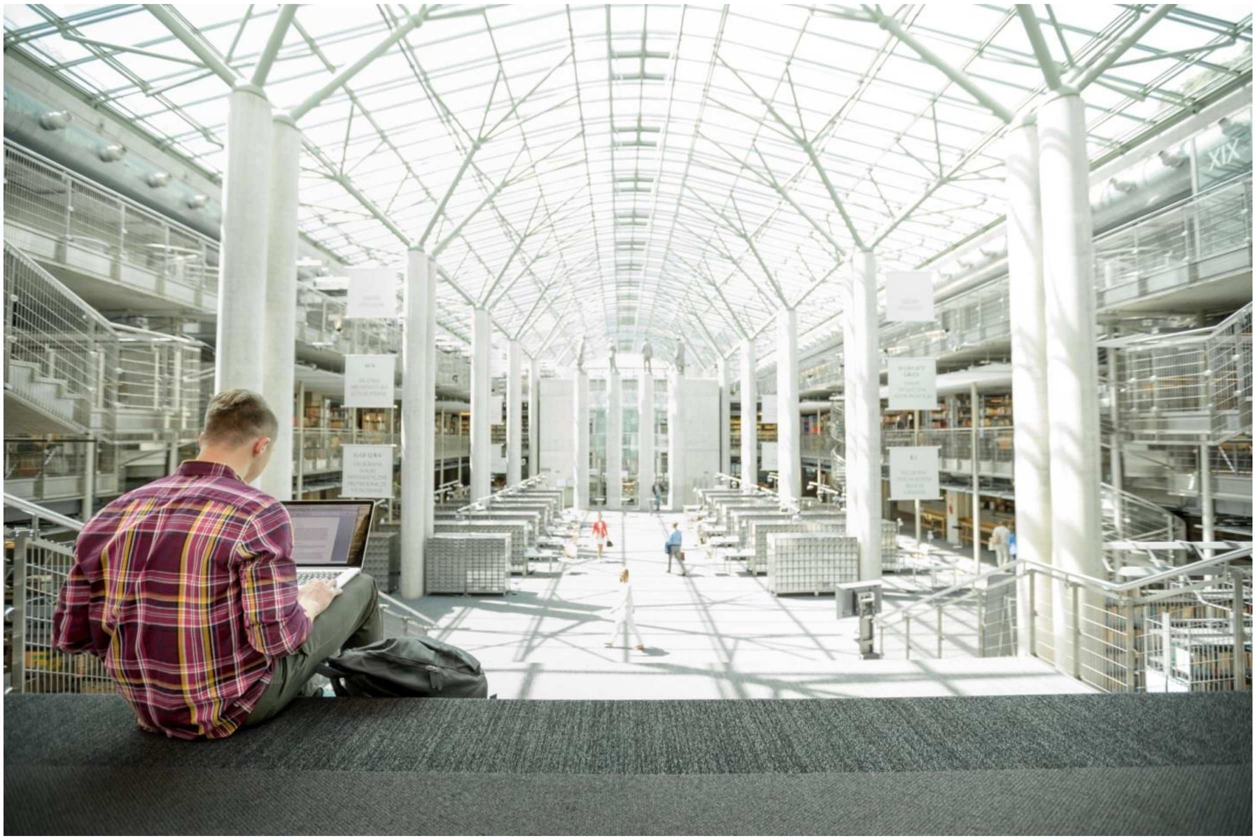

FlexPod Datacenter with VMware 6.5 Update1 and Cisco ACI 3.1

Deployment Guide for FlexPod Datacenter with VMware 6.5 Update1, Cisco ACI 3.1, and NetApp AFF A-Series

Last Updated: April 2, 2018

About the Cisco Validated Design Program

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

Table of Contents

Executive Summary

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

This document describes the Cisco and NetApp® FlexPod Datacenter with Cisco UCS Manager unified software release 3.2, Cisco Application Centric Infrastructure (ACI) 3.1(1i), NetApp ONTAP 9.3, and VMware vSphere 6.5 Update 1. Cisco UCS Manager (UCSM) 3.2 provides consolidated support for all the current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series, including Cisco UCS B200M5 servers. FlexPod Datacenter with Cisco UCS unified software release 3.2, and VMware vSphere 6.5 Update 1 is a predesigned, best-practice data center architecture built on Cisco Unified Computing System (UCS), Cisco Nexus® 9000 family of switches, Cisco Application Policy Infrastructure Controller (APIC), and NetApp All Flash FAS (AFF).

This document primarily focuses on deploying VMware vSphere 6.5 Update 1 on FlexPod Datacenter using iSCSI and NFS storage protocols. The Appendix section covers the delta changes on the configuration steps using the Fiber Channel (FC) storage protocol for the same deployment model.

FC storage traffic does not flow through the ACI Fabric and is not covered by the ACI policy model.

Solution Overview

Introduction

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

Audience

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

Purpose of this Document

This document provides a step-by-step configuration and implementation guidelines for the FlexPod Datacenter with Cisco UCS Fabric Interconnects, NetApp AFF, and Cisco ACI solution. This document primarily focuses on deploying VMware vSphere 6.5 Update 1 on FlexPod Datacenter using iSCSI and NFS storage protocols. The Appendix section covers the delta changes on the configuration steps using FC storage protocol for the same deployment model.

What's New?

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

- Support for the Cisco UCS 3.2 unified software release, Cisco UCS B200-M5 servers, Cisco UCS B200-M4 servers, and Cisco UCS C220-M4 servers
- Support for Cisco ACI version 3.1(1i)
- Support for the latest release of NetApp ONTAP® 9.3
- NFS, iSCSI, and FC storage design
- Validation of VMware vSphere 6.5 Update 1

Solution Design

Architecture

FlexPod architecture is highly modular, or pod-like. Although each customer's FlexPod unit might vary in its exact configuration, after a FlexPod unit is built, it can easily be scaled as requirements and demands change. This includes both scaling up (adding additional resources within a FlexPod unit) and scaling out (adding additional FlexPod units). Specifically, FlexPod is a defined set of hardware and software that serves as an integrated foundation for all virtualization solutions. FlexPod validated with VMware vSphere 6.5 Update 1 includes NetApp All Flash FAS storage, Cisco ACI® networking, Cisco Unified Computing System (Cisco UCS®), VMware vCenter and VMware ESXi in a single package. The design is flexible enough that the networking, computing, and storage can fit in a single data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

The reference architectures detailed in this document highlight the resiliency, cost benefit, and ease of deployment across multiple storage protocols. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

Figure 1 shows the VMware vSphere built on FlexPod components and its physical cabling with the Cisco UCS 6332-16UP Fabric Interconnects. The Nexus 9336PQ switches shown serve as spine switches in the Cisco ACI Fabric Spine-Leaf Architecture, while the Nexus 9332PQ switches serve as 40GE leaf switches. The Nexus 93180LC-EX 40GE leaf switch is also supported but was not validated. The Cisco APICs shown attach to the ACI Fabric with 10GE connections. This attachment can be accomplished with either other leaf switches in the fabric with 10 GE ports, such as the Nexus 93180YC-EX. The APIC cannot be directly attached to the Nexus 9332s. The APICs could be directly attached to Nexus 93180LC-EX switches with either 10GE breakout cables or QSFP to SFP+ Adapters (QSAs). This design has end-to-end 40 Gb Ethernet connections from Cisco UCS Blades or Cisco UCS C-Series rackmount servers, to a pair of Cisco UCS Fabric Interconnects, to Cisco Nexus 9000 switches, and through to NetApp AFF A300. These 40 GE paths carry NFS, iSCSI, and Virtual Machine (VM) traffic that has Cisco ACI policy applied. This infrastructure option can be expanded by connecting 16G FC or 10G FCoE links between the Cisco UCS Fabric Interconnects and the NetApp AFF A300 as shown below, or introducing a pair of Cisco MDS switches between the Cisco UCS Fabric Interconnects and the NetApp AFF A300 to provide FC block-level shared storage access. Note that FC/FCoE storage access does not have ACI policy applied. The FC configuration shown below is covered in the appendix of this document, but the FCoE and MDS options are also supported. The reference architecture reinforces the "wire-once" strategy, because the additional storage can be introduced into the existing architecture without a need for re-cabling from the hosts to the Cisco UCS Fabric Interconnects.

Physical Topology

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

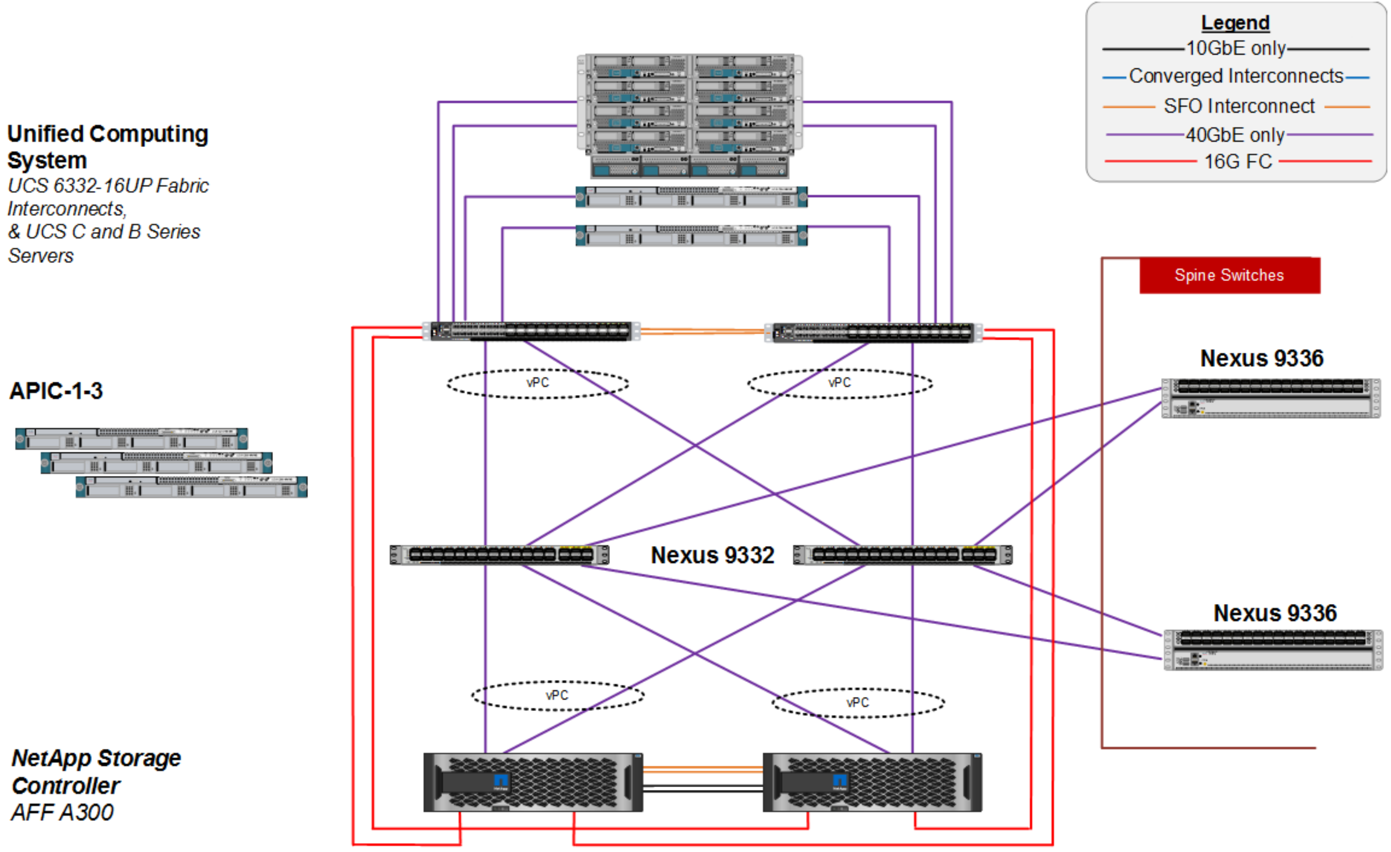

The reference 40Gb based hardware configuration includes:

- Three Cisco APICs
- Two Cisco Nexus 9336PQ fixed spine switches
- Two Cisco Nexus 9332PQ leaf switches
- Two Cisco UCS 6332-16UP fabric interconnects
- One chassis of Cisco UCS blade servers (Cisco UCS B200M4 and M5)
- Two Cisco UCS C220M4 rack servers
- One NetApp AFF A300 (HA pair) running ONTAP with disk shelves and solid state drives (SSD)

A 10GE-based design with Cisco UCS 6200 Fabric Interconnects is also supported, but not covered in this deployment Guide. All systems and fabric links feature redundancy and provide end-to-end high availability. For server virtualization, this deployment includes VMware vSphere 6.5 Update 1. Although this is the base design, each of the components can be scaled flexibly to support specific business requirements. For example, more (or different) blades and chassis could be deployed to increase compute capacity, additional disk shelves could be deployed to improve I/O capacity and throughput, or special hardware or software features could be added to introduce new features.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	- Cisco UCS Fabric Interconnects 6200 and 6300 Series. - Cisco UCS B-200 M5, B-200 M4, UCS C-220 M4	- 3.2(3a) * (Infrastructure & Server Bundle)	Initial Validation on 3.2(2d) Includes the Cisco UCS-IOM 2304 Cisco UCS Manager, Cisco UCS VIC 1340 and Cisco UCS VIC 1385
Network	Cisco APIC	3.1(1i)	
	Cisco Nexus 9000 ACI	n9000-13.1(1i)	
Storage	NetApp AFF A300	ONTAP 9.3	
Software	Cisco UCS Manager	3.2(3a) *	Initial Validation on 3.2(2d)
	VMware vSphere	6.5 Update 1 *	VMware vSphere Patches ESXi650-201803401-BG and ESXi650-201803402-BG applied after initial validation VMware vCenter upgraded to 6.5U1g after initial validation
	Cisco VIC nenic Driver	1.0.13.0	
	Cisco VIC fnic Driver	1.6.0.36	

* Initial validation was completed using earlier versions of Cisco UCS and VMware releases. However, Cisco and VMware released Speculative Execution vulnerability (Spectre & Meltdown) patches in updated software releases (shown below) after validation was complete. These patches and releases were installed and limited runs of validation tests were performed to check for continued behavior. Cisco recommends and supports the updated releases and patches for this CVD.

- Cisco UCS Manager 3.2(3a)
- VMware vCenter 6.5U1g
- ESXi 6.5U1 Cisco Custom ISO with patches ESXi650-201803401-BG and ESXi650-201803402-BG applied

Configuration Guidelines

This document provides details on configuring a fully redundant, highly available reference model for a FlexPod unit with NetApp ONTAP storage. Therefore, reference is made to the component being configured with each step, as either 01 or 02 or A and B. In this CVD we have used node01 and node02 to identify the two NetApp storage controllers provisioned in this deployment model. Similarly, Cisco Nexus A and Cisco Nexus B refer to the pair of Cisco Nexus switches configured. Likewise the Cisco UCS Fabric Interconnects are also configured in the same way. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts with -01, -02 for the hostnames to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the `network port vlan create` command:

Usage:

```
network port vlan create ?

 [-node] <nodename>                   Node

 { [-vlan-name] {<netport>|<ifgrp>}   VLAN Name

 |  -port {<netport>|<ifgrp>}         Associated Network Port

 [-vlan-id] <integer> }               Network Switch VLAN Identifier
```

Example:

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

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

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out-of-Band-Mgmt	VLAN for out-of-band management interfaces	3911
IB-MGMT	VLAN for in-band management interfaces	118/219/318/419/VS
Native-VLAN	VLAN to which untagged frames are assigned	2
Foundation-NFS-VLAN	VLAN for NFS traffic	3050/3150
vMotion-VLAN	VLAN designated for the movement of VMs from one physical host to another.	3000
Foundation-iSCSI-A	VLAN for iSCSI Boot on Fabric A	3010/3110
Foundation-iSCSI-B	VLAN for iSCSI Boot on Fabric B	3020/3120

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

Table 3 Virtual Machines

Virtual Machine Description	Host Name
Active Directory (AD)	ACI-FP-AD1, ACI-FP-AD2
VMware vCenter	fpv-vc

Virtual Machine Description	Host Name
NetApp Virtual Storage Console (VSC)	fpv-vsc

Physical Infrastructure

FlexPod Cabling

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

The tables in this section contain details for the prescribed and supported configuration of the NetApp AFF A300 running NetApp ONTAP® 9.3.

For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT) and Cisco FlexPod documents on cisco.com. Please log in to access these documents. Use cisco.com log in credentials to access FlexPod documents and NetApp support account to access the NetApp tool..

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps. Make sure to use the cabling directions in this section as a guide.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 2 details the cable connections used in the validation lab for the 40Gb end-to-end with additional 16 Gb Fibre Channel topology based on the Cisco UCS 6332-16UP Fabric Interconnect. Four 16Gb links connect directly to the NetApp AFF controllers from the Cisco UCS Fabric Interconnects. An additional 1Gb management connection is required for an out-of-band network switch apart from the FlexPod infrastructure. Cisco UCS fabric interconnects and Cisco Nexus switches are connected to the out-of-band network switch, and each NetApp AFF controller has two connections to the out-of-band network switch.

Figure 2 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnect

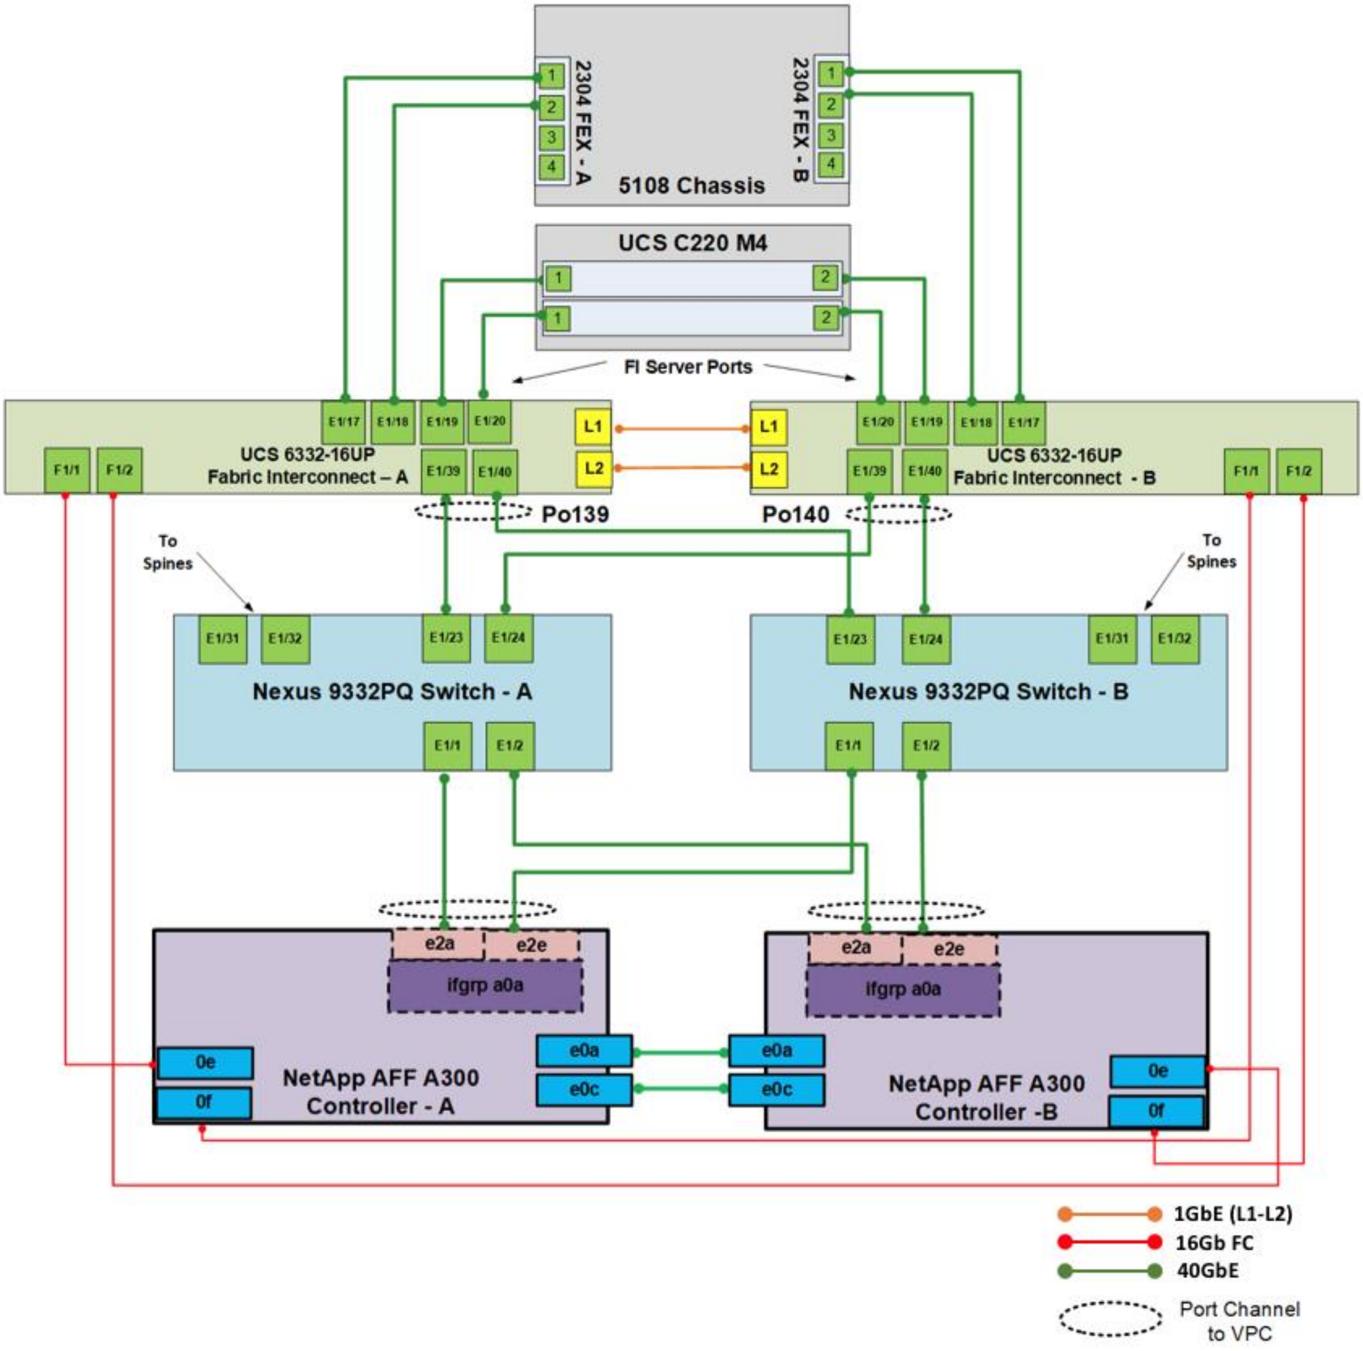

Infrastructure Servers Prerequisites

Active Directory DC/DNS

Production environments at most customers' locations might have an active directory and DNS infrastructure configured. In this document we have assumed an existing AD domain controller and an AD integrated DNS server role running on the same server, which is available in our lab environment. We will configure two additional AD/DNS servers connected to the Core Services End Point Group (EPG) in the ACI Fabric. These AD/DNS servers will be configured as additional Domain Controllers in the same domain as the prerequisite AD/DNS server.

Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco ACI fabric for use in a FlexPod environment and is written where the FlexPod components are added to an existing Cisco ACI fabric in several new ACI tenants. Required fabric setup is verified, but previous configuration of the ACI fabric is assumed.

Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

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

In ACI, both spine and leaf switches are configured using APIC, individual configuration of the switches is not required. Cisco 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) Verification

This sub-section verifies the setup of the Cisco APIC. Cisco recommends a cluster of at least 3 APICs controlling an ACI Fabric.

1. Log into the APIC GUI using a web browser, by browsing to the out of band IP address configured for APIC. Login with the admin user id and password.

In this validation, Google Chrome was used as the web browser. It might take a few minutes before APIC GUI is available after the initial setup.

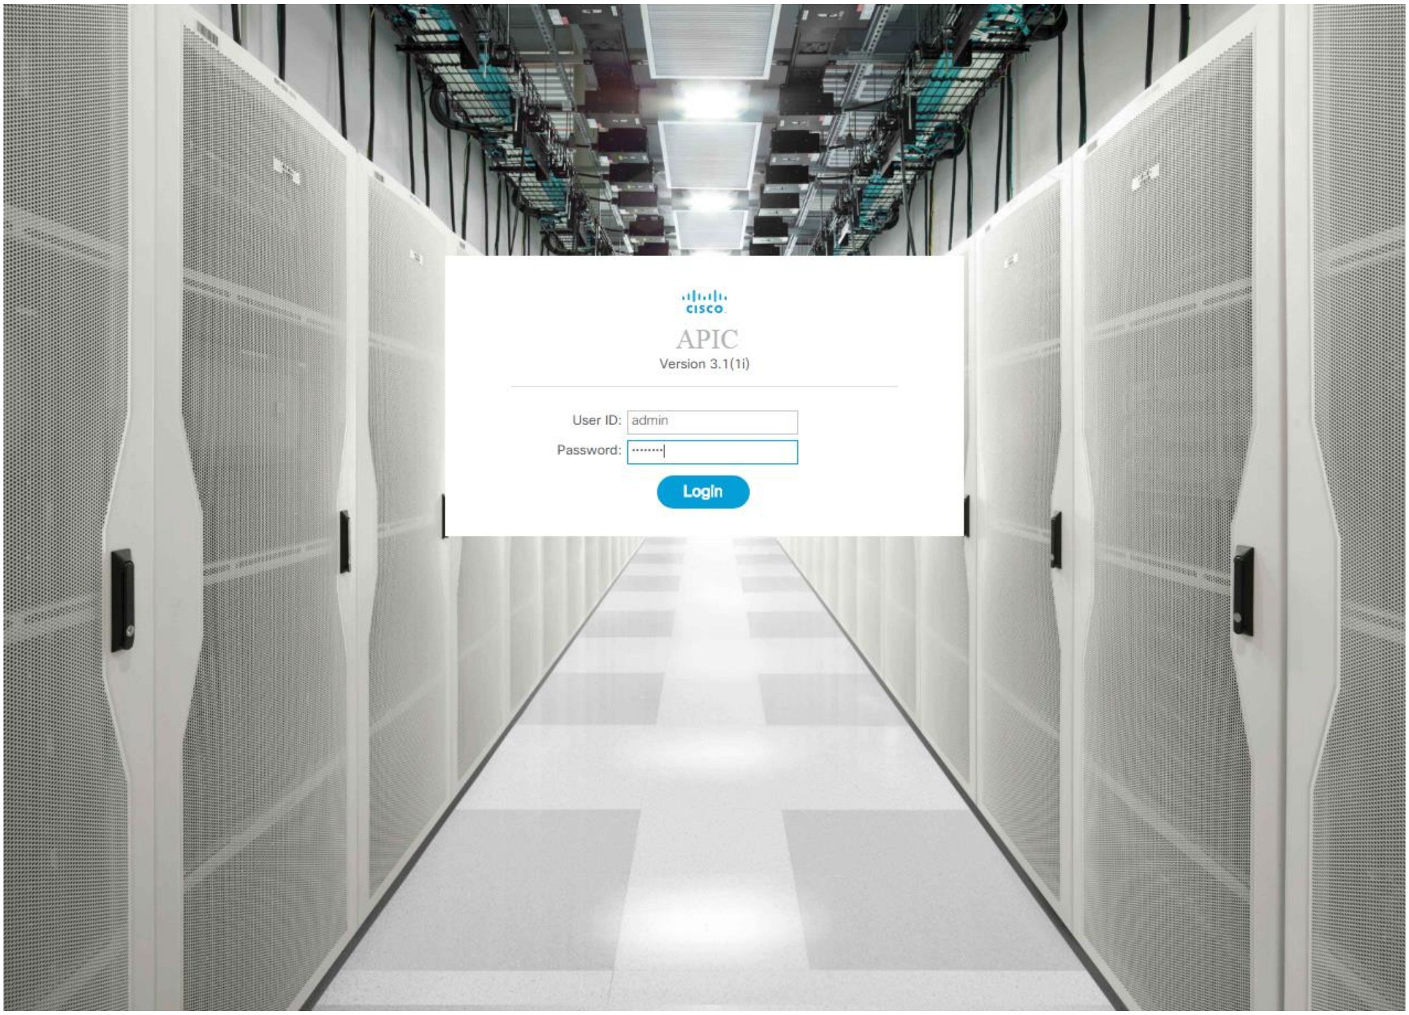

APIC

Version 3.1(1i)

User ID: admin

Password: ••••••••

2. Take the appropriate action to close any warning or information screens.

3. At the top in the APIC home page, select the **System** tab followed by **Controllers**.

4. On the left, select the **Controllers** folder. Verify that at least 3 APICs are available and have redundant connections to the fabric.

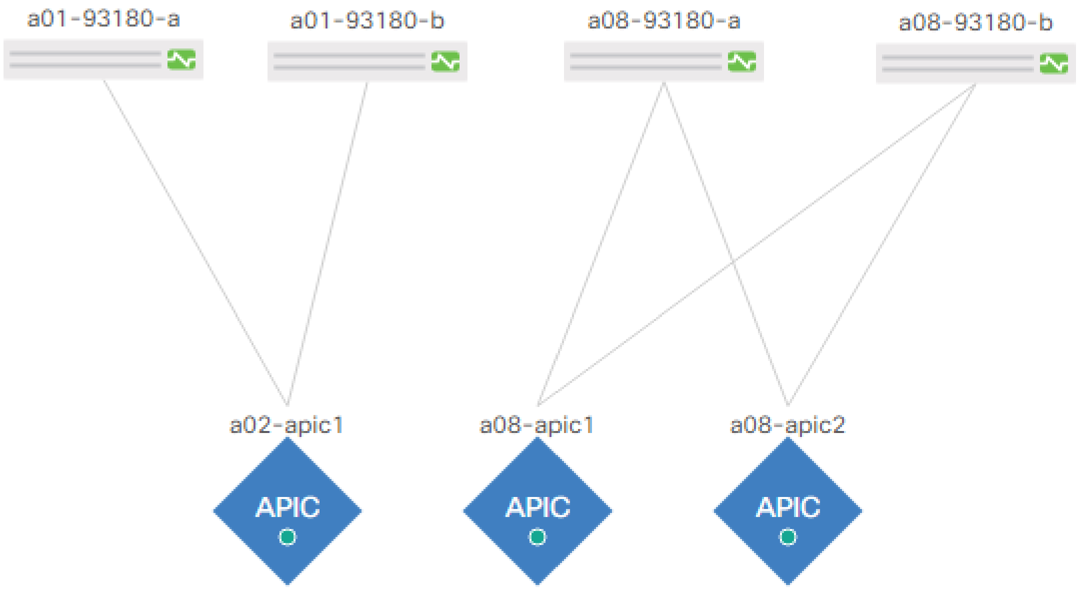

Cisco ACI Fabric Discovery

This section details the steps for adding the two Nexus 9332PQ leaf switches to the Fabric. This procedure is assuming that a FlexPod with dedicated leaves is being added to the fabric. If the two Nexus 9332s have already been added to the fabric, continue on to the next section. These switches are automatically discovered in the ACI Fabric and are manually assigned node IDs. To add Nexus 9332PQ leaf switches to the Fabric, complete the following steps:

1. At the top in the APIC home page, select the Fabric tab and make sure Inventory under Fabric is selected.

2. In the left pane, select and expand Fabric Membership.

3. The two 9332 Leaf Switches will be listed on the Fabric Membership page with Node ID 0 as shown:

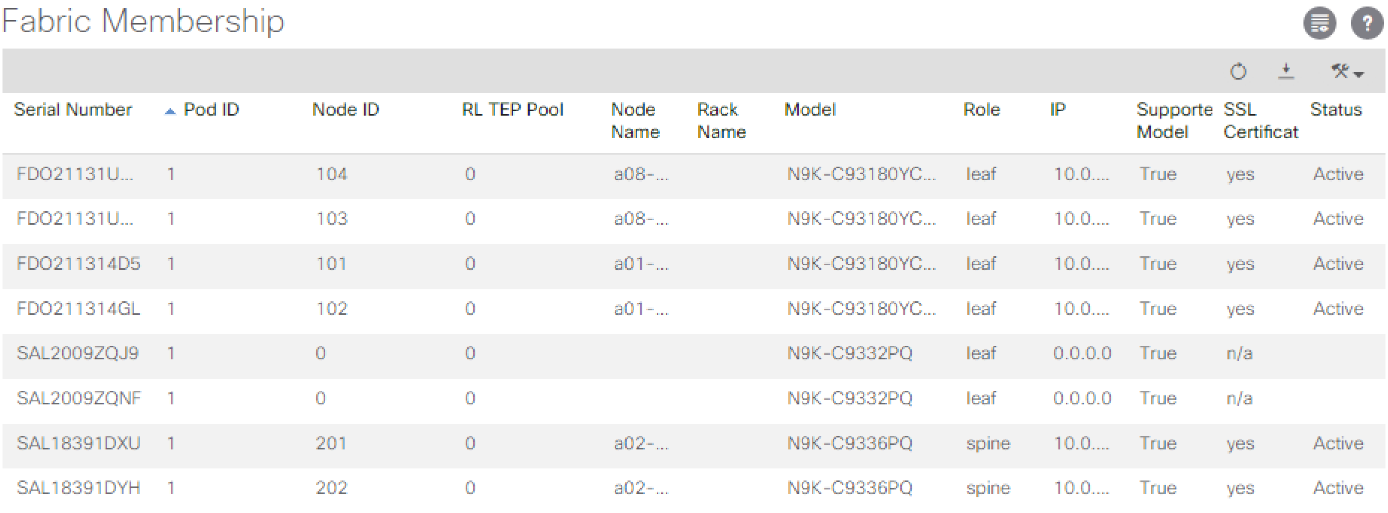

Fabric Membership

Serial Number	Pod ID	Node ID	RL TEP Pool	Node Name	Rack Name	Model	Role	IP	Supporte Model	SSL Certificat	Status
FDO21131U...	1	104	0	a08-...		N9K-C93180YC...	leaf	10.0....	True	yes	Active
FDO21131U...	1	103	0	a08-...		N9K-C93180YC...	leaf	10.0....	True	yes	Active
FDO211314D5	1	101	0	a01-...		N9K-C93180YC...	leaf	10.0....	True	yes	Active
FDO211314GL	1	102	0	a01-...		N9K-C93180YC...	leaf	10.0....	True	yes	Active
SAL2009ZQJ9	1	0	0			N9K-C9332PQ	leaf	0.0.0.0	True	n/a	
SAL2009ZQNF	1	0	0			N9K-C9332PQ	leaf	0.0.0.0	True	n/a	
SAL18391DXU	1	201	0	a02-...		N9K-C9336PQ	spine	10.0....	True	yes	Active
SAL18391DYH	1	202	0	a02-...		N9K-C9336PQ	spine	10.0....	True	yes	Active

4. Connect to the two Nexus 9332 leaf switches using serial consoles and login in as admin with no password (press enter). Use **show inventory** to get the leaf's serial number.

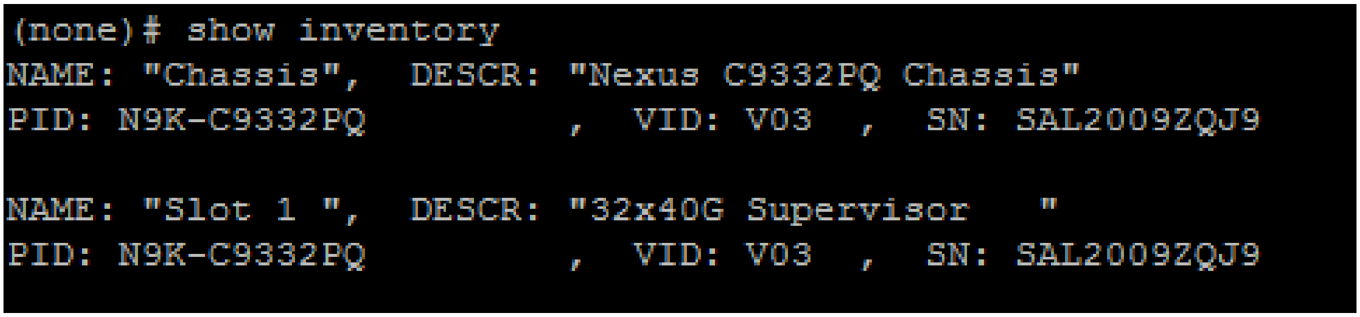

```
(none)# show inventory
NAME: "Chassis",  DESCR: "Nexus C9332PQ Chassis"
PID: N9K-C9332PQ           ,  VID: V03  ,  SN: SAL2009ZQJ9

NAME: "Slot 1 ",  DESCR: "32x40G Supervisor   "
PID: N9K-C9332PQ           ,  VID: V03  ,  SN: SAL2009ZQJ9
```

5. Match the serial numbers from the leaf listing to determine the A and B switches under Fabric Membership.

6. In the APIC GUI, under Fabric Membership, double-click the A leaf in the list. Enter a Node ID and a Node Name for the Leaf switch and click **Update**.

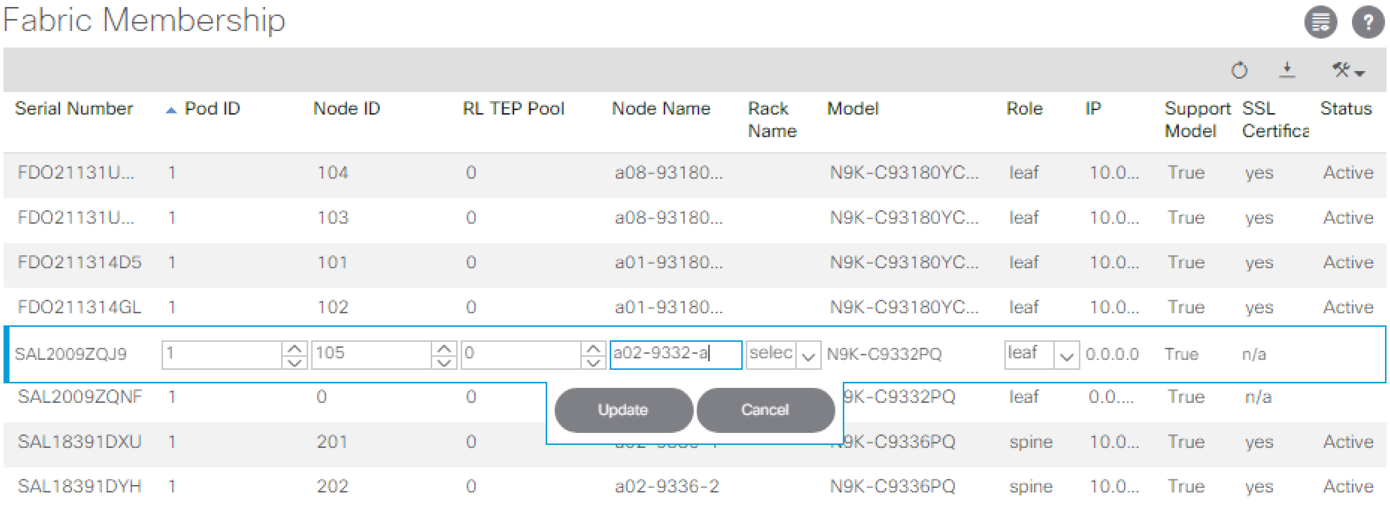

Fabric Membership

Serial Number	Pod ID	Node ID	RL TEP Pool	Node Name	Rack Name	Model	Role	IP	Support Model	SSL Certifica	Status
FDO21131U...	1	104	0	a08-93180...		N9K-C93180YC...	leaf	10.0...	True	yes	Active
FDO21131U...	1	103	0	a08-93180...		N9K-C93180YC...	leaf	10.0...	True	yes	Active
FDO211314D5	1	101	0	a01-93180...		N9K-C93180YC...	leaf	10.0...	True	yes	Active
FDO211314GL	1	102	0	a01-93180...		N9K-C93180YC...	leaf	10.0...	True	yes	Active
SAL2009ZQJ9	1	105	0	a02-9332-a	selec	N9K-C9332PQ	leaf	0.0.0.0	True	n/a	
SAL2009ZQNF	1	0	0			9K-C9332PQ	leaf	0.0....	True	n/a	
SAL18391DXU	1	201	0	[illegible]		9K-C9336PQ	spine	10.0...	True	yes	Active
SAL18391DYH	1	202	0	a02-9336-2		N9K-C9336PQ	spine	10.0...	True	yes	Active

7. Repeat step 6 for the B leaf in the list.

8. Click **Topology** in the left pane. The discovered ACI Fabric topology will appear. It may take a few minutes for the Nexus 9332 Leaf switches to appear and you will need to click the refresh button for the complete topology to appear. You may also need to move the switches around to get the arrangement that you desire.

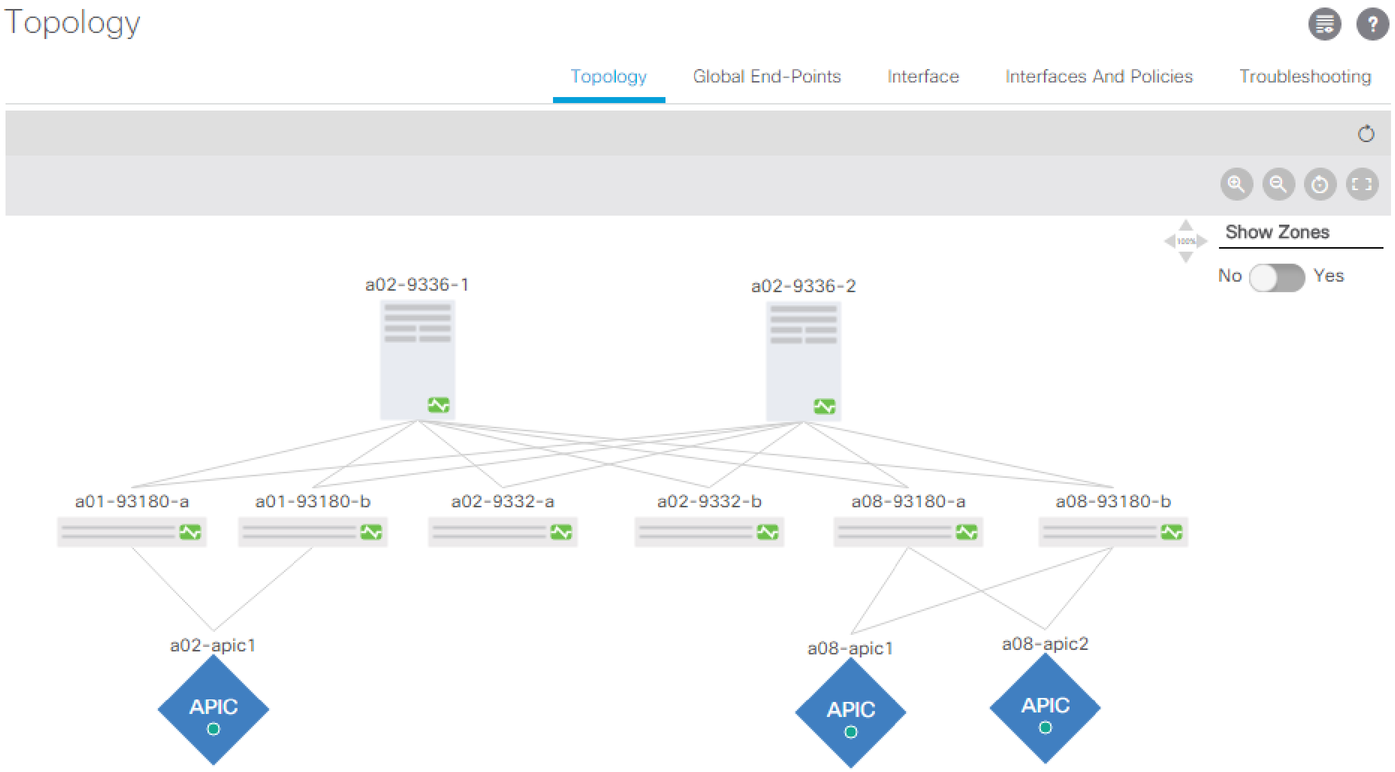

The topology shown in the screenshot above is the topology of the validation lab fabric containing 6 leaf switches, 2 spine switches, and 3 APICs. Notice that the APICs are connected to 10GE ports. Customer topology will vary depending on number and type of devices. Cisco recommends a cluster of at least 3 APICs in a production environment.

Initial ACI Fabric Setup Verification

This section details the steps for the initial setup of the Cisco ACI Fabric, where the software release is validated, out of band management IPs are assigned to the new leaves, NTP setup is verified, and the fabric BGP route reflectors are verified.

Software Upgrade

To upgrade the software, complete the following steps:

1. In the APIC GUI, at the top select **Admin** > **Firmware**.

2. This document was validated with ACI software release 3.1(1i). Select Fabric Node Firmware in the left pane under Firmware Management. All switches should show the same firmware release and the release version should be at minimum n9000-13.1(1i). The switch software version should also correlate with the APIC version.

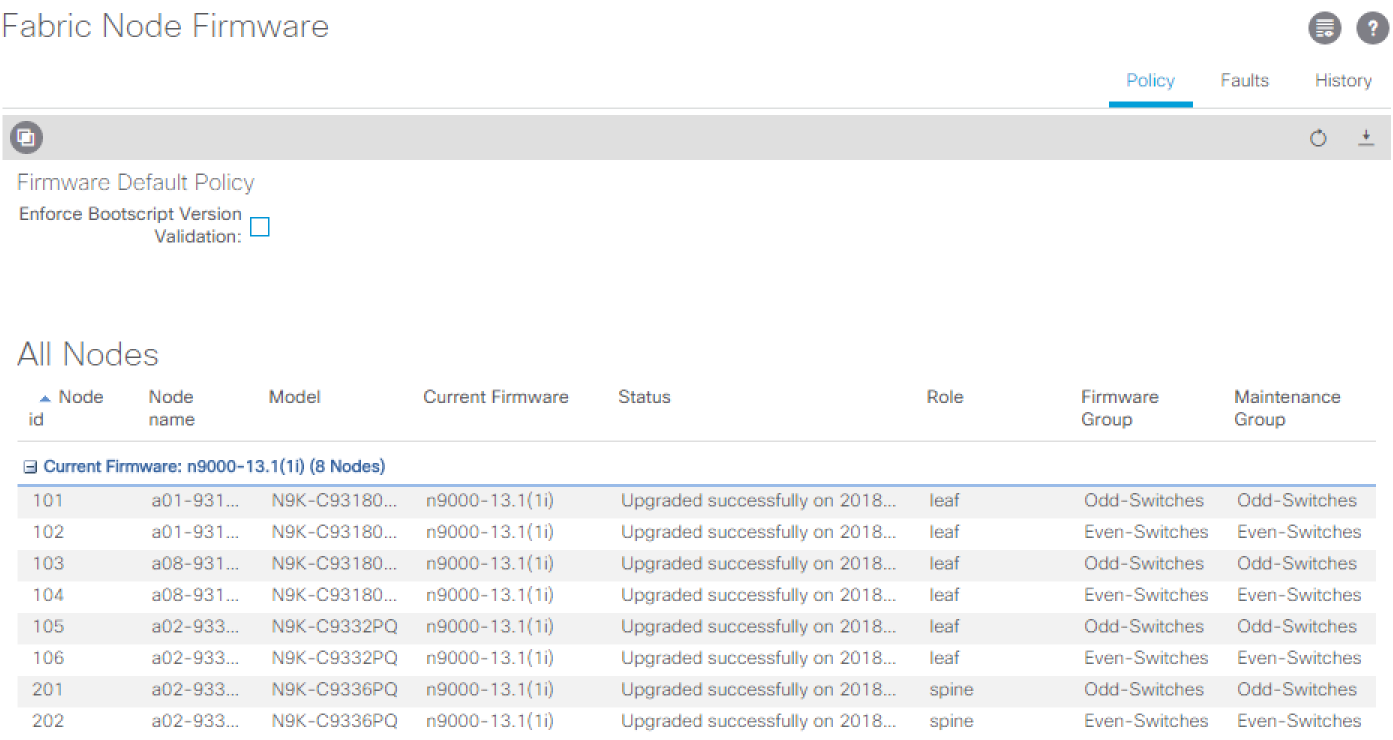

Fabric Node Firmware

Firmware Default Policy

Enforce Bootscript Version Validation: □

All Nodes

▲ Node id	Node name	Model	Current Firmware	Status	Role	Firmware Group	Maintenance Group
⊟ Current Firmware: n9000-13.1(1i) (8 Nodes)							
101	a01-931...	N9K-C93180...	n9000-13.1(1i)	Upgraded successfully on 2018...	leaf	Odd-Switches	Odd-Switches
102	a01-931...	N9K-C93180...	n9000-13.1(1i)	Upgraded successfully on 2018...	leaf	Even-Switches	Even-Switches
103	a08-931...	N9K-C93180...	n9000-13.1(1i)	Upgraded successfully on 2018...	leaf	Odd-Switches	Odd-Switches
104	a08-931...	N9K-C93180...	n9000-13.1(1i)	Upgraded successfully on 2018...	leaf	Even-Switches	Even-Switches
105	a02-933...	N9K-C9332PQ	n9000-13.1(1i)	Upgraded successfully on 2018...	leaf	Odd-Switches	Odd-Switches
106	a02-933...	N9K-C9332PQ	n9000-13.1(1i)	Upgraded successfully on 2018...	leaf	Even-Switches	Even-Switches
201	a02-933...	N9K-C9336PQ	n9000-13.1(1i)	Upgraded successfully on 2018...	spine	Odd-Switches	Odd-Switches
202	a02-933...	N9K-C9336PQ	n9000-13.1(1i)	Upgraded successfully on 2018...	spine	Even-Switches	Even-Switches

3. Click **Admin > Firmware > Controller Firmware**. If all APICs are not at the same release at a minimum of 3.1(1i), follow the Cisco APIC Management, Installation, Upgrade, and Downgrade Guide to upgrade both the APICs and switches if the APICs are not at a minimum release of 3.1(1i) and the switches are not at n9000-13.1(1i).

Setting Up Out-of-Band Management IP Addresses for New Leaf Switches

To set up out-of-band management IP addresses, complete the following steps:

1. To add out-of-band management interfaces for all the switches in the ACI Fabric, select **Tenants > mgmt**.
2. Expand Tenant mgmt on the left. Right-click Node Management Addresses and select Create Static Node Management Addresses.
3. Enter the node number range for the new leaf switches (105-106 in this example).
4. Select the checkbox for Out-of-Band Addresses.
5. Select default for Out-of-Band Management EPG.
6. Considering that the IPs will be applied in a consecutive range of two IPs, enter a starting IP address and net-mask in the Out-Of-Band IPV4 Address field.
7. Enter the out of band management gateway address in the Gateway field.
8. Click SUBMIT, then click YES.
9. On the left, expand Node Management Addresses and select Static Node Management Addresses. Verify the mapping of IPs to switching nodes.

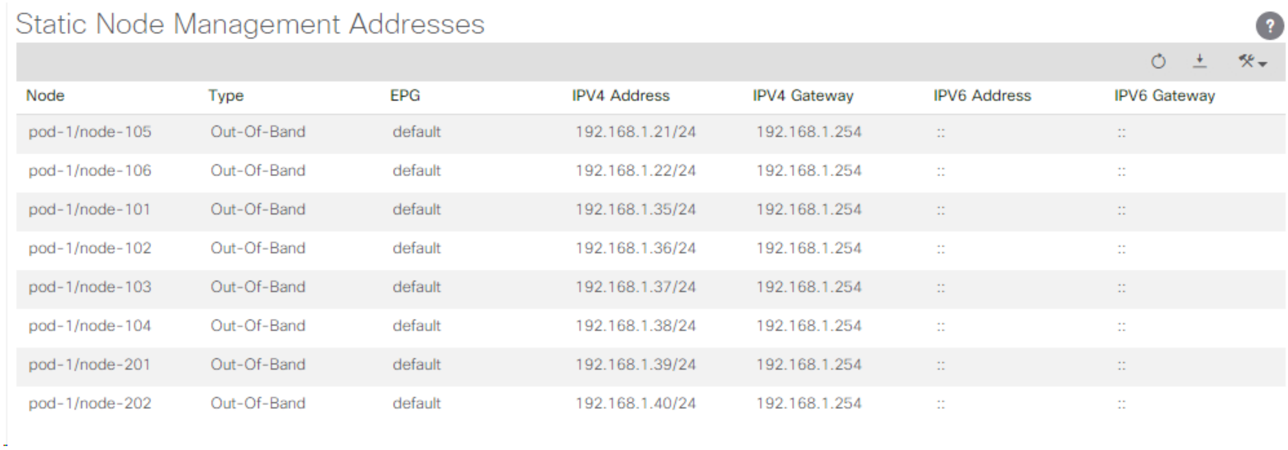

Static Node Management Addresses

Node	Type	EPG	IPV4 Address	IPV4 Gateway	IPV6 Address	IPV6 Gateway
pod-1/node-105	Out-Of-Band	default	192.168.1.21/24	192.168.1.254	::	::
pod-1/node-106	Out-Of-Band	default	192.168.1.22/24	192.168.1.254	::	::
pod-1/node-101	Out-Of-Band	default	192.168.1.35/24	192.168.1.254	::	::
pod-1/node-102	Out-Of-Band	default	192.168.1.36/24	192.168.1.254	::	::
pod-1/node-103	Out-Of-Band	default	192.168.1.37/24	192.168.1.254	::	::
pod-1/node-104	Out-Of-Band	default	192.168.1.38/24	192.168.1.254	::	::
pod-1/node-201	Out-Of-Band	default	192.168.1.39/24	192.168.1.254	::	::
pod-1/node-202	Out-Of-Band	default	192.168.1.40/24	192.168.1.254	::	::

10. Direct out-of-band access to the switches is now available using SSH.

Verifying Time Zone and NTP Server

This procedure will allow customers to verify setup of an NTP server for synchronizing the fabric time. To verify the time zone and NTP server set up, complete the following steps:

1. To verify NTP setup in the fabric, select and expand Fabric > Fabric Policies > Pod Policies > Policies > Date and Time.

2. Select default. In the Datetime Format - default pane, verify the correct Time Zone is selected and that Offset State is enabled. Adjust as necessary and click Submit and Submit Changes.

3. On the left, select Policy default. Verify that at least one NTP Server is listed.

4. If desired, select **enabled** for Server State to enable the ACI fabric switches as NTP servers. Click **Submit**.

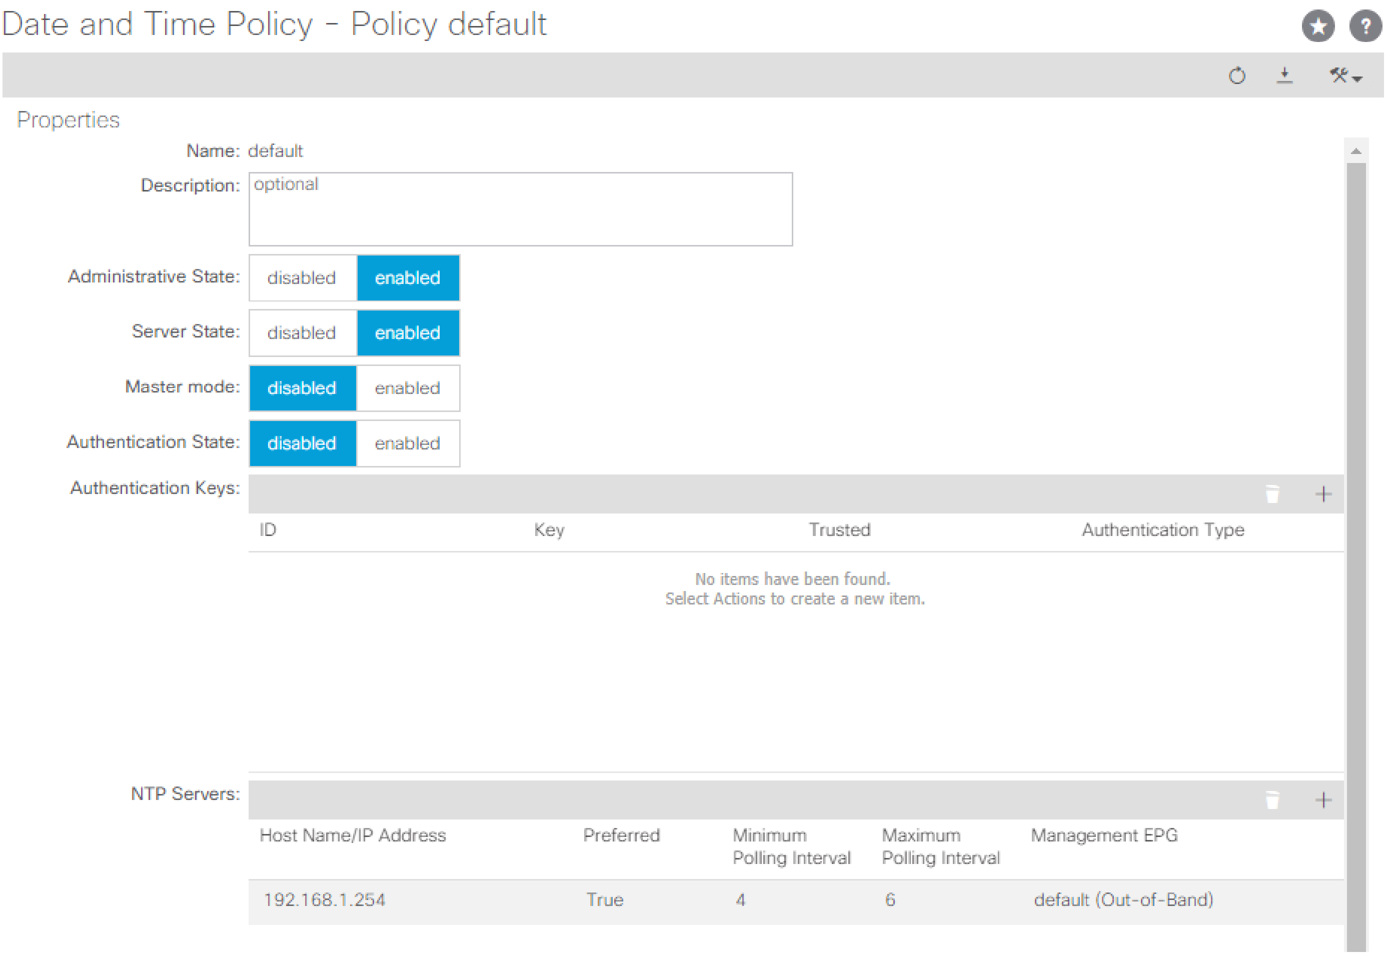

Date and Time Policy - Policy default

Properties

Name: default

Description: optional

Administrative State: disabled [✓ enabled]

Server State: disabled [✓ enabled]

Master mode: [✓ disabled] enabled

Authentication State: [✓ disabled] enabled

Authentication Keys:

ID	Key	Trusted	Authentication Type

No items have been found.
Select Actions to create a new item.

NTP Servers:

Host Name/IP Address	Preferred	Minimum Polling Interval	Maximum Polling Interval	Management EPG
192.168.1.254	True	4	6	default (Out-of-Band)

5. If necessary, 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.

Verifying Domain Name Servers

To verify optional DNS in the ACI fabric, complete the following steps:

1. Select and expand Fabric > Fabric Policies > Global Policies > DNS Profiles > default.

2. Verify the DNS Providers and DNS Domains.

3. If necessary, in the Management EPG drop-down, select the default (Out-of-Band) Management EPG. 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.

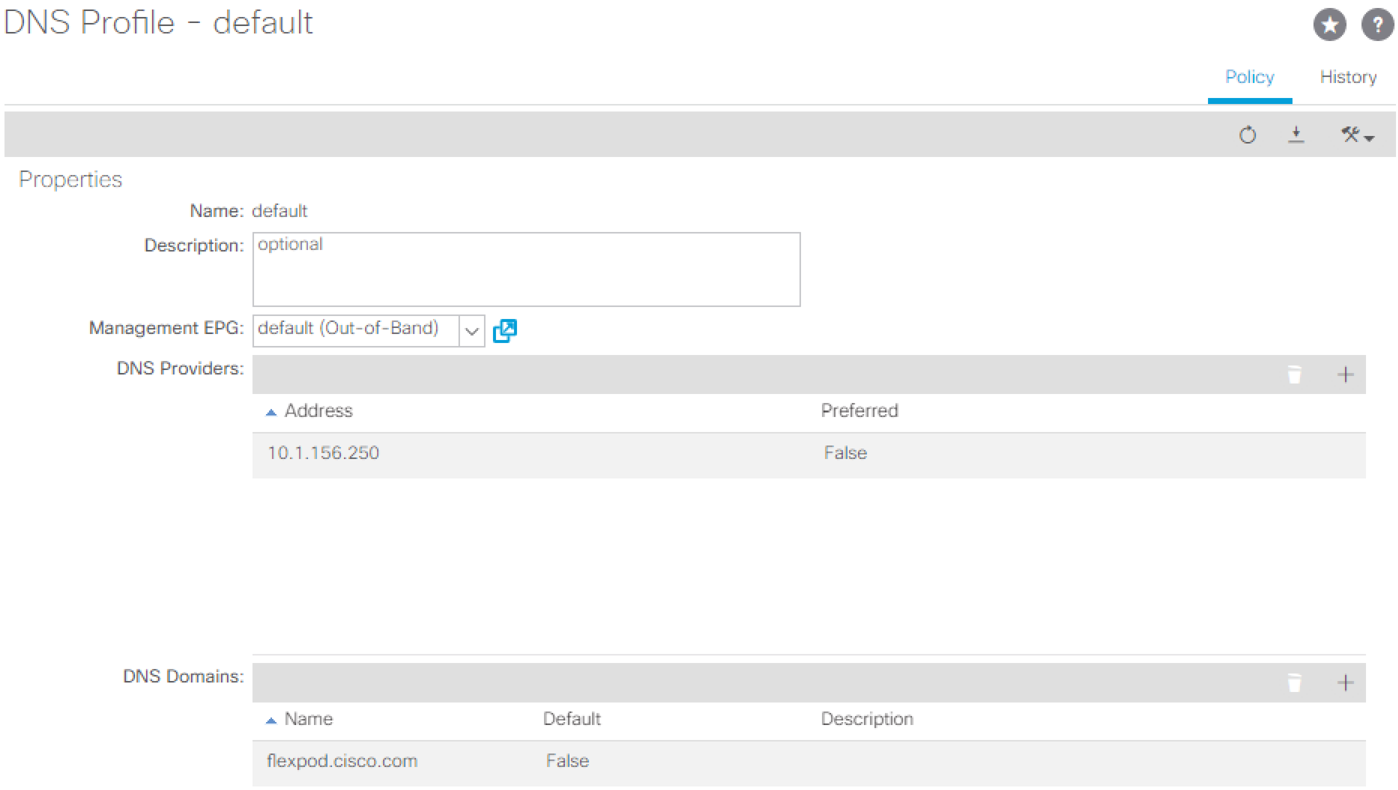

DNS Profile - default

Properties

Name: default

Description: optional

Management EPG: default (Out-of-Band)

DNS Providers:

Address	Preferred
10.1.156.250	False

DNS Domains:

Name	Default	Description
flexpod.cisco.com	False	

Verifying BGP Route Reflectors

In this ACI deployment, both the spine switches should be set up as BGP route-reflectors to distribute the leaf routes throughout the fabric. To verify the BGP Route Reflector, complete the following steps:

1. Select and expand System > System Settings > BGP Route Reflector.

2. Verify that a unique Autonomous System Number has been selected for this ACI fabric. If necessary, 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.

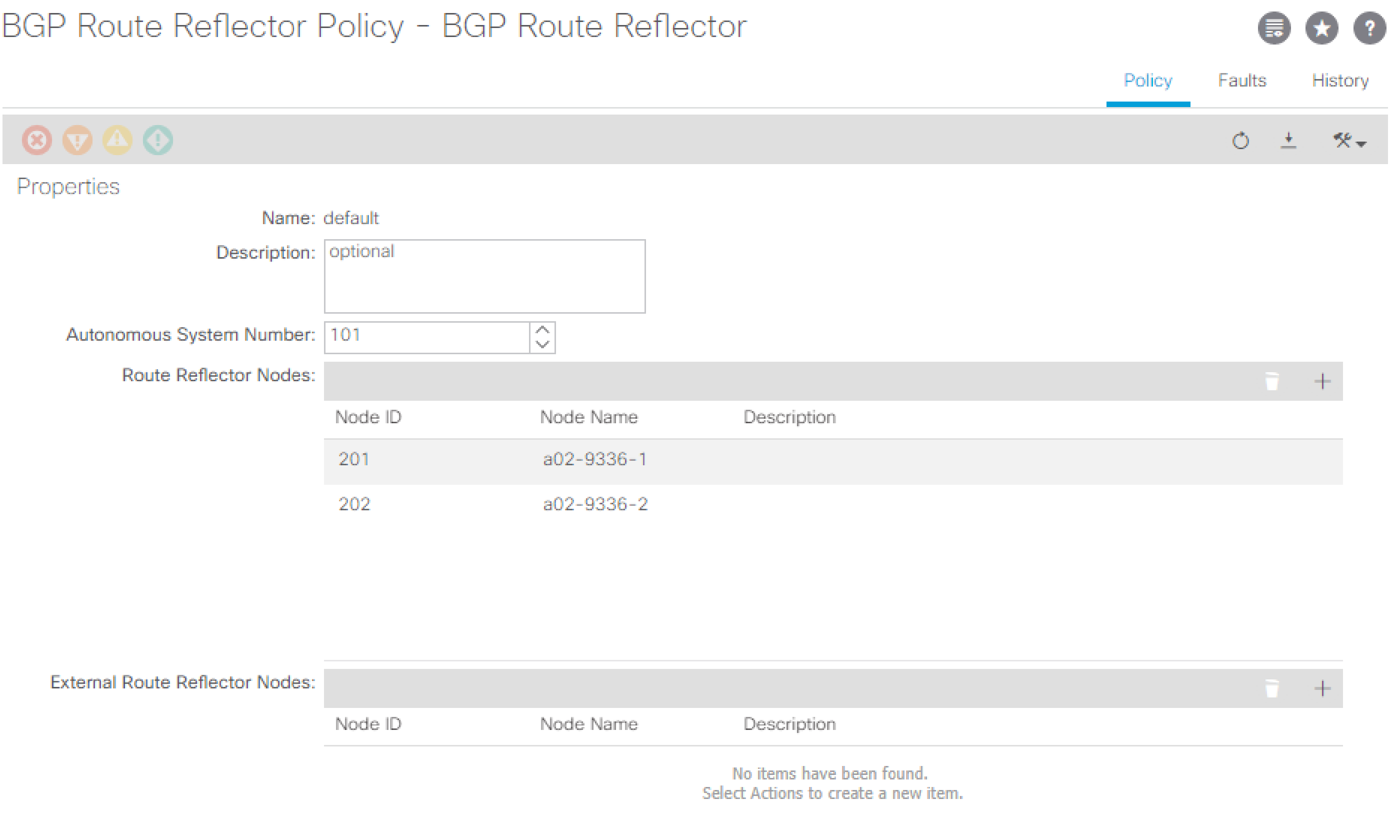

BGP Route Reflector Policy - BGP Route Reflector

Properties

Name: default

Description: optional

Autonomous System Number: 101

Route Reflector Nodes:

Node ID	Node Name	Description
201	a02-9336-1	
202	a02-9336-2	

External Route Reflector Nodes:

Node ID	Node Name	Description

No items have been found.
Select Actions to create a new item.

3. To verify the BGP Route Reflector has been enabled, select and expand Fabric > Fabric Policies > Pod Policies > Policy Groups. Under Policy Groups make sure a policy group has been created and select it. The BGP Route Reflector Policy field should show “default.”

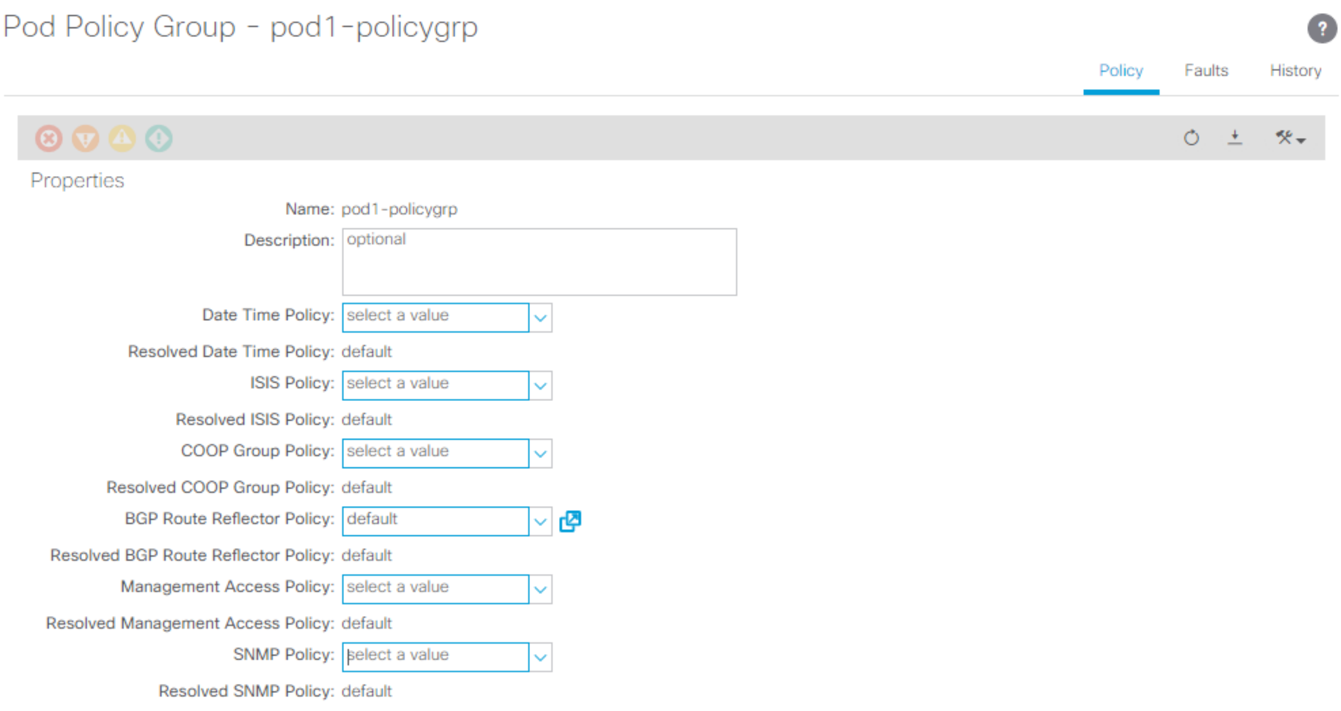

Pod Policy Group - pod1-policygrp

Properties

Name: pod1-policygrp

Description: optional

Date Time Policy: select a value

Resolved Date Time Policy: default

ISIS Policy: select a value

Resolved ISIS Policy: default

COOP Group Policy: select a value

Resolved COOP Group Policy: default

BGP Route Reflector Policy: default

Resolved BGP Route Reflector Policy: default

Management Access Policy: select a value

Resolved Management Access Policy: default

SNMP Policy: select a value

Resolved SNMP Policy: default

4. If a Policy Group has not been created, on the left, right-click Policy Groups under Pod Policies and select Create Pod Policy Group. In the Create Pod Policy Group window, name the Policy Group pod1-policygrp. Select the default BGP Route Reflector Policy. Click SUBMIT to complete creating the Policy Group.

5. On the left expand Profiles under Pod Policies and select Pod Profile default > default.

6. Verify that the pod1-policygrp or the Fabric Policy Group identified above is selected. If the Fabric Policy Group is not selected, view the drop-down list to select it and click Submit.

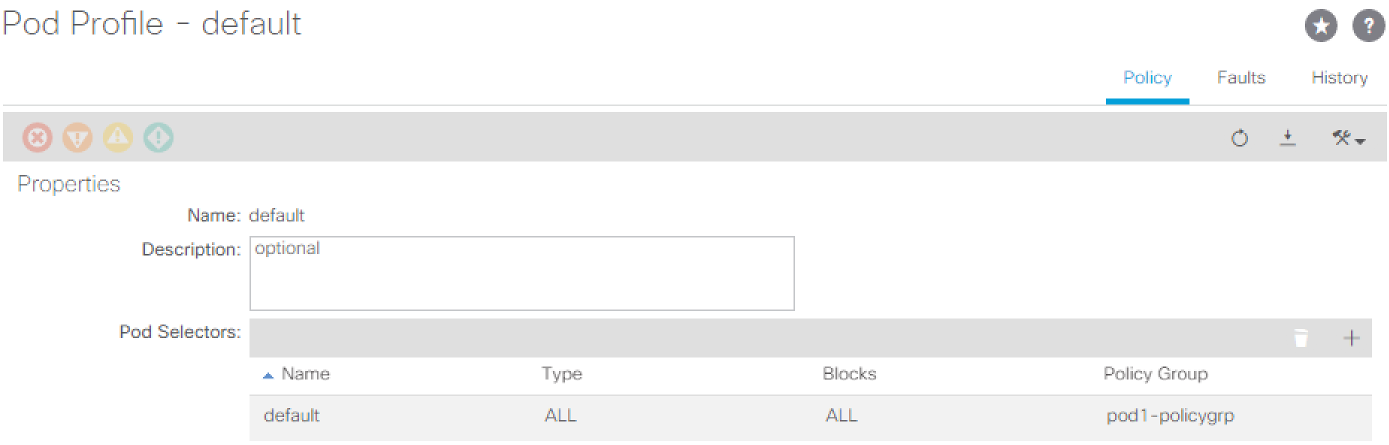

Name: default

Description: optional

Pod Selectors:

Name	Type	Blocks	Policy Group
default	ALL	ALL	pod1-policygrp

Verifying Fabric Wide Enforce Subnet Check for IP & MAC Learning

In this ACI deployment, Enforce Subnet Check for IP & MAC Learning should be enabled. To verify this setting, complete the following steps:

1. Select and expand System > System Settings > Fabric Wide Setting.

2. Ensure that Enforce Subnet Check is selected. If Enforce Subnet Check is not selected, select it and click Submit.

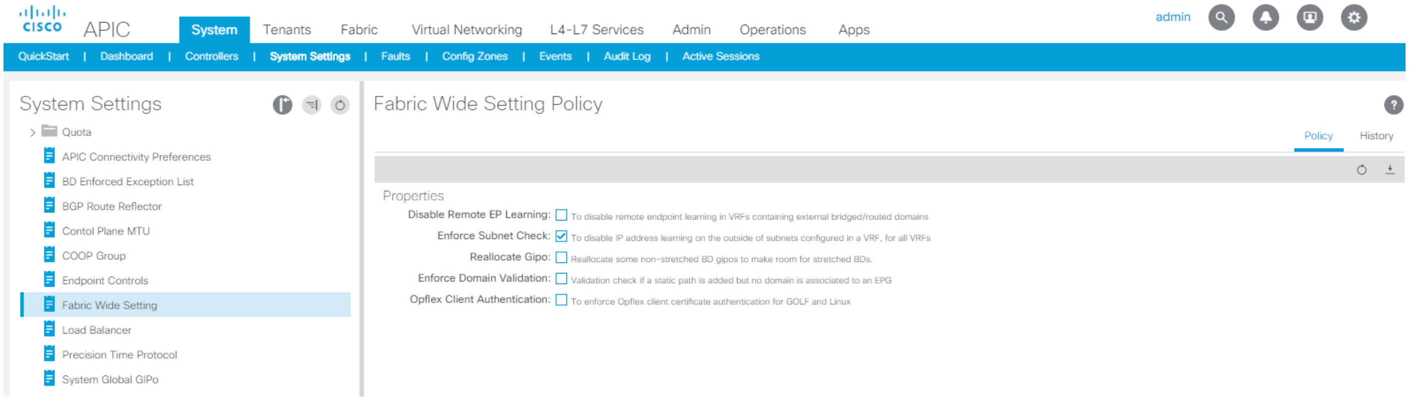

Disable Remote EP Learning: □ To disable remote endpoint learning in VRFs containing external bridged/routed domains

Enforce Subnet Check: ☑ To disable IP address learning on the outside of subnets configured in a VRF, for all VRFs

Reallocate Gipo: □ Reallocate some non-stretched BD gipos to make room for stretched BDs.

Enforce Domain Validation: □ Validation check if a static path is added but no domain is associated to an EPG

Opflex Client Authentication: □ To enforce Opflex client certificate authentication for GOLF and Linux

Verifying CoS Preservation Setting

In this FlexPod with ACI deployment, CoS Preservation should be turned off, complete the following steps to verify this setting:

1. Select and expand Fabric > Access Policies > Global Policies > QOS Class.

2. Ensure that Preserve QOS is not selected.

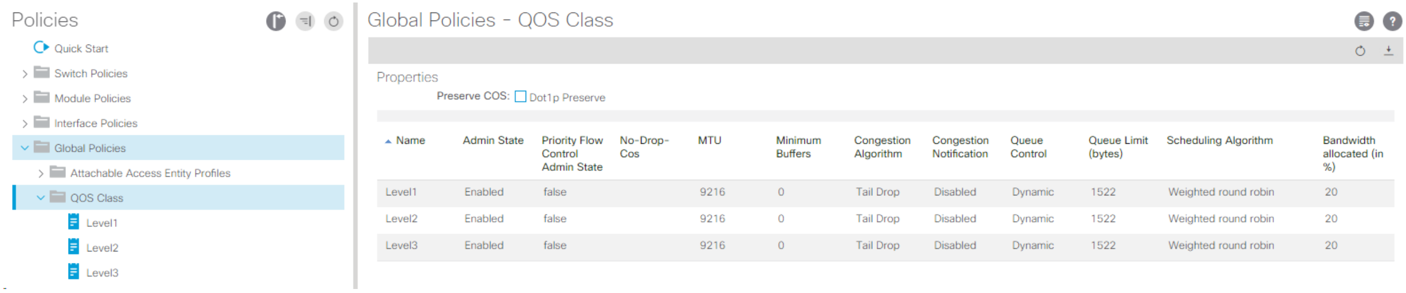

Global Policies - QOS Class

Properties

Preserve COS: □ Dot1p Preserve

Name	Admin State	Priority Flow Control Admin State	No-Drop-Cos	MTU	Minimum Buffers	Congestion Algorithm	Congestion Notification	Queue Control	Queue Limit (bytes)	Scheduling Algorithm	Bandwidth allocated (in %)
Level1	Enabled	false		9216	0	Tail Drop	Disabled	Dynamic	1522	Weighted round robin	20
Level2	Enabled	false		9216	0	Tail Drop	Disabled	Dynamic	1522	Weighted round robin	20
Level3	Enabled	false		9216	0	Tail Drop	Disabled	Dynamic	1522	Weighted round robin	20

In some shared deployments where other solutions in addition to FlexPod are being used, it may be necessary to enable this setting. The reason it is disabled in FlexPod is that NetApp AFF/FAS storage controllers inject CoS 4 on all tagged VLAN interfaces. Since FlexPod normally treats all UCS traffic as Best Effort (CoS 0), this setting is preferred. If it is desired to use other QoS classes than Best Effort in the UCS, then Preserve QOS should be enabled to preserve CoS marking on traffic going between the fabric interconnects. Note that the CoS 4 setting from the NetApp Storage Controllers will also be preserved and that should be considered when planning the QoS setup.

Fabric Access Policy Setup

This section details the steps to create various access policies creating parameters for CDP, LLDP, LACP, etc. These policies are used during vPC and VM domain creation. In an existing fabric, these policies may already exist. The existing policies can be used if configured the same way as listed. To define fabric access policies, complete the following steps:

1. Log into the APIC AGUI.

2. In the APIC UI, select and expand Fabric > Access Policies > Interface Policies > Policies.

Create Link Level Policies

This procedure will create link level policies for setting up the 1Gbps, 10Gbps, and 40Gbps link speeds. To create the link level policies, complete the following steps:

1. In the left pane, right-click Link Level and select **Create Link Level Policy**.

2. Name the policy as 1Gbps-Auto and select the 1Gbps Speed.

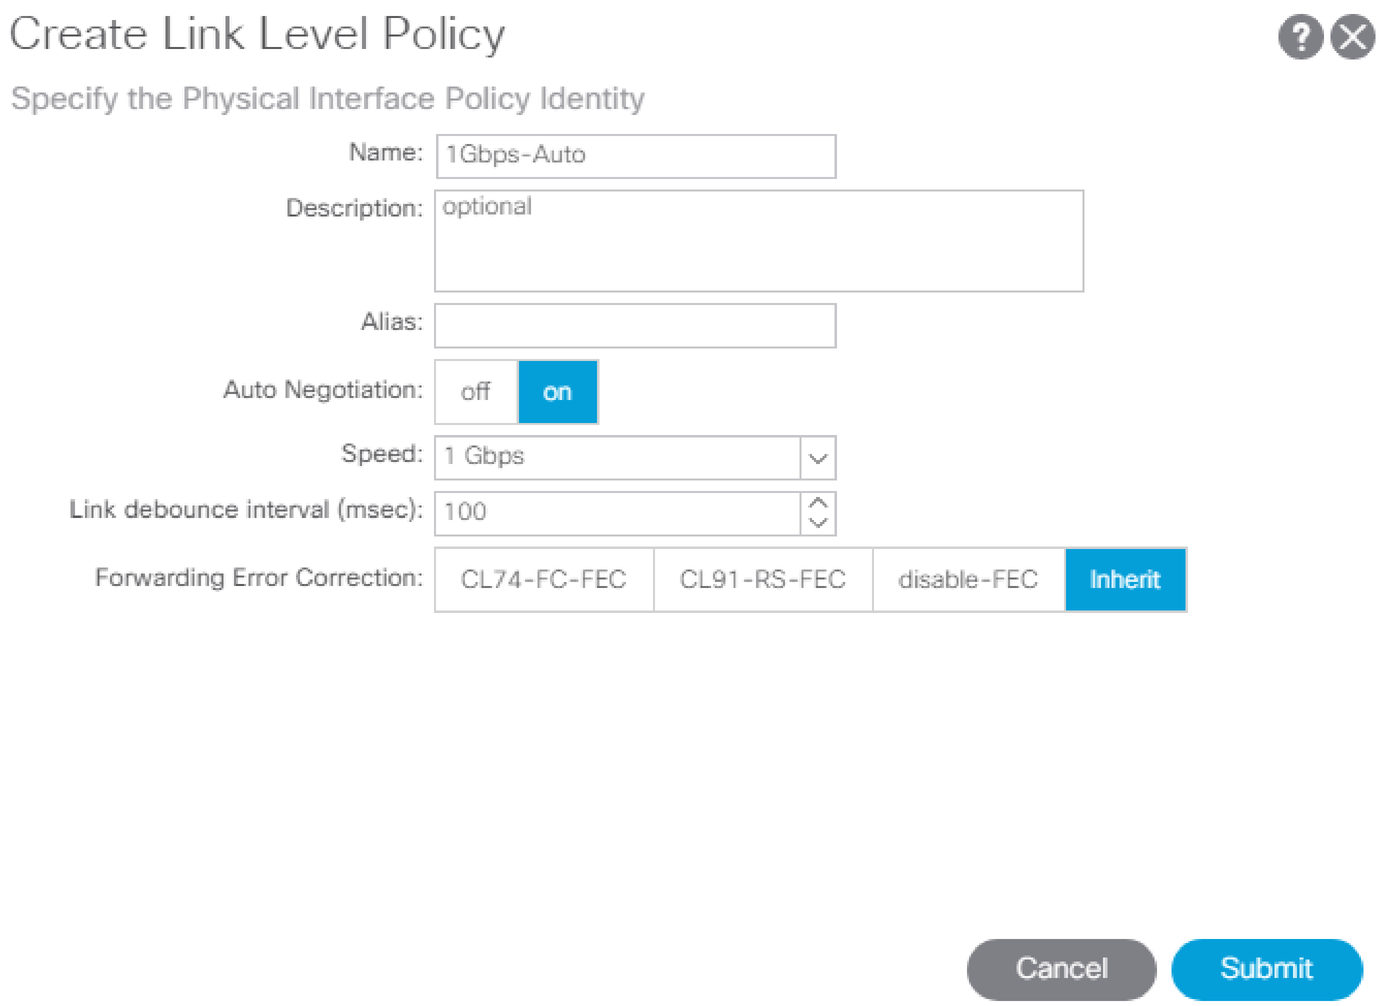

Create Link Level Policy

Specify the Physical Interface Policy Identity

Name: 1Gbps-Auto

Description: optional

Alias:

Auto Negotiation: off [✓ on]

Speed: 1 Gbps

Link debounce interval (msec): 100

Forwarding Error Correction: CL74-FC-FEC CL91-RS-FEC disable-FEC [✓ Inherit]

3. Click **Submit** to complete creating the policy.

4. In the left pane, right-click Link Level and select **Create Link Level Policy**.

5. Name the policy 10Gbps-Auto and select the 10Gbps Speed.

6. Click **Submit** to complete creating the policy.

7. In the left pane, right-click Link Level and select **Create Link Level Policy**.

8. Name the policy 40Gbps-Auto and select the 40Gbps Speed.

9. Click **Submit** to complete creating the policy.

Create CDP Policy

This procedure creates policies to enable or disable CDP on a link. To create a CDP policy, complete the following steps:

1. In the left pane, right-click CDP interface and select **Create CDP Interface Policy**.

2. Name the policy as CDP-Enabled and enable the Admin State.

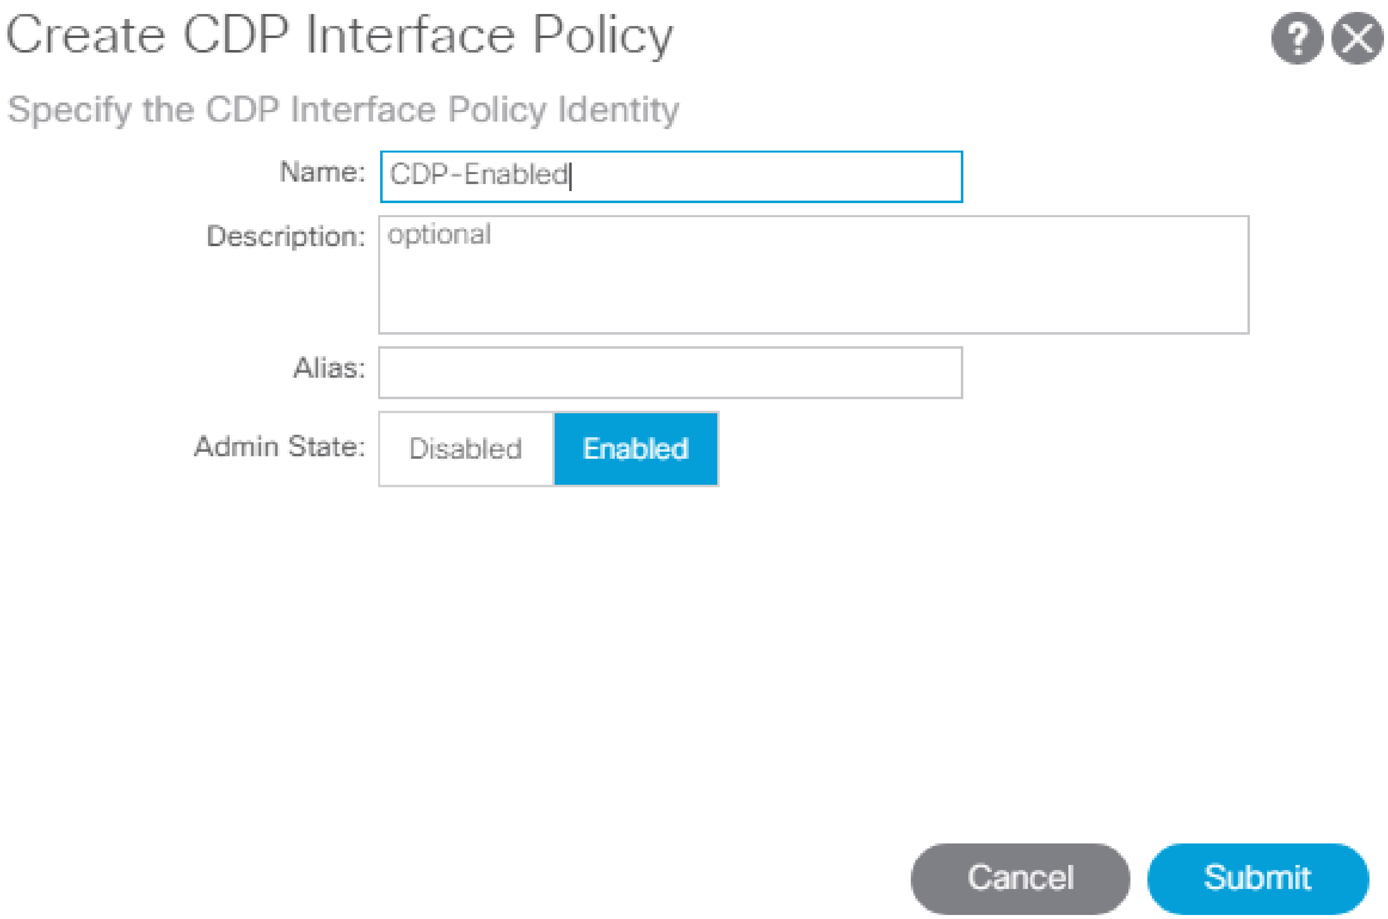

Create CDP Interface Policy

Specify the CDP Interface Policy Identity

Name: CDP-Enabled

Description: optional

Alias:

Admin State: Disabled [✓ Enabled]

3. Click **Submit** to complete creating the policy.

4. In the left pane, right-click the CDP Interface and select **Create CDP Interface Policy**.

5. Name the policy CDP-Disabled and disable the Admin State.

6. Click **Submit** to complete creating the policy.

Create LLDP Interface Policies

This procedure will create policies to enable or disable LLDP on a link. To create an LLDP Interface policy, complete the following steps:

1. In the left pane, right-click LLDP Interface and select **Create LLDP Interface Policy**.

2. Name the policy as LLDP-Enabled and enable both Transmit State and Receive State.

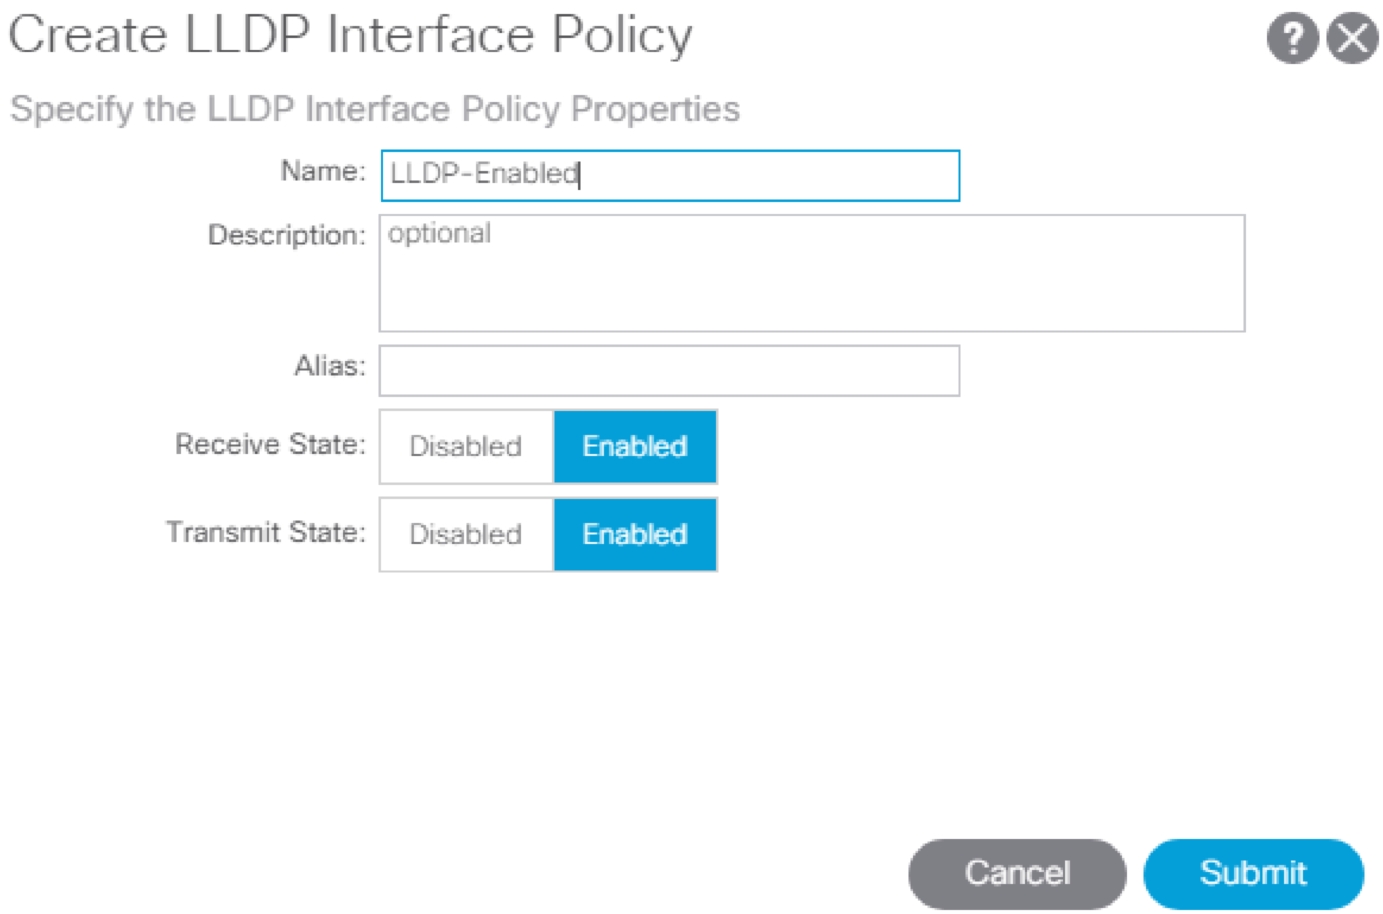

Create LLDP Interface Policy

Specify the LLDP Interface Policy Properties

Name: LLDP-Enabled

Description: optional

Alias:

Receive State: Disabled [✓ Enabled]

Transmit State: Disabled [✓ Enabled]

3. Click **Submit** to complete creating the policy.

4. In the left, right-click the LLDP Interface and select **Create LLDP Interface Policy**.

5. Name the policy as LLDP-Disabled and disable both the Transmit State and Receive State.

6. Click **Submit** to complete creating the policy.

Create Port Channel Policy

This procedure will create policies to set LACP active mode configuration and the MAC-Pinning mode configuration. To create the Port Channel policy, complete the following steps:

1. In the left pane, right-click Port Channel and select **Create Port Channel Policy**.

2. Name the policy as LACP-Active and select LACP Active for the Mode. Do not change any of the other values.

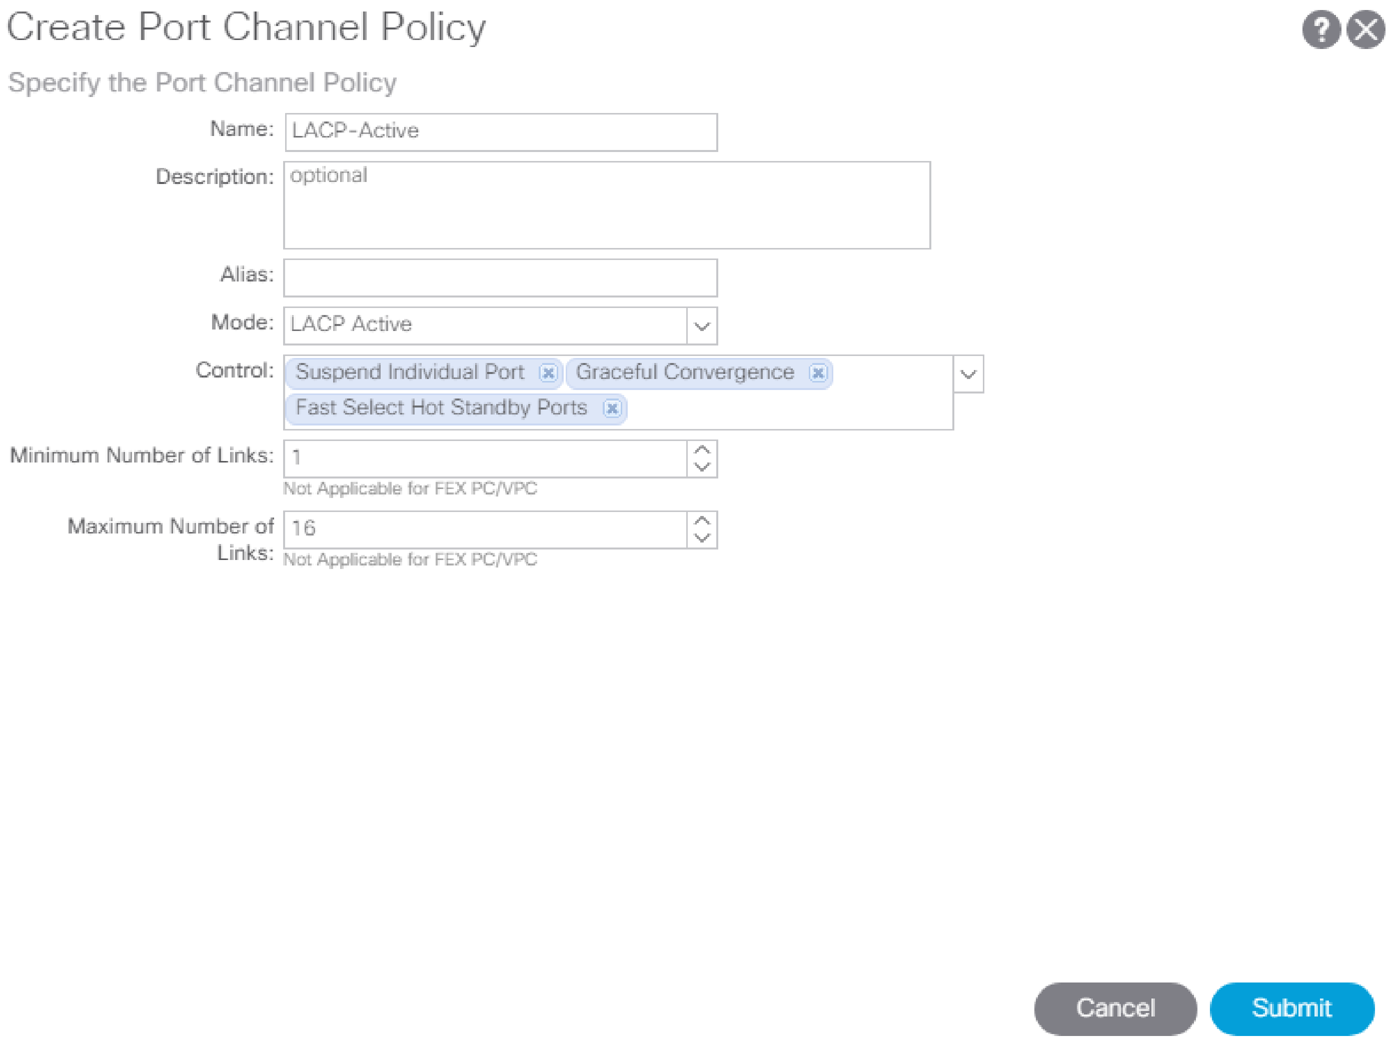

Create Port Channel Policy

Specify the Port Channel Policy

Name: LACP-Active

Description: optional

Alias:

Mode: LACP Active

Control: Suspend Individual Port, Graceful Convergence, Fast Select Hot Standby Ports

Minimum Number of Links: 1

Not Applicable for FEX PC/VPC

Maximum Number of Links: 16

Not Applicable for FEX PC/VPC

3. Click **Submit** to complete creating the policy.

4. In the left pane, right-click Port Channel and select **Create Port Channel Policy**.

5. Name the policy as MAC-Pinning and select MAC Pinning-Physical-NIC-load for the Mode. Do not change any of the other values.

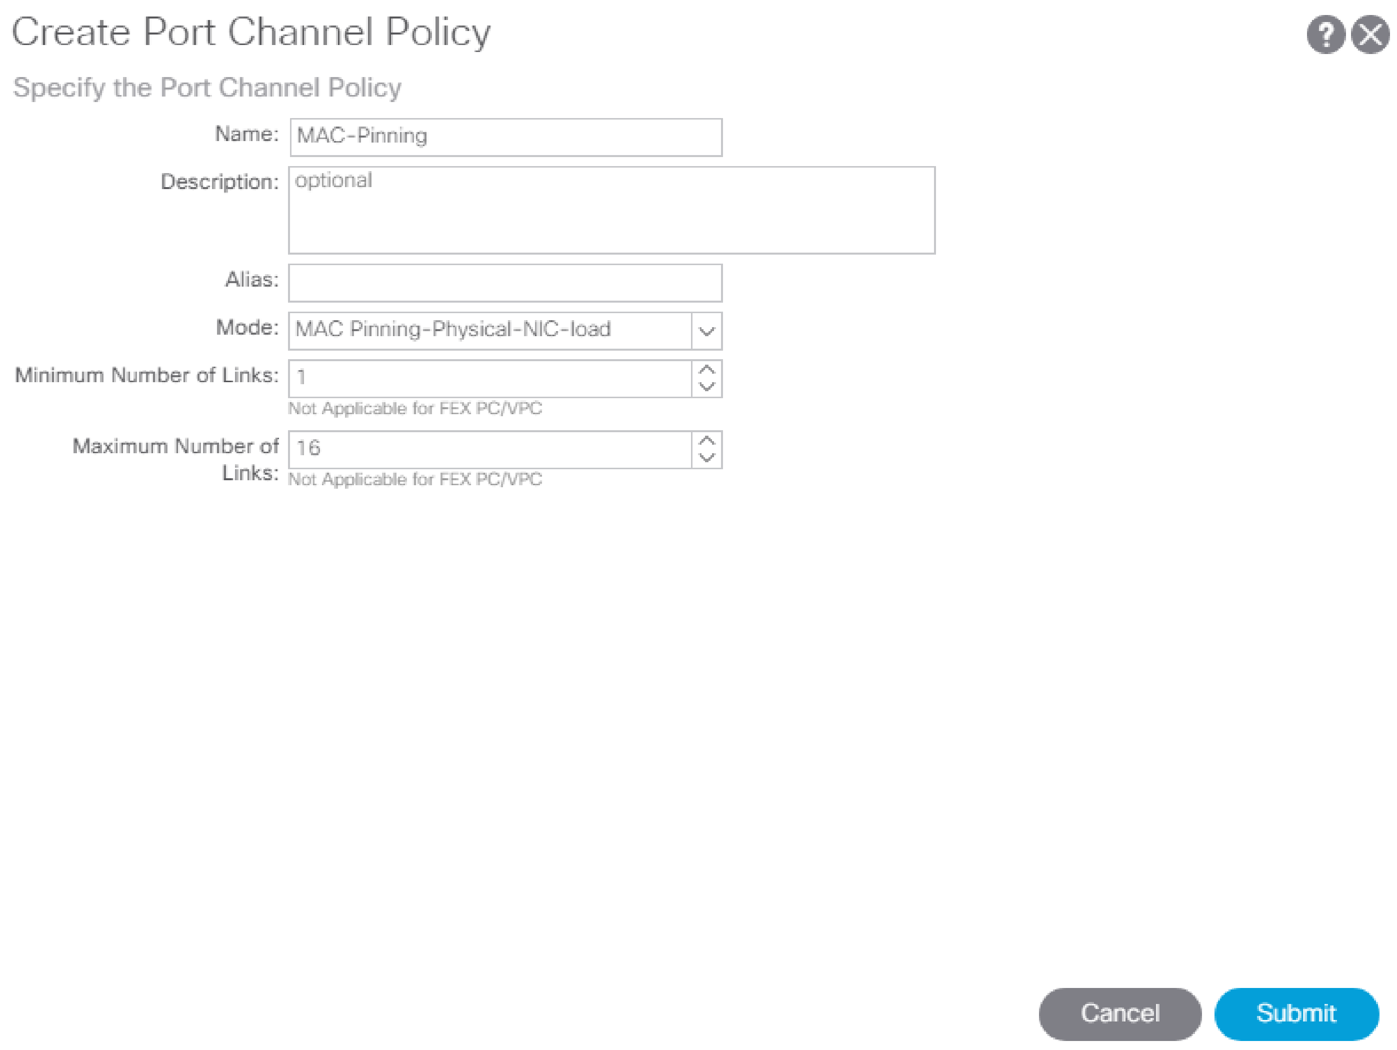

Create Port Channel Policy

Specify the Port Channel Policy

Name: MAC-Pinning

Description: optional

Alias:

Mode: MAC Pinning-Physical-NIC-load

Minimum Number of Links: 1
Not Applicable for FEX PC/VPC

Maximum Number of Links: 16
Not Applicable for FEX PC/VPC

6. Click **Submit** to complete creating the policy.

7. In the left pane, right-click Port Channel and select **Create Port Channel Policy**.

Create BPDU Filter/Guard Policies

This procedure will create policies to enable or disable BPDU filter and guard. To create a BPDU filter/Guard policy, complete the following steps:

1. In the left pane, right-click Spanning Tree Interface and select **Create Spanning Tree Interface Policy**.

2. Name the policy as BPDU-FG-Enabled and select both the BPDU filter and BPDU Guard Interface Controls.

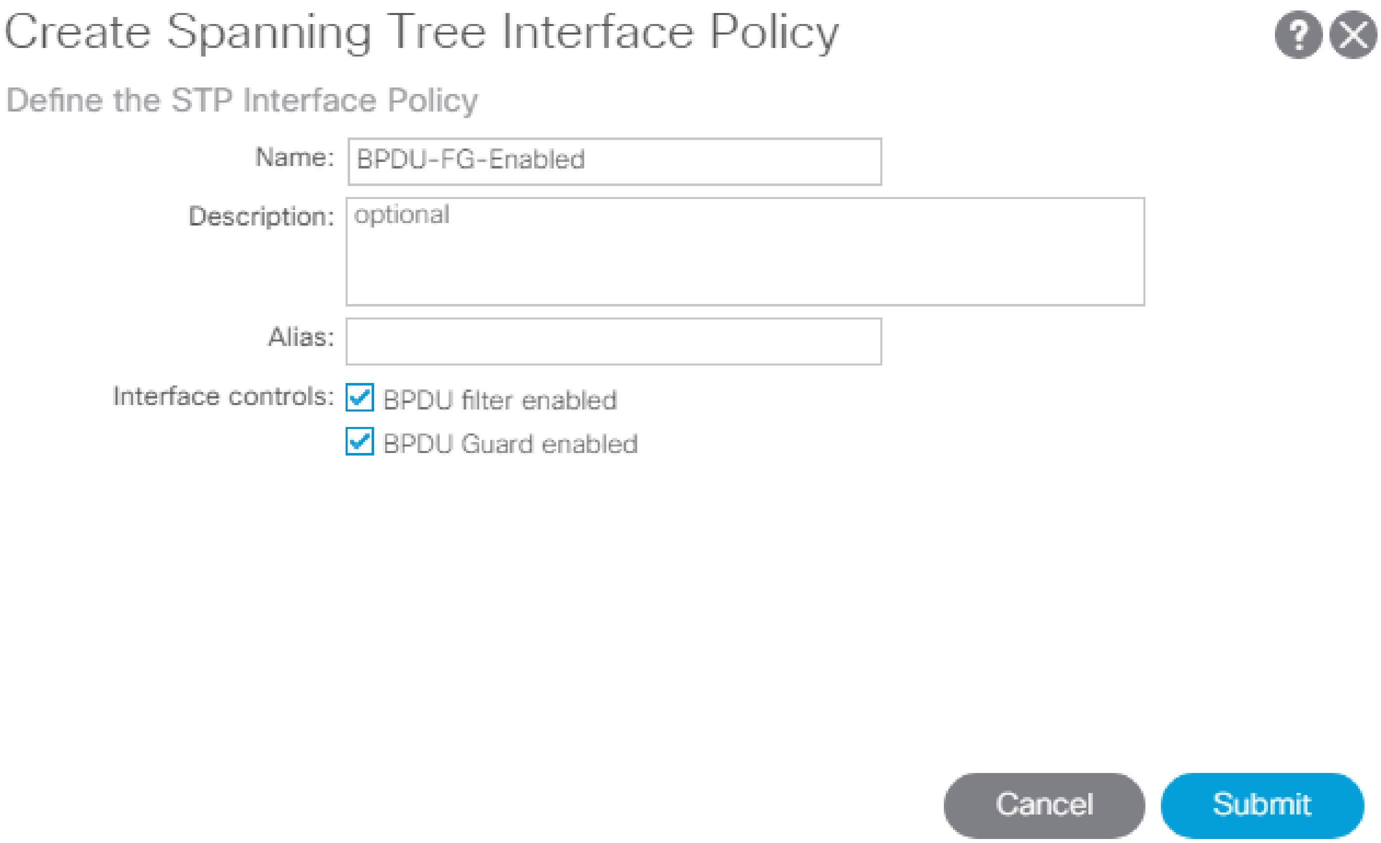

Create Spanning Tree Interface Policy

Define the STP Interface Policy

Name: BPDU-FG-Enabled

Description: optional

Alias:

Interface controls: ☑ BPDU filter enabled

☑ BPDU Guard enabled

3. Click **Submit** to complete creating the policy.

4. In the left pane, right-click Spanning Tree Interface and select **Create Spanning Tree Interface Policy**.

5. Name the policy as BPDU-FG-Disabled and make sure both the BPDU filter and BPDU Guard Interface Controls are cleared.

6. Click **Submit** to complete creating the policy.

Create VLAN Scope Policy

To create policies to enable port local scope for all the VLANs, complete the following steps:

1. In the left pane, right-click the L2 Interface and select **Create L2 Interface Policy**.

2. Name the policy as **VLAN-Scope-Port-Local** and make sure **Port Local scope** is selected for VLAN Scope. Do not change any of the other values.

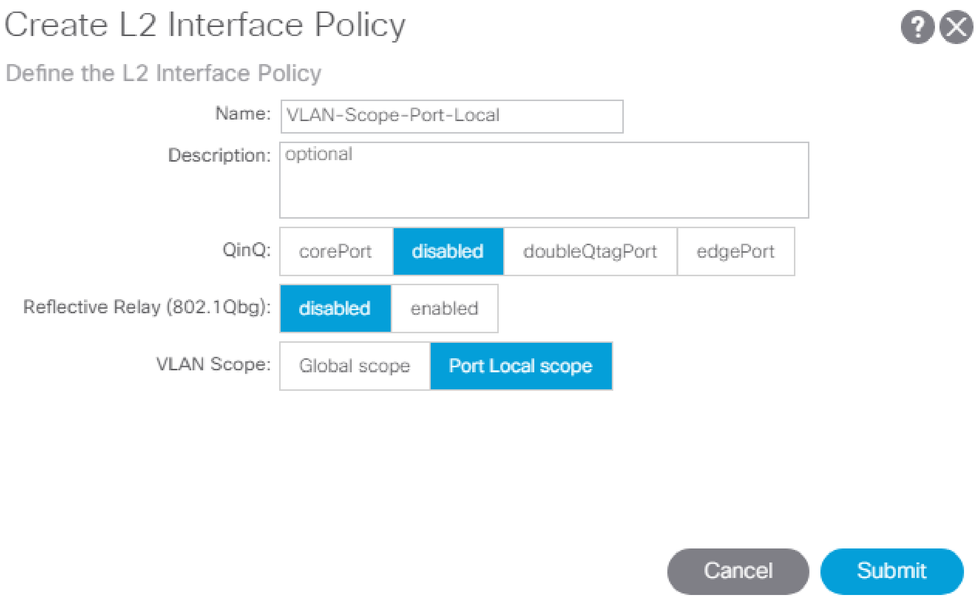

Create L2 Interface Policy

Define the L2 Interface Policy

Name: VLAN-Scope-Port-Local

Description: optional

QinQ: corePort [✓ disabled] doubleQtagPort edgePort

Reflective Relay (802.1Qbg): [✓ disabled] enabled

VLAN Scope: Global scope [✓ Port Local scope]

3. Click **Submit** to complete creating the policy.

4. Repeat above steps to create a **VLAN-Scope-Global Policy** and make sure **Global scope** is selected for VLAN Scope. Do not change any of the other values. See below.

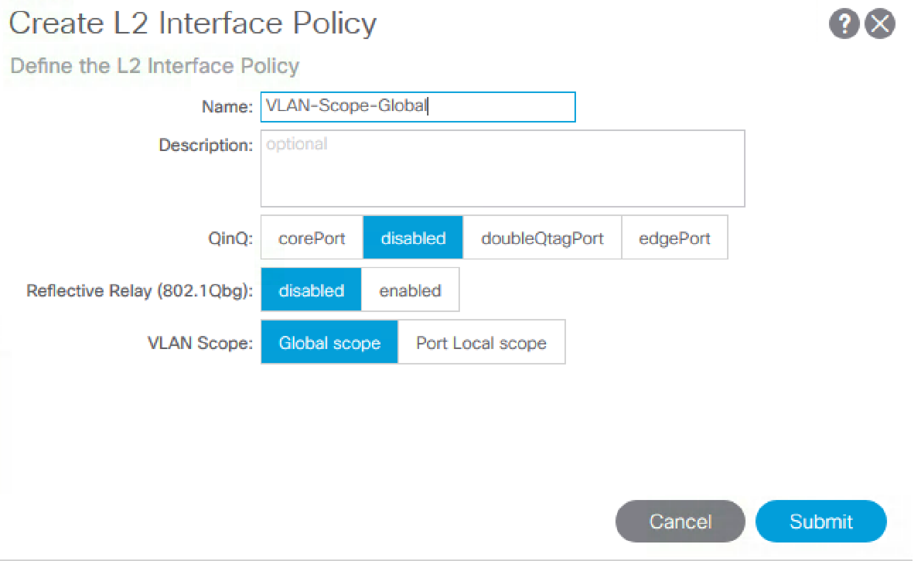

Create L2 Interface Policy

Define the L2 Interface Policy

Name: VLAN-Scope-Global

Description: optional

QinQ: corePort [✓ disabled] doubleQtagPort edgePort

Reflective Relay (802.1Qbg): [✓ disabled] enabled

VLAN Scope: [✓ Global scope] Port Local scope

Create Firewall Policy

To create policies to disable a firewall, complete the following steps:

1. In the left pane, right-click Firewall and select Create Firewall Policy.

2. Name the policy Firewall-Disabled and select Disabled for Mode. Do not change any of the other values.

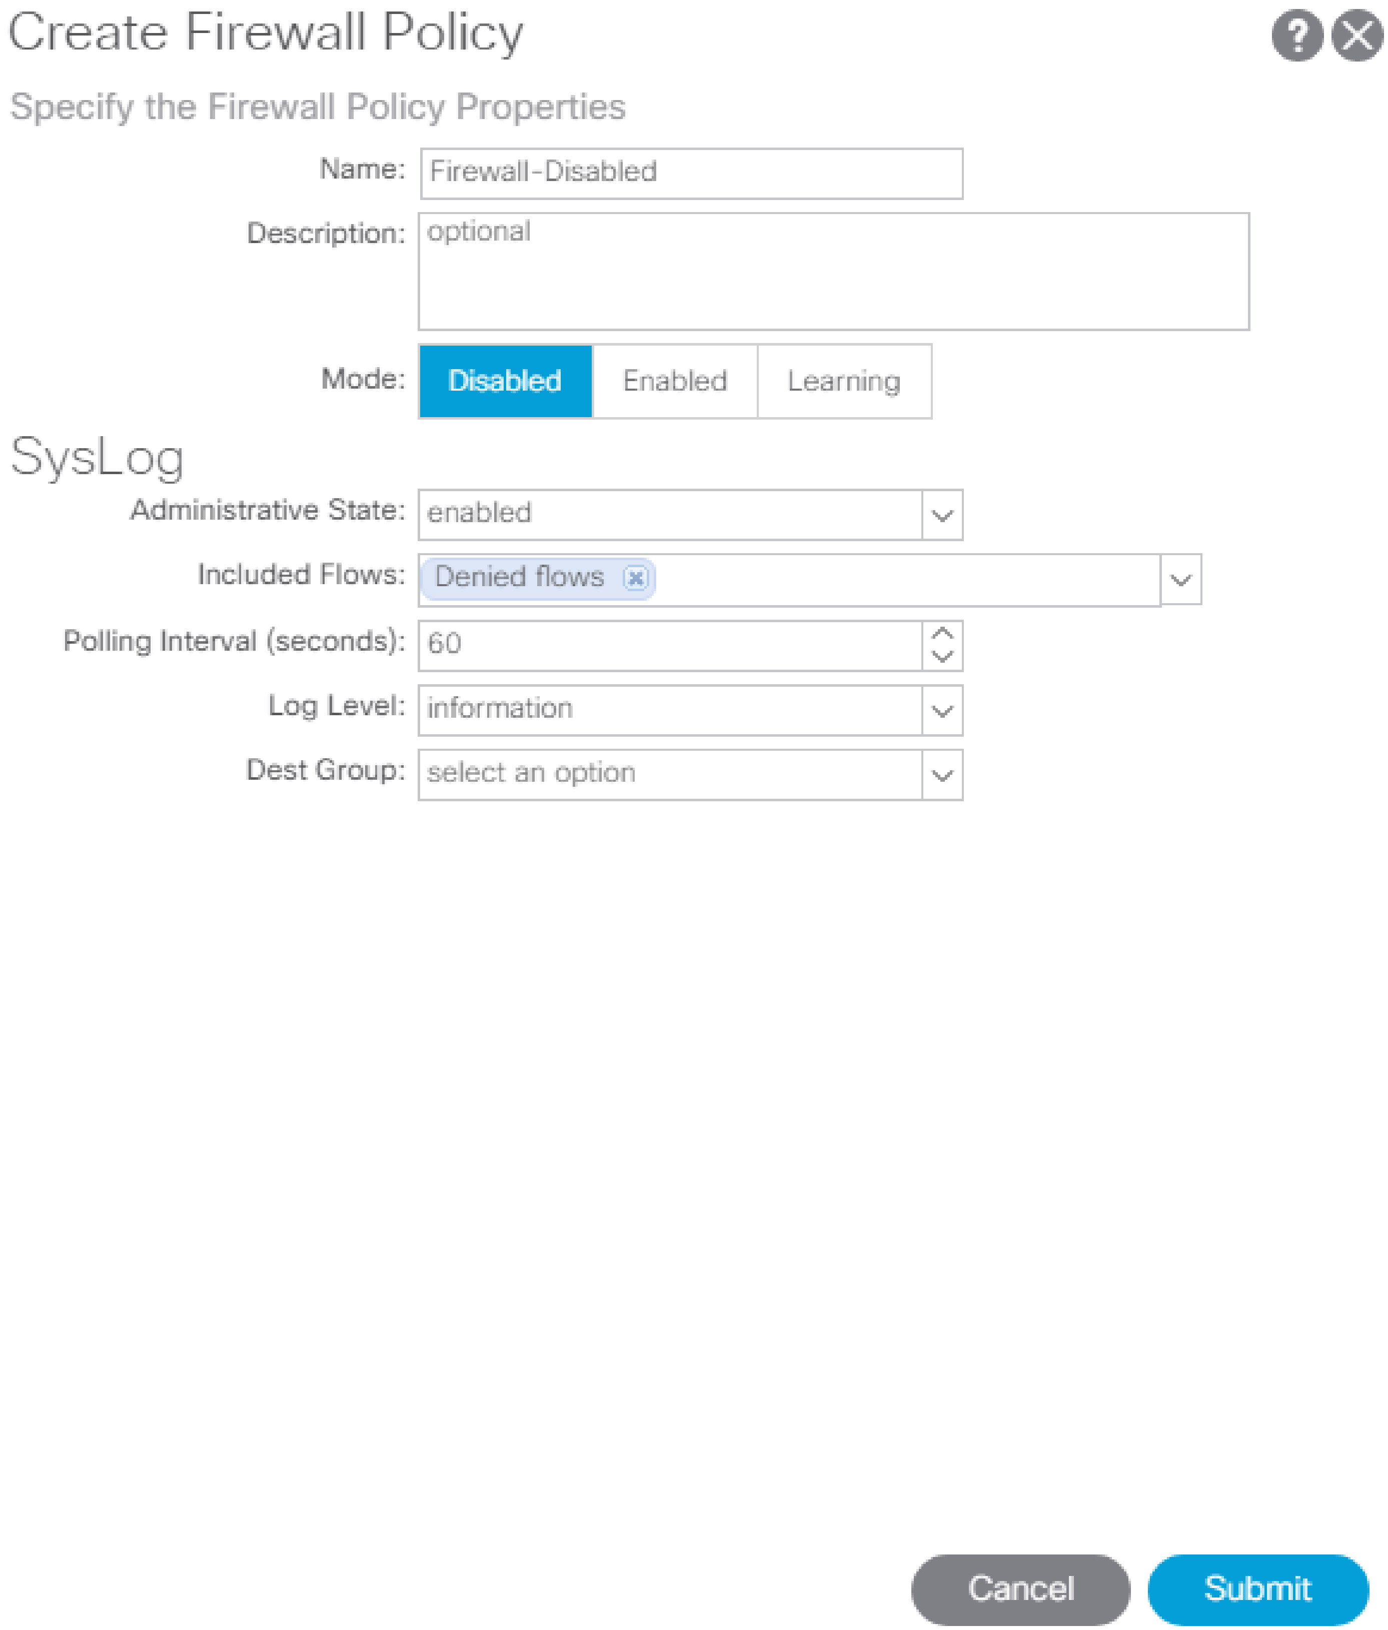

Create Firewall Policy

Specify the Firewall Policy Properties

Name: Firewall-Disabled

Description: optional

Mode: Disabled | Enabled | Learning

SysLog

Administrative State: enabled

Included Flows: Denied flows

Polling Interval (seconds): 60

Log Level: information

Dest Group: select an option

3. Click **Submit** to complete creating the policy.

Create Virtual Port Channels (vPCs)

This section details the steps to setup vPCs for connectivity to the In-Band Management Network, Cisco UCS, and NetApp Storage.

VPC - Management Switch

To setup a vPC for connectivity to the existing In-Band Management Network, complete the following steps:

This deployment guide covers the configuration for a single, pre-existing Cisco Nexus management switch. You can adjust the management configuration depending on your connectivity setup. The In-Band Management Network provides connectivity of Management Virtual Machines and Hosts in the ACI fabric to existing services on the In-Band Management network outside of the ACI fabric. Layer 3 connectivity outside of the ACI Fabric is assumed between the In-Band and Out-of-Band Management networks. This setup creates management networks that are physically isolated from tenant networks. In this validation, a 10GE vPC from two 10GE capable leaf switches in the fabric is connected to a port-channel on a Nexus 5K switch outside the fabric. Note that this vPC is not created on the Nexus 9332 leaves, but on other existing leaves that have 10GE ports.

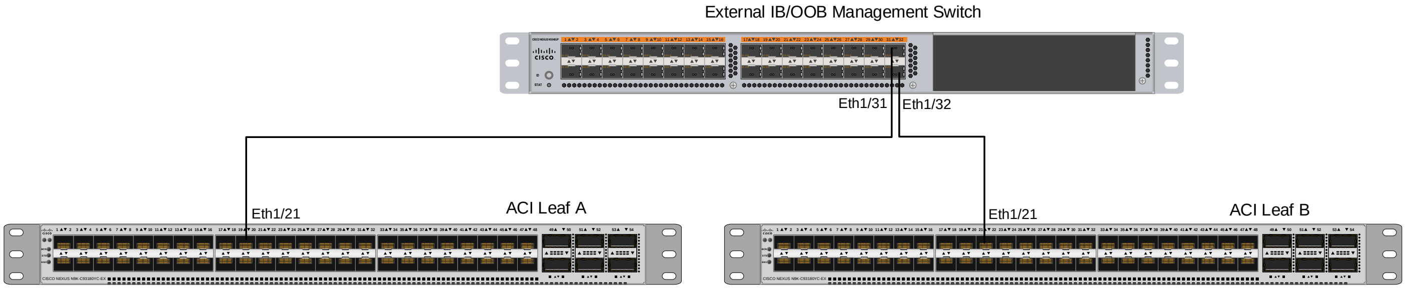

Table 4 VLAN for Incoming IB-MGMT

Name	VLAN
IB-MGMT	<118>

1. In the APIC GUI, at the top select Fabric > Access Policies > Quick Start.

2. In the right pane select Configure an interface, PC and VPC.

3. In the configuration window, configure a VPC domain between the 10GE capable leaf switches by clicking “+” under VPC Switch Pairs. If a VPC Domain already exists between the two switches being used for this vPC, skip to step 7.

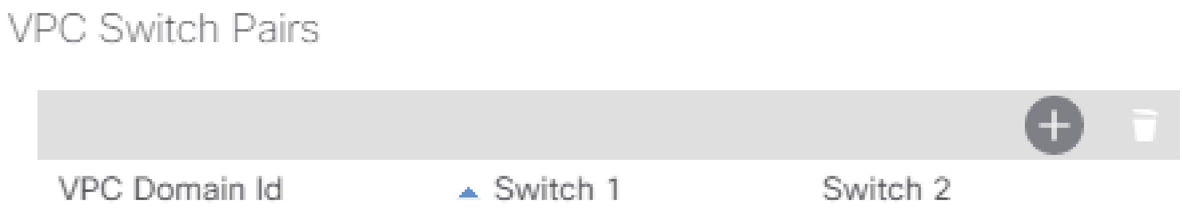

4. Enter a VPC Domain ID (1 in this example).

5. From the drop-down list, select Switch A and Switch B IDs to select the two leaf switches.

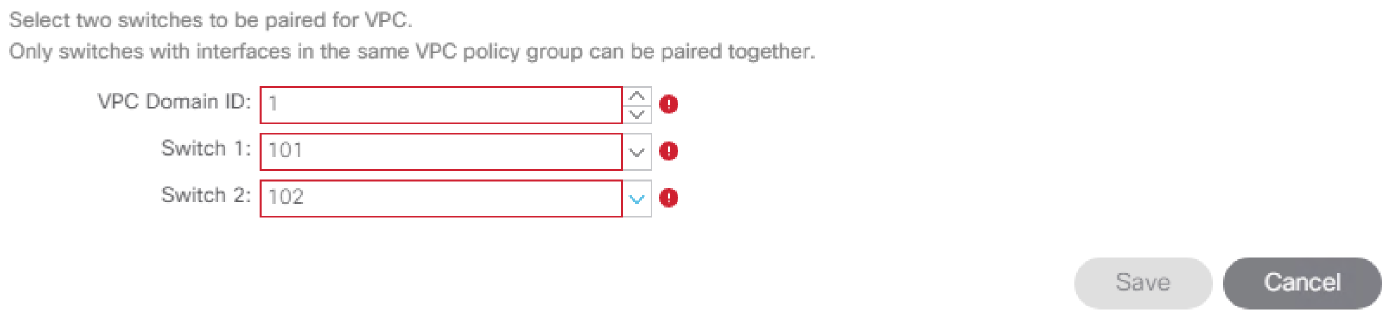

Select two switches to be paired for VPC.
Only switches with interfaces in the same VPC policy group can be paired together.

VPC Domain ID: 1

Switch 1: 101

Switch 2: 102

6. Click **SAVE.**

7. If a profile for the two leaf switches being used does not already exist under Configured Switch Interfaces, click the “+” under Configured Switch Interfaces. If the profile does exist, select it and proceed to step 10.

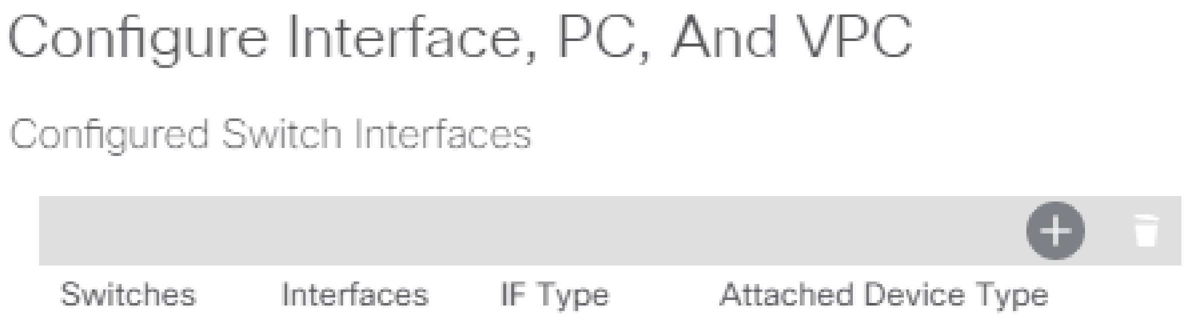

Configure Interface, PC, And VPC

Configured Switch Interfaces

8. From the Switches drop-down list on the right, select both the leaf switches being used for this vPC.

9. Leave the system generated Switch Profile Name in place.

10. Click the big green “+” on the right to configure switch interfaces.

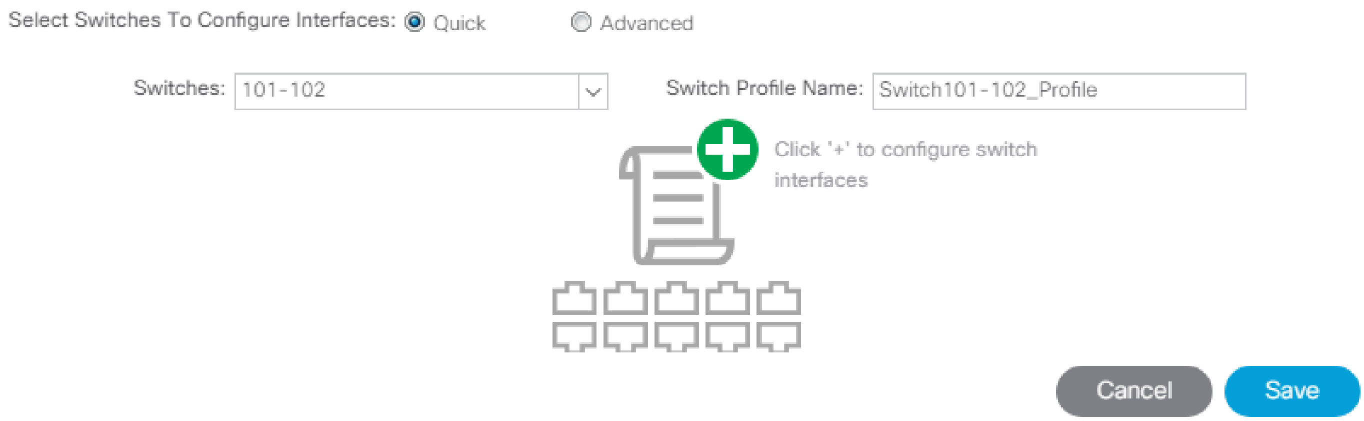

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: Switch101-102_Profile

Click '+' to configure switch interfaces

11. Configure various fields as shown in the figure below. In this screen shot, port 1/21 on both leaf switches is connected to a Nexus switch using 10Gbps links.

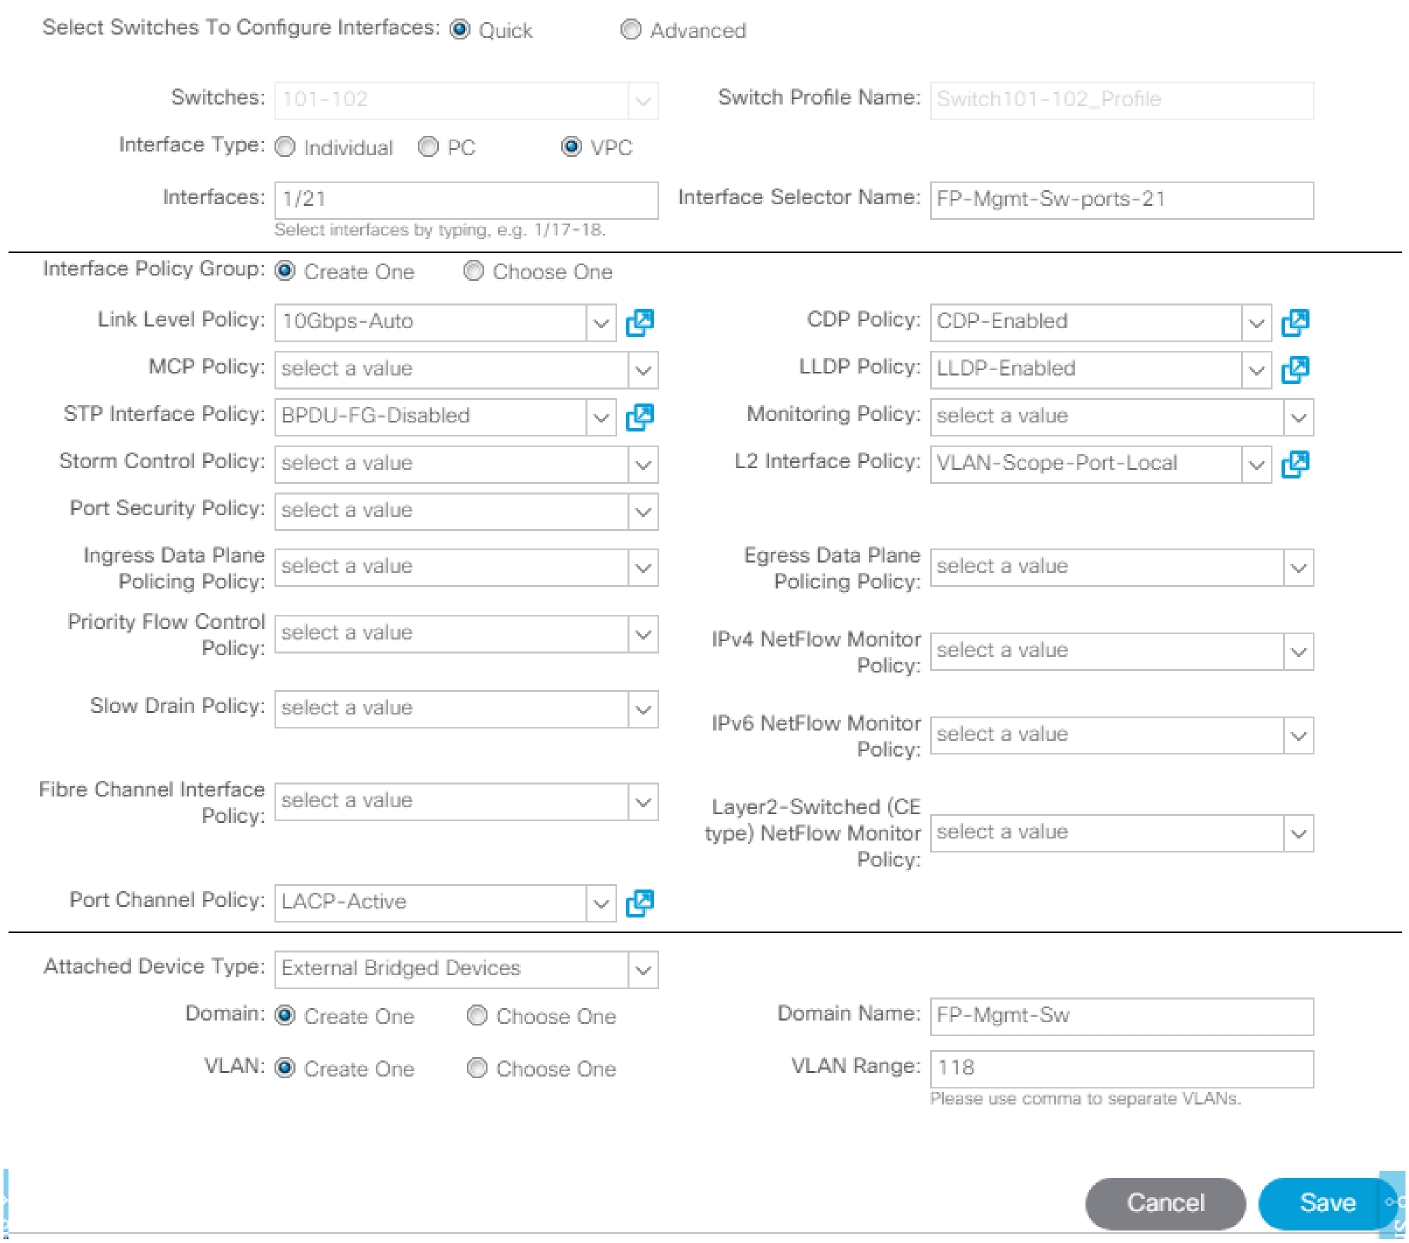

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: Switch101-102_Profile

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/21

Select interfaces by typing, e.g. 1/17-18.

Interface Selector Name: FP-Mgmt-Sw-ports-21

Interface Policy Group: ● Create One ○ Choose One

Link Level Policy: 10Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: select a value

LLDP Policy: LLDP-Enabled

STP Interface Policy: BPDU-FG-Disabled

Monitoring Policy: select a value

Storm Control Policy: select a value

L2 Interface Policy: VLAN-Scope-Port-Local

Port Security Policy: select a value

Ingress Data Plane Policing Policy: select a value

Egress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

IPv4 NetFlow Monitor Policy: select a value

Slow Drain Policy: select a value

IPv6 NetFlow Monitor Policy: select a value

Fibre Channel Interface Policy: select a value

Layer2-Switched (CE type) NetFlow Monitor Policy: select a value

Port Channel Policy: LACP-Active

Attached Device Type: External Bridged Devices

Domain: ● Create One ○ Choose One

Domain Name: FP-Mgmt-Sw

VLAN: ● Create One ○ Choose One

VLAN Range: 118

Please use comma to separate VLANs.

12. Click **Save.**

13. Click **Save** again to finish the configuring switch interfaces.

14. Click **Submit**.

To validate the configuration, log into the Nexus switch and verify the port-channel is up (`show port-channel summary`).

VPC – UCS Fabric Interconnects

Complete the following steps to setup vPCs for connectivity to the UCS Fabric Interconnects.

Figure 3 VLANs Configured for Cisco UCS

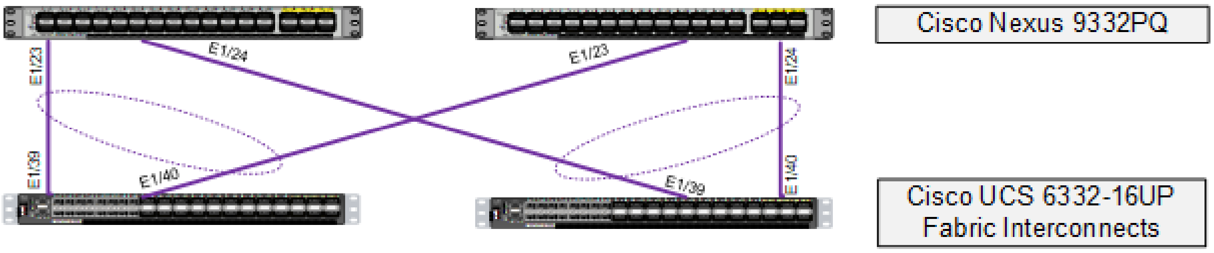

Table 5 VLANs for Cisco UCS Hosts

Name	VLAN
Native	<2>
Core-Services	<318>
AV-IB-Mgmt	<419>
AV-vMotion	<3000>
AV-Infra-NFS	<3150>
AV-Infra-iSCSI-A	<3110>
AV-Infra-iSCSI-B	<3120>

1. In the APIC GUI, select Fabric > Access Policies > Quick Start.

2. In the right pane, select Configure and interface, PC and VPC.

3. In the configuration window, configure a VPC domain between the 9332 leaf switches by clicking "+" under VPC Switch Pairs.

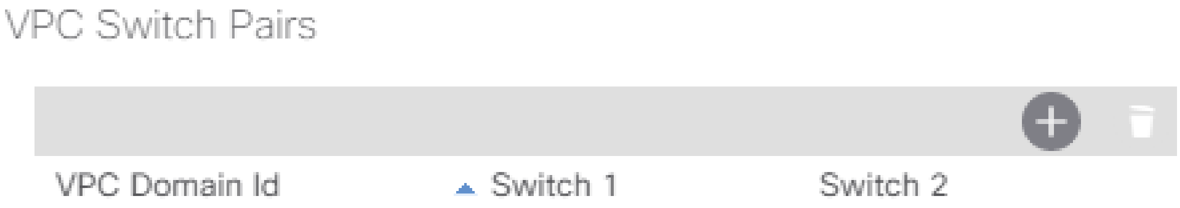

4. Enter a VPC Domain ID (10 in this example).

5. From the drop-down list, select 9332 Switch A and 9332 Switch B IDs to select the two leaf switches.

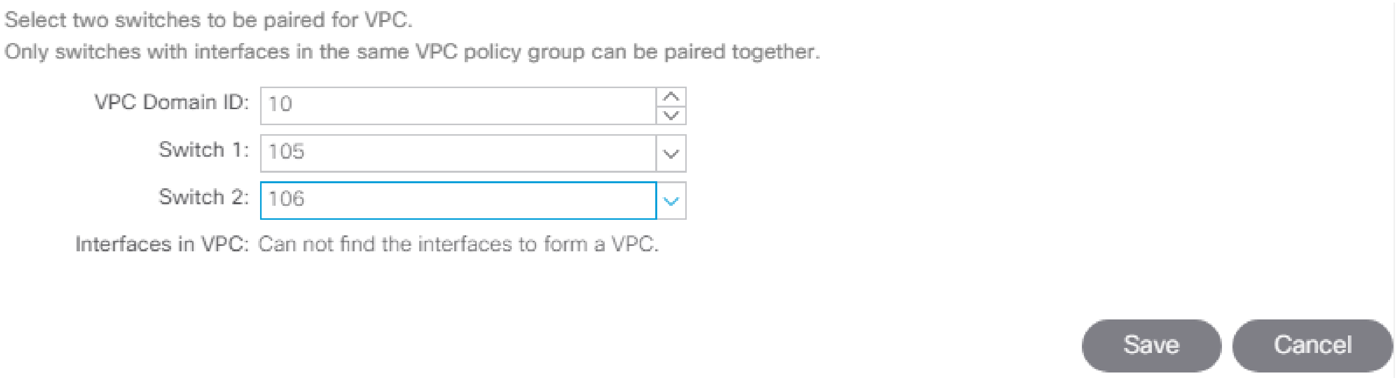

Select two switches to be paired for VPC.
Only switches with interfaces in the same VPC policy group can be paired together.

VPC Domain ID: 10

Switch 1: 105

Switch 2: 106

Interfaces in VPC: Can not find the interfaces to form a VPC.

6. Click **Save.**

7. Click the “+” under Configured Switch Interfaces.

8. Select the two Nexus 9332 switches under the Switches drop-down list.

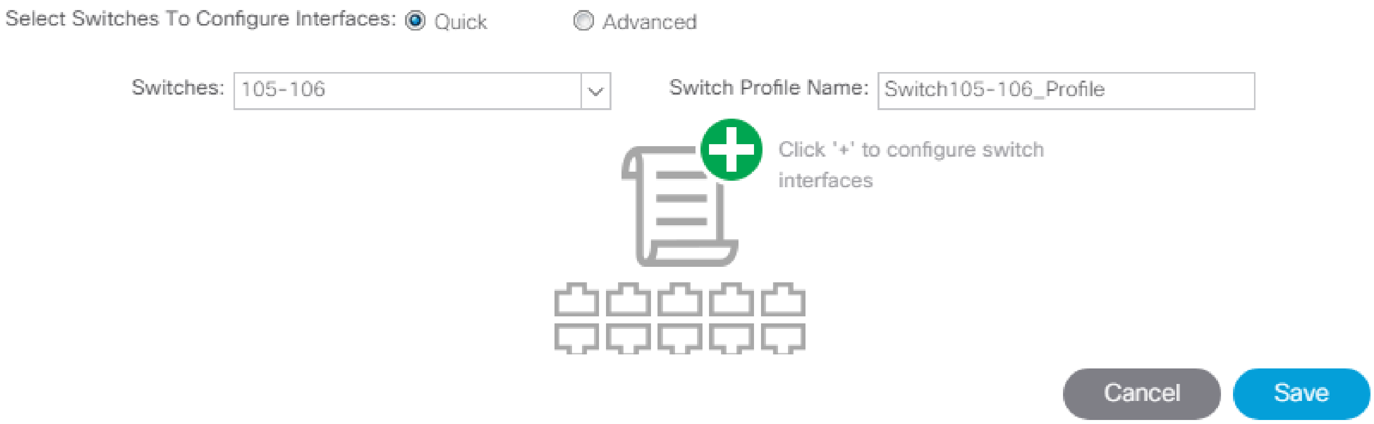

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 105-106

Switch Profile Name: Switch105-106_Profile

Click '+' to configure switch interfaces

9. Click ⊕ to add switch interfaces.

10. Configure various fields as shown in the figure below. In this screenshot, port 1/23 on both leaf switches is connected to UCS Fabric Interconnect A using 40Gbps links.

It is not recommended to use ports 1/25 and 1/26 for port-channels on a Nexus 9332PQ.

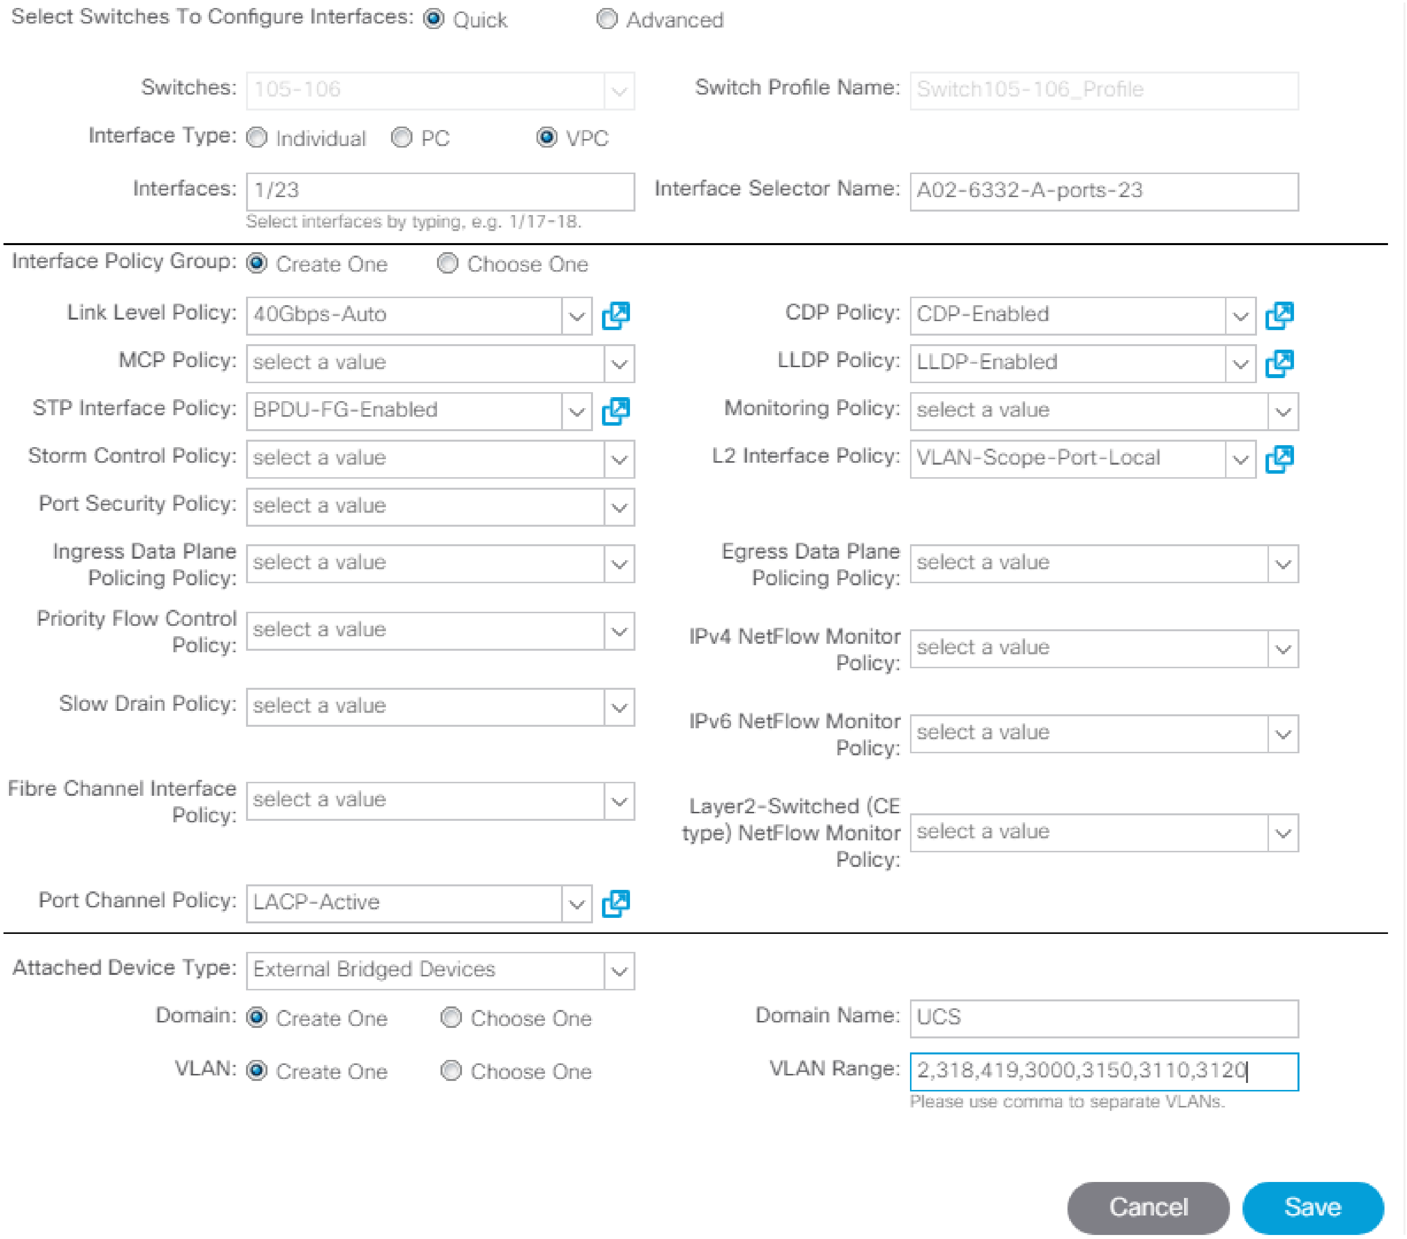

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 105-106

Switch Profile Name: Switch105-106_Profile

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/23

Select interfaces by typing, e.g. 1/17-18.

Interface Selector Name: A02-6332-A-ports-23

Interface Policy Group: ● Create One ○ Choose One

Link Level Policy: 40Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: select a value

LLDP Policy: LLDP-Enabled

STP Interface Policy: BPDU-FG-Enabled

Monitoring Policy: select a value

Storm Control Policy: select a value

L2 Interface Policy: VLAN-Scope-Port-Local

Port Security Policy: select a value

Ingress Data Plane Policing Policy: select a value

Egress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

IPv4 NetFlow Monitor Policy: select a value

Slow Drain Policy: select a value

IPv6 NetFlow Monitor Policy: select a value

Fibre Channel Interface Policy: select a value

Layer2-Switched (CE type) NetFlow Monitor Policy: select a value

Port Channel Policy: LACP-Active

Attached Device Type: External Bridged Devices

Domain: ● Create One ○ Choose One

Domain Name: UCS

VLAN: ● Create One ○ Choose One

VLAN Range: 2,318,419,3000,3150,3110,3120

Please use comma to separate VLANs.

11. Click **Save.**

12. Click **Save** again to finish the configuring switch interfaces.

13. Click **Submit**.

14. From the right pane, select Configure and interface, PC and VPC.

15. Select the switches configured in the last step under Configured Switch Interfaces.

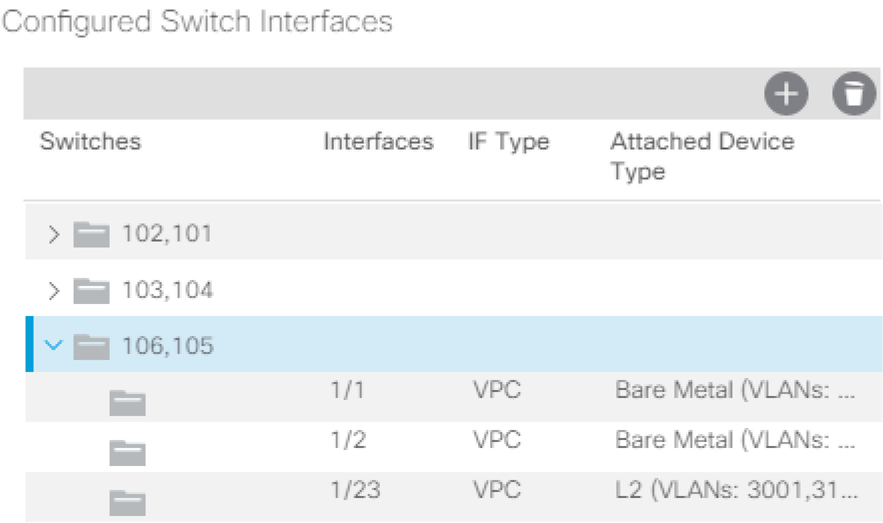

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
102,101			
103,104			
106,105			
	1/1	VPC	Bare Metal (VLANs: ...
	1/2	VPC	Bare Metal (VLANs: ...
	1/23	VPC	L2 (VLANs: 3001,31...

16. Click ➕ on the right to add switch interfaces.

17. Configure various fields as shown in the screenshot. In this screenshot, port 1/24 on both leaf switches is connected to UCS Fabric Interconnect B using 40Gbps links. Instead of creating a new domain, the External Bridged Device created in the last step (UCS) is attached to the FI-B as shown below.

It is not recommended to use ports 1/25 and 1/26 for port-channels on a Nexus 9332PQ.

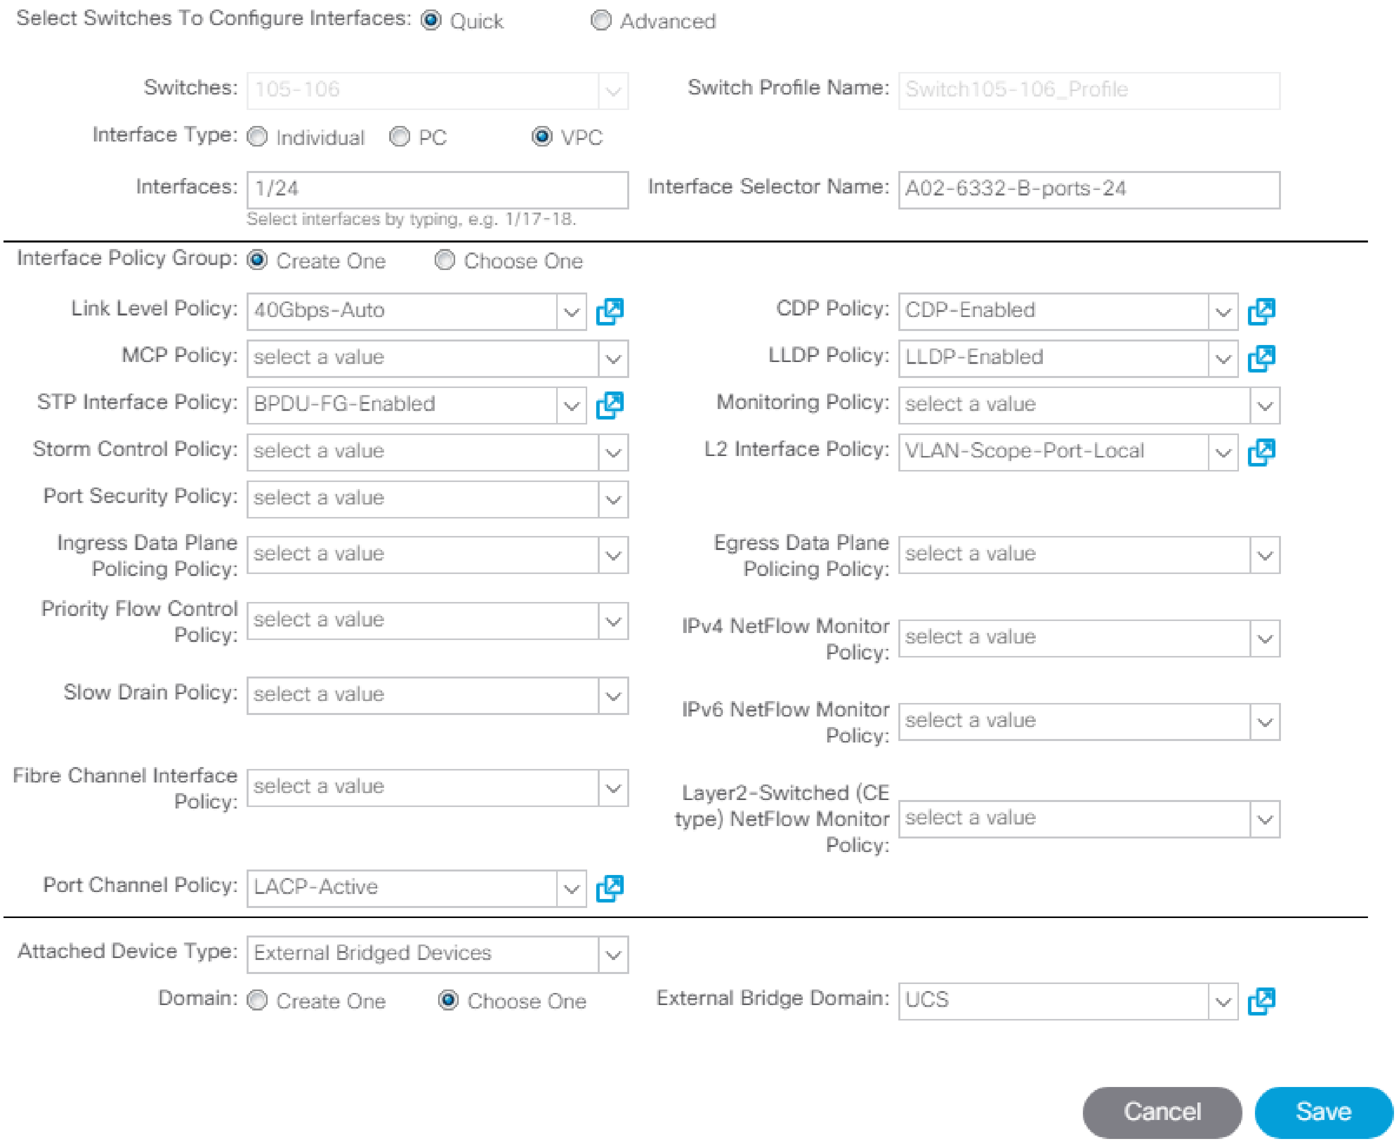

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 105-106

Switch Profile Name: Switch105-106_Profile

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/24

Select interfaces by typing, e.g. 1/17-18.

Interface Selector Name: A02-6332-B-ports-24

Interface Policy Group: ● Create One ○ Choose One

Link Level Policy: 40Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: select a value

LLDP Policy: LLDP-Enabled

STP Interface Policy: BPDU-FG-Enabled

Monitoring Policy: select a value

Storm Control Policy: select a value

L2 Interface Policy: VLAN-Scope-Port-Local

Port Security Policy: select a value

Ingress Data Plane Policing Policy: select a value

Egress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

IPv4 NetFlow Monitor Policy: select a value

Slow Drain Policy: select a value

IPv6 NetFlow Monitor Policy: select a value

Fibre Channel Interface Policy: select a value

Layer2-Switched (CE type) NetFlow Monitor Policy: select a value

Port Channel Policy: LACP-Active

Attached Device Type: External Bridged Devices

Domain: ○ Create One ● Choose One

External Bridge Domain: UCS

18. Click **Save.**

19. Click **Save** again to finish the configuring switch interfaces.

20. Click **Submit**.

21. **Optional:** Repeat this procedure to configure any additional UCS domains. For a uniform configuration, the External Bridge Domain (UCS) will be utilized for all the Fabric Interconnects.

VPC – NetApp AFF Cluster

Complete the following steps to setup vPCs for connectivity to the NetApp AFF storage controllers. The VLANs configured for NetApp are shown in Table 6 .

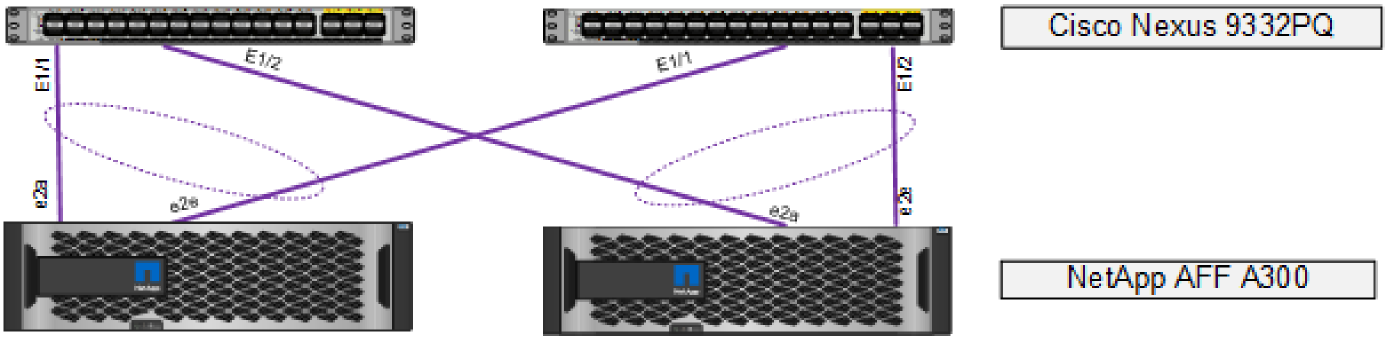

Table 6 VLANs for Storage

Name	VLAN
SVM-MGMT	<219>
NFS	<3050>
iSCSI-A	<3010>
iSCSI-B	<3020>

1. In the APIC GUI, select Fabric > Access Policies > Quick Start.

2. In the right pane, select Configure and interface, PC and VPC.

3. Select the paired Nexus 9332 switches configured in the last step under Configured Switch Interfaces.

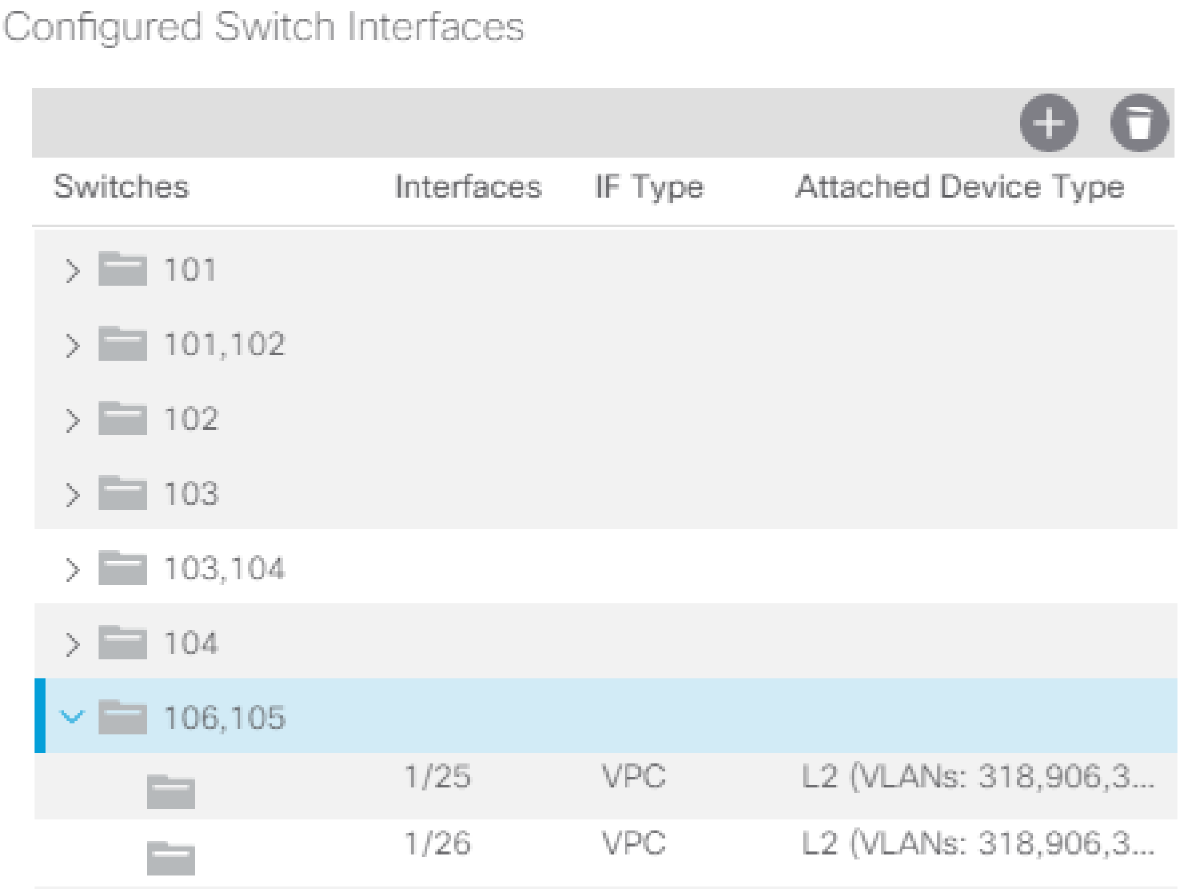

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
101			
101,102			
102			
103			
103,104			
104			
106,105			
	1/25	VPC	L2 (VLANs: 318,906,3...
	1/26	VPC	L2 (VLANs: 318,906,3...

4. Click ⊕ on the right to add switch interfaces.

5. Configure various fields as shown in the screenshot below. In this screen shot, port 1/1 on both leaf switches is connected to Storage Controller 1 using 40Gbps links.

