shelves

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

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

Clustered Data ONTAP 8.3.2

Complete the Configuration Worksheet

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

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [Clustered Data ONTAP 8.3 Software Setup Guide](#) to learn about the information required to configure ONTAP. Table 17 lists the information that you need to configure two ONTAP nodes. You should customize the cluster detail values with the information that is applicable to your deployment.

Table 17 ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster Node01 IP address	<node01-mgmt-ip>
Cluster Node01 netmask	<node01-mgmt-mask>
Cluster Node01 gateway	<node01-mgmt-gateway>
Cluster Node02 IP address	<node02-mgmt-ip>
Cluster Node02 netmask	<node02-mgmt-mask>
Cluster Node02 gateway	<node02-mgmt-gateway>
Data ONTAP 8.3.2 URL	<url-boot-software>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

autoboot

3. Press Ctrl-C when prompted to Press Ctrl-C for Boot Menu.



Note: If Data ONTAP 8.3.2 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.2 is the version being booted, continue with step 14.

4. To install new software, select option 7.

7

5. Enter *y* (Yes) to continue the installation.

y

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

e0M

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

y

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

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

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



Note: This web server must be pingable.

<boot-software-url>

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

Enter

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

y

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

y



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

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

Press Ctrl-C for Boot Menu

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

4

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

`y`

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

`y`



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

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press `Ctrl-C` when prompted to Press `Ctrl-C` for Boot Menu.

```
Ctrl-C
```



If Data ONTAP 8.3.2 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.2 is the version being booted, continue with step 14.

4. To install new software, select option 7.

```
7
```

5. Enter `y` (Yes) to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

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

```
y
```

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

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

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



This web server must be pingable.

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

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

```
Enter
```

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

```
y
```

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

```
y
```



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

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

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

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

```
4
```

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

```
y
```

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

```
y
```



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

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when Data ONTAP 8.3.2 boots on the node for the first time.

1. Follow the prompts to set up node 01.

```
Welcome to node setup.
```

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

"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.

Any changes you made before quitting will be saved.

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

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

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

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

For further information on AutoSupport, see:

<http://support.netapp.com/autosupport/>

Type yes to confirm and continue {yes}:yes

Enter the node management interface port [e0M]: Enter

Enter the node management interface IP address: <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

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

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

For System Setup, this node's management address is: <var-node01-mgmt-ip>.

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

2. Log in to the node with the admin user ID and no password.

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

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

```
Mode set to HA. Reboot node to activate HA.
```

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

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

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

Welcome to node setup.

You can enter the following commands at any time:

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

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

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

Any changes you made before quitting will be saved.

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

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

```
Enter the node management interface IP address [<node01-mgmt-ip>]: Enter
```

```
Enter the node management interface netmask [<node01-mgmt-mask>]: Enter
```

```
Enter the node management interface default gateway [<node01-mgmt-gateway>]:  
Enter
```

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

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

For System Setup, this node's management address is: <node01-mgmt-ip>.

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

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

Create Cluster on Node 01

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

Table 18 Cluster Create in ONTAP prerequisites.

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

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

```
cluster setup
```

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

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

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

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

Any changes you made before quitting will be saved.

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

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

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



Note: If a login prompt appears instead of the Cluster Setup wizard, start the wizard by logging in with the admin user and no password, and then enter the `cluster setup` command.

To create a new cluster, complete the following steps:

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

```
create
```

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

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

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

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

4. The system defaults are displayed. Enter `no` to not use the system defaults. Use the following prompts to configure the cluster ports.

```
Existing cluster interface configuration found:
```

Port	MTU	IP	Netmask
e0a	9000	169.254.118.102	255.255.0.0
e0b	9000	169.254.152.110	255.255.0.0
e0c	9000	169.254.191.92	255.255.0.0
e0d	9000	169.254.233.52	255.255.0.0

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

```
System Defaults:
```

```
Private cluster network ports [e0a,e0b,e0c,e0d].
```

```
Cluster port MTU values will be set to 9000.
```

```
Cluster interface IP addresses will be automatically generated.
```

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

5. The steps to create a cluster are displayed.

```
Enter the cluster administrators (username "admin") password: <password>
Retype the password: <password>
```

```
Step 1 of 5: Create a Cluster
```

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

```
List the private cluster network ports [e0a,e0b,e0c,e0d]: e0a,e0c
```

```
Enter the cluster ports' MTU size [9000]: Enter
Enter the cluster network netmask [255.255.0.0]: Enter
Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0a [169.254.49.117]: Enter
Generating a default IP address. This can take several minutes...
Enter the cluster interface IP address for port e0c [169.254.176.252]: Enter

Enter the cluster name: <clustername>
Enter the cluster base license key: <cluster-base-license-key>
Creating cluster <clustername>

Step 2 of 5: Add Feature License Keys

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

Enter an additional license key []: <nfs-license>
```



Note: The cluster is created. This can take a minute or two.



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

```
Enter the cluster management interface port [e0b]: e0i
Enter the cluster management interface IP address: <clustermgmt-ip>
Enter the cluster management interface netmask: <clustermgmt-mask>
Enter the cluster management interface default gateway [<clustermgmt-gateway>]:
<clustermgmt-gateway>
```

A cluster management interface on port e0i with IP address <clustermgmt-ip> has been created. You can use this address to connect to and manage the cluster.

6. Enter the DNS domain name.

```
Enter the DNS domain names: <dns-domain-name>
Enter the name server IP addresses: <nameserver1-ip>,<nameserver2-ip>
```



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

7. Set up the node.

```
Where is the controller located []: <node-location>
Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address [<node01-mgmt-ip>]: Enter
Enter the node management interface netmask [<node01-mgmt-mask>]: Enter
Enter the node management interface default gateway [<node01-mgmt-gateway>]:
Enter
```


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

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

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

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

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

To complete cluster setup, you must join each additional node to the cluster by running "cluster setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for information about additional system configuration tasks. You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address (<clustermgmt-ip>).

To access the command-line interface, connect to the cluster management IP address (for example, ssh admin@<clustermgmt-ip>).

<clustername>::>



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

Join Node 02 to Cluster

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

Table 19 Cluster Join in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
----------------	----------------------

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
Cluster management IP address	<clustermgmt-ip>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>

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

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

```
admin
```

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

```
cluster setup
```

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

```
Storage failover partners must be members of the same cluster.
```

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

```
Welcome to the cluster setup wizard.
```

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

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

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

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

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

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

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

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

```
{join}:
```



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

3. Run the following command to join a cluster:

```
join
```

4. ONTAP detects the existing cluster and attempts to join it. Follow the prompts to join the cluster.

```
Existing cluster interface configuration found:
```

Port	MTU	IP	Netmask
e0a	9000	169.254.1.79	255.255.0.0
e0b	9000	169.254.54.223	255.255.0.0
e0c	9000	169.254.100.157	255.255.0.0
e0d	9000	169.254.138.142	255.255.0.0

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

System Defaults:

Private cluster network ports [e0a,e0b,e0c,e0d].

Cluster port MTU values will be set to 9000.

Cluster interface IP addresses will be automatically generated.

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

Step 1 of 3: Join an Existing Cluster

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

List the private cluster network ports [e0a,e0b,e0c,e0d]: e0a,e0c

Enter the cluster ports' MTU size [9000]: Enter

Enter the cluster network netmask [255.255.0.0]: Enter

Generating a default IP address. This can take several minutes...

Enter the cluster interface IP address for port e0a [169.254.165.52]: Enter

Generating a default IP address. This can take several minutes...

Enter the cluster interface IP address for port e0c [169.254.13.182]: Enter

5. The steps to join a cluster are displayed.

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

Starting cluster support services ..

This node has joined the cluster <clustername>.

Step 2 of 3: Configure Storage Failover (SFO)

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

SFO is enabled.

Step 3 of 3: Set Up the Node

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

Notice: HA is configured in management.



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

6. Set up the node.

```
Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address [<node02-mgmt-ip>]: Enter
Enter the node management interface netmask [<node02-netmask>]: Enter
Enter the node management interface default gateway [<node02-gateway>]: Enter
```

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

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

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

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

Press enter to continue: Enter

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

To complete cluster setup, you must join each additional node to the cluster by running "cluster setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for information about additional system configuration tasks. You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address (<clustermgmt-ip>).

To access the command-line interface, connect to the cluster management IP address (for example, `ssh admin@<clustermgmt-ip>`).



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

Log in to Cluster

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

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

Zero All Spare Disks

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

```
disk zerospares
```



Disk autoassign should have assigned half of the connected SSDs to each node in the HA pair. Also, the first 48 SSDs should have been partitioned with Advanced Disk partitioning. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare disks can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Onboard UTA2 Ports Personality

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

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

```
ucadmin show
```

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

```

0g      cna      target    -      -      online
<node01>
0h      cna      target    -      -      online
<node02>
0e      cna      target    -      -      online
<node02>
0f      cna      target    -      -      online
<node02>
0g      cna      target    -      -      online
<node02>
0h      cna      target    -      -      online

```

8 entries were displayed.

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

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



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

Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, run the following command:

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

Set Up Management Broadcast Domain

To set up the default broadcast domain for management network interfaces by removing all ports that are not connected to the out-of-band management network, run the following commands. When the steps are complete, the default broadcast domain should only have the e0M and e0I ports from the two nodes. In this release of ONTAP, LIF failover groups are automatically setup so that they mirror the broadcast domains. The default broadcast domain is used for all out-of-band management interfaces.

```
broadcast-domain remove-ports -broadcast-domain Default -ports <node01>:e0b,
<node01>:e0d,<node01>:e0e,<node01>:e0f,<node01>:e0g,<node01>:e0h,<node01>:e0j,<no
de01>:e0k,<node01>:e0l,<node02>:e0b,<node02>:e0d,<node02>:e0e,<node02>:e0f,<node0
2>:e0g,<node02>:e0h,<node02>:e0j,<node02>:e0k,<node02>:e0l
broadcast-domain show
```

Set Up Service Processor Network Interface

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

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



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

Create Aggregates

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

To create new aggregates, complete the following steps:

1. Run the following commands:

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



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



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



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

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

```
node run <node01> aggr options aggr1_node01 nosnap on
node run <node02> aggr options aggr1_node02 nosnap on
```

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

```
node run <node01> snap delete -A -a -f aggr1_node01
```

```
node run <node02> snap delete -A -a -f aggr1_node02
```

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

```
aggr show
```

```
aggr rename -aggregate aggr0 -newname <node01-rootaggrname>
```

Verify Storage Failover

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

1. Verify the status of storage failover.

```
storage failover show
```



Note: Both node 01 and node 02 must be capable of performing a takeover. Execute step 2 if the nodes are not capable of performing a takeover. Otherwise, continue with step 3.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <node01> -enabled true
```



Note: Enabling failover on one node enables it for both nodes.

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



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

```
cluster ha show
```

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

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



Note: Do not run this command for clusters with more than two nodes because doing so causes problems with failover.

```
cluster ha modify -configured true
```

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

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

```
storage failover hwassist show
```



Note: The Monitor Status may show inactive, but this will resolve over time if the hwassist IPs are set correctly.

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

Disable Flow Control on 10GE Ports

NetApp recommends disabling flow control on all of the 10GE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps. After the steps are completed, only the 1GE ports should have full flow control.

1. Run the following commands to configure node 01:

```
network port modify -node <node01> -port e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0h -
flowcontrol-admin none
```

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

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

2. Run the following commands to configure node 02:

```
network port modify -node <node02> -port e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0h -
flowcontrol-admin none
```

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

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

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

Disable Unused FcoE Ports

Unused data FCoE ports on active interfaces should be disabled. To disable these ports, run the following commands. If using FCoE storage in this implementation, do not disable any FCoE ports connected to the Cisco UCS fabric interconnects.

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

Configure NTP

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

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

```
timezone <timezone>
```



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

2. To set the date for the cluster, run the following commands. It is only necessary to set the date if it is incorrect.

```
date
```

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



Note: The format for the date is `<[Century] [Year] [Month] [Day] [Hour] [Minute] . [Second]>`; for example, `201509081735.17`.

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

```
cluster time-service ntp server create -server <global-ntp-server-ip>
```

```
cluster time-service ntp server show
```



Note: The global NTP server IP should be reachable from the out-of-band management subnet.

Configure SNMP

To configure SNMP, complete the following steps:

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

```
snmp contact <storage-admin-email>
```

```
snmp location "<snmp-location>"
```

```
snmp init 1
```

```
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

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

Configure SNMPv1 Access

To configure SNMPv1 access, complete the following step:

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

```
snmp community add ro <snmp_community>
```

Configure AutoSupport

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

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

Enable Cisco Discovery Protocol

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

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

Create Broadcast Domains in ONTAP

To create data broadcast domains, run the following commands:

```
broadcast-domain create -broadcast-domain Infra-NFS -mtu 9000

broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500
broadcast-domain create -broadcast-domain Infra-iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra-iSCSI-B -mtu 9000
```



Note: If you are not setting up iSCSI boot or access to iSCSI application data LUNs, do not create the iSCSI broadcast domains.

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, run the following commands.

```
ifgrp create -node <node01> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <node01> -ifgrp a0a -port e0b
ifgrp add-port -node <node01> -ifgrp a0a -port e0d

ifgrp create -node <node02> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <node02> -ifgrp a0a -port e0b
ifgrp add-port -node <node02> -ifgrp a0a -port e0d

ifgrp show
```



Note: Since the corresponding ports have not been configured on the ACI leaf switches, the ifgrps do not show full active at this time.

Create VLANs

To create VLANs, complete the following steps:

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

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

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

```
broadcast-domain add-ports -broadcast-domain Infra-NFS -ports <node01>:a0a-
<storage-infra-nfs-vlan-id>, <node02>:a0a-<storage-infra-nfs-vlan-id>
```

2. Create in-band management VLAN ports and add them to the data broadcast domain.

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

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

3. Create iSCSI VLAN ports and add them to the data broadcast domain. If you are not setting up iSCSI boot or access to iSCSI application data LUNs, do not create the iSCSI VLAN ports.

```
network port vlan create -node <node01> -vlan-name a0a-<storage-infra-iscsi-A-
vlan-id>
network port vlan create -node <node01> -vlan-name a0a-<storage-infra-iscsi-B-
vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-infra-iscsi-A-
vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-infra-iscsi-B-
vlan-id>
```

```
broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <node01>:a0a-
<storage-infra-iscsi-A-vlan-id>,<node02>:a0a-<storage-infra-iscsi-A-vlan-id>
broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <node01>:a0a-
<storage-infra-iscsi-B-vlan-id>,<node02>:a0a-<storage-infra-iscsi-B-vlan-id>
```

```
broadcast-domain show
```

Create Storage Virtual Machine



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

To create the infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_node02 -
rootvolume-security-style unix
```

2. Remove unused SVM storage protocols from the list of `nfs`, `cifs`, `fc`, `iscsi`, and `ndmp`. In this example, we keep `nfs`, `fc`, and `iscsi`.

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

3. Add the two data aggregates to the Infra-SVM aggregate list so that the NetApp VSC can provision storage in those aggregates.

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

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

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

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

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

```
vserver nfs show
```

6. If either iSCSI boot or iSCSI LUN access is provided by this SVM, create the iSCSI service on this SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

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

```
iscsi show
```

7. If either FCoE boot or FCoE LUN access is provided by this SVM, create the FCP service on this SVM. This command also starts the FCP service and sets the FCP World Wide Node Name (WWNN) for the SVM.

```
fcv create -vserver Infra-SVM
```

```
fcv show
```

Create Load-Sharing Mirrors of SVM Root Volume

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

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

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

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

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

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

3. Create the mirroring relationships.

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

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

4. Initialize the mirroring relationship.

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

Configure HTTPS Access

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

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

```
set diag
```

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

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

```
security certificate show
```

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



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

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

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

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

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

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

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

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

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

7. Disable HTTP cluster management access. It is normal for some of these commands to return an error message stating that the entry does not exist.

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

8. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set admin
```

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



Note: The “|” symbols are part of the command above.

Configure NFSv3

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

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for the infrastructure nfs subnet CIDR address (for example, 192.168.150.0/24).

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

```
vserver export-policy rule show
```

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

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

Create NetApp FlexVol Volumes

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

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

To create a FlexVol volume, run the following commands:

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

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

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

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

Create Boot LUNs

To create two boot LUNs, run the following commands:

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

Adjust Storage Efficiency Settings

When volumes are created by default on NetApp AFF systems, inline compression and inline deduplication are enabled with no scheduled deduplication scans. In this section, a deduplication scan schedule is added to the volumes `infra_datastore_1` and `esxi_boot`, and inline deduplication is removed from the volume `infra_swap`.

To adjust the storage efficiency settings, complete the following steps:

1. Add a daily deduplication scan to the `infra_datastore_1` and `esxi_boot` volumes.

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

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

2. Remove inline deduplication from the `infra_swap` volume.

```
efficiency modify -vserver Infra-SVM -volume infra_swap -inline-dedupe false
```

```
efficiency show -instance
```

Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), complete the following step. If you are not setting up iSCSI boot or access to iSCSI application data LUNs, do not create the iSCSI LIFs.

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-node <node01> -home-port a0a-<storage-infra-iscsi-A-vlan-id> -address <node01-infra-iscsi-lif01a-ip> -netmask <storage-infra-iscsi-A-mask>
```

```
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -home-node <node01> -home-port a0a-<storage-infra-iscsi-B-vlan-id> -address <node01-infra-iscsi-lif01b-ip> -netmask <storage-infra-iscsi-B-mask>
```

```
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -home-node <node02> -home-port a0a-<storage-infra-iscsi-A-vlan-id> -address <node02-infra-iscsi-lif02a-ip> -netmask <storage-infra-iscsi-A-mask>
```

```
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-node <node02> -home-port a0a-<storage-infra-iscsi-B-vlan-id> -address <node02-infra-iscsi-lif02b-ip> -netmask <storage-infra-iscsi-B-mask>
```

```
network interface show -vserver Infra-SVM -lif iscsi*
```

Create FCoE LIFs

To create four FCoE LIFs (two on each node), run the following commands:

```
network interface create -vserver Infra-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node <node01> -home-port 0e
```

```
network interface create -vserver Infra-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node <node01> -home-port 0g
```



```
network interface create -vserver Infra-SVM -lif fcp_lif02a -role data -data-
protocol fcp -home-node <node02> -home-port 0e

network interface create -vserver Infra-SVM -lif fcp_lif02b -role data -data-
protocol fcp -home-node <node02> -home-port 0g

network interface show -vserver Infra-SVM -lif fcp*
```



Note: If you are not setting up FCoE boot or access to FCoE application data LUNs, do not create the FCoE LIFs.

Create NFS LIFs

To create NFS LIFs, run the following commands:

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

network interface create -vserver Infra-SVM -lif nfs_infra_datastore_1 -role data
-data-protocol nfs -home-node <node02> -home-port a0a-<storage-infra-nfs-vlan-id>
-address <nfs-lif-infra_datastore_1-ip> -netmask <nfs-lif-infra-mask> -status-
admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface show -vserver Infra-SVM -lif nfs*
```



Note: NetApp recommends creating a new LIF for each datastore.

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF to the in-band management network, complete the following steps.

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif svm-mgmt -role data -data-
protocol none -home-node <node02> -home-port a0a-<storage-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
```



Note: The SVM management IP in this step should be in the In-Band Management subnet.

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

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

network route show
```

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

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

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

Server Configuration

The following section details the Cisco UCS setup for a FlexPod with ACI. This section includes setup for both iSCSI boot and LUN access and FCoE boot and LUN access. Please skip any steps related to a storage protocol not being implemented in your FlexPod. The following table shows the End Point Group (EPG) VLANs, Subnets, and Bridge Domains used in this lab validation.

Table 20 Lab Validation Infrastructure (Foundation) Tenant Configuration

EPG	Storage VLAN	UCS VLAN	External VLAN	Subnet / Gateway	Bridge Domain
IB-MGMT	N/A	DVS	163	172.26.163.0/24 - L2	BD-common-Internal
Core-Services	N/A	363	163	172.26.163.10/24	BD-common-Internal
SVM-MGMT	263	N/A	163	172.26.163.0/24 - L2	BD-common-Internal
iSCSI-A	3010	3110	N/A	192.168.110.0/24 - L2	BD-iSCSI-A
iSCSI-B	3020	3120	N/A	192.168.120.0/24 - L2	BD-iSCSI-B
NFS-LIF	3050	N/A	N/A	192.168.150.0/24 - L2	BD-NFS
NFS-VMK	N/A	3150	N/A	192.168.150.0/24 - L2	BD-NFS
vMotion	N/A	3000	N/A	192.168.100.0/24 - L2	BD-Internal
VMware vDS Pool	N/A	1101-1120	N/A	Varies	Varies
ACI System VLAN for AVS	N/A	4093	N/A	Varies	Varies

Cisco UCS Base Configuration

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

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

Cisco UCS 6248 A

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

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

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

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

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

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

```
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  
Which switch fabric (A/B) []: A  
Enter the system name: <ucs-clustername>  
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: <nameserver1-ip>  
Configure the default domain name? (yes/no) [n]: y  
Default domain name: <dns-domain-name>  
Join centralized management environment (UCS Central)? (yes/no) [n]: n  
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no):  
yes  
2. Wait for the login prompt to make sure that the configuration has been saved.
```

Cisco UCS 6248 B

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

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

```
Enter the configuration method. (console/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 for the peer Fabric interconnect: <password>  
Physical switch Mgmt0 IP address: <ucsb-mgmt-ip>  
Apply and save the configuration (select 'no' if you want to re-enter)?  
(yes/no): y
```
2. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS Setup

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6248 fabric interconnect cluster address.

2. Under HTML, click the Launch UCS Manager link to launch the Cisco UCS Manager HTML5 User Interface.
3. When prompted, enter `admin` as the user name and enter the administrative password.
4. Click Login to log in to Cisco UCS Manager.
5. Respond to the popup on Anonymous Reporting and click OK.

Upgrade Cisco UCS Manager Software to Version 3.1(1h)

This document assumes the use of Cisco UCS 3.1(1h). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to version 3.1(1h), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Configure Cisco UCS Call Home

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

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

The screenshot displays the Cisco UCS Manager Admin configuration interface. The left-hand navigation pane is expanded to show the 'Admin' section under 'Communication Management'. The main configuration area is divided into several sections:

- Admin:** State is set to On. Switch Priority is set to Critical. Throttling is set to On.
- Contact Information:** Fields for Contact, Phone, Email, and Address are present but redacted.
- Ids:** Fields for Customer ID, Contract ID, and Site ID are present but redacted.
- Email Addresses:** Fields for From and Reply To are present but redacted.
- SMTP Server:** Host (IP Address or Hostname) is redacted. Port is set to 25.

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for out of band (mgmt0) server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:

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

▲ Create Block of IPv4 Addresses
✕

Create a Block of IPv4 Addresses ?

From : <input type="text" value="192.168.1.225"/>	Size : <input type="text" value="12"/>
Subnet Mask : <input type="text" value="255.255.255.0"/>	Default Gateway : <input type="text" value="192.168.1.254"/>
Primary DNS : <input type="text" value="0.0.0.0"/>	Secondary DNS : <input type="text" value="0.0.0.0"/>



Note: This block of IP addresses should be in the out of band management subnet.

5. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

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

Change Fabric Interconnect FC Mode to Switching

If you are providing FCoE boot or access to FCoE LUNs, the Cisco UCS Fabric Interconnects (FIs) must be put into FC Switching Mode. When in FC Switching Mode, the FIs function as Fibre Channel switches. When changing to FC Switching Mode, both FIs reboot immediately. This change should be done when no network traffic is running on either FI. If you are providing FCoE boot or access to FCoE LUNs complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. In the Actions pane, select Set FC Switching Mode.
4. Select Yes and OK.
5. Wait for both Fabric Interconnects to complete reboot (by monitoring the consoles) and log back into UCS Manager.
6. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
7. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
8. Verify the FC Mode is now Switch.

Edit Chassis Discovery Policy

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

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

Enable Server, Uplink, and Storage Ports

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Fixed Module.
4. Expand and select Ethernet Ports.
5. Select the ports that are connected to the chassis, Cisco 2232 FEX, and direct connect UCS C-Series servers, right-click them, and select "Configure as Server Port".

6. Click Yes to confirm server ports and click OK.
7. Verify that the ports connected to the chassis, C-series servers and to the Cisco 2232 FEX are now configured as Server ports by selecting Ethernet Ports.
8. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
9. Click Yes to confirm uplink ports and click OK.
10. Verify that the uplink ports are now configured as Network ports by selecting Ethernet Ports.
11. If you are not providing FCoE Boot or access to FCoE LUNs, continue to step 15.
12. Select the ports that are connected to FCoE ports on the NetApp Storage Controllers, right-click them, and select Configure as FCoE Storage Port.
13. Click Yes to confirm FCoE storage ports and click OK.
14. Verify that the FCoE Storage ports are now configured as Fcoe Storage ports by selecting Ethernet Ports.

The screenshot displays the 'Ethernet Ports' configuration window. On the left, a navigation tree shows the hierarchy: Equipment > Servers > LAN > Ethernet Ports. The main table lists 29 ports across 15 slots. The 'If Role' column indicates the configuration for each port, such as 'Server', 'Fcoe Storage', and 'Network'. The 'Overall Status' and 'Admin State' columns show the current operational state of each port.

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	1	00:2A:6A:62:55...	Server	Physical	Up	Enabled
1	0	2	00:2A:6A:62:55...	Server	Physical	Up	Enabled
1	0	3	00:2A:6A:62:55...	Server	Physical	Up	Enabled
1	0	4	00:2A:6A:62:55...	Server	Physical	Up	Enabled
1	0	5	00:2A:6A:62:55...	Server	Physical	Up	Enabled
1	0	6	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	7	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	8	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	9	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	10	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	11	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	12	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	13	00:2A:6A:62:55...	Fcoe Storage	Physical	Up	Enabled
1	0	14	00:2A:6A:62:55...	Fcoe Storage	Physical	Up	Enabled
1	0	15	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	16	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	17	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	18	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	19	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	20	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	21	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	22	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	23	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	24	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	25	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	26	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled
1	0	27	00:2A:6A:62:55...	Network	Physical	Up	Enabled
1	0	28	00:2A:6A:62:55...	Network	Physical	Up	Enabled
1	0	29	00:2A:6A:62:55...	Unconfigured	Physical	Sfp Not Pre...	Disabled

15. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
16. Expand Fixed Module.

17. Expand and select Ethernet Ports.
18. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX, right-click them, and select Configure as Server Port.
19. Click Yes to confirm server ports and click OK.
20. Verify that the ports connected to the chassis, C-series servers and to the Cisco 2232 FEX are now configured as Server ports by selecting Ethernet Ports.
21. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
22. Click Yes to confirm uplink ports and click OK.
23. Verify that the uplink ports are now configured as Network ports by selecting Ethernet Ports.
24. If you are not providing FCoE Boot or access to FCoE LUNs, continue to the next section, Acknowledge Cisco UCS Chassis and FEX.
25. Select the ports that are connected to FCoE ports on the NetApp Storage Controllers, right-click them, and select Configure as FCoE Storage Port.
26. Click Yes to confirm FCoE storage ports and click OK.
27. Verify that the FCoE Storage ports are now configured as FCoE Storage ports by selecting Ethernet Ports.

Acknowledge Cisco UCS Chassis and FEX

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.
4. Click Yes and then click OK to complete acknowledging the chassis.
5. If Nexus 2232 FEX are part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and select Acknowledge FEX.
7. Click “Yes” and then click OK to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus Switches

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

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



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

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

The screenshot shows a window titled "Create Port Channel" from the "Unified Computing System Manager". The main heading is "Set Port Channel Name". On the left, a navigation pane shows two steps: "1. ✓ Set Port Channel Name" and "2. Add Ports". The main area contains two input fields: "ID : 13" and "Name : Po-13-Nexus". At the bottom right, there are four buttons: "< Prev", "Next >", "Finish", and "Cancel".

8. Select the network uplink ports to be added to the port channel:
9. Click >> to add the ports to the port channel. Make sure both ports are added.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.

13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter Po-14-Nexus as the name of the port channel.
17. Click Next.
18. Select the network uplink ports to be added to the port channel:
19. Click >> to add the ports to the port channel. Make sure both ports are added.
20. Click Finish to create the port channel.
21. Click OK.

Create MAC Address Pools

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

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



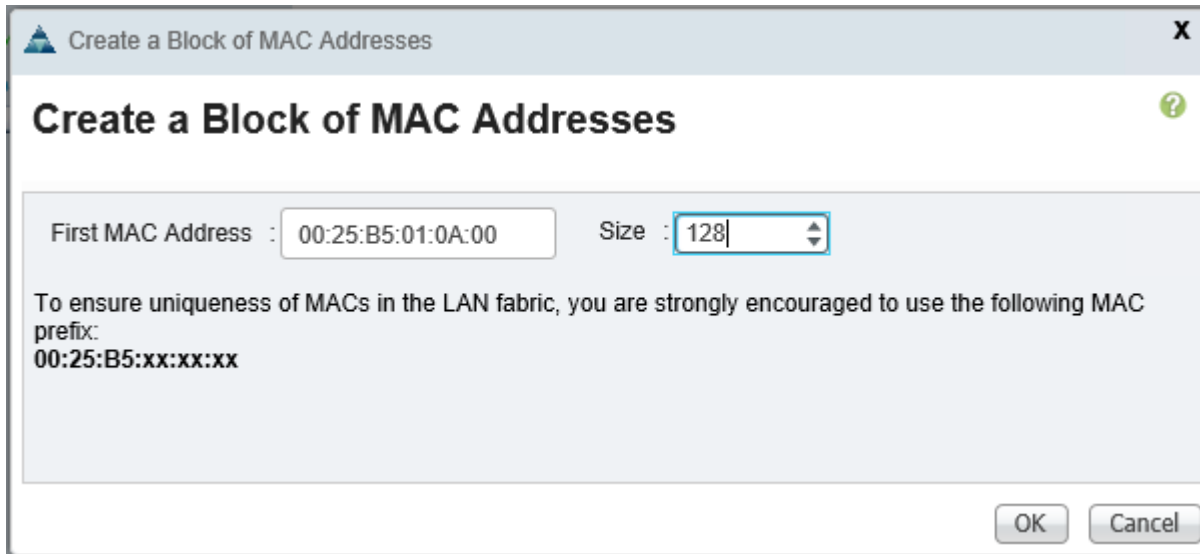
Note: In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC-Pool1-A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select the Sequential Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting MAC address.



Note: For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. It is recommended to not change the first three octets of the MAC address.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources. Remember that multiple Cisco VIC vNICs will be created on each server and each vNIC will be assigned a MAC address.



Create a Block of MAC Addresses

First MAC Address : 00:25:B5:01:0A:00 Size : 128

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

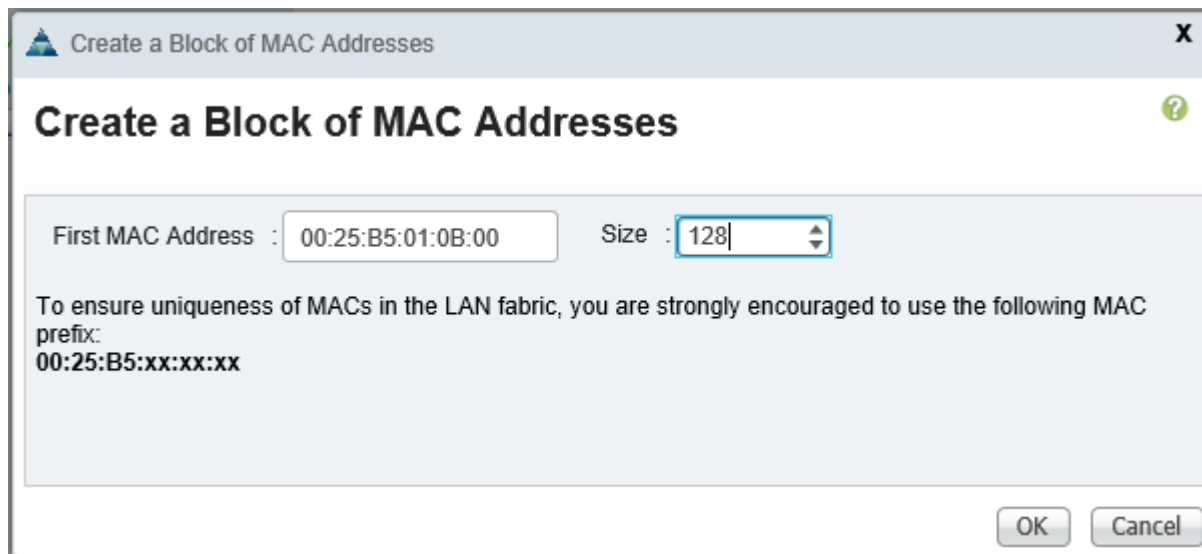
OK Cancel

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



Note: For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. It is recommended to not change the first three octets of the MAC address.

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



First MAC Address : 00:25:B5:01:0B:00 Size : 128

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

OK Cancel

24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Create a WWNN Address Pool

If you are providing FCoE boot or access to FCoE LUNs, create a World Wide Node Name (WWNN) pool by completing the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Select Create WWNN Pool to create the WWNN address pool.
5. Enter `wwnn-pool` as the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.
7. Select the Sequential Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting WWNN address.
11. Specify a size for the WWNN address pool that is sufficient to support the available blade or server resources. Each server will receive one WWNN.

Create WWN Block

From : 20:00:00:25:B5:01:00:00 Size : 128

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

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

OK Cancel

12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.

Create a WWPN Address Pools

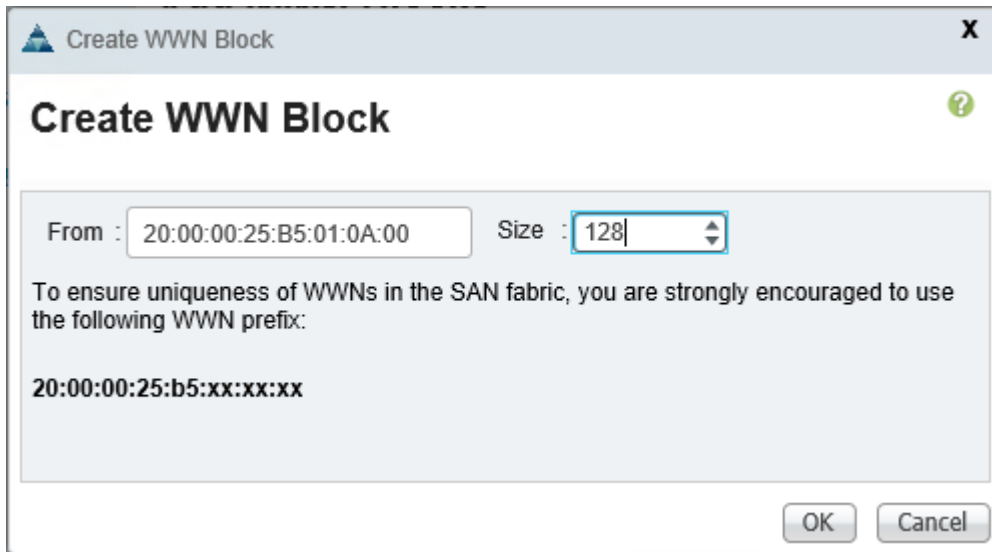
If you are providing FCoE boot or access to FCoE LUNs, create a World Wide Port Name (WWPN) pool for each SAN switching fabric by completing the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. Right-click WWPN Pools under the root organization.
4. Select Create WWPN Pool to create the first WWPN address pool.
5. Enter `WWPN-Pool-A` as the name of the WWPN pool.
6. Optional: Enter a description for the WWPN pool.
7. Select the Sequential Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting WWPN address.



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

11. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources. **Each server's Fabric A vHBA will receive one WWPN from this pool.**



Create WWN Block

From : 20:00:00:25:B5:01:0A:00 Size : 128

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

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

OK Cancel

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



Note: For the FlexPod solution, the recommendation is to place 0B in the next-to-last octet of the starting WWPN address to identify all of the WWPN addresses as fabric B addresses.

23. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources. **Each server's Fabric B vHBA will receive one WWPN from this pool.**

Create WWN Block

From : 20:00:00:25:B5:01:0B:00 Size : 128

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

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

OK Cancel

24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

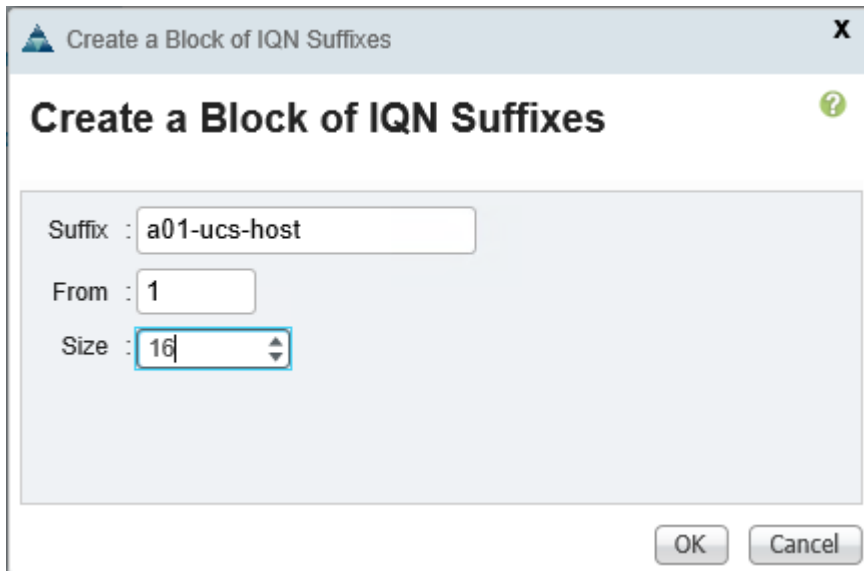
Create IQN Pools for iSCSI Boot

If you are providing iSCSI boot or access to iSCSI LUNs, configure the necessary IQN pools for the Cisco UCS environment, by completing the following steps.

1. In the UCS Manager, select the SAN tab on the left.
2. Select Pools > root.
3. Right-click **IQN Pools** under the root organization.
4. Select **Create IQN Suffix Pool** to create the IQN pool.
5. Enter `IQN-Pool1` for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter `iqn.1992-08.com.cisco` as the prefix
8. Select Sequential for Assignment Order.
9. Click **Next**.
10. Click **Add**.
11. Enter `ucs-host` as the suffix.
12. Enter 1 in the From field.

13. Specify a size of the IQN block sufficient to support the available server resources. Each server will receive one IQN.

14. Click OK.



The screenshot shows a dialog box titled "Create a Block of IQN Suffixes". It contains three input fields: "Suffix" with the value "a01-ucs-host", "From" with the value "1", and "Size" with the value "16". The "Size" field is a spinner control. At the bottom right, there are "OK" and "Cancel" buttons.

15. Click Finish.

16. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot

If you are providing iSCSI boot, these steps provide details for configuring the necessary IP pools iSCSI boot for the Cisco UCS environment. To create IP pools for iSCSI boot, complete the following steps:

1. In Cisco UCS Manager, select the LAN tab on the left.
2. Select Pools > root.



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

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

10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Enter the Subnet Mask.
12. Set the size to enough addresses to accommodate the servers.

The screenshot shows a dialog box titled "Create a Block of IPv4 Addresses". The dialog contains the following fields and values:

From :	192.168.110.100	Size :	16
Subnet Mask :	255.255.255.0	Default Gateway :	0.0.0.0
Primary DNS :	0.0.0.0	Secondary DNS :	0.0.0.0

Buttons: OK, Cancel

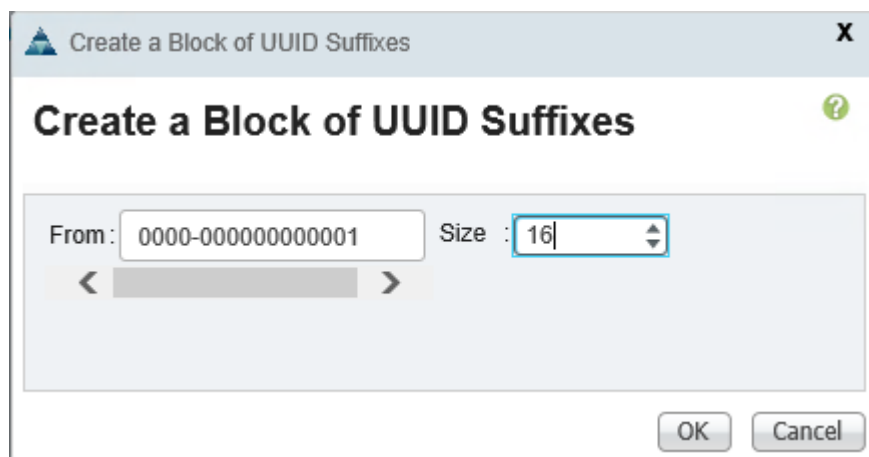
13. Click OK.
14. Click Next.
15. Click Finish.
16. Click OK in the confirmation message.
17. Right-click IP Pools under the root organization.
18. Select Create IP Pool to create the IP pool.
19. Enter `iSCSI-IP-Pool-B` for the name of the IP pool.
20. Optional: Enter a description of the IP pool.
21. Select Sequential for Assignment Order.
22. Click Next.
23. Click Add.
24. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.

25. Enter the Subnet Mask.
26. Set the size to enough addresses to accommodate the servers.
27. Click OK.
28. Click Next.
29. Click Finish.
30. Click OK in the confirmation message.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID-Pool` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Click Next.
9. Click Add to add a block of UUIDs.
10. Keep the From field at the default setting or specify a unique value.
11. Specify a size for the UUID block that is sufficient to support the available blade or server resources.



The screenshot shows a dialog box titled "Create a Block of UUID Suffixes". The dialog has a title bar with a close button (X) and a help icon (?). The main content area has a heading "Create a Block of UUID Suffixes" and a question mark icon. Below the heading, there are two input fields: "From:" with the value "0000-0000000000001" and "Size:" with the value "16". The "Size" field is a spinner control. Below these fields are left and right arrow buttons. At the bottom right, there are "OK" and "Cancel" buttons.

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

Create Server Pool

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



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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `Infra-Pool1` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the `Infra-Pool1` server pool.
9. Click Finish.
10. Click OK.

Create VLANs

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

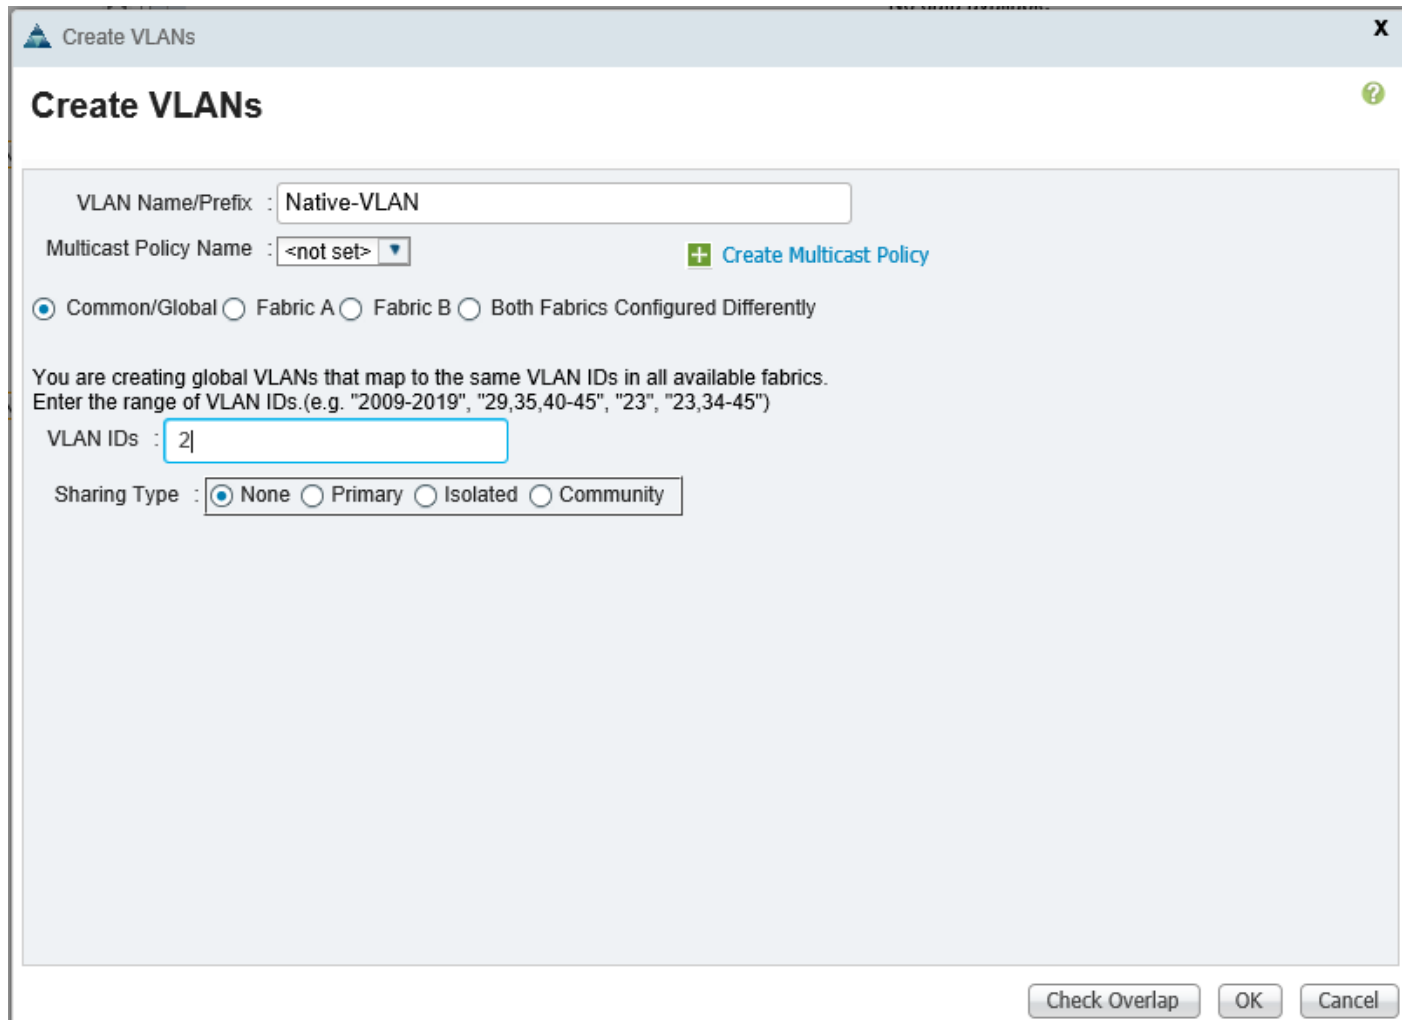
1. In Cisco UCS Manager, click the LAN tab in the navigation pane.



Note: In this procedure, up to seven unique VLANs are created and a range of 20 VLANs for the APIC-controlled VMware vDS is created. See Table 20

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `Native-VLAN` as the name of the VLAN to be used as the native VLAN.

6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.



The screenshot shows a 'Create VLANs' dialog box with the following fields and options:

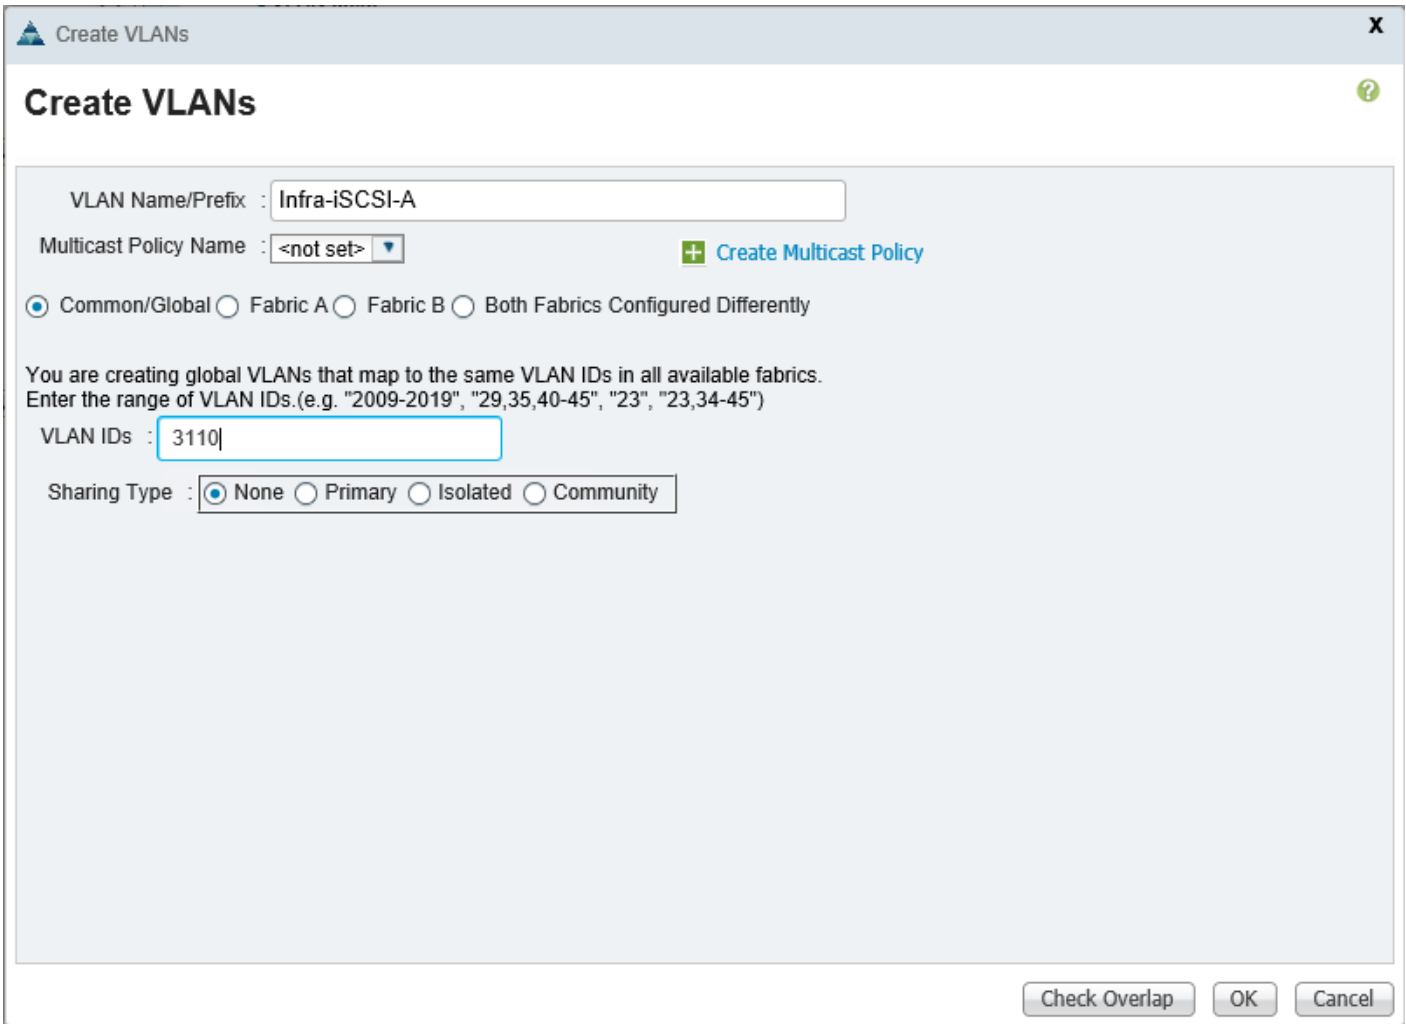
- VLAN Name/Prefix :** Native-VLAN
- Multicast Policy Name :** <not set> (with a dropdown arrow) and a '+ Create Multicast Policy' button.
- Scope:** Common/Global, Fabric A, Fabric B, Both Fabrics Configured Differently
- Instructions:** You are creating global VLANs that map to the same VLAN IDs in all available fabrics. Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")
- VLAN IDs :** 2
- Sharing Type :** None, Primary, Isolated, Community

Buttons at the bottom right: Check Overlap, OK, Cancel.

10. Click VLANs in the navigation pane. In the VLANs pane, right-click the newly created Native-VLAN and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs.
14. Enter Infra-iSCSI-A as the name of the VLAN to be used for the first iSCSI VLAN.
15. Keep the Common/Global option selected for the scope of the VLAN.

16. Enter the VLAN ID for the first iSCSI VLAN in the UCS

17. Click OK, then OK.



The screenshot shows a 'Create VLANs' dialog box with the following fields and options:

- VLAN Name/Prefix :** Infra-iSCSI-A
- Multicast Policy Name :** <not set> (with a dropdown arrow) and a '+ Create Multicast Policy' link.
- Scope:** Common/Global, Fabric A, Fabric B, Both Fabrics Configured Differently
- Instructions:** 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 :** 3110
- Sharing Type :** None, Primary, Isolated, Community

Buttons at the bottom: Check Overlap, OK, Cancel.

18. Right-click VLANs.

19. Select Create VLANs.

20. Enter Infra-iSCSI-B as the name of the VLAN to be used for the second iSCSI VLAN.

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

22. Enter the VLAN ID for the second iSCSI VLAN in UCS.

23. Click OK, then OK.

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

24. Right-click VLANs.
25. Select Create VLANs
26. Enter `IB-Mgmt` as the name of the VLAN to be used for in-band management traffic in the UCS.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the In-Band management VLAN ID.
29. Keep the Sharing Type as None.
30. Click OK, and then click OK again.

Create VLANs

VLAN Name/Prefix : IB-Mgmt

Multicast Policy Name : <not set> [+ 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 : 363

Sharing Type : None Primary Isolated Community

Check Overlap OK Cancel

31. Right-click VLANs.
32. Select Create VLANs.
33. Enter `Infra-NFS` as the name of the VLAN to be used for Infrastructure NFS in the UCS.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the NFS VLAN ID.
36. Keep the Sharing Type as None.
37. Click OK, and then click OK again.

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

38. Right-click VLANs.
39. Select Create VLANs.
40. Enter `vMotion` as the name of the VLAN to be used for VMware vMotion in the UCS.
41. Keep the Common/Global option selected for the scope of the VLAN.
42. Enter the vMotion VLAN ID.
43. Keep the Sharing Type as None.
44. Click OK, and then click OK again.

Create VLANs

VLAN Name/Prefix : vMotion

Multicast Policy Name : <not set> [+ 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 : 3000

Sharing Type : None Primary Isolated Community

[Check Overlap](#) [OK](#) [Cancel](#)

45. If you will be using the VMware vDS Virtual Distributed Switch in this FlexPod, complete steps 46-52.
If you will not be using the VMware vDS, continue at step 53.
46. Right-click VLANs.
47. Select Create VLANs.
48. Enter APIC-vDS- as the name of the VLANs to be used for the APIC-controlled VMware vDS.
49. Keep the Common/Global option selected for the scope of the VLAN.
50. Enter the VLAN ID range.
51. Keep the Sharing Type as None.
52. Click OK, and then click OK again.

Create VLANs

VLAN Name/Prefix : APIC-vDS-

Multicast Policy Name : <not set> [+ 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 : 1101-1120

Sharing Type : None Primary Isolated Community

53. If you are using the Cisco AVS Virtual Distributed Switch in VXLAN switching mode in this FlexPod, complete steps 54-60. If you will not be using the Cisco AVS, continue to the next section, Create VSANs and Configure FCoE Storage Ports.

54. Right-click VLANs.

55. Select Create VLANs.

56. Enter `ACI-System-VLAN` as the name of the VLANs to be used for OPFLEX communication between the ACI Fabric and the AVS Virtual Ethernet Module (VEM) VXLAN Tunnel Endpoints (VTEPs) on the UCS VMware ESXi Hosts.

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

58. Enter the VLAN ID for the ACI System VLAN.



Note: The ACI System VLAN will be set when the ACI APIC is setup. The recommended setting for this VLAN is 4093. VLAN 4094 can be set in the APIC, but cannot be set in Cisco UCS.

59. Keep the Sharing Type as None.

60. Click OK, and then click OK again.

Create VLANs

VLAN Name/Prefix : ACI-System-VLAN

Multicast Policy Name : <not set> [+ 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 : 4093

Sharing Type : None Primary Isolated Community

[Check Overlap](#) [OK](#) [Cancel](#)

Create VSANs and Configure FCoE Storage Ports

If you are providing FCoE boot or access to FCoE LUNs, to configure the necessary virtual local area networks (VSANs) and FCoE Storage Ports for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > SAN Cloud.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter VSAN-A as the name of the VSAN to be used as the Fabric A VSAN.
6. Select Enabled for FC Zoning.
7. Select Fabric A for the scope of the VSAN.

8. Enter a unique VSAN ID and FCoE VLAN ID.



Note: It is recommended to use a value other than 1 and to use the same value for both IDs.

9. Click OK, and then click OK again.

Create VSAN

Name : VSAN-A

FC Zoning Settings

FC Zoning : Disabled Enabled

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

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

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

VSAN ID : 101

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 : 101

OK Cancel

10. Select SAN > Storage Cloud.
11. Right-click VSANs.
12. Select Create Storage VSAN.
13. Enter VSAN-A as the name of the VSAN to be used as the Fabric A VSAN.
14. Select Enabled for FC Zoning.
15. Select Fabric A for the scope of the VSAN.
16. Enter a unique VSAN ID and FCoE VLAN ID. These IDs should be the same as what was entered above for VSAN-A in the SAN Cloud.

17. Click OK, and then click OK again.

Create Storage VSAN

Name : VSAN-A

FC Zoning Settings

FC Zoning : Disabled Enabled

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

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

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

VSAN ID : 101

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 : 101

OK Cancel

18. Select SAN > SAN Cloud.

19. Right-click VSANs.

20. Select Create VSAN.

21. Enter VSAN-B as the name of the VSAN to be used as the Fabric B VSAN.

22. Select Enabled for FC Zoning.

23. Select Fabric B for the scope of the VSAN.

24. Enter a unique VSAN ID and FCoE VLAN ID.



Note: It is recommended to use a value other than 1 and to use the same value for both IDs. Also, the Fabric B VSAN ID should be different than the Fabric A VSAN ID.

25. Click OK, and then click OK again.
26. Select SAN > Storage Cloud.
27. Right-click VSANs.
28. Select Create Storage VSAN.
29. Enter `vsan-B` as the name of the VSAN to be used as the Fabric B VSAN.
30. Select Enabled for FC Zoning.
31. Select Fabric B for the scope of the VSAN.
32. Enter a unique VSAN ID and FCoE VLAN ID. These IDs should be the same as what was entered above for VSAN-B in the SAN Cloud.
33. Click OK, and then click OK again.
34. Expand Storage Cloud > Fabric A > Storage FCoE Interfaces. Two interfaces should be shown.
35. Select the first interface that is connected to Storage Controller 01.
36. Fill in the user label with a description that describes the connection and port on the storage controller.
37. Click Save Changes and click OK.
38. Using the pull-down, select the VSAN-A VSAN.
39. Click Save Changes and click OK.

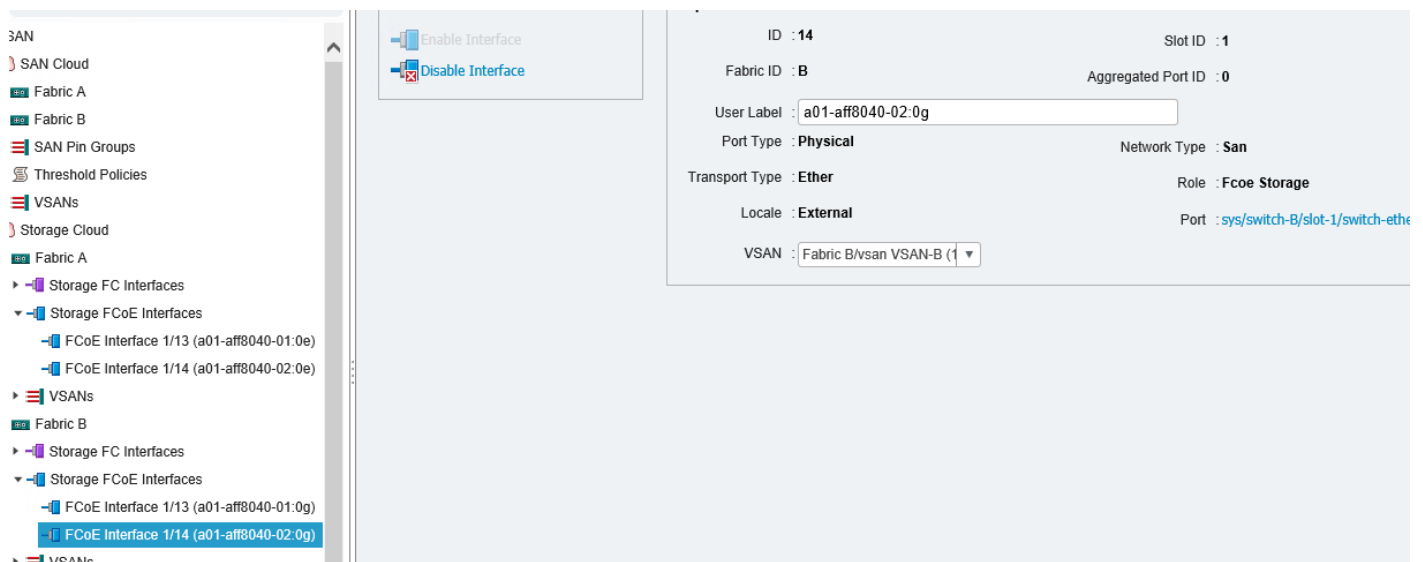
The screenshot displays the configuration page for a Storage FCoE Interface in the Network Configuration Manager (NCM). The interface is divided into two main sections: **Actions** and **Properties**.

Actions: Contains two buttons: "Enable Interface" (with a blue plus icon) and "Disable Interface" (with a blue minus icon).

Properties: Displays the following configuration details:

- ID :** 13
- Slot ID :** 1
- Fabric ID :** A
- Aggregated Port ID :** 0
- User Label :** a01-aff8040-01:0e
- Port Type :** Physical
- Network Type :** San
- Transport Type :** Ether
- Role :** Fcoe Storage
- Locale :** External
- Port :** sys/switch-A/slot-1/switch-
- VSAN :** Fabric A/vsan VSAN-A (1) ▼

40. Select the second interface that is connected to Storage Controller 02.
41. Fill in the user label with a description that describes the connection and port on the storage controller.
42. Click Save Changes and click OK.
43. Using the drop-down, select the VSAN-A VSAN.
44. Click Save Changes and click OK.
45. Expand Storage Cloud > Fabric B > Storage FCoE Interfaces. Two interfaces should be shown.
46. Select the first interface that is connected to Storage Controller 01.
47. Fill in the user label with a description that describes the connection and port on the storage controller.
48. Click Save Changes and click OK.
49. Using the pull-down, select the VSAN-B VSAN.
50. Click Save Changes and click OK.
51. Select the second interface that is connected to Storage Controller 02.
52. Fill in the user label with a description that describes the connection and port on the storage controller.
53. Click Save Changes and click OK.
54. Using the drop-down, select the VSAN-A VSAN.
55. Click Save Changes and click OK.

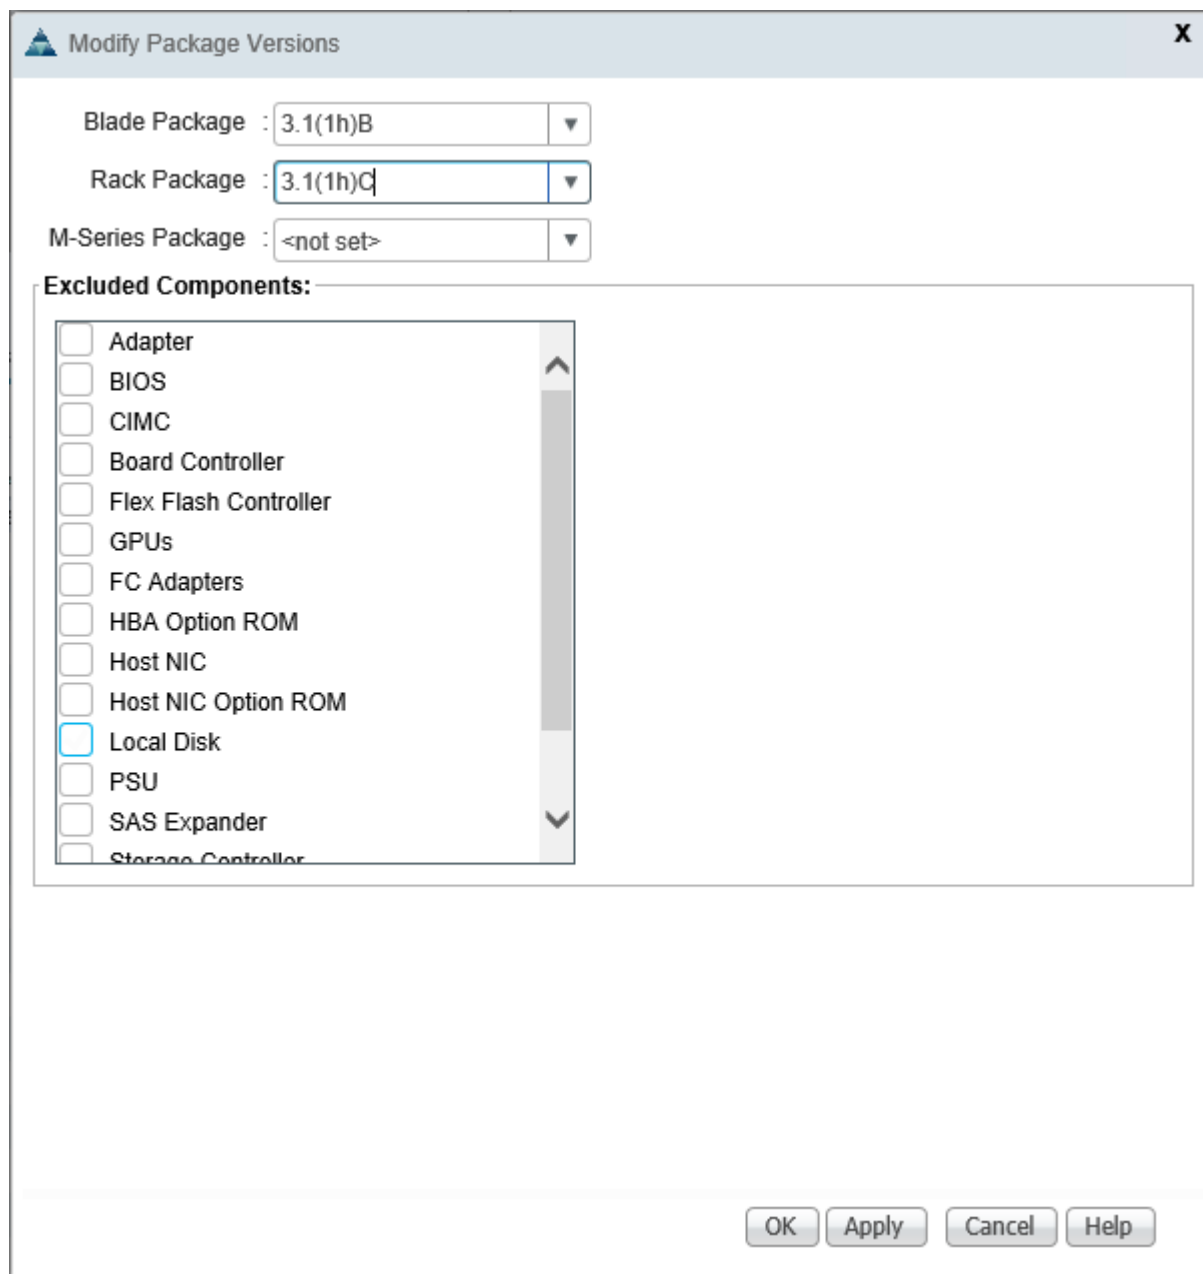


Modify Default Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for network adapters, 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, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select the default Host Firmware Package.
5. Under Actions, select Modify Package Versions.
6. Select the version 3.1(1h) for both the Blade and Rack Packages. Do not set a version for the M-Series Package.
7. Click OK to modify the host firmware package.
8. Click OK.



Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes in the bottom of the window.

6. Click Yes, then OK.

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

Create Local Disk Configuration Policy (Optional)

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

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

To create a local disk configuration policy for no local disks, complete the following steps:

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

Create Local Disk Configuration Policy

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : Disable Enable

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

FlexFlash RAID Reporting State : Disable Enable

OK Cancel

8. Click OK again.

Create Network Control Policy for Cisco Discovery Protocol and Link Layer Discovery Protocol

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

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

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

Create Network Control Policy

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

OK Cancel

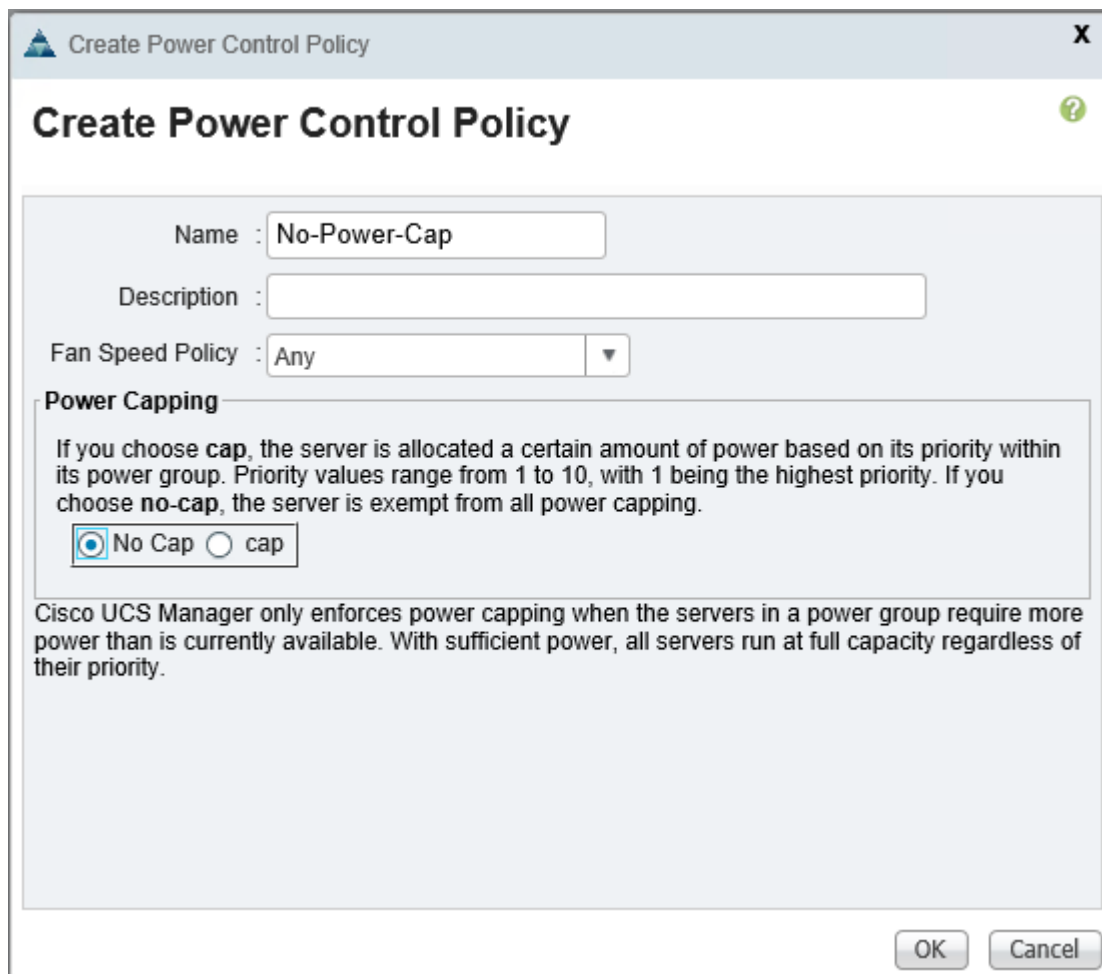
9. Click OK.

Create Power Control Policy

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

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

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



Create Power Control Policy

Name :

Description :

Fan Speed Policy :

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

No Cap cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

Create Server Pool Qualification Policy (Optional)

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



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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Enter UCSB-B200-M4 as the name for the policy.
6. Select Create Server PID Qualifications.

7. Select UCSB-B200-M4 as the PID.
8. Click OK.
9. Click OK to create the server pool policy qualification.

Create Server Pool Policy Qualification

Naming

Name :

Description :

This server pool policy qualification will apply to new or re-discovered servers. Existing servers are not qualified until they are re-discovered

Actions

- Create Adapter Qualifications
- Create Chassis/Server Qualifications
- Create Memory Qualifications
- Create CPU/Cores Qualifications
- Create Storage Qualifications
- Create Server PID Qualifications
- Create Power Group Qualifications
- Create Rack Qualifications

Qualifications

Name	Max	Model	From	To	Archite...	Speed
Server PID Qua...		UCSB-B200-M4				

Add Delete Info

Create Server BIOS Policy

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

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

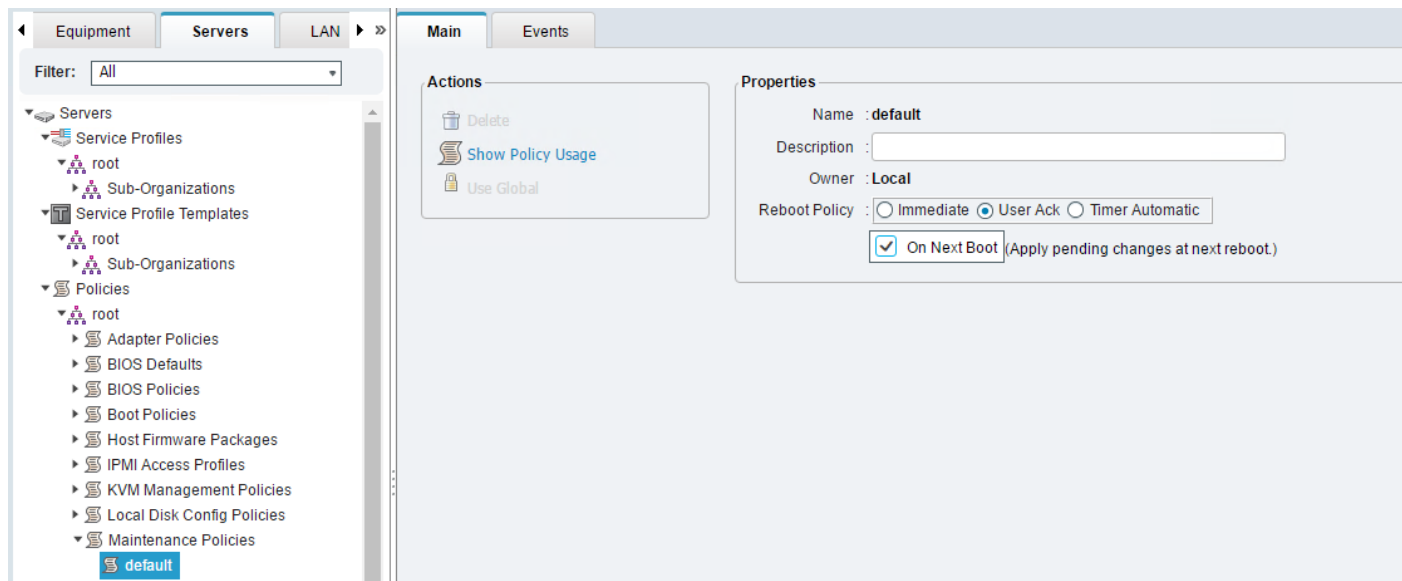


9. Click OK.

Update Default Maintenance Policy

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

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



Create vMedia Policy for VMware ESXi 6.0U1B Install Boot

It was stated at the beginning of this document that an HTTP web server should be setup and should contain necessary NetApp Data ONTAP and VMware software. The vMedia Policy created here will map the VMware ESXi 6.0u1b ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, complete the following steps:



Note: The VMware ESXi 6.0U1b ISO can be downloaded from:

https://software.cisco.com/download/release.html?mdfid=286290156&softwareid=286304568&release=6_0.U1b&reind=AVAILABLE.

