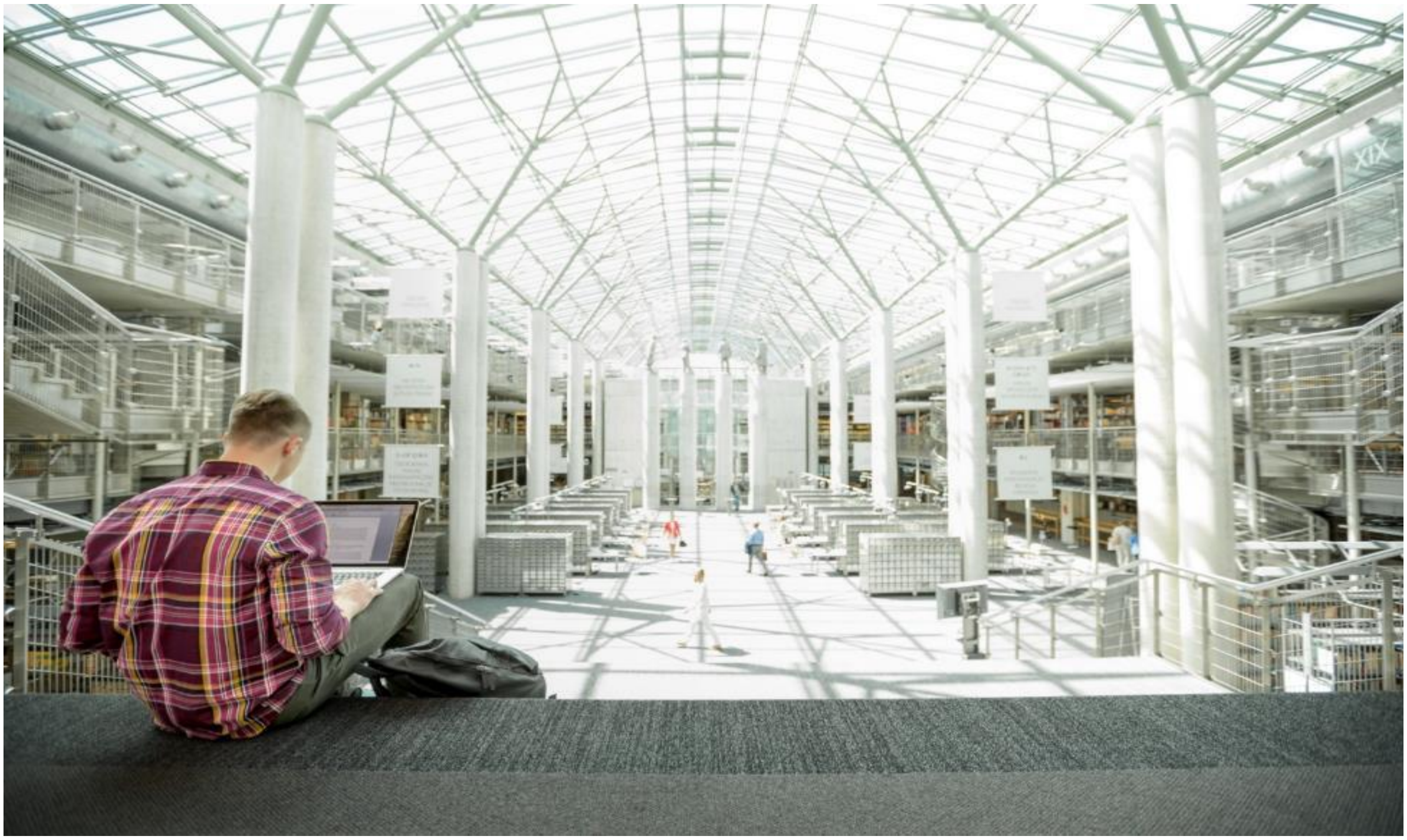

FlexPod Datacenter with VMware vSphere 7.0, Cisco VXLAN Single-Site Fabric, and NetApp ONTAP 9.7

Deployment Guide for FlexPod Datacenter with VMware vSphere 7.0, Cisco VXLAN BGP EVPN Single-Site Fabric, and NetApp ONTAP 9.7

Published: November 2020

In partnership with:

About the Cisco Validated Design Program

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

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

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

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

Contents

Executive Summary

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

FlexPod is a widely deployed architecture in today's on-premise, private cloud infrastructure and though cloud adoption is growing, businesses still have a need for private cloud infrastructure. To support the on-premise infrastructure, Enterprises also require a scalable data center network that is easy-to-manage. This FlexPod solution expands the existing portfolio of FlexPod solutions by enabling customers to deploy a standards-based, datacenter fabric that can be used in a heterogenous environment. The FlexPod infrastructure in this CVD incorporates a Cisco VXLAN BGP EVPN (Virtual Extensible LAN - Border Gateway Protocol - Ethernet VPN) network architecture to allow for greatly expanded network scale, with the potential to extend that network between locations as a contiguous fabric. This expanded FlexPod solution also includes the AI powered analytics of both Cisco Intersight and NetApp Active IQ from the base FlexPod design for infrastructure management and operational intelligence.

This document describes the deployment of the Cisco and NetApp® FlexPod Datacenter with NetApp ONTAP 9.7 on NetApp AFF A300 storage, Cisco UCS Manager unified software release 4.1(2) with 2nd Generation Intel Xeon Scalable Processors, VMware vSphere 7.0, and Cisco DCNM 11.4(1) managed Cisco VXLAN BGP EVPN design implemented on Cisco Nexus switches running NX-OS 9.3(5). Cisco UCS Manager (UCSM) 4.1(2) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 6400, 2200/2300/2400 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. Cisco DCNM 11 provides multi-tenant, multi-fabric (LAN, SAN) infrastructure management and automation that is optimized for large deployments though it can support smaller and more traditional network architectures as well. Also included are Cisco Intersight and NetApp Active IQ SaaS management platforms.

Solution Overview

Introduction

The industry trend in today's data center design is to move away from application silos and towards a shared infrastructure by using virtualization and pre-validated IT platforms to quickly deploy resources, 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. This FlexPod Datacenter solution with NetApp ONTAP 9.7, Cisco UCS unified software release 4.1(2), and VMware vSphere 7.0 is a predesigned, best-practice datacenter architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, and NetApp AFF A-Series storage arrays running ONTAP® 9.7.

Audience

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

Purpose of this Document

This document provides a step-by-step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS Fabric Interconnects, NetApp AFF storage, and a Cisco DCNM managed VXLAN BGP EVPN network fabric built using Cisco Nexus 9000 series switches.

What's New in this Release?

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

- A highly scalable, standards based VXLAN BGP EVPN data center fabric built using Cisco Nexus 9000 series switches
- Datacenter network deployed a managed as a single fabric using Cisco Data Center Network Manager (Cisco DCNM)-LAN Fabric Version 11.4(1)

This design also parallels the FlexPod Datacenter with VMware vSphere 7.0 CVD in highlighting the following recent features:

- Support for the Cisco UCS 4.1(2) unified software release, Cisco UCS B200-M5 and C220-M5 servers with 2nd Generation Intel Xeon Scalable Processors, and Cisco 1400 Series Virtual Interface Cards (VICs)
- Support for the latest Cisco UCS 6454 and 64108 (supported but not validated) Fabric Interconnects
- Support for the latest Cisco UCS 2408 Fabric Extender
- Addition of Cisco Intersight Software as a Service (SaaS) Management
- Support for the NetApp AFF A300 Storage Controller
- Support for the latest release of NetApp ONTAP® 9.7
- Support for NetApp Virtual Storage Console (VSC) 9.7
- Support for NetApp SnapCenter® and NetApp SnapCenter Plug-in for VMware vSphere 4.3.1

- Support for NetApp Active IQ Unified Manager 9.7
- iSCSI and NFS storage design
- Validation of VMware vSphere 7.0
- Unified Extensible Firmware Interface (UEFI) Secure Boot of VMware ESXi 7.0

Solution Design

VXLAN Single-Site FlexPod includes NetApp All Flash FAS storage, Cisco Nexus® networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package.

> Fibre Channel connectivity is not implemented within this architecture, but is not in conflict with the design, and can be considered a valid option to exist within a parallel SAN network as opposed to using iSCSI.

Figure 1 shows the FlexPod VXLAN Single-Site solution components and network connections for a configuration with the Cisco UCS 6454 Fabric Interconnects. This design has port-channeled 25 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects via the Cisco UCS 2408 Fabric Extenders, port-channeled 25 Gb Ethernet connections between the Cisco UCS C-Series rackmounts and the Cisco UCS Fabric Interconnects, and 100 Gb Ethernet connections between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000 series leaf and spine switches in the fabric, with 40 Gb Ethernet used between the Cisco Nexus 9000 and NetApp AFF A300 storage array. The reference architecture reinforces the " wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Topology

Figure 1. FlexPod with Cisco UCS 6454 Fabric Interconnects and NetApp AFF A300 Series

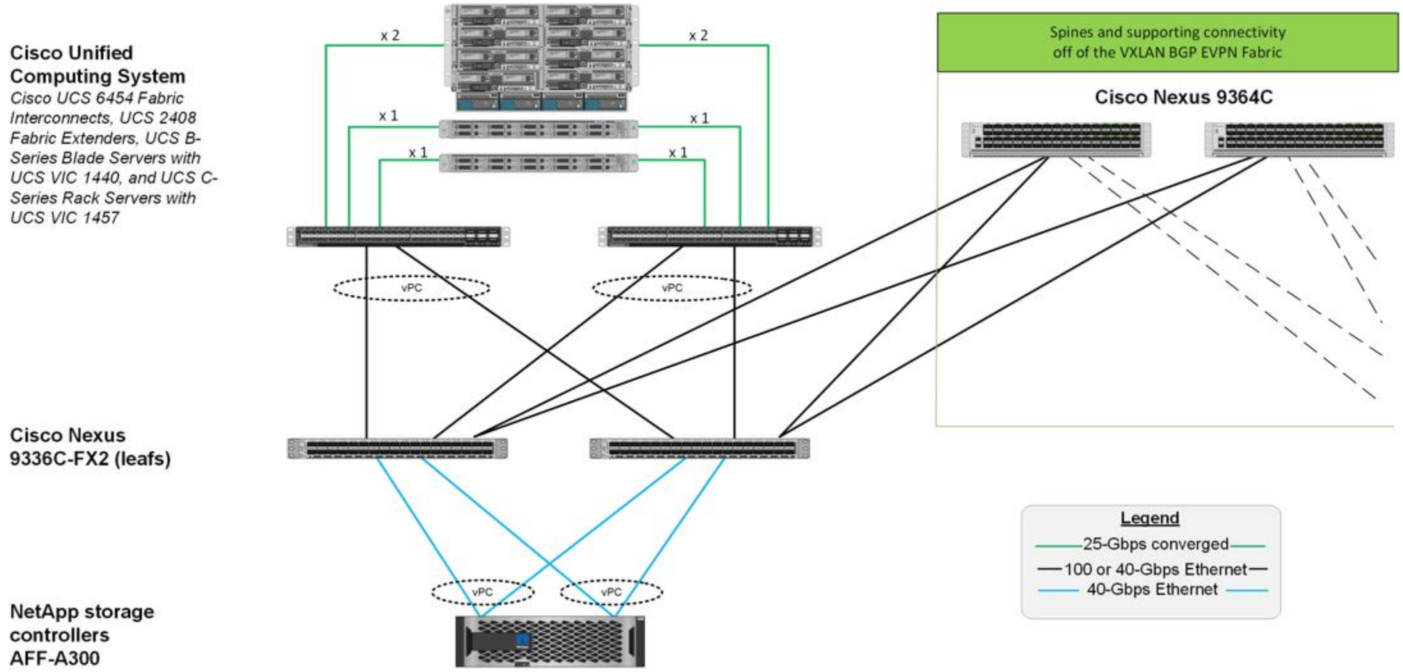

Cisco Unified Computing System
Cisco UCS 6454 Fabric Interconnects, UCS 2408 Fabric Extenders, UCS B-Series Blade Servers with UCS VIC 1440, and UCS C-Series Rack Servers with UCS VIC 1457
x 2
x 2
x 1
x 1
x 1
x 1
vPC
vPC
Spines and supporting connectivity off of the VXLAN BGP EVPN Fabric
Cisco Nexus 9364C
Cisco Nexus 9336C-FX2 (leafs)
vPC
vPC
NetApp storage controllers AFF-A300
Legend
25-Gbps converged
100 or 40-Gbps Ethernet
40-Gbps Ethernet

The reference hardware configuration includes:

- Two Cisco Nexus 9336C-FX2 leaf switches
- Two Cisco Nexus 9364C spine switches
- Two Cisco UCS 6454 fabric interconnects
- One NetApp AFF A300 (HA pair) running ONTAP 9.7

The FlexPod converged infrastructure will typically end at the connecting leaf switches. The Cisco Nexus 9364C spine switches are included for reference of the configuration required for deploying the fabric. In the Network Deployment section that will follow there is additional equipment that is brought up to include a set of Nexus 93180LC-EX border leafs within the fabric and a pair of Nexus 7K switches that represent the primary connection external to the fabric. The use of border leafs are a best practice, but should not be considered a requirement, and the Cisco Nexus 7Ks stand in as an example of existing network infrastructure with options varying depending on meeting the configuration requirements.

Deployment Hardware and Software

Table 1 lists the hardware components and software revisions used for validating this solution.

Table 1. Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6454	4.1(2a)	Includes the Cisco UCS Manager and Cisco UCS VIC 1455
Network Fabric	Cisco Nexus 9364C NX-OS	9.3(5)	Spine switches
	Cisco Nexus 9336C-FX2 NX-OS	9.3(5)	Leaf switches
Storage	NetApp AFF A300	ONTAP 9.7	
Software	Cisco UCS Manager	4.1(2)	
	VMware vSphere	7.0	
	VMware ESXi nenic Ethernet Driver	1.0.33.0	
	NetApp Virtual Storage Console (VSC) / VASA Provider Appliance	9.7.1	
	NetApp SnapCenter for vSphere	4.3.1	Includes SnapCenter Plug-in for VMware vSphere
	NetApp NFS Plug-in for VMware VAAI	1.1.2-3	
	NetApp Active IQ Unified Manager	9.7P1	
Management	Cisco Intersight	N/A	

Layer	Device	Image	Comments
	Cisco Data Center Network Manager (LAN Fabric)	11.4(1)	
	NetApp Active IQ	N/A	

Configuration Guidelines

This document explains how to configure a fully redundant, highly available configuration for a FlexPod unit with ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-01, VM-Host-Infra-02 to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

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

Example:

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

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

Table 2. Necessary VLANs

VLAN Name	VLAN Purpose	ID used in Validating this Document
Out-of-Band Mgmt	Out-of-band management interfaces	163
Site1-IB	In-band management interfaces	122
Common-Services	Example network for shared resources, used by vCenter and AD in this design	322
Native	untagged frames are assigned	2
iSCSI-A	iSCSI A traffic	3010
iSCSI-B	iSCSI B traffic	3020
Infra-NFS	Infrastructure NFS traffic	3050

VLAN Name	VLAN Purpose	ID used in Validating this Document
vMotion	VMware vMotion	3000
VM-Traffic-1	Production VM Interfaces	1001
VM-Traffic-2	Production VM Interfaces	1002
VM-Traffic-3	Production VM Interfaces	1003

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, a cabling diagram was used.

The cabling diagram in this section contains details for the prescribed and supported configuration of the NetApp AFF 300 running NetApp ONTAP® 9.7.

For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

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

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to NetApp Support.

Figure 2 details the cable connections used in the validation lab for the FlexPod topology based on the Cisco UCS 6454 fabric interconnect. 40/100Gb links connect the Cisco UCS Fabric Interconnects to and within the VXLAN fabric of the Cisco Nexus Switches, and 40Gb links connect the NetApp AFF controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has a connection to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (Site1-IB) Management Subnets.

Figure 2. FlexPod Cabling with Cisco UCS 6454 Fabric Interconnect

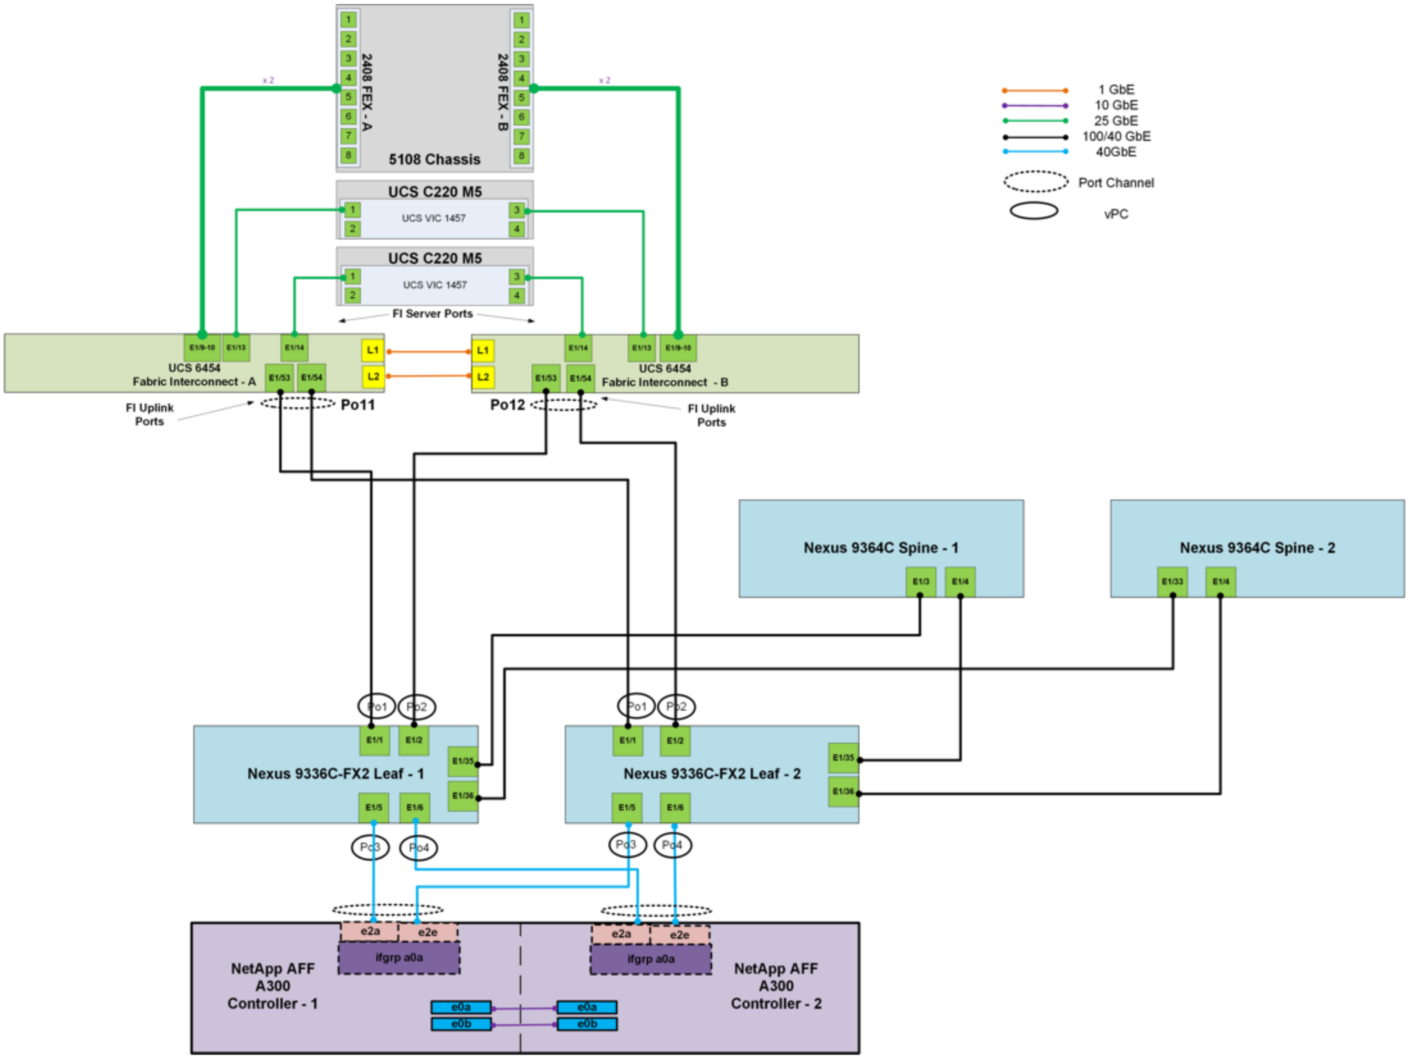

Solution Deployment - Network Fabric

This section provides a detailed step-by-step procedure for deploying a Cisco VXLAN BGP EVPN fabric to enable network connectivity between FlexPod storage, compute, and other components in the solution. The VXLAN fabric in this solution will deployed and managed by a Cisco Data Center Network Manager (Cisco DCNM). The network fabric will consist of single data center site with different models of Cisco Nexus spine and leaf switches. The network fabric used in this solution consists of a one pair of spine switches and two pairs of leaf switches. The Cisco UCS domains and NetApp storage arrays will connect to same leaf switch pair in this design. The other leaf switch pair in the design serve as Border leaf switch for connectivity outside the fabric. The separate leaf switch pairs for each role (access/TOR vs. border) ensures a scalable VXLAN fabric.

This design assumes a greenfield deployment of the VXLAN BGP EVPN fabric. Customers with an existing VXLAN fabric can use portions of the deployment discussed in this section to add new switches or align with this FlexPod design.

Deployment Overview

A high-level overview of the steps involved in deploying a single-site network fabric is provided below.

- **Physical Connectivity:** Complete the physical connectivity as outlined in the FlexPod Cabling section.
- **Cisco Nexus Switches - Base Setup and Configuration:** Bring-up Cisco Nexus switches with a minimal version of software that supports Cisco DCNM and VXLAN EVPN fabric and perform minimal setup and configuration so that they can be imported into the fabric by Cisco DCNM. The minimal configuration includes setting the Hostname, OOB Management IP and Gateway, Admin account and password, and setting boot variable for booting a valid image. The base setup and configuration is outside the scope of this document - please see relevant Nexus product documentation for how this can be done.
- **Out-of-Band (OOB) Management Connectivity:** Complete all the out-of-band management connectivity for the spine and leaf switches in the network fabric. Enabling OOB connectivity is outside the scope of this document - see relevant Nexus product documentation for setting this up.
- **Deploy Cisco DCNM**: Deploy Cisco DCNM LAN Fabric and enable OOB management connectivity to the spine and leaf switches in the fabric. Cisco DCNM will discover the switches, deploy the VXLAN BGP EVPN fabric and provide a centralized management portal for day-2 operation and management of the fabric. Deployment of Cisco DCNM is also outside the scope of this document – see Cisco DCNM 11.4(1) documentation on cisco.com for additional details.
- **Licensing:** Procure necessary licensing for Cisco DCNM and Nexus switches and configure the licenses before the available grace-period expires to fully utilize all services provided by this Cisco environment.
- **Deploy VXLAN BGP EVPN Fabric using Cisco DCNM:** Cisco DCNM’s Fabric Builder is used to configure and deploy the VXLAN BGP EVPN fabric in Site-A. To deploy the fabric configuration to the switches, the spine and leaf switches must be first discovered and added to Cisco DCNM. Cisco DCNM can will then deploy the IP underlay and VXLAN overlay across all the switches that make up the data center fabric in Site-A.
- **External or Outside Connectivity:** Enable connectivity from VXLAN fabric in Site-A to outside networks. In this design, these are any networks that are outside the VXLAN fabric in Site-A – they can be either internal or external to the Enterprise. In this design, this connectivity enabled reachability to key services host-

ed outside the fabric such as Microsoft Active Directory, DNS within the Enterprise, and services outside the Enterprise such as Cisco Intersight and Cisco Umbrella in the public cloud.

- **Access Layer Connectivity to NetApp Storage Cluster:** Enable access-layer connectivity from the VXLAN fabric in Site-A to the NetApp Storage infrastructure in the solution. The NetApp storage infrastructure in this solution consists of an AFF A300 storage array.
- **Access Layer Connectivity to Cisco UCS Domain:** Enable access-layer connectivity from the VXLAN fabric in Site-A to the Cisco UCS infrastructure in the solution. The Cisco UCS infrastructure in this solution consists of a pair of Cisco UCS Fabric Interconnects that connect to Cisco UCS B-series and C-series servers.
- **FlexPod Infrastructure Connectivity:** A dedicated tenant is defined in this design to enable the infrastructure connectivity in the FlexPod VSI solution. A FlexPod Foundation Tenant is configured to enable connectivity for FlexPod Compute and Storage infrastructure. In this design, the Foundation tenant will provide infrastructure connectivity to enable the FlexPod Virtual Server Infrastructure (VSI). This tenant is not used for applications workloads hosted on the FlexPod VSI, though it is used by management components such as VMware vCenter, NetApp VSC and so on. that is used to manage and operate the FlexPod VSI.
- **On-board multi-tier applications**: A separate application tenant is defined in the VXLAN fabric to meet the connectivity needs of the applications hosted on the FlexPod VSI. Expanded tenant separation is possible within Cisco UCS and NetApp storage, but is not discussed in depth within this design.

Deploy VXLAN BGP EVPN Fabric using Cisco DCNM

This section uses Cisco DCNM's LAN Fabric Builder to configure and deploy a VXLAN BGP EVPN fabric in Site-A (or Site-1). The LAN Fabric Builder in Cisco DCNM creates and manages a software-defined (SDN) fabric by selecting an existing fabric or by defining a new VXLAN fabric. The switches can be discovered and added to the fabric using Power On Auto Provisioning (POAP), or by directly importing switches (with a base configuration). You can then set the roles of the switches, pre-select the fabric settings, and then use one-click **Save & Deploy** to deploy the configuration and bring up a fully functional VXLAN BGP EVPN fabric that spans any number of spine and leaf switches.

Topology

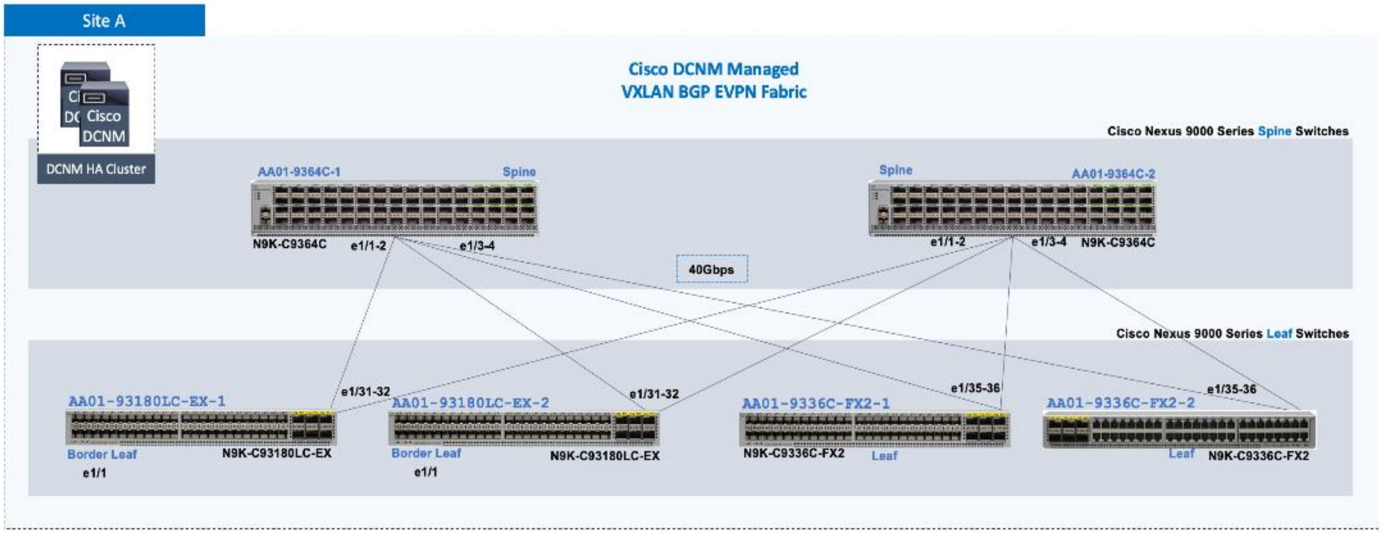

Topology figure above shows the connectivity of the Cisco Nexus 9364C Spines and Cisco Nexus 9336C-FX2 Leafs in the validation. Also pictured is the connection to a pair of Cisco Nexus 93180LC-EX Leafs that are used as border leaf switches for connectivity outside of the fabric.

Setup Information

The VXLAN BGP EVPN fabric configuration settings used for deploying the Site-A data center fabric is provided in the table below.

Table 3. Fabric Configuration Information – Site-A

Data Center	Parameters	Default Parameters	Notes
Fabric Name	Site-A	_	
Fabric Template	Easy_Fabric_11_1	_	
General Tab			
BGP ASN	65001	_	
NX-OS Software Image	9.3(5)	_	Optional (If Set, Image Version Check Enforced On All Switches. Images can be uploaded by going to Control > Image Upload)
Protocols Tab			

Data Center	Parameters	Default Parameters	Notes
Underlay Routing Protocol Tag	Site-A_UNDERLAY	UNDERLAY	
Resources Tab			
Underlay Routing Loopback IP Range	10.11.0.0/24	10.2.0.0/22	Optional (Default Values can be used as-is)
Underlay VTEP Loopback IP Range	10.11.1.0/24	10.3.0.0/22	Optional (Default Values can be used as-is)
Underlay RP Loopback IP Range	10.11.254.0/24	10.254.254.0/24	Optional (Default Values can be used as-is)
Underlay Subnet IP Range	10.11.3.0/22	10.4.0.0/16	Optional (Default Values can be used as-is)
Layer 2 VXLAN VNI Range	20000-24999	30000-49000	Optional (Default Values can be used as-is)
Layer 3 VXLAN VNI Range	30000-34999	50000-59000	Optional (Default Values can be used as-is)
Network VLAN Range	3000-3499	2300-2999	Optional (Default Values can be used as-is)
VRF VLAN Range	3500-3967	2000-2299	Optional (Default Values can be used as-is)
VRF Lite Deployment	ToExternalOnly	Manual	Optional (Default Values can be used as-is)
Auto Deploy Both	☑	□	Optional (Default Values can be used as-is)

Data Center	Parameters	Default Parameters	Notes
VRF Lite Subnet IP Range	10.11.99.0/24	10.33.0.0/16	Optional (Default Values can be used as-is)
VRF Lite Subnet Mask	30	30	Optional (Default Values can be used as-is)
Service Network VLAN Range	1500-1599	3000-3199	Optional (Default Values can be used as-is)
Manageability Tab			
NTP Server IPs	172.26.163.254	–	Optional
NTP Server VRFs	management	–	Optional
Configuration Backup Tab			
Hourly Fabric Backup	☑	□	Optional

The setup information for discovering the spine and leaf switches in the Site-A datacenter fabric is provided in the table below. Cisco DCNM also supports discovery and importing of fabric switches through Power-on-Auto-Provisioning(POAP) – however, POAP was not utilized in this CVD.

Table 4. Discovery Information - Site-A

Hostname	Switch Role	IP Address (OOB)	Notes
AA01-9364C-1	Spine	172.26.163.231/24	
AA01-9364C-2	Spine	172.26.163.232/24	
AA01-9336C-FX2-1	Leaf	172.26.163.223/24	Top-of-Rack (TOR) switch for access layer connectivity to Cisco UCS compute and NetApp storage
AA01-9336C-FX2-2	Leaf	172.26.163.224/24	Top-of-Rack (TOR) switch for access layer connectivity to Cisco UCS compute and NetApp storage
AA01-93180LC-EX-1-1	Border Leaf	172.26.163.221/24	

Hostname	Switch Role	IP Address (OOB)	Notes
AA01-93180LC-EX-1-2	Border Leaf	172.26.163.222/24	

The devices used in this design were pre-configured with a Hostname, Management IP address, username, password, and boot variable.

Create VXLAN Data Center Fabric in Site-A

To create the VXLAN BGP EVPN datacenter fabric in Site-A, go to the Setup Information section to follow these steps:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account with full access to the data center.

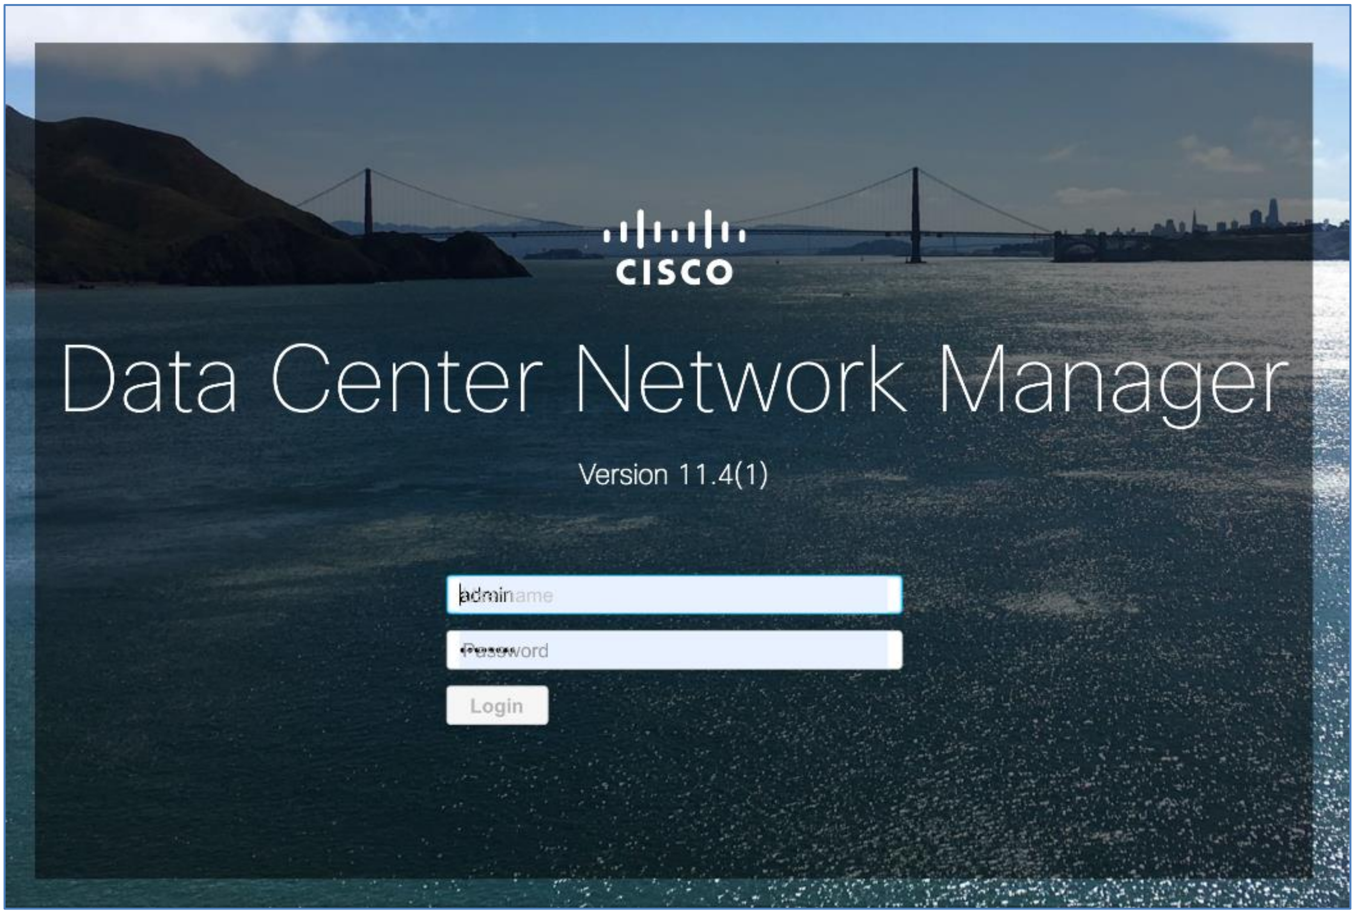

2. From the right window pane, select and click the **Fabric Builder** icon.

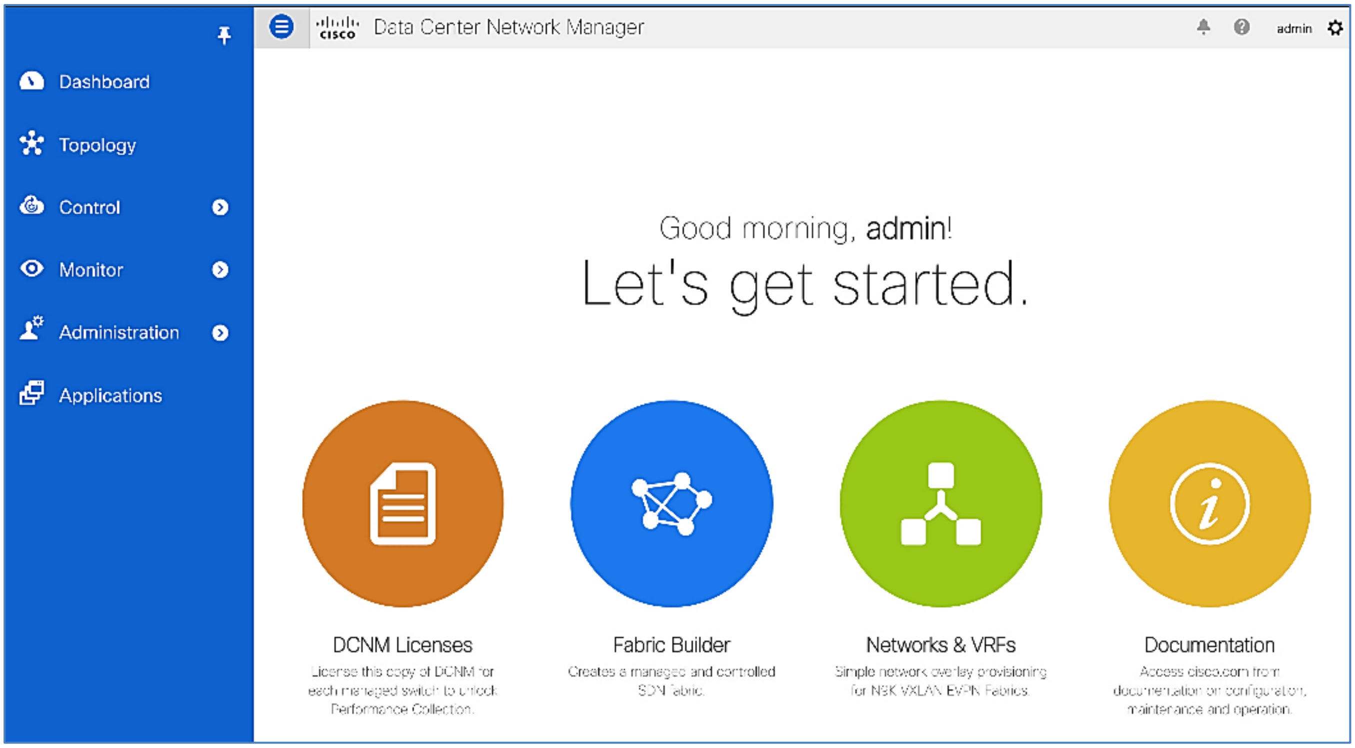

Good morning, admin!

Let's get started.

DCNM Licenses

License this copy of DCNM for each managed switch to unlock Performance Collection.

Fabric Builder

Creates a managed and controlled SDN fabric.

Networks & VRFs

Simple network overlay provisioning for N9K VXLAN EVPN Fabrics.

Documentation

Access cisco.com from documentation on configuration, maintenance and operation.

3. From the right window pane, click the **Create Fabric** icon.

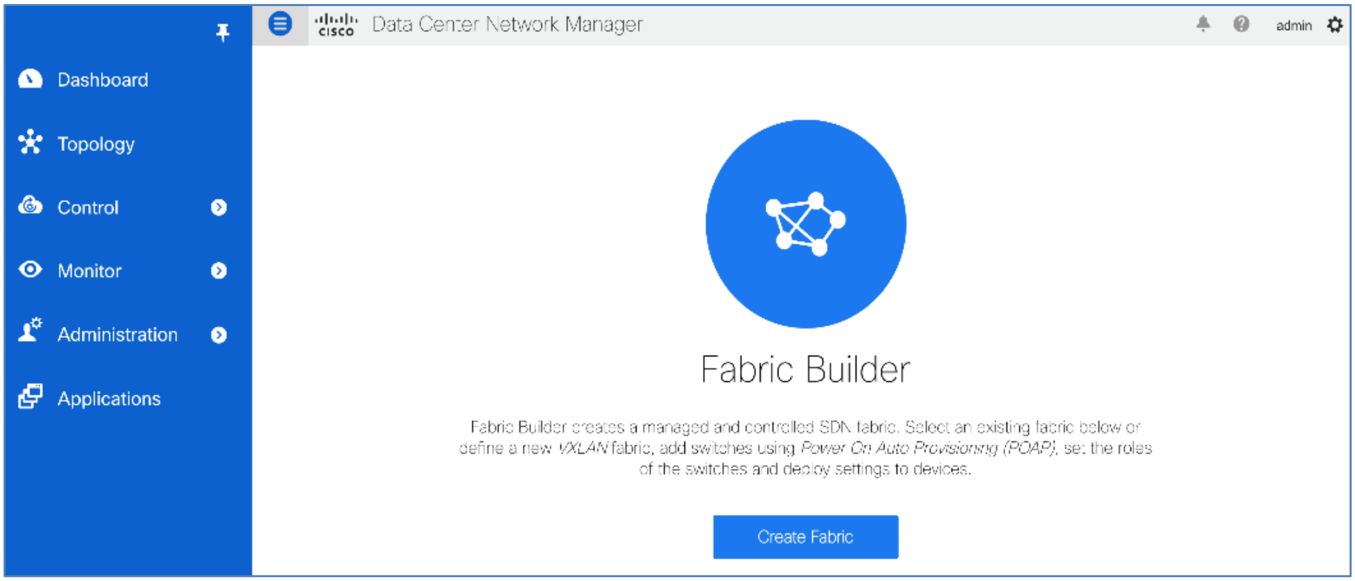

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

4. In the **Add Fabric** pop-up window, specify a **Fabric Name** and select the **Fabric Template** specified in Table 3 above from the drop-down list.

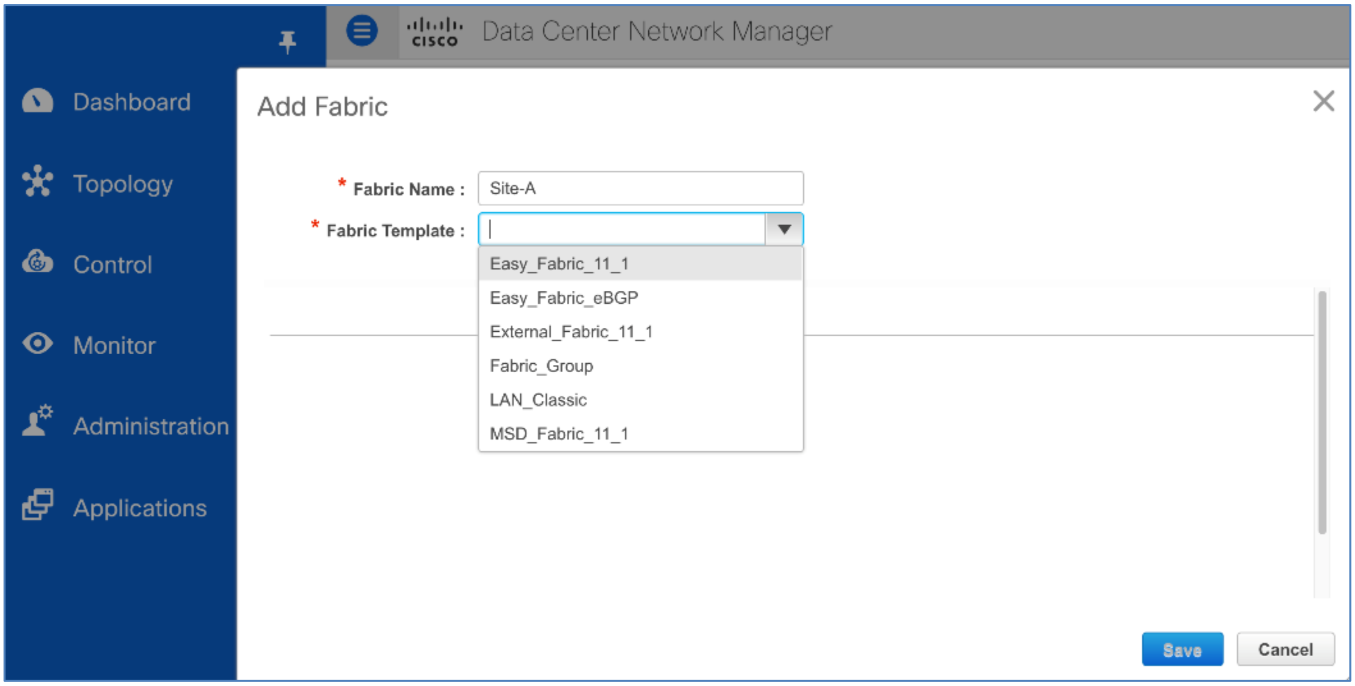

Add Fabric

* Fabric Name : Site-A

* Fabric Template :

5. The pop-up window will now expand to include multiple tabs for configuring the fabric. In the **General** tab, specify the **BGP ASN** information for Site-A and **NX-OS Software Image Version** (Optional) from the drop-down list.

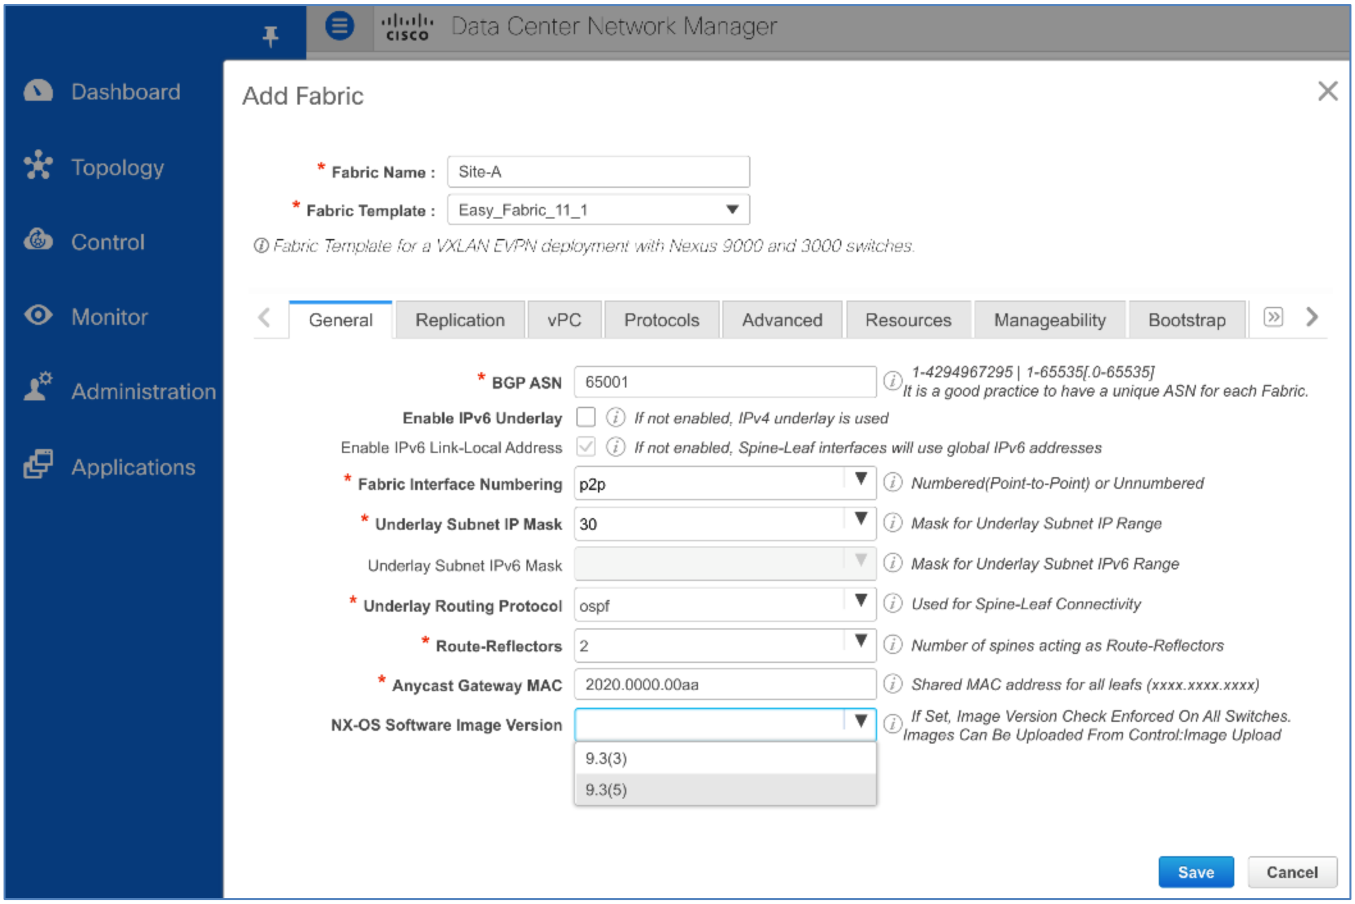

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

* BGP ASN: 65001 — 1-4294967295 | 1-65535[.0-65535] It is a good practice to have a unique ASN for each Fabric.

Enable IPv6 Underlay: □ — If not enabled, IPv4 underlay is used

Enable IPv6 Link-Local Address: ☑ — If not enabled, Spine-Leaf interfaces will use global IPv6 addresses

* Fabric Interface Numbering: p2p — Numbered(Point-to-Point) or Unnumbered

* Underlay Subnet IP Mask: 30 — Mask for Underlay Subnet IP Range

Underlay Subnet IPv6 Mask: — Mask for Underlay Subnet IPv6 Range

* Underlay Routing Protocol: ospf — Used for Spine-Leaf Connectivity

* Route-Reflectors: 2 — Number of spines acting as Route-Reflectors

* Anycast Gateway MAC: 2020.0000.00aa — Shared MAC address for all leafs (xxxx.xxxx.xxxx)

NX-OS Software Image Version: — If Set, Image Version Check Enforced On All Switches. Images Can Be Uploaded From Control:Image Upload

9.3(3)
9.3(5)

6. In the Replication tab, leave everything as-is. Alternatively, you can customize the default options selected by Cisco DCNM for Replication Mode, Multicast Group Subnet, Rendezvous-Points (RP), RP mode and other settings as needed.

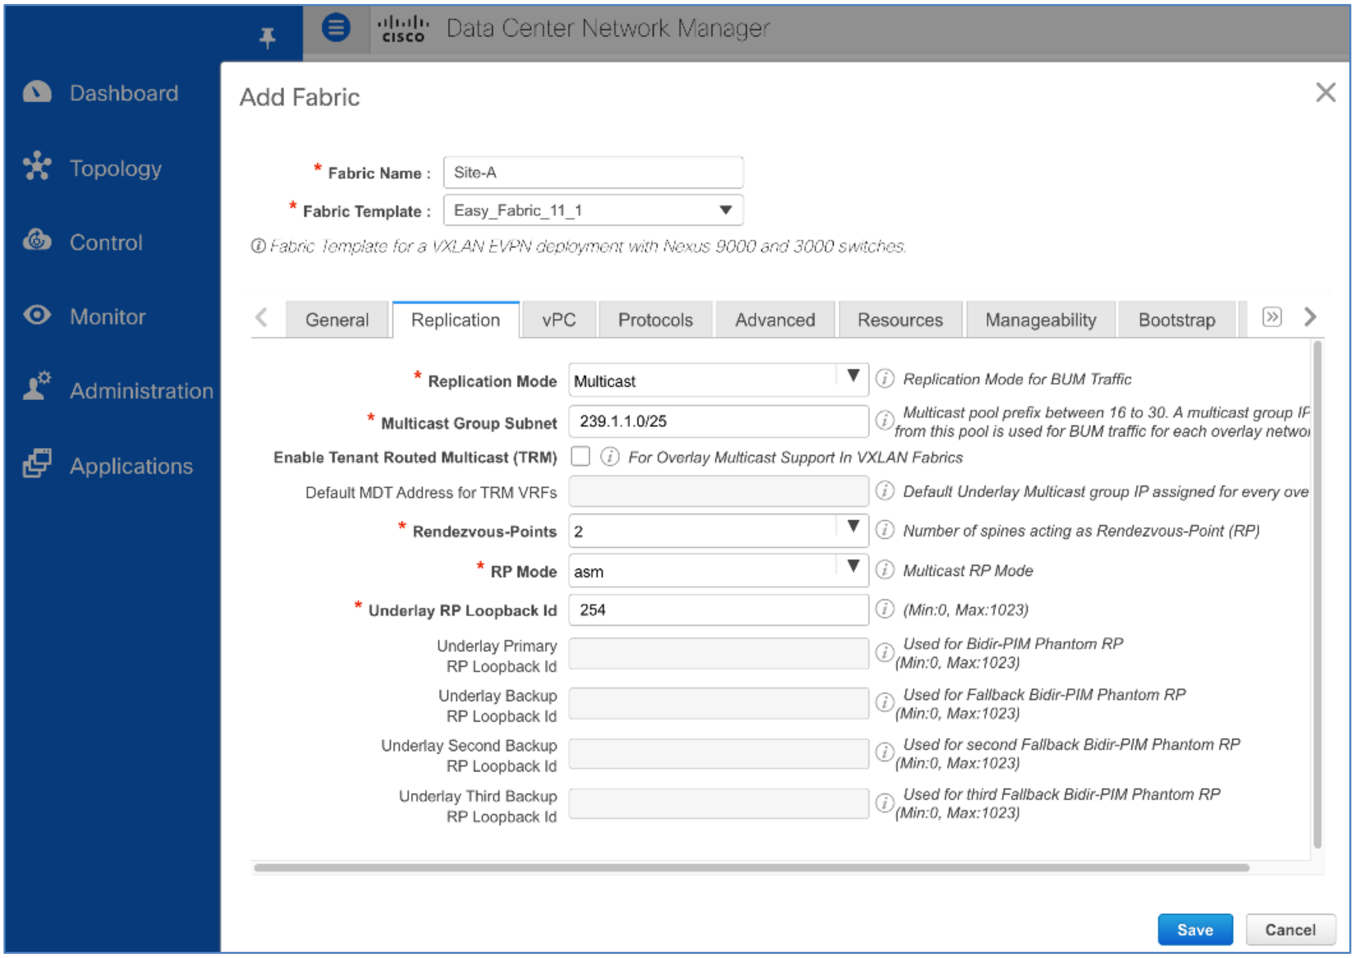

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

* Replication Mode: Multicast — *Replication Mode for BUM Traffic*

* Multicast Group Subnet: 239.1.1.0/25 — *Multicast pool prefix between 16 to 30. A multicast group IP from this pool is used for BUM traffic for each overlay netwo*

Enable Tenant Routed Multicast (TRM): □ — *For Overlay Multicast Support In VXLAN Fabrics*

Default MDT Address for TRM VRFs: — *Default Underlay Multicast group IP assigned for every ove*

* Rendezvous-Points: 2 — *Number of spines acting as Rendezvous-Point (RP)*

* RP Mode: asm — *Multicast RP Mode*

* Underlay RP Loopback Id: 254 — *(Min:0, Max:1023)*

Underlay Primary RP Loopback Id: — *Used for Bidir-PIM Phantom RP (Min:0, Max:1023)*

Underlay Backup RP Loopback Id: — *Used for Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)*

Underlay Second Backup RP Loopback Id: — *Used for second Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)*

Underlay Third Backup RP Loopback Id: — *Used for third Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)*

7. In the **vPC** tab, leave everything as-is. Alternatively, you can customize the selected default options and other settings as needed.

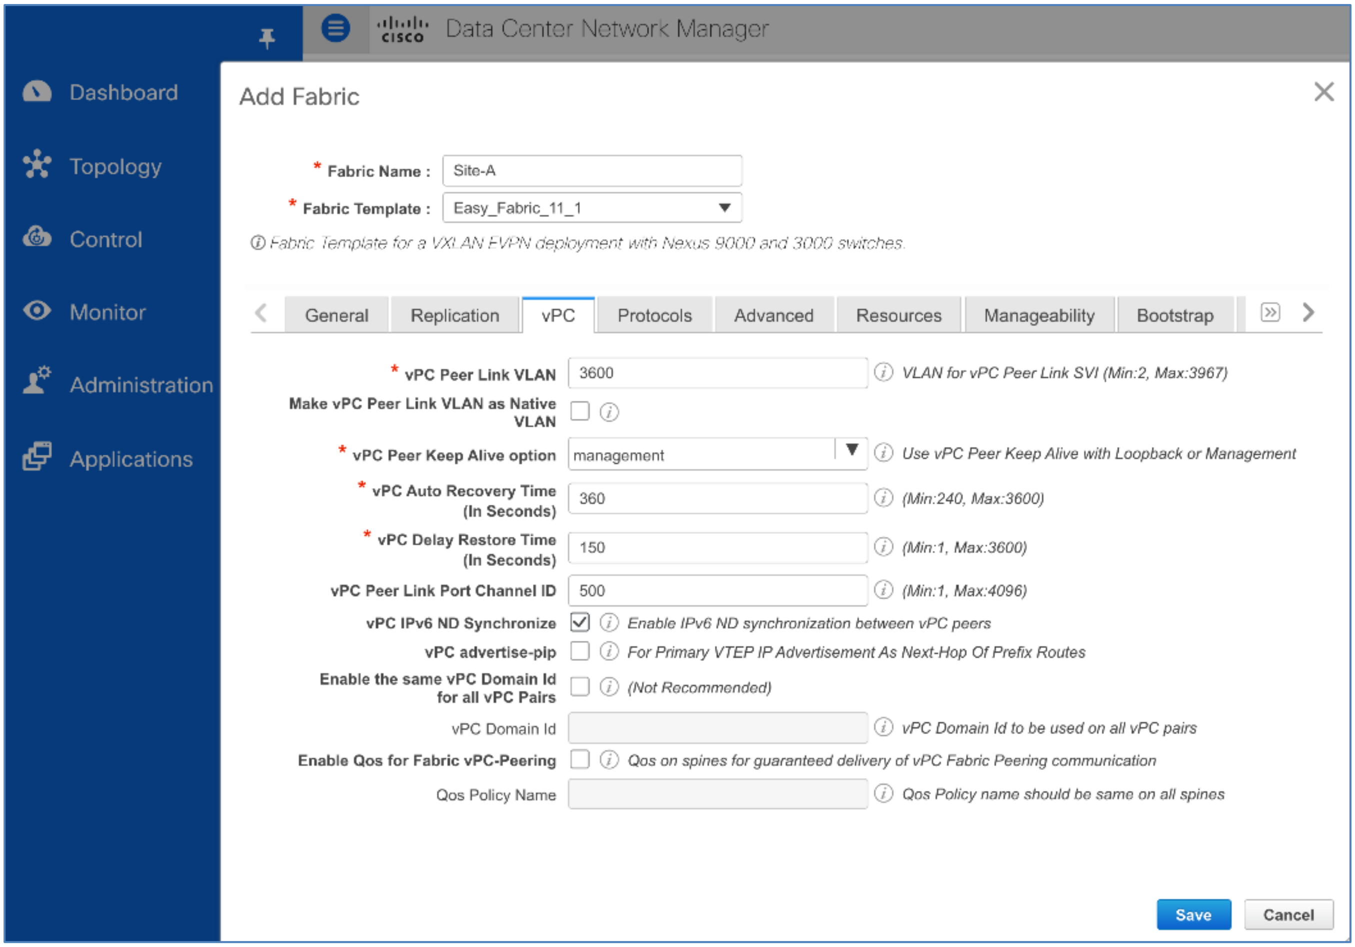

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

* vPC Peer Link VLAN	3600	VLAN for vPC Peer Link SVI (Min:2, Max:3967)
Make vPC Peer Link VLAN as Native VLAN	□	
* vPC Peer Keep Alive option	management	Use vPC Peer Keep Alive with Loopback or Management
* vPC Auto Recovery Time (In Seconds)	360	(Min:240, Max:3600)
* vPC Delay Restore Time (In Seconds)	150	(Min:1, Max:3600)
vPC Peer Link Port Channel ID	500	(Min:1, Max:4096)
vPC IPv6 ND Synchronize	☑	Enable IPv6 ND synchronization between vPC peers
vPC advertise-pip	□	For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes
Enable the same vPC Domain Id for all vPC Pairs	□	(Not Recommended)
vPC Domain Id		vPC Domain Id to be used on all vPC pairs
Enable Qos for Fabric vPC-Peering	□	Qos on spines for guaranteed delivery of vPC Fabric Peering communication
Qos Policy Name		Qos Policy name should be same on all spines

8. In the **Protocols** tab, specify the **Underlay Routing Protocol Tag** for Site-A. Leave everything else as-is. Alternatively, you can enable OSPF, ISIS and/or BGP authentication for additional security, enable Bi-directional Forward Detection (BFD) for quicker notification of failures to upper layer protocols such as BGP, OSPF, PIM and so on. You can also add additional customization for the iBGP configuration using the free-form templates provided at the bottom.

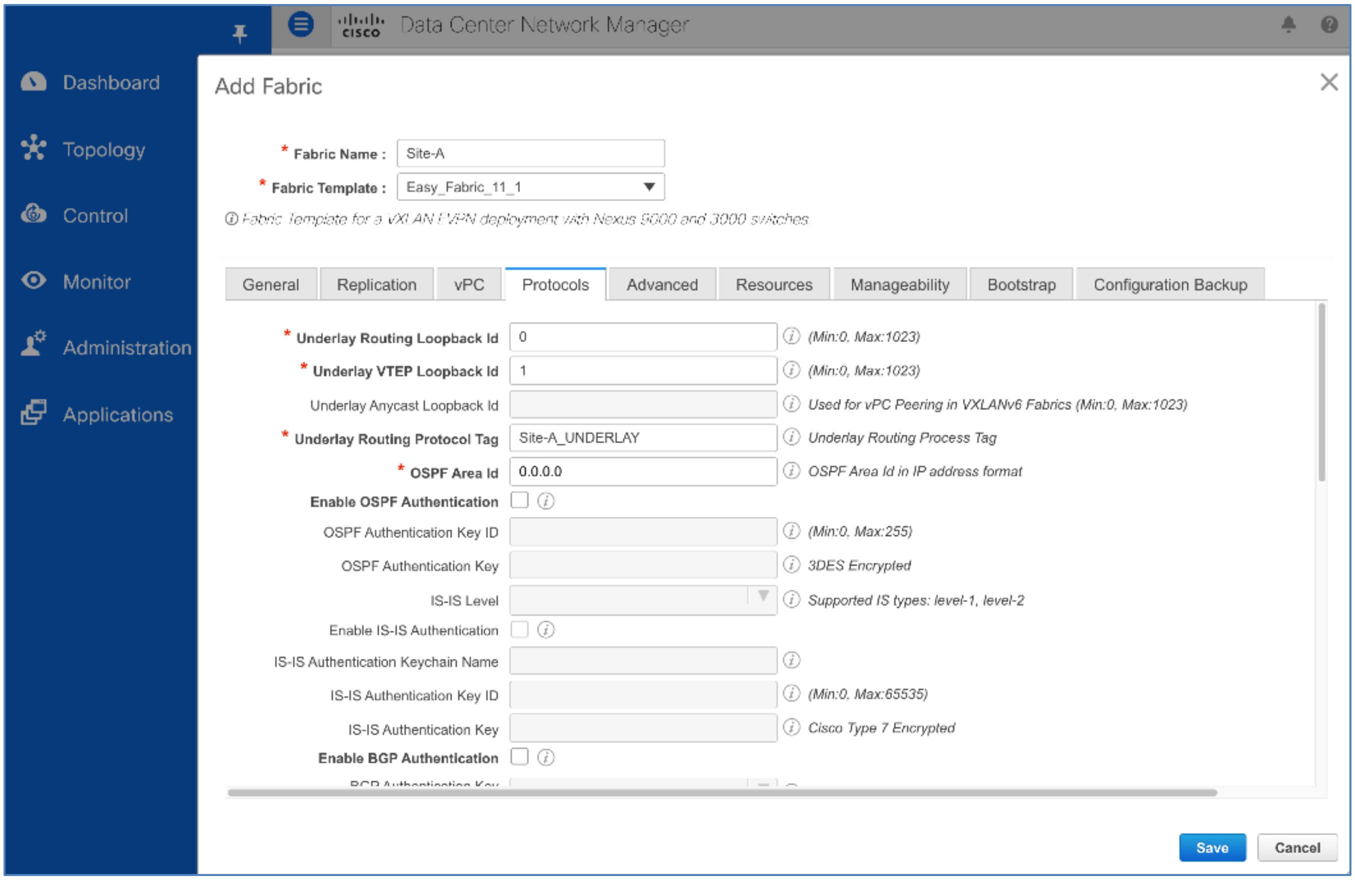

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches

* Underlay Routing Loopback Id: 0 ⓘ (Min:0, Max:1023)

* Underlay VTEP Loopback Id: 1 ⓘ (Min:0, Max:1023)

Underlay Anycast Loopback Id: ⓘ Used for vPC Peering in VXLANv6 Fabrics (Min:0, Max:1023)

* Underlay Routing Protocol Tag: Site-A_UNDERLAY ⓘ Underlay Routing Process Tag

* OSPF Area Id: 0.0.0.0 ⓘ OSPF Area Id in IP address format

Enable OSPF Authentication □ ⓘ

OSPF Authentication Key ID: ⓘ (Min:0, Max:255)

OSPF Authentication Key: ⓘ 3DES Encrypted

IS-IS Level: ⓘ Supported IS types: level-1, level-2

Enable IS-IS Authentication □ ⓘ

IS-IS Authentication Keychain Name: ⓘ

IS-IS Authentication Key ID: ⓘ (Min:0, Max:65535)

IS-IS Authentication Key: ⓘ Cisco Type 7 Encrypted

Enable BGP Authentication □ ⓘ

BGP Authentication Key

9. In the **Advanced** tab, leave everything as-is. Alternatively, you can customize the selected default options and other settings as needed. Note that the **Site Id** matches the **BGP ASN** specified in the **General** tab. Also, the **Interface MTU** is pre-configured to use a jumbo MTU of 9216 across all fabric links and on interfaces connecting to endpoints.

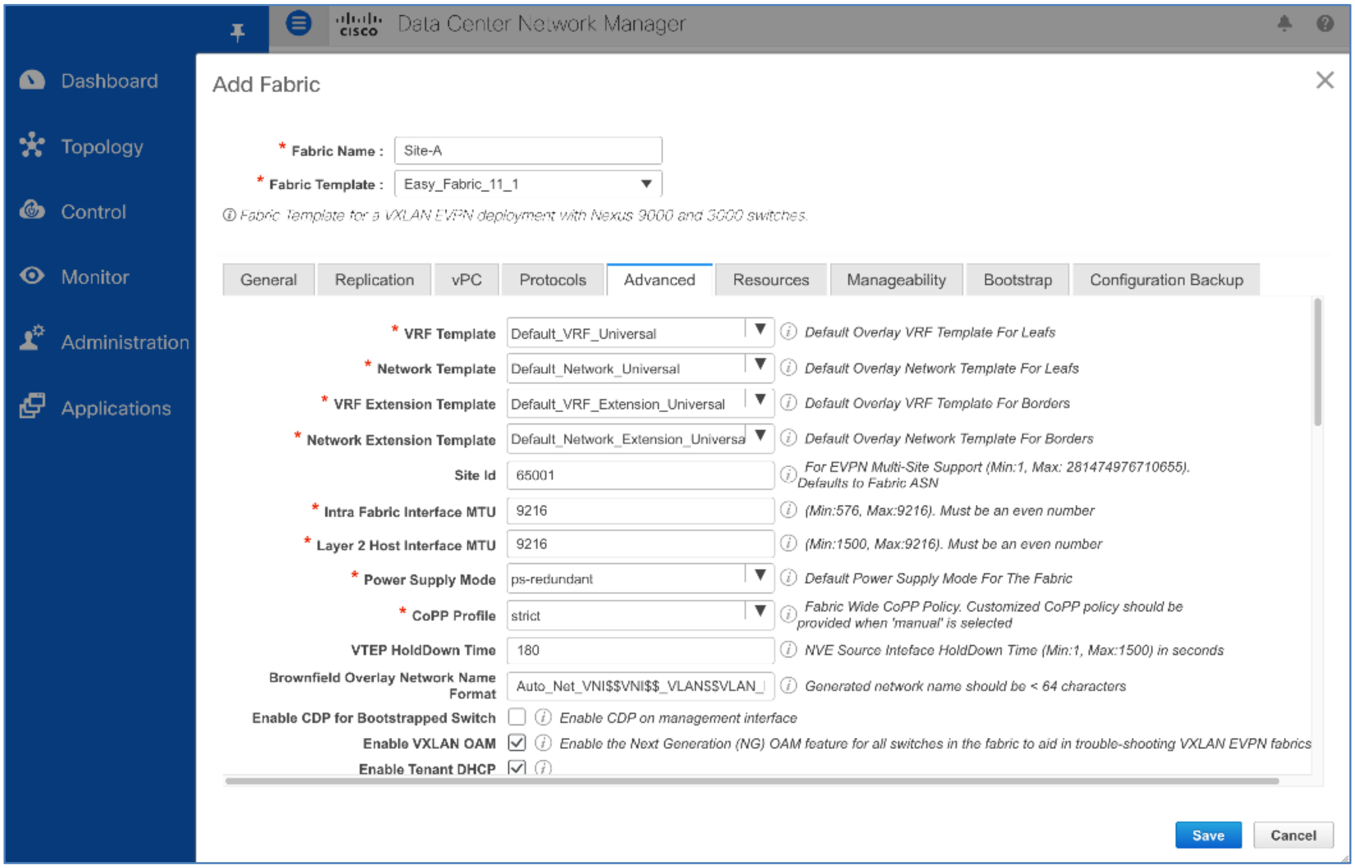

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

* VRF Template	Default_VRF_Universal	ⓘ Default Overlay VRF Template For Leafs
* Network Template	Default_Network_Universal	ⓘ Default Overlay Network Template For Leafs
* VRF Extension Template	Default_VRF_Extension_Universal	ⓘ Default Overlay VRF Template For Borders
* Network Extension Template	Default_Network_Extension_Universa	ⓘ Default Overlay Network Template For Borders
Site Id	65001	ⓘ For EVPN Multi-Site Support (Min:1, Max: 281474976710655). Defaults to Fabric ASN
* Intra Fabric Interface MTU	9216	ⓘ (Min:576, Max:9216). Must be an even number
* Layer 2 Host Interface MTU	9216	ⓘ (Min:1500, Max:9216). Must be an even number
* Power Supply Mode	ps-redundant	ⓘ Default Power Supply Mode For The Fabric
* CoPP Profile	strict	ⓘ Fabric Wide CoPP Policy. Customized CoPP policy should be provided when 'manual' is selected
VTEP HoldDown Time	180	ⓘ NVE Source Inteface HoldDown Time (Min:1, Max:1500) in seconds
Brownfield Overlay Network Name Format	Auto_Net_VNI$$VNI$$_VLAN$$VLAN_	ⓘ Generated network name should be < 64 characters
Enable CDP for Bootstrapped Switch	☐	ⓘ Enable CDP on management interface
Enable VXLAN OAM	☑	ⓘ Enable the Next Generation (NG) OAM feature for all switches in the fabric to aid in trouble-shooting VXLAN EVPN fabrics
Enable Tenant DHCP	☑	ⓘ

10. (Optional) In this deployment, the **Greenfield Cleanup Option** is changed from the default of Disable to Enable in order to speed up deployment in the Cisco lab – however, Cisco recommends that customers use the default option that will reload the switches during clean up in Greenfield deployments.

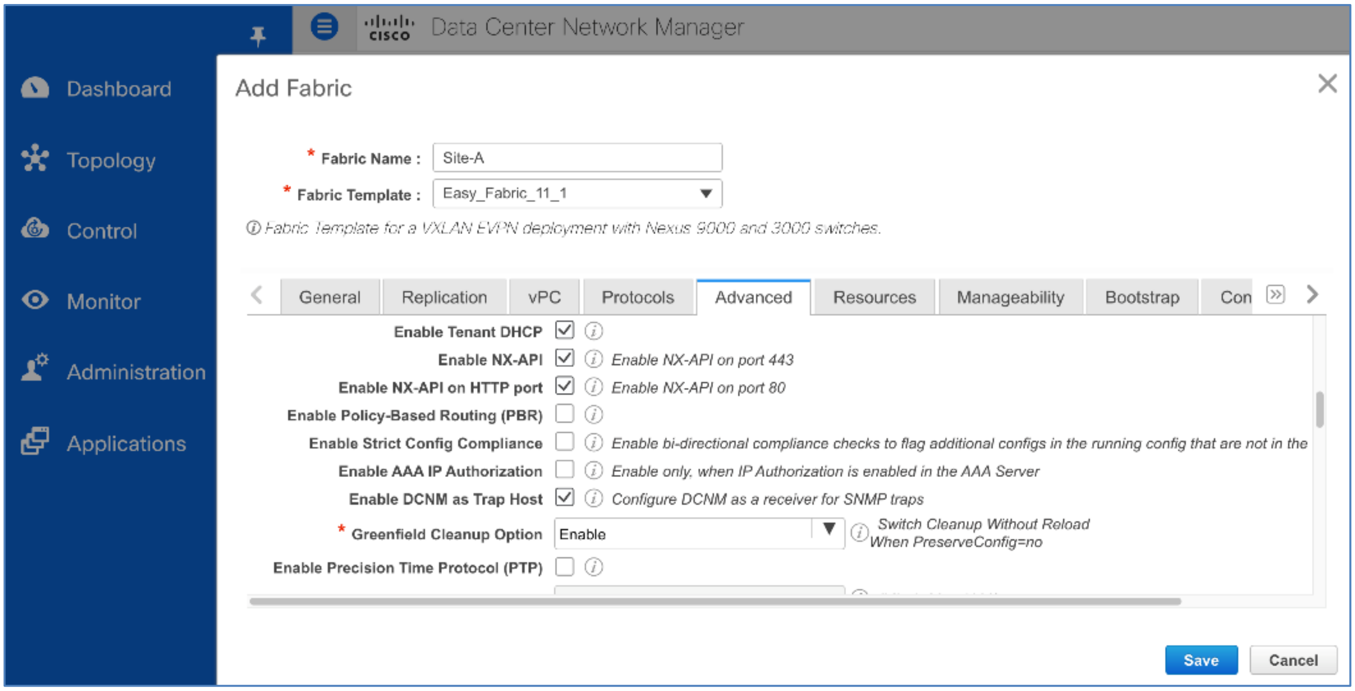

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

Enable Tenant DHCP ☑ ⓘ

Enable NX-API ☑ ⓘ *Enable NX-API on port 443*

Enable NX-API on HTTP port ☑ ⓘ *Enable NX-API on port 80*

Enable Policy-Based Routing (PBR) □ ⓘ

Enable Strict Config Compliance □ ⓘ *Enable bi-directional compliance checks to flag additional configs in the running config that are not in the*

Enable AAA IP Authorization □ ⓘ *Enable only, when IP Authorization is enabled in the AAA Server*

Enable DCNM as Trap Host ☑ ⓘ *Configure DCNM as a receiver for SNMP traps*

* Greenfield Cleanup Option: Enable ⓘ *Switch Cleanup Without Reload When PreserveConfig=no*

Enable Precision Time Protocol (PTP) □ ⓘ

11. (Optional) Customers can also enable PTP and Queueing on core facing interfaces as needed. Note this was not setup in this CVD. PTP is necessary when using Nexus operational tools such as Network Insights – Resources for a more precise and accurate timing of flows in the range of microseconds or nanoseconds.

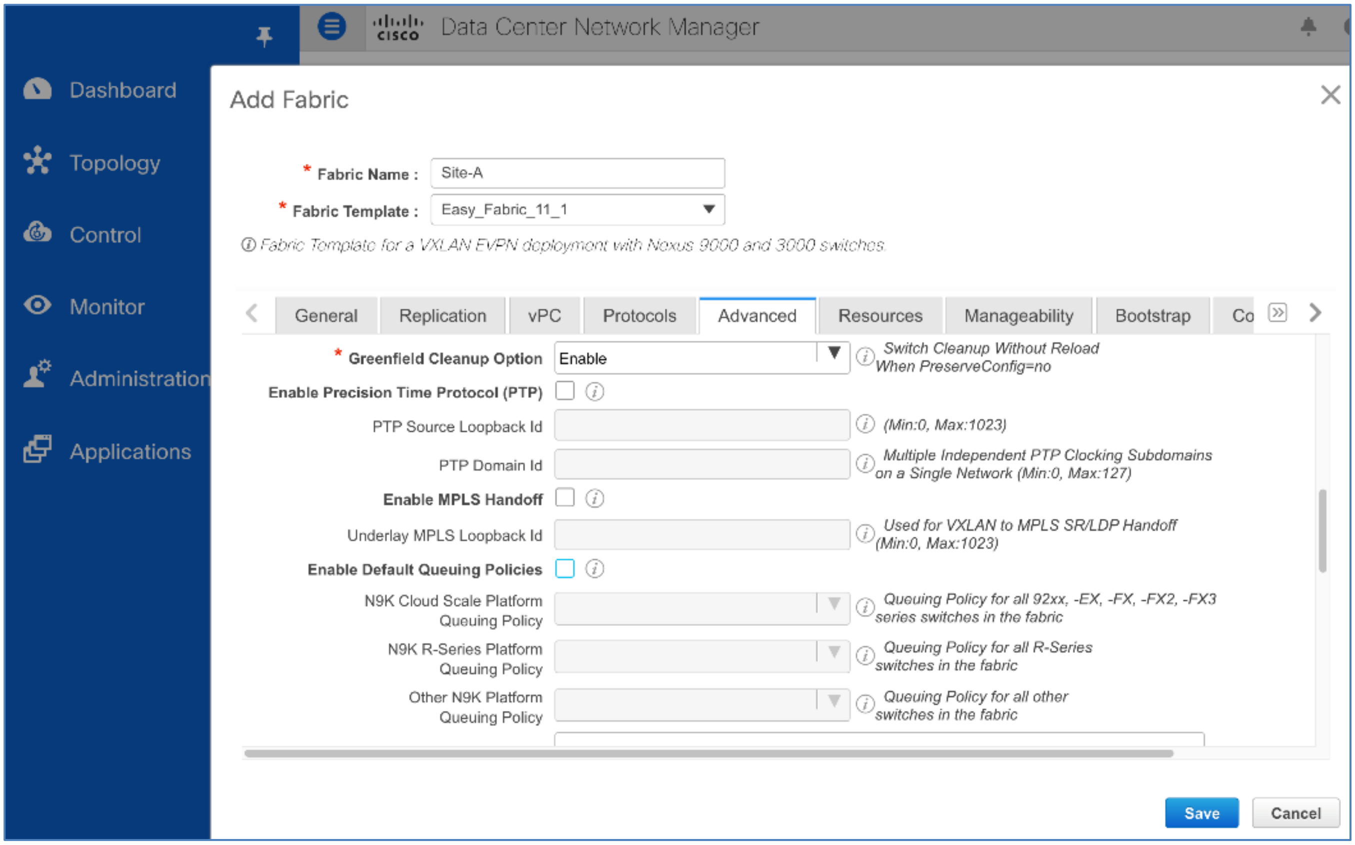

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

* Greenfield Cleanup Option: Enable ⓘ *Switch Cleanup Without Reload When PreserveConfig=no*

Enable Precision Time Protocol (PTP) □ ⓘ

PTP Source Loopback Id ⓘ *(Min:0, Max:1023)*

PTP Domain Id ⓘ *Multiple Independent PTP Clocking Subdomains on a Single Network (Min:0, Max:127)*

Enable MPLS Handoff □ ⓘ

Underlay MPLS Loopback Id ⓘ *Used for VXLAN to MPLS SR/LDP Handoff (Min:0, Max:1023)*

Enable Default Queuing Policies □ ⓘ

N9K Cloud Scale Platform Queuing Policy ⓘ *Queuing Policy for all 92xx, -EX, -FX, -FX2, -FX3 series switches in the fabric*

N9K R-Series Platform Queuing Policy ⓘ *Queuing Policy for all R-Series switches in the fabric*

Other N9K Platform Queuing Policy ⓘ *Queuing Policy for all other switches in the fabric*

12. (Optional) Cisco DCNM also allows for **Freeform** configurations that customers can use for additional configuration parameters as shown below.

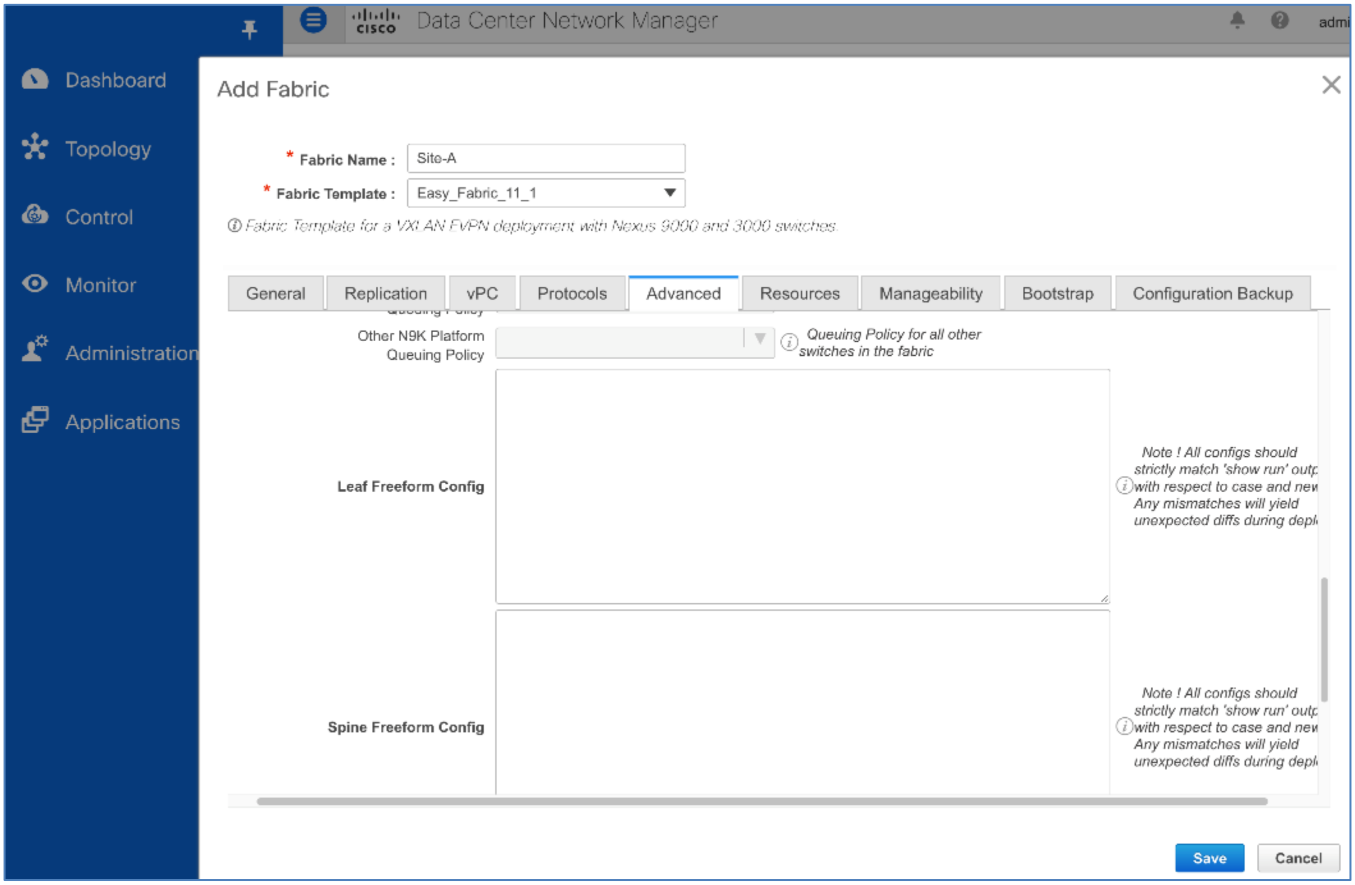

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

Other N9K Platform Queuing Policy — ⓘ Queuing Policy for all other switches in the fabric

Leaf Freeform Config — ⓘ Note ! All configs should strictly match 'show run' outp with respect to case and new Any mismatches will yield unexpected diffs during depl

Spine Freeform Config — ⓘ Note ! All configs should strictly match 'show run' outp with respect to case and new Any mismatches will yield unexpected diffs during depl

13. (Optional) In the **Resources** tab, you can specify the underlay loopbacks and subnets for various protocols. Cisco DCNM provides default values for these that can be used as-is. However, in this CVD, the parameters specified in the **Setup Information** section is used. Skip this step if using the default values. Otherwise, configure the **Underlay Routing Loopback IP Range, Underlay VTEP Loopback IP Range, Underlay RP Loopback IP Range** and **Underlay Subnet IP Range** for Site-A using the setup information.

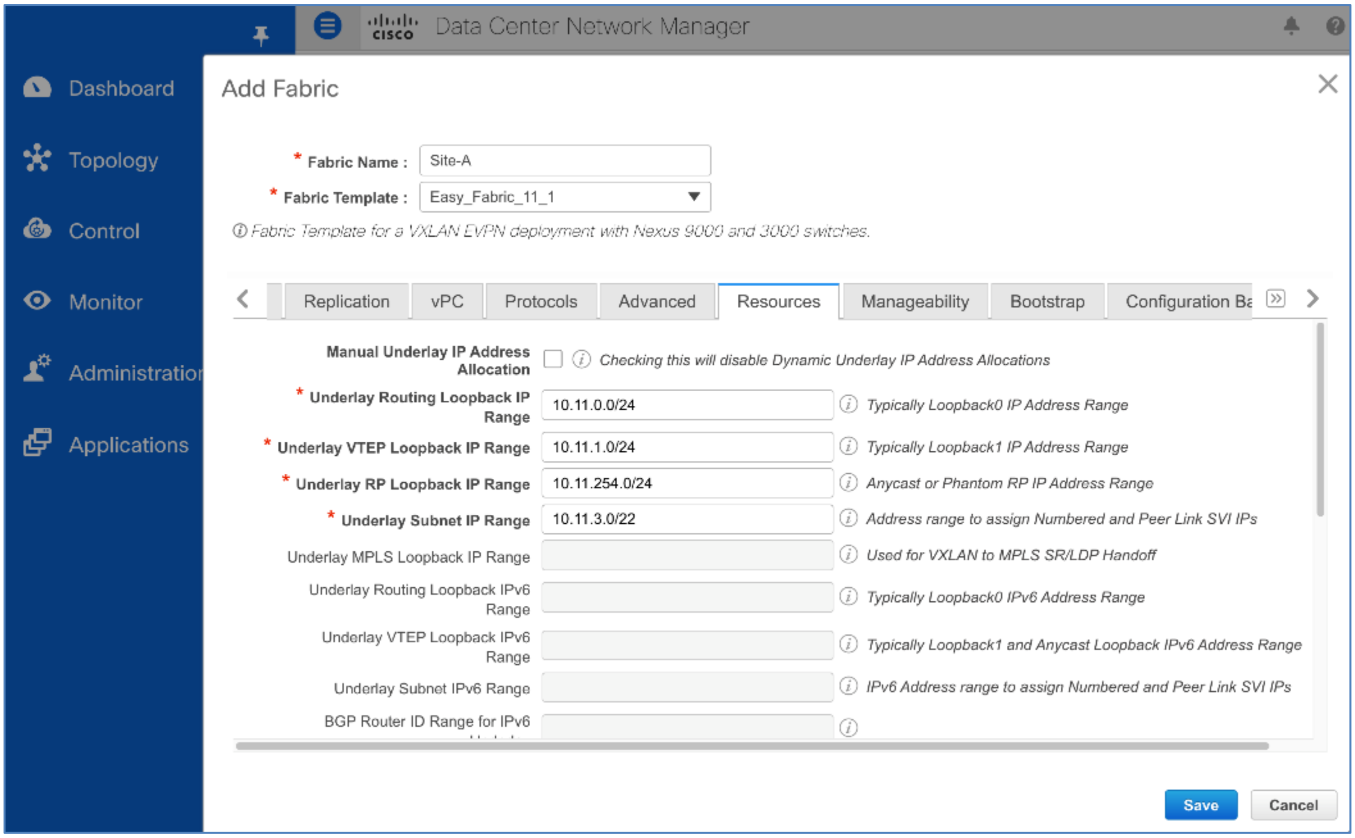

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

Manual Underlay IP Address Allocation □ Checking this will disable Dynamic Underlay IP Address Allocations

* Underlay Routing Loopback IP Range: 10.11.0.0/24 Typically Loopback0 IP Address Range

* Underlay VTEP Loopback IP Range: 10.11.1.0/24 Typically Loopback1 IP Address Range

* Underlay RP Loopback IP Range: 10.11.254.0/24 Anycast or Phantom RP IP Address Range

* Underlay Subnet IP Range: 10.11.3.0/22 Address range to assign Numbered and Peer Link SVI IPs

Underlay MPLS Loopback IP Range: Used for VXLAN to MPLS SR/LDP Handoff

Underlay Routing Loopback IPv6 Range: Typically Loopback0 IPv6 Address Range

Underlay VTEP Loopback IPv6 Range: Typically Loopback1 and Anycast Loopback IPv6 Address Range

Underlay Subnet IPv6 Range: IPv6 Address range to assign Numbered and Peer Link SVI IPs

BGP Router ID Range for IPv6

14. (Optional) In the **Resources** tab, you can specify the VXLAN Network IDs (VNID) and VLAN ranges for the access layer networks. Cisco DCNM provides default values for these that can be used as-is. However, in this CVD, the parameters specified in the **Setup Information** section is used. Skip this step if using the default values. Otherwise, configure the **Layer 2 VXLAN VNI Range, Layer 3 VXLAN VNI Range, Network VLAN Range, VRF VLAN Range** and **Service Network VLAN Range** for Site-A using the setup information.

The other parameters, namely VRF Lite Deployment, Auto Deploy Both, VRF Lite Subnet IP Range, and VRF Lite Subnet Mask can be specified now or can be updated in the External or Outside Connectivity section where they are used.

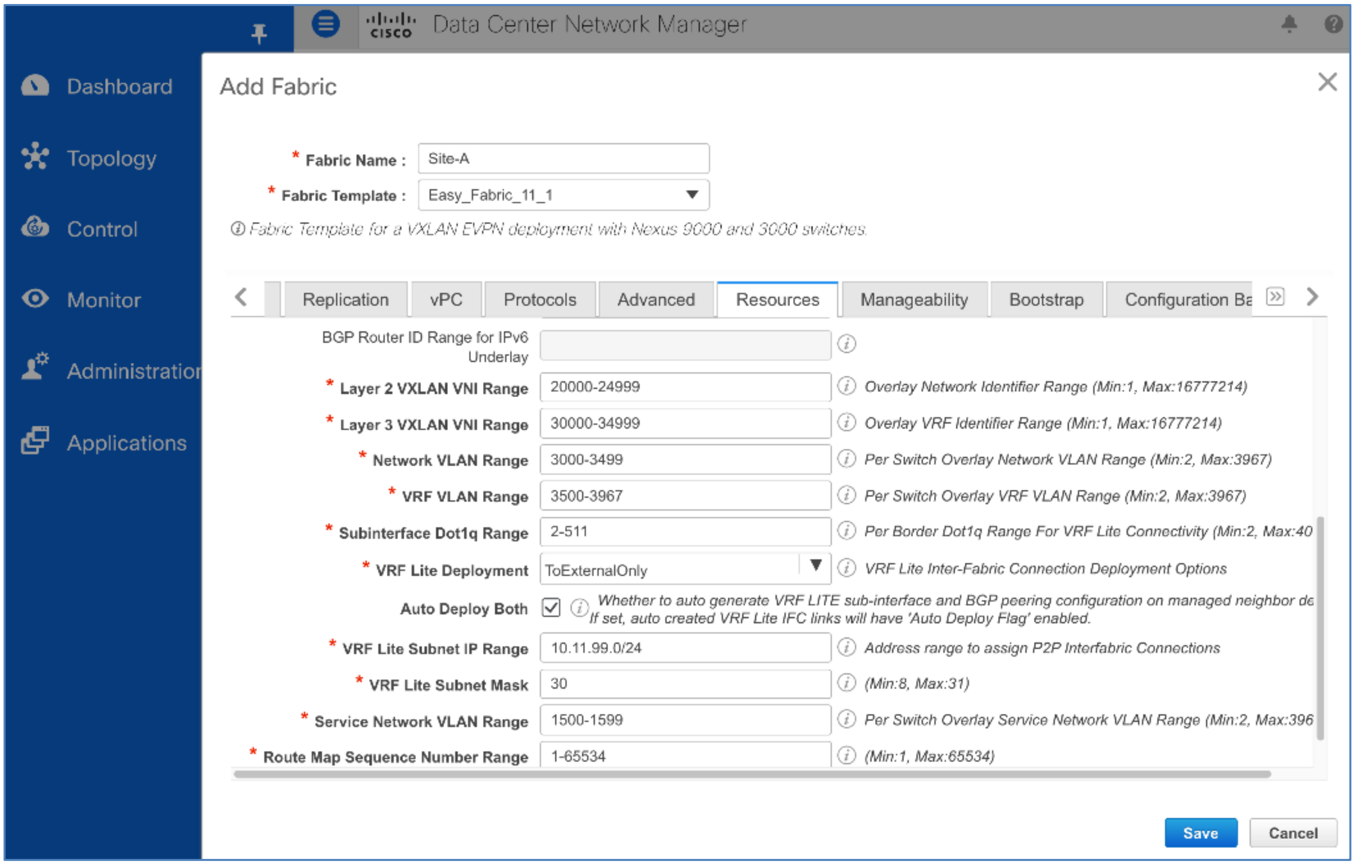

Add Fabric

* Fabric Name :	Site-A	
* Fabric Template :	Easy_Fabric_11_1 ▼	

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches

BGP Router ID Range for IPv6 Underlay		ⓘ
* Layer 2 VXLAN VNI Range	20000-24999	ⓘ Overlay Network Identifier Range (Min:1, Max:16777214)
* Layer 3 VXLAN VNI Range	30000-34999	ⓘ Overlay VRF Identifier Range (Min:1, Max:16777214)
* Network VLAN Range	3000-3499	ⓘ Per Switch Overlay Network VLAN Range (Min:2, Max:3967)
* VRF VLAN Range	3500-3967	ⓘ Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)
* Subinterface Dot1q Range	2-511	ⓘ Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:40
* VRF Lite Deployment	ToExternalOnly ▼	ⓘ VRF Lite Inter-Fabric Connection Deployment Options
Auto Deploy Both	☑	ⓘ Whether to auto generate VRF LITE sub-interface and BGP peering configuration on managed neighbor de If set, auto created VRF Lite IFC links will have 'Auto Deploy Flag' enabled.
* VRF Lite Subnet IP Range	10.11.99.0/24	ⓘ Address range to assign P2P Interfabric Connections
* VRF Lite Subnet Mask	30	ⓘ (Min:8, Max:31)
* Service Network VLAN Range	1500-1599	ⓘ Per Switch Overlay Service Network VLAN Range (Min:2, Max:396
* Route Map Sequence Number Range	1-65534	ⓘ (Min:1, Max:65534)

15. (Optional) In the **Manageability** tab, specify the NTP server and VRF for accessing the NTP servers. Other network infrastructure services such as DNS and Syslog servers can also be specified here.

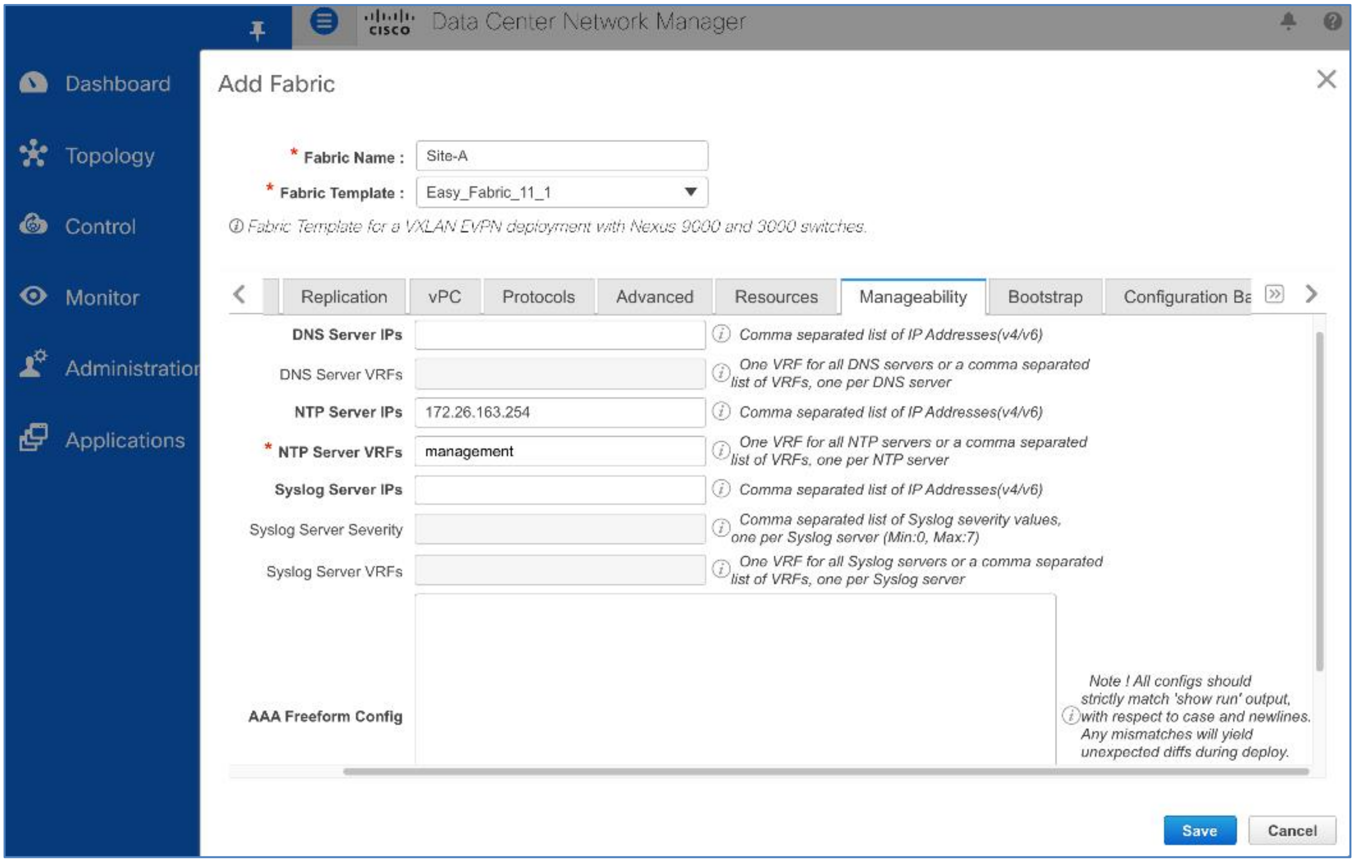

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

DNS Server IPs: ⓘ Comma separated list of IP Addresses(v4/v6)

DNS Server VRFs: ⓘ One VRF for all DNS servers or a comma separated list of VRFs, one per DNS server

NTP Server IPs: 172.26.163.254 ⓘ Comma separated list of IP Addresses(v4/v6)

* NTP Server VRFs: management ⓘ One VRF for all NTP servers or a comma separated list of VRFs, one per NTP server

Syslog Server IPs: ⓘ Comma separated list of IP Addresses(v4/v6)

Syslog Server Severity: ⓘ Comma separated list of Syslog severity values, one per Syslog server (Min:0, Max:7)

Syslog Server VRFs: ⓘ One VRF for all Syslog servers or a comma separated list of VRFs, one per Syslog server

AAA Freeform Config: ⓘ Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.

16. (Optional) In the **Bootstrap** tab, customers can specify bootstrap information if POAP is used to discover and import the switches into the fabric. This was not used in this CVD – proceed to the next tab.

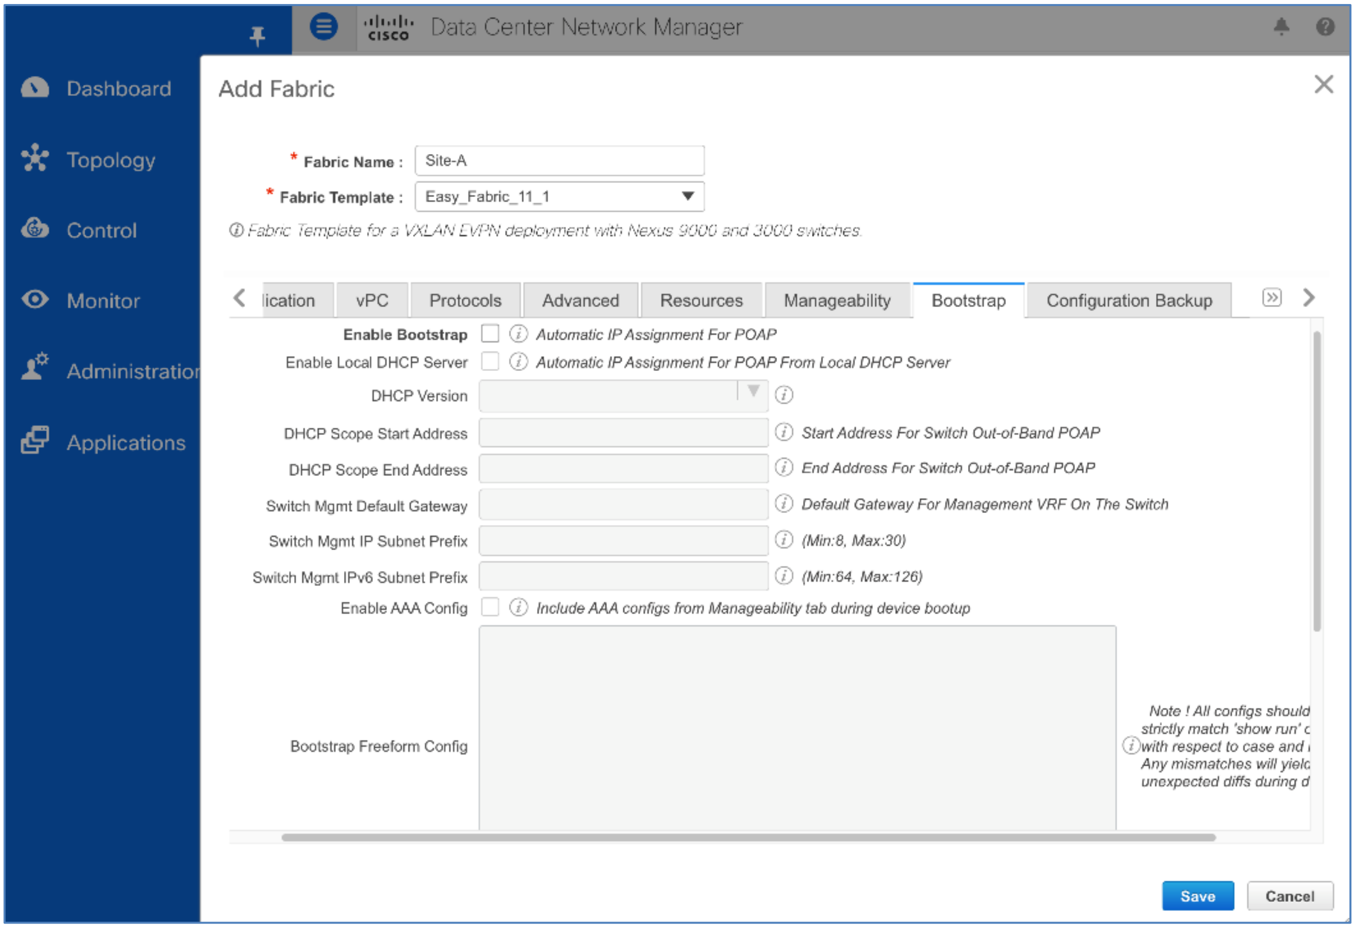

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

Enable Bootstrap □ ⓘ Automatic IP Assignment For POAP

Enable Local DHCP Server □ ⓘ Automatic IP Assignment For POAP From Local DHCP Server

DHCP Version ⓘ

DHCP Scope Start Address ⓘ Start Address For Switch Out-of-Band POAP

DHCP Scope End Address ⓘ End Address For Switch Out-of-Band POAP

Switch Mgmt Default Gateway ⓘ Default Gateway For Management VRF On The Switch

Switch Mgmt IP Subnet Prefix ⓘ (Min:8, Max:30)

Switch Mgmt IPv6 Subnet Prefix ⓘ (Min:64, Max:126)

Enable AAA Config □ ⓘ Include AAA configs from Manageability tab during device bootup

Bootstrap Freeform Config ⓘ Note ! All configs should strictly match 'show run' c with respect to case and i Any mismatches will yielc unexpected diffs during d

17. (Optional) In the **Configuration Backup** tab, specify a backup schedule for the fabric as shown below.

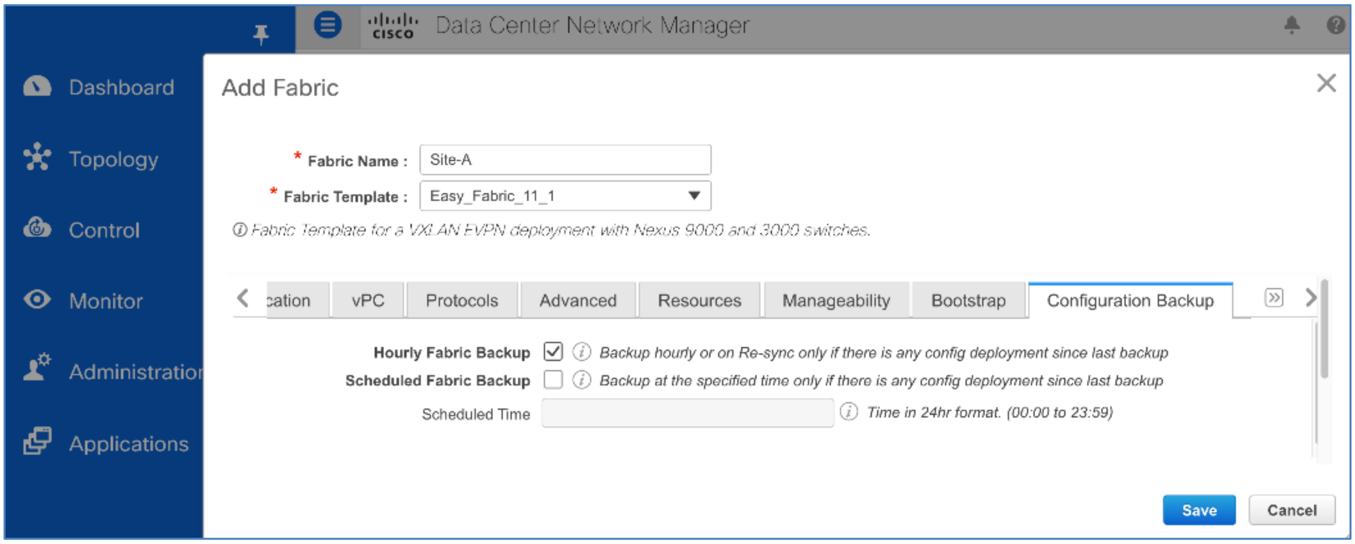

Add Fabric

* Fabric Name : Site-A

* Fabric Template : Easy_Fabric_11_1

ⓘ Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

Hourly Fabric Backup ☑ ⓘ Backup hourly or on Re-sync only if there is any config deployment since last backup

Scheduled Fabric Backup □ ⓘ Backup at the specified time only if there is any config deployment since last backup

Scheduled Time ⓘ Time in 24hr format. (00:00 to 23:59)

18. Click **Save** to save the fabric settings for the VXLAN Fabric in Site-A. You will get a pop-up on the right-bottom corner saying the Fabric was deployed successfully if the settings were saved. The saved settings

are merely the configuration intent at this stage – they will need to be deployed on the switches for it to take effect.

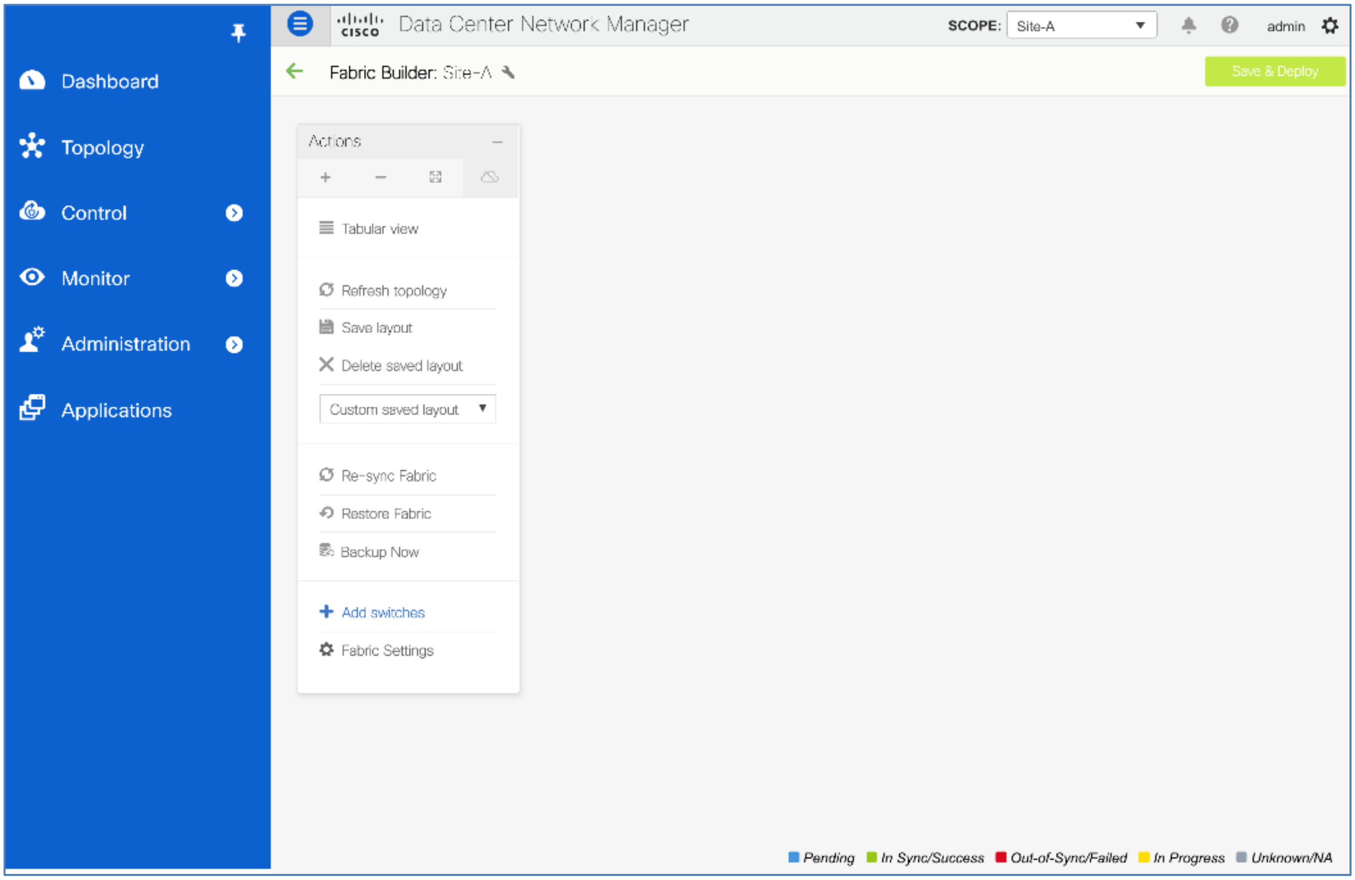

19. At this point, you can start adding switches to the VXLAN fabric. Note that you can use ⚙ **Fabric Settings** in the **Actions** menu at any time to modify the parameters – however, once switches have been added to the fabric, you will need to do a **Save & Deploy** (top-right corner) in order to save the settings and then to apply them to the switches in the fabric.

20. Proceed to the next section to discover and add switches to the Site-A datacenter fabric.

Always verify scope (for example, **SCOPE: Site-A**) in the top-right corner of the window when making changes or viewing the status to ensure that you are in the corrects datacenter or view.

Add Spine and Leaf switches to the VXLAN Fabric

As stated earlier, this design assumes a greenfield deployment where the fabric is built from the ground-up. Therefore, this section walks through the discovery, addition, and initial configuration of all switches to the Site-A fabric. For existing fabrics, customers can use relevant portions of this section to add switches to their fabric.

To add spine and leaf switches to the VXLAN fabric, follow these steps:

1. In the right-window pane, verify that the **SCOPE:** is **Site-A** in the drop-down list near the top-right corner. From the **Actions** menu, select and click **Add Switches**.

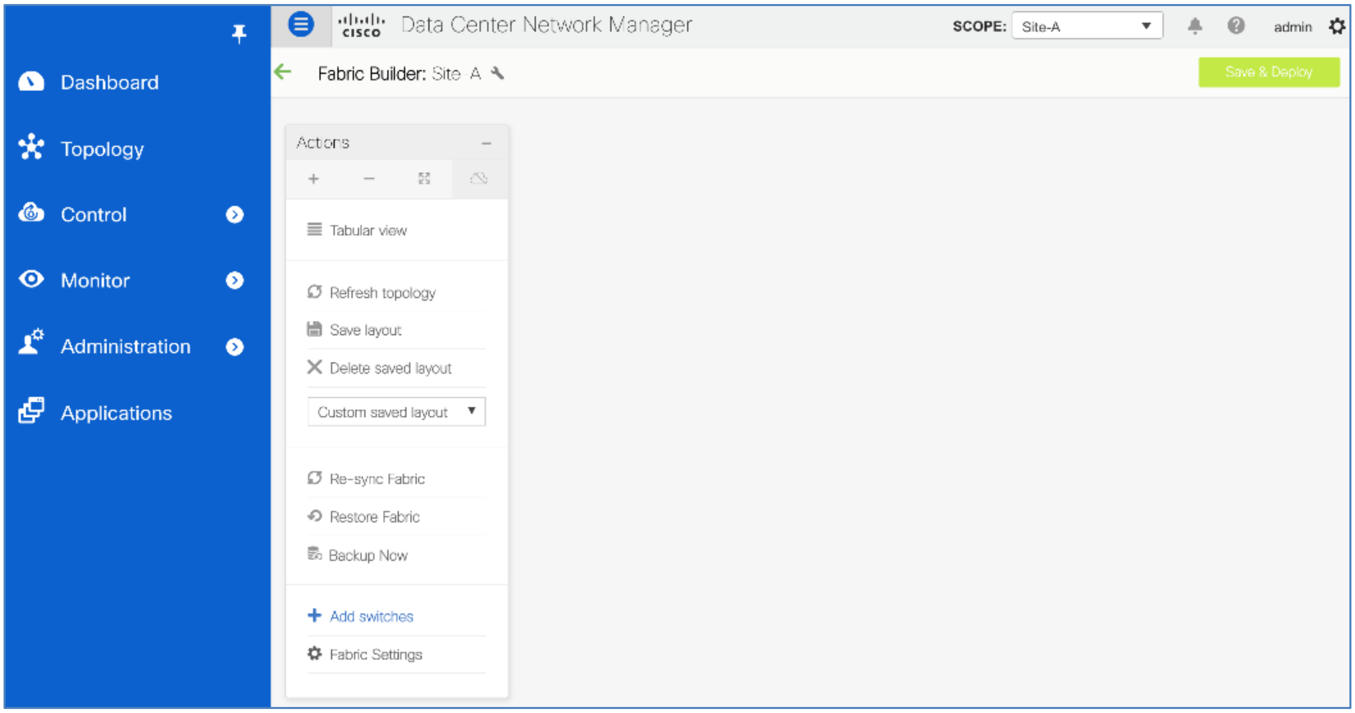

2. In the **Inventory Management** pop-up window, select the **Discover Existing Switches** tab. Note that you can also POAP on Cisco DCNM to discover and add switches to the VXLAN fabric. For the **Seed IP,** specify the IP address range of all switches that need to be discovered. For the **Username** and **Password,** specify the administrator username and password for the switches that you can use to log-on to the switches. For the **Max Hops**, specify '0' otherwise the discovery may take a long time to complete. For **Preserve Config**, select 'no' to clean up the configuration on the switches before adding them to the fabric.

The configuration on the switches are cleaned up before adding the switches as this CVD assumes a greenfield deployment. Cleaning up ensures there are no conflicts between the configuration deployed by Cisco DCNM and what is actually configured on the switch.

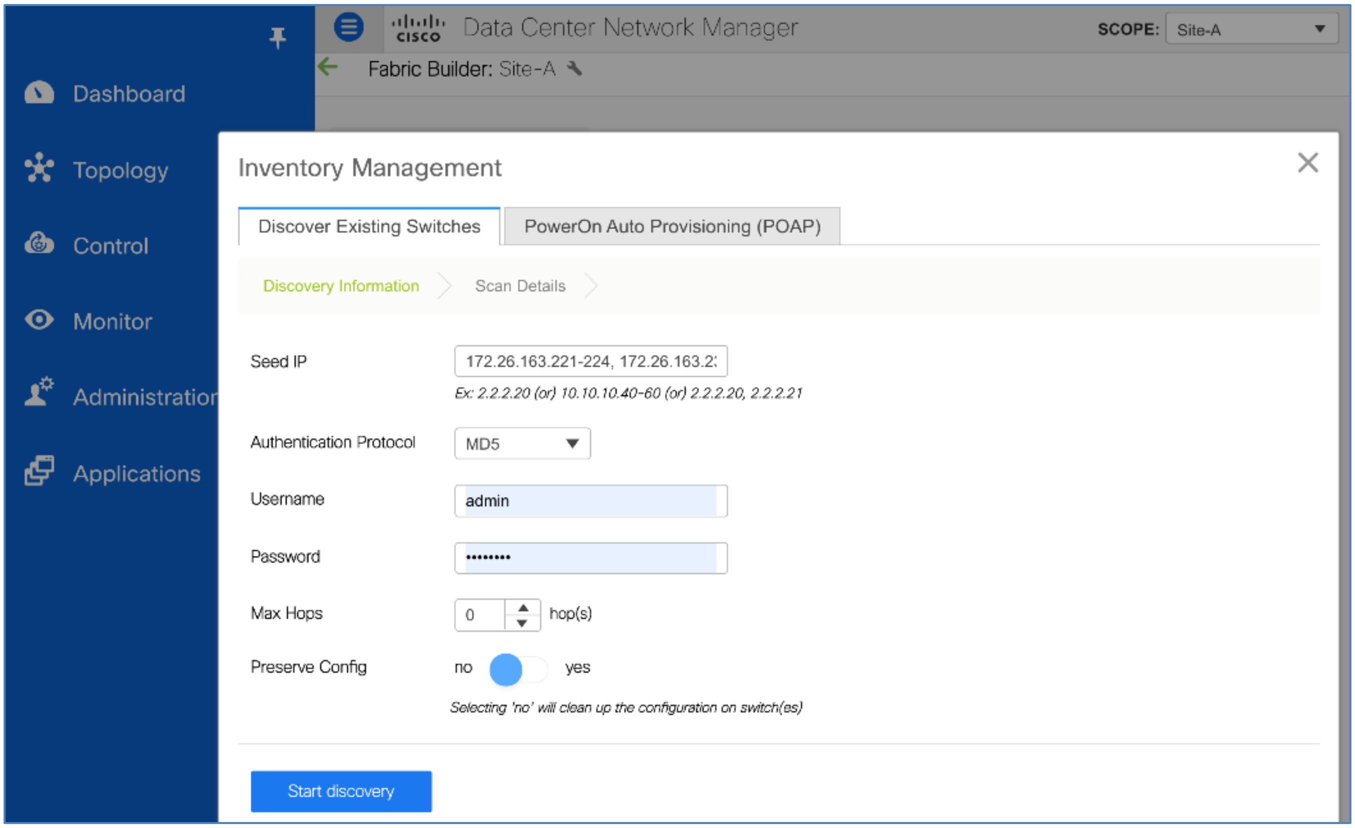

Inventory Management

Seed IP: 172.26.163.221-224, 172.26.163.2:

Ex: 2.2.2.20 (or) 10.10.10.40-60 (or) 2.2.2.20, 2.2.2.21

Authentication Protocol: MD5

Username: admin

Password: ••••••••

Max Hops: 0 hop(s)

Preserve Config: no ● yes

Selecting 'no' will clean up the configuration on switch(es)

3. Click the **Start discovery** button at the bottom to start the discovery process. You should see a spinning wheel in the red **Abort Request** button at the bottom as Cisco DCNM attempts to discover and find these switches.

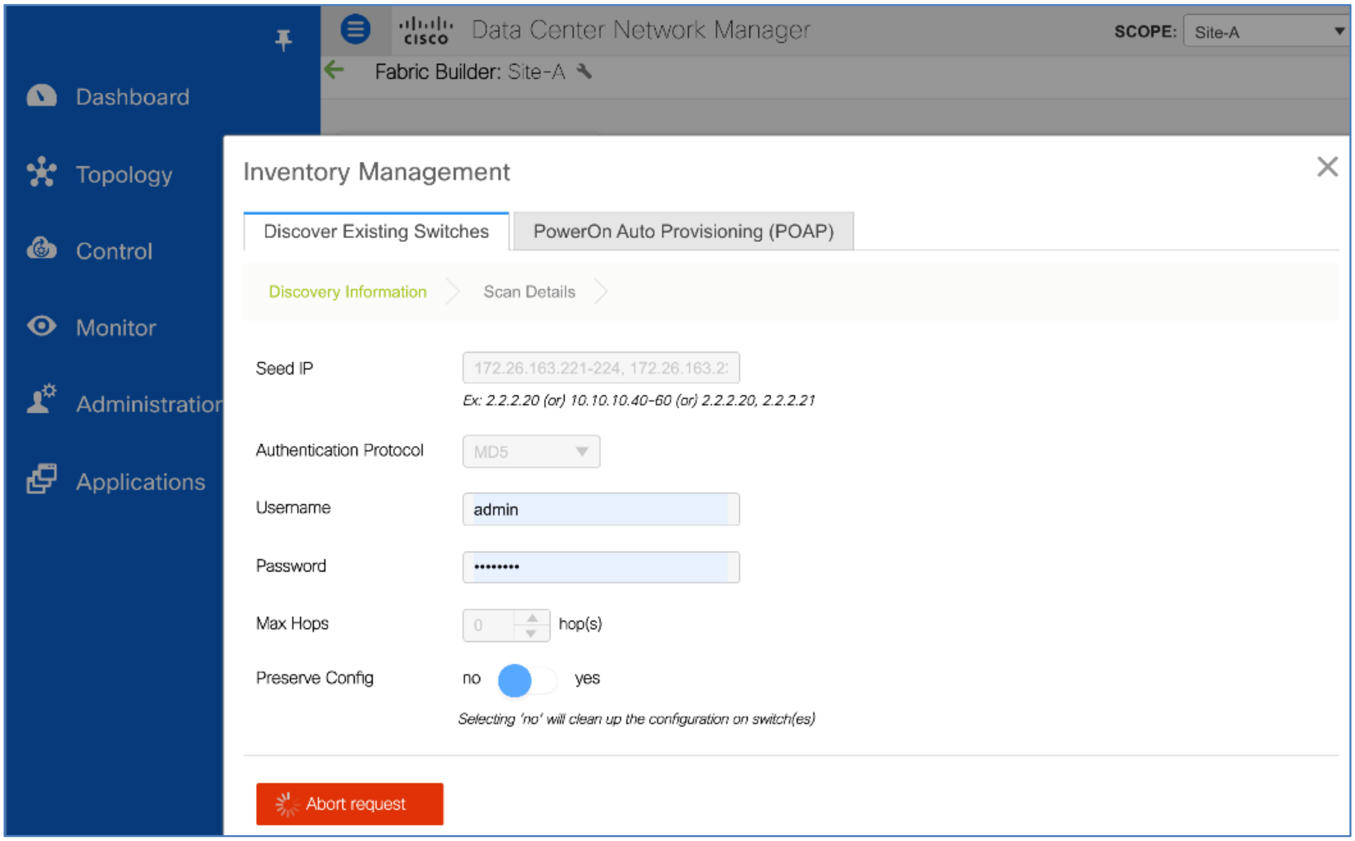

Seed IP: 172.26.163.221-224, 172.26.163.2

Ex: 2.2.2.20 (or) 10.10.10.40-60 (or) 2.2.2.20, 2.2.2.21

Authentication Protocol: MD5

Username: admin

Password: ••••••••

Max Hops: 0 hop(s)

Preserve Config: no ● yes

Selecting 'no' will clean up the configuration on switch(es)

4. Once the discovery completes, you will be provided with a list of switches that can be imported into the fabric. Use the checkboxes to the left of the switches to select the relevant switches. In this case, all switches in the list are selected using the checkbox to the left of the **Name** column. Click **Import into fabric**.

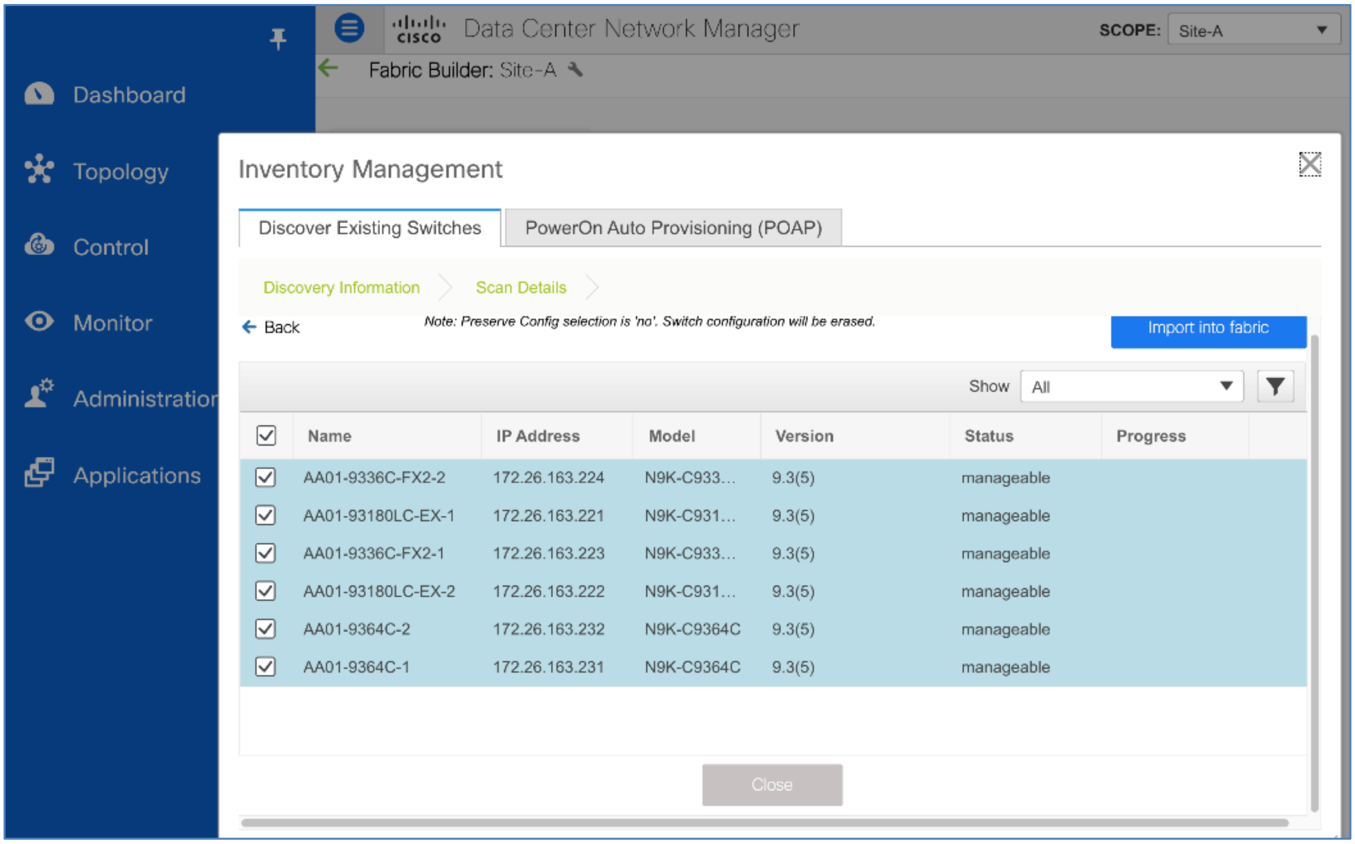

Note: Preserve Config selection is 'no'. Switch configuration will be erased.

☑	Name	IP Address	Model	Version	Status	Progress
☑	AA01-9336C-FX2-2	172.26.163.224	N9K-C933...	9.3(5)	manageable	
☑	AA01-93180LC-EX-1	172.26.163.221	N9K-C931...	9.3(5)	manageable	
☑	AA01-9336C-FX2-1	172.26.163.223	N9K-C933...	9.3(5)	manageable	
☑	AA01-93180LC-EX-2	172.26.163.222	N9K-C931...	9.3(5)	manageable	
☑	AA01-9364C-2	172.26.163.232	N9K-C9364C	9.3(5)	manageable	
☑	AA01-9364C-1	172.26.163.231	N9K-C9364C	9.3(5)	manageable	

5. Click **OK** in the Warning message that pops-up to confirm removal and cleanup of all configuration on the switches except for management connectivity.

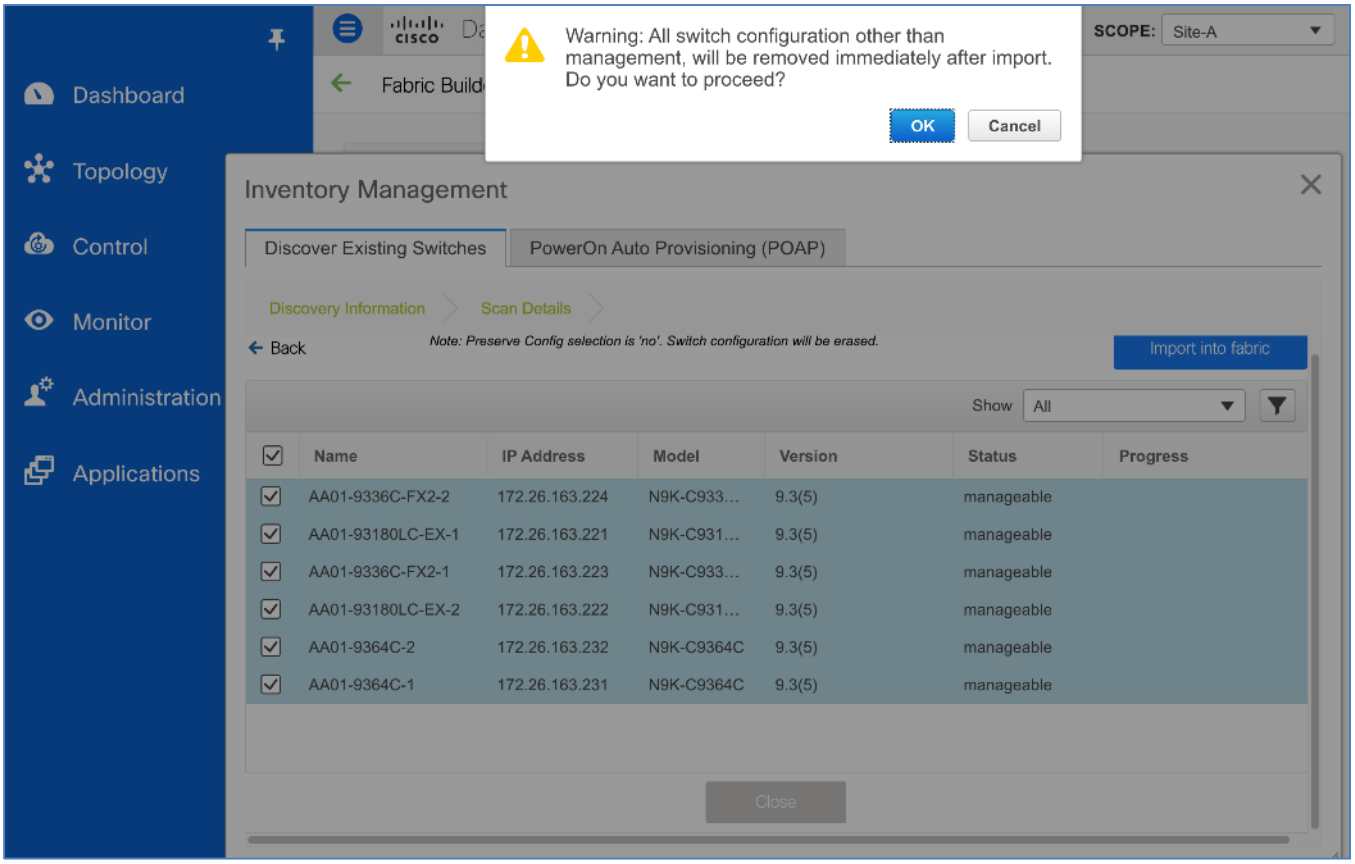

Warning: All switch configuration other than management, will be removed immediately after import. Do you want to proceed?

Inventory Management

Note: Preserve Config selection is 'no'. Switch configuration will be erased.

☑	Name	IP Address	Model	Version	Status	Progress
☑	AA01-9336C-FX2-2	172.26.163.224	N9K-C933…	9.3(5)	manageable	
☑	AA01-93180LC-EX-1	172.26.163.221	N9K-C931…	9.3(5)	manageable	
☑	AA01-9336C-FX2-1	172.26.163.223	N9K-C933…	9.3(5)	manageable	
☑	AA01-93180LC-EX-2	172.26.163.222	N9K-C931…	9.3(5)	manageable	
☑	AA01-9364C-2	172.26.163.232	N9K-C9364C	9.3(5)	manageable	
☑	AA01-9364C-1	172.26.163.231	N9K-C9364C	9.3(5)	manageable	

6. You can view the progress of the import in the **Progress** column for each switch being imported.

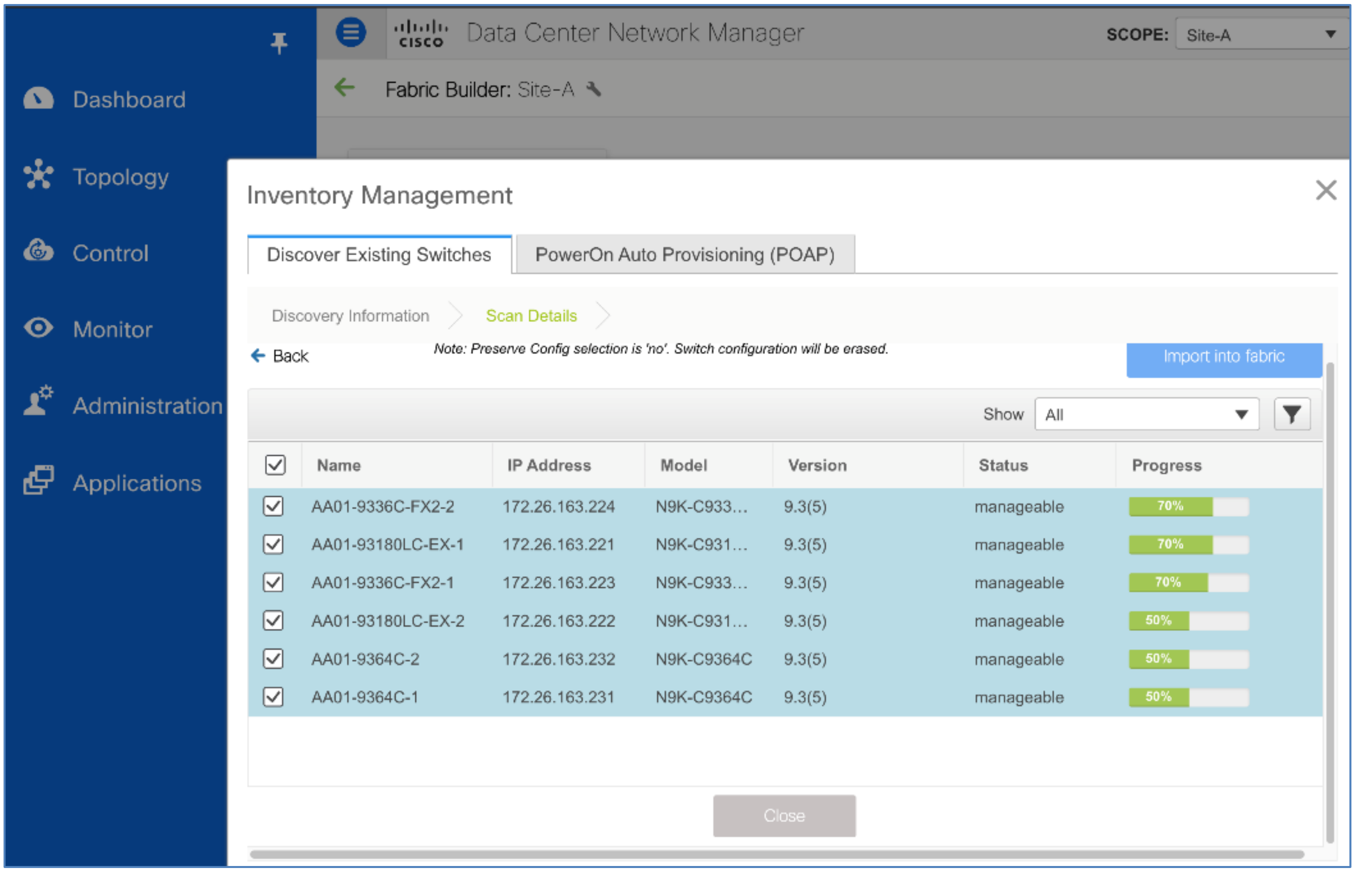

Note: Preserve Config selection is 'no'. Switch configuration will be erased.

☑	Name	IP Address	Model	Version	Status	Progress
☑	AA01-9336C-FX2-2	172.26.163.224	N9K-C933…	9.3(5)	manageable	70%
☑	AA01-93180LC-EX-1	172.26.163.221	N9K-C931…	9.3(5)	manageable	70%
☑	AA01-9336C-FX2-1	172.26.163.223	N9K-C933…	9.3(5)	manageable	70%
☑	AA01-93180LC-EX-2	172.26.163.222	N9K-C931…	9.3(5)	manageable	50%
☑	AA01-9364C-2	172.26.163.232	N9K-C9364C	9.3(5)	manageable	50%
☑	AA01-9364C-1	172.26.163.231	N9K-C9364C	9.3(5)	manageable	50%

7. Once imported is complete, the **Progress** column will show **done**. Click **Close**

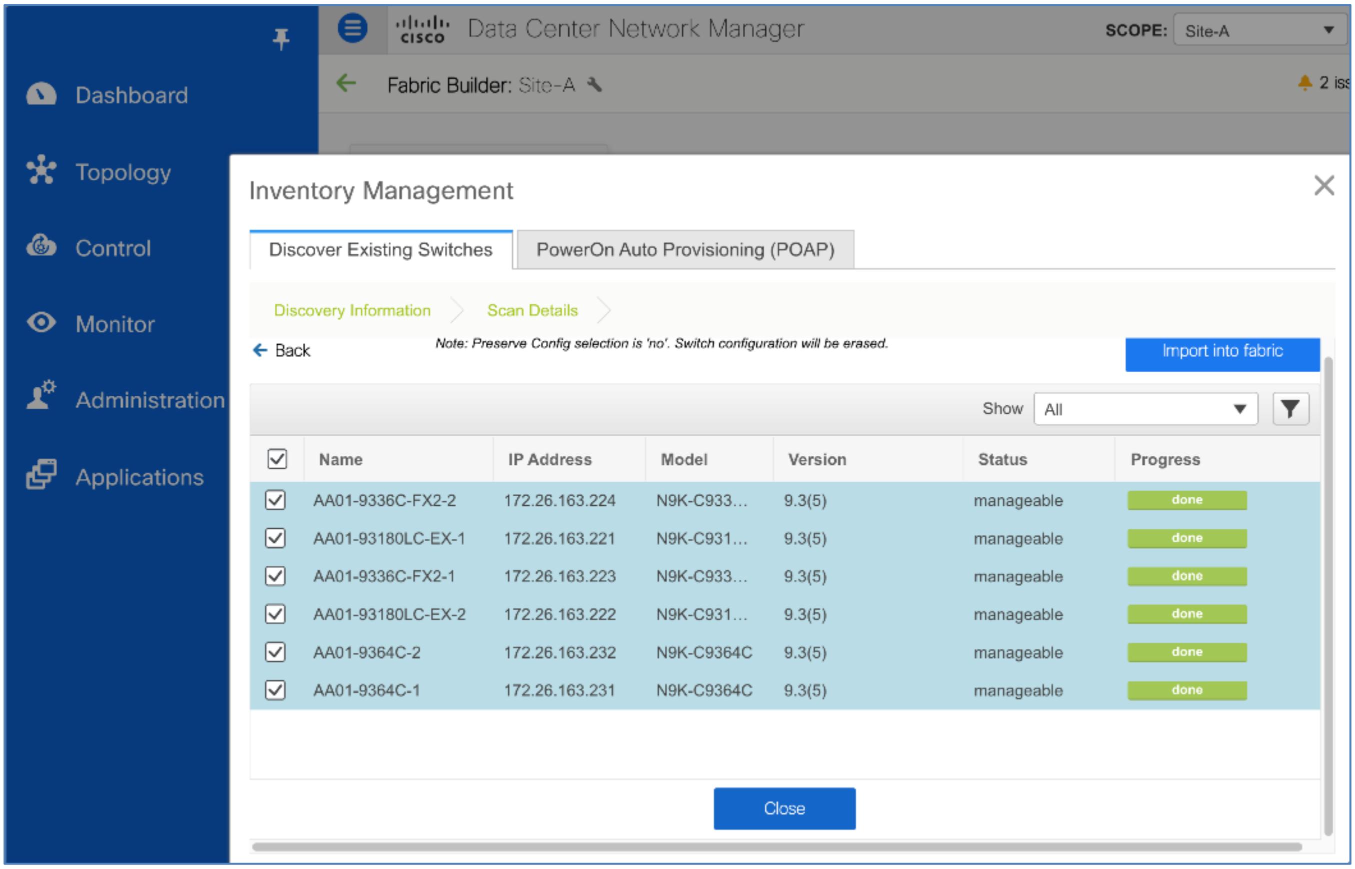

SCOPE: Site-A

Inventory Management

Note: Preserve Config selection is 'no'. Switch configuration will be erased.

Show All

☑	Name	IP Address	Model	Version	Status	Progress
☑	AA01-9336C-FX2-2	172.26.163.224	N9K-C933…	9.3(5)	manageable	done
☑	AA01-93180LC-EX-1	172.26.163.221	N9K-C931…	9.3(5)	manageable	done
☑	AA01-9336C-FX2-1	172.26.163.223	N9K-C933…	9.3(5)	manageable	done
☑	AA01-93180LC-EX-2	172.26.163.222	N9K-C931…	9.3(5)	manageable	done
☑	AA01-9364C-2	172.26.163.232	N9K-C9364C	9.3(5)	manageable	done
☑	AA01-9364C-1	172.26.163.231	N9K-C9364C	9.3(5)	manageable	done

8. You should now see the topology as shown below. Move the Actions menu/window to right-side for a better view of the topology.

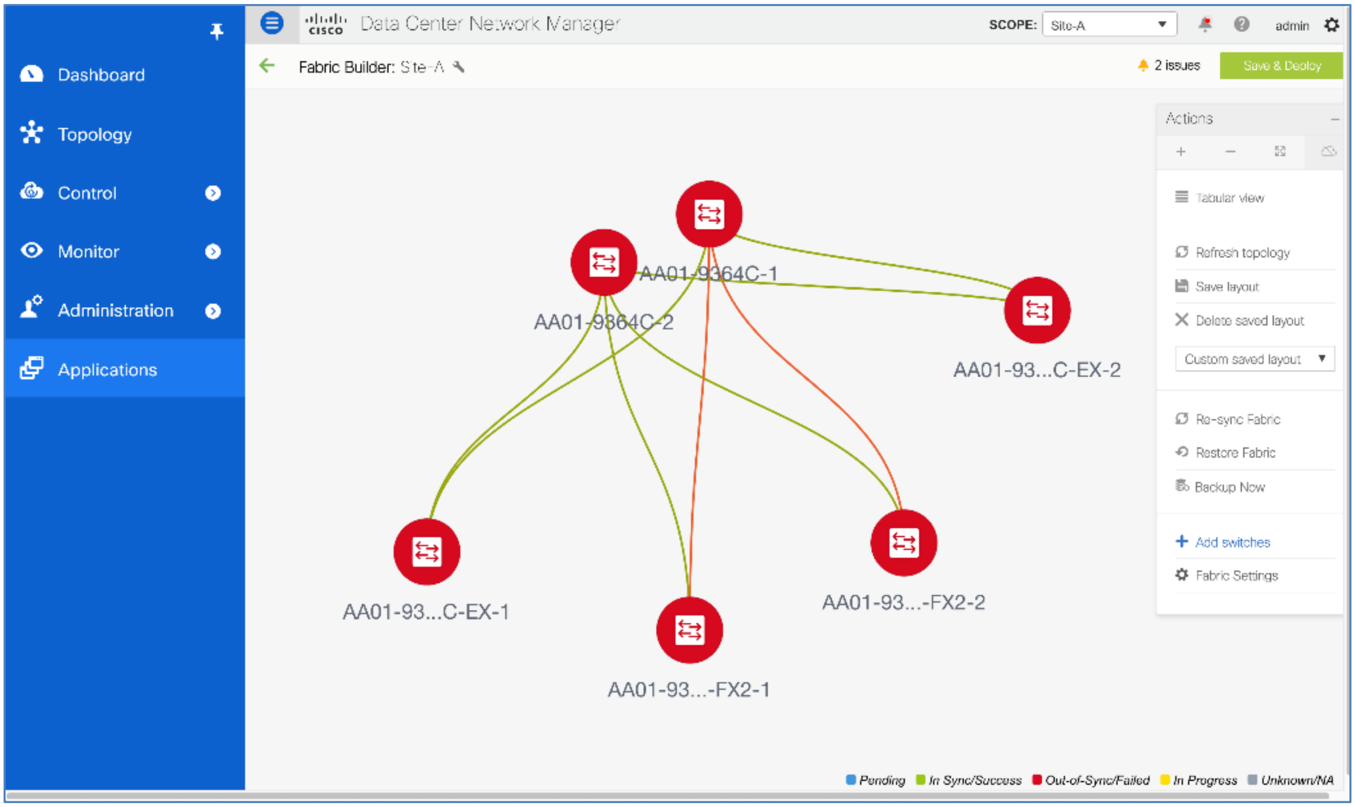

9. At this stage, all switches in the topology should be red to indicate they are in **Out-of-Sync/Failed** state. This is to be expected as the configuration on the switches (which have been wiped clean) do not match the Fabric Settings or the fabric configuration on Cisco DCNM. Also note that there may warnings or errors that show up, if any exist, it will show up to the left of the **Save & Deploy** button. Review the issues so they can be resolved. In this case, there are **2 issues Warnings**. The Fabric errors are warnings and indicate that a re-load of two of the switches should be done after a **Save & Deploy**. We will therefore wait on resolving these until after a **Save & Deploy** is done.

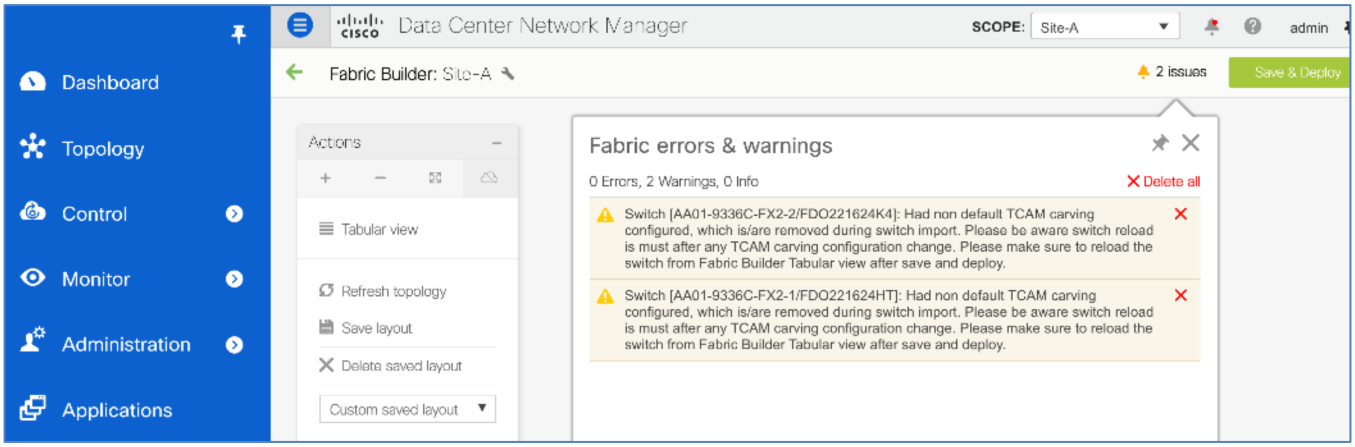

Fabric errors & warnings

0 Errors, 2 Warnings, 0 Info

Switch [AA01-9336C-FX2-2/FDO221624K4]: Had non default TCAM carving configured, which is/are removed during switch import. Please be aware switch reload is must after any TCAM carving configuration change. Please make sure to reload the switch from Fabric Builder Tabular view after save and deploy.

Switch [AA01-9336C-FX2-1/FDO221624HT]: Had non default TCAM carving configured, which is/are removed during switch import. Please be aware switch reload is must after any TCAM carving configuration change. Please make sure to reload the switch from Fabric Builder Tabular view after save and deploy.

10. From the **Actions** menu/window, select **Tabular view.** Verify that the **Discovery Status** is **ok.** For each switch, verify that **Role** is correct. Modify the role as needed. The role of 4 switches selected below needs to

be changed. To change the role, go back to the topology view by clicking on the green left arrow in the top left corner of the right window pane.

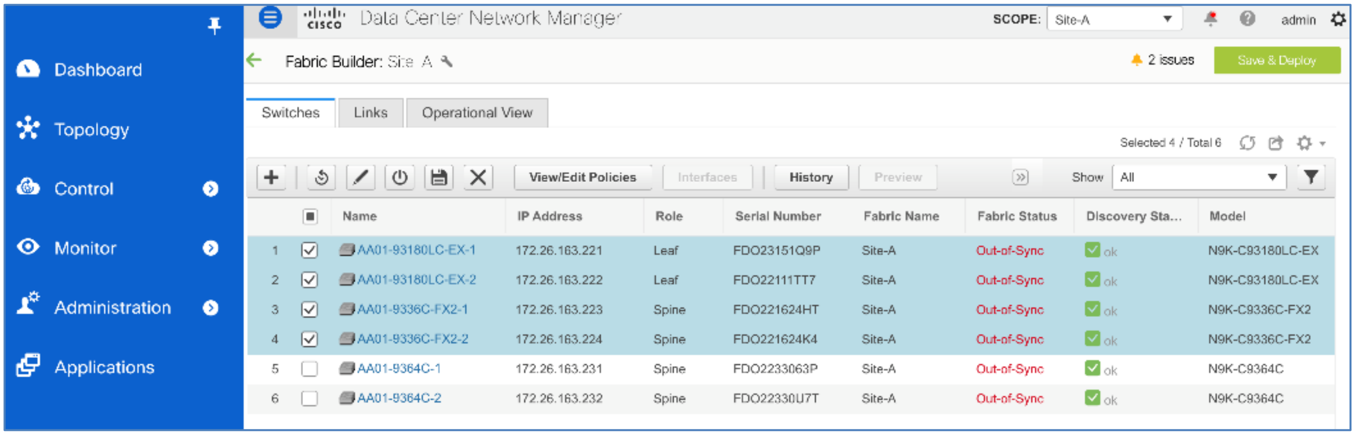

	☐	Name	IP Address	Role	Serial Number	Fabric Name	Fabric Status	Discovery Sta...	Model
1	☑	AA01-93180LC-EX-1	172.26.163.221	Leaf	FDO23151Q9P	Site-A	Out-of-Sync	✓ ok	N9K-C93180LC-EX
2	☑	AA01-93180LC-EX-2	172.26.163.222	Leaf	FDO22111TT7	Site-A	Out-of-Sync	✓ ok	N9K-C93180LC-EX
3	☑	AA01-9336C-FX2-1	172.26.163.223	Spine	FDO221624HT	Site-A	Out-of-Sync	✓ ok	N9K-C9336C-FX2
4	☑	AA01-9336C-FX2-2	172.26.163.224	Spine	FDO221624K4	Site-A	Out-of-Sync	✓ ok	N9K-C9336C-FX2
5	☐	AA01-9364C-1	172.26.163.231	Spine	FDO2233063P	Site-A	Out-of-Sync	✓ ok	N9K-C9364C
6	☐	AA01-9364C-2	172.26.163.232	Spine	FDO22330U7T	Site-A	Out-of-Sync	✓ ok	N9K-C9364C

11. In the topology view, from the **Actions** menu/window, select **Hierarchical** from the drop-down list. To change the role, select the switch and right-click to select **Set role** and then select the correct role for the switch from the list. A small window on the bottom right will pop-up to indicate with role change was successful.

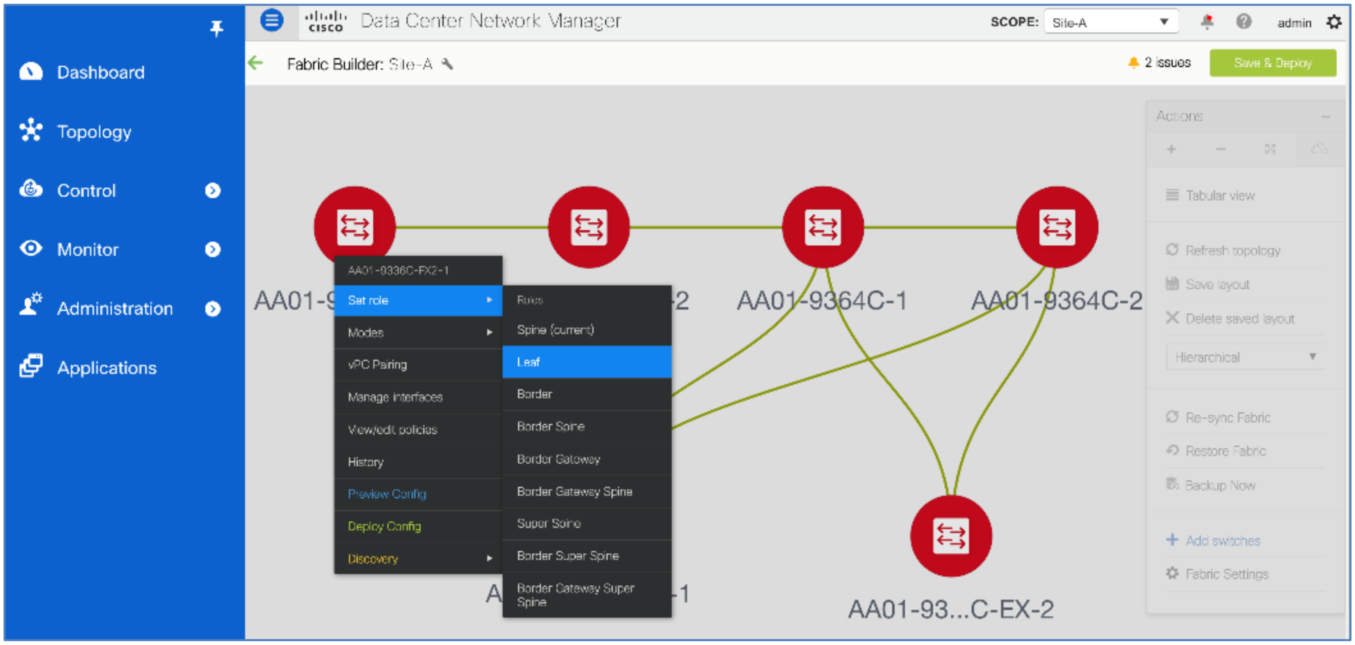

12. Repeat the previous step for all switches whose role needs to be changed.

13. Verify the roles have been changed, go back **Tabular view** from the **Actions** menu/window, and verify the **Role** column for each switch.

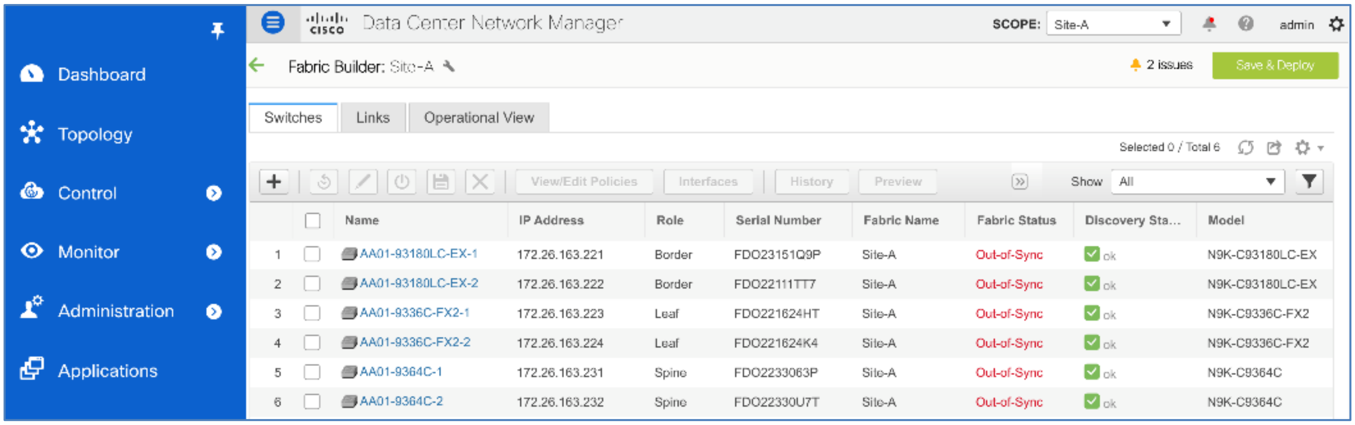

		Name	IP Address	Role	Serial Number	Fabric Name	Fabric Status	Discovery Sta...	Model
1	□	AA01-93180LC-EX-1	172.26.163.221	Border	FDO23151Q9P	Site-A	Out-of-Sync	✓ ok	N9K-C93180LC-EX
2	□	AA01-93180LC-EX-2	172.26.163.222	Border	FDO22111TT7	Site-A	Out-of-Sync	✓ ok	N9K-C93180LC-EX
3	□	AA01-9336C-FX2-1	172.26.163.223	Leaf	FDO221624HT	Site-A	Out-of-Sync	✓ ok	N9K-C9336C-FX2
4	□	AA01-9336C-FX2-2	172.26.163.224	Leaf	FDO221624K4	Site-A	Out-of-Sync	✓ ok	N9K-C9336C-FX2
5	□	AA01-9364C-1	172.26.163.231	Spine	FDO2233063P	Site-A	Out-of-Sync	✓ ok	N9K-C9364C
6	□	AA01-9364C-2	172.26.163.232	Spine	FDO22330U7T	Site-A	Out-of-Sync	✓ ok	N9K-C9364C

14. Go back to **Topology** view, from the **Actions** menu/window, select **Hierarchical** from the drop-down list. The topology view should now change based on the role of the devices. Select **Save layout** from the **Actions** menu.

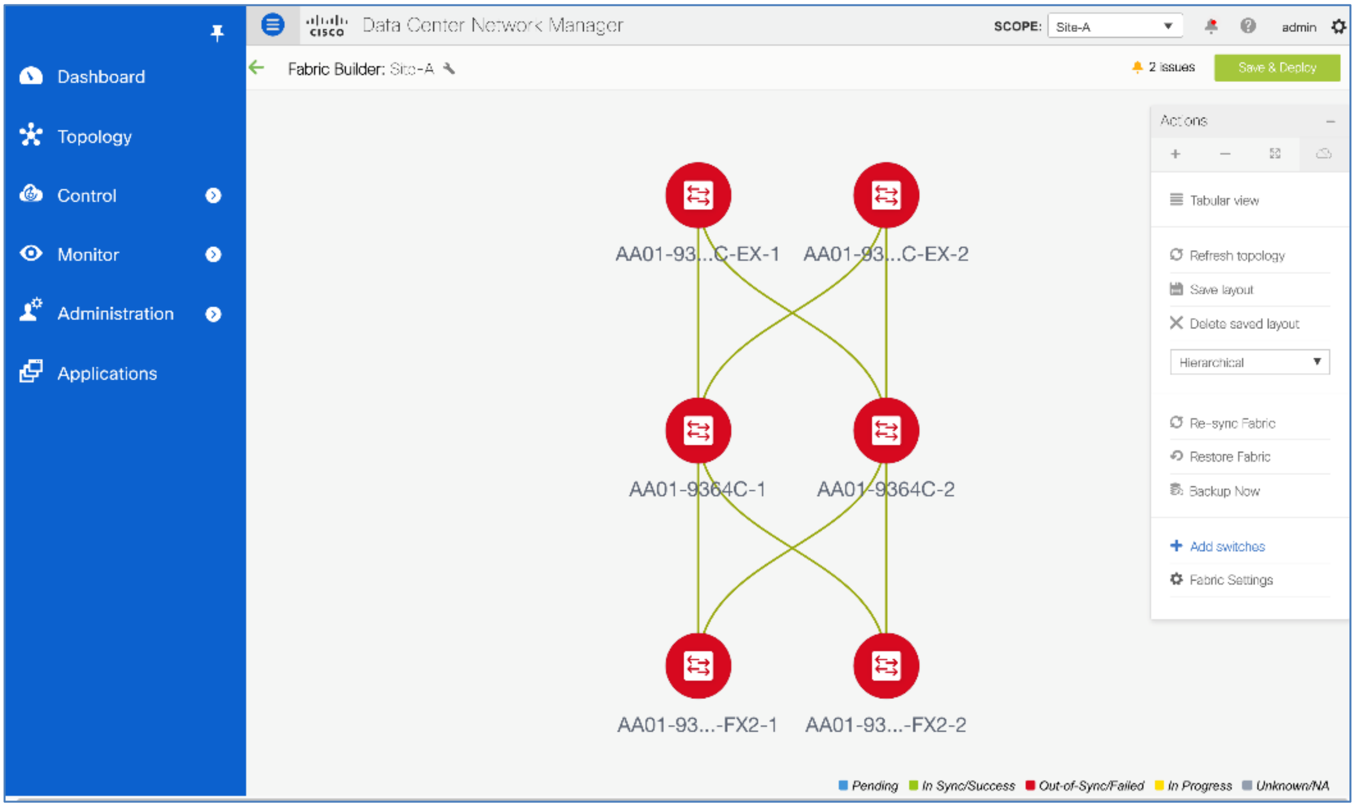

15. You are now ready to deploy the configuration to the switches in the fabric. Click **Save & Deploy**.

16. In the **Config Deployment** window, you can see the number of lines of configuration that will be deployed on each switch. The configuration deployed will vary depending on the role of the switch.

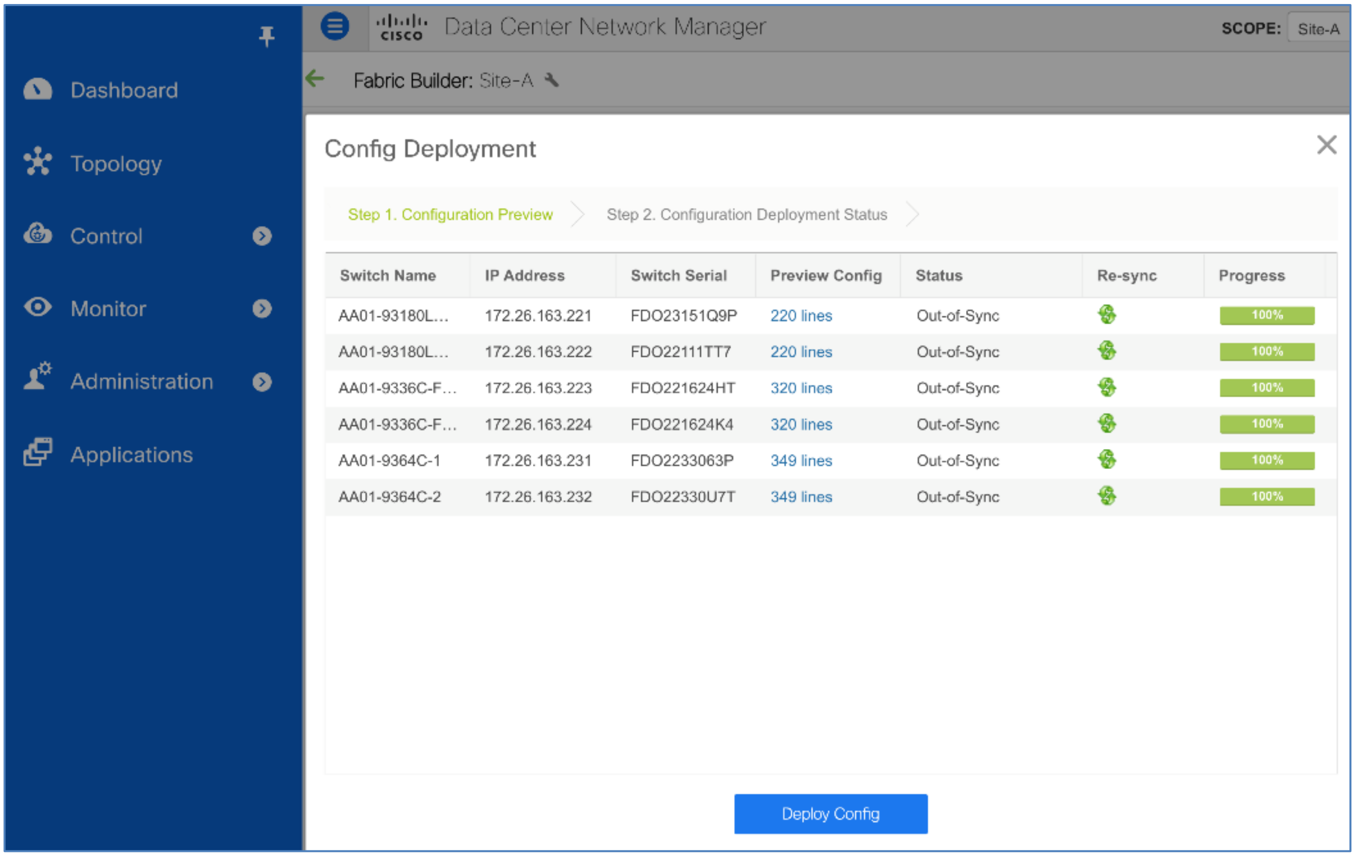

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
AA01-93180L…	172.26.163.221	FDO23151Q9P	220 lines	Out-of-Sync	✓	100%
AA01-93180L…	172.26.163.222	FDO22111TT7	220 lines	Out-of-Sync	✓	100%
AA01-9336C-F…	172.26.163.223	FDO221624HT	320 lines	Out-of-Sync	✓	100%
AA01-9336C-F…	172.26.163.224	FDO221624K4	320 lines	Out-of-Sync	✓	100%
AA01-9364C-1	172.26.163.231	FDO2233063P	349 lines	Out-of-Sync	✓	100%
AA01-9364C-2	172.26.163.232	FDO22330U7T	349 lines	Out-of-Sync	✓	100%

17. You can preview the configuration on each switch by clicking on the number of lines as shown below. In the **Preview Config – Switch (IP)** pop-up window, you can see the configuration that will be pushed to the switch in question. A partial view of the configuration is shown below.

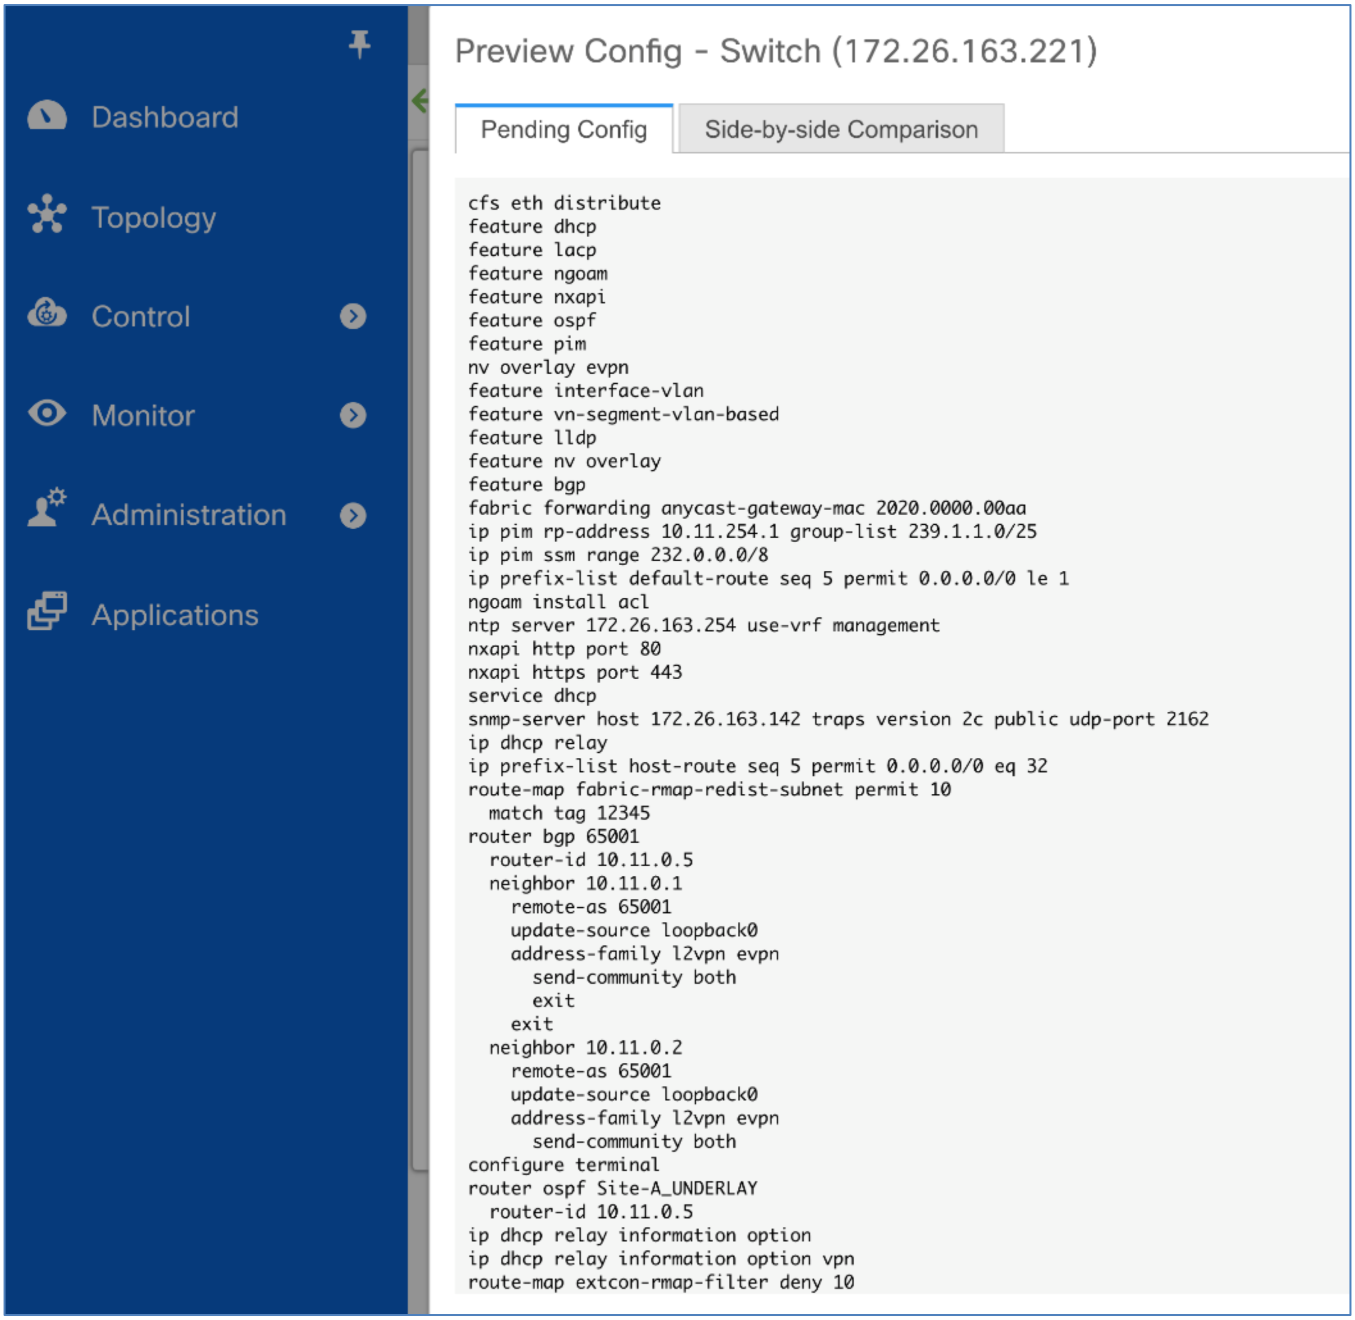

Preview Config - Switch (172.26.163.221)

```
cfs eth distribute
feature dhcp
feature lacp
feature ngoam
feature nxapi
feature ospf
feature pim
nv overlay evpn
feature interface-vlan
feature vn-segment-vlan-based
feature lldp
feature nv overlay
feature bgp
fabric forwarding anycast-gateway-mac 2020.0000.00aa
ip pim rp-address 10.11.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ngoam install acl
ntp server 172.26.163.254 use-vrf management
nxapi http port 80
nxapi https port 443
service dhcp
snmp-server host 172.26.163.142 traps version 2c public udp-port 2162
ip dhcp relay
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map fabric-rmap-redist-subnet permit 10
  match tag 12345
router bgp 65001
  router-id 10.11.0.5
  neighbor 10.11.0.1
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community both
      exit
    exit
  neighbor 10.11.0.2
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community both
configure terminal
router ospf Site-A_UNDERLAY
  router-id 10.11.0.5
ip dhcp relay information option
ip dhcp relay information option vpn
route-map extcon-rmap-filter deny 10
```

18. Exit the **Preview Config** pop-up window and click the **Deploy Config** button in the **Config Deployment** window.

19. In the **Config Deployment** window, you can view the deployment progress as shown below.

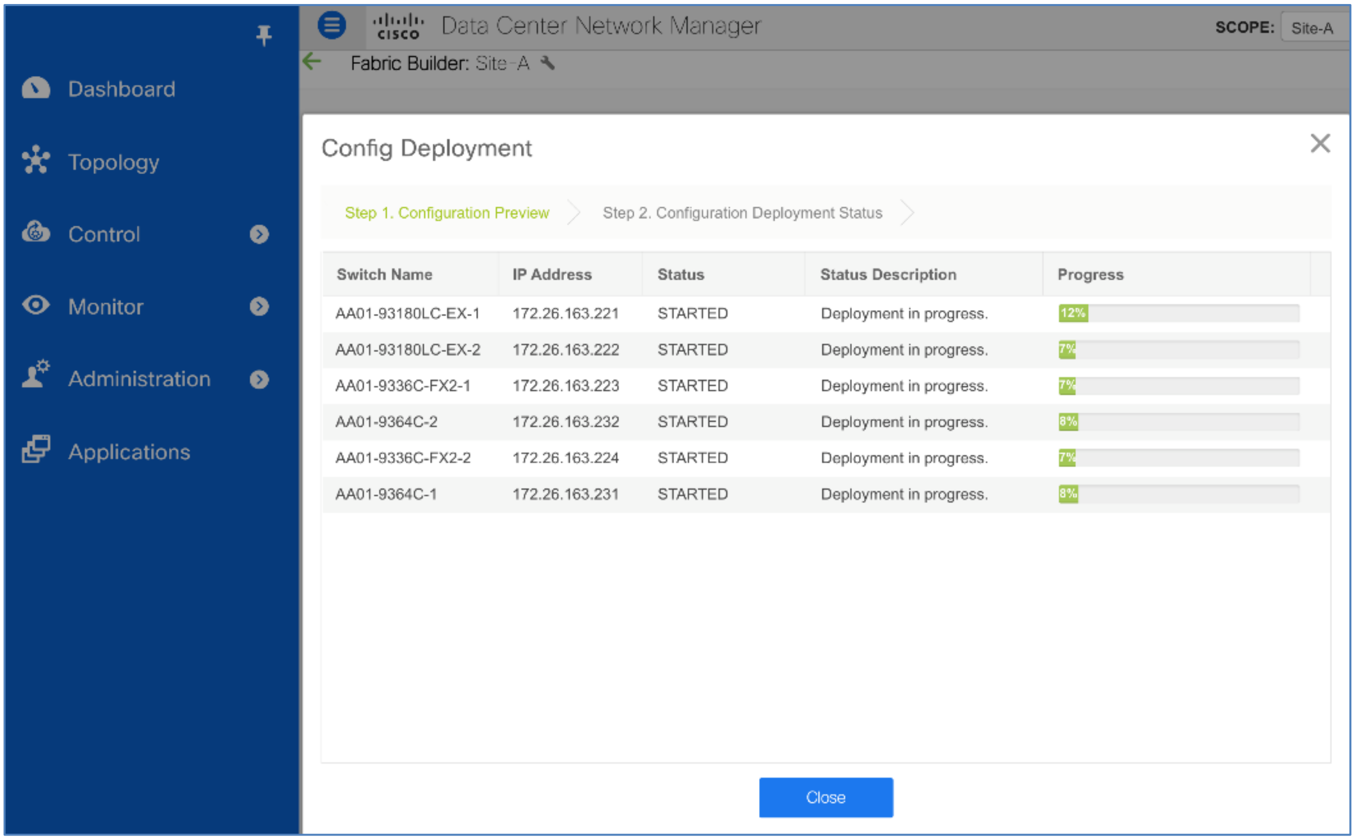

Config Deployment

Step 1. Configuration Preview

Step 2. Configuration Deployment Status

Switch Name	IP Address	Status	Status Description	Progress
AA01-93180LC-EX-1	172.26.163.221	STARTED	Deployment in progress.	12%
AA01-93180LC-EX-2	172.26.163.222	STARTED	Deployment in progress.	7%
AA01-9336C-FX2-1	172.26.163.223	STARTED	Deployment in progress.	7%
AA01-9364C-2	172.26.163.232	STARTED	Deployment in progress.	8%
AA01-9336C-FX2-2	172.26.163.224	STARTED	Deployment in progress.	7%
AA01-9364C-1	172.26.163.231	STARTED	Deployment in progress.	8%

20. The status will show **Deployed successfully** when complete. Click **Close.**

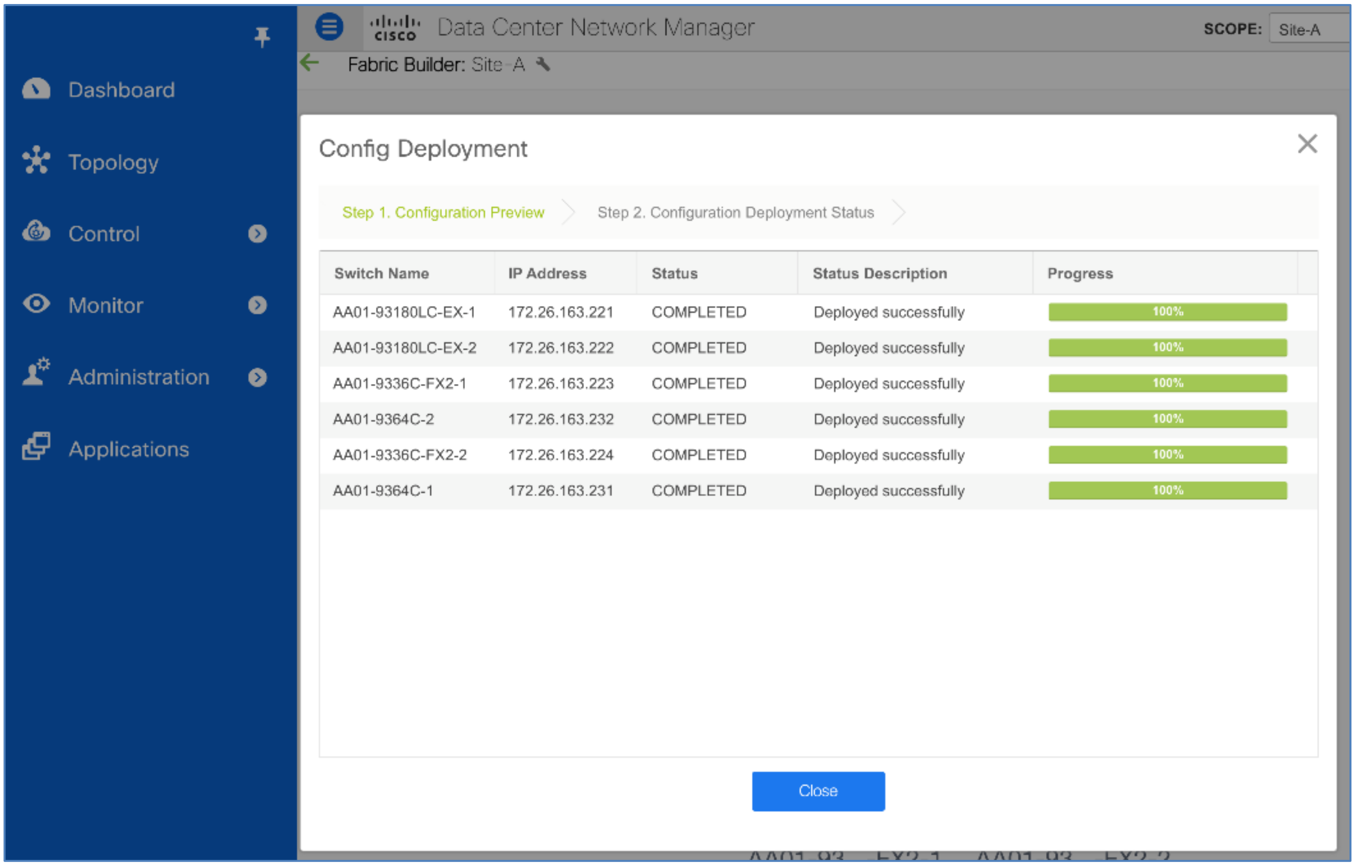

Switch Name	IP Address	Status	Status Description	Progress
AA01-93180LC-EX-1	172.26.163.221	COMPLETED	Deployed successfully	100%
AA01-93180LC-EX-2	172.26.163.222	COMPLETED	Deployed successfully	100%
AA01-9336C-FX2-1	172.26.163.223	COMPLETED	Deployed successfully	100%
AA01-9364C-2	172.26.163.232	COMPLETED	Deployed successfully	100%
AA01-9336C-FX2-2	172.26.163.224	COMPLETED	Deployed successfully	100%
AA01-9364C-1	172.26.163.231	COMPLETED	Deployed successfully	100%

21. The topology view should now show all switches in the green state indicating that the configurations between DCNM and the switches are in sync.

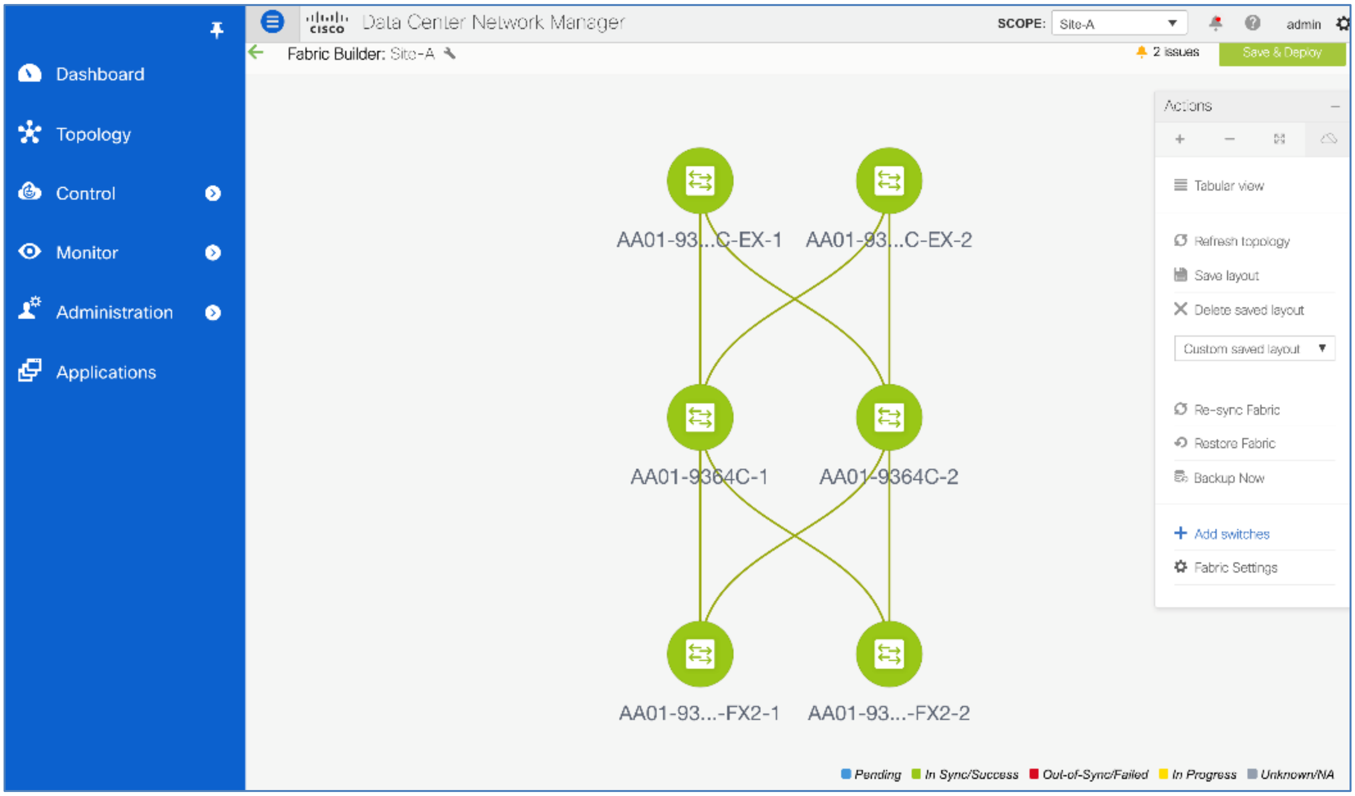

22. From the **Actions** menu/window, select **Tabular view** and for each switch, verify the **Role** and that the **Fabric Status** is **In-Sync.**

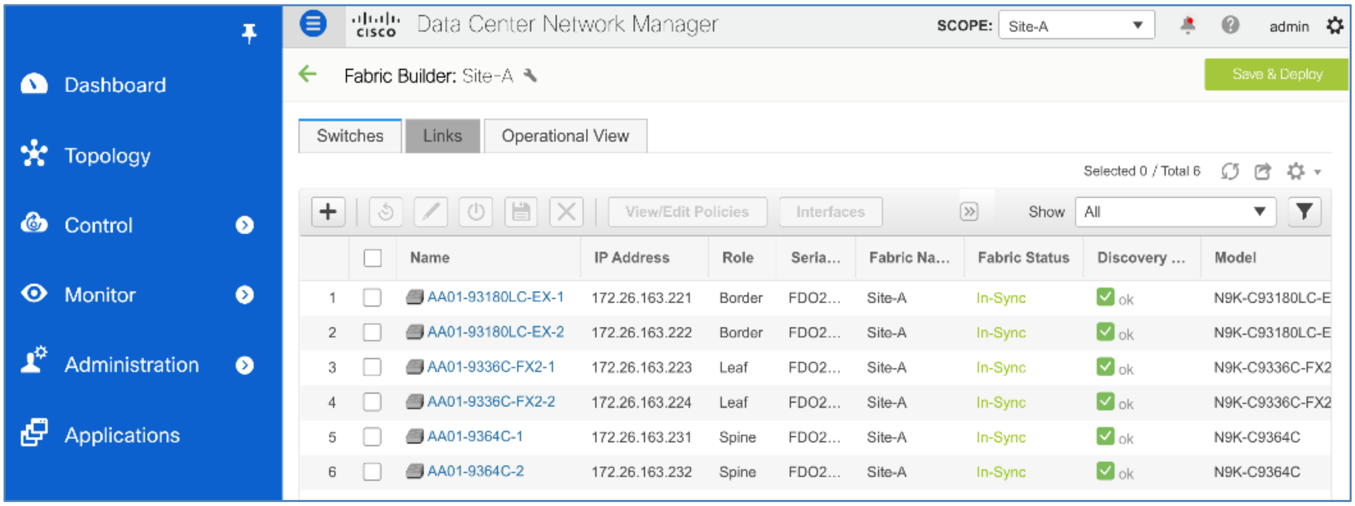

	□	Name	IP Address	Role	Seria...	Fabric Na...	Fabric Status	Discovery ...	Model
1	□	AA01-93180LC-EX-1	172.26.163.221	Border	FDO2...	Site-A	In-Sync	✓ ok	N9K-C93180LC-E
2	□	AA01-93180LC-EX-2	172.26.163.222	Border	FDO2...	Site-A	In-Sync	✓ ok	N9K-C93180LC-E
3	□	AA01-9336C-FX2-1	172.26.163.223	Leaf	FDO2...	Site-A	In-Sync	✓ ok	N9K-C9336C-FX2
4	□	AA01-9336C-FX2-2	172.26.163.224	Leaf	FDO2...	Site-A	In-Sync	✓ ok	N9K-C9336C-FX2
5	□	AA01-9364C-1	172.26.163.231	Spine	FDO2...	Site-A	In-Sync	✓ ok	N9K-C9364C
6	□	AA01-9364C-2	172.26.163.232	Spine	FDO2...	Site-A	In-Sync	✓ ok	N9K-C9364C

It is generally a good practice to use ⟳ **Refresh Topology** from the **Actions** menu/window in **Topology** view to verify the current state of the switches in the fabric especially when making change.