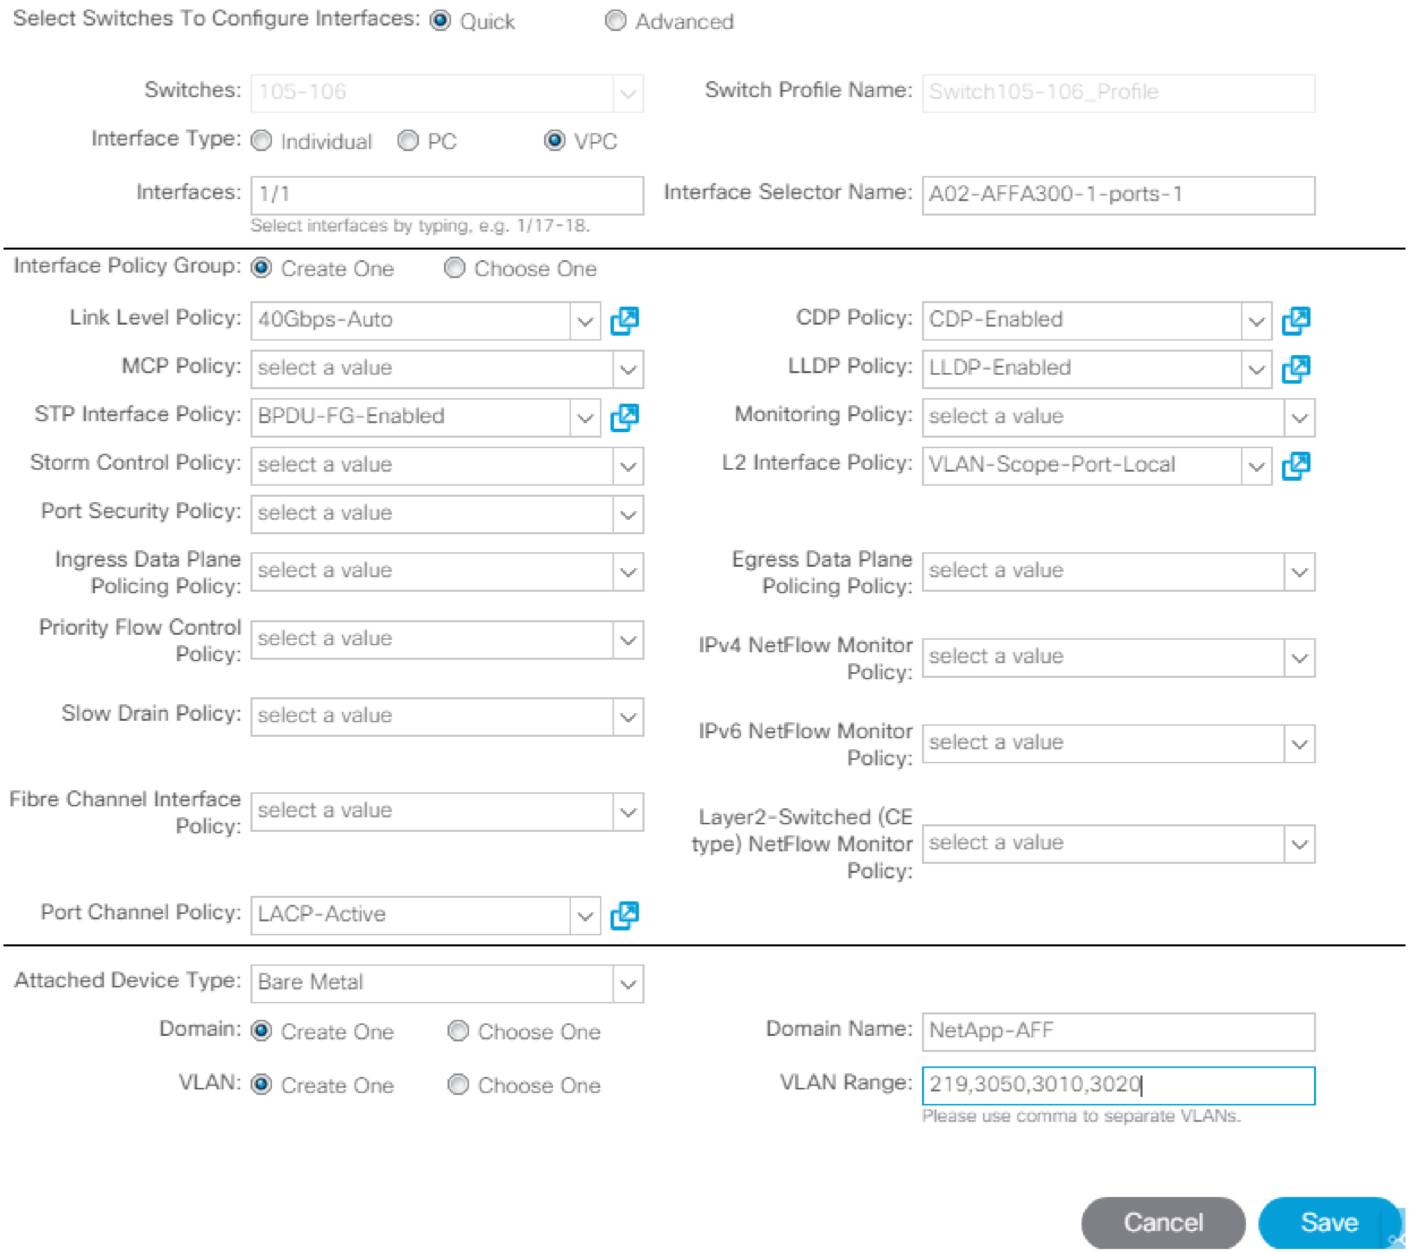

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 105-106

Switch Profile Name: Switch105-106_Profile

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/1

Select interfaces by typing, e.g. 1/17-18.

Interface Selector Name: A02-AFFA300-1-ports-1

Interface Policy Group: ● Create One ○ Choose One

Link Level Policy: 40Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: select a value

LLDP Policy: LLDP-Enabled

STP Interface Policy: BPDU-FG-Enabled

Monitoring Policy: select a value

Storm Control Policy: select a value

L2 Interface Policy: VLAN-Scope-Port-Local

Port Security Policy: select a value

Ingress Data Plane Policing Policy: select a value

Egress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

IPv4 NetFlow Monitor Policy: select a value

Slow Drain Policy: select a value

IPv6 NetFlow Monitor Policy: select a value

Fibre Channel Interface Policy: select a value

Layer2-Switched (CE type) NetFlow Monitor Policy: select a value

Port Channel Policy: LACP-Active

Attached Device Type: Bare Metal

Domain: ● Create One ○ Choose One

Domain Name: NetApp-AFF

VLAN: ● Create One ○ Choose One

VLAN Range: 219,3050,3010,3020

Please use comma to separate VLANs.

6. Click **Save.**

7. Click **Save** again to finish the configuring switch interfaces.

8. Click **Submit**.

9. From the right pane, select Configure and interface, PC and VPC.

10. Select the paired Nexus 9332 switches configured in the last step under Configured Switch Interfaces.

11. Click + to add switch interfaces.

12. Configure various fields as shown in the screenshot below. In this screenshot, port 1/2 on both leaf switches is connected to Storage Controller 2 using 40Gbps links. Instead of creating a new domain, the Bare Metal Device created in the previous step (NetApp-AFF) is attached to the storage controller 2 as shown below.

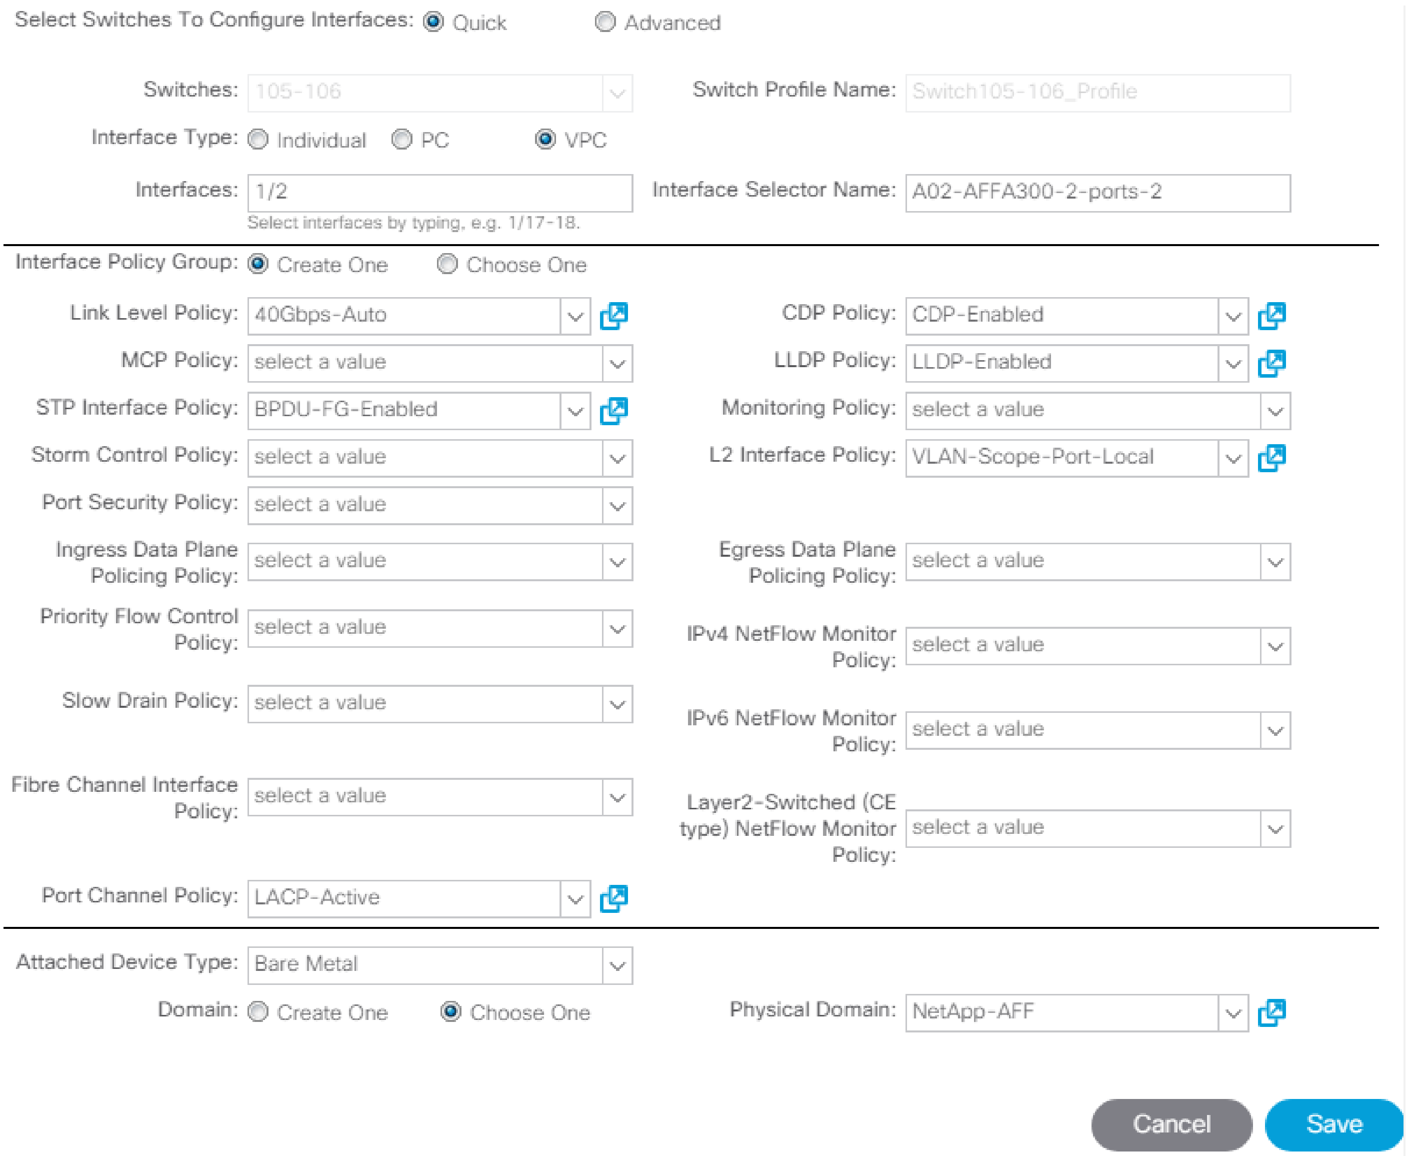

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 105-106

Switch Profile Name: Switch105-106_Profile

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/2

Select interfaces by typing, e.g. 1/17-18.

Interface Selector Name: A02-AFFA300-2-ports-2

Interface Policy Group: ● Create One ○ Choose One

Link Level Policy: 40Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: select a value

LLDP Policy: LLDP-Enabled

STP Interface Policy: BPDU-FG-Enabled

Monitoring Policy: select a value

Storm Control Policy: select a value

L2 Interface Policy: VLAN-Scope-Port-Local

Port Security Policy: select a value

Ingress Data Plane Policing Policy: select a value

Egress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

IPv4 NetFlow Monitor Policy: select a value

Slow Drain Policy: select a value

IPv6 NetFlow Monitor Policy: select a value

Fibre Channel Interface Policy: select a value

Layer2-Switched (CE type) NetFlow Monitor Policy: select a value

Port Channel Policy: LACP-Active

Attached Device Type: Bare Metal

Domain: ○ Create One ● Choose One

Physical Domain: NetApp-AFF

13. Click **Save.**

14. Click **Save** again to finish the configuring switch interfaces.

15. Click **Submit**.

16. **Optional:** Repeat this procedure to configure any additional NetApp AFF storage controllers. For a uniform configuration, the Bare Metal Domain (NetApp-AFF) will be utilized for all the Storage Controllers.

Configuring Common Tenant for In-Band Management Access

This section details the steps to setup in-band management access in the Tenant common. This design will allow all the other tenant EPGs to access the common management segment for Core Services VMs such as AD/DNS.

1. In the APIC GUI, select Tenants > common.
2. In the left pane, expand Tenant common and Networking.

Create VRFs

To create VRFs, complete the following steps:

1. Right-click VRFs and select **Create VRF**.
2. Enter vrf-FP-Common-IB-MGMT as the name of the VRF.
3. Uncheck Create A Bridge Domain.
4. Click Finish.

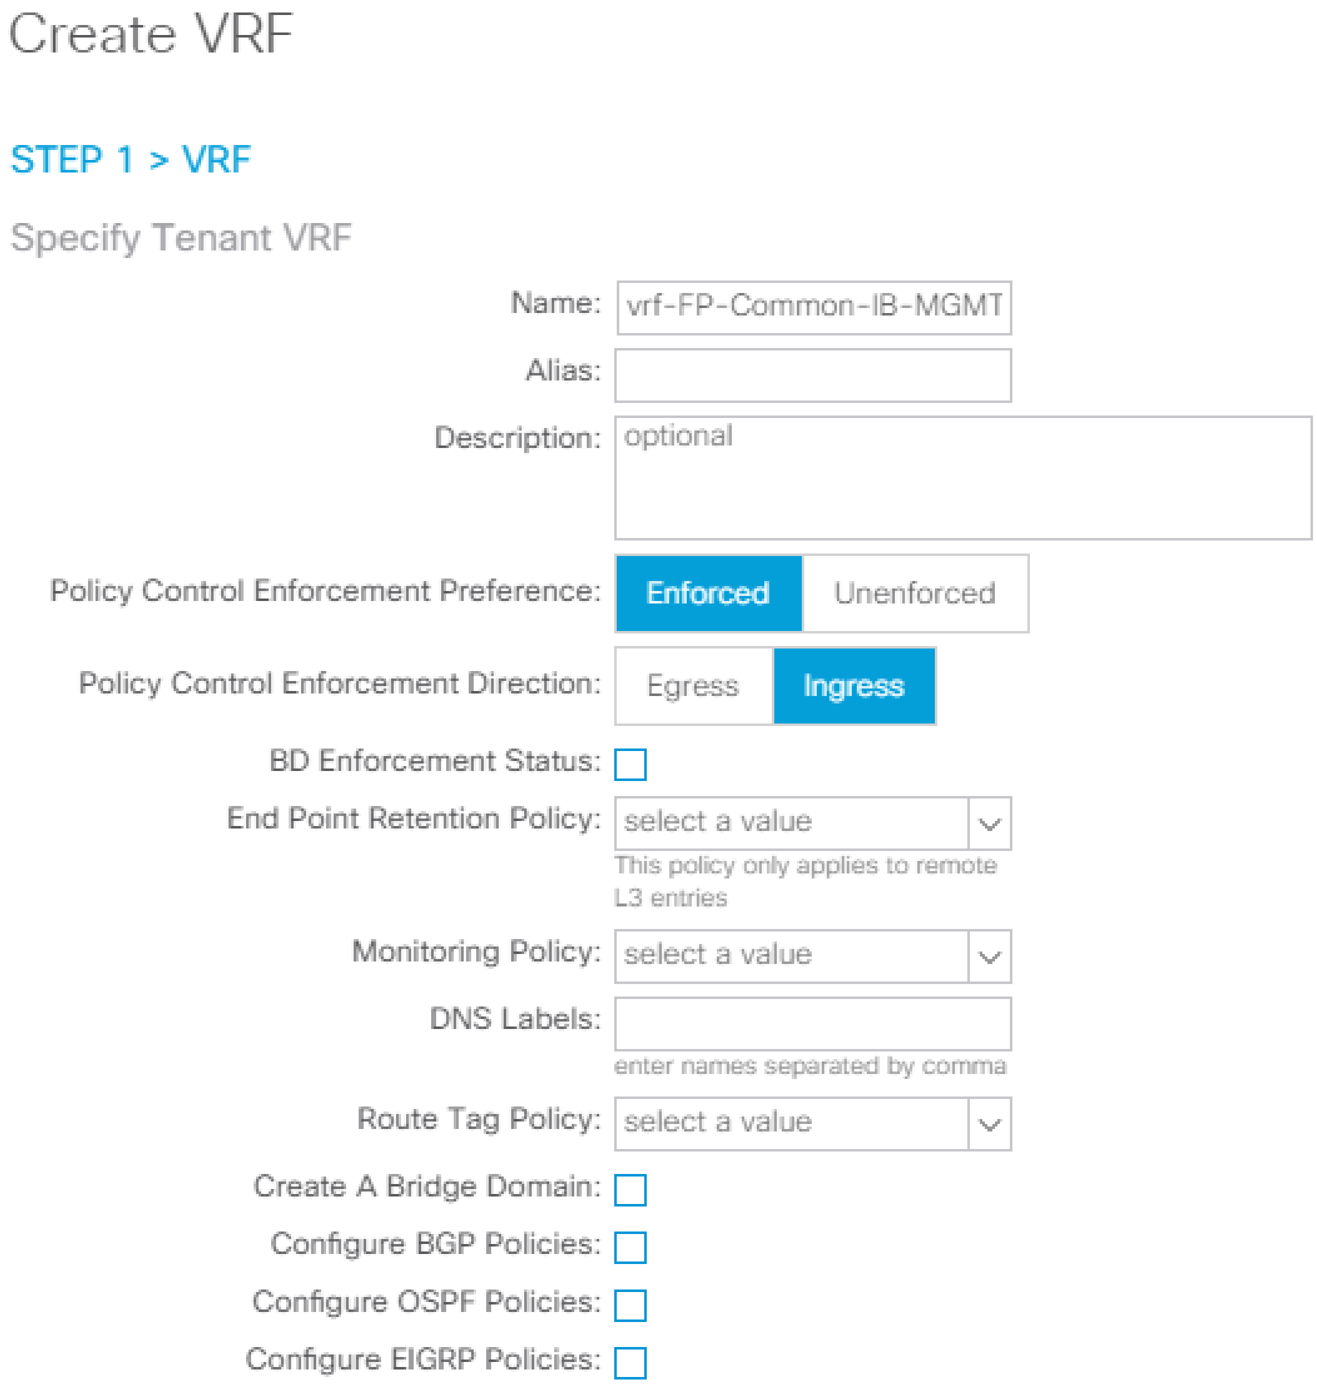

Create VRF

STEP 1 > VRF

Specify Tenant VRF

Name: vrf-FP-Common-IB-MGMT

Alias:

Description: optional

Policy Control Enforcement Preference: [✓ Enforced] Unenforced

Policy Control Enforcement Direction: Egress [✓ Ingress]

BD Enforcement Status: □

End Point Retention Policy: select a value

This policy only applies to remote L3 entries

Monitoring Policy: select a value

DNS Labels:

enter names separated by comma

Route Tag Policy: select a value

Create A Bridge Domain: □

Configure BGP Policies: □

Configure OSPF Policies: □

Configure EIGRP Policies: □

Create Bridge Domain

To create the incoming IB-MGMT Bridge domain for Core-Services, complete the following steps:

1. In the APIC GUI, select Tenants > common.
2. In the left pane, expand Tenant common and Networking.
3. Right-click Bridge Domains and select **Create Bridge Domain**.

4. Name the Bridge Domain as BD-FP-common-Core-Services.

5. Select vrf-FP-Common-IB-MGMT from the VRF drop-down list.

6. Select Custom under Forwarding and enable the flooding as shown in the screenshot below.

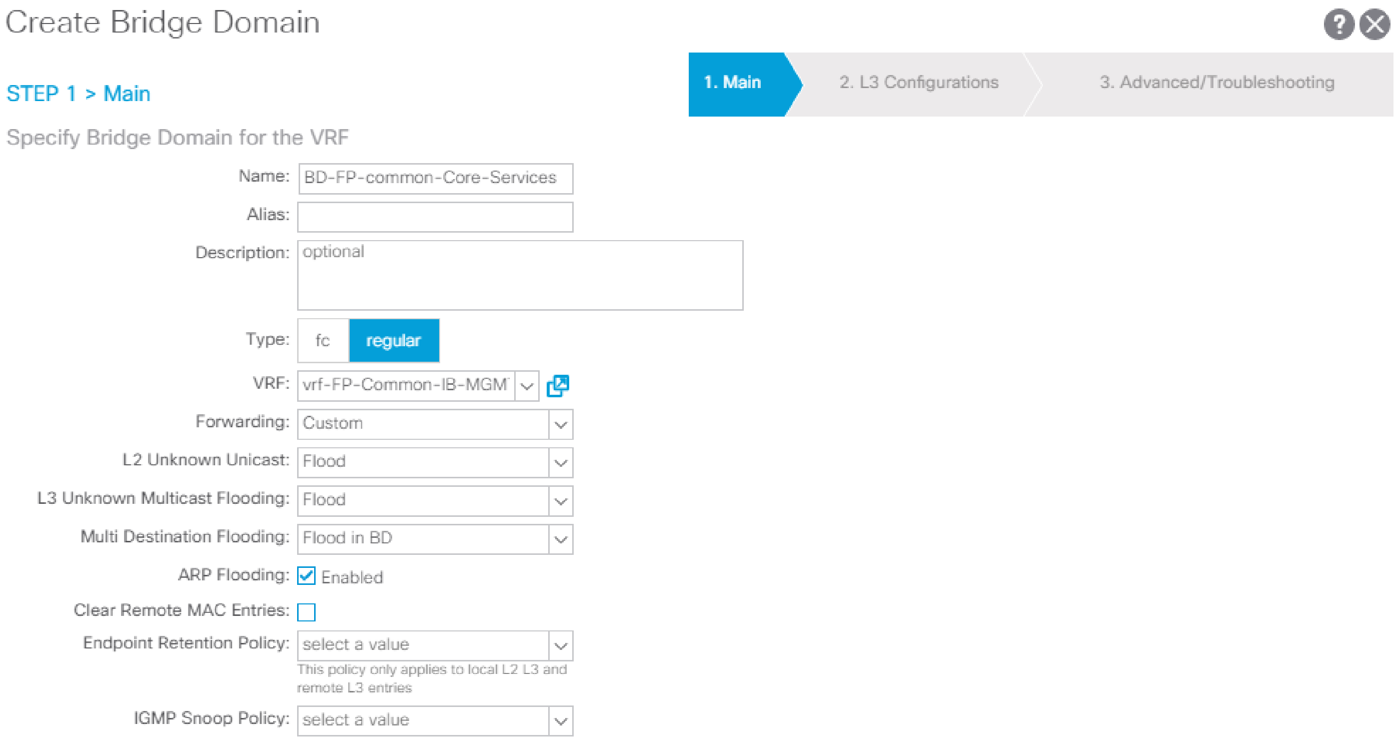

Create Bridge Domain

STEP 1 > Main

Specify Bridge Domain for the VRF

Name:	BD-FP-common-Core-Services
Alias:	
Description:	optional
Type:	fc [✓ regular]
VRF:	vrf-FP-Common-IB-MGM'
Forwarding:	Custom
L2 Unknown Unicast:	Flood
L3 Unknown Multicast Flooding:	Flood
Multi Destination Flooding:	Flood in BD
ARP Flooding:	☑ Enabled
Clear Remote MAC Entries:	□
Endpoint Retention Policy:	select a value This policy only applies to local L2 L3 and remote L3 entries
IGMP Snoop Policy:	select a value

7. Click **Next**.

8. Under L3 Configurations, make sure Limit IP Learning to Subnet is selected and select EP Move Detection Mode – GARP based detection. Select **Next**.

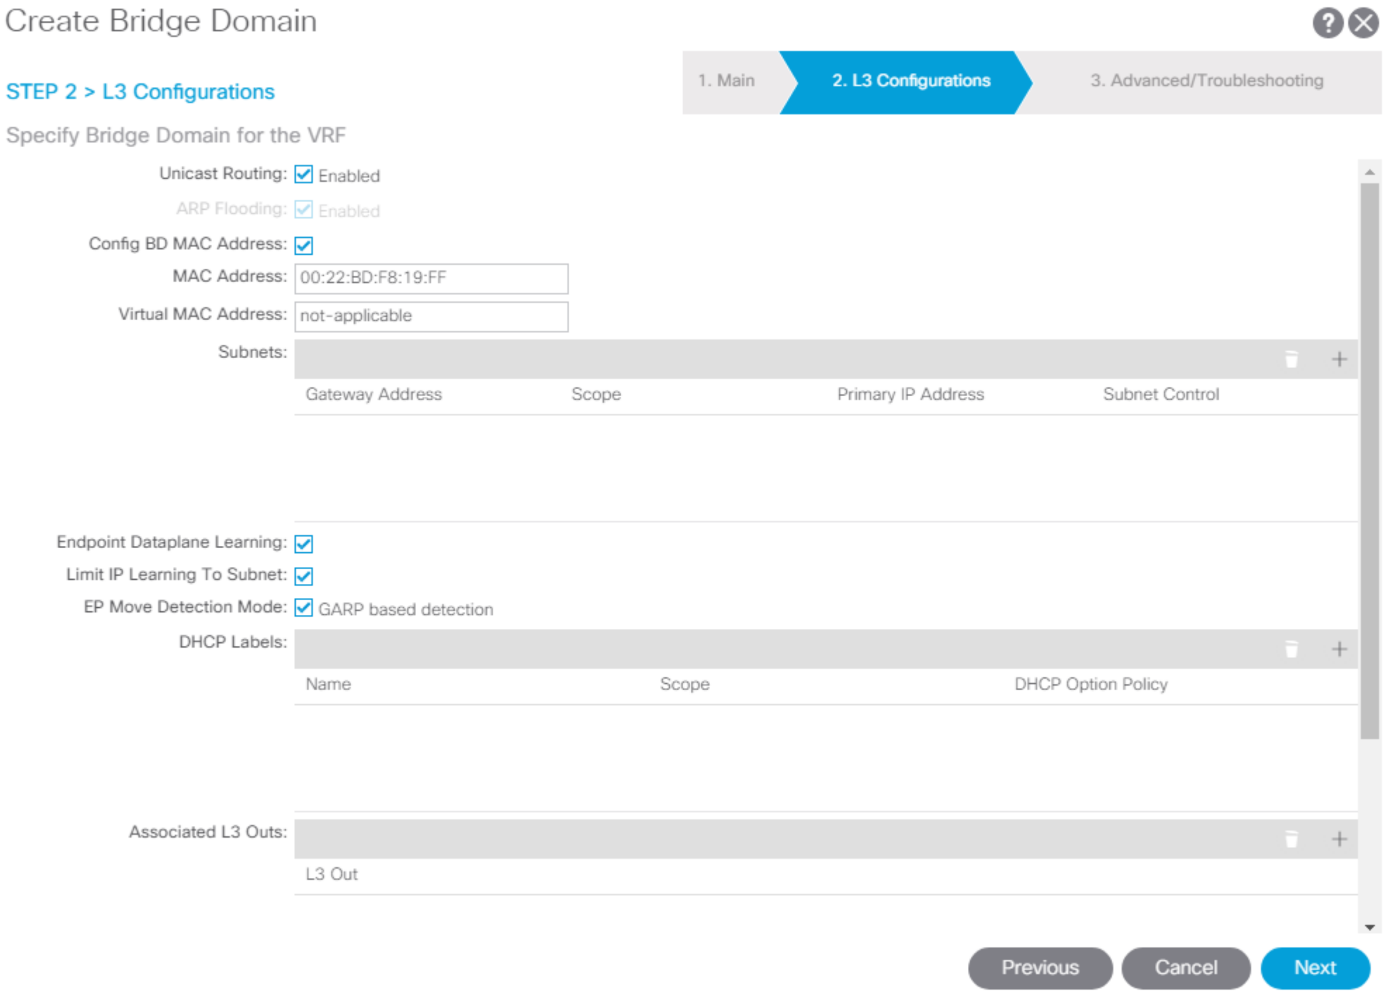

Create Bridge Domain

STEP 2 > L3 Configurations

Specify Bridge Domain for the VRF

Unicast Routing: ☑ Enabled

ARP Flooding: ☑ Enabled

Config BD MAC Address: ☑

MAC Address: 00:22:BD:F8:19:FF

Virtual MAC Address: not-applicable

Subnets:

Gateway Address	Scope	Primary IP Address	Subnet Control

Endpoint Dataplane Learning: ☑

Limit IP Learning To Subnet: ☑

EP Move Detection Mode: ☑ GARP based detection

DHCP Labels:

Name	Scope	DHCP Option Policy

Associated L3 Outs:

L3 Out

9. No changes are needed Advanced/Troubleshooting. Click **FINISH**.

Create Application Profile

To create an Application profile, complete the following steps:

1. In the APIC GUI, select **Tenants** > **common**.
2. In the left pane, expand Tenant common and Application Profiles.
3. Right-click the Application Profiles and select **Create Application Profiles**.
4. Enter Core-Services as the name of the application profile.

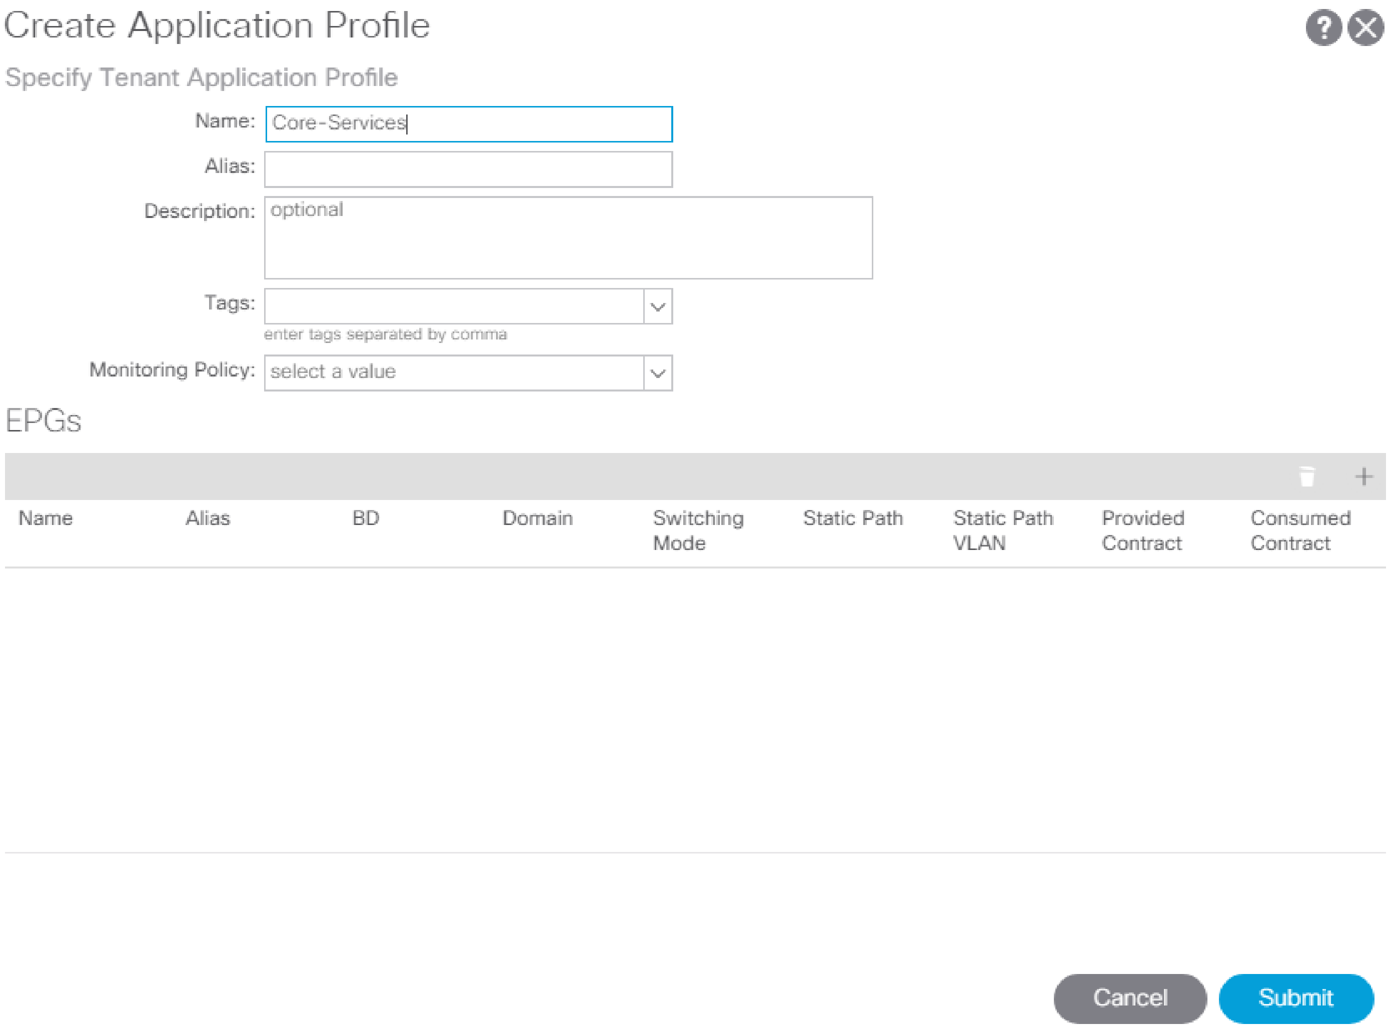

Create Application Profile

Specify Tenant Application Profile

Name: Core-Services

Alias:

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: select a value

EPGs

Name	Alias	BD	Domain	Switching Mode	Static Path	Static Path VLAN	Provided Contract	Consumed Contract

5. Click **Submit**.

Create EPG

To create the Core-Services EPG, complete the following steps:

1. Expand the Core-Services Application Profile and right-click Application EPGs.

2. Select Create Application EPG.

3. Enter Core-Services as the name of the EPG.

4. Select BD-FP-common-Core-Services from the drop-down list for Bridge Domain.

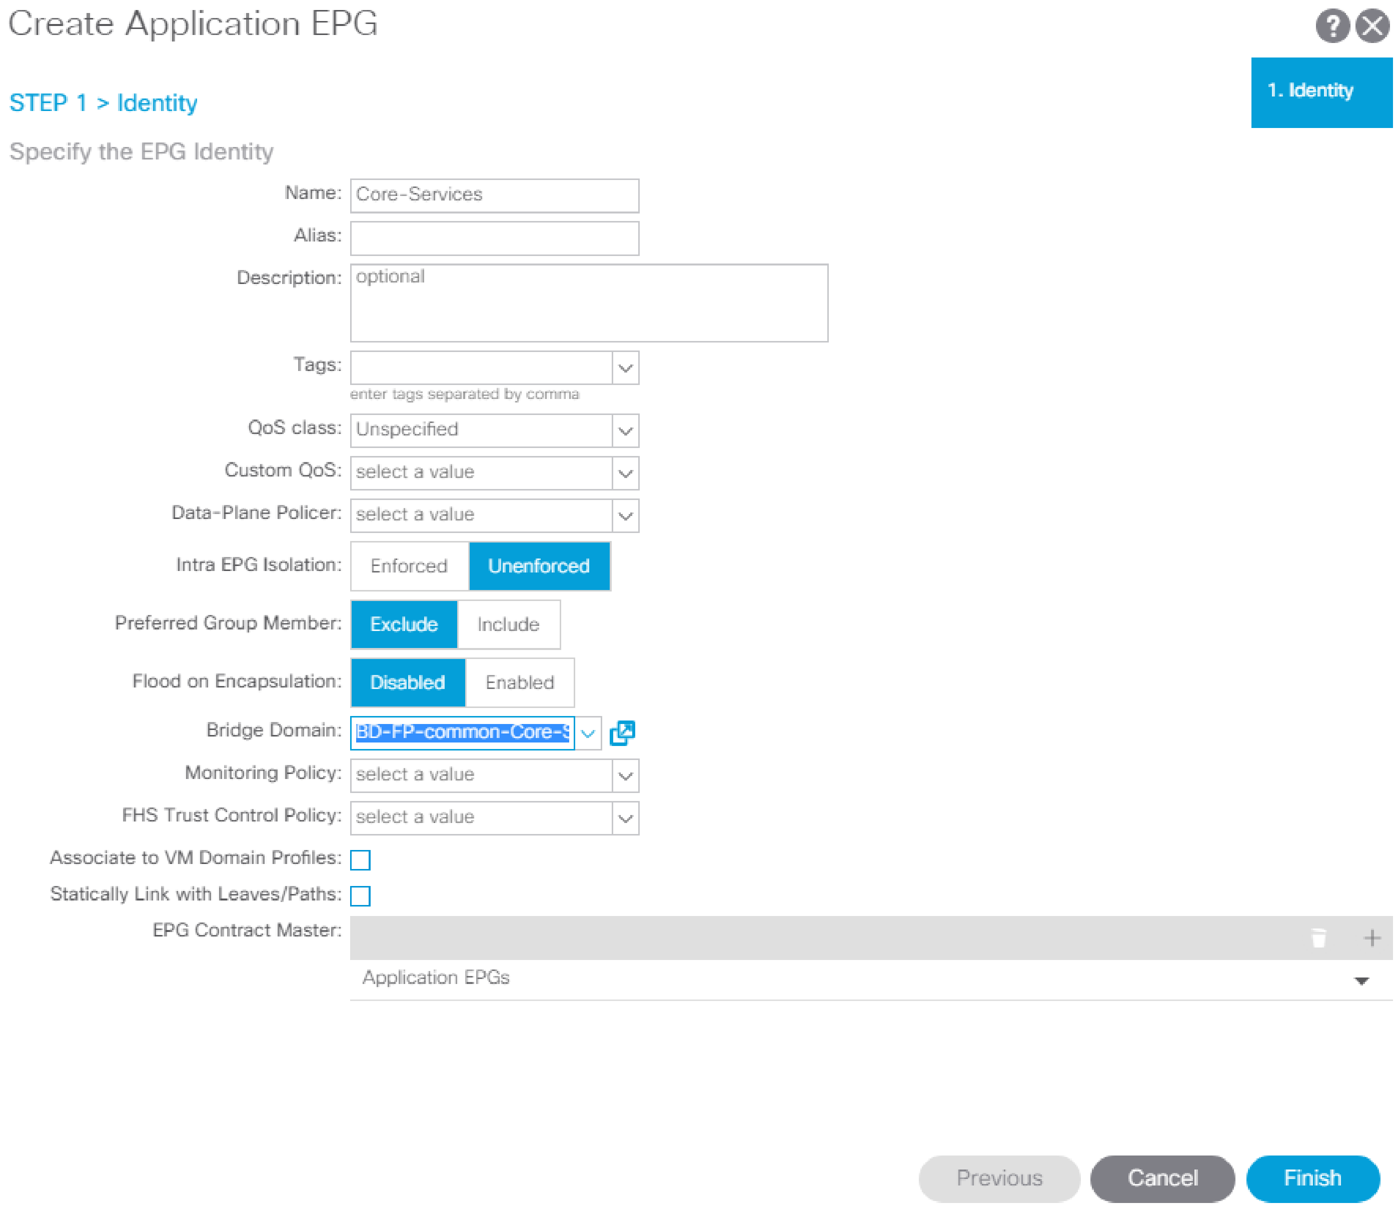

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name:	Core-Services
Alias:	
Description:	optional
Tags:	enter tags separated by comma
QoS class:	Unspecified
Custom QoS:	select a value
Data-Plane Policer:	select a value
Intra EPG Isolation:	Enforced [✓ Unenforced]
Preferred Group Member:	[✓ Exclude] Include
Flood on Encapsulation:	[✓ Disabled] Enabled
Bridge Domain:	BD-FP-common-Core-S
Monitoring Policy:	select a value
FHS Trust Control Policy:	select a value
Associate to VM Domain Profiles:	□
Statically Link with Leaves/Paths:	□
EPG Contract Master:	Application EPGs

5. Click **Finish**.

Set Domains for the EPG

To set Domains, complete the following steps:

1. Expand the newly create EPG and click **Domains**.

2. Right-click Domains and select Add L2 External Domain Association.

3. Select the FP-Mgmt-Sw as the L2 External Domain Profile.

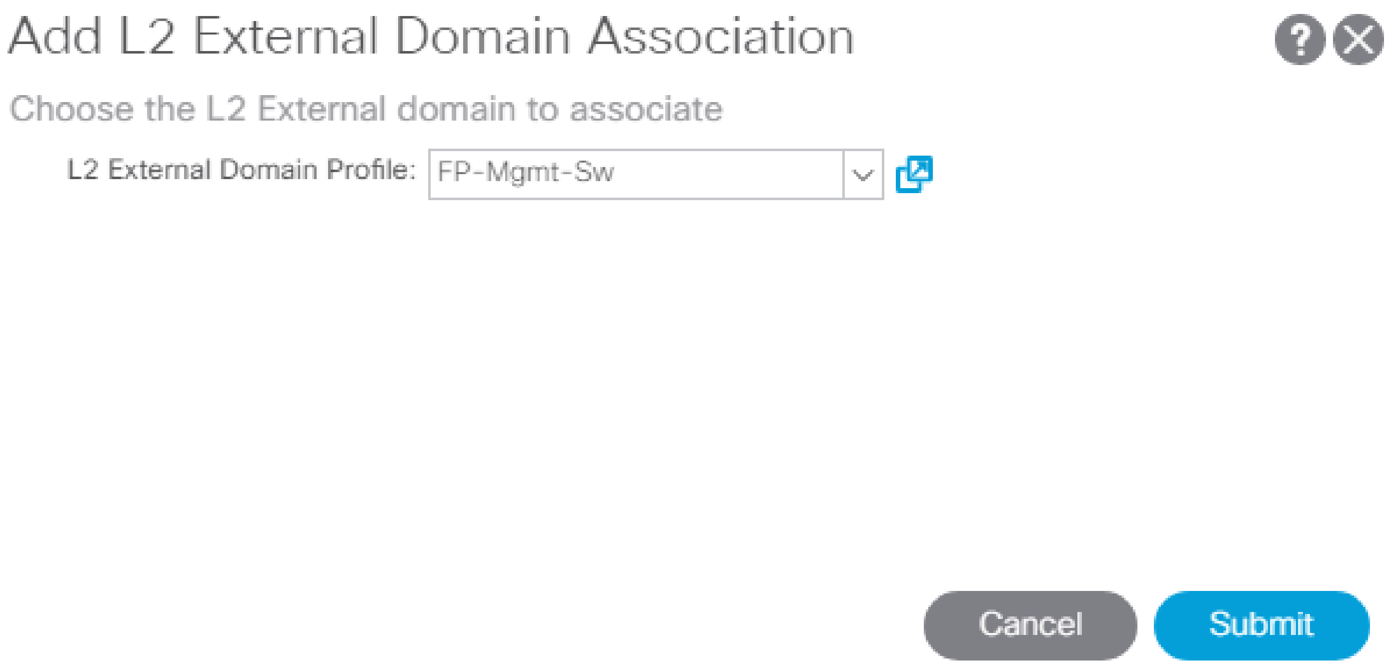

Add L2 External Domain Association

Choose the L2 External domain to associate

L2 External Domain Profile: FP-Mgmt-Sw

4. Click **Submit**.

5. Right-click Domains and select Add L2 External Domain Association.

6. Select the UCS as the L2 External Domain Profile.

7. Click **Submit**.

Set Static Ports for the EPG

To set Static Ports, complete the following steps:

1. In the left pane, right-click Static Ports.

2. Select Deploy Static EPG on PC, VPC, or Interface.

3. In the next screen, for the Path Type, select Virtual Port Channel and from the Path drop-down list, select the VPC for FP-Mgmt-Sw configured earlier.

4. Enter the external IB-MGMT VLAN under Port Encap.

5. Change Deployment Immediacy to Immediate.

6. Set the Mode to Trunk.

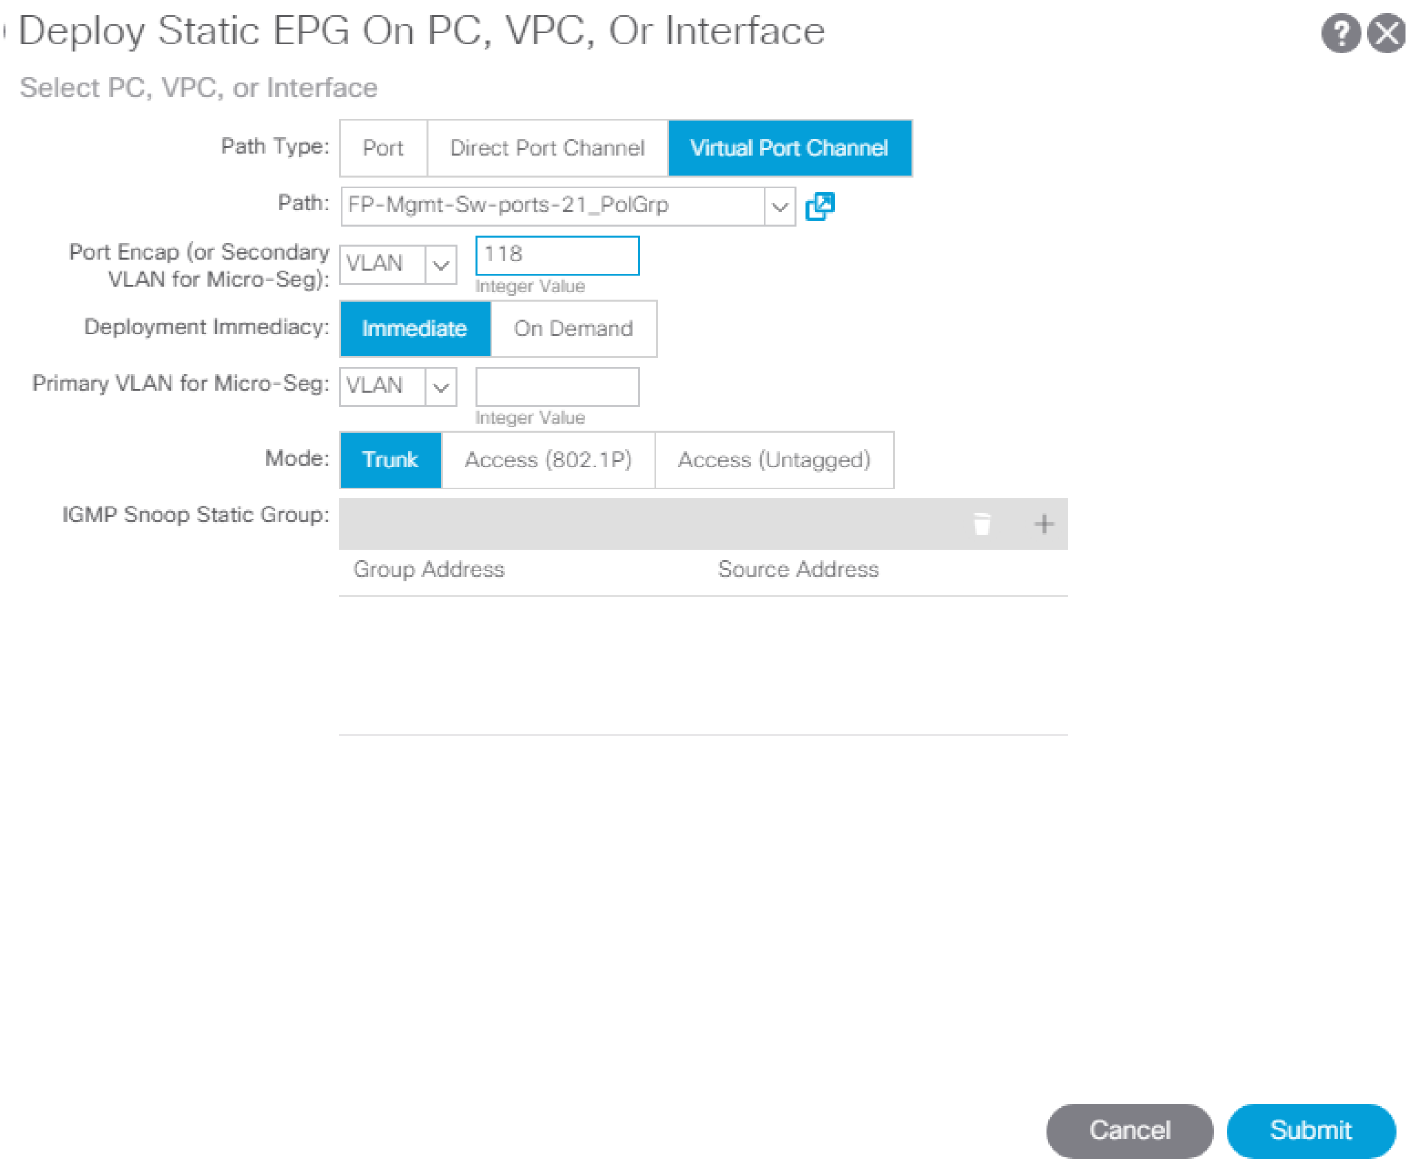

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Virtual Port Channel

Path: FP-Mgmt-Sw-ports-21_PolGrp

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 118

Deployment Immediacy: Immediate

Primary VLAN for Micro-Seg: VLAN

Mode: Trunk

IGMP Snoop Static Group:

7. Click **Submit**.

8. In the left pane, right-click Static Ports.

9. Select Deploy Static EPG on PC, VPC, or Interface.

10. In the next screen, for the Path Type, select Virtual Port Channel and from the Path drop-down list, select the VPC for UCS Fabric Interconnect A configured earlier.

11. Enter the UCS Core-Services VLAN under Port Encap.

This VLAN should be a different VLAN than the one entered above for the Management Switch.

12. Change Deployment Immediacy to Immediate.

13. Set the Mode to Trunk.

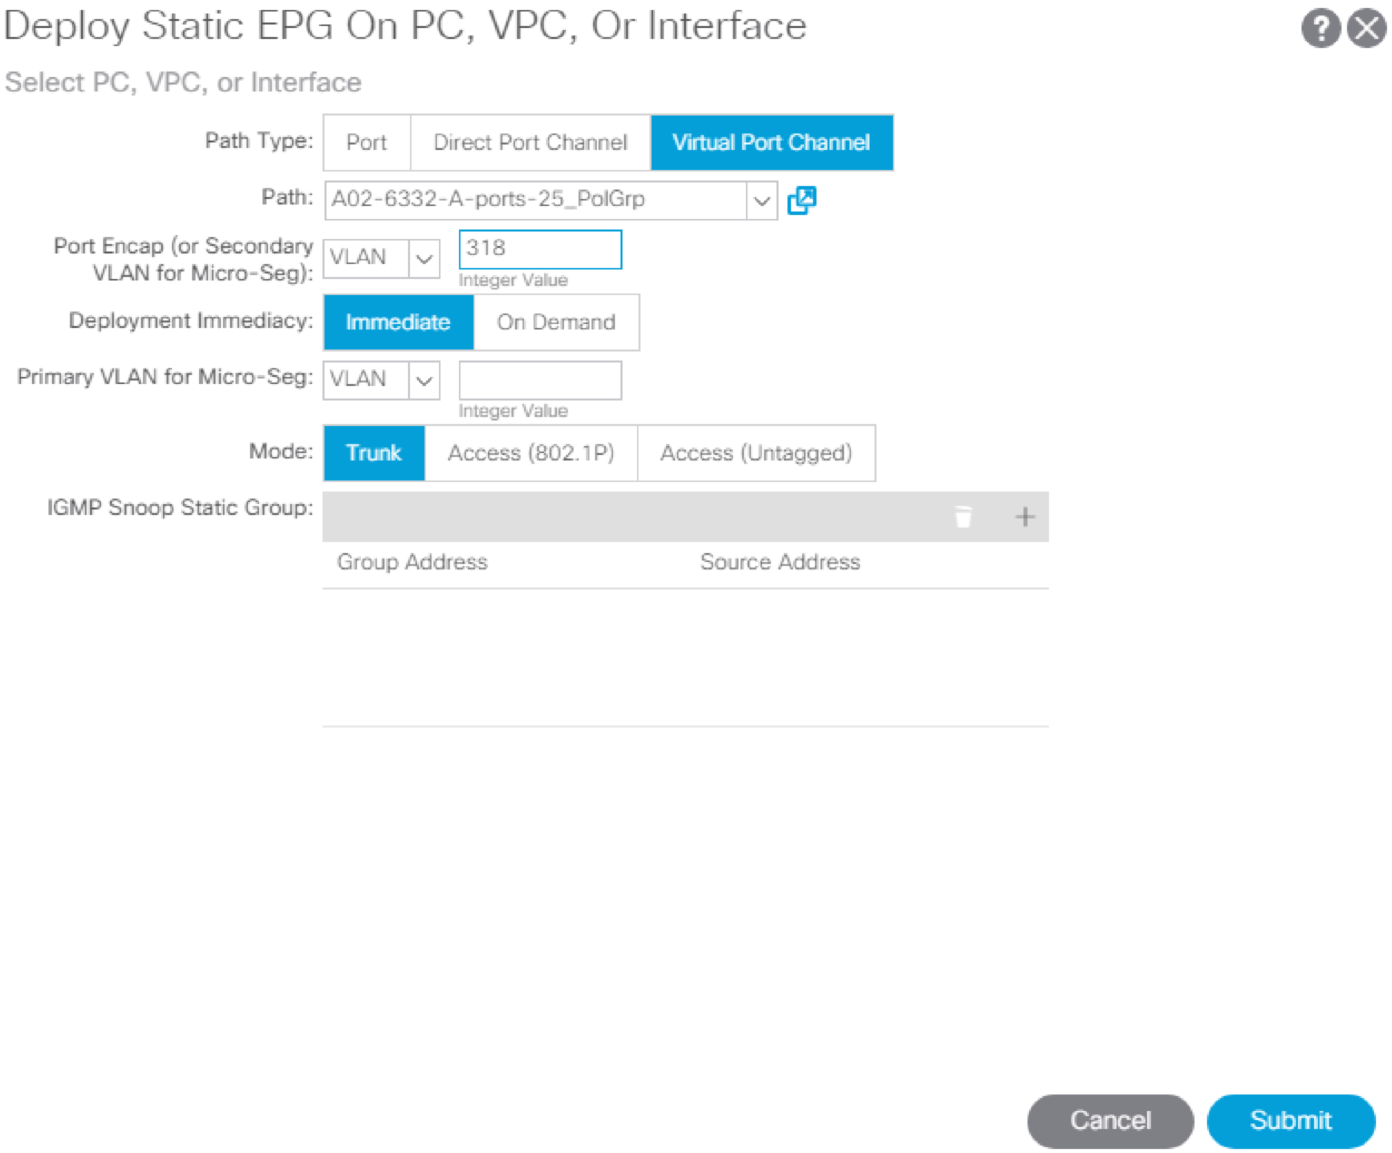

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port | Direct Port Channel | **Virtual Port Channel**

Path: A02-6332-A-ports-25_PolGrp

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 318 Integer Value

Deployment Immediacy: **Immediate** | On Demand

Primary VLAN for Micro-Seg: VLAN Integer Value

Mode: **Trunk** | Access (802.1P) | Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

14. Click Submit.

15. In the left pane, right-click Static Ports.

16. Select Deploy Static EPG on PC, VPC, or Interface.

17. In the next screen, for the Path Type, select Virtual Port Channel and from the Path drop-down list, select the VPC for UCS Fabric Interconnect B configured earlier.

18. Enter the UCS Core-Services VLAN under Port Encap.

This VLAN should be a different VLAN than the one entered above for the Management Switch.

19. Change Deployment Immediacy to Immediate.

20. Set the Mode to Trunk.

21. Click Submit.

Create EPG Subnet

A subnet gateway for this Core Services EPG provides Layer 3 connectivity to Tenant subnets. To create a EPG Subnet, complete the following steps:

1. In the left pane, right-click **Subnets** and select **Create EPG Subnet**.

2. In CIDR notation, enter an IP address and subnet mask to serve as the gateway within the ACI fabric for routing between the Core Services subnet and Tenant subnets. This IP should be different than the IB-MGMT subnet gateway. In this lab validation, 10.1.118.1/24 is the IB-MGMT subnet gateway and is configured externally to the ACI fabric. 10.1.118.254/24 will be used for the EPG subnet gateway. Set the Scope of the subnet to Shared between VRFs.

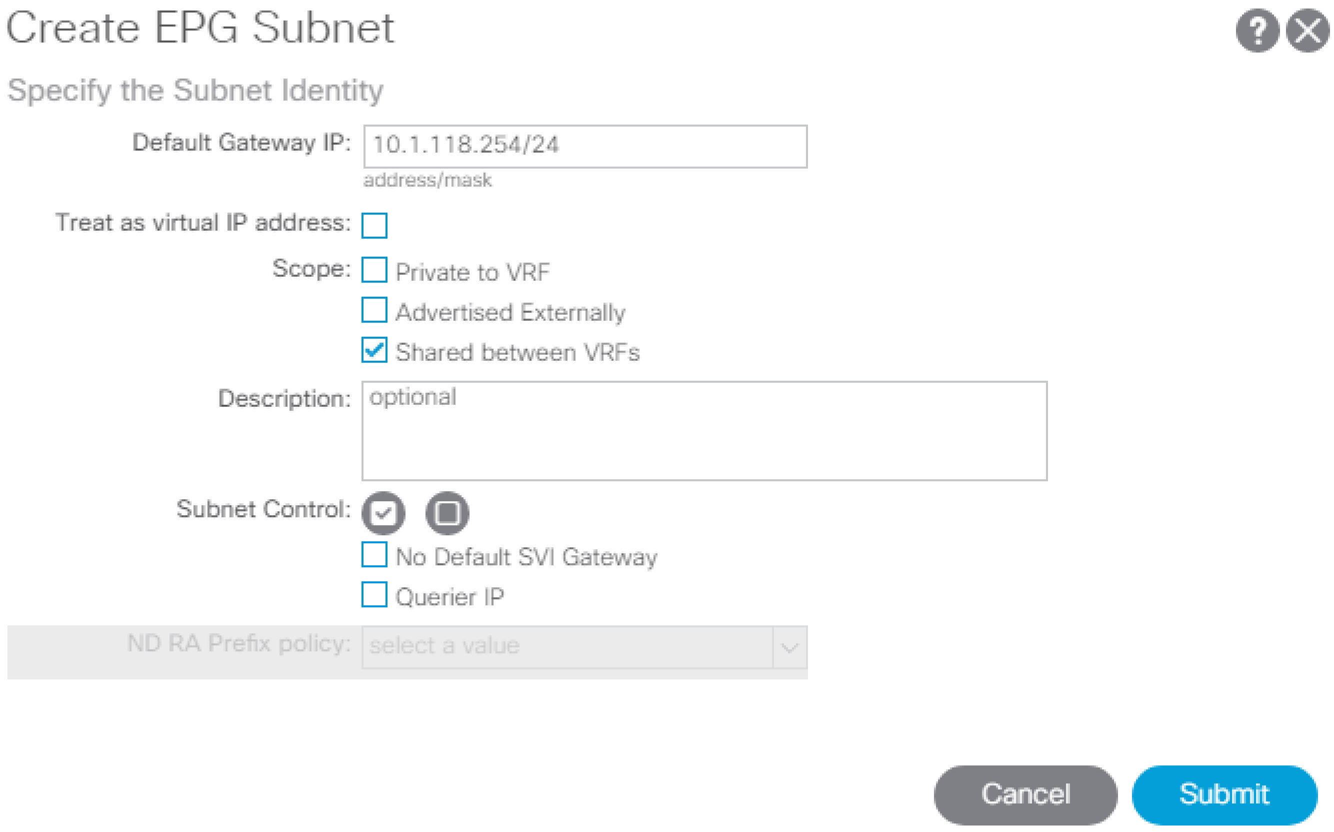

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP: 10.1.118.254/24
address/mask

Treat as virtual IP address: □

Scope: □ Private to VRF
□ Advertised Externally
☑ Shared between VRFs

Description: optional

Subnet Control:
□ No Default SVI Gateway
□ Querier IP

ND RA Prefix policy: select a value

3. Click **Submit** to create the Subnet.

Create Provided Contract

To create Provided Contract, complete the following steps:

1. In the left pane, right-click **Contracts** and select **Add Provided Contract**.

2. In the Add Provided Contract window, select **Create Contract** from the drop-down list.

3. Name the Contract FP-Allow-Common-Core-Services.

4. Set the scope to Global.

5. Click **+** to add a Subject to the Contract.

The following steps create a contract to allow all the traffic between various tenants and the common management segment. You are encouraged to limit the traffic by setting restrictive filters.

6. Name the subject Allow-All-Traffic.

7. Click **+** under Filter Chain to add a Filter.

8. From the drop-down Name list, select common/default.

9. In the Create Contract Subject window, click **Update** to add the Filter Chain to the Contract Subject.

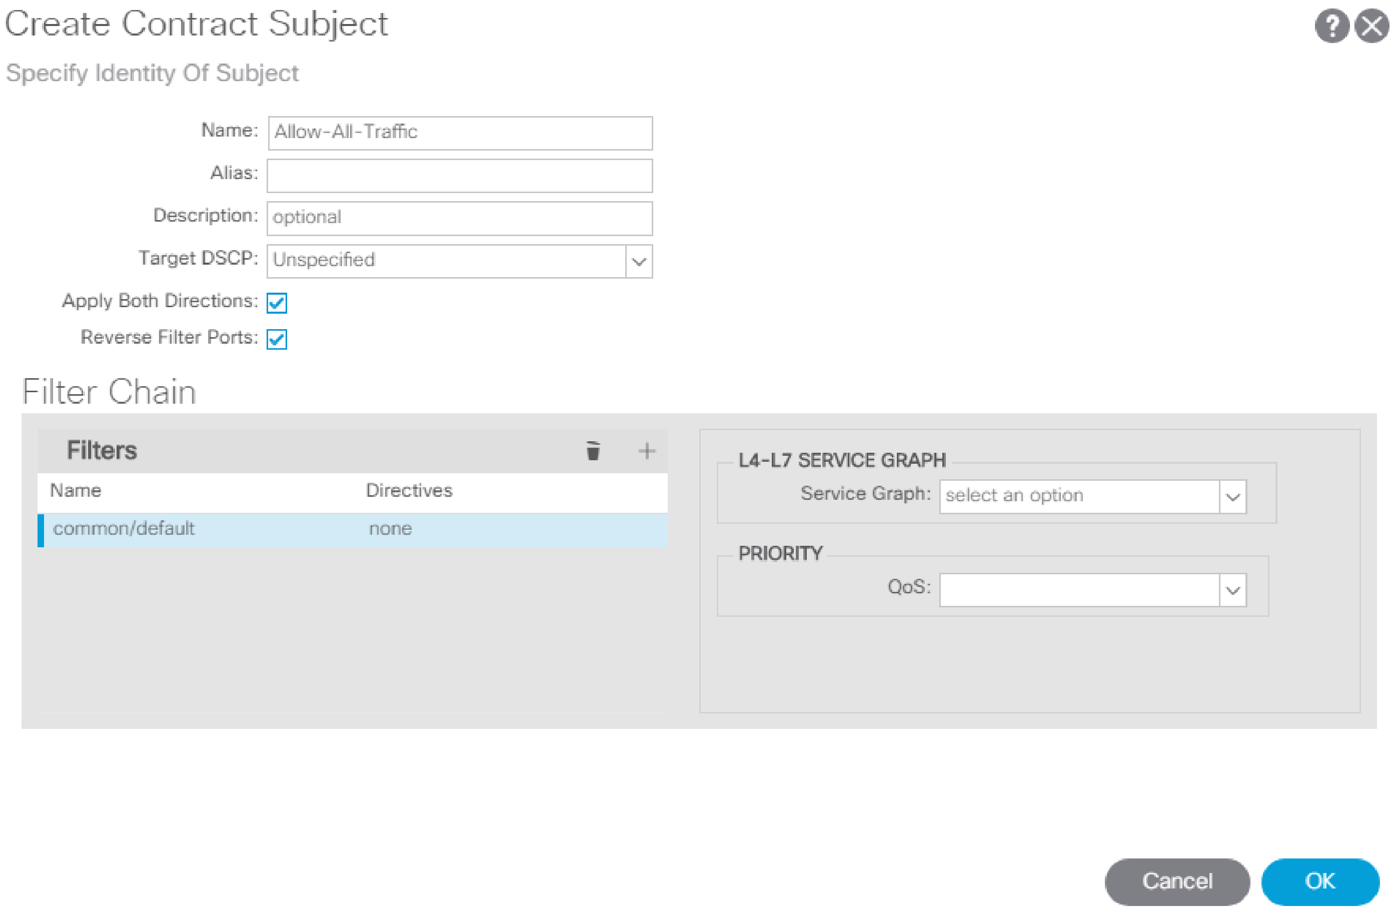

Create Contract Subject

Specify Identity Of Subject

Name: Allow-All-Traffic

Alias:

Description: optional

Target DSCP: Unspecified

Apply Both Directions: ☑

Reverse Filter Ports: ☑

Filter Chain

Filters

Name	Directives
common/default	none

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

10. Click **OK** to add the Contract Subject.

The Contract Subject Filter Chain can be modified later.

11. Click **Submit** to finish creating the Contract.

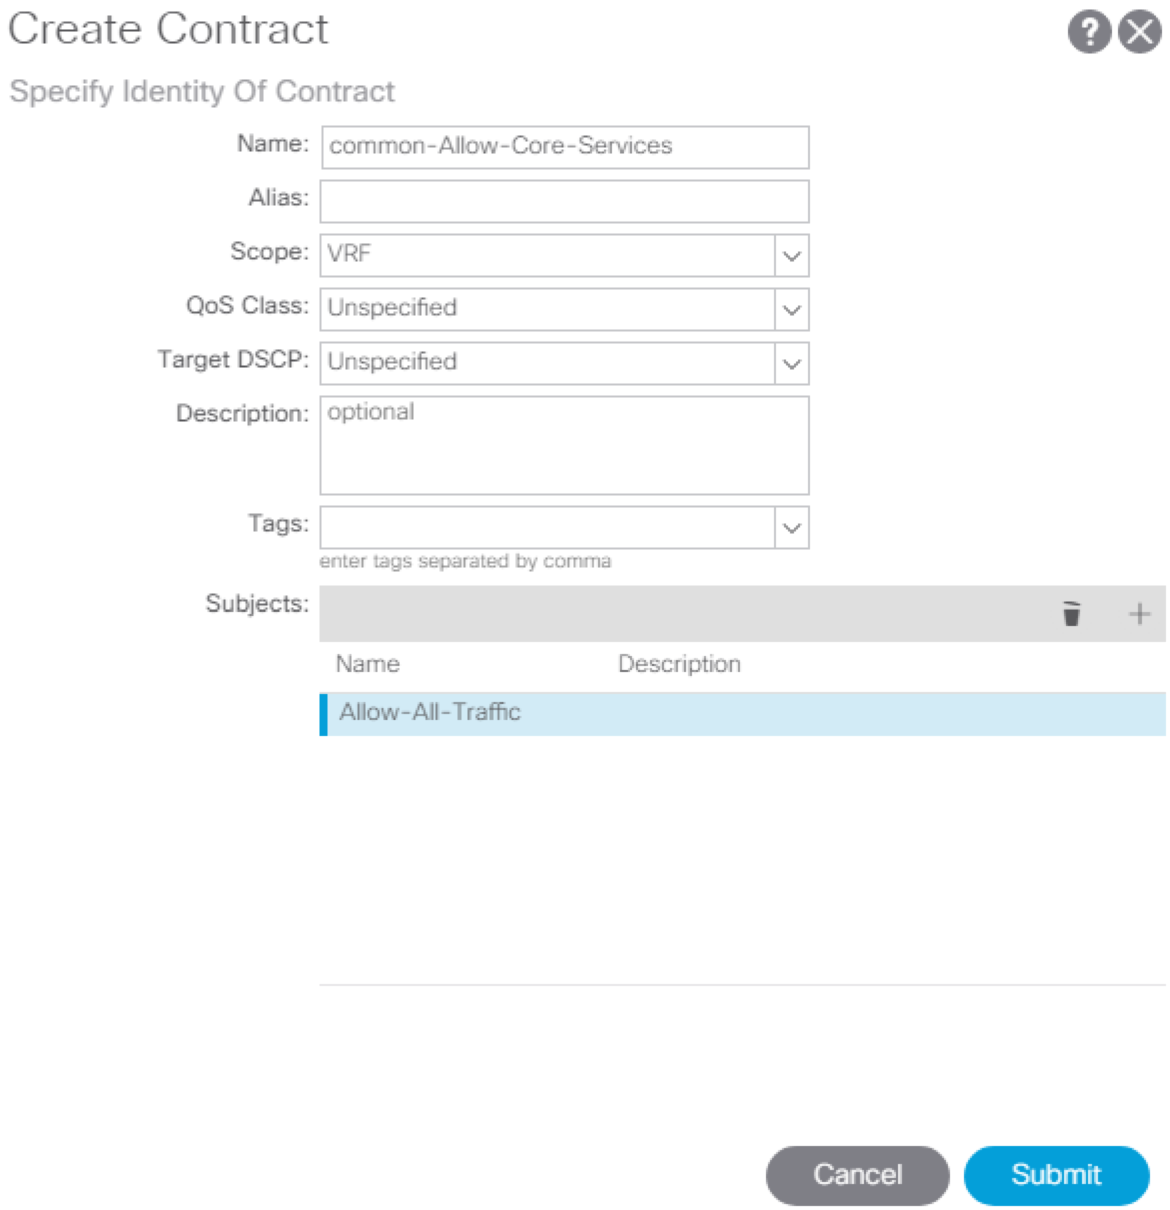

Create Contract

Specify Identity Of Contract

Name:	common-Allow-Core-Services
Alias:	
Scope:	VRF
QoS Class:	Unspecified
Target DSCP:	Unspecified
Description:	optional
Tags:	

enter tags separated by comma

Subjects:

Name	Description
Allow-All-Traffic	

12. Click **Submit** to finish adding a Provided Contract.

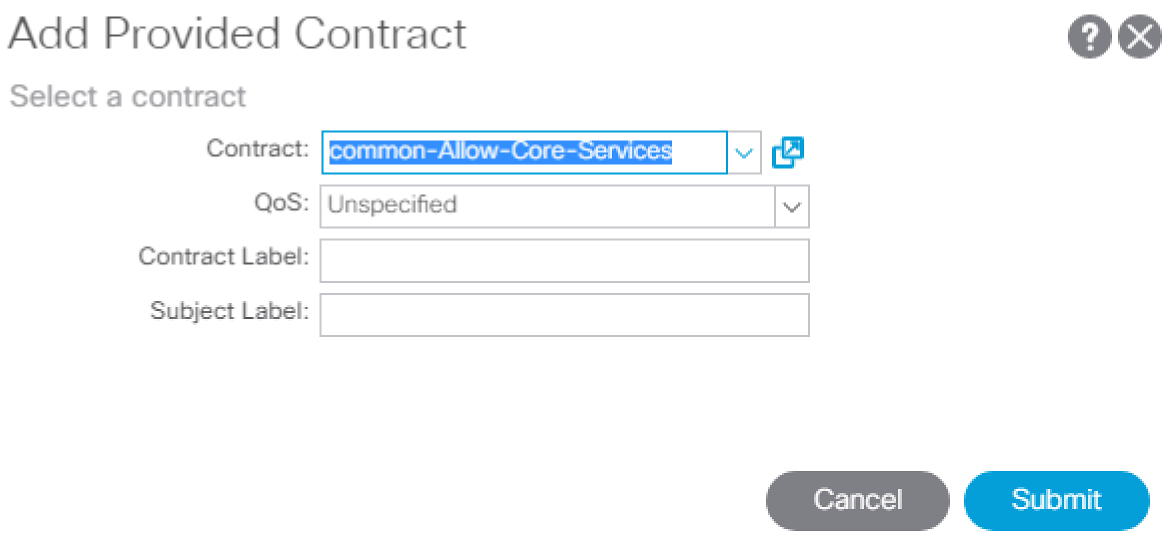

Add Provided Contract

Select a contract

Contract: common-Allow-Core-Services

QoS: Unspecified

Contract Label:

Subject Label:

Create Security Filters in Tenant Common

To create Security Filters for NFSv3 with NetApp Storage and for iSCSI, complete the following steps. This section can also be used to set up other filters necessary to your environment.

1. In the APIC GUI, at the top select Tenants > common.
2. On the left, expand Tenant common, Contracts, and Filters.
3. Right-click Filters and select Create Filter.
4. Name the filter Allow-All.
5. Click the + sign to add an Entry to the Filter.
6. Name the Entry Allow-All and select EtherType IP.
7. Leave the IP Protocol set at Unspecified.
8. Click UPDATE to add the Entry.

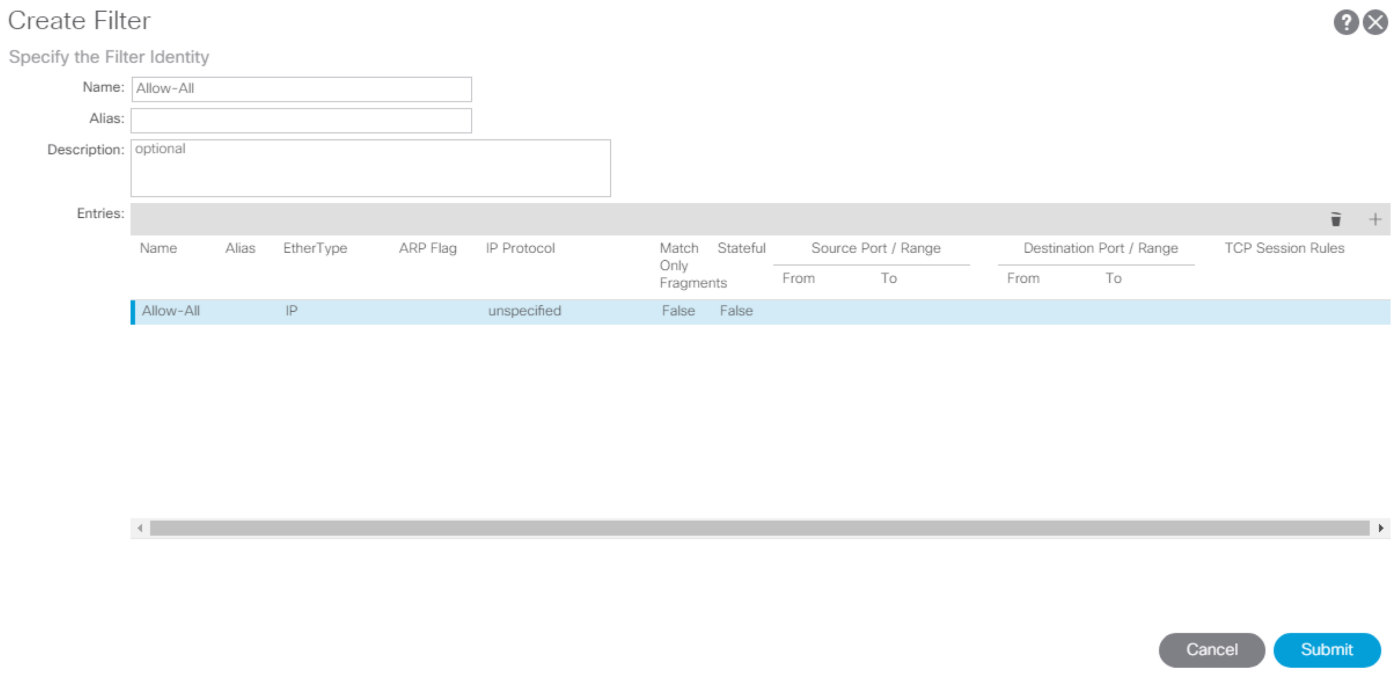

Create Filter

Specify the Filter Identity

Name: Allow-All

Alias:

Description: optional

Entries:

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
							From	To	From	To	
Allow-All		IP		unspecified	False	False					

9. Click SUBMIT to complete adding the Filter.

10. Right-click Filters and select Create Filter.

11. Name the filter NTAP-NVS-v3.

12. Click the + sign to add an Entry to the Filter.

13. Name the Entry tcp-111 and select EtherType IP.

14. 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.

15. Click UPDATE to add the Entry.

16. Click the + sign to add another Entry to the Filter.

17. Name the Entry tcp-635 and select EtherType IP.

18. 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.

19. Click UPDATE to add the Entry.

20. Click the + sign to add another Entry to the Filter.

21. Name the Entry tcp-2049 and select EtherType IP.

22. 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.

23. Click UPDATE to add the Entry.

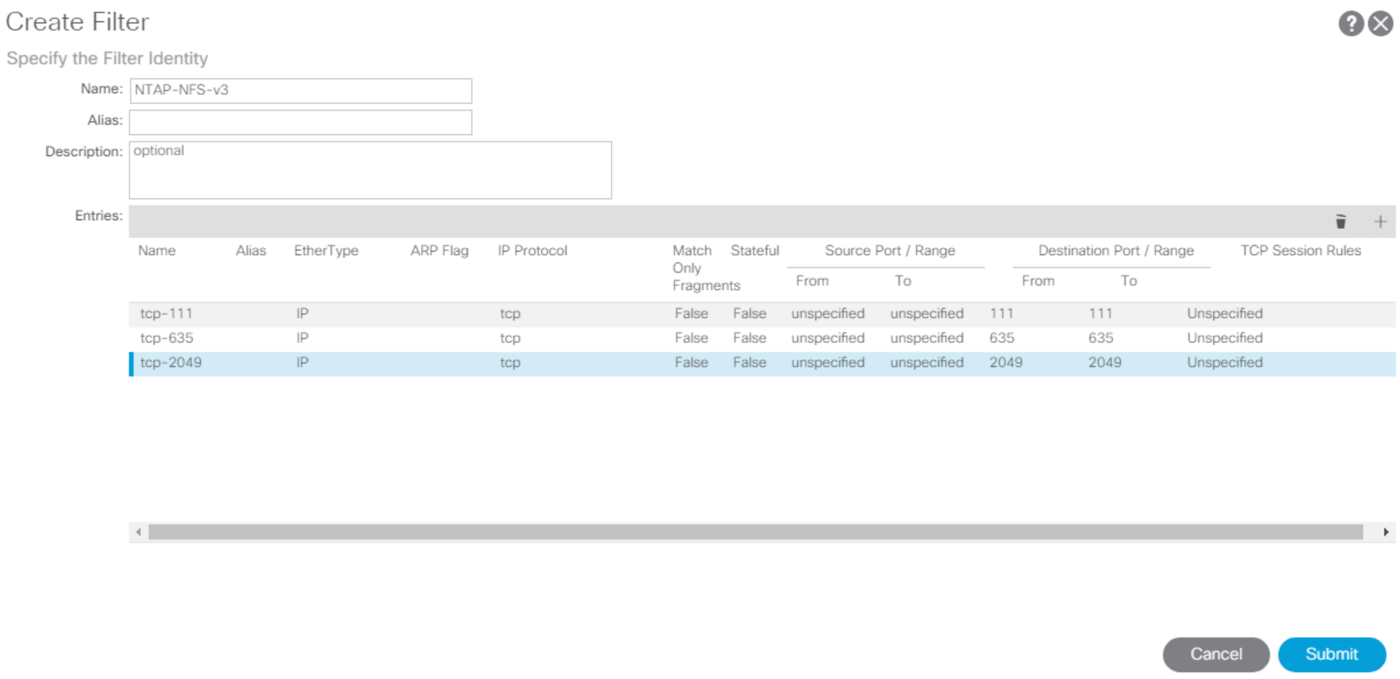

Create Filter

Specify the Filter Identity

Name: NTAP-NFS-v3

Alias:

Description: optional

Entries:

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	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

24. Click SUBMIT to complete adding the Filter.

25. Right-click Filters and select Create Filter.

26. Name the filter iSCSI.

27. Click the + sign to add an Entry to the Filter.

28. Name the Entry iSCSI and select EtherType IP.

29. 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.

30. Click UPDATE to add the Entry.

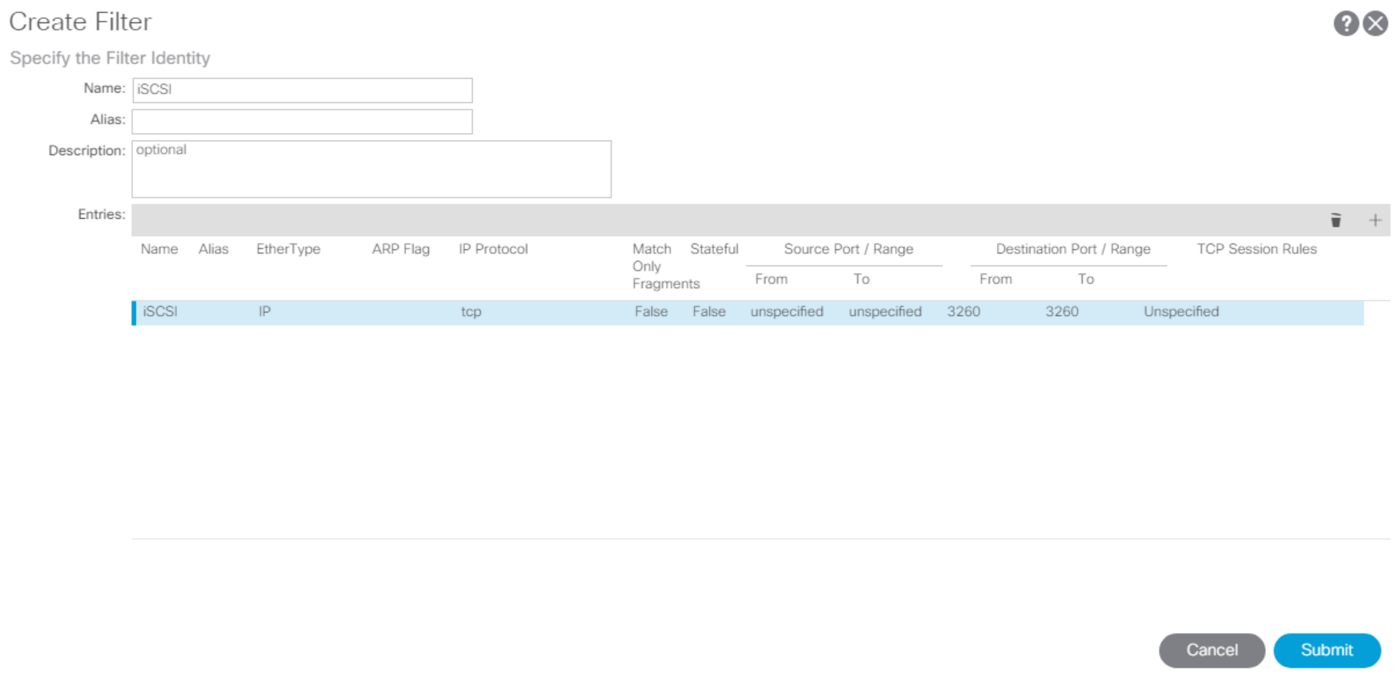

Create Filter

Specify the Filter Identity

Name: iSCSI

Alias:

Description: optional

Entries:

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
							From	To	From	To	
iSCSI		IP		tcp	False	False	unspecified	unspecified	3260	3260	Unspecified

31. Click SUBMIT to complete adding the Filter.

By adding these Filters to Tenant common, they can be used from within any Tenant in the ACI Fabric.

Deploy FPV-Foundation Tenant

This section details the steps for creating the FPV-Foundation Tenant in the ACI Fabric. This tenant will host infrastructure connectivity for the compute (VMware ESXi on UCS nodes) and the storage environments. To deploy the FPV-Foundation Tenant, complete the following steps:

1. In the APIC GUI, select Tenants > Add Tenant.

2. Name the Tenant FPV-Foundation.

3. For the VRF Name, enter FPV-Foundation. Keep the check box “Take me to this tenant when I click finish” checked.

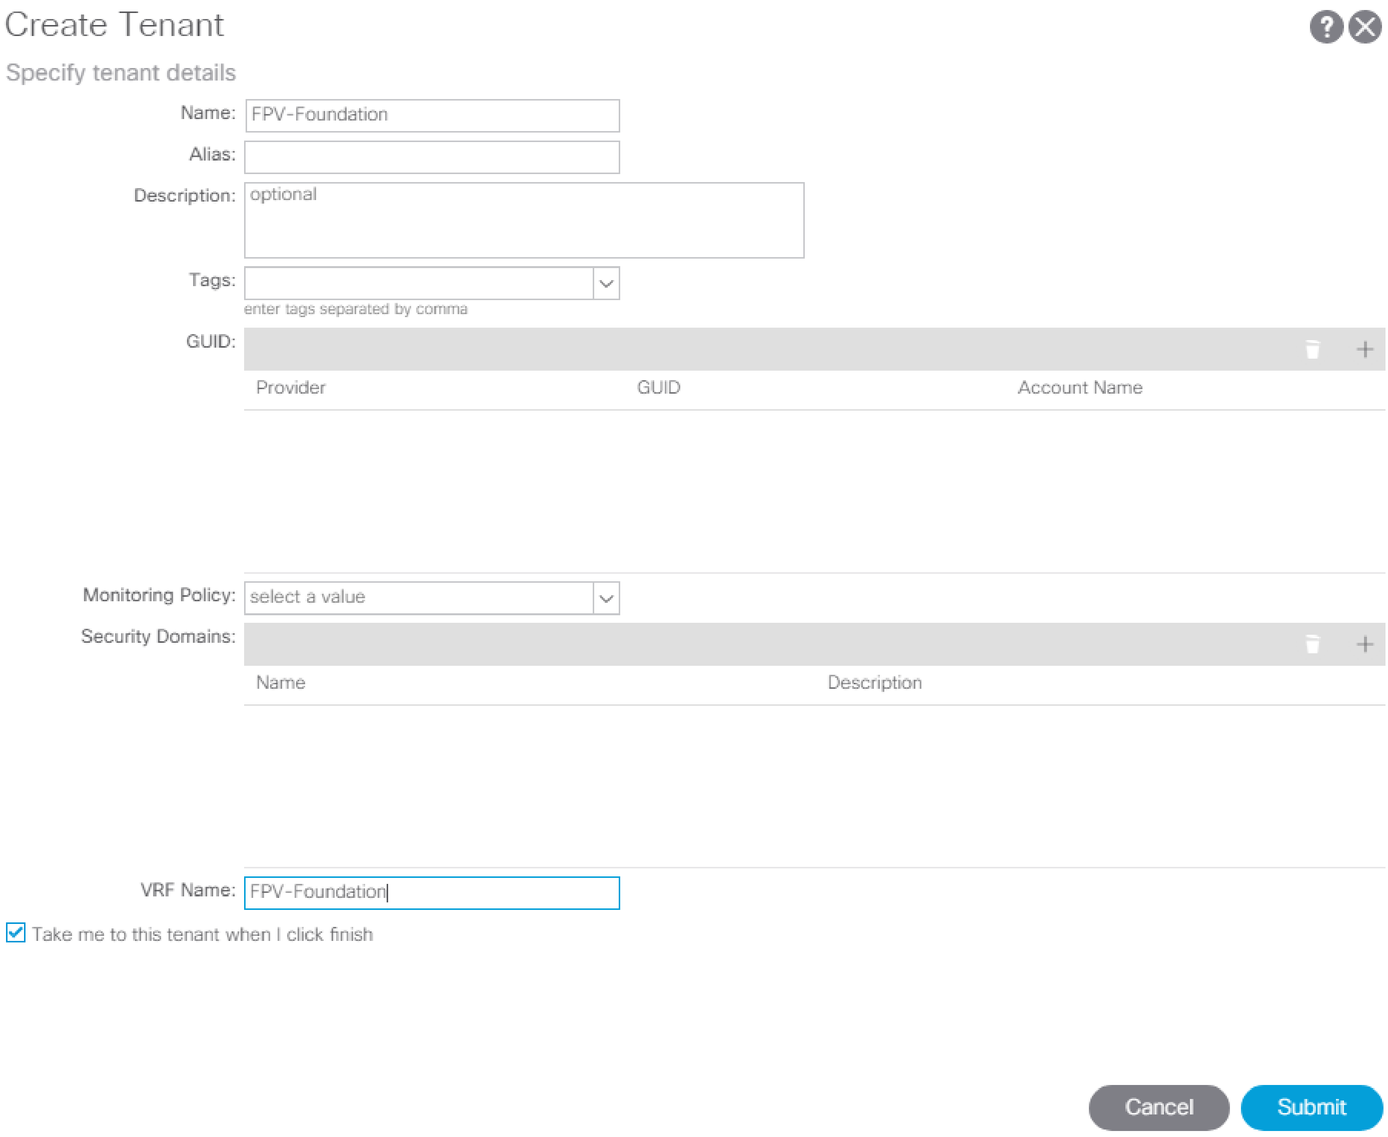

Create Tenant

Specify tenant details

Name: FPV-Foundation

Alias:

Description: optional

Tags:

enter tags separated by comma

GUID:

Provider	GUID	Account Name

Monitoring Policy: select a value

Security Domains:

Name	Description

VRF Name: FPV-Foundation

☑ Take me to this tenant when I click finish

4. Click **Submit** to finish creating the Tenant.

Create Bridge Domain

To create a Bridge Domain, complete the following steps:

1. In the left pane, expand Tenant FPV-Foundation and Networking.

2. Right-click Bridge Domains and select **Create Bridge Domain**.

3. Name the Bridge Domain BD-FPV-Foundation-Internal.

4. Select FPV-Foundation from the VRF drop-down list.

5. Select Custom under Forwarding and enable flooding.

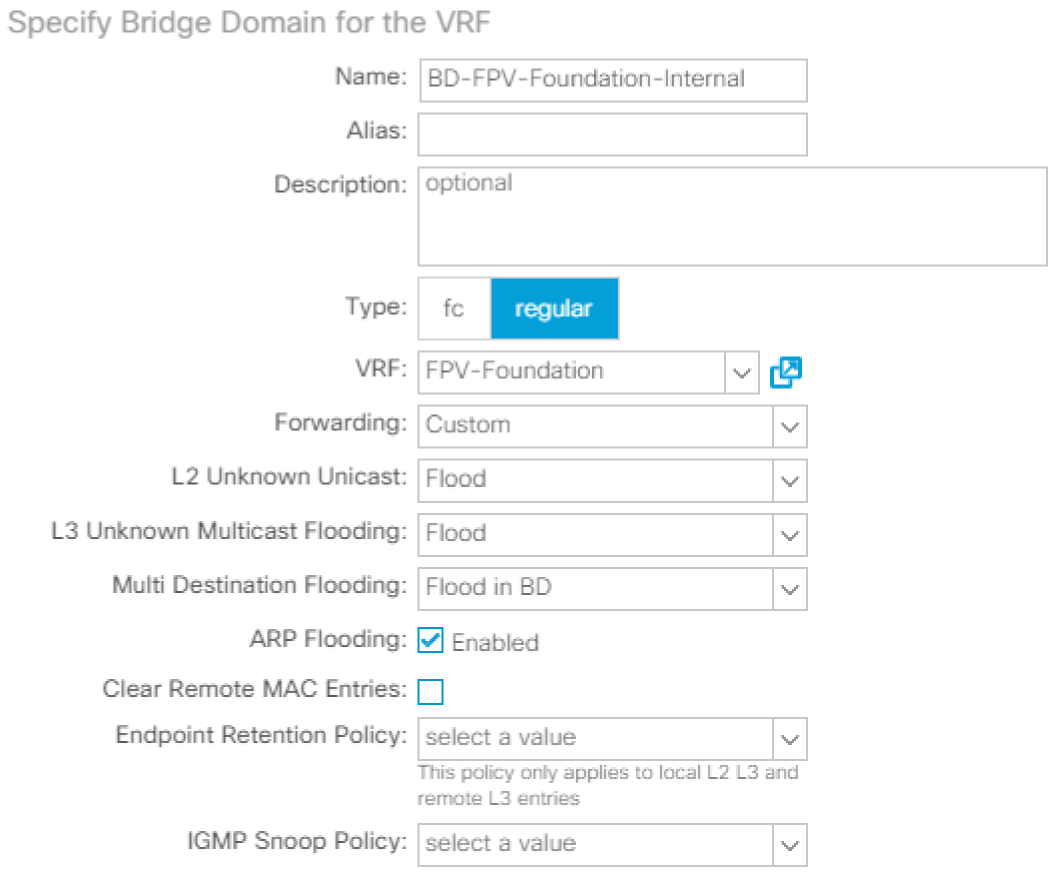

Specify Bridge Domain for the VRF

Name: BD-FPV-Foundation-Internal

Alias:

Description: optional

Type: fc [✓ regular]

VRF: FPV-Foundation

Forwarding: Custom

L2 Unknown Unicast: Flood

L3 Unknown Multicast Flooding: Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: ☑ Enabled

Clear Remote MAC Entries: □

Endpoint Retention Policy: select a value

This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: select a value

6. Click **Next**.

7. Under L3 Configurations, make sure Limit IP Learning to Subnet is selected and select EP Move Detection Mode – GARP based detection. Select **Next**.

8. No changes are needed for Advanced/Troubleshooting. Click **Finish** to finish creating the Bridge Domain.

Create Application Profile for IB-Management Access

To create an application profile for IB-Management Access, complete the following steps:

1. In the left pane, expand tenant FPV-Foundation, right-click Application Profiles and select **Create Application Profile**.

2. Name the Application Profile IB-MGMT and click **Submit** to complete adding the Application Profile.

Create EPG for IB-MGMT Access

This EPG will be used for VMware ESXi hosts and management virtual machines that are in the IB-MGMT subnet, but that do not provide ACI fabric Core Services. For example, AD server VMs could be placed in the Core Services EPG defined earlier to provide DNS services to tenants in the Fabric. Monitoring VMs can be placed in the IB-MGMT EPG and they will have access to the Core Services VMs, but will not be reachable from Tenant VMs.

To create the EPG for IB-MGMT access, complete the following steps:

1. In the left pane, expand the Application Profiles and right-click the IB-MGMT Application Profile and select **Create Application EPG**.

2. Name the EPG IB-MGMT.

3. From the Bridge Domain drop-down list, select Bridge Domain BD-FP-common-Core-Services from Tenant common.

Placing this EPG in the BD-FP-common-Core-Services Bridge Domain enables L2 connectivity between hosts and VMs placed in this EPG and hosts and VMs placed in the Core-Services EPG in Tenant common, including the Core-Services default gateway that is mapped into the Core-Services EPG as an External Bridged Device.

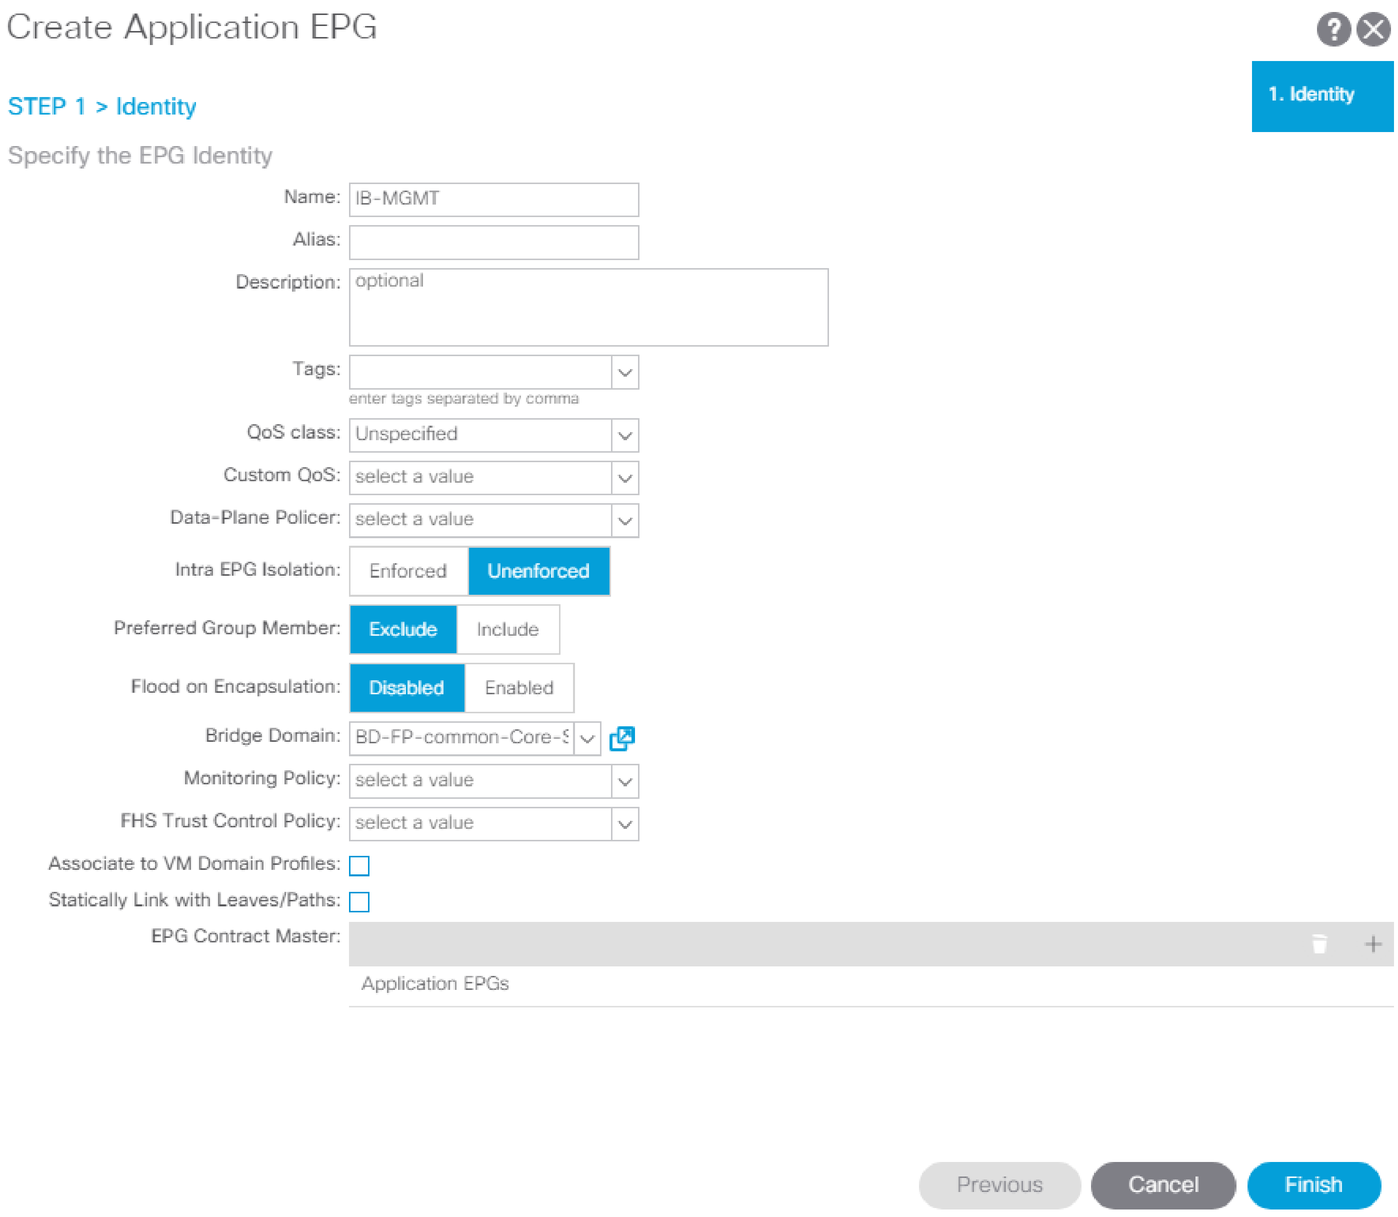

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name:	IB-MGMT
Alias:	
Description:	optional
Tags:	enter tags separated by comma
QoS class:	Unspecified
Custom QoS:	select a value
Data-Plane Policer:	select a value
Intra EPG Isolation:	Enforced [✓ Unenforced]
Preferred Group Member:	[✓ Exclude] Include
Flood on Encapsulation:	[✓ Disabled] Enabled
Bridge Domain:	BD-FP-common-Core-S
Monitoring Policy:	select a value
FHS Trust Control Policy:	select a value
Associate to VM Domain Profiles:	□
Statically Link with Leaves/Paths:	□
EPG Contract Master:	Application EPGs

4. Click **Finish** to complete creating the EPG.

5. In the left menu, expand the newly created EPG, right-click Domains and select **Add L2 External Domain Association**.

6. Select the UCS L2 External Domain Profile and click **Submit**.

7. In the left menu, right-click Static Ports and select **Deploy Static EPG on PC, VPC, or Interface**.

8. Select the Virtual Port Channel Path Type, then for Path select the vPC for the first UCS Fabric Interconnect.

9. For Port Encap leave VLAN selected and fill in the UCS IB-MGMT VLAN ID <419>.

10. Set the Deployment Immediacy to Immediate and click **Submit**.

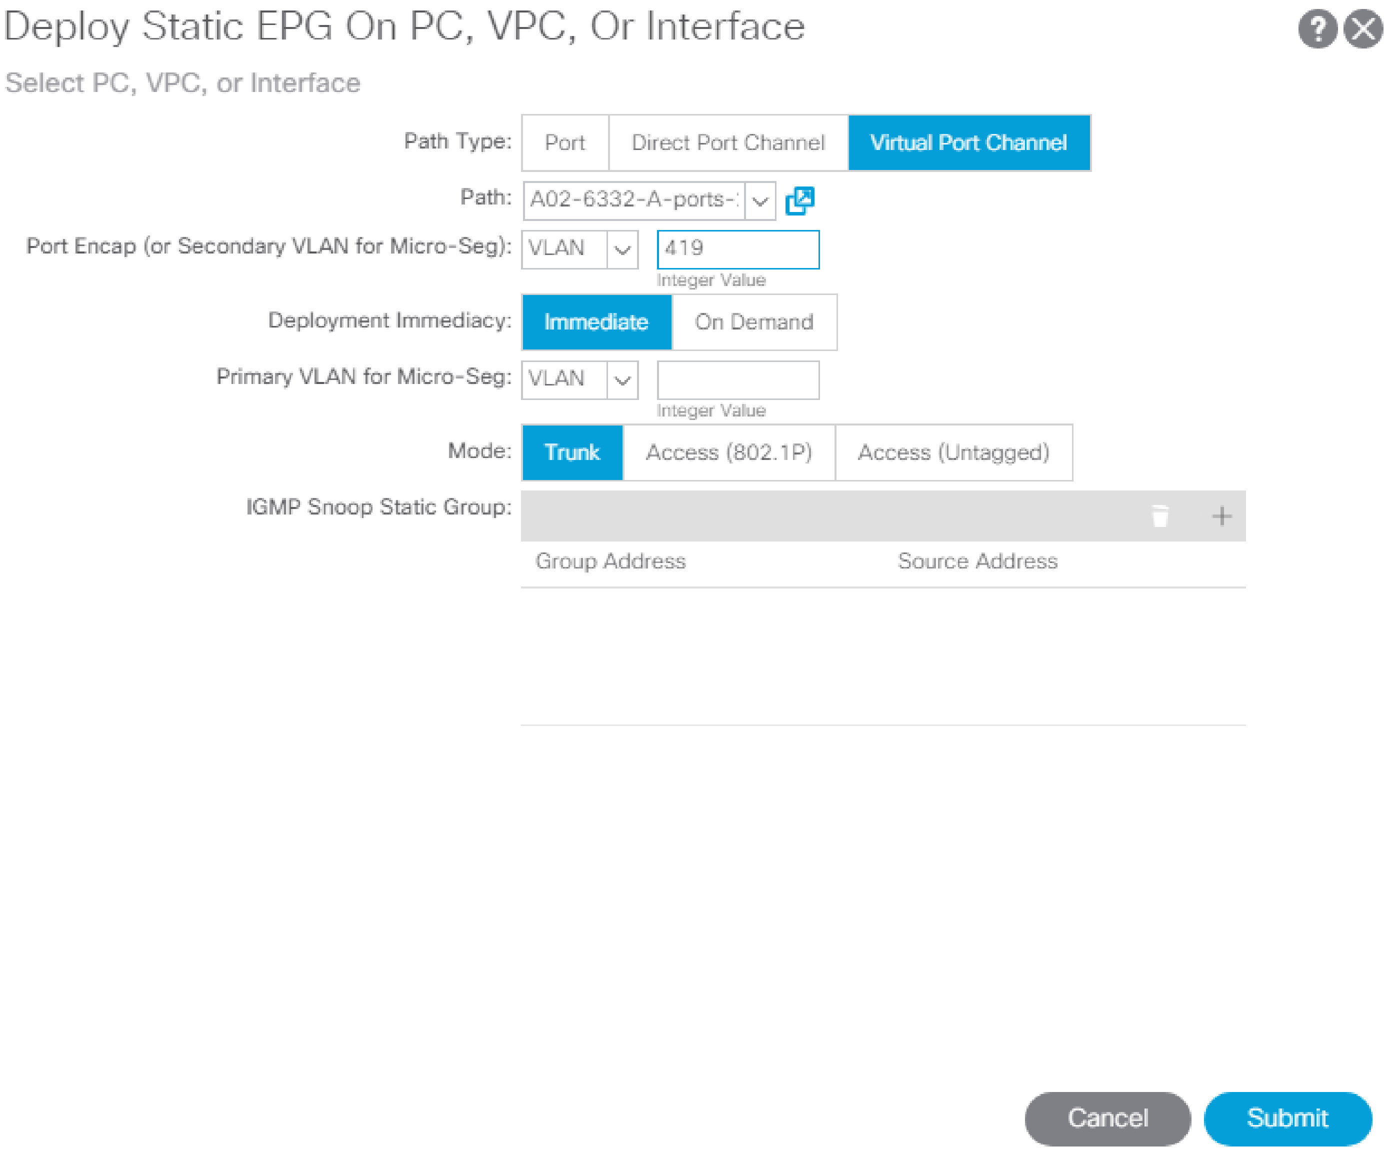

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port | Direct Port Channel | [✓ Virtual Port Channel]

Path: A02-6332-A-ports-:

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 419 Integer Value

Deployment Immediacy: [✓ Immediate] | On Demand

Primary VLAN for Micro-Seg: VLAN Integer Value

Mode: [✓ Trunk] | Access (802.1P) | Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

11. Repeat steps 7-10 to add the Static Port mapping for the second UCS Fabric Interconnect.

12. In the left menu, right-click Contracts and select **Add Consumed Contract.**

13. From the drop-down list for the Contract, select common-Allow-Core-Services from Tenant common.

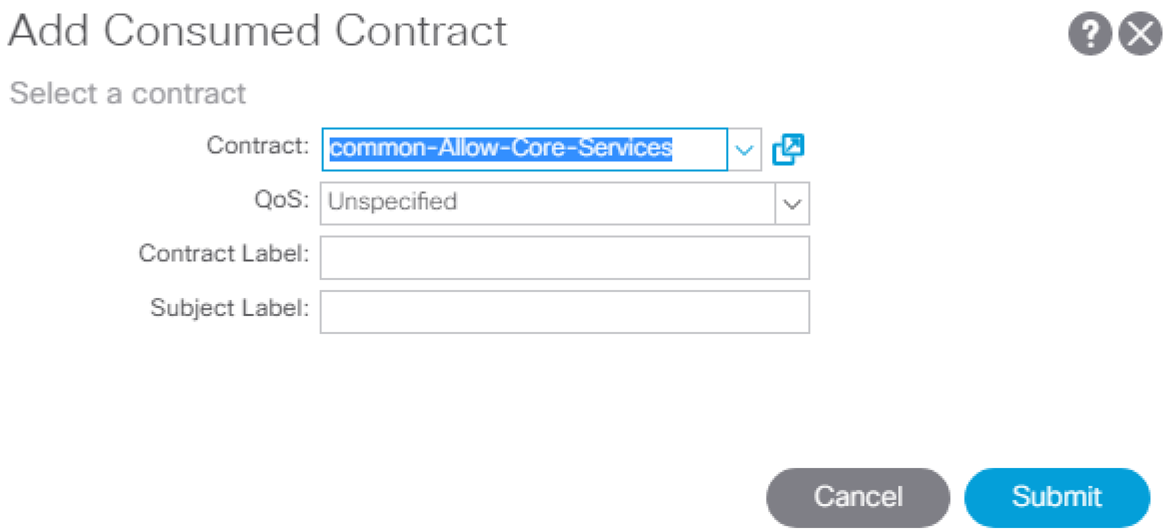

Add Consumed Contract

Select a contract

Contract: common-Allow-Core-Services

QoS: Unspecified

Contract Label:

Subject Label:

14. Click **Submit**.

This EPG will be utilized to provide ESXi hosts as well as the VMs that do not provide Core Services access to the existing in-band management network.

Create EPG for Infrastructure SVM-MGMT Access

This EPG will be used for VMs placed in the Core-Services EPG to reach the NetApp Infrastructure (FPV-Foundation-SVM) SVM management interface. This EPG will be placed in a different Bridge Domain than BD-FP-common-Core-Services so that multiple SVM management interfaces (with the same MAC address) can exist on the same storage controller. L3 will be used by the SVM-MGMT interface to communicate with the Core-Services VMs.

To create the EPG for SVM-MGMT access, complete the following steps:

1. In the left pane, expand the Application Profiles and right-click the IB-MGMT Application Profile and select **Create Application EPG**.

2. Name the EPG SVM-MGMT.

3. From the Bridge Domain drop-down list, select Bridge Domain BD-FPV-Foundation-Internal from Tenant FPV-Foundation.

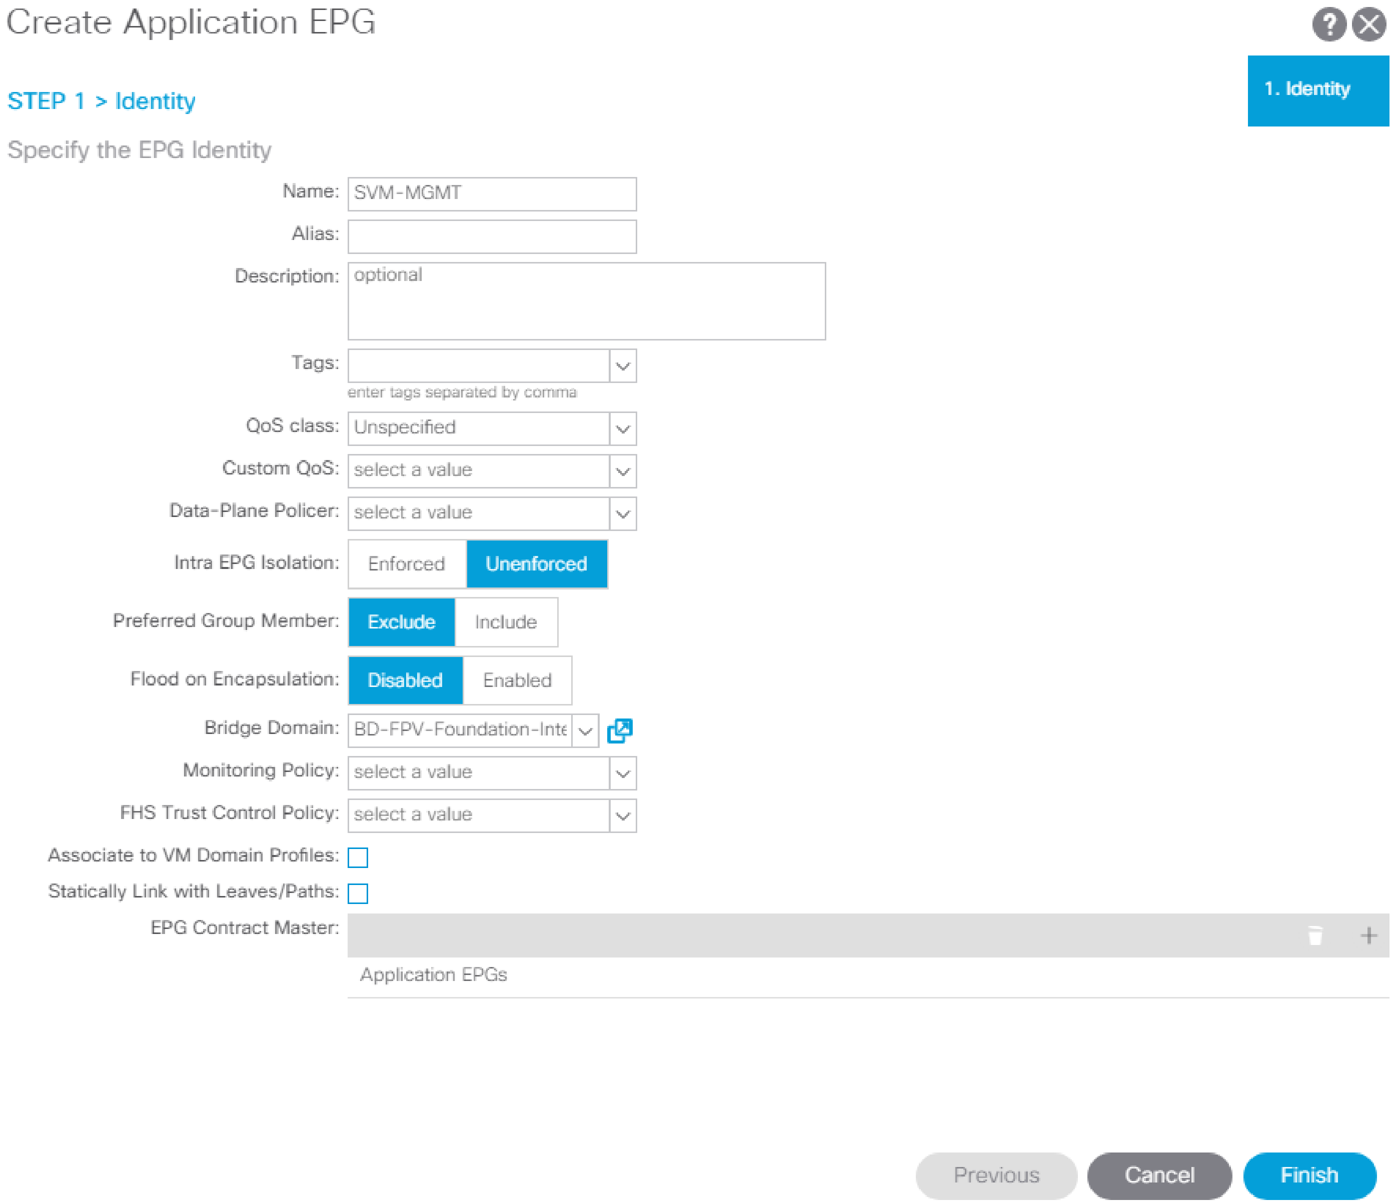

Create Application EPG

STEP 1 > Identity

1. Identity

Specify the EPG Identity

Name: SVM-MGMT

Alias:

Description: optional

Tags:
enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Flood on Encapsulation: [✓ Disabled] Enabled

Bridge Domain: BD-FPV-Foundation-Inte

Monitoring Policy: select a value

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

4. Click **Finish** to complete creating the EPG.

5. In the left menu, expand the newly created EPG, right-click Domains and select **Add Physical Domain Association**.

6. Select the NetApp-AFF Physical Domain Profile and click **Submit**.

7. In the left menu, right-click Static Ports and select **Deploy Static EPG on PC, VPC, or Interface**.

8. Select the Virtual Port Channel Path Type, then for Path select the vPC for the first NetApp AFF storage controller.

9. For Port Encap leave VLAN selected and fill in the storage SVM-MGMT VLAN ID <219>.

10. Set the Deployment Immediacy to Immediate and click **Submit**.

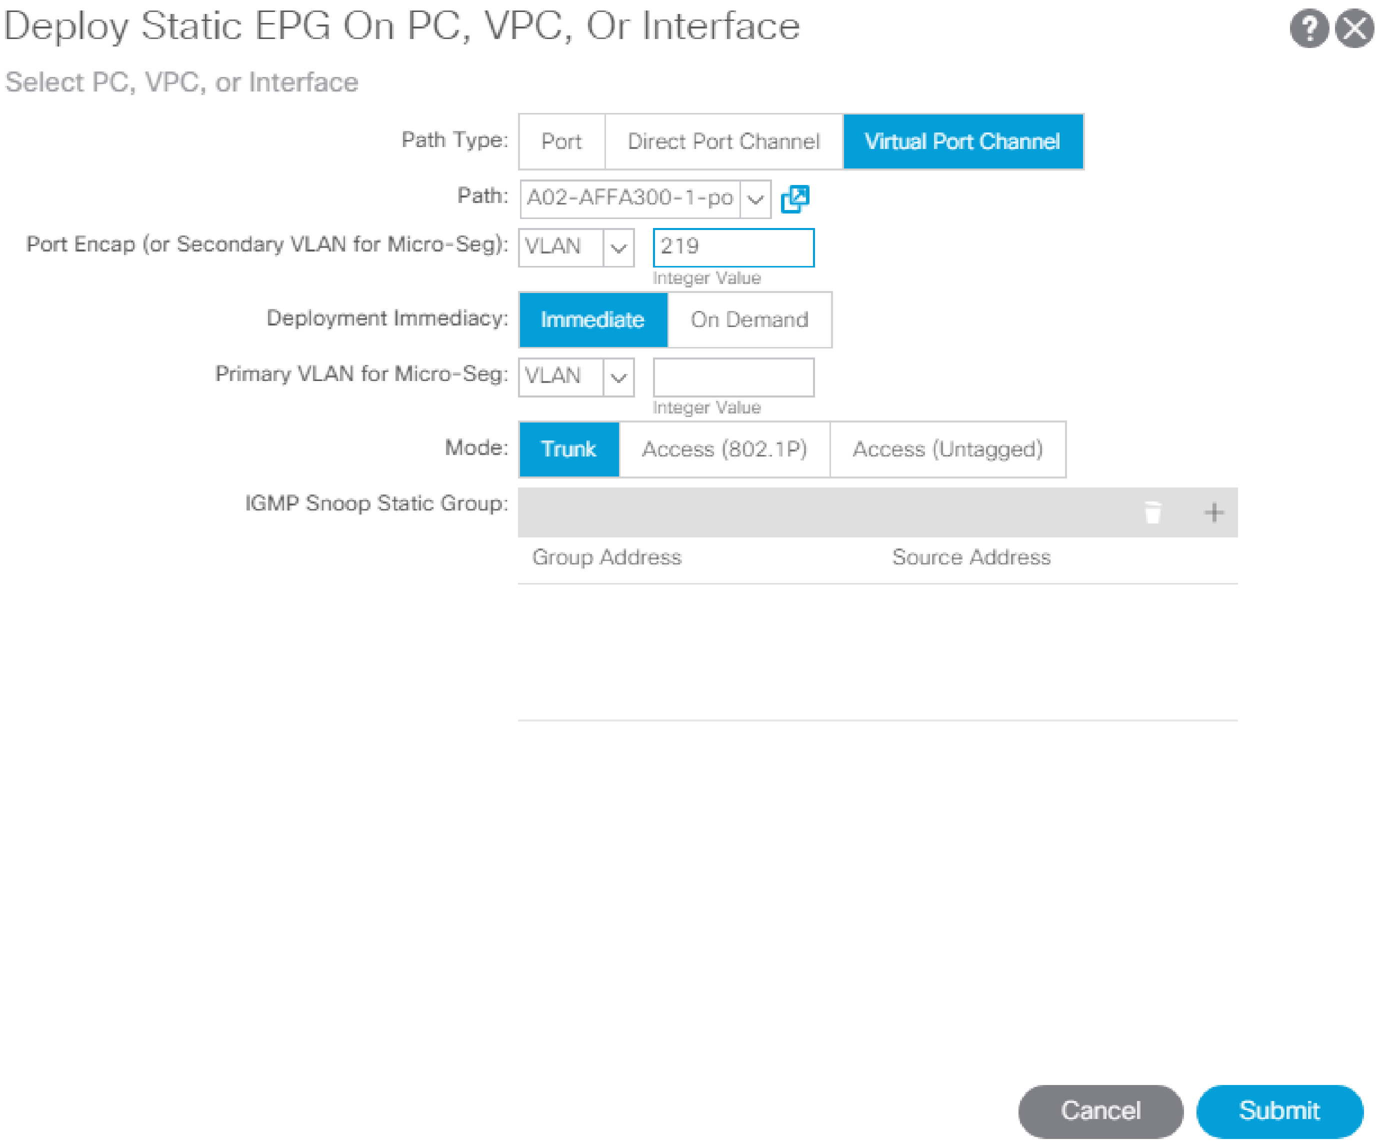

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Virtual Port Channel

Path: A02-AFFA300-1-po

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 219

Deployment Immediacy: Immediate

Primary VLAN for Micro-Seg: VLAN

Mode: Trunk

IGMP Snoop Static Group:

Group Address	Source Address

11. Repeat steps 7-10 to add the Static Port mapping for the second NetApp AFF storage controller.

12. In the left menu, right-click Subnets and select **Create EPG Subnet**.

13. Enter the Infrastructure SVM Management subnet gateway address and netmask <172.16.254.6/29> in the Default Gateway IP field.

14. Select only **Shared between VRFs** for the Scope and click **Submit** to add the subnet gateway.

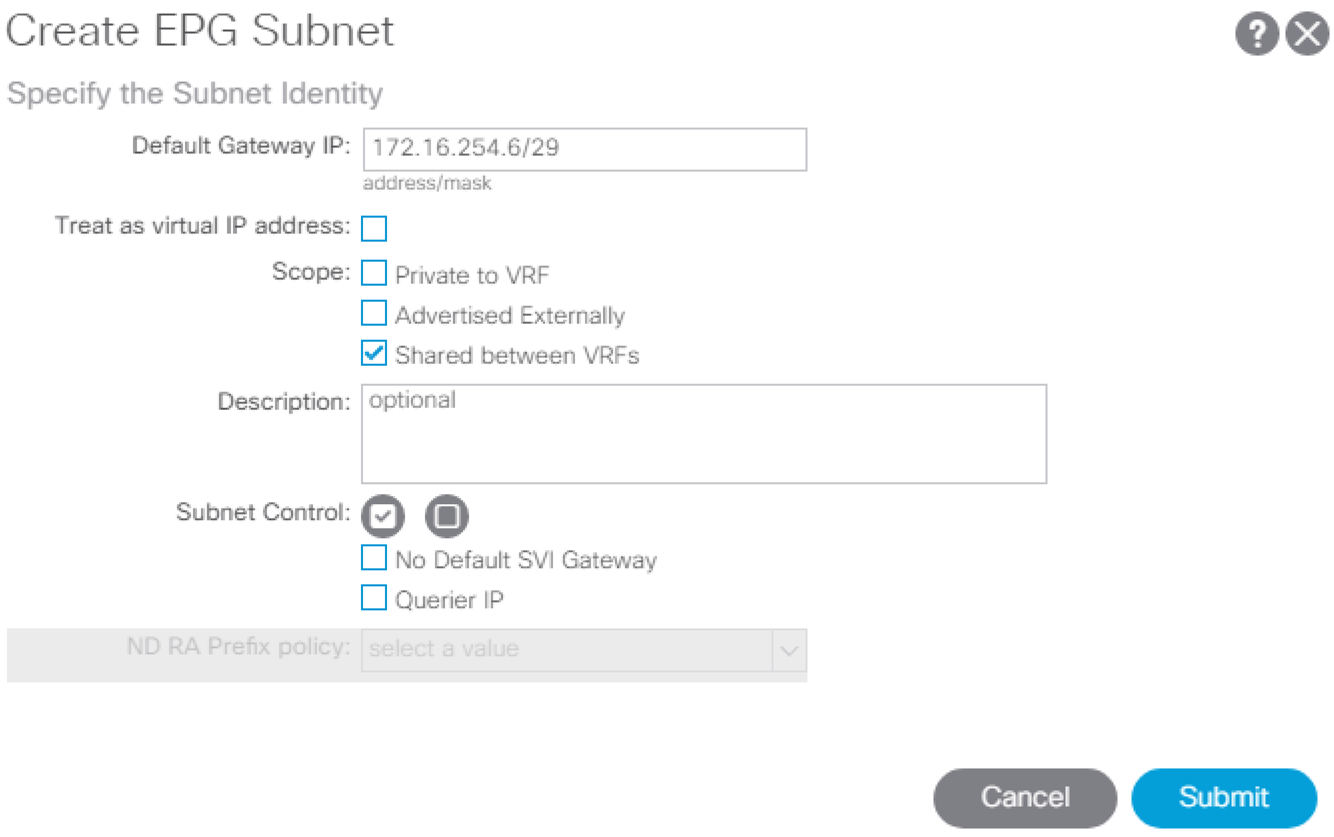

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP: 172.16.254.6/29
address/mask

Treat as virtual IP address: □

Scope: □ Private to VRF
□ Advertised Externally
☑ Shared between VRFs

Description: optional

Subnet Control:
□ No Default SVI Gateway
□ Querier IP

ND RA Prefix policy: select a value

15. In the left menu, right-click Contracts and select **Add Consumed Contract.**

16. From the drop-down list for the Contract, select common-Allow-Core-Services from Tenant common.

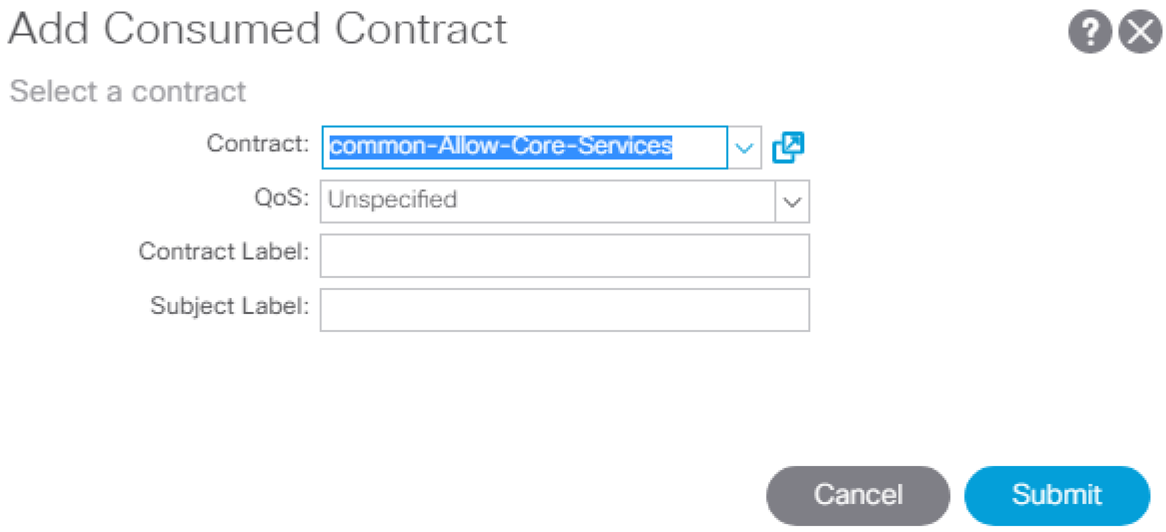

Add Consumed Contract

Select a contract

Contract: common-Allow-Core-Services

QoS: Unspecified

Contract Label:

Subject Label:

17. Click **Submit**.

Create Application Profile for Host Connectivity

To create an application profile for host connectivity, complete the following steps:

1. In the left pane, under the Tenant FP-Foundation, right-click Application Profiles and select **Create Application Profile**.

2. Name the Profile Host-Conn and click **Submit** to complete adding the Application Profile.

The following EPGs and the corresponding mappings will be created under this application profile.

Refer to Table 7 for the information required during the following configuration. Items marked by { } will need to be updated according to Table 7 . Note that since all storage interfaces on a single Interface Group on a NetApp AF-FA300 share the same MAC address, that different bridge domains must be used for each storage EPG.

Table 7 EPGs and mappings for AP-Host-Connectivity

EPG Name	Bridge Domain	Domain	Static Port – Compute	Static Port - Storage
vMotion	BD-FPV-Foundation-Internal	L2 External: UCS	VPC for all UCS FIs VLAN 3000	N/A
iSCSI-A	BD-FPV-Foundation-iSCSI-A	L2 External: UCS Physical: NetApp-AFF	VPC for all UCS FIs VLAN 3110	VPC for all NetApp AFFs VLAN 3010
iSCSI-B	BD-FPV-Foundation-iSCSI-B	L2 External: UCS Physical: NetApp-AFF	VPC for all UCS FIs VLAN 3120	VPC for all NetApp AFFs VLAN 3020
NFS	BD-FPV-Foundation-NFS	L2 External: UCS Physical: NetApp-AFF	VPC for all UCS FIs VLAN 3150	VPC for all NetApp AFFs VLAN 3050

Create Bridge Domains and EPGs

To create bridge domains and EPGs, complete the following steps:

1. For each row in the table above, if the Bridge Domain does not already exist, in the left pane, under Tenant FPV-Foundation, expand Networking > Bridge Domains.

2. Right-click Bridge Domains and select Create Bridge Domain.

3. Name the Bridge Domain {BD-FPV-Foundation-iSCSI-A}.

4. Select the FPV-Foundation VRF.

5. Select Custom for Forwarding and setup forwarding as shown in the screenshot.

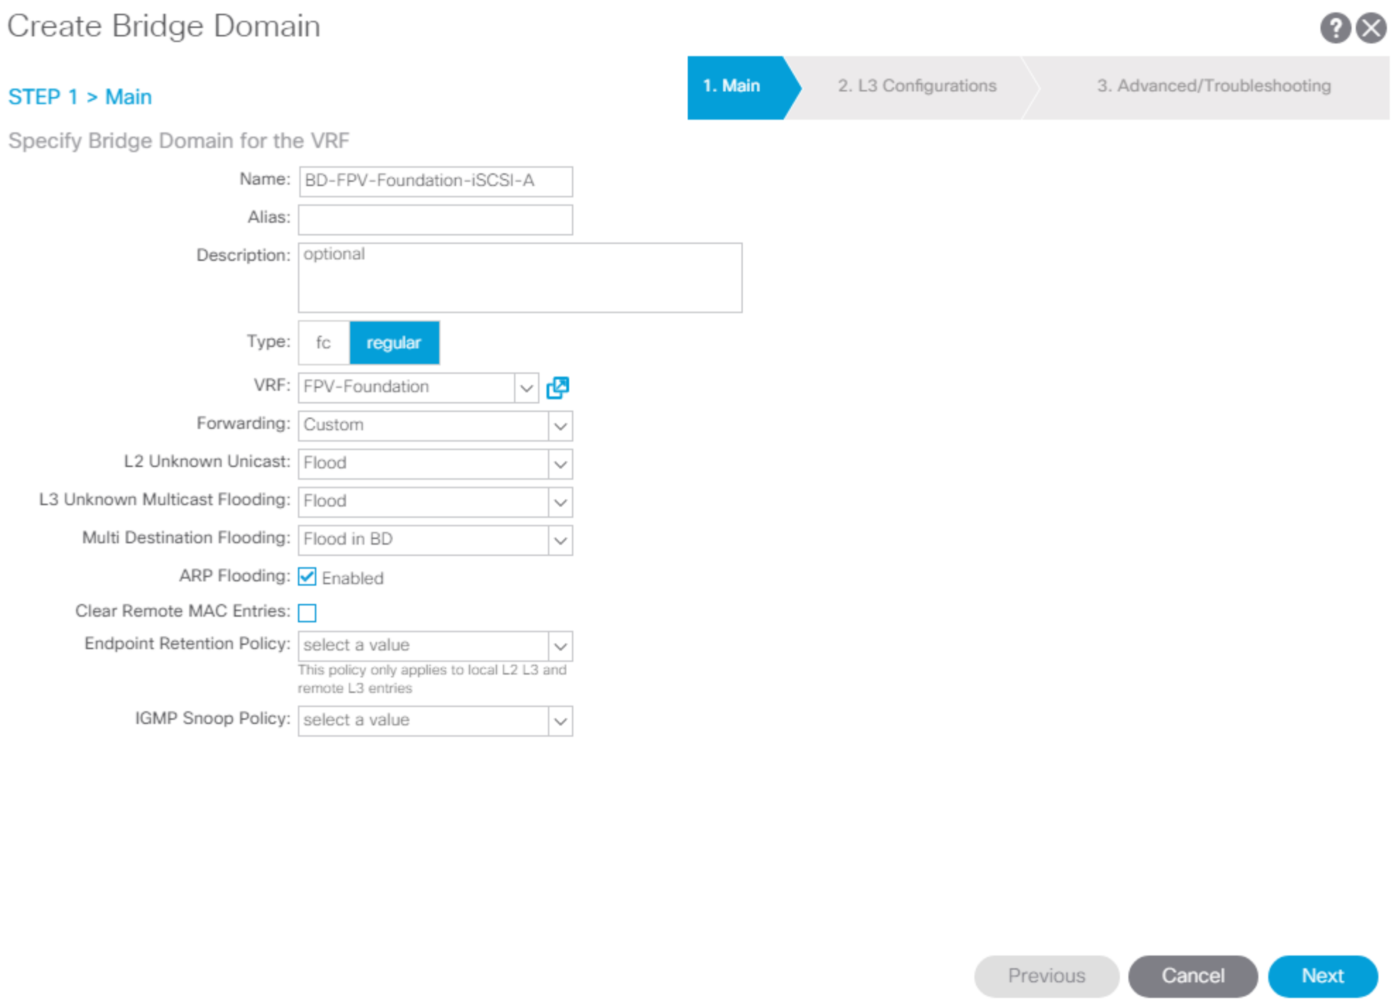

Create Bridge Domain

STEP 1 > Main

Specify Bridge Domain for the VRF

Name: BD-FPV-Foundation-iSCSI-A

Alias:

Description: optional

Type: fc [✓ regular]

VRF: FPV-Foundation

Forwarding: Custom

L2 Unknown Unicast: Flood

L3 Unknown Multicast Flooding: Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: ☑ Enabled

Clear Remote MAC Entries: □

Endpoint Retention Policy: select a value

This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: select a value

6. Click **Next**.

7. Under L3 Configurations, make sure Limit IP Learning to Subnet is selected and select EP Move Detection Mode – GARP based detection. Select **Next**.

8. No changes are needed for Advanced/Troubleshooting. Click **Finish** to finish creating the Bridge Domain.

9. In the left pane, expand Application Profiles > Host-Conn. Right-click Application EPGs and select **Create Application EPG**.

10. Name the EPG {vMotion}.

11. From the Bridge Domain drop-down list, select the Bridge Domain from the table.

12. Click **Finish** to complete creating the EPG.

13. In the left pane, expand the Application EPGs and EPG {vMotion}.

14. Right-click Domains and select Add L2 External Domain Association.

15. From the drop-down list, select the previously defined {UCS} L2 External Domain Profile.

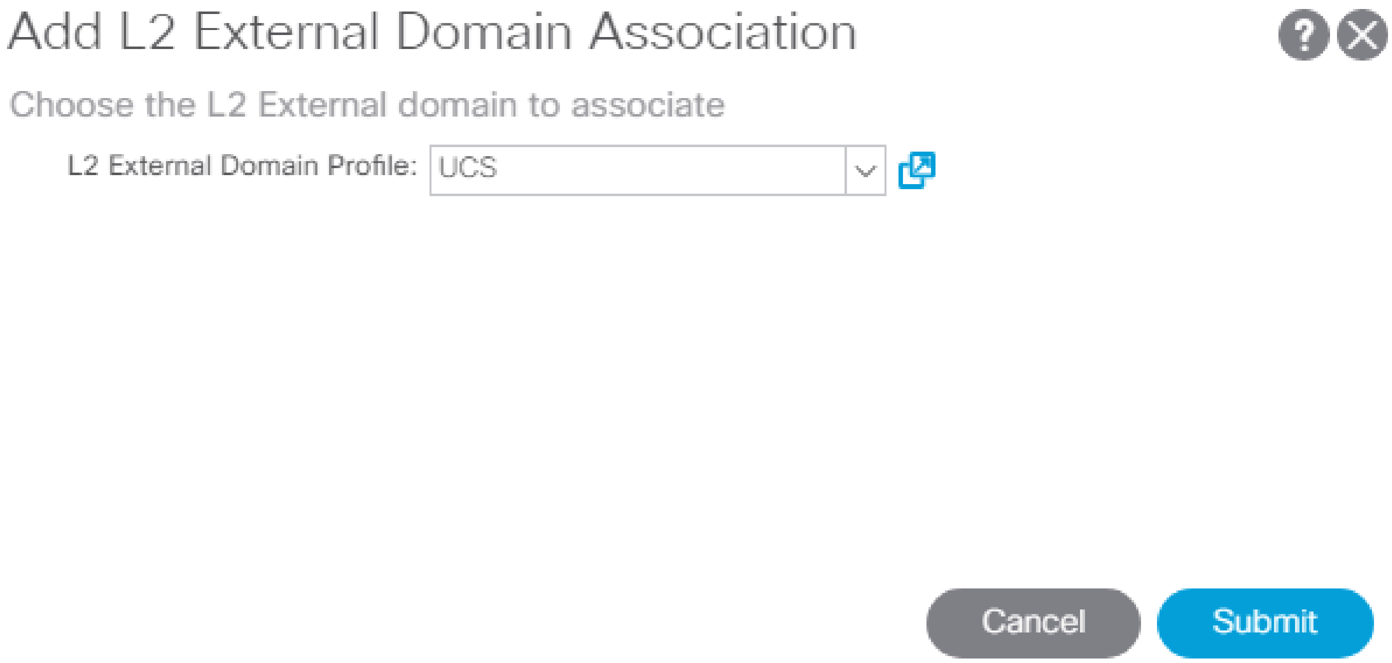

Add L2 External Domain Association

Choose the L2 External domain to associate

L2 External Domain Profile: UCS

16. Click **Submit** to complete the L2 External Domain Association.

17. Repeat the Domain Association steps (6-9) to add appropriate EPG specific domains from **Error! Reference source not found.**.

18. Right-click Static Ports and select **Deploy EPG on PC, VPC, or Interface**.

19. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.

20. From the drop-down list, select the appropriate VPCs.

21. Enter VLAN from **Error! Reference source not found.** {3000} for Port Encap.

22. Select Immediate for Deployment Immediacy and for Mode select Trunk.

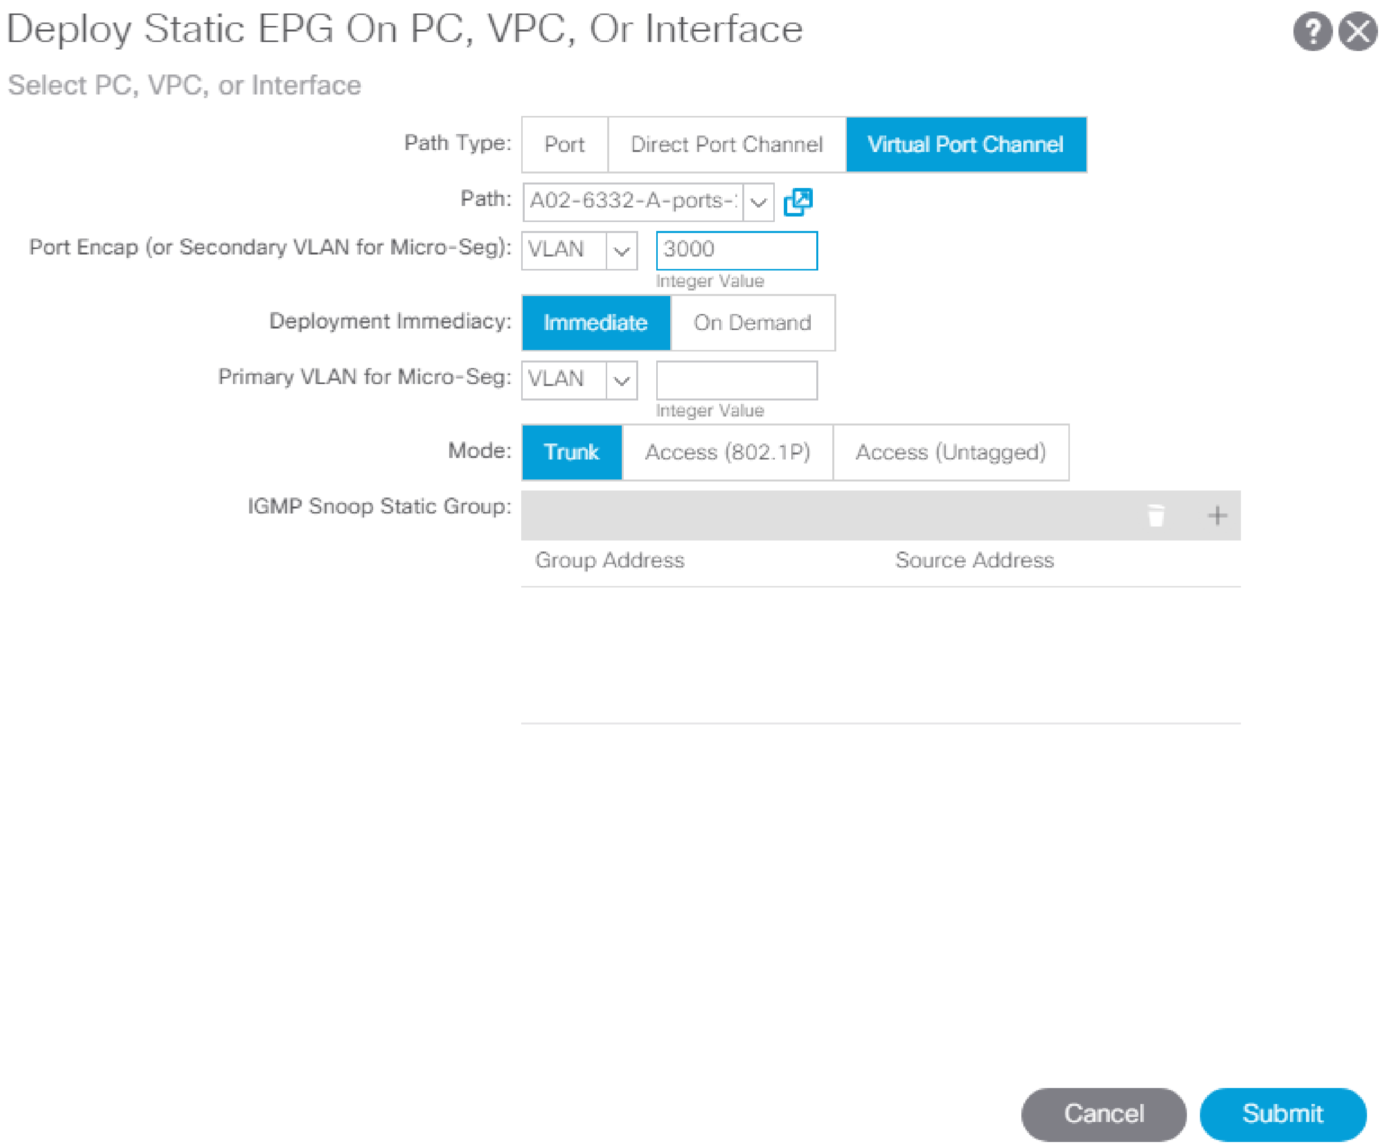

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port | Direct Port Channel | [✓ Virtual Port Channel]

Path: A02-6332-A-ports-:

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 3000 (Integer Value)

Deployment Immediacy: [✓ Immediate] | On Demand

Primary VLAN for Micro-Seg: VLAN (Integer Value)

Mode: [✓ Trunk] | Access (802.1P) | Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

23. Click **Submit** to complete adding the Static Path Mapping.

24. Repeat the above steps to add all the Static Path Mappings for the EPG listed in Table 7 .

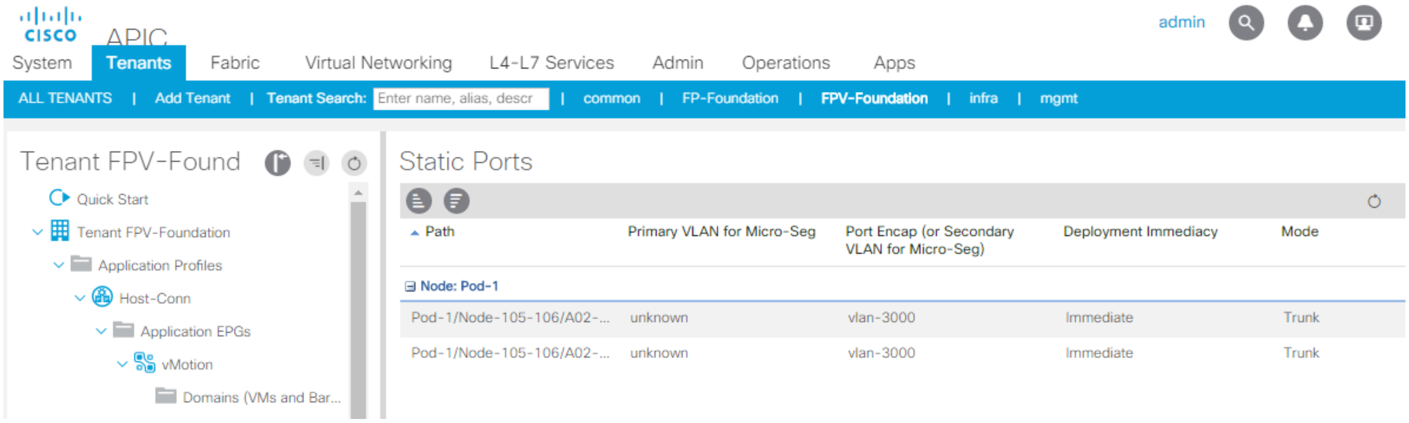

Static Ports

Path	Primary VLAN for Micro-Seg	Port Encap (or Secondary VLAN for Micro-Seg)	Deployment Immediacy	Mode
Node: Pod-1				
Pod-1/Node-105-106/A02-...	unknown	vlan-3000	Immediate	Trunk
Pod-1/Node-105-106/A02-...	unknown	vlan-3000	Immediate	Trunk

25. Optionally, add subnets for each EPG created, providing a default gateway that can be pinged for troubleshooting purposes.

Note that any of the storage EPGs could have been broken up into two EPGs (EPG-VMK and EPG-LIF) connected by contract. In that case, the two EPGs would need to be placed in the same Bridge Domain and one EPG would "provide" a contract (optionally filtered) and the other EPG would "consume" this contract. The two EPGs would then use L2 connectivity and it is not recommended to use subnets in this case.

Storage Configuration

Pursuant to best practices, NetApp recommends the following command on the LOADER prompt of the NetApp controllers to assist with LUN stability during copy operations. To access the LOADER prompt, connect to the controller via serial console port or Service Processor connection and press Ctrl-C to halt the boot process when prompted.

```
setenv bootarg.tmgr.disable_pit_hp 1
```

For more information about the workaround, see: http://nt-ap.com/2w6myr4 (requires Support login).

For more information about Windows Offloaded Data Transfers see: https://technet.microsoft.com/en-us/library/hh831628(v=ws.11).aspx.

NetApp All Flash FAS A300 Controllers

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities. Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the HWU application (requires a software subscription) at the NetApp Support site.

To access the HWU application to view the System Configuration guides, complete the following steps:

1. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

2. To compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers found in the AFF A300 Series product documentation at the NetApp Support site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported by the AFF A300 is available at the NetApp Support site.

For SAS disk shelves with NetApp storage controllers, refer to the SAS Disk Shelves Universal SAS and ACP Cabling Guide for proper cabling guidelines.

NetApp ONTAP 9.3

Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the ONTAP 9 Software Setup Guide. You must have access to the NetApp Support site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the ONTAP 9 Software Setup Guide to learn about configuring ONTAP. Table 8 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 8 ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Data ONTAP 9.3 URL	<url-boot-software>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.3 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.3 is the version being booted, select option 8 and y to reboot the node, then continue with step 14.

4. To install new software, select option 7.

5. Enter `y` to perform an upgrade.

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

7. Enter `y` to reboot now.

8. Enter the IP address, netmask, and default gateway for `e0M`.

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

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

This web server must be reachable.

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

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

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

13. Press Ctrl-C when the following message displays:

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

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

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with node 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.3 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.3 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.

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

7. Enter y to reboot now.

8. Enter the IP address, netmask, and default gateway for `e0M`.

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

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

This web server must be reachable.

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

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

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

12. Enter y to reboot the node.

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

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

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

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

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

16. Enter y to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with node 02 configuration while the disks for node 01 are zeroing.

Set Up Node

To set up a node, complete the following steps:

1. From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.3 boots on the node for the first time.

2. Follow the prompts to set up node 01:

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

You can return to cluster setup at any time by typing “cluster setup”.
To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem
occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

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

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

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

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

Table 9 Cluster Create in ONTAP Prerequisites

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

Cluster Detail	Cluster Detail Value
NTP server IP address	<ntp-ip>

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

4. Click Guided Setup on the Welcome screen.

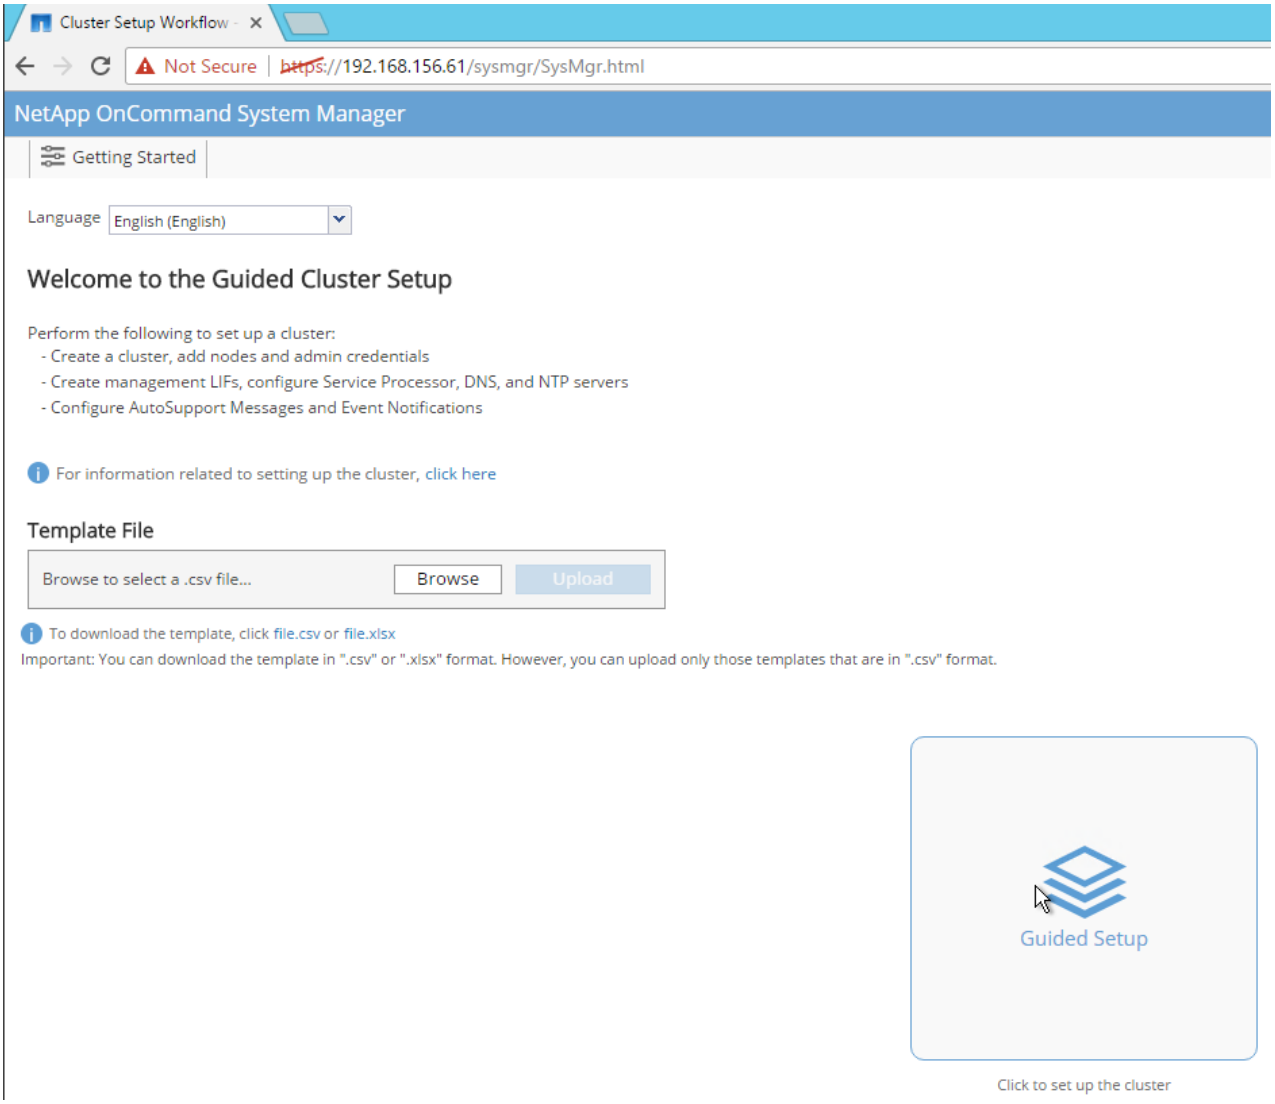

Cluster Setup Workflow - ×
Not Secure
~~https~~://192.168.156.61/sysmgr/SysMgr.html
NetApp OnCommand System Manager
Getting Started

Language English (English)

Welcome to the Guided Cluster Setup

Perform the following to set up a cluster:
- Create a cluster, add nodes and admin credentials
- Create management LIFs, configure Service Processor, DNS, and NTP servers
- Configure AutoSupport Messages and Event Notifications

For information related to setting up the cluster, click here

Template File

Browse to select a .csv file... Browse Upload

To download the template, click file.csv or file.xlsx

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

5. In the Cluster screen, complete the following steps:

   a. Enter the cluster and node names.

   b. Select the cluster configuration.

   c. Enter and confirm the password.

   d. Enter the cluster base and feature licenses.

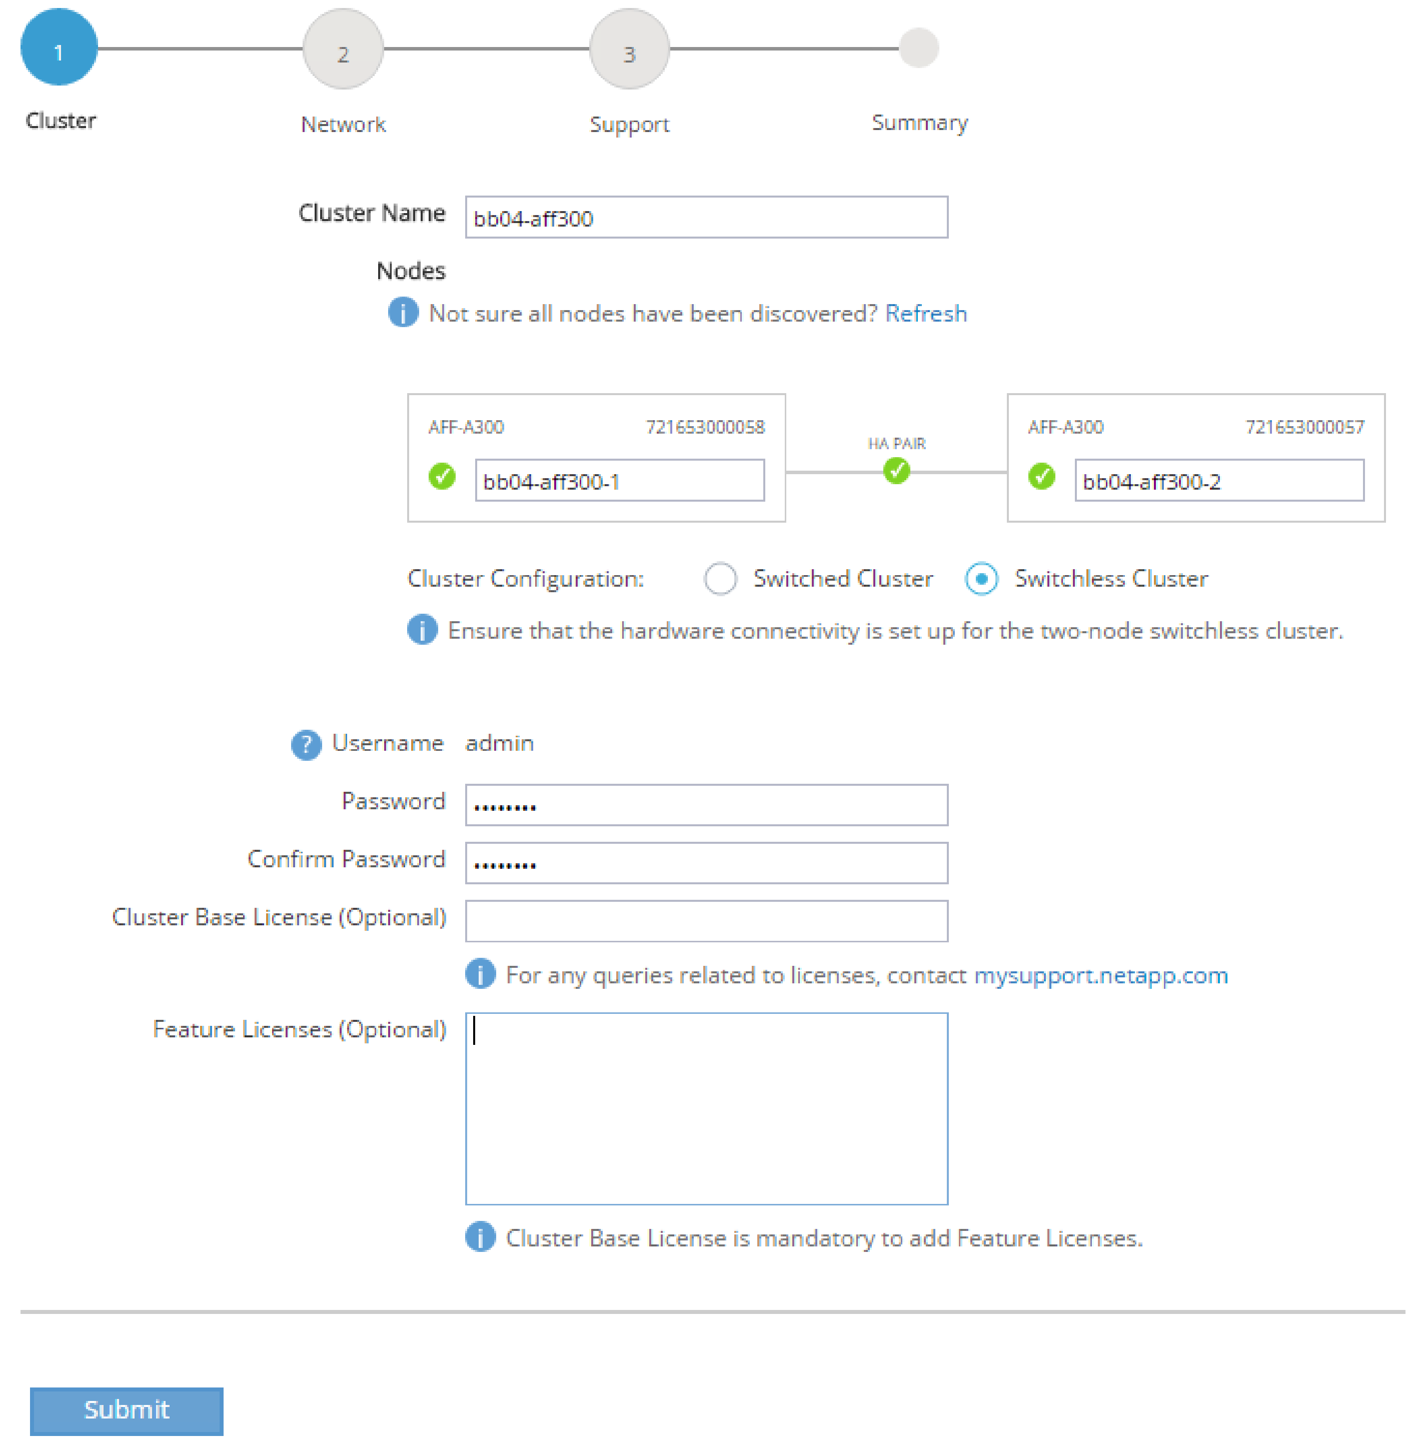

Cluster Name: bb04-aff300

Nodes

Not sure all nodes have been discovered? Refresh

AFF-A300	721653000058
✓ bb04-aff300-1	

HA PAIR ✓

AFF-A300	721653000057
✓ bb04-aff300-2	

Cluster Configuration: ○ Switched Cluster ● Switchless Cluster

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

Username: admin

Password: ••••••••

Confirm Password: ••••••••

Cluster Base License (Optional):

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

Feature Licenses (Optional):

Cluster Base License is mandatory to add Feature Licenses.

The nodes are discovered automatically, if they are not discovered, click the Refresh link. By default, the cluster interfaces are created on all new storage controllers shipped from the factory. If all the nodes are not discovered, then configure the cluster using the command line. Cluster license and feature licenses can also be installed after completing the cluster creation.

6. Click Submit.

7. On the network page, complete the following sections:

   a. Cluster Management

   – Enter the IP address, netmask, gateway and port details.

   b. Node Management

   – Enter the node management IP addresses and port details for all the nodes.

   c. Service Processor Management

   – Enter the IP addresses for all the nodes.

   d. DNS Details

   – Enter the DNS domain names and server address.

   e. NTP Details

   – Enter the primary and alternate NTP server.

8. Click Submit.

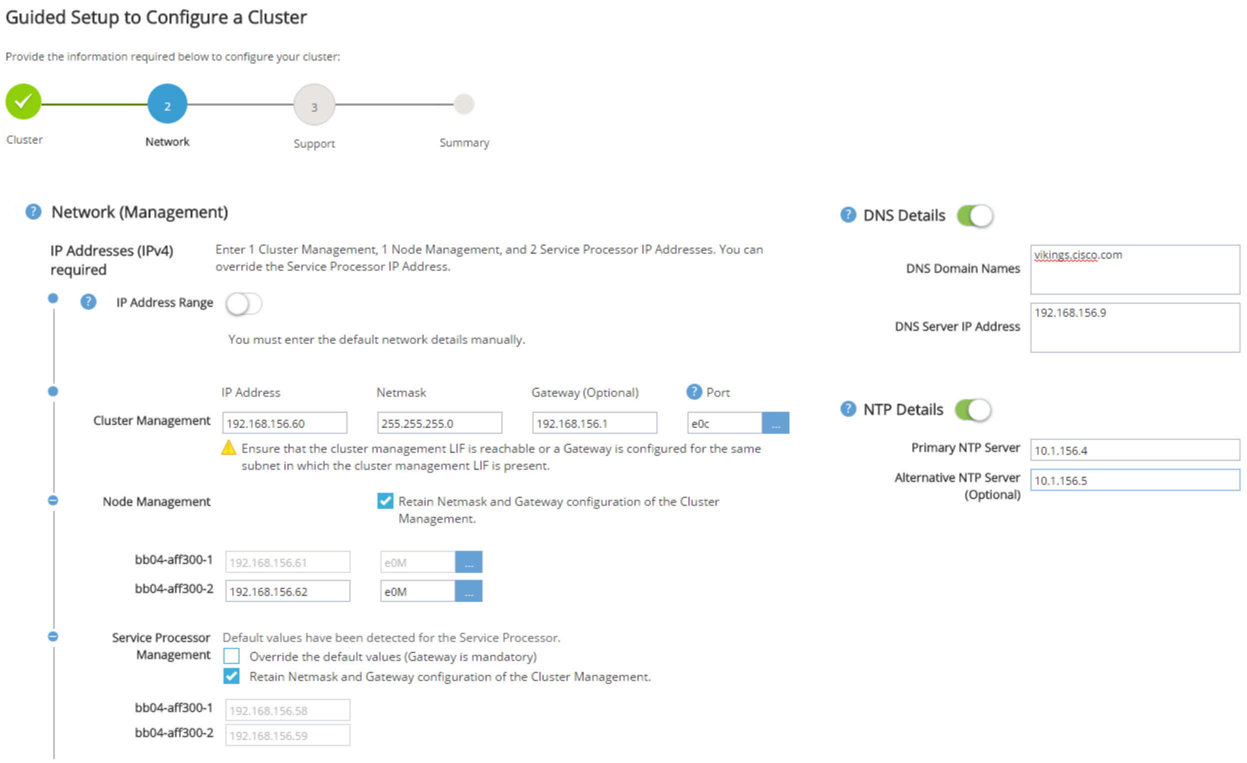

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Network (Management)

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

IP Address Range

You must enter the default network details manually.

	IP Address	Netmask	Gateway (Optional)	Port
Cluster Management	192.168.156.60	255.255.255.0	192.168.156.1	e0c

Ensure that the cluster management LIF is reachable or a Gateway is configured for the same subnet in which the cluster management LIF is present.

Node Management: ☑ Retain Netmask and Gateway configuration of the Cluster Management.

bb04-aff300-1	192.168.156.61	e0M
bb04-aff300-2	192.168.156.62	e0M

Service Processor Management: Default values have been detected for the Service Processor.

□ Override the default values (Gateway is mandatory)

☑ Retain Netmask and Gateway configuration of the Cluster Management.

bb04-aff300-1	192.168.156.58
bb04-aff300-2	192.168.156.59

DNS Details

DNS Domain Names: vikings.cisco.com

DNS Server IP Address: 192.168.156.9

NTP Details

Primary NTP Server: 10.1.156.4

Alternative NTP Server (Optional): 10.1.156.5

9. On the Support page, configure the AutoSupport and Event Notifications sections.

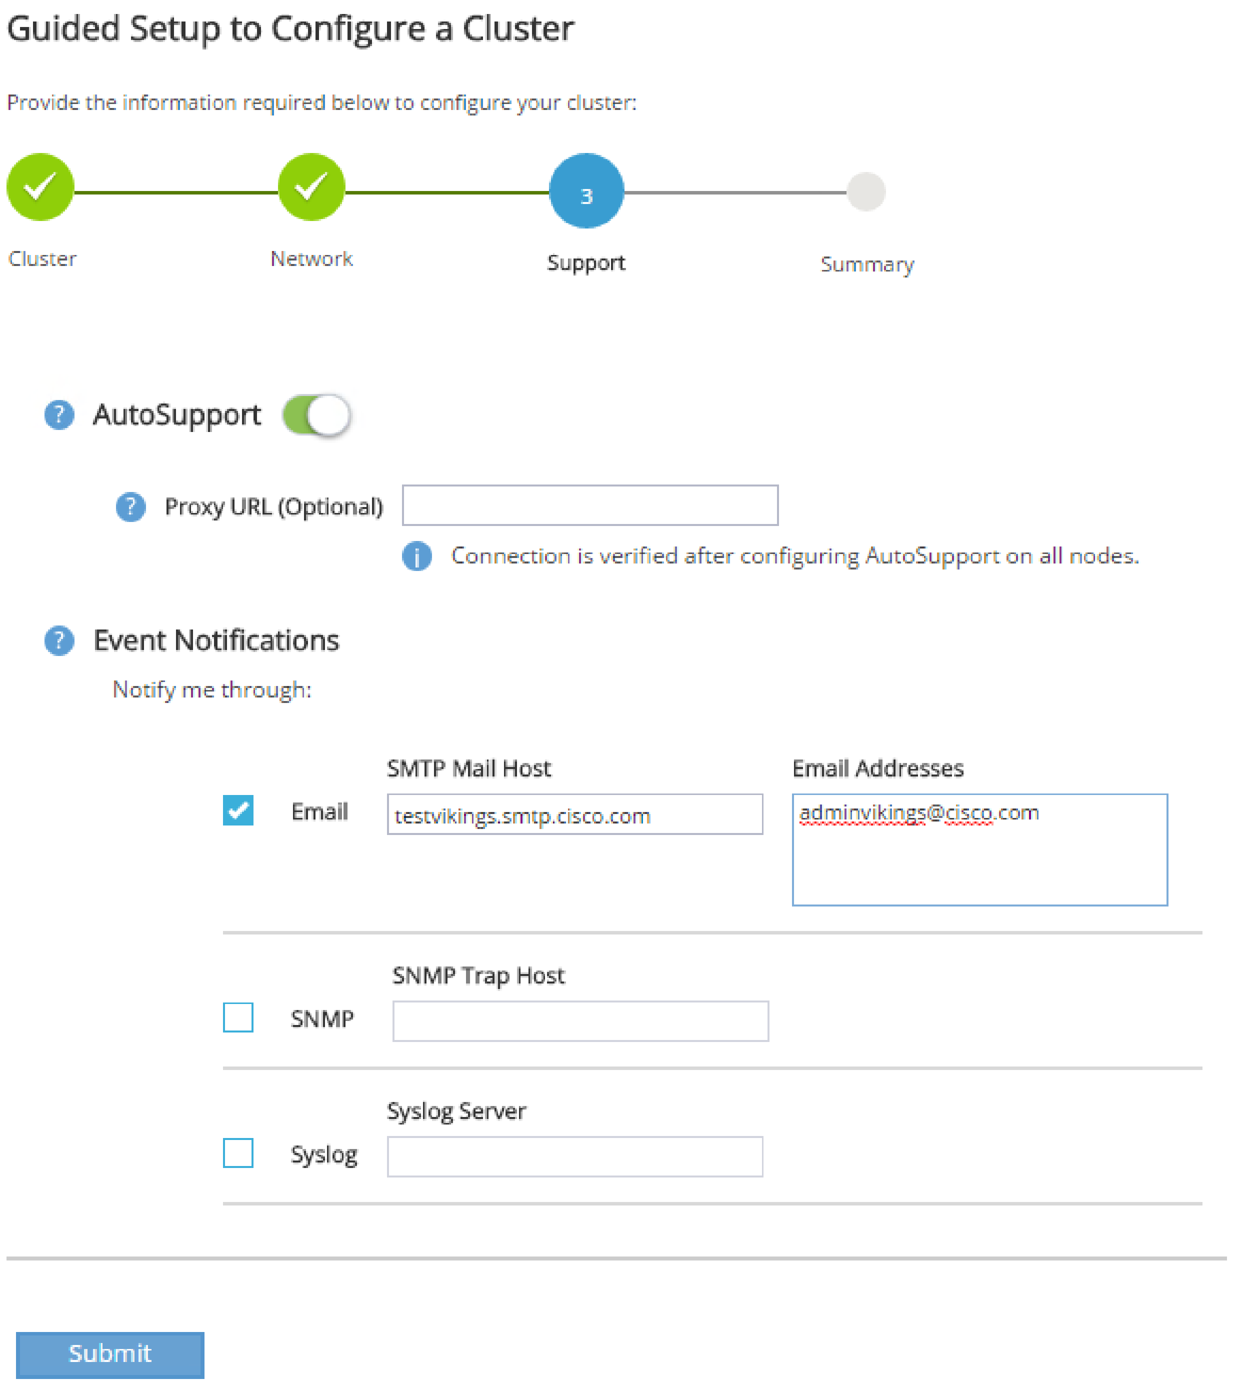

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

AutoSupport: ● On

Proxy URL (Optional):

Connection is verified after configuring AutoSupport on all nodes.