1. In Cisco UCS Manager, select the Servers tab.
2. Select Policies > root.
3. Right-click vMedia Policies.
4. Select Create vMedia Policy.
5. Name the policy ESXi-6.0u1b-HTTP.
6. **Enter "Mounts Cisco Custom ISO for ESXi 6.0u1b" in the Description field.**
7. Click Add.
8. Name the mount ESXi-6.0u1b-HTTP.
9. Select the CDD Device Type.
10. Select the HTTP Protocol.
11. Enter the IP Address of the web server.



Note: Since DNS server IPs were not entered into the KVM IP pool earlier, it is necessary to enter the IP of the web server instead of the host name.

12. Enter Vmware-ESXi-6.0.0-3380124-Custom-Cisco-6.0.1.2.iso as the Remote File name.

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

14. Click OK to create the vMedia Mount.

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

Create vNIC Templates

Up to 8 vNIC Templates will be created, depending on what is being provided by the FlexPod. To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps.

Create Infrastructure vNICs

1. In Cisco UCS Manager, select the LAN tab.
2. Select Policies > root.

3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter Infra-A as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.
9. Under VLANs, select IB-Mgmt VLAN, Native-VLAN, Infra-NFS, and vMotion VLAN.
10. Set Native-VLAN as the native VLAN.
11. Select vNIC Name for CDN Source.
12. Under MTU, enter 9000.
13. From the MAC Pool list, select MAC-Pool-A.
14. From the Network Control Policy list, select Enable-CDP-LLDP.
15. Click OK to complete creating the vNIC template.

Create vNIC Template

Create vNIC Template

Template Type : Initial Template Updating Template

VLANS

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	APIC-vDS-1120	<input type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	IB-Mgmt	<input type="radio"/>
<input type="checkbox"/>	Infra-iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	Infra-iSCSI-B	<input type="radio"/>
<input checked="" type="checkbox"/>	Infra-NFS	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

+ Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(128/128)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel

16. Click OK.

17. Right-click vNIC Templates.

18. Select Create vNIC Template.

19. Enter Infra-B as the vNIC template name.

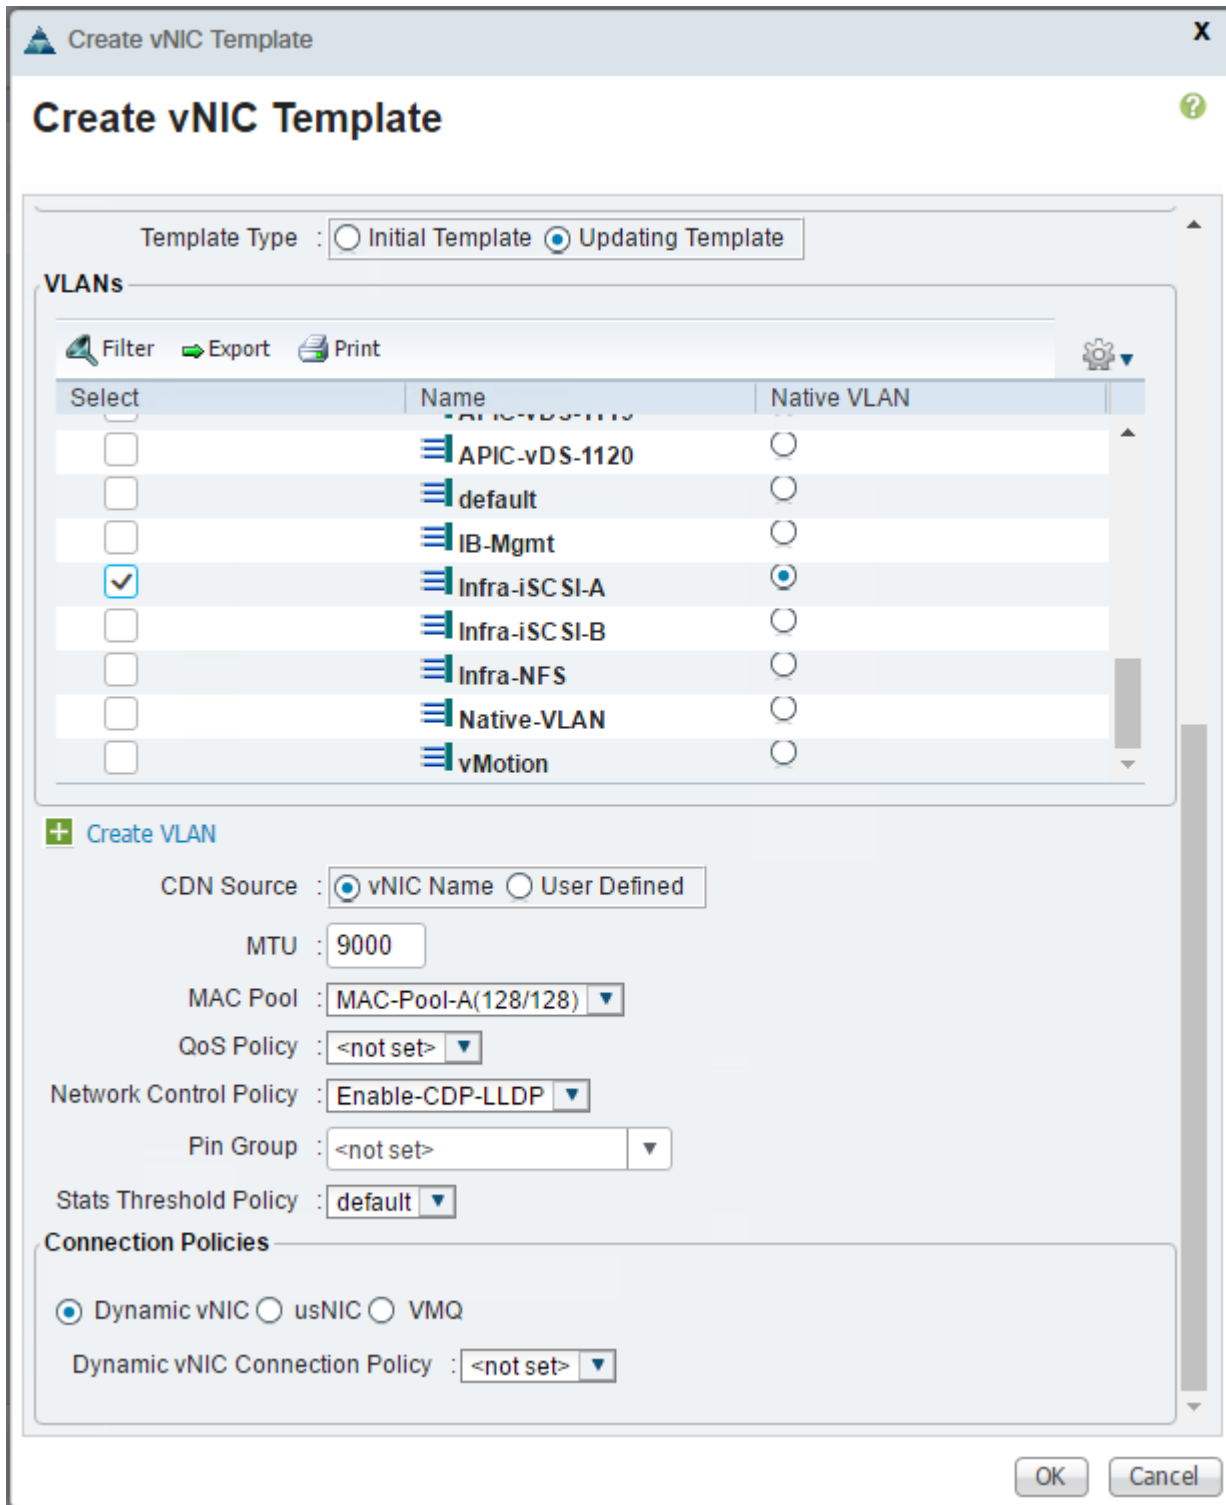
20. Select Fabric B for Fabric ID. Do not select the Enable Failover checkbox.
21. Under Target, make sure that the VM checkbox is not selected.
22. Select Updating Template for Template Type.
23. Under VLANs, select IB-Mgmt VLAN, Native-VLAN, Infra-NFS, and vMotion VLAN.
24. Set Native-VLAN as the native VLAN.
25. Select vNIC Name for the CDN Source.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select MAC-Pool-B.
28. From the Network Control Policy list, select Enable-CDP-LLDP.
29. Click OK to complete creating the vNIC template.
30. Click OK.

Create iSCSI vNICs

If you are providing iSCSI boot or providing access to iSCSI LUNs in this FlexPod, complete the steps in this section. Otherwise continue to the next section.

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter iSCSI-A as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.
9. Under VLANs, select Infra-iSCSI-A VLAN.
10. Set Infra-iSCSI-A as the native VLAN.
11. Select vNIC Name for the CDN Source.
12. Under MTU, enter 9000.
13. From the MAC Pool list, select MAC-Pool-A.

14. From the Network Control Policy list, select Enable-CDP-LLDP.
15. Click OK to complete creating the vNIC template.
16. Click OK.



17. Right-click vNIC Templates.

18. Select Create vNIC Template.
19. Enter iSCSI-B as the vNIC template name.
20. Select Fabric B as the Fabric ID. Do not select the Enable Failover checkbox.
21. Under Target, make sure that the VM checkbox is not selected.
22. Select Updating Template for Template Type.
23. Under VLANs, select Infra-iSCSI-B VLAN.
24. Set Infra-iSCSI-B as the native VLAN.
25. Select vNIC Name for the CDN Source.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select MAC-Pool-B.
28. From the Network Control Policy list, select Enable-CDP-LLDP.
29. Click OK to complete creating the vNIC template.
30. Click OK.

Create Data vNICs

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. If you are implementing the VMware vDS in this FlexPod, complete steps 3-31. Otherwise, continue to step 32.
3. Select Policies > root.
4. Right-click vNIC Templates.
5. Select Create vNIC Template.
6. Enter APIC-vDS-A as the vNIC template name.
7. Keep Fabric A selected.
8. Do not select the Enable Failover checkbox.
9. Under Target, make sure that the VM checkbox is not selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select the checkboxes for all APIC-vDS VLANs.
12. For CDN Source, select vNIC Name.

13. For MTU, enter 9000.
14. In the MAC Pool list, select MAC-Pool-A.
15. In the Network Control Policy list, select Enable-CDP-LLDP.
16. Click OK to create the vNIC template.
17. Click OK.

Create vNIC Template

Create vNIC Template

Template Type : Initial Template Updating Template

VLANS

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	ACI-System-VLAN	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1101	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1102	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1103	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1104	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1105	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1106	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1107	<input type="radio"/>

+ Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(128/128)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel

18. Right-click vNIC Templates.
19. Select Create vNIC Template
20. Enter APIC-vDS-B as the vNIC template name.
21. Select Fabric B.

22. Do not select the Enable Failover checkbox.
23. Under Target, make sure the VM checkbox is not selected.
24. Select Updating Template as the template type.
25. Under VLANs, select the checkboxes for all the APIC-vDS VLANs.
26. For CDN Source, select vNIC Name.
27. For MTU, enter 9000.
28. In the MAC Pool list, select MAC-Pool-B.
29. In the Network Control Policy list, select Enable-CDP-LLDP.
30. Click OK to create the vNIC template.

Create vNIC Template

Template Type : Initial Template Updating Template

VLANs

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	APIC-vDS-1113	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1114	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1115	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1116	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1117	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1118	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1119	<input type="radio"/>
<input checked="" type="checkbox"/>	APIC-vDS-1120	<input type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>

+ Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MAC-Pool-B(128/128)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel

31. Click OK.

32. If you are implementing the Cisco AVS in VXLAN Switching Mode in this FlexPod, complete steps 33-61. Otherwise, continue to the next section.

33. Select Policies > root.

34. Right-click vNIC Templates.
35. Select Create vNIC Template.
36. Enter APIC-AVS-A as the vNIC template name.
37. Keep Fabric A selected.
38. Do not select the Enable Failover checkbox.
39. Under Target, make sure that the VM checkbox is not selected.
40. Select Updating Template as the Template Type.
41. Under VLANs, select only the checkbox for the ACI-System-VLAN.
42. For CDN Source, select vNIC Name.
43. For MTU, enter 9000.
44. In the MAC Pool list, select MAC-Pool-A.
45. In the Network Control Policy list, select Enable-CDP-LLDP.
46. Click OK to create the vNIC template.
47. Click OK.

Create vNIC Template

Template Type : Initial Template Updating Template

VLANS

Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	ACI-System-VLAN	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1101	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1102	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1103	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1104	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1105	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1106	<input type="radio"/>
<input type="checkbox"/>	APIC-vDS-1107	<input type="radio"/>

+ Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(128/128)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel

48. Right-click vNIC Templates.
49. Select Create vNIC Template
50. Enter APIC-AVS-B as the vNIC template name.
51. Select Fabric B.

52. Do not select the Enable Failover checkbox.
53. Under Target, make sure the VM checkbox is not selected.
54. Select Updating Template as the template type.
55. Under VLANs, select only the checkbox for the ACI-System-VLAN.
56. For CDN Source, select vNIC Name.
57. For MTU, enter 9000.
58. In the MAC Pool list, select MAC-Pool-B.
59. In the Network Control Policy list, select Enable-CDP-LLDP.
60. Click OK to create the vNIC template.
61. Click OK.

Create LAN Connectivity Policies

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. If you are providing iSCSI Boot in this FlexPod, complete steps 3-74. Otherwise, continue to step 75.
3. Select LAN > Policies > root.
4. Right-click LAN Connectivity Policies.
5. Select Create LAN Connectivity Policy.
6. Enter iSCSI-Boot as the name of the policy.
7. Click the upper Add button to add a vNIC.
8. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.



Note: The two-digit number appended to the beginning of each vNIC name will work with Consistent Device Naming (CDN) to determine the vNIC/vmnic ordering placement on the VMware ESXi servers.

9. Select the Use vNIC Template checkbox.
10. In the vNIC Template list, select Infra-A.
11. In the Adapter Policy list, select VMWare.
12. Click OK to add this vNIC to the policy.

Create vNIC

Name : 00-Infra-A

Use vNIC Template :

+ Create vNIC Template

vNIC Template : Infra-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

13. Click the upper Add button to add another vNIC to the policy.
14. In the Create vNIC box, enter 01-Infra-B as the name of the vNIC.
15. Select the Use vNIC Template checkbox.
16. In the vNIC Template list, select Infra-B.
17. In the Adapter Policy list, select VMWare.
18. Click OK to add the vNIC to the policy.
19. Click the upper Add button to add a vNIC to the policy.
20. In the Create vNIC dialog box, enter 02-iSCSI-A as the name of the vNIC.
21. Select the Use vNIC Template checkbox.
22. In the vNIC Template list, select iSCSI-A.
23. In the Adapter Policy list, select VMWare.
24. Click OK to add this vNIC to the policy.

Create vNIC

Name : 02-iSCSI-A

Use vNIC Template :

+ Create vNIC Template

vNIC Template : iSCSI-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

25. Click the upper Add button to add a vNIC to the policy.
26. In the Create vNIC dialog box, enter 03-iSCSI-B as the name of the vNIC.
27. Select the Use vNIC Template checkbox.
28. In the vNIC Template list, select iSCSI-B.
29. In the Adapter Policy list, select VMWare.
30. Click OK to add this vNIC to the policy.
31. If the VMware vDS is being implemented in this FlexPod, complete steps 32-43. Otherwise, continue at step 44.
32. Click the upper Add button to add a vNIC.
33. In the Create vNIC dialog box, enter 04-APIC-vDS-A as the name of the vNIC.
34. Select the Use vNIC Template checkbox.
35. In the vNIC Template list, select APIC-vDS-A.
36. In the Adapter Policy list, select VMWare.
37. Click OK to add this vNIC to the policy.

Create vNIC

Name : 04-APIC-vDS-A

Use vNIC Template :

+ Create vNIC Template

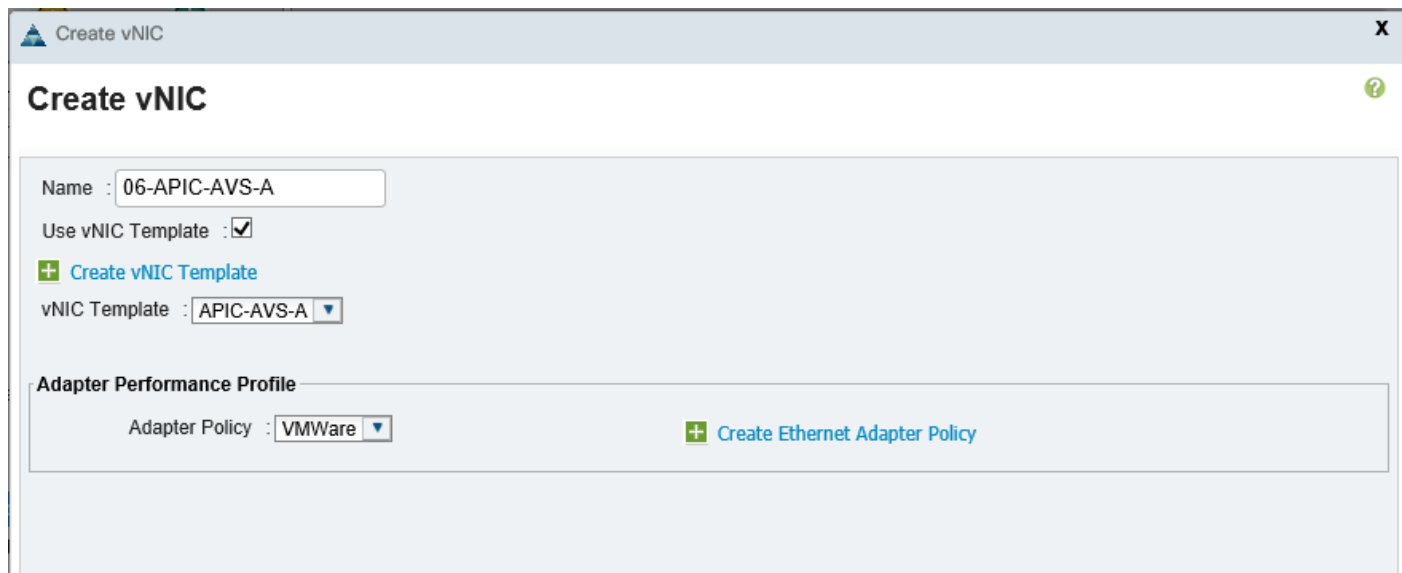
vNIC Template : APIC-vDS-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

38. Click the upper Add button to add another vNIC to the policy.
39. In the Create vNIC box, enter 05-APIC-vDS-B as the name of the vNIC.
40. Select the Use vNIC Template checkbox.
41. In the vNIC Template list, select APIC-vDS-B.
42. In the Adapter Policy list, select VMWare.
43. Click OK to add the vNIC to the policy.
44. If the Cisco AVS is being implemented in this FlexPod, complete steps 45-57. Otherwise, continue at step 58.
45. Click the upper Add button to add a vNIC.
46. In the Create vNIC dialog box, enter 06-APIC-AVS-A as the name of the vNIC.
47. Select the Use vNIC Template checkbox.
48. In the vNIC Template list, select APIC-AVS-A.
49. In the Adapter Policy list, select VMWare.
50. Click OK to add this vNIC to the policy.



Create vNIC

Name : 06-APIC-AVS-A

Use vNIC Template :

+ Create vNIC Template

vNIC Template : APIC-AVS-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

51. Click the upper Add button to add another vNIC to the policy.
52. In the Create vNIC box, enter 07-APIC-AVS-B as the name of the vNIC.
53. Select the Use vNIC Template checkbox.
54. In the vNIC Template list, select APIC-AVS-B.
55. In the Adapter Policy list, select VMWare.
56. Click OK to add the vNIC to the policy.
57. Verify that the proper vNICs have been created for your FlexPod Implementation.

Create LAN Connectivity Policy

Create LAN Connectivity Policy

Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 07-APIC-AVS-B	Derived	
vNIC 06-APIC-AVS-A	Derived	
vNIC 05-APIC-vDS-B	Derived	
vNIC 04-APIC-vDS-A	Derived	
vNIC 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

Delete Add Modify

► Add iSCSI vNICs

58. Expand the Add iSCSI vNICs section to add the iSCSI boot vNICs.

59. Click the lower Add button in the iSCSI vNIC section to define an iSCSI boot vNIC.

60. Enter iSCSI-A-Boot as the name of the vNIC.

61. Select 02-iSCSI-A for Overlay vNIC.

62. Set the iSCSI Adapter Policy to default.

63. Set the VLAN to Infra-iSCSI-A (native).

64. Leave the MAC Address set to None.

65. Click OK.

Create iSCSI vNIC

Name : iSCSI-A-Boot

Overlay vNIC : 02-iSCSI-A

iSCSI Adapter Policy : default [+ Create iSCSI Adapter Policy](#)

VLAN : Infra-iSCSI-A (native)

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

[+ Create MAC Pool](#)

66. Click the lower Add button in the iSCSI vNIC section to define an iSCSI boot vNIC.

67. Enter iSCSI-B-Boot as the name of the vNIC.

68. Set the Overlay vNIC to 03-iSCSI-B.

69. Set the iSCSI Adapter Policy to default.

70. Set the VLAN to Infra-iSCSI-B (native).

71. Leave the MAC Address set to None.

72. Click OK.

73. Verify that proper the iSCSI Boot vNICs have been created.

Create LAN Connectivity Policy

Create LAN Connectivity Policy

Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 07-APIC-AVS-B	Derived	
vNIC 06-APIC-AVS-A	Derived	
vNIC 05-APIC-vDS-B	Derived	
vNIC 04-APIC-vDS-A	Derived	
vNIC 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

Delete Add Modify

▼ Add iSCSI vNICs

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

Add Delete Modify

74. Click OK then OK again to create the LAN Connectivity Policy.
75. If you are providing FCoE Boot in this FlexPod, complete steps 76-130. Otherwise, continue in the next Section.
76. Select LAN > Policies > root.
77. Right-click LAN Connectivity Policies.
78. Select Create LAN Connectivity Policy.
79. Enter FCoE-Boot as the name of the policy.
80. Click the upper Add button to add a vNIC.

81. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.



Note: The two-digit number appended to the beginning of each vNIC name will work with Consistent Device Naming (CDN) to determine the vNIC/vmnic ordering placement on the VMware ESXi servers.

82. Select the Use vNIC Template checkbox.

83. In the vNIC Template list, select Infra-A.

84. In the Adapter Policy list, select VMWare.

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

Create vNIC

Name : 00-Infra-A

Use vNIC Template :

+ Create vNIC Template

vNIC Template : Infra-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

86. Click the upper Add button to add another vNIC to the policy.

87. In the Create vNIC box, enter 01-Infra-B as the name of the vNIC.

88. Select the Use vNIC Template checkbox.

89. In the vNIC Template list, select Infra-B.

90. In the Adapter Policy list, select VMWare.

91. Click OK to add the vNIC to the policy.

92. This LAN connectivity policy is for FCoE Boot but iSCSI LUN access can be provided, if you want to provide access to iSCSI LUNs in this FlexPod complete steps 93-104 to add iSCSI vNICs to this policy. Otherwise, continue with step 105.

93. Click the upper Add button to add a vNIC to the policy.

94. In the Create vNIC dialog box, enter 02-iSCSI-A as the name of the vNIC.

95. Select the Use vNIC Template checkbox.
96. In the vNIC Template list, select iSCSI-A.
97. In the Adapter Policy list, select VMWare.
98. Click OK to add this vNIC to the policy.

Create vNIC

Name : 02-iSCSI-A

Use vNIC Template :

[+ Create vNIC Template](#)

vNIC Template : iSCSI-A ▼

Adapter Performance Profile

Adapter Policy : VMWare ▼ [+ Create Ethernet Adapter Policy](#)

99. Click the upper Add button to add a vNIC to the policy.
100. In the Create vNIC dialog box, enter 03-iSCSI-B as the name of the vNIC.
101. Select the Use vNIC Template checkbox.
102. In the vNIC Template list, select iSCSI-B.
103. In the Adapter Policy list, select VMWare.
104. Click OK to add this vNIC to the policy.
105. If the VMware vDS is being implemented in this FlexPod, complete steps 106-117. Otherwise, continue at step 118.
106. Click the upper Add button to add a vNIC.
107. In the Create vNIC dialog box, enter 04-APIC-vDS-A as the name of the vNIC.
108. Select the Use vNIC Template checkbox.
109. In the vNIC Template list, select APIC-vDS-A.
110. In the Adapter Policy list, select VMWare.
111. Click OK to add this vNIC to the policy.

Create vNIC

Name : 04-APIC-vDS-A

Use vNIC Template :

+ Create vNIC Template

vNIC Template : APIC-vDS-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

112. Click the upper Add button to add another vNIC to the policy.
113. In the Create vNIC box, enter 05-APIC-vDS-B as the name of the vNIC.
114. Select the Use vNIC Template checkbox.
115. In the vNIC Template list, select APIC-vDS-B.
116. In the Adapter Policy list, select VMWare.
117. Click OK to add the vNIC to the policy.
118. If the Cisco AVS is being implemented in this FlexPod, complete steps 119-130. Otherwise, continue at step 131.
119. Click the upper Add button to add a vNIC.
120. In the Create vNIC dialog box, enter 06-APIC-AVS-A as the name of the vNIC.
121. Select the Use vNIC Template checkbox.
122. In the vNIC Template list, select APIC-AVS-A.
123. In the Adapter Policy list, select VMWare.
124. Click OK to add this vNIC to the policy.

Create vNIC

Name : 06-APIC-AVS-A

Use vNIC Template :

+ Create vNIC Template

vNIC Template : APIC-AVS-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Ethernet Adapter Policy

125. Click the upper Add button to add another vNIC to the policy.
126. In the Create vNIC box, enter 07-APIC-AVS-B as the name of the vNIC.
127. Select the Use vNIC Template checkbox.
128. In the vNIC Template list, select APIC-AVS-B.
129. In the Adapter Policy list, select VMWare.
130. Click OK to add the vNIC to the policy.
131. Verify that the proper vNICs have been created for your FlexPod Implementation.

Create LAN Connectivity Policy

Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 07-APIC-AVS-B	Derived	
vNIC 06-APIC-AVS-A	Derived	
vNIC 05-APIC-vDS-B	Derived	
vNIC 04-APIC-vDS-A	Derived	
vNIC 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

Delete Add Modify

[▶ Add iSCSI vNICs](#)

132. Click OK then OK again to create the policy.

Create vHBA Templates

If FCoE Boot or access to FCoE LUNs is being provided in this FlexPod, vHBA Templates must be defined in order to create vHBAs. Two vHBA Templates will be created, one for each SAN fabric. To create vHBA templates for the Cisco UCS environment, complete the following steps.

1. In Cisco UCS Manager, select the SAN tab.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter Fabric-A as the vHBA template name.
6. Leave Fabric A selected.
7. Select VSAN-A.
8. Select Initial Template for Template Type.
9. Select WWPN Pool WWPN-Pool-A
10. Click OK to complete creating the vHBA template.

Create vHBA Template

Name : Fabric-A

Description :

Fabric ID : A B

Select VSAN : VSAN-A

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-Pool-A(128/128)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

OK Cancel

11. Click OK.
12. Right-click vHBA Templates.
13. Select Create vHBA Template.
14. Enter Fabric-B as the vHBA template name.
15. Select Fabric B for the Fabric ID.
16. Select VSAN-B.
17. Select Initial Template for Template Type.
18. Select WWPN Pool WWPN-Pool-B.
19. Click OK to complete creating the vHBA template.
20. Click OK.

Create Storage Connection Policies

If FCoE Boot or access to FCoE LUNs is being provided in this FlexPod, Storage Connection Policies for LUNs in Infra-SVM should be created. To configure the necessary Infrastructure Storage Connection Policies, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > Policies > root.
3. Right-click Storage Connection Policies.
4. Select Create Storage Connection Policy.
5. Enter Infra-FCoE-A as the name of the policy.
6. Enter **“Zone LUNs from Storage Infra-SVM Fabric-A”** as the Description.
7. Select the Single Initiator Multiple Targets Zoning Type.
8. Click the Add button to add the first Target.
9. Enter the WWPN for LIF fcp_lif01a in SVM Infra-SVM. To get this value, log into the storage cluster **CLI and enter the command “network interface show -vserver Infra-SVM -lif fcp*”**.
10. Leave the Path set at A and select VSAN VSAN-A.
11. Click OK to complete creating the target.
12. Click the Add button to add the second Target.
13. Enter the WWPN for LIF fcp_lif02a in SVM Infra-SVM. To get this value, log into the storage cluster **CLI and enter the command “network interface show -vserver Infra-SVM -lif fcp*”**.
14. Leave the Path set at A and select VSAN VSAN-A.
15. Click OK to complete creating the target.

Create Storage Connection Policy

Create Storage Connection Policy

Name :

Description :

Zoning Type : None Single Initiator Single Target Single Initiator Multiple Targets

FC Target Endpoints

Filter Export Print

WWPN	Path	VSAN
20:01:00:A0:98:5B:48:16	A	VSAN-A
20:03:00:A0:98:5B:48:16	A	VSAN-A

+ Add Delete Info

OK Cancel

16. Click OK then OK again to complete creating the Storage Connection Policy.
17. Right-click Storage Connection Policies.
18. Select Create Storage Connection Policy.
19. Enter Infra-FCoE-B as the name of the policy.
20. Enter **“Zone LUNs from Storage Infra-SVM Fabric-B”** as the Description.
21. Select the Single Initiator Multiple Targets Zoning Type.
22. Click the Add button to add the first Target.
23. Enter the WWPN for LIF fcp_lif01b in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver Infra-SVM -lif fcp*”**.
24. Set the Path to B and select VSAN VSAN-B.

25. Click OK to complete creating the target.
26. Click the Add button to add the second Target.
27. Enter the WWPN for LIF fcp_lif02b in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command “**network interface show -vserver Infra-SVM -lif fcp***”.
28. Set the Path to B and select VSAN VSAN-B.
29. Click OK to complete creating the target.

Create Storage Connection Policy

Name :

Description :

Zoning Type : None Single Initiator Single Target Single Initiator Multiple Targets

FC Target Endpoints

Filter Export Print

WWPN	Path	VSAN
20:02:00:A0:98:5B:48:16	B	VSAN-B
20:04:00:A0:98:5B:48:16	B	VSAN-B

+ Add Delete Info

OK Cancel

30. Click OK then OK again to complete creating the Storage Connection Policy.

Create a SAN Connectivity Policy

If FCoE Boot or access to FCoE LUNs is being provided in this FlexPod, a SAN Connectivity Policy should be created. To configure the necessary Infrastructure SAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > Policies > root.
3. Right-click SAN Connectivity Policies.
4. Select Create SAN Connectivity Policy.
5. Enter Infra-FCoE as the name of the policy.
6. Enter **"Policy that Zones LUNs from Storage Infra-SVM"** as the Description.
7. Select the WWNN-Pool for WWNN Assignment.
8. Click the Add button to add a vHBA.
9. In the Create vHBA dialog box, enter Fabric-A as the name of the vHBA.
10. Select the Use vHBA Template checkbox.
11. In the vHBA Template list, select Fabric-A.
12. In the Adapter Policy list, select VMWare.
13. Click OK to add this vHBA to the policy.

Create vHBA

Name : Fabric-A

Use vHBA Template :

+ Create vHBA Template

vHBA Template : Fabric-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Fibre Channel Adapter Policy

OK Cancel

14. Click the Add button to add another vHBA to the policy.
15. In the Create vHBA box, enter Fabric-B as the name of the vHBA.
16. Select the Use vHBA Template checkbox.
17. In the vHBA Template list, select Fabric-B.
18. In the Adapter Policy list, select VMWare.
19. Click OK to add the vHBA to the policy.

Create SAN Connectivity Policy

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 Fabric-B	Derived
▶ vHBA Fabric-A	Derived

Delete Add Modify

OK Cancel

20. Click OK then OK again to complete creating the SAN Connectivity Policy.
21. In the list on the left under SAN Connectivity Policies, select the Infra-FCoE Policy.
22. In the center pane, select the vHBA Initiator Groups tab.
23. Select Add to add a vHBA Initiator Group.
24. Name the vHBA Initiator Group Fabric-A and select the Fabric-A vHBA Initiators.
25. Select the Infra-FCoE-A Storage Connection Policy.

Create vHBA Initiator Group



Create vHBA Initiator Group

vHBA Initiator Group

Name :

Description :

Select vHBA Initiators

Select	Name
<input checked="" type="checkbox"/>	 Fabric-A
<input type="checkbox"/>	 Fabric-B

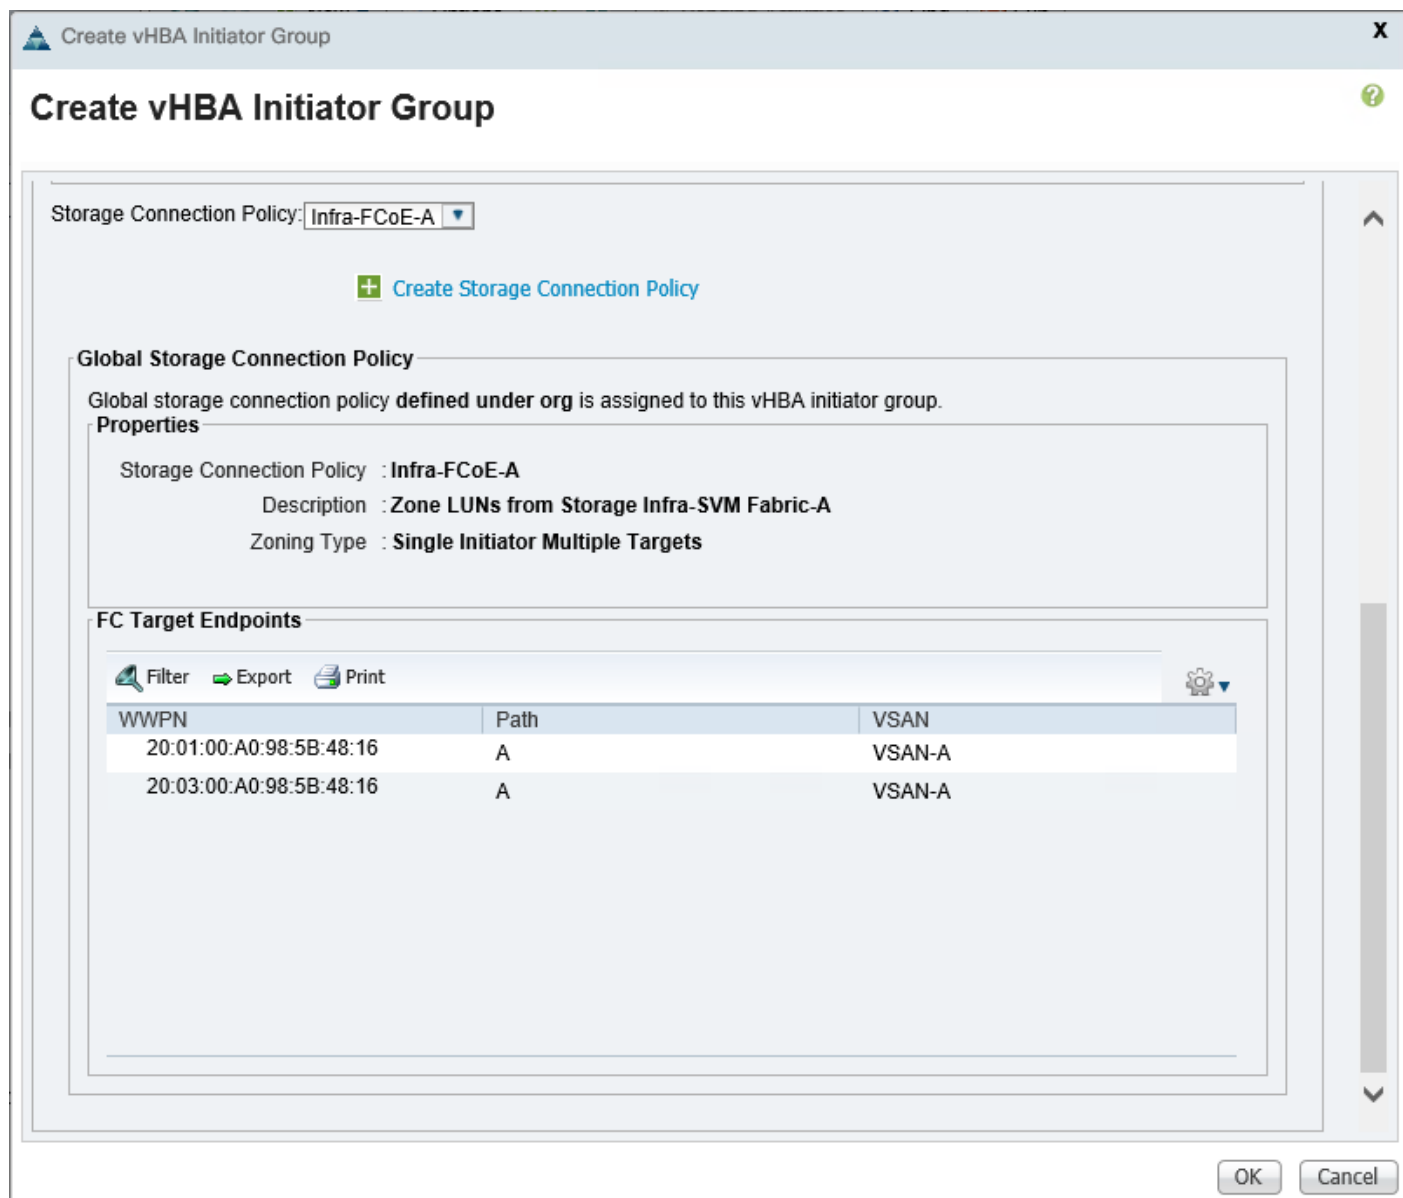
Storage Connection Policy:

[+ Create Storage Connection Policy](#)

Global Storage Connection Policy

Global storage connection policy **defined under org** is assigned to this vHBA initiator group

OK Cancel



26. Click OK then OK again to add this vHBA Initiator Group.
27. Select Add to add a vHBA Initiator Group.
28. Name the vHBA Initiator Group Fabric-B and select the Fabric-B vHBA Initiators.
29. Select the Infra-FCoE-B Storage Connection Policy.
30. Click OK then OK again to add this vHBA Initiator Group.

The screenshot shows the Cisco UCS Manager interface for configuring vHBA Initiator Groups. The left navigation pane is expanded to 'SAN > Policies > root > SAN Connectivity Policies > Infra-FCoE'. The main pane shows a table of vHBA Initiator Groups:

Name	Storage Connection Policy Name
Fabric-A	Infra-FCoE-A
Fabric-B	Infra-FCoE-B

Below the table are buttons for 'Add', 'Delete', and 'Info'. The 'Details' pane for 'Fabric-A' shows the following configuration:

- Actions:**
 - Modify vHBA Initiator Membership
 - Modify Storage Connection Policy
- Properties:**
 - Name: Fabric-A
 - Description:
- vHBA Initiators:**
 - Export
 - Print
 - Name: Fabric-A

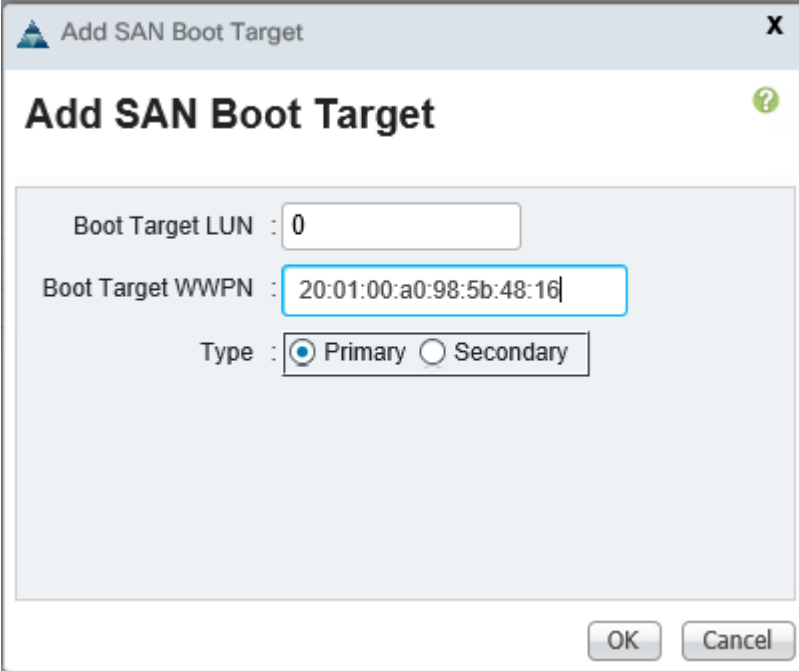
Create FCoE Boot Policy

If FCoE Boot is being provided in this FlexPod, an FCoE Boot Policy should be created. This procedure applies to a Cisco UCS environment in which two FCoE logical interfaces (LIFs) are on cluster node 1 (fcp_lif01a and fcp_lif01b) and two FCoE LIFs are on cluster node 2 (fcp_lif02a and fcp_lif02b). One boot policy is configured in this procedure. This policy configures the primary target to be fcp_lif01a.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter FCoE-Fabric-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
9. Expand the vHBAs section and select Add SAN Boot.

10. In the Add SAN Boot dialog box, enter Fabric-A.
11. Leave the Primary Type selected and click OK.
12. Select Add SAN Boot Target under vHBAs.
13. Enter the WWPN for LIF fcp_lif01a in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver Infra-SVM -lif fcp*”**.
14. Leave the Type set at Primary.



The screenshot shows a dialog box titled "Add SAN Boot Target". It contains the following fields and options:

- Boot Target LUN : 0
- Boot Target WWPN : 20:01:00:a0:98:5b:48:16
- Type : Primary Secondary

Buttons for "OK" and "Cancel" are located at the bottom right of the dialog.

15. Click OK.
16. Select Add SAN Boot Target under vHBAs.
17. Enter the WWPN for LIF fcp_lif02a in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver Infra-SVM -lif fcp*”**.
18. Leave the Type set at Secondary.
19. Click OK.
20. Under vHBAs, select Add SAN Boot.
21. In the Add SAN Boot dialog box, enter Fabric-B.
22. Leave the Secondary Type selected and click OK.
23. Select Add SAN Boot Target under vHBAs.

24. Enter the WWPN for LIF fcp_lif01b in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver Infra-SVM -lif fcp*”**.
25. Leave the Type set at Primary.
26. Click OK.
27. Select Add SAN Boot Target under vHBAs.
28. Enter the WWPN for LIF fcp_lif02b in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver Infra-SVM -lif fcp*”**.
29. Leave the Type set at Secondary.
30. Click OK.
31. Expand CIMC Mounted vMedia and select Add CIMC Mounted CD/DVD. This will allow VMware ESXi installation on a blank LUN when the vMedia Policy is enabled.

Create Boot Policy

Name : FCoE-Fabric-A

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

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

Name	Or...	vNIC/...	Type	WWN	LUN...	Slot N...	Boot...	Boot...	Descr...
Remote CD/DVD	1								
San	2								
▶ SAN Primary		Fabri...	Primary						
▶ SAN Secondary		Fabri...	Seco...						
CIMC Mounted CD/...	3								

Local Devices

vNICs

vHBAs

- Add SAN Boot
- Add SAN Boot Target

iSCSI vNICs

CIMC Mounted vMedia

- Add CIMC Mounted CD/DVD
- Add CIMC Mounted HDD

Move Up Move Down Delete

Set Uefi Boot Parameters

OK Cancel

32. Click OK then OK to complete creating the Boot Policy.

Create iSCSI Boot Policy

If iSCSI Boot is being provided in this FlexPod, an iSCSI Boot Policy should be created. This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi_lif01a and iscsi_lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi_lif02a and iscsi_lif02b). One boot policy is configured in this procedure. This policy configures the primary target to be iscsi_lif01a.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter iSCSI-Fabric-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
9. Expand the iSCSI vNICs section and select Add iSCSI Boot.
10. In the Add iSCSI Boot dialog box, enter iSCSI-A-Boot.
11. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter iSCSI-B-Boot.
14. Click OK.
15. Expand CIMC Mounted vMedia and select Add CIMC Mounted CD/DVD. This will allow VMware ESXi installation on a blank LUN when the vMedia Policy is enabled.

Create Boot Policy

Name : iSCSI-Fabric-A

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

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

Local Devices

vNICs

vHBAs

iSCSI vNICs

Add iSCSI Boot

CIMC Mounted vMedia

Add CIMC Mounted CD/DVD

Add CIMC Mounted HDD

Boot Order

Name	Or...	vNIC/...	Type	WWN	LUN...	Slot N...	Boot...	Boot...	Descr...
Remote CD/DVD	1								
iSCSI	2								
iSCSI		iSCSI...	Primary						
iSCSI		iSCSI...	Seco...						
CIMC Mounted C...	3								

Move Up Move Down Delete

Set Uefi Boot Parameters

OK Cancel

16. Click OK then OK again to save the boot policy.

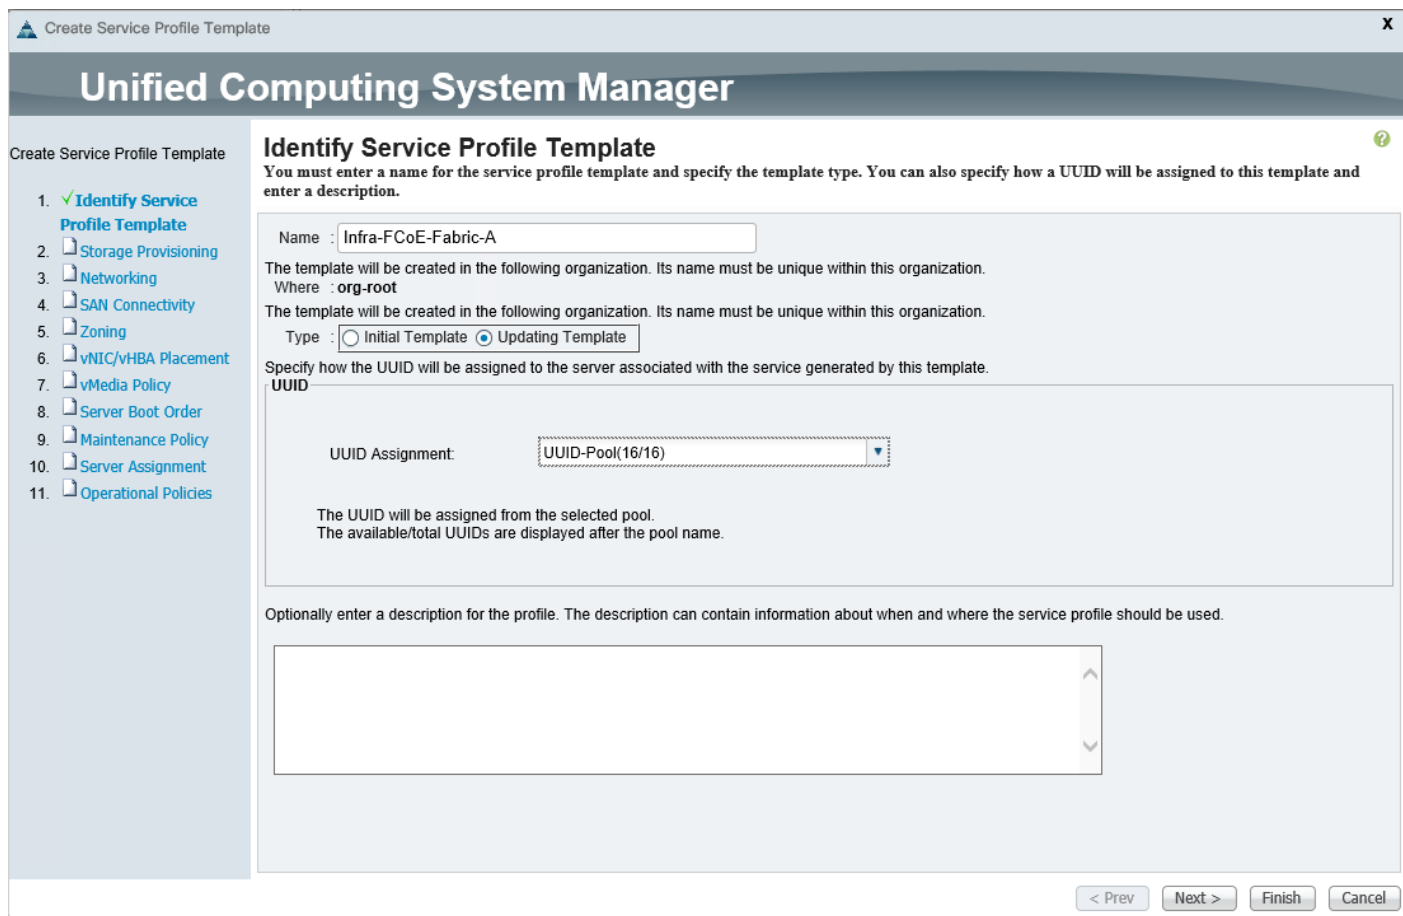
Create Infrastructure FCoE Boot Service Profile Template

If FCoE Boot is being provided in this FlexPod, an Infrastructure FCoE Boot Service Profile Template should be created. In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A FCoE boot.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.

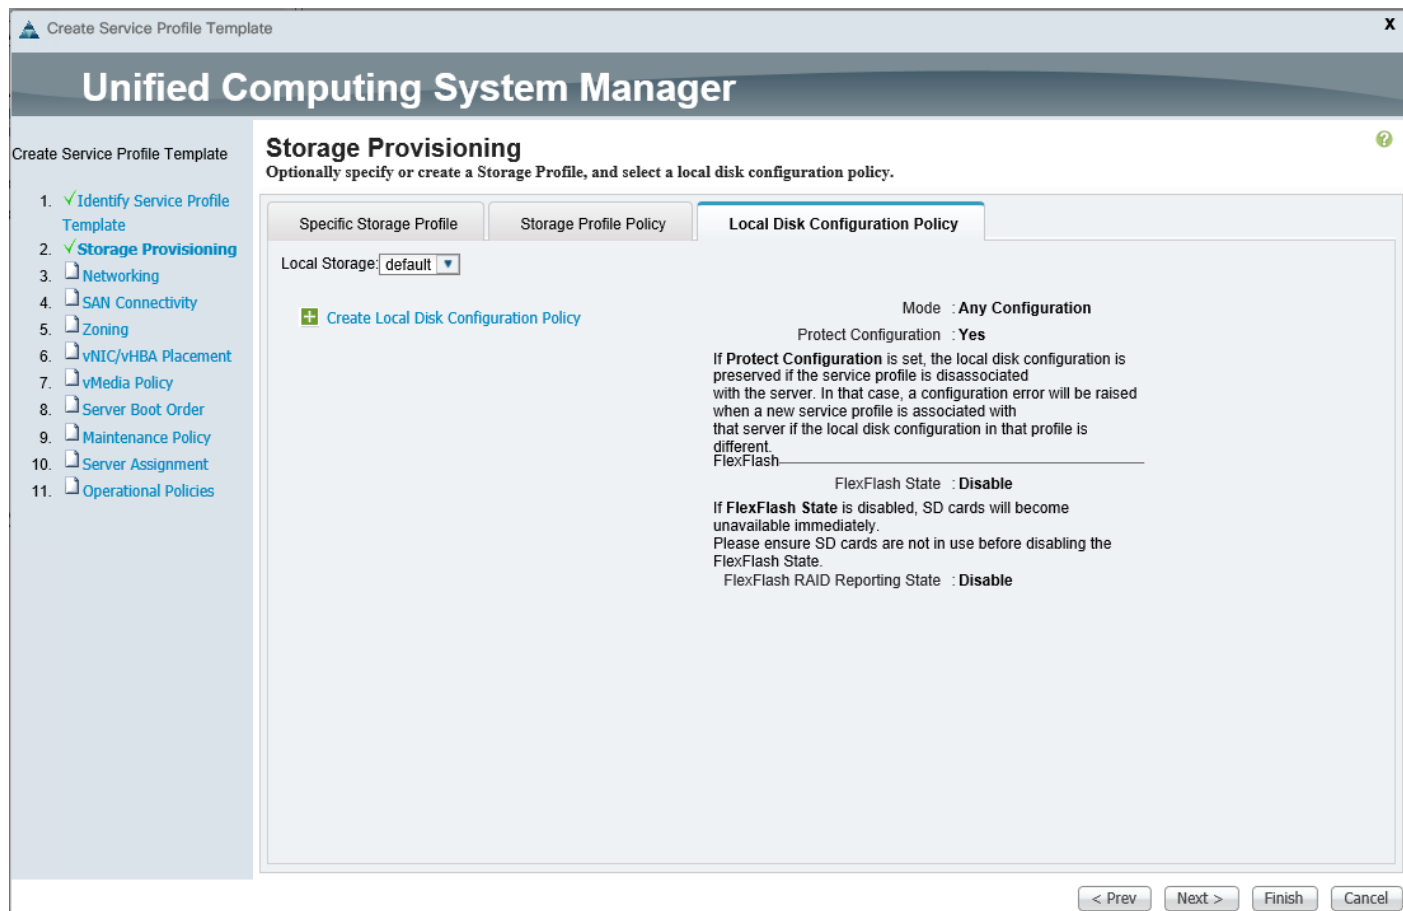
5. Enter Infra-FCoE-Fabric-A as the name of the service profile template. This service profile template is configured to boot from storage node 01 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.
8. Click Next.



Configure Storage Provisioning

To configure storage provisioning, complete the following steps:

1. Select the Local Disk Configuration Policy tab.
2. If you have servers with no physical disks, under the Local Disk Configuration Policy tab select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
3. Click Next.



Configure Networking Options

To configure networking options, complete the following steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select the FCoE-Boot LAN Connectivity Policy.
4. If providing access to iSCSI Application LUNs in this FlexPod, select IQN_Pool for Initiator Name Assignment.
5. Click Next.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) ▼

[+ Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

LAN Connectivity Policy : FCoE-Boot ▼ [+ Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment: IQN-Pool(16/16) ▼

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

Configure SAN Connectivity

To configure SAN connectivity, complete the following steps:

1. Select the `Use Connectivity Policy` option for the “How would you like to configure SAN connectivity?” field.
2. Select the `Infra-FCoE SAN Connectivity Policy`.
3. Click Next.

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The current step is 'SAN Connectivity', which is highlighted in the left-hand navigation pane. The main content area is titled 'SAN Connectivity' and includes the instruction: 'Optionally specify disk policies and SAN configuration information.' Below this, there is a question: 'How would you like to configure SAN connectivity?' with four radio button options: 'Simple', 'Expert', 'No vHBAs', and 'Use Connectivity Policy'. The 'Use Connectivity Policy' option is selected. Below the radio buttons, there is a dropdown menu for 'SAN Connectivity Policy' with 'Infra-FCoE' selected. To the right of the dropdown is a '+ Create SAN Connectivity Policy' button. At the bottom right of the wizard, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

SAN Connectivity

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

Configure Zoning

To configure zoning, complete the following steps:

1. Since a SAN Connectivity Policy and Storage Connection Policies are being used, no configuration is necessary at this point. Click Next.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Zoning

Specify zoning information

Zoning configuration involves the following steps:

1. Select vHBA Initiator(s) (vHBAs are created on storage page)
2. Select vHBA Initiator Group(s)
3. Add selected Initiator(s) to selected Initiator Group(s)

Select vHBA Initiators

Name
Fabric-A
Fabric-B

>> Add To >>

Select vHBA Initiator Groups

Name	Storage Connection Poli...
<ul style="list-style-type: none"> ▼ Fabric-A <ul style="list-style-type: none"> Storage Initiator Fabric-A ▼ Fabric-B <ul style="list-style-type: none"> Storage Initiator Fabric-B 	<ul style="list-style-type: none"> Infra-FCoE-A Infra-FCoE-B

Configure vNIC/HBA Placement

To configure vNIC/HBA placement, complete the following steps:

1. Make sure that Let System Perform Placement is selected. Since Consistent Device Naming (CDN) is being used, the system will use the vNIC names to order and place the vNICs and vHBAs.
2. Click Next.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

vNIC/vHBA Placement

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [+ Create Placement Policy](#)

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order
vHBA Fabric-A	Derived	Unspecified
vHBA Fabric-B	Derived	Unspecified
vNIC 00-Infra-A	Derived	Unspecified
vNIC 01-Infra-B	Derived	Unspecified
vNIC 02-iSCSI-A	Derived	Unspecified
vNIC 03-iSCSI-B	Derived	Unspecified
vNIC 04-APIC-vDS-A	Derived	Unspecified
vNIC 05-APIC-vDS-B	Derived	Unspecified

< Prev Next > Finish Cancel

Configure vMedia Policy

To configure the vMedia policy, complete the following steps:

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

Configure Server Boot Order

To configure the server boot order, complete the following steps:

1. Select FCoE-Fabric-A for Boot Policy.
2. Click Next to continue to the next section.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [+ Create Boot Policy](#)

Name : FCoE-Fabric-A
Description :

Reboot on Boot Order Change : **No**
Enforce vNIC/vHBA/iSCSI Name : **Yes**
Boot Mode : **Legacy**

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

Boot Order

[+](#) [-](#) [Filter](#) [Export](#) [Print](#)

Name	Order	vNIC/vH...	Type	WWN	LUN Na...	Slot Nu...	Boot Na...	Boot Path	Descript...
CIMC Mounted CD/DVD	3								
San	2								
Remote CD/DVD	1								

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

< Prev Next > Finish Cancel

Configure Maintenance Policy

To configure the maintenance policy, complete the following steps:

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

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The main title is 'Unified Computing System Manager'. The current step is 'Maintenance Policy', which is highlighted in the left-hand navigation pane. The main content area contains the following text: 'Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.' Below this, there is a section titled 'Maintenance Policy' with a dropdown menu set to 'default' and a '+ Create Maintenance Policy' button. The details for the selected policy are: Name : default, Description : , Reboot Policy : User Ack, and Config. Trigger State : On Next Boot. At the bottom right, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Maintenance Policy

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: default ▼ [+ Create Maintenance Policy](#)

Name : default
Description :
Reboot Policy : User Ack
Config. Trigger State : On Next Boot

< Prev Next > Finish Cancel

Configure Server Assignment

To configure the server assignment, complete the following steps:

1. In the Pool Assignment list, select `Infra-Pool1`.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management at the bottom of the page and select `default` from the Host Firmware list.
5. Click Next.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Server Assignment

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:

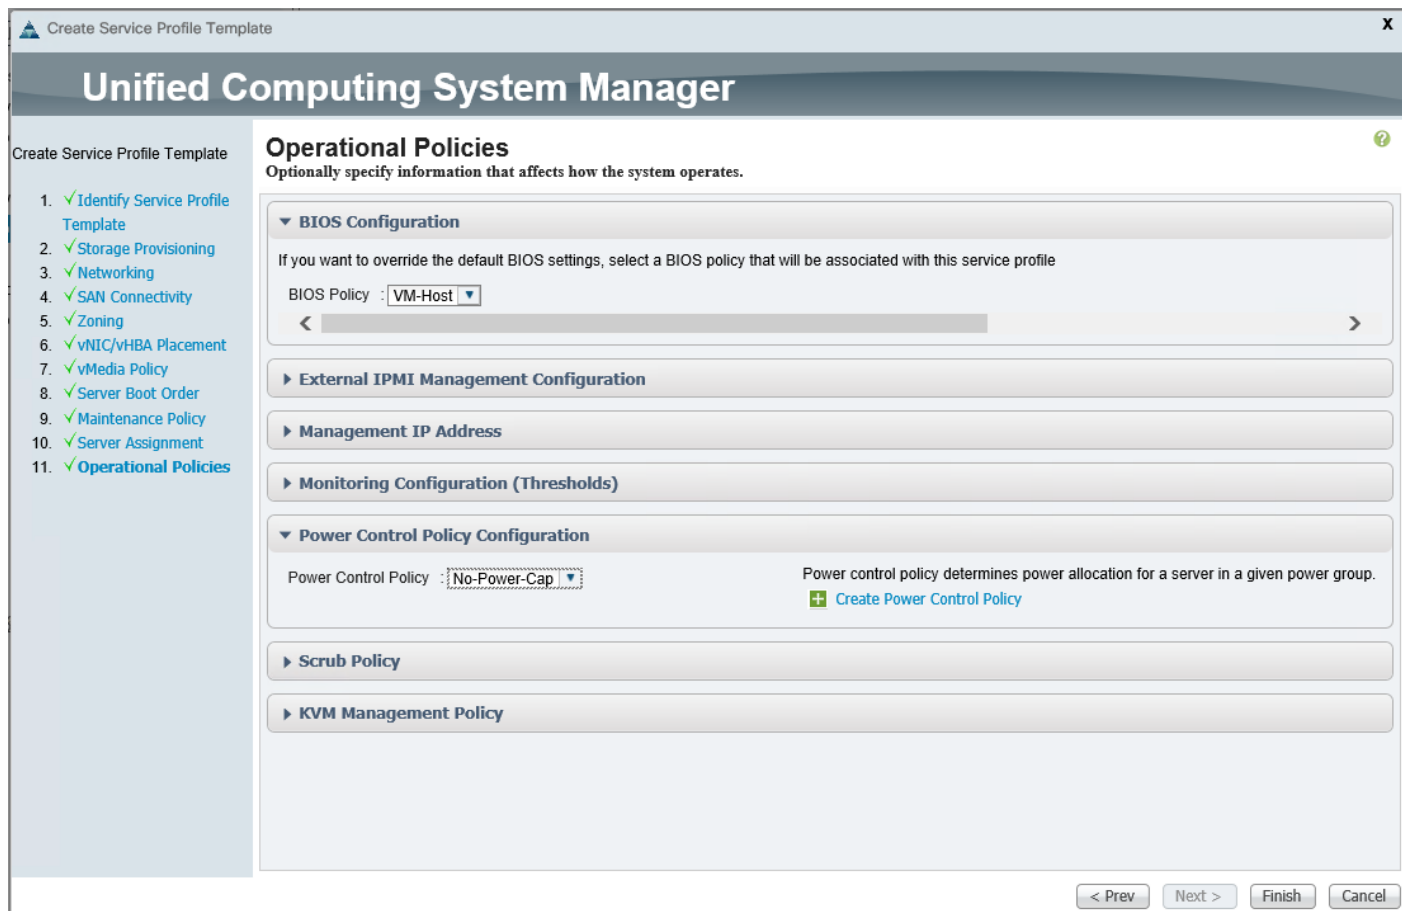
[+ Create Host Firmware Package](#)

< Prev Next > Finish Cancel

Configure Operational Policies

To configure the operational policies, complete the following steps:

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



3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.
5. Under Service Profile Templates > root, right-click the newly created Service Template Infra-FCoE-Fabric-A and select Create a Clone.
6. Name the clone Infra-FCoE-Fabric-A-vm.
7. Click OK, then OK again to create the clone.
8. Select the newly-cloned Infra-FCoE-Fabric-A-vm service Profile Template.
9. In the center pane, select the vMedia Policy tab.
10. Select Modify vMedia Policy.
11. Select the ESXi-6.0u1b-HTTP vMedia Policy and click OK.
12. Click OK then OK again to complete modifying the Service Profile Template.

Create iSCSI Boot Service Profile Template

If iSCSI Boot is being provided in this FlexPod, an Infrastructure iSCSI Boot Service Profile Template should be created. In this procedure, one service profile template for Infrastructure ESXi hosts is created for fabric A boot.

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

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

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The current step is 'Identify Service Profile Template'. The wizard is titled 'Create Service Profile Template' and has a close button (X) in the top right corner. The main title is 'Unified Computing System Manager'. On the left, there is a navigation pane with 11 steps: 1. Identify Service Profile Template (checked), 2. Storage Provisioning, 3. Networking, 4. SAN Connectivity, 5. Zoning, 6. vNIC/vHBA Placement, 7. vMedia Policy, 8. Server Boot Order, 9. Maintenance Policy, 10. Server Assignment, and 11. Operational Policies. The main content area is titled 'Identify Service Profile Template' and includes a help icon (question mark). Below the title, there is a text box for 'Name' containing 'Infra-iSCSI-Fabric-A'. Below that, there are two lines of text: 'The template will be created in the following organization. Its name must be unique within this organization. Where : org-root' and 'The template will be created in the following organization. Its name must be unique within this organization. Type : Initial Template Updating Template'. Below this, there is a section for 'Specify how the UUID will be assigned to the server associated with the service generated by this template.' with a sub-section 'UUID' containing a 'UUID Assignment' dropdown menu set to 'UUID-Pool(16/16)'. Below the dropdown, there is a note: 'The UUID will be assigned from the selected pool. The available/total UUIDs are displayed after the pool name.' At the bottom, there is a text box for 'Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.' and a set of navigation buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

Configure Storage Provisioning

To configure the storage provisioning, complete the following steps:

1. Select the Local Disk Configuration Policy tab.
2. If you have servers with no physical disks, select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
3. Click Next.



Configure Networking Options

To configure the networking options, complete the following steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. **Select the “Use Connectivity Policy” option to configure the LAN connectivity.**
3. Select the iSCSI-Boot LAN Connectivity Policy.
4. Select IQN_Pool for Initiator Name Assignment.
5. Click Next.

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The current step is 'Networking', which is highlighted in the left-hand navigation pane. The main content area is titled 'Networking' and includes the instruction 'Optionally specify LAN configuration information.' Below this, there are several configuration options:

- Dynamic vNIC Connection Policy:** A dropdown menu set to 'Select a Policy to use (no Dynamic vNIC Policy by default)'. A '+ Create Dynamic vNIC Connection Policy' button is visible below it.
- How would you like to configure LAN connectivity?** Radio buttons for 'Simple', 'Expert', 'No vNICs', and 'Use Connectivity Policy' (which is selected).
- LAN Connectivity Policy:** A dropdown menu set to 'iSCSI-Boot'. A '+ Create LAN Connectivity Policy' button is visible below it.
- Initiator Name:** A section with an 'Initiator Name Assignment' dropdown set to 'IQN-Pool(16/16)'. Below it is an 'Initiator Name' field with a '+ Create IQN Suffix Pool' button. A note states: 'The IQN will be assigned from the selected pool. The available/total IQNs are displayed after the pool name.'

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

Configure SAN Connectivity

To configure the SAN connectivity, complete the following steps:

1. If access to FCoE storage is not being provided in this FlexPod, select the **No vHBAs** option for the “How would you like to configure SAN connectivity?” field and **continue on to the next section.**
2. If access to FCoE storage is being provided in this FlexPod, select the **Use Connectivity Policy** option for the “How would you like to configure SAN connectivity?” field.
3. To provide access to FCoE LUNs in Infra-SVM, select the **Infra-FCoE SAN Connectivity Policy.**
4. Click **Next.**

Configure Zoning

1. It is not necessary to configure any Zoning options. Click **Next.**

Configure vNIC/HBA Placement

To configure the vNIC/HBA placement, complete the following steps:

1. In the “**Select Placement**” list, select the **Let System Perform Placement.**
2. Click **Next.**

vNIC/vHBA Placement
Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [+ Create Placement Policy](#)

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order
vNIC 00-Infra-A	Derived	Unspecified
vNIC 01-Infra-B	Derived	Unspecified
vNIC 02-iSCSI-A	Derived	Unspecified
vNIC 03-iSCSI-B	Derived	Unspecified
vNIC 04-APIC-vDS-A	Derived	Unspecified
vNIC 05-APIC-vDS-B	Derived	Unspecified
vNIC 06-APIC-AVS-A	Derived	Unspecified
vNIC 07-APIC-AVS-B	Derived	Unspecified

< Prev Next > Finish Cancel


Configure vMedia Policy

To configure the vMedia policy, complete the following steps:

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

Configure Server Boot Order

To configure the server boot order, complete the following steps:

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

`iscsi show -vserver Infra-SVM` **Note or copy the iSCSI target name for Infra-SVM.**

8. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from `Infra-SVM`
9. Enter the IP address of `iscsi_lif02a` for the IPv4 Address field. You can query the iSCSI LIF IP addresses by typing `network interface show -vserver Infra-SVM -lif iscsi*`.

Create iSCSI Static Target

iSCSI Target Name : e6bf8e00a0985b4700:vs.3

Priority : 1


Port : 3260

Authentication Profile : <not set> [+ Create iSCSI Authentication Profile](#)

IPv4 Address : 192.168.110.19

ID : 0

OK Cancel

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

▲ Set iSCSI Boot Parameters
✕

Set iSCSI Boot Parameters ?

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

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-A(12/12) ▼

IPv4 Address : 0.0.0.0

Subnet Mask : 255.255.255.0

Default Gateway : 0.0.0.0

Primary DNS : 0.0.0.0

Secondary DNS : 0.0.0.0

+ [Create IP Pool](#)

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

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addr...	LUN Id
iqn.1992-08.c...	1	3260		192.168.110.19	0
iqn.1992-08.c...	2	3260		192.168.110.18	0

+ Add
 🗑 Delete
📄 Info


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

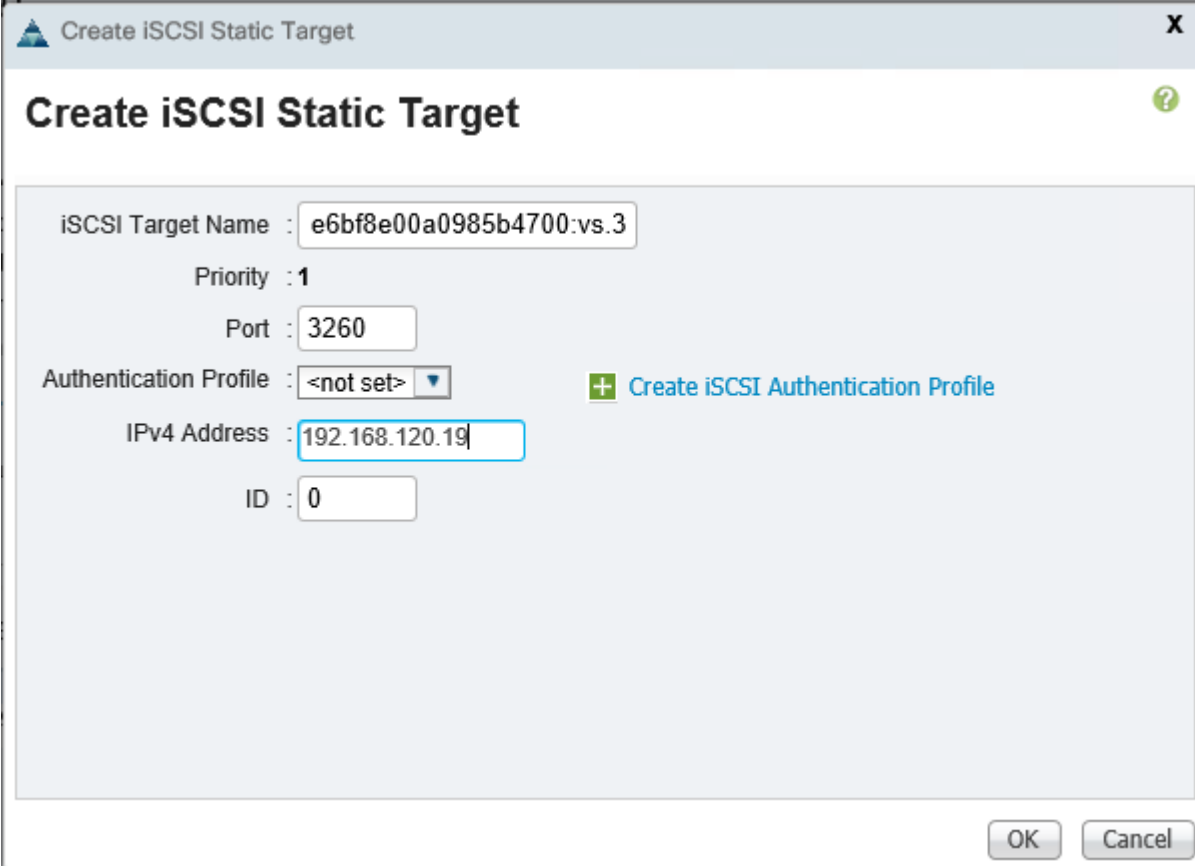
OK
Cancel

15. Click OK.

16. In the Boot Order pane, select iSCSI-B-Boot.

17. Click the Set iSCSI Boot Parameters button.

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




Create iSCSI Static Target

iSCSI Target Name : e6bf8e00a0985b4700:vs.3

Priority : 1


Port : 3260

Authentication Profile : <not set>  Create iSCSI Authentication Profile

IPv4 Address : 192.168.120.19

ID : 0

OK Cancel

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

Set iSCSI Boot Parameters
X

Set iSCSI Boot Parameters

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

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-B(12/12)

IPv4 Address : 0.0.0.0

Subnet Mask : 255.255.255.0

Default Gateway : 0.0.0.0

Primary DNS : 0.0.0.0

Secondary DNS : 0.0.0.0

[+ Create IP Pool](#)

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

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addr...	LUN Id
iqn.1992-08.c...	1	3260		192.168.120.19	0
iqn.1992-08.c...	2	3260		192.168.120.18	0

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

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

OK
Cancel

28. Click OK.

29. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

30. Click Next to continue to the next section.

Configure Maintenance Policy

To configure the maintenance policy, complete the following steps:

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



Configure Server Assignment

To configure the server assignment, complete the following steps:

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

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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

Server Assignment

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

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

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**

Operational Policies

Optionally specify information that affects how the system operates.