External or Outside Connectivity

In this design, connectivity outside or external to the VXLAN BGP EVPN fabric is necessary to access critical services in existing infrastructure that are outside the fabric and also to access cloud-based services. In this section, the external connectivity deployed will enable north-south traffic for access services such as Microsoft Active Directory and Domain Name System (DNS)

Topology

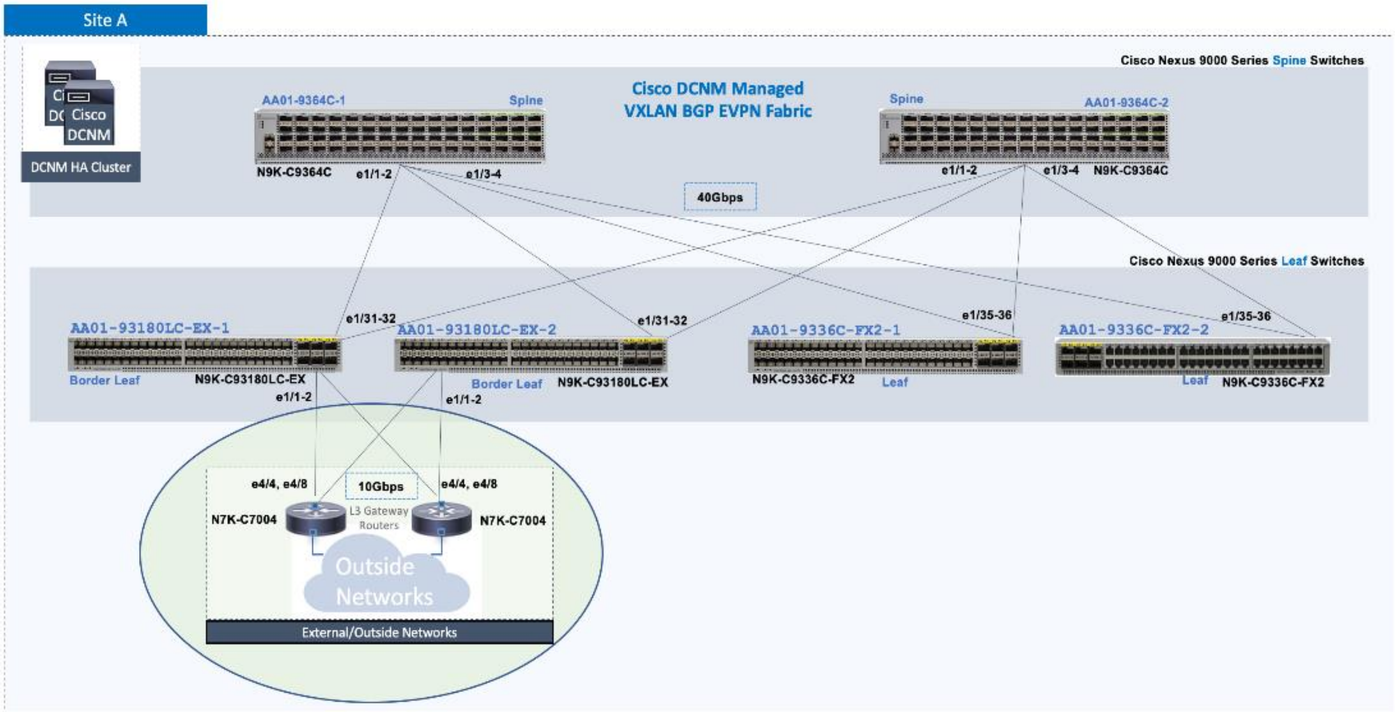

The Cisco Nexus 93180LC-EX border leaf switches for external access are positioned as a best practice being apart from the production FlexPod leaf switches in the infrastructure. If port availability and expansion needs are met with just the production leaf switches, this functionality could be combined.

Setup Information

The configuration parameters required for enabling external or outside connectivity from the Site-A datacenter fabric is provided in the table below.

Table 5. External or Outside Connectivity Parameters – Site-A

Data Center	Parameters	Default Parameters	Notes
Fabric Name	SiteA_External	_	
Fabric Template	External_Fabric_11_1	_	
General Tab			

Data Center	Parameters	Default Parameters	Notes
BGP AS#	65011	–	
Fabric Monitor Mode	☐	☑	Optional (If default is used, Cisco DCNM will not configure the external gateways)
Advanced Tab			
Enable AAA IP Authorization	☑	☐	Future DCNM releases will use this.
Resources Tab			
Sub-interface Dot1q Range	1101-1104	2-511	Optional (Default Values can be used as-is)
Underlay Routing Loopback IP Range	11.11.11.0/24	10.1.0.0/22	Optional (Default Values can be used as-is)
Configuration Backup Tab			
Hourly Fabric Backup	☑	☐	Optional

The setup information for discovering the external gateway switches in the outside/external fabric managed by Cisco DCNM is provided in the table below. Cisco DCNM also supports discovery and importing of external switches through Power-on-Auto-Provisioning(POAP) – however, POAP was not utilized in this CVD.

Table 6. Discovery Information – SiteA_External

Hostname	Switch Role	IP Address (OOB)	Notes
A07-7004-1-AA-East-Enterprise-1	External Gateway	172.26.163.115/24	
A07-7004-1-AA-East-Enterprise-2	External Gateway	172.26.163.116/24	

Switches used in this design were already configured with a Hostname, Management IP address, username, password, and boot variable.

The setup information for creating Inter-Fabric (IFC) links between the external fabric and Site-A datacenter fabric is provided in the table below.

Table 7. Inter-Fabric Link (IFC) Links between External Fabric and Site-A

Variable	Parameters	Notes
Link Sub-Type	VRF-Lite	
Link Template	ext_fabric_setup_11_1	
General		
Source IP Address/Mask	IFC#1: 10.11.99.5/30 IFC#2: 10.11.99.1/30 IFC#3: 10.11.99.13/30 IFC#4: 10.11.99.9/30	4 IFC Links (Auto-Deployed)
Destination IP	IFC#1: 10.11.99.6 IFC#2: 10.11.99.2 IFC#3: 10.11.99.14 IFC#4: 10.11.99.10	4 IFC Links (Auto-Deployed)
Auto Deploy Flag	☑	
Advanced		
Source Interface Description	IFC#1 → To AA-East-Enterprise-1: e4/4 IFC#2 → To AA-East-Enterprise-2: e4/4 IFC#3 → To AA-East-Enterprise-1: e4/8 IFC#4 → To AA-East-Enterprise-2: e4/8	4 IFC Links
Destination Interface Description	IFC#1 → To AA01-93180LC-EX-1: e1/1 IFC#2 → To AA01-93180LC-EX-1: e1/2 IFC#3 → To AA01-93180LC-EX-2: e1/1 IFC#4 → To AA01-93180LC-EX-2: e1/2	4 IFC Links

Create External Fabric in Cisco DCNM

In this design, a pair of Cisco Nexus 7000 series switches serve as gateways to networks outside or external to the VXLAN fabric. To achieve this connectivity, the Cisco Nexus 7000 series switches are imported into the fabric as Managed devices so that connectivity can be automatically provisioned for each tenant or VRF that is deployed within the fabric. Alternatively, Cisco Nexus 7000 series could be imported in ‘Monitored’ mode – in this case, the configuration on the Cisco Nexus 7000 series interfaces connecting to the VXLAN fabric would need to be done individually and manually by the network administrator.

To create the external fabric in Cisco DCNM, use the **Setup Information** provided above to follow these steps:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Fabric Builder**.

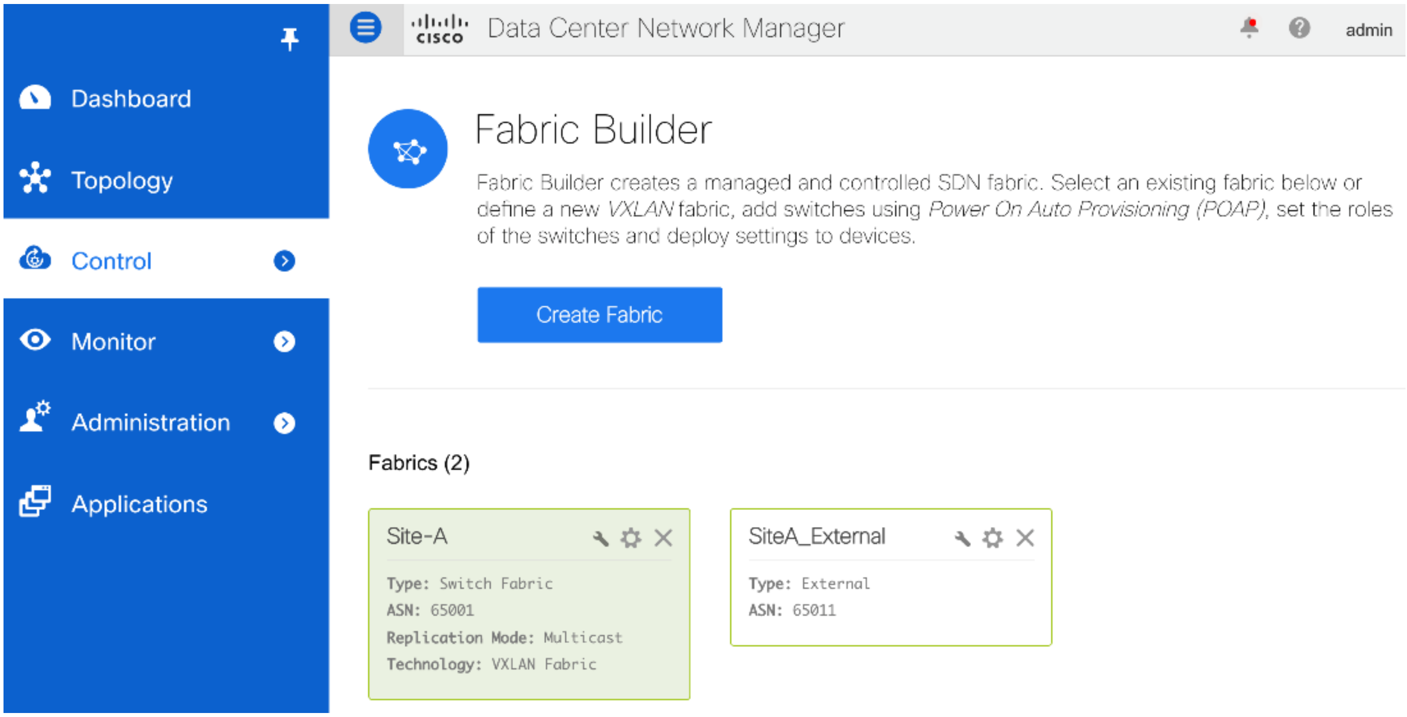

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (2)

Site-A

Type: Switch Fabric
ASN: 65001
Replication Mode: Multicast
Technology: VXLAN Fabric

SiteA_External

Type: External
ASN: 65011

3. From the right window pane, click the **Create Fabric** icon. In the **Add Fabric** pop-up window, specify a **Fabric Name** and select a **Fabric Template** from the drop-down list.

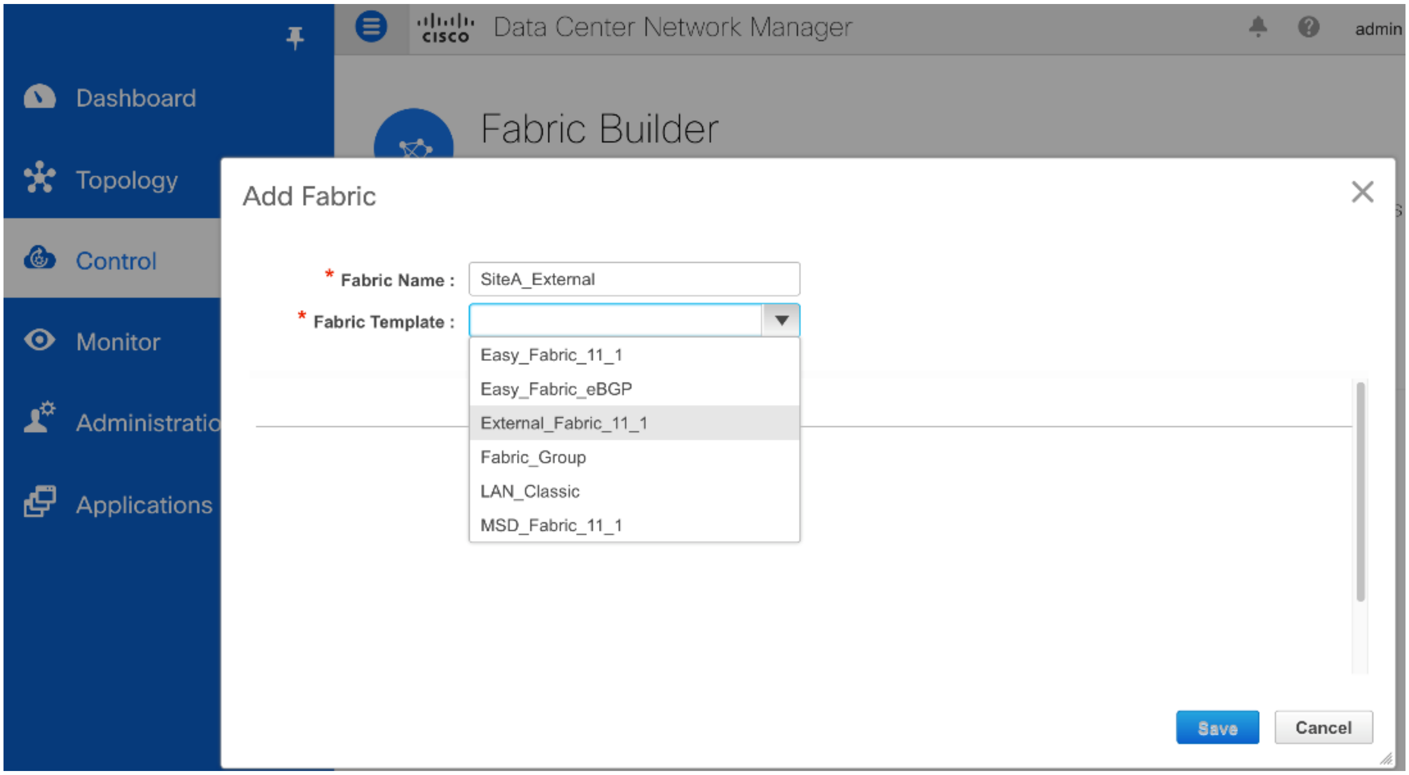

Add Fabric

* Fabric Name : SiteA_External

* Fabric Template :

Easy_Fabric_11_1
Easy_Fabric_eBGP
External_Fabric_11_1
Fabric_Group
LAN_Classic
MSD_Fabric_11_1

4. The pop-up window will now expand to include multiple tabs for configuring the external fabric. In the **General** tab, specify the **BGP AS#** information for the external fabric and deselect **Fabric Monitor Mode** so that the external fabric can be managed from Cisco DCNM.

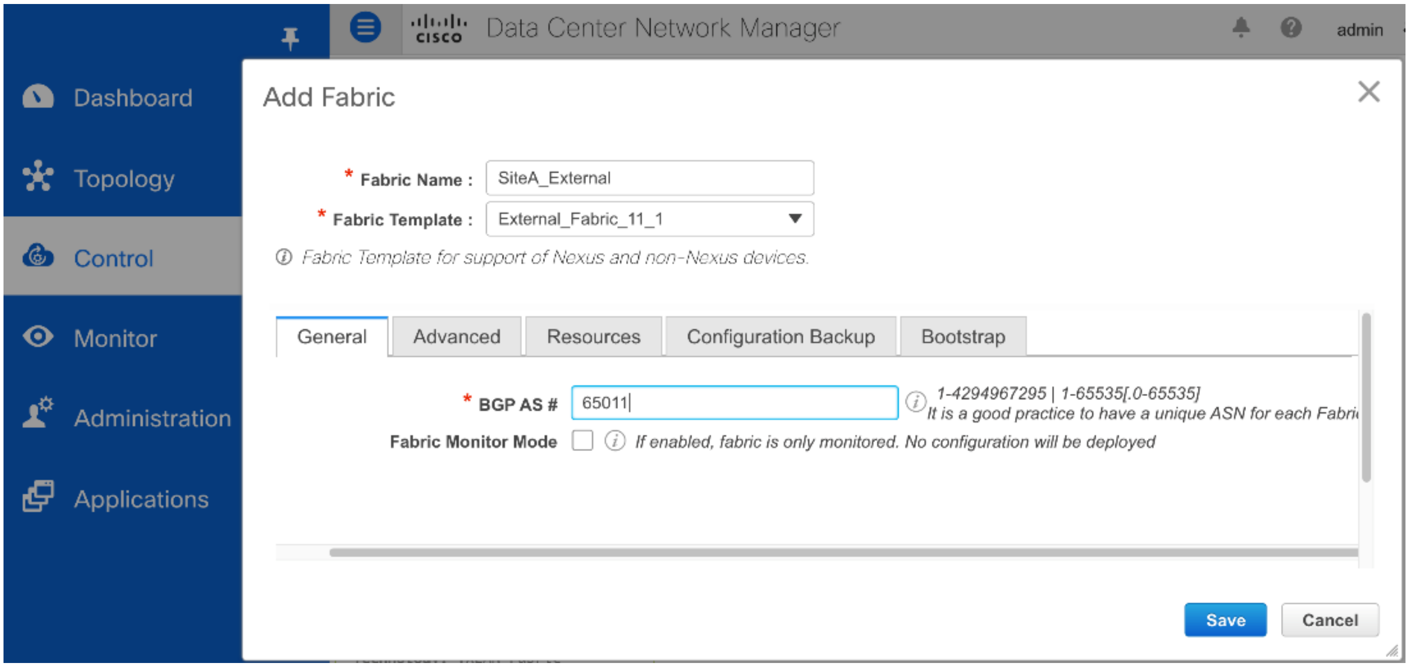

Add Fabric

* Fabric Name : SiteA_External

* Fabric Template : External_Fabric_11_1

Fabric Template for support of Nexus and non-Nexus devices.

* BGP AS # 65011 — 1-4294967295 | 1-65535[.0-65535]
It is a good practice to have a unique ASN for each Fabri

Fabric Monitor Mode □ — If enabled, fabric is only monitored. No configuration will be deployed

5. In the **Advanced** tab, select **Enable AAA IP Authorization** checkbox and leave everything else as-is.

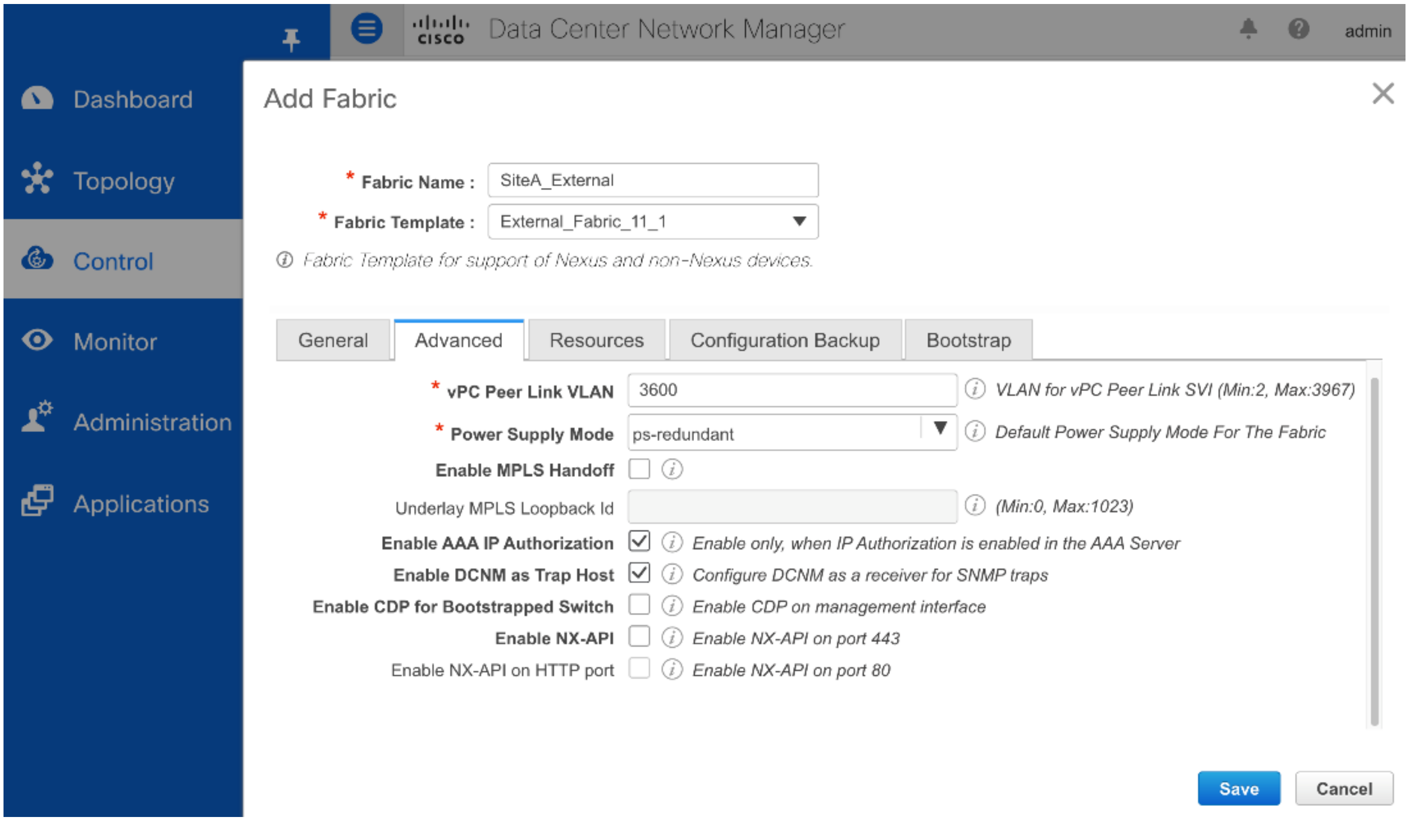

Add Fabric

* Fabric Name : SiteA_External

* Fabric Template : External_Fabric_11_1

Fabric Template for support of Nexus and non-Nexus devices.

* vPC Peer Link VLAN: 3600 — VLAN for vPC Peer Link SVI (Min:2, Max:3967)

* Power Supply Mode: ps-redundant — Default Power Supply Mode For The Fabric

Enable MPLS Handoff ☐

Underlay MPLS Loopback Id: (Min:0, Max:1023)

Enable AAA IP Authorization ☑ Enable only, when IP Authorization is enabled in the AAA Server

Enable DCNM as Trap Host ☑ Configure DCNM as a receiver for SNMP traps

Enable CDP for Bootstrapped Switch ☐ Enable CDP on management interface

Enable NX-API ☐ Enable NX-API on port 443

Enable NX-API on HTTP port ☐ Enable NX-API on port 80

6. (Optional) In the Resources tab, specify the Subinterface Dot1q Range and Underlay Routing Loopback IP Range.

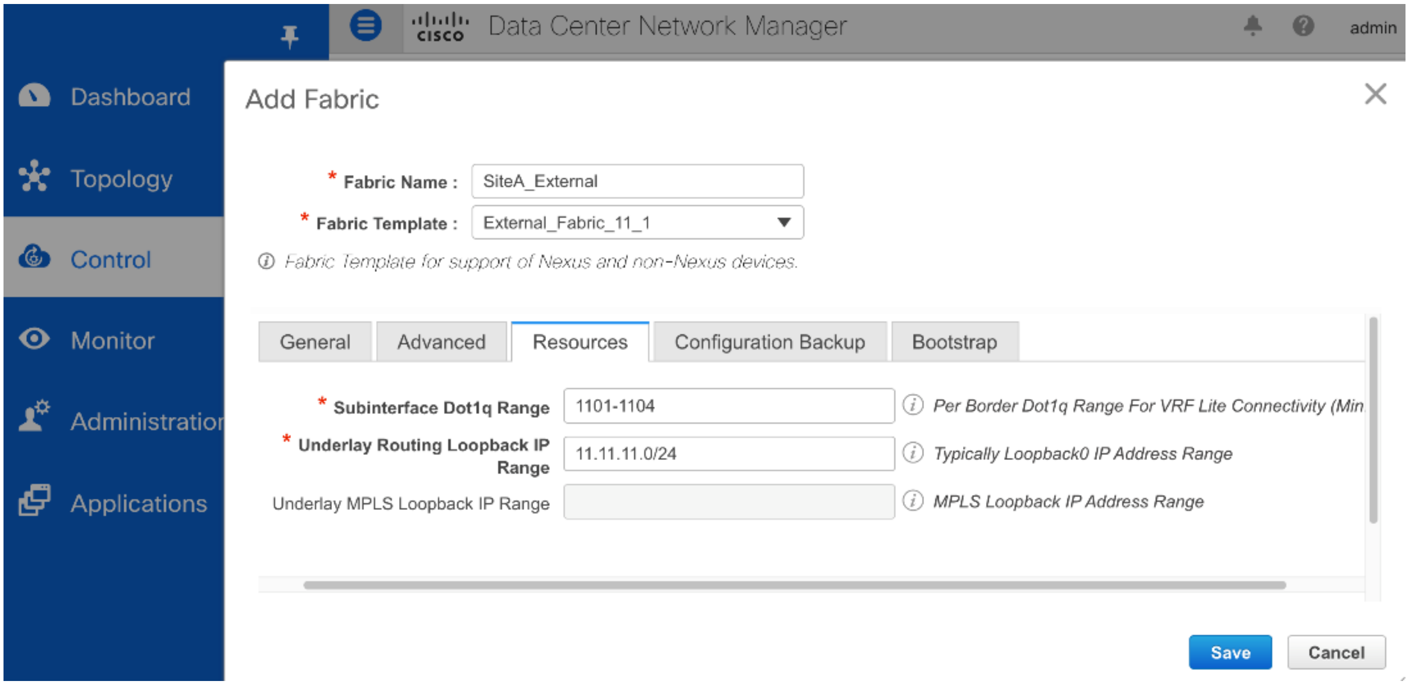

Add Fabric

* Fabric Name : SiteA_External

* Fabric Template : External_Fabric_11_1

Fabric Template for support of Nexus and non-Nexus devices.

* Subinterface Dot1q Range: 1101-1104 — Per Border Dot1q Range For VRF Lite Connectivity (Min

* Underlay Routing Loopback IP Range: 11.11.11.0/24 — Typically Loopback0 IP Address Range

Underlay MPLS Loopback IP Range: MPLS Loopback IP Address Range

7. (Optional) In the **Configuration Backup** tab, specify a backup schedule for the fabric as shown below.

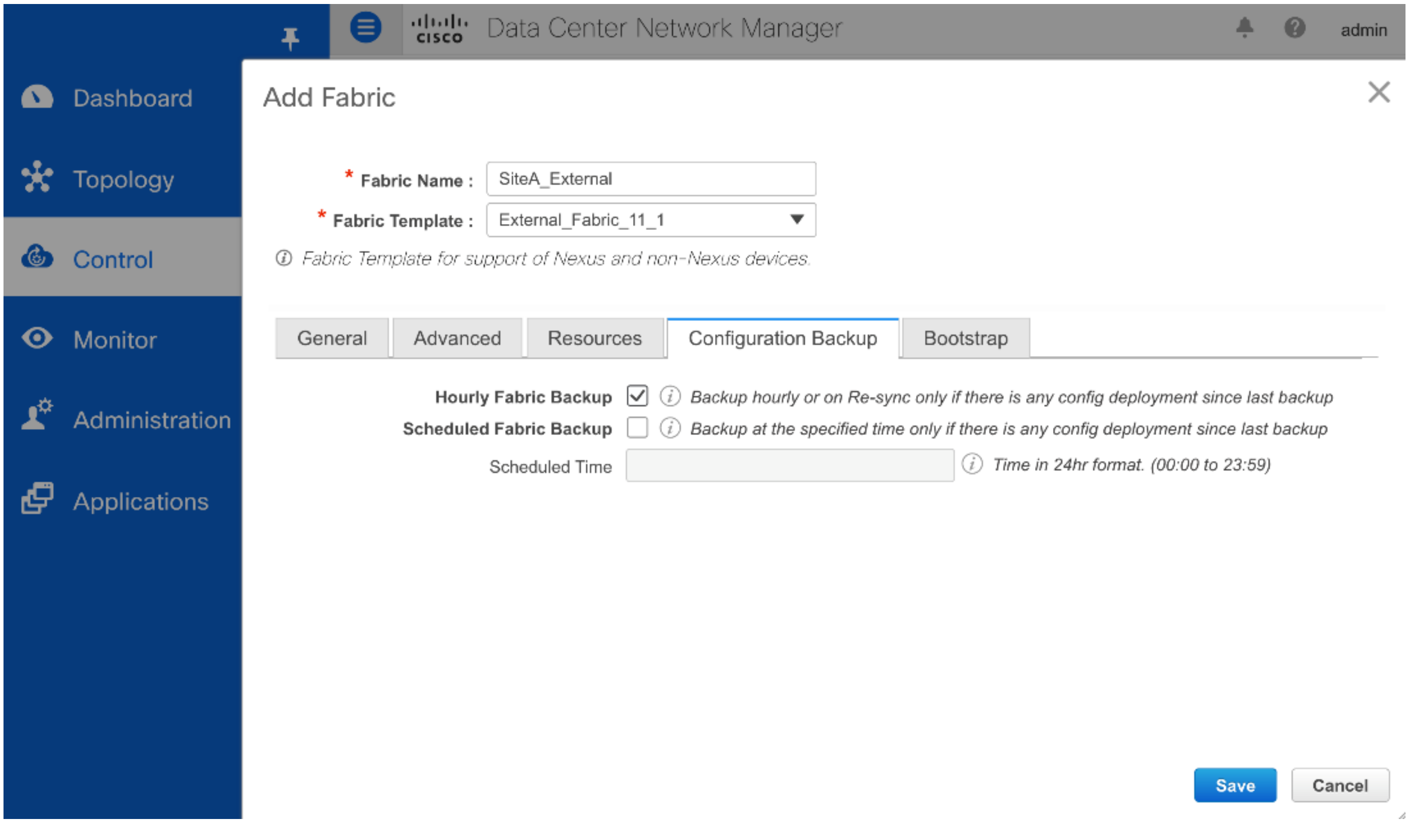

Add Fabric

* Fabric Name : SiteA_External

* Fabric Template : External_Fabric_11_1

Fabric Template for support of Nexus and non-Nexus devices.

Hourly Fabric Backup ☑ Backup hourly or on Re-sync only if there is any config deployment since last backup

Scheduled Fabric Backup □ Backup at the specified time only if there is any config deployment since last backup

Scheduled Time Time in 24hr format. (00:00 to 23:59)

8. (Optional) In the **Bootstrap** tab, customers can specify bootstrap information if POAP is used to discover and import the switches into the fabric. This was not used in this CVD.

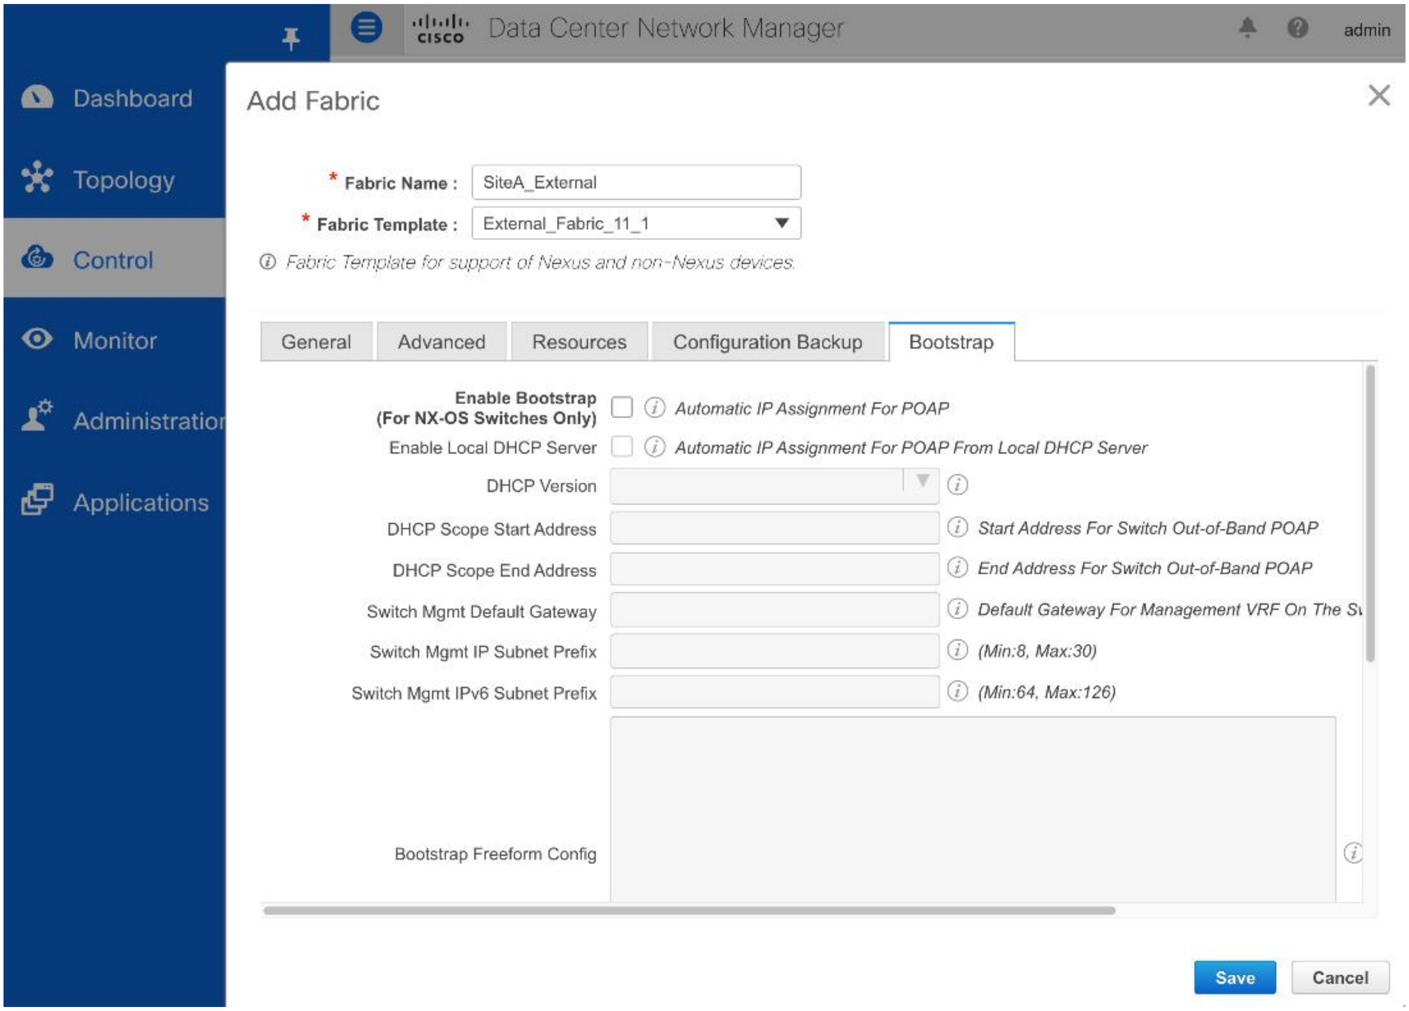

Add Fabric

* **Fabric Name :** SiteA_External

* **Fabric Template :** External_Fabric_11_1

Fabric Template for support of Nexus and non-Nexus devices.

Enable Bootstrap (For NX-OS Switches Only) □ *Automatic IP Assignment For POAP*

Enable Local DHCP Server □ *Automatic IP Assignment For POAP From Local DHCP Server*

DHCP Version

DHCP Scope Start Address *Start Address For Switch Out-of-Band POAP*

DHCP Scope End Address *End Address For Switch Out-of-Band POAP*

Switch Mgmt Default Gateway *Default Gateway For Management VRF On The Sv*

Switch Mgmt IP Subnet Prefix *(Min:8, Max:30)*

Switch Mgmt IPv6 Subnet Prefix *(Min:64, Max:126)*

Bootstrap Freeform Config

9. Click **Save** to save the fabric settings for the External Fabric in Site-A. You will get a pop-up on the right-bottom corner saying the Fabric was deployed successfully if the settings were saved. The saved settings are merely the configuration intent at this stage – they will need to be deployed on the switches for it to take effect.

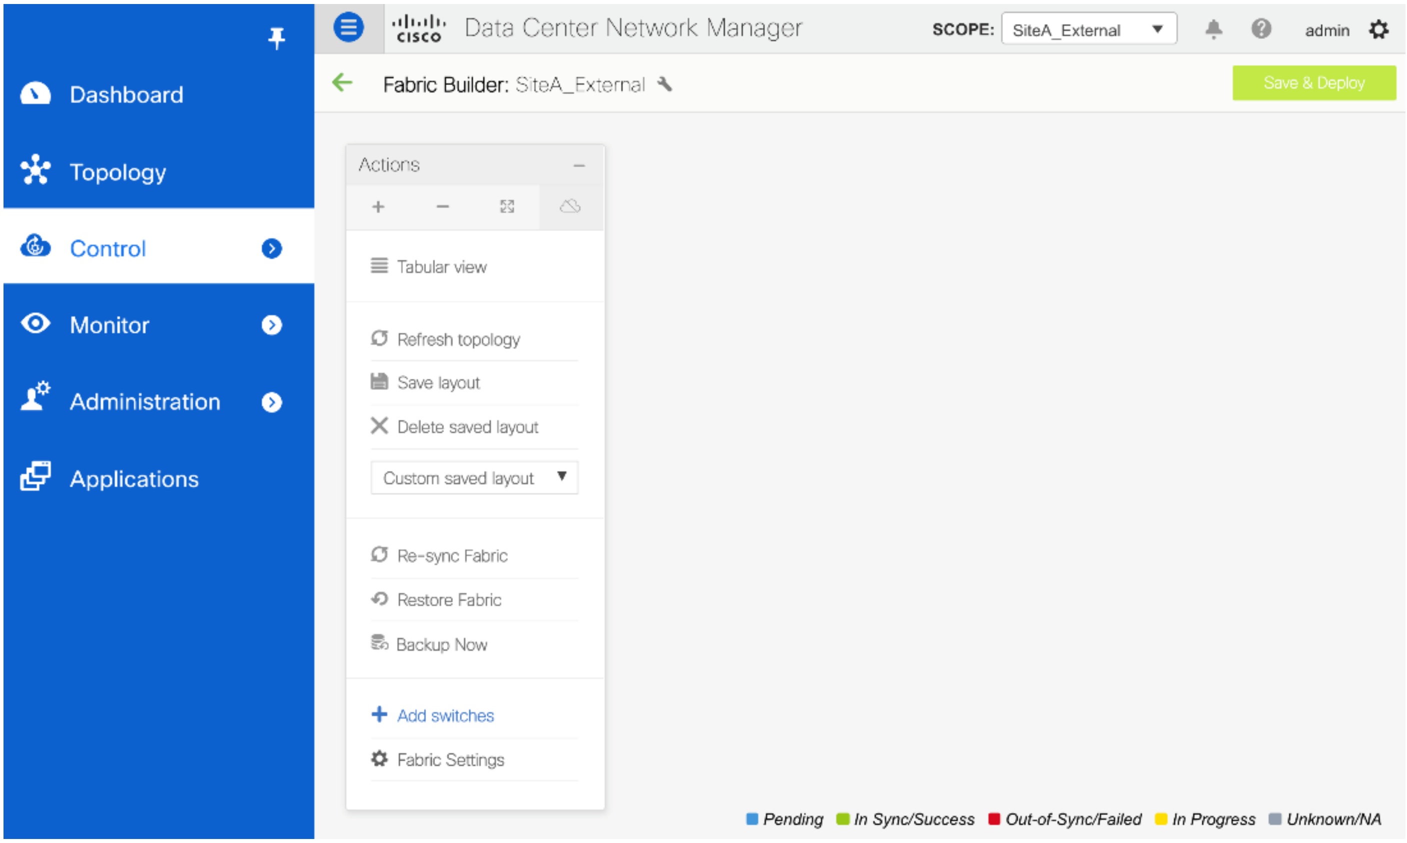

10. At this point, you can start adding switches to the External fabric. Note that you can also use ⚙ **Fabric Settings** in the **Actions** menu at any time to modify the parameters – however, once switches have been added to the fabric, you will need to do a **Save & Deploy** (top-right corner) in order to save the settings and then to apply them to the switches in the fabric.

11. Proceed to the next section to discover and add switches to the external fabric to connect to the Site-A data center fabric.

Add Gateway Switches to the External Fabric

As stated earlier, this design assumes a greenfield deployment where the fabric is built from the ground-up. Therefore, this section walks through the discovery, addition, and initial configuration of gateway switches to the external fabric to connect to Site-A data center fabric.

To add gateway switches to the external fabric, follow these steps:

1. In the right-window pane, verify that the **SCOPE:** is **SiteA_External** in the drop-down list near the top-right corner. From the **Actions** menu, select and click **Add Switches**.

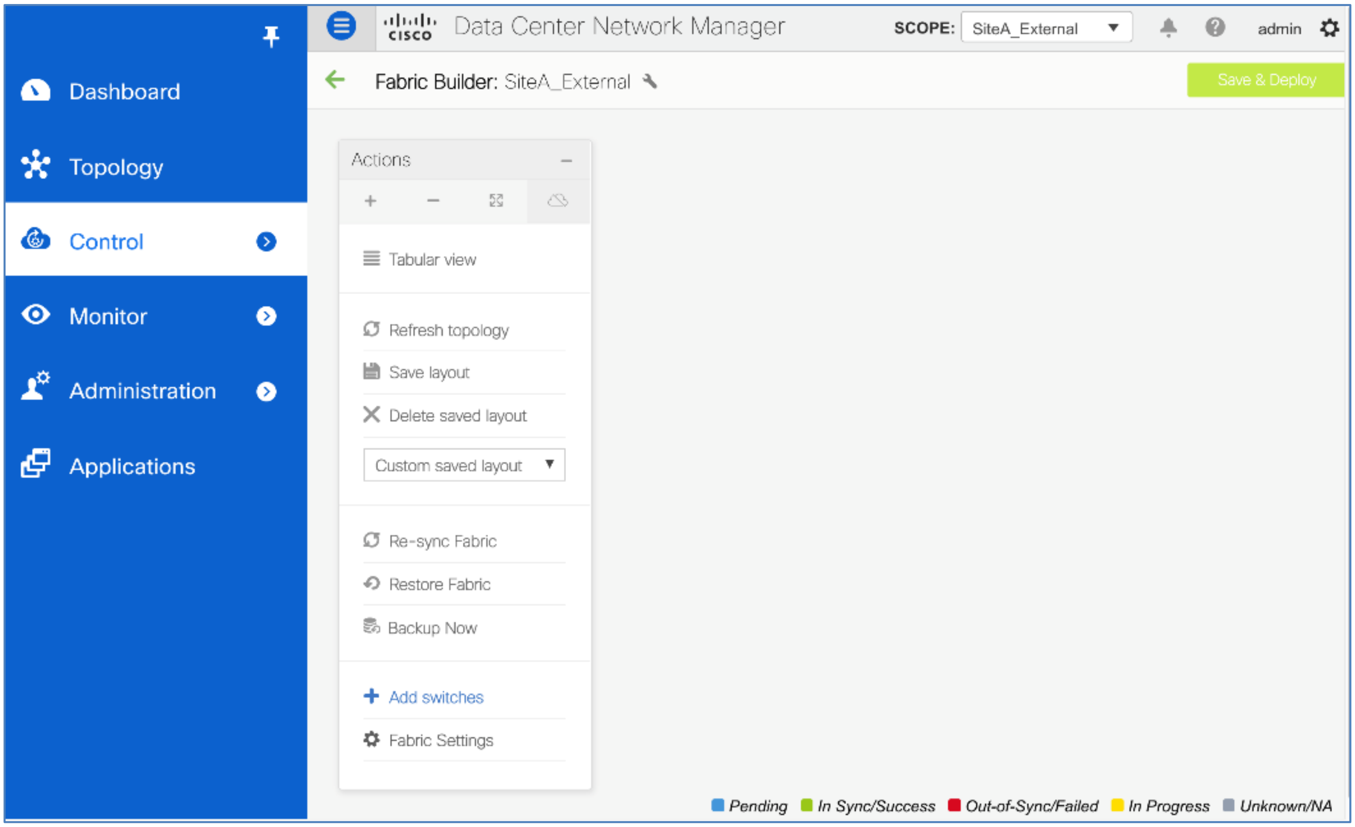

2. In the **Inventory Management** pop-up window, select the **Discover Existing Switches** tab. Note that you can also POAP on Cisco DCNM to discover and add switches to the VXLAN fabric. For the **Seed IP,** specify the IP address range of the switches that need to be discovered. For the **Username** and **Password,** specify the administrator username and password for the switches that you can use to log-on to the switches. For the **Max Hops**, specify ‘0’ to minimize the discovery time.

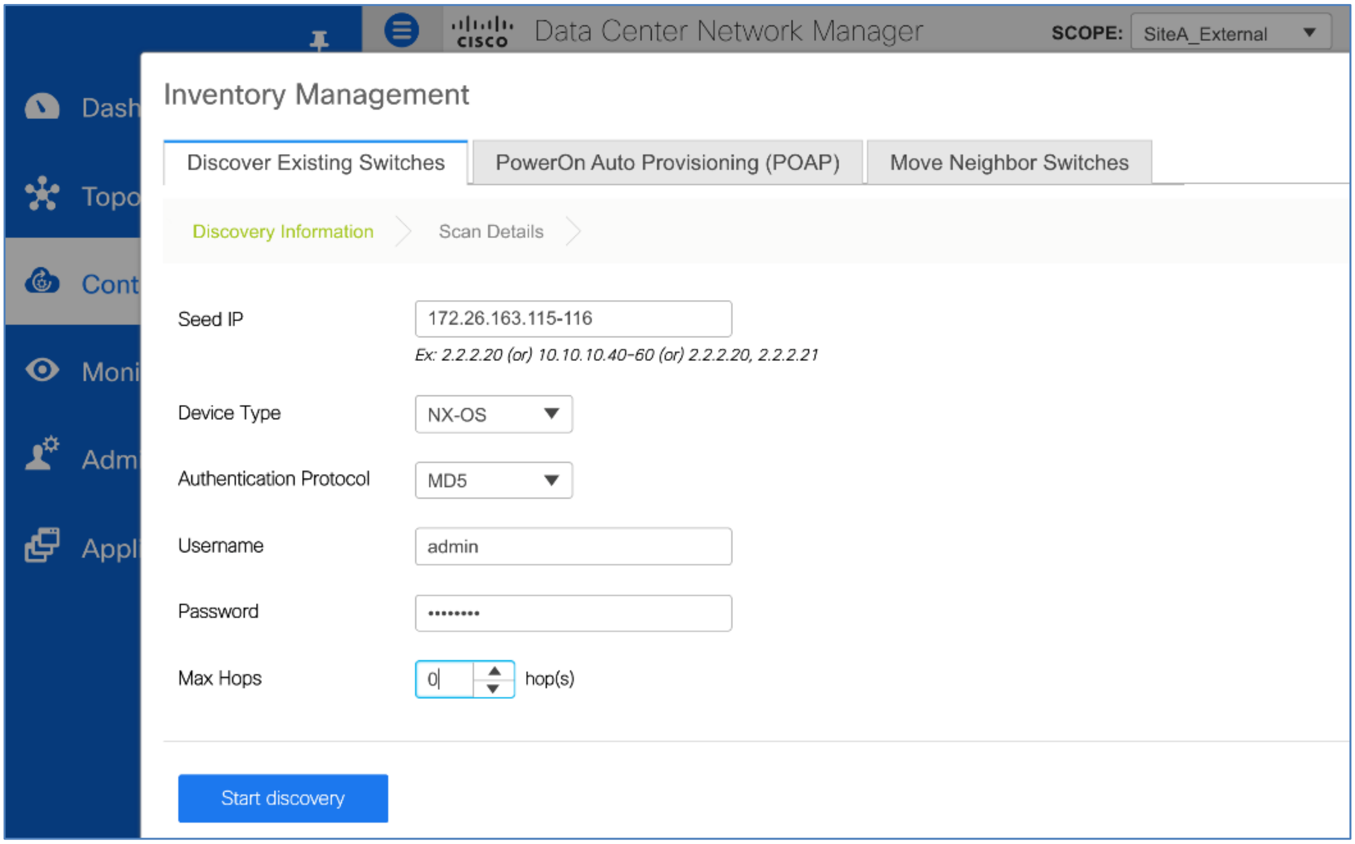

Seed IP: 172.26.163.115-116

Ex: 2.2.2.20 (or) 10.10.10.40-60 (or) 2.2.2.20, 2.2.2.21

Device Type: NX-OS

Authentication Protocol: MD5

Username: admin

Password: ••••••••

Max Hops: 0 hop(s)

3. Click **Start discovery** to start the discovery process. You should see a spinning wheel in the red **Abort Request** button at the bottom as Cisco DCNM attempts to discover and find these switches. Once the discovery completes, you will be provided with a list of switches that can be imported into the fabric.

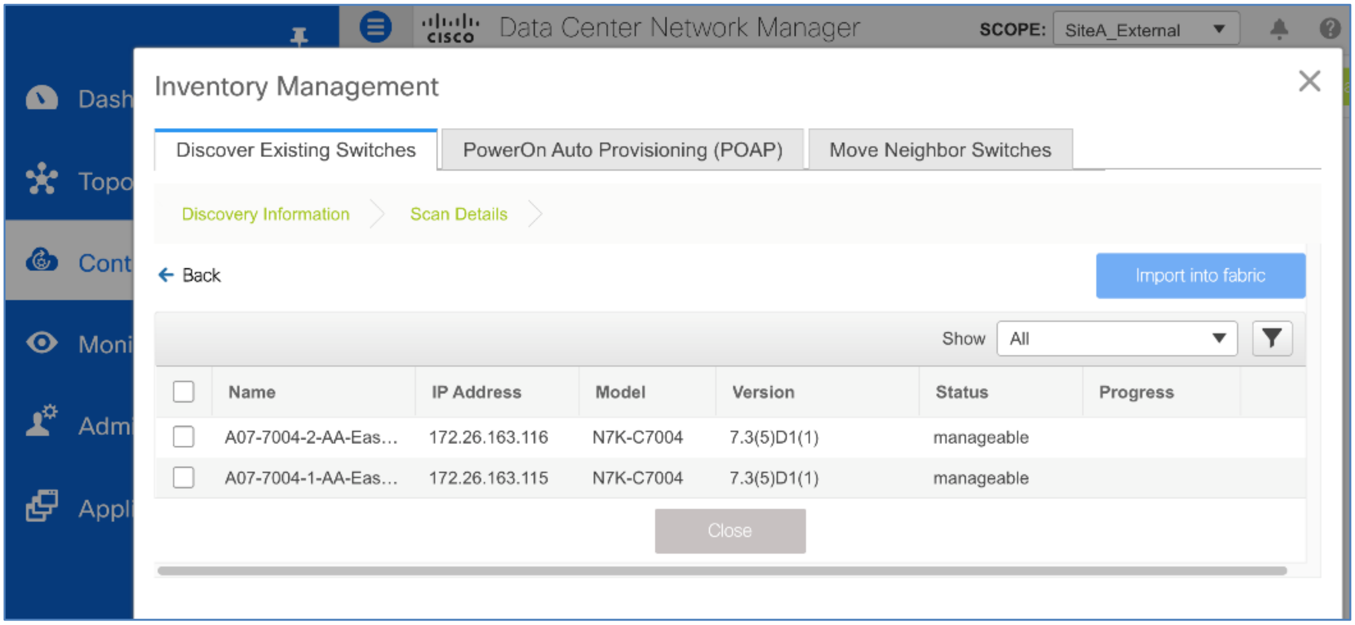

□	Name	IP Address	Model	Version	Status	Progress
□	A07-7004-2-AA-Eas...	172.26.163.116	N7K-C7004	7.3(5)D1(1)	manageable	
□	A07-7004-1-AA-Eas...	172.26.163.115	N7K-C7004	7.3(5)D1(1)	manageable	

4. Use checkboxes to the left of the switches to select the switches that should be imported into this fabric. In this case, all switches in the list are selected. Click **Import into fabric**.

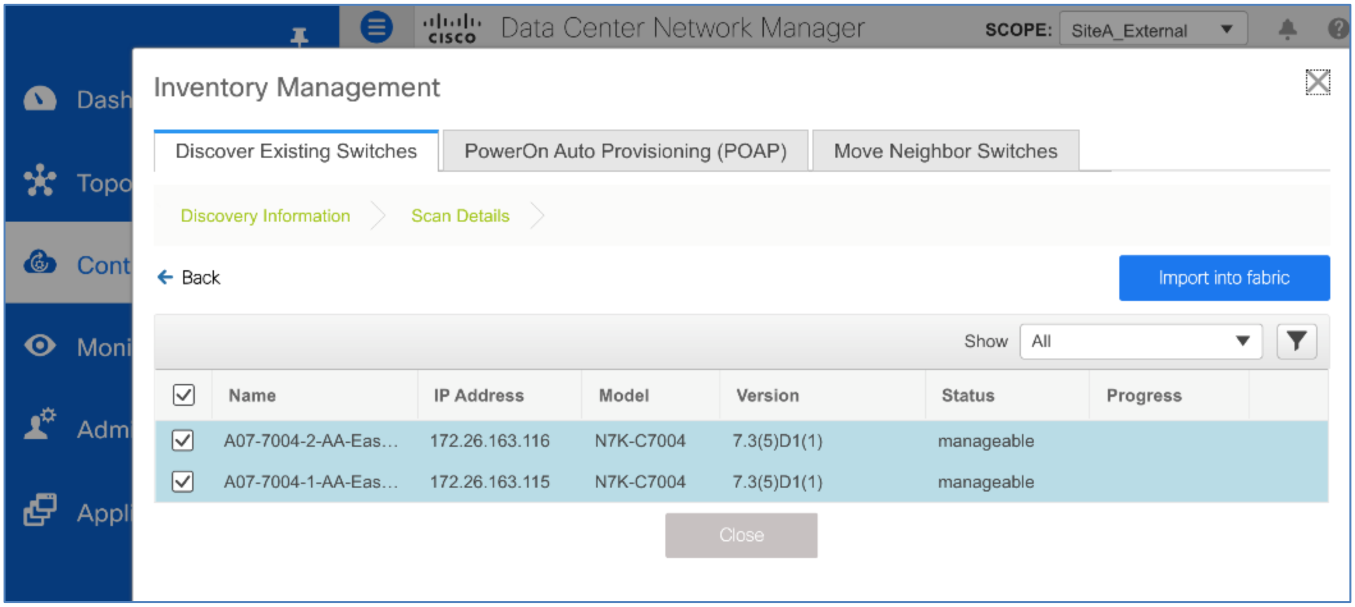

☑	Name	IP Address	Model	Version	Status	Progress
☑	A07-7004-2-AA-Eas…	172.26.163.116	N7K-C7004	7.3(5)D1(1)	manageable	
☑	A07-7004-1-AA-Eas…	172.26.163.115	N7K-C7004	7.3(5)D1(1)	manageable	

5. You can see the progress of the import in the **Progress** column for each switch being imported.

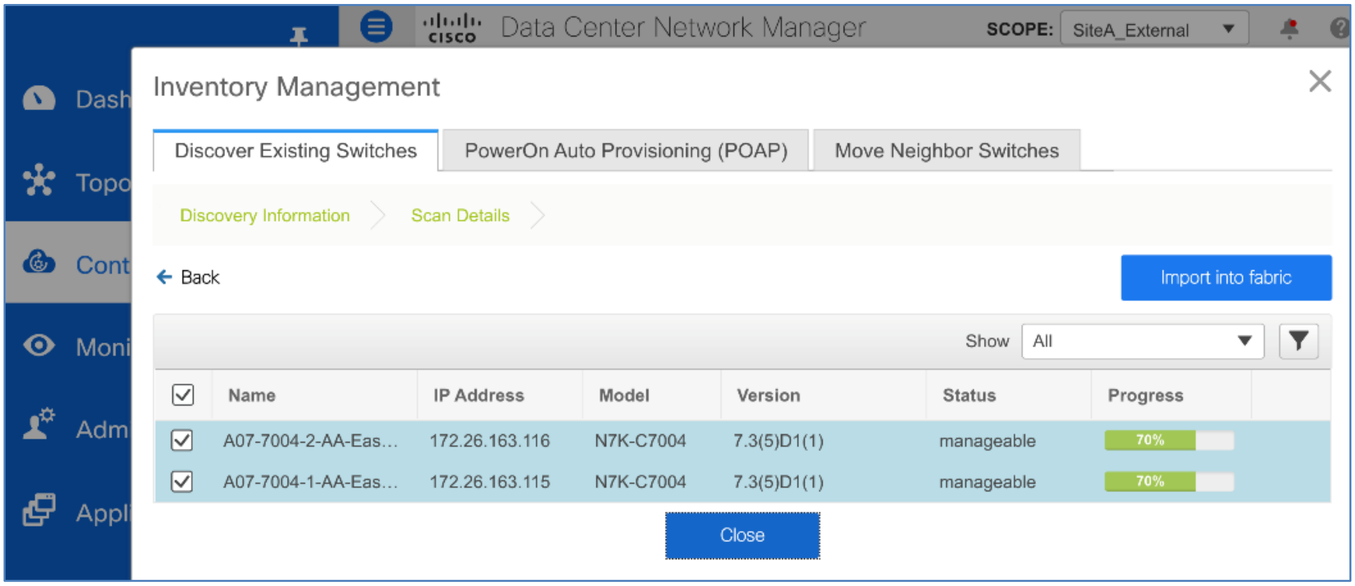

☑	Name	IP Address	Model	Version	Status	Progress
☑	A07-7004-2-AA-Eas…	172.26.163.116	N7K-C7004	7.3(5)D1(1)	manageable	70%
☑	A07-7004-1-AA-Eas…	172.26.163.115	N7K-C7004	7.3(5)D1(1)	manageable	70%

6. Once imported is complete, the **Progress** column will show **done**. Click **Close.**

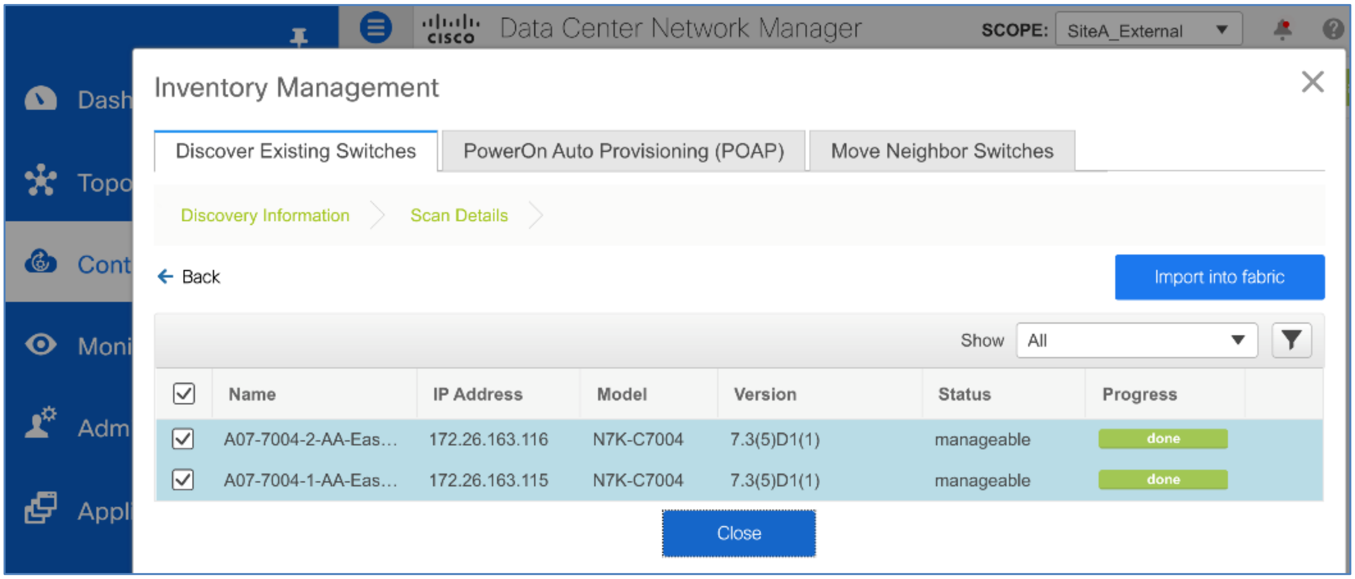

Inventory Management

Discover Existing Switches | PowerOn Auto Provisioning (POAP) | Move Neighbor Switches

Discovery Information > Scan Details >

← Back

Import into fabric

Show All

☑	Name	IP Address	Model	Version	Status	Progress
☑	A07-7004-2-AA-Eas…	172.26.163.116	N7K-C7004	7.3(5)D1(1)	manageable	done
☑	A07-7004-1-AA-Eas…	172.26.163.115	N7K-C7004	7.3(5)D1(1)	manageable	done

Close

7. For each switch, verify that the **Discovery Status** is **ok.** Also, verify the **Role** of each switch and modify as needed. To change the role, go back to the topology view by clicking on the green left arrow in the top left corner of the right window pane.

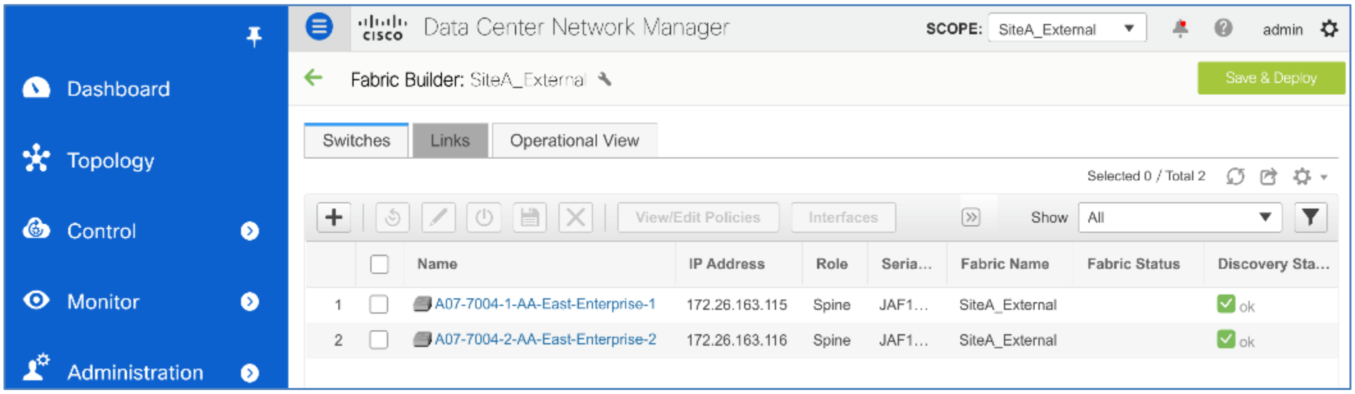

Fabric Builder: SiteA_External

Save & Deploy

Switches | Links | Operational View

Selected 0 / Total 2

View/Edit Policies | Interfaces | Show All

	☐	Name	IP Address	Role	Seria…	Fabric Name	Fabric Status	Discovery Sta…
1	☐	A07-7004-1-AA-East-Enterprise-1	172.26.163.115	Spine	JAF1…	SiteA_External		✓ ok
2	☐	A07-7004-2-AA-East-Enterprise-2	172.26.163.116	Spine	JAF1…	SiteA_External		✓ ok

8. From the **Topology** view, change the role of the external gateway switches to **Edge Router** role. Select the first switch from the topology and right-click. From the menu, select **Set role** and then **Edge Router** from the roles list. A small window will pop-up on the bottom right to confirm that the role change was successful.

Cisco recommends the Edge Router role to set up a VRF-lite Inter-Fabric Connection (IFC) from a Border device to an Edge device, which is what this design uses.

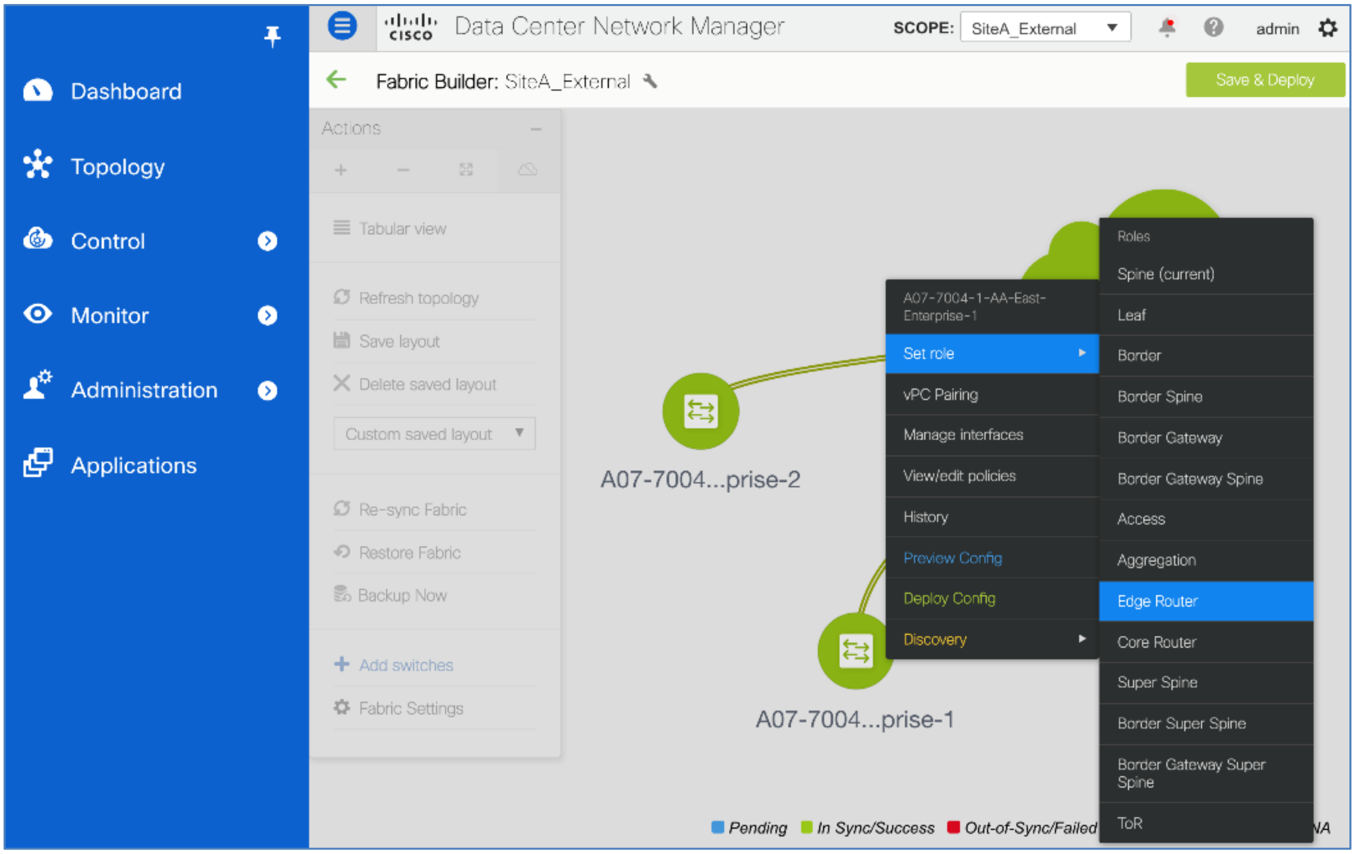

9. Repeat step 18 for the second switch.

10. Verify the roles have changed. From the **Actions** menu/window, click **Tabular view.** For each switch, verify that the **Role** change is correct. You are now ready to save and apply the configuration changes to the switches. Click the **Save & Deploy** button to apply the changes to both switches in the list.

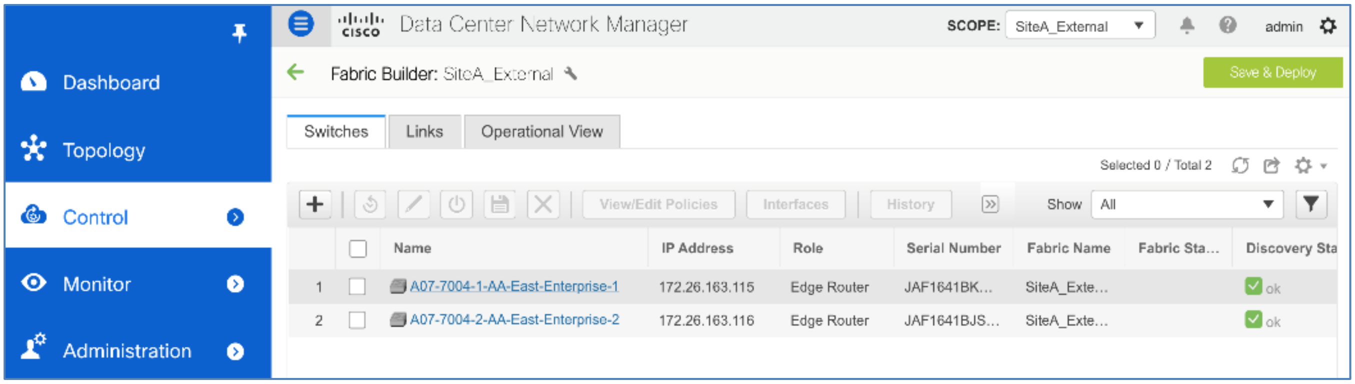

		Name	IP Address	Role	Serial Number	Fabric Name	Fabric Sta...	Discovery Sta
	☐							
1	☐	A07-7004-1-AA-East-Enterprise-1	172.26.163.115	Edge Router	JAF1641BK...	SiteA_Exte...		✓ ok
2	☐	A07-7004-2-AA-East-Enterprise-2	172.26.163.116	Edge Router	JAF1641BJS...	SiteA_Exte...		✓ ok

11. In the **Config Deployment** pop-up window, you can typically see the number of lines of configuration that will be deployed on each switch. The configuration deployed will vary depending on the role of the switch. In this case, zero lines of configuration is deployed. However, the configuration will occur once the inter-fabric links are deployed between External Fabric and Site-A in a later step. Click the **Deploy Config** button.

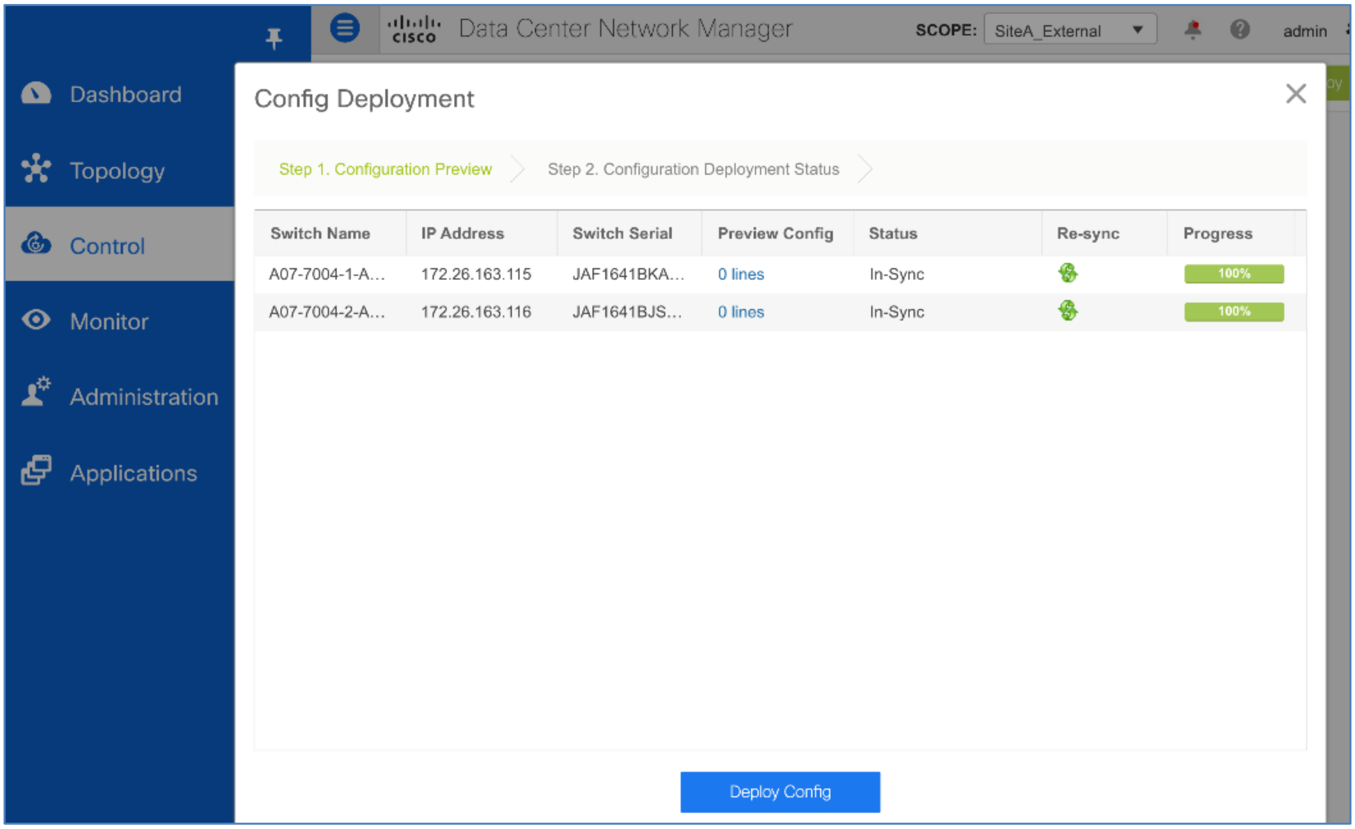

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
A07-7004-1-A...	172.26.163.115	JAF1641BKA...	0 lines	In-Sync	⟳	100%
A07-7004-2-A...	172.26.163.116	JAF1641BJS...	0 lines	In-Sync	⟳	100%

12. In the **Config Deployment** window, the deployment should complete and go to a **COMPLETED** status. Click **Close**.

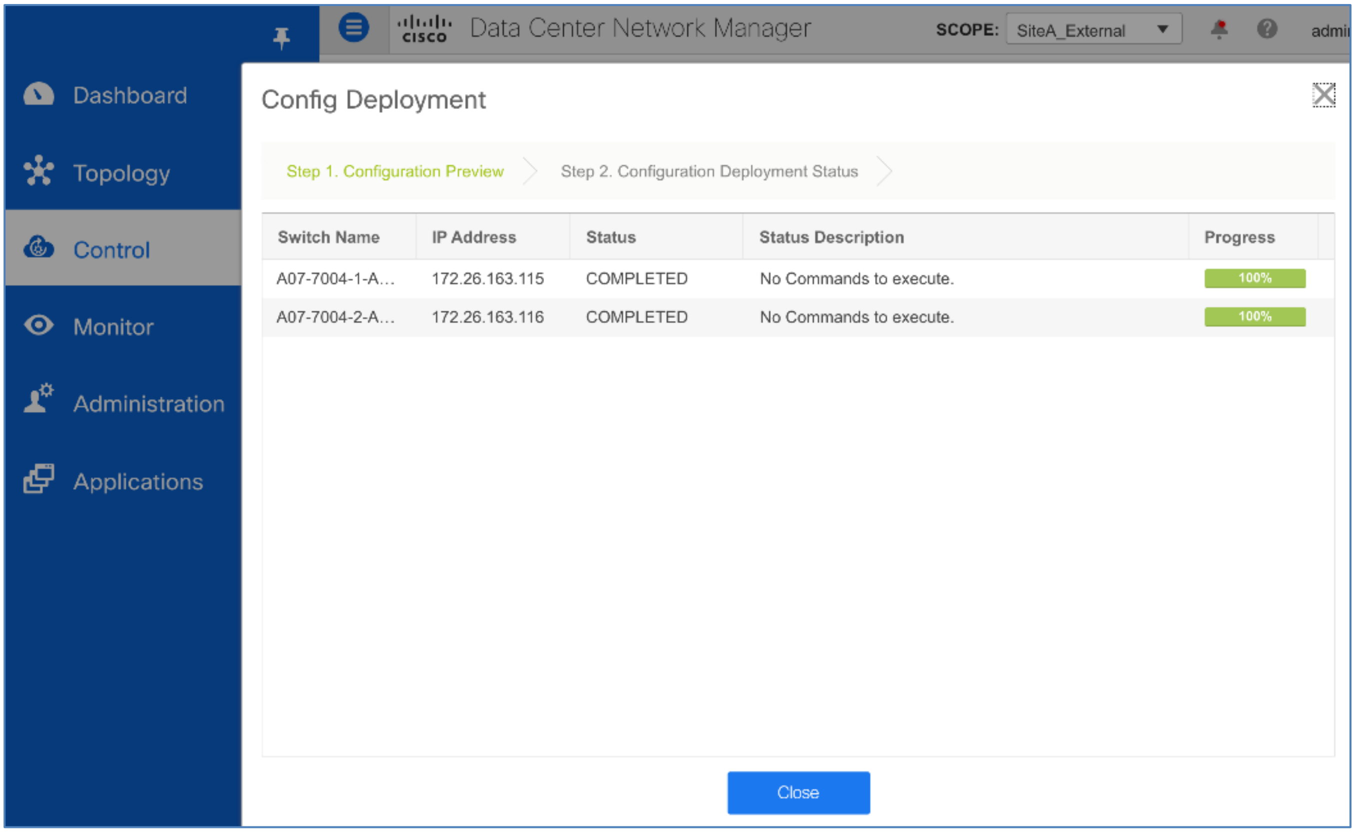

Config Deployment

Step 1. Configuration Preview · Step 2. Configuration Deployment Status

Switch Name	IP Address	Status	Status Description	Progress
A07-7004-1-A...	172.26.163.115	COMPLETED	No Commands to execute.	100%
A07-7004-2-A...	172.26.163.116	COMPLETED	No Commands to execute.	100%

13. Go to Topology view. From the **Actions** menu/window, select **Hierarchical** from the drop-down list. The topology view should now change based on the role of the devices. Select **Save layout** from the **Actions** menu.

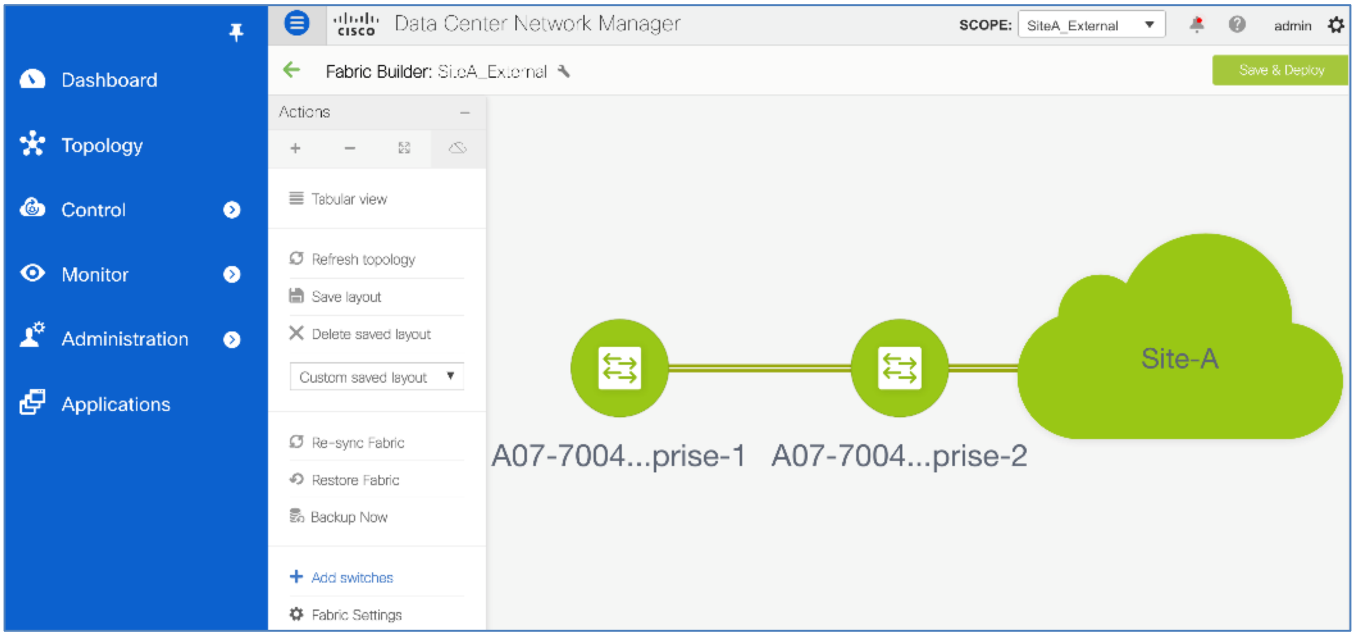

14. The next step is to deploy the Inter-Fabric Connections (IFC) between the external fabric and Site-A.

Deploy Inter-Fabric Connections between External Fabric and Site-A

In this design, the Inter-Fabric Connections (IFCs) are auto-deployed and configured. From the external fabric, you can view and delete IFCs, but you cannot create/edit/deploy them – you must do this from Site-A. To verify that IFCs are discovered and deployed between the External fabric and Site-A data center fabric, follow these steps:

1. From the left navigation menu, select **Control > Fabric Builder**. Select and click **Site-A** fabric from the two fabrics listed.

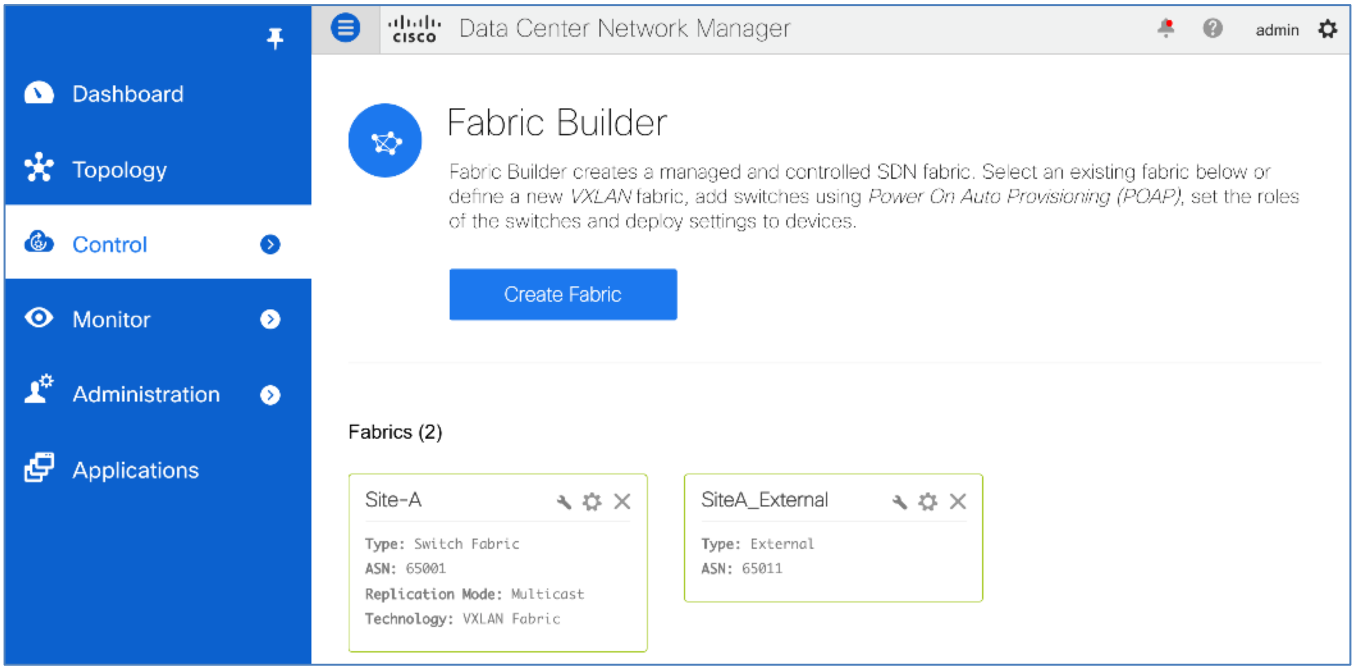

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (2)

Site-A

```
Type: Switch Fabric
ASN: 65001
Replication Mode: Multicast
Technology: VXLAN Fabric
```

SiteA_External

```
Type: External
ASN: 65011
```

2. From the Site-A topology view, in the **Actions** menu/window, select **Hierarchical** from the drop-down list. The topology view should now change based on the role of the devices. Select **Save layout** from the **Actions** menu.

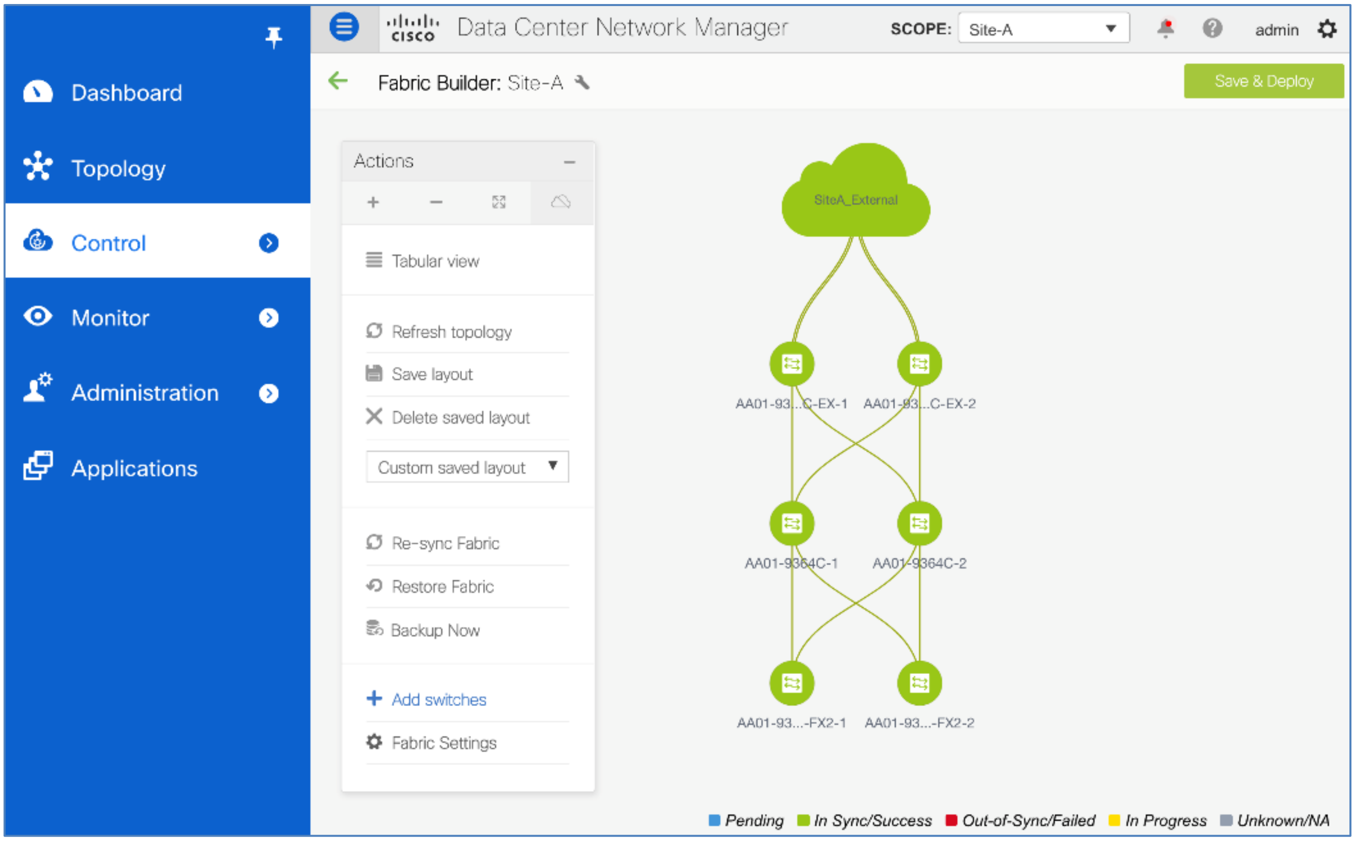

3. From the **Actions** menu/window, select **Tabular view.** Select the **Links** tab. Click **Fabric Name** to sort and find the 4 IFC links used in this design – they are automatically deployed by Cisco DCNM. The **Fabric Name** will have both fabrics in the name as shown below. Verify that each IFC link has a status of **Link Present**, **Admin State** of **Up:Up** and **Oper State** of **Up:Up**.

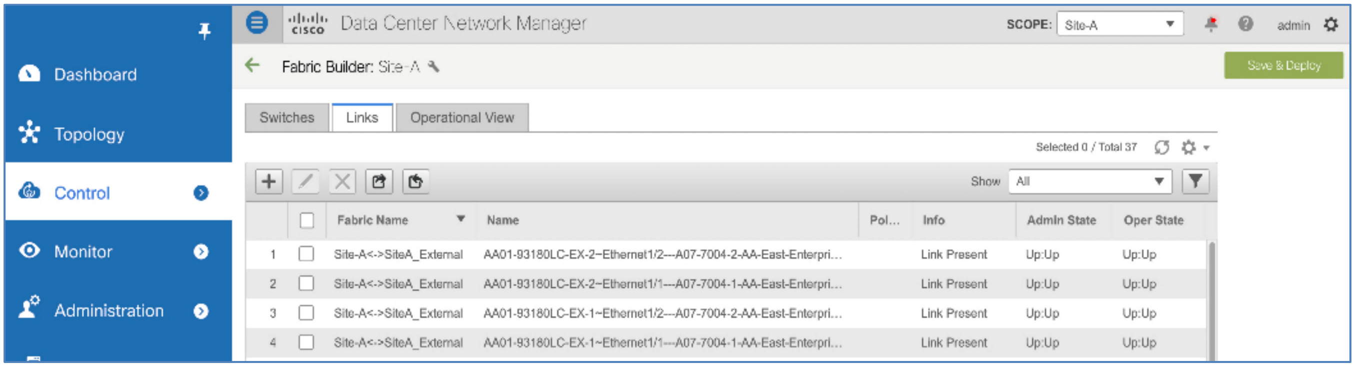

	□	Fabric Name	Name	Pol...	Info	Admin State	Oper State
1	□	Site-A<->SiteA_External	AA01-93180LC-EX-2~Ethernet1/2---A07-7004-2-AA-East-Enterpri...		Link Present	Up:Up	Up:Up
2	□	Site-A<->SiteA_External	AA01-93180LC-EX-2~Ethernet1/1---A07-7004-1-AA-East-Enterpri...		Link Present	Up:Up	Up:Up
3	□	Site-A<->SiteA_External	AA01-93180LC-EX-1~Ethernet1/2---A07-7004-2-AA-East-Enterpri...		Link Present	Up:Up	Up:Up
4	□	Site-A<->SiteA_External	AA01-93180LC-EX-1~Ethernet1/1---A07-7004-1-AA-East-Enterpri...		Link Present	Up:Up	Up:Up

4. Select one of the IFC links by selecting the checkbox to the left of the link.

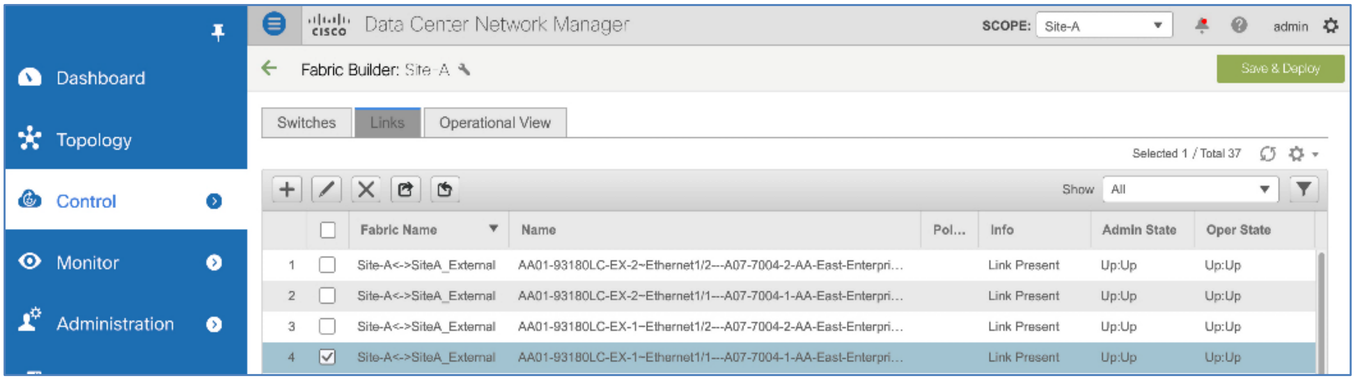

	☐	Fabric Name	Name	Pol...	Info	Admin State	Oper State
1	☐	Site-A<->SiteA_External	AA01-93180LC-EX-2~Ethernet1/2---A07-7004-2-AA-East-Enterpri...		Link Present	Up:Up	Up:Up
2	☐	Site-A<->SiteA_External	AA01-93180LC-EX-2~Ethernet1/1---A07-7004-1-AA-East-Enterpri...		Link Present	Up:Up	Up:Up
3	☐	Site-A<->SiteA_External	AA01-93180LC-EX-1~Ethernet1/2---A07-7004-2-AA-East-Enterpri...		Link Present	Up:Up	Up:Up
4	☑	Site-A<->SiteA_External	AA01-93180LC-EX-1~Ethernet1/1---A07-7004-1-AA-East-Enterpri...		Link Present	Up:Up	Up:Up

5. Click the ✎ to edit the IFC link. In the **Link Management – Edit Link** pop-up window, under the **Link Profile > General** section, note that the **Source IP Address/Mask** and **Destination IP** are not configured. Also, the **Auto Deploy Flag** is disabled**.** All other fields are populated as shown below. Click **X** to close this window.

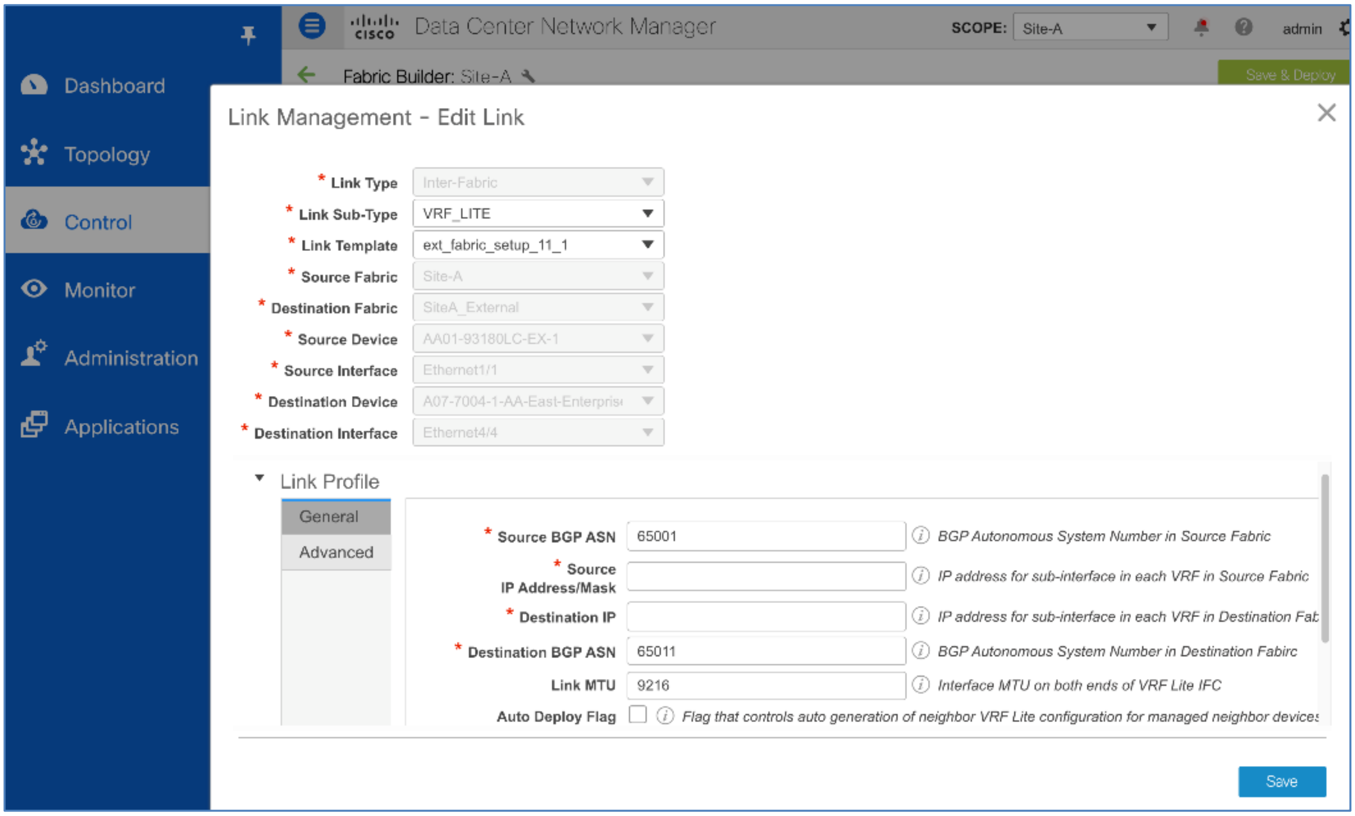

Link Management - Edit Link

* Link Type: Inter-Fabric
* Link Sub-Type: VRF_LITE
* Link Template: ext_fabric_setup_11_1
* Source Fabric: Site-A
* Destination Fabric: SiteA_External
* Source Device: AA01-93180LC-EX-1
* Source Interface: Ethernet1/1
* Destination Device: A07-7004-1-AA-East-Enterprise
* Destination Interface: Ethernet4/4

Link Profile

General
Advanced

* Source BGP ASN: 65001 — BGP Autonomous System Number in Source Fabric
* Source IP Address/Mask: — IP address for sub-interface in each VRF in Source Fabric
* Destination IP: — IP address for sub-interface in each VRF in Destination Fab
* Destination BGP ASN: 65011 — BGP Autonomous System Number in Destination Fabric
Link MTU: 9216 — Interface MTU on both ends of VRF Lite IFC
Auto Deploy Flag: ☐ — Flag that controls auto generation of neighbor VRF Lite configuration for managed neighbor devices

6. From the **Links** tab view, click the **Save & Deploy** button.

7. From the **Config Deployment** pop-up window, click the **Deploy Config** button.

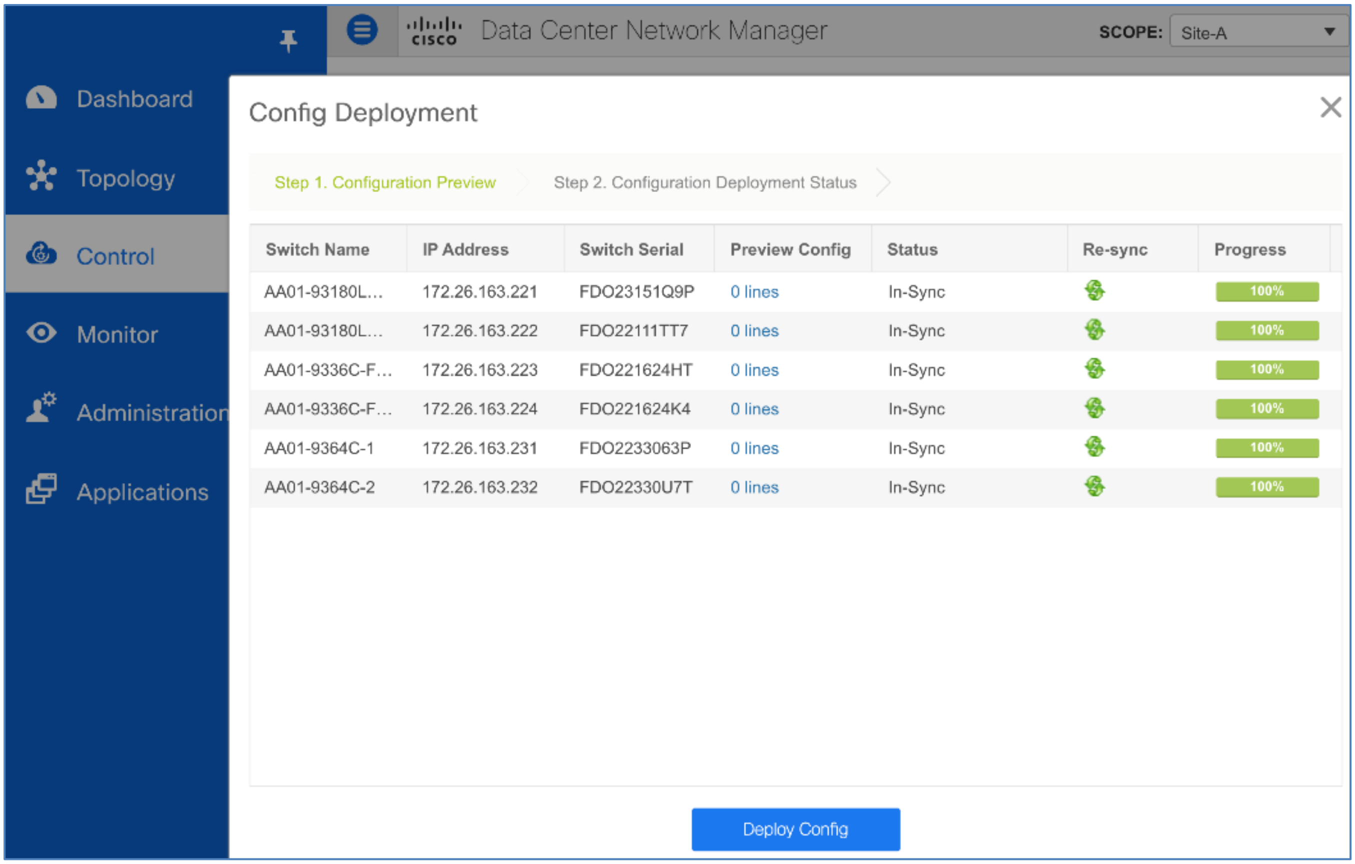

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
AA01-93180L…	172.26.163.221	FDO23151Q9P	0 lines	In-Sync	⟳	100%
AA01-93180L…	172.26.163.222	FDO22111TT7	0 lines	In-Sync	⟳	100%
AA01-9336C-F…	172.26.163.223	FDO221624HT	0 lines	In-Sync	⟳	100%
AA01-9336C-F…	172.26.163.224	FDO221624K4	0 lines	In-Sync	⟳	100%
AA01-9364C-1	172.26.163.231	FDO2233063P	0 lines	In-Sync	⟳	100%
AA01-9364C-2	172.26.163.232	FDO22330U7T	0 lines	In-Sync	⟳	100%

8. When the deployment completes and the status is **COMPLETED**, click the **Close** button to close this window.

9. From the **Links** tab view, select the same IFC link as in step 5 above and click the ✎ to edit the IFC link. In the **Link Management – Edit Link** pop-up window, under the **Link Profile > General** section, note that the **Source IP Address/Mask** and **Destination IP** are now configured. Also, the **Auto Deploy Flag** is now enabled**.** The **Save & Deploy** from Step 6 applied the Fabric Settings for VRF-Lite and auto-deployed the necessary configuration. Click **X** to close the window.

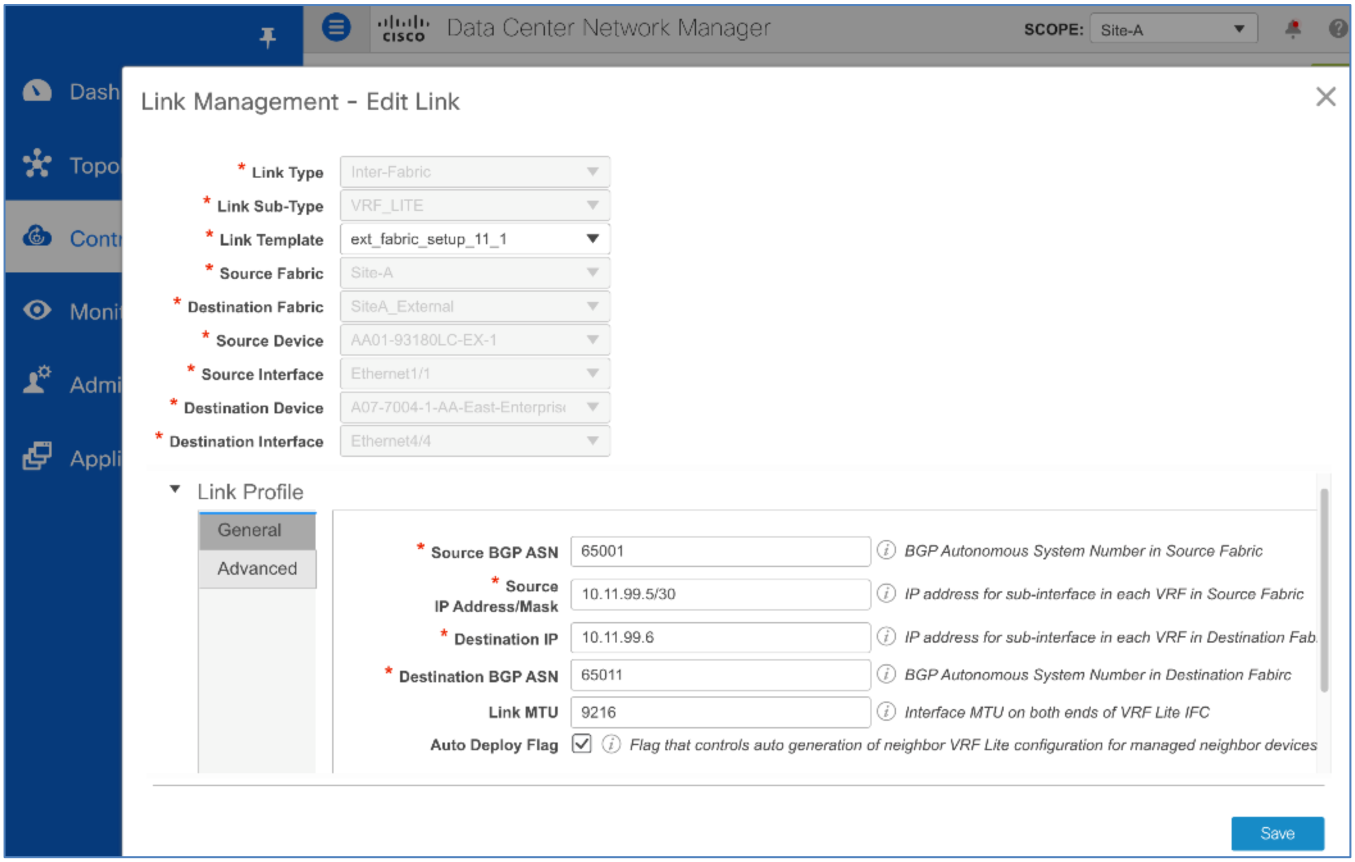

Link Management - Edit Link

* Link Type	Inter-Fabric
* Link Sub-Type	VRF_LITE
* Link Template	ext_fabric_setup_11_1
* Source Fabric	Site-A
* Destination Fabric	SiteA_External
* Source Device	AA01-93180LC-EX-1
* Source Interface	Ethernet1/1
* Destination Device	A07-7004-1-AA-East-Enterpris
* Destination Interface	Ethernet4/4

Link Profile

General | Advanced

* Source BGP ASN	65001	BGP Autonomous System Number in Source Fabric
* Source IP Address/Mask	10.11.99.5/30	IP address for sub-interface in each VRF in Source Fabric
* Destination IP	10.11.99.6	IP address for sub-interface in each VRF in Destination Fab
* Destination BGP ASN	65011	BGP Autonomous System Number in Destination Fabirc
Link MTU	9216	Interface MTU on both ends of VRF Lite IFC
Auto Deploy Flag	☑	Flag that controls auto generation of neighbor VRF Lite configuration for managed neighbor devices

10. You can verify the remaining three IFC links – they should all be configured also after the **Save & Deploy**.

11. For each IFC link, in the Link Profile > Advanced section, configure the Source Interface Description and the Destination Interface Description as shown below.

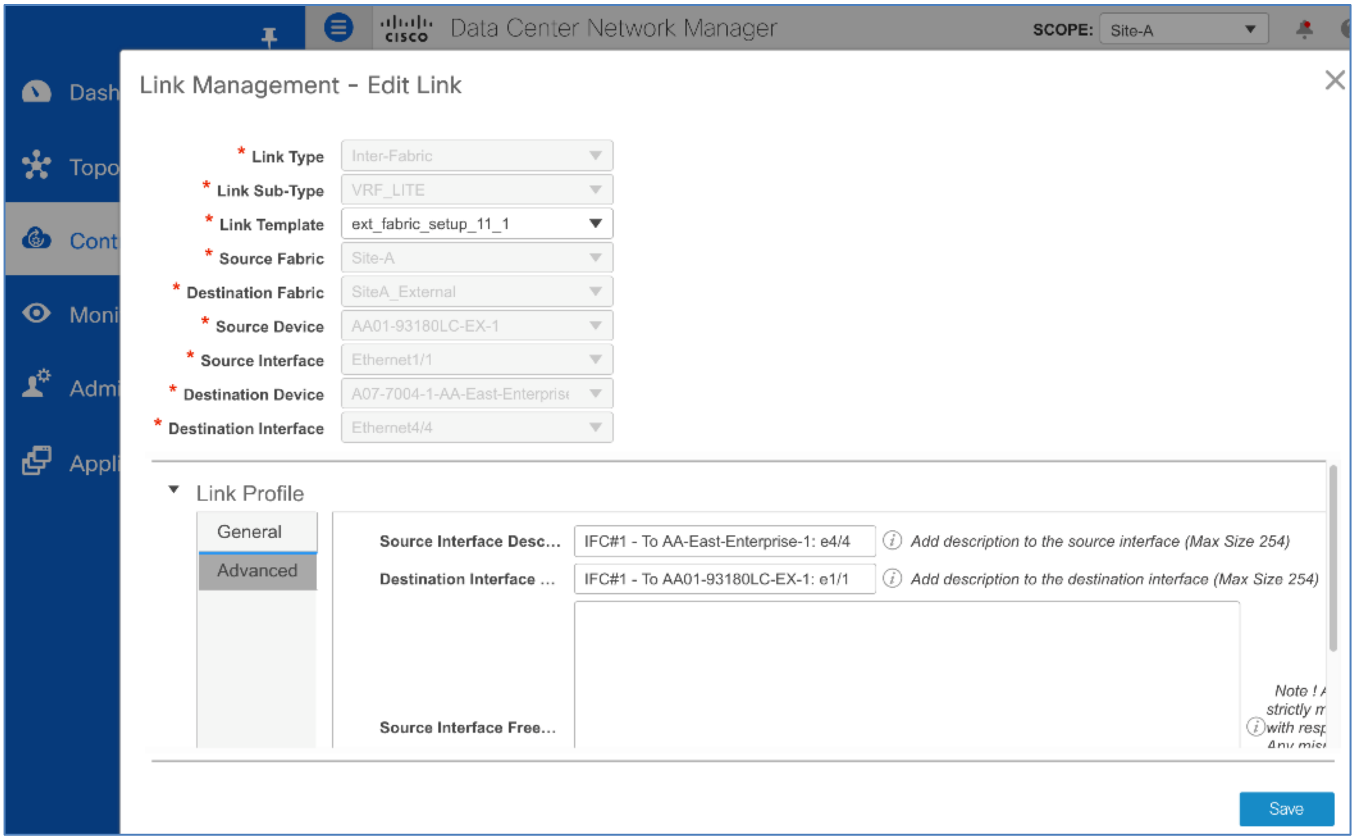

Link Management - Edit Link

* Link Type	Inter-Fabric
* Link Sub-Type	VRF_LITE
* Link Template	ext_fabric_setup_11_1
* Source Fabric	Site-A
* Destination Fabric	SiteA_External
* Source Device	AA01-93180LC-EX-1
* Source Interface	Ethernet1/1
* Destination Device	A07-7004-1-AA-East-Enterprise
* Destination Interface	Ethernet4/4

Link Profile

Source Interface Desc...	IFC#1 - To AA-East-Enterprise-1: e4/4	Add description to the source interface (Max Size 254)
Destination Interface ...	IFC#1 - To AA01-93180LC-EX-1: e1/1	Add description to the destination interface (Max Size 254)
Source Interface Free...		

12. Click **Save** to save the settings for the first IFC link.

13. Repeat steps 11 and 12 for the remaining 3 IFC links.

14. Go to **Switch** tab. The Border Leaf switches should be in **Pending** state**.** Click on the **Save & Deploy** button**.**

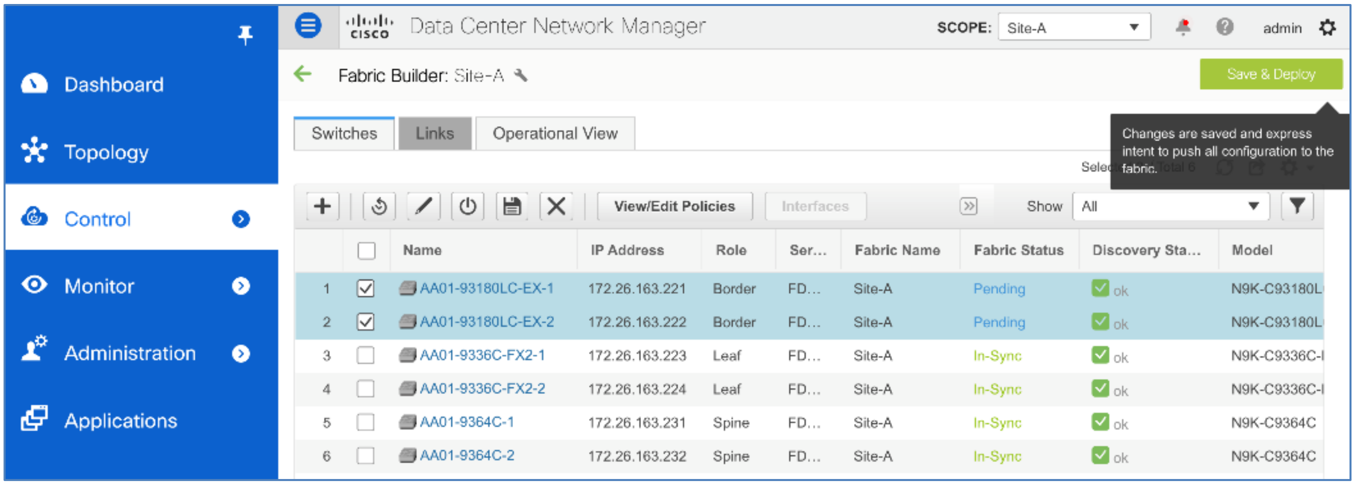

Fabric Builder: Site-A

Changes are saved and express intent to push all configuration to the fabric.

Show: All

		Name	IP Address	Role	Ser...	Fabric Name	Fabric Status	Discovery Sta...	Model
1	☑	AA01-93180LC-EX-1	172.26.163.221	Border	FD...	Site-A	Pending	✓ ok	N9K-C93180L
2	☑	AA01-93180LC-EX-2	172.26.163.222	Border	FD...	Site-A	Pending	✓ ok	N9K-C93180L
3	□	AA01-9336C-FX2-1	172.26.163.223	Leaf	FD...	Site-A	In-Sync	✓ ok	N9K-C9336C-I
4	□	AA01-9336C-FX2-2	172.26.163.224	Leaf	FD...	Site-A	In-Sync	✓ ok	N9K-C9336C-I
5	□	AA01-9364C-1	172.26.163.231	Spine	FD...	Site-A	In-Sync	✓ ok	N9K-C9364C
6	□	AA01-9364C-2	172.26.163.232	Spine	FD...	Site-A	In-Sync	✓ ok	N9K-C9364C

15. In the **Config Deployment** pop-up window, the Border switches should be **Out-of-Sync,** with about 11 lines of configuration change. Click on the '**11 lines'** to view the exact changes that will be deployed. Click the **Deploy Config** button.

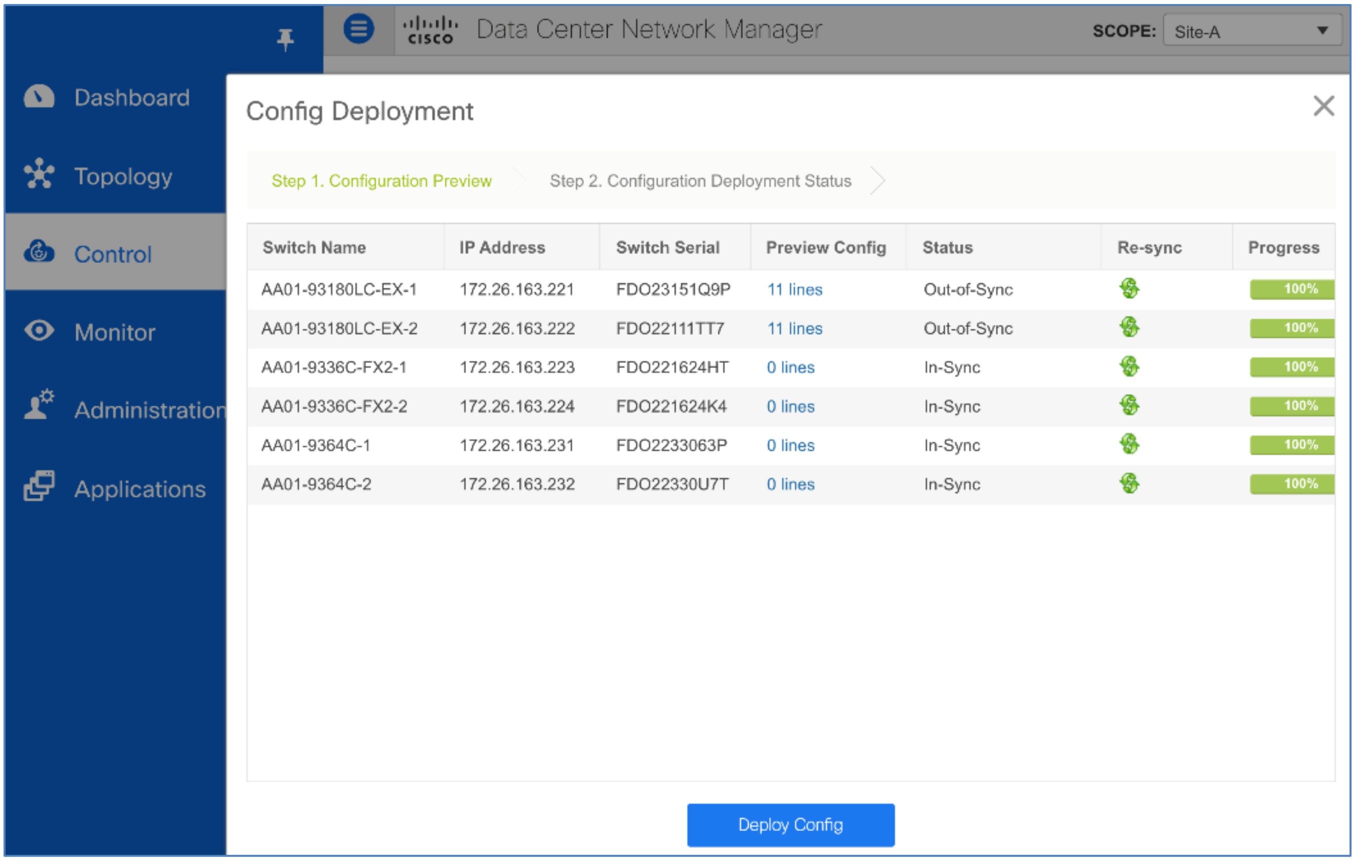

Config Deployment

Step 1. Configuration Preview | Step 2. Configuration Deployment Status

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
AA01-93180LC-EX-1	172.26.163.221	FDO23151Q9P	11 lines	Out-of-Sync		100%
AA01-93180LC-EX-2	172.26.163.222	FDO22111TT7	11 lines	Out-of-Sync		100%
AA01-9336C-FX2-1	172.26.163.223	FDO221624HT	0 lines	In-Sync		100%
AA01-9336C-FX2-2	172.26.163.224	FDO221624K4	0 lines	In-Sync		100%
AA01-9364C-1	172.26.163.231	FDO2233063P	0 lines	In-Sync		100%
AA01-9364C-2	172.26.163.232	FDO22330U7T	0 lines	In-Sync		100%

16. When the deployment completes, click the **Close** button to close the window. The switches should now be in an **In-Sync** state.

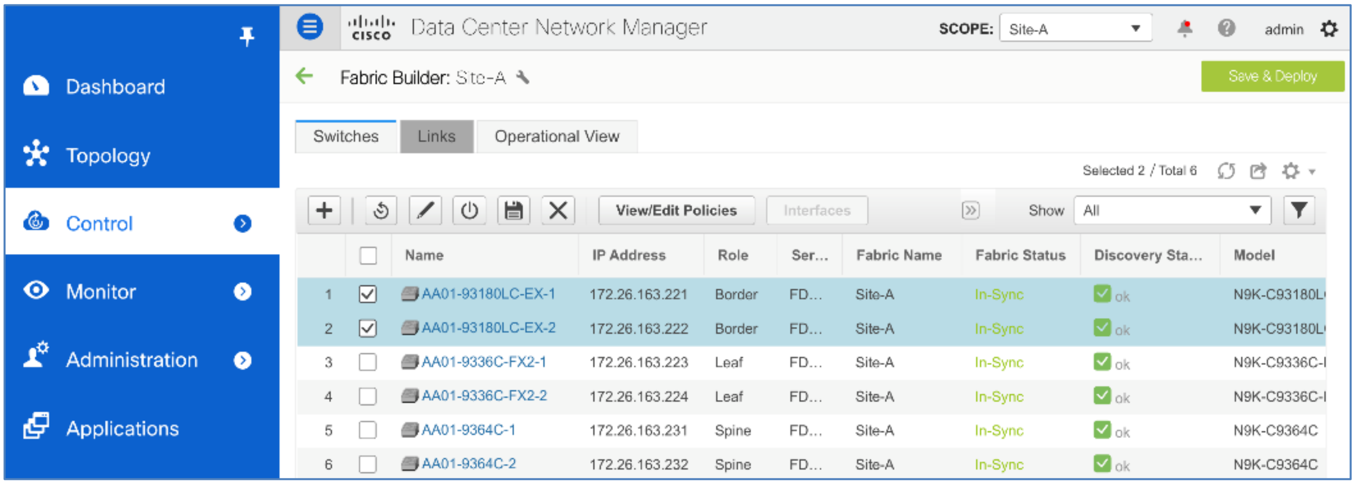

Fabric Builder: Site-A

Selected 2 / Total 6

		Name	IP Address	Role	Ser...	Fabric Name	Fabric Status	Discovery Sta...	Model
1	☑	AA01-93180LC-EX-1	172.26.163.221	Border	FD...	Site-A	In-Sync	✓ ok	N9K-C93180L
2	☑	AA01-93180LC-EX-2	172.26.163.222	Border	FD...	Site-A	In-Sync	✓ ok	N9K-C93180L
3	□	AA01-9336C-FX2-1	172.26.163.223	Leaf	FD...	Site-A	In-Sync	✓ ok	N9K-C9336C-I
4	□	AA01-9336C-FX2-2	172.26.163.224	Leaf	FD...	Site-A	In-Sync	✓ ok	N9K-C9336C-I
5	□	AA01-9364C-1	172.26.163.231	Spine	FD...	Site-A	In-Sync	✓ ok	N9K-C9364C
6	□	AA01-9364C-2	172.26.163.232	Spine	FD...	Site-A	In-Sync	✓ ok	N9K-C9364C

17. From the top-right corner of the window, for **Scope:** , select **SiteA_External** from the drop-down list.

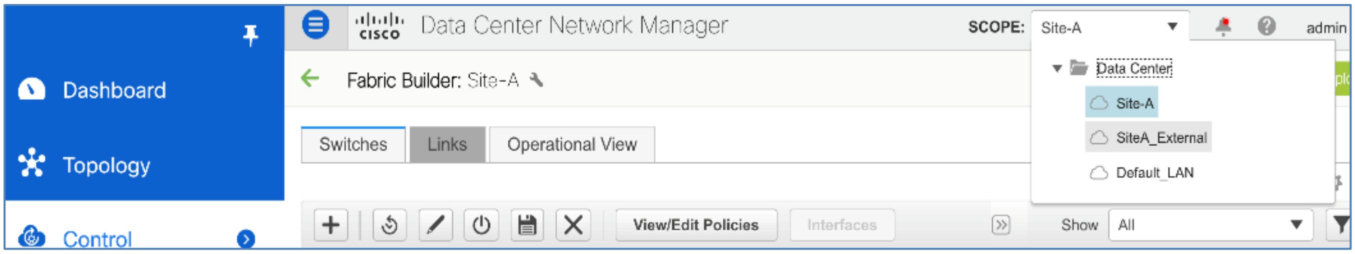

SCOPE: Site-A

18. Note that the external gateways are in a **Pending** state. Click the **Save & Deploy** button.

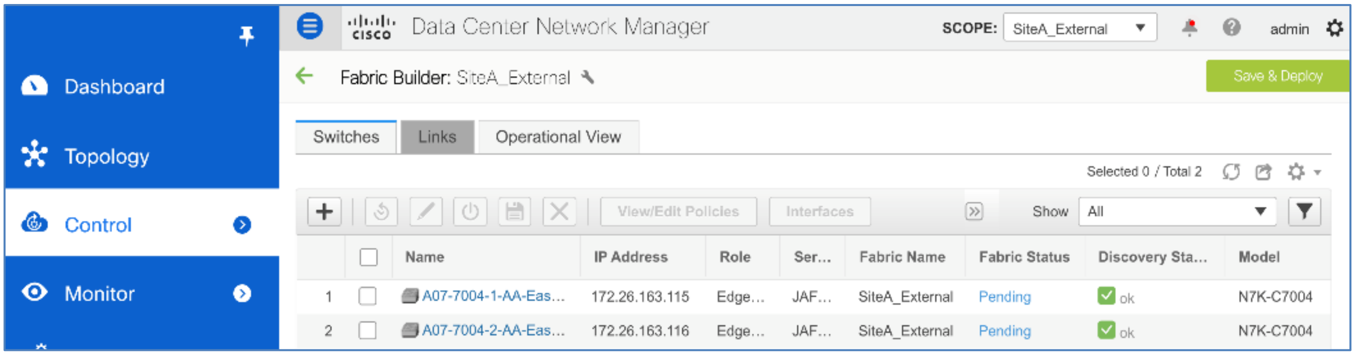

SCOPE: SiteA_External

Selected 0 / Total 2

	□	Name	IP Address	Role	Ser...	Fabric Name	Fabric Status	Discovery Sta...	Model
1	□	A07-7004-1-AA-Eas...	172.26.163.115	Edge...	JAF...	SiteA_External	Pending	✓ ok	N7K-C7004
2	□	A07-7004-2-AA-Eas...	172.26.163.116	Edge...	JAF...	SiteA_External	Pending	✓ ok	N7K-C7004

19. In the **Config Deployment** pop-up window, the External Gateway switches should be **Out-of-Sync,** with about 11 lines of configuration change. Click on the '**11 lines'** to view the exact changes that will be deployed. Click the **Deploy Config** button.

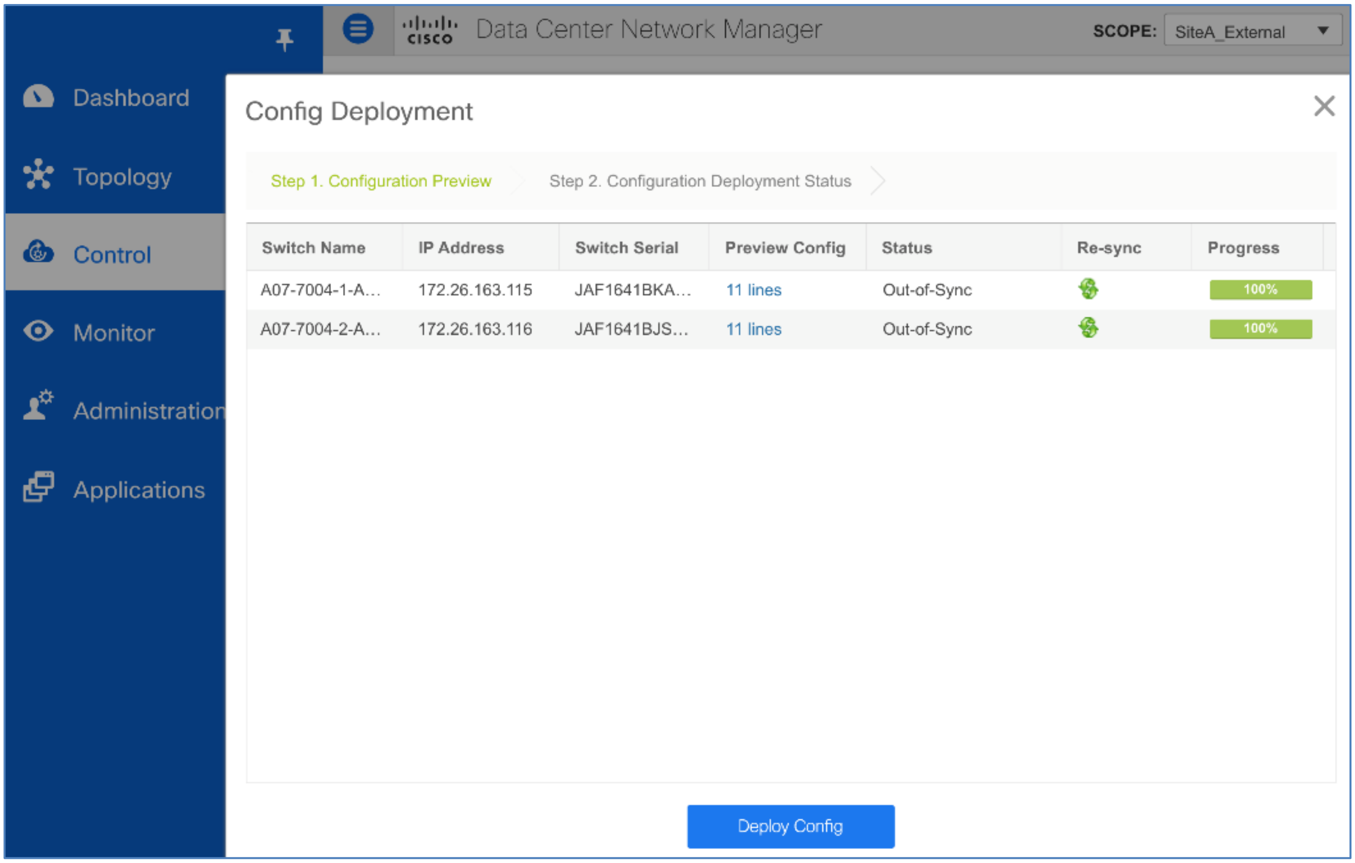

Config Deployment

Step 1. Configuration Preview | Step 2. Configuration Deployment Status

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
A07-7004-1-A…	172.26.163.115	JAF1641BKA…	11 lines	Out-of-Sync	↻	100%
A07-7004-2-A…	172.26.163.116	JAF1641BJS…	11 lines	Out-of-Sync	↻	100%

Deploy Config

20. When the deployment completes, click the **Close** button to close the window. The switches should now be in an **In-Sync** state.

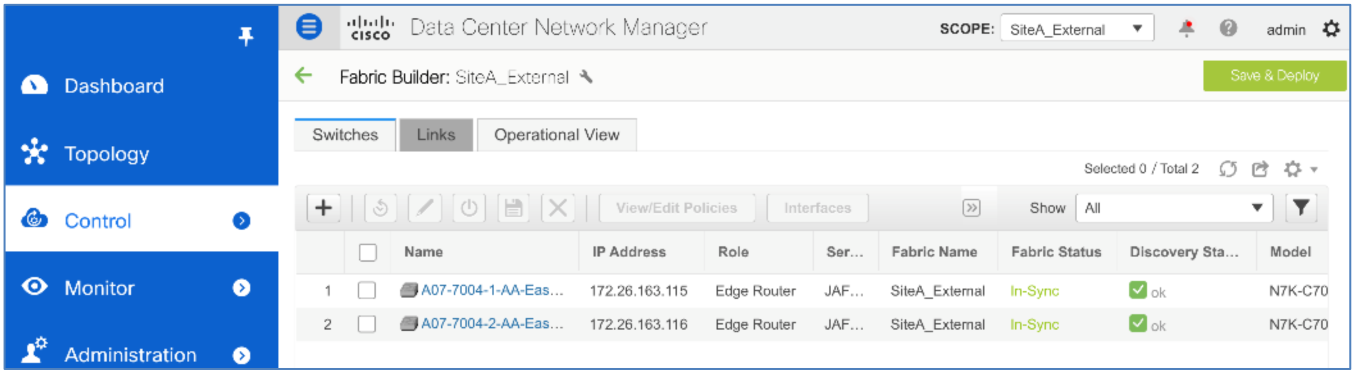

Fabric Builder: SiteA_External

Selected 0 / Total 2

	☐	Name	IP Address	Role	Ser…	Fabric Name	Fabric Status	Discovery Sta…	Model
1	☐	A07-7004-1-AA-Eas…	172.26.163.115	Edge Router	JAF…	SiteA_External	In-Sync	✓ ok	N7K-C70
2	☐	A07-7004-2-AA-Eas…	172.26.163.116	Edge Router	JAF…	SiteA_External	In-Sync	✓ ok	N7K-C70

You are now ready to deploy the infrastructure networks that FlexPod requires.

Enable Access Layer Connectivity to Cisco UCS Domain

In this section, access-layer connectivity is enabled from the VXLAN fabric in Site-A to the Cisco UCS infrastructure used in the FlexPod solution. The Cisco UCS infrastructure consists of a pair of Cisco UCS Fabric Interconnects that connects to Cisco UCS B-series and C-series servers.

Topology

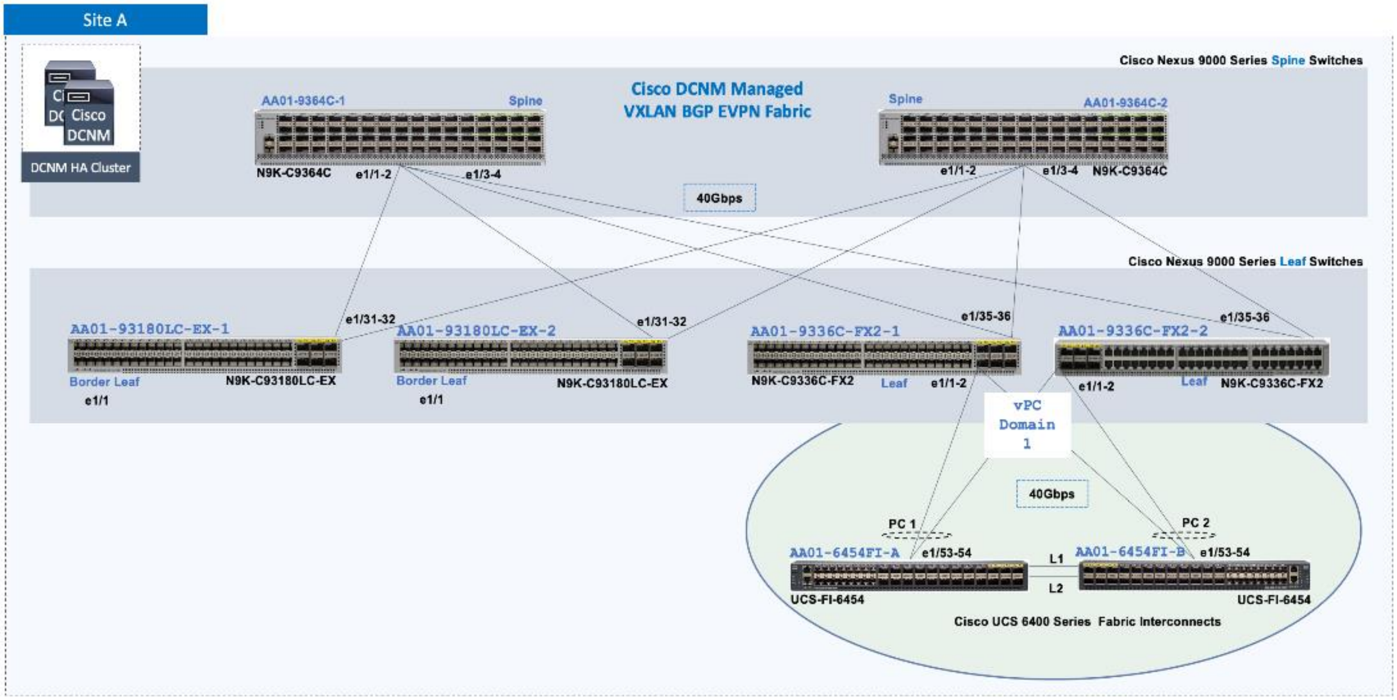

Setup Information

The configuration parameters for enabling access-layer connectivity to Cisco UCS Domain in Site-A data center fabric is provided below.

Table 8. Access Layer Switches – To Cisco UCS Domain

Hostname	Switch Role	IP Address (OOB)	Notes
AA01-9336C-FX2-1	Leaf	172.26.163.223/24	
AA01-9336C-FX2-2	Leaf	172.26.163.224/24	

Table 9. Access Layer Connectivity – To Cisco UCS Domain

Access Layer Connection	Parameters	Notes
Type	Virtual Port-Channel (vPC)	Using Virtual Peer-Links (requires hardware support)
vPC Pair	AA01-9336C-FX2-1---AA01-9336C-FX2-2	
vPC to Cisco UCS FI-A		
Peer-1 Member Interfaces	Ethernet 1/1	

Access Layer Connection	Parameters	Notes
Peer-2 Member Interfaces	Ethernet 1/1	
Peer-1 PO Description	To FXV-AA01-UCS6454FI-A: e1/53	
Peer-2 PO Description	To FXV-AA01-UCS6454FI-A: e1/54	
vPC to Cisco UCS FI-B		
Peer-1 Member Interfaces	Ethernet 1/2	
Peer-2 Member Interfaces	Ethernet 1/2	
Peer-1 PO Description	To FXV-AA01-UCS6454FI-B: e1/53	
Peer-2 PO Description	To FXV-AA01-UCS6454FI-B: e1/54	

Deployment Steps

To enable access-layer connectivity from Site-A data center fabric to the Cisco UCS domain, follow these steps using the **Setup Information** provided above:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Fabric Builder**. Click on the **Site-A** fabric.

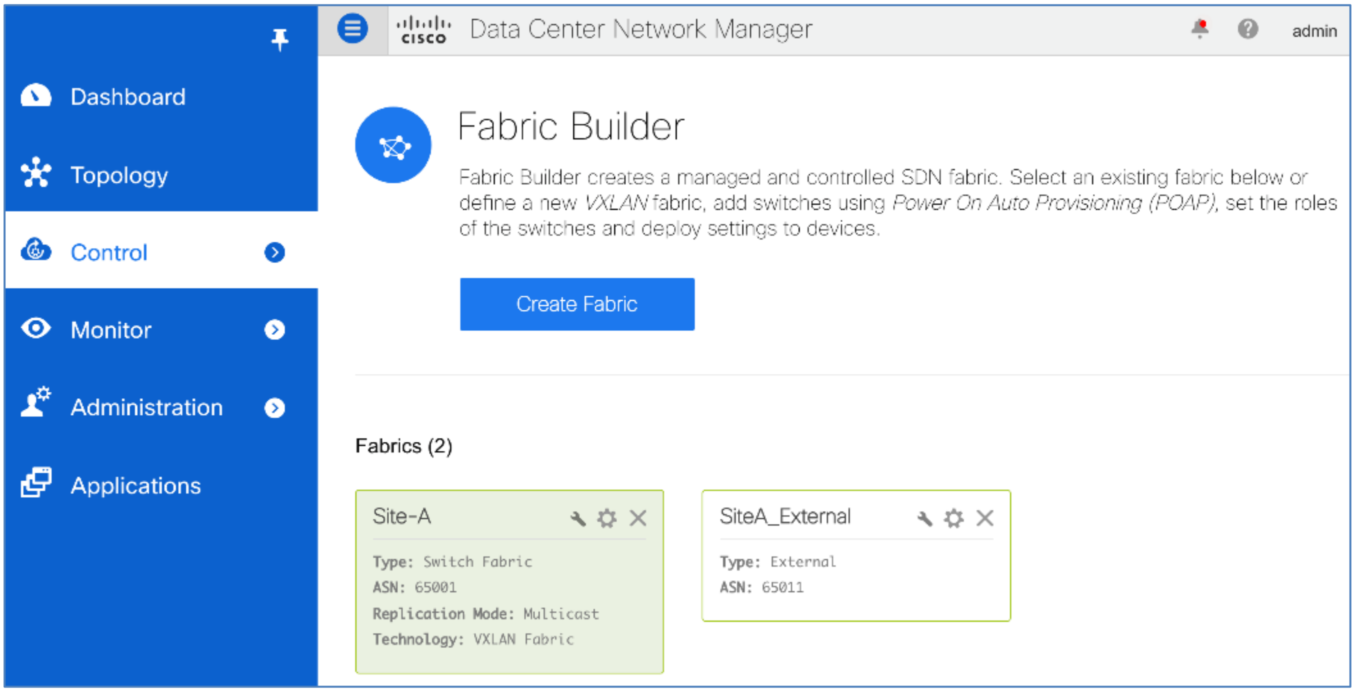

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (2)

Site-A
Type: Switch Fabric
ASN: 65001
Replication Mode: Multicast
Technology: VXLAN Fabric

SiteA_External
Type: External
ASN: 65011

3. From the right window pane, select one of the Leaf switches that connect to the Cisco UCS domain.

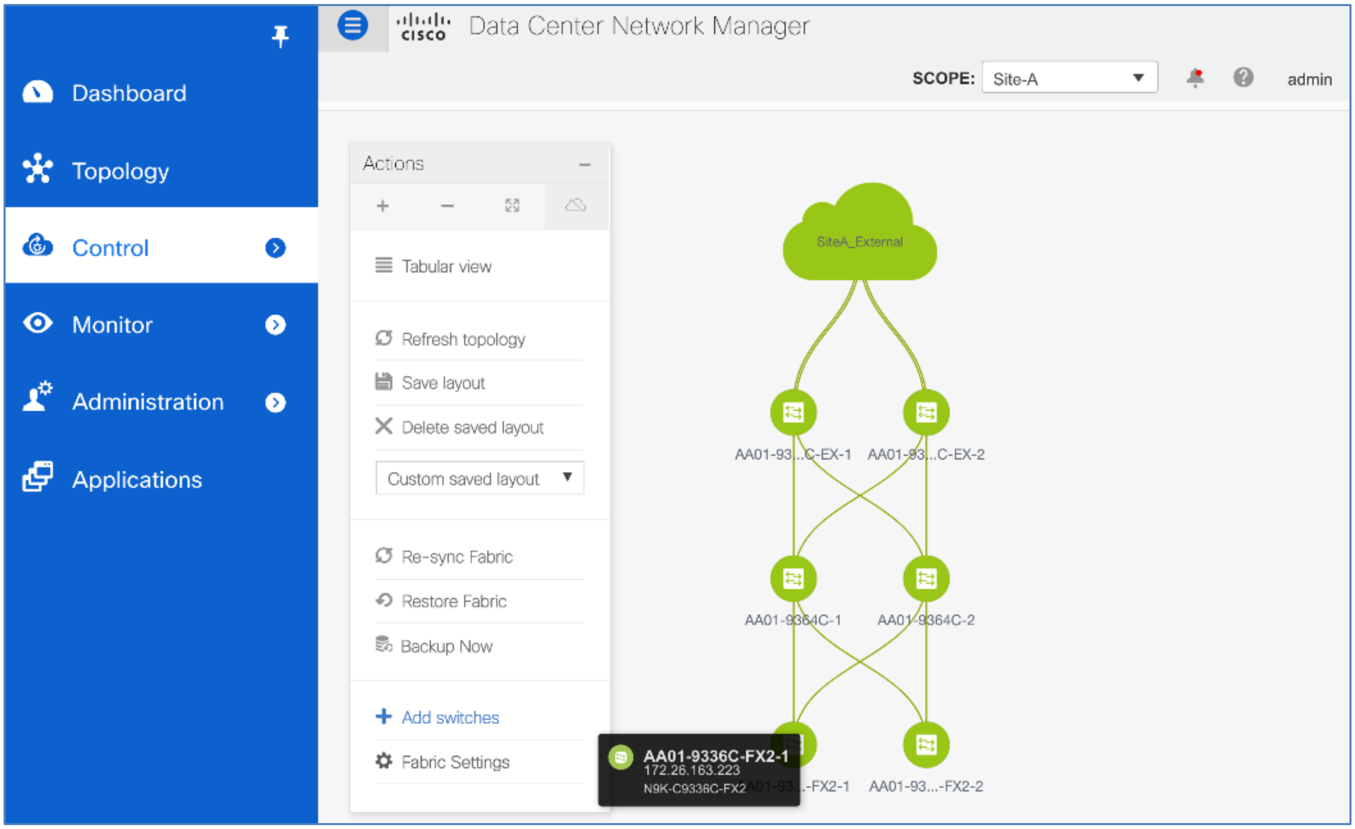

4. Right-click and select **vPC Pairing** from the list.

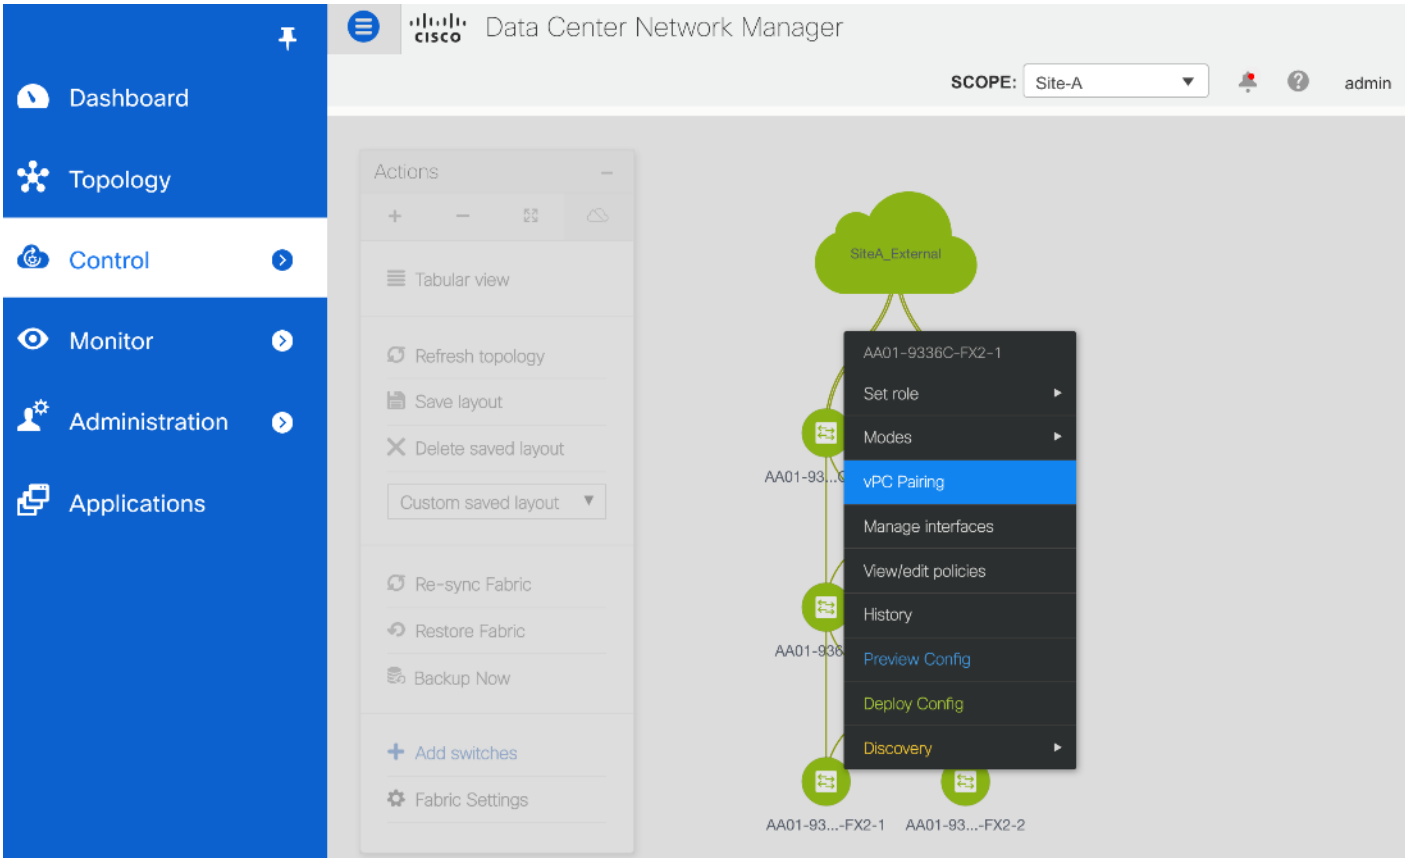

5. In the vPC peer for Leaf switch pop-up window, enable the checkbox next to **Use Virtual Peer link** and select the radio button for peer Leaf switch that will be part of the vPC pair for the vPC going to the Cisco UCS Fabric Interconnects in the UCS domain.

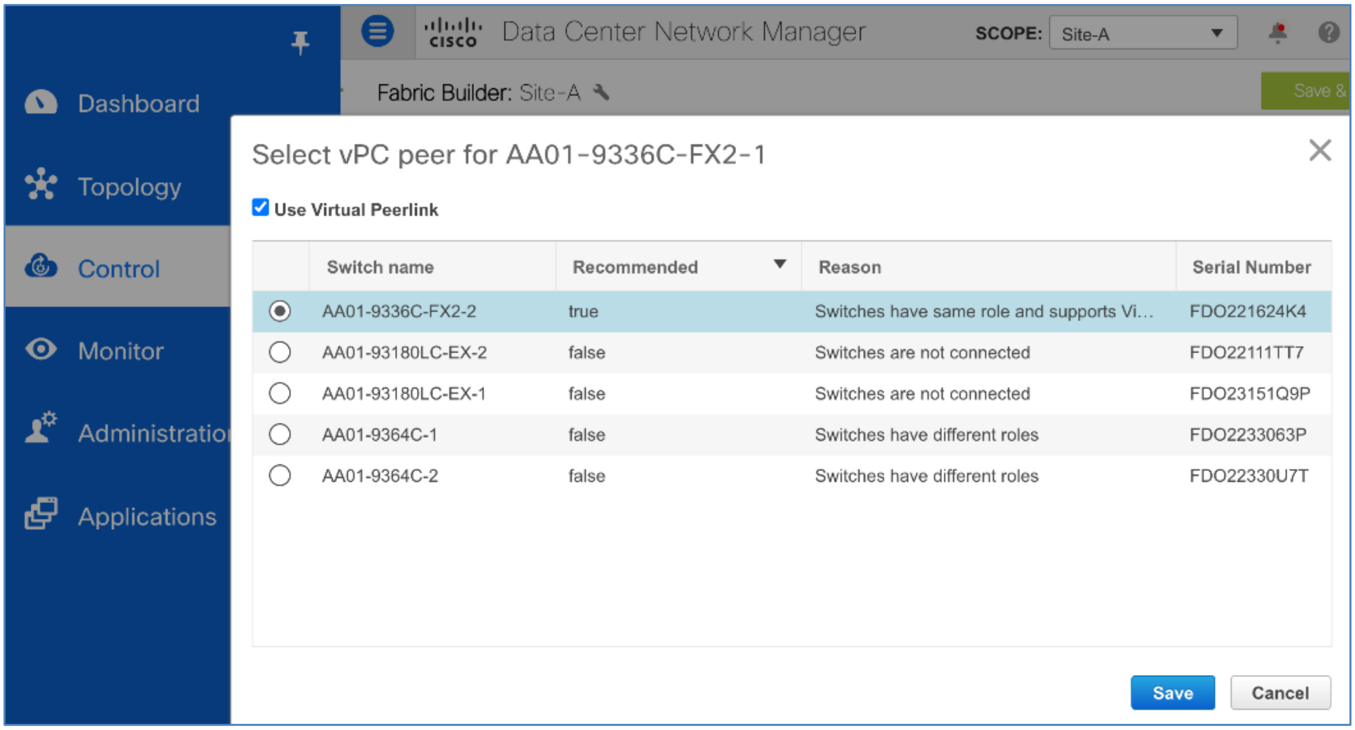

Select vPC peer for AA01-9336C-FX2-1

☑ Use Virtual Peerlink

	Switch name	Recommended	Reason	Serial Number
●	AA01-9336C-FX2-2	true	Switches have same role and supports Vi...	FDO221624K4
○	AA01-93180LC-EX-2	false	Switches are not connected	FDO22111TT7
○	AA01-93180LC-EX-1	false	Switches are not connected	FDO23151Q9P
○	AA01-9364C-1	false	Switches have different roles	FDO2233063P
○	AA01-9364C-2	false	Switches have different roles	FDO22330U7T

6. Click **Save**. A small pop-up window will show up on the right-bottom corner of the window to indicate whether the vPC pairing is successful. Note that the two leaf switches in the vPC pair are now grouped together in the topology view. Click the **Save & Deploy** button to deploy the vPC pairing.

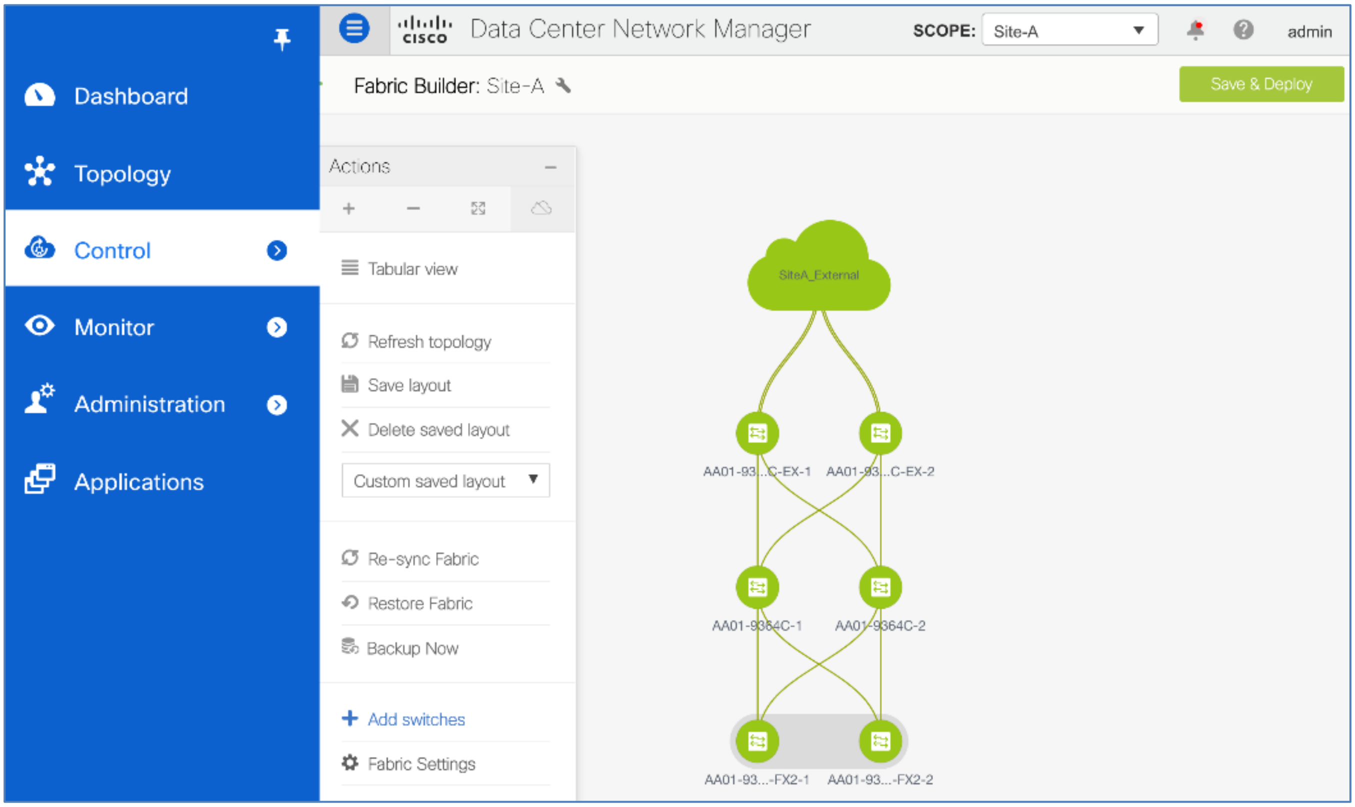

7. In the **Config Deployment** pop-up window, note that the leaf switches are **Out-of-Sync**, with **56 lines** of configuration to be deployed.

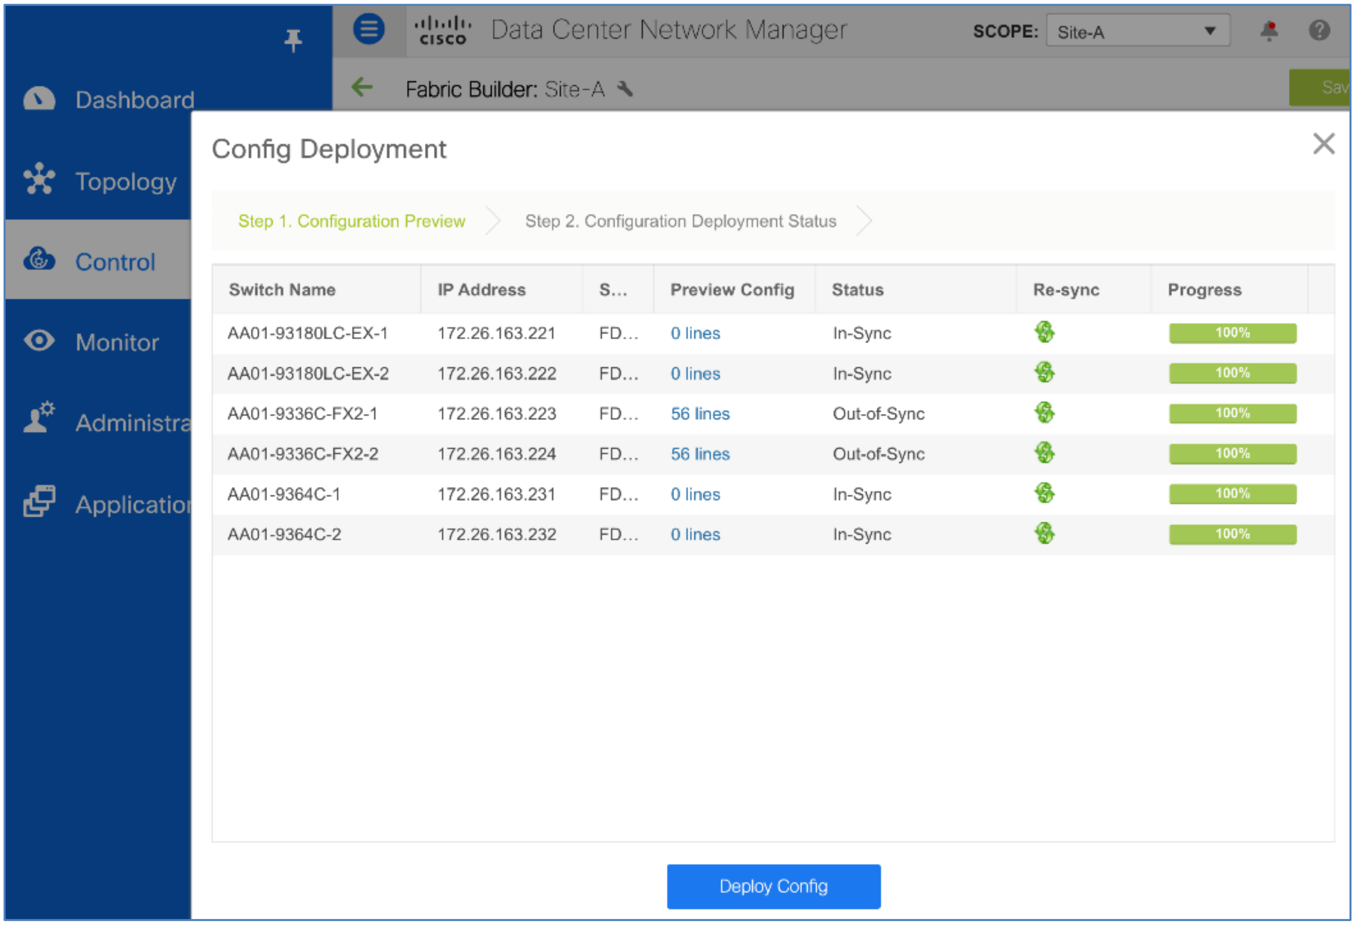

Config Deployment

Step 1. Configuration Preview | Step 2. Configuration Deployment Status

Switch Name	IP Address	S...	Preview Config	Status	Re-sync	Progress
AA01-93180LC-EX-1	172.26.163.221	FD...	0 lines	In-Sync		100%
AA01-93180LC-EX-2	172.26.163.222	FD...	0 lines	In-Sync		100%
AA01-9336C-FX2-1	172.26.163.223	FD...	56 lines	Out-of-Sync		100%
AA01-9336C-FX2-2	172.26.163.224	FD...	56 lines	Out-of-Sync		100%
AA01-9364C-1	172.26.163.231	FD...	0 lines	In-Sync		100%
AA01-9364C-2	172.26.163.232	FD...	0 lines	In-Sync		100%

Deploy Config

8. Click on the **56 lines** for one of the switches to preview the pending configuration on that switch.

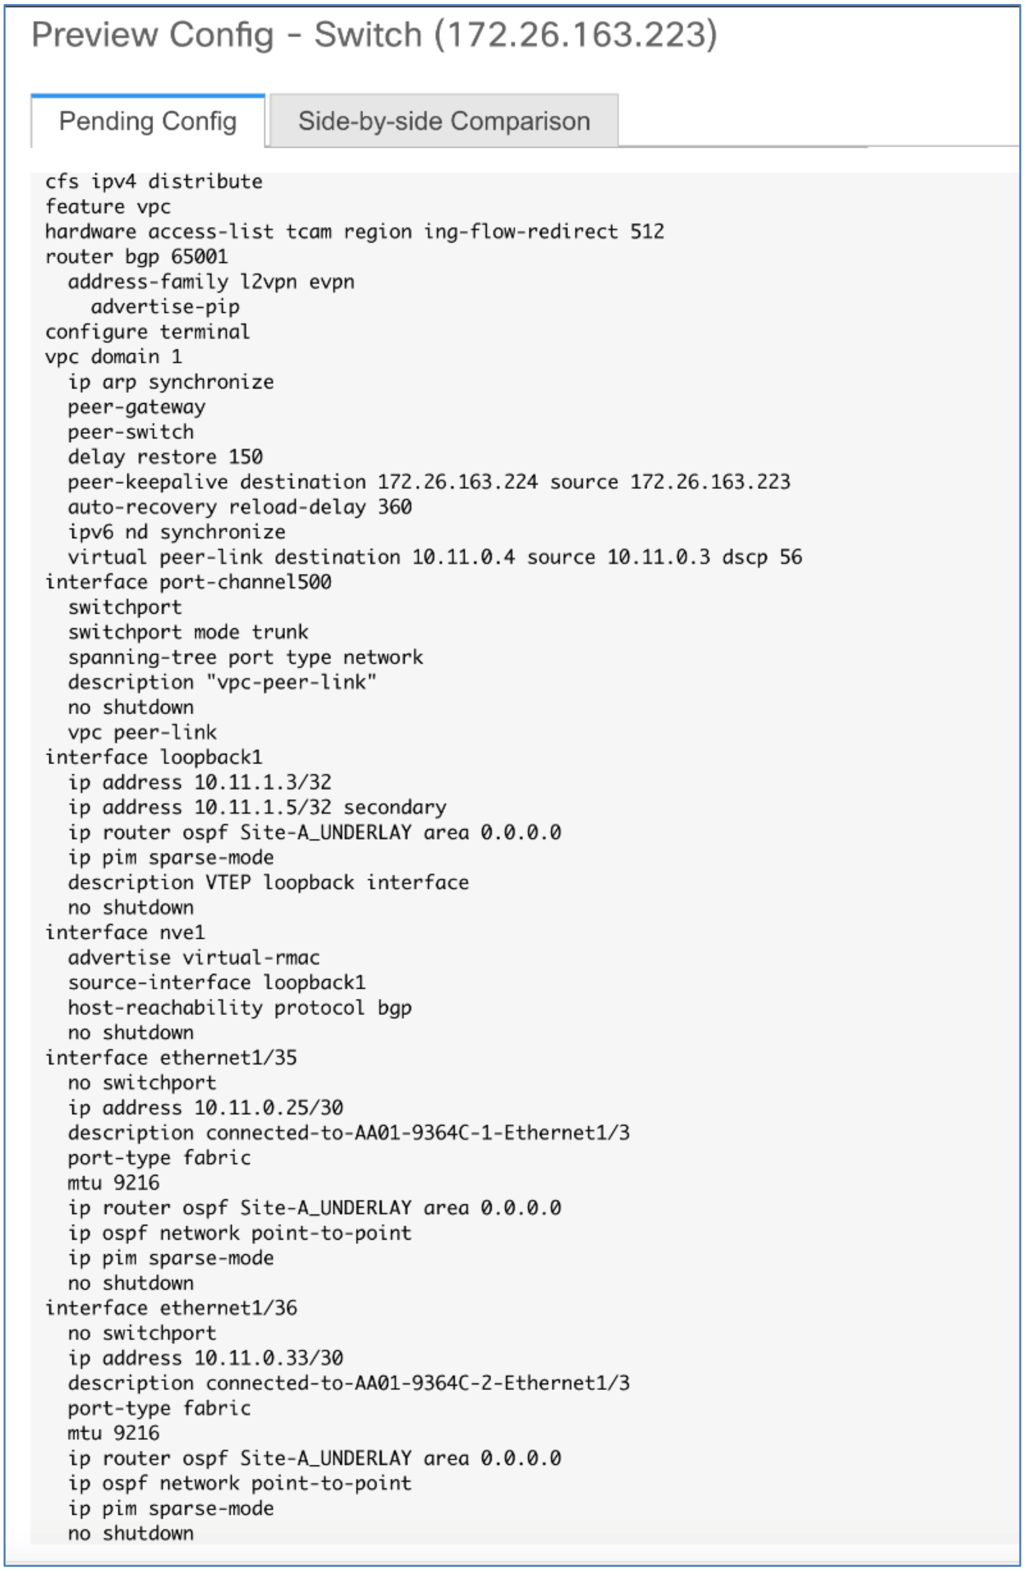

```
cfs ipv4 distribute
feature vpc
hardware access-list tcam region ing-flow-redirect 512
router bgp 65001
  address-family l2vpn evpn
    advertise-pip
configure terminal
vpc domain 1
  ip arp synchronize
  peer-gateway
  peer-switch
  delay restore 150
  peer-keepalive destination 172.26.163.224 source 172.26.163.223
  auto-recovery reload-delay 360
  ipv6 nd synchronize
  virtual peer-link destination 10.11.0.4 source 10.11.0.3 dscp 56
interface port-channel500
  switchport
  switchport mode trunk
  spanning-tree port type network
  description "vpc-peer-link"
  no shutdown
  vpc peer-link
interface loopback1
  ip address 10.11.1.3/32
  ip address 10.11.1.5/32 secondary
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  description VTEP loopback interface
  no shutdown
interface nve1
  advertise virtual-rmac
  source-interface loopback1
  host-reachability protocol bgp
  no shutdown
interface ethernet1/35
  no switchport
  ip address 10.11.0.25/30
  description connected-to-AA01-9364C-1-Ethernet1/3
  port-type fabric
  mtu 9216
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip ospf network point-to-point
  ip pim sparse-mode
  no shutdown
interface ethernet1/36
  no switchport
  ip address 10.11.0.33/30
  description connected-to-AA01-9364C-2-Ethernet1/3
  port-type fabric
  mtu 9216
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip ospf network point-to-point
  ip pim sparse-mode
  no shutdown
```

9. Click the **X** to close the preview window. Click on **Deploy Config** button to deploy the configuration.

10. When the deployment completes successfully with a **COMPLETED** status, click on the **Close** button to close the window. Now you can start configuring access layer connectivity to Cisco UCS domain.

11. Select one of the Leaf switches that connect to the Cisco UCS domain. Right-click and select **Manage Interfaces** from the list. In the **Manage Interfaces** pop-up window, bring the **Neighbors** column into view by dragging it from the far-right end and move it next to the **Reason** column. Note the interfaces on the leaf switches in the vPC pair that connect to the first UCS FI and the port numbers they connect to.

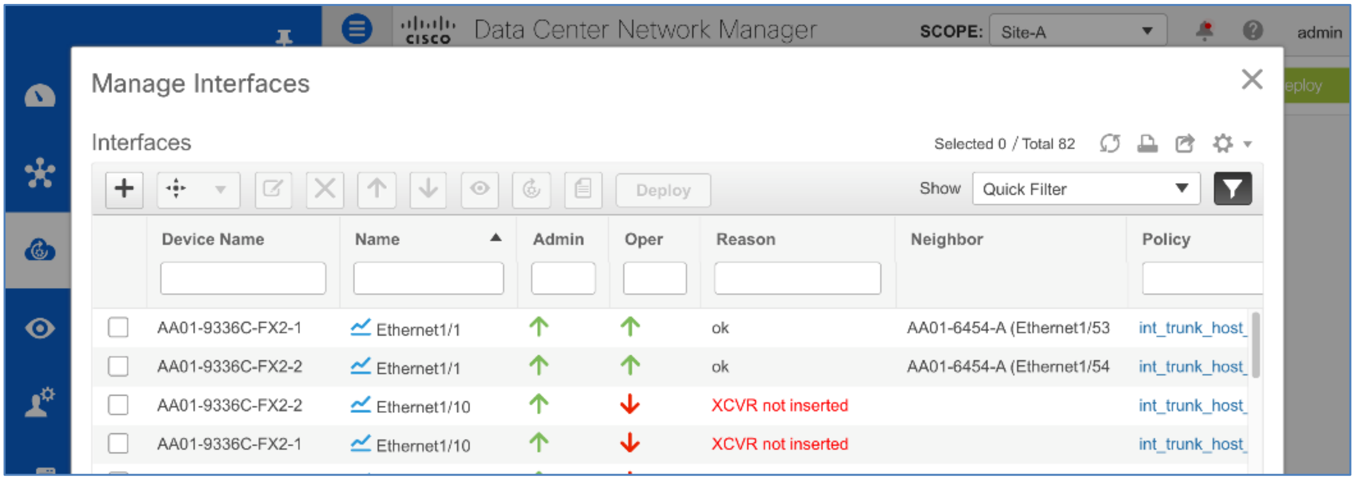

	Device Name	Name	Admin	Oper	Reason	Neighbor	Policy
□	AA01-9336C-FX2-1	Ethernet1/1	↑	↑	ok	AA01-6454-A (Ethernet1/53	int_trunk_host_
□	AA01-9336C-FX2-2	Ethernet1/1	↑	↑	ok	AA01-6454-A (Ethernet1/54	int_trunk_host_
□	AA01-9336C-FX2-2	Ethernet1/10	↑	↓	XCVR not inserted		int_trunk_host_
□	AA01-9336C-FX2-1	Ethernet1/10	↑	↓	XCVR not inserted		int_trunk_host_

12. Click the **[+]** button from the menu above. In the **Add Interface** pop-up window, specify the **Type** from the drop-down list.

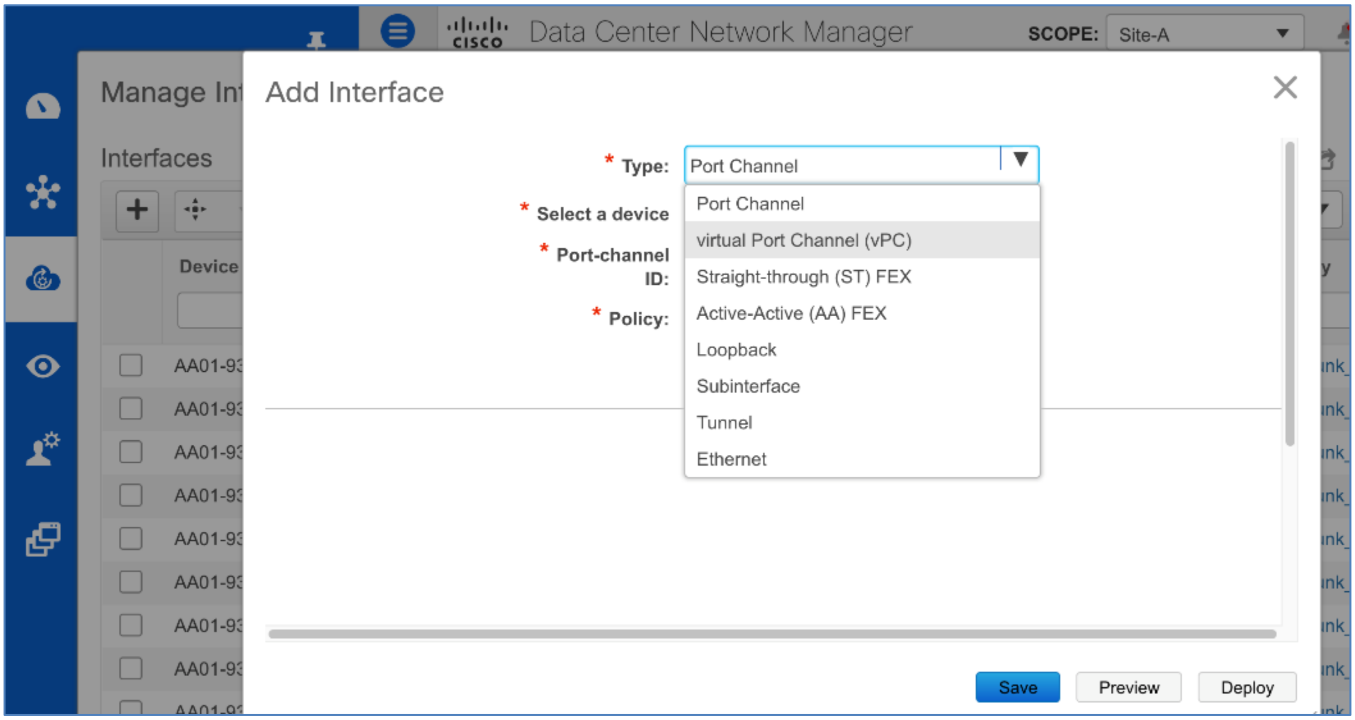

* Type: Port Channel

* Select a device

* Port-channel ID:

* Policy:

13. The menu changes to the reflect the options for the Interface **Type** selected. **Select a vPC pair** from the drop-down list.

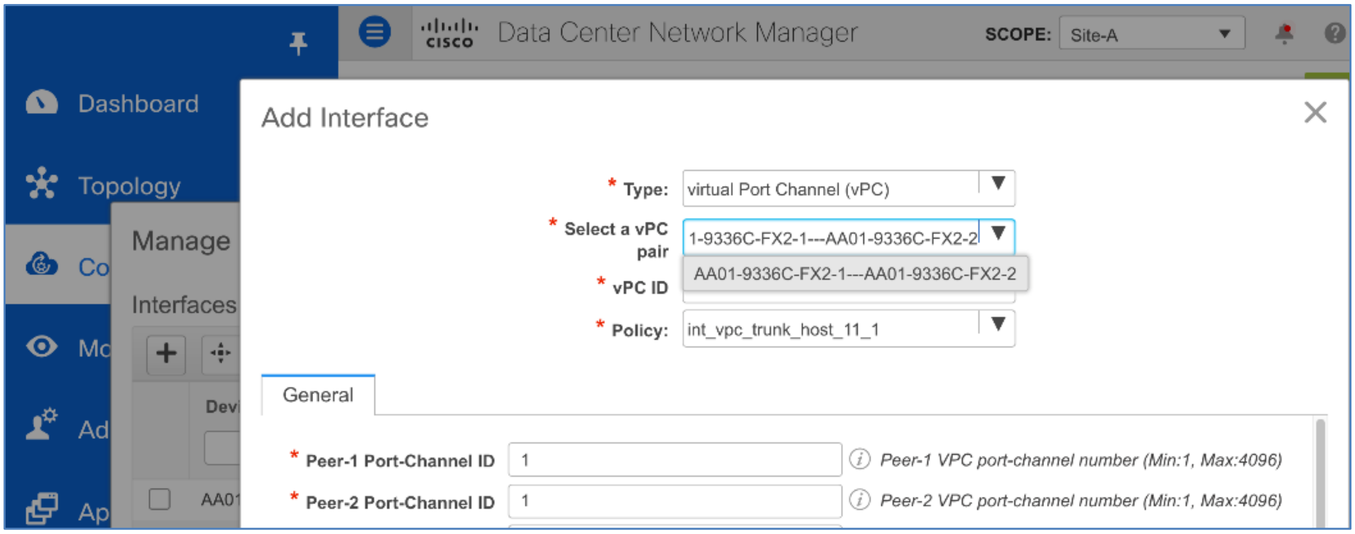

* Type: virtual Port Channel (vPC)

* Select a vPC pair: 1-9336C-FX2-1---AA01-9336C-FX2-2

AA01-9336C-FX2-1---AA01-9336C-FX2-2

* vPC ID

* Policy: int_vpc_trunk_host_11_1

* Peer-1 Port-Channel ID: 1 — *Peer-1 VPC port-channel number (Min:1, Max:4096)*

* Peer-2 Port-Channel ID: 1 — *Peer-2 VPC port-channel number (Min:1, Max:4096)*

14. Specify the Peer-1 Member Interfaces, Peer-2 Member Interfaces, Peer-1 PO Description, and Peer-2 PO Description. Leave all other fields as-is.

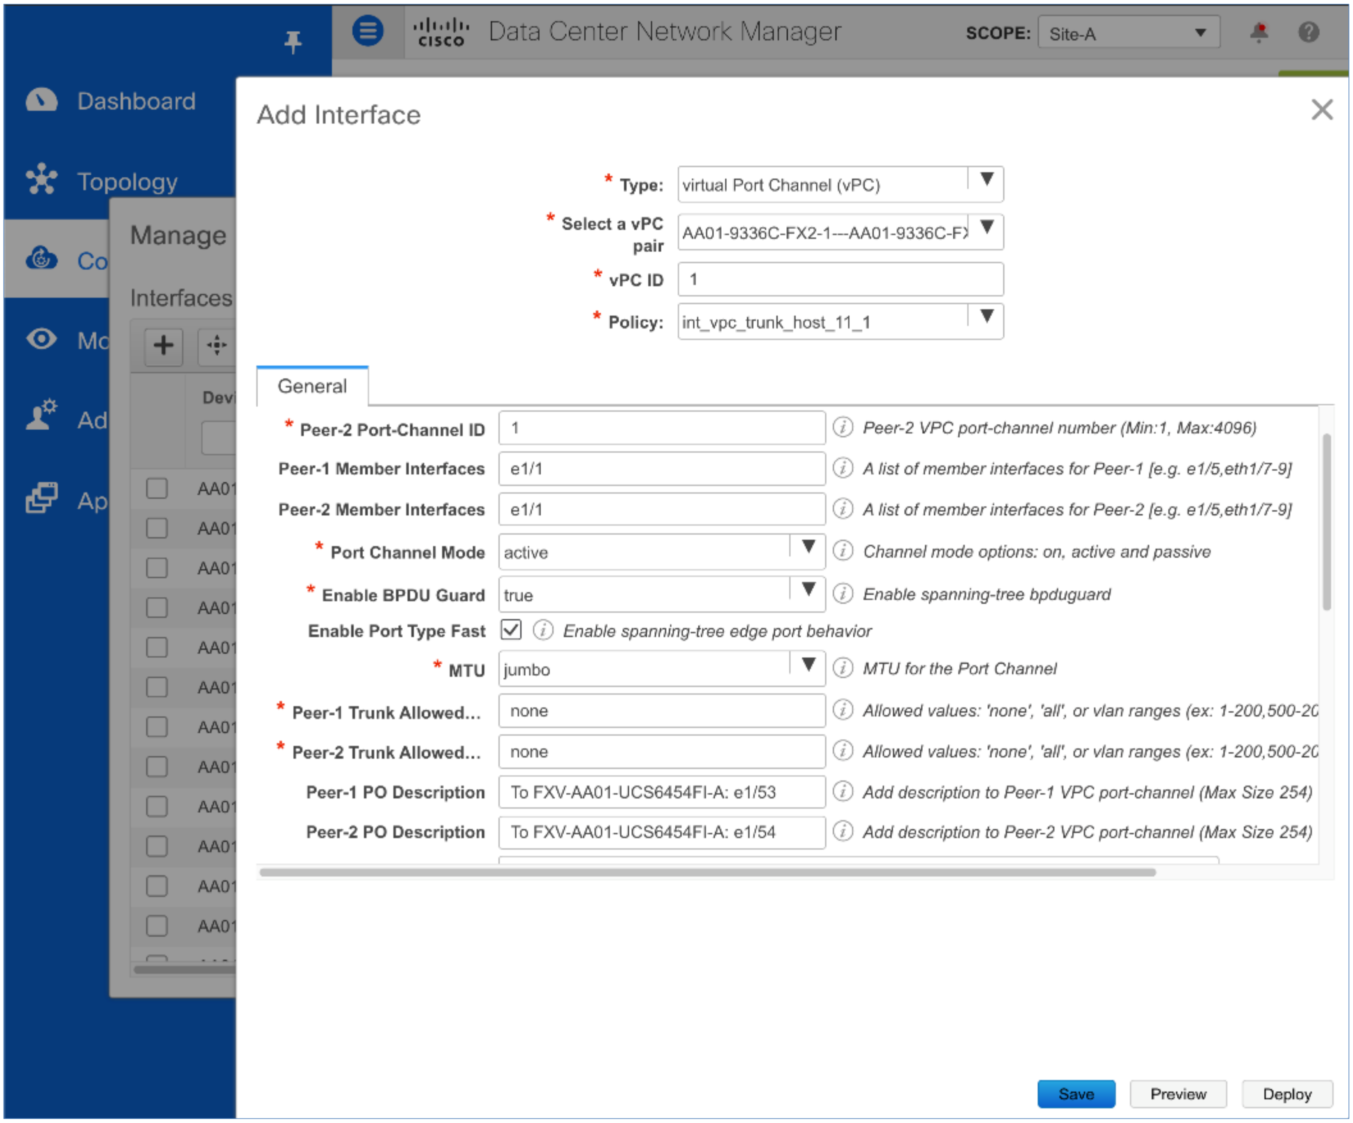

Add Interface

* Type:	virtual Port Channel (vPC)
* Select a vPC pair	AA01-9336C-FX2-1---AA01-9336C-F)
* vPC ID	1
* Policy:	int_vpc_trunk_host_11_1

General

* Peer-2 Port-Channel ID	1	Peer-2 VPC port-channel number (Min:1, Max:4096)
Peer-1 Member Interfaces	e1/1	A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]
Peer-2 Member Interfaces	e1/1	A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]
* Port Channel Mode	active	Channel mode options: on, active and passive
* Enable BPDU Guard	true	Enable spanning-tree bpduguard
Enable Port Type Fast	☑	Enable spanning-tree edge port behavior
* MTU	jumbo	MTU for the Port Channel
* Peer-1 Trunk Allowed...	none	Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-20
* Peer-2 Trunk Allowed...	none	Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-20
Peer-1 PO Description	To FXV-AA01-UCS6454FI-A: e1/53	Add description to Peer-1 VPC port-channel (Max Size 254)
Peer-2 PO Description	To FXV-AA01-UCS6454FI-A: e1/54	Add description to Peer-2 VPC port-channel (Max Size 254)

15. Click **Save.** You can also **Preview** configuration for the vPC to Cisco UCS FI-A using the **Preview** button. The preview will display the pending configuration for each switch in the vPC pair – use the drop-down list to select the second switch.

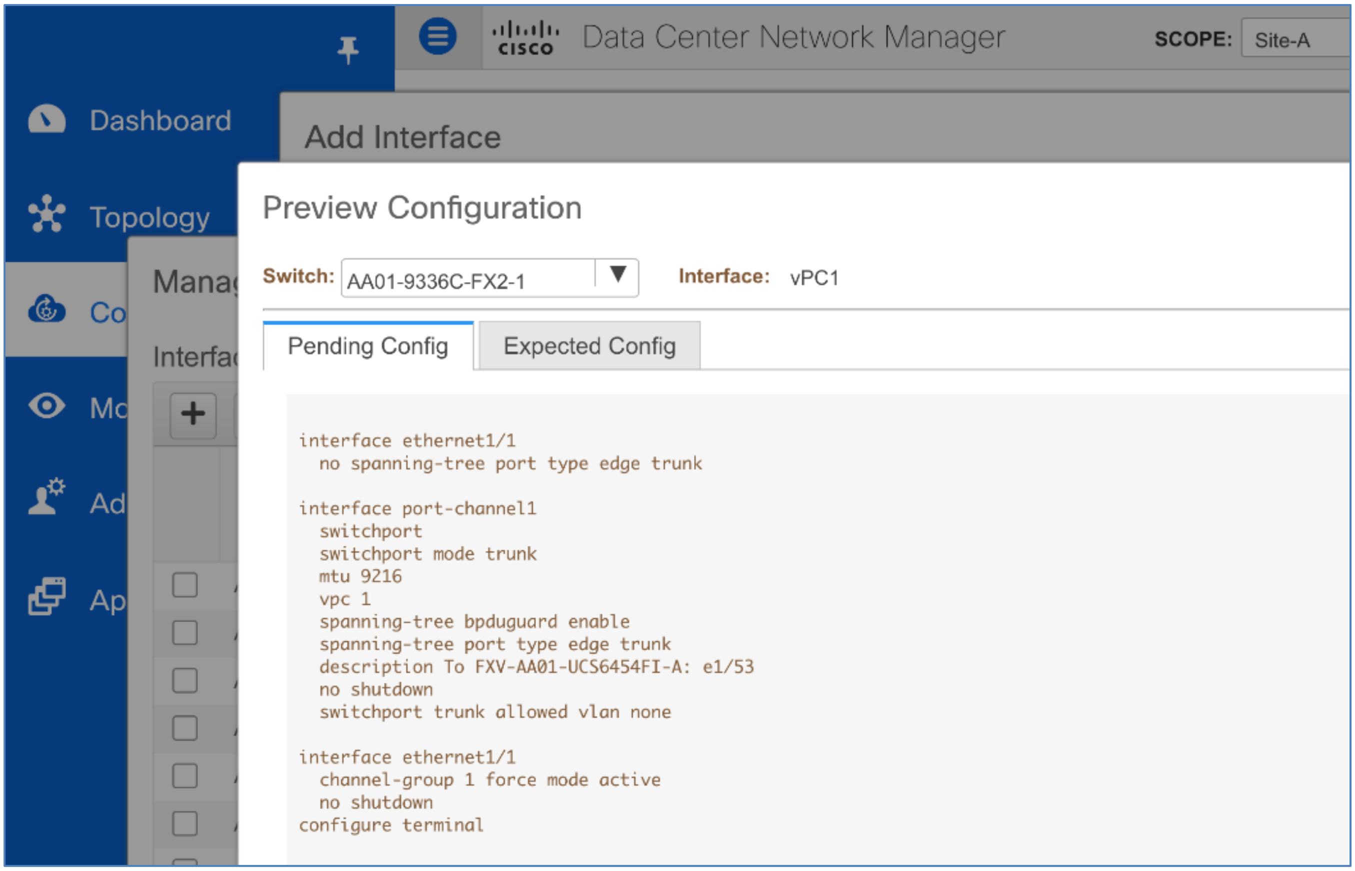

Preview Configuration

Switch: AA01-9336C-FX2-1 Interface: vPC1

Pending Config | Expected Config

```
interface ethernet1/1
  no spanning-tree port type edge trunk

interface port-channel1
  switchport
  switchport mode trunk
  mtu 9216
  vpc 1
  spanning-tree bpduguard enable
  spanning-tree port type edge trunk
  description To FXV-AA01-UCS6454FI-A: e1/53
  no shutdown
  switchport trunk allowed vlan none

interface ethernet1/1
  channel-group 1 force mode active
  no shutdown
configure terminal
```

16. Click **Deploy** to deploy the vPC configuration from the Leaf switch pair in the VXLAN fabric to the first Cisco UCS Fabric Interconnect (FI-A). Click **OK** in the pop-up window.