Event Notifications

Notify me through:

☑ Email	SMTP Mail Host: testvikings.smtp.cisco.com	Email Addresses: adminvikings@cisco.com
□ SNMP	SNMP Trap Host:	
□ Syslog	Syslog Server:	

10. Click Submit.

11. On the Summary page, review the configuration details.

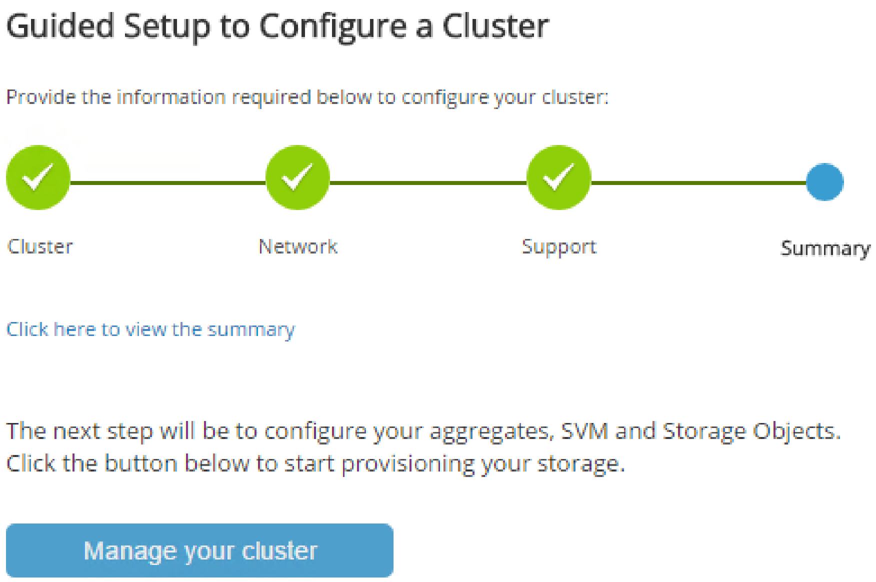

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Click here to view the summary

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

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, assume that it is on the same subnet.

Login to the Cluster

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

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

Zero All Spare Disks

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

```
disk zerospares.
```

Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an All Flash FAS configuration. Disk autoassign should have assigned one data partition to each node in an HA pair.

If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Onboard Unified Target Adapter 2 Port Personality

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

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

```
ucadmin show
                            Current  Current    Pending  Pending    Admin
Node            Adapter  Mode     Type       Mode     Type       Status
------------    -------  -------  ---------  -------  ---------  -----------
<st-node01>
                0e       fc       target     -        -          online
<st-node01>
                0f       fc       target     -        -          online
<st-node01>
                0g       cna      target     -        -          online
<st-node01>
                0h       cna      target     -        -          online
<st-node02>
                0e       fc       target     -        -          online
<st-node02>
                0f       fc       target     -        -          online
<st-node02>
                0g       cna      target     -        -          online
<st-node02>
                0h       cna      target     -        -          online
8 entries were displayed.
```

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

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

The ports must be offline to run this command. To take an adapter offline, run the `fcp adapter modify –node <home-node-of-the-port> -adapter <port-name> -state down` command. Ports must be converted in pairs (for example, `0e` and `0f`).

After conversion, a reboot is required. After reboot, bring the ports online by running `fcp adapter modify –node <home-node-of-the-port> -adapter <port-name> -state up`.

Set Auto-Revert on Cluster Management

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

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

Run the following command:

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

Set Up Management Broadcast Domain

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

```
broadcast-domain remove-ports -broadcast-domain Default -ports <st-node01>:e2a,<st-node01>:e2e,<st-
node02>:e2a,<st-node01>:e2e

broadcast-domain show
```

Set Up Service Processor Network Interface

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

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

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

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

Create Aggregates

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

To create new aggregates, run the following commands:

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

You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate. For all-flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions.

For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type. Start with five disks initially; you can add disks to an aggregate when additional storage is required. In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all but one remaining disk (spare) assigned to the controller.

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

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

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

Verify Storage Failover

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

1. Verify the status of the storage failover.

```
storage failover show
```

Both <st-node01> and <st-node02> must be able to perform a takeover. Continue with step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

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

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

4. Continue with step 5 if high availability is configured.

Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

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

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

Disable Flow Control on 10GbE and 40GbE Ports

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

1. Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port e0c,e0d,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <st-node02> -port e0c,e0d,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
```

```
Do you want to continue? {y|n}: y
network port show –fields flowcontrol-admin
```

Disable Unused FCoE Capability on CNA Ports

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

```
fcp adapter modify -node <st-node01> -adapter 0g -status-admin down
fcp adapter modify -node <st-node01> -adapter 0h -status-admin down
fcp adapter modify -node <st-node02> -adapter 0g -status-admin down
fcp adapter modify -node <st-node02> -adapter 0h -status-admin down
fcp adapter show –fields status-admin
```

Configure Network Time Protocol

If NTP was not configured during guided setup, it can be configured via the CLI as follows:

1. Set the time zone for the cluster.

```
timezone <timezone>
```

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

2. Set the date for the cluster.

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

The format for the date is <[Century][Year][Month][Day][Hour][Minute].[Second]> (for example, 201703231549.30).

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

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

Configure Simple Network Management Protocol

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

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

```
system snmp contact <snmp-contact>
system snmp location “<snmp-location>”
system 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.

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

Configure SNMPv1 Access

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

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

Configure AutoSupport

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

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

Enable Cisco Discovery Protocol and Link Layer 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
```

To enable the Link Layer Discovery Protocol (LLDP) on the NetApp storage controllers, run the following command:

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

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

Create Jumbo Frame MTU Broadcast Domains in ONTAP

To create a data broadcast domain with an MTU of 9000 for NFS on ONTAP, run the following command:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
```

If using iSCSI, create two iSCSI broadcast domains with an MTU of 9000, with the following command:

```
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Create Interface Groups

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

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

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

ifgrp show
```

Create VLANs

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

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

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

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

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

If using iSCSI, create two iSCSI VLANs and add them to the corresponding broadcast domains:

```
network port vlan create -node <st-node01> -vlan-name a0a-<Infra_iSCSI-A-VLAN>

network port vlan create -node <st-node02> -vlan-name a0a-<Infra_iSCSI-A-VLAN>

broadcast-domain add-ports -broadcast-domain iSCSI-A -ports <st-node01>:a0a-<Infra_iSCSI-A-VLAN>,<st-
node02>:a0a-<Infra_iSCSI-A-VLAN>

network port vlan create -node <st-node01> -vlan-name a0a-<Infra_iSCSI-B-VLAN>

network port vlan create -node <st-node02> -vlan-name a0a-<Infra_iSCSI-B-VLAN>

broadcast-domain add-ports -broadcast-domain iSCSI-B -ports <st-node01>:a0a-<Infra_iSCSI-B-VLAN>,<st-
node02>:a0a-<Infra_iSCSI-B-VLAN>
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

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

2. Remove the unused data protocols (CIFS and NDMP) from the SVM.

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

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

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

Create the NFS Service

You can enable and configure NFS servers on storage virtual machines (SVMs) with NetApp FlexVol® volumes to let NFS clients access files on your cluster. To do so, complete the following step:

Create the NFS service.

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

Enable VMware vStorage for NFS in ONTAP

To enable VMware vStorage for NFS in NetApp ONTAP, run the following command:

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

Create Load-Sharing Mirrors of SVM Root Volume

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

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

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

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

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

3. Create the mirroring relationships.

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

4. Initialize the mirroring relationship.

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

Create Block Protocol Service(s)

If the deployment is using FCP, create the FCP service on each SVM using the following command. This command also starts the FCP service and sets the worldwide name (WWN) for the SVM.

```
fcp create -vserver Infra-SVM

fcp show
```

If the deployment is using iSCSI, create the iSCSI service on each SVM using the following command. This command also starts the iSCSI service and sets the IQN for the SVM.

```
iscsi create -vserver Infra-SVM

iscsi show
```

The licenses for FCP and iSCSI must be installed before the services can be started. If the license(s) weren't installed during cluster setup, install them before this step.

Configure HTTPS Access

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

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

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

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

```
security certificate show
```

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

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

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

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

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

5. To obtain the values for the parameters required in step 5 (`<cert-ca>` and `<cert-serial>`), run the security certificate show command.

6. Enable each certificate that was just created by using the –server-enabled true and –client-enabled false parameters. Use TAB completion to aid in the completion of these commands:

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

7. Disable HTTP cluster management access:

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

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

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

```
set –privilege admin
vserver services web modify –name spi|ontapi|compat –vserver * -enabled true
```

Add Rules to Default Export Policy to Allow ESXi Host Access

To allow the ESXI hosts access to the NFS volumes, add rules to the default export policy by running the following command:

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -
clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false

vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 2 -protocol nfs -
clientmatch <<var_esxi_host2_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false

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

Create NetApp FlexVol Volumes

```
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 ESXi Host Boot LUNs

To create ESXi host boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun fpv-esxi-01 -size 15GB -ostype vmware -space-
reserve disabled

lun create -vserver Infra-SVM -volume esxi_boot -lun fpv-esxi-02 -size 15GB -ostype vmware -space-
reserve disabled
```

Create SAN LIFs

If using FCP, run the following commands to create four FC LIFs (two on each node):

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

network interface create -vserver Infra-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node
<st-node01> -home-port 0f –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node
<st-node02> -home-port 0e –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node
<st-node02> -home-port 0f –status-admin up
```

If using iSCSI, run the following commands to create four iSCSI LIFs (two on each node):

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<Infra_iSCSI-A-VLAN> –address <iscsi_lif01a_ip> -netmask
<iscsi_lif01a_mask> –status-admin up

network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<Infra_iSCSI-B-VLAN> –address <iscsi_lif01b_ip> -netmask
<iscsi_lif01b_mask> –status-admin up

network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -home-
node <st-node02> -home-port a0a-<Infra_iSCSI-A-VLAN> –address <iscsi_lif02a_ip> -netmask
<iscsi_lif02a_mask> –status-admin up

network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-
node <st-node02> -home-port a0a-<InfraiSCSI-B-VLAN> –address <iscsi_lif02b_ip> -netmask
<iscsi_lif02b_mask> –status-admin up
```

Create NFS LIFs

To create NFS LIFs, run the following commands:

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

network interface create -vserver Infra-SVM -lif nfs_lif02 -role data -data-protocol nfs -home-node
<st-node02> -home-port a0a-<infra-nfs-vlan-id> –address <node02-nfs_lif02-ip> -netmask <node02-
```

```
nfs_lif02-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface show
```

Add Infrastructure SVM Administrator

To add an infrastructure SVM administrator and an SVM administration LIF in the out-of-band management network, complete the following steps:

1. Create a network interface.

```
network interface create -vserver Infra-SVM -lif svm-mgmt -role data -data-protocol none -home-node
<st-node02> -home-port e0c -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -status-admin up –
failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```

The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

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

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

network route show
```

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

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

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

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

Server Configuration

Cisco UCS Base Configuration

This FlexPod deployment describes the configuration steps for the Cisco UCS 6332-16UP Fabric Interconnects (FI) in a design that will support iSCSI boot to the NetApp AFF through the Cisco ACI Fabric. An alternative Fibre Channel (FC) boot delta configuration is detailed in Appendix - FC Solution. If FC boot is desired, execute the following procedure not executing iSCSI-related steps and then execute the procedure in the appendix.

Table 1 Lab Validation Infrastructure (FPV-Foundation) Tenant Configuration

EPG	Storage VLAN	UCS VLAN	External VLAN	Subnet / Gateway	Bridge Domain
IB-MGMT	N/A	419/DVS	N/A	10.1.118.1/24	BD-FP-common-Core-Services
Core-Services	N/A	318/DVS	118	10.1.118.1/254 10.1.118.254/24	BD-FP-common-Core-Services
SVM-MGMT	219	N/A	N/A	172.16.254.6/29	BD-FPV-Foundation-Internal
iSCSI-A	3010	3110	N/A	192.168.10.0/24 – L2	BD- FPV-Foundation-iSCSI-A
iSCSI-B	3020	3120	N/A	192.168.20.0/24 – L2	BD- FPV-Foundation-iSCSI-B
NFS	3050	3150	N/A	192.168.50.0/24 – L2	BD- FPV-Foundation-NFS
vMotion	N/A	3000/DVS	N/A	192.168.100.0/24 – L2	BD-Internal
Native	N/A	2	N/A	N/A	N/A
VMware vDS Pool	N/A	1100-1199	N/A	Varies	Varies
ACI System VLAN for AVE	N/A	4093	N/A	Varies	Varies

Perform Initial Setup

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

Cisco UCS Fabric Interconnect A

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

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

```
Enter the configuration method: gui

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

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

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

2. Using a supported web browser, connect to http://`<ucsa-mgmt-ip>`, accept the security prompts, and click the ‘Express Setup’ link.

3. Select Initial Setup and click Submit.

4. Select Enable clustering, Fabric A, and IPv4.

5. Fill in the Virtual IP Address with the UCS cluster IP.

6. Completely fill in the System setup section. For system name, use the overall UCS system name. For the Mgmt IP Address, use `<ucsa-mgmt-ip>`.

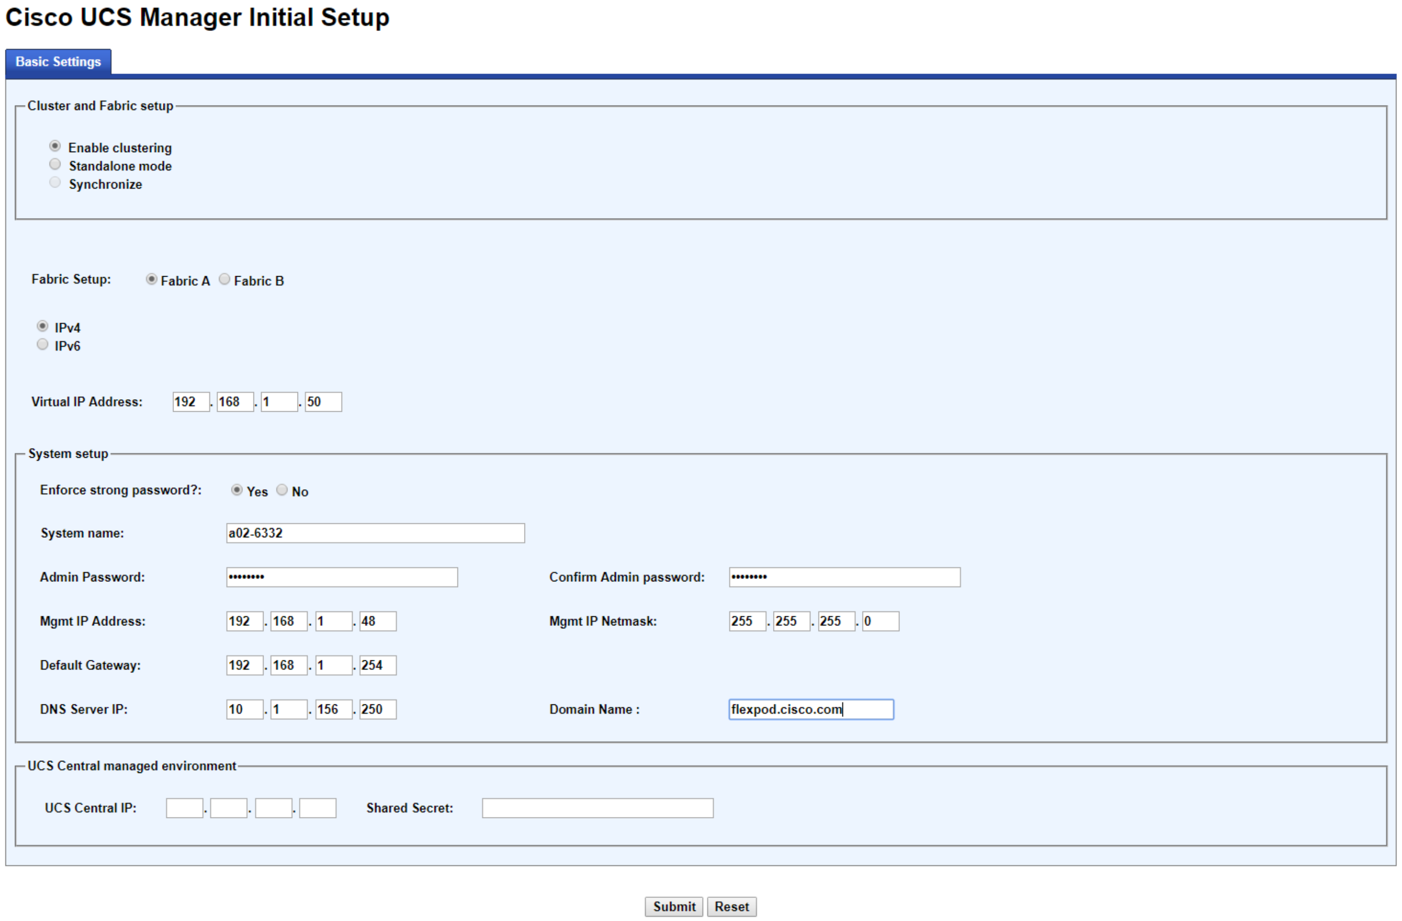

Cisco UCS Manager Initial Setup

Cluster and Fabric setup

● Enable clustering
○ Standalone mode
○ Synchronize

Fabric Setup: ● Fabric A ○ Fabric B

● IPv4
○ IPv6

Virtual IP Address: 192 . 168 . 1 . 50

System setup

Enforce strong password?: ● Yes ○ No

System name: a02-6332

Admin Password: ••••••••

Confirm Admin password: ••••••••

Mgmt IP Address: 192 . 168 . 1 . 48

Mgmt IP Netmask: 255 . 255 . 255 . 0

Default Gateway: 192 . 168 . 1 . 254

DNS Server IP: 10 . 1 . 156 . 250

Domain Name : flexpod.cisco.com

UCS Central managed environment

UCS Central IP: . . .

Shared Secret:

7. Click **Submit**.

Cisco UCS Fabric Interconnect B

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

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

```
Enter the configuration method: gui

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

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

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

2. Using a supported web browser, connect to http://`<ucsb-mgmt-ip>`, accept the security prompts, and click the ‘Express Setup’ link.

3. Under System setup, enter the Admin Password entered above and click **Submit**.

4. Enter `<ucsb-mgmt-ip>` for the Mgmt IP Address and click **Submit**.

Cisco UCS Setup

Log in to Cisco UCS Manager

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

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

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

2. Click the Launch UCS Manager link under HTML to launch Cisco UCS Manager.

3. If prompted to accept security certificates, accept as necessary.

4. When prompted, enter admin as the user name and enter the administrative password.

5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 3.2(3a)

This document assumes the use of Cisco UCS 3.2(3a) release. The testing for this CVD was initially completed using Cisco UCS 3.2(2d), prior to the release of UCS 3.2(3a). However, when Spectre and Meltdown patches were released in Cisco UCS 3.2(3a), the testbed was upgraded and a limited set of tests were run to ensure consistency with previous validation. To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.2(3a), refer to Cisco UCS Manager Install and Upgrade Guides.

Anonymous Reporting

To create anonymous reporting, complete the following step:

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

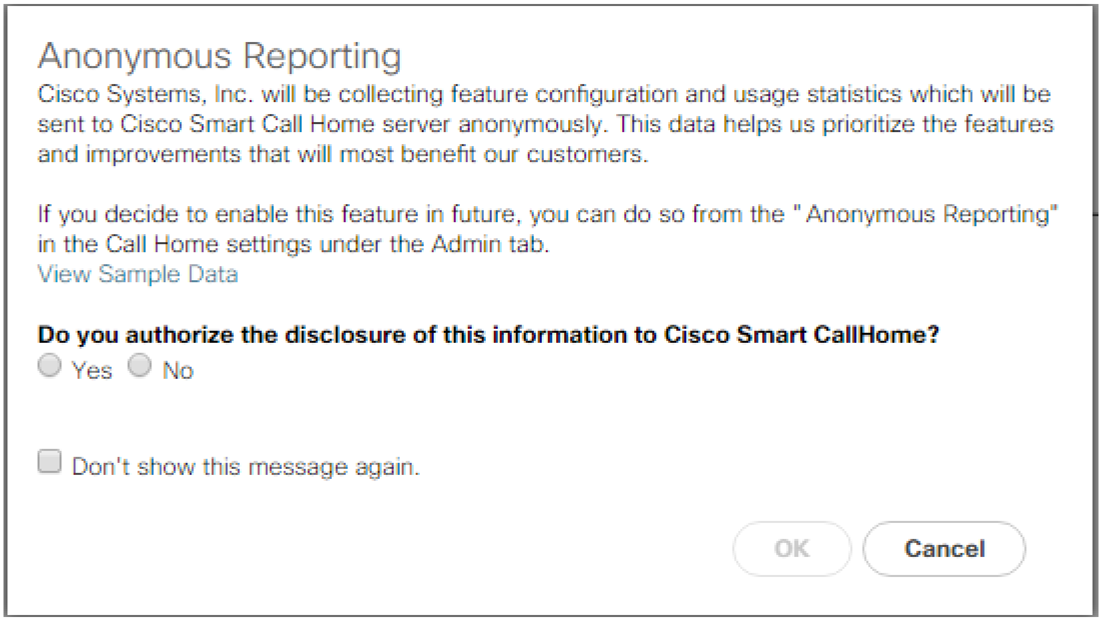

Anonymous Reporting

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

If you decide to enable this feature in future, you can do so from the "Anonymous Reporting" in the Call Home settings under the Admin tab.

View Sample Data

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

○ Yes ○ No

□ Don't show this message again.

Configure Cisco UCS Call Home

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

1. In Cisco UCS Manager, click the Admin icon on the left.

2. Select All > Communication Management > Call Home.

3. Change the State to On.

4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

Add Block of IP Addresses for KVM Access

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

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

2. Expand Pools > root > IP Pools.

3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.

4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information.

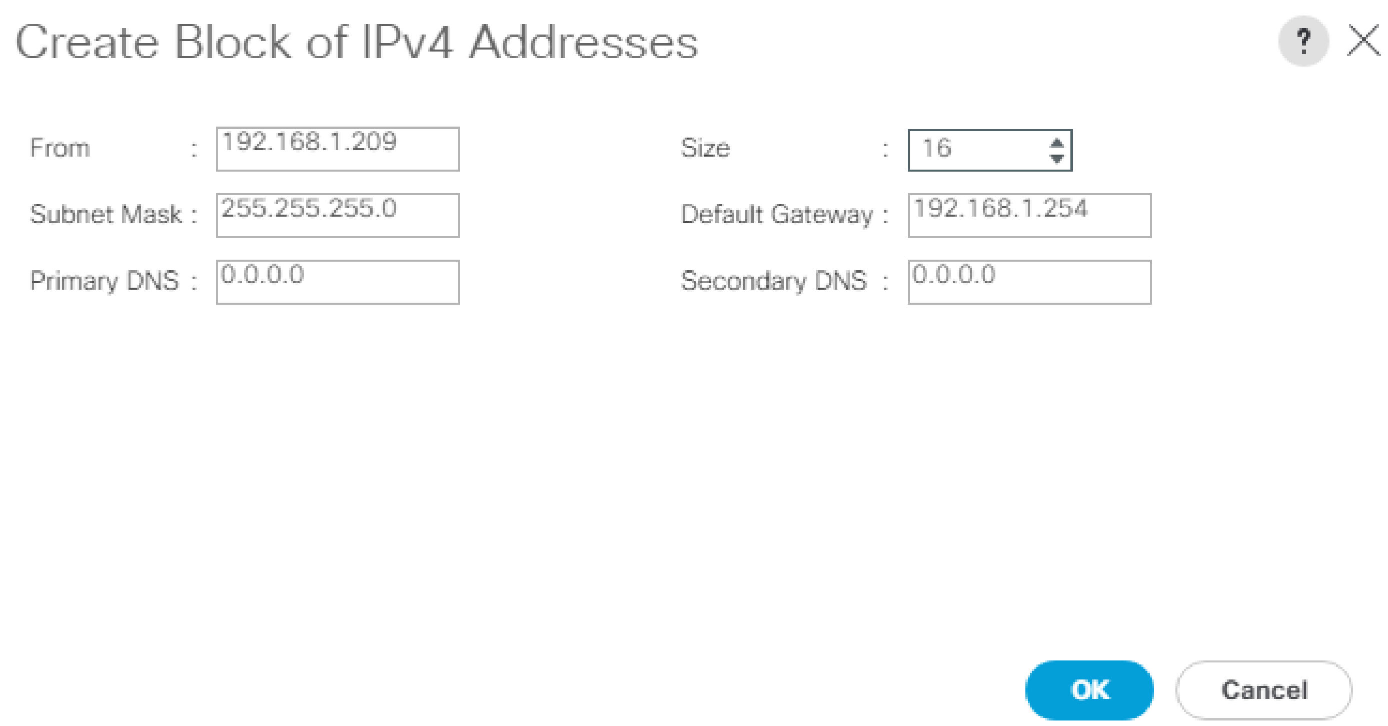

Create Block of IPv4 Addresses

From :	192.168.1.209	Size :	16
Subnet Mask :	255.255.255.0	Default Gateway :	192.168.1.254
Primary DNS :	0.0.0.0	Secondary DNS :	0.0.0.0

5. Click OK to create the block.

6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, click the Admin icon on the left.

2. Expand All > Time Zone Management.

3. Select Timezone.

4. In the Properties pane, select the appropriate time zone in the Timezone menu.

5. Click Save Changes, and then click OK.

6. Click Add NTP Server.

7. Enter <leaf-a-mgmt-ip> and click OK. Click OK on the confirmation.

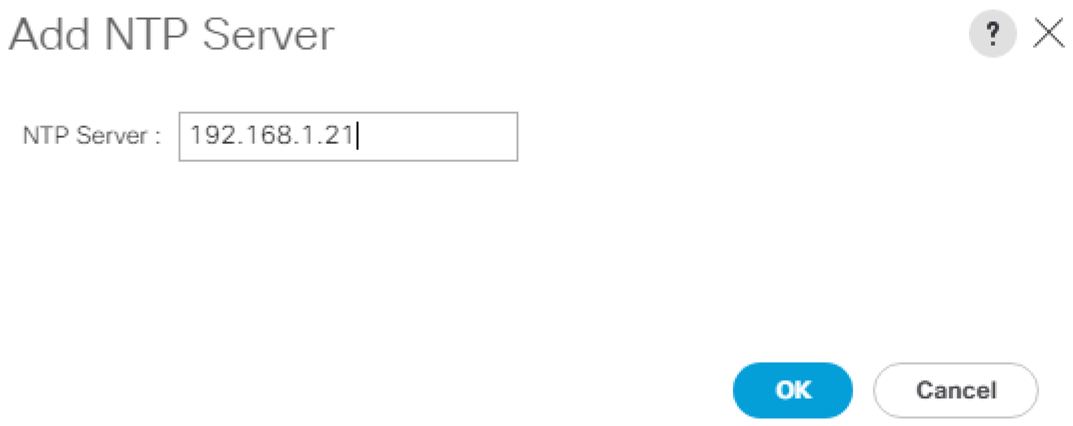

Add NTP Server

NTP Server : 192.168.1.21

8. Add <leaf-b-mgmt-ip> for the second NTP server.

Edit Policy to Automatically Discover Server Ports

If the UCS Port Auto-Discovery Policy is enabled, server ports will be discovered automatically. To enable the Port Auto-Discovery Policy, complete the following steps:

1. In Cisco UCS Manager, click the Equipment icon on the left and select Equipment in the second list.

2. In the right pane, click the Policies tab.

3. Under Policies, select the Port Auto-Discovery Policy tab.

4. Under Properties, set Auto Configure Server Port to Enabled.

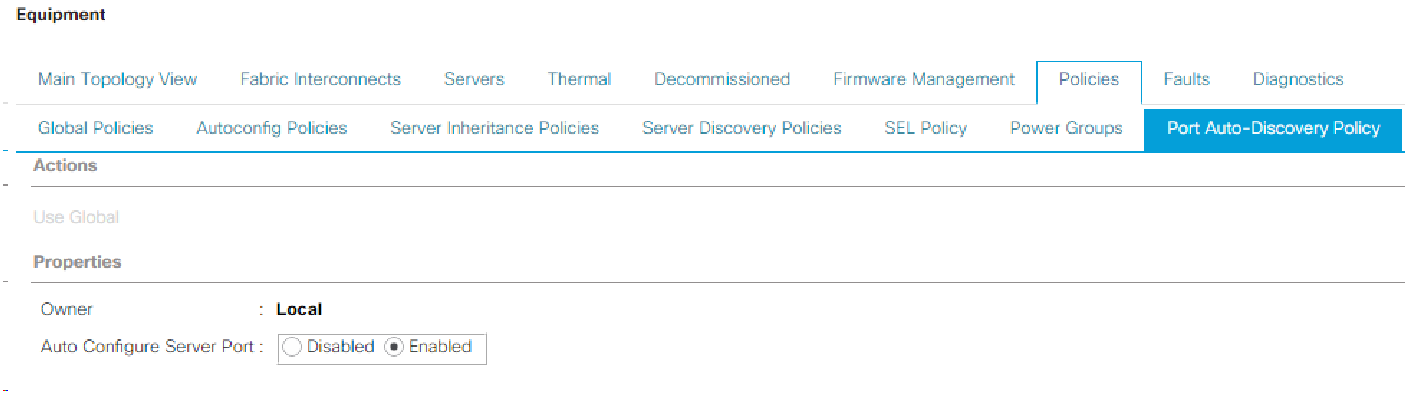

Owner : Local

Auto Configure Server Port : ○ Disabled ● Enabled

5. Click Save Changes.

6. Click OK.

Edit Chassis Discovery Policy

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

1. In Cisco UCS Manager, click the Equipment icon on the left and select Equipment in the second list.

2. In the right pane, click the Policies tab.

3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.

4. Set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G. For Backplane Speed Preference, select 40G if you have servers with VIC1340s and Port Expander cards. Otherwise, select 4x10G.

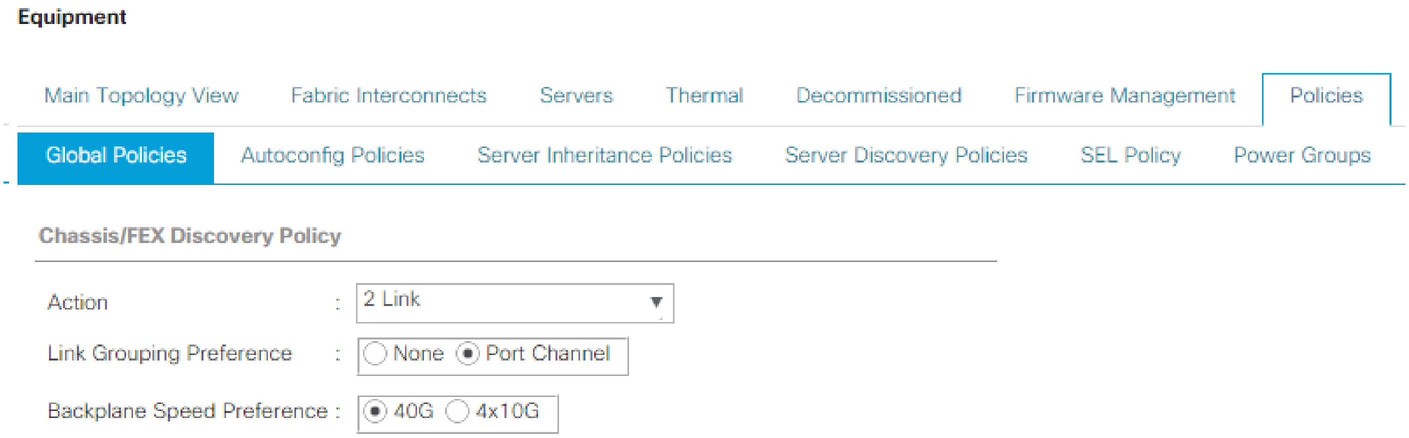

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : ○ None ● Port Channel

Backplane Speed Preference : ● 40G ○ 4x10G

5. Click Save Changes.

6. Click OK.

Verify Server and Enable Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment icon on the left.

2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.

3. Expand and select Ethernet Ports.

4. On the right, verify that the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers are configured as Server ports. If any Server ports are not configured correctly, right-click them, and select “Configure as Server Port.” Click Yes to confirm server ports and click OK.

In lab testing, for C220M4 servers with VIC 1385 PCIE cards, it has been necessary at times to manually configure Server ports.

5. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.

6. Select the ports that are connected to the Cisco Nexus 9332 switches, right-click them, and select Configure as Uplink Port.

The last 6 ports (ALE) of the Cisco UCS 6332 and UCS 6332-16UP FIs require the use of active (optical) or AOC cables when connected to a Nexus 9332. It may also be necessary to remove and reinsert these cables on the switch end to get them to come up the first time.

7. Click Yes to confirm uplink ports and click OK.
8. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
9. Expand and select Ethernet Ports.
10. On the right, verify that the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers are configured as Server ports. If any Server ports are not configured correctly, right-click them, and select "Configure as Server Port." Click Yes to confirm server ports and click OK.

In lab testing, for Cisco C220 M4 servers with VIC 1385 PCIE cards, it has been necessary at times to manually configure Server ports.

11. Verify that the ports connected to the chassis, Cisco UCS C-series servers and Cisco FEX are now configured as server ports.
12. Select the ports that are connected to the Cisco Nexus 9332 switches, right-click them, and select Configure as Uplink Port.

The last 6 ports (ALE) of the Cisco UCS 6332 and Cisco UCS 6332-16UP FIs require the use of active (optical) or AOC cables when connected to a Nexus 9332. It may also be necessary to remove and reinsert these cables on the switch end to get them to come up the first time.

13. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

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

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

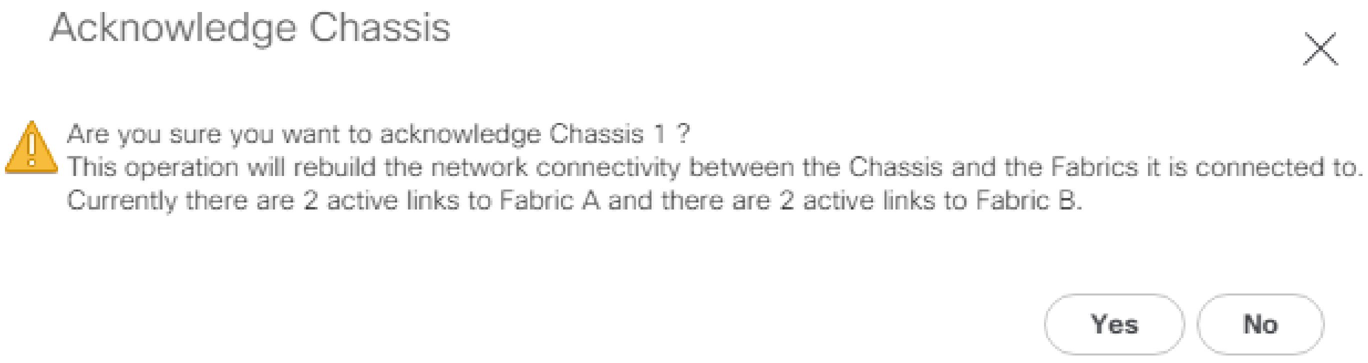

Acknowledge Chassis

Are you sure you want to acknowledge Chassis 1 ?
This operation will rebuild the network connectivity between the Chassis and the Fabrics it is connected to.
Currently there are 2 active links to Fabric A and there are 2 active links to Fabric B.

4. Click Yes and then click OK to complete acknowledging the chassis.
5. If Nexus 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.

Re-Acknowledge Any Inaccessible Cisco UCS C-Series Servers

If any Cisco UCS C-Series servers show an Inaccessible Status, complete the following steps:

1. In Cisco UCS Manager, click the Equipment icon on the left.

2. Under Equipment > Rack Mounts, expand Servers.

3. If any of the servers have a status of Inaccessible, right-click the server and select Server Maintenance. Select Re-Acknowledge and click OK. Click Yes and OK. The server should then be Discovered properly.

Create Uplink Port Channels to Cisco Nexus 9332 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 icon on the left.

In this procedure, two port channels are created: one from fabric A to both Cisco Nexus 9332 switches and one from fabric B to both Cisco Nexus 9332 switches.

2. Under LAN > LAN Cloud, expand the Fabric A tree.

3. Right-click Port Channels.

4. Select Create Port Channel.

5. Enter `139` as the unique ID of the port channel.

6. Enter `Po139-ACI` as the name of the port channel.

7. Click Next.

8. Select the ports connected to the Nexus switches to be added to the port channel:

   a. Click >> to add the ports to the port channel.

   b. Click Finish to create the port channel.

   c. Click OK.

9. Expand Port Channels and select Port-Channel 139. Since the vPC has already been configured in the ACI fabric, this port channel should come up. Note that it may take a few minutes for the port channel to come up.

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

11. Right-click Port Channels.

12. Select Create Port Channel.

13. Enter `140` as the unique ID of the port channel.

14. Enter `Po140-ACI` as the name of the port channel.

15. Click Next.

16. Select the ports connected to the Nexus switches to be added to the port channel:

    a. Click >> to add the ports to the port channel.
    b. Click Finish to create the port channel.
    c. Click OK.

17. Expand Port Channels and select Port-Channel 140. Since the vPC has already been configured in the ACI fabric, this port channel should come up. Note that it may take a few minutes for the port channel to come up.

Create an Organization for this FlexPod

To create a UCS Organization to contain unique parameters for this particular FlexPod, complete the following steps on Cisco UCS Manager.

1. Select the Servers icon on the left.

2. Under Servers > Service-Profiles > root, right-click Sub-Organizations and select Create Organization.

3. Name the Organization `FPV-FlexPod`, enter an optional Description, and click OK.

4. Click OK for the confirmation.

Create an IQN Pool for iSCSI Boot

To configure the necessary IQN pool for the Cisco UCS environment, complete the following steps on Cisco UCS Manager.

1. Select the SAN icon 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-Pool` for the name of the IQN pool.

6. Optional: Enter a description for the IQN pool.

7. Enter `iqn.2010-11.com.flexpod` for the Prefix.

8. Select **Sequential** for Assignment Order.

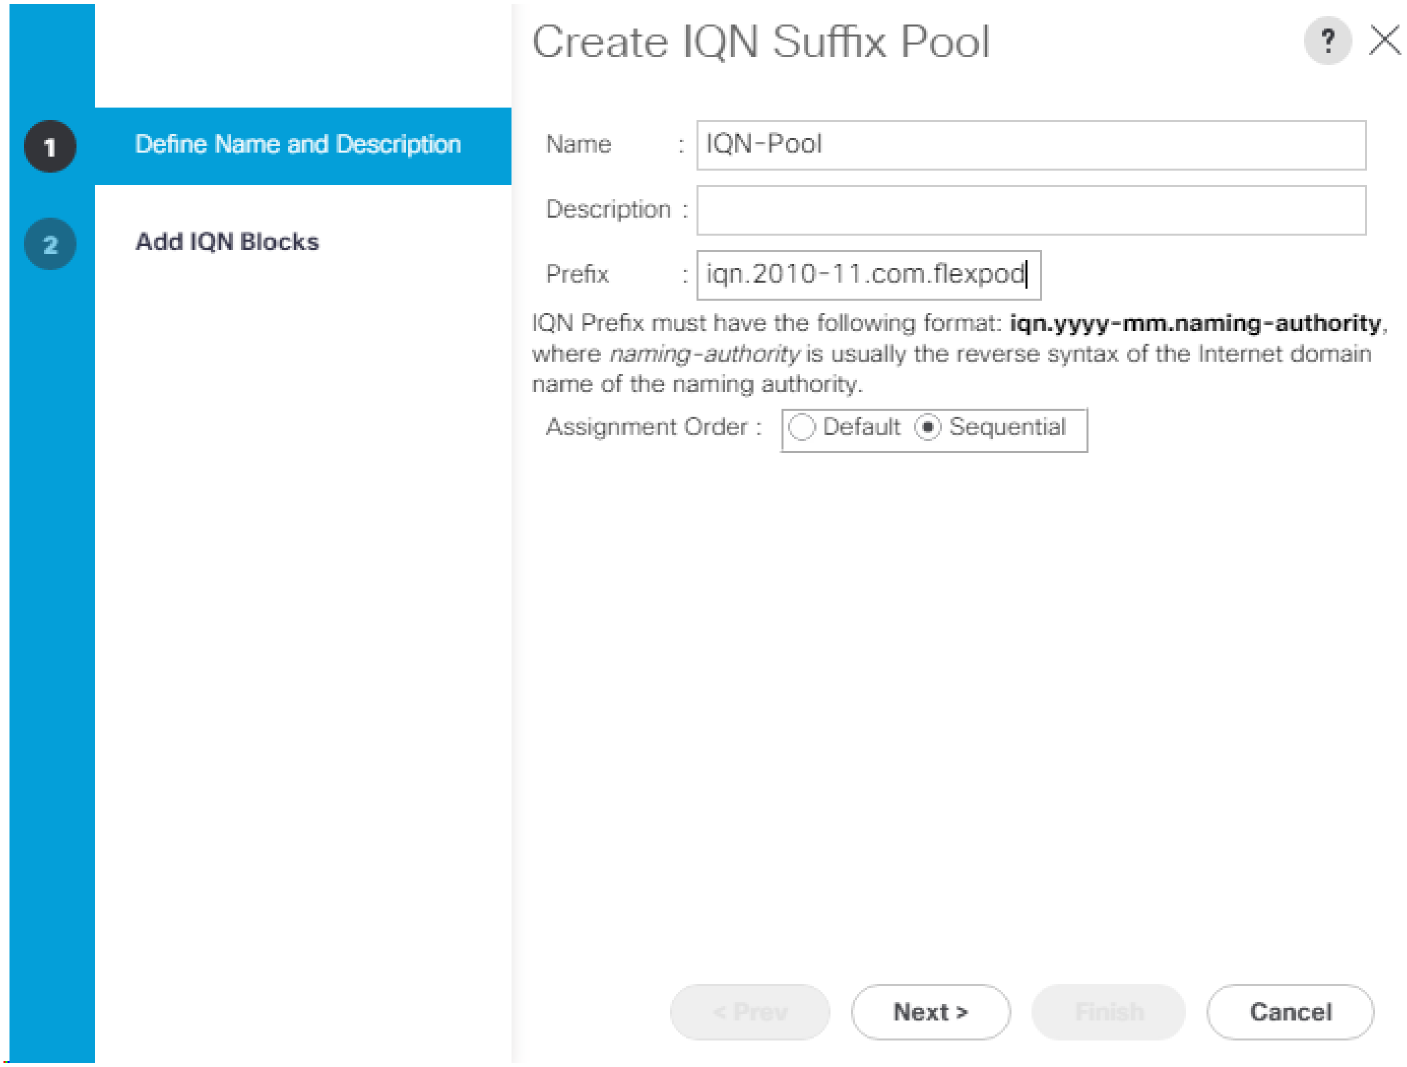

Create IQN Suffix Pool

Name : IQN-Pool

Description :

Prefix : iqn.2010-11.com.flexpod

IQN Prefix must have the following format: **iqn.yyyy-mm.naming-authority**, where *naming-authority* is usually the reverse syntax of the Internet domain name of the naming authority.

Assignment Order : ○ Default ● Sequential

9. Click Next.

10. Click Add.

11. Enter a name to identify the individual UCS host for the Suffix.

12. Enter 1 for the From field.

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

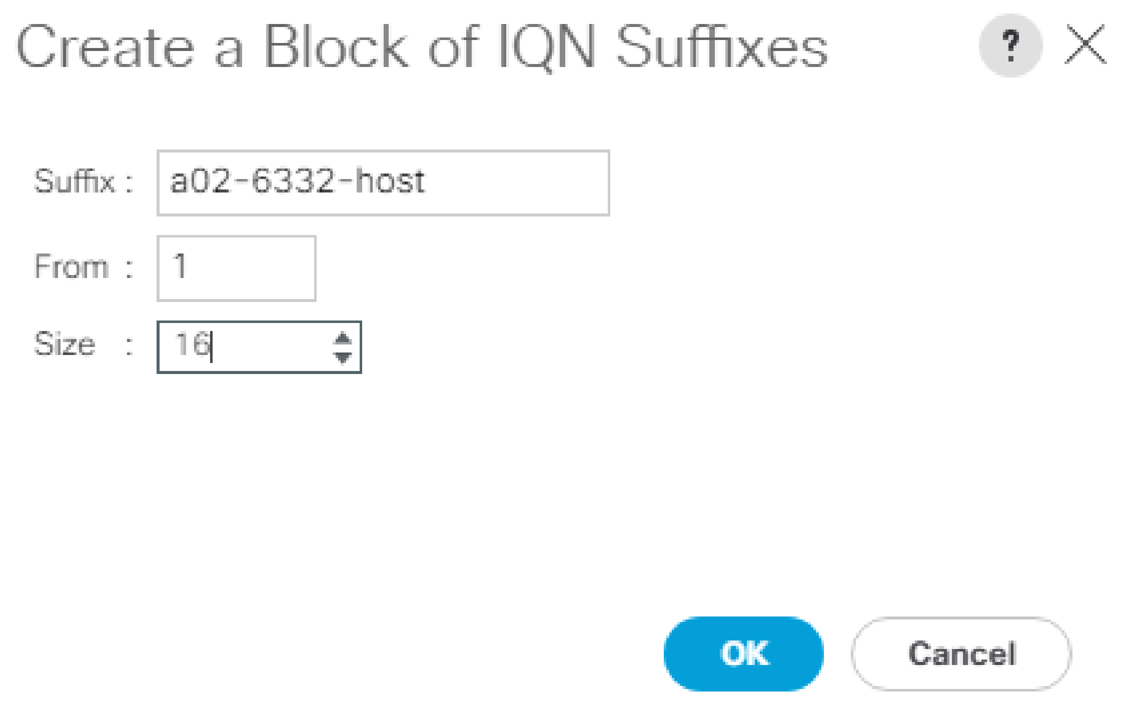

Create a Block of IQN Suffixes

Suffix : a02-6332-host

From : 1

Size : 16

14. Click OK.

15. Click Finish and OK to complete creating the IQN pool.

Create iSCSI Boot IP Address Pools

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

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

2. Select and expand Pools > root > Sub-Organizations > FPV-FlexPod.

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

4. Right-click IP Pools under the FPV-FlexPod organization.

5. Select Create IP Pool to create the IP pool.

6. Enter `iSCSI-IP-Pool-A` as the name of the first IP pool.

7. Optional: Enter a description for the IP pool.

8. Select **Sequential** for Assignment Order.

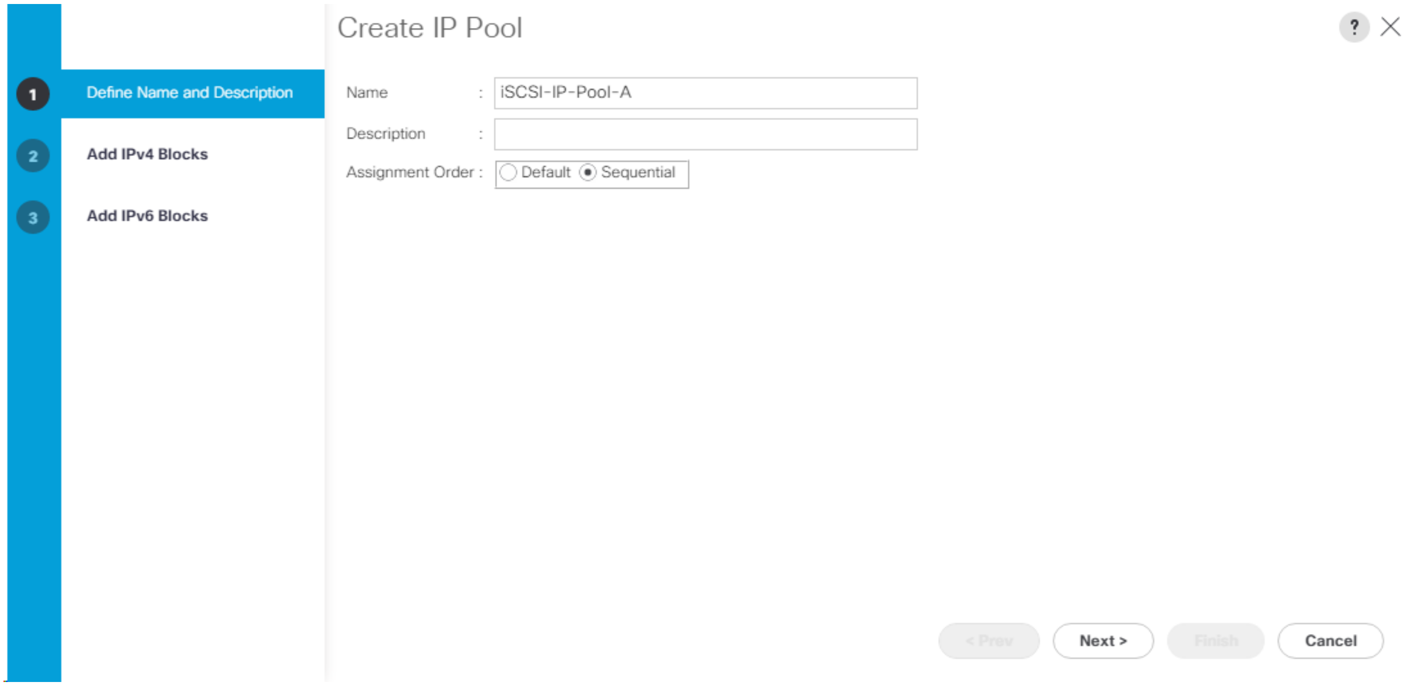

Create IP Pool

Name : iSCSI-IP-Pool-A

Description :

Assignment Order : ○ Default ● Sequential

9. Click Next.

10. Click Add to add a Block of IPs to the pool.

11. Specify a starting IP address and subnet mask in the subnet <192.168.10.0/24> for iSCSI boot on Fabric A. It is not necessary to specify the Default Gateway or DNS server addresses.

12. Specify a size for the IP pool that is sufficient to support the available blade or server resources.

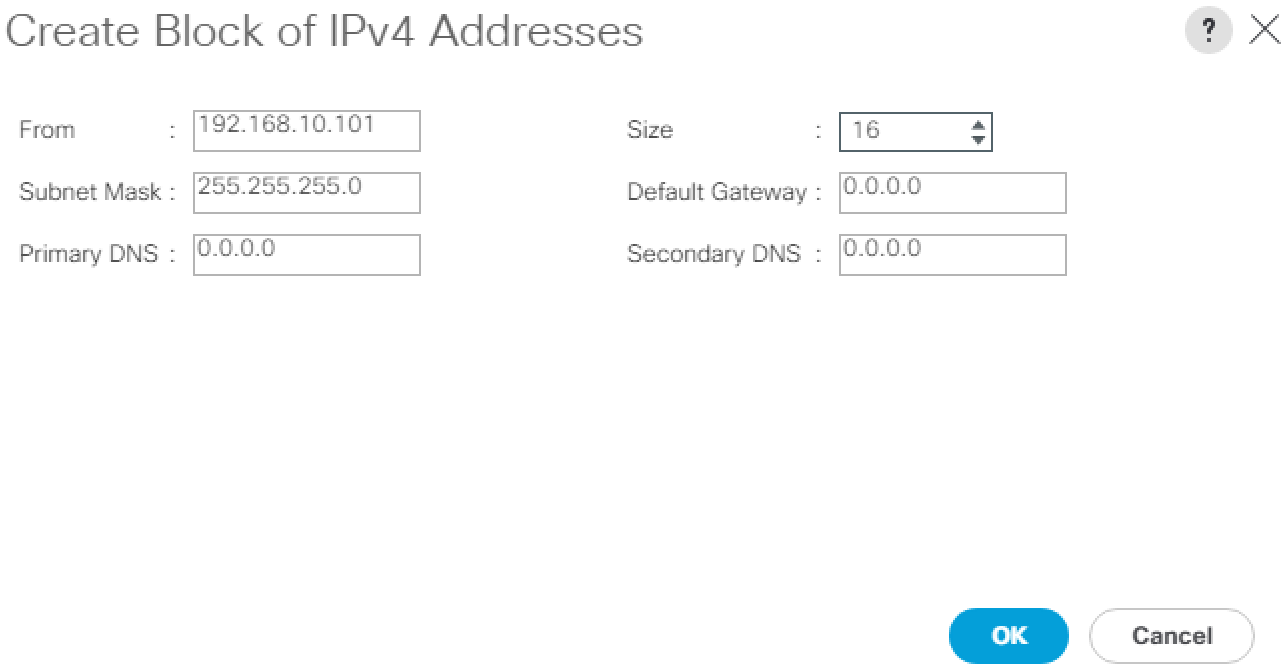

Create Block of IPv4 Addresses

From : 192.168.10.101

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

13. Click OK.

14. Click Next.

15. Click Finish.

16. In the confirmation message, click OK.

17. Right-click IP Pools under the FPV-FlexPod organization.

18. Select Create IP Pool to create the IP pool.

19. Enter `iSCSI-IP-Pool-B` as the name of the second IP pool.

20. Optional: Enter a description for the IP pool.

21. Select **Sequential** for Assignment Order

22. Click Next.

23. Click Add to add a Block of IPs to the pool.

24. Specify a starting IP address and subnet mask in the subnet <192.168.20.0/24> for iSCSI boot on Fabric B. It is not necessary to specify the Default Gateway or DNS server addresses.

25. Specify a size for the IP pool that is sufficient to support the available blade or server resources.

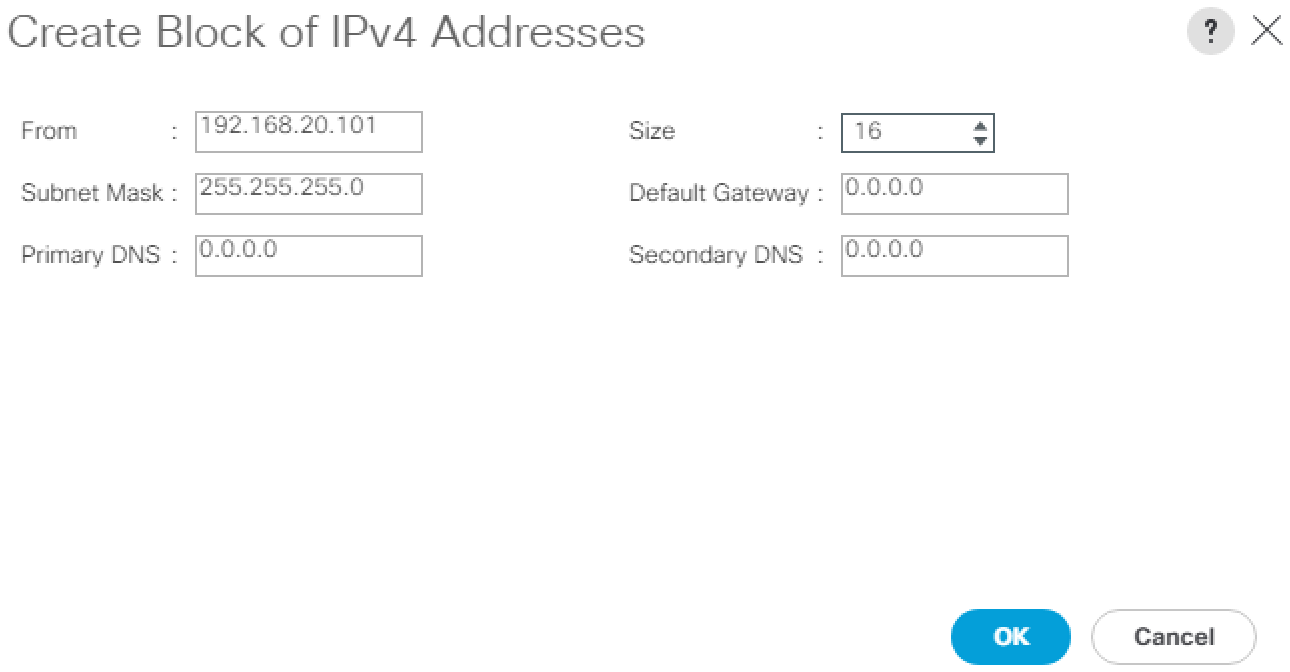

Create Block of IPv4 Addresses

From : 192.168.20.101

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

26. Click OK.

27. Click Next.

28. Click Finish.

29. In the confirmation message, 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 icon on the left.

2. Select Pools > root.

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

3. Right-click MAC Pools under the root organization.

4. Select Create MAC Pool to create the MAC address pool.

5. Enter `MAC-Pool-A` as the name of the MAC pool.

6. Optional: Enter a description for the MAC pool.

7. Select **Sequential** as the option for Assignment Order.

8. Click Next.

9. Click Add.

10. Specify a starting MAC address.

For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have also embedded the cabinet number (A2) information giving us 00:25:B5:A2:0A:00 as our first MAC address.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources assuming that multiple vNICs can be configured on each server.

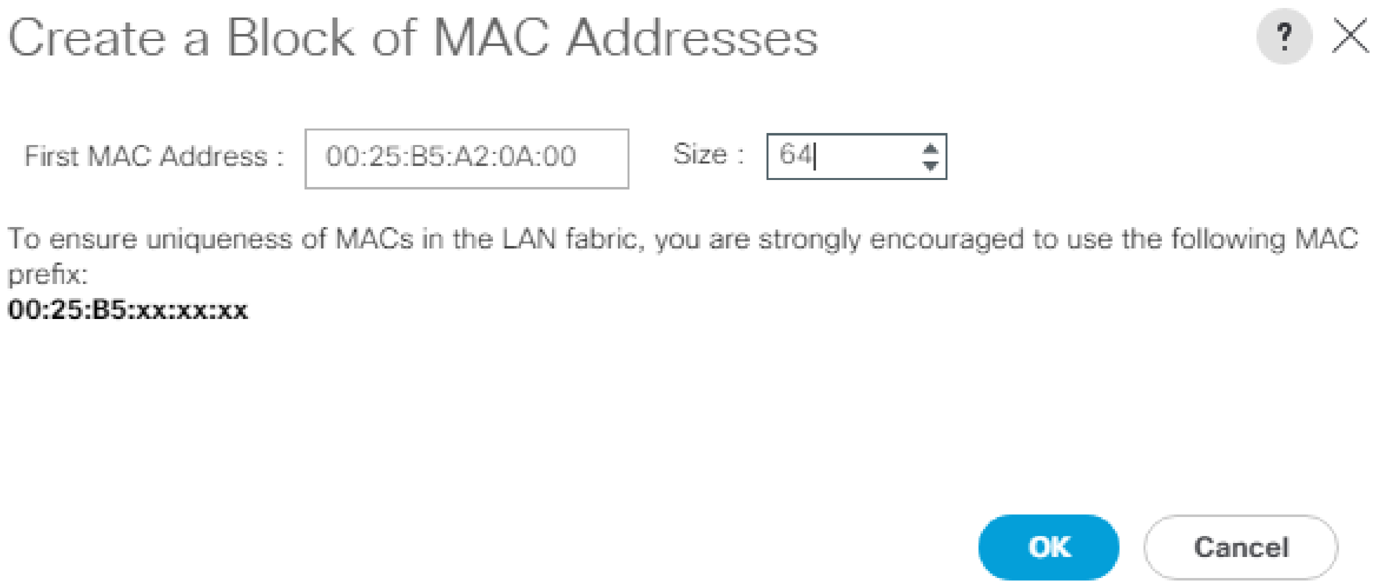

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:A2:0A:00

Size : 64

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

12. Click OK.

13. Click Finish.

14. In the confirmation message, click OK.

15. Right-click MAC Pools under the root organization.

16. Select Create MAC Pool to create the MAC address pool.

17. Enter `MAC-Pool-B` as the name of the MAC pool.

18. Optional: Enter a description for the MAC pool.

19. Select **Sequential** as the option for Assignment Order.

20. Click Next.

21. Click Add.

22. Specify a starting MAC address.

For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have also embedded the cabinet number (A2) information giving us 00:25:B5:A2:0A:00 as our first MAC address.

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

24. Click OK.

25. Click Finish.

26. In the confirmation message, click OK.

Create 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 icon on the left.

2. Select Pools > root.

3. Right-click UUID Suffix Pools.

4. Select Create UUID Suffix Pool.

5. Enter `UUID-Pool` as the name of the UUID suffix pool.

6. Optional: Enter a description for the UUID suffix pool.

7. Keep the prefix at the derived option.

8. Select **Sequential** for the Assignment Order.

9. Click Next.

10. Click Add to add a block of UUIDs.

11. Keep the From field at the default setting. Optionally, specify identifiers such as UCS location.

12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

13. Click OK.

14. Click Finish.

15. Click OK.

Create Server Pool

To configure the necessary server pool for the VMware management environment, complete the following steps:

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

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

2. Expand Pools > root > Sub-Organizations > FPV-FlexPod.

3. Right-click Server Pools under the FPV-FlexPod Organization.

4. Select Create Server Pool.

5. Enter FPV-MGMT-Pool as the name of the server pool.

6. Optional: Enter a description for the server pool.

7. Click Next.

8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the FPV-MGMT-Pool server pool.

9. Click Finish.

10. Click OK.

Create VLANs

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

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

In this procedure, eight unique and 100 sequential VLANs are created. See Table 2

Note that the ACI System VLAN is being added here for the set of Cisco VIC vNICs that will be used for the APIC-Integrated Virtual Switch. Although this VLAN is not necessary for the VMware vDS, it is being put in for future usage with the Cisco ACI Virtual Edge (AVE) virtual switch.

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 <2>.

8. Keep the Sharing Type as None.

9. Click OK and then click OK again.

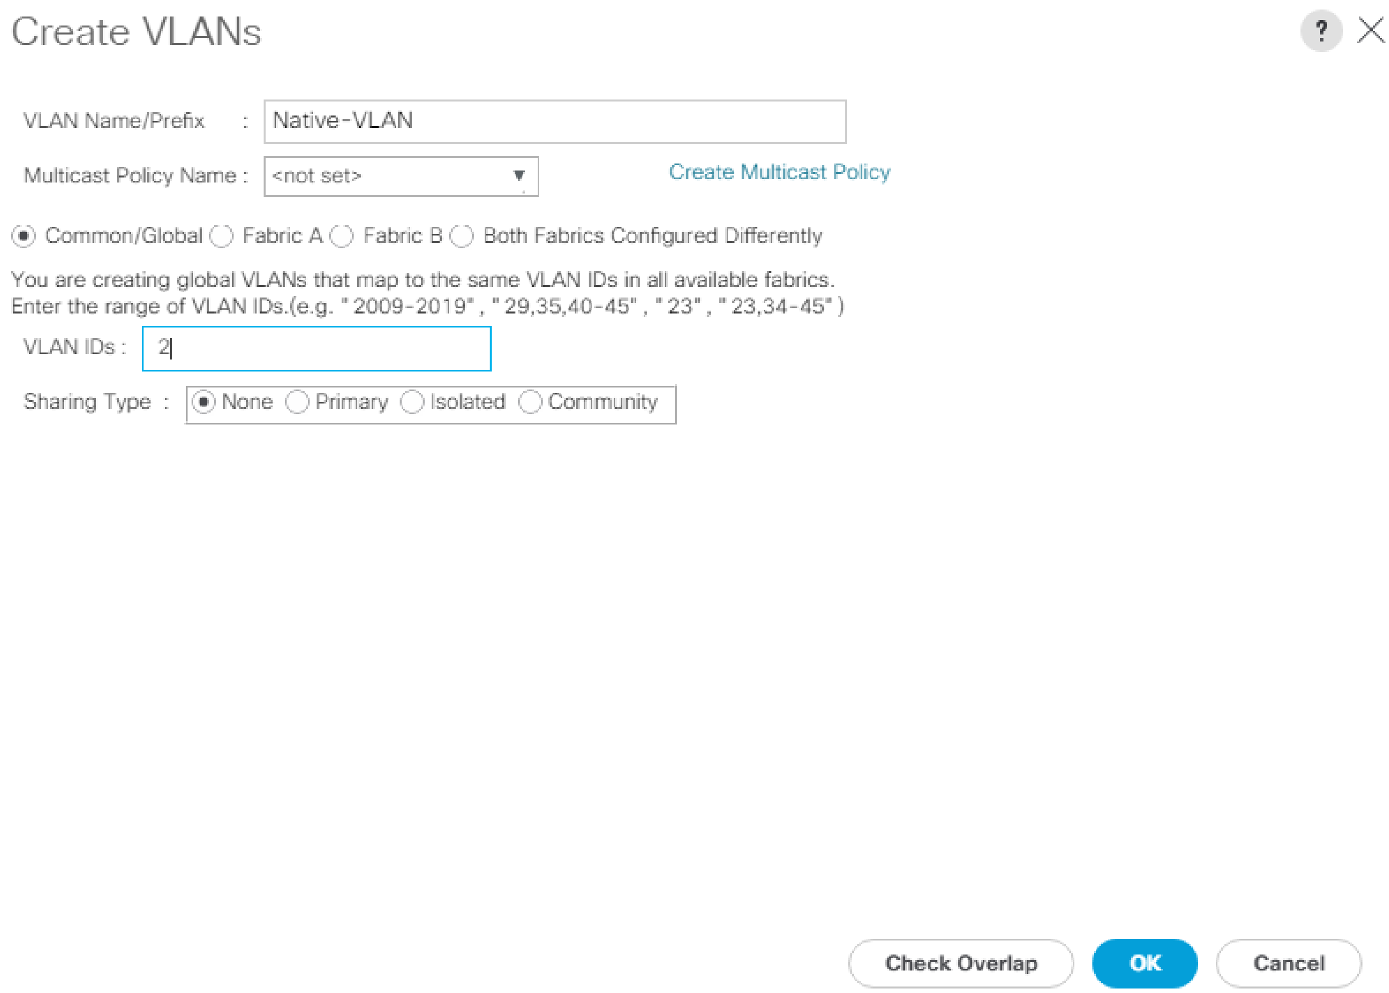

Create VLANs

VLAN Name/Prefix : Native-VLAN

Multicast Policy Name : <not set>

● 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 : 2

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select Set as Native VLAN.

11. Click Yes and then click OK.

12. Right-click VLANs.

13. Select Create VLANs

14. Enter `FP-Core-Services` as the name of the VLAN to be used for management traffic.

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

16. Enter the UCS Core-Services VLAN ID <318>.

17. Keep the Sharing Type as None.

18. Click OK, and then click OK again.

19. Right-click VLANs.

20. Select Create VLANs.

21. Enter `FPV-IB-MGMT` as the name of the VLAN to be used for management traffic.

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

23. Enter the UCS In-Band management VLAN ID <419>.

24. Keep the Sharing Type as None.

25. Click OK, and then click OK again.

26. Right-click VLANs.

27. Select Create VLANs.

28. Enter `FPV-Foundation-NFS` as the name of the VLAN to be used for infrastructure NFS.

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

30. Enter the UCS Infrastructure NFS VLAN ID <3150>.

31. Keep the Sharing Type as None.

32. Click OK, and then click OK again.

33. Right-click VLANs.

34. Select Create VLANs.

35. Enter `FPV-Foundation-iSCSI-A` as the name of the VLAN to be used for UCS Fabric A iSCSI boot.

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

37. Enter the UCS Fabric A iSCSI boot VLAN ID <3110>.

38. Keep the Sharing Type as None.

39. Click OK, and then click OK again.

40. Right-click VLANs.

41. Select Create VLANs.

42. Enter `FPV-Foundation-iSCSI-B` as the name of the VLAN to be used for UCS Fabric B iSCSI boot.

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

44. Enter the UCS Fabric B iSCSI boot VLAN ID <3120>.

45. Keep the Sharing Type as None.

46. Click OK, and then click OK again.

47. Right-click VLANs.

48. Select Create VLANs.

49. Enter `FPV-vMotion` as the name of the VLAN to be used for VMware vMotion.

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

51. Enter the FPV-vMotion VLAN ID <3000>.

52. Keep the Sharing Type as None.

53. Click OK, and then click OK again.

54. Right-click VLANs.

55. Select Create VLANs.

56. Enter `ACI-System-VLAN` as the name of the VLAN to be used for OpFlex communication to the VMware Virtual Switch.

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

58. Enter the ACI System VLAN ID <4093>.

> The ACI system VLAN ID can be determined by using ssh to connect to the APIC CLI and typing "ifconfig | grep bondo". You should see a bondo.xxxx interface listed. The xxxx is the ACI system VLAN id.

59. Keep the Sharing Type as None.

60. Click OK, and then click OK again.

61. Right-click VLANs.

62. Select Create VLANs.

63. Enter `FPV-vSwitch-Pool` as the prefix for this VLAN pool.

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

65. Enter a range of 100 VLANs for VLAN ID. <1100-1199> was used in this validation.

66. Keep the Sharing Type as None.

67. Click OK and then click OK again.

Modify Default Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To specify the UCS 3.2(3a) release for the Default firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

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

2. Select Policies > root.

3. Expand Host Firmware Packages.

4. Select default.

5. In the Actions pane, select Modify Package Versions.

6. Select the version 3.2(3a)B for the Blade Package, and 3.2(3a)C (optional) for the Rack Package.

7. Leave Excluded Components with only Local Disk selected.

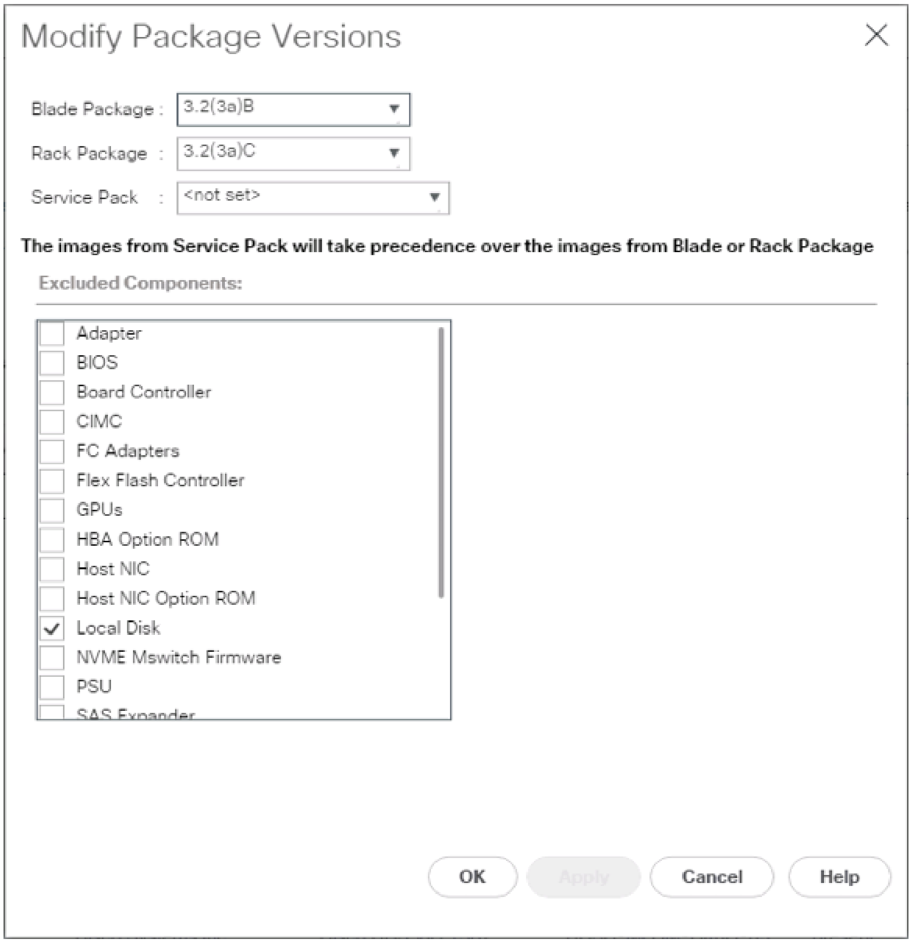

Modify Package Versions

Blade Package : 3.2(3a)B

Rack Package : 3.2(3a)C

Service Pack : <not set>

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

Excluded Components:

- ☐ Adapter
- ☐ BIOS
- ☐ Board Controller
- ☐ CIMC
- ☐ FC Adapters
- ☐ Flex Flash Controller
- ☐ GPUs
- ☐ HBA Option ROM
- ☐ Host NIC
- ☐ Host NIC Option ROM
- ☑ Local Disk
- ☐ NVME Mswitch Firmware
- ☐ PSU
- ☐ SAS Expander

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

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable the base quality of service in the Cisco UCS fabric, complete the following steps:

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

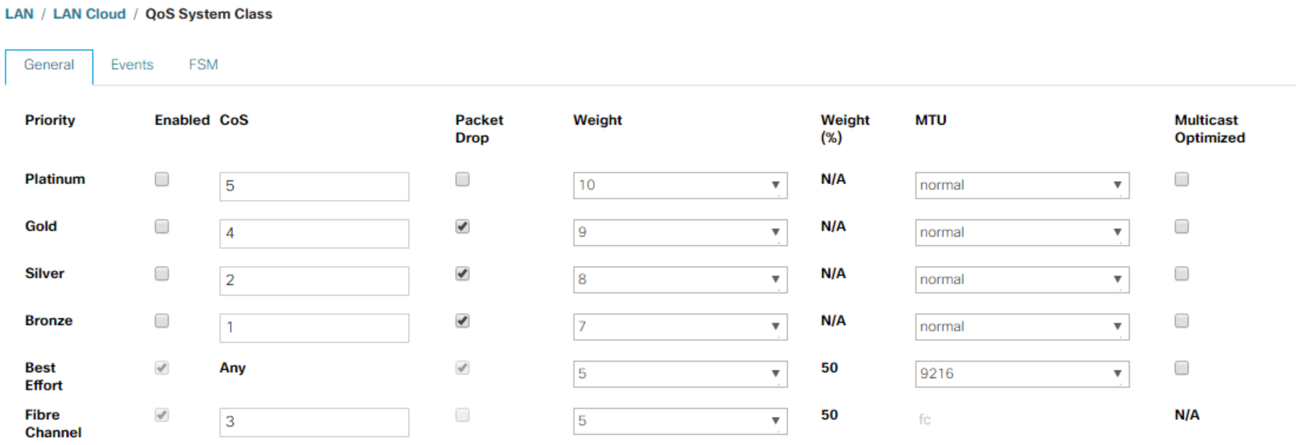

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	□	5	□	10	N/A	normal	□
Gold	□	4	☑	9	N/A	normal	□
Silver	□	2	☑	8	N/A	normal	□
Bronze	□	1	☑	7	N/A	normal	□
Best Effort	☑	Any	☑	5	50	9216	□
Fibre Channel	☑	3	□	5	50	fc	N/A

Create Local Disk Configuration Policy (Optional)

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

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

To create a local disk configuration policy, complete the following steps:

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

2. Select Policies > root.

3. Right-click Local Disk Config Policies.

4. Select Create Local Disk Configuration Policy.

5. Enter `SAN-Boot` as the local disk configuration policy name.

6. Change the mode to No Local Storage.

7. Click OK to create the local disk configuration policy.

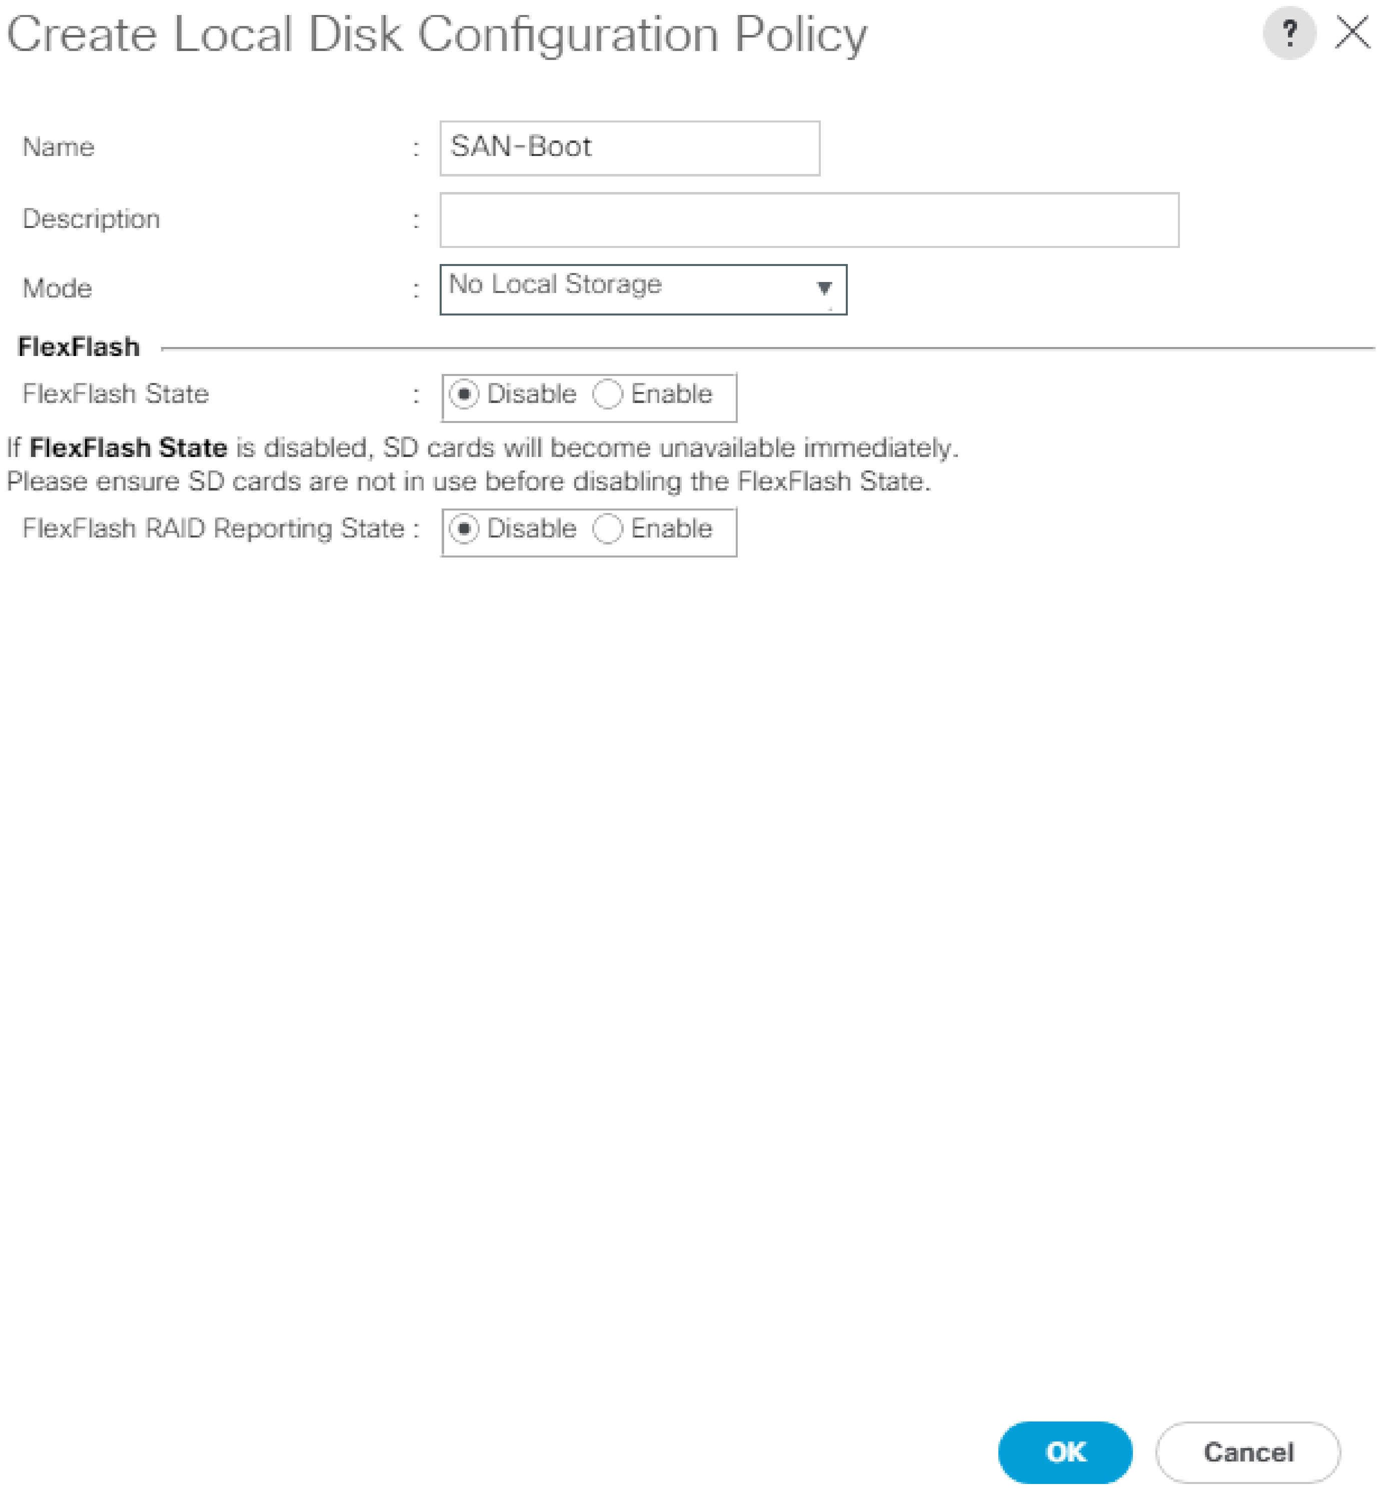

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : ● Disable ○ Enable

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

FlexFlash RAID Reporting State : ● Disable ○ Enable

8. Click OK.

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

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

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable-CDP-LLDP` as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

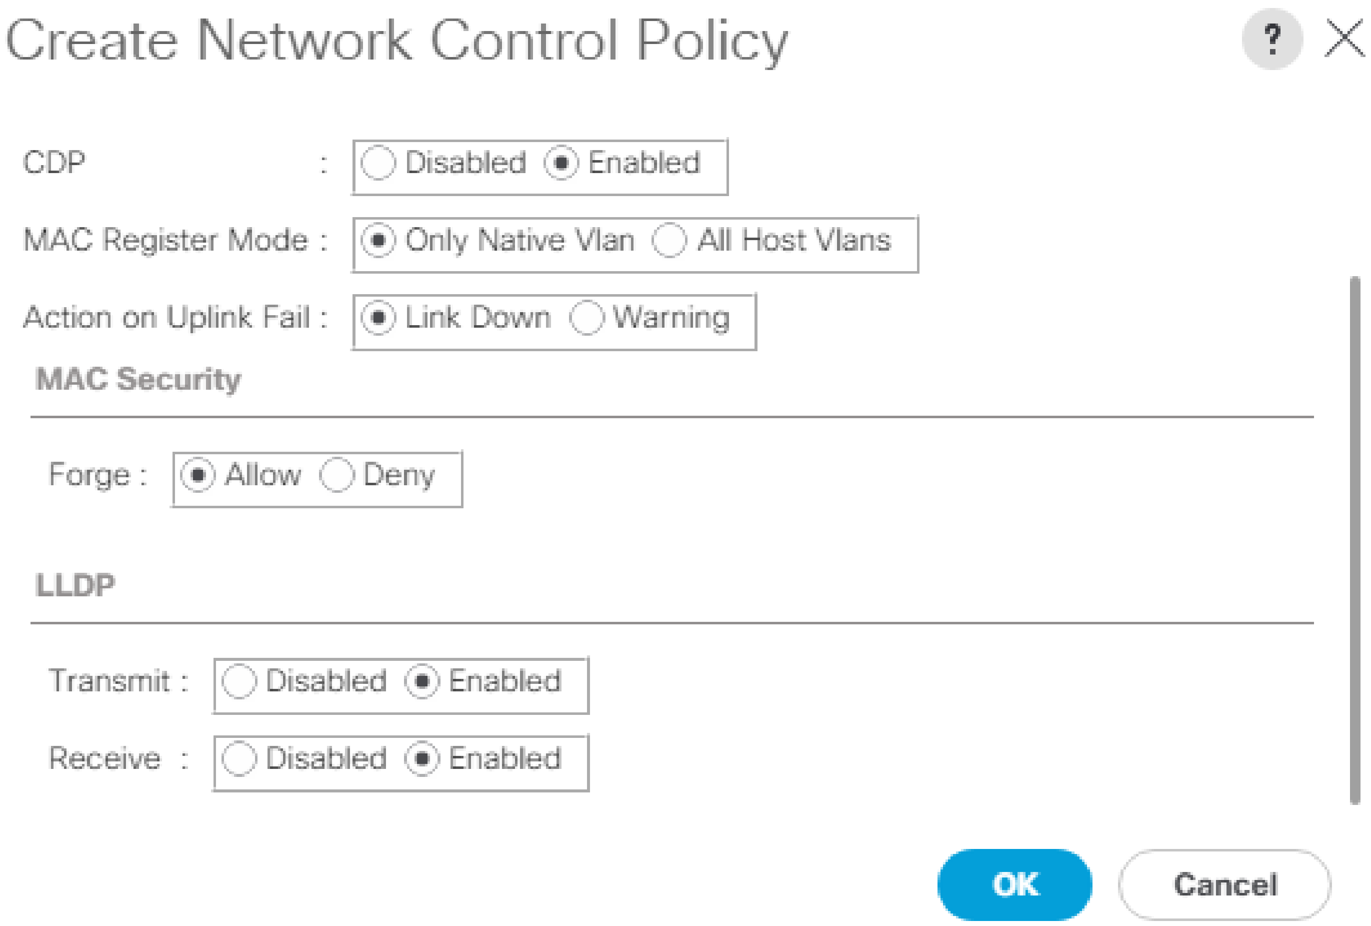

CDP : ○ Disabled ● Enabled

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

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 icon on the left.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter `No-Power-Cap` as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

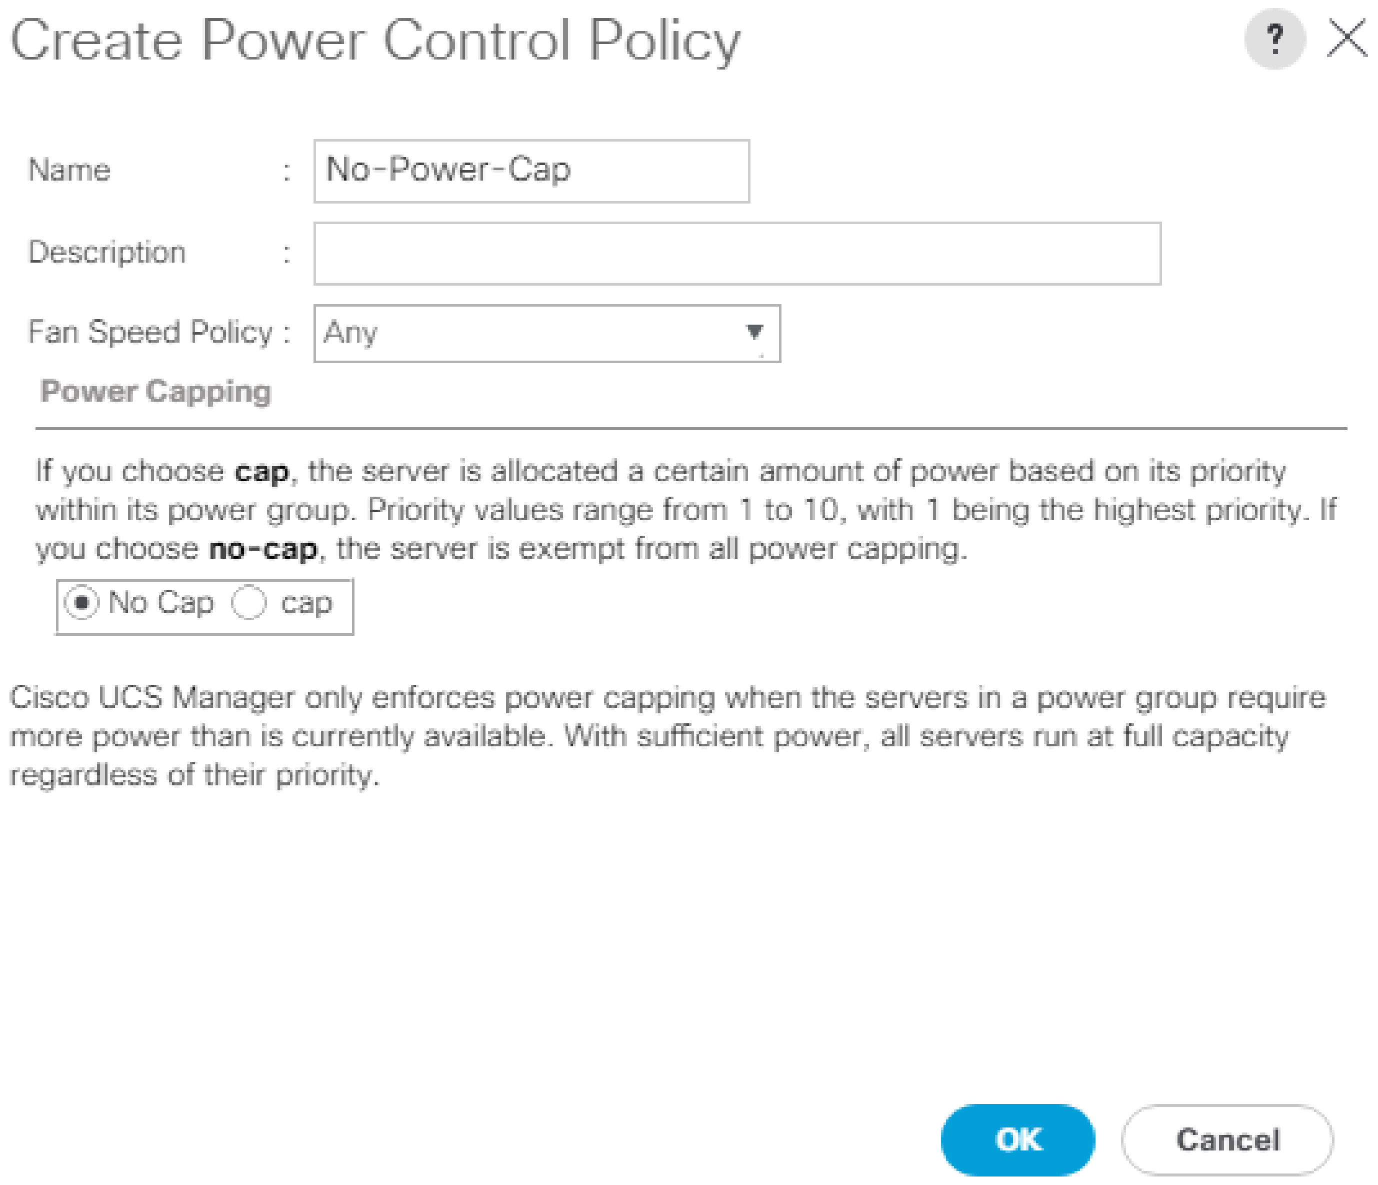

Create Power Control Policy

Name : No-Power-Cap

Description :

Fan Speed Policy : Any

Power Capping

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

● No Cap ○ cap

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

Create Server Pool Qualification Policy (Optional)

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

This example creates a policy to select Cisco UCS B200 M5 Servers.

1. In Cisco UCS Manager, click the Servers icon on the left.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.

5. Name the policy `UCSB-B200-M5`.

6. Select Create Server PID Qualifications.

7. Using the drop-down list, select `UCSB-B200-M5` as the PID.

8. Click OK to create the Server PID qualification.

9. Click OK to create the policy then OK for the confirmation.

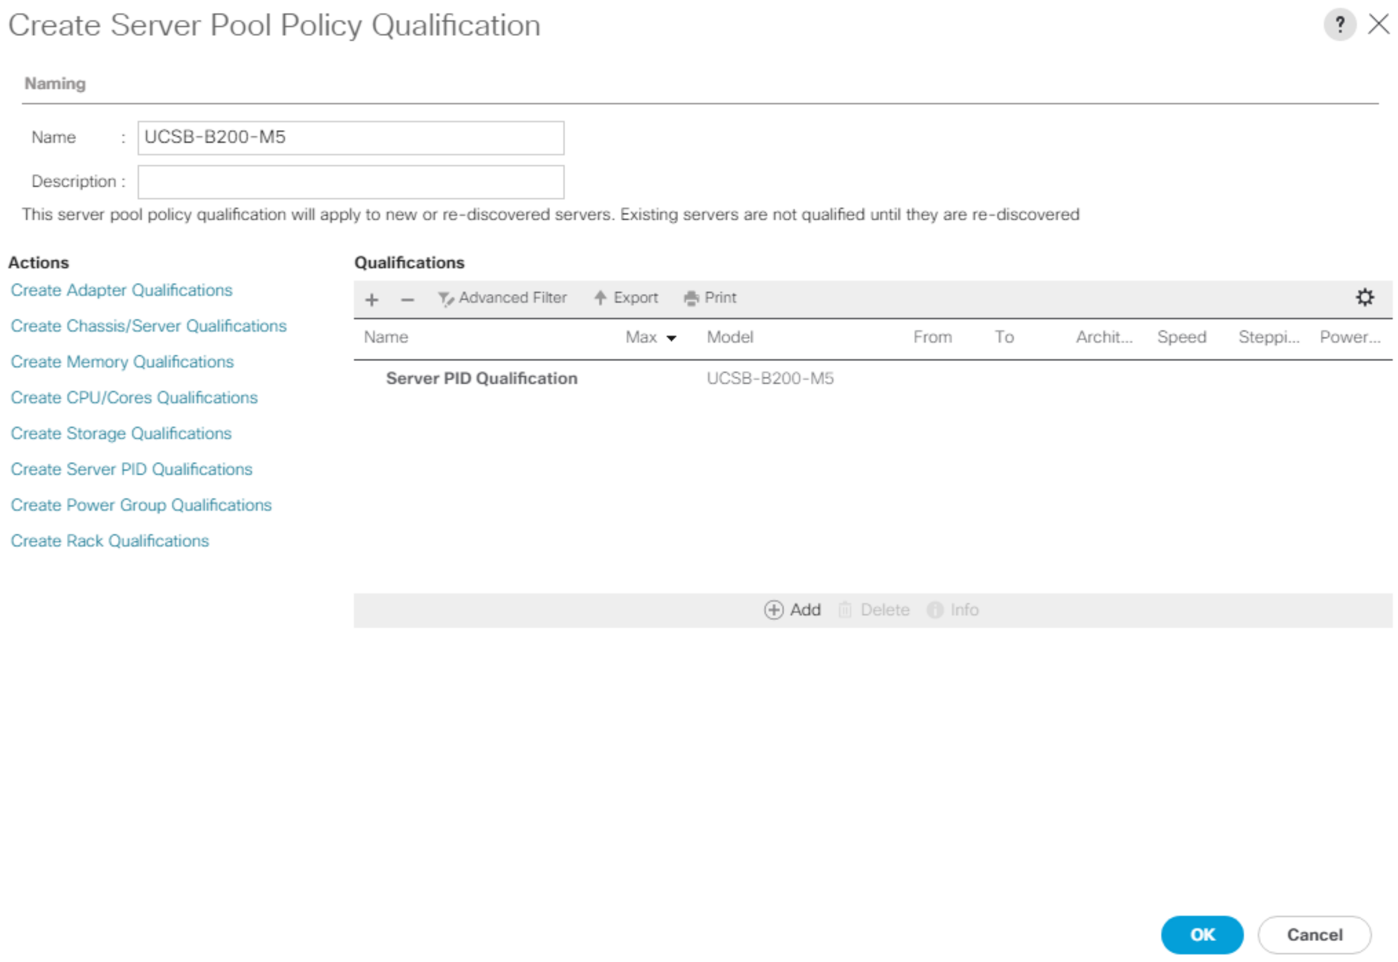

Create Server Pool Policy Qualification

Naming

Name : UCSB-B200-M5

Description :

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

Qualifications

Name	Max	Model	From	To	Archit...	Speed	Steppi...	Power...
Server PID Qualification		UCSB-B200-M5						

?
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
+
−
Advanced Filter
Export
Print
Add
Delete
Info
OK
Cancel

Create Server BIOS Policy

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

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

2. Select Policies > root.

3. Right-click BIOS Policies.

4. Select Create BIOS Policy.

5. Enter `Virtual-Host` as the BIOS policy name.

6. Click OK then OK again.

7. Expand BIOS Policies and select Virtual-Host.

8. On the right, change the Quiet Boot setting to Disabled.

9. Change CDN Control to Enabled.

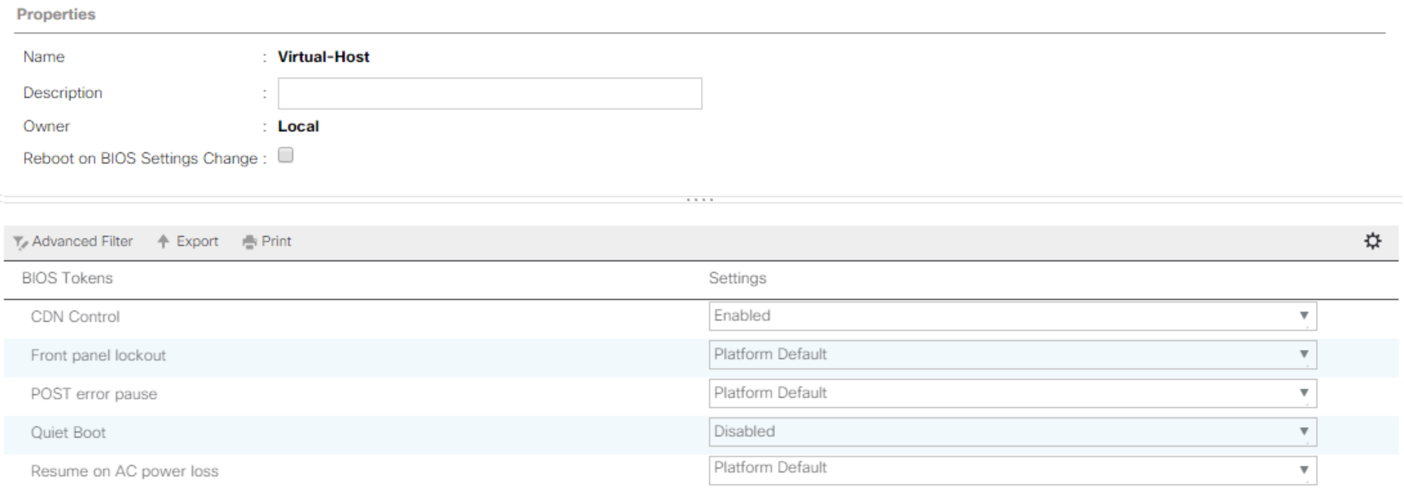

Properties

Name : **Virtual-Host**

Description :

Owner : **Local**

Reboot on BIOS Settings Change : □

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

10. Click Save Changes and OK.

11. Click the Advanced tab and then the Processor tab.

12. Set the following within the Processor tab:

    a. Processor C State -> Disabled
    b. Processor C1E -> Disabled
    c. Processor C3 Report -> Disabled
    d. Processor C7 Report -> Disabled
    e. Energy Performance -> Performance
    f. Frequency Floor Override -> Enabled
    g. DRAM Clock Throttling -> Performance

13. Click Save Changes and OK.

14. Click the RAS Memory tab and select:

    a. LV DDR Mode -> Performance-Mode

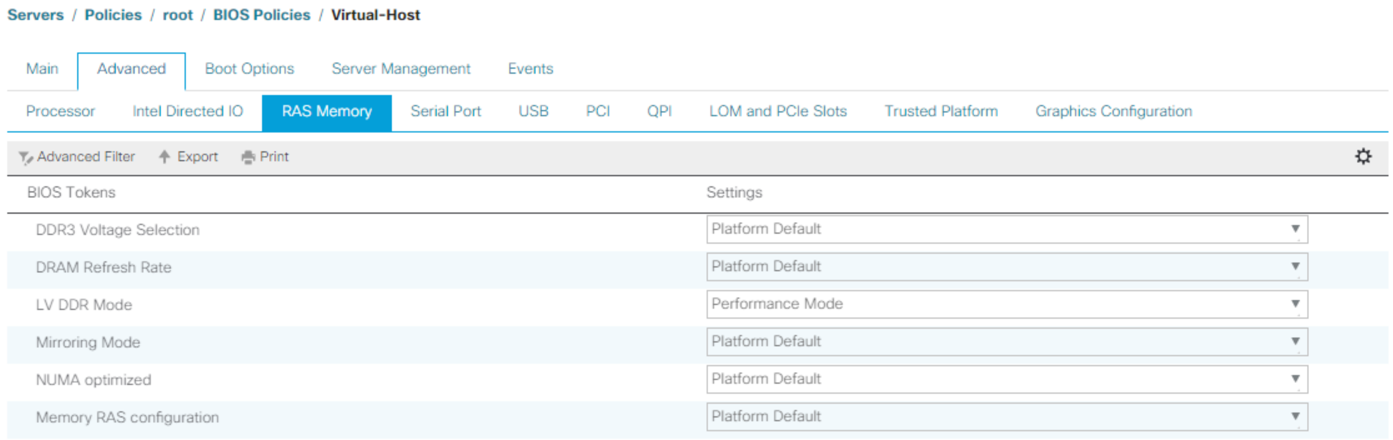

BIOS Tokens	Settings
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

15. Click Save Changes and OK.

Create High Traffic VMware Adapter Policy

To create the option VMware-High-Traffic Ethernet Adapter policy to provide higher vNIC performance, complete the following steps:

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

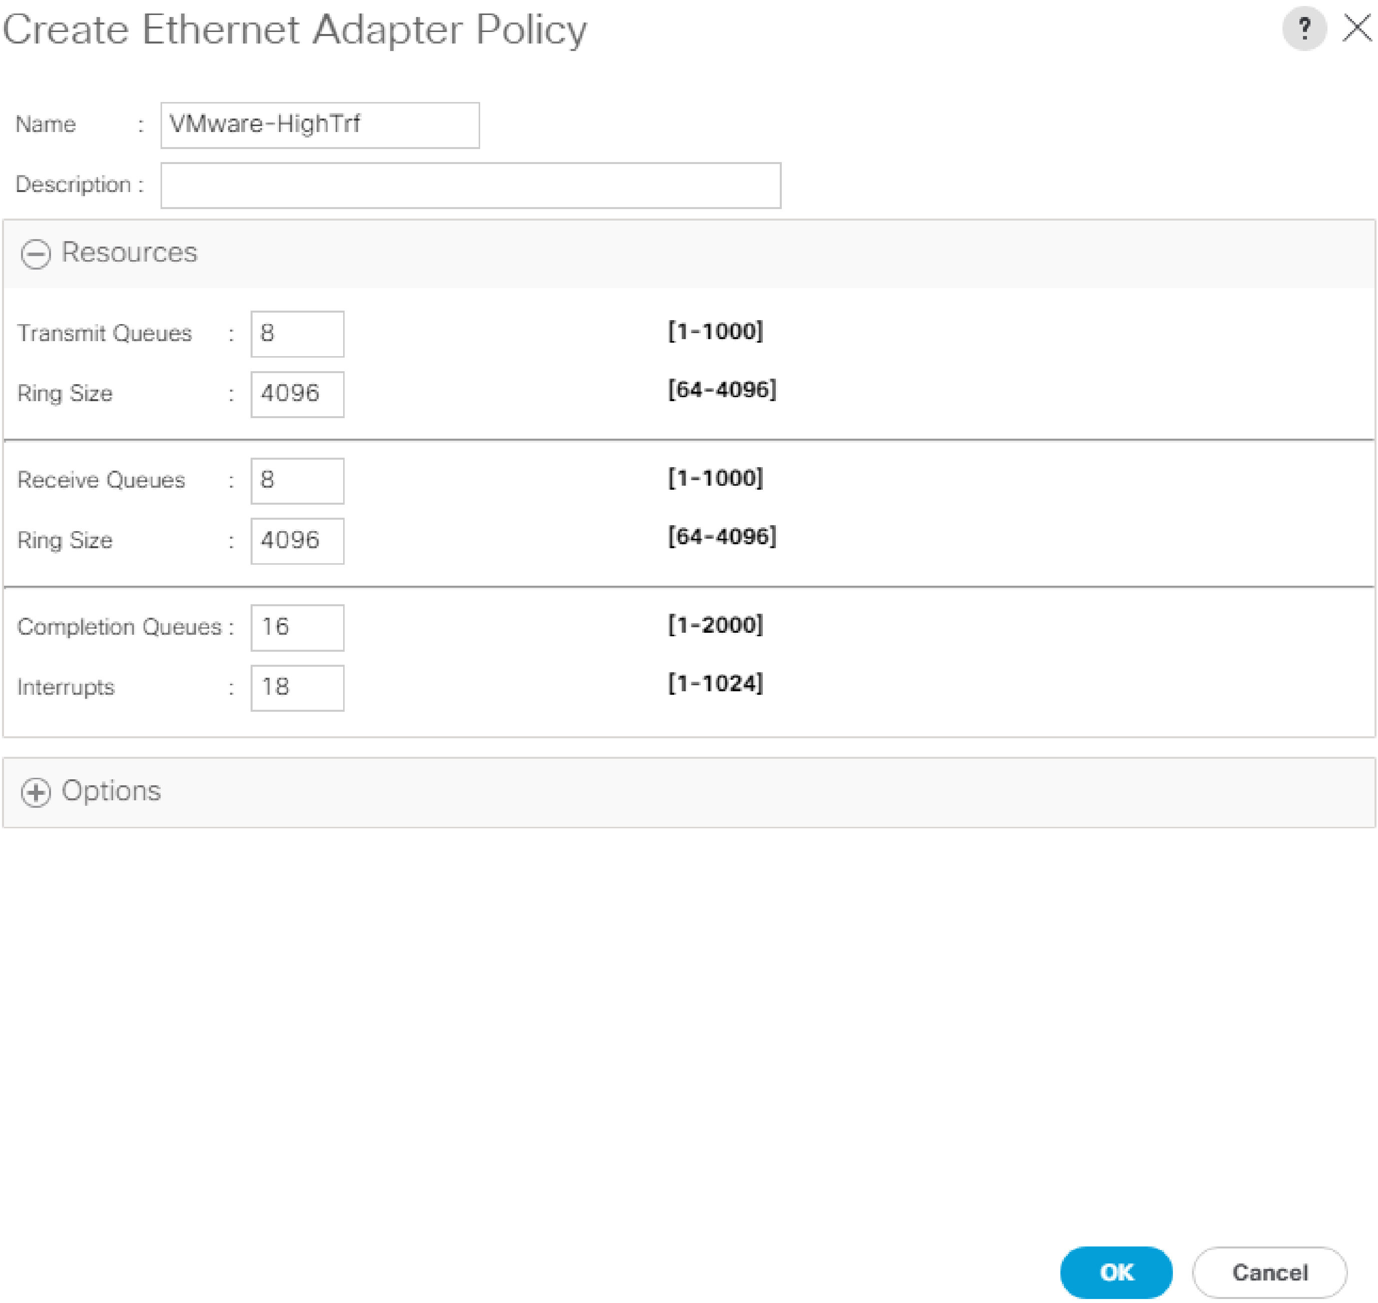

Create Ethernet Adapter Policy

Name :	VMware-HighTrf	
Description :		

⊖ Resources

Transmit Queues :	8	[1-1000]
Ring Size :	4096	[64-4096]
Receive Queues :	8	[1-1000]
Ring Size :	4096	[64-4096]
Completion Queues :	16	[1-2000]
Interrupts :	18	[1-1024]

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

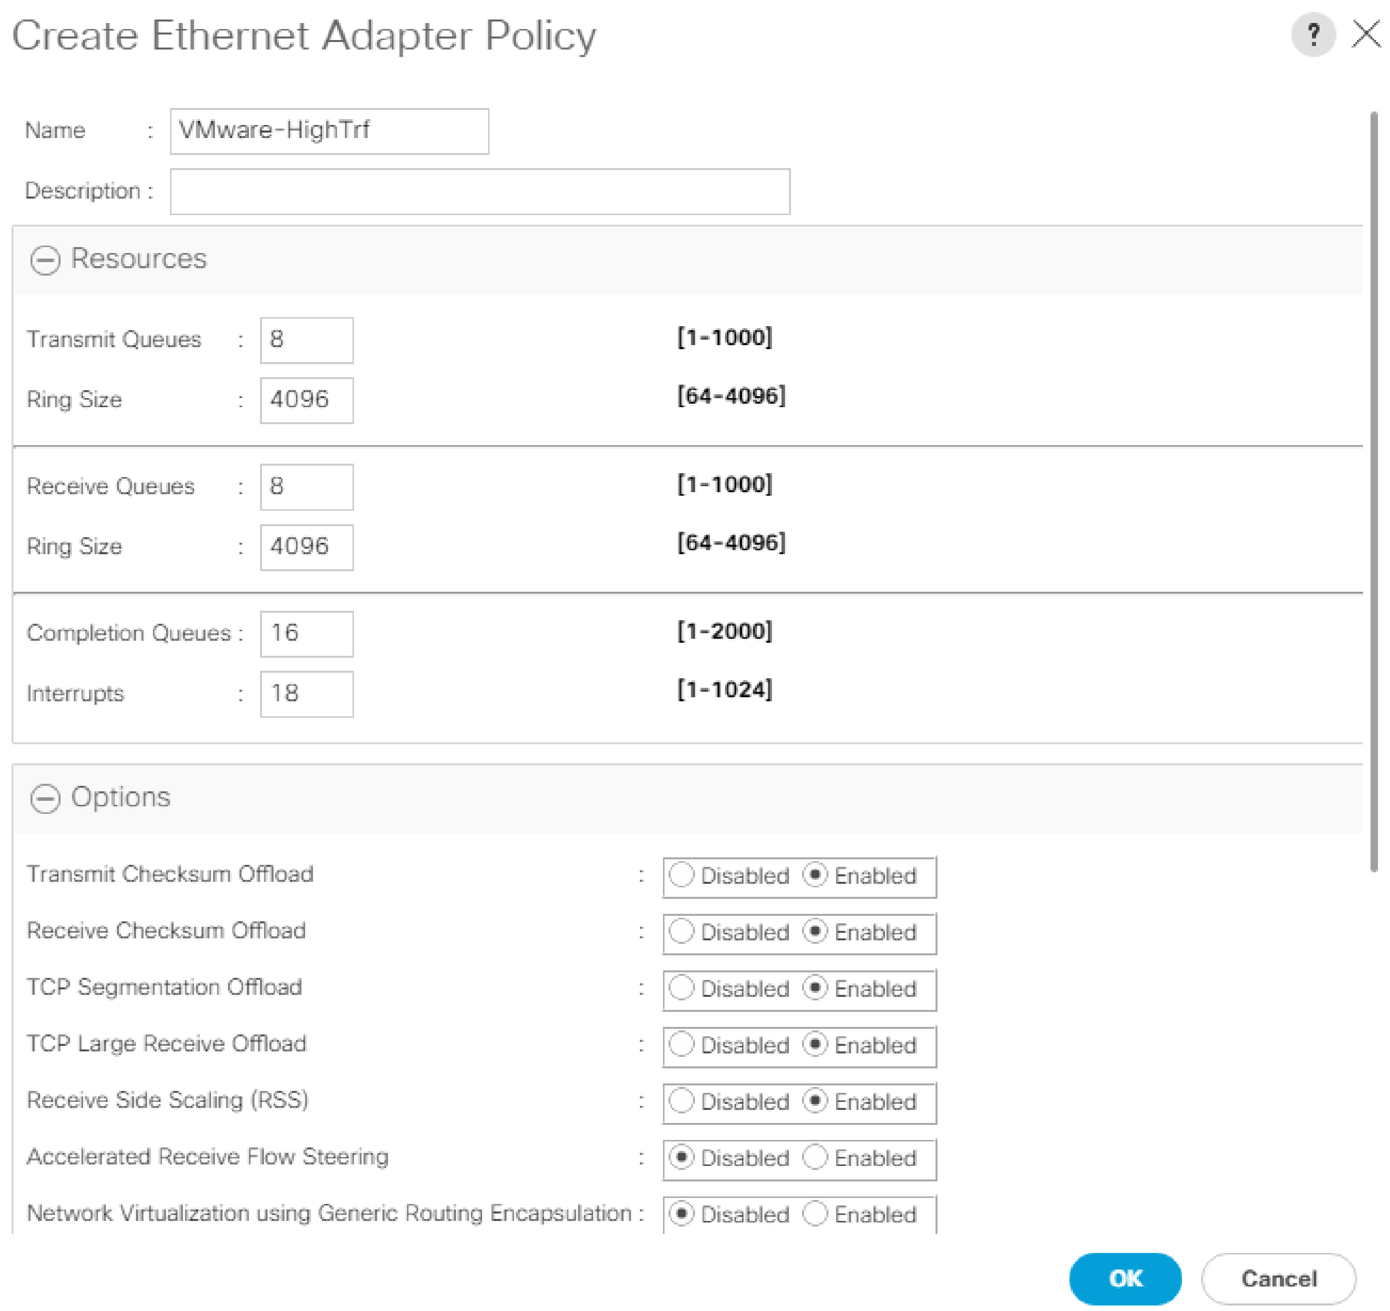

Create Ethernet Adapter Policy

Name : VMware-HighTrf

Description :

Resources

Transmit Queues : 8 [1-1000]

Ring Size : 4096 [64-4096]

Receive Queues : 8 [1-1000]

Ring Size : 4096 [64-4096]

Completion Queues : 16 [1-2000]

Interrupts : 18 [1-1024]

Options

Transmit Checksum Offload : ○ Disabled ● Enabled

Receive Checksum Offload : ○ Disabled ● Enabled

TCP Segmentation Offload : ○ Disabled ● Enabled

TCP Large Receive Offload : ○ Disabled ● Enabled

Receive Side Scaling (RSS) : ○ Disabled ● Enabled

Accelerated Receive Flow Steering : ● Disabled ○ Enabled

Network Virtualization using Generic Routing Encapsulation : ● Disabled ○ Enabled

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

Update the Default Maintenance Policy

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

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

2. Select Policies > root.

3. Select Maintenance Policies > default.

4. Change the Reboot Policy to User Ack.

5. Select "On Next Boot" to delegate maintenance windows to server administrators.

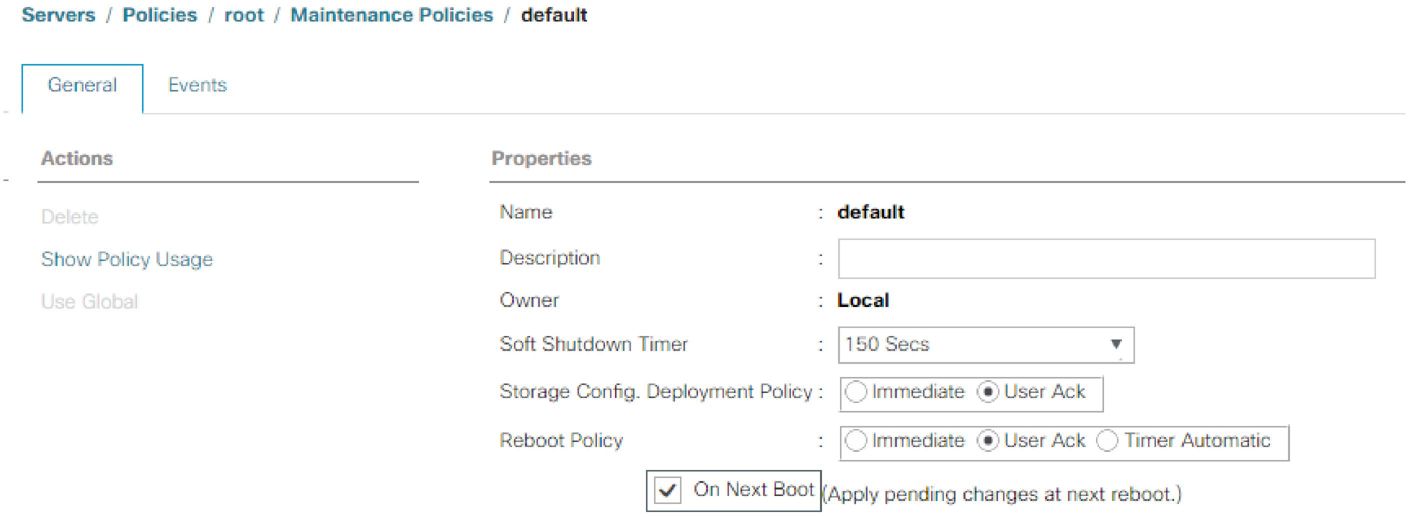

Servers / Policies / root / Maintenance Policies / default

Name : **default**

Description :

Owner : **Local**

Soft Shutdown Timer : 150 Secs ▼

Storage Config. Deployment Policy : ○ Immediate ● User Ack

Reboot Policy : ○ Immediate ● User Ack ○ Timer Automatic

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

6. Click Save Changes.

7. Click OK to accept the change.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 6 vNIC Templates will be created.

Create Infrastructure vNICs Templates

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

2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.

3. Right-click vNIC Templates under FPV-FlexPod.

4. Select Create vNIC Template.

5. Enter `Infra-A` as the vNIC template name.

6. Keep Fabric A selected.

7. Do not select the Enable Failover checkbox.

8. Select Primary Template for Redundancy Type.

9. Leave the Peer Redundancy Template set to <not set>.

10. Under Target, make sure that only the Adapter checkbox is selected.

11. Select Updating Template as the Template Type.

12. Under VLANs, select the checkboxes for FP-Core-Services, FPV-Foundation-NFS, FPV-IB-MGMT, FPV-vMotion, and Native-VLAN.

13. Set Native-VLAN as the native VLAN.

14. Select vNIC Name for the CDN Source.

15. For MTU, enter 9000.

16. In the MAC Pool list, select MAC-Pool-A.

17. In the Network Control Policy list, select Enable-CDP-LLDP.

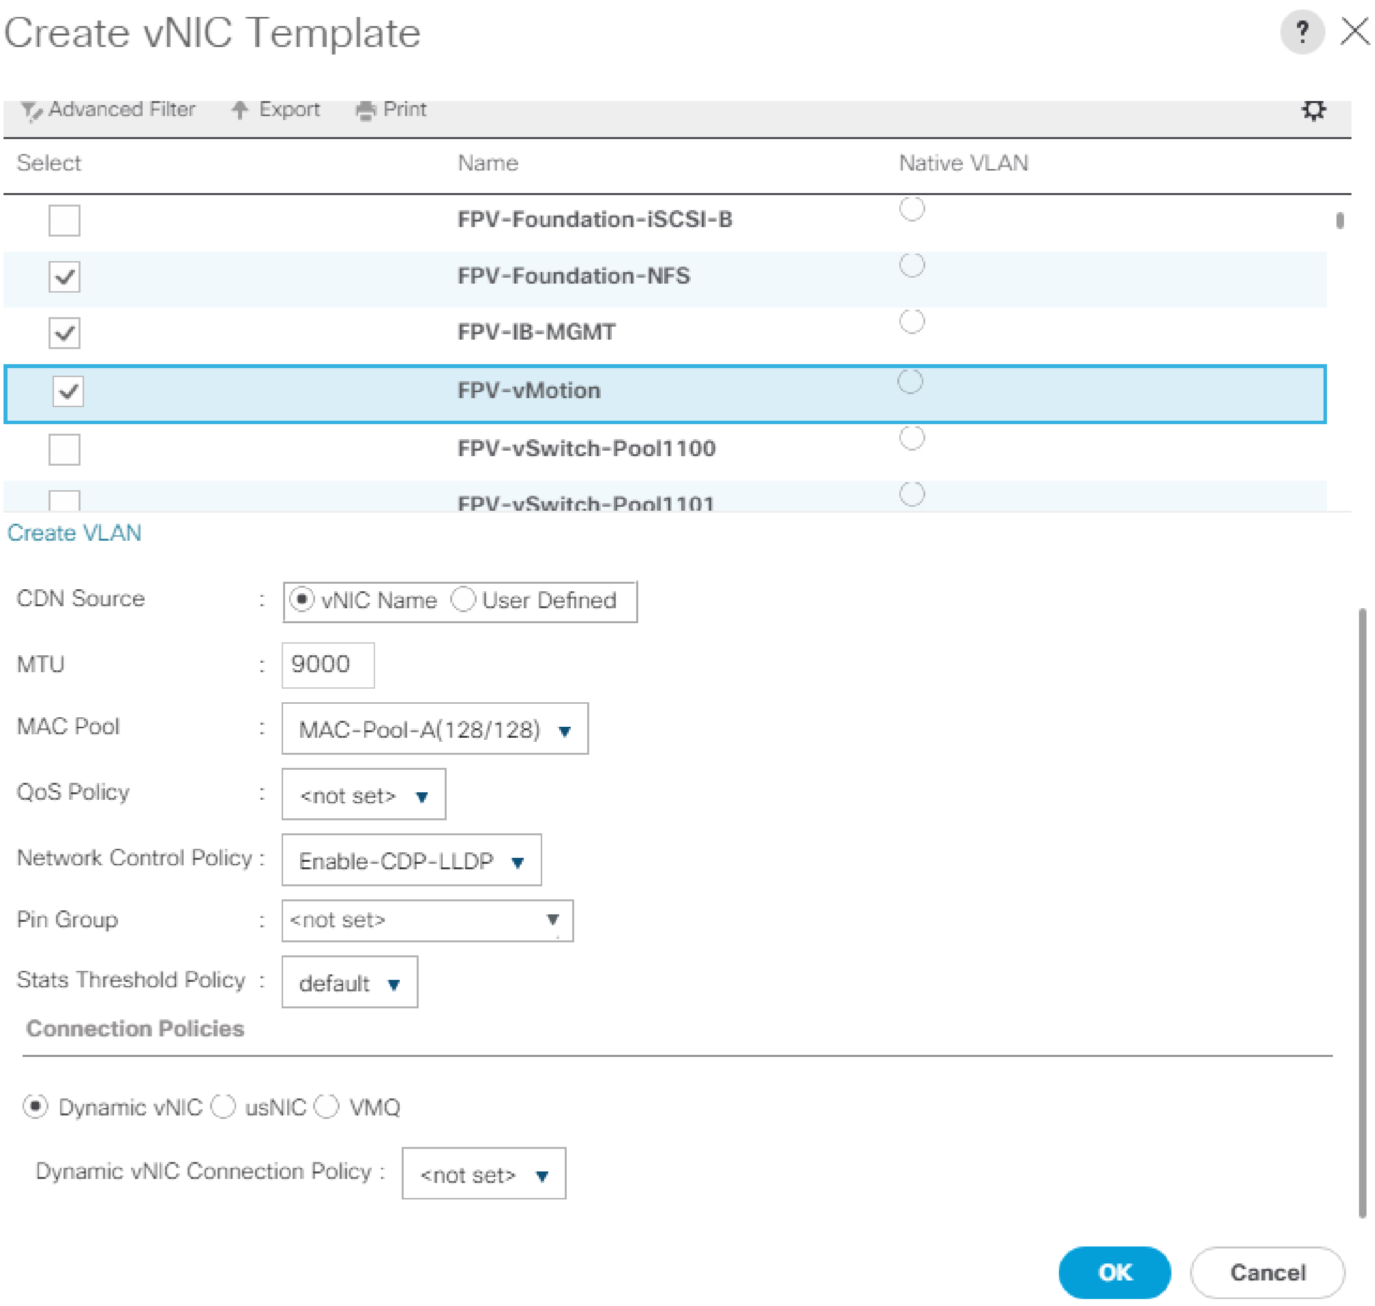

Create vNIC Template

Select	Name	Native VLAN
□	FPV-Foundation-iSCSI-B	○
☑	FPV-Foundation-NFS	○
☑	FPV-IB-MGMT	○
☑	FPV-vMotion	○
□	FPV-vSwitch-Pool1100	○
□	FPV-vSwitch-Pool1101	○

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>

18. Click OK to create the vNIC template.
19. Click OK.

Create the secondary redundancy template Infra-B:

1. Select the LAN icon on the left.
2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.
3. Right-click vNIC Templates under FPV-FlexPod.
4. Select Create vNIC Template.
5. Enter `Infra-B` as the vNIC template name.
6. Select Fabric B.
7. Do not elect the Enable Failover checkbox.
8. Set Redundancy Type to Secondary Template.
9. Select Infra-A for the Peer Redundancy Template.
10. In the MAC Pool list, select `MAC-Pool-B`. The MAC Pool is all that needs to be selected for the Secondary Template.
11. Click OK to create the vNIC template.
12. Click OK.

Create iSCSI Boot vNICs

To create iSCSI Boot vNICs, complete the following steps:

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.
3. Right-click vNIC Templates under FPV-FlexPod.
4. Select Create vNIC Template.
5. Enter `iSCSI-A` as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select No Redundancy for Redundancy Type.
9. Under Target, make sure that only the Adapter checkbox is selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select the checkbox for FPV-Foundation-iSCSI-A.

12. Set FPV-Foundation-iSCSI-A as the native VLAN.

13. Select vNIC Name for the CDN Source.

14. For MTU, enter 9000.

15. In the MAC Pool list, select MAC-Pool-A.

16. In the Network Control Policy list, select Enable-CDP-LLDP.

17. Click OK to create the vNIC template.

18. Click OK.

Create the Infra-iSCSI-B template:

1. Select the LAN icon on the left.

2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.

3. Right-click vNIC Templates under FPV-FlexPod.

4. Select Create vNIC Template.

5. Enter `iSCSI-B` as the vNIC template name.

6. Select Fabric B.

7. Do not elect the Enable Failover checkbox.

8. Select No Redundancy for Redundancy Type.

9. Under Target, make sure that only the Adapter checkbox is selected.

10. Select Updating Template as the Template Type.

11. Under VLANs, select the checkbox for FPV-Foundation-iSCSI-B.

12. Set FPV-Foundation-iSCSI-B as the native VLAN.

13. Select vNIC Name for the CDN Source.

14. For MTU, enter 9000.

15. In the MAC Pool list, select MAC-Pool-B.

16. In the Network Control Policy list, select Enable-CDP-LLDP.

17. Click OK to create the vNIC template.

18. Click OK.

Create vNIC Templates for APIC-Integrated Virtual Switch

To create vNIC templates for APIC-controlled virtual switch, complete the following steps:

Note that the ACI System VLAN is being added here for the set of Cisco VIC vNICs that will be used for the APIC-Integrated Virtual Switch. Although this VLAN is not necessary for the VMware vDS, it is being put in for future usage with the Cisco ACI Virtual Edge (AVE) virtual switch.

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.
3. Right-click vNIC Templates under FPV-FlexPod.
4. Select Create vNIC Template.
5. Enter `APIC-VS-A` as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Select Updating Template as the Template Type.
12. Under VLANs, select the checkboxes for `ACI-System-VLAN` and all 100 Virtual-Switch-Pool VLANs.
13. Do not set a native VLAN.
14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-Pool-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.

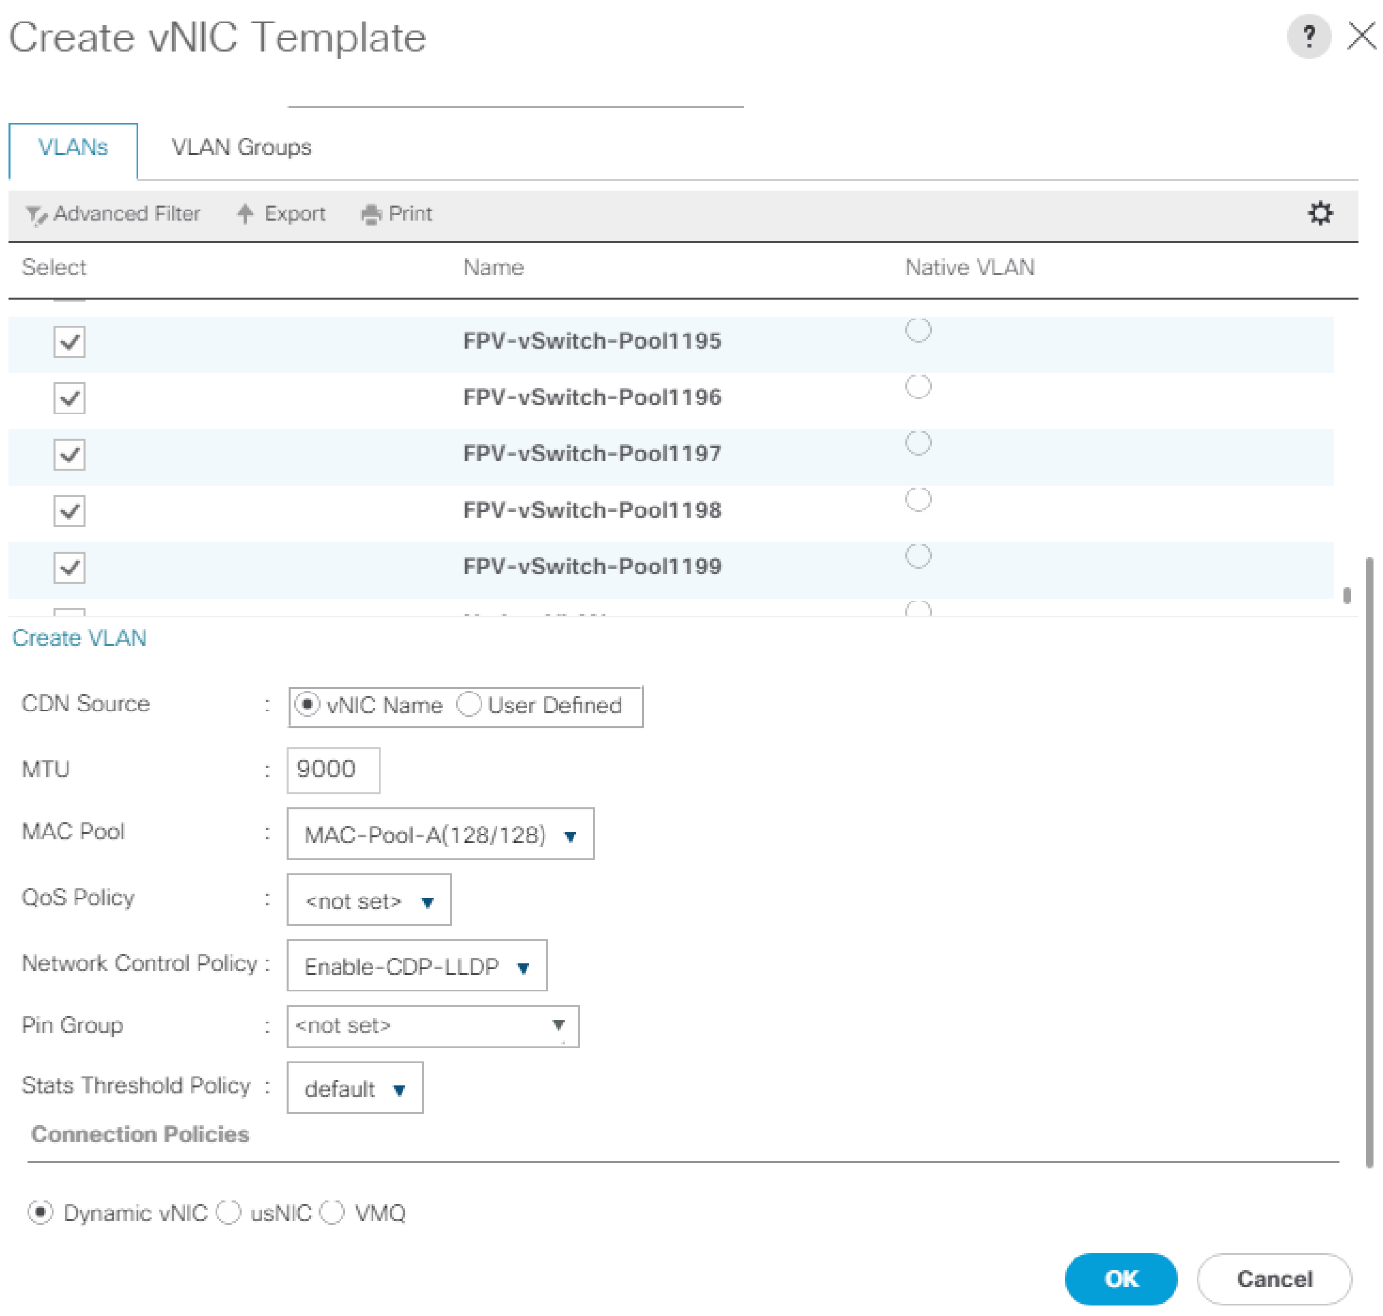

Create vNIC Template

Select	Name	Native VLAN
☑	FPV-vSwitch-Pool1195	○
☑	FPV-vSwitch-Pool1196	○
☑	FPV-vSwitch-Pool1197	○
☑	FPV-vSwitch-Pool1198	○
☑	FPV-vSwitch-Pool1199	○

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

18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template APIC-VS-B:

1. Select the LAN icon on the left.

2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.

3. Right-click vNIC Templates under FPV-FlexPod.

4. Select Create vNIC Template:

   a. Enter `APIC-VS-B` as the vNIC template name.
   b. Select Fabric B.
   c. Do not elect the Enable Failover checkbox.
   d. Set Redundancy Type to Secondary Template.
   e. Select APIC-VS-A for the Peer Redundancy Template.
   f. In the MAC Pool list, select `MAC-Pool-B`. The MAC Pool is all that needs to be selected for the Secondary Template.
   g. Click OK to create the vNIC template.
5. Click OK.

Create LAN Connectivity Policy for iSCSI Boot

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

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.
3. Right-click LAN Connectivity Policies under FPV-FlexPod.
4. Select Create LAN Connectivity Policy.
5. Enter `iSCSI-Boot` as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter `00-Infra-A` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Infra-A.
10. In the Adapter Policy list, select VMware.
11. Click OK to add this vNIC to the policy.

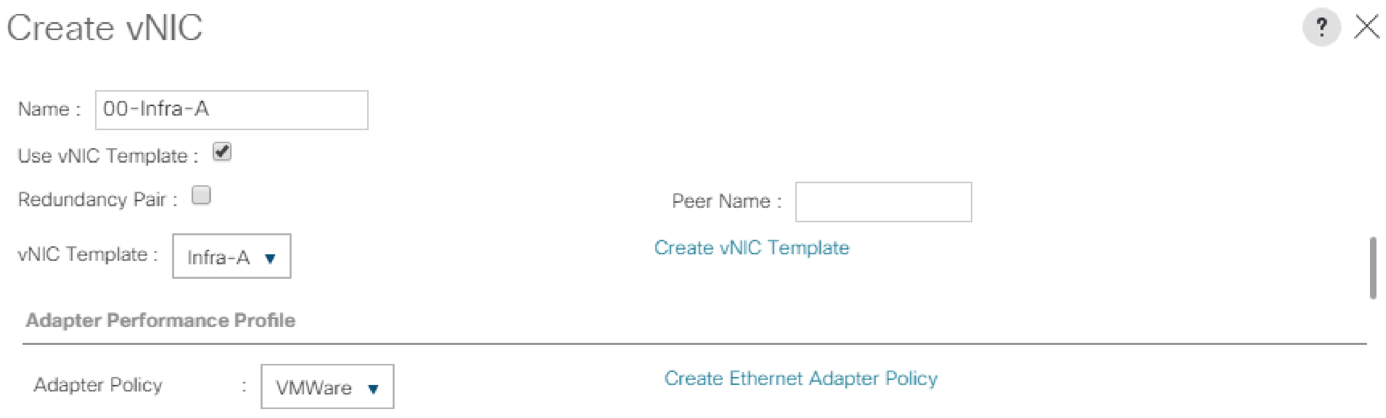

Create vNIC

Name : 00-Infra-A

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : Infra-A ▼

Adapter Performance Profile

Adapter Policy : VMWare ▼

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

13. In the Create vNIC box, enter `01-Infra-B` as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select Infra-B.
16. In the Adapter Policy list, select VMware.
17. Click OK to add the vNIC to the policy.
18. Click the upper Add button to add another vNIC to the policy.
19. In the Create vNIC box, enter `02-iSCSI-A` as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select iSCSI-A.
22. In the Adapter Policy list, select VMware.
23. Click OK to add the vNIC to the policy.
24. Click the upper Add button to add another vNIC to the policy.
25. In the Create vNIC box, enter `03-iSCSI-B` as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select iSCSI-B.
28. In the Adapter Policy list, select VMware.
29. Click OK to add the vNIC to the policy.
30. Click the upper Add button to add another vNIC to the policy.
31. In the Create vNIC box, enter `04-APIC-VS-A` as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select APIC-VS-A.
34. In the Adapter Policy list, select VMware. Optionally, select the VMware-HighTrf Adapter Policy.
35. Click OK to add the vNIC to the policy.
36. Click the upper Add button to add another vNIC to the policy.
37. In the Create vNIC box, enter `05-APIC-VS-B` as the name of the vNIC.
38. Select the Use vNIC Template checkbox.
39. In the vNIC Template list, select APIC-VS-B.
40. In the Adapter Policy list, select VMware. Optionally, select the VMware-HighTrf Adapter Policy.

41. Click OK to add the vNIC to the policy.
42. Expand the Add iSCSI vNICs section.
43. Click the lower Add button to add an iSCSI boot vNIC to the policy.
44. In the Create iSCSI vNIC box, enter `iSCSI-Boot-A` as the name of the vNIC.
45. Select 02-iSCSI-A for the Overlay vNIC.
46. Select the default iSCSI Adapter Policy.
47. FPV-Foundation-iSCSI-A (native) should be selected as the VLAN.
48. Do not select anything for MAC Address Assignment.

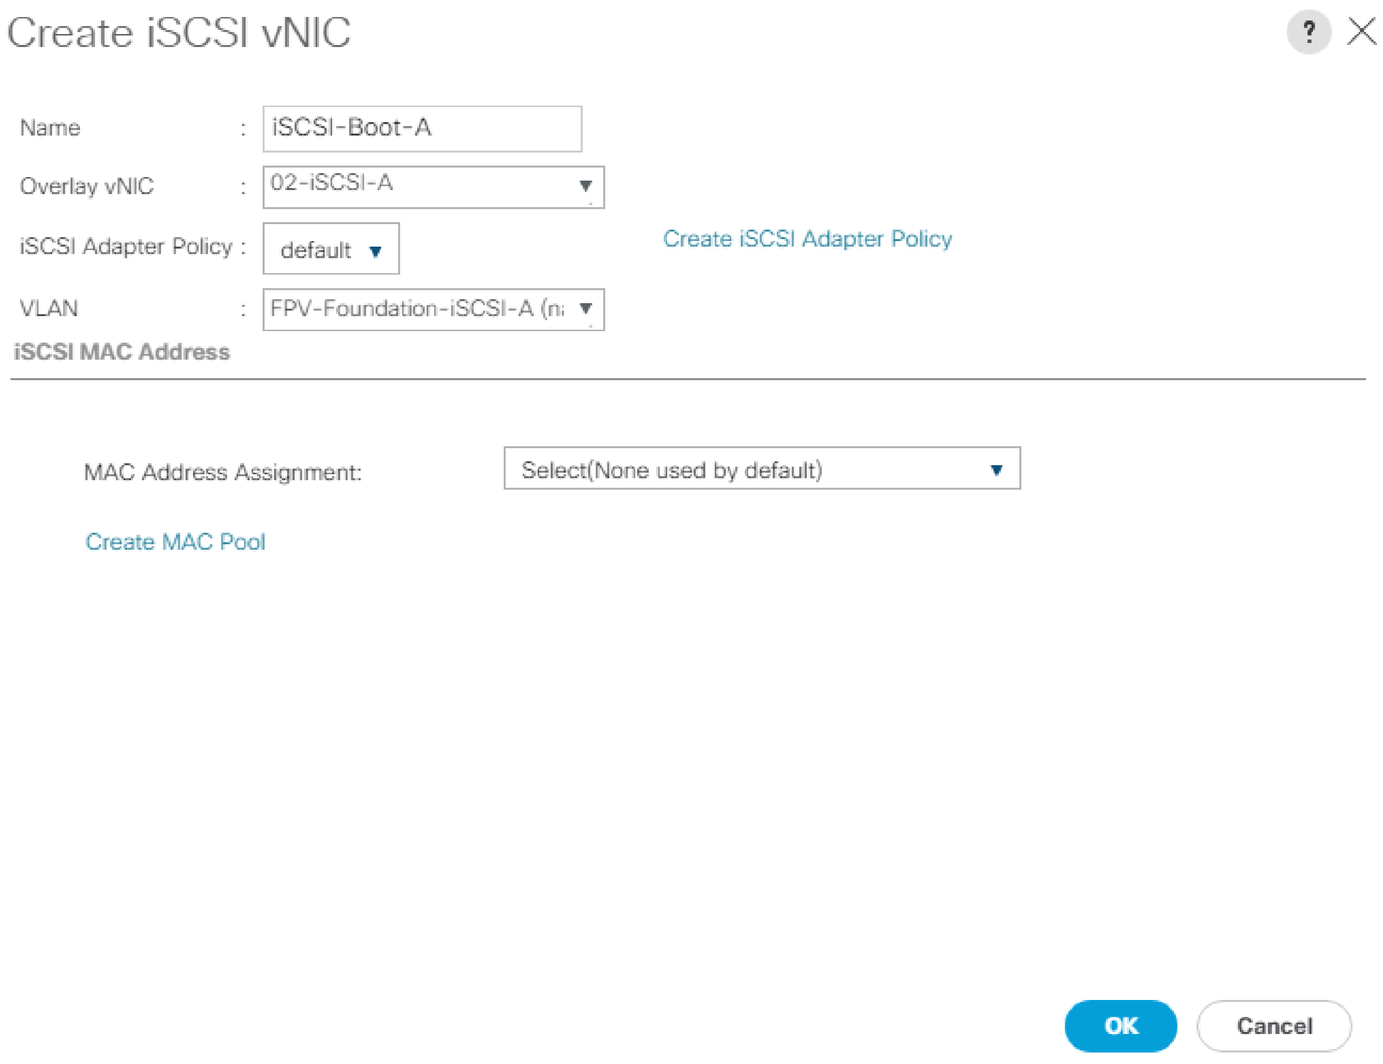

Create iSCSI vNIC

Name : iSCSI-Boot-A

Overlay vNIC : 02-iSCSI-A

iSCSI Adapter Policy : default

VLAN : FPV-Foundation-iSCSI-A (n

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

49. Click OK to add the vNIC to the policy.
50. Click the lower Add button to add an iSCSI boot vNIC to the policy.

51. In the Create iSCSI vNIC box, enter `iSCSI-Boot-B` as the name of the vNIC.

52. Select 03-iSCSI-B for the Overlay vNIC.

53. Select the default iSCSI Adapter Policy.

54. FPV-Foundation-iSCSI-B (native) should be selected as the VLAN.

55. Do not select anything for MAC Address Assignment.

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

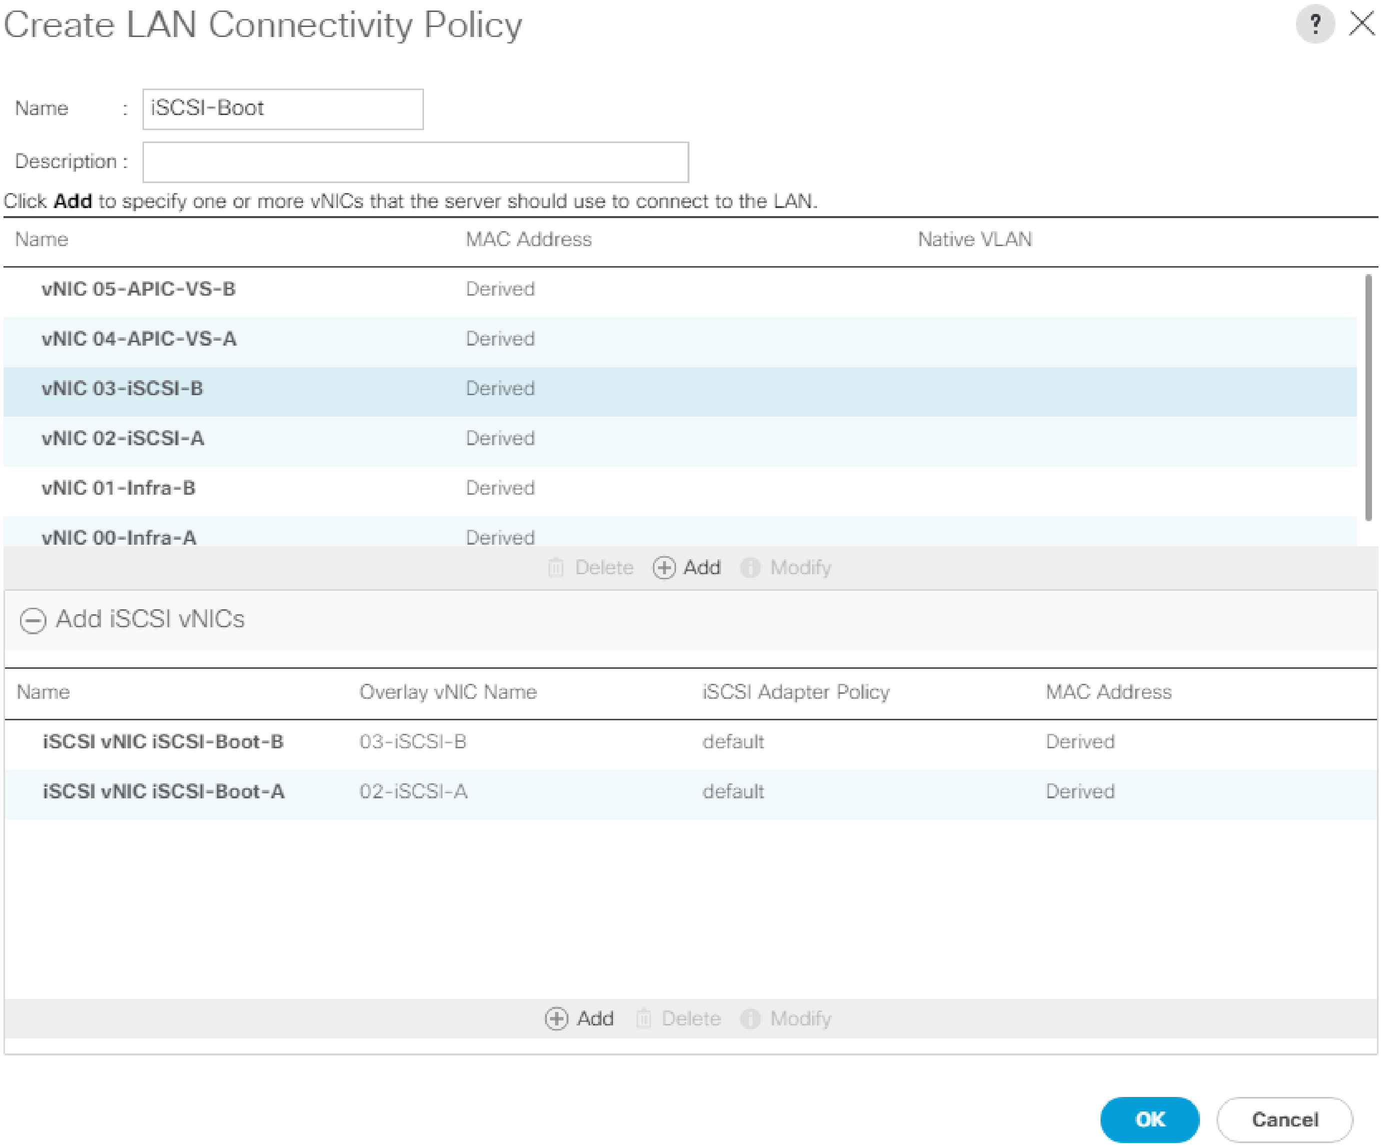

Create LAN Connectivity Policy

Name : iSCSI-Boot

Description :

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

Name	MAC Address	Native VLAN
vNIC 05-APIC-VS-B	Derived	
vNIC 04-APIC-VS-A	Derived	
vNIC 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

Add iSCSI vNICs

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

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

Create iSCSI Boot Policies

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on storage cluster node 1 (iscsi_lif01a and iscsi_lif01b) and two iSCSI LIFs are on storage cluster node 2 (iscsi_lif02a and iscsi_lif02b).

This boot policy configures the primary target to be iscsi_lif01a with four SAN paths.

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

1. In Cisco UCS Manager, click the Servers icon on the left.
2. Expand Policies > root > Sub-Organizations > FPV-FlexPod.
3. Right-click Boot Policies under FPV-FlexPod.
4. Select Create Boot Policy.
5. Enter `iSCSI-Boot` 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 list and select `Remote CD/DVD`.
9. Expand the iSCSI vNICs drop-down list and select Add iSCSI Boot.
10. Enter `iSCSI-Boot-A` in the iSCSI vNIC field.

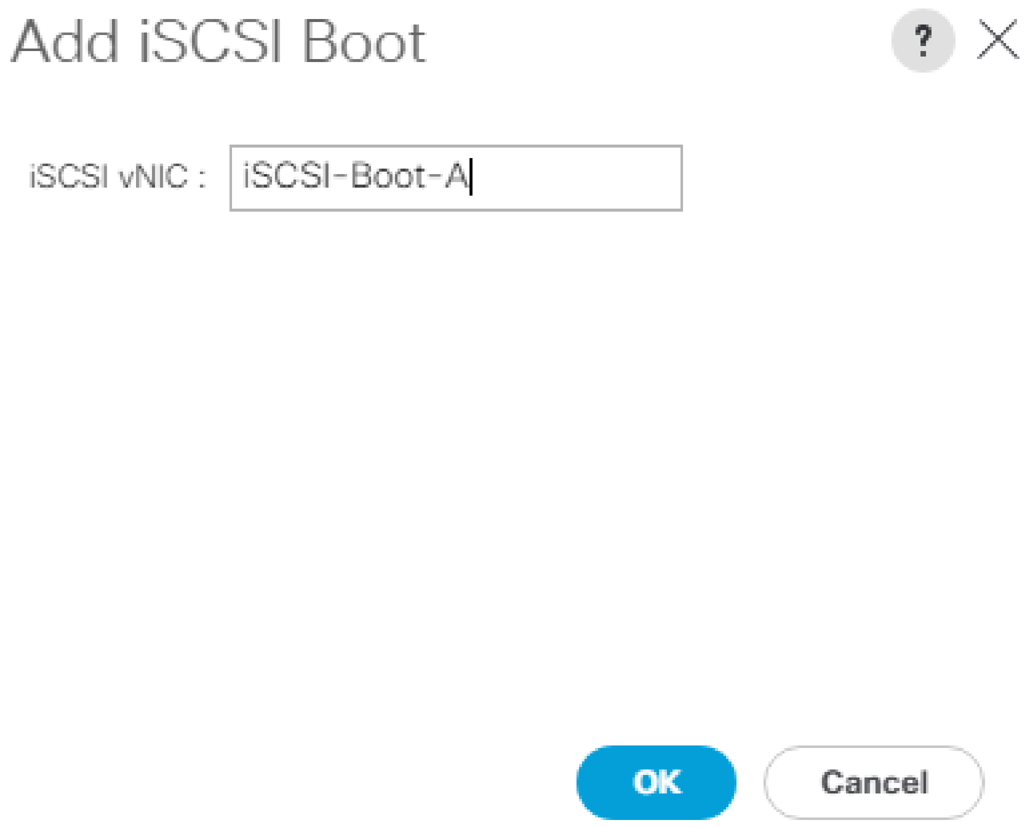

Add iSCSI Boot

iSCSI vNIC : iSCSI-Boot-A

11. Click OK.

12. From the iSCSI vNICs drop-down list, select Add iSCSI Boot.

13. Enter `iSCSI-Boot-B` in the iSCSI vNIC field.

14. Click OK.

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

Create iSCSI Boot Service Profile Template

In this procedure, one service profile template is created for Fabric A boot.

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

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

2. Expand Service Profile Templates > root > Sub-Organizations > FPV-FlexPod.

3. Select and right-click FPV-FlexPod.

4. Select Create Service Profile Template to open the Create Service Profile Template wizard.

5. Enter `iSCSI-Boot-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.

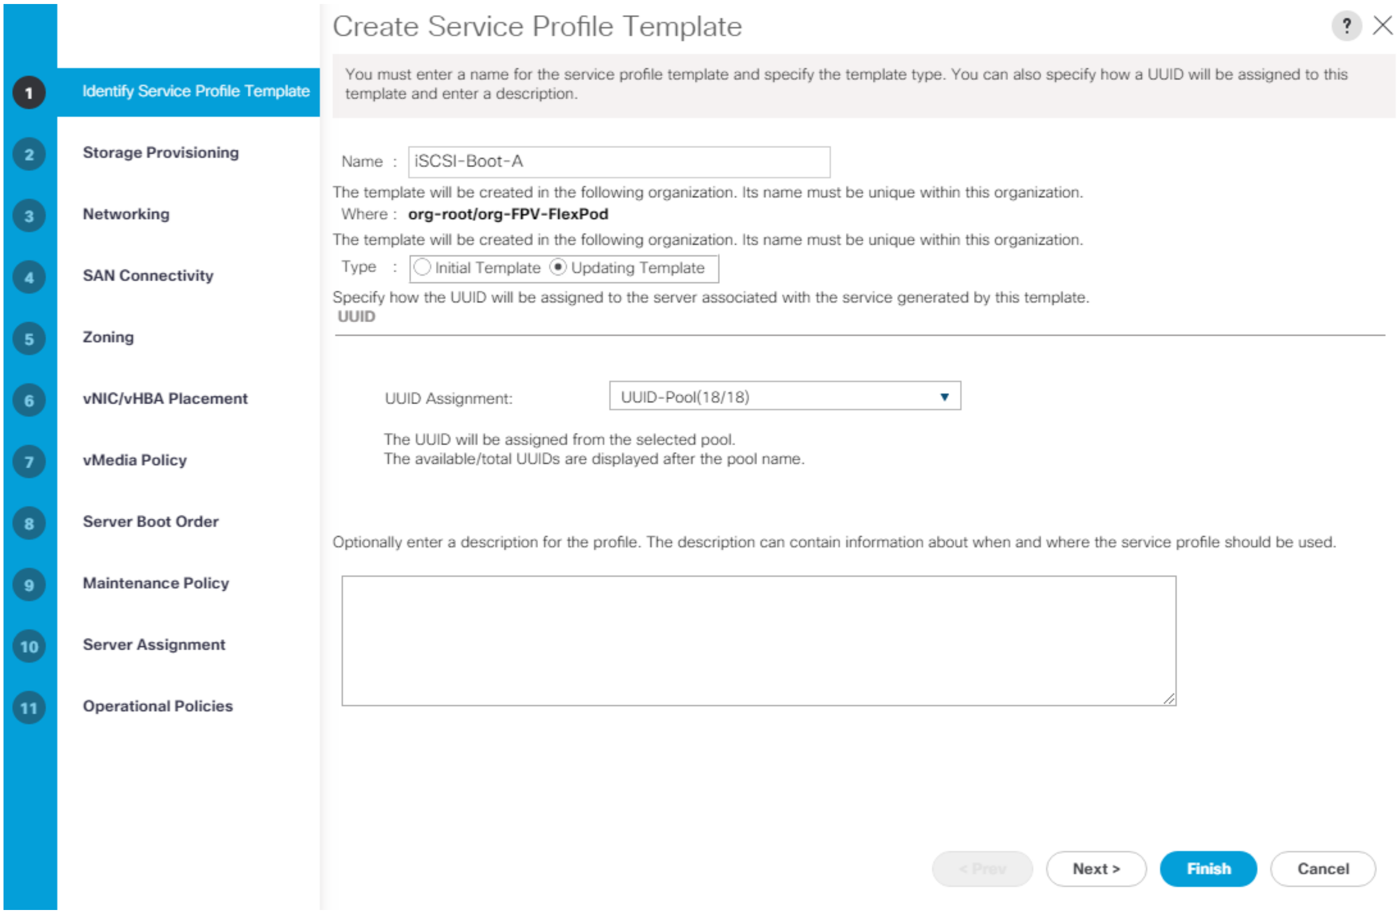

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name : iSCSI-Boot-A

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-FPV-FlexPod**

The template will be created in the following organization. Its name must be unique within this organization.

Type : ○ Initial Template ● Updating Template

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

UUID

UUID Assignment: UUID-Pool(18/18)

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

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

8. Click Next.

Configure Storage Provisioning

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

2. Click Next.

Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.

2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.

3. Select `iSCSI-Boot` from the LAN Connectivity Policy drop-down list.

4. Select `IQN-Pool` from the Initiator Name Assignment drop-down list.

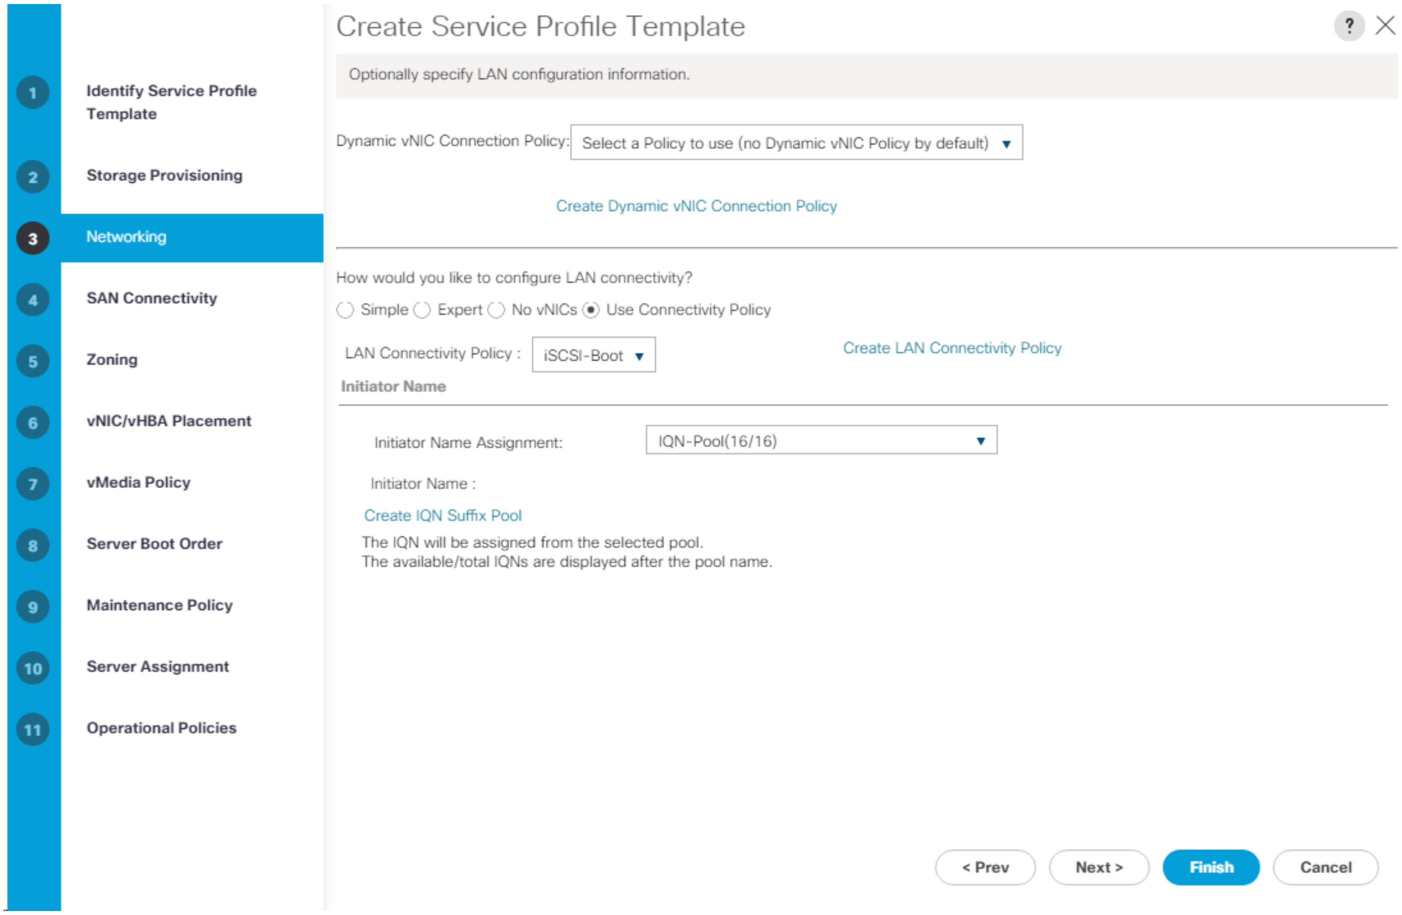

Create Service Profile Template

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 : iSCSI-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.

5. Click Next.

Configure SAN Connectivity Options

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

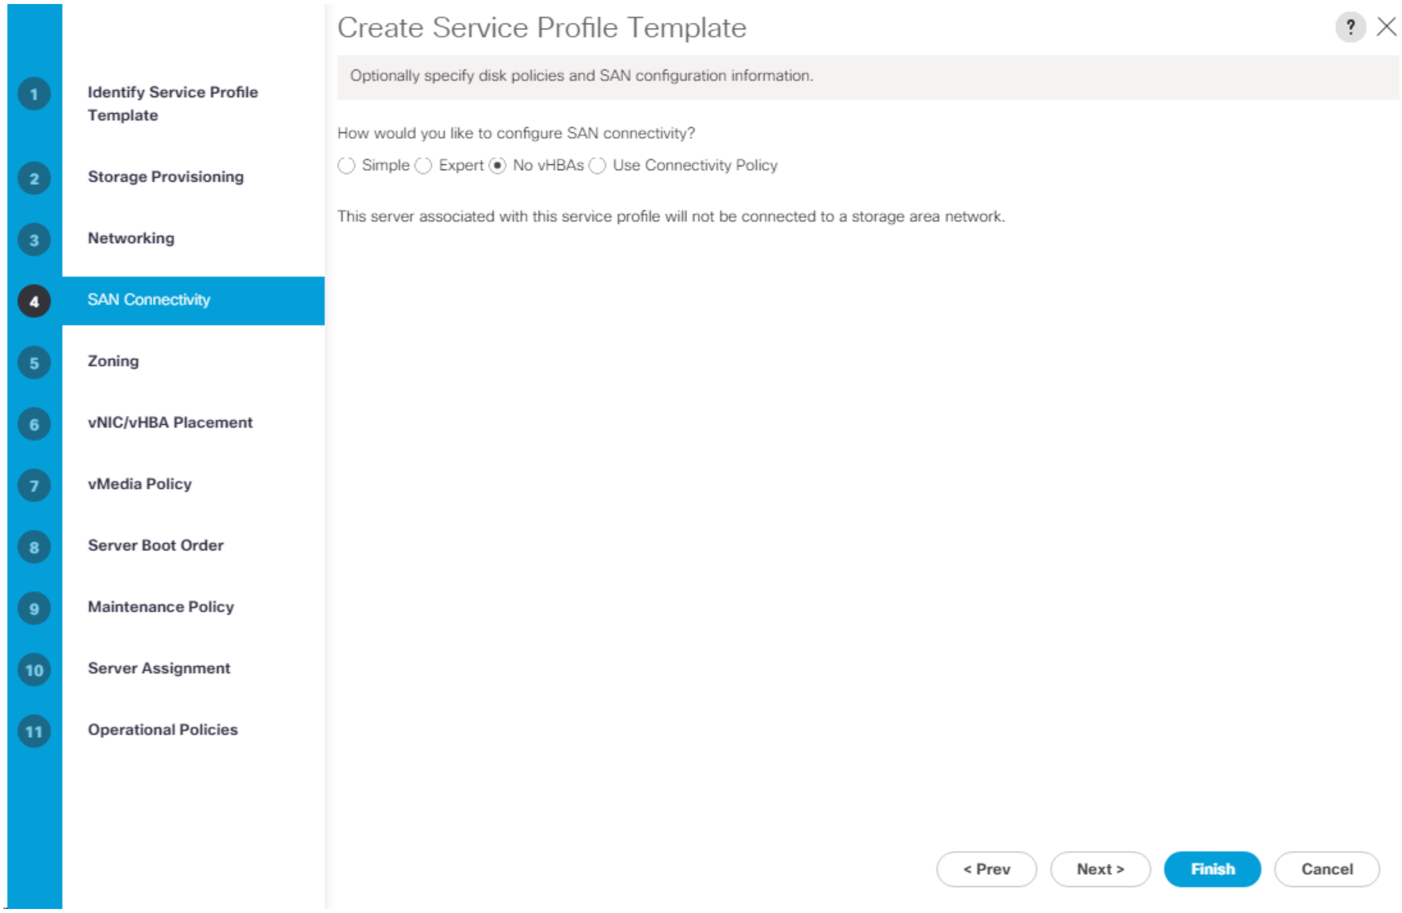

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

○ Simple ○ Expert ● No vHBAs ○ Use Connectivity Policy

This server associated with this service profile will not be connected to a storage area network.

2. Click Next.

Configure Zoning Options

1. Configure no Zoning Options and click Next.

Configure vNIC/HBA Placement

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

Configure vMedia Policy

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

Configure Server Boot Order

1. Select `iSCSI-Boot` for Boot Policy.
2. Under Boot Order, expand Boot Order and select the iSCSI-Boot-A row.
3. Select the Set iSCSI Boot Parameters button.
4. Select `iSCSI-IP-Pool-A` for the Initiator IP Address Policy.

5. Scroll to the bottom of the window and click Add.

6. Enter the IQN (Target Name) from the FPV-Foundation-SVM iSCSI Target Name. To get this IQN, ssh into the storage cluster interface and type "iscsi show".

7. For IPv4 address, enter the IP address of `iscsi_lif02a` from the FPV-Foundation-SVM. To get this IP, ssh into the storage cluster interface and type "network interface show –vserver FPV-Foundation-SVM".

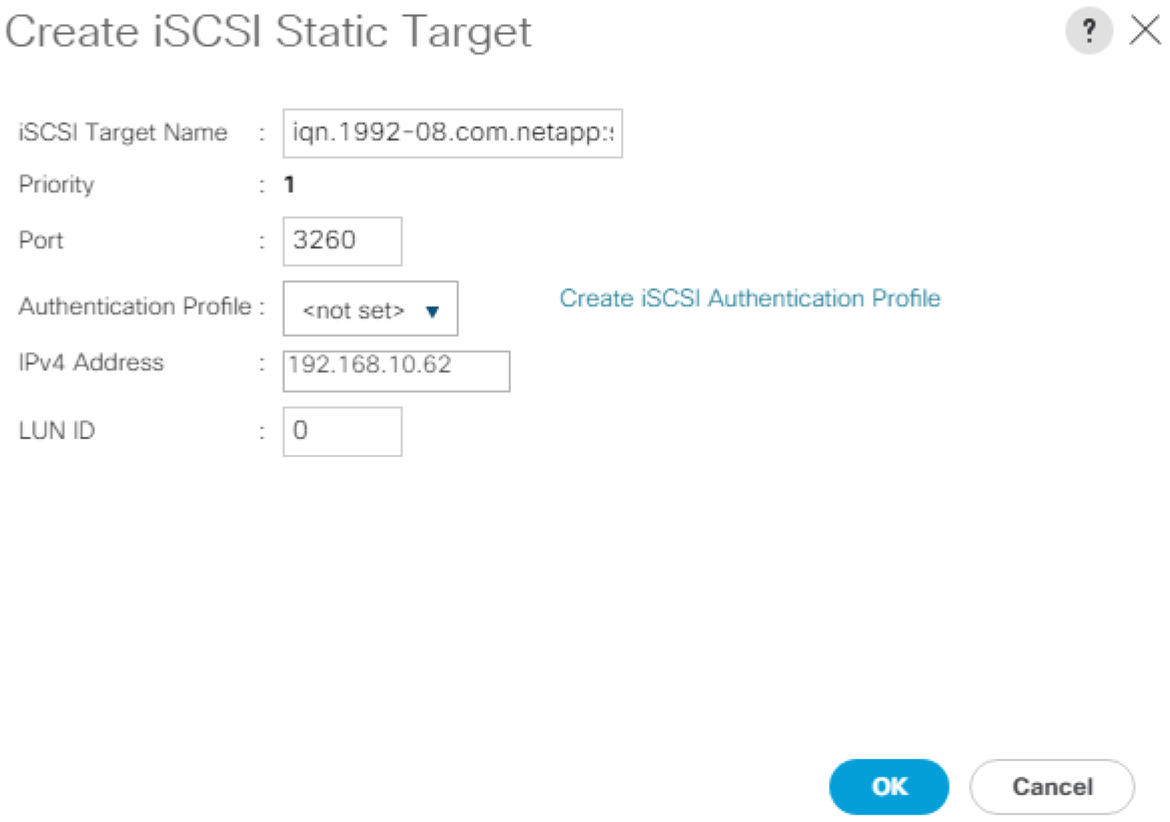

Create iSCSI Static Target

iSCSI Target Name : iqn.1992-08.com.netapp::

Priority : 1

Port : 3260

Authentication Profile : <not set>

IPv4 Address : 192.168.10.62

LUN ID : 0

8. Click OK to complete configuring the iSCSI target.

9. Click Add to add a second target.

10. Enter the IQN (Target Name) from the FPV-Foundation-SVM iSCSI Target Name. To get this IQN, ssh into the storage cluster interface and type "iscsi show".