BIOS Configuration

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

BIOS Policy :

External IPMI Management Configuration

Management IP Address

Monitoring Configuration (Thresholds)

Power Control Policy Configuration

Power Control Policy :

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

[+ Create Power Control Policy](#)

Scrub Policy

KVM Management Policy

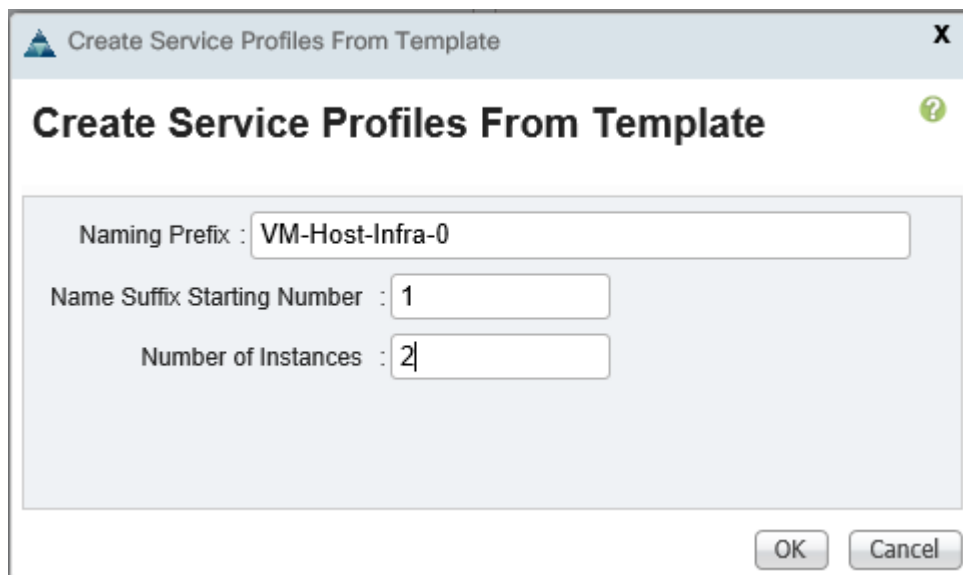
< Prev Next > Finish Cancel

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.
5. Under Service Profile Templates > root, right-click the newly created Service Template Infra-iSCSI-Fabric-A and select Create a Clone.
6. Name the clone Infra-iSCSI-Fabric-A-vM.
7. Click OK, then OK again to create the clone.
8. Select the newly-cloned Infra-iSCSI-Fabric-A-vM service Profile Template.
9. Under Properties, select the vMedia Policy tab.
10. Select Modify vMedia Policy.
11. Select the ESXi-6.0u1b-HTTP vMedia Policy and click OK.
12. Click OK to complete modifying the Service Profile Template.

Create Service Profiles

Service Profiles can now be created from the Service Profile Templates created above. First, create the **Service Profile from the “-vM” template with the vMedia Policy. With this policy, the server will boot into the VMware ESXi Installation.** Once VMware ESXi has been installed on the boot LUN, the Service Profile can be bound to the template **without the “-vM” suffix, removing the vMedia mounted ESXi Installation mount.** To create service profiles from the service profile template, complete the following steps:

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



8. Click OK in the confirmation message.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations. A following section in this document will also address adding servers for application tenants.

Gather Necessary Information

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

Table 21 iSCSI LIFs for iSCSI IQN

Vserver	iSCSI Target IQN
Infra-SVM	



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

Table 22 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN
VM-Host-Infra-01	
VM-Host-Infra-02	



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

Table 23 Table 7 vHBA WWPNs for fabric A and fabric B

Cisco UCS Service Profile Name	WWPN
VM-Host-Infra-01	Fabric-A
	Fabric-B
VM-Host-Infra-02	Fabric-A
	Fabric-B



Note: To gather the vHBA WWPN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root. Click each service profile and **then click the vHBAs.** The WWPNs are shown in the center pane.

Storage Configuration - SAN Boot

Clustered Data ONTAP iSCSI Boot Storage Setup

Create igroups

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

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

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

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



Use the values listed in Table 21 and Table 22 for the ION information.



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

Clustered Data ONTAP FCoE Boot Storage Setup

Create igroups

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

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol fcp -ostype
vmware -initiator <vm-host-infra-01-fabric-a-wwpn>,<vm-host-infra-01-fabric-a-
wwpn>
```

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol fcp -ostype
vmware -initiator <vm-host-infra-02-fabric-a-wwpn>,<vm-host-infra-02-fabric-a-
wwpn>
```

```
igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol fcp -ostype vmware
-initiator <vm-host-infra-01-fabric-a-wwpn>,<vm-host-infra-01-fabric-a-wwpn>,<vm-
host-infra-02-fabric-a-wwpn>,<vm-host-infra-02-fabric-a-wwpn>
```



Note: Use the values listed in Table 23 for the WWPN information.



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

Map Boot LUNs to igroups

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

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

```
lun show -m
```

Cisco ACI Fabric Configuration

The following section provides a detailed procedure for configuring the Cisco ACI Fabric and Infrastructure (Foundation) Tenant for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration. The following table shows the End Point Group (EPG) VLANs, Subnets, and Bridge Domains used in this lab validation.

Table 24 Lab Validation Infrastructure (Foundation) Tenant Configuration

EPG	Storage VLAN	UCS VLAN	External VLAN	Subnet / Gateway	Bridge Domain
IB-MGMT	N/A	DVS	163	172.26.163.0/24 - L2	BD-common-Internal
Core-Services	N/A	363	163	172.26.163.10/24	BD-common-Internal
SVM-MGMT	263	N/A	163	172.26.163.0/24 - L2	BD-common-Internal
iSCSI-A	3010	3110	N/A	192.168.110.0/24 - L2	BD-iSCSI-A
iSCSI-B	3020	3120	N/A	192.168.120.0/24 - L2	BD-iSCSI-B
NFS-LIF	3050	N/A	N/A	192.168.150.0/24 - L2	BD-NFS
NFS-VMK	N/A	3150	N/A	192.168.150.0/24 - L2	BD-NFS
vMotion	N/A	3000	N/A	192.168.100.0/24 - L2	BD-Internal
VMware vDS Pool	N/A	1101-1120	N/A	Varies	Varies
ACI System VLAN for AVS	N/A	4093	N/A	Varies	Varies

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as covered in 0.

In ACI, both spine and leaf switches are configured using APIC, individual configuration of the switches is not required. The APIC discovers the ACI infrastructure switches using LLDP and acts as the central control and management point for the entire configuration.

Cisco Application Policy Infrastructure Controller (APIC) Setup

This section details the setup of the Cisco APIC. Cisco recommends a cluster of at least 3 APICs controlling an ACI Fabric. To setup the Cisco APIC, complete the following steps:

1. On the back of the first APIC, connect the M port and the Eth1-1 port to the out of band management switch. The M port will be **used for connectivity to the server's Cisco Integrated Management Controller (CIMC)** and the Eth1-1 port will provide SSH access to the APIC CLI.
2. Using the supplied KVM dongle cable, connect a keyboard and monitor to the first APIC. Power on the machine, and using <F8> enter the CIMC Configuration Utility. Configure the CIMC with an IP address on the out of band management subnet. Make sure Dedicated NIC Mode and No NIC Redundancy are selected to put the CIMC interface on the M port. Also set the CIMC password.
3. Save the CIMC configuration using <F10> and use <ESC> to exit the configuration tool.
4. Using a web browser, browse to <https://<cimc-ip-address>>.

5. Make sure the CIMC Version is 2.0(3i). If it is not, go to [Cisco UCS C220 M3 downloads](#) and download the version 2.0(3i) Cisco UCS Host Upgrade Utility and upgrade all server firmware following the [Cisco Host Upgrade Utility 2.0\(3\) User Guide](#).
6. Login with the admin used id and the password entered in the CIMC setup.
7. From the Server tab on the left, select Summary and click Launch KVM.
8. The KVM Application will launch. Press <Enter> to bring up the APIC setup utility.

```

File View Macros Tools Power Virtual Media Help
i=B++o o =..o |
+-----+
ssh-dss AAAAB3NzaC1kc3MAAACBAJfFc79NpS2nJ3c0NyZHW7LqootexyCPQTZHEBX0aLZ5a3J+AgNF
wkJEE1lnCSg3wu0s/YNn+foQfInxa00DyeB8FtenXAY/5e4/PROnffAUADADCF8tqkd261a4I30Jfy6B
D+EoP25NC2j/DhPkAhsUm+gPjlgam+tLw05c0aP/AAAFQDZsdSH3++0MM7v+4k2muSef93EPQAAAIA
g1dR4Z65iG94MGDygLXUugp3TlJgR9s9q79Nzd3dfPXaXPyiET3EDhrm9F016bibJvkXDhiz7SHC8M5M
7zfV09KJP1znpIQeM99/0Ml+cV+GCa+u6jmWglvin+s62MWhD5J2jtDe1Ewk9KwKkkfR1T4SdjuYF8n+
n/kvAuV00gAAAIB/FmjGN1aHXBCv16vUWPoXILTC0wbuBqoQ47Jz5K9JR8HnSV1RA8nBUPookPJm7Bj9
zBkHFpRHegq2sULu06uY/jPJ4qgtUGe5LGEs9ruIWQU4WKiqcMALp5qe5xd8hQ5RU9Xgld/HyKuFgUAR
JdDmVI7WM0reLxzPqw3FNKZGGw==

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

*Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.

Press Enter at anytime to assume the default values. Use ctrl-c
at anytime to restart from the beginning.

Cluster configuration ...
Enter the fabric name [ACI Fabric1]:

```

9. Press <Enter> to accept the default fabric name. This value can be changed if desired.
10. Press <Enter> to select the default value for Enter the number of controllers in the fabric. While the fabric can operate with a single APIC, 3 APICs are recommended for redundancy.
11. Enter the controller number currently being set up under Enter the controller ID (1-3). Please remember only controller number 1 will allow you to setup the admin password. Remaining controllers and switches sync their passwords to the admin password set on the controller 1.
12. Enter the controller name or press <Enter> to accept the default.
13. Press <Enter> to select the default pool under Enter the address pool for TEP addresses. If this subnet is already in use, please enter a different range.
14. Enter the VLAN ID for the fabric's infra network or the fabric's system VLAN. Do not enter 4094 because that VLAN is reserved in the Cisco UCS. A recommended VLAN ID for this VLAN is 4093. If

you set up networking for the Cisco AVS in the preceding UCS section, use the same VLAN used there for the ACI-System-VLAN.

15. Press <Enter> to select the default address pool for Bridge Domain (BD) multicast addresses.
16. Press <Enter> to disable IPv6 for the Out-of-Band Mgmt Interface.
17. Enter an IP and subnet length in the out of band management subnet for the Out-of-band management interface.
18. Enter the gateway IP address of the out of band management subnet.
19. Press <Enter> to select auto speed/duplex mode.
20. Press <Enter> to enable strong passwords.
21. Enter the password for the admin user.
22. Reenter this password.
23. The full configuration is shown. If all values are correct, press <Enter> not to edit the configuration.
24. The APIC will continue configuration and continue bootup until the login: prompt appears.
25. Repeat the above steps for all APIC controllers adjusting as necessary.

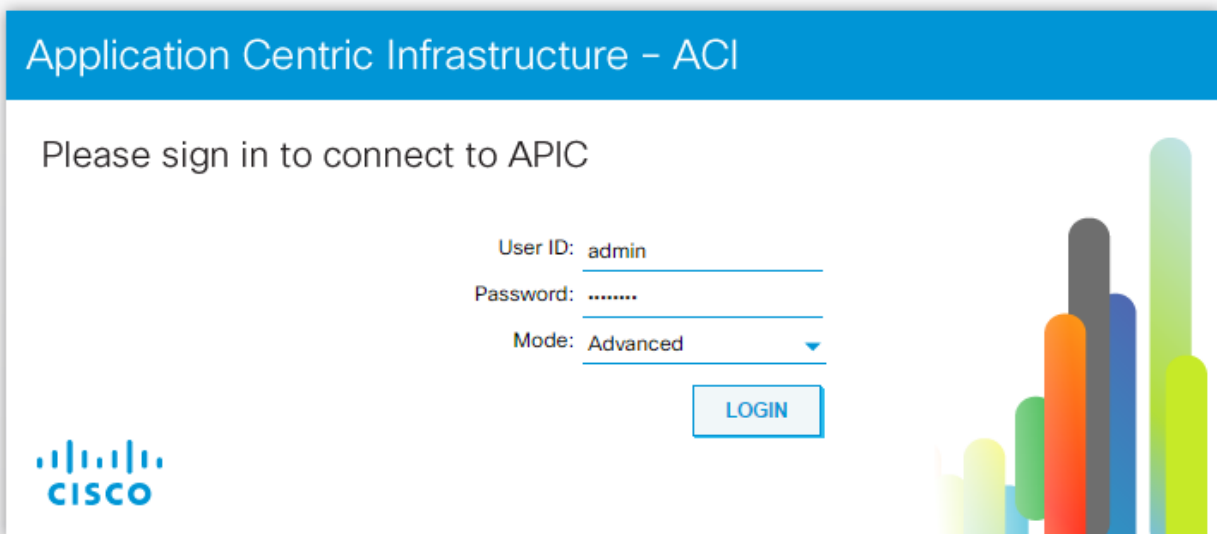
Cisco ACI Fabric Discovery

This section details the steps for Cisco ACI Fabric Discovery, where leaves, spines and APICs are automatically discovered in the ACI Fabric and assigned node ids. Cisco recommends a cluster of at least 3 APICs controlling an ACI Fabric.

1. Log into the APIC Advanced GUI using a web browser, by browsing to <https://<apic1-IP>>. Select the Advanced Mode and login with the admin user id and password.



Note: In this validation, Google Chrome was used as the web browser. Make sure the Advanced GUI choice is shown.



Version 1.3(2f)

© 2012-2016 Cisco Systems, Inc. All rights reserved.

The copyrights to certain works contained in this software are owned by other third parties and used and distributed under license. Certain components of this software are licensed under the GNU General Public License (GPL) version 2.0 or the GNU GPL 2.0 and LGPL 2.1

[Terms and Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks](#)

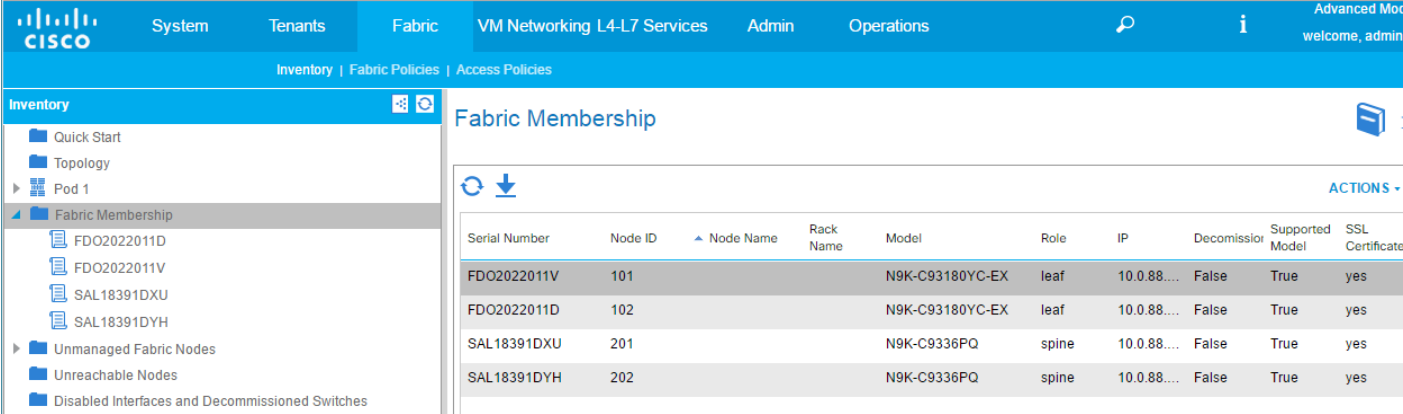
2. Select No to close the Warning.
3. At the top, select Fabric.
4. On the left, select and expand Fabric Membership.
5. Connect to the two leaves and two spines using serial consoles and login in as admin with no password. Use show inventory to get the leaf serial number.

```
(none)# show inventory
```

```
NAME: "Chassis", DESCR: "Nexus C93180YC-EX Chassis"
```

```
PID: N9K-C93180YC-EX , VID: V01 , SN: FDO2022011V
```

6. Match up the serial numbers from the leaf listing to determine whether leaf 1 or leaf 2 appears under Fabric Membership.
7. In the APIC GUI, under Fabric Membership, double click the leaf in the list. Enter a Node ID (101 or 102 **recommended**), and a **Node Name** that will become the leaf's switch name. Click **Update**.
8. As the fabric discovery continues, both spines and leaves will start appearing under Fabric Membership. It may be necessary to click the refresh button to see new items in the list. Repeat steps 4-6 to assign Node IDs and Node Names to these switches. It is recommended to use 100-based Node IDs for leaves and 200-based Node IDs for spines. Continue this process until all switches have been assigned Node IDs and Node Names. All switches will also receive IPs in the TEP address space assigned in the APIC.



The screenshot shows the Cisco APIC GUI with the 'Fabric Membership' section active. The left sidebar shows the 'Inventory' tree with 'Fabric Membership' selected. The main area displays a table of discovered nodes.

Serial Number	Node ID	Node Name	Rack Name	Model	Role	IP	Decommissioned	Supported Model	SSL Certificate
FDO2022011V	101			N9K-C93180YC-EX	leaf	10.0.88....	False	True	yes
FDO2022011D	102			N9K-C93180YC-EX	leaf	10.0.88....	False	True	yes
SAL18391DXU	201			N9K-C9336PQ	spine	10.0.88....	False	True	yes
SAL18391DYH	202			N9K-C9336PQ	spine	10.0.88....	False	True	yes

9. Click **Topology**. The discovered ACI Fabric topology will appear. It may take a few minutes and you will need to click the refresh button for the complete topology to appear. This topology shows 2 leaves, 2 spines, and 2 APICs. 2 APICs were used in this lab validation. Cisco recommends a cluster of at least 3 APICs in a production environment.

Initial ACI Fabric Setup

This section details the steps for initial setup of the Cisco ACI Fabric, where the software release is validated, out of band management IPs are assigned to the leaves and spines, NTP is setup, and the fabric BGP route reflectors are set up.

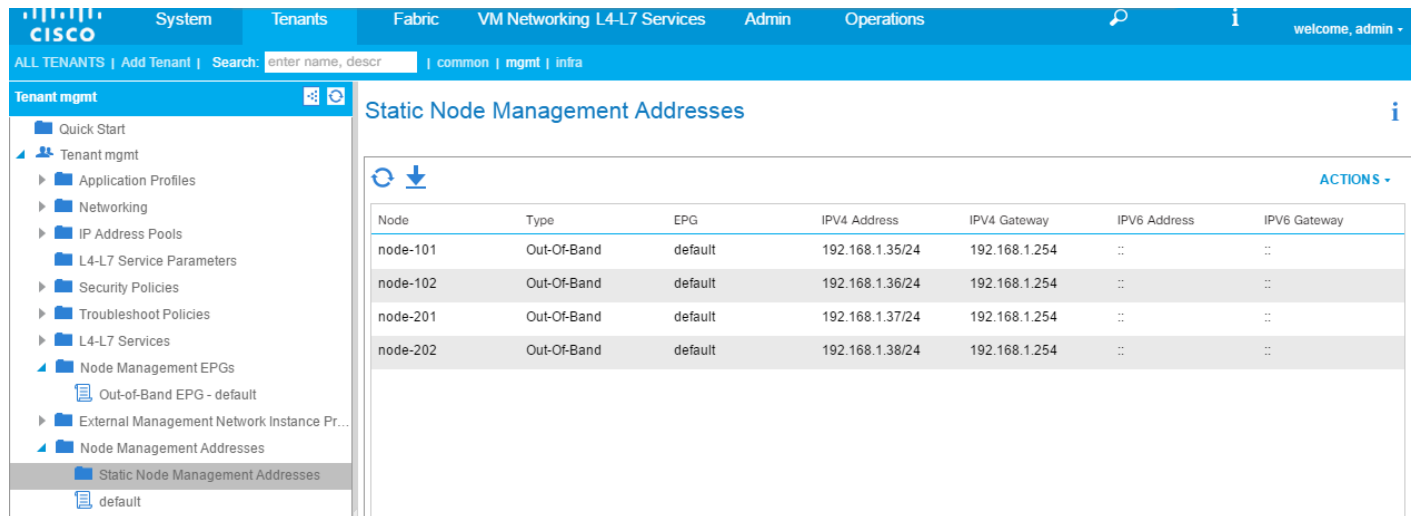
1. In the APIC Advanced GUI, at the top select Admin > Firmware.
2. This document was validated with ACI software release 1.3(2f). Select Fabric Node Firmware on the left. All switches should have the same release of firmware and should at a minimum be at release n9000-11.3 (2f). The switch software version should also match the APIC version. Also, the Default Firmware Version should at a minimum be at release n9000-11.3 (2f).

The screenshot shows the Cisco ACI Fabric Configuration interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'VM Networking L4-L7 Services', 'Admin', and 'Operations'. The left sidebar shows 'Firmware Management' with sub-items: 'Quick Start', 'Fabric Node Firmware', 'Controller Firmware', 'Catalog Firmware', 'Firmware Repository', and 'Download Tasks'. The main content area is titled 'Fabric Node Firmware' and includes a 'Policy' tab. Under 'Firmware Default Policy', the 'Default Firmware Version' is set to 'n9000-11.3(2f)'. Below this, a table titled 'All Nodes' displays the following data:

Node id	Node name	Model	Current Firmware	Status	Role	Firmware Group	Maintenance Group
Current Firmware: n9000-11.3(2f) (4 Nodes)							
101	a01-9318...	N9K-C93180...	n9000-11.3(2f)	Upgraded successfully on 201...	leaf		
102	a01-9318...	N9K-C93180...	n9000-11.3(2f)	Upgraded successfully on 201...	leaf		
201	a02-9336-1	N9K-C9336PQ	n9000-11.3(2f)	Upgraded successfully on 201...	spine		
202	a02-9336-2	N9K-C9336PQ	n9000-11.3(2f)	Upgraded successfully on 201...	spine		

3. If the software releases are not at the minimum level, do not match in all switches, or you cannot set the Default Firmware Version, follow the [Cisco APIC Controller and Switch Software Upgrade and Downgrade Guide](#) to upgrade both the APICs and switches to a minimum release of 1.3(2f).
4. If the APICs have not already been upgraded, click on Admin > Firmware > Controller Firmware. If all APICs are not at the same release at a minimum of 1.3(2f), follow the [Cisco APIC Controller and Switch Software Upgrade and Downgrade Guide](#) to upgrade both the APICs and switches to a minimum release of 1.3(2f).
5. To add out of band management interfaces for all the switches in the ACI Fabric, select Tenants > mgmt.
6. Expand Tenant mgmt on the left. Right-click Node Management Addresses and select Create Static Node Management Addresses.
7. Enter the node number range (101-102) for the leaf switches.
8. Select the checkbox for Out-of-Band Addresses.
9. Select default for Out-of-Band Management EPG.
10. Considering that the IPs will be applied in a consecutive range of two IPs, enter a starting IP address and netmask in the Out-Of-Band IPV4 Address field.
11. Enter the out of band management gateway address in the Gateway field.
12. Click SUBMIT, then click YES.
13. On the left, right-click Node Management Addresses and select Create Static Node Management Addresses.
14. Enter the node number range (201-202) for the spine switches.

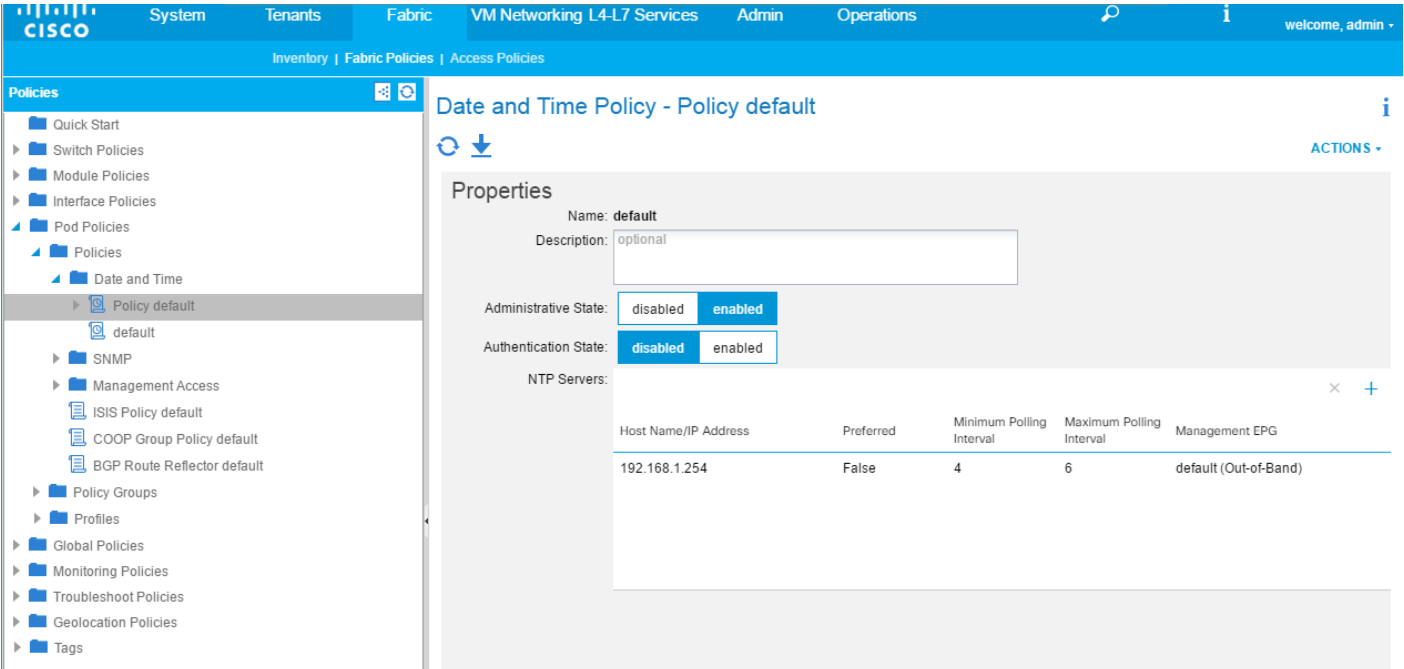
15. Select the checkbox for Out-of-Band Addresses.
16. Select default for Out-of-Band Management EPG.
17. Considering that the IPs will be applied in a consecutive range of two IPs, enter a starting IP address and netmask in the Out-Of-Band IPV4 Address field.
18. Enter the out of band management gateway address in the Gateway field.
19. Click SUBMIT, then click YES.
20. On the left, expand Node Management Addresses and select Static Node Management Addresses. Verify the mapping of IPs to switching nodes.



Node	Type	EPG	IPV4 Address	IPV4 Gateway	IPV6 Address	IPV6 Gateway
node-101	Out-Of-Band	default	192.168.1.35/24	192.168.1.254	::	::
node-102	Out-Of-Band	default	192.168.1.36/24	192.168.1.254	::	::
node-201	Out-Of-Band	default	192.168.1.37/24	192.168.1.254	::	::
node-202	Out-Of-Band	default	192.168.1.38/24	192.168.1.254	::	::

21. On the left, expand Security Policies and select Out-of-Band Contracts.
22. Right-click Out-of-Band Contracts and select Create Out-of-Band Contract.
23. Name the contract oob-default.
24. Click the + sign to the right of Subjects.
25. Name the Subject oob-default.
26. Click the + sign to the right of Filters to add a filter to the Filter Chain.
27. Use the drop-down to select the common/default filter. Click UPDATE.
28. Click OK to complete Creating the Contract Subject.
29. Click SUBMIT to complete creating the contract.
30. On the left, expand Node Management EPGs and select Out-of-Band EPG - default.
31. Click the + sign to the right of Provided Out-of-Band Contracts. Select the mgmt/oob-default OOB Contract and click UPDATE.

32. At the bottom right, click SUBMIT.
33. You should now be able to connect to any of the switches with ssh.
34. To set up NTP in the fabric, select and expand Fabric > Fabric Policies > Pod Policies > Policies > Date and Time.
35. Select default. In the Datetime Format - default pane, use the drop-down to select the correct Time Zone. Select the appropriate Display Format and Offset State. Click SUBMIT.
36. On the left, select Policy default.
37. On the right use the + sign to add NTP servers accessible on the out of band management subnet. Enter an IP address accessible on the out of band management subnet and select the default (Out-of-Band) Management EPG. Click Submit to add the NTP server. Repeat this process to add all NTP servers.



The screenshot displays the Cisco ACI GUI configuration page for the 'Date and Time Policy - Policy default'. The left-hand navigation pane shows a tree structure under 'Policies' > 'Date and Time' > 'Policy default', with 'default' selected. The main content area is titled 'Date and Time Policy - Policy default' and includes a 'Properties' section with the following details:

- Name: default
- Description: optional
- Administrative State: disabled / **enabled**
- Authentication State: disabled / **enabled**

Below the properties is the 'NTP Servers' section, which contains a table with one server entry:

Host Name/IP Address	Preferred	Minimum Polling Interval	Maximum Polling Interval	Management EPG
192.168.1.254	False	4	6	default (Out-of-Band)

38. To configure optional DNS in the ACI fabric, select and expand Fabric > Fabric Policies > Global Policies > DNS Profiles > default.
39. In the Management EPG drop-down, select the default (Out-of-Band) Management EPG.
40. Use the + signs to the right of DNS Providers and DNS Domains to add DNS servers and the DNS domain name. Note that the DNS servers should be reachable from the out of band management subnet. Click SUBMIT to complete the DNS configuration.
41. To configure the BGP Route Reflector, which will be used to distribute Layer 3 routes within the ACI fabric, select and expand Fabric > Fabric Policies > Pod Policies > Policies > BGP Route Reflector default.

42. Select a unique Autonomous System Number for this ACI fabric. Use the + sign on the right to add the two spines to the list of Route Reflector Nodes. Click SUBMIT to complete configuring the BGP Route Reflector.

The screenshot displays the Cisco ACI GUI for configuring a BGP Route Reflector Policy. The breadcrumb navigation shows 'Inventory | Fabric Policies | Access Policies'. The left sidebar is titled 'Policies' and includes a tree view with categories like 'Quick Start', 'Switch Policies', 'Module Policies', 'Interface Policies', 'Pod Policies', and 'Monitoring Policies'. The 'BGP Route Reflector default' policy is selected. The main panel shows the 'Properties' section for this policy, with the following details:

- Name: default
- Description: optional
- Autonomous System Number: 101
- Route Reflector Nodes: A table with two entries.

Node ID	Node Name	Description
201	a02-9336-1	
202	a02-9336-2	

43. To enable the BGP Route Reflector, on the left just under the BGP Route Reflector default, right-click Policy Groups under Pod Policies and select Create Pod Policy Group.
44. In the Create Pod Policy Group window, name the Policy Group pod1-policygrp.
45. Select the default BGP Route Reflector Policy.

Create Pod Policy Group

Specify the Policy Group properties

Name:

Description:

Date Time Policy:

ISIS Policy:

COOP Group Policy:

BGP Route Reflector Policy:

Management Access Policy:

SNMP Policy:

46. Click SUBMIT to complete creating the Policy Group.

47. On the left expand Profiles under Pod Policies and select default.

48. Using the drop-down, select the pod1-policygrp Fabric Policy Group.

The screenshot shows the Cisco ACI GUI interface. The top navigation bar includes System, Tenants, Fabric, VM Networking L4-L7 Services, Admin, and Operations. The breadcrumb trail is Inventory | Fabric Policies | Access Policies. The left sidebar shows a tree view under 'Policies' with 'Pod Policies' expanded to show 'Profiles' and 'default' selected. The main content area is titled 'Pod Selector - default' and shows the 'Properties' section with the following configuration: Name: default, Description: optional, Type: ALL, and Fabric Policy Group: pod1-policygrp. The 'Policy' tab is active, and there are 'Faults' and 'History' tabs visible.

49. Click SUBMIT to complete selecting this Policy Group.

Fabric Access Policy Setup

This section details the steps to create various access policies creating parameters for CDP, LLDP, LACP, etc. These policies will be used during vPC and VM domain creation. To define fabric access policies, complete the following steps:

1. In the APIC Advanced GUI, select and expand Fabric > Access Policies > Interface Policies > Policies.
2. On the left, right-click Link Level and select Create Link Level Policy.
3. Name the policy 1Gbps-Auto and select the 1Gbps Speed.

Create Link Level Policy
i
✕

Specify the Physical Interface Policy Identity

Name:

Description:

Label:

Auto Negotiation:

Speed: ▼

Link debounce interval (msec): ▲▼

Forwarding Error Correction:

4. Click SUBMIT to complete creating the policy.
5. On the left, right-click Link Level and select Create Link Level Policy.
6. Name the policy 10Gbps-Auto and select the 10Gbps Speed.
7. Click SUBMIT to complete creating the policy.

8. On the left, right-click CDP Interface and select Create CDP Interface Policy.
9. Name the policy CDP-Enabled and select the Enabled Admin State.

Create CDP Interface Policy

Specify the CDP Interface Policy Identity

Name: CDP-Enabled

Description: optional

Label:

Admin State: Disabled Enabled

10. Click SUBMIT to complete creating the policy.
11. On the left, right-click CDP Interface and select Create CDP Interface Policy.
12. Name the policy CDP-Disabled and select the Disabled Admin State.
13. Click SUBMIT to complete creating the policy.
14. On the left, right-click LLDP Interface and select Create LLDP Interface Policy.
15. Name the policy LLDP-Enabled and select Enabled for both the Transmit and Receive State.

Create LLDP Interface Policy i X

Specify the LLDP Interface Policy Properties

Name:

Description:

Label:

Receive State: Disabled Enabled

Transmit State: Disabled Enabled

16. Click SUBMIT to complete creating the policy.
17. On the left, right-click LLDP Interface and select Create LLDP Interface Policy.
18. Name the policy LLDP-Disabled and select Disabled for both the Transmit and Receive State.
19. Click SUBMIT to complete creating the policy.
20. On the left, right-click Port Channel and select Create Port Channel Policy.
21. Name the policy LACP-Active and select LACP Active for the Mode. Do not change any of the other values.

Create Port Channel Policy

Specify the Port Channel Policy

Name: LACP-Active

Description: optional

Label:

Mode: LACP Active

Control:

- Fast Select Hot Standby Ports
- Graceful Convergence
- Load Defer Member Ports
- Suspend Individual Port

Minimum Number of Links: 1
Not Applicable for FEX PC/VPC

Maximum Number of Links: 16
Not Applicable for FEX PC/VPC

22. Click SUBMIT to complete creating the policy.

23. On the left, right-click Port Channel and select Create Port Channel Policy.

24. Name the policy MAC-Pinning and select MAC Pinning for the Mode. Do not change any of the other values.

Create Port Channel Policy

Specify the Port Channel Policy

Name: **MAC-Pinning**

Description: optional

Label:

Mode: **MAC Pinning**

Minimum Number of Links: **1**

Maximum Number of Links: **16**

Not Applicable for FEX PC/VPC

Not Applicable for FEX PC/VPC

SUBMIT **CANCEL**

25. Click SUBMIT to complete creating the policy.

26. On the left, right-click Spanning Tree Interface and select Create Spanning Tree Interface Policy.

27. Name the policy BPDU-Filter-Guard and select both the BPDU filter and BPDU Guard Interface Controls.

Create Spanning Tree Interface Policy

Define the STP Interface Policy

Name:

Description:

Label:

Interface controls: BPDU filter enabled
 BPDU Guard enabled

28. Click SUBMIT to complete creating the policy.

29. On the left, right-click Spanning Tree Interface and select Create Spanning Tree Interface Policy.

30. Name the policy No-BPDU-Filter-Guard and make sure both the BPDU filter and BPDU Guard Interface Controls are cleared.

Create Spanning Tree Interface Policy

Define the STP Interface Policy

Name:

Description:

Label:

Interface controls: BPDU filter enabled
 BPDU Guard enabled

31. Click SUBMIT to complete creating the policy.

32. On the left, right-click L2 Interface and select Create L2 Interface Policy.

33. Name the policy VLAN-Scope-Global and make sure Global scope is selected.

Create L2 Interface Policy

Define the L2 Interface Policy

Name:

Description:

VLAN Scope: Global scope Port Local scope

34. Click SUBMIT to complete creating the policy.

35. On the left, right-click Firewall and select Create Firewall Policy.

36. Name the policy Firewall-Disabled and select Disabled for Mode. Do not change any of the other values.

Create Firewall Policy
i X

Specify the Firewall Policy Properties

Name: **Firewall-Disabled**

Description:

Mode: Disabled Enabled Learning

SysLog

Administrative State: enabled ▼

Included Flows: **Denied flows**

Polling Interval (seconds): 60 ▲▼

Log Level: information ▼

Dest Group: select an option ▼

SUBMIT
CANCEL

37. Click SUBMIT to complete creating the policy.

Create Virtual Port Channels (vPCs)

This section details the steps for setup of vPCs for connectivity to NetApp Storage, the Cisco UCS, and the In-Band Management network. To create the virtual port channels, complete the following steps:

1. In the APIC Advanced GUI, at the top select Fabric > Inventory > Topology.
2. On the right, select Configure.
3. Click ADD SWITCHES. Select both leaves and select ADD SELECTED.
4. On both switches, select the port connected for your In-Band Management connection.



Note: We are assuming connecting a vPC here to a port channel on another switch where the In-Band Management VLAN is connected.

5. At the bottom right, select CONFIGURE VPC.

- For Policy Group Name, enter VPC-IB-MGMT-In. Select the appropriate policies as shown in the screenshot.



Note: In this lab validation, the In-Band Management VLAN was connected to a Cisco Catalyst switch on two 1GE trunk ports.

BACK TO SUMMARY

CONFIGURING VPC

■ Port Channel ■ VPC ■ L2 Interface
■ L3 ■ Conn. to Fex ■ Selected

a01-93180-1 (Node-101) ✕

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a01-93180-2 (Node-102) ✕

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

Policy Group Name: VPC-IB-MGMT-In

Description: optional

Link Level Policy: 1Gbps-Auto ▼

CDP Policy: CDP-Enabled ▼

MCP Policy: default ▼

LLDP Policy: LLDP-Disabled ▼

STP Interface Policy: No-BPDU-Filter-Guard ▼

Egress Data Plane Policing Policy: default ▼

Ingress Data Plane Policing Policy: default ▼

Port Channel Policy: LACP-Active ▼

Storm Control Interface Policy: default ▼

L2 Interface Policy: VLAN-Scope-Global ▼

Attached Entity Profile: select an option ▼

- For the Attached Entity Policy drop-down, select Create Attachable Access Entity Profile.
- Name the profile AEP-IB-MGMT-In. Click the + sign to the right to add an L2 External Domain.
- In the drop-down, select Create Layer 2 Domain.
- In the Create Layer 2 Domain window, name the Domain L2-IB-MGMT-In.
- In the VLAN Pool drop-down, select create VLAN Pool.
- In the create VLAN Pool window, name the VLAN Pool VP-IB-MGMT. Select Static Allocation.
- Click the + sign to add an Encapsulation Block.
- Enter the incoming In-Band Management VLAN for both the From and To parts of the Range. Select Static Allocation.

Create Ranges

Specify the Encap Block Range

Type: **VLAN**

Range: 163 - 163
From To

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

OK CANCEL

15. Click OK to complete creating the range.

Create VLAN Pool

Specify the Pool identity

Name: VP-IB-MGMT

Description: optional

Allocation Mode: Dynamic Allocation Static Allocation

Encap Blocks: × +

VLAN Range	Allocation Mode
[163]	Static Allocation

SUBMIT CANCEL

16. Click SUBMIT to complete creating the VLAN Pool.

Create Layer 2 Domain
i ✕

Specify the Layer 2 Domain

Name: L2-IB-MGMT-In

VLAN Pool: VP-IB-MGMT(static) ▼ 📄

SUBMIT
CANCEL

17. Click SUBMIT to complete creating the Layer 2 Domain.

Create Attachable Access Entity Profile
i ✕

Specify the name, domains and infrastructure encaps

Name: AEP-IB-MGMT-In

Description:

Enable Infrastructure VLAN:

Domains (VMM, Physical or External) To Be Associated To ✕ +

Interfaces:	Domain Profile	Encapsulation
	<u>L2-IB-MGMT-In (L2)</u> ▼	

UPDATE
CANCEL

SUBMIT
CANCEL

18. Click UPDATE and SUBMIT to complete creating the Attachable Access Entity Profile.

19. Click APPLY CHANGES to complete creating the vPC. Click OK for the confirmation message.

20. On both switches at the top of the screen, select the port connected for the first NetApp Storage Controller.

21. At the bottom right, select CONFIGURE VPC.

22. For Policy Group Name, enter VPC-<node01>. Select the appropriate policies as shown in the screenshot.

CONFIGURING VPC

Legend: Port Channel (light blue), L3 (light blue), VPC (green), Conn. to Fex (orange), L2 Interface (red), Selected (blue)

a01-93180-1 (Node-101)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a01-93180-2 (Node-102)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

VPC

Policy Group Name: VPC-a01-aff8040-01

Description: optional

Link Level Policy: 10Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: default

LLDP Policy: LLDP-Disabled

STP Interface Policy: BPDU-Filter-Guard

Egress Data Plane Policing Policy: default

Ingress Data Plane Policing Policy: default

Port Channel Policy: LACP-Active

Storm Control Interface Policy: default

L2 Interface Policy: VLAN-Scope-Global

Attached Entity Profile: select an option

23. For the Attached Entity Policy drop-down, select Create Attachable Access Entity Profile.

24. Name the profile AEP-NTAP. Click the + sign to the right to add a Physical Domain.

25. In the drop-down, select Create Physical Domain.

26. In the Create Physical Domain window, name the Domain PD-NTAP.

27. In the VLAN Pool drop-down, select create VLAN Pool.

28. In the create VLAN Pool window, name the VLAN Pool VP-NTAP. Select Static Allocation.

29. Click the + sign to add an Encapsulation Block.

30. Enter the Storage Infrastructure NFS VLAN for both the From and To parts of the Range. Select Static Allocation.

Create Ranges
i ✕

Specify the Encap Block Range

Type: **VLAN**

Range: 3050 - 3050
From To

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

OK
CANCEL

31. Click OK to complete creating the range.

32. Repeat steps 28-30 to add encapsulation blocks for the Storage iSCSI VLANs (if iSCSI is being used in this FlexPod) and the Infrastructure SVM Management VLAN.

Create VLAN Pool
i ✕

Specify the Pool identity

Name: **VP-NTAP**

Description:

Allocation Mode: Dynamic Allocation Static Allocation

Encap Blocks: ✕ +

VLAN Range	Allocation Mode
[3050]	Static Allocation
[3010]	Static Allocation
[3020]	Static Allocation
[263]	Static Allocation


SUBMIT
CANCEL



33. Click SUBMIT to complete creating the VLAN Pool.

Create Physical Domain

Specify the domain name and the VLAN Pool

Name: PD-NTAP

VLAN Pool: VP-NTAP(static) 

Security Domains:  

Select	Name	Description

SUBMIT **CANCEL**

34. Click SUBMIT to complete creating the Physical Domain.



Create Attachable Access Entity Profile

Specify the name, domains and infrastructure encaps


Name: AEP-NTAP

Description: optional

Enable Infrastructure VLAN:

Domains (VMM, Physical or External) To Be Associated To  

Interfaces: Domain Profile Encapsulation

PD-NTAP (Physical)	
--------------------	--

UPDATE **CANCEL**

SUBMIT **CANCEL**

35. Click UPDATE and SUBMIT to complete creating the Attachable Access Entity Profile.

36. Click APPLY CHANGES to complete creating the vPC. Click OK for the confirmation message.

37. On both switches at the top of the screen, select the port connected for the second NetApp Storage Controller.
38. At the bottom right, select CONFIGURE VPC.
39. For Policy Group Name, enter VPC-<node02>. Select the appropriate policies and Attached Entity Profile as shown in the screenshot.

The screenshot displays the 'CONFIGURING VPC' interface. At the top left is a 'BACK TO SUMMARY' button. A legend at the top right identifies colors: Port Channel (light blue), L3 (light blue), VPC (green), Conn. to Fex (orange), L2 Interface (red), and Selected (dark blue). Two node grids are shown: 'a01-93180-1 (Node-101)' and 'a01-93180-2 (Node-102)'. Each grid has two rows of 16 ports. In Node-101, port 15 is green, 23 is green, and 16 is dark blue. In Node-102, port 15 is green, 23 is green, and 16 is dark blue. Below the grids is a 'VPC' section with a 'Policy Group Name' field containing 'VPC-a01-aff8040-02' and a 'Description' field containing 'optional'. A list of policies is shown with dropdown menus and copy icons: Link Level Policy (10Gbps-Auto), CDP Policy (CDP-Enabled), MCP Policy (default), LLDP Policy (LLDP-Disabled), STP Interface Policy (BPDU-Filter-Guard), Egress Data Plane Policing Policy (default), Ingress Data Plane Policing Policy (default), Port Channel Policy (LACP-Active), Storm Control Interface Policy (default), L2 Interface Policy (VLAN-Scope-Global), and Attached Entity Profile (AEP-NTAP).

40. Click APPLY CHANGES to complete creating the vPC. Click OK for the confirmation message.
41. On both switches at the top of the screen, select the port connected for the UCS Fabric Interconnect A.
42. At the bottom right, select CONFIGURE VPC.
43. For Policy Group Name, enter VPC-<ucs-fi-a>. Select the appropriate policies as shown in the screenshot.

BACK TO SUMMARY

CONFIGURING VPC

Port Channel
VPC
L2 Interface

L3
Conn. to Fex
Selected

a01-93180-1 (Node-101) ✕

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a01-93180-2 (Node-102) ✕

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

Policy Group Name: VPC-a01-6248-a

Description: optional

Link Level Policy: 10Gbps-Auto ✕

CDP Policy: CDP-Enabled ✕

MCP Policy: default ✕

LLDP Policy: LLDP-Enabled ✕

STP Interface Policy: BPDU-Filter-Guard ✕

Egress Data Plane Policing Policy: default ✕

Ingress Data Plane Policing Policy: default ✕

Port Channel Policy: LACP-Active ✕

Storm Control Interface Policy: default ✕

L2 Interface Policy: VLAN-Scope-Global ✕

Attached Entity Profile: select an option ✕

44. For the Attached Entity Policy drop-down, select Create Attachable Access Entity Profile.
45. Name the profile AEP-UCS. Click the + sign to the right to add a Physical Domain.
46. In the drop-down, select Create Physical Domain.
47. In the Create Physical Domain window, name the Domain PD-UCS.
48. In the VLAN Pool drop-down, select create VLAN Pool.
49. In the create VLAN Pool window, name the VLAN Pool VP-UCS. Select Static Allocation.
50. Click the + sign to add an Encapsulation Block.
51. Enter the UCS Infrastructure NFS VLAN for both the From and To parts of the Range. Select Static Allocation.

Create Ranges
i ✕

Specify the Encap Block Range

Type: **VLAN**

Range: 3150 - 3150
From To

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

OK
CANCEL

52. Click OK to complete creating the range.

53. Repeat steps 50-52 to add encapsulation blocks for the UCS iSCSI VLANs (if iSCSI is being used in this FlexPod), the UCS IB-MGMT/Core-Services VLAN, the vMotion VLAN, and the ACI Fabric System VLAN (if Cisco AVS is being used in this FlexPod).

Create VLAN Pool
i ✕

Specify the Pool identity

Name: **VP-UCS**

Description:

Allocation Mode: Dynamic Allocation Static Allocation

Encap Blocks: ✕ +

VLAN Range	Allocation Mode
[3100]	Static Allocation
[3110]	Static Allocation
[3120]	Static Allocation
[363]	Static Allocation
[3000]	Static Allocation
[1000]	Static Allocation

SUBMIT
CANCEL

54. Click SUBMIT to complete creating the VLAN Pool.

Create Physical Domain

Specify the domain name and the VLAN Pool

Name: PD-UCS

VLAN Pool: VP-UCS(static)

Security Domains:

Select	Name	Description
<input type="checkbox"/>		

SUBMIT **CANCEL**

55. Click SUBMIT to complete creating the Physical Domain.

Create Attachable Access Entity Profile

Specify the name, domains and infrastructure encaps

Name: AEP-UCS

Description: optional

Enable Infrastructure VLAN:

Domains (VMM, Physical or External) To Be Associated To

Interfaces:	Domain Profile	Encapsulation
	PD-UCS (Physical)	

UPDATE **CANCEL**

SUBMIT **CANCEL**

56. Click UPDATE and SUBMIT to complete creating the Attachable Access Entity Profile.

57. Click APPLY CHANGES to complete creating the vPC. Click OK for the confirmation message.

58. On both switches at the top of the screen, select the port connected to the second UCS Fabric Interconnect.
59. At the bottom right, select CONFIGURE VPC.
60. For Policy Group Name, enter VPC-<ucs-fi-b>. Select the appropriate policies and Attached Entity Profile as shown in the screenshot.

The screenshot displays the 'CONFIGURING VPC' interface. At the top, there is a 'BACK TO SUMMARY' button and a legend for port types: Port Channel (light blue), L3 (light blue), VPC (green), Conn. to Fex (orange), L2 Interface (red), and Selected (dark blue). Below the legend, two node configurations are shown: 'a01-93180-1 (Node-101)' and 'a01-93180-2 (Node-102)'. Each node has a grid of ports from 01 to 48. In Node-101, port 20 is selected (dark blue), and ports 15, 17, 19, and 23 are highlighted in green. In Node-102, port 20 is selected (dark blue), and ports 15, 17, 19, and 23 are highlighted in green. Below the node configurations, the 'VPC' configuration page is shown. It includes a 'Policy Group Name' field with the value 'VPC-a01-6248-b' and a 'Description' field with the value 'optional'. The configuration is divided into two columns of policy settings, each with a dropdown menu and a copy icon:

- Link Level Policy: 10Gbps-Auto
- CDP Policy: CDP-Enabled
- MCP Policy: default
- LLDP Policy: LLDP-Enabled
- STP Interface Policy: BPDU-Filter-Guard
- Egress Data Plane Policing Policy: default
- Ingress Data Plane Policing Policy: default
- Port Channel Policy: LACP-Active
- Storm Control Interface Policy: default
- L2 Interface Policy: VLAN-Scope-Global
- Attached Entity Profile: AEP-UCS

61. Click APPLY CHANGES to complete creating the vPC. Click OK for the confirmation message.

Create In-Band Management External Bridged Network and Core-Services EPG

This section details the steps for setup of the In-Band Management External Bridged Network in Tenant common. This setup will allow the In-Band Management network to be bridged into the ACI fabric.

1. In the APIC Advanced GUI, at the top select Tenants > common.
2. On the left, expand Tenant common and Networking.
3. Right-click External Bridged Networks and select Create Bridged Outside.
4. Name the Bridged Outside BO-IB-MGMT-In.

5. Select L2-IB-MGMT-In for the External Bridged Domain.
6. Use the drop-down next to Bridge Domain to select Create Bridge Domain.
7. Name the Bridge Domain BD-common-Internal.
8. Select the common/default VRF. Leave optimize selected for Forwarding. Select default for both the End Point Retention Policy and the IGMP Snoop Policy.

Create Bridge Domain


STEP 1 > Main

1. Main 2. L3 Configurations 3. Advanced/Troubleshooting


Specify Bridge Domain for the VRF


Name:

Description:

VRF: 

Forwarding:

End Point Retention Policy: 
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: 

9. At the bottom right, select NEXT.
10. No changes are needed for the L3 Configurations.
11. Select NEXT.
12. Select default for the Monitoring Policy.
13. Click FINISH.
14. Back in the Create Bridged Outside window, enter vlan-<external-IB-MGMT-VLAN> for Encap.
15. Select the VPC Path Type.
16. Using the Path drop-down, select VPC-IB-MGMT-In. Select Add. The VPC should appear in the list.

Create Bridged Outside

i X

STEP 1 > Identity 1. Identity 2. External EPG Networks

Configure the Bridged Outside

Name:

Description:

Tags:
enter tags separated by comma

External Bridged Domain: +

Bridge Domain: +

Encap:
e.g., vlan-1

Nodes And Interfaces Protocol Profiles

Path Type: Port PC VPC

Path: +

ADD

Node-101-102/VPC-IB-MGMT-In x [Clear All](#)

17. Click NEXT.

18. Select the + sign to add an External EPG Network.

19. Name the External Network EN-IB-MGMT and click OK.

20. Click FINISH.

21. On the left, expand Security Policies and select Contracts.

22. Right-click Contracts and select Create Contract.
23. Name the Contract common-Allow-IB-MGMT.
24. Select the Global Scope.
25. Click the + sign to add a Subject to the Contract.
26. Name the subject Allow-IB-MGMT.
27. Click the + under Filter Chain to add a Filter.
28. Click the drop-down, then click the + sign to add a Filter Identity.
29. Name the Filter Identity Allow-All.
30. Click the + sign to add an Entry to the Filter.
31. Name the Entry Allow-All and select the IP EtherType. Leave the IP Protocol set at Unspecified.
32. Click UPDATE.

Create Filter
i
✕

Specify the Filter Identity

Name:

Description:

Entries: ✕ +

Name	EtherType	ARP Flag	IP Protocol ▲	Match Only Fragmen	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
						From	To	From	To	
Allow-All	IP		unspecified	False	False					

33. Click SUBMIT to add the Filter.
34. In the Create Contract Subject window, click UPDATE to add the Filter Chain to the Contract Subject.

Create Contract Subject

Specify Identity Of Subject

Name: Allow-IB-MGMT

Description: optional

Target DSCP: unspecified

Apply Both Directions:

Reverse Filter Ports:

Filter Chain

Filters
Name
common/Allow-All

L4-L7 SERVICE GRAPH
Service Graph: select an option

PRIORITY
QoS:

OK CANCEL

35. Click OK to add the Contract Subject.



Note: The Contract Subject's Filter Chain can be modified later to make the contract more restrictive.

Create Contract

Specify Identity Of Contract

Name:

Scope:

QoS Class:

Target DSCP:

Description:

Subjects: × +

Name	Description
Allow-IB-MGMT	

36. Click SUBMIT to finish creating the Contract.

37. On the left, expand Networking, External Bridged Networks, BO-IB-MGMT-In, and Networks.

38. Select EN-IB-MGMT.

39. Click the + sign to the right of Provided Contracts to add a Provided Contract.

40. Under Name, select common/common-Allow-IB-MGMT.

41. Click UPDATE to add the Provided Contract.

The screenshot displays the Cisco ACI GUI for configuring an External Network Instance Profile. The main title is "External Network Instance Profile - EN-IB-MGMT". The configuration is under the "Policy" tab. The "Properties" section includes fields for Name (EN-IB-MGMT), Description (optional), Tags, Label, and QoS Class (Unspecified). The configuration status is "applied". Below this, there are three tables: "Provided Contracts", "Consumed Contracts", and "Taboo Contracts". The "Provided Contracts" table has one entry: "common-Allow-IB-..." with Tenant "common", Type "Contract", QoS Class "Unspecified", Match Type "AtleastOne", and State "formed". The other two tables are empty, with a message: "No items have been found. Select Actions to create a new item." At the bottom right, there are buttons for "SHOW USAGE", "SUBMIT", and "RESET".

42. Click SUBMIT to finish completion of the In-Band Management External Bridged Network.

43. Under Tenant common, right-click Application Profiles and select Create Application Profile.

44. Name the Tenant Application Profile `IB-MGMT` and click SUBMIT.

45. On the left expand Application Profiles and right-click `IB-MGMT`. Select Create Application EPG.

46. Name the EPG `Core-Services`.



Note: The `Core-Services` EPG will come from the VMware ESXi hosts in the Cisco UCS and can have both VMkernel ports and Virtual Machine interfaces in the In-Band Management subnet. These VMkernel ports and VMs will have access to the external network through the External Bridged Network just created. Tenant VMs will also be able to reach these VMkernel ports and VMs to receive `Core-Services` such as DNS or vCenter access.

47. Leave the Intra EPG Isolation set to Unenforced and select the `common/BD-common-Internal Bridge Domain`. Select the default Monitoring Policy.

Create Application EPG

STEP 1 > Identity

1. Identity

Specify the EPG Identity

Name: Core-Services

Description: optional

Tags:

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: common/BD-common-

Monitoring Policy: default

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

PREVIOUS FINISH CANCEL

48. Click FINISH.

49. On the left, expand IB-MGMT, Application EPGs, and EPG Core-Services.

50. Right-click Domains and select Add Physical Domain Association.

51. Select PD-UCS for the Physical Domain Profile. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: PD-UCS

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

SUBMIT CANCEL

52. Select SUBMIT to finish adding the Physical Domain Association.
53. On the left, under EPG Core-Services, right-click Static Bindings (Paths) and select Deploy Static EPG on PC, VPC, or Interface.
54. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.
55. Using the Path drop-down, select the VPC for UCS Fabric Interconnect A.
56. Enter `vlan-<ucs-IB-MGMT-VLAN>` for Port Encap.
57. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface i X

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ▾ 📄

Primary VLAN:

Port Encap:
 For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

58. Click SUBMIT to complete adding the Static Path Mapping.

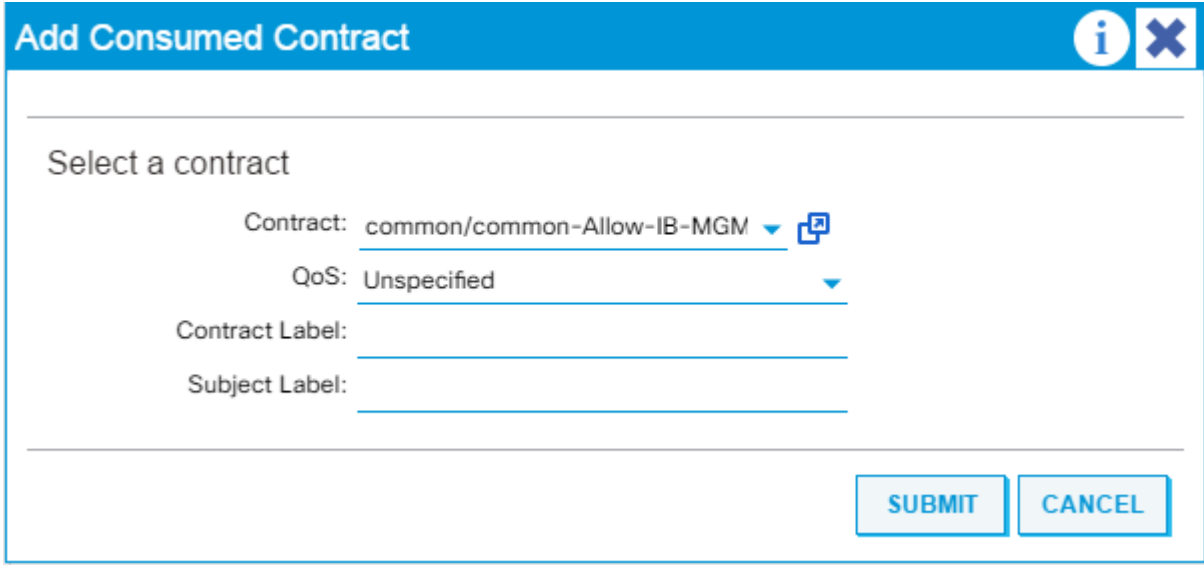
59. Repeat steps 53-58 to add the Static Path Mapping for UCS Fabric Interconnect B.

The screenshot shows the Cisco ACI GUI interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'VM Networking L4-L7 Services', 'Admin', and 'Operations'. The left sidebar shows a tree view with 'EPG Core-Services' expanded, and 'Static Bindings (Paths)' selected. The main content area displays a table of static bindings for 'Node-101-102'.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-363	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-363	Immediate	Trunk

60. On the left, under EPG Core-Services, right-click Contracts and select Add Consumed Contract.

61. In the Add Consumed Contract window, use the drop-down to select the common/common-Allow-IB-MGMT Contract.



Add Consumed Contract

Select a contract

Contract: common/common-Allow-IB-MGM

QoS: Unspecified

Contract Label: _____

Subject Label: _____

SUBMIT **CANCEL**

62. Click SUBMIT to add the Consumed Contract.



Note: The VMs that get assigned ports in the Core-Services EPG can now access the In-Band Management subnet mapped into the ACI Fabric.

63. Right-click Contracts and select Add Provided Contract.
64. In the Contract drop-down, select Create Contract.
65. Name the Contract common-Allow-Core-Services.
66. Select the Global Scope to allow the contract to be consumed in any tenant.
67. Click the + sign to add a Subject to the Contract.
68. Name the Subject Allow-Core-Services.
69. Under Filter Chain, click the + sign to add a Filter.
70. Select the drop-down, then select the Allow-All Filter from Tenant common.
71. Click UPDATE.

Create Contract Subject



Specify Identity Of Subject

Name: Allow-Core-ServicesDescription: optionalTarget DSCP: unspecifiedApply Both Directions: Reverse Filter Ports:

Filter Chain

Filters	
Name	
common/Allow-All	

L4-L7 SERVICE GRAPH
Service Graph: select an option

PRIORITY
QoS:

OK

CANCEL

72. Click OK to complete adding the Contract Subject.

Create Contract

Specify Identity Of Contract

Name:

Scope:

QoS Class:

Target DSCP:

Description:

Subjects:

Name	Description
Allow-Core-Services	

73. Click SUBMIT to complete creating the Contract.

74. Click SUBMIT to complete adding the Provided Contract.

75. Under EPG Core-Services, right-click Subnets and select Create EPG Subnet.

76. In the Create EPG Subnet window, add a gateway address with mask for Core-Services VMs to use to reach tenant VMs connected to Core-Services via Layer3. This gateway address resides in the ACI Fabric and should not be the gateway of the In-Band Management subnet.

77. For Scope, only select Shared between VRFs.

Create EPG Subnet i X

Specify the Subnet Identity

Default Gateway IP: 172.26.163.10/24
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 Querier IP

ND RA Prefix policy:

78. Click SUBMIT to complete creating the EPG Subnet.



Note: VMs assigned to ports in the Core-Services EPG can now be reached by contract and L3 from VMs in any tenant.

Create Security Filters in Tenant common

This section details the steps for creation of Security Filters for NFS v3 with NetApp Storage and for iSCSI. This section can also be used to set up other filters necessary to your environment. To create the security filters in tenant common, complete the following steps:

1. In the APIC Advanced GUI, at the top select Tenants > common.
2. On the left, expand Tenant common, Security Policies, and Filters.
3. Right-click Filters and select Create Filter.
4. Name the filter NTAP-NFS-v3.
5. Click the + sign to add an Entry to the Filter.
6. Name the Entry tcp-111 and select EtherType IP.

7. Select the TCP IP Protocol and enter 111 for From and To under the Destination Port / Range by backspacing over Unspecified and entering the number.
8. Click UPDATE to add the Entry.
9. Click the + sign to add another Entry to the Filter.
10. Name the Entry tcp-635 and select EtherType IP.
11. Select the TCP IP Protocol and enter 635 for From and To under the Destination Port / Range by backspacing over Unspecified and entering the number.
12. Click UPDATE to add the Entry.
13. Click the + sign to add the final Entry to the Filter.
14. Name the Entry tcp-2049 and select EtherType IP.
15. Select the TCP IP Protocol and enter 2049 for From and To under the Destination Port / Range by backspacing over Unspecified and entering the number.
16. Click UPDATE to add the Entry.

Create Filter
i
✕

Specify the Filter Identity

Name:

Description:

Entries: ✕ +

Name	EtherType	ARP Flag	IP Protocol	Match Only Fragmen	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
						From	To	From	To	
tcp-111	IP		tcp	False	False	unspecified	unspecified	111	111	Unspecified
tcp-635	IP		tcp	False	False	unspecified	unspecified	635	635	Unspecified
tcp-2049	IP		tcp	False	False	unspecified	unspecified	2049	2049	Unspecified

17. Click SUBMIT to complete adding the Filter.
18. Right-click Filters and select Create Filter.
19. Name the filter iSCSI.
20. Click the + sign to add an Entry to the Filter.
21. Name the Entry iSCSI and select EtherType IP.

22. Select the TCP IP Protocol and enter 3260 for From and To under the Destination Port / Range by backspacing over Unspecified and entering the number.
23. Click UPDATE to add the Entry.

i X
Create Filter

Specify the Filter Identity

Name:

Description:

Entries: x +

Name	EtherType	ARP Flag	IP Protocol	Match Only	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
						From	To	From	To	
iSCSI	IP		tcp	False	False	unspecified	unspecified	3260	3260	Unspecified

24. Click SUBMIT to complete adding the Filter.



Note: By adding these Filters to Tenant common, they can be used from within any Tenant in the ACI Fabric.

CISCO
welcome, admin

System Tenants Fabric VM Networking L4-L7 Services Admin Operations

ALL TENANTS | Add Tenant | Search: | common | mgmt | infra

Tenant common

- Quick Start
- Tenant common
 - Application Profiles
 - Networking
 - L4-L7 Service Parameters
 - Security Policies
 - Contracts
 - Taboo Contracts
 - Imported Contracts
 - Out-Of-Band Contracts
 - Filters
 - Allow-All
 - NTAP-NFS-v3
 - arp
 - default
 - est
 - iSCSI
 - icmp
 - Troubleshoot Policies
 - Monitoring Policies
 - L4-L7 Services

Security Policies - Filters

Name	Entries	Description
Allow-All	Allow-All (IP)	
arp	arp (ARP)	
default	default	
est	est (tcp, Rule: Established)	
icmp	icmp (icmp)	
iSCSI	iSCSI (tcp, Destination: 3260)	
NTAP-NFS-v3	tcp-111 (tcp, Destination: 111) tcp-2049 (tcp, Destination: 2049) tcp-635 (tcp, Destination: 635)	

Deploy Infrastructure (Foundation) Tenant

This section details the steps for creation of the Foundation Tenant in the ACI Fabric. This tenant will host infrastructure connectivity between the compute (VMware on Cisco UCS) and the storage (NetApp) environments. A corresponding Infra-SVM has already been created on the NetApp storage to align with this tenant. To deploy the Foundation Tenant, complete the following steps:

1. In the APIC Advanced GUI, at the top select Tenants > Add Tenant.
2. Name the Tenant `Foundation`. Select the default Monitoring Policy.
3. For the VRF Name, also enter `Foundation`. Leave the Take me to this tenant when I click finish checkbox checked.

Create Tenant
i X

Specify tenant details

Name:

Description:

Tags:
enter tags separated by comma

Monitoring Policy:

Security Domains:

Select	Name	Description

VRF Name:

Take me to this tenant when I click finish

4. Click SUBMIT to finish creating the Tenant.
5. If you are using iSCSI Boot or providing iSCSI LUN access in the FlexPod Infrastructure, complete steps 6-39. Otherwise, continue to step 40.
6. On the left under Tenant Foundation, right-click Application Profiles and select Create Application Profile.

7. Name the Application Profile iSCSI, select the default Monitoring Policy, and click SUBMIT to complete adding the Application Profile.
8. On the left, expand Application Profiles and iSCSI.
9. Right-click Application EPGs and select Create Application EPG.
10. Name the EPG iSCSI-A. Leave Intra EPG Isolation Unenforced.
11. Use the Bridge Domain drop-down to select Create Bridge Domain.
12. Name the Bridge Domain BD-iSCSI-A.
13. Select the Foundation/Foundation VRF.
14. Use the Forwarding drop-down to select Custom.
15. Select Flood for the L2 Unknown Unicast and default for the End Point Retention Policy and IGMP Snoop Policy.

Create Bridge Domain


STEP 1 > Main


1. Main 2. L3 Configurations 3. Advanced/Troubleshooting


Specify Bridge Domain for the VRF


Name:


Description:


VRF: 


Forwarding: 

L2 Unknown Unicast: 

L3 Unknown Multicast Flooding: 

Multi Destination Flooding: 

End Point Retention Policy: 
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: 

16. At the bottom right, click NEXT.
17. Make sure Unicast Routing is Enabled and click NEXT.
18. Select the default Monitoring Policy and click FINISH.

i X
Create Application EPG

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation:

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

19. Select the default Monitoring Policy and click FINISH to complete creating the EPG.
20. On the left, expand Application EPGs and EPG iSCSI-A. Right-click Domains and select Add Physical Domain Association.
21. Using the drop-down, select the PD-NTAP Physical Domain Profile.
22. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

i X
Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile:

Deploy Immediacy:

Resolution Immediacy:

23. Click SUBMIT to complete the Physical Domain Association.

24. Repeat steps 20-23 to add the PD-UCS Physical Domain Association.



Note: In this deployment for iSCSI, we are adding both the NetApp LIF endpoints and the VMware VMkernel (VMK) endpoints in a single EPG. This method allows unrestricted communication within the EPG. We also had the choice to put the LIFs in one EPG and the VMKs in a second EPG and connect them with a filtered contract. We will deploy NFS that way next.

Domain Profile	Domain Type	Deployment Immediacy	Resolution Immediacy	State	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Allow Micro-Segmentation
PD-NTAP	Physical Domain	Immediate	Immediate	formed			False
PD-UCS	Physical Domain	Immediate	Immediate	formed			False

25. Right-click Static-Bindings (Paths) and select Deploy EPG on PC, VPC, or Interface.

26. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.

27. Using the Path drop-down, select the VPC for NetApp Storage Controller 01.

28. Enter `vlan-<storage-iscsi-a-vlan>` for Port Encap.

29. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface
i X

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ▼ 📄

Primary VLAN: For example, vlan-1

Port Encap: For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

30. Click SUBMIT to complete adding the Static Path Mapping.
31. Repeat steps 25-30 to add the Static Path Mapping for NetApp Storage Controller 02.
32. Right-click Static-Bindings (Paths) and select Deploy EPG on PC, VPC, or Interface.
33. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.
34. Using the Path drop-down, select the VPC for UCS Fabric Interconnect A.
35. Enter `vlan-<ucs-iscsi-A-VLAN>` for Port Encap.
36. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Primary VLAN:
For example, vlan-1

Port Encap:
For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

37. Click SUBMIT to complete adding the Static Path Mapping.

38. Repeat steps 32-37 to add the Static Path Mapping for UCS Fabric Interconnect B.

The screenshot shows the Cisco ACI GUI interface. The top navigation bar includes 'System', 'Tenants', 'Fabric', 'VM Networking L4-L7 Services', 'Admin', and 'Operations'. The left sidebar shows the 'Tenant Foundation' tree with 'Static Bindings (Paths)' selected. The main content area displays a table of static bindings for 'Node-101-102'.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-3110	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-3110	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-01		vlan-3010	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-3010	Immediate	Trunk

39. Repeat steps 9-38 to build the iSCSI-B EPG. Make sure to create a separate Bridge Domain for this EPG and use the iSCSI-B VLAN IDs.

The screenshot shows the Cisco ACI GUI with the following components:

- Navigation Bar:** System, Tenants, Fabric, VM Networking L4-L7 Services, Admin, Operations. User: welcome_admin.
- Search Bar:** ALL TENANTS | Add Tenant | Search: enter name, descr | common | Foundation | infra | mgmt
- Left Sidebar (Tenant Foundation):**
 - Quick Start
 - Tenant Foundation
 - Application Profiles
 - ISCSI
 - Application EPGs
 - EPG ISCSI-A
 - EPG ISCSI-B
 - Domains (VMs and Bare-Meta..)
 - Static Bindings (Paths)
 - Static Bindings (Leaves)
 - Contracts
 - Static EndPoint
 - Subnets
 - L4-L7 Virtual IPs

- Main Content Area (Static Bindings (Paths)):**
- Buttons: Refresh, Download, Add, Print
- Table:

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-3120	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-3120	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-01		vlan-3020	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-3020	Immediate	Trunk
- ACTIONS -

40. On the left, under Tenant Foundation, right-click Application Profiles and select Create Application Profile.
41. Name the Profile `NFS`, select the default Monitoring Policy, and click SUBMIT.
42. Right-click the NFS Application Profile and select Create Application EPG.
43. Name the EPG `NFS-LIF` and leave Intra EPG Isolation set at Unenforced.
44. Use the Bridge Domain drop-down to select Create Bridge Domain.
45. Name the Bridge Domain `BD-NFS` and select the Foundation/Foundation VRF.



Note: It is important to create a new Bridge Domain for each traffic VLAN coming from the NetApp Storage Controllers. All of the VLAN interfaces on a given NetApp Interface Group share the same MAC address, and separating to different bridge domains in the ACI Fabric allows all the traffic to be forwarded properly.

46. For Forwarding, select Custom and select Flood for L2 Unknown Unicast. Select default for the End Point Retention Policy and the IGMP Snoop Policy.



Create Bridge Domain

STEP 1 > Main

1. Main

2. L3 Configurations

3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name: Description: VRF: Forwarding: L2 Unknown Unicast: L3 Unknown Multicast Flooding: Multi Destination Flooding: End Point Retention Policy: This policy only applies to local L2 L3 and remote L3 entriesIGMP Snoop Policy:

47. At the bottom right, click NEXT.

48. Make sure Unicast Routing is enabled and click NEXT.

49. Select the default Monitoring Policy and click FINISH.

Create Application EPG

STEP 1 > Identity

1. Identity

Specify the EPG Identity

Name: NFS-LIF

Description: optional

Tags: enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: Foundation/BD-NFS

Monitoring Policy: default

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

PREVIOUS FINISH CANCEL

50. Select the default Monitoring Policy and click FINISH to complete creating the EPG.
51. On the left expand NFS, Application EPGs, and EPG NFS-LIF.
52. Right-click Domains and select Add Physical Domain Association.
53. Select the PD-NTAP Physical Domain Profile.
54. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: PD-NTAP

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

SUBMIT CANCEL

55. Click SUBMIT to complete adding the Physical Domain Association.
56. Right-click Static Bindings (Paths) and select Deploy Static EPG on PC, VPC, or Interface.
57. Select the Virtual Port Channel Path Type.
58. Using the Path drop-down, select the VPC for NetApp Storage Controller 01.
59. For Port Encap, enter `vlan-<storage-Infra-NFS-VLAN>`.
60. Select Immediate for Deployment Immediacy and Trunk for Mode.

i X

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ▾ 📄

Primary VLAN:

For example, vlan-1

Port Encap:

For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

61. Click SUBMIT to finish adding the EPG Static Binding.

62. Repeat steps 56-61 for the Static Path to NetApp Storage Controller 02.

i
welcome, admin ▾

System Tenants Fabric VM Networking L4-L7 Services Admin Operations

ALL TENANTS | Add Tenant | Search: | common | Foundation | mgmt | infra

Tenant Foundation

- Quick Start
- Tenant Foundation
 - Application Profiles
 - NFS
 - Application EPGs
 - EPG NFS-LIF
 - Domains (VMs and Bare-Meta..
 - Static Bindings (Paths)
 - Static Bindings (Leaves)
 - Contracts
 - Static EndPoint
 - Subnets
 - L4-L7 Virtual IPs

Static Bindings (Paths)

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-aff8040-01		vlan-3050	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-3050	Immediate	Trunk

63. On the left under EPG NFS-LIF, right-click Contracts and select Add Provided Contract.

64. In the Add Provided Contract window, use the Contract drop-down to select Create Contract.

65. Name the contract Allow-NFS. Leave the Scope set at VRF.

66. Click the + sign to add a Contract Subject.
67. Name the subject Allow-NFS.
68. Click the + sign to add a Filter to the Filter Chain.
69. Click the drop-down and select NTAP-NFS-v3 from Tenant common.
70. Click UPDATE.