17. Repeat steps 1-16 to create, preview, and deploy the vPC to Cisco UCS FI-B.

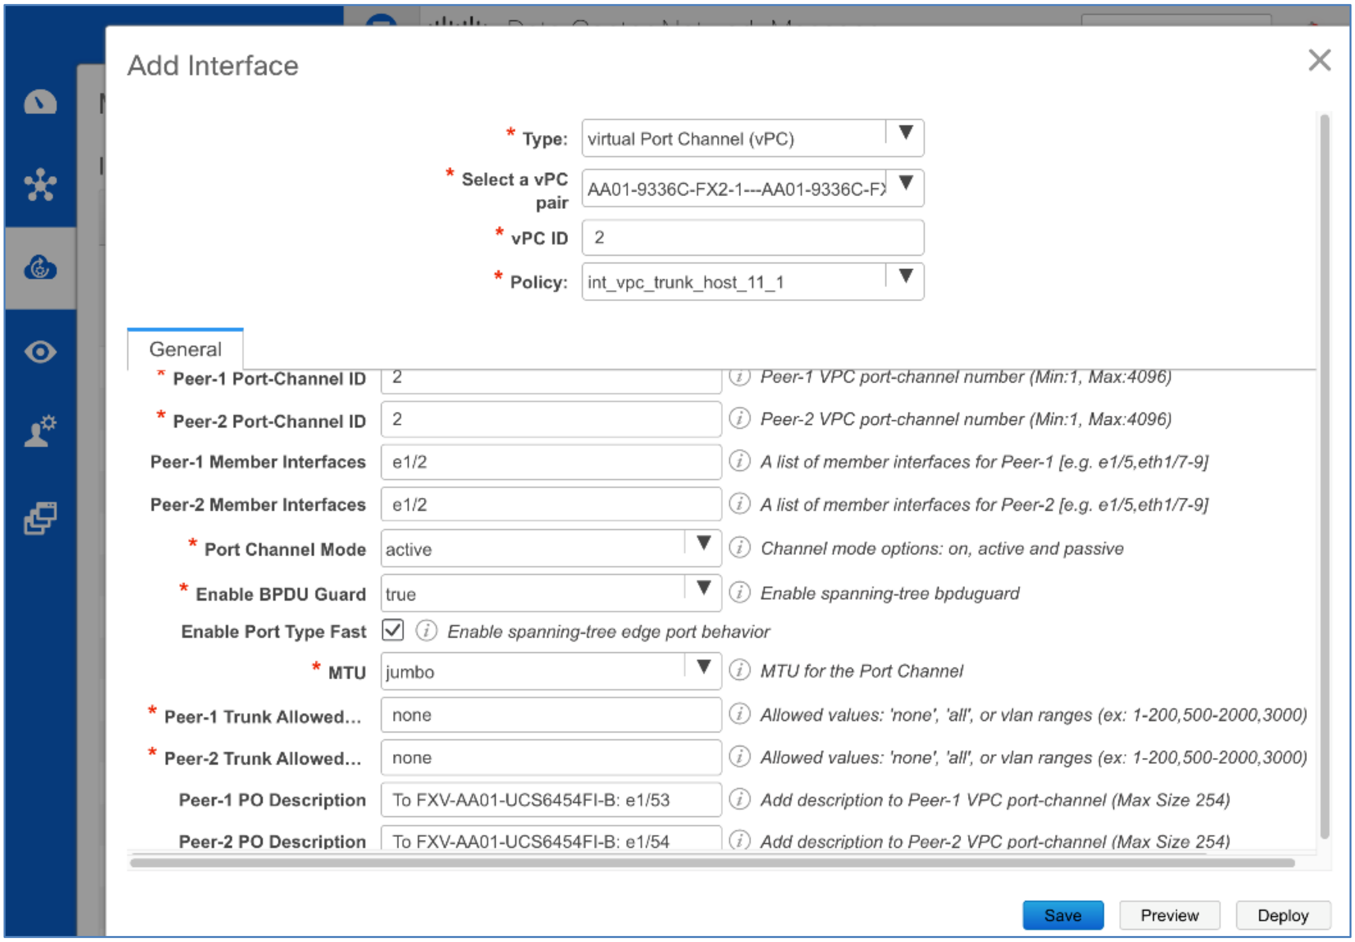

Add Interface

* Type: virtual Port Channel (vPC)

* Select a vPC pair: AA01-9336C-FX2-1---AA01-9336C-F)

* vPC ID: 2

* Policy: int_vpc_trunk_host_11_1

General

* Peer-1 Port-Channel ID	2	*Peer-1 VPC port-channel number (Min:1, Max:4096)*
* Peer-2 Port-Channel ID	2	*Peer-2 VPC port-channel number (Min:1, Max:4096)*
Peer-1 Member Interfaces	e1/2	*A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]*
Peer-2 Member Interfaces	e1/2	*A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]*
* Port Channel Mode	active	*Channel mode options: on, active and passive*
* Enable BPDU Guard	true	*Enable spanning-tree bpduguard*
Enable Port Type Fast	☑	*Enable spanning-tree edge port behavior*
* MTU	jumbo	*MTU for the Port Channel*
* Peer-1 Trunk Allowed...	none	*Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)*
* Peer-2 Trunk Allowed...	none	*Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)*
Peer-1 PO Description	To FXV-AA01-UCS6454FI-B: e1/53	*Add description to Peer-1 VPC port-channel (Max Size 254)*
Peer-2 PO Description	To FXV-AA01-UCS6454FI-B: e1/54	*Add description to Peer-2 VPC port-channel (Max Size 254)*

18. Click **Save.** Click on **Preview** and **Deploy** to preview and deploy the configuration for the vPC from the leaf switches in the VXLAN fabric to the second Cisco UCS Fabric Interconnect (FI-B). Click **OK** in the pop-up window. Click the **X** to close the **Manage Interfaces** window.

19. From the left navigation bar, select **Control > Fabric > Fabric Builder.** Select **Site-A** fabric. In the Topology view, address any issues that are highlighted next to the **Save & Deploy** button.

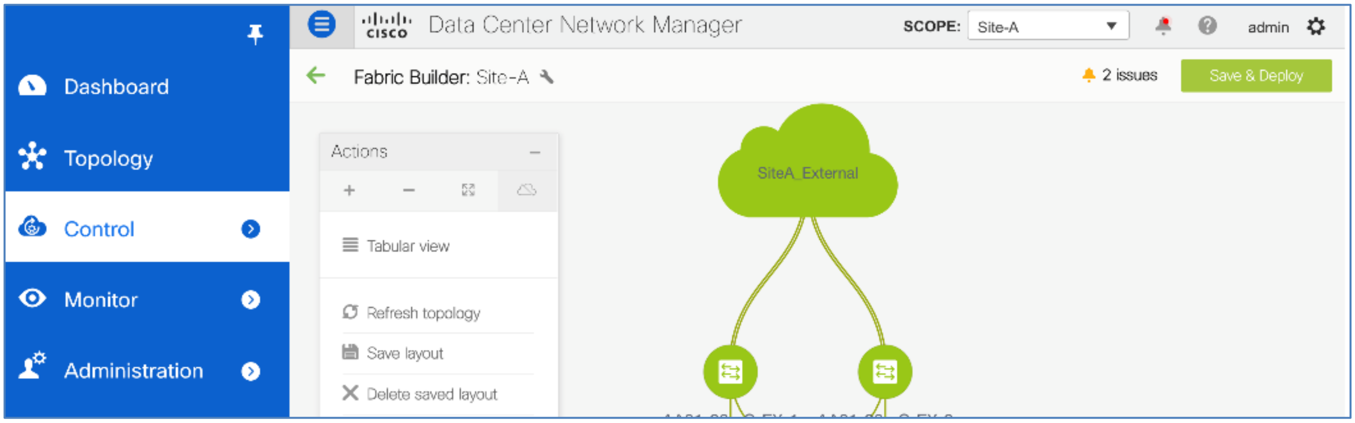

20. There are **2 issues** in this deployment. Click on the issues to get more information.

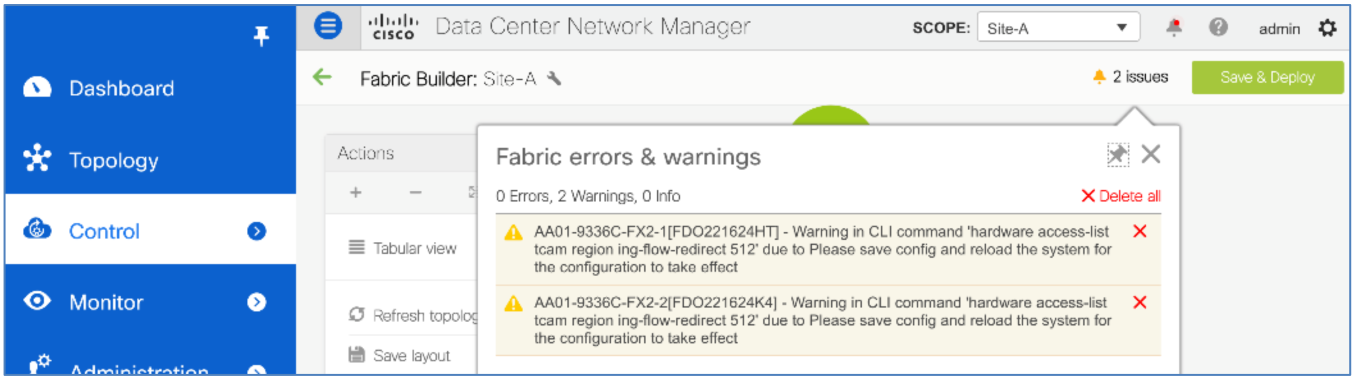

Fabric errors & warnings

0 Errors, 2 Warnings, 0 Info ✕ Delete all

⚠ AA01-9336C-FX2-1[FDO221624HT] - Warning in CLI command 'hardware access-list tcam region ing-flow-redirect 512' due to Please save config and reload the system for the configuration to take effect

⚠ AA01-9336C-FX2-2[FDO221624K4] - Warning in CLI command 'hardware access-list tcam region ing-flow-redirect 512' due to Please save config and reload the system for the configuration to take effect

21. To address the warnings, from the **Action** menu/window, click on **Tabular view.** Select all switches and click on the floppy drive icon to save the configuration on all switches.

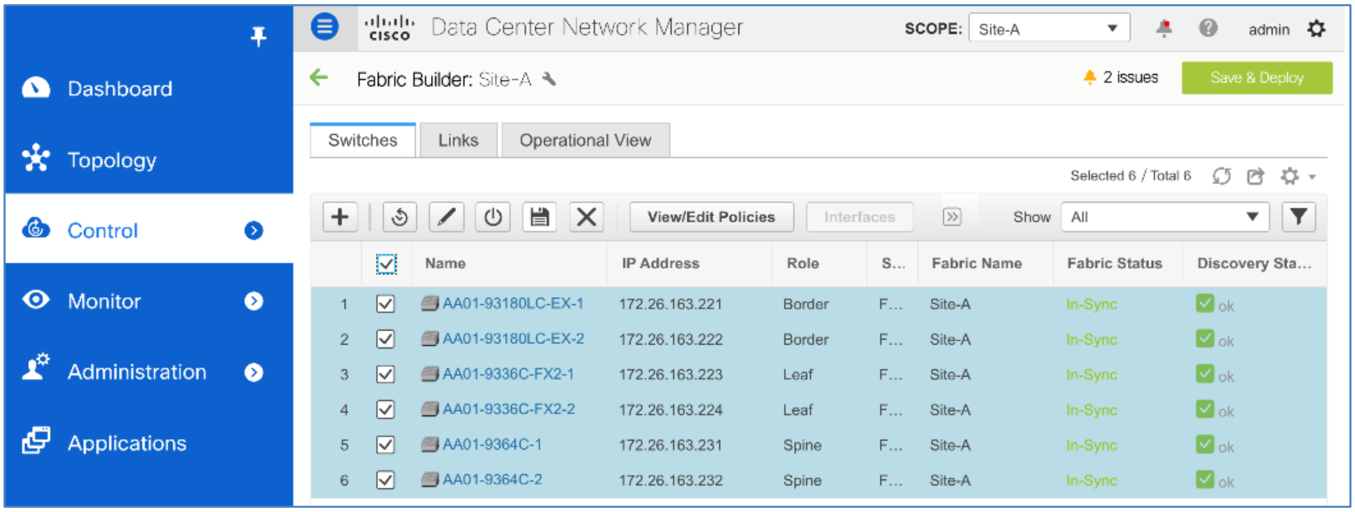

	☑	Name	IP Address	Role	S...	Fabric Name	Fabric Status	Discovery Sta...
1	☑	AA01-93180LC-EX-1	172.26.163.221	Border	F...	Site-A	In-Sync	✓ ok
2	☑	AA01-93180LC-EX-2	172.26.163.222	Border	F...	Site-A	In-Sync	✓ ok
3	☑	AA01-9336C-FX2-1	172.26.163.223	Leaf	F...	Site-A	In-Sync	✓ ok
4	☑	AA01-9336C-FX2-2	172.26.163.224	Leaf	F...	Site-A	In-Sync	✓ ok
5	☑	AA01-9364C-1	172.26.163.231	Spine	F...	Site-A	In-Sync	✓ ok
6	☑	AA01-9364C-2	172.26.163.232	Spine	F...	Site-A	In-Sync	✓ ok

22. Click **Close** when the save completes successfully.

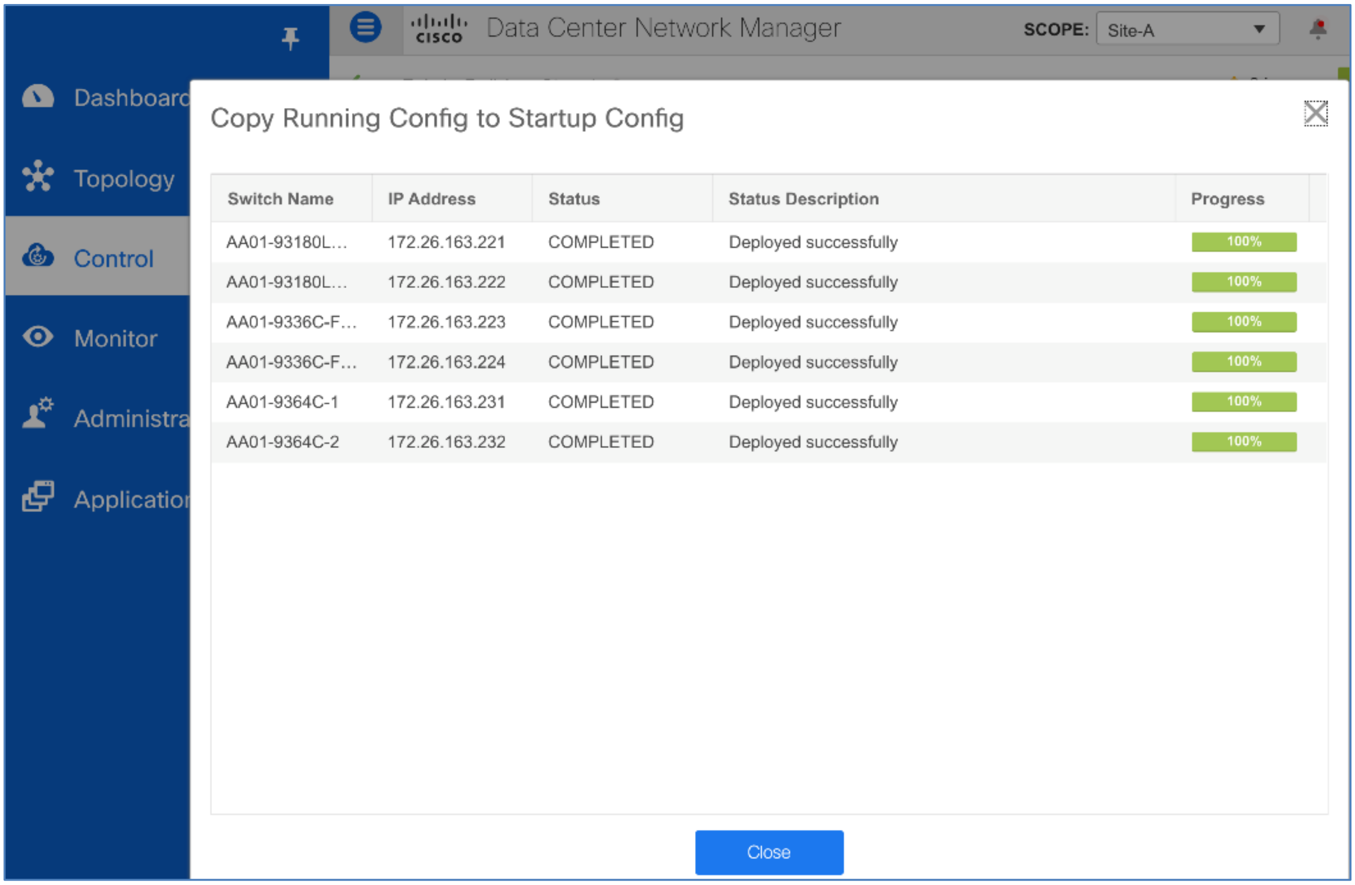

Copy Running Config to Startup Config

Switch Name	IP Address	Status	Status Description	Progress
AA01-93180L…	172.26.163.221	COMPLETED	Deployed successfully	100%
AA01-93180L…	172.26.163.222	COMPLETED	Deployed successfully	100%
AA01-9336C-F…	172.26.163.223	COMPLETED	Deployed successfully	100%
AA01-9336C-F…	172.26.163.224	COMPLETED	Deployed successfully	100%
AA01-9364C-1	172.26.163.231	COMPLETED	Deployed successfully	100%
AA01-9364C-2	172.26.163.232	COMPLETED	Deployed successfully	100%

23. Deselect all the switches and select only the switches that need to be reloaded per the Warning. Click on the Power icon to reboot the switches. Click **OK** in the pop-up window. Monitor the **Discovery Status** column for a status as the switch reboots or access the console of the switch in question directly.

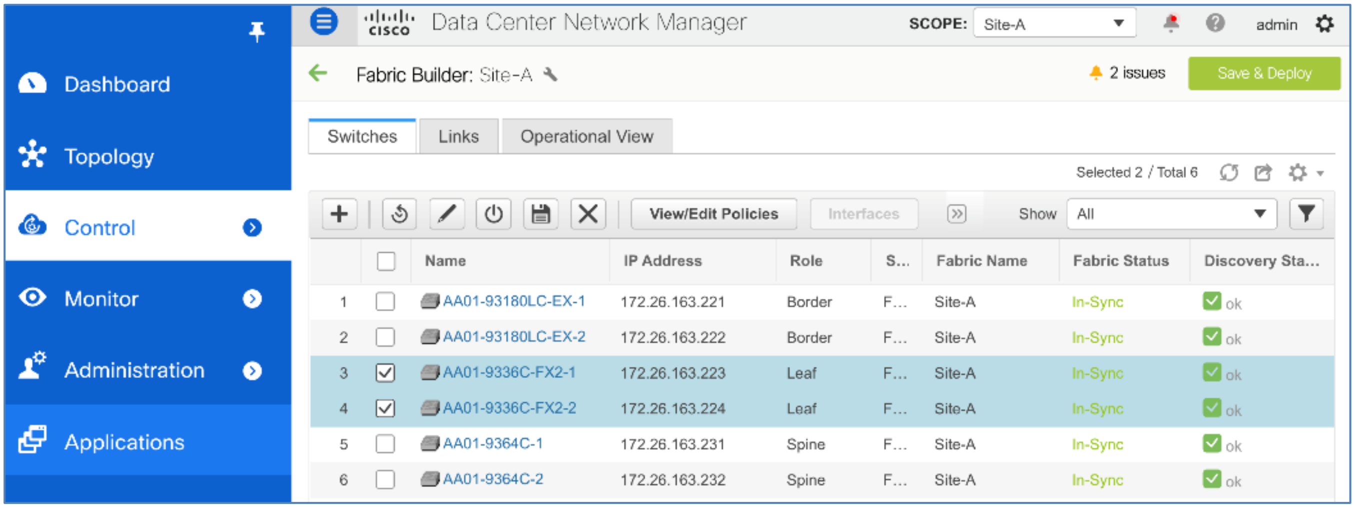

Fabric Builder: Site-A

Selected 2 / Total 6

	☐	Name	IP Address	Role	S…	Fabric Name	Fabric Status	Discovery Sta…
1	☐	AA01-93180LC-EX-1	172.26.163.221	Border	F…	Site-A	In-Sync	✓ ok
2	☐	AA01-93180LC-EX-2	172.26.163.222	Border	F…	Site-A	In-Sync	✓ ok
3	☑	AA01-9336C-FX2-1	172.26.163.223	Leaf	F…	Site-A	In-Sync	✓ ok
4	☑	AA01-9336C-FX2-2	172.26.163.224	Leaf	F…	Site-A	In-Sync	✓ ok
5	☐	AA01-9364C-1	172.26.163.231	Spine	F…	Site-A	In-Sync	✓ ok
6	☐	AA01-9364C-2	172.26.163.232	Spine	F…	Site-A	In-Sync	✓ ok

24. When the reboot completes after a few minutes, verify that the vPC is in a **consistent** state and the port-channel is up and operational.

25. From the left navigation bar, select **Control > Fabric > Interfaces.** Filter on the **Name** to view the vPCs deployed. Select the **Quick Filter** from the drop-down list next to **Show** to see the boxes for filtering under each column. Confirm that the vPC are in a **Consistent** state – see **Reason** column.

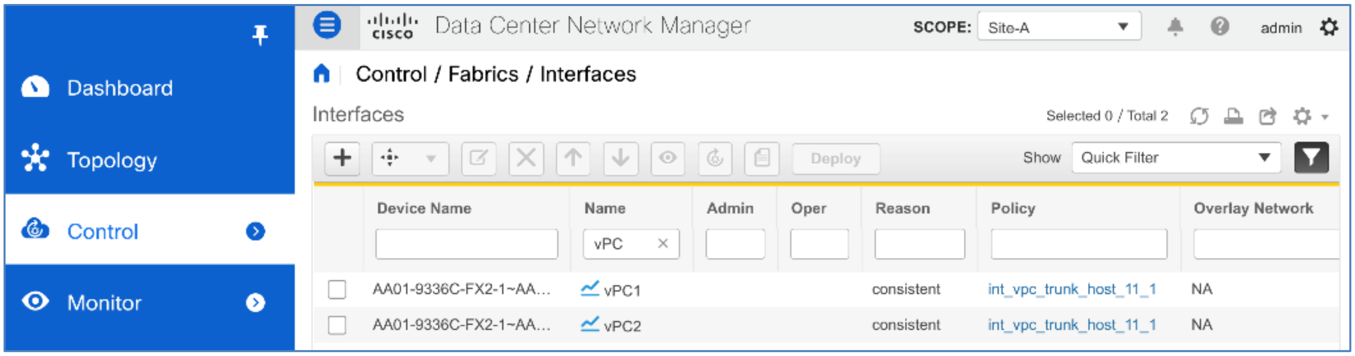

SCOPE: Site-A

Control / Fabrics / Interfaces

Interfaces — Selected 0 / Total 2

Show: Quick Filter

	Device Name	Name	Admin	Oper	Reason	Policy	Overlay Network
		vPC					
□	AA01-9336C-FX2-1~AA...	vPC1			consistent	int_vpc_trunk_host_11_1	NA
□	AA01-9336C-FX2-1~AA...	vPC2			consistent	int_vpc_trunk_host_11_1	NA

26. Filter on the **Name** to view the Port-Channels in the above vPCs. Note that the **Admin** and **Oper** status are up. Scroll to the right to see additional columns. Verify that the status is Green.

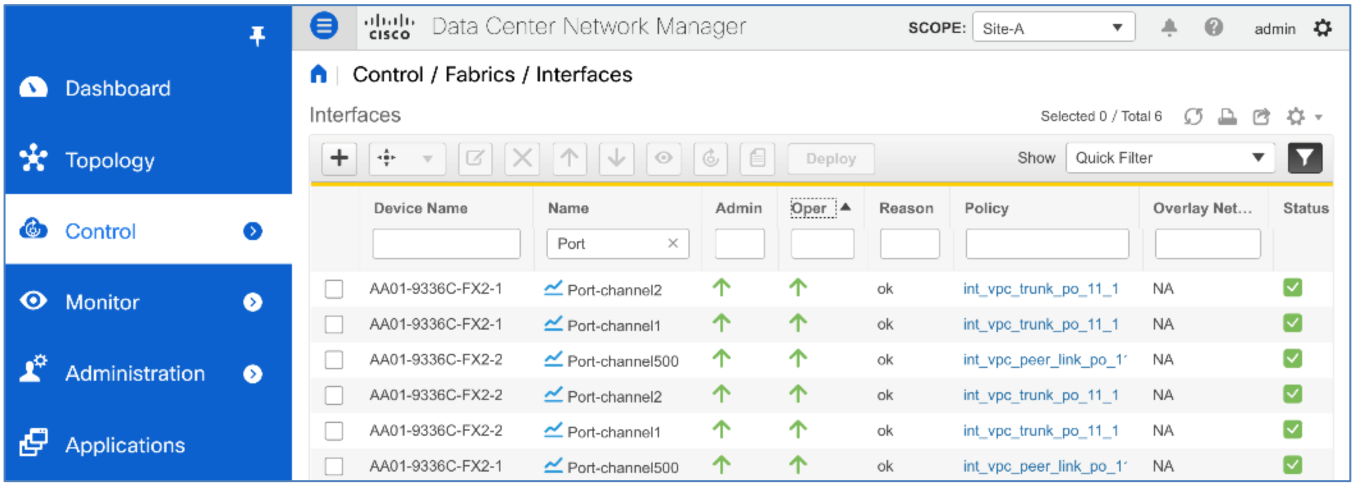

SCOPE: Site-A

Control / Fabrics / Interfaces

Interfaces — Selected 0 / Total 6

Show: Quick Filter

	Device Name	Name	Admin	Oper	Reason	Policy	Overlay Net...	Status
		Port						
□	AA01-9336C-FX2-1	Port-channel2	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-1	Port-channel1	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel500	↑	↑	ok	int_vpc_peer_link_po_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel2	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel1	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-1	Port-channel500	↑	↑	ok	int_vpc_peer_link_po_1	NA	✓

27. From the left navigation bar, select **Control > Fabric > Fabric Builder.** Select **Site-A** fabric. From the **Actions** menu/window, select **Backup Now** to back up the Site-A fabric. Repeat the same for the SiteA_External fabric by changing the **Scope:** to **SiteA_External** from the drop-down list.

Enable Access Layer Connectivity to NetApp Storage Cluster

In this section, access-layer connectivity is enabled from the VXLAN fabric in Site-A to the NetApp Storage infrastructure in the solution. The NetApp storage infrastructure in this solution consists of an AFF A300 storage array.

Topology

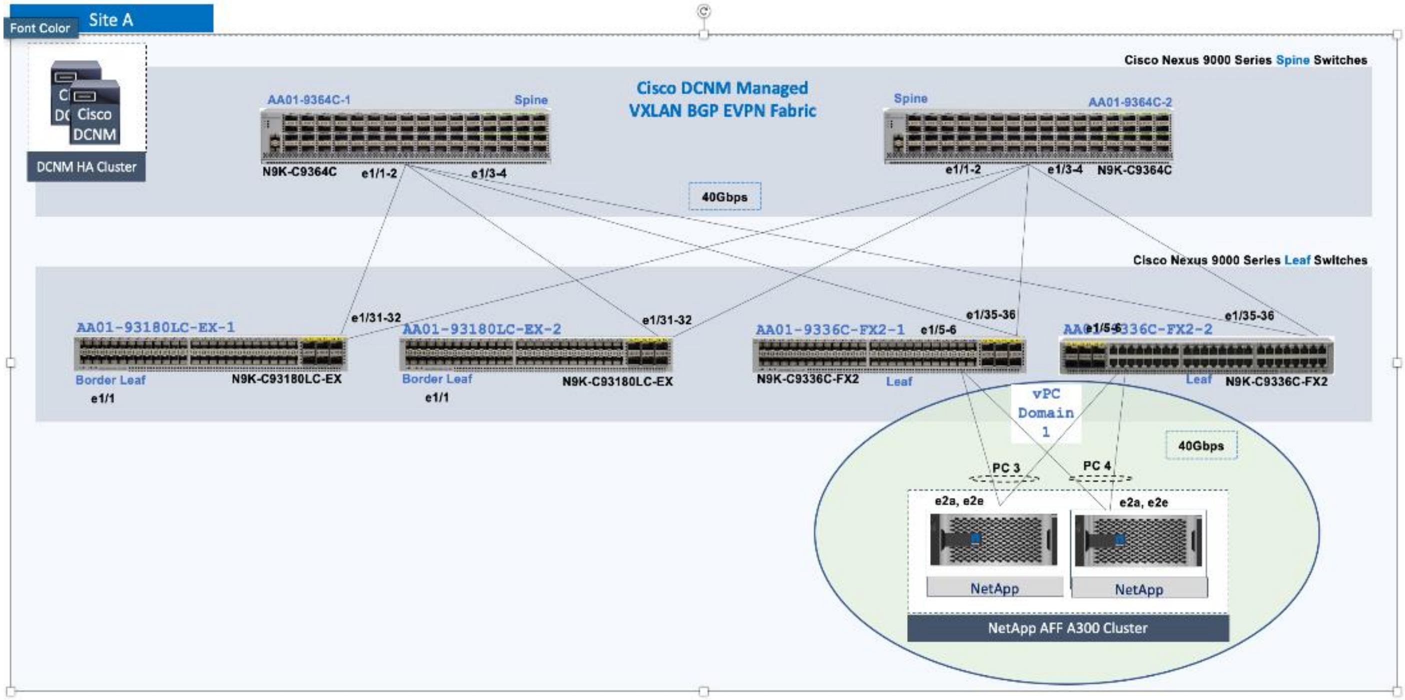

Setup Information

The configuration parameters for enabling access-layer connectivity to NetApp storage cluster in Site-A data center fabric is provided below.

Table 10. Access Layer Switches – To NetApp Storage Cluster

Hostname	Switch Role	IP Address (OOB)	Notes
AA01-9336C-FX2-1	Leaf	172.26.163.223/24	
AA01-9336C-FX2-2	Leaf	172.26.163.224/24	

Table 11. Access Layer Connectivity – To NetApp Storage Cluster

Access Layer Connection	Parameters	Notes
Type	Virtual Port-Channel (vPC)	Using Virtual Peer-Links (requires hardware support)
vPC Pair	AA01-9336C-FX2-1---AA01-9336C-FX2-2	
vPC to NetApp Controller-A		
Peer-1 Member Interfaces	Ethernet 1/5	

Access Layer Connection	Parameters	Notes
Peer-2 Member Interfaces	Ethernet 1/5	
Peer-1 PO Description	To FXV-BB09-A300-2-01: e2a	
Peer-2 PO Description	To FXV-BB09-A300-2-01: e2e	
vPC to NetApp Controller-B		
Peer-1 Member Interfaces	Ethernet 1/6	
Peer-2 Member Interfaces	Ethernet 1/6	
Peer-1 PO Description	To FXV-BB09-A300-2-02: e2a	
Peer-2 PO Description	To FXV-BB09-A300-2-02: e2e	

Deployment Steps

To enable access-layer connectivity from Site-A data center fabric to the NetApp Storage Cluster, follow these steps using the **Setup Information** provided above:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Fabric Builder**. Click on the **Site-A** fabric.

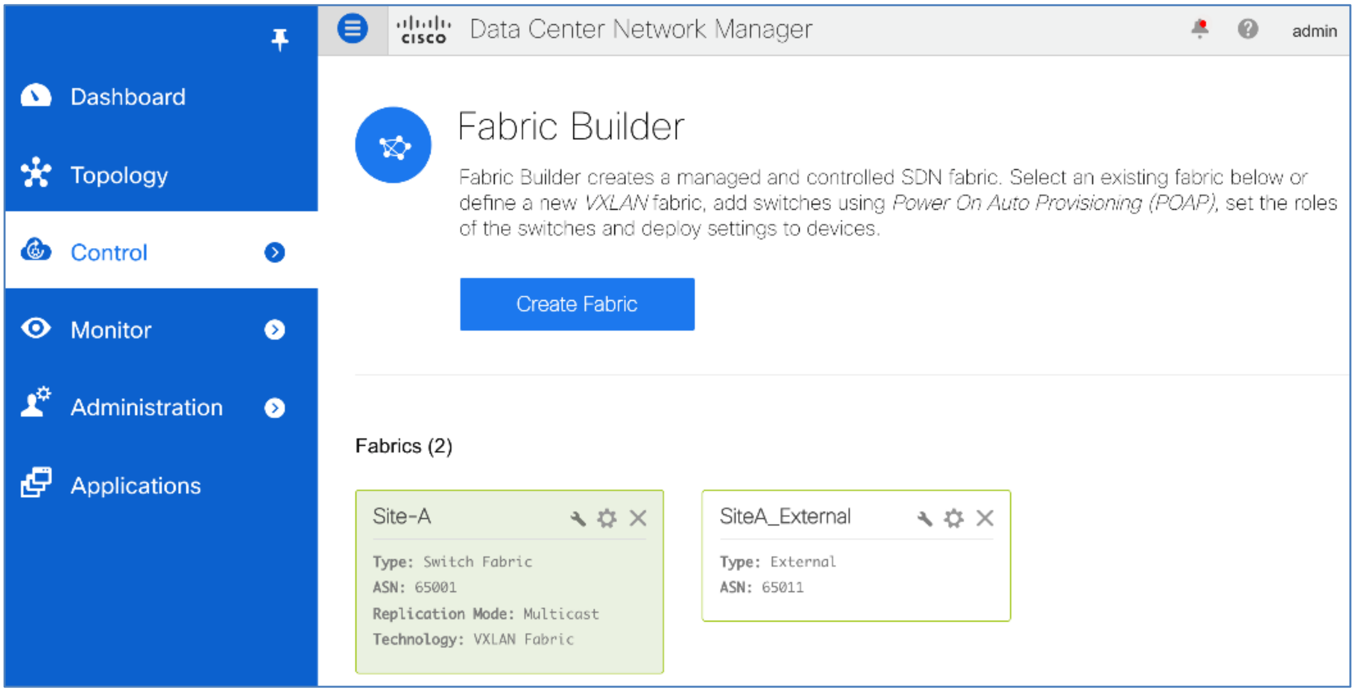

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (2)

Site-A
Type: Switch Fabric
ASN: 65001
Replication Mode: Multicast
Technology: VXLAN Fabric

SiteA_External
Type: External
ASN: 65011

3. From the right window pane, select one of the Leaf switches in the vPC pair that connect to the NetApp Storage Cluster. Right-click and select **Manage Interfaces** from the list.

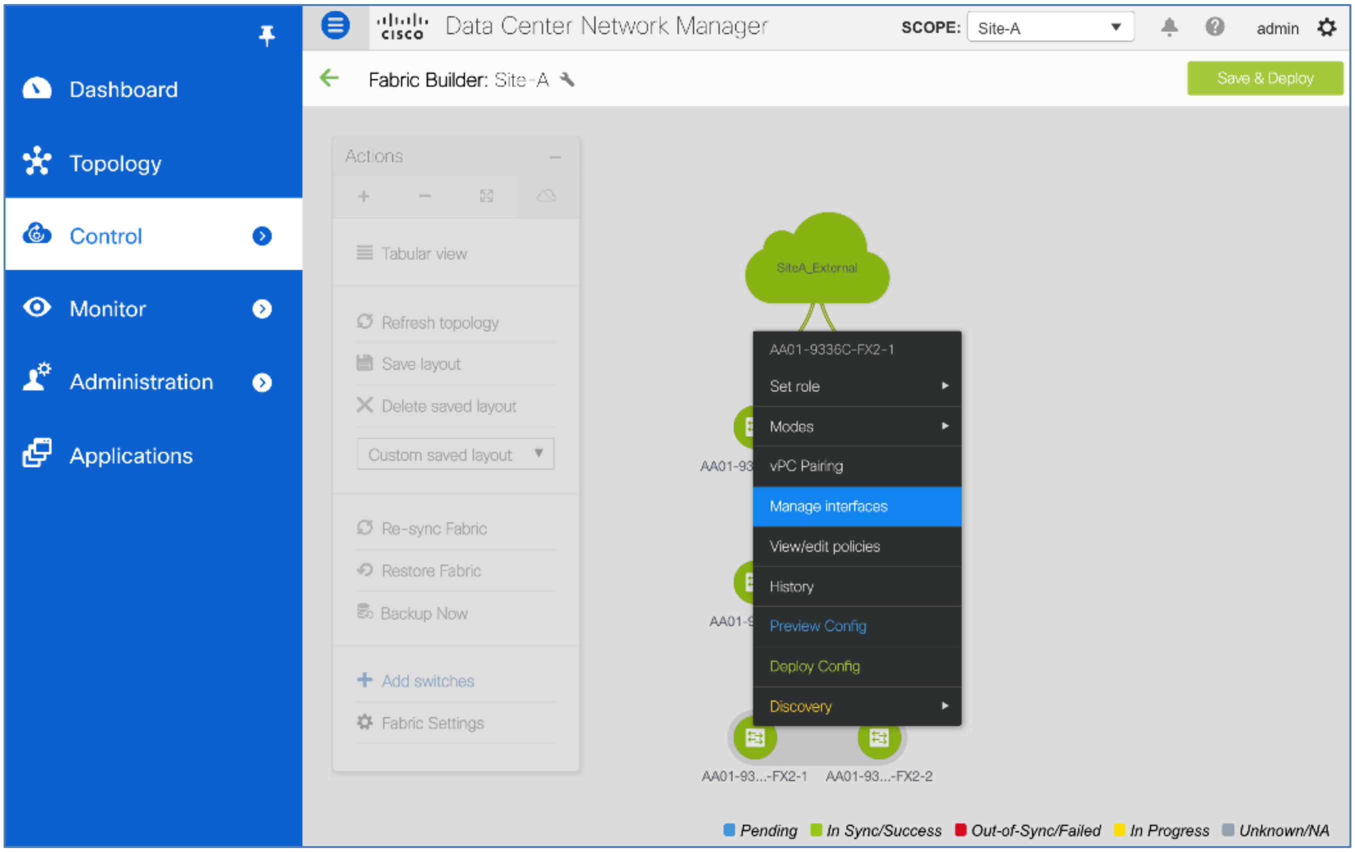

4. In the **Manage Interface** pop-up window, bring the **Neighbors** column into view by dragging it from the far-right end and move it next to the **Reason** column. Sort the interfaces based on the **Neighbor** column. Note the interfaces on the leaf switches in the vPC pair that connect to the first NetApp controller and the port numbers they connect to.

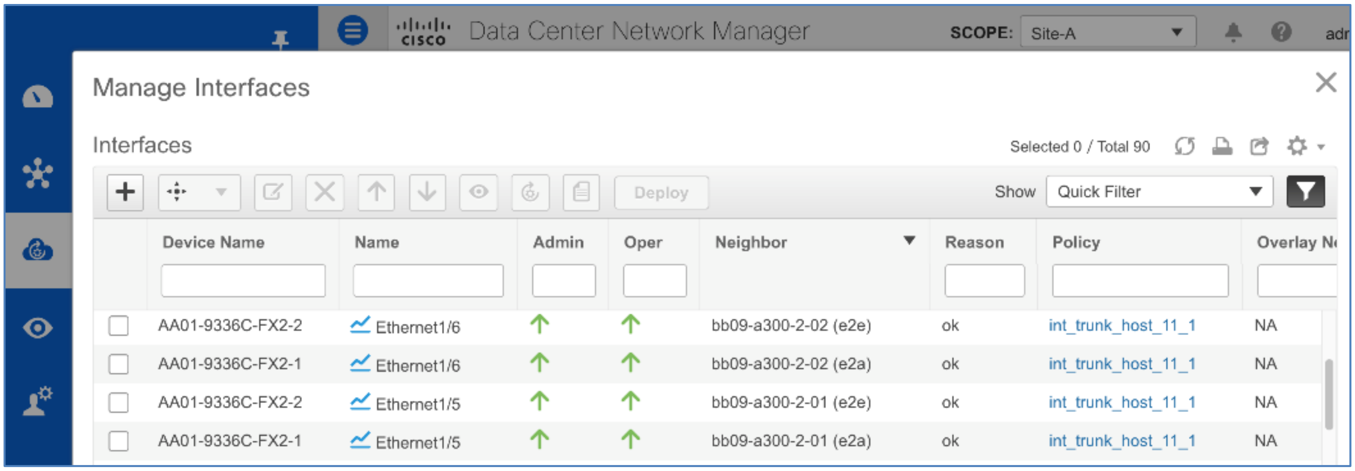

	Device Name	Name	Admin	Oper	Neighbor	Reason	Policy	Overlay N
□	AA01-9336C-FX2-2	Ethernet1/6	↑	↑	bb09-a300-2-02 (e2e)	ok	int_trunk_host_11_1	NA
□	AA01-9336C-FX2-1	Ethernet1/6	↑	↑	bb09-a300-2-02 (e2a)	ok	int_trunk_host_11_1	NA
□	AA01-9336C-FX2-2	Ethernet1/5	↑	↑	bb09-a300-2-01 (e2e)	ok	int_trunk_host_11_1	NA
□	AA01-9336C-FX2-1	Ethernet1/5	↑	↑	bb09-a300-2-01 (e2a)	ok	int_trunk_host_11_1	NA

5. Click the **[+]** button from the menu above. In the **Add Interface** pop-up window, specify the **Type** from the drop-down list.

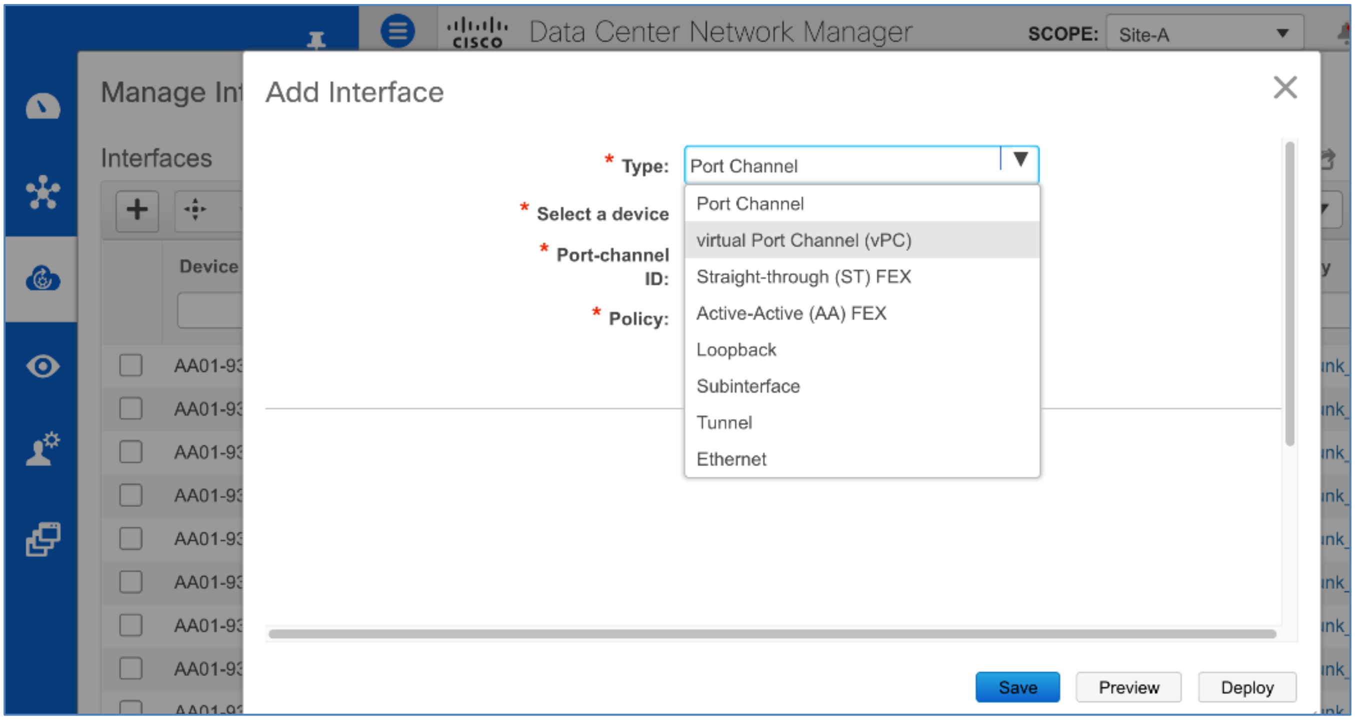

* Type: Port Channel

* Select a device

* Port-channel ID:

* Policy:

6. The menu changes to the reflect the options for the Interface **Type** selected. **Select a vPC pair** from the drop-down list.

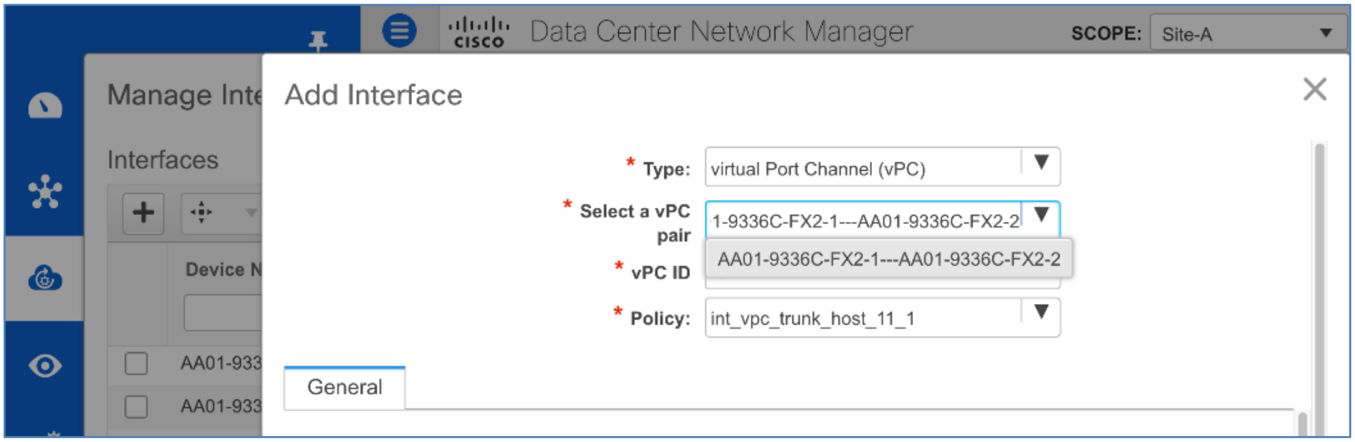

* Type: virtual Port Channel (vPC)

* Select a vPC pair: 1-9336C-FX2-1---AA01-9336C-FX2-2

* vPC ID

* Policy: int_vpc_trunk_host_11_1

7. Specify the Peer-1 Member Interfaces, Peer-2 Member Interfaces, Peer-1 PO Description, and Peer-2 PO Description. Leave all other fields as-is.

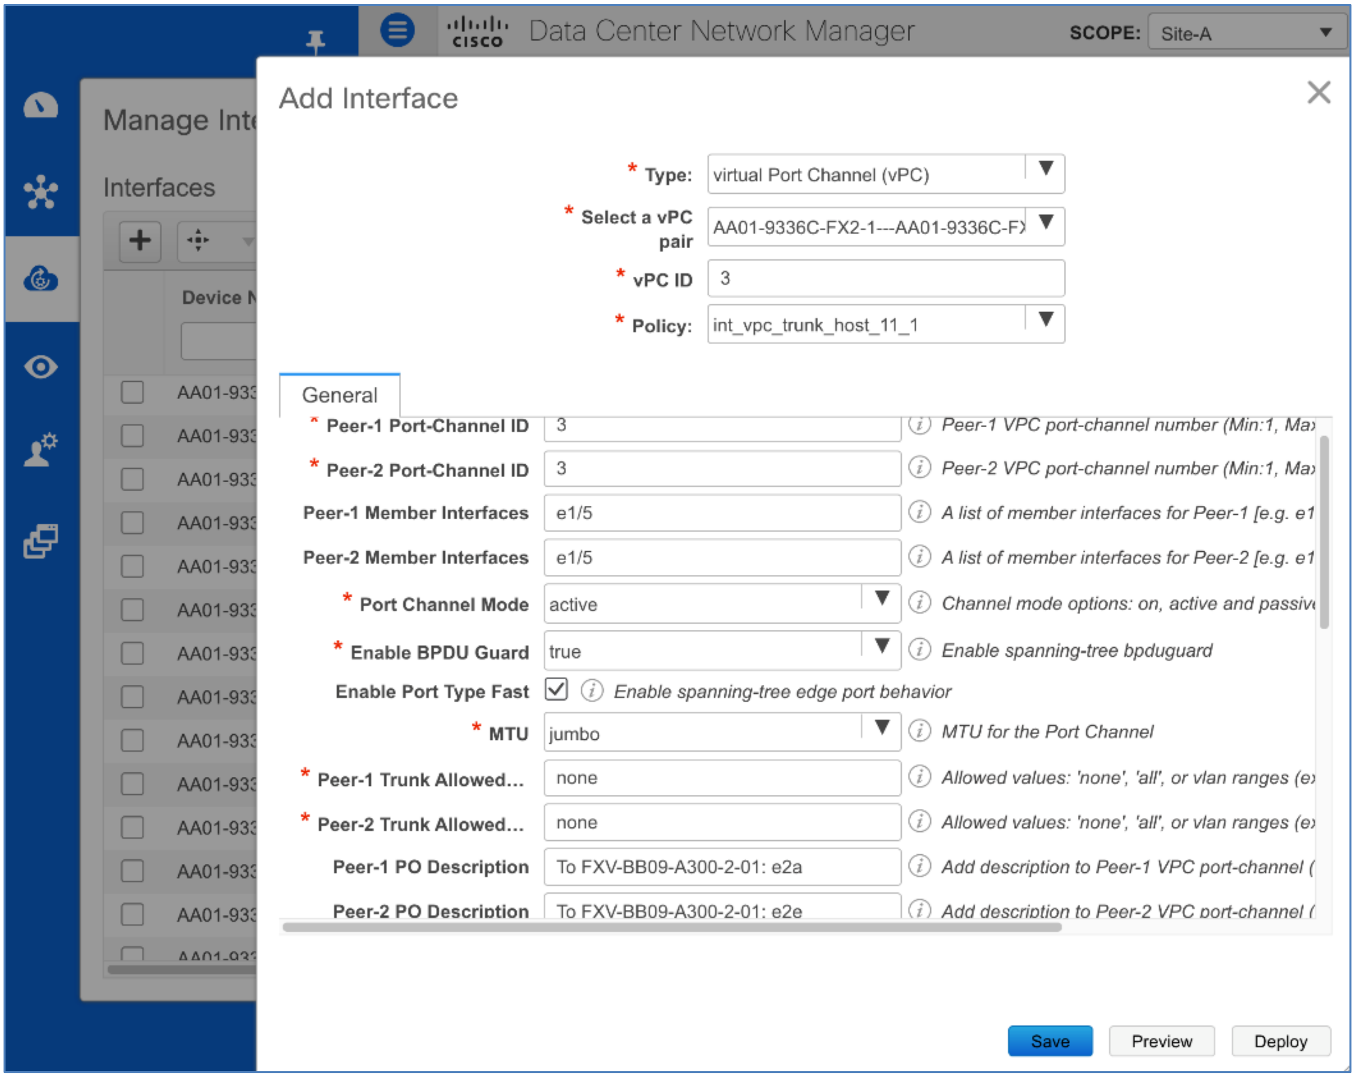

Add Interface

* Type: virtual Port Channel (vPC)

* Select a vPC pair: AA01-9336C-FX2-1---AA01-9336C-F)

* vPC ID: 3

* Policy: int_vpc_trunk_host_11_1

General

* Peer-1 Port-Channel ID	3	Peer-1 VPC port-channel number (Min:1, Ma)
* Peer-2 Port-Channel ID	3	Peer-2 VPC port-channel number (Min:1, Ma)
Peer-1 Member Interfaces	e1/5	A list of member interfaces for Peer-1 [e.g. e1
Peer-2 Member Interfaces	e1/5	A list of member interfaces for Peer-2 [e.g. e1
* Port Channel Mode	active	Channel mode options: on, active and passive
* Enable BPDU Guard	true	Enable spanning-tree bpduguard
Enable Port Type Fast	☑	Enable spanning-tree edge port behavior
* MTU	jumbo	MTU for the Port Channel
* Peer-1 Trunk Allowed...	none	Allowed values: 'none', 'all', or vlan ranges (e)
* Peer-2 Trunk Allowed...	none	Allowed values: 'none', 'all', or vlan ranges (e)
Peer-1 PO Description	To FXV-BB09-A300-2-01: e2a	Add description to Peer-1 VPC port-channel (
Peer-2 PO Description	To FXV-BB09-A300-2-01: e2e	Add description to Peer-2 VPC port-channel (

8. Click **Save.** You can also **Preview** configuration for the vPC to the first NetApp controller using the **Preview** button. The preview will display the pending configuration for each switch in the vPC pair – use the drop-down list to select the second switch.

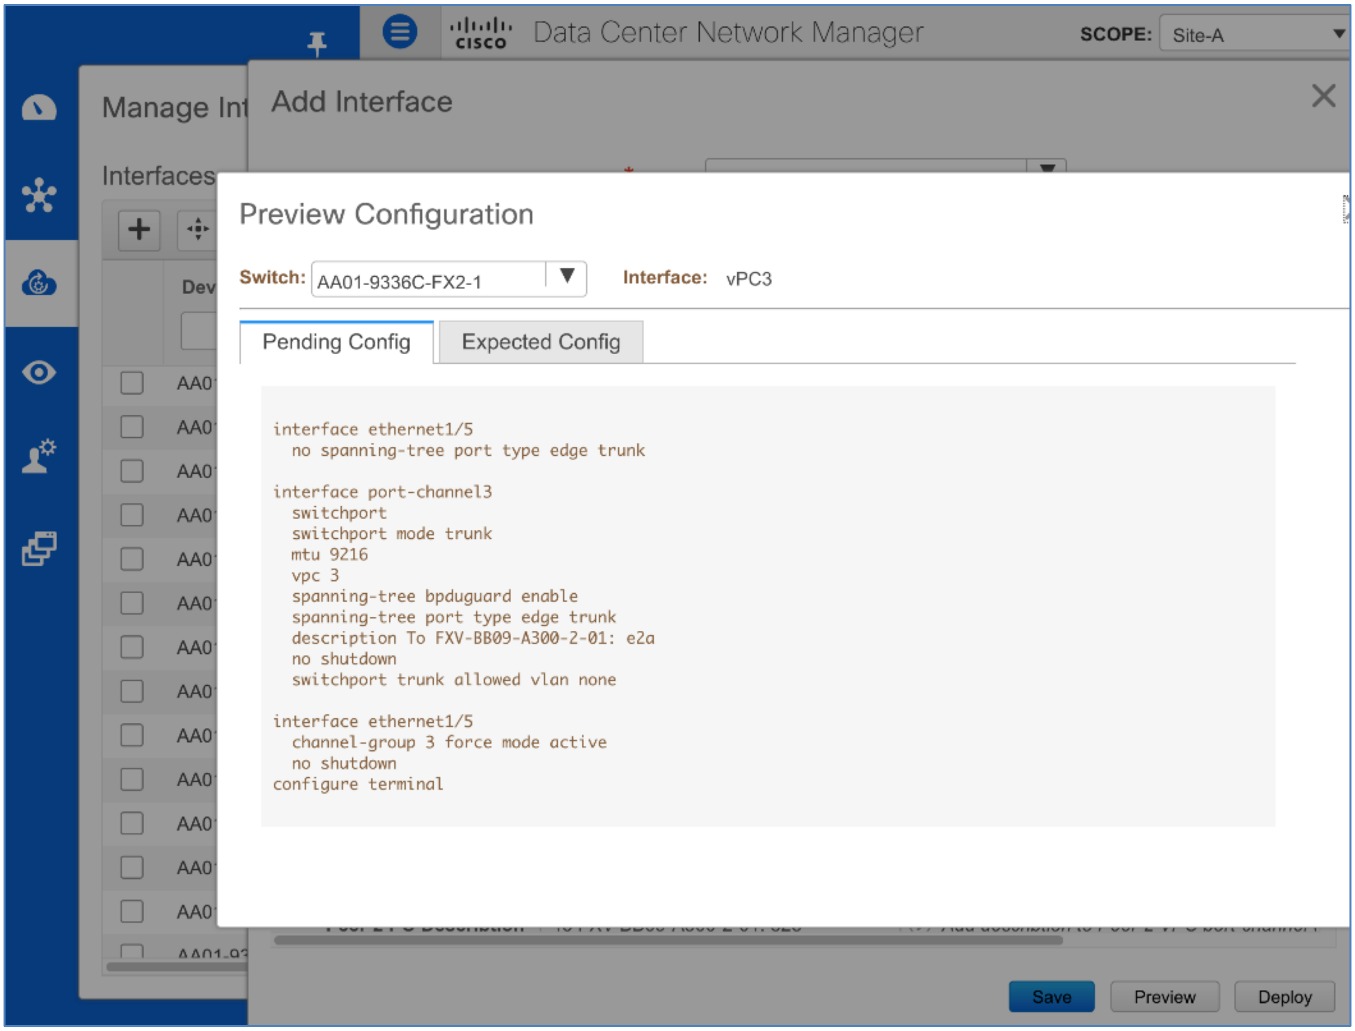

Preview Configuration

Switch: AA01-9336C-FX2-1

Interface: vPC3

Pending Config | Expected Config

```
interface ethernet1/5
  no spanning-tree port type edge trunk

interface port-channel3
  switchport
  switchport mode trunk
  mtu 9216
  vpc 3
  spanning-tree bpduguard enable
  spanning-tree port type edge trunk
  description To FXV-BB09-A300-2-01: e2a
  no shutdown
  switchport trunk allowed vlan none

interface ethernet1/5
  channel-group 3 force mode active
  no shutdown
configure terminal
```

9. Click **X** to close the window. Click **Deploy** to deploy the vPC configuration from the Leaf switch pair in the VXLAN fabric to the first NetApp controller. Click **OK** in the pop-up window.

10. Repeat steps 1-9 to create, preview, and deploy the vPC to the second NetApp controller.

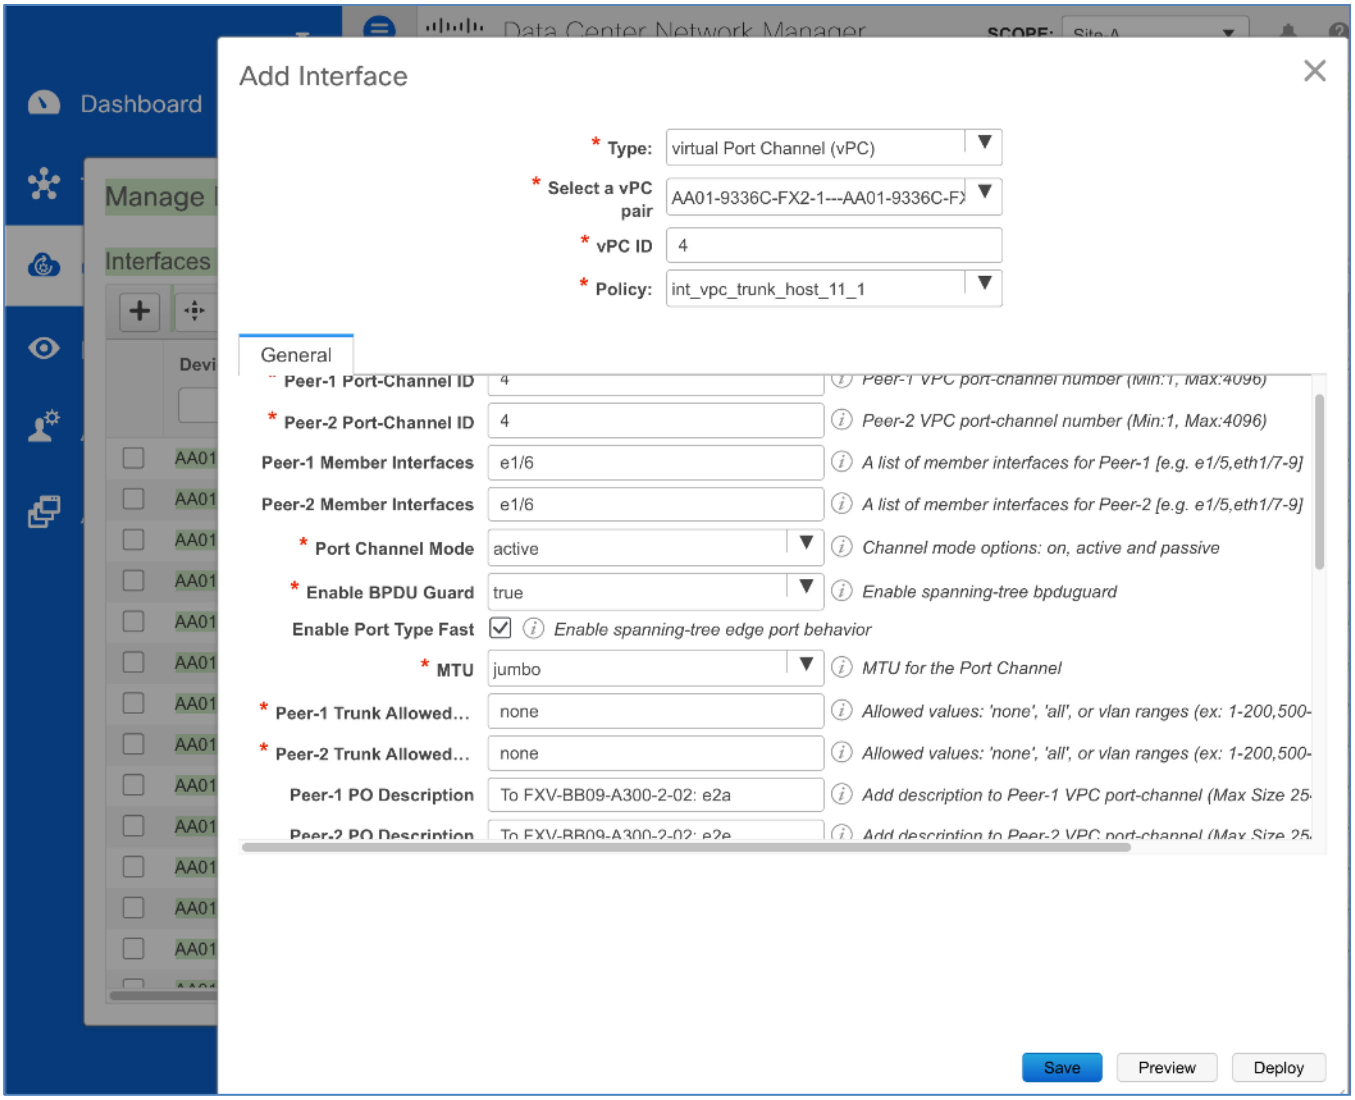

Add Interface

* Type: virtual Port Channel (vPC)

* Select a vPC pair: AA01-9336C-FX2-1---AA01-9336C-F)

* vPC ID: 4

* Policy: int_vpc_trunk_host_11_1

General

* Peer-1 Port-Channel ID: 4 — Peer-1 VPC port-channel number (Min:1, Max:4096)

* Peer-2 Port-Channel ID: 4 — Peer-2 VPC port-channel number (Min:1, Max:4096)

Peer-1 Member Interfaces: e1/6 — A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]

Peer-2 Member Interfaces: e1/6 — A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]

* Port Channel Mode: active — Channel mode options: on, active and passive

* Enable BPDU Guard: true — Enable spanning-tree bpduguard

Enable Port Type Fast: ☑ — Enable spanning-tree edge port behavior

* MTU: jumbo — MTU for the Port Channel

* Peer-1 Trunk Allowed...: none — Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-

* Peer-2 Trunk Allowed...: none — Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-

Peer-1 PO Description: To FXV-BB09-A300-2-02: e2a — Add description to Peer-1 VPC port-channel (Max Size 25

Peer-2 PO Description: To FXV-BB09-A300-2-02: e2e — Add description to Peer-2 VPC port-channel (Max Size 25

11. Click **Save.** Click on **Preview** and **Deploy** to preview and deploy the configuration for the vPC from the leaf switches in the VXLAN fabric to the second Cisco UCS Fabric Interconnect (FI-B). Click **OK** in the pop-up window. Click **X** to close the **Manage Interfaces** window.

12. From the left navigation bar, select **Control > Fabric > Fabric Builder.** Select **Site-A** fabric. In the Topology view, see if any issues that are highlighted next to the **Save & Deploy** button.

13. From the **Actions** menu/window, click on **Tabular view.** Select all switches and click on the floppy drive icon to save the configuration on all switches.

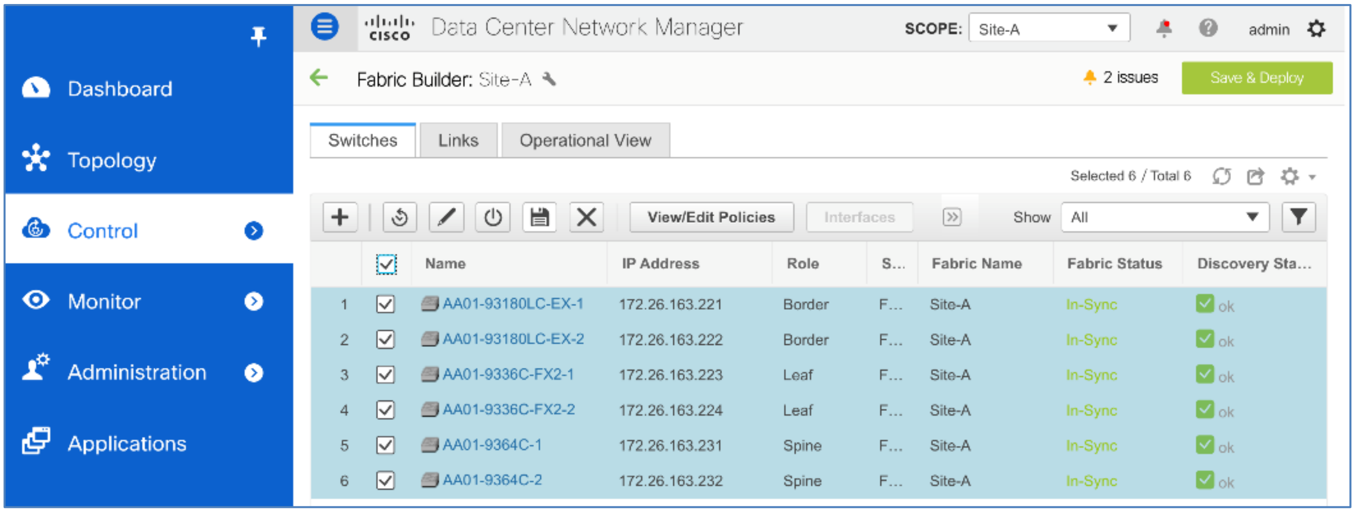

	☑	Name	IP Address	Role	S...	Fabric Name	Fabric Status	Discovery Sta...
1	☑	AA01-93180LC-EX-1	172.26.163.221	Border	F...	Site-A	In-Sync	✓ ok
2	☑	AA01-93180LC-EX-2	172.26.163.222	Border	F...	Site-A	In-Sync	✓ ok
3	☑	AA01-9336C-FX2-1	172.26.163.223	Leaf	F...	Site-A	In-Sync	✓ ok
4	☑	AA01-9336C-FX2-2	172.26.163.224	Leaf	F...	Site-A	In-Sync	✓ ok
5	☑	AA01-9364C-1	172.26.163.231	Spine	F...	Site-A	In-Sync	✓ ok
6	☑	AA01-9364C-2	172.26.163.232	Spine	F...	Site-A	In-Sync	✓ ok

14. Click **Close** when the save completes successfully.

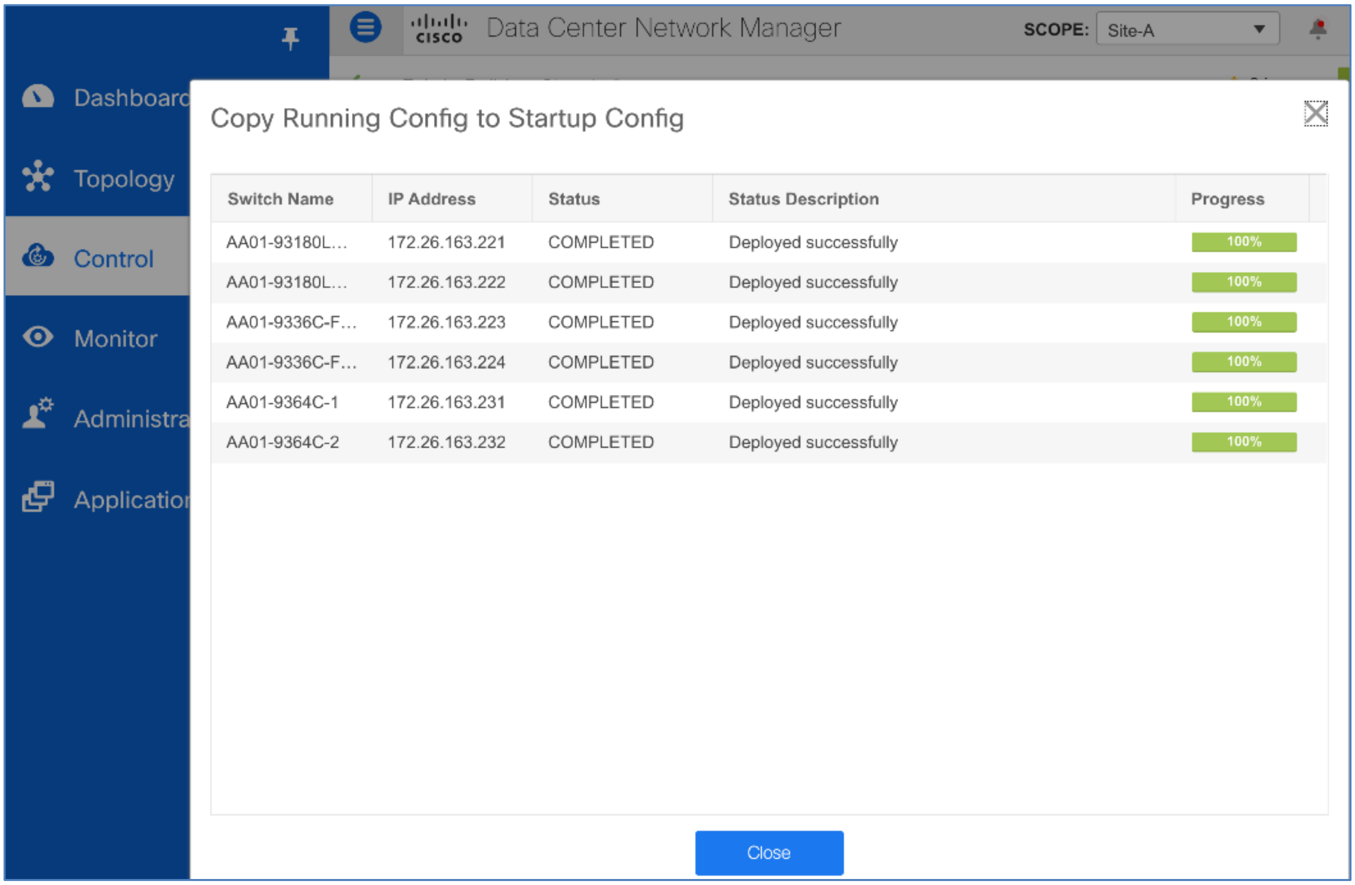

Copy Running Config to Startup Config

Switch Name	IP Address	Status	Status Description	Progress
AA01-93180L...	172.26.163.221	COMPLETED	Deployed successfully	100%
AA01-93180L...	172.26.163.222	COMPLETED	Deployed successfully	100%
AA01-9336C-F...	172.26.163.223	COMPLETED	Deployed successfully	100%
AA01-9336C-F...	172.26.163.224	COMPLETED	Deployed successfully	100%
AA01-9364C-1	172.26.163.231	COMPLETED	Deployed successfully	100%
AA01-9364C-2	172.26.163.232	COMPLETED	Deployed successfully	100%

15. Verify that the vPC is in a **consistent** state and the port-channel is up and operational.

16. From the left navigation bar, select **Control > Fabric > Interfaces.** Filter on the **Name** to view the vPCs deployed. Select the **Quick Filter** from the drop-down list next to **Show** to see the boxes for filtering under each column. Confirm that the vPC are in a **Consistent** state – see **Reason** column.

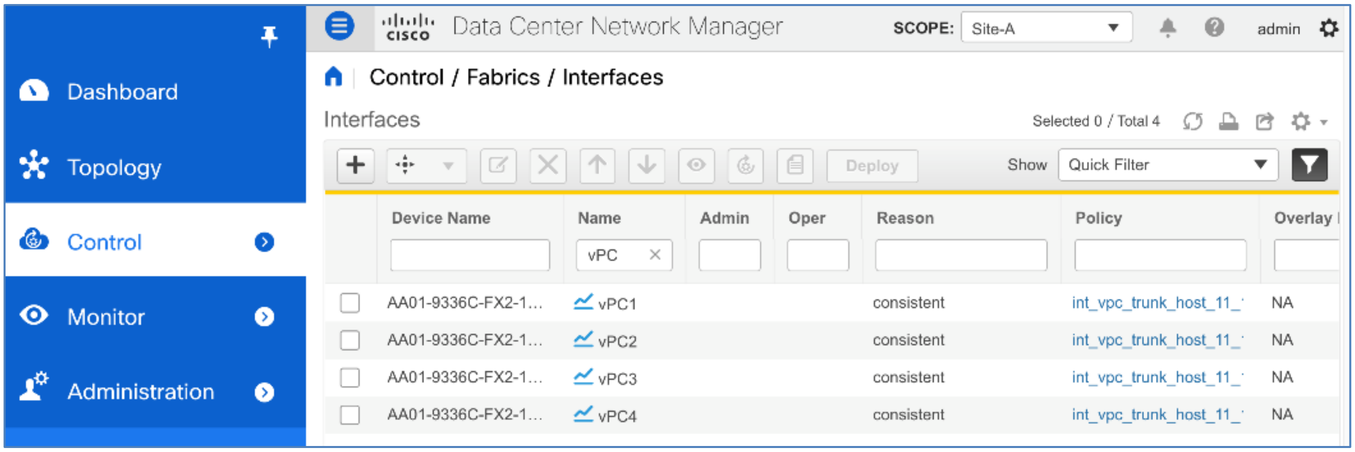

SCOPE: Site-A

Control / Fabrics / Interfaces

Interfaces

Selected 0 / Total 4

Show Quick Filter

	Device Name	Name	Admin	Oper	Reason	Policy	Overlay I
		vPC					
□	AA01-9336C-FX2-1...	vPC1			consistent	int_vpc_trunk_host_11_'	NA
□	AA01-9336C-FX2-1...	vPC2			consistent	int_vpc_trunk_host_11_'	NA
□	AA01-9336C-FX2-1...	vPC3			consistent	int_vpc_trunk_host_11_'	NA
□	AA01-9336C-FX2-1...	vPC4			consistent	int_vpc_trunk_host_11_'	NA

17. Filter on the **Name** to view the Port-Channels in the above vPCs. Note that the **Admin** and **Oper** status are up. Scroll to the right to see additional columns. Verify that the status is Green.

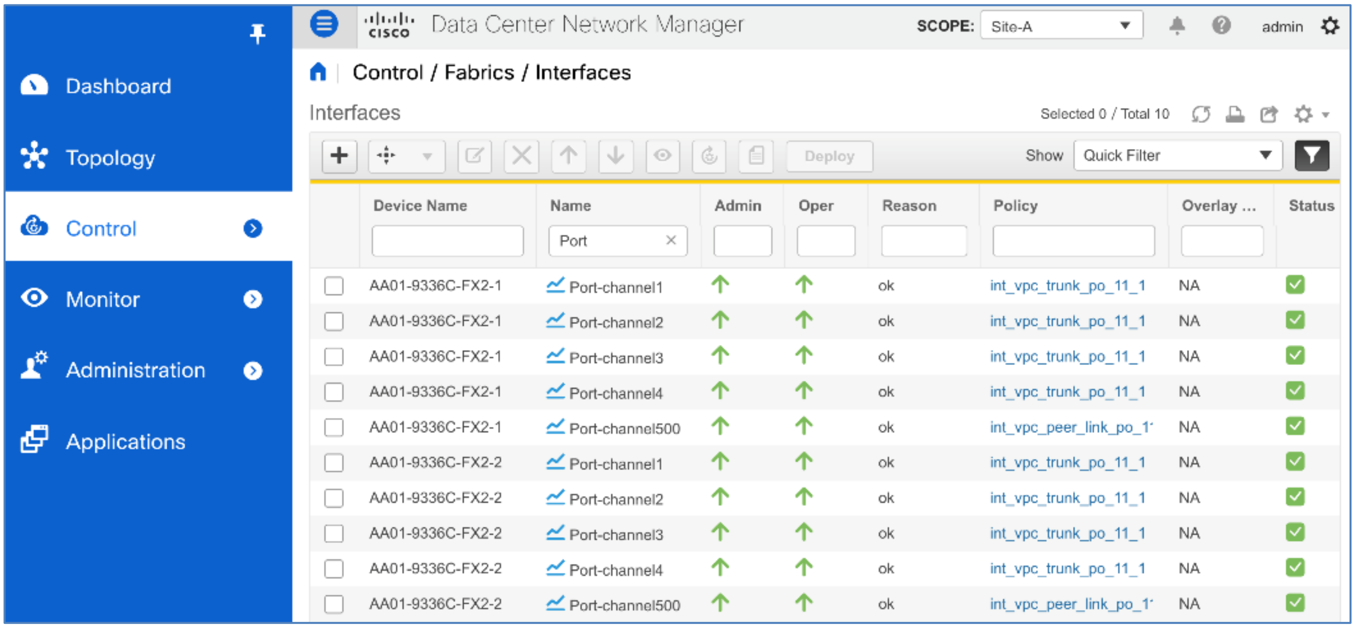

SCOPE: Site-A

Control / Fabrics / Interfaces

Interfaces

Selected 0 / Total 10

Show Quick Filter

	Device Name	Name	Admin	Oper	Reason	Policy	Overlay ...	Status
		Port						
□	AA01-9336C-FX2-1	Port-channel1	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-1	Port-channel2	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-1	Port-channel3	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-1	Port-channel4	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-1	Port-channel500	↑	↑	ok	int_vpc_peer_link_po_1'	NA	✓
□	AA01-9336C-FX2-2	Port-channel1	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel2	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel3	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel4	↑	↑	ok	int_vpc_trunk_po_11_1	NA	✓
□	AA01-9336C-FX2-2	Port-channel500	↑	↑	ok	int_vpc_peer_link_po_1'	NA	✓

18. From the left navigation bar, select **Control > Fabric > Fabric Builder.** Select **Site-A** fabric. From the **Actions** menu/window, select **Backup Now** to back up the Site-A fabric.

Enable Network Connectivity for FlexPod Infrastructure

To enable access to FlexPod infrastructure resources, namely compute and storage, the corresponding infrastructure networks must be first deployed in the VXLAN fabric in order to bring up the compute and storage infrastructure. The FlexPod infrastructure is isolated using a dedicated tenant/VRF (FPV-Foundation_VRF). Connectivity to external networks is also enabled directly from within FPV-Foundation_VRF. This tenant is not used for applications workloads hosted on the FlexPod Virtual Server Infrastructure (VSI) though it is used by management components such as VMware vCenter, NetApp VSC and so on. that is used to manage and operate the FlexPod VSI.

Setup Information

The configuration parameters for deploying the FlexPod infrastructure networks in Site-A data center fabric are provided below.

Table 12. Data Center Information

Scope	Site-A

Table 13. Infrastructure Tenant/VRF

VRF Name	VRF VLAN Name	VRF Interface Description	VRF Description
FPV-Foundation_VRF	FPV_Foundation_VRF_VLAN	FPV_Foundation_VRF_Interface	FPV_Foundation_VRF

Table 14. Infrastructure Networks - FPV-Foundation_VRF

Network Name	VLAN	VLAN Name	Forwarding	IP Subnet /Gateway*	VXLAN Network ID (VNID)	Notes
FPV-iSCSI-A_Network	3010	FPV-iSCSI-A_VLAN	Layer 2 Only	192.168.10.0/24	20000	ARP Suppression – N/A in L2-only mode
FPV-iSCSI-B_Network	3020	FPV-iSCSI-B_VLAN	Layer 2 Only	192.168.20.0/24	20001	“
FPV-InfraNFS_Network	3050	FPV-InfraNFS_VLAN	Layer 2 Only	192.168.50.0/24	20002	“
FPV-InBand-SiteA_Network	122	FPV-InBand-SiteA_VLAN FPV-InBand-SiteA_Interface	Layer 3	10.1.171.254/24*	20003	In-Band Management Network (e.g. ESXi hosts)
FPV-vMotion_Network	3000	FPV-vMotion_VLAN	Layer 2 Only	192.168.10.0/24	20004	“
FPV-CommonServices_Network	322	FPV-CommonServices_VLAN	Layer 3	10.3.171.254/24*	20005	Hosts VMware vCenter and NetApp VSC

* Gateway IP is configured only for L3 Forwarding and the default gateway is in the VXLAN Fabric

Deploy FlexPod Infrastructure Tenant in Cisco DCNM

To create the FlexPod Infrastructure Tenant in Cisco DCNM, use the **Setup Information** provided above to follow these steps:

1. Use a browser to navigate to Cisco DCNM’s GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > VRFs**. Click **OK** in the pop-window that complains about an **Unsupported Fabric Data Center selected**. Use **Scope:** to change the scope to **Site-A.**

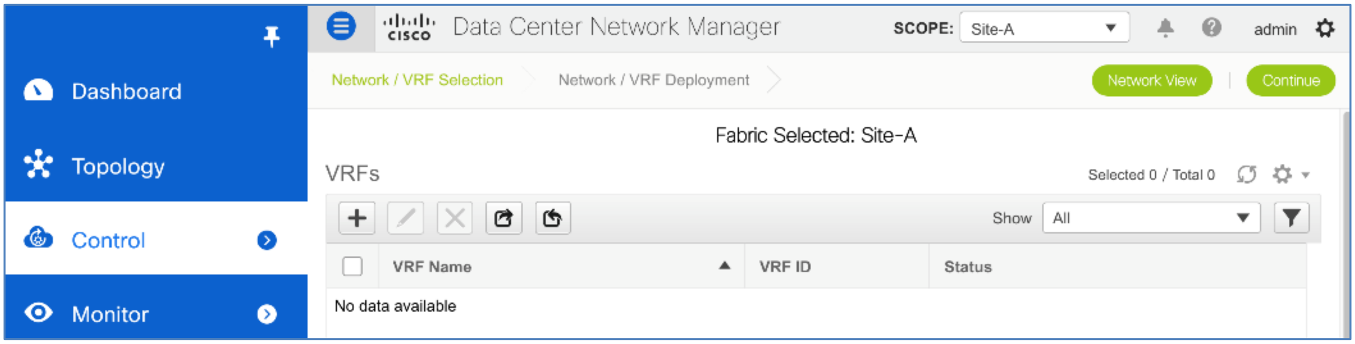

Fabric Selected: Site-A

☐	VRF Name ▲	VRF ID	Status
No data available			

3. Click on the **[+]** icon to deploy a new Tenant VRF for the FlexPod infrastructure traffic. Specify a **VRF VLAN Name, VRF Interface Description and VRF Description**. Click the **Create VRF** button.

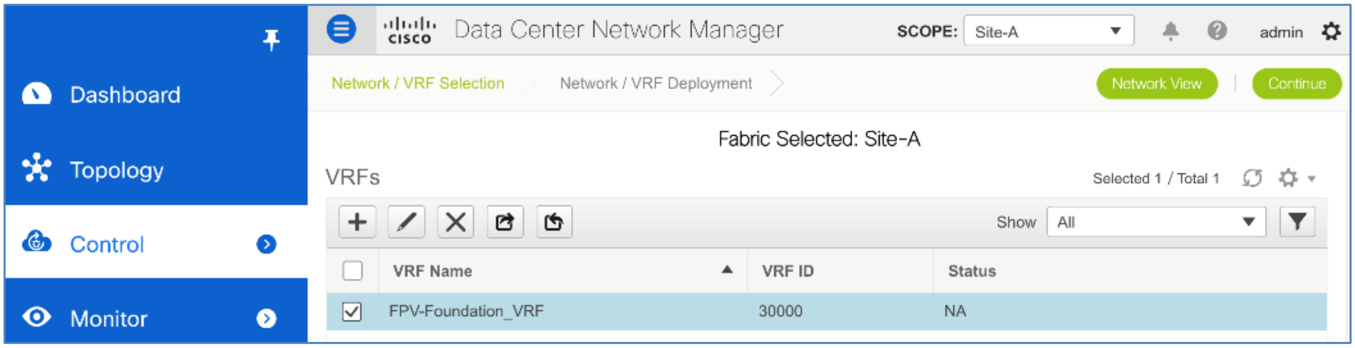

Fabric Selected: Site-A

☐	VRF Name ▲	VRF ID	Status
☑	FPV-Foundation_VRF	30000	NA

4. A small pop-up box will appear in the bottom-right corner to confirm that the VRF was created successfully. Click the **Continue** button.

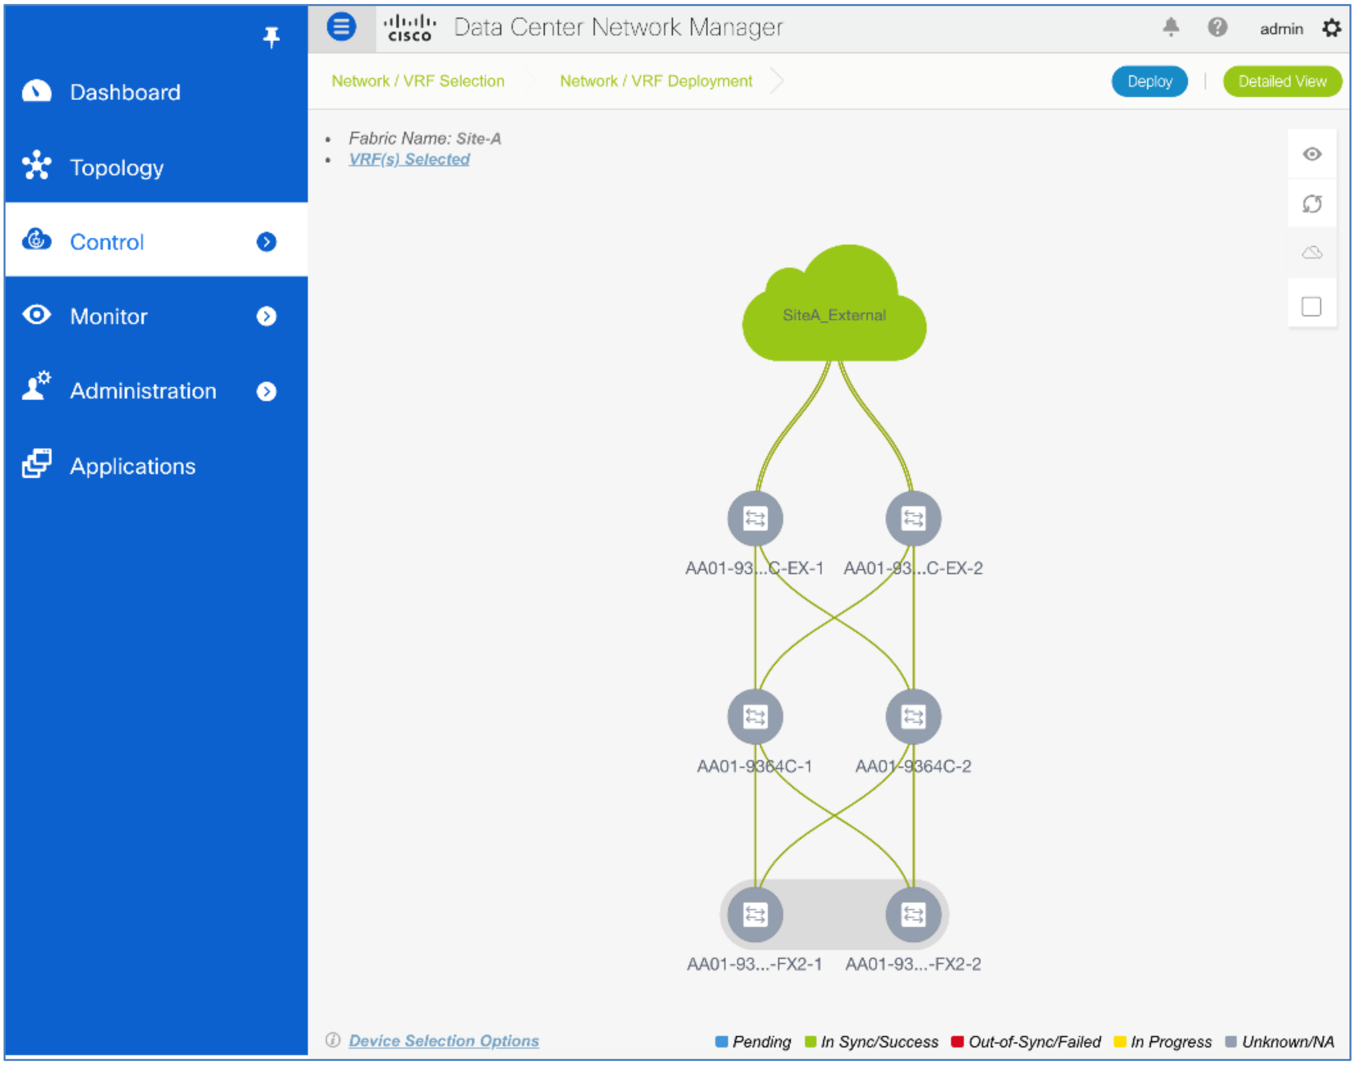

5. Click the **Detailed View** button.

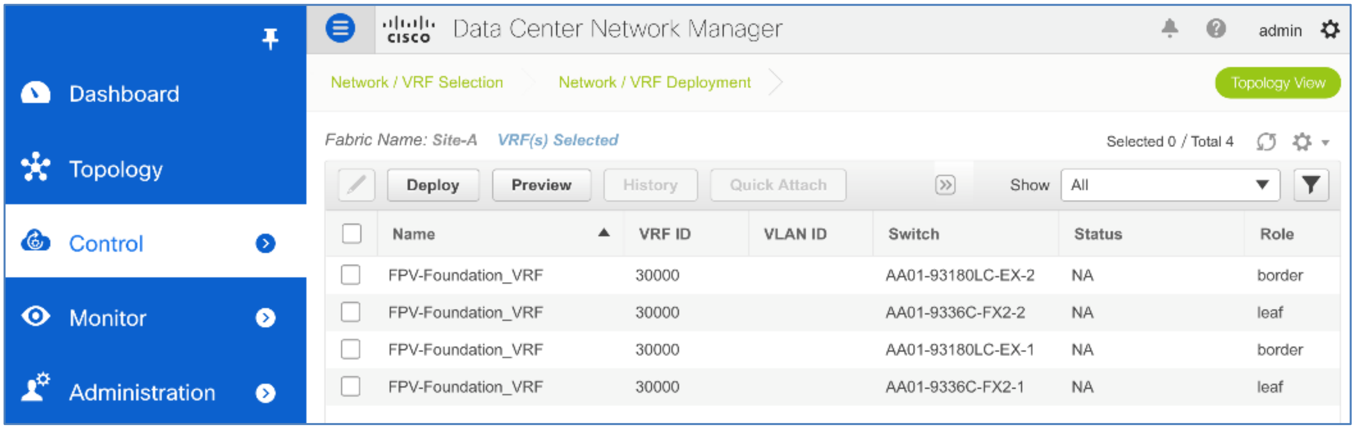

Fabric Name: Site-A VRF(s) Selected

Selected 0 / Total 4

□	Name	VRF ID	VLAN ID	Switch	Status	Role
□	FPV-Foundation_VRF	30000		AA01-93180LC-EX-2	NA	border
□	FPV-Foundation_VRF	30000		AA01-9336C-FX2-2	NA	leaf
□	FPV-Foundation_VRF	30000		AA01-93180LC-EX-1	NA	border
□	FPV-Foundation_VRF	30000		AA01-9336C-FX2-1	NA	leaf

6. Select the checkbox for all Leaf and Border switches in the list. Click the **Quick Attach** button.

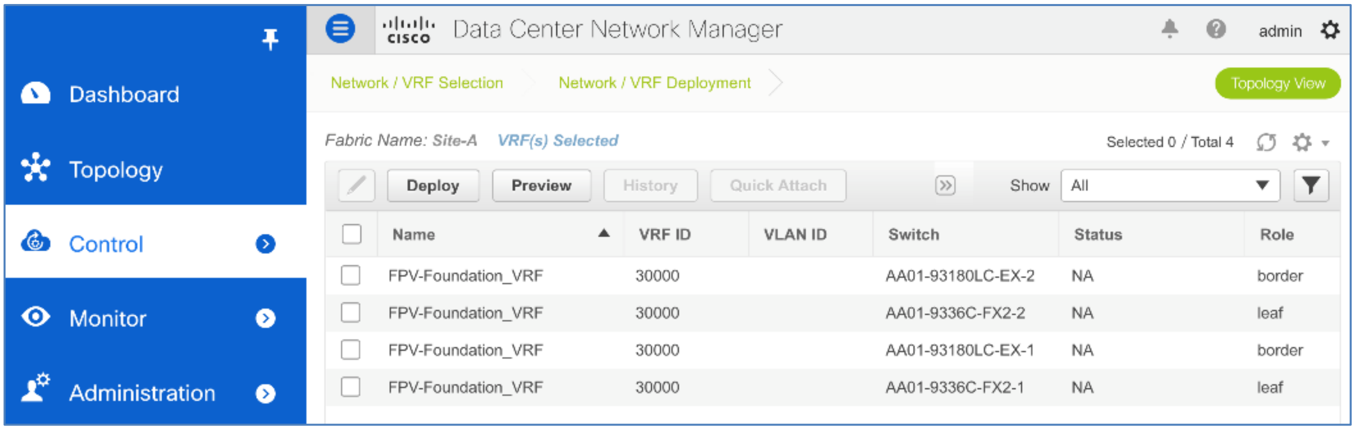

Fabric Name: Site-A VRF(s) Selected

Selected 0 / Total 4

□	Name	VRF ID	VLAN ID	Switch	Status	Role
□	FPV-Foundation_VRF	30000		AA01-93180LC-EX-2	NA	border
□	FPV-Foundation_VRF	30000		AA01-9336C-FX2-2	NA	leaf
□	FPV-Foundation_VRF	30000		AA01-93180LC-EX-1	NA	border
□	FPV-Foundation_VRF	30000		AA01-9336C-FX2-1	NA	leaf

7. Click **OK.**

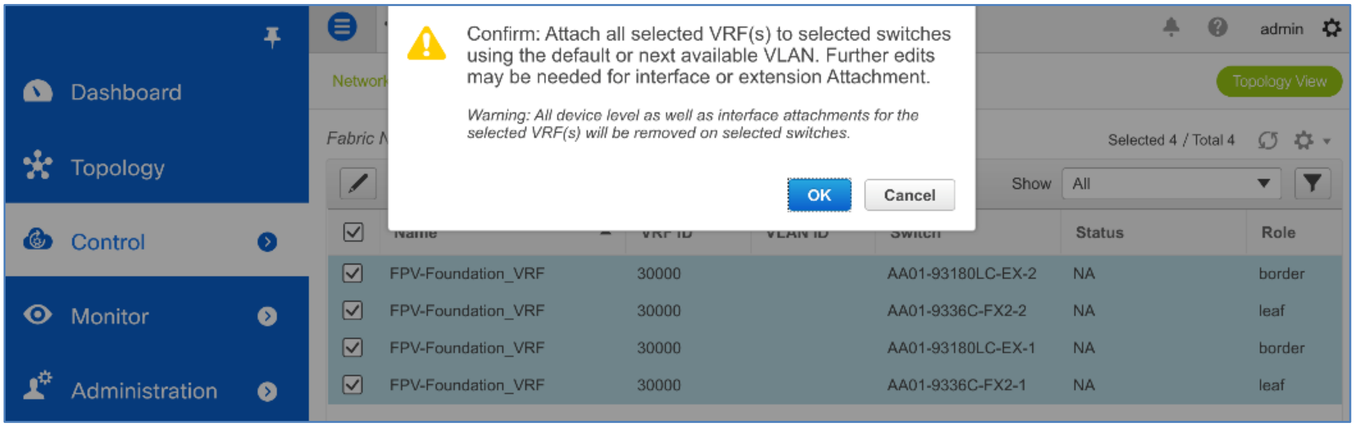

Confirm: Attach all selected VRF(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected VRF(s) will be removed on selected switches.

Selected 4 / Total 4

☑	Name	VRF ID	VLAN ID	Switch	Status	Role
☑	FPV-Foundation_VRF	30000		AA01-93180LC-EX-2	NA	border
☑	FPV-Foundation_VRF	30000		AA01-9336C-FX2-2	NA	leaf
☑	FPV-Foundation_VRF	30000		AA01-93180LC-EX-1	NA	border
☑	FPV-Foundation_VRF	30000		AA01-9336C-FX2-1	NA	leaf

8. Click the **Preview** button to preview the pending configuration on all the Leaf and Border switches. Click **X** to close the **Preview Configuration** window.

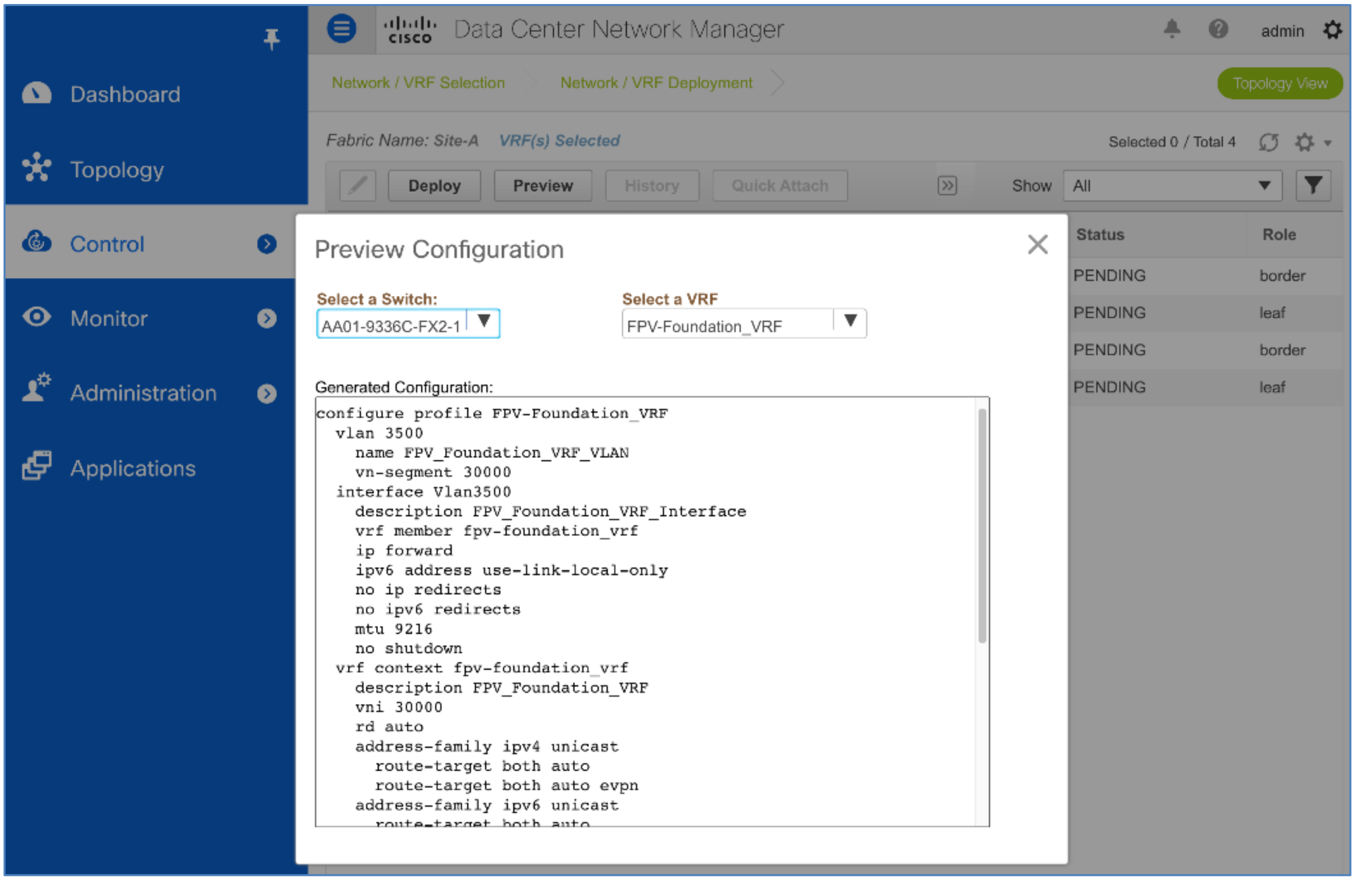

Fabric Name: Site-A VRF(s) Selected

Preview Configuration

Select a Switch: AA01-9336C-FX2-1

Select a VRF: FPV-Foundation_VRF

Generated Configuration:

```
configure profile FPV-Foundation_VRF
  vlan 3500
    name FPV_Foundation_VRF_VLAN
    vn-segment 30000
  interface Vlan3500
    description FPV_Foundation_VRF_Interface
    vrf member fpv-foundation_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-foundation_vrf
    description FPV_Foundation_VRF
    vni 30000
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
```

9. Click the **Deploy** button. Once the configuration is deployed, the **Status** go from **PENDING** to **IN PROGRESS** to **DEPLOYED.** Click the **Topology View** button.

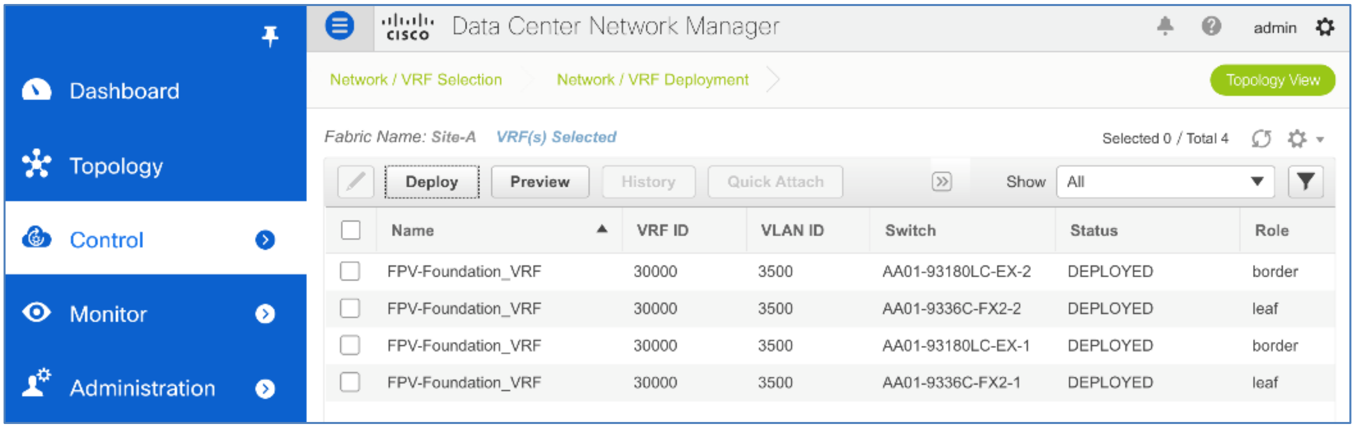

Fabric Name: Site-A VRF(s) Selected

□	Name	VRF ID	VLAN ID	Switch	Status	Role
□	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-2	DEPLOYED	border
□	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-2	DEPLOYED	leaf
□	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-1	DEPLOYED	border
□	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-1	DEPLOYED	leaf

10. In the **Topology View** to see the where the selected VRF is deployed in the Site-A topology.

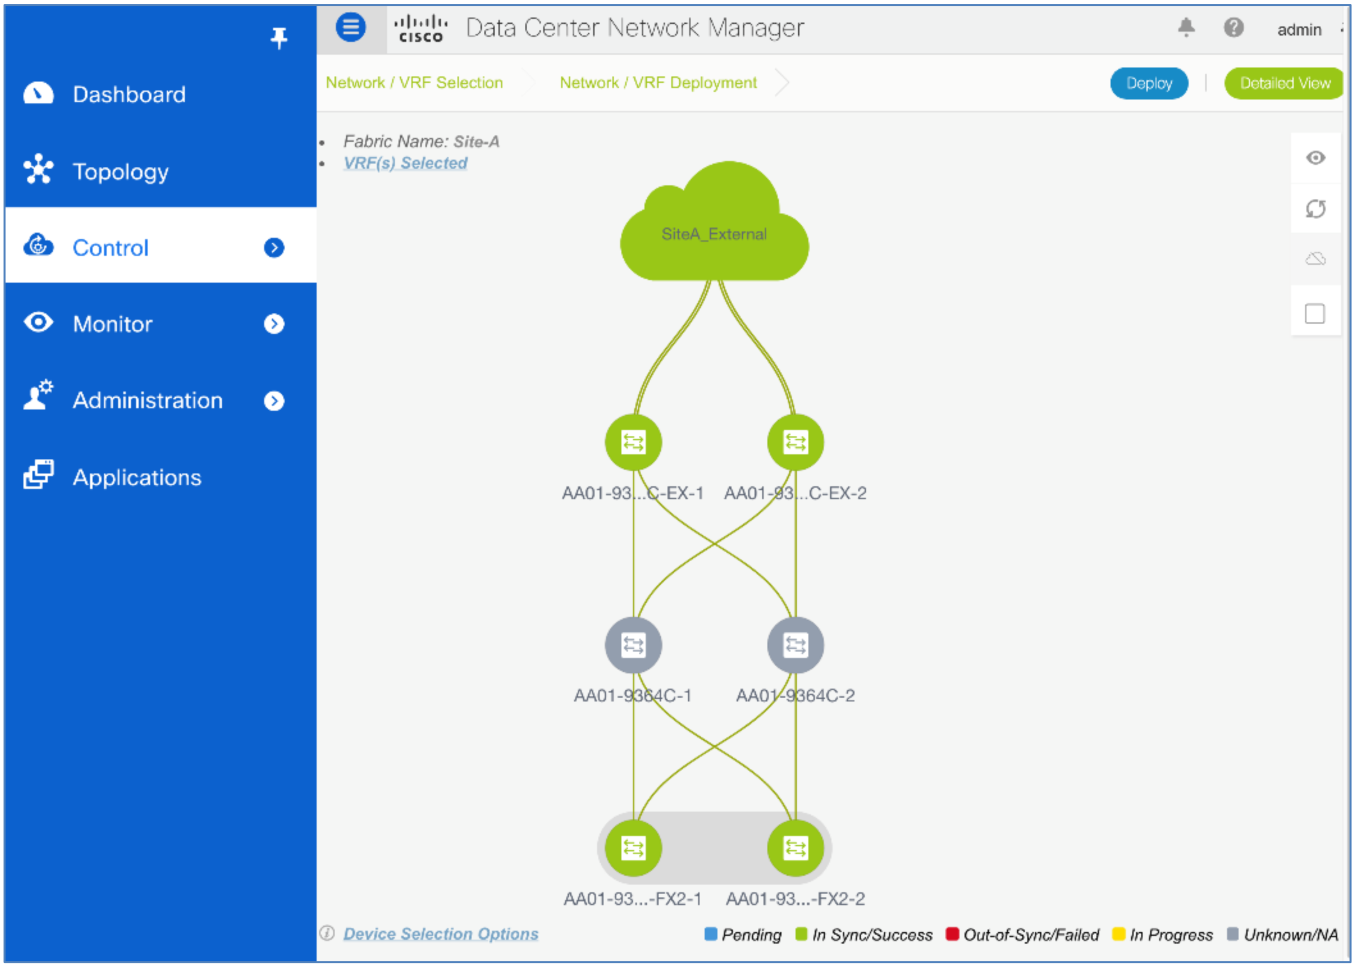

Deploy FlexPod Infrastructure Networks

To create the FlexPod Infrastructure Tenant networks in Cisco DCNM, use the **Setup Information** provided above to complete the steps in the upcoming sections.

Deploy FlexPod Storage Networks

The FlexPod storage networks, namely iSCSI-A, iSCSI-B and NFS networks are deployed in this design in **Layer 2 Only** mode with no gateway defined in the VXLAN fabric. To deploy the FlexPod Storage Networks, follow these steps:

1. Use a browser to navigate to Cisco DCNM’s GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**.

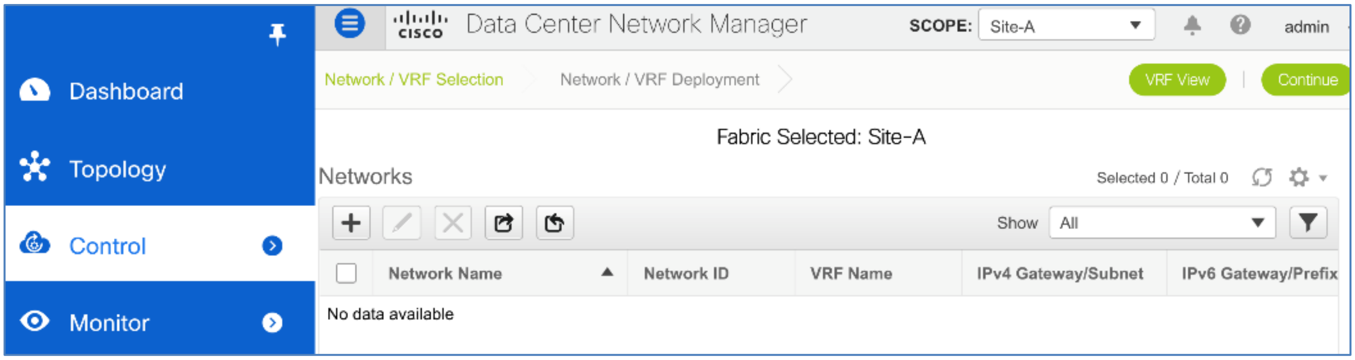

SCOPE: Site-A

Fabric Selected: Site-A

Networks

Selected 0 / Total 0

Show All

☐	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix

No data available

3. Click on the **[+]** icon to deploy a new Tenant network for the FlexPod infrastructure traffic. Specify a **Network Name.** Select the checkbox for **Layer 2 Only** mode. Specify a **VLAN ID**. In the **Network Profile > General** section, specify a **VLAN** name. Leave everything else as-is.

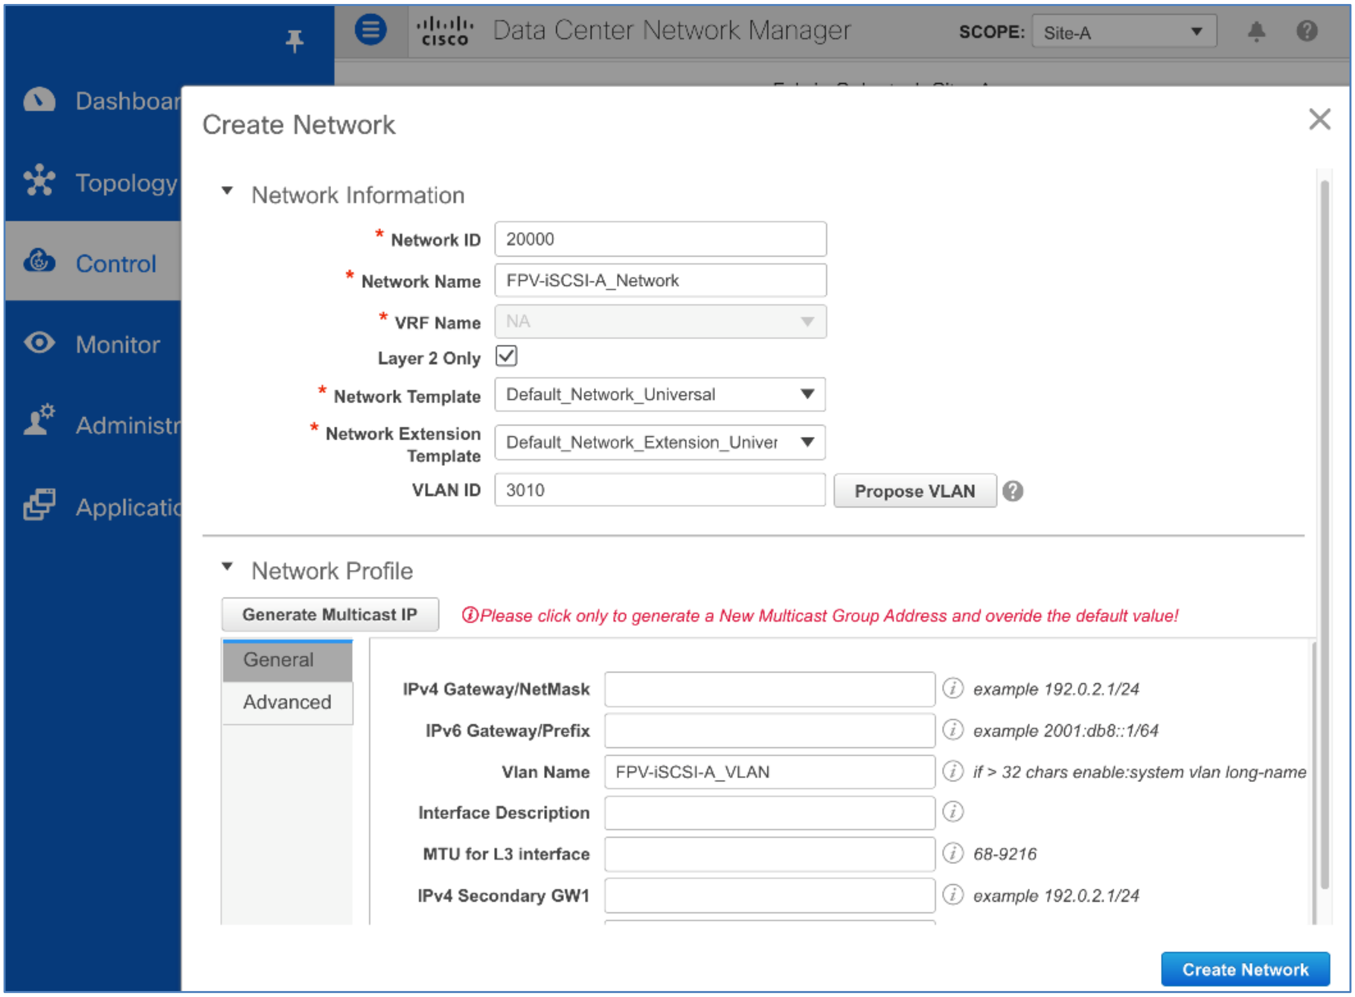

SCOPE: Site-A

Create Network

Network Information

* Network ID: 20000

* Network Name: FPV-iSCSI-A_Network

* VRF Name: NA

Layer 2 Only: ☑

* Network Template: Default_Network_Universal

* Network Extension Template: Default_Network_Extension_Univer

VLAN ID: 3010

Network Profile

ⓘPlease click only to generate a New Multicast Group Address and overide the default value!

IPv4 Gateway/NetMask: ⓘ example 192.0.2.1/24

IPv6 Gateway/Prefix: ⓘ example 2001:db8::1/64

Vlan Name: FPV-iSCSI-A_VLAN ⓘ if > 32 chars enable:system vlan long-name

Interface Description: ⓘ

MTU for L3 interface: ⓘ 68-9216

IPv4 Secondary GW1: ⓘ example 192.0.2.1/24

4. In the **Network Profile > Advanced** section, leave everything as-is.

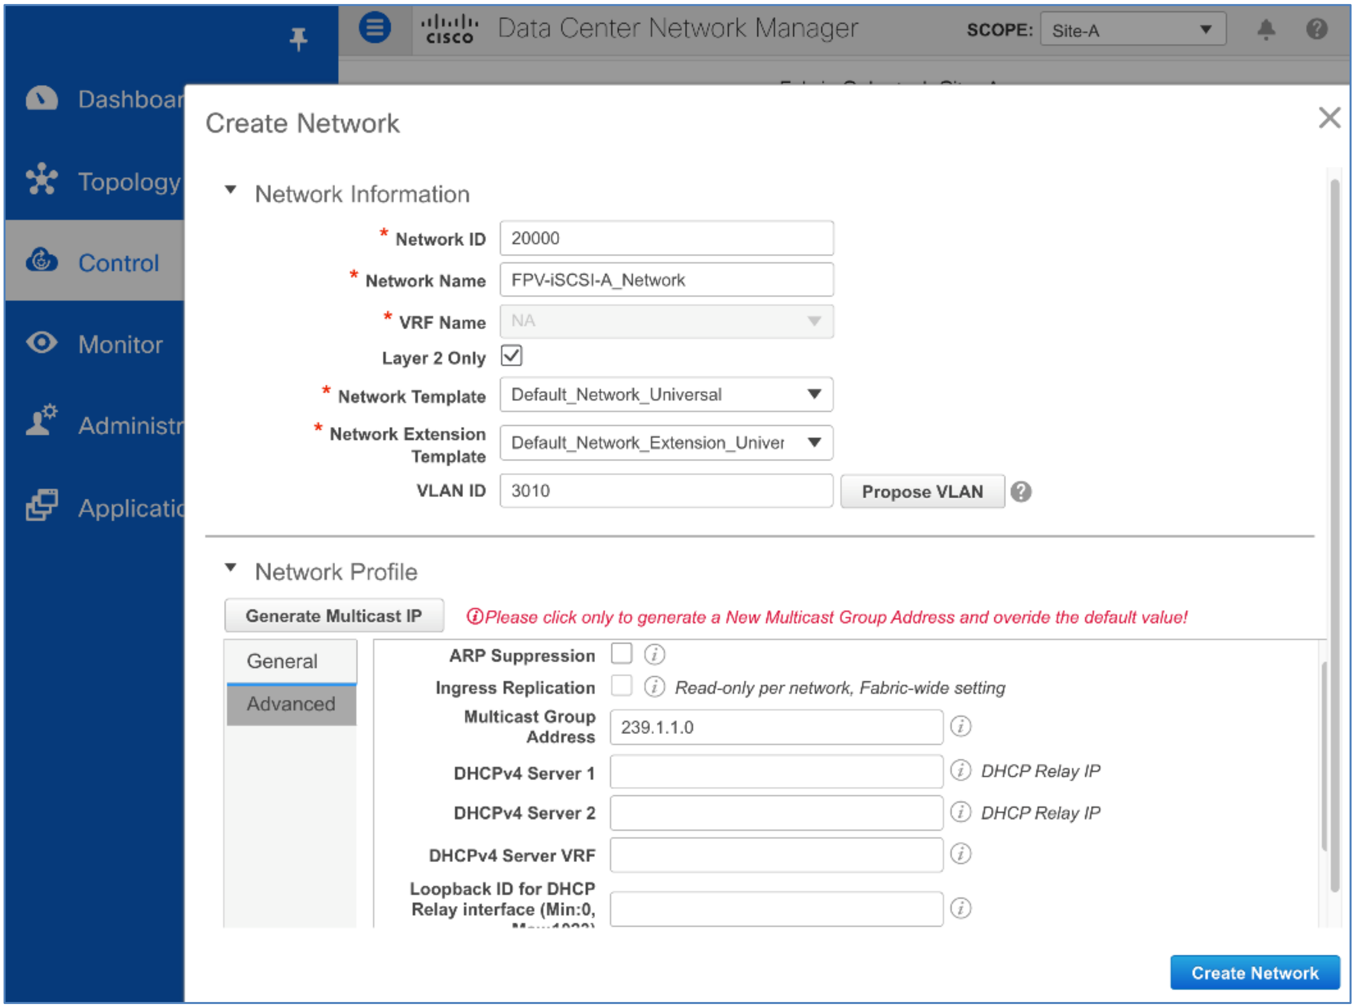

Create Network

Network Information

* Network ID: 20000

* Network Name: FPV-iSCSI-A_Network

* VRF Name: NA

Layer 2 Only: ☑

* Network Template: Default_Network_Universal

* Network Extension Template: Default_Network_Extension_Univer

VLAN ID: 3010

Network Profile

ⓘPlease click only to generate a New Multicast Group Address and overide the default value!

ARP Suppression: □

Ingress Replication: □ Read-only per network, Fabric-wide setting

Multicast Group Address: 239.1.1.0

DHCPv4 Server 1: DHCP Relay IP

DHCPv4 Server 2: DHCP Relay IP

DHCPv4 Server VRF:

Loopback ID for DHCP Relay interface (Min:0,

5. Click the **Create Network** button.

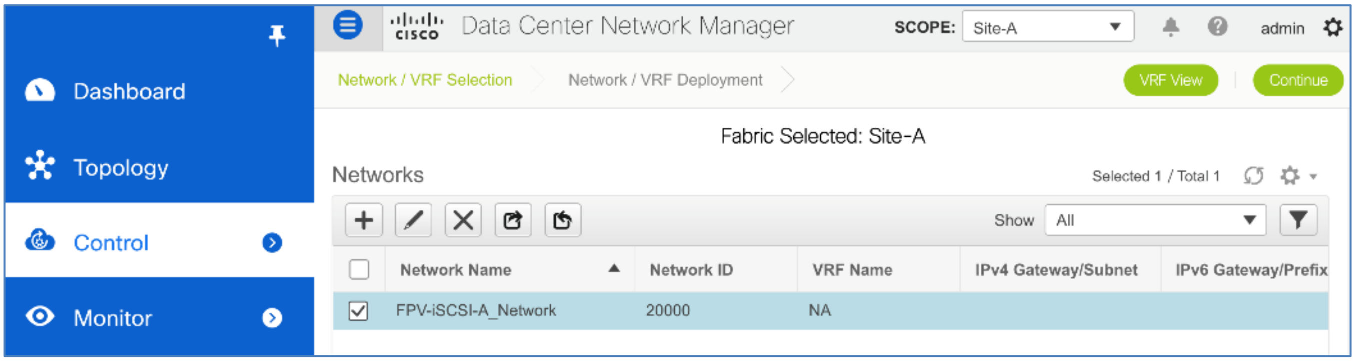

Fabric Selected: Site-A

Networks

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☑	FPV-iSCSI-A_Network	20000	NA		

6. Click the **Continue** button.

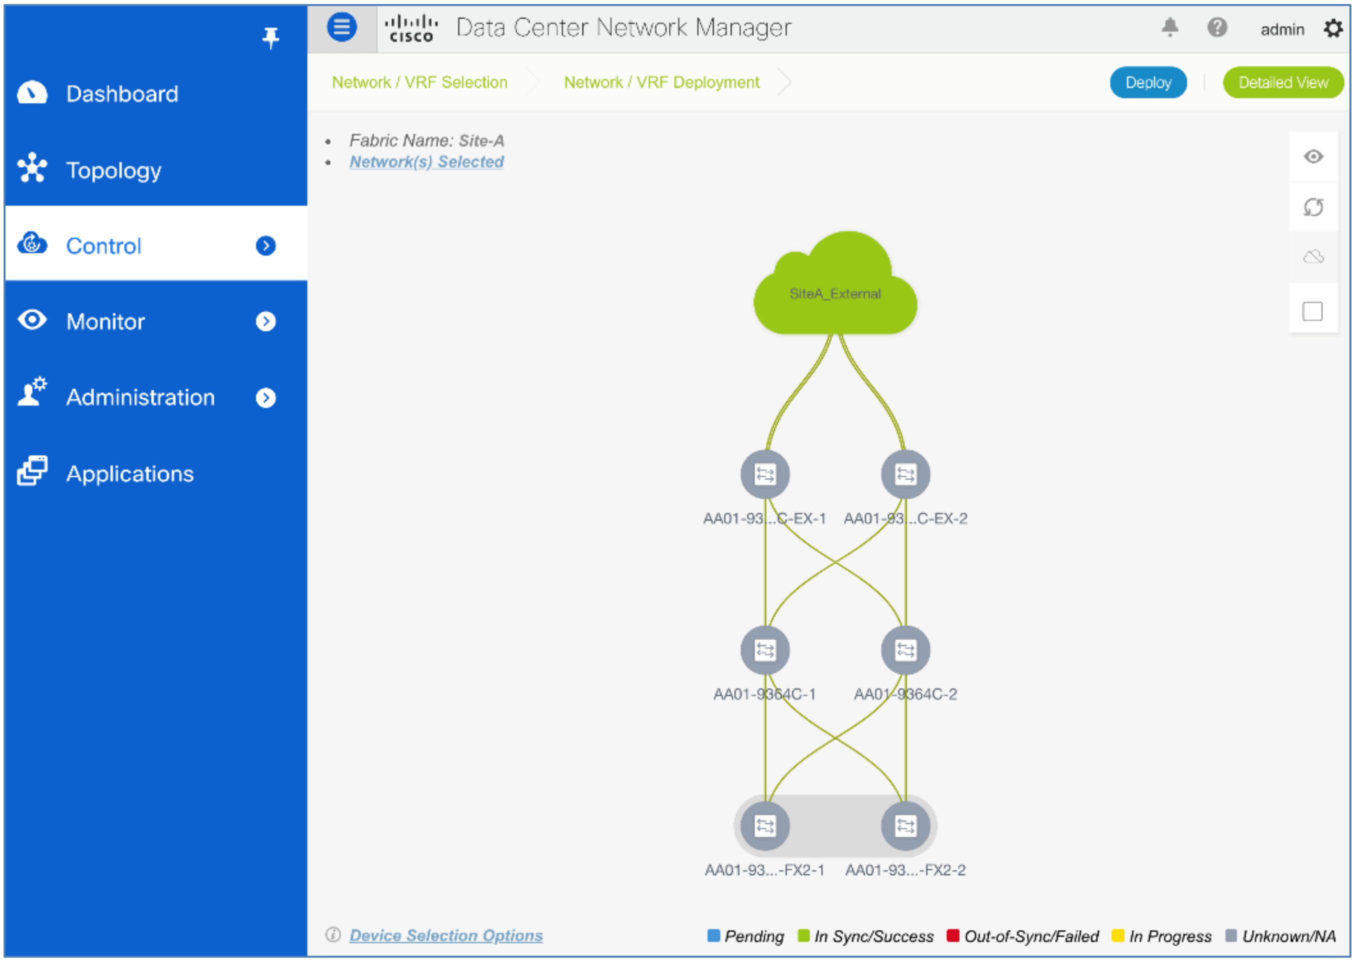

7. Click the **Detailed View** button. Select the Leaf switches where these networks need to be deployed. Click the **Quick Attach** button. Click **OK.**

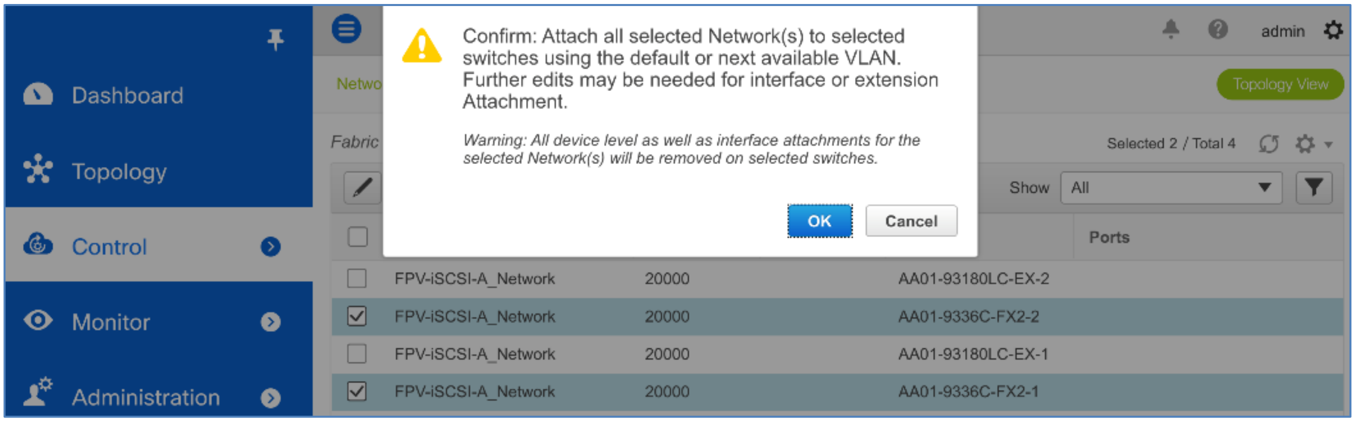

Confirm: Attach all selected Network(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected Network(s) will be removed on selected switches.

OK Cancel

□				Ports
□	FPV-iSCSI-A_Network	20000	AA01-93180LC-EX-2	
☑	FPV-iSCSI-A_Network	20000	AA01-9336C-FX2-2	
□	FPV-iSCSI-A_Network	20000	AA01-93180LC-EX-1	
☑	FPV-iSCSI-A_Network	20000	AA01-9336C-FX2-1	

8. Click **Preview** to view pending changes. Click the **X** to close the window.

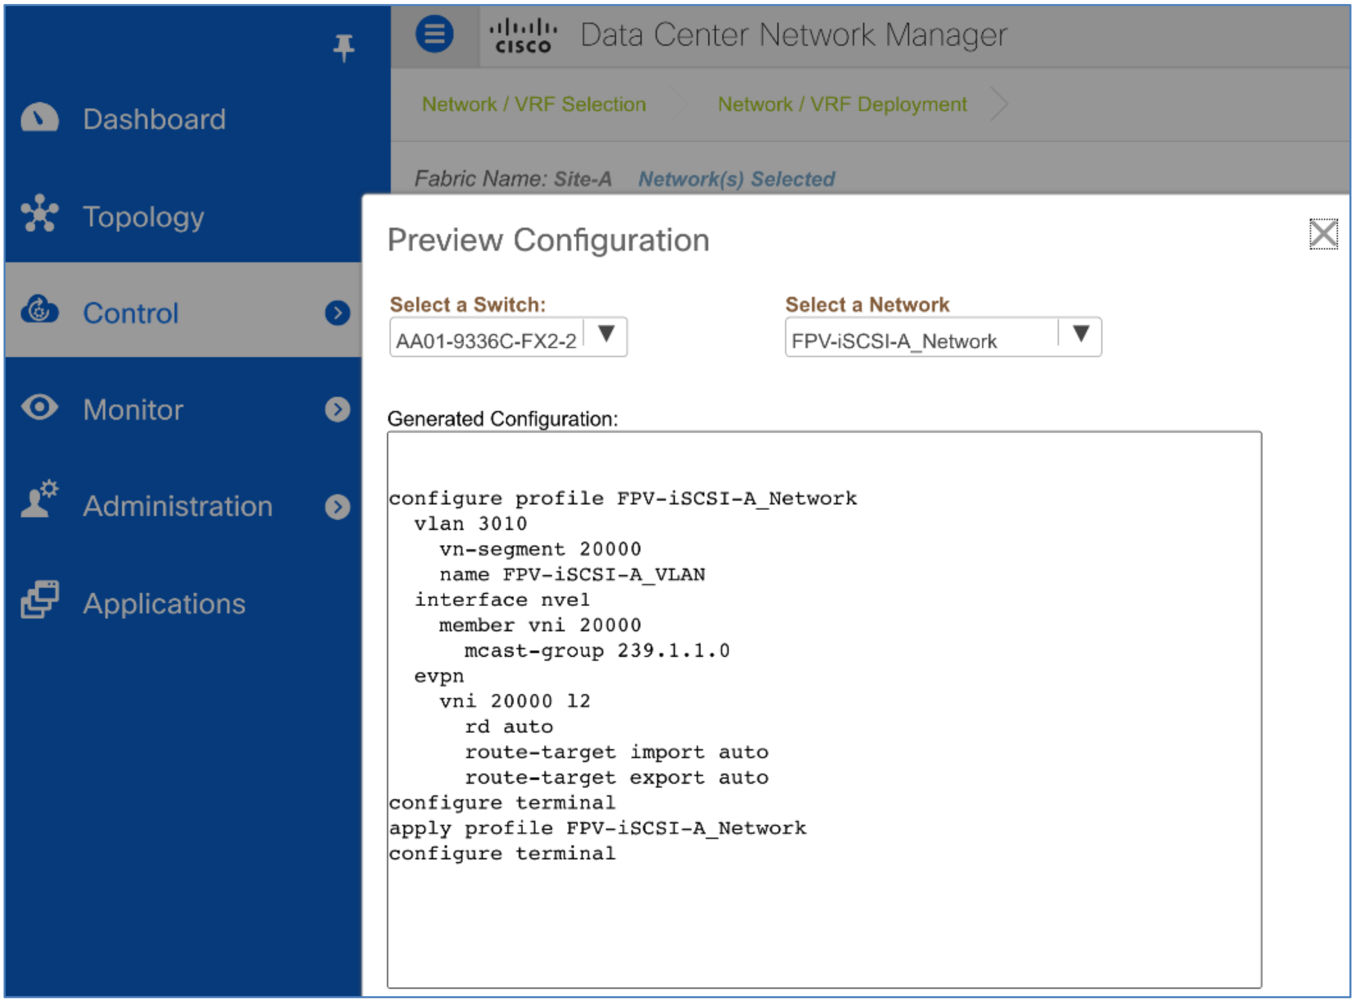

Select a Switch: AA01-9336C-FX2-2

Select a Network: FPV-iSCSI-A_Network

Generated Configuration:

```
configure profile FPV-iSCSI-A_Network
  vlan 3010
    vn-segment 20000
    name FPV-iSCSI-A_VLAN
  interface nve1
    member vni 20000
      mcast-group 239.1.1.0
  evpn
    vni 20000 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
apply profile FPV-iSCSI-A_Network
configure terminal
```

9. Click the **Deploy** button. The status should go from **PENDING** to **DEPLOYED** in the **Status** column for the two Leaf switches. Scroll to the right as needed to see all columns in this view.

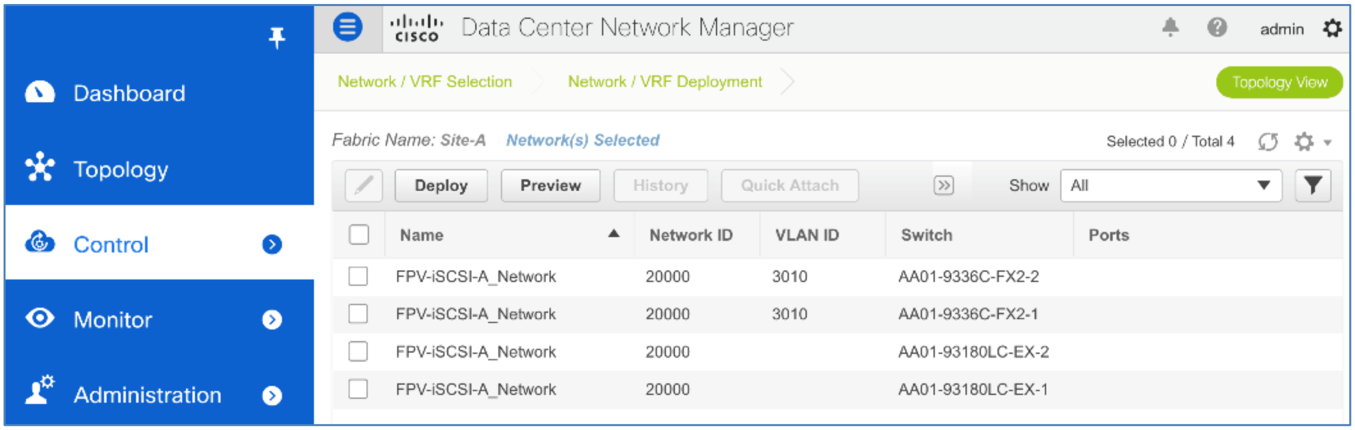

	Name	Network ID	VLAN ID	Switch	Ports
□	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-2	
□	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	
□	FPV-iSCSI-A_Network	20000		AA01-93180LC-EX-2	
□	FPV-iSCSI-A_Network	20000		AA01-93180LC-EX-1	

10. Repeat steps 1-9 to deploy the second iSCSI network.

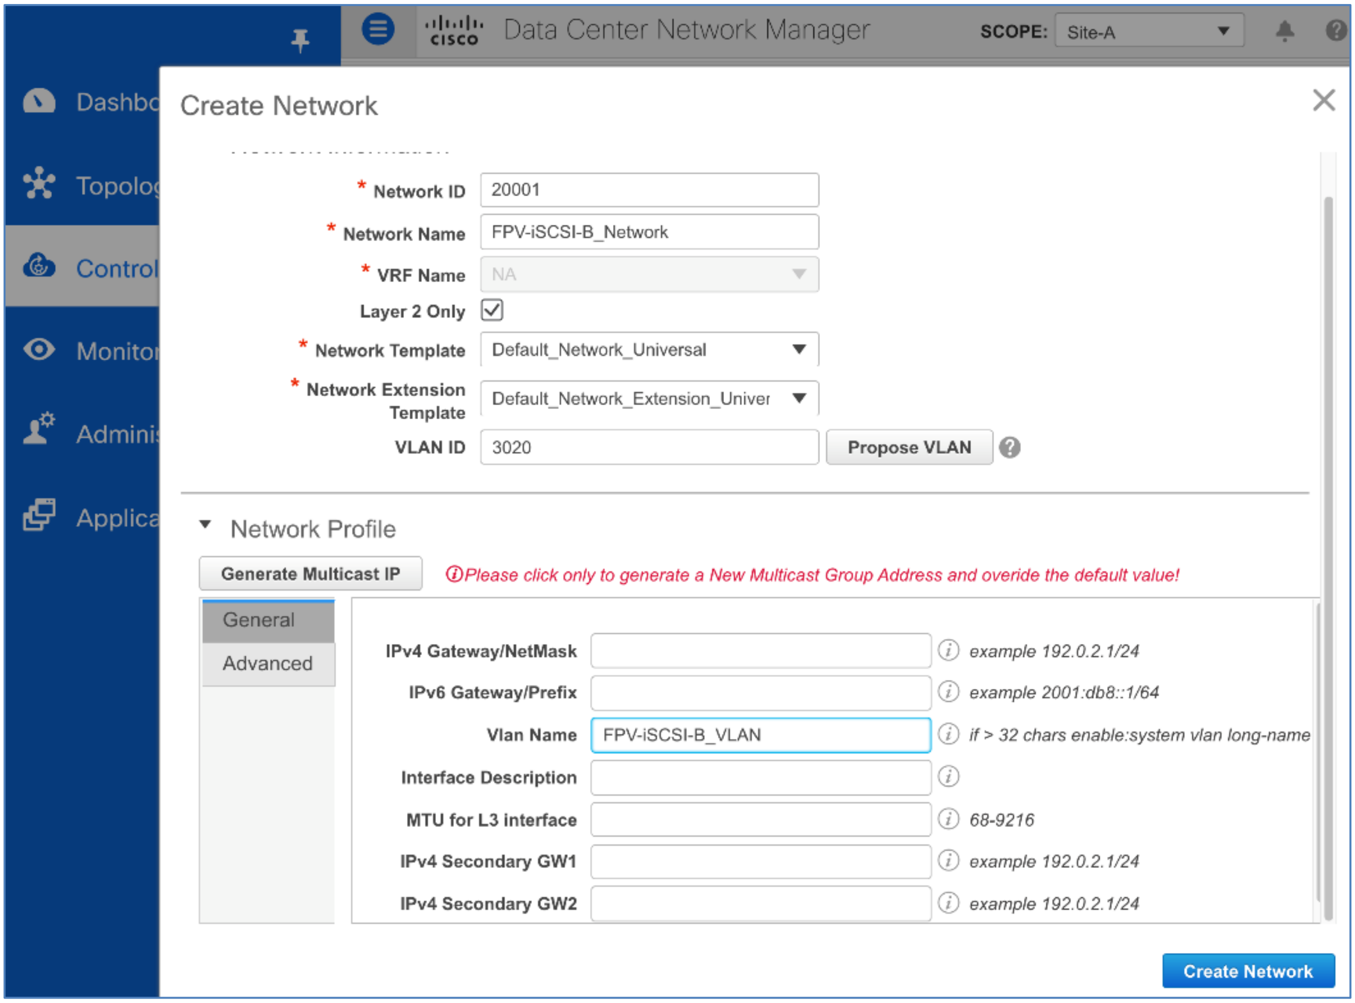

Create Network

* Network ID	20001
* Network Name	FPV-iSCSI-B_Network
* VRF Name	NA
Layer 2 Only	☑
* Network Template	Default_Network_Universal
* Network Extension Template	Default_Network_Extension_Univer
VLAN ID	3020

Propose VLAN

Network Profile

Generate Multicast IP

ⓘPlease click only to generate a New Multicast Group Address and overide the default value!

General

Advanced

IPv4 Gateway/NetMask		example 192.0.2.1/24
IPv6 Gateway/Prefix		example 2001:db8::1/64
Vlan Name	FPV-iSCSI-B_VLAN	if > 32 chars enable:system vlan long-name
Interface Description		
MTU for L3 interface		68-9216
IPv4 Secondary GW1		example 192.0.2.1/24
IPv4 Secondary GW2		example 192.0.2.1/24

Create Network

11. Click **Deploy** to deploy the configuration. The status should go from **PENDING** to **IN PROGRESS** to **DEPLOYED** in the **Status** column for the two Leaf switches.

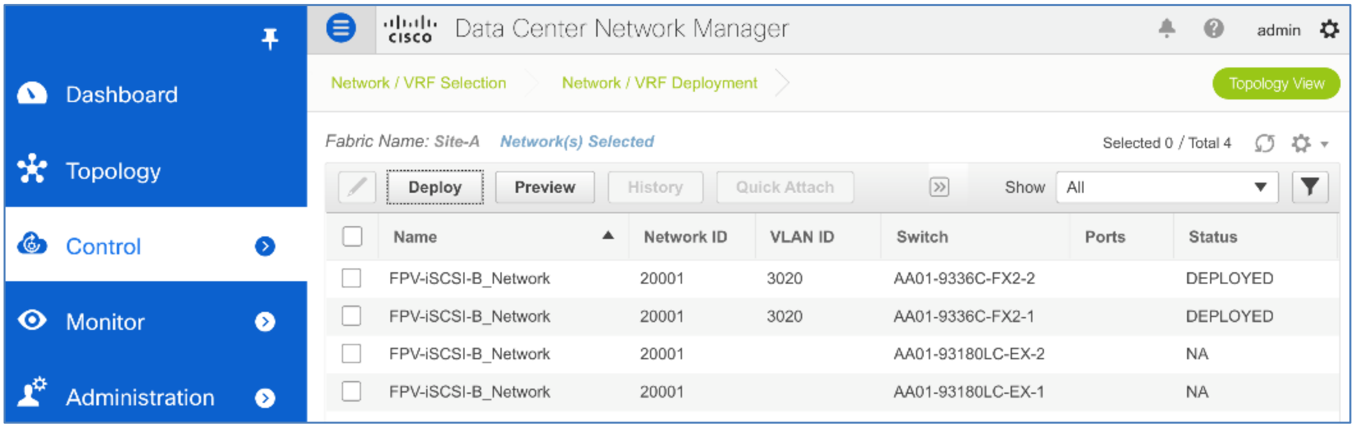

Fabric Name: Site-A *Network(s) Selected*

Selected 0 / Total 4

	Name	Network ID	VLAN ID	Switch	Ports	Status
□	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-2		DEPLOYED
□	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1		DEPLOYED
□	FPV-iSCSI-B_Network	20001		AA01-93180LC-EX-2		NA
□	FPV-iSCSI-B_Network	20001		AA01-93180LC-EX-1		NA

12. In the **Topology View** to see the where the selected network is deployed in the Site-A topology.

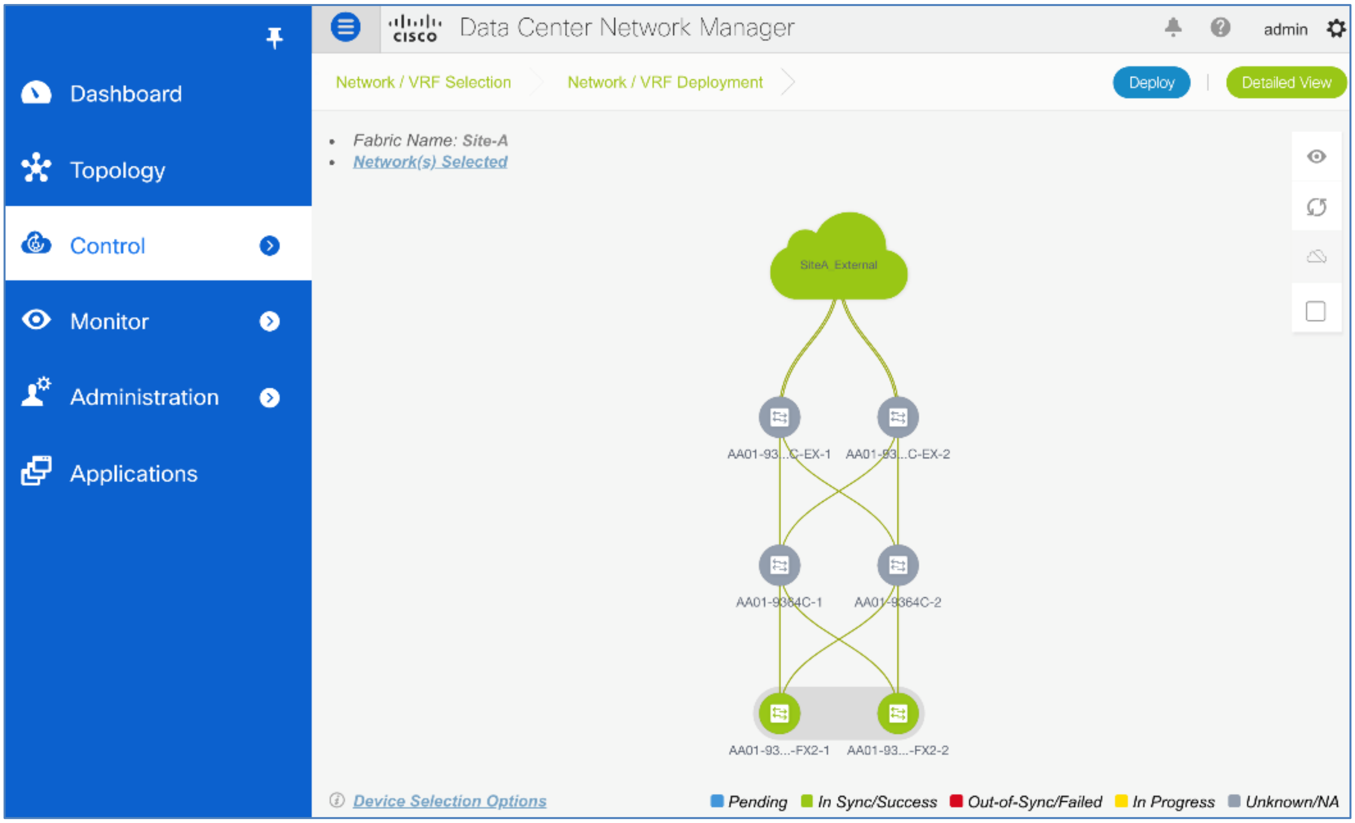

- *Fabric Name: Site-A*
- *Network(s) Selected*

🟦 *Pending* 🟩 *In Sync/Success* 🟥 *Out-of-Sync/Failed* 🟨 *In Progress* ⬜ *Unknown/NA*

13. Repeat steps 1-12 to deploy the NFS network.

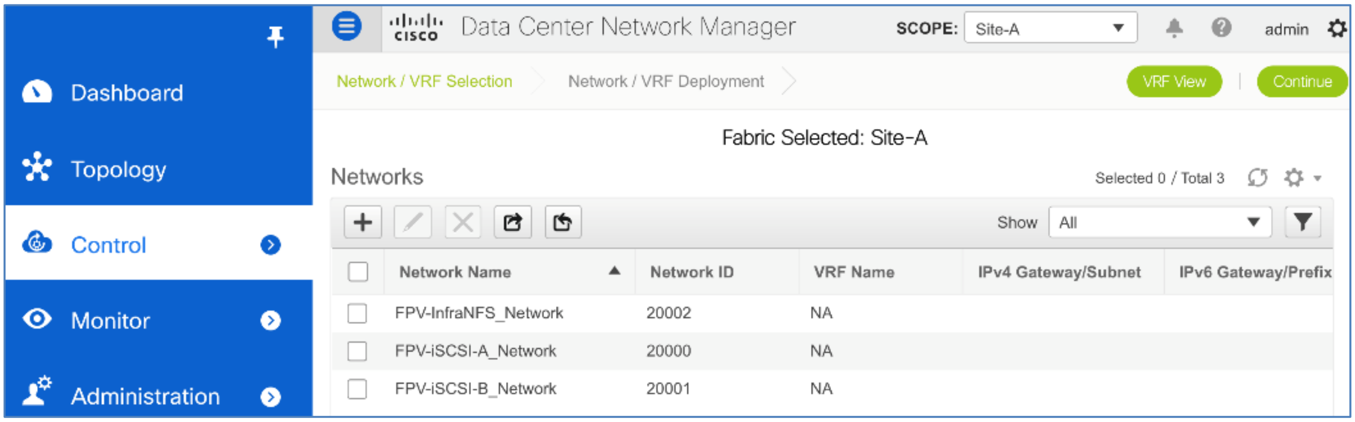

SCOPE: Site-A

Fabric Selected: Site-A

Selected 0 / Total 3

Show All

□	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
□	FPV-InfraNFS_Network	20002	NA		
□	FPV-iSCSI-A_Network	20000	NA		
□	FPV-iSCSI-B_Network	20001	NA		

Deploy FlexPod In-Band Management Network

The FlexPod In-Band Management network is deployed in this design in Layer 3 mode where the traffic is Layer 3 forwarded by the fabric and the gateway is a distributed anycast gateway in the VXLAN fabric. This is unlike the previous storage networks that are deployed in **Layer 2 Only** mode with no gateway defined in the fabric.

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**.

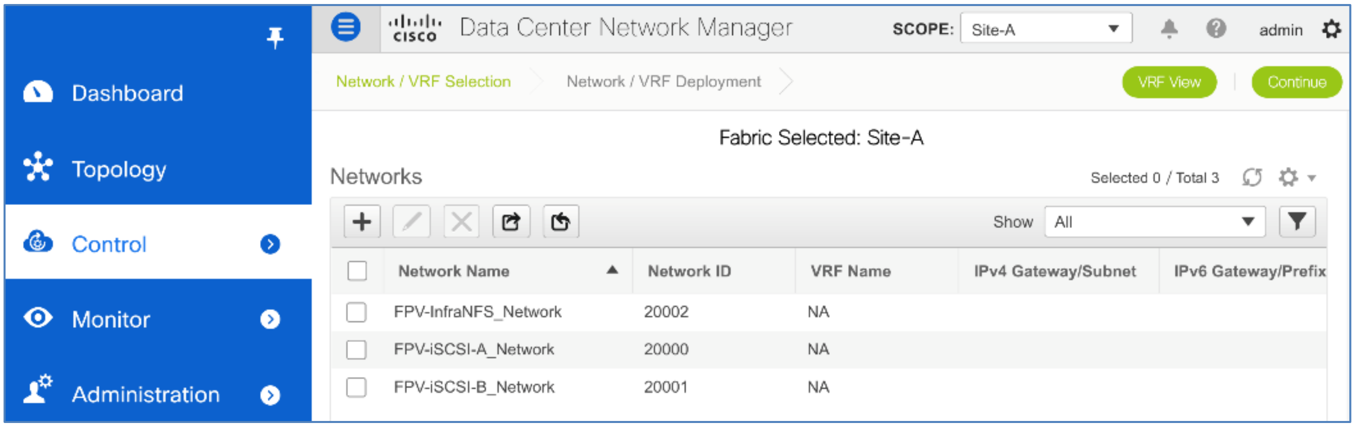

Fabric Selected: Site-A

Networks

Selected 0 / Total 3

☐	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☐	FPV-InfraNFS_Network	20002	NA		
☐	FPV-iSCSI-A_Network	20000	NA		
☐	FPV-iSCSI-B_Network	20001	NA		

3. Click on the **[+]** icon to deploy a new Tenant network for the FlexPod infrastructure traffic. Specify a **Network Name, VRF Name.** In this deployment, we are specifying the **VLAN ID** we specifically want to use but you can optionally let DCNM pick up one from the defined pool in **Fabric Settings.** Specify a **VLAN ID**. In the **Network Profile > General** section of the window, specify a **IPv4 Gateway/Network, VLAN Name, Interface Description** and **MTU**. Leave everything else as-is.

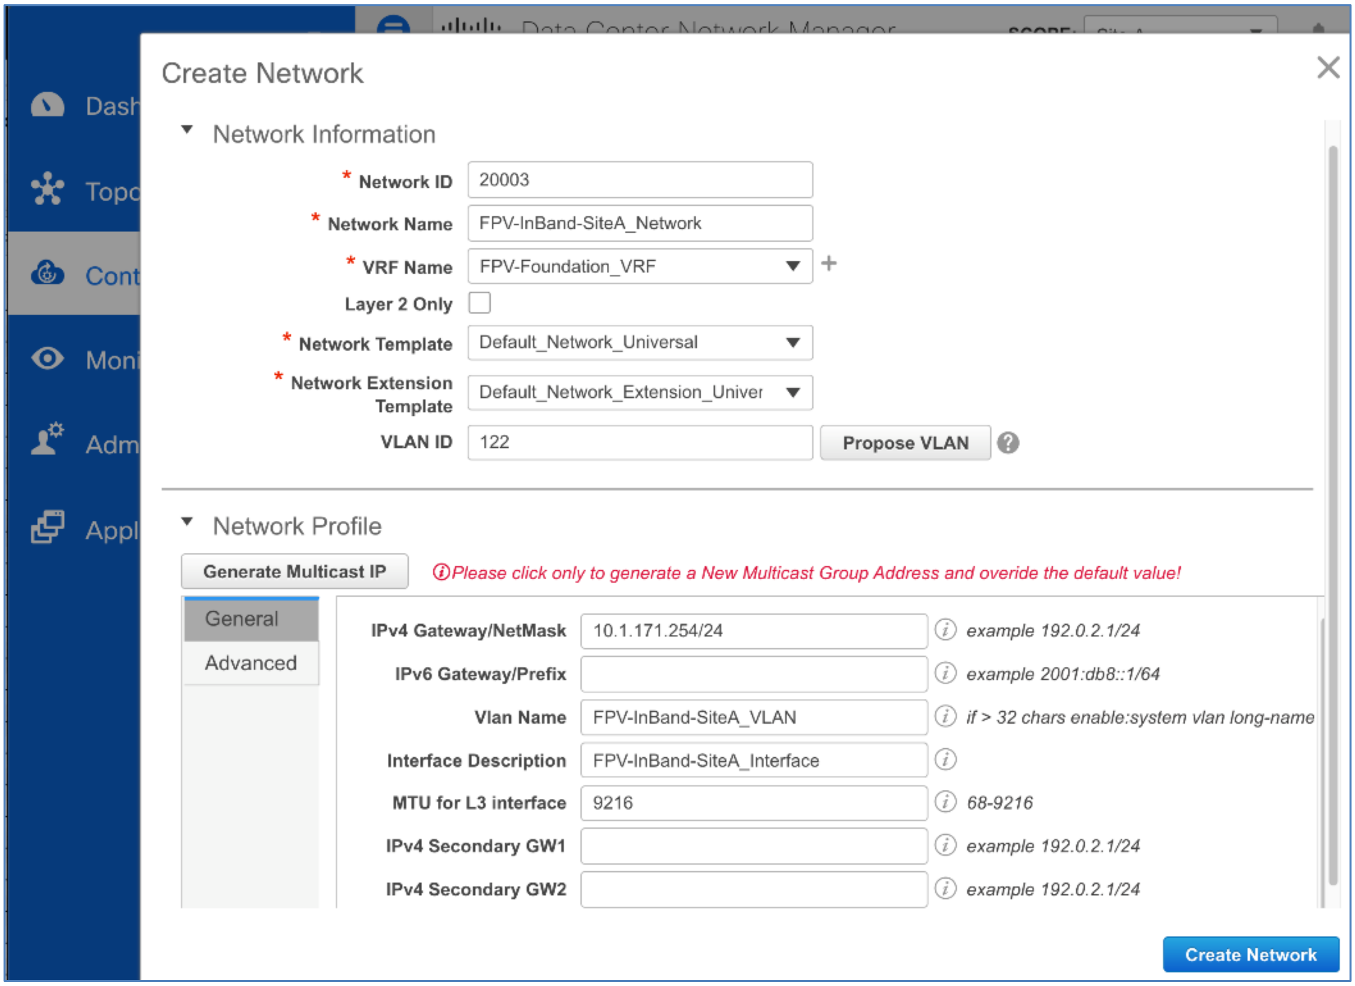

Create Network

Network Information

* Network ID	20003
* Network Name	FPV-InBand-SiteA_Network
* VRF Name	FPV-Foundation_VRF
Layer 2 Only	☐
* Network Template	Default_Network_Universal
* Network Extension Template	Default_Network_Extension_Univer
VLAN ID	122

Network Profile

ⓘPlease click only to generate a New Multicast Group Address and overide the default value!

IPv4 Gateway/NetMask	10.1.171.254/24	example 192.0.2.1/24
IPv6 Gateway/Prefix		example 2001:db8::1/64
Vlan Name	FPV-InBand-SiteA_VLAN	if > 32 chars enable:system vlan long-name
Interface Description	FPV-InBand-SiteA_Interface	
MTU for L3 interface	9216	68-9216
IPv4 Secondary GW1		example 192.0.2.1/24
IPv4 Secondary GW2		example 192.0.2.1/24

4. In the **Network Profile > Advanced** section of the window, enable **ARP Suppression.** Leave everything as-is.

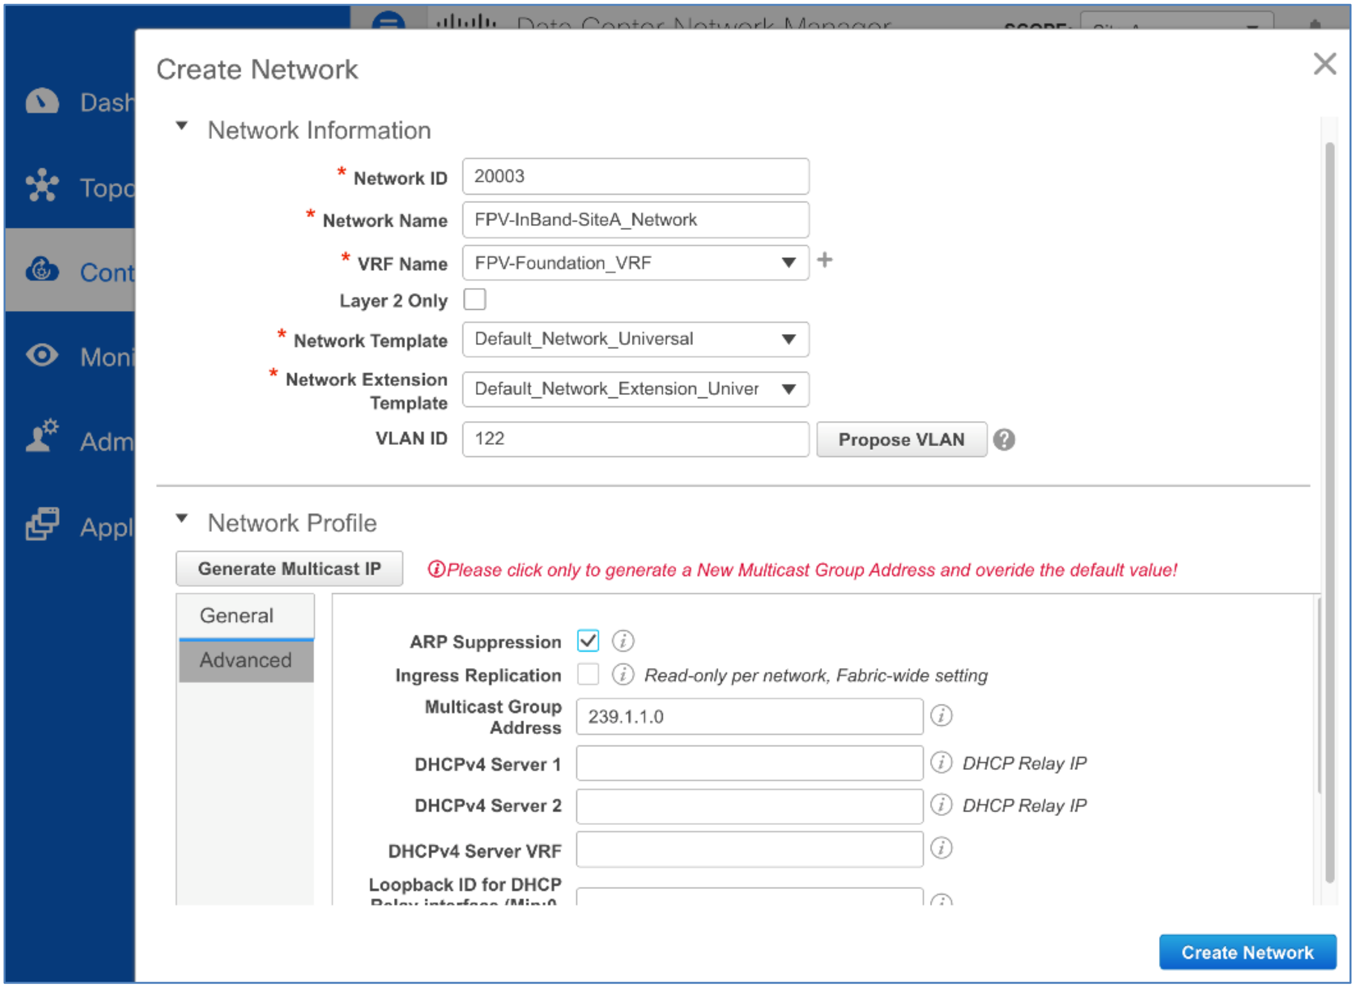

Create Network

Network Information

* Network ID	20003
* Network Name	FPV-InBand-SiteA_Network
* VRF Name	FPV-Foundation_VRF
Layer 2 Only	□
* Network Template	Default_Network_Universal
* Network Extension Template	Default_Network_Extension_Univer
VLAN ID	122

Network Profile

ⓘPlease click only to generate a New Multicast Group Address and overide the default value!

ARP Suppression	☑	
Ingress Replication	□	Read-only per network, Fabric-wide setting
Multicast Group Address	239.1.1.0	
DHCPv4 Server 1		DHCP Relay IP
DHCPv4 Server 2		DHCP Relay IP
DHCPv4 Server VRF		
Loopback ID for DHCP Relay interface (Min:0		

5. Click the **Create Network** button.

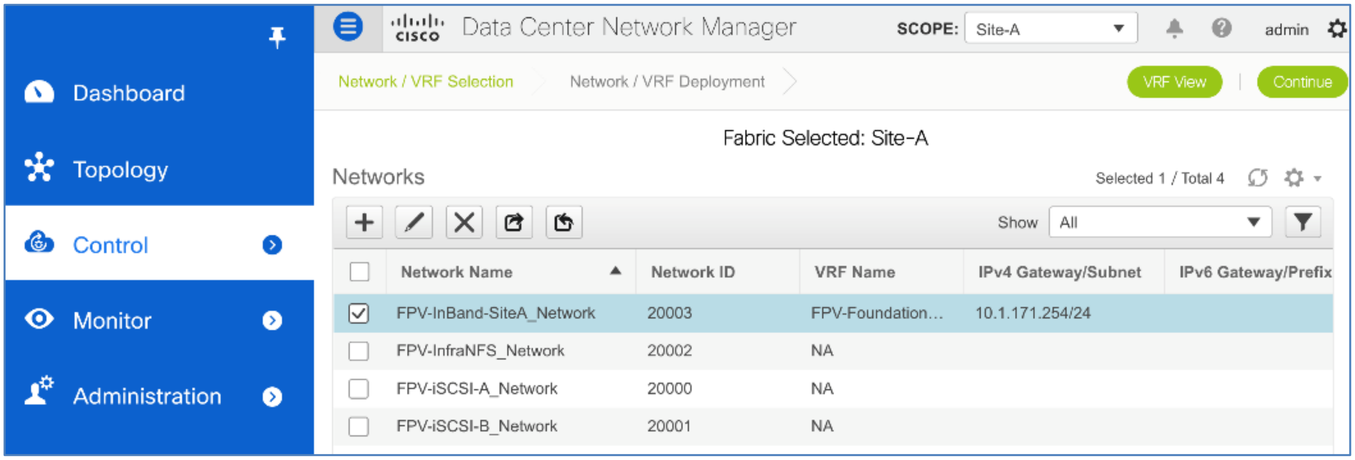

Fabric Selected: Site-A

Networks

Selected 1 / Total 4

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☑	FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
□	FPV-InfraNFS_Network	20002	NA		
□	FPV-iSCSI-A_Network	20000	NA		
□	FPV-iSCSI-B_Network	20001	NA		

6. Click the **Continue** button. Click the **Detailed View** button.

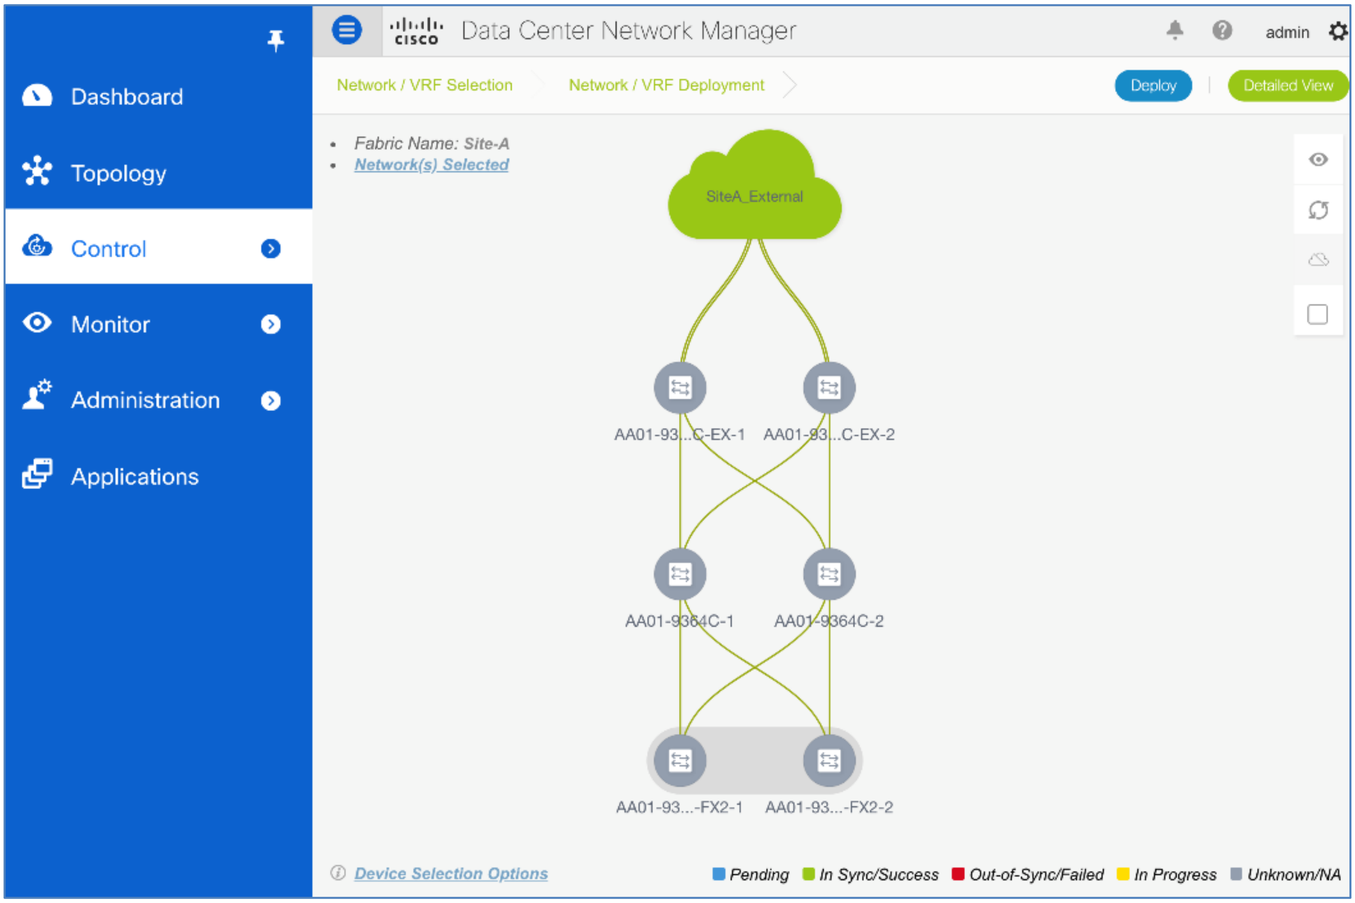

- *Fabric Name: Site-A*
- *Network(s) Selected*

7. Select the Leaf switches where these networks need to be deployed. Click the **Quick Attach** button.

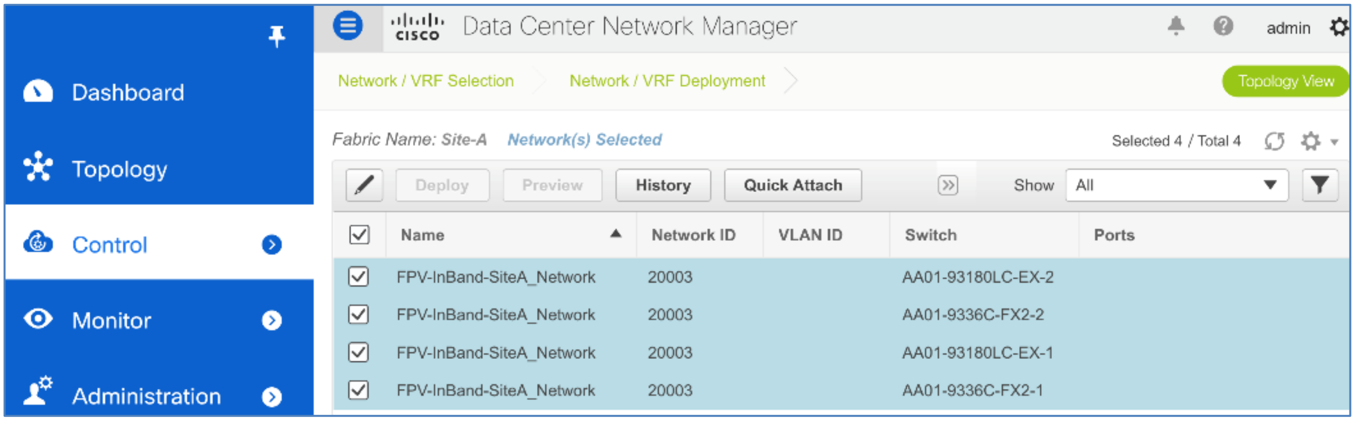

Fabric Name: Site-A *Network(s) Selected*

Selected 4 / Total 4

☑	Name ▲	Network ID	VLAN ID	Switch	Ports
☑	FPV-InBand-SiteA_Network	20003		AA01-93180LC-EX-2	
☑	FPV-InBand-SiteA_Network	20003		AA01-9336C-FX2-2	
☑	FPV-InBand-SiteA_Network	20003		AA01-93180LC-EX-1	
☑	FPV-InBand-SiteA_Network	20003		AA01-9336C-FX2-1	

8. Click **OK.**

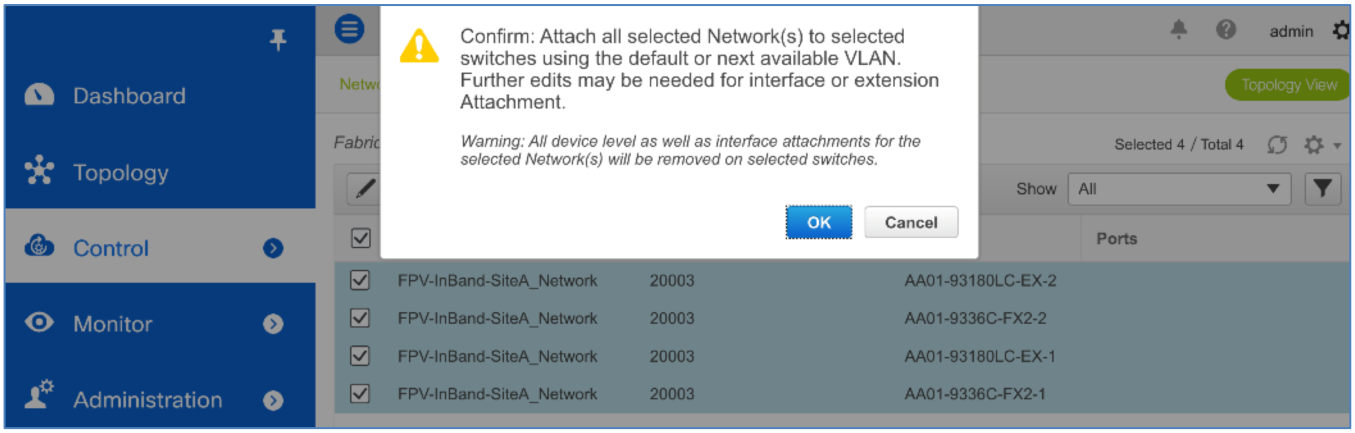

Confirm: Attach all selected Network(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected Network(s) will be removed on selected switches.

Selected 4 / Total 4

☑	FPV-InBand-SiteA_Network	20003	AA01-93180LC-EX-2
☑	FPV-InBand-SiteA_Network	20003	AA01-9336C-FX2-2
☑	FPV-InBand-SiteA_Network	20003	AA01-93180LC-EX-1
☑	FPV-InBand-SiteA_Network	20003	AA01-9336C-FX2-1

9. Click the **Preview** button to view pending changes. Click the **X** to close the window.

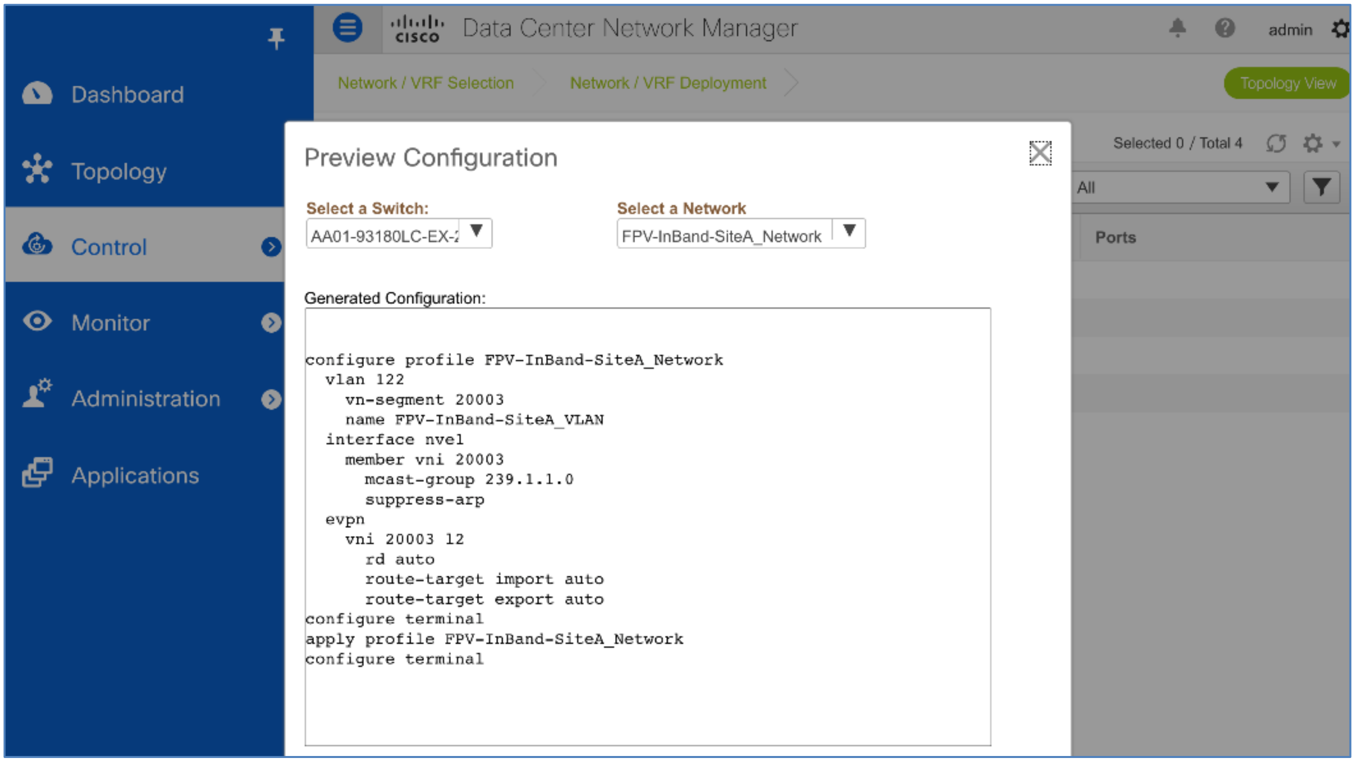

Preview Configuration

Select a Switch: AA01-93180LC-EX-2

Select a Network: FPV-InBand-SiteA_Network

Generated Configuration:

```
configure profile FPV-InBand-SiteA_Network
  vlan 122
    vn-segment 20003
    name FPV-InBand-SiteA_VLAN
  interface nvel
    member vni 20003
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 20003 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
apply profile FPV-InBand-SiteA_Network
configure terminal
```

10. Click the **Deploy** button.

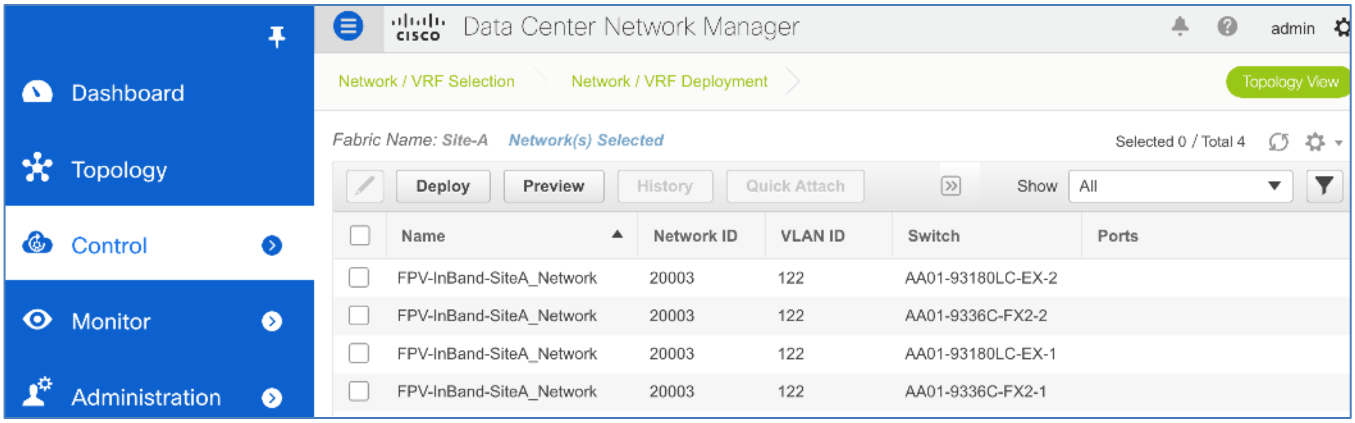

Fabric Name: Site-A Network(s) Selected

Selected 0 / Total 4

Name	Network ID	VLAN ID	Switch	Ports
FPV-InBand-SiteA_Network	20003	122	AA01-93180LC-EX-2	
FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-2	
FPV-InBand-SiteA_Network	20003	122	AA01-93180LC-EX-1	
FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1	

11. Click the **Topology View** button to view where the selected network is deployed in the fabric topology.

Deploy FlexPod vMotion Network

The FlexPod vMotion network is deployed in this design in **Layer 2 Only** mode with no gateway defined in the VXLAN fabric.

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**.

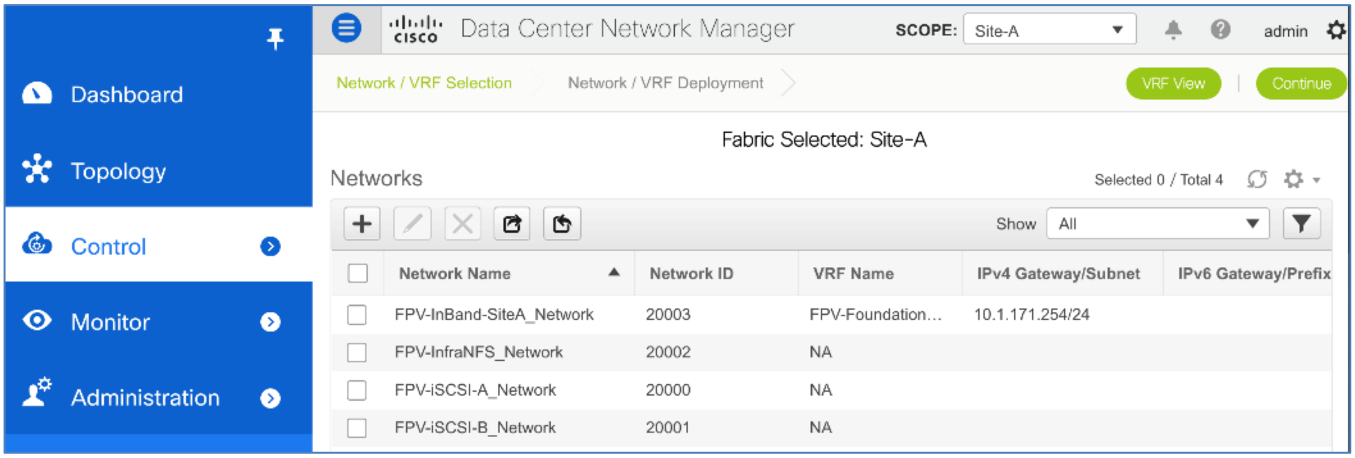

Fabric Selected: Site-A

Networks

Selected 0 / Total 4

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
FPV-InfraNFS_Network	20002	NA		
FPV-iSCSI-A_Network	20000	NA		
FPV-iSCSI-B_Network	20001	NA		

3. Click on the **[+]** icon to deploy a new Tenant network for the FlexPod infrastructure traffic. Specify a **Network Name.** Select the checkbox for **Layer 2 Only** mode. Specify a **VLAN ID**. In the **Network Profile > General** section, specify a **VLAN** name. Leave everything else as-is.

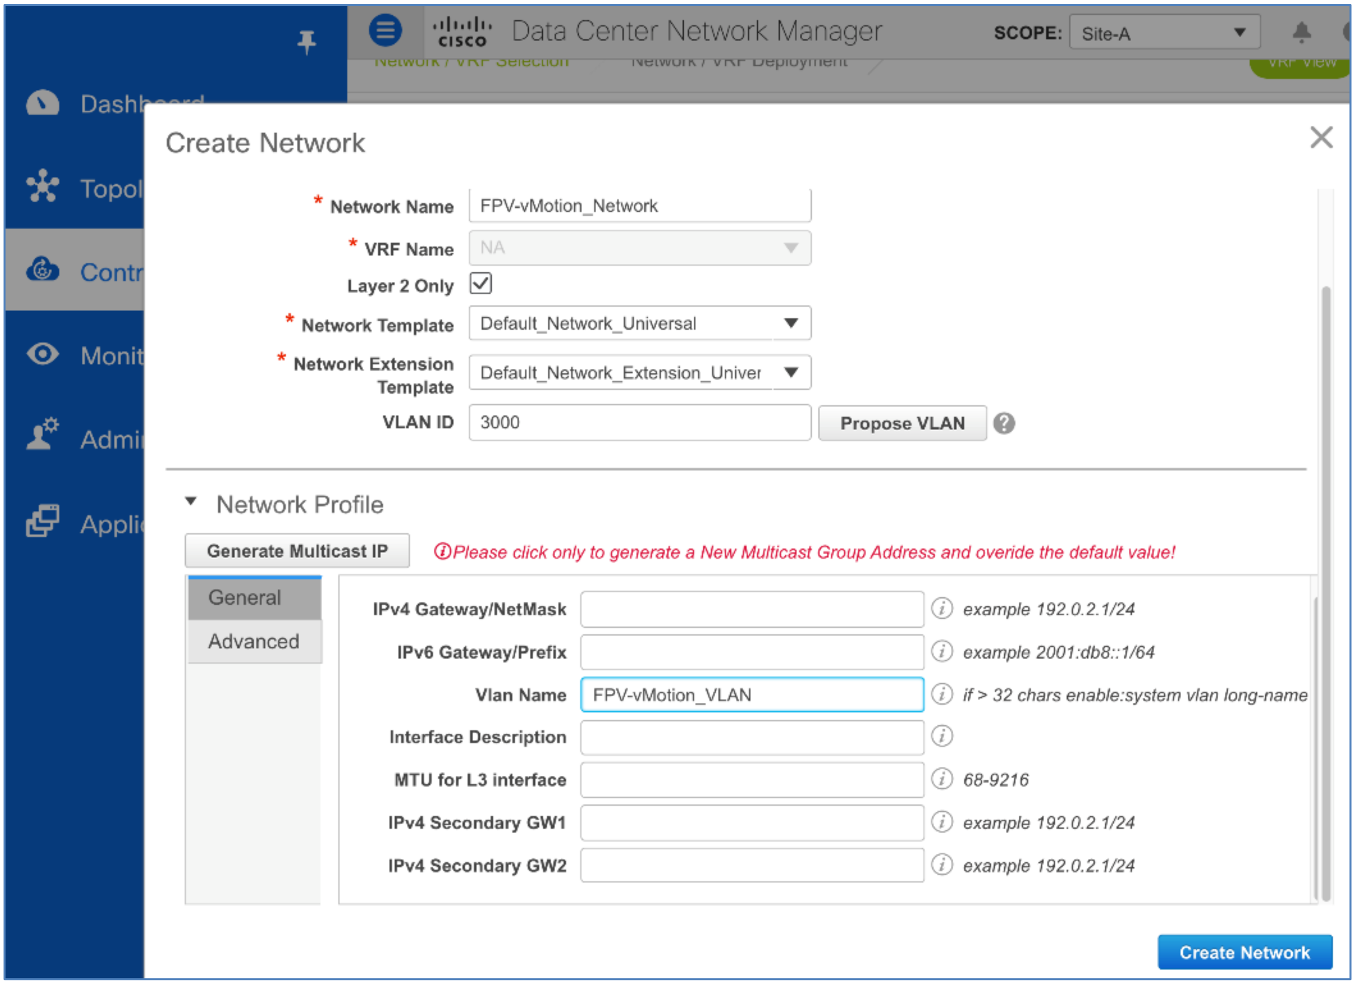

Create Network

* Network Name: FPV-vMotion_Network

* VRF Name: NA

Layer 2 Only: ☑

* Network Template: Default_Network_Universal

* Network Extension Template: Default_Network_Extension_Univer

VLAN ID: 3000

Network Profile

ⓘPlease click only to generate a New Multicast Group Address and overide the default value!