11. For IPv4 address, enter the IP address of `iscsi_lif01a` from the FPV-Foundation-SVM. To get this IP, ssh into the storage cluster interface and type "network interface show –vserver FPV-Foundation-SVM".

12. Click OK to complete configuring the iSCSI target.

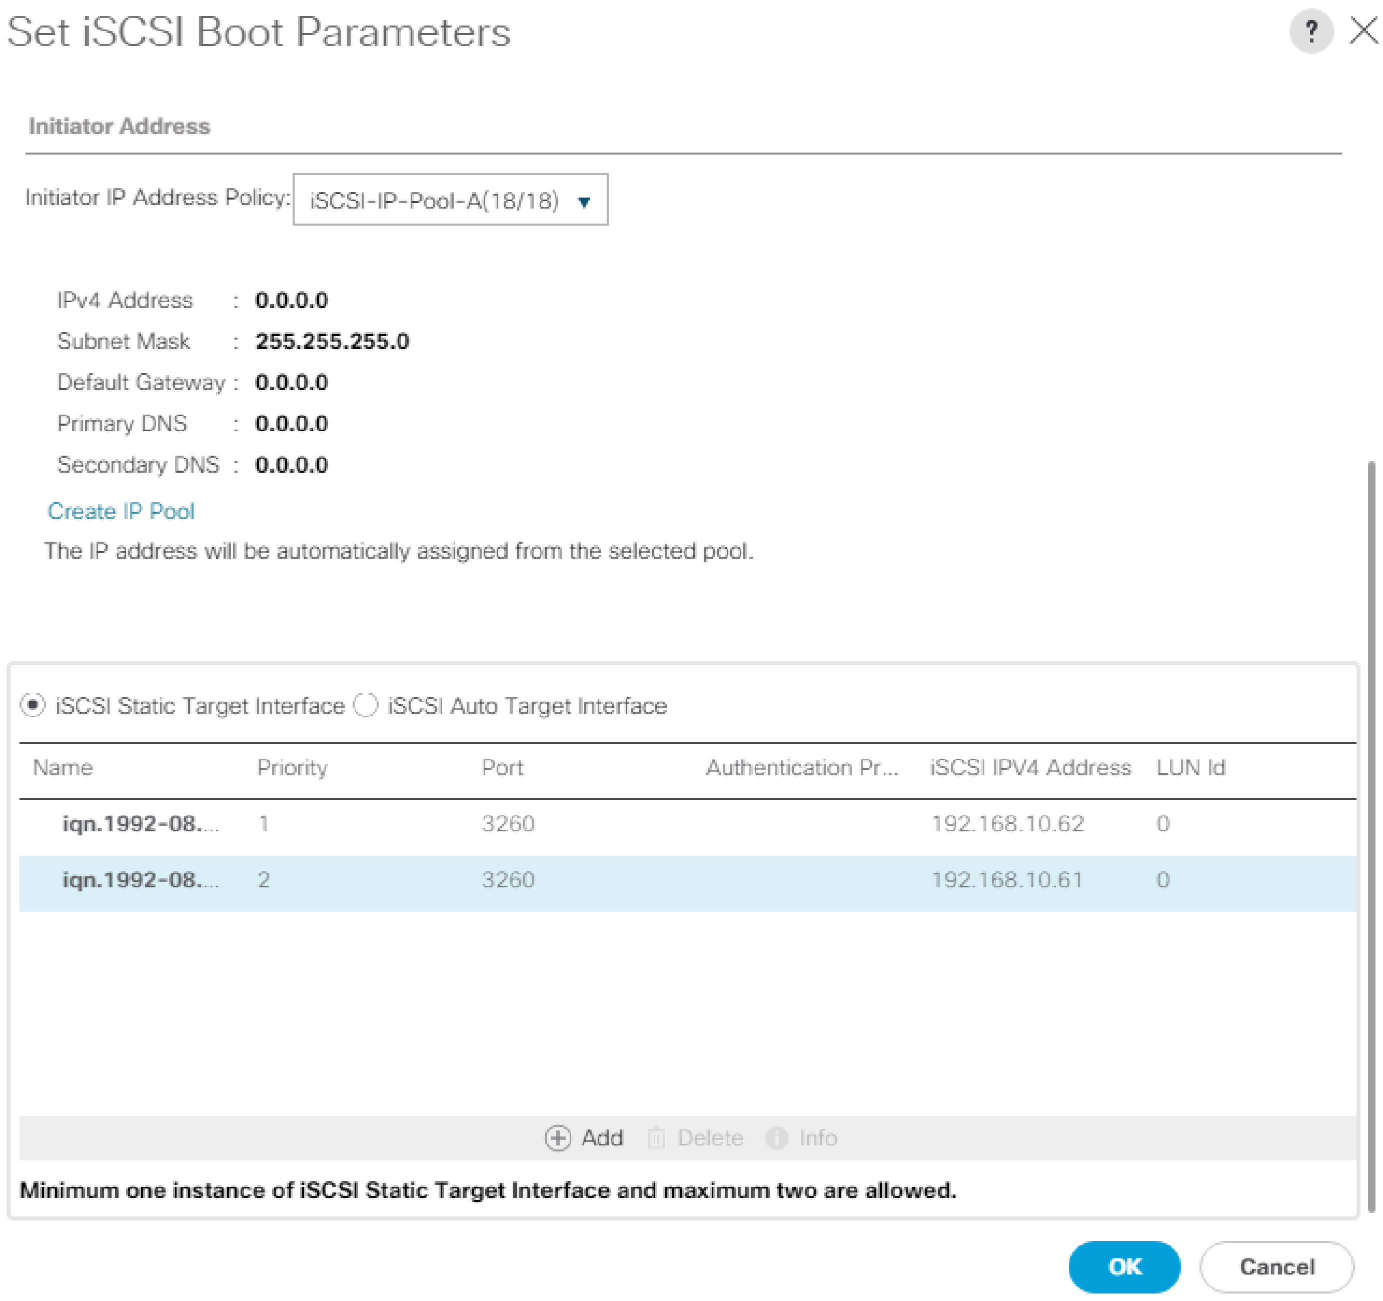

Set iSCSI Boot Parameters

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-A(18/18)

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 Address	LUN Id
iqn.1992-08....	1	3260		192.168.10.62	0
iqn.1992-08....	2	3260		192.168.10.61	0

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

13. Click OK to complete setting the iSCSI Boot Parameters for Fabric A Boot.

14. Under Boot Order, select the iSCSI-Boot-B row.

15. Select the Set iSCSI Boot Parameters button.

16. Select `iSCSI-IP-Pool-B` for the Initiator IP Address Policy.

17. Scroll to the bottom of the window and click Add.

18. Enter the IQN (Target Name) from the FPV-Foundation-SVM iSCSI Target Name. To get this IQN, ssh into the storage cluster interface and type “iscsi show.”

19. For IPv4 address, enter the IP address of `iscsi_lif02b` from the FPV-Foundation-SVM. To get this IP, ssh into the storage cluster interface and type “network interface show –vserver FPV-Foundation-SVM.”

20. Click OK to complete configuring the iSCSI target.

21. Click Add to add a second target.

22. Enter the IQN (Target Name) from the FPV-Foundation-SVM iSCSI Target Name. To get this IQN, ssh into the storage cluster interface and type “iscsi show.”

23. For IPv4 address, enter the IP address of `iscsi_lif01b` from the FPV-Foundation-SVM. To get this IP, ssh into the storage cluster interface and type “network interface show –vserver FPV-Foundation-SVM.”

24. Click OK to complete configuring the iSCSI target.

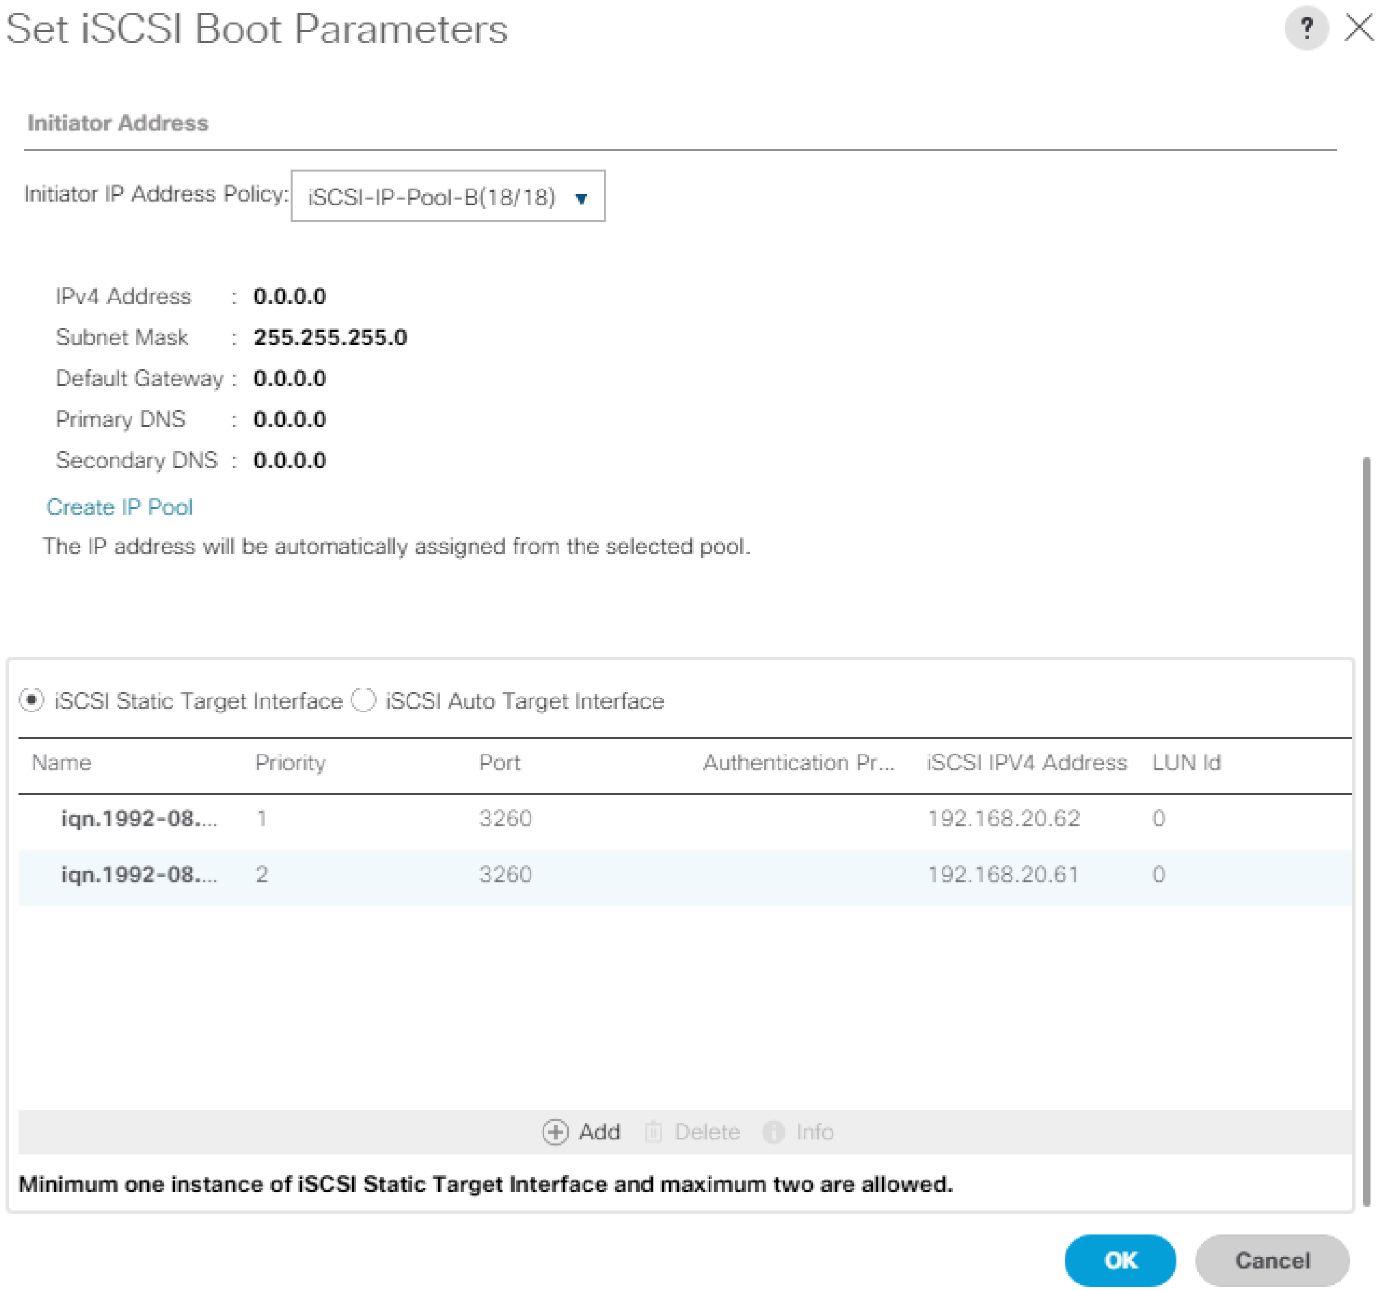

Set iSCSI Boot Parameters

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-B(18/18)

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 Address	LUN Id
iqn.1992-08....	1	3260		192.168.20.62	0
iqn.1992-08....	2	3260		192.168.20.61	0

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

25. Click OK to complete setting the iSCSI Boot Parameters for Fabric A Boot.

26. Click Next.

Configure Maintenance Policy

To configure the Maintenance Policy, complete the following steps:

1. Change the Maintenance Policy to default.

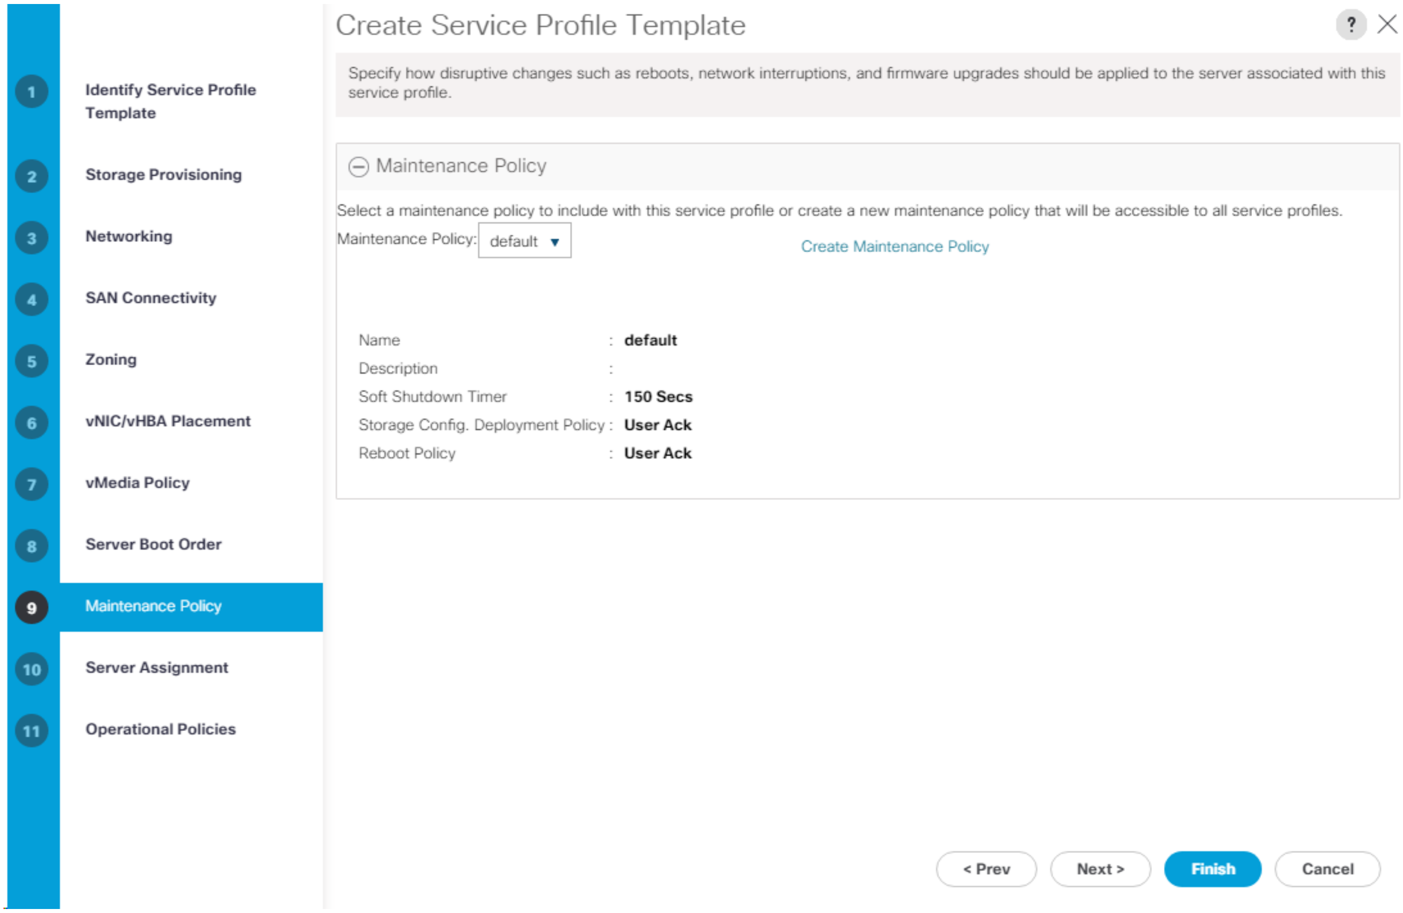

Create Service Profile Template

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

Maintenance Policy

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

Maintenance Policy: default ▾ Create Maintenance Policy

Name	: default
Description	:
Soft Shutdown Timer	: 150 Secs
Storage Config. Deployment Policy	: User Ack
Reboot Policy	: User Ack

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select `FPV-MGMT-Pool`.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select “`UCSB-B200-M5`” for the Server Pool Qualification.
4. Expand Firmware Management at the bottom of the page and select the default policy.

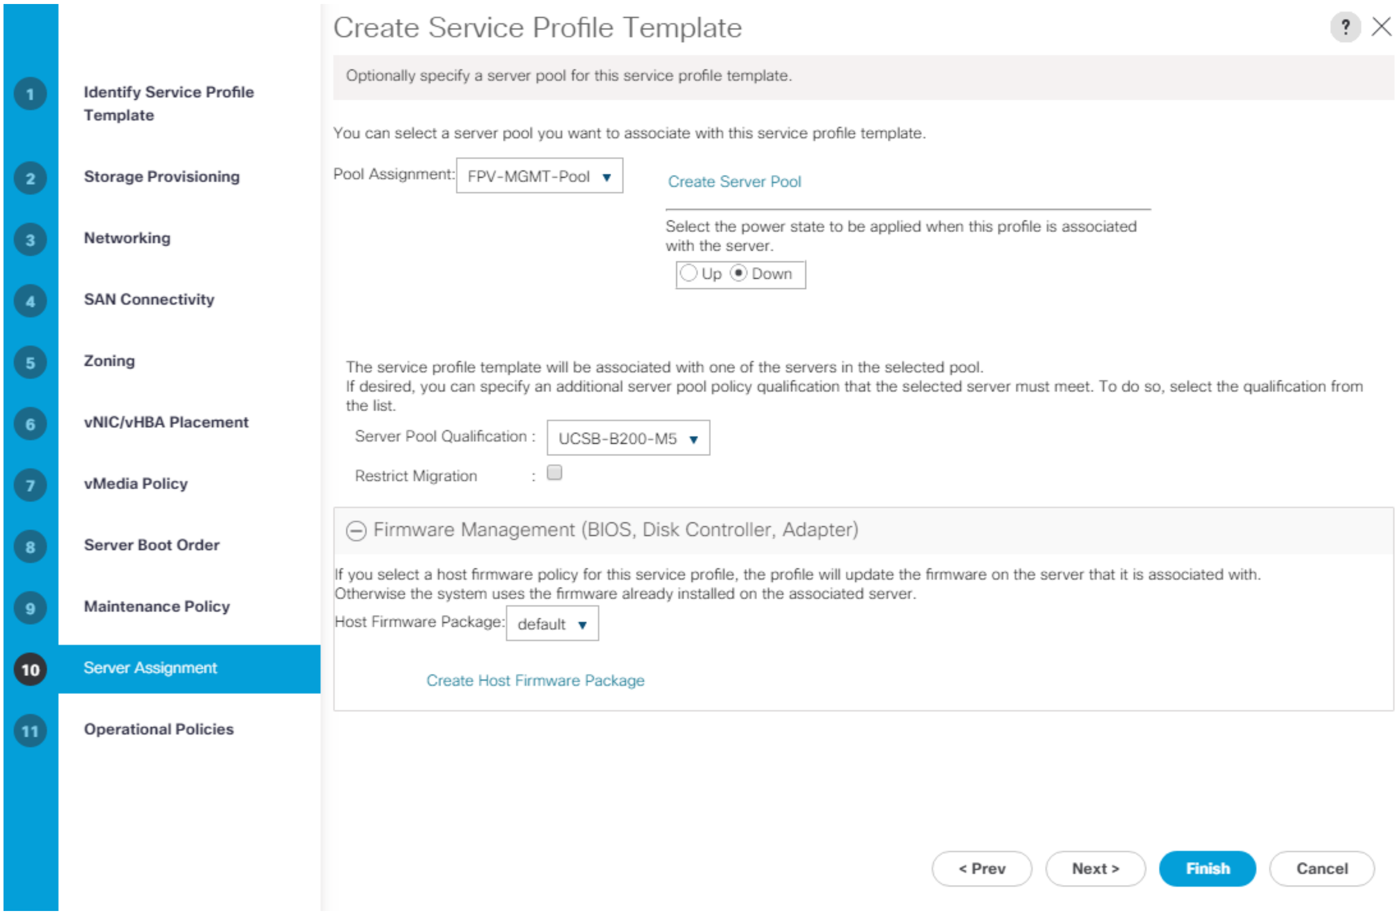

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: FPV-MGMT-Pool ▾

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 : UCSB-B200-M5 ▾

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: default ▾

Create Host Firmware Package

5. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. In the BIOS Policy list, select Virtual-Host.

2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

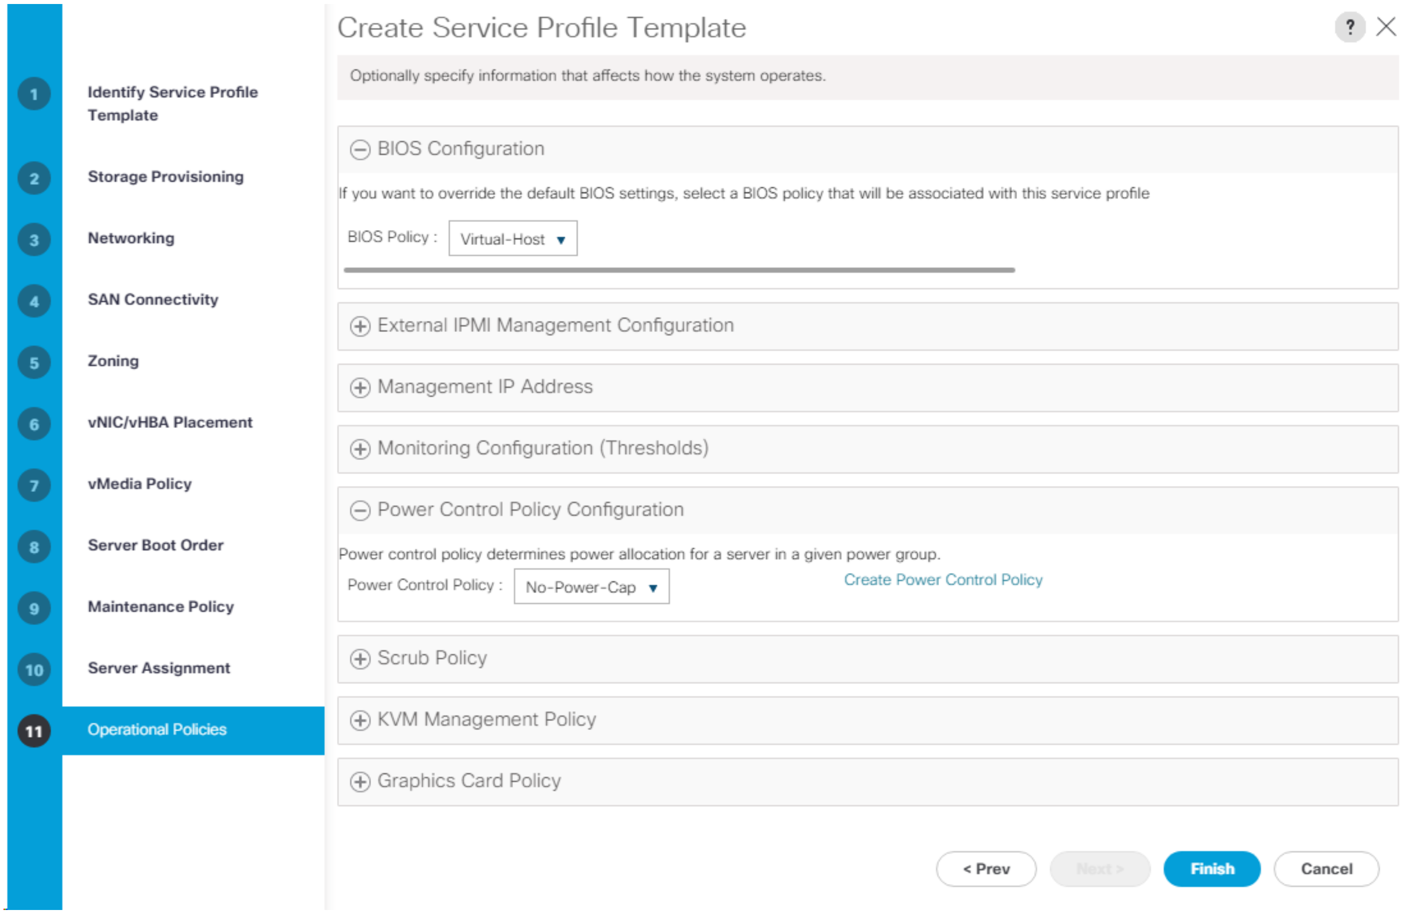

Create Service Profile Template

Optionally specify information that affects how the system operates.

BIOS Configuration

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

BIOS Policy : Virtual-Host

External IPMI Management Configuration
Management IP Address
Monitoring Configuration (Thresholds)

Power Control Policy Configuration

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

Power Control Policy : No-Power-Cap

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create Service Profiles

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

1. Connect to UCS Manager and click the Servers icon on the left.

2. Select Service Profile Templates > root > Sub-Organizations > FPV-FlexPod > Service Template iSCSI-Boot-A.

3. Right-click iSCSI-Boot-A and select Create Service Profiles from Template.

4. Enter `fpv-esxi-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.

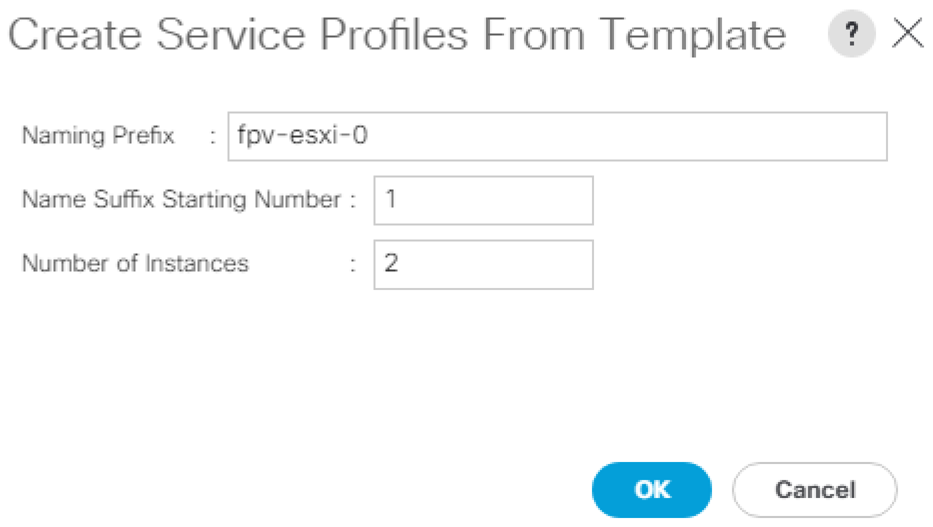

Create Service Profiles From Template

Naming Prefix : fpv-esxi-0

Name Suffix Starting Number : 1

Number of Instances : 2

8. Click OK in the confirmation message.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 2 and Table 3 .

Table 2 iSCSI LIFs for iSCSI IQN

SVM	Target: IQN
Foundation-FPV-SVM	

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

Table 3 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN	Variables
fpv-esxi-01		<fpv-esxi-01-iqn>
fpv-esxi-02		<fpv-esxi-02-iqn>

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

Storage Configuration – SAN Boot

NetApp ONTAP Boot Storage Setup

Create igroups

To create igroups, run the following commands:

```
igroup create -vserver FPV-Foundation-SVM -igroup fpv-esxi-01 -protocol iscsi -ostype vmware -
initiator <fpv-esxi-01-iqn>
igroup create -vserver FPV-Foundation-SVM -igroup fpv-esxi-02 -protocol iscsi -ostype vmware -
initiator <fpv-esxi-02-iqn>
igroup create -vserver FPV-Foundation-SVM -igroup MGMT-Hosts-All -protocol iscsi -ostype vmware -
initiator <fpv-esxi-01-iqn>,<fpv-esxi-02-iqn>
igroup show -vserver FPV-Foundation-SVM
```

1. To get the management host IQNs, log in to Cisco UCS Manager and click the Servers icon on the left.

2. Select the host Service Profile. The host IQN is listed under the iSCSI vNICs tab on the right.

Map Boot LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-01 -igroup fpv-esxi-01 -lun-id 0
lun map -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-02 -igroup fpv-esxi-02 -lun-id 0
lun show -vserver FPV-Foundation-SVM -m
```

VMware vSphere 6.5 Update 1 Setup

VMware ESXi 6.5 Update 1

This section provides detailed instructions for installing VMware ESXi 6.5 Update 1 in a FlexPod environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and mapped CD/DVD in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 6.5 Update 1

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

1. Click the following link: VMware vSphere Hypervisor (ESXi) 6.5U1.
2. You will need a user id and password on vmware.com to download this software.
3. Download the .iso file.

Log in to Cisco UCS 6300 Fabric Interconnect

Cisco UCS Manager

The Cisco UCS IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click Servers on the left.
7. Select the `fpv-esxi-01` Service Profile.
8. On the right, under the General tab, click the >> to the right of KVM Console.
9. Follow the prompts to launch the Java-based KVM console.
10. Select the `fpv-esxi-02` Service Profile.
11. On the right, under the General tab, click the >> to the right of KVM Console.

12. Follow the prompts to launch the Java-based KVM console.

Set Up VMware ESXi Installation

ESXi Hosts fpv-esxi-01 and fpv-esxi-02

Skip this section if you are using vMedia policies. ISO file will already be connected to KVM.

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.

Install ESXi

ESXi Hosts fpv-esxi-01 and fpv-esxi-02

To install VMware ESXi to the iSCSI-bootable LUN of the hosts, complete the following steps on each host:

1. Boot the server by selecting Boot Server and clicking OK, then click OK two more times.
2. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
5. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. After the installation is complete, press Enter to reboot the server. The mapped iso will be automatically unmapped.

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 steps in the following subsections.

ESXi Host fpv-esxi-01 and fpv-esxi-02

To configure each ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select Troubleshooting Options and press Enter.
4. Select Enable ESXi Shell and press Enter.
5. Select Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Select the Configure Management Network option and press Enter.
8. Select Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field.

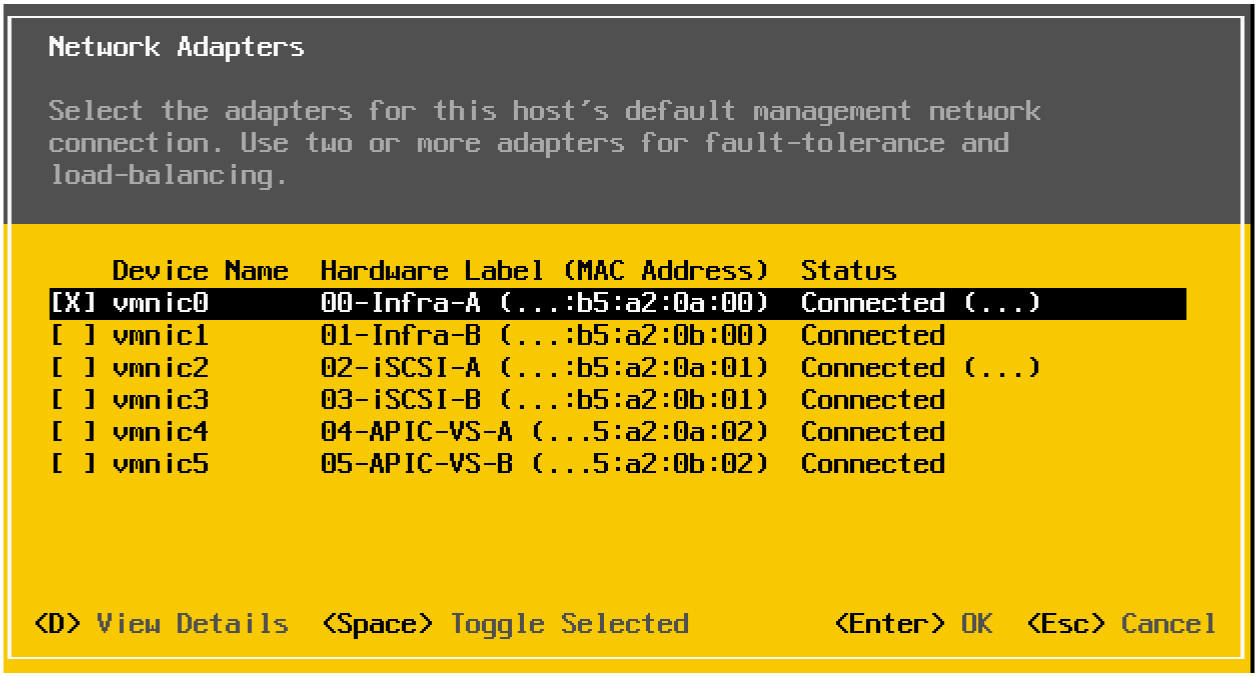

```
Network Adapters

Select the adapters for this host's default management network
connection. Use two or more adapters for fault-tolerance and
load-balancing.

    Device Name  Hardware Label (MAC Address)  Status
[X] vmnic0       00-Infra-A (...:b5:a2:0a:00)  Connected (...)
[ ] vmnic1       01-Infra-B (...:b5:a2:0b:00)  Connected
[ ] vmnic2       02-iSCSI-A (...:b5:a2:0a:01)  Connected (...)
[ ] vmnic3       03-iSCSI-B (...:b5:a2:0b:01)  Connected
[ ] vmnic4       04-APIC-VS-A (...5:a2:0a:02)  Connected
[ ] vmnic5       05-APIC-VS-B (...5:a2:0b:02)  Connected

<D> View Details  <Space> Toggle Selected      <Enter> OK  <Esc> Cancel
```

In lab testing, examples have been seen where the vmnic and device ordering do not match. If this is the case, use the Consistent Device Naming (CDN) to note which vmnics are mapped to which vNICs and adjust the up-coming procedure accordingly.

10. Use the arrow keys and spacebar to highlight and select vmnic1.

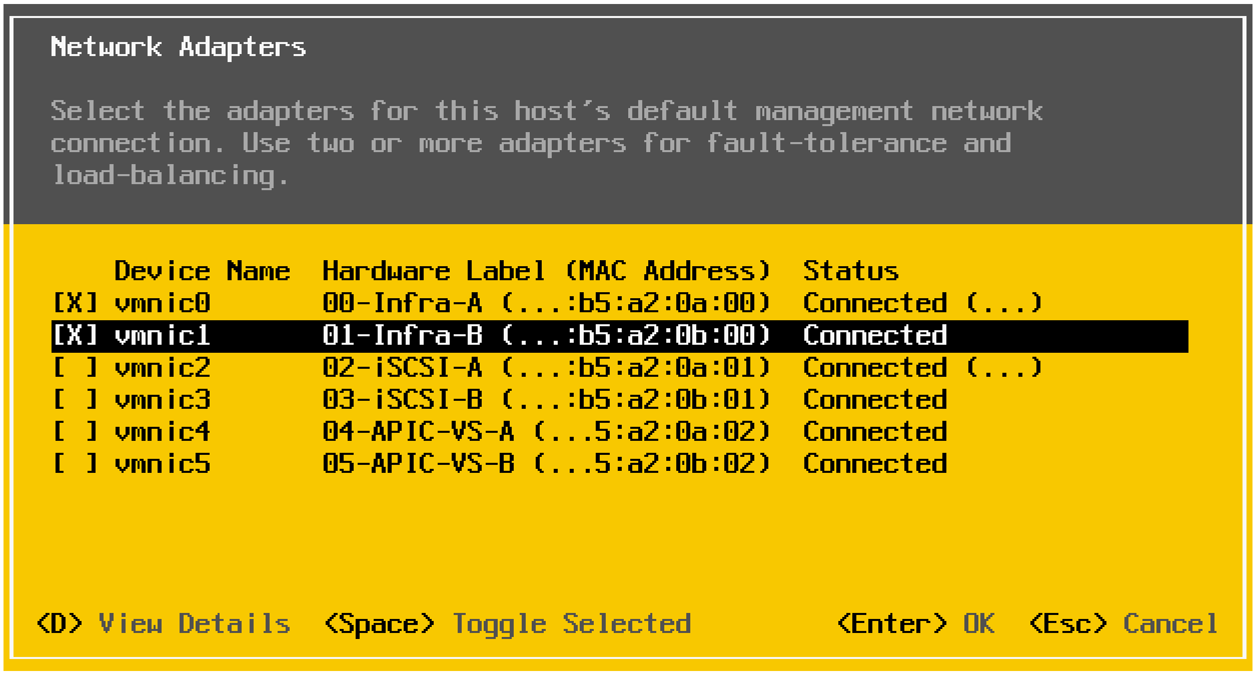

```
Network Adapters

Select the adapters for this host's default management network
connection. Use two or more adapters for fault-tolerance and
load-balancing.

    Device Name  Hardware Label (MAC Address)  Status
[X] vmnic0       00-Infra-A (...:b5:a2:0a:00)  Connected (...)
[X] vmnic1       01-Infra-B (...:b5:a2:0b:00)  Connected
[ ] vmnic2       02-iSCSI-A (...:b5:a2:0a:01)  Connected (...)
[ ] vmnic3       03-iSCSI-B (...:b5:a2:0b:01)  Connected
[ ] vmnic4       04-APIC-VS-A (...5:a2:0a:02)  Connected
[ ] vmnic5       05-APIC-VS-B (...5:a2:0b:02)  Connected

<D> View Details  <Space> Toggle Selected        <Enter> OK  <Esc> Cancel
```

11. Press Enter.

12. Select the VLAN (Optional) option and press Enter.

13. Enter the UCS `<ib-mgmt-vlan-id>` <419> and press Enter.

14. Select IPv4 Configuration and press Enter.

15. Select the Set static IPv4 address and network configuration option by using the space bar.

16. Enter the IP address for managing the ESXi host.

17. Enter the subnet mask for the ESXi host.

18. Enter the default gateway for the ESXi host.

19. Press Enter to accept the changes to the IP configuration.

20. Select the DNS Configuration option and press Enter.

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

21. Enter the IP address of the primary DNS server.

22. Optional: Enter the IP address of the secondary DNS server.

23. Enter the fully qualified domain name (FQDN) for the ESXi host.

24. Press Enter to accept the changes to the DNS configuration.

25. Press Esc to exit the Configure Management Network menu. Enter Y to Apply changes and restart management network.

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, press Enter again once the test has completed, review environment if there is a failure.

28. Re-select the Configure Management Network and press Enter.

29. Select the IPv6 Configuration option and press Enter.

30. Using the spacebar, select Disable IPv6 (restart required) and press Enter.

31. Press Esc to exit the Configure Management Network submenu.

32. Press Y to confirm the changes and reboot the ESXi host.

Log in to VMware ESXi Hosts by Using VMware Host Client

ESXi Host fpv-esxi-01

To log in to the `fpv-esxi-01` ESXi host by using the VMware Host Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the `fpv-esxi-01` management IP address. Respond to any security prompts.

2. Enter `root` for the user name.

3. Enter the root password.

4. Click Login to connect.

5. Repeat this process to log into `fpv-esxi-02` in a separate browser tab or window.

Set Up VMkernel Ports and Virtual Switch

ESXi Host fpv-esxi-01 and fpv-esxi-02

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

1. From the Host Client, select Networking on the left.

2. In the center pane, select the Virtual switches tab.

3. Highlight vSwitch0.

4. Select Edit settings.

5. Change the MTU to 9000.

6. Expand NIC teaming and highlight vmnic1. Select Mark active.

7. Click Save.

8. Select Networking on the left.

9. In the center pane, select the Virtual switches tab.

10. Highlight iScsiBootvSwitch.
11. Select Edit settings.
12. Change the MTU to 9000
13. Click Save.
14. Select Add standard virtual switch.
15. Name the vSwitch `iScsiBootvSwitch-B`.
16. Set the MTU to 9000.
17. Select vmnic3 for the Uplink.
18. Click Add.
19. Select the VMkernel NICs tab.
20. Highlight vmk1 `iScsiBootPG`.
21. Select Edit settings.
22. Change the MTU to 9000.
23. Expand IPv4 settings and change the IP address to an address outside of the UCS iSCSI-IP-Pool-A.

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

24. Click Save.
25. Select Add VMkernel NIC
26. Specify a New port group name of `iScsiBootPG-B`.
27. Select `iScsciBootvSwitch-B` for Virtual switch.
28. Set the MTU to 9000. Do not enter a VLAN ID since the iSCSI-B VLAN is also the native VLAN on this vNIC.
29. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask within the Configuration.

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

30. Click Create.
31. Select Add VMkernel NIC.
32. Specify a New port group name of `VMkernel-Infra-NFS`.

33. Select vSwitch0 for Virtual switch.

34. Enter the UCS Foundation Tenant NFS VLAN id <3150>.

35. Set the MTU to 9000.

36. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask in the Foundation Tenant NFS subnet <192.168.50.0/24>.

37. Click Create.

38. Select Add VMkernel NIC

39. Specify a New port group name of `VMkernel-vMotion`.

40. Select vSwitch0 for Virtual switch.

41. Enter the UCS vMotion VLAN id <3000>.

42. Set the MTU to 9000.

43. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask in the vMotion subnet <192.168.100.0/24>.

44. Select the vMotion stack for the TCP/IP stack.

45. Click Create.

46. Optionally, if you have 40GE vNICs in this FlexPod, create two more vMotion VMkernel ports in the same subnet and VLAN. These will need to be in new port groups.

47. On the left, select Networking, then select the Port groups tab.

48. In the center pane, right-click VM Network and select Remove.

49. Click Remove to complete removing the port group.

50. In the center pane, select Add port group.

51. Name the port group `IB-MGMT Network`, enter `<ib-mgmt-vlan-id>` <419> in the VLAN ID field, and make sure Virtual switch vSwitch0 is selected.

52. Click Add to finalize the edits for the IB-MGMT Network.

53. In the center pane, select Add port group.

54. Name the port group `Core-Services Network`, enter `<core-services-vlan-id>` <318> in the VLAN ID field, and make sure Virtual switch vSwitch0 is selected.

55. Click Add to finalize the edits for the Core-Services Network.

56. Highlight the VMkernel-vMotion Port group and select Edit settings.

57. Expand NIC teaming and select the radio button next to Override failover order.

58. Select vmnic0 and then select Mark standby to pin vMotion traffic to UCS Fabric Interconnect B (vmnic1) with failover.

59. Click Save.

60. Repeat steps 56-59 to pin all vMotion traffic to UCS Fabric Interconnect B (vmnic1) with failover.

61. Select the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:

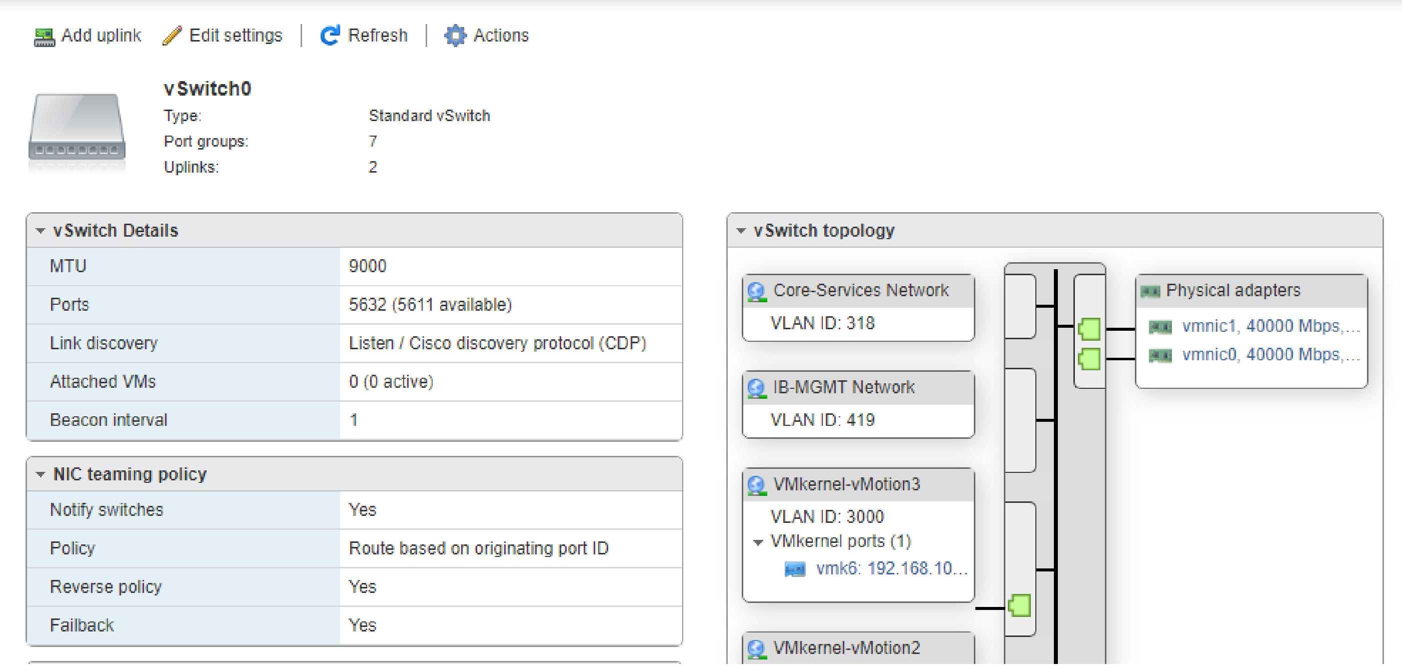

vSwitch0

Type:	Standard vSwitch
Port groups:	7
Uplinks:	2

vSwitch Details

MTU	9000
Ports	5632 (5611 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

NIC teaming policy

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

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

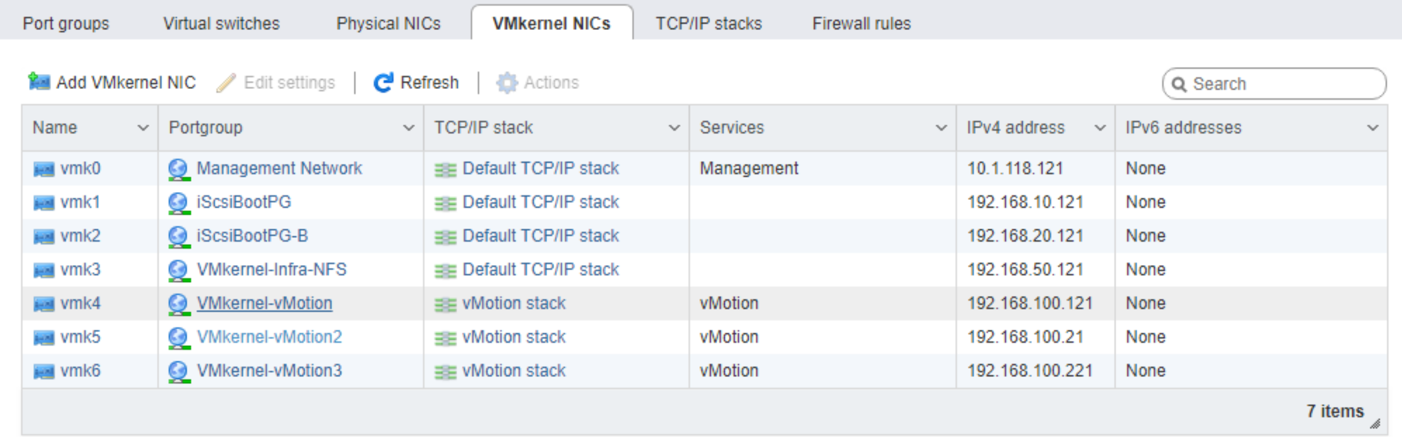

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.118.121	None
vmk1	iScsiBootPG	Default TCP/IP stack		192.168.10.121	None
vmk2	iScsiBootPG-B	Default TCP/IP stack		192.168.20.121	None
vmk3	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.121	None
vmk4	VMkernel-vMotion	vMotion stack	vMotion	192.168.100.121	None
vmk5	VMkernel-vMotion2	vMotion stack	vMotion	192.168.100.21	None
vmk6	VMkernel-vMotion3	vMotion stack	vMotion	192.168.100.221	None

7 items

Setup iSCSI Multipathing

ESXi Hosts fpv-esxi-01 and fpv-esxi-02

To setup the iSCSI multipathing on the ESXi host fpv-esxi-01 and fpv-esxi-02, complete the following steps:

1. From each Host Client, select Storage on the left.
2. In the center pane, click Adapters.
3. Select the iSCSI software adapter and click Configure iSCSI.
4. Under Dynamic targets, click Add dynamic target.
5. Enter the IP Address of NetApp storage iscsi_lif01a and press Enter.
6. Repeat putting the ip address of iscsi_lif01b, iscsi_lif02a, iscsi_lif02b.
7. Click Save configuration.

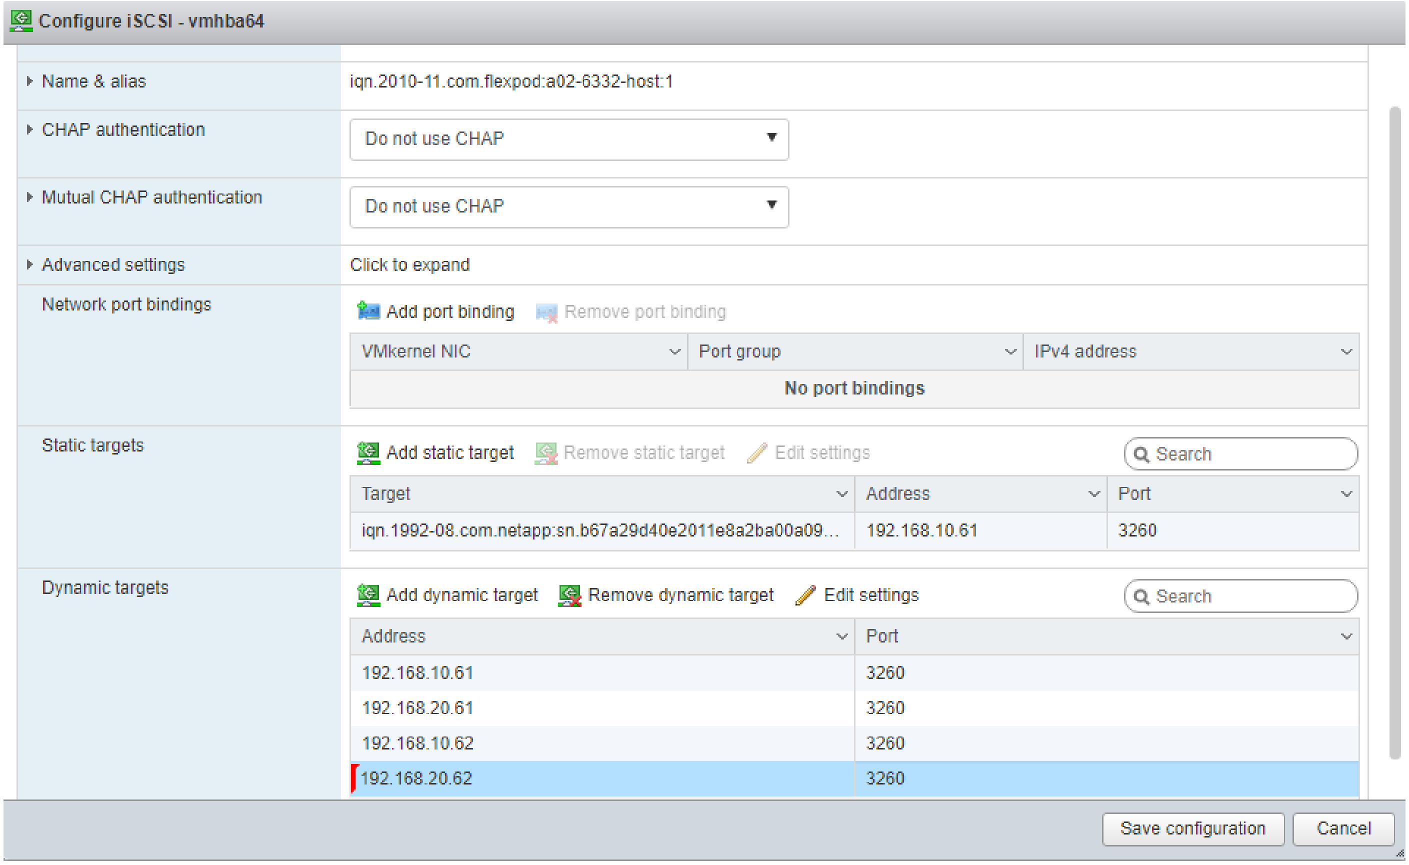

Configure iSCSI - vmhba64

Name & alias: iqn.2010-11.com.flexpod:a02-6332-host:1

CHAP authentication: Do not use CHAP

Mutual CHAP authentication: Do not use CHAP

Advanced settings: Click to expand

Network port bindings

VMkernel NIC	Port group	IPv4 address
No port bindings		

Static targets

Target	Address	Port
iqn.1992-08.com.netapp:sn.b67a29d40e2011e8a2ba00a09...	192.168.10.61	3260

Dynamic targets

Address	Port
192.168.10.61	3260
192.168.20.61	3260
192.168.10.62	3260
192.168.20.62	3260

To get all the iscsi_lif IP addresses, login to NetApp storage cluster management interface and type "network interface show."

The host will automatically rescan the storage adapter and 4 targets will be added to Static targets. This can be verified by selecting Configure iSCSI again.

Mount Required Datastores

ESXi Hosts fpv-esxi-01 and fpv-esxi-02

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

1. From the Host Client, select Storage on the left.

2. In the center pane, select the Datastores tab.

3. In the center pane, select New Datastore to add a new datastore.

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

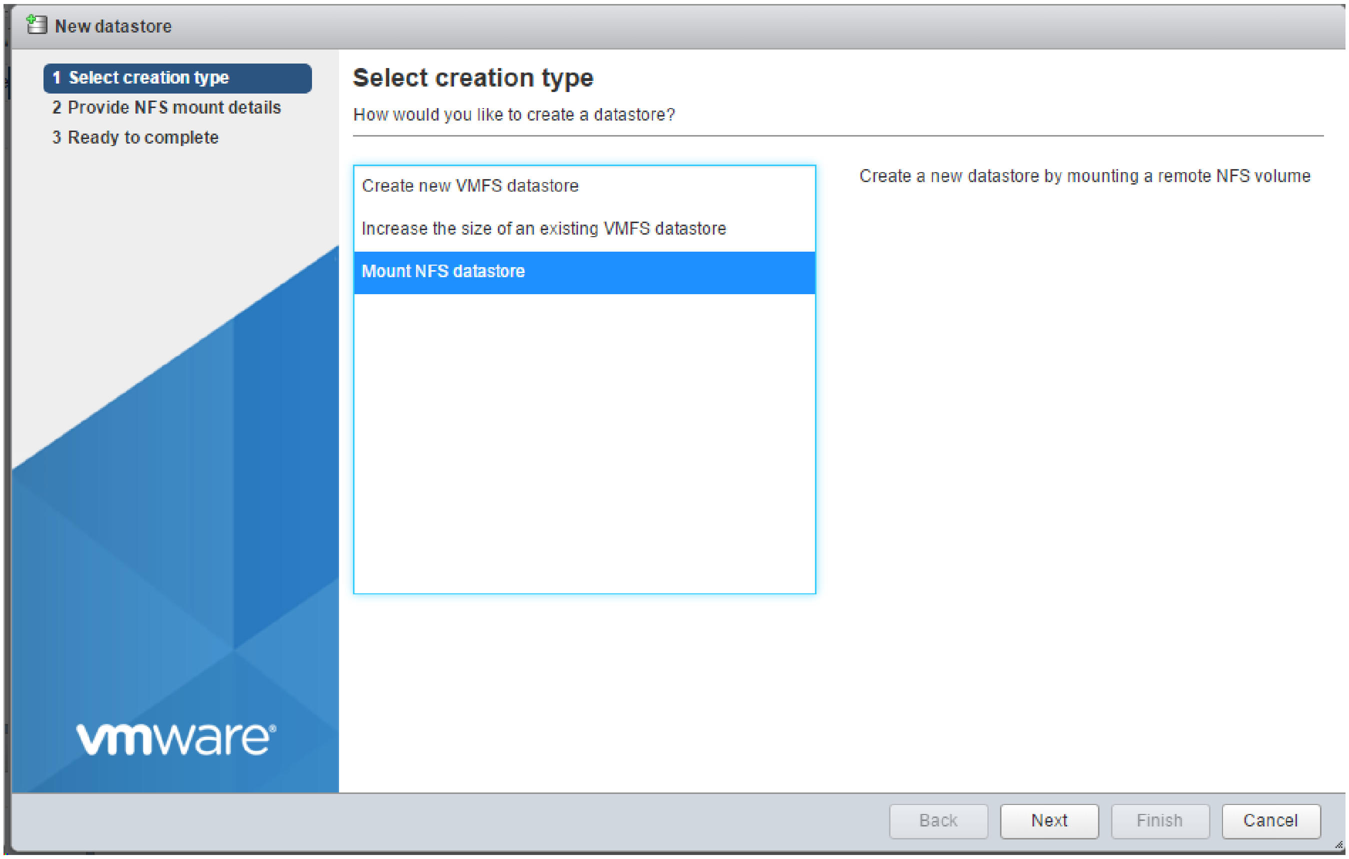

Select creation type

How would you like to create a datastore?

Create new VMFS datastore
Increase the size of an existing VMFS datastore
Mount NFS datastore

Create a new datastore by mounting a remote NFS volume

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

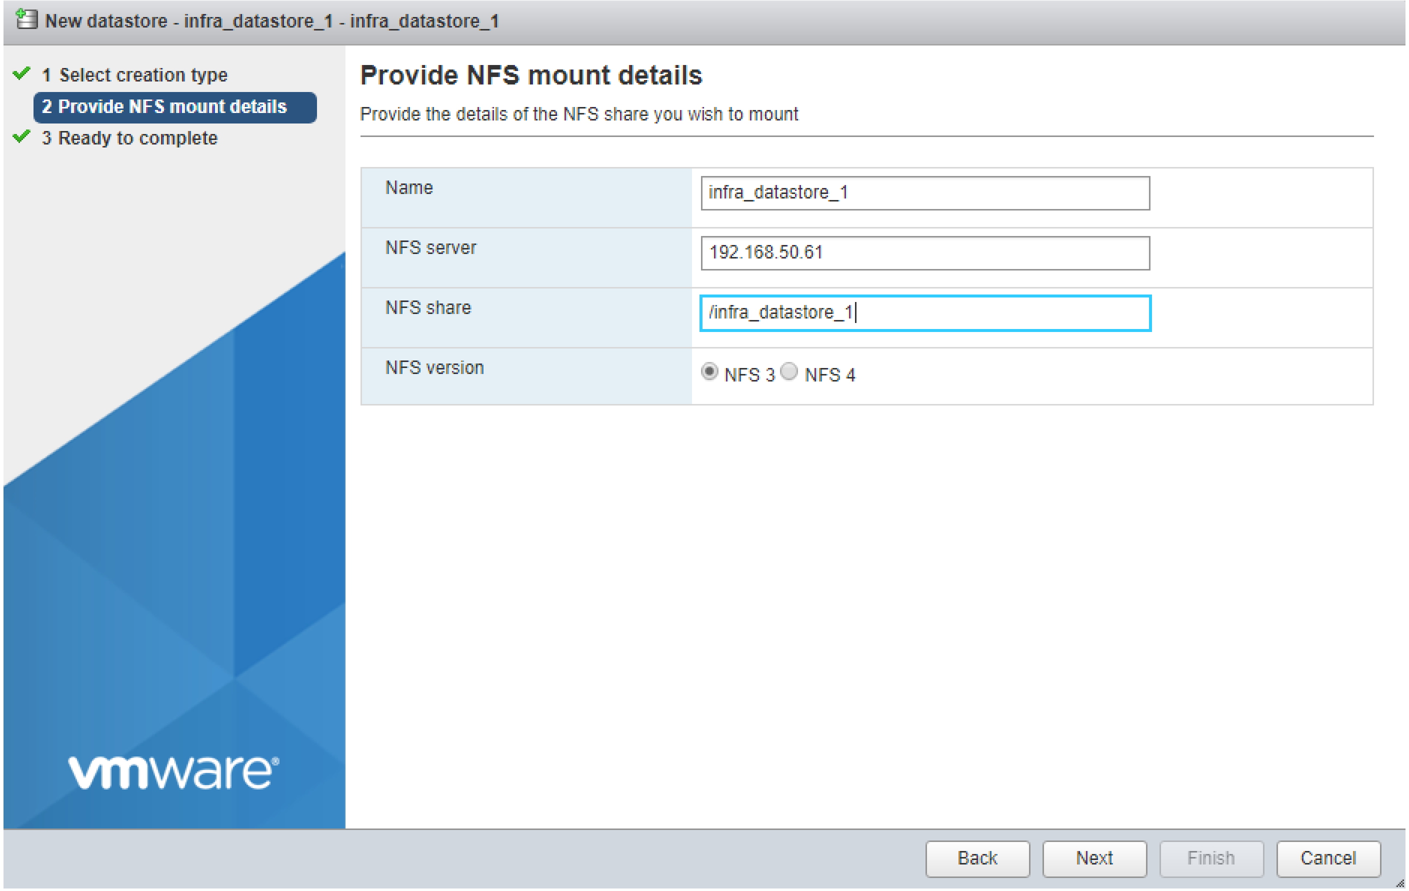

New datastore - infra_datastore_1 - infra_datastore_1

Provide NFS mount details

Provide the details of the NFS share you wish to mount

Name	infra_datastore_1
NFS server	192.168.50.61
NFS share	/infra_datastore_1
NFS version	● NFS 3 ○ NFS 4

6. Click Finish. The datastore should now appear in the datastore list.

7. In the center pane, select New Datastore to add a new datastore.

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

9. Input infra_datastore_2 for the datastore name. Input the IP address for the `nfs_lif02` LIF for the NFS server. Input /infra_datastore_2 for the NFS share. Leave the NFS version set at NFS 3. Click Next.

10. Click Finish. The datastore should now appear in the datastore list.

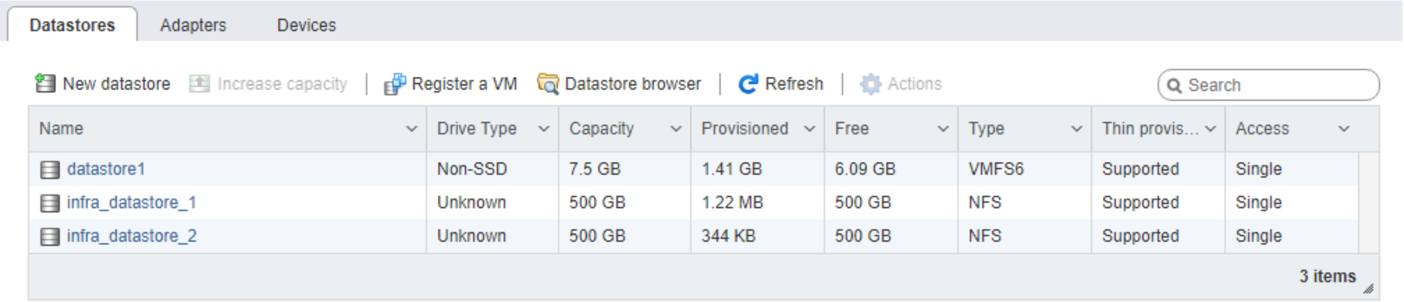

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provis...	Access
datastore1	Non-SSD	7.5 GB	1.41 GB	6.09 GB	VMFS6	Supported	Single
infra_datastore_1	Unknown	500 GB	1.22 MB	500 GB	NFS	Supported	Single
infra_datastore_2	Unknown	500 GB	344 KB	500 GB	NFS	Supported	Single

3 items

11. Mount both datastores on both ESXi hosts.

Configure NTP on ESXi Hosts

ESXi Hosts fpv-esxi-01 and fpv-esxi-02

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From the Host Client, select Manage on the left.

2. In the center pane, select the Time & Date tab.

3. Click Edit settings.

4. Make sure Use Network Time Protocol (enable NTP client) is selected.

5. Use the drop-down list to select Start and stop with host.

6. Enter the two ACI leaf switch management and NTP addresses in the NTP servers box separated by a comma.

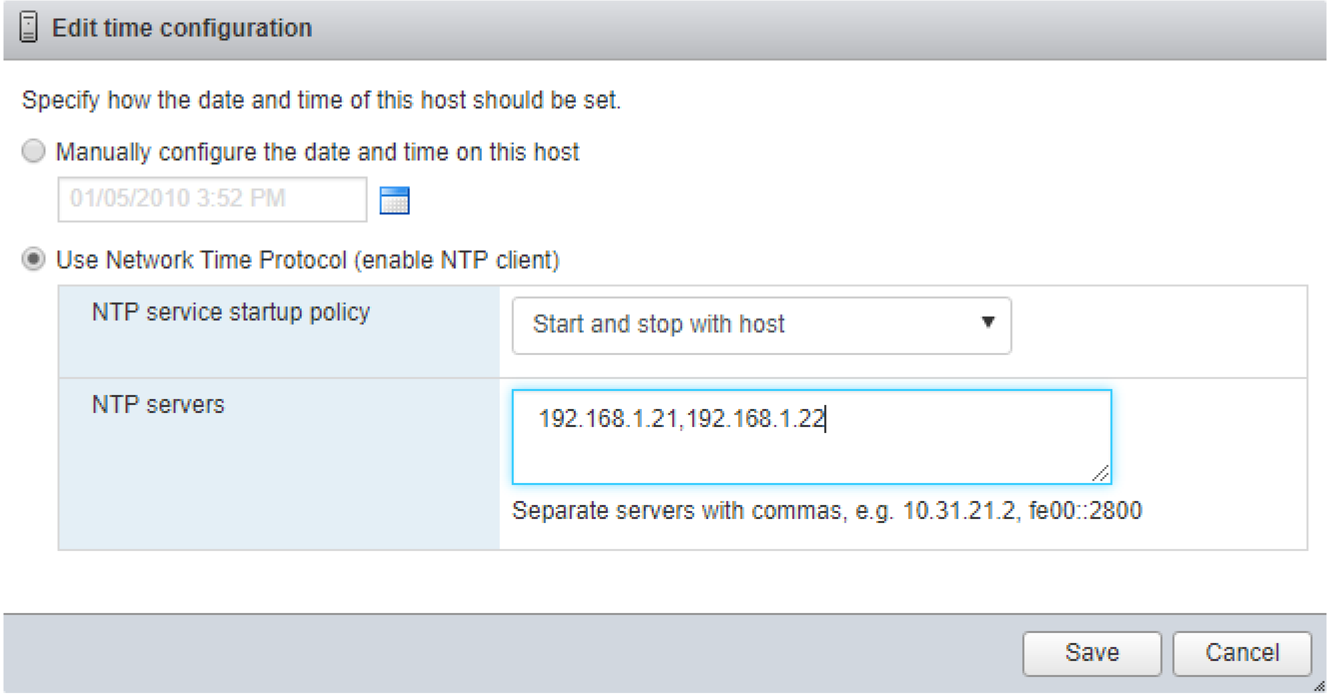

Edit time configuration

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

○ Manually configure the date and time on this host

01/05/2010 3:52 PM

● Use Network Time Protocol (enable NTP client)

NTP service startup policy	Start and stop with host
NTP servers	192.168.1.21,192.168.1.22 Separate servers with commas, e.g. 10.31.21.2, fe00::2800

7. Click Save to save the configuration changes.

8. Select Actions > NTP service > Start.

9. Verify that NTP service is now running and the clock is now set to approximately the correct time.

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

Install VMware ESXi Patches and Drivers

ESXi Hosts fpv-esxi-01 and fpv-esxi-02

To install necessary ESXi patches and updated device drivers, complete the following steps on each host:

1. Download VMware ESXi patch ESXi650-201712001 from https://my.vmware.com/group/vmware/patch#search.

2. Download the Cisco UCS VIC nenic driver version 1.0.13.0 from https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI65-CISCO-NENIC-10130&productId=614. Extract the offline_bundle.zip file from the download.

3. Download the Cisco UCS VIC fnic driver version 1.6.0.36 from https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESX60-CISCO-FNIC-16036&productId=491. Extract the offline_bundle.zip file from the download.

4. Download the NetApp NFS Plug-in for VMware VAAI 1.1.2 offline bundle from https://mysupport.netapp.com/NOW/download/software/nfs_plugin_vaai_esxi5.5/1.1.2/.

5. From the Host Client, select Storage on the left.

6. In the center pane, select infra_datastore_1 and then select Datastore browser.

7. In the Datastore browser, select Create directory and create a Drivers folder in the datastore.

8. Select the Drivers folder.

9. Use Upload to upload the four downloaded items above to the Drivers folder on infra_datastore_1. Since infra_datastore_1 is accessible to both ESXi hosts, this upload only needs to be done on the first host.

10. Use ssh to connect to each ESXi host as the root user.

11. Enter the following commands on each host.

```
cd /vmfs/volumes/infra_datastore_1/Drivers

ls

esxcli software vib update -d /vmfs/volumes/infra_datastore_1/Drivers/ESXi650-201712001.zip

esxcli software vib update -d /vmfs/volumes/infra_datastore_1/Drivers/VMW-ESX-6.5.0-nenic-1.0.13.0-
offline_bundle-7098243.zip

esxcli software vib update -d /vmfs/volumes/infra_datastore_1/Drivers/fnic_driver_1.6.0.36-
offline_bundle-6806699.zip

esxcli software vib install -d /vmfs/volumes/infra_datastore_1/Drivers/NetAppNasPlugin.v23.zip

reboot
```

VMware vCenter 6.5 Update 1

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.5 Update 1 Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Building the VMware vCenter Server Appliance

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

1. Locate and copy the VMware-VCSA-all-6.5.0-7515524.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 6.5 vCenter Server Appliance.

2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012).

3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appears.

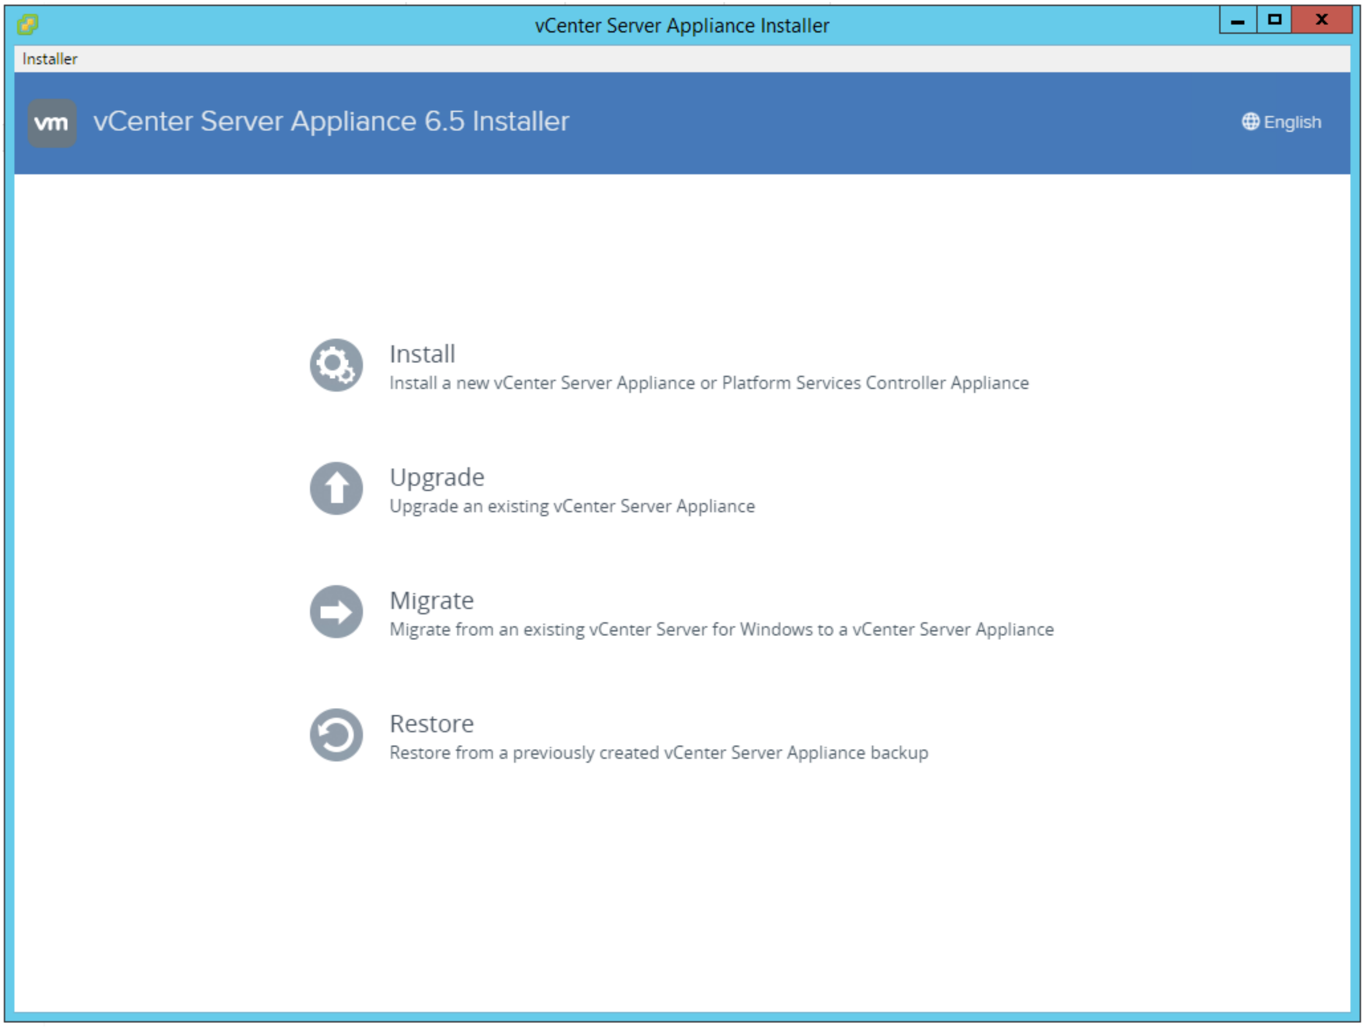
vCenter Server Appliance Installer

Installer

vCenter Server Appliance 6.5 Installer

Install
Install a new vCenter Server Appliance or Platform Services Controller Appliance

Upgrade
Upgrade an existing vCenter Server Appliance

Migrate
Migrate from an existing vCenter Server for Windows to a vCenter Server Appliance

Restore
Restore from a previously created vCenter Server Appliance backup

4. Click Install to start the vCenter Server Appliance deployment wizard.

5. Click Next in the Introduction section.

6. Read and accept the license agreement and click Next.

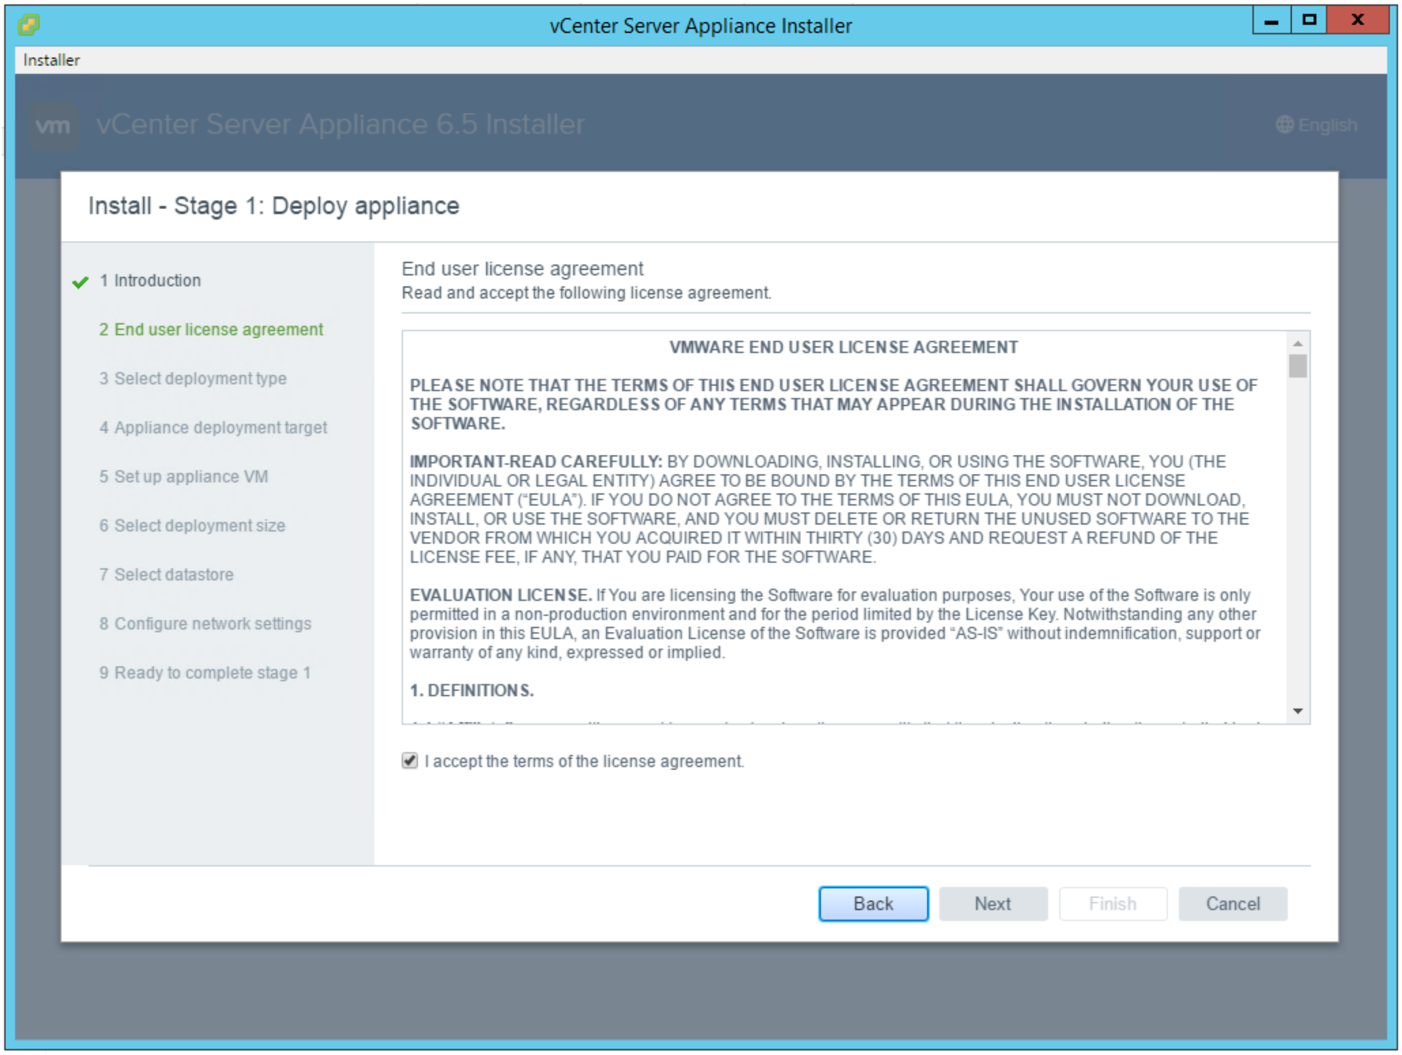

vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
Install - Stage 1: Deploy appliance
✓ 1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

End user license agreement

Read and accept the following license agreement.

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE VENDOR FROM WHICH YOU ACQUIRED IT WITHIN THIRTY (30) DAYS AND REQUEST A REFUND OF THE LICENSE FEE, IF ANY, THAT YOU PAID FOR THE SOFTWARE.

EVALUATION LICENSE. If You are licensing the Software for evaluation purposes, Your use of the Software is only permitted in a non-production environment and for the period limited by the License Key. Notwithstanding any other provision in this EULA, an Evaluation License of the Software is provided "AS-IS" without indemnification, support or warranty of any kind, expressed or implied.

1. DEFINITIONS.

☑ I accept the terms of the license agreement.

7. In the "Select deployment type" section, select vCenter Server with an Embedded Platform Services Controller and click Next.

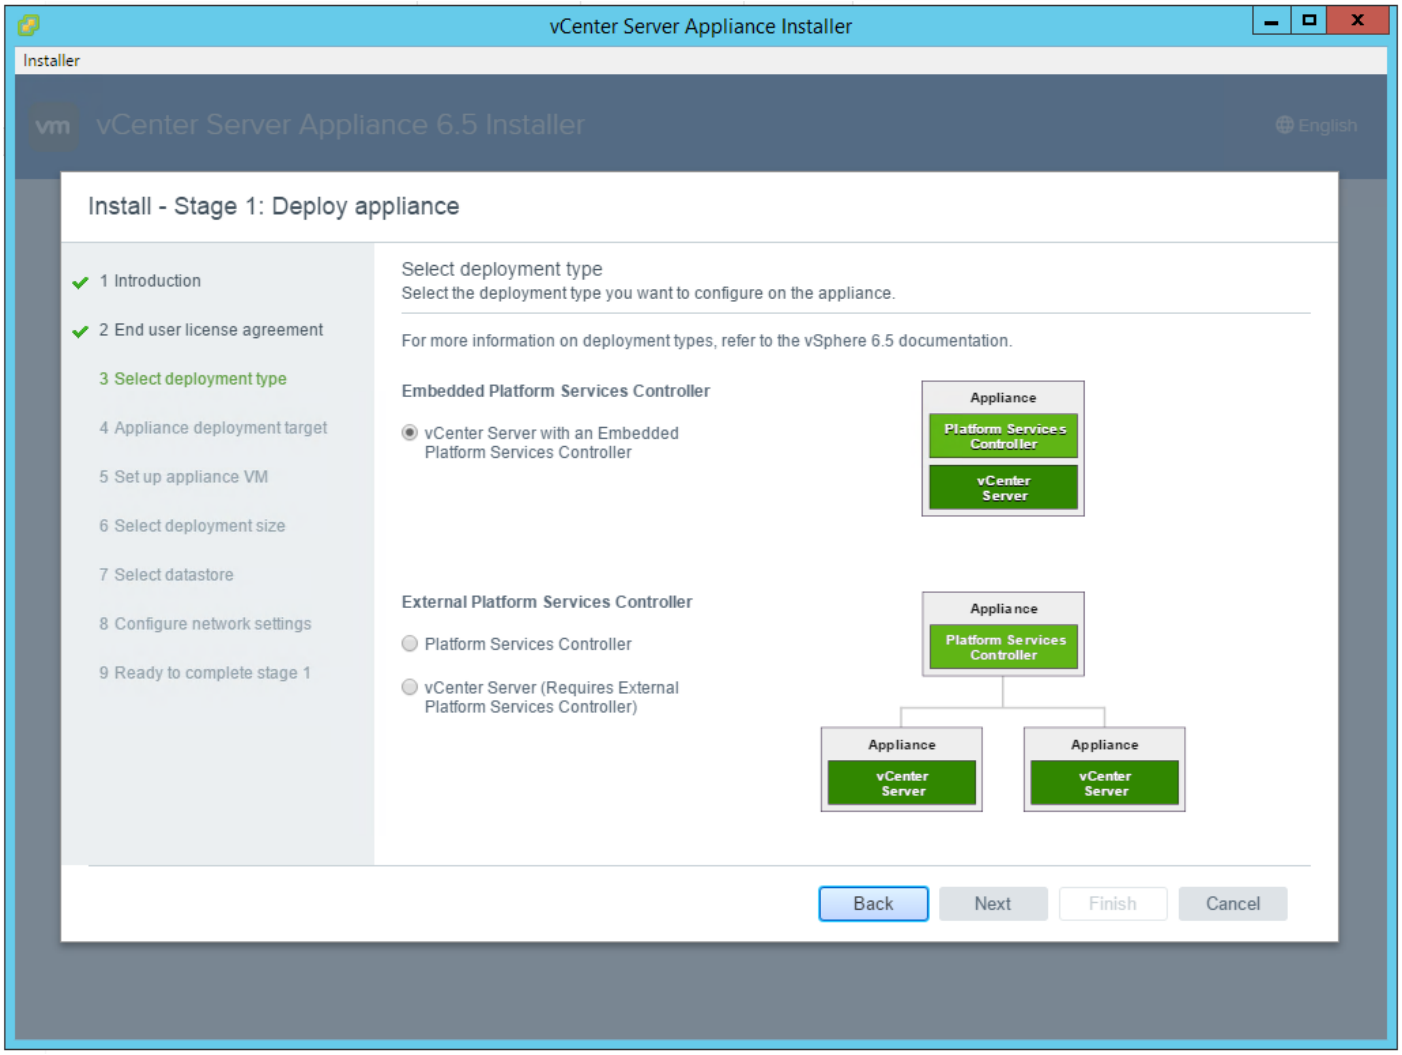
vCenter Server Appliance Installer
Installer
vm
vCenter Server Appliance 6.5 Installer
English
Appliance
Platform Services Controller
vCenter Server
Appliance
Platform Services Controller
Appliance
vCenter Server
Appliance
vCenter Server
Back
Next
Finish
Cancel

Install - Stage 1: Deploy appliance

- ✓ 1 Introduction
- ✓ 2 End user license agreement
- 3 Select deployment type
- 4 Appliance deployment target
- 5 Set up appliance VM
- 6 Select deployment size
- 7 Select datastore
- 8 Configure network settings
- 9 Ready to complete stage 1

Select deployment type

Select the deployment type you want to configure on the appliance.

For more information on deployment types, refer to the vSphere 6.5 documentation.

Embedded Platform Services Controller

● vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

○ Platform Services Controller

○ vCenter Server (Requires External Platform Services Controller)

8. In the “Appliance deployment target”, enter the ESXi host name or IP address for fpv-esxi-01, User name and Password. Click Next.

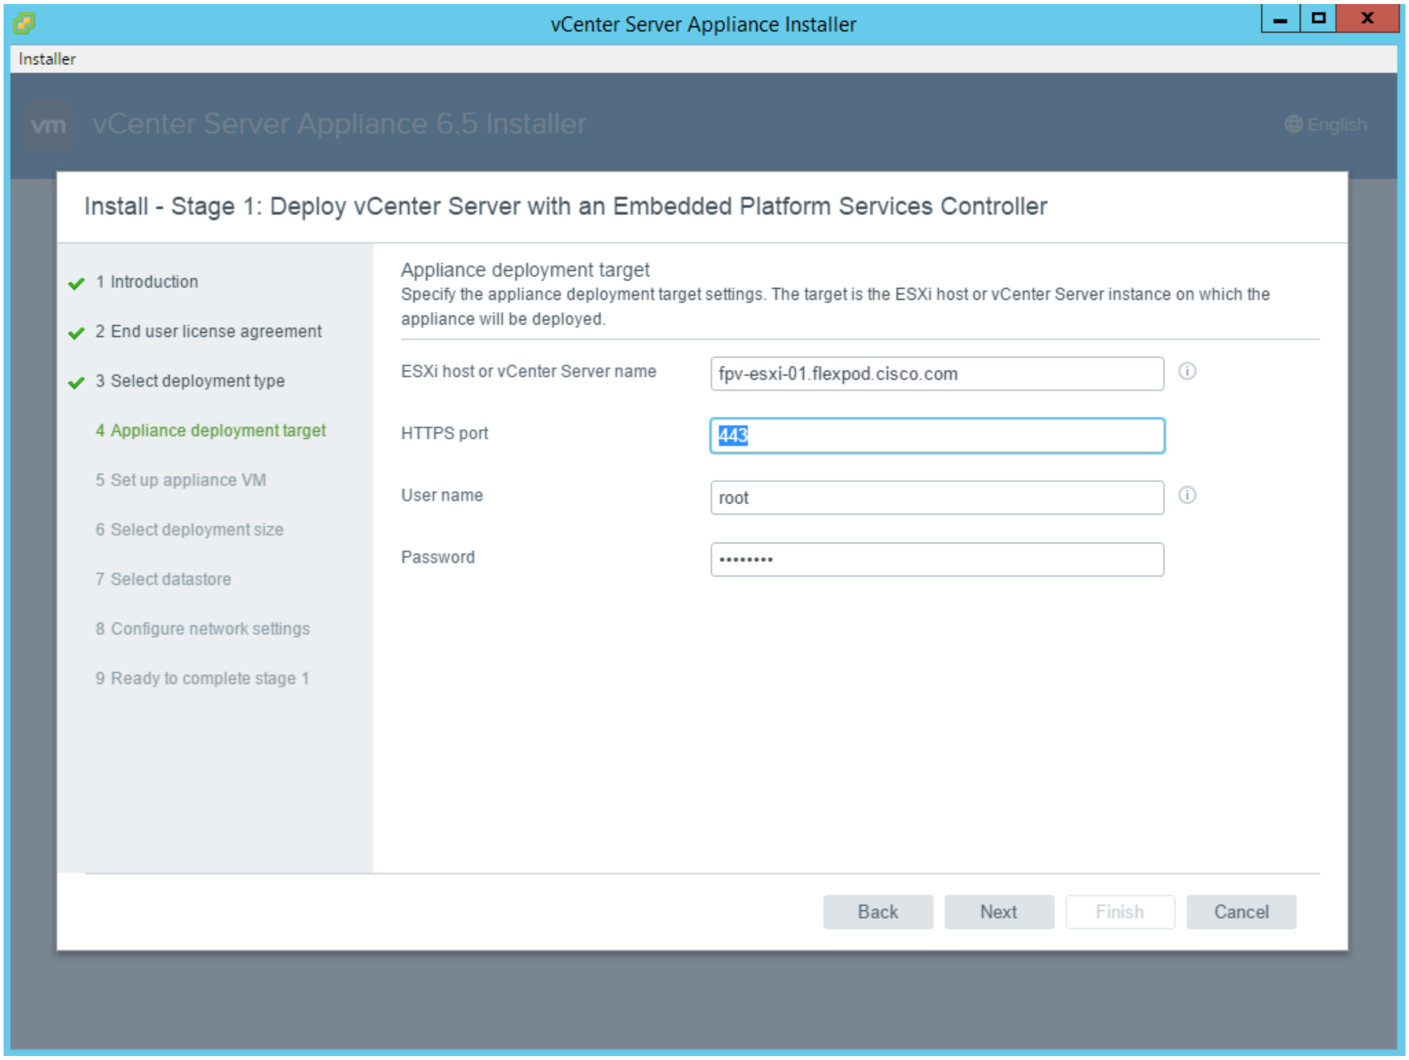

vCenter Server Appliance Installer
Installer
vCenter Server Appliance 6.5 Installer
English

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

- ✓ 1 Introduction
- ✓ 2 End user license agreement
- ✓ 3 Select deployment type
- 4 Appliance deployment target
- 5 Set up appliance VM
- 6 Select deployment size
- 7 Select datastore
- 8 Configure network settings
- 9 Ready to complete stage 1

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

ESXi host or vCenter Server name	fpv-esxi-01.flexpod.cisco.com
HTTPS port	443
User name	root
Password	••••••••

9. Click Yes to accept the certificate.

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

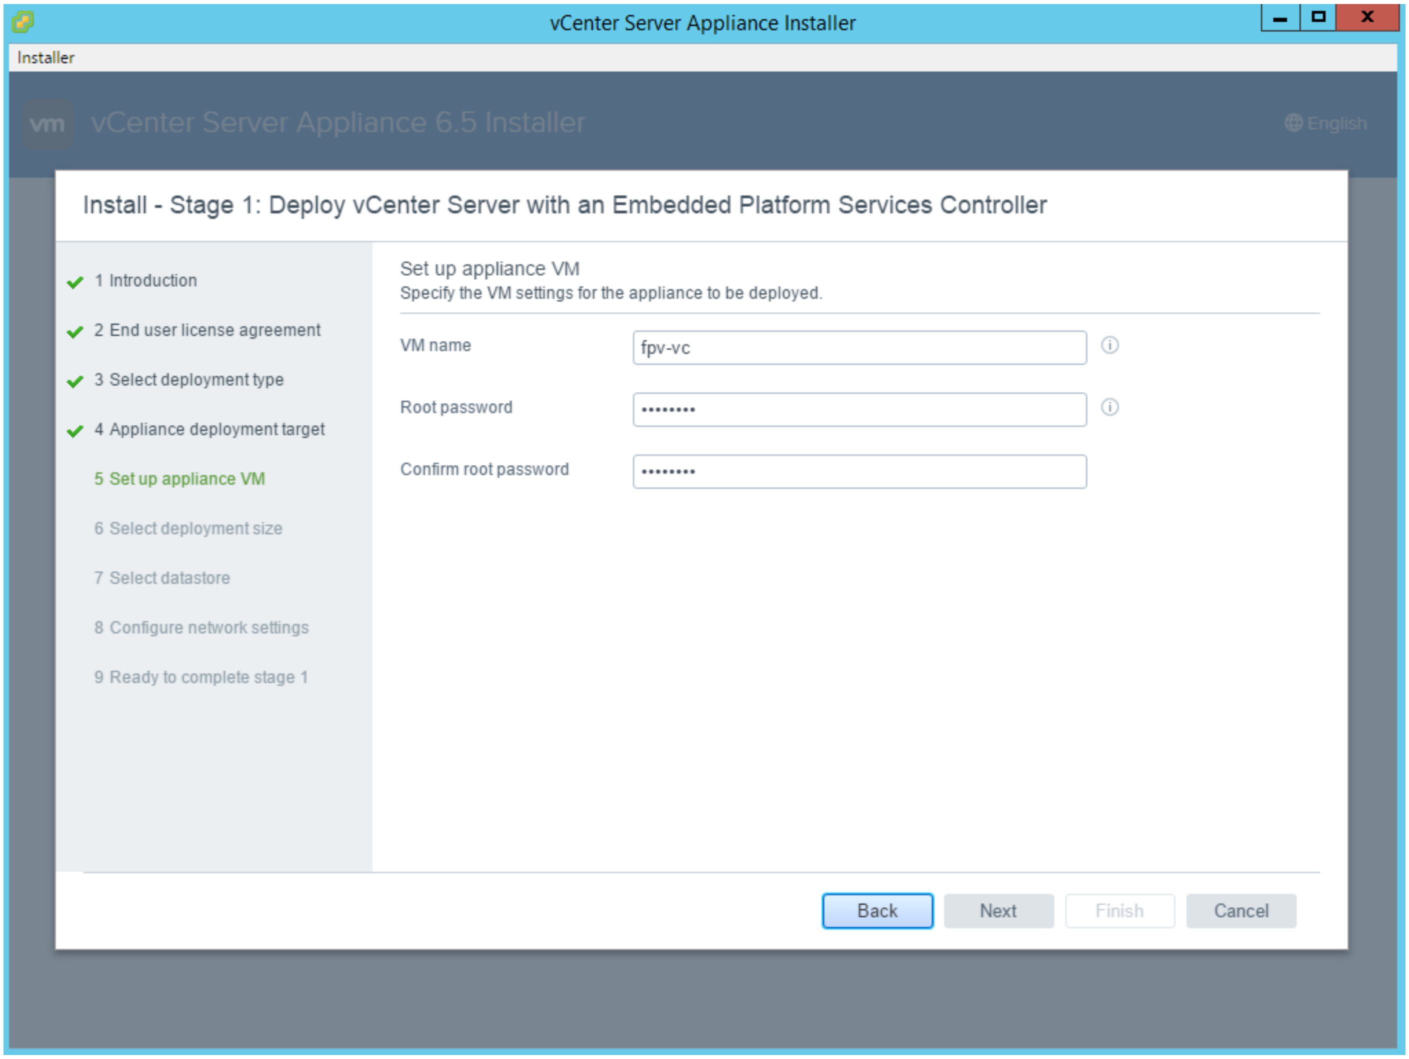
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
✓ 4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name	fpv-vc
Root password	••••••••
Confirm root password	••••••••

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

12. Click Next.

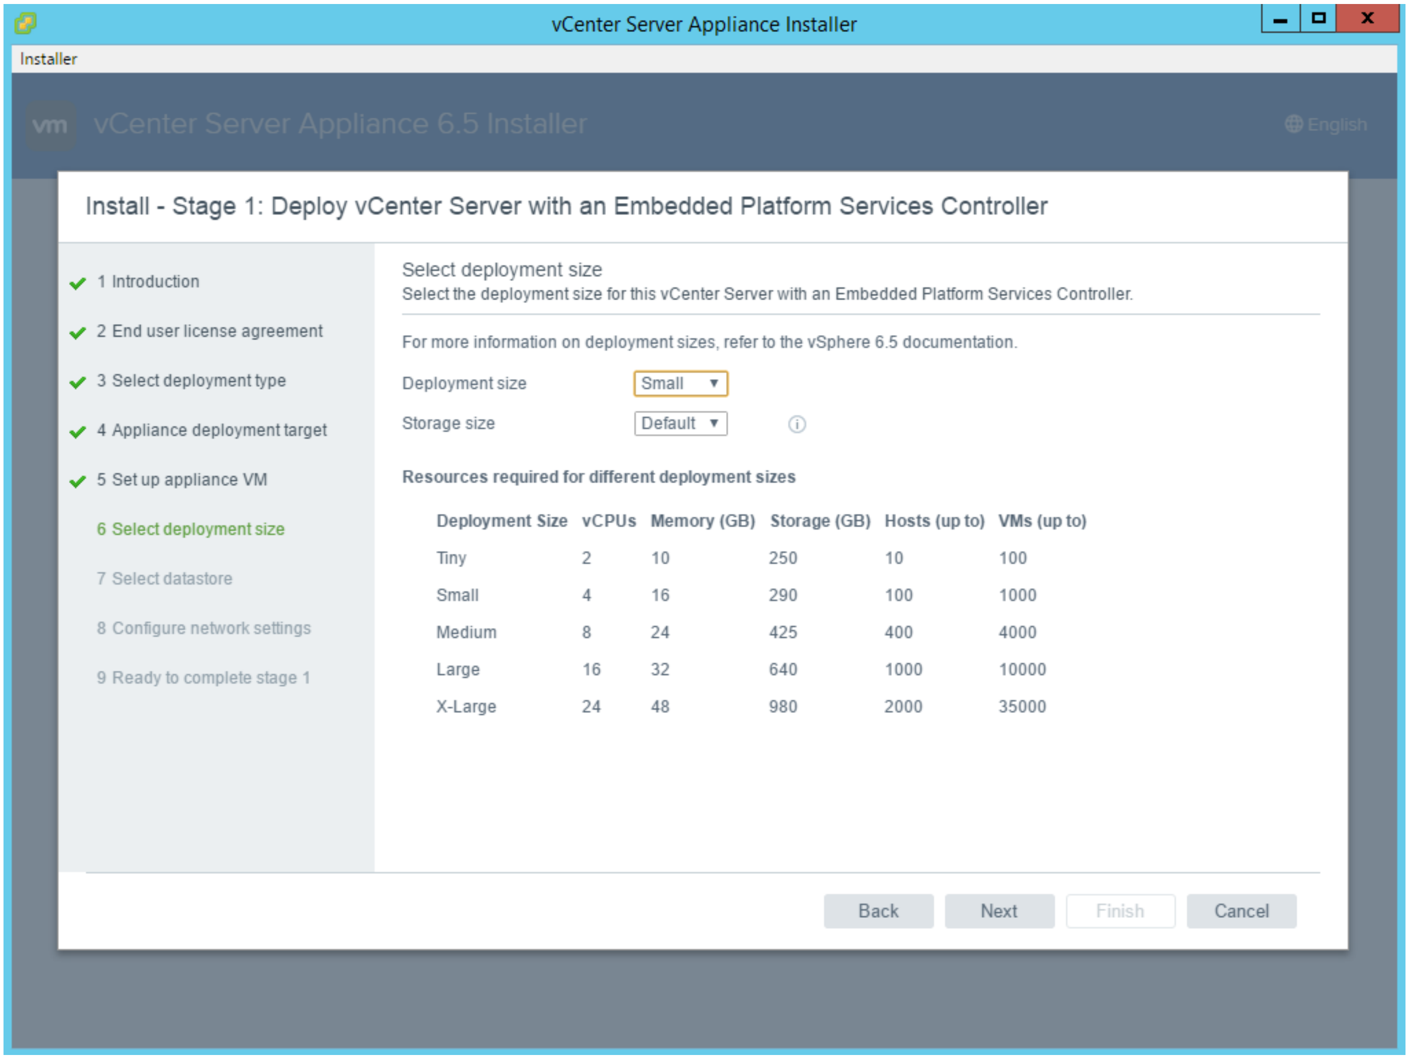

vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
✓ 4 Appliance deployment target
✓ 5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

Select deployment size

Select the deployment size for this vCenter Server with an Embedded Platform Services Controller.

For more information on deployment sizes, refer to the vSphere 6.5 documentation.

Deployment size: Small

Storage size: Default

Resources required for different deployment sizes

Deployment Size	vCPUs	Memory (GB)	Storage (GB)	Hosts (up to)	VMs (up to)
Tiny	2	10	250	10	100
Small	4	16	290	100	1000
Medium	8	24	425	400	4000
Large	16	32	640	1000	10000
X-Large	24	48	980	2000	35000

13. Select the infra_datastore_2. Click Next.

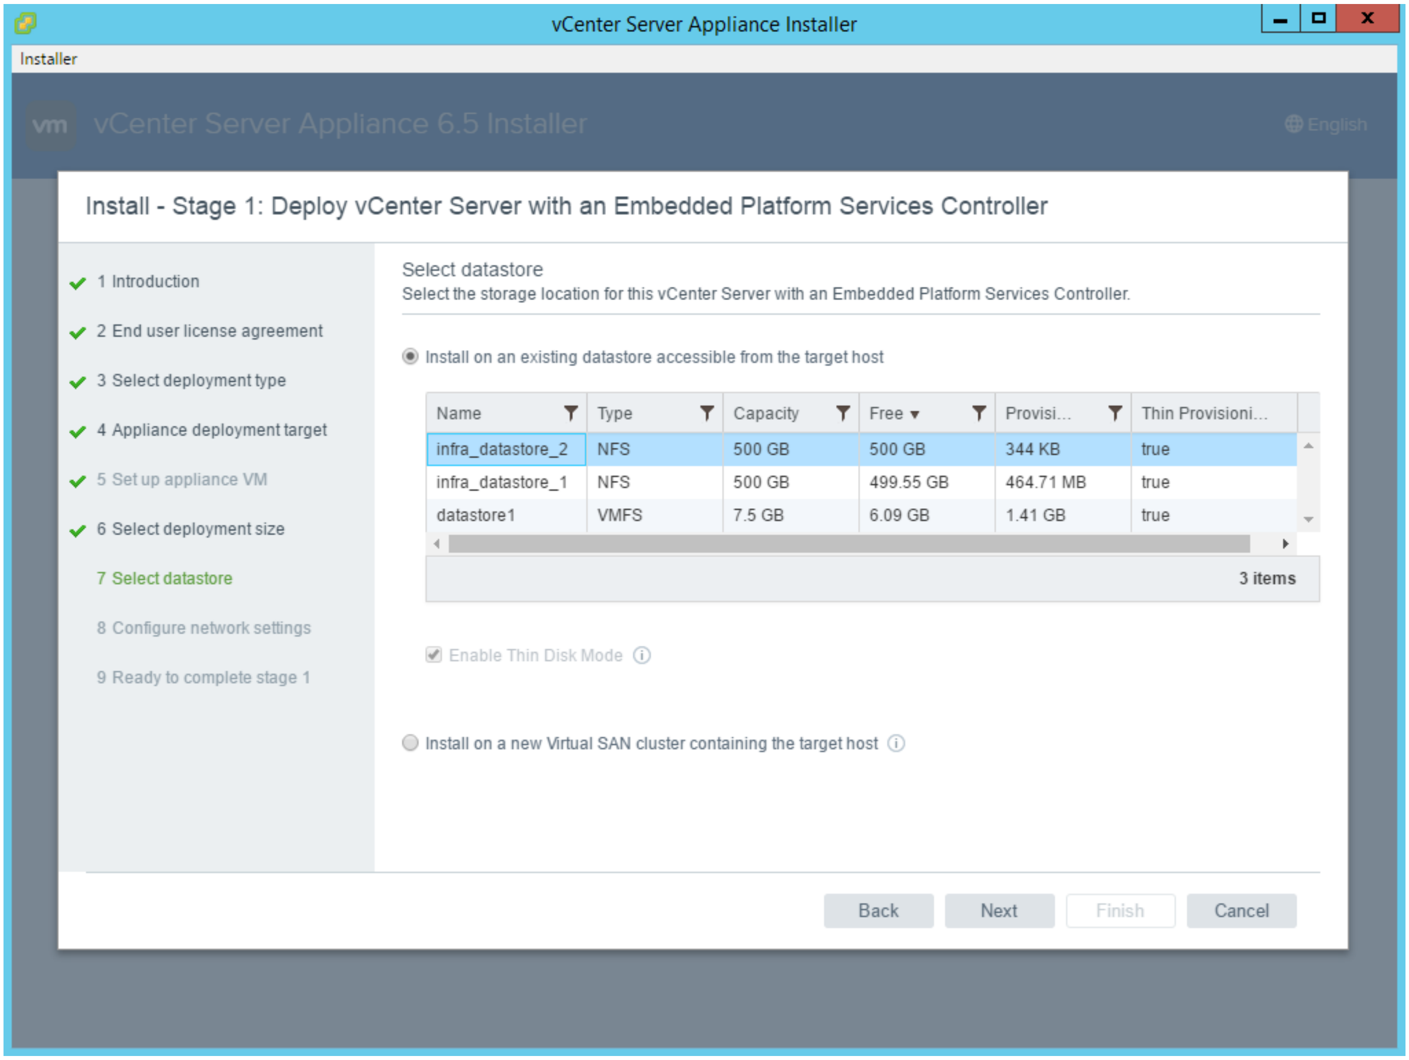

vCenter Server Appliance Installer

Installer

vCenter Server Appliance 6.5 Installer

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

- ✓ 1 Introduction
- ✓ 2 End user license agreement
- ✓ 3 Select deployment type
- ✓ 4 Appliance deployment target
- ✓ 5 Set up appliance VM
- ✓ 6 Select deployment size
- 7 Select datastore
- 8 Configure network settings
- 9 Ready to complete stage 1

Select datastore

Select the storage location for this vCenter Server with an Embedded Platform Services Controller.

● Install on an existing datastore accessible from the target host

Name	Type	Capacity	Free	Provisi...	Thin Provisioni...
infra_datastore_2	NFS	500 GB	500 GB	344 KB	true
infra_datastore_1	NFS	500 GB	499.55 GB	464.71 MB	true
datastore1	VMFS	7.5 GB	6.09 GB	1.41 GB	true

3 items

☑ Enable Thin Disk Mode

○ Install on a new Virtual SAN cluster containing the target host

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

    a. Choose a Network: Core-Services Network

    b. IP version: IPV4

    c. IP assignment: static

    d. System name: <vcenter-fqdn>

    e. IP address: <vcenter-ip>

    f. Subnet mask or prefix length: <vcenter-subnet-mask>

    g. Default gateway: <vcenter-gateway>

    h. DNS Servers: <dns-server>

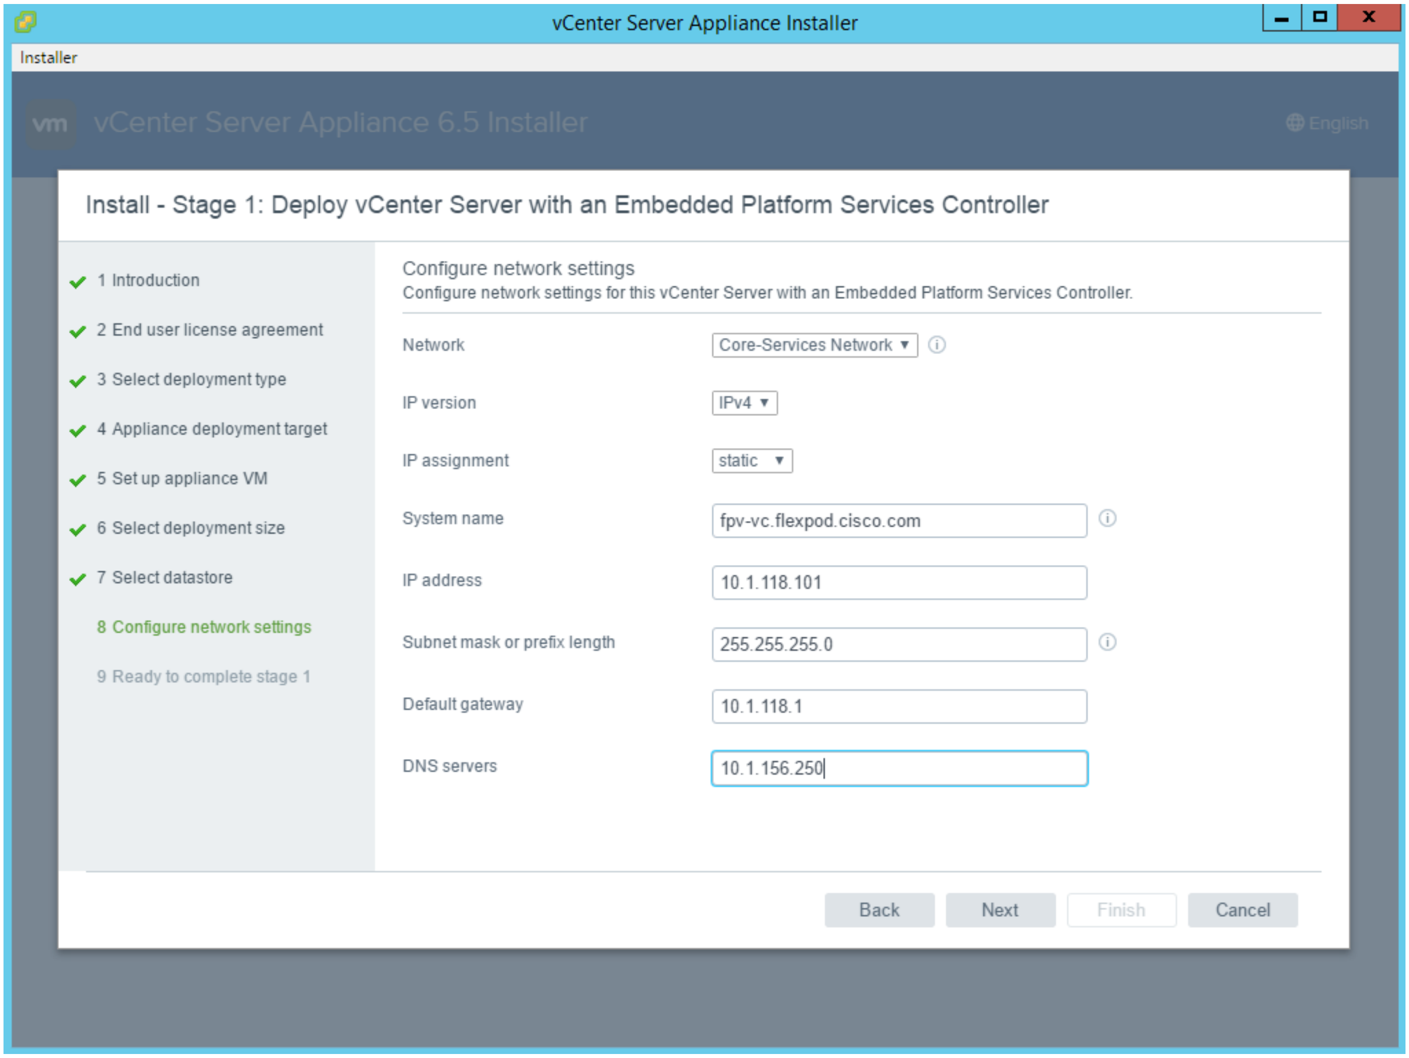

vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
✓ 4 Appliance deployment target
✓ 5 Set up appliance VM
✓ 6 Select deployment size
✓ 7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

Configure network settings

Configure network settings for this vCenter Server with an Embedded Platform Services Controller.

Network	Core-Services Network ▾
IP version	IPv4 ▾
IP assignment	static ▾
System name	fpv-vc.flexpod.cisco.com
IP address	10.1.118.101
Subnet mask or prefix length	255.255.255.0
Default gateway	10.1.118.1
DNS servers	10.1.156.250

15. Click Next.

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

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

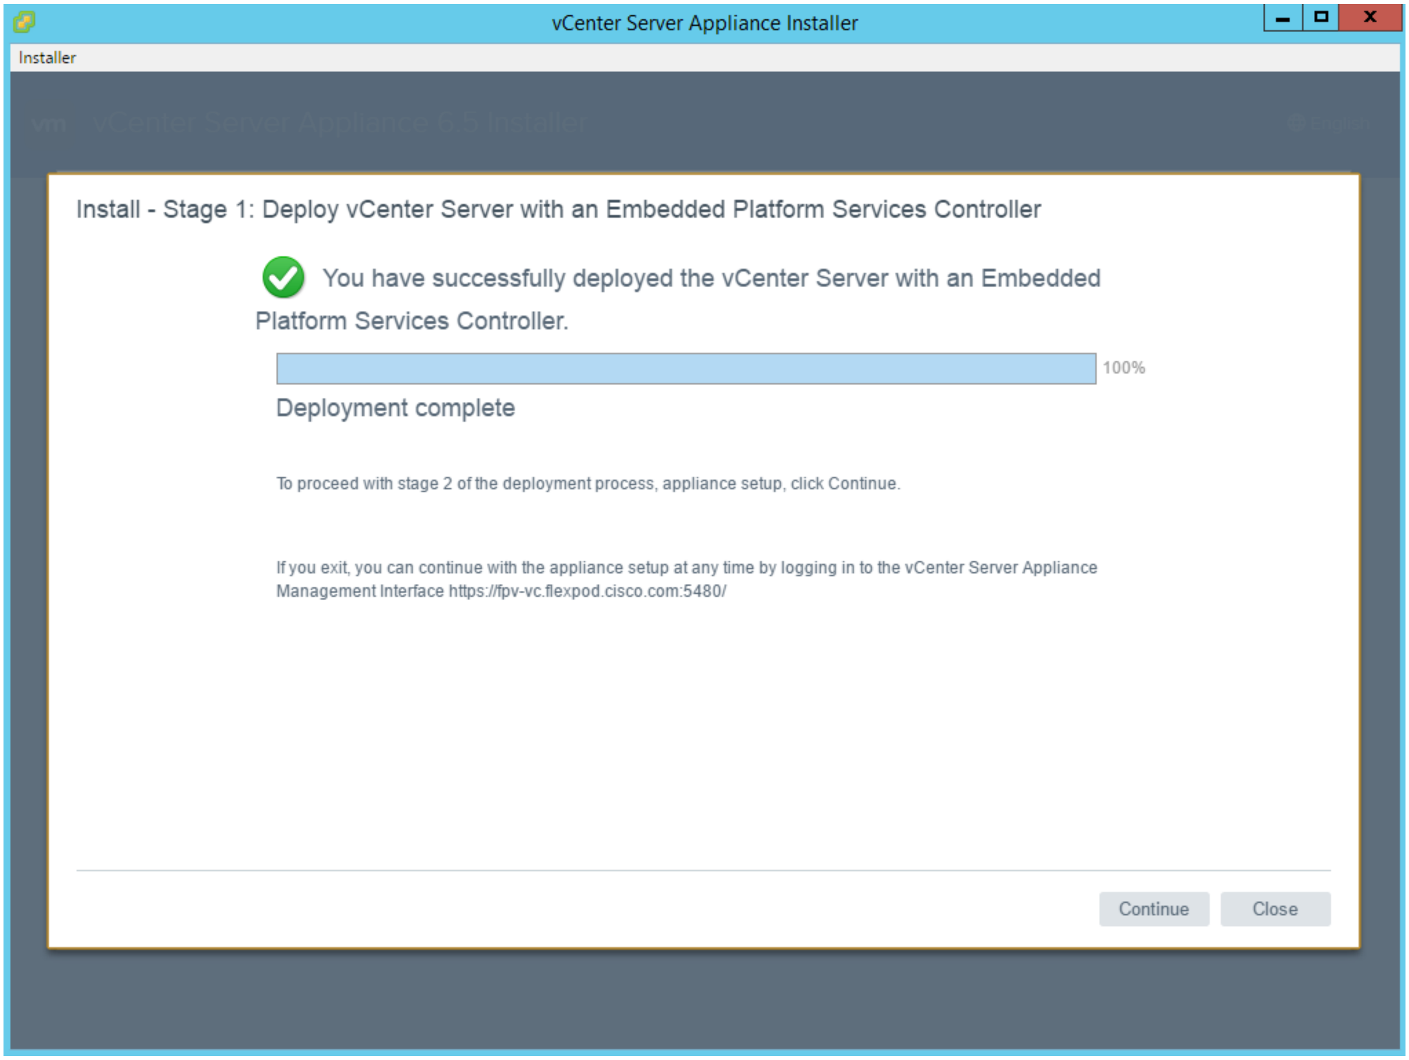
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

You have successfully deployed the vCenter Server with an Embedded Platform Services Controller.

100%

Deployment complete

To proceed with stage 2 of the deployment process, appliance setup, click Continue.

If you exit, you can continue with the appliance setup at any time by logging in to the vCenter Server Appliance Management Interface https://fpv-vc.flexpod.cisco.com:5480/

18. Click Continue to proceed with stage 2 configuration.
19. Click Next.
20. In the Appliance Configuration, configure the below settings:
    a. Time Synchronization Mode: Synchronize time with NTP servers.
    b. NTP Servers: <leaf-a-ntp-ip>, <leaf-b-ntp-ip>
    c. SSH access: Enabled.

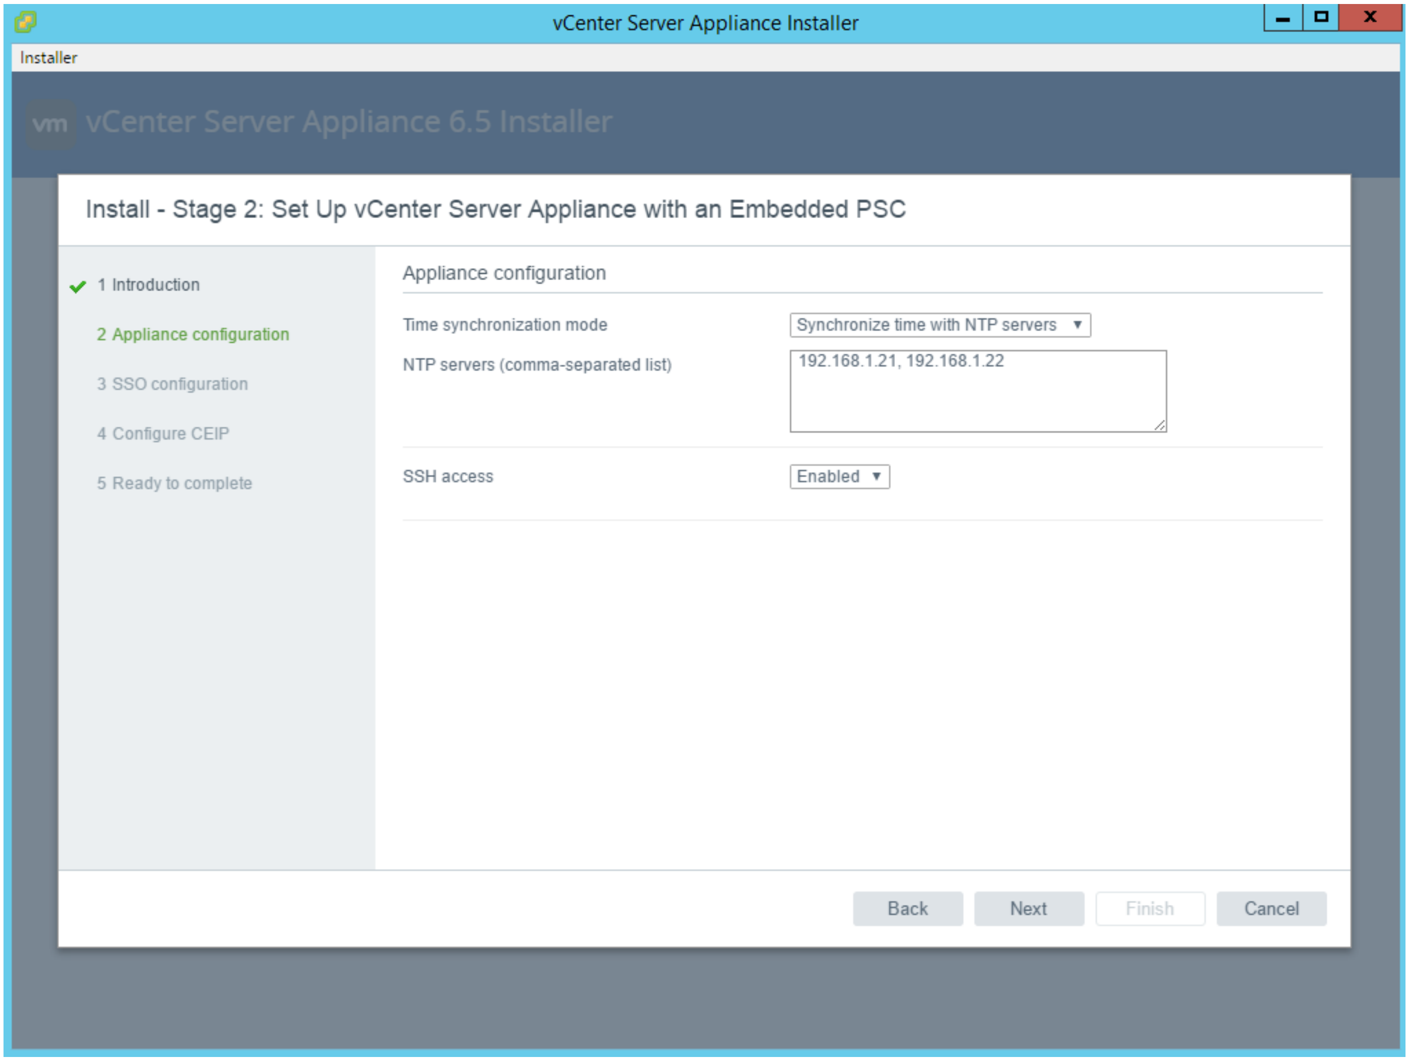
vCenter Server Appliance Installer
Installer
vCenter Server Appliance 6.5 Installer

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Appliance configuration

Time synchronization mode: Synchronize time with NTP servers

NTP servers (comma-separated list): 192.168.1.21, 192.168.1.22

SSH access: Enabled

21. Click Next.

22. Complete the SSO configuration as shown below.

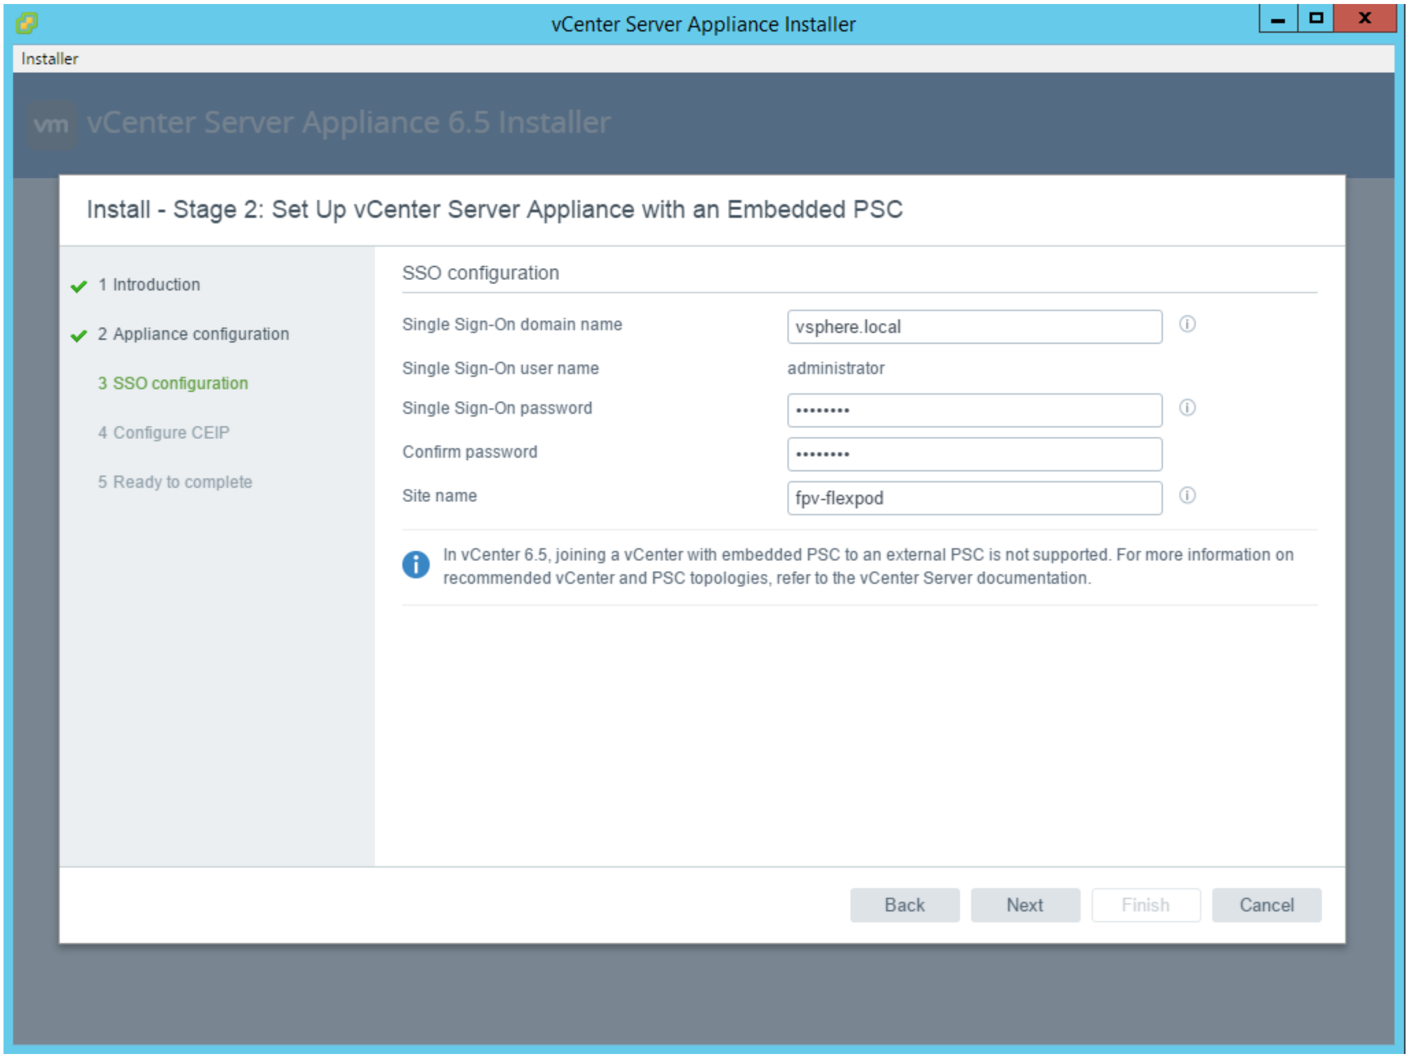
vCenter Server Appliance Installer
Installer
vCenter Server Appliance 6.5 Installer

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

✓ 1 Introduction
✓ 2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete

SSO configuration

Single Sign-On domain name	vsphere.local
Single Sign-On user name	administrator
Single Sign-On password	••••••••
Confirm password	••••••••
Site name	fpv-flexpod

In vCenter 6.5, joining a vCenter with embedded PSC to an external PSC is not supported. For more information on recommended vCenter and PSC topologies, refer to the vCenter Server documentation.

23. Click Next.

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

25. Click Next.

26. Review the configuration and click Finish.

27. Click OK.

28. Click Close.

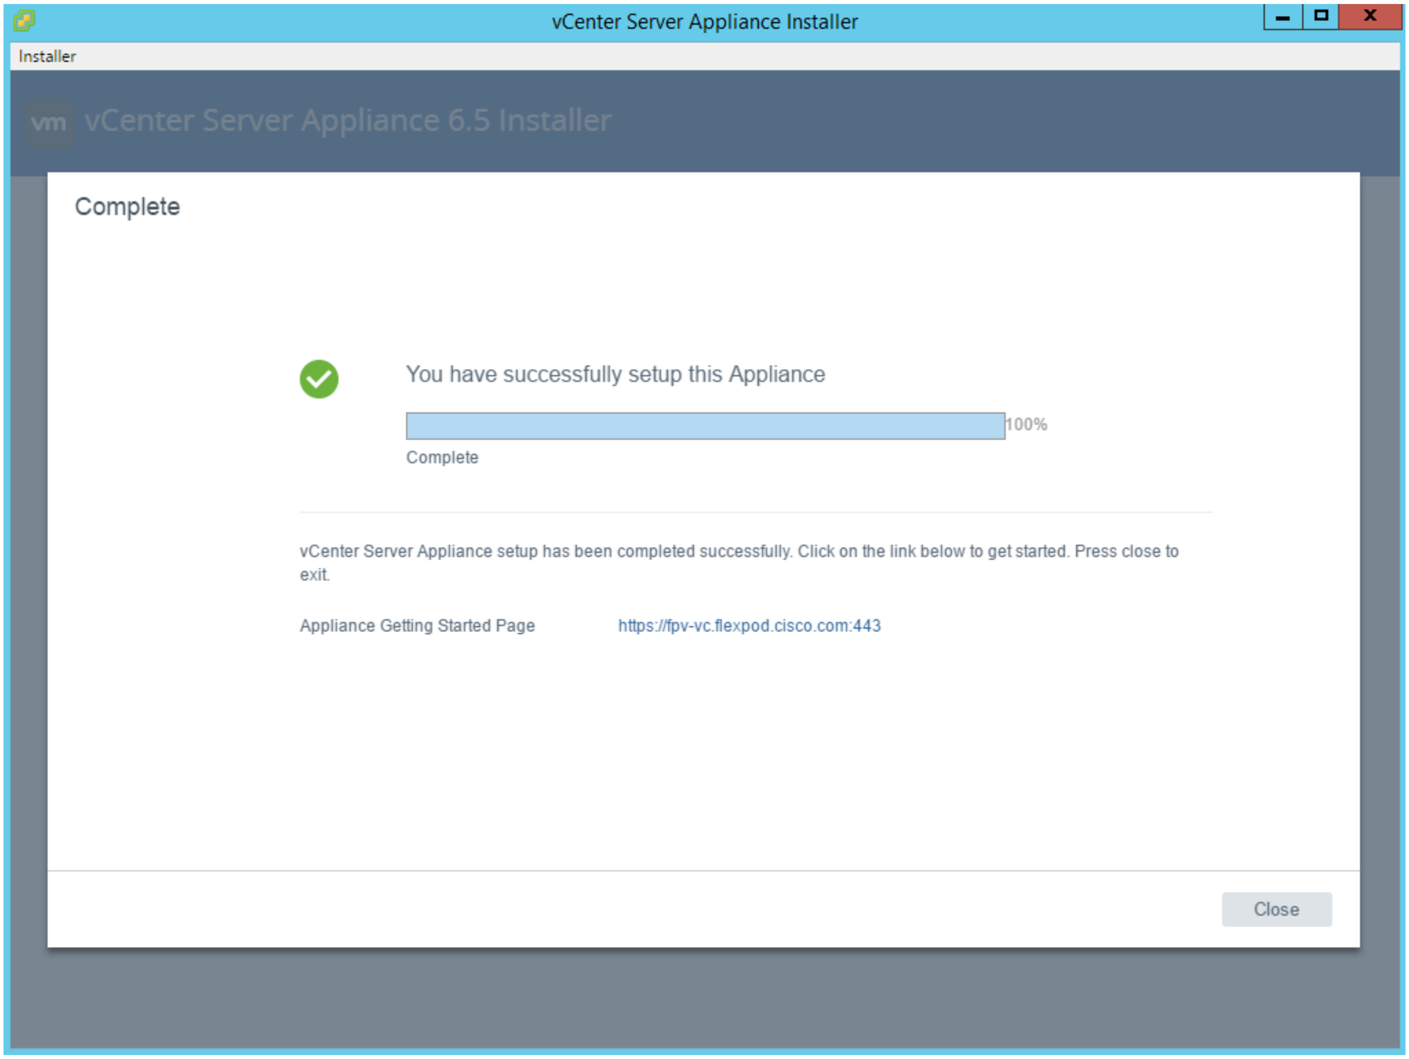
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
Complete

You have successfully setup this Appliance

100%

Complete

vCenter Server Appliance setup has been completed successfully. Click on the link below to get started. Press close to exit.

Appliance Getting Started Page https://fpv-vc.flexpod.cisco.com:443

Setting Up VMware vCenter Server

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

1. Using a web browser, navigate to https://<vcenter-ip>/vsphere-client.
2. Log in using the Single Sign-On username (Administrator@vsphere.local) and password created during the vCenter installation.
3. Click “Create Datacenter” in the center pane.
4. Type a name for the FlexPod Datacenter <FPV-FlexPod-DC> in the Datacenter name field.
5. Click OK.

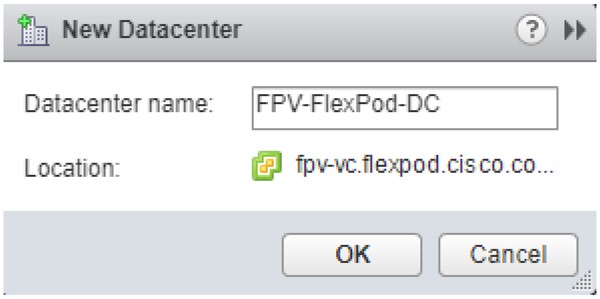

Datacenter name:	FPV-FlexPod-DC
Location:	fpv-vc.flexpod.cisco.co...

6. Right-click the data center just created and select New Cluster.

7. Name the cluster FPV-Foundation.

8. Check the box to turn on DRS. Leave the default values.

9. Check the box to turn on vSphere HA. Leave the default values.

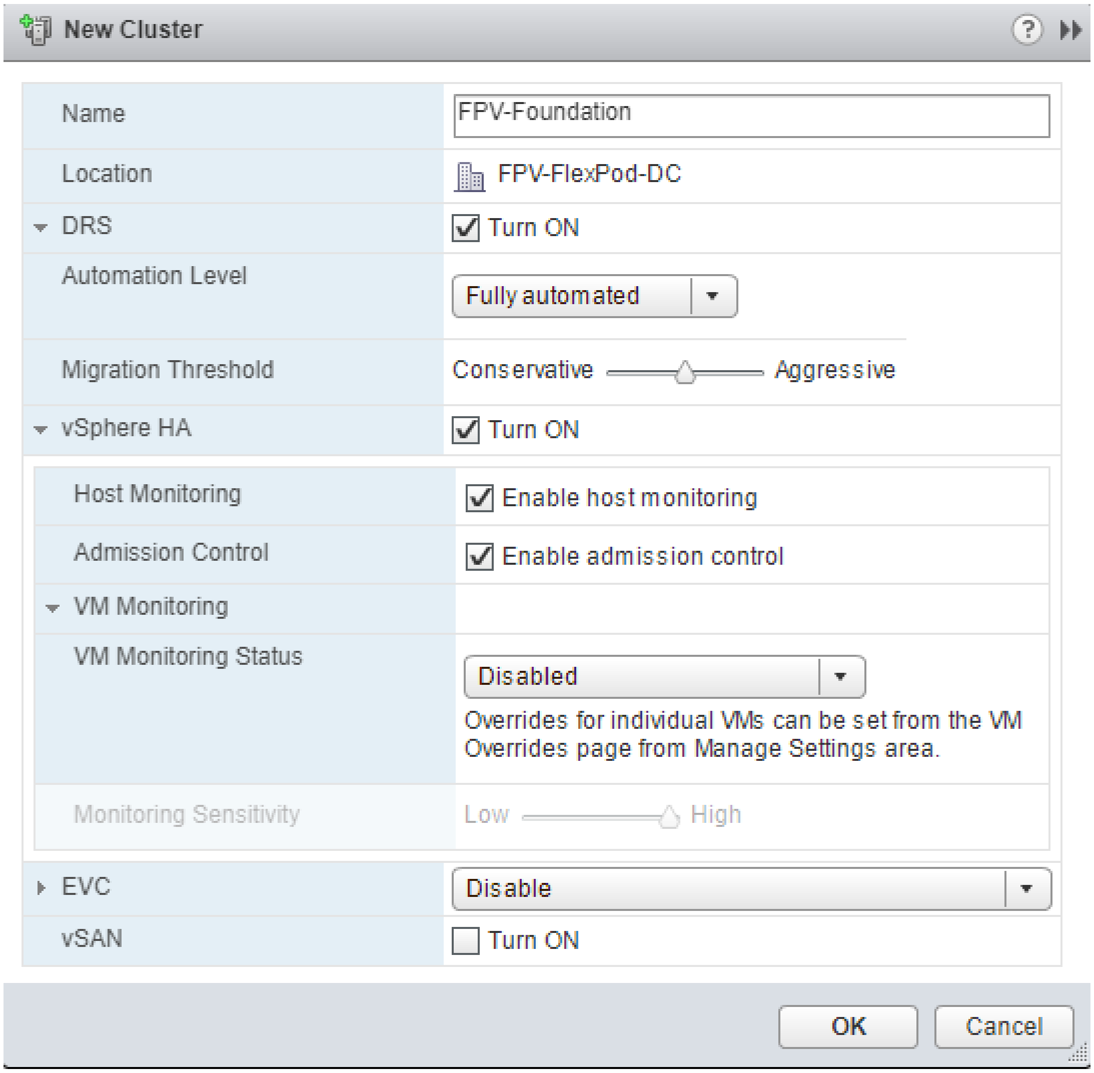

Name	FPV-Foundation
Location	FPV-FlexPod-DC
DRS	☑ Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative — Aggressive
vSphere HA	☑ Turn ON
Host Monitoring	☑ Enable host monitoring
Admission Control	☑ Enable admission control
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low — High
EVC	Disable
vSAN	□ Turn ON

10. Click OK to create the new cluster.

11. On the left pane, expand the Datacenter.

12. Right click the FPV-Foundation cluster and select Add Host.
13. In the Host field, enter either the IP address or the FQDN name of one of the VMware ESXi hosts. Click Next.
14. Type root as the user name and the root password. Click Next to continue.
15. Click Yes to accept the certificate.
16. Review the host details and click Next to continue.
17. Assign a license or leave in evaluation mode and click Next to continue.
18. Click Next to continue.
19. Click Next to continue.
20. Review the configuration parameters. Click Finish to add the host.

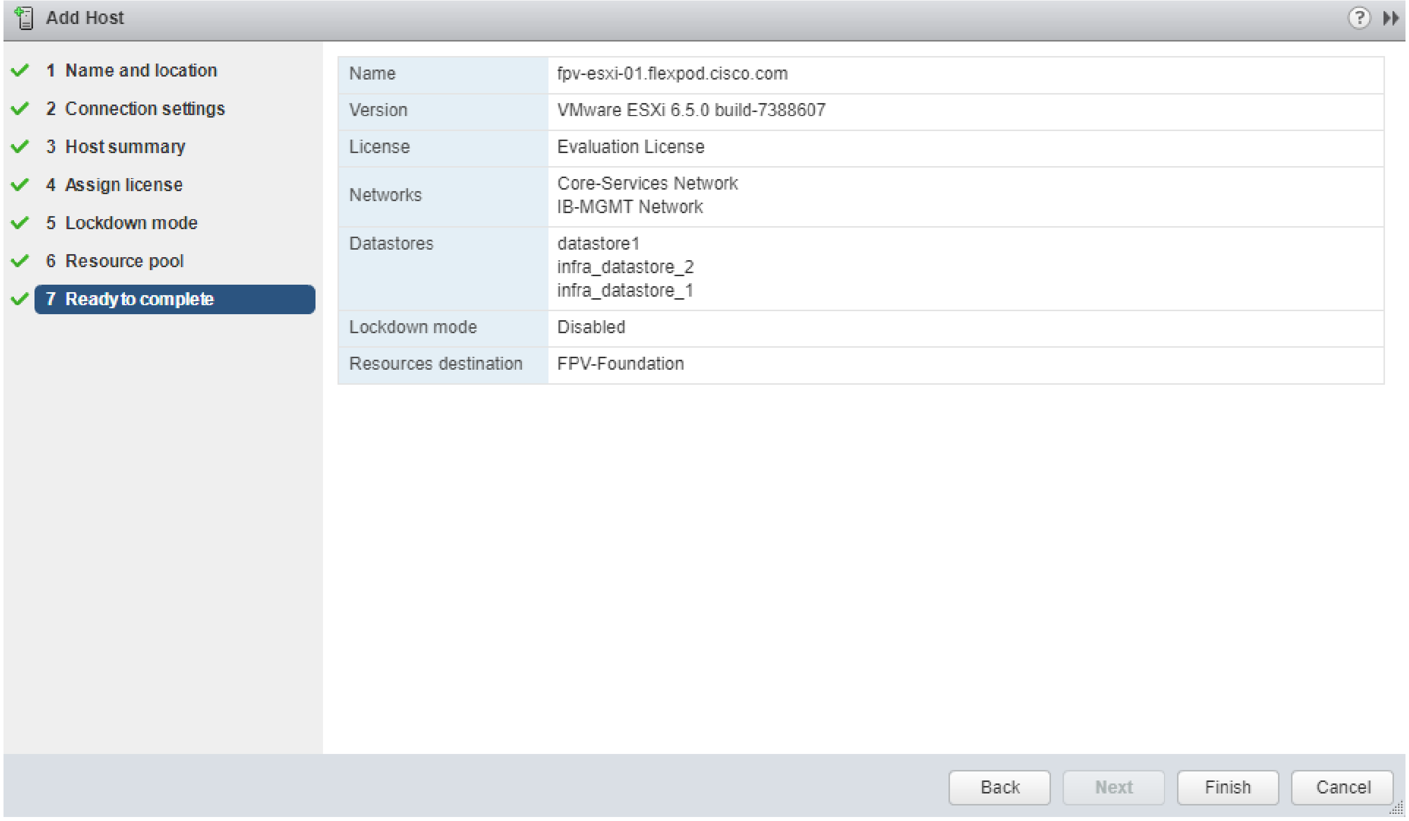

Name	fpv-esxi-01.flexpod.cisco.com
Version	VMware ESXi 6.5.0 build-7388607
License	Evaluation License
Networks	Core-Services Network IB-MGMT Network
Datastores	datastore1 infra_datastore_2 infra_datastore_1
Lockdown mode	Disabled
Resources destination	FPV-Foundation

21. Repeat the steps 12 to 20 to add the remaining VMware ESXi hosts to the cluster.

Two VMware ESXi hosts are added to the cluster.

Add AD User Authentication to vCenter (Optional)

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

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).

2. Connect to https://<vcenter-ip>, and select Log in to vSphere Web Client.

3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.

4. Navigate to Home. In the center pane, select System Configuration under Administration.

5. On the left, select Nodes and under Nodes select the vCenter.

6. In the center pane, select the manage tab, and within the Settings select Active Directory and click Join.

7. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Click OK.

8. On the left, right-click the vCenter and select Reboot.

9. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.

10. Log back into the vCenter Web Client.

11. In the center pane, select System Configuration under Administration.

12. On the left, select Nodes and under Nodes select the vCenter.

13. In the center pane under the Manage tab, select Active Directory. Make sure your Active Directory Domain is listed.

14. Navigate back to the vCenter Home.

15. In the center pane under Administration, select Roles.

16. On the left under Single Sign-On, select Configuration.

17. In the center pane, select the Identity Sources tab.

18. Click the green + sign to add an Identity Source.

19. Select the Active Directory (Integrated Windows Authentication) Identity source type and click Next.

20. Your AD domain name should be filled in. Leave Use machine account selected and click Next.

21. Click Finish to complete adding the AD domain as an Identity Source.

22. Your AD domain should now appear in the Identity Sources list.

23. On the left, under Single Sign-On, select Users and Groups.

24. In the center pane, select your AD domain for the Domain.

25. Make sure the FlexPod Admin user setup in step 1 appears in the list.

26. On the left under Administration, select Global Permissions.

27. Select the Manage tab, and click the green + sign to add a User or Group.

28. In the Global Permission Root - Add Permission window, click Add.

29. In the Select Users/Groups window, select your AD Domain.

30. Under Users and Groups, select either the FlexPod Admin user or the Domain Admins group.

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

31. Click Add. Click Check names to verify correctness of the names. Click OK to acknowledge the correctness of the names.

32. Click OK to add the selected User or Group.

33. Verify the added User or Group is listed under Users and Groups and the Administrator role is assigned.

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, for example flexadmin@domain.

ESXi Dump Collector setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. To setup the ESXi Dump Collector, complete the following steps:

1. Log into the vSphere web client as Administrator@vsphere.local.

2. In the vSphere web client, select Home.

3. In the center pane, click System Configuration.

4. In the left pane, select Services.

5. Under services, click VMware vSphere ESXi Dump Collector.

6. In the center pane, click the green start icon to start the service.

7. In the Actions menu, click Edit Startup Type.

8. Select Automatic.

9. Click OK.

10. Connect to each ESXi host via ssh as root

11. Run the following commands:

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

Add APIC-Integrated vSphere Distributed Switch (vDS)

The APIC-Integrated vDS is an integration between the Cisco ACI fabric and VMware allowing EPGs to be created in the ACI fabric and pushed into the vDS as port groups. The Virtual Machine Manager (VMM) domain in the APIC is configured with a pool of VLANs (100 in this validation) that are used as EPGs are assigned as port groups by associating the VMM domain with the EPG. These VLANs are already assigned to the UCS server vNICs and piped through Cisco UCS.

Create Virtual Machine Manager (VMM) Domain in APIC

APIC GUI

1. In the APIC GUI, select Virtual Networking > Inventory.
2. On the left, expand VMM Domains > VMware.
3. Right-click VMware and select Create vCenter Domain.
4. Name the Virtual Switch fpv-vc-vDS. Leave VMware vSphere Distributed Switch selected.
5. Select the UCS_AttEntityP Associated Attachable Entity Profile.
6. Under VLAN Pool, select Create VLAN Pool.
7. Name the VLAN Pool VP-fpv-vc-vDS. Leave Dynamic Allocation selected.
8. Click the “+” to add a block of VLANs to the pool.
9. Enter the VLAN range <1100-1199> and click OK.

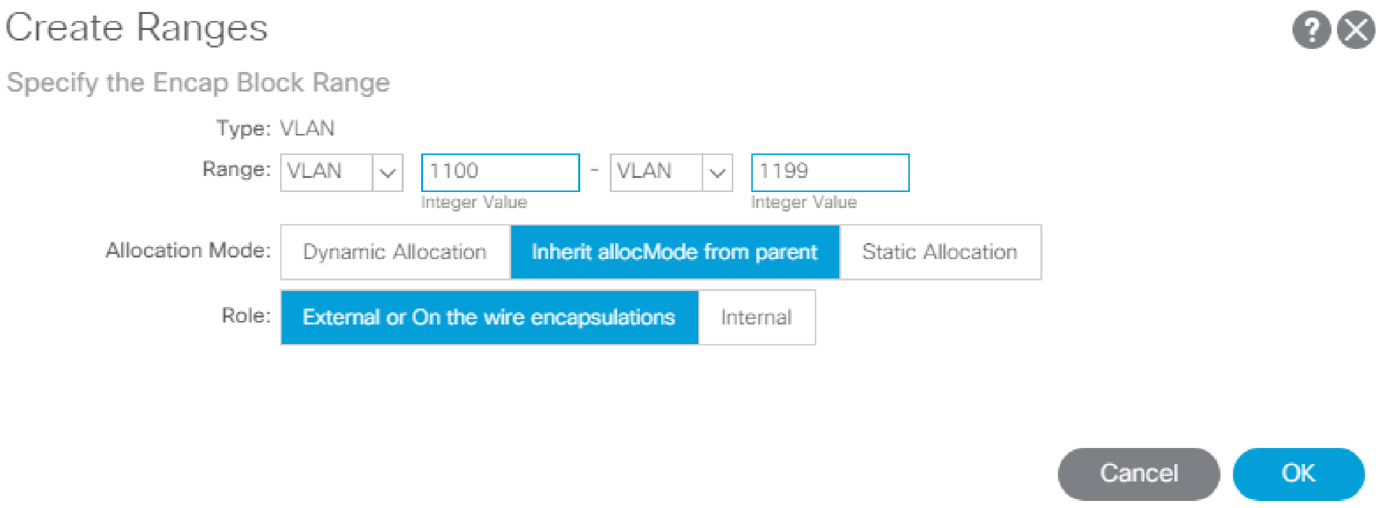

Create Ranges

Specify the Encap Block Range

Type: VLAN

Range: VLAN 1100 - VLAN 1199

Integer Value

Integer Value

Allocation Mode: Dynamic Allocation | Inherit allocMode from parent | Static Allocation

Role: External or On the wire encapsulations | Internal

10. Click Submit to complete creating the VLAN Pool.

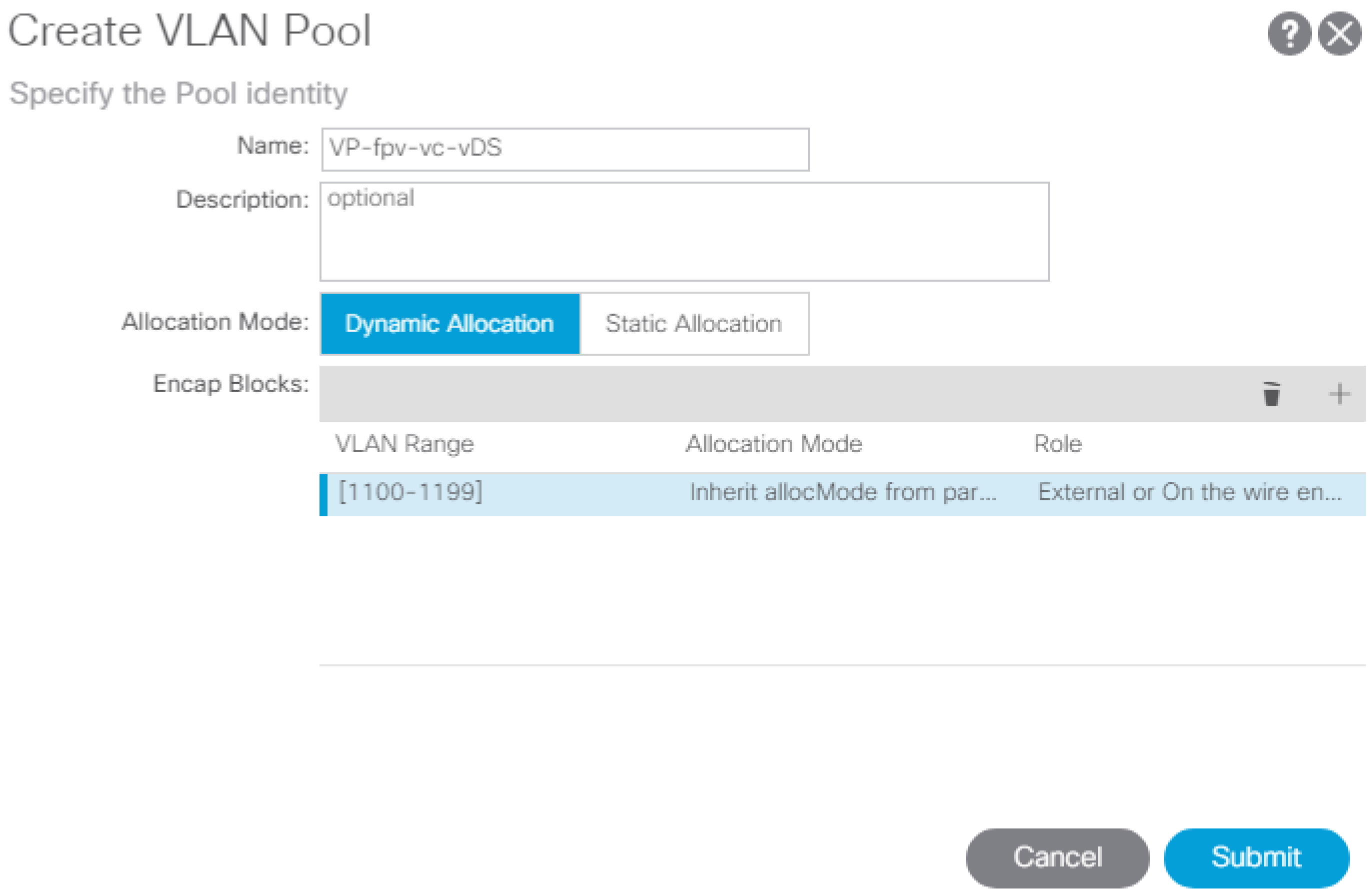

Create VLAN Pool

Specify the Pool identity

Name: VP-fpv-vc-vDS

Description: optional

Allocation Mode: Dynamic Allocation | Static Allocation

Encap Blocks:

VLAN Range	Allocation Mode	Role
[1100-1199]	Inherit allocMode from par...	External or On the wire en...

11. Click the “+” to the right of vCenter Credentials to add credentials for the vCenter.

12. For name, enter the vCenter hostname <fpv-vc>. For the username and password, the FlexPod Admin AD account created above can be used.

13. Click OK to complete creating the vCenter credentials.

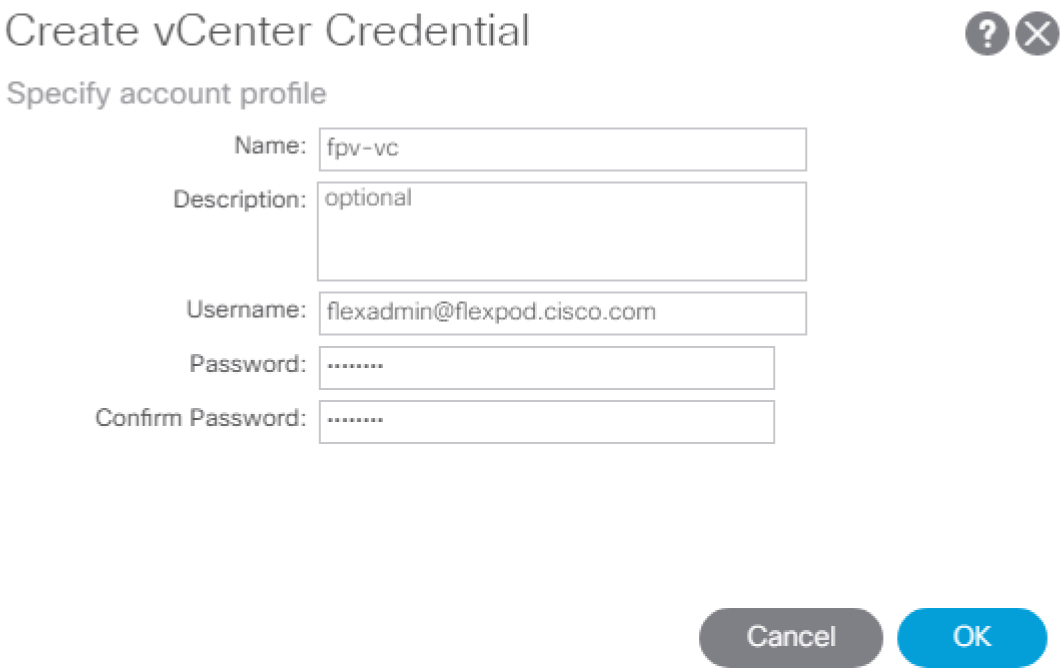

Create vCenter Credential

Specify account profile

Name: fpv-vc

Description: optional

Username: flexadmin@flexpod.cisco.com

Password: ••••••••

Confirm Password: ••••••••

14. Click the “+” to the right of vCenter to add the vCenter linkage.

15. Enter the vCenter hostname for Name. Enter the vCenter FQDN or IP address.

16. Select DVS Version 6.5. Enable Stats Collection.

17. For Datacenter, enter the exact Datacenter name specified in vCenter.

18. Do not select a Management EPG.

19. For Associated Credential, select the vCenter credentials entered in step 13.

20. Click OK to complete the vCenter linkage.

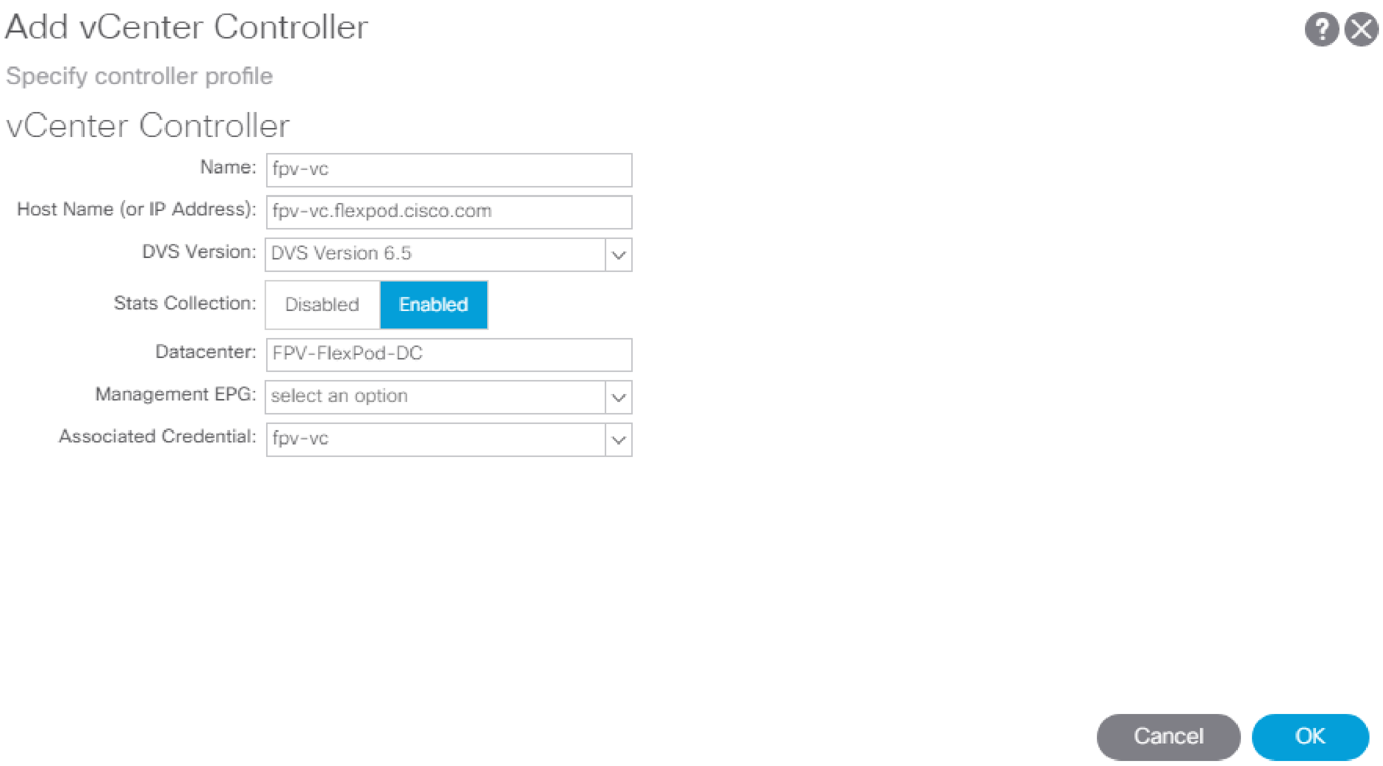

Add vCenter Controller

Specify controller profile

vCenter Controller

Name:	fpv-vc
Host Name (or IP Address):	fpv-vc.flexpod.cisco.com
DVS Version:	DVS Version 6.5
Stats Collection:	Disabled [✓ Enabled]
Datacenter:	FPV-FlexPod-DC
Management EPG:	select an option
Associated Credential:	fpv-vc

21. For Port Channel Mode, select MAC Pinning-Physical-NIC-load.

22. For vSwitch Policy, select LLDP.

23. Click Submit to complete Creating the vCenter Domain.

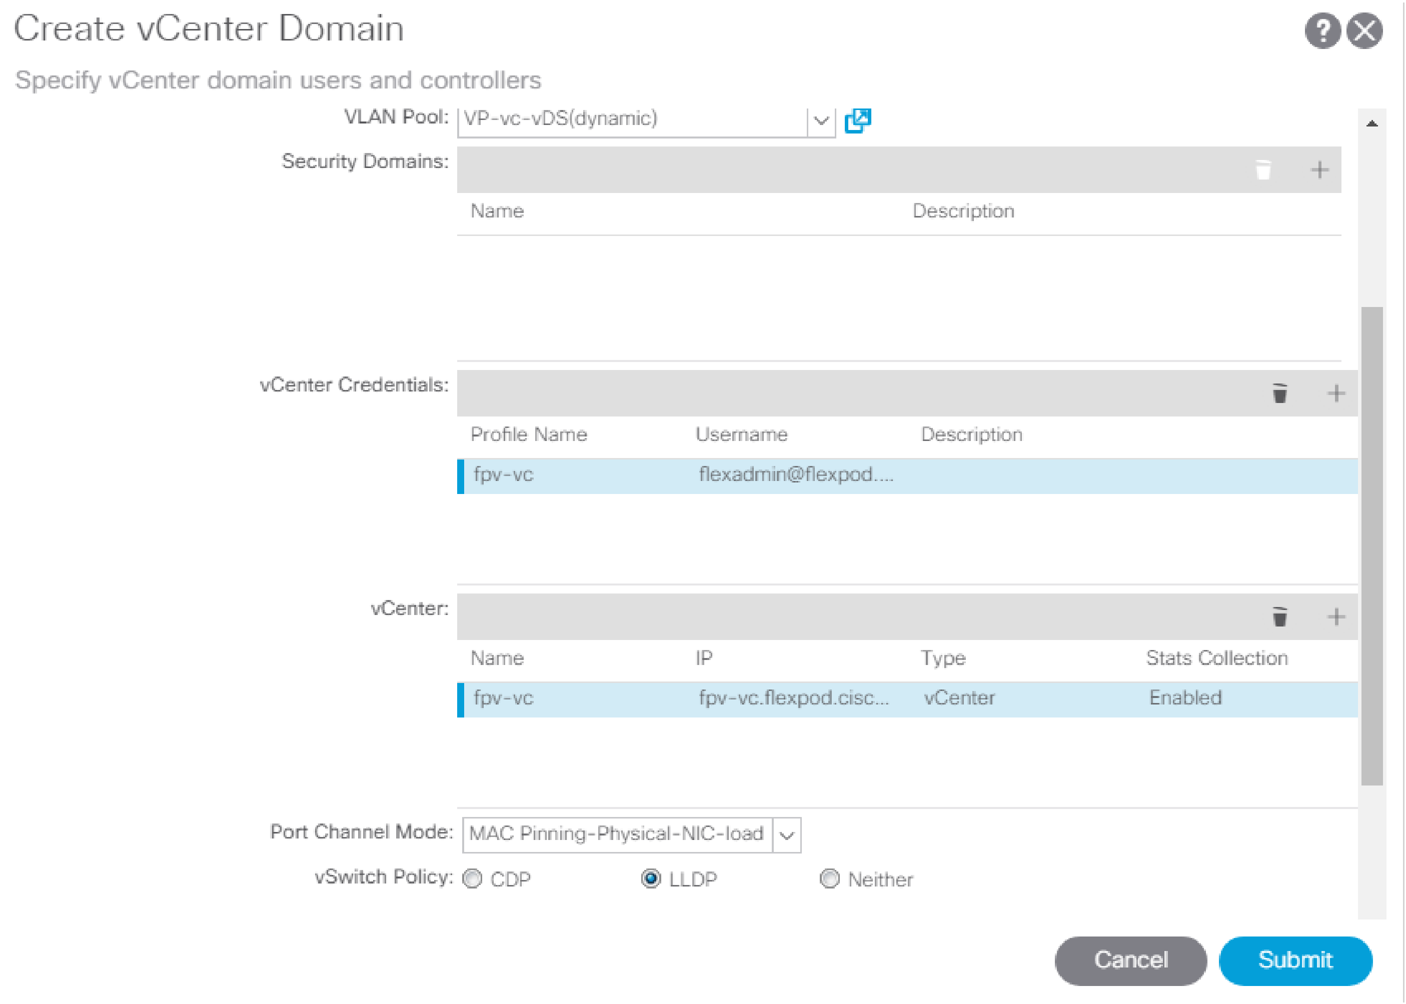

Create vCenter Domain

Specify vCenter domain users and controllers

VLAN Pool: VP-vc-vDS(dynamic)

Security Domains:

Name	Description

vCenter Credentials:

Profile Name	Username	Description
fpv-vc	flexadmin@flexpod....	

vCenter:

Name	IP	Type	Stats Collection
fpv-vc	fpv-vc.flexpod.cisc...	vCenter	Enabled

Port Channel Mode: MAC Pinning-Physical-NIC-load

vSwitch Policy: ○ CDP ● LLDP ○ Neither

The vDS should now appear in vCenter.

24. In the APIC GUI, select Tenants > common.

25. Under Tenant common, expand Application Profiles > Core-Services > Application EPGs > Core-Services.

26. Under the Core-Services EPG, right-click Domains and select Add VMM Domain Association.

27. Use the drop-down list to select the fpv-vc-vDS VMM Domain Profile. Select Immediate Deploy Immediacy and change no other values. Click Submit to create the Core-Services port group in the vDS.

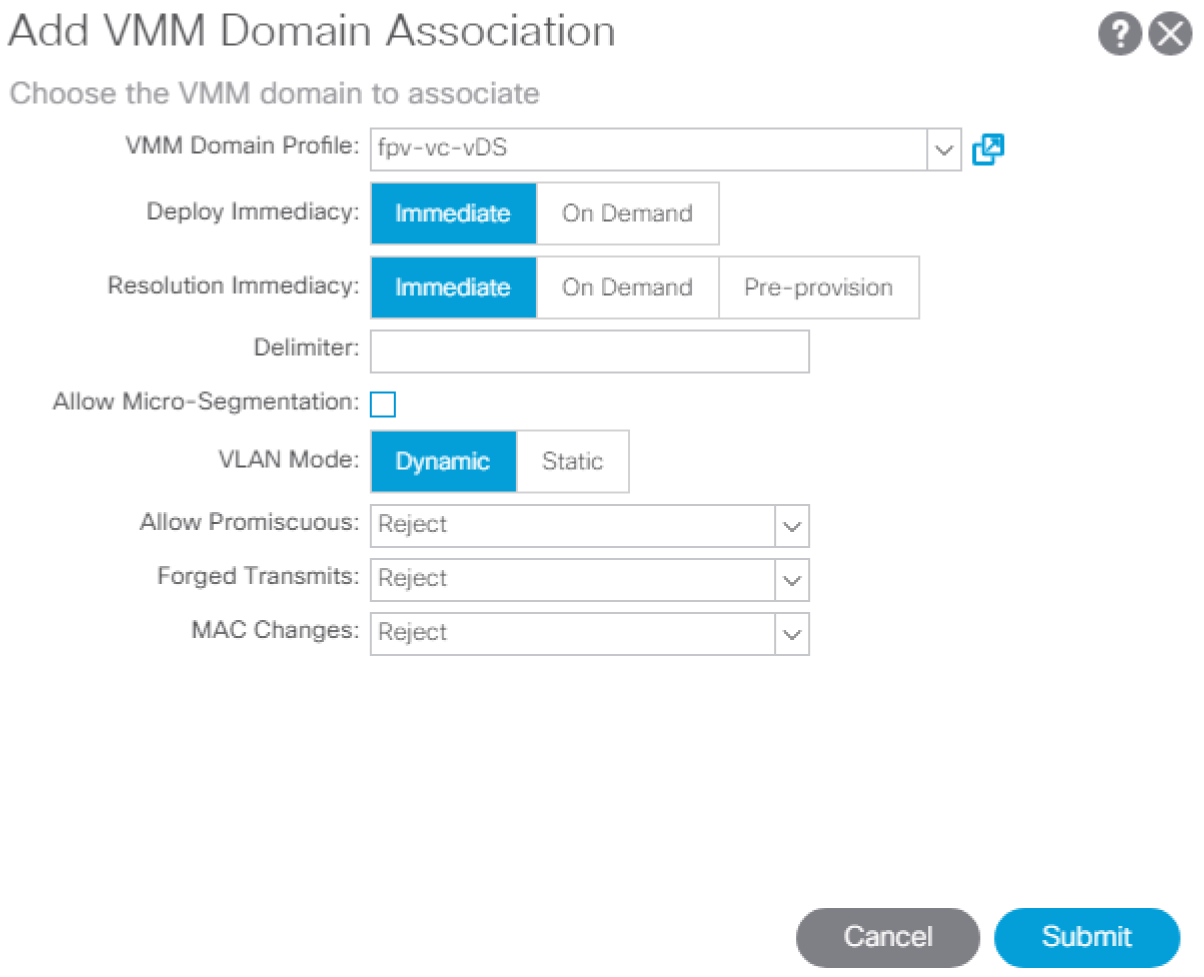

Add VMM Domain Association

Choose the VMM domain to associate

VMM Domain Profile: fpv-vc-vDS

Deploy Immediacy: [✓ Immediate] On Demand

Resolution Immediacy: [✓ Immediate] On Demand Pre-provision

Delimiter:

Allow Micro-Segmentation: □

VLAN Mode: [✓ Dynamic] Static

Allow Promiscuous: Reject

Forged Transmits: Reject

MAC Changes: Reject

28. Repeat steps 24-27 to add the VMM Domain Association to the IB-MGMT EPG in Tenant FPV-Foundation.

Add Management ESXi Hosts to APIC-Integrated vDS

vSphere Web Client

1. Log into the vCenter vSphere Web Client as the FlexPod Admin user.
2. Under the Navigator on the left, select the Networking icon.
3. Expand the vCenter, Datacenter, and vDS folder. Right-click the vDS and select Add and Manage Hosts.
4. Select Add hosts and click Next.
5. Click the ✚ to add New hosts.
6. Select both hosts and click OK.
7. Click Next.
8. Select only Manage physical adapters and click Next.
9. On both hosts, assign vmnic4 as uplink1 and vmnic 5 as uplink2. Click Next.

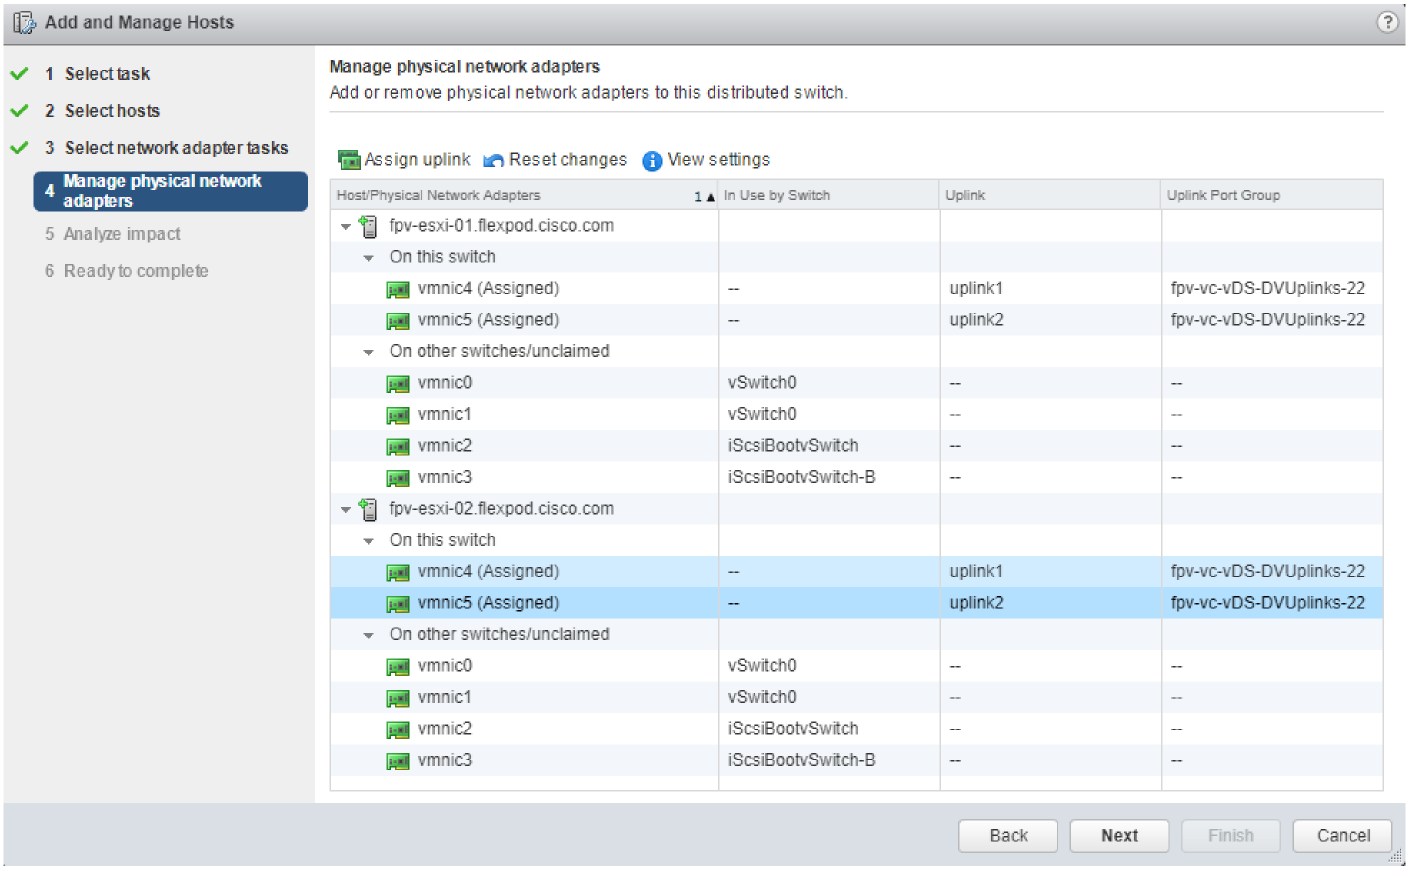

1 Select task
2 Select hosts
3 Select network adapter tasks
4 Manage physical network adapters
5 Analyze impact
6 Ready to complete

Manage physical network adapters

Add or remove physical network adapters to this distributed switch.

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
fpv-esxi-01.flexpod.cisco.com			
On this switch			
vmnic4 (Assigned)	--	uplink1	fpv-vc-vDS-DVUplinks-22
vmnic5 (Assigned)	--	uplink2	fpv-vc-vDS-DVUplinks-22
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	iScsiBootvSwitch	--	--
vmnic3	iScsiBootvSwitch-B	--	--
fpv-esxi-02.flexpod.cisco.com			
On this switch			
vmnic4 (Assigned)	--	uplink1	fpv-vc-vDS-DVUplinks-22
vmnic5 (Assigned)	--	uplink2	fpv-vc-vDS-DVUplinks-22
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	iScsiBootvSwitch	--	--
vmnic3	iScsiBootvSwitch-B	--	--

10. Click Next and Finish to complete adding the two Management hosts to the vDS. VMs can now be assigned to the Core-Services and IB-MGMT port groups in the vDS.

It is not recommended to assign the vCenter VM to the Core-Services port group in the vDS or to migrate the management hosts' Infrastructure NFS VMkernel ports to the vDS.

Cisco ACI vCenter Plug-in

The Cisco ACI vCenter plug-in is a user interface that allows you to manage the ACI fabric from within the vSphere Web client. This allows the VMware vSphere Web Client to become a single pane of glass to configure both VMware vCenter and the ACI fabric. The Cisco ACI vCenter plug-in empowers virtualization administrators to define network connectivity independently of the networking team while sharing the same infrastructure. No configuration of in-depth networking is done through the Cisco ACI vCenter plug-in. Only the elements that are relevant to virtualization administrators are exposed.

Cisco ACI vCenter Plug-in Installation

To begin the plug-in installation on a Windows system, complete the following steps:

To complete the installation of the ACI vCenter Plug-in, VMware PowerCLI 6.5 Release 1 must be installed on the Windows administration workstation. VMware PowerCLI 6.5 Release 1 can be downloaded from https://my.vmware.com/group/vmware/details?downloadGroup=PCLI650R1&productId=614.

1. Connect to: https://<apic-ip>/vcplugin.

2. Follow the Installation instructions on that web page to complete plug-in installation.
3. Logout and log back into the vSphere Web Client.
4. Select Home. The Cisco ACI Fabric plugin should be available under Inventories.
5. Select the Cisco ACI Fabric plugin.
6. In the center pane, select Connect vSphere to your ACI Fabric.
7. Click Yes to add a new ACI Fabric.
8. Enter one APIC IP address or FQDN and uncheck Use Certificate.
9. Enter the admin Username and Password. Click OK.
10. Click OK to confirm the addition of the other APICs.

Cisco UCS Manager Plug-in for VMware vSphere Web Client

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

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

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

Cisco UCS Manager Plug-in Installation

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

1. Download the plugin and registration tool from: https://software.cisco.com/download/release.html?mdfid=286282669&catid=282558030&softwareid=286282010&release=2.0.3&relind=AVAILABLE&rellifecycle=&reltype=latest.
2. Place the downloaded ucs-vcplugin-2.0.3.zip file onto the web server used for hosting the ONTAP software and VMware ESXi ISO.
3. Unzip the Cisco_UCS_Plugin_Registration_Tool_1_1_3.zip and open the executable file within it.
4. Leave Register Plugin selected for the Action, and fill in:
   - IP/Hostname

- Username
- Password
- URL that plugin has been uploaded

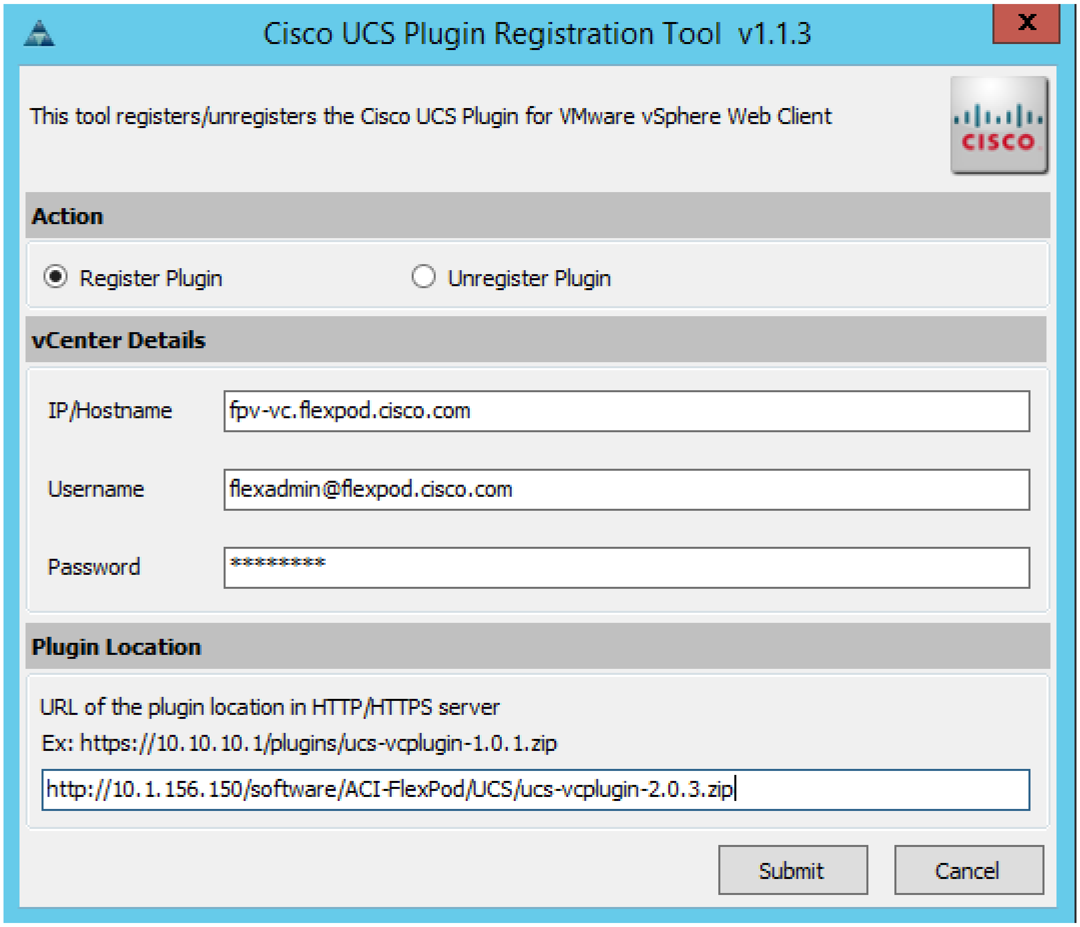

This tool registers/unregisters the Cisco UCS Plugin for VMware vSphere Web Client

Action

● Register Plugin ○ Unregister Plugin

vCenter Details

IP/Hostname: fpv-vc.flexpod.cisco.com

Username: flexadmin@flexpod.cisco.com

Password: *********

Plugin Location

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

http://10.1.156.150/software/ACI-FlexPod/UCS/ucs-vcplugin-2.0.3.zip

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

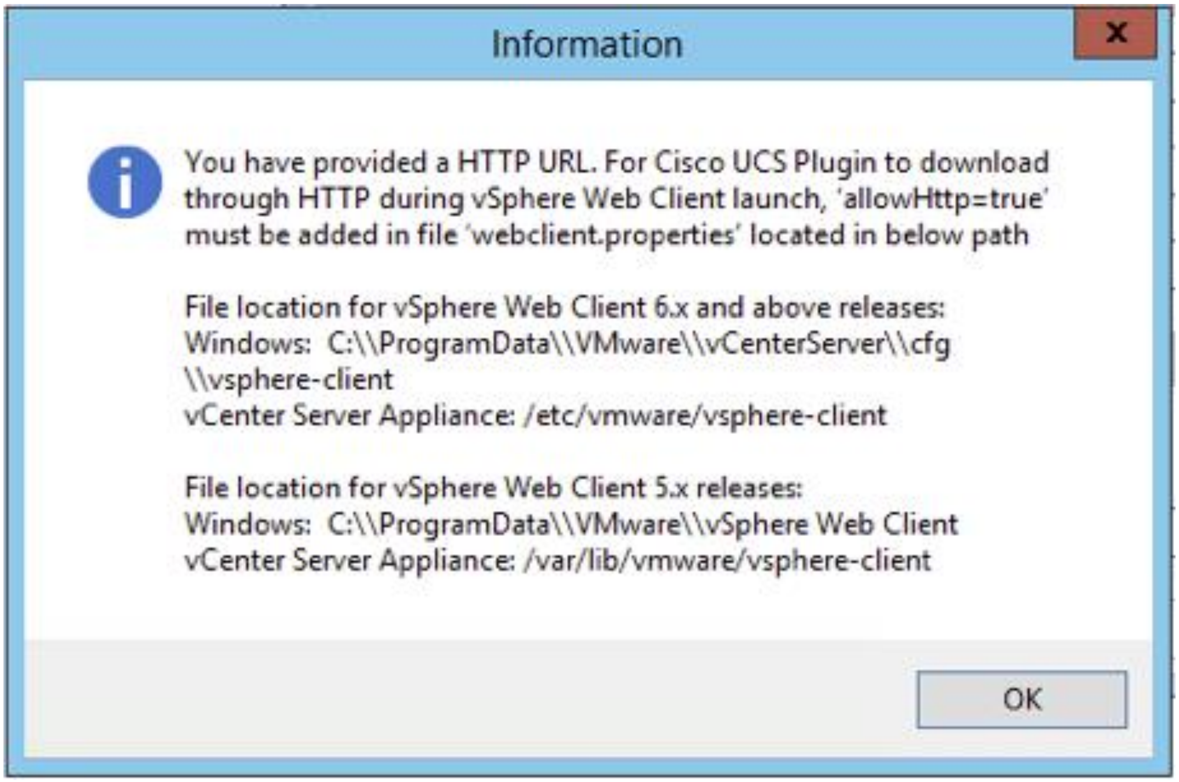

You have provided a HTTP URL. For Cisco UCS Plugin to download through HTTP during vSphere Web Client launch, 'allowHttp=true' must be added in file 'webclient.properties' located in below path

File location for vSphere Web Client 6.x and above releases:
Windows: C:\\ProgramData\\VMware\\vCenterServer\\cfg\\vsphere-client
vCenter Server Appliance: /etc/vmware/vsphere-client

File location for vSphere Web Client 5.x releases:
Windows: C:\\ProgramData\\VMware\\vSphere Web Client
vCenter Server Appliance: /var/lib/vmware/vsphere-client

6. Take care of this issue after the plugin has been registered, click OK to close the Information dialogue box.

7. Click OK on the confirmation message. Then click Cancel to close the registration tool.

8. To resolve the change needed for the HTTP download of the vSphere Web Client launch, connect to the VCSA with ssh using the root account and type:

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

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

9. Reboot the VCSA.

Cisco UCS Domain Registration

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

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

1. Login to the vSphere Web Client with the FlexPod Admin user id.

2. Select Home from the Navigator or pull-down options, and click the Cisco UCS icon appearing in the Administration section.

3. Click the Register button and provide the following options in the Register UCS Domain dialogue box that appears:

   - UCS Hostname/IP
   - Username
   - Password
   - Port (if different than 443)
   - Leave SSL selected and click the Visible to All users option

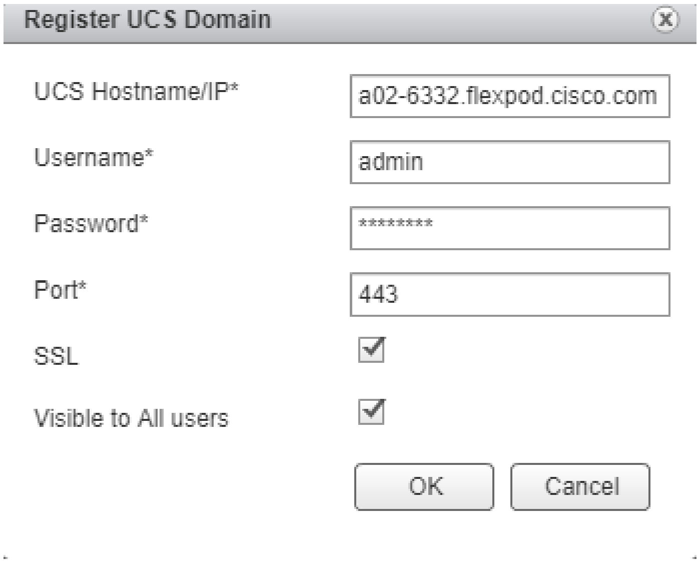

Register UCS Domain

UCS Hostname/IP*: a02-6332.flexpod.cisco.com

Username*: admin

Password*: ********

Port*: 443

SSL ☑

Visible to All users ☑

4. Click OK to register the UCS Domain.

Using the Cisco UCS vCenter Plugin

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

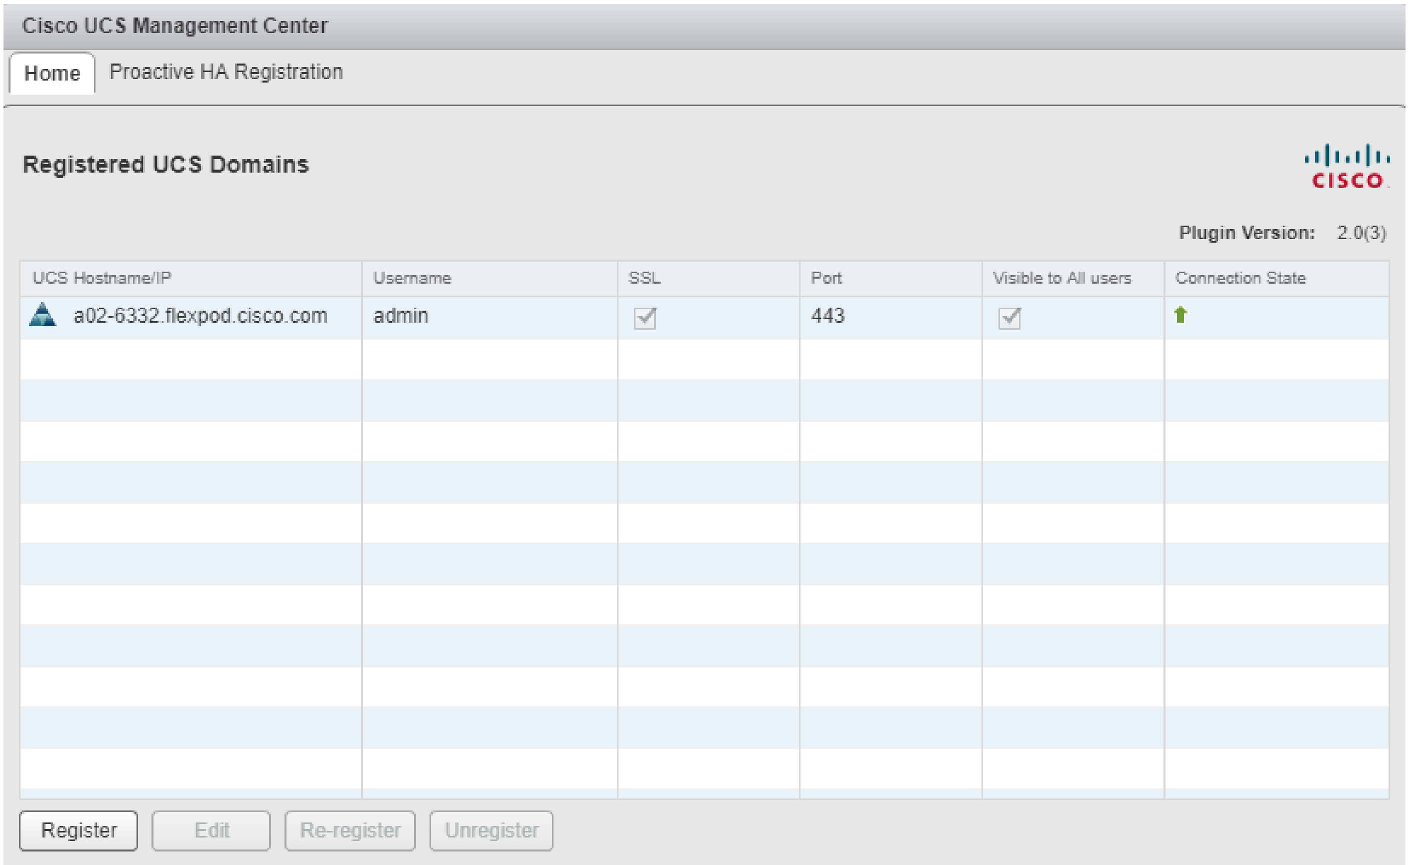

Registered UCS Domains

Plugin Version: 2.0(3)

UCS Hostname/IP	Username	SSL	Port	Visible to All users	Connection State
a02-6332.flexpod.cisco.com	admin	☑	443	☑	↑

This will display a view of the components associated to the domain:

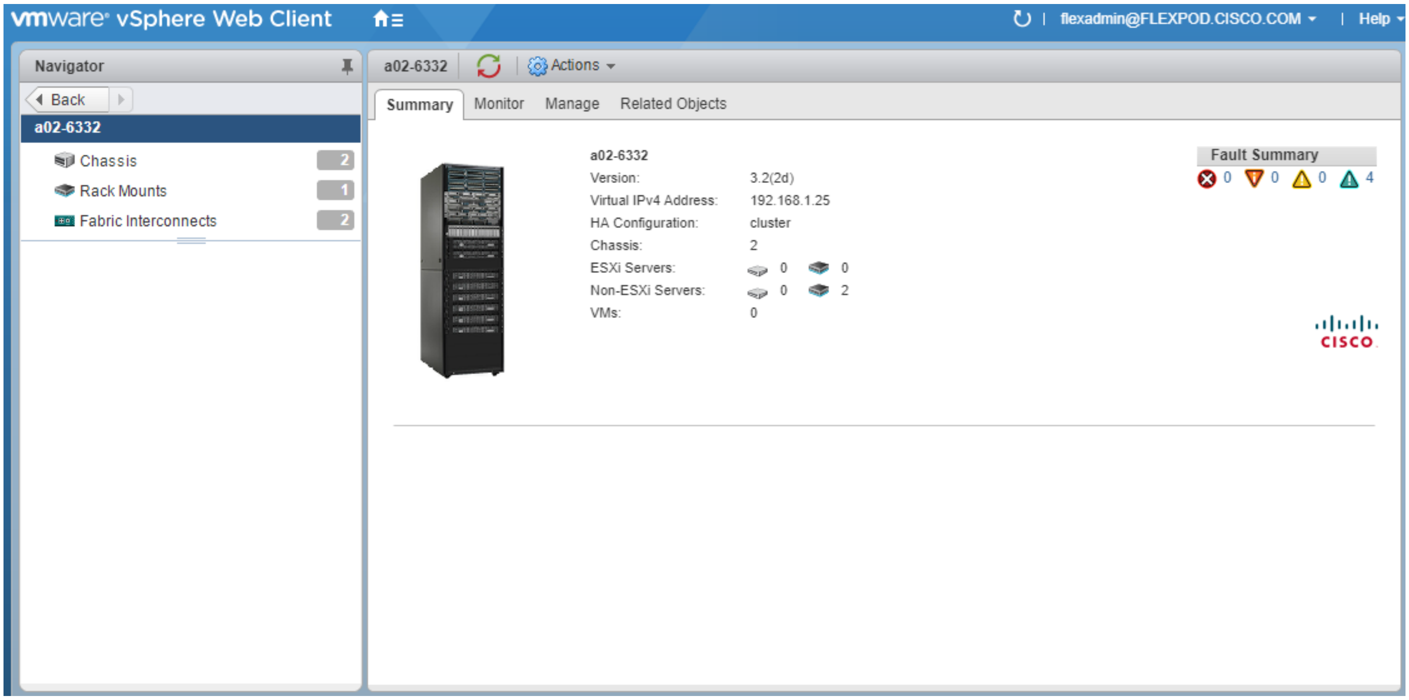

a02-6332

Version:	3.2(2d)
Virtual IPv4 Address:	192.168.1.25
HA Configuration:	cluster
Chassis:	2
ESXi Servers:	0 0
Non-ESXi Servers:	0 2
VMs:	0

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

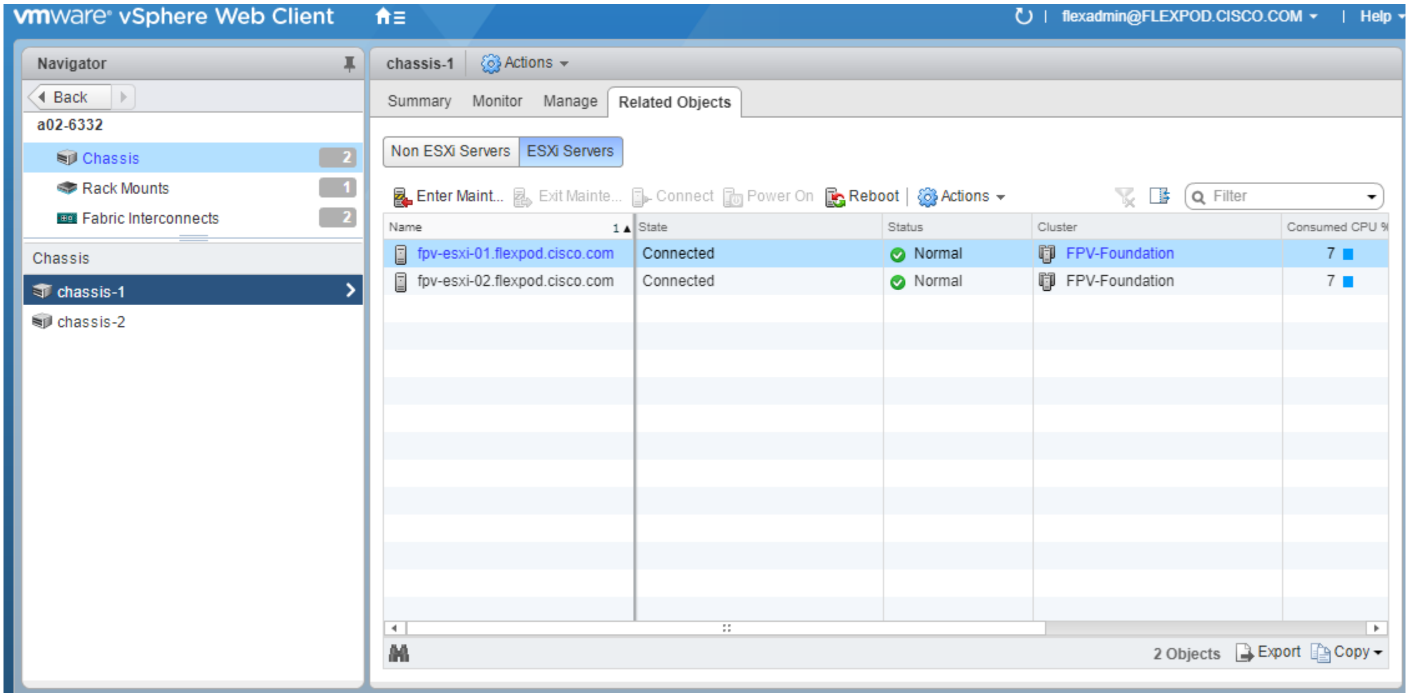

Name	State	Status	Cluster	Consumed CPU %
fpv-esxi-01.flexpod.cisco.com	Connected	Normal	FPV-Foundation	7
fpv-esxi-02.flexpod.cisco.com	Connected	Normal	FPV-Foundation	7

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

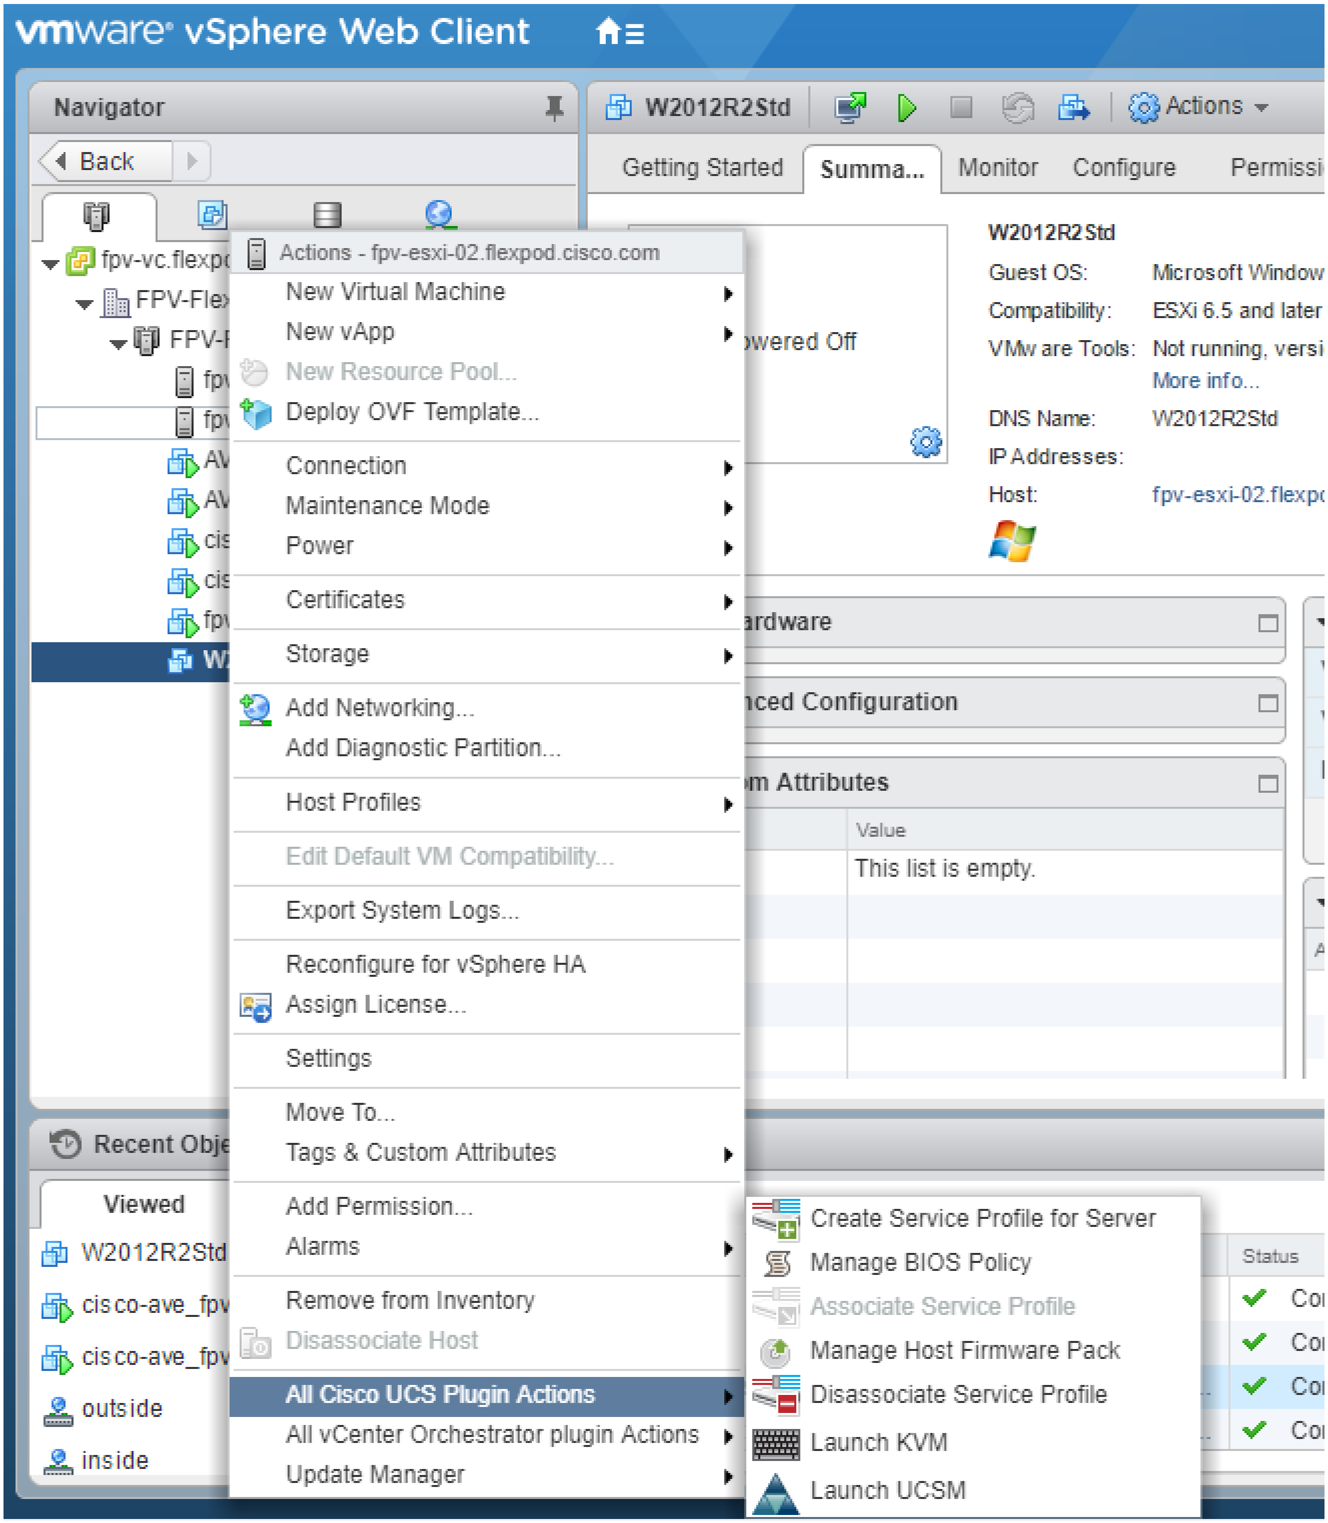

For the installation instructions and usage information, please refer to the Cisco UCS Manager Plug-in for VMware vSphere Web Client User Guide.

Build a Windows Server 2016 Virtual Machine for Cloning

A Windows Server 2016 virtual machine can be built in the VMware environment and cloned to create other virtual machines for management and tenant functions.

1. In the VMware vSphere Web Client, build a virtual machine named Win2016-DC-GUI, in one of the NFS datastores, compatible with ESXi6.5 and later, with Microsoft Windows Server 2016, and the following hardware:
   - 2 CPUs
   - 4096 MB Memory
   - 120 GB, Thin Provisioned Hard Disk
   - VMXNET 3 Network Adapter on the IB-MGMT Port Group on the vDS
2. Edit the settings of the VM and enable the Force BIOS setup VM option.
3. Boot the VM and open a console Window. Map a Windows Server 2016 installation iso to the CD drive. Set the BIOS boot order to Boot from CD-ROM Drive before Hard Drive. Save and Exit the BIOS.
4. Install Windows Server 2016 Datacenter with Desktop Experience on the VM, assign it an IP address and hostname, do not join the VM to the Windows Domain, and install all Windows Updates on the VM.
5. Shut down the VM.

Build Windows Active Directory Servers for ACI Fabric Core Services

Two Windows Server 2016 virtual machines will be cloned from the Win2016-DC-GUI VM and provisioned as Active Directory (AD) Domain Controllers in the existing AD Domain.

1. Create two clones of the Win2016-DC-GUI VM connected to the Core-Services EPG in ACI tenant common by right-clicking the Win2016-DC-GUI VM in vCenter and selecting Clone > Clone to Virtual Machine. Place one of these VMs in infra_datastore_1 on fpv-esxi-01 and the other in infra_datastore_2 on fpv-esxi-02.
2. Boot each clone and sysprep it. Then assign the VM an IP address and hostname. Do not join the VM to the AD domain.
3. Install Active Directory Domain Services on each VM and make it a Domain Controller and DNS server in the AD domain. Ensure the DNS server Forwarders are set correctly.
4. Add a persistent route to each VM to route to the tenant IP address space (172.16.0.0/16 in this validation) through the Core-Services EPG gateway address (10.1.118.254):

   ```
   route ADD -p 172.16.0.0 MASK 255.255.0.0 10.1.118.254