Create Contract Subject

Specify Identity Of Subject

Name:

Description:

Target DSCP:

Apply Both Directions:

Reverse Filter Ports:

Filter Chain

Filters
Name
common/NTAP-NFS-v3

L4-L7 SERVICE GRAPH
Service Graph:

PRIORITY
QoS:



Note: Optionally, add ICMP to the filter chain to allow ping in this contract for troubleshooting purposes.

71. Click OK to complete the Contract Subject.

Create Contract

Specify Identity Of Contract

Name:

Scope:

QoS Class:

Target DSCP:

Description:

Subjects: × +

Name	Description
Allow-NFS	

72. Click SUBMIT to complete creating the Contract.

Add Provided Contract

Select a contract

Contract:

QoS:

Contract Label:

Subject Label:

73. Click SUBMIT to complete Adding the Provided Contract.
74. Right-click Application EPGs under the NFS Application Profile and select Create Application EPG.
75. Name the EPG NFS-VMK and leave Intra EPG Isolation set at Unenforced.
76. Use the Bridge Domain drop-down to select Foundation/BD-NFS. Select the default Monitoring Policy.

Create Application EPG

STEP 1 > Identity 1. Identity

Specify the EPG Identity

Name: NFS-VMK

Description: optional

Tags: enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: Foundation/BD-NFS

Monitoring Policy: default

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

77. Click FINISH to complete creating the EPG.
78. On the left expand NFS, Application EPGs, and EPG NFS-VMK.
79. Under EPG NFS-VMK, right-click Domains and select Add Physical Domain Association.
80. Select the PD-UCS Physical Domain Profile.
81. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: PD-UCS

Deploy Immediacy: **Immediate** On Demand

Resolution Immediacy: **Immediate** On Demand Pre-provision

SUBMIT **CANCEL**

82. Click SUBMIT to complete adding the Physical Domain Association.

83. Under EPG NFS-VMK, right-click Static Bindings (Paths) and select Deploy Static EPG on PC, VPC, or Interface.

84. Select the Virtual Port Channel Path Type.

85. Using the Path drop-down, select the VPC for UCS Fabric Interconnect A.

86. For Port Encap, enter `vlan-<ucs-Infra-NFS-VLAN>`.

87. Select Immediate for Deployment Immediacy and Trunk for Mode.

i X

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ↕

Primary VLAN: For example, vlan-1

Port Encap: For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

88. Click SUBMIT to finish adding the EPG Static Binding.

89. Repeat steps 83-88 for the Static Path to UCS Fabric Interconnect B.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-3150	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-3150	Immediate	Trunk

90. On the left under EPG NFS-VMK, right-click Contracts and select Add Consumed Contract.

91. In the Add Consumed Contract window, use the Contract drop-down to select Foundation/Allow-NFS.

Add Consumed Contract i X

Select a contract

Contract: Foundation/Allow-NFS ▼

QoS: Unspecified ▼

Contract Label: _____

Subject Label: _____

SUBMIT
CANCEL

92. Click SUBMIT to complete adding the Consumed Contract.



Note: We have now put the mapping in place to map the Infrastructure NFS VMK ports to the Infrastructure NFS LIFs on the NetApp Storage with a filter in place that only allows NFS V3 to pass.

93. On the left, under Tenant Foundation, right-click Application Profiles and select Create Application Profile.

94. Name the Profile `IB-MGMT`, set the Monitoring Policy to default and click SUBMIT.

95. Right-click the IB-MGMT Application Profile and select Create Application EPG.

96. Name the EPG `Infra-SVM-MGMT` and leave Intra EPG Isolation set at Unenforced.

97. Use the Bridge Domain drop-down to select `common/BD-common-Internal`. Select the default Monitoring policy.



Note: We used the `BD-common-Internal` Bridge Domain that we used earlier for both the L2 Bridged External IB-MGMT Network and the Core-Services EPG. The reason we are using this bridge domain is that we will be connecting all three of these entities at L2 and will be using the same subnet for all three, meaning they all have to be in the same bridge domain. If we were using a contract and EPG subnets to connect to EPGs at L3, such as when we connect Tenant EPGs to the Core-Services EPG, they do not have to be in the same bridge domain.

i X
Create Application EPG

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation: Enforced Unenforced

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

98. Click FINISH to complete creating the EPG.
99. On the left expand IB-MGMT, Application EPGs, and EPG Infra-SVM-MGMT.
100. Right-click Domains and select Add Physical Domain Association.
101. Select the PD-NTAP Physical Domain Profile.
102. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

i X
Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile:

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

103. Click SUBMIT to complete adding the Physical Domain Association.
104. Right-click Static Bindings (Paths) and select Deploy Static EPG on PC, VPC, or Interface.
105. Select the Virtual Port Channel Path Type.
106. Using the Path drop-down, select the VPC for NetApp Storage Controller 01.
107. For Port Encap, enter `vlan-<storage-IB-MGMT-VLAN>`.
108. Select Immediate for Deployment Immediacy and Trunk for Mode.

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

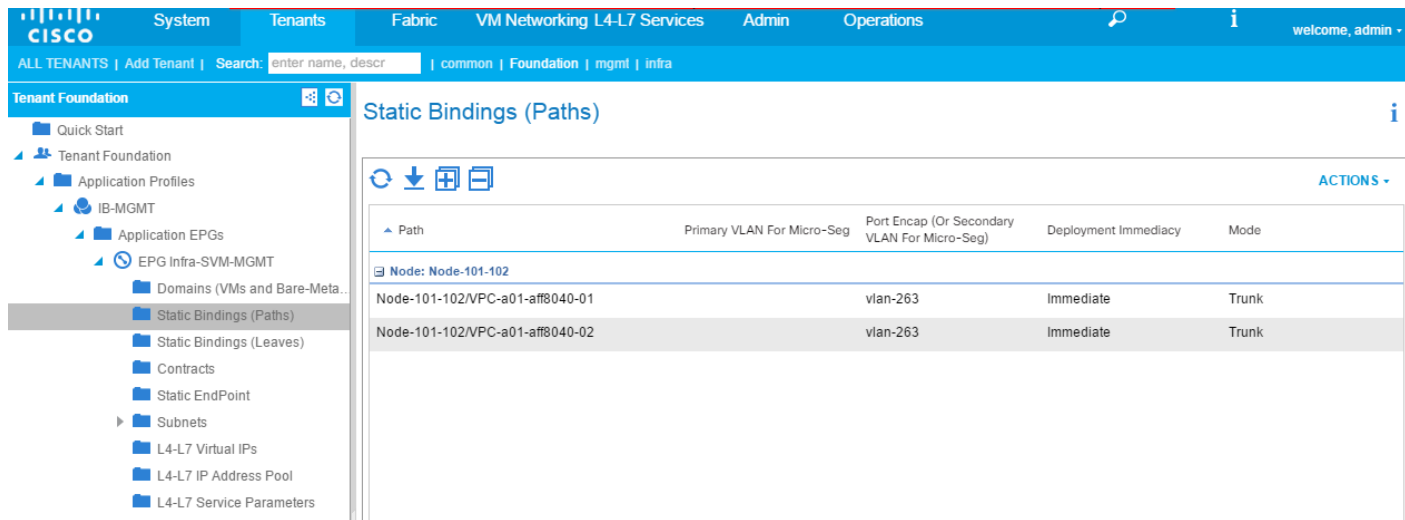
Primary VLAN:
For example, vlan-1

Port Encap:
For example, vlan-1

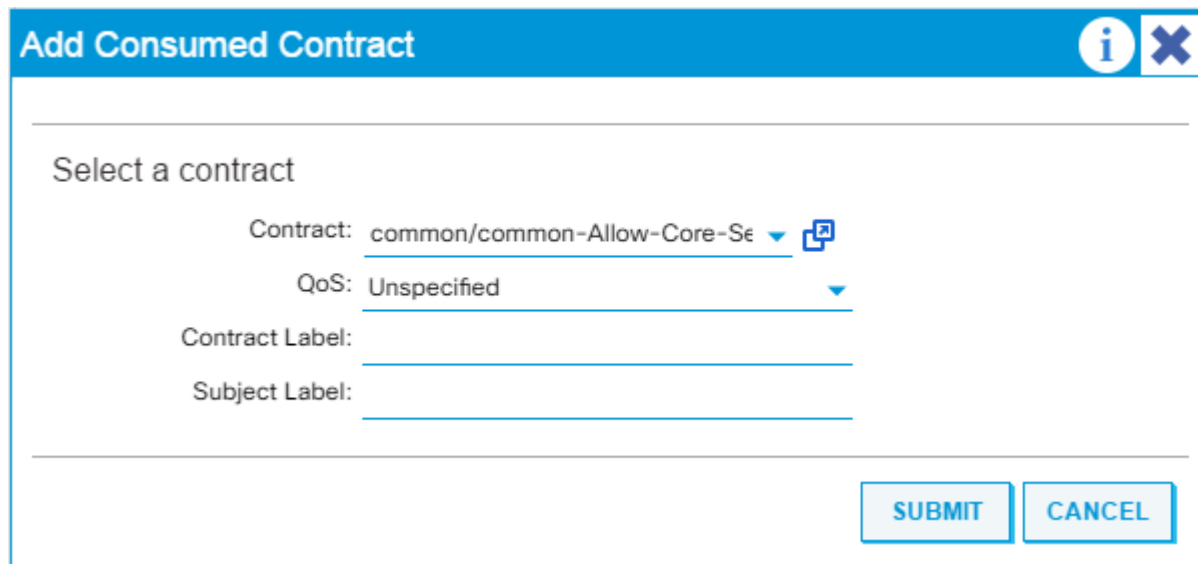
Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

109. Click SUBMIT to finish adding the EPG Static Binding.
110. Repeat steps 104-109 for the Static Path to NetApp Storage Controller 02.



111. On the left under EPG Infra-SVM-MGMT, right-click Contracts and select Add Consumed Contract.
112. In the Add Consumed Contract window, use the Contract drop-down to select common/common-Allow-Core-Services.



113. Click SUBMIT to complete adding the Consumed Contract.
114. On the left under EPG Infra-SVM-MGMT, right-click Contracts and select Add Consumed Contract.
115. In the Add Consumed Contract window, use the Contract drop-down to select common/common-Allow-IB-MGMT.
116. Click SUBMIT to complete adding the Consumed Contract.



Note: Now the mapping is in place to link the NetApp Storage Infra-SVM's management interface both to Core-Services and to the mapped-in IB-MGMT subnet. Since the storage is already setup, the Infra-SVM Management Interface should be reachable from the IB-MGMT subnet.

117. On the left, right-click Application Profiles and select Create Application Profile.
118. Name the Application Profile `vMotion`, select the default Monitoring Policy, and click SUBMIT to complete adding the Application Profile.
119. On the left, expand Application Profiles and `vMotion`.
120. Right-click Application EPGs and select Create Application EPG.
121. Name the EPG `vMotion`. Leave Intra EPG Isolation Unenforced.
122. Use the Bridge Domain drop-down to select Create Bridge Domain.
123. Name the Bridge Domain `BD-Internal`.
124. Select the Foundation/Foundation VRF.
125. For Forwarding, select Custom and select Flood for L2 Unknown Unicast. Select default for the End Point Retention Policy and the IGMP Snoop Policy.
126. At the bottom right, click NEXT.
127. Make sure Unicast Routing is Enabled and click NEXT.
128. Select the default Monitoring Policy and click FINISH.

i X
Create Application EPG

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation:

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

129. Select the default Monitoring Policy and click FINISH to complete creating the EPG.
130. On the left, expand Application EPGs and EPG vMotion. Right-click Domains and select Add Physical Domain Association.
131. Using the drop-down, select the PD-UCS Physical Domain Profile.
132. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.
133. Click SUBMIT to complete the Physical Domain Association.
134. Right-click Static-Bindings (Paths) and select Deploy EPG on PC, VPC, or Interface.
135. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.
136. Using the Path drop-down, select the VPC for UCS Fabric Interconnect A.
137. Enter `vlan-<ucs-vMotion-VLAN>` for Port Encap.
138. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Primary VLAN:
For example, vlan-1

Port Encap:
For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

139. Click SUBMIT to complete adding the Static Path Mapping.

140. Repeat steps 134-139 to add the Static Path Mapping for UCS Fabric Interconnect B.

The screenshot shows the Cisco ACI GUI interface. The top navigation bar includes tabs for System, Tenants, Fabric, VM Networking L4-L7 Services, Admin, and Operations. The left sidebar shows the Tenant Foundation tree with 'Static Bindings (Paths)' selected. The main content area displays a table of static bindings for Node-101-102.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-3000	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-3000	Immediate	Trunk

VMware vSphere 6.0 U1b Setup

VMware ESXi 6.0 U1b

This section provides detailed instructions for installing VMware ESXi 6.0 U1b in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

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

This procedure assumes that a vMedia policy has been created and linked to the ESXi Installation ISO, and has been attached to the Cisco UCS Server Service Profile.

Log in to Cisco UCS 6200 Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Under HTML, click the Launch UCS Manager link.
3. When prompted, enter `admin` as the user name and enter the administrative password.
4. To log in to Cisco UCS Manager, click Login.
5. From the main menu, click the Servers tab.
6. Select Servers > Service Profiles > root > `VM-Host-Infra-01`.
7. Right-click `VM-Host-Infra-01` and select KVM Console.
8. If prompted to accept an Unencrypted KVM session, accept as necessary.
9. Select Servers > Service Profiles > root > `VM-Host-Infra-02`.
10. Right-click `VM-Host-Infra-02` and select KVM Console.
11. If prompted to accept an Unencrypted KVM session, accept as necessary.
12. Boot each server by selecting Boot Server and clicking OK. Click OK again.

Install ESXi

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

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

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
4. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, press Enter to reboot the server.
9. After reboot, in UCS Manager, select each Service Profile and select Bind to a Template under the General Tab. Bind the Service Profile to the template without the vMedia Policy to stop the mounting of the ESXi ISO to the server.

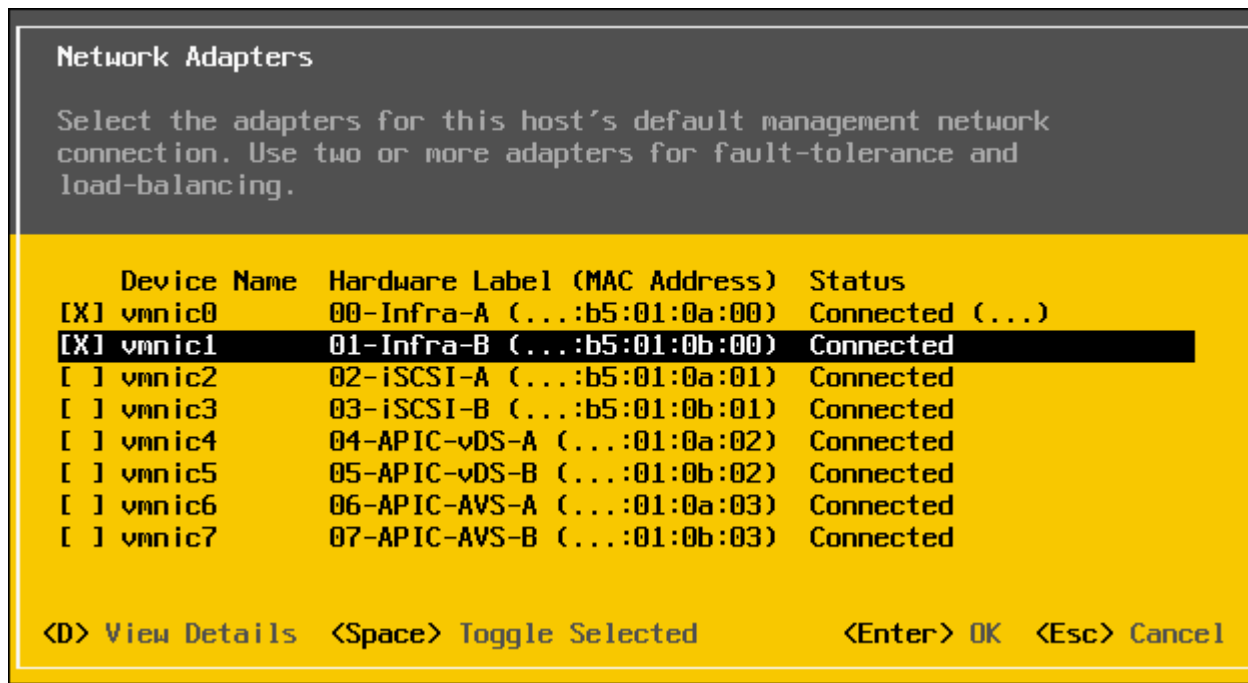
Set Up Management Networking for ESXi Hosts

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

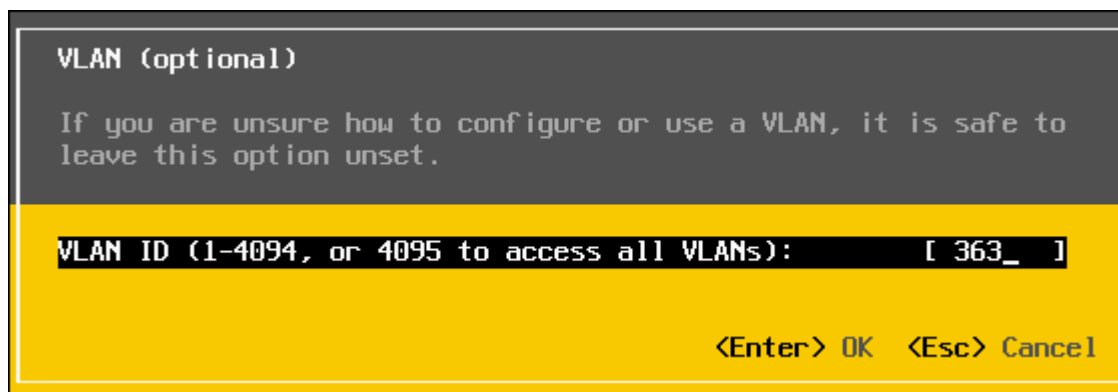
ESXi Host VM-Host-Infra-01

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

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select the Configure Management Network option and press Enter.
4. Select Network Adapters and press Enter.
5. Ensure that the vmnic mapping aligns with the UCS vNIC mapping.
6. Highlight `vmnic1` and select by pressing the Space Bar.



7. Press Enter to save and exit the Network Adapters window.
8. Select the VLAN (Optional) option and press Enter.
9. Enter the <ucs-ib-mgmt-vlan> and press Enter.



10. From the Configure Management Network menu, select IPv4 Configuration and press Enter.
11. Select the Set Static IP Address and Network Configuration option by using the Space Bar.
12. Enter the IP address for managing the first ESXi host: <vm-host-infra-01-mgmt-ip>.
13. Enter the subnet mask for the first ESXi host.
14. Enter the default gateway for the first ESXi host.
15. Press Enter to accept the changes to the IP configuration.

16. Select the IPv6 Configuration option and press Enter.
17. Using the Space Bar, select Disable IPv6 (restart required) and press Enter.
18. Select the DNS Configuration option and press Enter.



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

19. Enter the IP address of the primary DNS server.
20. Optional: Enter the IP address of the secondary DNS server.
21. Enter the fully qualified domain name (FQDN) for the first ESXi host.
22. Press Enter to accept the changes to the DNS configuration.
23. Press Esc to exit the Configure Management Network submenu.
24. Press Y to confirm the changes and reboot the host.
25. The ESXi host reboots. After reboot, press F2 and log back in as root.
26. Select Test Management Network to verify that the management network is set up correctly and press Enter.
27. Press Enter to run the test.
28. Press Enter to exit the window.
29. Press Esc to log out of the VMware console.
30. Repeat this procedure for ESXi host VM-Host-Infra-02.



Note: The ESXi host's management port maps into the Core-Services EPG in the ACI fabric. The Management Network test verifies that Core-Services and the mapping of the In-Band Management network into the ACI fabric are correctly set up and that the contract between these end points is correct.

Download VMware vSphere Client

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

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



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

Log in to VMware ESXi Hosts by Using VMware vSphere Client

ESXi Host VM-Host-Infra-01

To log in to the `VM-Host-Infra-01` ESXi host by using the VMware vSphere Client, complete the following steps:

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

ESXi Host VM-Host-Infra-02

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

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-Infra-02` as the host you are trying to connect to: `<vm-host-infra-02-mgmt-ip>`.
2. Enter `root` for the user name.
3. Enter the root password.

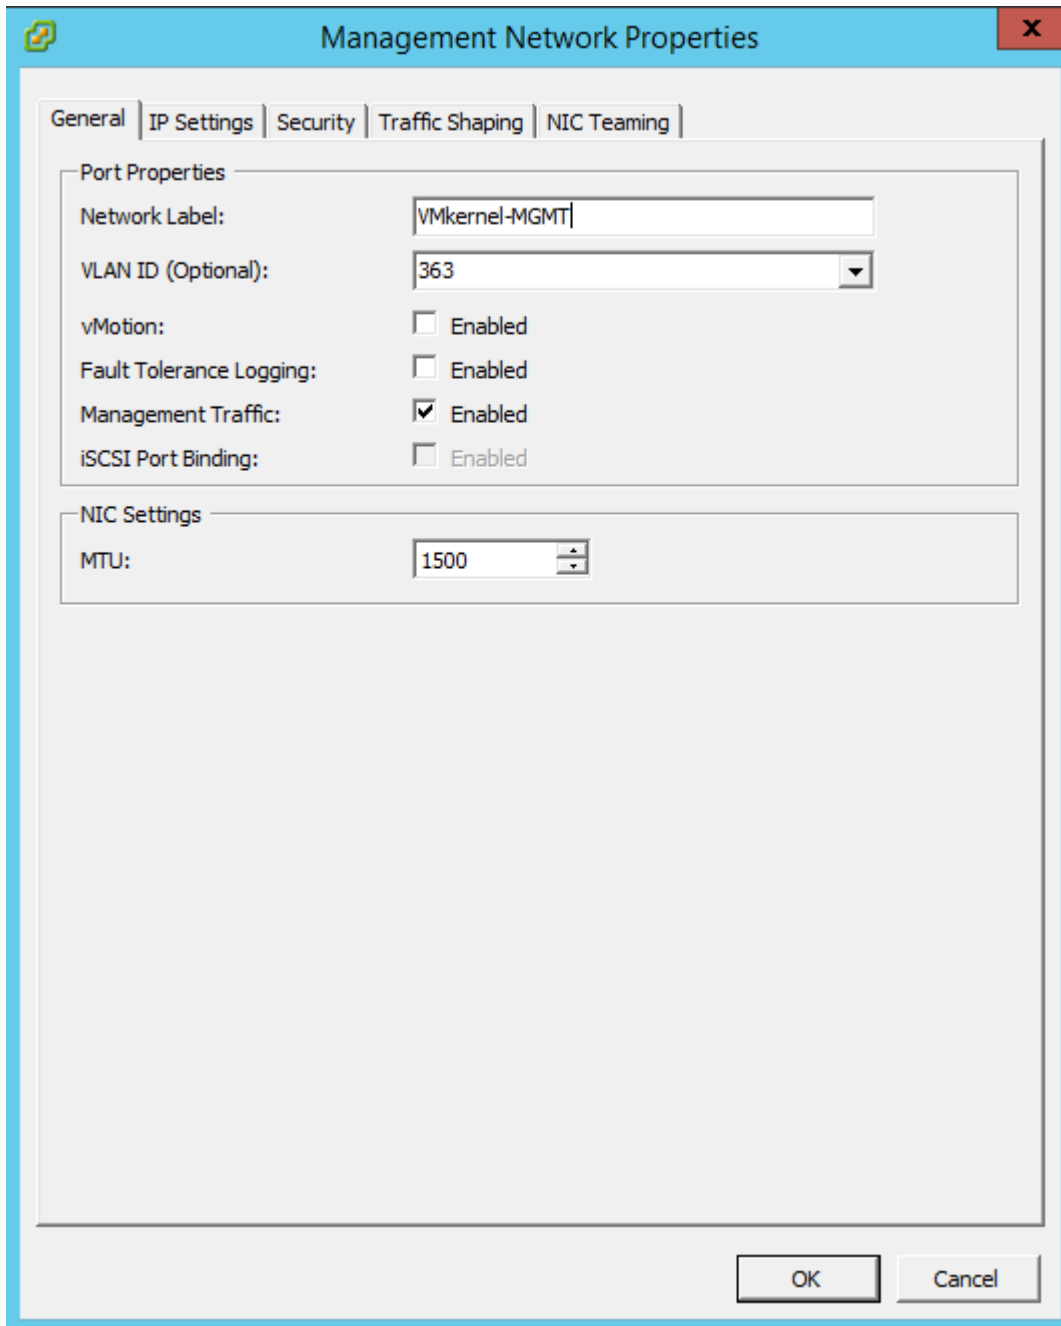
Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

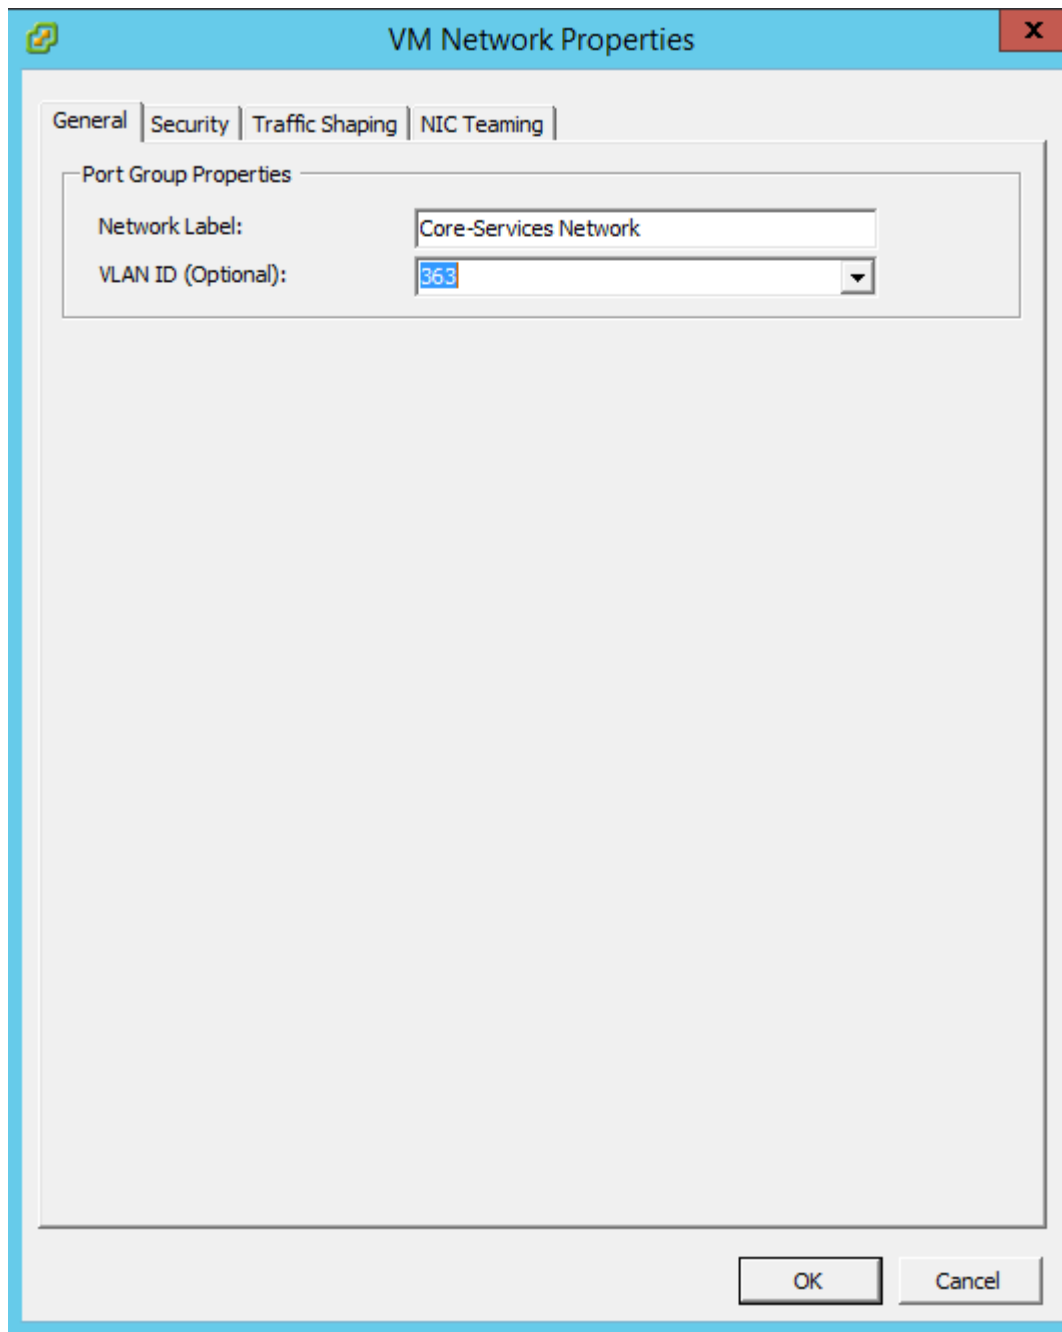
To set up the VMkernel ports and the virtual switches on the `VM-Host-Infra-01` ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of `vSwitch0`, click Properties.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to 9000.
7. Select the NIC Teaming tab.
8. Move `vmnic1` from Standby Adapters to Active Adapters, leaving both `vmnic0` and `vmnic1` as Active Adapters.
9. Click OK.

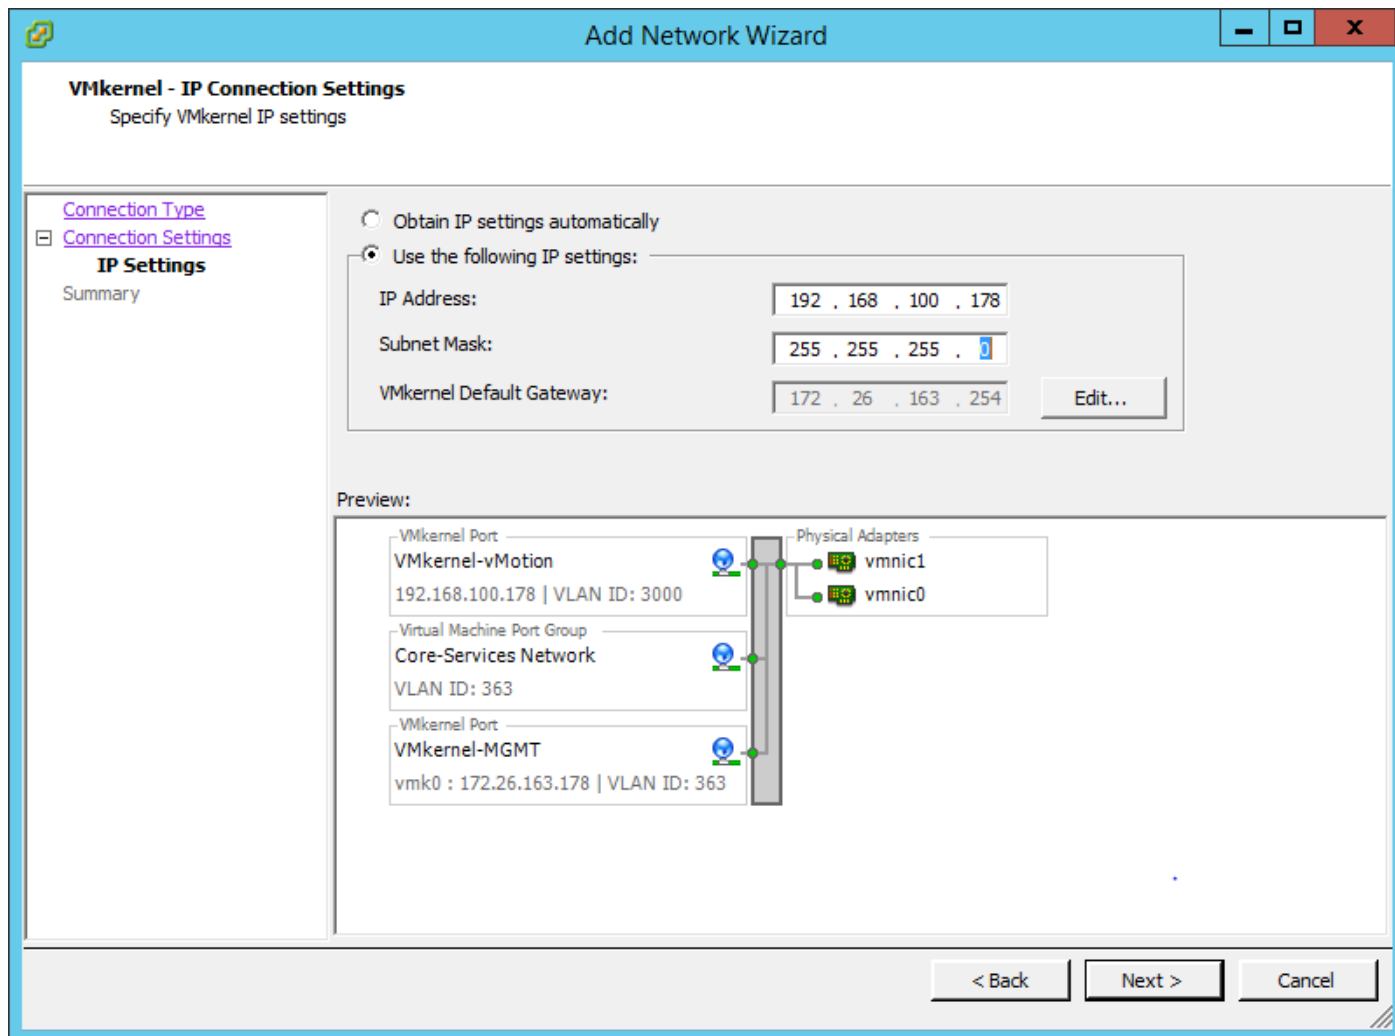
10. Select the Management Network configuration and click Edit.
11. Change the network label to `VMkernel-MGMT` and make sure the Management Traffic checkbox is checked.



12. Click OK to finalize the edits for VMkernel-MGMT.
13. Select the VM Network configuration and click Edit.
14. Change the network label to `Core-Services Network` and enter `<ucs-ib-mgmt-vlan>` in the VLAN ID (Optional) field.



15. Click OK to finalize the edits for Core-Services Network.
16. Click Add to add a VMkernel port.
17. Select VMkernel and click Next.
18. Label the network `vmkernel1-vMotion` and enter the `<ucs-vMotion-VLAN>` VLAN ID.
19. Select the Use this port group for vMotion checkbox and click Next.
20. Enter the ESXi host's vMotion IP address and Subnet Mask.



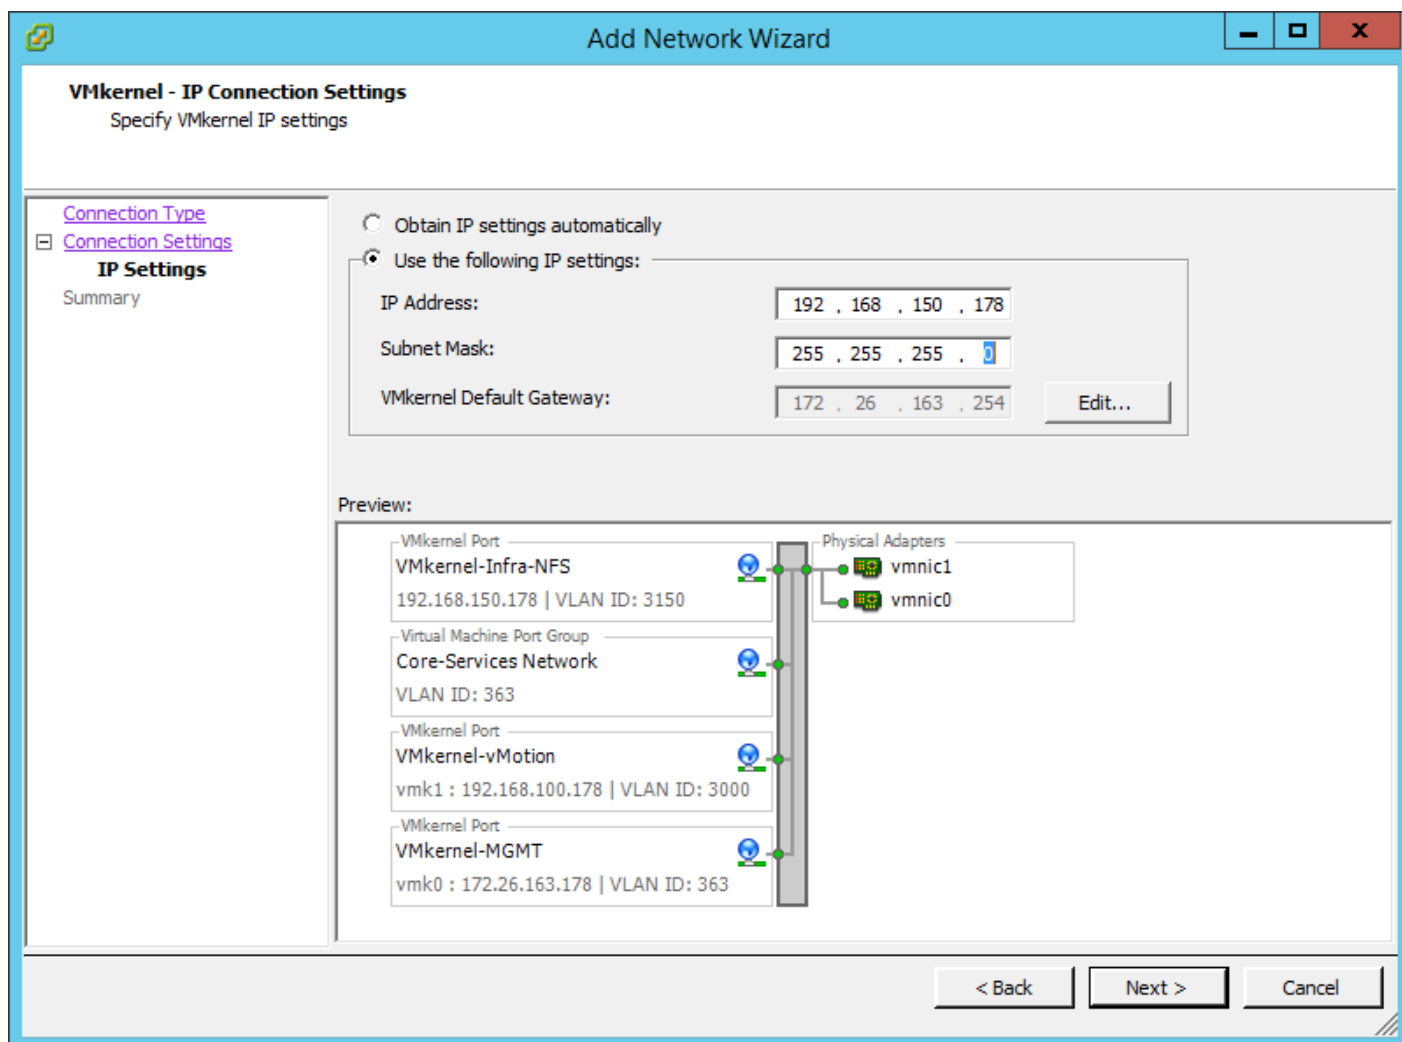
21. Click Next.
22. Click Finish.
23. Select the VMkernel-vMotion configuration and click Edit.
24. Change the MTU to 9000 and click OK.
25. Click Add to add a VMkernel port.
26. Select VMkernel and click Next.
27. Label the network `VMkernel-Infra-NFS` and enter the `<ucs-Infra-NFS-VLAN>` VLAN ID.



Note: It is important to place the Infrastructure NFS VMkernel port on vSwitch0 and not later place it on a vDS. vCenter will have its storage mapped on this Infrastructure NFS port group, and in case of reboot, a vDS cannot operate properly without vCenter present. You could end up in a scenario where vCenter cannot be brought up because the port group it sits on is on the vDS which will not come up until vCenter is present.

28. Click Next.

29. Enter the ESXi host's Infrastructure NFS IP address and Subnet Mask.



30. Click Next.

31. Click Finish.

32. Select the VMkernel-Infra-NFS configuration and click Edit.

33. Change the MTU to 9000 and click OK.

34. Click Close.

35. If this ESXi Host is booted by FCoE and you did not configure iSCSI vNICs in the UCS configuration, proceed to step 109.

36. If this ESXi Host is not iSCSI booted, iSCSI vSwitches and the VMware iSCSI initiator need to be added. Complete steps 37-72. Otherwise proceed to step 73.

37. In the Networking screen select Add Networking.
38. In the popup, select VMkernel to add a VMkernel port in the Infrastructure iSCSI-A subnet. Click Next.
39. Select vmnic2 and click Next.
40. Label the Network `VMkernel-Infra-iSCSI-A`. Do not insert the VLAN ID.



Note: It is important to not set a VLAN ID here because the iSCSI VLAN was set as the Native VLAN of the vNIC and these iSCSI packets should come from the vSwitch without a VLAN tag.

41. Click Next.
42. **Enter an IP Address for this ESXi host's Infra iSCSI-A interface.**



Note: This IP should be in the Infra iSCSI-A subnet, but should not overlap with the iSCSI-A IP Pool set in the Cisco UCS or any of the iSCSI-A LIF IPs on the NetApp Storage controllers.

43. Click Next.
44. Click Finish.
45. Click Properties to the right of the newly created vSwitch1.
46. Select the vSwitch configuration and click Edit.
47. Change the MTU to 9000 and click OK.
48. Select the VMkernel-Infra-iSCSI-A configuration and click Edit.
49. Change the MTU to 9000 and click OK.
50. Click Close.
51. In the Networking screen select Add Networking.
52. In the popup, select VMkernel to add a VMkernel port in the Infrastructure iSCSI-B subnet. Click Next.
53. Select vmnic3 and click Next.
54. Label the Network `VMkernel-Infra-iSCSI-B`. Do not insert the VLAN ID.



Note: It is important to not set a VLAN ID here because the iSCSI VLAN was set as the Native VLAN of the vNIC and these iSCSI packets should come from the vSwitch without a VLAN tag.

55. Click Next.

56. Enter an IP Address for this ESXi host's Infra iSCSI-B interface.



Note: This IP should be in the Infra iSCSI-B subnet, but should not overlap with the iSCSI-B IP Pool set in the Cisco UCS or any of the iSCSI-B LIF IPs on the NetApp Storage controllers.

57. Click Next.

58. Click Finish.

59. Click Properties to the right of the newly created vSwitch2.

60. Select the vSwitch configuration and click Edit.

61. Change the MTU to 9000 and click OK.

62. Select the VMkernel-Infra-iSCSI-B configuration and click Edit.

63. Change the MTU to 9000 and click OK.

64. Click Close.

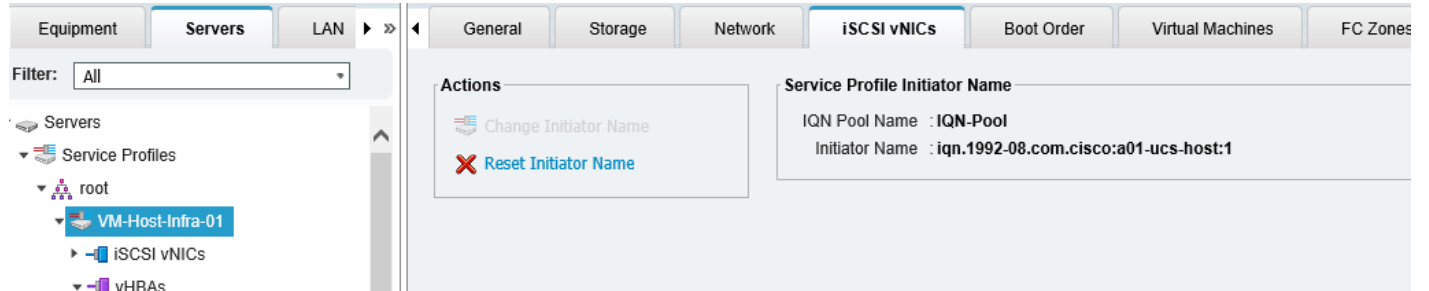
The screenshot displays the VMware vSphere Configuration interface, specifically the Networking section for a vSphere Standard Switch. The interface is divided into Hardware and Software panes on the left, and a main configuration area on the right. The main area shows three vSwitches: vSwitch0, vSwitch1, and vSwitch2. Each vSwitch is connected to physical adapters (vmnic1, vmnic2, vmnic3) and has several VMkernel ports configured. The VMkernel ports are listed with their IP addresses and VLAN IDs.

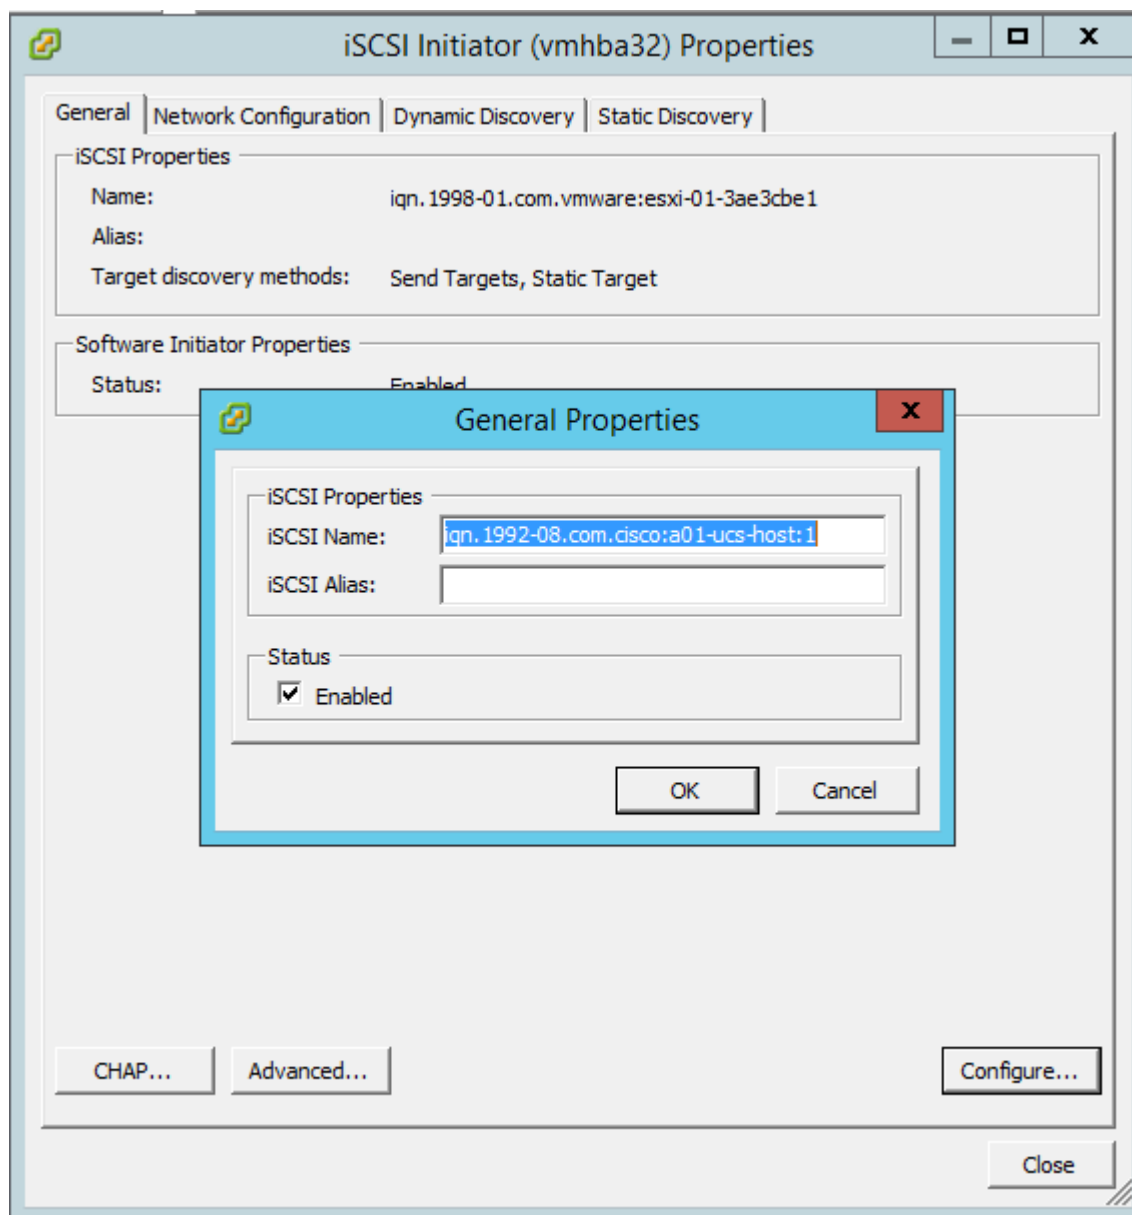
Standard Switch	VMkernel Port	IP Address	VLAN ID
vSwitch0	Core-Services Network	-	363
	VMkernel-Infra-NFS	192.168.150.178	3150
	VMkernel-vMotion	192.168.100.178	3000
vSwitch1	VMkernel-Infra-iSCSI-A	192.168.110.178	-
vSwitch2	VMkernel-Infra-iSCSI-B	192.168.120.178	-

65. On the left, in the Hardware Pane, select Storage Adapters.

66. To add the VMware Software iSCSI Adapter, click Add.

67. Make sure Add Software iSCSI Adapter is selected and click OK.
68. Click OK on the confirmation message.
69. Select the newly added iSCSI Software Adapter and click Properties.
70. In the lower right, click Configure.
71. Change the iSCSI Name to the IQN specified in the UCS Service Profile for this ESXi host. To get this IQN, select the Service Profile in Cisco UCS Manager and select the iSCSI vNICs tab. The IQN is the Initiator Name.





72. Click OK and Close to complete configuration of the Software iSCSI Adapter.

73. If this host is not iSCSI booted, proceed to step 102.

74. In the Hardware pane on the left, select Networking.

75. On the right side of `iScsiBootvSwitch`, click Properties.

76. Select the vSwitch configuration and click Edit.

77. Change the MTU to 9000.

78. Click OK.

79. Select `iScsiBootPG` and click Edit.

80. Change the Network Label to `VMkernel-Infra-iSCSI-A`.
81. Do not set a VLAN for this interface. Change the MTU to 9000.
82. Click Ok.
83. Click Close.
84. In the vSphere Standard Switch view, click Add Networking.
85. Select VMkernel and click Next.
86. Select Create a vSphere standard switch to create a new vSphere standard switch.
87. Select the check boxes for the network adapter vmnic3.
88. Click Next.
89. Change the network label to `VMkernel-Infra-iSCSI-B`. Do not set a VLAN for this interface.
90. Click Next.
91. Enter the IP address and the subnet mask for the iSCSI VLAN B interface for `VM-Host-Infra-01`.



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

92. Click Next.
93. Click Finish.
94. On the right side of `vswitch1`, click Properties.
95. Select the vSwitch configuration and click Edit.
96. Change the MTU to 9000.
97. Click OK.
98. Select `VMkernel-Infra-iSCSI-B` and click Edit.
99. Change the MTU to 9000.
100. Click Ok.
101. Click Close.

The screenshot displays the VMware vSphere Configuration interface for a vSphere Standard Switch. The left-hand pane is divided into 'Hardware' and 'Software' sections. The 'Hardware' section includes links for Health Status, Processors, Memory, Storage, Networking (selected), Storage Adapters, Network Adapters, Advanced Settings, and Power Management. The 'Software' section includes links for Licensed Features, Time Configuration, DNS and Routing, Authentication Services, Virtual Machine Startup/Shutdown, Virtual Machine Swapfile Location, Security Profile, Host Cache Configuration, System Resource Reservation, Agent VM Settings, and Advanced Settings.

The main configuration area shows three standard switches:

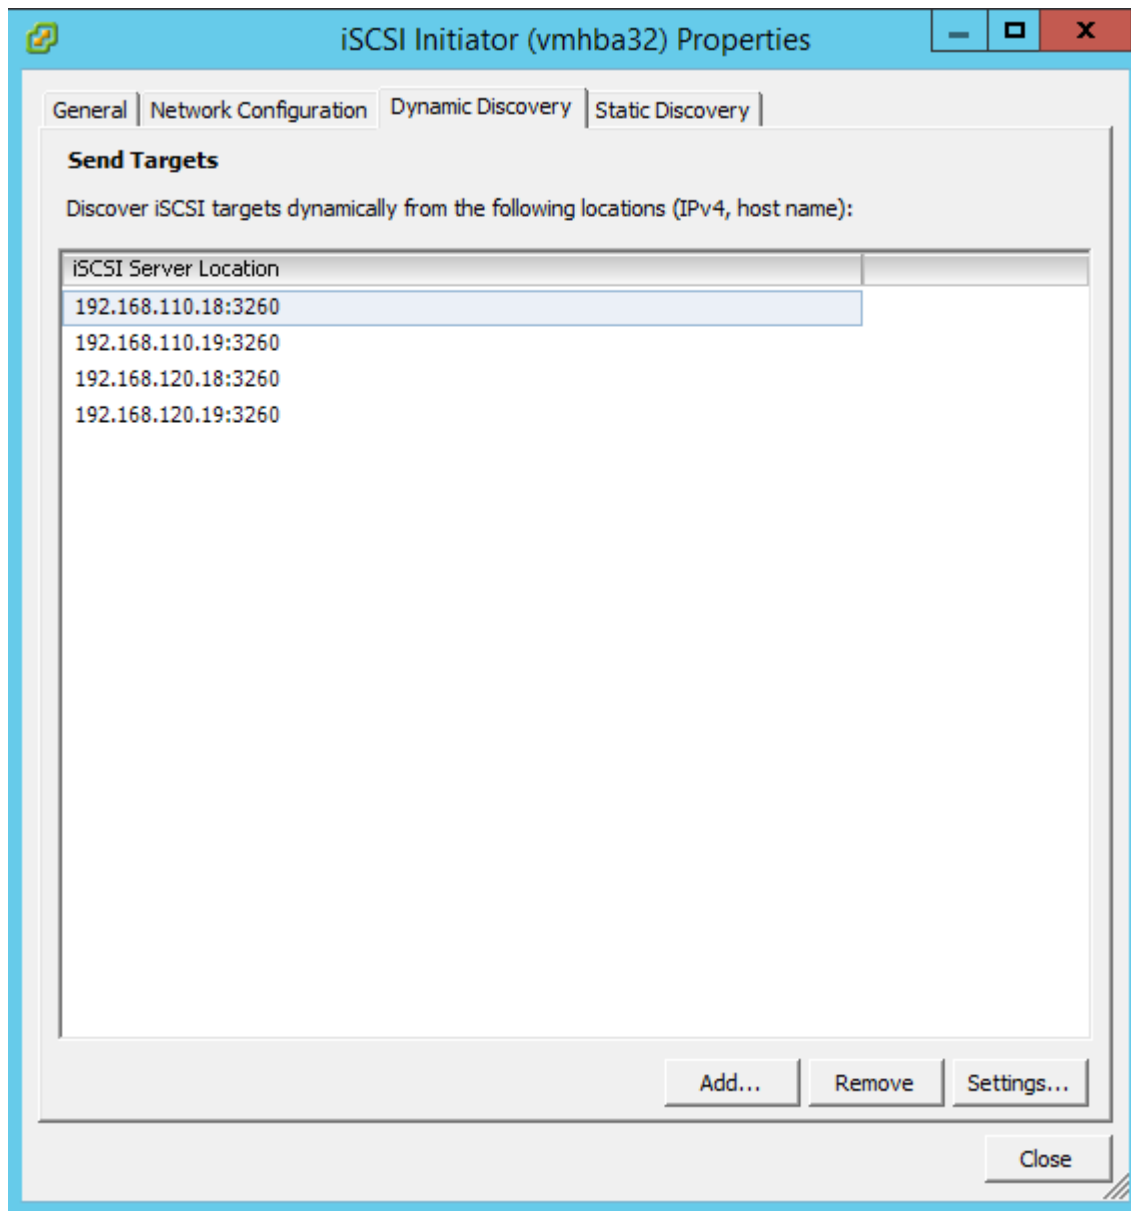
- Standard Switch: vSwitch0**: Contains a Virtual Machine Port Group 'Core-Services Network' (VLAN ID: 363) connected to physical adapters 'vmnic1' and 'vmnic0' (both 40000 Full). It also lists VMkernel ports: 'VMkernel-Infra-NFS' (vmk2: 192.168.150.178 | VLAN ID: 3150), 'VMkernel-vMotion' (vmk1: 192.168.100.178 | VLAN ID: 3000), and 'VMkernel-MGMT' (vmk0: 172.26.163.178 | VLAN ID: 363).
- Standard Switch: vSwitch1**: Contains a VMkernel Port 'VMkernel-Infra-iSCSI-A' (vmk3: 192.168.110.178) connected to physical adapter 'vmnic2' (40000 Full).
- Standard Switch: vSwitch2**: Contains a VMkernel Port 'VMkernel-Infra-iSCSI-B' (vmk4: 192.168.120.178) connected to physical adapter 'vmnic3' (40000 Full).

102. In the Hardware Pane on the left select Storage Adapters.
103. Select the iSCSI Software Adapter, and select Properties.
104. In the iSCSI Initiator Properties window, select the Dynamic Discovery tab.
105. Click Add. Enter the IP Address of iscsi_lif01a from NetApp Storage SVM Infra-SVM.



Note: To get this IP address, ssh into the Storage Cluster CLI and type “network interface show -vserver Infra-SVM -lif iscsi*”.

106. Repeat the previous step adding the IPs for iscsi_lif01b, iscsi_lif02a, and iscsi_lif02b.



107. Click Close.
108. Click Yes to Rescan the host bus adapter.
109. Repeat this entire procedure section for ESXi Host VM-Host-Infra-02.

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

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

- [fnic Driver version 1.6.0.25](#)
- [enic Driver version 2.3.0.7](#)

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

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

1. From each vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. From Resources > Storage, right-click datastore1 and select Browse Datastore.
4. Click the fourth button and select Upload File.
5. Navigate to the saved location for the downloaded VIC drivers and select `fnic_driver_1.6.0.25-3741467.zip`.
6. Click Open and Yes to upload the file to datastore1.
7. Click the fourth button and select Upload File.
8. Navigate to the saved location for the downloaded VIC drivers and select `ESXi60-enic-2.3.0.7-3642661.zip`.
9. Click Open and Yes to upload the file to datastore1.
10. Make sure the files have been uploaded to both ESXi hosts.
11. In the ESXi host vSphere Client, select the Configuration tab.
12. In the Software pane, select Security Profile.
13. To the right of Services, click Properties.
14. Select SSH and click Options at the bottom right.
15. Click Start and OK.



Note: When the SSH service is started this way, it will not be restarted on server reboot.

16. Click OK to close the window.
17. Ensure SSH is started on each host.
18. From the management workstation, start an ssh session to each ESXi host. Login as root with the root password.
19. At the command prompt, run the following commands to account for each host

```
esxcli software vib update -d /vmfs/volumes/datastore1/fnic_driver_1.6.0.25-offline_bundle-3741467.zip
```

```
esxcli software vib update -d /vmfs/volumes/datastore1/ESXi60-enic-2.3.0.7-offline_bundle-3642661.zip
```


reboot

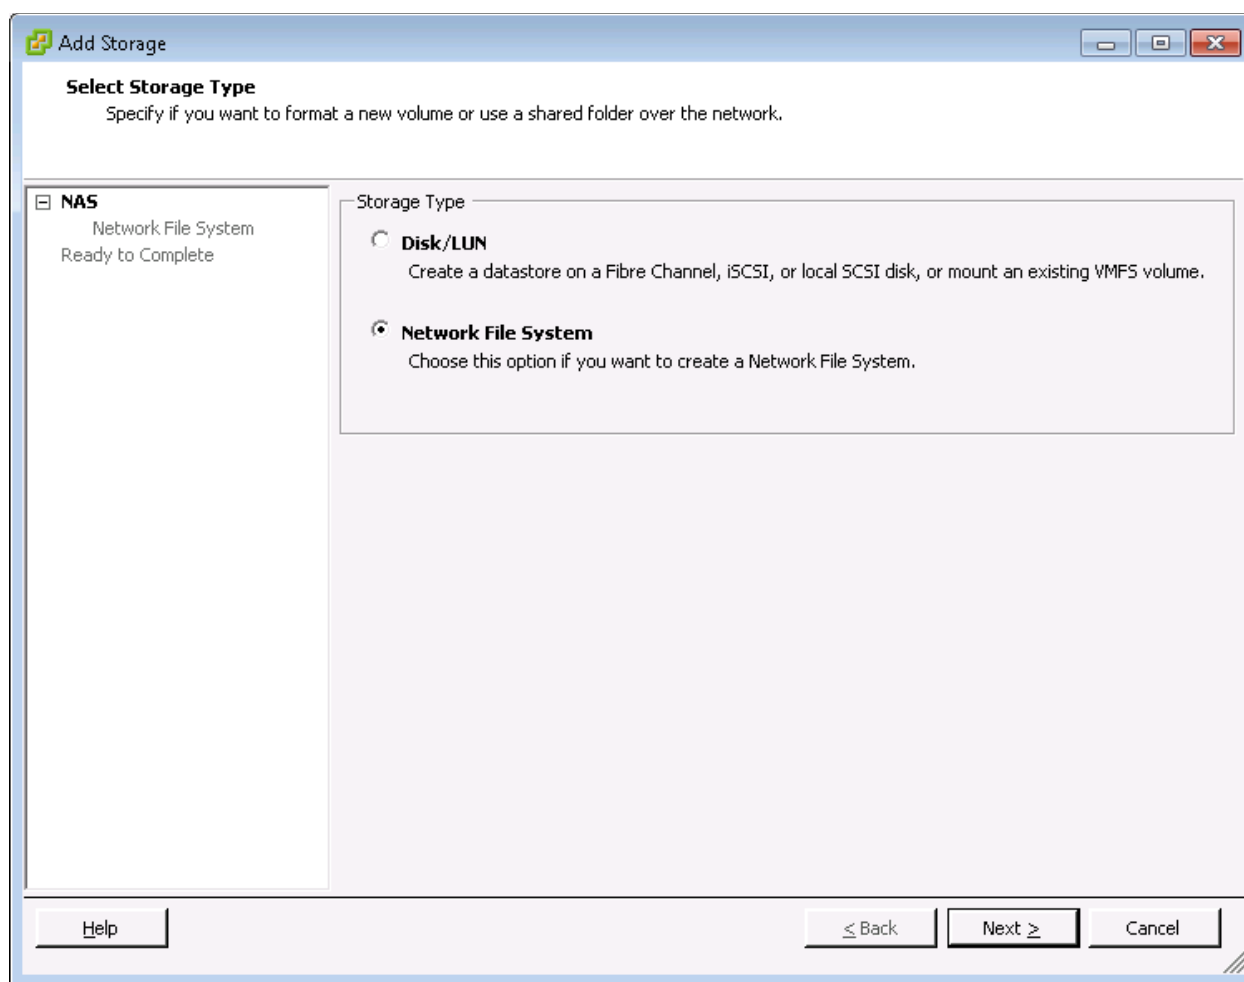
20. After each host has rebooted, log back into each host with vSphere Client.

Mount Required Datastores

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

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

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



5. Select Network File System and click Next.
6. The wizard prompts for the location of the NFS export. Enter the IP address for NetApp Storage LIF `nfs_infra_datastore_1`.



Note: To get the IP addresses of the NFS LIFs, ssh into the Storage Cluster CLI and enter “network interface show -lif nfs*”.

7. Enter /infra_datastore_1 as the path for the NFS export.
8. Confirm that the Mount NFS read only checkbox is not selected.
9. Enter infra_datastore_1 as the datastore name.

10. To continue with the NFS datastore creation, click Next.
11. To finalize the creation of the NFS datastore, click Finish.
12. From the Datastores area, click Add Storage to open the Add Storage wizard.
13. Select Network File System and click Next.
14. The wizard prompts for the location of the NFS export. Enter the IP address for LIF nfs_infra_swap.

15. Enter `/infra_swap` as the path for the NFS export.
16. Confirm that the Mount NFS read only checkbox is not selected.
17. Enter `infra_swap` as the datastore name.
18. To continue with the NFS datastore creation, click Next.
19. To finalize the creation of the NFS datastore, click Finish.
20. Mount both datastores on both ESXi hosts.

Configure NTP on ESXi Hosts

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

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

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click Time Configuration in the Software pane.
4. Click Properties at the upper-right side of the window.
5. At the bottom of the Time Configuration dialog box, click NTP Client Enabled.
6. At the bottom of the Time Configuration dialog box, click Options.
7. In the NTP Daemon (ntpd) Options dialog box, complete the following steps:
 - a. Click General in the left pane and select Start and stop with host.
 - b. Click NTP Settings in the left pane and click Add.
8. In the Add NTP Server dialog box, enter `<ntp-server-ip>` as the IP address of the NTP server and click OK.
9. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes checkbox and click OK.
10. Click OK.
11. In the Time Configuration dialog box, verify that the clock is now set to approximately the correct time.



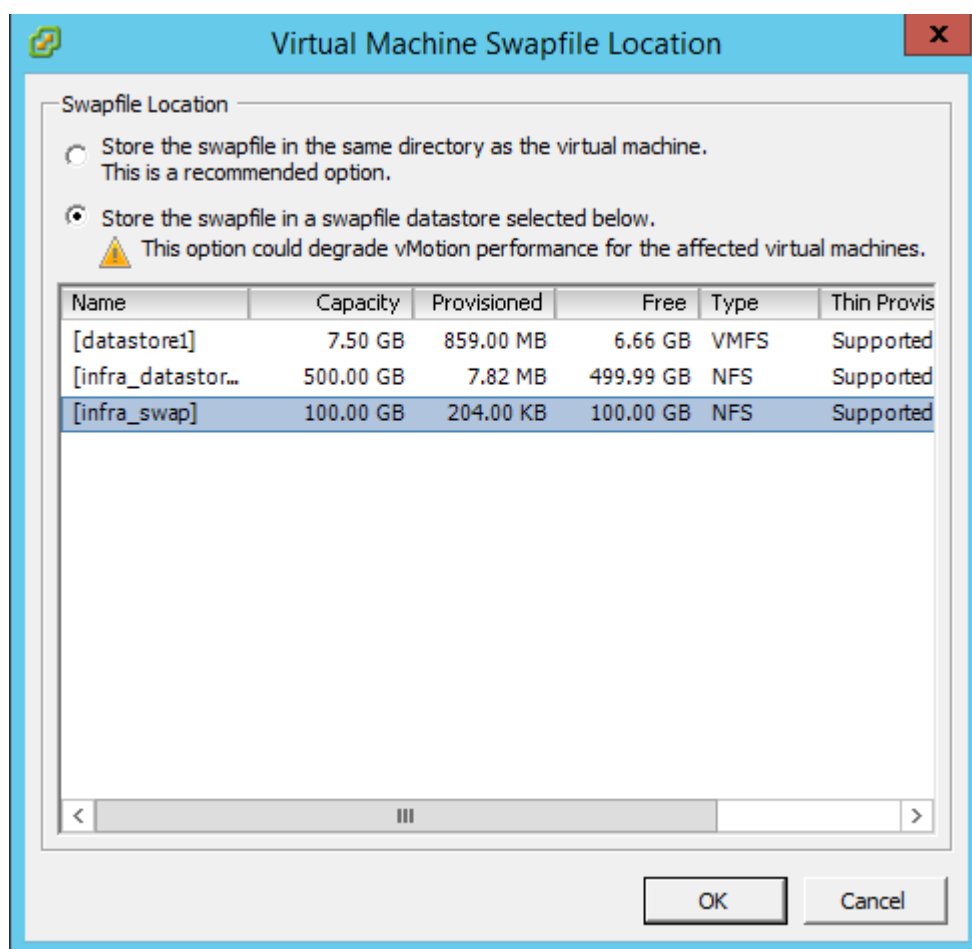
Note: The NTP server time may vary slightly from the host time.

Move VM Swap File Location

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

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

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



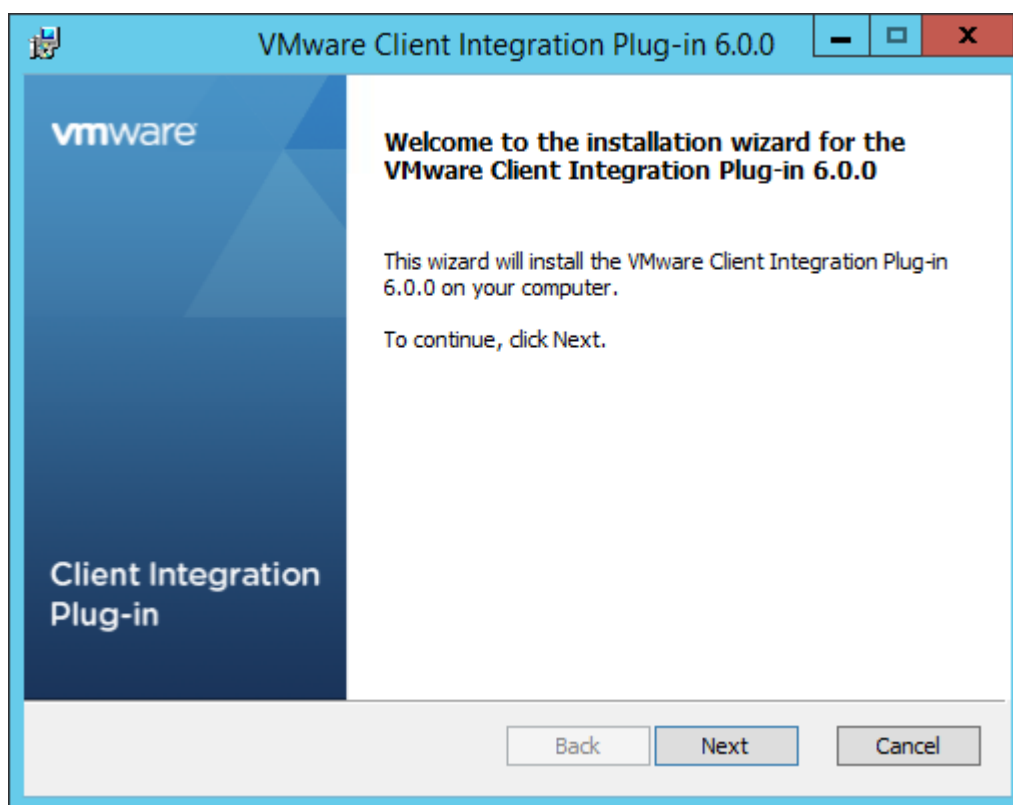
7. Click OK to finalize moving the swap file location.

VMware vCenter 6.0U1b

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.0U1b Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Install the Client Integration Plug-in

1. Download the .iso installer for the version 6.0U1b vCenter Server Appliance and Client Integration Plug-in.
2. Mount the ISO image to the Windows virtual machine (management workstation) on which you want to install the Client Integration Plug-In to deploy the vCenter Server Appliance. This can be done in Windows Server 2012 by copying the VMware-VCSA-all-6.0.0-3343019.iso to the desktop, then right-clicking and selecting Mount.
3. In the mounted iso directory, navigate to the vcsa directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe. The Client Integration Plug-in installation wizard appears.




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

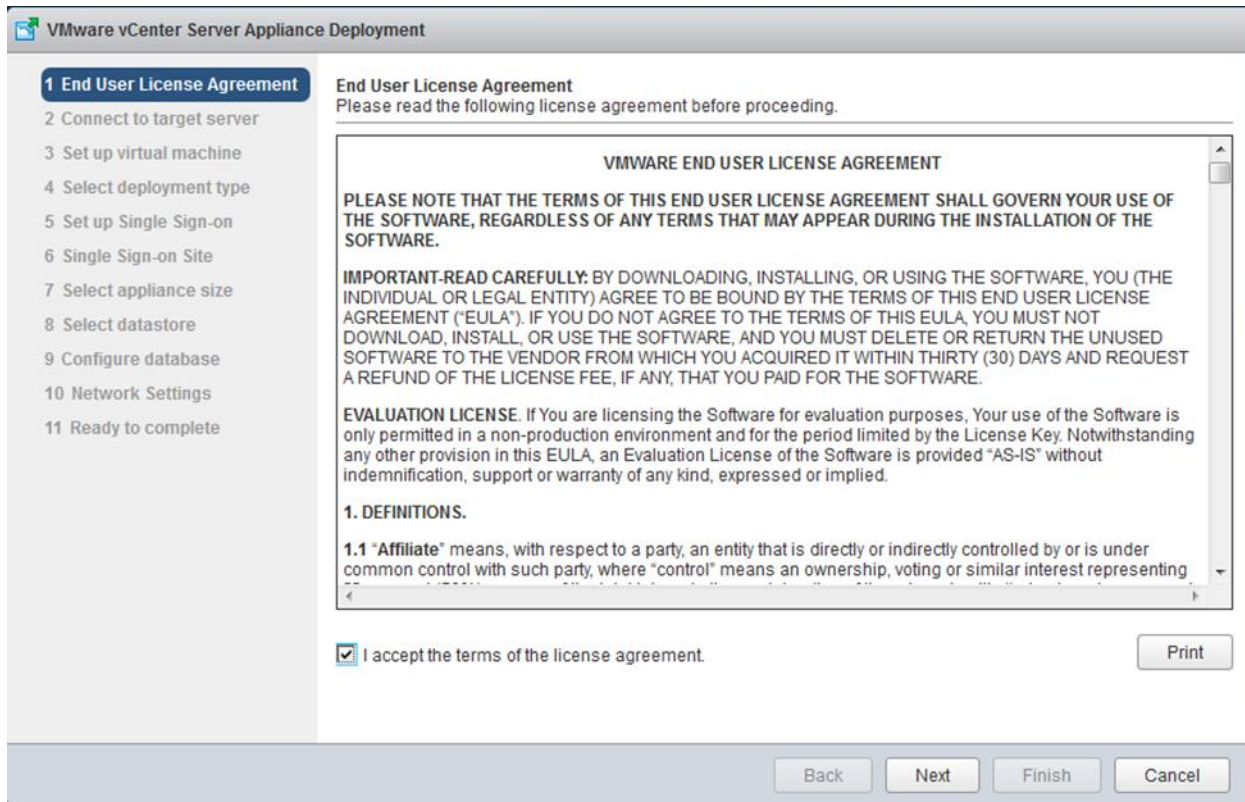
Building the VMware vCenter Server Appliance

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

1. In the mounted iso top-level directory, double-click vcsa-setup.html.
2. Allow the plug-in to run on the browser when prompted.
3. On the Home page, click Install to start the vCenter Server Appliance deployment wizard.