IPv4 Gateway/NetMask		example 192.0.2.1/24
IPv6 Gateway/Prefix		example 2001:db8::1/64
Vlan Name	FPV-vMotion_VLAN	if > 32 chars enable:system vlan long-name
Interface Description		
MTU for L3 interface		68-9216
IPv4 Secondary GW1		example 192.0.2.1/24
IPv4 Secondary GW2		example 192.0.2.1/24

4. In the **Network Profile > Advanced** section, leave everything as-is.

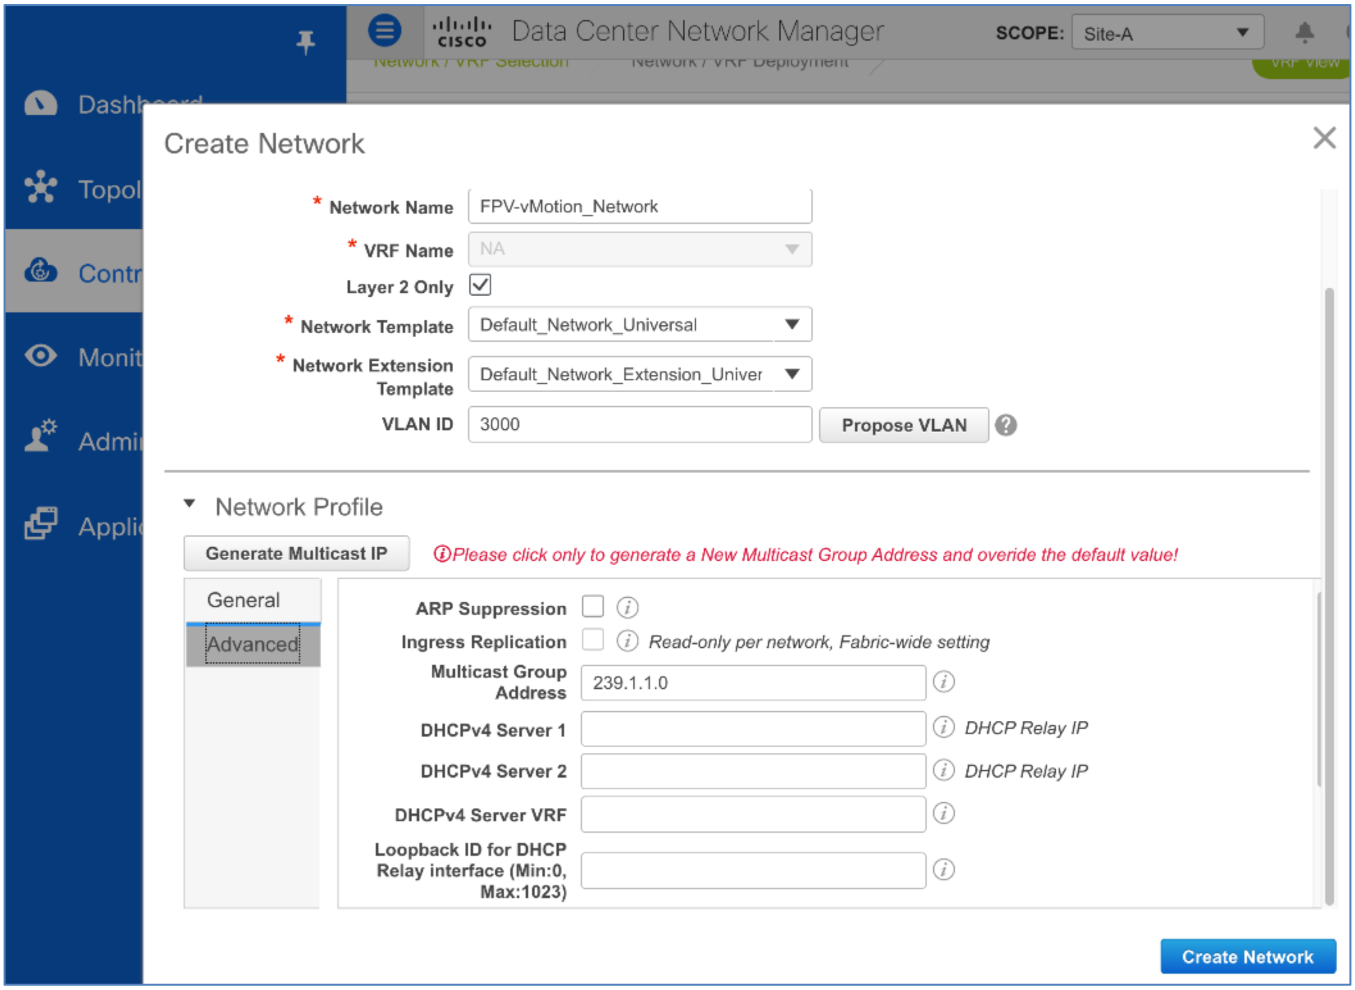

Create Network

* Network Name: FPV-vMotion_Network

* VRF Name: NA

Layer 2 Only: ☑

* Network Template: Default_Network_Universal

* Network Extension Template: Default_Network_Extension_Univer

VLAN ID: 3000

Propose VLAN

Network Profile

Generate Multicast IP

Please click only to generate a New Multicast Group Address and overide the default value!

General | Advanced

ARP Suppression: ☐

Ingress Replication: ☐ *Read-only per network, Fabric-wide setting*

Multicast Group Address: 239.1.1.0

DHCPv4 Server 1: *DHCP Relay IP*

DHCPv4 Server 2: *DHCP Relay IP*

DHCPv4 Server VRF:

Loopback ID for DHCP Relay interface (Min:0, Max:1023):

Create Network

5. Click the **Create Network** button.

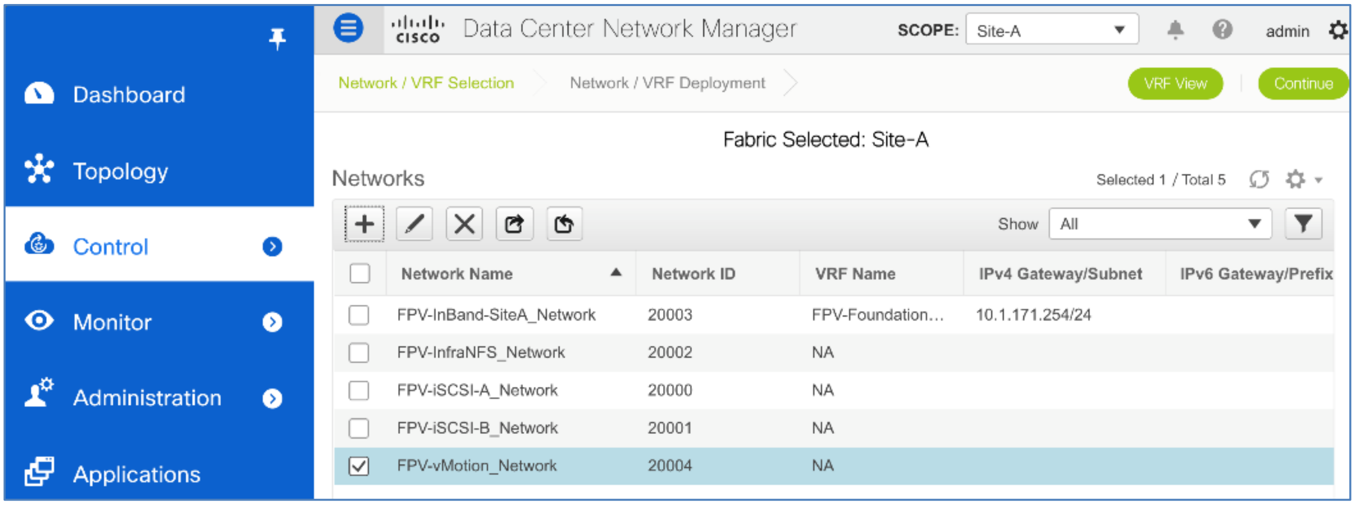

Fabric Selected: Site-A

Networks

Selected 1 / Total 5

Show All

☐	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☐	FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
☐	FPV-InfraNFS_Network	20002	NA		
☐	FPV-iSCSI-A_Network	20000	NA		
☐	FPV-iSCSI-B_Network	20001	NA		
☑	FPV-vMotion_Network	20004	NA		

6. Click the **Continue** button.

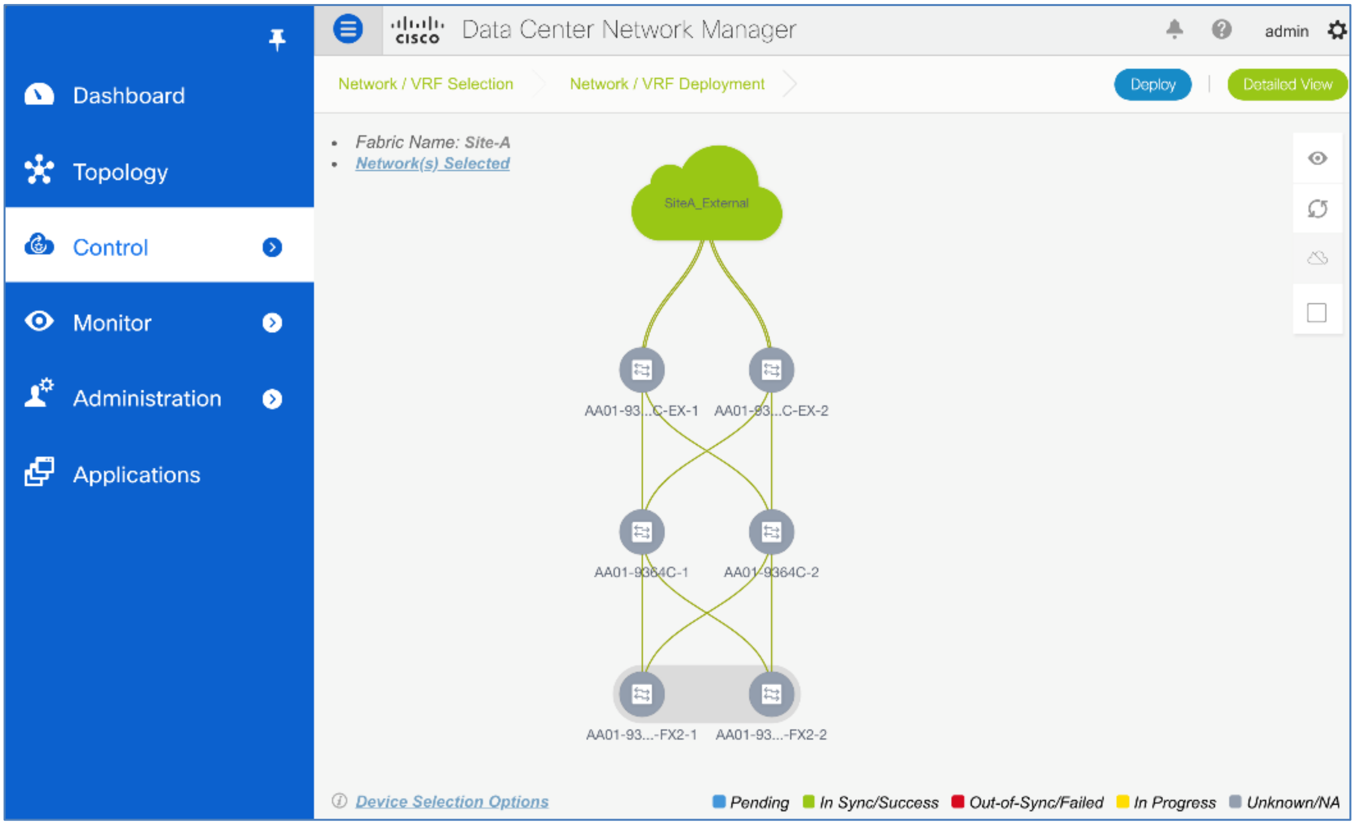

- Fabric Name: Site-A
- Network(s) Selected

7. Click the **Detailed View** button. Select the Leaf switches where these networks need to be deployed. Click the **Quick Attach** button.

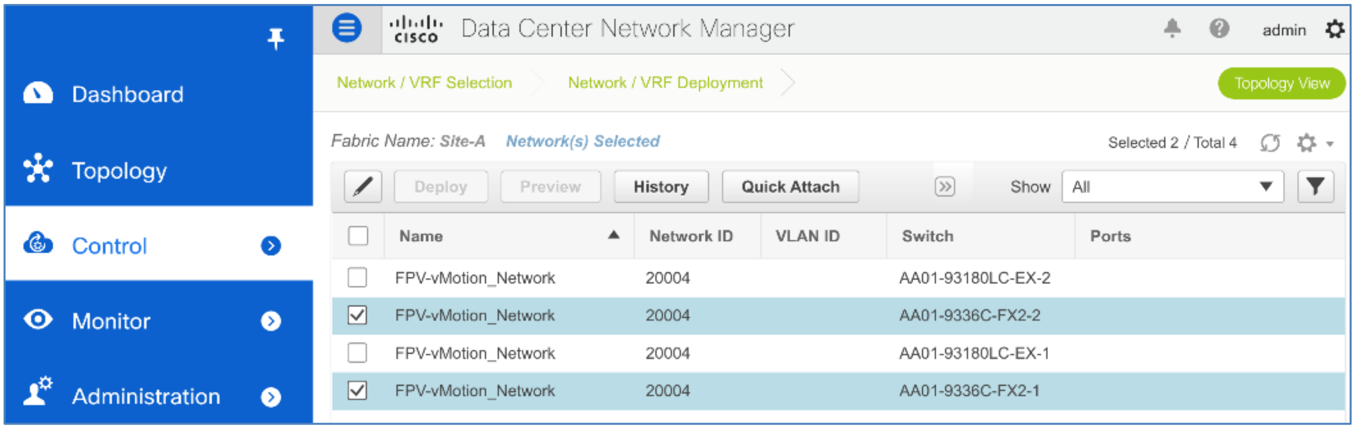

Fabric Name: Site-A Network(s) Selected

Selected 2 / Total 4

□	Name	Network ID	VLAN ID	Switch	Ports
□	FPV-vMotion_Network	20004		AA01-93180LC-EX-2	
☑	FPV-vMotion_Network	20004		AA01-9336C-FX2-2	
□	FPV-vMotion_Network	20004		AA01-93180LC-EX-1	
☑	FPV-vMotion_Network	20004		AA01-9336C-FX2-1	

8. Click **OK.**

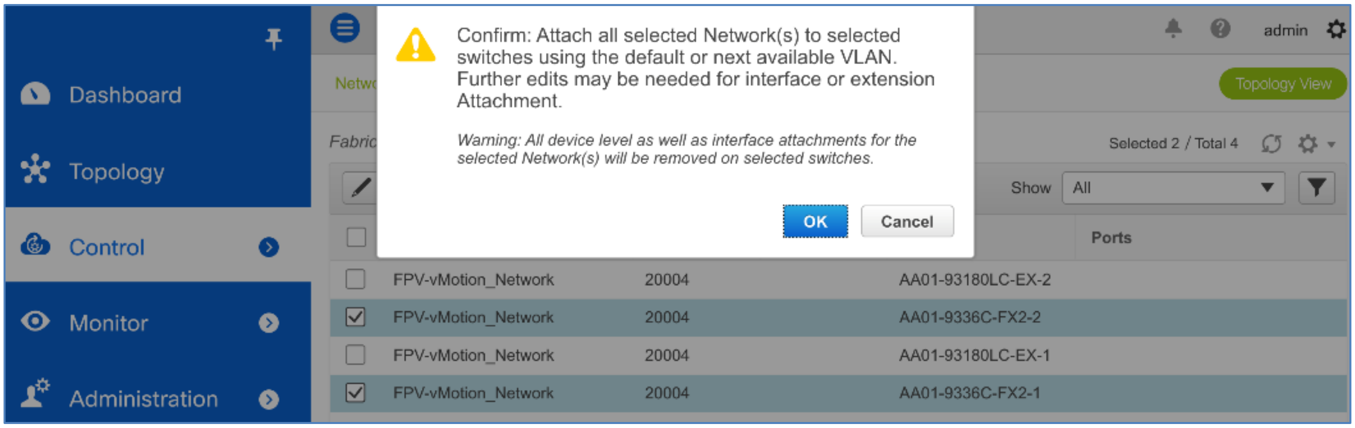

Confirm: Attach all selected Network(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected Network(s) will be removed on selected switches.

□	FPV-vMotion_Network	20004	AA01-93180LC-EX-2
☑	FPV-vMotion_Network	20004	AA01-9336C-FX2-2
□	FPV-vMotion_Network	20004	AA01-93180LC-EX-1
☑	FPV-vMotion_Network	20004	AA01-9336C-FX2-1

9. Click the **Preview** button to view pending changes. Click the **X** to close the window.

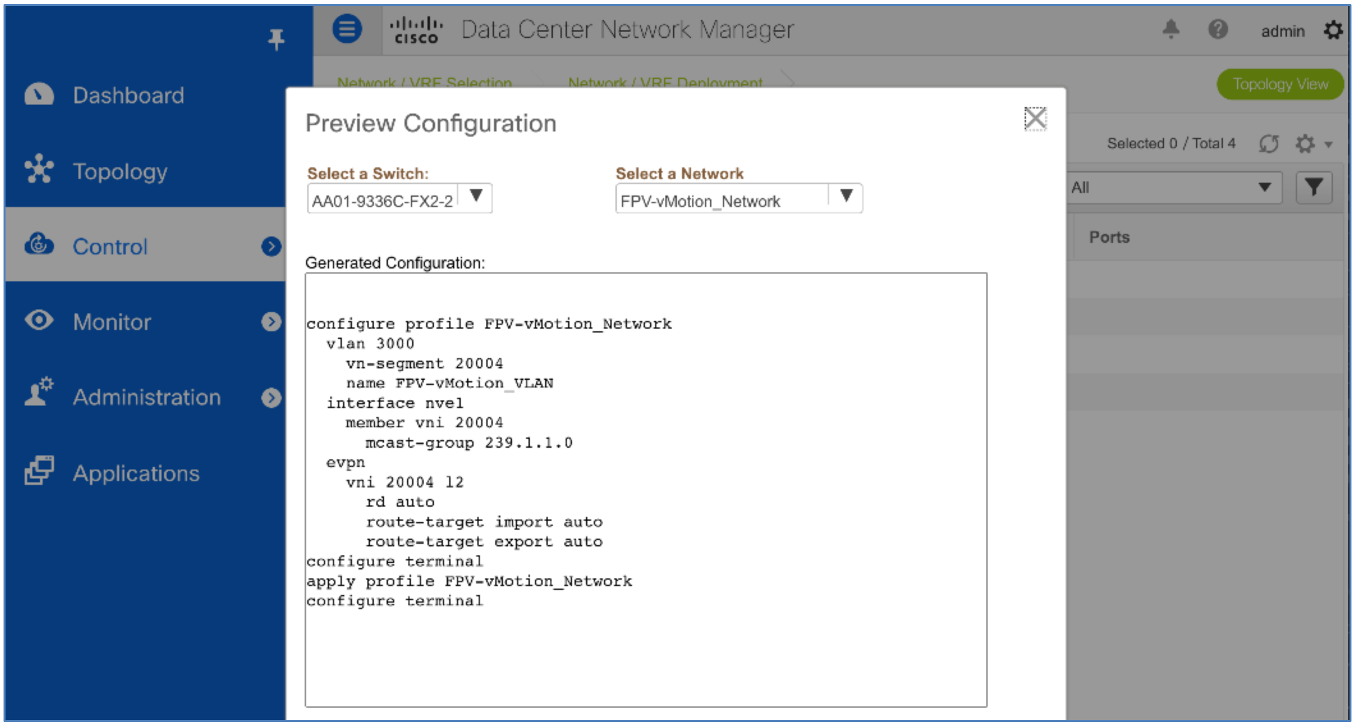

Preview Configuration

Select a Switch: AA01-9336C-FX2-2

Select a Network: FPV-vMotion_Network

Generated Configuration:

```
configure profile FPV-vMotion_Network
  vlan 3000
    vn-segment 20004
    name FPV-vMotion_VLAN
  interface nve1
    member vni 20004
      mcast-group 239.1.1.0
  evpn
    vni 20004 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
apply profile FPV-vMotion_Network
configure terminal
```

10. Click the **Deploy** button. The status should go from **PENDING** to **IN PROGRESS** to **DEPLOYED** in the **Status** column for the two Leaf switches. Scroll to the right as needed to see all columns in this view. Click the **Topology View** button.

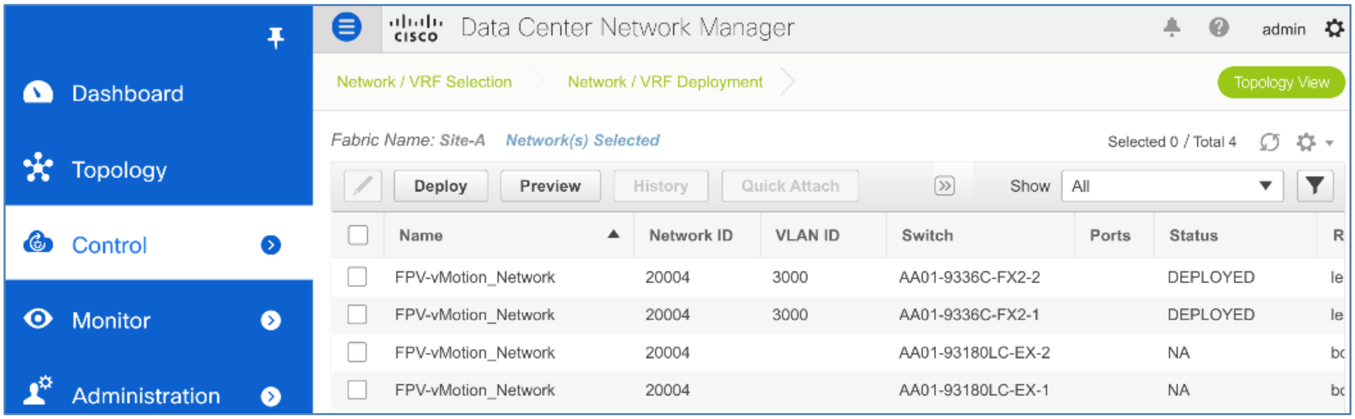

Fabric Name: Site-A   Network(s) Selected   Selected 0 / Total 4

□	Name	Network ID	VLAN ID	Switch	Ports	Status	R
□	FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-2		DEPLOYED	le
□	FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-1		DEPLOYED	le
□	FPV-vMotion_Network	20004		AA01-93180LC-EX-2		NA	bc
□	FPV-vMotion_Network	20004		AA01-93180LC-EX-1		NA	bc

11. In the **Topology View** to see the where the selected network is deployed in the Site-A topology.

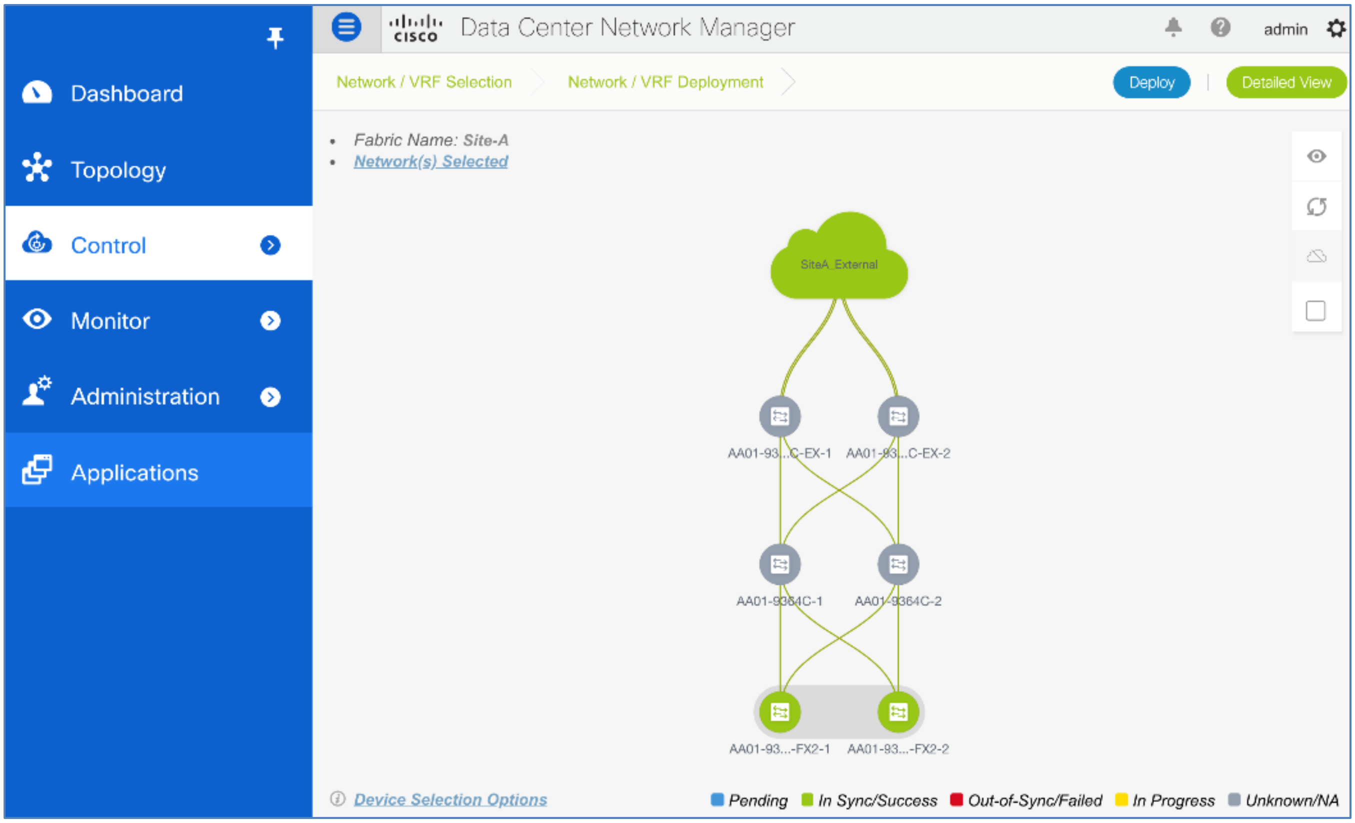

- Fabric Name: Site-A
- Network(s) Selected

🟦 Pending 🟩 In Sync/Success 🟥 Out-of-Sync/Failed 🟨 In Progress ⬜ Unknown/NA

Deploy FlexPod Infrastructure Network for Common Services

The FlexPod Common Services network is deployed in this design in Layer 3 mode where the traffic is Layer 3 forwarded by the fabric and the gateway is a distributed anycast gateway in the VXLAN fabric. This network is used to host common infrastructure services such as VMware vCenter.

To deploy the FlexPod infrastructure network for Common Services, follow these steps:

1. Use a browser to navigate to Cisco DCNM’s GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**.

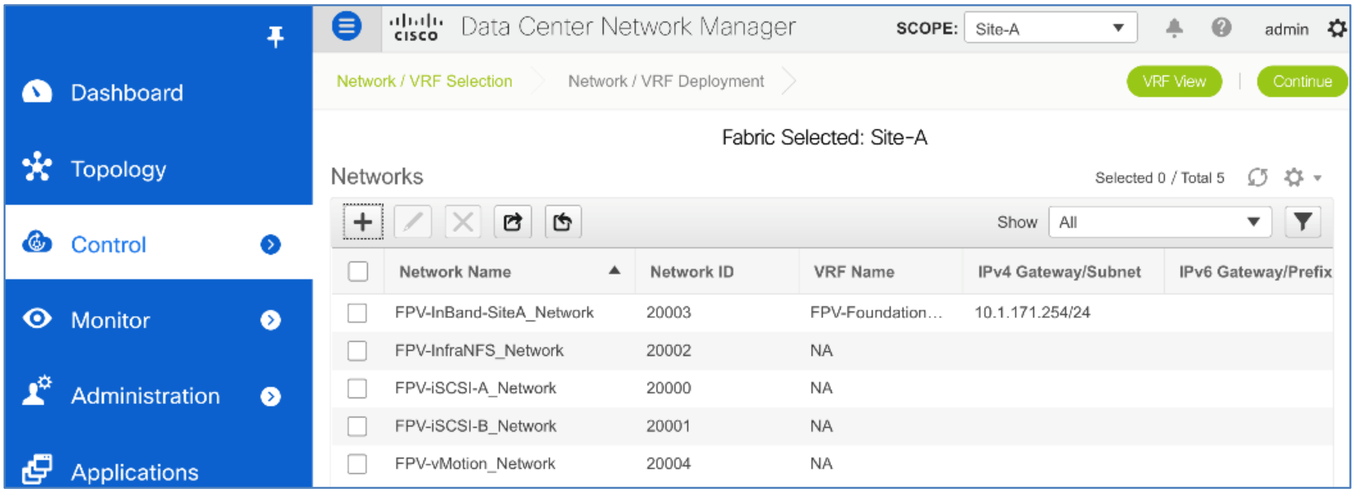

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
□	FPV-InBand-SiteA_Network	20003	FPV-Foundation...	10.1.171.254/24	
□	FPV-InfraNFS_Network	20002	NA		
□	FPV-iSCSI-A_Network	20000	NA		
□	FPV-iSCSI-B_Network	20001	NA		
□	FPV-vMotion_Network	20004	NA		

3. Click on the **[+]** icon to deploy a new Tenant network for the FlexPod infrastructure traffic. Specify a **Network Name, VRF Name.** In this deployment, we are specifying the **VLAN ID** we specifically want to use but you can optionally let DCNM pick up one from the defined pool in **Fabric Settings.** Specify a **VLAN ID**. In the **Network Profile > General** section of the window, specify a **IPv4 Gateway/Network, VLAN Name, Interface Description** and **MTU**. Leave everything else as-is.

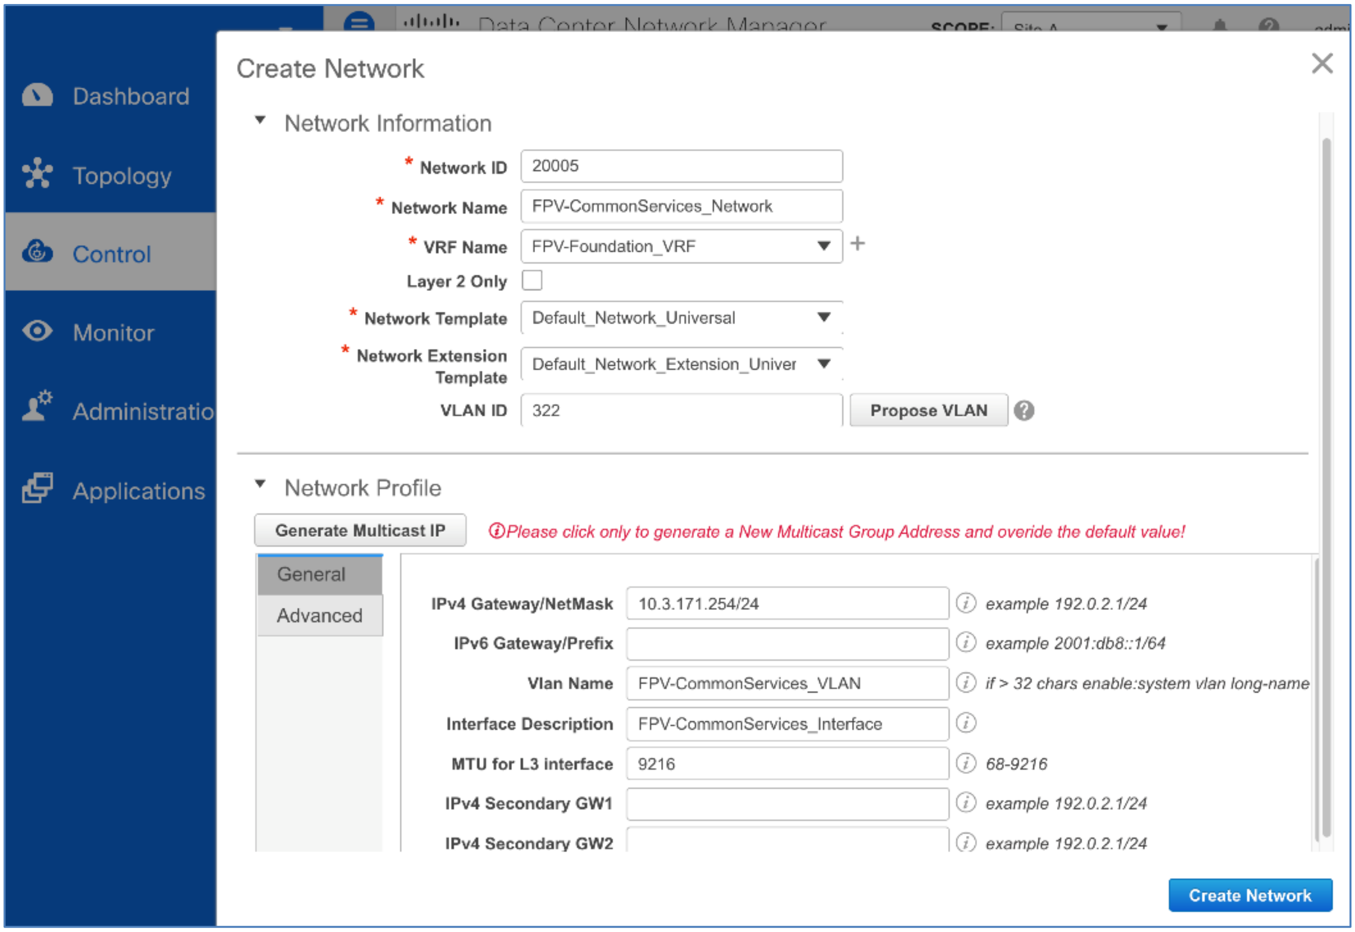

Create Network

Network Information

* Network ID	20005
* Network Name	FPV-CommonServices_Network
* VRF Name	FPV-Foundation_VRF
Layer 2 Only	□
* Network Template	Default_Network_Universal
* Network Extension Template	Default_Network_Extension_Univer
VLAN ID	322

Network Profile

Please click only to generate a New Multicast Group Address and overide the default value!

IPv4 Gateway/NetMask	10.3.171.254/24	example 192.0.2.1/24
IPv6 Gateway/Prefix		example 2001:db8::1/64
Vlan Name	FPV-CommonServices_VLAN	if > 32 chars enable:system vlan long-name
Interface Description	FPV-CommonServices_Interface	
MTU for L3 interface	9216	68-9216
IPv4 Secondary GW1		example 192.0.2.1/24
IPv4 Secondary GW2		example 192.0.2.1/24

4. In the **Network Profile > Advanced** section of the window, enable **ARP Suppression.** Leave everything as-is.

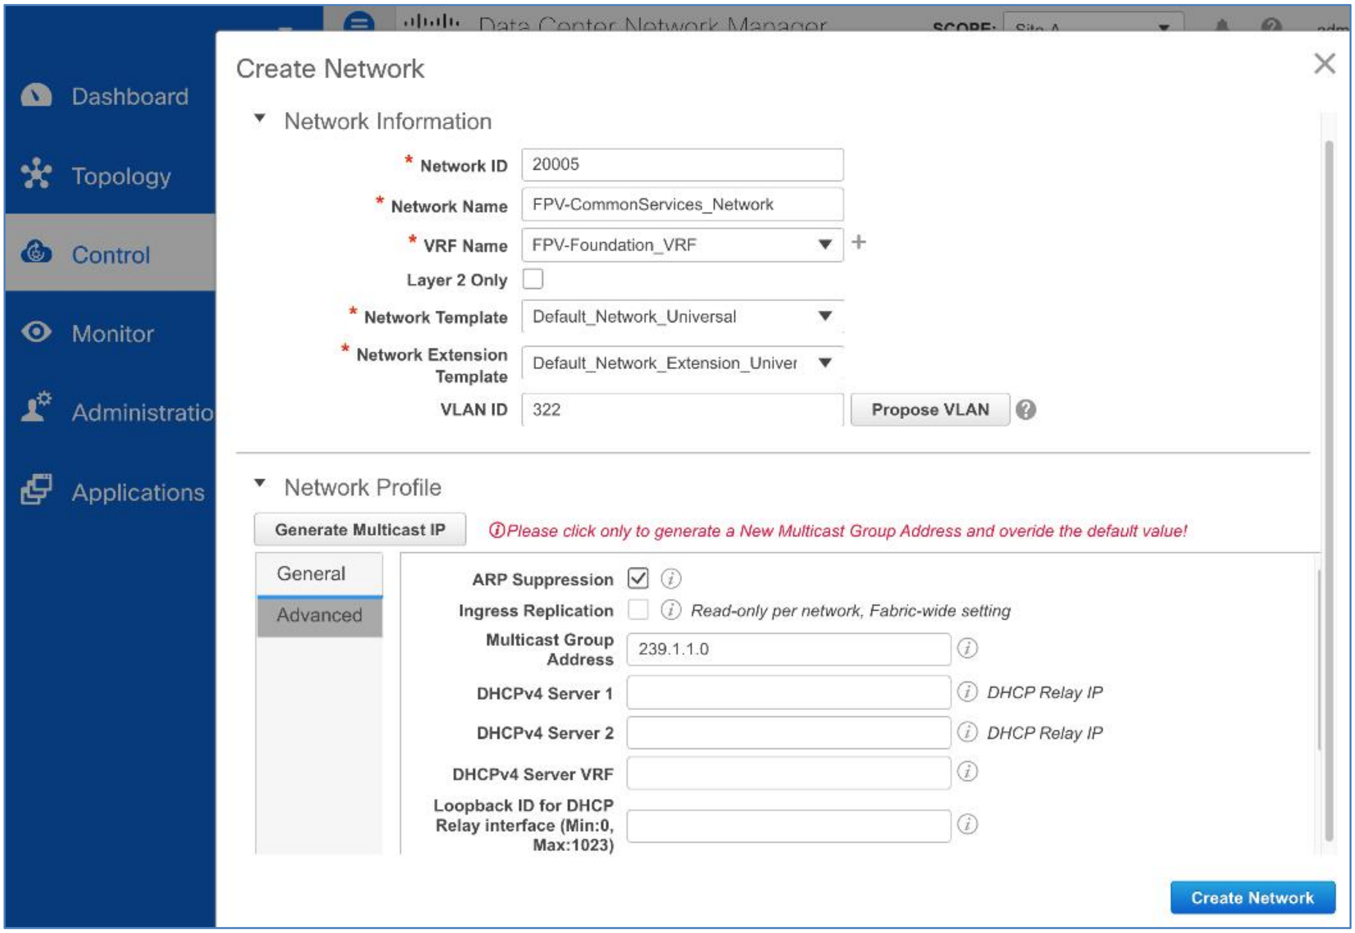

Create Network

Network Information

* Network ID	20005
* Network Name	FPV-CommonServices_Network
* VRF Name	FPV-Foundation_VRF
Layer 2 Only	□
* Network Template	Default_Network_Universal
* Network Extension Template	Default_Network_Extension_Univer
VLAN ID	322

Network Profile

Please click only to generate a New Multicast Group Address and overide the default value!

ARP Suppression	☑
Ingress Replication	□ *Read-only per network, Fabric-wide setting*
Multicast Group Address	239.1.1.0
DHCPv4 Server 1	*DHCP Relay IP*
DHCPv4 Server 2	*DHCP Relay IP*
DHCPv4 Server VRF	
Loopback ID for DHCP Relay interface (Min:0, Max:1023)	

5. Click the **Create Network** button. Click the **Continue** button.

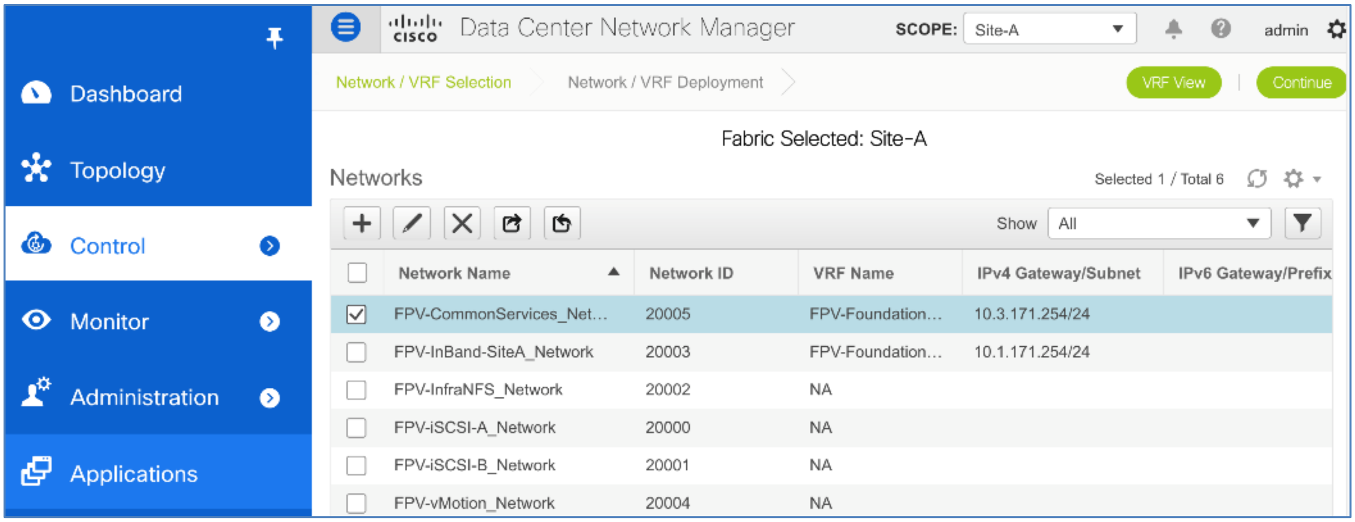

Fabric Selected: Site-A

Networks

Selected 1 / Total 6

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☑	FPV-CommonServices_Net…	20005	FPV-Foundation…	10.3.171.254/24	
□	FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
□	FPV-InfraNFS_Network	20002	NA		
□	FPV-iSCSI-A_Network	20000	NA		
□	FPV-iSCSI-B_Network	20001	NA		
□	FPV-vMotion_Network	20004	NA		

6. Click the **Detailed View** button.

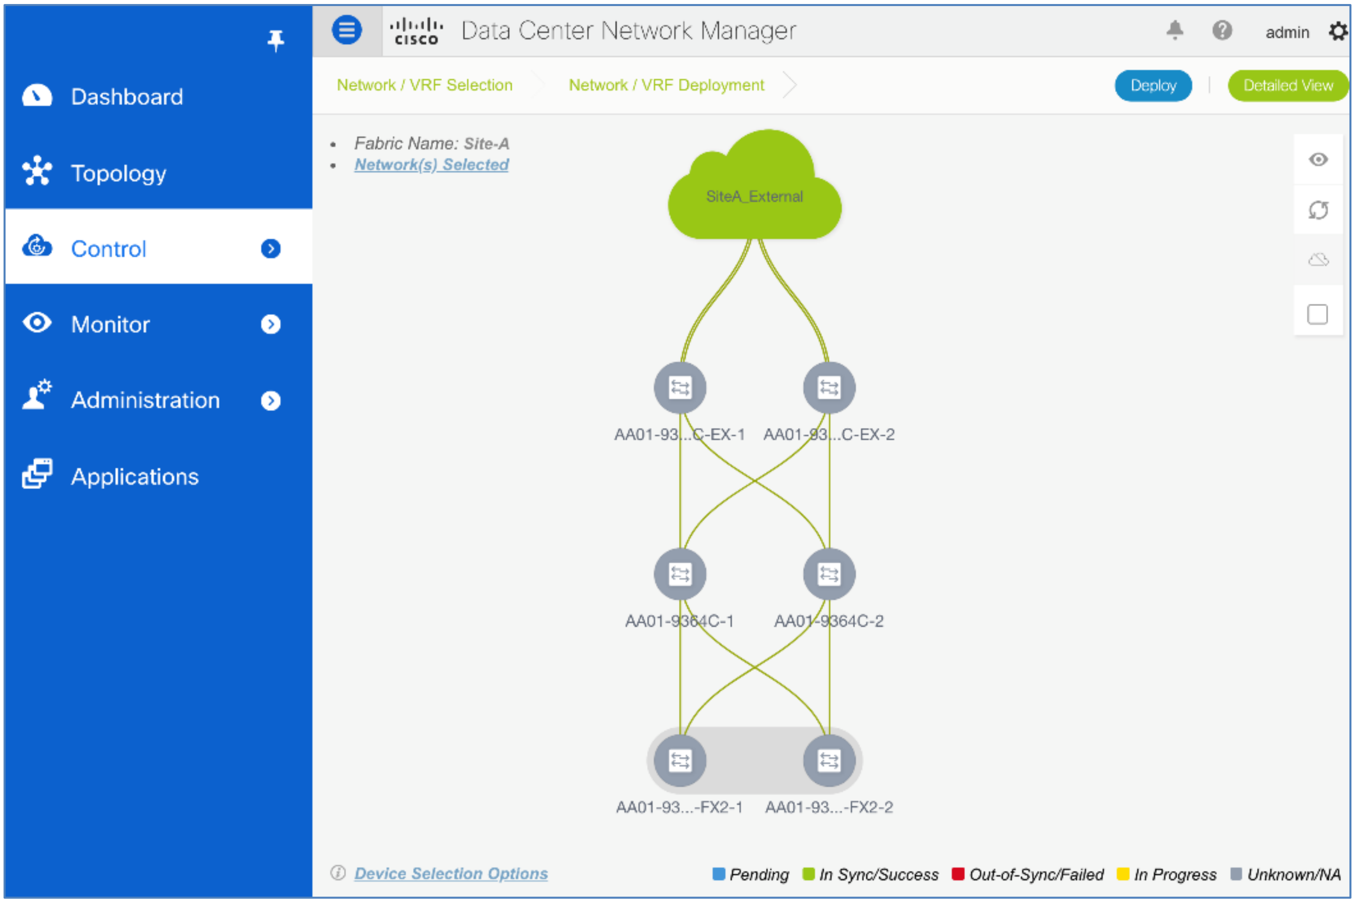

7. Select the Leaf switches where these networks need to be deployed. Click the **Quick Attach** button.

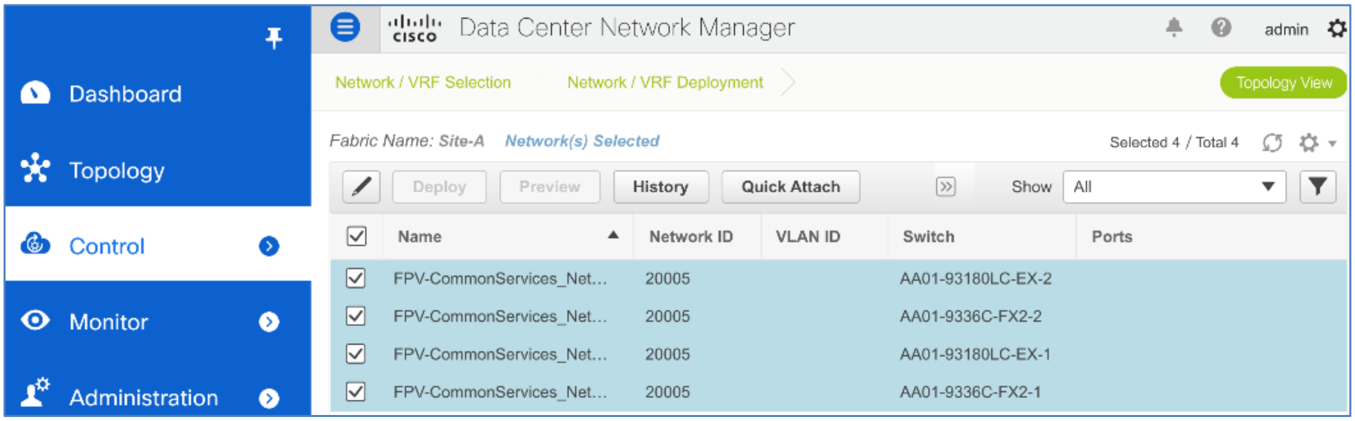

Fabric Name: Site-A Network(s) Selected

Selected 4 / Total 4

☑	Name	Network ID	VLAN ID	Switch	Ports
☑	FPV-CommonServices_Net...	20005		AA01-93180LC-EX-2	
☑	FPV-CommonServices_Net...	20005		AA01-9336C-FX2-2	
☑	FPV-CommonServices_Net...	20005		AA01-93180LC-EX-1	
☑	FPV-CommonServices_Net...	20005		AA01-9336C-FX2-1	

8. Click **OK.**

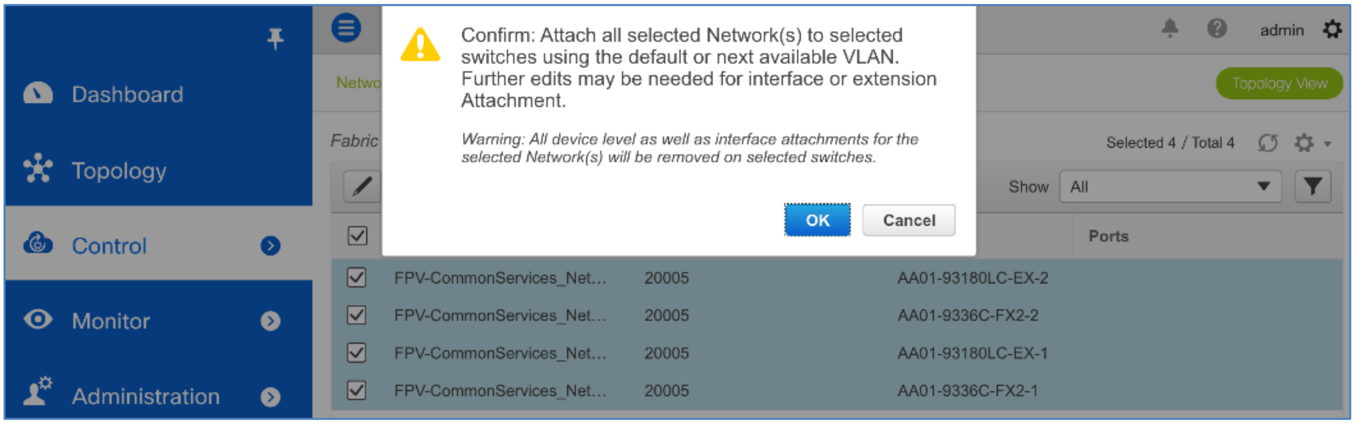

Confirm: Attach all selected Network(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected Network(s) will be removed on selected switches.

Fabric

Selected 4 / Total 4

☑	FPV-CommonServices_Net...	20005	AA01-93180LC-EX-2
☑	FPV-CommonServices_Net...	20005	AA01-9336C-FX2-2
☑	FPV-CommonServices_Net...	20005	AA01-93180LC-EX-1
☑	FPV-CommonServices_Net...	20005	AA01-9336C-FX2-1

9. Click the **Preview** button to view pending changes. Click the **X** to close the window.

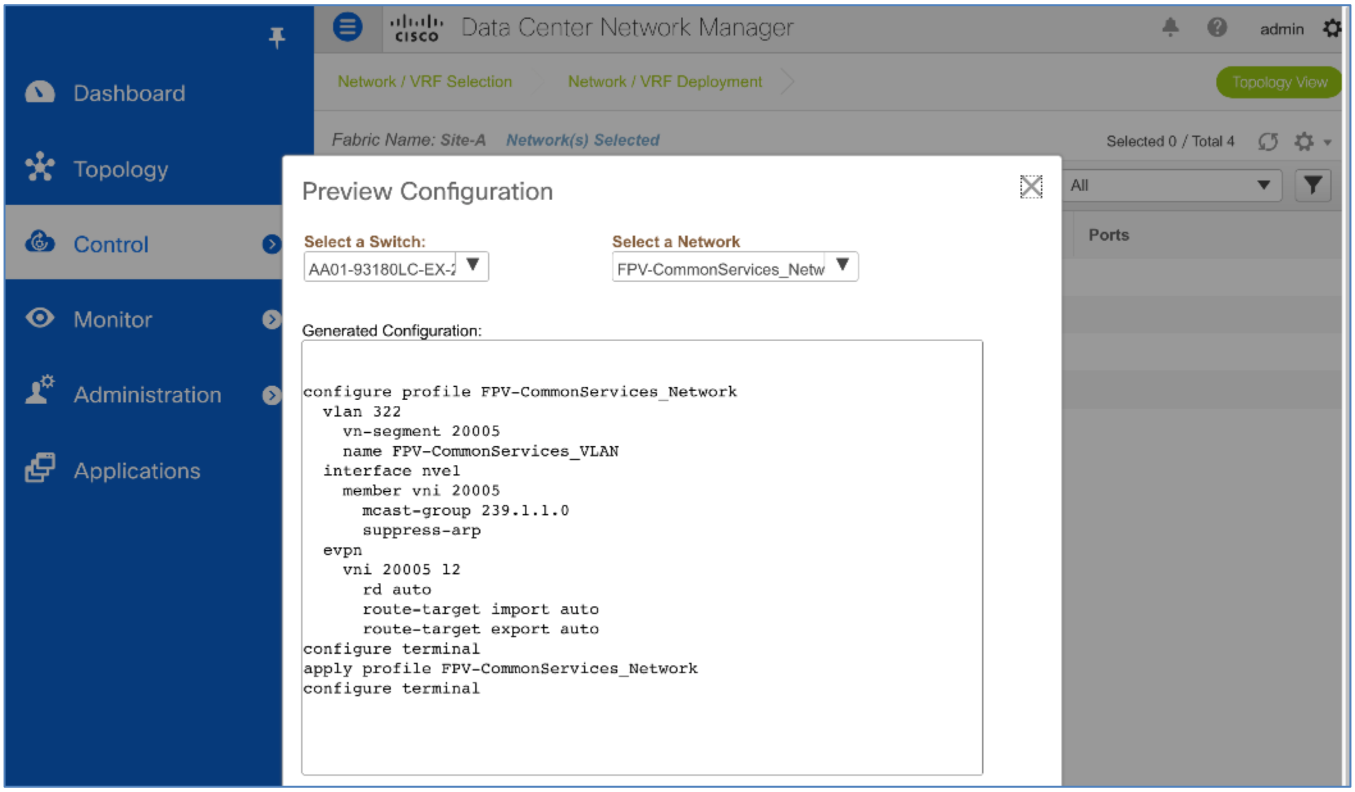

Fabric Name: Site-A *Network(s) Selected*

Selected 0 / Total 4

Preview Configuration

Select a Switch: AA01-93180LC-EX-2

Select a Network: FPV-CommonServices_Netw

Generated Configuration:

```
configure profile FPV-CommonServices_Network
  vlan 322
    vn-segment 20005
    name FPV-CommonServices_VLAN
  interface nve1
    member vni 20005
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 20005 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
apply profile FPV-CommonServices_Network
configure terminal
```

10. Click the **Deploy** button. Verify the status is **DEPLOYED**. Scroll to the right as needed to see all columns in this view.

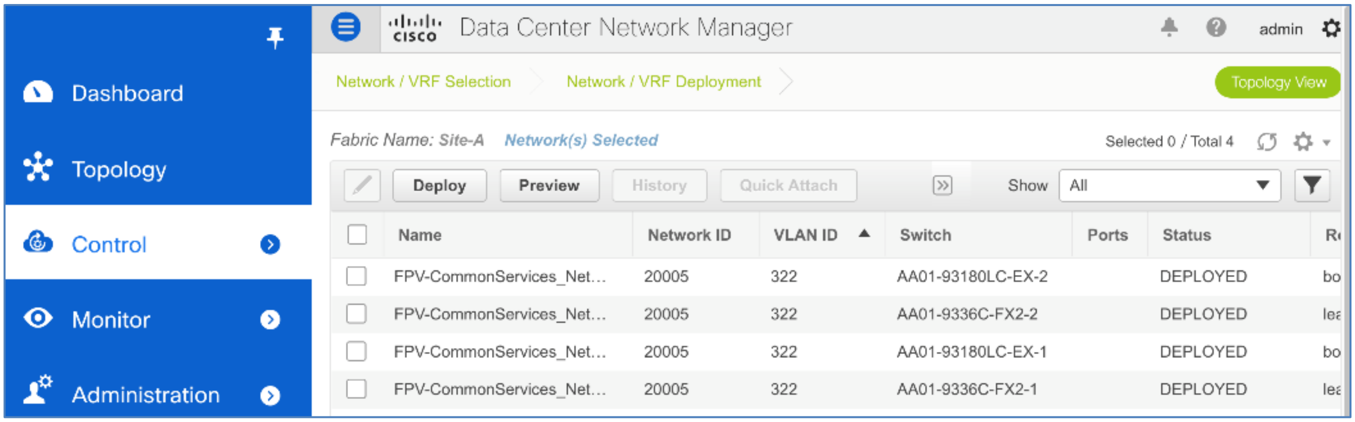

Fabric Name: Site-A *Network(s) Selected*

Selected 0 / Total 4

□	Name	Network ID	VLAN ID ▲	Switch	Ports	Status	Re
□	FPV-CommonServices_Net...	20005	322	AA01-93180LC-EX-2		DEPLOYED	bo
□	FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-2		DEPLOYED	lea
□	FPV-CommonServices_Net...	20005	322	AA01-93180LC-EX-1		DEPLOYED	bo
□	FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-1		DEPLOYED	lea

11. Click the **Topology View** button to view where the selected network is deployed in the fabric topology.

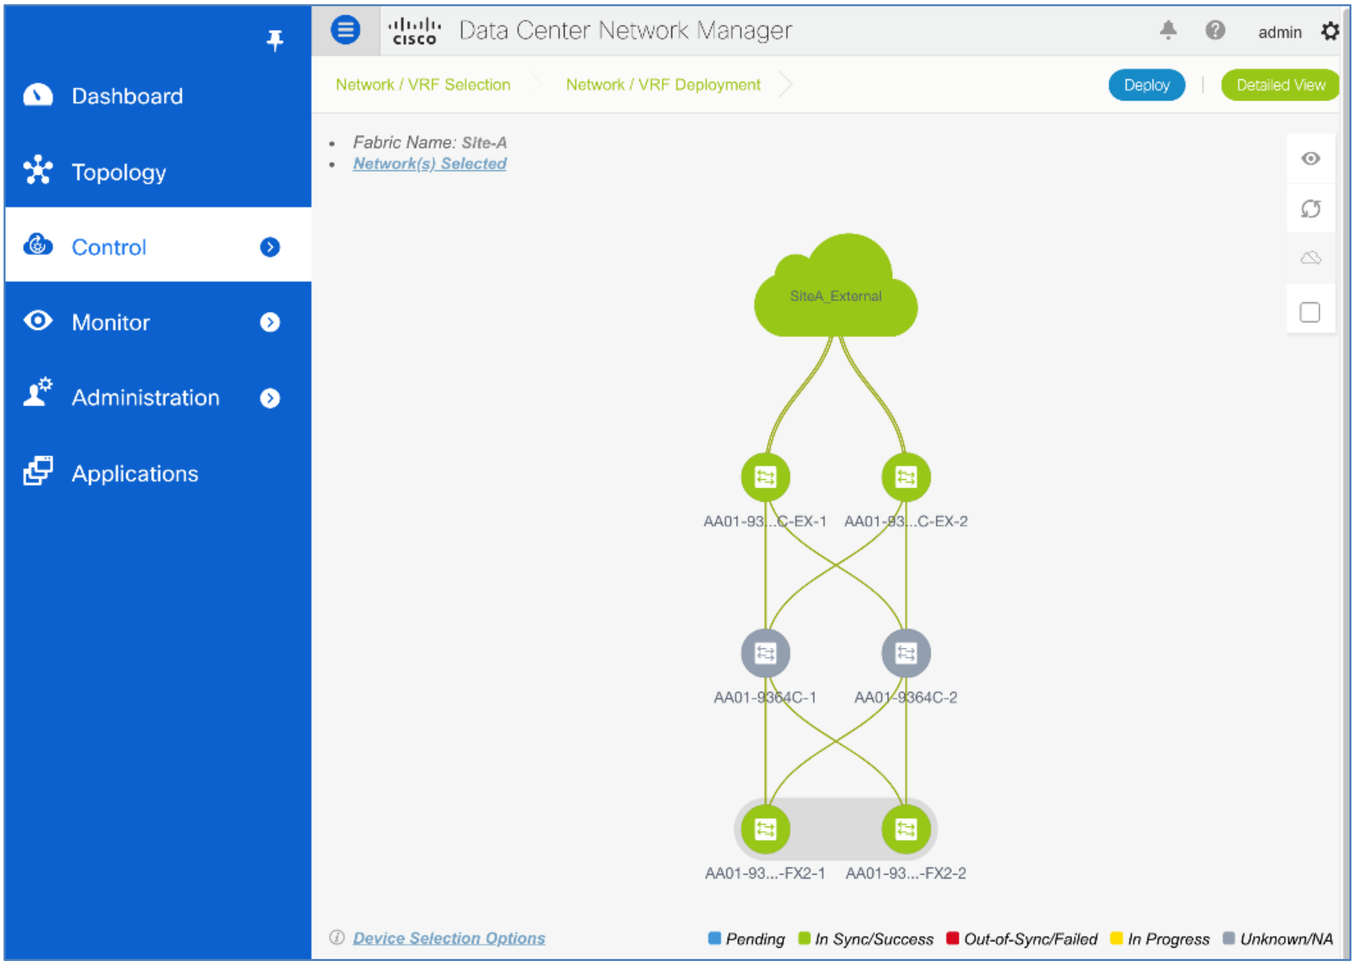

- *Fabric Name: Site-A*
- *Network(s) Selected*

Enable Access-Layer Connectivity to FlexPod Infrastructure Networks

To enable FlexPod infrastructure networks on access-layer connections to Cisco UCS domain and NetApp storage cluster, complete the steps outlined in the upcoming sections.

Enable Infrastructure Networks on Access-Layer Connections to Cisco UCS Domain

To enable infrastructure networks on the access-layer connections to Cisco UCS domain, follow these steps:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**. Click **OK** to exit any pop-ups that come up. Select the correct scope in the drop-down list next to **Scope:** in the top-right corner of the window.

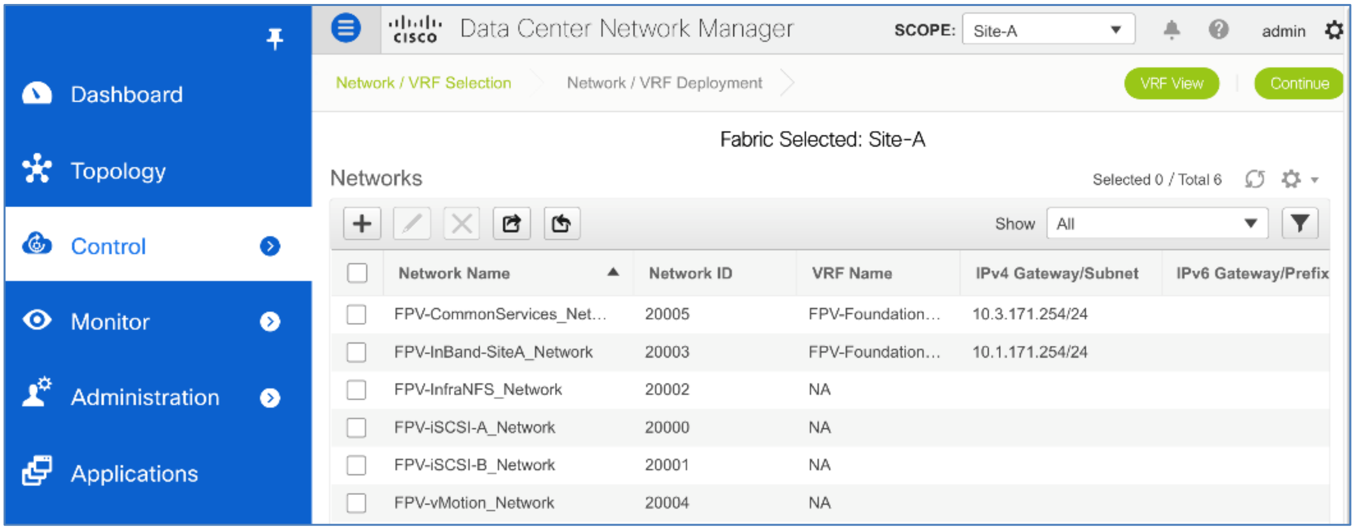

SCOPE: Site-A

Fabric Selected: Site-A

Networks

Selected 0 / Total 6

Show All

☐	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☐	FPV-CommonServices_Net…	20005	FPV-Foundation…	10.3.171.254/24	
☐	FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
☐	FPV-InfraNFS_Network	20002	NA		
☐	FPV-iSCSI-A_Network	20000	NA		
☐	FPV-iSCSI-B_Network	20001	NA		
☐	FPV-vMotion_Network	20004	NA		

3. Select the networks that need to be enabled on the access-layer connection to Cisco UCS Domain. All networks are selected in this case. Click the **Continue** button.

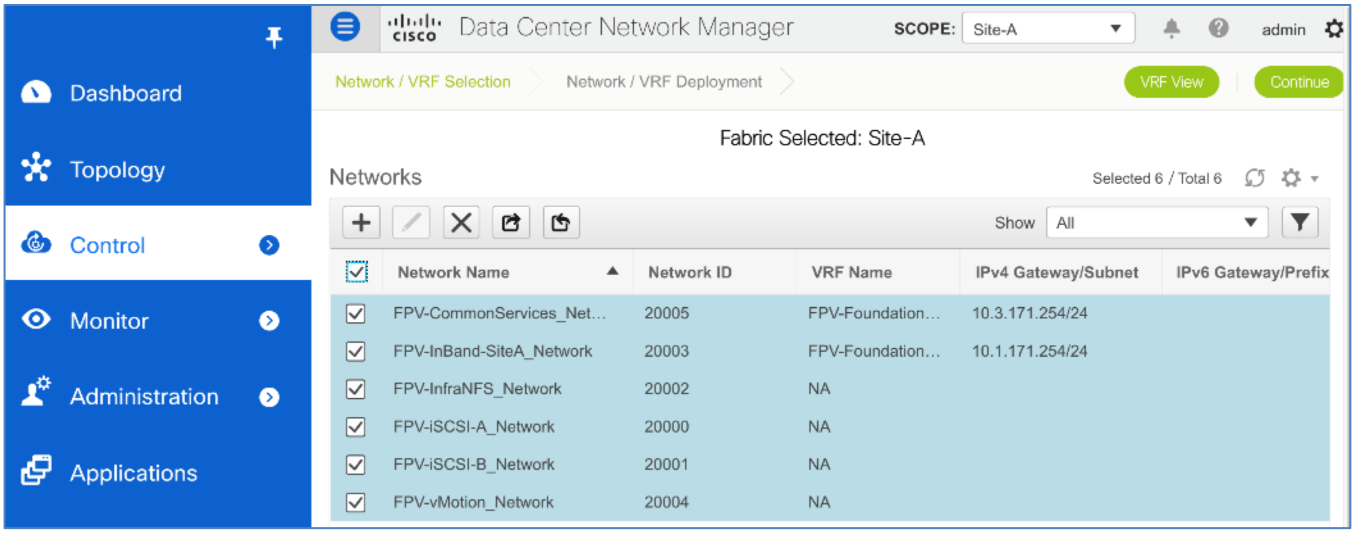

SCOPE: Site-A

Fabric Selected: Site-A

Networks

Selected 6 / Total 6

Show All

☑	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
☑	FPV-CommonServices_Net…	20005	FPV-Foundation…	10.3.171.254/24	
☑	FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
☑	FPV-InfraNFS_Network	20002	NA		
☑	FPV-iSCSI-A_Network	20000	NA		
☑	FPV-iSCSI-B_Network	20001	NA		
☑	FPV-vMotion_Network	20004	NA		

4. Click the **Detailed View** button.

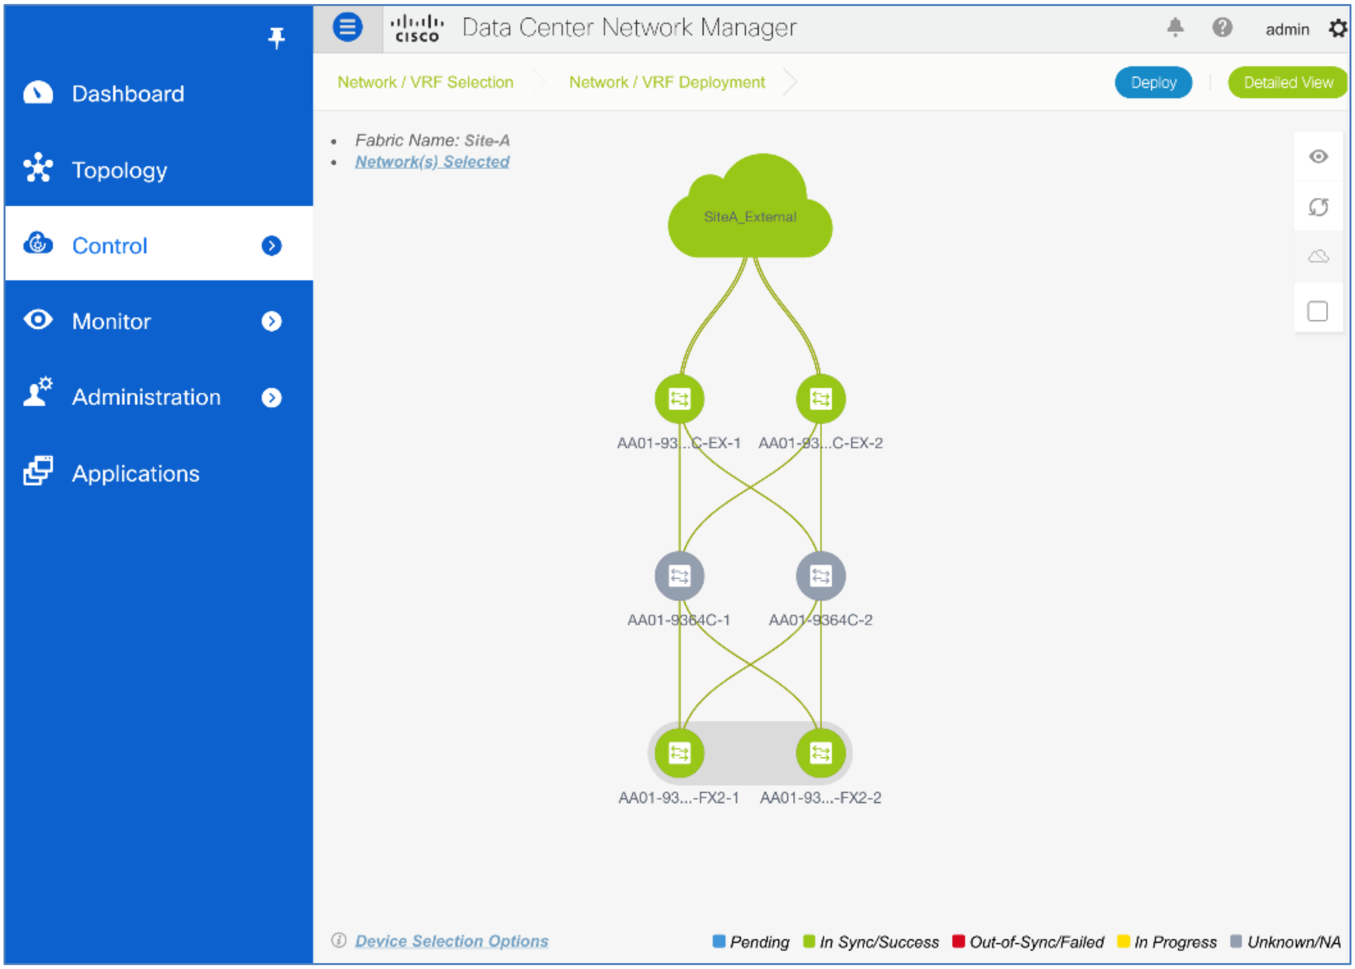

5. Select **Quick Filter** from the drop-down list next to **Show** from the top menu. This will expose the filter box above every column.

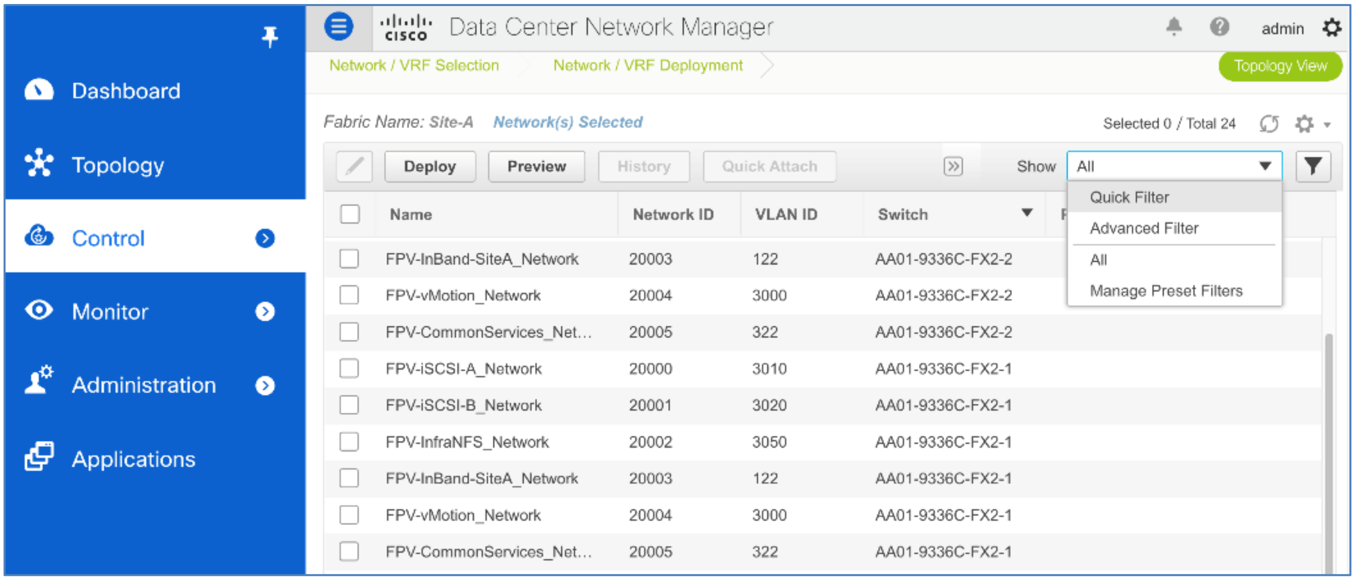

Name	Network ID	VLAN ID	Switch
FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-2
FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-2
FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-2
FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1
FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1
FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1
FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1
FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-1
FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-1

6. Filter based on the access-layer leaf switch that connects to the Cisco UCS Domain. Select the first switch in the vPC pair to the Cisco UCS domain. Select all networks that need to be deployed. Click on the pencil icon to edit the previously deployed networks on the first leaf switch in the vPC pair.

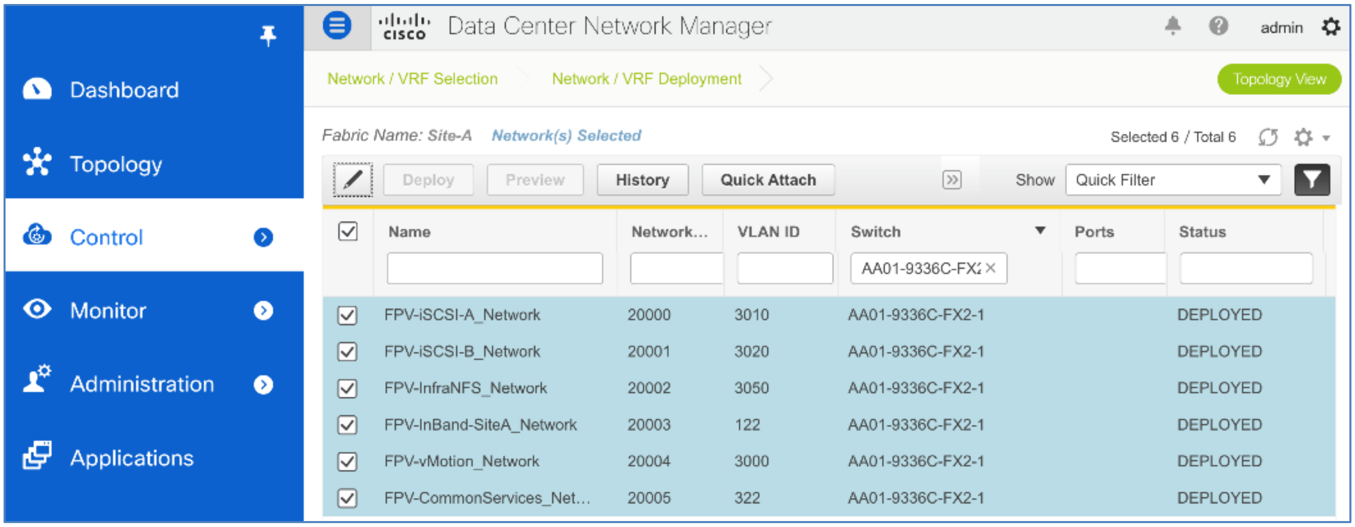

Fabric Name: Site-A *Network(s) Selected* Selected 6 / Total 6

☑	Name	Network...	VLAN ID	Switch	Ports	Status
				AA01-9336C-FX2 ×		
☑	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-1		DEPLOYED

7. You should see a network tab for each network and **both** leaf switches listed though only one switch was selected in the previous. This is because the switches are part of a vPC pair – a configuration that gets applied to one will get applied to both. Note the box … in the **Interfaces** column.

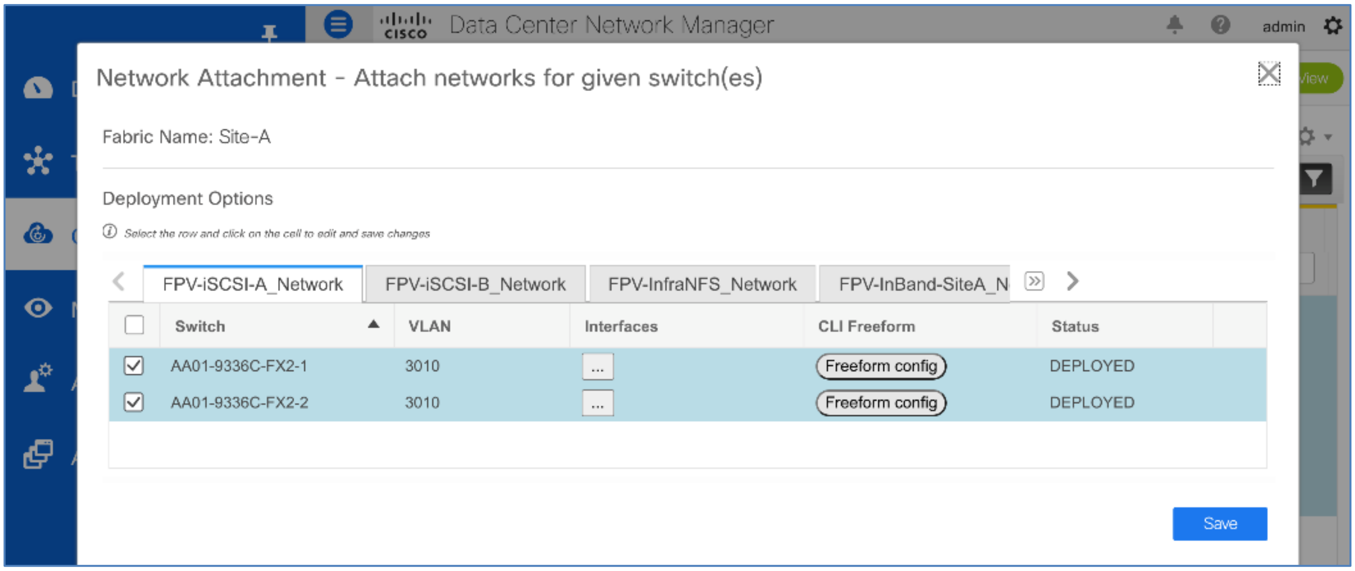

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ Select the row and click on the cell to edit and save changes

☐	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	3010	…	Freeform config	DEPLOYED
☑	AA01-9336C-FX2-2	3010	…	Freeform config	DEPLOYED

8. For the first network, click the box … in the **Interfaces** column next to the first switch listed in the leaf switch pair.

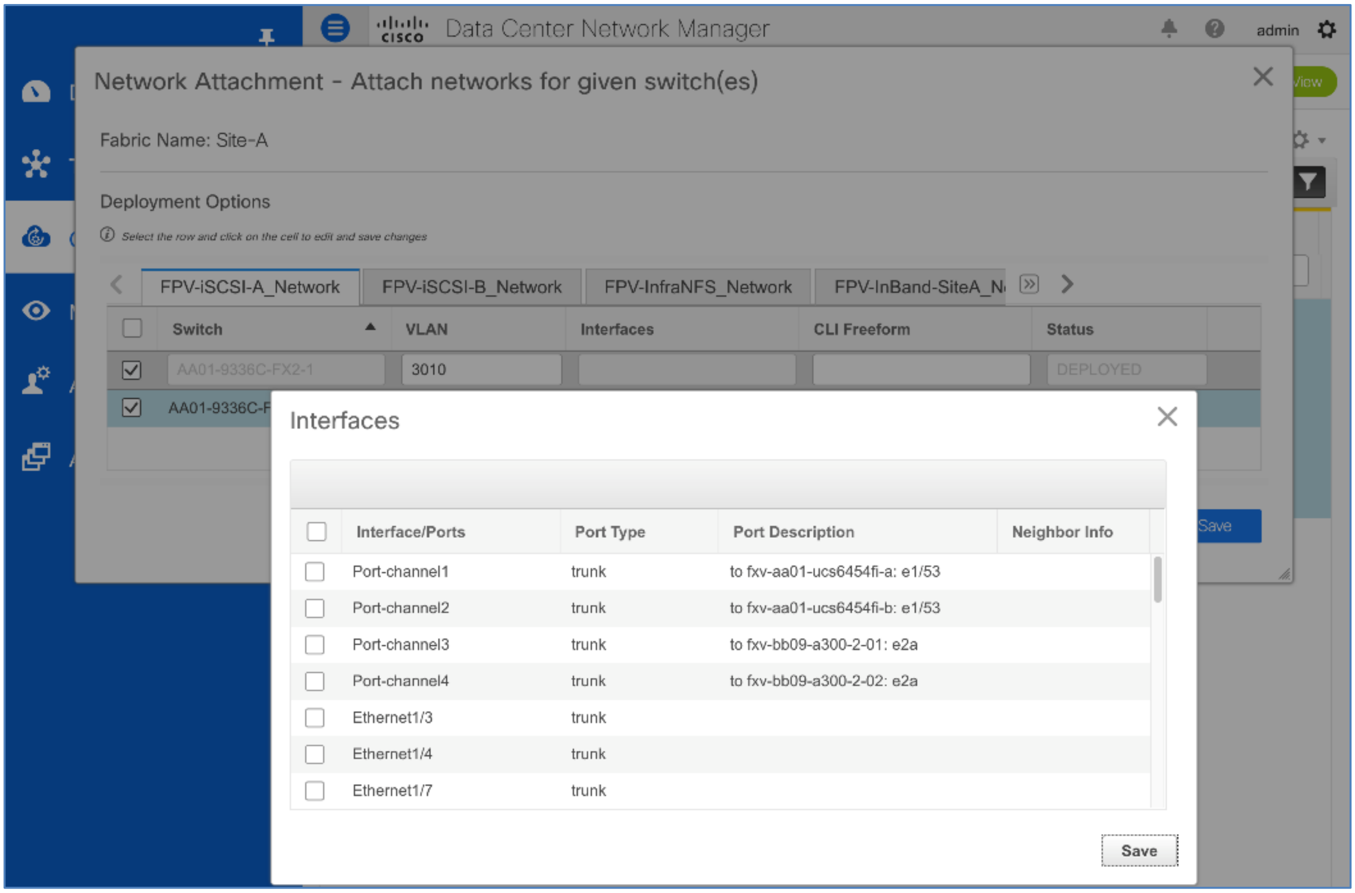

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

☐	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	3010			DEPLOYED
☑	AA01-9336C-F				

Interfaces

☐	Interface/Ports	Port Type	Port Description	Neighbor Info
☐	Port-channel1	trunk	to fxv-aa01-ucs6454fi-a: e1/53	
☐	Port-channel2	trunk	to fxv-aa01-ucs6454fi-b: e1/53	
☐	Port-channel3	trunk	to fxv-bb09-a300-2-01: e2a	
☐	Port-channel4	trunk	to fxv-bb09-a300-2-02: e2a	
☐	Ethernet1/3	trunk		
☐	Ethernet1/4	trunk		
☐	Ethernet1/7	trunk		

9. In the **Interfaces** pop-up window, select both port-channels that go to both Cisco UCS Fabric Interconnects in the Cisco UCS Domain.

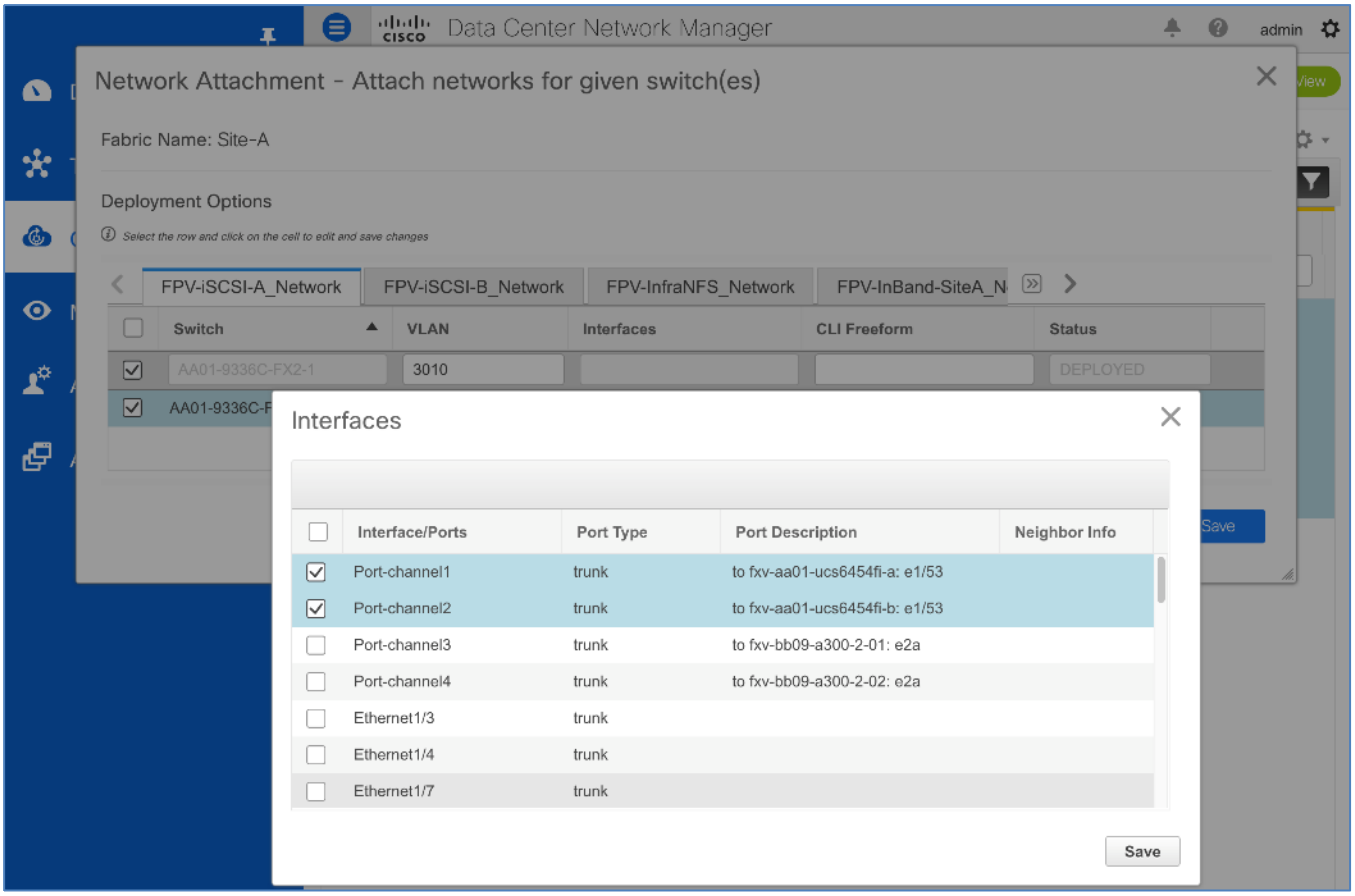

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ *Select the row and click on the cell to edit and save changes*

☐	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	3010			DEPLOYED
☑	AA01-9336C-F				

Interfaces

☐	Interface/Ports	Port Type	Port Description	Neighbor Info
☑	Port-channel1	trunk	to fxv-aa01-ucs6454fi-a: e1/53	
☑	Port-channel2	trunk	to fxv-aa01-ucs6454fi-b: e1/53	
☐	Port-channel3	trunk	to fxv-bb09-a300-2-01: e2a	
☐	Port-channel4	trunk	to fxv-bb09-a300-2-02: e2a	
☐	Ethernet1/3	trunk		
☐	Ethernet1/4	trunk		
☐	Ethernet1/7	trunk		

10. Click **Save.** Note that the interfaces are now populated for the first switch. Click **Save** again.

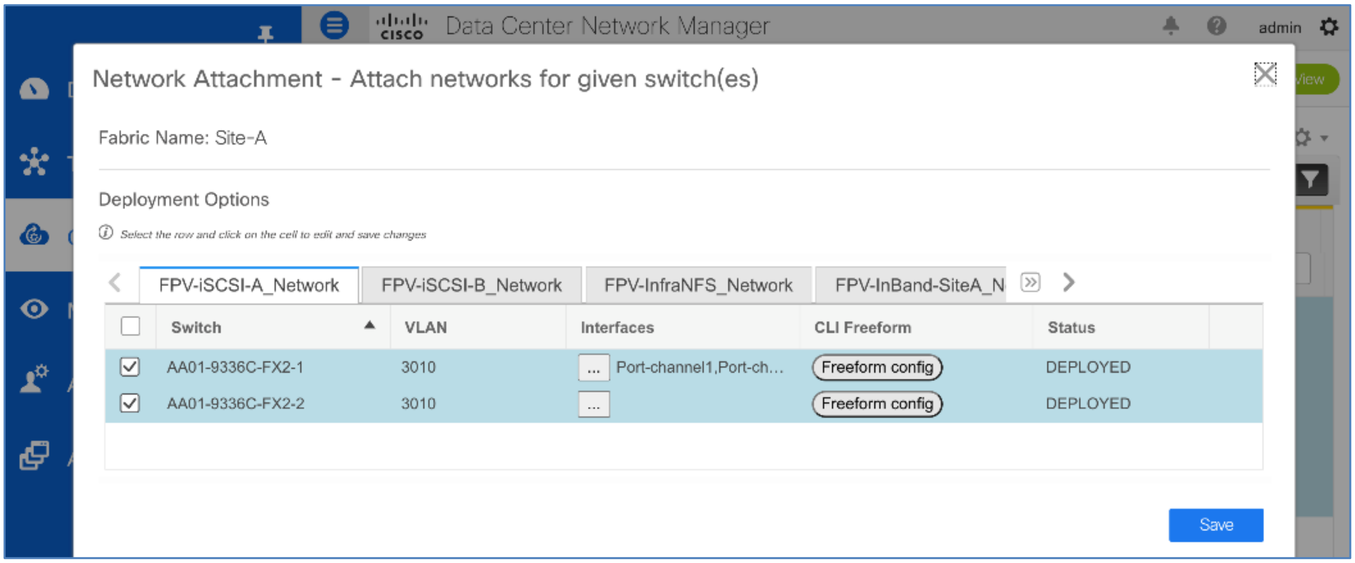

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ *Select the row and click on the cell to edit and save changes*

☐	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	3010	... Port-channel1,Port-ch...	Freeform config	DEPLOYED
☑	AA01-9336C-FX2-2	3010	...	Freeform config	DEPLOYED

11. The first network for the first switch now lists the **Ports** and the **Status** is now **PENDING.** Click the pencil icon again.

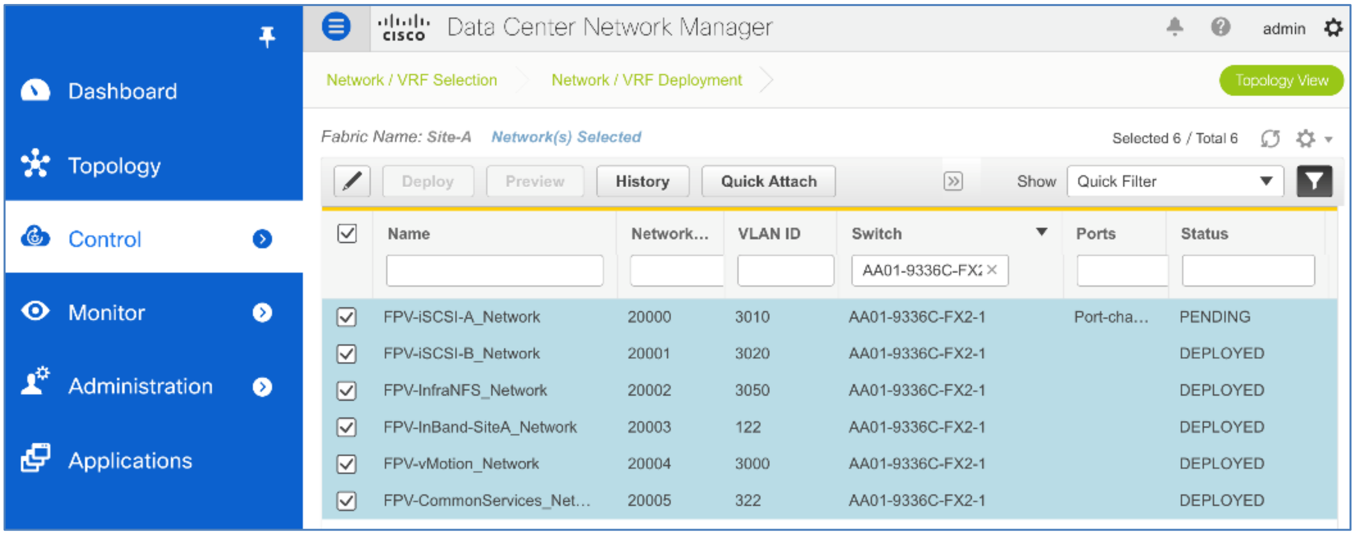

Fabric Name: Site-A Network(s) Selected

Selected 6 / Total 6

☑	Name	Network...	VLAN ID	Switch	Ports	Status
				AA01-9336C-FX2 ×		
☑	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	Port-cha...	PENDING
☑	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-1		DEPLOYED
☑	FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-1		DEPLOYED

12. You will now see that the port/interface information is now populated for both leaf switches though the configuration was only done for one switch. Cisco DCNM automatically configures both leaf switches in the same leaf switch vPC pair.

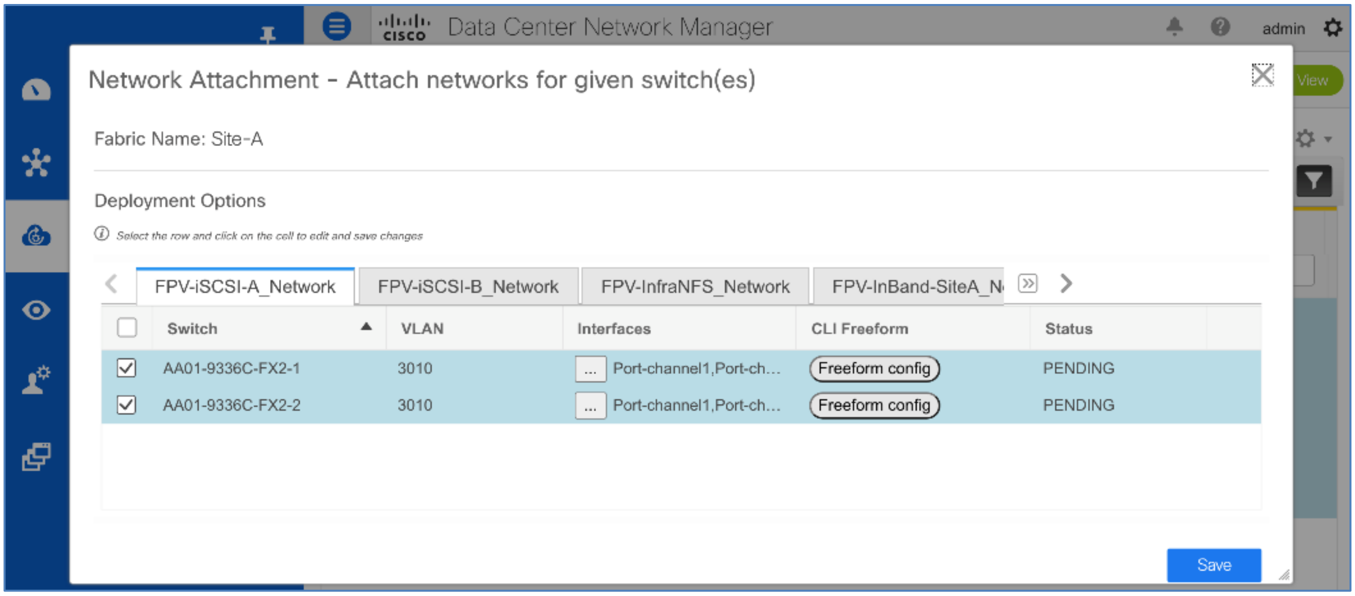

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ Select the row and click on the cell to edit and save changes

□	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	3010	... Port-channel1,Port-ch...	Freeform config	PENDING
☑	AA01-9336C-FX2-2	3010	... Port-channel1,Port-ch...	Freeform config	PENDING

13. For each network tab, repeat steps 7-9 to enable these networks on the access-layer connections to Cisco UCS Domain. Click **Save**. All switches are now in **PENDING** state. If you click on the pencil icon again, you will see that the ports or interfaces for the peer leaf switches are now configured as well.

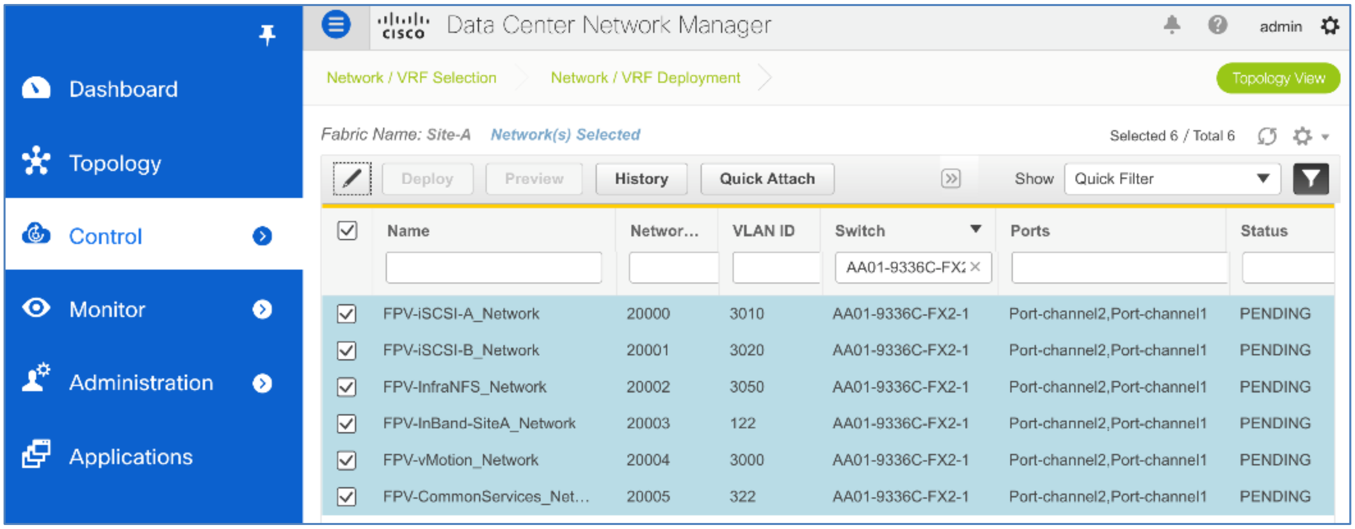

Fabric Name: Site-A *Network(s) Selected*

Selected 6 / Total 6

☑	Name	Networ...	VLAN ID	Switch	Ports	Status
				AA01-9336C-FX2 ×		
☑	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	Port-channel2,Port-channel1	PENDING
☑	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1	Port-channel2,Port-channel1	PENDING
☑	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1	Port-channel2,Port-channel1	PENDING
☑	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1	Port-channel2,Port-channel1	PENDING
☑	FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-1	Port-channel2,Port-channel1	PENDING
☑	FPV-CommonServices_Net...	20005	322	AA01-9336C-FX2-1	Port-channel2,Port-channel1	PENDING

14. Deselect the checkbox next to all switches. Click the **Preview** button. Note that all networks/VLANs are being enabled on the access layer connections to Cisco UCS domain. Click the **X** to close the Preview window.

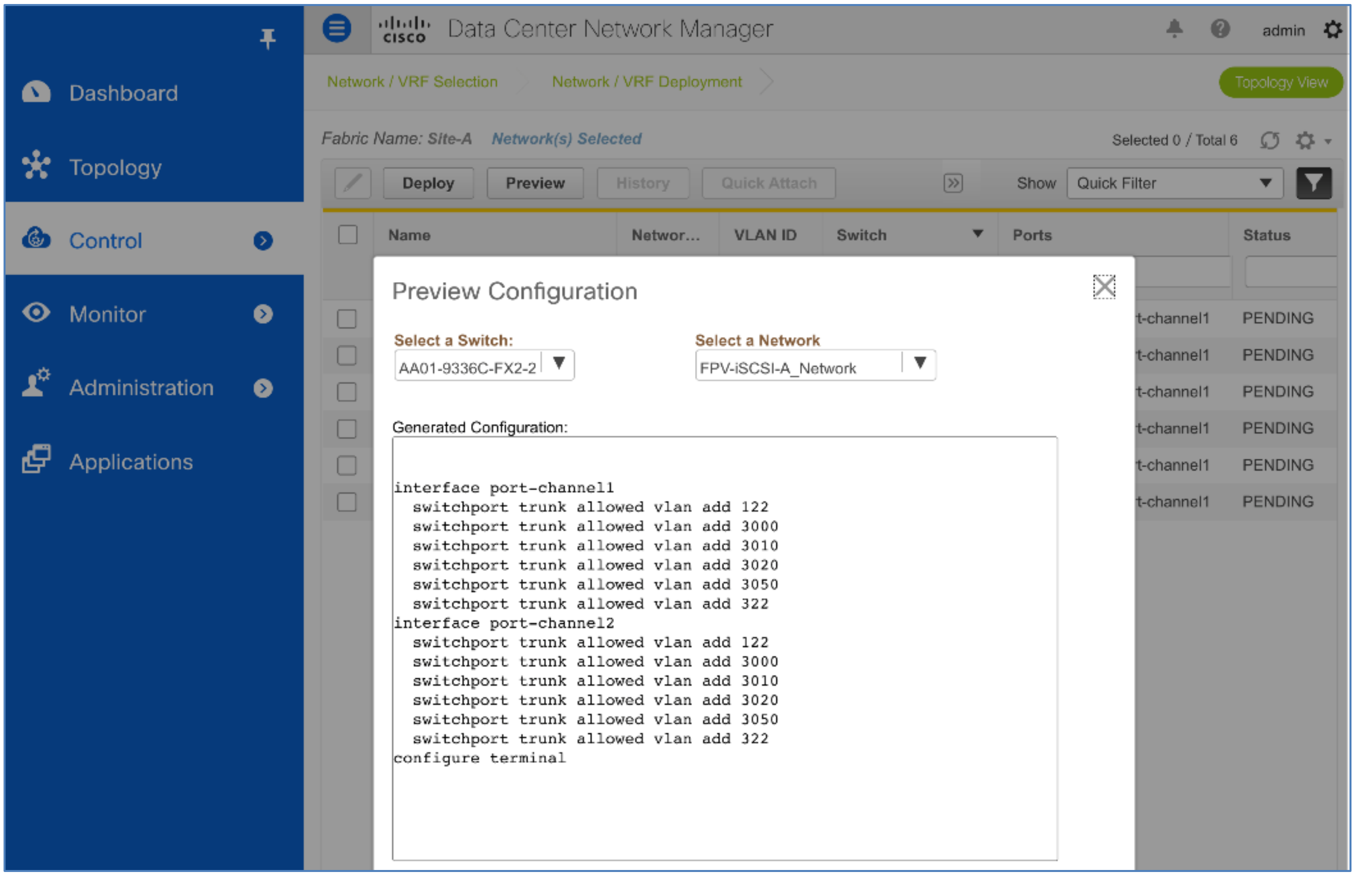

Fabric Name: Site-A *Network(s) Selected*

Selected 0 / Total 6

Preview Configuration

Select a Switch: AA01-9336C-FX2-2

Select a Network: FPV-iSCSI-A_Network

Generated Configuration:

```
interface port-channel1
  switchport trunk allowed vlan add 122
  switchport trunk allowed vlan add 3000
  switchport trunk allowed vlan add 3010
  switchport trunk allowed vlan add 3020
  switchport trunk allowed vlan add 3050
  switchport trunk allowed vlan add 322
interface port-channel2
  switchport trunk allowed vlan add 122
  switchport trunk allowed vlan add 3000
  switchport trunk allowed vlan add 3010
  switchport trunk allowed vlan add 3020
  switchport trunk allowed vlan add 3050
  switchport trunk allowed vlan add 322
configure terminal
```

15. Click the **Deploy** button. The status should change from **PENDING** to **IN PROGRESS** to **DEPLOYED.**

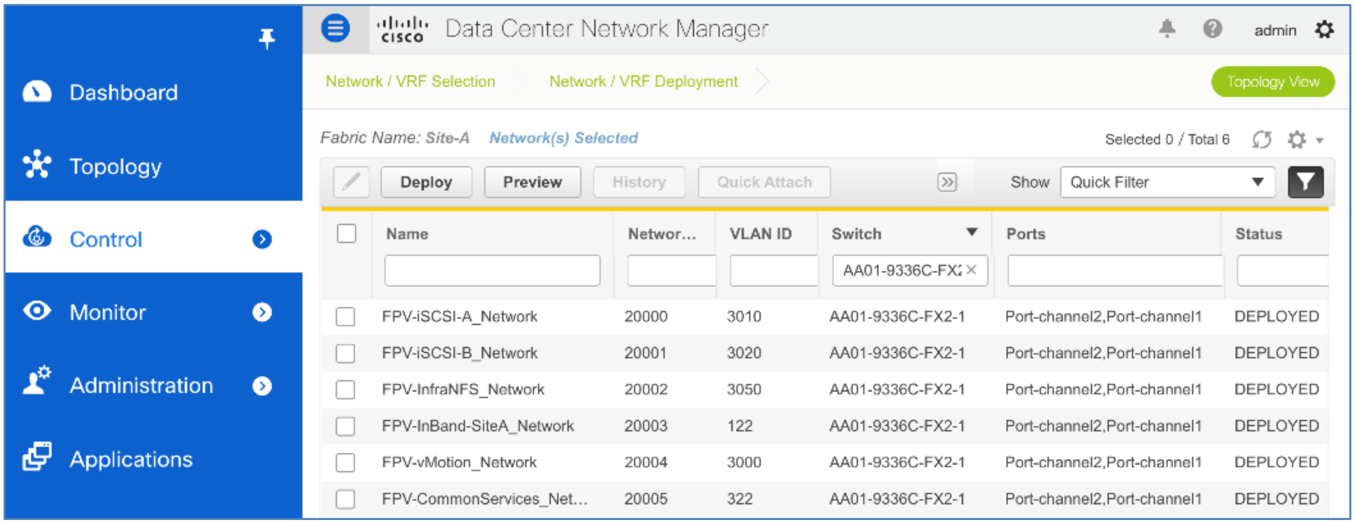

Fabric Name: Site-A ***Network(s) Selected***

Selected 0 / Total 6

□	Name	Networ...	VLAN ID	Switch	Ports	Status
				AA01-9336C-FX2		
□	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
□	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
□	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
□	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
□	FPV-vMotion_Network	20004	3000	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
□	FPV-CommonServices_Net…	20005	322	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED

Enable Infrastructure Networks on Access-Layer Connections to NetApp Storage Cluster

To enable infrastructure networks on the access-layer connections to the NetApp storage cluster, follow these steps:

1. Repeat steps 1-5 from the previous section. However, select the following infrastructure networks from the list below that needs to be enabled on the access-layer connection to the NetApp storage cluster.

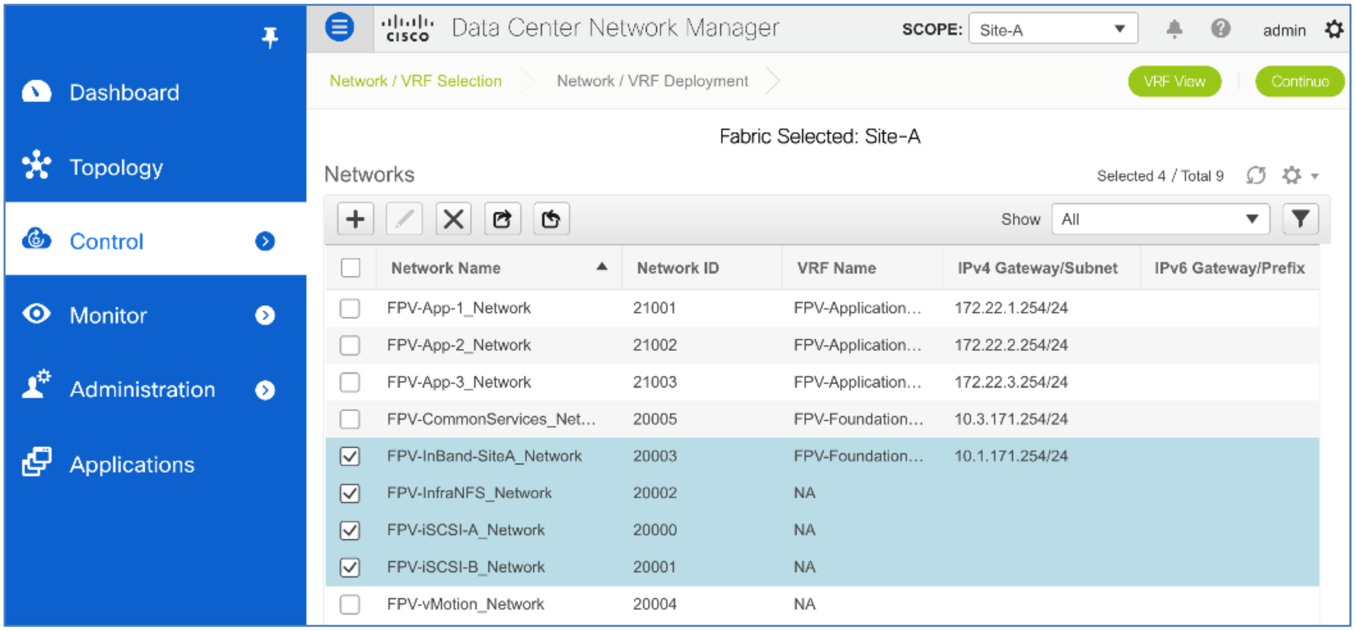

Fabric Selected: Site-A

Networks

Selected 4 / Total 9

□	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
□	FPV-App-1_Network	21001	FPV-Application…	172.22.1.254/24	
□	FPV-App-2_Network	21002	FPV-Application…	172.22.2.254/24	
□	FPV-App-3_Network	21003	FPV-Application…	172.22.3.254/24	
□	FPV-CommonServices_Net…	20005	FPV-Foundation…	10.3.171.254/24	
☑	FPV-InBand-SiteA_Network	20003	FPV-Foundation…	10.1.171.254/24	
☑	FPV-InfraNFS_Network	20002	NA		
☑	FPV-iSCSI-A_Network	20000	NA		
☑	FPV-iSCSI-B_Network	20001	NA		
□	FPV-vMotion_Network	20004	NA		

2. Filter on the first leaf switch in the vPC pair to the NetApp Storage cluster.

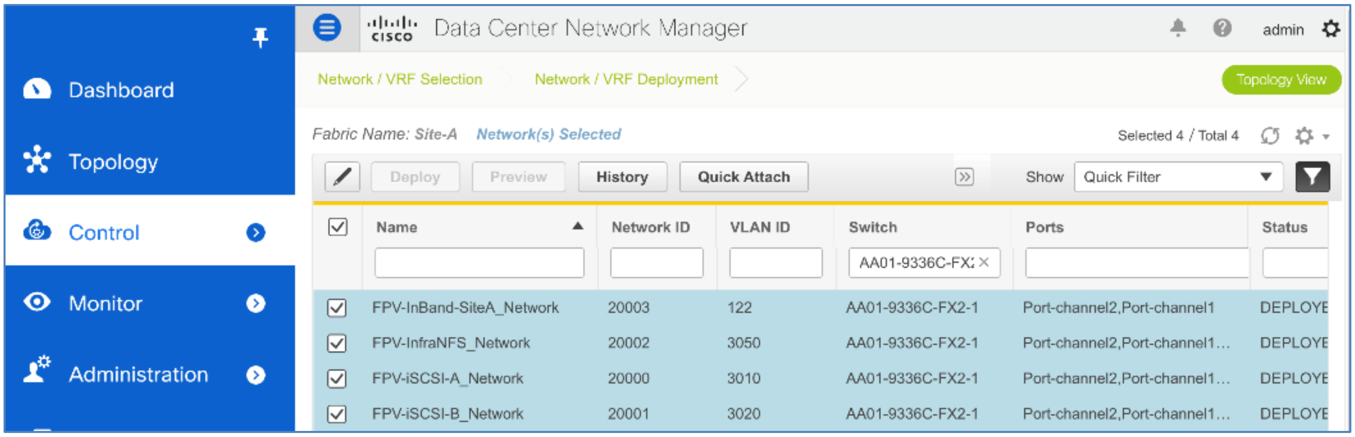

Fabric Name: Site-A *Network(s) Selected*

Selected 4 / Total 4

☑	Name	Network ID	VLAN ID	Switch	Ports	Status
				AA01-9336C-FX:		
☑	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYE
☑	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1	Port-channel2,Port-channel1…	DEPLOYE
☑	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	Port-channel2,Port-channel1…	DEPLOYE
☑	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1	Port-channel2,Port-channel1…	DEPLOYE

3. Click on the pencil icon from the menu.

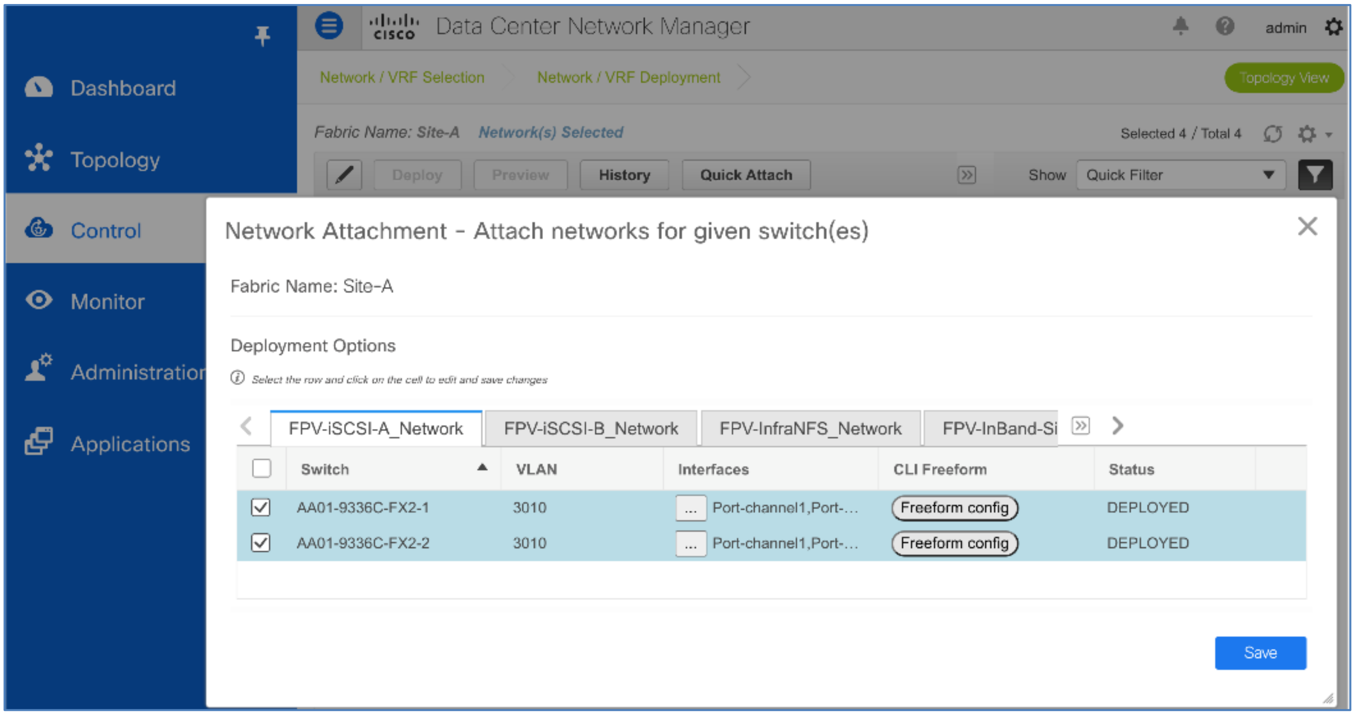

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

FPV-iSCSI-A_Network | FPV-iSCSI-B_Network | FPV-InfraNFS_Network | FPV-InBand-Si

□	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	3010	… Port-channel1,Port-…	Freeform config	DEPLOYED
☑	AA01-9336C-FX2-2	3010	… Port-channel1,Port-…	Freeform config	DEPLOYED

4. For the first network, click the box … in the **Interfaces** column next to the first switch listed in the leaf switch pair. Select the **Interfaces/Ports** that connect to the NetApp Storage cluster. In this case, two port-channels going to Cisco UCS domain were already configured for these networks – however, two additional port-channels going to NetApp had to be selected in this step.

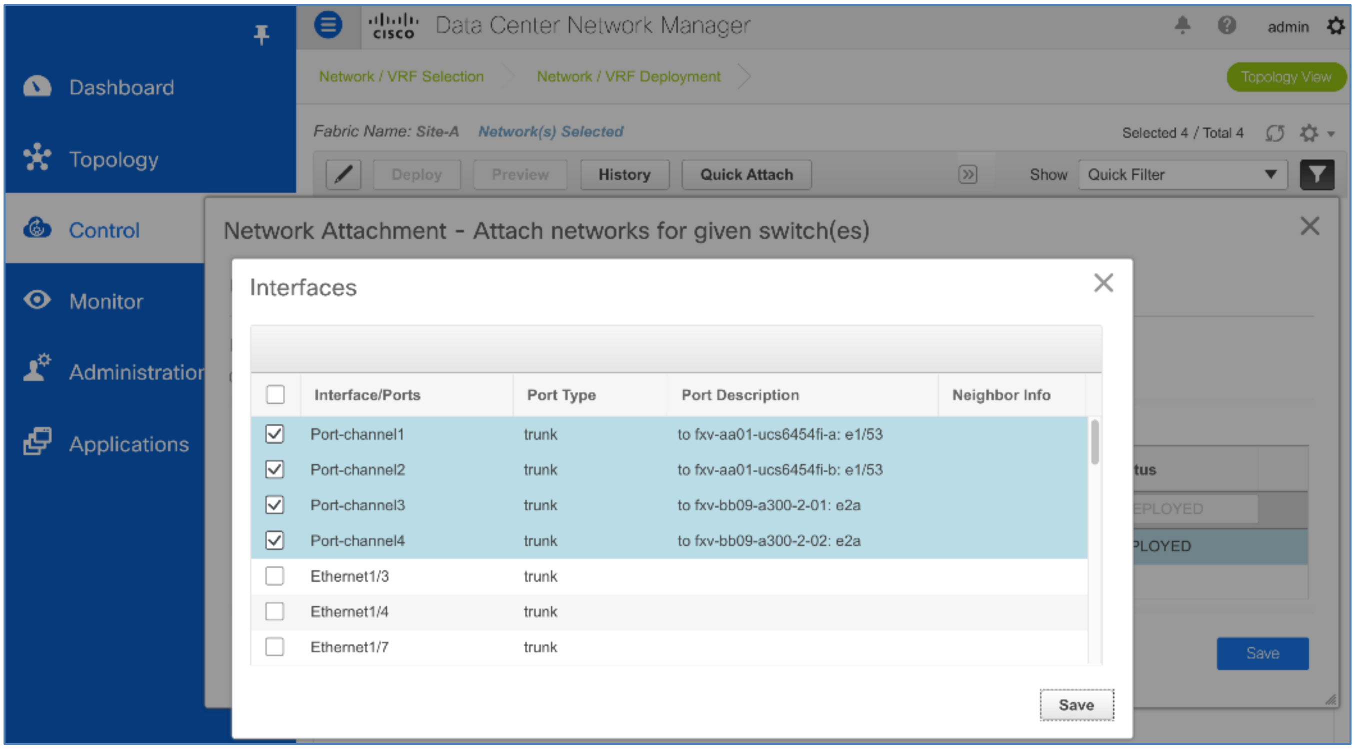

Fabric Name: Site-A *Network(s) Selected* Selected 4 / Total 4

Network Attachment - Attach networks for given switch(es)

Interfaces

☐	Interface/Ports	Port Type	Port Description	Neighbor Info
☑	Port-channel1	trunk	to fxv-aa01-ucs6454fi-a: e1/53	
☑	Port-channel2	trunk	to fxv-aa01-ucs6454fi-b: e1/53	
☑	Port-channel3	trunk	to fxv-bb09-a300-2-01: e2a	
☑	Port-channel4	trunk	to fxv-bb09-a300-2-02: e2a	
☐	Ethernet1/3	trunk		
☐	Ethernet1/4	trunk		
☐	Ethernet1/7	trunk		

Save

5. Click **Save**.

6. Repeat steps 4-5 for each network tab. Click **Save.** Note that the status of these networks are in **PENDING** state at this stage. Scroll to the right as needed to see all columns available in this view.

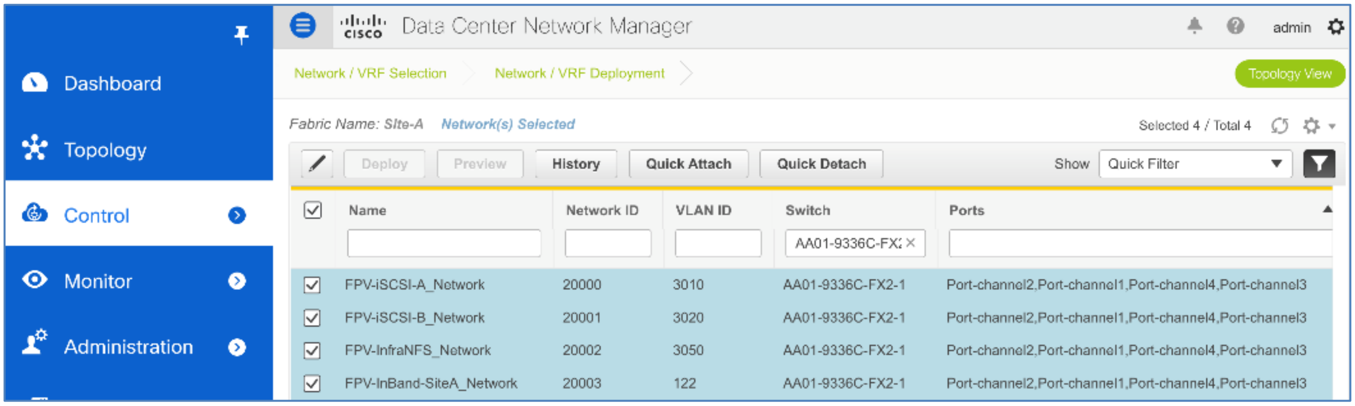

Fabric Name: Site-A *Network(s) Selected* Selected 4 / Total 4

☑	Name	Network ID	VLAN ID	Switch	Ports
				AA01-9336C-FX2 ×	
☑	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3
☑	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3
☑	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3
☑	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3

7. Deselect all networks. Click the **Preview** button to see the pending changes. Click the **X** to close the window.

8. Click the **Deploy** button. The status should go from **PENDING** to **IN PROGRESS** to **DEPLOYED.** Scroll to the right as needed to see all columns available in this view.

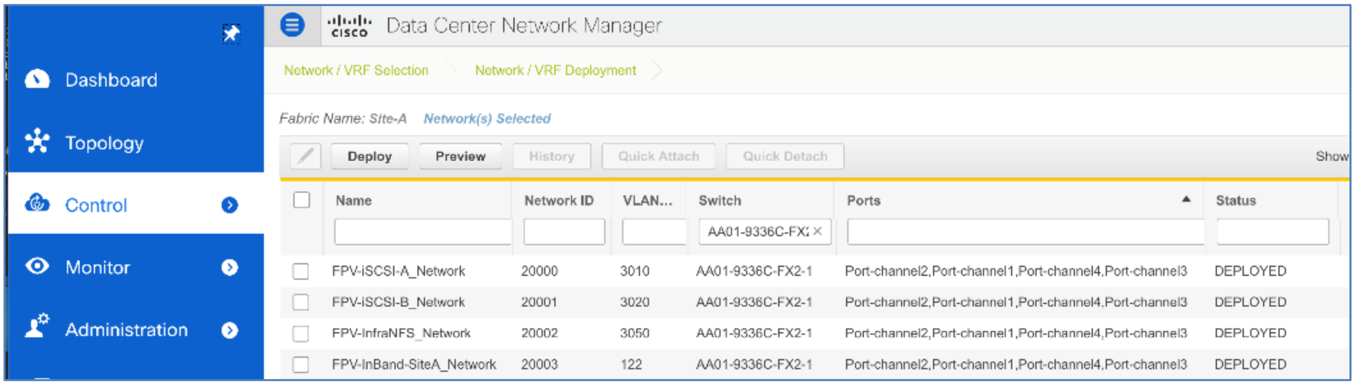

Fabric Name: Site-A *Network(s) Selected*

☐	Name	Network ID	VLAN...	Switch	Ports	Status
				AA01-9336C-FX2		
☐	FPV-iSCSI-A_Network	20000	3010	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED
☐	FPV-iSCSI-B_Network	20001	3020	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED
☐	FPV-InfraNFS_Network	20002	3050	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED
☐	FPV-InBand-SiteA_Network	20003	122	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED

Enable External Connectivity for FlexPod Infrastructure Networks

To enable access to external/outside networks from the FlexPod infrastructure tenant, follow these steps:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > VRFs**. Click **OK** to exit any pop-ups that come up. Select the correct scope in the drop-down list next to **Scope:** in the top-right corner of the window. Select the Tenants that need access to external or outside networks. Click the **Continue** button.

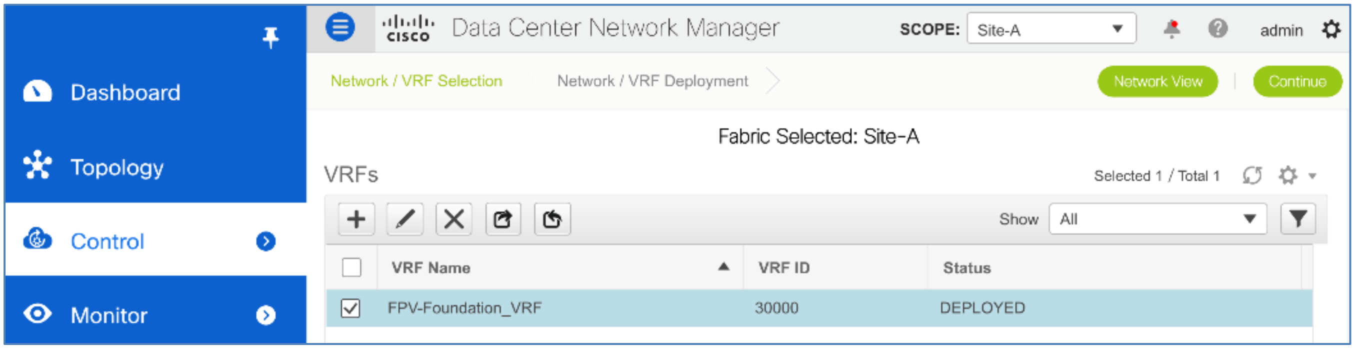

SCOPE: Site-A

Fabric Selected: Site-A

Show: All

☐	VRF Name	VRF ID	Status
☑	FPV-Foundation_VRF	30000	DEPLOYED

3. Click the **Detailed View** button.

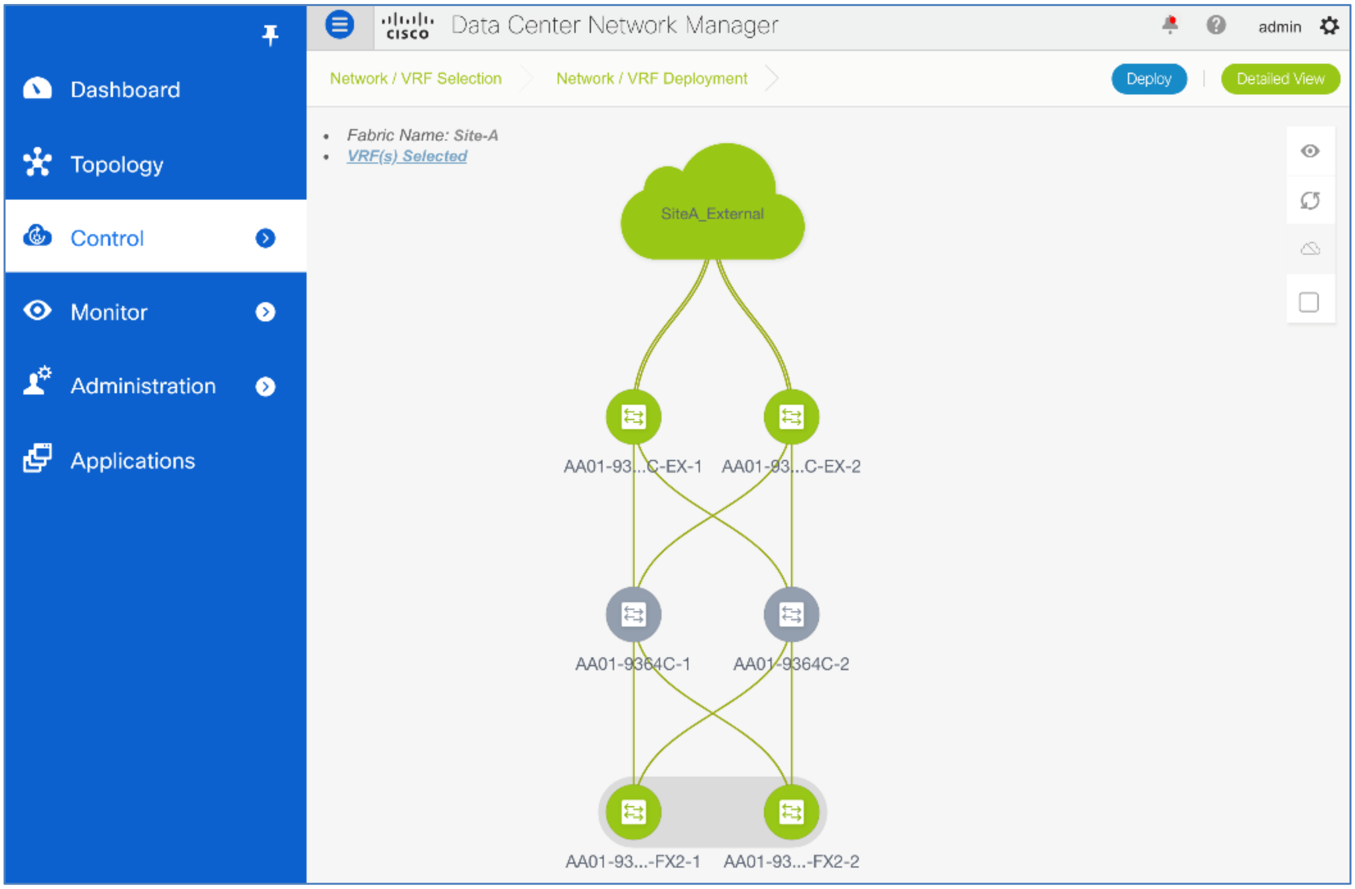

4. Select and click the arrow on the **Switch** column to sort based on hostname. Select the two border switches that provide connectivity to external/outside networks. Click on the pencil icon to edit the Tenant VRF.

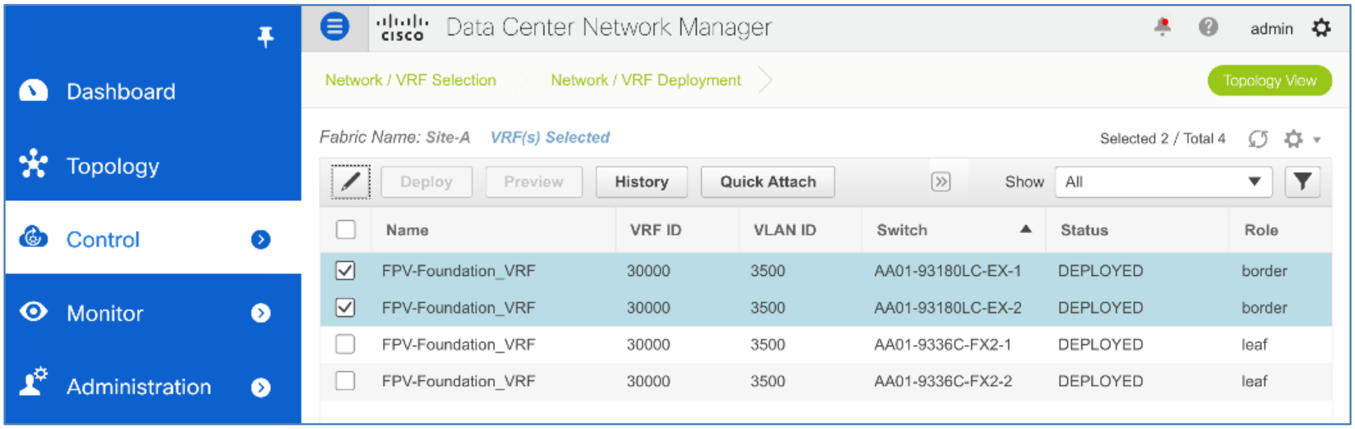

☐	Name	VRF ID	VLAN ID	Switch ▲	Status	Role
☑	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-1	DEPLOYED	border
☑	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-2	DEPLOYED	border
☐	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-1	DEPLOYED	leaf
☐	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-2	DEPLOYED	leaf

5. In the **VRF Extension Attachment – Attach extensions** pop-up window, click on the icon in the **Extend** column for the first Border switch. Scroll to the right as needed to see all columns available in this view.

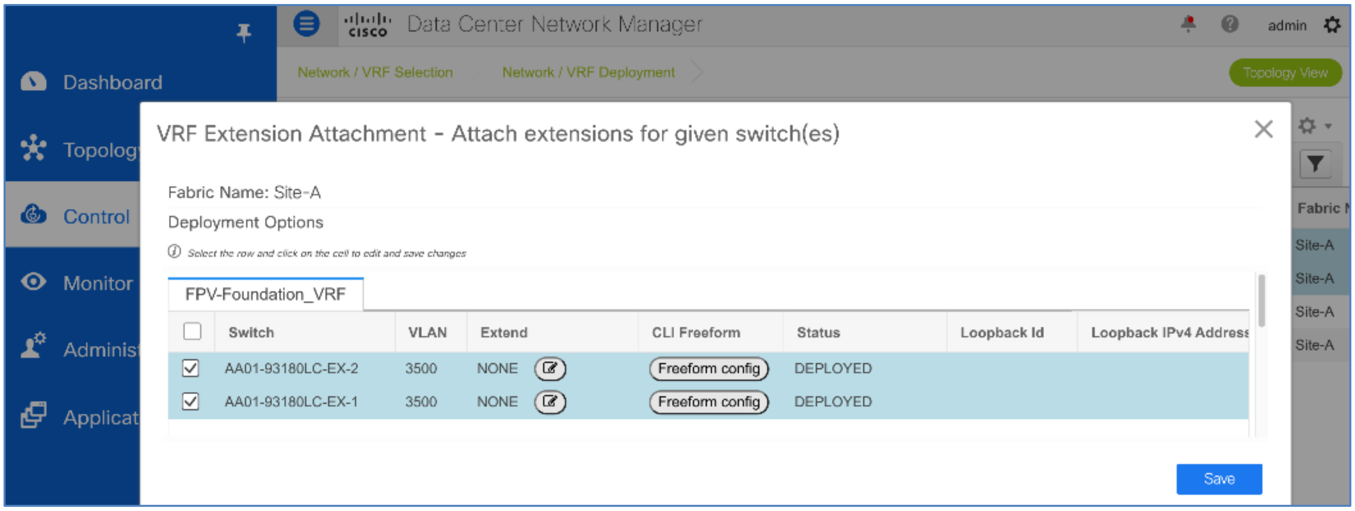

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

FPV-Foundation_VRF

□	Switch	VLAN	Extend	CLI Freeform	Status	Loopback Id	Loopback IPv4 Address
☑	AA01-93180LC-EX-2	3500	NONE	Freeform config	DEPLOYED		
☑	AA01-93180LC-EX-1	3500	NONE	Freeform config	DEPLOYED		

6. From the **Extend** column, select **VRF-Lite** from the drop-down list.

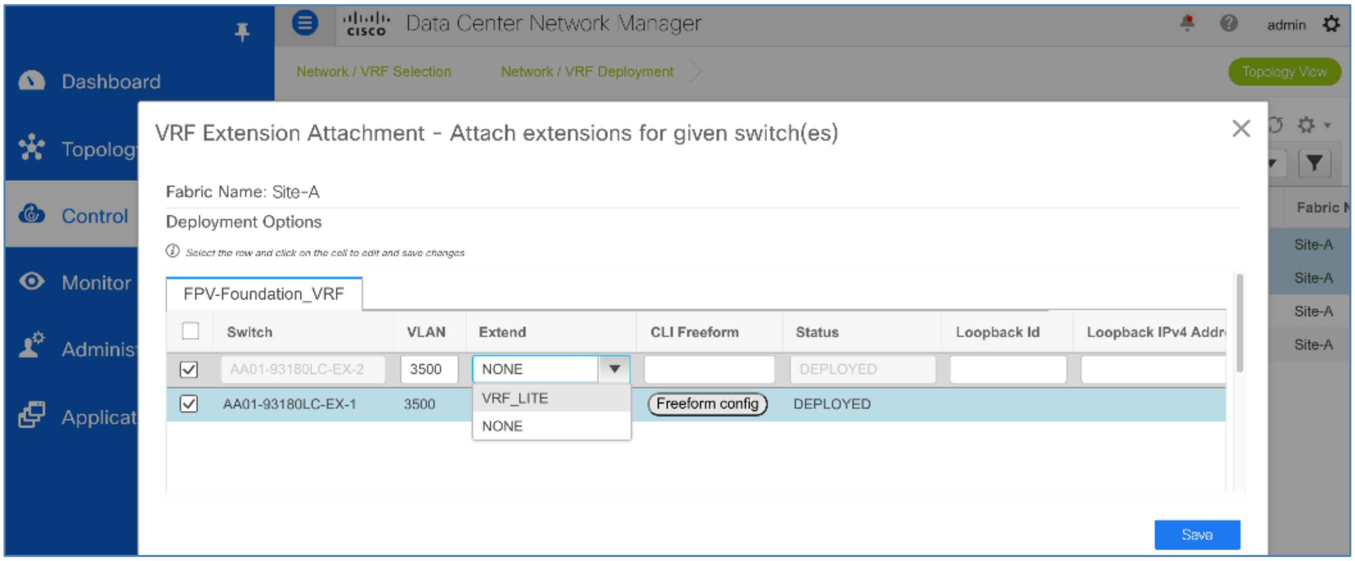

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

FPV-Foundation_VRF

□	Switch	VLAN	Extend	CLI Freeform	Status	Loopback Id	Loopback IPv4 Addr
☑	AA01-93180LC-EX-2	3500	NONE VRF_LITE NONE		DEPLOYED		
☑	AA01-93180LC-EX-1	3500		Freeform config	DEPLOYED		

7. The window will now expand to include the interfaces that can be used to extend the VRF using VRF-Lite to the External Gateway. In the **Extension Details** section, select all relevant switches and interfaces where the VRF should be extended. Scroll to the right as needed to see all columns available in this view.

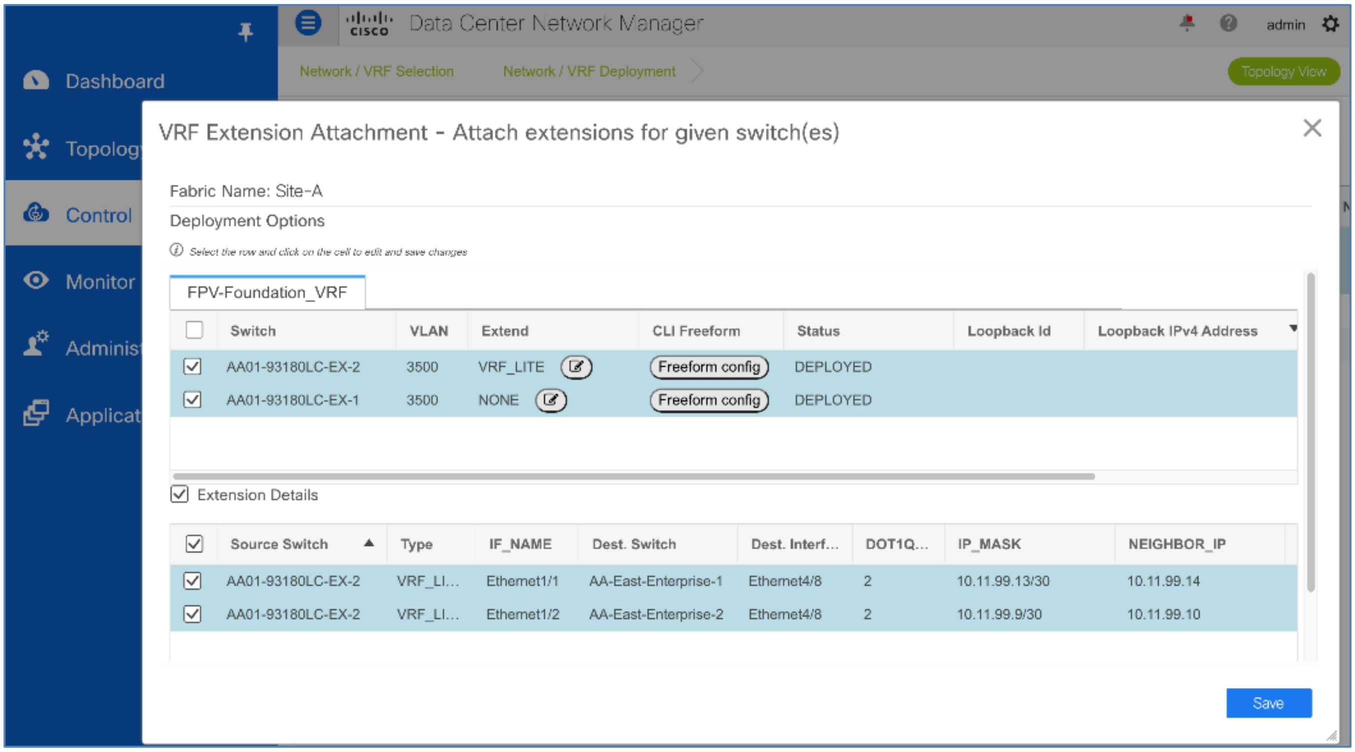

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ Select the row and click on the cell to edit and save changes

FPV-Foundation_VRF

□	Switch	VLAN	Extend	CLI Freeform	Status	Loopback Id	Loopback IPv4 Address
☑	AA01-93180LC-EX-2	3500	VRF_LITE	Freeform config	DEPLOYED		
☑	AA01-93180LC-EX-1	3500	NONE	Freeform config	DEPLOYED		

☑ Extension Details

☑	Source Switch	Type	IF_NAME	Dest. Switch	Dest. Interf...	DOT1Q...	IP_MASK	NEIGHBOR_IP
☑	AA01-93180LC-EX-2	VRF_LI...	Ethernet1/1	AA-East-Enterprise-1	Ethernet4/8	2	10.11.99.13/30	10.11.99.14
☑	AA01-93180LC-EX-2	VRF_LI...	Ethernet1/2	AA-East-Enterprise-2	Ethernet4/8	2	10.11.99.9/30	10.11.99.10

8. From the **PEER_VRF_NAME** column, specify the name that should be used for the VRF in the External Gateway switch that connects to this Border switch. Repeat this step for each Border switch in the list.

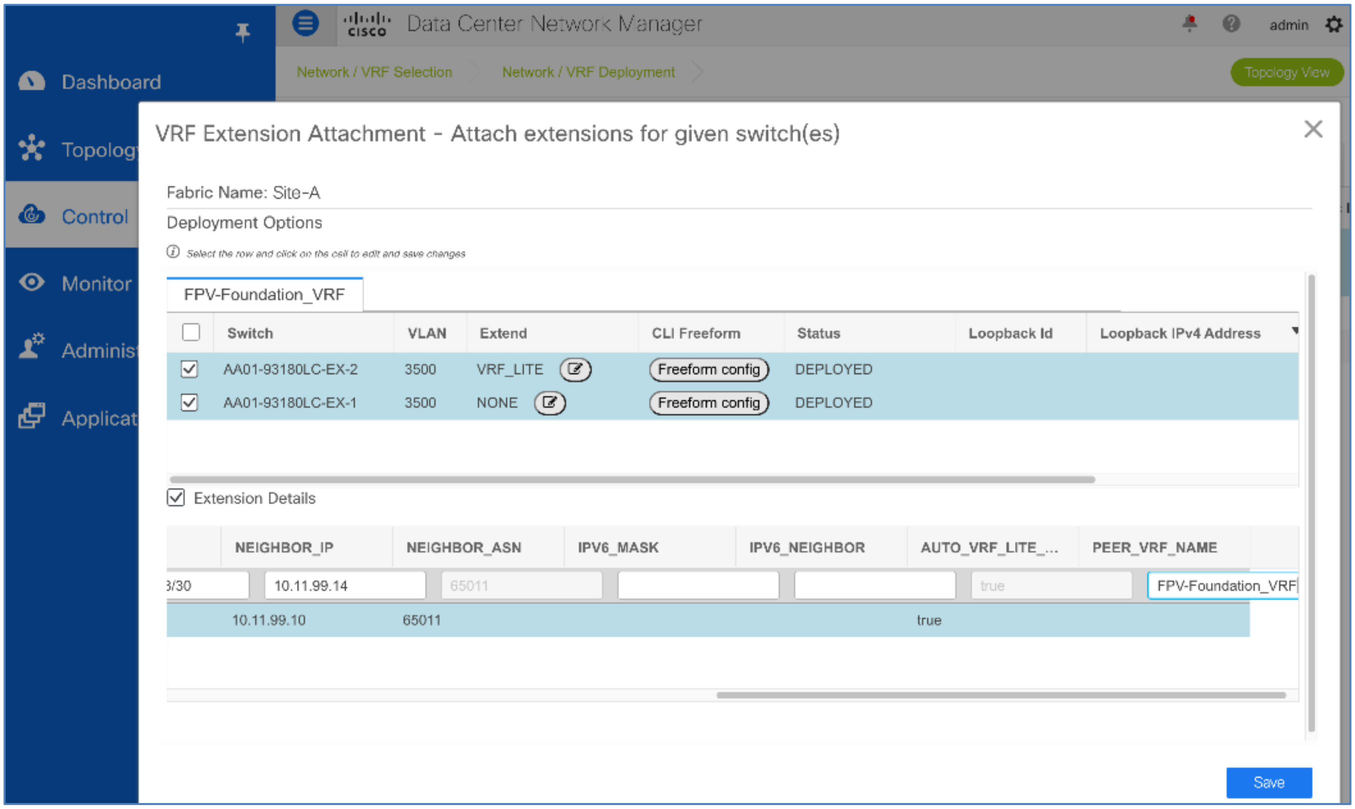

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ Select the row and click on the cell to edit and save changes

FPV-Foundation_VRF

□	Switch	VLAN	Extend	CLI Freeform	Status	Loopback Id	Loopback IPv4 Address
☑	AA01-93180LC-EX-2	3500	VRF_LITE	Freeform config	DEPLOYED		
☑	AA01-93180LC-EX-1	3500	NONE	Freeform config	DEPLOYED		

☑ Extension Details

	NEIGHBOR_IP	NEIGHBOR_ASN	IPV6_MASK	IPV6_NEIGHBOR	AUTO_VRF_LITE_...	PEER_VRF_NAME
3/30	10.11.99.14	65011			true	FPV-Foundation_VRF
	10.11.99.10	65011			true	

9. Repeat steps 5-8 for the second Border switch in the list.

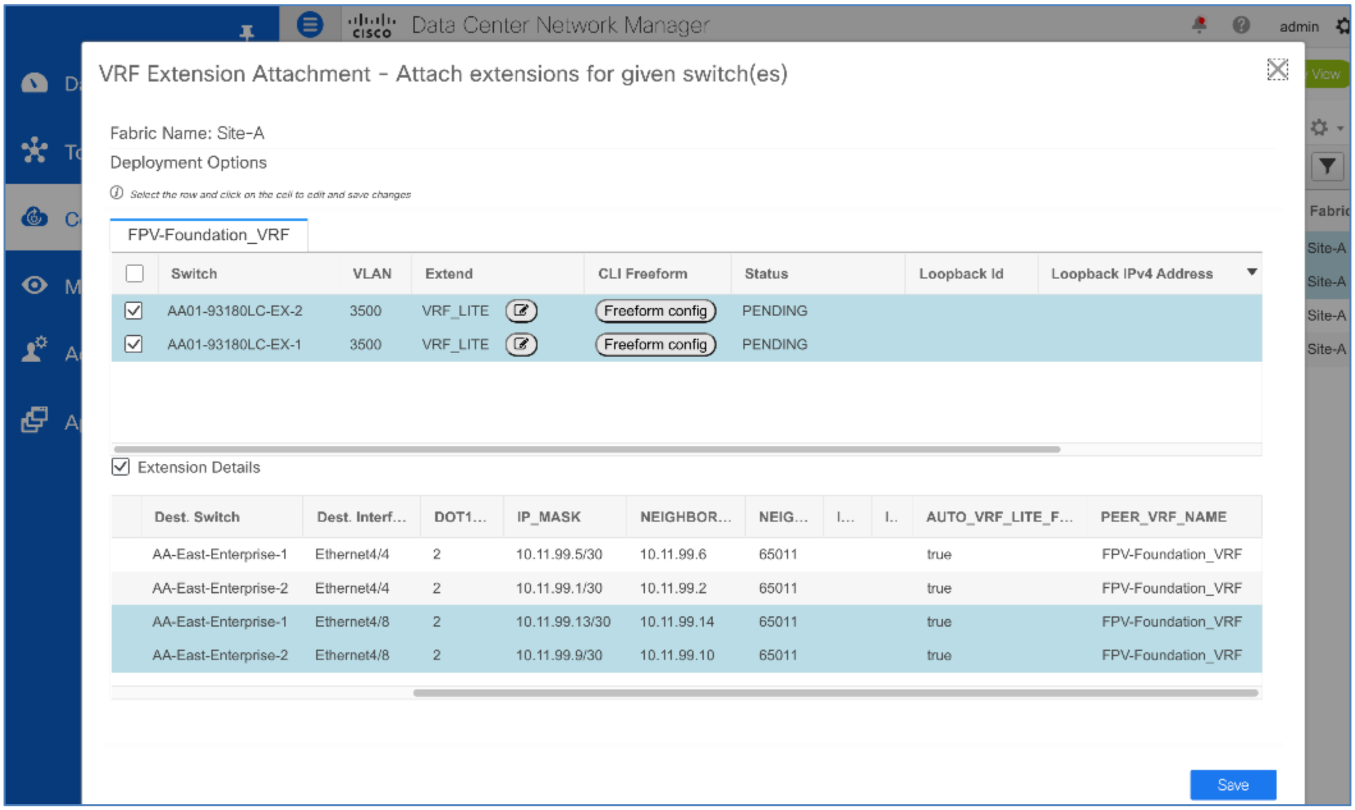

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

FPV-Foundation_VRF

□	Switch	VLAN	Extend	CLI Freeform	Status	Loopback Id	Loopback IPv4 Address
☑	AA01-93180LC-EX-2	3500	VRF_LITE	Freeform config	PENDING		
☑	AA01-93180LC-EX-1	3500	VRF_LITE	Freeform config	PENDING		

☑ Extension Details

Dest. Switch	Dest. Interf...	DOT1...	IP_MASK	NEIGHBOR...	NEIG...	I...	I..	AUTO_VRF_LITE_F...	PEER_VRF_NAME
AA-East-Enterprise-1	Ethernet4/4	2	10.11.99.5/30	10.11.99.6	65011			true	FPV-Foundation_VRF
AA-East-Enterprise-2	Ethernet4/4	2	10.11.99.1/30	10.11.99.2	65011			true	FPV-Foundation_VRF
AA-East-Enterprise-1	Ethernet4/8	2	10.11.99.13/30	10.11.99.14	65011			true	FPV-Foundation_VRF
AA-East-Enterprise-2	Ethernet4/8	2	10.11.99.9/30	10.11.99.10	65011			true	FPV-Foundation_VRF

Save

10. Click **Save**. The VRF configuration on the Border switches are now in a **PENDING** state.

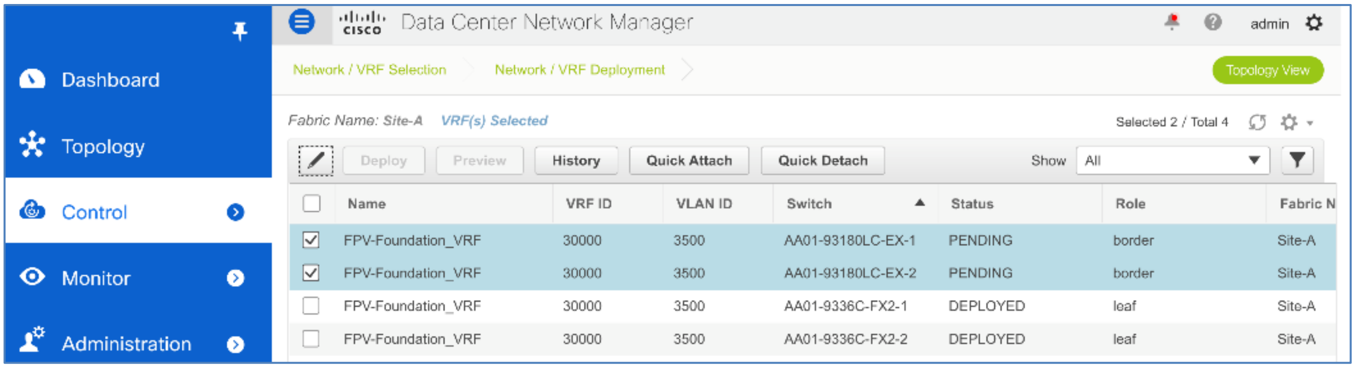

Fabric Name: Site-A VRF(s) Selected

Selected 2 / Total 4

Show All

□	Name	VRF ID	VLAN ID	Switch	Status	Role	Fabric N
☑	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-1	PENDING	border	Site-A
☑	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-2	PENDING	border	Site-A
□	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-1	DEPLOYED	leaf	Site-A
□	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-2	DEPLOYED	leaf	Site-A

11. Deselect both Border switches. Click the **Preview** button to view the pending configuration changes.

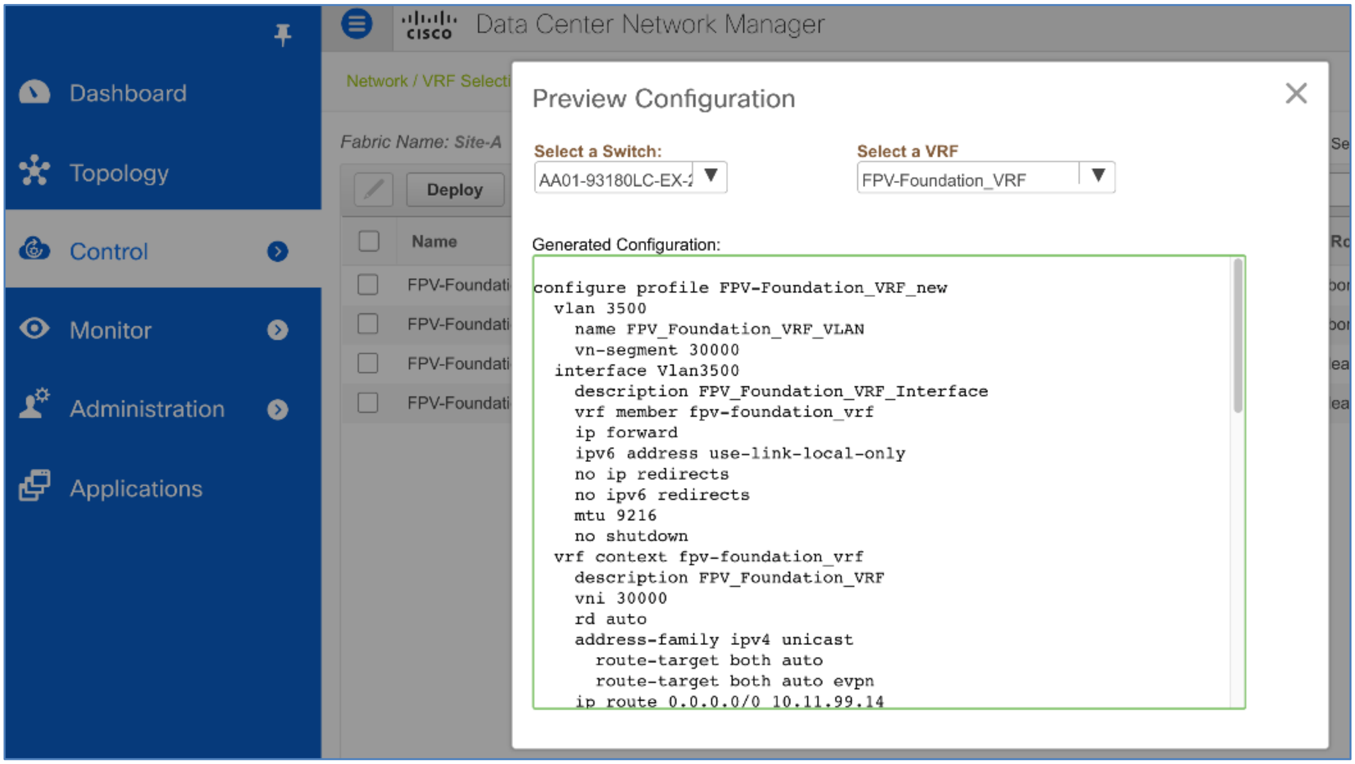

Preview Configuration

Select a Switch: AA01-93180LC-EX-

Select a VRF: FPV-Foundation_VRF

Generated Configuration:

```
configure profile FPV-Foundation_VRF_new
  vlan 3500
    name FPV_Foundation_VRF_VLAN
    vn-segment 30000
  interface Vlan3500
    description FPV_Foundation_VRF_Interface
    vrf member fpv-foundation_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-foundation_vrf
    description FPV_Foundation_VRF
    vni 30000
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    ip route 0.0.0.0/0 10.11.99.14
```

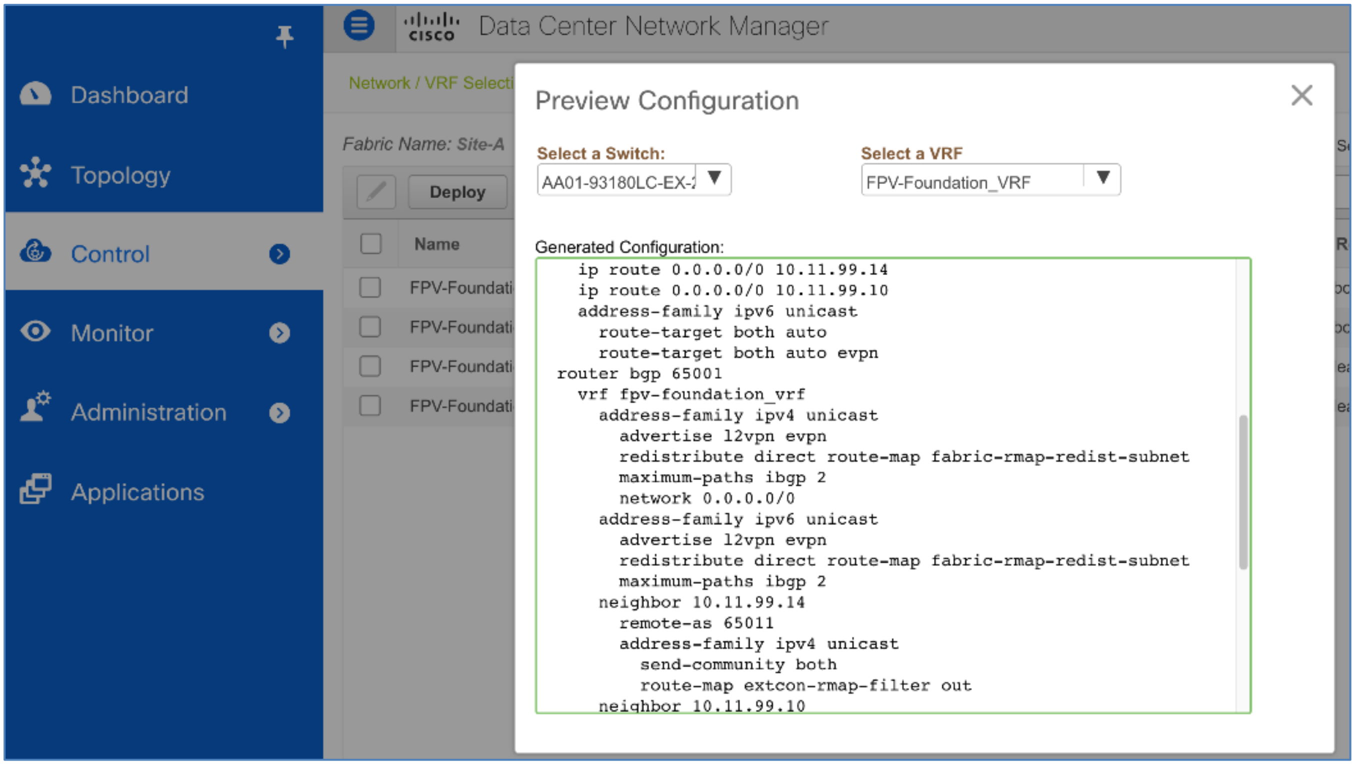

Preview Configuration

Select a Switch: AA01-93180LC-EX-

Select a VRF: FPV-Foundation_VRF

Generated Configuration:

```
    ip route 0.0.0.0/0 10.11.99.14
    ip route 0.0.0.0/0 10.11.99.10
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  router bgp 65001
    vrf fpv-foundation_vrf
      address-family ipv4 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
        network 0.0.0.0/0
      address-family ipv6 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
      neighbor 10.11.99.14
        remote-as 65011
        address-family ipv4 unicast
          send-community both
          route-map extcon-rmap-filter out
      neighbor 10.11.99.10
```

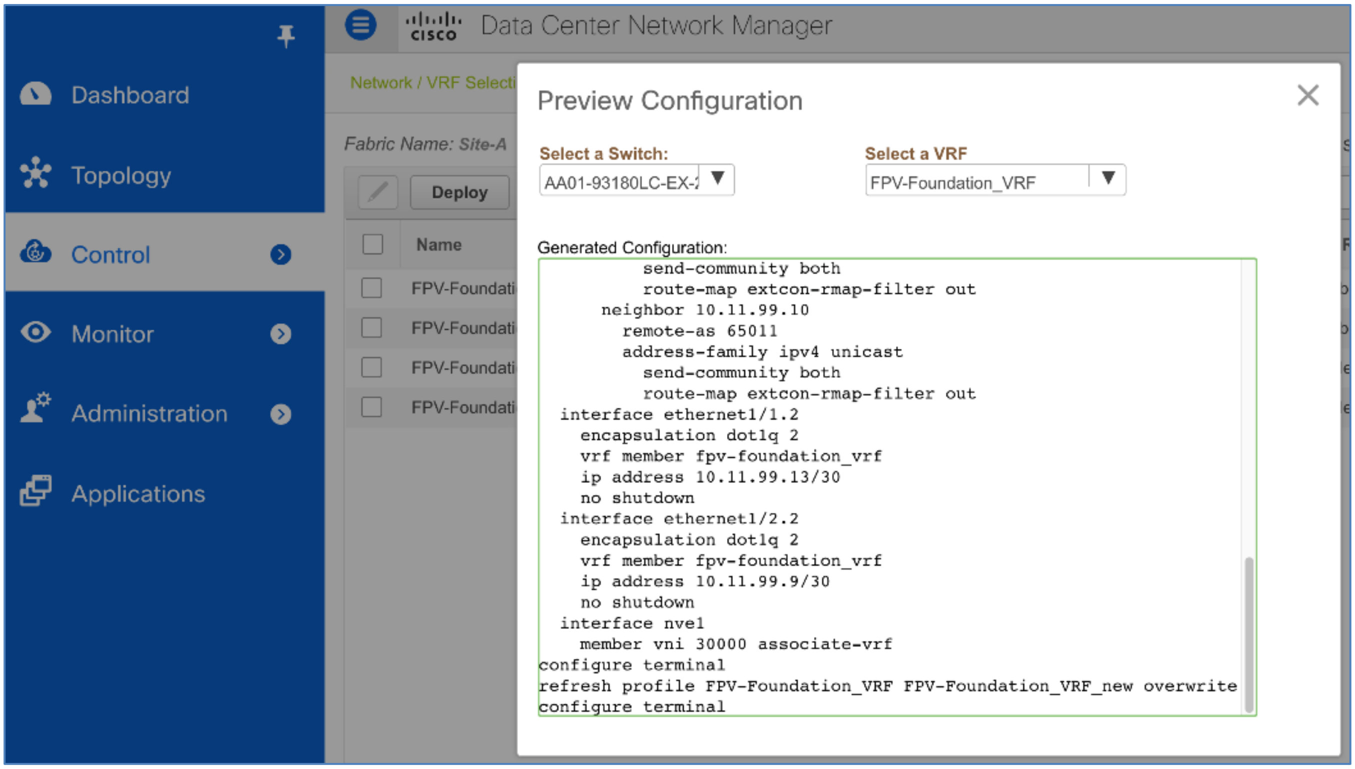

Preview Configuration

Select a Switch: AA01-93180LC-EX-

Select a VRF: FPV-Foundation_VRF

Generated Configuration:

```
          send-community both
          route-map extcon-rmap-filter out
      neighbor 10.11.99.10
        remote-as 65011
        address-family ipv4 unicast
          send-community both
          route-map extcon-rmap-filter out
  interface ethernet1/1.2
    encapsulation dot1q 2
    vrf member fpv-foundation_vrf
    ip address 10.11.99.13/30
    no shutdown
  interface ethernet1/2.2
    encapsulation dot1q 2
    vrf member fpv-foundation_vrf
    ip address 10.11.99.9/30
    no shutdown
  interface nve1
    member vni 30000 associate-vrf
configure terminal
refresh profile FPV-Foundation_VRF FPV-Foundation_VRF_new overwrite
configure terminal
```

12. Click the **X** to close the **Preview** window. Click the **Deploy** button. The **Status** column should go from **PENDING** to **IN PROGRESS** to **DEPLOYED**.

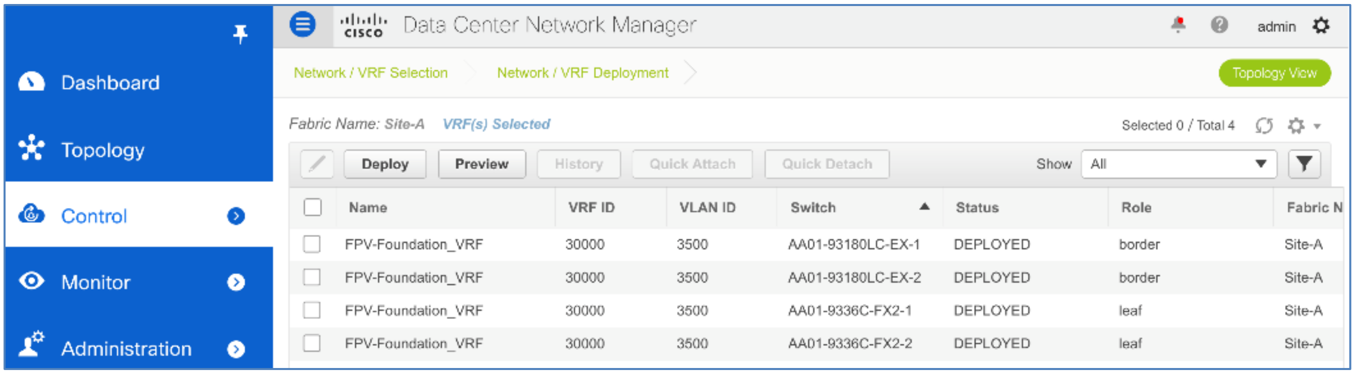

Fabric Name: Site-A *VRF(s) Selected*

	Name	VRF ID	VLAN ID	Switch	Status	Role	Fabric N
□	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-1	DEPLOYED	border	Site-A
□	FPV-Foundation_VRF	30000	3500	AA01-93180LC-EX-2	DEPLOYED	border	Site-A
□	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-1	DEPLOYED	leaf	Site-A
□	FPV-Foundation_VRF	30000	3500	AA01-9336C-FX2-2	DEPLOYED	leaf	Site-A

13. From the left navigation bar, select **Control > Fabrics > Fabric Builder**.

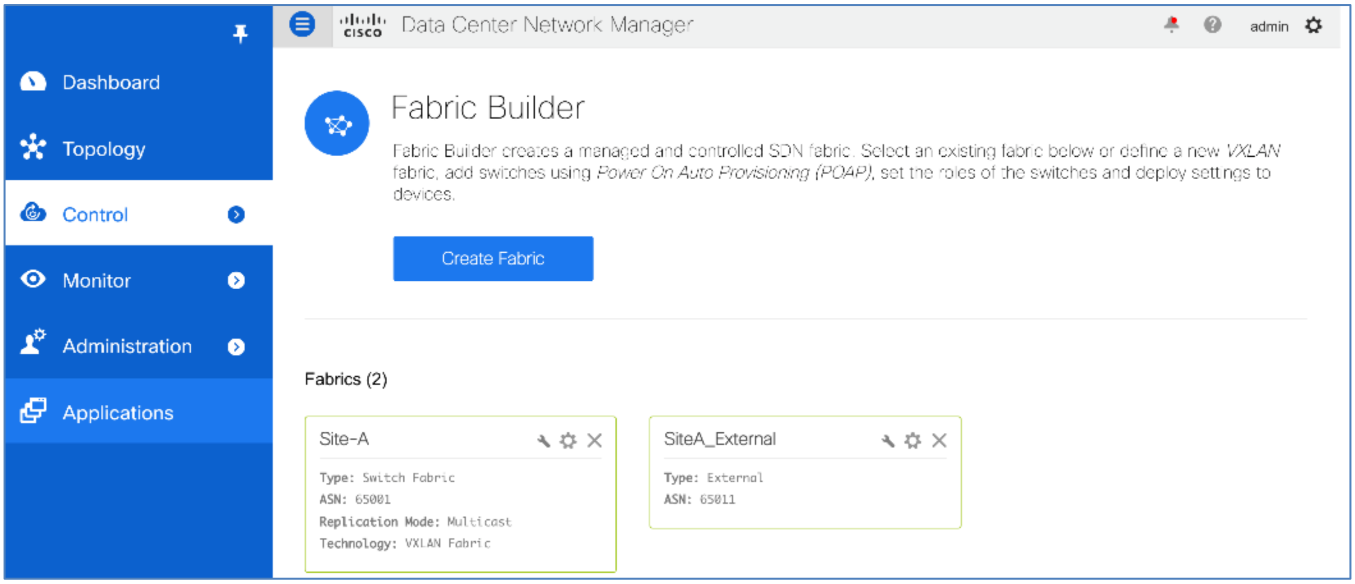

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (2)

Site-A

Type: Switch Fabric
ASN: 65001
Replication Mode: Multicast
Technology: VXLAN Fabric

SiteA_External

Type: External
ASN: 65011

14. Select the external fabric from the list. Note that the external gateways are **Out-of-Sync/Failed** state.

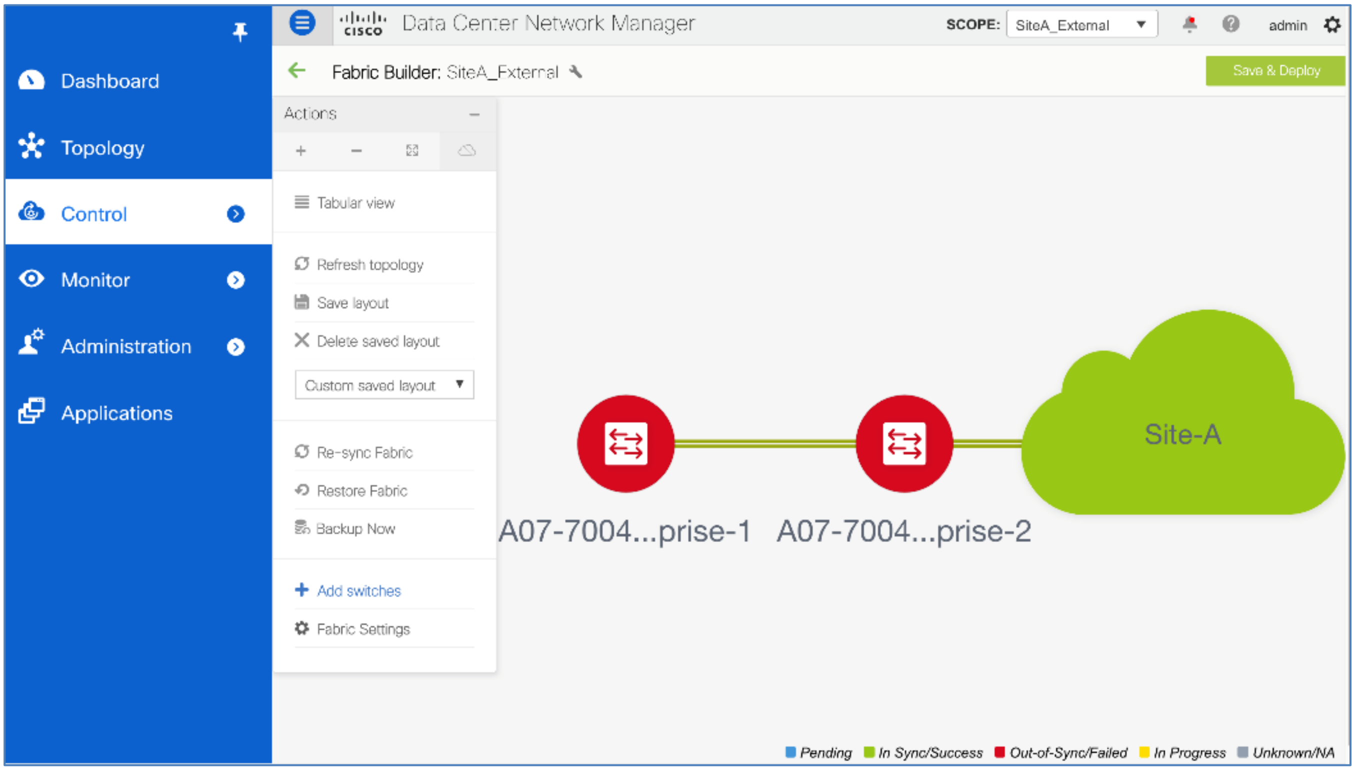

SCOPE: SiteA_External

Fabric Builder: SiteA_External

15. Click the **Save & Deploy** button. Note that there are **'30 lines'** of changes.

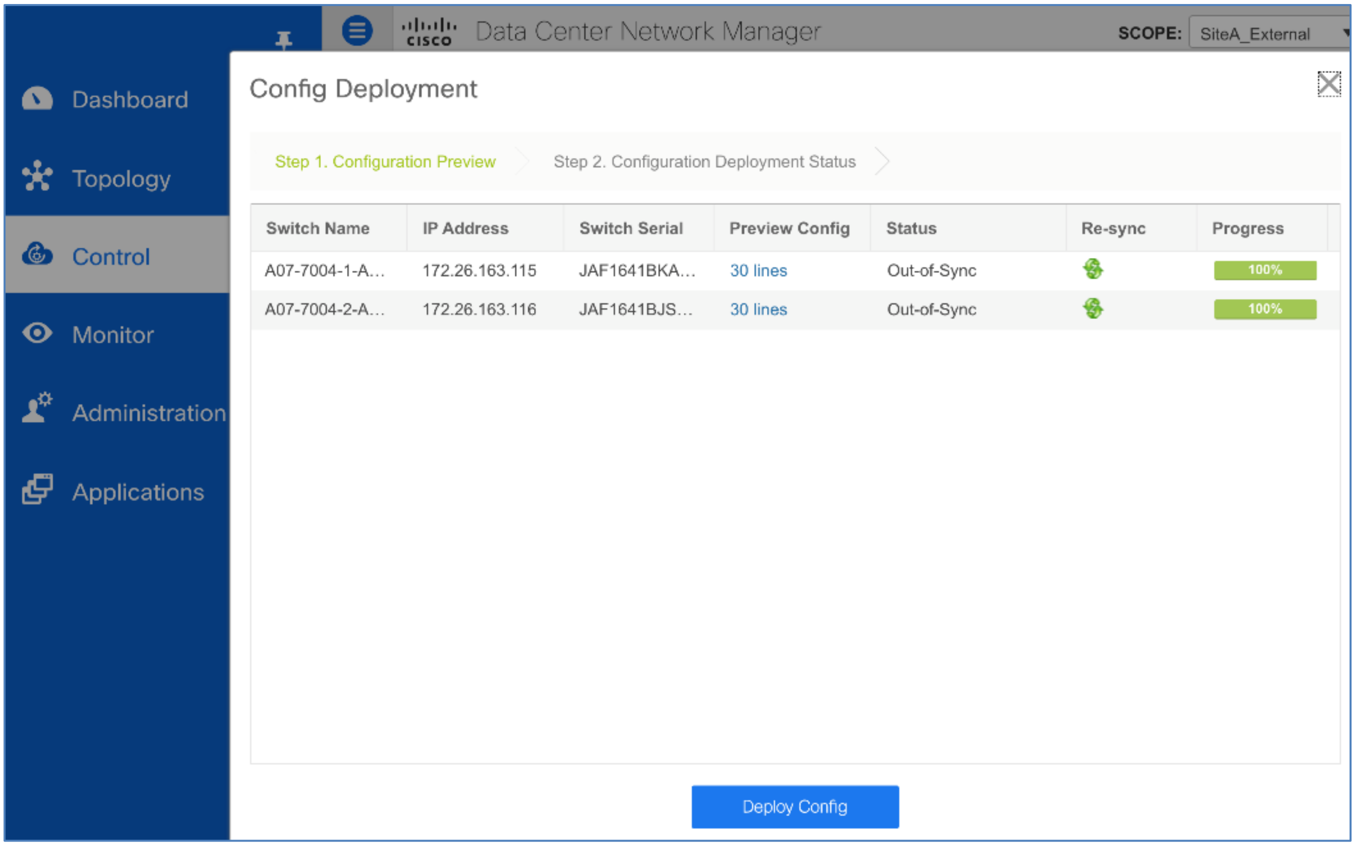

Config Deployment

Step 1. Configuration Preview | Step 2. Configuration Deployment Status

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
A07-7004-1-A…	172.26.163.115	JAF1641BKA…	30 lines	Out-of-Sync		100%
A07-7004-2-A…	172.26.163.116	JAF1641BJS…	30 lines	Out-of-Sync		100%

16. Click on the **’30 lines’** to preview the pending changes.

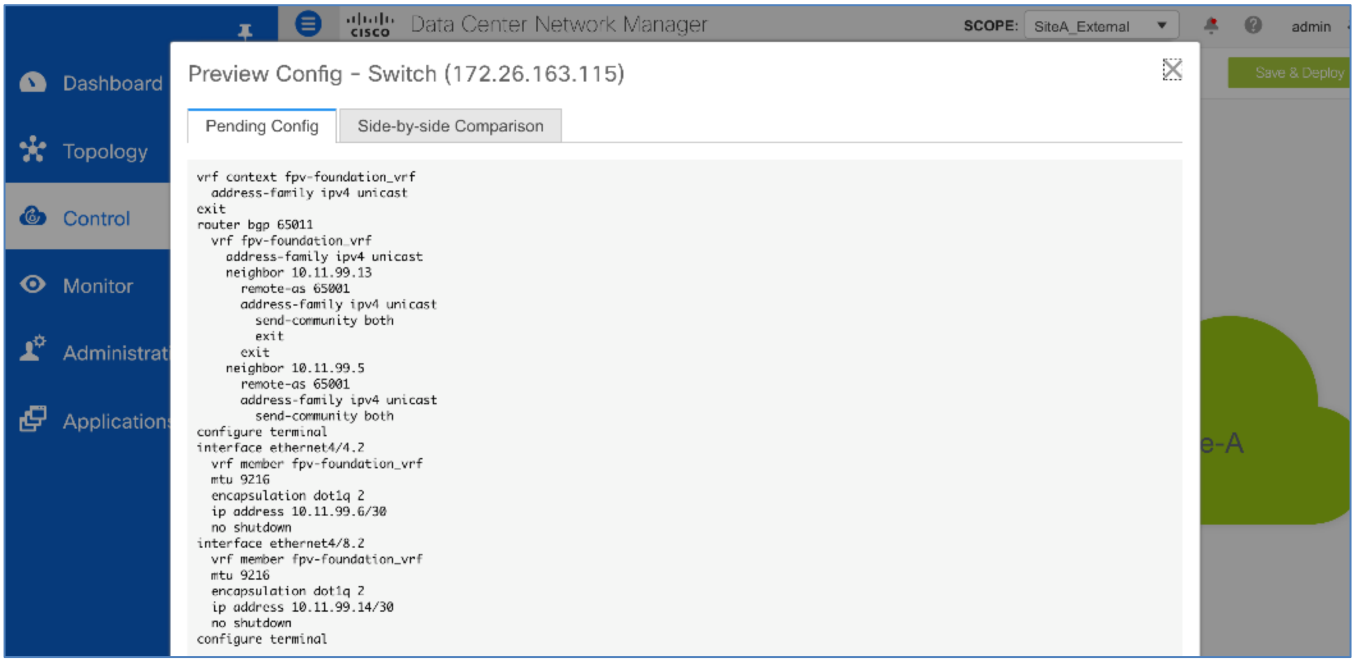

Preview Config - Switch (172.26.163.115)

```
vrf context fpv-foundation_vrf
  address-family ipv4 unicast
exit
router bgp 65011
  vrf fpv-foundation_vrf
    address-family ipv4 unicast
    neighbor 10.11.99.13
      remote-as 65001
      address-family ipv4 unicast
        send-community both
        exit
      exit
    neighbor 10.11.99.5
      remote-as 65001
      address-family ipv4 unicast
        send-community both
configure terminal
interface ethernet4/4.2
  vrf member fpv-foundation_vrf
  mtu 9216
  encapsulation dot1q 2
  ip address 10.11.99.6/30
  no shutdown
interface ethernet4/8.2
  vrf member fpv-foundation_vrf
  mtu 9216
  encapsulation dot1q 2
  ip address 10.11.99.14/30
  no shutdown
configure terminal
```

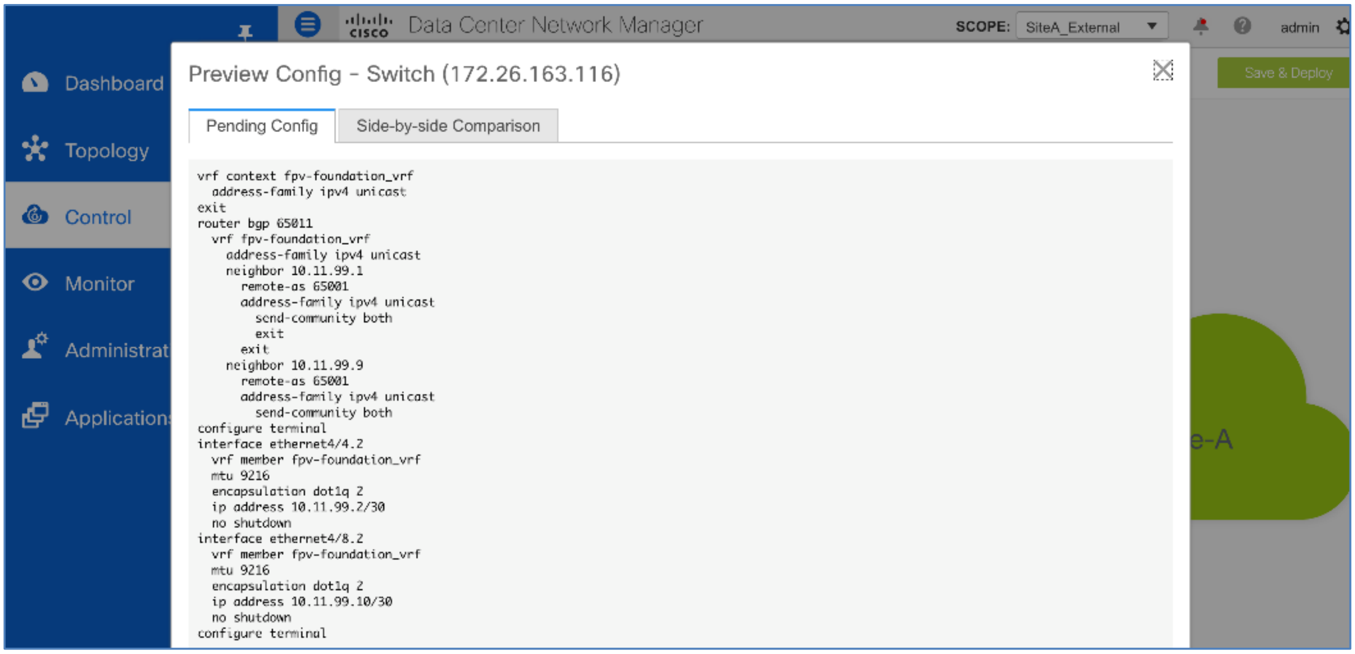

Preview Config - Switch (172.26.163.116)

```
vrf context fpv-foundation_vrf
  address-family ipv4 unicast
exit
router bgp 65011
  vrf fpv-foundation_vrf
    address-family ipv4 unicast
    neighbor 10.11.99.1
      remote-as 65001
      address-family ipv4 unicast
        send-community both
        exit
      exit
    neighbor 10.11.99.9
      remote-as 65001
      address-family ipv4 unicast
        send-community both
configure terminal
interface ethernet4/4.2
  vrf member fpv-foundation_vrf
  mtu 9216
  encapsulation dot1q 2
  ip address 10.11.99.2/30
  no shutdown
interface ethernet4/8.2
  vrf member fpv-foundation_vrf
  mtu 9216
  encapsulation dot1q 2
  ip address 10.11.99.10/30
  no shutdown
configure terminal
```

17. Click the **X** to close the **Preview** window. Click the **Deploy Config** button.

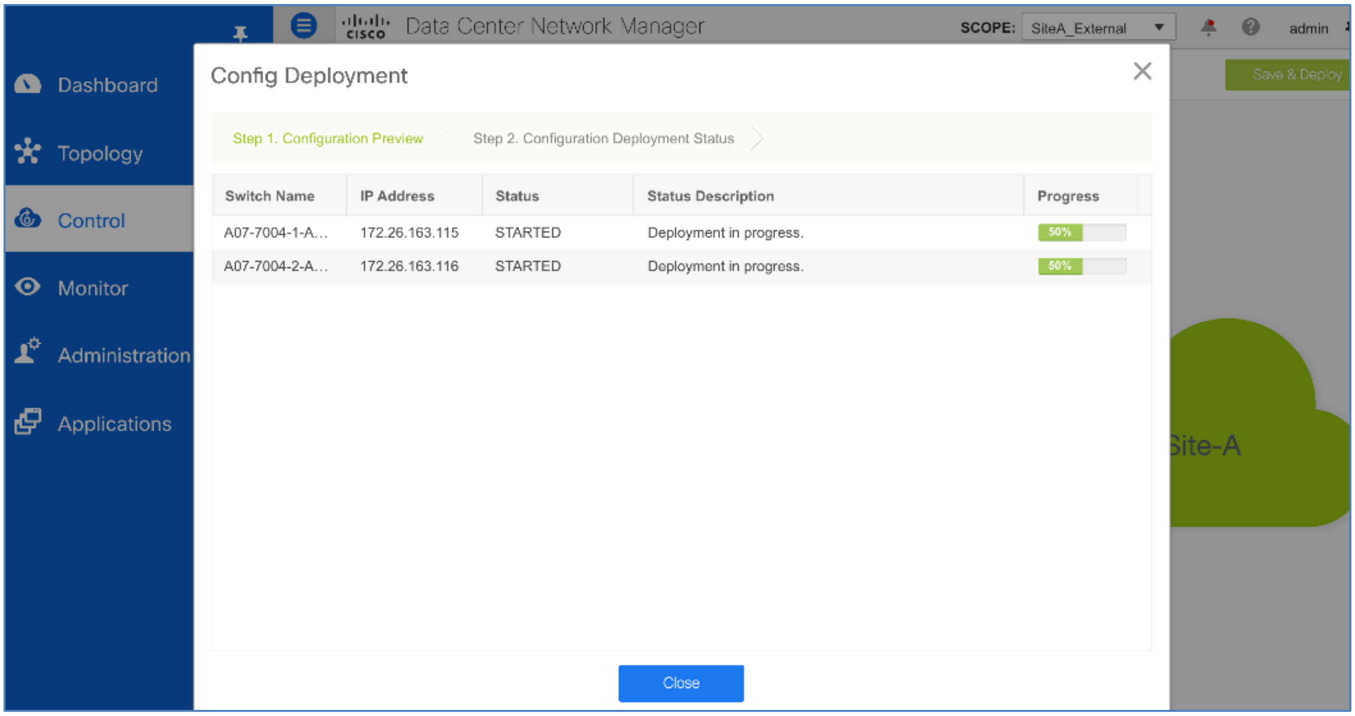

Config Deployment

Switch Name	IP Address	Status	Status Description	Progress
A07-7004-1-A...	172.26.163.115	STARTED	Deployment in progress.	50%
A07-7004-2-A...	172.26.163.116	STARTED	Deployment in progress.	50%

18. When the deployment completes and the **Status** is **COMPLETED**, click the **Close** button to close the **Config Deployment** pop-up window.