   route print
   ```
5. Reset the DNS in all existing VMs, servers, and hardware components to point to the two just-created DNS servers.
6. The two new AD servers can be placed in a separate site if the original AD server is in a different subnet and will not be reachable from tenant subnets.

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 where <core-services-EPG-gateway> = 10.1.118.254:

```
route –p ADD 172.16.0.0 MASK 255.255.0.0 10.1.118.254

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.

The VMware console can be used to connect to the CLI instead of ssh.

2. Type the following commands:

```
shell

echo “[Route]” >> /etc/systemd/network/10-eth0.network

echo “Gateway=10.1.118.254” >> /etc/systemd/network/10-eth0.network

echo “Destination=172.16.0.0/16” >> /etc/systemd/network/10-eth0.network

cat etc/systemd/network/10-eth0.network
```

3. The file should look like the following:

```
[Match]

Name=eth0


[Network]

Gateway=10.1.118.1

Address=10.1.118.101/24

DHCP=no


[DHCP]

UseDNS=false
```

```
[Route]

Gateway=10.1.118.254

Destination=172.16.0.0/16
```

4. Type the following commands:

```
systemctl restart systemd-networkd

ip route show
```

Adding the Supernet Route in VMware ESXi

To add a persistent Supernet Route in VMware ESXi, if placing the ESXi management VMkernel port in the Core-Services network, complete the following steps:

1. Using an ssh client, connect to the VMware ESXi CLI and login as root.

SSH will need to be enabled in the VMware ESXi Host Security Settings.

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 10.1.118.254 --network 172.16.0.0/16

esxcfg-route -l
```

NetApp FlexPod Management Tools Setup

NetApp Virtual Storage Console 7.1 Deployment Procedure

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

Virtual Storage Console 7.1 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.3:

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

Install Virtual Storage Console Appliance 7.1

To install the VSC 7.1 Appliance, complete the following steps:

1. Connect to the NetApp Support Site, and download the NetApp Virtual Storage Console 7.1 full installation package at https://mysupport.netapp.com/NOW/download/software/vsc_win/7.1/.
2. Log in to the vSphere Web Client as the FlexPod Admin user.
3. Go to Home > Hosts & Clusters.
4. Right-click the Datacenter and select Deploy OVF Template.
5. Click Browse and browse to and select the downloaded unified-virtual-applicance-for-vsc-vp-sra-7.1.ova file. Click Open.

The downloaded file should be on a local desktop disk.

6. Click Next.
7. Name the appliance and make sure the Datacenter is selected. Click Next.

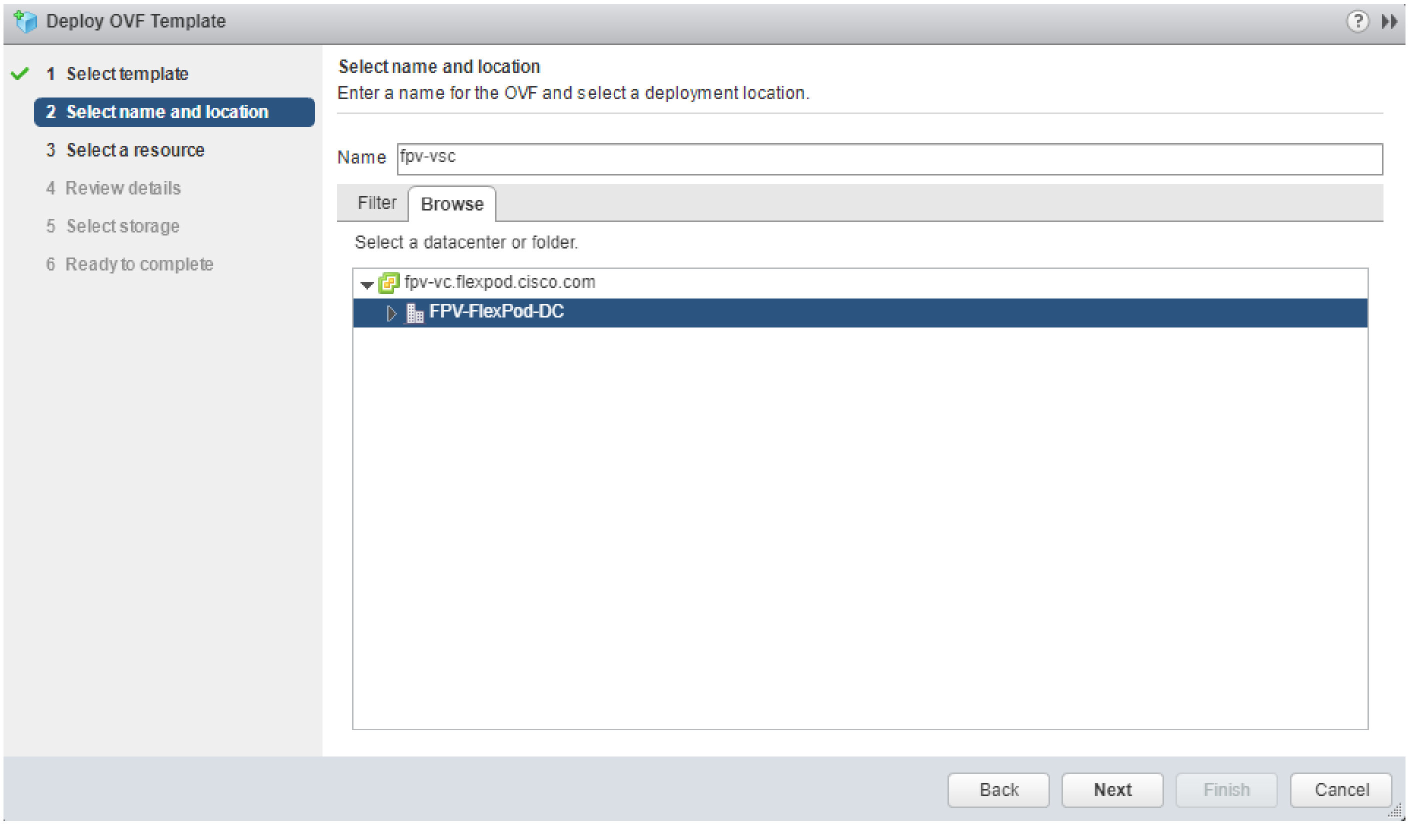

Deploy OVF Template

Select name and location

Enter a name for the OVF and select a deployment location.

Name: fpv-vsc

Select a datacenter or folder.

- fpv-vc.flexpod.cisco.com
  - FPV-FlexPod-DC

8. Select the FPV-Foundation ESXi cluster and click Next.

9. Review the details and click Next.

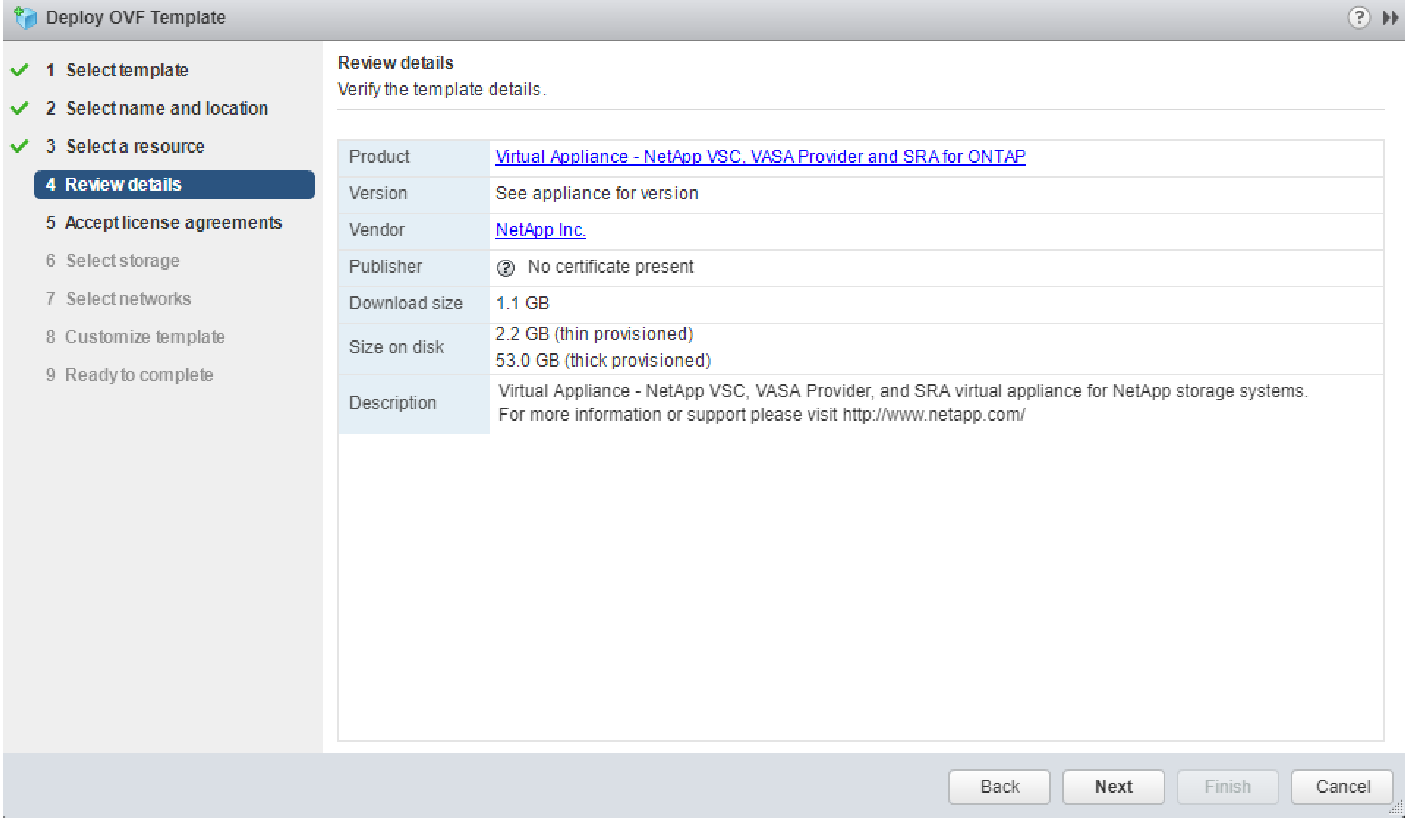

Review details

Verify the template details.

Product	Virtual Appliance - NetApp VSC, VASA Provider and SRA for ONTAP
Version	See appliance for version
Vendor	NetApp Inc.
Publisher	? No certificate present
Download size	1.1 GB
Size on disk	2.2 GB (thin provisioned) 53.0 GB (thick provisioned)
Description	Virtual Appliance - NetApp VSC, VASA Provider, and SRA virtual appliance for NetApp storage systems. For more information or support please visit http://www.netapp.com/

10. Click Accept to accept the license and click Next.

11. Select the Thin provision virtual disk format and one of the NFS datastores. Click Next.

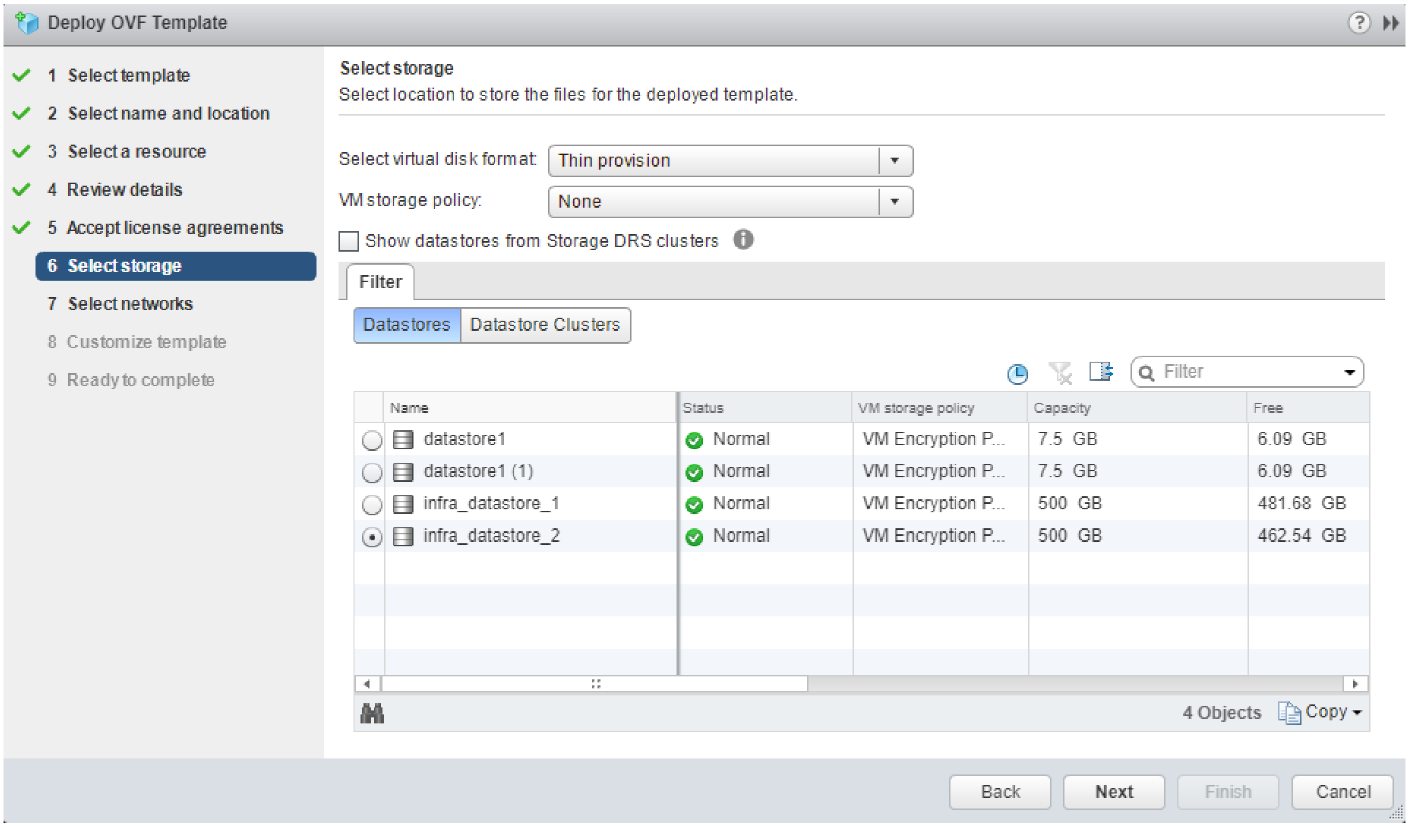

Deploy OVF Template

1 Select template
2 Select name and location
3 Select a resource
4 Review details
5 Accept license agreements
6 Select storage
7 Select networks
8 Customize template
9 Ready to complete

Select storage

Select location to store the files for the deployed template.

Select virtual disk format: Thin provision

VM storage policy: None

□ Show datastores from Storage DRS clusters

Filter

Datastores | Datastore Clusters

	Name	Status	VM storage policy	Capacity	Free
○	datastore1	Normal	VM Encryption P...	7.5 GB	6.09 GB
○	datastore1 (1)	Normal	VM Encryption P...	7.5 GB	6.09 GB
○	infra_datastore_1	Normal	VM Encryption P...	500 GB	481.68 GB
●	infra_datastore_2	Normal	VM Encryption P...	500 GB	462.54 GB

4 Objects

12. Select the Core-Services Port Group from Tenant common on the vDS. Click Next.

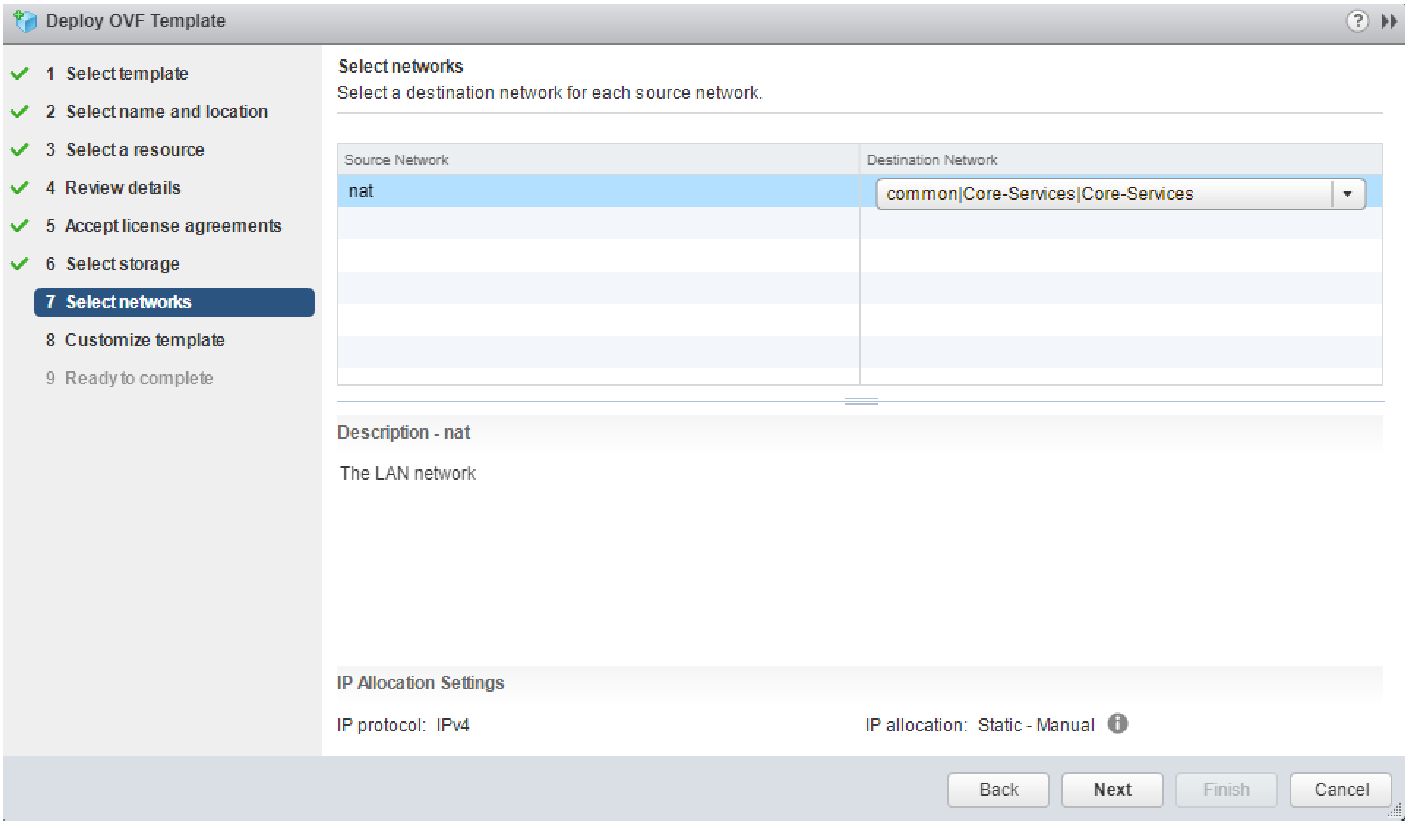

Deploy OVF Template

1 Select template
2 Select name and location
3 Select a resource
4 Review details
5 Accept license agreements
6 Select storage
7 Select networks
8 Customize template
9 Ready to complete

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	common|Core-Services|Core-Services

Description - nat

The LAN network

IP Allocation Settings

IP protocol: IPv4

IP allocation: Static - Manual

13. Expand and fill in all network information and vCenter registration information. Assign the VM an IP address in the infrastructure IB-MGMT subnet. Click Next.

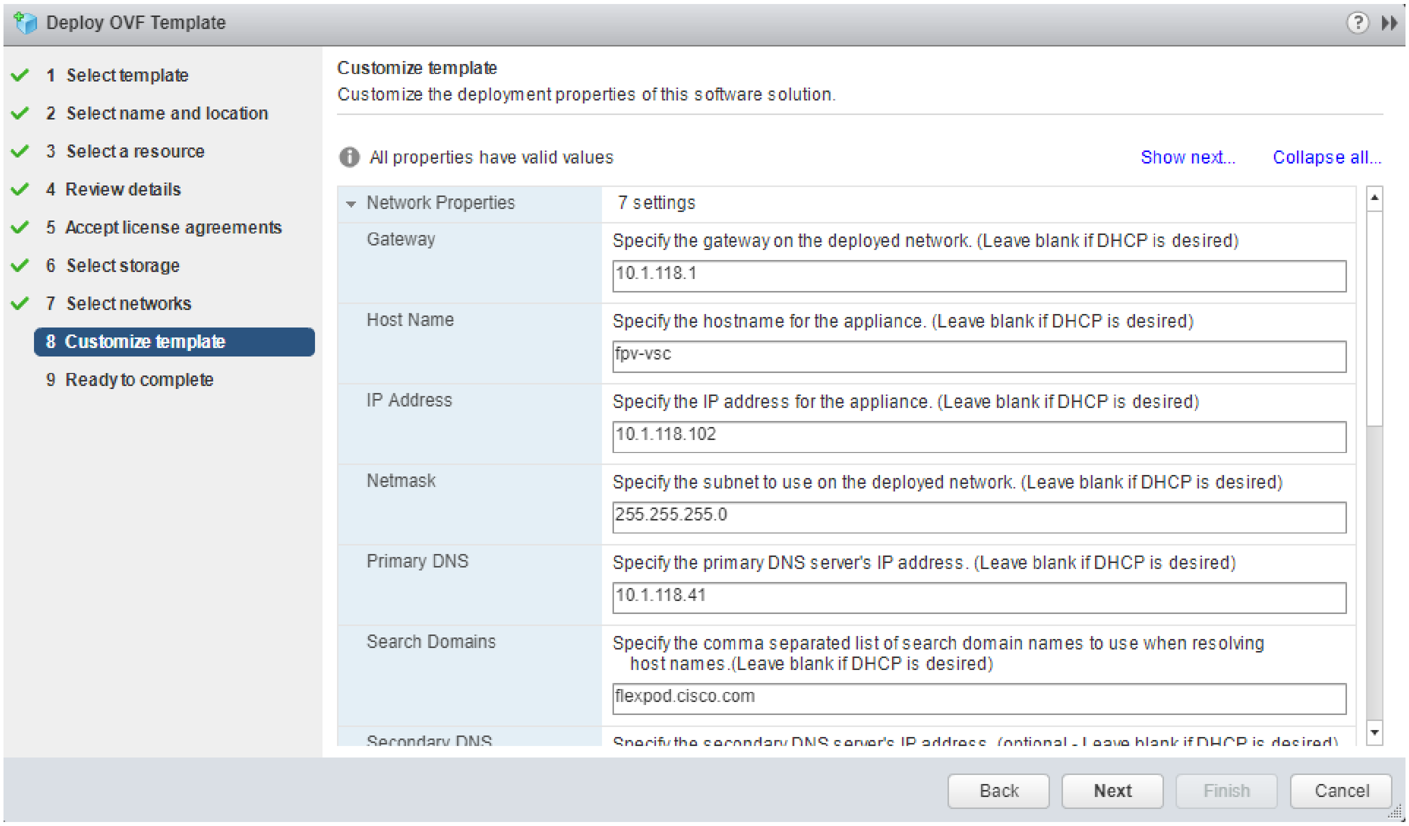

Deploy OVF Template

- ✓ 1 Select template
- ✓ 2 Select name and location
- ✓ 3 Select a resource
- ✓ 4 Review details
- ✓ 5 Accept license agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template
- 9 Ready to complete

Customize template
Customize the deployment properties of this software solution.

All properties have valid values — Show next... Collapse all...

Network Properties	7 settings
Gateway	Specify the gateway on the deployed network. (Leave blank if DHCP is desired) 10.1.118.1
Host Name	Specify the hostname for the appliance. (Leave blank if DHCP is desired) fpv-vsc
IP Address	Specify the IP address for the appliance. (Leave blank if DHCP is desired) 10.1.118.102
Netmask	Specify the subnet to use on the deployed network. (Leave blank if DHCP is desired) 255.255.255.0
Primary DNS	Specify the primary DNS server's IP address. (Leave blank if DHCP is desired) 10.1.118.41
Search Domains	Specify the comma separated list of search domain names to use when resolving host names.(Leave blank if DHCP is desired) flexpod.cisco.com
Secondary DNS	Specify the secondary DNS server's IP address. (optional - Leave blank if DHCP is desired)

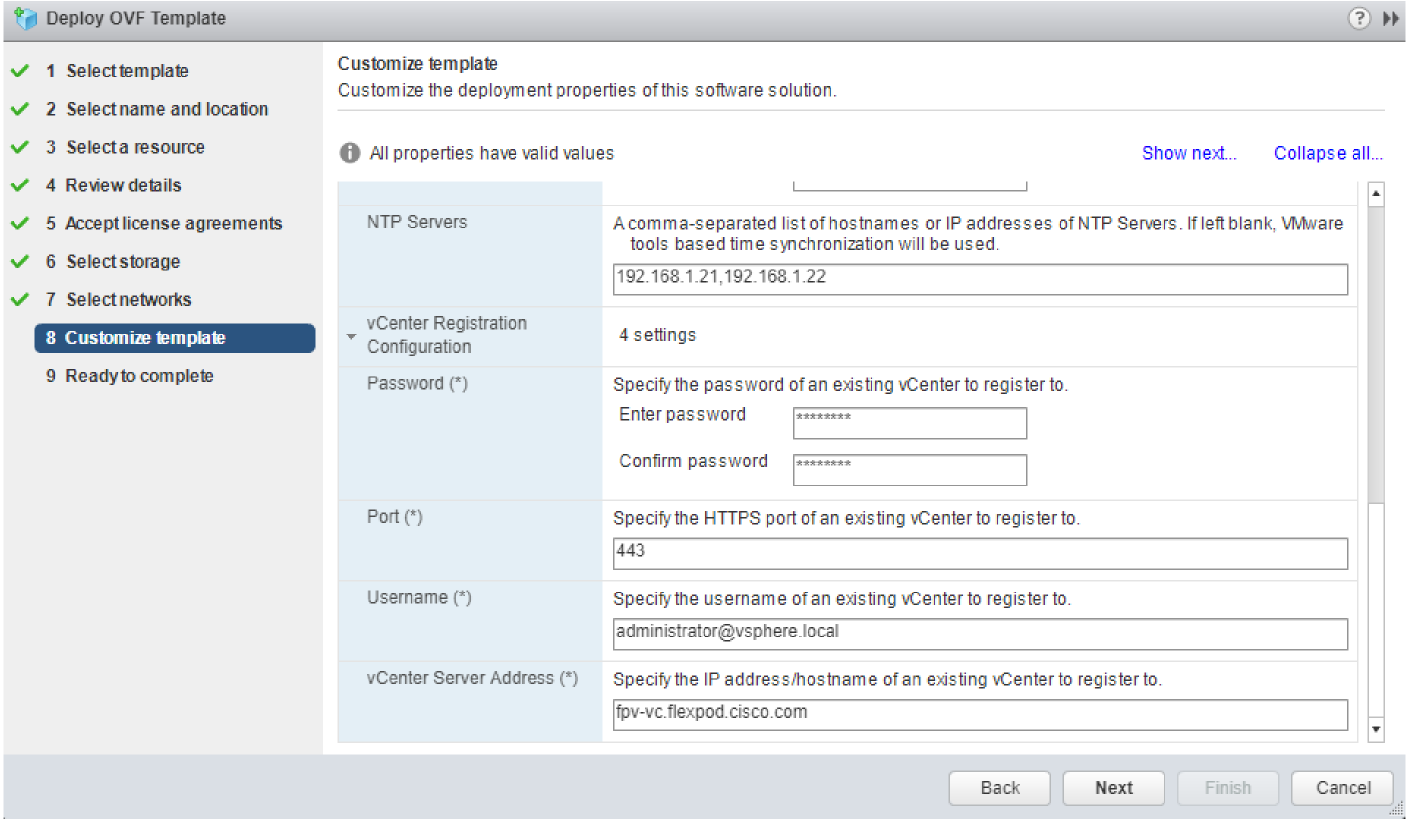

Deploy OVF Template

- ✓ 1 Select template
- ✓ 2 Select name and location
- ✓ 3 Select a resource
- ✓ 4 Review details
- ✓ 5 Accept license agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template
- 9 Ready to complete

Customize template
Customize the deployment properties of this software solution.

All properties have valid values — Show next... Collapse all...

NTP Servers	A comma-separated list of hostnames or IP addresses of NTP Servers. If left blank, VMware tools based time synchronization will be used. 192.168.1.21,192.168.1.22
vCenter Registration Configuration	4 settings
Password (*)	Specify the password of an existing vCenter to register to. Enter password ******** Confirm password ********
Port (*)	Specify the HTTPS port of an existing vCenter to register to. 443
Username (*)	Specify the username of an existing vCenter to register to. administrator@vsphere.local
vCenter Server Address (*)	Specify the IP address/hostname of an existing vCenter to register to. fpv-vc.flexpod.cisco.com

14. Review the installation information. Click Finish to install the appliance.

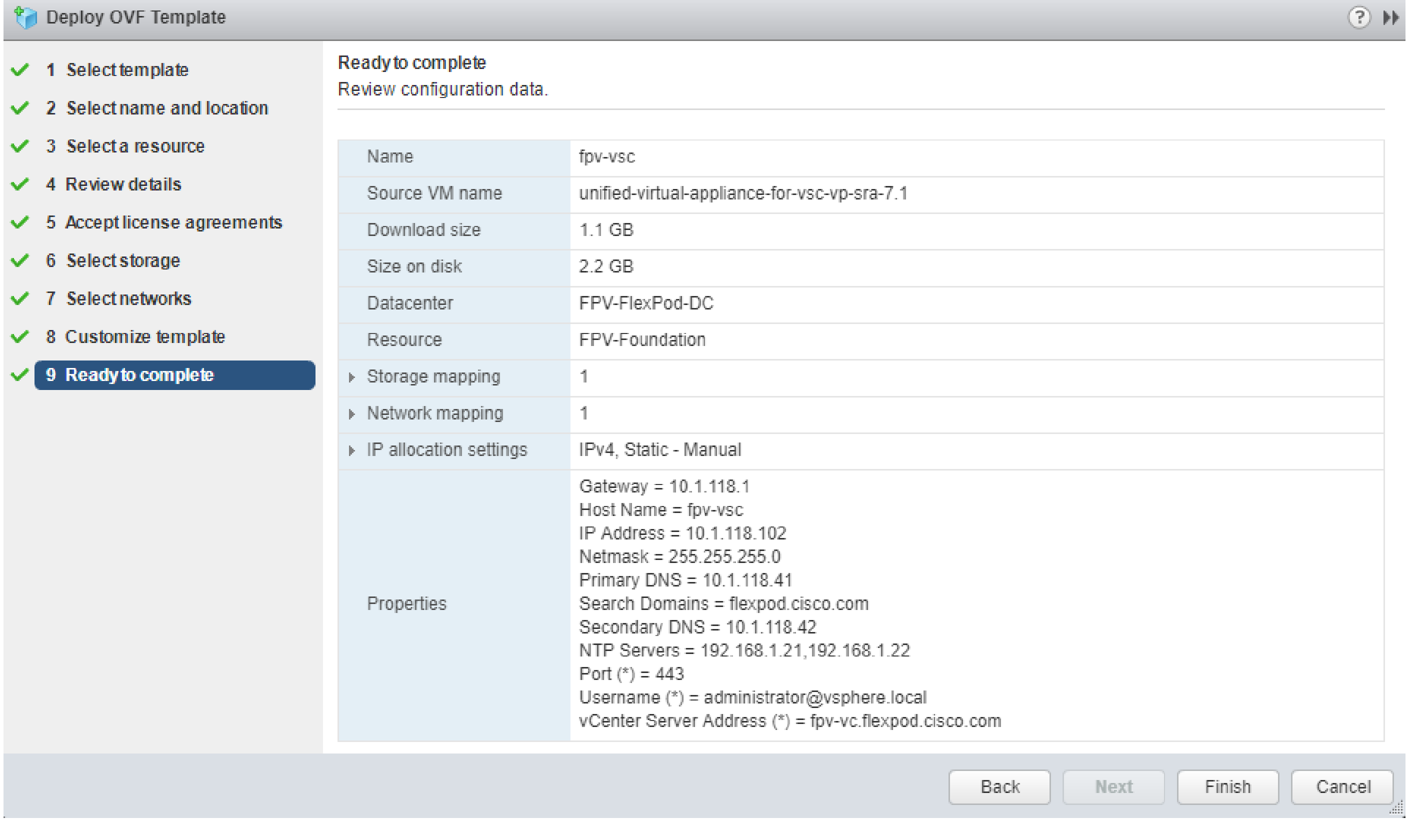

Ready to complete

Review configuration data.

Name	fpv-vsc
Source VM name	unified-virtual-appliance-for-vsc-vp-sra-7.1
Download size	1.1 GB
Size on disk	2.2 GB
Datacenter	FPV-FlexPod-DC
Resource	FPV-Foundation
Storage mapping	1
Network mapping	1
IP allocation settings	IPv4, Static - Manual
Properties	Gateway = 10.1.118.1 Host Name = fpv-vsc IP Address = 10.1.118.102 Netmask = 255.255.255.0 Primary DNS = 10.1.118.41 Search Domains = flexpod.cisco.com Secondary DNS = 10.1.118.42 NTP Servers = 192.168.1.21,192.168.1.22 Port (*) = 443 Username (*) = administrator@vsphere.local vCenter Server Address (*) = fpv-vc.flexpod.cisco.com

15. Once the OVF deployment completes, power on the VM.

16. Right-click the deployed VM and select Guest OS > Install VMware Tools.

17. Click Mount. VMware Tools will automatically install on the appliance and it will reboot.

18. Log out and log back into the vSphere Web Client.

Add the Supernet Route to the VSC Appliance and Secure the Appliance

Since the NetApp VSC Appliance was placed in the Core-Services network, it will need to have the Supernet route added to communicate with tenant VMs. Complete the following steps.

1. From the vSphere Web Client, right-click the VSC VM and select Open Console.
2. Log into the appliance using the “maint” user id and “admin123” password.
3. Enter “2” to go to System Configuration.
4. Enter “3” to Change ‘maint’ user password.
5. Enter the “admin123” password and then enter and confirm a secure password.
6. Enter “b” to return to the Main Menu.
7. Enter “3” to go to Network Configuration.

8. Enter “6” to Change static routes.

9. Enter “1” to Add route.

10. Enter the Supernet destination address <172.16.0.0/16>.

11. Enter the Supernet gateway address <10.1.118.254>.

12. Enter “y” to confirm the settings.

13. Press Enter.

14. Enter “7” to Commit changes.

15. Enter “y” to commit the changes.

16. Press Enter.

17. Enter “5” to Display static routes.

18. Press Enter.

19. Enter “x” to exit the console.

20. Close the console window.

Install NetApp NFS VAAI Plug-in

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

1. Download the NetApp NFS Plug-in 1.1.2 for VMware `.vib` file from the NFS Plugin Download to the local desktop.

2. Rename the downloaded file `NetAppNasPlugin.vib`.

3. In the vSphere Web Client, open Home > Virtual Storage Console.

4. On the left, select NFS VAAI Tools.

5. In the center pane, click Select File and browse to the downloaded NetAppNasPlugin.vib. Click Open.

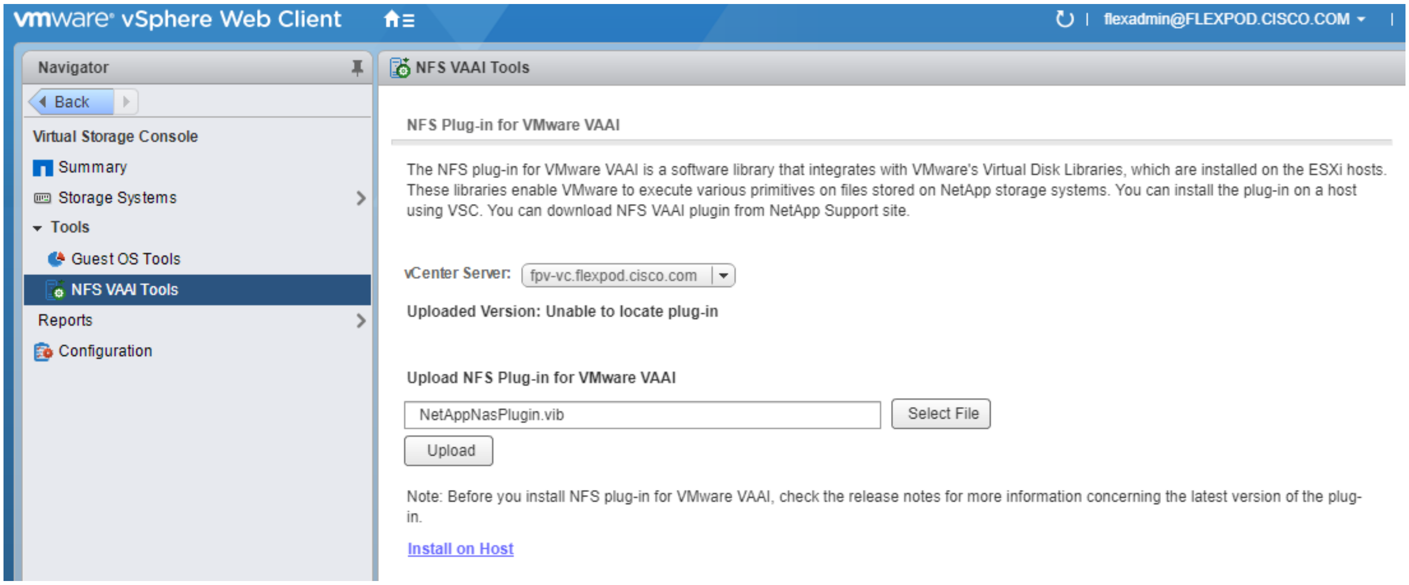

NFS Plug-in for VMware VAAI

The NFS plug-in for VMware VAAI is a software library that integrates with VMware's Virtual Disk Libraries, which are installed on the ESXi hosts. These libraries enable VMware to execute various primitives on files stored on NetApp storage systems. You can install the plug-in on a host using VSC. You can download NFS VAAI plugin from NetApp Support site.

vCenter Server: fpv-vc.flexpod.cisco.com

Uploaded Version: Unable to locate plug-in

Upload NFS Plug-in for VMware VAAI

NetAppNasPlugin.vib

Note: Before you install NFS plug-in for VMware VAAI, check the release notes for more information concerning the latest version of the plug-in.

6. Click Upload to upload the plugin to the VSC appliance.

7. The Install on Host link can now be used to install the plugin on the ESXi hosts.

Note that the NFS Plug-in for VMware VAAI has already been installed on the two management hosts in this document.

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. In Virtual Storage Console, on the left select Storage Systems. Under the Objects tab, click Actions > Modify.

2. In the IP Address/Hostname field, enter the storage cluster management IP or FQDN. Enter admin for the user name and the admin password for password. Confirm Use TLS to Connect to This Storage System is selected. Click OK.

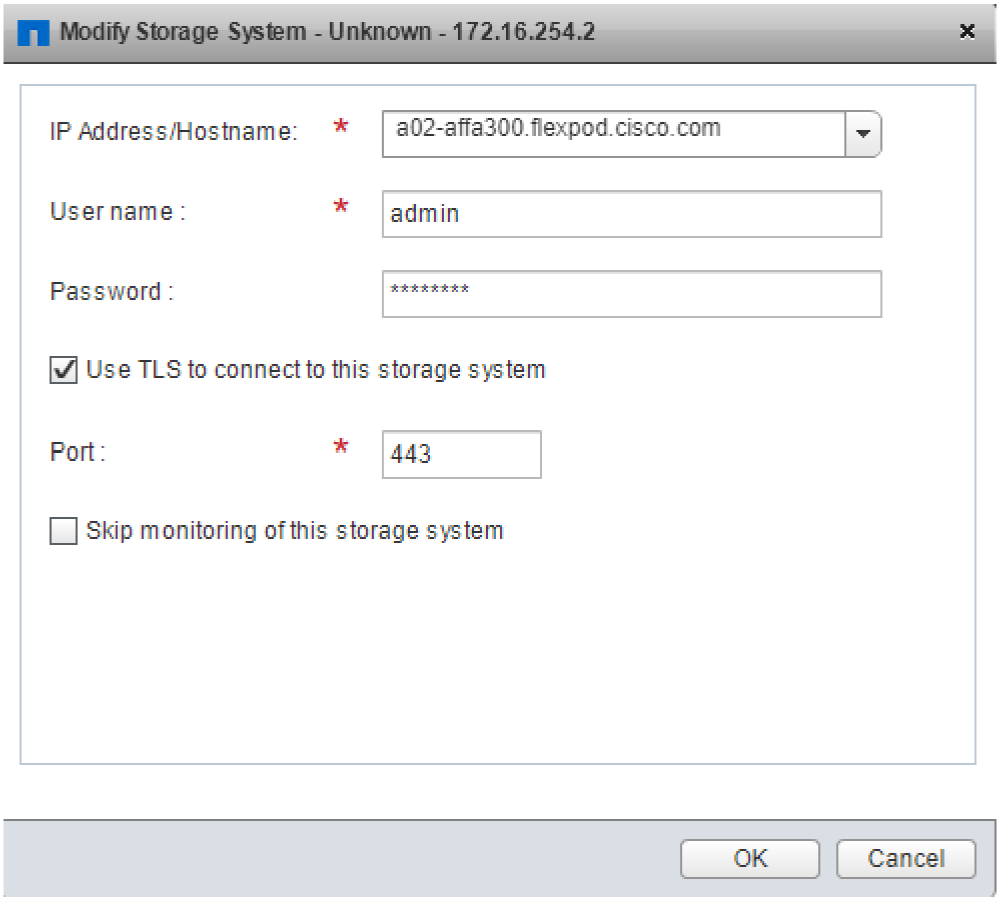

Modify Storage System - Unknown - 172.16.254.2

IP Address/Hostname: * a02-affa300.flexpod.cisco.com

User name : * admin

Password : ********

☑ Use TLS to connect to this storage system

Port : * 443

□ Skip monitoring of this storage system

3. Click OK to accept the controller privileges.

4. 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.

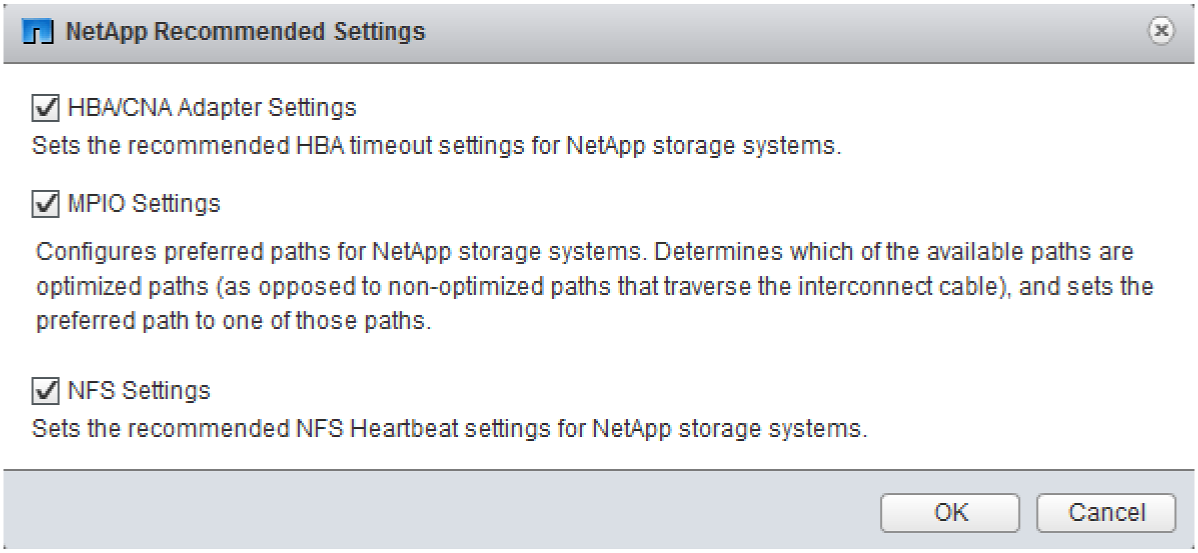

NetApp Recommended Settings

☑ HBA/CNA Adapter Settings

Sets the recommended HBA timeout settings for NetApp storage systems.

☑ MPIO Settings

Configures preferred paths for NetApp storage systems. Determines which of the available paths are optimized paths (as opposed to non-optimized paths that traverse the interconnect cable), and sets the preferred path to one of those paths.

☑ NFS Settings

Sets the recommended NFS Heartbeat settings for NetApp storage systems.

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

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

3. Click OK.

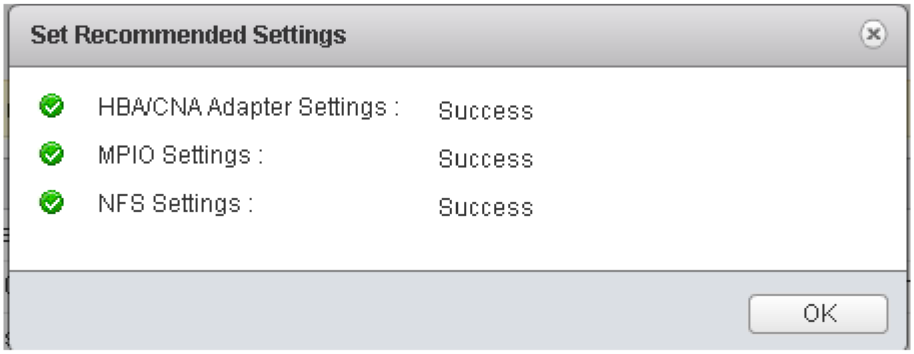

Set Recommended Settings

HBA/CNA Adapter Settings : Success

MPIO Settings : Success

NFS Settings : Success

4. 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.

5. Once all hosts have had their Optimized Settings set, these settings can be verified in the Virtual Storage Console tool.

Virtual Storage Console 7.1 Provisioning Datastores

Using VSC, the administrator can provision an NFS, FC or iSCSI datastore and attach it to a single host or multiple hosts in the cluster. The following steps illustrate provisioning a datastore and attaching it to the cluster.

Provisioning an NFS Datastore

To provision the NFS datastore, complete the following steps:

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

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

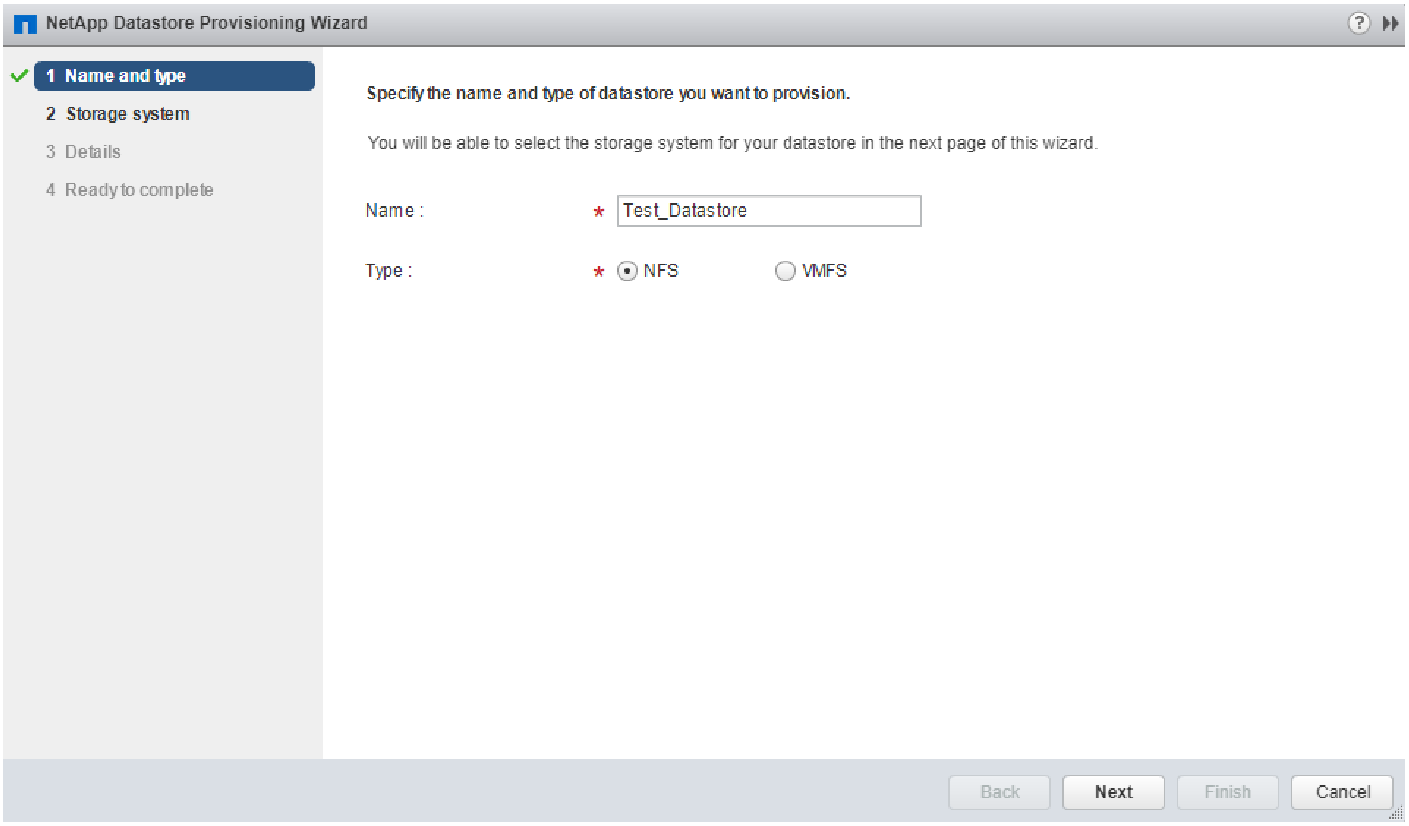

NetApp Datastore Provisioning Wizard

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

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

Name : * Test_Datastore

Type : * ● NFS ○ VMFS

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

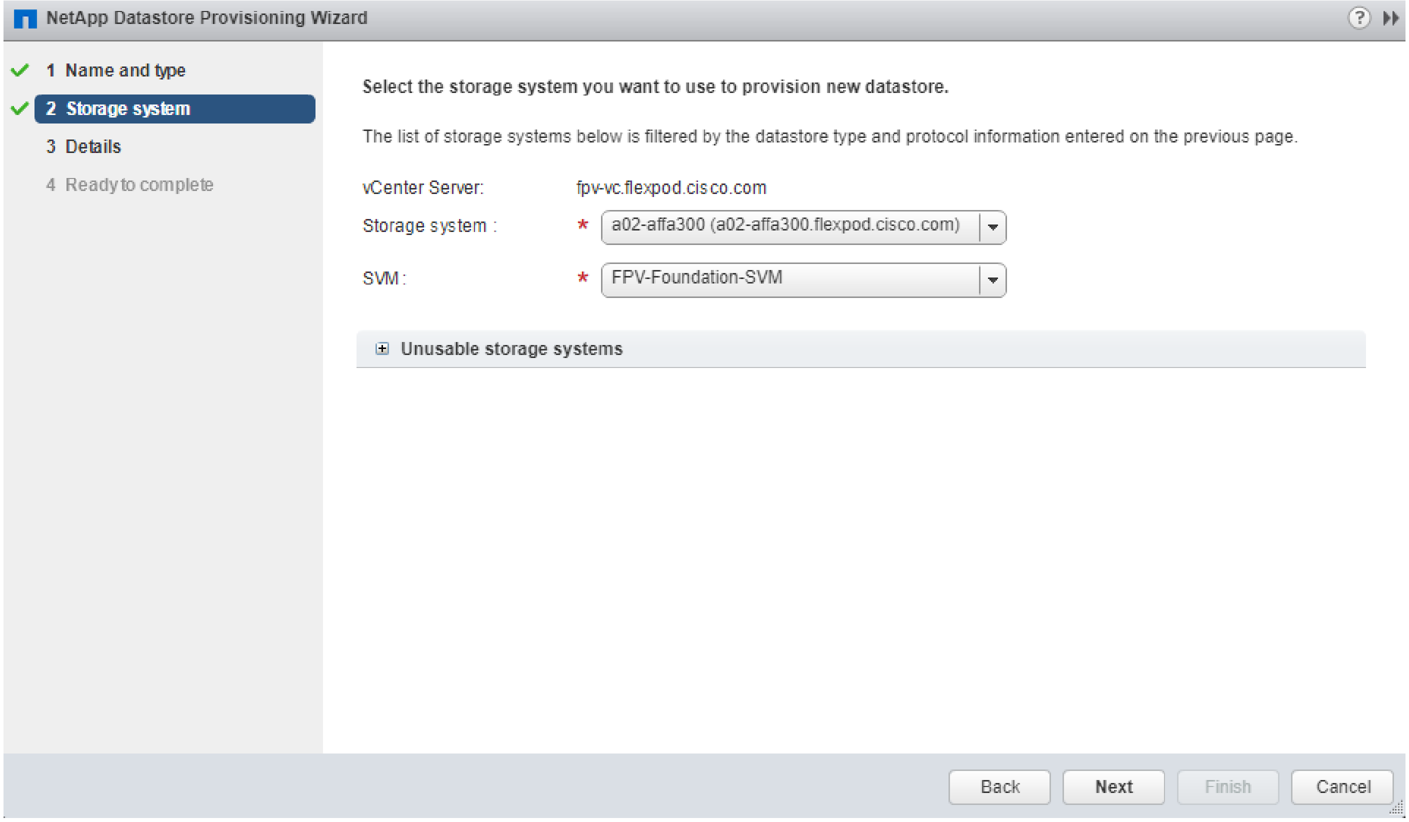

Select the storage system you want to use to provision new datastore.

The list of storage systems below is filtered by the datastore type and protocol information entered on the previous page.

vCenter Server: fpv-vc.flexpod.cisco.com

Storage system : * a02-affa300 (a02-affa300.flexpod.cisco.com)

SVM : * FPV-Foundation-SVM

6. Select whether to Thin provision the datastore, enter the size of the datastore and select the aggregate name.

7. Click Next.

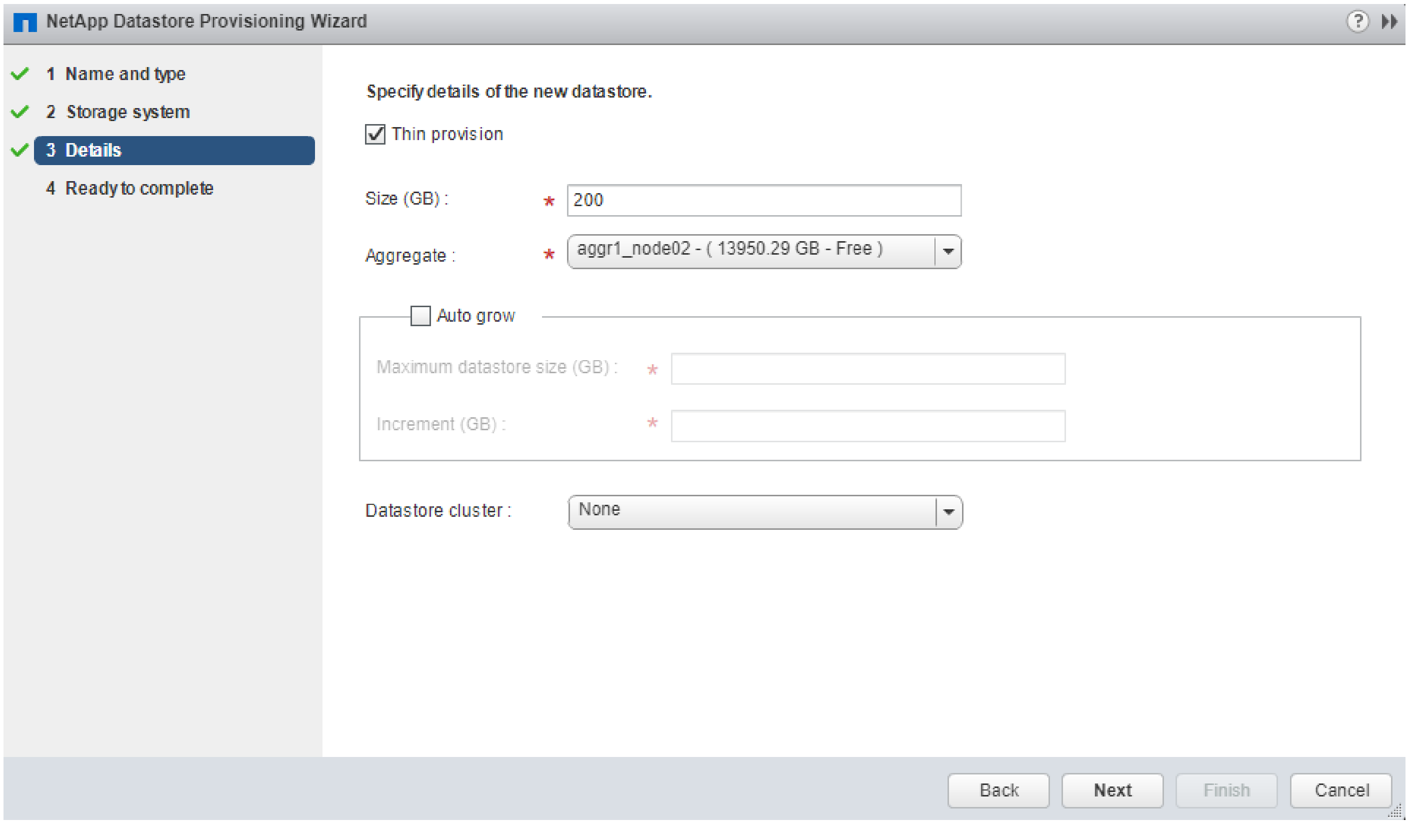

NetApp Datastore Provisioning Wizard

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

Specify details of the new datastore.

☑ Thin provision

Size (GB) : * 200

Aggregate : * aggr1_node02 - (13950.29 GB - Free)

□ Auto grow

Maximum datastore size (GB) : *

Increment (GB) : *

Datastore cluster : None

8. Review the details and click Finish.

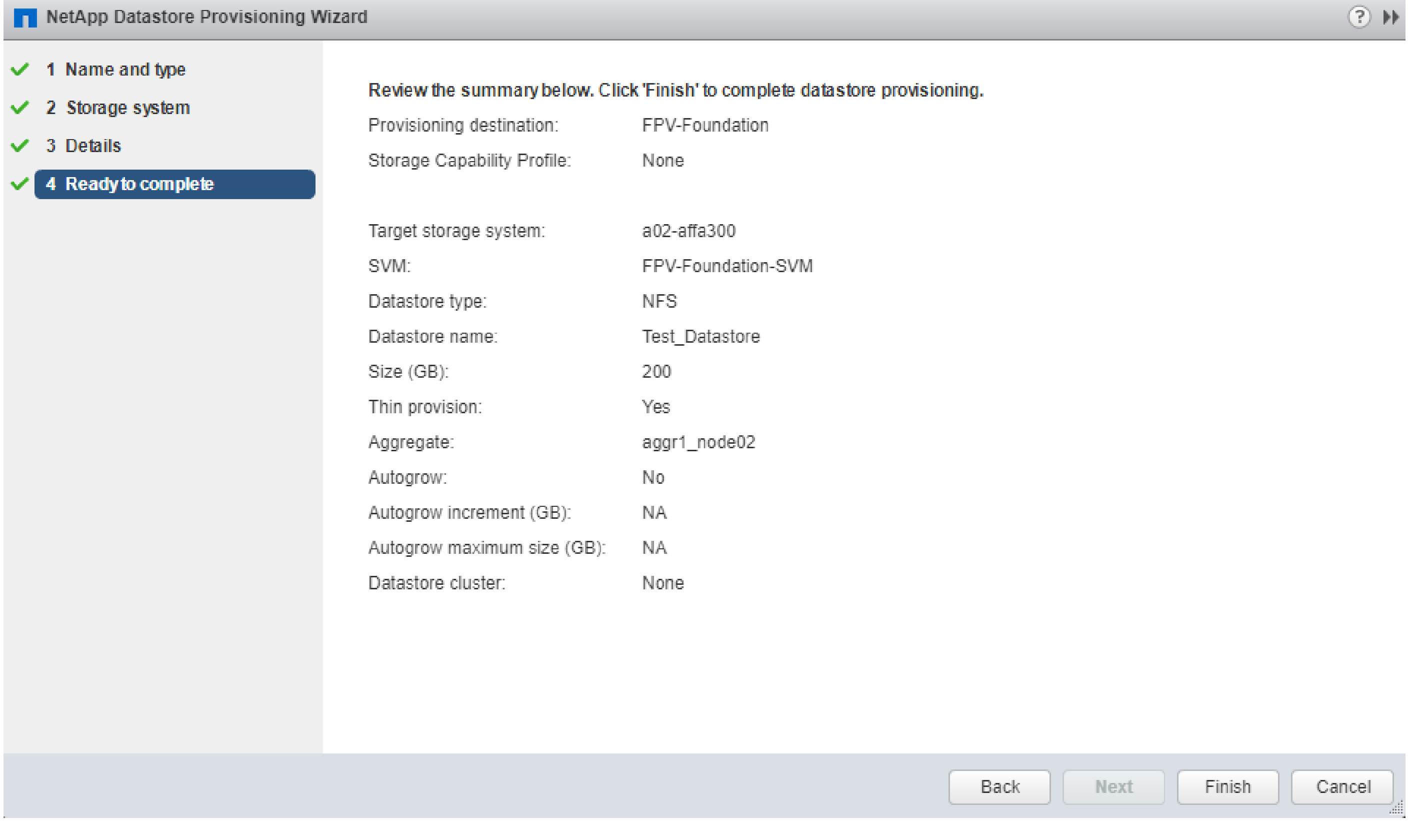

NetApp Datastore Provisioning Wizard

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

Review the summary below. Click 'Finish' to complete datastore provisioning.

Provisioning destination:	FPV-Foundation
Storage Capability Profile:	None
Target storage system:	a02-affa300
SVM:	FPV-Foundation-SVM
Datastore type:	NFS
Datastore name:	Test_Datastore
Size (GB):	200
Thin provision:	Yes
Aggregate:	aggr1_node02
Autogrow:	No
Autogrow increment (GB):	NA
Autogrow maximum size (GB):	NA
Datastore cluster:	None

9. Click Ok.

The datastore will be created and mounted on all the hosts in the cluster.

Provisioning iSCSI Datastore

To provision an iSCSI datastore, complete the following steps:

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

2. Enter the datastore name and select the type as VMFS. For VMFS protocol, select iSCSI.

3. Click Next.

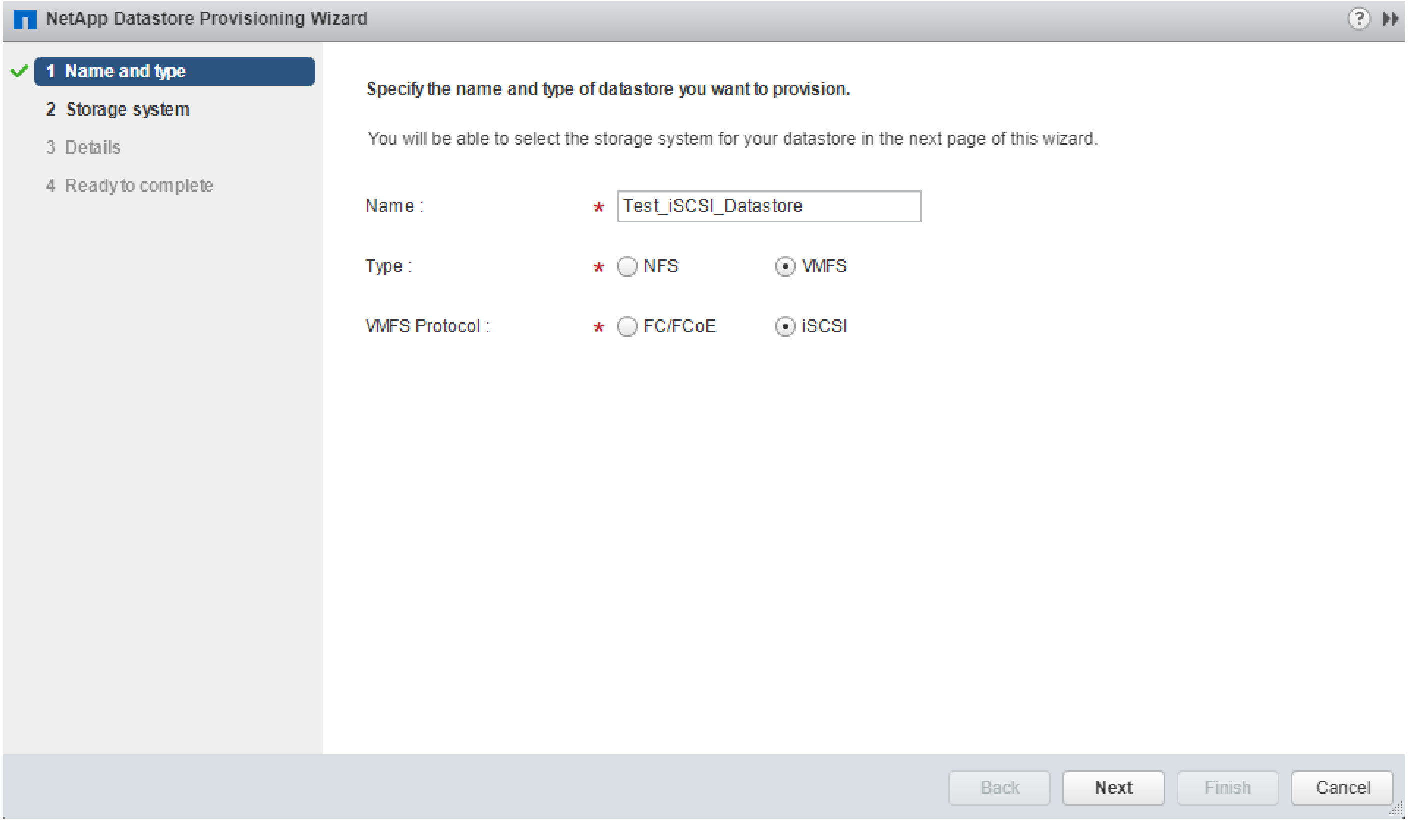

NetApp Datastore Provisioning Wizard

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

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

Name : * Test_iSCSI_Datastore

Type : * ○ NFS ● VMFS

VMFS Protocol : * ○ FC/FCoE ● iSCSI

4. Select the cluster name for the Storage system and the desired SVM to create the datastore. In this example, FPV-Foundation-SVM is selected.

5. Click Next.

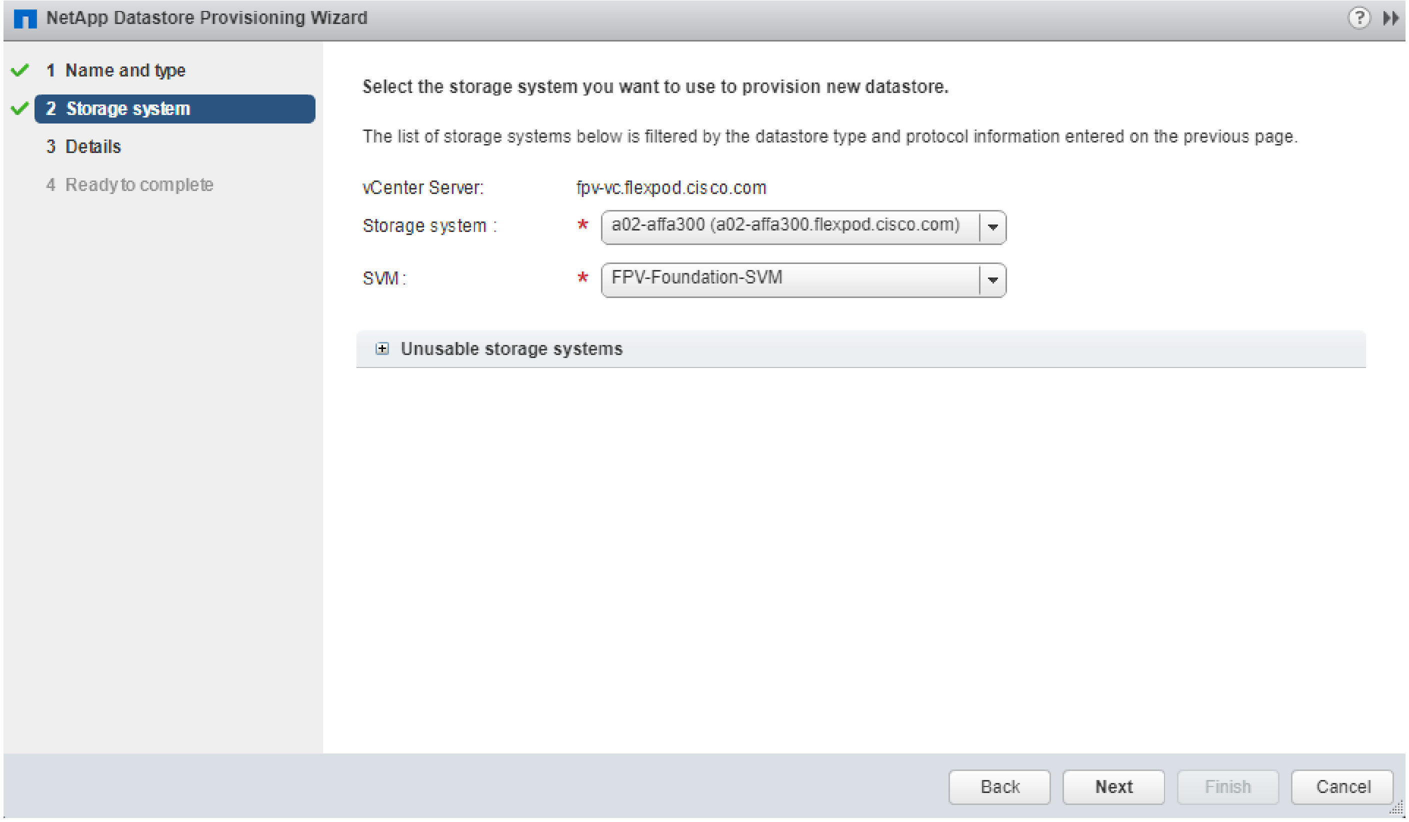

Select the storage system you want to use to provision new datastore.

The list of storage systems below is filtered by the datastore type and protocol information entered on the previous page.

vCenter Server: fpv-vc.flexpod.cisco.com

Storage system : * a02-affa300 (a02-affa300.flexpod.cisco.com)

SVM : * FPV-Foundation-SVM

6. Select whether to Thin provision the datastore, and enter the size of the datastore. Select the “Create new volume” check box and select the aggregate name.

7. Click Next.

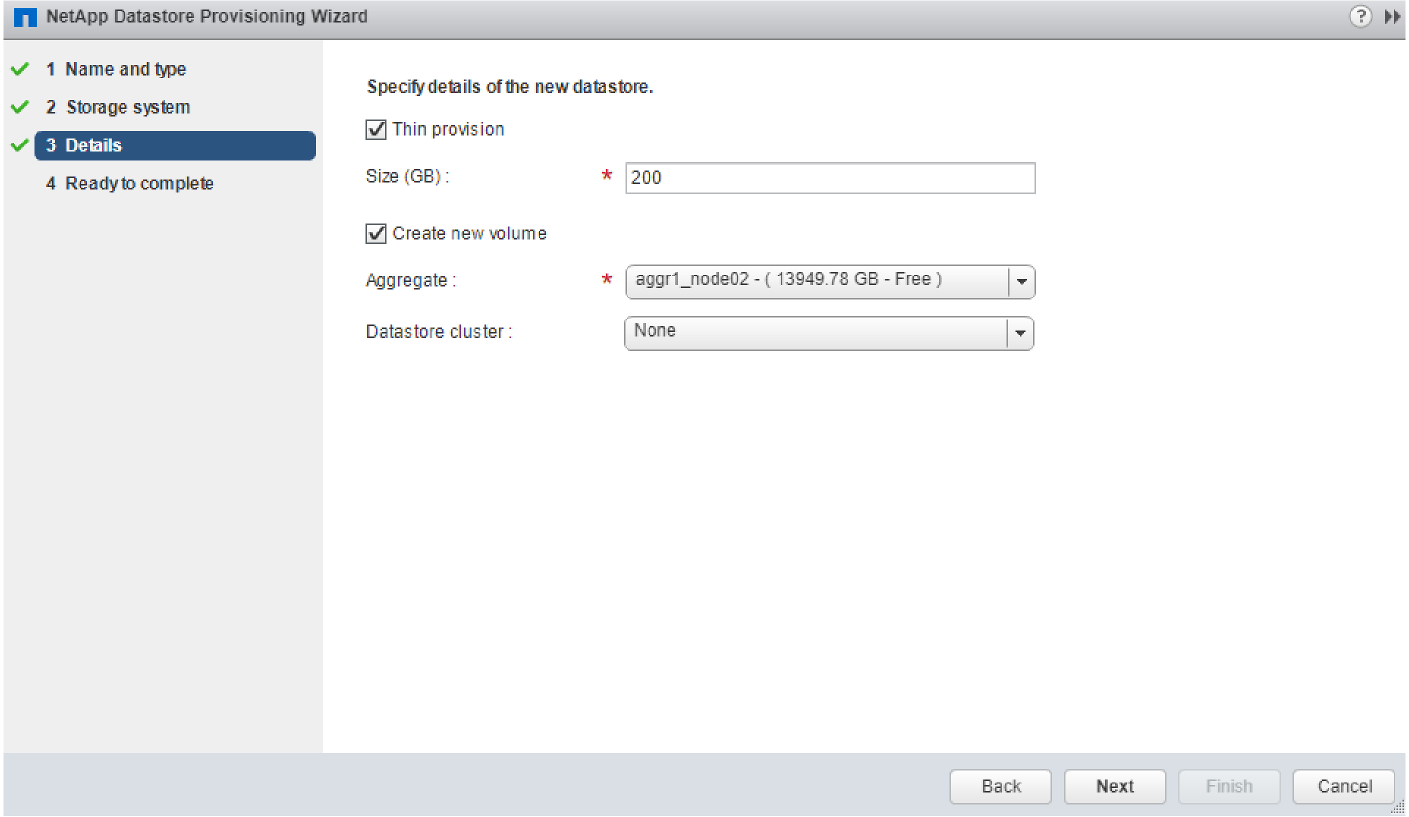

Specify details of the new datastore.

☑ Thin provision

Size (GB) : * 200

☑ Create new volume

Aggregate : * aggr1_node02 - (13949.78 GB - Free)

Datastore cluster : None

8. Review the details and click Finish.

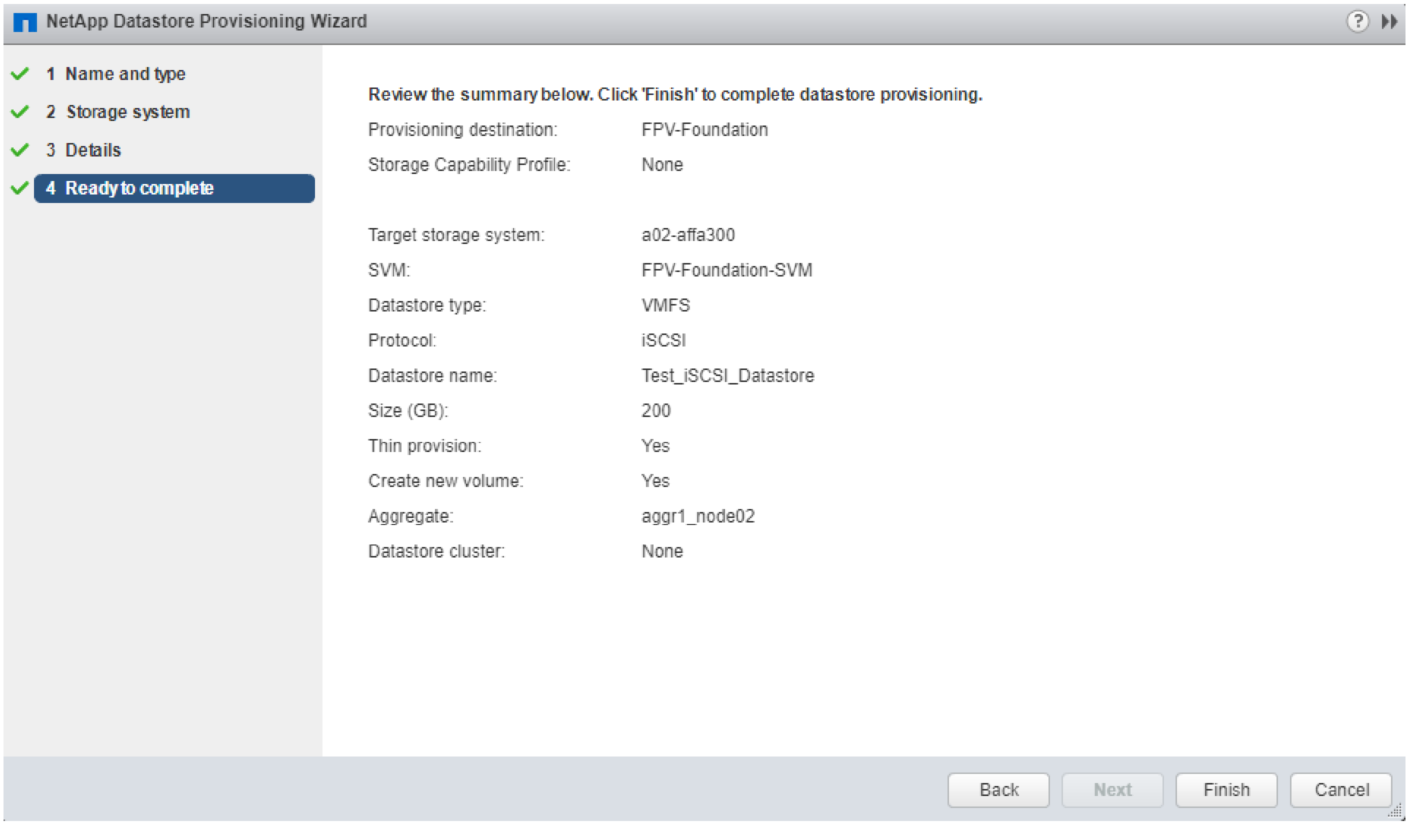

Review the summary below. Click 'Finish' to complete datastore provisioning.

Provisioning destination:	FPV-Foundation
Storage Capability Profile:	None
Target storage system:	a02-affa300
SVM:	FPV-Foundation-SVM
Datastore type:	VMFS
Protocol:	iSCSI
Datastore name:	Test_iSCSI_Datastore
Size (GB):	200
Thin provision:	Yes
Create new volume:	Yes
Aggregate:	aggr1_node02
Datastore cluster:	None

9. Click Ok.

The datastore will be created and mounted on all the hosts in the cluster.

NetApp SnapCenter

NetApp SnapCenter is a Windows application, with the following prerequisites:

- The server it resides on must be in an Active Directory domain.
- The server it resides on must have .NET 4.5.2 or later installed.
- The server it resides on should have at least 32 GB DRAM. (8 GB is the absolute minimum; less than 32 GB will generate a warning during installation.)
- It requires a service user account in the Active Directory domain.

Installing SnapCenter

To install SnapCenter, complete the following steps:

1. Download the SnapCenter installer (SnapCenter4.0.exe) from:

   https://mysupport.netapp.com/NOW/download/software/snapcenter/4.0/download.shtml

2. Double-click SnapCenter4.0.exe to start the installation.

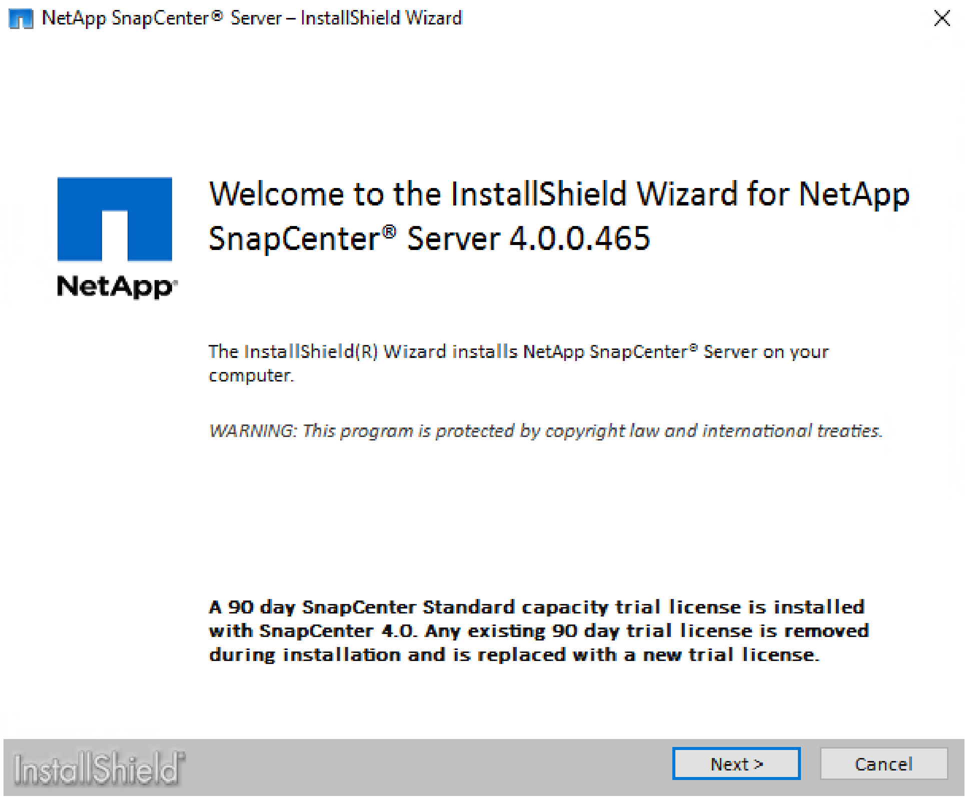

NetApp SnapCenter® Server – InstallShield Wizard

Welcome to the InstallShield Wizard for NetApp SnapCenter® Server 4.0.0.465

The InstallShield(R) Wizard installs NetApp SnapCenter® Server on your computer.

WARNING: This program is protected by copyright law and international treaties.

A 90 day SnapCenter Standard capacity trial license is installed with SnapCenter 4.0. Any existing 90 day trial license is removed during installation and is replaced with a new trial license.

3. Click Next on the Welcome screen.

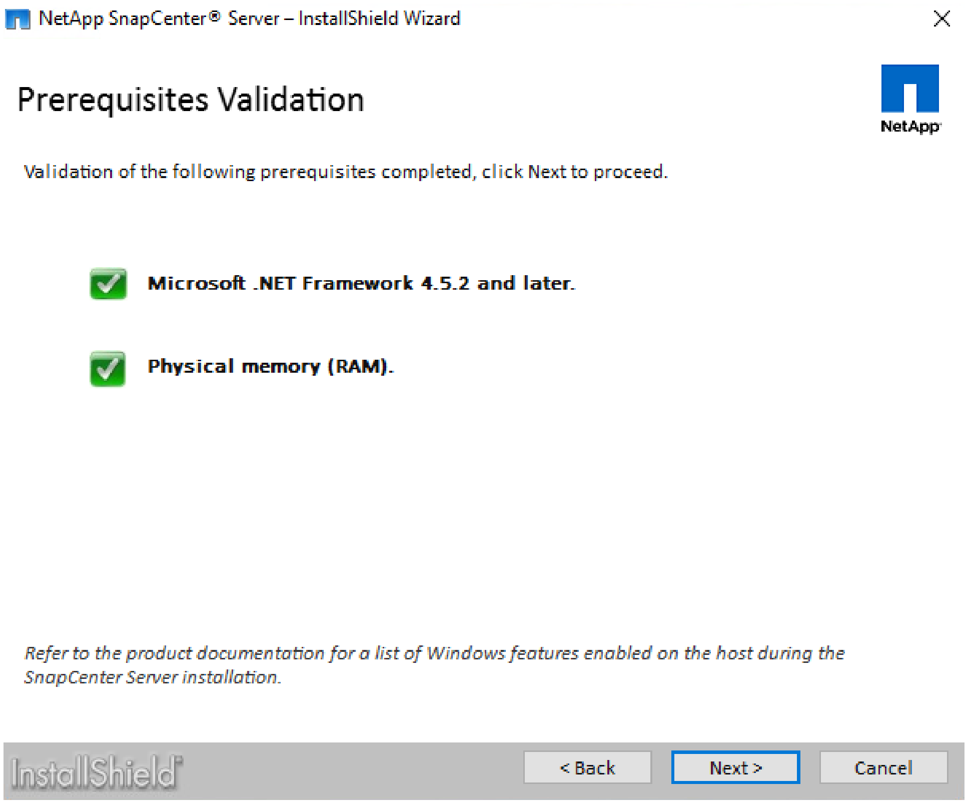

NetApp SnapCenter® Server – InstallShield Wizard

Prerequisites Validation

Validation of the following prerequisites completed, click Next to proceed.

✓ Microsoft .NET Framework 4.5.2 and later.

✓ Physical memory (RAM).

Refer to the product documentation for a list of Windows features enabled on the host during the SnapCenter Server installation.

4. Click Next on the Prerequisites Validation screen.

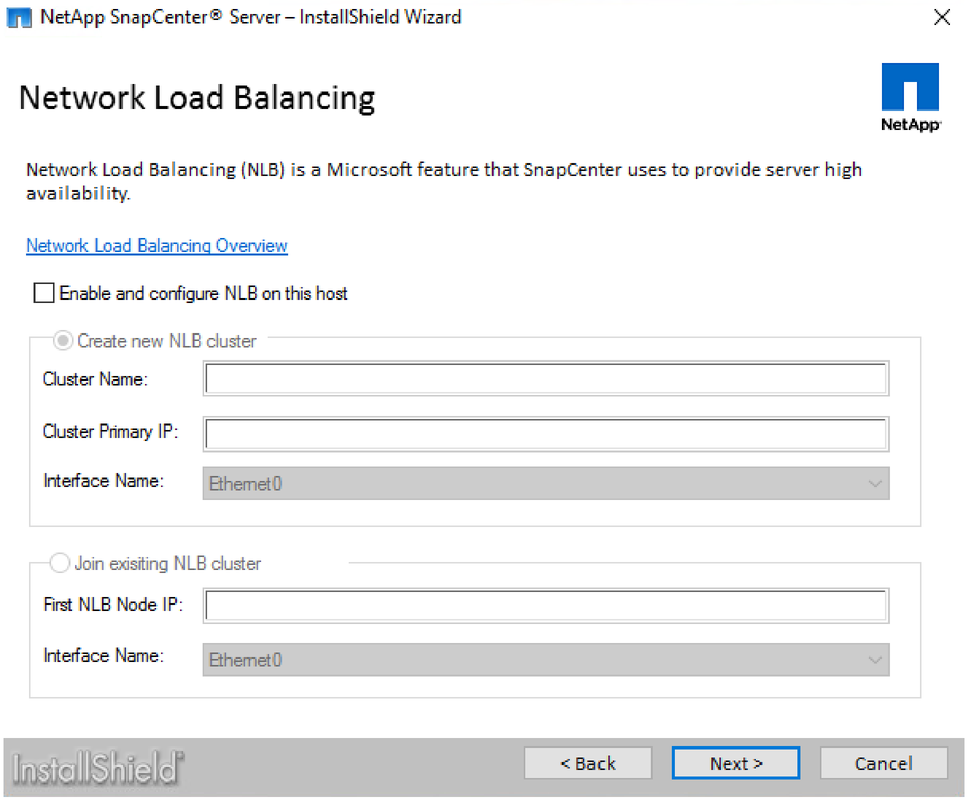

NetApp SnapCenter® Server – InstallShield Wizard

Network Load Balancing

Network Load Balancing (NLB) is a Microsoft feature that SnapCenter uses to provide server high availability.

Network Load Balancing Overview

□ Enable and configure NLB on this host

● Create new NLB cluster

Cluster Name:

Cluster Primary IP:

Interface Name: Ethernet0

○ Join exisiting NLB cluster

First NLB Node IP:

Interface Name: Ethernet0

5. Click Next on the Network Load Balancing screen.

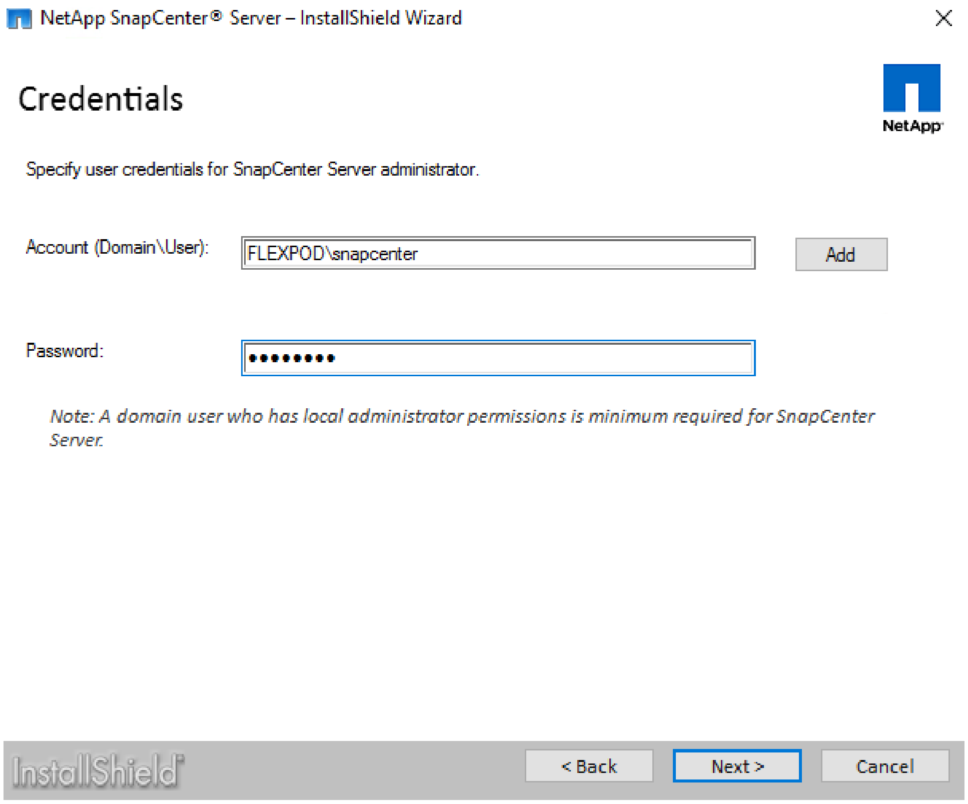
NetApp SnapCenter® Server – InstallShield Wizard

Credentials

Specify user credentials for SnapCenter Server administrator.

Account (Domain\User): FLEXPOD\snapcenter

Password: ••••••••

Note: A domain user who has local administrator permissions is minimum required for SnapCenter Server.

6. Enter the credentials for the service account on the Credentials screen, and click Next.

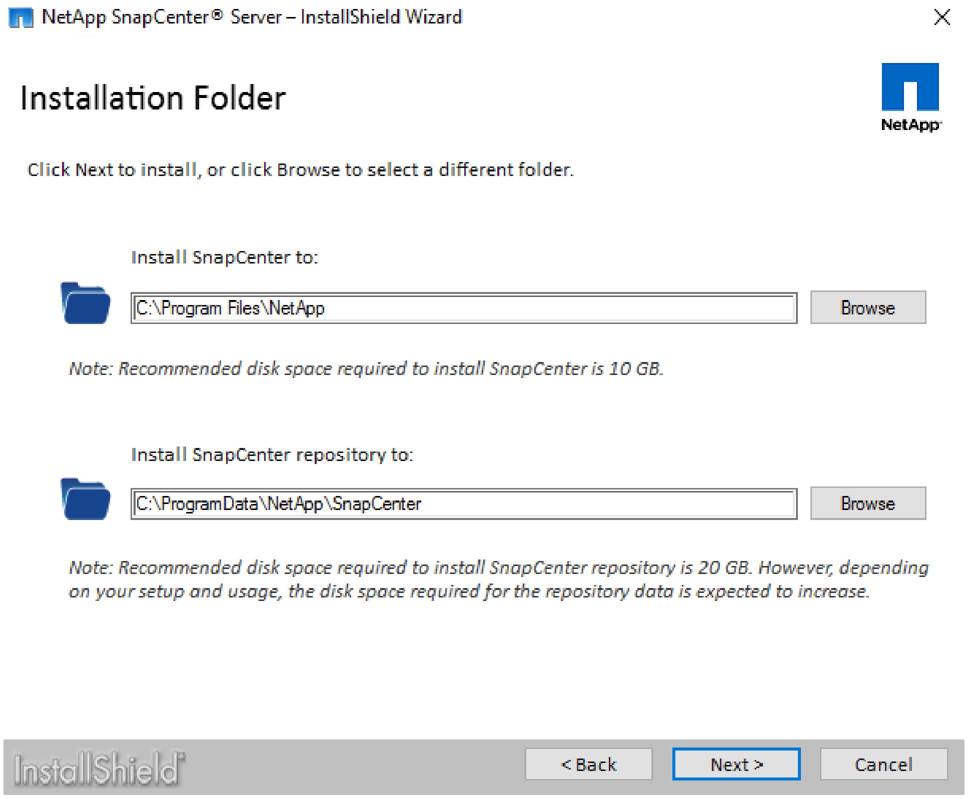
NetApp SnapCenter® Server – InstallShield Wizard

Installation Folder

Click Next to install, or click Browse to select a different folder.

Install SnapCenter to: C:\Program Files\NetApp

Note: Recommended disk space required to install SnapCenter is 10 GB.

Install SnapCenter repository to: C:\ProgramData\NetApp\SnapCenter

Note: Recommended disk space required to install SnapCenter repository is 20 GB. However, depending on your setup and usage, the disk space required for the repository data is expected to increase.

7. Click Next on the Installation folder screen.

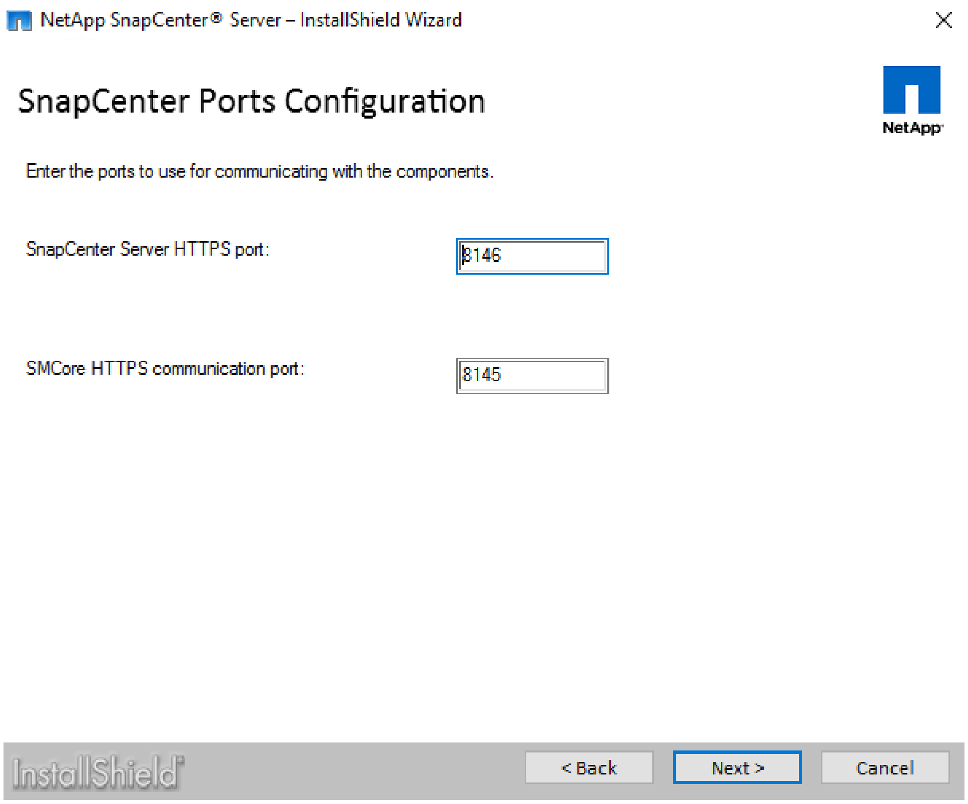

NetApp SnapCenter® Server – InstallShield Wizard

SnapCenter Ports Configuration

Enter the ports to use for communicating with the components.

SnapCenter Server HTTPS port: 8146

SMCore HTTPS communication port: 8145

8. Click Next on the SnapCenter Ports Configuration screen.

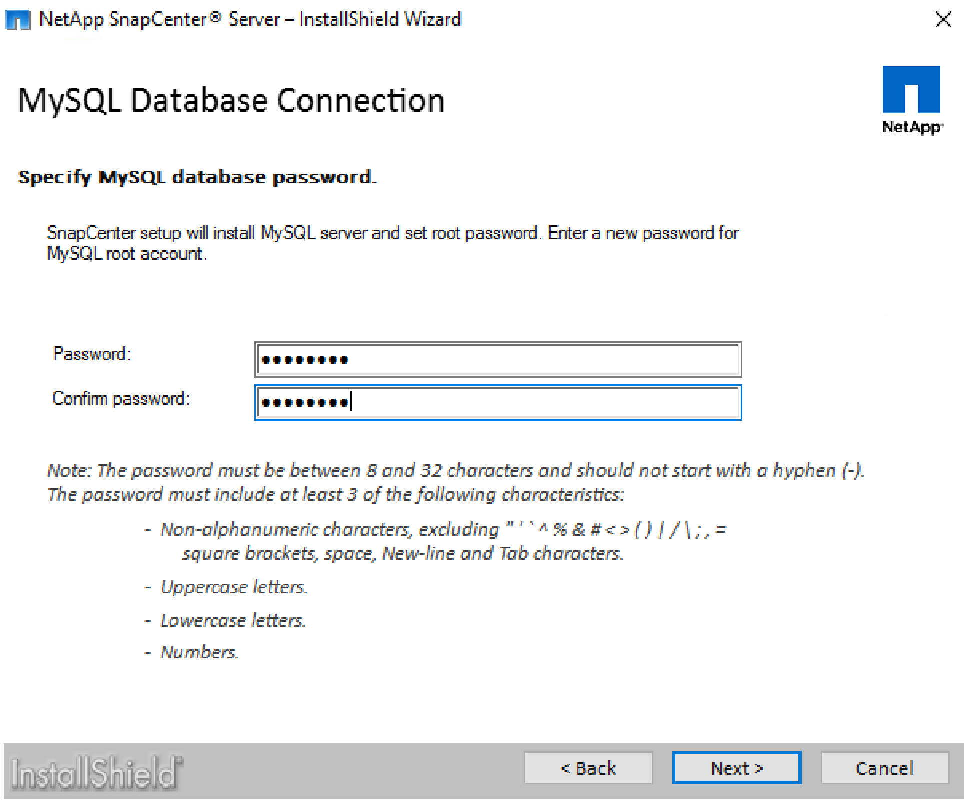

NetApp SnapCenter® Server – InstallShield Wizard

MySQL Database Connection

Specify MySQL database password.

SnapCenter setup will install MySQL server and set root password. Enter a new password for MySQL root account.

Password: ●●●●●●●●●

Confirm password: ●●●●●●●●●

Note: The password must be between 8 and 32 characters and should not start with a hyphen (-). The password must include at least 3 of the following characteristics:

- *Non-alphanumeric characters, excluding " ' ` ^ % & # < > ( ) | / \ ; , = square brackets, space, New-line and Tab characters.*
- *Uppercase letters.*
- *Lowercase letters.*
- *Numbers.*

9. Enter and confirm a password for the MySQL root account. Click Next.

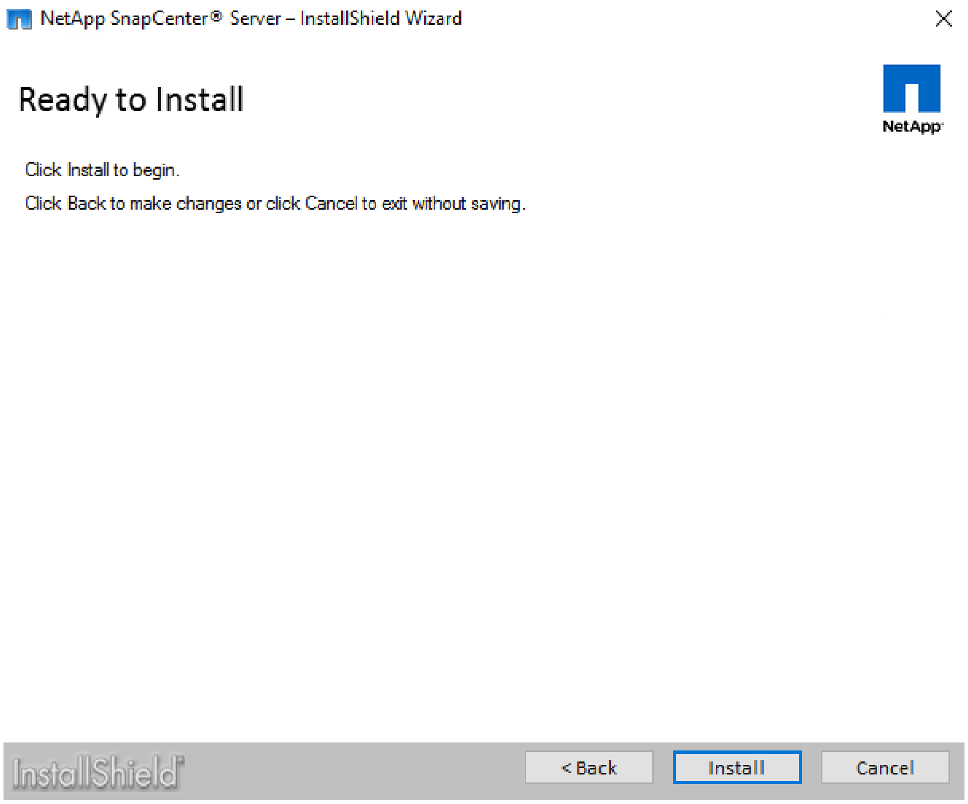

NetApp SnapCenter® Server – InstallShield Wizard

Ready to Install

Click Install to begin.

Click Back to make changes or click Cancel to exit without saving.

10. Click Next. The software will now start installing.

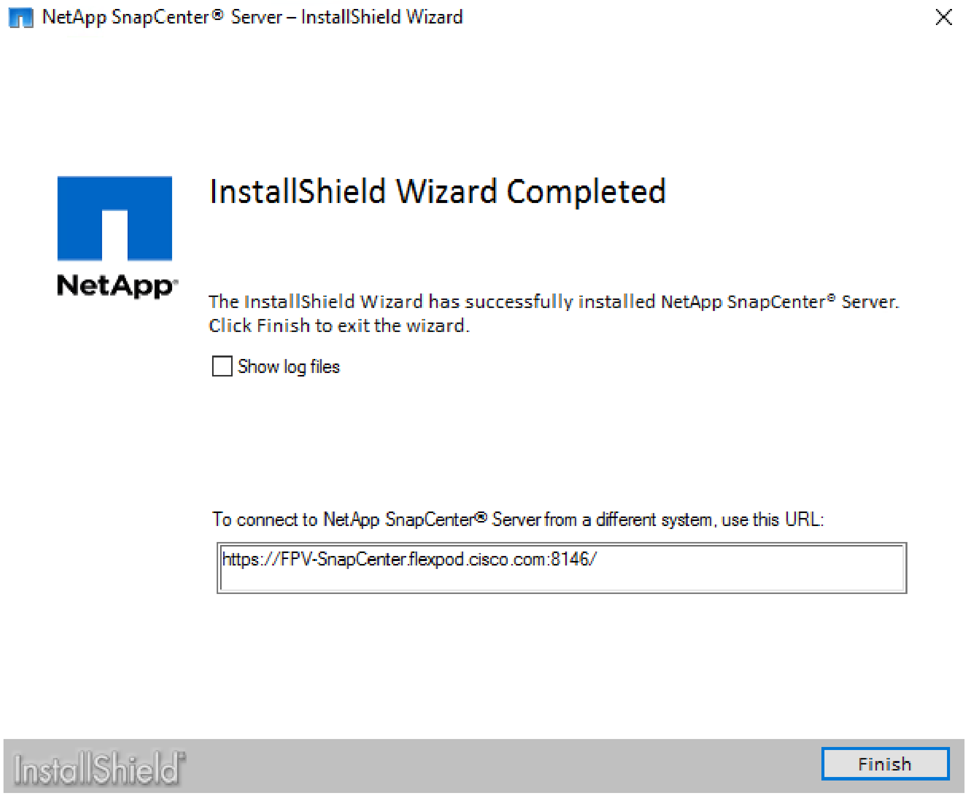

NetApp SnapCenter® Server – InstallShield Wizard

InstallShield Wizard Completed

The InstallShield Wizard has successfully installed NetApp SnapCenter® Server. Click Finish to exit the wizard.

□ Show log files

To connect to NetApp SnapCenter® Server from a different system, use this URL:

https://FPV-SnapCenter.flexpod.cisco.com:8146/

11. After the installation completes, connect to the SnapCenter server via a web browser, at the URL listed on the InstallShield Wizard Completed screen. Click Finish to close the installer.

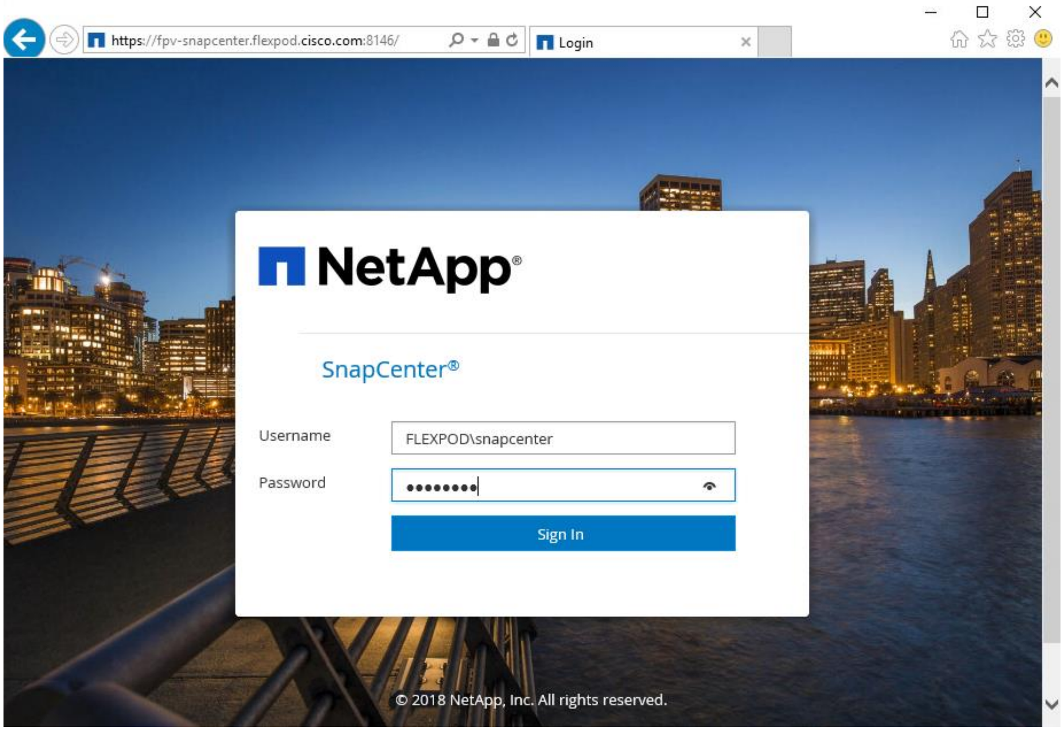

https://fpv-snapcenter.flexpod.cisco.com:8146/

Username: FLEXPOD\snapcenter

Password: ••••••••

© 2018 NetApp, Inc. All rights reserved.

12. Log in to the SnapCenter server using the service account credentials.

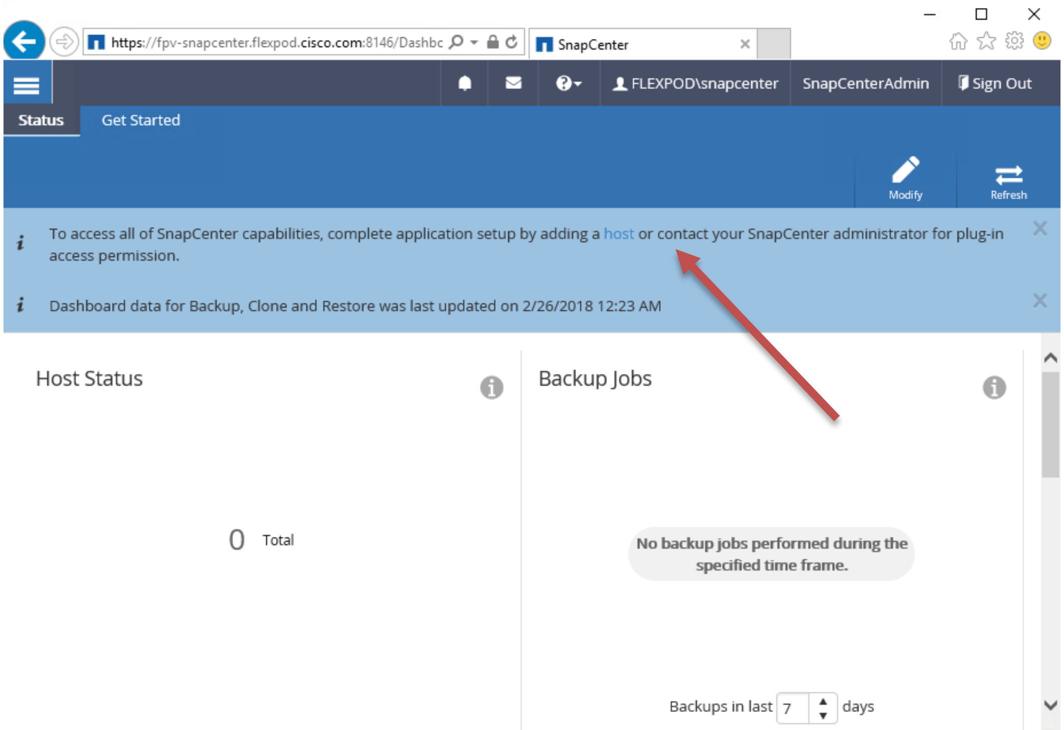

https://fpv-snapcenter.flexpod.cisco.com:8146/Dashbc

To access all of SnapCenter capabilities, complete application setup by adding a host or contact your SnapCenter administrator for plug-in access permission.

Dashboard data for Backup, Clone and Restore was last updated on 2/26/2018 12:23 AM

Host Status: 0 Total

Backup Jobs: No backup jobs performed during the specified time frame.

Backups in last 7 days

13. Click the “host” link (indicated by the arrow) to start adding your ESXi hosts.

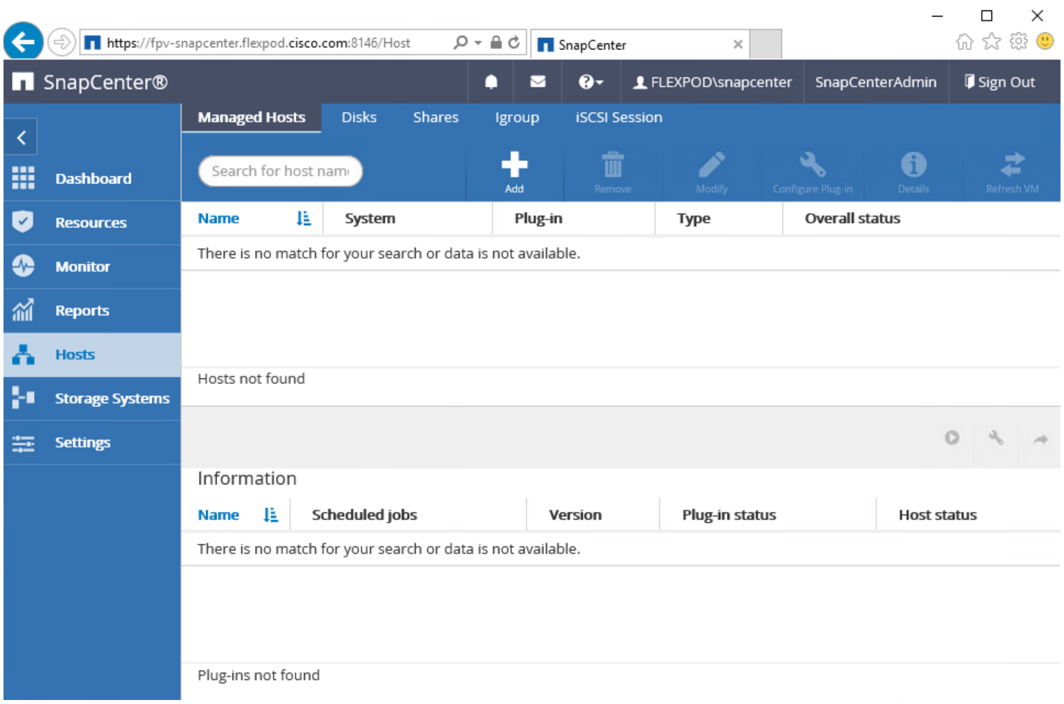

Name	System	Plug-in	Type	Overall status

There is no match for your search or data is not available.

Hosts not found

Information

Name	Scheduled jobs	Version	Plug-in status	Host status

There is no match for your search or data is not available.

Plug-ins not found

14. Click the Add button to add hosts.

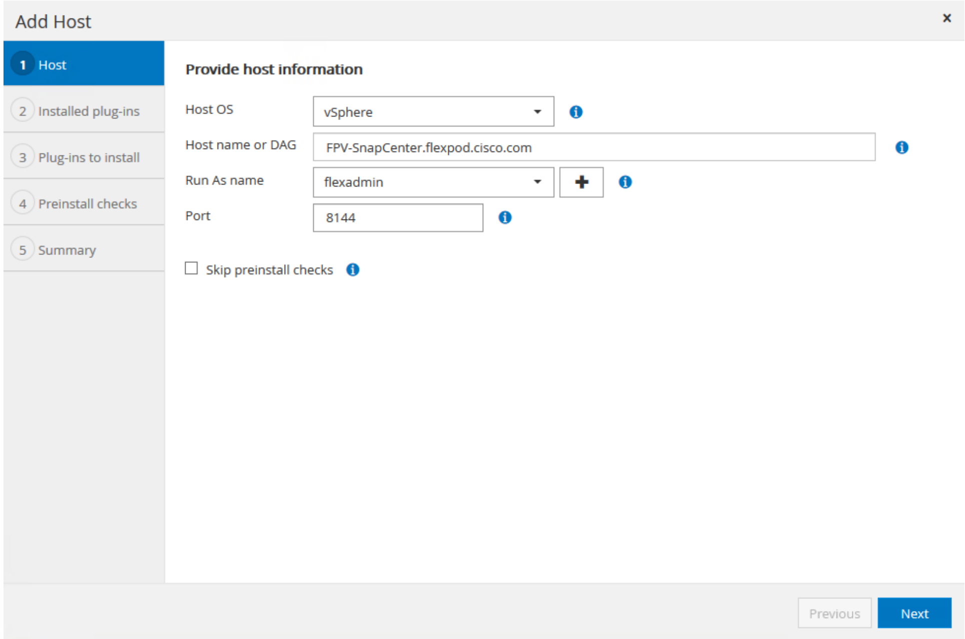

Provide host information

Host OS	vSphere
Host name or DAG	FPV-SnapCenter.flexpod.cisco.com
Run As name	flexadmin
Port	8144

□ Skip preinstall checks

15. In the first Add Host page (Host), select vSphere for Host OS; accept the default value (the hostname of the SnapCenter server) for the Host name field, and a domain admin for the Run As name field. If this is the first host for this instance of SnapCenter, you will need to add the domain admin credentials via the + button to the right of the Run As name drop-down list.

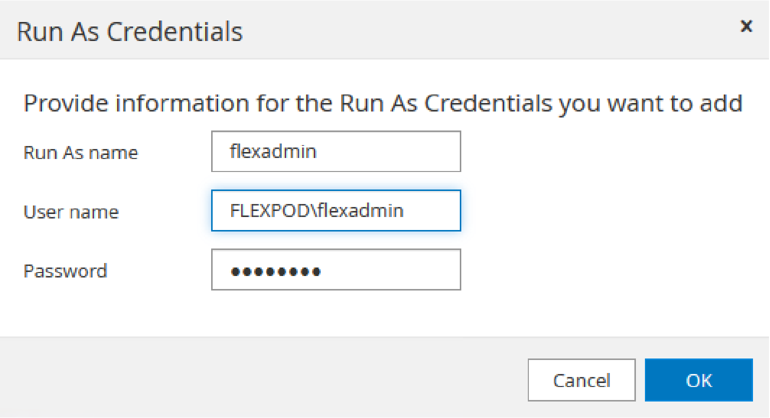

Run As Credentials

Provide information for the Run As Credentials you want to add

Run As name	flexadmin
User name	FLEXPOD\flexadmin
Password	••••••••

16. Back in the first Add Host page, click Next.

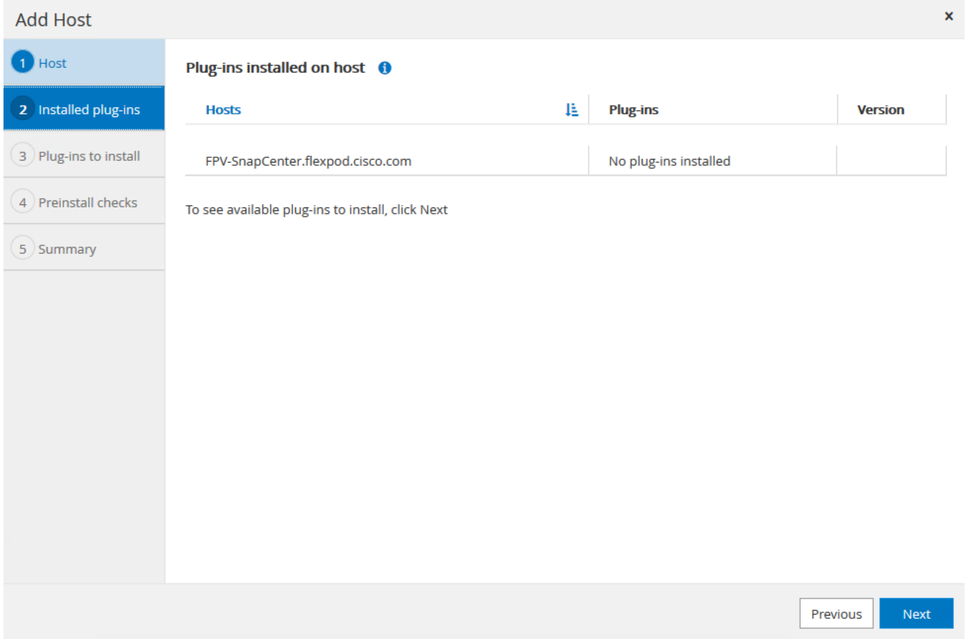

Add Host

Plug-ins installed on host

Hosts	Plug-ins	Version
FPV-SnapCenter.flexpod.cisco.com	No plug-ins installed	

To see available plug-ins to install, click Next

17. In the second Add Host page (Installed plug-ins), click Next.

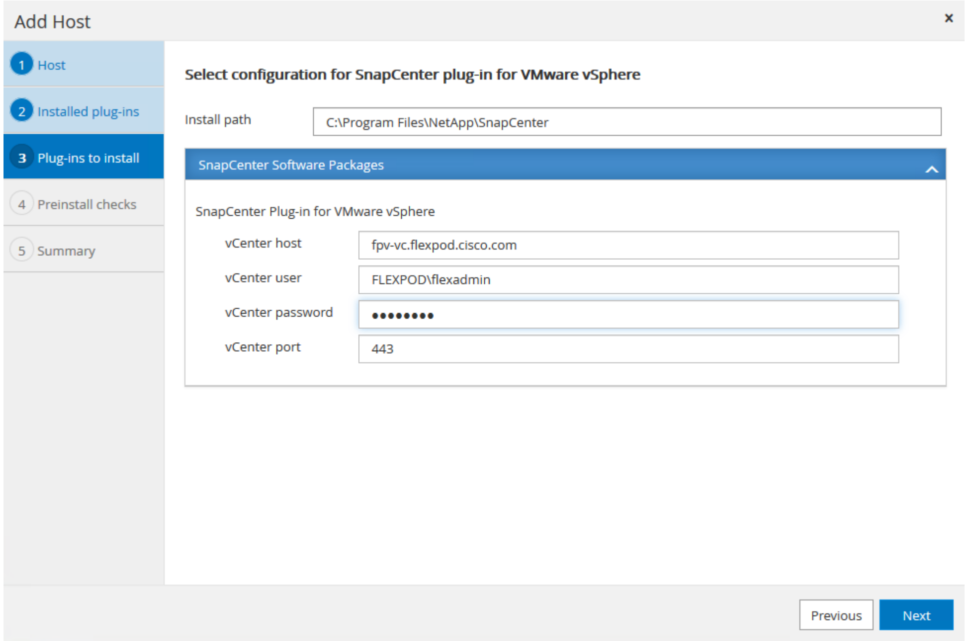

Select configuration for SnapCenter plug-in for VMware vSphere

Install path: C:\Program Files\NetApp\SnapCenter

SnapCenter Software Packages

SnapCenter Plug-in for VMware vSphere

vCenter host	fpv-vc.flexpod.cisco.com
vCenter user	FLEXPOD\flexadmin
vCenter password	••••••••
vCenter port	443

18. In the third Add Host page (Plug-ins to install), enter the vCenter credentials, and click Next. The plug-in installer will run the Preinstall checks.

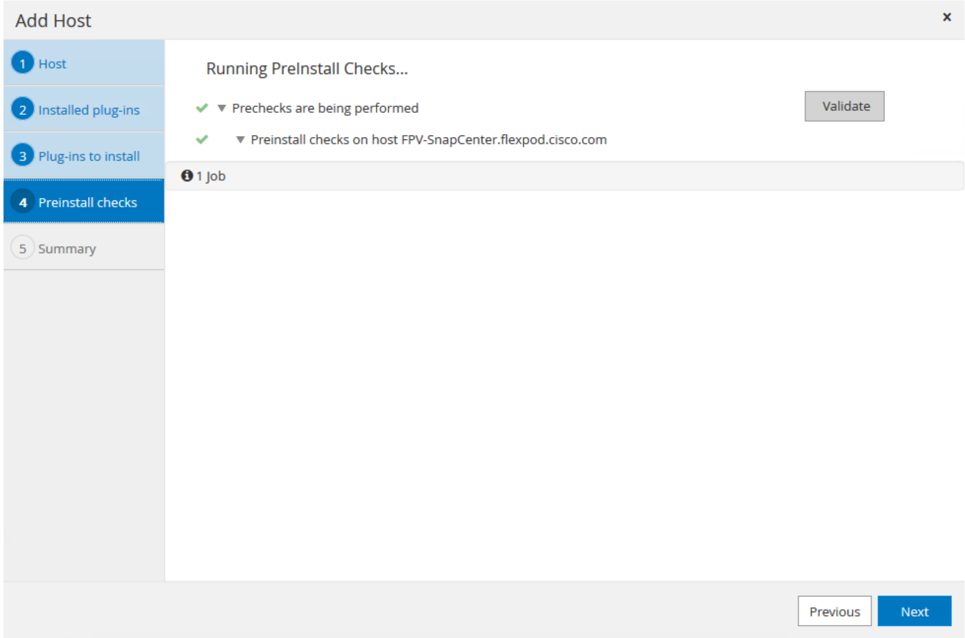

Running PreInstall Checks...

✓ ▼ Prechecks are being performed

✓ ▼ Preinstall checks on host FPV-SnapCenter.flexpod.cisco.com

1 Job

19. If any errors are displayed in the fourth Add Host page (Preinstall checks), correct the problem listed, and click Validate to run the Preinstall checks again. When the checks run without error, click Next.

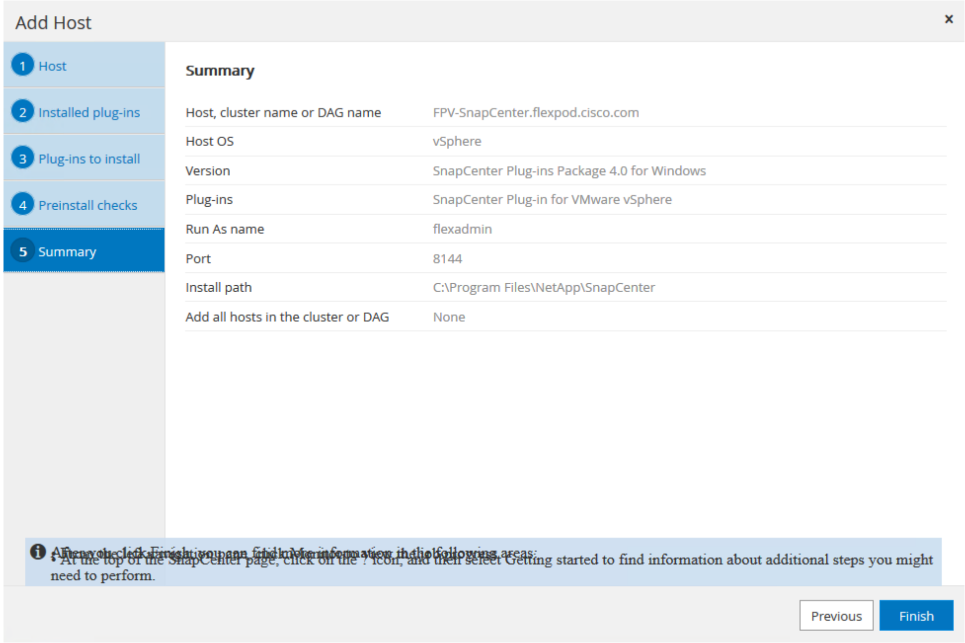

Summary

Host, cluster name or DAG name	FPV-SnapCenter.flexpod.cisco.com
Host OS	vSphere
Version	SnapCenter Plug-ins Package 4.0 for Windows
Plug-ins	SnapCenter Plug-in for VMware vSphere
Run As name	flexadmin
Port	8144
Install path	C:\Program Files\NetApp\SnapCenter
Add all hosts in the cluster or DAG	None

After you click Finish you can find more information in the following areas:

- At the top of the SnapCenter page, click on the ? icon, and then select Getting started to find information about additional steps you might need to perform.

20. Click Finish.

21. Log in to vCenter via the vSphere Web Client. The SnapCenter Plug-in for VMware vSphere should now be in the Inventories section of the Home page. (If you were logged in to vCenter during the SnapCenter plug-in installation, log out of vCenter and then log back in to load the newly installed plug-in.) Click the SnapCenter plug-in icon.

22. In the left Navigator pane of the VMware vSphere web client in vCenter for the Plug-in for VMware vSphere, click Storage Systems.

23. On the Storage Systems page, click Add Storage System.

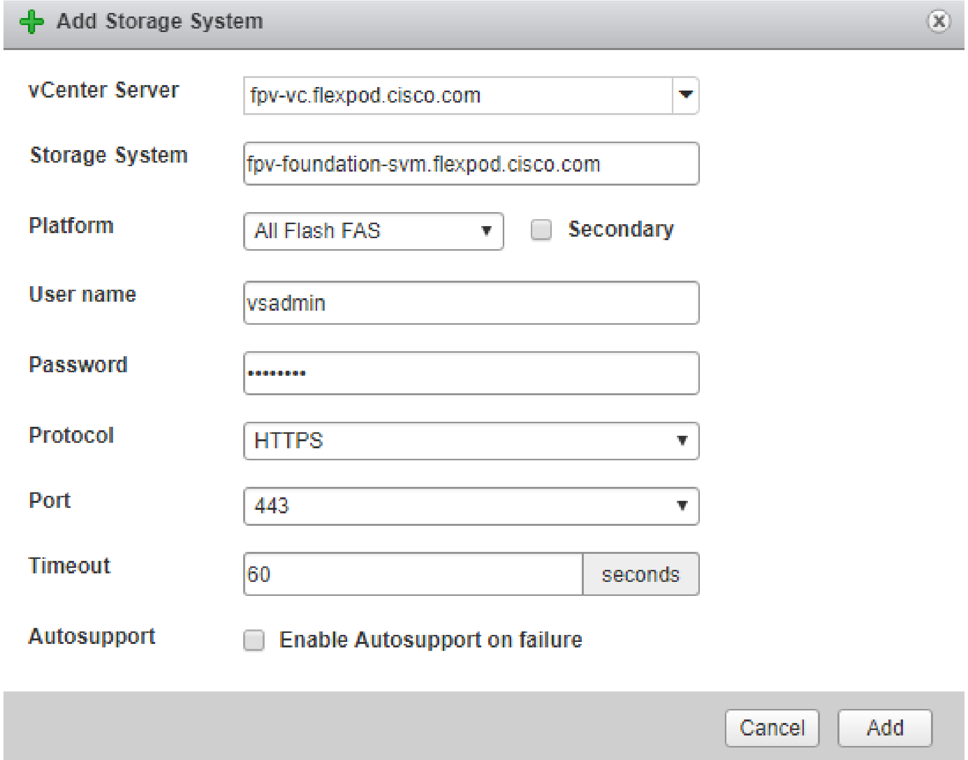

Add Storage System

vCenter Server	fpv-vc.flexpod.cisco.com
Storage System	fpv-foundation-svm.flexpod.cisco.com
Platform	All Flash FAS □ Secondary
User name	vsadmin
Password	••••••••
Protocol	HTTPS
Port	443
Timeout	60 seconds
Autosupport	□ Enable Autosupport on failure

24. In the Add Storage System dialog box, enter the SVM information. Click Add.

25. Repeat the previous two steps as needed to add other SVMs.

26. SnapCenter backups have two components: Policies (which define things like frequency, retention, etc.) and Resource Groups (define which VMs and datastores are backed up). To create a policy, in the left Navigator pane of the VMware vSphere web client in vCenter for the Plug-in for VMware vSphere, click Policies.

27. In the Policies page, click New Policy in the toolbar.

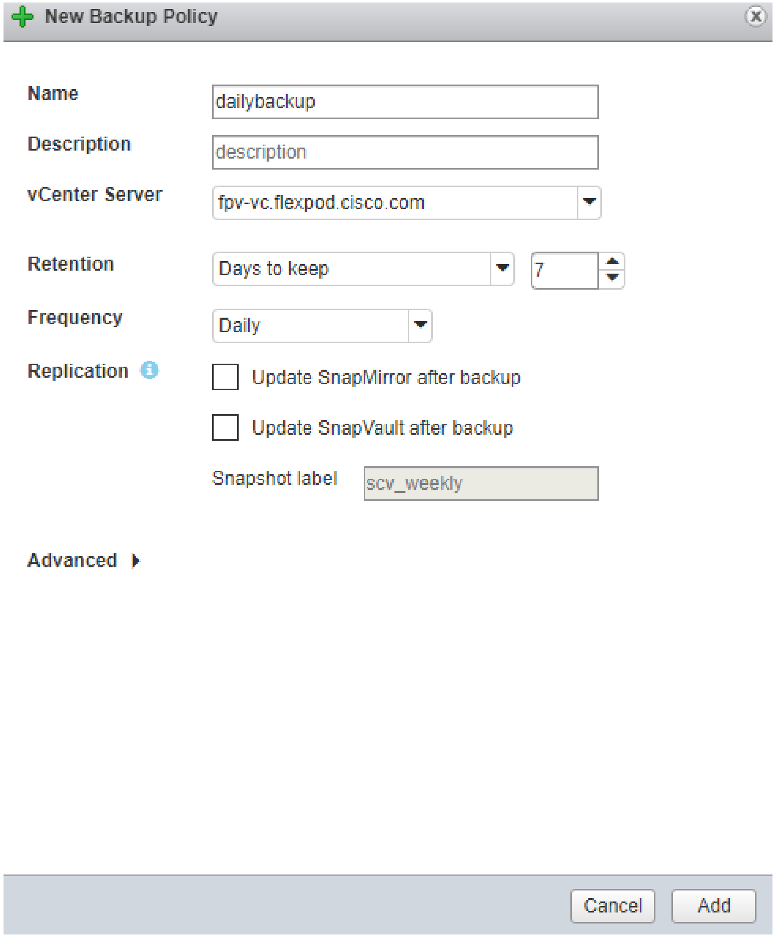

New Backup Policy

Name: dailybackup

Description: description

vCenter Server: fpv-vc.flexpod.cisco.com

Retention: Days to keep 7

Frequency: Daily

Replication

□ Update SnapMirror after backup

□ Update SnapVault after backup

Snapshot label: scv_weekly

28. In the New Backup Policy page, enter the desired parameters. Click Add.

29. A resource group can contain VMs, and SAN and NAS datastores; it cannot contain VSAN or VVOL datastores. To create a resource group, in the left Navigator pane of the VMware vSphere web client in vCenter for the Plug-in for VMware vSphere, click Resource Groups and then click (Create Resource Group).

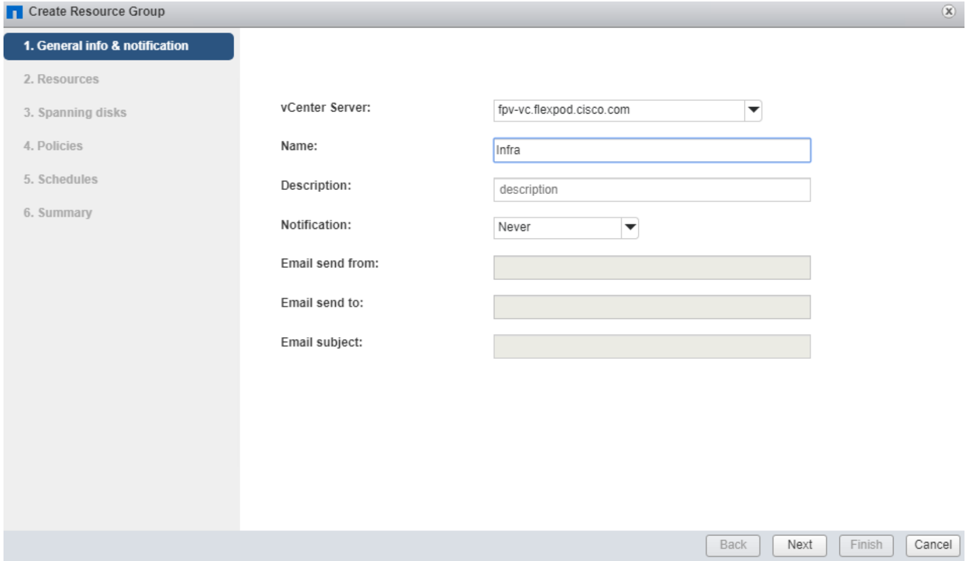

Create Resource Group

vCenter Server: fpv-vc.flexpod.cisco.com

Name: Infra

Description: description

Notification: Never

Email send from:

Email send to:

Email subject:

30. Enter the desired info and click Next.

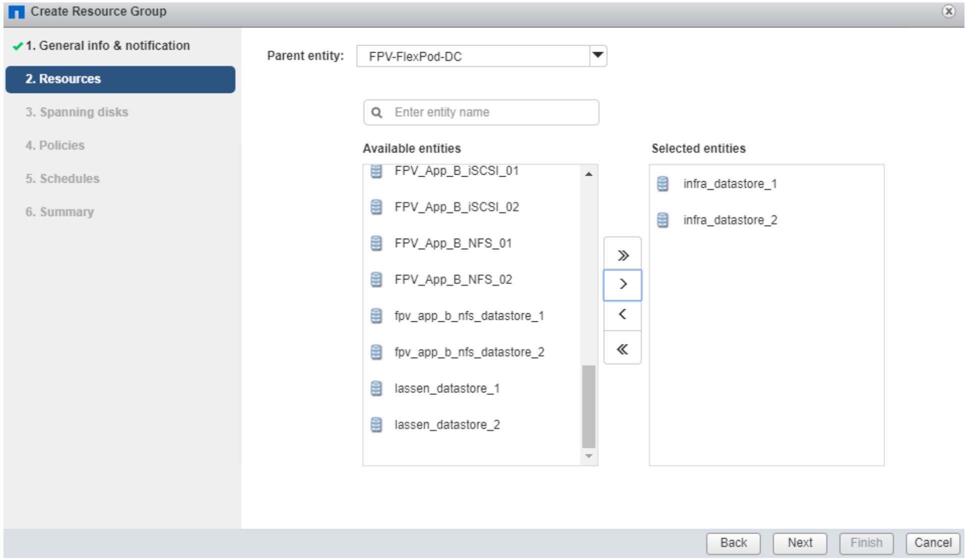

Parent entity: FPV-FlexPod-DC

Enter entity name

Available entities

- FPV_App_B_iSCSI_01
- FPV_App_B_iSCSI_02
- FPV_App_B_NFS_01
- FPV_App_B_NFS_02
- fpv_app_b_nfs_datastore_1
- fpv_app_b_nfs_datastore_2
- lassen_datastore_1
- lassen_datastore_2

Selected entities

- infra_datastore_1
- infra_datastore_2

31. Select the desired resources to back up and click Next.

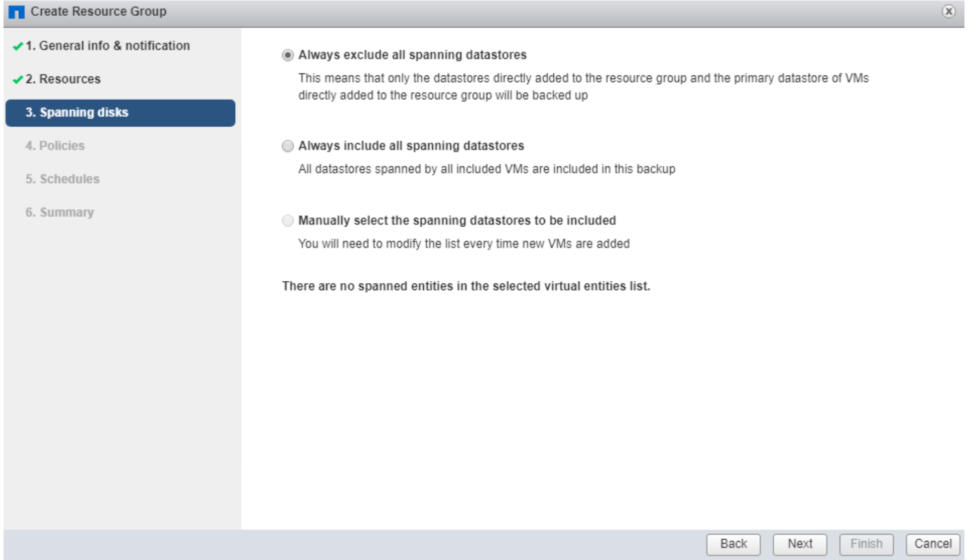

● Always exclude all spanning datastores

This means that only the datastores directly added to the resource group and the primary datastore of VMs directly added to the resource group will be backed up

○ Always include all spanning datastores

All datastores spanned by all included VMs are included in this backup

○ Manually select the spanning datastores to be included

You will need to modify the list every time new VMs are added

There are no spanned entities in the selected virtual entities list.

32. Select the desired treatment for spanning disks (VMDKs that span multiple datastores) and click Next.

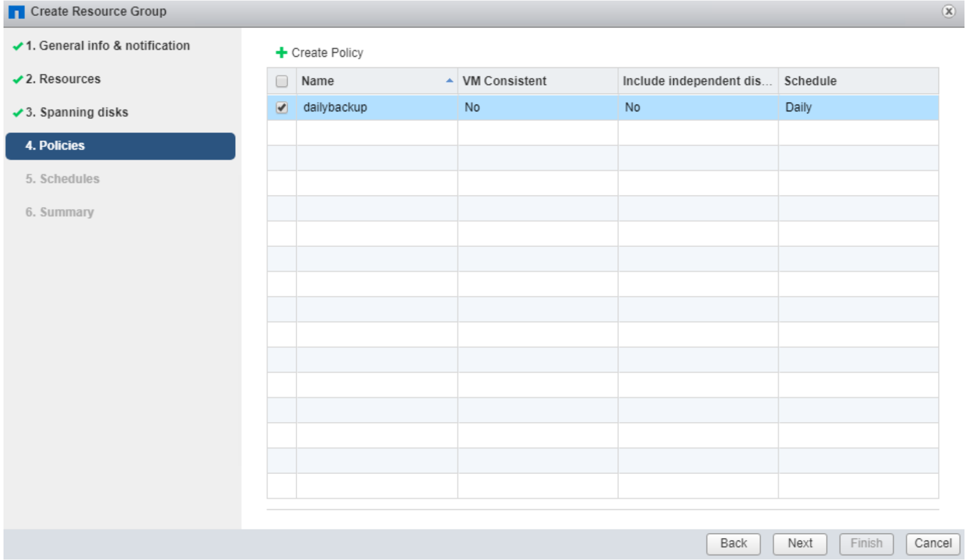

□	Name	VM Consistent	Include independent dis...	Schedule
☑	dailybackup	No	No	Daily

33. Select the backup policy and click Next.

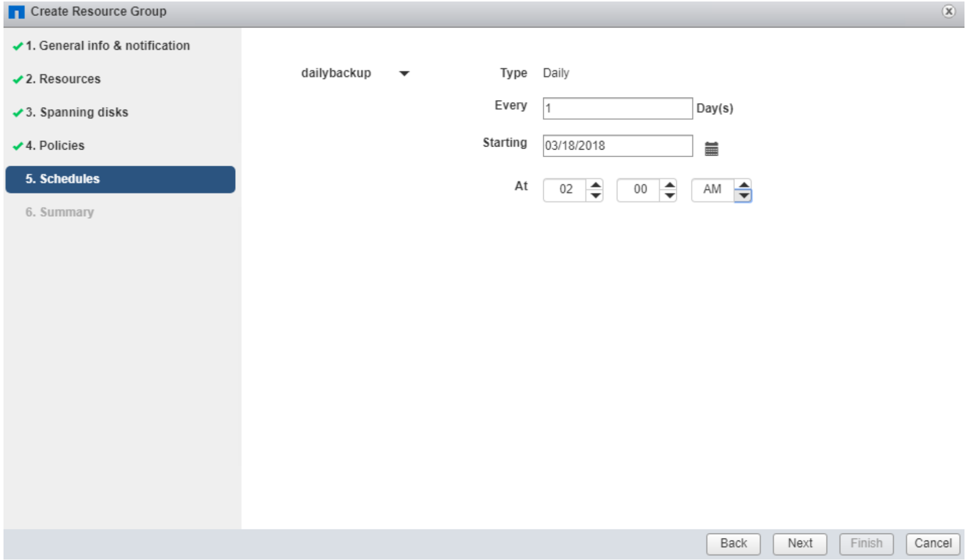

dailybackup

Type Daily

Every 1 Day(s)

Starting 03/18/2018

At 02 00 AM

34. Edit schedule parameters as desired and click Next.

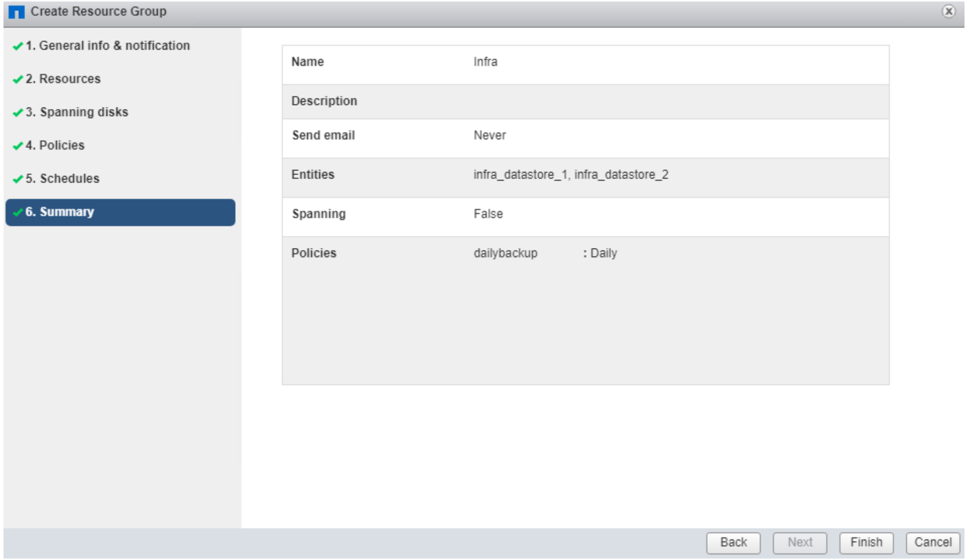

Name	Infra
Description	
Send email	Never
Entities	infra_datastore_1, infra_datastore_2
Spanning	False
Policies	dailybackup : Daily

35. Review the parameters and click Finish.

OnCommand Unified Manager 7.2

To use the OVA deployment method to deploy OnCommand Unified Manager as a virtual machine in vSphere, complete the following steps:

1. Download and review the OnCommand Unified Manager 7.2 Installation and Setup Guide

2. Download the OnCommand Unified Manager version 7.2P1 OVA file for VMware vSphere from http://mysupport.netapp.com to the local management machine.

3. In vSphere Web Client, select the FPV-Foundation cluster in the Navigator, then select Deploy OVF Template from the Actions menu.

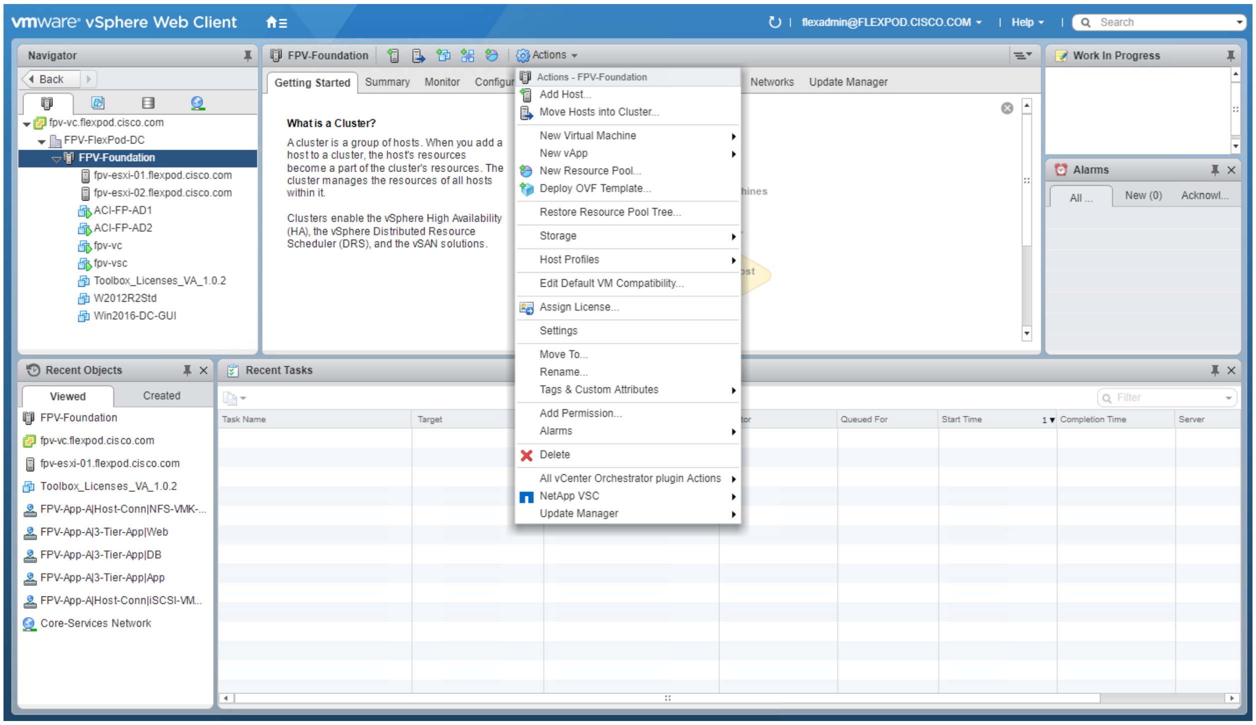

What is a Cluster?

A cluster is a group of hosts. When you add a host to a cluster, the host's resources become a part of the cluster's resources. The cluster manages the resources of all hosts within it.

Clusters enable the vSphere High Availability (HA), the vSphere Distributed Resource Scheduler (DRS), and the vSAN solutions.

4. Select the Local file radio button, then Browse to navigate to the OVA file. Select the file, click Open and the click Next.

5. On the name and location page, enter the desired VM name, select the FPV-FlexPod-DC datacenter as the installation location, and click Next.

6. On the resource page, select the FPV-Foundation cluster as the resource to run the VM on.

7. On the Review details page, ignore the warning about advanced configuration options, and click Next to continue.

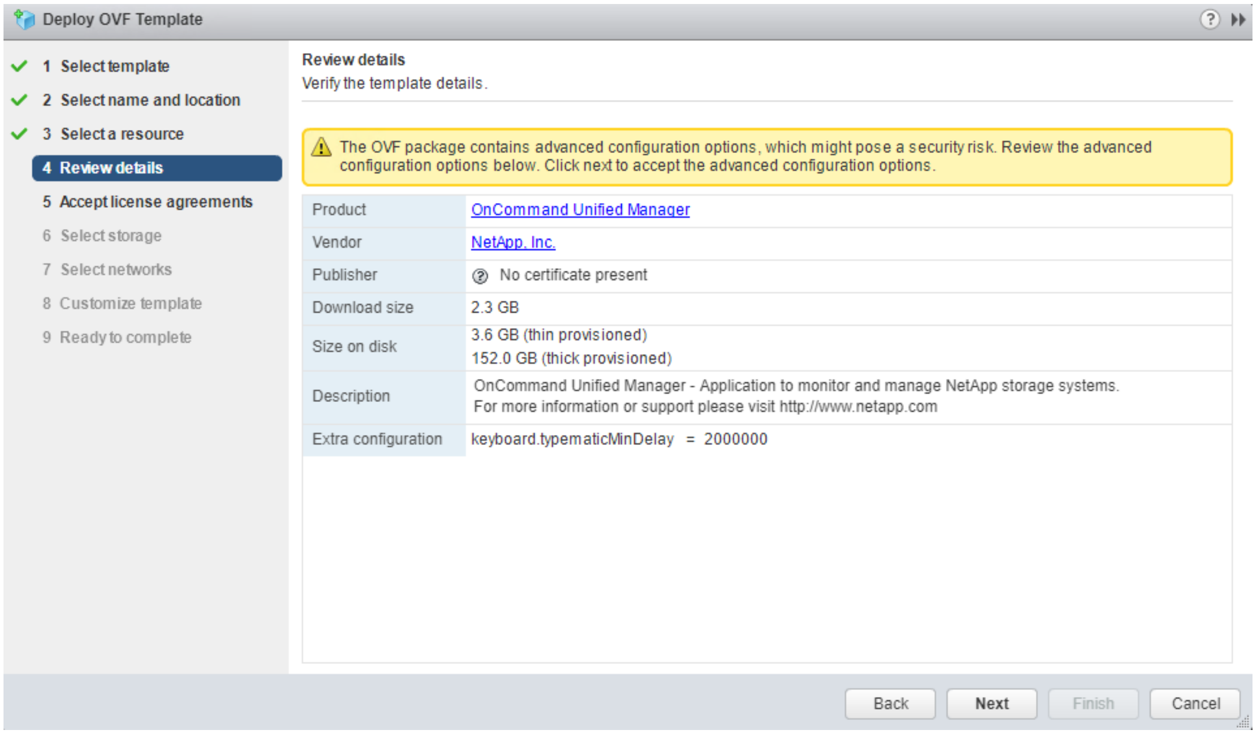

Review details

Verify the template details.

The OVF package contains advanced configuration options, which might pose a security risk. Review the advanced configuration options below. Click next to accept the advanced configuration options.

Product	OnCommand Unified Manager
Vendor	NetApp, Inc.
Publisher	No certificate present
Download size	2.3 GB
Size on disk	3.6 GB (thin provisioned) 152.0 GB (thick provisioned)
Description	OnCommand Unified Manager - Application to monitor and manage NetApp storage systems. For more information or support please visit http://www.netapp.com
Extra configuration	keyboard.typematicMinDelay = 2000000

8. On the Accept license agreement page, note the post-deployment steps. Click Accept, then click Next.

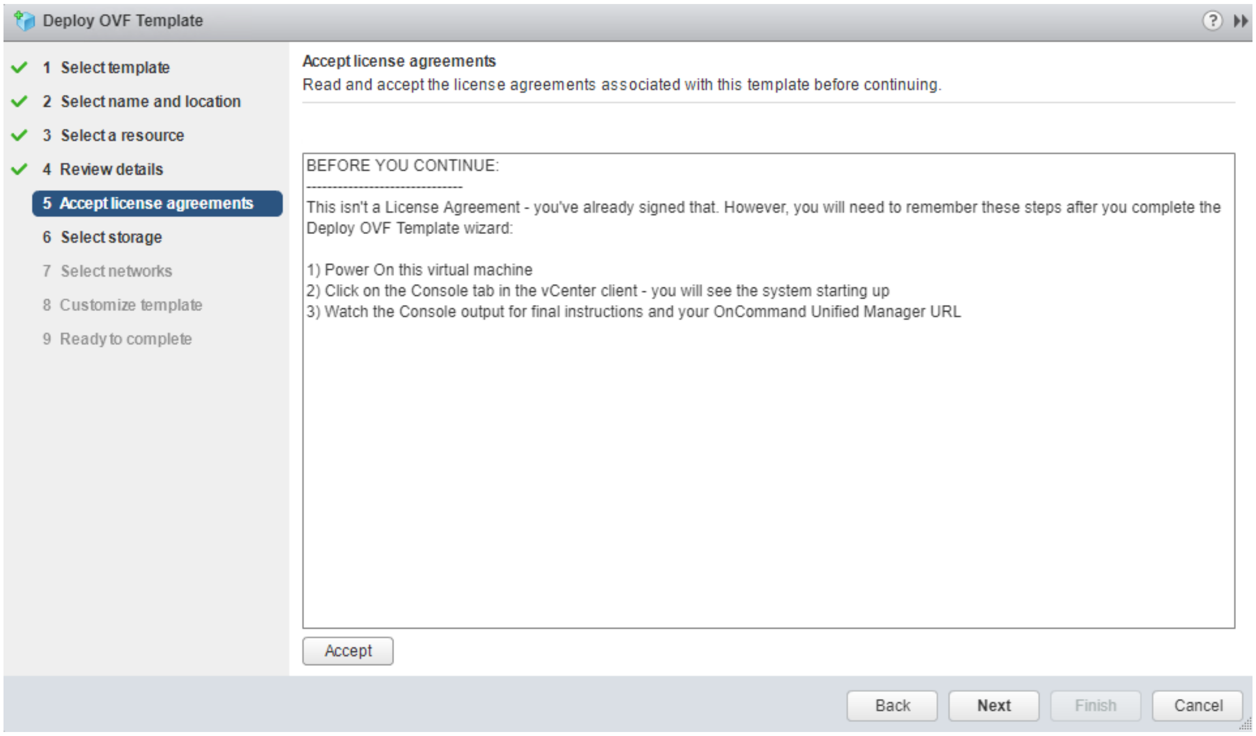

Accept license agreements

Read and accept the license agreements associated with this template before continuing.