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



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

The screenshot shows the 'VMware vCenter Server Appliance Deployment' wizard. The left sidebar lists 11 steps, with '2 Connect to target server' selected. The main area is titled 'Connect to target server' and contains the following fields and instructions:

- Connect to target server**
Specify the ESXi host or vCenter Server on which to deploy the vCenter Server Appliance.
- FQDN or IP Address:**
- User name:** ⓘ
- Password:**

⚠ Before proceeding, if the target is an ESXi host:

- Make sure the ESXi host is not in lock down mode or maintenance mode.
- When deploying to a vSphere Distributed Switch (VDS), the appliance must be deployed to an ephemeral portgroup. After deployment, it can be moved to a static or dynamic portgroup.

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

6. Click Next.
7. Click Yes to accept the certificate.
8. Enter the Appliance name and password details in the “Set up virtual machine” page.

The screenshot shows the VMware vCenter Server Appliance Deployment wizard. The title bar reads "VMware vCenter Server Appliance Deployment". On the left, a navigation pane lists steps 1 through 11. Steps 1 and 2 are completed, indicated by green checkmarks. Step 3, "Set up virtual machine", is the current step and is highlighted with a blue background. Steps 4 through 11 are listed below it. The main area of the wizard is titled "Set up virtual machine" and contains the instruction "Specify virtual machine settings for the vCenter Server Appliance to be deployed." Below this instruction are four input fields: "Appliance name:" with the value "vc", "OS user name:" with the value "root", "OS password:" with a masked password of ten dots, and "Confirm OS password:" with a masked password of ten dots. Each input field has a small information icon to its right. At the bottom right of the wizard, there are four buttons: "Back", "Next", "Finish", and "Cancel".

VMware vCenter Server Appliance Deployment

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

Set up virtual machine
Specify virtual machine settings for the vCenter Server Appliance to be deployed.

Appliance name: ⓘ

OS user name:

OS password: ⓘ

Confirm OS password:

Back Next Finish Cancel

9. Click Next.

10. In the "Select deployment type" page, choose "Install vCenter Server with an embedded Platform Services Controller".

VMware vCenter Server Appliance Deployment

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

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

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

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

Embedded Platform Services Controller

Install vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

Install Platform Services Controller

Install vCenter Server (Requires External Platform Services Controller)

Back Next Finish Cancel

11. Click Next.

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

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

VMware vCenter Server Appliance Deployment

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

Set up Single Sign-on (SSO)

Create or join a SSO domain. An SSO configuration cannot be changed after deployment.

Create a new SSO domain
 Join an SSO domain in an existing vCenter 6.0 platform services controller

vCenter SSO User name: administrator

vCenter SSO Password: ⓘ

Confirm password:

SSO Domain name: ⓘ

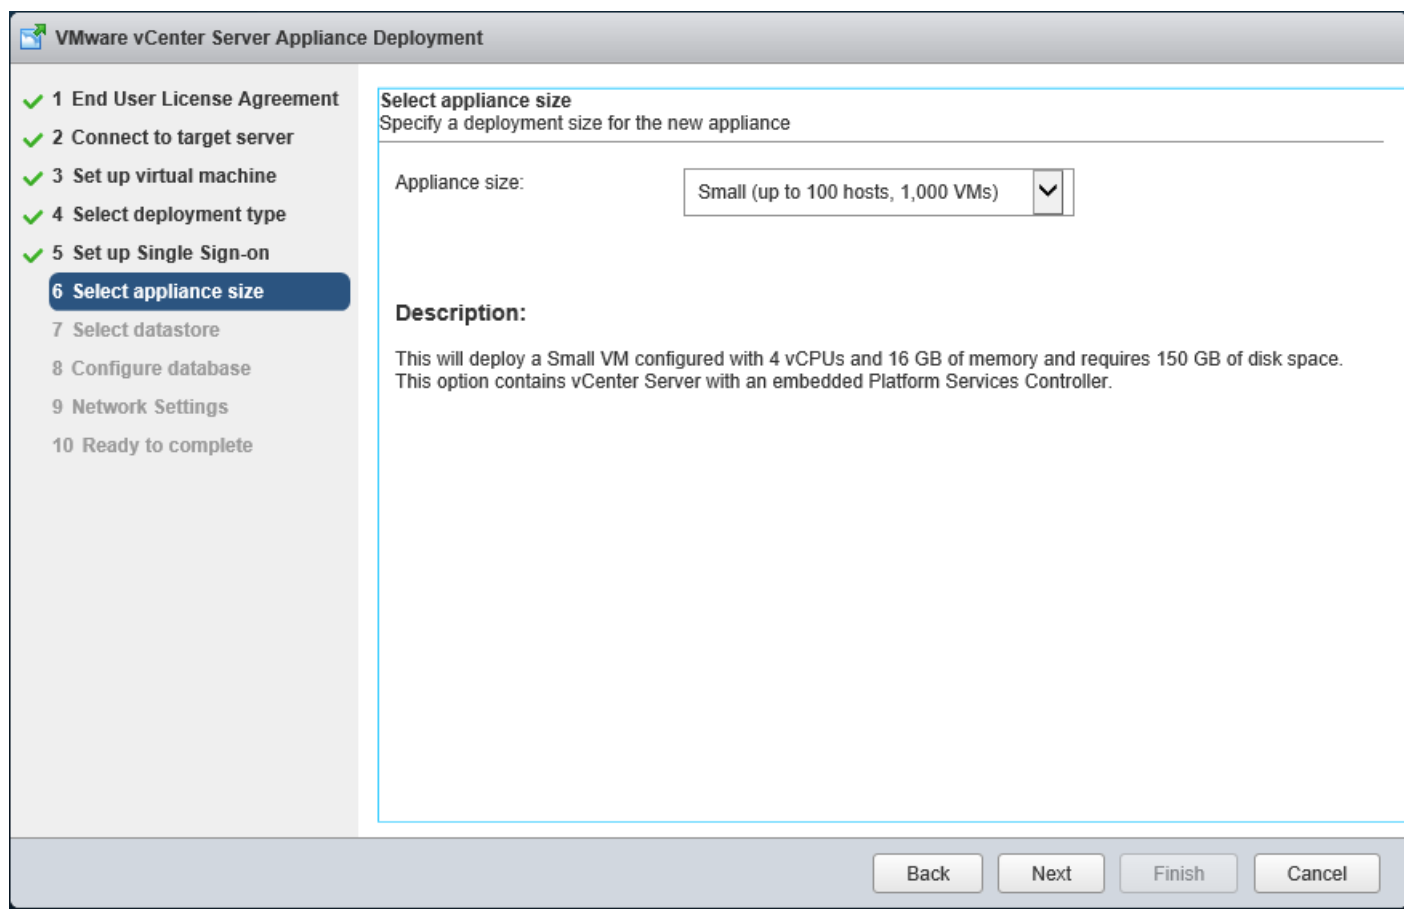
SSO Site name: ⓘ

⚠ Before proceeding, make sure that the vCenter Single Sign-On domain name used is different than your Active Directory domain name.

Back Next Finish Cancel

14. Click Next.

15. Select the appliance size. For example, “Small (up to 100 hosts, 1,000 VMs)”.



16. Click Next.

17. In the "Select datastore" page, choose `infra_datastore_1`. Select the checkbox for Enable Thin Disk Mode.

VMware vCenter Server Appliance Deployment

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

Select datastore
Select the storage location for this deployment

The following datastores are accessible. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Type	Capacity	Free	Provisioned	Thin Provisioni...
datastore1	VMFS	7.5 GB	6.66 GB	0.84 GB	true
infra_datastore...	NFS	500 GB	499.99 GB	0.01 GB	true
infra_swap	NFS	100 GB	100 GB	0 GB	true

Enable Thin Disk Mode ⓘ

Back Next Finish Cancel

18. Click next.

19. Select Use an **embedded database** in the “Configure database” page. Click Next.

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

- Choose a Network: Core-Services Network
- IP address family: IPV4
- Network type: static
- Network address: <vcenter-ip>
- System name: <vcenter-fqdn>
- Subnet mask: <vcenter-netmask>
- Network gateway: <vcenter-gateway>
- Network DNS Servers
- Configure time sync: Use NTP servers

- Enable SSH


21. Review the configuration and click Finish.

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

Setting Up VMware vCenter Server

To setup the VMware vCenter server, complete the following steps:

1. Using a web browser, navigate to <https://<vcenter-ip>:5480>.
2. Log in as root, with the root password entered above in the vCenter installation.
3. On the left select Access. Use the Edit button on the right to change the Access settings according to your local security policy.
4. On the left select Time. Use the appropriate Edit buttons on the right to set the timezone and make any needed adjustments to the Time Synchronization settings.
5. On the left select Administration. Make any needed changes on the right to either the root password or Password expiry settings according to your local security policy.
6. Logout of the VMware vCenter Server Appliance setup interface.
7. Using a web browser, navigate to <https://<vcenter-ip>>.



The screenshot shows the VMware vSphere Web Client interface. At the top is the VMware logo. Below it, the 'Getting Started' section provides instructions on how to access vSphere remotely using the vSphere Web Client, including a link to log in and a link to vSphere Documentation. In the center, there is a 3D illustration of vCenter Servers. On the right side, there are links for Administrators (Web-Based Datastore Browser) and Developers (vSphere Web Services SDK).

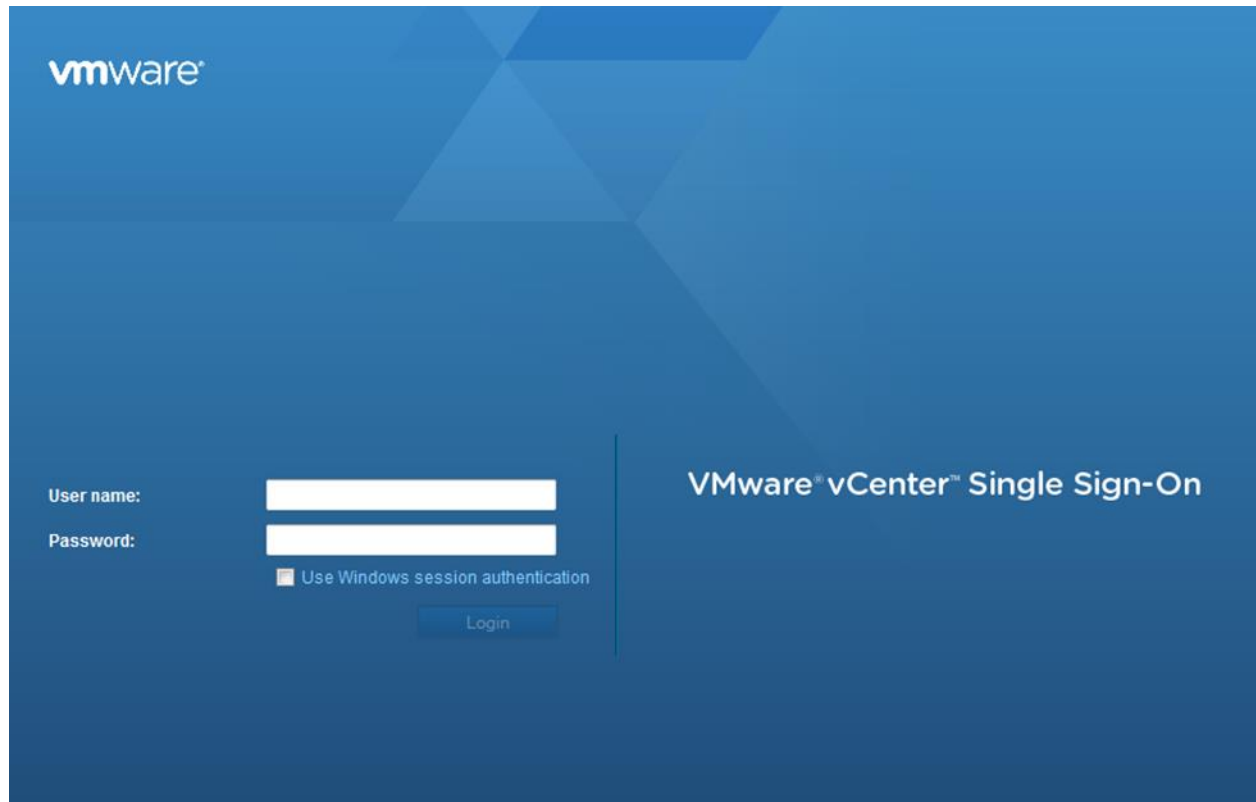
Getting Started
To access vSphere remotely, use the vSphere Web Client.
[Log in to vSphere Web Client](#)
For help, see [vSphere Documentation](#)

For Administrators
Web-Based Datastore Browser
Use your web browser to find and download files (for example, virtual machine and virtual disk files).
[Browse datastores in the vSphere inventory](#)

For Developers
vSphere Web Services SDK
Learn about our latest SDKs, Toolkits, and APIs for managing VMware ESX, ESXi, and VMware vCenter. Get sample code, reference documentation, participate in our Forum Discussions, and view our latest Sessions and Webinars.
[Learn more about the Web Services SDK](#)
[Browse objects managed by vSphere](#)
[Download trusted root CA certificates](#)

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

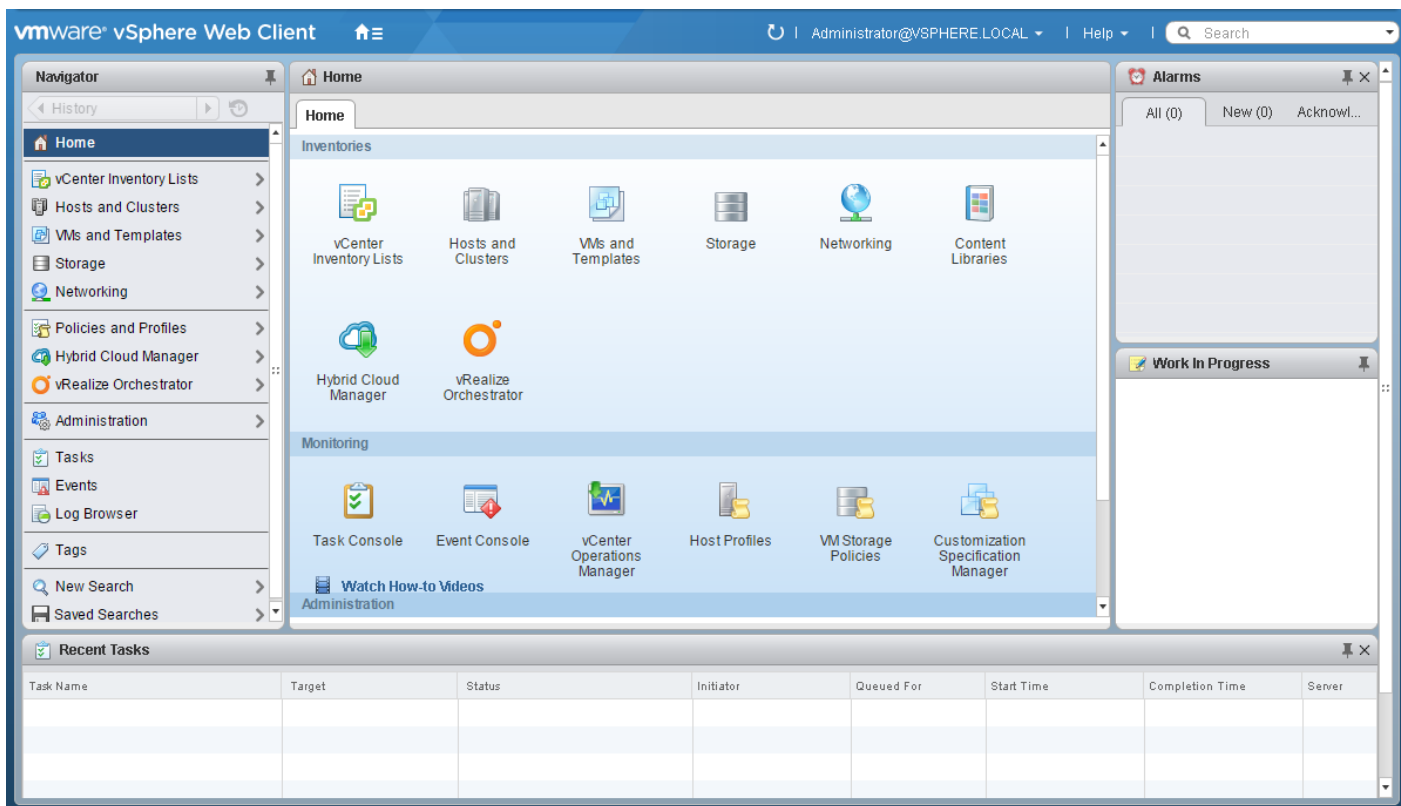
8. Click Log in to vSphere Web Client.



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

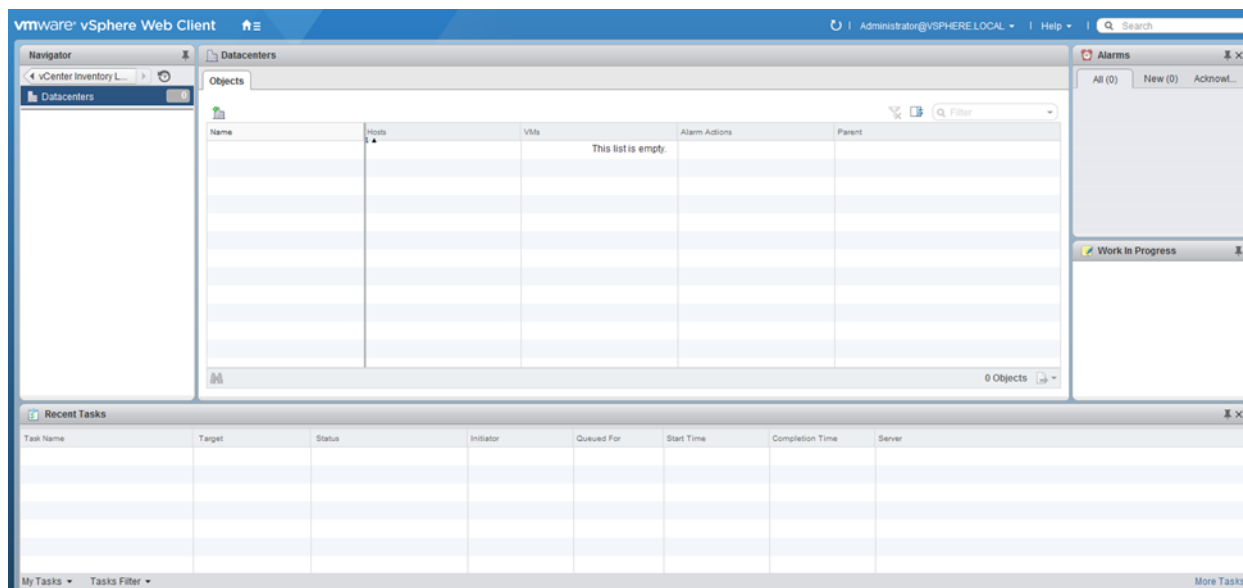


Note: In this lab validation, Google Chrome was used to connect to vCenter server.



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

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




12. To create a Data center, click the icon in the center pane that has a green plus symbol above it.

13. Type "FlexPod-DC" in the Datacenter name field.

14. Select the vCenter Name.

15. Click OK.



16. Right-click the data center FlexPod-DC in the list in the center pane, then select New Cluster from the drop-down.

17. Name the cluster FlexPod-MGMT.

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

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

Name	FlexPod-MGMT
Location	FlexPod-DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	
Admission Control Status	Admission control will prevent powering on VMs that violate availability constraints <input checked="" type="checkbox"/> Enable admission control
Policy	Specify the type of the policy that admission control should enforce. <input checked="" type="radio"/> Host failures cluster tolerates: 1 <input type="radio"/> Percentage of cluster resources reserved as failover spare capacity: Reserved failover CPU capacity: 25 % CPU Reserved failover Memory capacity: 25 % Memory
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low ——— High
EVC	Disable
Virtual SAN	<input type="checkbox"/> Turn ON

OK Cancel

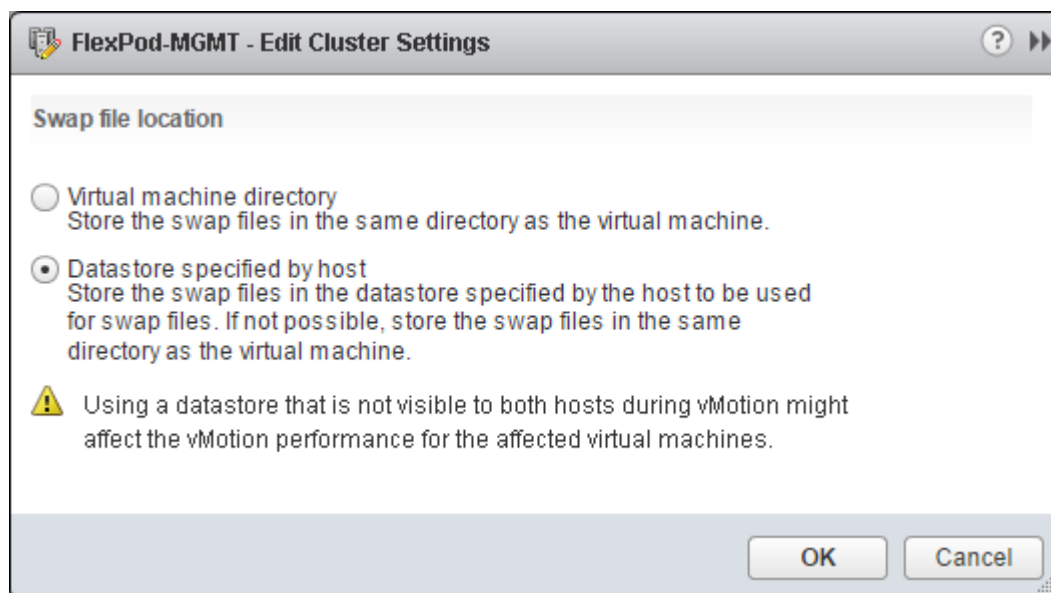
20. Click OK to create the new cluster.

21. On the left pane, double click the “FlexPod-DC” Datacenter.

22. Click Clusters.

23. Under the Clusters pane, right click “FlexPod-MGMT”.

24. Select Settings.
25. In the center pane, click Edit to the right of General.
26. Select Datastore specified by host for the Swap file location.



27. Click OK.
28. On the left, right-click FlexPod-MGMT.
29. Click Add Host.
30. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.
31. Type root as the user name and the root password. Click Next to continue.
32. Click Yes to accept the certificate.
33. Review the host details and click Next to continue.
34. Assign a license or keep the Evaluation License and click Next to continue.
35. Click Next to continue.
36. Click Next to continue.
37. Review the configuration parameters, then click Finish to add the host.
38. Repeat steps 28 to 37 to add the remaining VMware ESXi host to the cluster.

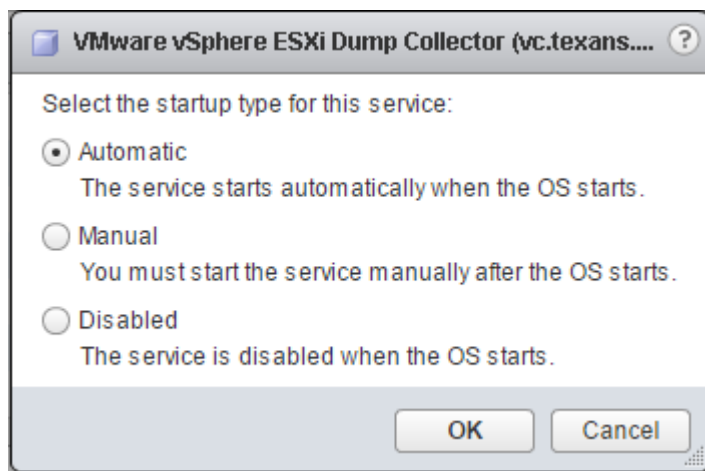


Note: Two VMware ESXi hosts are added to the cluster.

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI and then using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance.

1. In the vSphere web client, select Home.
2. In the center pane, click System Configuration.
3. In the left hand pane, select Services and select VMware vSphere ESXi Dump Collector.
4. In the Actions menu, choose Start.
5. In the Actions menu, click Edit Startup Type.
6. Select Automatic.



7. Click OK.
8. Select Home > Hosts and Clusters.
9. Make sure that FlexPod-DC and FlexPod-MGMT are expanded on the left.
10. For each ESXi host that is iSCSI-booted, right-click the host and select Settings. Scroll down and select Security Profile. Scroll down to Services and select Edit. Select SSH and click Start. Click OK.
11. SSH to each ESXi host that is iSCSI booted. Use root for the user id and the root password. Type the following commands:

```
esxcli system coredump network set --interface-name vmk0 --server-ipv4 <vcenter-  
ip> --server-port 6500
```

```
esxcli system coredump network set --enable true
```

```
esxcli system coredump network check
```

12. Step 10 above can be reversed to turn off SSH on any host servers that have the SSH alarm.

Add Active Directory (AD) Servers to Core-Services Network

It has been assumed in this deployment that preferably an AD Infrastructure or minimally a DNS Infrastructure exists and is reachable from both the Out-of-Band and In-Band Management networks. To make DNS and other AD services such as AD Authentication available as Core Services to FlexPod with ACI Tenants, server virtual machines must be built and added to replication for your AD or DNS Infrastructure. At least two server VMs are recommended with an interface in the Core-Services Network port-group. Detailed steps are not provided here, but this is a very necessary step to offer proper Core Services in this environment. If you have added new AD or DNS servers here, you should modify the DNS servers on the ESXi hosts. You can also modify the DNS servers on the vCenter Services Appliance using Networking on the left in the <https://<vcenter-ip>:5480> interface.

Add AD User Authentication to vCenter (Optional)

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

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).
2. Connect to <https://<vcenter-ip>>, and select Log in to vSphere Web Client.
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. In the center pane, select System Configuration under Administration.
5. On the left, select Nodes and under Nodes select the vCenter.
6. In the center pane, select the manage tab, and within the Settings select Active Directory and click Join.
7. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Click OK.
8. On the left, right-click the vCenter and select Reboot.
9. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.
10. Log back into the vCenter Web Client.
11. In the center pane, select System Configuration under Administration.
12. On the left, select Nodes and under Nodes select the vCenter.
13. In the center pane under the Manage tab, select Active Directory. Make sure your Active Directory Domain is listed.
14. Navigate back to the vCenter Home.

15. In the center pane under Administration, select Roles.
16. On the left under Single Sign-On, select Configuration.
17. In the center pane, select the Identity Sources tab.
18. Click the green + sign to add an Identity Source.
19. Select the Active Directory (Integrated Windows Authentication) Identity source type.
20. Your AD domain name should be filled in. Leave Use machine account selected and click OK.
21. Your AD domain should now appear in the Identity Sources list.
22. On the left, under Single Sign-On, select Users and Groups.
23. In the center pane, select your AD domain for the Domain.
24. Make sure the FlexPod Admin user setup in step 1 appears in the list.
25. On the left under Administration, select Global Permissions.
26. Select the Manage tab, and click the green + sign to add an User or Group.
27. In the Global Permission Root - Add Permission window, click Add.
28. In the Select Users/Groups window, select your AD Domain.
29. Under Users and Groups, select either the FlexPod Admin user or the Domain Admins group.



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

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

Add vSphere Distributed Switch (vDS)

In this FlexPod, you have the choice of using the APIC-controlled VMware vDS or the APIC-controlled Cisco AVS in VXLAN switching mode. The vDS is a Distributed Virtual Switch (DVS) that uses VLANs for network separation and is included in vSphere with Enterprise Plus licensing. If you are installing the vDS in this FlexPod, complete the following steps:

Add vDS in APIC

APIC Advanced GUI

To add the vDS in the APIC Advanced GUI, complete the following steps:

1. Log into the APIC Advanced GUI using the admin user.
2. At the top, click on VM Networking.
3. On the left, select VMware.
4. On the right, click the + sign to add a vCenter Domain.
5. In the Create vCenter Domain window, enter a Virtual Switch Name. A suggested name is <vcenter-name>-vDS. Make sure VMware vSphere Distributed Switch is selected.
6. Select the AEP-UCS Associated Attachable Entity Profile to associate the vDS with the UCS Physical Domain.
7. Use the VLAN Pool drop-down to select Create VLAN Pool.
8. In the Create VLAN Pool window, name the pool VP-<vcenter-name>-vDS.
9. Make sure Dynamic Allocation is selected. Click the + sign to add a VLAN range.
10. In the Create Ranges window, enter the VLAN range that was entered in the Cisco UCS for the APIC-vDS VLANs. Select the Dynamic Allocation Mode.

Create Ranges

Specify the Encap Block Range

Type: **VLAN**

Range: 1101 - 1120
From To

Allocation Mode: **Dynamic Allocation** Inherit allocMode from parent Static Allocation

OK CANCEL

11. Click OK to create the VLAN range.

Create VLAN Pool

Specify the Pool identity

Name: VP-vc-vDS

Description: optional

Allocation Mode: **Dynamic Allocation** Static Allocation

Encap Blocks: × +

VLAN Range	Allocation Mode
[1101-1120]	Dynamic Allocation

SUBMIT CANCEL

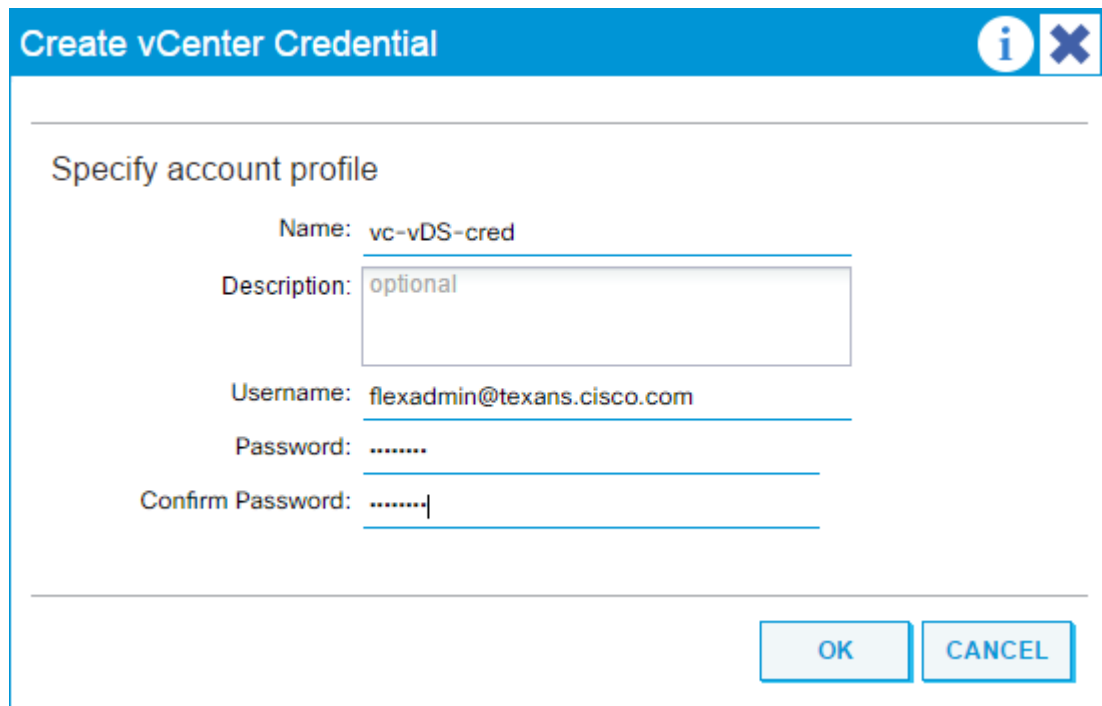
12. Click SUBMIT to create the VLAN Pool.

13. Click the + sign to the right of vCenter Credentials.

14. In the Create vCenter Credential window, for vCenter Credential Name, enter <vcenter-name>-vDS-cred.

15. For Username enter the FlexPod Admin name with the AD Domain name appended.

16. Enter and confirm the password for FlexPod Admin.



Create vCenter Credential

Specify account profile

Name: **vc-vDS-cred**

Description: optional

Username: **flexadmin@texans.cisco.com**

Password:

Confirm Password:

OK CANCEL

17. Click OK to complete adding the credential.
18. Click the + sign on the right of vCenter/vShield to add the vCenter server for APIC to vCenter communication.
19. In the Add vCenter/vShield Controller window, select Type vCenter.
20. Enter a name for the vCenter.
21. Enter the vCenter IP Address or Host Name.
22. For DVS Version, select DVS Version 6.0
23. Enable Stats Collection.
24. For Datacenter, enter the exact vCenter Datacenter name (FlexPod-DC).
25. Do not select a Management EPG.
26. For vCenter Credential Name, select <vcenter-name>-vDS-cred.

Add vCenter/vShield Controller

Specify controller profile

Type: vCenter
 vCenter + vShield

vCenter Controller

Name:

Host Name (or IP Address):

DVS Version:

Stats Collection:

Datacenter:

Management EPG:

Associated Credential:

27. Click OK to add the vCenter Controller.
28. Back in the Create vCenter Domain Window, select the MAC Pinning+ Port Channel Mode.
29. Select both the CDP and LLDP vSwitch Policy.
30. Select the Disabled Firewall Mode.
31. Click SUBMIT to complete creating the vCenter Domain and adding the vDS.
32. At the top, select VM Networking.
33. On the left, expand VMware and select the vDS. Verify that AEP-UCS appears as the Associated Attachable Entity Profile.

The screenshot displays the Cisco vSphere Web Client interface for configuring a vDS. The main content area shows the 'Properties' section for the 'vc-vDS' entity. The configuration details are as follows:

- Name:** vc-vDS
- Virtual Switch:** Distributed Switch
- Associated Attachable Entity Profiles:** AEP-UCS
- Encapsulation:** VLAN mode
- VLAN Pool:** VP-vc-vDS(dynamic)
- Security Domains:** No Security Domains Discovered
- vCenter Credentials:**

Profile Name	Username	Description
vc-vDS-cred	flexadmin@texans.cisco...	

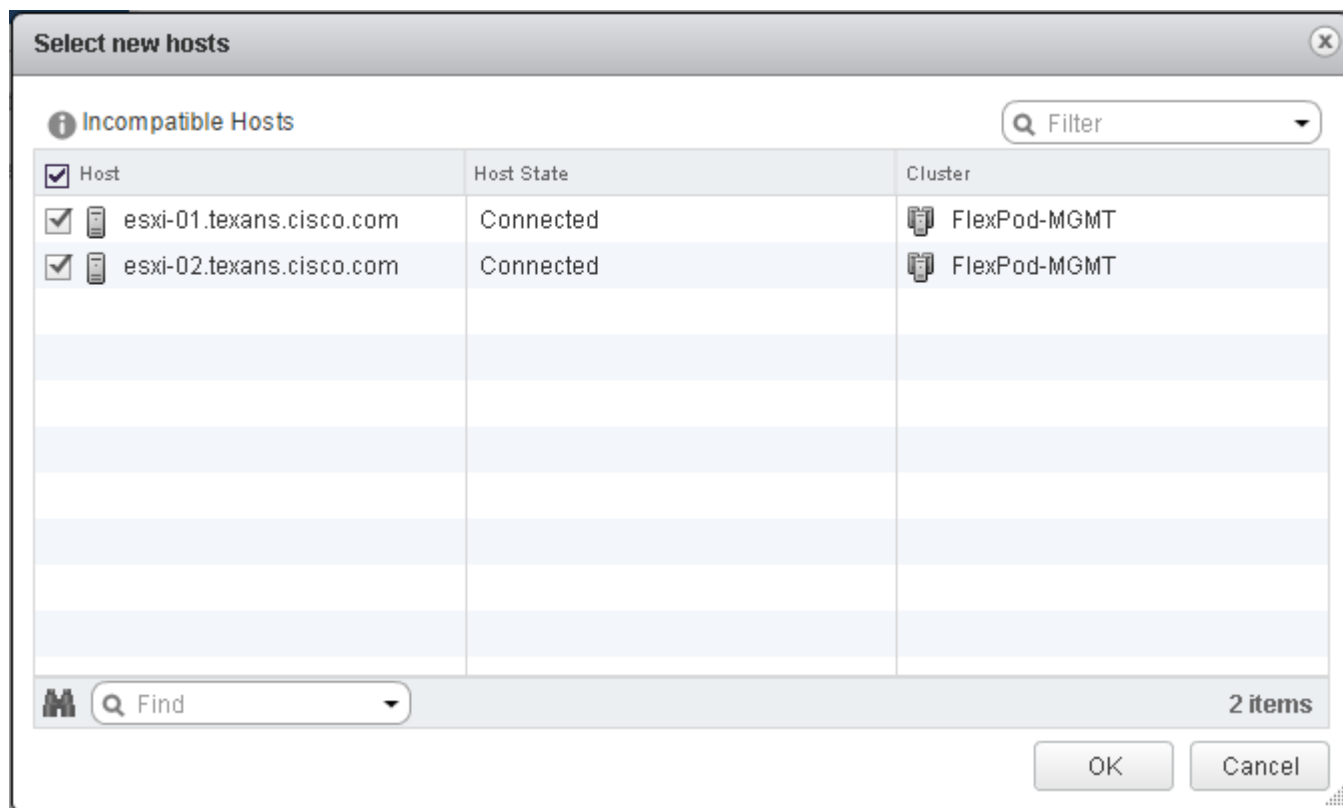
At the bottom right of the configuration pane, there are three buttons: SHOW USAGE, SUBMIT, and RESET.

Add VMware ESXi Host Servers to vDS

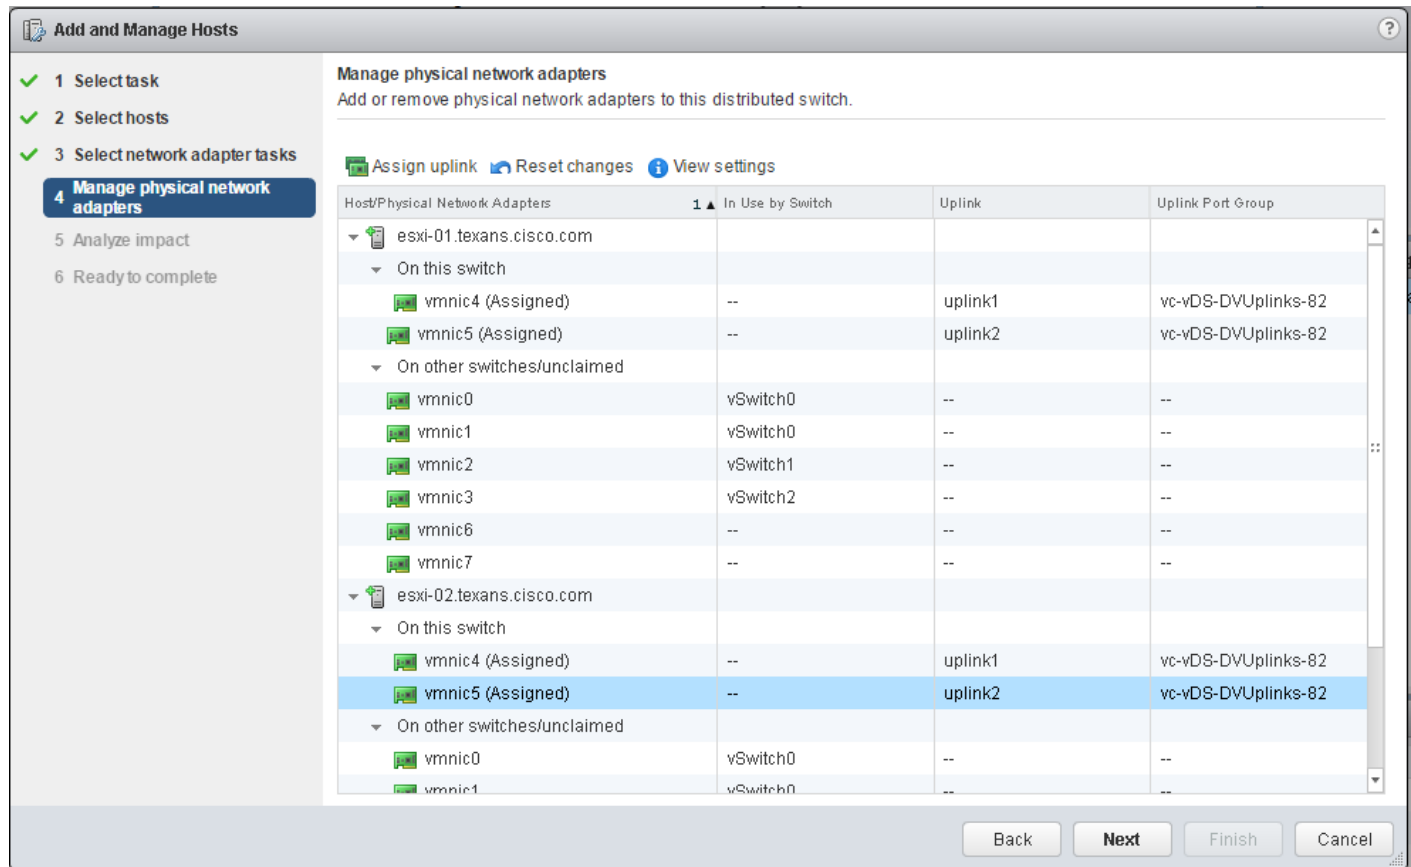
vSphere Web Client

To add the two management VMware ESXi Hosts to the vDS, complete the following steps:

1. Log into the vSphere Web Client as the FlexPod Admin.
2. From the Home screen, select Networking under Inventories.
3. On the left, expand the Datacenter and the vDS folder. Select the vDS switch icon.
4. In the center pane under Basic Tasks, select Add and manage hosts.
5. In the Add and Manage Hosts window, make sure Add hosts is selected and click Next.
6. Click the + sign to add New hosts.
7. In the Select new hosts window, select both of the FlexPod-MGMT hosts.



8. Click OK to complete the host selection.
9. Click Next.
10. Select only "Manage physical adapters". We are not migrating any VMkernel adapters to the vDS.
11. Click Next.
12. On the hosts, select the appropriate vmnics, click Assign uplink, and click OK. Repeat this process until all four vmnics (2 per host) have been assigned. If you have iSCSI vNICs, these will be vmnic4 and vmnic5. If you do not have iSCSI vNICs, these should be vmnic2 and vmnic3.



13. Click Next.

14. Verify that these changes will have no impact and click Next.

15. Click Finish to complete adding the ESXi hosts to the vDS.

16. With the vDS selected on the left, in the center pane select the Related Objects tab.

17. Under Related Objects, select the Hosts tab. Verify the two ESXi hosts are now part of the vDS.

Create In-Band Management Port-Profile on vDS

APIC Advanced GUI

To create an In-Band Management VMware Port-Profile on the vDS that has access to both the Core-Services VMs and the mapped in In-Band Management subnet, complete the following steps:

1. In the APIC Advanced GUI, select Tenants > Foundation.
2. On the left, expand Tenant Foundation, then Application Profiles and IB-MGMT.
3. Under IB-MGMT, right-click Application EPGs and select Create Application EPG.
4. Name the EPG IB-MGMT and make sure Intra EPG Isolation is set to Unenforced.

5. Use the Bridge Domain drop-down to select the common/BD-common-Internal Bridge Domain. Select the default Monitoring Policy.

Create Application EPG

STEP 1 > Identity

1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation:

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

6. Click FINISH to complete creating the EPG.
7. Under the IB-MGMT Application Profile, select and expand Application EPGs and EPG IB-MGMT.
8. Under EPG IB-MGMT, right-click Domains (VMs and Bare-Metal) and select Add VMM Domain Association.
9. In the Add VMM Domain Association window, select the VMware vDS for the VMM Domain Profile.
10. For Deploy and Resolution Immediacy, select Immediate. You do not normally need to change the remaining fields.

Add VMM Domain Association

Choose the VMM domain to associate

VMM Domain Profile: VMware/vc-vDS

Deploy Immediacy: **Immediate** On Demand

Resolution Immediacy: **Immediate** On Demand Pre-provision

VLAN Mode: **Dynamic** Static

Allow Micro-Segmentation:

Allow Promiscuous: Reject

Forged Transmits: Reject

MAC Changes: Reject

SUBMIT **CANCEL**

11. Click SUBMIT to complete the VMM Domain Association.
12. Under EPG IB-MGMT, right-click Contracts and select Add Consumed Contract.
13. In the Add Consumed Contract window, select the common/common-Allow-IB-MGMT contract.
14. Click Submit to complete adding the consumed contract.
15. Under EPG IB-MGMT, right-click Contracts and select Add Consumed Contract.
16. In the Add Consumed Contract window, select the common/common-Allow-Core-Services contract.
17. Click Submit to complete adding the consumed contract.

Tenant Name	Contract Name	Contract Type	Provided / Consumed	QoS Class	State	Label	Subject Label
Contract Type: Contract							
common	common-Allow-Core-Services	Contract	Consumed	Unspecified	formed		
common	common-Allow-IB-MGMT	Contract	Consumed	Unspecified	formed		

18. At the top, select VM Networking.

19. On the left, expand VMware, the vDS, Controllers, the vCenter, the DVS - vDS, and Portgroups.

20. Select the Foundation|IB-MGMT|IB-MGMT Portgroup.

21. Verify the assigned VLAN under Properties.

Properties

Name: Foundation|IB-MGMT|IB-MGMT

Primary VLAN for Micro-Seg: **unknown**

Port Encap (or Secondary VLAN for Micro-Seg): **vlan-1108**

VM Name	Name	State	MAC	IP Address
No items have been found. Select Actions to create a new item.				



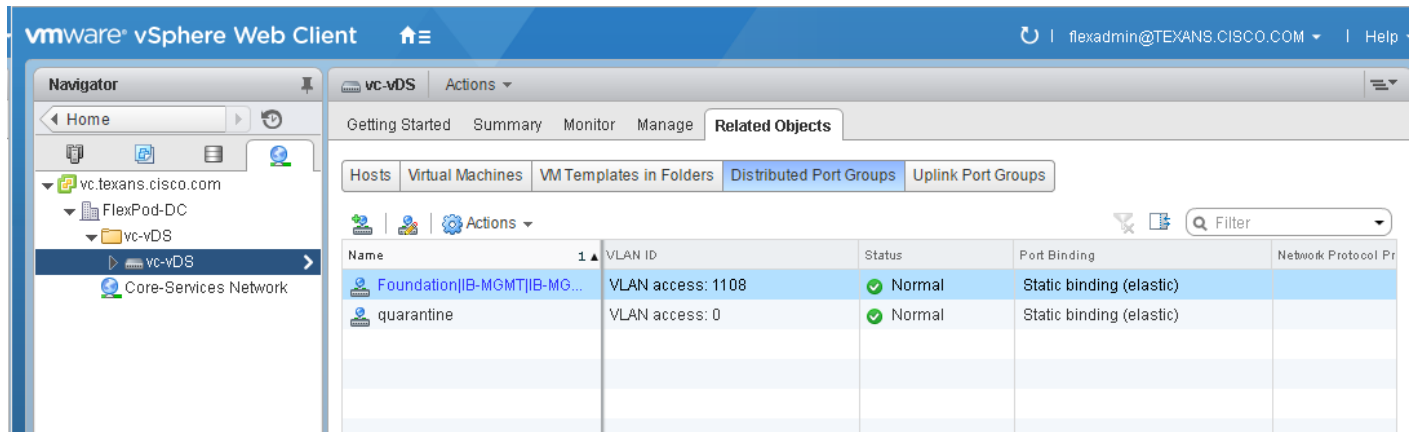
Note: The IB-MGMT port profile can now be assigned to a VM.

22. In the vSphere Web Client, from the Home screen, select Networking under Inventories.

23. On the left, expand the vCenter, the Datacenter, and the vDS folder and select the vDS switch icon.

24. On the right, under Related Objects, select Distributed Port Groups.

25. In the list of port groups, select Foundation|IB-MGMT|IB-MGMT. Again, verify the assigned VLAN.



Add Cisco Application Virtual Switch (AVS)

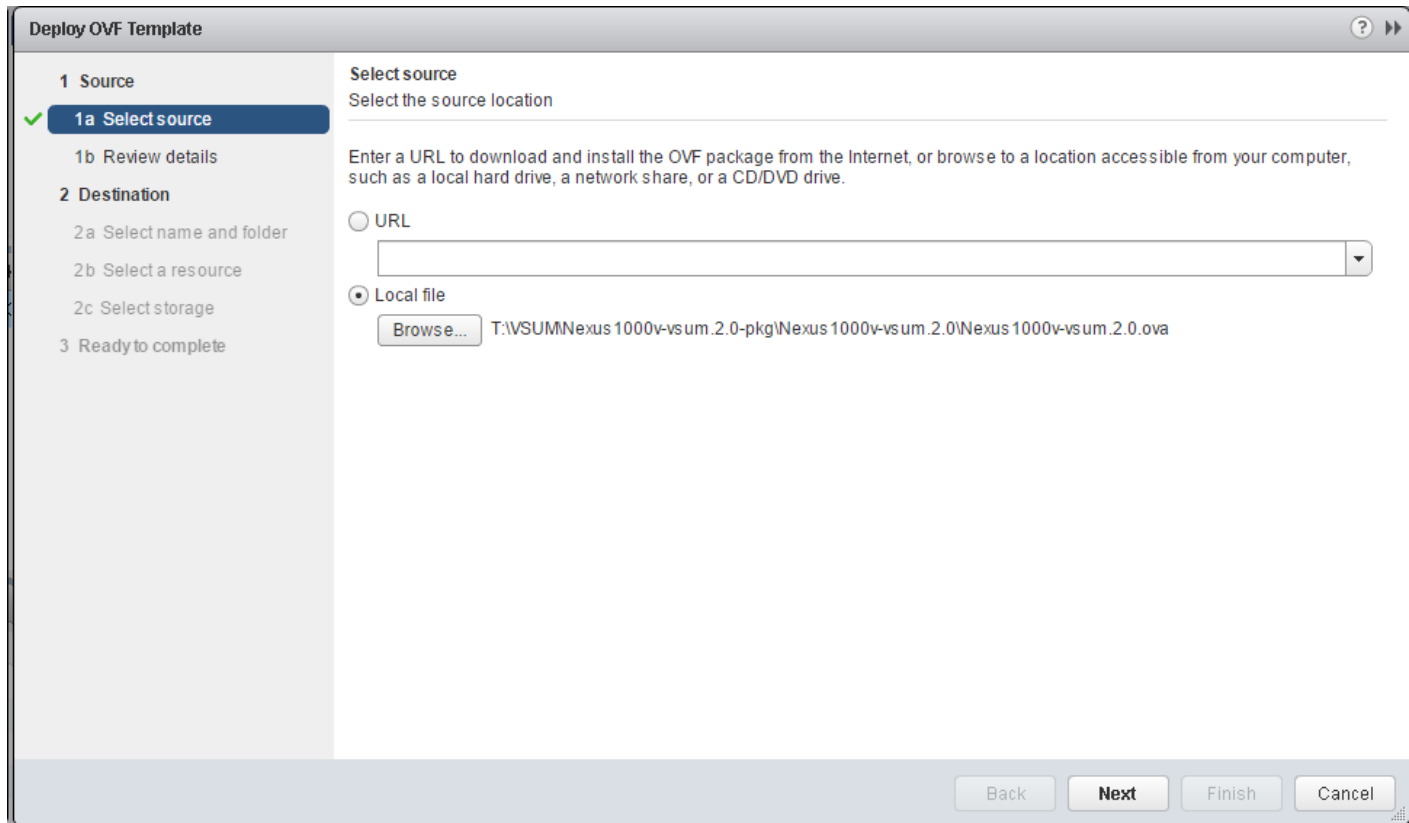
In this FlexPod, you have the choice of using the APIC-controlled VMware vDS or the APIC-controlled Cisco AVS in VXLAN switching mode. The AVS in VXLAN switching mode is a Vendor Specific Distributed Virtual Switch (DVS) that uses VXLANs for network separation and is offered at no additional cost with VMware Enterprise Plus licensing. If you are installing the AVS in this FlexPod, complete the following steps:

Install Cisco Virtual Switch Update Manager (VSUM) Virtual Appliance

vSphere Web Client

To install VSUM into your FlexPod Management Cluster, complete the following steps:

1. Download and unzip the VSUM Release 2.0 .zip file from [Cisco VSUM 2.0 Download](#).
2. In the Nexus100v-vsum.2.0.pkg folder that is unzipped from the downloaded zip, unzip the Nexus1000v-vsum.2.0.zip file.
3. Log into vSphere Web Client as the FlexPod Admin user.
4. From the Home screen, on the left, select VMs and Templates.
5. Select the vCenter on the left and using the Actions drop-down in the center pane, select Deploy OVF Template.
6. If a Security Prompt pops up, click Allow to allow the Client Integration Plugin to run.
7. In the Deploy OVF Template window, select Local file, then Browse and browse to the Nexus1000v-vsum.2.0.ova file downloaded and unzipped above.
8. Select the file and click Open.



9. Click Next.

10. Review the details and click Next.

11. Click the Accept button to accept the License Agreement and click Next.

12. Give the VM a name and select the FlexPod-DC datacenter. Click Next.

13. Select the FlexPod-MGMT ESXi Cluster and click Next.

14. Select infra_datastore_1 and make sure the Thin Provision virtual disk format is selected. Click Next.

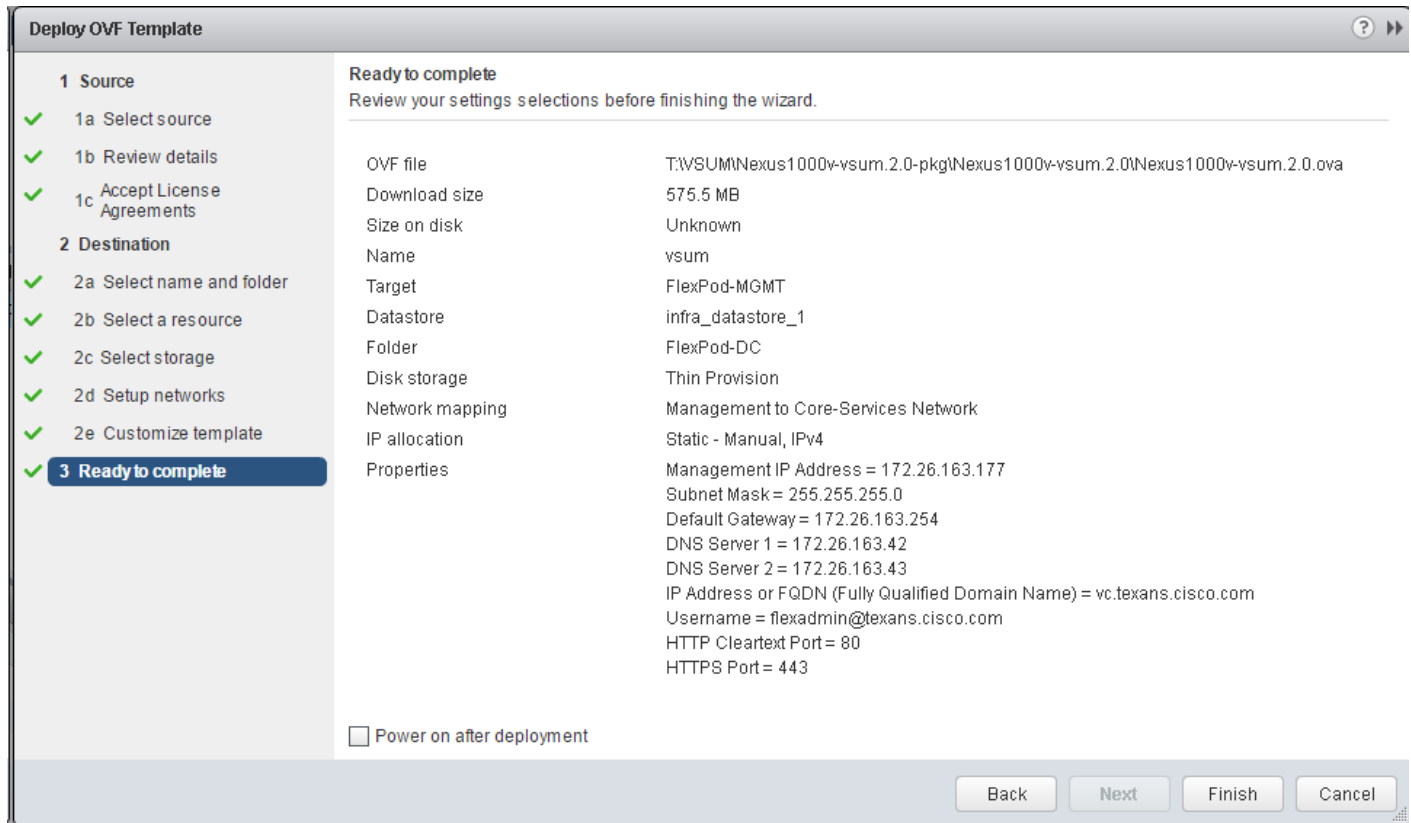
15. Make sure the Core-Services Network is chosen and click Next.

16. Fill in all IP, DNS, and vCenter properties for the VSUM Appliance and click Next.

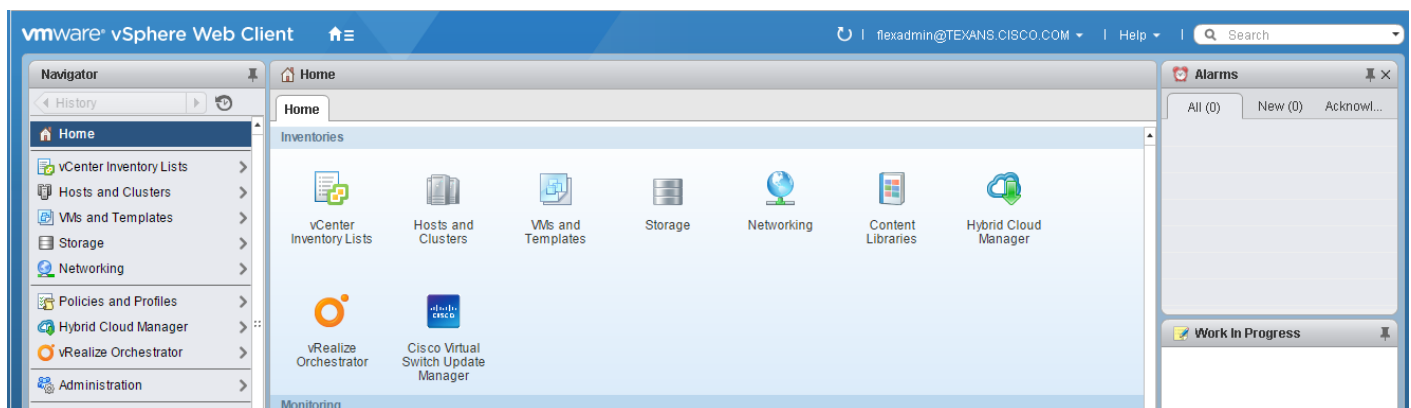


Note: The VSUM IP address should be in the IB-MGMT subnet.

17. Review all values and click Finish to complete the deployment of the VSUM Appliance.



18. On the left, expand the vCenter and Datacenter. Right-click the VSUM VM and select Power > Power On.
19. Right-click the VSUM VM again and select Open Console. When a login prompt appears, close the console.
20. Log out and Log back in to the vSphere Web Client.
21. Verify that Cisco Virtual Switch Update Manager now appears in the center pane under Inventories.



Add Cisco AVS in APIC

APIC Advanced GUI

To add the AVS in the APIC GUI, complete the following steps:

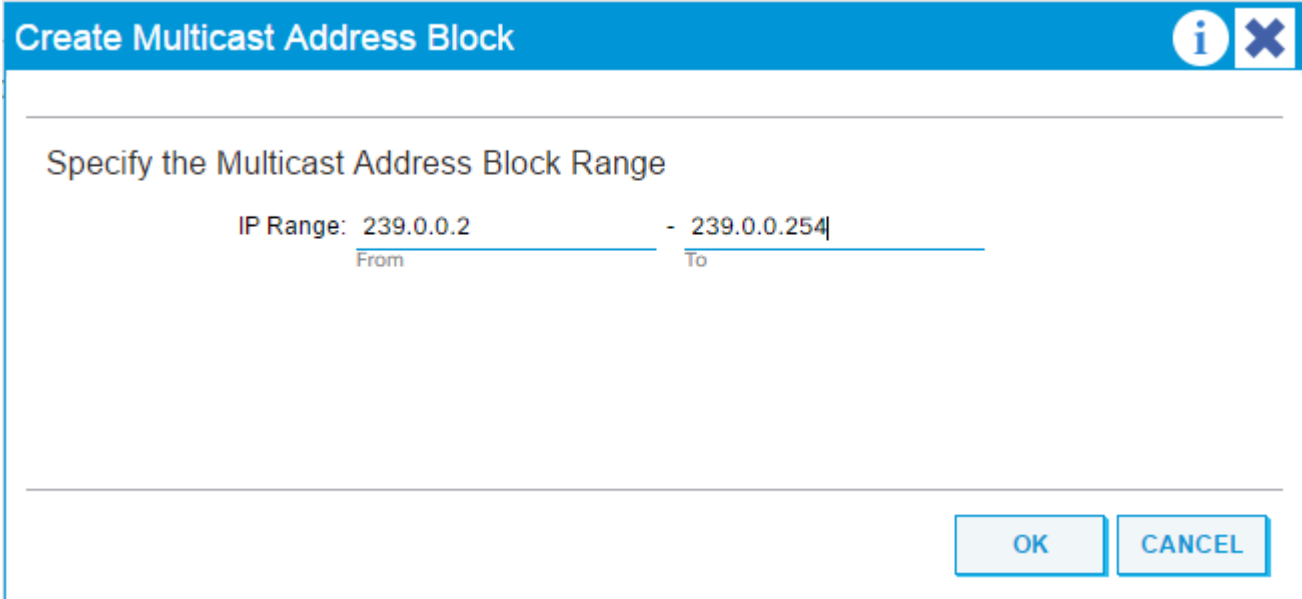
1. Log into the APIC Advanced GUI using the admin user.
2. At the top, select Fabric > Access Policies.
3. On the left, expand Global Policies and Attachable Access Entity Profiles. Select AEP-UCS.
4. Select the checkbox next to Enable Infrastructure VLAN.

The screenshot shows the Cisco APIC GUI for the 'Attachable Access Entity Profile - AEP-UCS'. The left sidebar shows the navigation tree with 'Global Policies' > 'Attachable Access Entity Profiles' > 'AEP-UCS' selected. The main content area shows the 'Properties' section with the following details:

- Name: AEP-UCS
- Description: optional
- Enable Infrastructure VLAN:
- Domains (VMM, Physical or External) Associated to Interfaces:

Name	State
PD-UCS (Physical)	formed
vc-vDS (Vmm-VMware)	formed

5. Click SUBMIT to complete modification of the AEP.
6. At the top, select VM Networking. On the left, select VMware.
7. From VM Networking > Inventory > VMware, on the right, click the + sign to add a vCenter Domain.
8. In the Create vCenter Domain window, enter a Virtual Switch Name. A suggested name is <vcenter-name>-AVS. Select the Cisco AVS Virtual Switch.
9. For Switching Preference, select Local Switching.
10. For Encapsulation, select VXLAN.
11. For Associated Attachable Entity Profile, select AEP-UCS.
12. For the AVS Fabric-Wide Multicast Address, enter a Multicast Address. A suggested entry is 239.0.0.1.
13. For the Pool of Multicast Addresses (one per-EPG), use the drop-down to select Create Multipath Address Pool.
14. Name the pool MAP-<vcenter-name>-AVS.
15. Click the + sign to create an Address Block.
16. Enter a multicast address IP Range. A suggested range is 239.0.0.2 to 239.0.0.254.



Create Multicast Address Block

Specify the Multicast Address Block Range

IP Range: 239.0.0.2 - 239.0.0.254
From To

OK CANCEL

17. Click OK to complete creating the Multicast Address Block.

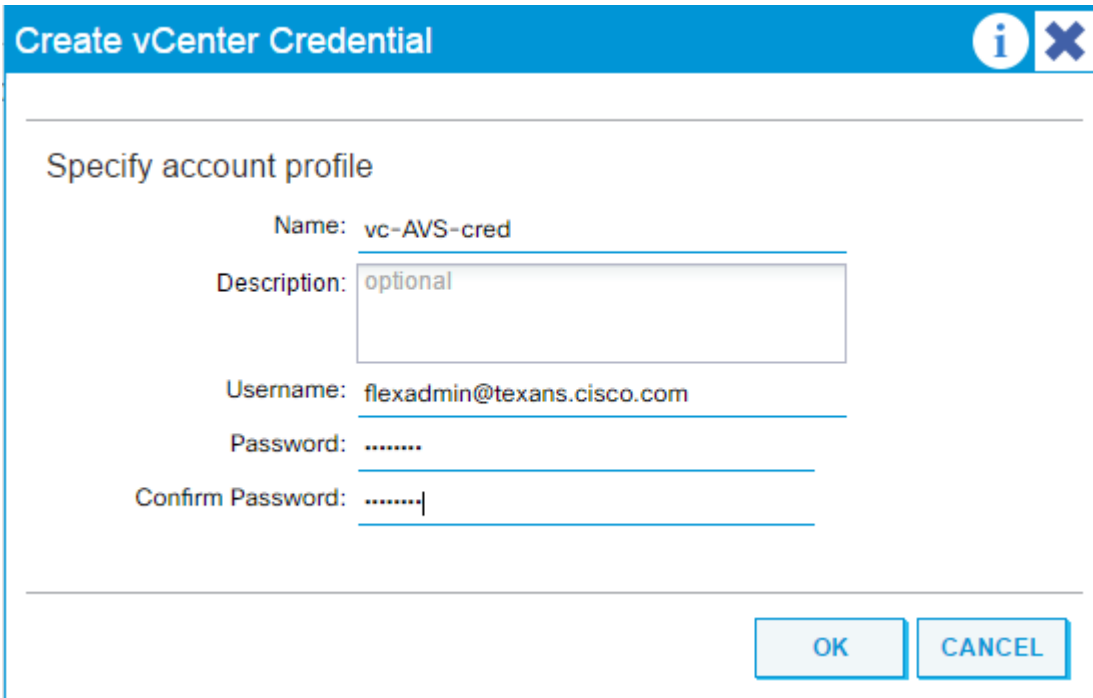
18. Click SUBMIT to complete creating the Multicast Address Pool.

19. Click the + sign to the right of vCenter Credentials to add vCenter Credentials.

20. In the Create vCenter Credential window, name the account profile <vcenter-name>-AVS-cred.

21. For Username, enter the FlexPod Admin user with the AD domain appended.

22. Enter and confirm the FlexPod Admin password.



Create vCenter Credential

Specify account profile

Name: vc-AVS-cred

Description: optional

Username: flexadmin@texans.cisco.com

Password:

Confirm Password:

OK CANCEL

23. Click OK to complete adding the vCenter credentials.

24. Click the + sign on the right of vCenter to add the vCenter server for APIC to vCenter communication.
25. In the Create vCenter Controller window, enter <vcenter-name> for Name.
26. Enter the vCenter IP Address or Host Name.
27. For DVS Version, select DVS Version 6.0
28. For Datacenter, enter the exact vCenter Datacenter name (FlexPod-DC).
29. Do not select a Management EPG.
30. For Associated Credential, select <vcenter-name>-AVS-cred.

Create vCenter Controller

Specify controller profile

Type: vCenter

Name: vc

Host Name (or IP Address): 172.26.163.175

DVS Version: DVS Version 6.0

Datacenter: FlexPod-DC

Management EPG: select an option

Associated Credential: vc-AVS-cred

OK CANCEL

31. Click OK to complete adding the vCenter Controller.
32. For Port Channel Mode, select MAC Pinning+.
33. For vSwitch Policy, select CDP and LLDP. Do not select BPDU Guard or BPDU Filter.
34. For Firewall Mode, select Disabled.

Specify vCenter domain users and controllers

Virtual Switch Name:

Virtual Switch: **Cisco AVS**

Switching Preference: **Local Switching**

Encapsulation: VLAN VXLAN

Associated Attachable Entity Profile:

AVS Fabric-Wide Multicast Address:
Must Use a Multicast Address different from the Pool of Multicast Addresses.

Pool of Multicast Addresses (one per-EPG):

Security Domains:

Name	Description

vCenter Credentials:

Profile Name	Username	Description
vc-AVS-cred	flexadmin@texans...	

vCenter:

Name	IP	Type	Stats Collection
vc	172.26.163.175	vCenter	Disabled

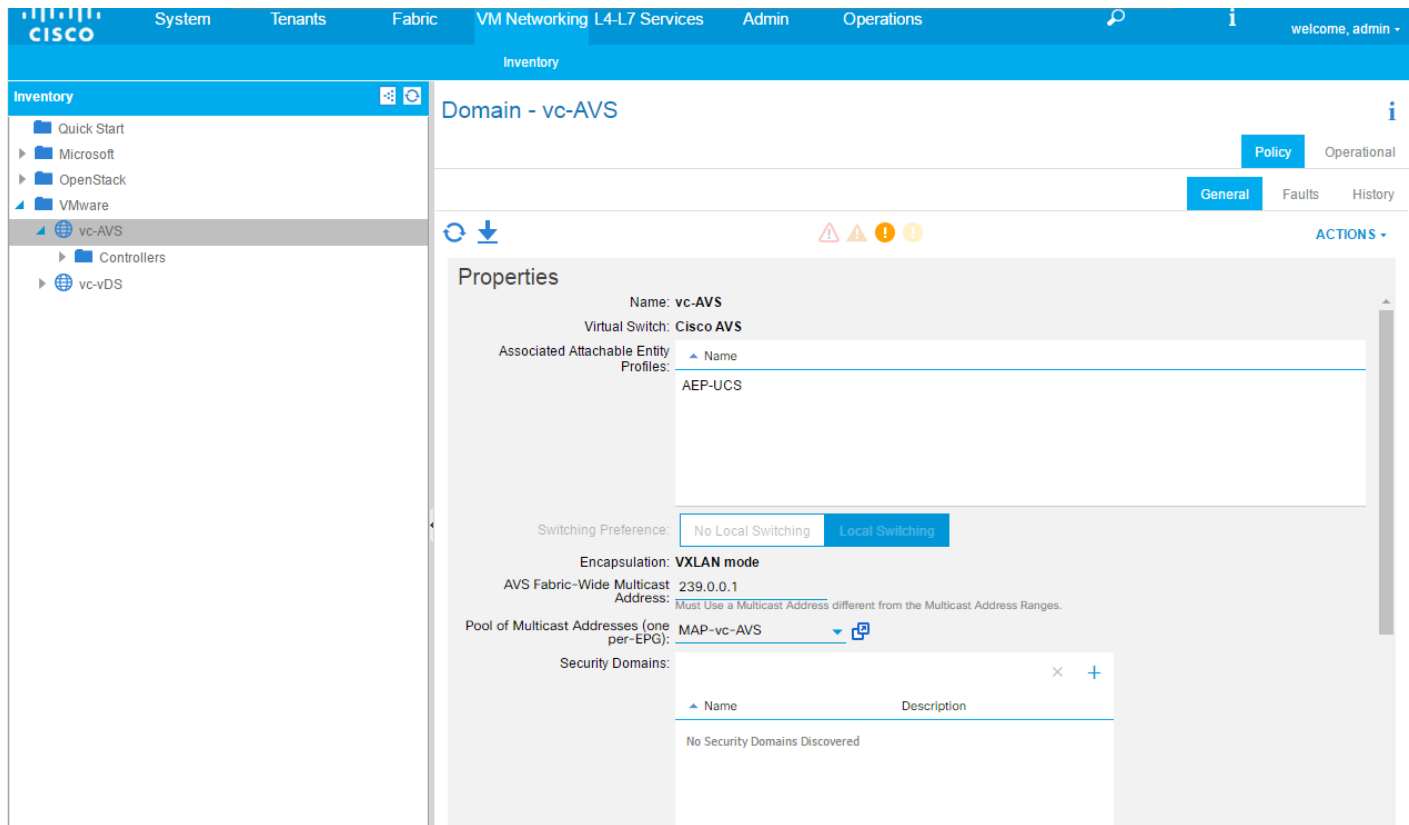
Port Channel Mode:

vSwitch Policy: CDP LLDP BPDU Guard BPDU Filter

Firewall Mode:

35. Click SUBMIT to add the vCenter Domain and Cisco AVS.

36. On the left, expand VMware and select the AVS. Verify that AEP-UCS is specified for Associated Attachable Entity Profiles.



Add VMware ESXi Host Servers to AVS

vSphere Web Client

To add the two management VMware ESXi Hosts to the vDS, complete the following steps:

1. Download Cisco AVS version 5.2(1)SV3(1.25), by going to [Cisco AVS Download](#) and navigating to version 5.2(1)SV3(1.25). Download the CiscoAVS_1.25-5.2.1.SV3.1.25-pkg.zip file, but do not unzip it.
2. Log into the vSphere Web Client as the FlexPod Admin.
3. From the Home screen, select Cisco Virtual Switch Update Manager under Inventories.
4. Under Basic Tasks, select AVS.
5. Under Image Tasks, select Upload.
6. On the right under Upload switch image, click Upload.
7. Click Choose File.
8. Navigate to the CiscoAVS_1.25-5.2.1.SV3.1.25-pkg.zip file, select it and click Open.

The screenshot shows the Cisco Virtual Switch Update Manager interface. The top header is blue with the Cisco logo and the text "Virtual Switch Update Manager". Below this is a grey bar with the text "Virtual Switch Image File Uploader". The main content area is divided into two columns. The left column is titled "Getting started" and contains two paragraphs of text. The right column is titled "File Upload" and contains a "Select file:" section with a "Choose File" button and a file name "CiscoAVS_1...25-pkg.zip". Below this, the file details are displayed: "File Name: CiscoAVS_1.25-5.2.1.SV3.1.25-pkg.zip", "File Size: 50.84MB", and "File Type: Cisco AVS". At the bottom right of the "File Upload" section are two buttons: "Upload" and "Cancel".