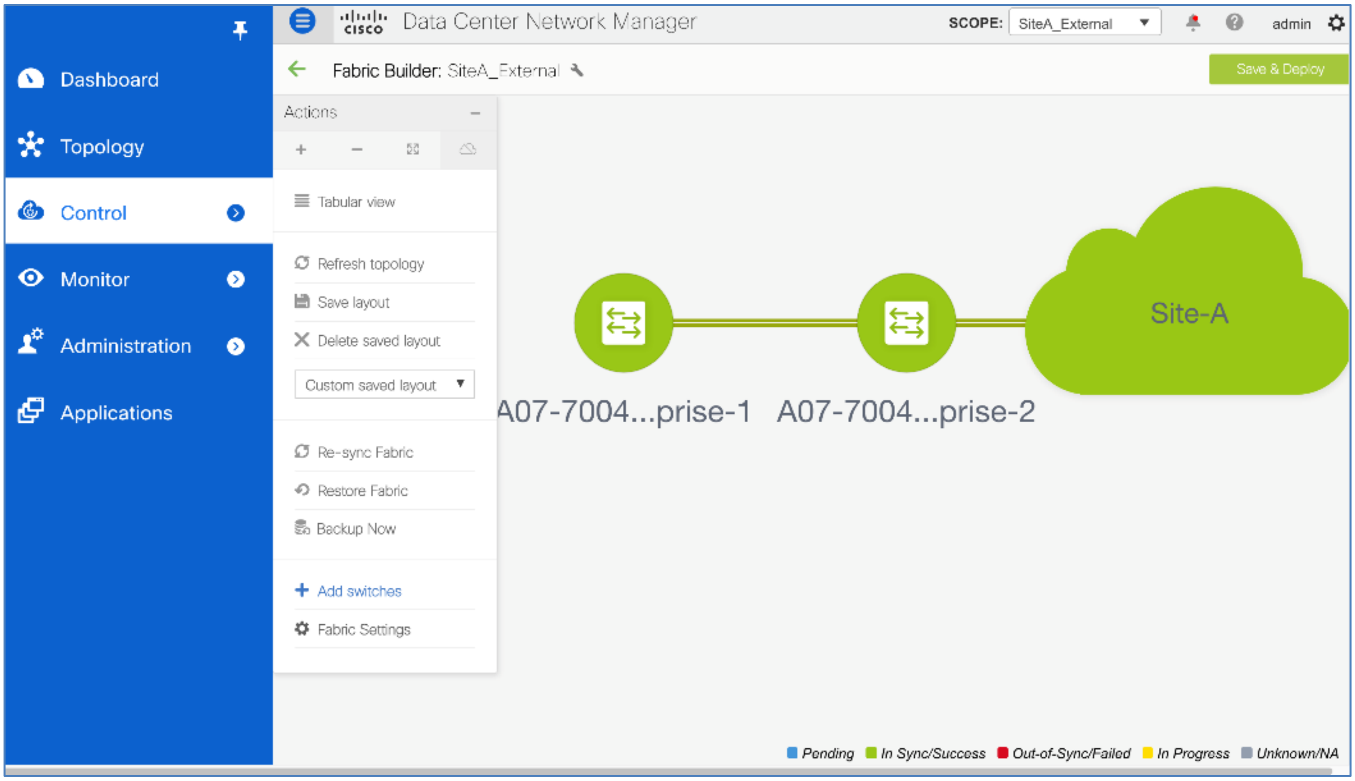

19. The external fabric switches are back to an **In Sync/Success** state.

Enable Network Connectivity for FlexPod Applications

In this design, Applications are deployed in a dedicated Tenant, separate from the FlexPod infrastructure Tenant. To enable access to FlexPod Application Tenant VMs hosted on the FlexPod infrastructure, the Application Tenant and networks must be first deployed in the VXLAN fabric. In this design, the FlexPod Application traffic is part a separate tenant/VRF (FPV-Application_VRF), dedicated to Application traffic. This tenant is used by the applications workloads hosted on the FlexPod Virtual Server Infrastructure (VSI).

Setup Information

The configuration parameters for deploying the FlexPod infrastructure networks in Site-A datacenter fabric are provided below.

Table 15. Data Center Information

Scope	Site-A

Table 16. Application Tenant/VRF

VRF Name	VRF VLAN Name	VRF Interface Description	VRF Description
FPV-Application_VRF	FPV_Application_VRF_VLAN	FPV_Application_VRF_Interface	FPV_Application_VRF

Table 17. Application Networks (FPV-Application_VRF)

VLAN Name	VLAN	VLAN Name	Forwarding	IP Subnet/Gateway*	VXLAN Network ID (VNID)	Notes
FPV-App-1_Network	1001	FPV-App-1_VLAN FPV-App-1_Interface	Layer 3	172.22.1.254/24	21001	MTU = 9216
FPV-App-2_Network	1002	FPV-App-2_VLAN FPV-App-2_Interface	Layer 3	172.22.2.254/24	21002	MTU = 9216
FPV-App-3_Network	1003	FPV-App-3_VLAN FPV-App-3_Interface	Layer 3	172.22.3.254/24	21003	MTU = 9216

* Gateway IP is specified only for L3 Forwarding and when the default gateway is in the VXLAN Fabric

Table 18. Application Storage Networks (FPV-Application_VRF)

Network Name	VLAN	VLAN Name	Forwarding	IP Subnet /Gateway*	VXLAN Network ID (VNID)	Notes
FPV-App1-NFS_Network	3051	FPV-App1-NFS_VLAN	Layer 2 Only	192.168.51.0/24	21004	App-1 NFS
FPV-App2-iSCSI-A_Network	3012	FPV-App2-iSCSI-A_VLAN	Layer 2 Only	192.168.12.0/24	21005	App-2 iSCSI
FPV-App2-iSCSI-B_Network	3022	FPV-App2-iSCSI-B_VLAN	Layer 2 Only	192.168.22.0/24	21006	App-2 iSCSI
FPV-App3-NFS_Network	3053	FPV-App3-NFS_VLAN	Layer 2 Only	192.168.53.0/24	21007	App-3 NFS
FPV-App3-iSCSI-A_Network	3013	FPV-App3-iSCSI-A_VLAN	Layer 2 Only	192.168.13.0/24	21008	App-3 iSCSI
FPV-App3-iSCSI-B_Network	3023	FPV-App3-iSCSI-B_VLAN	Layer 2 Only	192.168.23.0/24	21009	App-3 iSCSI

* Gateway IP is specified only for L3 Forwarding and when the default gateway is in the VXLAN Fabric

Deploy FlexPod Application Tenant in Cisco DCNM

To create the FlexPod Application Tenant in Cisco DCNM, use the **Setup Information** provided above to follow these steps:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > VRFs**. Click **OK** in the pop-window that complains about an **Unsupported Fabric Data Center selected**. Use **Scope:** to change the scope to **Site-A.**

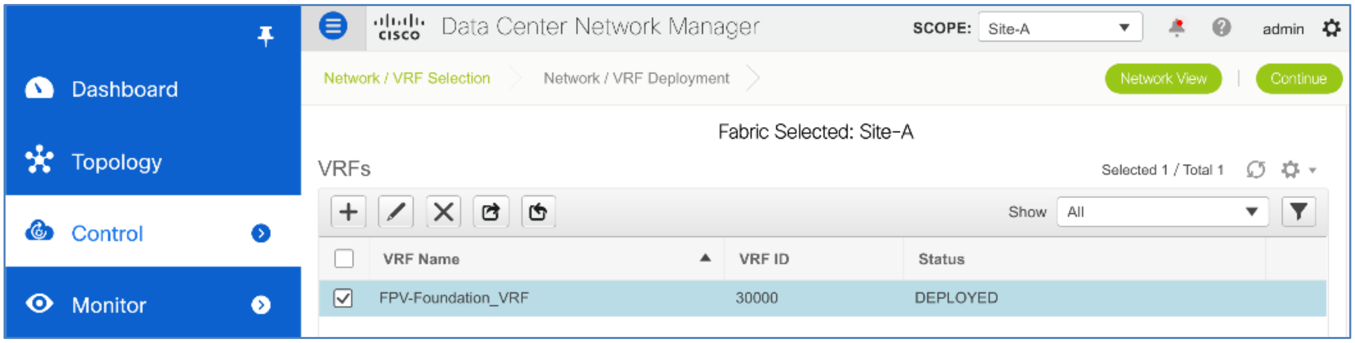

☐	VRF Name	VRF ID	Status
☑	FPV-Foundation_VRF	30000	DEPLOYED

3. Click on the **[+]** icon to deploy a new Tenant VRF for the FlexPod infrastructure traffic. Specify a **VRF VLAN Name, VRF Interface Description and VRF Description**. Leave everything else as-is.

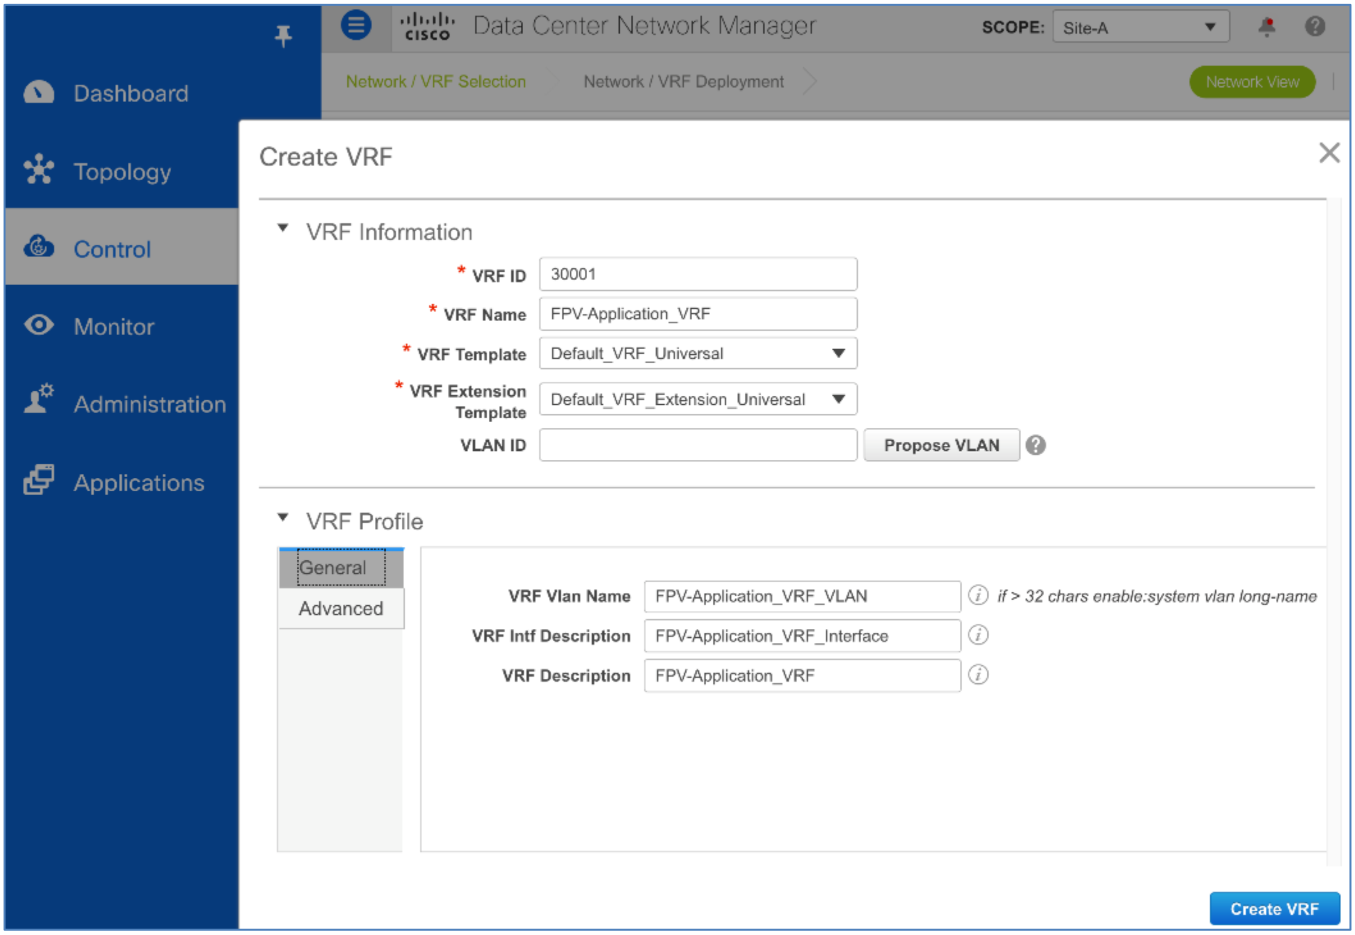

* VRF ID	30001
* VRF Name	FPV-Application_VRF
* VRF Template	Default_VRF_Universal
* VRF Extension Template	Default_VRF_Extension_Universal
VLAN ID	
VRF Vlan Name	FPV-Application_VRF_VLAN (if > 32 chars enable:system vlan long-name)
VRF Intf Description	FPV-Application_VRF_Interface
VRF Description	FPV-Application_VRF

4. Click the **Create VRF** button. A small pop-up box will appear in the bottom-right corner to confirm that the VRF was created successfully.

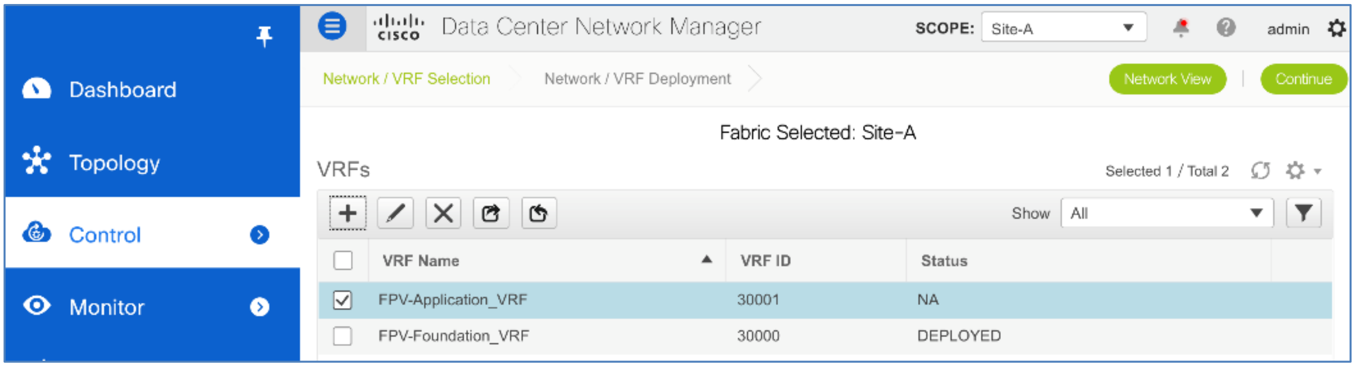

Fabric Selected: Site-A

VRFs

Selected 1 / Total 2

	VRF Name	VRF ID	Status
☑	FPV-Application_VRF	30001	NA
□	FPV-Foundation_VRF	30000	DEPLOYED

5. Click the **Continue** button.

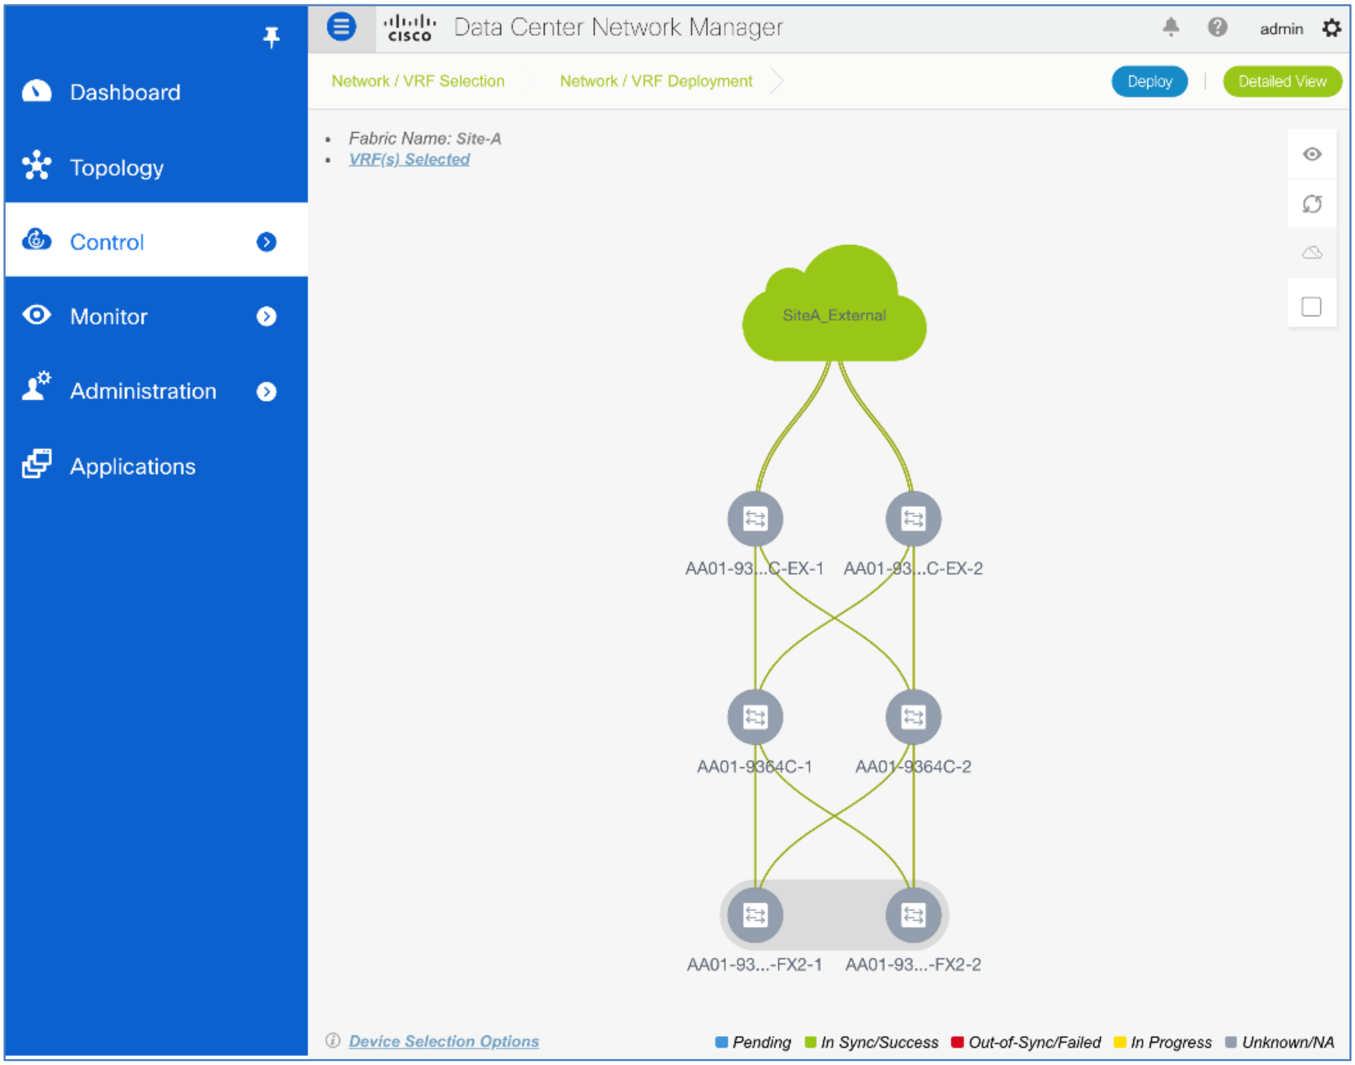

- Fabric Name: Site-A
- VRF(s) Selected

6. Click the **Detailed View** button.

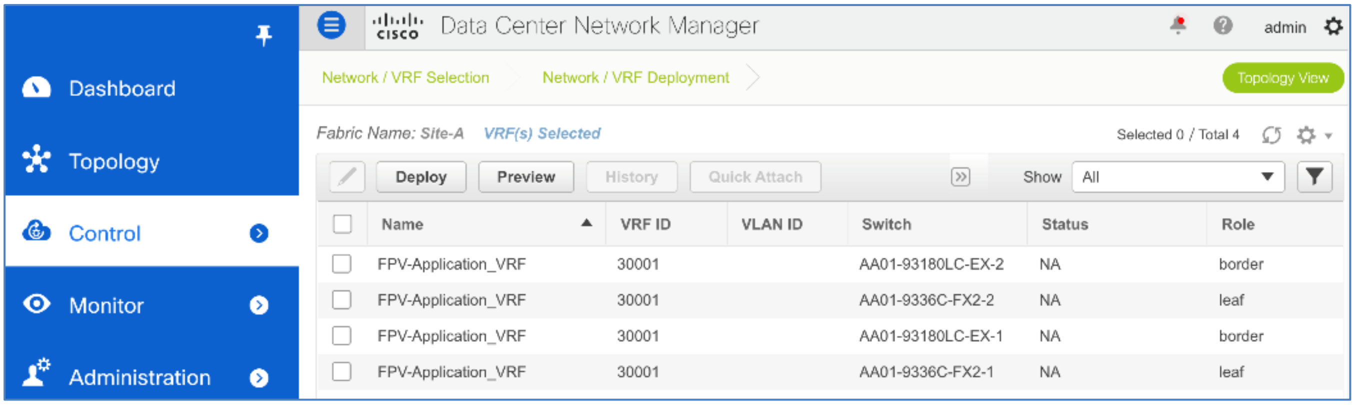

Fabric Name: Site-A *VRF(s) Selected*

Selected 0 / Total 4

□	Name	VRF ID	VLAN ID	Switch	Status	Role
□	FPV-Application_VRF	30001		AA01-93180LC-EX-2	NA	border
□	FPV-Application_VRF	30001		AA01-9336C-FX2-2	NA	leaf
□	FPV-Application_VRF	30001		AA01-93180LC-EX-1	NA	border
□	FPV-Application_VRF	30001		AA01-9336C-FX2-1	NA	leaf

7. Select the checkbox for all Leaf and Border switches in the list. Click the **Quick Attach** button.

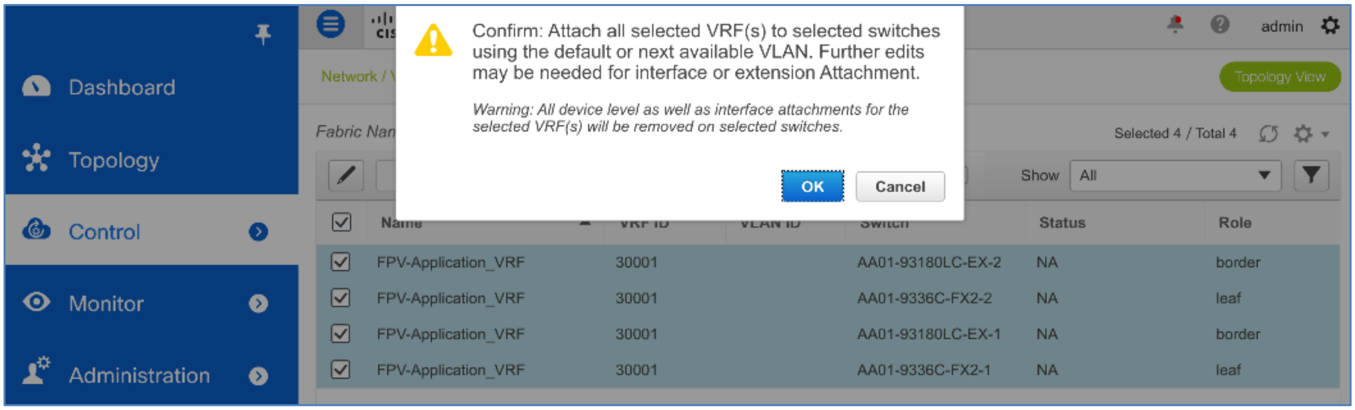

Confirm: Attach all selected VRF(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected VRF(s) will be removed on selected switches.

Selected 4 / Total 4

☑	Name	VRF ID	VLAN ID	Switch	Status	Role
☑	FPV-Application_VRF	30001		AA01-93180LC-EX-2	NA	border
☑	FPV-Application_VRF	30001		AA01-9336C-FX2-2	NA	leaf
☑	FPV-Application_VRF	30001		AA01-93180LC-EX-1	NA	border
☑	FPV-Application_VRF	30001		AA01-9336C-FX2-1	NA	leaf

8. Click **OK.**

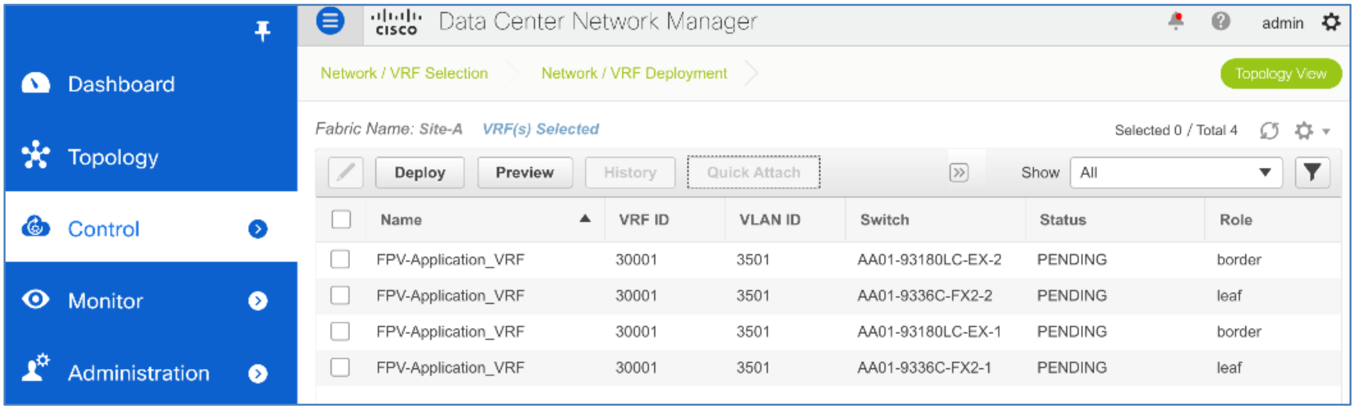

Fabric Name: Site-A *VRF(s) Selected*

Selected 0 / Total 4

□	Name	VRF ID	VLAN ID	Switch	Status	Role
□	FPV-Application_VRF	30001	3501	AA01-93180LC-EX-2	PENDING	border
□	FPV-Application_VRF	30001	3501	AA01-9336C-FX2-2	PENDING	leaf
□	FPV-Application_VRF	30001	3501	AA01-93180LC-EX-1	PENDING	border
□	FPV-Application_VRF	30001	3501	AA01-9336C-FX2-1	PENDING	leaf

9. Click the **Preview** button to preview the pending configuration on all the Leaf and Border switches.

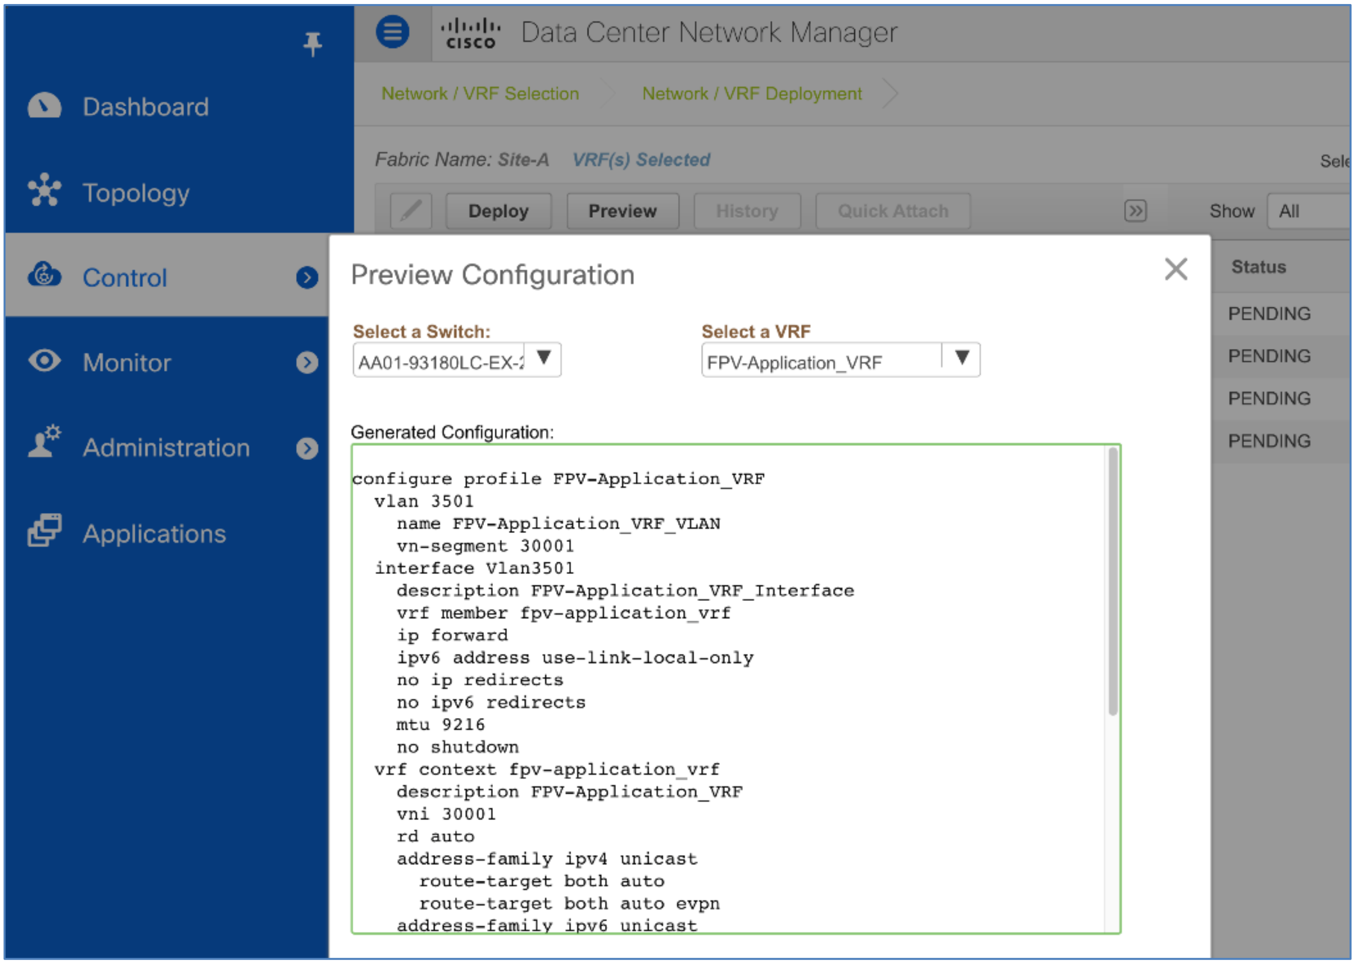

Select a Switch: AA01-93180LC-EX-2

Select a VRF: FPV-Application_VRF

Generated Configuration:

```
configure profile FPV-Application_VRF
  vlan 3501
    name FPV-Application_VRF_VLAN
    vn-segment 30001
  interface Vlan3501
    description FPV-Application_VRF_Interface
    vrf member fpv-application_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-application_vrf
    description FPV-Application_VRF
    vni 30001
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
```

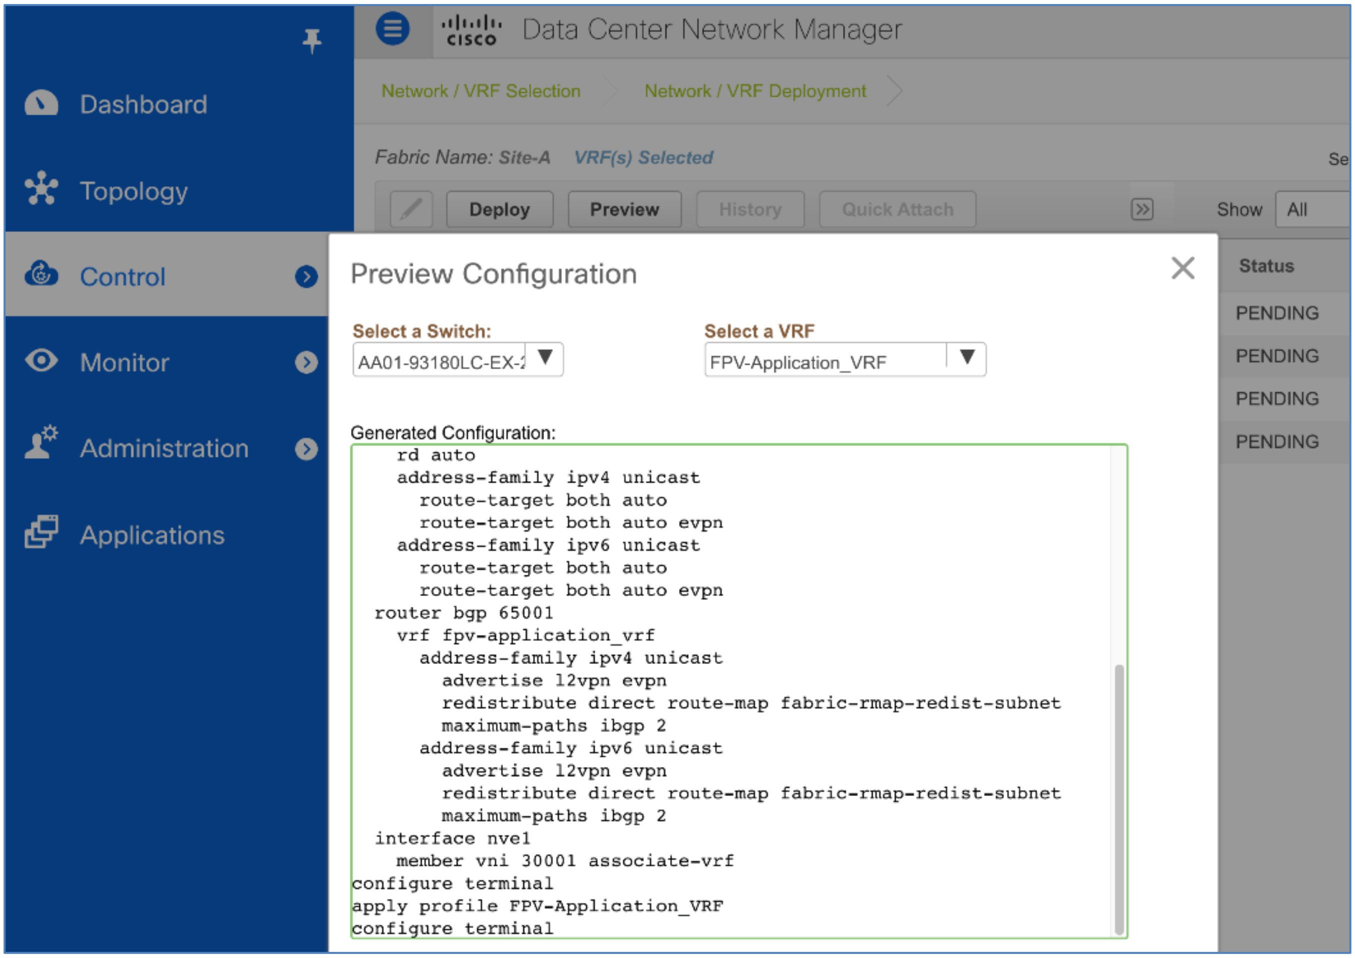

Fabric Name: Site-A VRF(s) Selected

Select a Switch: AA01-93180LC-EX-2

Select a VRF: FPV-Application_VRF

Generated Configuration:

```
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  router bgp 65001
    vrf fpv-application_vrf
      address-family ipv4 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
      address-family ipv6 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
  interface nve1
    member vni 30001 associate-vrf
configure terminal
apply profile FPV-Application_VRF
configure terminal
```

10. Click the **X** to close the **Preview Configuration** window. Click the **Deploy** button. Once the configuration is deployed, the **Status** go from **PENDING** to **IN PROGRESS** to **DEPLOYED.** Click the **Topology View** button.

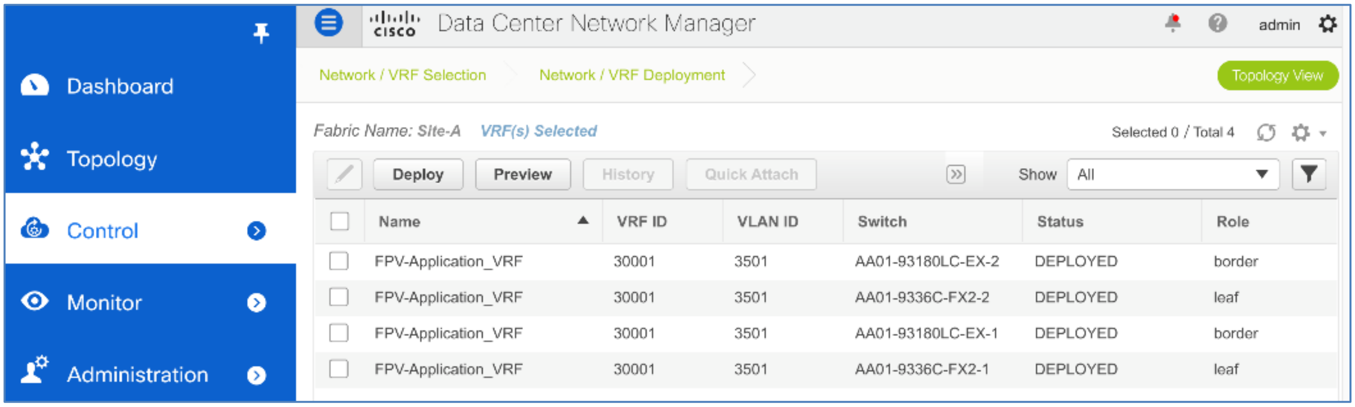

Fabric Name: Site-A VRF(s) Selected

Selected 0 / Total 4

Name	VRF ID	VLAN ID	Switch	Status	Role
FPV-Application_VRF	30001	3501	AA01-93180LC-EX-2	DEPLOYED	border
FPV-Application_VRF	30001	3501	AA01-9336C-FX2-2	DEPLOYED	leaf
FPV-Application_VRF	30001	3501	AA01-93180LC-EX-1	DEPLOYED	border
FPV-Application_VRF	30001	3501	AA01-9336C-FX2-1	DEPLOYED	leaf

11. In the **Topology View** to see the where the selected VRF is deployed in the Site-A topology.

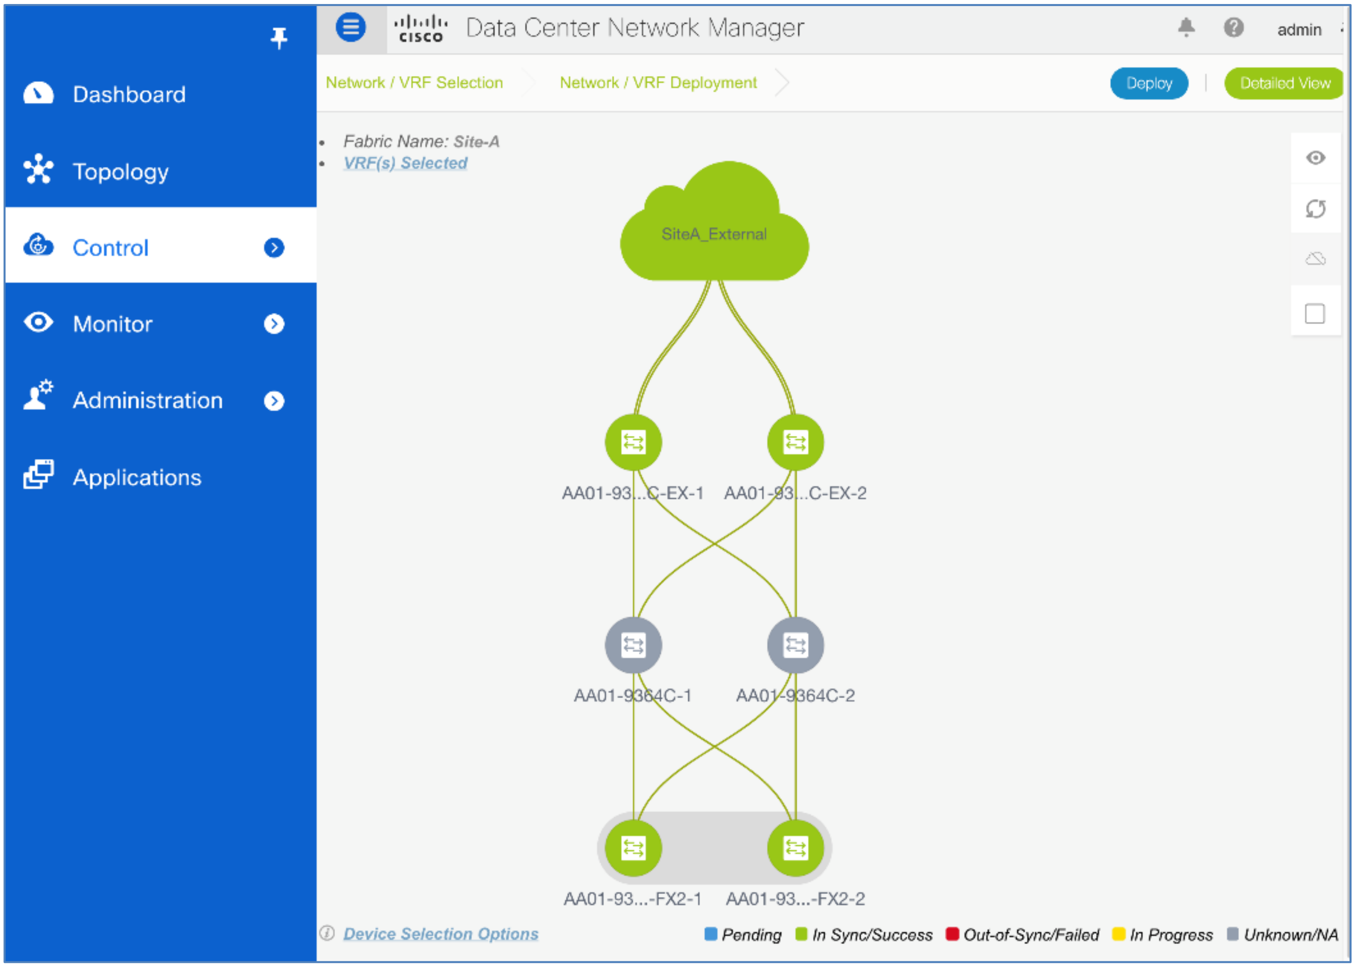

- Fabric Name: Site-A
- VRF(s) Selected

Deploy FlexPod Application Networks

The Applications networks are deployed in this design in Layer 3 mode where the traffic is Layer 3 forwarded by the fabric and the gateway is a distributed anycast gateway in the VXLAN fabric.

To create the FlexPod Application Tenant networks in Cisco DCNM, use the **Setup Information** provided above and follow these steps:

1. Use a browser to navigate to Cisco DCNM's GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**.

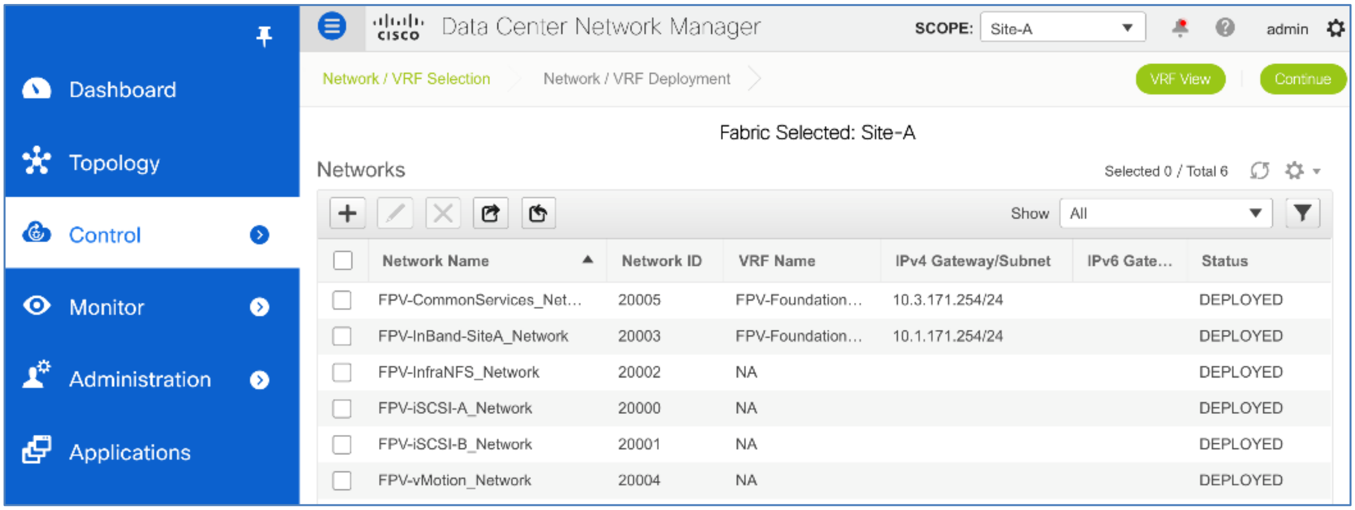

Fabric Selected: Site-A

Networks

Selected 0 / Total 6

Show All

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gate...	Status
□	FPV-CommonServices_Net...	20005	FPV-Foundation...	10.3.171.254/24		DEPLOYED
□	FPV-InBand-SiteA_Network	20003	FPV-Foundation...	10.1.171.254/24		DEPLOYED
□	FPV-InfraNFS_Network	20002	NA			DEPLOYED
□	FPV-iSCSI-A_Network	20000	NA			DEPLOYED
□	FPV-iSCSI-B_Network	20001	NA			DEPLOYED
□	FPV-vMotion_Network	20004	NA			DEPLOYED

3. Click on the **[+]** icon to deploy a new Tenant network for the FlexPod infrastructure traffic. Specify a **Network Name, VRF Name.** In this deployment, we are specifying the **VLAN ID** we specifically want to use but you can optionally let DCNM pick up one from the defined pool in **Fabric Settings.** Specify a **VLAN ID**. In the **Network Profile > General** section of the window, specify a **IPv4 Gateway/Network, VLAN Name, Interface Description** and **MTU**. Leave everything else as-is.

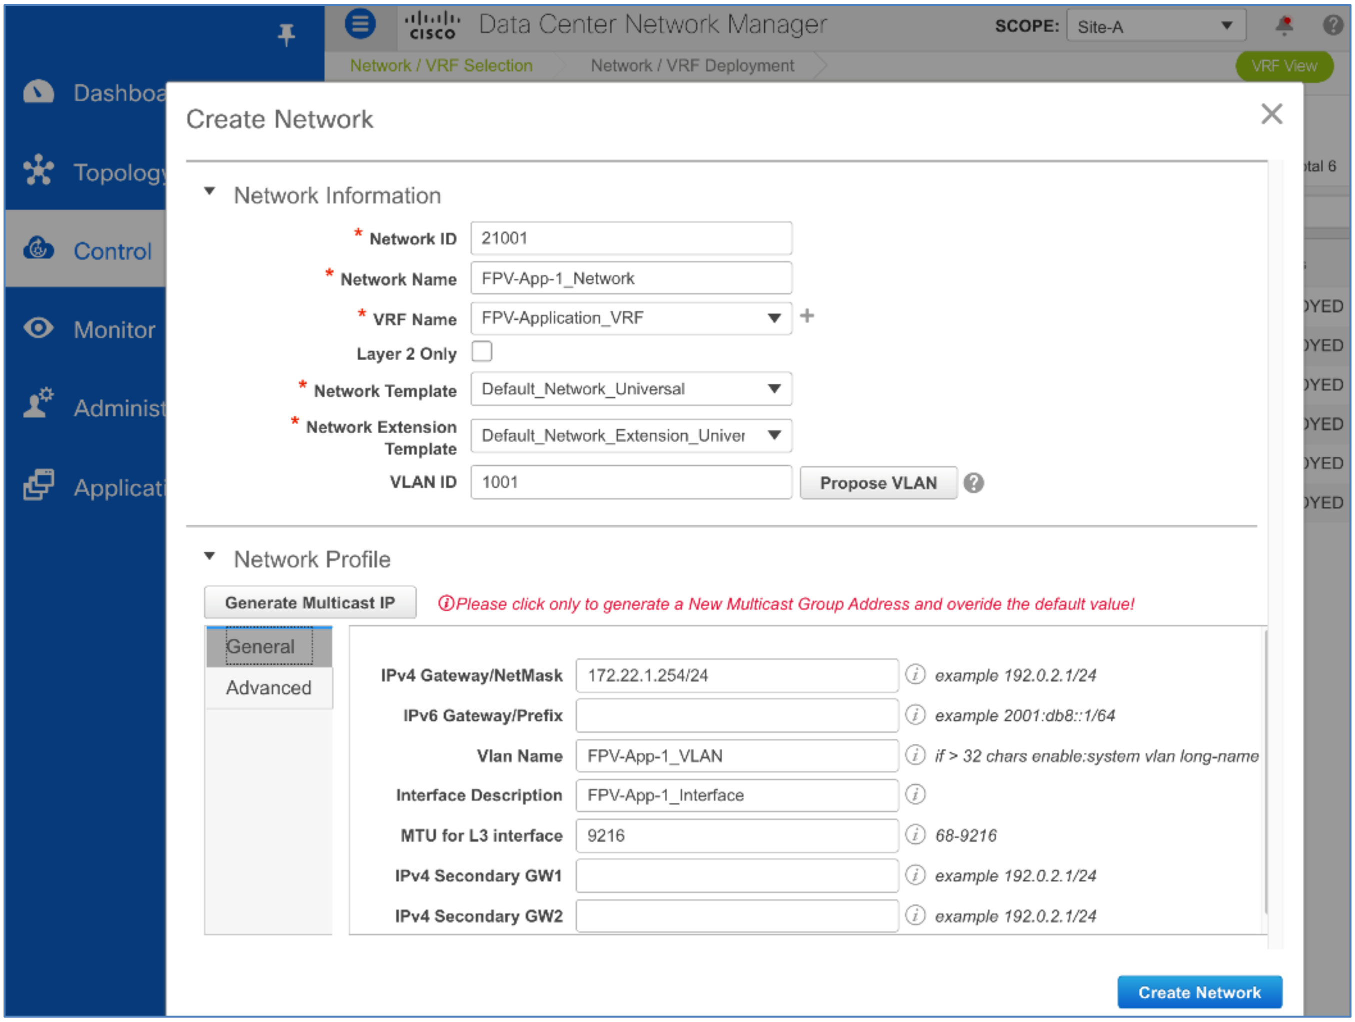

Create Network

Network Information

* Network ID	21001
* Network Name	FPV-App-1_Network
* VRF Name	FPV-Application_VRF
Layer 2 Only	□
* Network Template	Default_Network_Universal
* Network Extension Template	Default_Network_Extension_Univer
VLAN ID	1001

Network Profile

Please click only to generate a New Multicast Group Address and overide the default value!

IPv4 Gateway/NetMask	172.22.1.254/24	example 192.0.2.1/24
IPv6 Gateway/Prefix		example 2001:db8::1/64
Vlan Name	FPV-App-1_VLAN	if > 32 chars enable:system vlan long-name
Interface Description	FPV-App-1_Interface	
MTU for L3 interface	9216	68-9216
IPv4 Secondary GW1		example 192.0.2.1/24
IPv4 Secondary GW2		example 192.0.2.1/24

4. In the **Network Profile > Advanced** section of the window, enable **ARP Suppression.** Leave everything as-is.

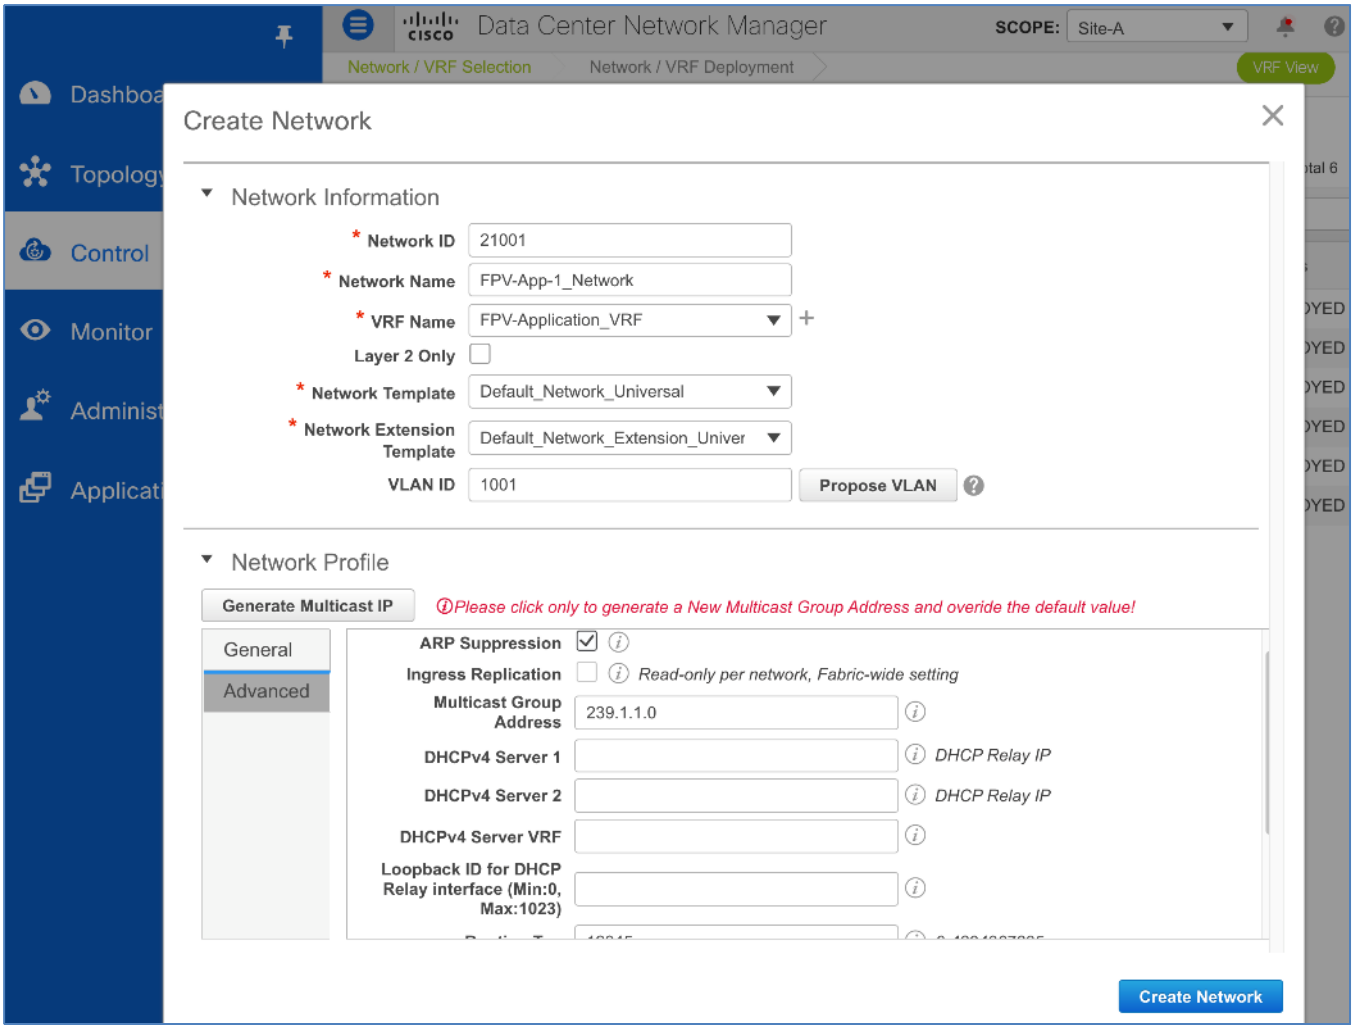

Create Network

Network Information

* Network ID: 21001

* Network Name: FPV-App-1_Network

* VRF Name: FPV-Application_VRF

Layer 2 Only: □

* Network Template: Default_Network_Universal

* Network Extension Template: Default_Network_Extension_Univer

VLAN ID: 1001 Propose VLAN

Network Profile

Generate Multicast IP

Please click only to generate a New Multicast Group Address and overide the default value!

General | Advanced

ARP Suppression: ☑

Ingress Replication: □ *Read-only per network, Fabric-wide setting*

Multicast Group Address: 239.1.1.0

DHCPv4 Server 1: *DHCP Relay IP*

DHCPv4 Server 2: *DHCP Relay IP*

DHCPv4 Server VRF:

Loopback ID for DHCP Relay interface (Min:0, Max:1023):

Create Network

5. Click the **Create Network** button.

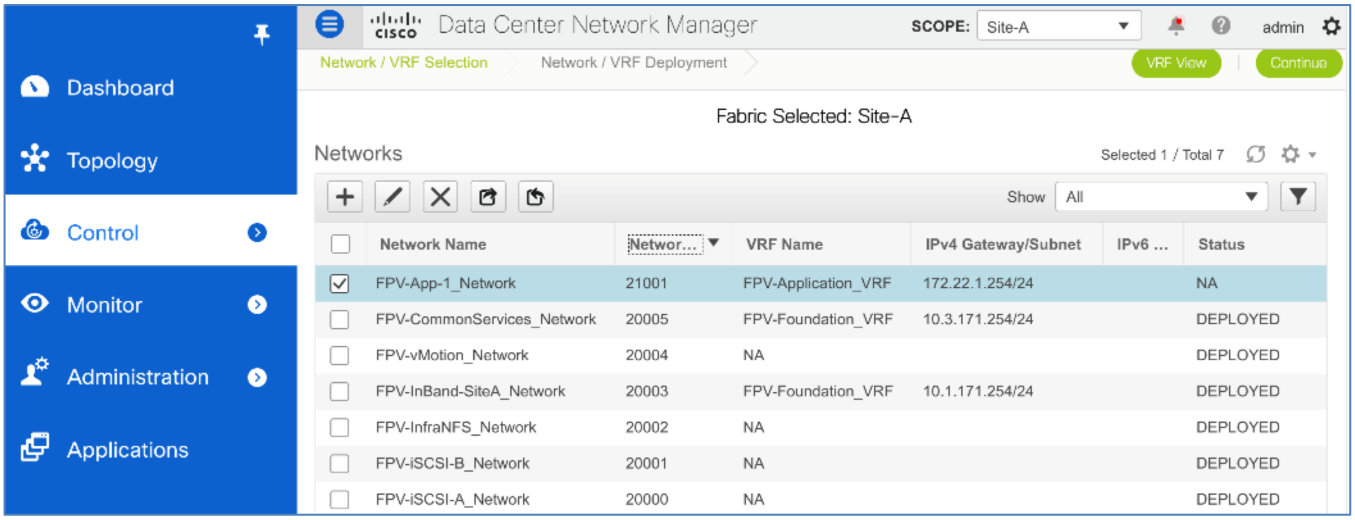

Fabric Selected: Site-A

Networks

Selected 1 / Total 7

	Network Name	Networ...	VRF Name	IPv4 Gateway/Subnet	IPv6 ...	Status
☑	FPV-App-1_Network	21001	FPV-Application_VRF	172.22.1.254/24		NA
□	FPV-CommonServices_Network	20005	FPV-Foundation_VRF	10.3.171.254/24		DEPLOYED
□	FPV-vMotion_Network	20004	NA			DEPLOYED
□	FPV-InBand-SiteA_Network	20003	FPV-Foundation_VRF	10.1.171.254/24		DEPLOYED
□	FPV-InfraNFS_Network	20002	NA			DEPLOYED
□	FPV-iSCSI-B_Network	20001	NA			DEPLOYED
□	FPV-iSCSI-A_Network	20000	NA			DEPLOYED

6. Click the **Continue** button.

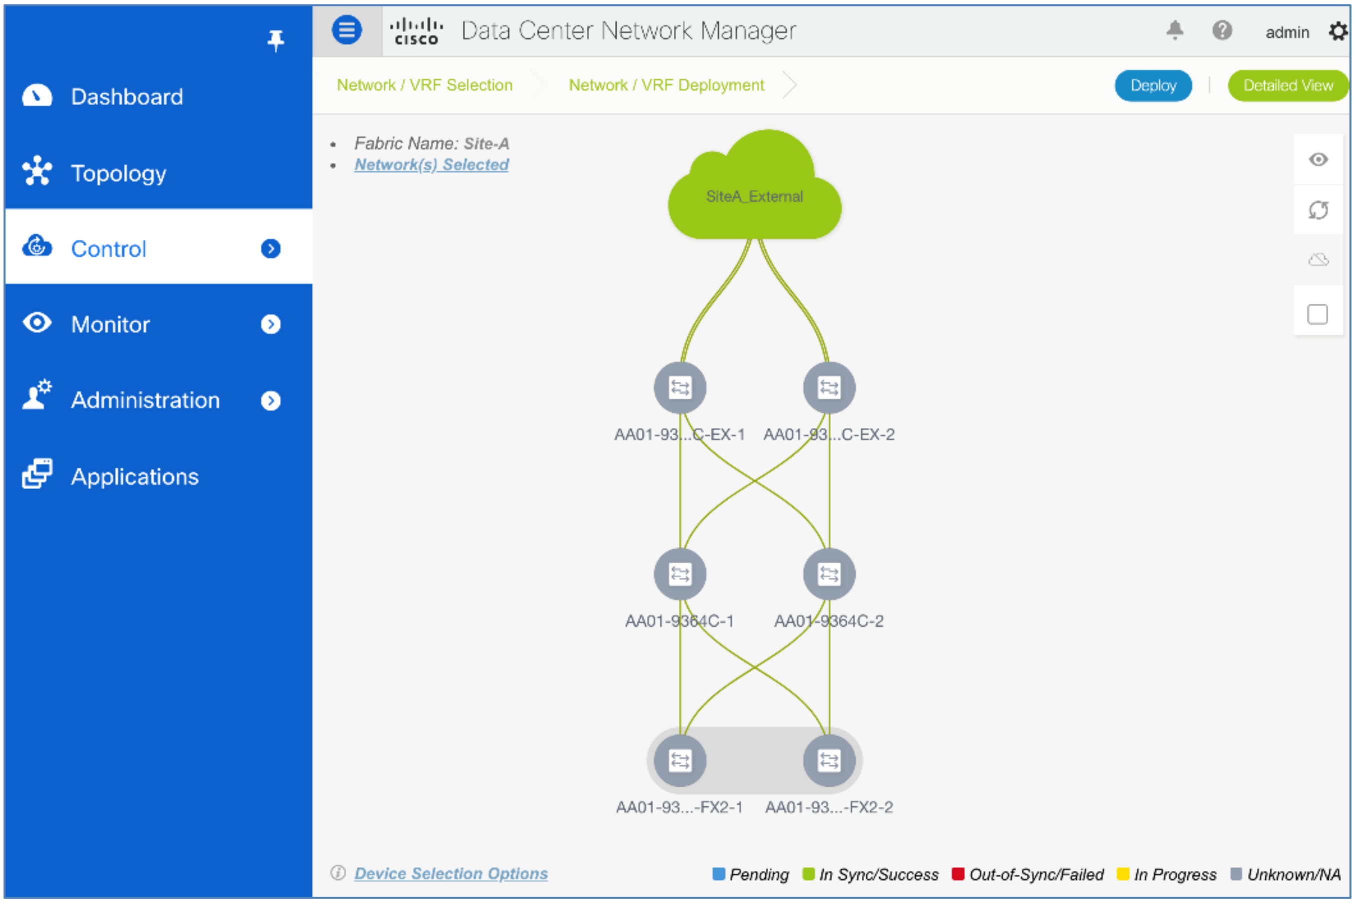

7. Click the **Detailed View** button. Select the Leaf switches where these networks need to be deployed. Click the **Quick Attach** button. You can decide which switches to deploy this network on.

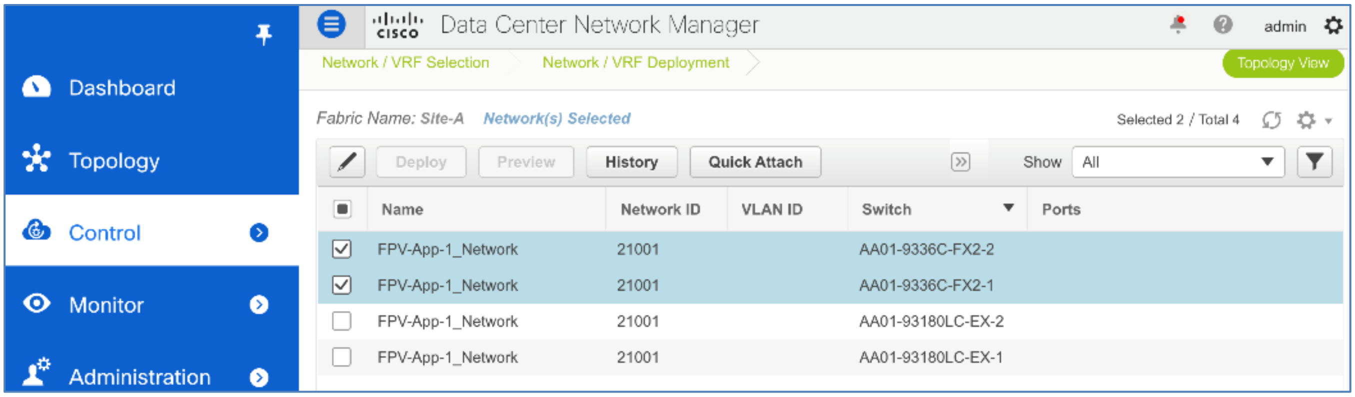

■	Name	Network ID	VLAN ID	Switch	Ports
☑	FPV-App-1_Network	21001		AA01-9336C-FX2-2	
☑	FPV-App-1_Network	21001		AA01-9336C-FX2-1	
□	FPV-App-1_Network	21001		AA01-93180LC-EX-2	
□	FPV-App-1_Network	21001		AA01-93180LC-EX-1	

8. Click **OK.**

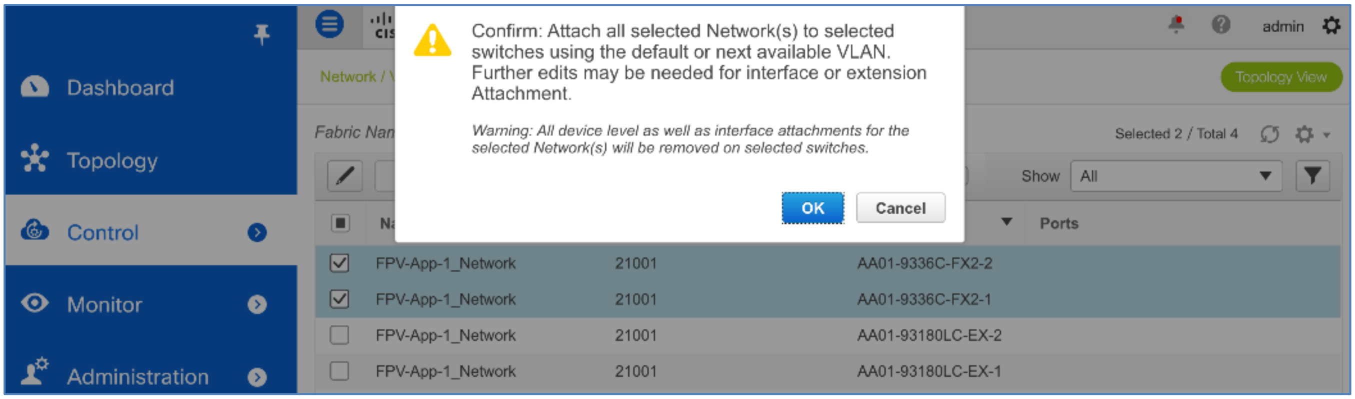

Confirm: Attach all selected Network(s) to selected switches using the default or next available VLAN. Further edits may be needed for interface or extension Attachment.

Warning: All device level as well as interface attachments for the selected Network(s) will be removed on selected switches.

	Na		Ports
☑	FPV-App-1_Network	21001	AA01-9336C-FX2-2
☑	FPV-App-1_Network	21001	AA01-9336C-FX2-1
□	FPV-App-1_Network	21001	AA01-93180LC-EX-2
□	FPV-App-1_Network	21001	AA01-93180LC-EX-1

9. Click the **Preview** button to view pending changes.

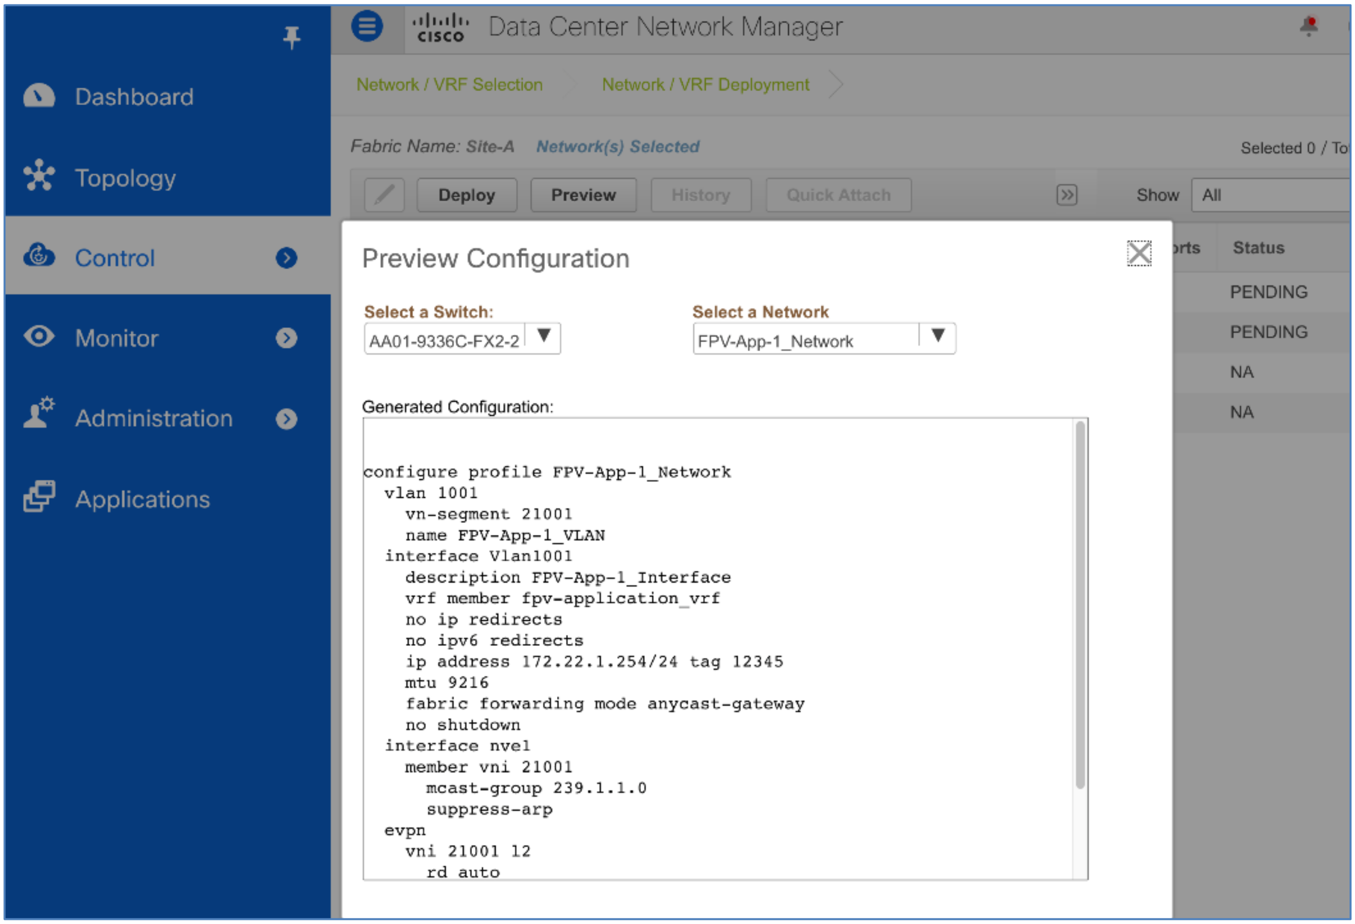

Fabric Name: Site-A *Network(s) Selected*

Preview Configuration

Select a Switch: AA01-9336C-FX2-2

Select a Network FPV-App-1_Network

Generated Configuration:

```
configure profile FPV-App-1_Network
  vlan 1001
    vn-segment 21001
    name FPV-App-1_VLAN
  interface Vlan1001
    description FPV-App-1_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.1.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21001
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21001 l2
      rd auto
```

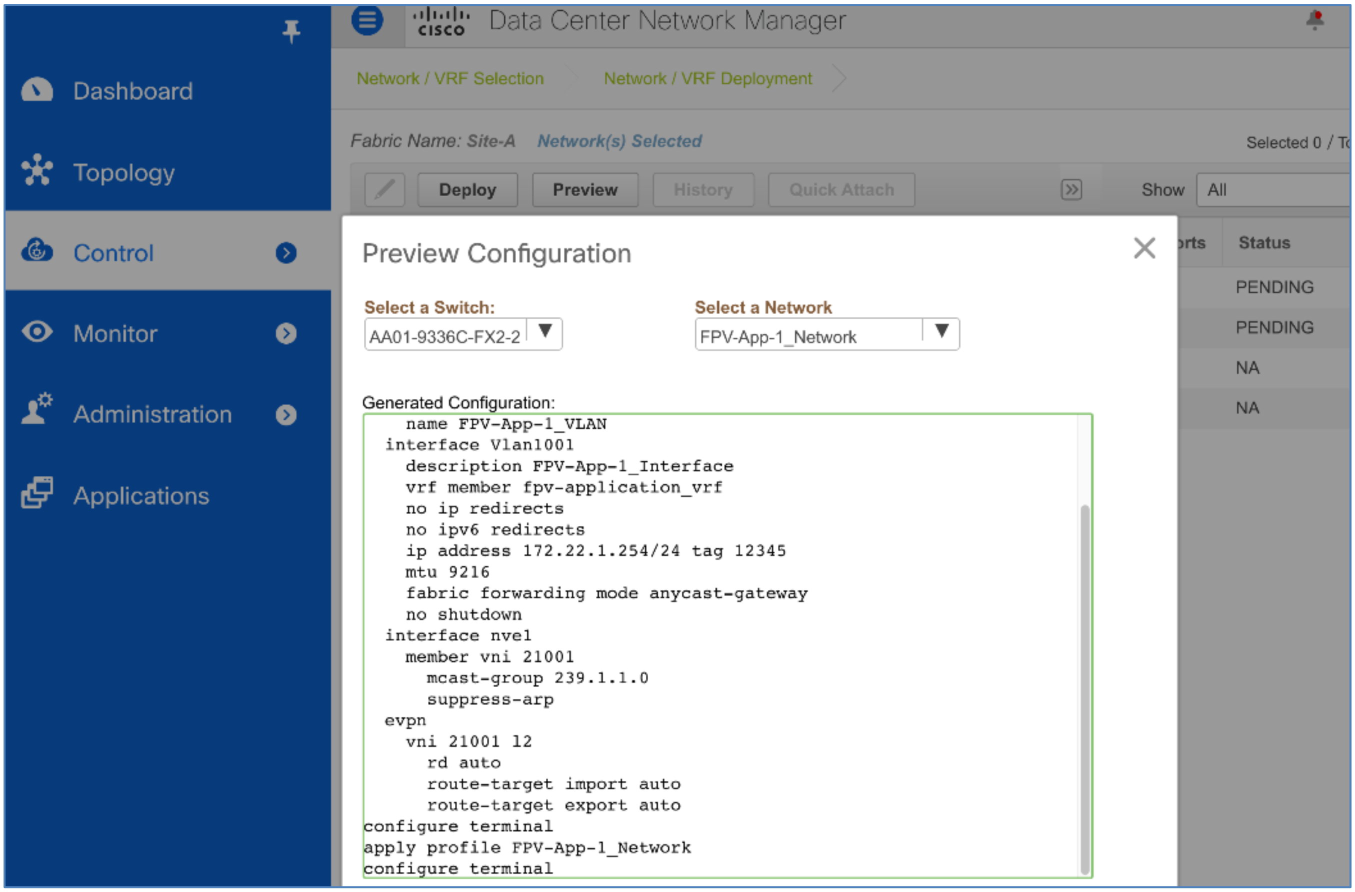

Select a Switch: AA01-9336C-FX2-2

Select a Network: FPV-App-1_Network

Generated Configuration:

```
    name FPV-App-1_VLAN
  interface Vlan1001
    description FPV-App-1_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.1.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21001
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21001 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
apply profile FPV-App-1_Network
configure terminal
```

10. Click **X** to close the **Preview** window. Click the **Deploy** button.

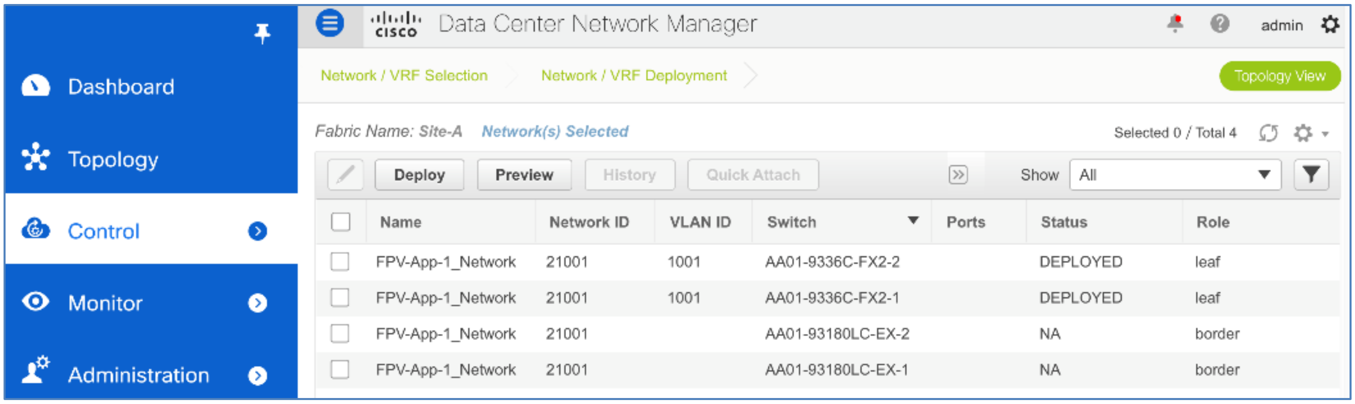

Fabric Name: Site-A Network(s) Selected Selected 0 / Total 4

□	Name	Network ID	VLAN ID	Switch	Ports	Status	Role
□	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-2		DEPLOYED	leaf
□	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1		DEPLOYED	leaf
□	FPV-App-1_Network	21001		AA01-93180LC-EX-2		NA	border
□	FPV-App-1_Network	21001		AA01-93180LC-EX-1		NA	border

11. Click the **Topology View** button to view where the selected network is deployed in the fabric topology.

12. Repeat steps 1-11 to deploy remaining Application networks.

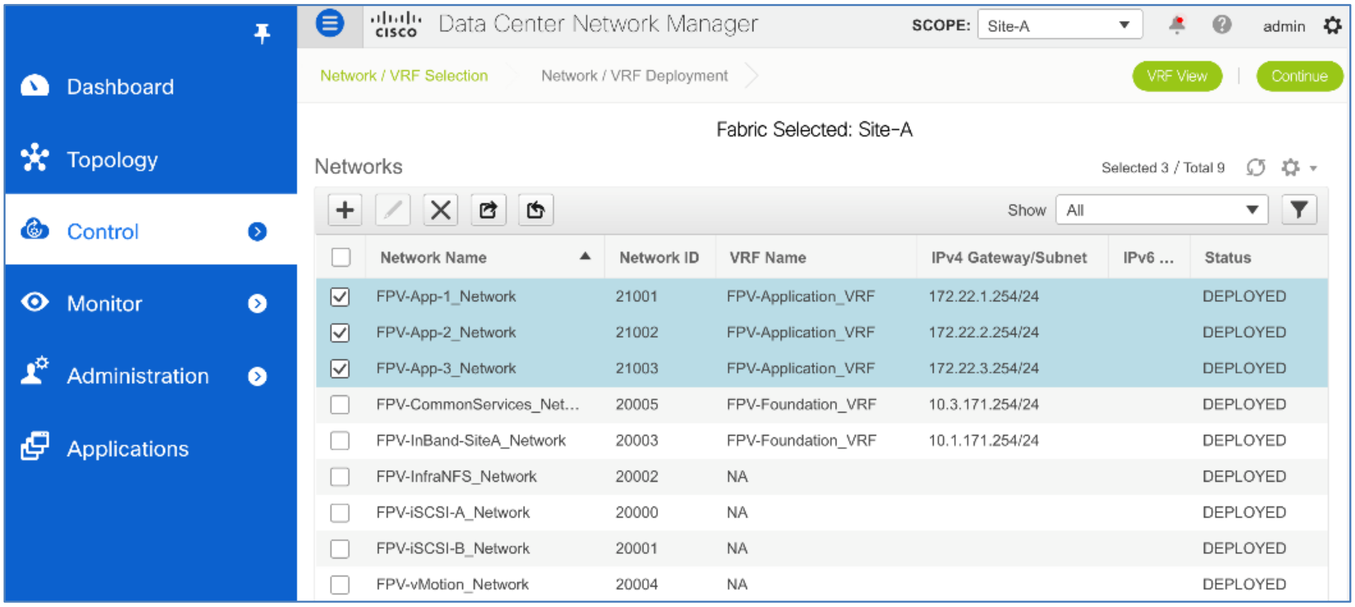

SCOPE: Site-A

Fabric Selected: Site-A

Networks — Selected 3 / Total 9

Show: All

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 ...	Status
☑	FPV-App-1_Network	21001	FPV-Application_VRF	172.22.1.254/24		DEPLOYED
☑	FPV-App-2_Network	21002	FPV-Application_VRF	172.22.2.254/24		DEPLOYED
☑	FPV-App-3_Network	21003	FPV-Application_VRF	172.22.3.254/24		DEPLOYED
□	FPV-CommonServices_Net...	20005	FPV-Foundation_VRF	10.3.171.254/24		DEPLOYED
□	FPV-InBand-SiteA_Network	20003	FPV-Foundation_VRF	10.1.171.254/24		DEPLOYED
□	FPV-InfraNFS_Network	20002	NA			DEPLOYED
□	FPV-iSCSI-A_Network	20000	NA			DEPLOYED
□	FPV-iSCSI-B_Network	20001	NA			DEPLOYED
□	FPV-vMotion_Network	20004	NA			DEPLOYED

Enable Access-Layer Connectivity to FlexPod Application Networks

To enable FlexPod Application networks on access-layer connections to Cisco UCS domain and NetApp storage cluster, complete the steps outlined in the upcoming sections.

Enable Application Networks on Access-Layer Connections to Cisco UCS Domain

To enable FlexPod Applications networks on the access-layer connections to Cisco UCS domain, follow these steps:

1. Use a browser to navigate to Cisco DCNM’s GUI. Log in using an administrator account.

2. From the left navigation bar, select **Control > Fabrics > Networks**. Click **OK** to exit any pop-ups that come up. Select the correct scope in the drop-down list next to **Scope:** in the top-right corner of the window. Select the networks that need to be enabled on the access-layer connection to the Cisco UCS Domain.

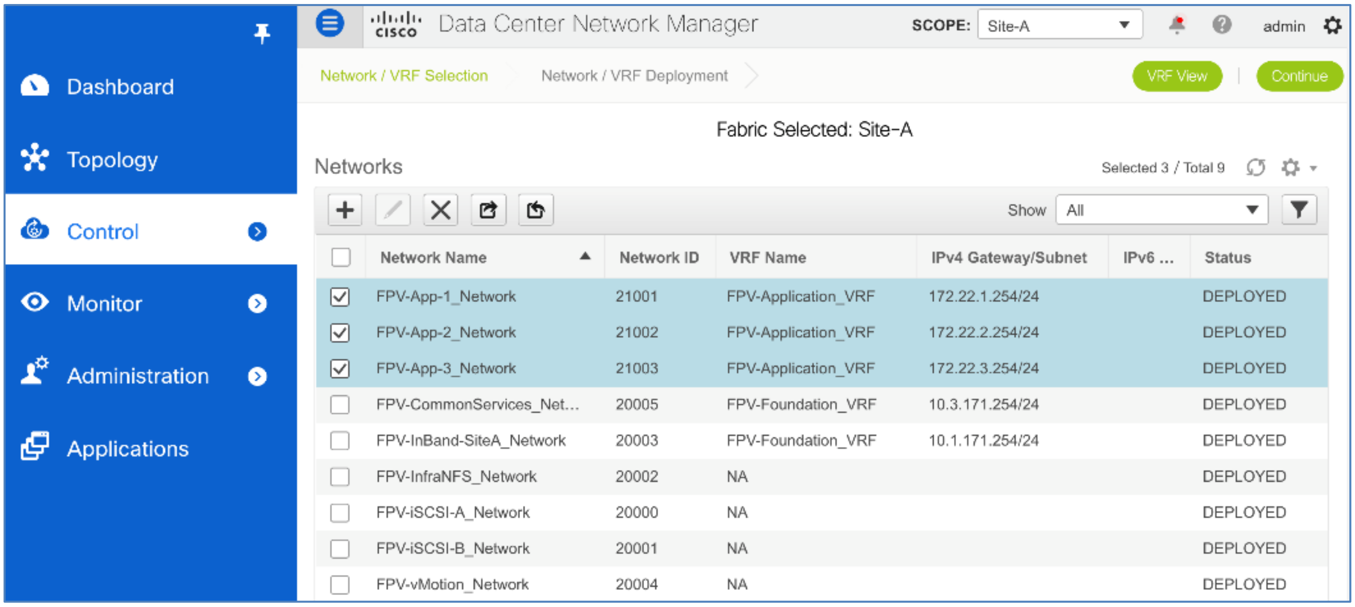

Fabric Selected: Site-A

Networks

Selected 3 / Total 9

□	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 ...	Status
☑	FPV-App-1_Network	21001	FPV-Application_VRF	172.22.1.254/24		DEPLOYED
☑	FPV-App-2_Network	21002	FPV-Application_VRF	172.22.2.254/24		DEPLOYED
☑	FPV-App-3_Network	21003	FPV-Application_VRF	172.22.3.254/24		DEPLOYED
□	FPV-CommonServices_Net...	20005	FPV-Foundation_VRF	10.3.171.254/24		DEPLOYED
□	FPV-InBand-SiteA_Network	20003	FPV-Foundation_VRF	10.1.171.254/24		DEPLOYED
□	FPV-InfraNFS_Network	20002	NA			DEPLOYED
□	FPV-iSCSI-A_Network	20000	NA			DEPLOYED
□	FPV-iSCSI-B_Network	20001	NA			DEPLOYED
□	FPV-vMotion_Network	20004	NA			DEPLOYED

3. Click the **Continue** button. Click the **Detailed View** button.

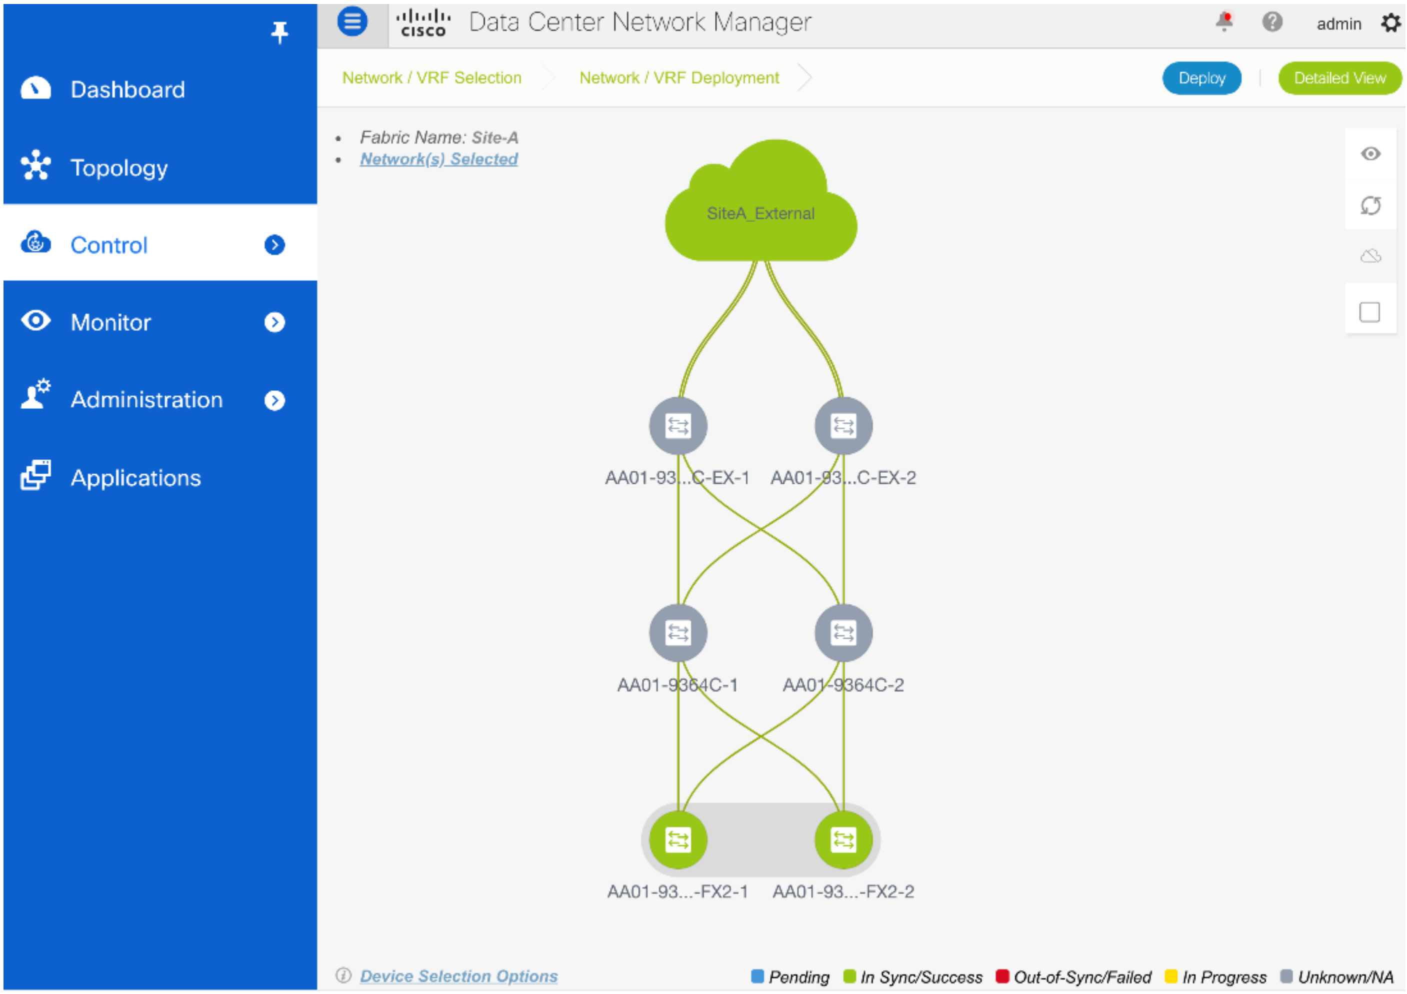

4. Select **Quick Filter** from the drop-down list next to **Show** from the top menu. This will expose the filter box above every column. Filter based on the access-layer leaf switch that connects to the Cisco UCS Domain. Select the first switch in the vPC pair to the Cisco UCS domain. Select all networks that need to be deployed. Click on the pencil icon to edit the previously deployed networks on the first leaf switch in the vPC pair.

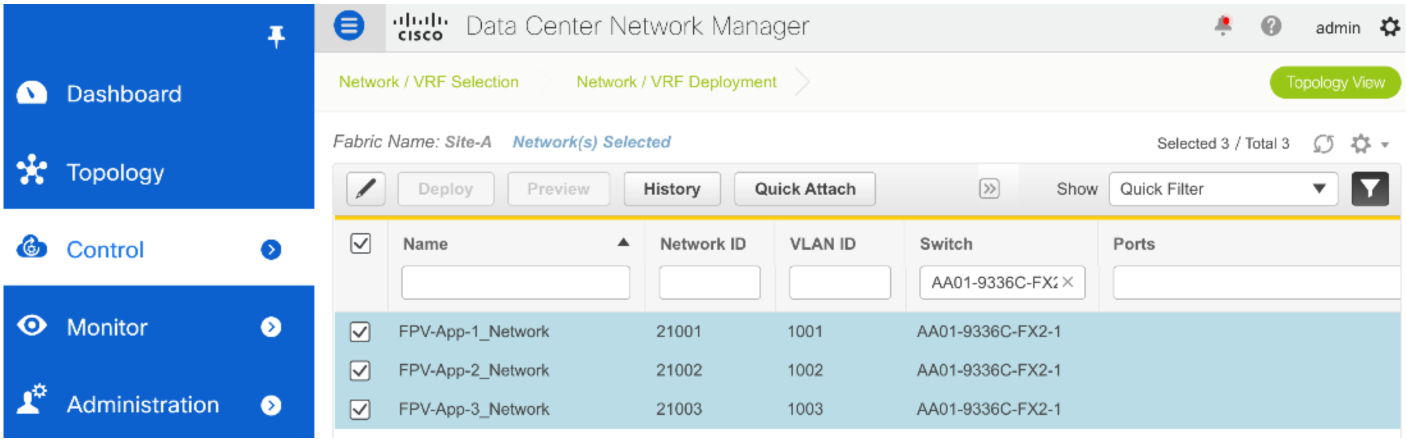

Fabric Name: Site-A   Network(s) Selected

Selected 3 / Total 3

Show: Quick Filter

☑	Name	Network ID	VLAN ID	Switch	Ports
				AA01-9336C-FX2	
☑	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	
☑	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	
☑	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	

5. You should see a network tab for each network and **both** leaf switches listed though only one switch was selected in the previous. This is because the switches are part of a vPC pair – a configuration that gets applied to one will get applied to both. Note the box … in the **Interfaces** column.

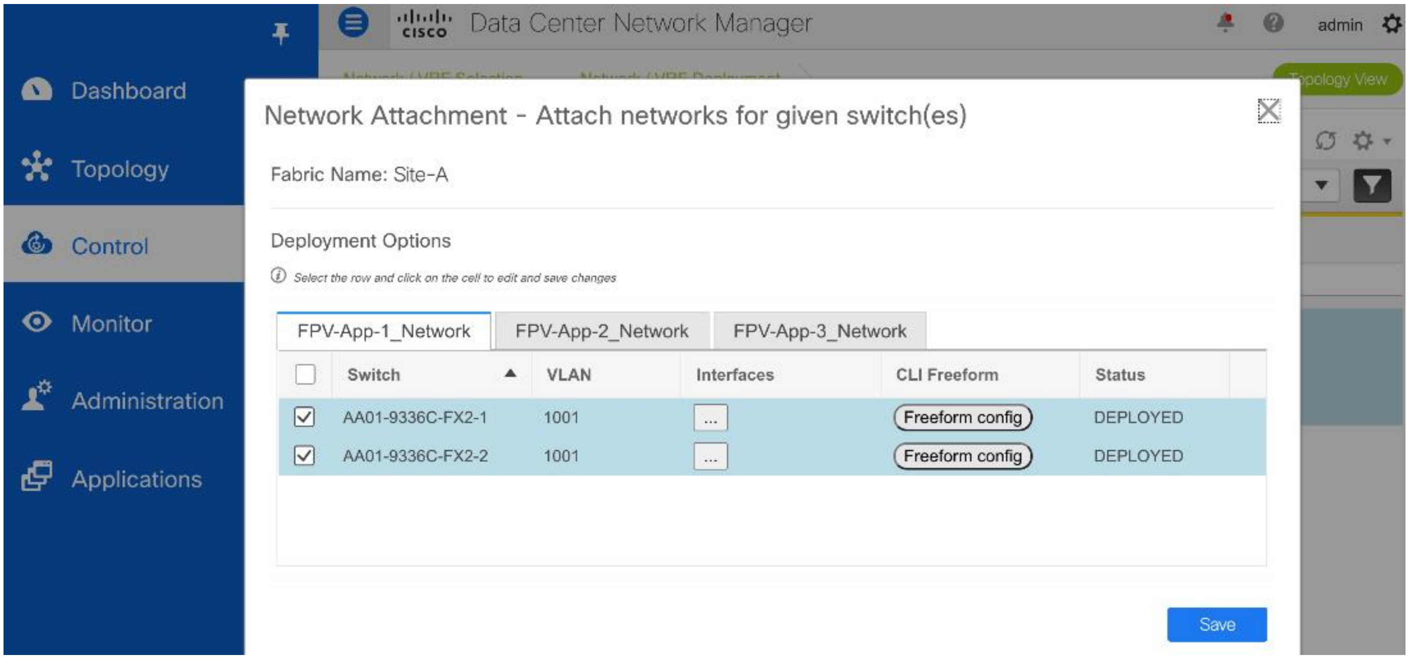

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

FPV-App-1_Network | FPV-App-2_Network | FPV-App-3_Network

□	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	1001	…	Freeform config	DEPLOYED
☑	AA01-9336C-FX2-2	1001	…	Freeform config	DEPLOYED

6. In the first network tab, click the box … in the **Interfaces** column next to the first switch listed. In the **Interfaces** pop-up window, select both port-channels that go to both Cisco UCS Fabric Interconnects in the Cisco UCS Domain.

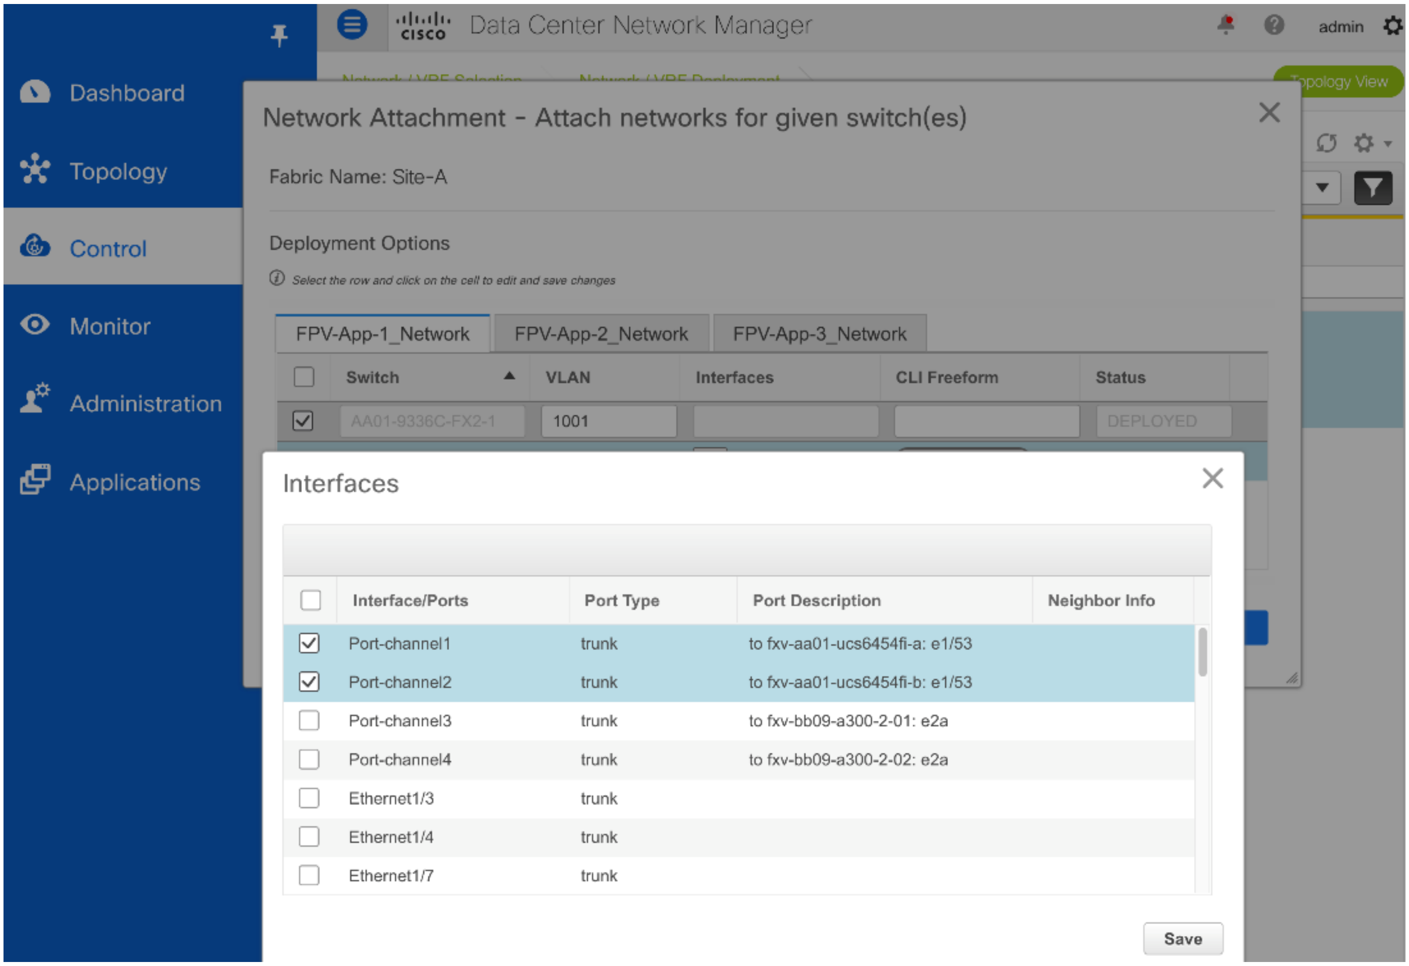

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

ⓘ *Select the row and click on the cell to edit and save changes*

FPV-App-1_Network | FPV-App-2_Network | FPV-App-3_Network

□	Switch ▲	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	1001			DEPLOYED

Interfaces

□	Interface/Ports	Port Type	Port Description	Neighbor Info
☑	Port-channel1	trunk	to fxv-aa01-ucs6454fi-a: e1/53	
☑	Port-channel2	trunk	to fxv-aa01-ucs6454fi-b: e1/53	
□	Port-channel3	trunk	to fxv-bb09-a300-2-01: e2a	
□	Port-channel4	trunk	to fxv-bb09-a300-2-02: e2a	
□	Ethernet1/3	trunk		
□	Ethernet1/4	trunk		
□	Ethernet1/7	trunk		

Save

7. Click **Save.** Note that the interfaces are now populated for the first switch. Click **Save** again. The first network for the first switch now lists the **Ports** and the **Status** is now **PENDING.**

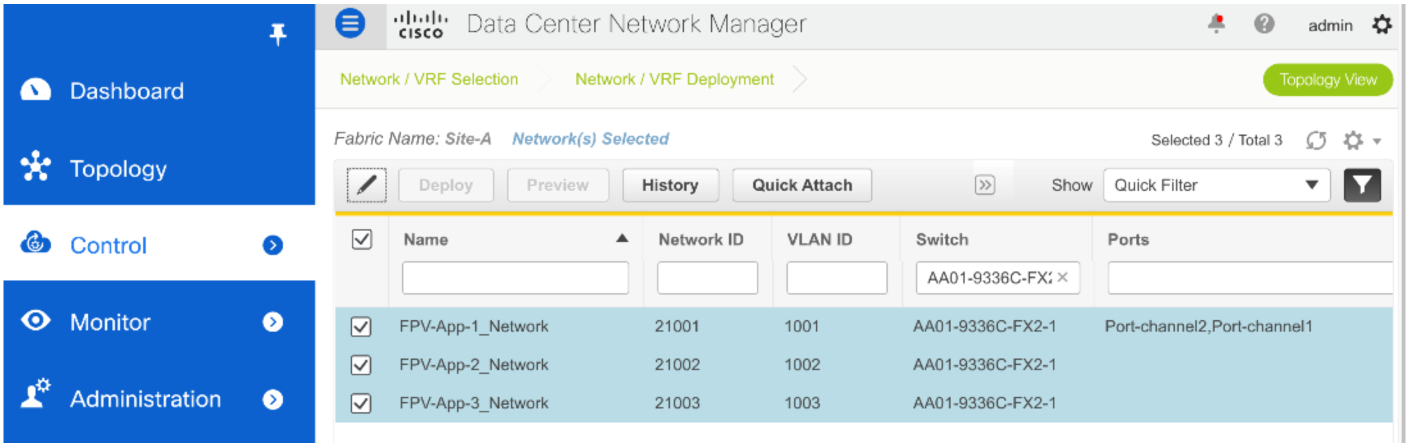

Fabric Name: Site-A *Network(s) Selected*

Selected 3 / Total 3

☑	Name ▲	Network ID	VLAN ID	Switch	Ports
				AA01-9336C-FX2 ×	
☑	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	Port-channel2,Port-channel1
☑	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	
☑	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	

8. Click the pencil icon again.

9. You will now see that the port/interface information is now populated for both leaf switches though the configuration was only done for one switch. Cisco DCNM automatically configures both leaf switches in the same leaf switch vPC pair.

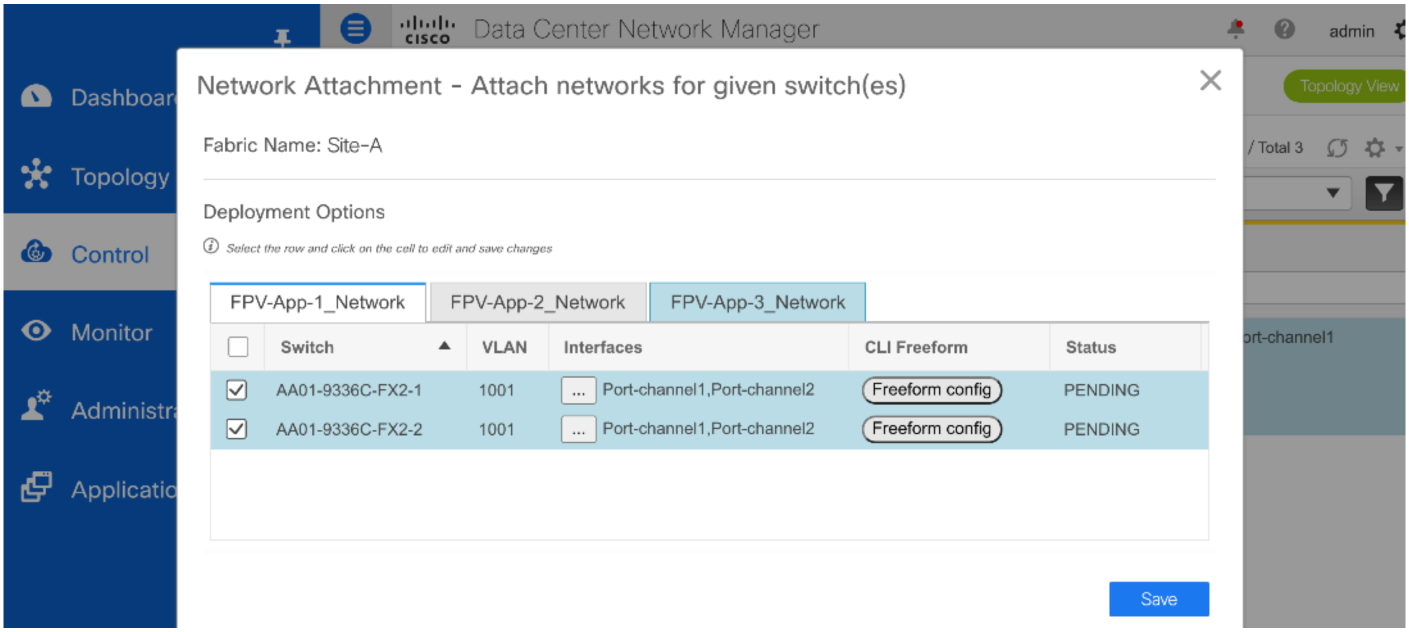

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

FPV-App-1_Network | FPV-App-2_Network | FPV-App-3_Network

☐	Switch ▲	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	1001	... Port-channel1,Port-channel2	Freeform config	PENDING
☑	AA01-9336C-FX2-2	1001	... Port-channel1,Port-channel2	Freeform config	PENDING

10. For each network tab, repeat steps 7-9 to enable these networks on the access-layer connections to Cisco UCS Domain. Click **Save**. All switches are now in **PENDING** state. If you click on the pencil icon again, you will see that the ports or interfaces for the peer leaf switches are also configured now.

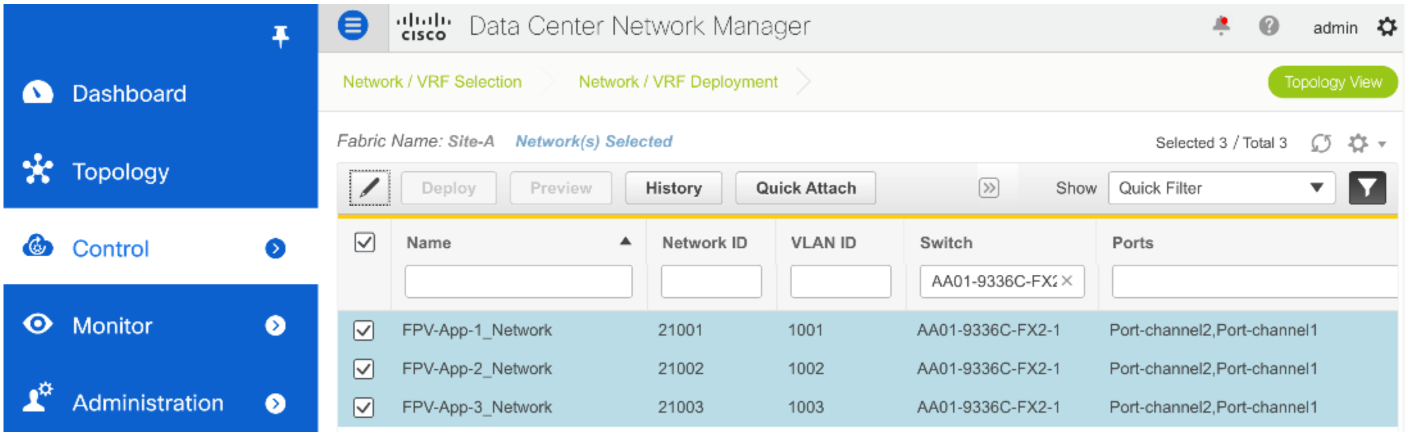

Fabric Name: Site-A ***Network(s) Selected***

☑	Name ▲	Network ID	VLAN ID	Switch	Ports
				AA01-9336C-FX2 ×	
☑	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	Port-channel2,Port-channel1
☑	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	Port-channel2,Port-channel1
☑	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	Port-channel2,Port-channel1

11. Deselect the checkbox next to all switches. Click the **Preview** button. Note that all networks/VLANs are being enabled on the access layer connections to Cisco UCS domain. Click the **X** to close the Preview window.

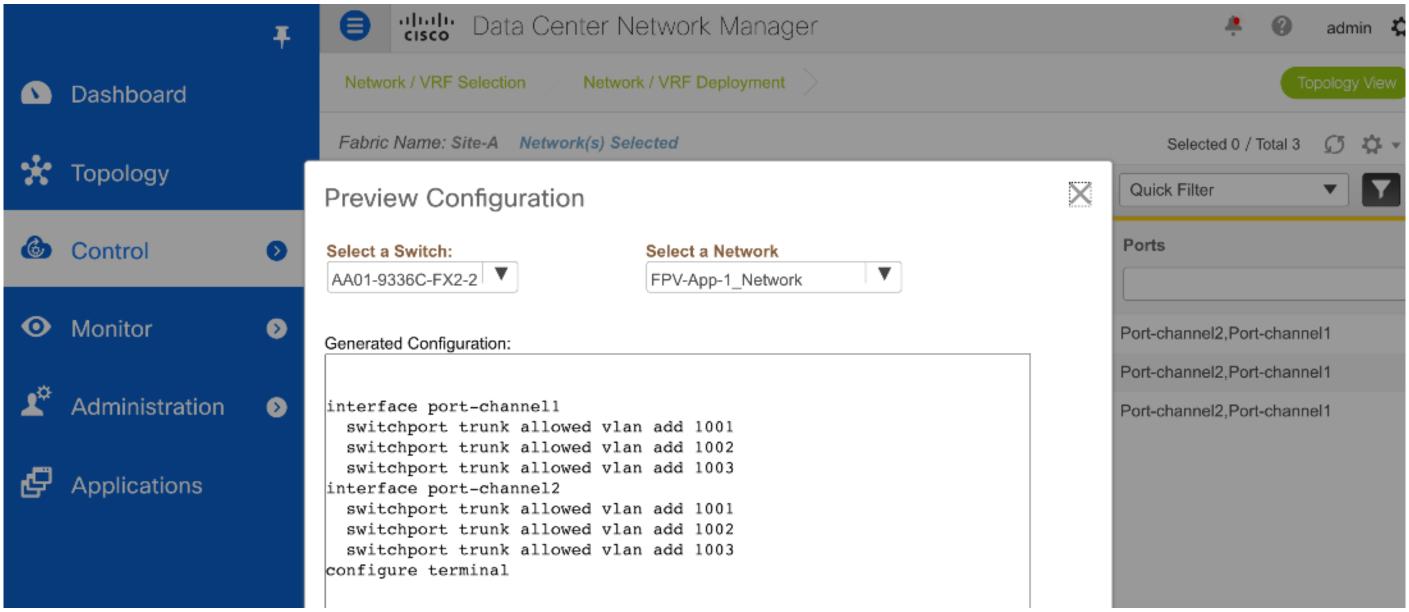

Preview Configuration

Select a Switch: AA01-9336C-FX2-2

Select a Network: FPV-App-1_Network

Generated Configuration:

```
interface port-channel1
  switchport trunk allowed vlan add 1001
  switchport trunk allowed vlan add 1002
  switchport trunk allowed vlan add 1003
interface port-channel2
  switchport trunk allowed vlan add 1001
  switchport trunk allowed vlan add 1002
  switchport trunk allowed vlan add 1003
configure terminal
```

12. Click the **Deploy** button. The status should change from **PENDING** to **IN PROGRESS** to **DEPLOYED.**

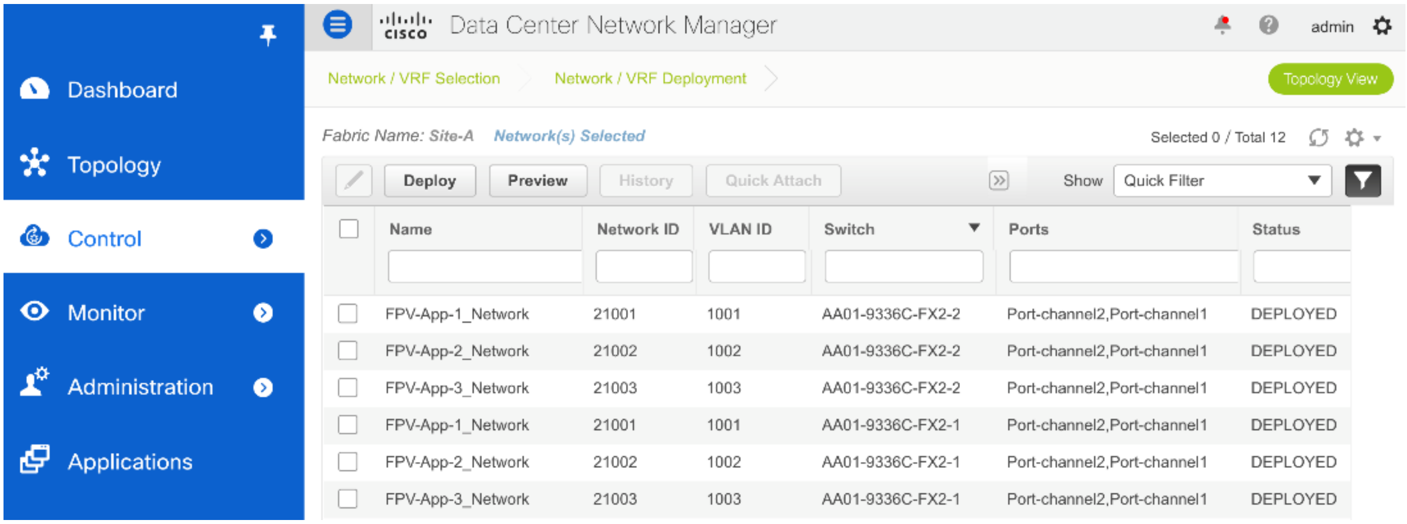

Fabric Name: Site-A Network(s) Selected Selected 0 / Total 12

☐	Name	Network ID	VLAN ID	Switch	Ports	Status
☐	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-2	Port-channel2,Port-channel1	DEPLOYED
☐	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-2	Port-channel2,Port-channel1	DEPLOYED
☐	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-2	Port-channel2,Port-channel1	DEPLOYED
☐	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
☐	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED
☐	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED

Enable Applications Networks on Access-Layer Connections to NetApp Storage Cluster

To enable FlexPod Applications networks on the access-layer connections to the NetApp storage cluster, follow these steps:

1. Repeat steps 1-5 from the previous section.

2. Filter on the first leaf switch in the vPC pair to the NetApp Storage cluster.

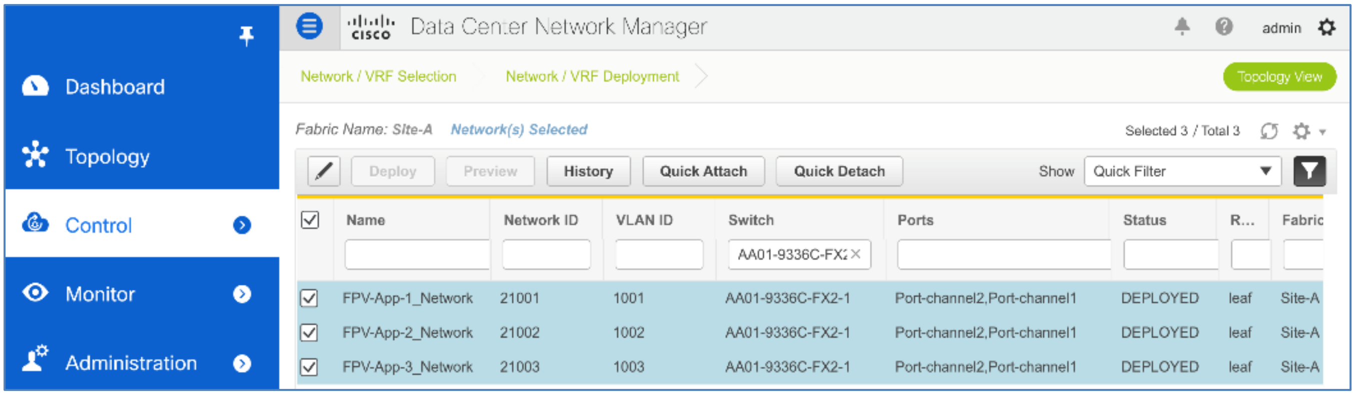

Fabric Name: Site-A Network(s) Selected

Selected 3 / Total 3

☑	Name	Network ID	VLAN ID	Switch	Ports	Status	R...	Fabric
				AA01-9336C-FX2				
☑	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED	leaf	Site-A
☑	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED	leaf	Site-A
☑	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	Port-channel2,Port-channel1	DEPLOYED	leaf	Site-A

3. Click on the pencil icon from the menu.

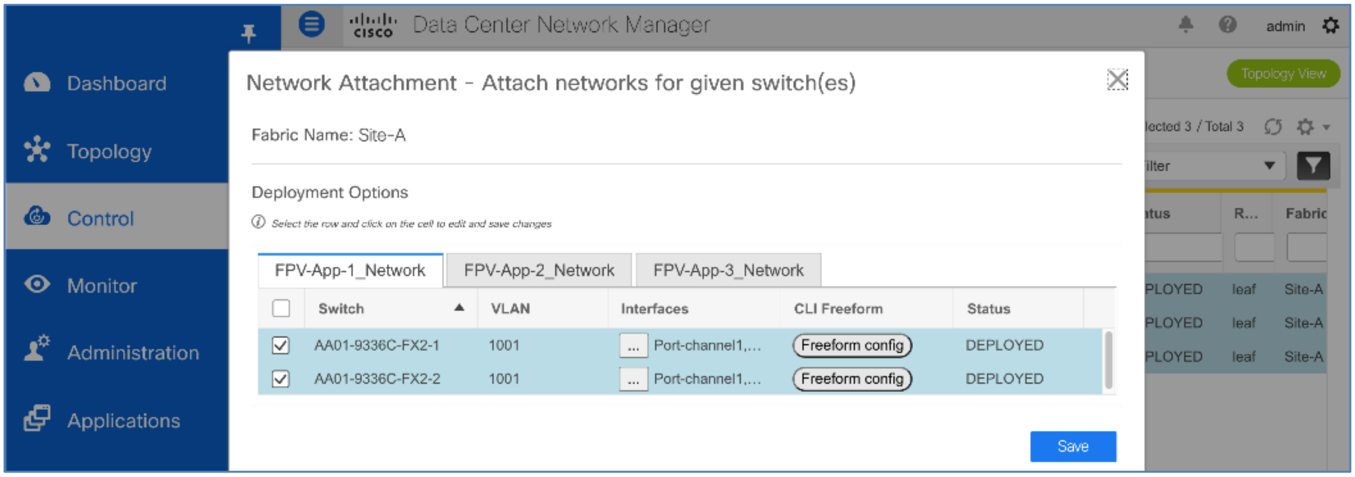

Network Attachment - Attach networks for given switch(es)

Fabric Name: Site-A

Deployment Options

Select the row and click on the cell to edit and save changes

☐	Switch	VLAN	Interfaces	CLI Freeform	Status
☑	AA01-9336C-FX2-1	1001	... Port-channel1,...	Freeform config	DEPLOYED
☑	AA01-9336C-FX2-2	1001	... Port-channel1,...	Freeform config	DEPLOYED

4. For the first network, click the box ... in the **Interfaces** column next to the first switch listed in the leaf switch pair. Select the **Interfaces/Ports** that connect to the NetApp Storage cluster. In this case, two port-channels going to Cisco UCS domain were already configured for these networks – however, two additional port-channels going to NetApp had to be selected in this step.

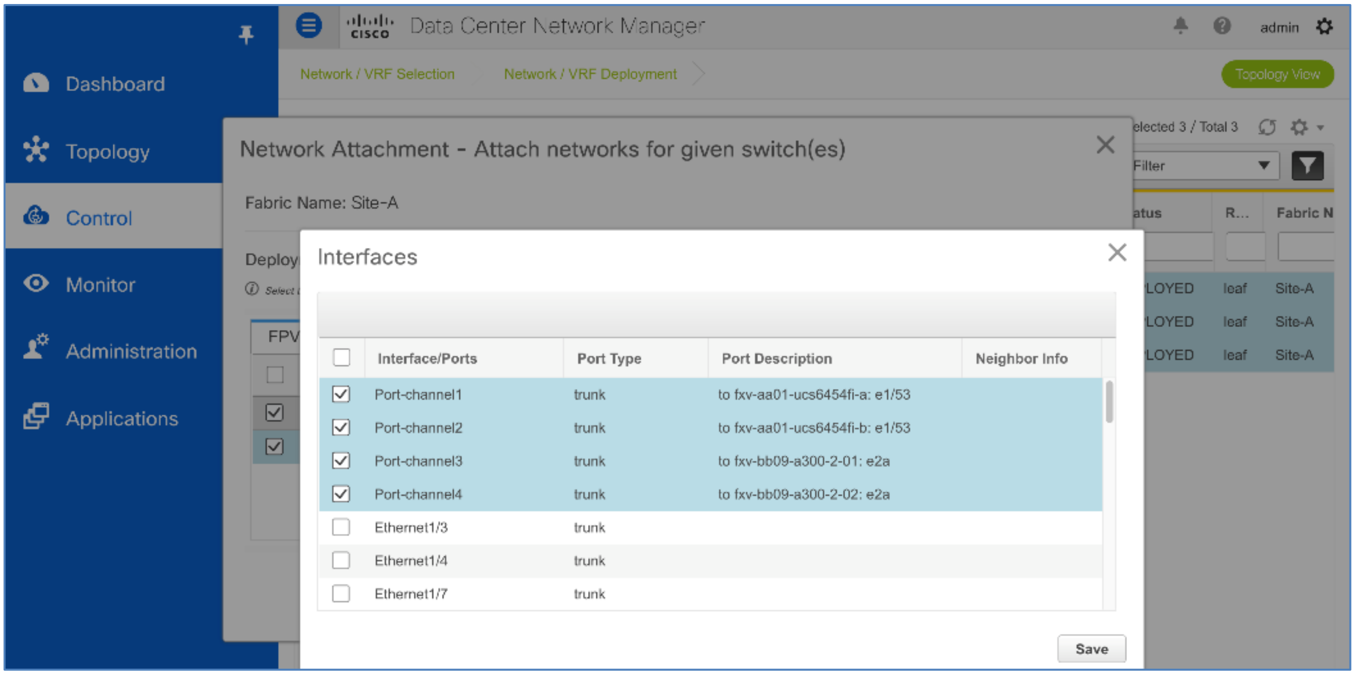

Interfaces

☐	Interface/Ports	Port Type	Port Description	Neighbor Info
☑	Port-channel1	trunk	to fxv-aa01-ucs6454fi-a: e1/53	
☑	Port-channel2	trunk	to fxv-aa01-ucs6454fi-b: e1/53	
☑	Port-channel3	trunk	to fxv-bb09-a300-2-01: e2a	
☑	Port-channel4	trunk	to fxv-bb09-a300-2-02: e2a	
☐	Ethernet1/3	trunk		
☐	Ethernet1/4	trunk		
☐	Ethernet1/7	trunk		

5. Click **Save**.

6. Repeat steps 4-5 for each network tab. Click **Save.** Note that the status of these networks are in **PENDING** state at this stage. Scroll to the right as needed to see all columns available in this view.

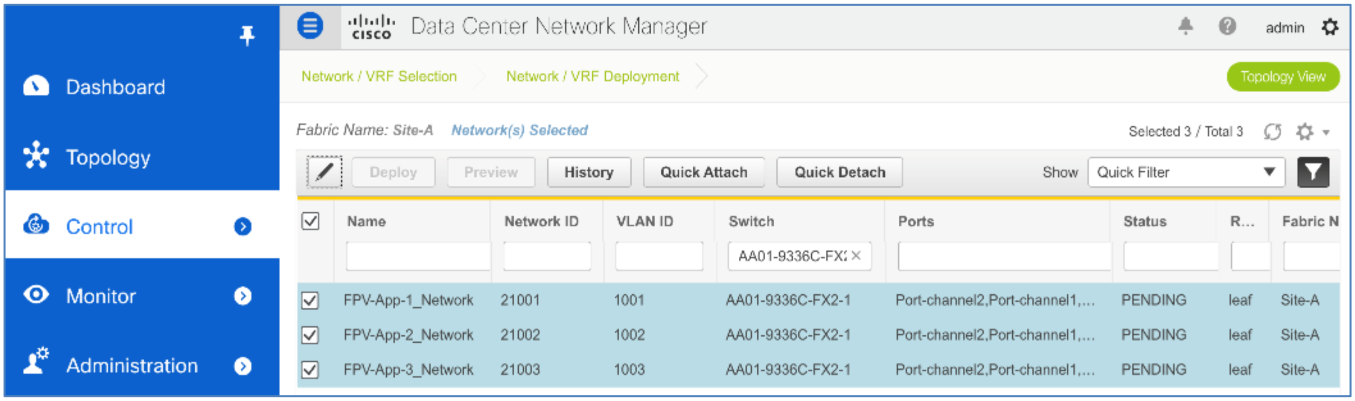

Fabric Name: Site-A *Network(s) Selected*

☑	Name	Network ID	VLAN ID	Switch	Ports	Status	R...	Fabric N
				AA01-9336C-FX: ×				
☑	FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	Port-channel2,Port-channel1,...	PENDING	leaf	Site-A
☑	FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	Port-channel2,Port-channel1,...	PENDING	leaf	Site-A
☑	FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	Port-channel2,Port-channel1,...	PENDING	leaf	Site-A

7. Deselect all networks. Click the **Preview** button to see the pending changes. Click the **X** to close the window.

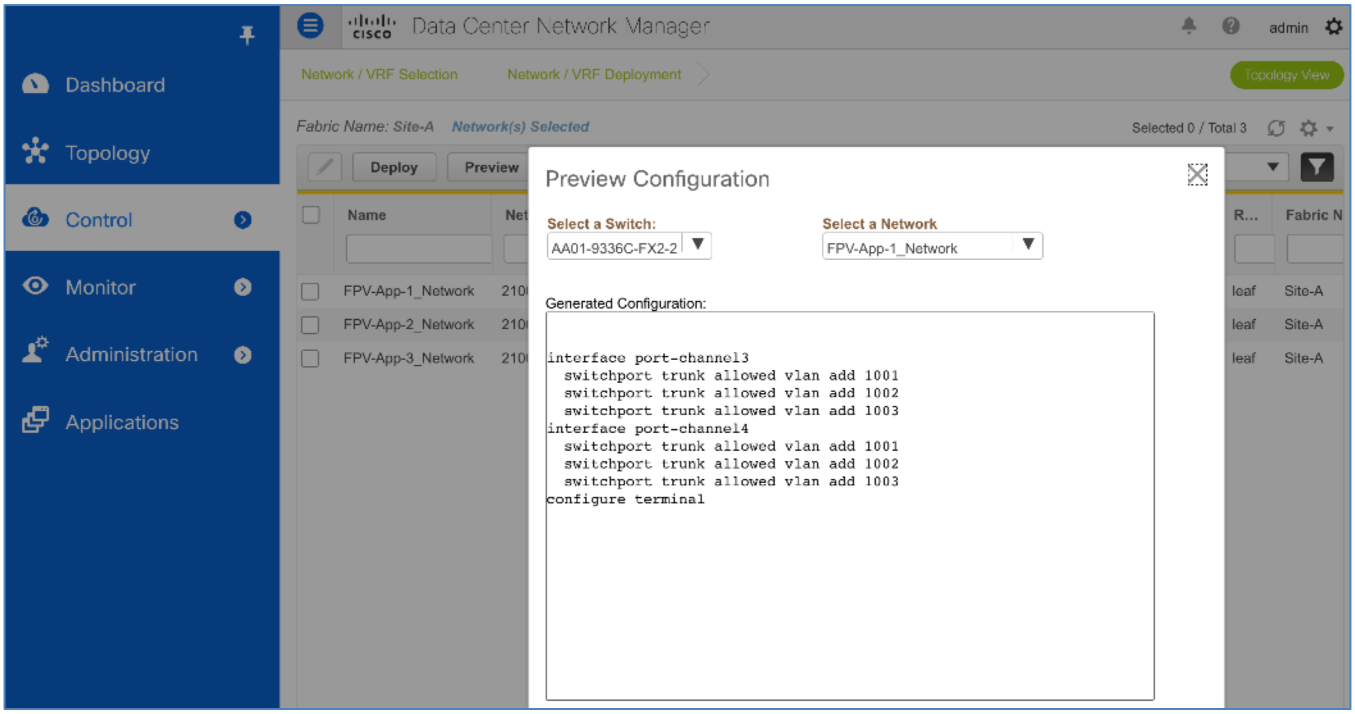

Fabric Name: Site-A Network(s) Selected

Selected 0 / Total 3

Name	R...	Fabric N
FPV-App-1_Network	leaf	Site-A
FPV-App-2_Network	leaf	Site-A
FPV-App-3_Network	leaf	Site-A

Preview Configuration

Select a Switch: AA01-9336C-FX2-2

Select a Network: FPV-App-1_Network

Generated Configuration:

```
interface port-channel3
  switchport trunk allowed vlan add 1001
  switchport trunk allowed vlan add 1002
  switchport trunk allowed vlan add 1003
interface port-channel4
  switchport trunk allowed vlan add 1001
  switchport trunk allowed vlan add 1002
  switchport trunk allowed vlan add 1003
configure terminal
```

8. Click the **Deploy** button. The status should go from **PENDING** to **IN PROGRESS** to **DEPLOYED.** Scroll to the right as needed to see all columns available in this view.

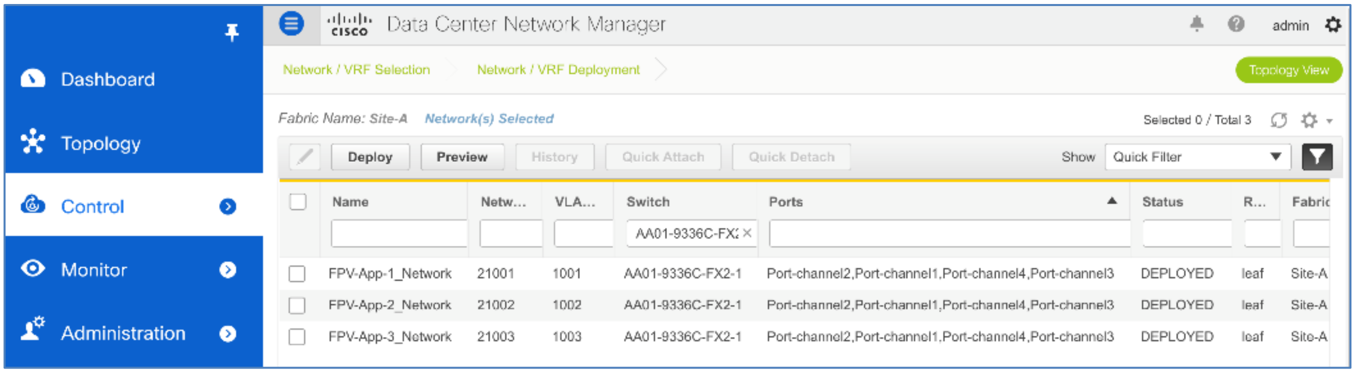

Fabric Name: Site-A Network(s) Selected

Selected 0 / Total 3

Name	Netw...	VLA...	Switch	Ports	Status	R...	Fabric
			AA01-9336C-FX2				
FPV-App-1_Network	21001	1001	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED	leaf	Site-A
FPV-App-2_Network	21002	1002	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED	leaf	Site-A
FPV-App-3_Network	21003	1003	AA01-9336C-FX2-1	Port-channel2,Port-channel1,Port-channel4,Port-channel3	DEPLOYED	leaf	Site-A

Deploy FlexPod Application Storage Networks

The FlexPod Applications deployed in the Application Tenant (FPV-Application-Tenant) may require storage access to iSCSI or NFS volumes hosted on the NetApp cluster. In this deployment, the storage networks shown in Table 18 are deployed.

To deploy storage access for the Applications hosted on the FlexPod infrastructure, follow these steps:

1. Deploy Application storage networks from Cisco DCNM using the setup information in Table 18 – use the procedures outlined here.

2. Enable Application storage networks on the access-layer connection to Cisco UCS Domain - use the procedures outlined here.

3. Enable Application storage networks on the access-layer connection to NetApp Storage - use the procedures outlined here.

Solution Deployment - Storage

NetApp All Flash FAS A300 Controllers

See the following section (NetApp Hardware Universe) for planning the physical location of the storage systems:

- Site Preparation
- System Connectivity Requirements
- Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements
- AFF Series Systems

NetApp Hardware Universe

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

To confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install, follow these steps found at the NetApp Support site.

To configure the HWU, follow these steps:

1. Access the HWU application to view the System Configuration guides. Click the Platforms menu to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

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

Controllers

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

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

NetApp ONTAP 9.7

Complete Configuration Worksheet

Before running the setup script, complete the Cluster setup worksheet in the ONTAP 9 Documentation Center. You must have access to the NetApp Support site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the Software setup section of the ONTAP 9 Documentation Center to learn about configuring ONTAP. Table 18 lists the information needed to configure two ONTAP nodes. Customize the cluster-detail values with the information applicable to your deployment.

Table 19. ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
ONTAP 9.7 URL	<url-boot-software>

Configure Node 01

To configure node 01, follow these steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.7 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.7 is the version being booted, 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 continue the installation.

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

7. Enter **n** to skip the reboot

8. Choose option 7 from the menu: **Install new software first**

9. Enter **y** to continue the installation

10. Enter the IP address, netmask, and default gateway for **e0M**.

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

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

This web server must be pingable from node 01

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

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

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

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

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

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

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

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

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

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

18. Enter yes to erase all the data on the disks.

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

Configure Node 02

To configure node 02, follow these steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.7 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.7 is the version being booted, 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 continue the installation..

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

7. Enter **n** to skip the reboot

8. Choose option 7 from the menu: **Install new software first**

9. Enter y to continue the installation.

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

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

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

This web server must be pingable from node 2

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

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

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

14. Enter 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.

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

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

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

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

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

18. Enter **yes** to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.7 boots on the node for the first time.

To set up a node, follow these steps:

1. Follow the prompts to set up node 01.

```
Welcome to node setup.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the setup wizard.
     Any changes you made before quitting will be saved.

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

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

Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created

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

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

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

Table 20. Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-sp-ip>
Node 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>

Cluster Detail	Cluster Detail Value
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>
NTP server B IP address	<switch-b-ntp-ip>
SNMPv3 User	<snmp-v3-usr>
SNMPv3 Authentication Protocol	<snmp-v3-auth-proto>
SNMPv3 Privacy Protocol	<snmpv3-priv-proto>

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

3. In the Initialize Storage System screen, follow these steps:

   a. Enter the cluster name and administrator password.

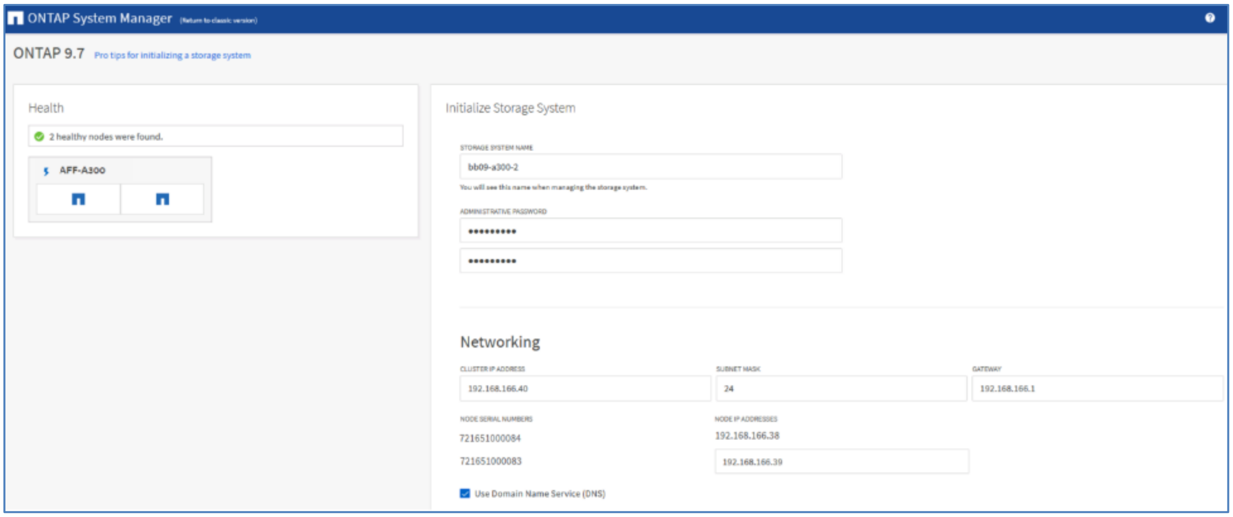

ONTAP System Manager (Return to classic version)
ONTAP 9.7
Pro tips for initializing a storage system
Health
AFF-A300
Initialize Storage System

✓ 2 healthy nodes were found.

STORAGE SYSTEM NAME: bb09-a300-2

You will see this name when managing the storage system.

ADMINISTRATIVE PASSWORD: •••••••••

•••••••••

Networking

CLUSTER IP ADDRESS	SUBNET MASK	GATEWAY
192.168.166.40	24	192.168.166.1

NODE SERIAL NUMBERS	NODE IP ADDRESSES
721651000084	192.168.166.38
721651000083	192.168.166.39

☑ Use Domain Name Service (DNS)

   b. Under Networking section enter Cluster IP, subnet mask and gateway address followed by node1 and node2 IP address.

   c. Enter the DNS domain names and name server address.

   d. Enter the primary and alternate NTP server.

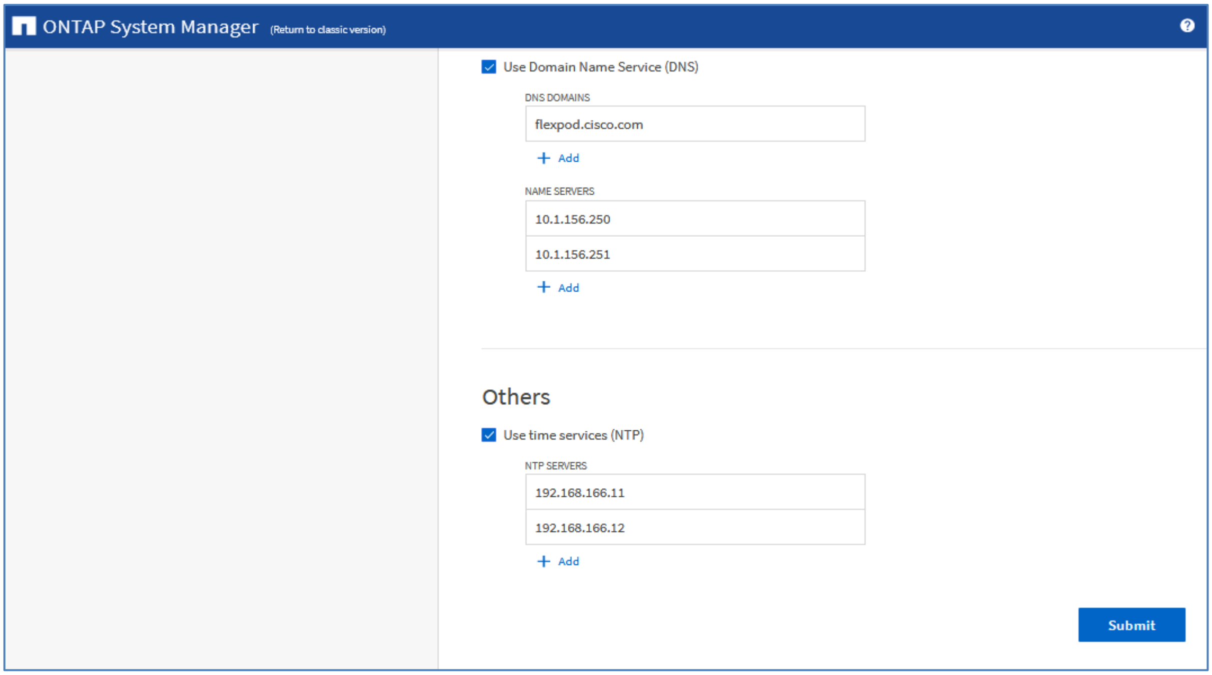

☑ Use Domain Name Service (DNS)

DNS DOMAINS: flexpod.cisco.com

NAME SERVERS: 10.1.156.250, 10.1.156.251

Others

☑ Use time services (NTP)

NTP SERVERS: 192.168.166.11, 192.168.166.12

4. Click Submit.

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

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

Cluster license and feature licenses can also be installed after completing cluster creation.

The cluster setup is triggered now.

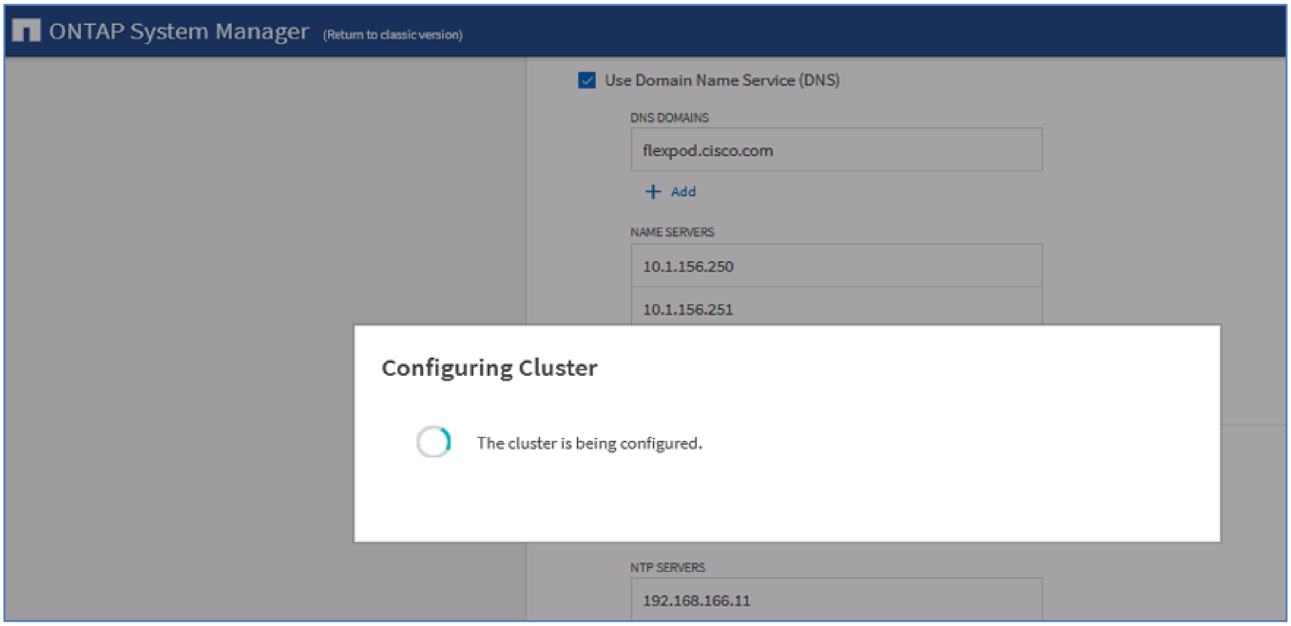

☑ Use Domain Name Service (DNS)

DNS DOMAINS

flexpod.cisco.com

NAME SERVERS

10.1.156.250

10.1.156.251

Configuring Cluster

The cluster is being configured.

NTP SERVERS

192.168.166.11

5. Cluster setup complete message will pop up and the page will be redirected to System Manager.

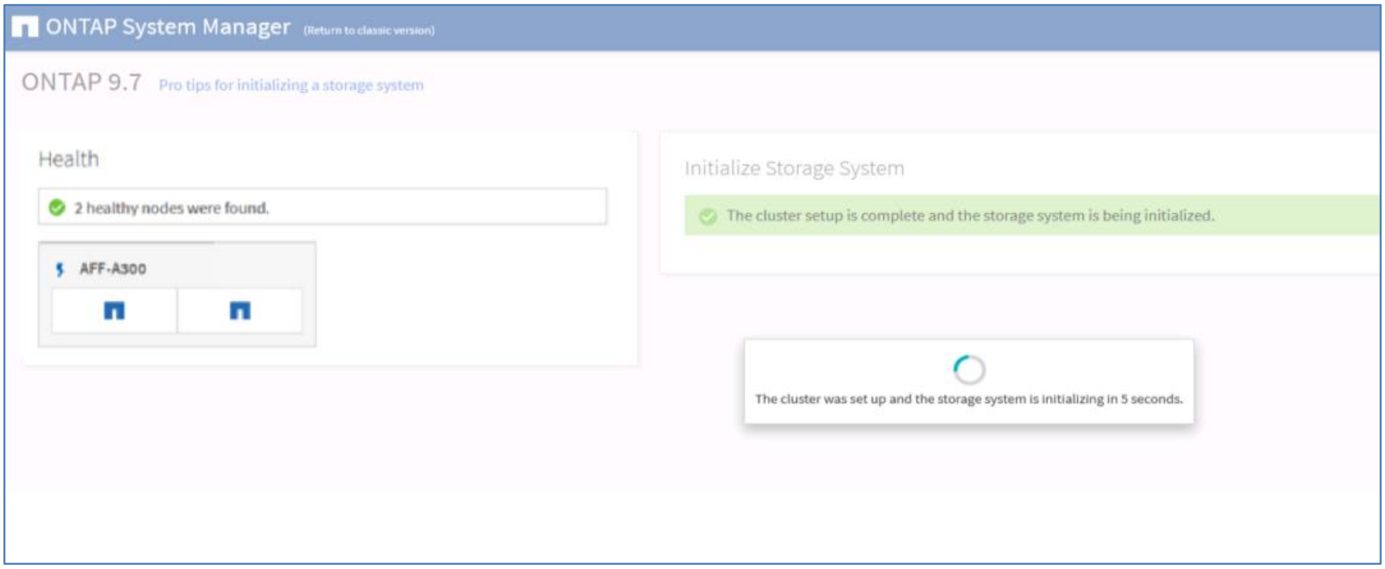

ONTAP System Manager (Return to classic version)
ONTAP 9.7
Pro tips for initializing a storage system
Health
AFF-A300
Initialize Storage System

2 healthy nodes were found.

The cluster setup is complete and the storage system is being initialized.

The cluster was set up and the storage system is initializing in 5 seconds.

6. Login to System Manager and Under Cluster click Overview to see the Node details.

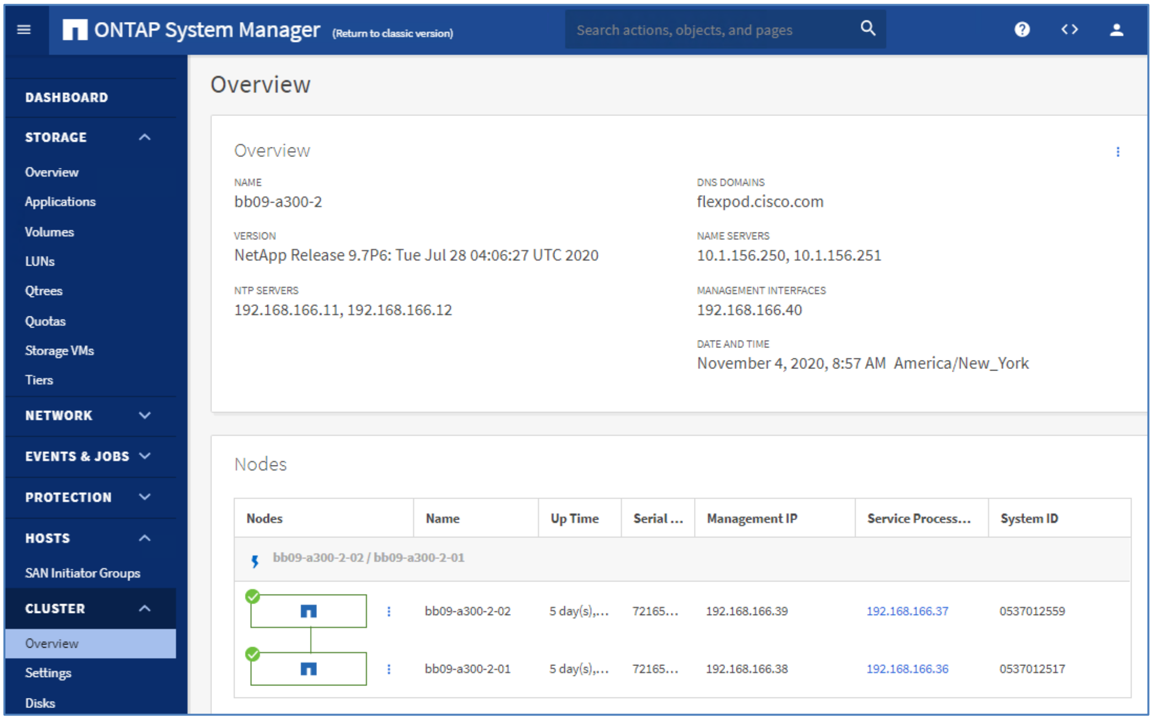

Overview

Overview

NAME
bb09-a300-2

VERSION
NetApp Release 9.7P6: Tue Jul 28 04:06:27 UTC 2020

NTP SERVERS
192.168.166.11, 192.168.166.12

DNS DOMAINS
flexpod.cisco.com

NAME SERVERS
10.1.156.250, 10.1.156.251

MANAGEMENT INTERFACES
192.168.166.40

DATE AND TIME
November 4, 2020, 8:57 AM America/New_York

Nodes

Nodes	Name	Up Time	Serial ...	Management IP	Service Process...	System ID
bb09-a300-2-02 / bb09-a300-2-01						
	bb09-a300-2-02	5 day(s),...	72165...	192.168.166.39	192.168.166.37	0537012559
	bb09-a300-2-01	5 day(s),...	72165...	192.168.166.38	192.168.166.36	0537012517

7. To enable and configure AutoSupport, expand Cluster and click on Settings and click More options.

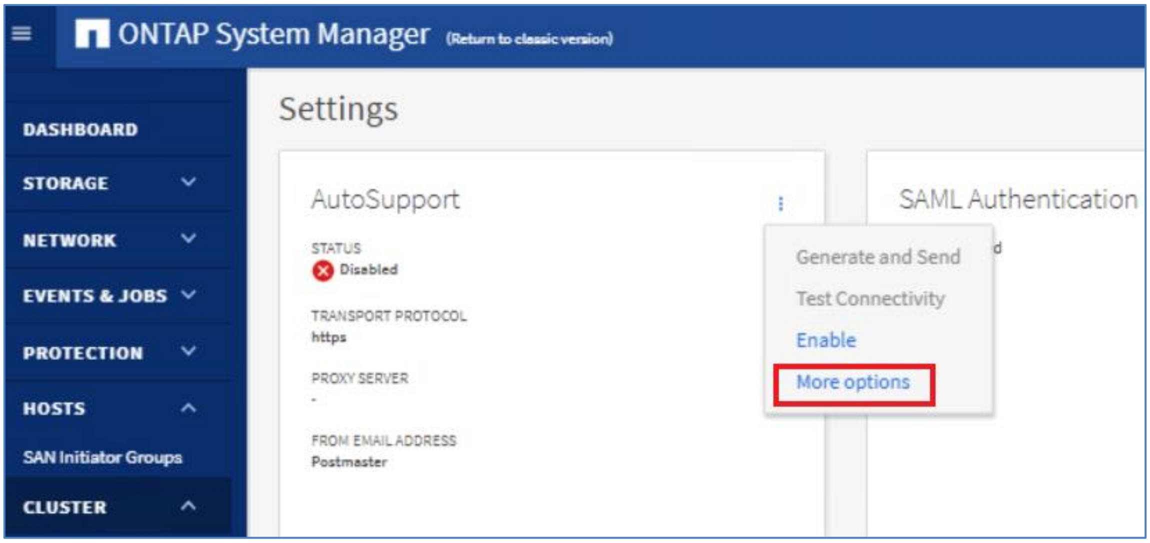

Settings

AutoSupport

STATUS
Disabled

TRANSPORT PROTOCOL
https

PROXY SERVER
-

FROM EMAIL ADDRESS
Postmaster

8. Click Edit to change the transport protocol and provide the details.

9. Click Save.

10. In the EMAIL section provide the FROM and RECEPIENTS address and click Save.

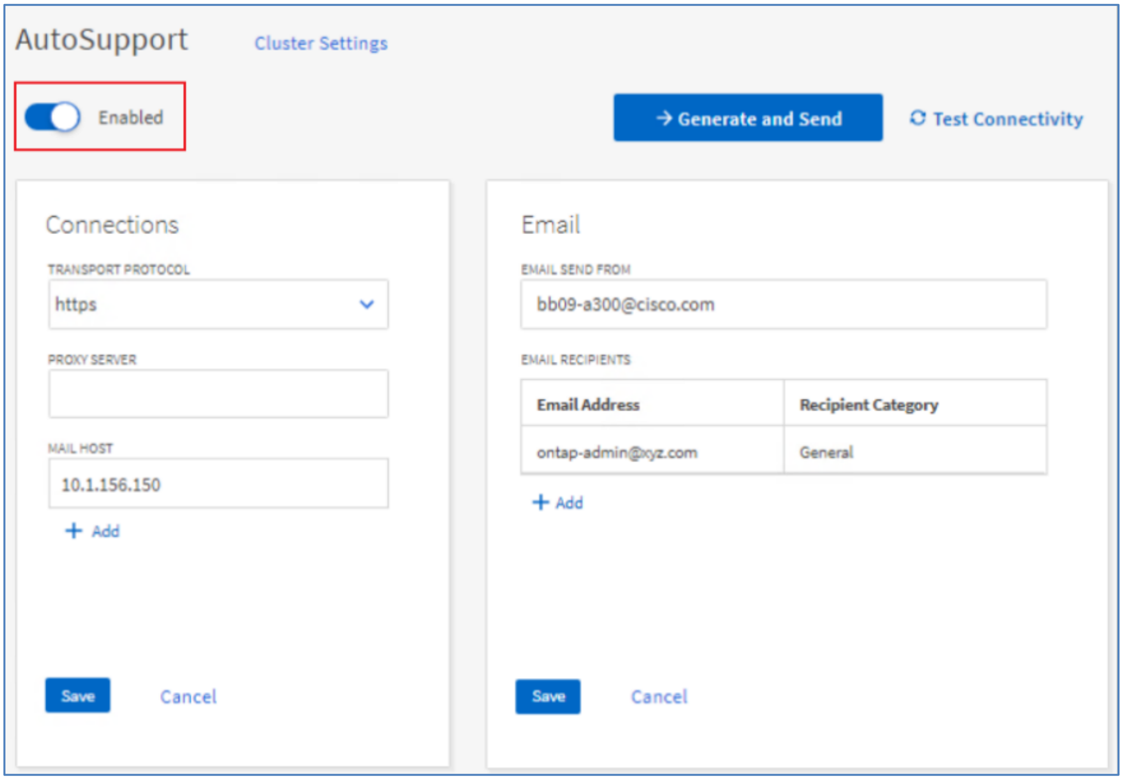

AutoSupport

Enabled

Connections

TRANSPORT PROTOCOL: https

PROXY SERVER:

MAIL HOST: 10.1.156.150

Email

EMAIL SEND FROM: bb09-a300@cisco.com

EMAIL RECIPIENTS

Email Address	Recipient Category
ontap-admin@xyz.com	General

11. Click on Cluster Settings near AutoSupport

12. Click on the right arrow in the License section.

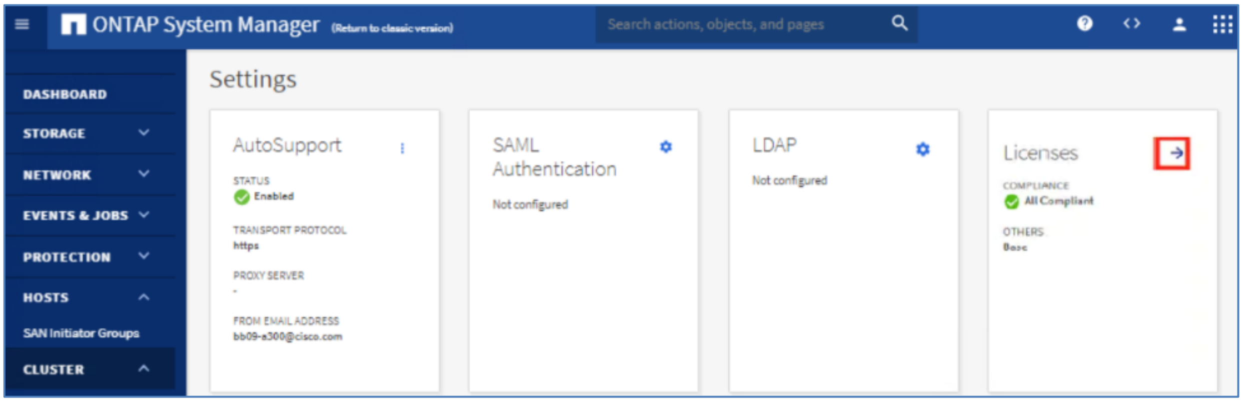

Settings

AutoSupport

STATUS: Enabled

TRANSPORT PROTOCOL: https

PROXY SERVER: -

FROM EMAIL ADDRESS: bb09-a300@cisco.com

SAML Authentication

Not configured

LDAP

Not configured

Licenses

COMPLIANCE: All Compliant

OTHERS: Base

13. Click Add to add the required License to the cluster and enter the license keys in a comma separated list.

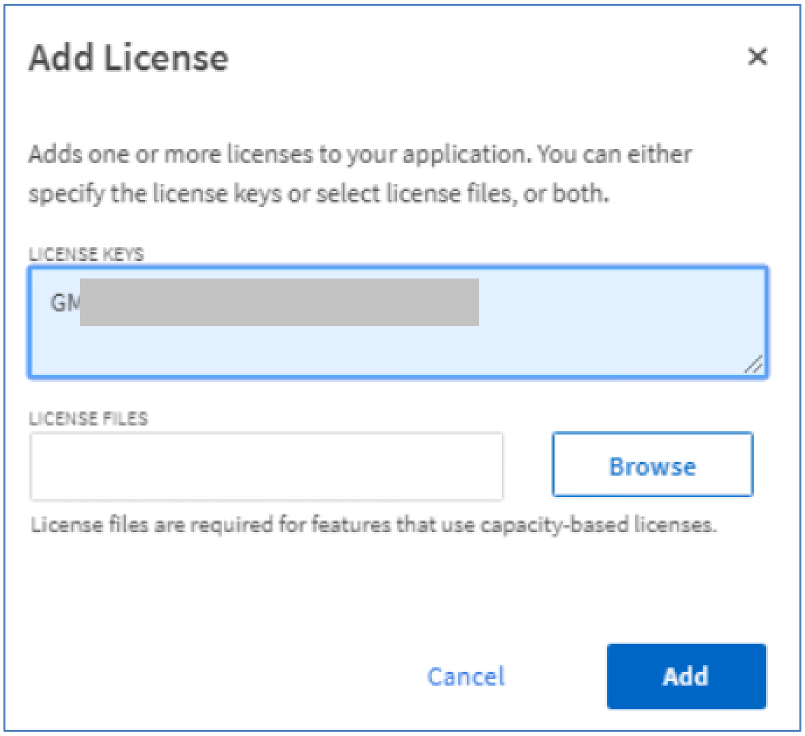

Adds one or more licenses to your application. You can either specify the license keys or select license files, or both.

LICENSE KEYS

GM[redacted]

LICENSE FILES

License files are required for features that use capacity-based licenses.

14. Click Storage and then click Tiers to configure the storage aggregates.

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

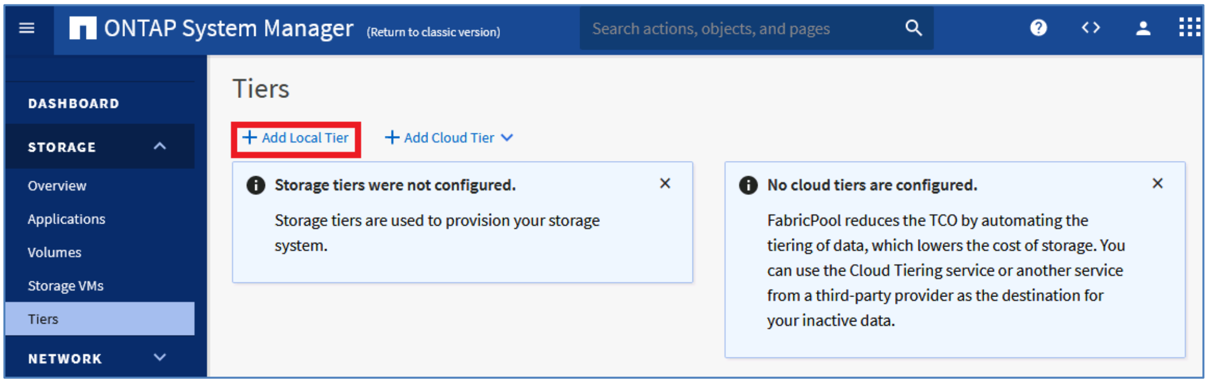

Storage tiers were not configured.

Storage tiers are used to provision your storage system.

No cloud tiers are configured.

FabricPool reduces the TCO by automating the tiering of data, which lowers the cost of storage. You can use the Cloud Tiering service or another service from a third-party provider as the destination for your inactive data.

16. Expand the recommendation details and click Save.

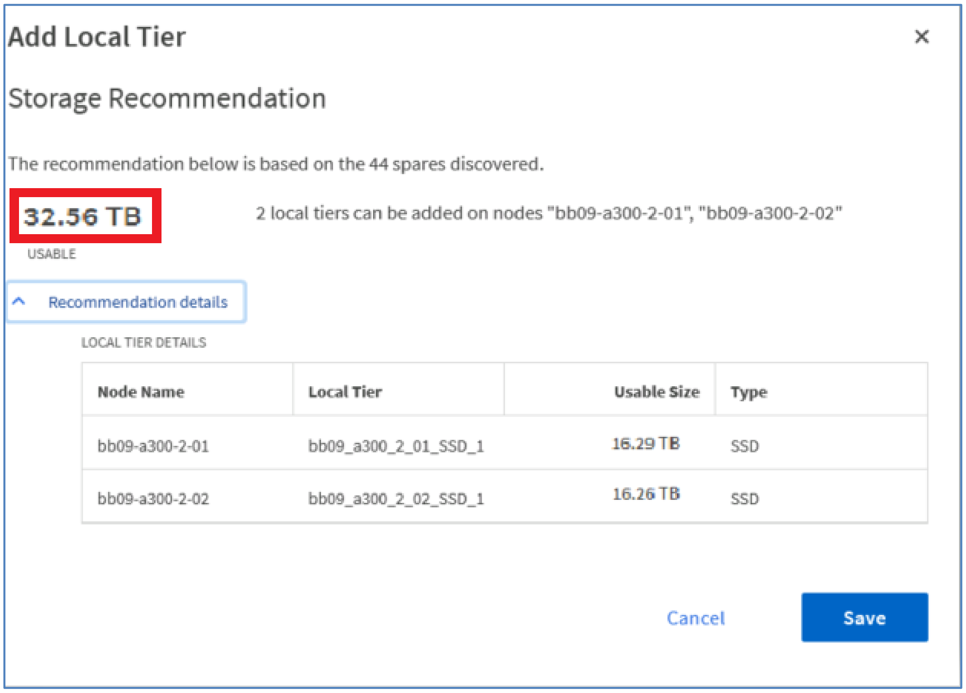

Add Local Tier

Storage Recommendation

The recommendation below is based on the 44 spares discovered.

32.56 TB
USABLE

2 local tiers can be added on nodes "bb09-a300-2-01", "bb09-a300-2-02"

LOCAL TIER DETAILS

Node Name	Local Tier	Usable Size	Type
bb09-a300-2-01	bb09_a300_2_01_SSD_1	16.29 TB	SSD
bb09-a300-2-02	bb09_a300_2_02_SSD_1	16.26 TB	SSD

Log into the Cluster

To log into the cluster, follow these steps:

1. Open an SSH connection to either the cluster IP or the host name.

2. Log in to the admin user with the password you provided earlier.

Verify Storage Failover

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

1. Verify the status of the storage failover.

```
storage failover show
```

Both <st-node01> and <st-node02> must be capable of performing a takeover. Continue with step 2 if the nodes are capable of performing a takeover.

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

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

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

4. Continue with step 5 if high availability is not configured.

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

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

6. Verify that hardware assist is correctly configured.

```
storage failover hwassist show
```

Set Auto-Revert on Cluster Management

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

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

Run the following command:

```
net interface modify -vserver <vservername> -lif <mgmtlif> -auto-revert true
```

Zero All Spare Disks

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

```
disk zerospares
```

Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an AFF configuration. Disk autoassign should have assigned one data partition to each node in an HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the **`disk option modify`** command. Spare partitions can then be moved from one node to another by running the **`disk removeowner`** and **`disk assign`** commands.

Set Up Service Processor Network Interface

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

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

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

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

Create Auto-provisioned Aggregates

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

```
bb09-a300-2::*> storage aggregate auto-provision -verbose
Per node summary of new aggregates to create, discovered spares, and also
remaining spare disks and partitions after aggregate creation:
                    New    Total New -Discovered Spare- -Remaining Spare-
Node               Aggrs  Usable Size  Disks  Partitions  Disks Partitions
------------------ ----- ------------ ------ ----------- ------ ----------
bb09-a300-2-01           1      16.29TB      0          24      0      1
bb09-a300-2-02           1      16.26TB      0          24      0      1
------------------ ----- ------------ ------ ----------- ------ ----------
Total:                 2      32.56TB      0          48      0       2

New data aggregates to create with counts of
disks and partitions to be used:
                                                -Devices To Use-
Node               New Data Aggregate            Usable Size Disks Partitions
------------------ ---------------------------- ------------ ----- ----------
bb09-a300-2-01       bb09_a300_2_01_SSD_1           16.29TB     0         23
bb09-a300-2-02       bb09_a300_2_02_SSD_1           16.26TB     0         23

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

RAID Group In New                   Disk       Usable Disk Or   ---Count---
Data Aggregate To Be Created           Type         Size Partition Data Parity
-------------------------------------- ------ ---------- --------- ---- ------
/bb09_a300_2_01_SSD_1/plex0/rg0        SSD  894.3GB partition   21      2
/bb09_a300_2_02_SSD_1/plex0/rg1        SSD  894.3GB partition   21      2

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

```
               Disk         Device Disk Or   Remaining

Node             Type    Usable Size Partition    Spares

------------------ ------ ------------ --------- ---------

bb09-a300-2-01       SSD    894.3GB partition         1

bb09-a300-2-02       SSD    894.3GB partition         1


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


Info: Aggregate auto provision has started. Use the "storage aggregate

      show-auto-provision-progress" command to track the progress.
```

Create Aggregates Manually (Optional)

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

To create new aggregates, run the following commands:

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

You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate.

For all-flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.

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

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

Remove Ports from Default Broadcast Domain

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

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

network port broadcast-domain show
```

Disable Flow Control on 10GbE and 40GbE Ports

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

1. Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port 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 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
```

Enable Cisco Discovery Protocol

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

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

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

Enable Link-layer Discovery Protocol on all Ethernet Ports

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

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

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

Create Management Broadcast Domain

If the management interfaces are required to be on a separate VLAN, create a new broadcast domain for those interfaces by running the following command:

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

Create NFS Broadcast Domain

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

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

Create iSCSI Broadcast Domain

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

```
network port broadcast-domain create -broadcast-domain Infra-iSCSI-A -mtu 9000
network port broadcast-domain create -broadcast-domain Infra-iSCSI-B -mtu 9000
```

Create Interface Groups

To create the LACP interface groups for the 40GbE data interfaces, run the following commands:

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

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

Create VLANs

To create VLANs, follow these steps:

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

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

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

network port vlan show
```

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

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

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

3. Create the iSCSI VLAN ports for the iSCSI LIFs on each storage controller

```
network port vlan create –node <st-node01> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create –node <st-node02> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create –node <st-node01> -vlan-name a0a-<infra-iscsi-b-vlan-id>
network port vlan create –node <st-node02> -vlan-name a0a-<infra-iscsi-b-vlan-id>
```

4. To add each of the iSCSI VLAN ports to the corresponding broadcast domain, run the following commands:

```
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-node01>:a0a-<infra-iscsi-
a-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-node01>:a0a-<infra-iscsi-
b-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-node02>:a0a-<infra-iscsi-
a-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-node02>:a0a-<infra-iscsi-
b-vlan-id>
network port broadcast-domain show
```

Configure Network Time Protocol

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

1. Set the time zone for the cluster.

```
timezone <timezone>
```

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

2. Set the date for the cluster.

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

The format for the date is <[Century][Year][Month][Day][Hour][Minute].[Second]> (for example, 202009271549.30).

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

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

Configure Simple Network Management Protocol

To configure the Simple Network Management Protocol (SNMP), follow these steps:

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

```
snmp contact <snmp-contact>
snmp location “<snmp-location>”
snmp init 1
options snmp.enable on
```

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

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

Configure SNMPv3 Access

SNMPv3 offers advanced security by using encryption and passphrases. The SNMPv3 user can run SNMP utilities from the traphost using the authentication and privacy settings that you specify. To configure SNMPv3 access, run the following commands:

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

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

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

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

Enter the authentication protocol password again:

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

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:
```

For additional detail refer to the SNMP Configuration Express Guide

Create SVM

To create an infrastructure SVM, follow these steps:

1. Run the vserver create command.

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

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

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

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

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

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

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

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

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

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

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, follow these steps:

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

```
volume create -vserver Infra-SVM -volume infra_svm_root_m01 -aggregate aggr1_node01 -size 1GB -type DP

volume create -vserver Infra-SVM -volume infra_svm_root_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

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

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

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-SVM:infra_svm_root_m01 -type
LS -schedule 15min
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-SVM:infra_svm_root_m02 -type
LS -schedule 15min
```

4. Initialize the mirroring relationship.

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

Create Block Protocol (iSCSI) Service

Run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

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

Configure HTTPS Access

To configure secure access to the storage controller, follow these steps:

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

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

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

```
security certificate show
```

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

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

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

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

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

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

```
 security certificate show
```

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

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

7. Disable HTTP cluster management access.

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

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

8. Change back to the normal admin privilege level and verify that the system logs are available in a web browser.

```
set –privilege admin
```

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

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

Configure NFSv3

To configure NFSv3 on the SVM, follow these steps:

1. Create a new rule for the infrastructure NFS subnet in the default export policy.

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

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

```
volume modify –vserver Infra-SVM –volume infra_svm_root –policy default
```

Create FlexVol Volumes

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

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

To create a FlexVol volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node01 -size 1TB -state online -
policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_datastore_2 -aggregate aggr1_node02 -size 1TB -state online -
policy default -junction-path /infra_datastore_2 -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 200GB -state online -policy
default -space-guarantee none -percent-snapshot-space 0

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

If SnapCenter will be used to back up the infra datastores volume, add “-snapshot-policy none” to the end of the volume create command for the infra datastores volume.