```
BEFORE YOU CONTINUE:
------------------------------
This isn't a License Agreement - you've already signed that. However, you will need to remember these steps after you complete the
Deploy OVF Template wizard:

1) Power On this virtual machine
2) Click on the Console tab in the vCenter client - you will see the system starting up
3) Watch the Console output for final instructions and your OnCommand Unified Manager URL
```

9. On the storage page, select infra_datastore_1, and click Next to continue.

10. On the Select networks page, select common | Core-Services |Core-Services as the Destination Network, and click Next.

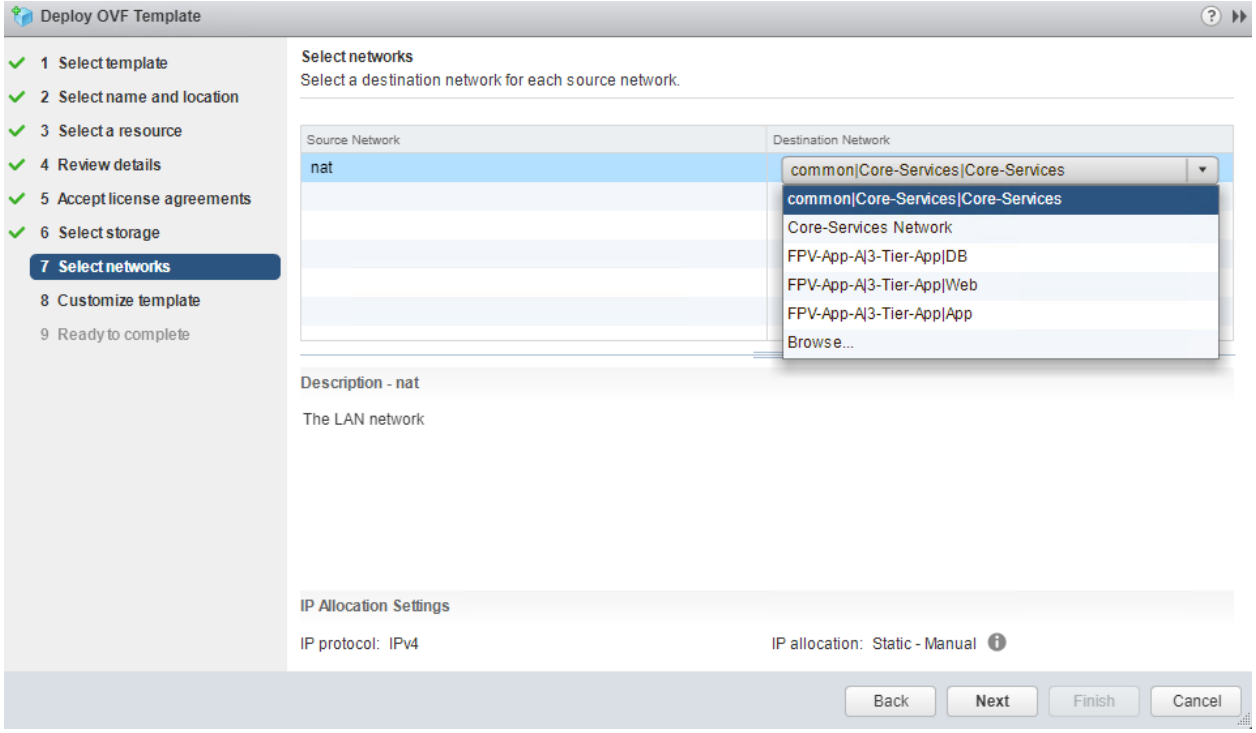

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	common|Core-Services|Core-Services

common|Core-Services|Core-Services
Core-Services Network
FPV-App-A|3-Tier-App|DB
FPV-App-A|3-Tier-App|Web
FPV-App-A|3-Tier-App|App
Browse...

Description - nat

The LAN network

IP Allocation Settings

IP protocol: IPv4

IP allocation: Static - Manual

11. On the Customize template page, add deployment-specific parameters, and click Next.

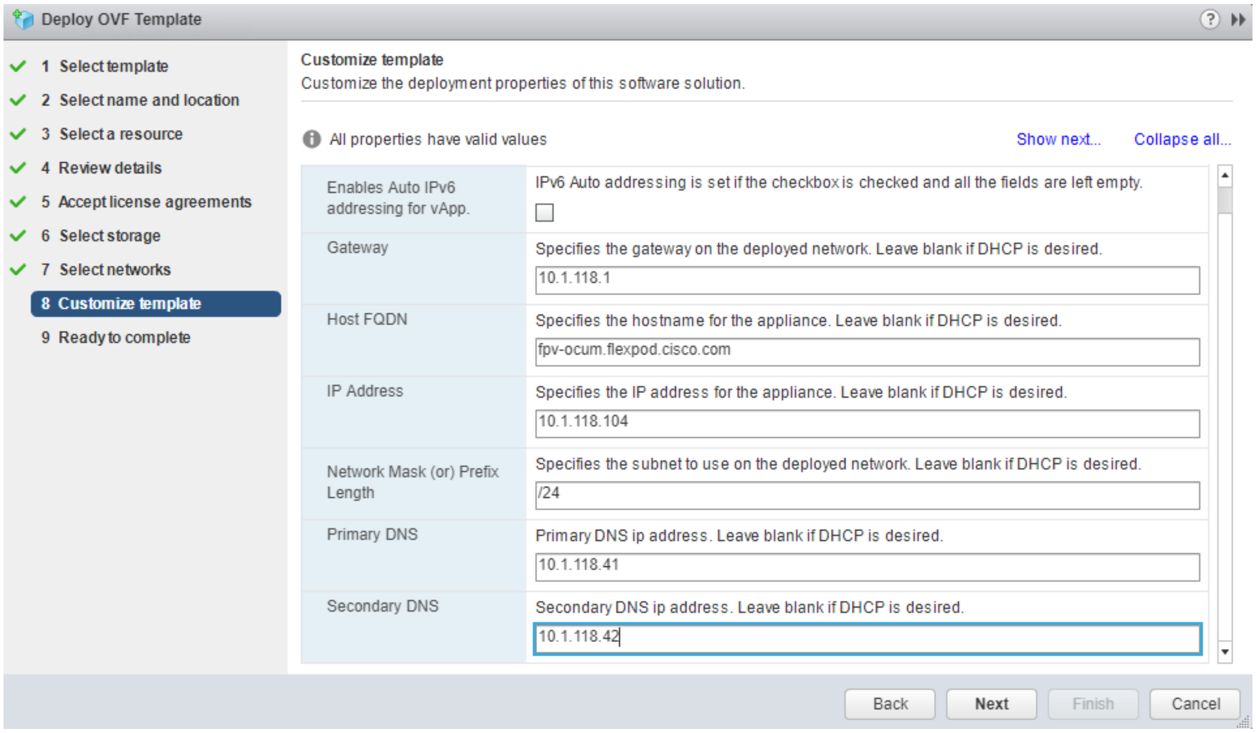

Customize template

Customize the deployment properties of this software solution.

All properties have valid values

Enables Auto IPv6 addressing for vApp.	IPv6 Auto addressing is set if the checkbox is checked and all the fields are left empty. ☐
Gateway	Specifies the gateway on the deployed network. Leave blank if DHCP is desired. 10.1.118.1
Host FQDN	Specifies the hostname for the appliance. Leave blank if DHCP is desired. fpv-ocum.flexpod.cisco.com
IP Address	Specifies the IP address for the appliance. Leave blank if DHCP is desired. 10.1.118.104
Network Mask (or) Prefix Length	Specifies the subnet to use on the deployed network. Leave blank if DHCP is desired. /24
Primary DNS	Primary DNS ip address. Leave blank if DHCP is desired. 10.1.118.41
Secondary DNS	Secondary DNS ip address. Leave blank if DHCP is desired. 10.1.118.42

12. Review the deployment parameters, and click Finish to deploy the VM.

13. After the VM deployment completes, edit the VM’s settings in vSphere Web Client to mount the VMware tools ISO. Open the VM console in vSphere Web Client, and power on the VM; the installation of VMware tools will execute automatically.

14. Set the time parameters as prompted.

15. Create the maintenance user as prompted.

16. Using a web browser, connect to the URL listed on the console. For the first login, use the maintenance user credentials.

17. Enter e-mail and NTP settings and click Next.

18. Enable AutoSupport and click Next.

19. Click Save and Continue.

20. Click Add to add a new storage cluster.

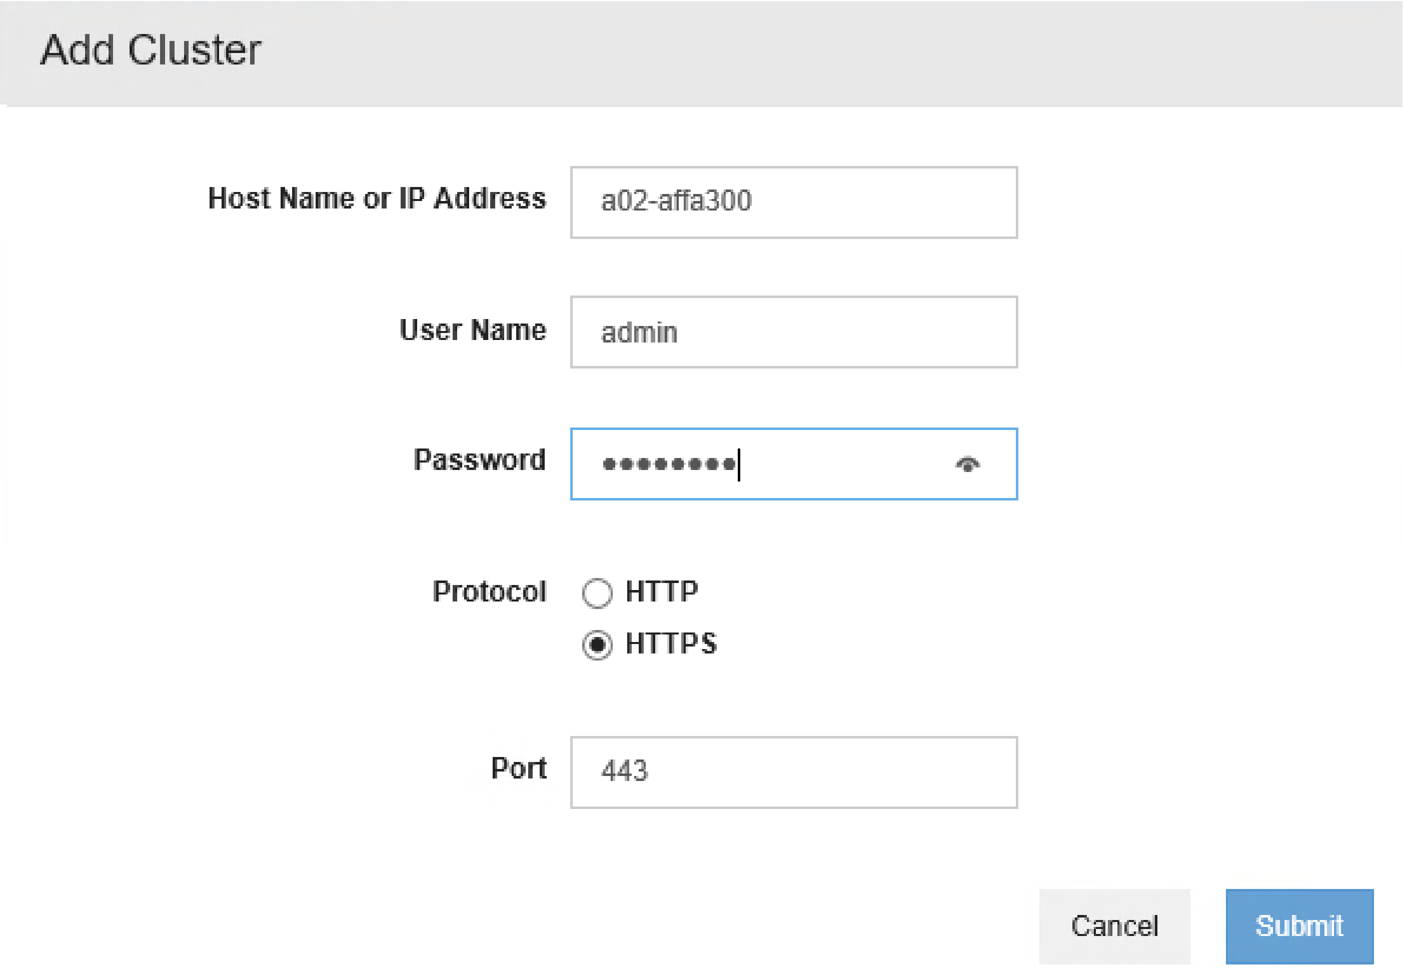

Add Cluster

Host Name or IP Address: a02-affa300

User Name: admin

Password: ●●●●●●●●●

Protocol: ○ HTTP ● HTTPS

Port: 443

21. Storage information can then be viewed by logging into OnCommand Unified Manager at https://OncommandUnifiedManagerHostname as shown below:

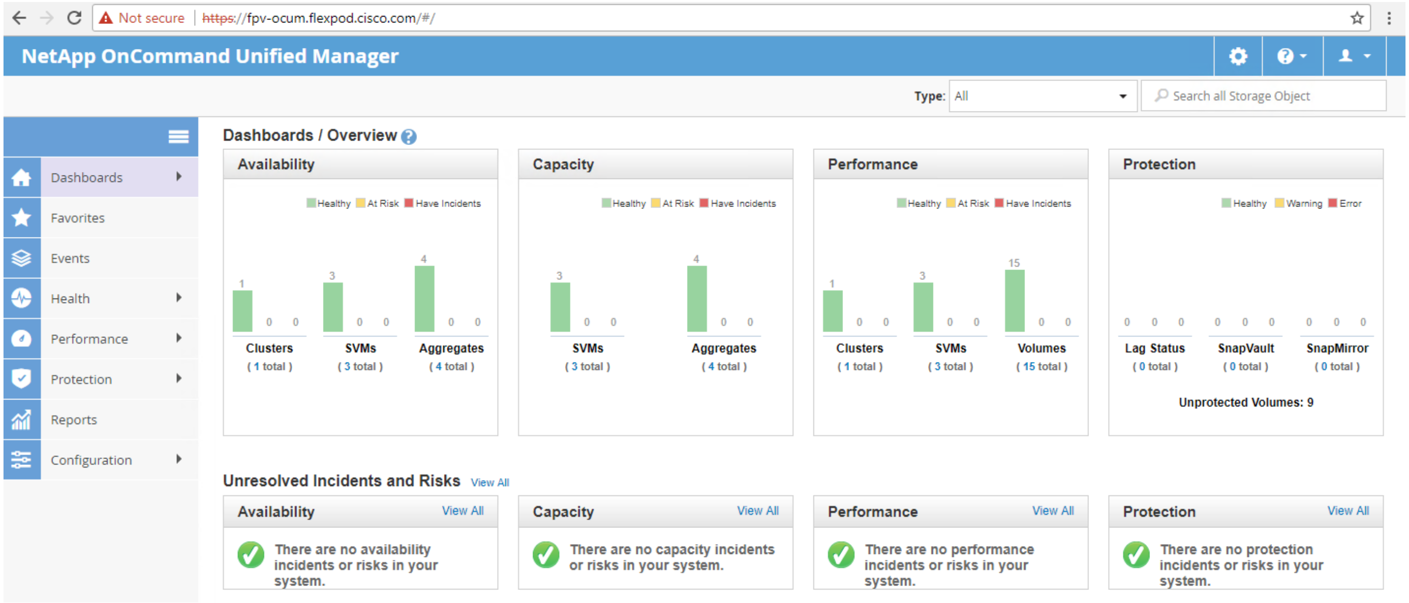

There are no availability incidents or risks in your system.

There are no capacity incidents or risks in your system.

There are no performance incidents or risks in your system.

There are no protection incidents or risks in your system.

Sample Tenant Setup

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 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 Nexus 7000 core routers using OSPF.
- The Nexus 7000s' default gateway is shared with the ACI fabric using OSPF.
- Each of the two Nexus 7000s is connected to each of the two Nexus 9000 leaf switches.
- Sub-interfaces are configured and used for external connectivity.
- The Nexus 7000s are configured to originate and send a default route to the Nexus 9000 leaf switches.
- This Shared Layer 3 Out was set up on a set of 10GE leaves that were part of the ACI fabric and not the 9332 leaves (which were also part of the fabric) used in this validation.

Figure 4 ACI Shared Layer 3 Out Connectivity Details

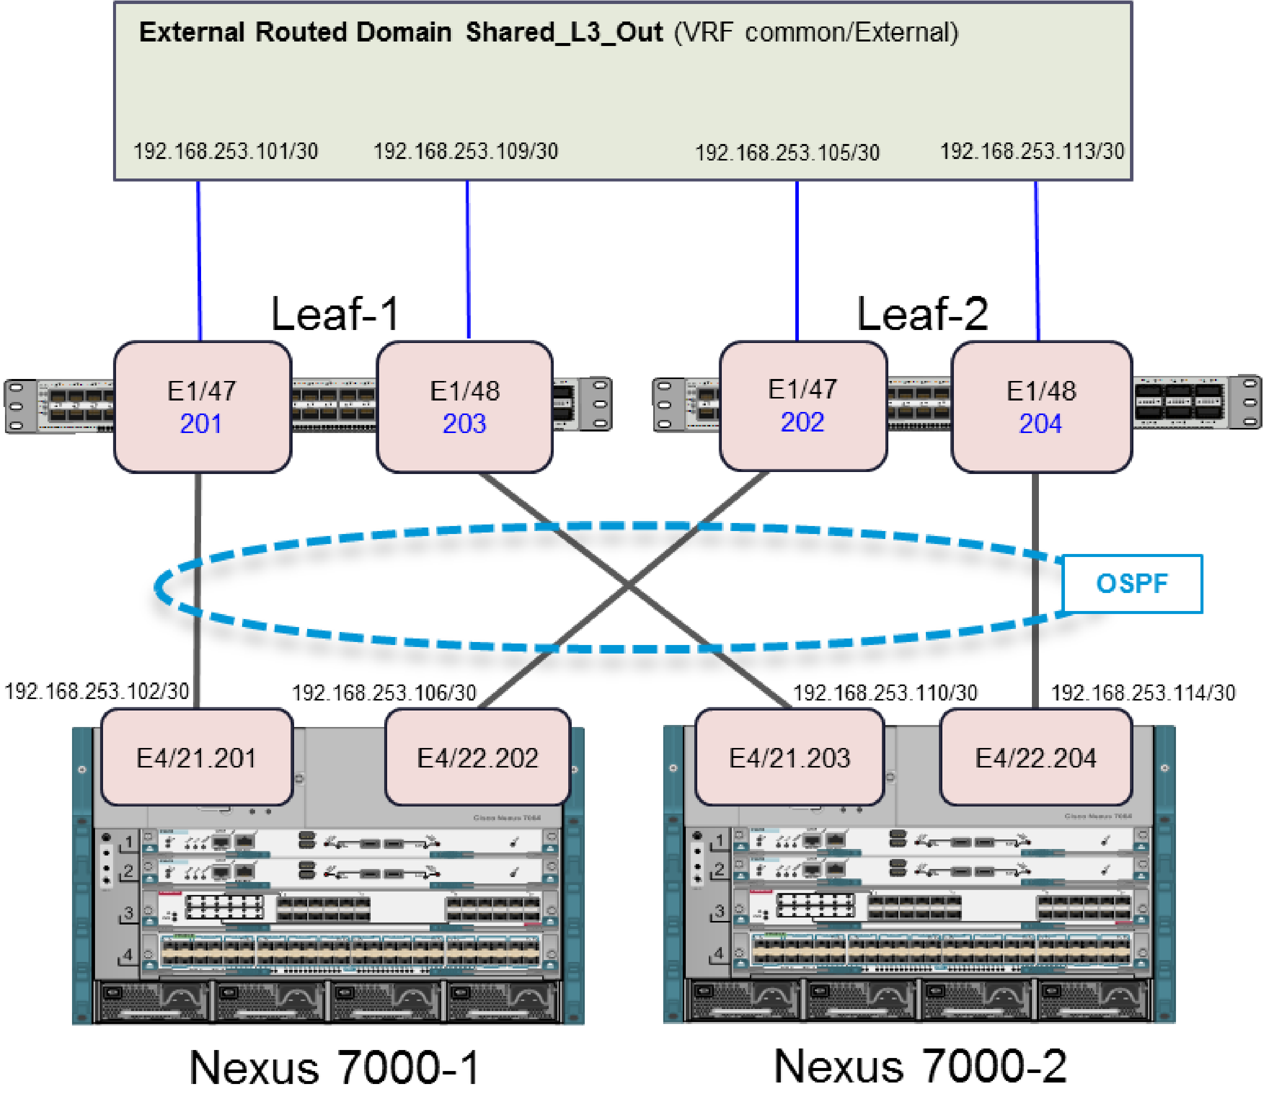

Configuring the Nexus 7000s for ACI Connectivity (Sample)

The following configuration is a sample from the virtual device contexts (VDCs) from two Nexus 7004s. Interfaces and a default route from the two Nexus 7000s also needs to be set up, but is not shown here because this is set up according to the customer's security policy.

Nexus 7004-1 VDC

```
feature ospf


vlan 100

  name OSPF-Peering


interface Vlan100
```

```
  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
    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
```

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 GUI

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 list to select Create Attachable Entity Profile.
6. Name the Profile `AEP-Shared-L3-Out` and click Next.

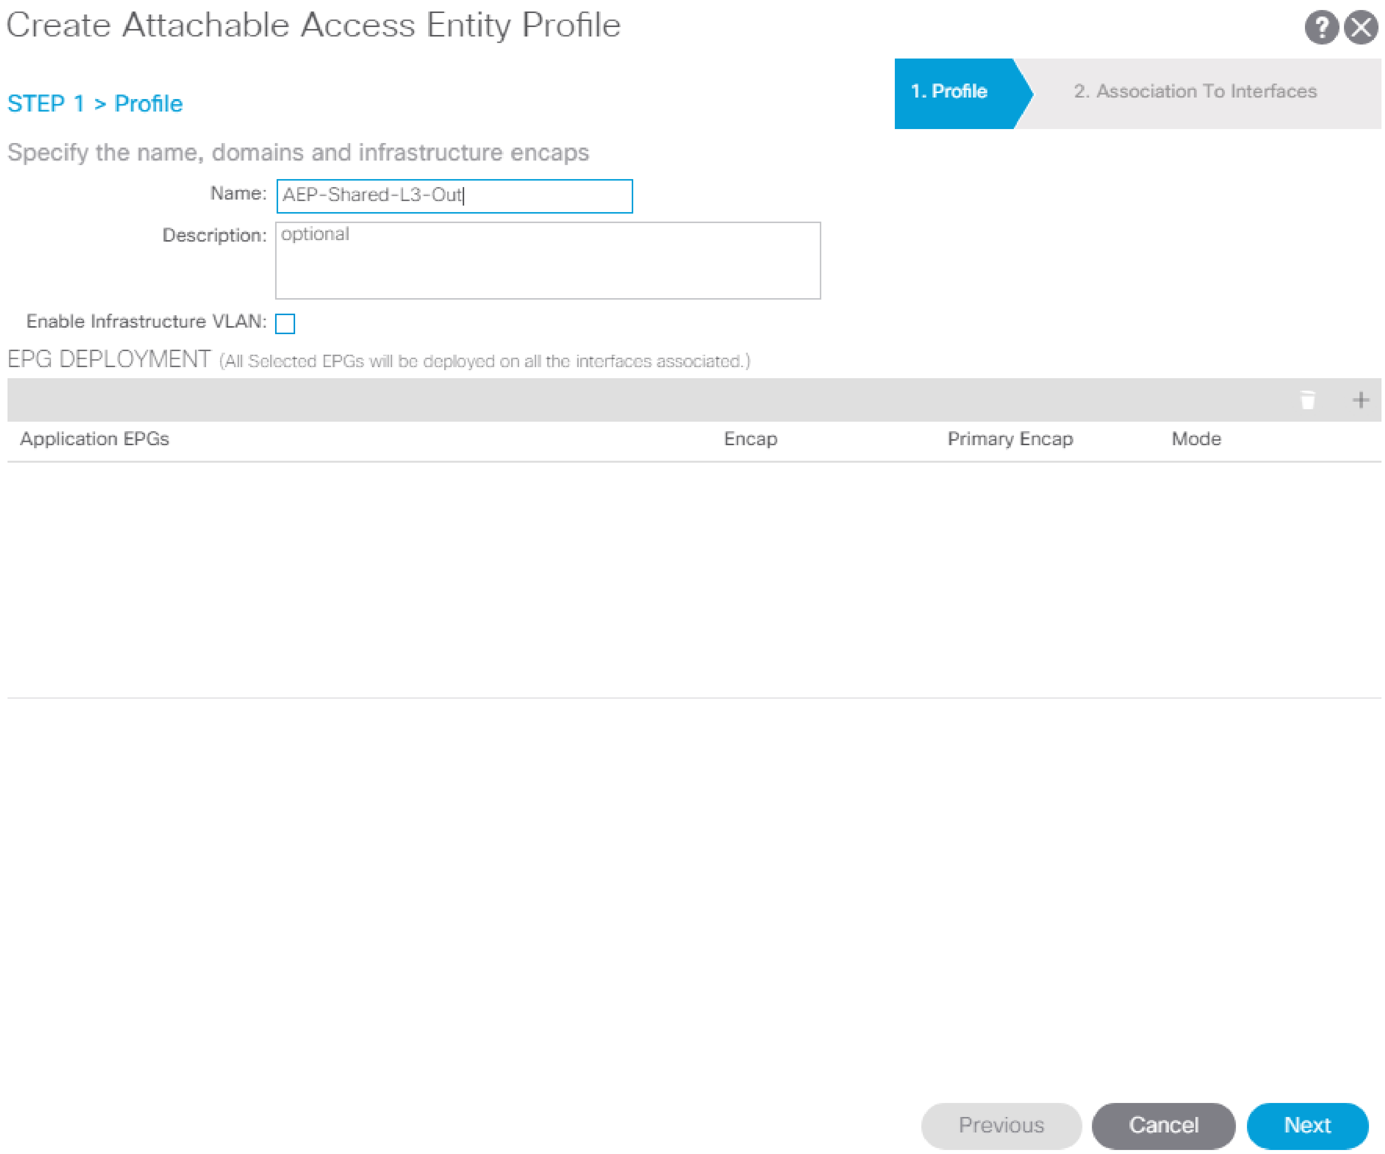

Create Attachable Access Entity Profile

STEP 1 > Profile

Specify the name, domains and infrastructure encaps

Name: AEP-Shared-L3-Out

Description: optional

Enable Infrastructure VLAN: □

EPG DEPLOYMENT (All Selected EPGs will be deployed on all the interfaces associated.)

Application EPGs	Encap	Primary Encap	Mode

7. Click Finish to continue without specifying interfaces.

8. Back in the Create Layer 3 Domain window, use the VLAN Pool drop-down list 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.

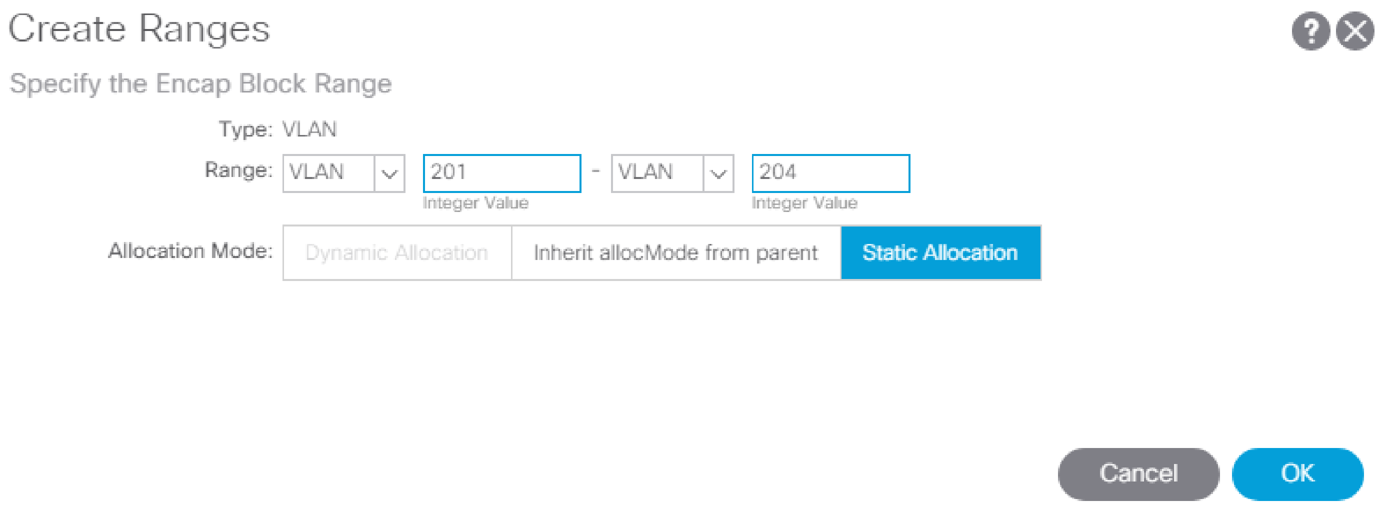

Create Ranges

Specify the Encap Block Range

Type: VLAN

Range: VLAN 201 - VLAN 204

Allocation Mode: Static Allocation

12. Click OK to complete adding the VLAN range.

13. Click Submit to complete creating the VLAN Pool.

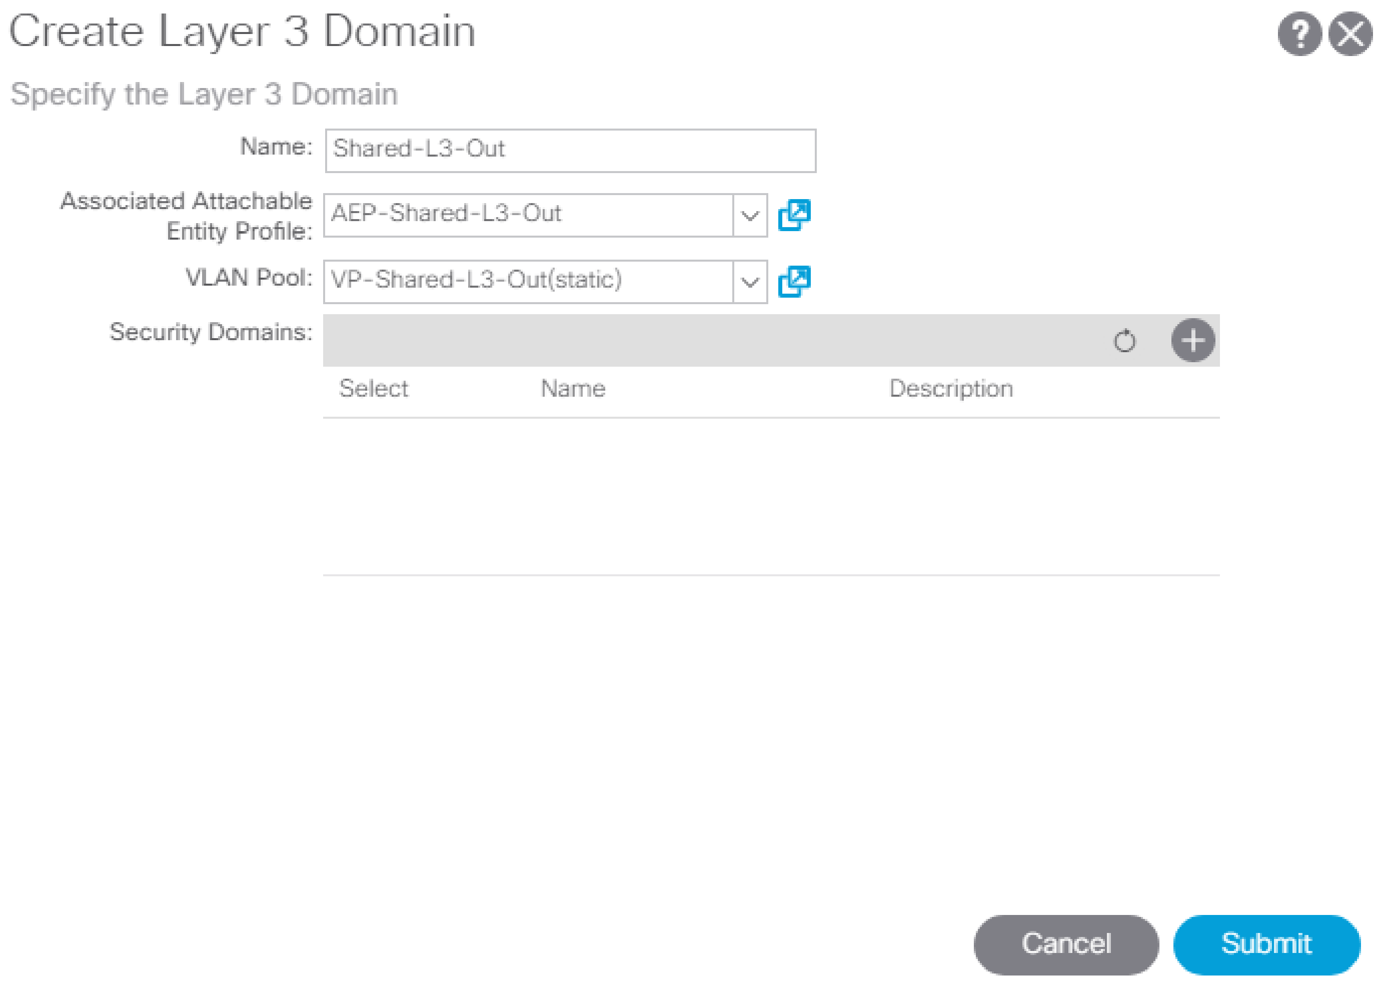

Create Layer 3 Domain

Specify the Layer 3 Domain

Name: Shared-L3-Out

Associated Attachable Entity Profile: AEP-Shared-L3-Out

VLAN Pool: VP-Shared-L3-Out(static)

Security Domains:

Select	Name	Description

14. Click Submit to complete creating the Layer 3 Domain.

15. At the top, select Fabric > Access Policies.

16. On the left, select Quick Start. Under Steps, select Configure an interface, PC, or VPC.

17. If an Interface, PC, or vPC has already been configured on the leaf pair being used here, select that switch pair in the list on the left and skip to step 19. Otherwise, in the center pane, click the green plus sign to select switches.

18. Using the Switches pull-down, select the two leaf switches connected to the Nexus 7000s and click away from the list to get the two switches filled in next to Switches. The Switch Profile Name will be automatically filled in.

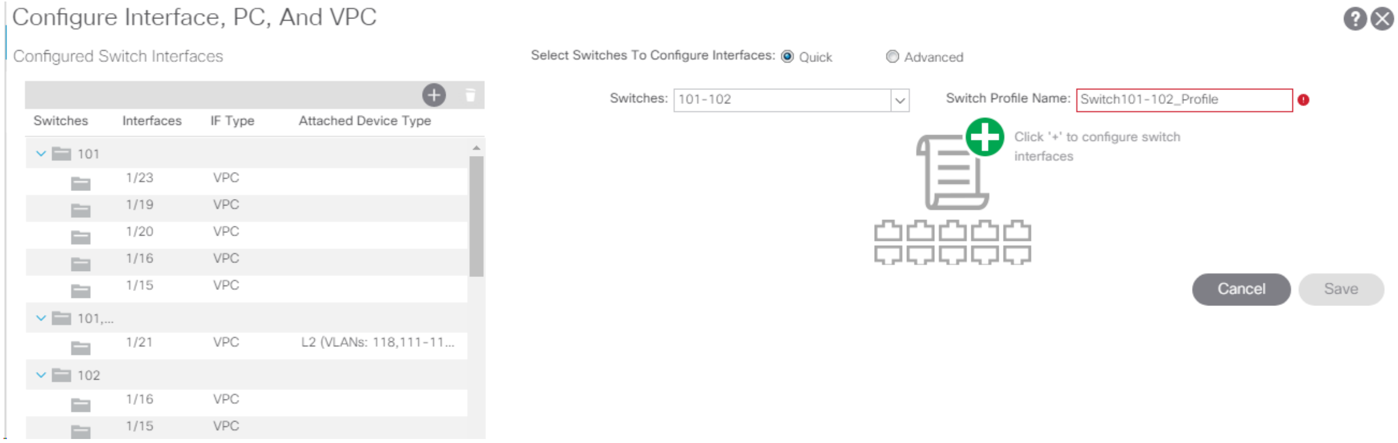

Configure Interface, PC, And VPC

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
101			
	1/23	VPC	
	1/19	VPC	
	1/20	VPC	
	1/16	VPC	
	1/15	VPC	
101,...			
	1/21	VPC	L2 (VLANs: 118,111-11...
102			
	1/16	VPC	
	1/15	VPC	

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: Switch101-102_Profile

Click '+' to configure switch interfaces

19. Click the green plus sign to configure switch interfaces.

20. Next to interfaces, enter the 2-port identifiers for the ports connected to the Nexus 7000s and used for Shared-L3-Out. It is important to use the same two ports on each leaf. Fill in the policies, Attached Device Type, and External Route Domain as shown below.

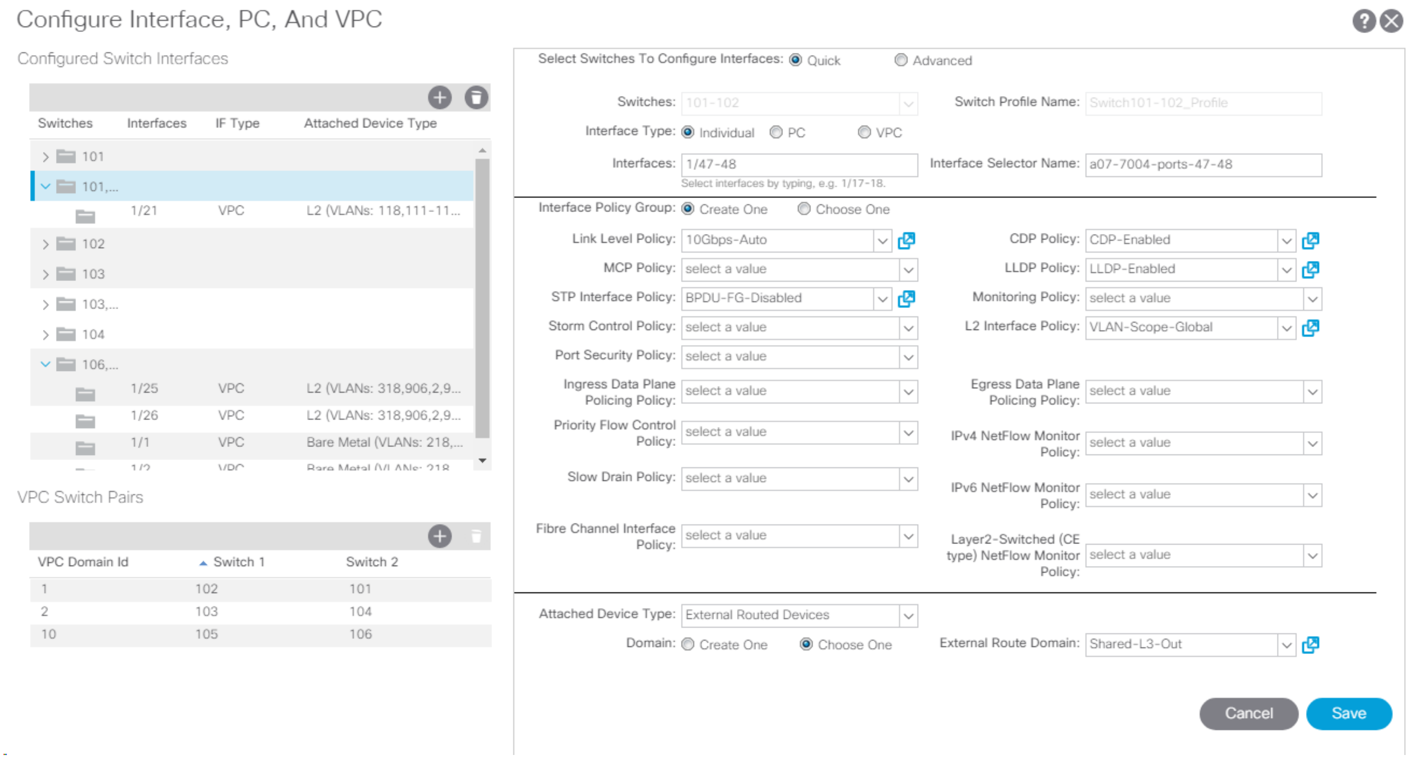

Configure Interface, PC, And VPC

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
101			
101,...			
	1/21	VPC	L2 (VLANs: 118,111-11...
102			
103			
103,...			
104			
106,...			
	1/25	VPC	L2 (VLANs: 318,906,2,9...
	1/26	VPC	L2 (VLANs: 318,906,2,9...
	1/1	VPC	Bare Metal (VLANs: 218,...
	1/2	VPC	Bare Metal (VLANs: 218

VPC Switch Pairs

VPC Domain Id	Switch 1	Switch 2
1	102	101
2	103	104
10	105	106

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: Switch101-102_Profile

Interface Type: ● Individual ○ PC ○ VPC

Interfaces: 1/47-48

Select interfaces by typing, e.g. 1/17-18.

Interface Selector Name: a07-7004-ports-47-48

Interface Policy Group: ● Create One ○ Choose One

Link Level Policy: 10Gbps-Auto

CDP Policy: CDP-Enabled

MCP Policy: select a value

LLDP Policy: LLDP-Enabled

STP Interface Policy: BPDU-FG-Disabled

Monitoring Policy: select a value

Storm Control Policy: select a value

L2 Interface Policy: VLAN-Scope-Global

Port Security Policy: select a value

Ingress Data Plane Policing Policy: select a value

Egress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

IPv4 NetFlow Monitor Policy: select a value

Slow Drain Policy: select a value

IPv6 NetFlow Monitor Policy: select a value

Fibre Channel Interface Policy: select a value

Layer2-Switched (CE type) NetFlow Monitor Policy: select a value

Attached Device Type: External Routed Devices

Domain: ○ Create One ● Choose One

External Route Domain: Shared-L3-Out

21. On the lower right, click Save. Click Save again and then click Submit.

22. At the top, select Tenants > common.

23. On the left, expand Tenant common and Networking.

24. Right-click VRFs and select create VRF.

Name the VRF `common-FP-External`. Select default for both the End Point Retention Policy and Monitoring Policy.

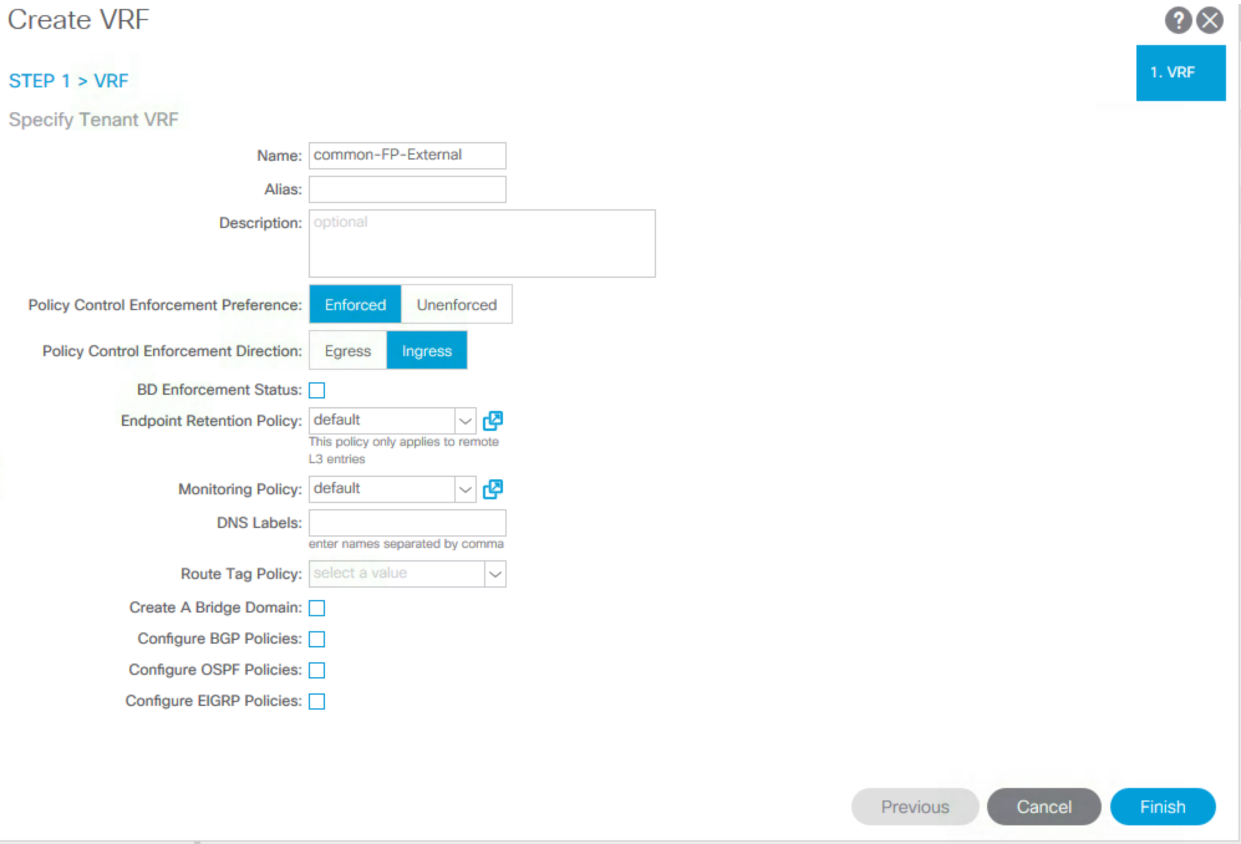

Create VRF

STEP 1 > VRF

Specify Tenant VRF

Name: common-FP-External

Alias:

Description: optional

Policy Control Enforcement Preference: [✓ Enforced] Unenforced

Policy Control Enforcement Direction: Egress [✓ Ingress]

BD Enforcement Status: □

Endpoint Retention Policy: default

This policy only applies to remote L3 entries

Monitoring Policy: default

DNS Labels:

enter names separated by comma

Route Tag Policy: select a value

Create A Bridge Domain: □

Configure BGP Policies: □

Configure OSPF Policies: □

Configure EIGRP Policies: □

25. Click Finish to complete creating the VRF.

26. On the left, right-click External Routed Networks and select Create Routed Outside.

27. Name the Routed Outside `Shared-L3-Out`.

28. Select the checkbox next to OSPF.

29. Enter 0.0.0.10 (configured in the Nexus 7000s) as the OSPF Area ID.

30. Using the VRF drop-down list, select common/`common-FP-External`.

31. Using the External Routed Domain drop-down list, select `Shared-L3-Out`.

32. Click the + sign to the right of Nodes and Interfaces Protocol Profiles to add a Node Profile.

33. Name the Node Profile Nodes-101-102 for the Nexus 9000 Leaf Switches.

34. Click the + sign to the right of Nodes to add a Node.

35. In the select Node window, select Leaf switch 101.

36. Provide a Router ID IP address that will also be used as the Loopback Address (192.168.254.101).

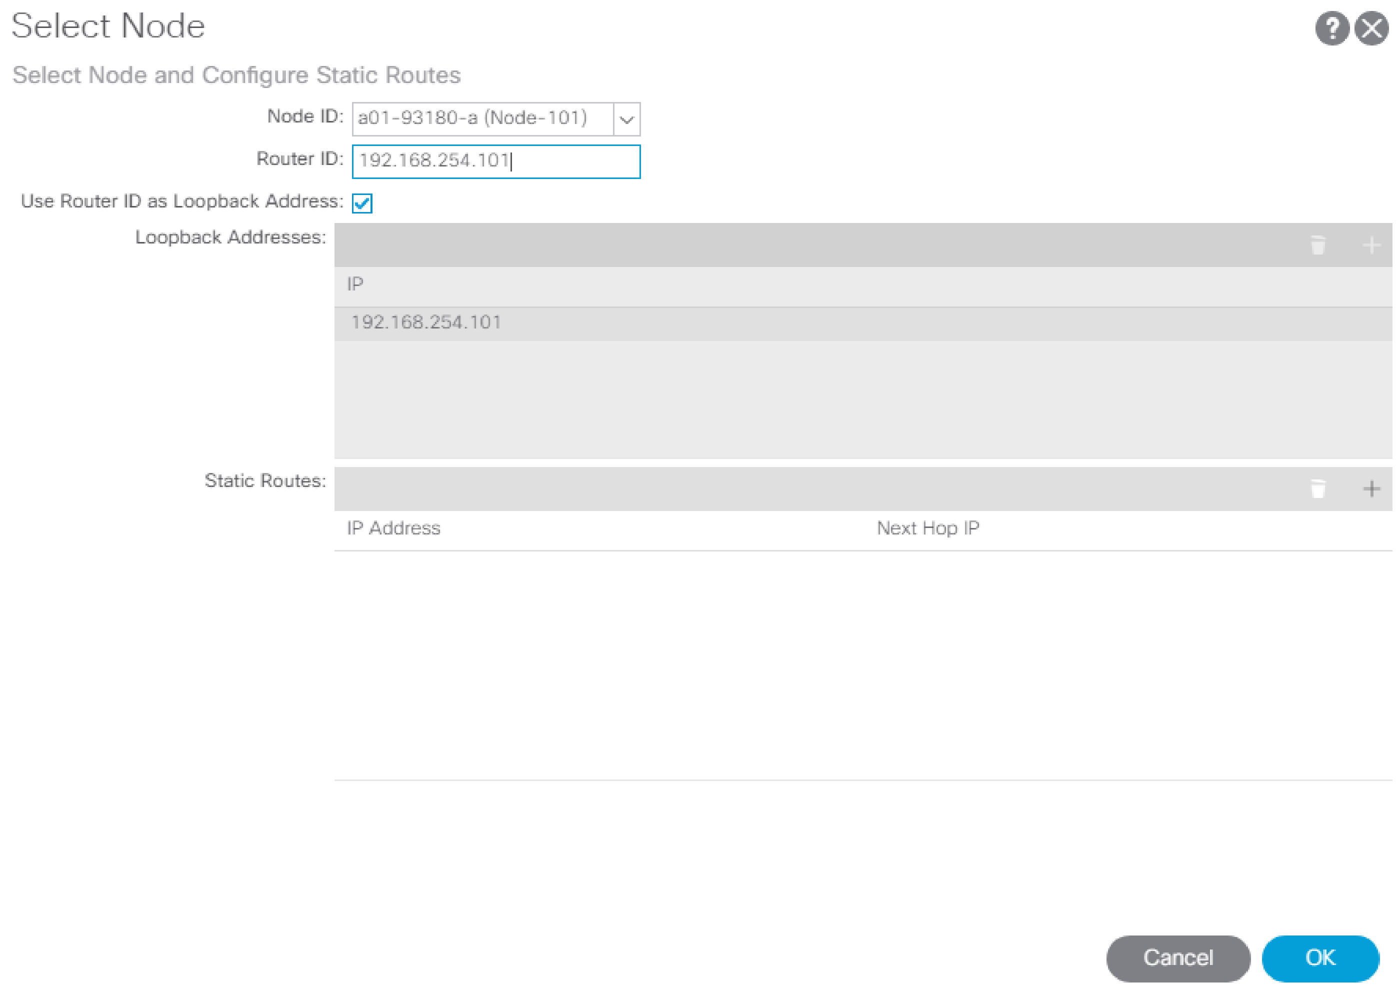

Select Node

Select Node and Configure Static Routes

Node ID: a01-93180-a (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

37. Click OK to complete selecting the Node.

38. Click the + sign to the right of Nodes to add a Node.

39. In the select Node window, select Leaf switch 102.

40. Provide a Router ID IP address that will also be used as the Loopback Address (192.168.254.102).

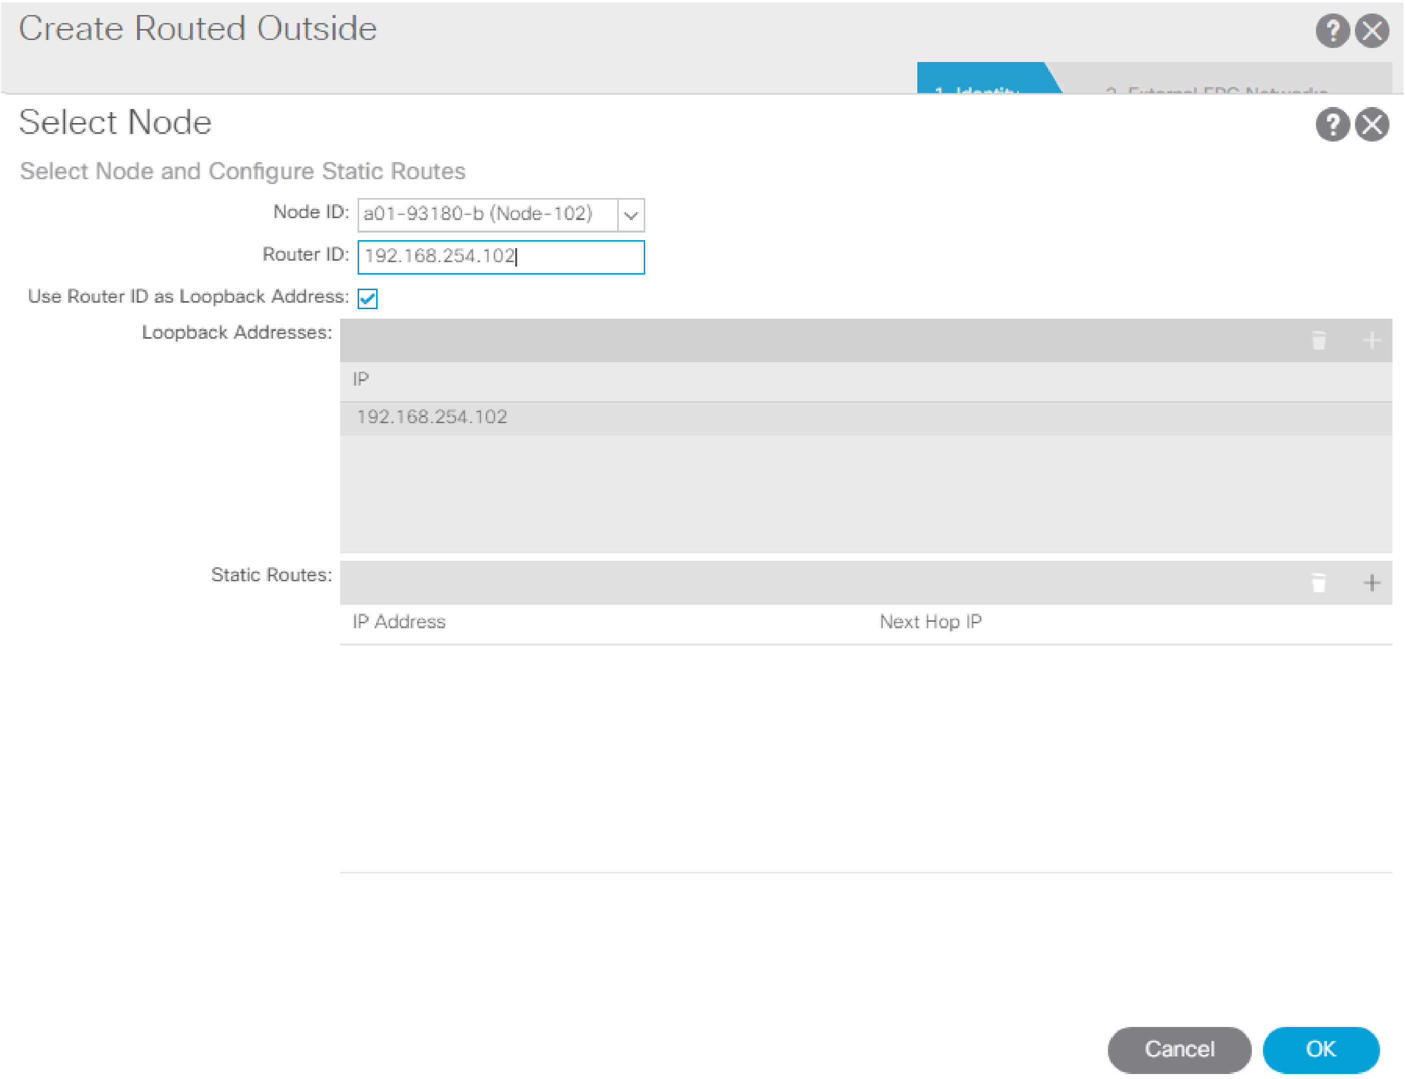

Create Routed Outside

Select Node

Select Node and Configure Static Routes

Node ID: a01-93180-b (Node-102)

Router ID: 192.168.254.102

Use Router ID as Loopback Address: ☑

Loopback Addresses:

IP
192.168.254.102

Static Routes:

IP Address	Next Hop IP

41. Click OK to complete selecting the Node.

42. Click the + sign to the right of OSPF Interface Profiles to create an OSPF Interface Profile.

43. Name the profile OIP-Nodes-101-102.

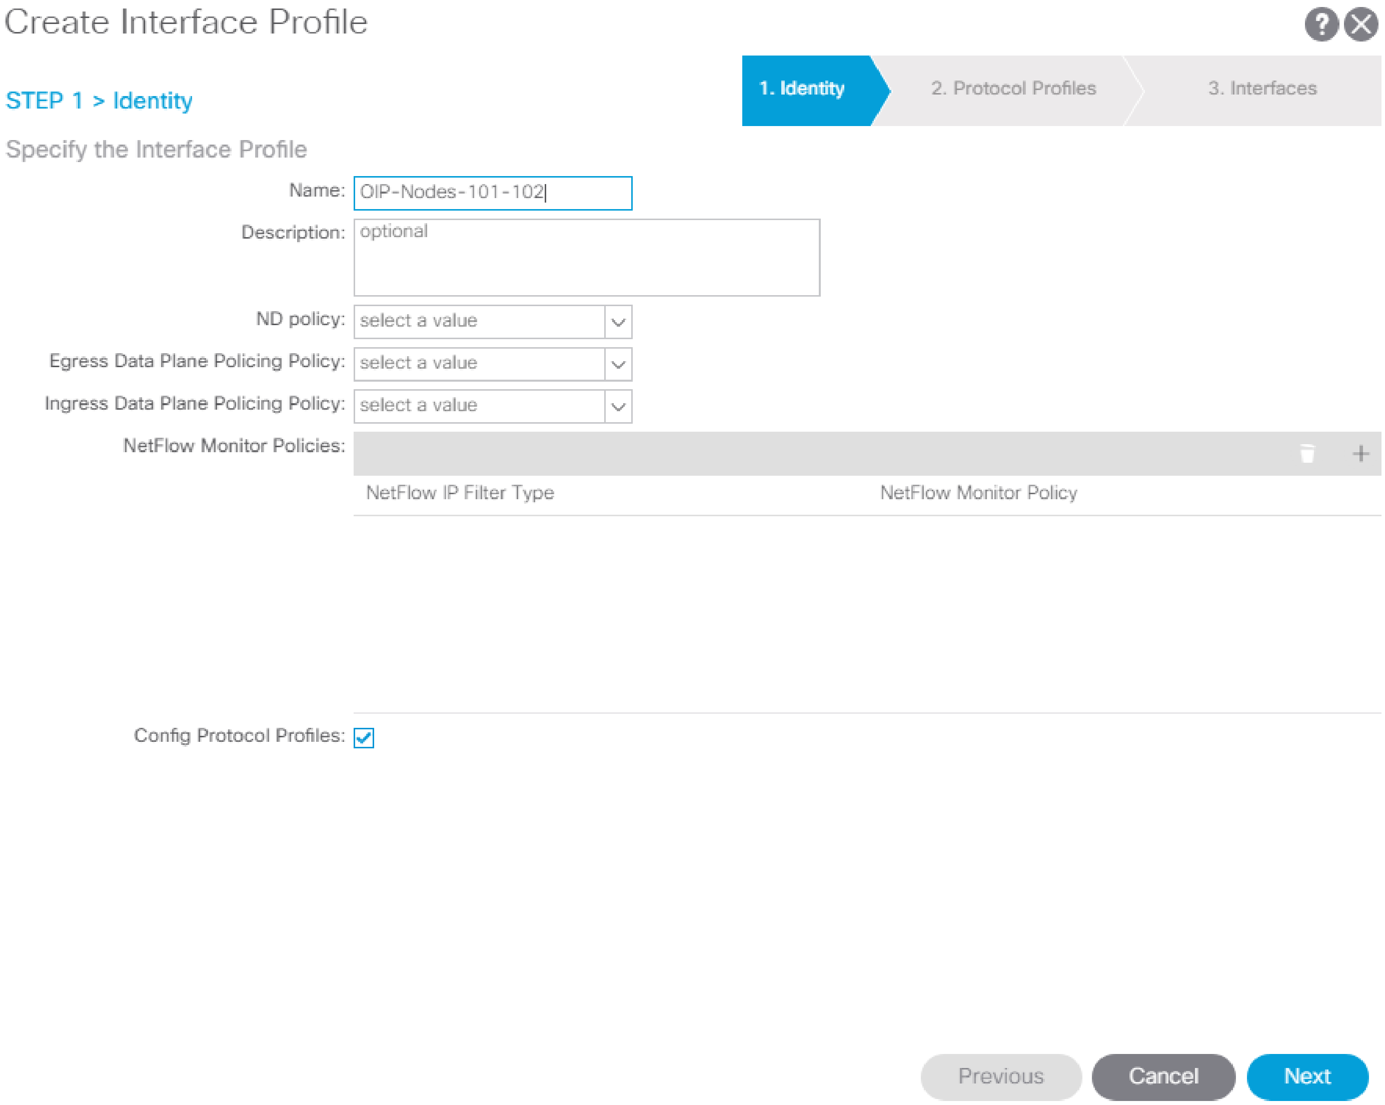

Create Interface Profile

STEP 1 > Identity

Specify the Interface Profile

Name: OIP-Nodes-101-102

Description: optional

ND policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

NetFlow Monitor Policies:

NetFlow IP Filter Type	NetFlow Monitor Policy

Config Protocol Profiles: ☑

44. Click Next.

45. Using the OSPF Policy drop-down list, select Create OSPF Interface Policy.

46. Name the policy To-7K.

47. Select the Point-to-Point Network Type.

48. Select the Advertise subnet and MTU ignore Interface Controls.

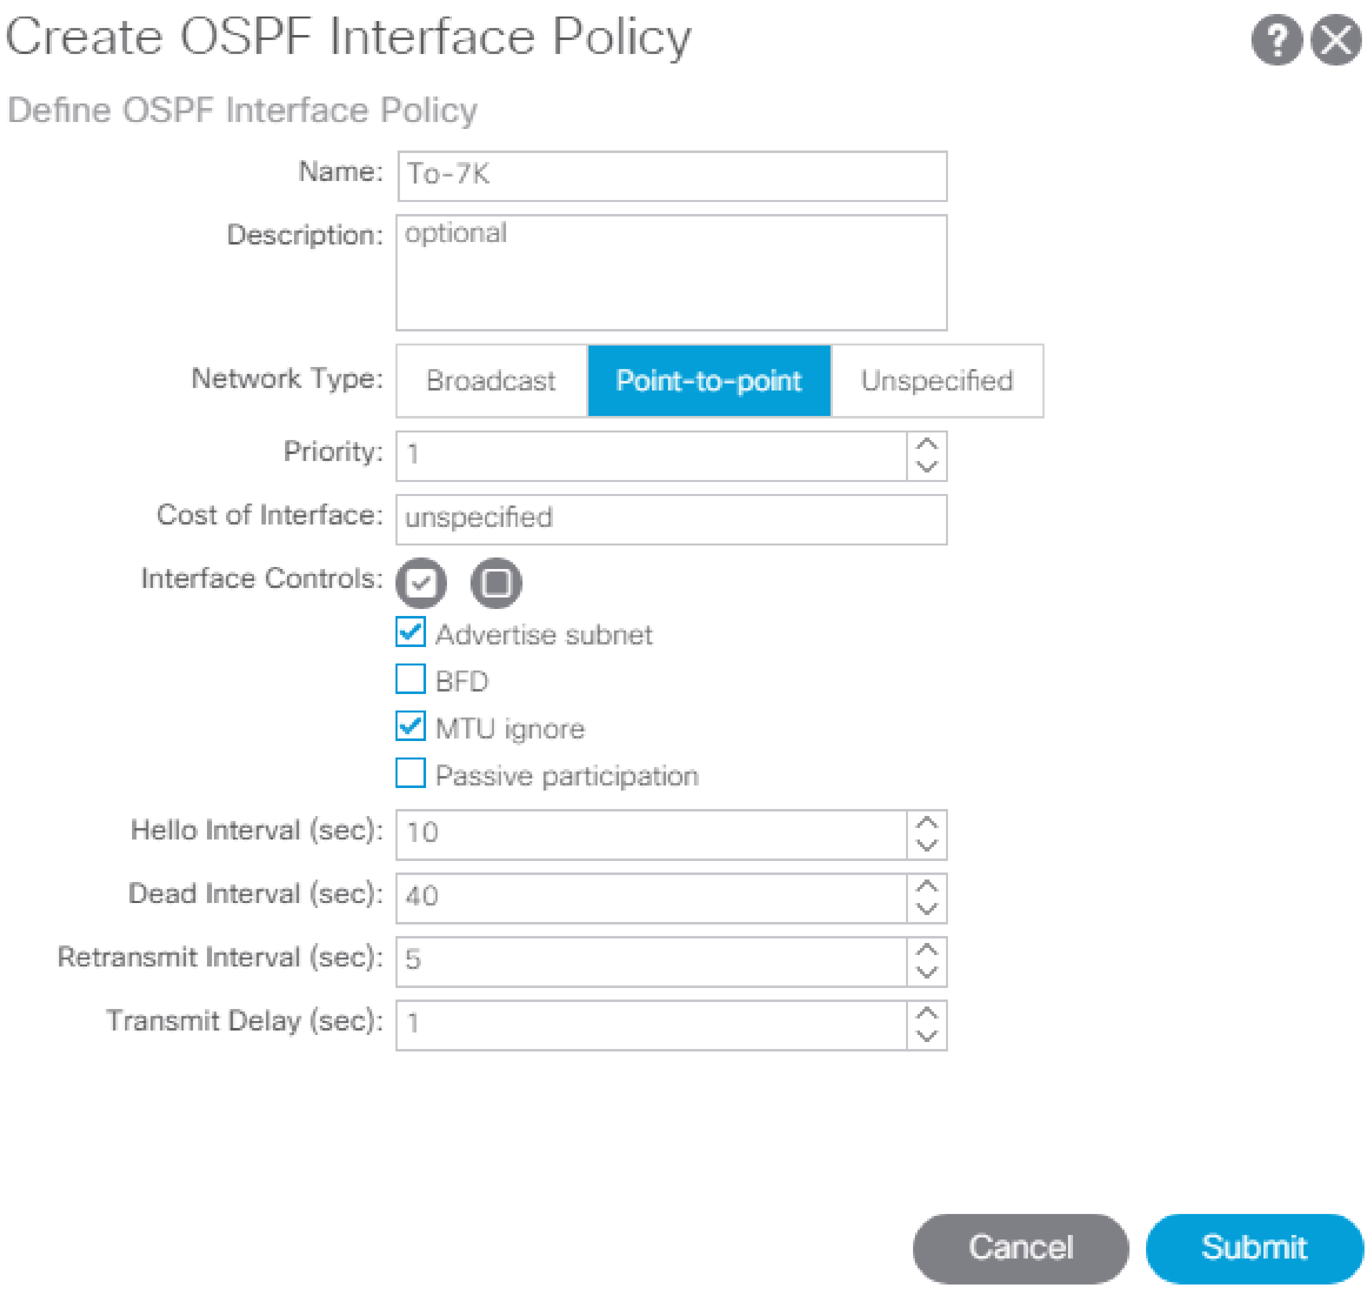

Create OSPF Interface Policy

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

49. Click SUBMIT to complete creating the policy.

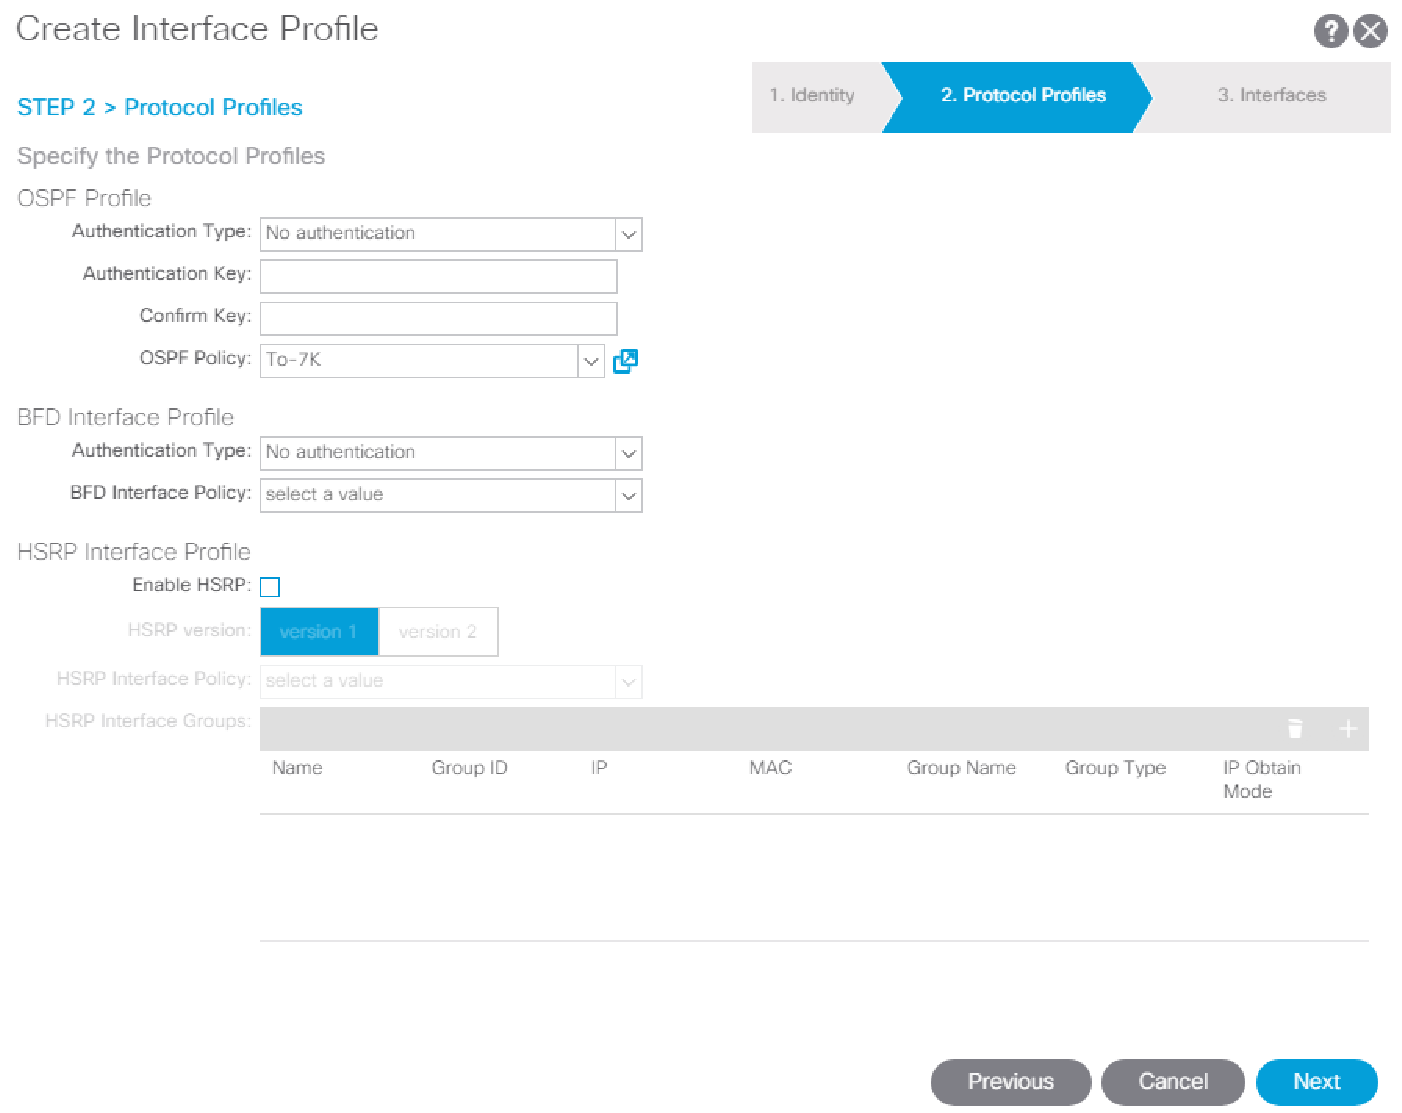

Create Interface Profile

STEP 2 > Protocol Profiles

Specify the Protocol Profiles

OSPF Profile

Authentication Type: No authentication

Authentication Key:

Confirm Key:

OSPF Policy: To-7K

BFD Interface Profile

Authentication Type: No authentication

BFD Interface Policy: select a value

HSRP Interface Profile

Enable HSRP: □

HSRP version: version 1 | version 2

HSRP Interface Policy: select a value

HSRP Interface Groups:

Name	Group ID	IP	MAC	Group Name	Group Type	IP Obtain Mode

50. Click Next.

51. Select Routed Sub-Interface under Interfaces.

52. Click the + sign to the right of Routed Sub-Interfaces to add a routed sub-interface.

53. In the Select Routed Sub-Interface window, select the interface on Node 101 that is connected to Nexus 7000-1.

54. Enter VLAN 201 for Encap.

55. Enter the IPv4 Primary Address (192.168.253.101/30)

56. Leave the MTU set to inherit.

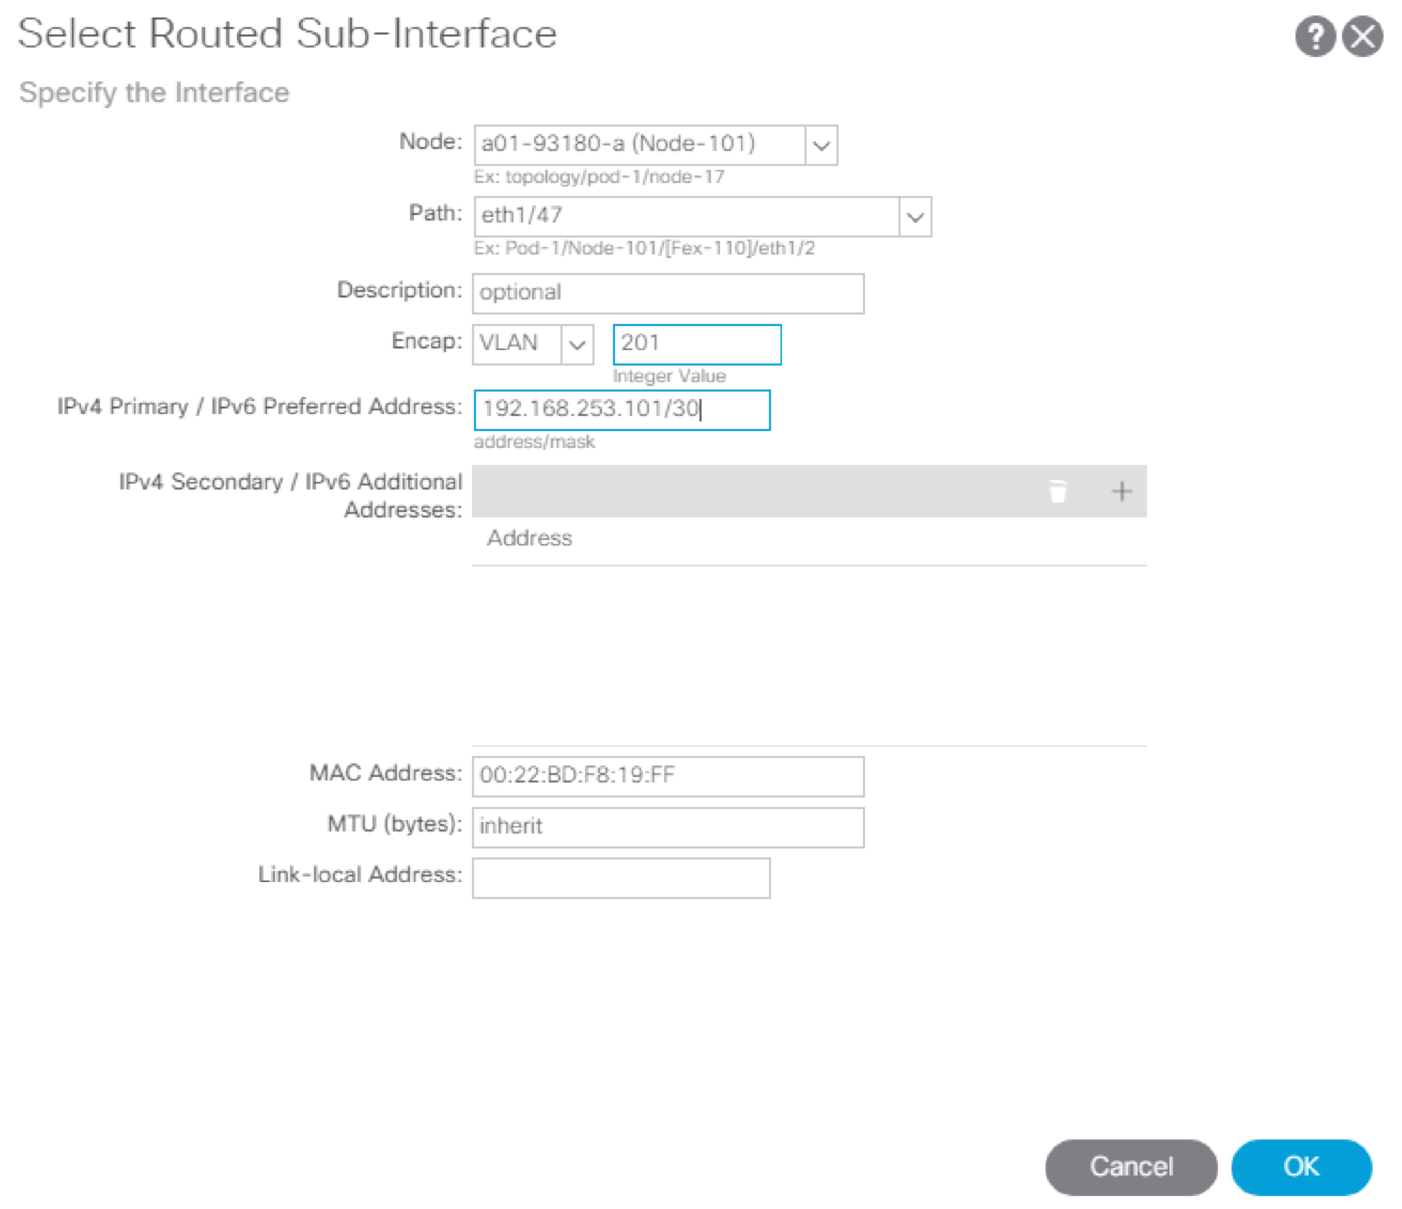

Select Routed Sub-Interface

Specify the Interface

Node:	a01-93180-a (Node-101)
Path:	eth1/47
Description:	optional
Encap:	VLAN 201
IPv4 Primary / IPv6 Preferred Address:	192.168.253.101/30
IPv4 Secondary / IPv6 Additional Addresses:	Address
MAC Address:	00:22:BD:F8:19:FF
MTU (bytes):	inherit
Link-local Address:	

Ex: topology/pod-1/node-17
Ex: Pod-1/Node-101/[Fex-110]/eth1/2
Integer Value
address/mask
Cancel
OK

57. Click OK to complete creating the routed sub-interface.

58. 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).

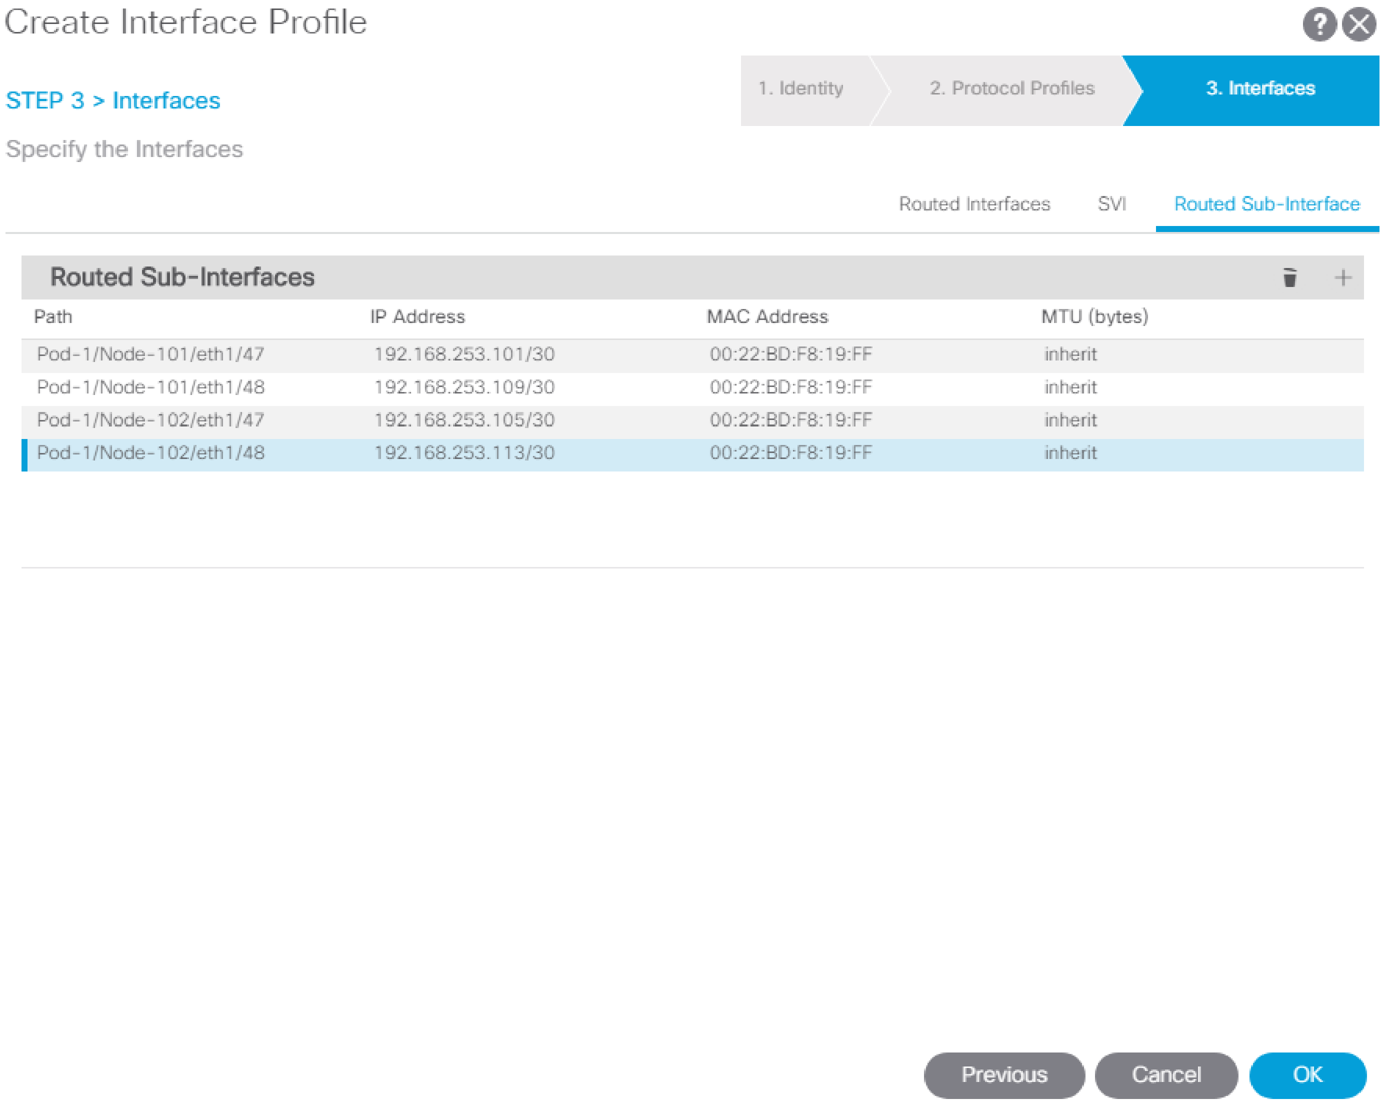

STEP 3 > Interfaces

Specify the Interfaces

Routed Sub-Interfaces

Path	IP Address	MAC Address	MTU (bytes)
Pod-1/Node-101/eth1/47	192.168.253.101/30	00:22:BD:F8:19:FF	inherit
Pod-1/Node-101/eth1/48	192.168.253.109/30	00:22:BD:F8:19:FF	inherit
Pod-1/Node-102/eth1/47	192.168.253.105/30	00:22:BD:F8:19:FF	inherit
Pod-1/Node-102/eth1/48	192.168.253.113/30	00:22:BD:F8:19:FF	inherit

59. Click OK to complete creating the Node Interface Profile.

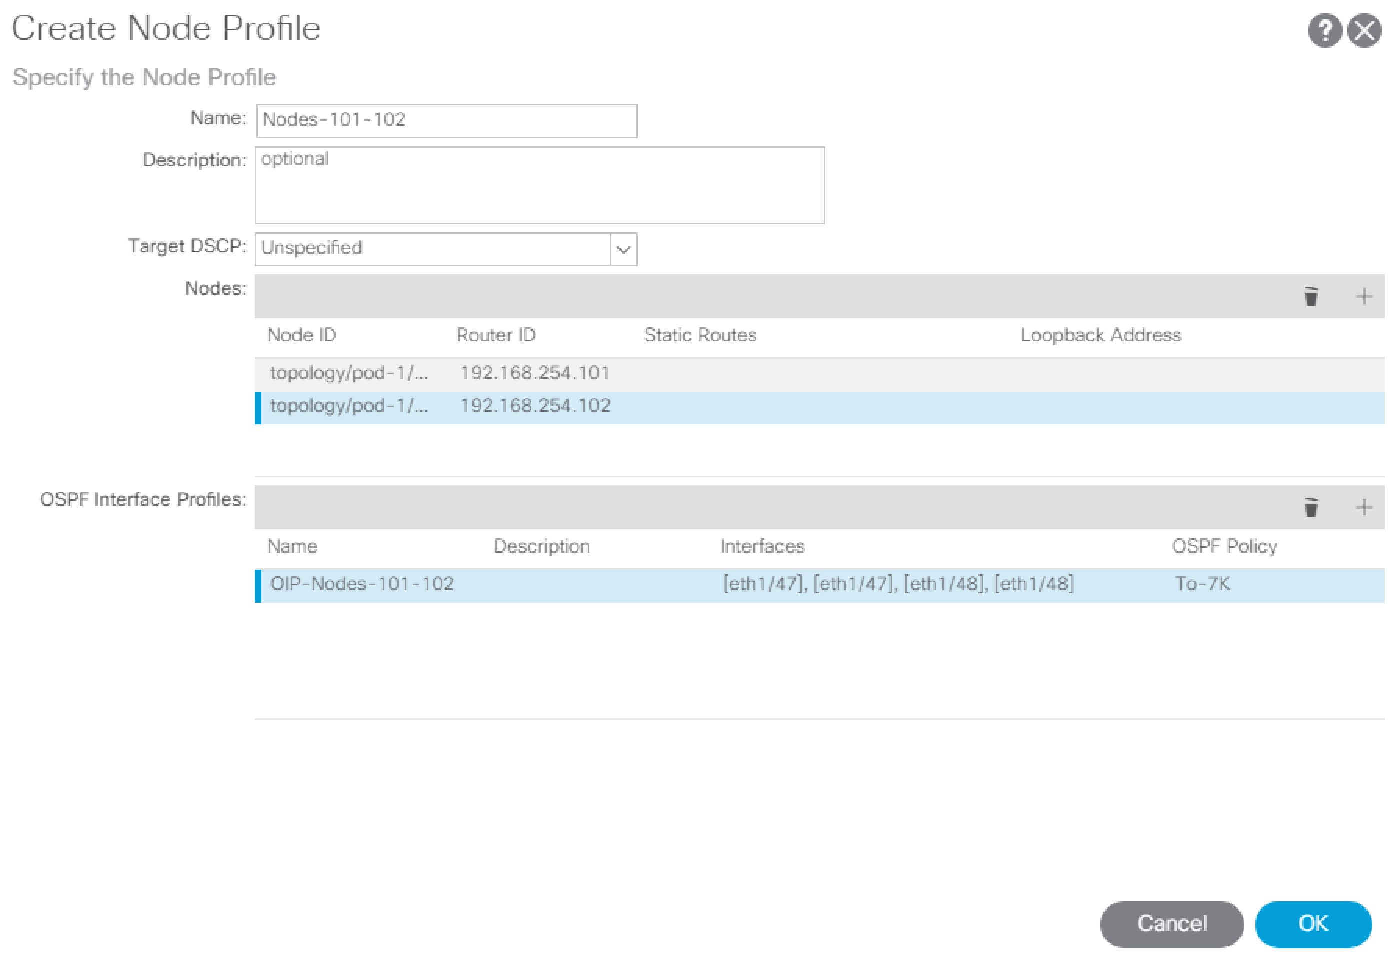

Create Node Profile

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/...	192.168.254.101		
topology/pod-1/...	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

60. Click OK to complete creating the Node Profile.

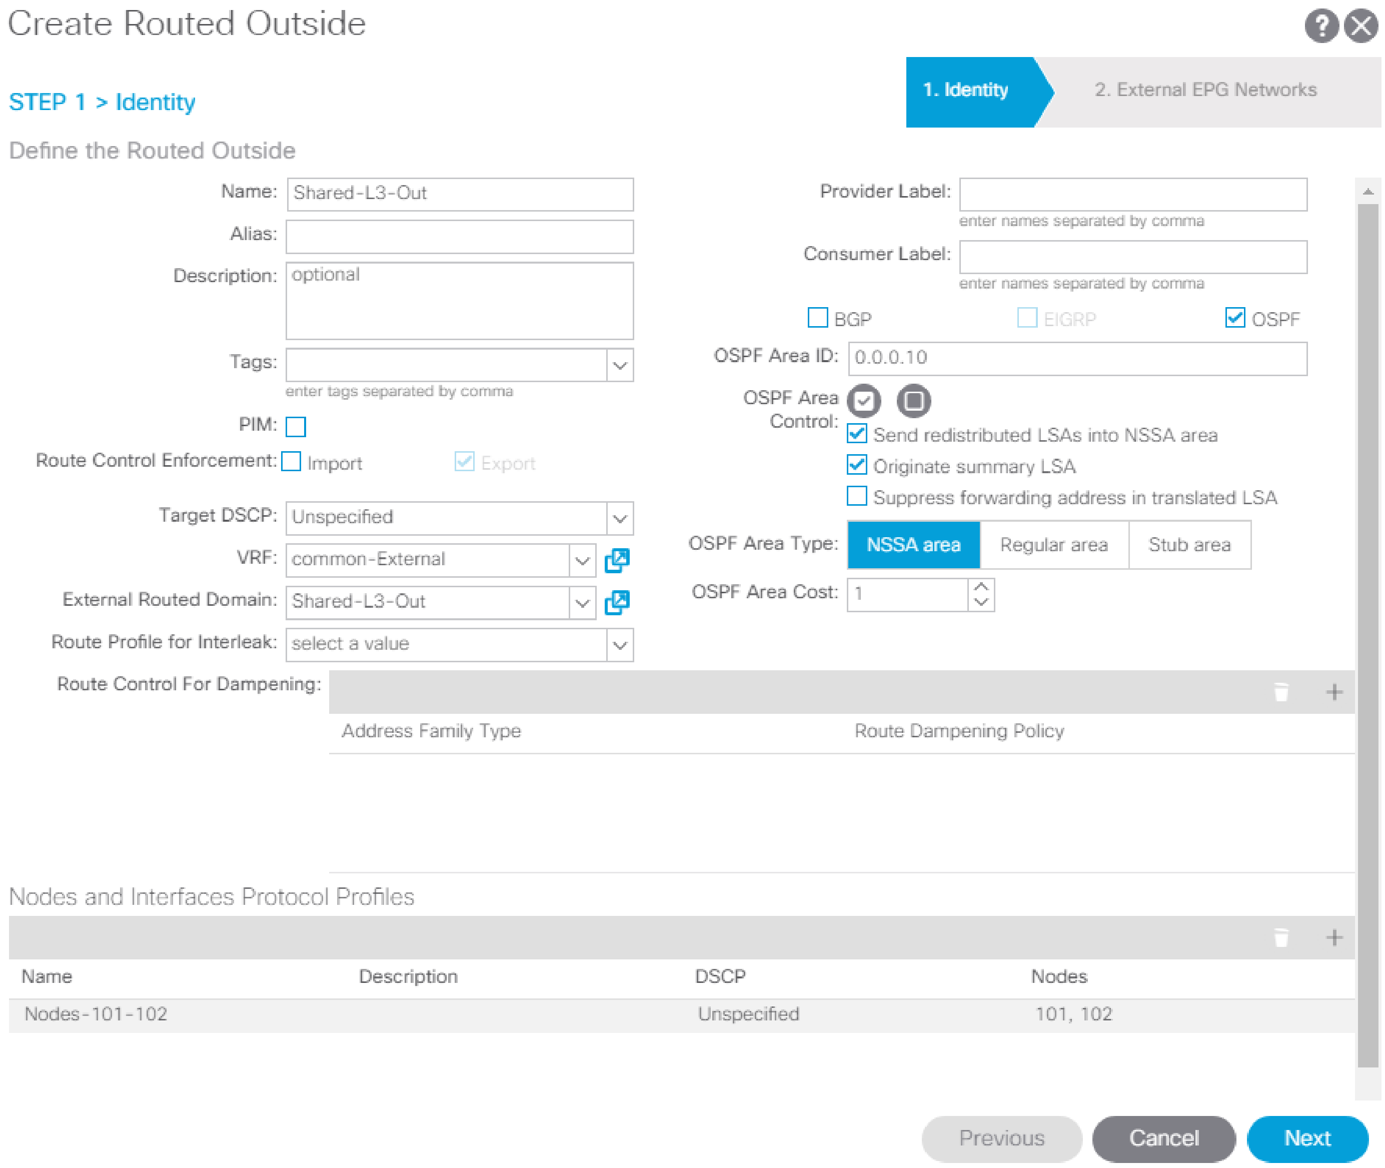

Create Routed Outside

STEP 1 > Identity

Define the Routed Outside

Name: Shared-L3-Out

Alias:

Description: optional

Tags:

enter tags separated by comma

PIM: □

Route Control Enforcement: □ Import ☑ Export

Target DSCP: Unspecified

VRF: common-External

External Routed Domain: Shared-L3-Out

Route Profile for Interleak: select a value

Route Control For Dampening:

Address Family Type	Route Dampening Policy

Provider Label:

enter names separated by comma

Consumer Label:

enter names separated by comma

□ BGP □ EIGRP ☑ OSPF

OSPF Area ID: 0.0.0.10

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: 1

Nodes and Interfaces Protocol Profiles

Name	Description	DSCP	Nodes
Nodes-101-102		Unspecified	101, 102

61. Click Next.

62. Click the + sign under External EPG Networks to create and External EPG Network.

63. Name the External Network Default-Route.

64. Click the + sign to add a Subnet.

65. 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.

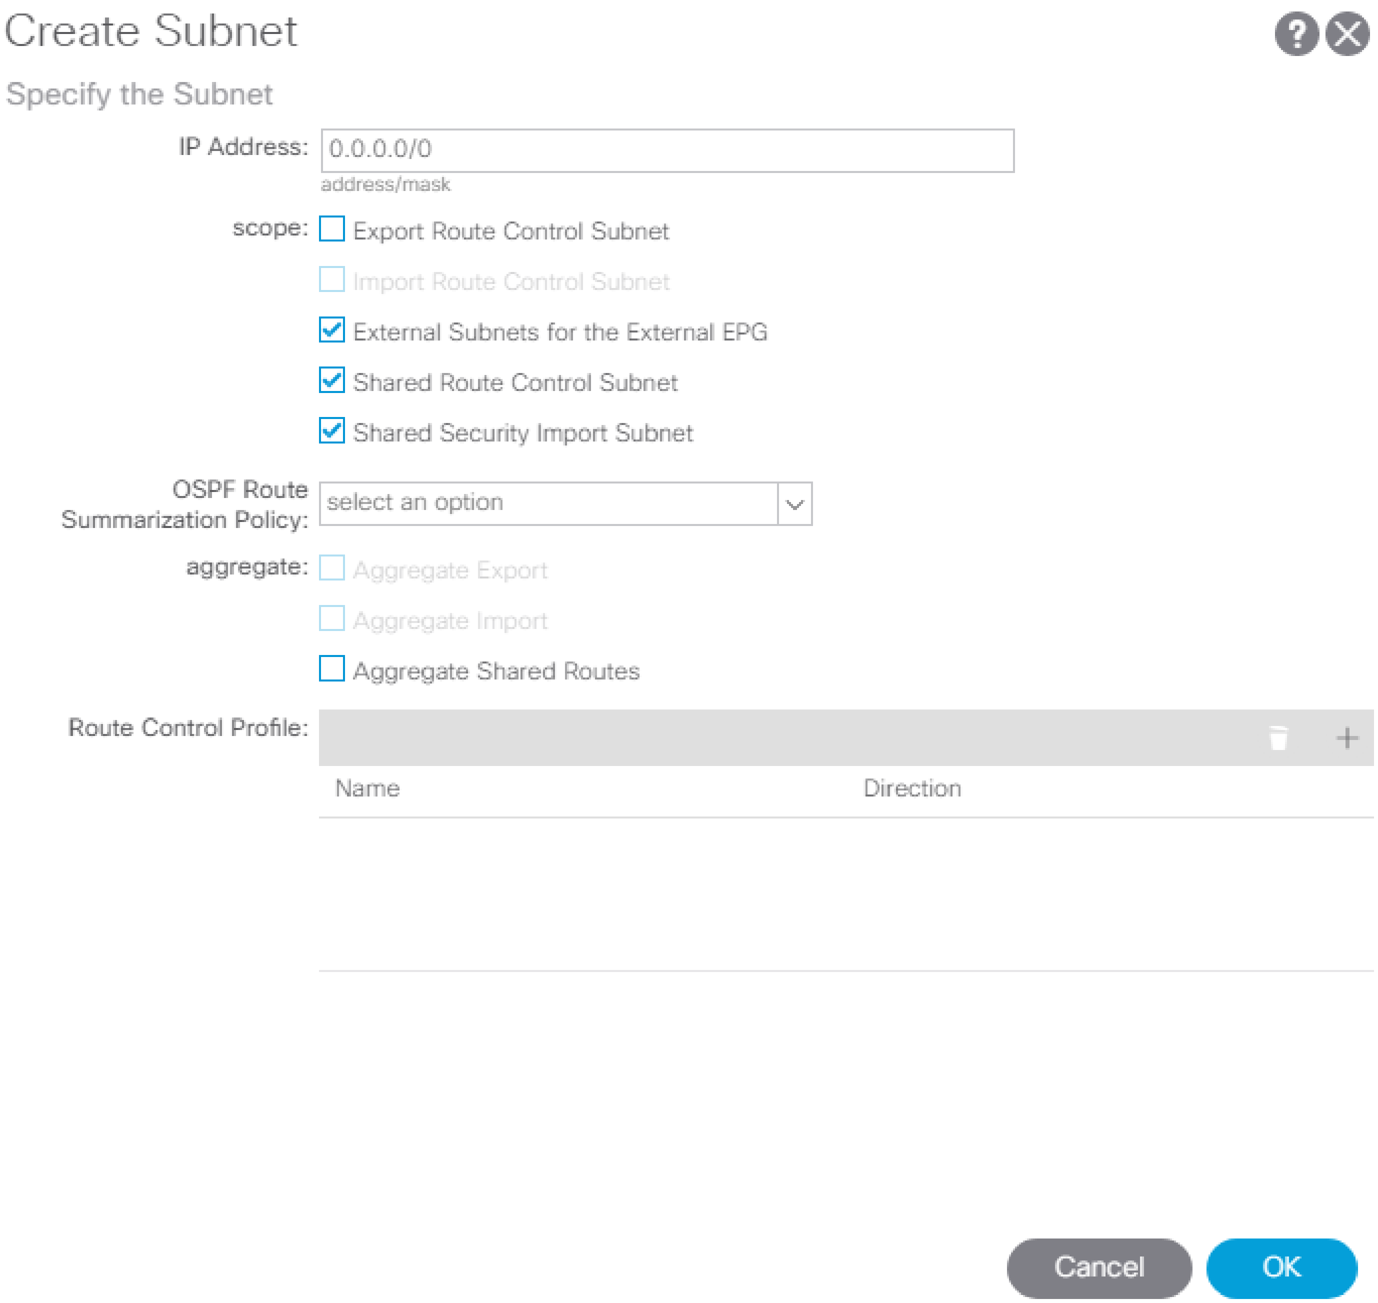

Create Subnet

Specify the Subnet

IP Address: 0.0.0.0/0

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: select an option

aggregate:
☐ Aggregate Export
☐ Aggregate Import
☐ Aggregate Shared Routes

Route Control Profile:

Name	Direction

66. Click OK to complete creating the subnet.

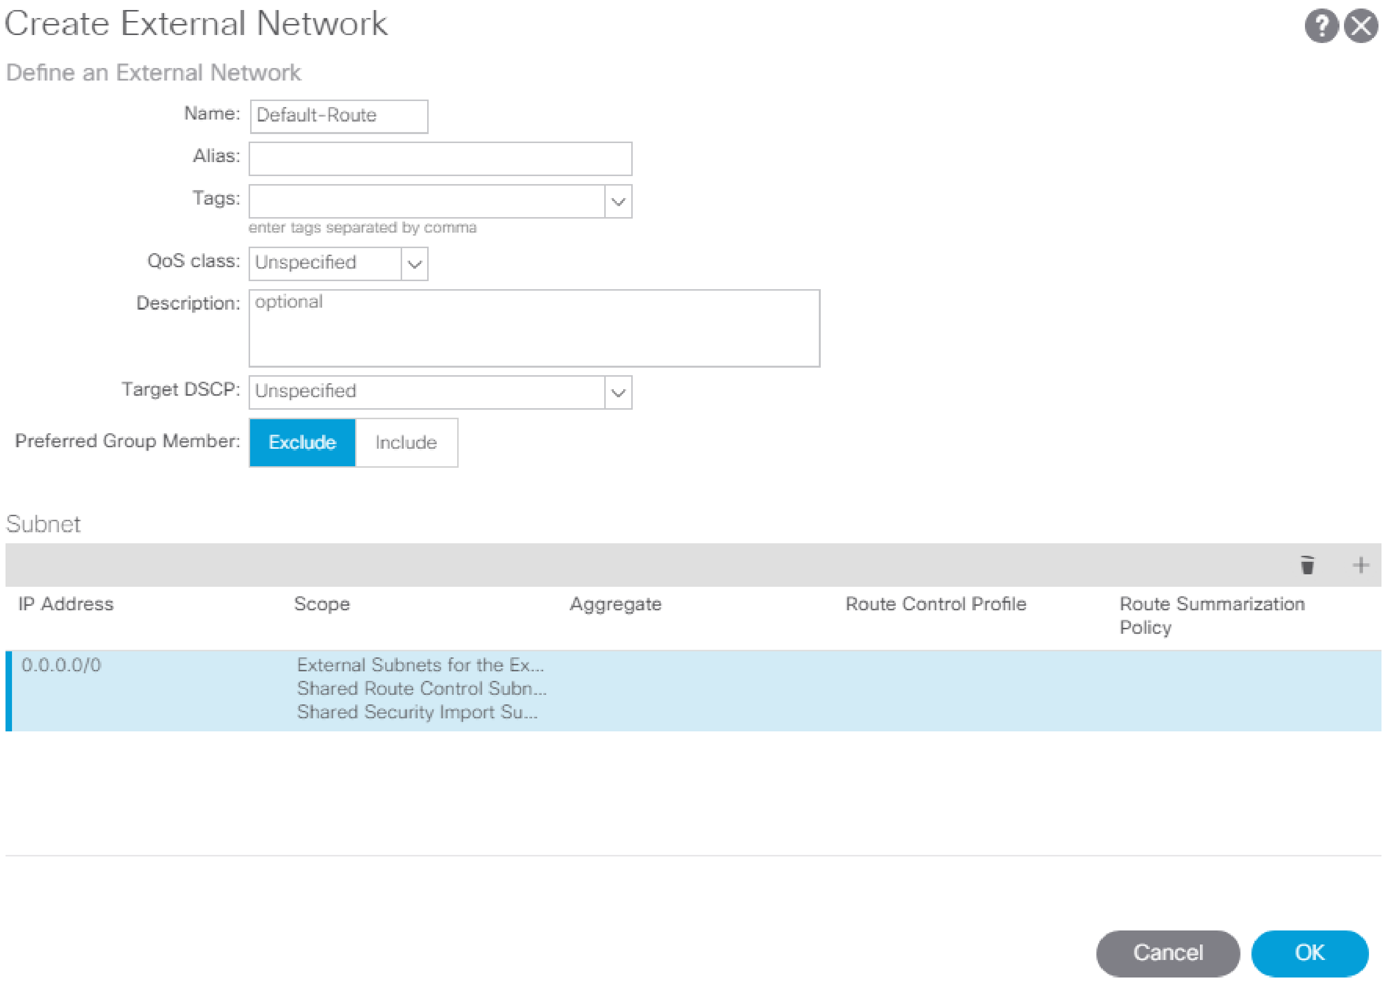

Create External Network

Define an External Network

Name:	Default-Route
Alias:	
Tags:	enter tags separated by comma
QoS class:	Unspecified
Description:	optional
Target DSCP:	Unspecified
Preferred Group Member:	Exclude / Include

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 Su...			

67. Click OK to complete creating the external network.

68. Click Finish to complete creating the Shared-L3-Out.

69. On the left, right-click Contracts and select Create Contract.

70. Name the contract Allow-Shared-L3-Out.

71. Select the Global Scope to allow the contract to be consumed from all tenants.

72. Click the + sign to the right of Subjects to add a contract subject.

73. Name the subject Allow-All.

74. Click the + sign to the right of Filters to add a filter.

75. Use the drop-down list to select the Allow-All filter from Tenant common.

76. Click Update.

77. Click OK to complete creating the contract subject.

78. Click Submit to complete creating the contract.

79. On the left, expand Tenant common, Networking, External Routed Networks, Shared-L3-Out, and Networks. Select Default-Route.

80. On the right, under Policy, select Contracts.

81. Click the + sign to the right of Provided Contracts to add a Provided Contract.

82. Select the common/Allow-Shared-L3-Out contract and click Update.

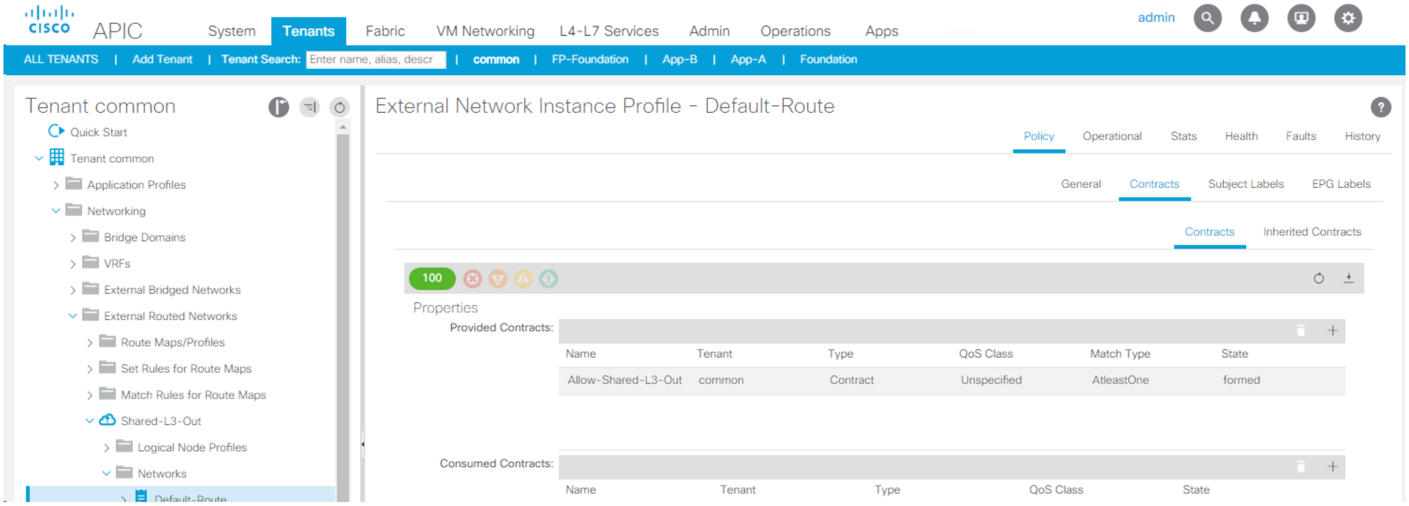

Provided Contracts:

Name	Tenant	Type	QoS Class	Match Type	State
Allow-Shared-L3-Out	common	Contract	Unspecified	AtleastOne	formed

Consumed Contracts:

Name	Tenant	Type	QoS Class	State

Tenant EPGs can now consume the Allow-Shared-L3-Out contract and connect outside of the fabric. More restrictive contracts can be built and provided here for more restrictive access to the outside.

Lab Validation Tenant Configuration

The following table lists the VLANs, Subnets, and Bridge Domains for the sample App-A Tenant set up as part of this lab validation:

Table 4 Lab Validation Tenant FPV-App-A Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI	3014	Virtual Switch	192.168.14.0/24 – L2	BD-iSCSI
NFS	3054	Virtual Switch	192.168.54.0/24 – L2	BD-NFS
SVM-MGMT	220	Virtual Switch	172.18.254.22/29	BD-Internal
Web	N/A	Virtual Switch	172.16.0.254/24	BD-Internal
App	N/A	Virtual Switch	172.16.1.254/24	BD-Internal
DB	N/A	Virtual Switch	172.16.2.254/24	BD-Internal

Deploy ACI Application (FPV-App-A) Tenant

This section details the steps for creation of the FPV-App-A Sample Tenant in the ACI Fabric. This tenant will host application connectivity between the compute (ESXi Server on UCS) and the storage (NetApp) environments. This tenant will also host the three application tiers of the sample three-tier application. A corresponding FPV-App-A-SVM will be created on the NetApp storage to align with this tenant. To deploy the FPV-App-A Tenant, complete the following steps:

Note that in this validation, only one iSCSI VLAN was used. It is assumed that the Tenant host servers will boot from LUNs in the Foundation SVM using the infrastructure iSCSI interfaces. The Tenant iSCSI interfaces will be placed on the VMware vDS.

APIC GUI

1. In the APIC GUI, select Fabric > Access Policies.
2. On the left, expand Pools and VLAN.
3. Select the storage VLAN Pool created earlier (NetApp-AFF_vlans).
4. In the center pane, click the + sign to add an encapsulation block.
5. Enter `<storage-FPV-App-A-NFS-VLAN>` for the From and To fields.
6. Select Static Allocation.
7. Leave Role set to External or On the wire encapsulations.

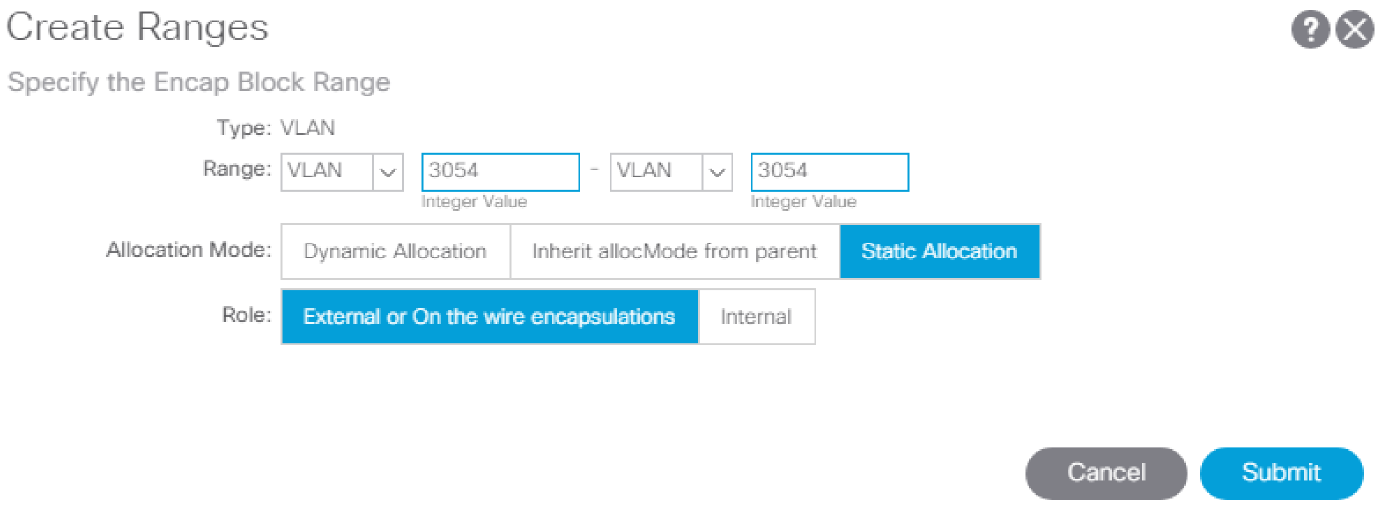

Create Ranges

Specify the Encap Block Range

Type: VLAN

Range: VLAN 3054 (Integer Value) - VLAN 3054 (Integer Value)

Allocation Mode: Dynamic Allocation | Inherit allocMode from parent | [✓ Static Allocation]

Role: [✓ External or On the wire encapsulations] | Internal

8. Click Submit to add the Encap Block Range.
9. In the center pane, click the + sign to add another encapsulation block.
10. Enter `<storage-FPV-App-A-SVM-MGMT-VLAN>` for the From and To fields.
11. Select Static Allocation.
12. Click Submit to add the Encap Block Range.
13. If iSCSI LUN access is being provided by the FPV-App-A tenant, complete steps 14-18. Otherwise, continue at step 19.
14. In the center pane, click the + sign to add an encapsulation block.
15. Enter `<storage-FPV-App-A-iSCSI-VLAN>` for the From and To fields.

16. Select Static Allocation.

17. Leave Role set to External or On the wire encapsulations.

18. Click SUBMIT to add the Encap Block Range.

19. At the top select Tenants > Add Tenant.

20. Name the Tenant `FPV-App-A`. Select the default Monitoring Policy.

21. For the VRF Name, also enter `FPV-App-A`. Leave the Take me to this tenant when I click finish checkbox checked.

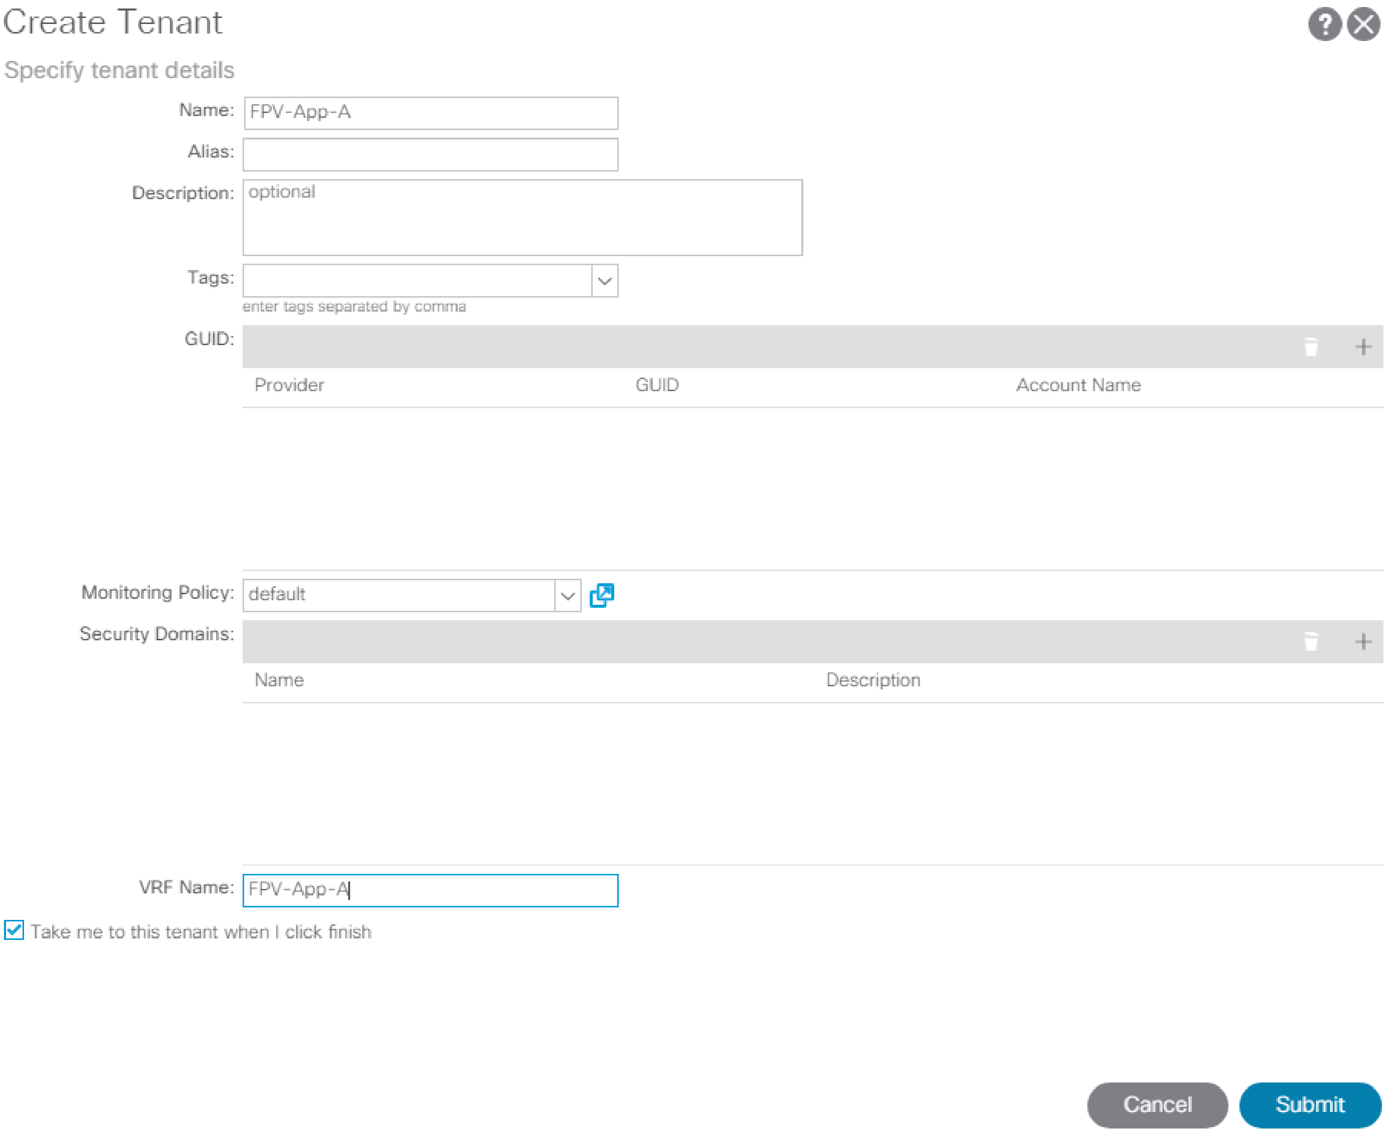

Create Tenant

Specify tenant details

Name: FPV-App-A

Alias:

Description: optional

Tags:

enter tags separated by comma

GUID:

Provider	GUID	Account Name

Monitoring Policy: default

Security Domains:

Name	Description

VRF Name: FPV-App-A

☑ Take me to this tenant when I click finish

22. Click Submit to finish creating the Tenant.

23. On the left under Tenant FPV-App-A, right-click Application Profiles and select Create Application Profile.

24. Name the Application Profile Host-Conn select the default Monitoring Policy, and click Submit to complete adding the Application Profile.

25. If you are using providing iSCSI LUN access from this tenant, complete steps 26-65. Otherwise, continue to step 66.

26. On the left, expand Application Profiles and Host-Conn.

27. Right-click Application EPGs and select Create Application EPG.

28. Name the EPG `iSCSI-LIF`. Leave Intra EPG Isolation Unenforced.

29. Use the Bridge Domain drop-down list to select Create Bridge Domain.

30. Name the Bridge Domain `BD-iSCSI`.

31. Select the FPV-App-A VRF.

32. Use the Forwarding drop-down list to select Custom.

33. Select Flood for the L2 Unknown Unicast and default for the End Point Retention Policy and IGMP Snoop Policy.

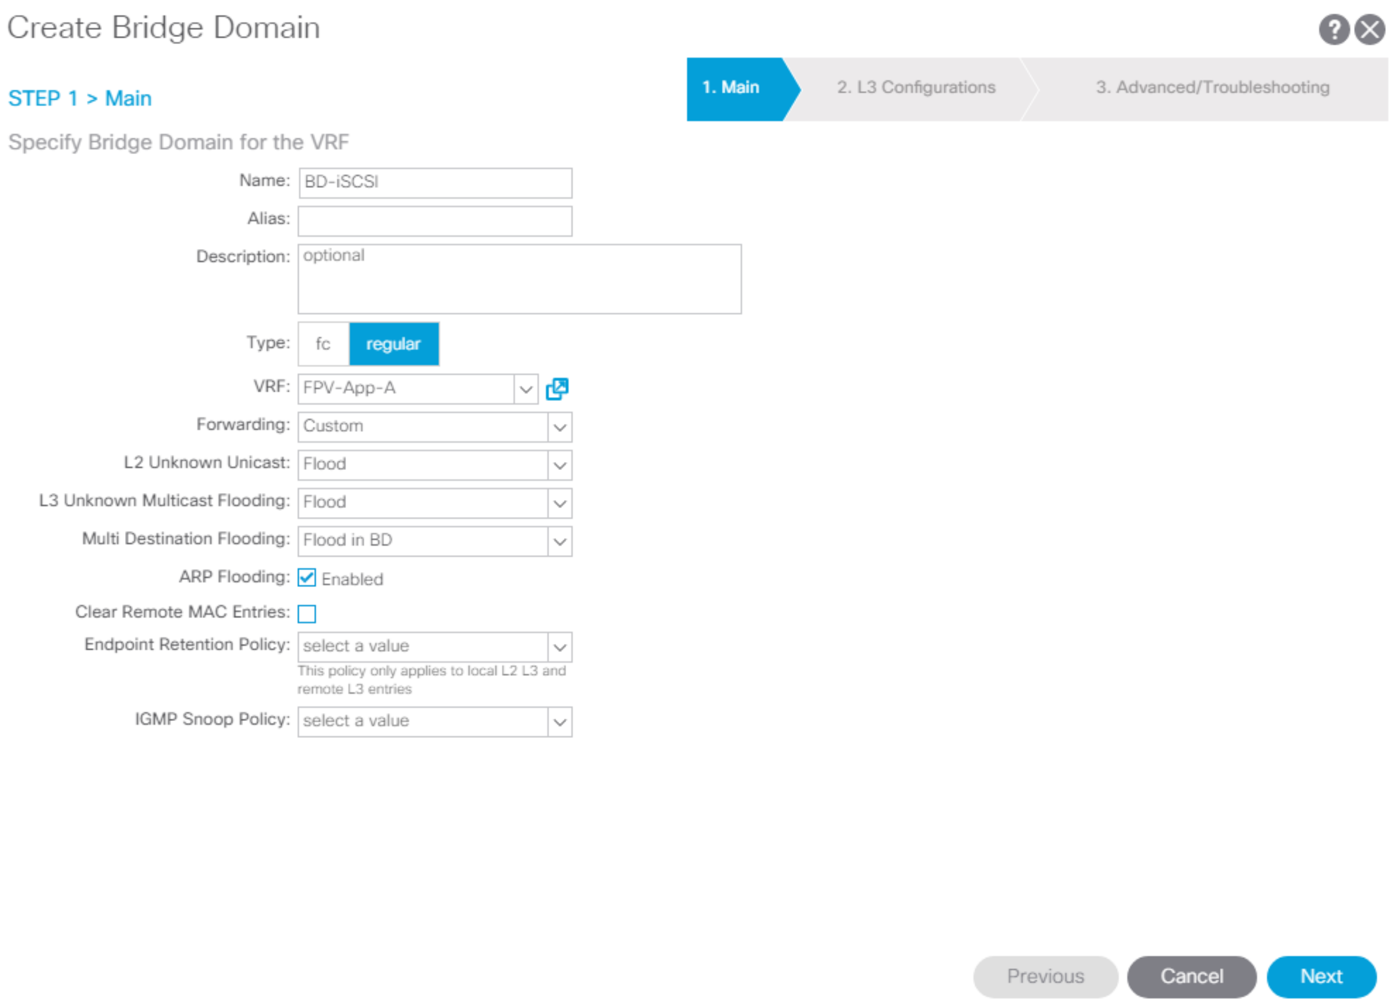

Create Bridge Domain

STEP 1 > Main

Specify Bridge Domain for the VRF

Name:	BD-iSCSI
Alias:	
Description:	optional
Type:	fc [✓ regular]
VRF:	FPV-App-A
Forwarding:	Custom
L2 Unknown Unicast:	Flood
L3 Unknown Multicast Flooding:	Flood
Multi Destination Flooding:	Flood in BD
ARP Flooding:	☑ Enabled
Clear Remote MAC Entries:	□
Endpoint Retention Policy:	select a value This policy only applies to local L2 L3 and remote L3 entries
IGMP Snoop Policy:	select a value

34. At the bottom right, click Next.

35. Under L3 Configurations, make sure Limit IP Learning to Subnet is selected and select EP Move Detection Mode – GARP based detection. Select **Next**.

36. Select the default Monitoring Policy and click Finish.

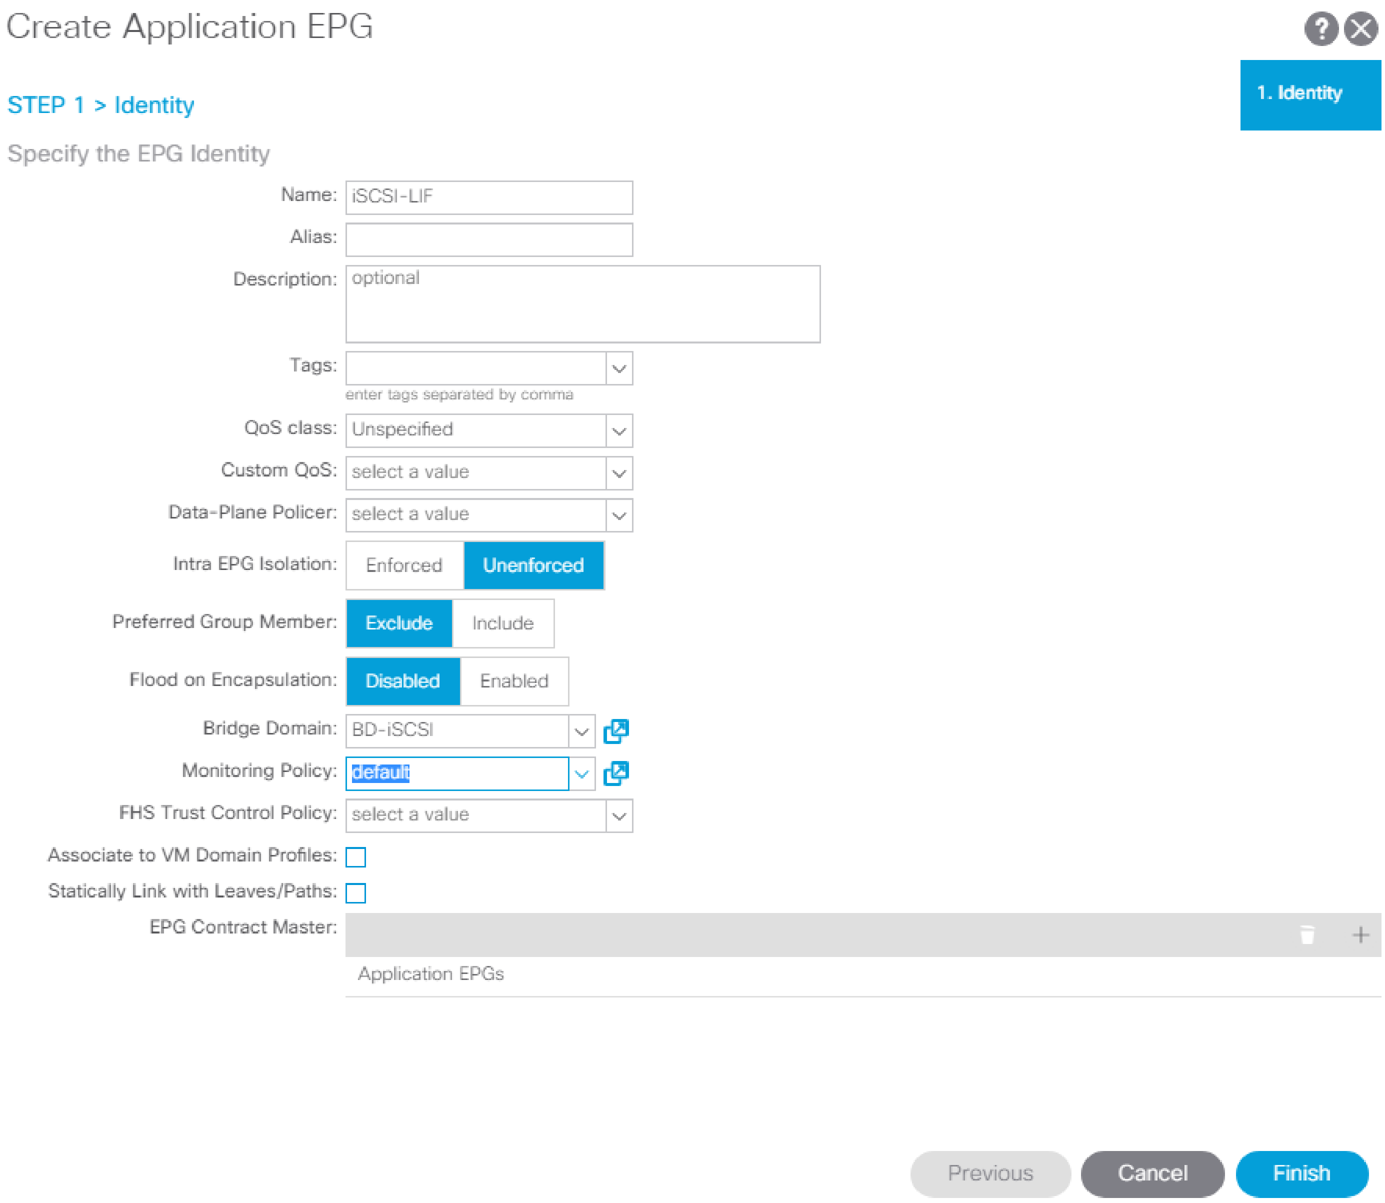

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: iSCSI-LIF
Alias:
Description: optional
Tags:
enter tags separated by comma
QoS class: Unspecified
Custom QoS: select a value
Data-Plane Policer: select a value
Intra EPG Isolation: Enforced [✓ Unenforced]
Preferred Group Member: [✓ Exclude] Include
Flood on Encapsulation: [✓ Disabled] Enabled
Bridge Domain: BD-iSCSI
Monitoring Policy: default
FHS Trust Control Policy: select a value
Associate to VM Domain Profiles: □
Statically Link with Leaves/Paths: □
EPG Contract Master:
Application EPGs

37. Select the default Monitoring Policy and click Finish to complete creating the EPG.

38. On the left, expand Application EPGs and EPG iSCSI-LIF. Right-click Domains and select Add Physical Domain Association.

39. Using the drop-down list, select the NetApp-AFF Physical Domain Profile.

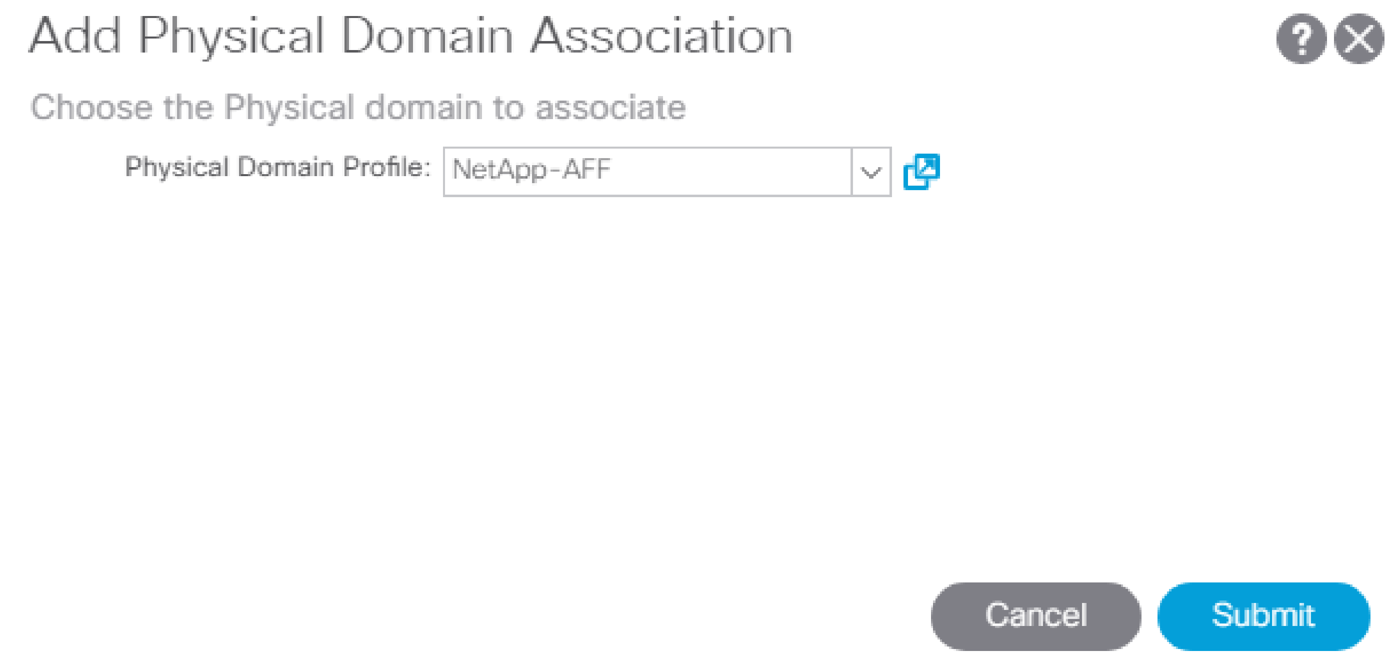

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: NetApp-AFF

40. Click Submit to complete the Physical Domain Association.

41. Right-click Static-Ports and select Deploy Static EPG on PC, VPC, or Interface.

42. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.

43. Using the Path drop-down list, select the VPC for NetApp Storage Controller 01.

44. Enter `<storage-FPV-App-A-iSCSI-VLAN>` for Port Encap.

45. Select the Immediate Deployment Immediacy and the Trunk Mode.

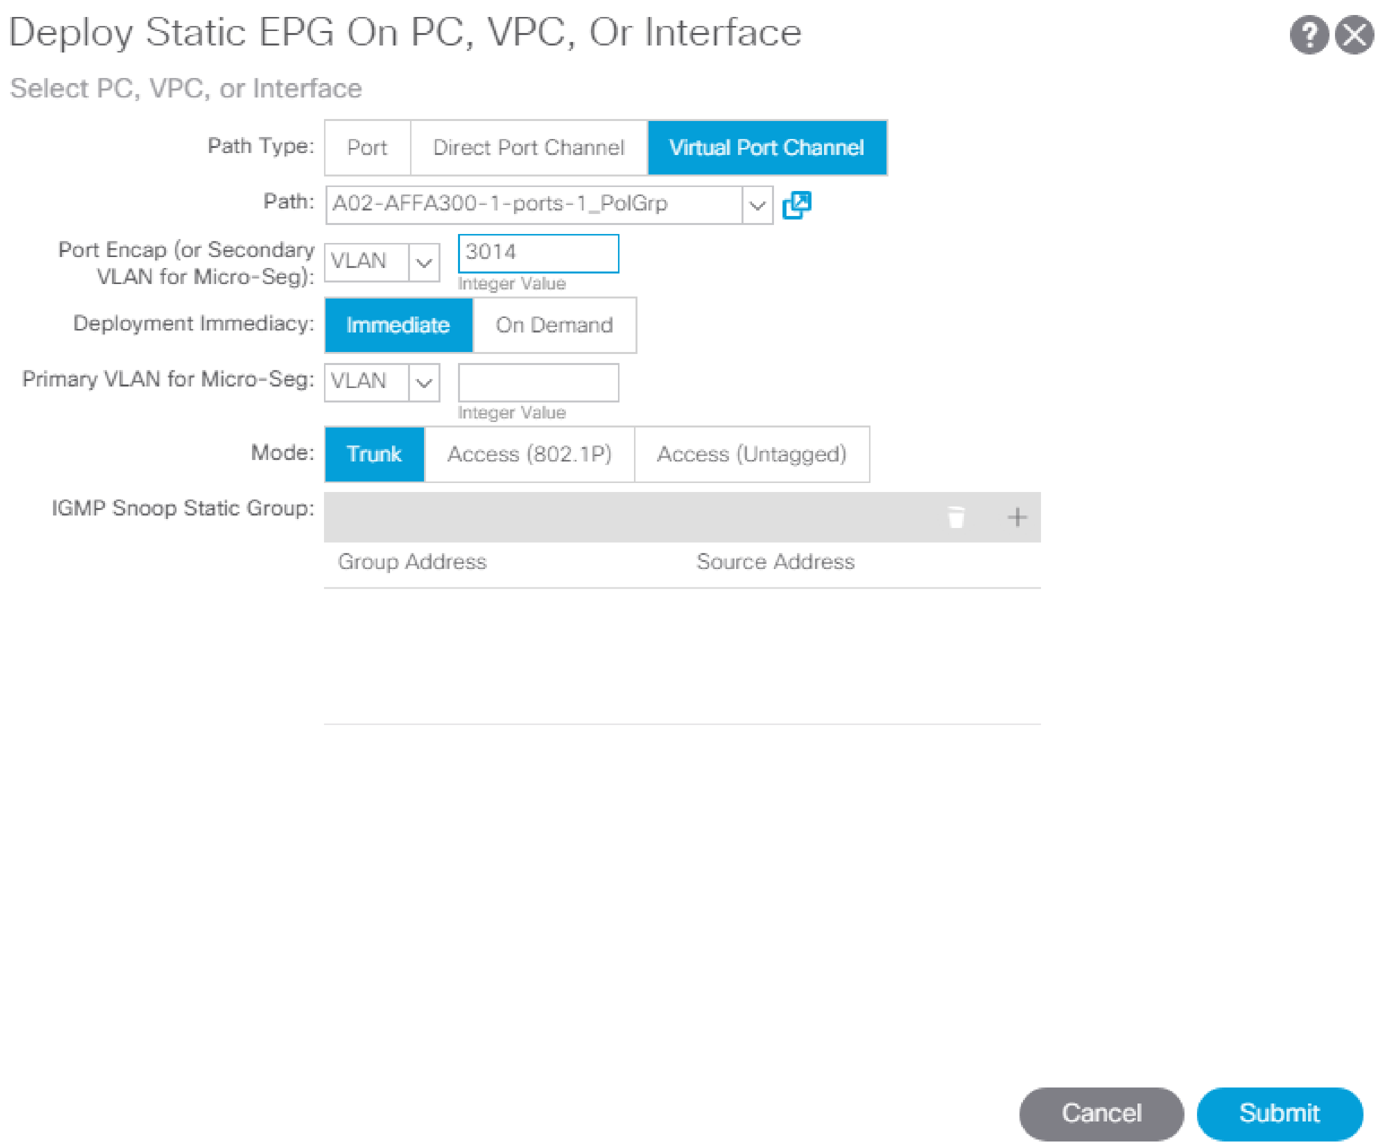

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port | Direct Port Channel | [✓ Virtual Port Channel]

Path: A02-AFFA300-1-ports-1_PolGrp

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 3014 (Integer Value)

Deployment Immediacy: [✓ Immediate] | On Demand

Primary VLAN for Micro-Seg: VLAN (Integer Value)

Mode: [✓ Trunk] | Access (802.1P) | Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

46. Click Submit to complete adding the Static Port Mapping.

47. Repeat steps 41-46 to add the Static Path Mapping for NetApp Storage Controller 02.

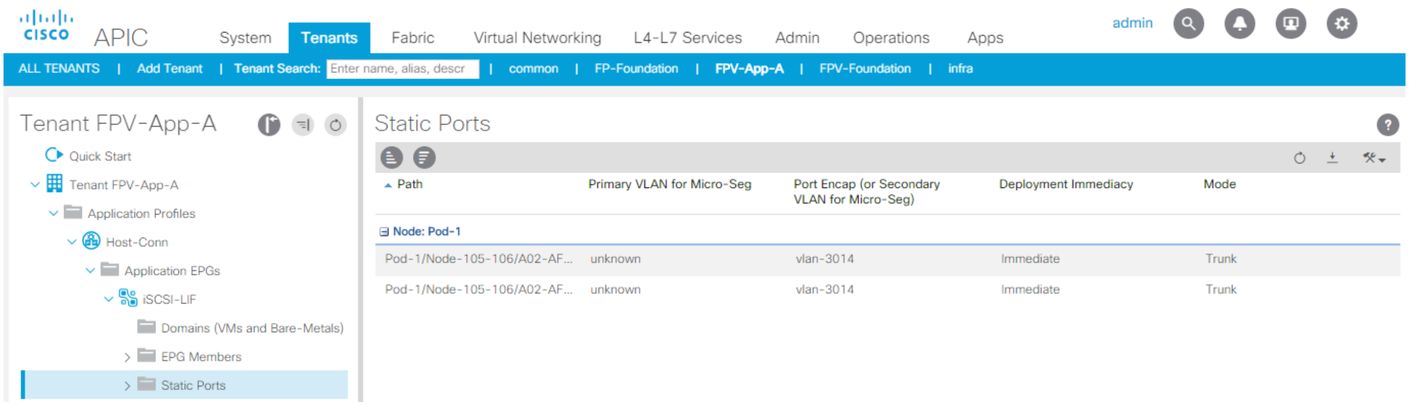

Static Ports

Path	Primary VLAN for Micro-Seg	Port Encap (or Secondary VLAN for Micro-Seg)	Deployment Immediacy	Mode
Node: Pod-1				
Pod-1/Node-105-106/A02-AF...	unknown	vlan-3014	Immediate	Trunk
Pod-1/Node-105-106/A02-AF...	unknown	vlan-3014	Immediate	Trunk

48. Under EPG iSCSI-LIF, right-click Contracts and select Add Provided Contract.

49. Use the drop-down list to select Create Contract.

50. Name the contract Allow-iSCSI. Set the Scope of the contract to Tenant. Click the “+” to add a Subject.

51. Name the Subject Allow-iSCSI. Click the “+” in the Filter Chain to add a filter.

52. Using the drop-down list, select the iSCSI filter from Tenant common. Click Update.

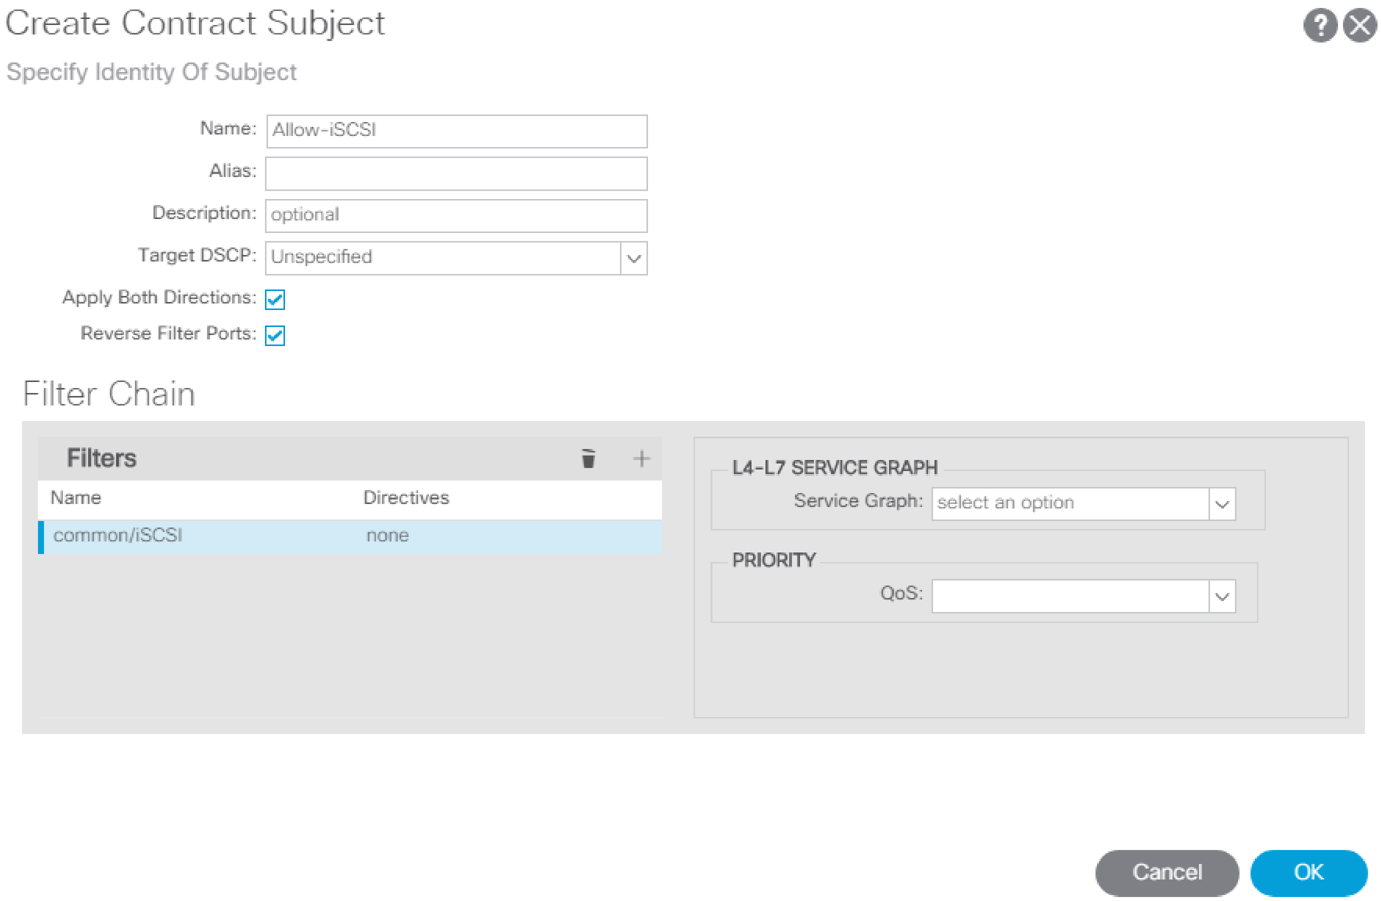

Create Contract Subject

Specify Identity Of Subject

Name: Allow-iSCSI

Alias:

Description: optional

Target DSCP: Unspecified

Apply Both Directions: ☑

Reverse Filter Ports: ☑

Filter Chain

Filters

Name	Directives
common/iSCSI	none

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

Optionally, add ICMP to the filter chain to allow ping in this contract for troubleshooting purposes.

53. Click OK to complete creating the Contract Subject.

54. Click Submit to complete creating the Contract.

55. Click Submit to complete adding the Provided Contract.

56. Right-click Application EPGs and select Create Application EPG.

57. Name the EPG `iSCSI-VMK-VM`. Leave Intra EPG Isolation Unenforced.

58. Use the Bridge Domain drop-down list to select the tenant `BD-iSCSI` Bridge Domain.

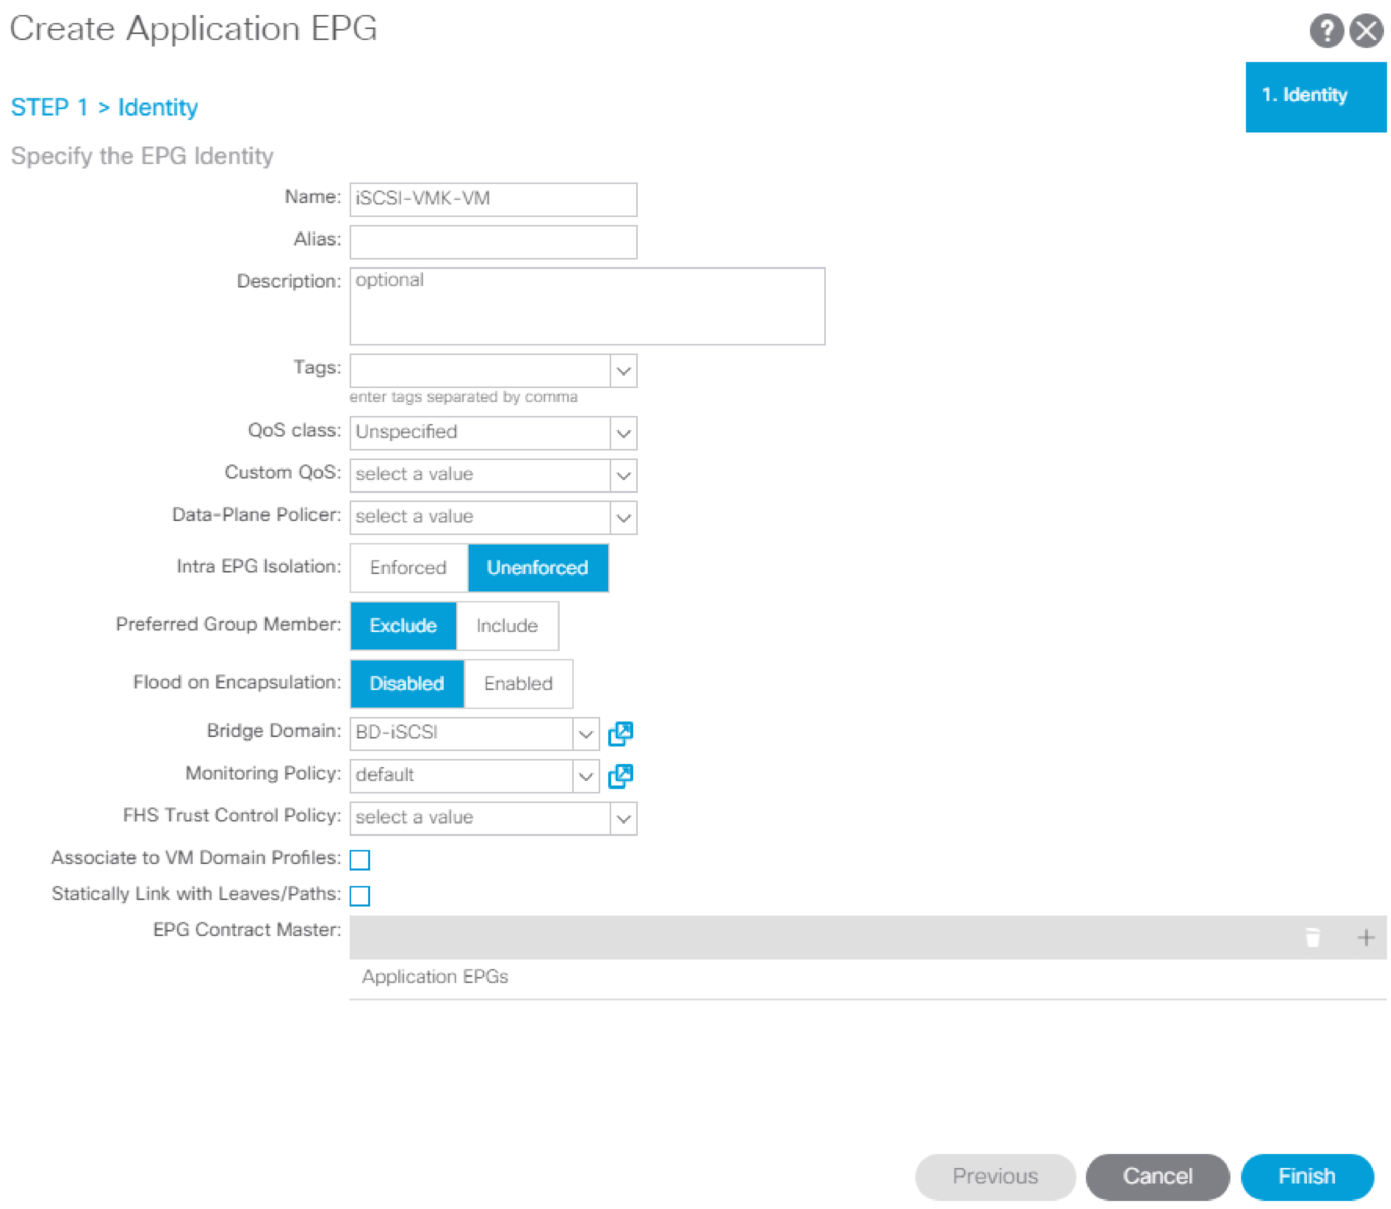

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: iSCSI-VMK-VM

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Flood on Encapsulation: [✓ Disabled] Enabled

Bridge Domain: BD-iSCSI

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

59. Select the default Monitoring Policy and click Finish to complete creating the EPG.

60. On the left, expand Application EPGs and EPG iSCSI-VMK-VM. Right-click Domains and select Add VMM Domain Association.

61. From the drop-down list, select fpv-vc-vDS. Set Deploy Immediacy to Immediate.

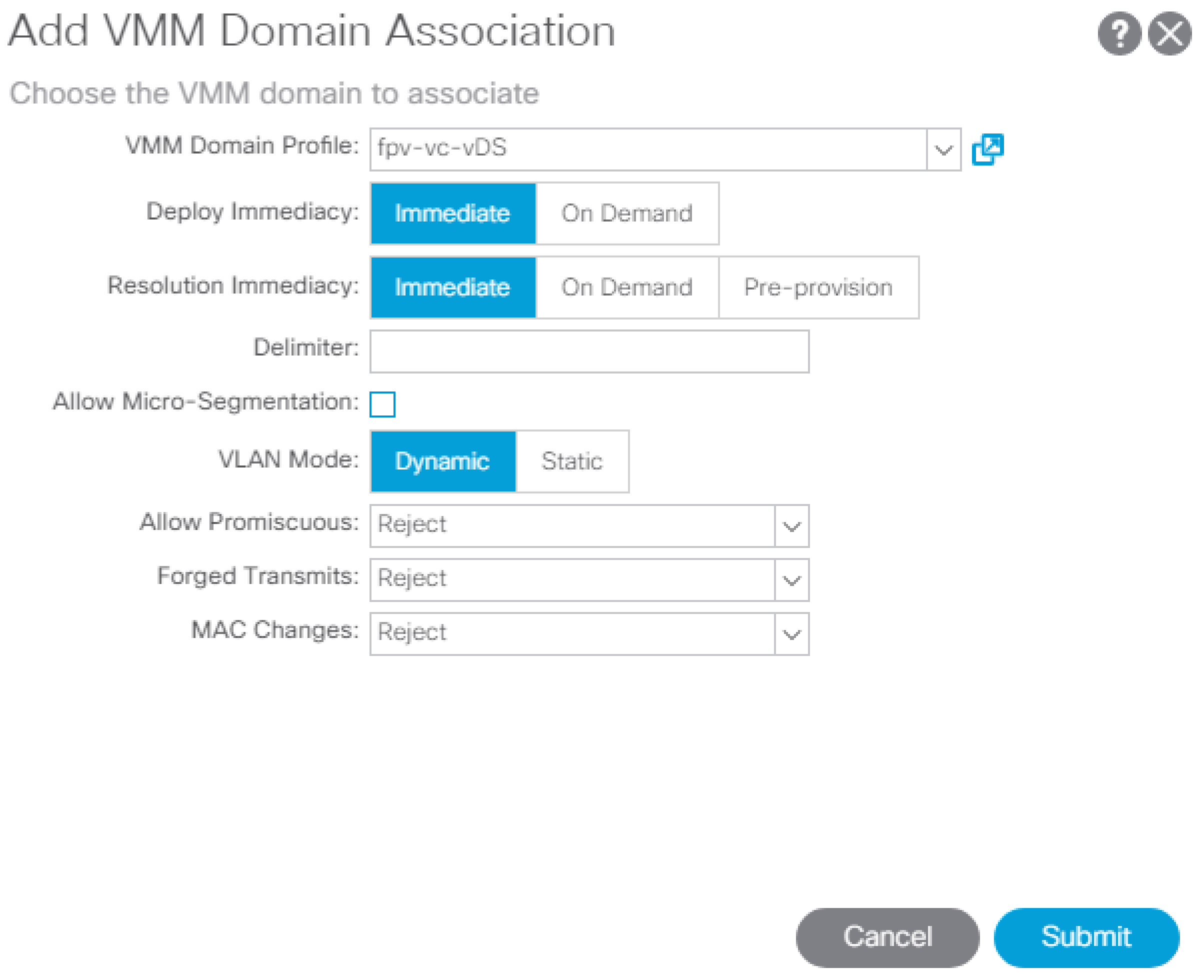

Add VMM Domain Association

Choose the VMM domain to associate

VMM Domain Profile: fpv-vc-vDS

Deploy Immediacy: [✓ Immediate] On Demand

Resolution Immediacy: [✓ Immediate] On Demand Pre-provision

Delimiter:

Allow Micro-Segmentation: □

VLAN Mode: [✓ Dynamic] Static

Allow Promiscuous: Reject

Forged Transmits: Reject

MAC Changes: Reject

62. Click Submit to complete the VMM Domain Association.

63. Under EPG iSCSI-VMK-VM, right-click Contracts and select Add Consumed Contract.

64. Use the drop-down list to select Allow-iSCSI within the tenant.

65. Click Submit to complete adding the Consumed Contract.

In this deployment for iSCSI, you are putting the LIFs in one EPG and the Host/VM Interfaces in a second EPG and connecting them with a filtered contract. You also could have added both the NetApp LIF endpoints and the Host/VM Interface endpoints in a single EPG. This method would have allowed unrestricted communication within the EPG.

66. On the left, under Tenant FPV-App-A, right-click the Host-Conn Application Profile and select Create Application EPG.

67. Name the EPG `NFS-LIF` and leave Intra EPG Isolation set at Unenforced.

68. Use the Bridge Domain drop-down list to select Create Bridge Domain.

69. Name the Bridge Domain `BD-NFS` and select the FPV-App-A VRF.

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.

70. For Forwarding, select Custom and select Flood for L2 Unknown Unicast. Select default for the End Point Retention Policy and the IGMP Snoop Policy.

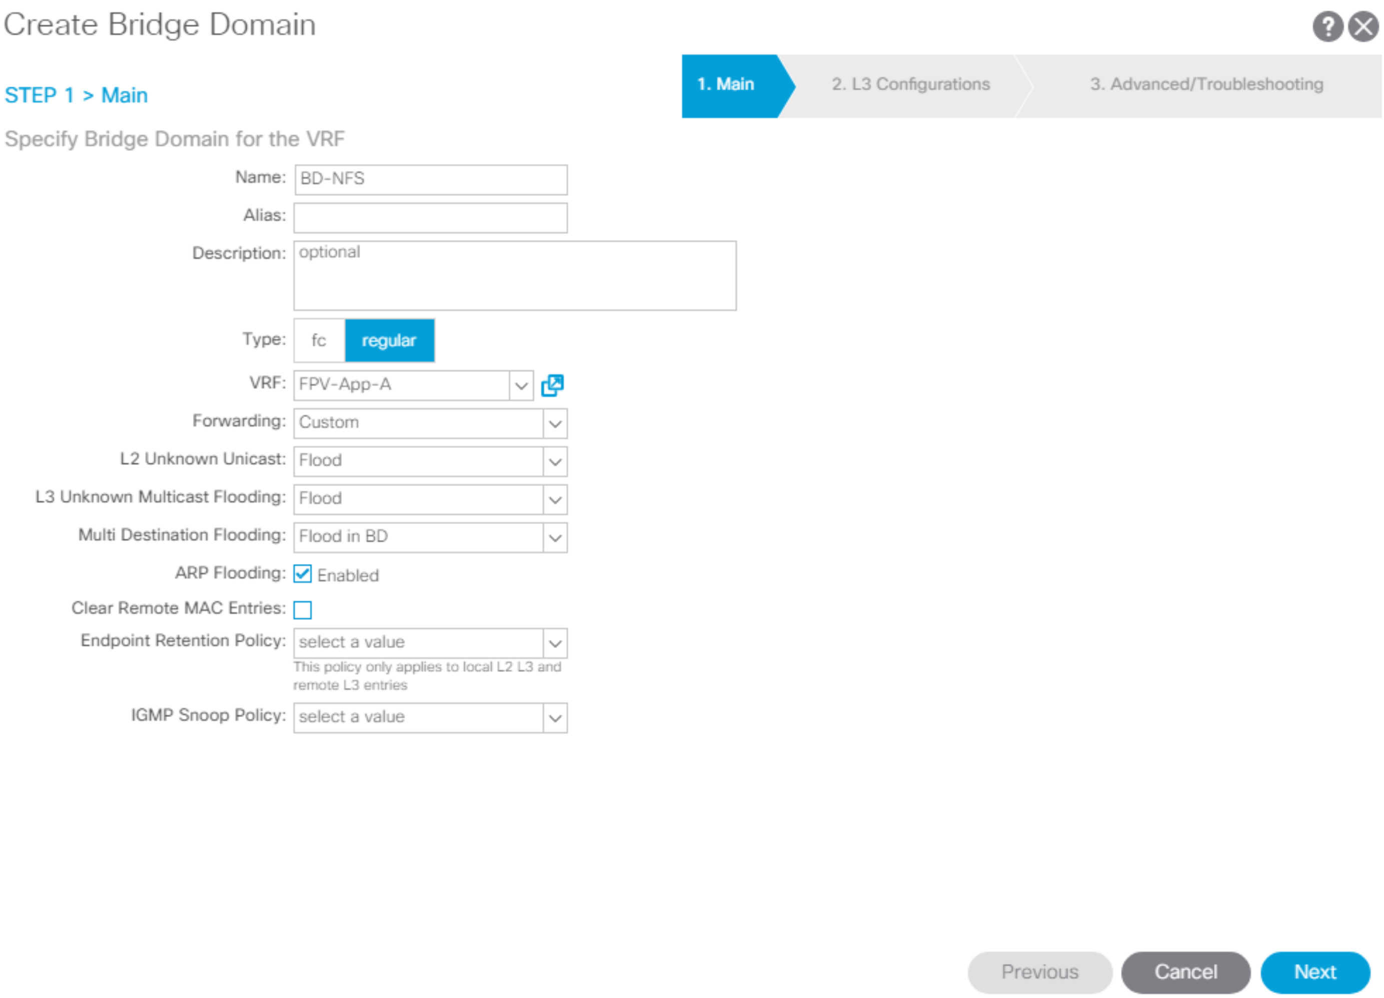

Create Bridge Domain

STEP 1 > Main

Specify Bridge Domain for the VRF

Name: BD-NFS

Alias:

Description: optional

Type: fc [✓ regular]

VRF: FPV-App-A

Forwarding: Custom

L2 Unknown Unicast: Flood

L3 Unknown Multicast Flooding: Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: ☑ Enabled

Clear Remote MAC Entries: □

Endpoint Retention Policy: select a value

This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: select a value

71. At the bottom right, click Next.

72. Under L3 Configurations, make sure Limit IP Learning to Subnet is selected and select EP Move Detection Mode – GARP based detection. Select **Next**.

73. Select the default Monitoring Policy and click Finish.

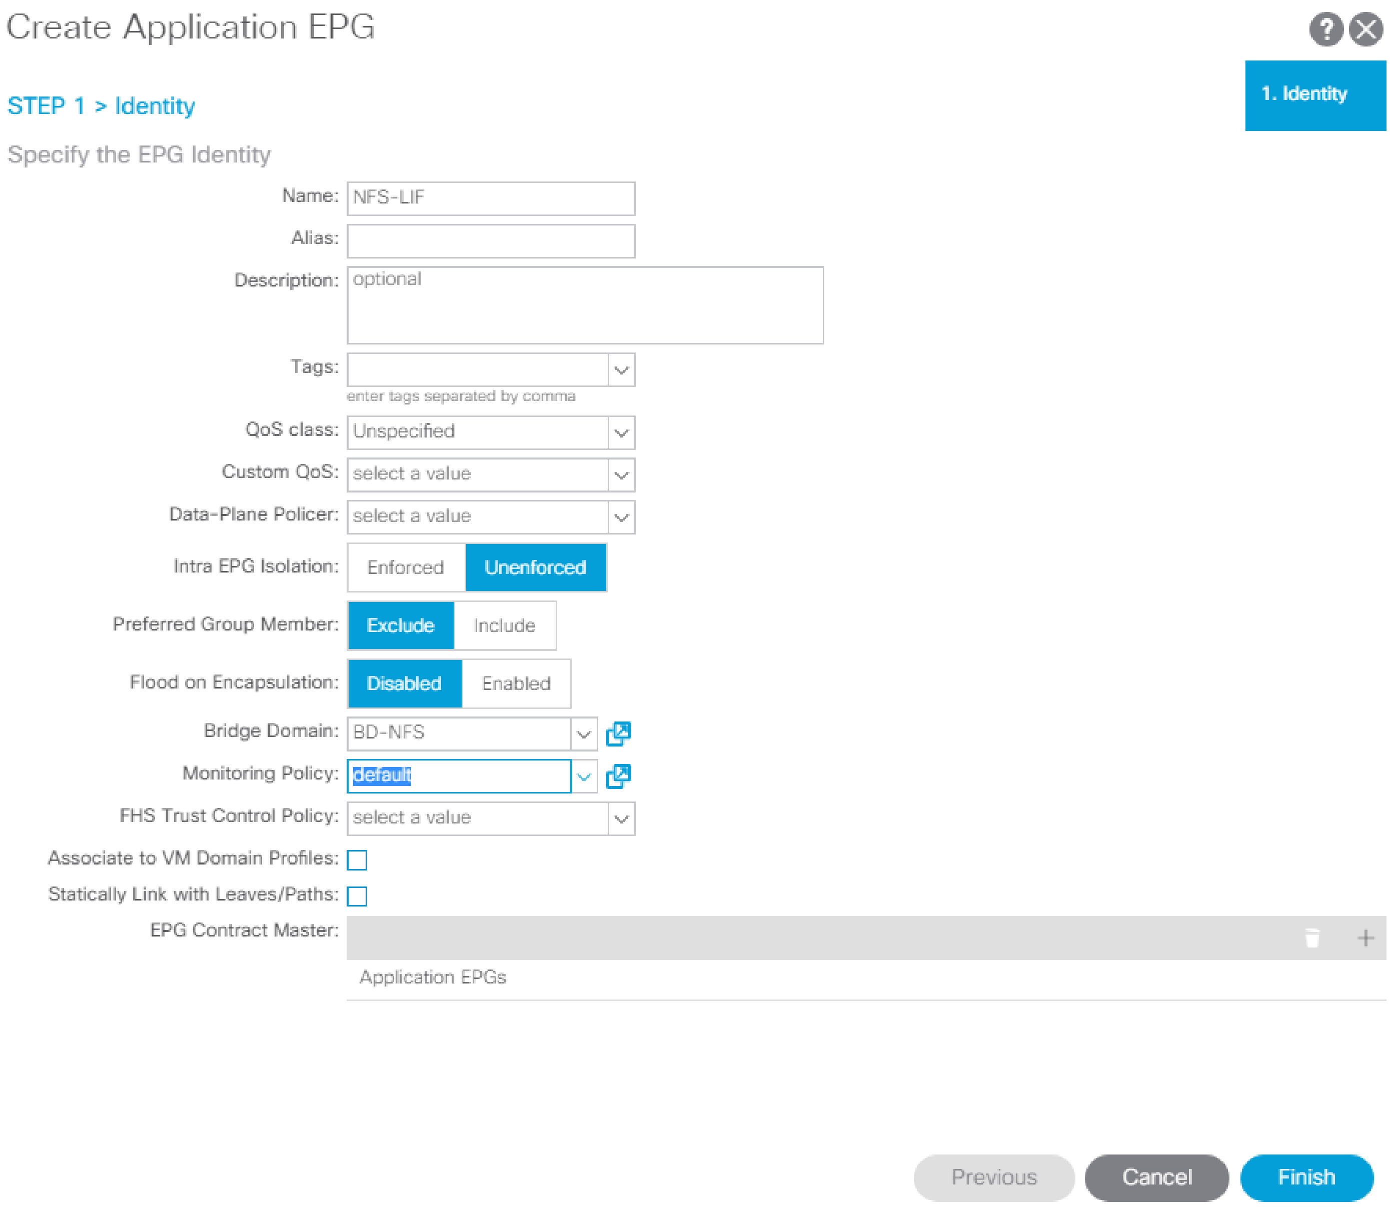

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: NFS-LIF

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Flood on Encapsulation: [✓ Disabled] Enabled

Bridge Domain: BD-NFS

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

74. Select the default Monitoring Policy and click Finish to complete creating the EPG.

75. On the left expand Host-Conn, Application EPGs, and EPG NFS-LIF.

76. Right-click Domains and select Add Physical Domain Association.

77. Select the NetApp-AFF Physical Domain Profile.

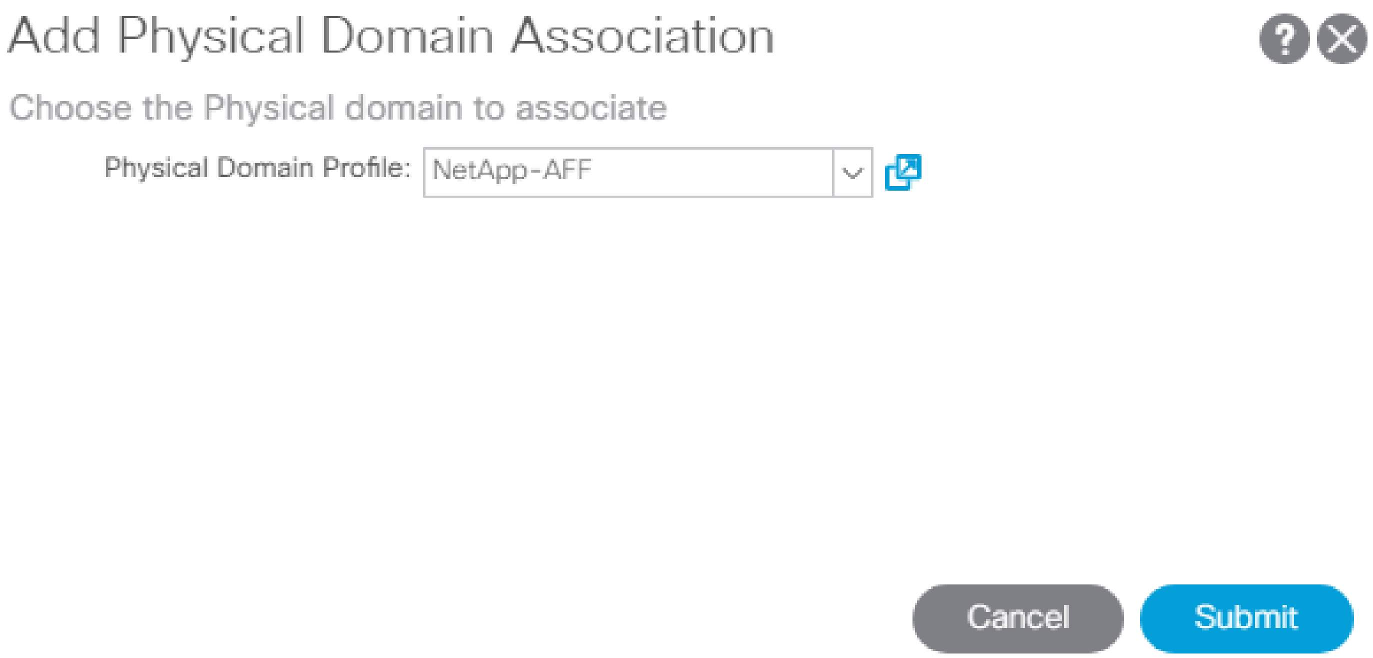

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: NetApp-AFF

78. Click Submit to compete adding the Physical Domain Association.

79. Right-click Static Ports and select Deploy Static EPG on PC, VPC, or Interface.

80. Select the Virtual Port Channel Path Type.

81. Using the Path drop-down list, select the VPC for NetApp Storage Controller 01.

82. For Port Encap, enter `<storage-FPV-App-A-NFS-VLAN>`.

83. Select Immediate for Deployment Immediacy and Trunk for Mode.

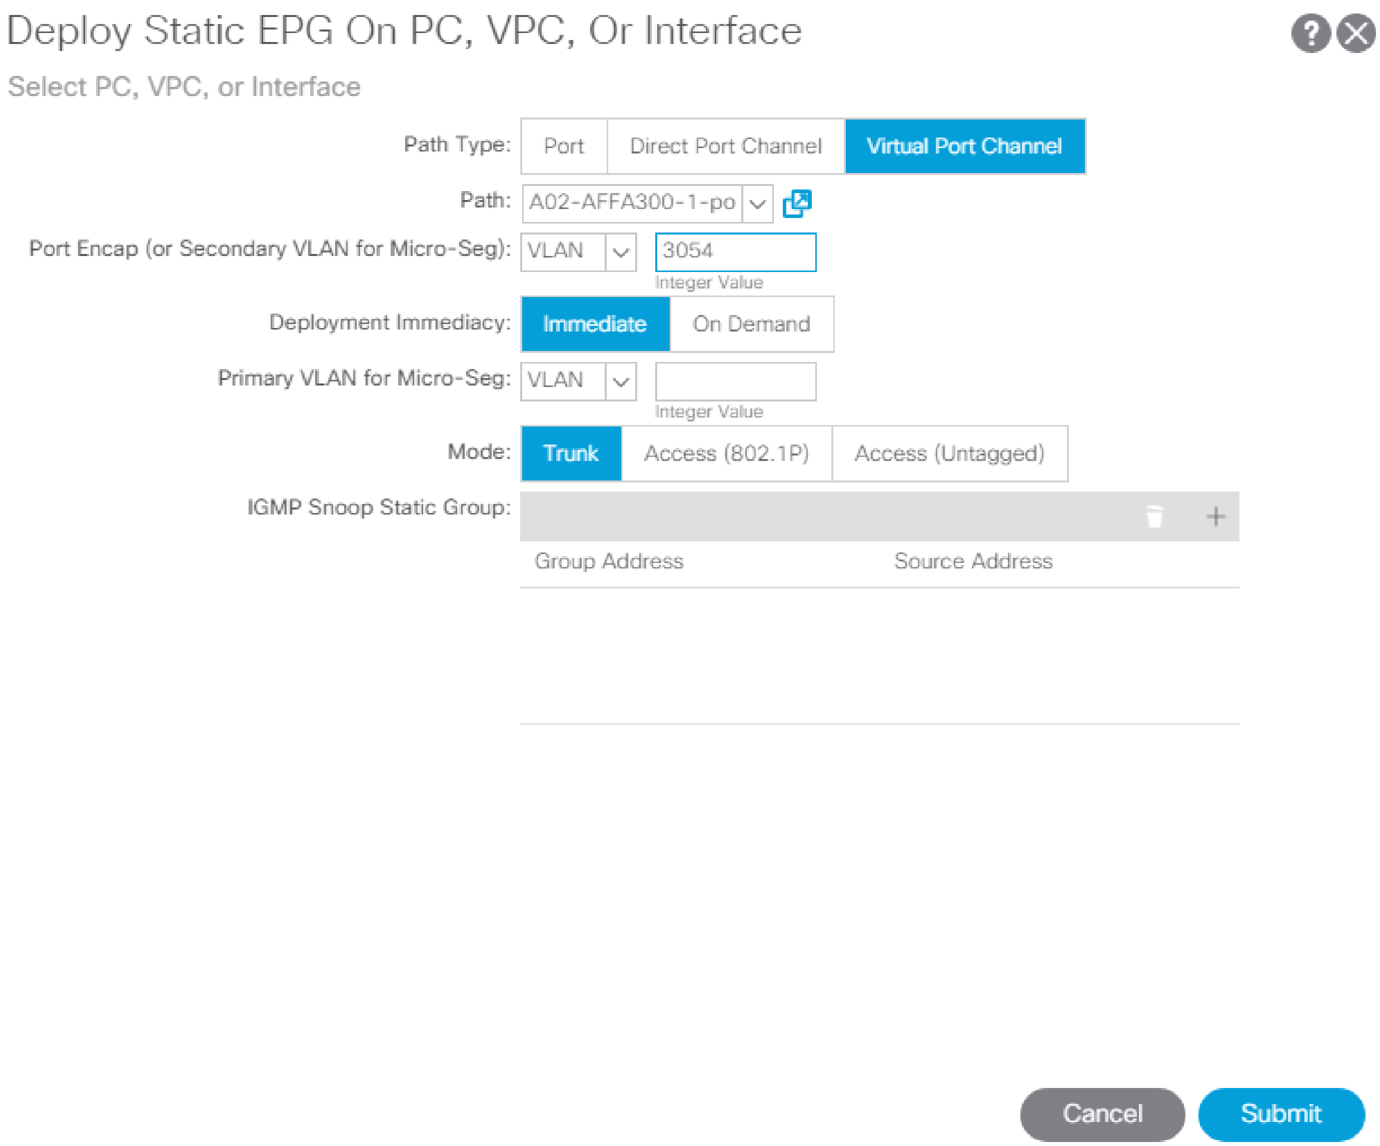

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port | Direct Port Channel | [✓ Virtual Port Channel]

Path: A02-AFFA300-1-po

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 3054

Deployment Immediacy: [✓ Immediate] | On Demand

Primary VLAN for Micro-Seg: VLAN

Mode: [✓ Trunk] | Access (802.1P) | Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

84. Click Submit to finish adding the EPG Static Binding.

85. Repeat steps 79-84 for the Static Port for NetApp Storage Controller 02.

86. On the left under EPG NFS-LIF, right-click Contracts and select Add Provided Contract.

87. In the Add Provided Contract window, use the Contract drop-down list to select Create Contract.

88. Name the contract `Allow-NFS`. Set the Scope to Tenant.

89. Click the “+” sign to add a Contract Subject.

90. Name the subject `Allow-NFS`.

91. Click the “+” sign to add a Filter to the Filter Chain.

92. Click the drop-down list and select NTAP-NFS-v3 from Tenant common.

93. Click Update.

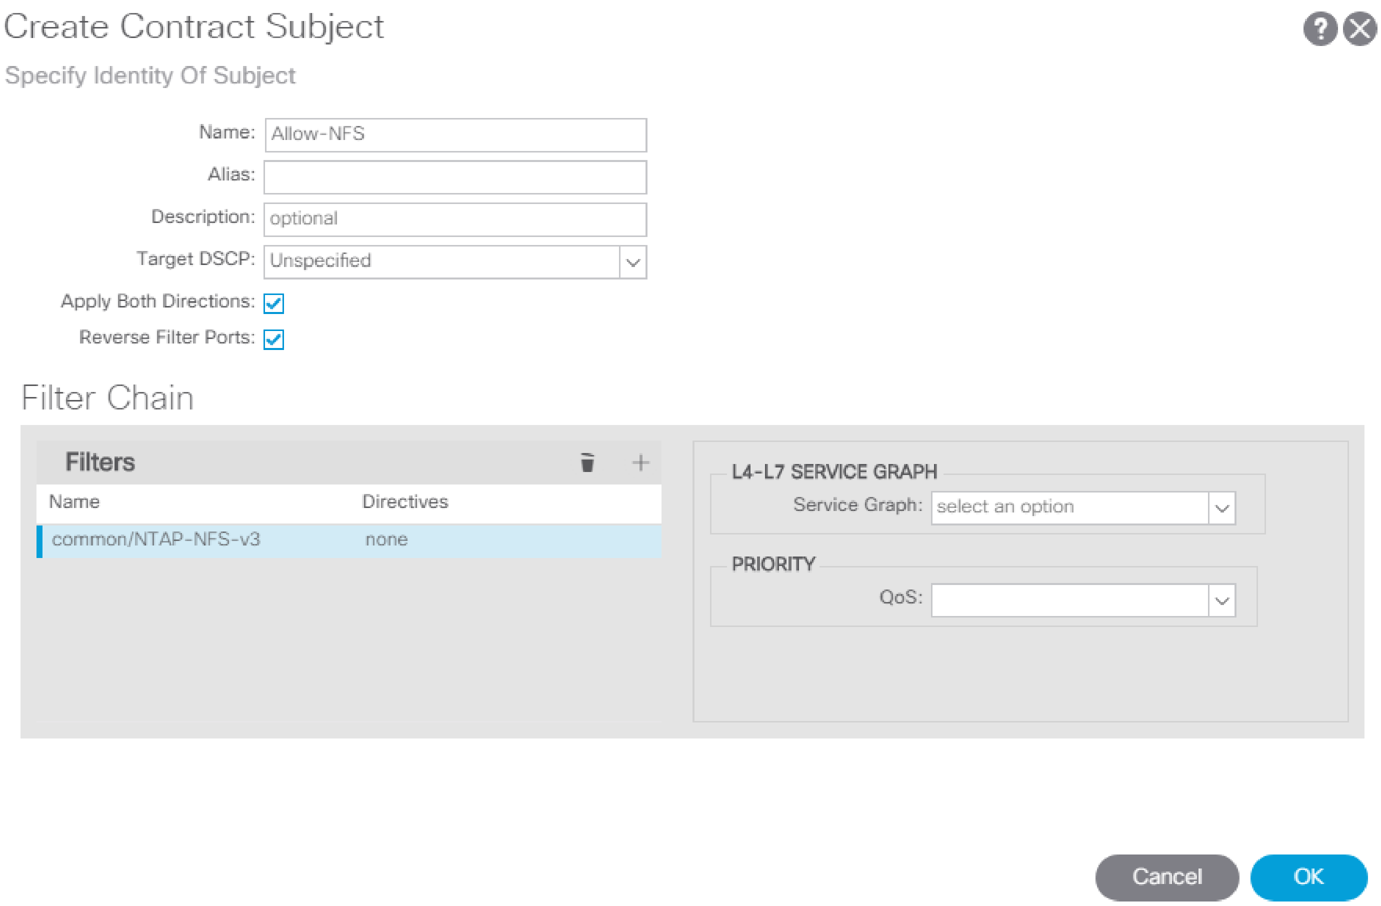

Create Contract Subject

Specify Identity Of Subject

Name: Allow-NFS

Alias:

Description: optional

Target DSCP: Unspecified

Apply Both Directions: ☑

Reverse Filter Ports: ☑

Filter Chain

Filters

Name	Directives
common/NTAP-NFS-v3	none

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

Optionally, add ICMP to the filter chain to allow ping in this contract for troubleshooting purposes.

94. Click OK to complete the Contract Subject.

95. Click Submit to complete creating the Contract.

96. Click Submit to complete Adding the Provided Contract.

97. Right-click Application EPGs under the Host-Conn Application Profile and select Create Application EPG.

98. Name the EPG `NFS-VMK-VM` and leave Intra EPG Isolation set at Unenforced.

99. Use the Bridge Domain drop-down list to select BD-NFS in the same tenant. Select the default Monitoring Policy.

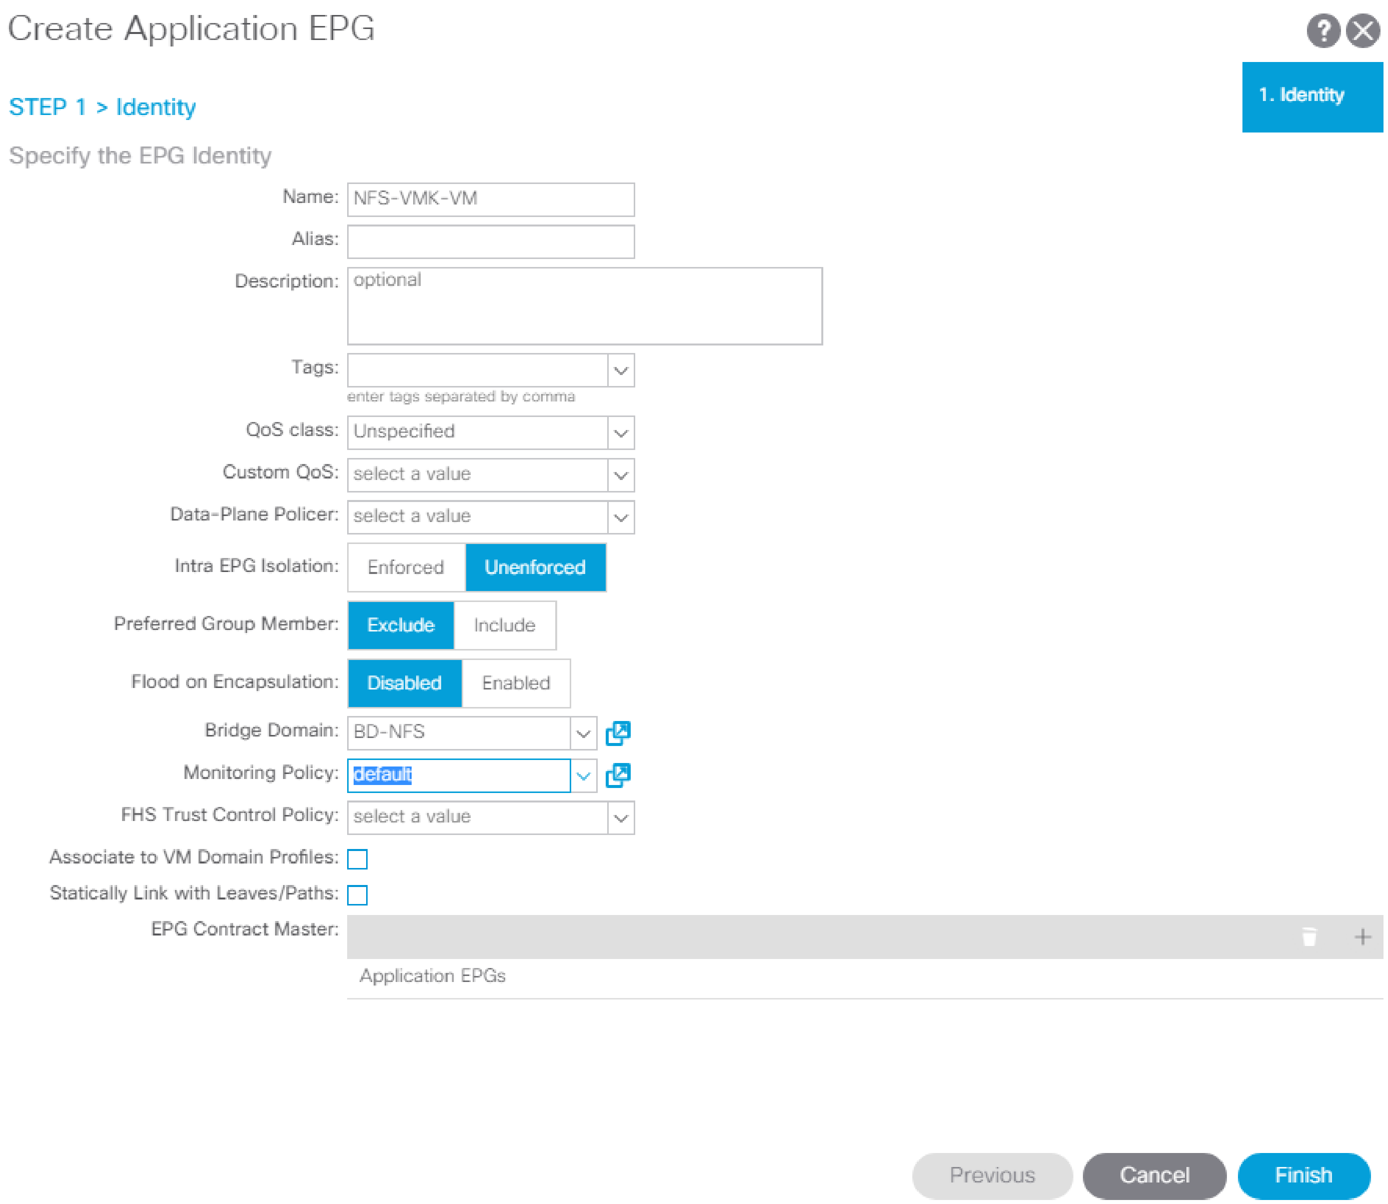

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: NFS-VMK-VM

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Flood on Encapsulation: [✓ Disabled] Enabled

Bridge Domain: BD-NFS

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

100. Click Finish to complete creating the EPG.

101. On the left expand Host-Conn, Application EPGs, and EPG NFS-VMK-VM.

102. Under EPG NFS-VMK-VM, right-click Domains and select Add VMM Domain Association.

103. Select the fpv-vc-vDS VMM Domain Profile.

104. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

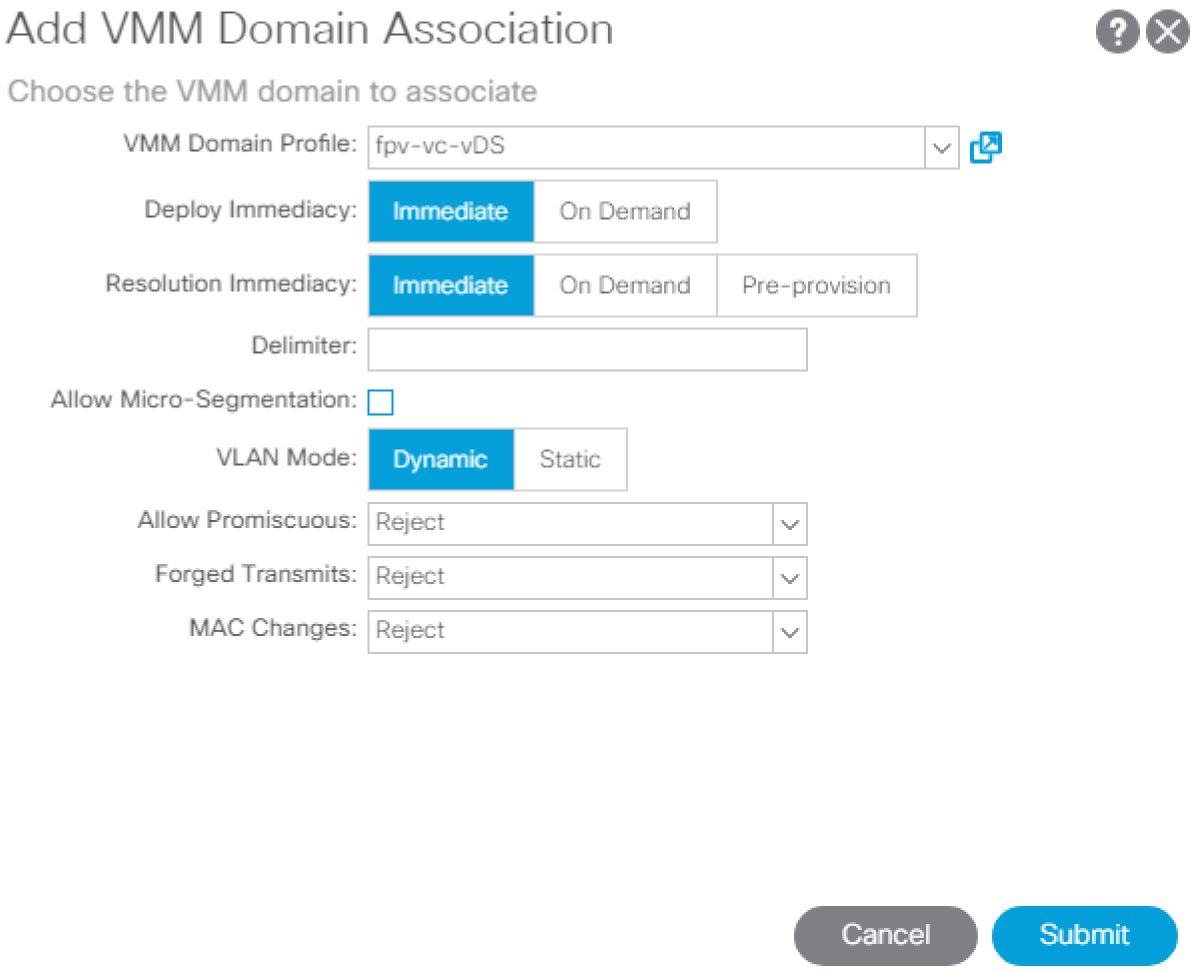

Add VMM Domain Association

Choose the VMM domain to associate

VMM Domain Profile: fpv-vc-vDS

Deploy Immediacy: [✓ Immediate] On Demand

Resolution Immediacy: [✓ Immediate] On Demand Pre-provision

Delimiter:

Allow Micro-Segmentation: □

VLAN Mode: [✓ Dynamic] Static

Allow Promiscuous: Reject

Forged Transmits: Reject

MAC Changes: Reject

105. Click Submit to complete adding the VMM Domain Association.

106. On the left under EPG NFS-VMK-VM, right-click Contracts and select Add Consumed Contract.

107. In the Add Consumed Contract window, from the Contract drop-down list, select Allow-NFS in the FPV-App-A Tenant.

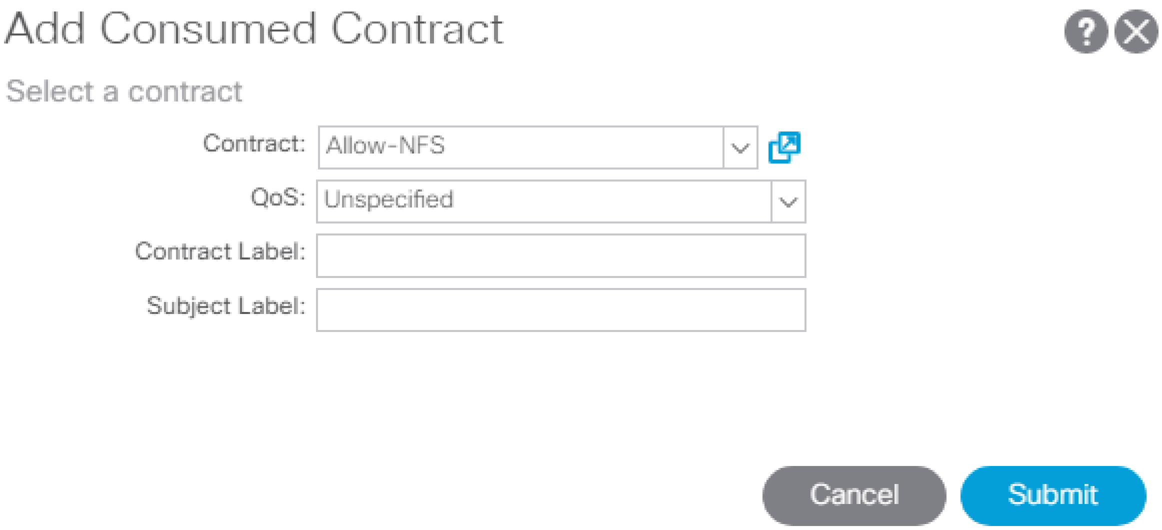

Add Consumed Contract

Select a contract

Contract: Allow-NFS

QoS: Unspecified

Contract Label:

Subject Label:

108. Click Submit to complete adding the Consumed Contract.

109. On the left, under Tenant FPV-App-A, right-click Application Profiles and select Create Application Profile.

110. Name the Profile `MGMT`, select the default Monitoring Policy, and click SUBMIT.

111. Right-click the MGMT Application Profile and select Create Application EPG.

112. Name the EPG `SVM–MGMT` and leave Intra EPG Isolation set at Unenforced.

113. Use the Bridge Domain drop-down list to select create Bridge Domain.

114. Name the Bridge Domain BD-Internal and select the FPV-App-A VRF.

115. Leave Forwarding set at Optimize, select the default End Point Retention Policy and IGMP Snoop Policy.

116. Click Next.

117. Make sure Unicast Routing is enabled and click Next.

118. Select the default Monitoring Policy and click Finish to complete creating the Bridge Domain.

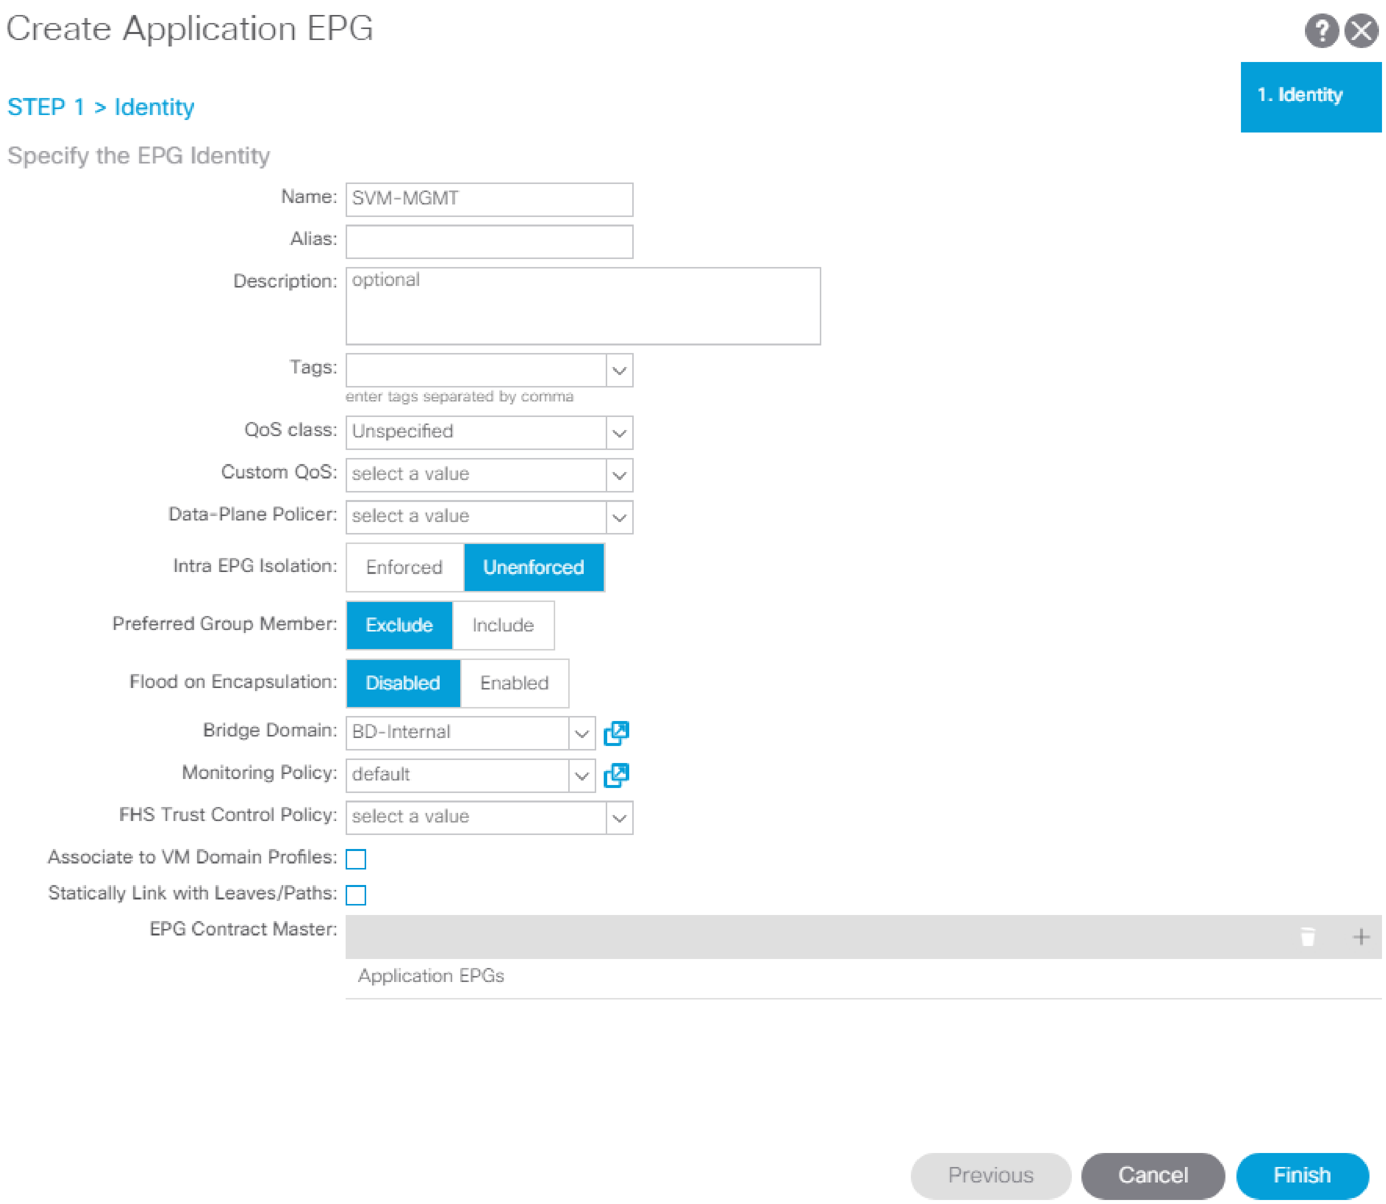

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: SVM-MGMT

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Flood on Encapsulation: [✓ Disabled] Enabled

Bridge Domain: BD-Internal

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

119. Select the default Monitoring Policy and click Finish to complete creating the EPG.

120. On the left expand MGMT, Application EPGs, and EPG SVM-MGMT.

121. Right-click Domains and select Add Physical Domain Association.

122. Select the NetApp-AFF Physical Domain Profile.

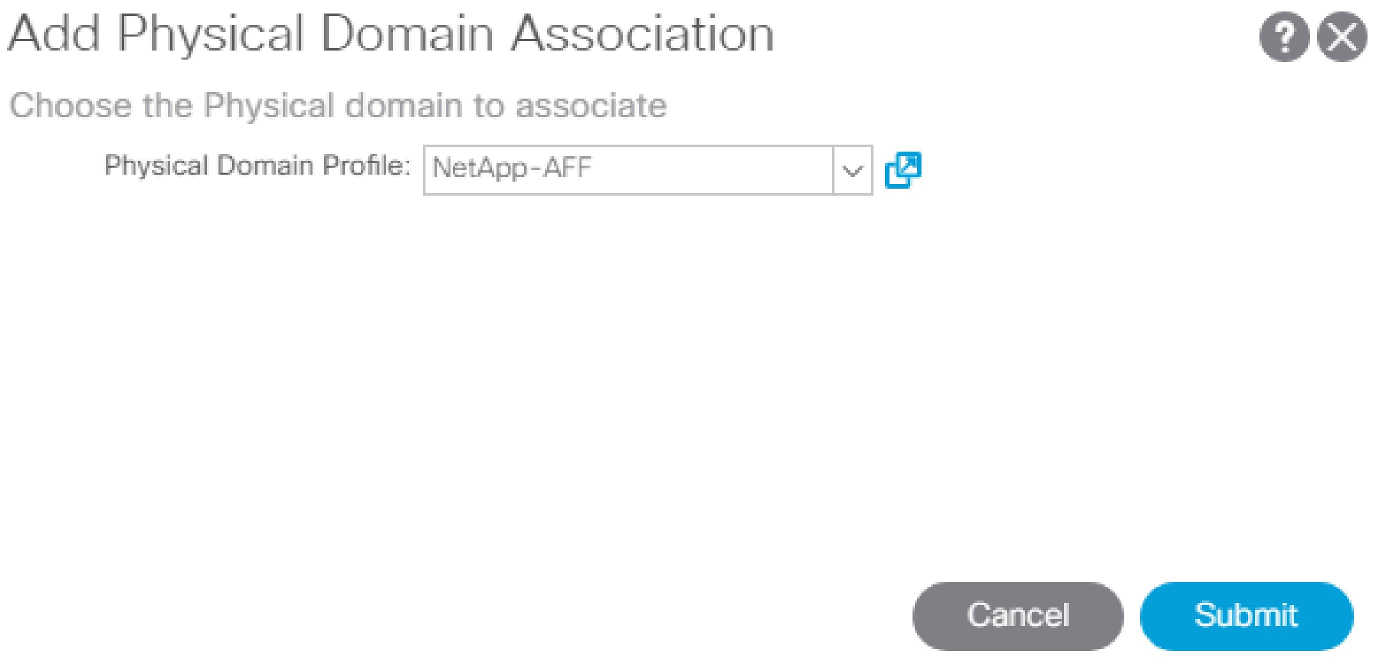

Add Physical Domain Association

Choose the Physical domain to associate

Physical Domain Profile: NetApp-AFF

123. Click Submit to compete adding the Physical Domain Association.

124. Right-click Static Ports and select Deploy Static EPG on PC, VPC, or Interface.

125. Select the Virtual Port Channel Path Type.

126. Using the Path drop-down list, select the VPC for NetApp Storage Controller 01.

127. For Port Encap, enter `<storage-FPV-App-A-SVM-MGMT-VLAN>`.

128. Select Immediate for Deployment Immediacy and Trunk for Mode.

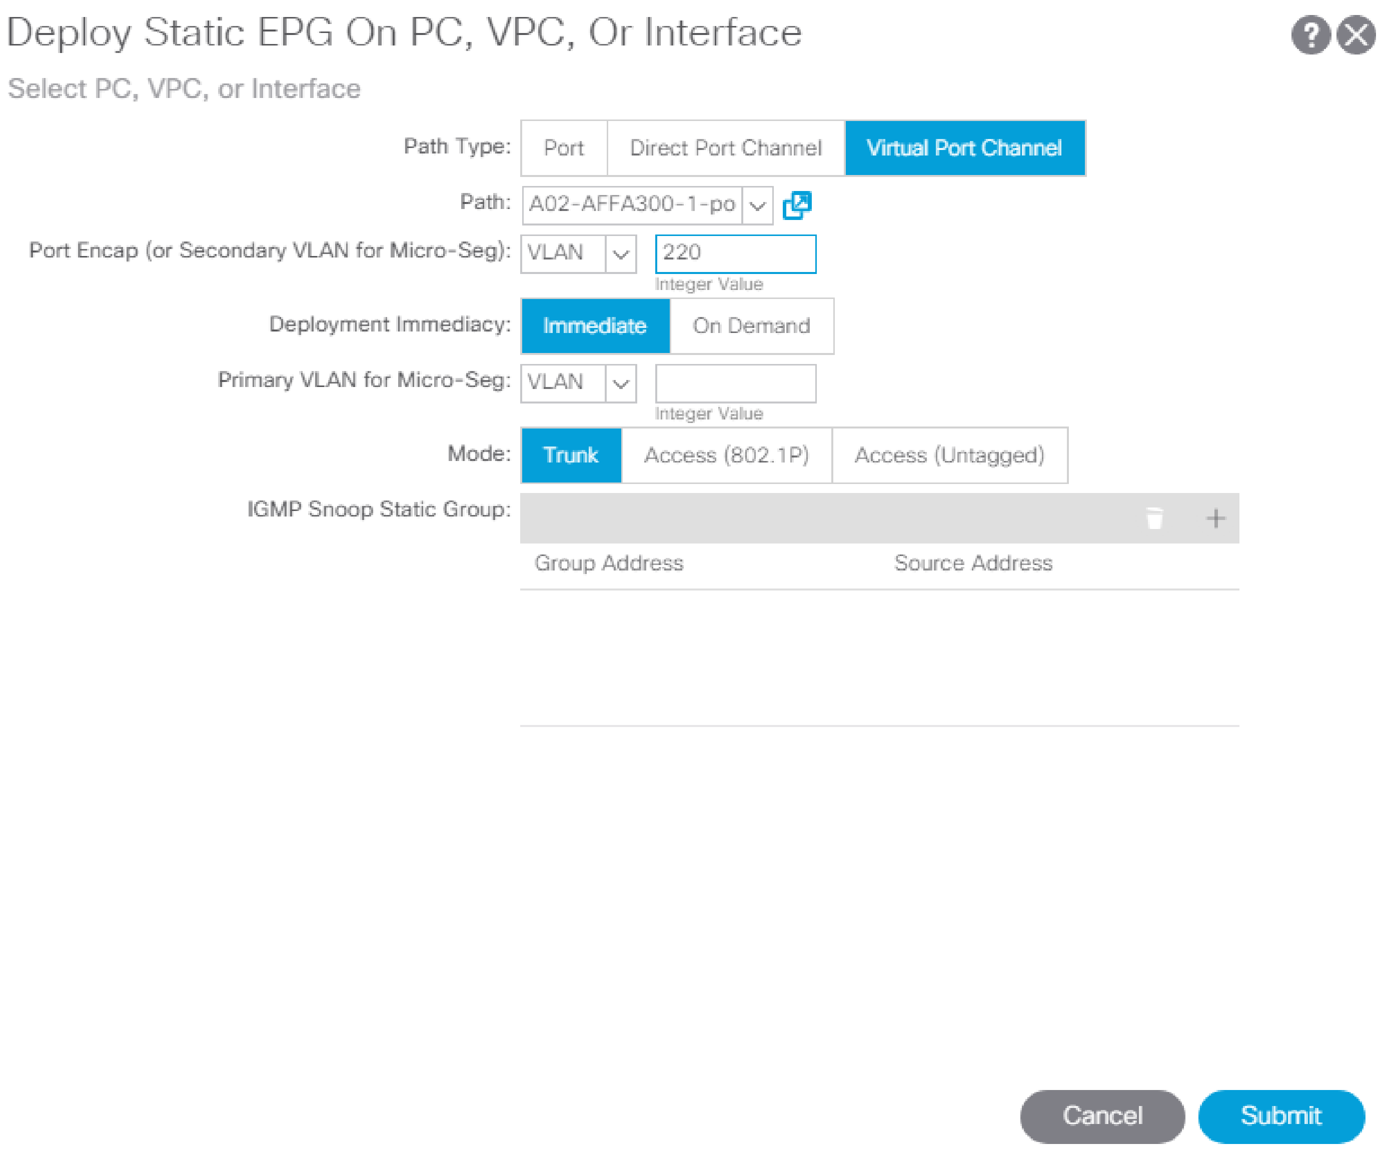

Deploy Static EPG On PC, VPC, Or Interface

Select PC, VPC, or Interface

Path Type: Port | Direct Port Channel | [✓ Virtual Port Channel]

Path: A02-AFFA300-1-po

Port Encap (or Secondary VLAN for Micro-Seg): VLAN 220

Deployment Immediacy: [✓ Immediate] | On Demand

Primary VLAN for Micro-Seg: VLAN

Mode: [✓ Trunk] | Access (802.1P) | Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

129. Click Submit to finish adding the EPG Static Binding.

130. Repeat steps 124-129 for the Static Port Mapping to NetApp Storage Controller 02.

131. On the left under EPG SVM-MGMT, right-click Contracts and select Add Provided Contract.

132. In the Add Provided Contract window, use the Contract drop-down list to select Create Contract.

133. Name the contract `Allow-SVM-MGMT`. Set the Scope set to Tenant.

134. Click the “+” sign to add a Contract Subject.

135. Name the subject `Allow-All`. Click the “+” sign to add a filter.

136. Use the drop-down list to select the Allow-All filter from Tenant common. Click Update.

137. Click OK to complete adding the Contract Subject.

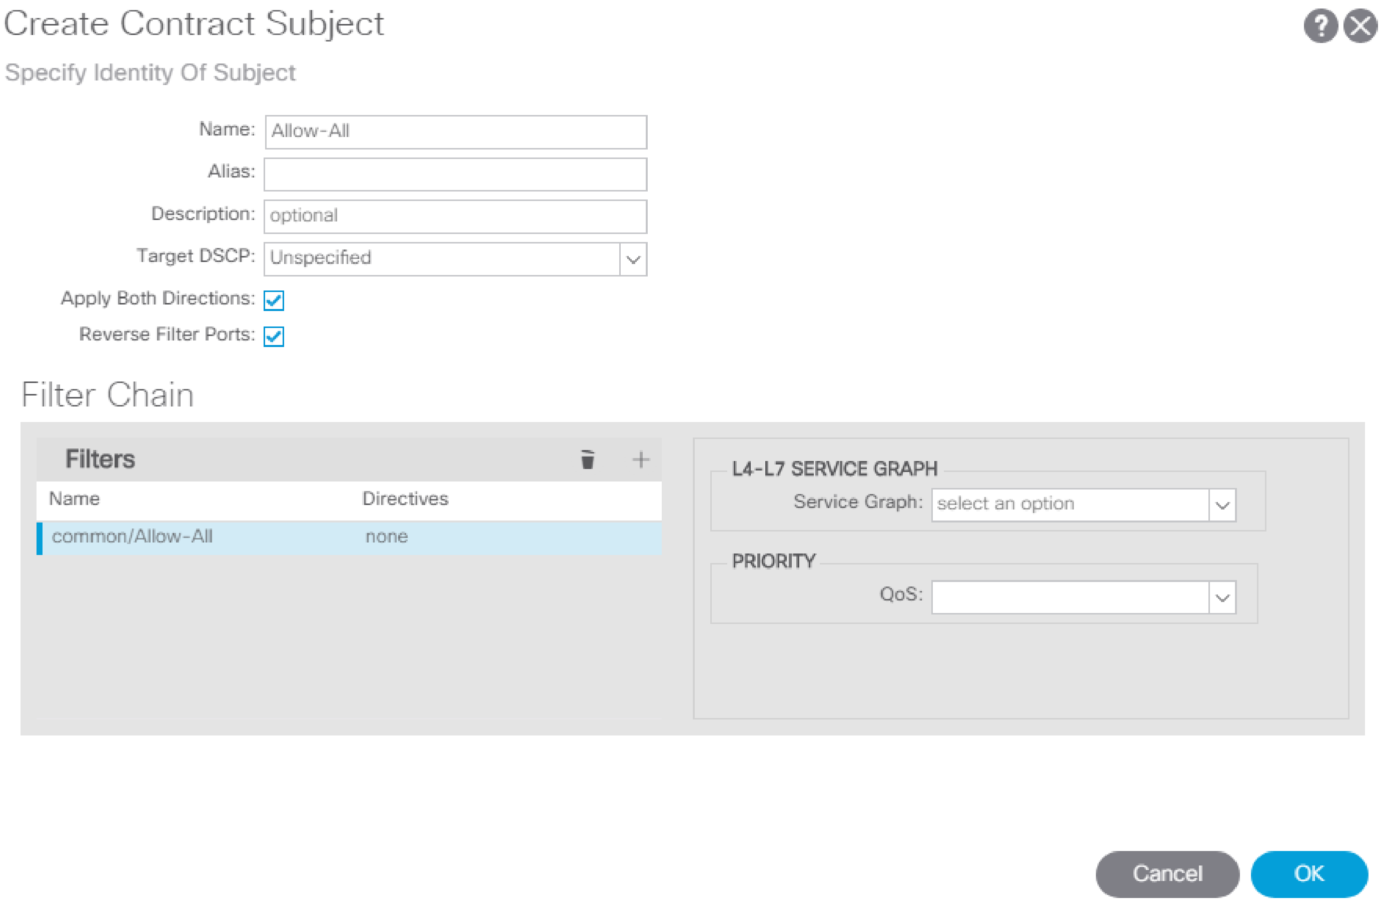

Create Contract Subject

Specify Identity Of Subject

Name: Allow-All

Alias:

Description: optional

Target DSCP: Unspecified

Apply Both Directions: ☑

Reverse Filter Ports: ☑

Filter Chain

Filters

Name	Directives
common/Allow-All	none

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

138. Click Submit to complete creating the Contract.

139. Click Submit to complete adding the Provided Contract.

140. On the left under EPG SVM-MGMT, right-click Subnets and select Create EPG Subnet.

141. For the Default Gateway IP, enter the SVM gateway IP address and mask for the FPV-App-A tenant.

142. Select only the Shared between VRFs scope.

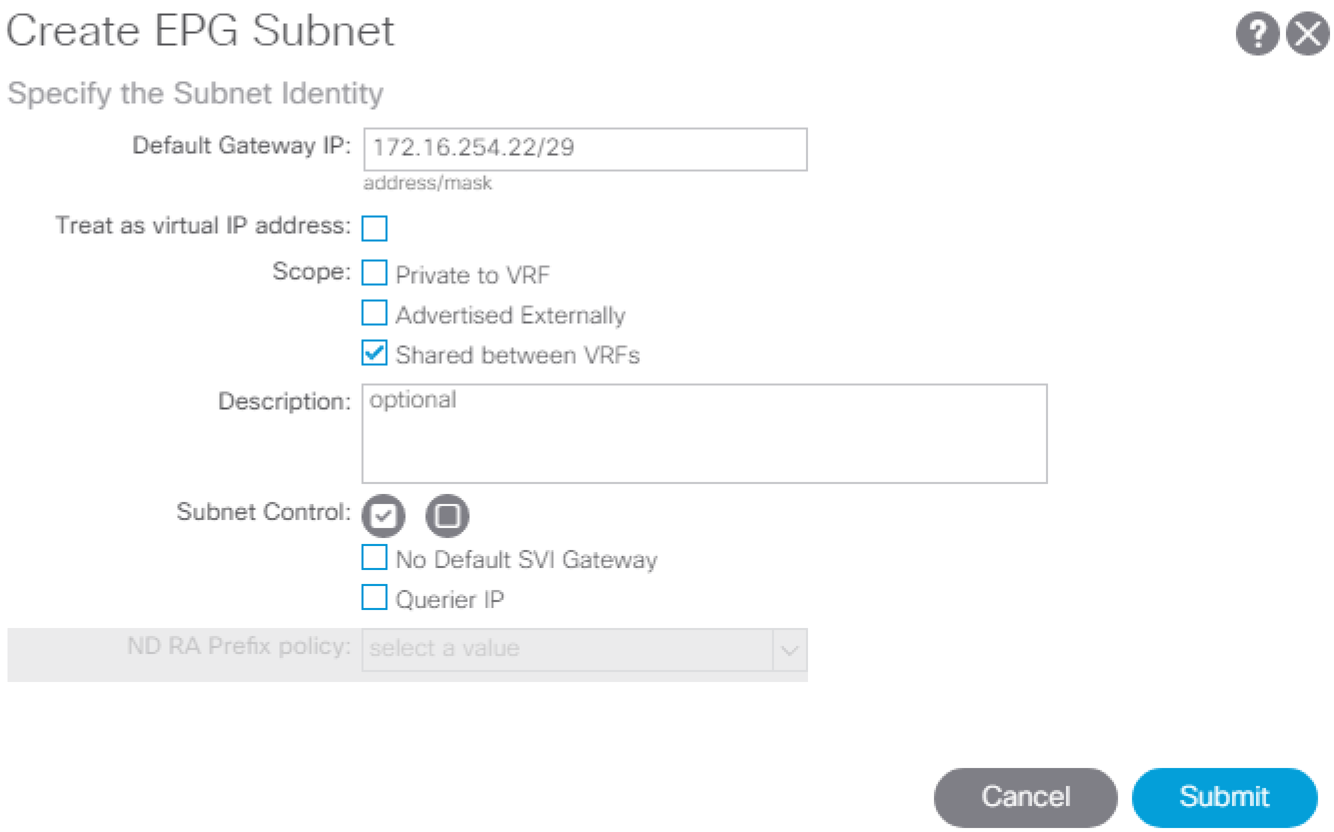

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP: 172.16.254.22/29
address/mask

Treat as virtual IP address: □

Scope: □ Private to VRF
□ Advertised Externally
☑ Shared between VRFs

Description: optional

Subnet Control: □ No Default SVI Gateway
□ Querier IP

ND RA Prefix policy: select a value

143. Click Submit to complete adding the EPG subnet.

144. On the left under EPG SVM-MGMT, right-click Contracts and select Add Consumed Contract.

145. In the Add Consumed Contract window, use the Contract drop-down list to select common-Allow -Core-Services in Tenant common.

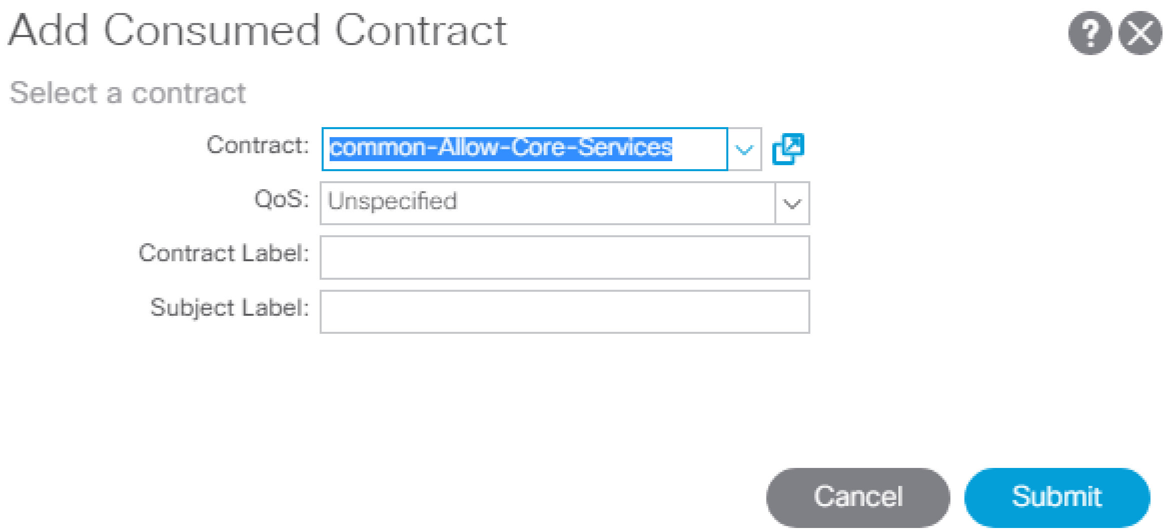

Add Consumed Contract

Select a contract

Contract: common-Allow-Core-Services

QoS: Unspecified

Contract Label:

Subject Label:

146. Click Submit to complete adding the Consumed Contract.

147. On the left, right-click Application Profiles and select Create Application Profile.

148. Name the Application Profile `3-Tier-App` and select the default Monitoring Policy.

149. Click Submit to complete creating the Application Profile.

150. Expand 3-Tier-App, right-click Application EPGs under 3-Tier-App and select Create Application EPG.

151. Name the EPG `Web` and leave Intra EPG Isolation set at Unenforced.

152. Use the Bridge Domain drop-down list to select BD-Internal in the current tenant. Select the default Monitoring Policy.

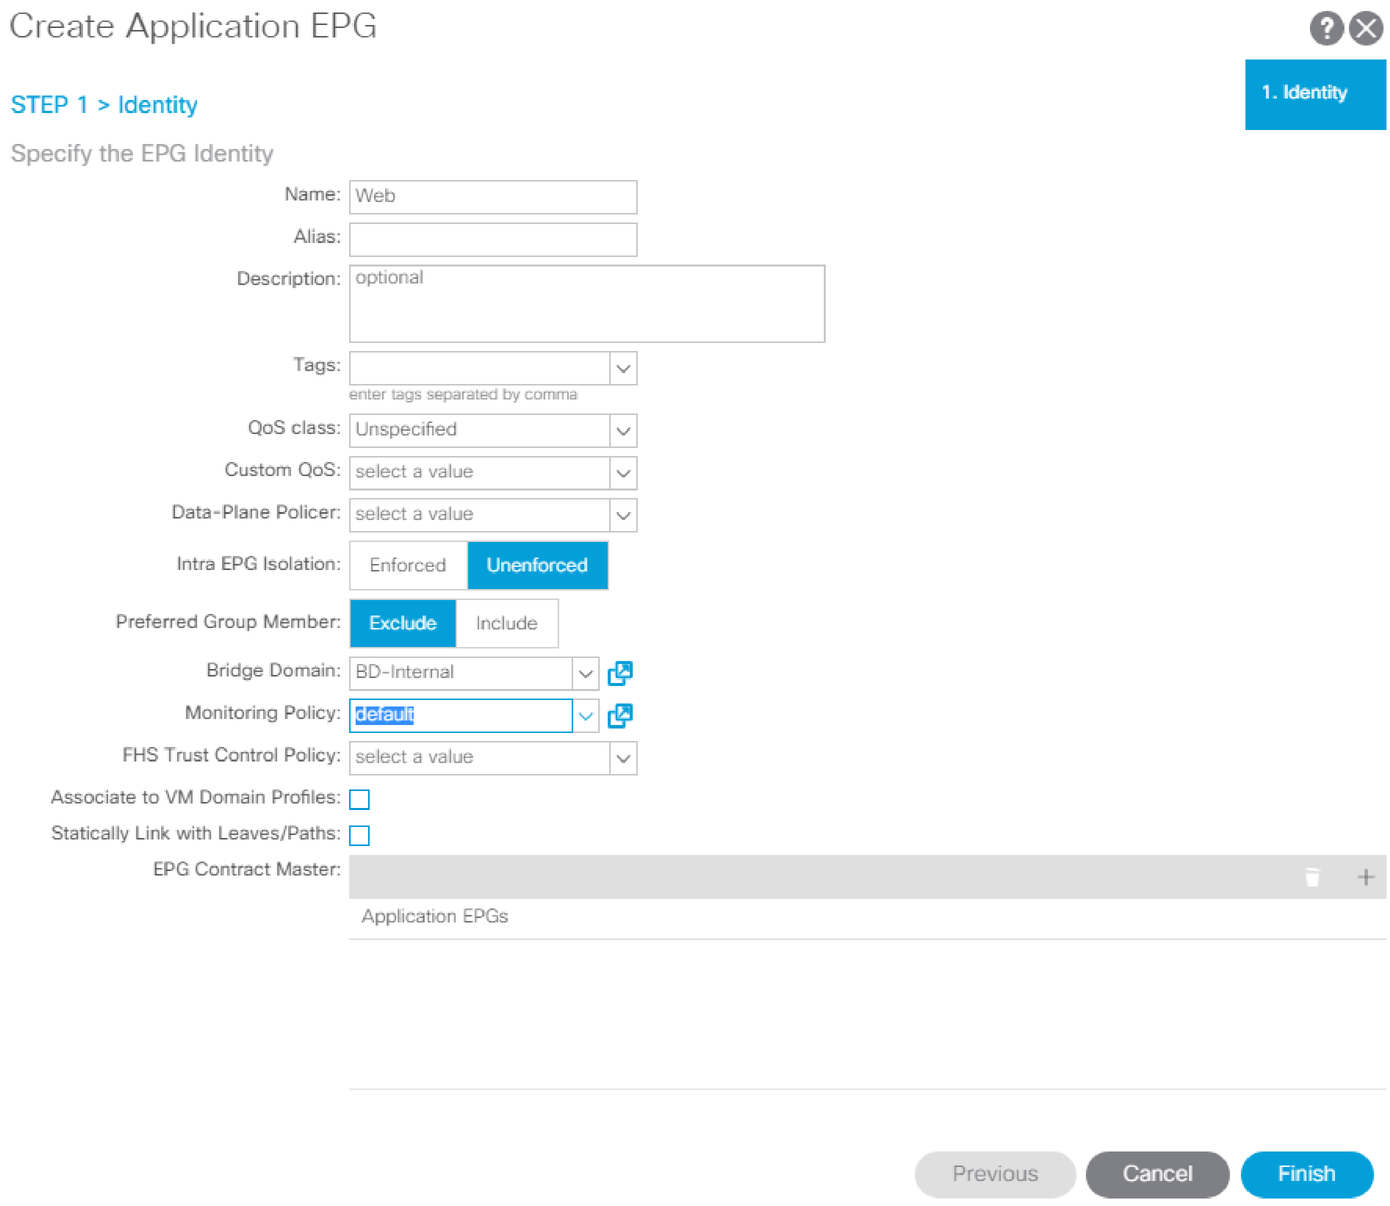

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: Web

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Bridge Domain: BD-Internal

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

153. Click Finish to complete creating the EPG.

154. On the left expand 3-Tier-App, Application EPGs, and EPG Web.

155. Under EPG Web, right-click Domains and select Add VMM Domain Association.

156. Select the fpv-vc-vDS VMM Domain Profile.

157. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. Leave VLAN Mode set at Dynamic.

158. Click Submit to compete adding the VMM Domain Association.

159. On the left under EPG Web, right-click Contracts and select Add Provided Contract.

160. In the Add Provided Contract window, use the Contract drop-down list to select Create Contract.

161. Name the Contract `Allow-Web-App`. Select the Application Profile Scope.

162. Click the "+" sign to add a Contract Subject.

163. Name the subject Allow-All.

164. Click the "+" sign to add a Contract filter.

165. Use the drop-down list to select the Allow-All filter from Tenant common. Click Update.

166. Click OK to complete creating the Contract Subject.

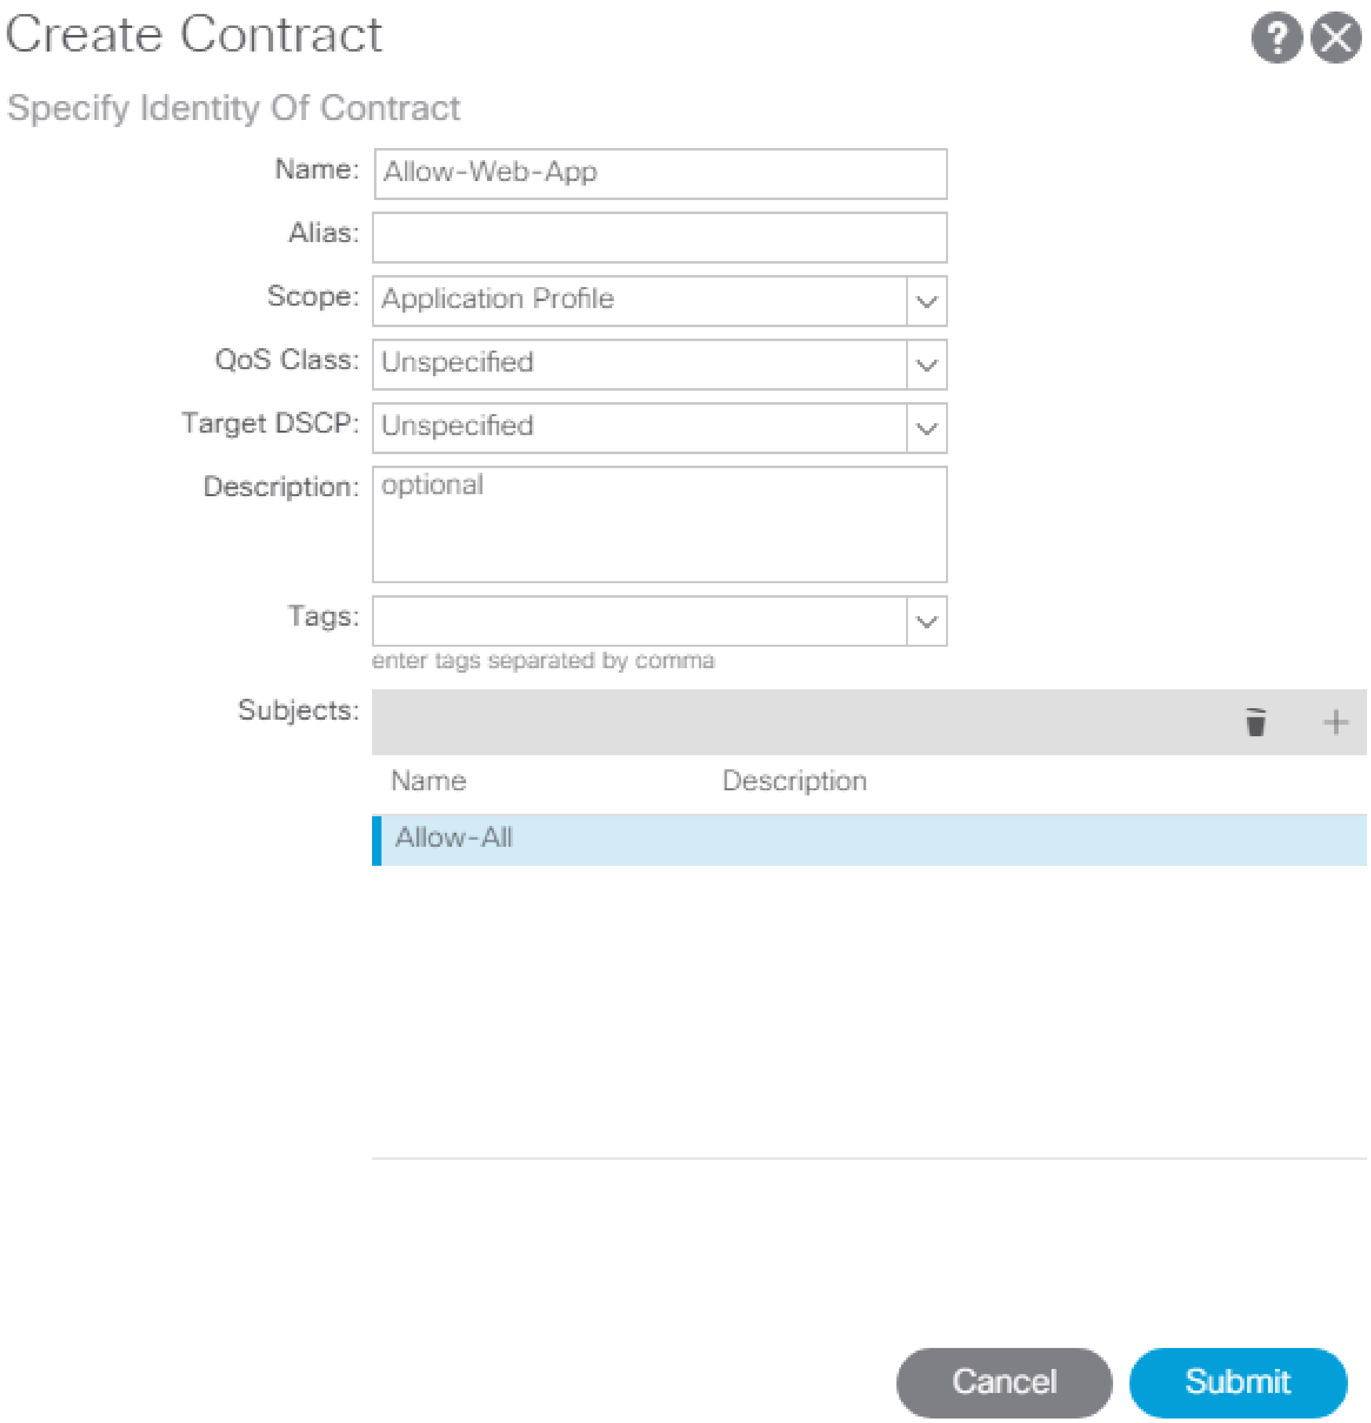

Create Contract

Specify Identity Of Contract

Name:	Allow-Web-App
Alias:	
Scope:	Application Profile
QoS Class:	Unspecified
Target DSCP:	Unspecified
Description:	optional
Tags:	enter tags separated by comma

Subjects:

Name	Description
Allow-All	

167. Click Submit to complete creating the Contract.

168. Click Submit to complete adding the Provided Contract.

169. Right-click Contracts and select Add Consumed Contract.

170. In the Add Consumed Contract window, use the Contract drop-down list to select the common/Allow-FP-Shared-L3-Out contract.

171. Click Submit to complete adding the Consumed Contract.

172. Optionally, repeat steps 169-171 to add the common/common-Allow-Core-Services Consumed Contract.

173. On the left under EPG Web, right-click Subnets and select Create EPG Subnet.

174. 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.

175. For scope, select Advertised Externally and Shared between VRFs.

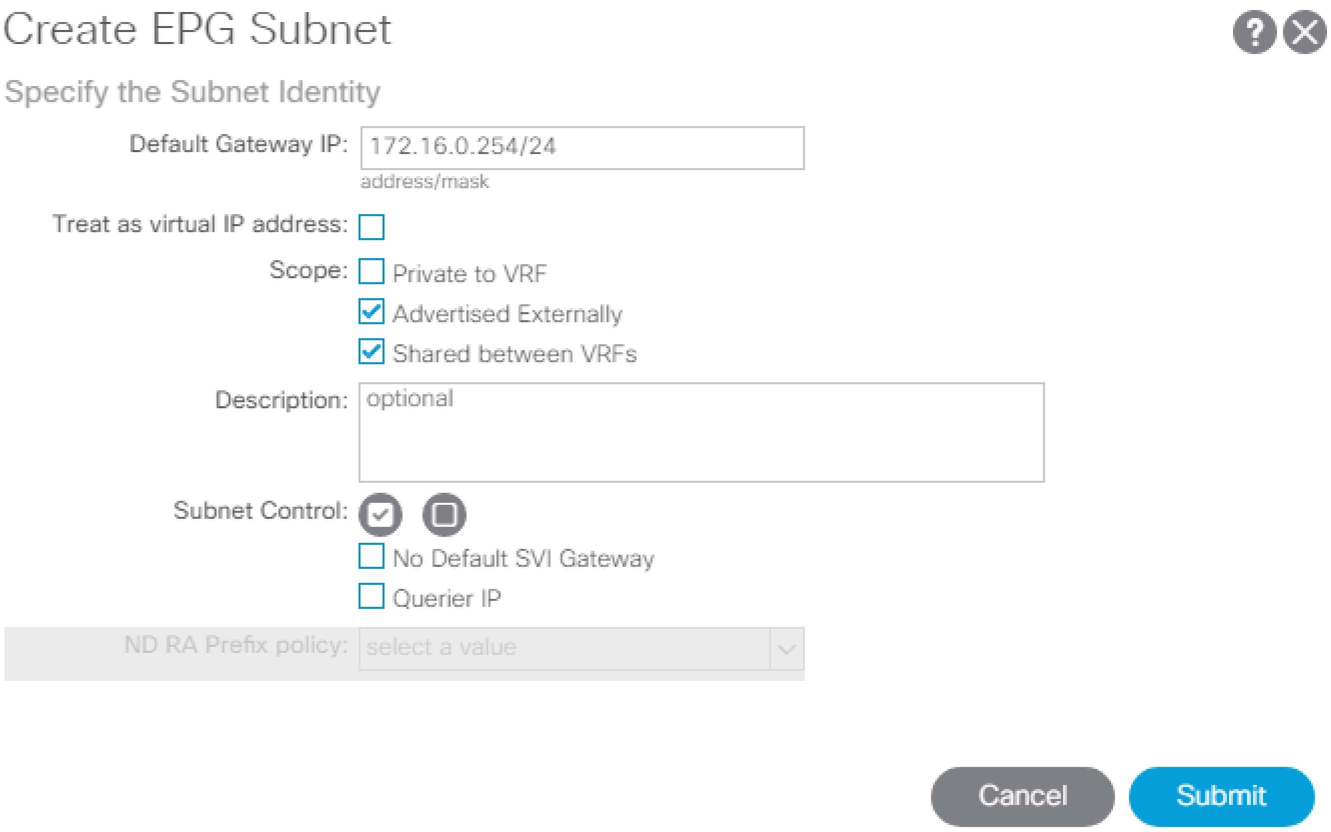

Create EPG Subnet

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: optional

Subnet Control:
□ No Default SVI Gateway
□ Querier IP

ND RA Prefix policy: select a value

176. Click Submit to complete creating the EPG Subnet.

177. Right-click Application EPGs under 3-Tier-App and select Create Application EPG.

178. Name the EPG `App` and leave Intra EPG Isolation set at Unenforced.

179. Use the Bridge Domain drop-down list to select BD-Internal within the current Tenant. Select the default Monitoring Policy.

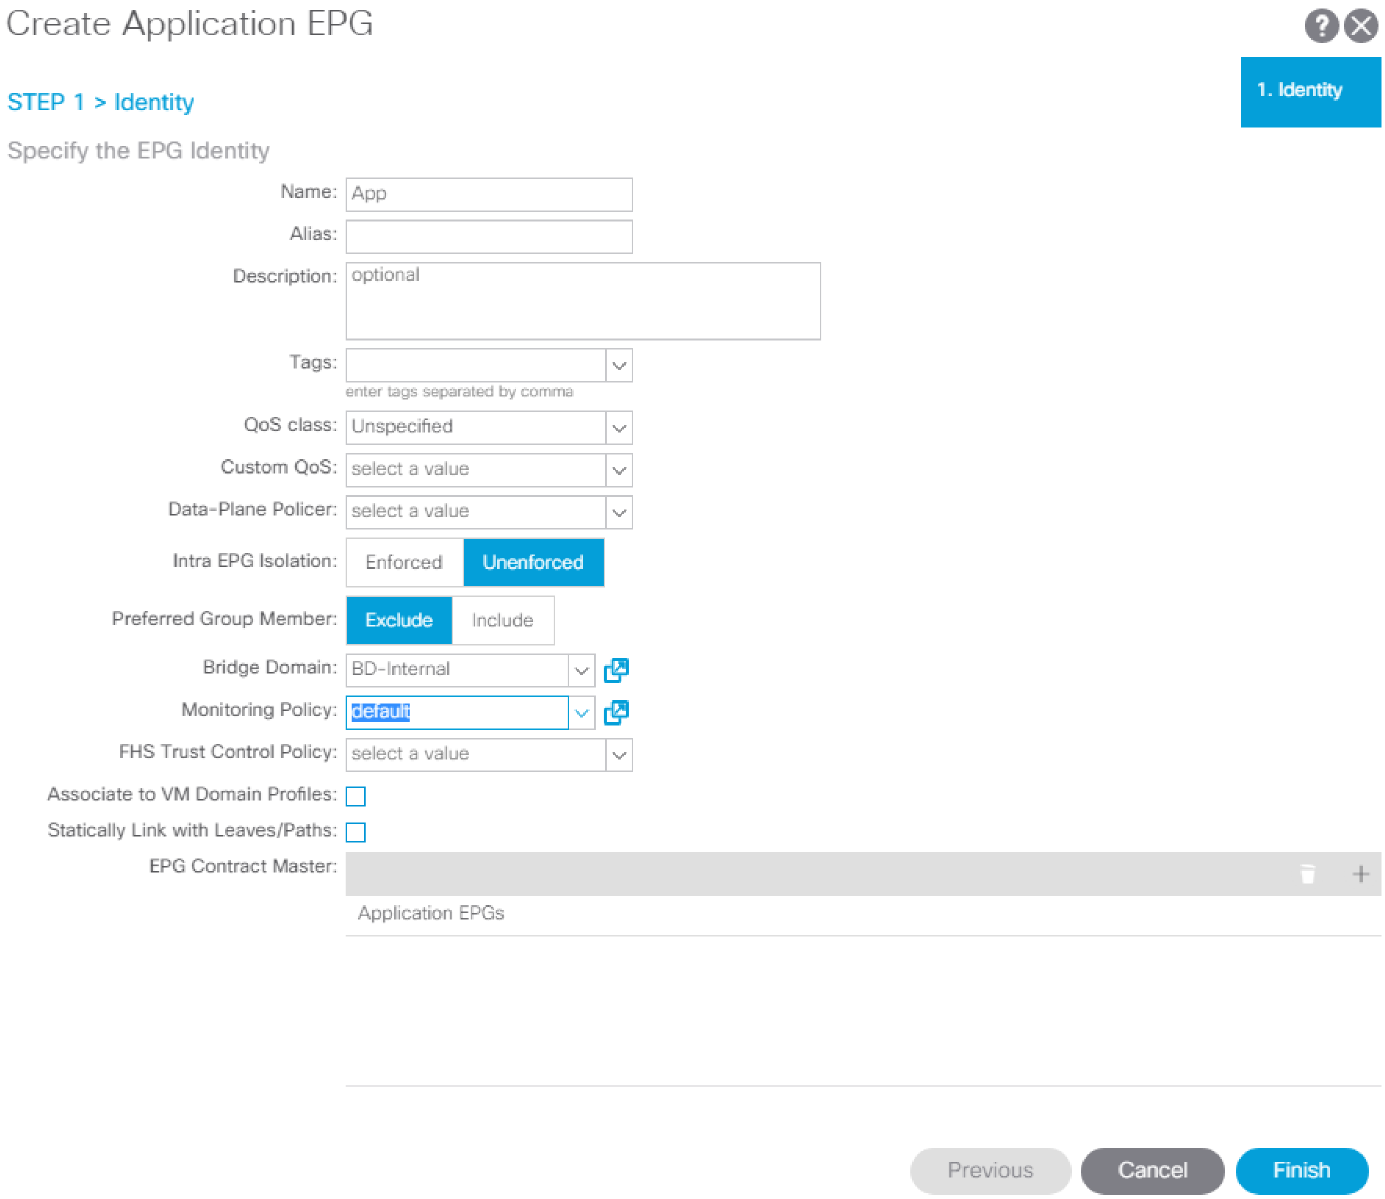

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: App

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Bridge Domain: BD-Internal

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

180. Click Finish to complete creating the EPG.

181. On the left expand 3-Tier-App, Application EPGs, and EPG App.

182. Under EPG App, right-click Domains and select Add VMM Domain Association.

183. Select the fpv-vc-vDS VMM Domain Profile.

184. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. Select the Dynamic VLAN Mode.

185. Click Submit to compete adding the VMM Domain Association.

186. On the left under EPG App, right-click Contracts and select Add Provided Contract.

187. In the Add Provided Contract window, use the Contract drop-down list to select Create Contract.

188. Name the Contract `Allow-App-DB`. Select the Application Profile Scope.

189. Click the "+" sign to add a Contract Subject.

190. Name the subject `Allow-All`.

191. Click the "+" sign to add a Contract filter.

192. Use the drop-down list to select the Allow-All filter from Tenant common. Click Update.

193. Click OK to complete creating the Contract Subject.

194. Click Submit to complete creating the Contract.

195. Click Submit to complete adding the Provided Contract.

196. Right-click Contracts and select Add Consumed Contract.

197. In the Add Consumed Contract window, use the Contract drop-down list to select the Allow-Web-App contract in the current tenant.

198. Click Submit to complete adding the Consumed Contract.

199. Optionally, repeat steps 196-198 to add the common/common-Allow -Core-Services Consumed Contract.

200. On the left under EPG App, right-click Subnets and select Create EPG Subnet.

201. 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.

202. 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.

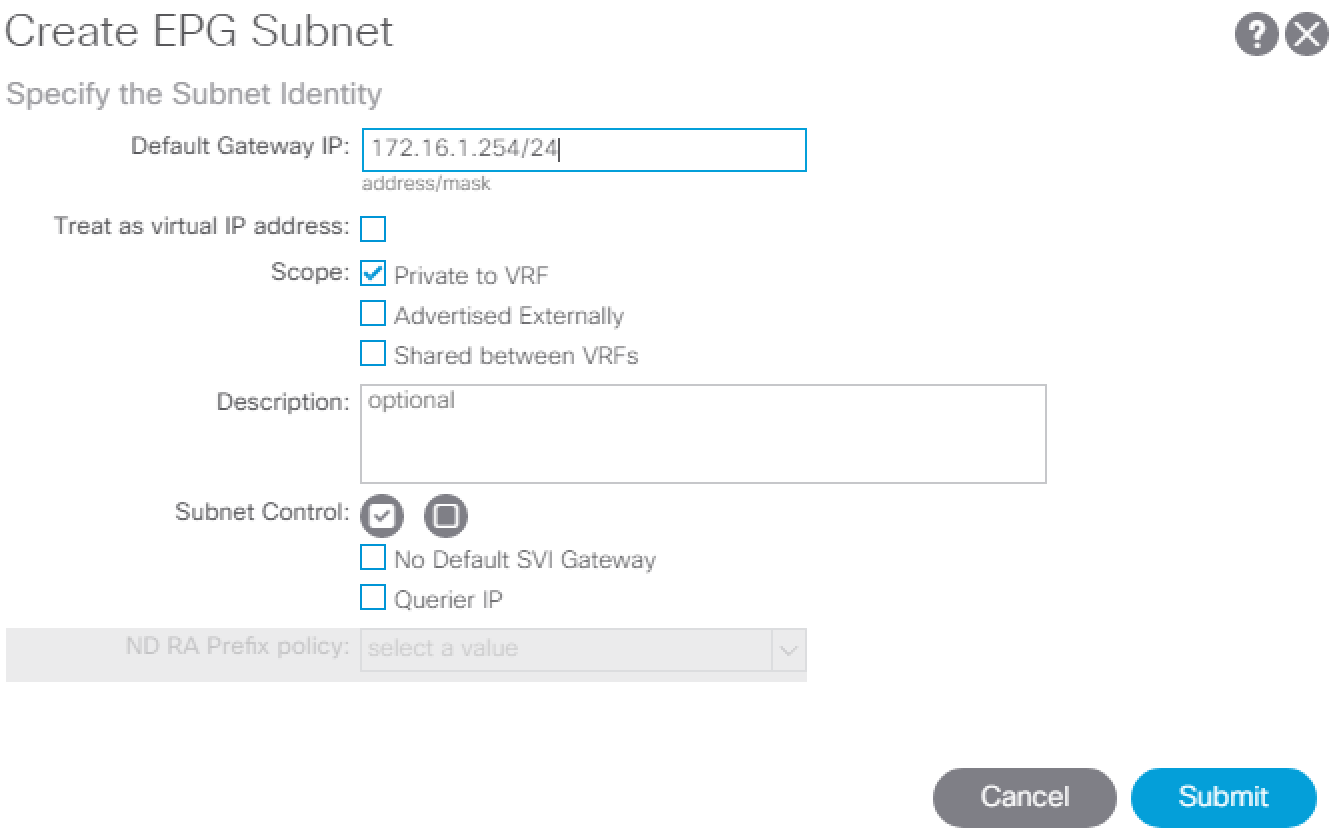

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP: 172.16.1.254/24
address/mask

Treat as virtual IP address: □

Scope: ☑ Private to VRF
□ Advertised Externally
□ Shared between VRFs

Description: optional

Subnet Control:
□ No Default SVI Gateway
□ Querier IP

ND RA Prefix policy: select a value

203. Click Submit to complete creating the EPG Subnet.

204. Right-click Application EPGs under 3-Tier-App and select Create Application EPG.

205. Name the EPG `DB` and leave Intra EPG Isolation set at Unenforced.

206. Use the Bridge Domain drop-down list to select BD-Internal in the current tenant. Select the default Monitoring Policy.

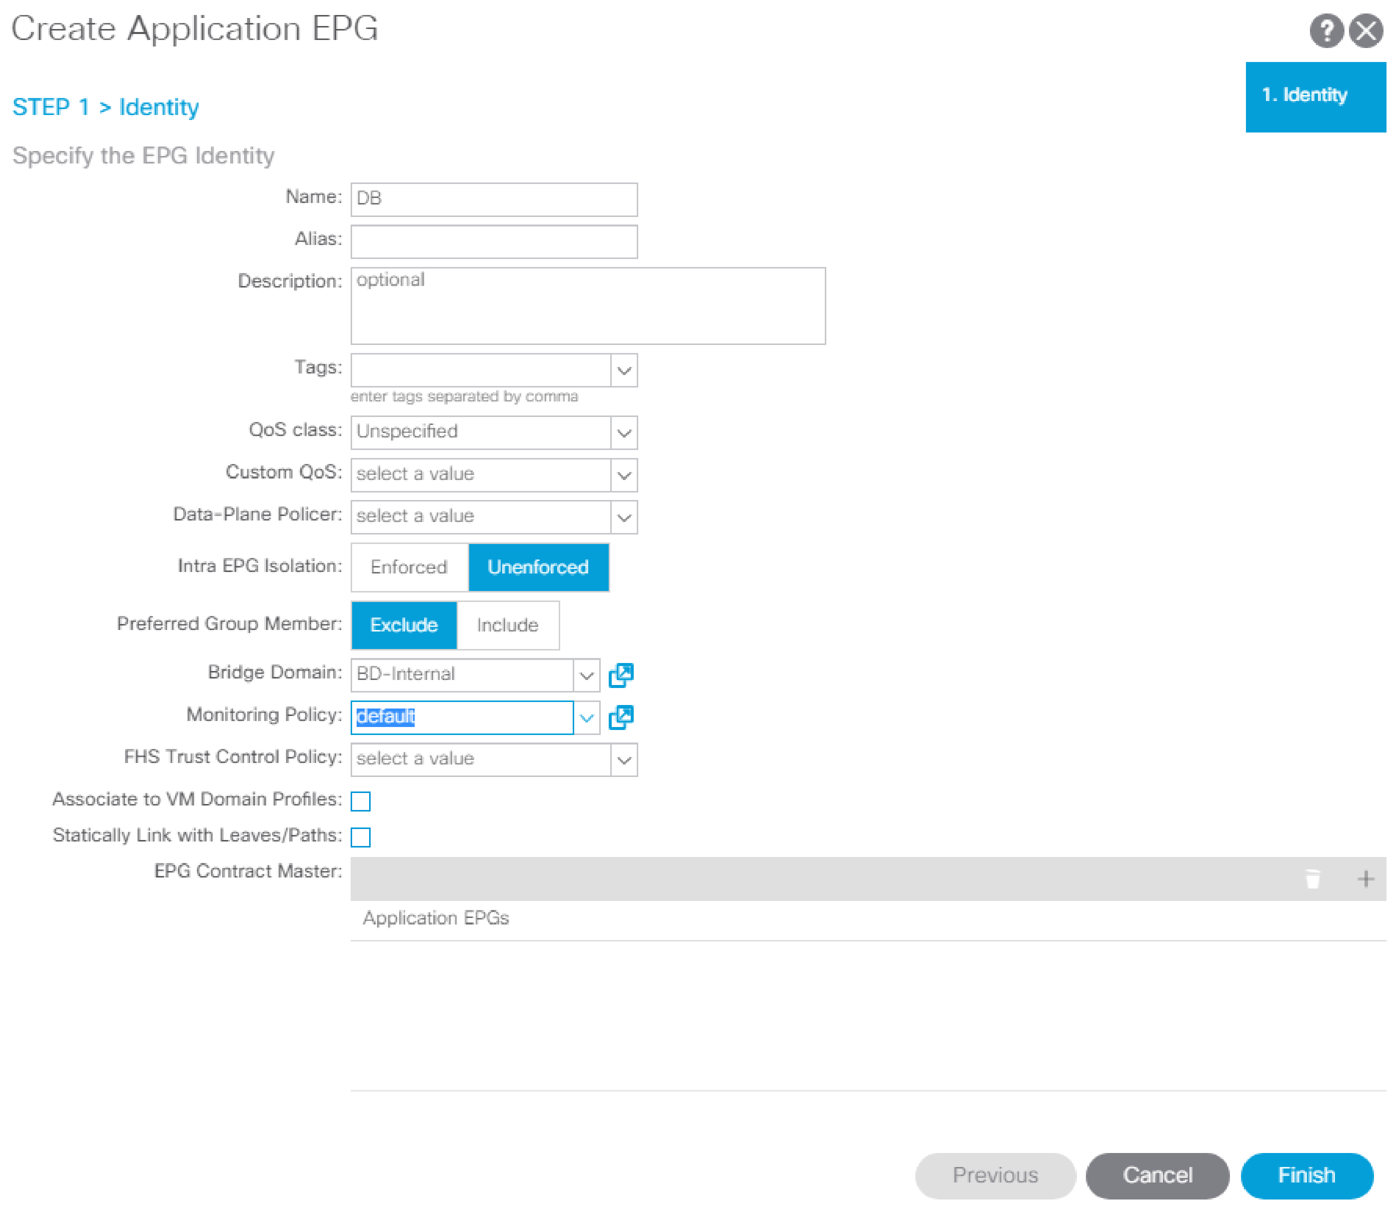

Create Application EPG

STEP 1 > Identity

Specify the EPG Identity

Name: DB

Alias:

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select a value

Data-Plane Policer: select a value

Intra EPG Isolation: Enforced [✓ Unenforced]

Preferred Group Member: [✓ Exclude] Include

Bridge Domain: BD-Internal

Monitoring Policy: default

FHS Trust Control Policy: select a value

Associate to VM Domain Profiles: □

Statically Link with Leaves/Paths: □

EPG Contract Master:

Application EPGs

207. Click Finish to complete creating the EPG.

208. On the left expand 3-Tier-App, Application EPGs, and EPG DB.

209. Under EPG DB, right-click Domains and select Add VMM Domain Association.

210. Select the fpv-vc-vDS VMM Domain Profile.

211. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. Select the Dynamic VLAN Mode.

212. Click Submit to compete adding the VMM Domain Association.

213. On the left under EPG DB, right-click Contracts and select Add Consumed Contract.

214. In the Add Consumed Contract window, use the Contract drop-down list to select the Allow-App-DB contract in the current tenant.

215. Click Submit to complete adding the Consumed Contract.

216. Repeat steps 213-215 to add the common/common-Allow-Core-Services (optional) and Allow-SVM-MGMT in the current tenant Consumed Contracts.

217. On the left under EPG DB, right-click Subnets and select Create EPG Subnet.

218. 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.

219. If the common-Allow-Core-Services contract was consumed, select only the Shared between VRFs scope. Otherwise, select only the Private to VRF scope.

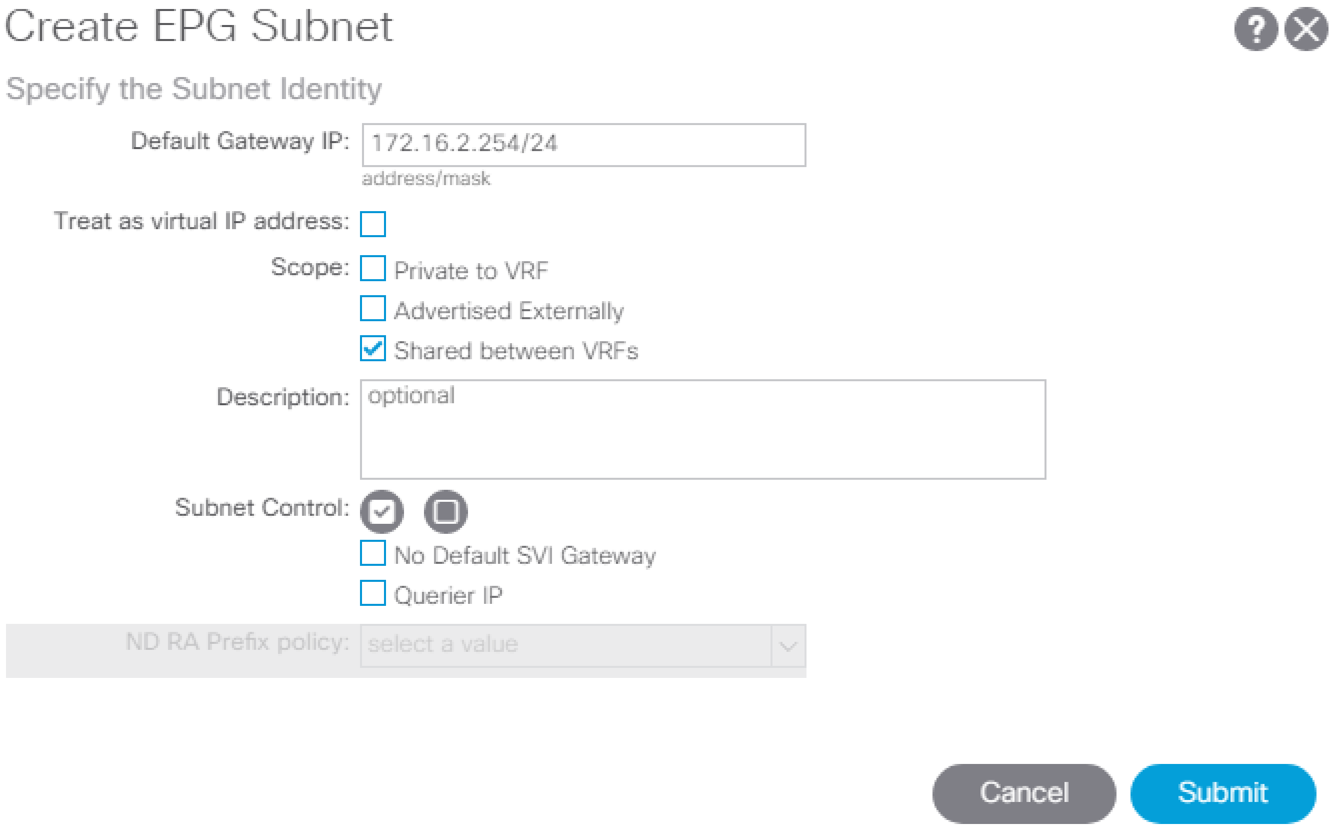

Create EPG Subnet

Specify the Subnet Identity

Default Gateway IP: 172.16.2.254/24
address/mask

Treat as virtual IP address: □

Scope: □ Private to VRF
□ Advertised Externally
☑ Shared between VRFs

Description: optional

Subnet Control:
□ No Default SVI Gateway
□ Querier IP

ND RA Prefix policy: select a value

220. Click Submit to complete creating the EPG Subnet.

Configure Tenant Storage

This section describes the procedure for deploying a NetApp storage SVM for a tenant named FPV-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.

Storage Cluster CLI

Create Tenant IPspace

To create the tenant IPspace, run the following commands:

```
ipspace create -ipspace FPV-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 FPV-App-A -broadcast-domain FPV-App-A-NFS -mtu 9000
broadcast-domain create -ipspace FPV-App-A -broadcast-domain FPV-App-A-SVM-MGMT -mtu 1500
broadcast-domain create -ipspace FPV-App-A -broadcast-domain FPV-App-A-iSCSI -mtu 9000
broadcast-domain show -ipspace FPV-App-A
```

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-FPV-App-A-nfs-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-FPV-App-A-nfs-vlan-id>

broadcast-domain add-ports -ipspace FPV-App-A -broadcast-domain FPV-App-A-NFS -ports <node01>:a0a-
<storage-FPV-App-A-nfs-vlan-id>, <node02>:a0a-<storage-FPV-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-FPV-App-A-svm-mgmt-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-FPV-App-A-svm-mgmt-vlan-id>

broadcast-domain add-ports -ipspace FPV-App-A -broadcast-domain FPV-App-A-SVM-MGMT -ports
<node01>:a0a-<storage-FPV-App-A-svm-mgmt-vlan-id>, <node02>:a0a-<storage-FPV-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-FPV-App-A-iscsi-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-FPV-App-A-iscsi-vlan-id>

broadcast-domain add-ports -ipspace FPV-App-A -broadcast-domain FPV-App-A-iSCSI -ports <node01>:a0a-
<storage-FPV-App-A-iscsi-vlan-id>,<node02>:a0a-<storage-FPV-App-A-iscsi-vlan-id>

broadcast-domain show -ipspace FPV-App-A
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver FPV-App-A-SVM -rootvolume rootvol -aggregate aggr1_node02 -rootvolume-
security-style unix -ipspace FPV-App-A
```

2. Remove the unused data protocols (CIFS, NDMP, and optionally FCP) from the SVM.

```
vserver remove-protocols -vserver FPV-App-A-SVM -protocols cifs,ndmp,fcp
```

3. Add the two data aggregates to the FPV-App-A-SVM aggregate list.

```
vserver modify -vserver FPV-App-A-SVM -aggr-list aggr1_node01,aggr1_node02
```

Setup SVM Management Access

1. Create the SVM Management interface.

```
network interface create -vserver FPV-App-A-SVM -lif SVM-MGMT -role data -data-protocol none -home-
node <st-node-02> -home-port a0a-<FPV-App-A-SVM-svm-mgmt-vlan-id> -address <svm-mgmt-ip> -netmask
<svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true
```

2. Place an entry in your AD DNS server for the FPV-App-A-SVM management interface.

3. Create SVM default route.

```
network route create -vserver FPV-App-A-SVM -destination 0.0.0.0/0 -gateway 172.16.254.22

network route show
Vserver             Destination     Gateway         Metric
------------------- --------------- --------------- ------
FPV-App-A-SVM
                    0.0.0.0/0       172.16.254.22   20
FPV-Foundation-SVM
                    0.0.0.0/0       172.16.254.6    20
a02-affa300
                    0.0.0.0/0       192.168.1.254   20
3 entries were displayed.
```

4. Create snapdrive SVM user and unlock vsadmin SVM user.

```
security login create -user snapdrive -application http -authentication-method password -role vsadmin
-vserver FPV-App-A-SVM

Please enter a password for user 'snapdrive':
Please enter it again:

security login create -user snapdrive -application ontapi -authentication-method password -role
vsadmin -vserver FPV-App-A-SVM

security login password -username vsadmin -vserver FPV-App-A-SVM

Enter a new password:
Enter it again:

security login unlock -username vsadmin -vserver FPV-App-A-SVM
```

Create the NFS Service

You can enable and configure NFSv3 servers on storage virtual machines (SVMs) with NetApp FlexVol® volumes to let NFS clients access files on your cluster. It is a best practice to configure the DNS service on an SVM. To configure DNS and NFS, complete the following steps:

1. Configure the DNS for your SVM.

```
dns create -vserver FPV-App-A-SVM -domains <domain_name> -name-servers
<dns_server1_ip>,<dns_server2_ip>
```

The SVM Management EPG was connected to the Core Services EPG by contract, allowing access to the DNS servers. For more than one DNS server, separate the IPs by comma.

2. Create the NFS service.

```
vserver nfs create -vserver FPV-App-A-SVM –udp disabled
```

3. Add support for the NetApp NFS VAAI plugin.

```
vserver nfs modify -vserver FPV-App-A-SVM –vstorage enabled
```

4. Modify SVM root volume to use default NFS export policy.

```
volume modify -vserver FPV-App-A-SVM -volume rootvol -policy default
```

Modify NFS Export Policy

To modify the default NFS export policy that limits access to devices in the NFS subnet, run the following command:

```
export-policy rule create -vserver FPV-App-A-SVM -policyname default -clientmatch <tenant-nfs-subnet-
cidr> -protocol nfs -rorule sys -rwrule sys -superuser sys -allow-suid false
```

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 FPV-App-A-SVM –volume rootvol_m01 –aggregate aggr1_node01 –size 1GB –type DP
volume create –vserver FPV-App-A-SVM –volume rootvol_m02 –aggregate aggr1_node02 –size 1GB –type DP
```

2. Create the mirroring relationships.

```
snapmirror create –source-path FPV-App-A-SVM:rootvol –destination-path FPV-App-A-SVM:rootvol_m01 –
type LS -schedule 15min
snapmirror create –source-path FPV-App-A-SVM:rootvol –destination-path FPV-App-A-SVM:rootvol_m02 –
type LS -schedule 15min
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set –source-path FPV-App-A-SVM:rootvol
snapmirror show
```

Create Block Protocol Service(s)

If the deployment is using iSCSI and iSCSI is being deployed in the tenant SVM, create the iSCSI service on each SVM using the following command. This command also starts the iSCSI service and sets the IQN for the SVM.

```
iscsi create -vserver FPV-App-A-SVM

iscsi show
```

If the deployment is using FCP, create the FCP service on each SVM using the following command. This command also starts the FCP service and sets the worldwide name (WWN) for the SVM.

```
fcp create -vserver FPV-App-A-SVM

fcp show
```

The licenses for FCP and iSCSI must be installed before the services can be started. If the license(s) weren't installed during cluster setup, install them before this step.

Configure HTTPS Access

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

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

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

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

```
security certificate show
```

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

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

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

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

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

5. To obtain the values for the parameters required in step 5 (`<cert-ca>` and `<cert-serial>`), run the security certificate show command.

6. Enable each certificate that was just created by using the –server-enabled true and –client-enabled false parameters. Use TAB completion to aid in the completion of these commands.

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

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

```
set -privilege admin
vserver services web modify –name spi|ontapi|compat –vserver * -enabled true
```

Create NetApp FlexVol Volumes

```
volume create -vserver FPV-App-A-SVM -volume fpv_app_a_nfs_datastore_1 -aggregate aggr1_node01 -size
500GB -state online -policy default -security-style unix -junction-path /fpv_app_a_nfs_datastore_1  -
space-guarantee none -percent-snapshot-space 0

volume create -vserver FPV-App-A-SVM -volume fpv_app_a_nfs_datastore_2 -aggregate aggr1_node02 -size
500GB -state online -policy default -security-style unix -junction-path /fpv_app_a_nfs_datastore_2  -
space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path FPV-App-A-SVM:rootvol
```

Schedule Deduplication

On NetApp All Flash FAS systems, deduplication is enabled by default. To schedule deduplication, complete the following steps:

1. After the volumes are created, assign to them a once-a-day deduplication schedule:

```
efficiency modify –vserver FPV-App-A-SVM –volume fpv_app_a_nfs_datastore_1 –schedule sun-sat@0
efficiency modify –vserver FPV-App-A-SVM –volume fpv_app_a_nfs_datastore_1 –schedule sun-sat@0
```

Create SAN LIFs

If using iSCSI, run the following commands to create four iSCSI LIFs (two on each node):

```
network interface create -vserver FPV-App-A-SVM -lif iscsi_lif01 -role data -data-protocol iscsi -
home-node <st-node01> -home-port a0a-<iSCSI-VLAN> –address <iscsi-lif01-ip> -netmask <iscsi-lif01-
mask> –status-admin up

network interface create -vserver FPV-App-A-SVM -lif iscsi_lif02 -role data -data-protocol iscsi -
home-node <st-node02> -home-port a0a-<iSCSI-VLAN> –address <iscsi-lif02-ip> -netmask <iscsi-lif02-
mask> –status-admin up
```

If using FCP, run the following commands to create four FC LIFs (two on each node):

```
network interface create -vserver FPV-App-A-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-
node <st-node01> -home-port 0e –status-admin up

network interface create -vserver FPV-App-A-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-
node <st-node01> -home-port 0f –status-admin up

network interface create -vserver FPV-App-A-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-
node <st-node02> -home-port 0e –status-admin up

network interface create -vserver FPV-App-A-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-
node <st-node02> -home-port 0f –status-admin up
```

Create NFS LIFs

To create NFS LIFs, run the following commands:

```
network interface create -vserver FPV-App-A-SVM -lif nfs_lif01 -role data -data-protocol nfs -home-
node <st-node01> -home-port a0a-<tenant-nfs-vlan-id> –address <node01-nfs-lif01-ip> -netmask <node01-
nfs-lif01-mask> -status-admin up –failover-policy broadcast-domain-wide –firewall-policy data –auto-
revert true

network interface create -vserver FPV-App-A-SVM -lif nfs_lif02 -role data -data-protocol nfs -home-
node <st-node02> -home-port a0a-<tenant-nfs-vlan-id> –address <node02-nfs-lif02-ip> -netmask <node02-
nfs-lif02-mask> -status-admin up –failover-policy broadcast-domain-wide –firewall-policy data –auto-
revert true

network interface show
```

Add 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 FPV-App-A -vserver FPV-App-A-SVM -max-throughput INF
vserver modify -vserver FPV-App-A-SVM -qos-policy-group FPV-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 FPV-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. An ESXi cluster can also be set up to host more than one tenant. 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 design, the same Cisco UCS service profile template that was used for the management ESXi hosts can be used to generate new service profiles for the tenant hosts. In order to accommodate, different server requirements and to allow the usage of different server pools for tenant hosts, the management host server profile template will be cloned and modified to use a different server pool. All procedures in this section are completed using the Cisco UCS Manager HTML 5 User Interface.

Add New Tenant-Specific Server Pool

Since a clone of the management Service Profile Template will be used for a tenant, a new tenant-specific server pool can be created and mapped in the new cloned Service Profile Template. Create this pool.

Create Tenant Service Profile Template

It is recommended to clone the existing management Service Profile Template for the servers running the applications in the new tenant. These templates can be created by cloning the existing Service Profile Template and modifying it by adding the new tenant-specific server pool.

Create New Service Profiles for Tenant Servers

Using the cloned and modified Tenant Service Profile Template, 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 5 iSCSI LIFs for iSCSI IQN

SVM	iSCSI Target IQN

SVM	iSCSI Target IQN
FPV-Foundation-SVM	
FPV-App-A-SVM	

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

Table 6 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN
fpv-esxi-03	
fpv-esxi-04	

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.”

Configure Storage SAN Boot for the Tenant

This section details the steps for setting up SAN boot for the tenant ESXi Host servers.

ESXi Boot LUN in FPV-Foundation-SVM Tenant Hosts

1. From the cluster management node SSH connection, enter the following:

```
lun create -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-03 -size 15GB -ostype vmware -
space-reserve disabled

lun create -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-04 -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 FPV-Foundation-SVM –igroup fpv-esxi-03 –protocol iscsi –ostype vmware –
initiator <fpv-esxi-03-iqn>

igroup create –vserver FPV-Foundation-SVM –igroup fpv-esxi-04 –protocol iscsi –ostype vmware –
initiator <fpv-esxi-04-iqn>

igroup create –vserver FPV-App-A-SVM –igroup fpv-esxi-03 –protocol iscsi –ostype vmware –initiator
<fpv-esxi-03-iqn>

igroup create –vserver FPV-App-A-SVM –igroup fpv-esxi-04 –protocol iscsi –ostype vmware –initiator
<fpv-esxi-04-iqn>

igroup create –vserver FPV-App-A-SVM –igroup FPV-App-A-Host-All –protocol iscsi –ostype vmware –
initiator <fpv-esxi-03-iqn>,<fpv-esxi-04-iqn>
```

Use the values listed in Table 8 and Table 9 for the IQN information.

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 FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-03-igroup fpv-esxi-03 -lun-id 0

lun map -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-04-igroup fpv-esxi-04 -lun-id 0

lun show -m
```

Install and Configure VMware ESXi on Tenant Hosts

To install and configure VMware ESXi on tenant hosts, complete the following steps:

1. Using the UCS KVM's ISO mapping capability, follow the section of this document titled VMware vSphere 6.5 Update 1 Setup to install and configure the FPV-App-A tenant ESXi hosts. Configure the tenant ESXi hosts in the same IB-MGMT subnet as the management hosts. It is not necessary to add the FPV-Foundation tenant NFS VMkernel ports unless it is desired for the tenant hosts to have access to the infrastructure or management datastores. Access to the infrastructure datastores may be desired since the patches and drivers were placed in infra_datastore_1.

2. Add the ESXi hosts to vCenter in a new FPV-App-A cluster, configure iSCSI coredump if necessary, and add the hosts to the VMware vDS.

3. The main difference in the installation will be that tenant VMkernel Interfaces for iSCSI and NFS will need to be installed in port groups on the fpv-vc-vDS. Note that only one iSCSI VLAN is being used for tenant iSCSI, instead of the two VLANs used for iSCSI boot.

4. To add VMkernel interfaces for FPV-App-A NFS and iSCSI, in VMware vSphere Web Client, on the left select Hosts and Clusters and select the first tenant ESXi hosts. In the center pane, select the Configure tab. Select Network > VMkernel adapters.

5. Select to add a VMkernel adapter. Leave VMkernel Network Adapter selected and click Next.

6. Click Browse, and select the tenant NFS port group on the VMware vDS. Click OK. Click Next.

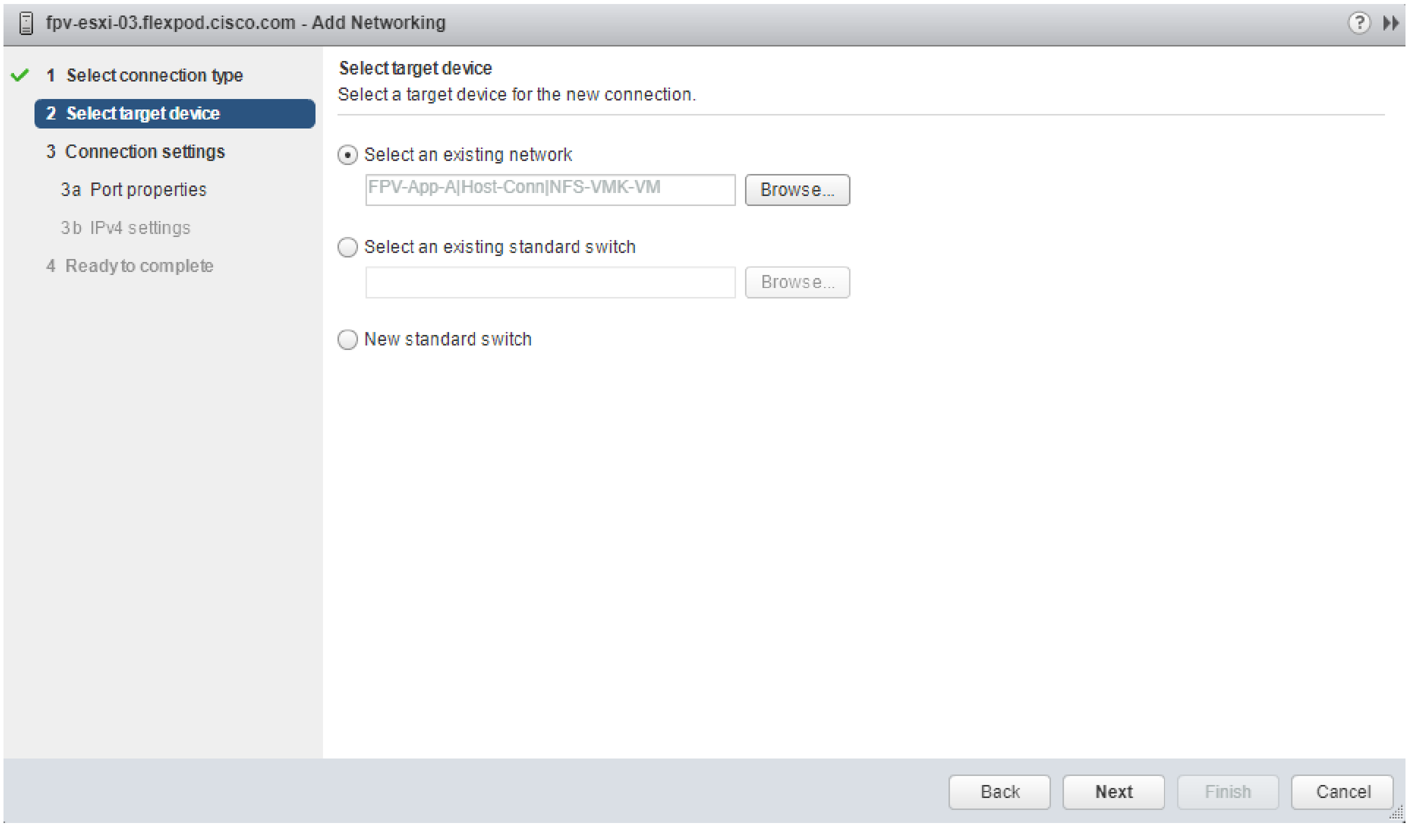

fpv-esxi-03.flexpod.cisco.com - Add Networking

✓ 1 Select connection type
2 Select target device
3 Connection settings
3a Port properties
3b IPv4 settings
4 Ready to complete

Select target device
Select a target device for the new connection.

● Select an existing network
FPV-App-A|Host-Conn|NFS-VMK-VM

○ Select an existing standard switch

○ New standard switch

7. Click Next.

8. Select to use static IPv4 settings and enter an IP address and netmask in the tenant NFS subnet. Click Next.

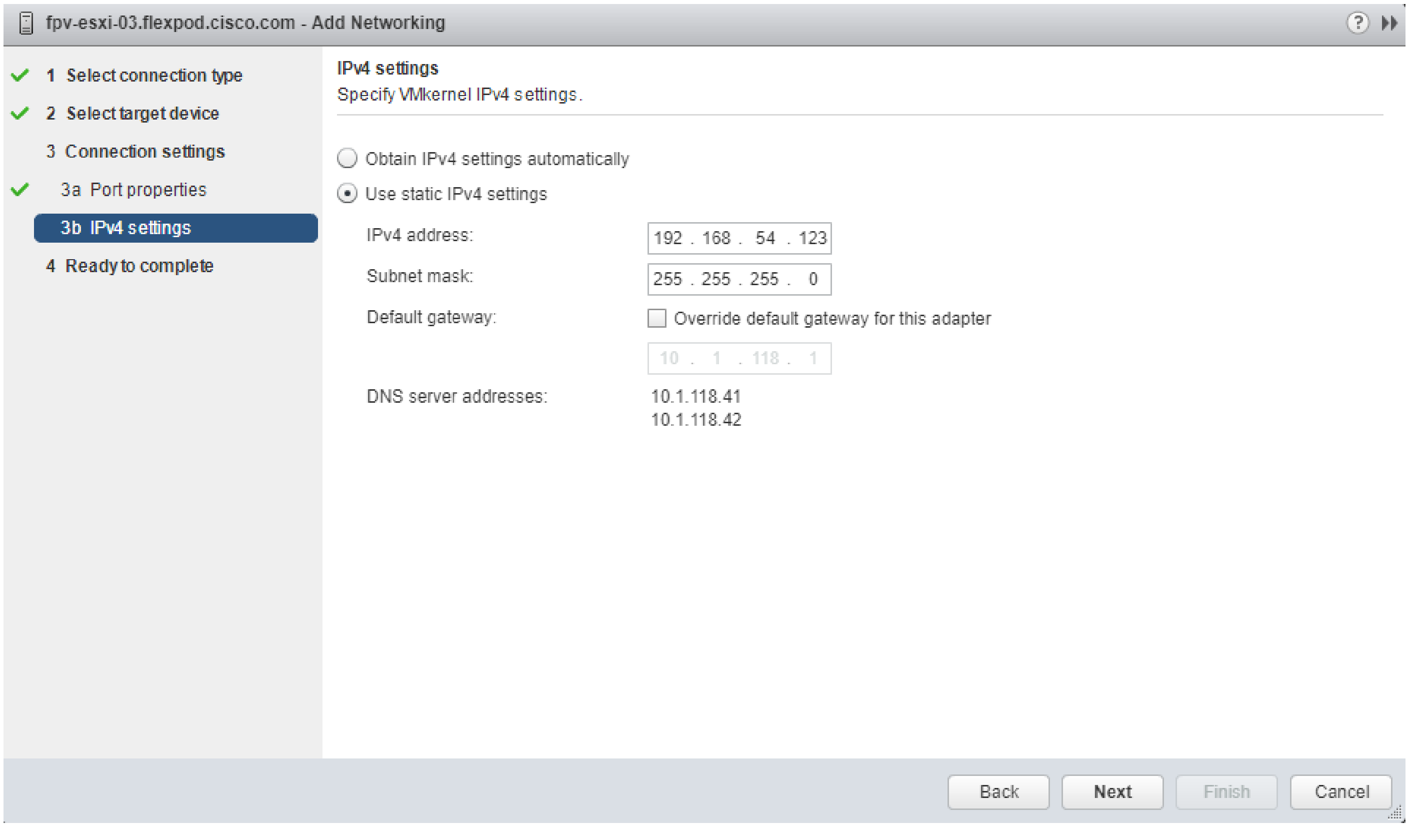

fpv-esxi-03.flexpod.cisco.com - Add Networking

IPv4 settings

Specify VMkernel IPv4 settings.

○ Obtain IPv4 settings automatically

● Use static IPv4 settings

IPv4 address: 192 . 168 . 54 . 123

Subnet mask: 255 . 255 . 255 . 0

Default gateway: □ Override default gateway for this adapter

10 . 1 . 118 . 1

DNS server addresses: 10.1.118.41
10.1.118.42

9. Review the settings and click Finish to complete creating the VMkernel port.

10. In the center pane, highlight the newly-created VMkernel port and select ✎ to edit the VMkernel port settings.

11. Select NIC settings on the left and change the MTU to 9000. Click OK.

12. Repeat this procedure to add tenant NFS and iSCSI VMkernel ports to all tenant ESXi hosts.

13. To add the two tenant iSCSI targets, from the host Configure tab in vSphere Web Client, select Storage > Storage Adapters. Select the iSCSI Software Adapter and select the Targets tab.

14. Under Targets, make sure Dynamic Discovery is selected and click Add.

15. Enter the IP address of iscsi_lif01 in the FPV-App-A-SVM and click OK. Repeat to add the IP address of iscsi_lif02.

16. Click to rescan all SAN storage devices. Click OK.

17. The new targets can now be verified under Targets > Static Discovery.

18. Repeat this procedure to add the iSCSI targets to all tenant ESXi hosts.

19. To mount the two tenant NFS datastores, from the host Configure tab in vSphere Web Client, select Storage > Datastores. Click to add a datastore.

20. Select NFS and click Next.

21. Leave NFS 3 selected and click Next.

22. Enter the Datastore name, the configured Junction Path of the volume for the Folder, and the IP address of the NFS LIF in the tenant SVM on the storage node where the volume is located. Click Next.

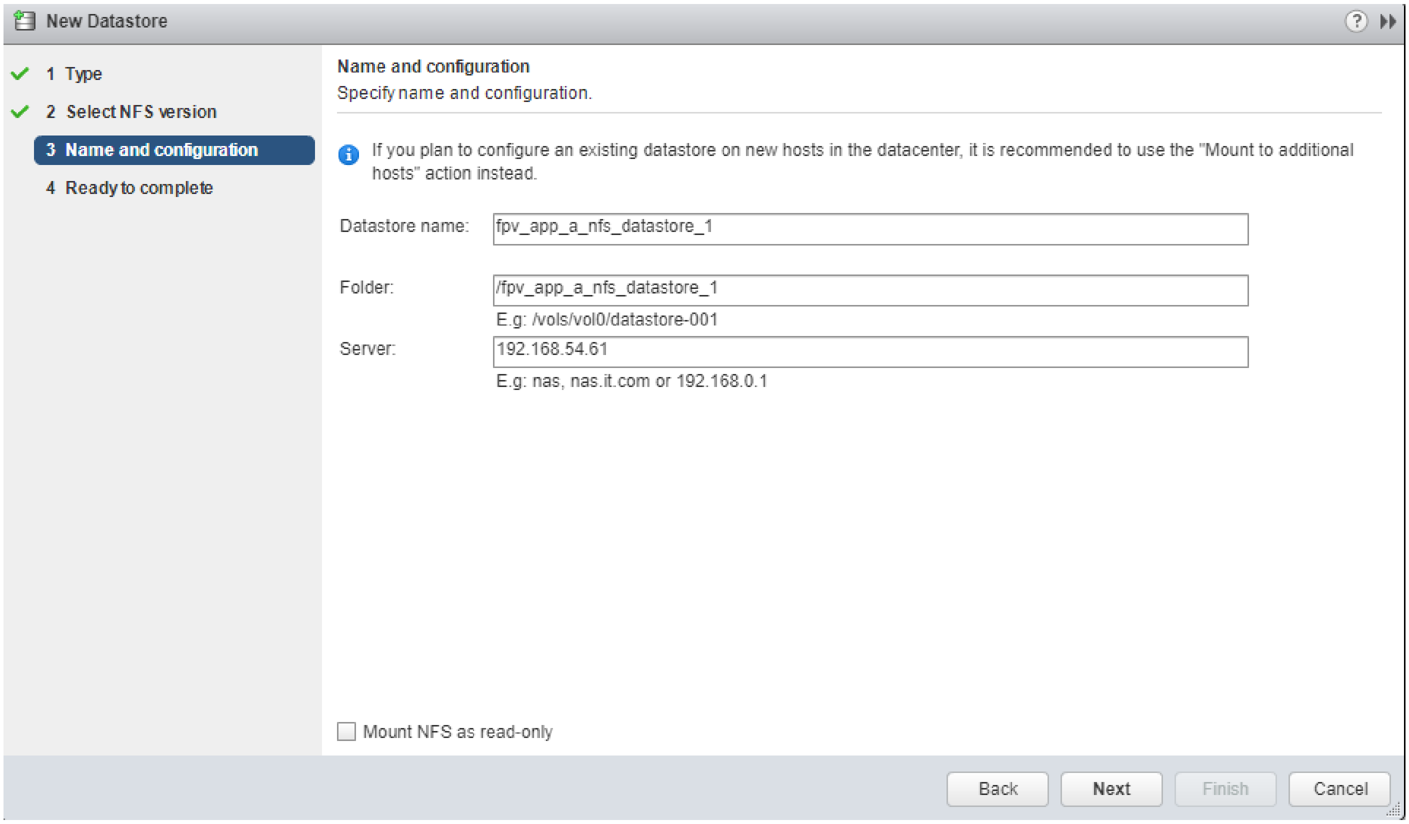

New Datastore

Name and configuration

Specify name and configuration.

If you plan to configure an existing datastore on new hosts in the datacenter, it is recommended to use the "Mount to additional hosts" action instead.

Datastore name: fpv_app_a_nfs_datastore_1

Folder: /fpv_app_a_nfs_datastore_1

E.g: /vols/vol0/datastore-001

Server: 192.168.54.61

E.g: nas, nas.it.com or 192.168.0.1

□ Mount NFS as read-only

23. Review the settings and click Finish to complete mounting the datastore. Repeat this procedure to mount the second tenant NFS datastore.

24. Now that the two tenant NFS datastores have been mounted on one host, NetApp VSC can be used to mount these datastores on the remaining tenant hosts. Right-click the first host and select NetApp VSC > Update Host and Storage Data. Click OK two times. Wait for the NetApp Storage Discovery to complete.

25. Right-click one of the remaining hosts and select NetApp VSC > Mount Datastores. Select only the two tenant NFS datastores and click OK. Repeat this procedure for all remaining tenant hosts.

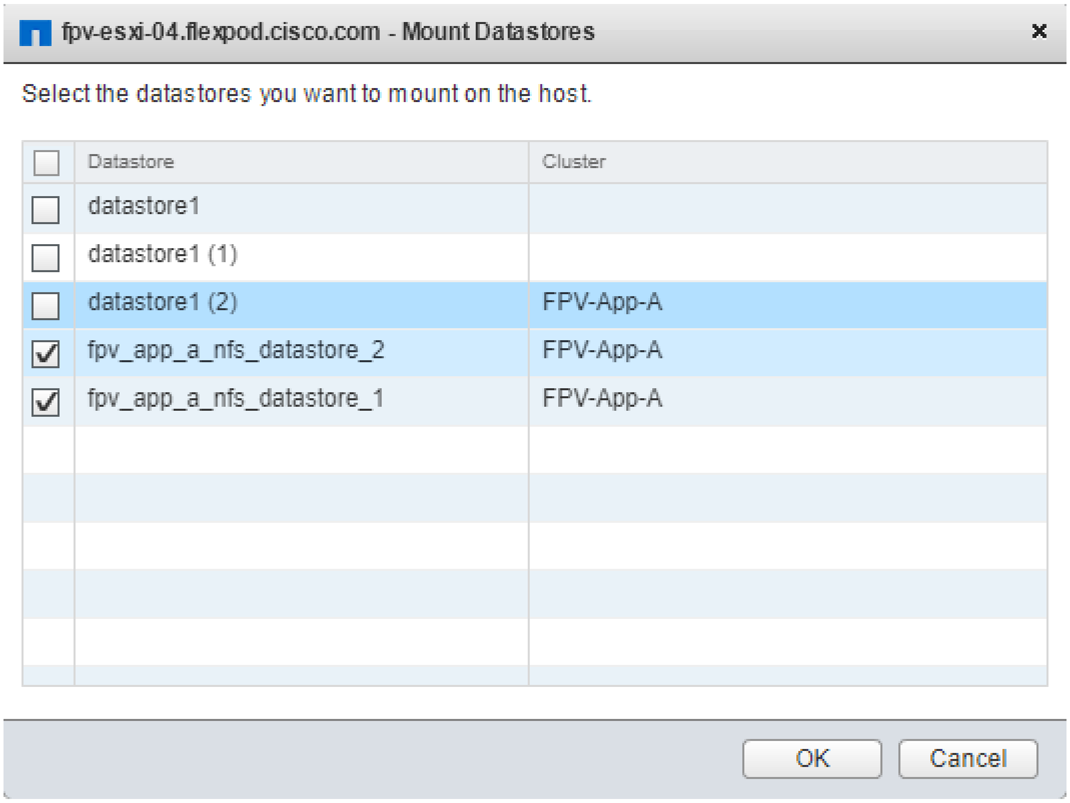

fpv-esxi-04.flexpod.cisco.com - Mount Datastores

Select the datastores you want to mount on the host.

☐	Datastore	Cluster
☐	datastore1	
☐	datastore1 (1)	
☐	datastore1 (2)	FPV-App-A
☑	fpv_app_a_nfs_datastore_2	FPV-App-A
☑	fpv_app_a_nfs_datastore_1	FPV-App-A

26. Using NetApp VSC, configure Optimal Storage Settings on the VMware ESXi hosts, put them in Maintenance Mode, and reboot them.

27. You are now ready to install software into the tenant and begin running a workload. Because FCoE zoning or iSCSI targets to the FPV-App-A-SVM LIFs are in place, SAN storage can be provisioned with NetApp VSC. New NFS storage can also be provisioned with VSC.

28. Use NetApp SnapCenter to create backup schedules for tenant datastores.

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. Any subnet connected to Core-Services or the Shared-L3 Out must have IP addresses unique within the ACI fabric. Notice below that the storage protocol subnets and the App tier subnet have overlapping IPs with FPV-App-A. The remaining subnets are unique since they connect to either Core-Services or the Shared-L3-Out. The second tenant built in this lab validation had the following characteristics and was built with the same contract structure as the FPV-App-A tenant:

Table 7   Lab Validation Tenant App-B Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI	3015	vDS	192.168.114.0/24 - L2	BD-iSCSI
NFS	3055	vDS	192.168.54.0/24 - L2	BD-NFS
SVM-MGMT	222	N/A	172.16.254.30/29	BD-Internal
Web	N/A	vDS	172.16.3.254/24	BD-Internal
App	N/A	vDS	172.16.1.254/24	BD-Internal
DB	N/A	vDS	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, two Cisco ASAv virtual firewall devices were connected to the ACI Fabric using the APIC-integrated VMware vDS. Inside and Outside vDS port groups were created using EPGs in Tenant common. When the L4-7 Service Graph with the ASAv instance is deployed within a tenant, the new port groups are created in the tenant and the ASAv network interfaces are moved to the new port groups. The ASAv firewalls were deployed between the tenant Web EPG in the 3-Tier App and the Fabric Shared L3 Out interface. VLAN Stitching does not make use of device packages. Instead the firewall was configured using the firewall CLI and ASDM interfaces. The detailed VLANs and IP subnet addresses are listed in Table 8

Table 8 Tenant L4-L7 VLAN Stitching Details

Tenant	Outside VLAN	Outside Firewall IP (ASAv)	Outside Subnet Gateway (ACI)	Inside VLAN	Inside Firewall IP (ASAv)	Web EPG Gateway (ACI)
App-A	vDS	172.16.252.1/30	172.16.252.2/30	vDS	172.16.0.1/24	172.16.0.254/24
App-B	vDS	172.16.252.5/30	172.16.252.6/30	vDS	172.16.3.1.24	172.16.3.254/24

Provide Management connectivity to Cisco ASAv

This section details the configuration necessary to provide management connectivity to Cisco ASAv. The management interfaces on the Cisco ASAv VM are added to the SVM MGMT EPG and managed from either a Core Services EPG or DB Tier EPG.

The SVM MGMT EPG was created in an earlier section titled: Create EPG for Infrastructure SVM-MGMT Access. In this setup, management was done from the Core-Services EPG to which SVM-MGMT already has access to through the established contract.

To provide management connectivity to ASAv from Core-Services EPG, complete the following steps:

APIC GUI

1. From the Cisco APIC GUI, at the top, select Tenants > FPV-Foundation.
2. On the left, expand Tenant FPV-Foundation and select Application Profiles.
3. Expand Application Profiles and select Application Profile `IB-MGMT`.
4. Select and expand Application EPGs.
5. Select `SVM-MGMT` EPG. Right-click and select Add VMM Domain Association from the list.
6. Use the drop-down list to select the `fpv-vc-vDS` VMM Domain Profile. Select Immediate Deploy Immediacy and change no other values. Click Submit to create the `SVM-MGMT` port group in the vDS.

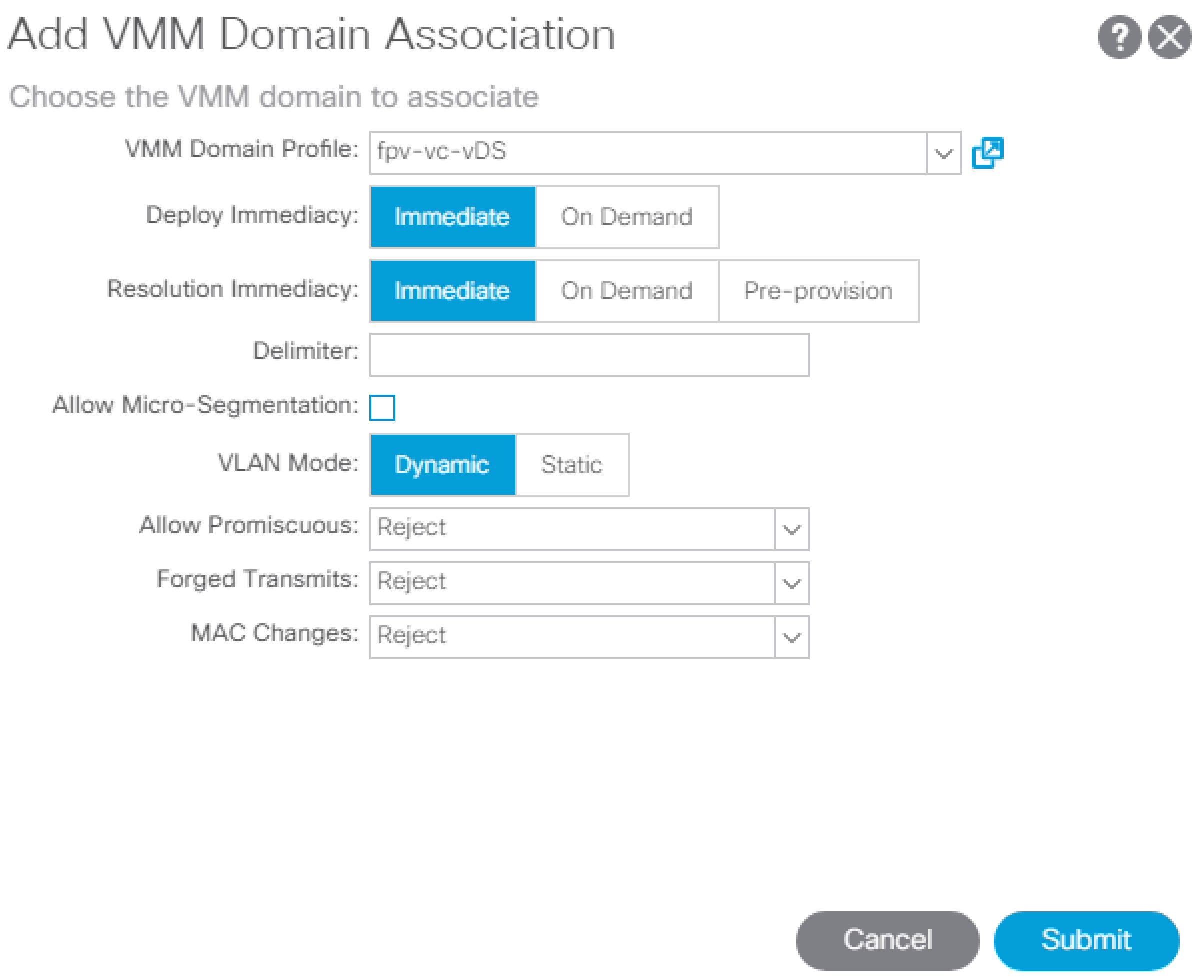

Add VMM Domain Association

Choose the VMM domain to associate

VMM Domain Profile: fpv-vc-vDS

Deploy Immediacy: Immediate

Resolution Immediacy: Immediate

Delimiter:

Allow Micro-Segmentation: □

VLAN Mode: Dynamic

Allow Promiscuous: Reject

Forged Transmits: Reject

MAC Changes: Reject

Deploy Inside and Outside EPGs in Tenant common

This section details setup of the Inside and Outside placeholder EPGs for the ASAv's used in this lab validation.

APIC GUI

1. From the Cisco APIC GUI, at the top, select Tenants > common.
2. On the left, expand Tenant common and Application Profiles.
3. Right-click Application Profiles and select Create Application Profile.
4. Name the Application Profile `ASAv-Deployment`. Under EPGs, click the + to add an EPG.
5. Name the first EPG `Inside`. Use the drop-down list under BD to select Create Bridge Domain.
6. Name the Bridge Domain `common-ASAv`. Use the drop-down list to select the `default` VRF in Tenant common.
7. Leave Forwarding set to Optimize and click Next.

8. No L3 Configurations changes are necessary. Click Next.

9. No Advanced/Troubleshooting changes are necessary. Click Finish.

10. Back in the Create Application Profile window, use the Domain drop-down list to select the `fpv-vc-vDS` VMM domain. Click Update.

11. Under EPGs, click the + to add an EPG.

12. Name the second EPG `Outside`. Use the drop-down list under BD to select the `common-ASAv` Bridge Domain.

13. Use the Domain drop-down list to select the `fpv-vc-vDS` VMM domain. Click Update.

14. Click Submit to complete creating the Application Profile.

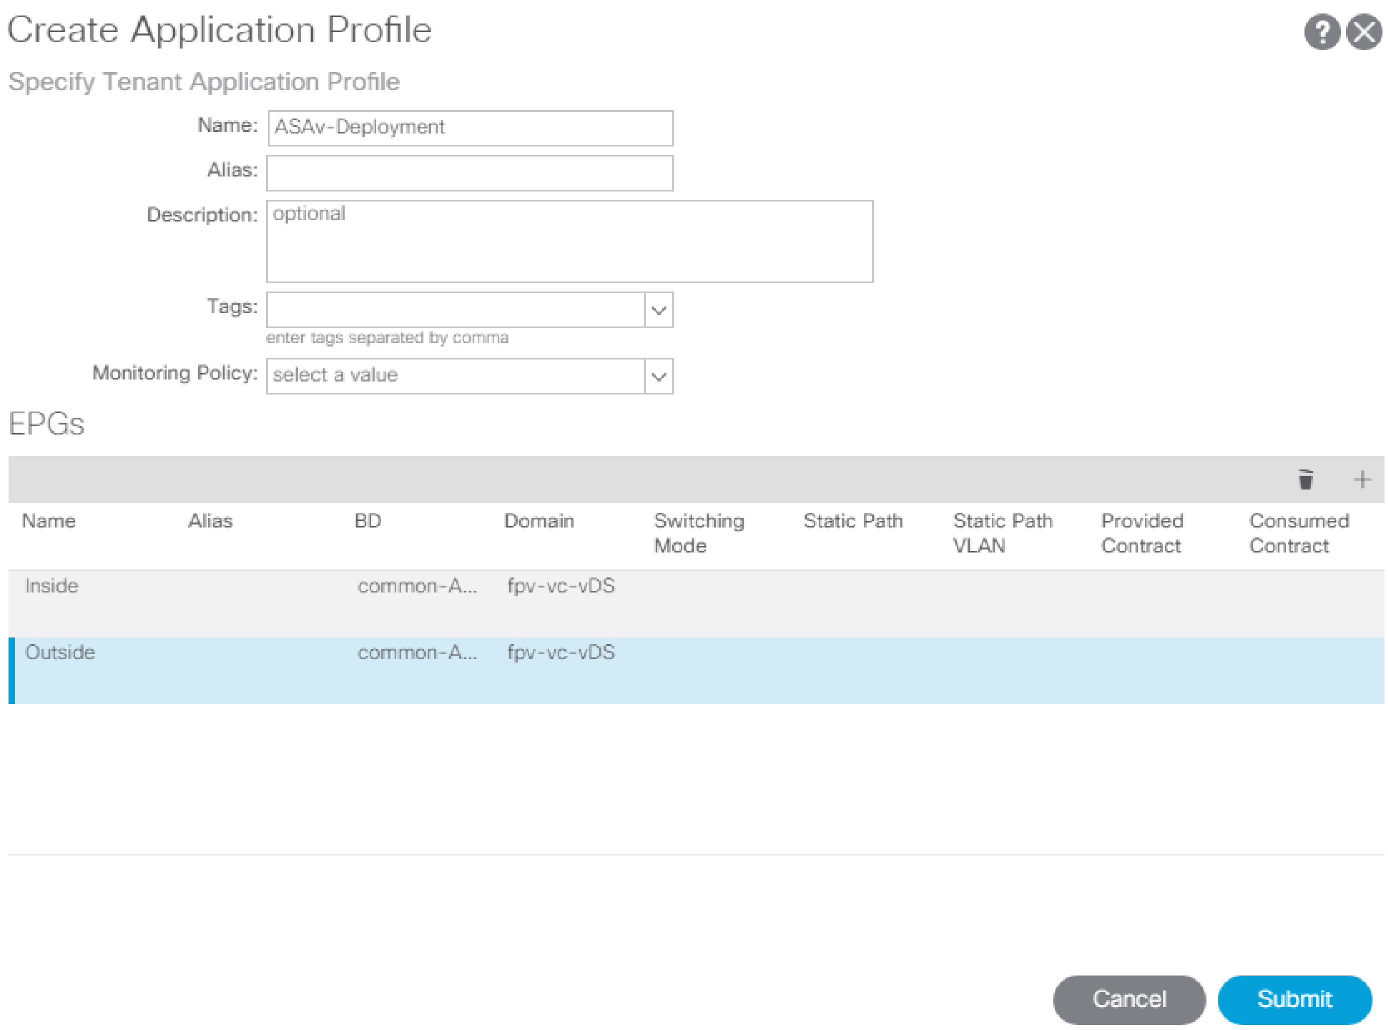

Create Application Profile

Specify Tenant Application Profile

Name: ASAv-Deployment

Alias:

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: select a value

EPGs

Name	Alias	BD	Domain	Switching Mode	Static Path	Static Path VLAN	Provided Contract	Consumed Contract
Inside		common-A...	fpv-vc-vDS					
Outside		common-A...	fpv-vc-vDS					

Create Tenant Firewall Outside Bridge Domain and Subnet

This section details setup of the bridge domain and associated subnet to be used for the ASAv firewall context **outside** interface. Since the firewall's outside interface is being connected to the Shared L3 Out that resides in Tenant common, the Outside Bridge Domain is also created in Tenant common.

The ASAv's Inside and Outside Interfaces must connect to different Bridge Domains when in use. Above, the two interfaces were placed in the same Bridge Domain, but when the L4-7 Service Graph is deployed, those interfaces will be moved to separate Bridge Domains.

APIC GUI

1. From the Cisco APIC 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-FPV-App-A-FW-Web-Out`.

5. Using the VRF drop-down list, select `common-FP-External`.

6. Leave Forwarding set to optimize and click Next.

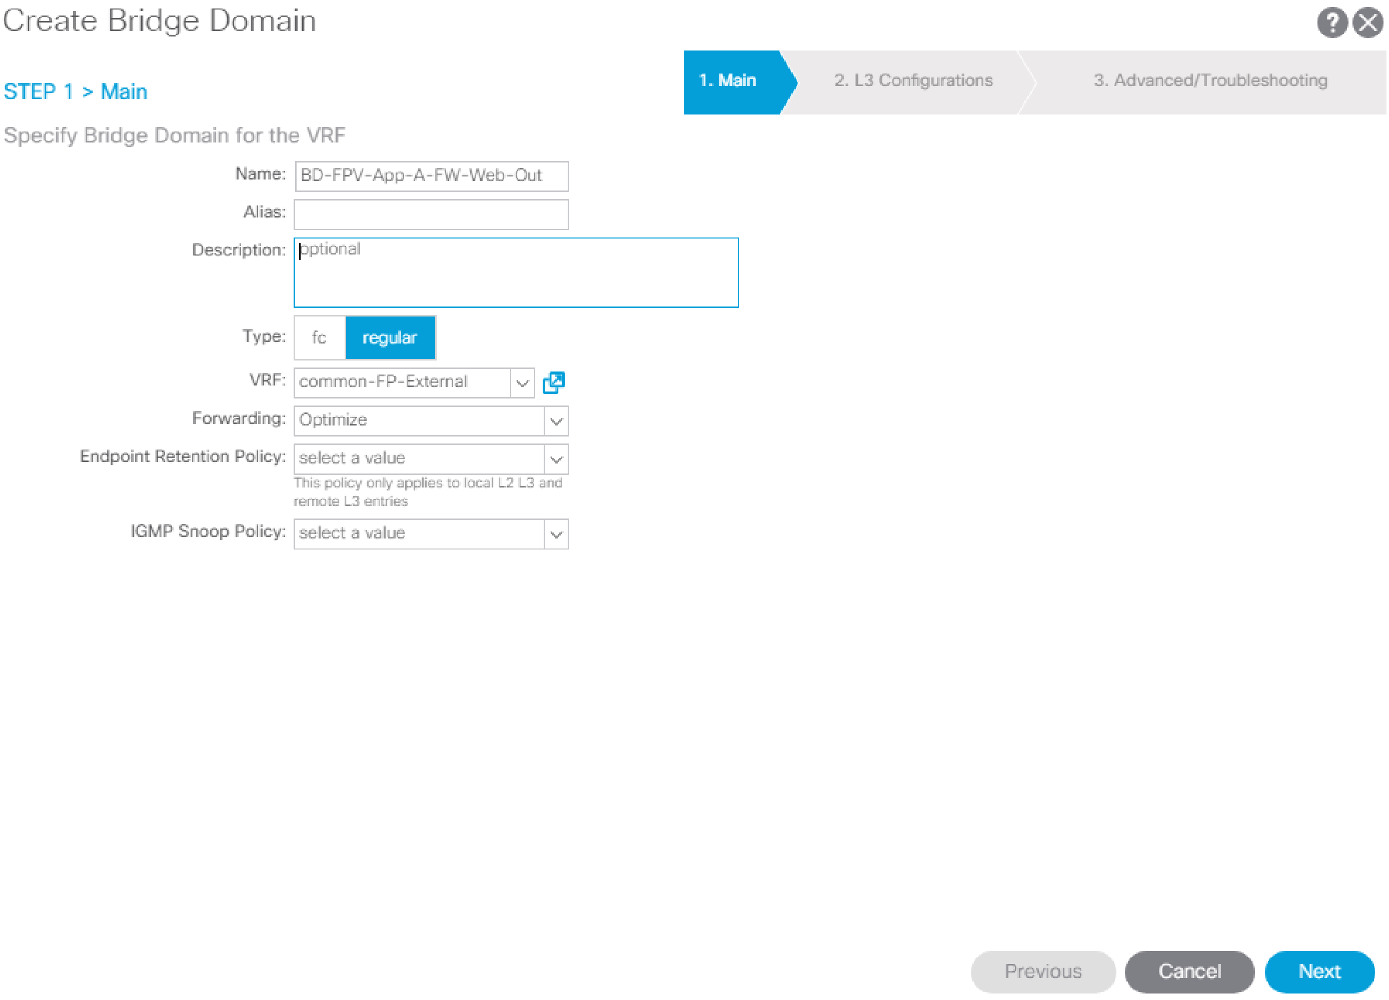

Create Bridge Domain

STEP 1 > Main

Specify Bridge Domain for the VRF

Name: BD-FPV-App-A-FW-Web-Out

Alias:

Description: optional

Type: fc regular

VRF: common-FP-External

Forwarding: Optimize

Endpoint Retention Policy: select a value

This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: select a value

7. Click the + to the right of Subnets to add a bridge domain subnet.

8. 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.