Getting started

Virtual Switch Image File Uploader feature provides a provision to dynamically upload switch images. User can upload all required Nexus 1000v and AVS switch version images only once from local file system.

User can continue to do the different operations for Nexus 1000v and AVS switch. Once user uploads the Nexus 1000v switch software version once, it resides in VSUM repository. User can import the switch images from Cisco.com or [here](#).

File Upload

Select file:

Choose File CiscoAVS_1...25-pkg.zip

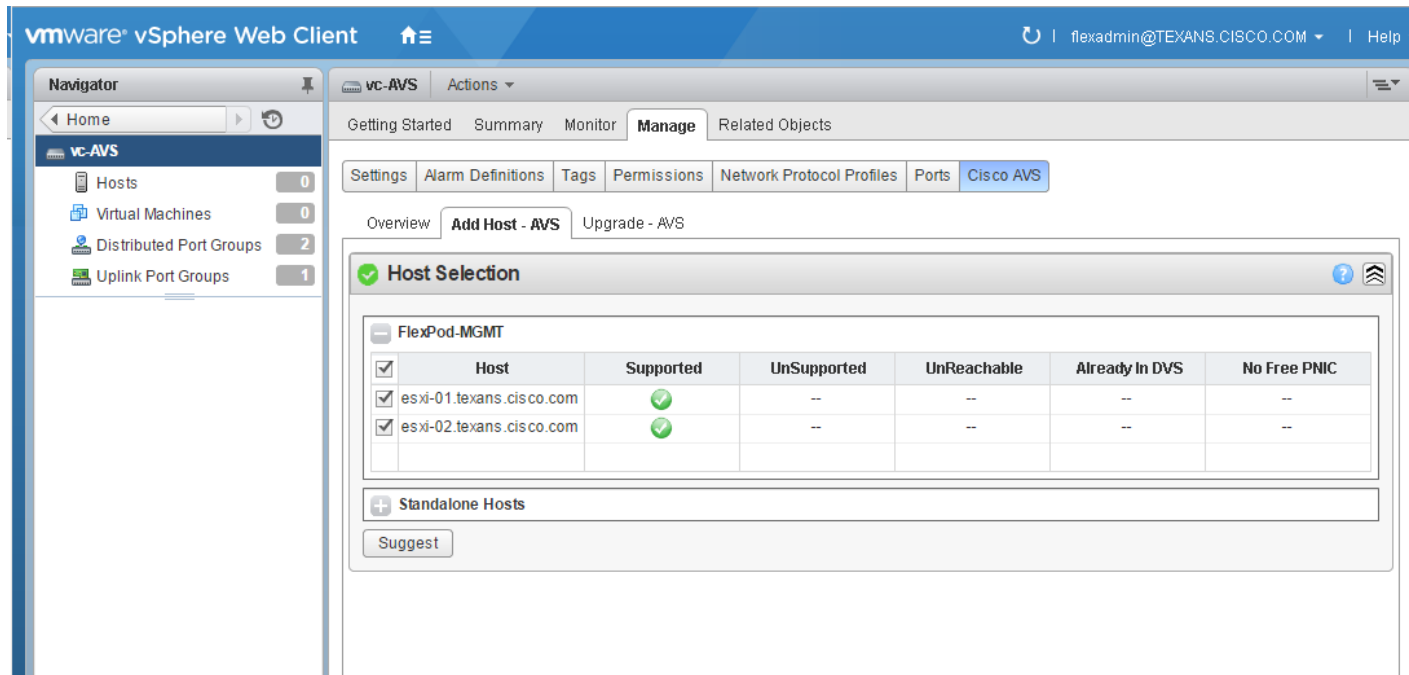
File Name: CiscoAVS_1.25-5.2.1.SV3.1.25-pkg.zip

File Size: 50.84MB

File Type: Cisco AVS

Upload Cancel

9. Click Upload to upload the file.
10. Click OK.
11. Close the Cisco Virtual Switch Update Manager tab in the browser and return to vSphere Web Client.
12. Click Refresh at the lower right. CiscoAVS_1.25 should now appear in the list of Manage Uploaded switch images.
13. On the left, under Basic Tasks, select AVS.
14. Click Configure.
15. On the right, select the FlexPod-DC Datacenter.
16. Select the AVS for the Distributed Virtual Switch and click Manage.
17. In the center pane, under Manage, select the Cisco AVS tab.
18. In the center pane, select the Add Host - AVS tab.
19. Using the drop-down, select the 5.2(1)SV3(1.25) Target Version. Click Show Host.
20. Expand FlexPod-MGMT and select both ESXi hosts.



21. Click Suggest.

22. Under PNIC Selection, select the two vmnics per host set up for AVS. At this point in the deployment, these should be the only vmnics not already assigned.

23. Click Finish to install the Virtual Ethernet Module (VEM) on each host and add the host to the AVS.

24. On the left, select Hosts. The two ESXi hosts should now show up as part of the AVS.



Note: You may need to click refresh to see the newly added hosts.

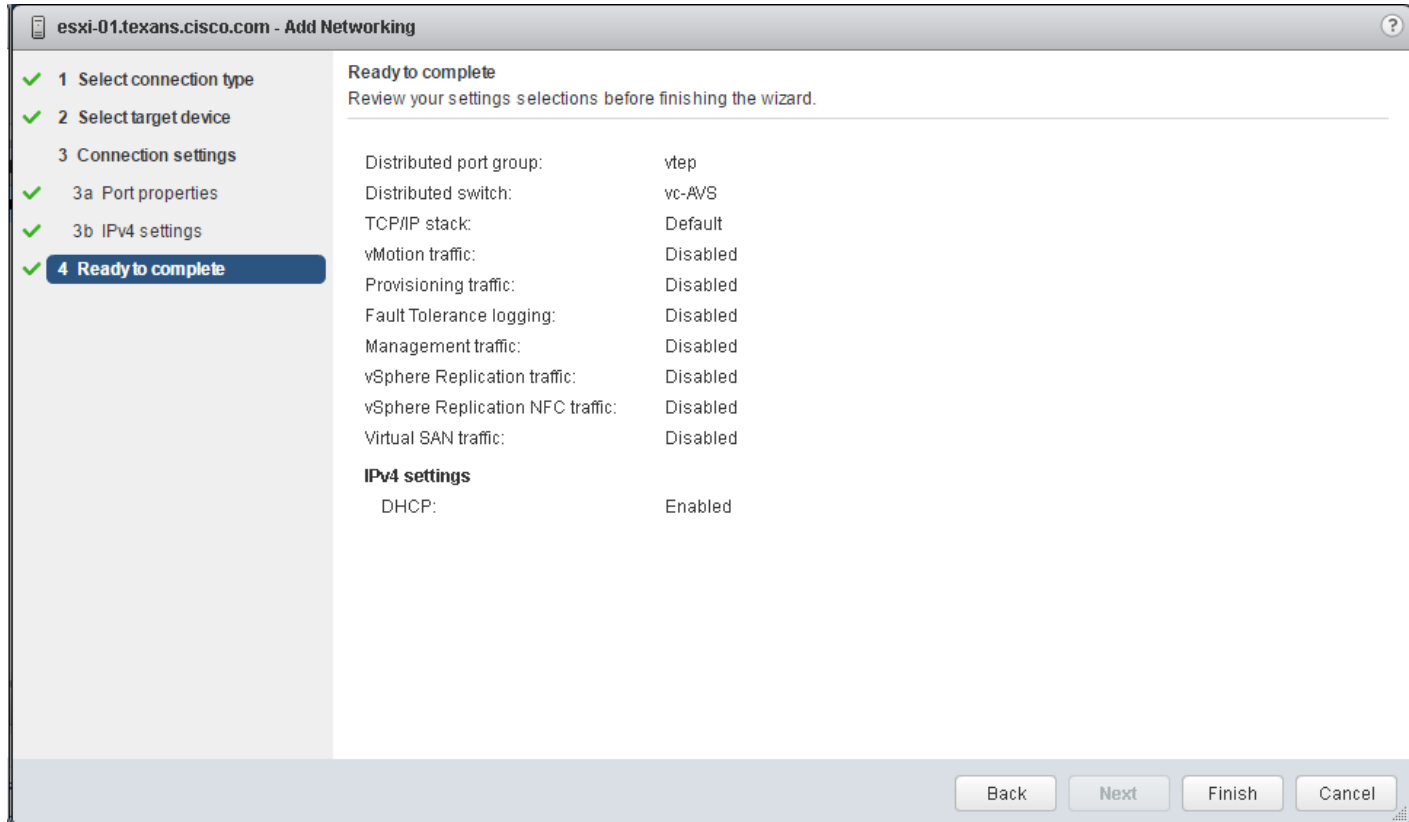
Add Second VXLAN Tunnel Endpoint (VTEP) to Each ESXi Host for Load Balancing

vSphere Web Client

To add a second VTEP to each ESXi host in the Cisco AVS for load balancing, complete the following steps:

1. In the vSphere Web Client, from the Home screen, select Hosts and Clusters.
2. On the left, expand the vCenter, Datacenter, and Cluster. Select the first ESXi host.
3. In the center pane, under the Manage tab, select the Networking tab. Select VMkernel adapters.
4. In the list of VMkernel ports, make sure the vtep VMkernel port has been assigned an IP address in the ACI Fabric system subnet (10.0.0.0/16 by default).
5. Click the icon with the + sign to add a VMkernel port.
6. In the Add Networking window, make sure VMkernel Network Adapter is selected and click Next.
7. Leave Select an existing network selected and click Browse.

8. Select vtep and click OK.
9. Make sure vtep is now in the text box and click Next.
10. Make sure the Default TCP/IP stack is selected. Do not enable any services. Click Next.
11. Leave Obtain IPv4 settings automatically selected and click Next.



12. Click Finish to complete adding the VTEP.
13. Verify that the just added VTEP obtains an IP address in the same subnet as the first VTEP.
14. Repeat this procedure to add a VTEP to the second ESXi host.

Create In-Band Management Port-Profile on AVS

APIC Advanced GUI

To create an In-Band Management VMware Port-Profile on the AVS that has access to both the Core-Services VMs and the mapped in In-Band Management subnet, complete the following steps:

1. In the APIC Advanced GUI, select Tenants > Foundation.
2. On the left, expand Tenant Foundation, then Application Profiles and IB-MGMT.
3. Under IB-MGMT, right-click Application EPGs and select Create Application EPG.
4. Name the EPG IB-MGMT-AVS and make sure Intra EPG Isolation is set to Unenforced.

- Use the Bridge Domain drop-down to select the common/BD-common-Internal Bridge Domain. Select the default Monitoring Policy.

Create Application EPG i X

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation:

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

- Click FINISH to complete creating the EPG.
- Under Application Profile IB-MGMT, select and expand Application EPGs and EPG IB-MGMT-AVS.
- Under EPG IB-MGMT-AVS, right-click Domains (VMs and Bare-Metal) and select Add VMM Domain Association.
- In the Add VMM Domain Association window, select the AVS for the VMM Domain Profile.
- For Deploy and Resolution Immediacy, select Immediate. Do not add a value for Port Encap, a VXLAN will automatically be assigned.

Add VMM Domain Association

Choose the VMM domain to associate

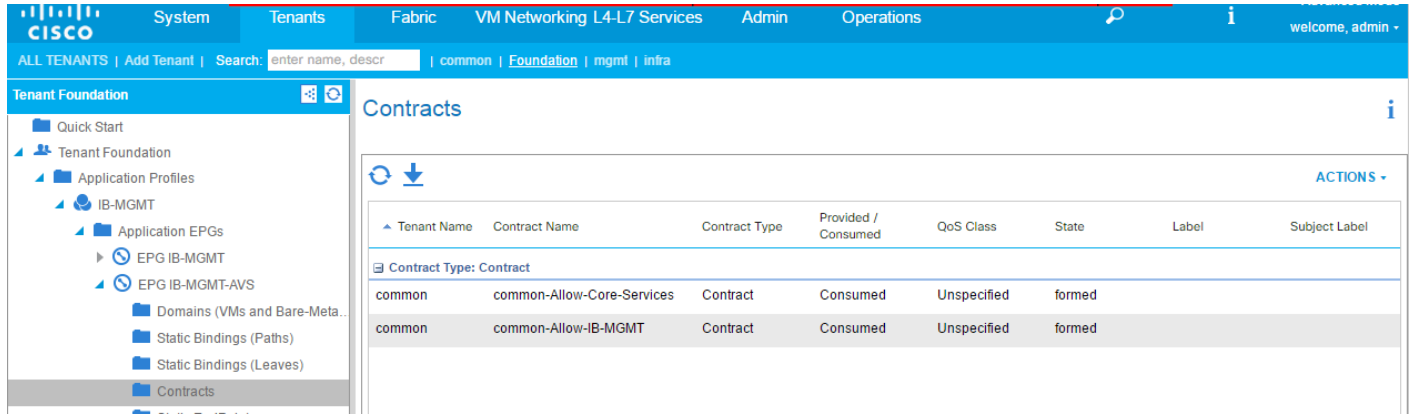
VMM Domain Profile: VMware/vc-AVS

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

Port Encap: _____
For example, vlan-1

11. Click SUBMIT to complete the VMM Domain Association.
12. Under EPG IB-MGMT-AVS, right-click Contracts and select Add Consumed Contract.
13. In the Add Consumed Contract window, select the common/common-Allow-IB-MGMT contract.
14. Click SUBMIT to complete adding the consumed contract.
15. Under EPG IB-MGMT-AVS, right-click Contracts and select Add Consumed Contract.
16. In the Add Consumed Contract window, select the common/common-Allow-Core-Services contract.
17. Click SUBMIT to complete adding the consumed contract.

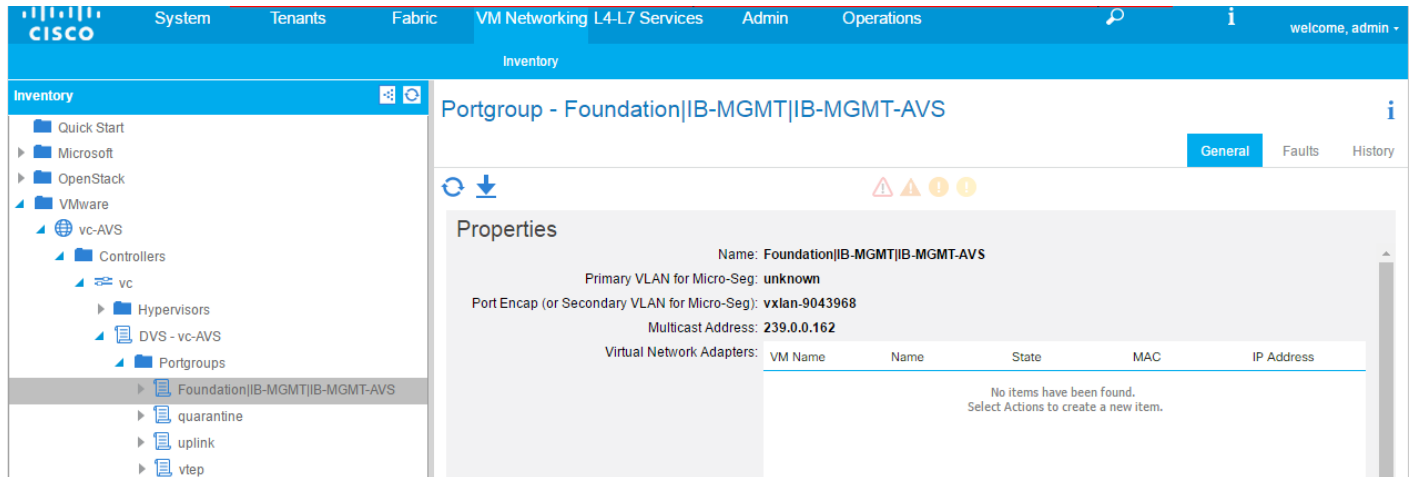


18. At the top, select VM Networking.

19. On the left, expand VMware, the AVS, Controllers, the vCenter, the DVS - AVS, and Portgroups.

20. Select the Foundation|IB-MGMT|IB-MGMT-AVS Portgroup.

21. Verify the assigned VXLAN and Multicast Address under Properties.



Note: The IB-MGMT-AVS port profile can now be assigned to a VM.

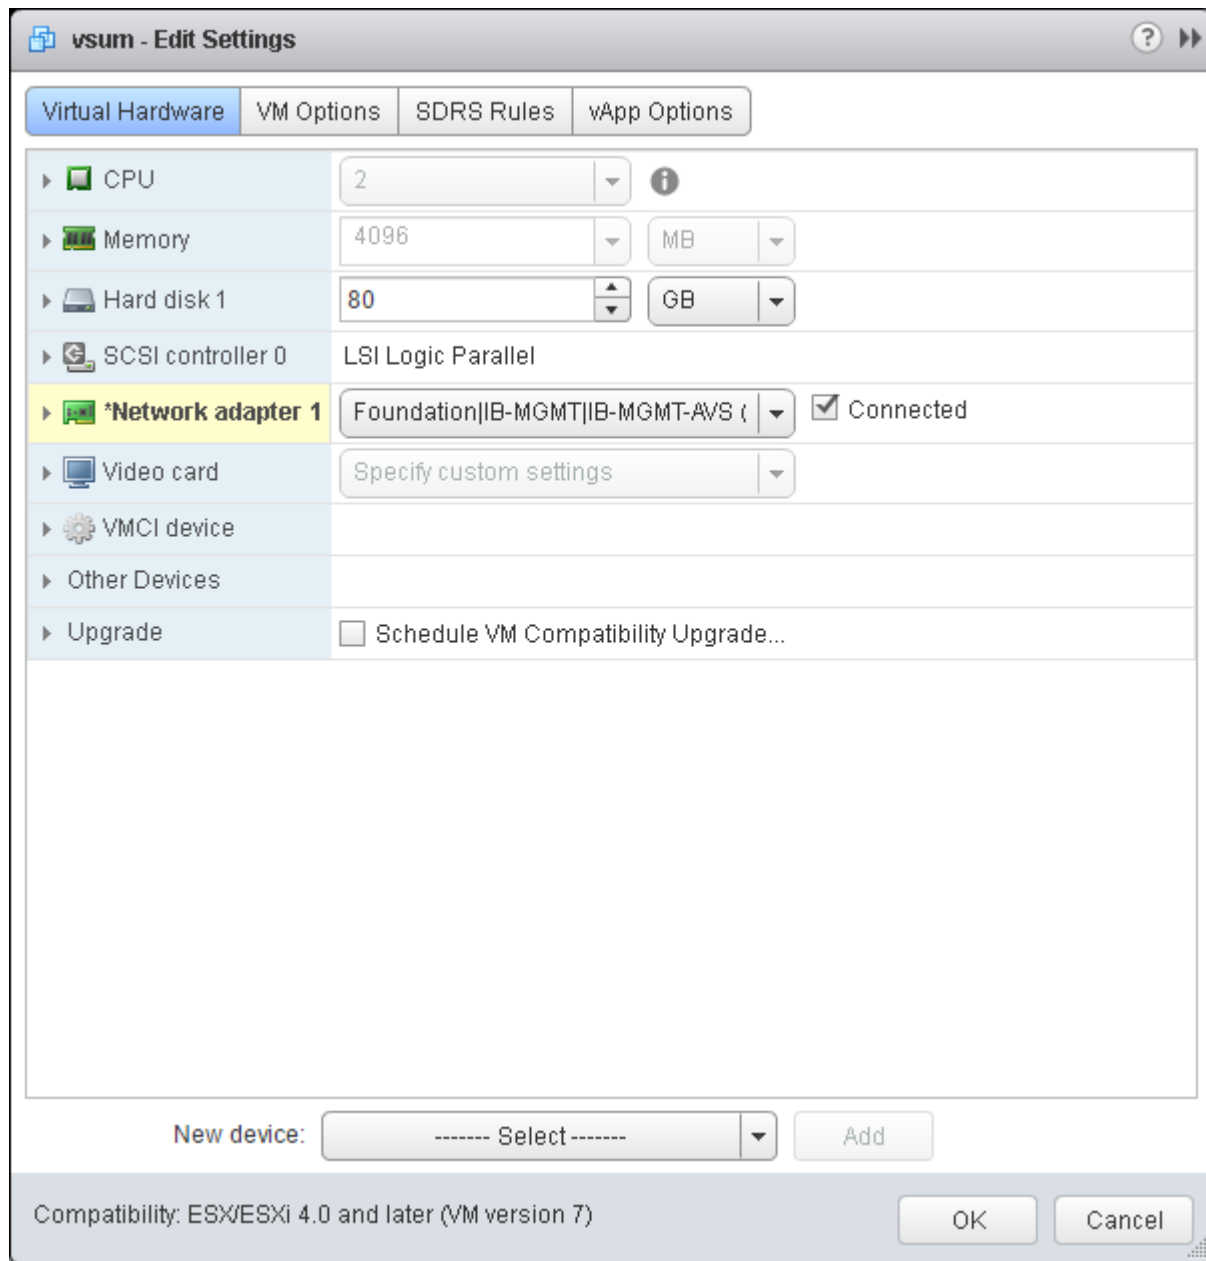
22. Select the vtep Portgroup.

23. Verify that the assigned VLAN matches the ACI Fabric System VLAN (4093 in this validation). Also, verify the four VTEPs have assigned IP addresses in the ACI Fabric System Subnet (10.0.0.0/16 by default).

The screenshot shows the Cisco Inventory interface for a portgroup named 'vtep'. The left sidebar shows the navigation tree with 'vtep' selected under 'Portgroups'. The main content area displays the 'Properties' for 'vtep', including its name, primary VLAN (unknown), port encapsulation (vlan-4093), and a table of management network adapters.

Server Name	Name	State	MAC	IP Address
esxi-01.texans.c...	vmk6	Up	00:50:56:60:CE:...	10.0.216.125
esxi-02.texans.c...	vmk6	Up	00:50:56:68:13:67	10.0.104.95
esxi-02.texans.c...	vmk5	Up	00:50:56:60:65:16	10.0.216.126
esxi-01.texans.c...	vmk5	Up	00:50:56:6B:E2:...	10.0.216.127

24. In the vSphere Web Client, from the Home screen, select Hosts and Clusters under Inventories.
25. On the left, expand the vCenter, the Datacenter, and the ESXi Cluster and select the VSUM VM.
26. Right-click the VSUM VM and select Edit Settings.
27. For Network adapter 1, use the drop-down to select the Foundation|IB-MGMT|IB-MGMT-AVS port group.



28. Click OK to complete the VM network adapter change.



Note: By placing the VSUM network adapter in this port group, it can reach the Core-Services VMs and the IB-MGMT bridged L2 network, but it is not a Core-Services VM and it is not necessary for tenants to reach it.



Note: If you install both the VMware vDS and Cisco AVS in your FlexPod, you will need to put contracts between EPGs from the two virtual switches in order to connect VMs from the port groups in the two virtual switches. Note also that in this case, one IB-MGMT EPG can be setup and bound to both VMM domains. In this case, VMs with interfaces assigned in either of both port groups will all be in the same EPG and will

have unrestricted communication. Even though the two port groups have the same name, these duplicate names are not a problem in VMware vCenter.

FlexPod Management Tools Setup

NetApp Virtual Storage Console 6.2P2 Deployment Procedure

This section describes the deployment procedures for the NetApp VSC.

Virtual Storage Console 6.2 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.3.2:

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

Install Virtual Storage Console 6.2P2

To install the VSC 6.2P2 software, complete the following steps:

1. Build a VSC VM with Windows Server 2012 R2, 4GB of RAM, two CPUs, and one virtual network interface in the `Core-Services Network` port group. The virtual network interface should be a VMXNET 3 adapter.
2. Bring up the VM, install VMware Tools, assign the IP address and gateway in the IB-MGMT subnet, and join the machine to the Active Directory domain.
3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the User Interfaces and Infrastructure Feature on the VM.
4. Install all Windows updates on the VM.
5. Log in to the VSC VM as the FlexPod Admin user using the VMware console.
6. From the VMware console on the VSC VM, download the x64 version of [Virtual Storage Console 6.2P2](#) from the [NetApp Support](#) site.
7. Right-click the `vsc-6.2P2-win64.exe` file downloaded in step 5 and select Run as Administrator.
8. Select the appropriate language and click OK.
9. On the Installation wizard Welcome page, click Next.
10. Select the checkbox to accept the message and click Next.



Note: The Backup and Recovery capability requires an additional license.

11. Click Next to accept the default installation location.



12. Click Install.



13. Click Finish.

Register Virtual Storage Console with vCenter Server

To register the VSC with the vCenter Server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open <https://localhost:8143/Register.html> in Internet Explorer.
2. Click Continue to This Website (Not Recommended).
3. In the Plug-in Service Information section, select the local IP address of the VSC VM.
4. In the vCenter Server Information section, enter the host name or IP address, the user name (FlexPod admin user or root), and the user password for the vCenter Server. Click Register to complete the registration.

vSphere Plugin Registration

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

Plugin service information	
Host name or IP Address:	<input type="text" value="172.26.163.174"/> ▼
vCenter Server information	
Host name or IP Address:	<input type="text" value="172.26.163.175"/>
Port:	<input type="text" value="443"/>
User name:	<input type="text" value="flexadmin@texans.cisco.com"/>
User password:	<input type="password" value="••••••••"/> 🔒

Register

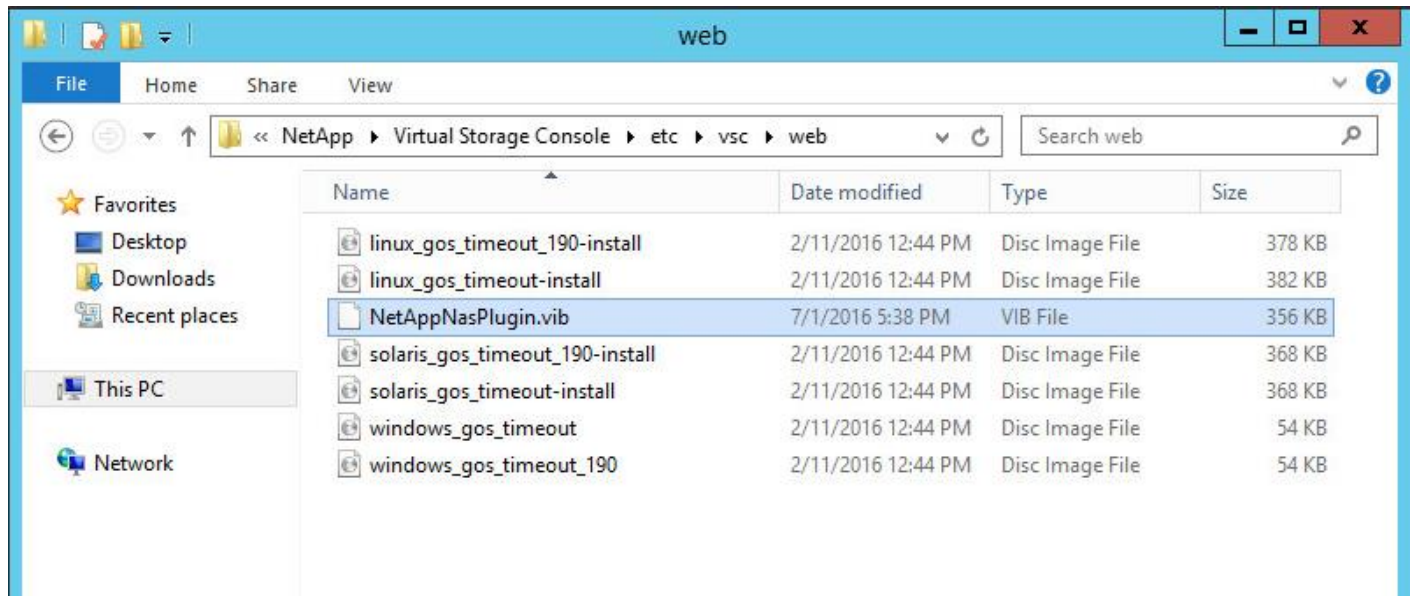
5. Upon successful registration, storage controller discovery begins automatically.

Install NetApp NFS VAAI Plug-in

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

1. Onto the VSC VM, download the NetApp NFS Plug-in 1.1.0 for VMware .vib file from the [NFS Plugin Download](#).
2. Rename the downloaded file NetAppNasPlugin.vib.

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



Discover and Add Storage Resources

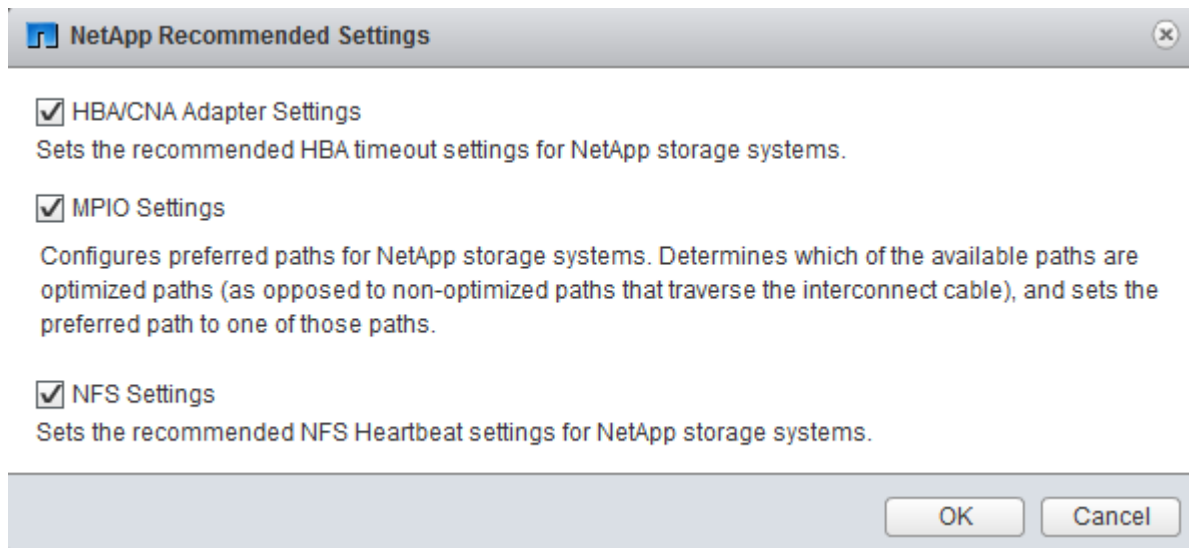
To discover storage resources for the Monitoring and Host Configuration capability and the Provisioning and Cloning capability, complete the following steps:

1. Using the vSphere web client, log in to the vCenter Server as the FlexPod admin user. If the vSphere web client was previously opened, close it and then reopen it.
2. In the Home screen, click the Home tab and click Virtual Storage Console.
3. Select Storage Systems. Under the Objects tab, click Actions > Modify.
4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name and the admin password for password. Confirm that Use SSL to Connect to This Storage System is selected. Click OK.
5. Click OK to accept the controller privileges.
6. Wait for the Storage Systems to update. You may need to click Refresh to complete this update.

Optimal Storage Settings for ESXi Hosts

VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

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

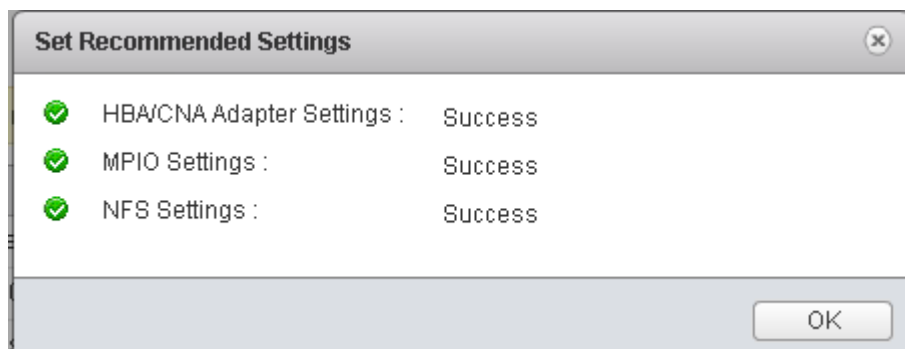


2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings.



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

3. Click OK.



4. From the Home screen in the vSphere Web Client, select Virtual Storage Console.
5. On the left under Virtual Storage Console, select NFS VAAI Tools.
6. Make sure that NFS Plug-in for VMware VAAI Version 1.1.0-0 is shown.
7. Click Install on Host.
8. Select both ESXi hosts and click Install.
9. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.

Virtual Storage Console 6.2P2 Backup and Recovery

Prerequisites for Use of Backup and Recovery Capability

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

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

Backup and Recovery Configuration

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

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



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

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

5. Select any options to include in the backup.

Schedule Backup

✓ 1 Details
✓ 2 Options
3 Spanned Entities
4 Scripts
5 Schedule and Retention
6 Credentials and Alerts
7 Summary

Select the options you want to include along with this backup job.

- Initiate SnapVault update
- Initiate SnapMirror update
- Perform VMware consistency snapshot
- Include datastores with independent disks

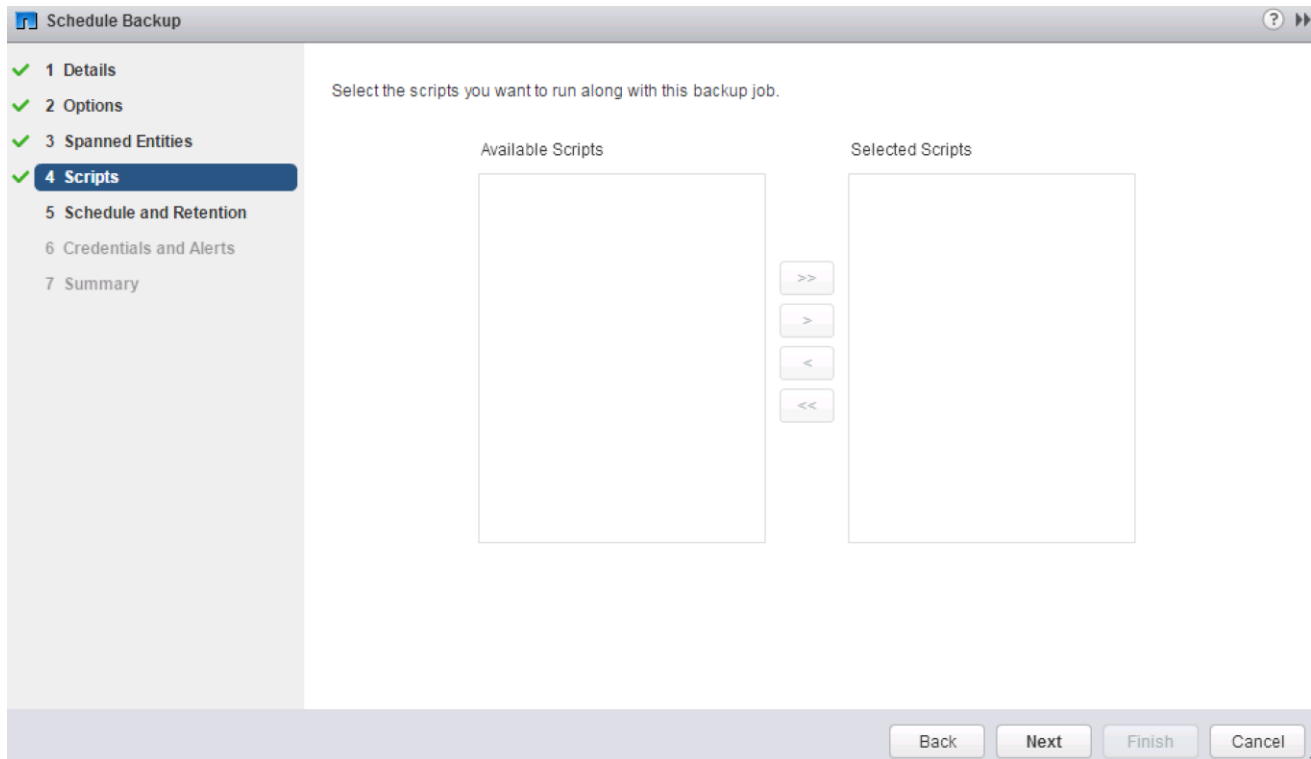
SnapVault integration in VSC is supported for Clustered Data ONTAP 8.2 or higher.

Back Next Finish Cancel

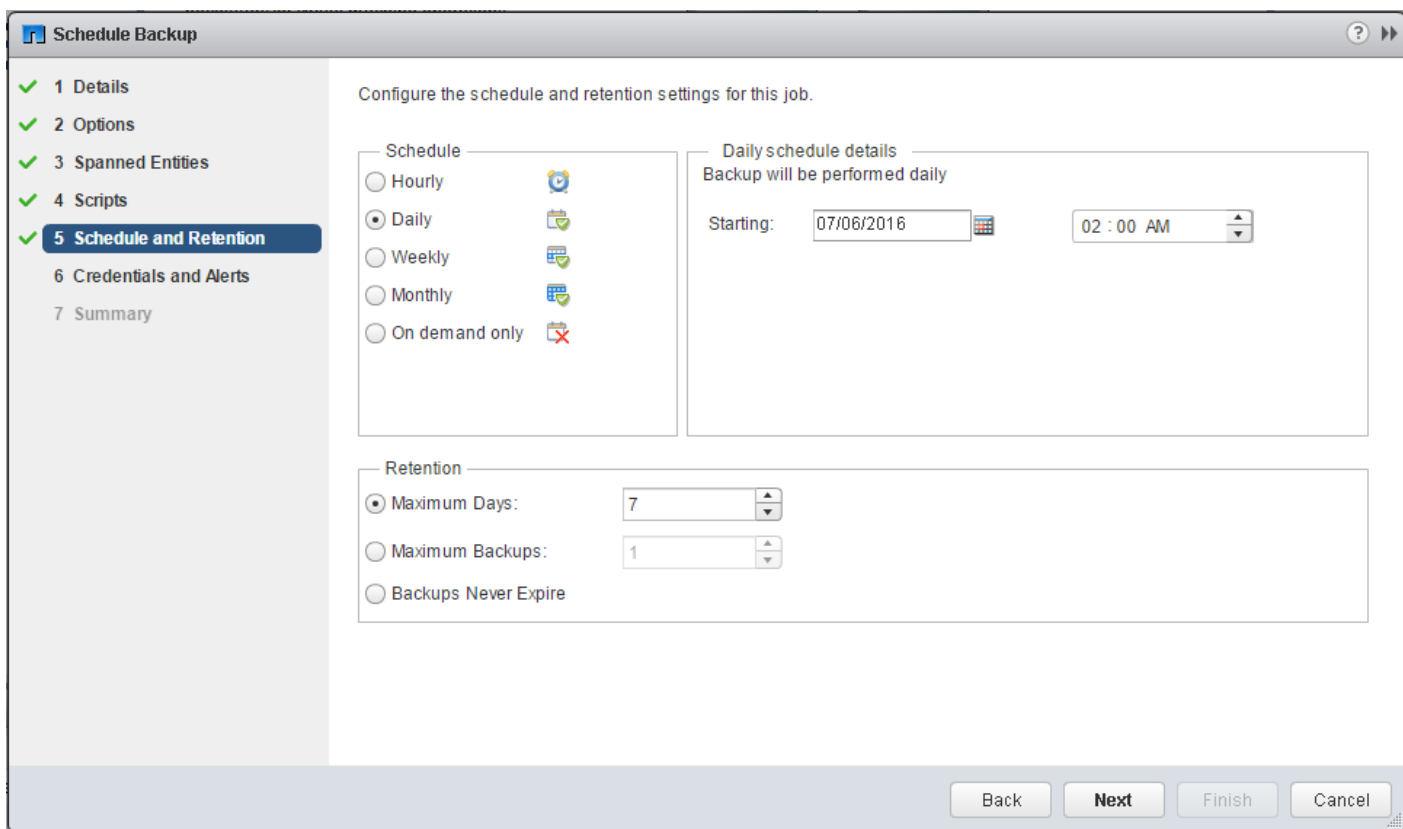


Note: For consistent VM snapshots, select Perform VMware consistency snapshot to make a VMware snapshot of each VM just before the NetApp snapshot is taken. The VMware snapshots is then deleted after the NetApp snapshot is taken.

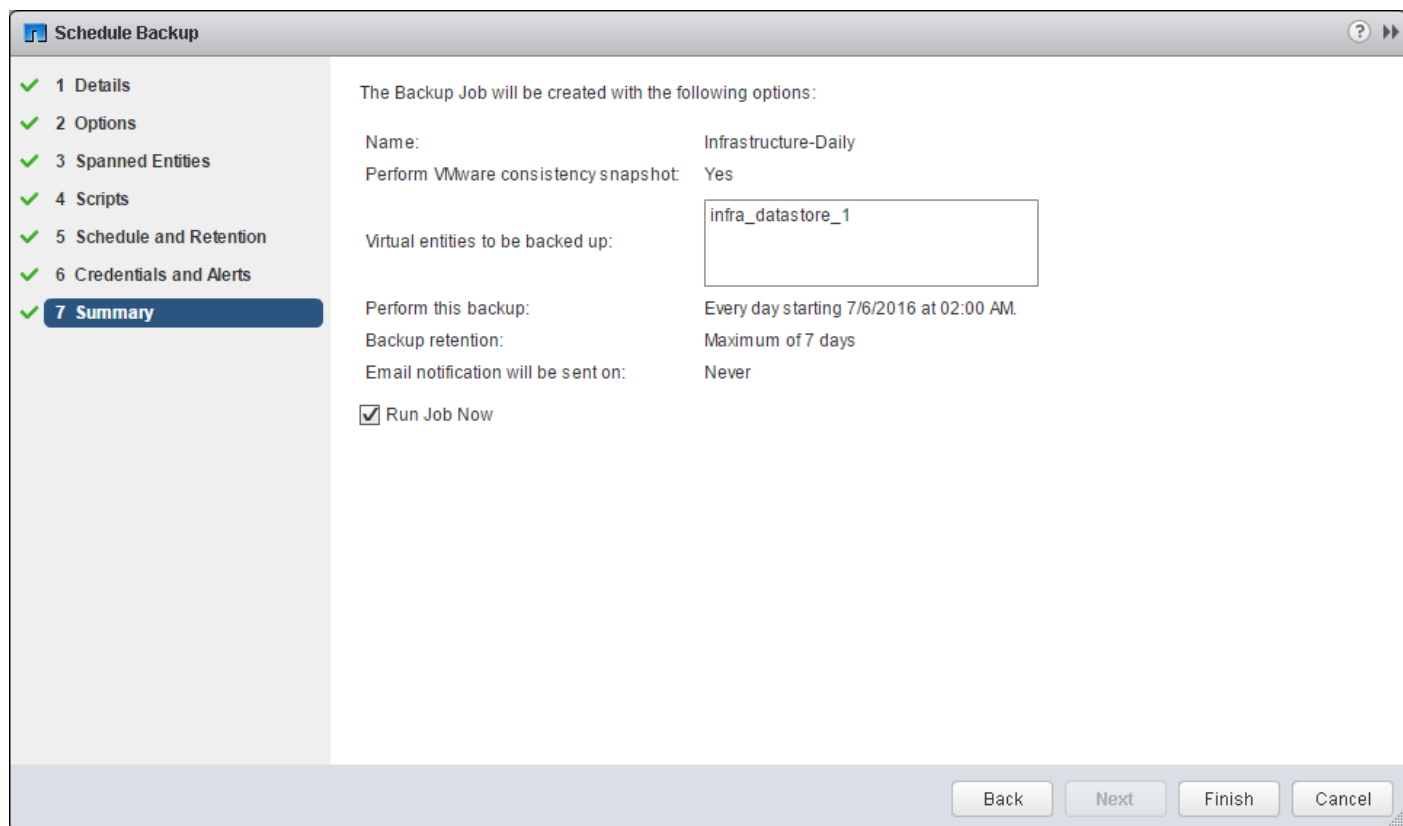
6. Click Next on the Options screen.
7. Click Next on the Spanned Entities screen.
8. Select one or more backup scripts if available, and click Next in the Scripts screen.



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



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



The screenshot shows the 'Schedule Backup' wizard in a summary view. On the left, a navigation pane lists seven steps: 1 Details, 2 Options, 3 Spanned Entities, 4 Scripts, 5 Schedule and Retention, 6 Credentials and Alerts, and 7 Summary (which is highlighted). The main area displays the following configuration:


The Backup Job will be created with the following options:

Name:	Infrastructure-Daily
Perform VMware consistency snapshot:	Yes
Virtual entities to be backed up:	infra_datastore_1
Perform this backup:	Every day starting 7/6/2016 at 02:00 AM.
Backup retention:	Maximum of 7 days
Email notification will be sent on:	Never

Run Job Now

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

12. Review the summary page and click Finish. If you want to run the job immediately, select the Run Job Now option and then click Finish.
13. Click OK.



14. You can also create other backup jobs with overlapping schedules. For example, you can create weekly or monthly backups that overlay daily backups.
15. On the storage cluster interface, automatic Snapshot copies of the volume can now be disabled because NetApp VSC is now handling scheduled backups. To do so, enter the following command:

```
volume modify -vserver Infra-SVM -volume infra_datastore_1 -snapshot-policy none
```

16. Also, to delete any existing automatic Snapshot copies that have been created on the volume, enter the following command:

```
volume snapshot show -vserver Infra-SVM -volume infra_datastore_1
```

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



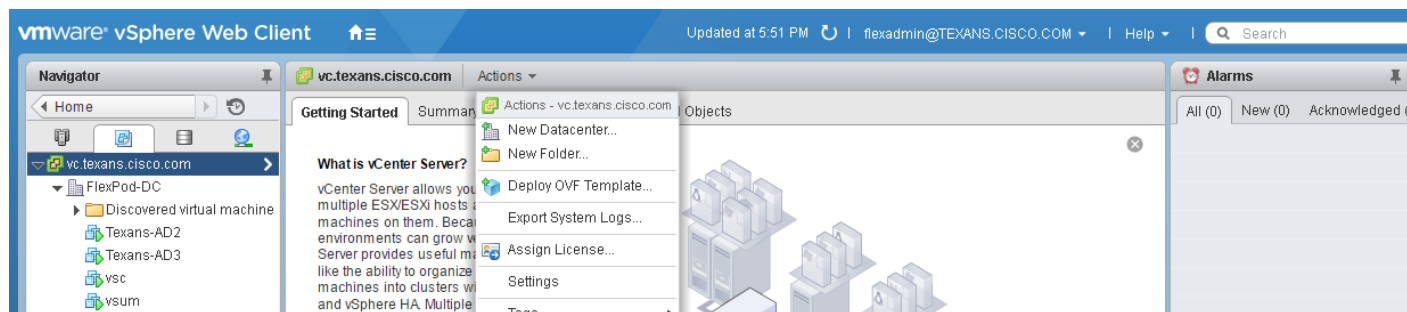
Note: The wildcard character * can be used in snapshot names in the previous command.

OnCommand Performance Manager 2.1

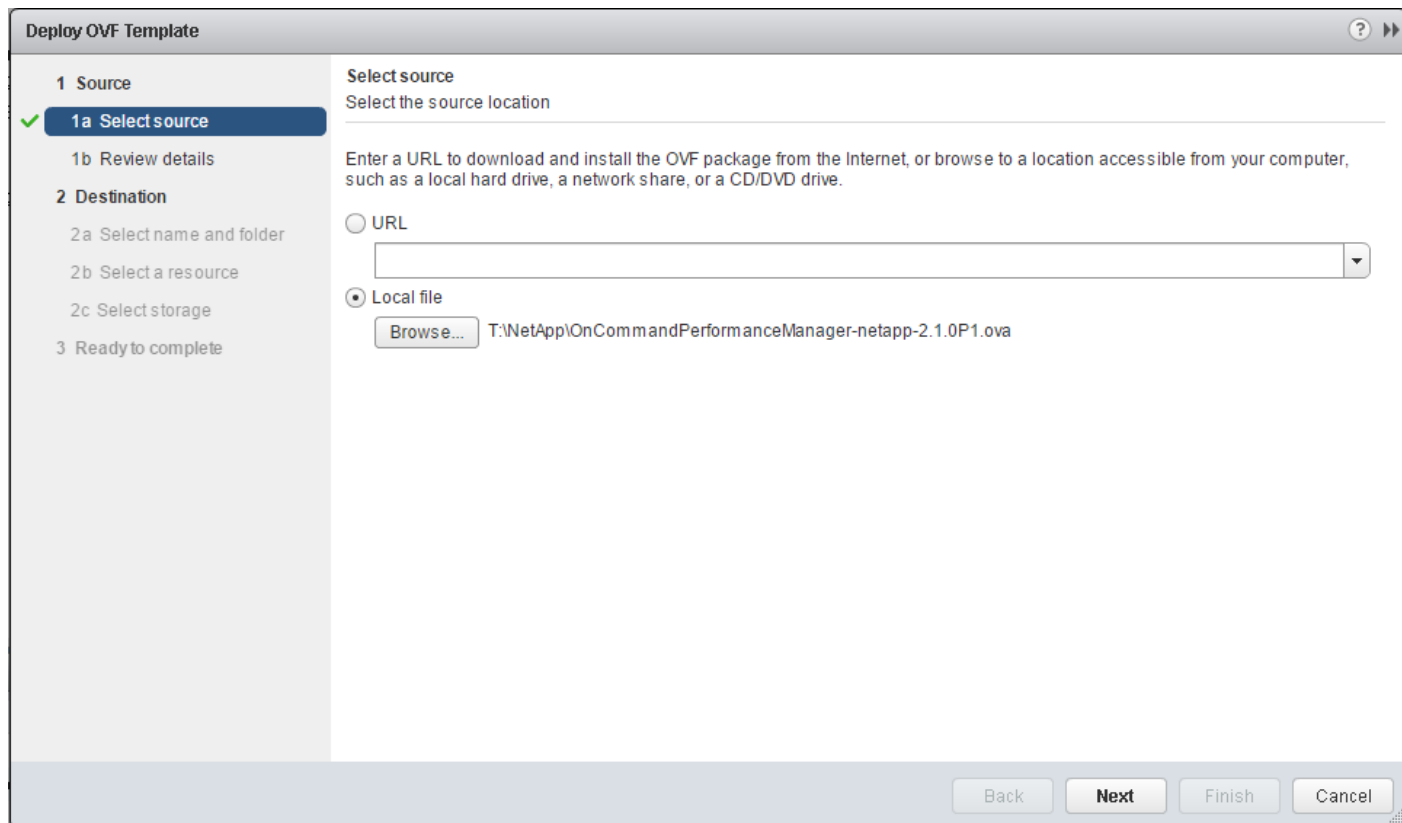
OnCommand Performance Manager Open Virtualization Format (OVF) Deployment

To install the OnCommand Performance Manager, complete the following steps:

1. Download and review the [OnCommand Performance Manager 2.1 Installation and Administration Guide for VMware Virtual Appliances](#).
2. Download OnCommand Performance Manager version 2.1 ([OnCommandPerformanceManager-netapp-2.1.0P1.ova](#)). Click Continue at the bottom of the page and follow the prompts to complete the installation.
3. Log in to the vSphere Web Client. Go to Home > VMs and Templates.
4. At the top of the center pane, click Actions > Deploy OVF Template.

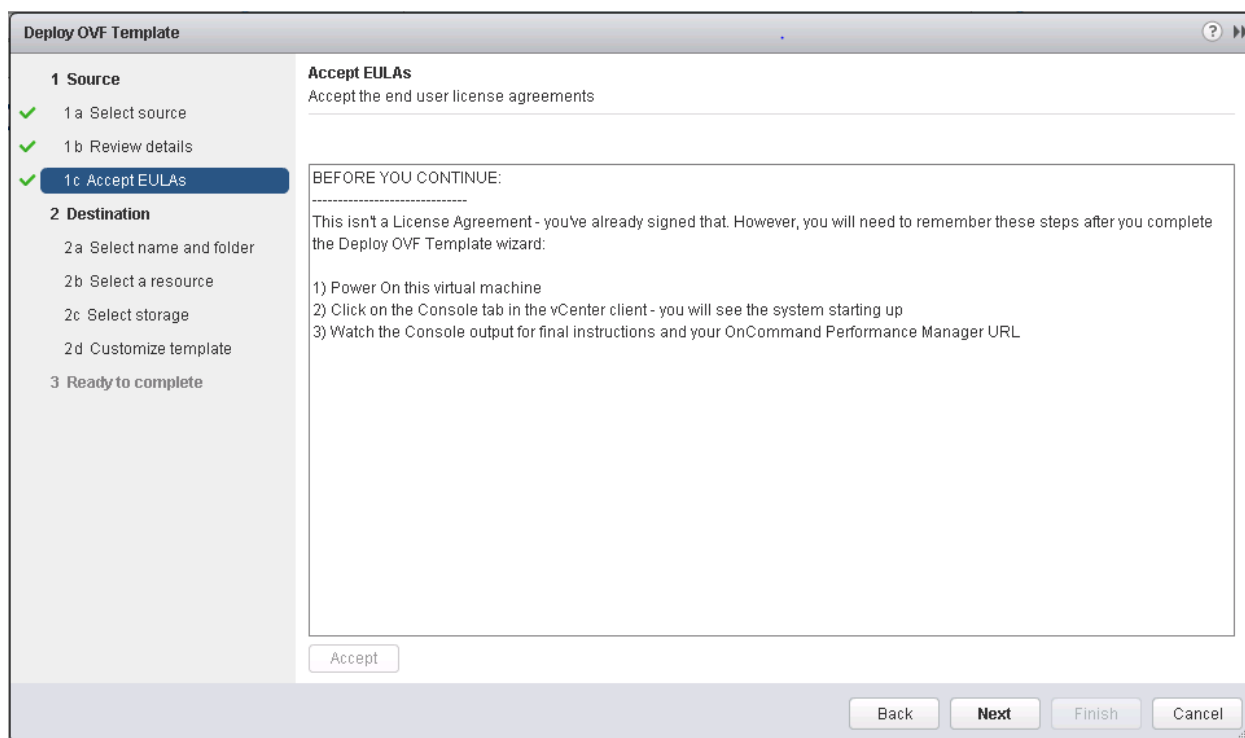


5. Browse to the OnCommandPerformanceManager-netapp-2.1.0P1.ova file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.



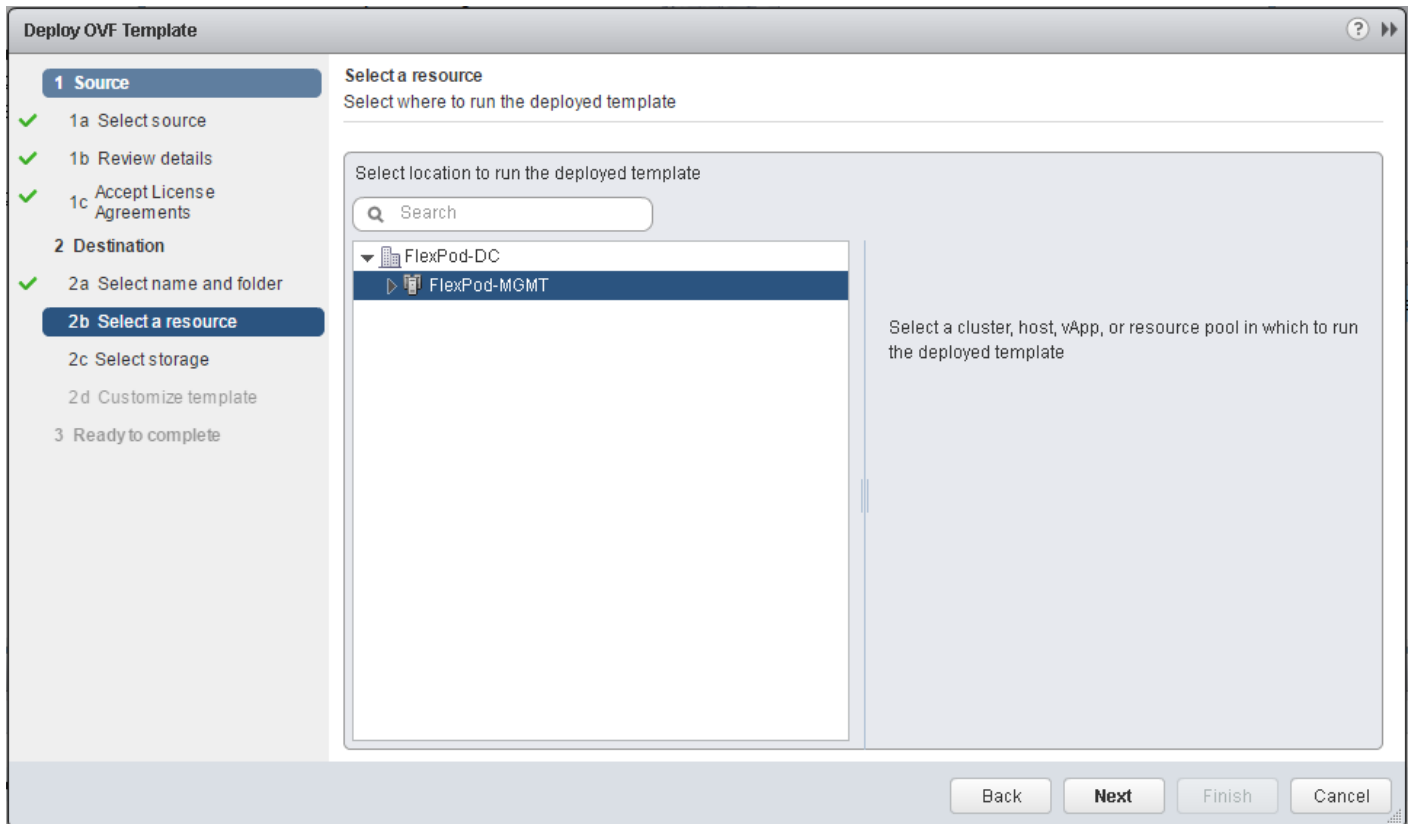
6. Review the details and click Next.

7. Click Accept to accept the agreement. Click Next.

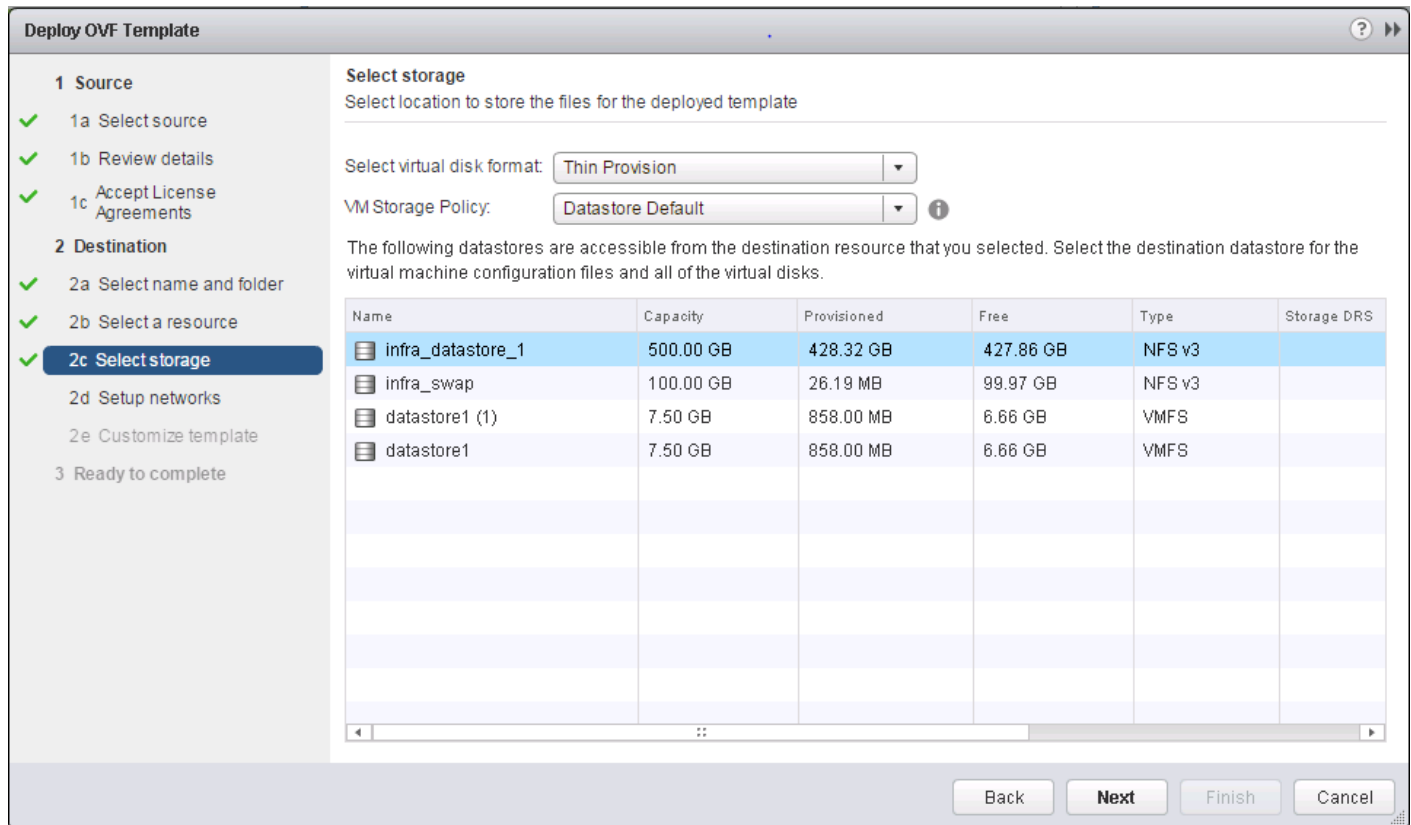


8. Enter the name of the VM and select the FlexPod-DC folder to hold the VM. Click Next.

9. Select FlexPod-MGMT within the FlexPod-DC datacenter as the destination compute resource pool to host the VM. Click Next to continue.



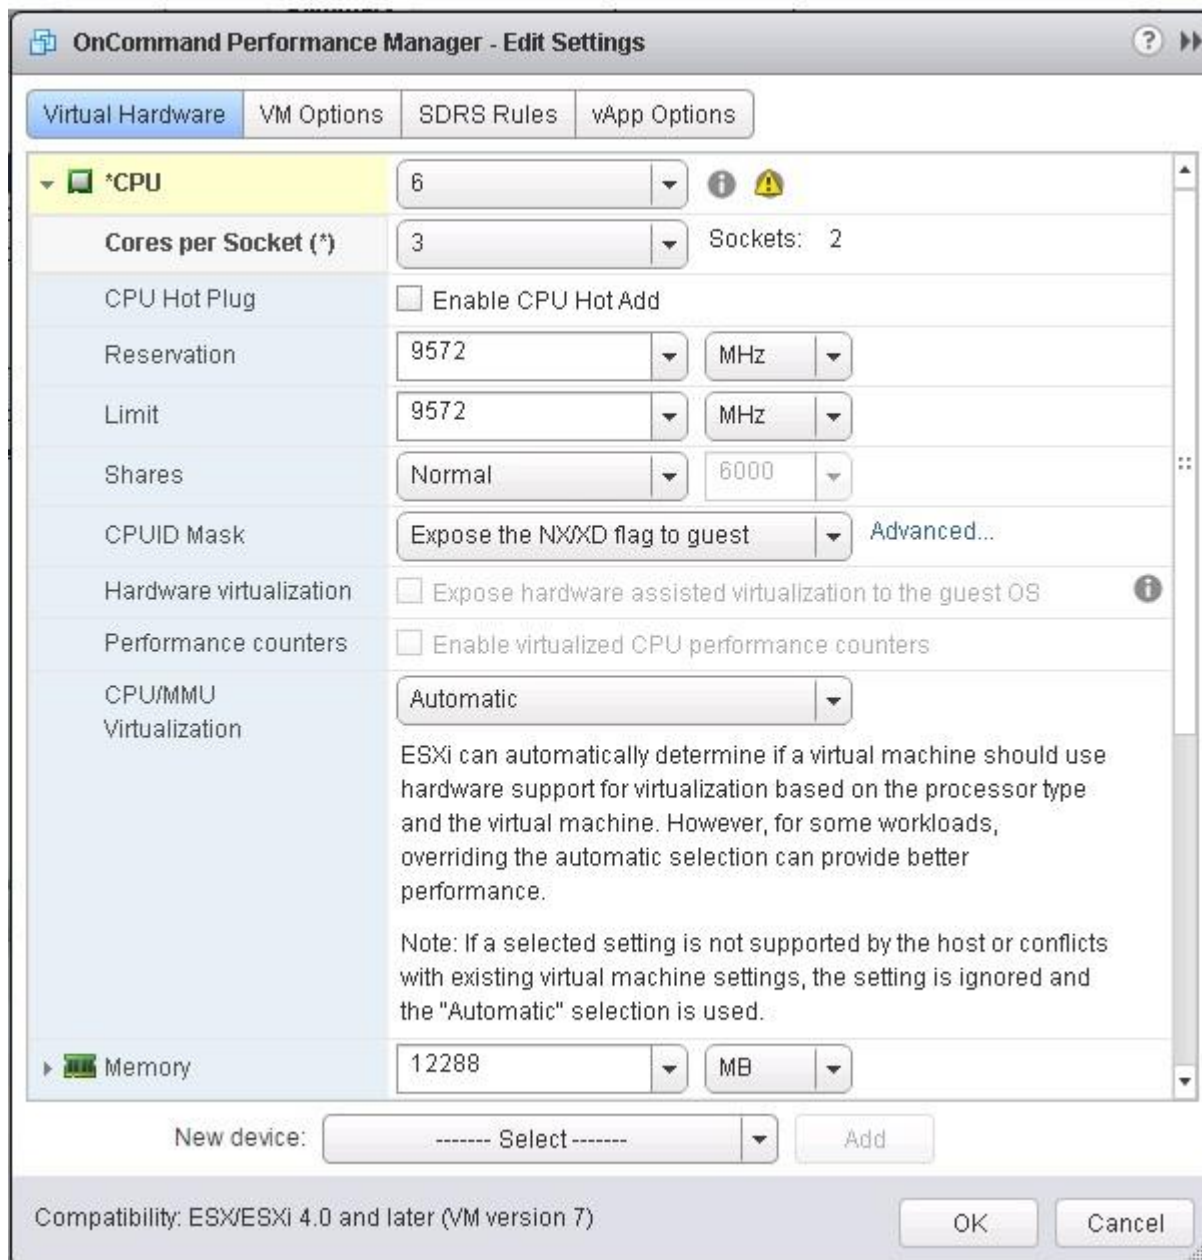
10. Select `infra_datastore_1` as the storage target for the VM and select Thin Provision as the virtual disk format. Click Next.



11. Select Foundation|IB-MGMT|IB-MGMT OR Foundation|IB-MGMT|IB-MGMT-AVS as the destination network to the nat source network. It is not necessary for this VM to be in the Core-Services Network. Click Next.
12. Do not enable DHCP. Fill out the details for the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS. Click Next to continue.
13. Deselect Power On After Deployment.
14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.
15. In the left pane, navigate to Home > Hosts and Clusters. Expand the FlexPod_Management cluster and select the newly created OnCommand Performance Manager VM. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.
16. Expand the CPU options, and complete the following steps:
 - a. The minimum required CPU Reservation is 9572MHz. Determine the CPU frequency of the host server.
 - b. Set the number of CPUs to the number of CPUs required (9572 / the CPU Frequency of the host rounded up to the next even number).
 - c. Set the number of Cores per Socket where the Sockets number on the right matches the number of CPU sockets in the host. For example, if a host has two CPUs operating at a

speed of 1999MHz, then the VM needs six virtual CPUs ($9572 / 1999 = 4.79$ - rounded to 6 virtual CPUs). If the host has two physical CPU sockets, then allocate three Cores per Socket.

Use the [OnCommand Performance Manager 2.1 Installation and Administration Guide for VMware Virtual Appliances](#) for guidance on these settings.



17. Click OK to accept the changes.

18. Right-click the VM in the left-hand pane. Click Power On.

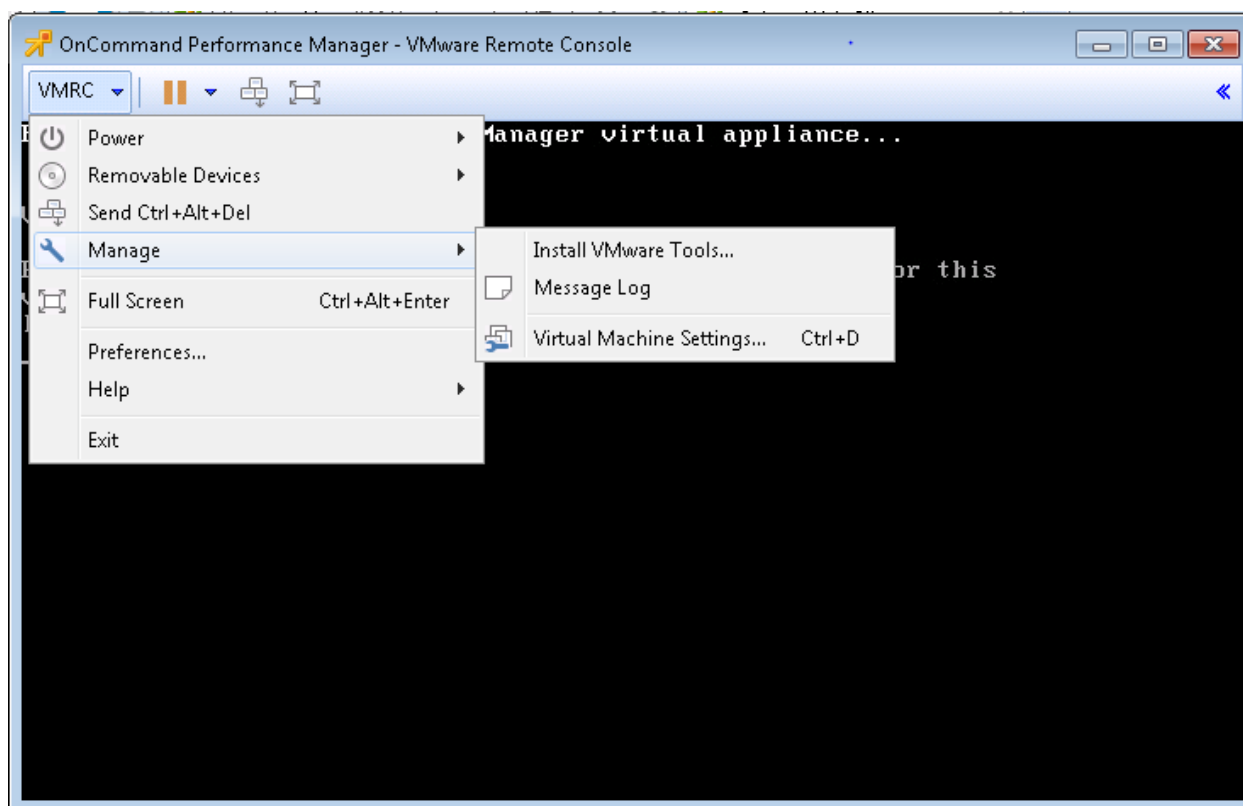
OnCommand Performance Manager Basic Setup

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.



Note: It may be necessary to download and install the Remote Console at this point.

2. After VMware Tools Installation comes up in the VMware Remote Console window, select VMRC > Manage > Install VMware Tools. VMware Tools installs in the VM.



3. Set up OnCommand Performance Manager by answering the following questions in the console window:

Geographic area: <<Enter your geographic location>>

Time zone: <<Select the city or region corresponding to your time zone>>

These commands complete the network configuration checks, generate SSL certificates, and start the OnCommand Performance Manager services.

4. To Create a Maintenance User account, run the following commands:



Note: The maintenance user manages and maintains the settings on the OnCommand Performance Manager virtual appliance.

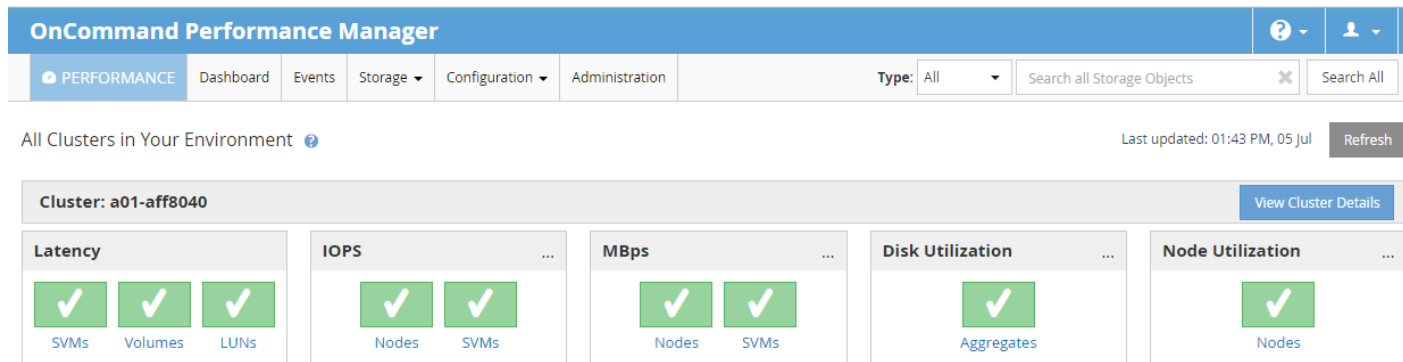
```
Username : admin
```

```
Enter new UNIX password: <password>
```

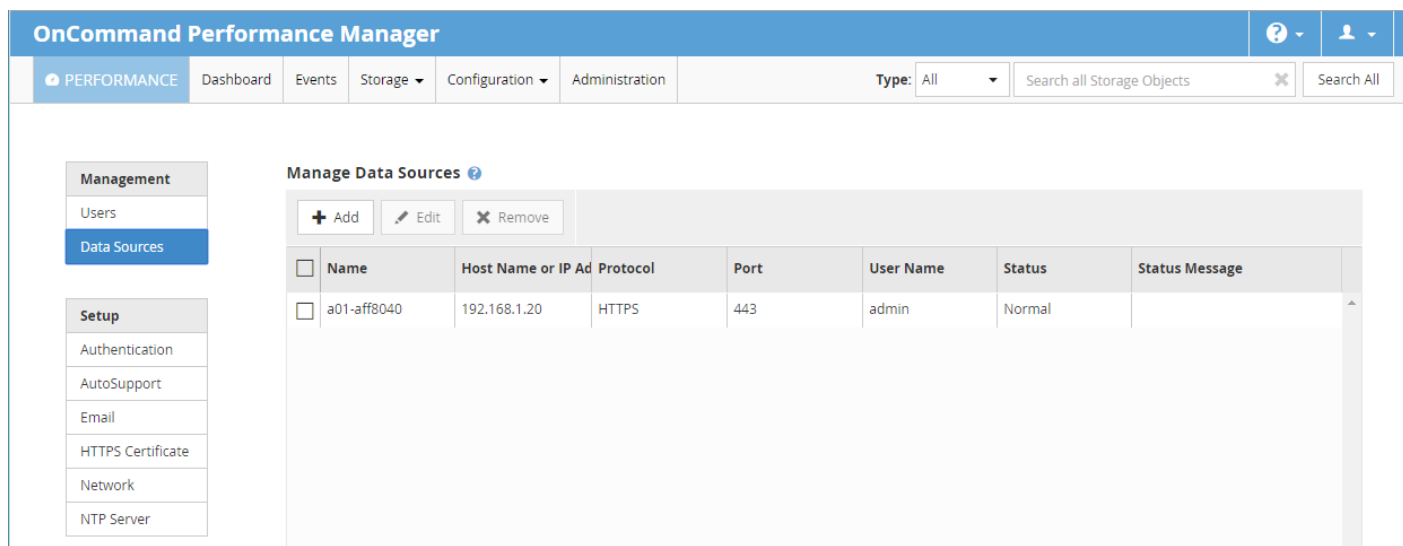
```
Retype new UNIX password: <password>
```

5. With a web browser, navigate to the OnCommand Performance Manager using the URL `https://<oncommand-pm-ip>`.

6. Log in using the Maintenance User account (admin) credentials.
7. Enter a Maintenance User e-mail address, SMTP mail server information, and the NTP server IP address. Click Save.
8. Select **Yes** to enable AutoSupport capabilities. Click Save.
9. Click Save to not change the admin password.
10. Enter the storage cluster host name or IP address, the storage cluster admin user name, and the storage cluster admin password. Click Add Cluster, then click Save to complete setup. It may take up to 15 minutes for the cluster to be visible in OnCommand Performance Manager.