Create Boot LUNs

To create boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -size 32GB -ostype vmware -space-reserve
disabled

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

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-03 -size 32GB -ostype vmware -space-reserve
disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-04 -size 32GB -ostype vmware -space-reserve
disabled

lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-05 -size 32GB -ostype vmware -space-reserve
disabled
```

Modify Volume Efficiency

On NetApp All Flash FAS systems, deduplication is enabled by default. To disable the efficiency policy on the infra_swap volume, run the following command:

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

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
```

Create iSCSI LIFs

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

```
network interface create -vserver Infra-SVM -lif iscsi-lif-1a -role data -data-protocol iscsi -home-node <st-
node01> -home-port a0a-<infra-iscsi-a-vlan-id> -address <st-node01-infra-iscsi-a-ip> -netmask <infra-iscsi-a-
mask> -status-admin up
network interface create -vserver Infra-SVM -lif iscsi-lif-1b -role data -data-protocol iscsi -home-node <st-
node01> -home-port a0a-<infra-iscsi-b-vlan-id> -address <st-node01-infra-iscsi-b-ip> -netmask <infra-iscsi-b-
mask> -status-admin up
network interface create -vserver Infra-SVM -lif iscsi-lif-2a -role data -data-protocol iscsi -home-node <st-
node02> -home-port a0a-<infra-iscsi-a-vlan-id> -address <st-node02-infra-iscsi-a-ip> -netmask <infra-iscsi-a-
mask> -status-admin up
network interface create -vserver Infra-SVM -lif iscsi-lif-2b -role data -data-protocol iscsi -home-node <st-
node02> -home-port a0a-<infra-iscsi-b-vlan-id> -address <st-node02-infra-iscsi-b-ip> -netmask <infra-iscsi-b-
mask> -status-admin up
network interface show
```

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the in-band management network, follow these steps:

1. Run the following commands:

```
network interface create –vserver Infra-SVM –lif svm-mgmt –role data –data-protocol none –home-node <st-
node02> -home-port  a0a-<ib-mgmt-vlan-id> –address <svm-mgmt-ip> –netmask <svm-mgmt-mask> -status-admin up –
failover-policy broadcast-domain-wide –firewall-policy mgmt –auto-revert true
```

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

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

network route show
```

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

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

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

Configure and Test AutoSupport

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

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

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

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

Solution Deployment – Compute

Cisco UCS Base Configuration

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

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

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

Cisco UCS Fabric Interconnect A

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

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

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

Enter the management mode. (ucsm/intersight)? ucsm

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

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

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

Enter the password for "admin": <password>
Confirm the password for "admin": <password>

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

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

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

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

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

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

Cluster IPv4 address : <ucs-cluster-ip>

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

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

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

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

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

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

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

Cisco UCS Fabric Interconnect B

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

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

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

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

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

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

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

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

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

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

Cisco UCS Setup

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (Cisco UCS) environment, follow these steps:

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

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

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

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

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

5. Click Login to log into Cisco UCS Manager.

Anonymous Reporting

To enable anonymous reporting, follow this step:

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

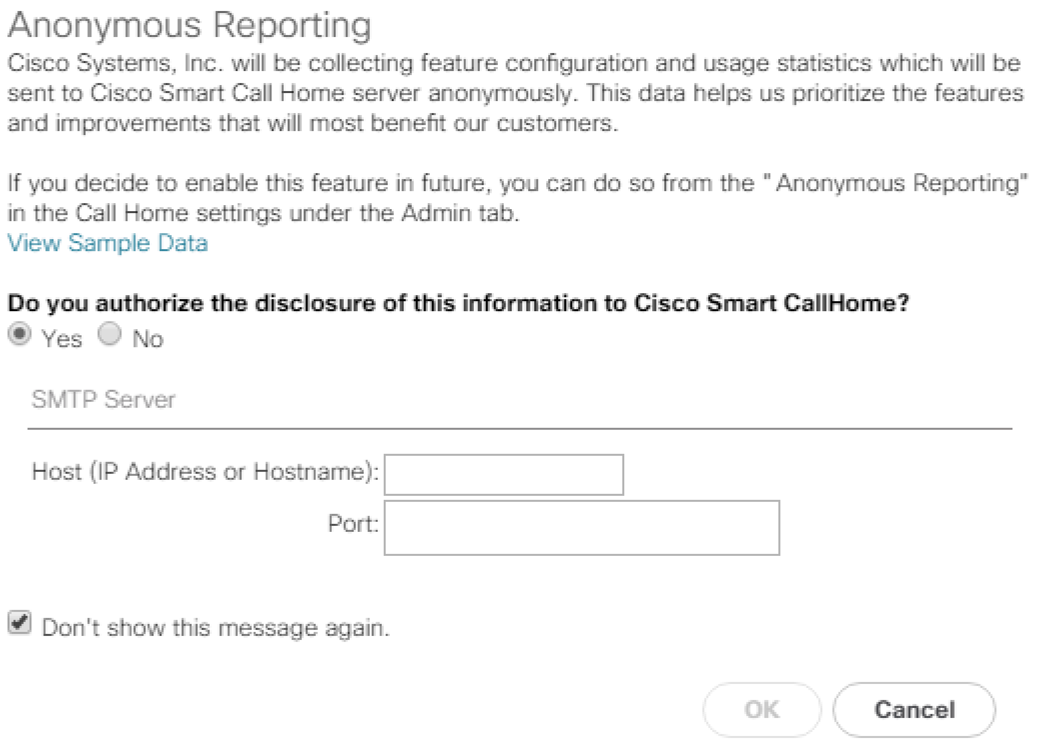

Anonymous Reporting

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

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

View Sample Data

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

● Yes ○ No

SMTP Server

Host (IP Address or Hostname):

Port:

☑ Don't show this message again.

Upgrade Cisco UCS Manager Software to Version 4.1(2a)

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

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

Configure Cisco UCS Call Home

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

1. In Cisco UCS Manager, click Admin.

2. Choose All > Communication Management > Call Home.

3. Change the State to On.

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

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Cisco Nexus switches, follow these steps:

1. In Cisco UCS Manager, click Admin.

2. Expand All > Time Zone Management.

3. Choose Timezone.

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

5. Click Save Changes and then click OK.

6. Click Add NTP Server.

7. Enter <ntp-server> and click OK. Click OK on the confirmation.

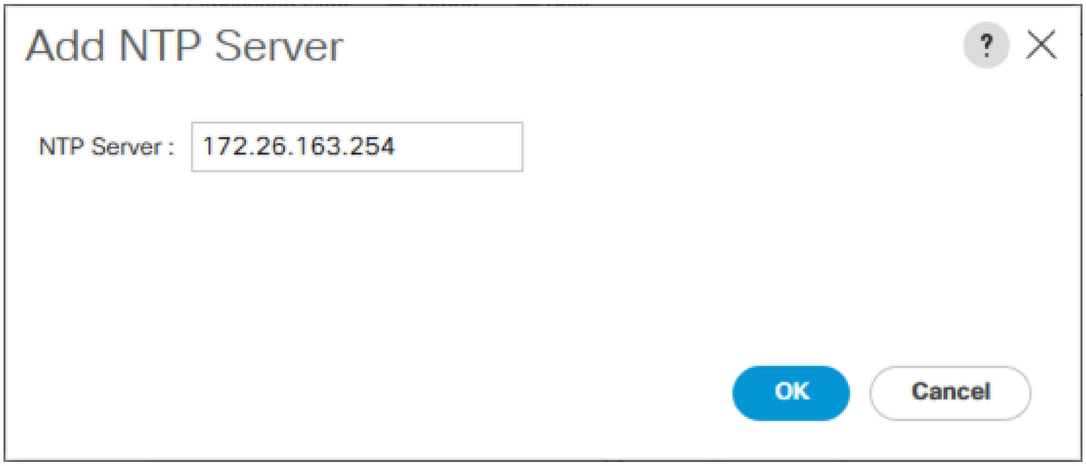

Add NTP Server

NTP Server : 172.26.163.254

8. Click OK to close the window.

Add Additional DNS Server(s)

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

1. In Cisco UCS Manager, click Admin.

2. Expand All > Communications Management.

3. Choose DNS Management.

4. In the Properties pane, choose Specify DNS Server.

5. Enter the IP address of the additional DNS server.

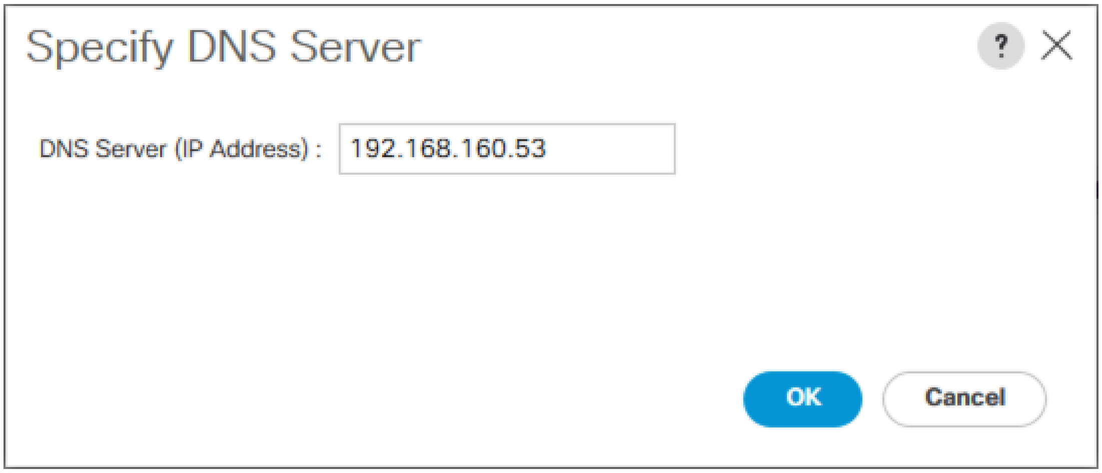
Specify DNS Server

DNS Server (IP Address) : 192.168.160.53

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

Add an Additional Administrative User

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

1. In Cisco UCS Manager, click Admin.

2. Expand User Management > User Services > Locally Authenticated Users.

3. Right-click Locally Authenticated Users and choose Create User.

4. In the Create User fields it is only necessary to fill in the Login ID, Password, and Confirm Password fields. Fill in the Create User fields according to your local security policy.

5. Leave the Account Status field set to Active.

6. Set Account Expires according to your local security policy.

7. Under Roles, choose admin.

8. Leave Password Required selected for the SSH Type field.

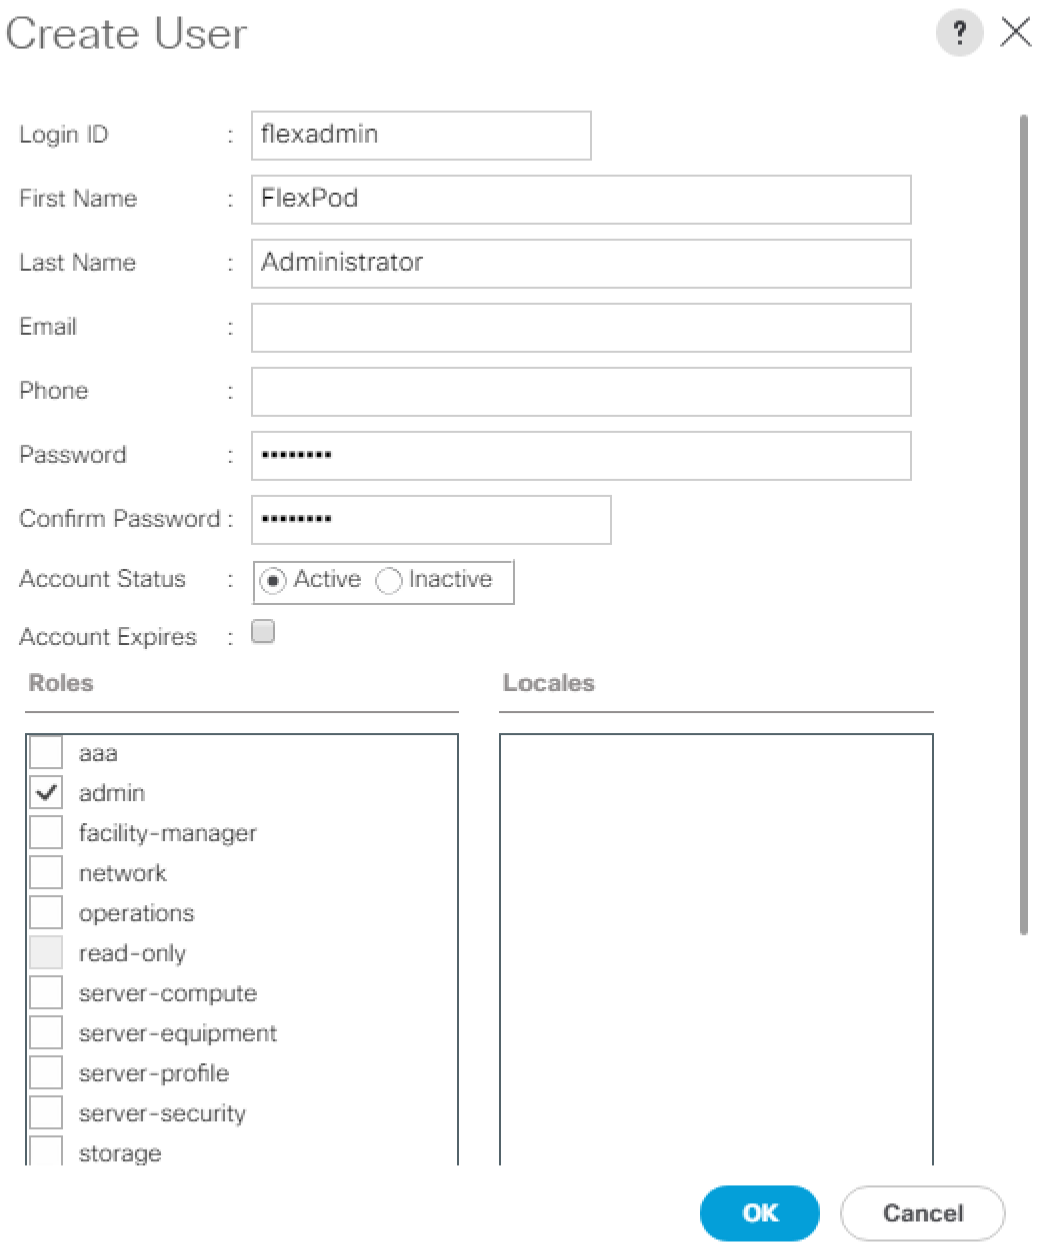

Create User

Login ID : flexadmin

First Name : FlexPod

Last Name : Administrator

Email :

Phone :

Password : •••••••••

Confirm Password : •••••••••

Account Status : ● Active ○ Inactive

Account Expires : □

Roles

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

Locales

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

Edit Global Policies

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

1. In Cisco UCS Manager, click Equipment and choose the Policies tab, and select the Global Policies sub-tab.

2. Set the Chassis/FEX Discovery Policy to match the minimum number of ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.

If varying numbers of links between chassis and the Fabric Interconnects will be used, set Action to 2 Link, the minimum recommended number of links for a FlexPod.

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

3. Scroll down to Info Policy and choose Enabled for Action.

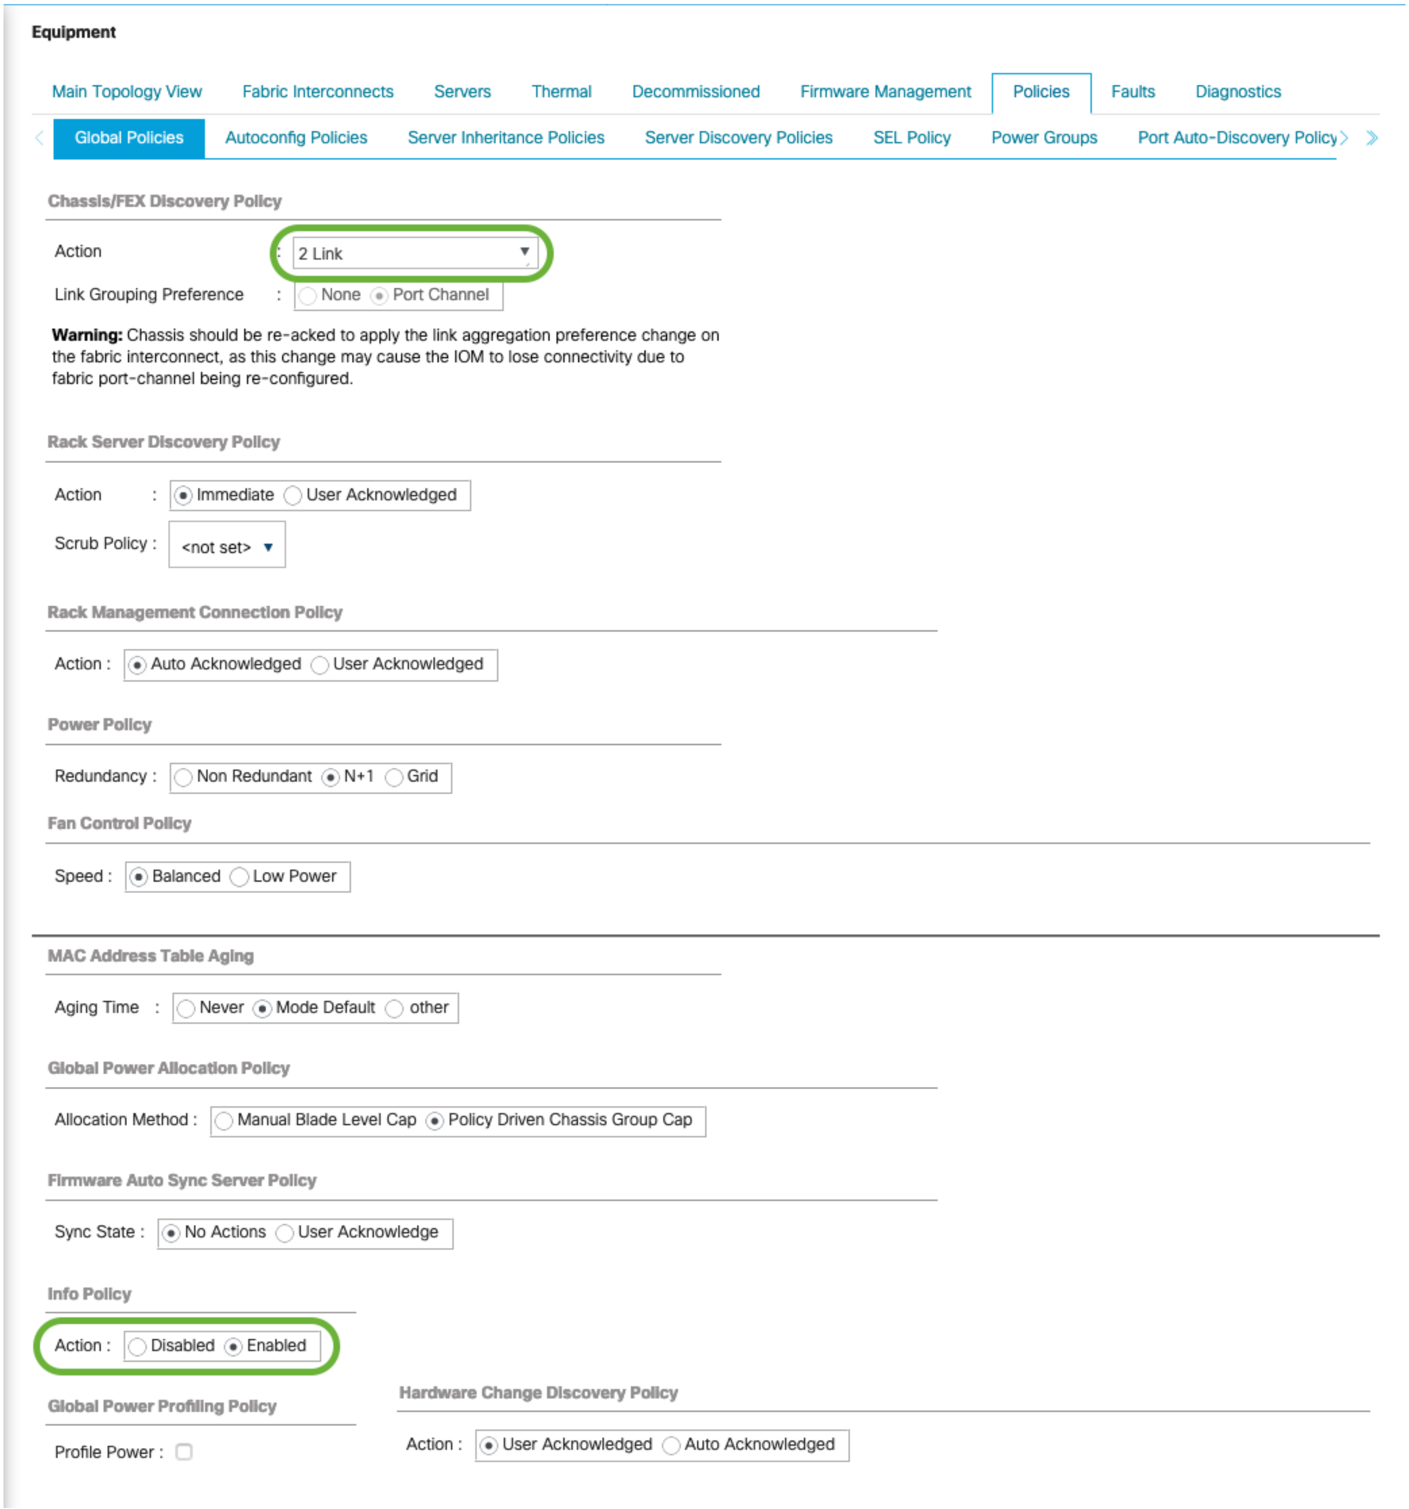

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : ○ None ● Port Channel

Warning: Chassis should be re-acked to apply the link aggregation preference change on the fabric interconnect, as this change may cause the IOM to lose connectivity due to fabric port-channel being re-configured.

Rack Server Discovery Policy

Action : ● Immediate ○ User Acknowledged

Scrub Policy : <not set>

Rack Management Connection Policy

Action : ● Auto Acknowledged ○ User Acknowledged

Power Policy

Redundancy : ○ Non Redundant ● N+1 ○ Grid

Fan Control Policy

Speed : ● Balanced ○ Low Power

MAC Address Table Aging

Aging Time : ○ Never ● Mode Default ○ other

Global Power Allocation Policy

Allocation Method : ○ Manual Blade Level Cap ● Policy Driven Chassis Group Cap

Firmware Auto Sync Server Policy

Sync State : ● No Actions ○ User Acknowledge

Info Policy

Action : ○ Disabled ● Enabled

Global Power Profiling Policy

Profile Power : □

Hardware Change Discovery Policy

Action : ● User Acknowledged ○ Auto Acknowledged

4. Click Save Changes and then click OK.

Enable Port Auto-Discovery Policy

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

1. In Cisco UCS Manager, click Equipment, choose All > Equipment in the Navigation Pane, and choose the Policies tab.

2. Under Port Auto-Discovery Policy, set Auto Configure Server Port to Enabled.

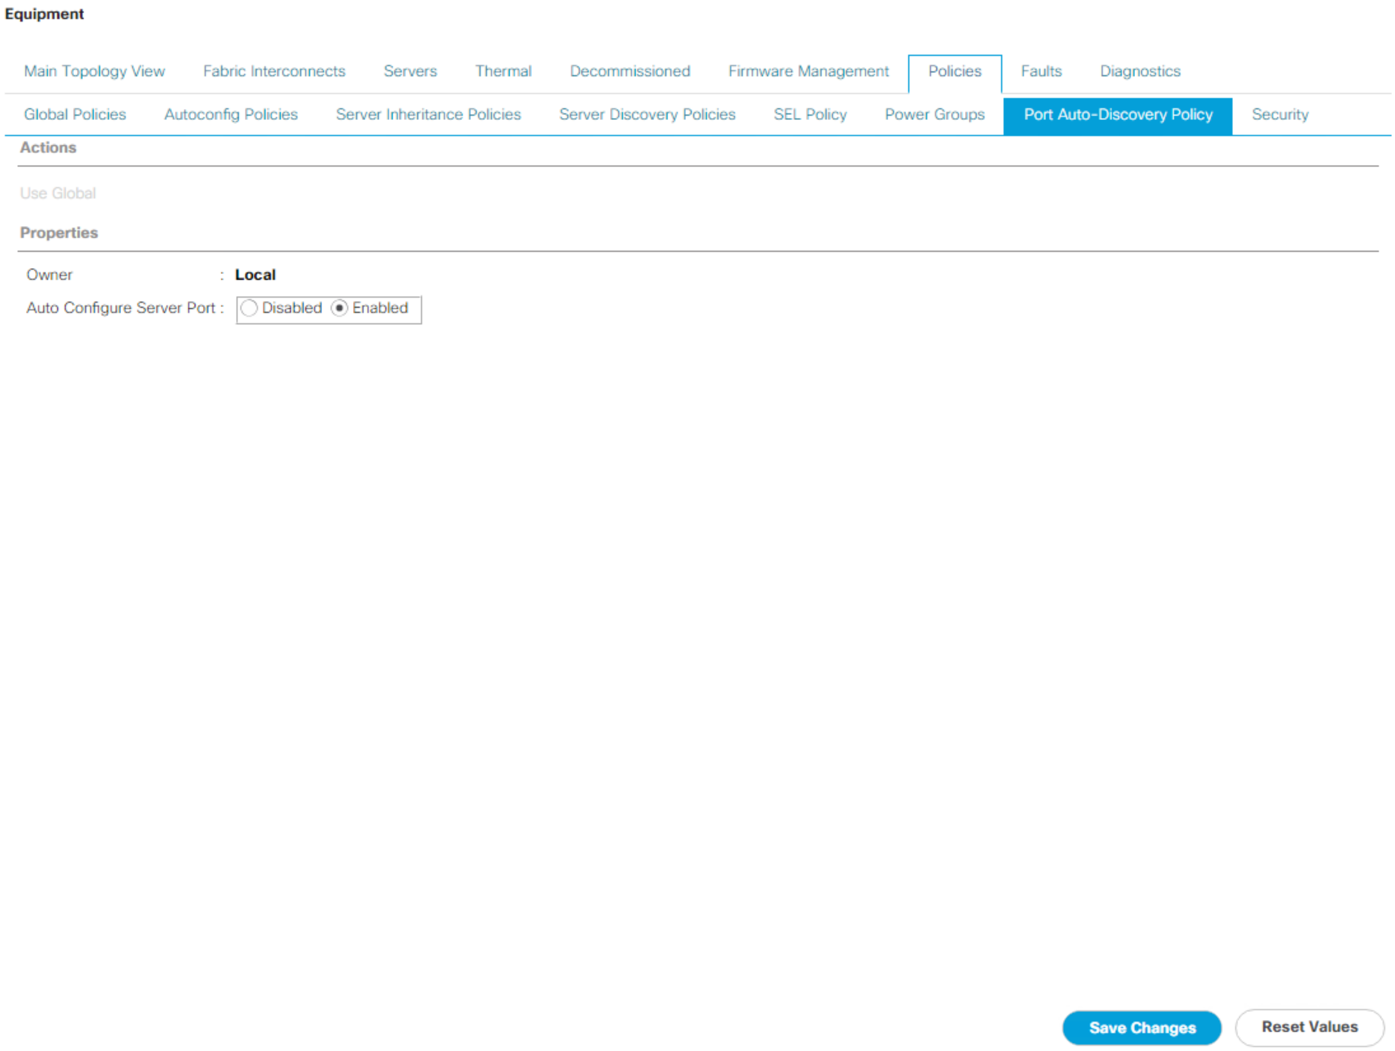

Owner : **Local**

Auto Configure Server Port : ○ Disabled ● Enabled

3. Click Save Changes and then click OK.

Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, click Equipment.

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

3. Expand and choose Ethernet Ports.

4. Verify that all ports connected to UCS chassis and rack mounts are configured as Server ports and have a status of Up.

5. If any rack mount ports are missing, choose the ports that are connected to Cisco FEXes and direct connect Cisco UCS C-Series servers, right-click them, and choose Configure as Server Port.

6. Click Yes to confirm server ports and click OK.

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

8. Choose the ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.

9. Click Yes to confirm uplink ports and click OK.

10. Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.

11. Expand and choose Ethernet Ports.

12. Verify that all ports connected to UCS chassis and rack mounts are configured as Server ports and have a status of Up.

13. If any rack mount ports are missing, choose the ports that are connected to Cisco FEXes and direct connect C-series servers, right-click them, and choose Configure as Server Port.

14. Click Yes to confirm server ports and click OK.

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

16. Choose the ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.

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

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external FEX modules, follow these steps:

1. In Cisco UCS Manager, click Equipment.

2. Expand Chassis and choose each chassis that is listed.

3. Right-click each chassis and choose Acknowledge Chassis.

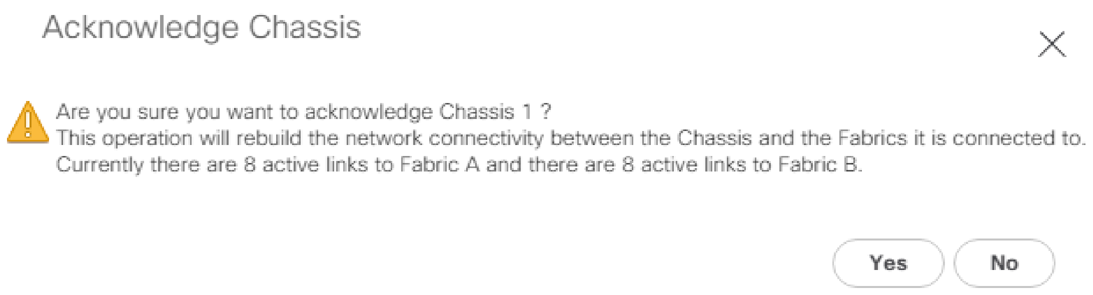
Acknowledge Chassis

Are you sure you want to acknowledge Chassis 1 ?
This operation will rebuild the network connectivity between the Chassis and the Fabrics it is connected to.
Currently there are 8 active links to Fabric A and there are 8 active links to Fabric B.

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

5. If Nexus FEXes are part of the configuration, expand Rack Mounts and FEX.

6. Right-click each FEX that is listed and choose Acknowledge FEX.

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

Create an Organization

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

1. In Cisco UCS Manager, click Servers.

2. In the Navigation Pane, expand Servers > Service Profiles.

3. Right-click root under Service Profiles and choose Create Organization.

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

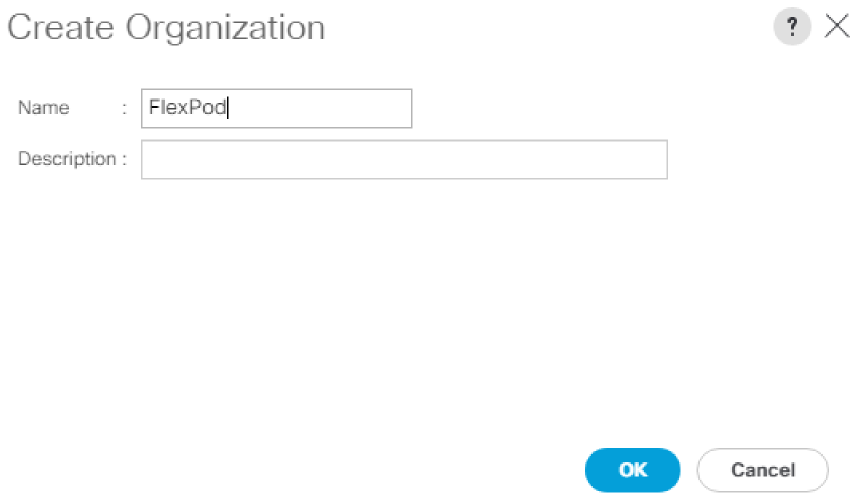
Create Organization

Name : FlexPod

Description :

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

Add Block of IP Addresses for KVM Access

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

1. In Cisco UCS Manager, click LAN.

2. Expand Pools > root > IP Pools.

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

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

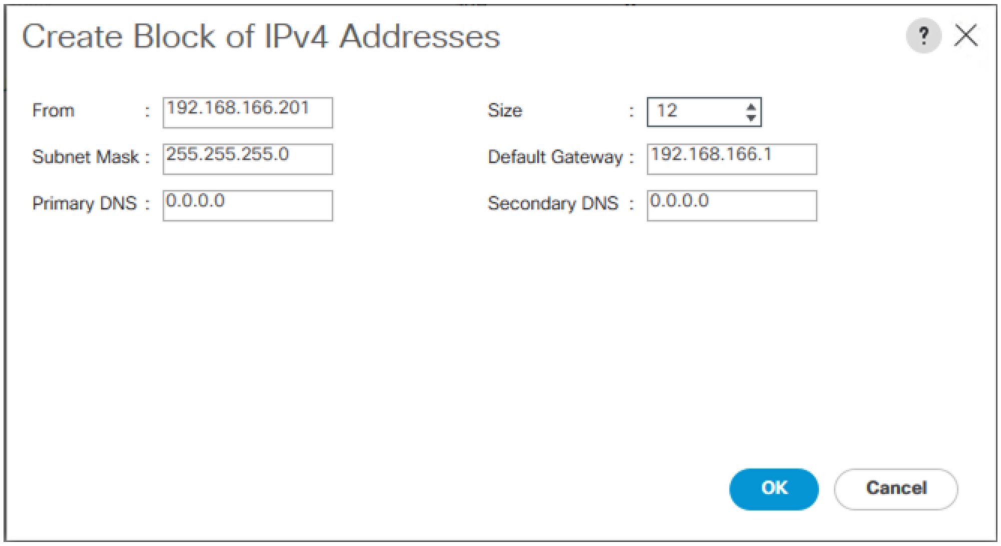

Create Block of IPv4 Addresses

From : 192.168.166.201

Size : 12

Subnet Mask : 255.255.255.0

Default Gateway : 192.168.166.1

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.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools for iSCSI boot for the Cisco UCS environment, follow these steps:

The IP Pools for iSCSI Boot are created here in the root organization, assuming that all UCS servers will be booted from the NetApp Infrastructure SVM. If servers will be booted from tenant SVMs with UCS tenant organizations, consider creating the IP Pools for iSCSI Boot in the tenant organization.

1. In Cisco UCS Manager, click LAN.

2. Expand Pools > root.

3. Right-click IP Pools.

4. Choose Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Choose Sequential for the assignment order.
8. Click Next.
9. Click Add to add a block of IP addresses.
10. In the From field, enter the beginning of the range to assign as-iSCSi boot IP addresses on Fabric A.
11. Set the size to enough addresses to accommodate the servers.
12. Enter the appropriate Subnet Mask.
13. Click OK.
14. Click Next.
15. Click Finish and OK to complete creating the Fabric A iSCSI IP Pool.
16. Right-click IP Pools.
17. Choose Create IP Pool.
18. Enter iSCSI-IP-Pool-B as the name of IP pool.
19. Optional: Enter a description for the IP pool.
20. Choose Sequential for the assignment order.
21. Click Next.
22. Click Add to add a block of IP addresses.
23. In the From field, enter the beginning of the range to assign as-iSCSi IP addresses on Fabric B.
24. Set the size to enough addresses to accommodate the servers.
25. Enter the appropriate Subnet Mask.
26. Click OK.
27. Click Next.
28. Click Finish and OK to complete creating the Fabric B iSCSI IP Pool.

Create Uplink Port Channels to Cisco Nexus Switches

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

1. In Cisco UCS Manager, click LAN.

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

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

3. Right-click Port Channels under Fabric A.

4. Choose Create Port Channel.

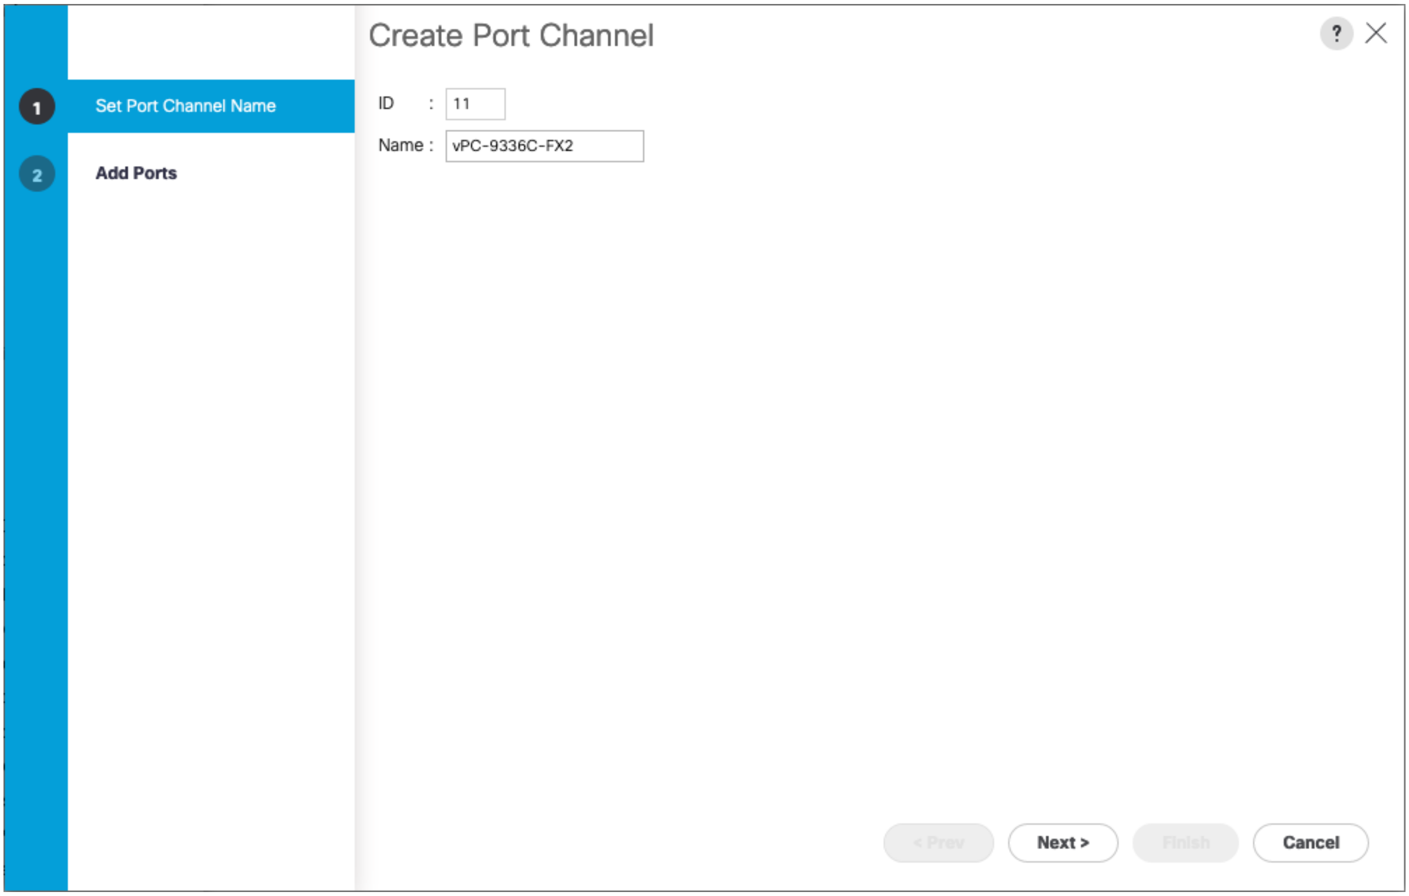

Create Port Channel

ID : 11

Name : vPC-9336C-FX2

5. Enter `11` as the unique ID of the port channel.

> The Port Channel IDs in this example correspond to the first ports of upstream interface members of the Nexus leafs implementing the vPC, where 11 represents 1/1 on the switch. This is optional but can be helpful in correlating the Port Channel to the vPC.

6. Enter `vPC-9336C-FX2` as the name of the port channel.

7. Click Next.

8. Choose the uplink ports connected to the Nexus switches to be added to the port channel.

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

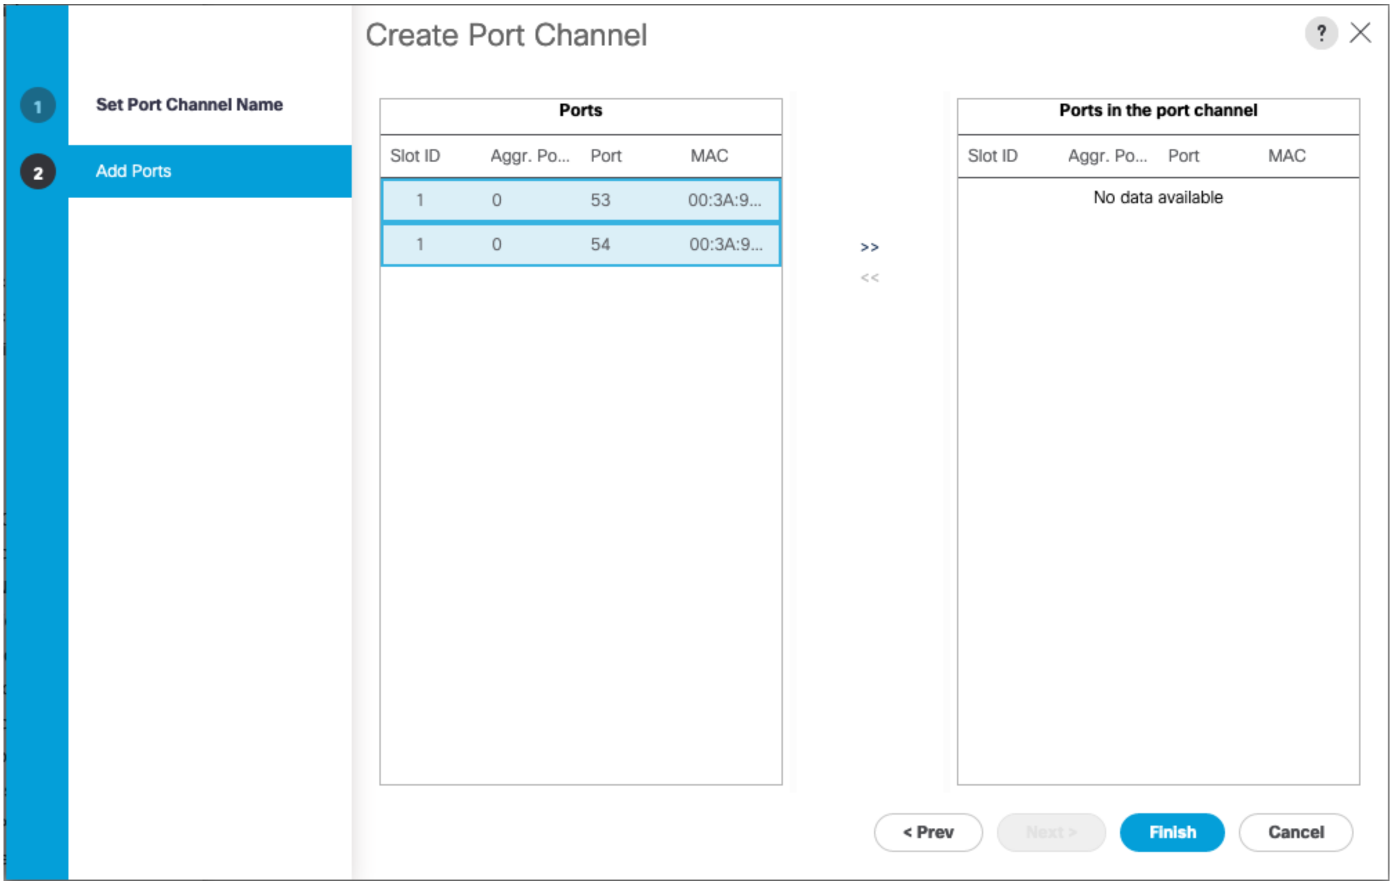

Ports

Slot ID	Aggr. Po...	Port	MAC
1	0	53	00:3A:9...
1	0	54	00:3A:9...

Ports in the port channel

Slot ID	Aggr. Po...	Port	MAC

No data available

10. Click Finish to create the port channel.

11. Click OK.

12. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, choose Port-Channel 11. Ensure Auto is selected for the Admin Speed. After a few minutes, verify that the Overall Status is Up, and the Operational Speed is correct.

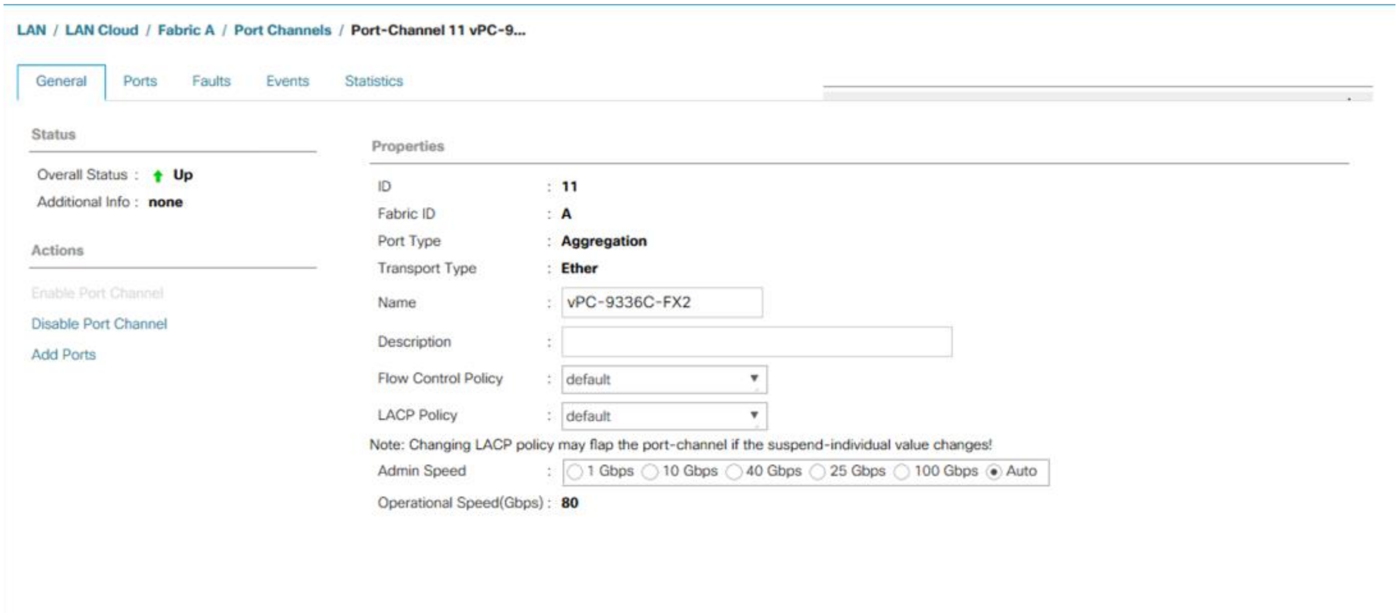

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 11 vPC-9...

Status

Overall Status : Up

Additional Info : none

Actions

Properties

ID : 11

Fabric ID : A

Port Type : Aggregation

Transport Type : Ether

Name : vPC-9336C-FX2

Description :

Flow Control Policy : default

LACP Policy : default

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

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

Operational Speed(Gbps) : 80

13. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.

14. Right-click Port Channels under Fabric B.

15. Choose Create Port Channel.

16. Enter `12` as the unique ID of the port channel.

17. Enter `vPC-9336C-FX2` as the name of the port channel.

18. Click Next.

19. Choose the ports connected to the Nexus switches to be added to the port channel:

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

21. Click Finish to create the port channel.

22. Click OK.

23. In the navigation pane, under LAN > LAN Cloud > Fabric B > Port Channels, choose Port-Channel 12. Ensure Auto is selected for the Admin Speed. After a few minutes, verify that the Overall Status is Up, and the Operational Speed is correct.

Add UDLD to Uplink Port Channels

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

1. In Cisco UCS Manager, click LAN.

2. Expand Policies > LAN Cloud > UDLD Link Policy.

3. Right-click UDLD Link Policy and choose Create UDLD Link Policy.

4. Name the Policy UDLD-Normal and choose Enabled for the Admin State and Normal for the Mode.

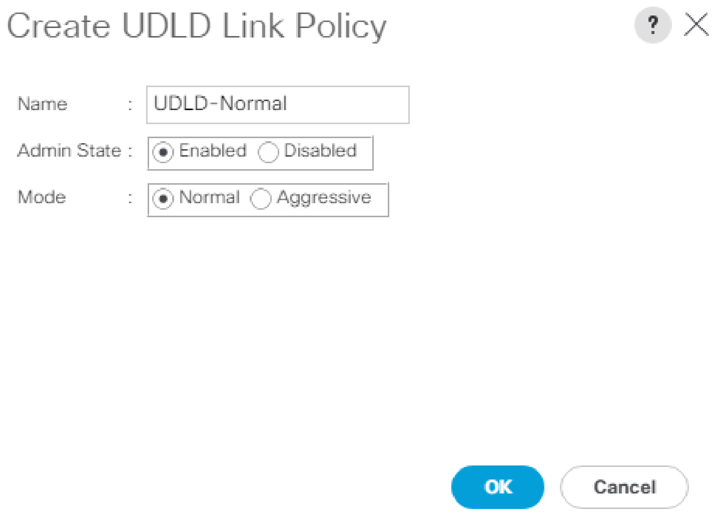

Create UDLD Link Policy

Name : UDLD-Normal

Admin State : ● Enabled ○ Disabled

Mode : ● Normal ○ Aggressive

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

6. Expand Policies > LAN Cloud > Link Profile.

7. Right-click Link Profile and choose Create Link Profile.

8. Name the Profile UDLD-Normal and choose the UDLD-Normal Link Policy created above.

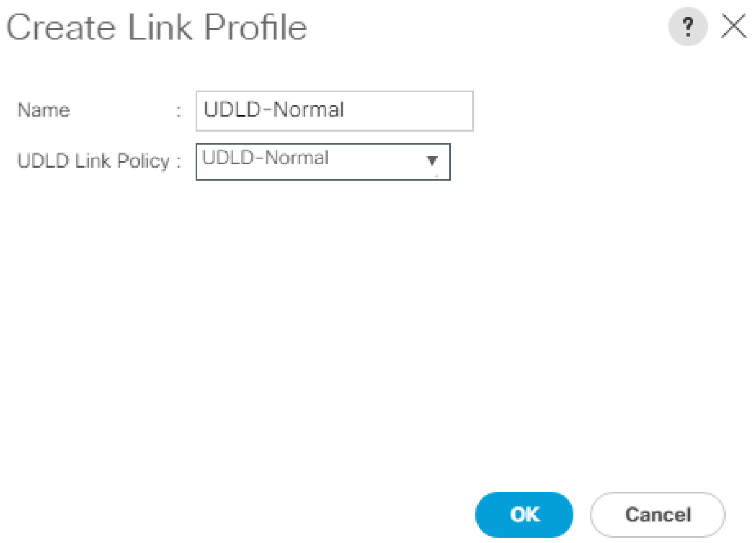

Create Link Profile

Name : UDLD-Normal

UDLD Link Policy : UDLD-Normal

9. Click OK, then click OK again to complete creating the profile.

10. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 11. Choose the first Eth Interface under Port-Channel 11. From the drop-down list, choose the UDLD-Normal Link Profile created above, click Save Changes and OK. Repeat this process for each Eth Interface under Port-Channel 11 and for each Eth Interface under Port-Channel 12 on Fabric B.

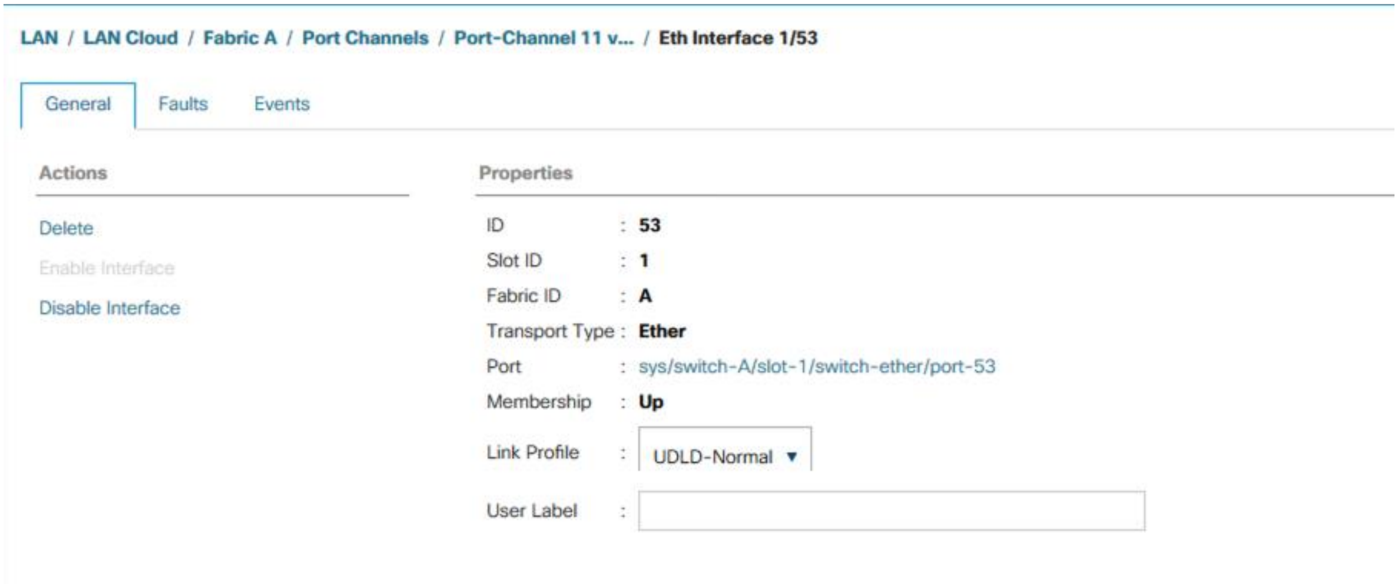
LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 11 v... / Eth Interface 1/53
General
Faults
Events
Actions
Delete
Enable Interface
Disable Interface
Properties

ID : 53
Slot ID : 1
Fabric ID : A
Transport Type : Ether
Port : sys/switch-A/slot-1/switch-ether/port-53
Membership : Up
Link Profile : UDLD-Normal ▾
User Label :

Set Jumbo Frames in Cisco UCS Fabric

Jumbo Frames are used in FlexPod for the NFS and iSCSI storage protocols. The typical best practice in FlexPod has been to set the MTU of the Best Effort QoS System Class in Cisco UCS Manager to 9216 for Jumbo Frames. In the Cisco UCS 6454 Fabric Interconnect with Cisco UCS Manager version 4.0 software the MTU for the Best Effort QoS System Class is fixed at normal and cannot be changed. With this setting of normal in the 6454, Jumbo Frames can pass through the Cisco UCS fabric without being dropped. In Cisco UCS Manager version 4.1, the MTU for the Best Effort QoS System Class is again settable. To configure jumbo frames in the Cisco UCS fabric, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand LAN > LAN Cloud > QoS System Class.

3. In the right pane, click the General tab.

4. On the Best Effort row, enter 9216 in the box under the MTU column.

5. Click Save Changes.

6. Click OK.

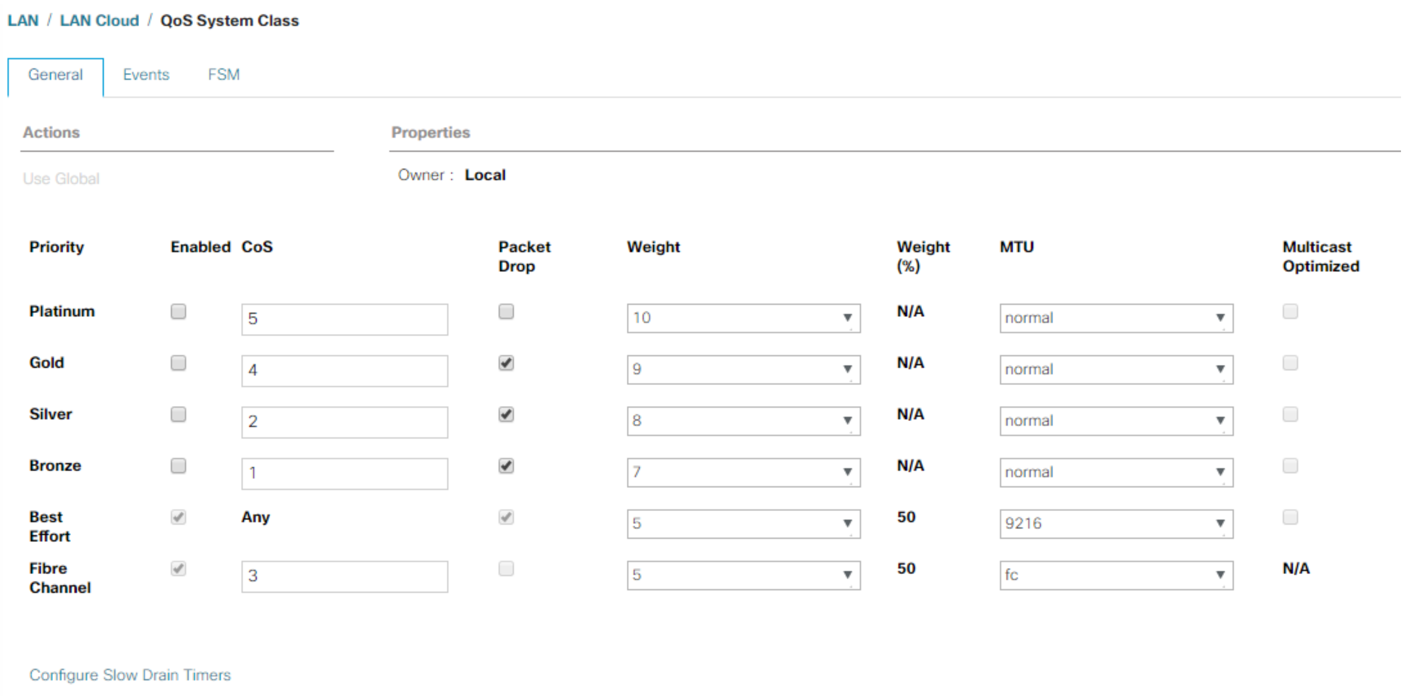

LAN / LAN Cloud / QoS System Class

Owner : Local

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

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

Create VLANs

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

1. In Cisco UCS Manager, click LAN.

In this procedure, five unique VLANs are created. See Table2.

2. Expand LAN > LAN Cloud.

3. Right-click VLANs.

4. Choose Create VLANs.

5. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.

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

7. Enter the native VLAN ID.

8. Keep the Sharing Type as None.

9. Click OK and then click OK again.

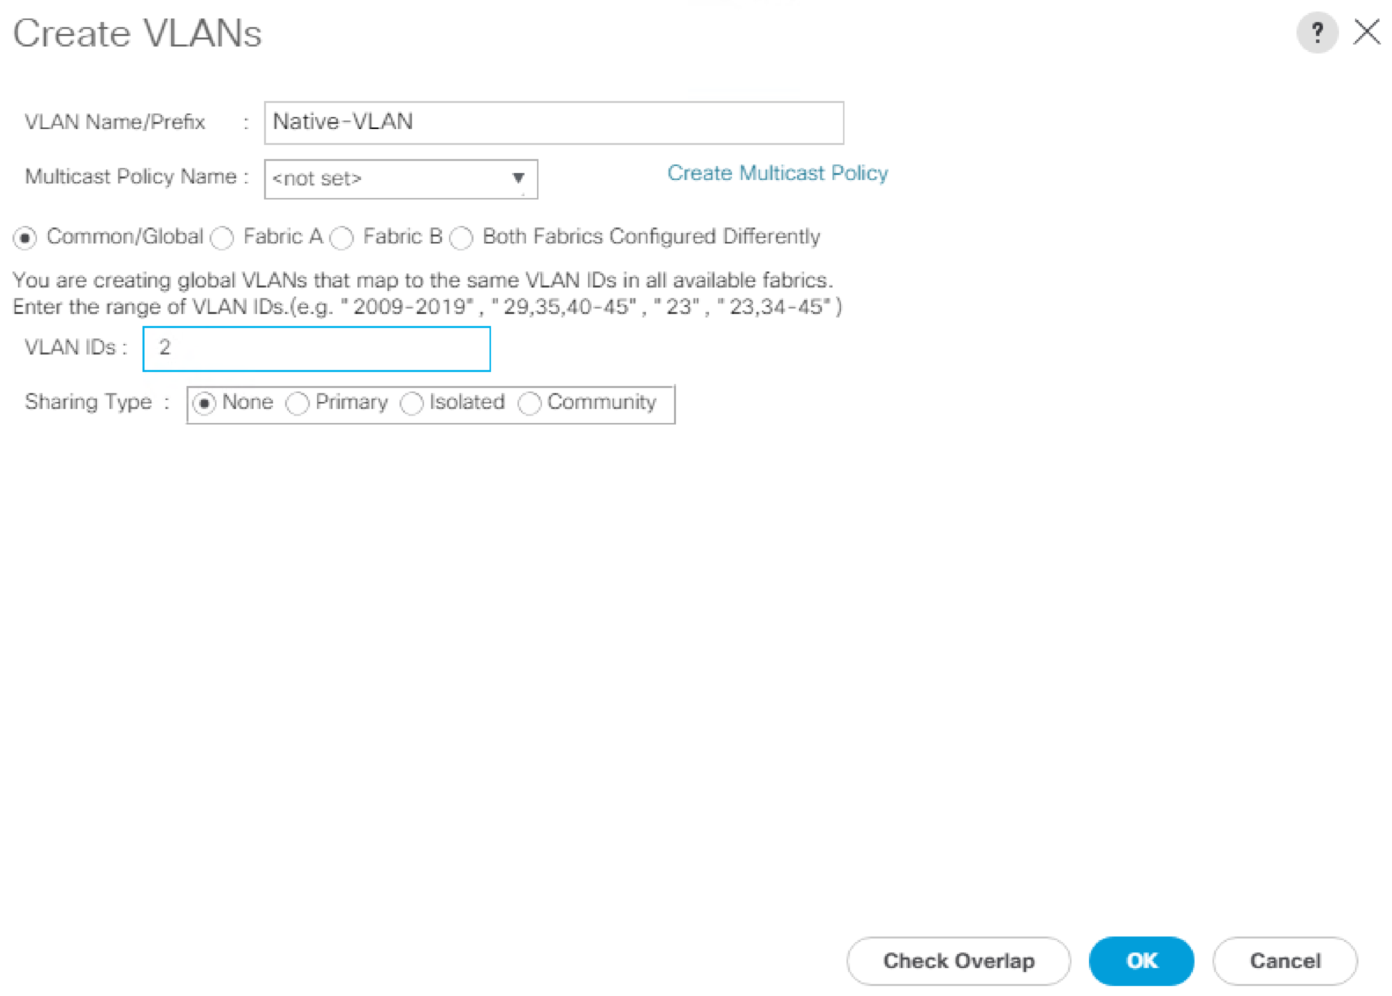

Create VLANs

VLAN Name/Prefix : 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 choose Set as Native VLAN.

11. Click Yes and then click OK.

12. Right-click VLANs.

13. Choose Create VLANs

14. Enter Site1-IB as the name of the VLAN to be used for management traffic (ESXi Hosts, site specific infrastructure).

Modify these VLAN names as necessary for your environment.

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

16. Enter the In-Band management VLAN ID.

17. Keep the Sharing Type as None.

18. Click OK, and then click OK again.

19. Right-click VLANs.

20. Choose Create VLANs.

21. Enter Common-IB as the name of the VLAN to be used for Common internal infrastructure (that may be relevant to application networks or differing locations)

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

23. Enter the Common Infrastructure VLAN ID.

24. Keep the Sharing Type as None.

25. Click OK, and then click OK again.

26. Right-click VLANs.

27. Choose Create VLANs.

28. Enter iSCSI-A as the name of the VLAN to be used for iSCSI-A traffic.

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

30. Enter the Infrastructure iSCSI-A VLAN ID.

31. Keep the Sharing Type as None.

32. Click OK, and then click OK again.

33. Right-click VLANs.

34. Choose Create VLANs.

35. Enter iSCSI-B as the name of the VLAN to be used for iSCSI-B traffic.

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

37. Enter the Infrastructure iSCSI-B VLAN ID.

38. Keep the Sharing Type as None.

39. Click OK, and then click OK again.

40. Right-click VLANs.

41. Choose Create VLANs.

42. Enter Infra-NFS as the name of the VLAN to be used for NFS.

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

44. Enter the Infrastructure NFS VLAN ID.

45. Keep the Sharing Type as None.

46. Click OK, and then click OK again.

47. Right-click VLANs.

48. Choose Create VLANs.

49. Enter vMotion as the name of the VLAN to be used for vMotion.

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

51. Enter the vMotion VLAN ID.

52. Keep the Sharing Type as None.

53. Click OK and then click OK again.

54. Choose Create VLANs.

55. Enter VM-Traffic as the name of the VLAN to be used for VM Traffic, or optionally add a "-" to use as a prefix for multiple VLANs created.

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

57. Enter the VM-Traffic VLAN ID, or a range to use if creating multiple VLANs.

58. Keep the Sharing Type as None.

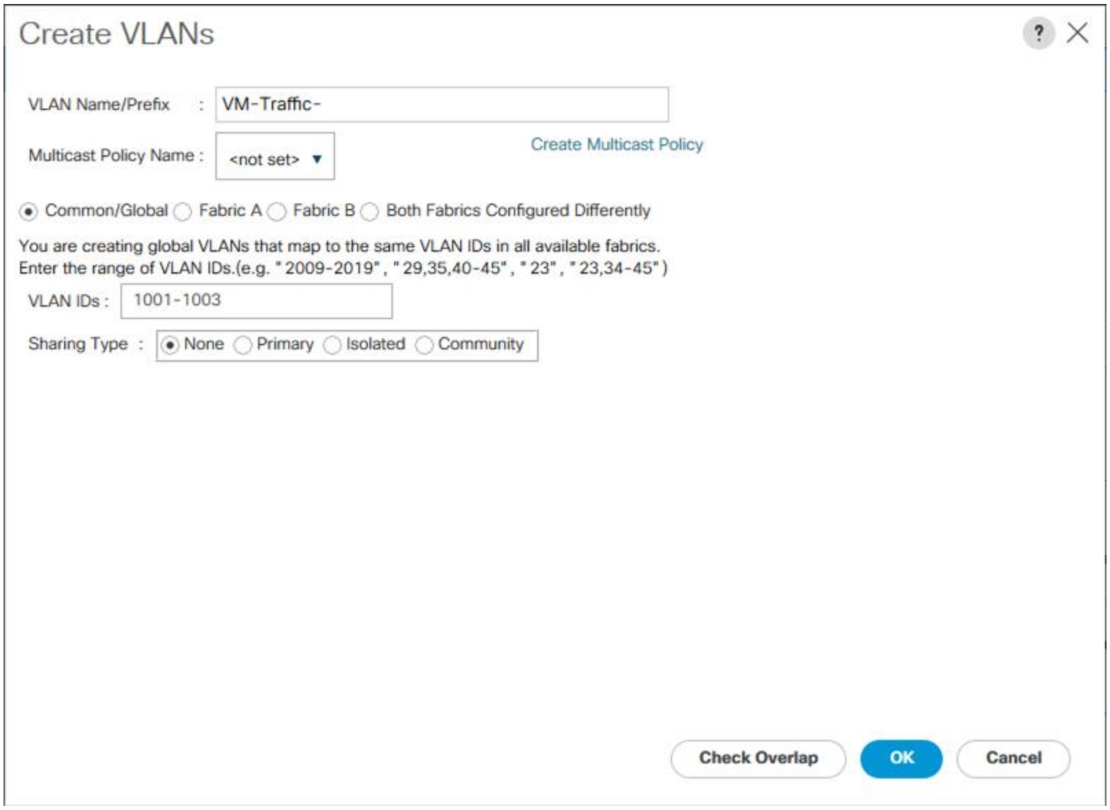

Create VLANs

VLAN Name/Prefix : VM-Traffic-

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 : 1001-1003

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

59. Click OK and then click OK again.

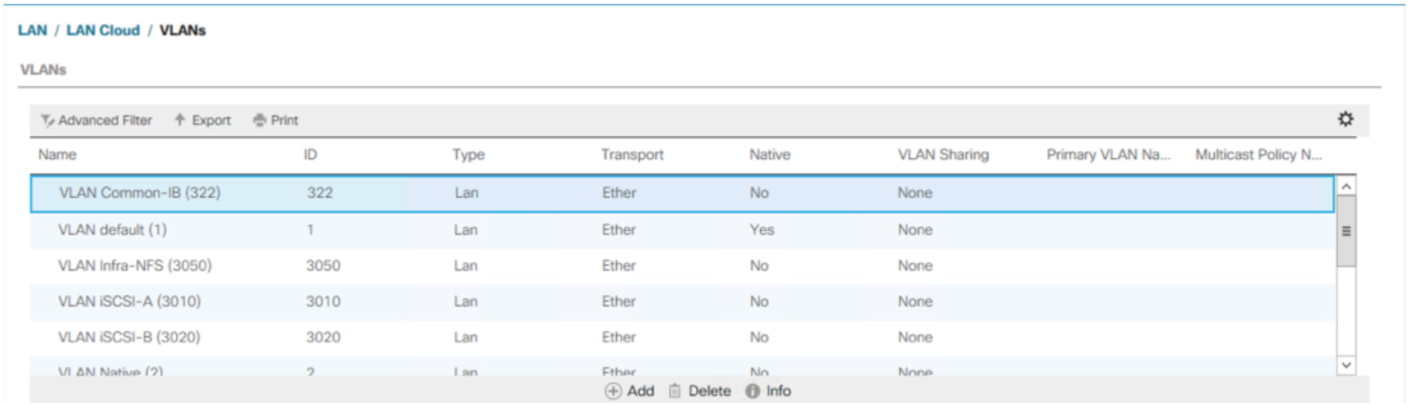

LAN / LAN Cloud / VLANs

VLANs

Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Na...	Multicast Policy N...
VLAN Common-IB (322)	322	Lan	Ether	No	None		
VLAN default (1)	1	Lan	Ether	Yes	None		
VLAN Infra-NFS (3050)	3050	Lan	Ether	No	None		
VLAN iSCSI-A (3010)	3010	Lan	Ether	No	None		
VLAN iSCSI-B (3020)	3020	Lan	Ether	No	None		
VLAN Native (2)	2	Lan	Ether	No	None		

Create MAC Address Pools

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

1. In Cisco UCS Manager, click LAN.

2. Expand Pools > root.

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

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

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

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

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

7. Choose Sequential as the option for Assignment Order.

8. Click Next.

9. Click Add.

10. Specify a starting MAC address.

For the FlexPod solution, the recommendation is to place A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the example of also embedding the cabinet number information giving us `00:25:B5:A0:1A:00` as our first MAC address.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created. In this example, with the MAC block modification, a maximum of 256 addresses are available.

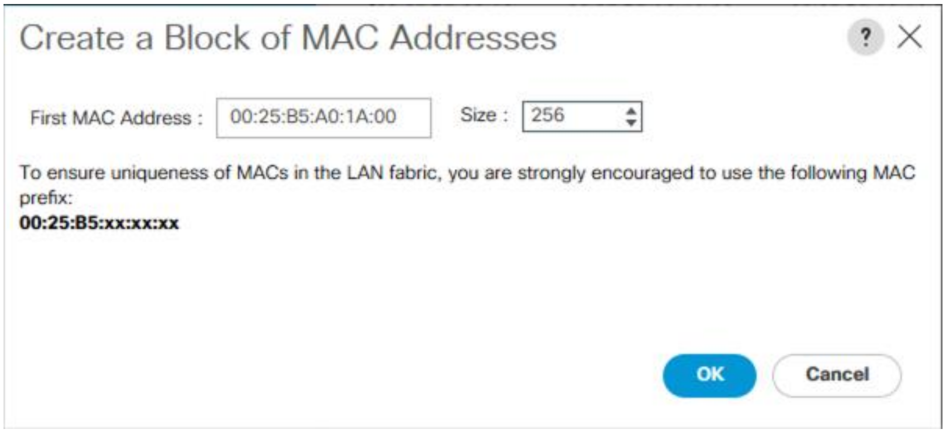

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:A0:1A:00

Size : 256

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

12. Click OK.

13. Click Finish.

14. In the confirmation message, click OK.

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

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

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

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

19. Choose Sequential as the option for Assignment Order.

20. Click Next.

21. Click Add.

22. Specify a starting MAC address.

> For the FlexPod solution, it is recommended to place B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward our example of also embedding the cabinet number information giving us `00:25:B5:A0:1B:00` as our first MAC address.

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created. In this example, with the MAC block modification, a maximum of 256 addresses are available.

24. Click OK.

25. Click Finish.

26. In the confirmation message, click OK.

Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)

To create a network control policy that enables CDP and LLDP on server virtual network controller (vNIC) ports, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand Policies > root.

3. Right-click Network Control Policies.

4. Choose Create Network Control Policy.

5. Enter Enable-CDP-LLDP as the policy name.

6. For CDP, choose the Enabled option.

7. For LLDP, scroll down and choose Enabled for both Transmit and Receive.

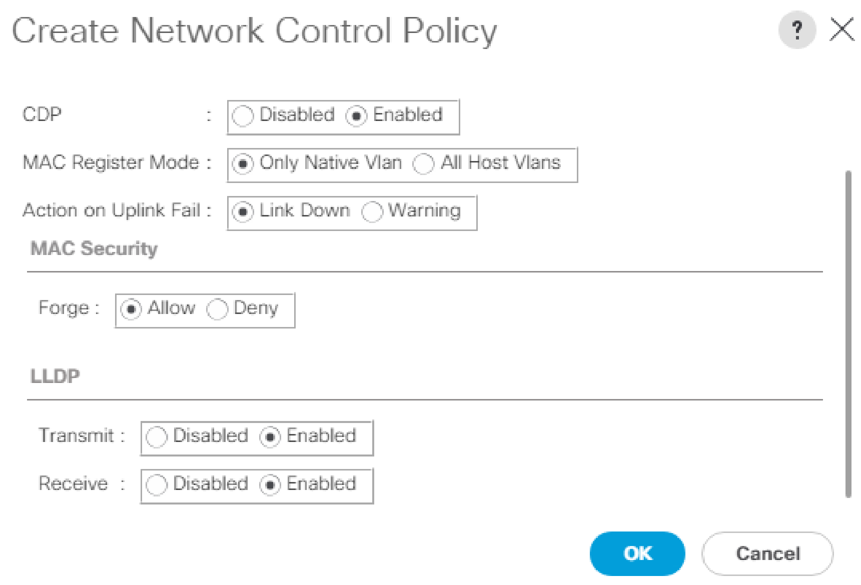

Create Network Control Policy

CDP : ○ Disabled ● Enabled

MAC Register Mode : ● Only Native Vlan ○ All Host Vlans

Action on Uplink Fail : ● Link Down ○ Warning

MAC Security

Forge : ● Allow ○ Deny

LLDP

Transmit : ○ Disabled ● Enabled

Receive : ○ Disabled ● Enabled

8. Click OK to create the network control policy.

9. Click OK.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates within the FlexPod organization, follow these steps. A total of 6 vNIC Templates will be created. Two of the vNIC templates (vSwitch0-A and vSwitch0-B) will be created for vNICs to connect to VMware ESXi vSwitch0. vSwitch0 will have port groups for the IB-MGMT, Infra-NFS, vMotion, and VM-Traffic VLANs. The third and fourth vNIC templates (vDS0-A and vDS0-B) will be created for vNICs to connect to the VMware Virtual Distributed Switch (vDS0). The fifth and sixth vNIC templates will be used for the iSCSI A and iSCSI B.

The vDS will have port groups for the vMotion and VM-Traffic VLANs. Having the vMotion VMkernel on the vDS will allow the QoS marking of vMotion packets to occur within the vDS if QoS policies need to be applied to vMotion in the future. Any tenant or application VLANs can be placed on the vDS in the future.

If QoS policy application to the vMotion packets will not be considered at some point, the vMotion VLAN can instead be placed on the vSwitch0-A and vSwitch0-B vNIC templates and handled as standard port groups in vSwitch0 alongside the rest of the infrastructure traffic.

To create the infrastructure vNIC templates, follow these steps:

1. In Cisco UCS Manager, click LAN.

2. Expand Policies > root > Sub-Organizations > FlexPod.

3. Under the FlexPod Organization, right-click vNIC Templates.

4. Choose Create vNIC Template.

5. Enter vSwitch0-A as the vNIC template name.

6. Keep Fabric A selected.

7. Do not select the Enable Failover checkbox.

8. Choose Primary Template for Redundancy Type.

9. Leave the Peer Redundancy Template set to <not set>.

10. Under Target, make sure that only the Adapter checkbox is selected.

11. Choose Updating Template as the Template Type.

12. Under VLANs, choose the checkboxes for Common-IB, Site1-IB, Infra-NFS, and Native-VLAN VLANs.

13. Set Native-VLAN as the native VLAN.

14. Leave vNIC Name selected for the CDN Source.

15. For MTU, enter 9000.

16. In the MAC Pool list, choose MAC-Pool-A.

17. In the Network Control Policy list, choose CDP-LLDP.

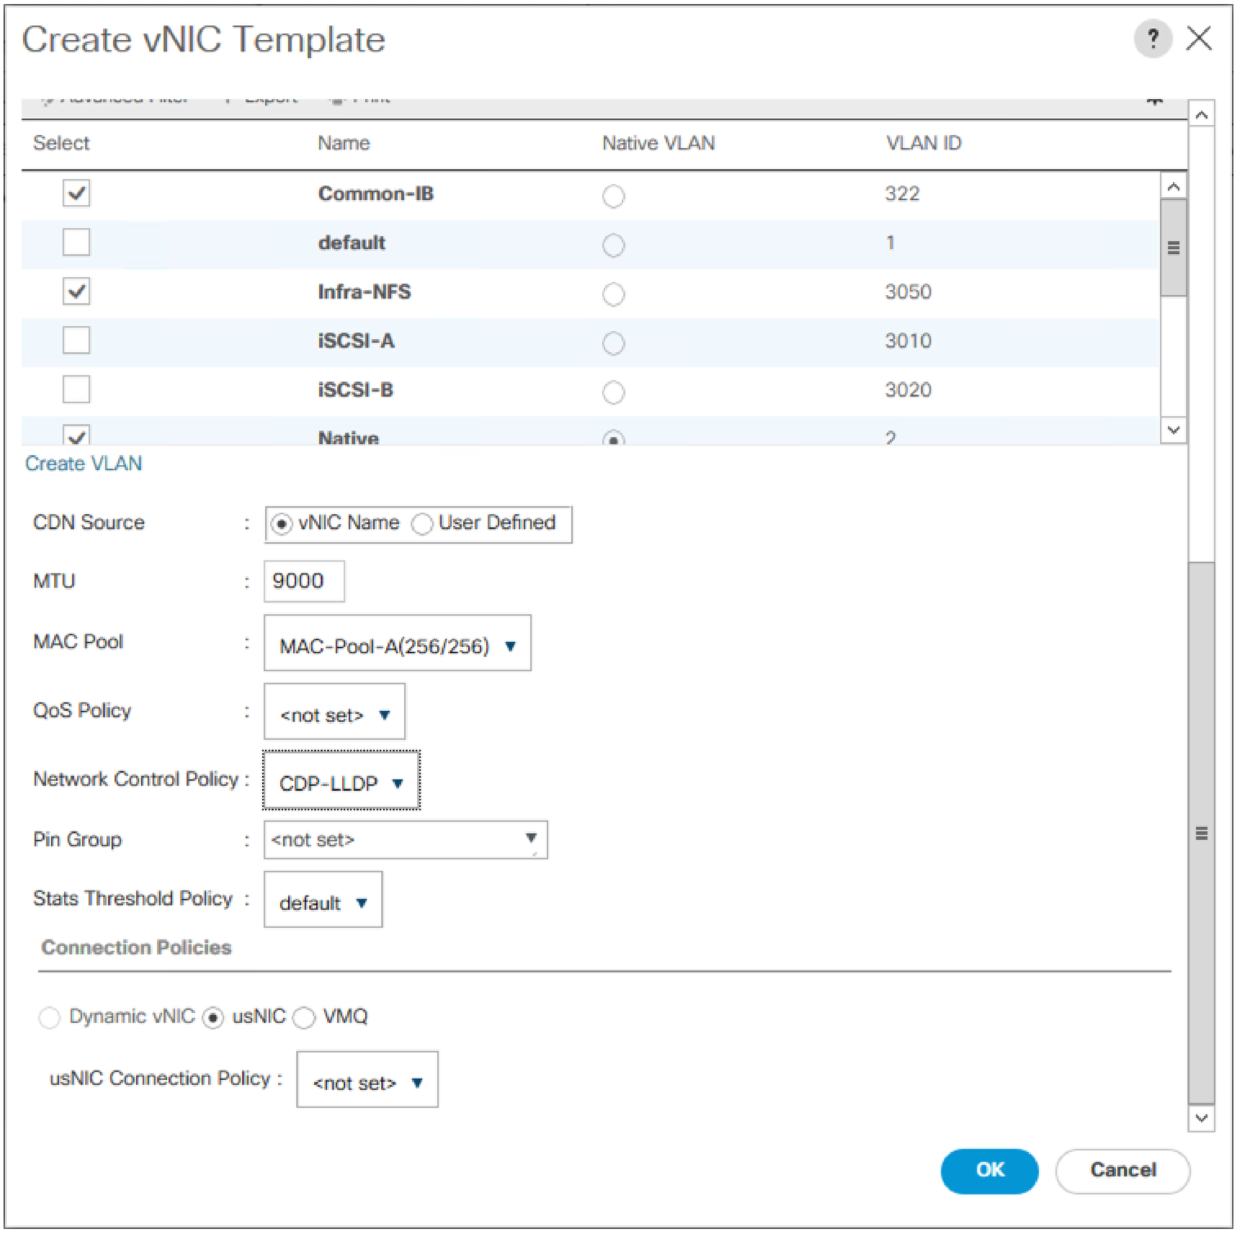

Select	Name	Native VLAN	VLAN ID
☑	Common-IB	○	322
□	default	○	1
☑	Infra-NFS	○	3050
□	iSCSI-A	○	3010
□	iSCSI-B	○	3020
☑	Native	●	2

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(256/256)

QoS Policy : <not set>

Network Control Policy : CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

○ Dynamic vNIC ● usNIC ○ VMQ

usNIC Connection Policy : <not set>

18. Click OK to create the vNIC template.

19. Click OK.

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

21. Choose Create vNIC Template.

22. Enter vSwitch0-B as the vNIC template name.

23. Choose Fabric B.

24. Do not select the Enable Failover checkbox.

25. Set Redundancy Type to Secondary Template.

26. Choose vSwitch0-A for the Peer Redundancy Template.

27. In the MAC Pool list, choose MAC-Pool-B.

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

28. Click OK to create the vNIC template.

29. Click OK.

30. Under the FlexPod Organization, right-click vNIC Templates.

31. Choose Create vNIC Template.

32. Enter vDS0-A as the vNIC template name.

33. Keep Fabric A selected.

34. Do not select the Enable Failover checkbox.

35. Choose Primary Template for Redundancy Type.

36. Leave the Peer Redundancy Template set to <not set>.

37. Under Target, make sure that only the Adapter checkbox is selected.

38. Choose Updating Template as the Template Type.

39. Under VLANs, choose the checkboxes for vMotion, any configured VM-Traffic VLANs, and the Native VLAN.

40. Set Native-VLAN as the native VLAN.

41. Leave vNIC Name selected for the CDN Source.

42. For MTU, enter 9000.

43. In the MAC Pool list, choose MAC-Pool-A.

44. In the Network Control Policy list, choose CDP-LLDP.

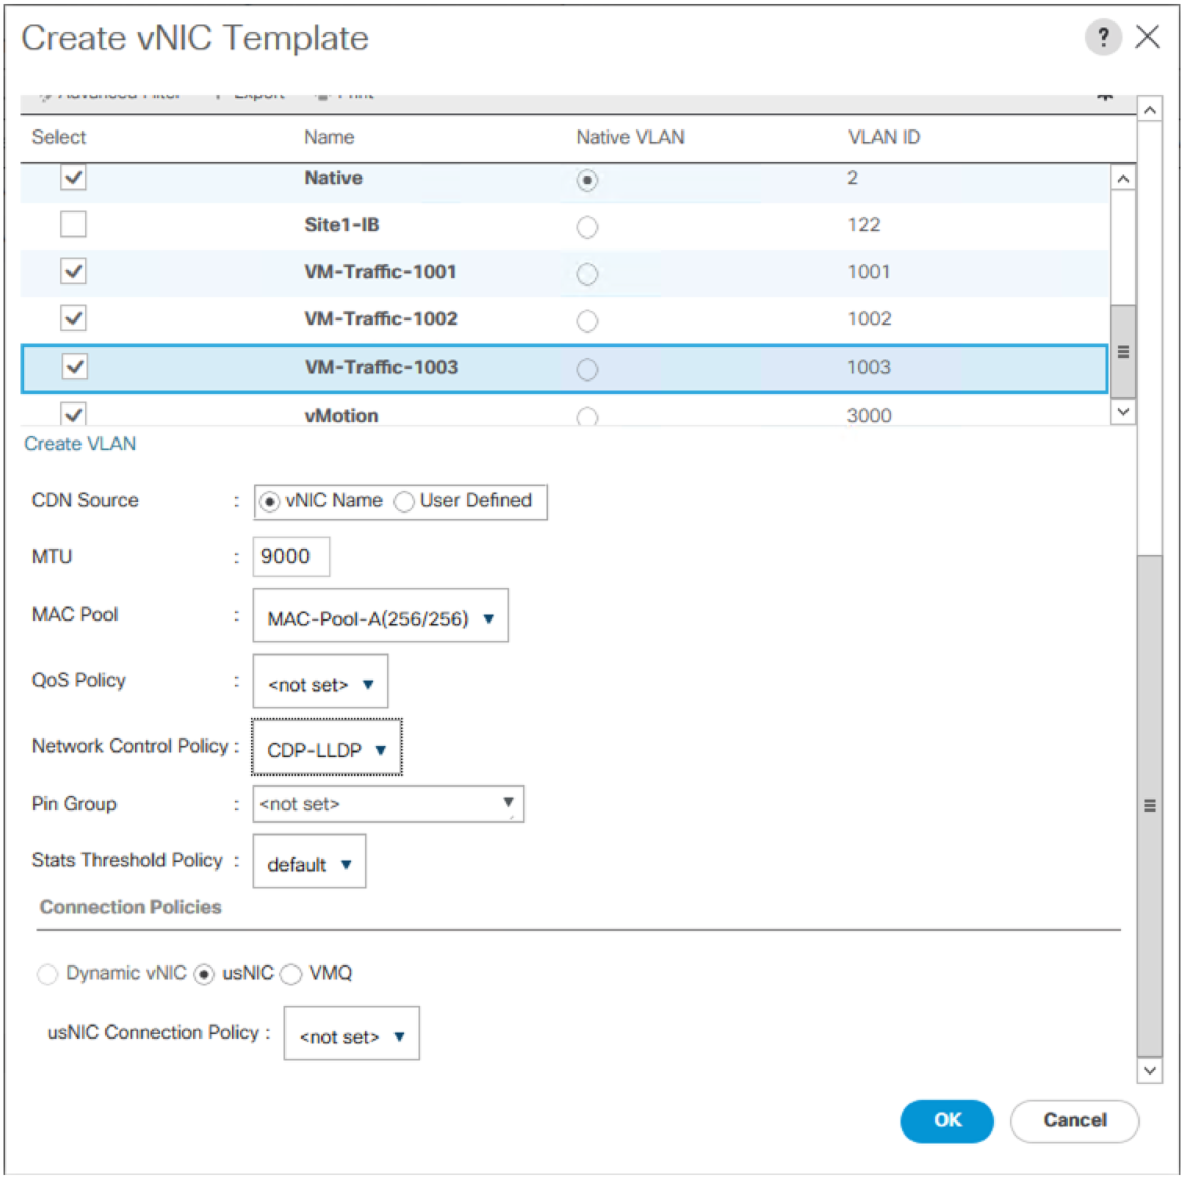

Create vNIC Template

Select	Name	Native VLAN	VLAN ID
☑	Native	●	2
□	Site1-IB	○	122
☑	VM-Traffic-1001	○	1001
☑	VM-Traffic-1002	○	1002
☑	VM-Traffic-1003	○	1003
☑	vMotion	○	3000

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(256/256)

QoS Policy : <not set>

Network Control Policy : CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

○ Dynamic vNIC ● usNIC ○ VMQ

usNIC Connection Policy : <not set>

45. Click OK to create the vNIC template.

46. Click OK.

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

48. Choose Create vNIC Template

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

50. Choose Fabric B.

51. Do not select the Enable Failover checkbox.

52. Set Redundancy Type to Secondary Template.

53. Choose vDS0-A for the Peer Redundancy Template.

54. In the MAC Pool list, choose MAC-Pool-B.

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

55. Click OK to create the vNIC template.

56. Click OK.

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

58. Choose Create vNIC Template

59. Enter iSCSI-A as the vNIC template name.

60. Choose Fabric A. Do not choose the Enable Failover checkbox.

61. Leave Redundancy Type set at No Redundancy.

62. Under Target, make sure that only the Adapter checkbox is selected.

63. Choose Updating Template for Template Type.

64. Under VLANs, choose only iSCSI-A.

65. Choose iSCSI-A as the native VLAN.

66. Leave vNIC Name set for the CDN Source.

67. Under MTU, enter 9000.

68. From the MAC Pool list, choose MAC-Pool-A.

69. From the Network Control Policy list, choose CDP-LLDP.

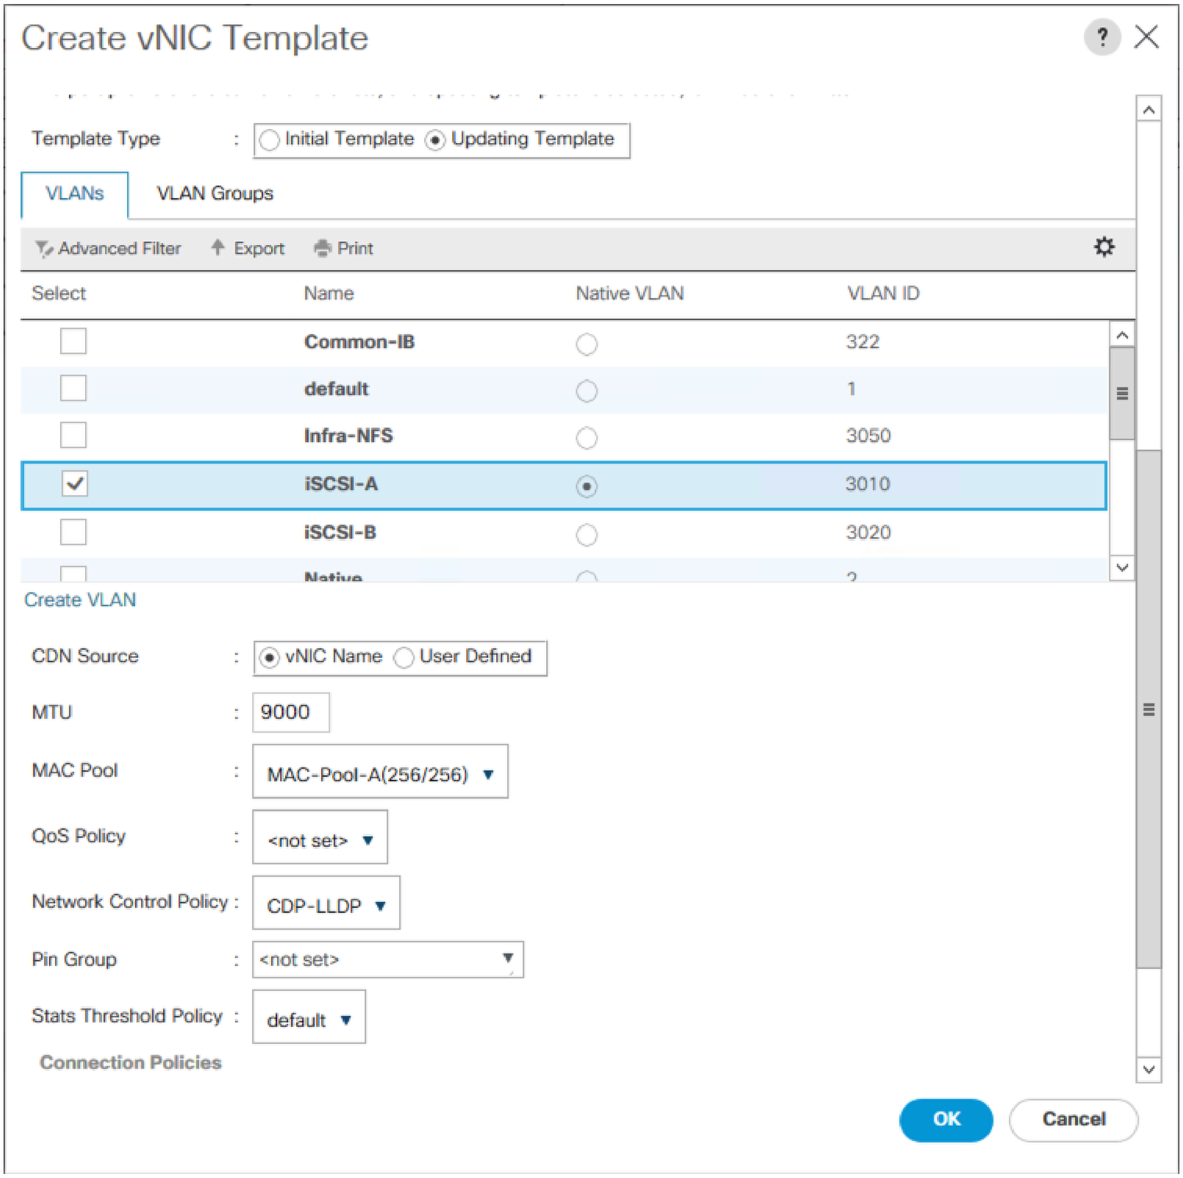

Create vNIC Template

Template Type : ○ Initial Template ● Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN	VLAN ID
□	Common-IB	○	322
□	default	○	1
□	Infra-NFS	○	3050
☑	iSCSI-A	●	3010
□	iSCSI-B	○	3020
□	Native	○	2

Create VLAN

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MAC-Pool-A(256/256)

QoS Policy : <not set>

Network Control Policy : CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

70. Click OK to complete creating the vNIC template.

71. Click OK.

72. Right-click vNIC Templates.

73. Choose Create vNIC Template.

74. Enter iSCSI-B as the vNIC template name.

75. Choose Fabric B. Do not choose the Enable Failover checkbox.

76. Leave Redundancy Type set at No Redundancy.

77. Under Target, make sure that only the Adapter checkbox is selected.

78. Choose Updating Template for Template Type.

79. Under VLANs, choose only iSCSI-B.

80. Choose Infra-iSCSI-B as the native VLAN.

81. Leave vNIC Name set for the CDN Source.

82. Under MTU, enter 9000.

83. From the MAC Pool list, choose MAC-Pool-B.

84. From the Network Control Policy list, choose CDP-LLDP.

85. Click OK to complete creating the vNIC template.

86. Click OK.

Create High Traffic VMware Adapter Policy

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

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

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Right-click Adapter Policies and choose Create Ethernet Adapter Policy.

4. Name the policy VMware-HighTrf.

5. Expand Resources and set the values as shown below.

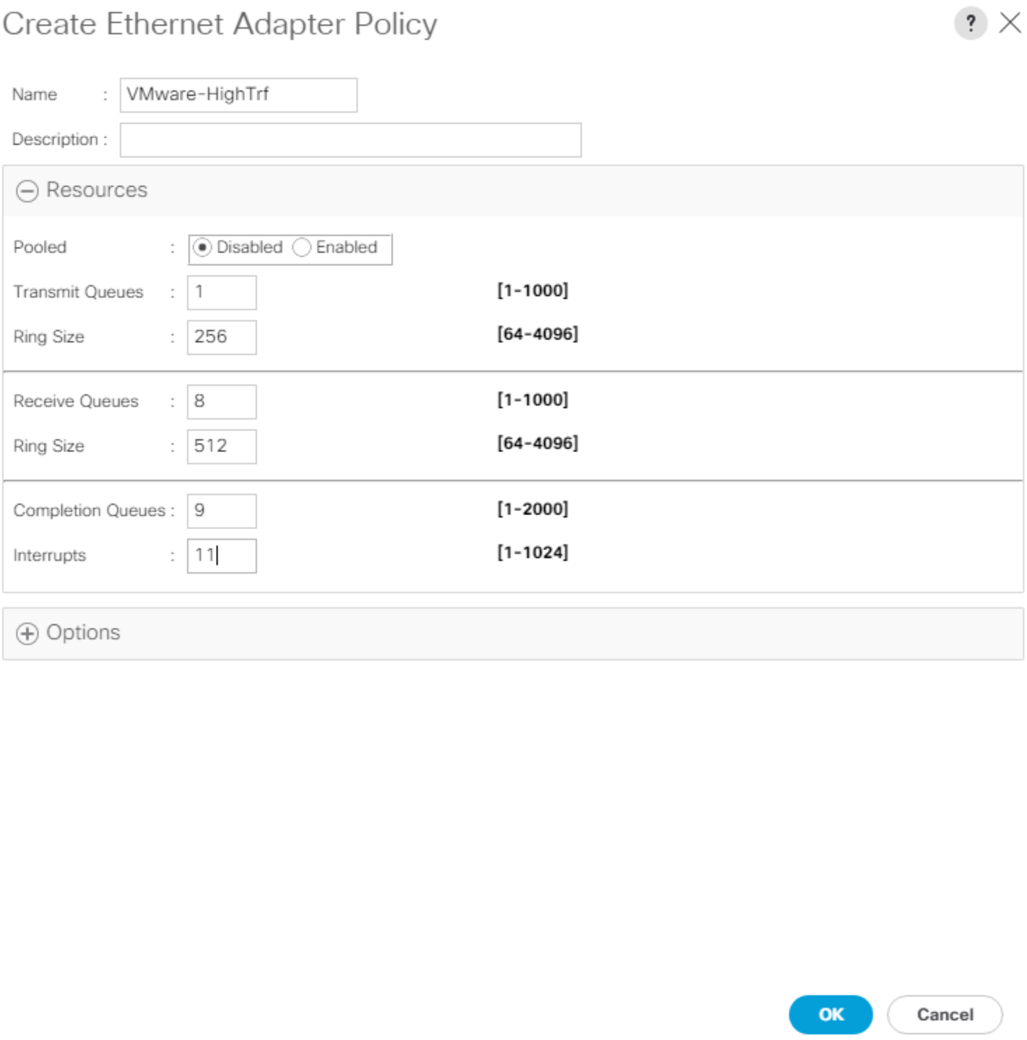

Create Ethernet Adapter Policy

Name	VMware-HighTrf	
Description		

⊖ Resources

Pooled	● Disabled ○ Enabled	
Transmit Queues	1	[1-1000]
Ring Size	256	[64-4096]
Receive Queues	8	[1-1000]
Ring Size	512	[64-4096]
Completion Queues	9	[1-2000]
Interrupts	11	[1-1024]

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

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

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

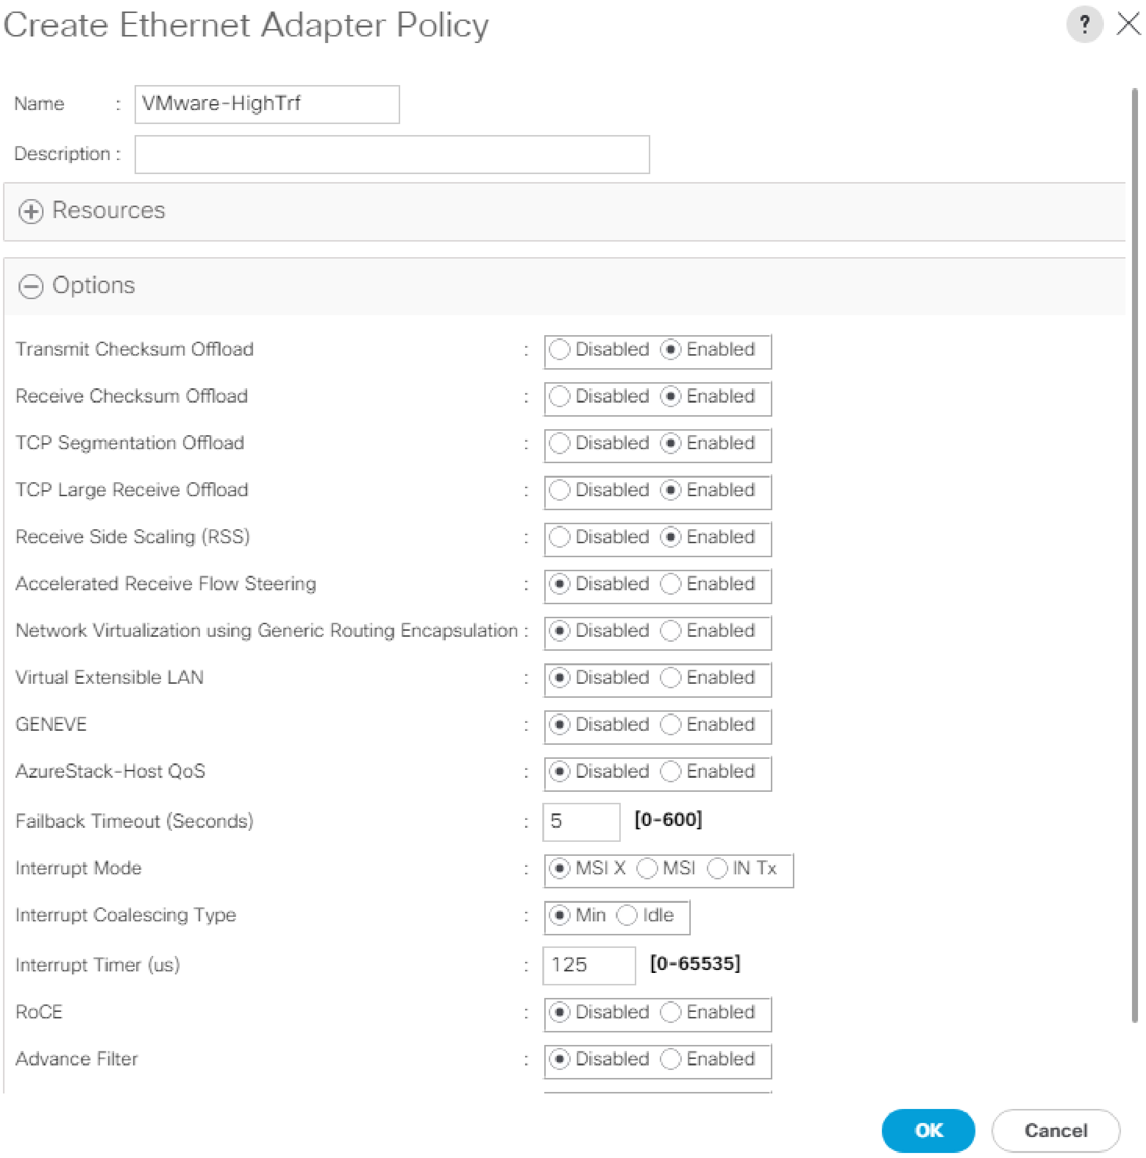

Create Ethernet Adapter Policy

Name : VMware-HighTrf

Description :

⊕ Resources

⊖ Options

Transmit Checksum Offload : ○ Disabled ● Enabled

Receive Checksum Offload : ○ Disabled ● Enabled

TCP Segmentation Offload : ○ Disabled ● Enabled

TCP Large Receive Offload : ○ Disabled ● Enabled

Receive Side Scaling (RSS) : ○ Disabled ● Enabled

Accelerated Receive Flow Steering : ● Disabled ○ Enabled

Network Virtualization using Generic Routing Encapsulation : ● Disabled ○ Enabled

Virtual Extensible LAN : ● Disabled ○ Enabled

GENEVE : ● Disabled ○ Enabled

AzureStack-Host QoS : ● Disabled ○ Enabled

Failback Timeout (Seconds) : 5 [0-600]

Interrupt Mode : ● MSI X ○ MSI ○ IN Tx

Interrupt Coalescing Type : ● Min ○ Idle

Interrupt Timer (us) : 125 [0-65535]

RoCE : ● Disabled ○ Enabled

Advance Filter : ● Disabled ○ Enabled

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

Create LAN Connectivity Policy for iSCSI Boot

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

1. In Cisco UCS Manager, click LAN.

2. Expand LAN > Policies > root > Sub-Organizations > FlexPod Organization.

3. Right-click LAN Connectivity Policies under the FlexPod Organization.

4. Choose Create LAN Connectivity Policy.

5. Enter iSCSI-Boot as the name of the policy.

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

7. On the left under LAN > Policies > root > Sub-Organizations > FlexPod Organization > LAN Connectivity Policies, choose iSCSI-Boot.

8. Click the Add button to add a vNIC.

9. In the Create vNIC dialog box, enter 00-vSwitch0-A as the name of the vNIC.

10. Choose the Use vNIC Template checkbox.

11. In the vNIC Template list, choose vSwitch0-A.

12. In the Adapter Policy list, choose VMWare.

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

14. Click Save Changes and OK.

15. Click the Add button to add another vNIC to the policy.

16. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.

17. Choose the Use vNIC Template checkbox.

18. In the vNIC Template list, choose vSwitch0-B.

19. In the Adapter Policy list, choose VMWare.

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

21. Click Save Changes and OK.

22. Click the Add button to add a vNIC.

23. In the Create vNIC dialog box, enter 02-vDS0-A as the name of the vNIC.

24. Choose the Use vNIC Template checkbox.

25. In the vNIC Template list, choose vDS0-A.

26. In the Adapter Policy list, choose VMWare-HighTrf.

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

28. Click Save Changes and OK.

29. Click the Add button to add another vNIC to the policy.

30. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.

31. Choose the Use vNIC Template checkbox.

32. In the vNIC Template list, choose vDS0-B.

33. In the Adapter Policy list, choose VMWare-HighTrf.

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

35. Click Save Changes and OK.

36. Click the Add button to add a vNIC.

37. In the Create vNIC dialog box, enter 04-iSCSI-A as the name of the vNIC.

38. Choose the Use vNIC Template checkbox.

39. In the vNIC Template list, choose iSCSI-A.

40. In the Adapter Policy list, choose VMWare.

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

42. Click Save Changes and OK.

43. Click Add to add a vNIC to the policy.

44. In the Create vNIC dialog box, enter 05-iSCSI-B as the name of the vNIC.

45. Choose the Use vNIC Template checkbox.

46. In the vNIC Template list, choose iSCSI-B.

47. In the Adapter Policy list, choose VMWare.

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

49. Click Save Changes and OK.

50. Expand Add iSCSI vNICs.

51. Choose Add in the Add iSCSI vNICs section.

52. Set the name to iSCSI-Boot-A.

53. Choose 04-iSCSI-A as the Overlay vNIC.

54. Set the iSCSI Adapter Policy to default.

55. Leave the VLAN set to Infra-iSCSI-A (native).

56. Leave the MAC Address set to None.

57. Click OK.

58. Click Save Changes and OK.

59. Choose Add in the Add iSCSI vNICs section.

60. Set the name to iSCSI-Boot-B.

61. Choose 05-iSCSI-B as the Overlay vNIC.

62. Set the iSCSI Adapter Policy to default.

63. Leave the VLAN set to Infra-iSCSI-B (native).

64. Leave the MAC Address set to None.

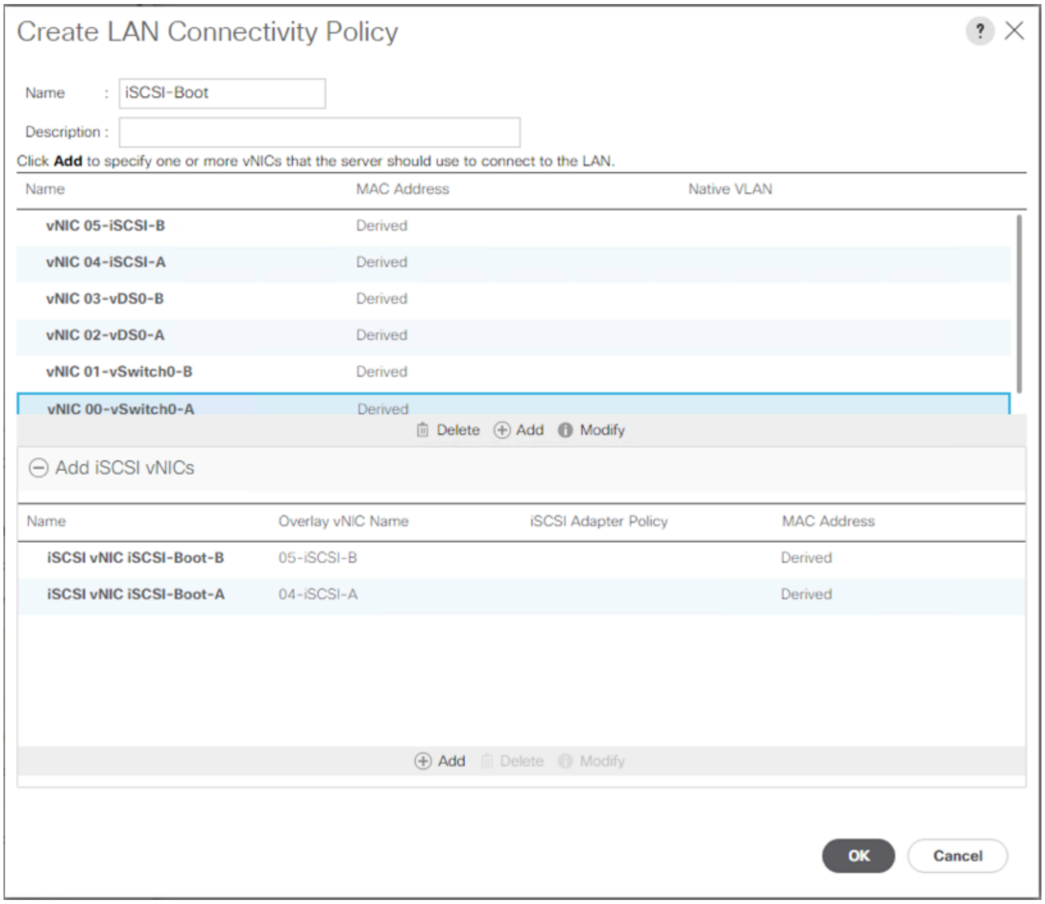

Create LAN Connectivity Policy

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-iSCSI-B	Derived	
vNIC 04-iSCSI-A	Derived	
vNIC 03-vDS0-B	Derived	
vNIC 02-vDS0-A	Derived	
vNIC 01-vSwitch0-B	Derived	
vNIC 00-vSwitch0-A	Derived	

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-B	05-iSCSI-B		Derived
iSCSI vNIC iSCSI-Boot-A	04-iSCSI-A		Derived

65. Click OK then click OK again.

Create IQN Pools for iSCSI Boot

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

1. In Cisco UCS Manager, click SAN.
2. Expand Pools > root.
3. Right-click IQN Pools.
4. Choose Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter iqn.2010-11.com.flexpod as the prefix.
8. Choose Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter ucs-host as the suffix.

> If multiple Cisco UCS domains are being used, a more specific IQN suffix may need to be used.

12. Enter 1 in the From field.
13. Specify the size of the IQN block sufficient to support the available server resources.

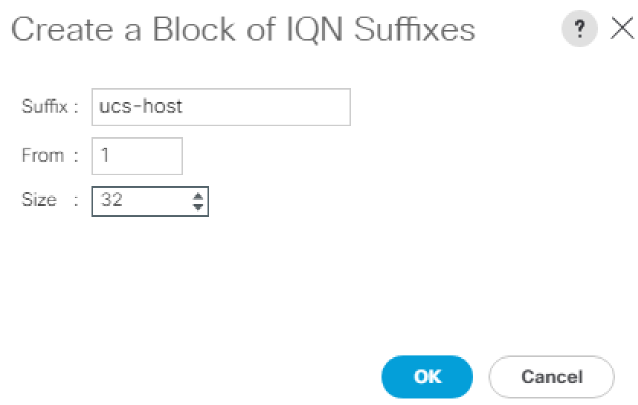

Create a Block of IQN Suffixes

Suffix : ucs-host

From : 1

Size : 32

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

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment in the FlexPod Organization, follow these steps:

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

1. In Cisco UCS Manager, click Servers.
2. Expand Pools > root > Sub-Organizations > FlexPod.
3. Right-click Server Pools under the FlexPod Organization.
4. Choose Create Server Pool.
5. Enter Infra-Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Choose three (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra-Pool server pool.

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

9. Click Finish.
10. Click OK.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click Servers.
2. Expand Pools > root.
3. Right-click UUID Suffix Pools.
4. Choose Create UUID Suffix Pool.
5. Enter UUID-Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Choose Sequential for the Assignment Order.
9. Click Next.

10. Click Add to add a block of UUIDs.

11. Keep the From field at the default setting.

12. Specify a size for the UUID block that is sufficient to support the available blade or server resources and the number of Service Profiles that will be created.

13. Click OK.

14. Click Finish.

15. Click OK.

Modify Default Host Firmware Package

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

To modify the default firmware management policy in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Expand Host Firmware Packages.

4. Choose default.

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

6. Choose version 4.1(2a) for both the Blade and Rack Packages.

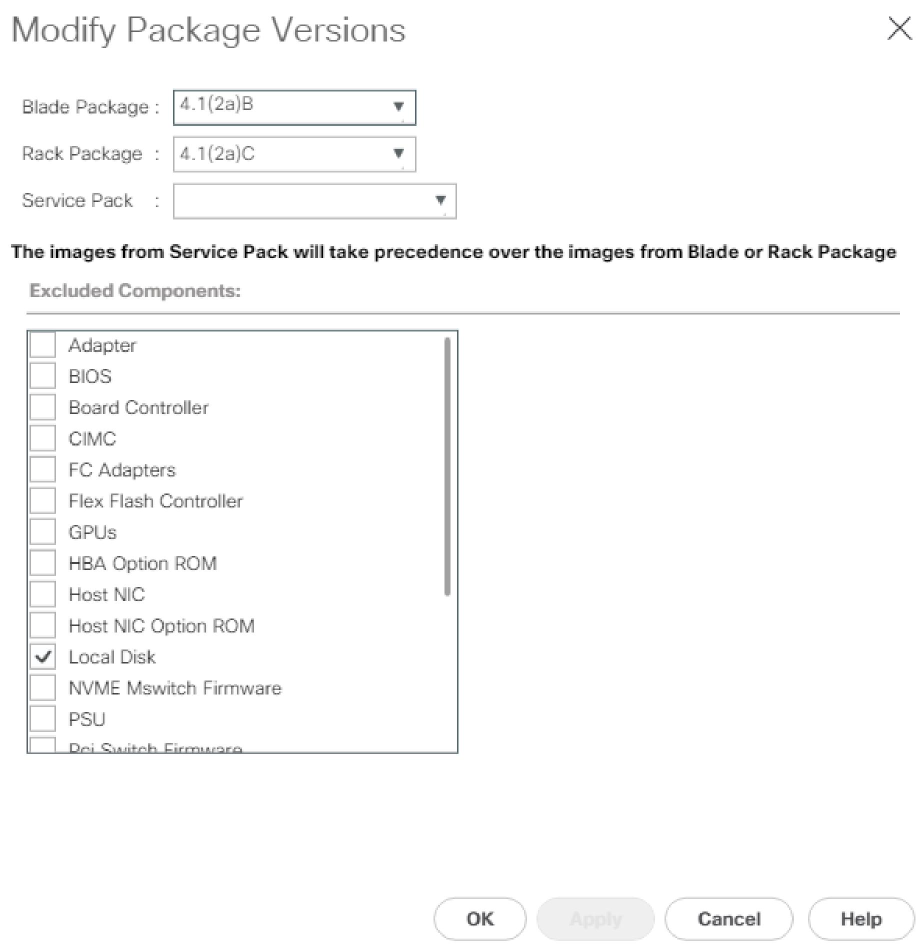

Modify Package Versions

Blade Package : 4.1(2a)B

Rack Package : 4.1(2a)C

Service Pack :

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
- □ Pci Switch Firmware

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

Create Local Disk Configuration Policy (Optional)

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

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

To create a local disk configuration policy, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Right-click Local Disk Config Policies.

4. Choose Create Local Disk Configuration Policy.

5. Enter SAN-Boot as the local disk configuration policy name.

6. Change the mode to No Local Storage.

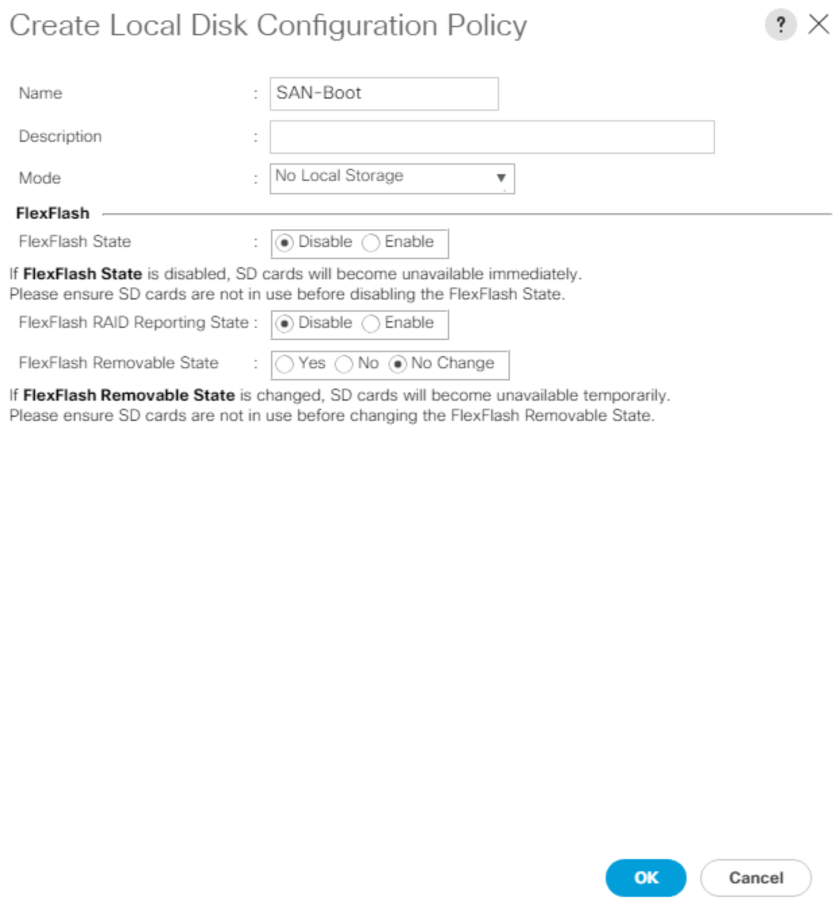

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : ● Disable ○ Enable

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

FlexFlash RAID Reporting State : ● Disable ○ Enable

FlexFlash Removable State : ○ Yes ○ No ● No Change

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

7. Click OK to create the local disk configuration policy.

8. Click OK.

Create Power Control Policy

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

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Right-click Power Control Policies.

4. Choose Create Power Control Policy.

5. Enter No-Power-Cap as the power control policy name.

6. Change the power capping setting to No Cap.

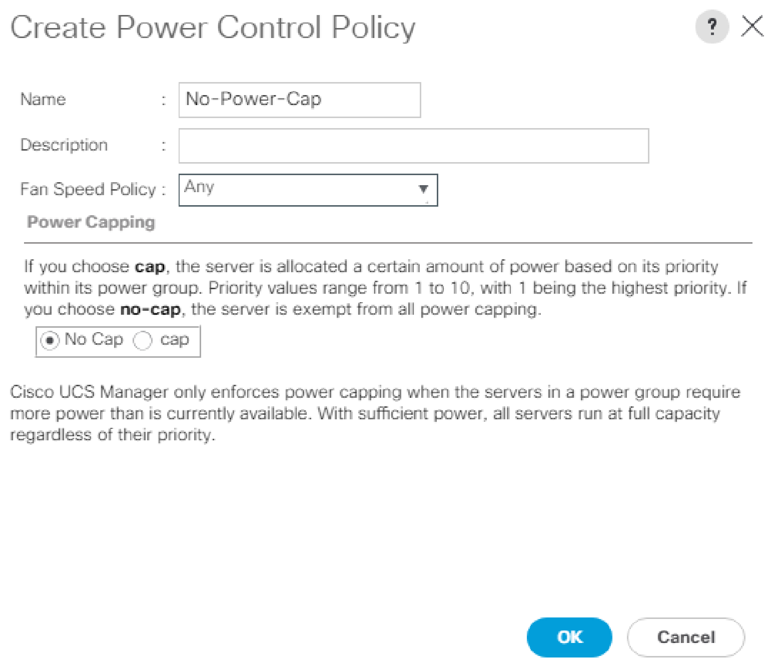

Create Power Control Policy

Name : 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.

7. Click OK to create the power control policy.

8. Click OK.

Create Server Pool Qualification Policy (Optional)

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

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

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Right-click Server Pool Policy Qualifications.

4. Choose Create Server Pool Policy Qualification.

5. Name the policy UCS-B200-M5.

6. Choose Create Server PID Qualifications.

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

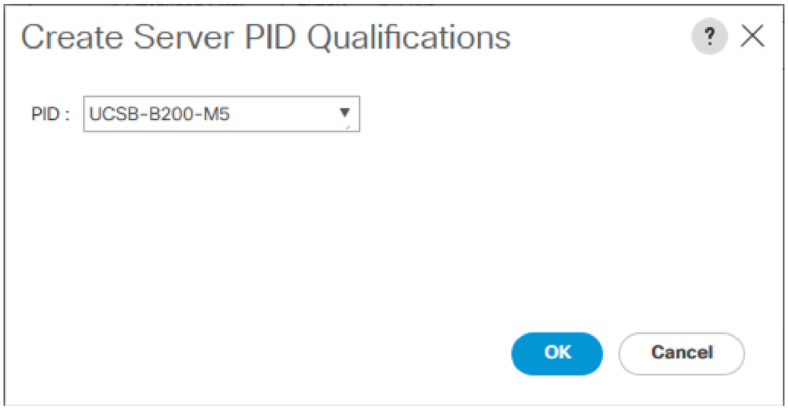

PID : UCSB-B200-M5

8. Click OK

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

10. Click OK to create the policy then OK for the confirmation.

Update the Default Maintenance Policy

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

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root.

3. Choose Maintenance Policies > default.

4. Change the Reboot Policy to User Ack.

5. Choose “On Next Boot” to delegate maintenance windows to server administrators.

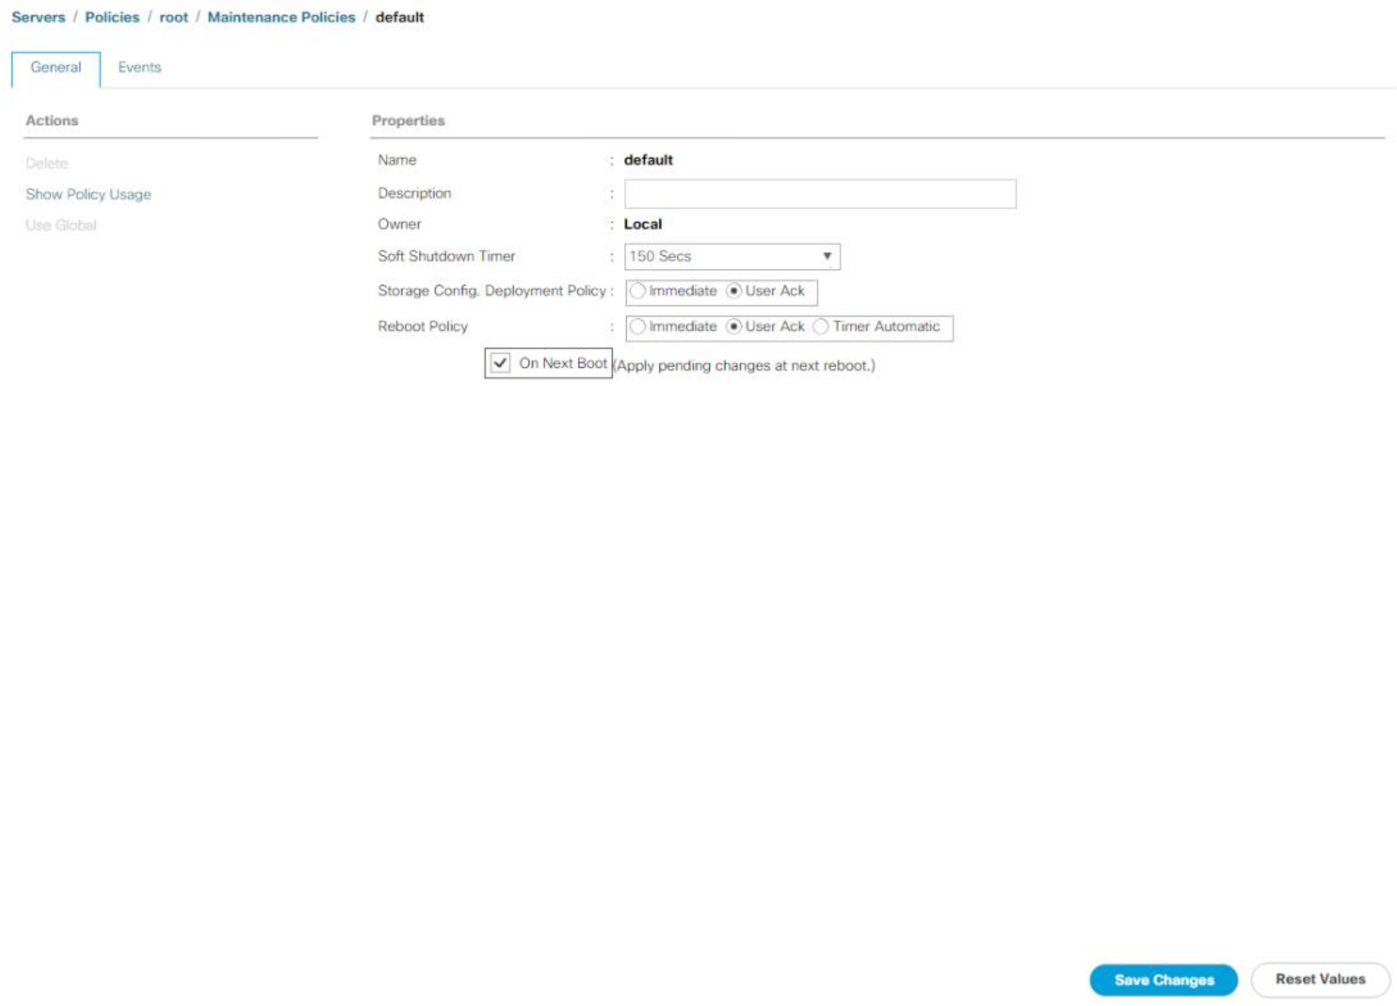

Servers / Policies / root / Maintenance Policies / **default**

Properties

Name : **default**

Description :

Owner : **Local**

Soft Shutdown Timer : 150 Secs ▼

Storage Config. Deployment Policy : ○ Immediate ● User Ack

Reboot Policy : ○ Immediate ● User Ack ○ Timer Automatic

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

6. Click Save Changes.

7. Click OK to accept the changes.

Create vMedia Policy for VMware ESXi 7.0 ISO Install Boot

In the NetApp ONTAP setup steps, an HTTP web server is required, which is used for hosting ONTAP as well as VMware software. The vMedia Policy created will map the Cisco Custom ISO for UCS 4.1.2a to the Cisco UCS server in order to boot the ESXi installation. To create this policy, follow these steps:

1. In Cisco UCS Manager, choose Servers.

2. Expand Policies > root.

3. Right-click vMedia Policies.

4. Choose Create vMedia Policy.

5. Name the policy ESXi-7.0-HTTP.

6. Enter “Mounts Cisco Custom ISO ESXi7 for UCS 4.1(2a)” in the Description field.

7. Click Add to add a vMedia Mount.

8. Name the mount ESXi-7.0-HTTP.

9. Choose the CDD Device Type.

10. Choose the HTTP Protocol.

11. Enter the IP Address of the web server.

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

12. Enter VMware-ESXi-7.0.0-16324942-Custom-Cisco-4.1.2a.iso as the Remote File name.

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

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

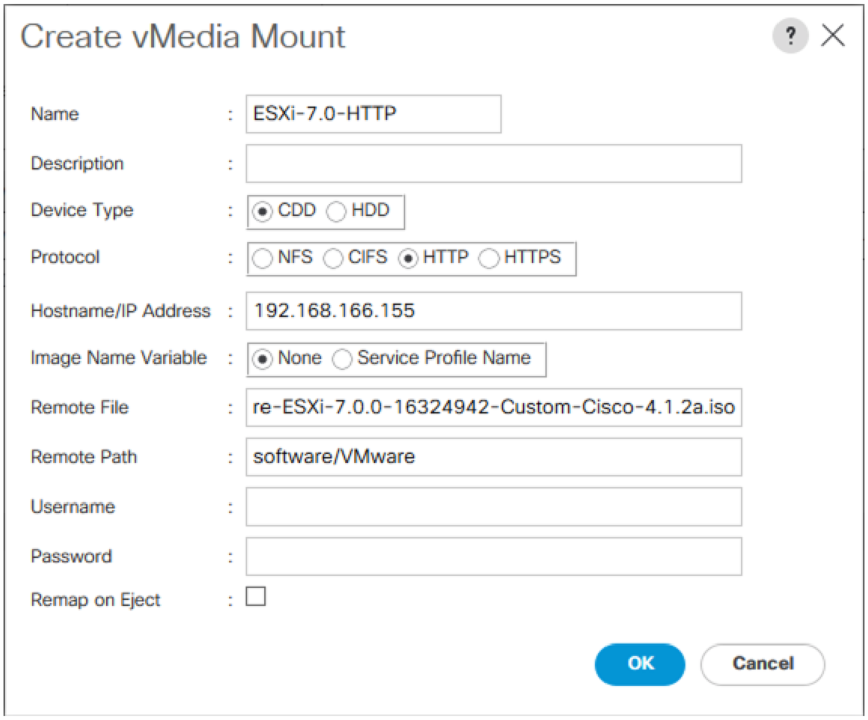

Create vMedia Mount

Name :	ESXi-7.0-HTTP
Description :	
Device Type :	● CDD ○ HDD
Protocol :	○ NFS ○ CIFS ● HTTP ○ HTTPS
Hostname/IP Address :	192.168.166.155
Image Name Variable :	● None ○ Service Profile Name
Remote File :	re-ESXi-7.0.0-16324942-Custom-Cisco-4.1.2a.iso
Remote Path :	software/VMware
Username :	
Password :	
Remap on Eject :	□

14. Click OK to create the vMedia Mount.

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

For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi host. On first boot the host will boot into the ESXi installer since the iSCSI mounted disk is empty. After ESXi is installed, the vMedia will not be referenced as long as the boot disk is accessible.

Create Server BIOS Policy

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

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root > Sub-Organizations > FlexPod.

3. Right-click BIOS Policies under FlexPod Organization.

4. Choose Create BIOS Policy.

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

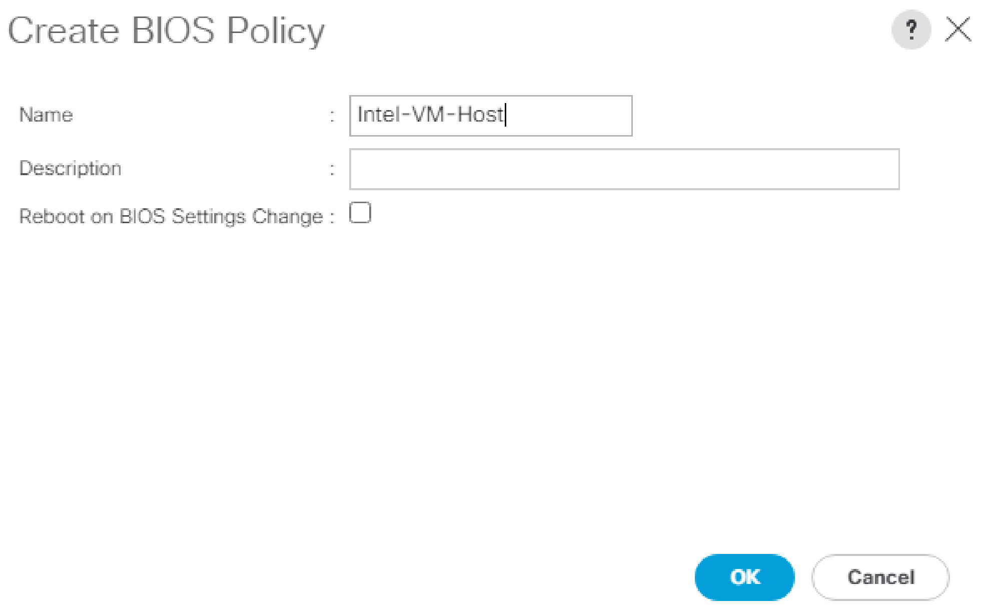

Create BIOS Policy

Name : Intel-VM-Host

Description :

Reboot on BIOS Settings Change : □

6. Click OK, then click OK again to create the BIOS Policy.

7. Under the FlexPod Organization, expand BIOS Policies and choose the newly created BIOS Policy. Set the following within the Main tab of the Policy:

   a. CDN Control -> Enabled

   b. Quiet Boot -> Disabled

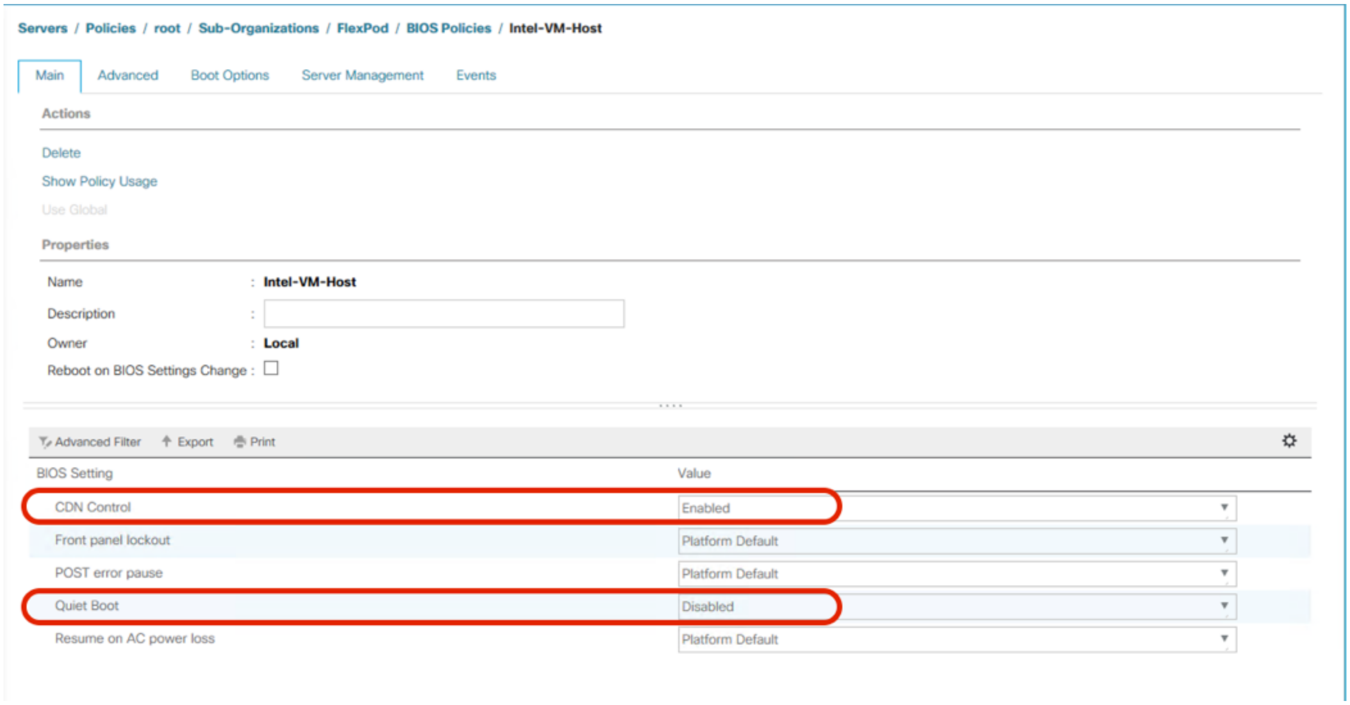

Servers / Policies / root / Sub-Organizations / FlexPod / BIOS Policies / Intel-VM-Host

Name : **Intel-VM-Host**

Description :

Owner : **Local**

Reboot on BIOS Settings Change : □

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

8. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab. Set the following within the Processor tab:

   a. Processor C State -> Disabled

   b. Processor C1E -> Disabled

   c. Processor C3 Report -> Disabled

   d. Processor C6 Report -> Disabled

   e. Processor C7 Report -> Disabled

   f. Power Technology -> Custom

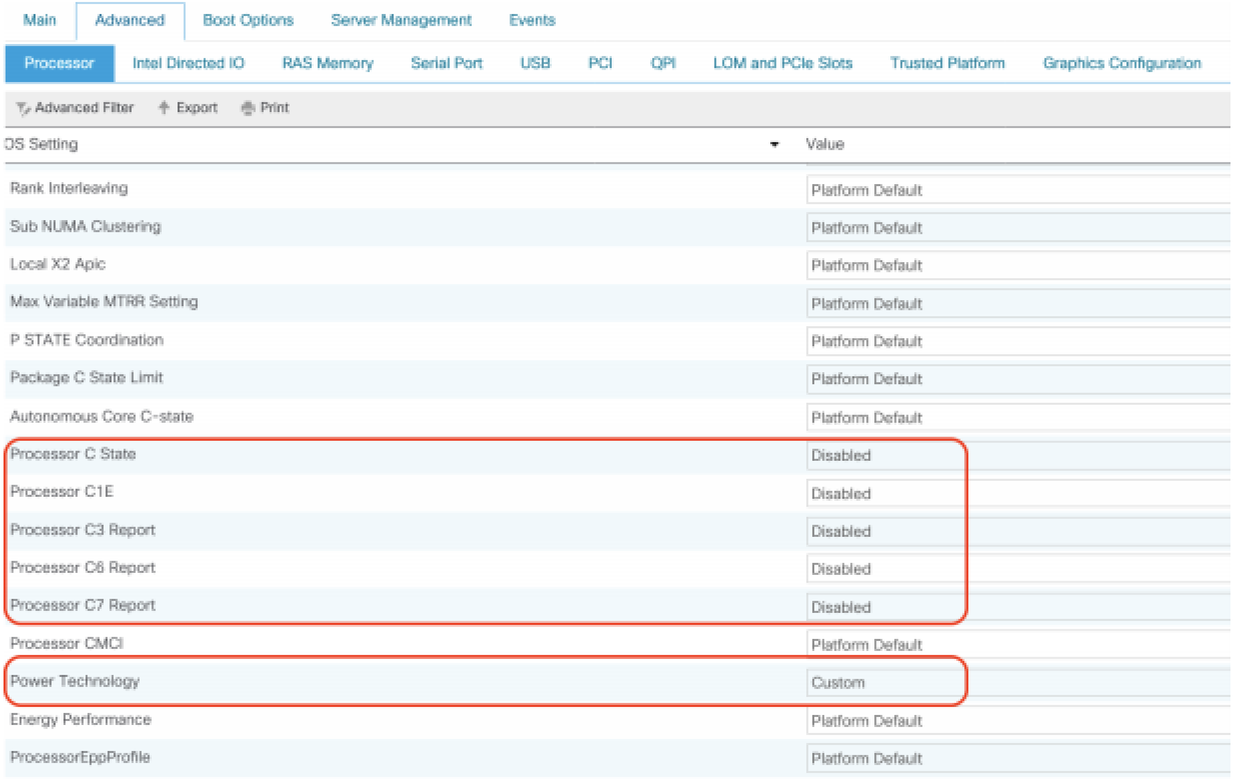

OS Setting	Value
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Autonomous Core C-state	Platform Default
Processor C State	Disabled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

9. Click the RAS Memory tab, and choose:

    a. NVM Performance Setting -> Balanced Profile

    b. Memory RAS configuration -> Maximum Performance

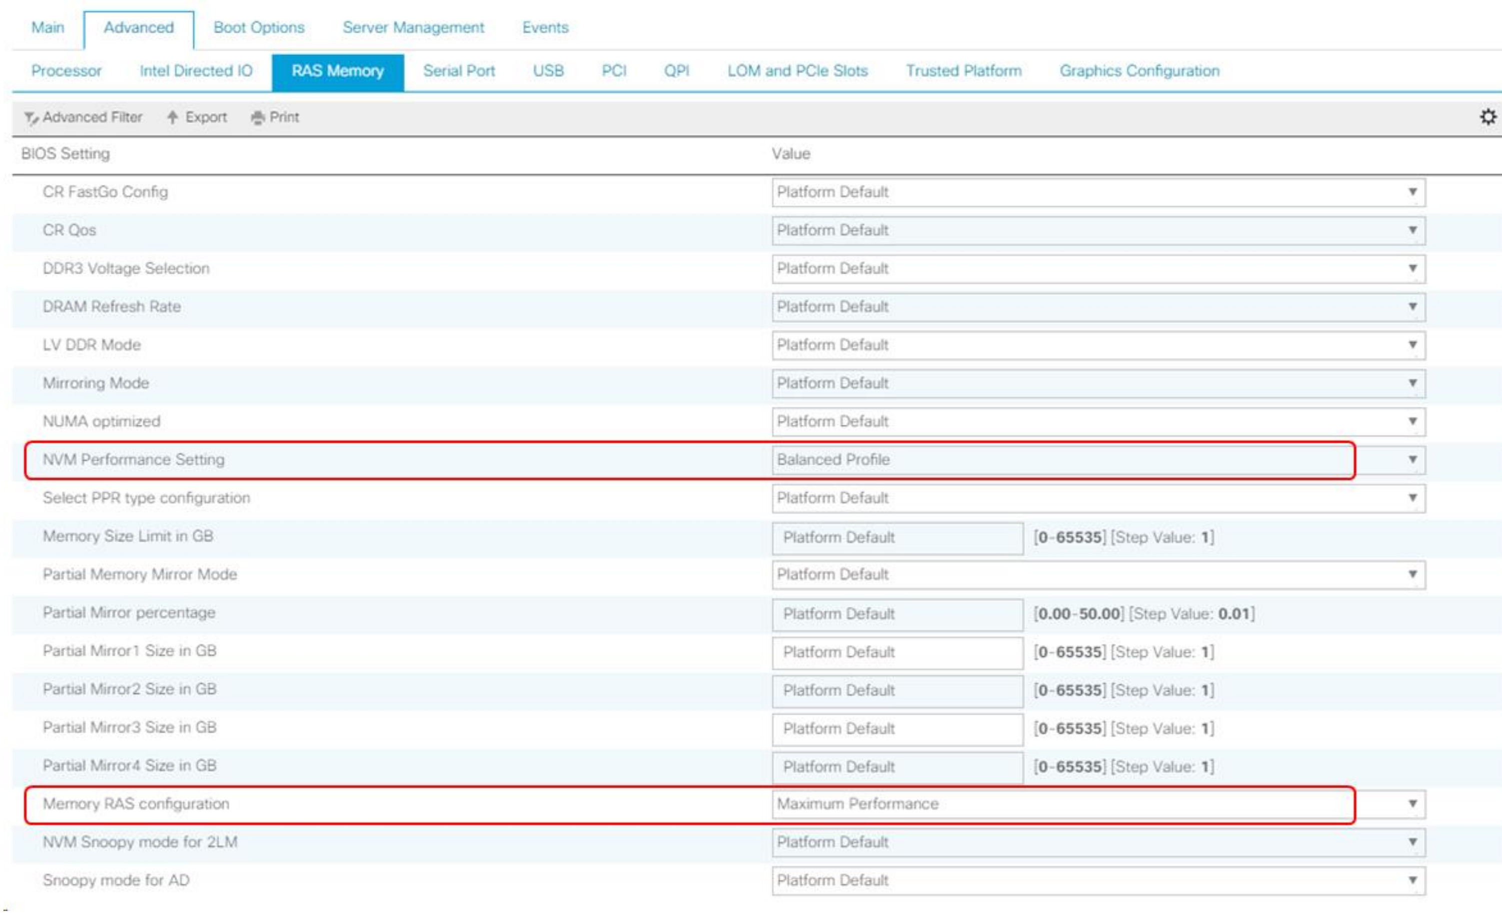

BIOS Setting	Value
CR FastGo Config	Platform Default
CR Qos	Platform Default
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Platform Default
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
NVM Performance Setting	Balanced Profile
Select PPR type configuration	Platform Default
Memory Size Limit in GB	Platform Default [0-65535] [Step Value: 1]
Partial Memory Mirror Mode	Platform Default
Partial Mirror percentage	Platform Default [0.00-50.00] [Step Value: 0.01]
Partial Mirror1 Size in GB	Platform Default [0-65535] [Step Value: 1]
Partial Mirror2 Size in GB	Platform Default [0-65535] [Step Value: 1]
Partial Mirror3 Size in GB	Platform Default [0-65535] [Step Value: 1]
Partial Mirror4 Size in GB	Platform Default [0-65535] [Step Value: 1]
Memory RAS configuration	Maximum Performance
NVM Snoopy mode for 2LM	Platform Default
Snoopy mode for AD	Platform Default

10. Click Save Changes.

11. Click OK.

Create iSCSI Boot Policy

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

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

To create a boot policy for the Cisco UCS environment within the FlexPod Organization, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Policies > root > Sub-Organizations > FlexPod Organization.

3. Right-click Boot Policies under the FlexPod Organization.

4. Choose Create Boot Policy.

5. Enter Boot-iSCSI-A as the name of the boot policy.

6. Optional: Enter a description for the boot policy.

7. Do not choose the Reboot on Boot Order Change checkbox.

8. Choose the Uefi Boot Mode.

9. Check the checkbox for Boot Security.

10. Expand the Local Devices drop-down list and click Add Remote CD/DVD.

11. Expand the iSCSI vNICs drop-down list and click Add iSCSI Boot.

12. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.

13. Click OK.

14. Choose Add iSCSI Boot.

15. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-B.

16. Click OK.

17. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

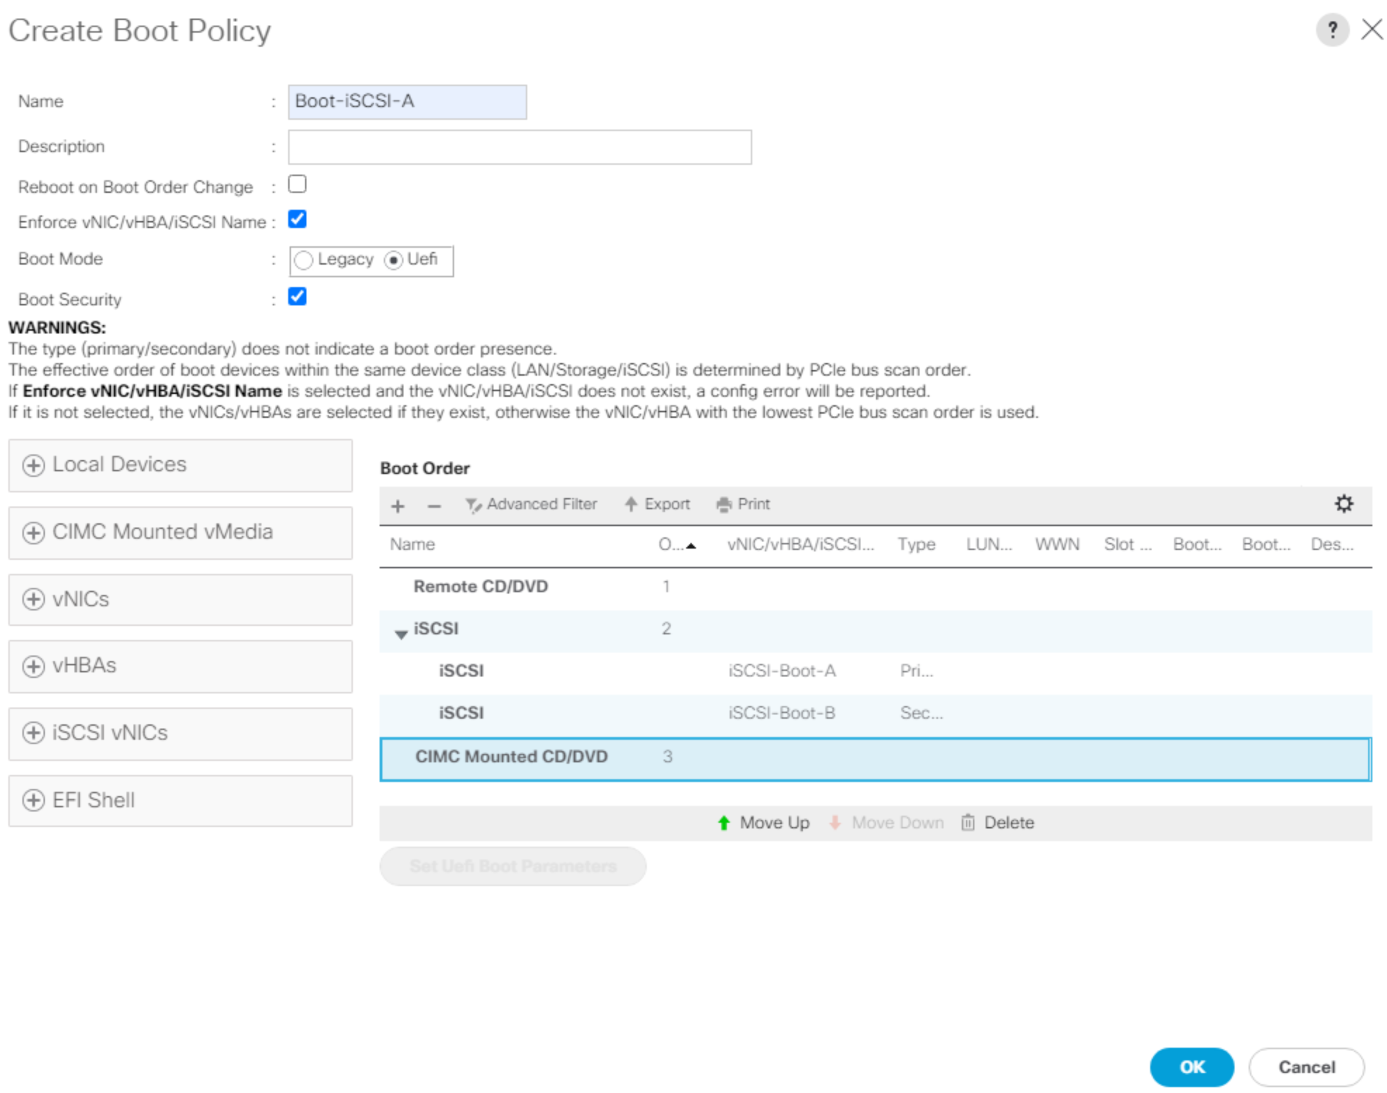

Create Boot Policy

Name : Boot-iSCSI-A

Description :

Reboot on Boot Order Change : □

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ○ Legacy ● Uefi

Boot Security : ☑

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

Boot Order

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

18. Expand iSCSI and select iSCSI Target Primary. Select Set Uefi Boot Parameters.

For Cisco UCS B200 M5 and Cisco UCS C220 M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for M4 and earlier servers, VMware ESXi 7.0 will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

19. Fill-in the Set Uefi Boot Parameters exactly as shown in the following screenshot:

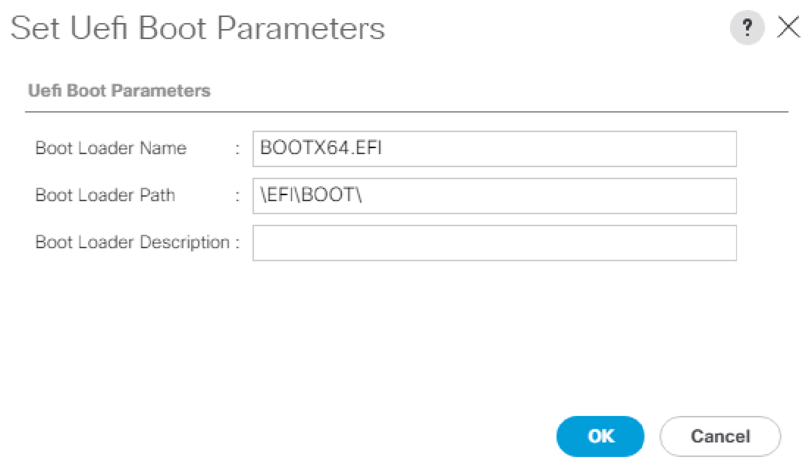

Set Uefi Boot Parameters

Uefi Boot Parameters

Boot Loader Name : BOOTX64.EFI

Boot Loader Path : \EFI\BOOT\

Boot Loader Description :

20. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target and click OK for the confirmation.

21. Repeat this process to set Uefi Boot Parameters for the Secondary iSCSI Boot Targets.

22. Click OK then click OK again to create the policy.

Create iSCSI Boot Service Profile Template

In this procedure, one service profile template for Infrastructure ESXi hosts within the FlexPod Organization is created for Fabric A boot.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Service Profile Templates > root > Sub-Organizations > FlexPod Organization.

3. Right-click the FlexPod Organization.

4. Choose Create Service Profile Template to open the Create Service Profile Template wizard.

5. Enter Intel-VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.

6. Choose the Updating Template option.

7. Under UUID Assignment, choose UUID-Pool.

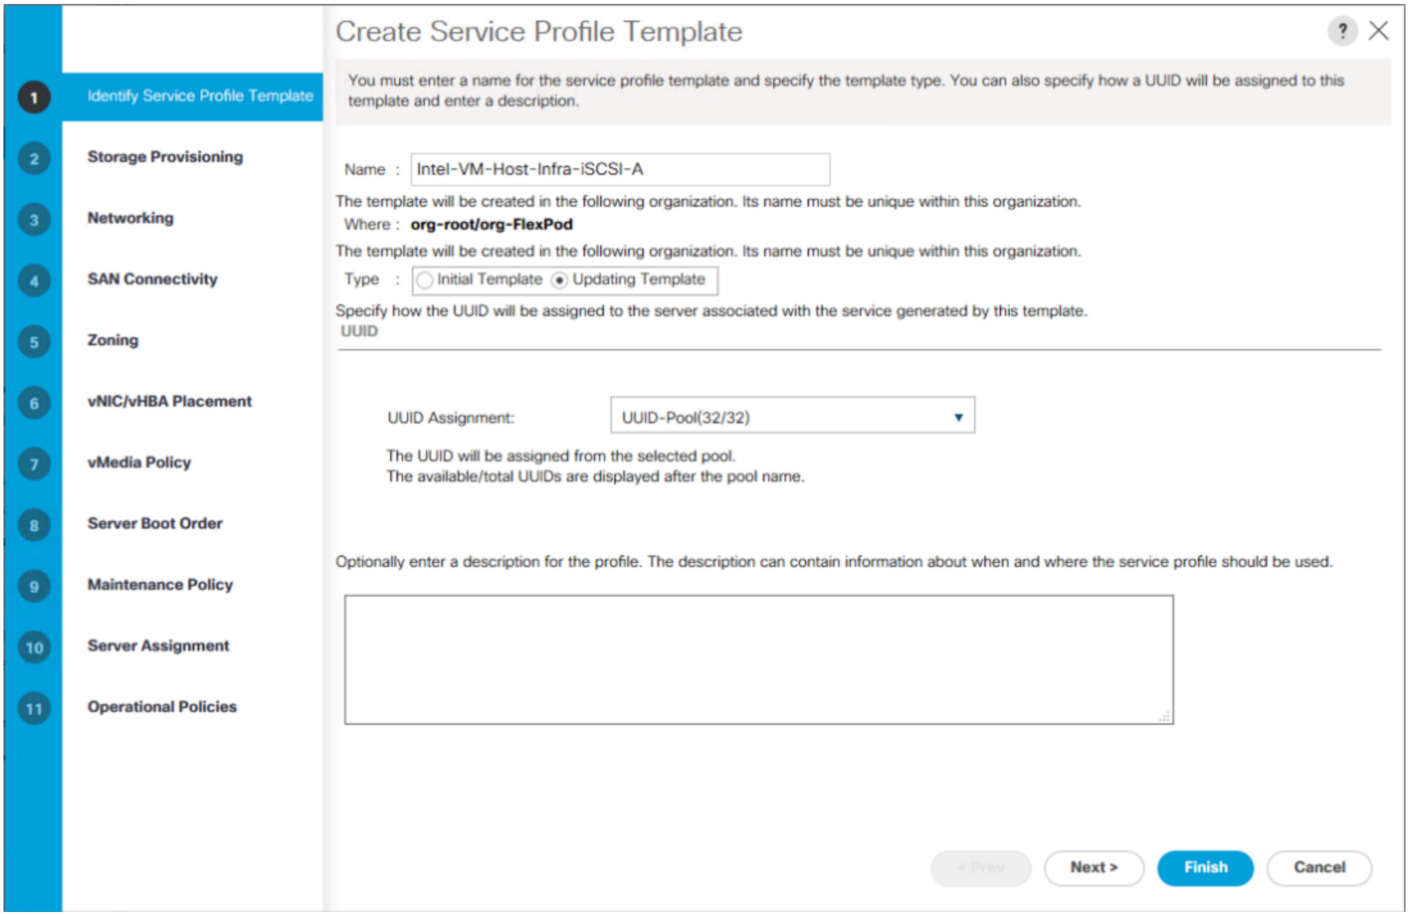

Create Service Profile Template

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 : Intel-VM-Host-Infra-iSCSI-A

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-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(32/32)

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

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

8. Click Next.

Configure Storage Provisioning

To configure the storage provisioning, follow these steps:

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy tab and choose the SAN-Boot Local Storage Policy. Otherwise, choose the default Local Storage Policy.

2. Click Next.

Configure Networking Options

To configure the network options, follow these steps:

1. Choose the “Use Connectivity Policy” option to configure the LAN connectivity.

2. Choose iSCSI-Boot from the LAN Connectivity Policy drop-down list.

3. Choose IQN_Pool in Initiator Name Assignment.

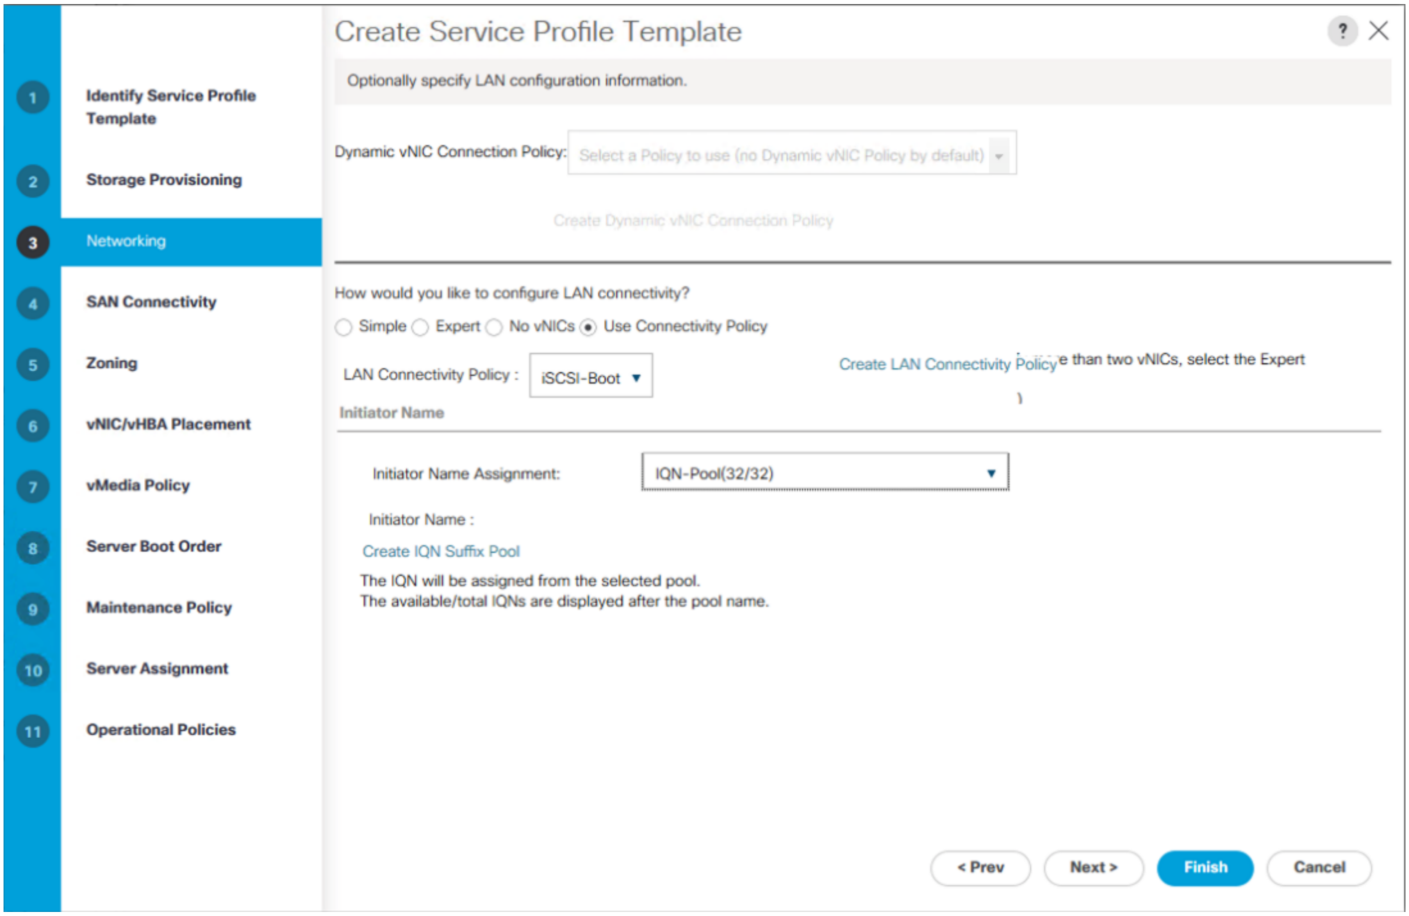

Optionally specify LAN configuration information.

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 : iSCSI-Boot

e than two vNICs, select the Expert

Initiator Name

Initiator Name Assignment: IQN-Pool(32/32)

Initiator Name :

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

4. Click Next.

Configure Storage Options

To configure the storage options, follow these steps:

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

2. Click Next.

Configure Zoning Options

To configure the zoning options, follow this step:

1. Make no changes and click Next.

Configure vNIC/HBA Placement

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

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

2. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

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

Configure Server Boot Order

To configure the server boot orders, follow these steps:

1. Choose Boot-iSCSI-A for Boot Policy.

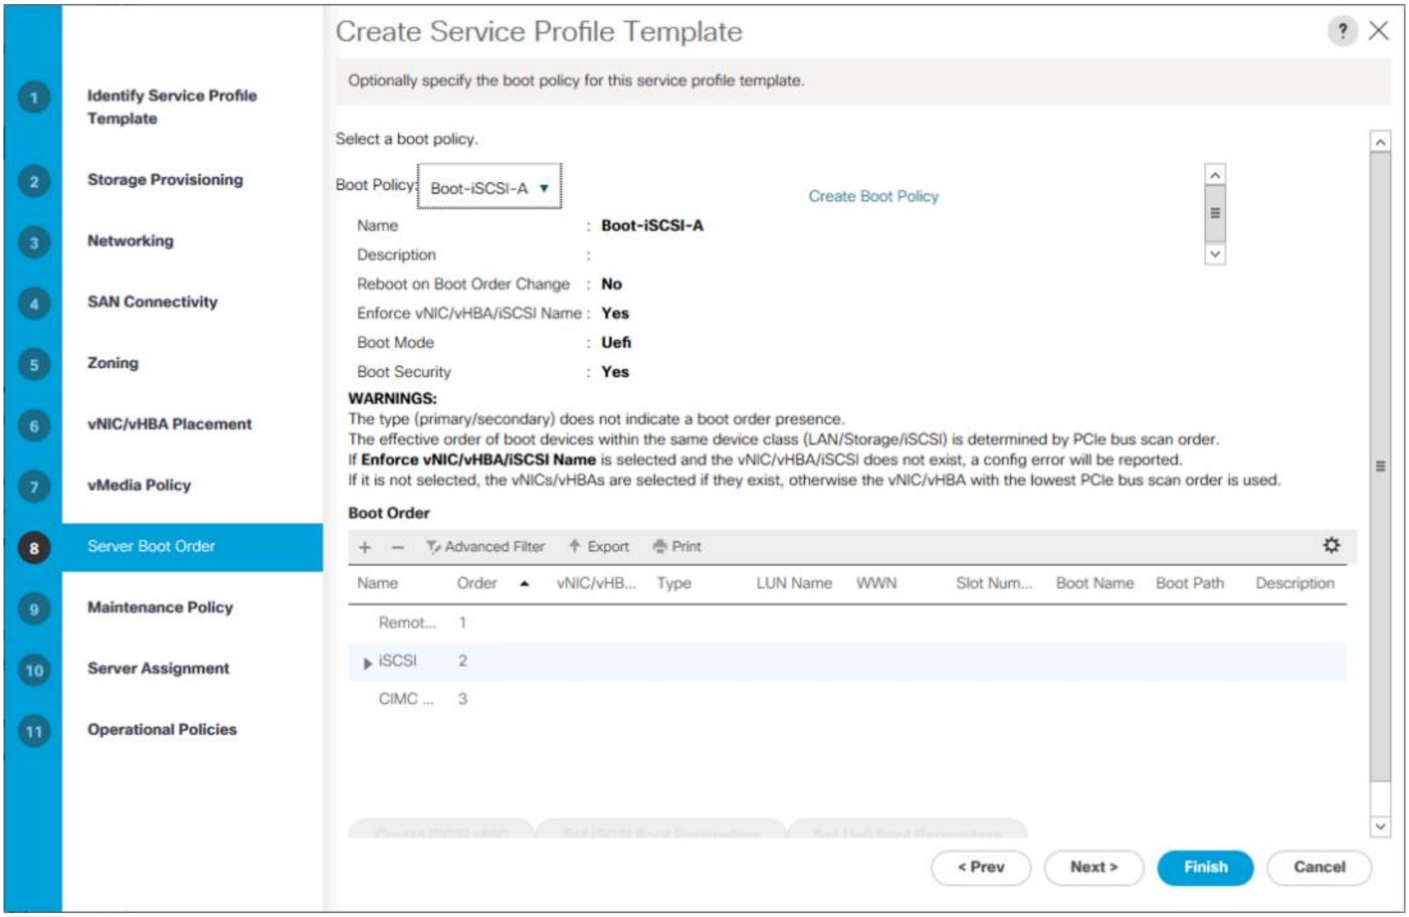

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Boot-iSCSI-A ▾

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

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

Boot Order

Name	Order ▲	vNIC/vHB...	Type	LUN Name	WWN	Slot Num...	Boot Name	Boot Path	Description
Remot...	1								
▸ iSCSI	2								
CIMC ...	3								

2. In the Boot order, expand iSCSI and choose iSCSI-Boot-A.
3. Click Set iSCSI Boot Parameters.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.

5. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.

6. Set iSCSI-IP-Pool-A as the “Initiator IP address Policy.”

7. Choose iSCSI Static Target Interface option.

8. Click Add.

9. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, log into the storage cluster management interface and run the “iscsi show” command”.

10. Enter the IP address of iscsi-lif-1a for the IPv4 Address field.

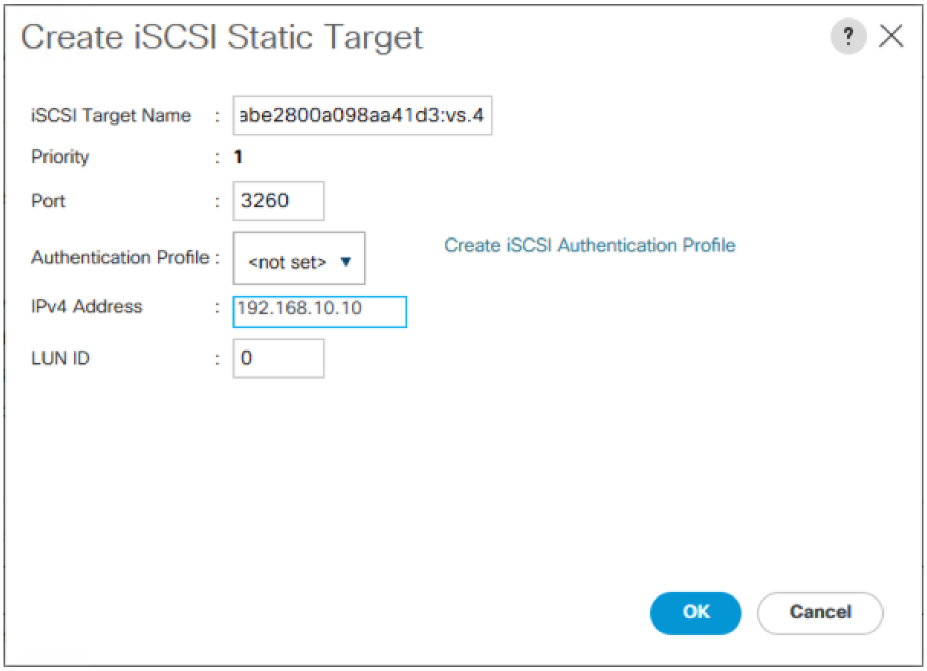

Create iSCSI Static Target

iSCSI Target Name : abe2800a098aa41d3:vs.4
Priority : 1
Port : 3260
Authentication Profile : <not set>
IPv4 Address : 192.168.10.10
LUN ID : 0

11. Click OK to add the iSCSI static target.

12. Click Add.

13. Enter the iSCSI Target Name.

14. Enter the IP address of iscsi-lif-2a for the IPv4 Address field.

15. Click OK to add the iSCSI static target.

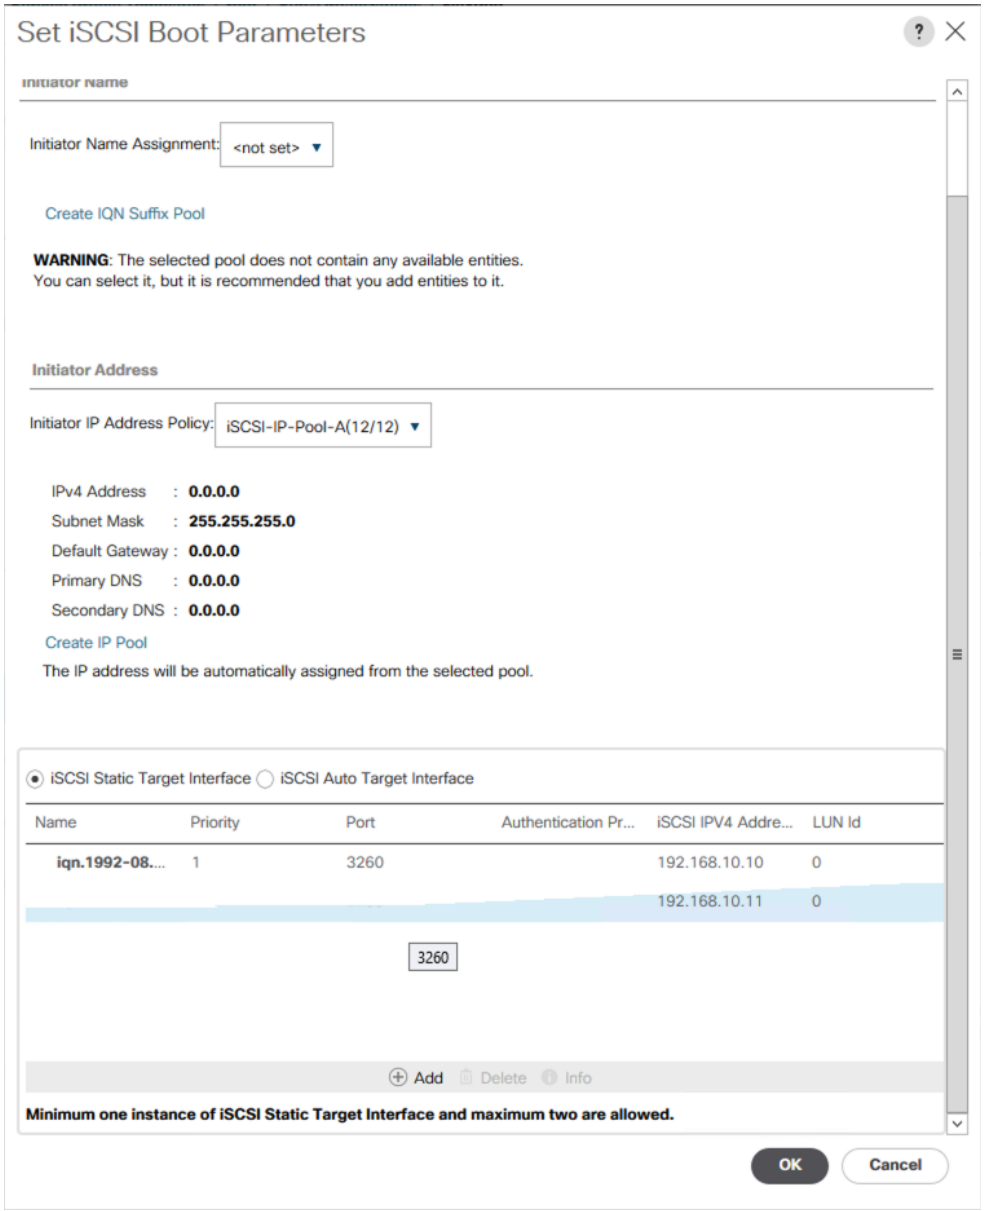

Set iSCSI Boot Parameters

Initiator Name

Initiator Name Assignment: <not set> ▾

Create IQN Suffix Pool

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

Initiator Address

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

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

Create IP Pool

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

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...	LUN Id
iqn.1992-08....	1	3260		192.168.10.10	0
				192.168.10.11	0

3260

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

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

17. In the Boot order, choose iSCSI-Boot-B.

18. Click Set iSCSI Boot Parameters.

19. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.

20. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.

21. Set iSCSI-IP-Pool-B as the “Initiator IP address Policy”.

22. Choose the iSCSI Static Target Interface option.

23. Click Add.

24. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, login into storage cluster management interface and run “iscsi show” command”.

25. Enter the IP address of iscsi-lif-01b for the IPv4 Address field.

26. Click OK to add the iSCSI static target.

27. Click Add.

28. Enter the iSCSI Target Name.

29. Enter the IP address of iscsi-lif-02b for the IPv4 Address field.

30. Click OK to add the iSCSI static target.

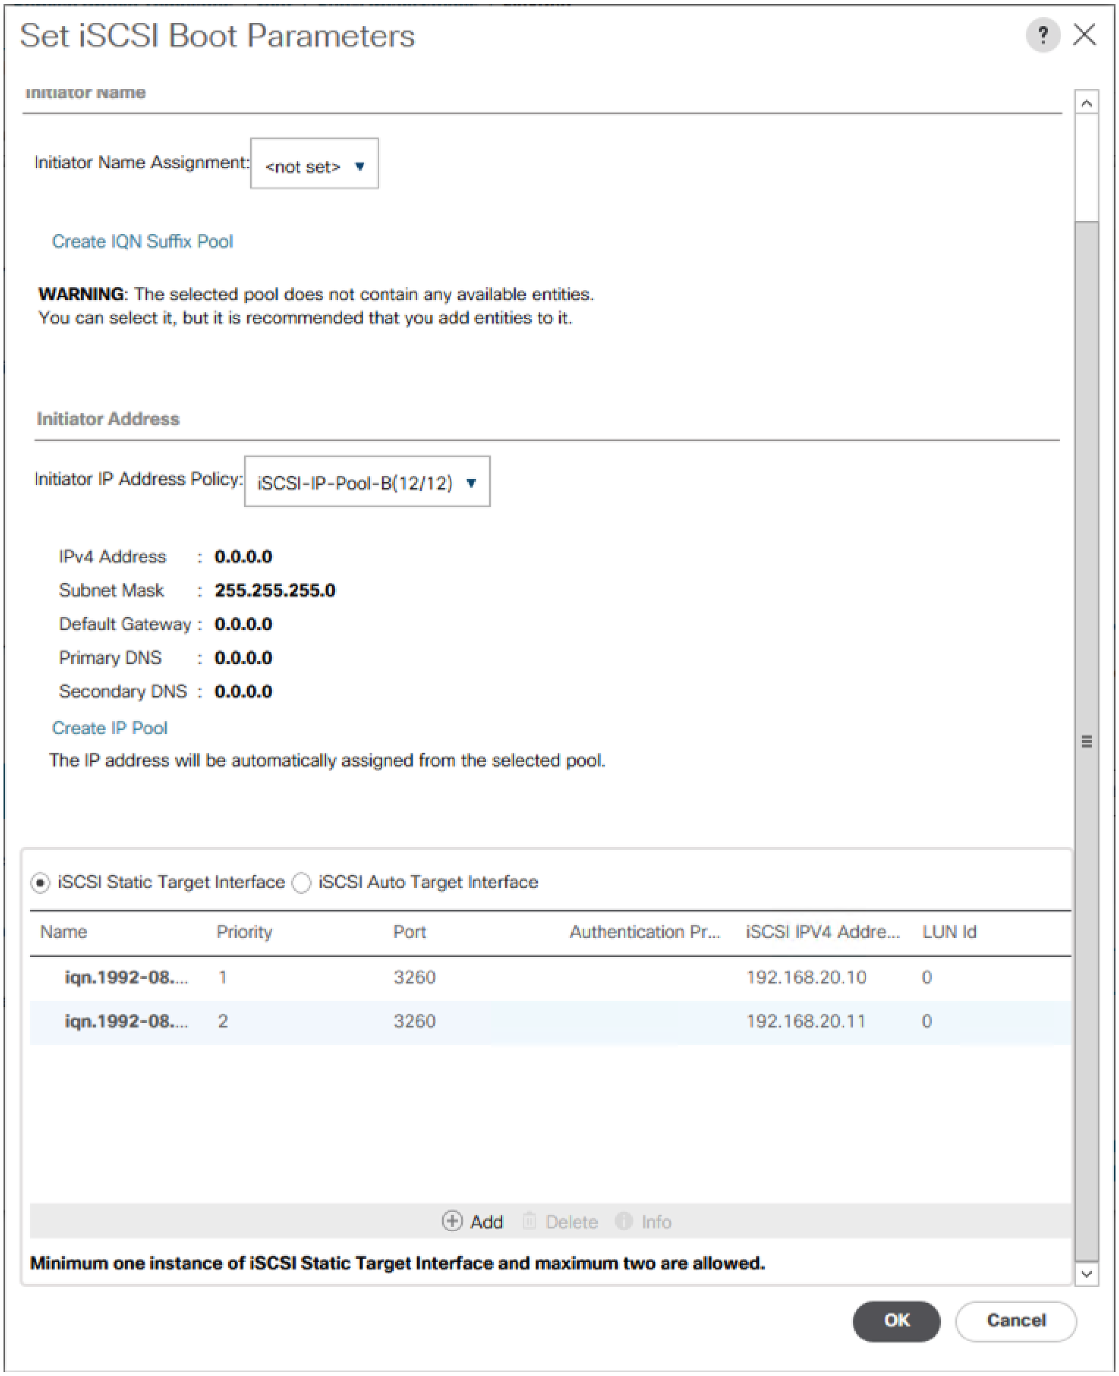

Set iSCSI Boot Parameters

Initiator Name

Initiator Name Assignment: <not set>

Create IQN Suffix Pool

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

Initiator Address

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

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

Create IP Pool

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

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...	LUN Id
iqn.1992-08....	1	3260		192.168.20.10	0
iqn.1992-08....	2	3260		192.168.20.11	0

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

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

32. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.

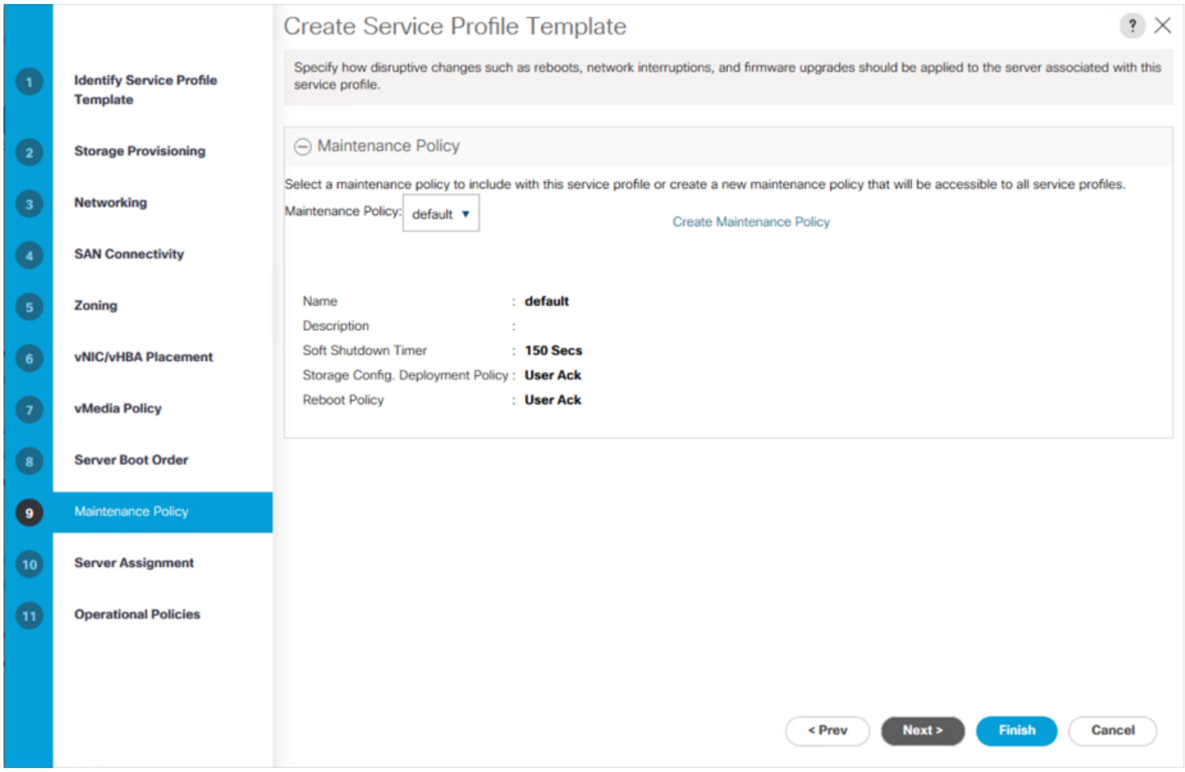

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

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

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, choose Infra-Pool.
2. Choose Down as the power state to be applied when the profile is associated with the server.
3. Optional: choose “UCS-B200M5” for the Server Pool Qualification to select only UCS B200M5 servers in the pool.
4. Expand Firmware Management at the bottom of the page and choose the default policy.

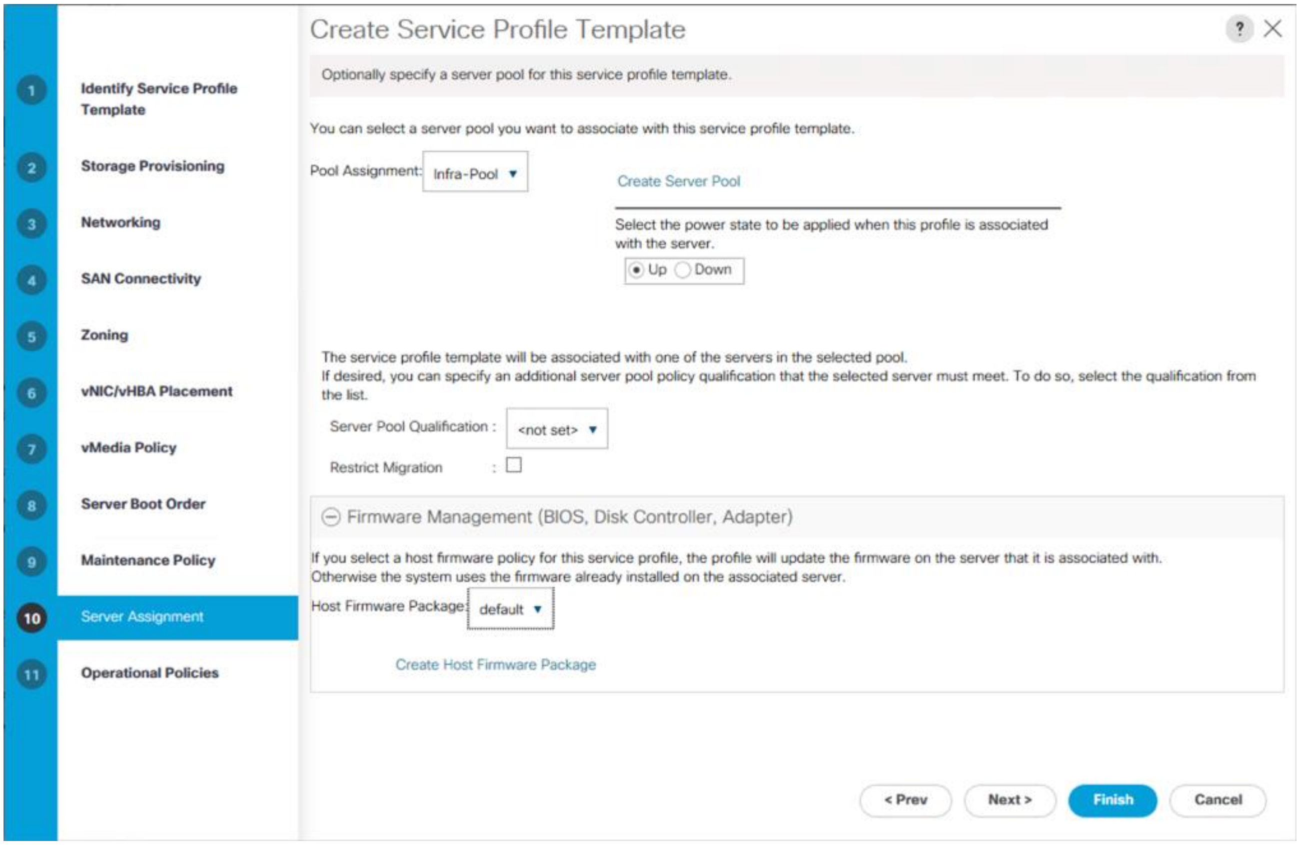

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: Infra-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 : <not set> ▾

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 ▾

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, choose Intel-VM-Host.

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

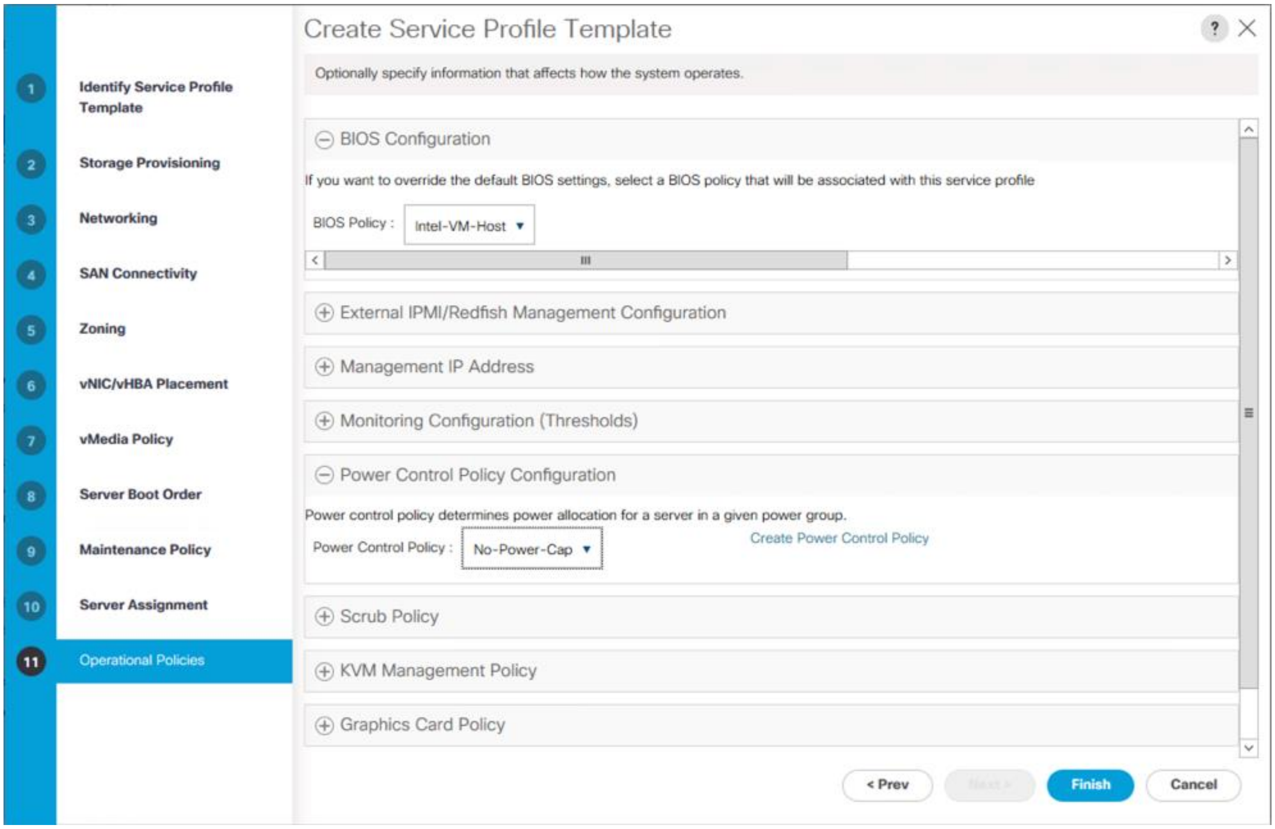

Create Service Profile Template

Optionally specify information that affects how the system operates.

BIOS Configuration

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

BIOS Policy : Intel-VM-Host

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 vMedia-Enabled Service Profile Template

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

1. Connect to Cisco UCS Manager and click Servers.

2. Choose Service Profile Templates > root > Sub-Organizations > FlexPod Organization > Service Template Intel-VM-Host-Infra-iSCSI-A.

3. Right-click Intel-VM-Host-Infra-iSCSI-A and click Create a Clone.

4. Name the clone Intel-VM-Host-Infra-iSCSI-A-vM and click OK then click OK again to create the clone.

5. Choose the newly created Intel-VM-Host-Infra-iSCSI-A-vM and choose the vMedia Policy tab.

6. Click Modify vMedia Policy.

7. Choose the ESXi-7.0-HTTP vMedia Policy and click OK.

8. Click OK to confirm.

Create Service Profiles

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

1. Connect to Cisco UCS Manager and click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > FlexPod Organization.
3. Right-click the appropriate vMedia-enabled template and choose Create Service Profiles from Template.
4. For Naming Prefix, enter VM-Host-Infra-0.
5. For Name Suffix Starting Number, enter 1.
6. For Number of Instances, enter 5.

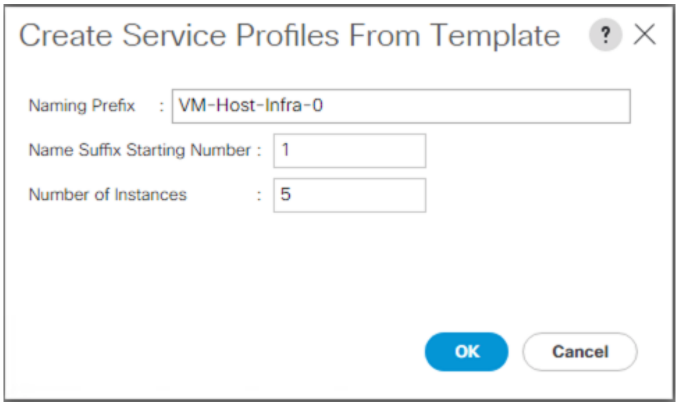

Create Service Profiles From Template

Naming Prefix : VM-Host-Infra-0

Name Suffix Starting Number : 1

Number of Instances : 5

Previously created server pool was for the 3 Cisco UCS B200s, the generated Service Profiles for the remaining Cisco UCS C220s were manually associated. Server pool qualifications and server associations should be created to address what is appropriate to the customers environment.

7. Click OK to create the service profiles.
8. Click OK in the confirmation message.
9. When VMware ESXi 7.0 has been installed on the hosts, the host Service Profiles can be bound to the corresponding non-vMedia-enabled Service Profile Template to remove the vMedia Mapping from the host.

Storage Configuration - Boot LUNs and Igroups

Gather Required Information

Table 21. iSCSI IQN for SVM

SVM Name	SVM Target IQN
Infra-SVM	

To obtain the iSCSI IQN, run **iscsi show** command on the storage cluster management interface.

Table 22. iSCSI IQN for SVM

Cisco UCS Service Profile Name	iSCSI IQN	Variable
VM-Host-Infra-01		<vm-host-infra-01-iqn>
VM-Host-Infra-02		<vm-host-infra-02-iqn>
VM-Host-Infra-03		<vm-host-infra-03-iqn>
VM-Host-Infra-04		<vm-host-infra-04-iqn>
VM-Host-Infra-05		<vm-host-infra-05-iqn>

To obtain the iSCSI vNIC IQN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the “iSCSI vNICs” tab on the top right. The “Initiator Name” is displayed at the top of the page under the “Service Profile Initiator Name.”

Create igroups

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

```
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-01 -protocol iscsi -ostype vmware -initiator <vm-host-infra-01-iqn>
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-02 -protocol iscsi -ostype vmware -initiator <vm-host-infra-02-iqn>
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-03 -protocol iscsi -ostype vmware -initiator <vm-host-infra-03-iqn>
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-04 -protocol iscsi -ostype vmware -initiator <vm-host-infra-04-iqn>
lun igroup create -vserver Infra-SVM -igroup vm-host-infra-05 -protocol iscsi -ostype vmware -initiator <vm-host-infra-05-iqn>

lun igroup show -protocol iscsi
```

Use the values listed in Table 20 and Table 21 for the IQN information.

Map Boot LUNs to igroups

To map the boot LUNs to igroups use the following commands:

```
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -igroup vm-host-infra-01 -lun-id
0
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/ VM-Host-Infra-02 -igroup vm-host-infra-02 -lun-id
0
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/ VM-Host-Infra-03 -igroup vm-host-infra-03 -lun-id

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/ VM-Host-Infra-04 -igroup vm-host-infra-04 -lun-id

lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/ VM-Host-Infra-05 -igroup vm-host-infra-05 -lun-id
0
```

Solution Deployment - VMware vSphere

VMware ESXi 7.0

This section provides detailed instructions for installing VMware ESXi 7.0 in a FlexPod environment. After the procedures are completed, three booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download ESXi 7.0 from VMware

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

1. Click the following link: Cisco Custom ISO for UCS 4.1.2a.

2. You will need a user id and password on vmware.com to download this software.

3. Download the .iso file.

Log into Cisco UCS 6454 Fabric Interconnect

Cisco UCS Manager

The Cisco UCS IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log into the Cisco UCS environment, follow these steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.

2. Click the Launch UCS Manager link to launch the HTML 5 UCS Manager GUI.

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

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

5. To log in to Cisco UCS Manager, click Login.

6. From the main menu, click Servers.

7. Choose Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-01.

8. In the Actions pane, click KVM Console.

9. Follow the prompts to launch the HTML5 KVM console.

10. Choose Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-02.

11. In the Actions pane, click KVM Console.

12. Follow the prompts to launch the HTML5 KVM console.

13. Choose Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-03.

14. In the Actions pane, click KVM Console.

15. Follow the prompts to launch the HTML5 KVM console.

Set Up VMware ESXi Installation

All ESXi Hosts

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

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

1. In the KVM window, click Virtual Media.

2. Choose Activate Virtual Devices.

3. If prompted to accept an Unencrypted KVM session, accept as necessary.

4. Click Virtual Media and choose Map CD/DVD.

5. Browse to the ESXi installer ISO image file and click Open.

6. Click Map Device.

7. Click the KVM Console tab to monitor the server boot.

Install ESXi

All ESXi Hosts

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

1. Boot the server by selecting Boot Server in the KVM and click OK, then click OK again.

2. On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.

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

3. After the installer is finished loading, press Enter to continue with the installation.

4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

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

5. Choose the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

6. Choose the appropriate keyboard layout and press Enter.

7. Enter and confirm the root password and press Enter.

8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.

9. After the installation is complete, press Enter to reboot the server.

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

10. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, follow these steps on each ESXi host:

All ESXi Hosts

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

1. After the server has finished rebooting, in the UCS KVM console, press F2 to customize VMware ESXi.

2. Log in as root, enter the corresponding password, and press Enter to log in.

3. Use the down arrow key to choose Troubleshooting Options and press Enter.

4. Choose Enable ESXi Shell and press Enter.

5. Choose Enable SSH and press Enter.

6. Press Esc to exit the Troubleshooting Options menu.

7. Choose the Configure Management Network option and press Enter.

8. Choose Network Adapters and press Enter.

9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field. If the numbers do not match, note the mapping of vmnic ports to vNIC ports for later use.

10. Using the spacebar, choose vmnic1.

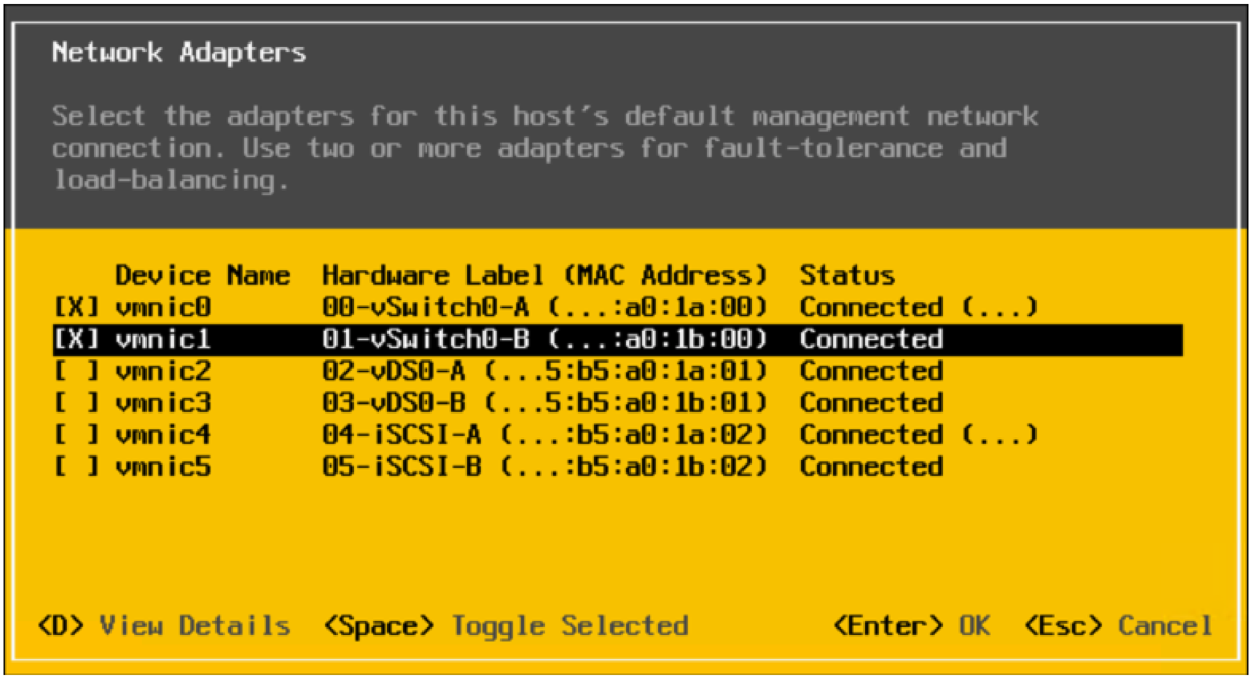

```
Network Adapters

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-vSwitch0-A (...:a0:1a:00)  Connected (...)
[X] vmnic1       01-vSwitch0-B (...:a0:1b:00)  Connected
[ ] vmnic2       02-vDS0-A (...5:b5:a0:1a:01)  Connected
[ ] vmnic3       03-vDS0-B (...5:b5:a0:1b:01)  Connected
[ ] vmnic4       04-iSCSI-A (...:b5:a0:1a:02)  Connected (...)
[ ] vmnic5       05-iSCSI-B (...:b5:a0:1b: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 upcoming procedure accordingly.

11. Press Enter.

12. Choose the VLAN (Optional) option and press Enter.

13. Enter the <ib-mgmt-vlan-id> and press Enter.

14. Choose IPv4 Configuration and press Enter.

15. Choose the “Set static IPv4 address and network configuration” option by using the arrow keys and space bar.

16. Move to the IPv4 Address field and enter the IP address for managing the ESXi host.

17. Move to the Subnet Mask field and enter the subnet mask for the ESXi host.

18. Move to the Default Gateway field and enter the default gateway for the ESXi host.

19. Press Enter to accept the changes to the IP configuration.

20. Choose the IPv6 Configuration option and press Enter.

21. Using the spacebar, choose Disable IPv6 (restart required) and press Enter.

22. Choose the DNS Configuration option and press Enter.

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

23. Using the spacebar, choose "Use the following DNS server addresses and hostname:"

24. Move to the Primary DNS Server field and enter the IP address of the primary DNS server.

25. Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.

26. Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.