9. For Scope, select only Advertised Externally. Click OK to complete creating the subnet.

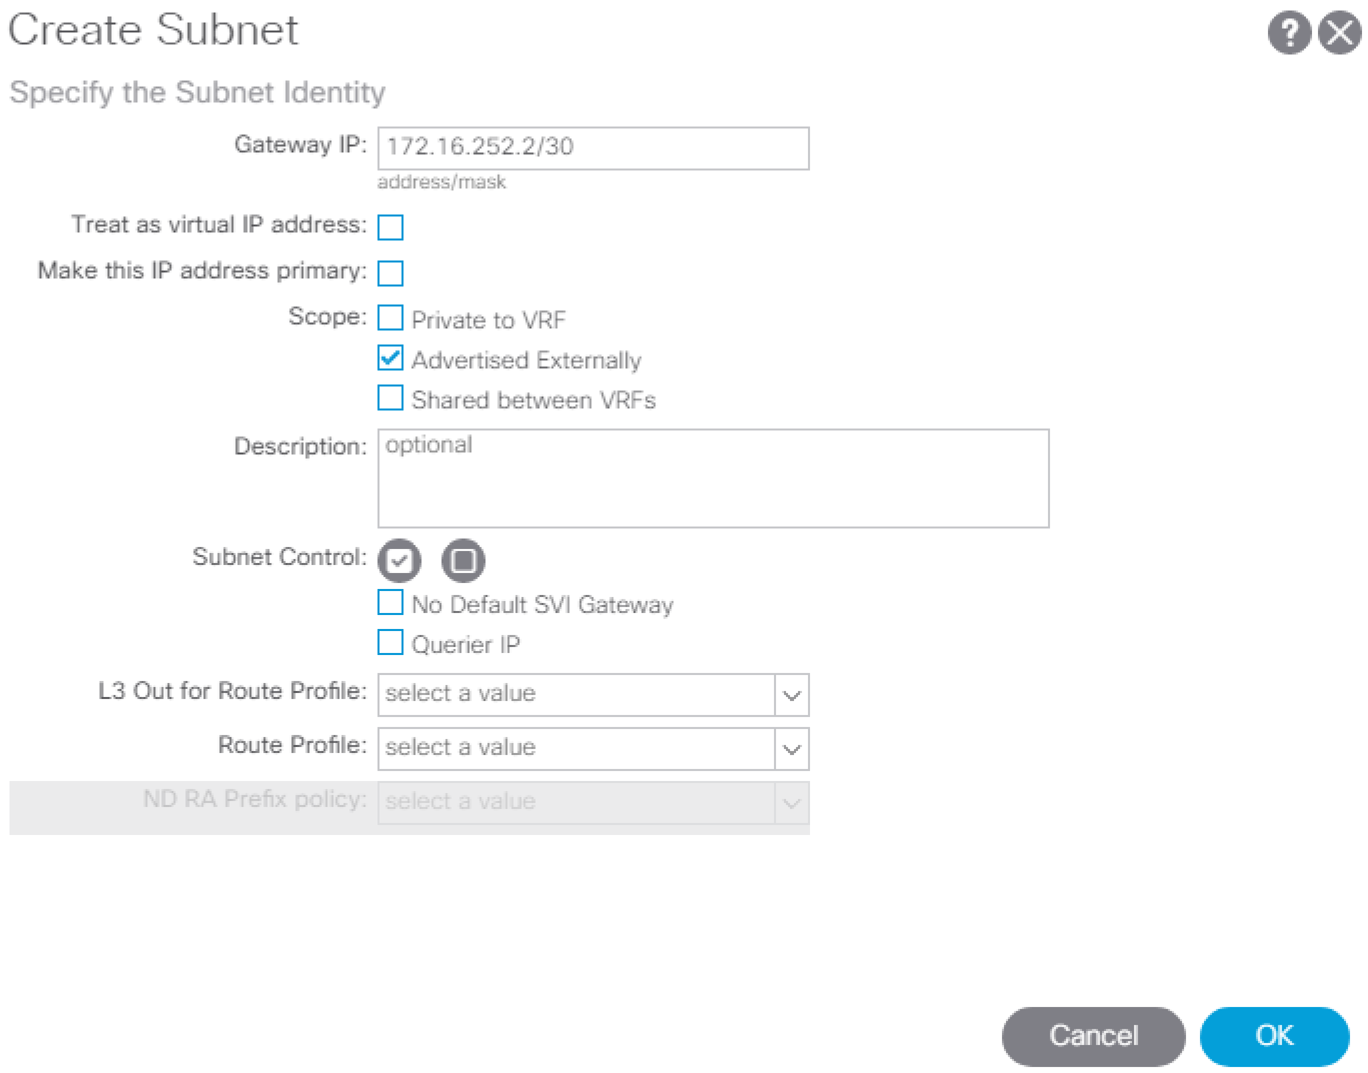

Create Subnet

Specify the Subnet Identity

Gateway IP: 172.16.252.2/30 (address/mask)

Treat as virtual IP address: □

Make this IP address primary: □

Scope:
- □ Private to VRF
- ☑ Advertised Externally
- □ Shared between VRFs

Description: optional

Subnet Control:
- □ No Default SVI Gateway
- □ Querier IP

L3 Out for Route Profile: select a value

Route Profile: select a value

ND RA Prefix policy: select a value

10. Click OK to complete creating the subnet.

11. Click the + to the right of Associated L3 Outs to add the Shared L3 Out.

12. Using the drop-down list, select `FP-Shared-L3-Out` and click Update.

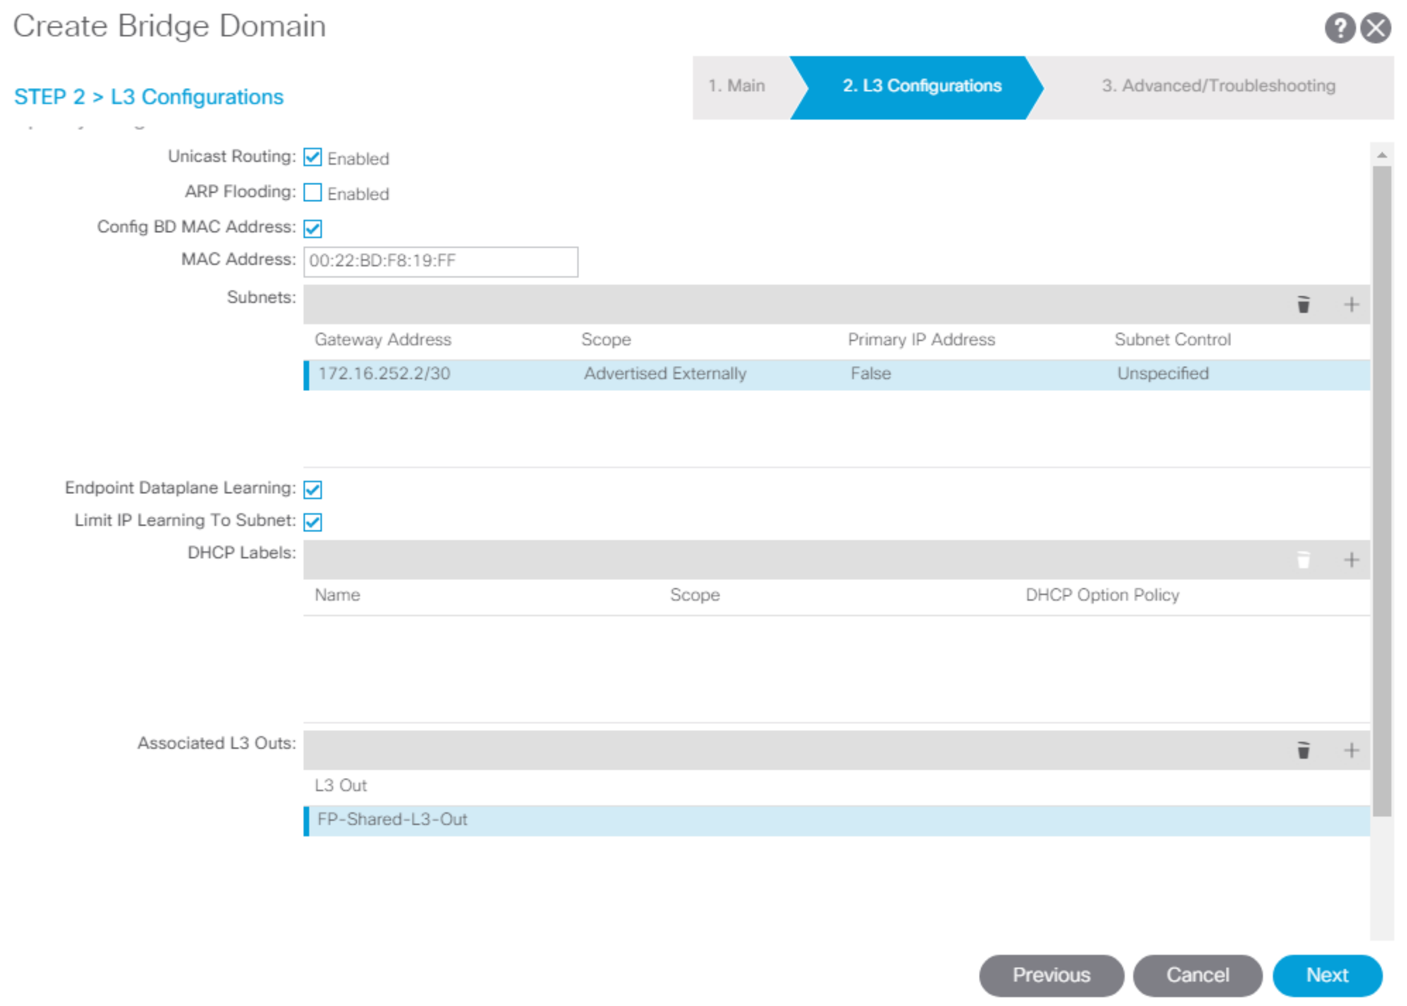

Create Bridge Domain

STEP 2 > L3 Configurations

Unicast Routing: ☑ Enabled

ARP Flooding: □ Enabled

Config BD MAC Address: ☑

MAC Address: 00:22:BD:F8:19:FF

Subnets:

Gateway Address	Scope	Primary IP Address	Subnet Control
172.16.252.2/30	Advertised Externally	False	Unspecified

Endpoint Dataplane Learning: ☑

Limit IP Learning To Subnet: ☑

DHCP Labels:

Name	Scope	DHCP Option Policy

Associated L3 Outs:

L3 Out
FP-Shared-L3-Out

13. Click Next.

14. Click Finish to complete creating the bridge domain.

Deploy ASAv VM from OVF

To setup of ASAv VM in the ACI Tenant, complete the following steps:

VMware vCenter Web Client

1. Download and expand the latest version of the Cisco ASAv .zip file from Cisco.com.

2. Using Windows Explorer, navigate to the folder with the files extracted from the .zip file.

3. Create a Not Needed folder and move the asav-esxi.mf and asav-esxi.ovf files to this folder. Five files should be left in this folder.

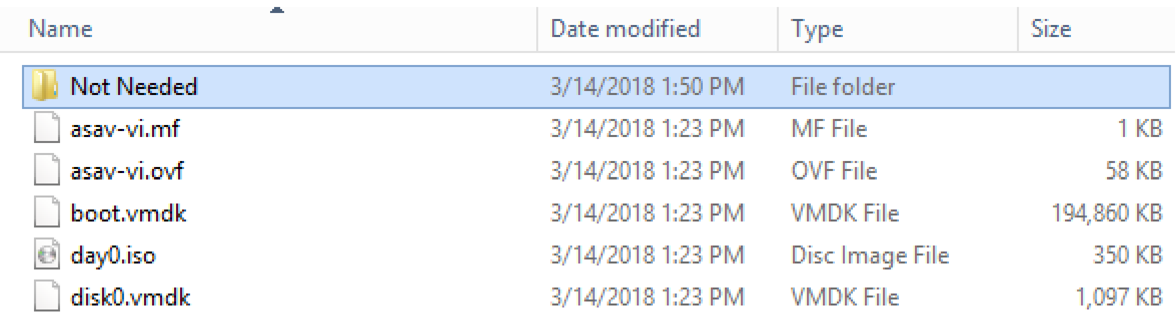

Name	Date modified	Type	Size
Not Needed	3/14/2018 1:50 PM	File folder	
asav-vi.mf	3/14/2018 1:23 PM	MF File	1 KB
asav-vi.ovf	3/14/2018 1:23 PM	OVF File	58 KB
boot.vmdk	3/14/2018 1:23 PM	VMDK File	194,860 KB
day0.iso	3/14/2018 1:23 PM	Disc Image File	350 KB
disk0.vmdk	3/14/2018 1:23 PM	VMDK File	1,097 KB

4. In the VMware vSphere Web Client, under Hosts and Clusters right-click the FPV-App-A cluster and select Deploy OVF Template.

5. In the Deploy OVF Template Window, click Browse and navigate to the folder with the extracted ASAv OVF files. Select the five files and select Open.

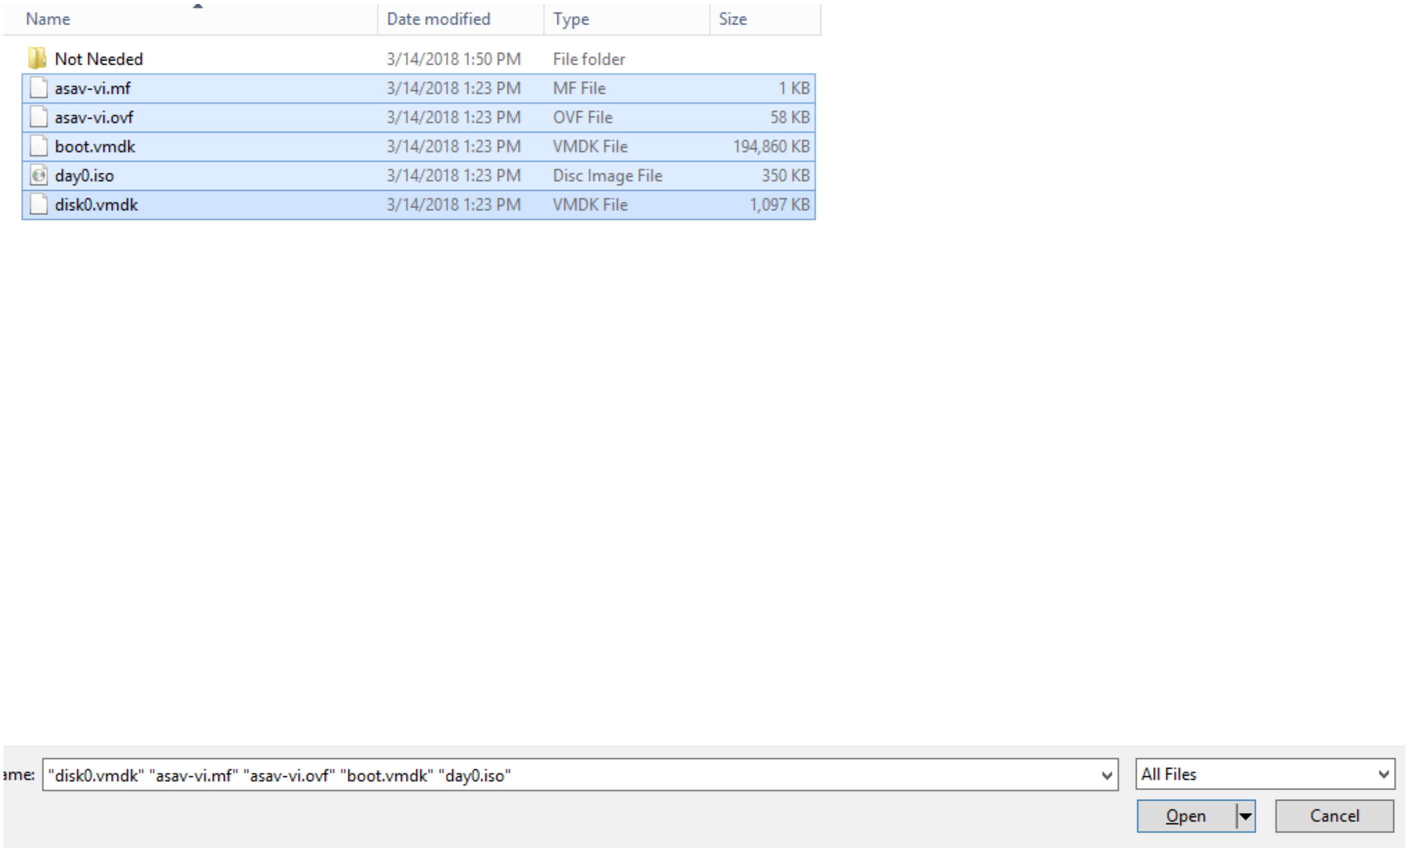

Name	Date modified	Type	Size
Not Needed	3/14/2018 1:50 PM	File folder	
asav-vi.mf	3/14/2018 1:23 PM	MF File	1 KB
asav-vi.ovf	3/14/2018 1:23 PM	OVF File	58 KB
boot.vmdk	3/14/2018 1:23 PM	VMDK File	194,860 KB
day0.iso	3/14/2018 1:23 PM	Disc Image File	350 KB
disk0.vmdk	3/14/2018 1:23 PM	VMDK File	1,097 KB

ame: "disk0.vmdk" "asav-vi.mf" "asav-vi.ovf" "boot.vmdk" "day0.iso"

All Files

6. Back in the Deploy OVF Template window, click Next.

7. Name the ASAv VM and select the folder to store the VM in. Click Next.

8. Select the ESXI cluster to deploy the VM on and click Next.

9. Review the warning and click Next.

10. Accept both License Agreements and click Next.

11. Select the ASAv configuration that you are licensed for and click Next.

12. Select the Thin provision virtual disk format and a datastore in the FPV-App-A tenant. Click Next.

13. Under Select networks, place GigabitEthernet0-0 in the common | ASAv-Deployment |Outside port group, Management0-0 in the FPV-App-A|MGMT|SVM-MGMT port group, and GigabitEthernet0-1 in the common | ASAv-Deployment | Inside port group. Place the remaining interfaces in the quarantine port group. Click Next.

In this sample deployment, the ASAv's management address is being placed in the Tenant's SVM-MGMT EPG which could have contracts allowing it to reach the Tenant's DB EPG and the Core-Services EPG. Placing the ASA's Management interface here allows it to be managed from either the DB EPG or from a Core-Services VM.

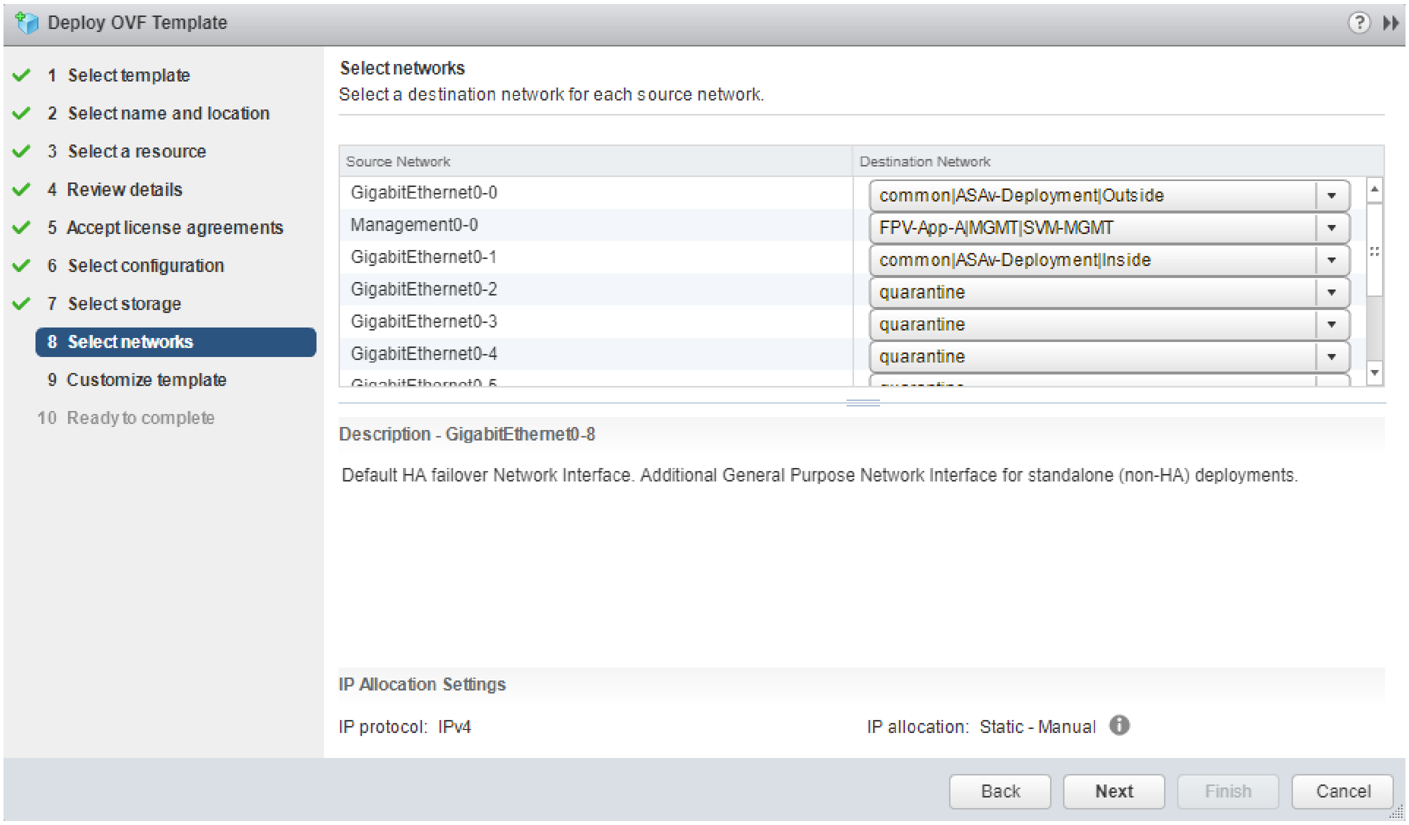

Select networks

Select a destination network for each source network.

Source Network	Destination Network
GigabitEthernet0-0	common|ASAv-Deployment|Outside
Management0-0	FPV-App-A|MGMT|SVM-MGMT
GigabitEthernet0-1	common|ASAv-Deployment|Inside
GigabitEthernet0-2	quarantine
GigabitEthernet0-3	quarantine
GigabitEthernet0-4	quarantine

Description - GigabitEthernet0-8

Default HA failover Network Interface. Additional General Purpose Network Interface for standalone (non-HA) deployments.

IP Allocation Settings

IP protocol: IPv4

IP allocation: Static - Manual

14. Under Customize template, fill in all relevant values including Firewall Mode –routed and click Next.

15. Review all of the deployment information and click Finish to deploy the Appliance.

16. Power on the ASAv, do any required configuration, and assign a license. A sample ASAv NAT configuration is shown in the appendix of this document.

It may be necessary to temporarily connect the EPG the ASAv Management Interface is in to the Shared L3 Out in order to contact the Cisco software license servers. Create Tenant L4-L7 Device and Service Graph

This section details setup of the L4-L7 Device and Service Graph in the ACI Tenant.

APIC GUI

1. From the Cisco APIC GUI, at the top, select Tenants > FPV-App-A.
2. On the left, expand Tenant FPV-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-FP-Shared-L3-Out contract and select Delete to remove this contract association.
5. On the left, expand Subnets and select the EPG subnet.
6. 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.
7. Click Submit to complete modifying the subnet. Click Submit Changes if a warning pops up.
8. On the left, expand Tenant FPV-App-A > Services > L4-L7.
9. Right-click Devices and select Create L4-L7 Devices.
10. In the Create L4-L7 Devices window, uncheck the Managed checkbox.
11. Name the Device FPV-App-A-ASAv. Select the Firewall Service Type.
12. Select the VIRTUAL Device Type.
13. For the VMM Domain, select fpv-vc-vDS. Select the Single Node View.
14. Leave Context Aware set to Single.
15. Under Device 1, use the drop-down list to select the fpv-vc/FPV-App-A-ASAv VM. Click + to add the first Device Interface.
16. Name the Device Interface Outside. Use the vNIC drop-down list to select Network adapter 2.
17. Click Update.
18. Click + to add the second Device Interface.
19. Name the Device Interface Inside. Use the vNIC drop-down list to select Network adapter 3.
20. Click Update.
21. Under Cluster, click + to add a Concrete Interface.
22. Name the interface Outside. Use the drop-down list to select Device1/Outside.
23. Click Update.
24. Under Cluster, click + to add the second Concrete Interface.

25. Name the interface Inside. Use the drop-down list to select Device1/Inside.

26. Click Update.

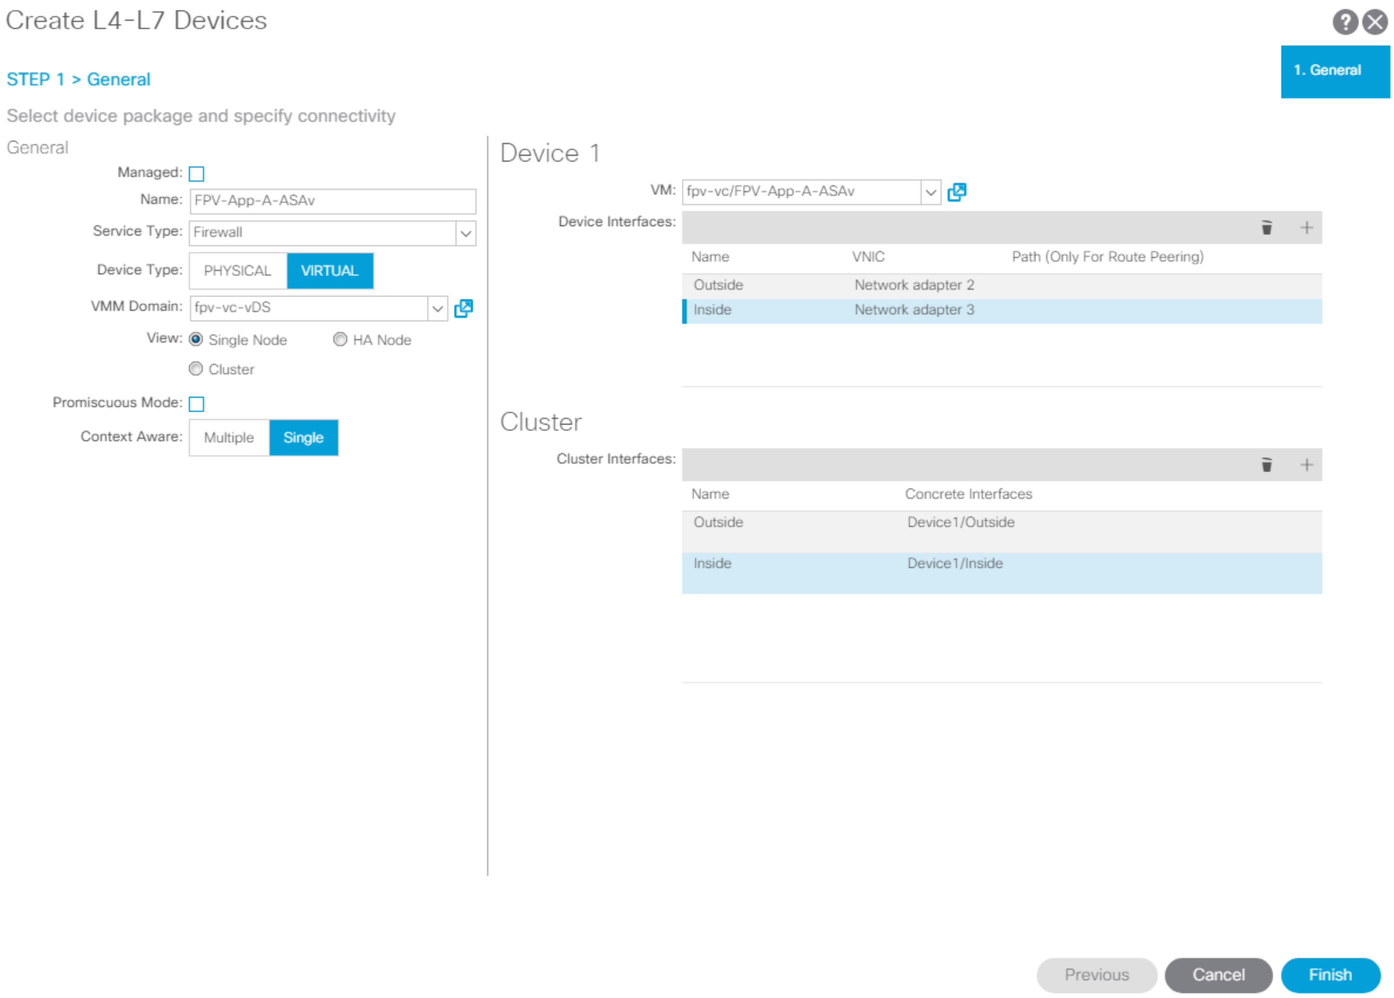

Create L4-L7 Devices

STEP 1 > General

Select device package and specify connectivity

General

Managed: □

Name: FPV-App-A-ASAv

Service Type: Firewall

Device Type: PHYSICAL [✓ VIRTUAL]

VMM Domain: fpv-vc-vDS

View: ● Single Node ○ HA Node ○ Cluster

Promiscuous Mode: □

Context Aware: Multiple [✓ Single]

Device 1

VM: fpv-vc/FPV-App-A-ASAv

Device Interfaces:

Name	VNIC	Path (Only For Route Peering)
Outside	Network adapter 2	
Inside	Network adapter 3	

Cluster

Cluster Interfaces:

Name	Concrete Interfaces
Outside	Device1/Outside
Inside	Device1/Inside

27. Click Finish to complete creating the L4-L7 Device.

28. Right-click Service Graph Templates and select Create L4-L7 Service Graph Template.

29. In the Create L4-L7 Service Graph Template window, name the Graph FPV-App-A-ASAv.

30. Make sure Graph Type is set to Create A New Graph.

31. On the left drag the FPV-App-A-ASAv (Firewall) icon to between the two EPGs on the right.

32. Select the Routed Firewall.

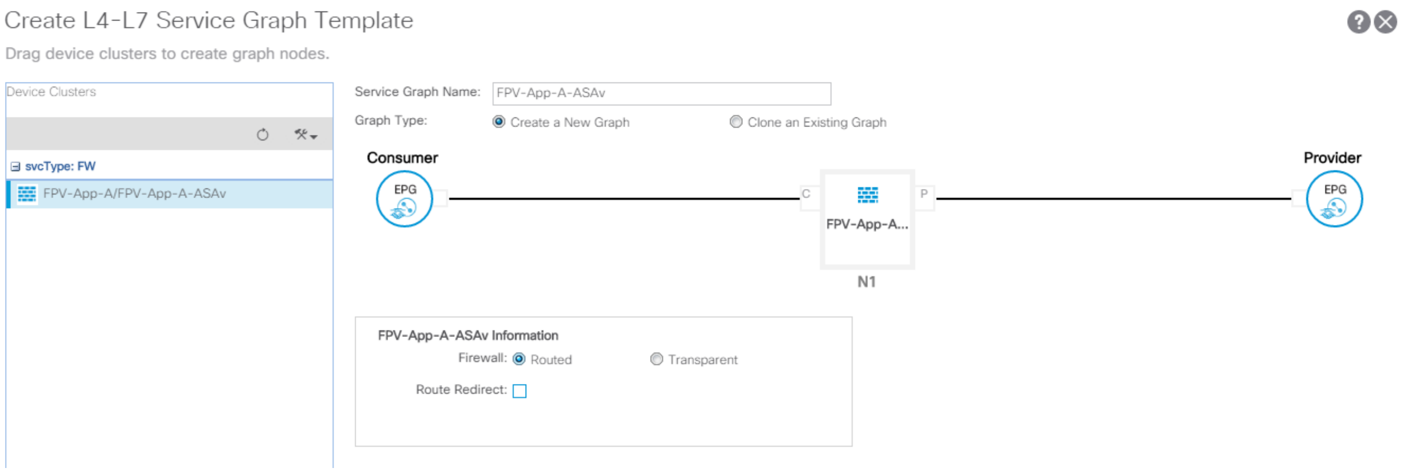

Create L4-L7 Service Graph Template

Drag device clusters to create graph nodes.

Device Clusters

svcType: FW

FPV-App-A/FPV-App-A-ASAv

Service Graph Name: FPV-App-A-ASAv

Graph Type: ● Create a New Graph ○ Clone an Existing Graph

FPV-App-A-ASAv Information

Firewall: ● Routed ○ Transparent

Route Redirect: □

33. Click Submit to complete creating the Service Graph Template.

34. On the left, expand Service Graph Templates and select the FPV-App-A-ASAv Template.

35. Right-click the FPV-App-A-ASAv 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 list to select common/FP-Shared-L3-Out/FP-Default-Route.

37. Use the Provider EPG drop-down list to select FPV-App-A/3-Tier-App/epg-Web.

These EPG selections place the firewall between the Shared-L3-Out and FPV-App-A Web EPGs.

38. Under Contract Information, leave Create A New Contract selected and name the contract Allow-FPV-App-A-FW-L3-Web

It is important that this contract have a unique name within the ACI Fabric.

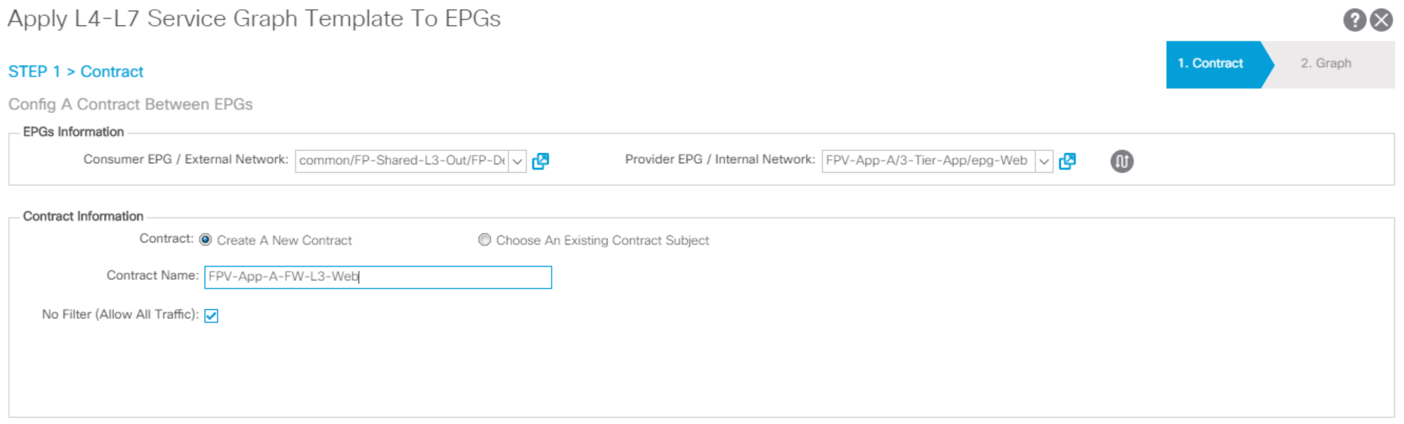

Apply L4-L7 Service Graph Template To EPGs

STEP 1 > Contract

Config A Contract Between EPGs

EPGs Information

Consumer EPG / External Network: common/FP-Shared-L3-Out/FP-De

Provider EPG / Internal Network: FPV-App-A/3-Tier-App/epg-Web

Contract Information

Contract: ● Create A New Contract ○ Choose An Existing Contract Subject

Contract Name: FPV-App-A-FW-L3-Web

No Filter (Allow All Traffic): ☑

39. Click Next.

40. Under the Consumer Connector, use the BD drop-down list to select common/BD-FPV-App-A-FW-Web-Out.

41. Under Consumer Connector use the Cluster Interface drop-down list to select Outside.

42. Under Provider Connector use the Cluster Interface drop-down list to select Inside.

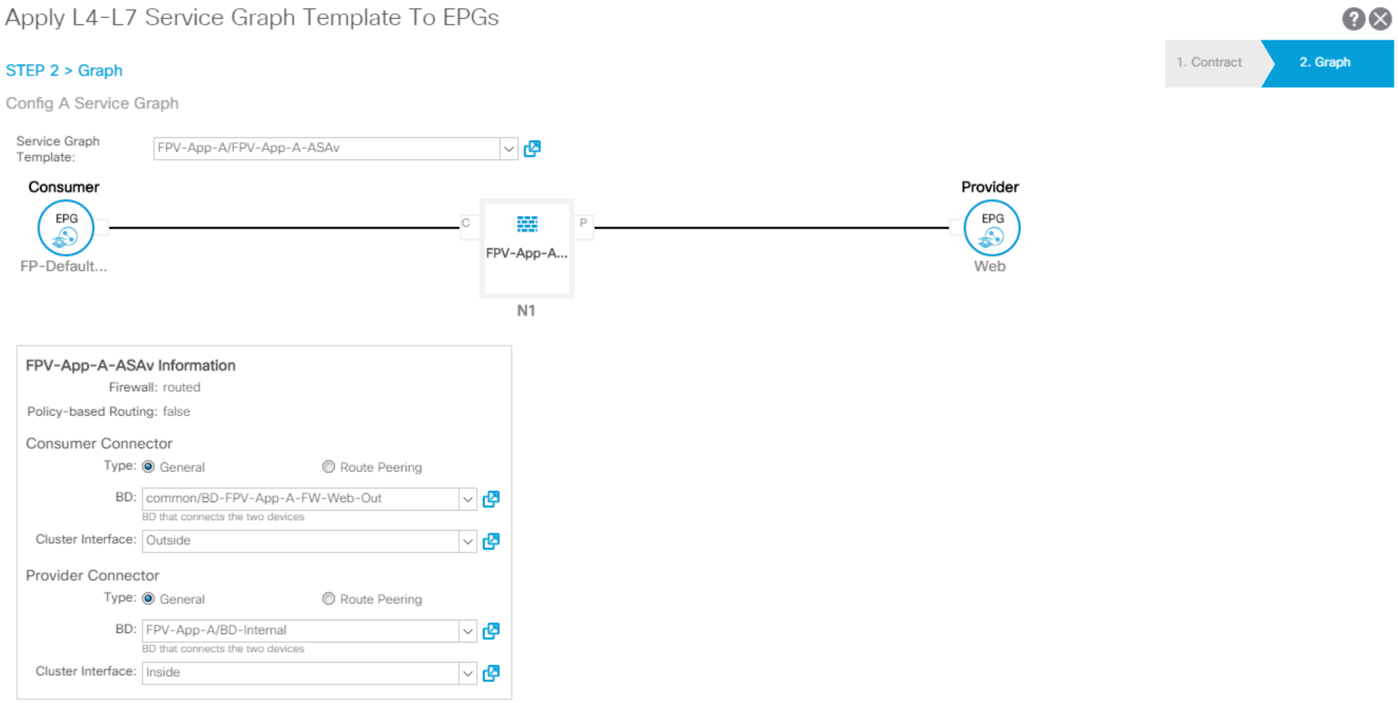

Apply L4-L7 Service Graph Template To EPGs

STEP 2 > Graph

Config A Service Graph

Service Graph Template: FPV-App-A/FPV-App-A-ASAv

FPV-App-A-ASAv Information

Firewall: routed

Policy-based Routing: false

Consumer Connector

Type: ● General ○ Route Peering

BD: common/BD-FPV-App-A-FW-Web-Out

BD that connects the two devices

Cluster Interface: Outside

Provider Connector

Type: ● General ○ Route Peering

BD: FPV-App-A/BD-Internal

BD that connects the two devices

Cluster Interface: Inside

43. Click Finish to compete applying the Service Graph Template.

44. On the left, expand Deployed Graph Instances and expand the deployed instance.

45. Select Function Node – N1.

46. Verify that the Function Connectors display values for Encap and Class ID.

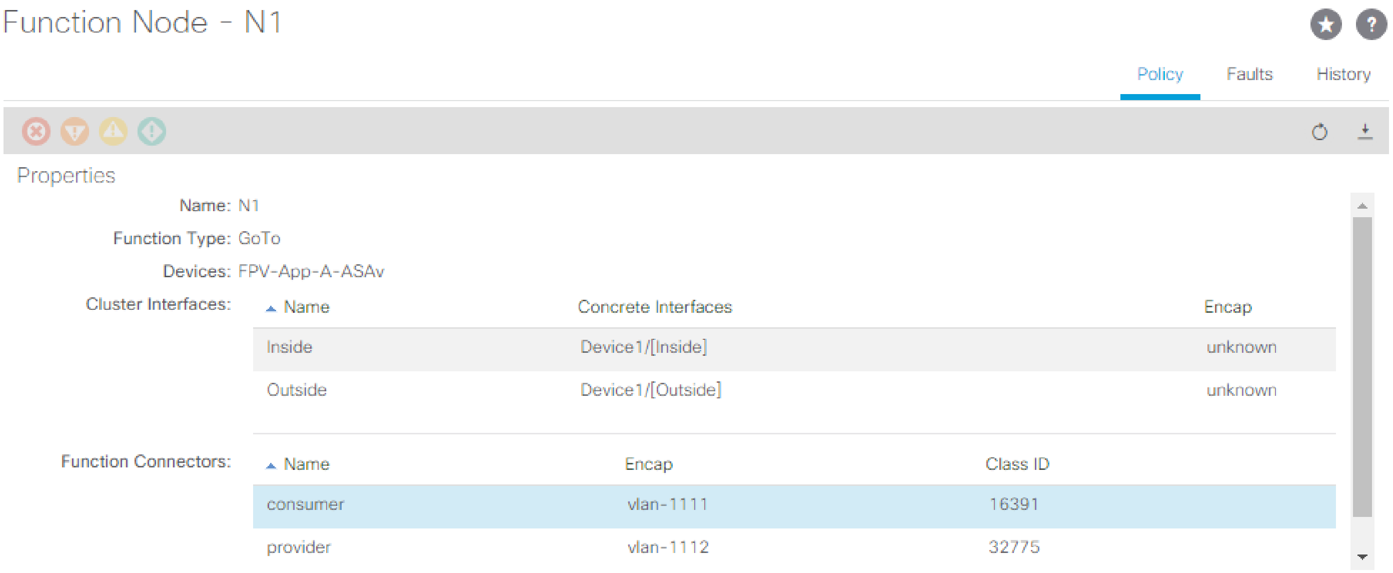

Function Node - N1

Properties

Name: N1

Function Type: GoTo

Devices: FPV-App-A-ASAv

Cluster Interfaces:

Name	Concrete Interfaces	Encap
Inside	Device1/[Inside]	unknown
Outside	Device1/[Outside]	unknown

Function Connectors:

Name	Encap	Class ID
consumer	vlan-1111	16391
provider	vlan-1112	32775

47. Configure the ASAv firewall device context as a routed firewall with NAT from the Inside to the Outside interface. A sample configuration is in the Appendix. The outside interface should be configured in the subnet entered in the BD-FPV-App-A-FW-Web-Out bridge domain. The ASAv's Outside default gateway should be the

bridge domain’s gateway address. The Inside interface should be configured in the subnet entered in the FPV-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.

48. For VMs with interfaces in the Web EPG, set the default gateway to the ASAv’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 ASAv firewall for the second ACI tenant.

Appendix - FC Solution

This section details the configuration steps for the Cisco UCS 6332-16UP Fabric Interconnects (FI) in a design that will support direct connectivity to the NetApp AFF using 16 Gb/s Fibre Channel.

Figure 5 shows the VMware vSphere 6.5U1 built on FlexPod components and the network connections for a configuration with the Cisco UCS 6332-16UP Fabric Interconnects with storage FC connections directly connected to the fabric interconnect. This design has 40Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and C-Series rackmounts and the Cisco UCS Fabric Interconnect, and between the Cisco UCS Fabric Interconnects and Cisco Nexus 9000 switches. This design also has two 16Gb FC connections between each Cisco UCS Fabric Interconnect and NetApp AFF A300. FC zoning is done in the Cisco UCS Fabric Interconnect. This infrastructure is deployed to provide FC-booted hosts with file-level and block-level access to shared storage with use cases that do not require the Cisco MDS SAN connectivity or scale.

Figure 5 FlexPod with Cisco UCS 6332-16UP Fabric Interconnects and Cisco UCS Direct Connect SAN

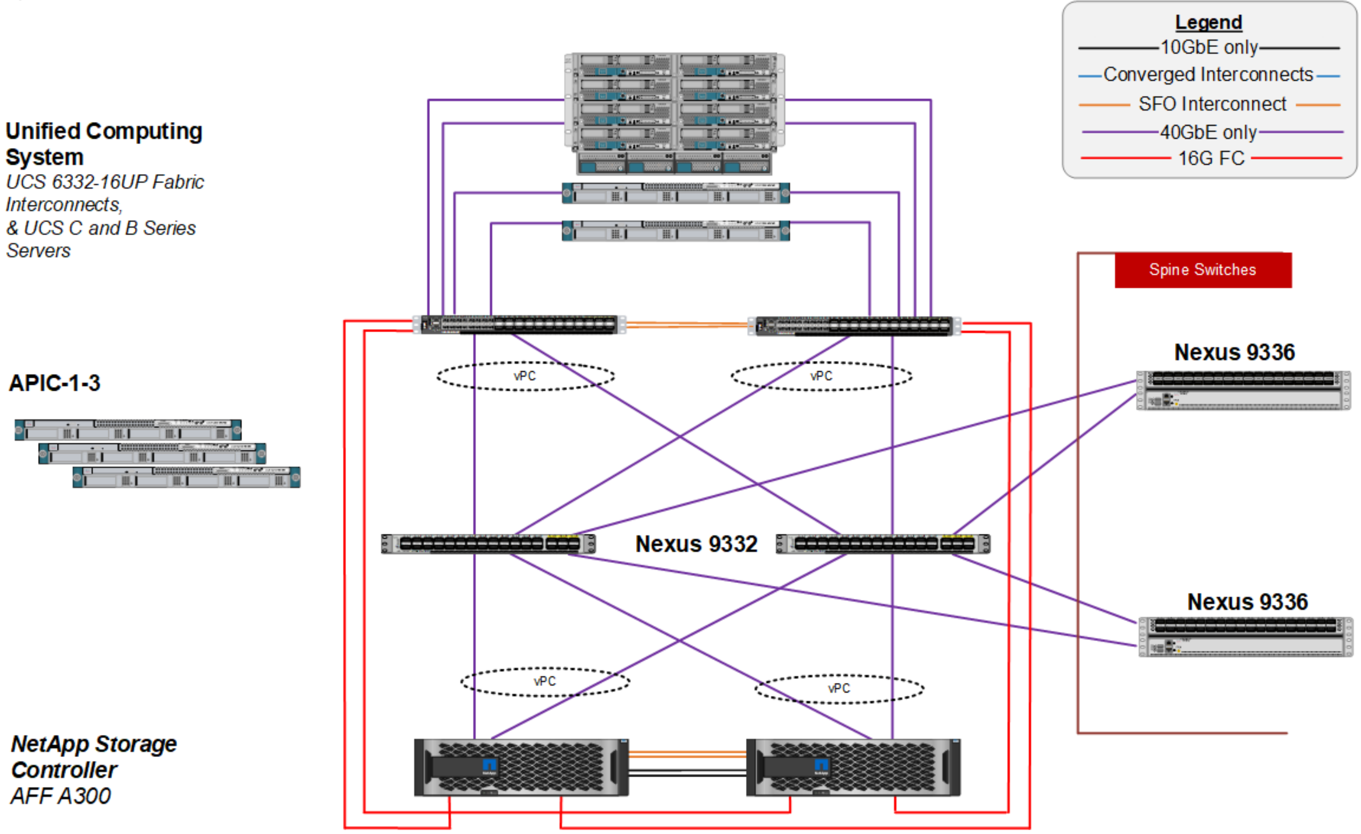

This Appendix only details the delta configuration required for the Storage, Cisco UCS, and VMware setup. As mentioned in previous sections of this document, any iSCSI steps can be skipped if iSCSI is not being implemented.

Storage Configuration

Set Onboard Unified Target Adapter 2 Port Personality

In order to use FC storage targets, FC ports must be configured on the storage. To set the personality of the onboard unified target adapter 2 (UTA2), complete the following steps for both controllers in an HA pair:

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

```
ucadmin show
                       Current  Current    Pending  Pending    Admin
Node          Adapter  Mode     Type       Mode     Type       Status
------------  -------  -------  ---------  -------  ---------  -----------
<st-node01>
              0e       fc       target     -        -          online
<st-node01>
              0f       fc       target     -        -          online
<st-node01>
              0g       fc       target     -        -          online
<st-node01>
              0h       fc       target     -        -          online
<st-node02>
              0e       fc       target     -        -          online
<st-node02>
              0f       fc       target     -        -          online
<st-node02>
              0g       fc       target     -        -          online
<st-node02>
              0h       fc       target     -        -          online
8 entries were displayed.
```

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

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

The ports must be offline to run this command. To take an adapter offline, run the `fcp adapter modify –node <home-node-of-the-port> -adapter <port-name> -state down` command. Ports must be converted in pairs (for example, `0e` and `0f`).

After conversion, a reboot is required. After reboot, bring the ports online by running `fcp adapter modify –node <home-node-of-the-port> -adapter <port-name> -state up`.

Add FCP Storage Protocol to Infrastructure SVM

Run the following command to add the FCP storage protocol to the Infrastructure SVM. It is assumed that an FCP license has been installed on each cluster node:

```
vserver add-protocols -vserver FPV-Foundation-SVM -protocols fcp
vserver show -fields allowed-protocols
```

Create FCP Storage Protocol in Infrastructure SVM

Run the following command to create the FCP storage protocol in the Infrastructure SVM. It is assumed that an FCP license has been installed on each cluster node:

```
fcp create -vserver FPV-Foundation-SVM
fcp show
```

Create FC LIFs

Run the following commands to create four FC LIFs (two on each node) in the Infrastructure SVM by using the previously configured FC ports:

```
network interface create -vserver FPV-Foundation-SVM -lif fcp_lif01a -role data -data-protocol fcp -
home-node <st-node01> -home-port 0e -status-admin up

network interface create -vserver FPV-Foundation-SVM -lif fcp_lif01b -role data -data-protocol fcp -
home-node <st-node01> -home-port 0f -status-admin up

network interface create -vserver FPV-Foundation-SVM -lif fcp_lif02a -role data -data-protocol fcp -
home-node <st-node02> -home-port 0e -status-admin up

network interface create -vserver FPV-Foundation-SVM -lif fcp_lif02b -role data -data-protocol fcp -
home-node <st-node02> -home-port 0f -status-admin up

network interface show -vserver FPV-Foundation-SVM -lif fcp*
```

Server Configuration

This section details the delta steps in the Cisco UCS setup to provide FC-based boot and storage.

Configure FC Unified Ports (UP) on Cisco UCS Fabric Interconnects

In order to use Fiber Channel Storage Ports for storage directly connected to the Cisco UCS fabric interconnects, 10G Ethernet ports need to be converted to 16G FC ports.

To convert the first six ports of the UCS 6322-16 UP to FC, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. In the center pane, select Configure Unified Ports. Click Yes to proceed.
4. The 6332-16UP requires ports to be converted in groups of 6 ports from the left. To convert the first six ports, move the slider to the right until ports 1-6 are highlighted.

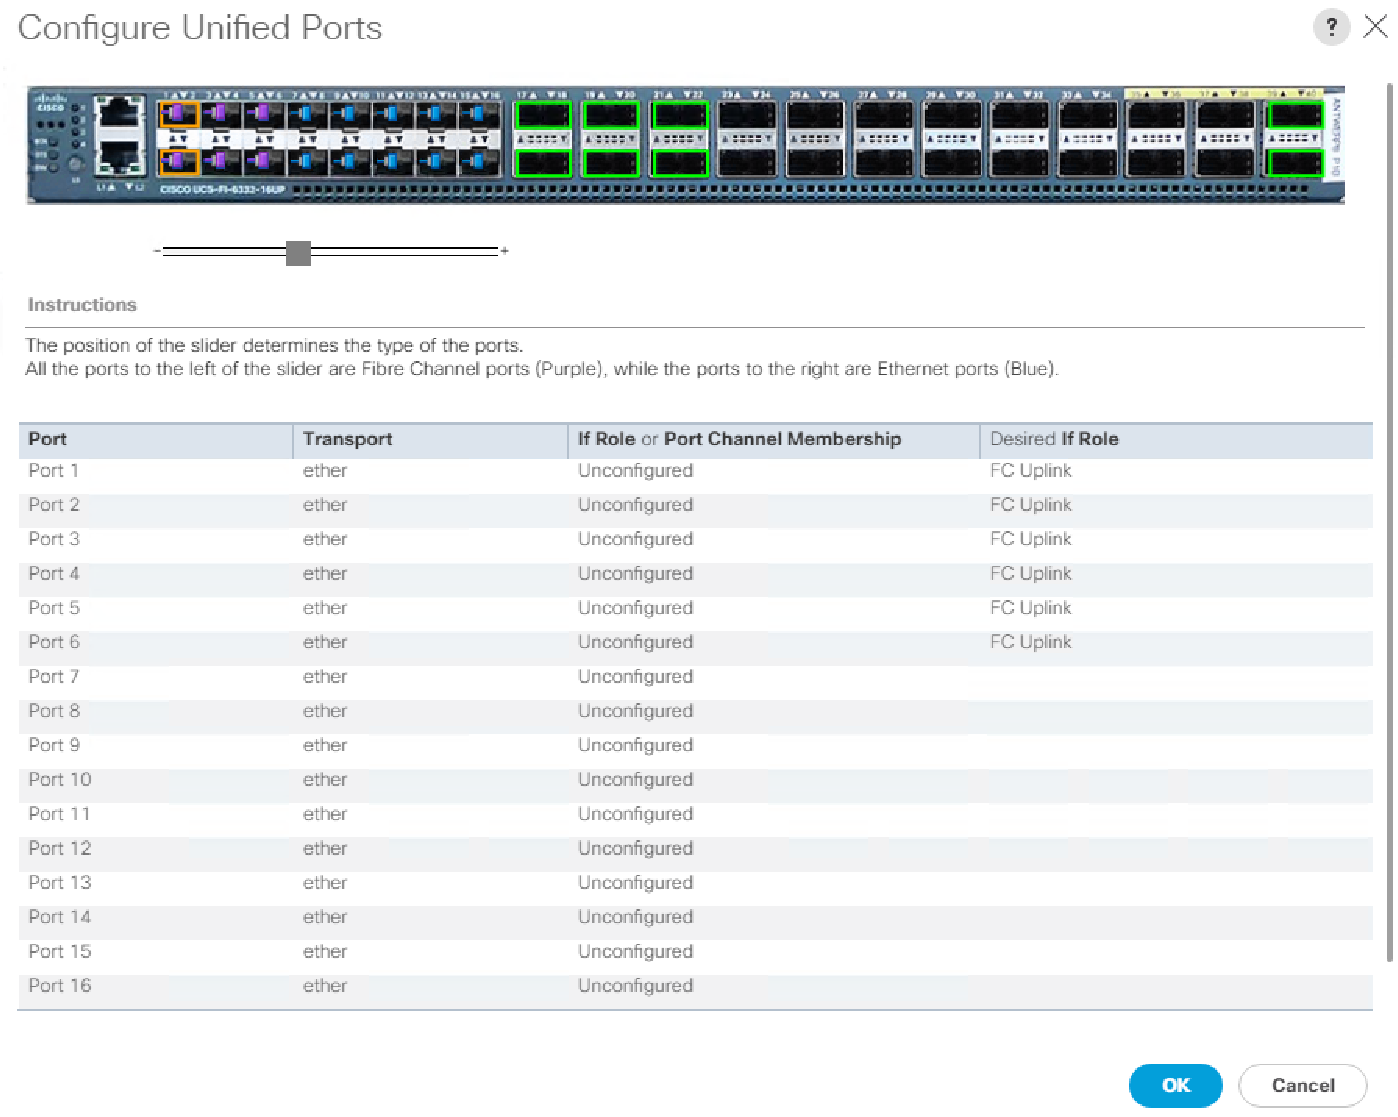

Configure Unified Ports

Instructions

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

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	FC Uplink
Port 6	ether	Unconfigured	FC Uplink
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Unconfigured	

5. Click OK, then click Yes, then click OK to convert the ports. The Fabric Interconnect will reboot. If this was the primary fabric interconnect, you will need to reconnect to Cisco UCS Manager. Wait until the reboot is complete and the FI is back in the UCS domain cluster. This can be checked by logging into the FI's CLI and typing `show cluster state`. Wait until the "HA Ready" state appears.

6. Repeat this process to convert the first six ports of FI B.

Place Cisco UCS Fabric Interconnects in Fiber Channel Switching Mode

In order to use Fiber Channel Storage Ports for storage directly connected to the Cisco UCS fabric interconnects, the fabric interconnects must be changed from fiber channel end host mode to fiber channel switching mode.

To place the fabric interconnects in fiber channel switching mode, complete the following steps:

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

2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).

3. In the center pane, select set FC Switching Mode. Click Yes and OK for the confirmation message.

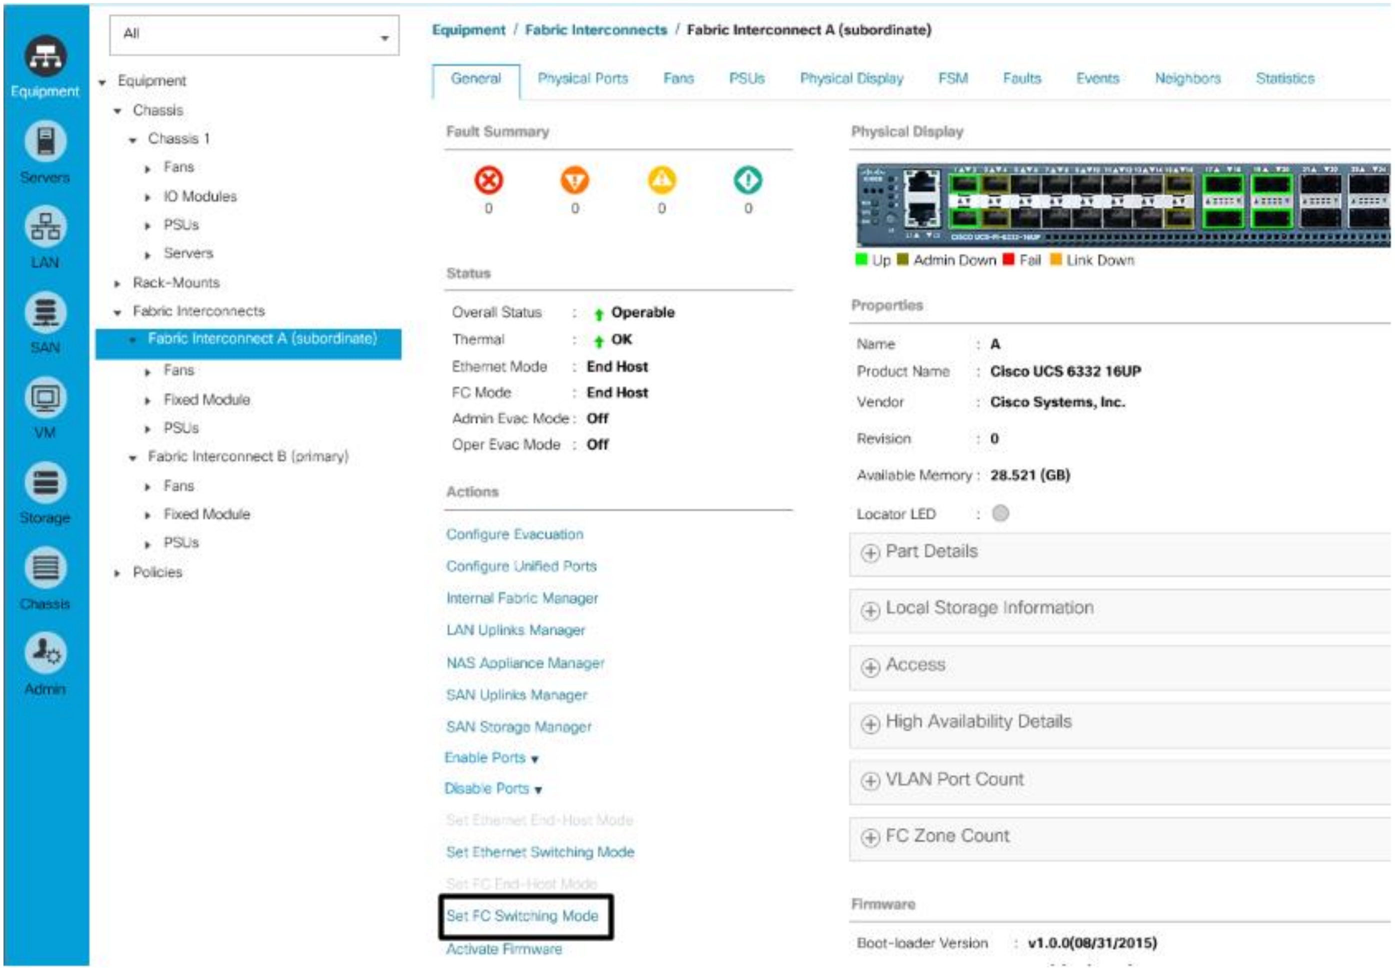

Overall Status : Operable
Thermal : OK
Ethernet Mode : End Host
FC Mode : End Host
Admin Evac Mode : Off
Oper Evac Mode : Off

Name : A
Product Name : Cisco UCS 6332 16UP
Vendor : Cisco Systems, Inc.
Revision : 0
Available Memory : 28.521 (GB)
Locator LED :

Boot-loader Version : v1.0.0(08/31/2015)

4. Wait for both Fabric Interconnects to reboot by monitoring the console ports and log back into Cisco UCS Manager.

It may be necessary to go to the Pending Activities window and select Reboot Fabric Interconnect to reboot the Primary Fabric Interconnect.

Create Storage VSANs

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

1. In Cisco UCS Manager, click the SAN icon on the left.

In this step, two VSANs are created, one for Fabric A and one for Fabric B.

2. Select and expand SAN > Storage Cloud.

3. Right-click VSANs.

4. Select Create Storage VSAN.

5. Enter `VSAN-A` as the name of the VSAN to be used for Fabric A.

6. Set FC Zoning to Enabled.

7. Select Fabric A.

8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

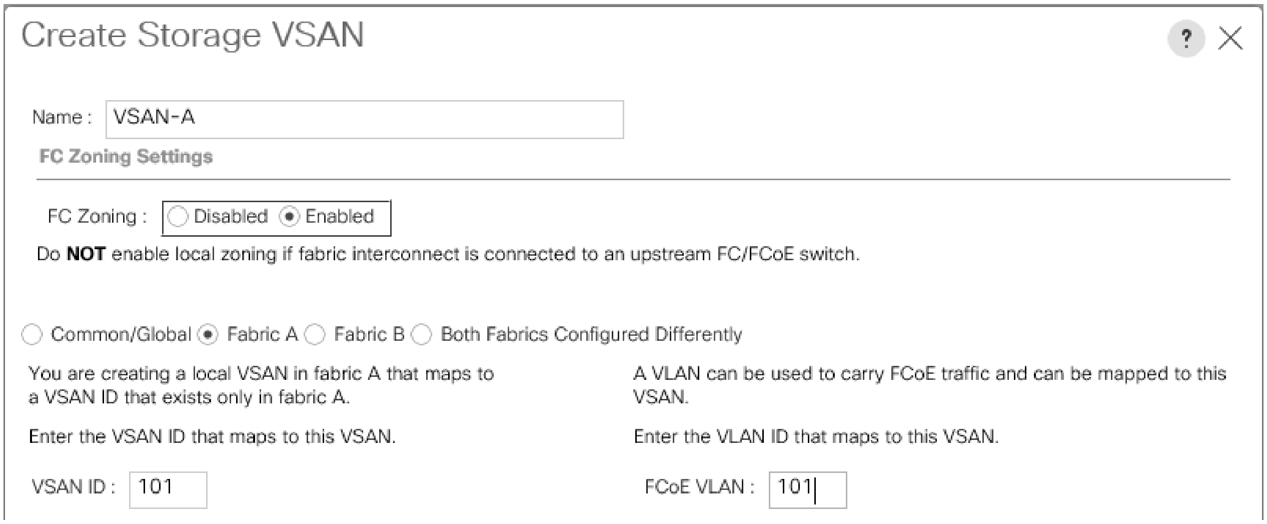

Create Storage VSAN

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

9. Click OK and then click OK again.

10. Under Storage Cloud, right-click VSANs.

11. Select Create Storage VSAN.

12. Enter `VSAN-B` as the name of the VSAN to be used for Fabric B.

13. Set FC Zoning to Enabled.

14. Select Fabric B.

15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID for Fabric B. It is recommended use the same ID for both parameters and to use something other than 1.

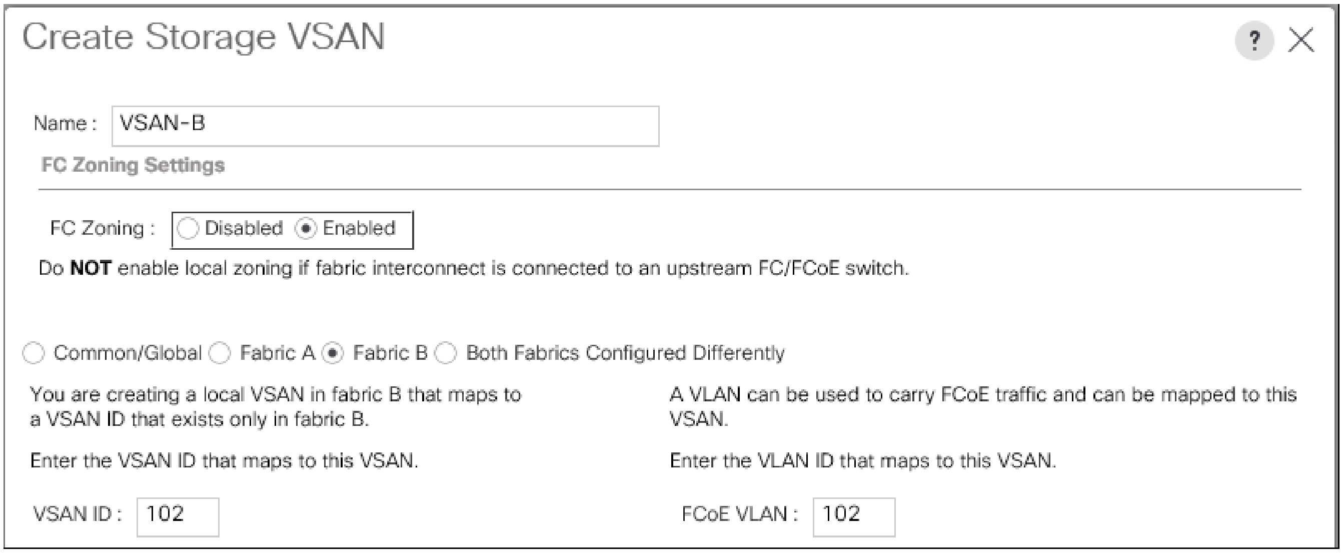

Create Storage VSAN

Name : VSAN-B

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 B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 102

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 : 102

16. Click OK, and then click OK again

Configure FC Storage Ports

To configure the necessary FCoE Storage port for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select and expand Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > FC Ports
3. Select the ports (1 and 2 for this document) that are connected to the NetApp array, right-click them, and select “Configure as FC Storage Port”
4. Click Yes to confirm and then click OK.
5. Verify that the ports connected to the NetApp array are now configured as FC Storage ports.
6. Select and expand Equipment > Fabric Interconnects > Fabric Interconnect B > Fixed Module > FC Ports
7. Select the ports (1 and 2 for this document) that are connected to the NetApp array, right-click them, and select “Configure as FC Storage Port.”
8. Click Yes to confirm and then click OK.
9. Verify that the ports connected to the NetApp array are now configured as FC Storage ports.

Assign VSANs to FC Storage Ports

To assign storage VSANs to FC Storage Ports, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select and expand SAN > Storage Cloud.
3. Expand Fabric A and Storage FC Interfaces.
4. Select the first FC Interface (1/1).

5. For User Label, enter the storage controller name and port. Click Save Changes, then click OK.

6. Use the drop-down list to select VSAN `VSAN-A`. Click Save Changes, then click OK.

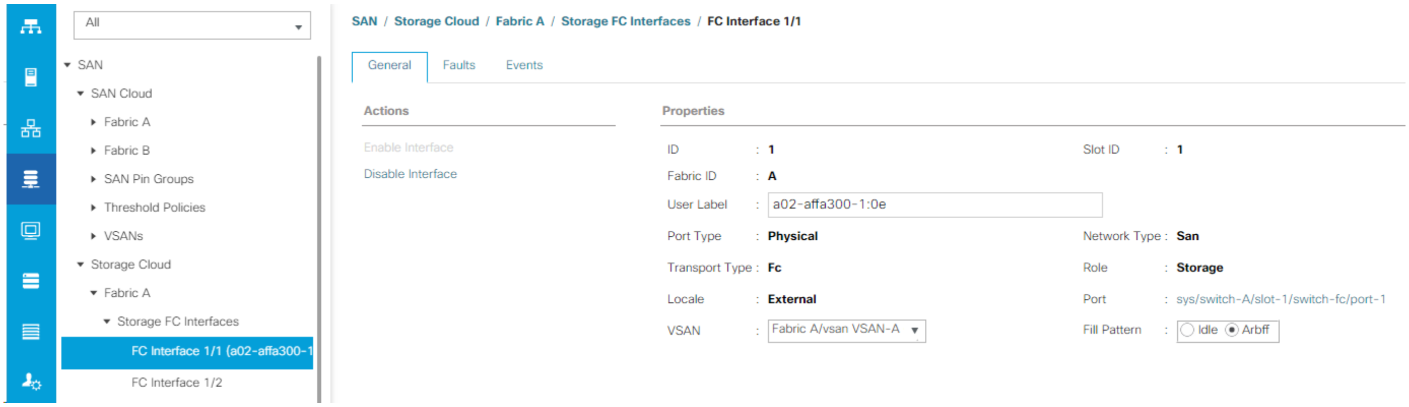

SAN / Storage Cloud / Fabric A / Storage FC Interfaces / FC Interface 1/1

ID : 1
Slot ID : 1
Fabric ID : A
User Label : a02-affa300-1:0e
Port Type : Physical
Network Type : San
Transport Type : Fc
Role : Storage
Locale : External
Port : sys/switch-A/slot-1/switch-fc/port-1
VSAN : Fabric A/vsan VSAN-A
Fill Pattern : ○ Idle ● Arbff

7. Select the second FC Interface (1/2).

8. For User Label, enter the storage controller name and port. Click Save Changes, then click OK.

9. Use the drop-down list to select VSAN `VSAN-A`. Click Save Changes, then click OK.

10. Expand Fabric B and Storage FC Interfaces.

11. Select the first FC Interface (1/1)

12. For User Label, enter the storage controller name and port. Click Save Changes, then click OK.

13. Use the drop-down list to select VSAN `VSAN-B`. Click Save Changes, then click OK.

14. Select the second FC Interface (1/2).

15. For User Label, enter the storage controller name and port. Click Save Changes, then click OK.

16. Use the drop-down list to select VSAN `VSAN-B`. Click Save Changes, then click OK.

Create a WWNN Pool for FC Boot

To configure the necessary WWNN pool for the Cisco UCS environment, complete the following steps on Cisco UCS Manager.

1. Select SAN on the left.

2. Select and expand Pools > root.

3. Right-click WWNN Pools under the root organization.

4. Select Create WWNN Pool to create the WWNN pool.

5. Enter `WWNN-Pool` for the name of the WWNN pool.

6. Optional: Enter a description for the WWNN pool.

7. Select **Sequential** for Assignment Order.

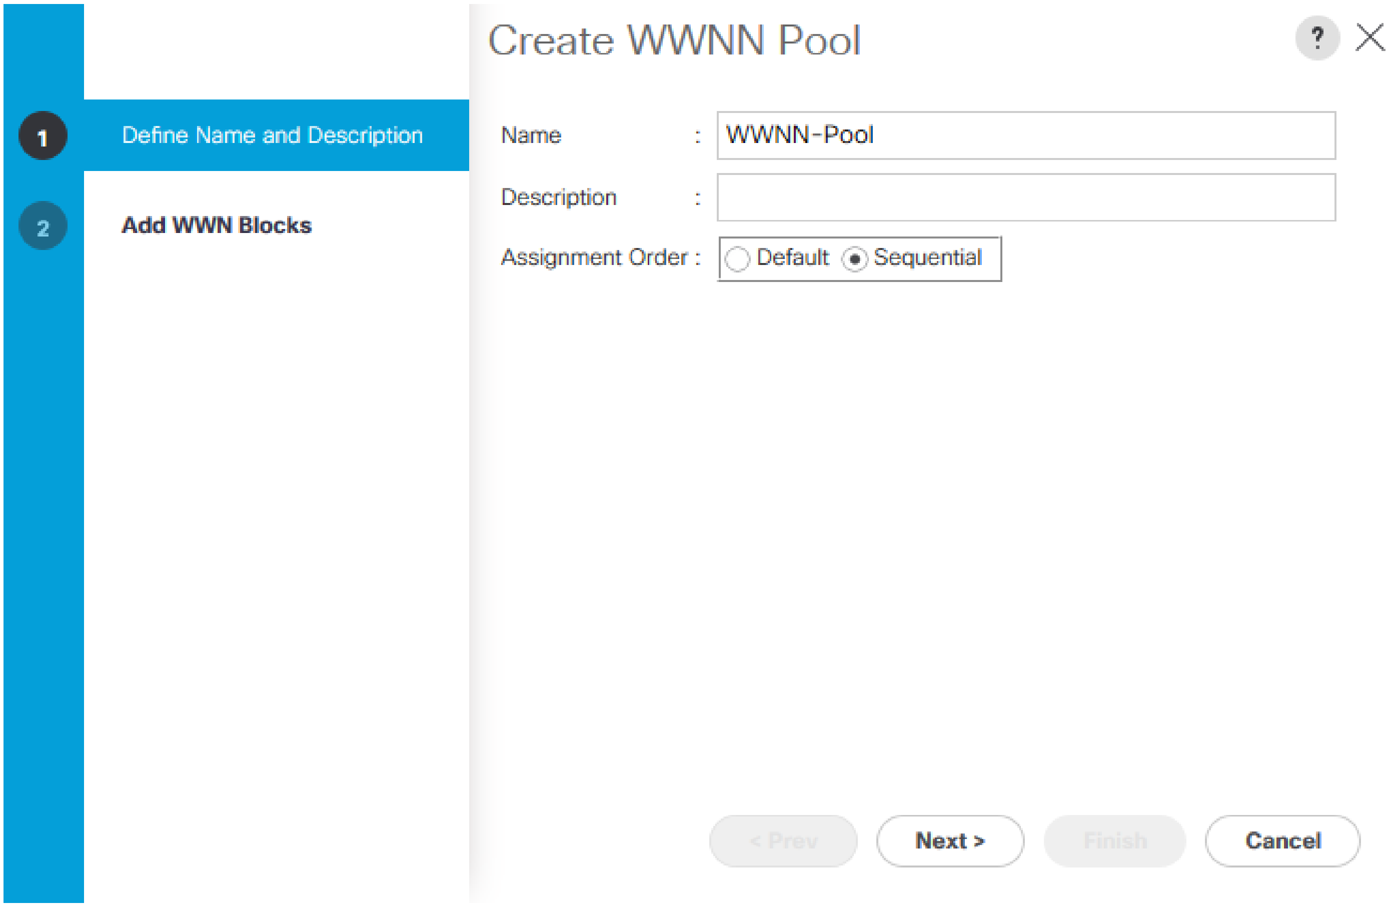

Name : WWNN-Pool

Description :

Assignment Order : ○ Default ● Sequential

8. Click Next.

9. Click Add.

10. Modify the From field as necessary for the UCS Environment and click OK.

Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the UCS domain. Within the From field in our example, the 6th octet was changed from 00 to A2 to represent as identifying information for this being in the UCS 6332 in the 2nd cabinet of the A row. Also, when having multiple UCS domains sitting in adjacency, it is important that these blocks, the WWNN, WWPN, and MAC hold differing values between each set.

---

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

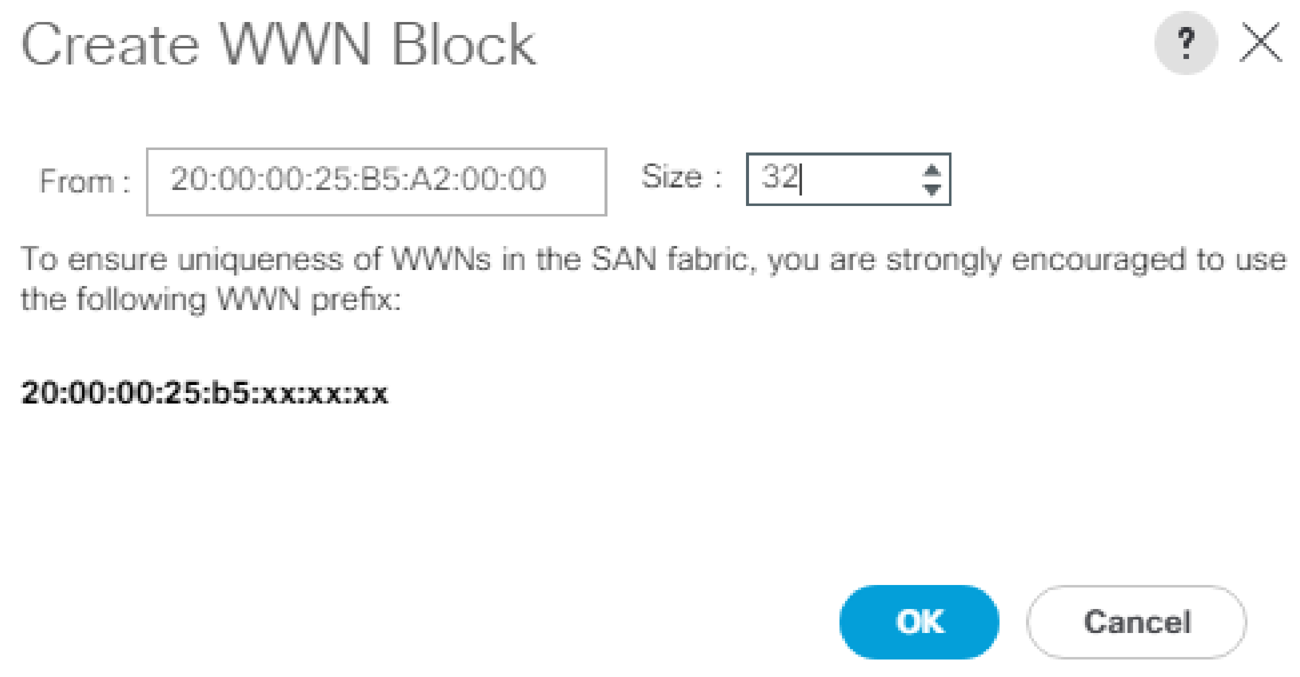

Create WWN Block

From : 20:00:00:25:B5:A2:00:00 Size : 32

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

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

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

Create WWPN Pools

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

1. In Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. In this procedure, two WWPN pools are created, one for each switching fabric.
4. Right-click WWPN Pools under the root organization.
5. Select Create WWPN Pool to create the WWPN pool.
6. Enter `WWPN-Pool-A` as the name of the WWPN pool.
7. Optional: Enter a description for the WWPN pool.
8. Select **Sequential** for Assignment Order.

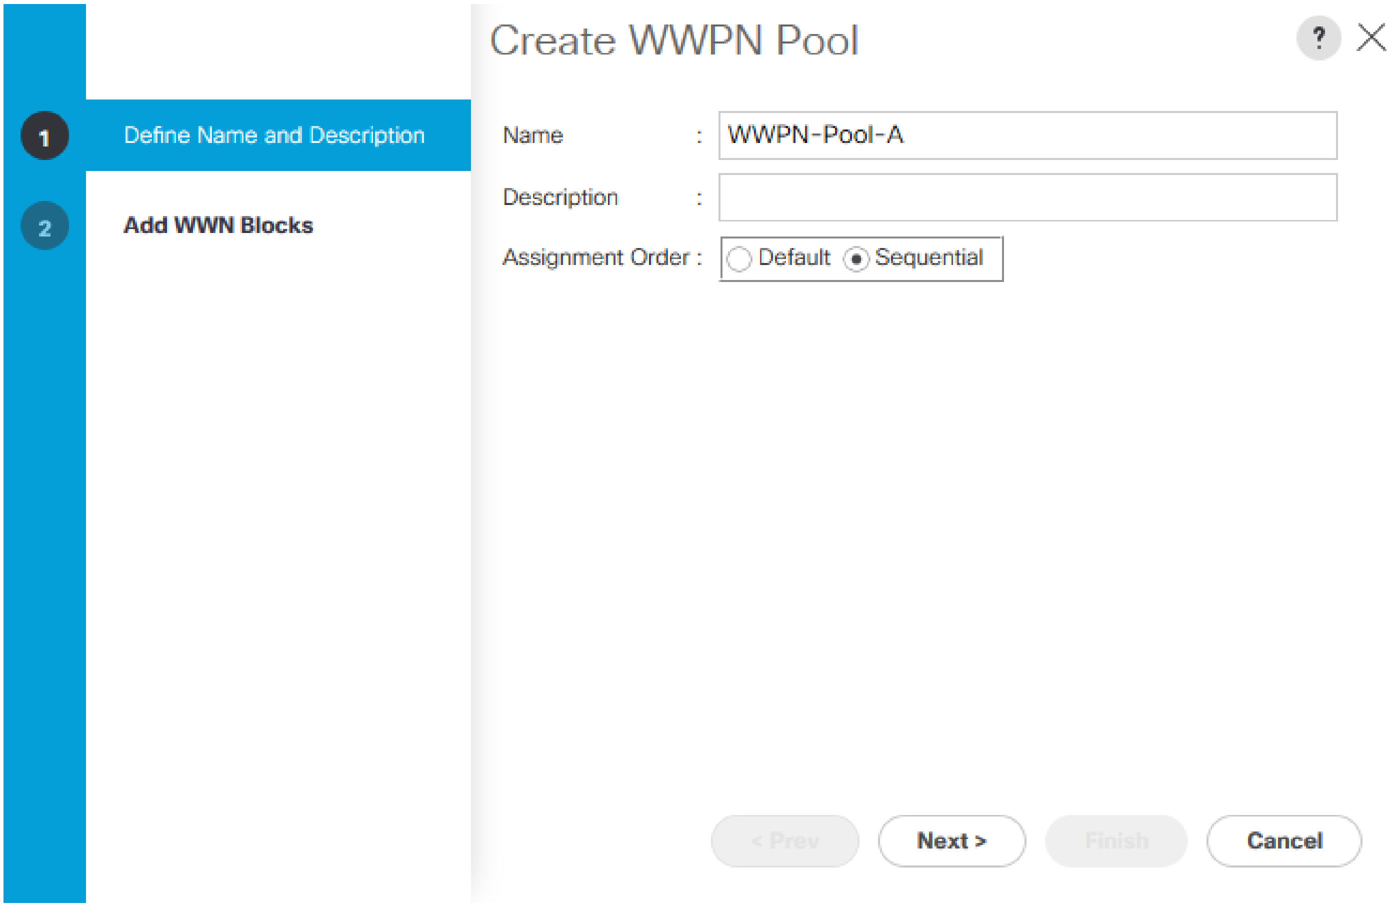

Create WWPN Pool

Name : WWPN-Pool-A

Description :

Assignment Order : ○ Default ● Sequential

9. Click Next.

10. Click Add.

11. Specify a starting WWPN

For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern you used for the WWNN you will see a WWPN block starting with `20:00:00:25:B5:A2:0A:00`

12. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources.

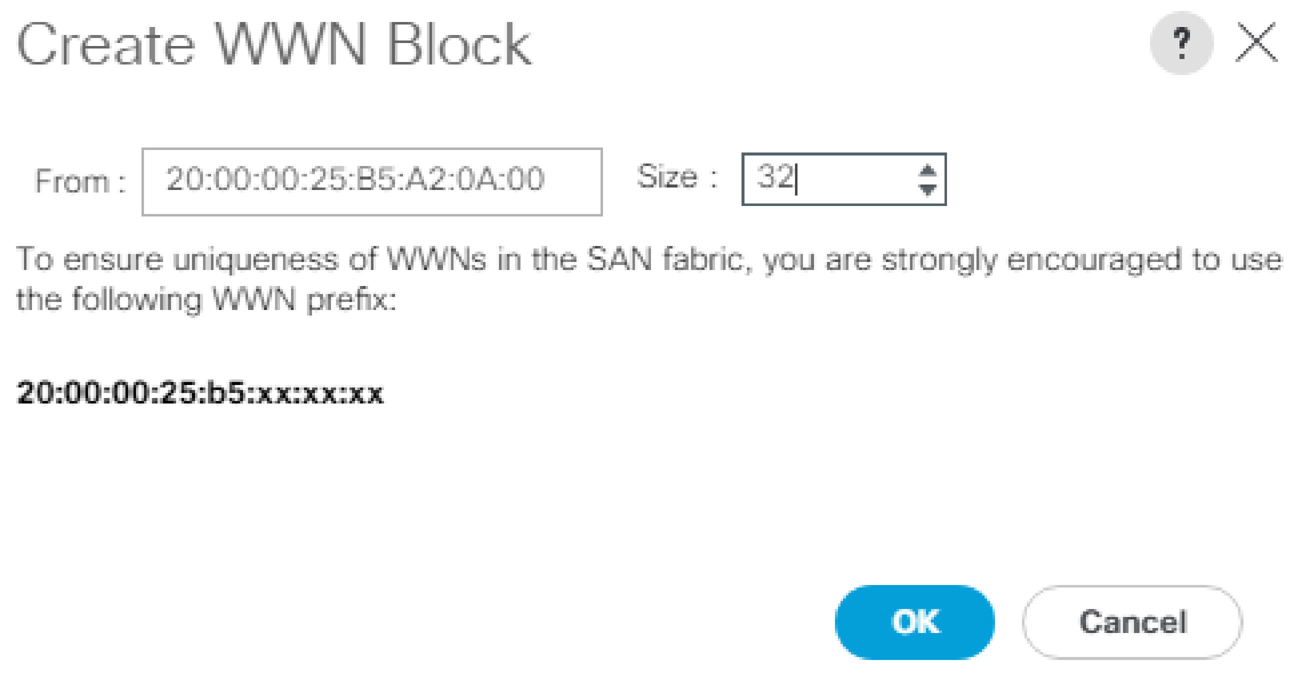

Create WWN Block

From : 20:00:00:25:B5:A2:0A:00

Size : 32

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

13. Click OK.

14. Click Finish.

15. In the confirmation message, click OK.

16. Right-click WWPN Pools under the root organization.

17. Select Create WWPN Pool to create the WWPN pool.

18. Enter `WWPN-Pool-B` as the name of the WWPN pool.

19. Optional: Enter a description for the WWPN pool.

20. Select **Sequential** for Assignment Order.

21. Click Next.

22. Click Add.

23. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:A2:0B:00`.

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

25. Click OK.

26. Click Finish.
27. In the confirmation message, click OK.

Create vHBA Templates

To create the necessary virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter `vHBA-Template-A` as the vHBA template name.
6. Keep Fabric A selected.
7. Leave Redundancy Type set to No Redundancy.
8. Select `VSAN-A`.
9. Leave Initial Template as the Template Type.
10. Select `WWPN-Pool-A` as the WWPN Pool.
11. Click OK to create the vHBA template.
12. Click OK

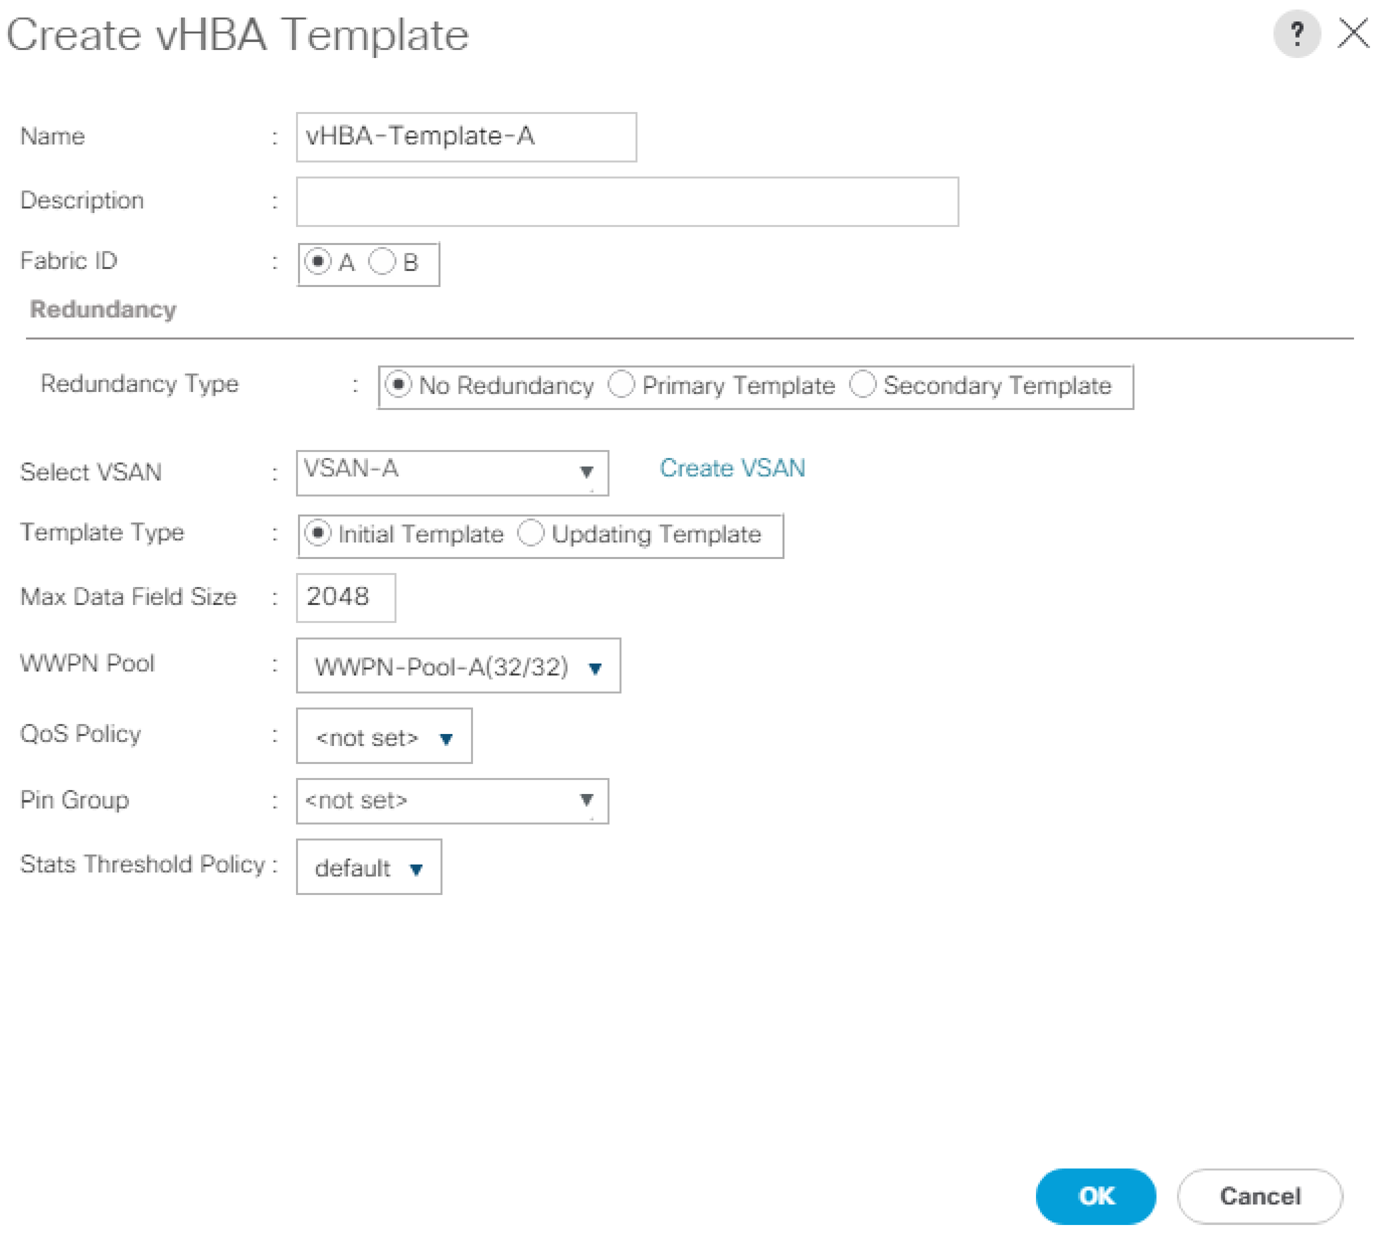

Create vHBA Template

Name : vHBA-Template-A

Description :

Fabric ID : ● A ○ B

Redundancy

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Select VSAN : VSAN-A ▼ Create VSAN

Template Type : ● Initial Template ○ Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-Pool-A(32/32) ▼

QoS Policy : <not set> ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

13. Right-click vHBA Templates.

14. Select Create vHBA Template.

15. Enter `vHBA-Template-B` as the vHBA template name.

16. Select Fabric B as the Fabric ID.

17. Leave Redundancy Type set to No Redundancy.

18. Select `VSAN-B`.

19. Leave Initial Template as the Template Type.

20. Select `WWPN-Pool-B` as the WWPN Pool.
21. Click OK to create the vHBA template.
22. Click OK.

Create SAN Connectivity Policy

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

1. In Cisco UCS Manager, click SAN on the left.
2. Select SAN > Policies > root > Sub-Organizations > FPV-FlexPod.
3. Select and right-click on SAN Connectivity Policies.
4. Select Create SAN Connectivity Policy.
5. Enter `FPV-App-B-FC-Bt` as the name of the policy.
6. Select the previously created `WWNN-Pool` for the WWNN Assignment.
7. Click Add at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter `Fabric-A` as the name of the vHBA.
9. Select the Use vHBA Template checkbox.
10. In the vHBA Template list, select `vHBA-Template-A`.
11. In the Adapter Policy list, select VMware.

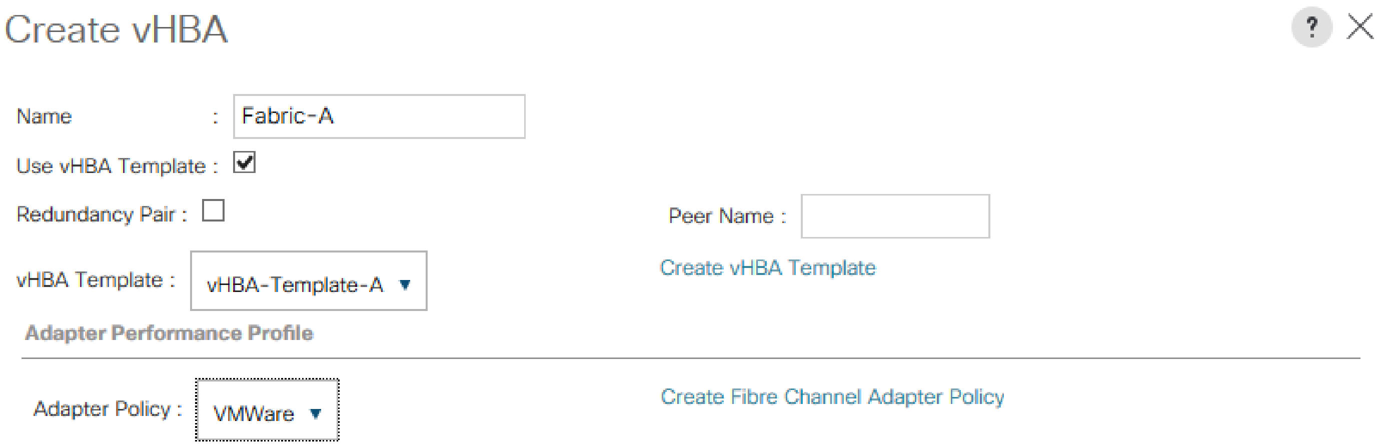

Create vHBA

Name : Fabric-A

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : vHBA-Template-A

Create vHBA Template
Adapter Performance Profile
Create Fibre Channel Adapter Policy

Adapter Policy : VMWare

12. Click OK.
13. Click the Add button at the bottom to add a second vHBA.
14. In the Create vHBA dialog box, enter `Fabric-B` as the name of the vHBA.
15. Select the Use vHBA Template checkbox.

16. In the vHBA Template list, select `vHBA-Template-B`.

17. In the Adapter Policy list, select VMware.

18. Click OK.

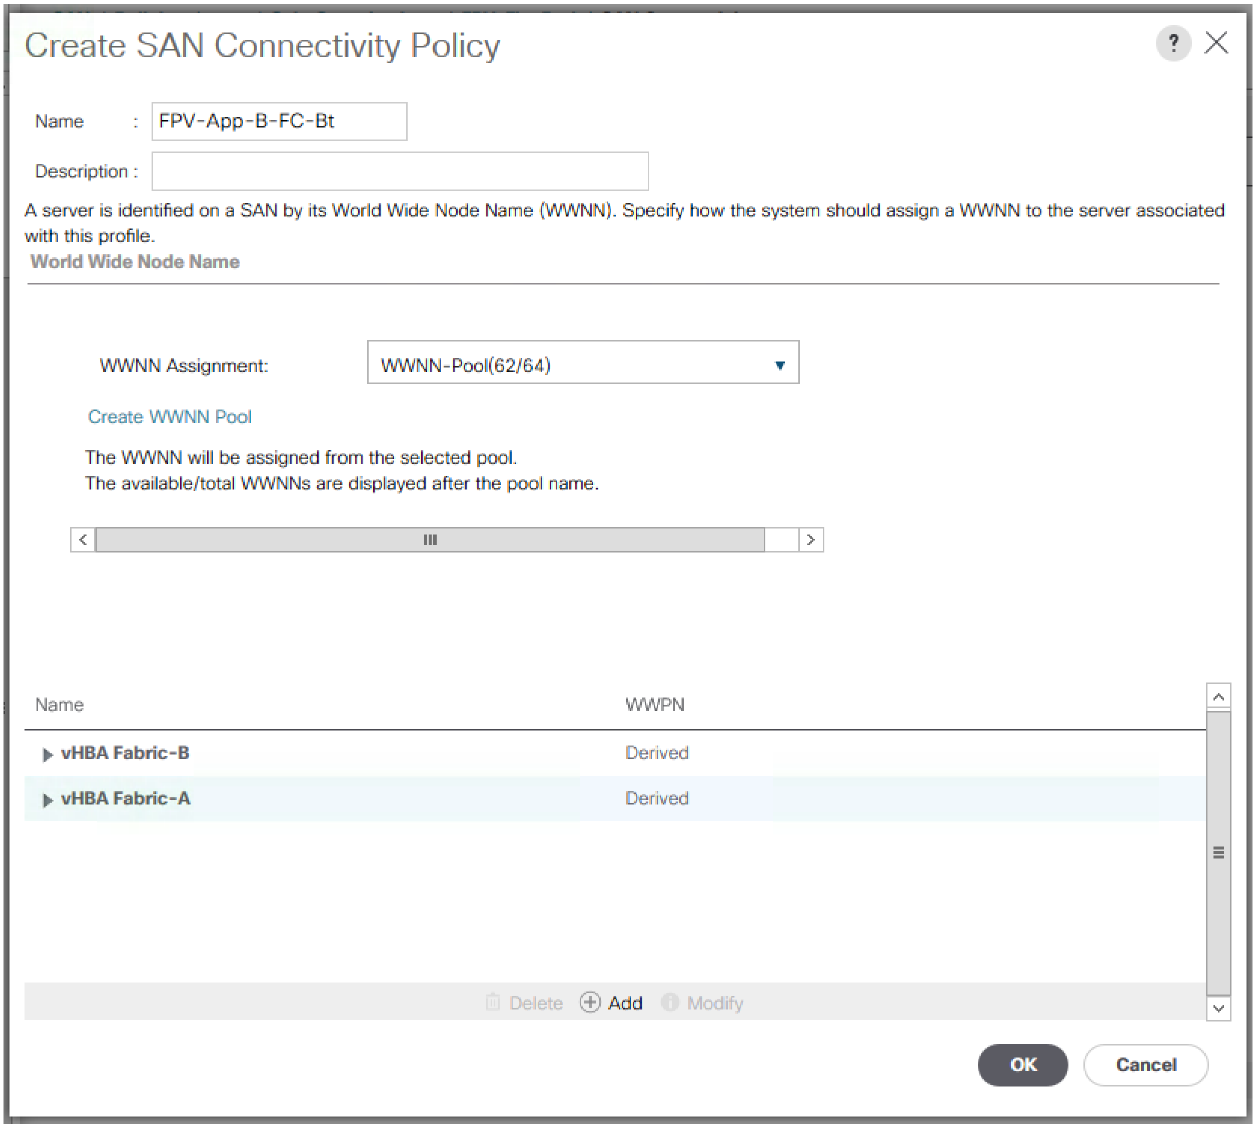

Create SAN Connectivity Policy

Name : FPV-App-B-FC-Bt

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: WWNN-Pool(62/64)

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

19. Click OK to create the SAN Connectivity Policy.

20. Click OK to confirm creation.

Create Server Pool

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

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

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

2. Select Pools > root > Sub-Organizations > FPV-FlexPod.
3. Select and right-click on Server Pools.
4. Select Create Server Pool.
5. Enter `FPV-App-B-Server-Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used in the cluster and click >> to add them to the `FPV-App-B-Server-Pool` server pool.
9. Click Finish.
10. Click OK.

Create LAN Connectivity Policy for FC Boot

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

1. In Cisco UCS Manager, click LAN on the left.
2. Select and expand LAN > Policies > root > Sub-Organizations > FPV-FlexPod.
3. Select and right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter `FC-Boot` as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter `00-Infra-A` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Infra-A.
10. In the Adapter Policy list, select VMware.
11. Click OK to add this vNIC to the policy.

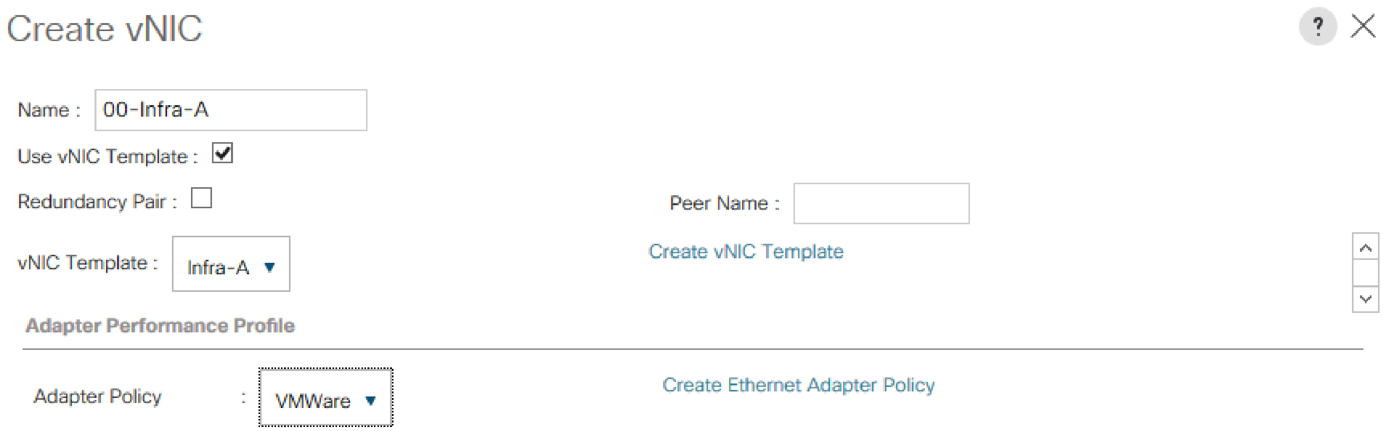

Create vNIC

Name : 00-Infra-A

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : Infra-A ▾

Adapter Performance Profile

Adapter Policy : VMWare ▾

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

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

14. Select the Use vNIC Template checkbox.

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

16. In the Adapter Policy list, select VMware.

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

18. Click the upper Add button to add a vNIC.

19. In the Create vNIC dialog box, enter `02-APIC-VS-A` as the name of the vNIC.

20. Select the Use vNIC Template checkbox.

21. In the vNIC Template list, select APIC-VS-A.

22. In the Adapter Policy list, select either VMware or VMware-HighTrf.

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

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

25. In the Create vNIC box, enter `03-APIC-VS-B` as the name of the vNIC.

26. Select the Use vNIC Template checkbox.

27. In the vNIC Template list, select APIC-VS-B.

28. In the Adapter Policy list, select either VMware or VMware-HighTrf.

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

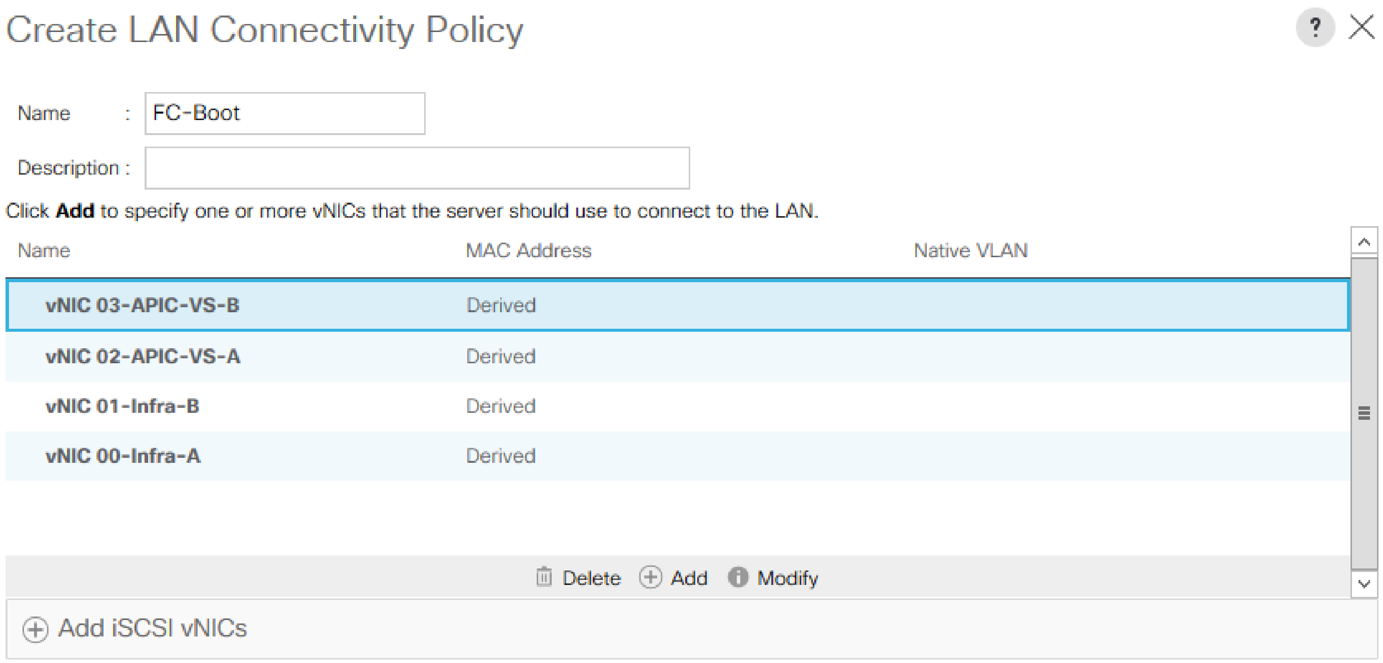

Create LAN Connectivity Policy

Name : FC-Boot

Description :

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

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

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

Create Boot Policy (FC Boot)

This procedure applies to a Cisco UCS environment in which two FC logical interfaces (LIFs) are on cluster node 1 (fcp_lif01a and fcp_lif01b) and two FC LIFs are on cluster node 2 (fcp_lif02a and fcp_lif02b).

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select and expand Servers > Policies > root > Sub-Organizations > FPV-FlexPod > Boot Policies.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `FC-Boot` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.

Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down list and select `Add Remote CD/DVD`.
9. Expand the vHBAs drop-down list and select Add SAN Boot.
10. Select the Primary for type field.

11. Enter `Fabric-A` in vHBA field.

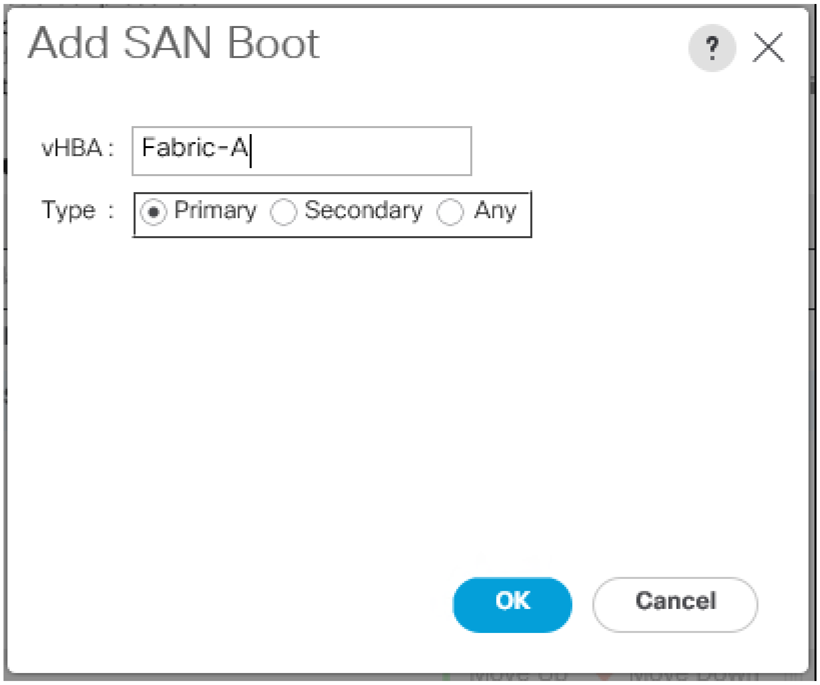

vHBA : Fabric-A

Type : ● Primary ○ Secondary ○ Any

12. Click OK.

13. From the vHBA drop-down list, select Add SAN Boot Target.

14. Keep 0 as the value for Boot Target LUN.

15. Enter the WWPN for fcp_lif01a.

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

---

16. Select Primary for the SAN boot target type.

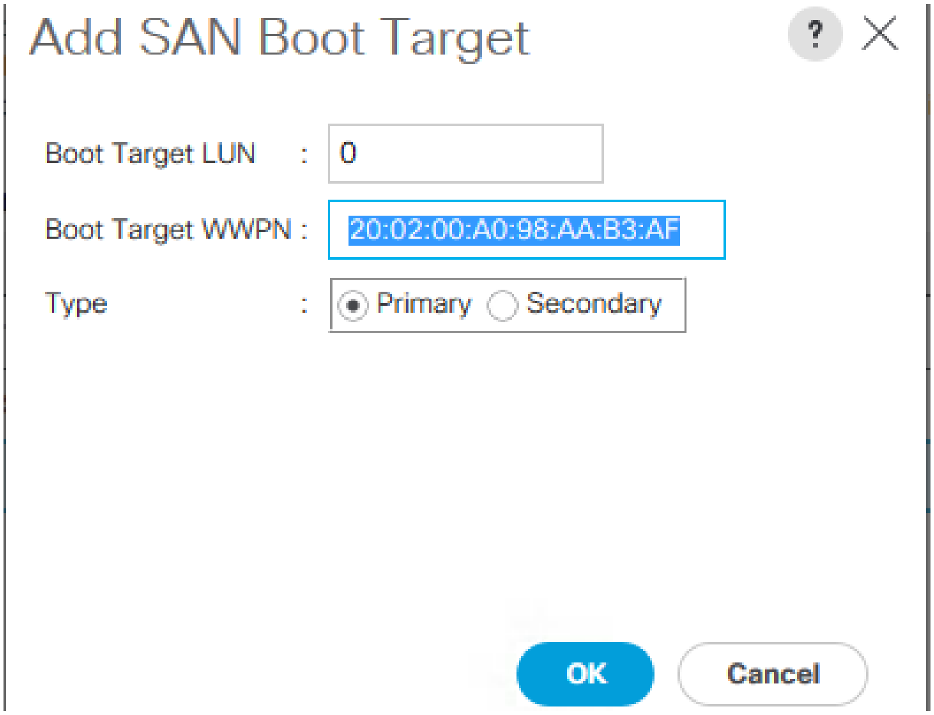

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:02:00:A0:98:AA:B3:AF

Type : ● Primary ○ Secondary

17. Click OK to add the SAN boot target.

18. From the vHBA drop-down list, select Add SAN Boot Target.

19. Enter 0 as the value for Boot Target LUN.

20. Enter the WWPN for fcp_lif02a.

21. Click OK to add the SAN boot target.

22. From the vHBA drop-down list, select Add SAN Boot.

23. In the Add SAN Boot dialog box, enter `Fabric-B` in the vHBA box.

24. The SAN boot type should automatically be set to Secondary.

25. Click OK to add the SAN boot.

26. From the vHBA drop-down list, select Add SAN Boot Target.

27. Keep 0 as the value for Boot Target LUN.

28. Enter the WWPN for fcp_lif01b.

29. Select Primary for the SAN boot target type.

30. Click OK to add the SAN boot target.

31. From the vHBA drop-down list, select Add SAN Boot Target.

32. Keep 0 as the value for Boot Target LUN.

33. Enter the WWPN for fcp_lif02b.

34. Click OK to add the SAN boot target.

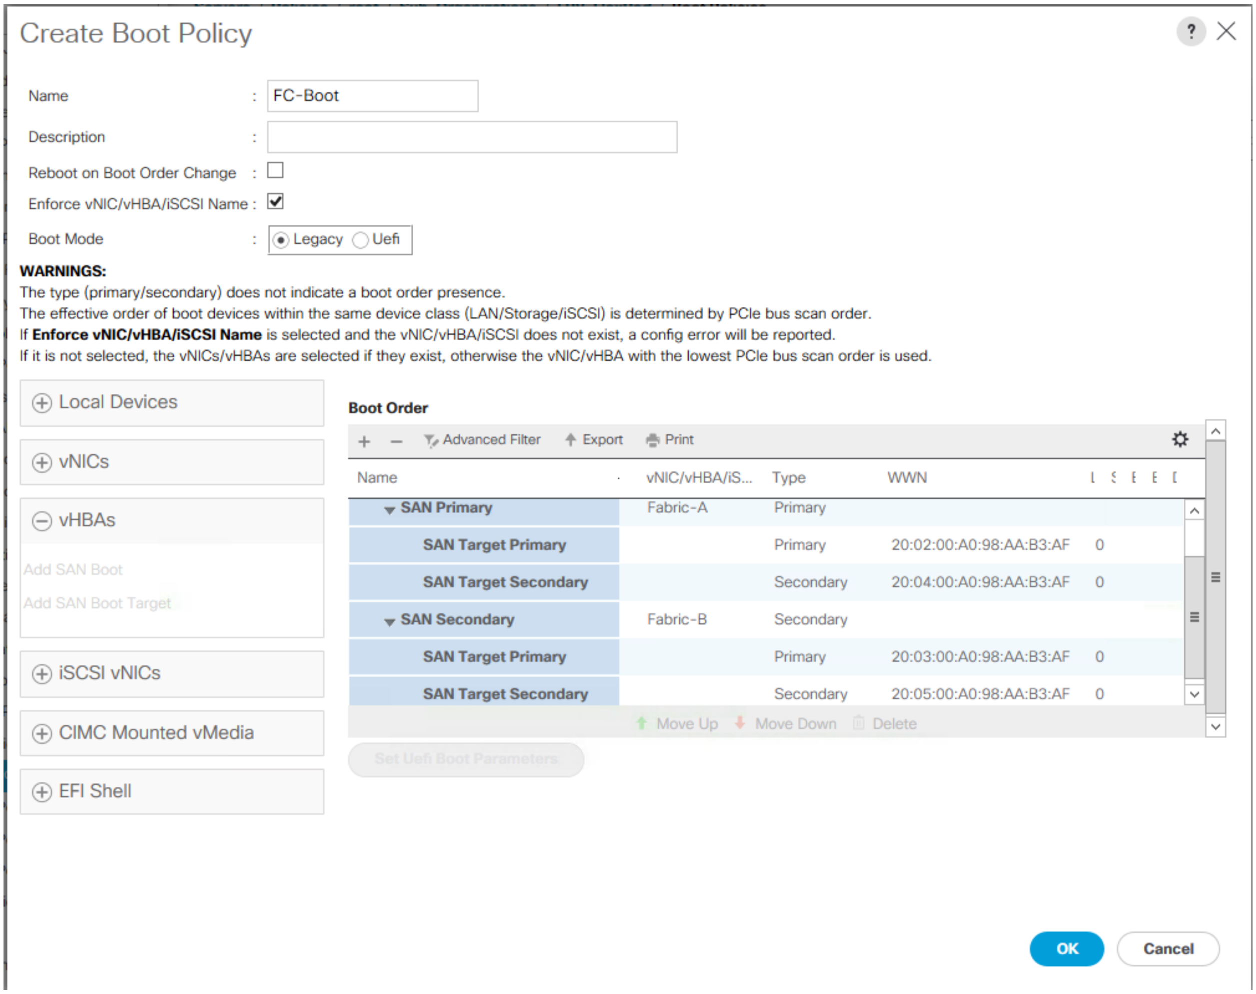

Create Boot Policy

Name :	FC-Boot
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.

Boot Order

Name	vNIC/vHBA/iS...	Type	WWN	L	S	E	E	D
SAN Primary	Fabric-A	Primary						
SAN Target Primary		Primary	20:02:00:A0:98:AA:B3:AF	0				
SAN Target Secondary		Secondary	20:04:00:A0:98:AA:B3:AF	0				
SAN Secondary	Fabric-B	Secondary						
SAN Target Primary		Primary	20:03:00:A0:98:AA:B3:AF	0				
SAN Target Secondary		Secondary	20:05:00:A0:98:AA:B3:AF	0				

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

Create Service Profile Template

In this procedure, a service profile template for VMware ESXi hosts are created for Fabric A boot.

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

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

2. Select Service Profile Templates > root > Sub-Organizations > FPV-FlexPod.

3. Select and right-click FPV-FlexPod.

4. Select Create Service Profile Template to open the Create Service Profile Template wizard.

5. Enter `FPV-App-B-FC-Boot` as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.

6. Select the “Updating Template” option.

7. Under UUID Assignment, select UUID-Pool as the UUID pool.

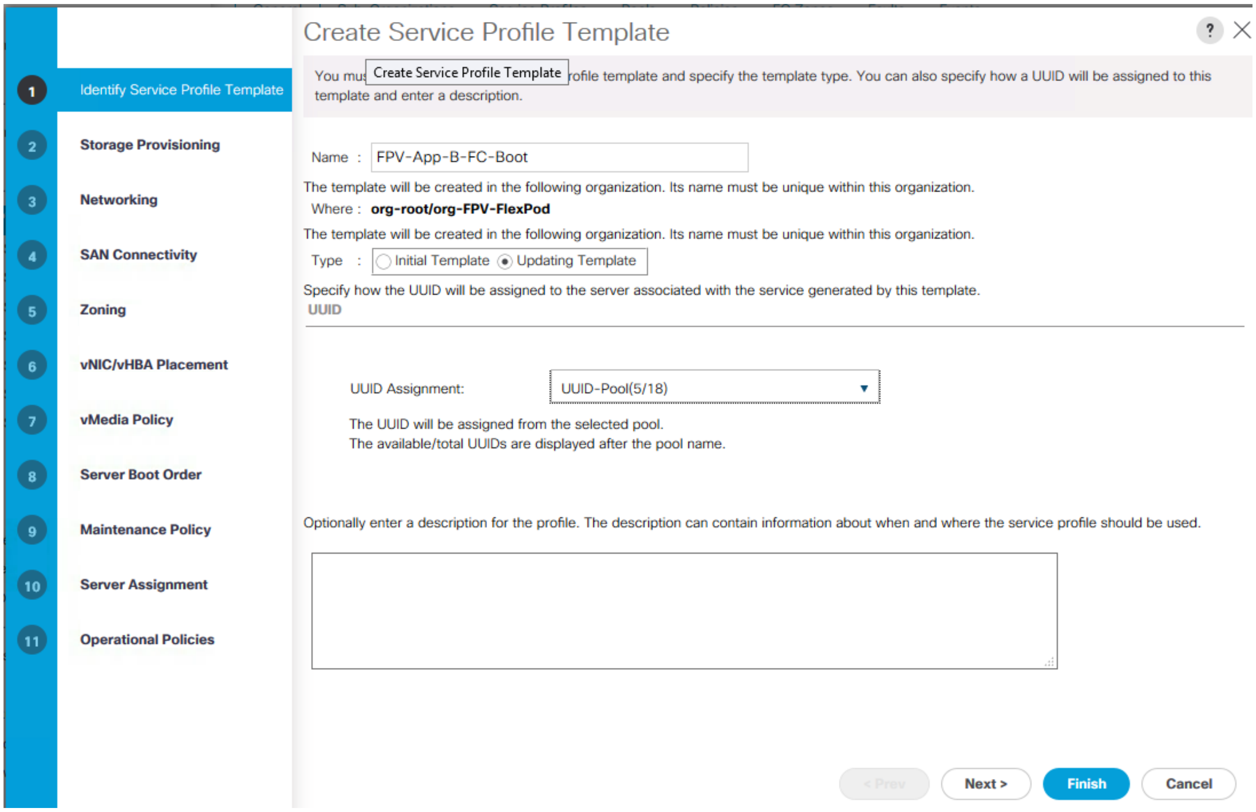

Create Service Profile Template

You mus[Create Service Profile Template]rofile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name : FPV-App-B-FC-Boot

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-FPV-FlexPod**

The template will be created in the following organization. Its name must be unique within this organization.

Type : ○ Initial Template ● Updating Template

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

UUID

UUID Assignment: UUID-Pool(5/18)

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

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

8. Click Next.

Configure Storage Provisioning

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

2. Click Next.

Configure Networking Options

1. Keep the setting at default for Dynamic vNIC Connection Policy.

2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.

3. Select FC-Boot from the LAN Connectivity Policy drop-down list.

4. Leave Initiator Name Assignment at <not set>.

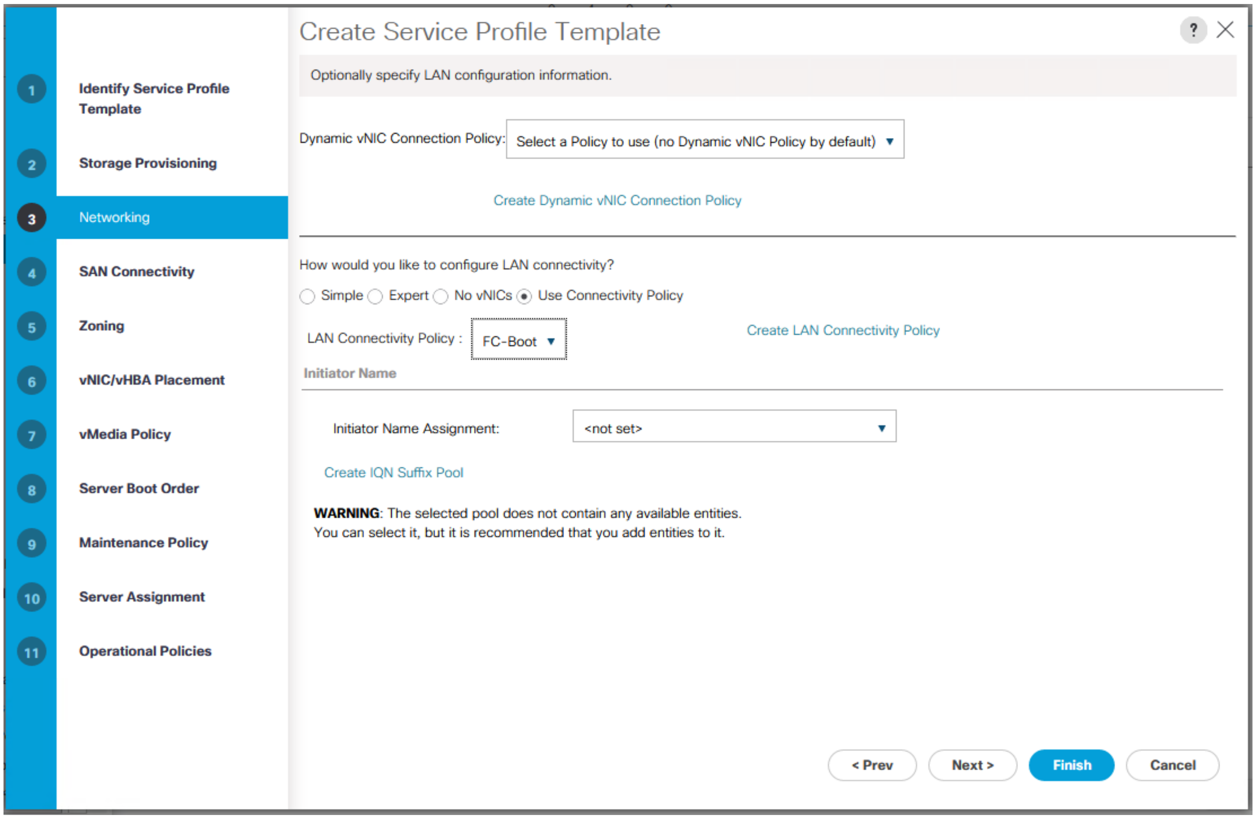

Create Service Profile Template

Optionally specify LAN configuration information.

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

How would you like to configure LAN connectivity?

○ Simple ○ Expert ○ No vNICs ● Use Connectivity Policy

LAN Connectivity Policy : FC-Boot ▾

Initiator Name

Initiator Name Assignment: <not set> ▾

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

5. Click Next.

Configure Storage Options

To configure storage options, 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 FPV-App-B-FC-Bt option from the SAN Connectivity Policy drop-down list.

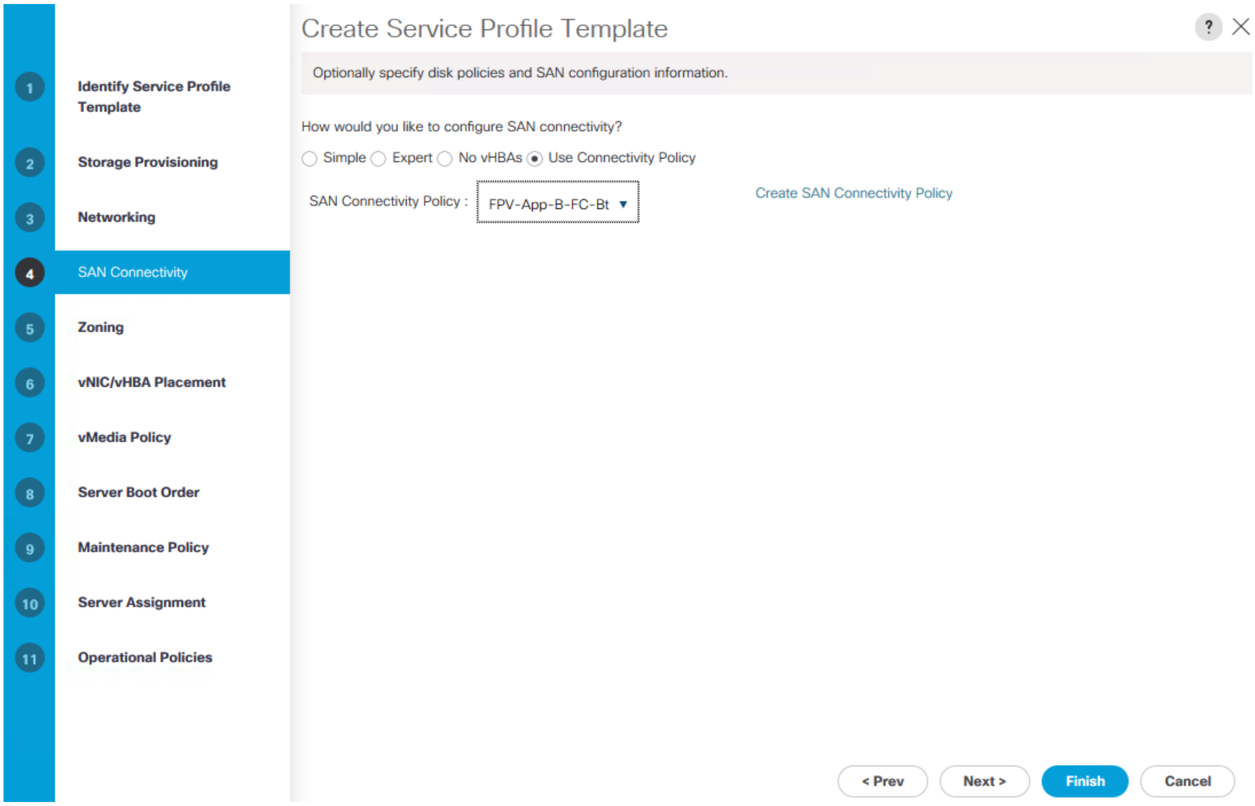

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

○ Simple ○ Expert ○ No vHBAs ● Use Connectivity Policy

SAN Connectivity Policy : FPV-App-B-FC-Bt ▾

3. Click Next.

Configure Zoning Options

1. Set no Zoning options and click Next.

Configure vNIC/HBA Placement

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

2. Click Next.

Configure vMedia Policy

1. Do not select a vMedia Policy.

2. Click Next.

Configure Server Boot Order

1. Select FC-Boot for Boot Policy.

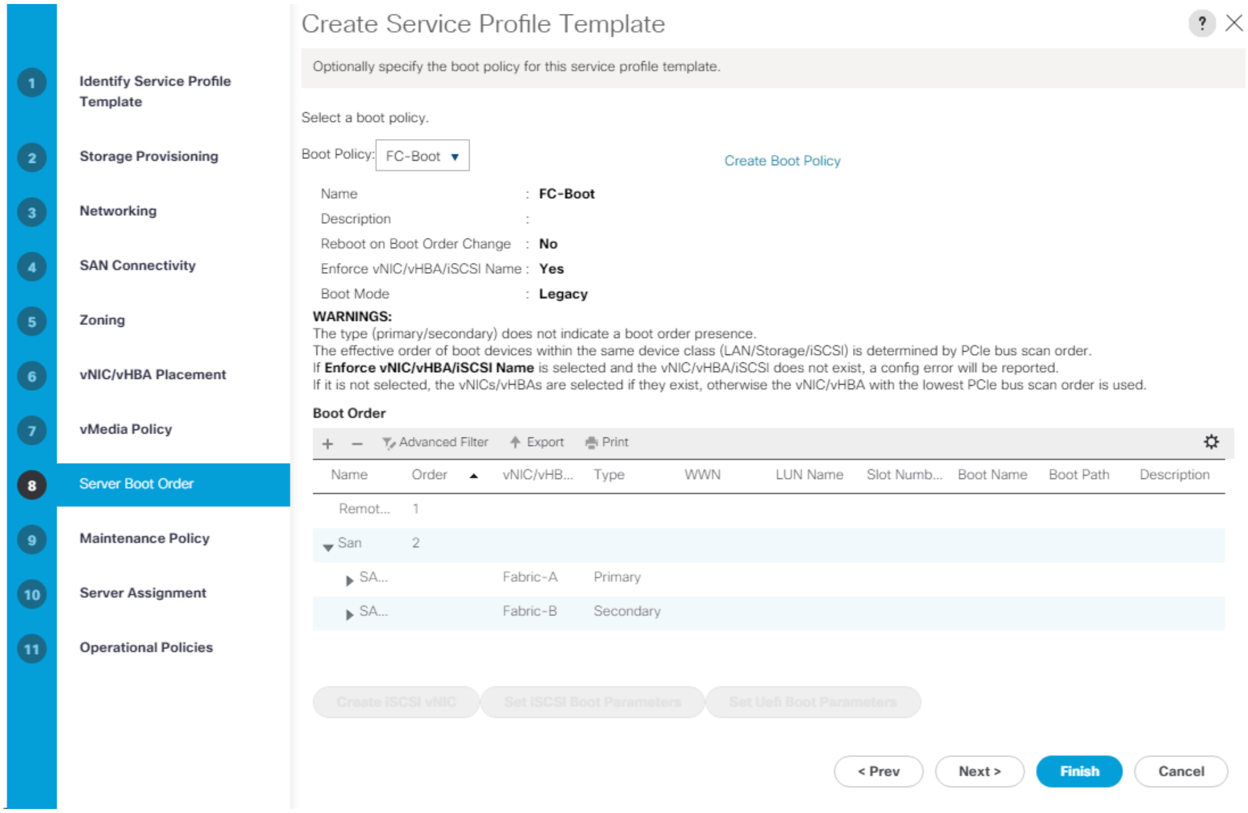

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: FC-Boot ▼

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

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

Boot Order

Name	Order ▲	vNIC/vHB...	Type	WWN	LUN Name	Slot Numb...	Boot Name	Boot Path	Description
Remot...	1								
▼ San	2								
▶ SA...		Fabric-A	Primary						
▶ SA...		Fabric-B	Secondary						

2. Click Next.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

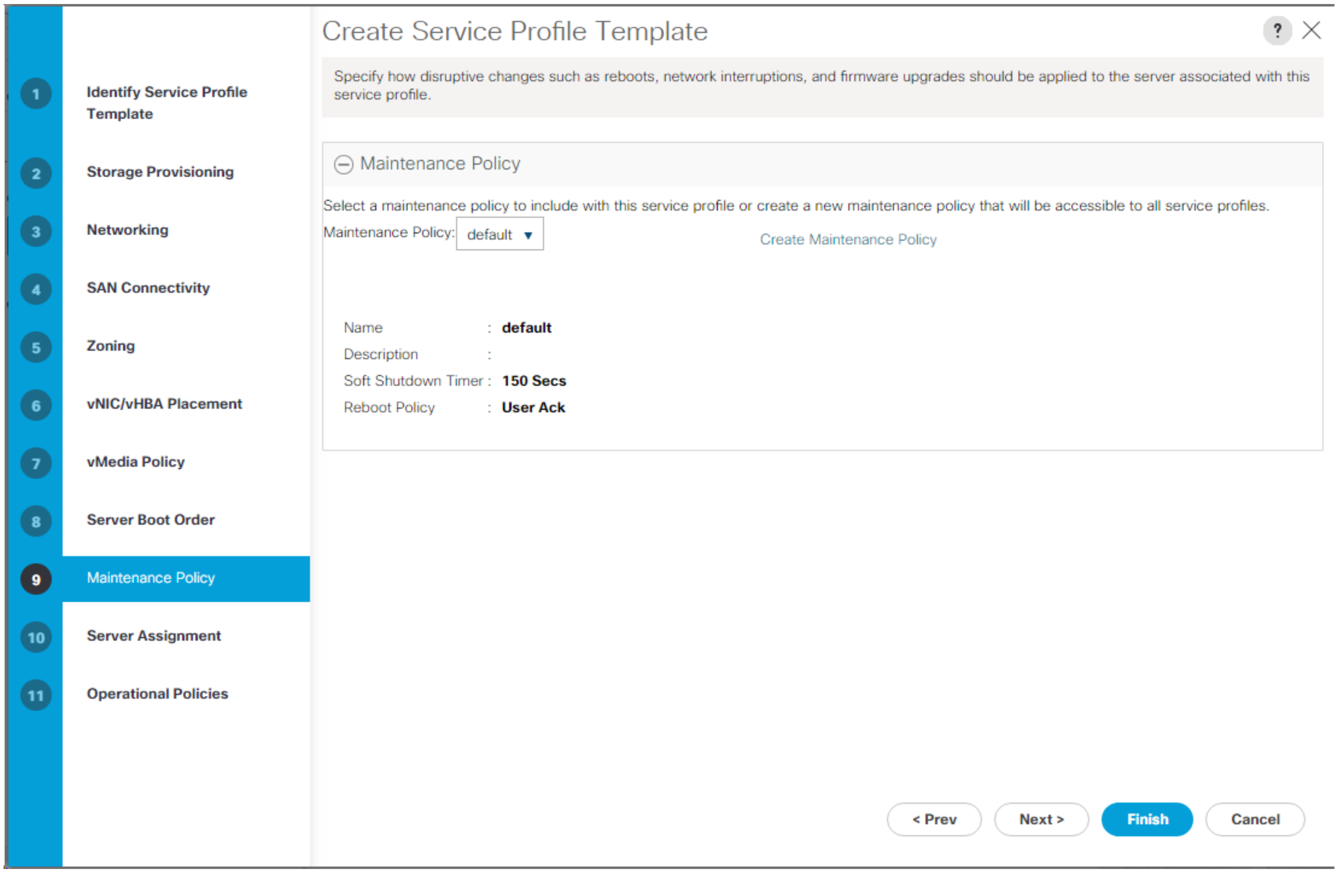

Create Service Profile Template

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

Maintenance Policy

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

Maintenance Policy: default

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

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select `FPV-App-B-Server-Pool`.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select “UCSB-200-M5” for the Server Pool Qualification.
4. Expand Firmware Management at the bottom of the page and select the default policy.

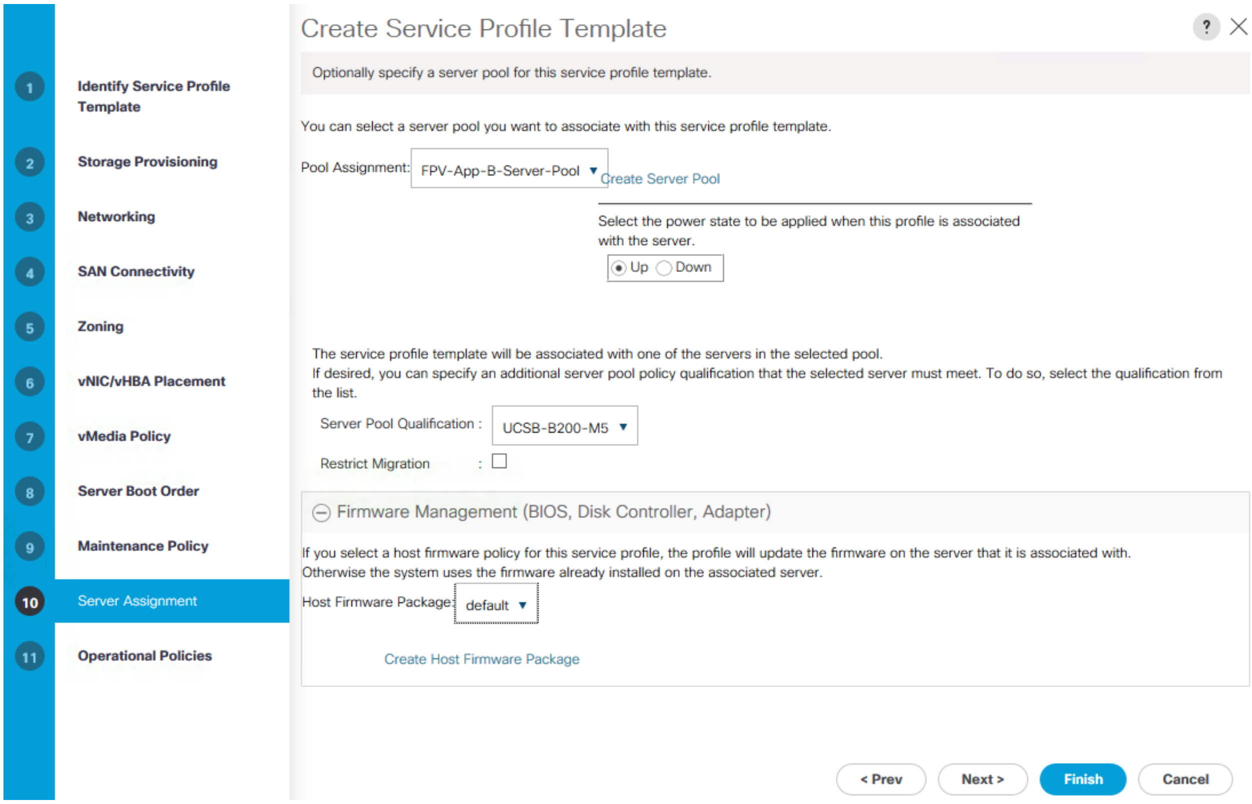

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: FPV-App-B-Server-Pool ▾ 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 : UCSB-B200-M5 ▾

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: default ▾

Create Host Firmware Package

5. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. In the BIOS Policy list, select Virtual-Host.

2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

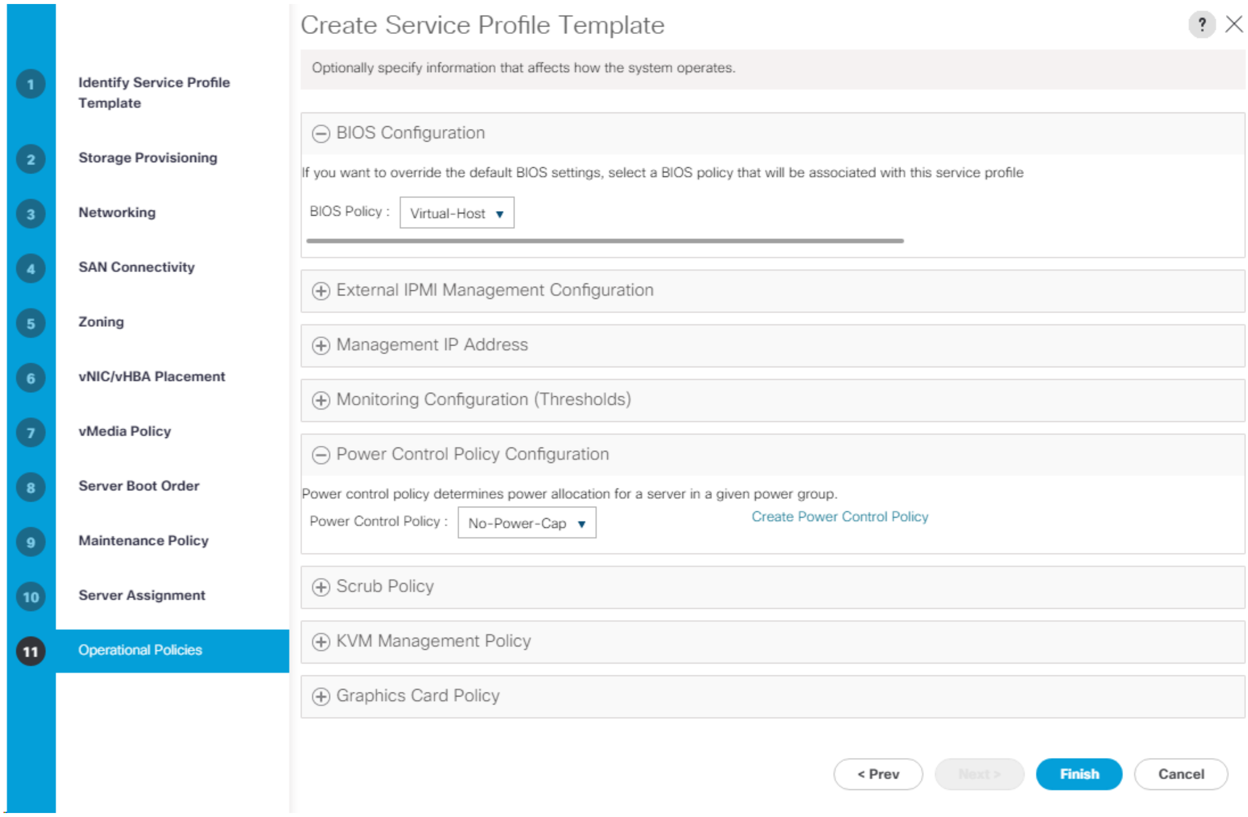

Create Service Profile Template

Optionally specify information that affects how the system operates.

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

BIOS Policy : Virtual-Host

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

Power Control Policy : No-Power-Cap

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create Service Profiles

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

1. Connect to UCS Manager and click Servers on the left.

2. Select Service Profile Templates > root > Sub-Organizations > FPV-FlexPod > Service Template `FPV-App-B-FC-Boot`.

3. Right-click `FPV-App-B-FC-Boot` and select Create Service Profiles from Template.

4. Enter `fpv-esxi-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.

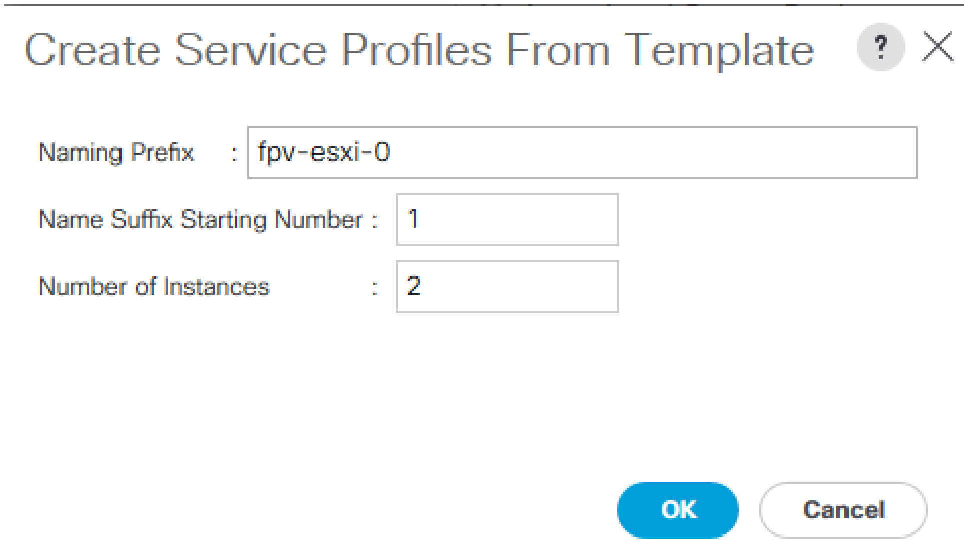

Create Service Profiles From Template

Naming Prefix : fpv-esxi-0

Name Suffix Starting Number : 1

Number of Instances : 2

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.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 9 and Table 10 .

Table 9 WWPNs from NetApp Storage

SVM	Adapter	Path	Target: WWPN
FPV-Foundation-SVM	fcp_lif01a	Fabric A	<fcp_lif01a-wwpn>
	fcp_lif01b	Fabric B	<fcp_lif01b-wwpn>
	fcp_lif02a	Fabric A	<fcp_lif02a-wwpn>
	fcp_lif02b	Fabric B	<fcp_lif02b-wwpn>

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

Table 10 WWPNs for UCS Service Profiles

Cisco UCS Service Profile Name	Path	Initiator WWPN
fpv-esxi-01	Fabric A	fpv-esxi-01-wwpna

Cisco UCS Service Profile Name	Path	Initiator WWPN
	Fabric B	fpv-esxi-01-wwpnb
fpv-esxi-02	Fabric A	fpv-esxi-02-wwpna
	Fabric B	fpv-esxi-02-wwpnb

To obtain the FC vHBA WWPN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the "Storage" tab, then "vHBAs" tab on the right. The WWPNs are displayed in the table at the bottom of the page.

Adding Direct Connected Tenant FC Storage

To add FC storage from an additional storage SVM, two storage connection policies, one for each fabric must be added in Cisco UCS Manager and attached to vHBA Initiator Groups in the SAN Connectivity Policy. These steps were not shown in the initial deployment above because it is not necessary to zone boot targets. Boot targets are automatically zoned in the fabric interconnect when zoning is enabled on the fabric VSAN. To add direct connected tenant FC storage from a tenant SVM, complete the following steps:

It is recommended to clone the Foundation service profile template to separate templates for each additional tenant. That way FC storage from the specific tenant SVM can be mapped to the specific tenant host servers.

Create Storage Connection Policies

In this procedure, one storage connection policy is created for each fabric.

To create the storage connection policies, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Right-click SAN > Policies > root > Sub-Organizations > FPV-FlexPod > Storage Connection Policies and select Create Storage Connection Policy.
3. Name the policy to indicate a tenant on Fabric A (`FPV-App-B-FC-A`).
4. Select the Single Initiator Multiple Targets Zoning Type.
5. Click Add to add a target.
6. Enter the WWPN of the first fabric A FC LIF (fcp_lif01a) in the tenant SVM connected to fabric interconnect A. Select Path A and VSAN `VSAN-A`. Click OK.
7. Click Add to add a target.
8. Enter the WWPN of the second fabric A FC LIF (fcp_lif02a) in the tenant SVM connected to fabric interconnect A. Select Path A and VSAN `VSAN-A`. Click OK.
9. Click OK then click OK again to complete adding the Storage Connection Policy.
10. Right-click SAN > Policies > root > Sub-Organizations > FPV-FlexPod > Storage Connection Policies and select Create Storage Connection Policy.

11. Name the policy to indicate a tenant on Fabric B(`FPV-App-B-FC-B`).

12. Select the Single Initiator Multiple Targets Zoning Type.

13. Click Add to add a target.

14. Enter the WWPN of the first fabric B FC LIF (fcp_lif01b) in the tenant SVM connected to fabric interconnect B. Select Path B and VSAN `VSAN-B`. Click OK.

15. Click Add to add a target.

16. Enter the WWPN of the second fabric B FC LIF (fcp_lif02b) in the tenant SVM connected to fabric interconnect B. Select Path B and VSAN `VSAN-B`. Click OK.

17. Click OK then OK again to complete adding the Storage Connection Policy.

Map Storage Connection Policies vHBA Initiator Groups in SAN Connectivity Policy

In this section, storage connection policies are mapped to vHBA initiator groups for each fabric.

To create the storage connection policy mappings, complete the following steps:

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

2. Select SAN > Policies > root > Sub-Organizations > FPV-FlexPod > SAN Connectivity Policies.

3. Create a duplicate of the `FPV-App-B-FC-Bt` SAN Connectivity Policy for the tenant. This policy will need to be mapped into the Tenant Service Profile Template.

4. Select the Tenant San Connectivity Policy.

5. In the center pane, select the vHBA Initiator Groups tab.

6. Click Add to add a vHBA Initiator Group.

7. Name the group `Fabric-A` and select the Fabric A Initiator.

8. Use the drop-down list to select the Tenant Fabric A Storage Connection Policy.

9. Click OK then click OK again to complete adding the Initiator Group.

10. Click Add to add a vHBA Initiator Group.

11. Name the group `Fabric-B` and select the Fabric B Initiator.

12. Use the drop-down list to select the Fabric B Storage Connection Policy.

13. Click OK and OK to complete adding the Initiator Group.

Create igroups

From the storage cluster CLI, to create igroups, run the following commands:

```
igroup create -vserver FPV-Foundation-SVM -igroup fpv-esxi-01 -protocol fcp -ostype vmware -initiator
<fpv-esxi-01-wwpna>,<fpv-esxi-01-wwpnb>
igroup create -vserver FPV-Foundation-SVM -igroup fpv-esxi-02 -protocol fcp -ostype vmware -initiator
<fpv-esxi-02-wwpna>,<fpv-esxi-02-wwpnb>
```

```
igroup create -vserver FPV-Foundation-SVM -igroup MGMT-Hosts-All -protocol fcp -ostype vmware –
initiator <fpv-esxi-01-wwpna>,<fpv-esxi-01-wwpnb>,<fpv-esxi-02-wwpna>,<fpv-esxi-02-wwpnb>
```

Map Boot LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-01 -igroup fpv-esxi-01 -lun-id 0

lun map -vserver FPV-Foundation-SVM -volume esxi_boot -lun fpv-esxi-02 -igroup fpv-esxi-02 -lun-id 0
```

For additional storage related tasks, please see the storage configuration portion of this document.

FC storage in the tenant SVM can be created and mapped with NetApp Virtual Storage Console (VSC).

Appendix - FlexPod Backups

Cisco UCS Backup

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

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

To create a backup using the Cisco UCS Manager GUI, complete the following steps:

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

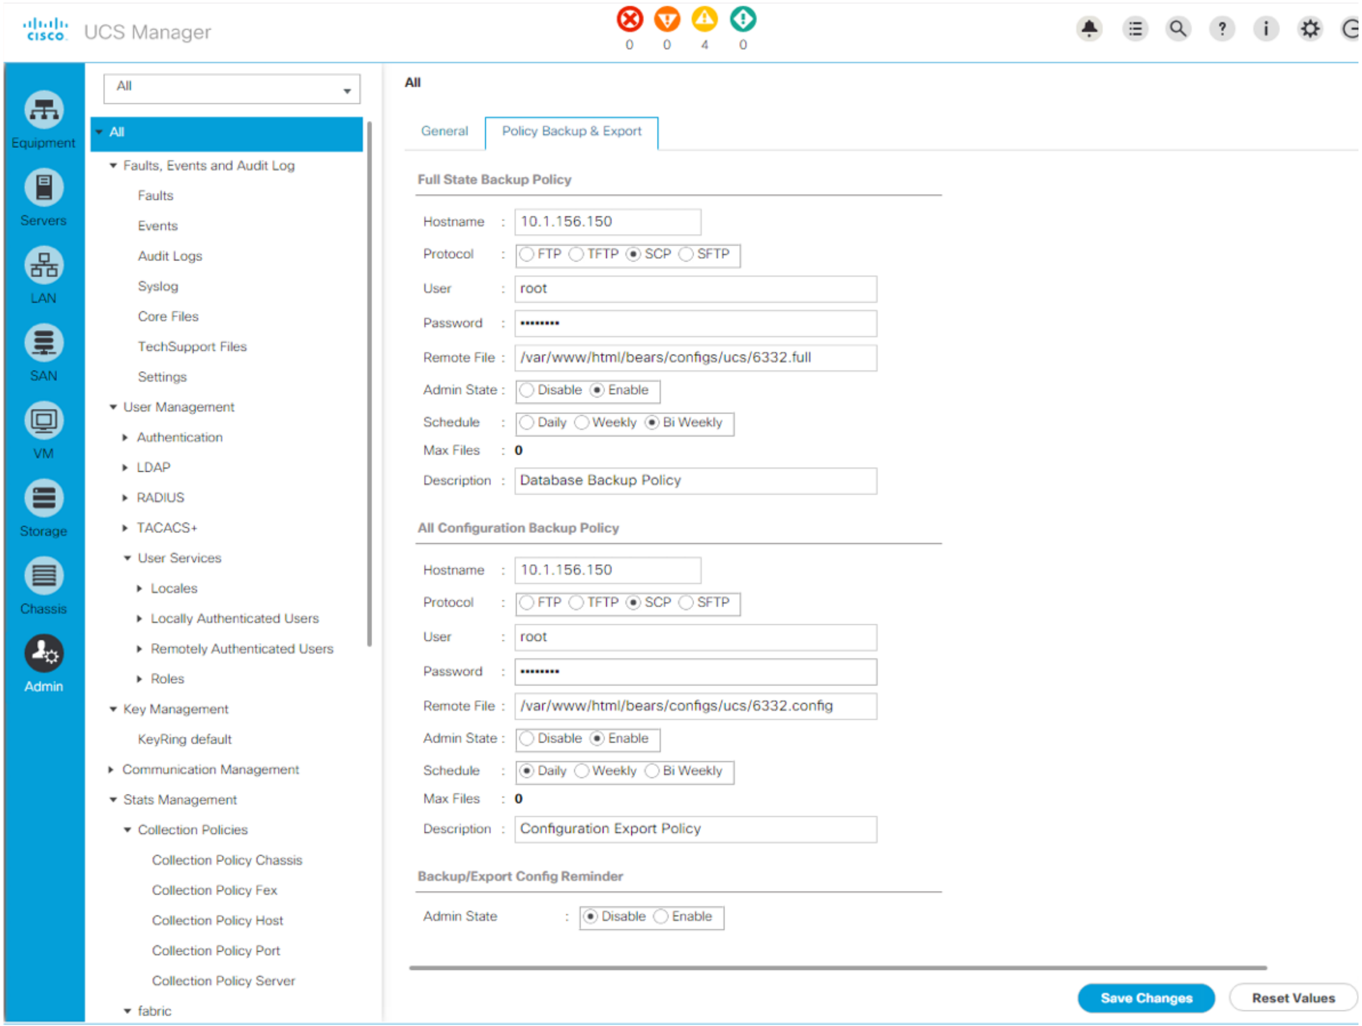

Full State Backup Policy

Hostname : 10.1.156.150

Protocol : ○ FTP ○ TFTP ● SCP ○ SFTP

User : root

Password : ••••••••

Remote File : /var/www/html/bears/configs/ucs/6332.full

Admin State : ○ Disable ● Enable

Schedule : ○ Daily ○ Weekly ● Bi Weekly

Max Files : 0

Description : Database Backup Policy

All Configuration Backup Policy

Hostname : 10.1.156.150

Protocol : ○ FTP ○ TFTP ● SCP ○ SFTP

User : root

Password : ••••••••

Remote File : /var/www/html/bears/configs/ucs/6332.config

Admin State : ○ Disable ● Enable

Schedule : ● Daily ○ Weekly ○ Bi Weekly

Max Files : 0

Description : Configuration Export Policy

Backup/Export Config Reminder

Admin State : ● Disable ○ Enable

4. Click Save Changes to create the Policy.

Appendix – Sample Cisco ASAv Configuration

The Cisco ASAv configuration used in App-A tenant is provided below as a starting point for customers to that want to leverage Cisco ASAv for enabling decentralized, tenant level firewall services.

```
 fpv-app-a-asav# show run
: Saved

:
: Serial Number: <removed>
: Hardware:   ASAv, 2048 MB RAM, CPU Xeon E5 series 2000 MHz
:
ASA Version 9.9(1)
!
hostname fpv-app-a-asav
domain-name flexpod.cisco.com
enable password $sha512$5000$lrcV/TDjJuYrdH0cQE6DFA==$rhHwYAoiVHG1bN0rfZxtBQ==
pbkdf2
xlate per-session deny tcp any4 any4
xlate per-session deny tcp any4 any6
xlate per-session deny tcp any6 any4
xlate per-session deny tcp any6 any6
xlate per-session deny udp any4 any4 eq domain
xlate per-session deny udp any4 any6 eq domain
xlate per-session deny udp any6 any4 eq domain
xlate per-session deny udp any6 any6 eq domain
!
license smart
 feature tier standard
 throughput level 1G
names

!
interface GigabitEthernet0/0
 description Outside Interface
 nameif Outside
 security-level 0
 ip address 172.16.252.1 255.255.255.252
!
interface GigabitEthernet0/1
 description Inside Interface
 nameif Inside
 security-level 100
 ip address 172.16.0.1 255.255.255.0
!
interface GigabitEthernet0/2
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 shutdown
 no nameif
 no security-level
 no ip address
```

```
!
interface GigabitEthernet0/4
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/6
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/7
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/8
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 management-only
 nameif management
 security-level 0
 ip address 172.16.254.19 255.255.255.248
!
ftp mode passive
clock timezone EST -5
clock summer-time EDT recurring
dns domain-lookup management
dns server-group DefaultDNS
 name-server 10.1.118.42 management
 name-server 10.1.118.41 management
 domain-name flexpod.cisco.com
object network Inside-subnet
 subnet 172.16.0.0 255.255.255.0
object network FPV-App-A-Web
 host 172.16.0.10
access-list Outside_acl extended permit ip any any
pager lines 20
logging asdm informational
mtu management 1500
mtu Outside 1500
mtu Inside 1500
no failover
icmp unreachable rate-limit 1 burst-size 1
no asdm history enable
```

```
arp timeout 14400
no arp permit-nonconnected
arp rate-limit 8192
!
object network Inside-subnet
 nat (Inside,Outside) dynamic interface
object network FPV-App-A-Web
 nat (Inside,Outside) static interface service tcp 3389 3389
access-group Outside_acl in interface Outside
route management 0.0.0.0 0.0.0.0 172.16.254.22 1
route Outside 0.0.0.0 0.0.0.0 172.16.252.2 1
timeout xlate 3:00:00
timeout pat-xlate 0:00:30
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 sctp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:05:00 absolute
timeout tcp-proxy-reassembly 0:01:00
timeout floating-conn 0:00:00
timeout conn-holddown 0:00:15
timeout igp stale-route 0:01:10
user-identity default-domain LOCAL
aaa authentication ssh console LOCAL
aaa authentication login-history
http server enable
http 0.0.0.0 0.0.0.0 management
no snmp-server location
no snmp-server contact
crypto ipsec security-association pmtu-aging infinite
crypto ca trustpoint _SmartCallHome_ServerCA
 no validation-usage
 crl configure
crypto ca trustpool policy
 auto-import
crypto ca certificate chain _SmartCallHome_ServerCA
 certificate ca 513fb9743870b73440418d30930699ff
    30820538 30820420 a0030201 02021051 3fb97438 70b73440 418d3093 0699ff30
    0d06092a 864886f7 0d01010b 05003081 ca310b30 09060355 04061302 55533117
    30150603 55040a13 0e566572 69536967 6e2c2049 6e632e31 1f301d06 0355040b
    13165665 72695369 676e2054 72757374 204e6574 776f726b 313a3038 06035504
    0b133128 63292032 30303620 56657269 5369676e 2c20496e 632e202d 20466f72
    20617574 686f7269 7a656420 75736520 6f6e6c79 31453043 06035504 03133c56
    65726953 69676e20 436c6173 73203320 5075626c 69632050 72696d61 72792043
    65727469 66696361 74696f6e 20417574 686f7269 7479202d 20473530 1e170d31
    33313033 31303030 3030305a 170d3233 31303330 32333539 35395a30 7e310b30
    09060355 04061302 5553311d 301b0603 55040a13 1453796d 616e7465 6320436f
    72706f72 6174696f 6e311f30 1d060355 040b1316 53796d61 6e746563 20547275
    7374204e 6574776f 726b312f 302d0603 55040313 2653796d 616e7465 6320436c
    61737320 33205365 63757265 20536572 76657220 4341202d 20473430 82012230
    0d06092a 864886f7 0d010101 05000382 010f0030 82010a02 82010100 b2d805ca
    1c742db5 175639c5 4a520996 e84bd80c f1689f9a 422862c3 a530537e 5511825b
    037a0d2f e17904c9 b4967719 81019459 f9bcf77a 9927822d b783dd5a 277fb203
    7a9c5325 e9481f46 4fc89d29 f8be7956 f6f7fdd9 3a68da8b 4b823341 12c3c83c
    ccd6967a 84211a22 04032717 8b1c6861 930f0e51 80331db4 b5ceeb7e d062acee
    b37b0174 ef6935eb cad53da9 ee9798ca 8daa440e 25994a15 96a4ce6d 02541f2a
    6a26e206 3a6348ac b44cd175 9350ff13 2fd6dae1 c618f59f c9255df3 003ade26
    4db42909 cd0f3d23 6f164a81 16fbf283 10c3b8d6 d855323d f1bd0fbd 8c52954a
    16977a52 2163752f 16f9c466 bef5b509 d8ff2700 cd447c6f 4b3fb0f7 02030100
    01a38201 63308201 5f301206 03551d13 0101ff04 08300601 01ff0201 00303006
    03551d1f 04293027 3025a023 a021861f 68747470 3a2f2f73 312e7379 6d63622e
    636f6d2f 70636133 2d67352e 63726c30 0e060355 1d0f0101 ff040403 02010630
    2f06082b 06010505 07010104 23302130 1f06082b 06010505 07300186 13687474
    703a2f2f 73322e73 796d6362 2e636f6d 306b0603 551d2004 64306230 60060a60
    86480186 f8450107 36305230 2606082b 06010505 07020116 1a687474 703a2f2f
    7777772e 73796d61 7574682e 636f6d2f 63707330 2806082b 06010505 07020230
    1c1a1a68 7474703a 2f2f7777 772e7379 6d617574 682e636f 6d2f7270 61302906
    03551d11 04223020 a41e301c 311a3018 06035504 03131153 796d616e 74656350
    4b492d31 2d353334 301d0603 551d0e04 1604145f 60cf6190 55df8443 148a602a
    b2f57af4 4318ef30 1f060355 1d230418 30168014 7fd365a7 c2ddecbb f03009f3
    4339fa02 af333133 300d0609 2a864886 f70d0101 0b050003 82010100 5e945649
    dd8e2d65 f5c13651 b603e3da 9e7319f2 1f59ab58 7e6c2605 2cfa81d7 5c231722
    2c3793f7 86ec85e6 b0a3fd1f e232a845 6fe1d9fb b9afd270 a0324265 bf84fe16
    2a8f3fc5 a6d6a393 7d43e974 21913528 f463e92e edf7f55c 7f4b9ab5 20e90abd
    e045100c 14949a5d a5e34b91 e8249b46 4065f422 72cd99f8 8811f5f3 7fe63382
    e6a8c57e fed008e2 25580871 68e6cda2 e614de4e 52242dfd e5791353 e75e2f2d
    4d1b6d40 15522bf7 87897812 816ed94d aa2d78d4 c22c3d08 5f87919e 1f0eb0de
    30526486 89aa9d66 9c0e760c 80f274d8 2af8b83a ced7d60f 11be6bab 14f5bd41
    a0226389 f1ba0f6f 2963662d 3fac8c72 c5fbc7e4 d40ff23b 4f8c29c7
  quit
telnet timeout 5
ssh stricthostkeycheck
ssh 0.0.0.0 0.0.0.0 management
ssh timeout 5
ssh key-exchange group dh-group1-sha1
console timeout 0
threat-detection basic-threat
threat-detection statistics access-list
no threat-detection statistics tcp-intercept
ntp server 10.1.118.1 source management
dynamic-access-policy-record DfltAccessPolicy
username admin password
$sha512$5000$x7Sww489lLTWs1uHDwoHQw==$e9Z/gZkxOMgF1l/zeSBr0A== pbkdf2 privilege 15
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum client auto
  message-length maximum 512
  no tcp-inspection
policy-map global_policy
 class inspection_default
  inspect ip-options
  inspect netbios
  inspect rtsp
  inspect sunrpc
  inspect tftp
  inspect xdmcp
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect esmtp
  inspect sqlnet
  inspect sip
  inspect skinny
policy-map type inspect dns migrated_dns_map_2
 parameters
  message-length maximum client auto
  message-length maximum 512
  no tcp-inspection
policy-map type inspect dns migrated_dns_map_1
 parameters
  message-length maximum client auto
  message-length maximum 512
  no tcp-inspection
!
service-policy global_policy global
prompt hostname context
no call-home reporting anonymous
call-home
 profile License
  destination address http
https://tools.cisco.com/its/service/oddce/services/DDCEService
  destination transport-method http
 profile CiscoTAC-1
  no active
  destination address http
https://tools.cisco.com/its/service/oddce/services/DDCEService
  destination address email callhome@cisco.com
  destination transport-method http
  subscribe-to-alert-group diagnostic
  subscribe-to-alert-group environment
  subscribe-to-alert-group inventory periodic monthly
  subscribe-to-alert-group configuration periodic monthly
  subscribe-to-alert-group telemetry periodic daily
Cryptochecksum:19ab2f84378a32842eba4ada89c54709
: end
   fpv-app-a-asav#
```

About the Authors

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

John has over seven years of experience in 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.

Dave Derry, Technical Marketing Engineer, Infrastructure and Cloud Engineering, NetApp

Dave Derry is a Technical Marketing Engineer in the Converged Infrastructure Engineering team at NetApp. He has been with NetApp since 2012, serving in a variety of engineering roles. Prior to that, he was an engineer at Cisco Systems for over ten years, in a variety of development and solution test roles.

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.
- Sanjeev Naldurgkar, Cisco Systems, Inc.
- Rajendra Yogendra, Cisco Systems, Inc.
- Archana Sharma, Cisco Systems, Inc.
- Melissa Palmer, NetApp
- Lindsey Street, NetApp