11. After the cluster is added it can be accessed by clicking on Administration > Manage Data Sources.

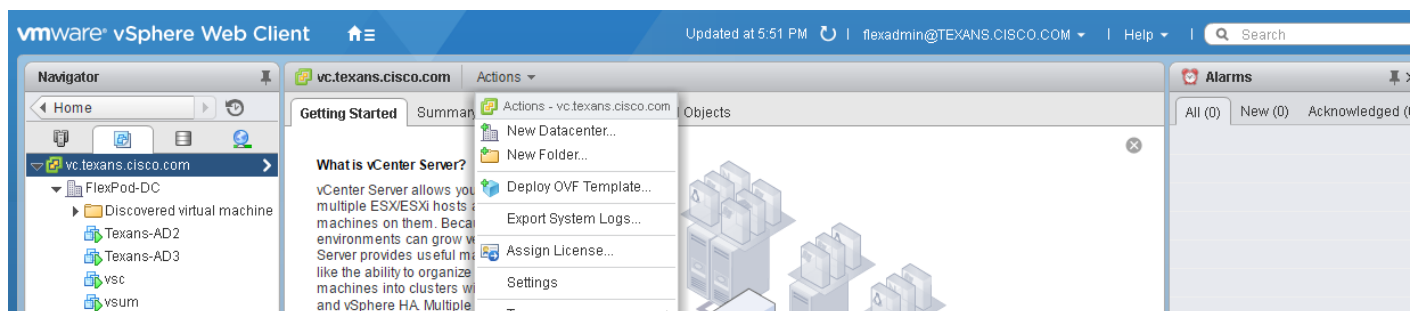


OnCommand Unified Manager 6.4

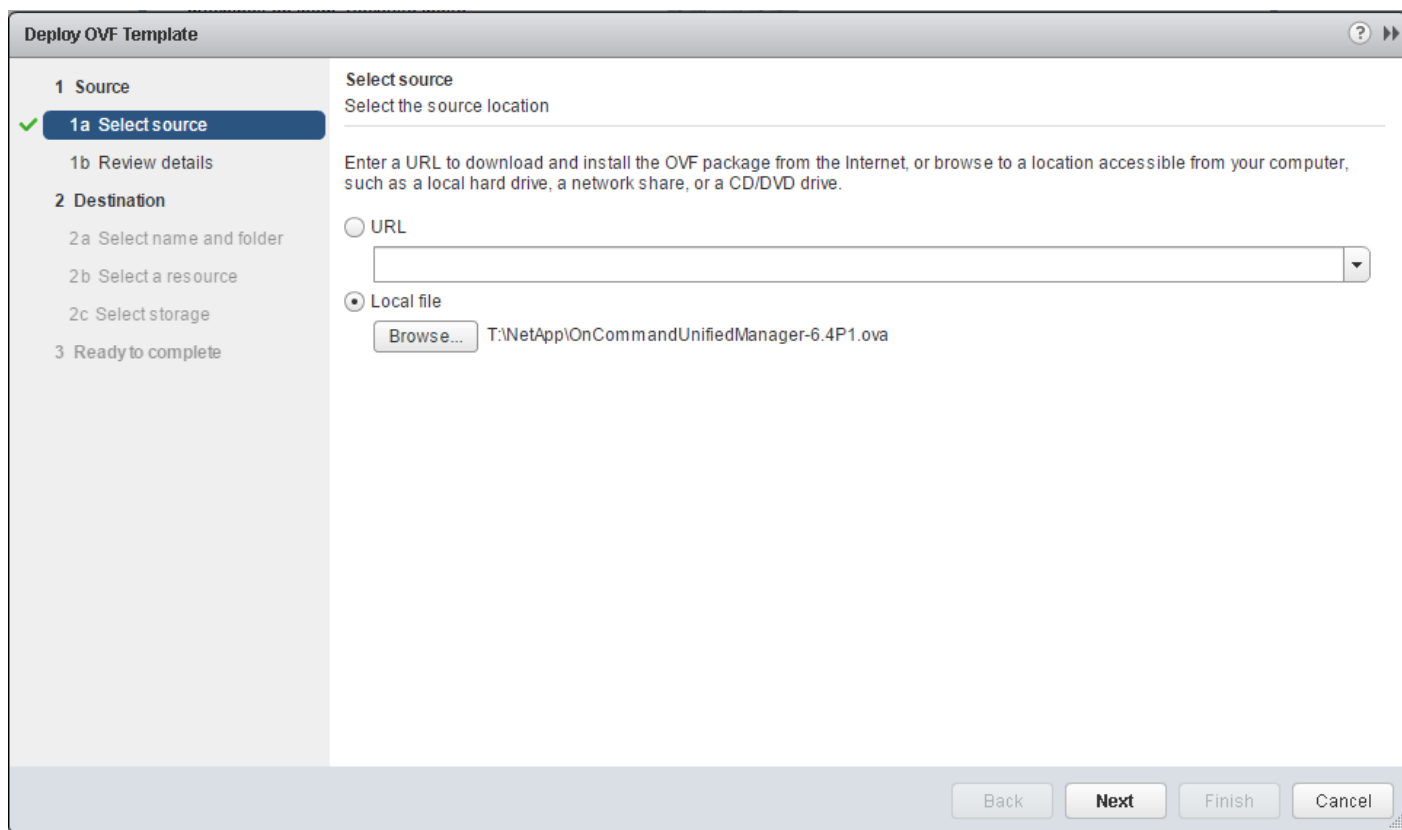
OnCommand Unified Manager OVF Deployment

To install the OnCommand Unified Manager, complete the following steps:

1. Download and review the [OnCommand Unified Manager Installation and Setup Guide](#).
2. Download OnCommand Unified Manager version 6.4 (OnCommandUnifiedManager-6.4P1.ova), from http://mysupport.netapp.com/NOW/download/software/oncommand_cdot/6.4/. Click CONTINUE at the bottom of the page and follow the prompts to download the .ova file.
3. Log in to the vSphere Web Client as the FlexPod Admin user. From the Home screen, select VMs and Templates.
4. At the top of the center pane, click Actions > Deploy OVF Template.

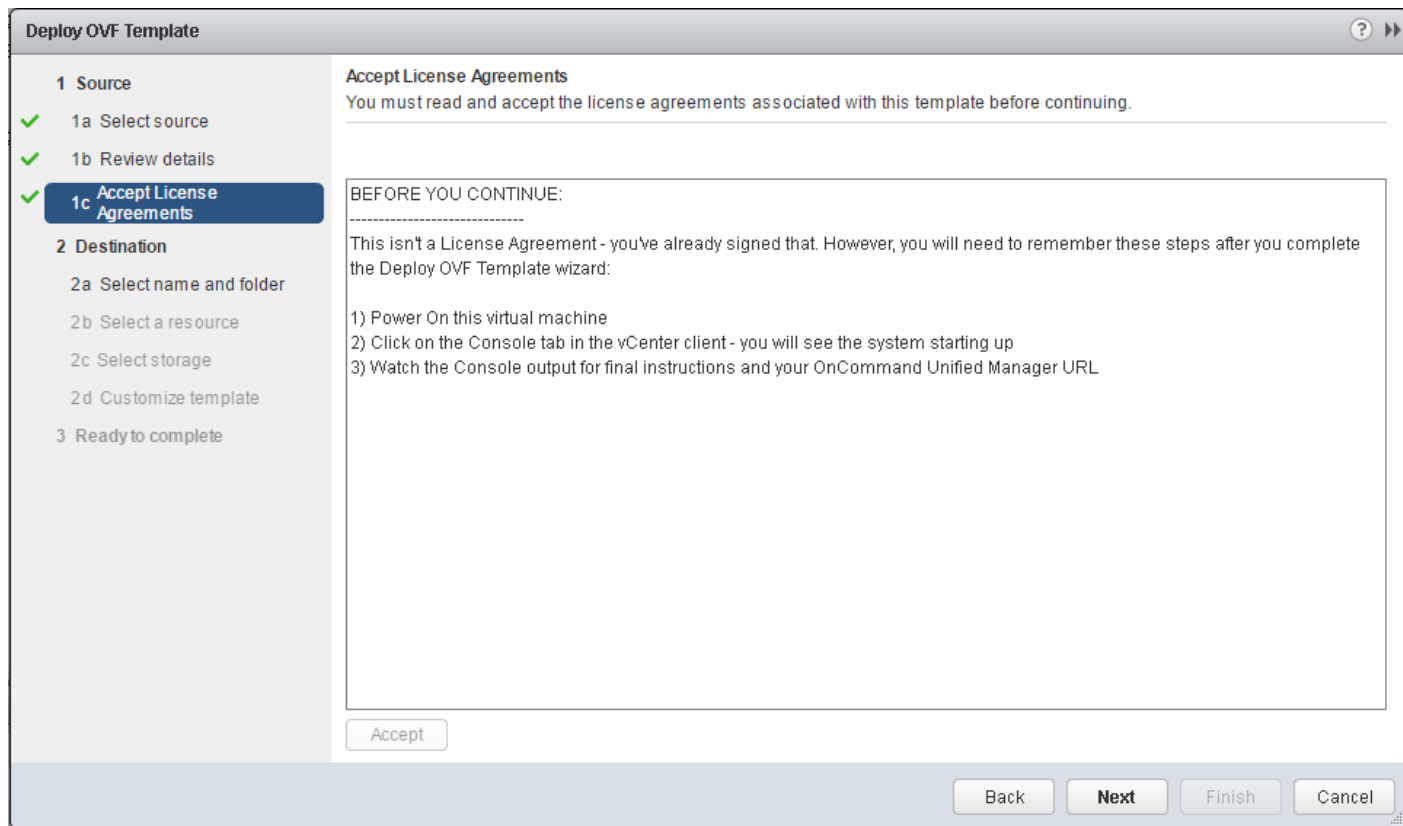


5. Browse to the .ova file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.



6. Click Next.

7. Click Accept to accept the agreement, and then click Next.



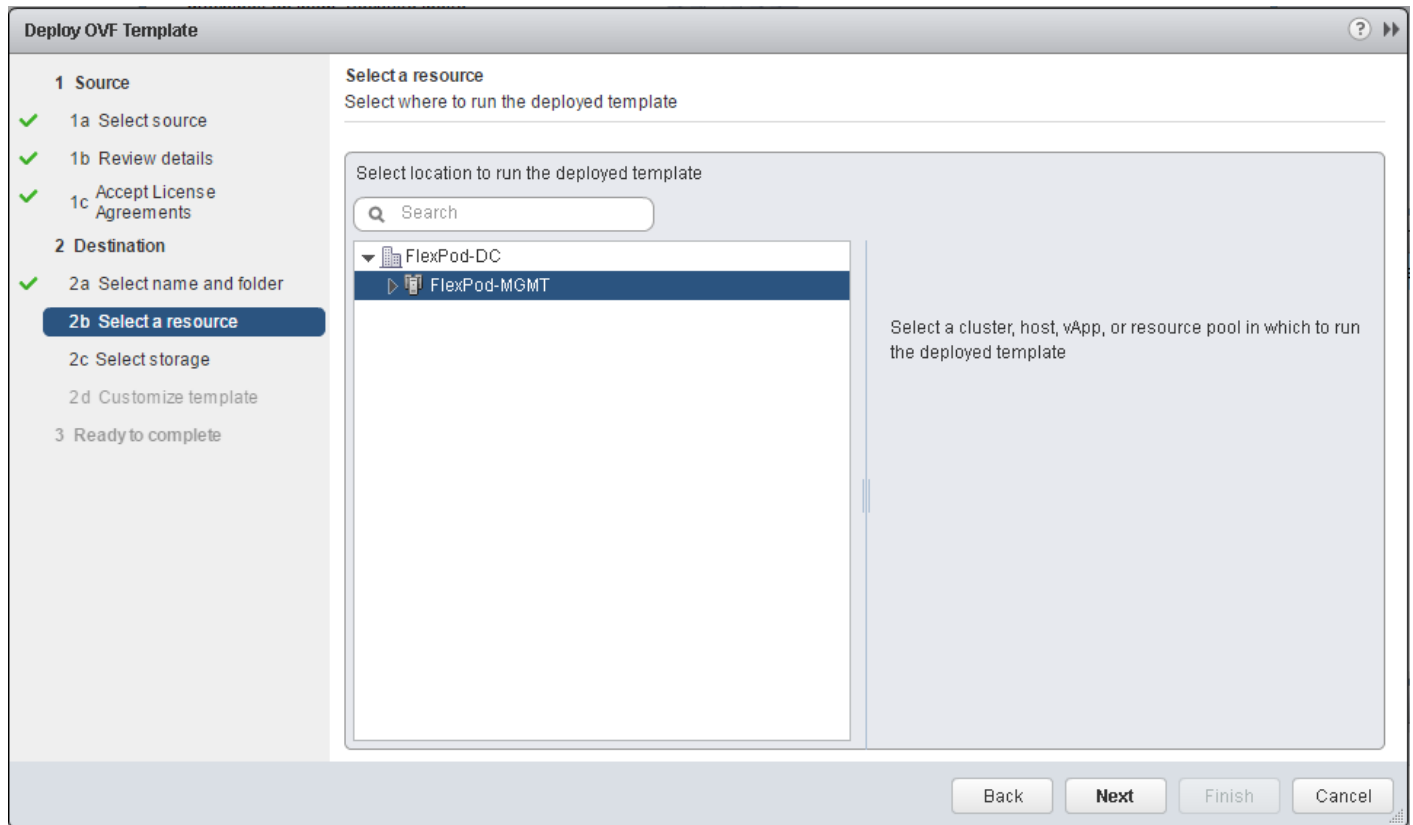
8. Enter the name of the VM and select the FlexPod-DC folder to hold the VM. Click Next.

The screenshot shows the 'Deploy OVF Template' wizard in vSphere. The left sidebar contains a progress list with the following items:

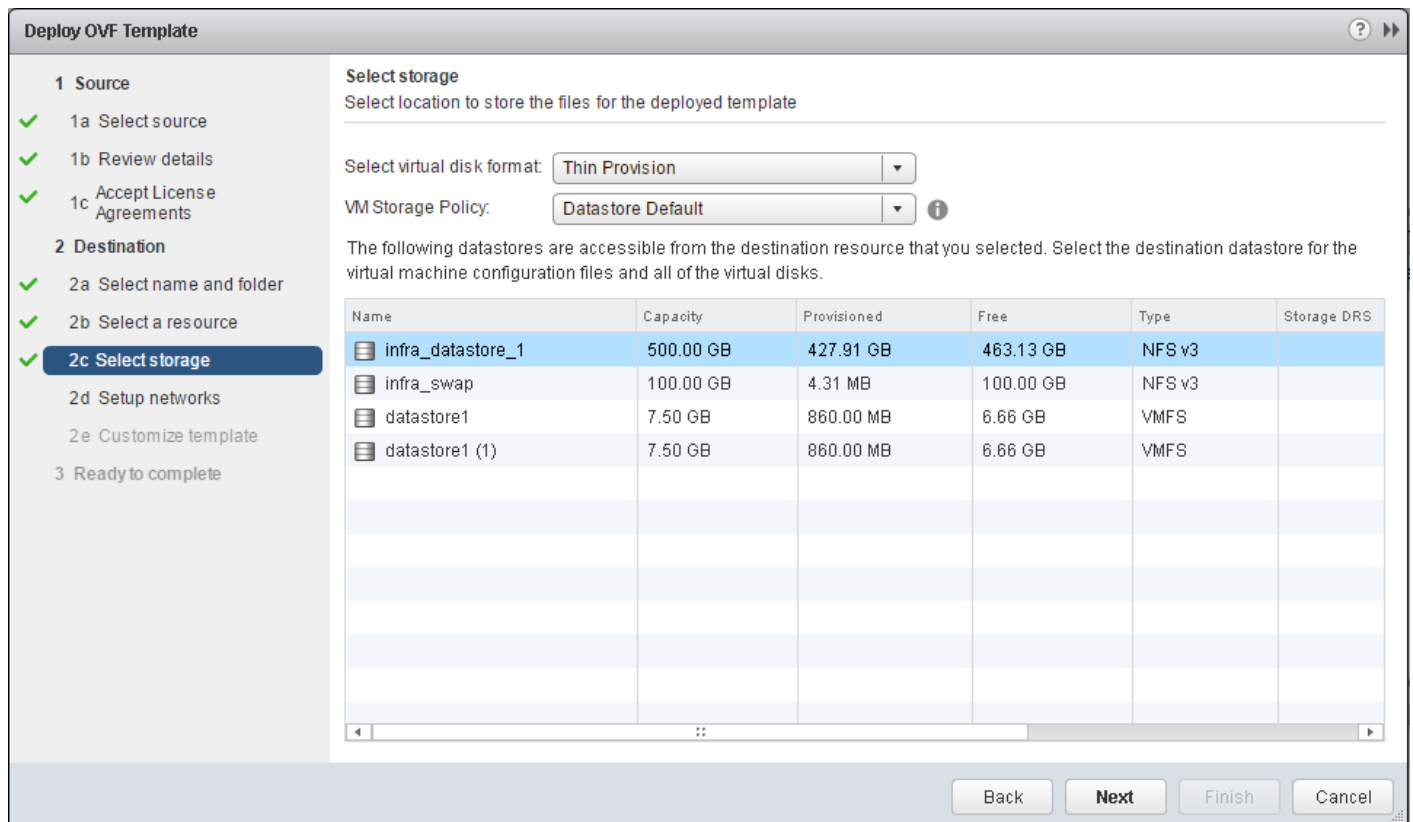
- 1 Source
 - 1a Select source
 - 1b Review details
 - 1c Accept License Agreements
- 2 Destination
 - 2a Select name and folder**
 - 2b Select a resource
 - 2c Select storage
 - 2d Customize template
- 3 Ready to complete

The main area is titled 'Select name and folder' and includes the instruction 'Specify a name and location for the deployed template'. A text field labeled 'Name:' contains the text 'OnCommand Unified Manager'. Below this is a tree view titled 'Select a folder or datacenter' with a search box. The tree view shows a folder structure: 'vc.texans.cisco.com' expanded to show 'FlexPod-DC'. The 'FlexPod-DC' folder is selected. To the right of the tree view, there is explanatory text: 'The folder you select is where the entity will be located, and will be used to apply permissions to it.' and 'The name of the entity must be unique within each vCenter Server VM folder.' At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

9. Select `FlexPod-MGMT` within the `FlexPod-DC` datacenter as the destination compute resource pool to host the VM. Click Next.



10. Select `infra_datastore_1` as the storage target for the VM and select Thin Provision as the virtual disk format. Click Next.



11. Select Foundation|IB-MGMT|IB-MGMT OR Foundation|IB-MGMT|IB-MGMT-AVS as the destination network to the nat source network. It is not necessary for this VM to be in the Core-Services Network. Click Next.
12. Fill out the details for the Host FQDN, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS. Click Next to continue.
13. Make sure Power on after deployment is cleared.
14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.

Deploy OVF Template

1 Source

- ✓ 1a Select source
- ✓ 1b Review details
- ✓ 1c Accept License Agreements

2 Destination

- ✓ 2a Select name and folder
- ✓ 2b Select a resource
- ✓ 2c Select storage
- ✓ 2d Setup networks
- ✓ 2e Customize template
- ✓ **3 Ready to complete**

Ready to complete
Review your settings selections before finishing the wizard.

OVF file	T:\NetApp\OnCommandUnifiedManager-6.4P1.ova
Download size	2.0 GB
Size on disk	4.7 GB
Name	OnCommand Unified Manager
Target	FlexPod-MGMT
Datastore	infra_datastore_1
Folder	FlexPod-DC
Disk storage	Thin Provision
Network mapping	nat to Foundation IB-MGMT IB-MGMT-AVS
IP allocation	Static - Manual, IPv4
Properties	Enables Auto IPv6 addressing for vApp. = False Host FQDN = ocum.texans.cisco.com IP Address = 172.26.163.172 Network Mask (or) Prefix Length = 255.255.255.0 Gateway = 172.26.163.254 Primary DNS = 172.26.163.42 Secondary DNS = 172.26.163.43

Power on after deployment

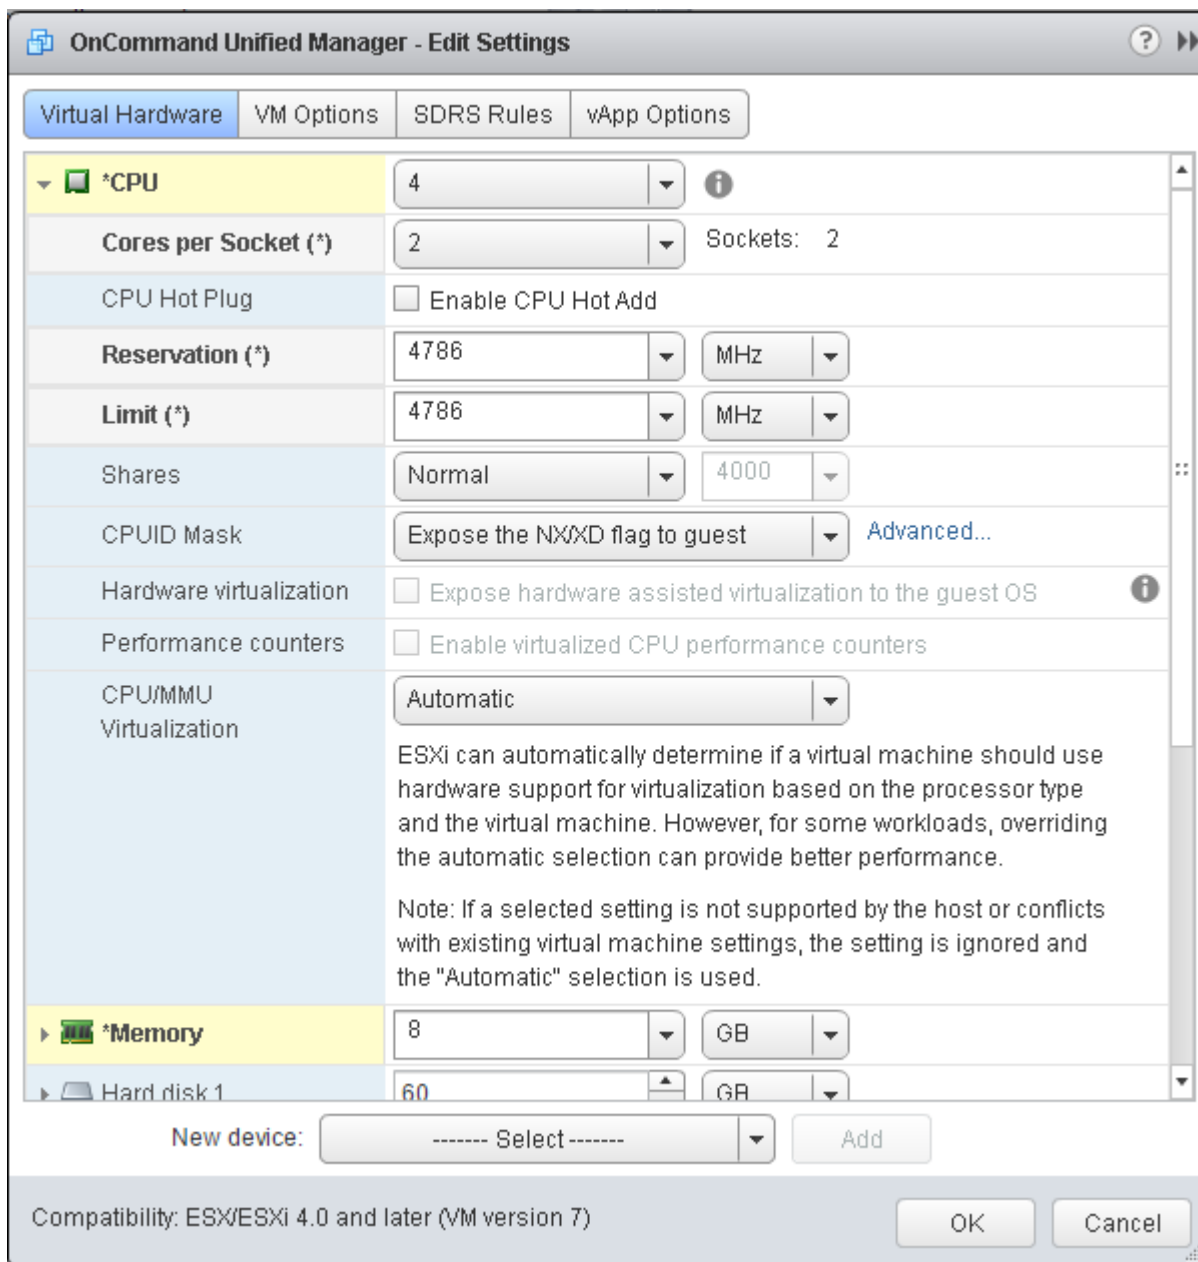
Back Next Finish Cancel

15. In the left pane, navigate to vCenter -> Virtual Machines. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.
16. Expand the CPU options, and complete the following steps:
 - a. The minimum required CPU Reservation can be lowered to 4786MHz. Set the CPU Reservation and Limit to 4786 MHz, and determine the CPU frequency of the host server.
 - b. Set the number of CPUs to the number of CPUs required (4786/CPU Frequency of host rounded up to the next even number).
 - c. Set the number of Cores per Socket where the Sockets number on the right matches the number of CPU sockets in the host. For example, if a host has two CPUs operating at a speed of 1995MHz, then the VM would need four virtual CPUs (4786 / 1995 = 2.40 - round-

ed to 4 virtual CPUs). If the host has two physical CPU sockets, set two cores per socket and the CPU reservation and limit to 4786 MHz.

- d. Set the amount of memory to 8GB.

Use the [OnCommand Unified Manager Installation and Setup Guide](#) for guidance on these settings.




- 17. Click OK to accept the changes.

- 18. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Unified Manager Basic Setup

- 1. Select the VM in the left-hand pane. In the center pane, select Open with Remote Console.

- In the VMware Remote Console (VMRC) window, select Manage > Install VMware Tools. VMware Tools installs in the VM.



- Set up OnCommand Unified Manager by answering the following questions in the console window:

Geographic area: <<Enter your geographic location>>

Time zone: <<Select the city or region corresponding to your time zone>>

These commands complete the network configuration, generate SSL certificates for HTTPS, and start the OnCommand Unified Manager services.

- To create a Maintenance User account, run the following commands:



Note: The maintenance user manages and maintains the settings on the OnCommand Unified Manager virtual appliance.

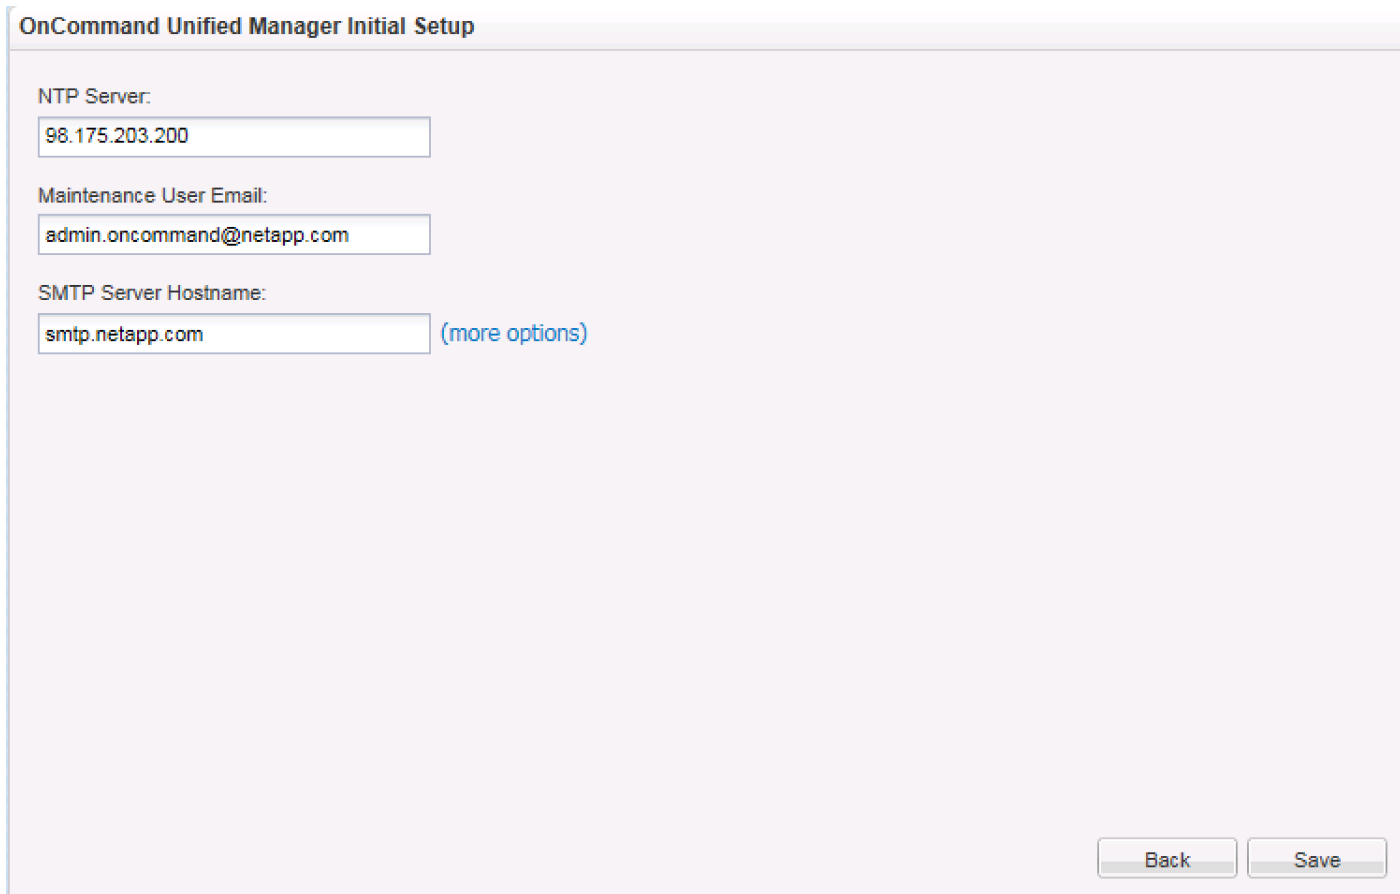
```
Username : admin
```

```
Enter new UNIX password: <password>
```

```
Retype new UNIX password: <password>
```

- With a web browser, navigate to the OnCommand Unified Manager using the URL `https:// <on-command-server-ip>`.
- Log in using the Maintenance User account credentials.
- Select `Yes` to enable AutoSupport capabilities.

8. Click Continue.
9. Provide the NTP Server IP address <ntp-server-ip>.
10. Provide the Maintenance User e-mail <storage-admin-email>.
11. Provide the SMTP Server Hostname.



The screenshot shows a web-based configuration form titled "OnCommand Unified Manager Initial Setup". The form contains three input fields with the following values:

- NTP Server:** 98.175.203.200
- Maintenance User Email:** admin.oncommand@netapp.com
- SMTP Server Hostname:** smtp.netapp.com (more options)

At the bottom right of the form, there are two buttons: "Back" and "Save".

12. Click Save.
13. Provide the Cluster Management IP address, User Name, Password, Protocol, and Port. Leave Application Instance for OnCommand Performance Manager set to None.

The screenshot shows the 'Add Cluster' form in NetApp OnCommand Unified Manager. The form includes the following fields and options:

- Host Name or IP Address:** 192.168.1.20
- User Name:** admin
- Password:** (masked with dots)
- Protocol:** HTTPS HTTP
- Port:** 443
- Link OnCommand Performance Manager with the Cluster (Recommended):** (indicated by an information icon)
- Select Application Instance:** None
- Save:** A blue button at the bottom right.

On the left side of the interface, there is a 'Get Started' section with the text: 'Welcome to OnCommand Unified Manager. Get started by adding the clusters you want to manage.'

14. Click Save. If asked to trust the certificate for the storage cluster, click Yes.



Note: The Cluster Add operation might take a couple of minutes.

15. On the left, select the storage cluster that was just added. Verify that the cluster has been added to OnCommand Unified Manager.

The screenshot shows the 'Managed Clusters' page in NetApp OnCommand Unified Manager. It displays a list of clusters and a detailed view of the selected cluster.

Managed Clusters (1)

Cluster ID	Host Name or IP Address	Cluster Health	Pairing Status
a01-aff8040	192.168.1.20	OK	Not Paired

Cluster: a01-aff8040

MONITORING STATUS

- Health: Poll completed
- Performance: Not available

CAPACITY

- 93.66% free
- Using 756.03 GB of 11.65 TB

PERFORMANCE

- Not available

Navigation buttons: Health, Performance (Setup)

Link OnCommand Performance Manager and OnCommand Unified Manager

To link the OnCommand Performance Manager and the OnCommand Unified Manager, complete the following steps:

1. Select the icon in the upper left-hand corner of the OnCommand Unified Manager web interface and select Health.
2. In the upper right, use the Administration drop-down to select Manage Users.
3. In the upper left, click Add to add a user.
4. Name the user `eventpub` and enter and confirm a password. Enter an e-mail address for this user that is different than the admin e-mail entered earlier for OnCommand Unified Manager. Select the Event Publisher Role.

Add User ? [X]

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

Type: Local User ▼

Name: eventpub

Password:

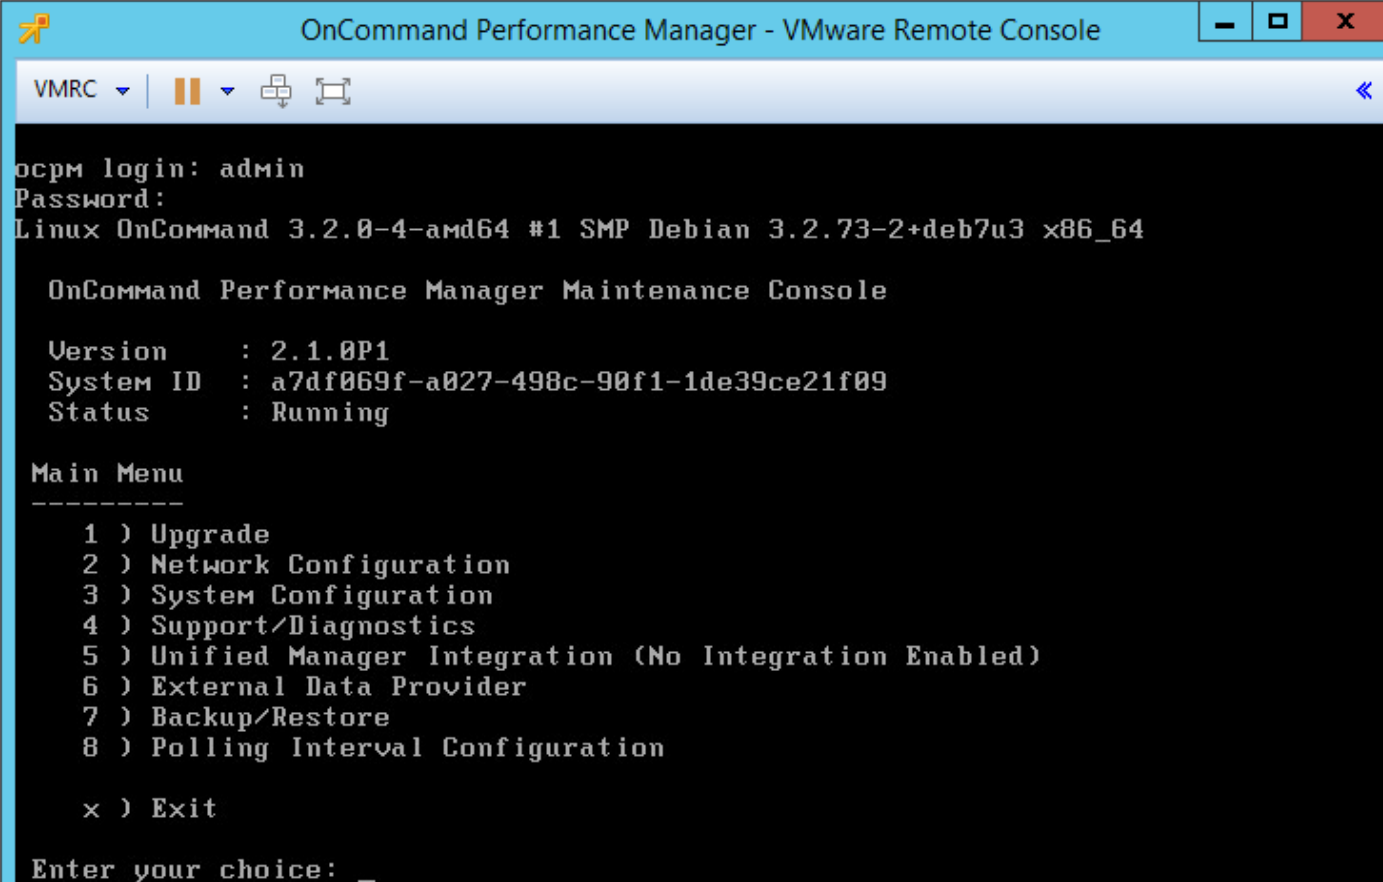
Confirm Password:

Email: eventpub@netapp.com

Role: Event Publisher ▼

[Add] [Cancel]

5. Click Add to add the user.
6. Back in the vSphere Web Client, select the OnCommand Performance Manager VM. In the center pane, select Open with Remote Console.
7. Log in to the OnCommand Performance Manager console with the admin user and password.



```
OnCommand Performance Manager - VMware Remote Console
VMRC | [Pause] [Print] [Fullscreen]
ocpm login: admin
Password:
Linux OnCommand 3.2.0-4-amd64 #1 SMP Debian 3.2.73-2+deb7u3 x86_64

OnCommand Performance Manager Maintenance Console

Version      : 2.1.0P1
System ID    : a7df069f-a027-498c-90f1-1de39ce21f09
Status       : Running

Main Menu
-----
1 ) Upgrade
2 ) Network Configuration
3 ) System Configuration
4 ) Support/Diagnostics
5 ) Unified Manager Integration (No Integration Enabled)
6 ) External Data Provider
7 ) Backup/Restore
8 ) Polling Interval Configuration

x ) Exit

Enter your choice: _
```

8. Select option 5 for Unified Manager Integration.
9. Select option 1 for Full Integration.
10. Select option 2 to Enable Full Integration.
11. Enter `y` to continue.
12. Enter the OnCommand Unified Manager IP address.
13. Enter `y` to accept the security certificate.
14. Enter `admin` for the Administrator Username.
15. Enter the admin password.
16. Enter the event publisher username (`eventpub`) entered above.
17. Enter the `eventpub` password.
18. Enter a unique name for the Performance Manager (for example, `ocpm-1`).
19. Review the settings and enter `y` if they are correct.

20. Press any key to restart the OnCommand Performance Manager service.
21. Press any key to continue.
22. Enter option 1 to Display the Full Integration Settings.
23. Press any key to continue.
24. Log out of the OnCommand Performance Manager console and return to the OnCommand Unified Manager web interface.
25. Using the icon in the upper left-hand corner, bring up the OnCommand Unified Manager Dashboard.
26. Select the storage cluster on the left. Verify that the Pairing Status is now Good and that performance numbers show up on the right under Performance.

The screenshot displays the NetApp OnCommand Unified Manager interface. The top navigation bar includes the title "NetApp OnCommand Unified Manager" and user profile icons. Below the navigation bar, a "CLUSTERS" tab is active. The main content area is divided into two sections: "Managed Clusters" on the left and a detailed view for cluster "a01-aff8040" on the right.

Managed Clusters:

Cluster ID	Cluster Health	Pairing Status	IP Address
a01-aff8040	OK	Good	192.168.1.20

Cluster: a01-aff8040 Details:

- MONITORING STATUS:** Health: Poll completed; Performance: Poll completed.
- CAPACITY:** 93.66% free; Using 756.17 GB of 11.65 TB.
- PERFORMANCE:** 85.32 IOPS; 2.6 MBps Throughput.

At the bottom of the detailed view, there are two tabs: "Health" and "Performance".

Sample Tenant Setup

Add Supernet Routes to Core-Services Devices

In this FlexPod with Cisco ACI lab validation, a Core-Services subnet was setup in Tenant common to allow Tenant VMs to access Core Services such as DNS, Active Directory Authentication, VMware vCenter, VMware ESXi, and NetApp Virtual Storage Console. Tenant VMs access the Core-Services devices over Layer 3 using their EPG subnet gateway. In this implementation, the Core-Services devices were setup connected by contract to the Bridged Layer 2 In Network that had a default gateway outside of the ACI Fabric. Since the Core-Services devices use this default gateway that is outside of the ACI Fabric, persistent, static routes must be placed in the Core-Services devices to reach the Tenant VMs.

To simplify this setup, all tenant VMs and devices connected to Core-Services had their IP subnets mapped from a range (172.16.0.0/16 in this deployment), allowing one Supernet route to be put into each Core-Services device. This section describes the procedure for deploying these Supernet routes to each type of Core-Services device.

Adding the Supernet Route in a Windows VM

To add a persistent Supernet Route in a Windows VM (AD servers and the NetApp VSC VM), open a command prompt with Administrator privileges in Windows and type the following command:

```
route -p ADD 172.16.0.0 MASK 255.255.0.0 <core-services-EPG-gateway>
route print
```

Adding the Supernet Route in the vCenter Server Appliance

To add a persistent Supernet Route in the VMware vCenter Server Appliance (VCSA) complete the following steps:

1. Using an ssh client, connect to the VCSA CLI and login as root.



Note: The VMware console can be used to connect to the CLI instead of ssh.

2. Type the following commands:

```
shell.set -enabled True
shell
echo 172.16.0.0 <core-services-EPG-gateway> 255.255.0.0 eth0 >
/etc/sysconfig/network/ifroute-eth0.
service network restart
route
```

Adding the Supernet Route in VMware ESXi

To add a persistent Supernet Route in VMware ESXi, complete the following steps:

1. Using an ssh client, connect to the VMware ESXi CLI and login as root.



Note: SSH will need to be enabled in the VMware ESXi Host Security Settings.



Note: This procedure can also be used to add VMkernel routes to VMware ESXi for routed storage protocols.

2. Type the following commands:

```
esxcli network ip route ipv4 add -gateway <core-services-EPG-gateway> --network 172.16.0.0/16
```

```
esxcfg-route -l
```

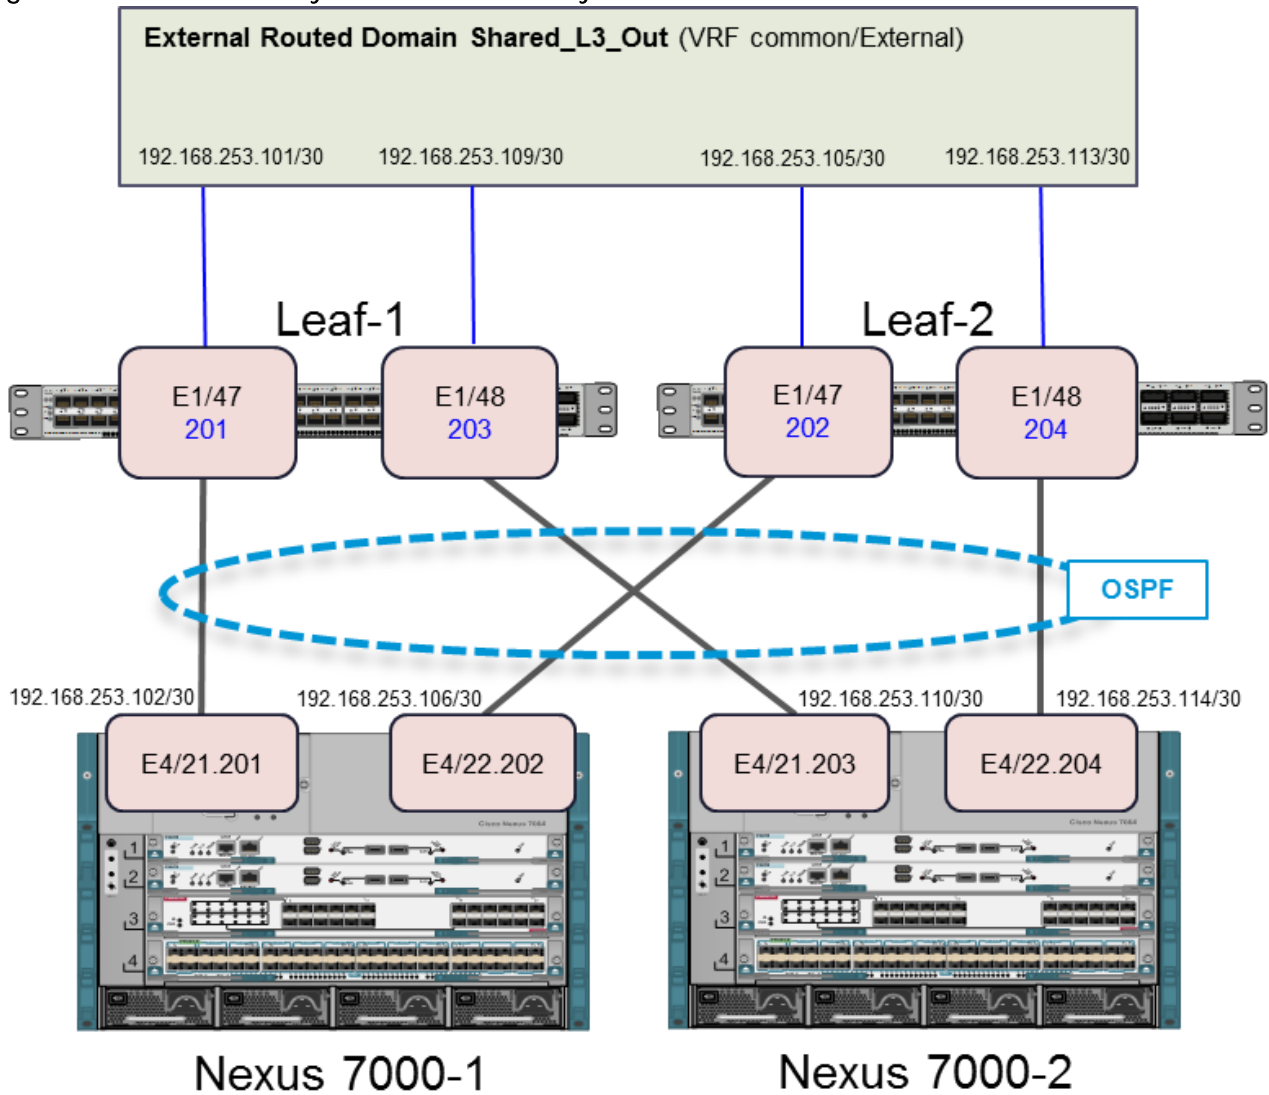
ACI Shared Layer 3 Out Setup

This section describes the procedure for deploying the ACI Shared Layer 3 Out. This external network is setup with a routing protocol and provides ACI tenants with a gateway to enter and leave the fabric.

This section provides a detailed procedure for setting up the Shared Layer 3 Out in Tenant common to existing Cisco Nexus 7000 core routers using sub-interfaces and VRF aware OSPF. Some highlights of this connectivity are:

- A new bridge domain and associated VRF is configured in Tenant common for external connectivity.
- **The shared Layer 3 Out created in Tenant common “provides” an external connectivity contract that can be “consumed” from any tenant.**
- Routes to tenant EPG subnets connected by contract are shared across VRFs with the Cisco Nexus 7000 core routers using OSPF.
- **The Cisco Nexus 7000s’ default gateway is shared with the ACI fabric using OSPF.**
- Each of the two Cisco Nexus 7000s is connected to each of the two Nexus 9000 leaf switches.
- Sub-interfaces are configured and used for external connectivity.
- The Cisco Nexus 7000s are configured to originate and send a default route to the Cisco Nexus 9000 leaf switches.

Figure 3 ACI Shared Layer 3 Out Connectivity Details



Configuring the Cisco Nexus 7000s for ACI Connectivity (Sample)

The following configuration is a sample from the virtual device contexts (VDCs) from two Cisco Nexus 7004s. Interfaces and a default route from the two Cisco Nexus 7000s also needs to be set up, but is not shown here because this would be set up according to customer security policy.

Cisco Nexus 7004-1 VDC

```
feature ospf

vlan 100
  name OSPF-Peering

interface Vlan100
```


Sample Tenant Setup

```
no shutdown
mtu 9216
no ip redirects
ip address 192.168.253.253/30
no ipv6 redirects
ip ospf mtu-ignore
ip router ospf 10 area 0.0.0.0

interface Ethernet4/21
no shutdown

interface Ethernet4/21.201
encapsulation dot1q 201
ip address 192.168.253.102/30
ip ospf network point-to-point
ip ospf mtu-ignore
ip router ospf 10 area 0.0.0.10
no shutdown

interface Ethernet4/22
no shutdown

interface Ethernet4/22.202
encapsulation dot1q 202
ip address 192.168.253.106/30
ip ospf cost 5
ip ospf network point-to-point
ip ospf mtu-ignore
ip router ospf 10 area 0.0.0.10
no shutdown

interface loopback0
ip address 192.168.254.3/32
```

Sample Tenant Setup

```
ip router ospf 10 area 0.0.0.0
```

```
router ospf 10
```

```
router-id 192.168.254.3
```

```
area 0.0.0.10 nssa no-summary default-information-originate no-redistribution
```

Cisco Nexus 7004-2 VDC

```
feature ospf
```

```
vlan 100
```

```
name OSPF-Peering
```

```
interface Vlan100
```

```
no shutdown
```

```
mtu 9216
```

```
no ip redirects
```

```
ip address 192.168.253.254/30
```

```
no ipv6 redirects
```

```
ip ospf mtu-ignore
```

```
ip router ospf 10 area 0.0.0.0
```

```
interface Ethernet4/21
```

```
no shutdown
```

```
interface Ethernet4/21.203
```

```
encapsulation dot1q 203
```

```
ip address 192.168.253.110/30
```

```
ip ospf cost 21
```

```
ip ospf network point-to-point
```

```
ip ospf mtu-ignore
```

```
ip router ospf 10 area 0.0.0.10
```

```
no shutdown
```

```
interface Ethernet4/22
```

```
no shutdown

interface Ethernet4/22.204
  encapsulation dot1q 204
  ip address 192.168.253.114/30
  ip ospf cost 30
  ip ospf network point-to-point
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
  no shutdown

interface loopback0
  ip address 192.168.254.4/32
  ip router ospf 10 area 0.0.0.0

router ospf 10
  router-id 192.168.254.4
  area 0.0.0.10 nssa no-summary default-information-originate no-redistribution
```

Configuring ACI Shared Layer 3 Out

ACI Advanced GUI

To configure the ACI Shared Layer 3 Out, complete the following steps:

1. At the top, select Fabric > Access Policies.
2. On the left, expand Physical and External Domains.
3. Right-click External Routed Domains and select Create Layer 3 Domain.
4. Name the Domain Shared-L3-Out.
5. Use the Associated Attachable Entity Profile drop-down to select Create Attachable Entity Profile.
6. Name the Profile AEP-Shared-L3-Out and click NEXT.

Create Attachable Access Entity Profile

1. Profile 2. Association To Interfaces

STEP 1 > Profile

Specify the name, domains and infrastructure encaps

Name:

Description:

Enable Infrastructure VLAN:

PREVIOUS NEXT CANCEL

7. Click FINISH to continue without specifying interfaces.
8. Back in the Create Layer 3 Domain window, use the VLAN Pool drop-down to select Create VLAN Pool.
9. Name the VLAN Pool VP-Shared-L3-Out and select Static Allocation.
10. Click the + sign to add and Encap Block.
11. In the Create Ranges window, enter the From and To VLAN IDs for the Shared-L3-Out VLAN range (201-204). Select Static Allocation.

15. At the top, select Fabric > Inventory.

16. On the left, select Topology. On the right, select Configure.

17. In the center pane, select ADD SWITCHES.

18. Using the shift key, select the two leaf switches and select ADD SELECTED.

19. On the two switches, select the 4 ports connected to the Nexus 7000s and used for Shared-L3-Out.

The screenshot displays a network configuration interface. At the top, there are buttons for 'BACK TO TOPOLOGY', 'ADD SWITCHES', and 'REFRESH'. A legend indicates 'Port Channel' (light blue), 'L3' (light blue), 'VPC' (green), and 'Conn. tc' (orange). Two switch configurations are shown: 'a01-93180-1 (Node-101)' and 'a01-93180-2 (Node-102)'. Each switch has a grid of 48 ports (2 rows of 24). In both switches, ports 15, 16, 19, and 20 are highlighted in green, indicating they are selected. Below the switches is a 'Summary' section. It features three tabs: 'Port Channels' (selected), 'Virtual Port Channels', and 'Fab'. The 'Virtual Port Channels' tab contains a table with the following data:

Switch 1	Ports	Switch 2	Ports
101	1/23	102	1/23
101	1/15	102	1/15
101	1/16	102	1/16
101	1/19	102	1/19
101	1/20	102	1/20

To the right of the table is a 'Fab' column with a 'Swi' sub-column and five buttons: 'CONFIGURE PORT', 'CLEAR PORT CONF.', 'CONFIGURE PC', 'DELETE PC', and 'CONFIGURE VPC'.

20. On the lower right, select CONFIGURE PORT.

21. Select the appropriate policies and AEP-Shared-L3-Out as shown in the screenshot below.

BACK TO SUMMARY

CONFIGURING INTERFACE

Port Channel
VPC
L2 Interface

L3
Conn. to Fex
Selected

a01-93180-1 (Node-101) ✕

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a01-93180-2 (Node-102) ✕

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

Interface

Description: optional

Link Level Policy: 10Gbps-Auto ▼

CDP Policy: CDP-Enabled ▼

MCP Policy: default ▼

LLDP Policy: LLDP-Enabled ▼

STP Interface Policy: No-BPDU-Filter-Guard ▼

Egress Data Plane Policing Policy: default ▼

Ingress Data Plane Policing Policy: default ▼

Storm Control Interface Policy: default ▼

L2 Interface Policy: VLAN-Scope-Global ▼

Attached Entity Profile: AEP-Shared-L3-i ▼

BACK TO SUMMARY
APPLY CHANGES

22. Select APPLY CHANGES to configure the ports. Click OK for the Success confirmation.

23. At the top, select Tenants > common.

24. On the left, expand Tenant common and Networking.

25. Right-click VRFs and select create VRF.

26. Name the VRF common-External. Select default for both the End Point Retention Policy and Monitoring Policy.

Create VRF

1. VRF 2. Bridge Domain

STEP 1 > VRF


Specify Tenant VRF


Name:

Description:

Policy Control Enforcement Preference: Enforced Unenforced

Policy Control Enforcement Direction: Egress Ingress

End Point Retention Policy: 
This policy only applies to remote L3 entries

Monitoring Policy: 

DNS Labels:
enter names separated by comma

Route Tag Policy:

Create A Bridge Domain:

Configure BGP Policies:

Configure OSPF Policies:

Configure EIGRP Policies:

PREVIOUS NEXT CANCEL

27. Leave Create A Bridge Domain selected and click NEXT.

28. Name the Bridge Domain BD-common-External. Leave all other values unchanged.

Create VRF

1. VRF 2. Bridge Domain

STEP 2 > Bridge Domain

Specify Bridge Domain for the VRF

Name:

Description:

Forwarding:

Config BD MAC Address:

MAC Address:

PREVIOUS FINISH CANCEL

29. Click FINISH to complete creating the VRF.
30. On the left, right-click External Routed Networks and select Create Routed Outside.
31. Name the Routed Outside Shared-L3-Out.
32. Select the checkbox next to OSPF.
33. Enter 0.0.0.10 (configured in the Nexus 7000s) as the OSPF Area ID.
34. Using the VRF drop-down, select common/common-External.
35. Using the External Routed Domain drop-down, select Shared-L3-Out.
36. Click the + sign to the right of Nodes and Interfaces Protocol Profiles to add a Node Profile.
37. Name the Node Profile Nodes-101-102 for the Nexus 9000 Leaf Switches.
38. Click the + sign to the right of Nodes to add a Node.

39. In the select Node window, select Leaf switch 101.

40. Provide a Router ID IP address that will also be used as the Loopback Address (192.168.254.101).

Select Node

Select Node and Configure Static Routes

Node ID: topology/pod-1/node-101

Router ID: 192.168.254.101

Use Router ID as Loopback Address:

Loopback Addresses:

IP
192.168.254.101

Static Routes:

IP Address	Next Hop IP
------------	-------------

OK CANCEL

41. Click OK to complete selecting the Node.

42. Click the + sign to the right of Nodes to add a Node.

43. In the select Node window, select Leaf switch 102.

44. Provide a Router ID IP address that will also be used as the Loopback Address (192.168.254.102).

Select Node

Select Node and Configure Static Routes

Node ID:

Router ID:

Use Router ID as Loopback Address:

Loopback Addresses:

IP
192.168.254.102

Static Routes:

IP Address	Next Hop IP
------------	-------------

45. Click OK to complete selecting the Node.
46. Click the + sign to the right of OSPF Interface Profiles to create an OSPF Interface Profile.
47. Name the profile OIP-Nodes-101-102.
48. Using the OSPF Policy drop-down, select Create OSPF Interface Policy.
49. Name the policy To-7K.
50. Select the Point-to-Point Network Type.
51. Select the Advertise subnet and MTU ignore Interface Controls.

Create OSPF Interface Policy

Define OSPF Interface Policy

Name: To-7K

Description: optional

Network Type: Broadcast Point-to-point Unspecified

Priority: 1

Cost of Interface: unspecified

Interface Controls: Advertise subnet
 BFD
 MTU ignore
 Passive participation

Hello Interval (sec): 10

Dead Interval (sec): 40

Retransmit Interval (sec): 5

Transmit Delay (sec): 1

52. Click SUBMIT to complete creating the policy.

53. Select Routed Sub-Interface under Interfaces.

54. Click the + sign to the lower right of Routed Sub-Interfaces to add a routed sub-interface.

55. In the Select Routed Sub-Interface window, select the interface on Node 101 that is connected to Nexus 7000-1.

56. Enter vlan-201 for Encap.

57. Enter the IPv4 Primary Address (192.168.253.101/30)

58. Leave the MTU set to inherit.

Select Routed Sub-Interface

Specify the Interface

Path: ▾

Encap:
For example, vlan-1

IPv4 Primary / IPv6 Preferred Address:
address/mask

IPv4 Secondary / IPv6 Additional Addresses:
Address

MAC Address:

MTU (bytes):

Link-local Address:

59. Click OK to complete creating the routed sub-interface.

60. Repeat steps 54-59 to add the second Leaf 1 interface (VLAN 203, IP 192.168.253.109/30), the first Leaf 2 interface (VLAN 202, IP 192.168.253.105/30), and the second Leaf 2 interface (VLAN 204, IP 192.168.253.113/30).

Create Interface Profile

Specify the Interface Profile

Name:

Description:

ND policy:

Egress Data Plane Policing Policy:

Ingress Data Plane Policing Policy:

OSPF Profile

Authentication Type:

Authentication Key:

Confirm Key:

OSPF Policy:

BFD Interface Profile

Authentication Type:

BFD Interface Policy:

Interfaces

Routed Interfaces SVI Routed Sub-Interface

Routed Sub-Interfaces × +

Path	Encap	IP Address	MAC Address	MTU (bytes)
Node-101/eth1/47	vlan-201	192.168.253.101/30	00:22:BD:F8:19:FF	inherit
Node-101/eth1/48	vlan-203	192.168.253.109/30	00:22:BD:F8:19:FF	inherit
Node-102/eth1/47	vlan-202	192.168.253.105/30	00:22:BD:F8:19:FF	inherit
Node-102/eth1/48	vlan-204	192.168.253.113/30	00:22:BD:F8:19:FF	inherit

OK
CANCEL

61. Click OK to complete creating the Node Interface Profile.

Create Node Profile

Specify the Node Profile

Name: **Nodes-101-102**

Description: optional

Target DSCP: unspecified

Nodes: ✕ +

Node ID	Router ID	Static Routes	Loopback Address
topology/pod-1/n...	192.168.254.101		
topology/pod-1/n...	192.168.254.102		

OSPF Interface Profiles: ✕ +

Name	Description	Interfaces	OSPF Policy
OIP-Nodes-101-102		[eth1/47], [eth1/47], [eth1/48], [eth1/48]	To-7K

OK CANCEL

62. Click OK to complete creating the Node Profile.

Create Routed Outside
i X

STEP 1 > Identity

1. Identity
2. External EPG Networks

Define the Routed Outside

Name:

Description:

Tags:
enter tags separated by comma

Route Control Enforcement: Import Export

Target DSCP:

VRF: +

External Routed Domain: +

Route Profile for Interleak:

Route Control For Dampening: x +

Address Family Type	Route Dampening Policy

BGP EIGRP

OSPF

OSPF Area ID:

OSPF Area Control: Send redistributed LSAs into NSSA area
 Originate summary LSA
 Suppress forwarding address in translated LSA

OSPF Area Type: NSSA area Regular area Stub area

OSPF Area Cost:

Nodes And Interfaces Protocol Profiles x +

Name	Description	DSCP	Nodes
Nodes-101-102		Unspecified	101, 102

PREVIOUS
NEXT
CANCEL

63. Click NEXT.

64. Click the + sign to create an External EPG Network.

65. Name the External Network Default-Route.

66. Click the + sign to add a Subnet.

67. Enter 0.0.0.0/0 as the IP Address. Select the checkboxes for External Subnets for the External EPG, Shared Route Control Subnet, and Shared Security Import Subnet.

Create Subnet

Specify the Subnet

IP Address:
address/mask

scope: Export Route Control Subnet
 Import Route Control Subnet
 External Subnets for the External EPG
 Shared Route Control Subnet
 Shared Security Import Subnet

OSPF Route Summarization Policy:

aggregate: Aggregate Export
 Aggregate Import
 Aggregate Shared Routes

Route Control Profile:

Name	Direction

68. Click OK to complete creating the subnet.

Create External Network
i
✕

Define an External Network

Name: Default-Route

Tags: ▼
enter tags separated by comma

QoS class: Unspecified ▼

Description:

Target DSCP: unspecified

Subnet ✕ +

IP Address	Scope	Aggregate	Route Control Profile	Route Summarization Policy
0.0.0.0/0	External Subnets for the Ex...	Shared Route Control Subn...	Shared Security Import Sub...	

69. Click OK to complete creating the external network.

70. Click FINISH to complete creating the Shared-L3-Out.

71. On the left, right-click Security Policies and select Create Contract.

72. Name the contract Allow-Shared-L3-Out.

73. Select the Global Scope to allow the contract to be consumed from all tenants.

74. Click the + sign to the right of Subjects to add a contract subject.

75. Name the subject Allow-All.

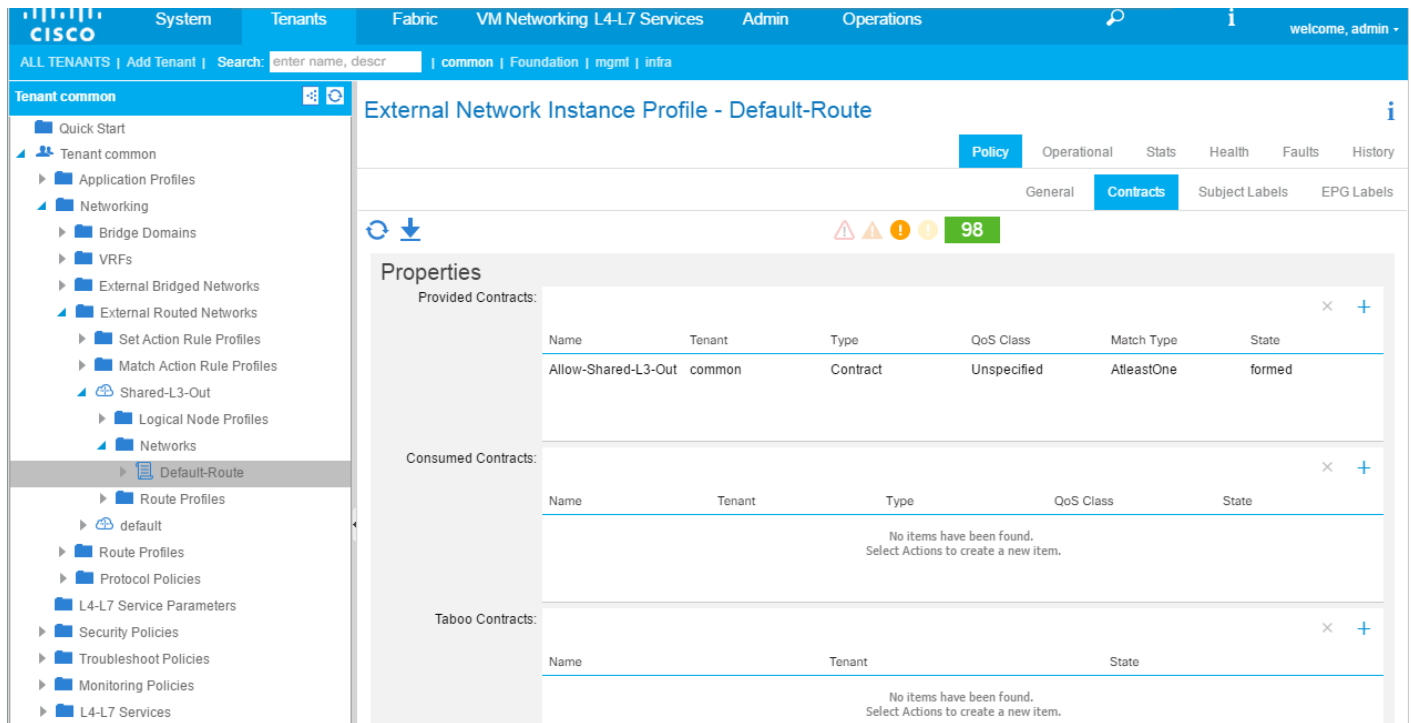
76. Click the + sign to the right of Filters to add a filter.

77. Use the drop-down to select the Allow-All filter from Tenant common.

78. Click UPDATE.

79. Click OK to complete creating the contract subject.

80. Click SUBMIT to complete creating the contract.
81. On the left, expand Tenant common, Networking, External Routed Networks, Shared-L3-Out, and Networks. Select Default-Route.
82. On the right, under Policy, select Contracts.
83. Click the + sign to the right of Provided Contracts to add a Provided Contract.
84. Select the common/Allow-Shared-L3-Out contract and click UPDATE.



Note: Tenant EPGs can now consume the Allow-Shared-L3-Out contract and connect outside of the fabric. Note that more restrictive contracts can be built and provided here for more restrictive access to the outside.

Lab Validation Tenant Configuration

The following table shows the VLANs, Subnets, and Bridge Domains for the sample App-A Tenant set up as part of this lab validation:

Table 25 Lab Validation Tenant Contracts App-A Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI-A	3011	3111	192.168.111.0/24 - L2	BD-iSCSI-A
iSCSI-B	3021	3121	192.168.121.0/24 - L2	BD-iSCSI-B
NFS-LIF	3051	N/A	192.168.151.0/24 - L2	BD-NFS
NFS-VMK	N/A	DVS	192.168.151.0/24 - L2	BD-NFS
SVM-MGMT	264	N/A	172.16.254.6/29	BD-Internal
Web	N/A	DVS	172.16.0.254/24	BD-Internal

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
App	N/A	DVS	172.16.1.254/24	BD-Internal
DB	N/A	DVS	172.16.2.254/24	BD-Internal

Configure Tenant Storage

This section describes the procedure for deploying a NetApp storage SVM for a tenant named App-A. In this section, VLAN interface ports, a separate IPspace, the tenant SVM, storage protocols within the SVM, tenant logical interfaces (LIFs), and tenant data volumes are deployed. All procedures in this section are completed using a SSH connection to the storage cluster CLI.

Create Tenant IPspace

To create the tenant IPspace, run the following commands:

```
ipSPACE create -ipSPACE App-A
ipSPACE show
```

Create Tenant Broadcast Domains in ONTAP

To create data broadcast domains in the tenant IPspace, run the following commands. If you are not setting up access to iSCSI application data LUNs in this tenant, do not create the iSCSI broadcast domains.

```
broadcast-domain create -ipSPACE App-A -broadcast-domain App-A-NFS -mtu 9000
broadcast-domain create -ipSPACE App-A -broadcast-domain App-A-SVM-MGMT -mtu 1500
broadcast-domain create -ipSPACE App-A -broadcast-domain App-A-iSCSI-A -mtu 9000
broadcast-domain create -ipSPACE App-A -broadcast-domain App-A-iSCSI-B -mtu 9000
broadcast-domain show -ipSPACE App-A
```



Note: If using the Cisco AVS in this FlexPod Implementation, set the MTU of the App-A-NFS broadcast domain to 8950 instead of 9000.

Create VLAN Interfaces

To create tenant-storage VLAN interfaces, complete the following steps:

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

```
network port vlan create -node <node01> -vlan-name a0a-<storage-App-A-nfs-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-App-A-nfs-vlan-id>

broadcast-domain add-ports -ipSPACE App-A -broadcast-domain App-A-NFS -ports
<node01>:a0a-<storage-App-A-nfs-vlan-id>, <node02>:a0a-<storage-App-A-nfs-vlan-id>
```

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

```
network port vlan create -node <node01> -vlan-name a0a-<storage-App-A-svm-mgmt-vlan-id>
```

```
network port vlan create -node <node02> -vlan-name a0a-  

<storage-App-A-svm-mgmt-vlan-id>
```

```
broadcast-domain add-ports -ip-space App-A -broadcast-domain App-A-SVM-MGMT -ports  

<node01>:a0a-  

<storage-App-A-svm-mgmt-vlan-id>, <node02>:a0a-  

<storage-App-A-svm-mgmt-vlan-id>
```

3. Create tenant iSCSI VLAN ports and add them to the data broadcast domain. If you are not setting up access to iSCSI application data LUNs in this tenant, do not create the iSCSI VLAN ports.

```
network port vlan create -node <node01> -vlan-name a0a-  

<storage-App-A-iscsi-A-vlan-id>
```

```
network port vlan create -node <node01> -vlan-name a0a-  

<storage-App-A-iscsi-B-vlan-id>
```

```
network port vlan create -node <node02> -vlan-name a0a-  

<storage-App-A-iscsi-A-vlan-id>
```

```
network port vlan create -node <node02> -vlan-name a0a-  

<storage-App-A-iscsi-B-vlan-id>
```

```
broadcast-domain add-ports -ip-space App-A -broadcast-domain App-A-iSCSI-A -ports  

<node01>:a0a-  

<storage-App-A-iscsi-A-vlan-id>, <node02>:a0a-  

<storage-App-A-iscsi-A-vlan-id>
```

```
broadcast-domain add-ports -ip-space App-A -broadcast-domain App-A-iSCSI-B -ports  

<node01>:a0a-  

<storage-App-A-iscsi-B-vlan-id>, <node02>:a0a-  

<storage-App-A-iscsi-B-vlan-id>
```

```
broadcast-domain show -ip-space App-A
```

Create Tenant Storage Virtual Machine

To create the tenant App-A SVM in the App-A IPspace, complete the following steps:



Note: The SVM is referred to as a Vserver (or `vserver`) in the GUI and CLI.

1. Run the `vserver create` command.

```
vserver create -vserver App-A-SVM -rootvolume rootvol -aggregate aggr1_node01 -  

rootvolume-security-style unix -ip-space App-A
```

2. Remove unused SVM storage protocols from the list of `nfs`, `cifs`, `fcp`, `iscsi`, and `ndmp`. In this example, we keep `nfs`, `fcp`, and `iscsi`.

```
vserver remove-protocols -vserver App-A-SVM -protocols cifs,ndmp
```

3. Add the two data aggregates to the App-A-SVM aggregate list for the NetApp VSC to be able to provision storage in those aggregates.

```
vserver modify -vserver App-A-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the NFS protocol in the App-A-SVM.

```
nfs create -vserver App-A-SVM -udp disabled
```

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

```
vserver nfs modify -vserver App-A-SVM -vstorage enabled
```

```
vserver nfs show -vserver App-A-SVM
```

6. If iSCSI LUN access is being provided by this SVM, create the iSCSI service on this SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM.

```
iscsi create -vserver App-A-SVM
```

```
iscsi show
```

7. If FCoE LUN access is provided by this SVM, create the FCP service on this SVM. This command also starts the FCP service and sets the FCP World Wide Node Name (WWNN) for the SVM.

```
fcp create -vserver App-A-SVM
```

```
fcp show
```

Create Load-Sharing Mirrors of SVM Root Volume

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

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

```
volume create -vserver App-A-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
```

```
volume create -vserver App-A-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create the mirroring relationships.

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

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

3. Initialize the mirroring relationship.

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

Configure HTTPS Access

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

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

```
set diag
```

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

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

```
security certificate show
```

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



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

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

```
Example: security certificate delete -vserver App-A-SVM -common-name App-A-SVM -ca App-A-SVM -type server -serial 5375EEF1D7AE5
```

4. To generate and install a self-signed certificate, run the following command as a one-time command. Generate a server certificate for App-A-SVM. Use TAB completion to aid in the completion of this command.