27. Press Enter to accept the changes to the DNS configuration.

28. Press Esc to exit the Configure Management Network submenu.

29. Press Y to confirm the changes and reboot the ESXi host.

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

All ESXi Hosts

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

1. From the ESXi console menu main screen, type Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of Static Macros.

2. Log in as root.

3. Type `esxcfg-vmknic –l` to get a detailed listing of interface vmk0. vmk0 should be a part of the "Management Network" port group. Note the IP address and netmask of vmk0.

4. To remove vmk0, type `esxcfg-vmknic –d "Management Network"`.

5. To re-add vmk0 with a random MAC address, type `esxcfg-vmknic –a –i <vmk0-ip> -n <vmk0-netmask> "Management Network"`.

6. Verify vmk0 has been re-added with a random MAC address by typing `esxcfg-vmknic –l`.

7. Tag vmk0 as the management interface by typing `esxcli network ip interface tag add -i vmk0 -t Management`.

8. When vmk0 was re-added, if a message popped up saying vmk1 was marked as the management interface, type `esxcli network ip interface tag remove -i vmk1 -t Management`.

9. Type `exit` to log out of the command line interface.

10. Type Ctrl-Alt-F2 to return to the ESXi console menu interface.

Install VMware and Cisco VIC Drivers for the ESXi Host

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

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

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

All ESXi Hosts

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

1. Using an SCP program such as WinSCP, copy the two offline bundles referenced above to the /tmp directory on each ESXi host.

2. Using a SSH tool such as PuTTY, ssh to each VMware ESXi host. Log in as root with the root password.

3. Type `cd /tmp`.

4. Run the following commands on each host:

```
esxcli software component apply -d /tmp/ucs-tool-esxi_1.1.5-1OEM.zip

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

reboot
```

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

```
esxcli software component list | grep ucs

esxcli software vib list | grep NetApp
```

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

The process for the first VMware host will be setup directly using the ESXi vSphere Web Client for that host. Subsequent hosts will be added to vCenter and configured with slightly differing steps through the vCenter client.

ESXi Host VM-Host-Infra-01

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

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.

2. Enter root for the User name.

3. Enter the root password.

4. Click Login to connect.

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

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

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

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

In these steps, the vMotion VMkernel will not be created as it will later be setup on the vDS after vCenter is in place. If vMotion is not going to be handled within the vDS and has been set to be in vSwitch0 by the vNIC templates created during the UCS setup, the required portgroup and VMkernel should be created in this section.

1. From the Host Client Navigator, choose Networking.
2. In the center pane, choose the Virtual switches tab.
3. Highlight the vSwitch0 line.
4. Choose Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, choose vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.
9. Click Save.
10. Choose Networking, then choose the Port groups tab.
11. In the center pane, right-click VM Network and choose Edit settings.
12. Name the port group Site1-IB Network and enter <site1-ib-vlan-id> in the VLAN ID field.
13. Click Save to finalize the edits for the Site1-IB Network.
14. Still within the Port groups tab select the Add port group option to create a Common-Services port group that will be used for vCenter and other virtual infrastructure components.
15. Specify the name and VLAN ID.
16. Click Add to create the port group.
17. At the top, choose the VMkernel NICs tab.

18. Click Add VMkernel NIC.

19. For New port group, enter VMkernel-Infra-NFS.

20. For Virtual switch, choose vSwitch0.

21. Enter <infra-nfs-vlan-id> for the VLAN ID.

22. Change the MTU to 9000.

23. Choose Static IPv4 settings and expand IPv4 settings.

24. Enter the ESXi host Infrastructure NFS IP address and netmask.

25. Leave TCP/IP stack set at Default TCP/IP stack and do not choose any of the Services.

26. Click Create.

27. Choose the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:

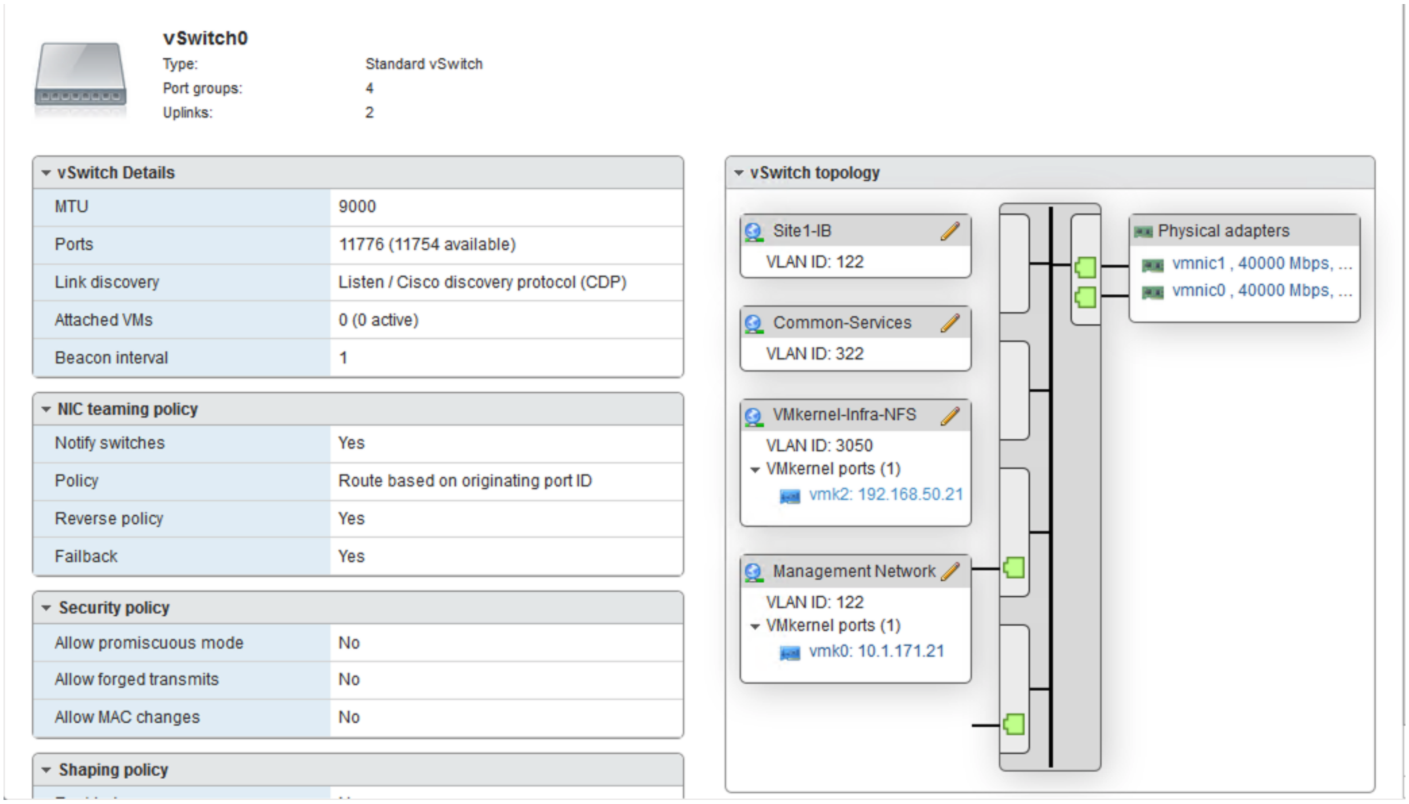

vSwitch0

Type:	Standard vSwitch
Port groups:	4
Uplinks:	2

▾ vSwitch Details

MTU	9000
Ports	11776 (11754 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

▾ NIC teaming policy

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

▾ Security policy

Allow promiscuous mode	No
Allow forged transmits	No
Allow MAC changes	No

▾ Shaping policy

▾ vSwitch topology

Site1-IB
VLAN ID: 122

Common-Services
VLAN ID: 322

VMkernel-Infra-NFS
VLAN ID: 3050
▾ VMkernel ports (1)
vmk2: 192.168.50.21

Management Network
VLAN ID: 122
▾ VMkernel ports (1)
vmk0: 10.1.171.21

Physical adapters
vmnic1 , 40000 Mbps, ...
vmnic0 , 40000 Mbps, ...

28. In the center pane, choose the Virtual switches tab.

29. Highlight the iScsiBootvSwitch line.

30. Choose Edit settings.

31. Change the MTU to 9000.

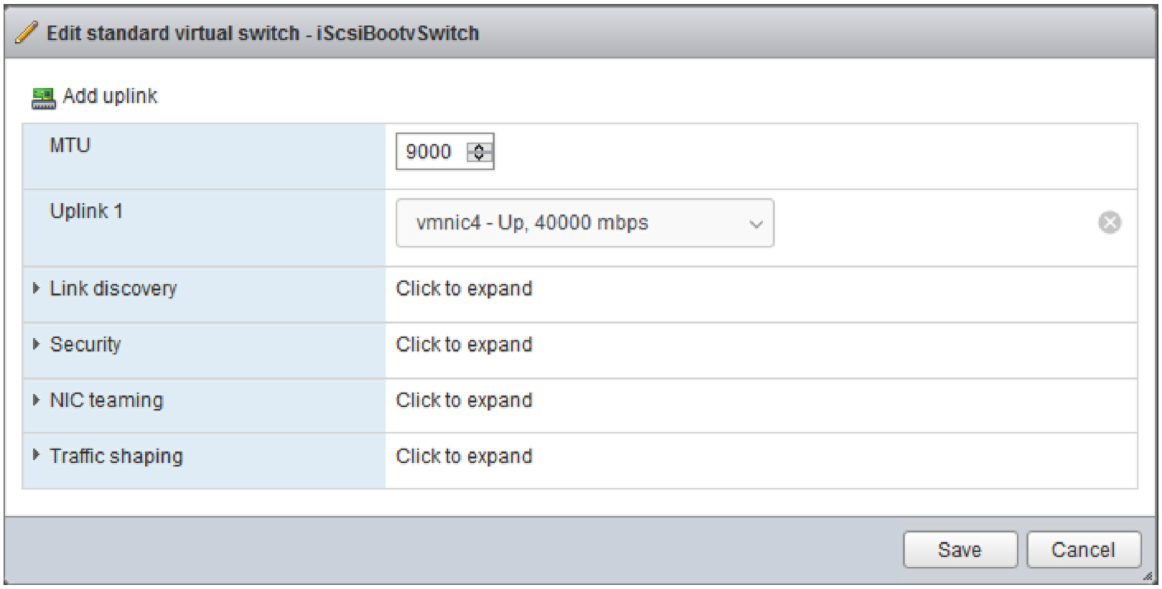

Edit standard virtual switch - iScsiBootvSwitch

Add uplink

MTU	9000
Uplink 1	vmnic4 - Up, 40000 mbps
Link discovery	Click to expand
Security	Click to expand
NIC teaming	Click to expand
Traffic shaping	Click to expand

32. Click Save to save the changes to iScsiBootvSwitch.

33. Choose Add standard virtual switch.

34. Name the switch vSwitch1.

35. Change the MTU to 9000.

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

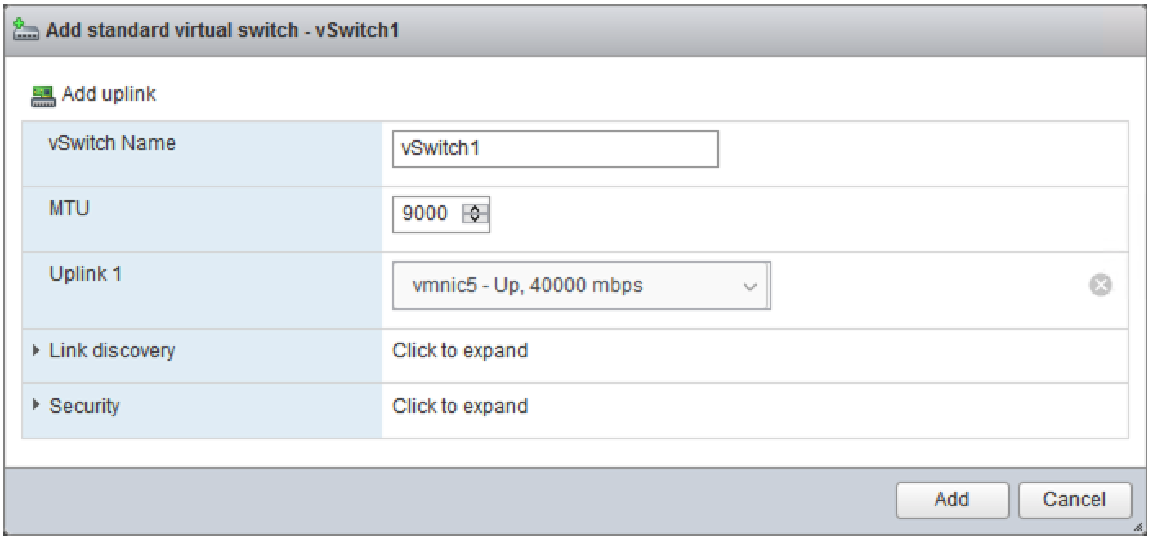

Add standard virtual switch - vSwitch1

Add uplink

vSwitch Name	vSwitch1
MTU	9000
Uplink 1	vmnic5 - Up, 40000 mbps
Link discovery	Click to expand
Security	Click to expand

37. Choose Add to add vSwitch1.

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

39. Highlight the iScsiBootPG line.

40. Choose Edit settings.

41. Change the MTU to 9000.

42. Expand IPv4 Settings and enter a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.

> It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments in Cisco UCS.

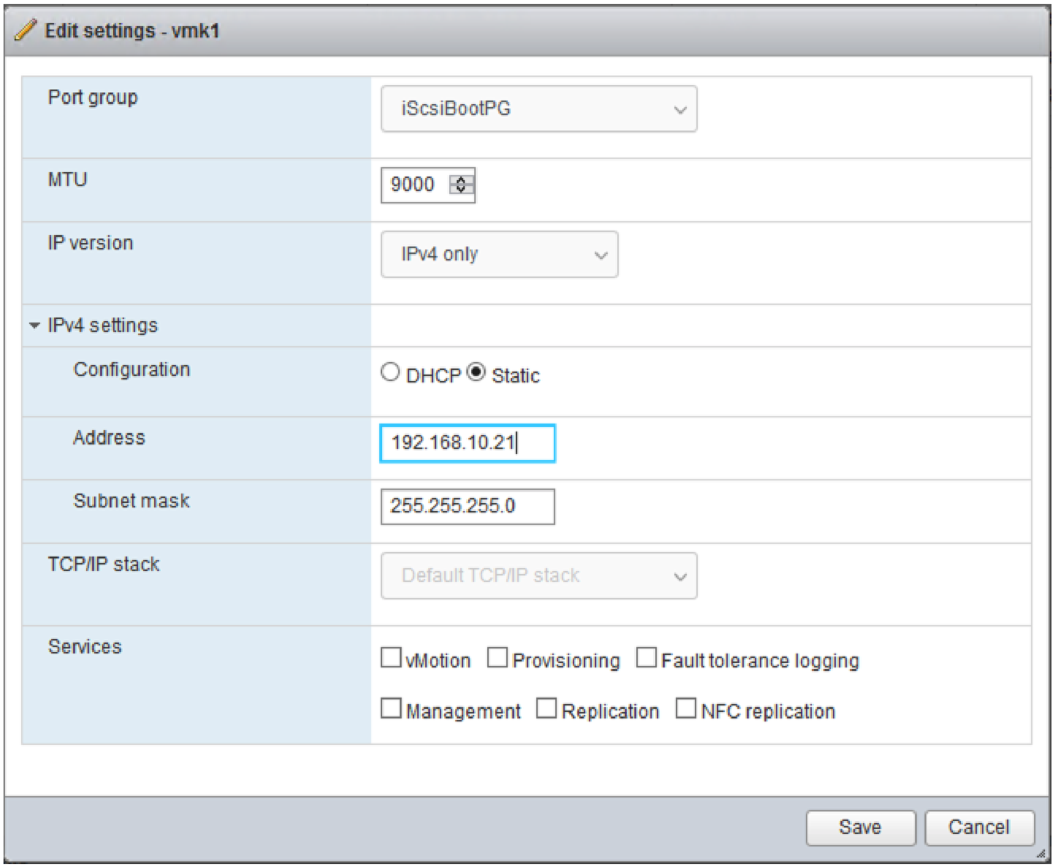

Edit settings - vmk1

Port group	iScsiBootPG
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	○ DHCP ● Static
Address	192.168.10.21
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	□ vMotion □ Provisioning □ Fault tolerance logging □ Management □ Replication □ NFC replication

43. Click Save to save the changes to iScsiBootPG VMkernel NIC.

44. Choose Add VMkernel NIC.

45. For New port group, enter iScsiBootPG-B.

46. For Virtual switch, use the pull-down to choose vSwitch1.

47. Change the MTU to 9000.

48. For IPv4 settings, choose Static.

49. Expand IPv4 Settings and enter a unique IP address in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B.

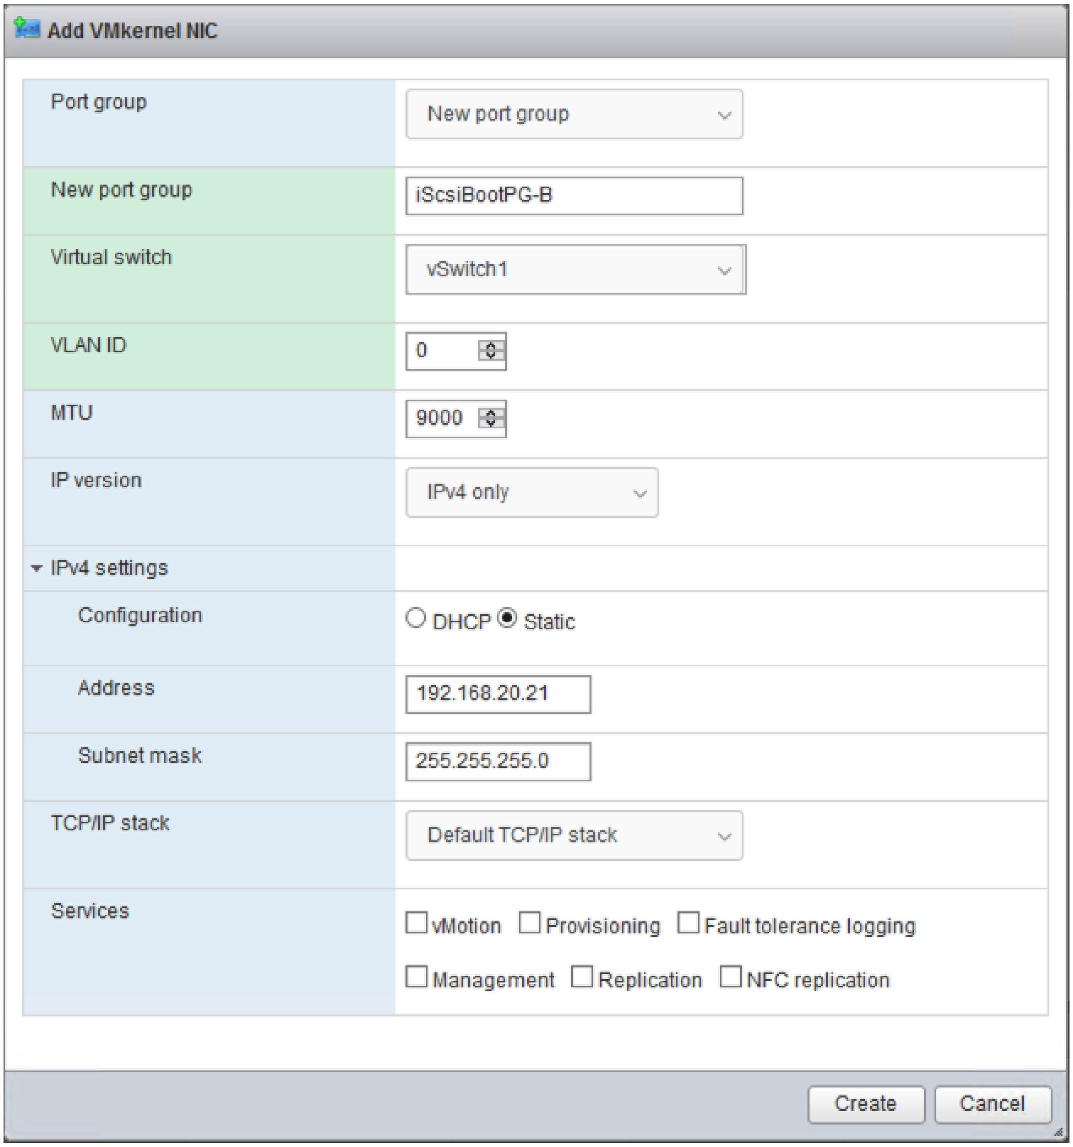

Add VMkernel NIC

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	vSwitch1
VLAN ID	0
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	○ DHCP ● Static
Address	192.168.20.21
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	□ vMotion □ Provisioning □ Fault tolerance logging □ Management □ Replication □ NFC replication

50. Click Create to complete creating the VMkernel NIC.

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

52. Highlight the iScsiBootPG line.

53. Choose Edit settings.

54. Change the Name to iScsiBootPG-A.

55. Click Save to complete editing the port group name.

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

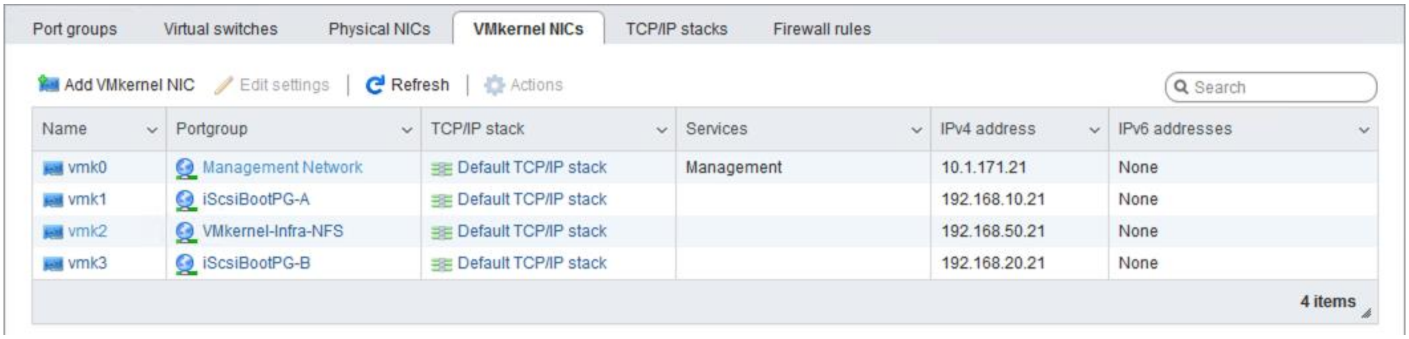

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.171.21	None
vmk1	iScsiBootPG-A	Default TCP/IP stack		192.168.10.21	None
vmk2	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.21	None
vmk3	iScsiBootPG-B	Default TCP/IP stack		192.168.20.21	None

Mount Required Datastores

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, choose Storage.

2. In the center pane, choose the Datastores tab.

3. In the center pane, choose New Datastore to add a new datastore.

4. In the New datastore popup, choose Mount NFS datastore and click Next.

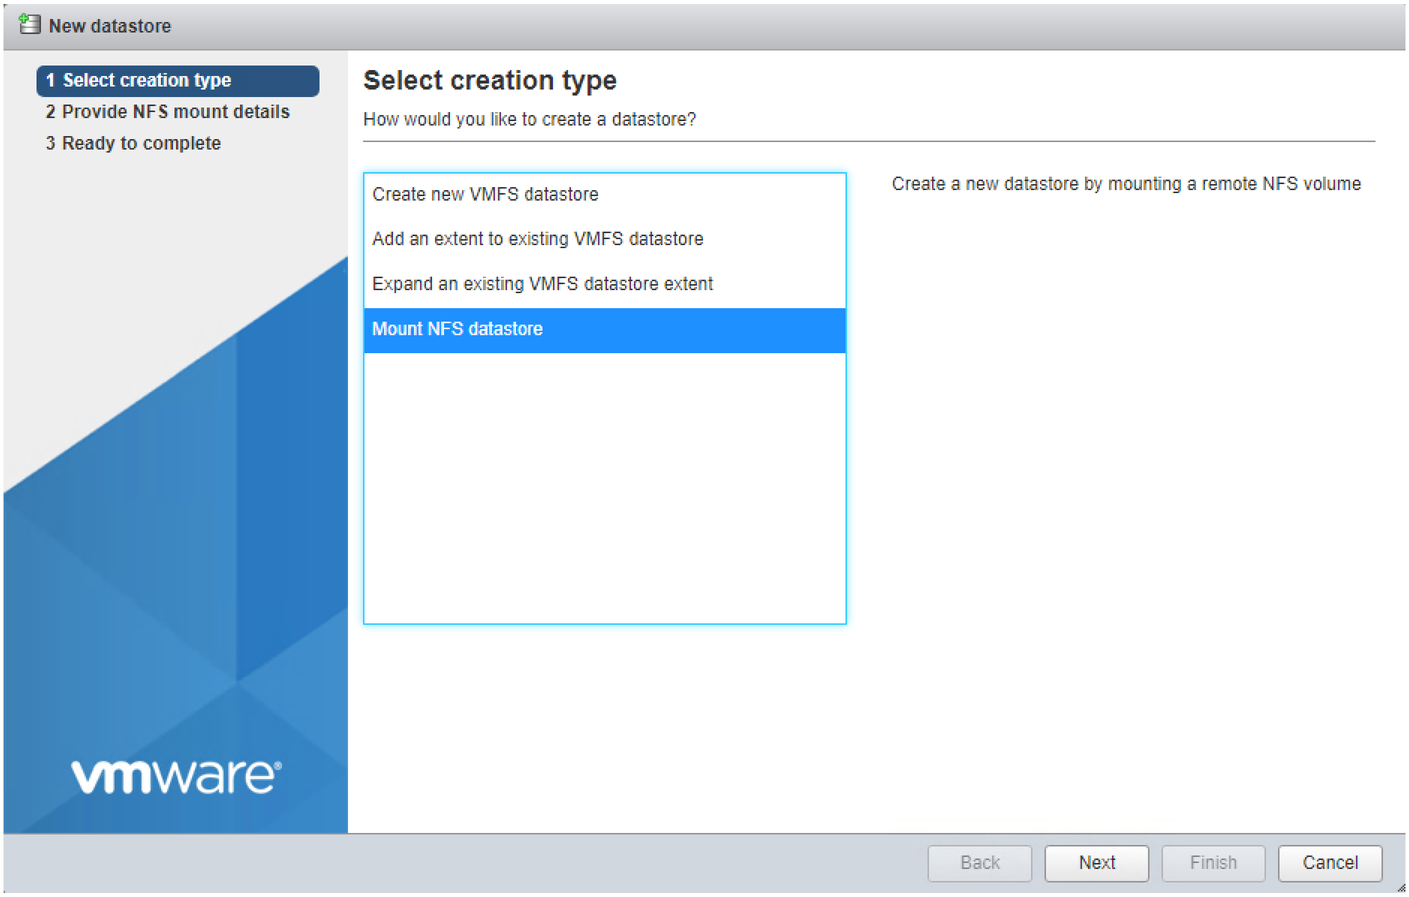

New datastore

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

Select creation type

How would you like to create a datastore?

Create new VMFS datastore
Add an extent to existing VMFS datastore
Expand an existing VMFS datastore extent
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-lif-01 LIF for the NFS server. Input /infra_datastore_1 for the NFS share. Leave the NFS version set at NFS 3. Click Next.

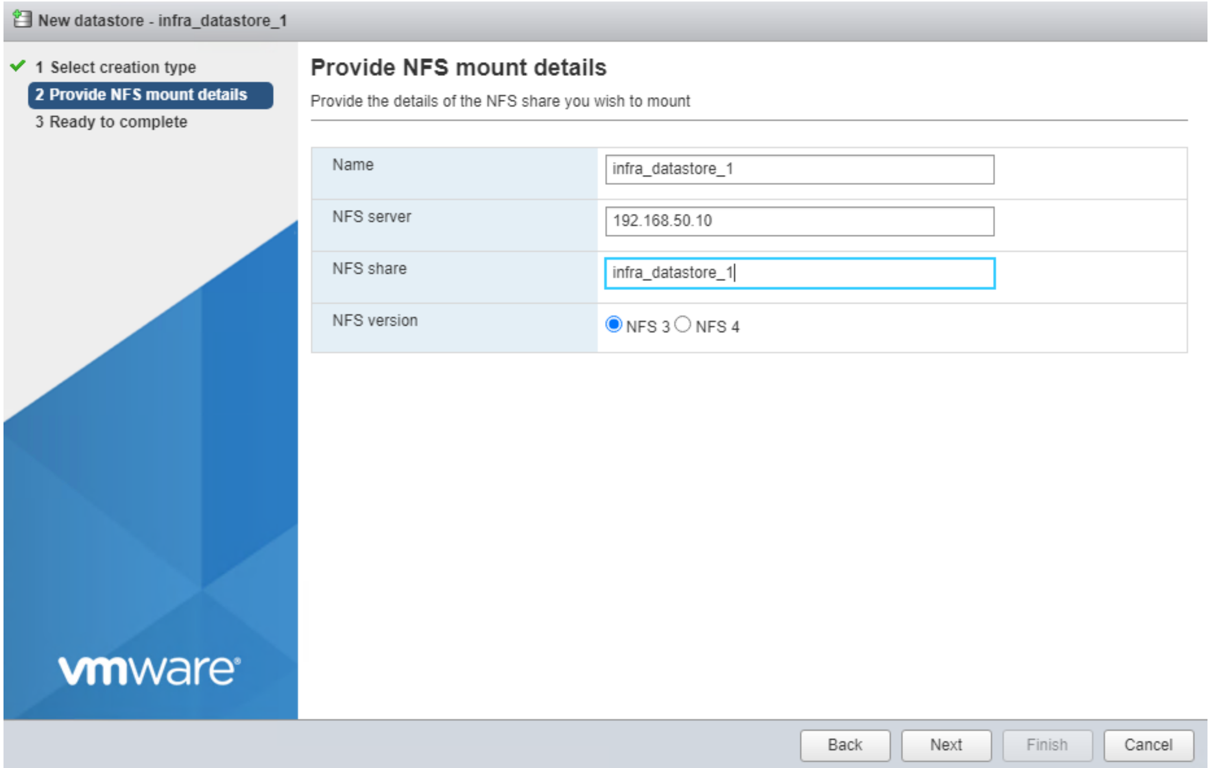

New datastore - infra_datastore_1

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

Provide NFS mount details

Provide the details of the NFS share you wish to mount

Name	infra_datastore_1
NFS server	192.168.50.10
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, choose New Datastore to add a new datastore.

8. In the New datastore popup, choose Mount NFS datastore and click Next.

9. Input infra_swap for the datastore name. Input the IP address for the nfs-lif-01 LIF for the NFS server. Input /infra_swap for the NFS share. Leave the NFS version set at NFS 3. Click Next.

10. Click Finish. The datastore should now appear in the datastore list.

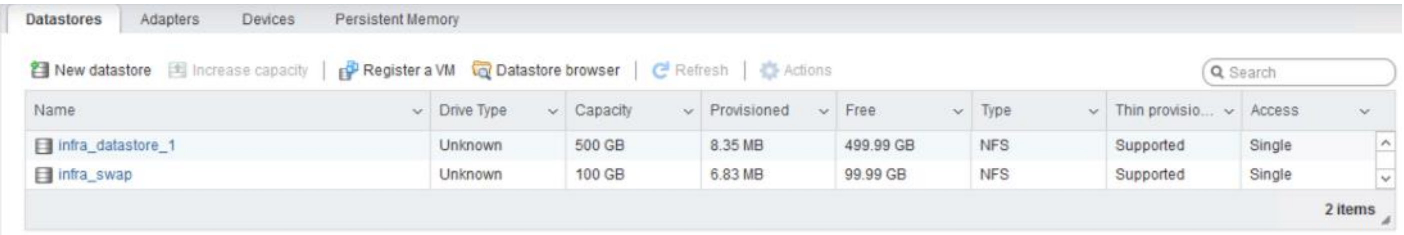

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisio...	Access
infra_datastore_1	Unknown	500 GB	8.35 MB	499.99 GB	NFS	Supported	Single
infra_swap	Unknown	100 GB	6.83 MB	99.99 GB	NFS	Supported	Single

2 items

Configure NTP on First ESXi Host

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, choose Manage.

2. In the center pane, choose System > Time & Date.

3. Click Edit NTP settings.

4. Make sure “Manually configure the date and time on this host and enter the approximate date and time.

5. Select Use Network Time Protocol (enable NTP client).

6. Use the drop-down list to choose Start and stop with host.

7. Enter the two Nexus switch NTP addresses in the NTP server(s) box separated by a comma.

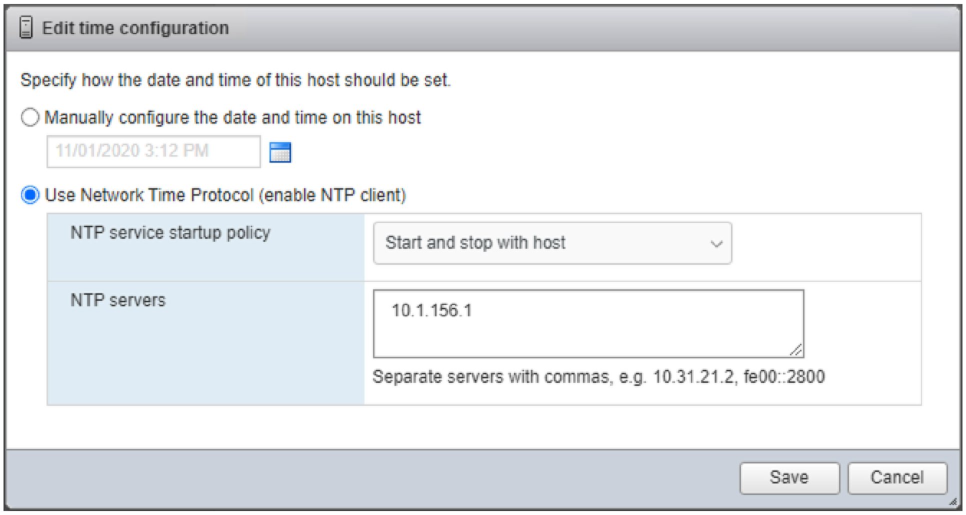

Edit time configuration

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

○ Manually configure the date and time on this host

11/01/2020 3:12 PM

● Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 10.1.156.1

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

8. Click Save to save the configuration changes.

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

Configure ESXi Host Swap

ESXi Host VM-Host-Infra-01

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

1. From the Host Client, choose Manage.

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

3. Click Edit settings.

4. Use the drop-down list to choose infra_swap. Leave all other settings unchanged.

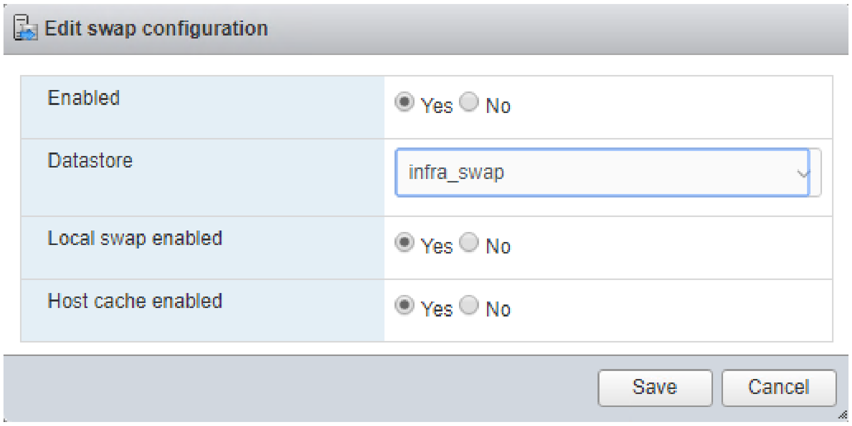

Edit swap configuration

Enabled	● Yes ○ No
Datastore	infra_swap
Local swap enabled	● Yes ○ No
Host cache enabled	● Yes ○ No

5. Click Save to save the configuration changes.

Configure Host Power Policy

ESXi Host VM-Host-Infra-01

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

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

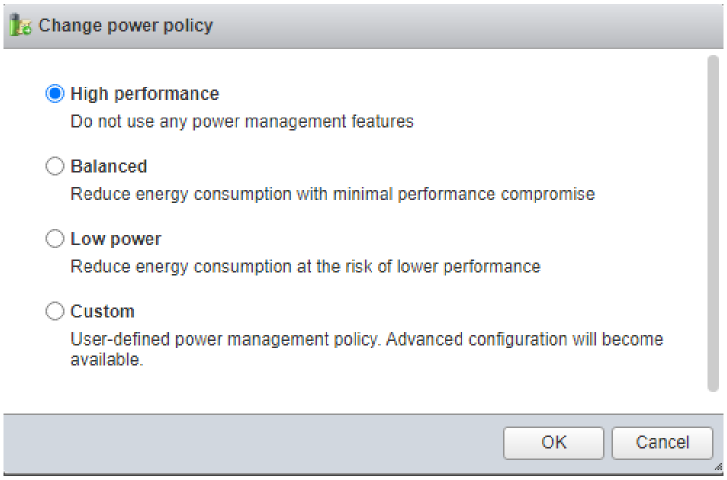

Change power policy

● High performance
Do not use any power management features

○ Balanced
Reduce energy consumption with minimal performance compromise

○ Low power
Reduce energy consumption at the risk of lower performance

○ Custom
User-defined power management policy. Advanced configuration will become available.

VMware vCenter 7.0d

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

Build the VMware vCenter Server Appliance

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

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

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

2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012 and above).

3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click `installer.exe`. The vCenter Server Appliance Installer wizard appears.

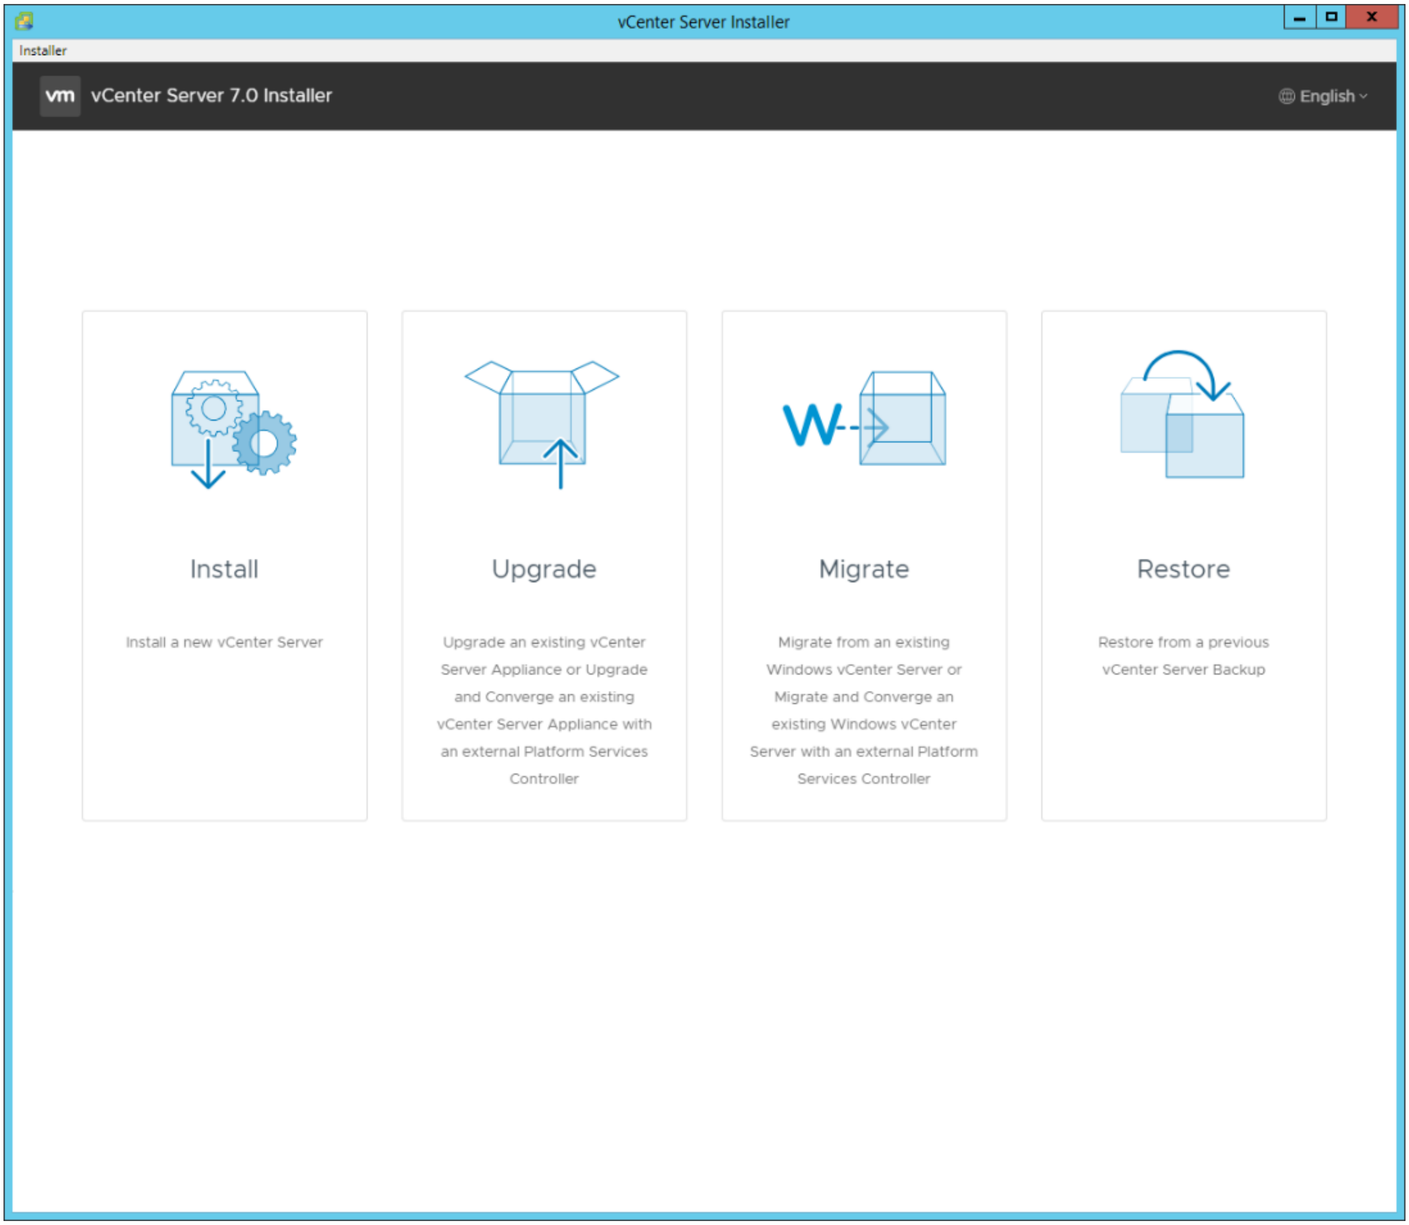

Install

Install a new vCenter Server

Upgrade

Upgrade an existing vCenter Server Appliance or Upgrade and Converge an existing vCenter Server Appliance with an external Platform Services Controller

Migrate

Migrate from an existing Windows vCenter Server or Migrate and Converge an existing Windows vCenter Server with an external Platform Services Controller

Restore

Restore from a previous vCenter Server 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.

7. In the “vCenter Server deployment target” window, enter the host name or IP address of the first ESXi host, User name (root) and Password. Click NEXT.

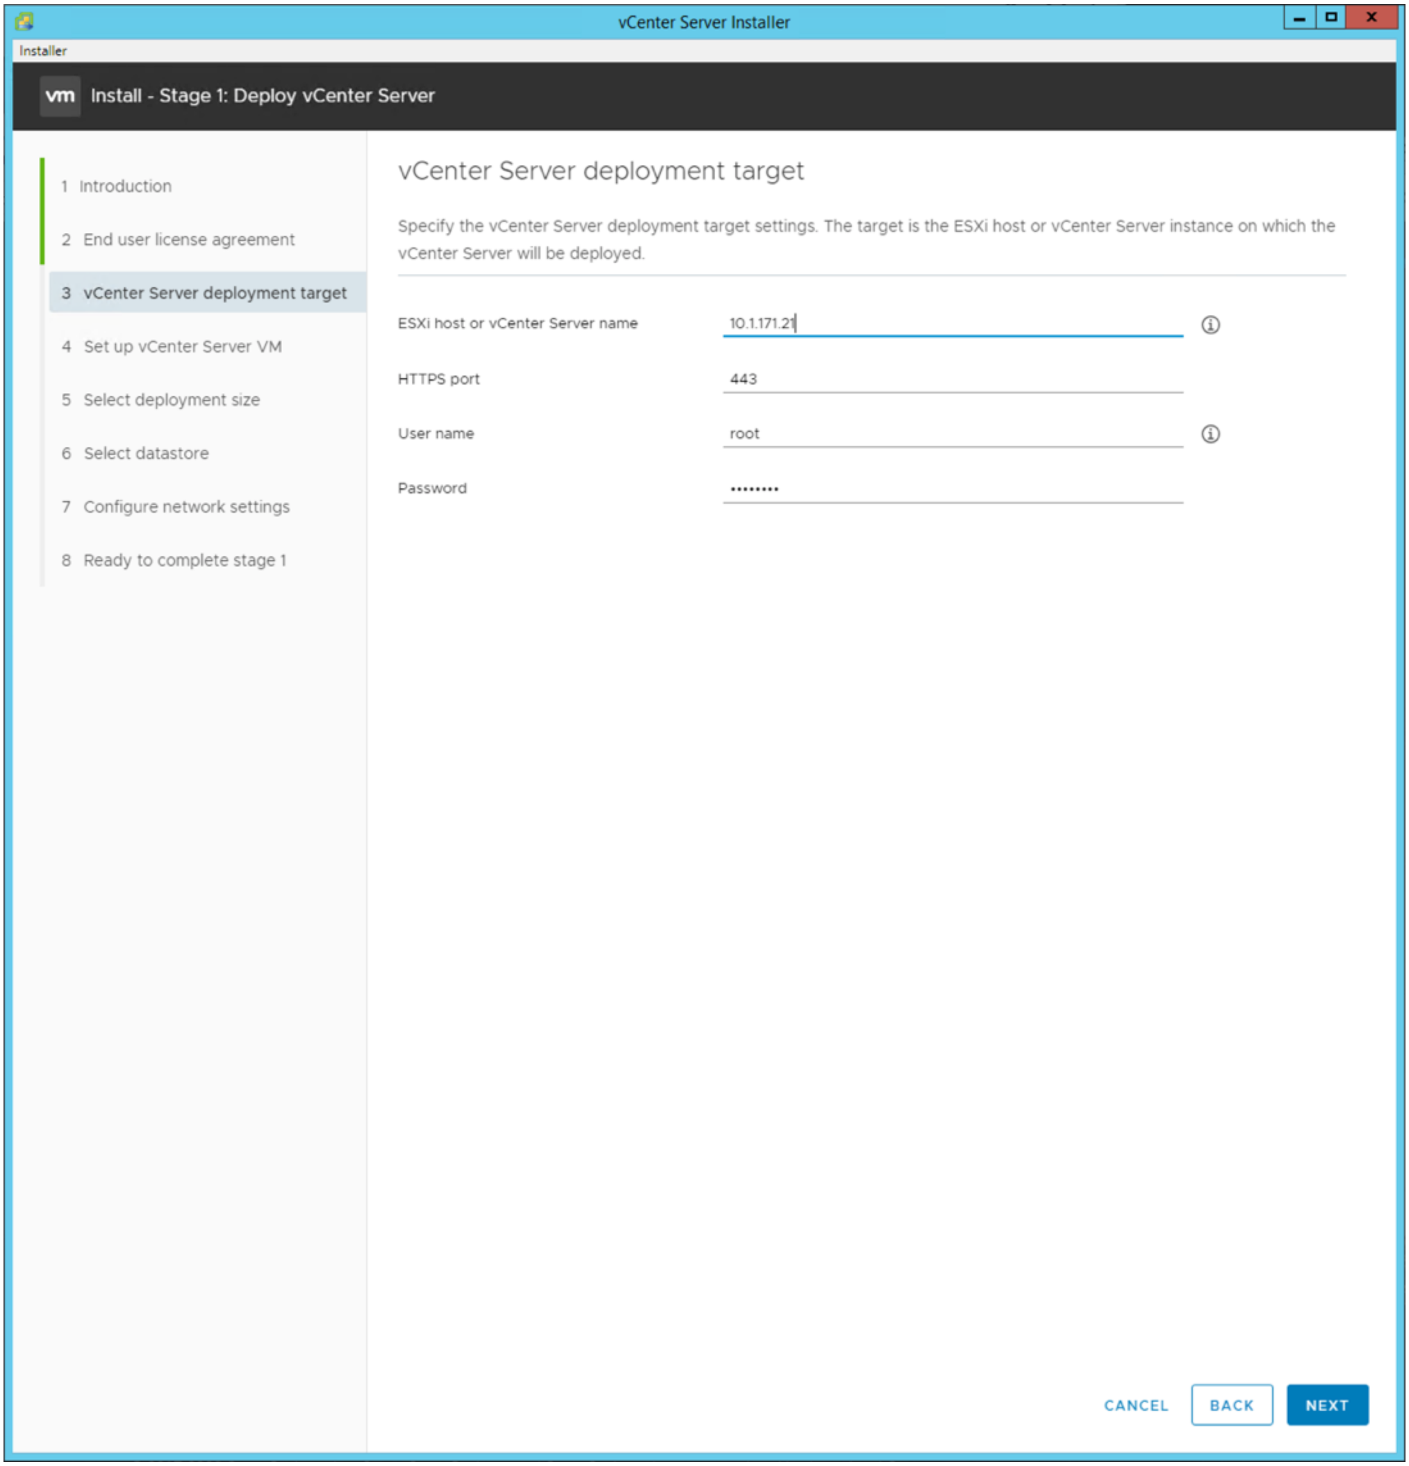

vCenter Server Installer

Installer

Install - Stage 1: Deploy vCenter Server

1 Introduction
2 End user license agreement
3 vCenter Server deployment target
4 Set up vCenter Server VM
5 Select deployment size
6 Select datastore
7 Configure network settings
8 Ready to complete stage 1

vCenter Server deployment target

Specify the vCenter Server deployment target settings. The target is the ESXi host or vCenter Server instance on which the vCenter Server will be deployed.

ESXi host or vCenter Server name: 10.1.171.21

HTTPS port: 443

User name: root

Password: ••••••••

8. Click YES to accept the certificate.

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

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

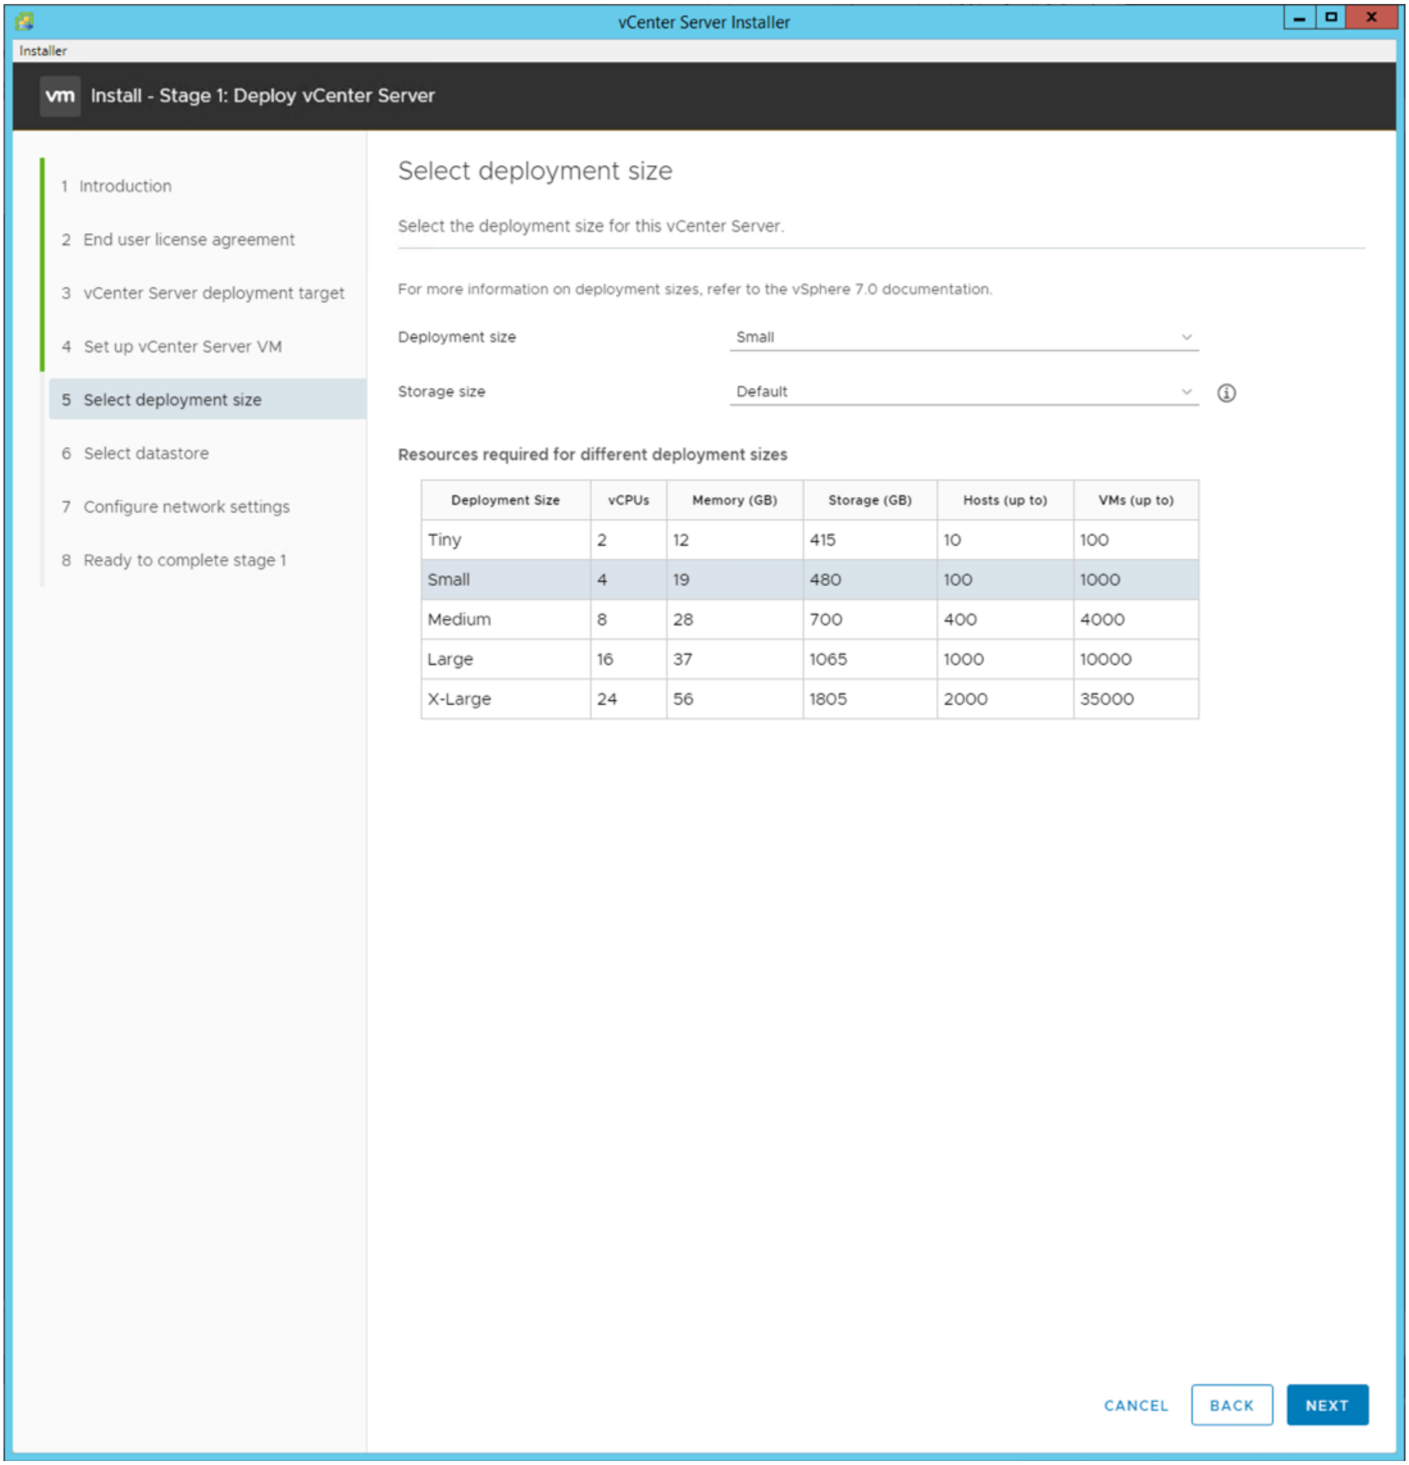

vCenter Server Installer
Installer
vm Install - Stage 1: Deploy vCenter Server
1 Introduction
2 End user license agreement
3 vCenter Server deployment target
4 Set up vCenter Server VM
5 Select deployment size
6 Select datastore
7 Configure network settings
8 Ready to complete stage 1

Select deployment size

Select the deployment size for this vCenter Server.

For more information on deployment sizes, refer to the vSphere 7.0 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	12	415	10	100
Small	4	19	480	100	1000
Medium	8	28	700	400	4000
Large	16	37	1065	1000	10000
X-Large	24	56	1805	2000	35000

11. Choose infra_datastore_1 for storage. Click NEXT.

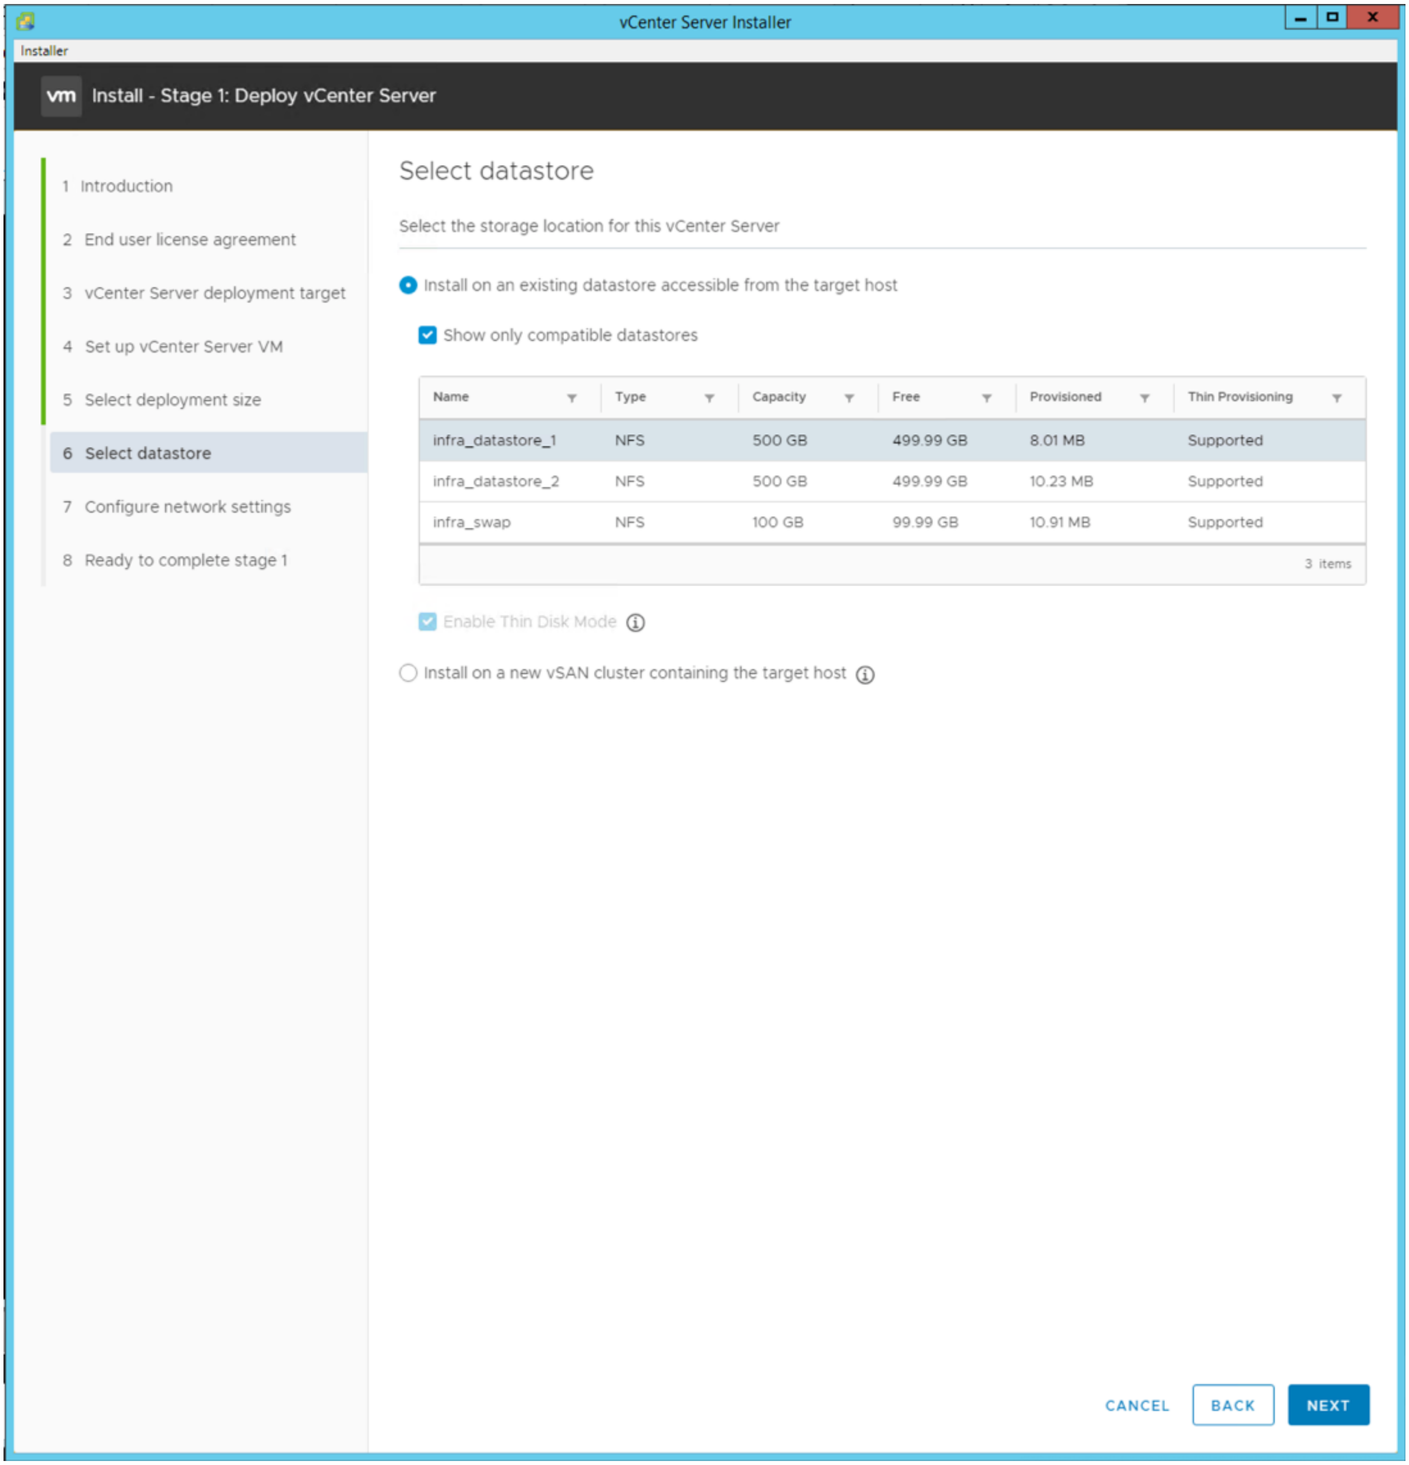

vCenter Server Installer
Installer
vm Install - Stage 1: Deploy vCenter Server
1 Introduction
2 End user license agreement
3 vCenter Server deployment target
4 Set up vCenter Server VM
5 Select deployment size
6 Select datastore
7 Configure network settings
8 Ready to complete stage 1
CANCEL
BACK
NEXT

Select datastore

Select the storage location for this vCenter Server

● Install on an existing datastore accessible from the target host

☑ Show only compatible datastores

Name	Type	Capacity	Free	Provisioned	Thin Provisioning
infra_datastore_1	NFS	500 GB	499.99 GB	8.01 MB	Supported
infra_datastore_2	NFS	500 GB	499.99 GB	10.23 MB	Supported
infra_swap	NFS	100 GB	99.99 GB	10.91 MB	Supported

3 items

☑ Enable Thin Disk Mode

○ Install on a new vSAN cluster containing the target host

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

    a. Choose a Network: Common-Services Network.

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

b. IP version: IPV4

c. IP assignment: static

d. FQDN: <vcenter-fqdn>

e. IP address: <vcenter-ip>

f. Subnet mask or prefix length: <vcenter-subnet-mask>

g. Default gateway: <vcenter-gateway>

h. DNS Servers: <dns-server>

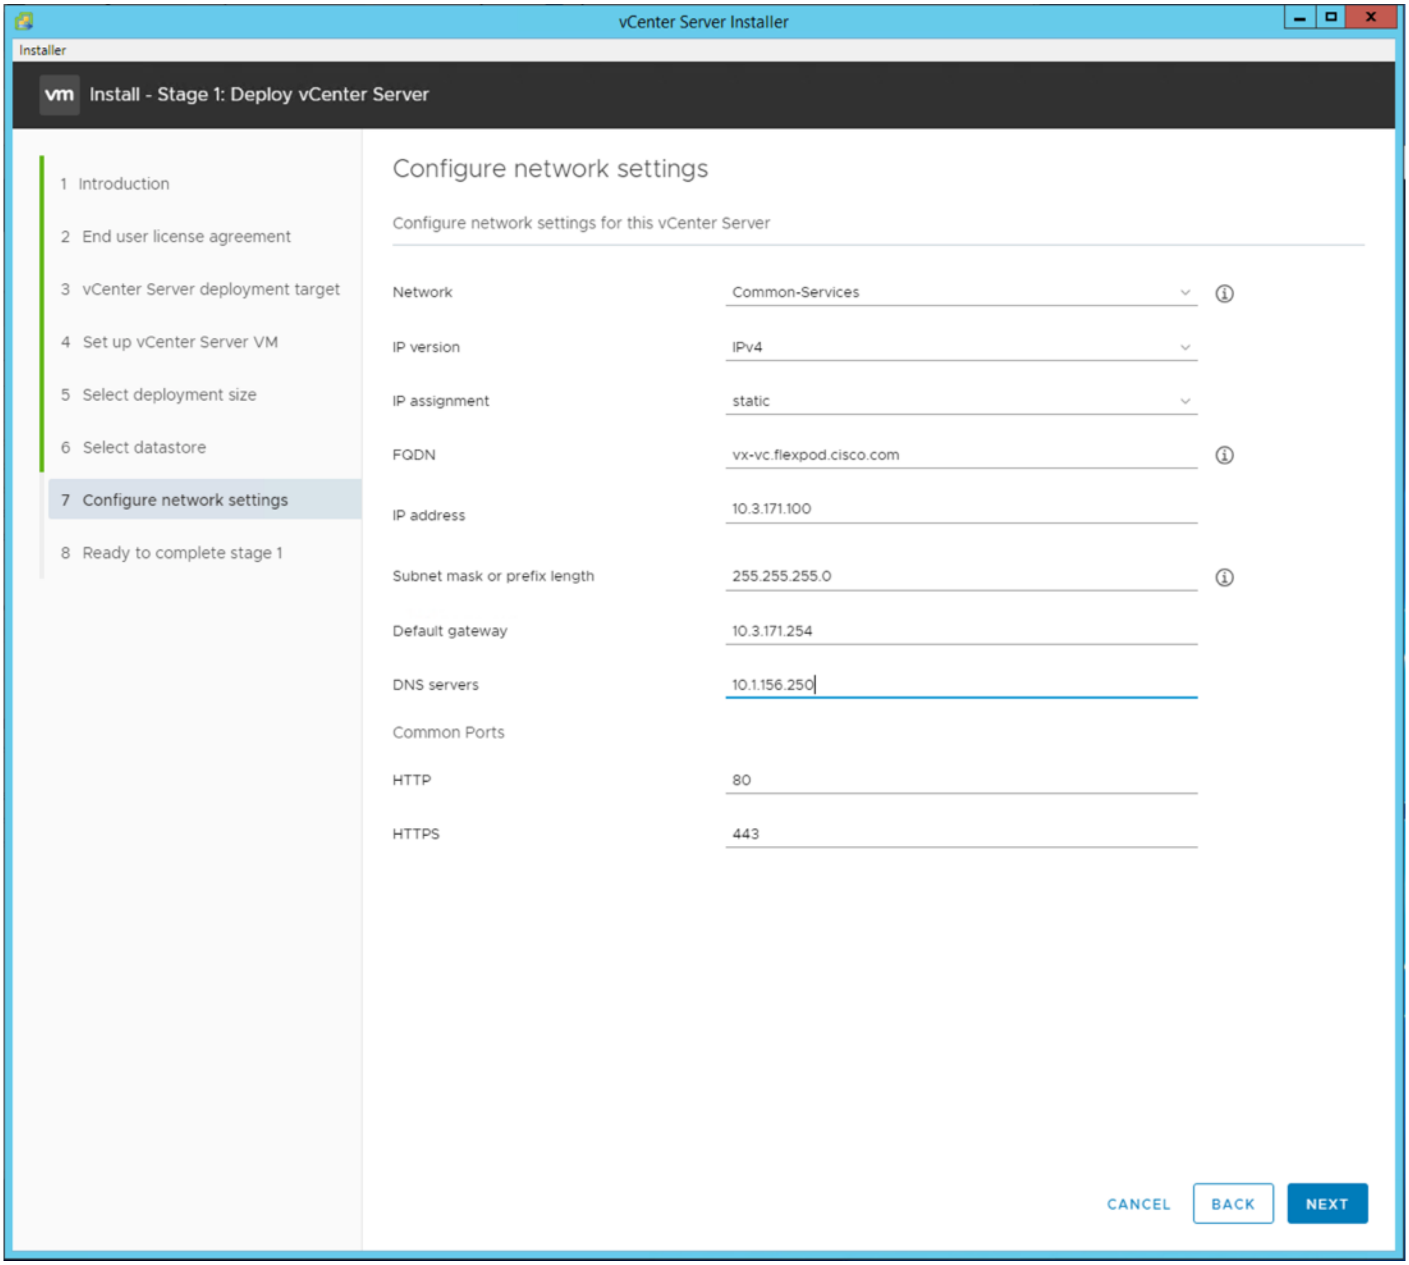
vCenter Server Installer
Installer
Install - Stage 1: Deploy vCenter Server
1 Introduction
2 End user license agreement
3 vCenter Server deployment target
4 Set up vCenter Server VM
5 Select deployment size
6 Select datastore
7 Configure network settings
8 Ready to complete stage 1
CANCEL
BACK
NEXT

Configure network settings

Configure network settings for this vCenter Server

Network	Common-Services
IP version	IPv4
IP assignment	static
FQDN	vx-vc.flexpod.cisco.com
IP address	10.3.171.100
Subnet mask or prefix length	255.255.255.0
Default gateway	10.3.171.254
DNS servers	10.1.156.250

Common Ports

HTTP	80
HTTPS	443

13. Click NEXT.

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

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

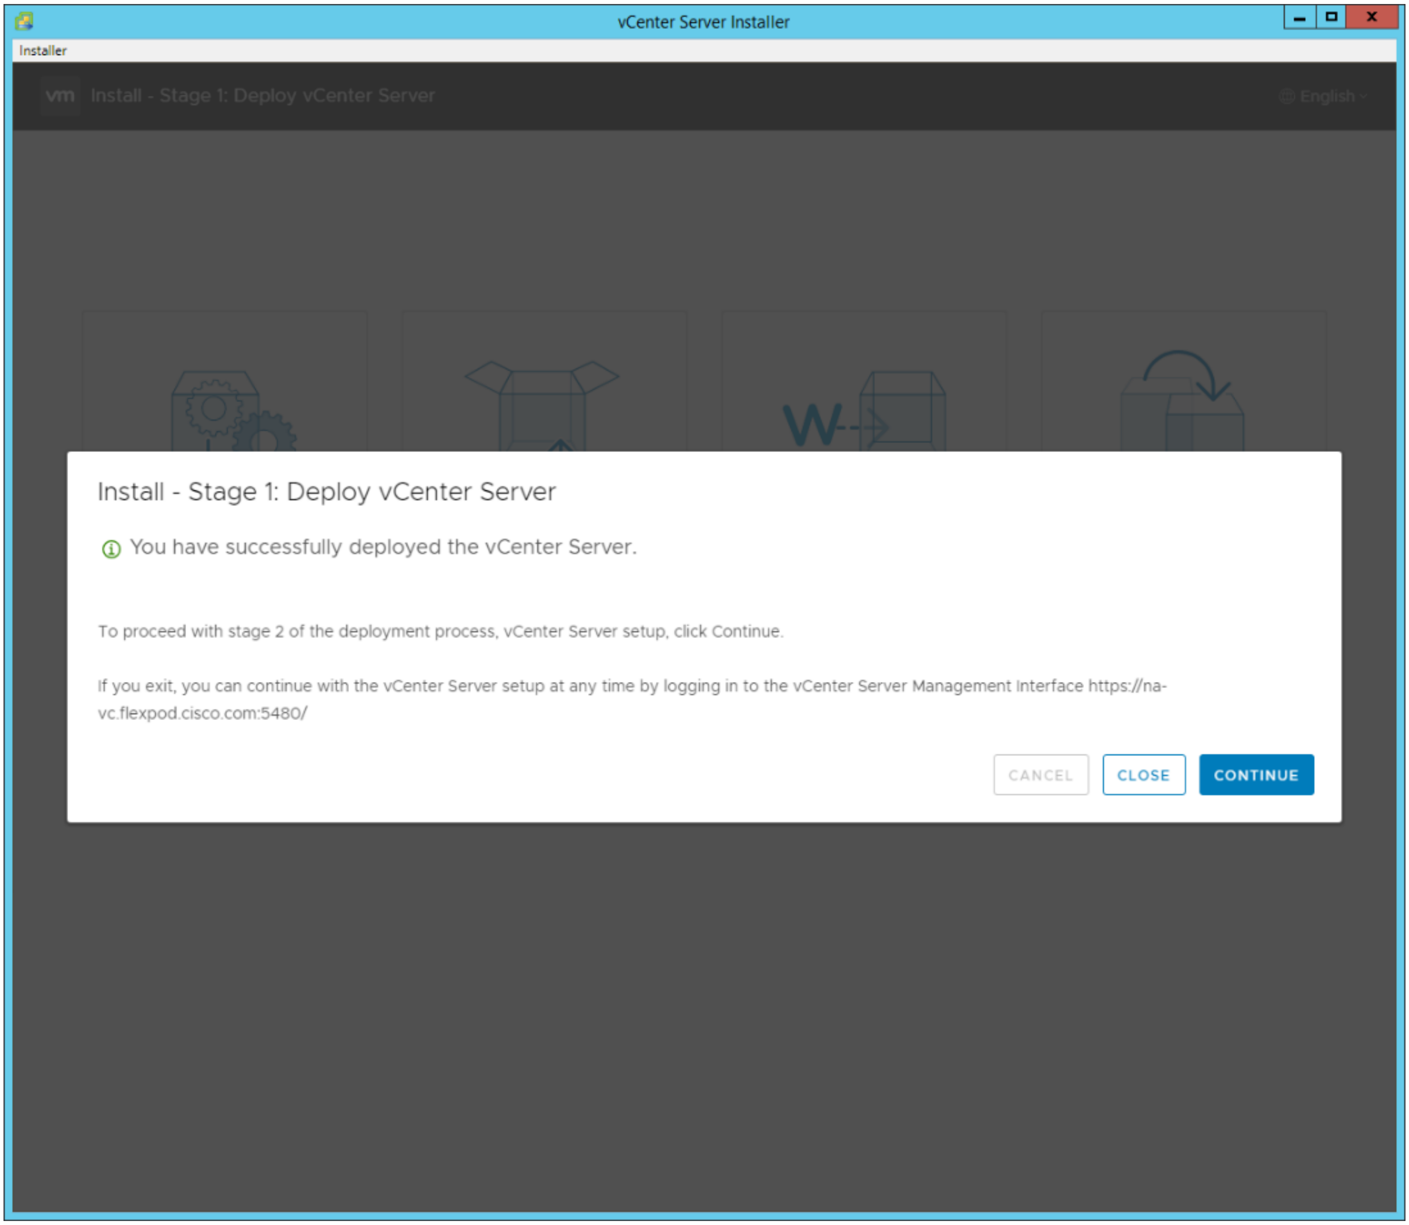

Install - Stage 1: Deploy vCenter Server

You have successfully deployed the vCenter Server.

To proceed with stage 2 of the deployment process, vCenter Server setup, click Continue.

If you exit, you can continue with the vCenter Server setup at any time by logging in to the vCenter Server Management Interface https://na-vc.flexpod.cisco.com:5480/

15. Click CONTINUE to proceed with stage 2 configuration.

16. Click NEXT.

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

    a. Time Synchronization Mode: Synchronize time with NTP servers.

    b. NTP Servers: <ntp-server>

    c. SSH access: Enabled.

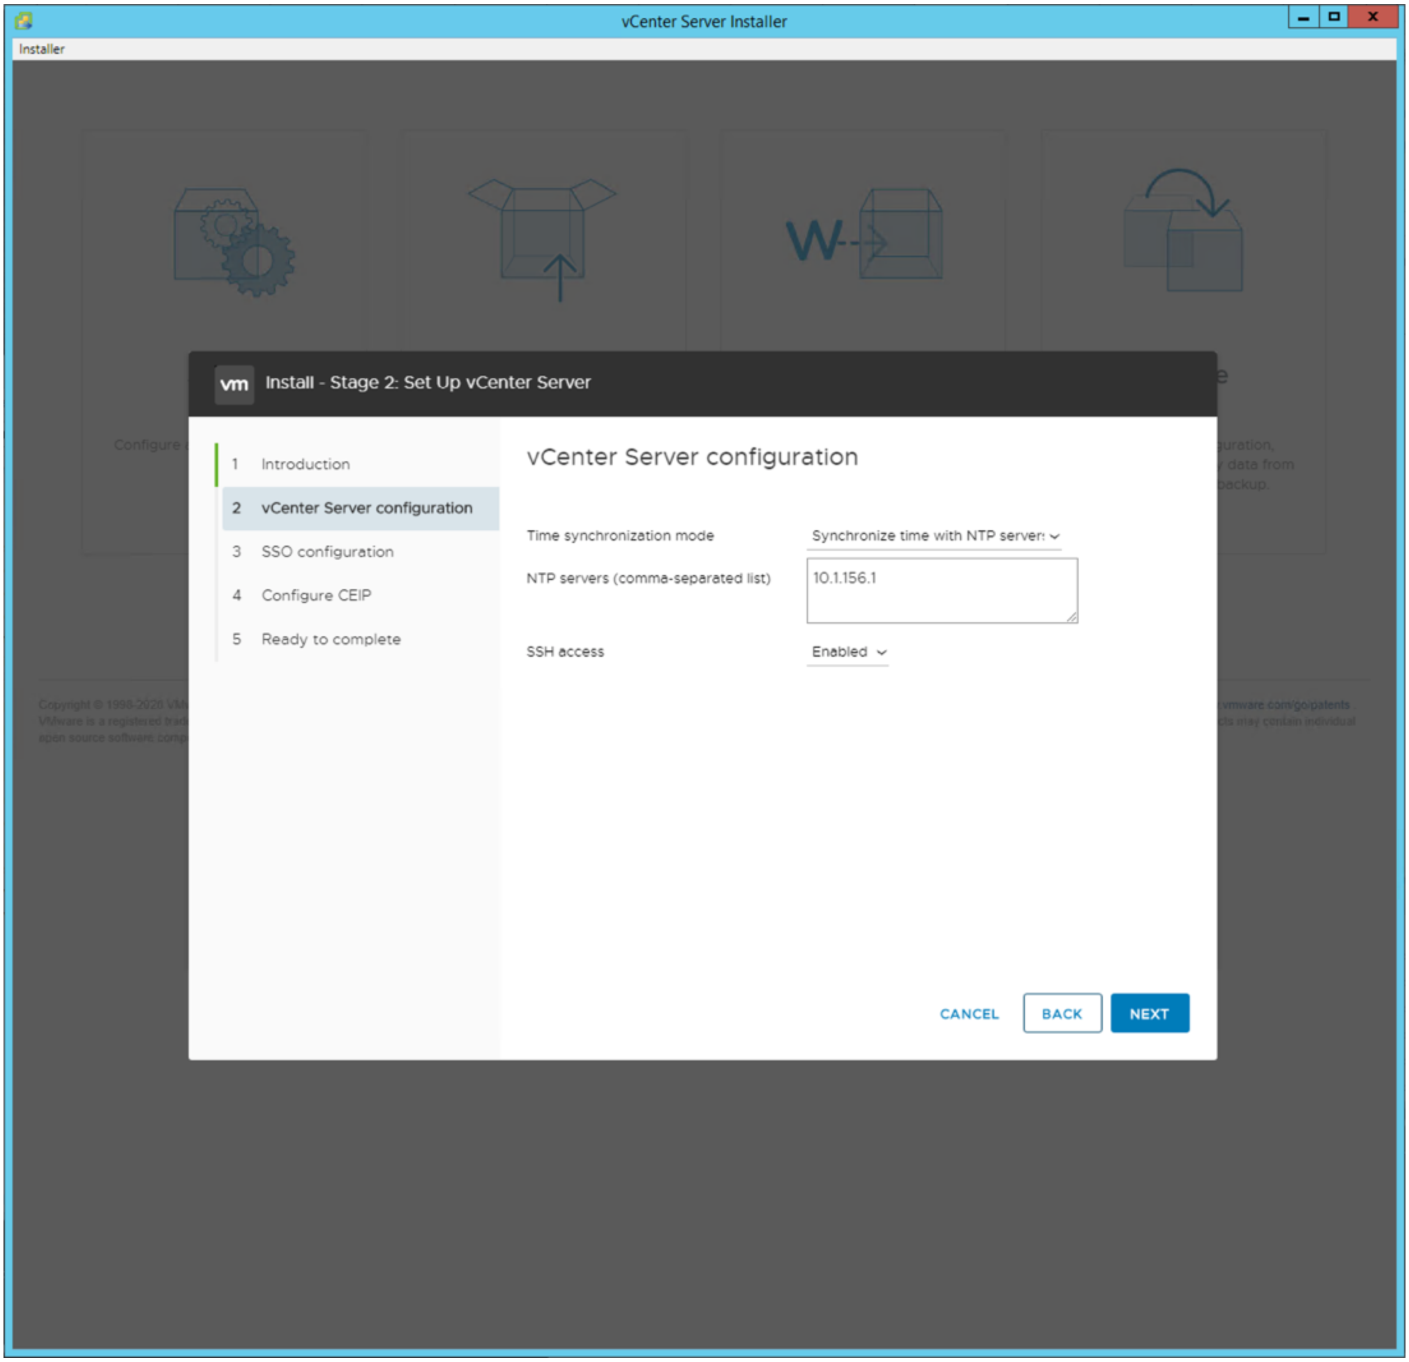

vCenter Server configuration

Time synchronization mode: Synchronize time with NTP servers

NTP servers (comma-separated list): 10.1.156.1

SSH access: Enabled

18. Click NEXT.

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

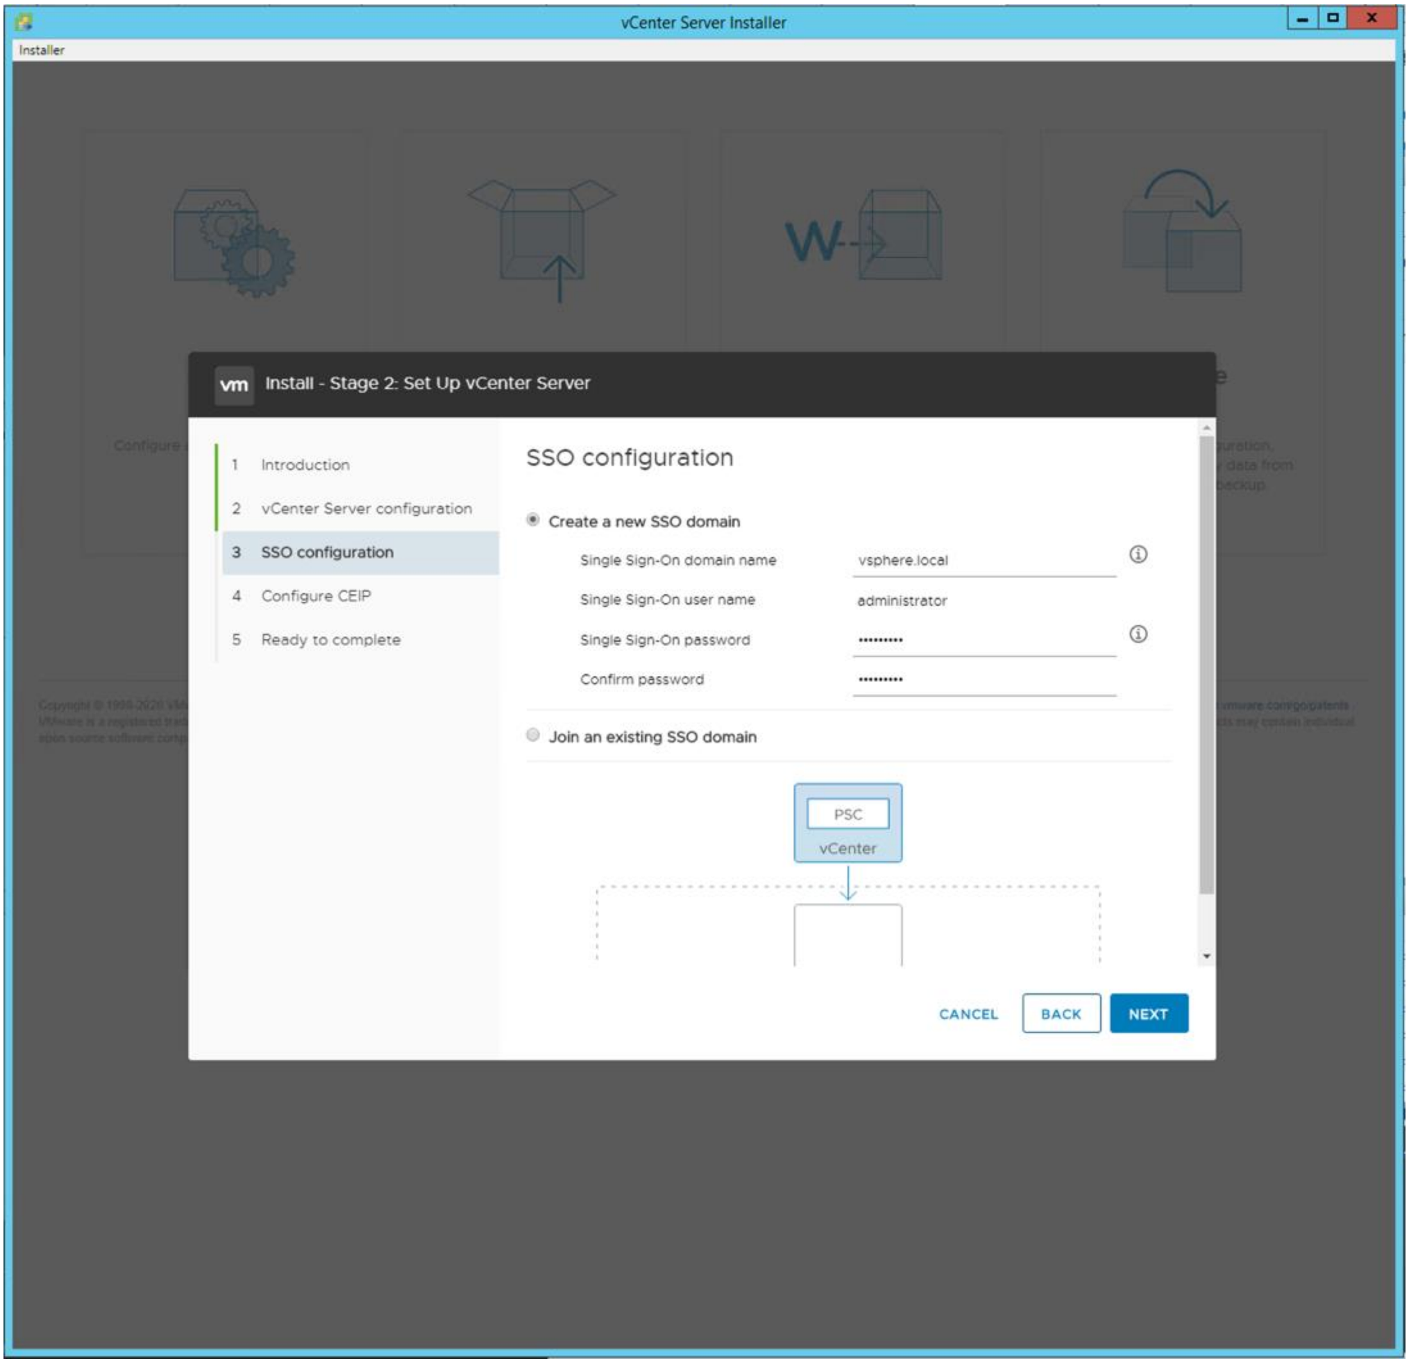

vm Install - Stage 2: Set Up vCenter Server

1 Introduction
2 vCenter Server configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete

SSO configuration

● Create a new SSO domain

Single Sign-On domain name	vsphere.local
Single Sign-On user name	administrator
Single Sign-On password	••••••••
Confirm password	••••••••

○ Join an existing SSO domain

20. Click NEXT.

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

22. Click NEXT.

23. Review the configuration and click FINISH.

24. Click OK.

vCenter Server setup will take a few minutes to complete.

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

Adjust vCenter CPU Settings

If a vCenter deployment size of Small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration. Cisco UCS B and C-Series servers are normally 2-socket servers. In this validation, the Small deployment size was selected and vCenter was setup for a 4-socket server. This setup will cause issues in the VMware ESXi cluster Admission Control. To resolve the Admission Control issue, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.

2. Enter root for the user name.

3. Enter the root password.

4. Click Login to connect.

5. On the left, choose Virtual Machines.

6. In the center pane, right-click the vCenter VM and choose Edit.

7. In the Edit settings window, expand CPU and check the value of Sockets.

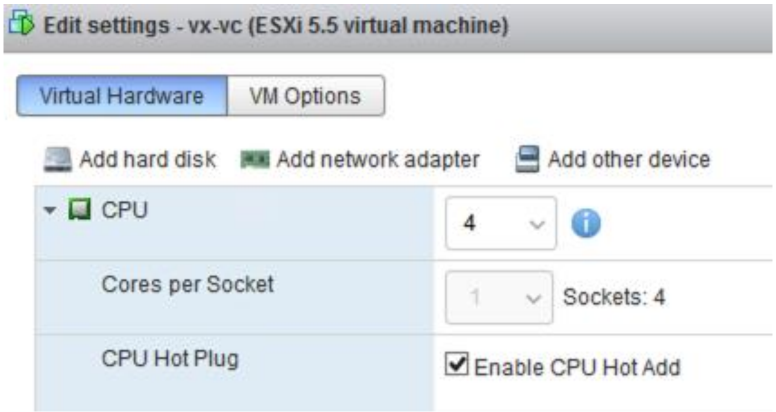

Edit settings - vx-vc (ESXi 5.5 virtual machine)

CPU: 4

Cores per Socket: 1 Sockets: 4

CPU Hot Plug: ☑ Enable CPU Hot Add

8. If the number of Sockets does not match your server configuration, it will need to be adjusted. Click Cancel.

9. If the number of Sockets needs to be adjusted:

    a. Right-click the vCenter VM and choose Guest OS > Shut down. Click Yes on the confirmation.
    b. Once vCenter is shut down, right-click the vCenter VM and choose Edit settings.
    c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to your server configuration (usually 2).

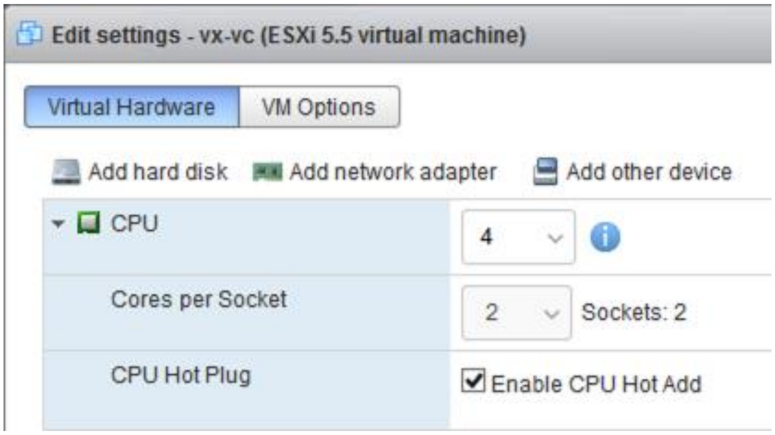
Edit settings - vx-vc (ESXi 5.5 virtual machine)

CPU: 4

Cores per Socket: 2 Sockets: 2

CPU Hot Plug: ☑ Enable CPU Hot Add

d. Click Save.

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

Setup VMware vCenter Server

To setup the VMware vCenter Server, follow these steps:

1. Using a web browser, navigate to https://<vcenter-ip-address>:5480. You will need to navigate security screens.

2. Log into the VMware vCenter Server Management interface as root with the root password set in the vCenter installation.

3. In the menu on the left, choose Time.

4. Choose EDIT to the right of Time zone.

5. Choose the appropriate Time zone and click SAVE.

6. In the menu on the left choose Administration.

7. According to your Security Policy, adjust the settings for the root user and password.

8. In the menu on the left choose Update.

9. Follow the prompts to STAGE AND INSTALL any available vCenter updates. In this validation, vCenter version 7.0.0.10700 was installed and did not require any updates at the time of installation.

10. In the upper right-hand corner of the screen, choose root > Logout to logout of the Appliance Management interface.

11. Using a web browser, navigate to https://<vcenter-fqdn>. You will need to navigate security screens.

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

12. Choose LAUNCH VSPHERE CLIENT (HTML5).

Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is the only option in vSphere 7 and will be used going forward.

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

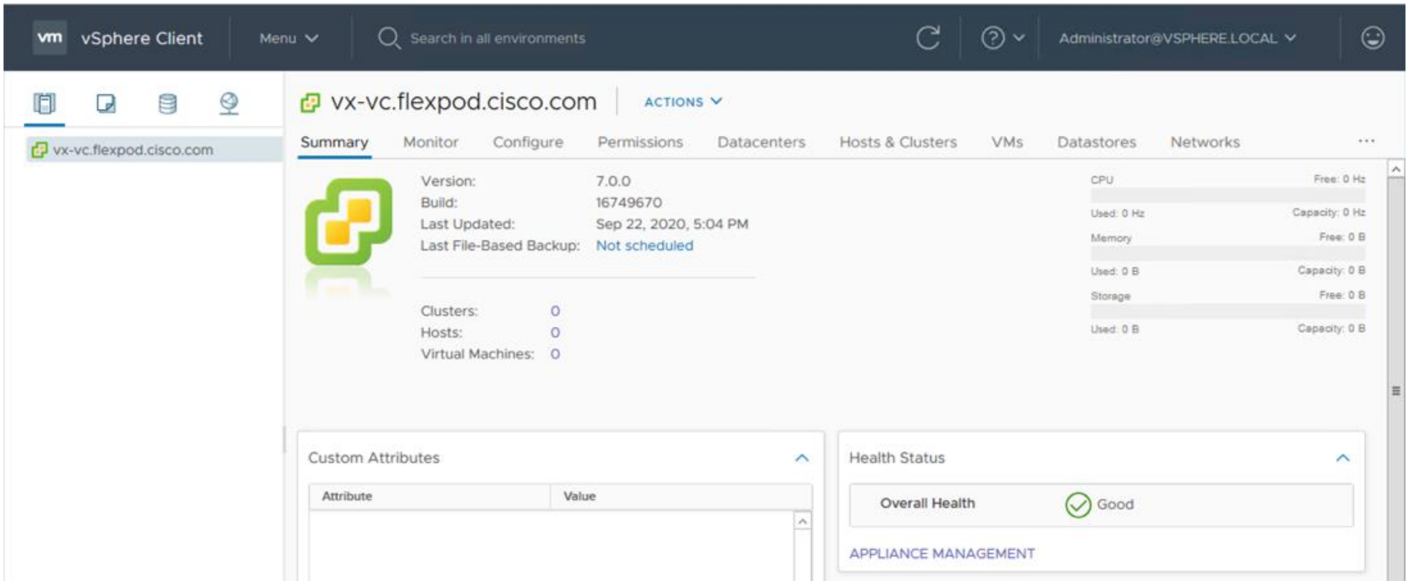

vx-vc.flexpod.cisco.com

Version: 7.0.0
Build: 16749670
Last Updated: Sep 22, 2020, 5:04 PM
Last File-Based Backup: Not scheduled

Clusters: 0
Hosts: 0
Virtual Machines: 0

CPU Free: 0 Hz
Used: 0 Hz Capacity: 0 Hz
Memory Free: 0 B
Used: 0 B Capacity: 0 B
Storage Free: 0 B
Used: 0 B Capacity: 0 B

Overall Health ✓ Good

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

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

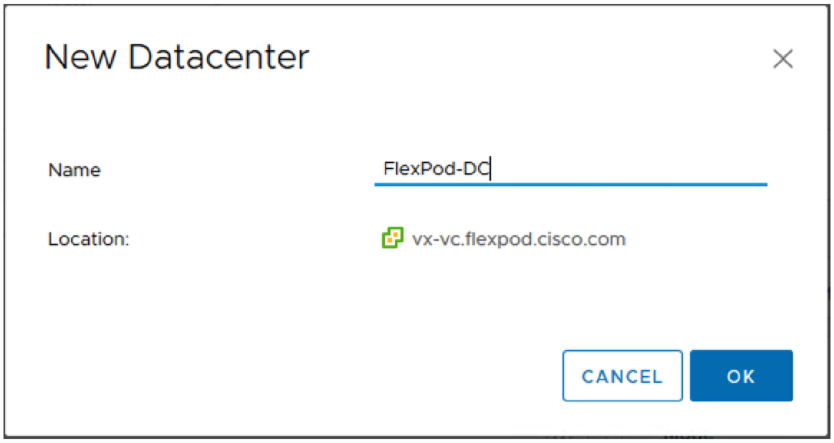

New Datacenter

Name FlexPod-DC

Location: vx-vc.flexpod.cisco.com

16. Click OK.

17. Expand the vCenter on the left.

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

19. Name the cluster FlexPod-Management.

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

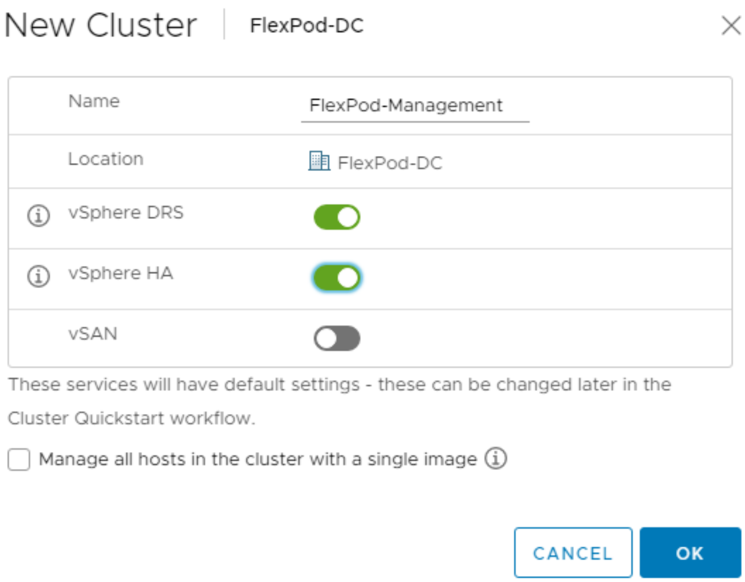

New Cluster | FlexPod-DC

Name	FlexPod-Management
Location	FlexPod-DC
vSphere DRS	On
vSphere HA	On
vSAN	Off

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

□ Manage all hosts in the cluster with a single image

21. Click OK to create the new cluster.

22. Right-click “FlexPod-Management” and choose Settings.

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

24. Choose Datastore specified by host and click OK.

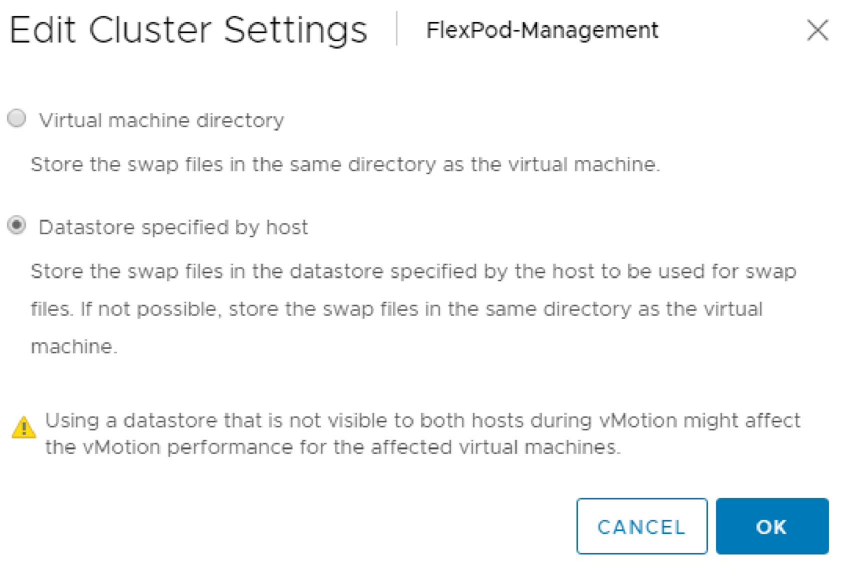

Edit Cluster Settings | FlexPod-Management

○ Virtual machine directory

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

● Datastore specified by host

Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine.

Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

25. Right-click “FlexPod-Management” and click Add Hosts.

26. In the IP address or FQDN field, enter either the IP address or the FQDN of the first VMware ESXi host. Enter root as the Username and the root password. Click NEXT.

27. In the Security Alert window, choose the host and click OK.

28. Verify the Host summary information and click NEXT.

29. Ignore warnings about the host being moved to Maintenance Mode if they appear and click FINISH to complete adding the host to the cluster.

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

31. In the list, right-click the added ESXi host and choose Settings.

32. In the left pane of the host under Virtual Machines, choose Swap File location.

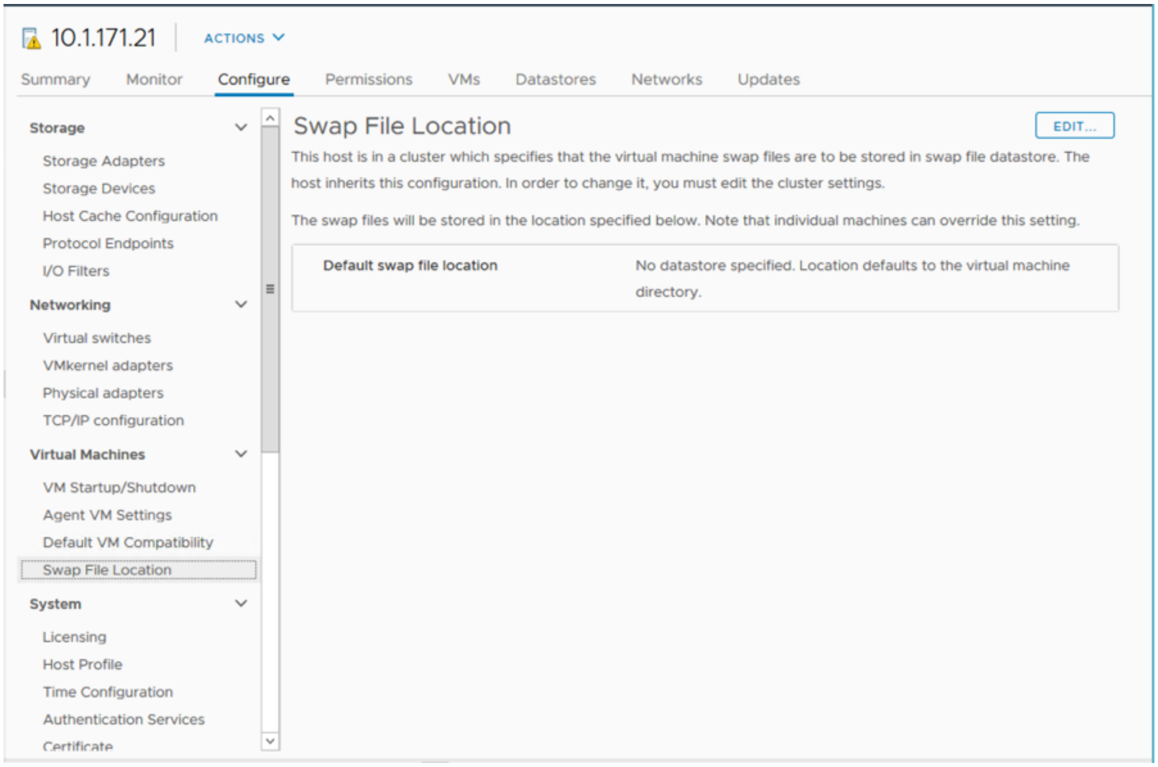

Swap File Location

This host is in a cluster which specifies that the virtual machine swap files are to be stored in swap file datastore. The host inherits this configuration. In order to change it, you must edit the cluster settings.

The swap files will be stored in the location specified below. Note that individual machines can override this setting.

Default swap file location: No datastore specified. Location defaults to the virtual machine directory.

33. Click EDIT.

34. Choose the infra_swap datastore and click OK.

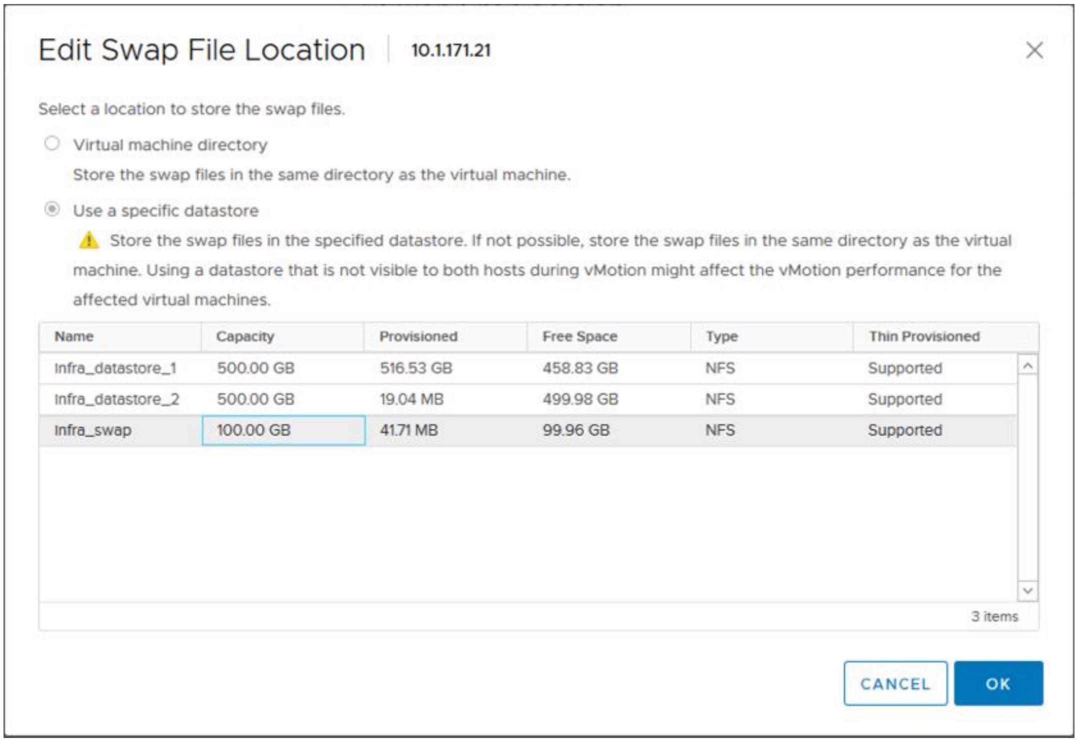

Select a location to store the swap files.

○ Virtual machine directory

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

● Use a specific datastore

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

Name	Capacity	Provisioned	Free Space	Type	Thin Provisioned
Infra_datastore_1	500.00 GB	516.53 GB	458.83 GB	NFS	Supported
Infra_datastore_2	500.00 GB	19.04 MB	499.98 GB	NFS	Supported
Infra_swap	100.00 GB	41.71 MB	99.96 GB	NFS	Supported

3 items

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

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

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

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

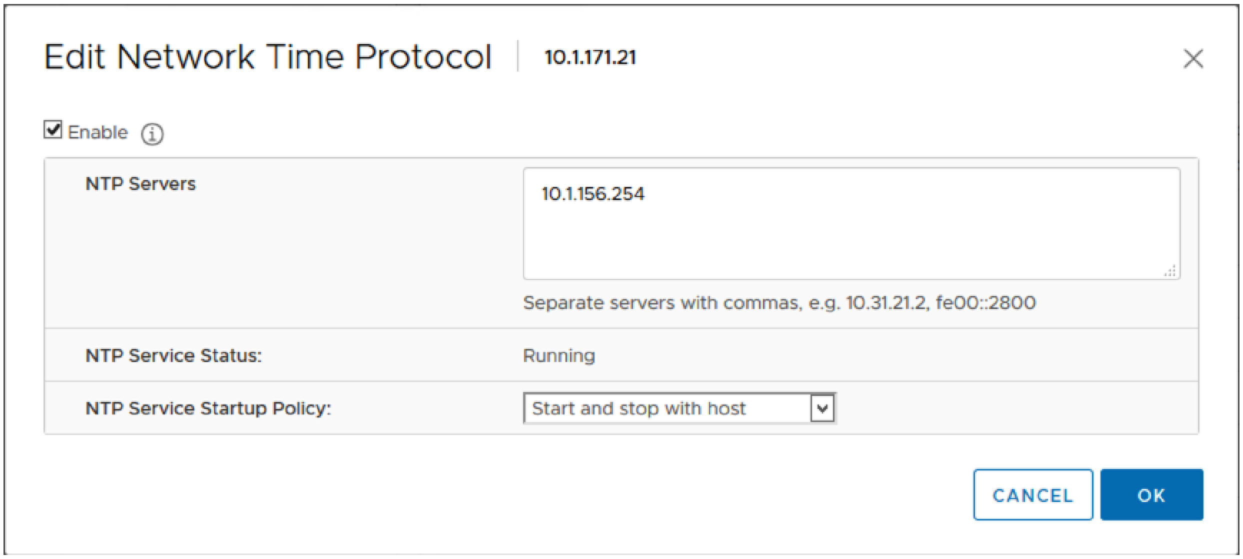

Edit Network Time Protocol | 10.1.171.21

☑ Enable

NTP Servers: 10.1.156.254

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

NTP Service Status: Running

NTP Service Startup Policy: Start and stop with host

39. In the list under Hardware, choose Overview. Scroll to the bottom and ensure the Power Management Active policy is High Performance. If the Power Management Active policy is not High Performance, to the right of Power Management, choose EDIT POWER POLICY. Choose High performance and click OK.

40. In the list under Storage, choose Storage Devices. Make sure the NETAPP Fibre Channel Disk LUN 0 or NETAPP iSCSI Disk LUN 0 is selected.

41. Choose the Paths tab.

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

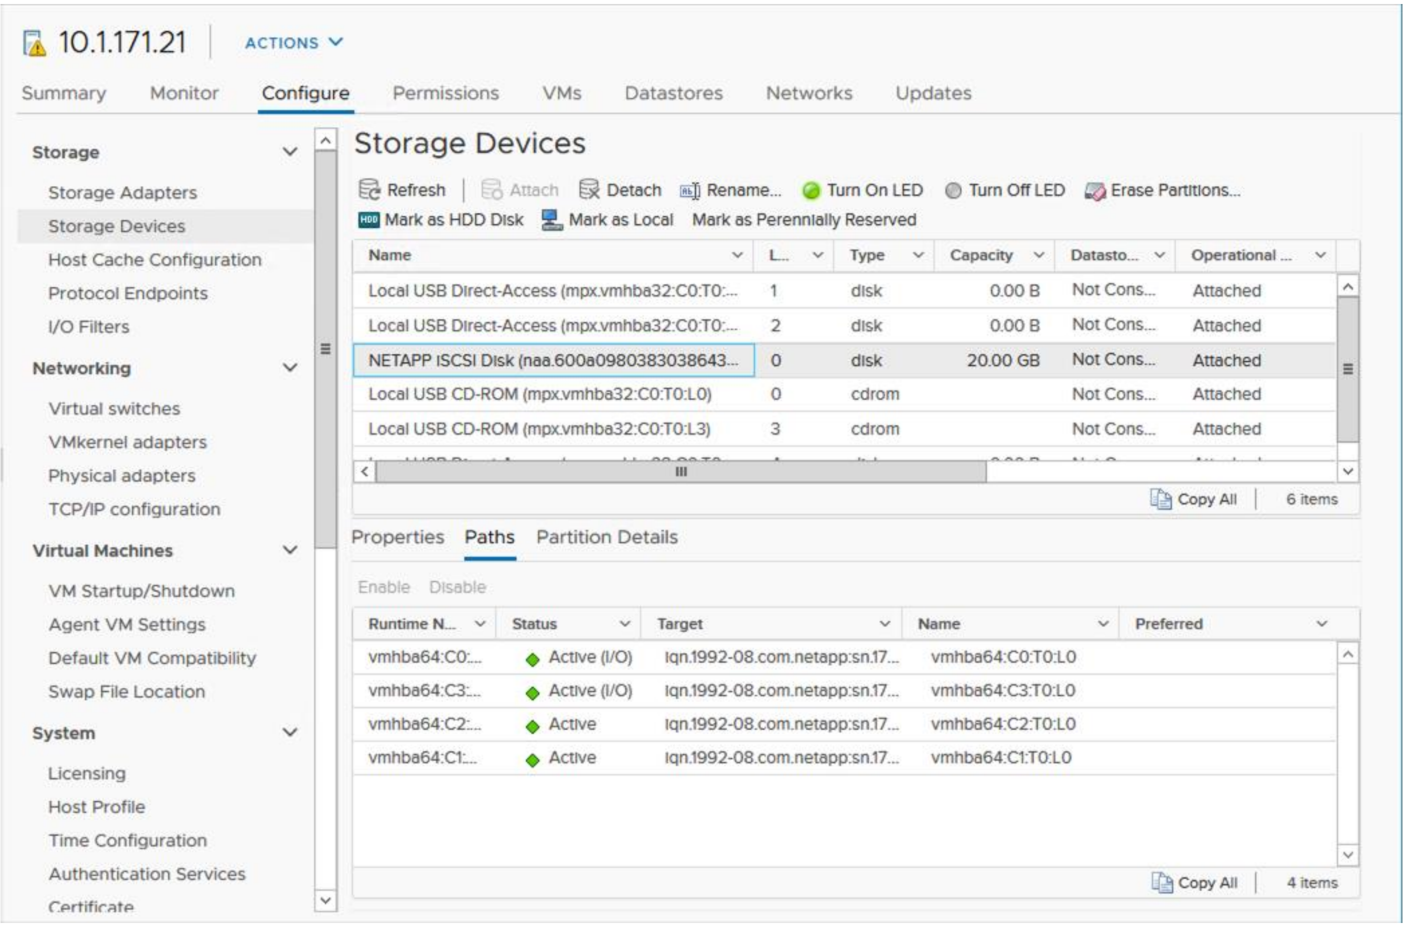

Name	L...	Type	Capacity	Datasto...	Operational ...
Local USB Direct-Access (mpx.vmhba32:C0:T0:...	1	disk	0.00 B	Not Cons...	Attached
Local USB Direct-Access (mpx.vmhba32:C0:T0:...	2	disk	0.00 B	Not Cons...	Attached
NETAPP iSCSI Disk (naa.600a0980383038643...	0	disk	20.00 GB	Not Cons...	Attached
Local USB CD-ROM (mpx.vmhba32:C0:T0:L0)	0	cdrom		Not Cons...	Attached
Local USB CD-ROM (mpx.vmhba32:C0:T0:L3)	3	cdrom		Not Cons...	Attached

Runtime N...	Status	Target	Name	Preferred
vmhba64:C0:...	◆ Active (I/O)	iqn.1992-08.com.netapp:sn.17...	vmhba64:C0:T0:L0	
vmhba64:C3:...	◆ Active (I/O)	iqn.1992-08.com.netapp:sn.17...	vmhba64:C3:T0:L0	
vmhba64:C2:...	◆ Active	iqn.1992-08.com.netapp:sn.17...	vmhba64:C2:T0:L0	
vmhba64:C1:...	◆ Active	iqn.1992-08.com.netapp:sn.17...	vmhba64:C1:T0:L0	

Add AD User Authentication to vCenter (Optional)

If an AD Infrastructure is set up in this FlexPod environment, you can setup in AD and authenticate from vCenter.

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

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).
2. Connect to https://<vcenter-ip> and choose LAUNCH VSPHERE CLIENT (HTML5).
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. Under Menu, choose Administration. In the list on the left, under Single Sign On, choose Configuration.
5. In the center pane, under Configuration, choose the Identity Provider tab.
6. In the list under Type, select Active Directory Domain.
7. Choose JOIN AD.

8. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Do not fill in an Organizational unit. Click JOIN.

9. Click Acknowledge.

10. In the list on the left under Deployment, choose System Configuration. Choose the radio button to choose the vCenter, then choose REBOOT NODE.

11. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.

12. Log back into the vCenter vSphere HTML5 Client as Administrator@vsphere.local.

13. Under Menu, choose Administration. In the list on the left, under Single Sign On, choose Configuration.

14. In the center pane, under Configuration, choose the Identity Provider tab. Under Type, select Identity Sources. Click ADD.

15. Make sure your Active Directory (Integrated Windows Authentication) is selected, your Windows Domain name is listed, and Use machine account is selected. Click ADD.

16. In the list select the Active Directory (Integrated Windows Authentication) Identity source type. If desired, select SET AS DEFAULT and click OK.

17. On the left under Access Control, choose Global Permissions.

18. In the center pane, click the + sign to add a Global Permission.

19. In the Add Permission window, choose your AD domain for the Domain.

20. On the User/Group line, enter either the FlexPod Admin username or the Domain Admins group. Leave the Role set to Administrator. Choose the Propagate to children checkbox.

The FlexPod Admin user was created in the Domain Admins group. The selection here depends on whether the FlexPod Admin user will be the only user used in this FlexPod or you would like to add other users later. By selecting the Domain Admins group, any user placed in that group in the AD domain will be able to login to vCenter as an Administrator.

21. Click OK to add the selected User or Group. The user or group should now appear in the Global Permissions list with the Administrator role.

22. Log out and log back into the vCenter HTML5 Client as the FlexPod Admin user. You will need to add the domain name to the user, for example, flexadmin@domain.

FlexPod VMware vSphere Distributed Switch (vDS)

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

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

A vMotion, and a VM-Traffic port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would need to have the corresponding VLANs added to the Cisco UCS LAN cloud, to the Cisco UCS vDS0-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC peer-link interfaces on the switches.

In this document, the infrastructure ESXi management VMkernel ports, the In-Band management interfaces including the vCenter management interface, and the infrastructure NFS VMkernel ports are left on vSwitch0 to facilitate bringing the virtual environment back up in the event it needs to be completely shut down. When the vMotion port group is created, it is pinned to Cisco UCS fabric B to reduce the need for vMotion traffic leaving the fabric interconnect. Pinning should be done in a vDS to ensure consistency across all ESXi hosts.

Configure the VMware vDS in vCenter

VMware vSphere Web Client

To configure the vDS, follow these steps:

1. After logging into the VMware vSphere HTML5 Client, choose Networking under Menu.
2. Right-click the FlexPod-DC datacenter and choose Distributed Switch > New Distributed Switch.
3. Give the Distributed Switch a descriptive name (vDS0) and click NEXT.
4. Make sure version 7.0.0 – ESXi 7.0 and later is selected and click NEXT.
5. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter VM-Traffic-1 for the Port group name. Click NEXT.
6. Review the information and click FINISH to complete creating the vDS.
7. Expand the FlexPod-DC datacenter and the newly created vDS. Choose the newly created vDS.
8. Right-click the VM-Traffic-1 port group and choose Edit Settings.
9. Choose VLAN on the left.
10. Choose VLAN for VLAN type and enter the VM-Traffic-1 VLAN ID. Click OK.
11. Right-click the vDS and choose Settings > Edit Settings.
12. In the Edit Settings window, choose Advanced on the left.
13. Change the MTU to 9000. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.

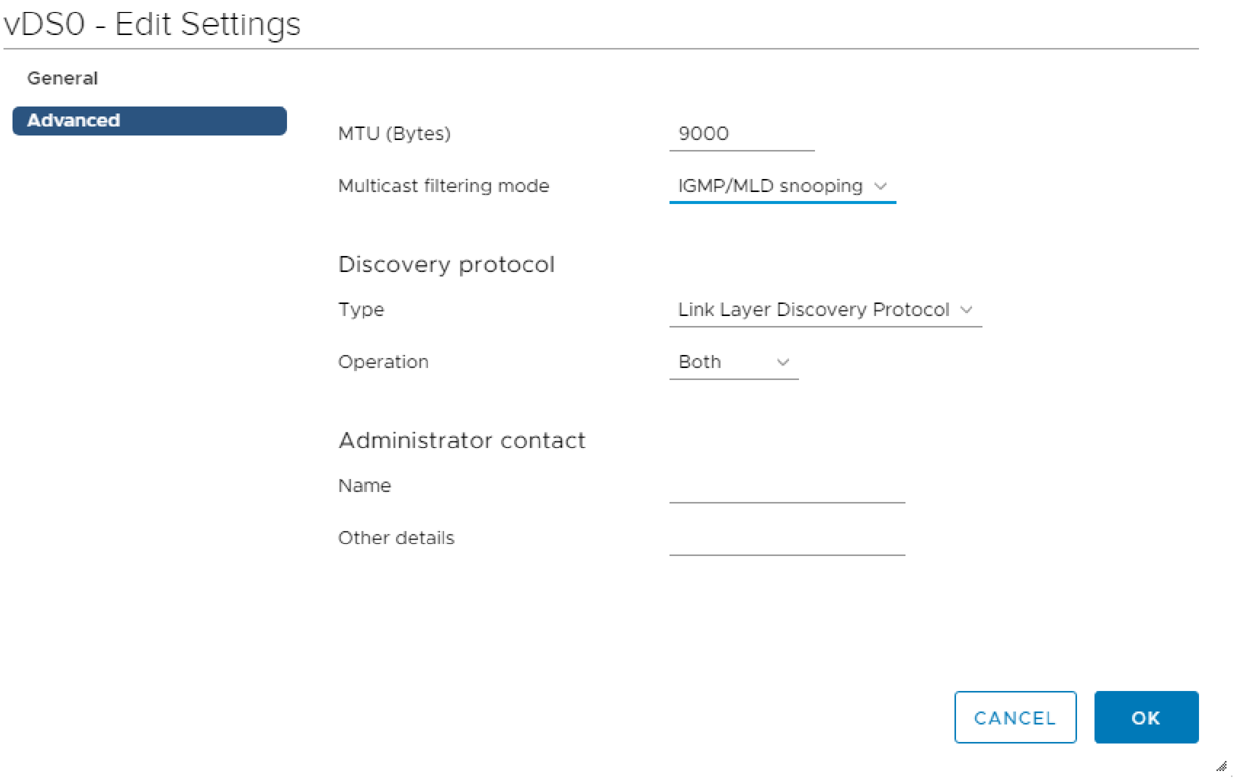

vDS0 - Edit Settings

MTU (Bytes) 9000

Multicast filtering mode IGMP/MLD snooping

Discovery protocol

Type Link Layer Discovery Protocol

Operation Both

Administrator contact

Name

Other details

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

15. Enter vMotion as the name and click NEXT.

16. Set the VLAN type to VLAN, enter the VLAN ID used for vMotion, click the Customize default policies configuration check box, and click NEXT.

17. Leave the Security options set to Reject and click NEXT.

18. Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.

19. Choose Uplink 1 from the list of Active uplinks and click the down arrow icon twice to place Uplink 1 in the list of Standby uplinks. This will pin all vMotion traffic to Cisco UCS Fabric Interconnect B except when a failure occurs.

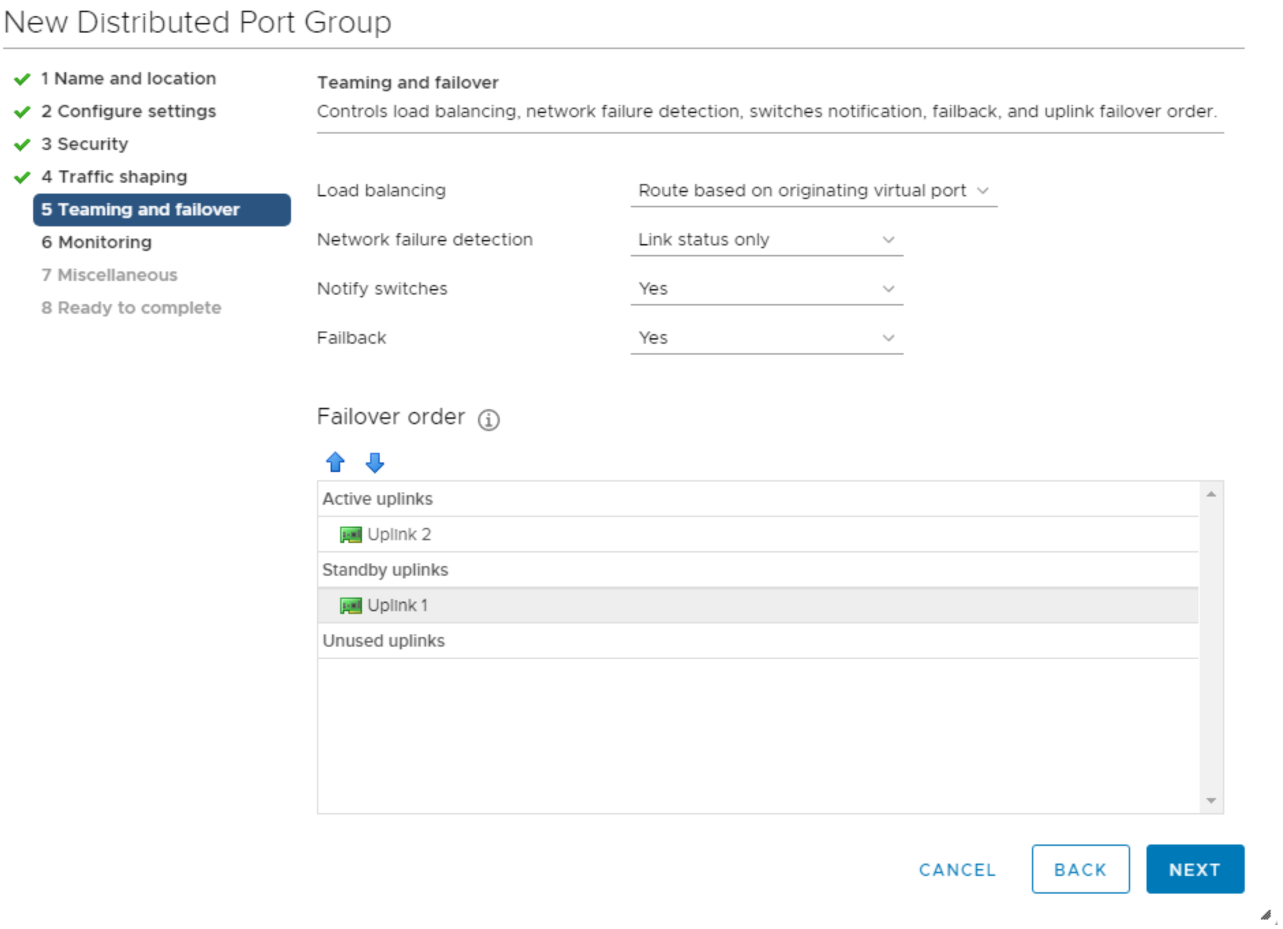

New Distributed Port Group

- ✓ 1 Name and location
- ✓ 2 Configure settings
- ✓ 3 Security
- ✓ 4 Traffic shaping
- 5 Teaming and failover
- 6 Monitoring
- 7 Miscellaneous
- 8 Ready to complete

Teaming and failover

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

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

Failover order

Active uplinks
Uplink 2
Standby uplinks
Uplink 1
Unused uplinks

20. Click NEXT.

21. Leave NetFlow disabled and click NEXT.

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

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

24. Repeat the addition of new distributed port groups for any additional application port groups, setting the active links as appropriate.

25. Right-click the vDS and choose Add and Manage Hosts.

26. Make sure Add hosts is selected and click NEXT.

27. Click the green + sign to add New hosts. Choose the one configured FlexPod Management host and click OK. Click NEXT.

28. Choose vmnic2 and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 and click Assign uplink. Choose Uplink 2 and click OK.

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

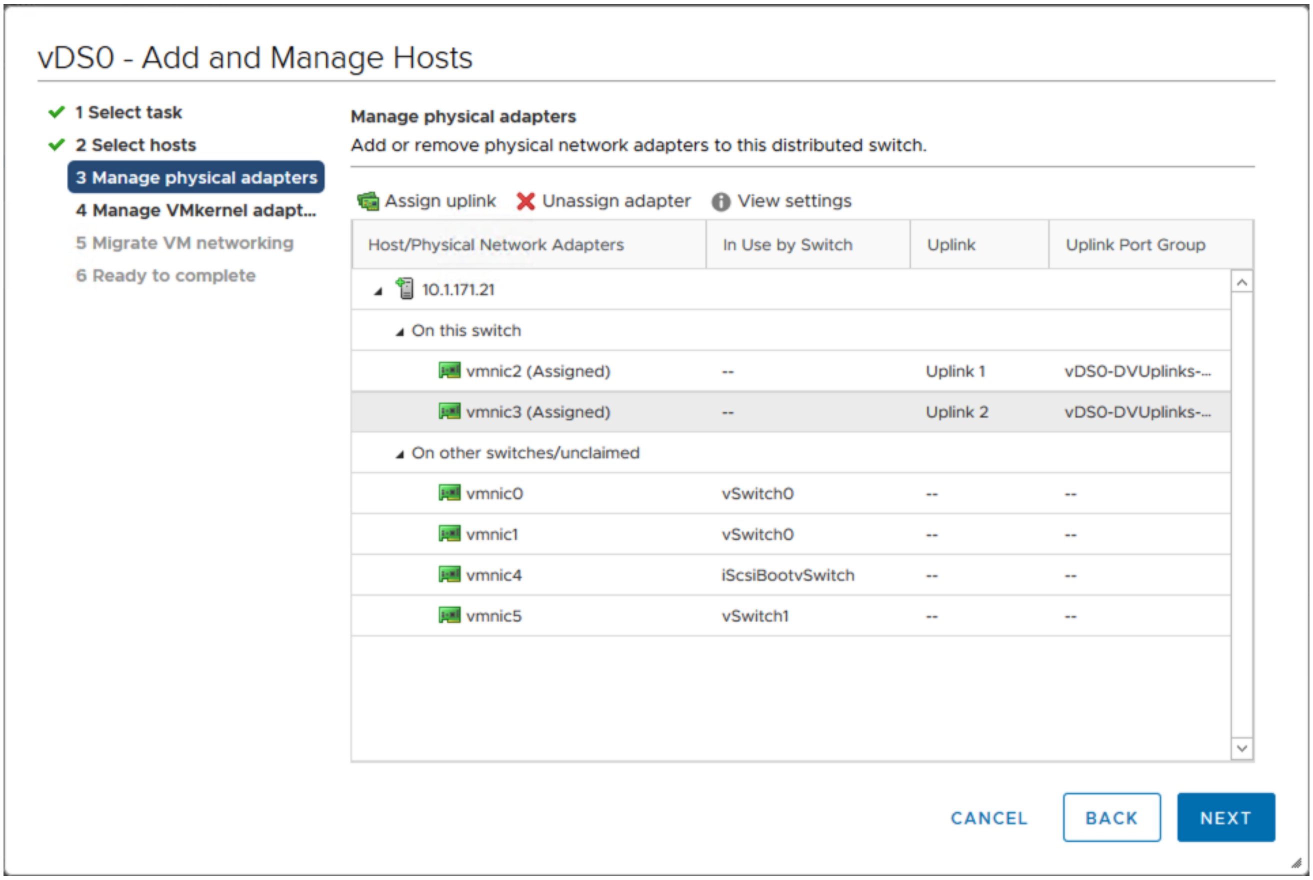

vDS0 - Add and Manage Hosts

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

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
10.1.171.21			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic4	iScsiBootvSwitch	--	--
vmnic5	vSwitch1	--	--

29. Click NEXT.

30. Do not migrate any VMkernel ports and click NEXT.

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

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

Add and Configure a VMware ESXi Host in vCenter

This section details the steps to add and configure an ESXi host in vCenter. This section assumes the host has had VMware ESXi 7.0b installed, the management IP address set, and the Cisco UCS Tool and NetApp NFS Plug-in for VMware VAAI installed. This procedure is initially being run on the second and third ESXi management hosts but can be run on any added ESXi host.

Add the ESXi Host to vCenter

ESXi Hosts created other than VM-Host-Infra-01

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

1. From the Home screen in the VMware vCenter HTML5 Interface, choose Menu > Hosts and Clusters.

2. Right-click the “FlexPod-Management” cluster and click Add Hosts.

3. In the IP address or FQDN field, enter either the IP address or the FQDN name of the configured VMware ESXi host. Also enter the user id (root) and associated password. If more than one host is being added, add the corresponding host information, optionally selecting “Use the same credentials for all hosts”. Click NEXT.

4. Choose all hosts being added and click OK to accept the certificate(s).

5. Review the host details and click NEXT to continue.

6. Review the configuration parameters and click FINISH to add the host(s).

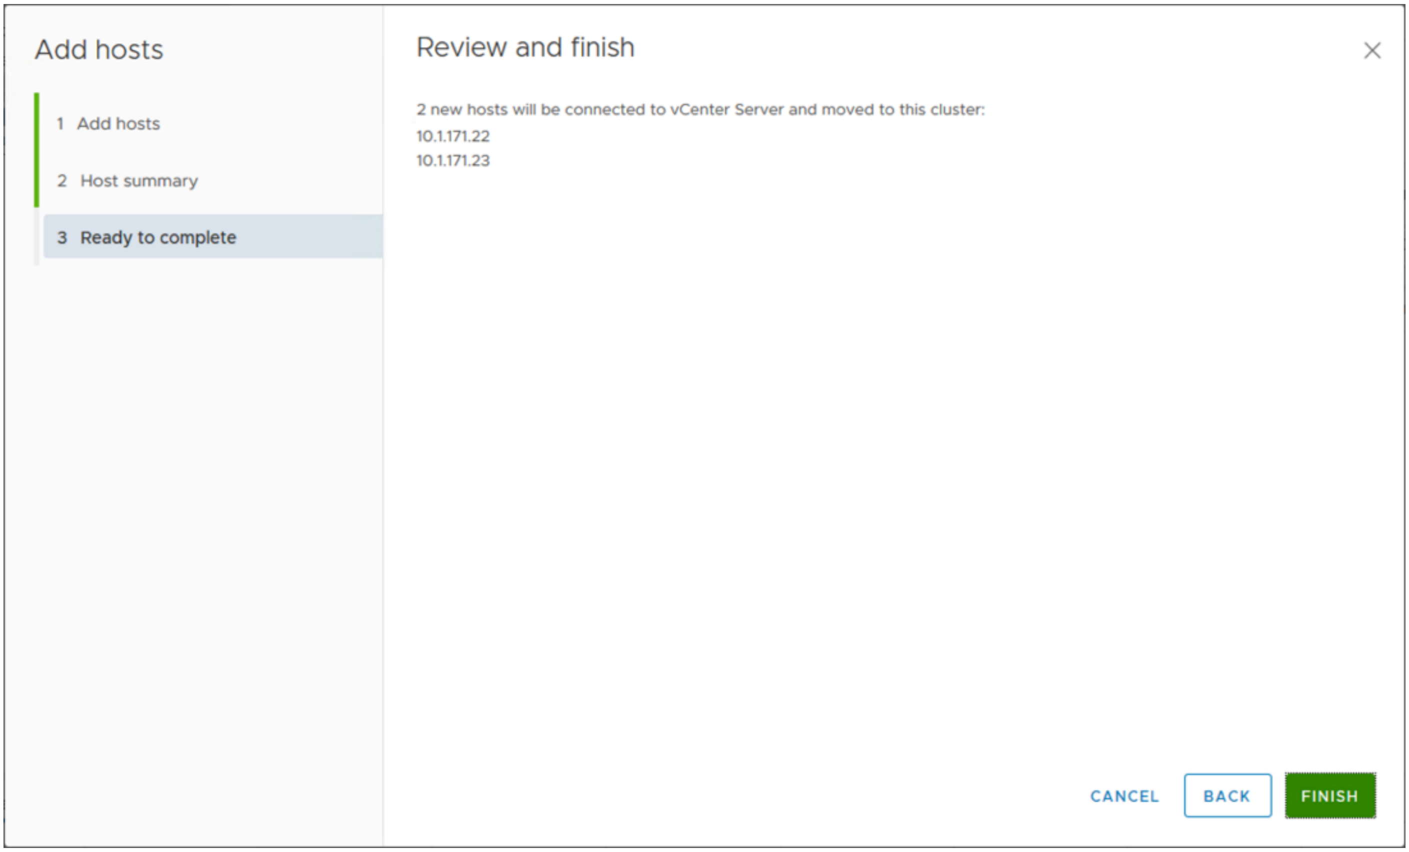

Review and finish

2 new hosts will be connected to vCenter Server and moved to this cluster:

10.1.171.22

10.1.171.23

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

Set Up VMkernel Ports and Virtual Switch

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

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

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, choose the Configure tab.

3. In the list, choose Virtual switches under Networking.

4. Expand Standard Switch: vSwitch0.

5. Choose EDIT to Edit settings.

6. Change the MTU to 9000.

7. Choose Teaming and failover located on the left.

8. In the Failover order section, use the arrow icons to move the vmnics until both are Active adapters.

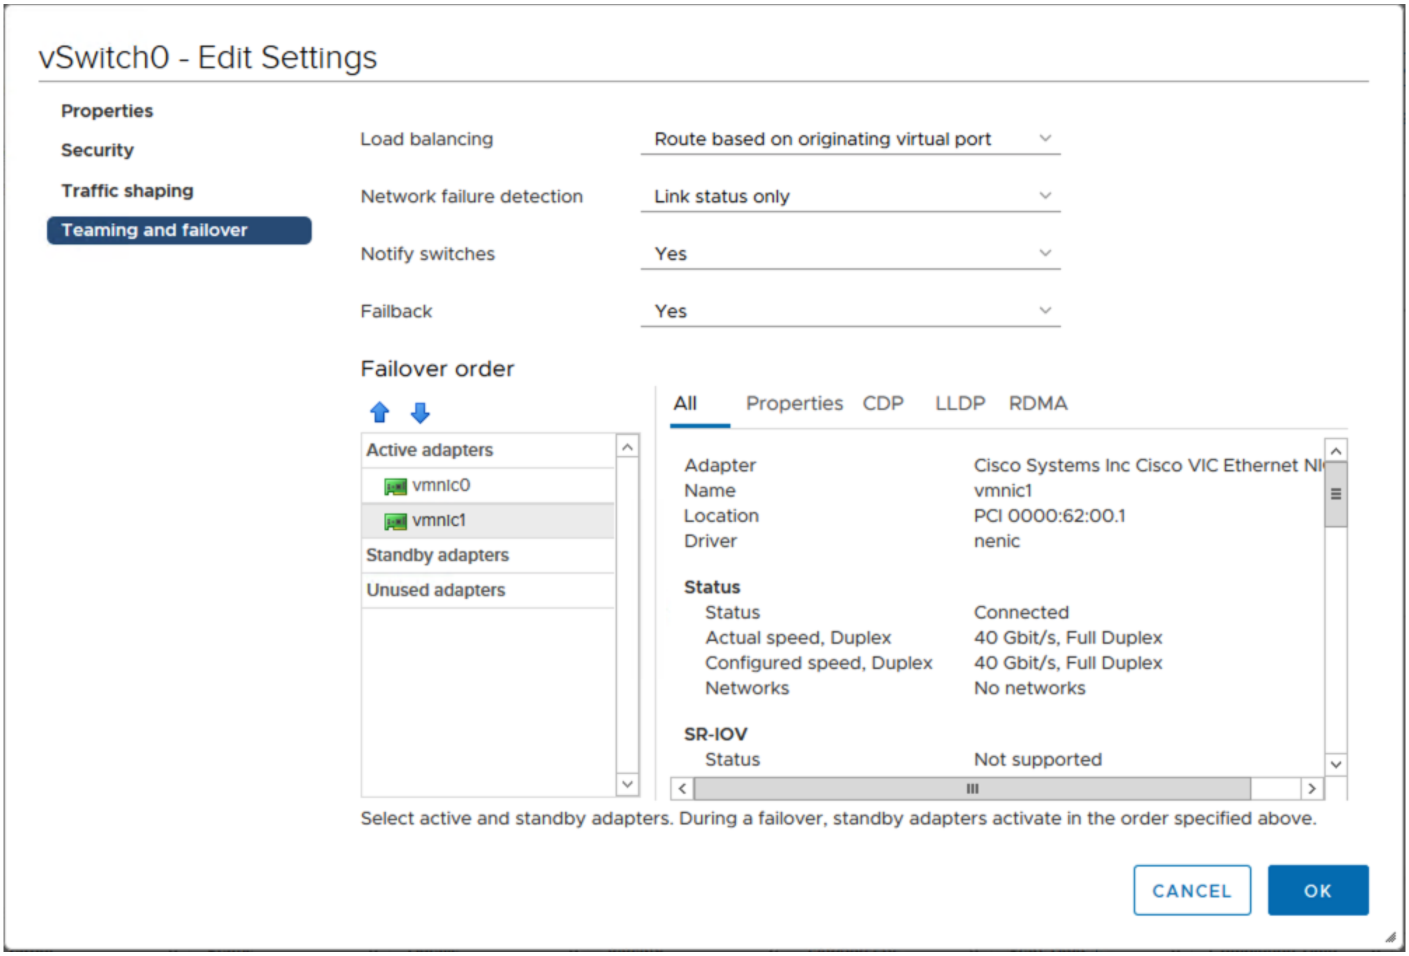

vSwitch0 - Edit Settings

Load balancing: Route based on originating virtual port

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active adapters
- vmnic0
- vmnic1

Standby adapters

Unused adapters

Adapter	Cisco Systems Inc Cisco VIC Ethernet NI
Name	vmnic1
Location	PCI 0000:62:00.1
Driver	nenic
Status	
Status	Connected
Actual speed, Duplex	40 Gbit/s, Full Duplex
Configured speed, Duplex	40 Gbit/s, Full Duplex
Networks	No networks
SR-IOV	
Status	Not supported

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

Speeds will appear as 20 Gbit/s for UCS blades not equipped with a port expander.

9. Click OK.

10. In the center pane, to the right of VM Network click ... > Edit Settings to edit settings.

11. Rename the port group to Site1-IB and enter <site1-ib-vlan-id> in the VLAN ID field.

12. Click OK to finalize the edits for the Site1-IB port group.

13. Still within the center pane, click Add Networking to create the Common-Services port group.

14. Select Virtual Machine Port Group for a Standard Switch and click NEXT.

15. Leave Select an existing standard switch selected with the default of vSwitch0 and click NEXT.

16. Provide the name Common-Services for the Network label and specify the <common-services-vlan-id> in the VLAN ID field and click NEXT.

17. Click FINISH to create the port group.

18. Located on the left under Networking, choose VMkernel adapters.

19. In the center pane, click Add Networking.

20. Make sure VMkernel Network Adapter is selected and click NEXT.

21. Choose an existing standard switch and click BROWSE. Choose vSwitch0 and click OK. Click NEXT.

22. For Network label, enter VMkernel-Infra-NFS.

23. Enter <infra-nfs-vlan-id> for the VLAN ID.

24. Choose Custom for MTU and make sure 9000 is entered.

25. Leave the Default TCP/IP stack selected and do not choose any of the Enabled services. Click NEXT.

26. Choose Use static IPv4 settings and enter the IPv4 address and subnet mask for the Infra-NFS VMkernel port for this ESXi host.

27. Click NEXT.

28. Review the settings and click FINISH to create the VMkernel port.

29. On the left under Networking, choose Virtual switches. Then expand vSwitch0. The properties for vSwitch0 should be similar to the following example:

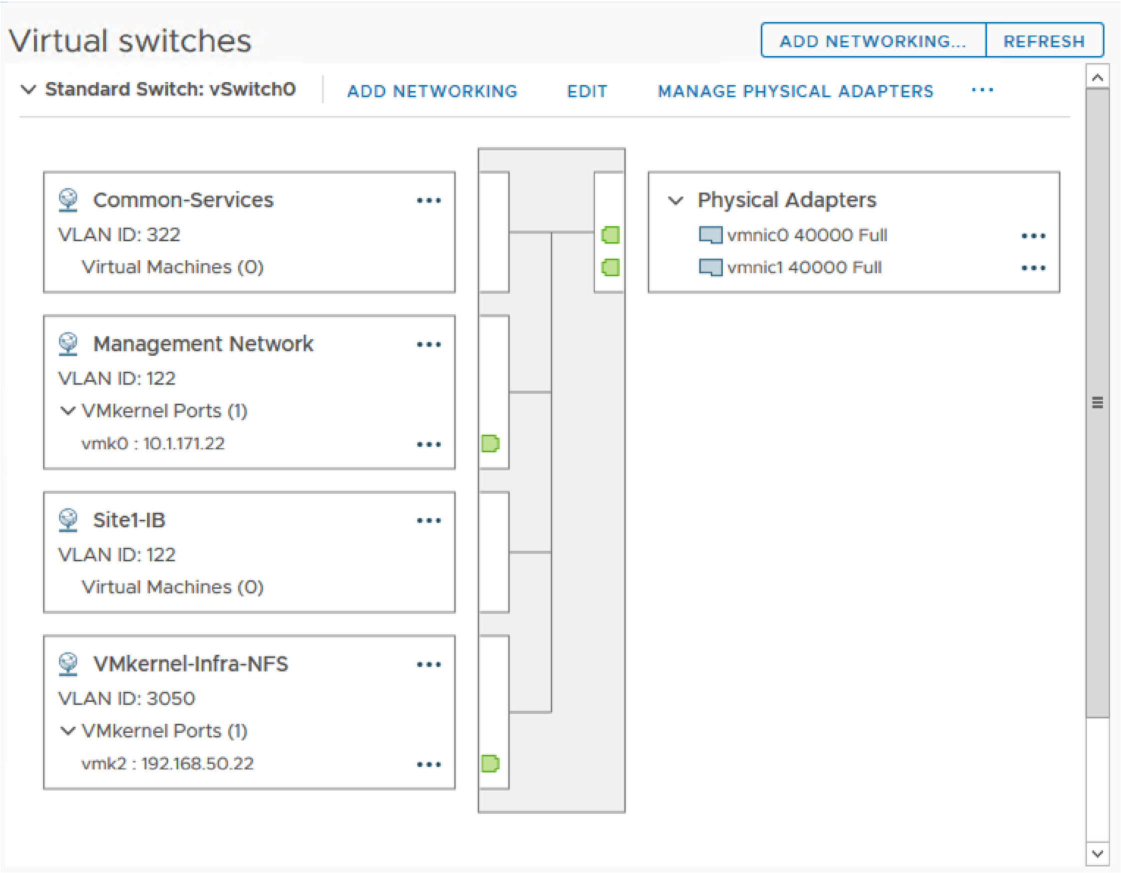

30. Repeat this procedure for all hosts being added.

Mount Required Datastores

ESXi Hosts created other than VM-Host-Infra-01

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

1. From the vCenter Home screen, choose Menu > Storage.
2. Located on the left, expand FlexPod-DC.
3. Located on the left, right-click infra_datastore_1 and choose Mount Datastore to Additional Hosts.
4. Choose the ESXi host(s) and click OK.

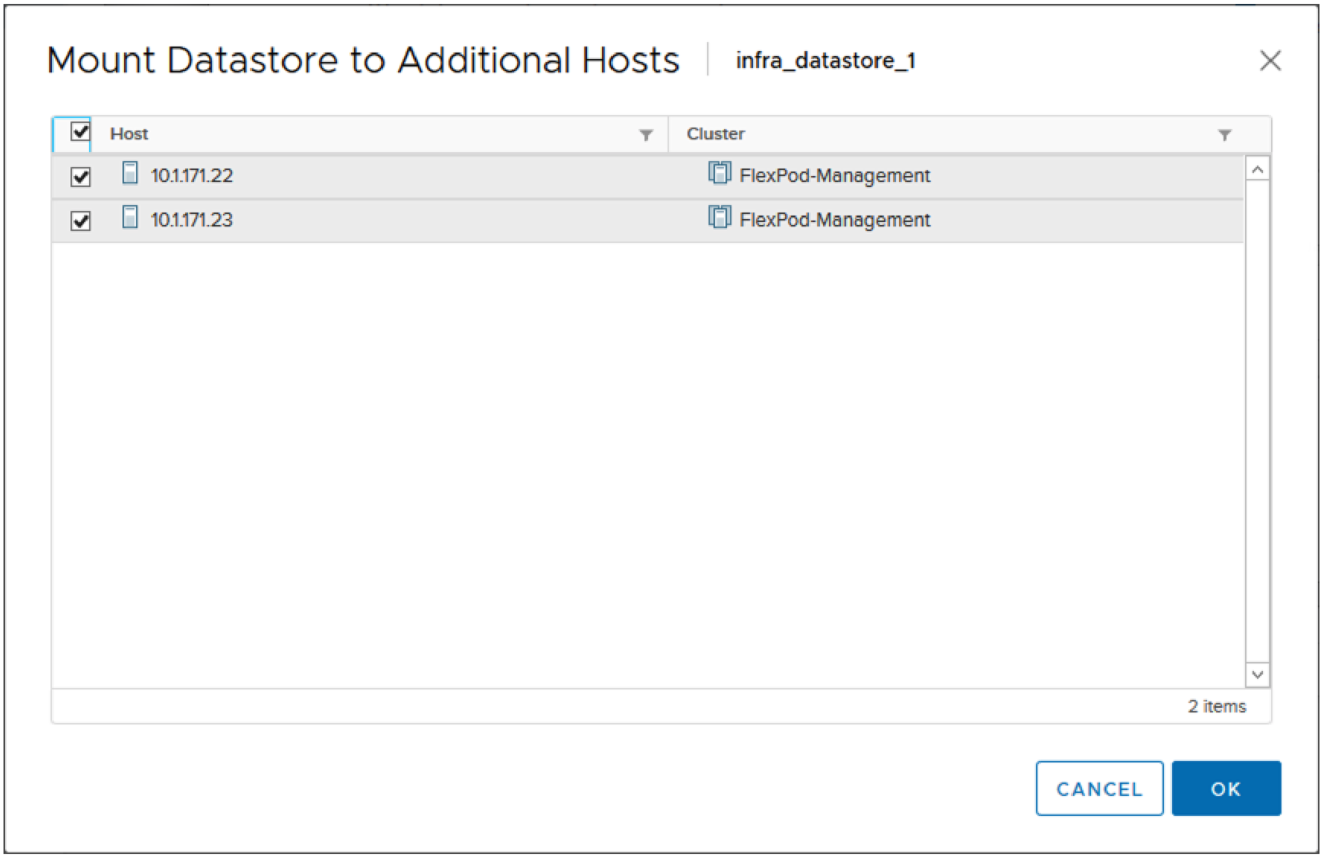

☑	Host	Cluster
☑	10.1.171.22	FlexPod-Management
☑	10.1.171.23	FlexPod-Management

2 items

5. Repeat steps 1-4 to mount the infra_swap and any additional datastores created to the ESXi host(s).

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

Configure NTP on ESXi Host

ESXi Hosts created other than VM-Host-Infra-01

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

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, choose the Configure tab.

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

4. To the right of Manual Time Configuration, click EDIT.

5. Set the correct local time and click OK.

6. To the right of Network Time Protocol, click EDIT.

7. Choose the Enable checkbox.

8. Enter the two Nexus switch NTP IP addresses in the NTP servers box separated by a comma.

9. Click the Start NTP Service checkbox.

10. Use the drop-down list to choose Start and stop with host.

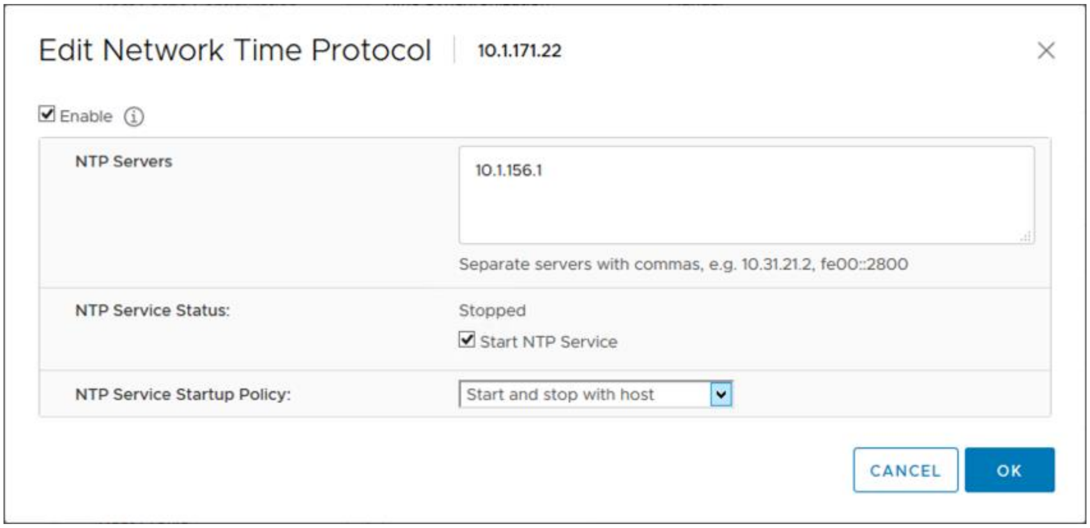

Edit Network Time Protocol | 10.1.171.22

☑ Enable

NTP Servers: 10.1.156.1

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

NTP Service Status: Stopped

☑ Start NTP Service

NTP Service Startup Policy: Start and stop with host

11. Click OK to save the configuration changes.

12. Verify that NTP service is now enabled and running and the clock is now set to approximately the correct time.

Configure ESXi Host Swap

ESXi Hosts created other than VM-Host-Infra-01

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

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, choose the Configure tab.

3. In the list under System, choose System Swap.

4. Located on the right, click EDIT.

5. Choose Can use datastore and use the drop-down list to choose infra_swap. Leave all other settings unchanged.

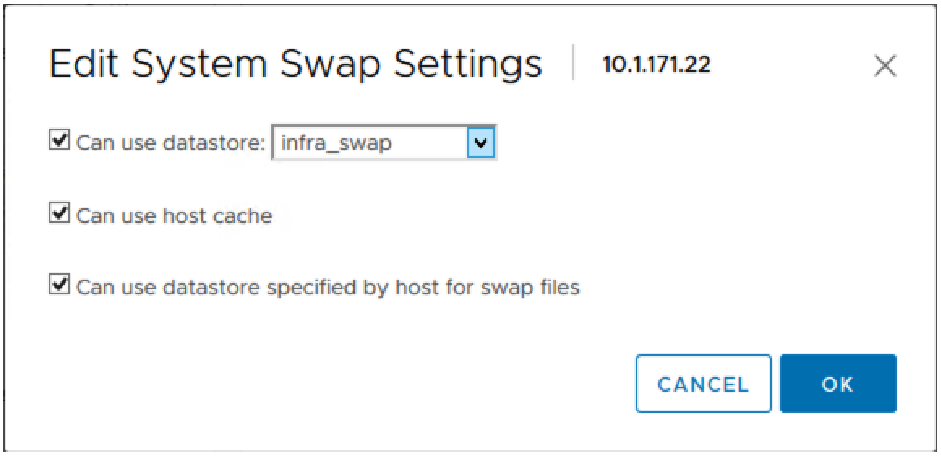

Edit System Swap Settings | 10.1.171.22

☑ Can use datastore: infra_swap

☑ Can use host cache

☑ Can use datastore specified by host for swap files

6. Click OK to save the configuration changes.

7. In the list under Virtual Machines, choose Swap File Location.

8. Located on the right, click EDIT.

9. Choose infra_swap and click OK.

Change ESXi Power Management Policy

ESXi Hosts created other than VM-Host-Infra-01

To change the ESXi power management policy, follow these steps:

1. In the list under Hardware, choose Overview. Scroll to the bottom and to the right of Power Management, choose EDIT POWER POLICY.

2. Choose High performance and click OK.

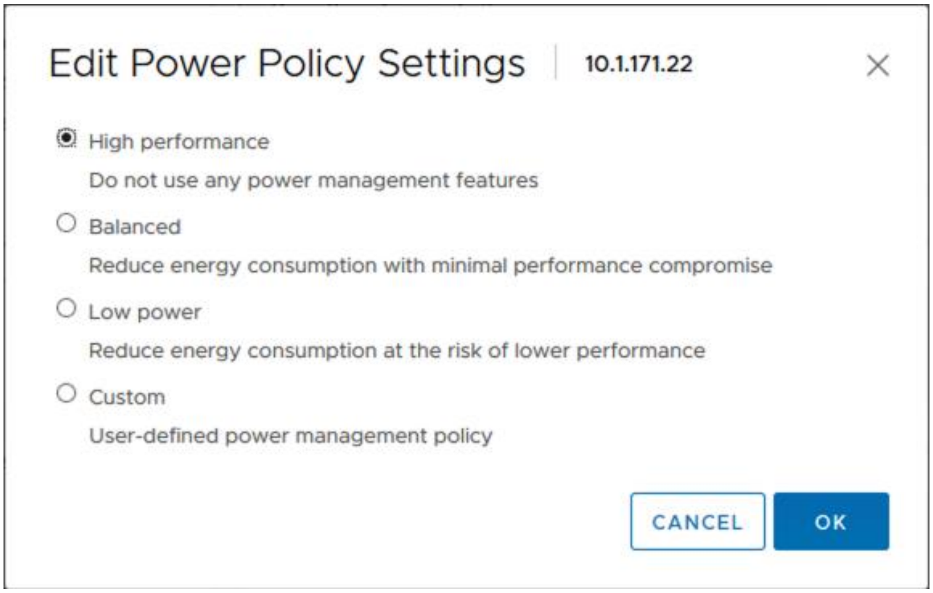

Edit Power Policy Settings | 10.1.171.22

● High performance
Do not use any power management features

○ Balanced
Reduce energy consumption with minimal performance compromise

○ Low power
Reduce energy consumption at the risk of lower performance

○ Custom
User-defined power management policy

Add iSCSI Configuration

All ESXi Hosts

To add the iSCSI configuration to the ESXi hosts, follow these steps:

1. In the vSphere HTML5 Client, under Hosts and Clusters, choose the ESXi host.

2. In the center pane, click Configure. In the list under Networking, select Virtual switches.

3. In the center pane, expand iScsiBootvSwitch. Click EDIT to edit settings for the vSwitch.

4. Change the MTU to 9000 and click OK.

5. Choose ... > Edit Settings to the right of iScsiBootPG. Change the Network label to iScsiBootPG-A and click OK.

6. Choose ... > Edit Settings to the right of the VMkernel Port IP address. Change the MTU to 9000.

7. Click IPv4 settings on the left. Change the IP address to a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.

> It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

8. Click OK.

9. In the upper right-hand corner, choose ADD NETWORKING to add another vSwitch.

10. Make sure VMkernel Network Adapter is selected and click NEXT.

11. Choose New standard switch and change the MTU to 9000. Click NEXT.

12. Choose ➕ to add an adapter. Make sure vmnic5 is highlighted and click OK. vmnic5 should now be under Active adapters. Click NEXT.

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

14. Choose Use static IPv4 settings. Enter a unique IP address and netmask in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B. Click NEXT.

15. Click FINISH to complete creating the vSwitch and the VMkernel port.

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

17. Choose the iSCSI Software Adapter and below, choose the Dynamic Discovery tab.

18. Click Add.

19. Enter the IP address of the storage controller's Infra-SVM LIF iscsi-lif-01a and click OK.

20. Repeat this process to add the IPs for iscsi-lif-02a, iscsi-lif-01b, and iscsi-lif-02b.

21. Under Storage Adapters, click Rescan Adapter to rescan the iSCSI Software Adapter.

22. Under Static Discovery, four static targets should now be listed.

23. Under Paths, four paths should now be listed with two of the paths having the "Active (I/O)" Status.

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

ESXi Hosts created other than VM-Host-Infra-01

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

1. After logging into the VMware vSphere HTML5 Client, choose Networking under Menu.

2. Right-click the vDS (vDS0) and click Add and Manage Hosts.

3. Make sure Add hosts is selected and click NEXT.

4. Click the green + sign to add New hosts. Choose the configured FlexPod Management host(s) and click OK. Click NEXT.

5. Choose vmnic2 on each host and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 on each host and click Assign uplink. Choose Uplink 2 and click OK. If more than one host is being connected to the vDS, use the Apply this uplink assignment to the rest of the hosts checkbox.

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

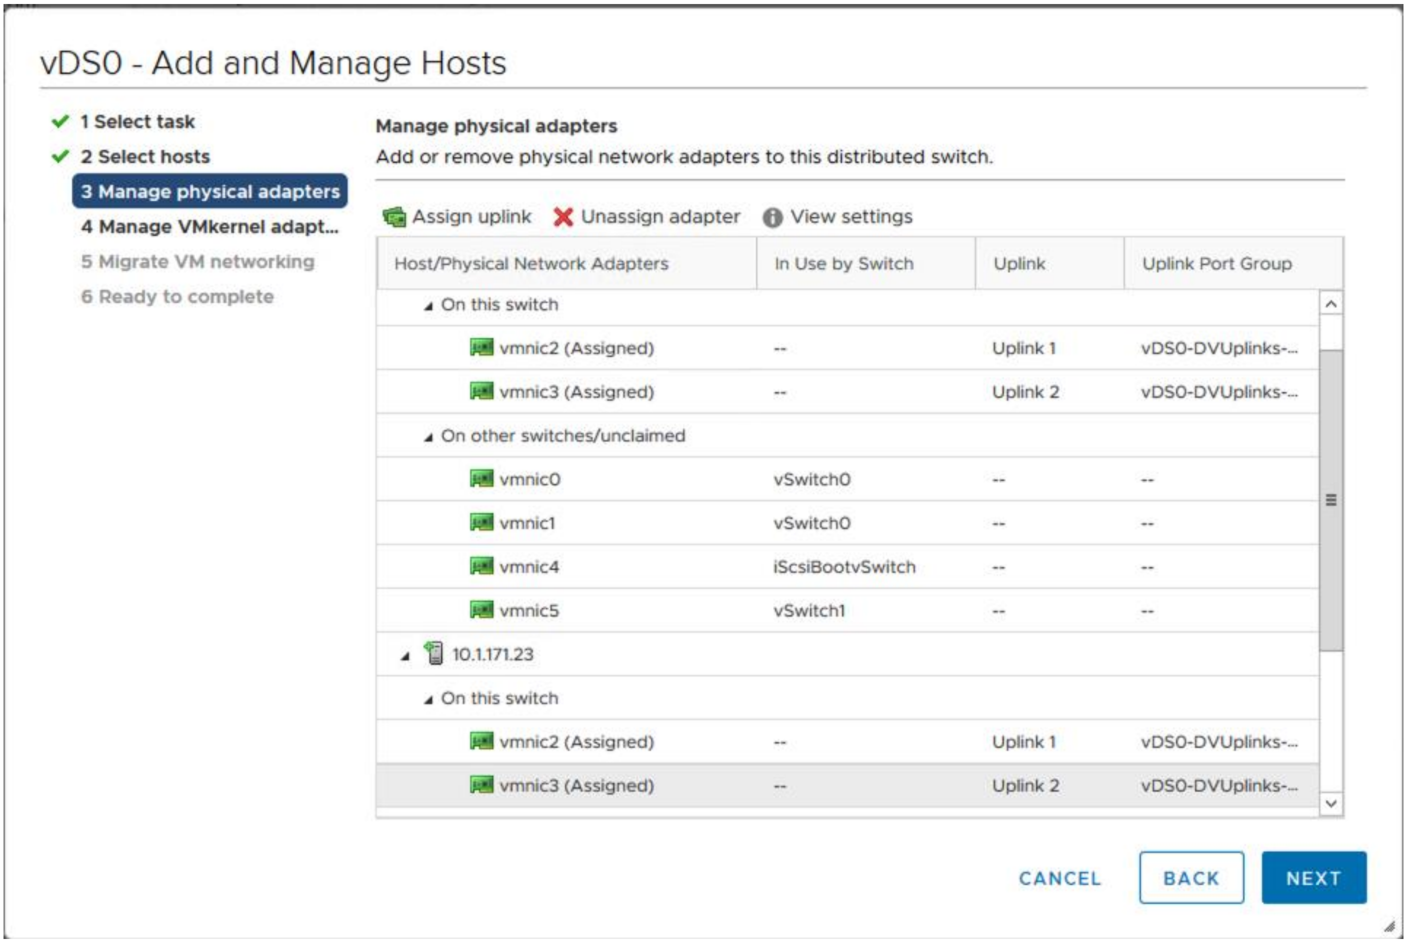

vDS0 - Add and Manage Hosts

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

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic4	iScsiBootvSwitch	--	--
vmnic5	vSwitch1	--	--
10.1.171.23			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...

6. Click NEXT.

7. Do not migrate any VMkernel ports and click NEXT.

8. Do not migrate any VM ports and click NEXT.

9. Click FINISH to complete adding the ESXi host(s) to the vDS.

Add the vMotion VMkernel Port(s) to the ESXi Host

All ESXi Hosts

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

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, click the Configure tab.

3. In the list under Networking, choose VMkernel adapters.

4. Choose Add Networking to Add host networking.

5. Make sure VMkernel Network Adapter is selected and click NEXT.

6. Choose BROWSE to the right of Select an existing network.

7. Choose vMotion on the vDS and click OK.

8. Click NEXT.

9. Make sure the Network label is vMotion with the vDS in parenthesis. From the drop-down list, select Custom for MTU and make sure the MTU is set to 9000. Choose the vMotion TCP/IP stack and click NEXT.

10. Choose Use static IPv4 settings and input the host's vMotion IPv4 address and Subnet mask.

11. Click NEXT.

12. Review the parameters and click FINISH to add the vMotion VMkernel port.

13. Optionally, repeat this process to add two more vMotion VMkernel ports.

NetApp Virtual Storage Console 9.7.1 Deployment Procedure

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

Virtual Storage Console 9.7.1 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run ONTAP 9.7.1 or above:

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

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

Table 23. Port Requirements for VSC

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

The requirements for deploying VSC are listed here.

Install Virtual Storage Console 9.7.1

To install the VSC 9.7.1 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

1. Login to vCenter and navigate to Hosts and Clusters.
2. Select ACTIONS for the FlexPod-DC datacenter and choose Deploy OVF Template.
3. Browse to the VSC OVA file downloaded from the NetApp Support site.

4. Enter the VM name and choose a datacenter or folder in which to deploy and click NEXT.

5. Choose a host cluster resource in which to deploy OVA and click NEXT.

6. Review the details and accept the license agreement.

7. Choose the infra_datastore_1 volume and choose the Thin Provision option for the virtual disk format.

8. From Select Networks, choose a destination network (Site1-IB) and click NEXT.

9. From Customize Template, enter the VSC administrator password, vCenter name or IP address and other configuration details and click NEXT.

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

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

12. Verify that VSC, VASA Provider, and SRA services are running after the deployment is completed.

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

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

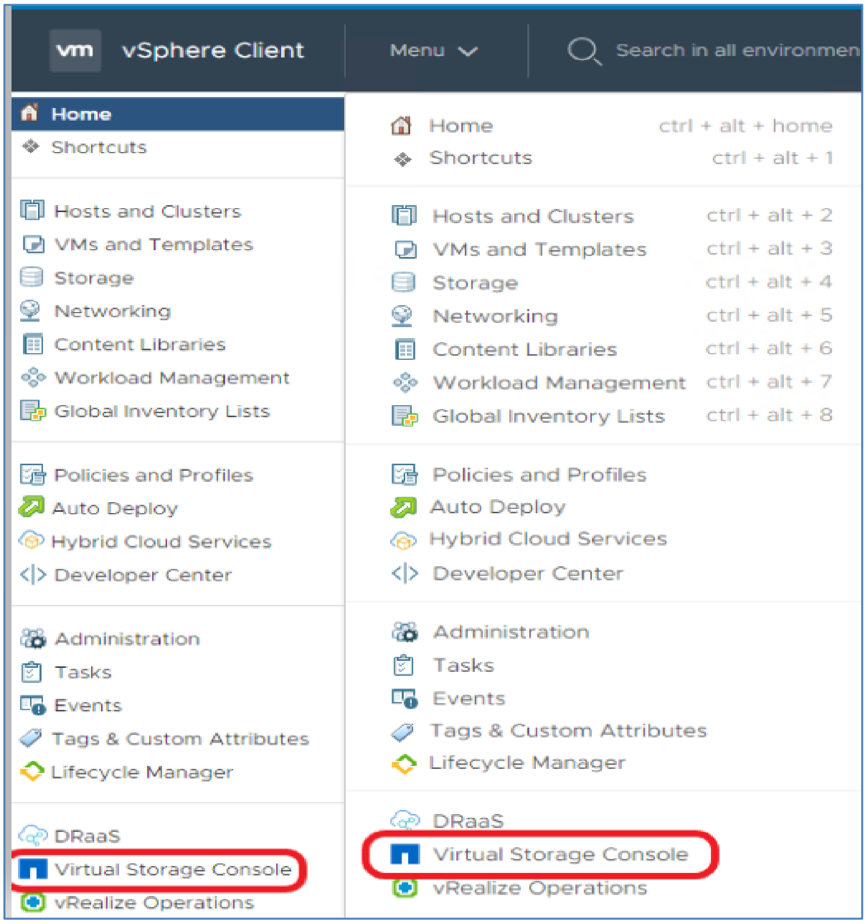

If the virtual appliance for VSC, VASA Provider, and SRA is not registered with any vCenter Server, use **https://appliance_ip:8143/Register.html** to register the VSC instance.

Download the NetApp NFS Plug-in for VAAI

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

1. Download the NetApp NFS Plug-in 1.1.2 for VMware .vib file from the NFS Plugin Download page and save it to your local machine or admin host.

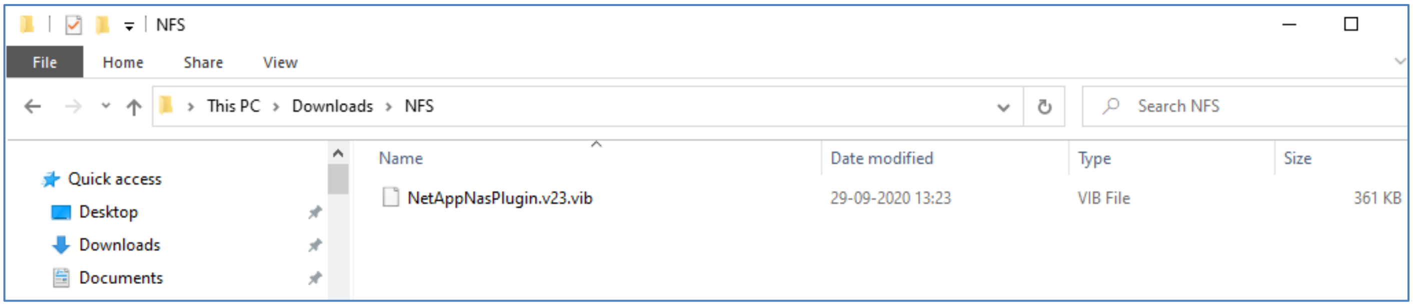

Name	Date modified	Type	Size
NetAppNasPlugin.v23.vib	29-09-2020 13:23	VIB File	361 KB

Install the NetApp NFS Plug-in for VAAI

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

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

1. Rename the .vib file that you downloaded from the NetApp Support Site to NetAppNasPlugin.vib to match the predefined name that VSC uses.

2. Click Settings in the VSC Getting Started page.

3. Click NFS VAAI Tools tab.

4. Click Change in the Existing version section.

5. Browse and choose the renamed .vib file, and then click Upload to upload the file to the virtual appliance.

6. In the Install on ESXi Hosts section, choose the ESXi host on which you want to install the NFS Plug-in for VAAI, and then click Install.

7. Reboot the ESXi host after the installation finishes.

Verify the VASA Provider

The VASA provider for ONTAP is enabled by default during the installation of the NetApp Virtual Storage Console (VSC) 9.7.1. To verify the VASA provider was enabled, follow these steps:

1. From the vSphere Client, click Menu > Virtual Storage Console. Click Settings.

2. Click Manage Capabilities in the Administrative Settings tab.

3. In the Manage Capabilities dialog box if not enabled, click Enable VASA Provider slider.

4. Enter the IP address of the virtual appliance for VSC, VASA Provider, and VMware Storage Replication Adapter (SRA) and the administrator password, and then click Apply.

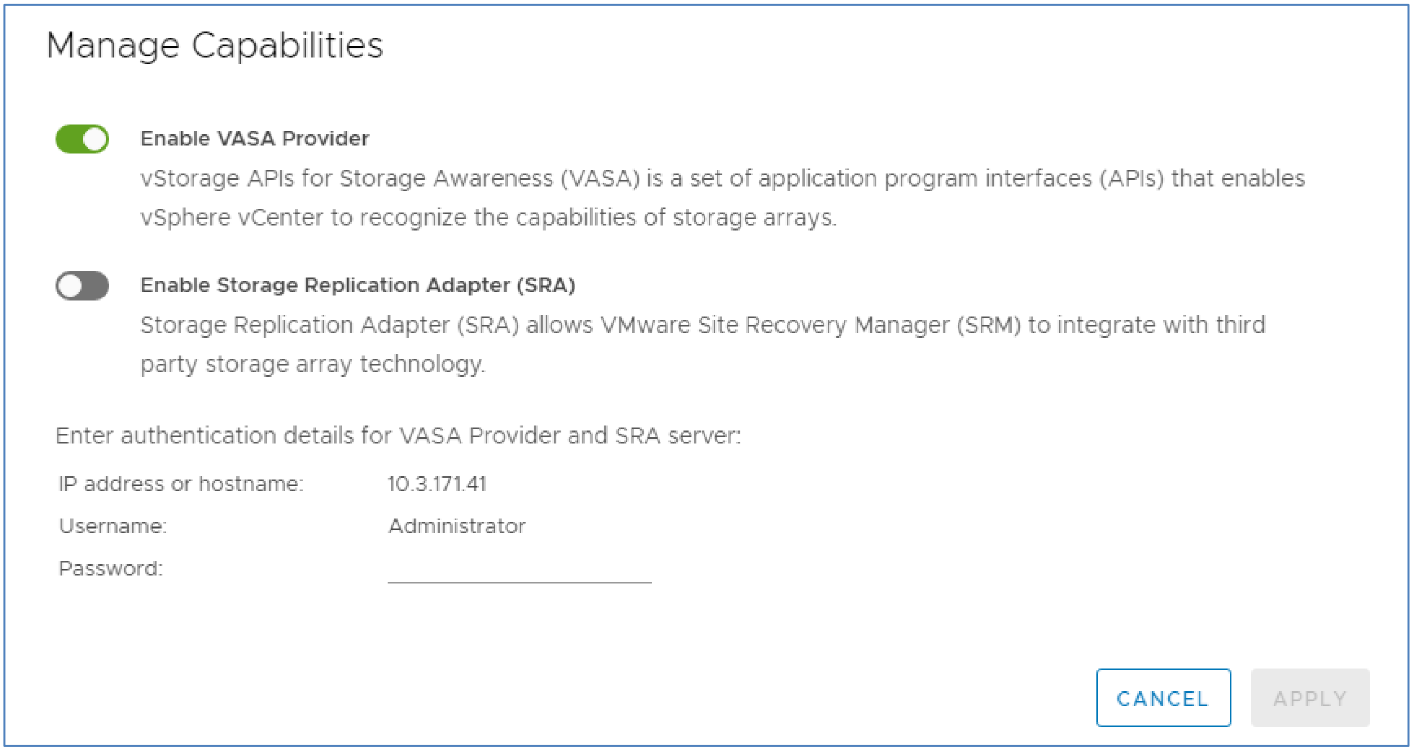

Manage Capabilities

Enable VASA Provider

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

Enable Storage Replication Adapter (SRA)

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

Enter authentication details for VASA Provider and SRA server:

IP address or hostname: 10.3.171.41

Username: Administrator

Password:

Discover and Add Storage Resources

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

1. Using the vSphere Web Client, log in to the vCenter Server as the FlexPod admin user. If the vSphere Web Client was previously opened, close the tab, and then reopen it.

2. In the Home screen, click the Home tab and click Virtual Storage Console.

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

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

3. Choose Storage Systems > Add.

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

5. Specify the vCenter Server instance where the storage will be located.

6. In the IP Address/Hostname field, enter the storage cluster management IP.

7. Confirm Port 443 to Connect to this storage system.

8. Enter admin for the user name and the admin password for the cluster.

9. Click Save to add the storage configuration to VSC.

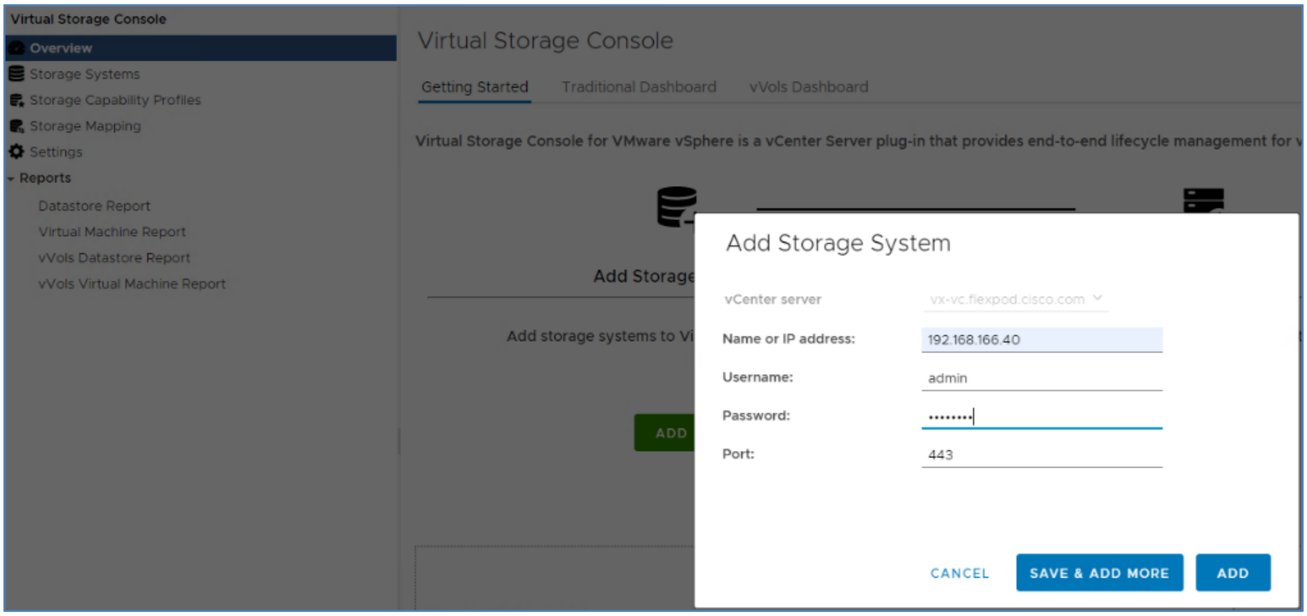

Virtual Storage Console for VMware vSphere is a vCenter Server plug-in that provides end-to-end lifecycle management for v

Add storage systems to Vi

Add Storage System

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

Name or IP address: 192.168.166.40

Username: admin

Password: ••••••••

Port: 443

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

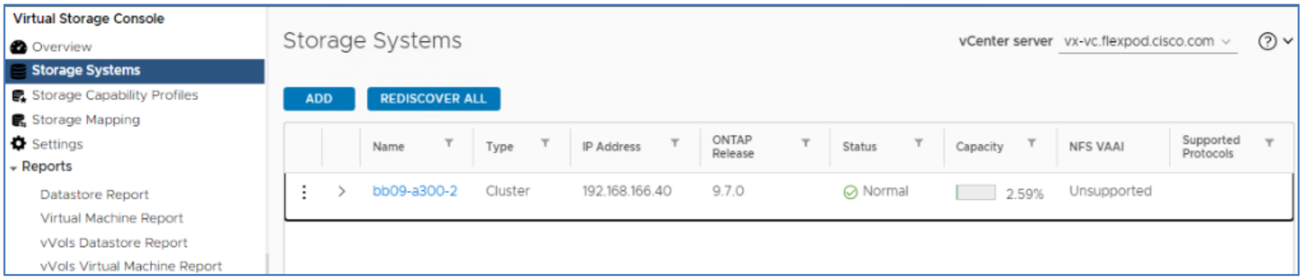

vCenter server: vx-vc.flexpod.cisco.com

Name	Type	IP Address	ONTAP Release	Status	Capacity	NFS VAAI	Supported Protocols
bb09-a300-2	Cluster	192.168.166.40	9.7.0	Normal	2.59%	Unsupported	

Discover the Cluster and SVMs

To Discover the cluster and SVMs with the cluster admin account, follow these steps:

1. From the vSphere Client Home page, click Hosts and Clusters.

2. Right-click the FlexPod-DC datacenter then click NetApp VSC > Update Host and Storage Data and click Yes.

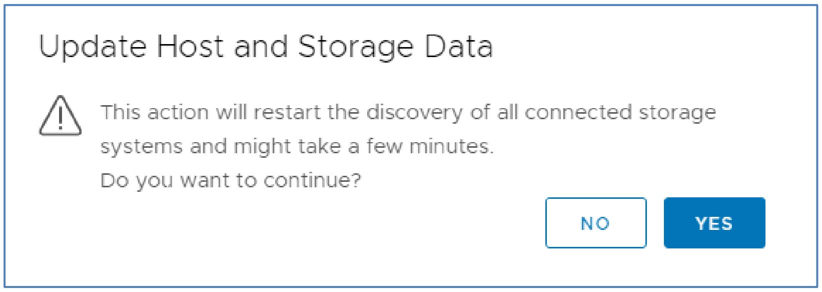

Update Host and Storage Data

This action will restart the discovery of all connected storage systems and might take a few minutes.
Do you want to continue?

3. Click OK.

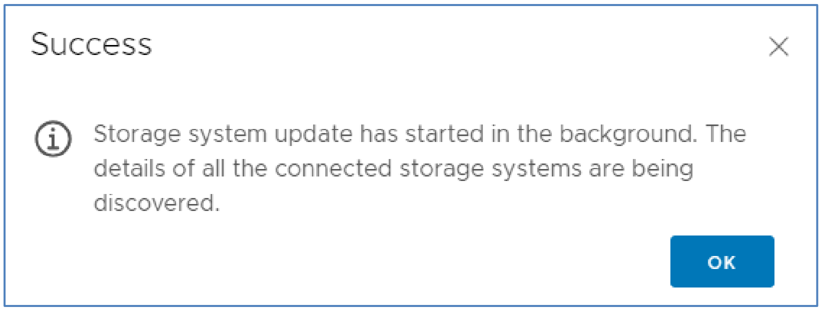

Success

Storage system update has started in the background. The details of all the connected storage systems are being discovered.

Optimal Storage Settings for ESXi Hosts

VSC enables the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, follow these steps:

1. From the VMware vSphere Web Client Home page, click vCenter > Hosts.
2. Choose a host and then click Actions > NetApp VSC > Set Recommended Values.
3. In the NetApp Recommended Settings dialog box, choose all the values for your system.

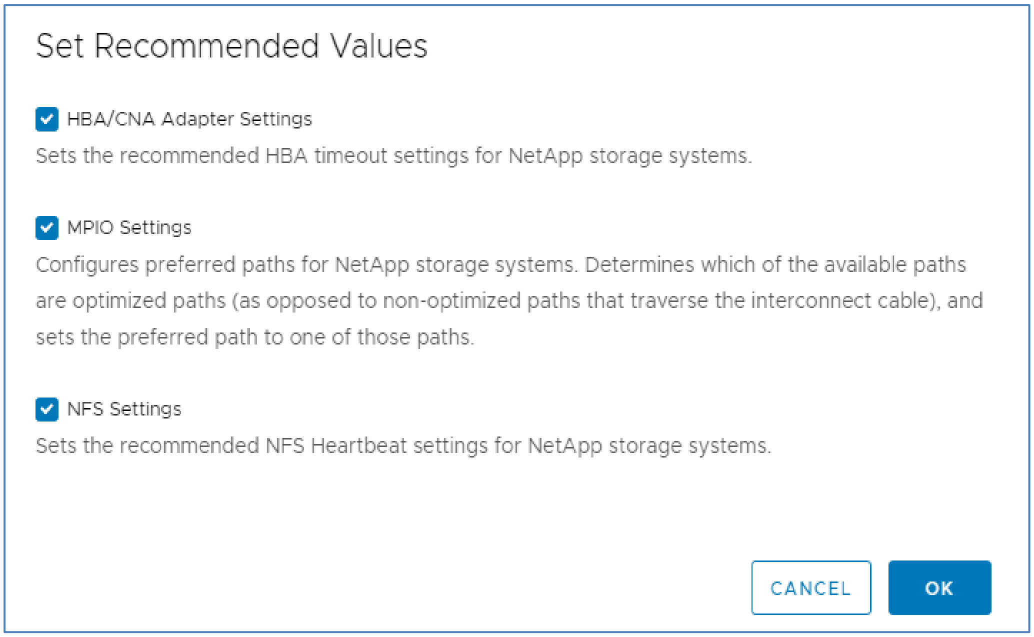

Set Recommended Values

☑ 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.

This functionality sets values for HBAs and converged network adapters (CNAs), sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS). A vSphere host reboot may be required after applying the settings.

4. Click OK.

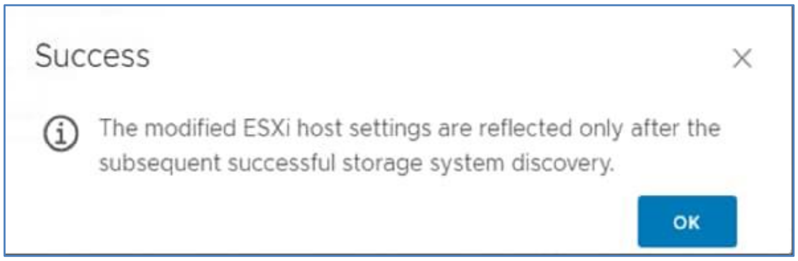

Success

The modified ESXi host settings are reflected only after the subsequent successful storage system discovery.

Virtual Storage Console 9.7.1 Provisioning Datastores

Using VSC, the administrator can provision an NFS, FC or iSCSI datastore and attach it to a single host or multiple hosts in the cluster. The following steps describe provisioning a datastore and attaching it to the cluster.

It is a NetApp best practice to use Virtual Storage Console (VSC) to provision datastores for the FlexPod infrastructure. When using VSC to create vSphere datastores, all NetApp storage best practices are implemented during volume creation and no additional configuration is needed to optimize performance of the datastore volumes.

Storage Capabilities

A storage capability is a set of storage system attributes that identifies a specific level of storage performance (storage service level), storage efficiency, and other capabilities such as encryption for the storage object that is associated with the storage capability.

Create the Storage Capability Profile

To leverage the automation features of VASA two primary components must first be configured. The Storage Capability Profile (SCP) and the VM Storage Policy. The Storage Capability Profile expresses a specific set of storage characteristics into one or more profiles used to when provisioning a Virtual Machine. The SCP is specified as part of VM Storage Policy which is specified when you deploy a virtual machine. NetApp Virtual Storage Console comes with two pre-configured Storage Capability Profiles- Platinum and Bronze.