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

```
Example: security certificate create -common-name app-a-svm.texans.cisco.com -type server -size 2048 -country US -state "North Carolina" -locality "RTP" -organization "Cisco" -unit "UCS" -email-addr "admin@texans.cisco.com" -expire-days 365 -protocol SSL -hash-function SHA256 -is-system-internal-certificate false -vserver App-A-SVM
```

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

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

```
Example: security ssl modify -vserver App-A-SVM -server-enabled true -client-enabled false -ca app-a-svm.texans.cisco.com -serial 5375F34974EB8 -common-name app-a-svm.texans.cisco.com
```

7. Change back to the normal admin privilege level and set up the system to enable SVM logs to be accessed from the web.

```
set admin
```

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

Configure NFSv3

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

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for the App-A NFS subnet CIDR address (for example, 192.168.151.0/24).

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

```
vserver export-policy rule show
```

2. Assign the FlexPod export policy to the App-A SVM root volume.

```
volume modify -vserver App-A-SVM -volume rootvol -policy default
```

Create FlexVol Volumes

To create a volume for the App-A NFS datastore, run the following commands:

```
volume create -vserver App-A-SVM -volume App_A_datastore_1 -aggregate
aggr1_node02 -size 500GB -state online -policy default -junction-path
/App_A_datastore_1 -space-guarantee none -percent-snapshot-space 0
```

```
snapmirror update-ls-set -source-path App-A-SVM:rootvol
```



Note: You are not creating a volume for a swap datastore here. Testing revealed that if the swap datastore was on a VMkernel port on a DVS, the swap datastore would no longer be set in the ESXi host settings on reboot. NetApp recommends using the `infra_swap` datastore with a VMkernel port on vSwitch0.

Adjust Storage Efficiency Settings

On NetApp All Flash FAS systems, volumes are created by default inline compression and inline deduplication is enabled with no scheduled deduplication scans. In this section a deduplication scan schedule is added to the volume `App_A_datastore_1`, and inline deduplication is removed from the volume `App_A_swap`. To adjust the storage efficiency settings, complete the following steps:

1. Add a daily deduplication scan to the `App_A_datastore_1` volume.

```
efficiency modify -vserver App-A-SVM -volume App_A_datastore_1 -schedule sun-
sat@0
efficiency show -instance -vserver App-A-SVM
```

Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), run the following commands.



Note: If you are not setting up access to iSCSI application data LUNs in this tenant, do not create the iSCSI LIFs.

```
network interface create -vserver App-A-SVM -lif iscsi_lif01a -role data -data-
protocol iscsi -home-node <node01> -home-port a0a-<storage-App-A-iscsi-A-vlan-id>
-address <node01-iscsi-App-A-lif01a-ip> -netmask <storage-App-A-iscsi-A-mask>
```

```
network interface create -vserver App-A-SVM -lif iscsi_lif01b -role data -data-
protocol iscsi -home-node <node01> -home-port a0a-<storage-App-A-iscsi-B-vlan-id>
-address <node01-App-A-iscsi-lif01b-ip> -netmask <storage-App-A-iscsi-B-mask>
```

```
network interface create -vserver App-A-SVM -lif iscsi_lif02a -role data -data-
protocol iscsi -home-node <node02> -home-port a0a-<storage-App-A-iscsi-A-vlan-id>
-address <node02-App-A-iscsi-lif02a-ip> -netmask <storage-App-A-iscsi-A-mask>
```

```
network interface create -vserver App-A-SVM -lif iscsi_lif02b -role data -data-
protocol iscsi -home-node <node02> -home-port a0a-<storage-App-A-iscsi-B-vlan-id>
-address <node02-App-A-iscsi-lif02b-ip> -netmask <storage-App-A-iscsi-B-mask>
```

```
network interface show -vserver App-A-SVM -lif iscsi*
```


Create FCoE LIFs

To create four FCoE LIFs (two on each node), run the following commands.



Note: If you are not setting up access to FCoE application data LUNs in this tenant, do not create the FCoE LIFs.

```
network interface create -vserver App-A-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node <node01> -home-port 0e
```

```
network interface create -vserver App-A-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node <node01> -home-port 0g
```

```
network interface create -vserver App-A-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node <node02> -home-port 0e
```

```
network interface create -vserver App-A-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node <node02> -home-port 0g
```

```
network interface show -vserver App-A-SVM -lif fcp*
```

Create NFS LIF

To create an NFS LIF in the App-A SVM, run the following commands:

```
network interface create -vserver App-A-SVM -lif nfs_App_A_datastore_1 -role data -data-protocol nfs -home-node <node02> -home-port a0a-<storage-App-A-nfs-vlan-id> -address <nfs-lif-App-A_datastore_1-ip> -netmask <nfs-lif-App-A-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true
```

```
network interface show -vserver App-A-SVM -lif nfs*
```



Note: NetApp recommends creating a new LIF for each datastore.



Note: You are not creating a LIF for a swap datastore here. Testing revealed that if the swap datastore was on a VMkernel port on a DVS, the swap datastore would no longer be set in the ESXi host settings on reboot. NetApp recommends using the `infra_swap` datastore with a VMkernel port on vSwitch0.

Add Tenant SVM Administrator

To add the tenant SVM administrator and SVM administration LIF to the SVM, complete the following steps.

1. Run the following commands:

```
network interface create -vserver App-A-SVM -lif svm-mgmt -role data -data-protocol none -home-node <node01> -home-port a0a-<storage-App-A-svm-mgmt-vlan-id> -address <App-A-svm-mgmt-ip> -netmask <App-A-svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```



Note: The SVM management IP in this step should be from a subnet reachable from the Core-Services subnet. In this validation, a Supernet route with destination IP range 172.16.0.0/16 was put into each Core-Services device. An example use case for this would be to allow the tenant SVM to do DNS lookups from the Core-Services DNS servers.

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

```
network route create -vserver App-A-SVM -destination 0.0.0.0/0 -gateway <App-A-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 App-A-SVM
Enter a new password: <password>
Enter it again: <password>
```

```
security login unlock -username vsadmin -vserver App-A-SVM
```

Add Quality of Service (QoS) Policy to Monitor Application Workload

To add a storage QoS policy to monitor both the IOPs and bandwidth delivered from the APP-A-SVM, complete the following steps:

1. Create the QoS policy-group to measure the SVM output without an upper limit.

```
qos policy-group create -policy-group App-A -vserver App-A-SVM -max-throughput
INF
```

```
vserver modify -vserver App-A-SVM -qos-policy App-A
```

2. Monitor the QoS policy group output.

```
qos statistics performance show
```

Configure Cisco UCS for the Tenant

This section describes procedures for deploying Cisco UCS Servers for a tenant named App-A. It is assumed in this FlexPod Deployment that a tenant is most likely an application or group of applications. Because of this assumption, it is assumed that a new set of ESXi servers will be setup for the tenant in a separate ESXi cluster. It is also assumed in this implementation that the new ESXi servers will be booted from the storage Infrastructure SVM, although server boot could be moved into the tenant SVM.

In this section, required additions to Cisco UCS are detailed, including adding tenant iSCSI VLANs and adding these VLANs to the iSCSI vNICs if iSCSI is being provided by the tenant, adding a new SAN Connectivity Policy to zone any FCoE interfaces created in the tenant Storage SVM, creating additional Service Profile Templates if necessary, and generating the new Service Profiles for the ESXi hosts for the tenant. All procedures in this section are completed using the Cisco UCS Manager HTML 5 User Interface.

Add Tenant iSCSI VLANs

If iSCSI LUN access is being provided by the App-A tenant, the iSCSI VLANs must be added to the Cisco UCS LAN Cloud. To add tenant iSCSI VLANs, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `App-A-iSCSI-A` as the name of the VLAN to be used for the first iSCSI VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the VLAN ID for the first iSCSI VLAN in the UCS
8. Click OK, then OK.

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

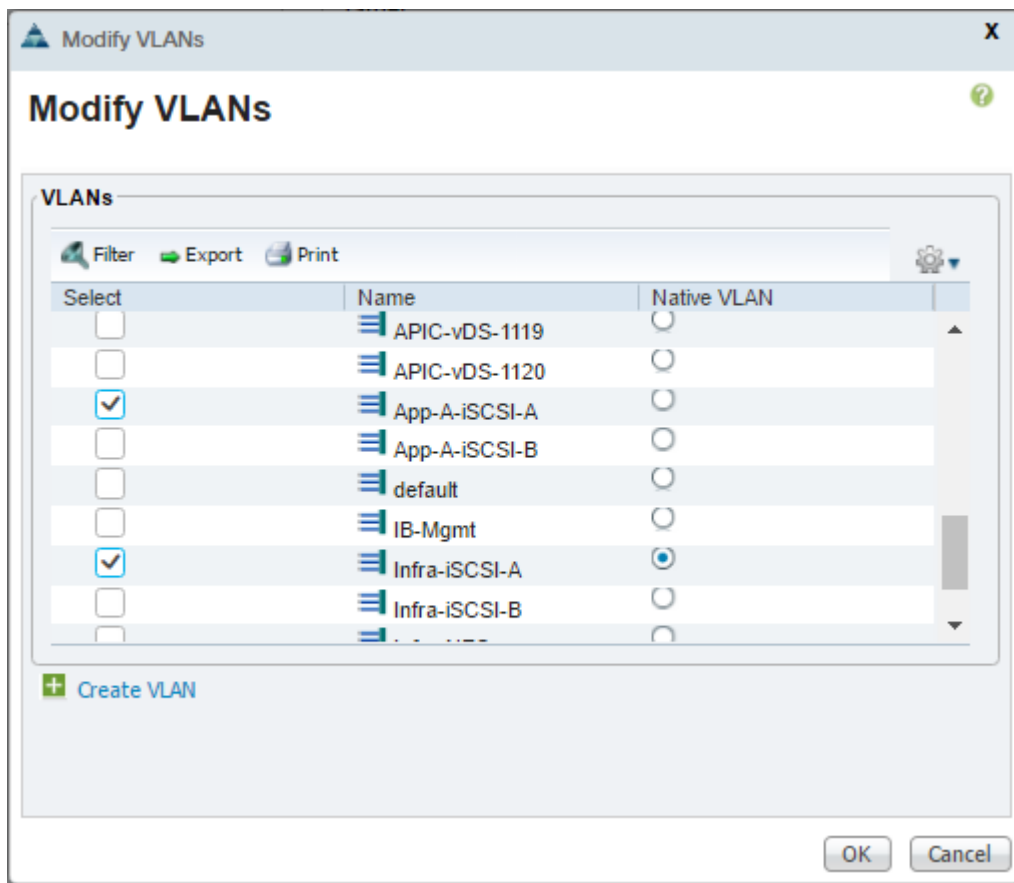
9. Right-click VLANs.

10. Select Create VLANs.
11. Enter App-A-iSCSI-B as the name of the VLAN to be used for the second iSCSI VLAN.
12. Keep the Common/Global option selected for the scope of the VLAN.
13. Enter the VLAN ID for the second iSCSI VLAN in the UCS.
14. Click OK, then OK.

Add Tenant iSCSI VLANs to iSCSI vNIC Templates

If iSCSI LUN access is being provided by the App-A tenant, the iSCSI VLANs must be added to the iSCSI vNIC Templates. To add tenant iSCSI VLANs to iSCSI vNIC templates, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > Policies > root > vNIC Templates > vNIC Template iSCSI-A.
3. Under Actions select Modify VLANs.
4. Using the checkbox, add the App-A-iSCSI-A VLAN to the template.



5. Click OK then OK again to complete adding the VLAN to the vNIC Template.
6. On the left, select vNIC Template iSCSI-B.

7. Under Actions select Modify VLANs.
8. Using the checkbox, add the App-A-iSCSI-B VLAN to the template.
9. Click OK then OK again to complete adding the VLAN to the vNIC Template.

Add Tenant SAN Connectivity Policy

If FCoE LUN access is being provided by the App-A tenant, a new SAN Connectivity Policy must be built to add zoning through a Storage Connection Policy for the storage FCoE LIFs in the Tenant SVM. Note that it is assumed that all servers are being SAN-booted from the Infrastructure SVM. It is not necessary to add boot targets to the Storage Connection Policy since boot targets from the boot policy are automatically zoned, only application LUN targets need to be added to the policy.

To add tenant SAN connectivity policy, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. On the left, select SAN > Policies > root > Storage Connection Policies
3. Right-click Storage Connection Policies.
4. Select Create Storage Connection Policy.
5. Enter App-A-FCoE-A as the name of the policy.
6. **Enter “Zone LUNs from Storage App-A-SVM Fabric-A” as the Description.**
7. Select the Single Initiator Multiple Targets Zoning Type.
8. Click the Add button to add the first Target.
9. Enter the WWPN for LIF fcp_lif01a in SVM App-A-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver App-A-SVM -lif fcp*”**.
10. Leave the Path set at A and select VSAN VSAN-A.
11. Click OK to complete creating the target.
12. Click the Add button to add the second Target.
13. Enter the WWPN for LIF fcp_lif02a in SVM Infra-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver App-A-SVM -lif fcp*”**.
14. Leave the Path set at A and select VSAN VSAN-A.
15. Click OK to complete creating the target.

Create Storage Connection Policy

Name :

Description :

Zoning Type : None Single Initiator Single Target Single Initiator Multiple Targets

FC Target Endpoints

Filter Export Print

WWPN	Path	VSAN
20:06:00:A0:98:5B:48:16	A	VSAN-A
20:08:00:A0:98:5B:48:16	A	VSAN-A

+ Add Delete Info

OK Cancel

16. Click OK then OK again to complete creating the Storage Connection Policy.

17. Right-click Storage Connection Policies.

18. Select Create Storage Connection Policy.

19. Enter App-A-FCoE-B as the name of the policy.

20. Enter **“Zone LUNs from Storage App-A-SVM Fabric-B”** as the Description.

21. Select the Single Initiator Multiple Targets Zoning Type.

22. Click the Add button to add the first Target.

23. Enter the WWPN for LIF fcp_lif01b in SVM App-A-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver App-A-SVM -lif fcp*”**.

24. Set the Path to B and select VSAN VSAN-B.

25. Click OK to complete creating the target.
26. Click the Add button to add the second Target.
27. Enter the WWPN for LIF fcp_lif02b in SVM App-A-SVM. To get this value, log into the storage cluster CLI and enter the command **“network interface show -vserver App-A-SVM -lif fcp*”**.
28. Set the Path to B and select VSAN VSAN-B.
29. Click OK to complete creating the target.
30. Click OK then OK again to complete creating the Storage Connection Policy.
31. On the left, right-click SAN Connectivity Policies and select Create SAN Connectivity Policy.
32. Enter App-A-FCoE as the name of the policy.
33. **Enter “Policy that Zones LUNs from Storage App-A-SVM” as the Description.**
34. Select the WWNN-Pool for WWNN Assignment.
35. Click the Add button to add a vHBA.
36. In the Create vHBA dialog box, enter Fabric-A as the name of the vHBA.
37. Select the Use vHBA Template checkbox.
38. In the vHBA Template list, select Fabric-A.
39. In the Adapter Policy list, select VMWare.
40. Click OK to add this vHBA to the policy.

Create vHBA

Name : Fabric-A

Use vHBA Template :

+ Create vHBA Template

vHBA Template : Fabric-A

Adapter Performance Profile

Adapter Policy : VMWare

+ Create Fibre Channel Adapter Policy

OK Cancel

41. Click the Add button to add another vHBA to the policy.
42. In the Create vHBA box, enter Fabric-B as the name of the vHBA.
43. Select the Use vHBA Template checkbox.
44. In the vHBA Template list, select Fabric-B.
45. In the Adapter Policy list, select VMWare.
46. Click OK to add the vHBA to the policy.

Create SAN Connectivity Policy

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 Fabric-B	Derived
▶ -vHBA Fabric-A	Derived

47. Click OK then OK again to complete creating the SAN Connectivity Policy.
48. In the list on the left under SAN Connectivity Policies, select the App-A-FCoE Policy.
49. In the center pane, select the vHBA Initiator Groups tab.
50. Select Add to add a vHBA Initiator Group.
51. Name the vHBA Initiator Group Fabric-A and select the Fabric-A vHBA Initiators.
52. Select the App-A-FCoE-A Storage Connection Policy.

Create vHBA Initiator Group



Create vHBA Initiator Group

vHBA Initiator Group

Name :

Description :

Select vHBA Initiators

Select	Name
<input checked="" type="checkbox"/>	 Fabric-A
<input type="checkbox"/>	 Fabric-B

Storage Connection Policy:

[+ Create Storage Connection Policy](#)

Global Storage Connection Policy

Global storage connection policy defined under org is assigned to this vHBA initiator group.

OK Cancel

Storage Connection Policy: **App-A-FCoE-A**

[+ Create Storage Connection Policy](#)

Global Storage Connection Policy

Global storage connection policy defined under org is assigned to this vHBA initiator group.

Properties

Storage Connection Policy : **App-A-FCoE-A**
 Description : **Zone LUNs from Storage App-A-SVM Fabric-A**
 Zoning Type : **Single Initiator Multiple Targets**

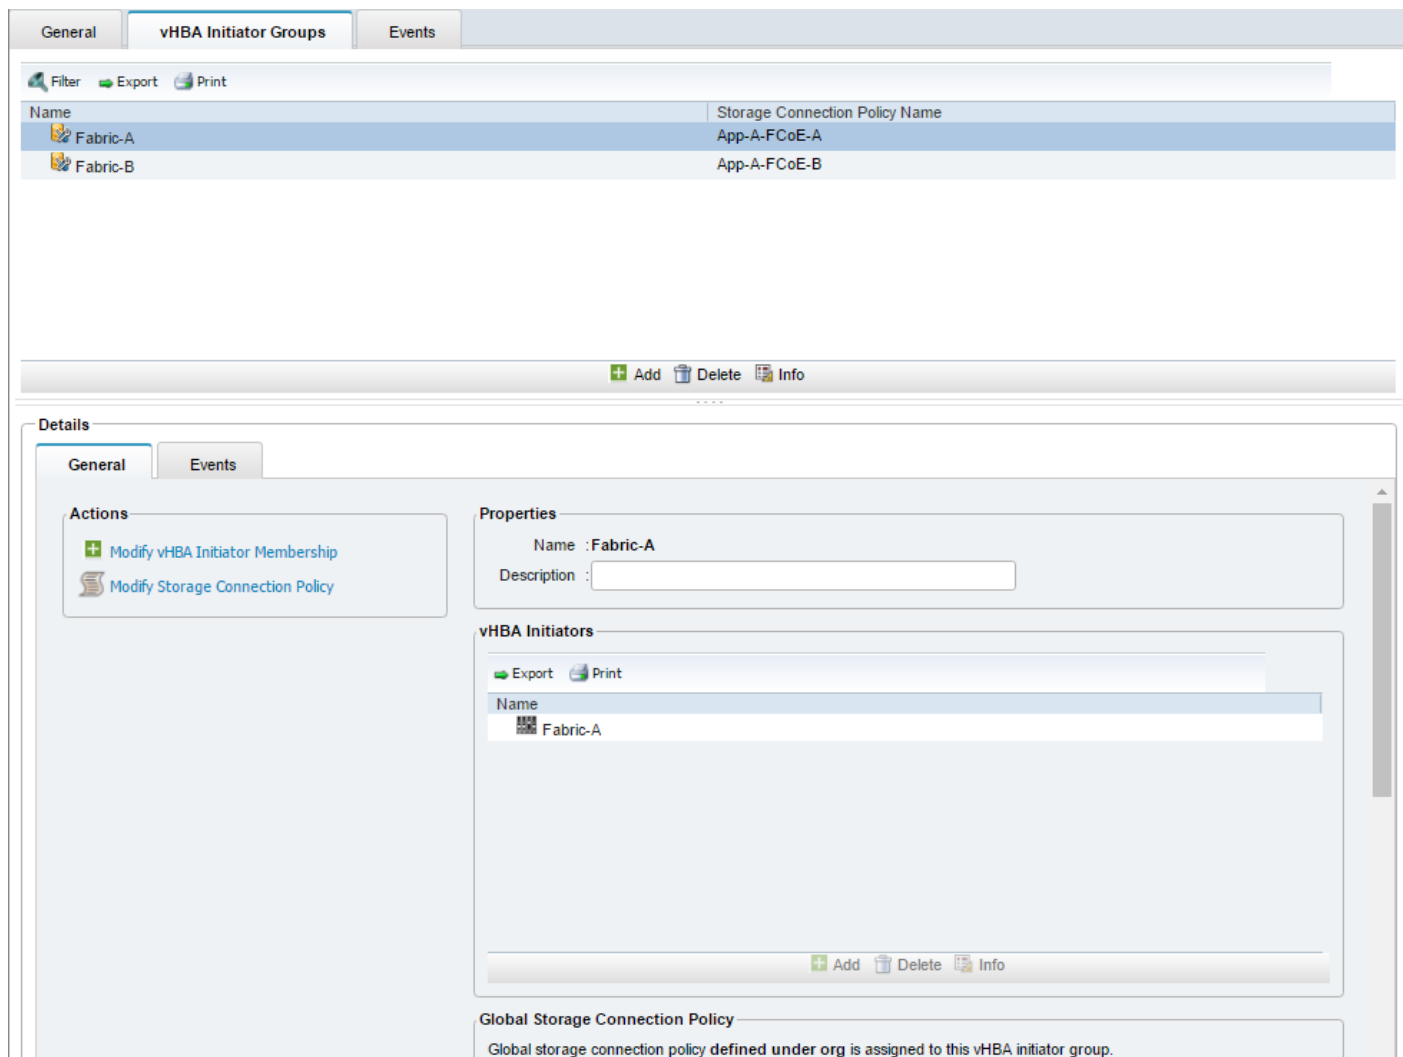
FC Target Endpoints

Filter Export Print

WWPN	Path	VSAN
20:06:00:A0:98:5B:48:16	A	VSAN-A
20:08:00:A0:98:5B:48:16	A	VSAN-A

OK Cancel

53. Click OK then OK again to add this vHBA Initiator Group.
54. Select Add to add a vHBA Initiator Group.
55. Name the vHBA Initiator Group Fabric-B and select the Fabric-B vHBA Initiators.
56. Select the App-A-FCoE-B Storage Connection Policy.
57. Click OK then OK again to add this vHBA Initiator Group.



Create Application-Specific Service Profile Templates

It is recommended to create new Service Profile Templates for the servers running the applications in the new tenant. These templates can be created by cloning the existing Service Profile Templates and modifying them with any necessary changes. Since the tenant-specific iSCSI VLANs were added to the iSCSI vNIC templates and any iSCSI-booted servers will continue to boot from LUNs in the Infrastructure Storage SVM, no changes to the LAN configuration and LAN Connectivity Policy are needed. If FCoE storage is being provided by the tenant, the new Storage Connectivity Policy should replace the Infrastructure Storage Connectivity Policy in the Service Profile Template. Since the FCoE boot targets specified in the boot policy are automatically zoned, replacing the SAN Connectivity Policy has no effect on FCoE boot.

Add New Application-Specific Server Pool

Since new Service Profile Templates for the servers running the applications in the new tenant are being created, a new tenant-specific server pool can be created and mapped in the new Service Profile Templates. Create this pool and map it in the new Service Profile Templates.

Create New Service Profiles for Application-Specific Servers

Using the cloned and modified Application-Specific Service Profile Templates, create Service Profiles associated to servers for the new tenant.

Gather Necessary Information

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

Table 26 iSCSI LIFs for iSCSI IQN.

Vserver	iSCSI Target IQN
Infra-SVM	
App-A-SVM	



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

Table 27 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN
VM-Host-App-A-01	
VM-Host-App-A-02	



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

Table 28 Table 7 vHBA WWPNs for Fabric A and Fabric B

Cisco UCS Service Profile Name	WWPN
VM-Host-App-A-01	Fabric-A
	Fabric-B
VM-Host-App-A-02	Fabric-A
	Fabric-B



Note: To gather the vHBA WWPN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root. Click each service profile and **then click the vHBAs.** The WWPNs are shown in the center pane.

Configure Storage SAN Boot for the Tenant

This section describes procedures for setting up SAN boot for the tenant ESXi Host servers.

Clustered VMware ESXi Boot LUNs in Infra-SVM

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

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

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

```
lun show
```

Clustered Data ONTAP iSCSI Boot Storage Setup

Create igroups

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

```
igroup create -vserver Infra-SVM -igroup VM-Host-App-A-01 -protocol iscsi -ostype
vmware -initiator <vm-host-App-A-01-iqn>
```

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol iscsi -ostype
vmware -initiator <vm-host-App-A-02-iqn>
```



Note: Use the values listed in Table 8 and Table 9 for the IQN information.



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

Clustered Data ONTAP FCoE Boot Storage Setup

Create igroups

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

```
igroup create -vserver Infra-SVM -igroup VM-Host-App-A-01 -protocol fcp -ostype
vmware -initiator <vm-host-App-A-01-fabric-a-wwpn>,<vm-host-App-A-01-fabric-b-
wwpn>
```

```
igroup create -vserver Infra-SVM -igroup VM-Host-App-A-02 -protocol fcp -ostype
vmware -initiator <vm-host-App-A-02-fabric-a-wwpn>,<vm-host-App-A-02-fabric-b-
wwpn>
```



Note: Use the values listed in Table 10 for the WWPN information.



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

Map Boot LUNs to igroups

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

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

lun show -m
```

Deploy ACI Application (App-A) Tenant

This section details the steps for creation of the App-A Sample Tenant in the ACI Fabric. This tenant will host application connectivity between the compute (VMware on UCS) and the storage (NetApp) environments. This tenant will also host the three application tiers of the sample three-tier application. A corresponding App-A-SVM has already been created on the NetApp storage to align with this tenant. To deploy the App-A Tenant, complete the following steps:

1. In the APIC Advanced GUI, select Fabric > Access Policies.
2. On the left, expand Pools and VLAN.
3. Select VLAN Pool VP-NTAP.
4. In the center pane, click the + sign to add an encapsulation block.
5. Enter <storage-App-A-NFS-VLAN> for the From and To fields.
6. Select Static Allocation.

Create Ranges i X

Specify the Encap Block Range

Type: **VLAN**

Range: 3051 - 3051

From To

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

SUBMIT
CANCEL

7. Click SUBMIT to add the Encap Block Range.
8. In the center pane, click the + sign to add an encapsulation block.

9. Enter <storage-App-A-SVM-MGMT-VLAN> for the From and To fields.
10. Select Static Allocation.
11. Click SUBMIT to add the Encap Block Range.
12. If iSCSI LUN access is being provided by the App-A tenant, complete steps 13-29. Otherwise, continue at step 30.
13. In the center pane, click the + sign to add an encapsulation block.
14. Enter <storage-App-A-iSCSI-A-VLAN> for the From and To fields.
15. Select Static Allocation.
16. Click SUBMIT to add the Encap Block Range.
17. In the center pane, click the + sign to add an encapsulation block.
18. Enter <storage-App-A-iSCSI-B-VLAN> for the From and To fields.
19. Select Static Allocation.
20. Click SUBMIT to add the Encap Block Range.
21. On the left, select VLAN Pool VP-UCS.
22. In the center pane, click the + sign to add an encapsulation block.
23. Enter <ucs-App-A-iSCSI-A-VLAN> for the From and To fields.
24. Select Static Allocation.
25. Click SUBMIT to add the Encap Block Range.
26. In the center pane, click the + sign to add an encapsulation block.
27. Enter <ucs-App-A-iSCSI-B-VLAN> for the From and To fields.
28. Select Static Allocation.
29. Click SUBMIT to add the Encap Block Range.
30. At the top select Tenants > Add Tenant.
31. Name the Tenant App-A. Select the default Monitoring Policy.
32. For the VRF Name, also enter App-A. Leave the Take me to this tenant when I click finish checkbox checked.

Create Tenant
i X

Specify tenant details

Name:

Description:

Tags:

Monitoring Policy:

Security Domains:

Select	Name	Description

VRF Name:

Take me to this tenant when I click finish

33. Click SUBMIT to finish creating the Tenant.
34. If you are using providing iSCSI LUN access from this tenant, complete steps 35-68. Otherwise, continue to step 69.
35. On the left under Tenant App-A, right-click Application Profiles and select Create Application Profile.
36. Name the Application Profile iSCSI, select the default Monitoring Policy, and click SUBMIT to complete adding the Application Profile.
37. On the left, expand Application Profiles and iSCSI.
38. Right-click Application EPGs and select Create Application EPG.
39. Name the EPG iSCSI-A. Leave Intra EPG Isolation Unenforced.
40. Use the Bridge Domain drop-down to select Create Bridge Domain.
41. Name the Bridge Domain BD-iSCSI-A.
42. Select the App-A/App-A VRF.
43. Use the Forwarding drop-down to select Custom.

44. Select Flood for the L2 Unknown Unicast and default for the End Point Retention Policy and IGMP Snoop Policy.

Create Bridge Domain

1. Main 2. L3 Configurations 3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name:

Description:

VRF:

Forwarding:

L2 Unknown Unicast:

L3 Unknown Multicast Flooding:

Multi Destination Flooding:

End Point Retention Policy:
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy:

PREVIOUS NEXT CANCEL

45. At the bottom right, click NEXT.

46. Make sure Unicast Routing is Enabled and click NEXT.

47. Select the default Monitoring Policy and click FINISH.

Create Application EPG

1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name: iSCSI-A

Description: optional

Tags:
enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: App-A/BD-iSCSI-A

Monitoring Policy: default

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

48. Select the default Monitoring Policy and click FINISH to complete creating the EPG.
49. On the left, expand Application EPGs and EPG iSCSI-A. Right-click Domains and select Add Physical Domain Association.
50. Using the drop-down, select the PD-NTAP Physical Domain Profile.
51. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: PD-NTAP

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

SUBMIT CANCEL

52. Click SUBMIT to complete the Physical Domain Association.

53. Repeat steps 49-52 to add the PD-UCS Physical Domain Association.



Note: In this deployment for iSCSI, we are adding both the NetApp LIF endpoints and the VMware VMkernel (VMK) endpoints in a single EPG. This method allows unrestricted communication within the EPG. We also had the choice to put the LIFs in one EPG and the VMKs in a second EPG and connect them with a filtered contract.

Domain Profile	Domain Type	Deployment Immediacy	Resolution Immediacy	State	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Allow Micro-Segmentation
PD-NTAP	Physical Domain	Immediate	Immediate	formed			False
PD-UCS	Physical Domain	Immediate	Immediate	formed			False

54. Right-click Static-Bindings (Paths) and select Deploy EPG on PC, VPC, or Interface.

55. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.

56. Using the Path drop-down, select the VPC for NetApp Storage Controller 01.

57. Enter `vlan-<storage-App-A-iSCSI-A-VLAN>` for Port Encap.

58. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface
i
✕

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ▼ 📄

Primary VLAN: For example, vlan-1

Port Encap: For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

59. Click SUBMIT to complete adding the Static Path Mapping.

60. Repeat steps 54-59 to add the Static Path Mapping for NetApp Storage Controller 02.

61. Right-click Static-Bindings (Paths) and select Deploy EPG on PC, VPC, or Interface.

62. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.

63. Using the Path drop-down, select the VPC for UCS Fabric Interconnect A.

64. Enter `vlan-<ucs-App-A-iSCSI-A-VLAN>` for Port Encap.

65. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel **Virtual Port Channel**

Path:

Primary VLAN: For example, vlan-1

Port Encap: For example, vlan-1

Deployment Immediacy: **Immediate** On Demand

Mode: **Trunk** Access (802.1P) Access (Untagged)

SUBMIT CANCEL

66. Click SUBMIT to complete adding the Static Path Mapping.

67. Repeat steps 61-66 to add the Static Path Mapping for UCS Fabric Interconnect B.

The screenshot shows the Cisco ACI GUI for Tenant App-A. The left sidebar shows the navigation tree with 'Static Bindings (Paths)' selected. The main content area displays a table of static bindings for Node-101-102.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-3111	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-3111	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-01		vlan-3011	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-3011	Immediate	Trunk

68. Repeat steps 38-67 to build the iSCSI-B EPG. Make sure to create a separate Bridge Domain for this EPG and use the App-A iSCSI-B VLAN IDs.

The screenshot shows the Cisco ACI GUI for Tenant App-A. The left navigation pane is expanded to show 'Static Bindings (Paths)'. The main content area displays a table of static bindings for Node-101-102.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a		vlan-3121	Immediate	Trunk
Node-101-102/VPC-a01-6248-b		vlan-3121	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-01		vlan-3021	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-3021	Immediate	Trunk

69. On the left, under Tenant App-A, right-click Application Profiles and select Create Application Profile.
70. Name the Profile `NFS`, select the default Monitoring Policy, and click SUBMIT.
71. Right-click the NFS Application Profile and select Create Application EPG.
72. Name the EPG `NFS-LIF` and leave Intra EPG Isolation set at Unenforced.
73. Use the Bridge Domain drop-down to select Create Bridge Domain.
74. Name the Bridge Domain `BD-NFS` and select the App-A/App-A VRF.



Note: It is important to create a new Bridge Domain for each traffic VLAN coming from the NetApp Storage Controllers. All of the VLAN interfaces on a given NetApp Interface Group share the same MAC address, and separating to different bridge domains in the ACI Fabric allows all the traffic to be forwarded properly.

75. For Forwarding, select Custom and select Flood for L2 Unknown Unicast. Select default for the End Point Retention Policy and the IGMP Snoop Policy.

Create Bridge Domain

STEP 1 > Main 1. Main 2. L3 Configurations 3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name: **BD-NFS**

Description:

VRF: **App-A/App-A**

Forwarding: **Custom**

L2 Unknown Unicast: **Flood**

L3 Unknown Multicast Flooding: **Flood**

Multi Destination Flooding: **Flood in BD**

End Point Retention Policy: **default**
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: **default**

76. At the bottom right, click NEXT.

77. Make sure Unicast Routing is enabled and click NEXT.

78. Select the default Monitoring Policy and click FINISH.

Create Application EPG

1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name: NFS-LIF

Description: optional

Tags: enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: App-A/BD-NFS

Monitoring Policy: select a value

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

PREVIOUS FINISH CANCEL

79. Select the default Monitoring Policy and click FINISH to complete creating the EPG.

80. On the left expand NFS, Application EPGs, and EPG NFS-LIF.

81. Right-click Domains and select Add Physical Domain Association.

82. Select the PD-NTAP Physical Domain Profile.

83. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: PD-NTAP

Deploy Immediacy: **Immediate** On Demand

Resolution Immediacy: **Immediate** On Demand Pre-provision

SUBMIT **CANCEL**

84. Click SUBMIT to complete adding the Physical Domain Association.
85. Right-click Static Bindings (Paths) and select Deploy Static EPG on PC, VPC, or Interface.
86. Select the Virtual Port Channel Path Type.
87. Using the Path drop-down, select the VPC for NetApp Storage Controller 01.
88. For Port Encap, enter `vlan-<storage-App-A-NFS-VLAN>`.
89. Select Immediate for Deployment Immediacy and Trunk for Mode.

i X

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ▼ 📄

Primary VLAN:

For example, vlan-1

Port Encap:

For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

90. Click SUBMIT to finish adding the EPG Static Binding.

91. Repeat steps 85-90 for the Static Path to NetApp Storage Controller 02.

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-aff8040-01		vlan-3051	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-3051	Immediate	Trunk

92. On the left under EPG NFS-LIF, right-click Contracts and select Add Provided Contract.

93. In the Add Provided Contract window, use the Contract drop-down to select Create Contract.

94. Name the contract Allow-NFS. Leave the Scope set at VRF.

95. Click the + sign to add a Contract Subject.

96. Name the subject Allow-NFS.

97. Click the + sign to add a Filter to the Filter Chain.

98. Click the drop-down and select NTAP-NFS-v3 from Tenant common.

99. Click UPDATE.

Create Contract Subject

Specify Identity Of Subject

Name: Allow-NFS

Description: optional

Target DSCP: unspecified

Apply Both Directions:

Reverse Filter Ports:

Filter Chain

Filters
Name
common/NTAP-NFS-v3

L4-L7 SERVICE GRAPH
Service Graph: select an option

PRIORITY
QoS:

OK CANCEL



Note: Optionally, add ICMP to the filter chain to allow ping in this contract for troubleshooting purposes.

100. Click OK to complete the Contract Subject.

Create Contract

Specify Identity Of Contract

Name:

Scope:

QoS Class:

Target DSCP:

Description:

Subjects:

Name	Description
Allow-NFS	

101. Click SUBMIT to complete creating the Contract.

Add Provided Contract

Select a contract

Contract:

QoS:

Contract Label:

Subject Label:

102. Click SUBMIT to complete Adding the Provided Contract.
103. Right-click Application EPGs under the NFS Application Profile and select Create Application EPG.
104. Name the EPG NFS-VMK and leave Intra EPG Isolation set at Unenforced.
105. Use the Bridge Domain drop-down to select App-A/BD-NFS. Select the default Monitoring Policy.

Create Application EPG

STEP 1 > Identity

1. Identity

Specify the EPG Identity

Name: NFS-VMK

Description: optional

Tags:

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: App-A/BD-NFS

Monitoring Policy: default

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

PREVIOUS FINISH CANCEL

106. Click FINISH to complete creating the EPG.
107. On the left expand NFS, Application EPGs, and EPG NFS-VMK.
108. Under EPG NFS-VMK, right-click Domains and select Add VMM Domain Association.
109. Select the VMM Domain Profile for the DVS you have installed.

110. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. If adding a VMware vDS association, select the Dynamic VLAN Mode.

Add VMM Domain Association i X

Choose the VMM domain to associate

VMM Domain Profile: VMware/vc-AVS ▼ [icon]

Deploy Immediacy: **Immediate** On Demand

Resolution Immediacy: **Immediate** On Demand Pre-provision

Port Encap: _____
For example, vlan-1

SUBMIT **CANCEL**

Add VMM Domain Association

Choose the VMM domain to associate

VMM Domain Profile: VMware/vc-vDS

Deploy Immediacy: **Immediate** On Demand

Resolution Immediacy: **Immediate** On Demand Pre-provision

VLAN Mode: **Dynamic** Static

Allow Micro-Segmentation:

Allow Promiscuous: Reject

Forged Transmits: Reject

MAC Changes: Reject

SUBMIT **CANCEL**

111. Click SUBMIT to complete adding the VMM Domain Association.
112. On the left under EPG NFS-VMK, right-click Contracts and select Add Consumed Contract.
113. In the Add Consumed Contract window, use the Contract drop-down to select App-A/Allow-NFS.

Add Consumed Contract

Select a contract

Contract: App-A/Allow-NFS

QoS: Unspecified

Contract Label: _____

Subject Label: _____

SUBMIT **CANCEL**

114. Click SUBMIT to complete adding the Consumed Contract.
115. On the left, under Tenant App-A, right-click Application Profiles and select Create Application Profile.
116. Name the Profile `svm-MGMT`, select the default Monitoring Policy, and click SUBMIT.
117. Right-click the SVM-MGMT Application Profile and select Create Application EPG.
118. Name the EPG `App-A-svm-MGMT` and leave Intra EPG Isolation set at Unenforced.
119. Use the Bridge Domain drop-down to select create Bridge Domain.
120. Name the Bridge Domain `BD-Internal` and select the `App-A/App-A` VRF.
121. Leave Forwarding set at optimize, select the default End Point Retention Policy and IGMP Snoop Policy.
122. Click NEXT.
123. Make sure Unicast Routing is enabled and click NEXT.
124. Select the default Monitoring Policy and click FINISH to complete creating the Bridge Domain.

Create Application EPG

STEP 1 > Identity

1. Identity

Specify the EPG Identity

Name: App-A-SVM-MGMT

Description: optional

Tags:
enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Intra EPG Isolation: Enforced Unenforced

Bridge Domain: App-A/BD-Internal

Monitoring Policy: default

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

PREVIOUS FINISH CANCEL

125. Select the default Monitoring Policy and click FINISH to complete creating the EPG.
126. On the left expand SVM-MGMT, Application EPGs, and EPG App-A-SVM-MGMT.
127. Right-click Domains and select Add Physical Domain Association.
128. Select the PD-NTAP Physical Domain Profile.
129. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: PD-NTAP

Deploy Immediacy: **Immediate** On Demand

Resolution Immediacy: **Immediate** On Demand Pre-provision

SUBMIT **CANCEL**

130. Click SUBMIT to complete adding the Physical Domain Association.
131. Right-click Static Bindings (Paths) and select Deploy Static EPG on PC, VPC, or Interface.
132. Select the Virtual Port Channel Path Type.
133. Using the Path drop-down, select the VPC for NetApp Storage Controller 01.
134. For Port Encap, enter `vlan-<storage-App-A-SVM-MGMT-VLAN>`.
135. Select Immediate for Deployment Immediacy and Trunk for Mode.

i X

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path: ▼ 📄

Primary VLAN: For example, vlan-1

Port Encap: For example, vlan-1

Deployment Immediacy: Immediate On Demand

Mode: Trunk Access (802.1P) Access (Untagged)

SUBMIT
CANCEL

136. Click SUBMIT to finish adding the EPG Static Binding.

137. Repeat steps 131-136 for the Static Path Mapping to NetApp Storage Controller 02.

CISCO
System Tenants Fabric VM Networking L4-L7 Services Admin Operations
welcome, admin

ALL TENANTS | Add Tenant | Search:
common | App-A | Foundation | mgmt | infra

Tenant App-A

- Quick Start
- Tenant App-A
 - Application Profiles
 - NFS
 - SVM-MGMT
 - Application EPGs
 - EPG App-A-SVM-MGMT
 - Domains (VMs and Bare-Meta..)
 - Static Bindings (Paths)
 - Static Bindings (Leaves)
 - Contracts
 - Static EndPoint
 - Subnets

Static Bindings (Paths)

Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-aff8040-01		vlan-264	Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02		vlan-264	Immediate	Trunk

138. On the left under EPG App-A-SVM-MGMT, right-click Contracts and select Add Provided Contract.

139. In the Add Provided Contract window, use the Contract drop-down to select Create Contract.
140. Name the contract `Allow-SVM-MGMT`. Leave the Scope set at VRF.
141. Click the + sign to add a Contract Subject.
142. Name the subject `Allow-All`. Click the + sign to add a filter.
143. Use the drop-down to select the Allow-All filter from Tenant common. Click UPDATE.
144. Click OK to complete adding the Contract Subject.

Create Contract

Specify Identity Of Contract

Name: `Allow-SVM-MGMT`

Scope: `VRF`

QoS Class: `Unspecified`

Target DSCP: `unspecified`

Description: `optional`

Subjects:

Name	Description
<code>Allow-All</code>	

SUBMIT CANCEL

145. Click SUBMIT to complete creating the Contract.
146. Click SUBMIT to complete adding the Provided Contract.
147. On the left under EPG `App-A-SVM-MGMT`, right-click Subnets and select Create EPG Subnet.

148. For the Default Gateway IP, enter the gateway IP address and mask that was entered earlier in the Tenant Storage Deployment for the App-A tenant.
149. If this EPG will be connected to Core-Services by contract, select only the Shared between VRFs scope. Otherwise, if the tenant SVM management interface will only be accessed from EPGs within the tenant, leave only the Private to VRF Scope selected.

Create EPG Subnet i X

Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 Querier IP

ND RA Prefix policy:

150. Click SUBMIT to complete adding the EPG subnet.
151. On the left, right-click Application Profiles and select Create Application Profile.
152. Name the Application Profile `Three-Tier-App` and select the default Monitoring Policy.
153. Click SUBMIT to complete creating the Application Profile.
154. Expand `Three-Tier-App`, right-click Application EPGs under `Three-Tier-App` and select Create Application EPG.
155. Name the EPG `web` and leave Intra EPG Isolation set at Unenforced.
156. Use the Bridge Domain drop-down to select `App-A/BD-Internal`. Select the default Monitoring Policy.

Create Application EPG
i X

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation:

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

157. Click FINISH to complete creating the EPG.
158. On the left expand Three-Tier-App, Application EPGs, and EPG Web.
159. Under EPG Web, right-click Domains and select Add VMM Domain Association.
160. Select the VMM Domain Profile for the DVS you have installed.
161. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. If adding a VMware vDS association, select the Dynamic VLAN Mode.
162. Click SUBMIT to complete adding the VMM Domain Association.
163. On the left under EPG Web, right-click Contracts and select Add Provided Contract.
164. In the Add Provided Contract window, use the Contract drop-down to select Create Contract.
165. Name the Contract `Allow-Web-App`. Select the Application Profile Scope.
166. Click the + sign to add a Contract Subject.

167. Name the subject Allow-All.
168. Click the + sign to add a Contract filter.
169. Use the drop-down to select the Allow-All filter from Tenant common. Click UPDATE.
170. Click OK to complete creating the Contract Subject.

Create Contract

Specify Identity Of Contract

Name:

Scope:

QoS Class:

Target DSCP:

Description:

Subjects: x +

Name	Description
Allow-All	

171. Click SUBMIT to complete creating the Contract.
172. Click SUBMIT to complete adding the Provided Contract.
173. Right-click Contracts and select Add Consumed Contract.
174. In the Add Consumed Contract window, use the Contract drop-down to select the common/Allow-Shared-L3-Out contract.
175. Click SUBMIT to complete adding the Consumed Contract.

176. Optionally, repeat steps 173-175 to add the common/Allow-Core-Services Consumed Contract.
177. On the left under EPG Web, right-click Subnets and select Create EPG Subnet.
178. For the Default Gateway IP, enter a gateway IP address and mask from a subnet in the Supernet (172.16.0.0/16) that was set up for assigning Tenant IP addresses.
179. For scope, select Advertised Externally and Shared between VRFs.

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP: 172.16.0.254/24
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 Querier IP

ND RA Prefix policy:

SUBMIT **CANCEL**

180. Click SUBMIT to complete creating the EPG Subnet.
181. Right-click Application EPGs under Three-Tier-App and select Create Application EPG.
182. Name the EPG `App` and leave Intra EPG Isolation set at Unenforced.
183. Use the Bridge Domain drop-down to select `App-A/BD-Internal`. Select the default Monitoring Policy.

i X
Create Application EPG

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation: Enforced Unenforced

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

184. Click FINISH to complete creating the EPG.
185. On the left expand Three-Tier-App, Application EPGs, and EPG App.
186. Under EPG Web, right-click Domains and select Add VMM Domain Association.
187. Select the VMM Domain Profile for the DVS you have installed.
188. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. If adding a VMware vDS association, select the Dynamic VLAN Mode.
189. Click SUBMIT to complete adding the VMM Domain Association.
190. On the left under EPG App, right-click Contracts and select Add Provided Contract.
191. In the Add Provided Contract window, use the Contract drop-down to select Create Contract.
192. Name the Contract Allow-App-DB. Select the Application Profile Scope.
193. Click the + sign to add a Contract Subject.

194. Name the subject `Allow-All`.
195. Click the + sign to add a Contract filter.
196. Use the drop-down to select the Allow-All filter from Tenant common. Click UPDATE.
197. Click OK to complete creating the Contract Subject.
198. Click SUBMIT to complete creating the Contract.
199. Click SUBMIT to complete adding the Provided Contract.
200. Right-click Contracts and select Add Consumed Contract.
201. In the Add Consumed Contract window, use the Contract drop-down to select the App-A/Allow-Web-App contract.
202. Click SUBMIT to complete adding the Consumed Contract.
203. Optionally, repeat steps 200-202 to add the common/Allow-Core-Services Consumed Contract.
204. On the left under EPG App, right-click Subnets and select Create EPG Subnet.
205. For the Default Gateway IP, enter a gateway IP address and mask from a subnet in the Supernet (172.16.0.0/16) that was set up for assigning Tenant IP addresses.
206. If this EPG was connected to Core-Services by contract, select only the Shared between VRFs scope. Otherwise, if the tenant SVM management interface will only be accessed from EPGs within the tenant, leave only the Private to VRF Scope selected.

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 Querier IP

ND RA Prefix policy:

207. Click SUBMIT to complete creating the EPG Subnet.
208. Right-click Application EPGs under Three-Tier-App and select Create Application EPG.
209. Name the EPG DB and leave Intra EPG Isolation set at Unenforced.
210. Use the Bridge Domain drop-down to select App-A/BD-Internal. Select the default Monitoring Policy.

Create Application EPG
i X

STEP 1 > Identity
1. Identity

Specify the EPG Identity

Name:

Description:

Tags:

QoS class:

Custom QoS:

Intra EPG Isolation: Enforced Unenforced

Bridge Domain:

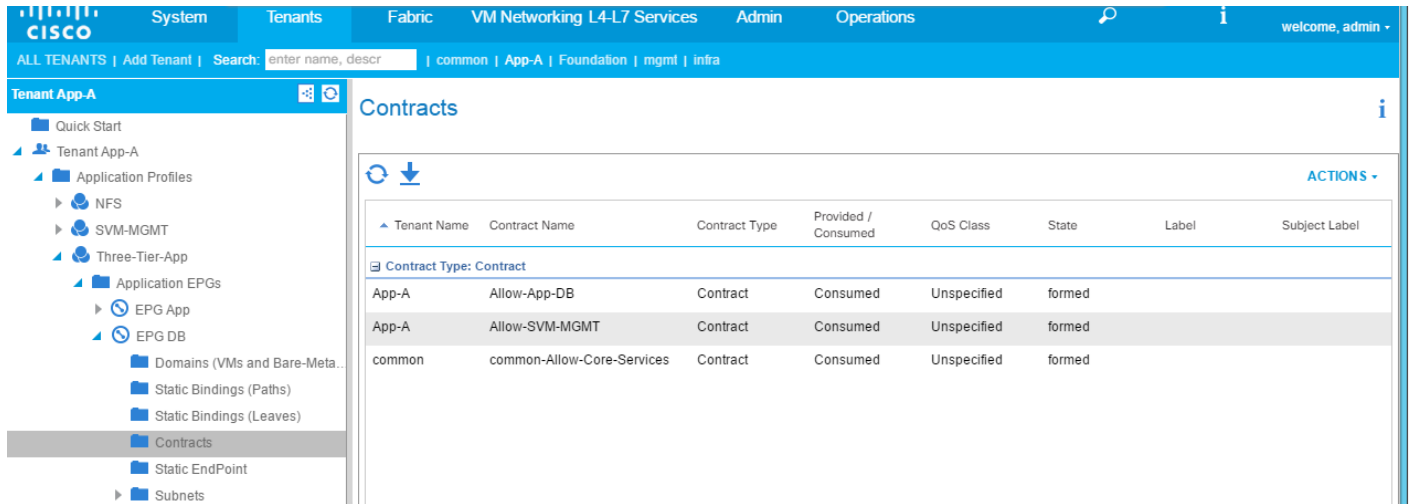
Monitoring Policy:

Associate to VM Domain Profiles:

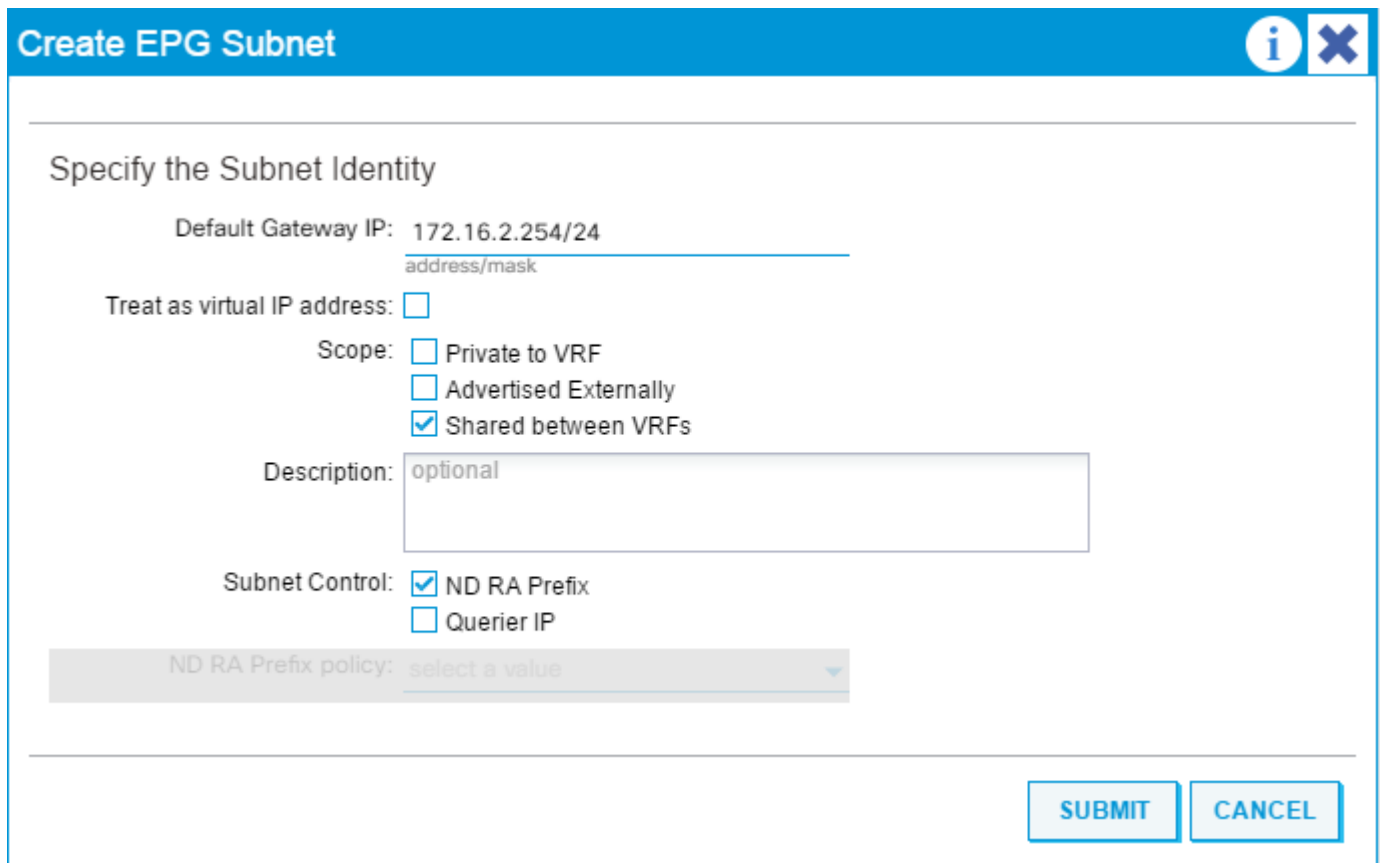
Statically Link with Leaves/Paths:

211. Click FINISH to complete creating the EPG.
212. On the left expand Three-Tier-App, Application EPGs, and EPG DB.
213. Under EPG DB, right-click Domains and select Add VMM Domain Association.
214. Select the VMM Domain Profile for the DVS you have installed.
215. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. If adding a VMware vDS association, select the Dynamic VLAN Mode.
216. Click SUBMIT to complete adding the VMM Domain Association.
217. On the left under EPG DB, right-click Contracts and select Add Consumed Contract.
218. In the Add Consumed Contract window, use the Contract drop-down to select the App-A/Allow-App-DB contract.
219. Click SUBMIT to complete adding the Consumed Contract.

- Repeat steps 217-219 to add the common/Allow-Core-Services and App-A/Allow-SVM-MGMT Consumed Contracts.



- On the left under EPG DB, right-click Subnets and select Create EPG Subnet.
- For the Default Gateway IP, enter a gateway IP address and mask from a subnet in the Supernet (172.16.0.0/16) that was set up for assigning Tenant IP addresses.
- Select only the Shared between VRFs scope.



- Click SUBMIT to complete creating the EPG Subnet.

Install and Configure VMware ESXi on Tenant Hosts

To install and configure VMware ESXi on tenant hosts, complete the following steps:

1. Using Virtual Media Enabled Service Profiles, follow the section of this document titled VMware vSphere 6.0 U1b Setup to install and configure the App-A tenant ESXi hosts. Configure the tenant ESXi hosts in the same IB-MGMT subnet as the Infrastructure hosts. Also, add a VMkernel port on vSwitch0 for Infra-NFS to allow the infra_swap datastore to be mounted as the ESXi swap datastore.
2. The main difference in the installation will be that the tenant iSCSI VMkernel Interfaces need to be installed as tagged VLAN port groups on the iSCSI vSwitches and the App-A iSCSI LIF IPs will need to be entered as dynamic targets in the VMware Software iSCSI Initiator. If using iSCSI boot, leave the Infrastructure iSCSI port groups set as untagged VLAN interfaces. The tenant NFS VMkernel Interfaces need to be installed on the DVS after the ESXi host is added to the DVS. If the NFS interfaces are placed on the Cisco AVS, set the MTU of these interfaces to 8950 instead of 9000. Add the new hosts to a new cluster in vCenter, then add the ESXi hosts to the DVS. Mount the infra_swap datastore and the datastore configured in the App-A SVM to the ESXi hosts. If necessary, configure the hosts to core dump to the vCenter ESXi dump collector.
3. Using NetApp VSC, configure Optimal Storage Settings on the VMware ESXi hosts and install the NetApp NFS VAAI Plugin.
4. You are now ready to install software into the tenant and begin running a workload. Because FCoE zoning or iSCSI targets to the App-A-SVM LIFs are in place, SAN storage can be provisioned with NetApp VSC. NFS storage can also be provisioned with VSC, but remember to create a LIF for the NFS datastore on the node where the datastore will be placed before provisioning the datastore with VSC. VSC will use the least used LIF on the node for the datastore.

Build a Second Tenant (Optional)

In this lab validation, a second tenant was built to demonstrate that multiple tenants could access and use the Shared-L3-Out and that tenants can be completely logically separated, but can also have overlapping IP address spaces. The second tenant built in this lab validation had the following characteristics and was built with the same contract structure as the App-A tenant:

Table 29 Lab Validation Tenant App-B Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI-A	3012	3112	192.168.111.0/24 - L2	BD-iSCSI-A
iSCSI-B	3022	3122	192.168.121.0/24 - L2	BD-iSCSI-B
NFS-LIF	3052	N/A	192.168.151.0/24 - L2	BD-NFS
NFS-VMK	N/A	DVS	192.168.151.0/24 - L2	BD-NFS
SVM-MGMT	265	N/A	172.16.254.14/29	BD-Internal
Web	N/A	DVS	172.16.3.254/24	BD-Internal
App	N/A	DVS	172.16.4.254/24	BD-Internal
DB	N/A	DVS	172.16.5.254/24	BD-Internal

Deploy L4-L7 VLAN Stitching in Sample Tenants

This procedure details a setup method to demonstrate the ACI L4-L7 VLAN Stitching feature with L4-L7 services devices. In this lab validation, a pair of Cisco ASA-5585-X firewall devices in a High Availability configuration was connected to the ACI Fabric with a pair of vPCs. To add these ASAs to the fabric, because of hardware location in the lab, a pair of Nexus 9372PX leaves was also added to the fabric. The addition of the 9372s is not shown in this procedure. The firewalls were connected to ports Eth1/47 and Eth1/48 on the Nexus 9372s/ VLAN Stitching does not make use of device packages. Instead the firewalls were configured using the firewall CLI and ASDM interfaces. Inside and Outside VLAN interfaces connecting to the firewalls were configured in the ACI fabric. The detailed VLANs and IP subnet addresses are shown in the table below:

Table 30 **Tenant L4-L7 VLAN Stitching Details**

Tenant	Outside VLAN	Outside Firewall IP (ASA)	Outside Subnet Gateway (ACI)	Inside VLAN	Inside Firewall IP (ASA)	Web EPG Gateway (ACI)
App-A	501	172.16.253.1/29	172.16.253.6/29	502	172.16.0.1/24	172.16.0.254/24
App-B	503	172.16.253.9/29	172.16.253.14/29	504	172.16.3.1.24	172.16.3.254/24

Deploy Sample Cisco ASA VPCs

This section details setup of virtual port-channels for the Cisco ASA-5585-Xs used in this lab validation.

APIC Advanced GUI

1. From the Cisco APIC Advanced GUI, at the top, select Fabric > Inventory.
2. On the left, select Topology. On the upper right, select Configure.
3. Select ADD SWITCHES.
4. Using the Shift key, select the two leaf switches the ASAs are attached to. Click ADD SELECTED.
5. On the two leaf switches, select the ports that are connected the first ASA. On the lower right, click CONFIGURE VPC.
6. Name the Policy Group VPC-<ASA-1-name>.
7. Select the appropriate policies as shown in the screenshot.

BACK TO SUMMARY

CONFIGURING VPC

Port Channel
VPC
L2 Interface

L3
Conn. to Fex
Selected

a08-9372-1 (Node-103)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a08-9372-2 (Node-104)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

VPC

Policy Group Name: VPC-a05-asa5585-1

Description: optional

Link Level Policy: 10Gbps-Auto

CDP Policy: CDP-Disabled

MCP Policy: default

LLDP Policy: LLDP-Disabled

STP Interface Policy: BPDU-Filter-Guard

Egress Data Plane Policing Policy: default

Ingress Data Plane Policing Policy: default

Port Channel Policy: LACP-Active

Storm Control Interface Policy: default

L2 Interface Policy: VLAN-Scope-Global

Attached Entity Profile: select an option

BACK TO SUMMARY
APPLY CHANGES

8. Using the Attached Entity Profile drop-down, select Create Attachable Access Entity Profile.
9. Name the profile AEP-ASA. Click the + sign to add a Physical Domain.
10. Using the drop-down, select Create Physical Domain.
11. In the Create Physical Domain window, name the Domain PD-ASA.
12. Using the VLAN Pool drop-down, select Create VLAN Pool.
13. In the Create VLAN Pool window, name the VLAN Pool VP-ASA.
14. Select Static Allocation. Click the + sign to add a VLAN range to the pool.
15. In the Create Ranges window, input the From and To values for VLANs to be used for the firewall In-side and Outside VLANs. Select Static Allocation.

Create Ranges

Specify the Encap Block Range

Type: **VLAN**

Range: 501 - 504
From To

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation



Note: In this lab validation, 4 VLANs were added for 2 tenants (2 per tenant).

16. Click OK to complete creating the VLAN range.
17. Click SUBMIT to complete creating the VLAN Pool.
18. Click SUBMIT to complete creating the Physical Domain.
19. Click UPDATE.
20. Click SUBMIT to complete creating the Attachable Access Entity Profile.

BACK TO SUMMARY

CONFIGURING VPC

Port Channel
VPC
L2 Interface

L3
Conn. to Fex
Selected

a08-9372-1 (Node-103)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a08-9372-2 (Node-104)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

VPC

Policy Group Name: VPC-a05-asa5585-1

Description: optional

Link Level Policy: 10Gbps-Auto

CDP Policy: CDP-Disabled

MCP Policy: default

LLDP Policy: LLDP-Disabled

STP Interface Policy: BPDU-Filter-Guard

Egress Data Plane Policing Policy: default

Ingress Data Plane Policing Policy: default

Port Channel Policy: LACP-Active

Storm Control Interface Policy: default

L2 Interface Policy: VLAN-Scope-Global

Attached Entity Profile: AEP-ASA

BACK TO SUMMARY
APPLY CHANGES

21. Click APPLY CHANGES to complete creating the VPC.
22. Click OK for the confirmation.
23. On the two leaf switches, select the ports that are connected the second ASA. On the lower right, click CONFIGURE VPC.
24. Name the Policy Group VPC-<ASA-2-name>.
25. Select the appropriate policies and Attached Entity Profile as shown in the screenshot.

CONFIGURING VPC

Legend: Port Channel (light blue), L3 (light blue), VPC (green), Conn. to Fex (orange), L2 Interface (red), Selected (dark blue)

a08-9372-1 (Node-103)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

a08-9372-2 (Node-104)

01	03	05	07	09	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
02	04	06	08	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48

VPC

Policy Group Name: VPC-a05-asa5585-2

Description: optional

Link Level Policy: 10Gbps-Auto

CDP Policy: CDP-Disabled

MCP Policy: default

LLDP Policy: LLDP-Disabled

STP Interface Policy: BPDU-Filter-Guard

Egress Data Plane Policing Policy: default

Ingress Data Plane Policing Policy: default

Port Channel Policy: LACP-Active

Storm Control Interface Policy: default

L2 Interface Policy: VLAN-Scope-Global

Attached Entity Profile: AEP-ASA

BACK TO SUMMARY **APPLY CHANGES**

26. Click APPLY CHANGES to complete creating the VPC.

27. Click OK for the confirmation.

Create Tenant Firewall Outside Bridge Domain and Subnet

This section details setup of the bridge domain and associated subnet to be used for the ASA firewall context outside interface.

APIC Advanced GUI

1. From the Cisco APIC Advanced GUI, at the top, select Tenants > common.
2. On the left, expand Tenant common and Networking.
3. Right-click Bridge Domains and select Create Bridge Domain.
4. Name the Bridge Domain BD-App-A-Firewall-Outside.

- Using the VRF drop-down, select common/common-External.
- Leave Forwarding set to optimize and select the default End Point Retention Policy and IGMP Snoop Policy.

Create Bridge Domain


STEP 1 > Main


1. Main | 2. L3 Configurations | 3. Advanced/Troubleshooting


Specify Bridge Domain for the VRF


Name: BD-App-A-Firewall-Outside

Description:

VRF: common/common-Extern 

Forwarding: optimize 

End Point Retention Policy: default 
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: default 

- Click NEXT.
- Click the + sign to the right of Subnets to add a bridge domain subnet.
- Put in a gateway IP and mask for the subnet to be used for the outside interface of the tenant firewall context. It is recommended that this subnet is in the Supernet used for tenants elsewhere in this document.
- For Scope, select only Advertised Externally.

Create Subnet

Specify the Subnet Identity

Gateway IP: 172.16.253.6/29
address/mask

Treat as virtual IP address:

Make this IP address primary:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 Querier IP

L3 Out for Route Profile:

Route Profile:

ND RA Prefix policy:

OK CANCEL

11. Click OK to complete creating the subnet.
12. Click the + sign to the right of Associated L3 Outs to add the Shared L3 Out.
13. Using the drop-down, select common/Shared-L3-Out and click UPDATE.

i X

Create Bridge Domain

STEP 2 > L3 Configurations 1. Main 2. L3 Configurations 3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Unicast Routing: Enabled
 ARP Flooding: Enabled
 Config BD MAC Address:
 MAC Address:

Subnets: x +

Gateway Address	Scope	Primary IP Address	Subnet Control
172.16.253.6/29	Advertised Externally	False	ND RA Prefix

Enforce subnet check for IP learning:

DHCP Labels: x +

Name	Scope	DHCP Option Policy

Associated L3 Outs: x +

L3 Out
Shared-L3-Out

L3 Out for Route Profile: ▼
 Route Profile: ▼
 Link-local IPv6 Address:

PREVIOUS
NEXT
CANCEL

14. Click NEXT.

15. Select the default Monitoring Policy and click FINISH to complete creating the bridge domain.

Create Tenant L4-L7 Device and Service Graph

This section details setup of L4-L7 Device and Service Graph in the ACI Tenant.

APIC Advanced GUI

1. From the Cisco APIC Advanced GUI, at the top, select Tenants > App-A.

2. On the left, expand Tenant App-A, Application Profiles, Three-Tier-App, Application EPGs, and EPG Web.
3. Under EPG Web, select Contracts.
4. In the center pane, right-click the Allow-Shared-L3-Out contract and select Delete to remove this contract association.
5. Click YES for the confirmation.
6. On the left, expand Subnets and select the EPG subnet.
7. In the center pane, uncheck Advertised Externally. If the Web EPG is connected to Core-Services, leave Shared between VRFs selected. Otherwise, select Private to VRF.
8. Click SUBMIT to complete modifying the subnet.
9. On the left, expand Tenant App-A and L4-L7 Services.
10. Right-click L4-L7 Devices and select Create L4-L7 Devices.
11. In the Create L4-L7 Devices window, uncheck the Managed checkbox.
12. Name the Device ASA-App-A-Context. Select the Firewall Service Type.
13. For the Physical Domain, select PD-ASA. Select the HA Cluster Mode.
14. For Function Type, select GoTo.
15. Under Device 1, click the + sign to add the Device Interface.
16. Name the Device asa-1. Use the drop-down to select Path Type VPC. Select the VPC for the first ASA.
17. Click UPDATE.
18. Under Device 2, click the + sign to add the Device Interface.
19. Name the Device asa-2. Use the drop-down to select Path Type VPC. Select the VPC for the second ASA.
20. Click UPDATE.
21. Under Cluster, click the + sign to add a Concrete Interface.
22. Name the interface outside. Use the drop-down to select both the asa-1 and asa-2 devices.
23. For Encap, input vlan<App-A-outside-VLAN-ID>. Click UPDATE.
24. Under Cluster, click the + sign to add a second Concrete Interface.
25. Name the interface inside. Use the drop-down to select both the asa-1 and asa-2 devices.

26. For Encap, input vlan<App-A-inside-VLAN-ID>. Click UPDATE.

Create L4-L7 Devices
i
✕

STEP 1 > General
1. General

Please select device package and enter connectivity information.

General

Managed:

Name: ASA-App-A-Context

Service Type: Firewall

Device Type: PHYSICAL VIRTUAL

Physical Domain: PD-ASA

Mode: Single Node HA Cluster

Function Type: GoTo

Device 1

Device Interfaces: ✕ +

Name	Path
asa-1	Node-103-104/PC-a05-asa5585-1

Device 2

Device Interfaces: ✕ +

Name	Path
asa-2	Node-103-104/PC-a05-asa5585-2

Cluster

Cluster Interfaces: ✕ +

Name	Concrete Interfaces	Encap
outside	Device1/asa-1,Device2/asa-2	vlan-501
inside	Device1/asa-1,Device2/asa-2	vlan-502

PREVIOUS
FINISH
CANCEL

27. Click FINISH to complete creating the L4-L7 Device.

28. Right-click L4-L7 Service Graph Templates and select Create L4-L7 Service Graph Template.

29. In the Create L4-L7 Service Graph Template window, name the Graph ASA-App-A-Context.

30. Make sure Graph Type is set to Create A New One.

31. On the left drag the ASA-App-A Context (Firewall) icon to between the two EPGs.

32. Select the Routed Firewall.

Create L4-L7 Service Graph Template

Drag device clusters to create graph nodes.

Device Clusters

- App-A /ASA-App-A-Context (Firewall)

Graph Name: ASA-App-A-Context

Graph Type: Create A New One Clone An Existing One

Routed Transparent'."/>

33. Click SUBMIT to complete creating the Service Graph Template.

34. On the left, expand L4-L7 Service Graph Templates and select the ASA-App-A-Context Template.

35. Right-click the ASA-App-A-Context Template and select Apply L4-L7 Service Graph Template.

36. In the Apply L4-L7 Service Graph Template to EPGs window, use the Consumer EPG drop-down to select common/Shared-L3-Out/epg-Default-Route.

37. Use the Provider EPG drop-down to select App-A/Three-Tier-App/epg-Web.



Note: These EPG selections place the firewall between the Shared-L3-Out and App-A Web EPGs.

38. Under Contract Information, leave Create A New Contract selected and name the contract Allow-App-A-Firewall.



Note: It is important that this contract have a unique name within the ACI Fabric.




Apply L4-L7 Service Graph Template To EPGs

STEP 1 > Contract

1. Contract 2.

Config A Contract Between EPGs

EPGs Information

Consumer EPG / External Network: common/Shared-L3-Out/epg-Def  Provider EPG / External Network: App-A/Three-Tier-App/epg-Web  

Contract Information

Contract: Create A New Contract Choose An Existing Contract Subject

Contract Name: Allow-App-A-Firewall

No Filter (Allow All Traffic):

PREVIOUS NEXT

39. Click NEXT.

40. Under Consumer Connector, use the BD drop-down to select common/BD-App-A-Firewall-Outside.

41. Under Consumer Connector use the Cluster Interface drop-down to select outside.

42. Under Provider Connector use the Cluster Interface drop-down to select inside.

Apply L4-L7 Service Graph Template To EPGs

STEP 2 > Graph

1. Contract 2. G

Config A Service Graph

The screenshot shows the configuration of a Service Graph Template. On the left, a 'Device Clusters' pane lists 'App-A /ASA-App-A-Context (Firewall)'. The main area displays the 'Graph Template: App-A/ASA-App-A-Context'. The graph shows a central function node 'ASA-App-A-Context' (N1) connected to a 'Consumer' EPG (Default-Route) and a 'Provider' EPG (Web). Below the graph, the configuration for the ASA-App-A-Context function node is shown:

ASA-App-A-Context Information
Firewall: routed

Consumer Connector
 Type: General Route Peering
 BD: common/BD-App-A-Firewall-Outside
The Bridge Domain that connects the two devices
 Cluster Interface: outside

Provider Connector
 Type: General Route Peering
 BD: App-A/BD-Internal
The Bridge Domain that connects the two devices
 Cluster Interface: inside

At the bottom right, there are buttons for 'PREVIOUS', 'FINISH', and 'C'.

43. Click FINISH to complete applying the Service Graph Template.
44. On the left, expand Deployed Graph Instances and Allow-App-A-Firewall ASA-App-A-Context.
45. Select Function Node - N1.
46. Verify that the Function Connectors display values for Encap and Class ID.

Function Node - N1



Policy

Faults

History



Properties

Name: N1

Function Type: GoTo

Devices: ASA-App-A-Context

Cluster Interfaces:

Name	Concrete Interfaces	Encap
inside	ASA-App-A-Context_Device_1[asa-1], ASA-App-A-Context_Device_2[...]	vlan-502
outside	ASA-App-A-Context_Device_1[asa-1], ASA-App-A-Context_Device_2[...]	vlan-501

Function Connectors:

Name	Encap	Class ID
consumer	vlan-501	32778
provider	vlan-502	16391

47. Configure the ASA firewall device context as a routed firewall with NAT from the inside to the outside interface. The outside interface should be configured as a sub-interface encapsulated in the VLAN used for the consumer or outside interface and in the subnet entered in the BD-App-A-Firewall-Outside bridge domain. **The ASA context's default gateway should be the bridge domain's gateway address.** The inside interface should be configured as a sub-interface encapsulated in the VLAN used for the provider or inside interface and in the subnet entered in the App-A Web EPG. You can also add NAT rules to the firewall to reach VMs in the Web EPG with services such as HTTP or RDP. The actual ASA configuration is not covered in this document.
48. For VMs with interfaces **in the Web EPG**, **set the default gateway to the ASA context's inside interface IP.** You will also need to add persistent static routes to the EPG gateway to reach Core-Services and the App EPG.
49. Starting with Create Tenant Firewall Outside Bridge Domain and Subnet, add the ASA context and firewall for the second ACI tenant.

About the Authors

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

John George recently moved to Cisco from Netapp and is focused on designing, developing, validating, and supporting the FlexPod Converged Infrastructure. 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.

Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Masters of Science in Information Security from East Carolina University.

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, Cisco Systems, Inc.
- Chris O' Brien, Cisco Systems, Inc.
- Ramesh Isaac, Cisco Systems, Inc.
- Nabil Fares, NetApp
- Melissa Palmer, NetApp