Adaptive QoS policies are not currently supported with VSC 9.7.1. Storage Capability Profiles (SCP) can still be created with Max IOPS and Min IOPS defined.

To review or edit one of the built-in profiles pre-configured with VSC 9.7.1 follow these steps:

1. In the NetApp Virtual Storage Console click Storage Capability Profiles.

2. Choose the Platinum Storage Capability Profile and choose Clone from the toolbar.

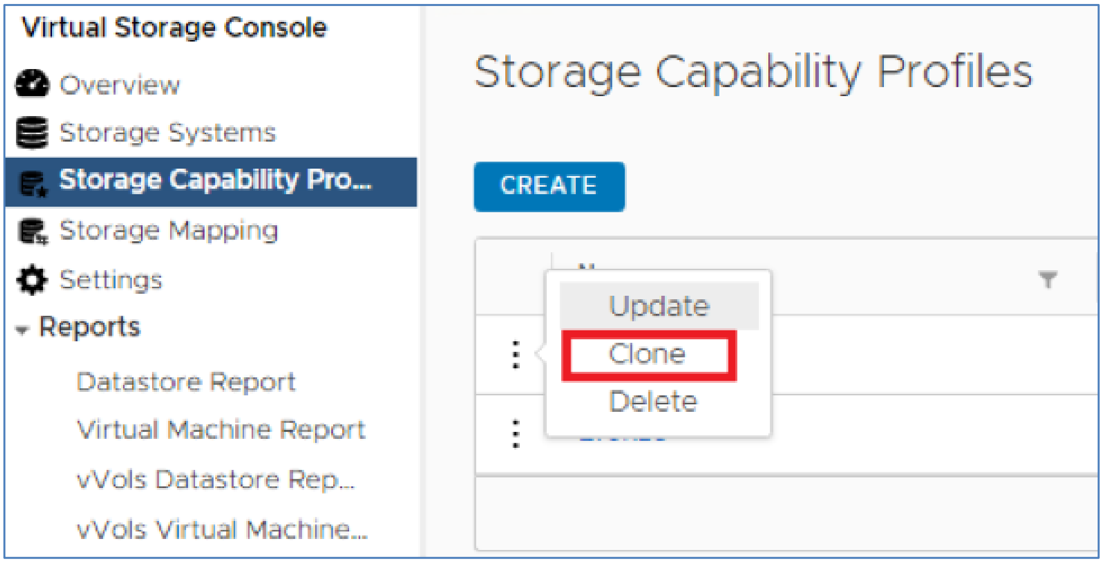

3. Enter a name for the cloned SCP and add a description if desired.

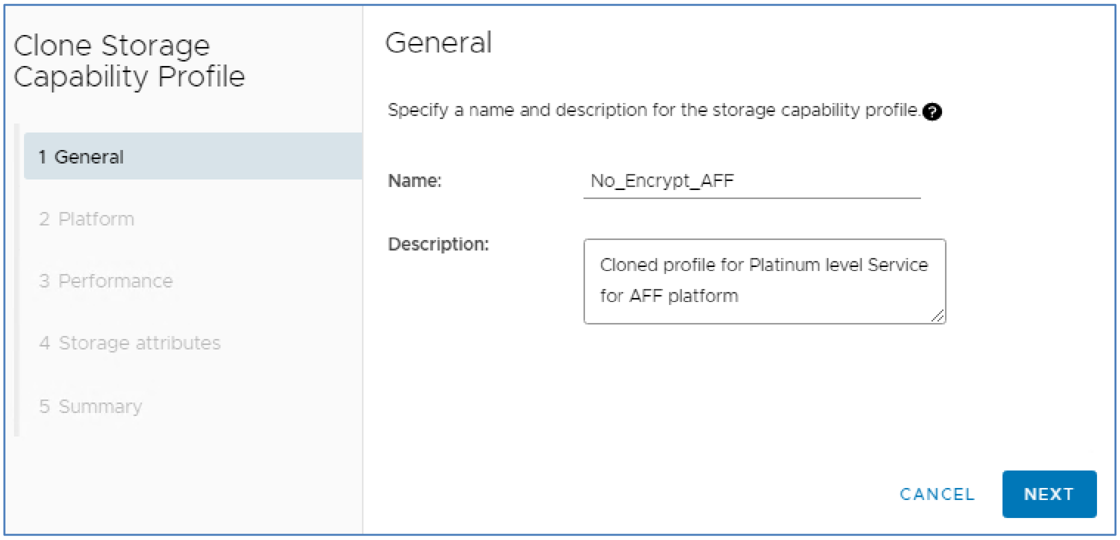
Clone Storage Capability Profile

General

Specify a name and description for the storage capability profile.

Name: No_Encrypt_AFF

Description: Cloned profile for Platinum level Service for AFF platform

4. Choose ALL Flash FAS(AFF) for the storage platform. Click Next.

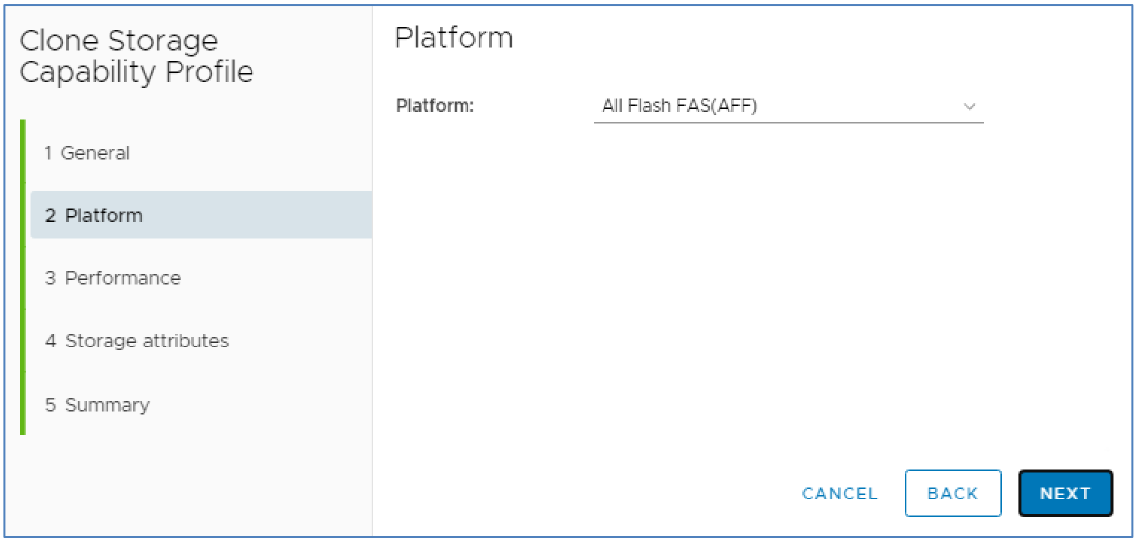
Clone Storage Capability Profile

Platform

Platform: All Flash FAS(AFF)

5. Choose None to allow unlimited performance or set a the desired minimum and maximum IOPS for the QoS policy group.

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

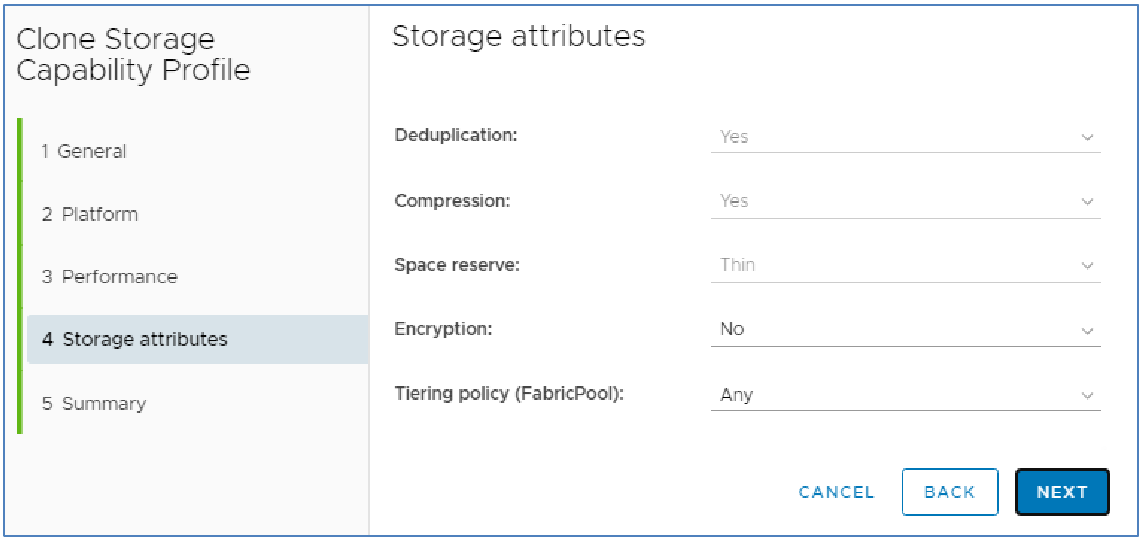

Clone Storage Capability Profile

1 General

2 Platform

3 Performance

4 Storage attributes

5 Summary

Storage attributes

Deduplication: Yes

Compression: Yes

Space reserve: Thin

Encryption: No

Tiering policy (FabricPool): Any

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

As a best practice it is always recommended to create clone and edit the Capability profile rather than changing the Default one.

Create a VM Storage Policy

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

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

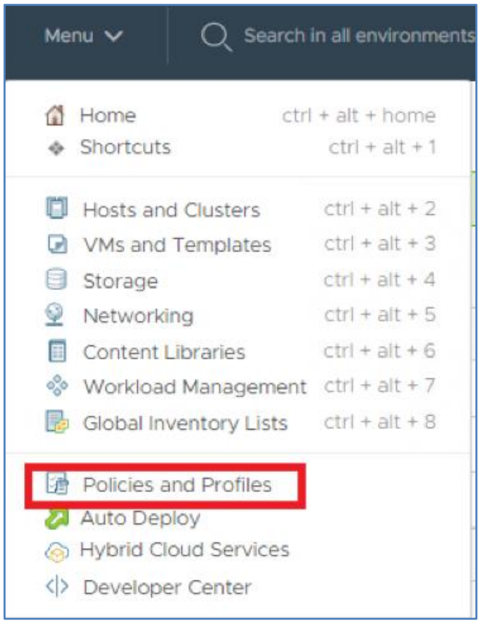

Menu
Search in all environments
Home ctrl + alt + home
Shortcuts ctrl + alt + 1
Hosts and Clusters ctrl + alt + 2
VMs and Templates ctrl + alt + 3
Storage ctrl + alt + 4
Networking ctrl + alt + 5
Content Libraries ctrl + alt + 6
Workload Management ctrl + alt + 7
Global Inventory Lists ctrl + alt + 8
Policies and Profiles
Auto Deploy
Hybrid Cloud Services
Developer Center

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

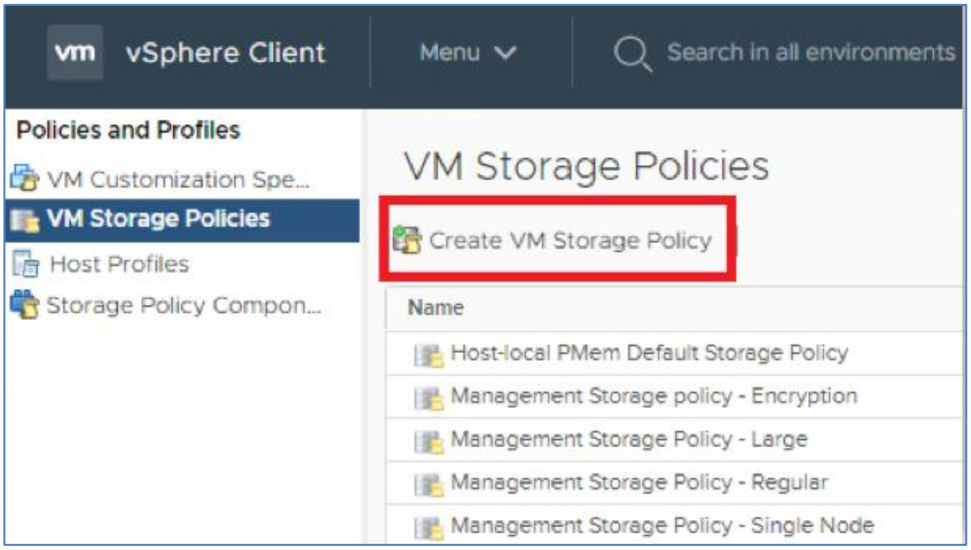
vm
vSphere Client
Menu
Search in all environments
Policies and Profiles
VM Customization Spe...
VM Storage Policies
Host Profiles
Storage Policy Compon...
VM Storage Policies
Create VM Storage Policy
Name
Host-local PMem Default Storage Policy
Management Storage policy - Encryption
Management Storage Policy - Large
Management Storage Policy - Regular
Management Storage Policy - Single Node

3. Create a name for the VM storage policy and enter a description and click NEXT.

4. Choose Enable rules for NetApp.clustered.Data.ONTAP.VP.VASA10 storage located under the Datastore specific rules section and click NEXT.

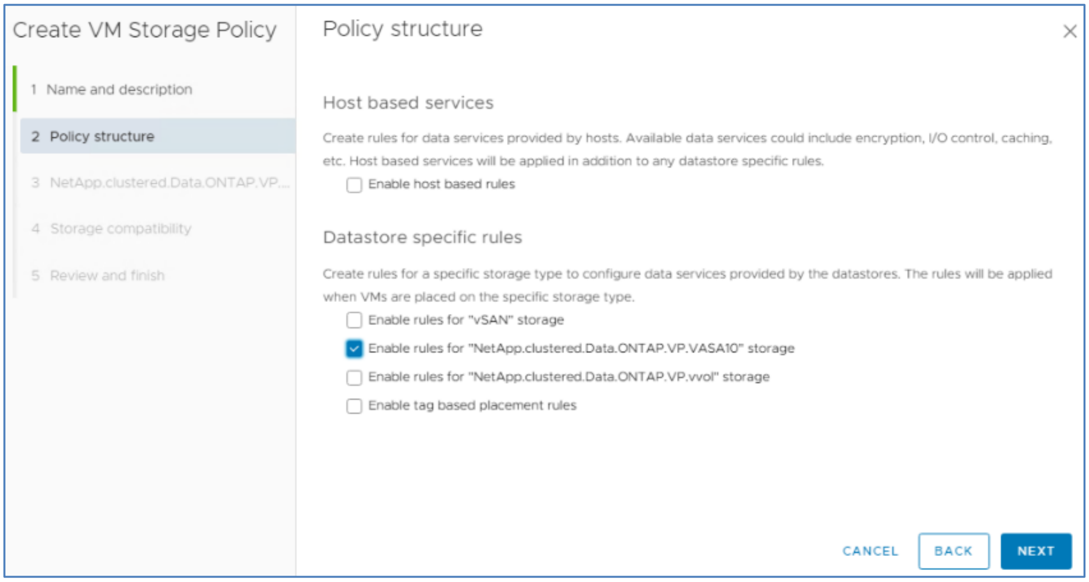

Policy structure

Host based services

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

☐ Enable host based rules

Datastore specific rules

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

☐ Enable rules for "vSAN" storage
☑ Enable rules for "NetApp.clustered.Data.ONTAP.VP.VASA10" storage
☐ Enable rules for "NetApp.clustered.Data.ONTAP.VP.vvol" storage
☐ Enable tag based placement rules

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

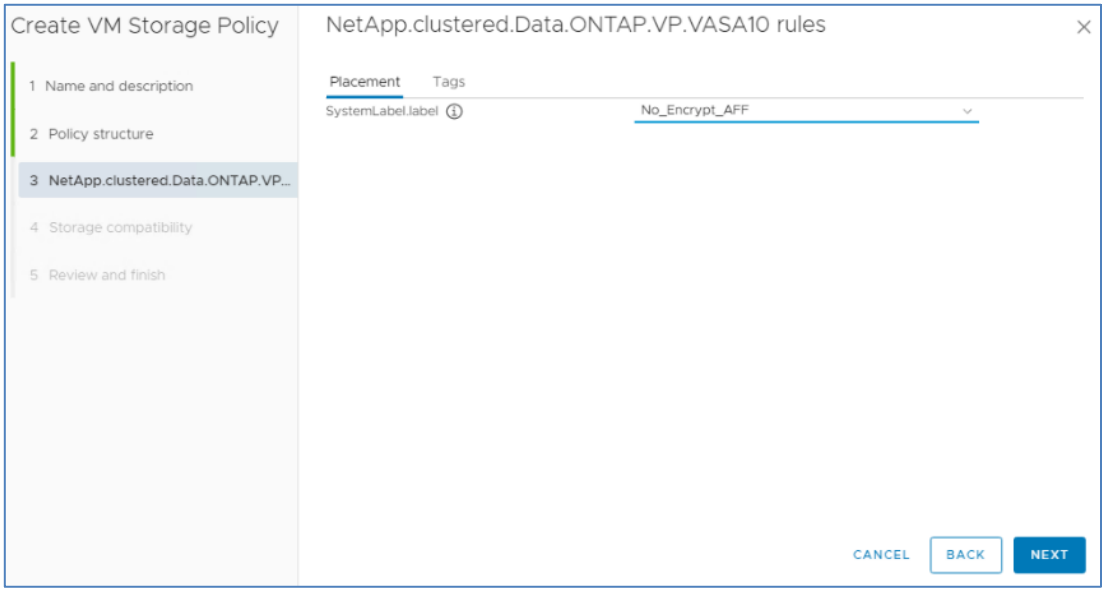

Create VM Storage Policy

1 Name and description

2 Policy structure

3 NetApp.clustered.Data.ONTAP.VP...

4 Storage compatibility

5 Review and finish

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

Placement Tags

SystemLabel.label No_Encrypt_AFF

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

7. Review the policy summary and click FINISH.

Provision NFS Datastore

To provision the NFS datastore, follow these steps:

1. From the Virtual Storage Console Home page, click Overview.

2. In the Getting Started tab, click Provision.

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

4. Choose the type as NFS and Enter the datastore name.

5. Provide the size of the datastore and the NFS Protocol.

6. Check the storage capability profile and click NEXT.

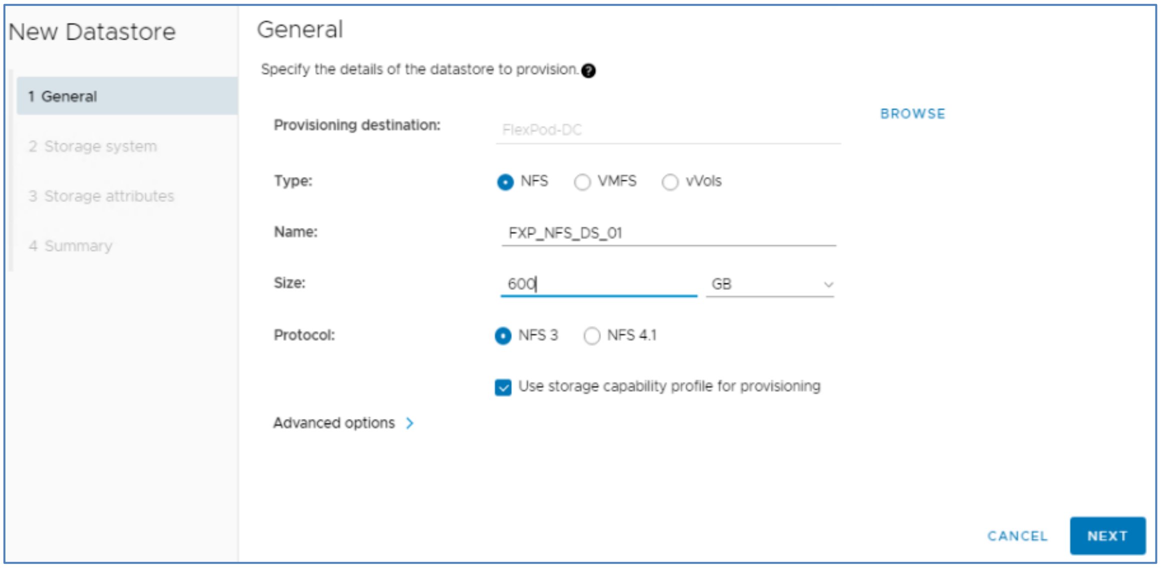

New Datastore

1 General

2 Storage system

3 Storage attributes

4 Summary

General

Specify the details of the datastore to provision.

Provisioning destination: FlexPod-DC

Type: ● NFS ○ VMFS ○ vVols

Name: FXP_NFS_DS_01

Size: 600 GB

Protocol: ● NFS 3 ○ NFS 4.1

☑ Use storage capability profile for provisioning

7. Choose the desired Storage Capability Profile, cluster name, and the desired SVM to create the datastore. In this example, the Infra-SVM is selected.

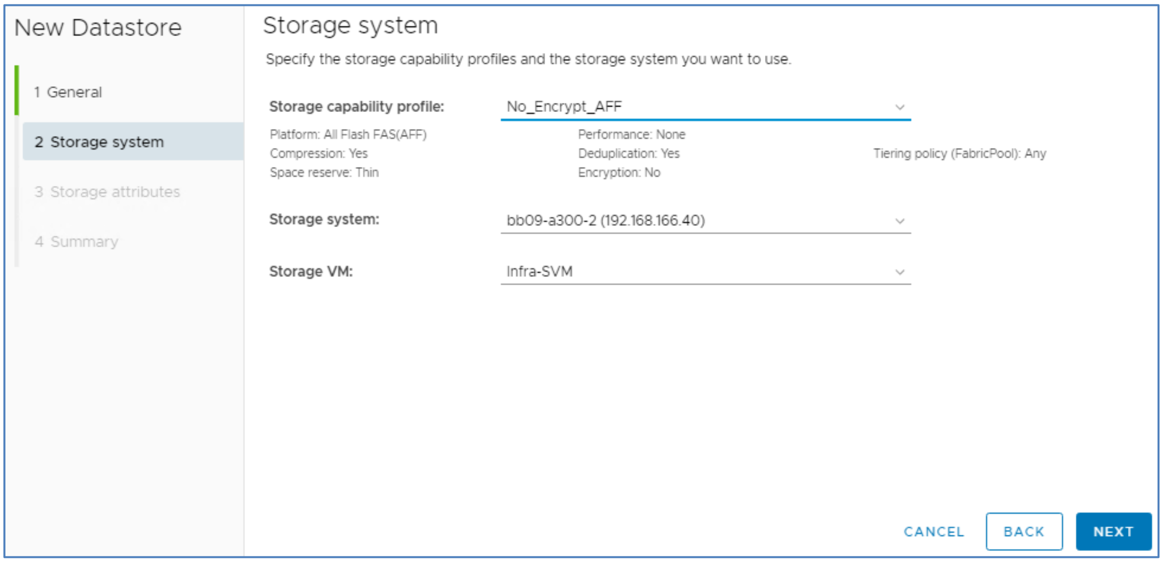

New Datastore

1 General

2 Storage system

3 Storage attributes

4 Summary

Storage system

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

Storage capability profile: No_Encrypt_AFF

Platform: All Flash FAS(AFF)
Compression: Yes
Space reserve: Thin

Performance: None
Deduplication: Yes
Encryption: No

Tiering policy (FabricPool): Any

Storage system: bb09-a300-2 (192.168.166.40)

Storage VM: Infra-SVM

8. Click NEXT.

9. Choose the aggregate name and click NEXT.

10. Review the Summary and click FINISH.

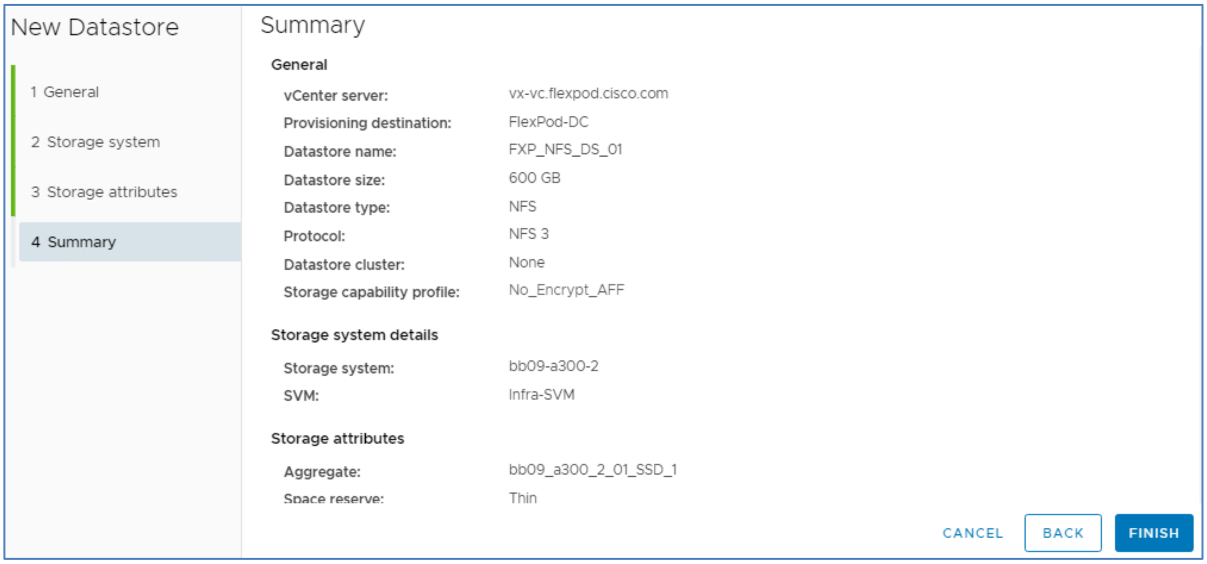

Summary

General

vCenter server: vx-vc.flexpod.cisco.com
Provisioning destination: FlexPod-DC
Datastore name: FXP_NFS_DS_01
Datastore size: 600 GB
Datastore type: NFS
Protocol: NFS 3
Datastore cluster: None
Storage capability profile: No_Encrypt_AFF

Storage system details

Storage system: bb09-a300-2
SVM: Infra-SVM

Storage attributes

Aggregate: bb09_a300_2_01_SSD_1
Space reserve: Thin

The datastore is created and mounted on the hosts in the cluster. Click Refresh from the vSphere Web Client to see the newly created datastore or it is also listed in the VSC home page > Traditional Dashboard > Datastores view. Also, VSC Home page > Reports > Datastore Report should be listing the newly created datastore.

Virtual Volumes(vVols)

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

Lab testing has shown that if a virtual machine (VM) has one or more disks in vVol datastores and the VM is migrated to another host, just at the end of the migration the VM can be stunned or frozen for 45 or more seconds. Total observed vMotion time noted as approximately 80 seconds.

Verify NDMP Vserver Scope Mode

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

1. View NDMP scope mode with the following command:

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

2. Disable NDMP node-scoped mode.

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

3. Enable NDMP services on the vserver.

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

Create the Storage Capability Profile

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

To create storage capability profile for the vVol datastore, follow these steps:

1. In the NetApp Virtual Storage Console click Storage Capability Profiles.

2. Choose the Platinum Storage Capability Profile and choose Clone from the toolbar.

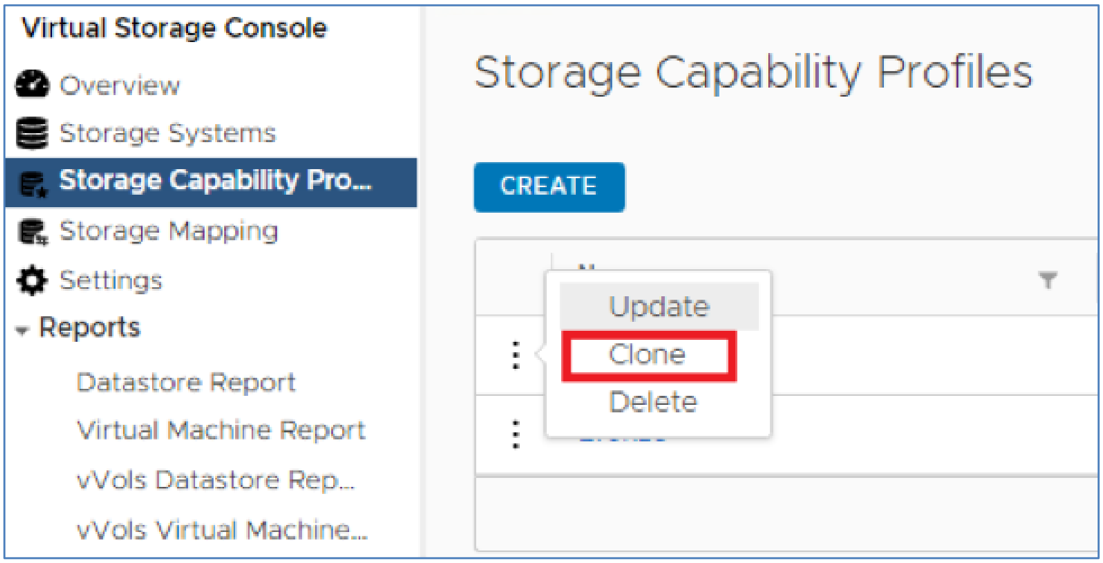

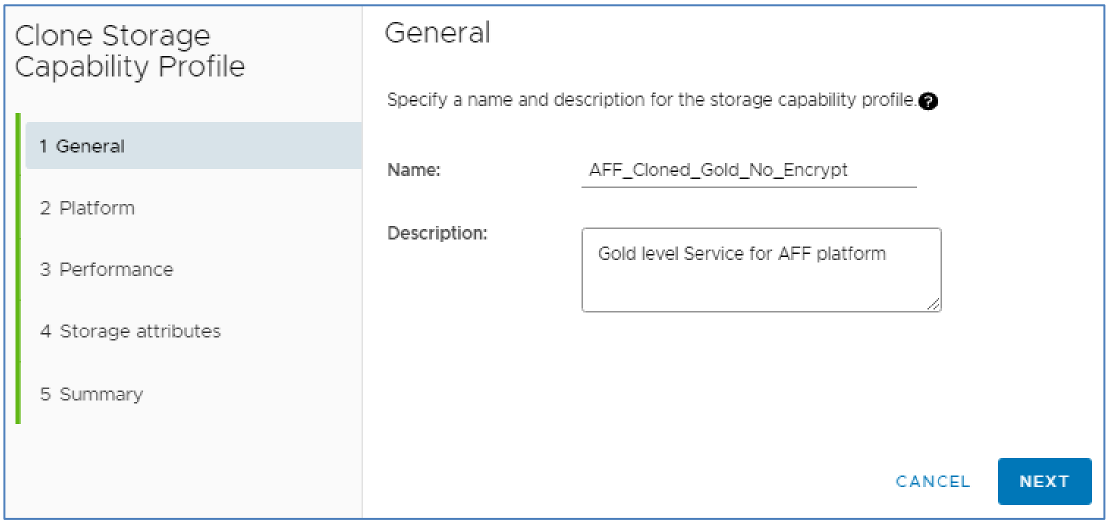

Specify a name and description for the storage capability profile.

Name: AFF_Cloned_Gold_No_Encrypt

Description: Gold level Service for AFF platform

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

4. Choose None to allow unlimited performance or set a the desired minimum and maximum IOPS for the QoS policy group. You can set the value for Max IOPS, which enables you to use the QoS functionality.

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

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

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

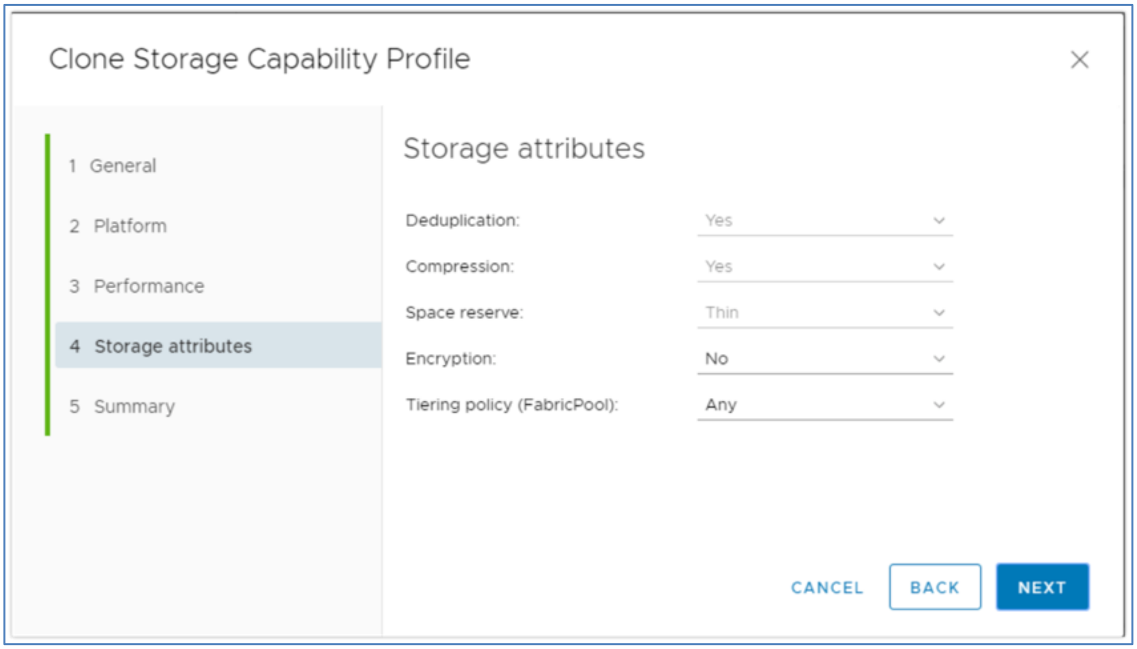

Clone Storage Capability Profile

Storage attributes

Deduplication: Yes

Compression: Yes

Space reserve: Thin

Encryption: No

Tiering policy (FabricPool): Any

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

Create a VM Storage Policy

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

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

2. Click Create VM Storage Policy.

3. Create a new name for the VM storage Policy and enter a description and click NEXT.

4. Choose Enable rules for NetApp.clustered.Data.ONTAP.VP.VASA.10 storage and NetApp.clustered.Data.ONTAP.VP.vvol storage, located under the Datastore specific rules section and click NEXT.

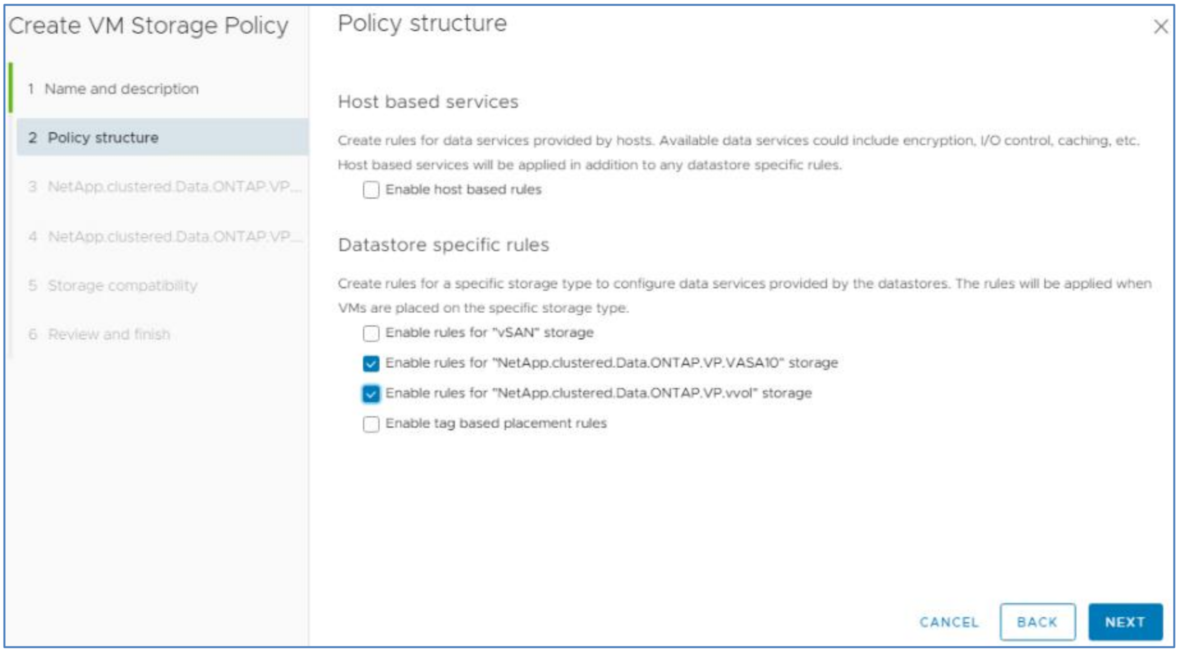

Policy structure

Host based services

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

□ Enable host based rules

Datastore specific rules

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

□ Enable rules for "vSAN" storage

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

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

□ Enable tag based placement rules

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

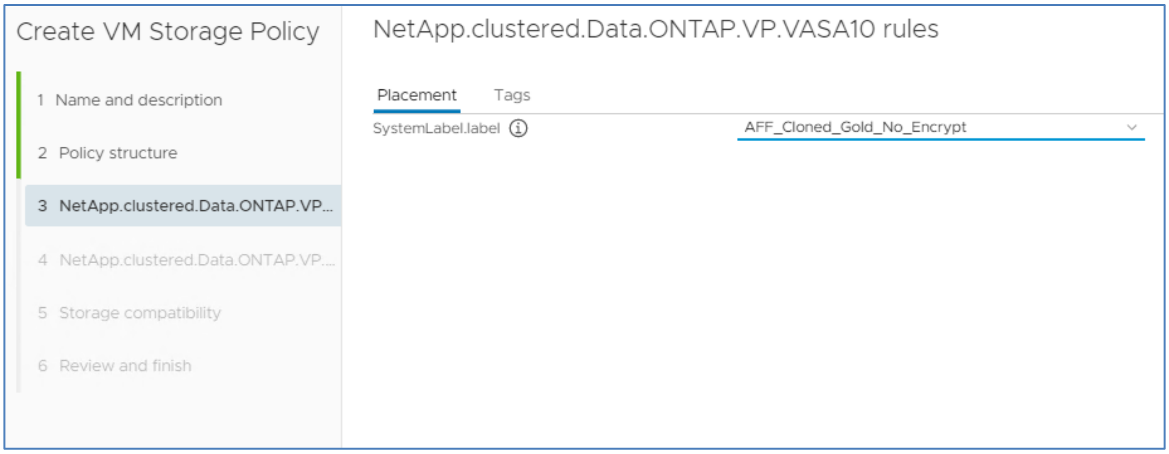

NetApp.clustered.Data.ONTAP.VP.VASA10 rules

SystemLabel.label AFF_Cloned_Gold_No_Encrypt

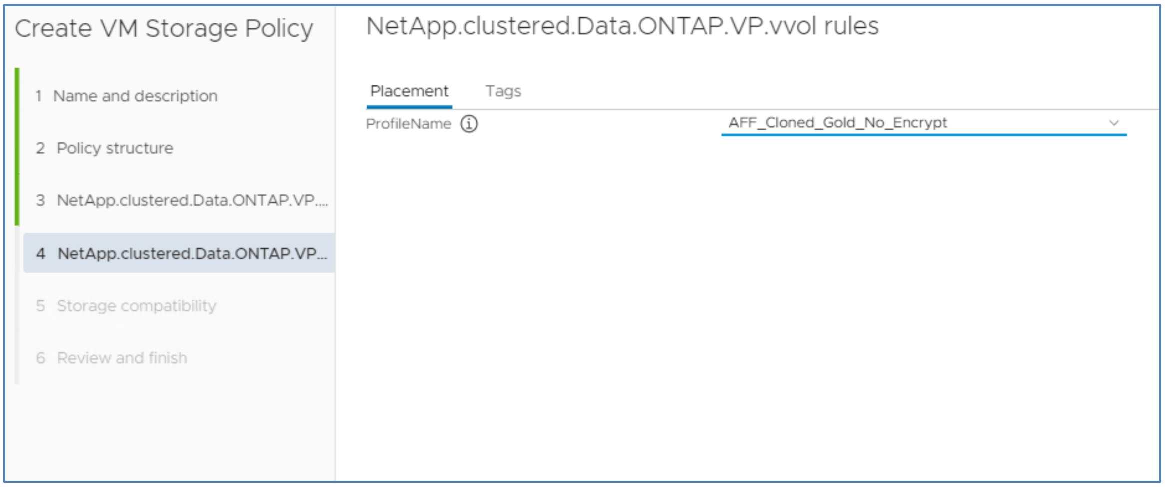

Create VM Storage Policy

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

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

Placement | Tags

ProfileName: AFF_Cloned_Gold_No_Encrypt

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

7. Review the Policy Summary and click Finish.

Provision a vVols Datastore

To provision the vVols datastore over NFS protocol, follow these steps:

1. From the Virtual Storage Console Home page, click Overview.

2. In the Getting Started tab, click Provision.

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

4. Choose the type as vVols and Enter the datastore name.

5. Select NFS for protocol and click Next.

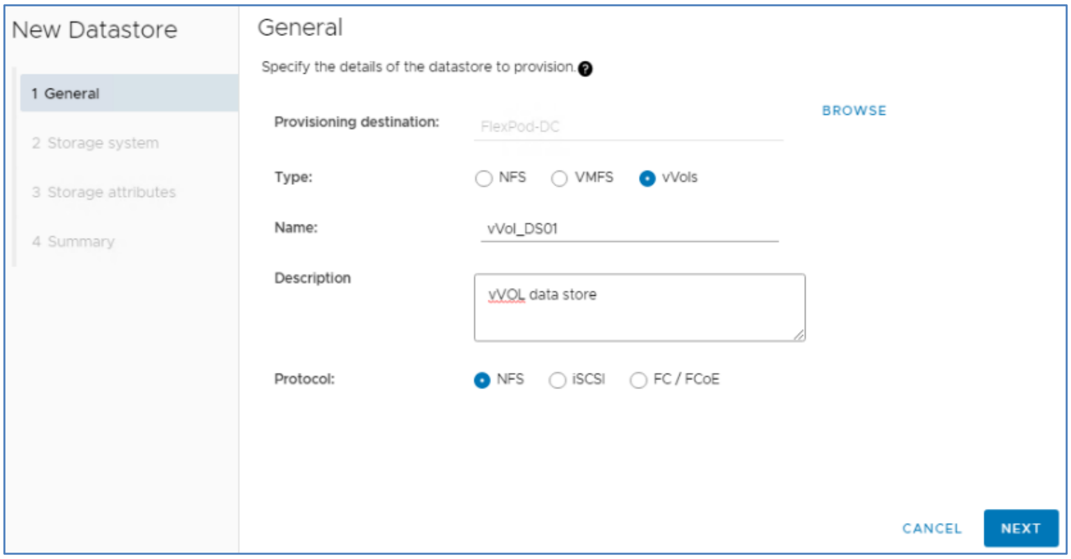

New Datastore

1 General
2 Storage system
3 Storage attributes
4 Summary

General

Specify the details of the datastore to provision.

Provisioning destination: FlexPod-DC BROWSE

Type: ○ NFS ○ VMFS ● vVols

Name: vVol_DS01

Description: vVOL data store

Protocol: ● NFS ○ iSCSI ○ FC / FCoE

CANCEL NEXT

6. Select the Storage capability profile created earlier for vVols.

7. Select the NFS storage server and the NetApp Storage SVM where the vVols needs to be created and click Next.

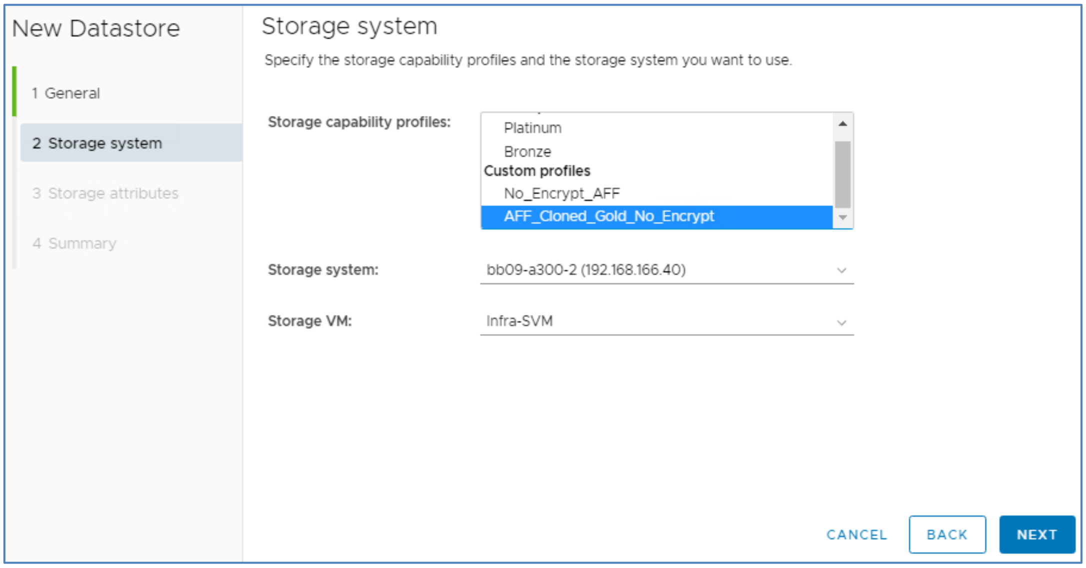

New Datastore

1 General
2 Storage system
3 Storage attributes
4 Summary

Storage system

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

Storage capability profiles: Platinum / Bronze / Custom profiles / No_Encrypt_AFF / AFF_Cloned_Gold_No_Encrypt (selected)

Storage system: bb09-a300-2 (192.168.166.40)

Storage VM: Infra-SVM

8. Create new vVols or select existing vVols.

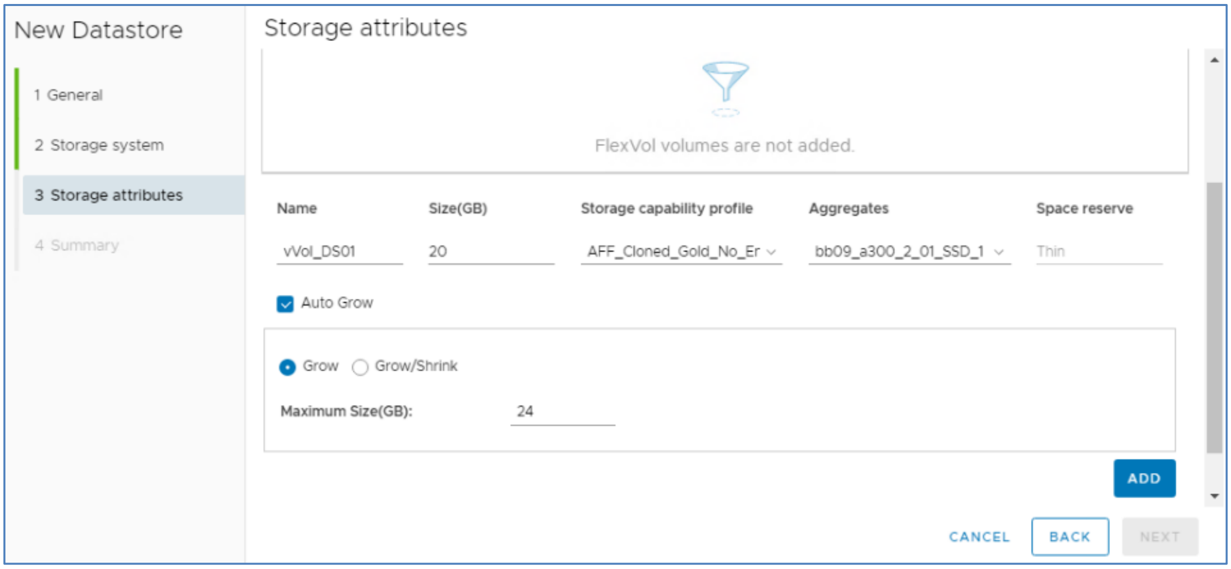

New Datastore

1 General
2 Storage system
3 Storage attributes
4 Summary

Storage attributes

FlexVol volumes are not added.

Name	Size(GB)	Storage capability profile	Aggregates	Space reserve
vVol_DS01	20	AFF_Cloned_Gold_No_Er	bb09_a300_2_01_SSD_1	Thin

☑ Auto Grow

● Grow ○ Grow/Shrink

Maximum Size(GB): 24

9. Check the storage capability profile and click ADD. The default storage profile should be added automatically.

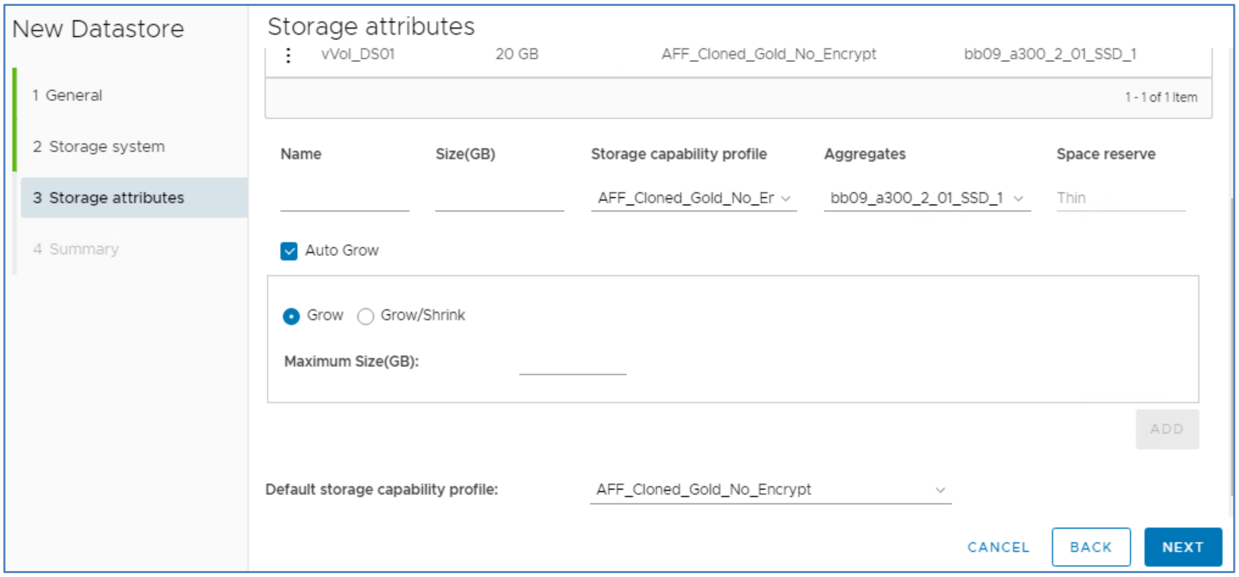

New Datastore

1 General
2 Storage system
3 Storage attributes
4 Summary

Storage attributes

vVol_DS01	20 GB	AFF_Cloned_Gold_No_Encrypt	bb09_a300_2_01_SSD_1

1 - 1 of 1 Item

Name	Size(GB)	Storage capability profile	Aggregates	Space reserve
		AFF_Cloned_Gold_No_Er	bb09_a300_2_01_SSD_1	Thin

☑ Auto Grow

● Grow ○ Grow/Shrink

Maximum Size(GB):

Default storage capability profile: AFF_Cloned_Gold_No_Encrypt

You can create multiple vVols for a datastore.

10. Check the storage capability profile and click NEXT.

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

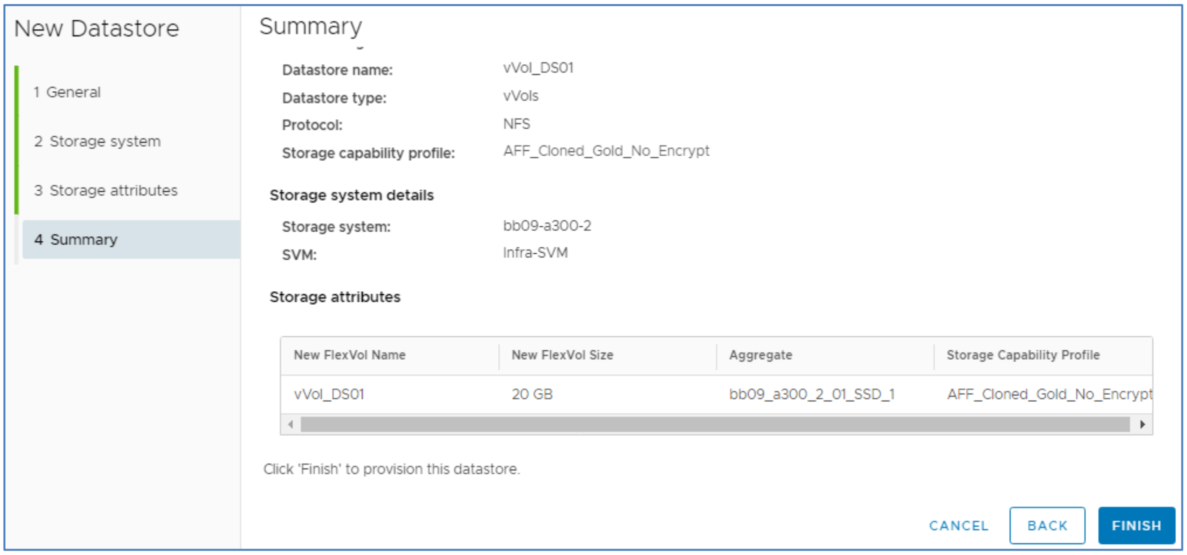

New Datastore

1 General
2 Storage system
3 Storage attributes
4 Summary

Summary

Datastore name: vVol_DS01
Datastore type: vVols
Protocol: NFS
Storage capability profile: AFF_Cloned_Gold_No_Encrypt

Storage system details

Storage system: bb09-a300-2
SVM: Infra-SVM

Storage attributes

New FlexVol Name	New FlexVol Size	Aggregate	Storage Capability Profile
vVol_DS01	20 GB	bb09_a300_2_01_SSD_1	AFF_Cloned_Gold_No_Encrypt

Click 'Finish' to provision this datastore.

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

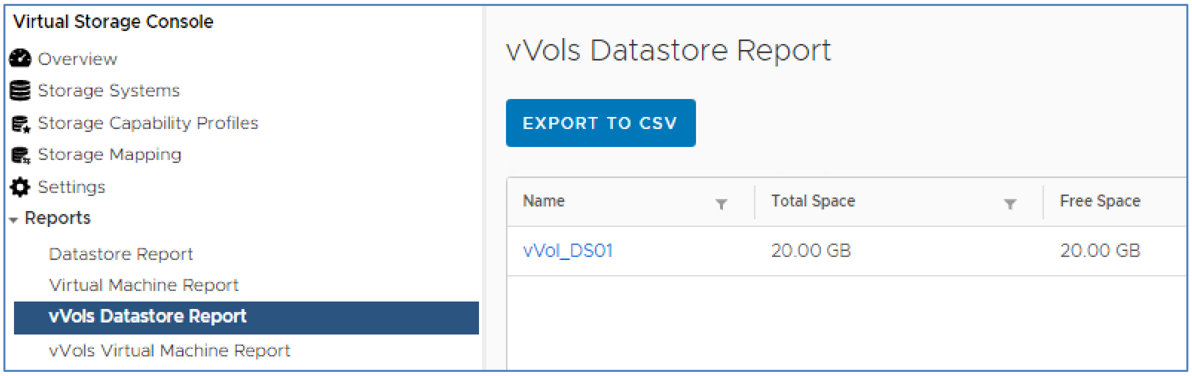

Name	Total Space	Free Space
vVol_DS01	20.00 GB	20.00 GB

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

NetApp SnapCenter 4.3.1

SnapCenter Software is a simple, centralized, scalable platform that provides application-consistent data protection for applications, databases, host file systems, and VMs running on ONTAP systems anywhere in the Hybrid Cloud.

NetApp SnapCenter Architecture

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

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

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

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

Install SnapCenter Plug-In for VMware vSphere 4.3.1

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

Host and Privilege Requirements for the SnapCenter Plug-In for VMware vSphere

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

- You must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance as a Linux VM.
- You should deploy the virtual appliance on the vCenter Server.
- You must not deploy the virtual appliance in a folder that has a name with special characters.
- You must deploy and register a separate, unique instance of the virtual appliance for each vCenter Server.

Table 24. Port Requirements

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

License Requirements for SnapCenter Plug-In for VMware vSphere

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

Table 25. SnapCenter Plug-in for VMware vSphere License Requirements

Product	License Requirement
ONTAP	SnapManager Suite: Used for backup operations One of these: SnapMirror or SnapVault (for secondary data protection regardless of the type of relationship)
ONTAP Primary Destinations	To perform protection of VMware VMs and datastores the following licenses should be installed: SnapRestore: used for restore operations FlexClone: used for mount and attach operations
ONTAP Secondary Destinations	To perform protection of VMware VMs and datastores only: FlexClone: used for mount and attach operations

Product	License Requirement
VMware	vSphere Standard, Enterprise, or Enterprise Plus A vSphere license is required to perform restore operations, which use Storage vMotion. vSphere Essentials or Essentials Plus licenses do not include Storage vMotion.

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

Download and Deploy the SnapCenter Plug-In for VMware vSphere 4.3.1

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

1. Download SnapCenter Plug-in for VMware vSphere OVA file from NetApp support site (https://mysupport.netapp.com).

2. From VMware vCenter, navigate to the VMs and Templates tab, right-click FlexPod-DC and choose Deploy OVF Template.

3. Specify the location of the OVF Template and click NEXT.

4. On the Select a name and folder page, enter a unique name and location for the VM and click NEXT to continue.

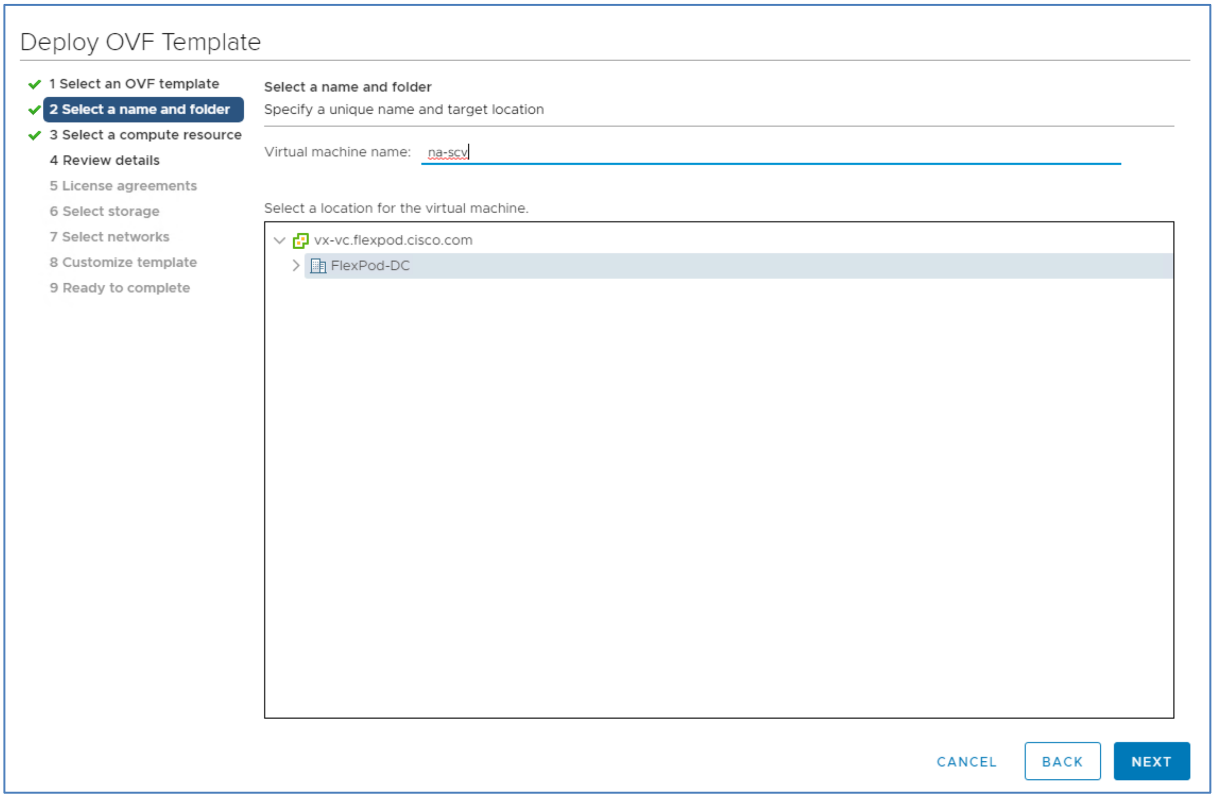

Deploy OVF Template

Select a name and folder

Specify a unique name and target location

Virtual machine name: na-scv

Select a location for the virtual machine.

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

5. On the Select a compute resource page, choose a resource where you want to run the deployed VM template, and click NEXT.

6. On the Review details page, verify the OVA template details and click NEXT.

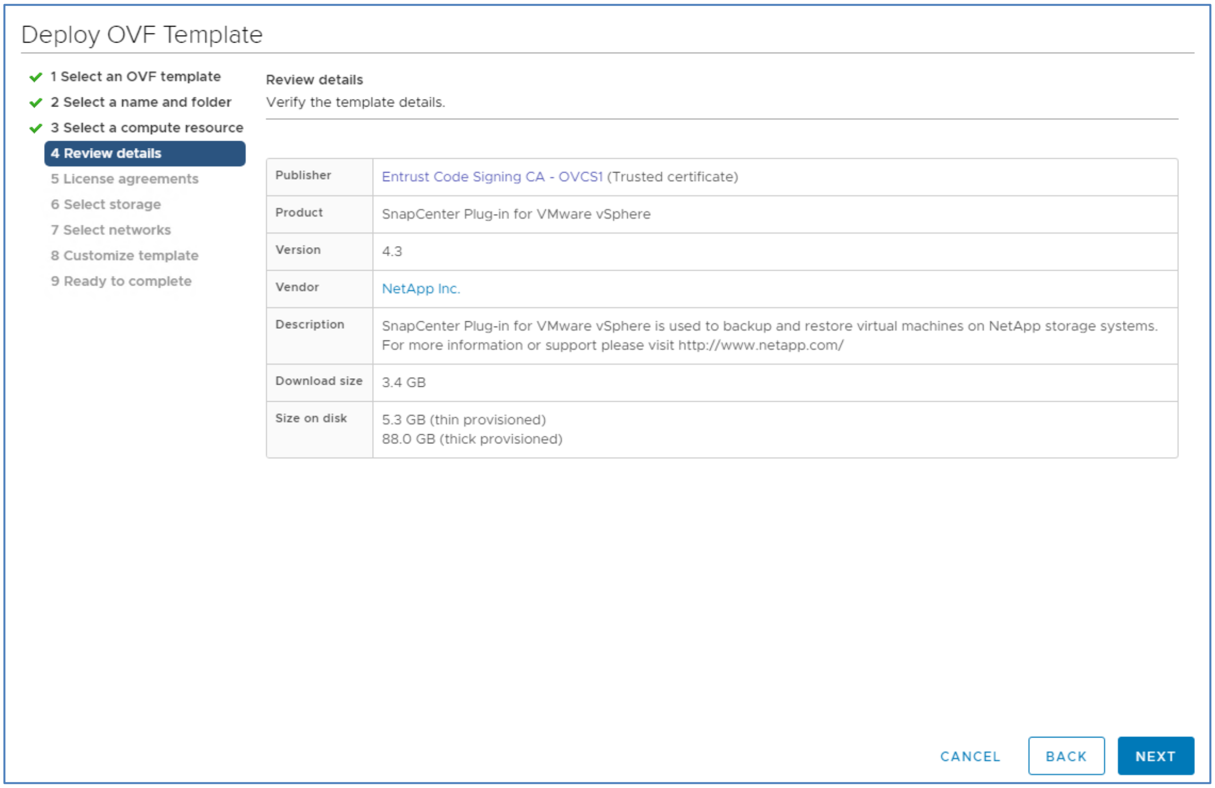

Review details

Verify the template details.

Publisher	Entrust Code Signing CA - OVCS1 (Trusted certificate)
Product	SnapCenter Plug-in for VMware vSphere
Version	4.3
Vendor	NetApp Inc.
Description	SnapCenter Plug-in for VMware vSphere is used to backup and restore virtual machines on NetApp storage systems. For more information or support please visit http://www.netapp.com/
Download size	3.4 GB
Size on disk	5.3 GB (thin provisioned) 88.0 GB (thick provisioned)

7. On the License agreements page, check the box I accept all license agreements.

8. On the Select storage page, change the datastore virtual disk format to Thin Provision and click NEXT.

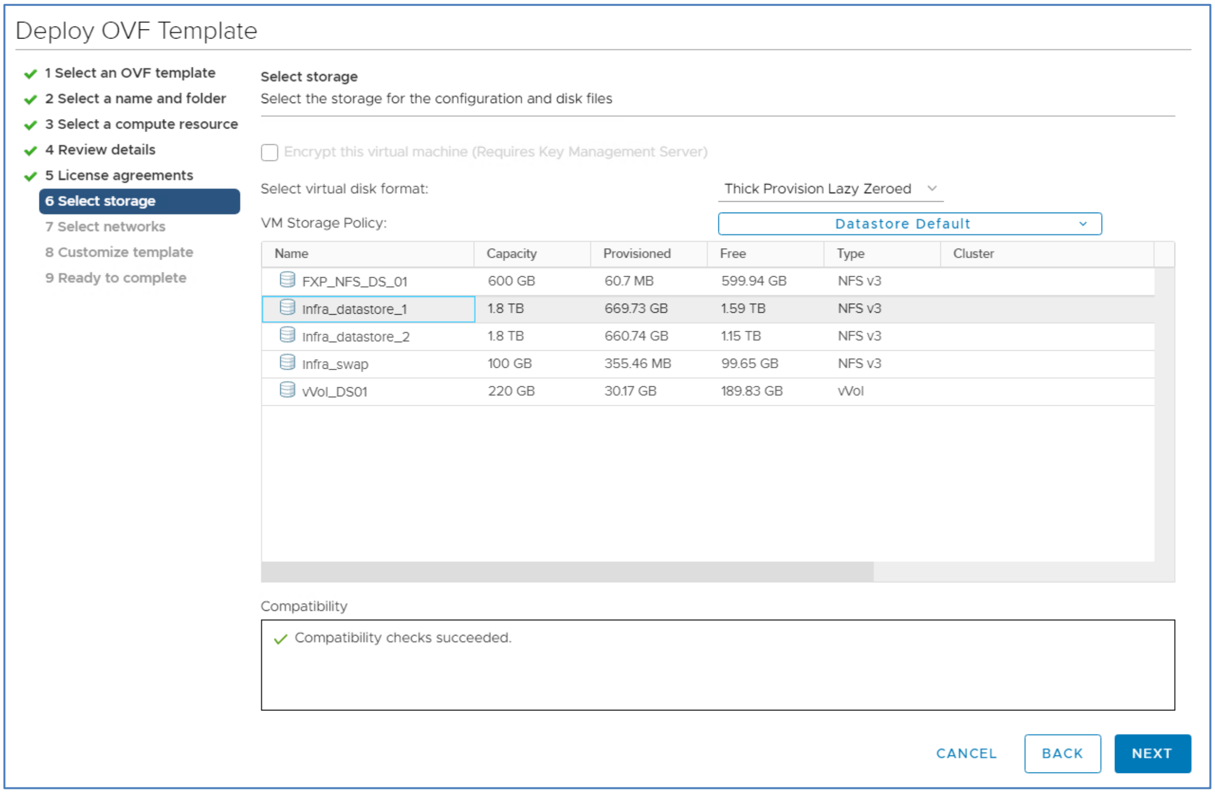

Deploy OVF Template

Select storage

Select the storage for the configuration and disk files

□ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
FXP_NFS_DS_01	600 GB	60.7 MB	599.94 GB	NFS v3	
Infra_datastore_1	1.8 TB	669.73 GB	1.59 TB	NFS v3	
Infra_datastore_2	1.8 TB	660.74 GB	1.15 TB	NFS v3	
Infra_swap	100 GB	355.46 MB	99.65 GB	NFS v3	
vVol_DS01	220 GB	30.17 GB	189.83 GB	vVol	

Compatibility

✓ Compatibility checks succeeded.

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

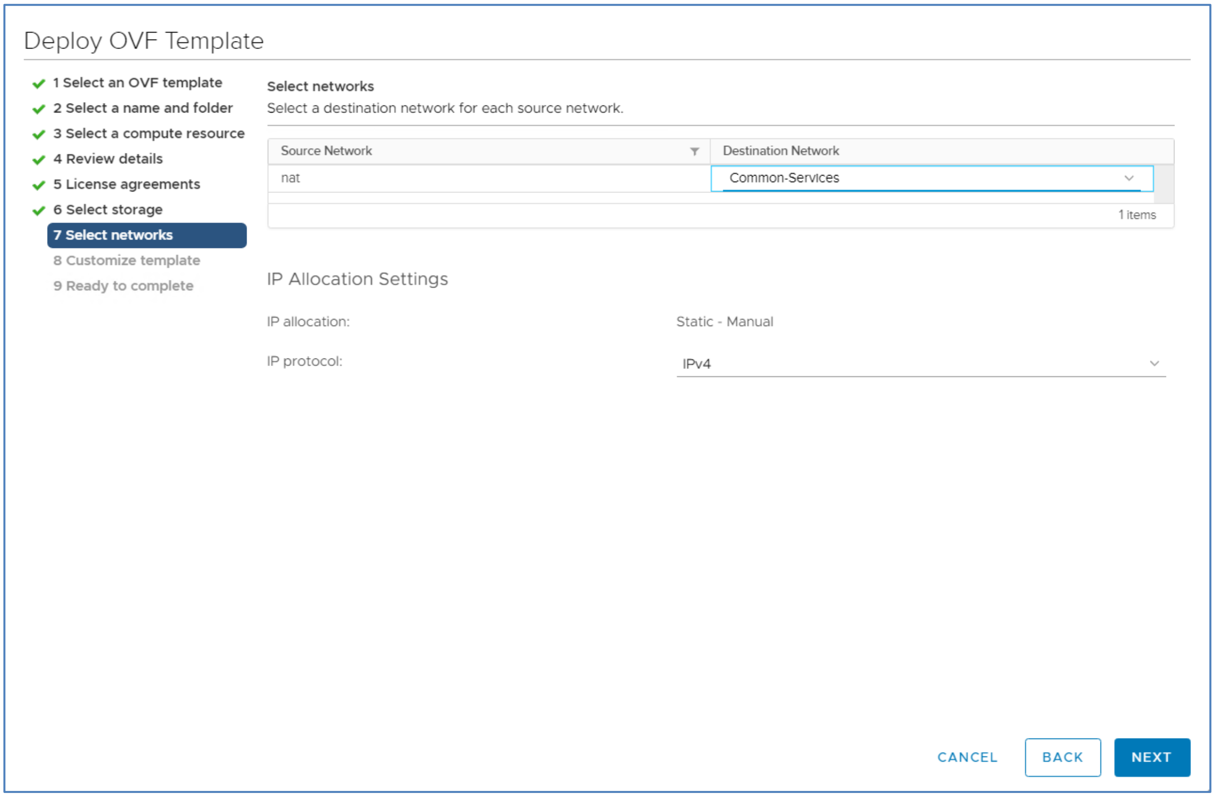

Deploy OVF Template

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	Common-Services

1 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

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

    a. In Register to existing vCenter, enter the vCenter credentials.

    b. In Create SnapCenter Plug-in for VMware vSphere credentials, enter the SnapCenter Plug-in for VMware vSphere credentials.

    c. In Create SCV credentials, create a username and password for the SCV maintenance user.

    d. In Setup Network Properties, enter the network information.

    e. In Setup Date and Time, choose the time zone where the vCenter is located.

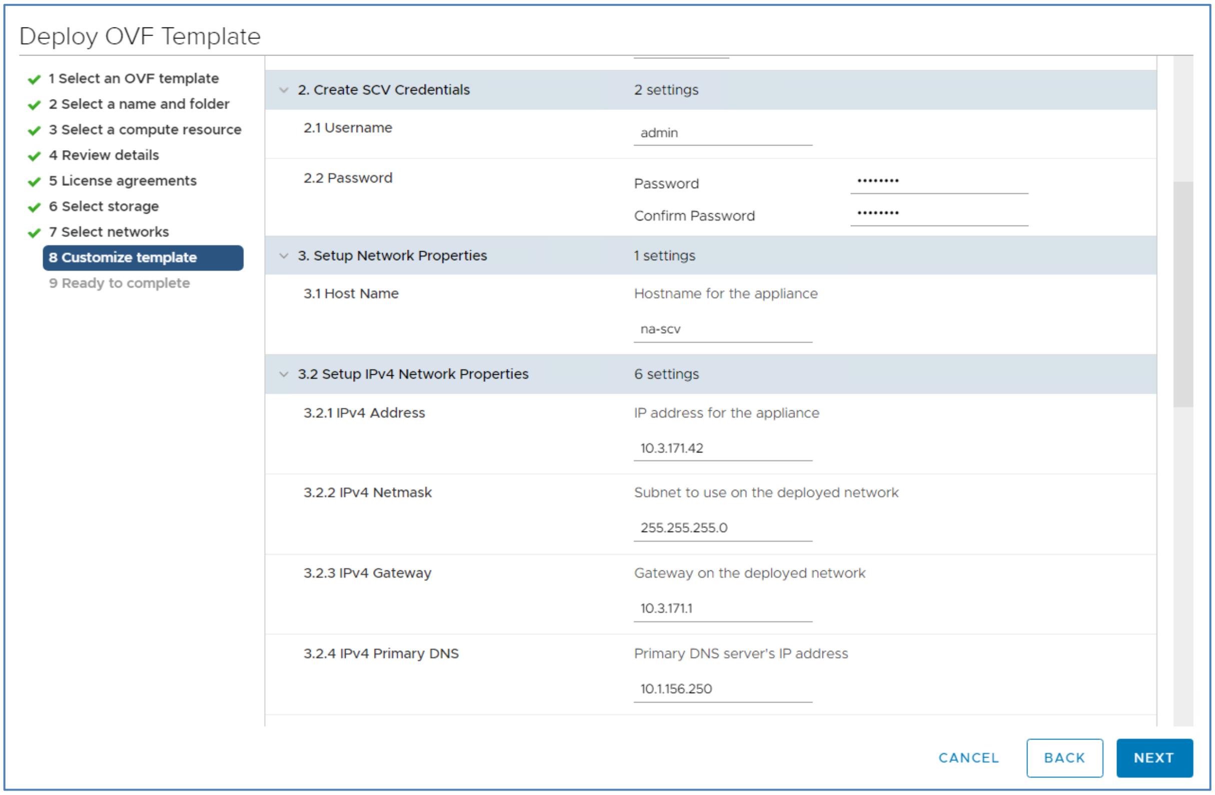

Deploy OVF Template

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

2. Create SCV Credentials — 2 settings

2.1 Username: admin

2.2 Password: Password •••••••• Confirm Password ••••••••

3. Setup Network Properties — 1 settings

3.1 Host Name: Hostname for the appliance: na-scv

3.2 Setup IPv4 Network Properties — 6 settings

3.2.1 IPv4 Address: IP address for the appliance: 10.3.171.42

3.2.2 IPv4 Netmask: Subnet to use on the deployed network: 255.255.255.0

3.2.3 IPv4 Gateway: Gateway on the deployed network: 10.3.171.1

3.2.4 IPv4 Primary DNS: Primary DNS server's IP address: 10.1.156.250

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

12. Navigate to the VM where the virtual appliance was deployed, then click the Summary tab, and then click the Power On box to start the virtual appliance.

13. While the virtual appliance is powering on, click Install VMware tools in the Yellow banner displayed in the summary tab of the appliance.

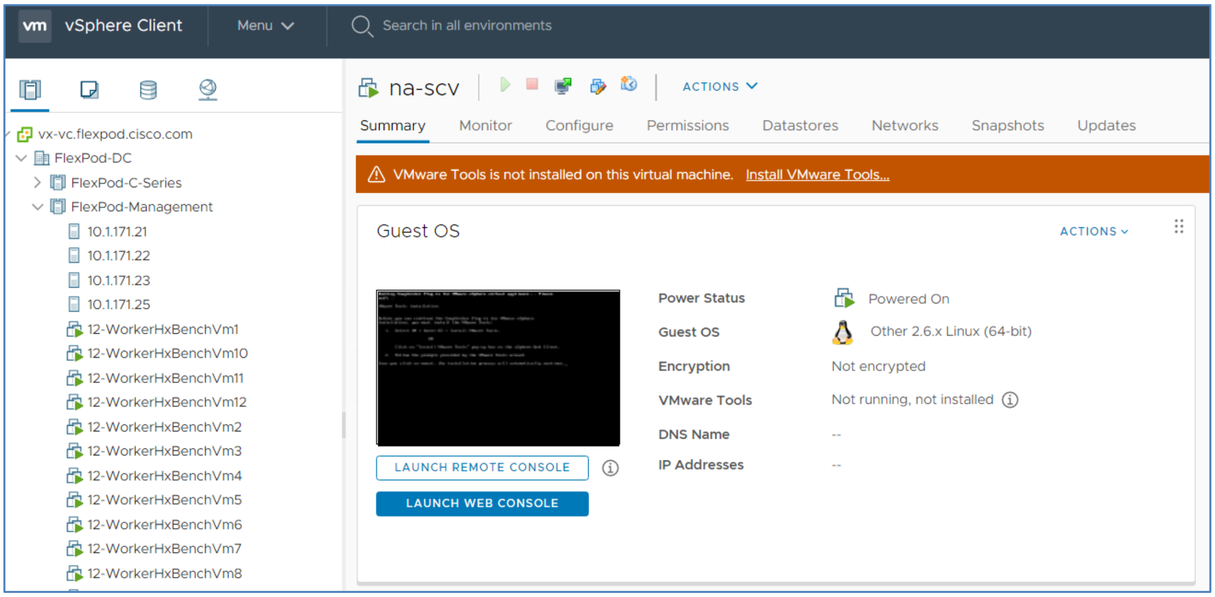

- vx-vc.flexpod.cisco.com
  - FlexPod-DC
    - FlexPod-C-Series
    - FlexPod-Management
      - 10.1.171.21
      - 10.1.171.22
      - 10.1.171.23
      - 10.1.171.25
      - 12-WorkerHxBenchVm1
      - 12-WorkerHxBenchVm10
      - 12-WorkerHxBenchVm11
      - 12-WorkerHxBenchVm12
      - 12-WorkerHxBenchVm2
      - 12-WorkerHxBenchVm3
      - 12-WorkerHxBenchVm4
      - 12-WorkerHxBenchVm5
      - 12-WorkerHxBenchVm6
      - 12-WorkerHxBenchVm7
      - 12-WorkerHxBenchVm8

na-scv

⚠ VMware Tools is not installed on this virtual machine. Install VMware Tools...

Guest OS

Power Status	Powered On
Guest OS	Other 2.6.x Linux (64-bit)
Encryption	Not encrypted
VMware Tools	Not running, not installed
DNS Name	--
IP Addresses	--

14. Log into SnapCenter Plug-in for VMware vSphere using the IP address displayed on the appliance console screen with the credentials that you provided in the deployment wizard. Verify on the Dashboard that the virtual appliance is successfully connected to vCenter and the SnapCenter Plug-in for VMware vSphere is successfully enabled and connected.

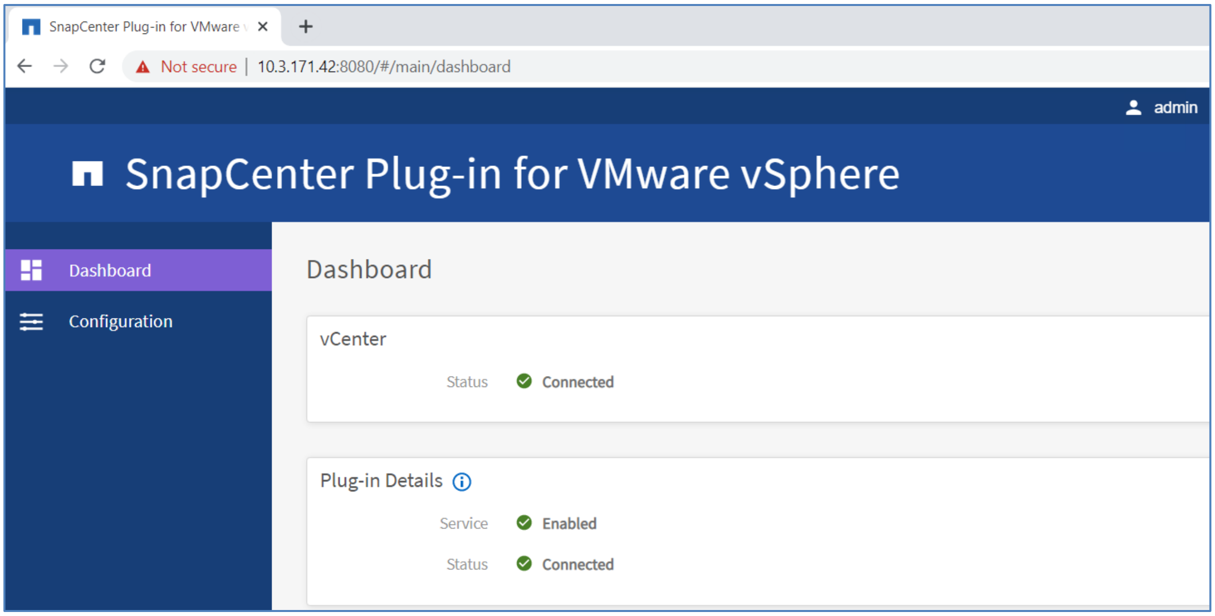
SnapCenter Plug-in for VMware
Not secure
10.3.171.42:8080/#/main/dashboard
admin
SnapCenter Plug-in for VMware vSphere
Dashboard
Configuration

Dashboard

vCenter

Status ✓ Connected

Plug-in Details

Service ✓ Enabled

Status ✓ Connected

SnapCenter Plug-In for VMware vSphere in vCenter Server

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

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

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

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

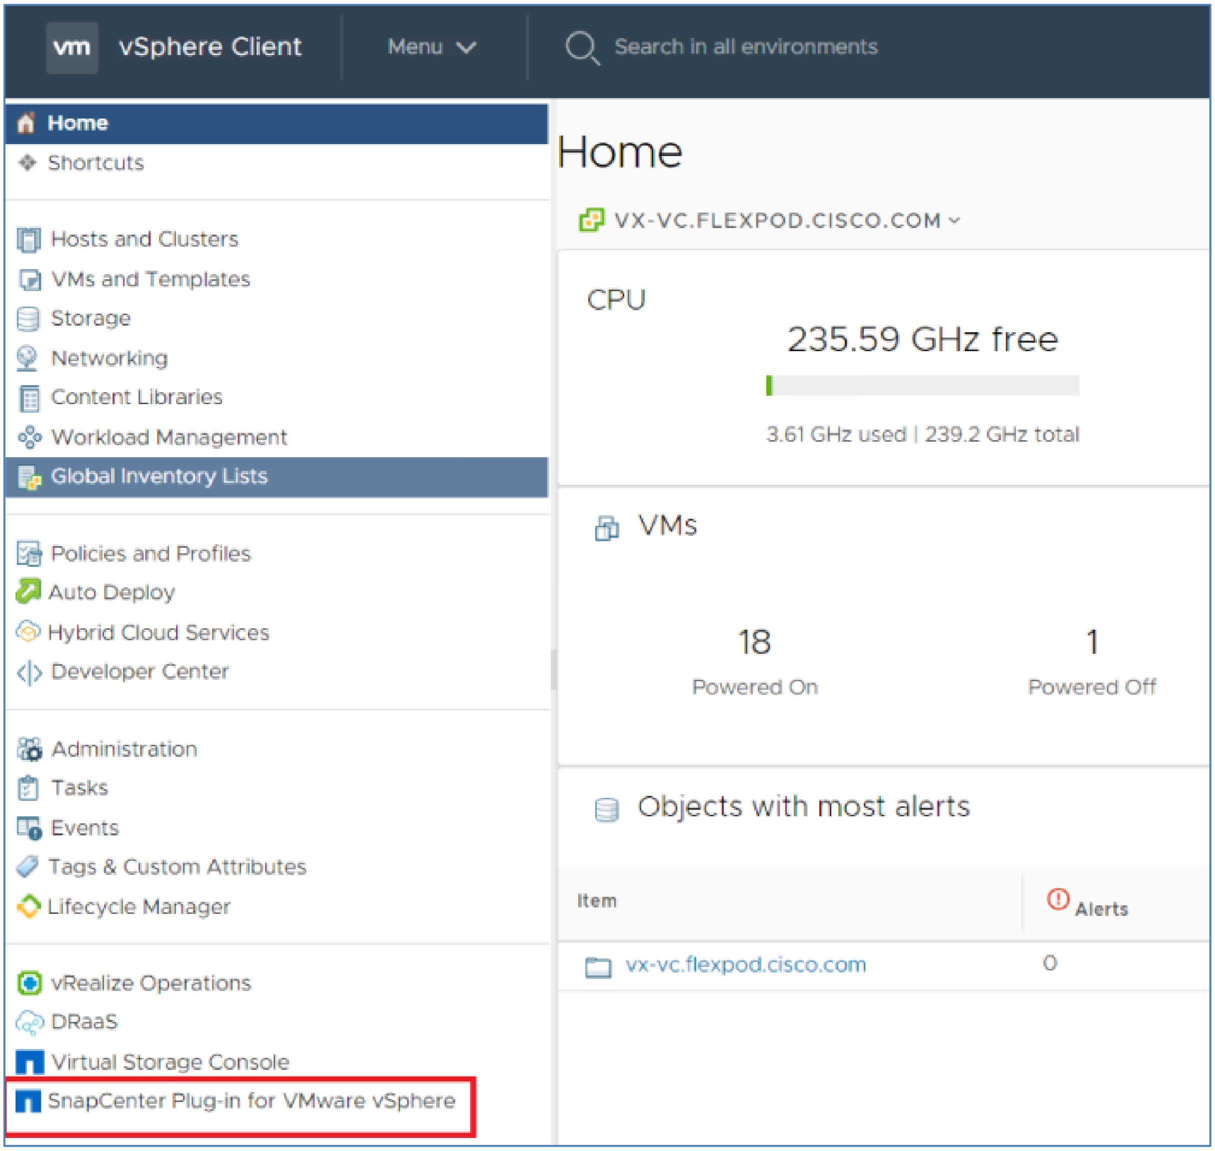

Add Storage System

To add storage systems, follow these steps:

1. Go to the Storage Systems tab.
2. Click Add Storage System to add a cluster or SVM.
3. Enter vCenter, Storage System, user credentials, and other required information.

Check the box for Log SnapCenter server events to syslog and Send AutoSupport Notification for failed operation to storage system.

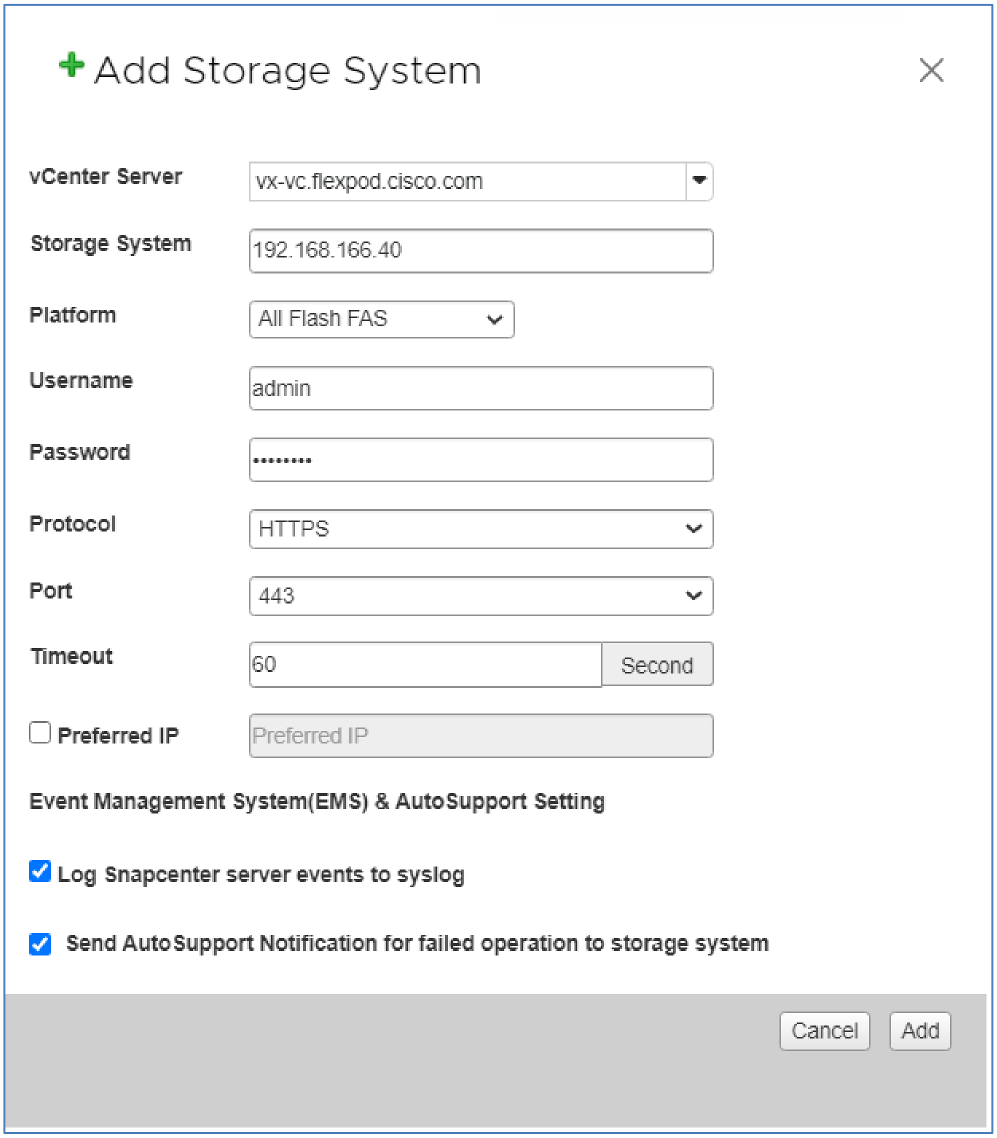

Add Storage System

vCenter Server	vx-vc.flexpod.cisco.com
Storage System	192.168.166.40
Platform	All Flash FAS
Username	admin
Password	••••••••
Protocol	HTTPS
Port	443
Timeout	60 Second
□ Preferred IP	Preferred IP

Event Management System(EMS) & AutoSupport Setting

☑ Log Snapcenter server events to syslog

☑ Send AutoSupport Notification for failed operation to storage system

Create Backup Policies for Virtual Machines and Datastores

To create backup policies for VMs and datastores, follow these steps:

1. In the left Navigator pane of the VMware vSphere Web Client, click Policies.

2. On the Policies page, click New Policy in the toolbar.

3. On the New Backup Policy page, follow these steps:

    a. Enter the policy name and a description.

    b. Enter the backups to keep.

    c. From the Frequency drop-down list, choose the backup frequency (hourly, daily, weekly, monthly, and on-demand only).

d. Expand the Advanced options and select VM Consistency and Include datastore with independent disks. Click Add.

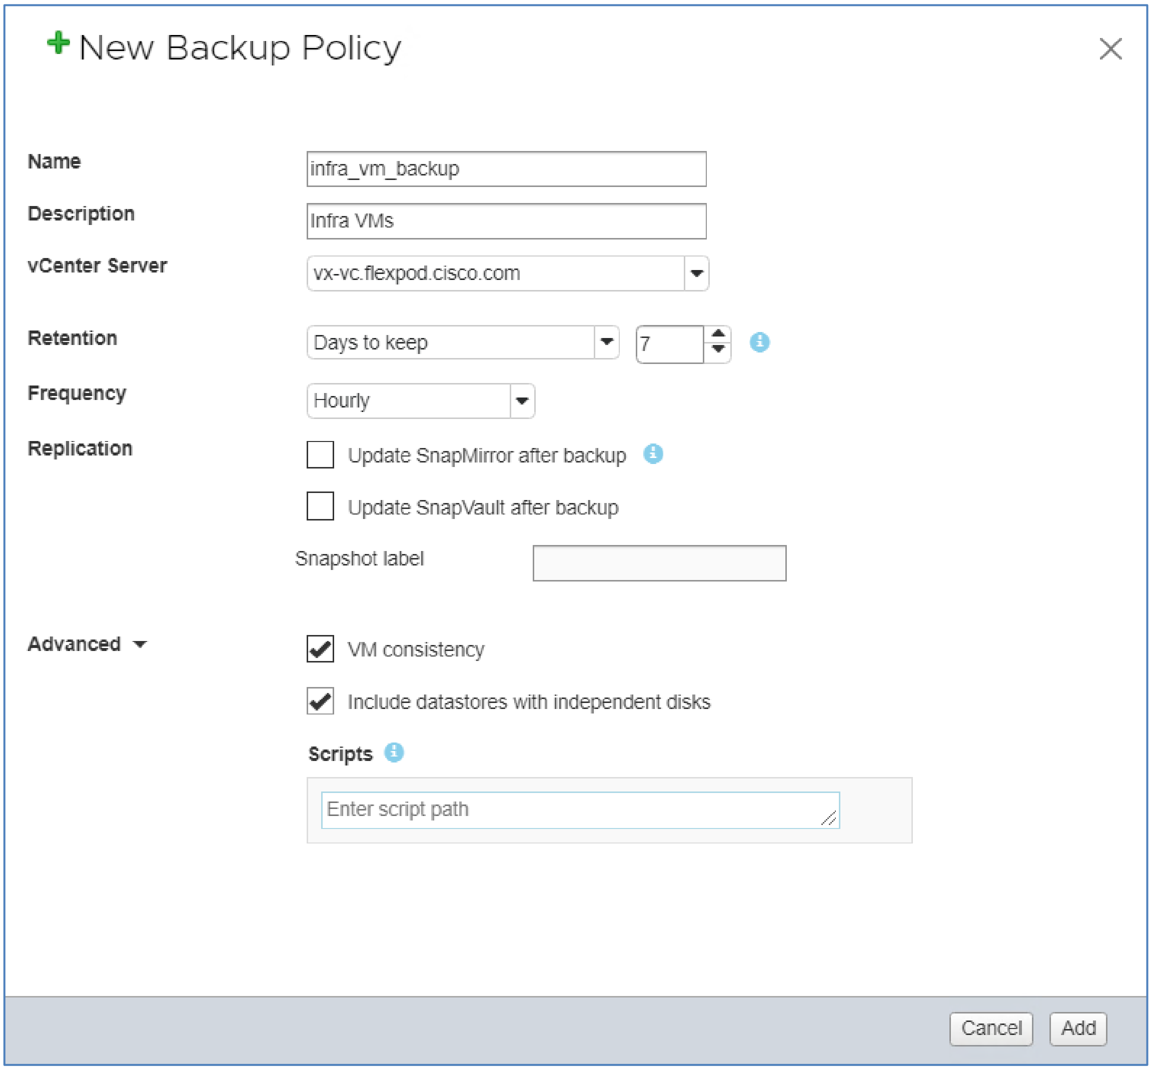

New Backup Policy

Name: infra_vm_backup

Description: Infra VMs

vCenter Server: vx-vc.flexpod.cisco.com

Retention: Days to keep 7

Frequency: Hourly

Replication:
- □ Update SnapMirror after backup
- □ Update SnapVault after backup

Snapshot label:

Advanced
- ☑ VM consistency
- ☑ Include datastores with independent disks

Scripts: Enter script path

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

Create Resource Groups

Resource groups are groups of virtual machines or datastores that are backed up together. A backup policy is associated with the resource group to back up the virtual machines and retain a certain number of backups as defined in the policy.

To create resource groups, follow these steps:

1. In the left Navigator pane of the SnapCenter Plug-in for VMware vSphere, click Resource Groups and then click Create Resource Group. This is the easiest way to create a resource group. However, you can also create a resource group with one resource by performing one of the following steps:

2. To create a resource group for one virtual machine, click VMs and Templates, right-click a virtual machine, choose NetApp SnapCenter from the drop-down list, and then choose Create Resource Group from the secondary drop-down list.

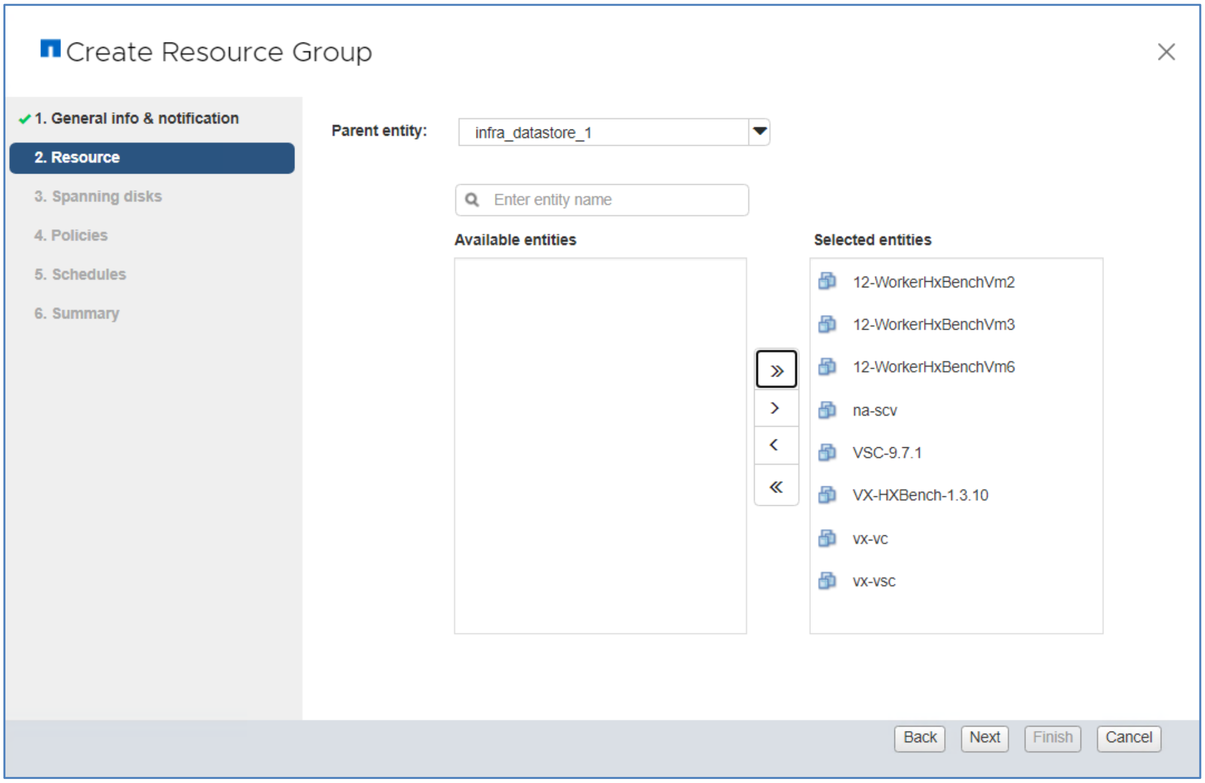

Create Resource Group

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

Parent entity: infra_datastore_1

Enter entity name

Available entities

Selected entities

- 12-WorkerHxBenchVm2
- 12-WorkerHxBenchVm3
- 12-WorkerHxBenchVm6
- na-scv
- VSC-9.7.1
- VX-HXBench-1.3.10
- vx-vc
- vx-vsc

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

3. In the General Info & Notification page, enter the resource group name and complete the notification settings. Click Next.

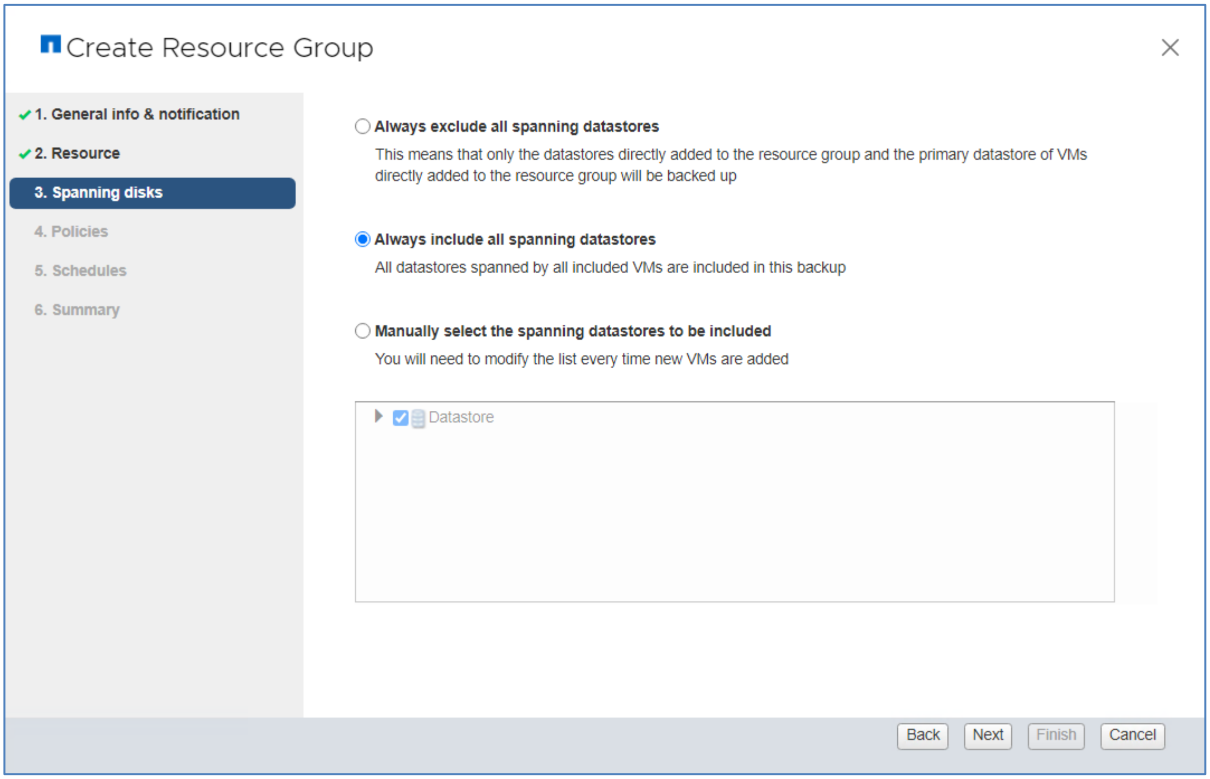

Create Resource Group

○ **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

☑ Datastore

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

4. Choose a datastore as the parent entity to create a resource group of virtual machines, and then choose the virtual machines from the available list. Click Next.

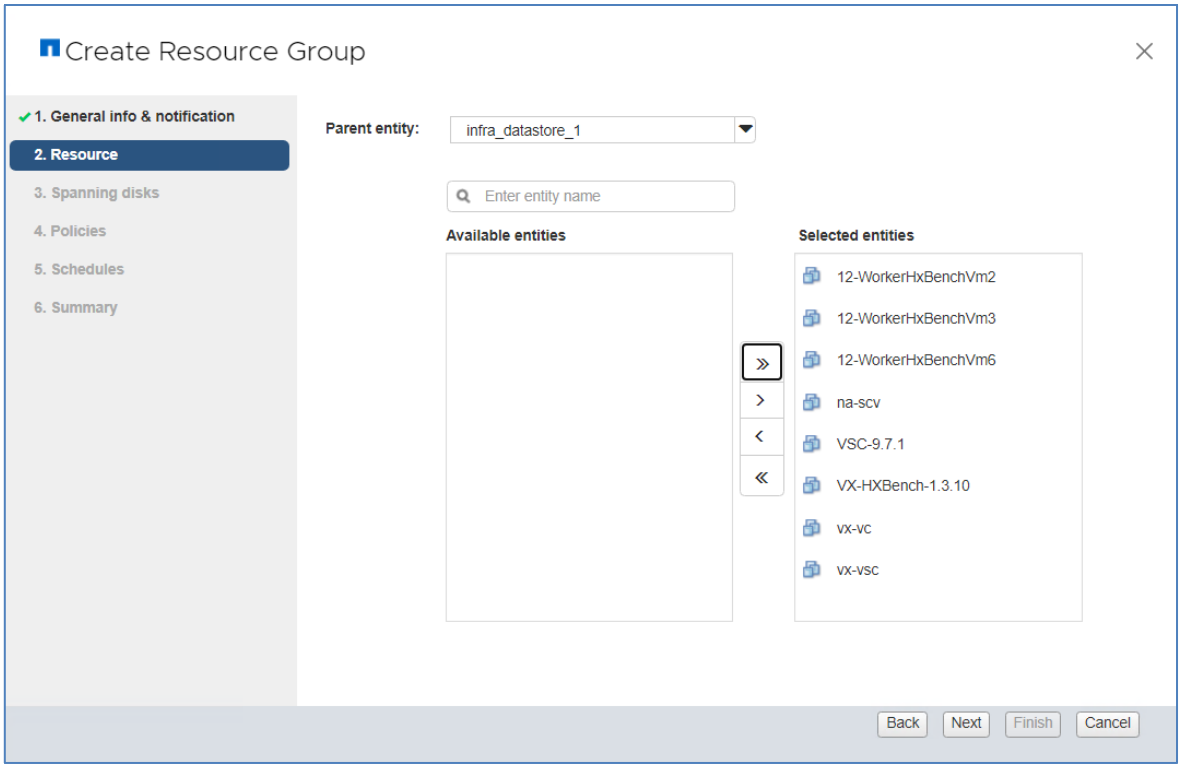

Create Resource Group

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

Parent entity: infra_datastore_1

Enter entity name

Available entities

Selected entities
- 12-WorkerHxBenchVm2
- 12-WorkerHxBenchVm3
- 12-WorkerHxBenchVm6
- na-scv
- VSC-9.7.1
- VX-HXBench-1.3.10
- vx-vc
- vx-vsc

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

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

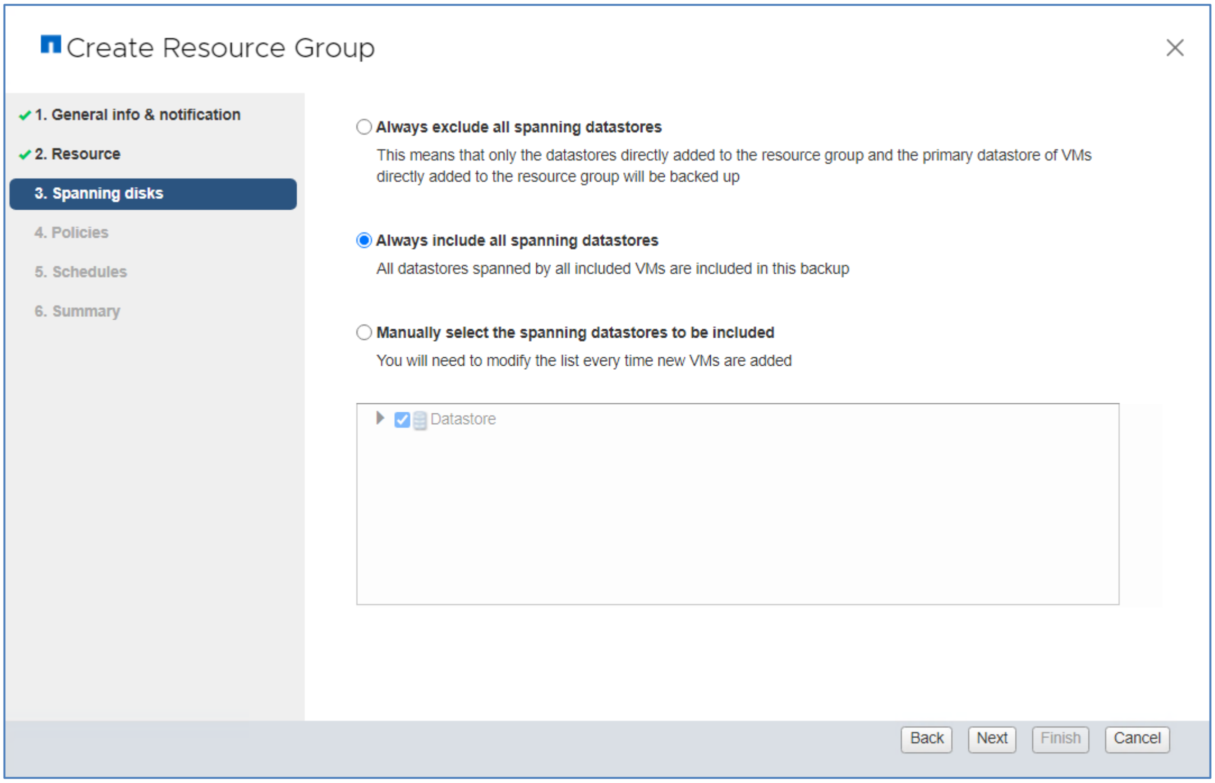

Create Resource Group

○ **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

☑ Datastore

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

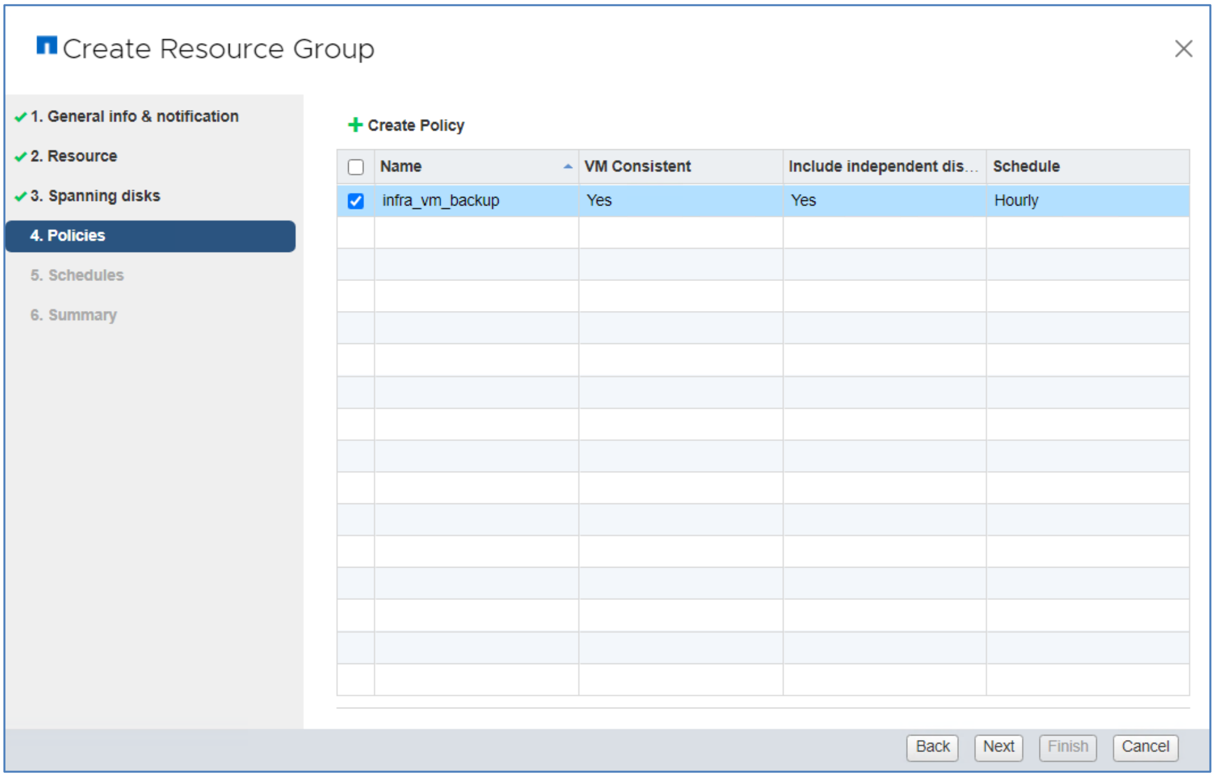

Create Resource Group

□	Name	VM Consistent	Include independent dis...	Schedule
☑	infra_vm_backup	Yes	Yes	Hourly

7. From the Schedules option, choose the schedule for each selected policy and click Next.

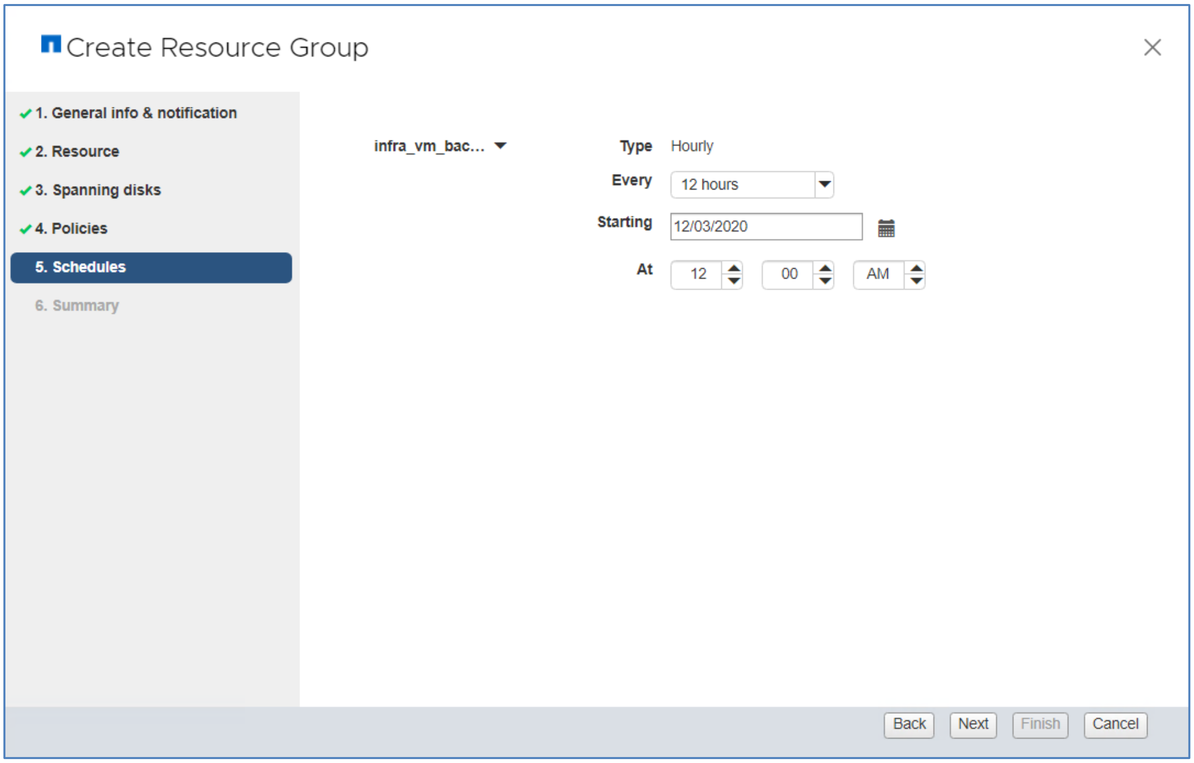

Type Hourly

Every 12 hours

Starting 12/03/2020

At 12 00 AM

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

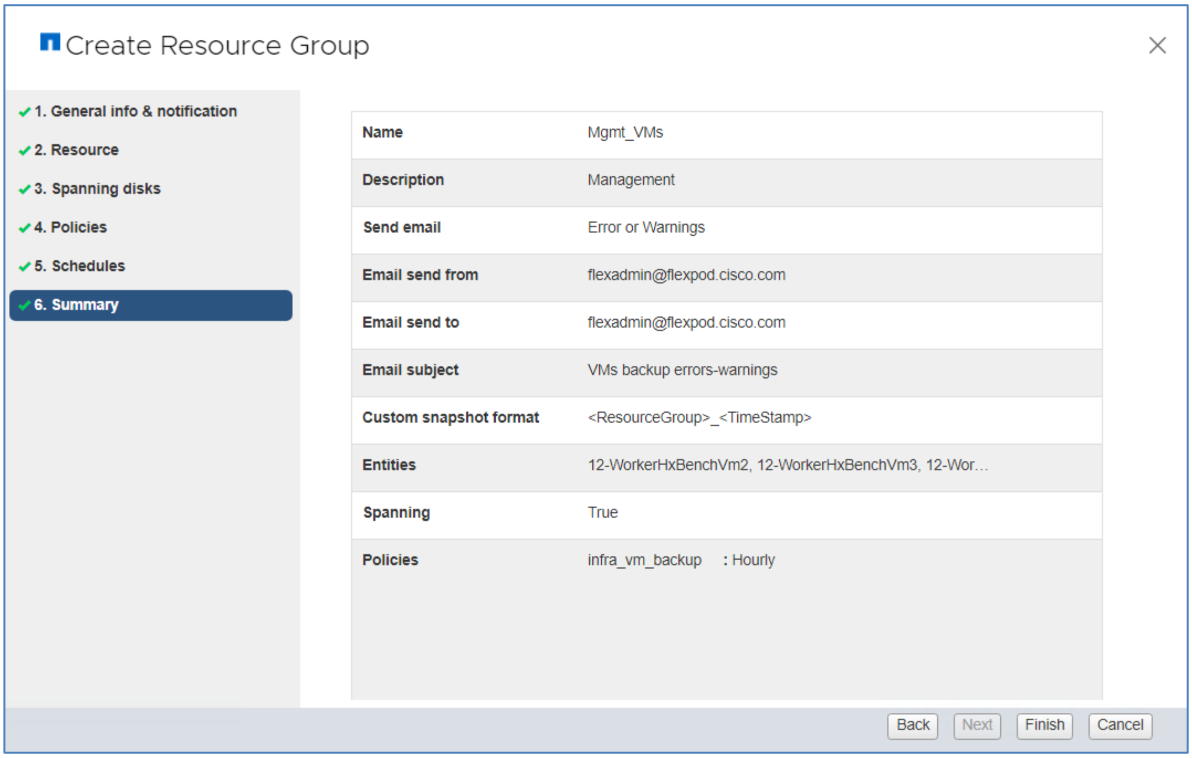

Name	Mgmt_VMs
Description	Management
Send email	Error or Warnings
Email send from	flexadmin@flexpod.cisco.com
Email send to	flexadmin@flexpod.cisco.com
Email subject	VMs backup errors-warnings
Custom snapshot format	<ResourceGroup>_<TimeStamp>
Entities	12-WorkerHxBenchVm2, 12-WorkerHxBenchVm3, 12-Wor...
Spanning	True
Policies	infra_vm_backup : Hourly

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

Backups of the virtual machines included in the resource group occurs according to the schedule of the policies associated with the resource group. To view the backups associated with each schedule, follow these steps:

1. Navigate to the VMs and Templates tab.

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

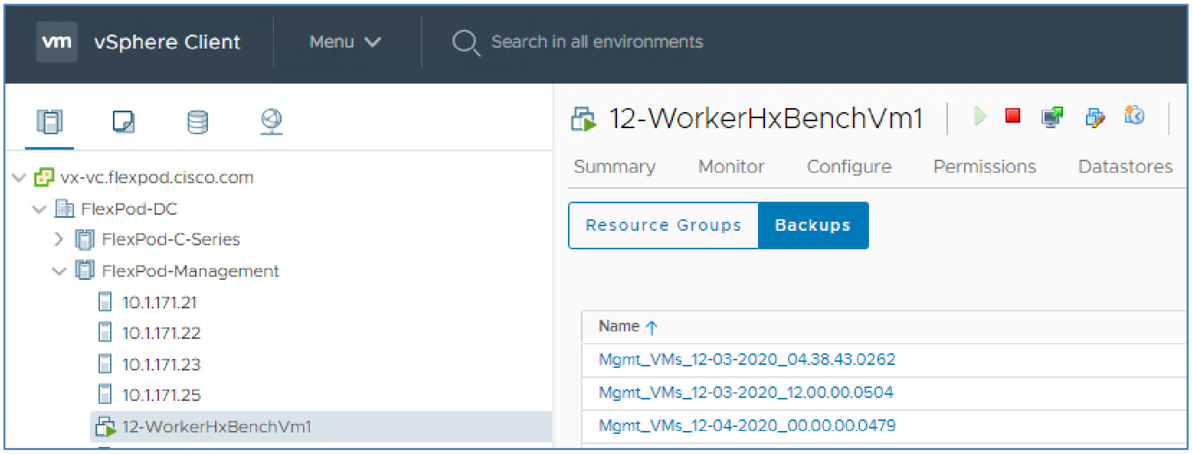

Name ↑
Mgmt_VMs_12-03-2020_04.38.43.0262
Mgmt_VMs_12-03-2020_12.00.00.0504
Mgmt_VMs_12-04-2020_00.00.00.0479

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

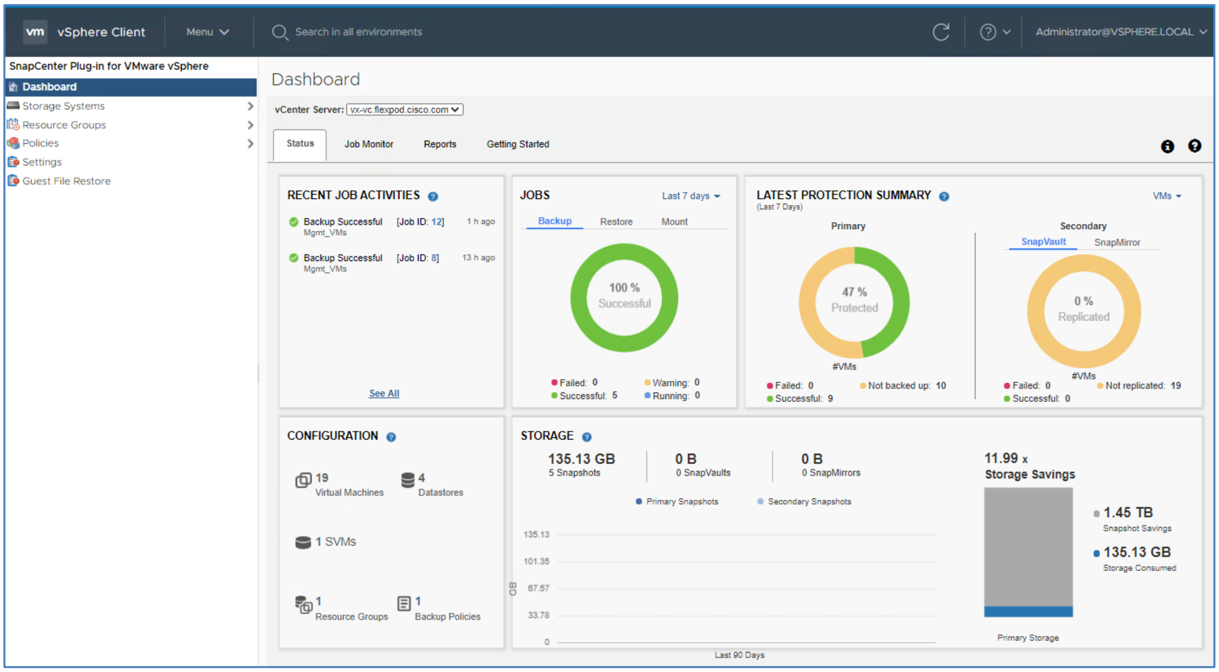

RECENT JOB ACTIVITIES

Backup Successful [Job ID: 12] 1 h ago
Mgmt_VMs

Backup Successful [Job ID: 8] 13 h ago
Mgmt_VMs

See All

JOBS — Last 7 days — Backup | Restore | Mount

100 % Successful

Failed: 0 · Warning: 0 · Successful: 5 · Running: 0

LATEST PROTECTION SUMMARY (Last 7 Days) — VMs

Primary: 47 % Protected, #VMs — Failed: 0 · Not backed up: 10 · Successful: 9

Secondary (SnapVault | SnapMirror): 0 % Replicated, #VMs — Failed: 0 · Not replicated: 19 · Successful: 0

CONFIGURATION

19 Virtual Machines · 4 Datastores · 1 SVMs · 1 Resource Groups · 1 Backup Policies

STORAGE

135.13 GB 5 Snapshots · 0 B 0 SnapVaults · 0 B 0 SnapMirrors

11.99 x Storage Savings

1.45 TB Snapshot Savings

135.13 GB Storage Consumed

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

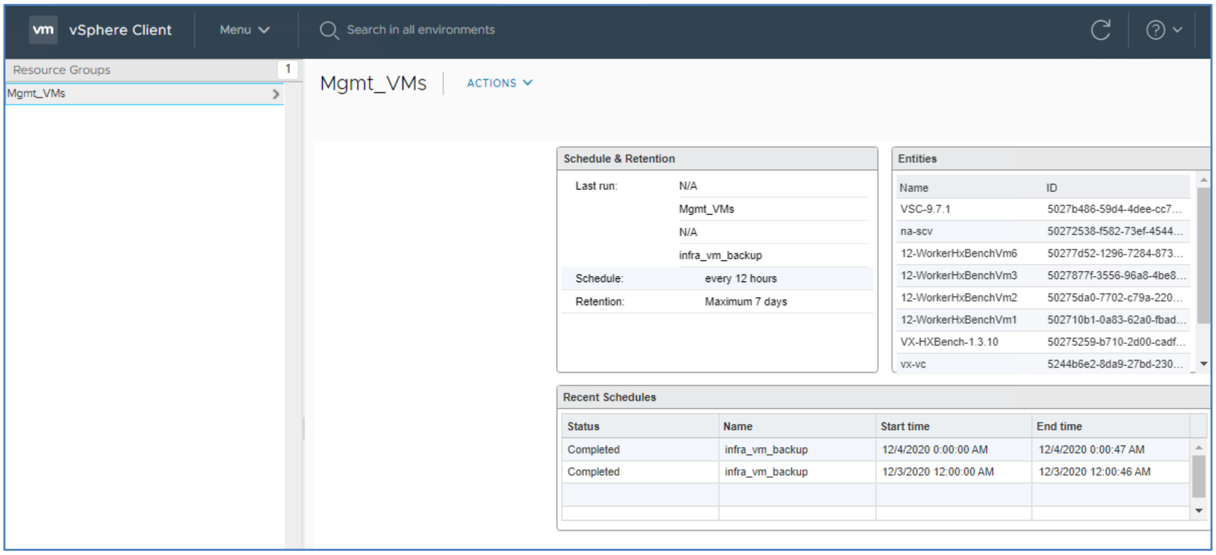

Mgmt_VMs

Schedule & Retention

Last run:	N/A
	Mgmt_VMs
	N/A
	infra_vm_backup
Schedule:	every 12 hours
Retention:	Maximum 7 days

Entities

Name	ID
VSC-9.7.1	5027b486-59d4-4dee-cc7...
na-scv	50272538-f582-73ef-4544...
12-WorkerHxBenchVm6	50277d52-1296-7284-873...
12-WorkerHxBenchVm3	5027877f-3556-96a8-4be8...
12-WorkerHxBenchVm2	50275da0-7702-c79a-220...
12-WorkerHxBenchVm1	502710b1-0a83-62a0-fbad...
VX-HXBench-1.3.10	50275259-b710-2d00-cadf...
vx-vc	5244b6e2-8da9-27bd-230...

Recent Schedules

Status	Name	Start time	End time
Completed	infra_vm_backup	12/4/2020 0:00:00 AM	12/4/2020 0:00:47 AM
Completed	infra_vm_backup	12/3/2020 12:00:00 AM	12/3/2020 12:00:46 AM

Create On-Demand Backup

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

1. From the VMs and Templates tab, choose a virtual machine contained in the resource group where you want to create an on-demand backup.

2. Click the More Objects tab and choose the Resource Groups tab from the toolbar to display the list of resource groups.

3. Right-click the resource group and click Run Now to run the backup immediately.

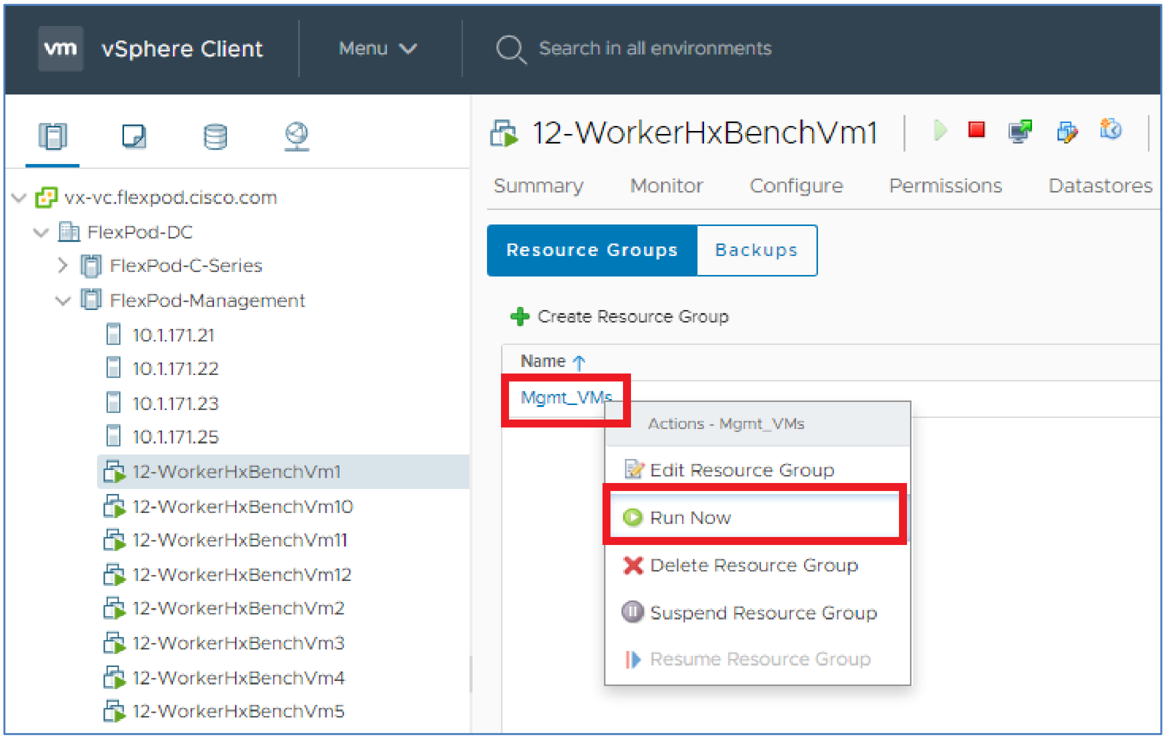

Restore from vCenter by Using SnapCenter Plug-In

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

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

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

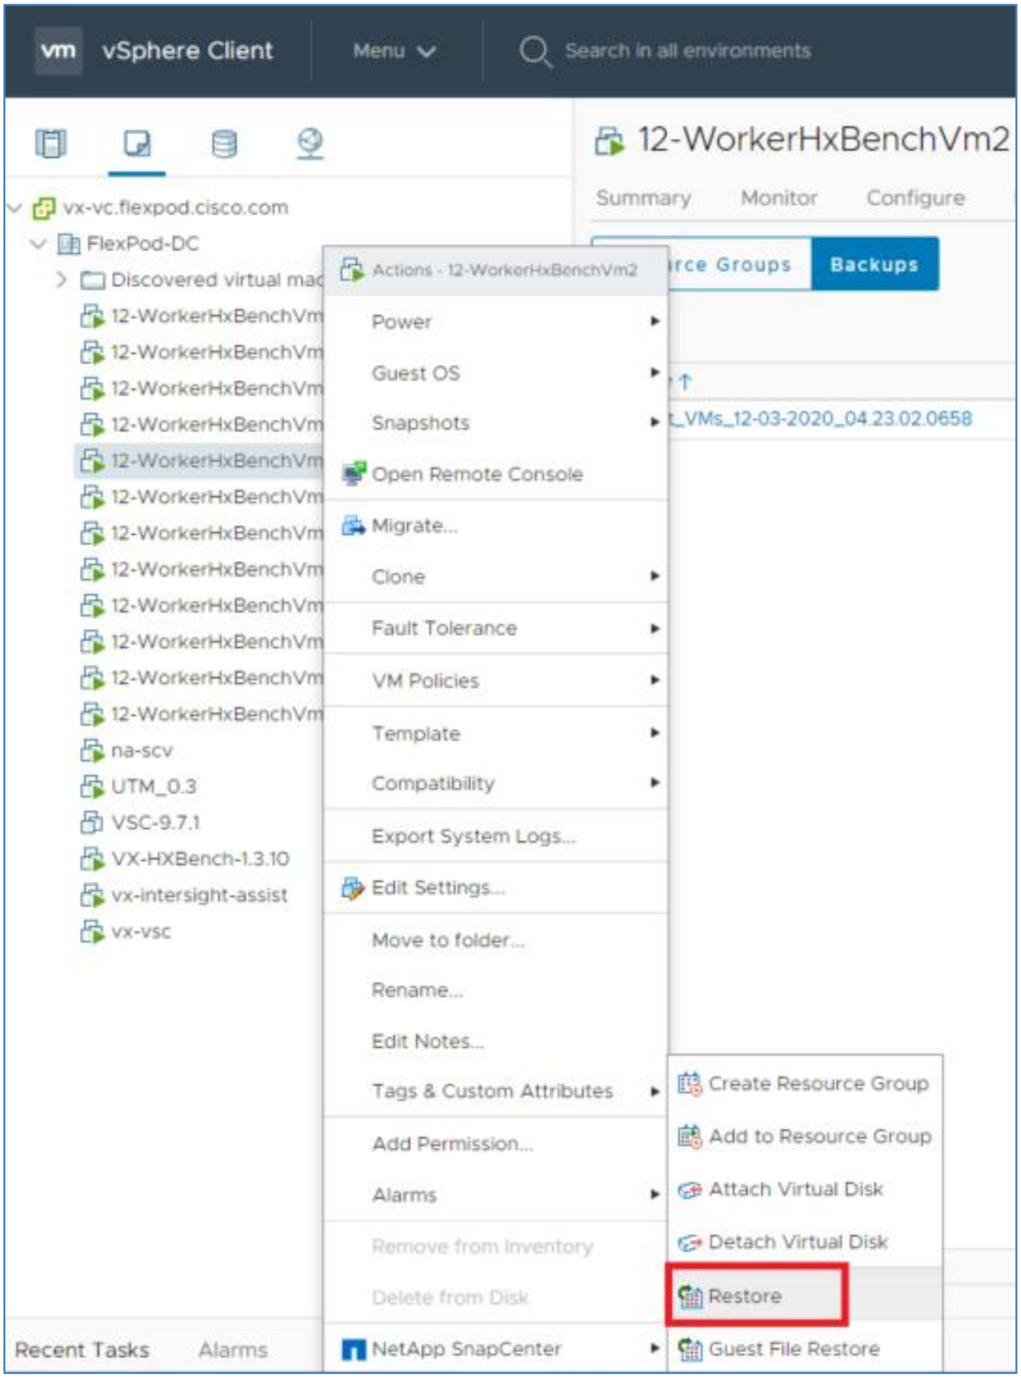

2. Choose a backup from which to restore. Click Next.

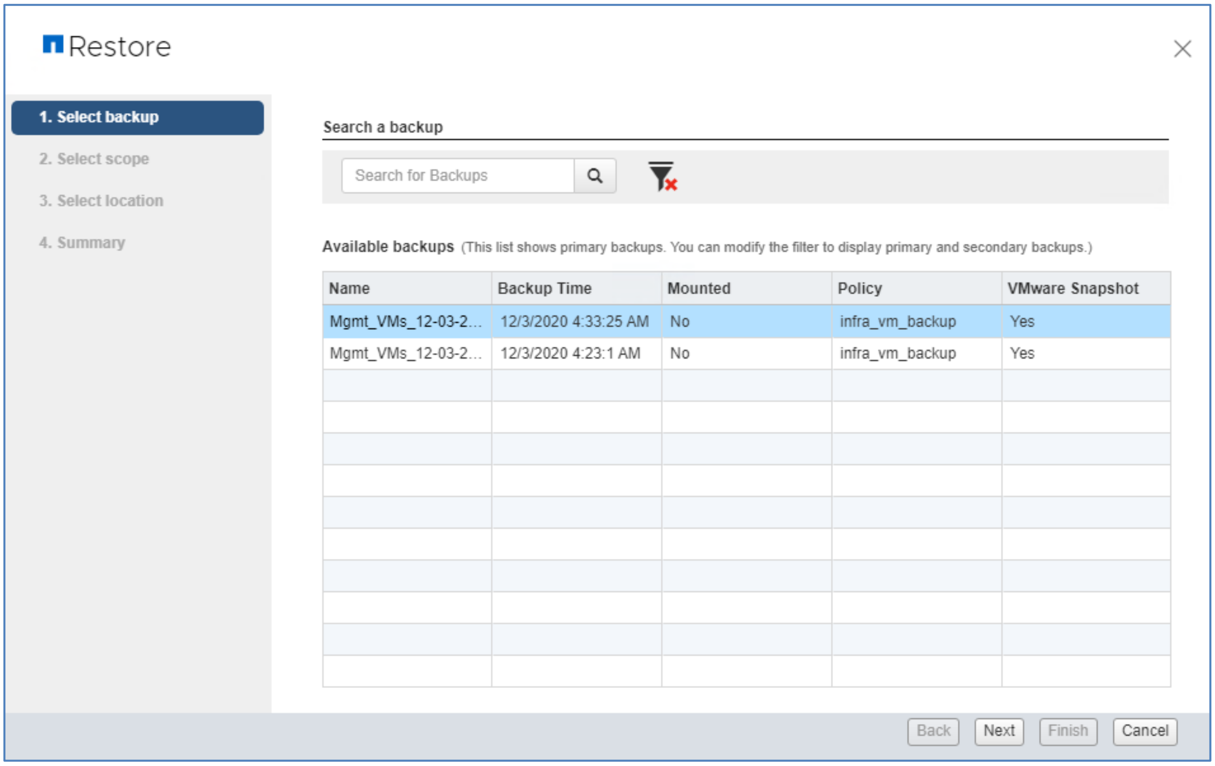

Available backups (This list shows primary backups. You can modify the filter to display primary and secondary backups.)

Name	Backup Time	Mounted	Policy	VMware Snapshot
Mgmt_VMs_12-03-2…	12/3/2020 4:33:25 AM	No	infra_vm_backup	Yes
Mgmt_VMs_12-03-2…	12/3/2020 4:23:1 AM	No	infra_vm_backup	Yes

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

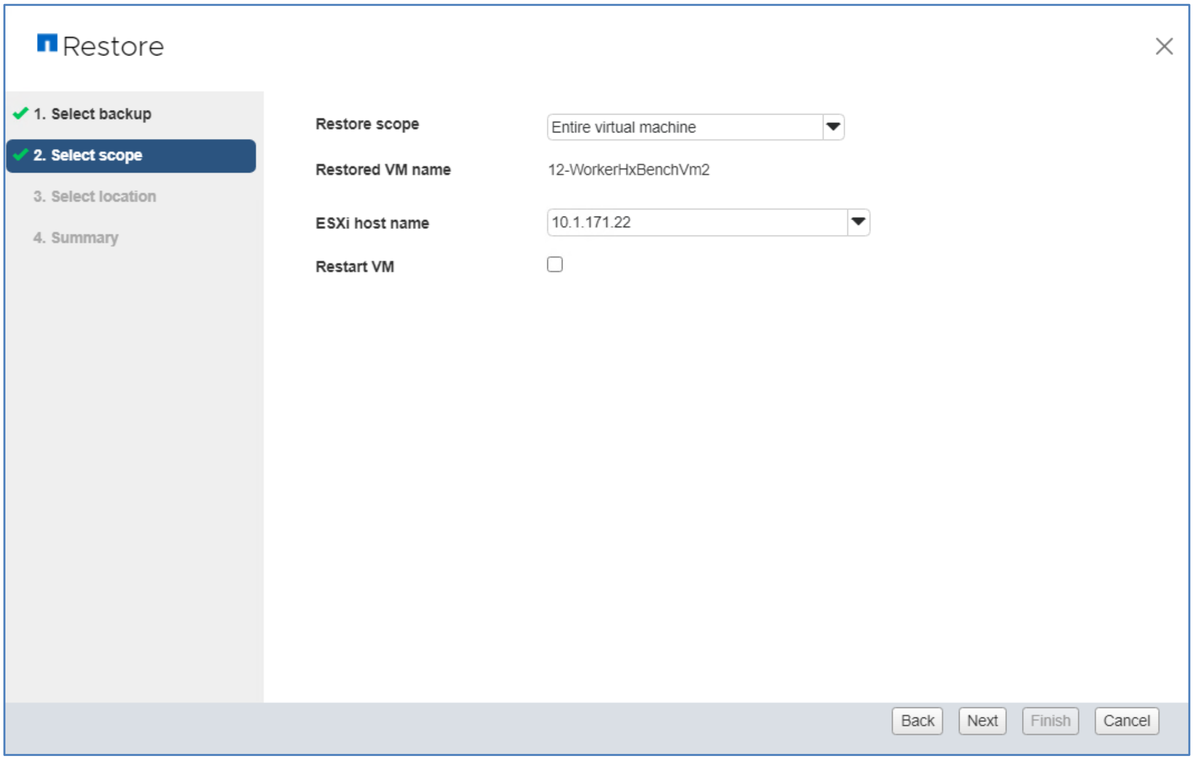

Restore scope	Entire virtual machine
Restored VM name	12-WorkerHxBenchVm2
ESXi host name	10.1.171.22
Restart VM	☐

4. Choose the destination datastore and click Next.

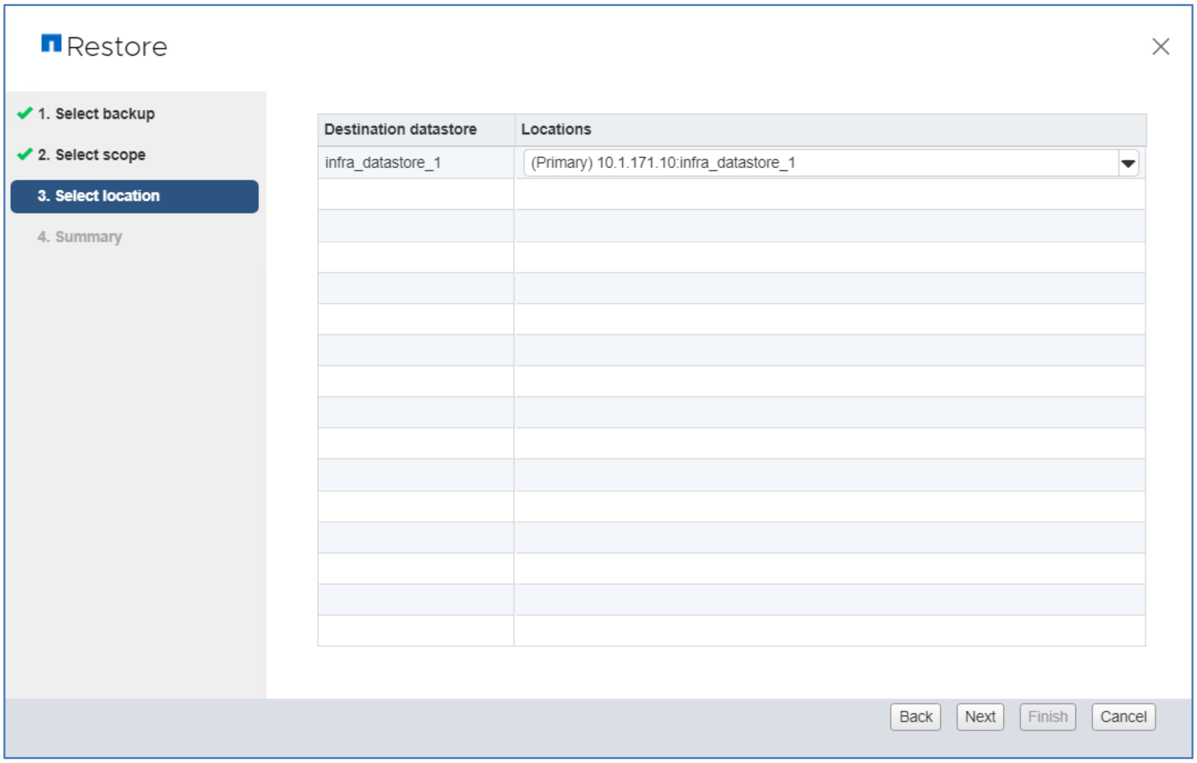

Destination datastore	Locations
infra_datastore_1	(Primary) 10.1.171.10:infra_datastore_1

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

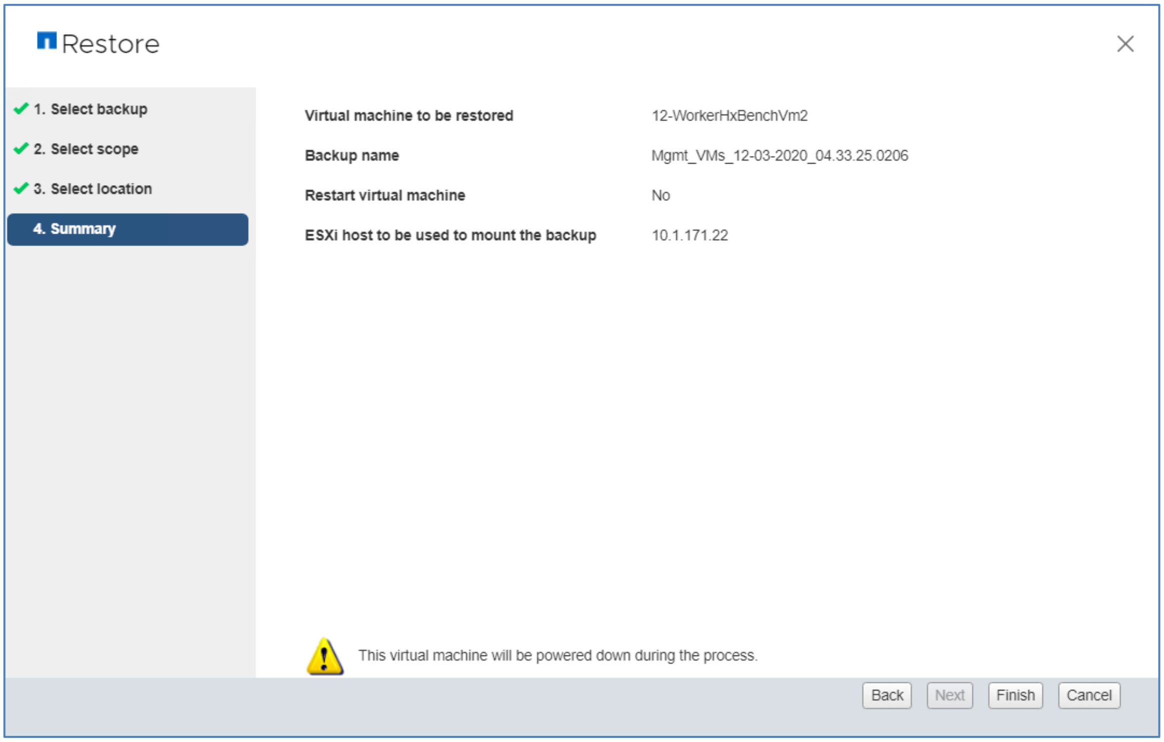

Virtual machine to be restored	12-WorkerHxBenchVm2
Backup name	Mgmt_VMs_12-03-2020_04.33.25.0206
Restart virtual machine	No
ESXi host to be used to mount the backup	10.1.171.22

This virtual machine will be powered down during the process.

Active IQ Unified Manager 9.7P1

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface. Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems.

This section describes the steps to deploy NetApp Active IQ Unified Manager 9.7P1 as a virtual appliance. The following table lists the recommended configuration for the virtual machine to install and run Active IQ Unified Manager

Table 26. Virtual Machine Configuration

Hardware Configuration	Recommended Settings
RAM	12 GB (minimum requirement 8 GB)
Processors	4 CPUs/ vCPUs
CPU Cycle Capacity	9572 MHz total (minimum requirement 9572 MHz)
Free Disk Space	5 GB (thin provisioned) 152 GB (thick provisioned)

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

To install Unified Manager 9.7P1, follow these steps:

1. Download NetApp Active IQ Unified Manager for VMware vSphere .ova file from Netapp support site.

2. From the VMware VCenter, click the VMs and Templates Tab, then click In Actions> Deploy OVF Template.

3. Specify the location of the OVF Template and click Next.

4. On the Select a name and folder page, enter a unique name for the VM or vApp, and select a deployment location, and then click Next.

5. On the Select a compute resource page, select a resource where you want to run the deployed VM template, and click Next.

6. On the Review details page, verify the .ova template details and click Next.

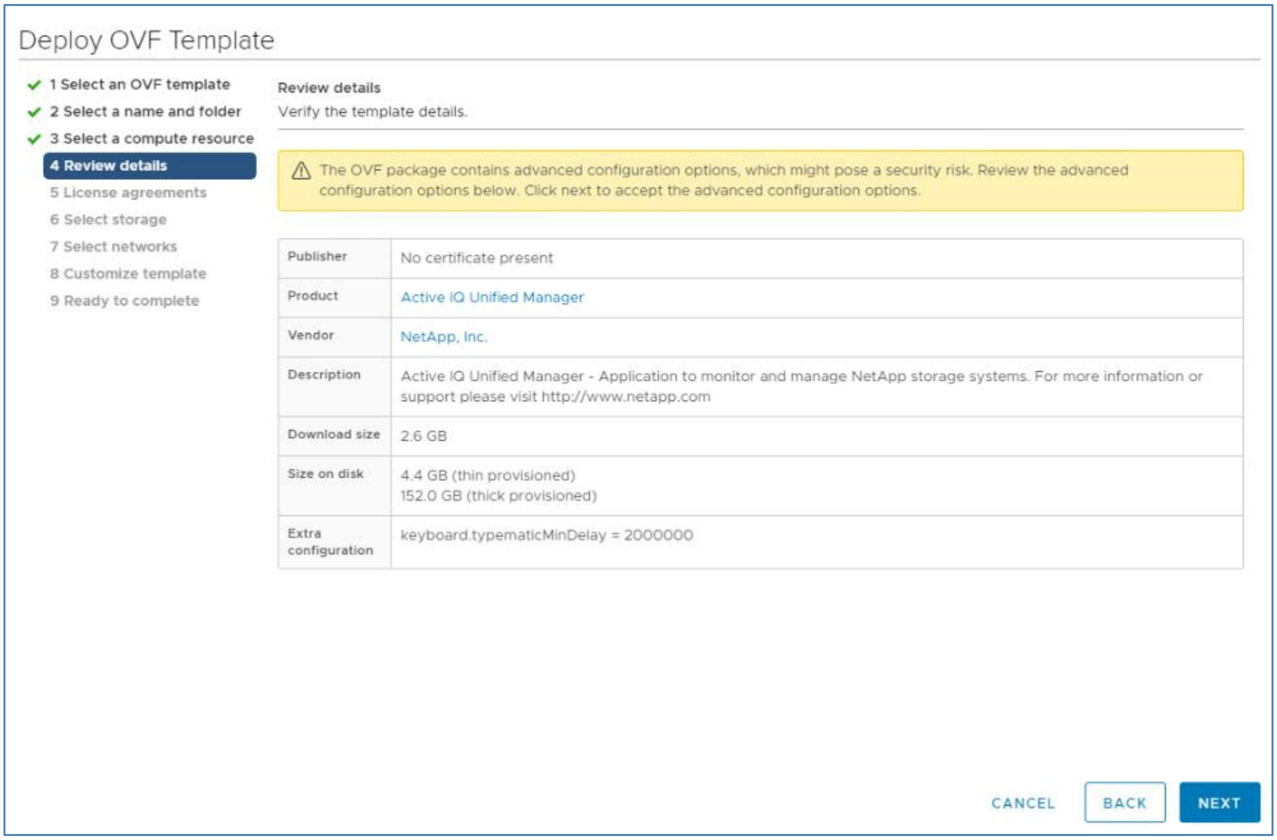

Deploy OVF Template

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

Review details

Verify the template details.

⚠ The OVF package contains advanced configuration options, which might pose a security risk. Review the advanced configuration options below. Click next to accept the advanced configuration options.

Publisher	No certificate present
Product	Active IQ Unified Manager
Vendor	NetApp, Inc.
Description	Active IQ Unified Manager - Application to monitor and manage NetApp storage systems. For more information or support please visit http://www.netapp.com
Download size	2.6 GB
Size on disk	4.4 GB (thin provisioned) 152.0 GB (thick provisioned)
Extra configuration	keyboard.typematicMinDelay = 2000000

7. On the License agreements page, select the checkbox for I accept all license agreements.

8. On the Select storage page, define where and how to store the files for the deployed OVF template.

9. Select the disk format for the VMDKs

10. Select a VM Storage Policy.

11. Select a datastore to store the deployed OVA template.

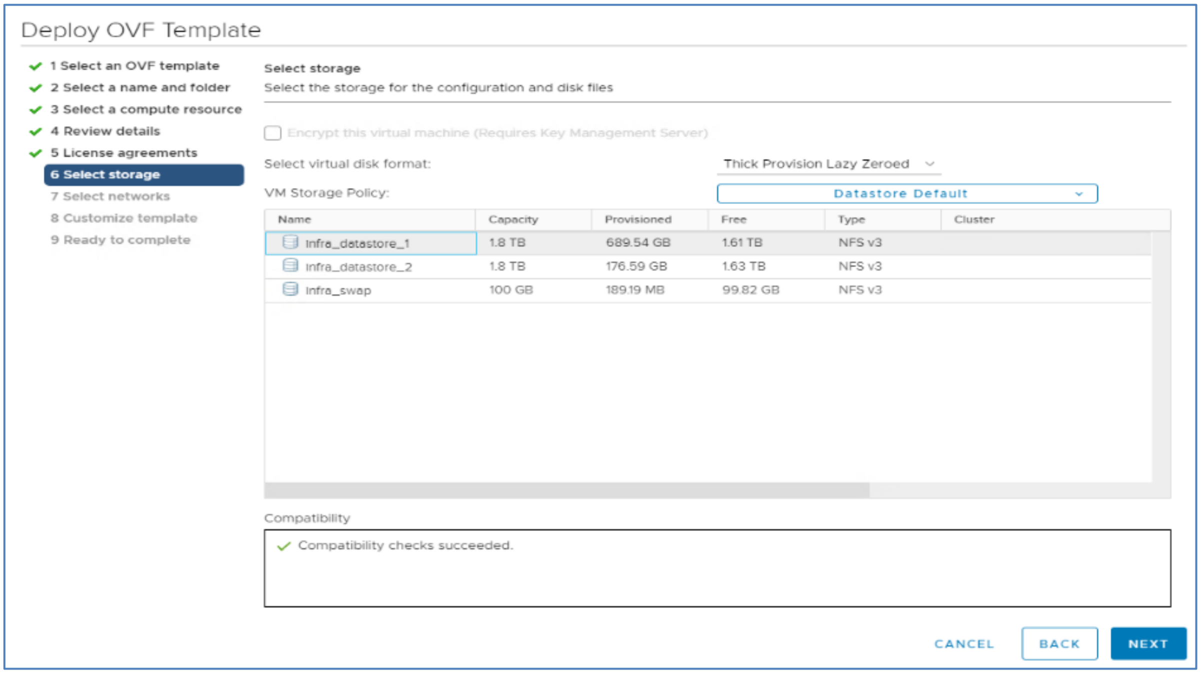

Select storage

Select the storage for the configuration and disk files

□ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
Infra_datastore_1	1.8 TB	689.54 GB	1.61 TB	NFS v3	
Infra_datastore_2	1.8 TB	176.59 GB	1.63 TB	NFS v3	
Infra_swap	100 GB	189.19 MB	99.82 GB	NFS v3	

Compatibility

✓ Compatibility checks succeeded.

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

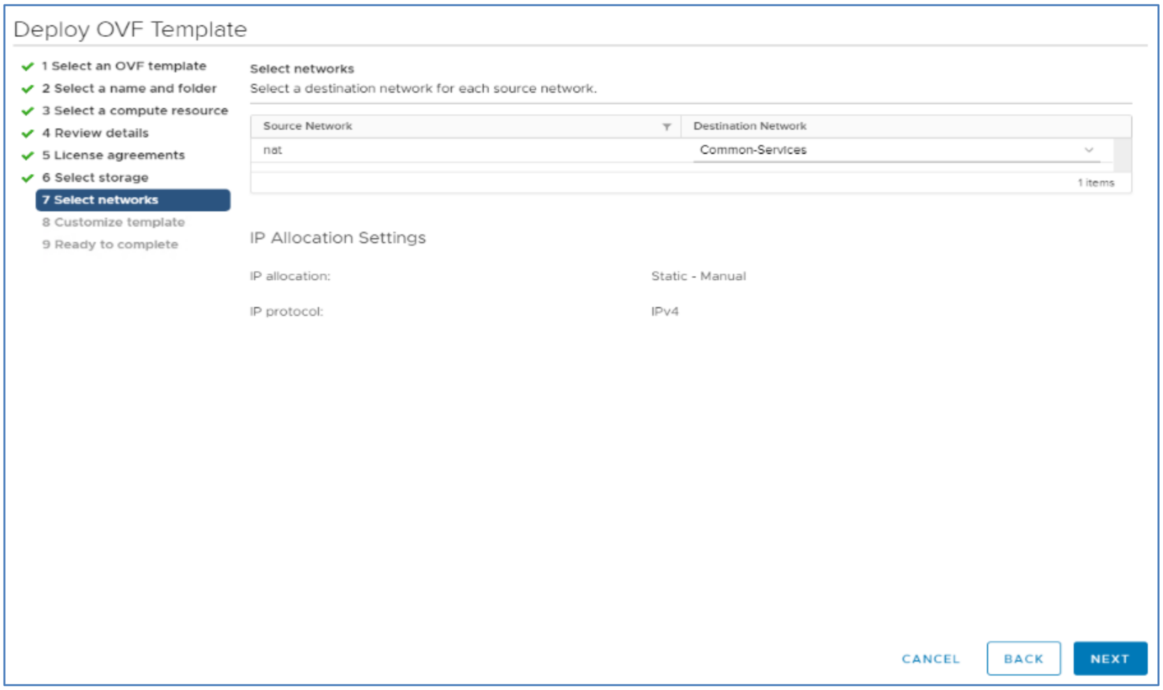

Select networks

Select a destination network for each source network.

Source Network	Destination Network
nat	Common-Services

1 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

13. On the Customize template page, provide network details.

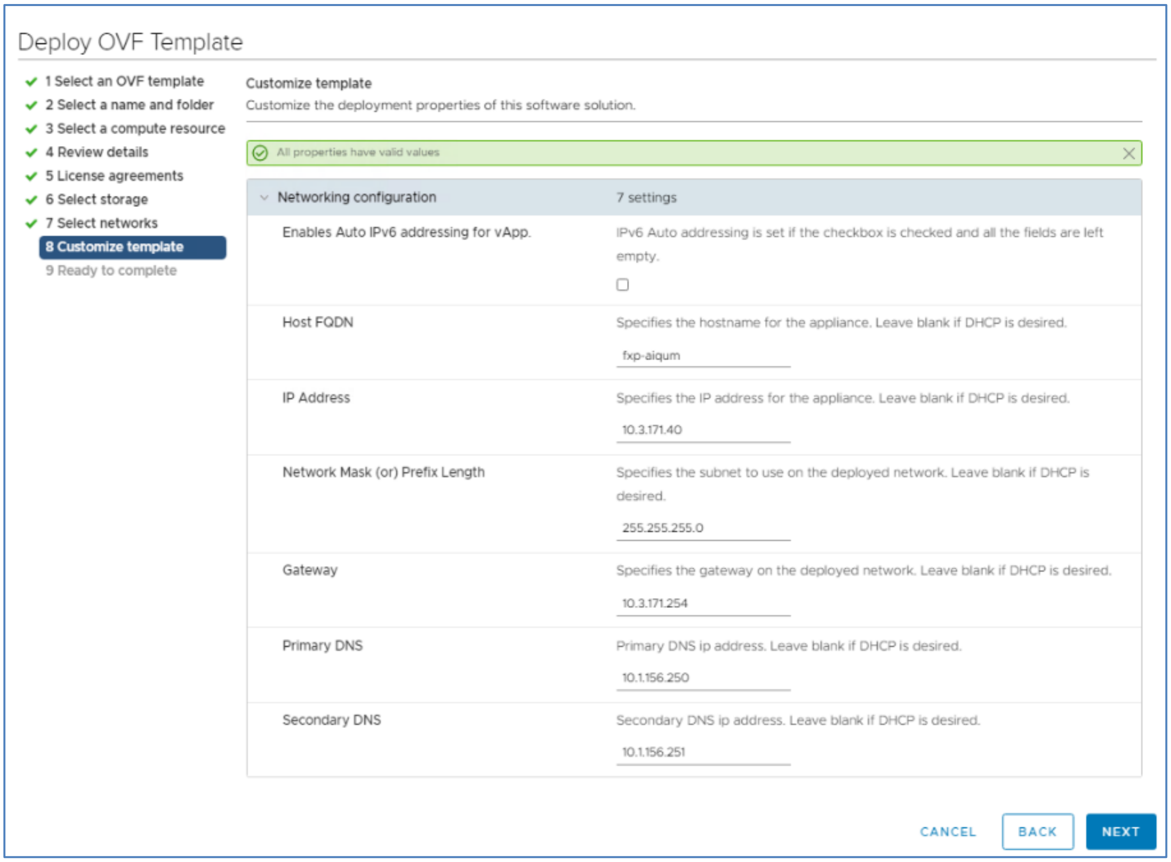

Customize template

Customize the deployment properties of this software solution.

All properties have valid values

Networking configuration 7 settings

Enables Auto IPv6 addressing for vApp. IPv6 Auto addressing is set if the checkbox is checked and all the fields are left empty. □

Host FQDN Specifies the hostname for the appliance. Leave blank if DHCP is desired. fxp-aiqum

IP Address Specifies the IP address for the appliance. Leave blank if DHCP is desired. 10.3.171.40

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

Gateway Specifies the gateway on the deployed network. Leave blank if DHCP is desired. 10.3.171.254

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

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

14. On the Ready to complete page, review the page and click Finish.

15. Navigate to the VM in the VMs and Templates Tab, then click the Summary tab, and then click the Power On box to start the virtual machine

16. While the virtual Machine is powering on, right-click the deployed virtual machine and then click Install VMware tools.

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

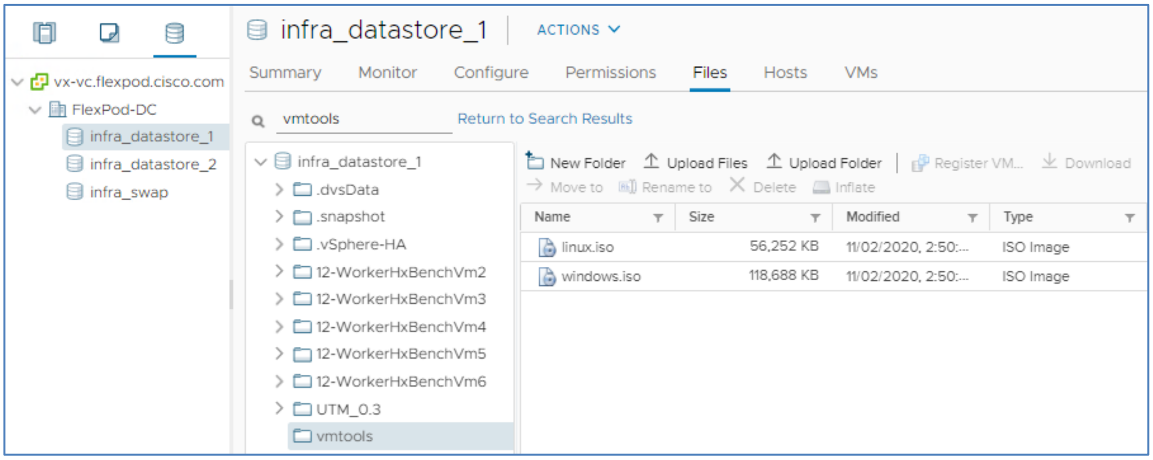

Name	Size	Modified	Type
linux.iso	56,252 KB	11/02/2020, 2:50:...	ISO Image
windows.iso	118,688 KB	11/02/2020, 2:50:...	ISO Image

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

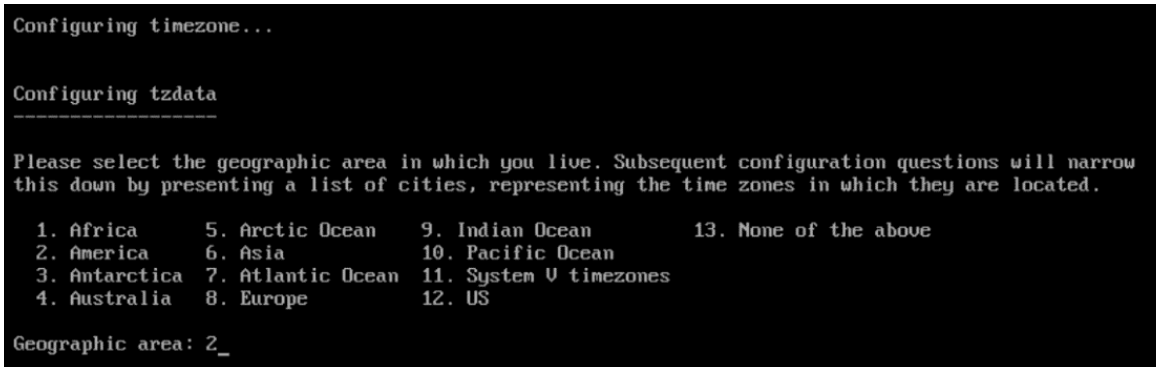

```
Configuring timezone...


Configuring tzdata
------------------

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

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

Geographic area: 2_
```

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

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

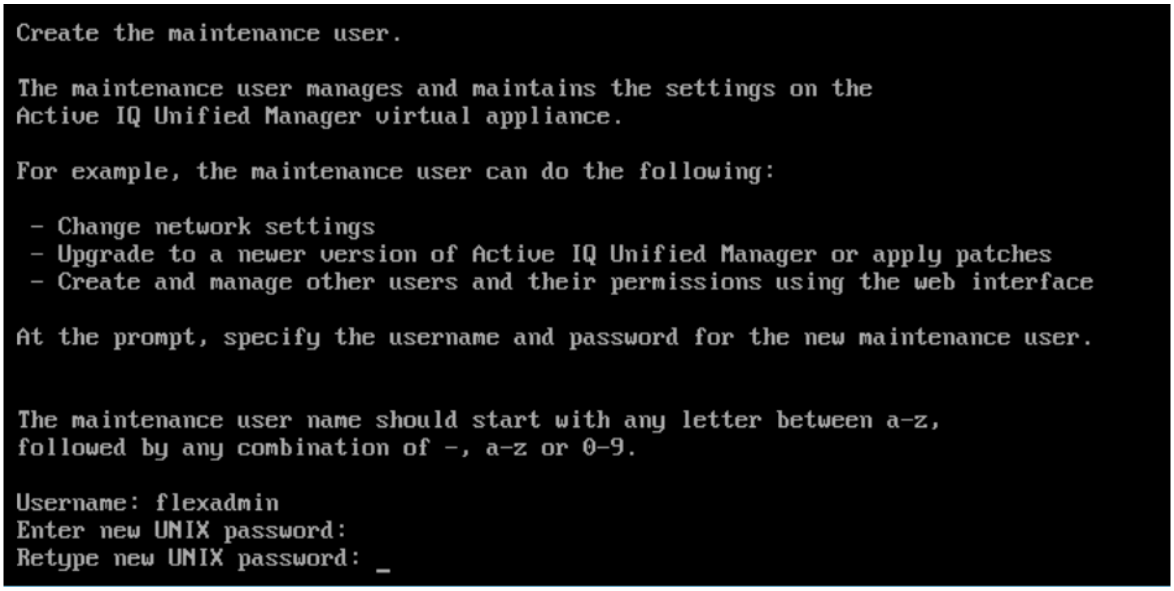

```
Create the maintenance user.

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

For example, the maintenance user can do the following:

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

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


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

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

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

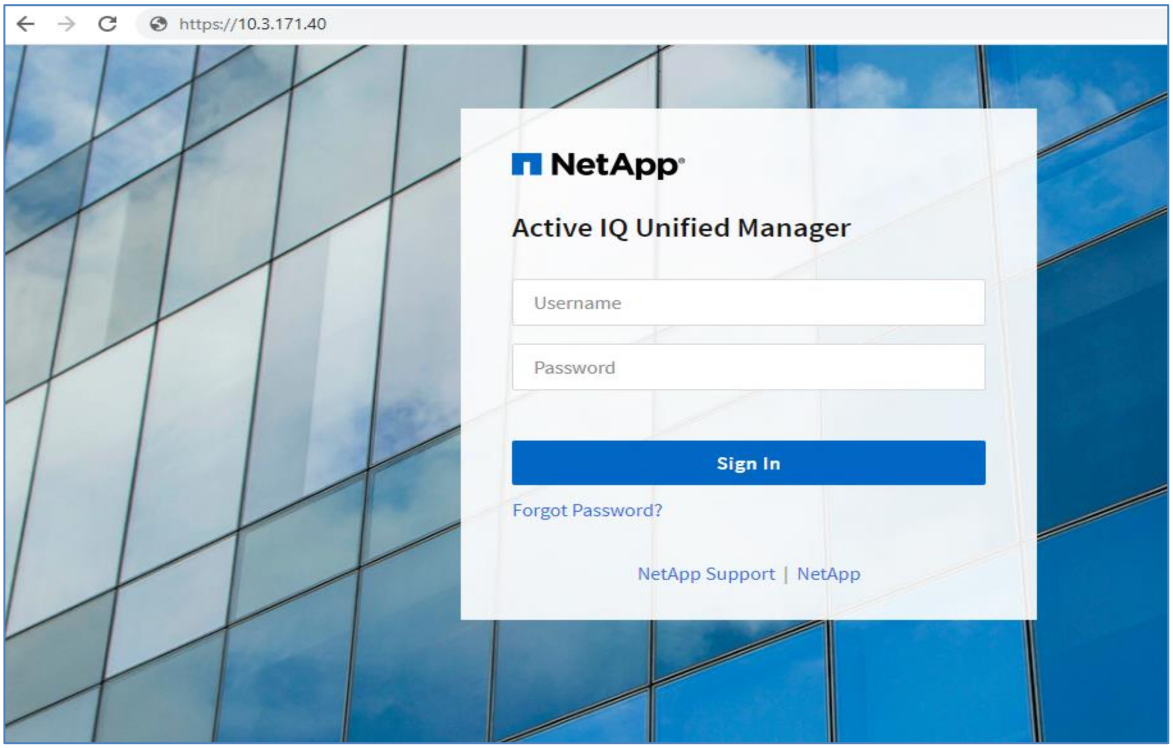

Configure Active IQ Unified Manager

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

1. Login to the Active IQ Unified Manager.

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

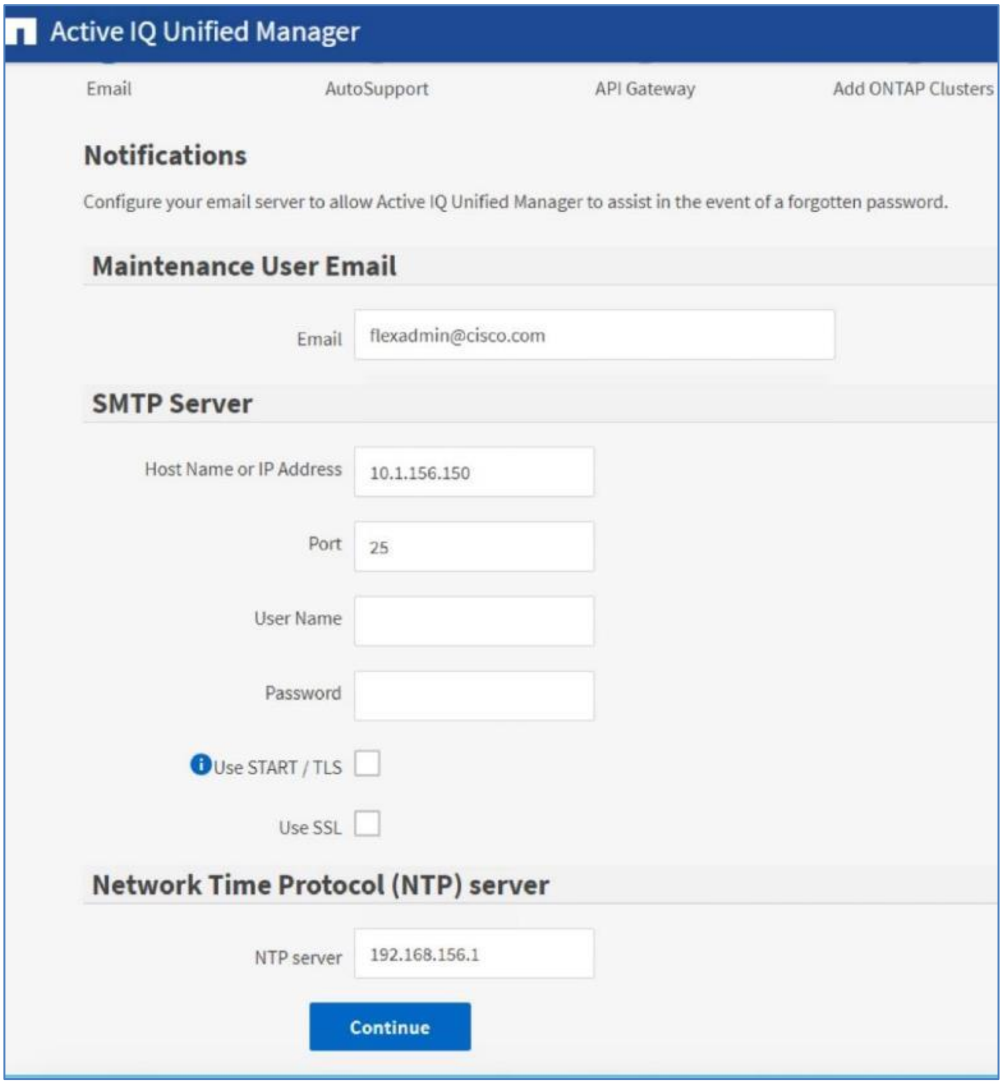

Notifications

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

Maintenance User Email

Email: flexadmin@cisco.com

SMTP Server

Host Name or IP Address: 10.1.156.150

Port: 25

User Name:

Password:

Use START / TLS □

Use SSL □

Network Time Protocol (NTP) server

NTP server: 192.168.156.1

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

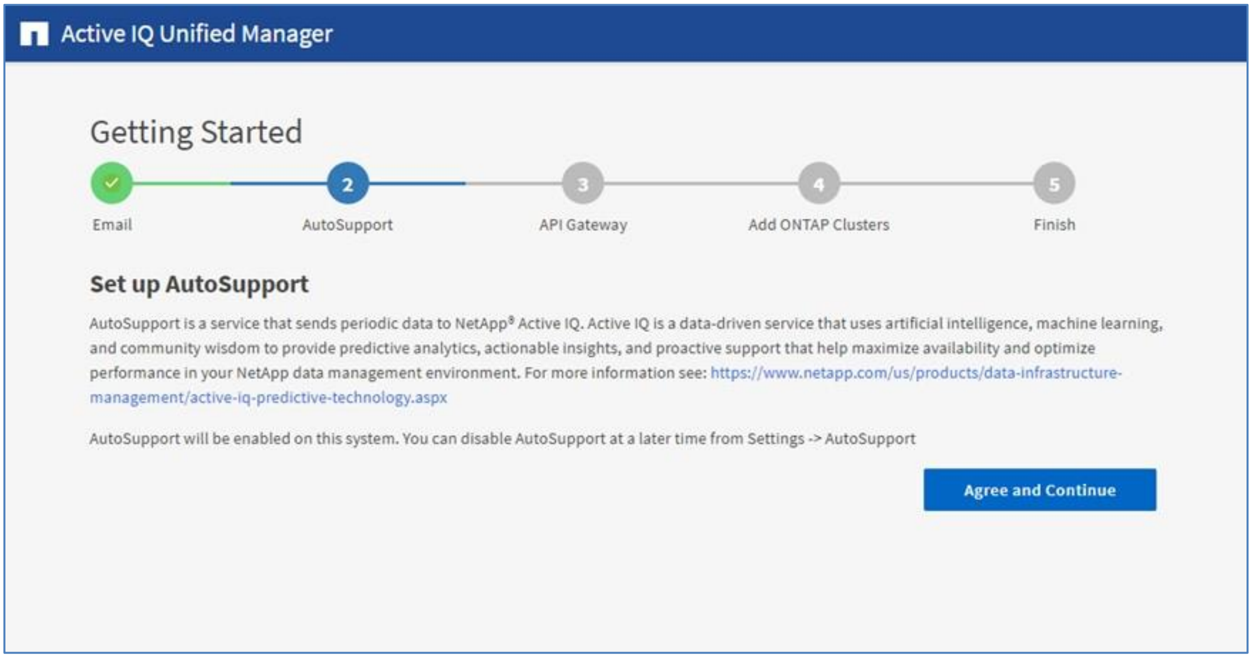

Active IQ Unified Manager

Getting Started

Set up AutoSupport

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

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

4. Configure AutoSupport for Unified Manager by clicking Continue.

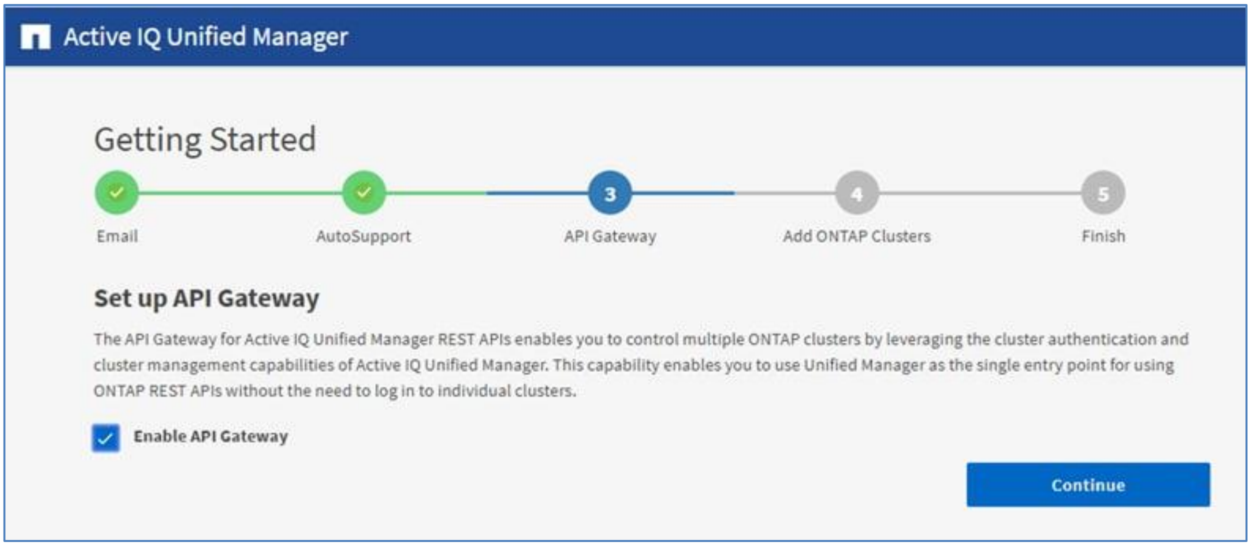

Active IQ Unified Manager

Getting Started

Set up API Gateway

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

☑ Enable API Gateway

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

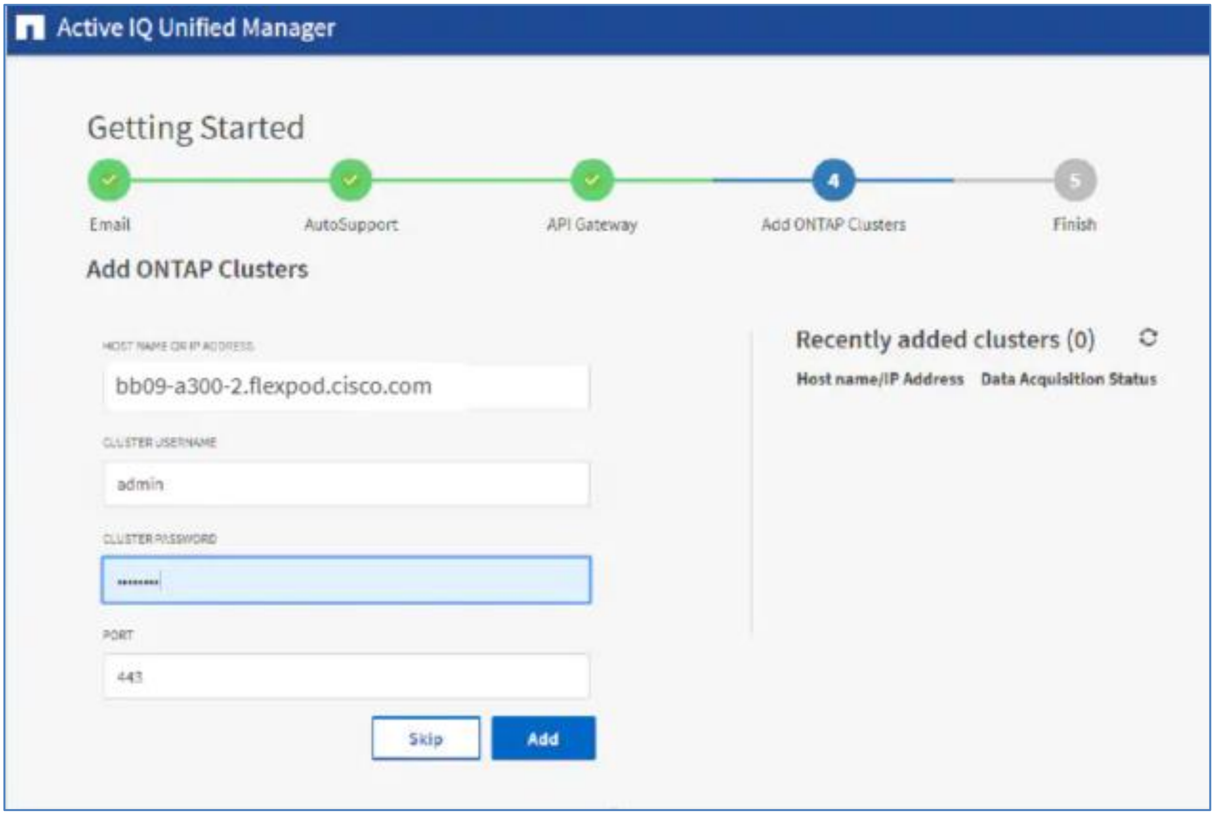

Add ONTAP Clusters

HOST NAME OR IP ADDRESS: bb09-a300-2.flexpod.cisco.com

CLUSTER USERNAME: admin

CLUSTER PASSWORD: ••••••••

PORT: 443

Recently added clusters (0)

Host name/IP Address	Data Acquisition Status

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

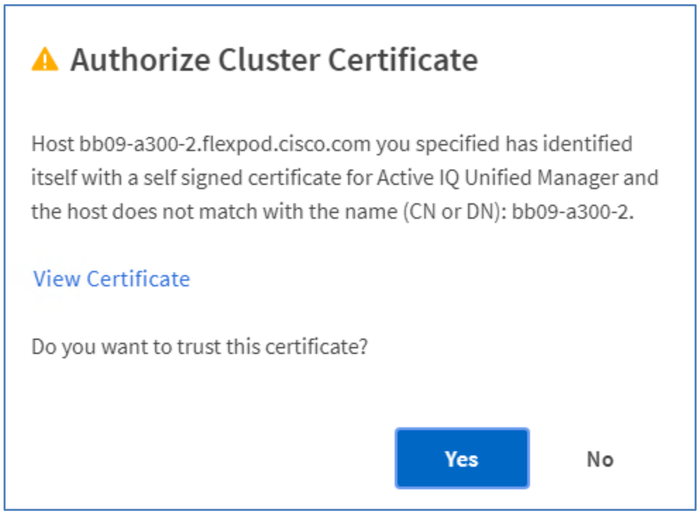

Authorize Cluster Certificate

Host bb09-a300-2.flexpod.cisco.com you specified has identified itself with a self signed certificate for Active IQ Unified Manager and the host does not match with the name (CN or DN): bb09-a300-2.

View Certificate

Do you want to trust this certificate?

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

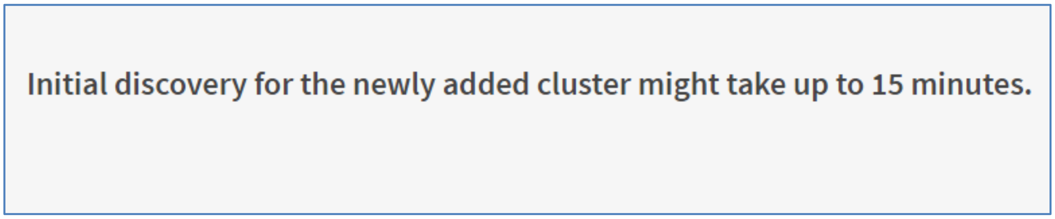

Initial discovery for the newly added cluster might take up to 15 minutes.

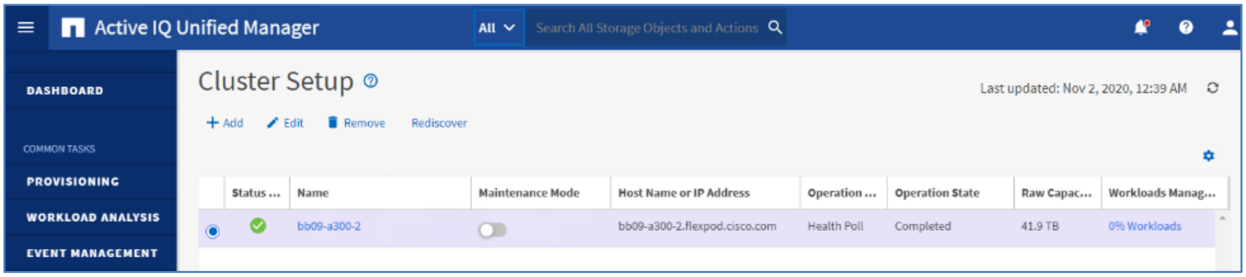

Last updated: Nov 2, 2020, 12:39 AM

Status ...	Name	Maintenance Mode	Host Name or IP Address	Operation ...	Operation State	Raw Capac...	Workloads Manag...
✓	bb09-a300-2	○	bb09-a300-2.flexpod.cisco.com	Health Poll	Completed	41.9 TB	0% Workloads

Add the vCenter Server to Active IQ Unified Manager

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

Before adding vCenter into Active IQ Unified Manager the log level of the vCenter server must be changed by following these steps:

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

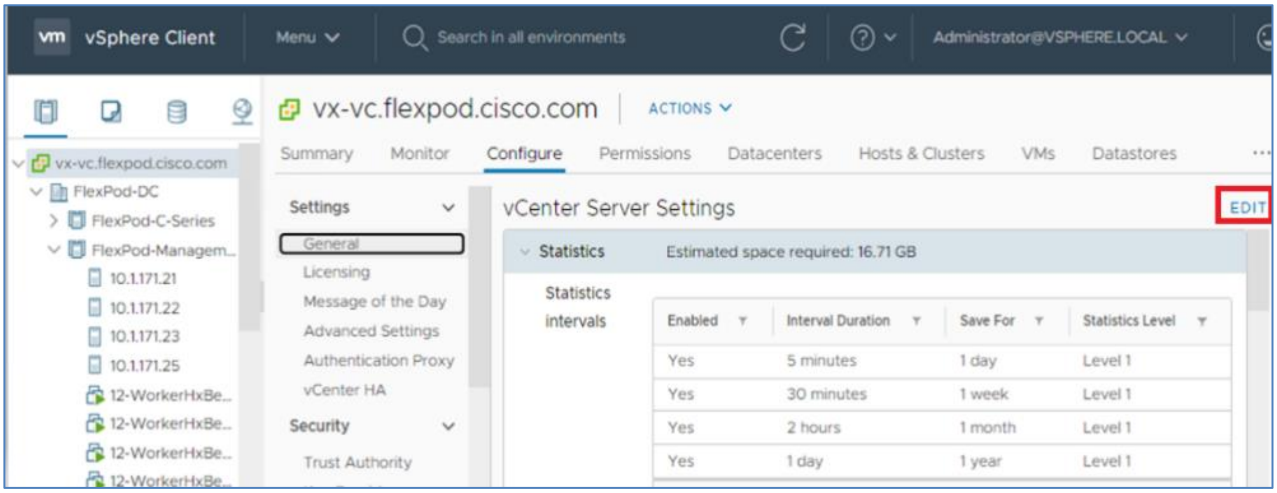

Estimated space required: 16.71 GB

Statistics intervals

Enabled	Interval Duration	Save For	Statistics Level
Yes	5 minutes	1 day	Level 1
Yes	30 minutes	1 week	Level 1
Yes	2 hours	1 month	Level 1
Yes	1 day	1 year	Level 1

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

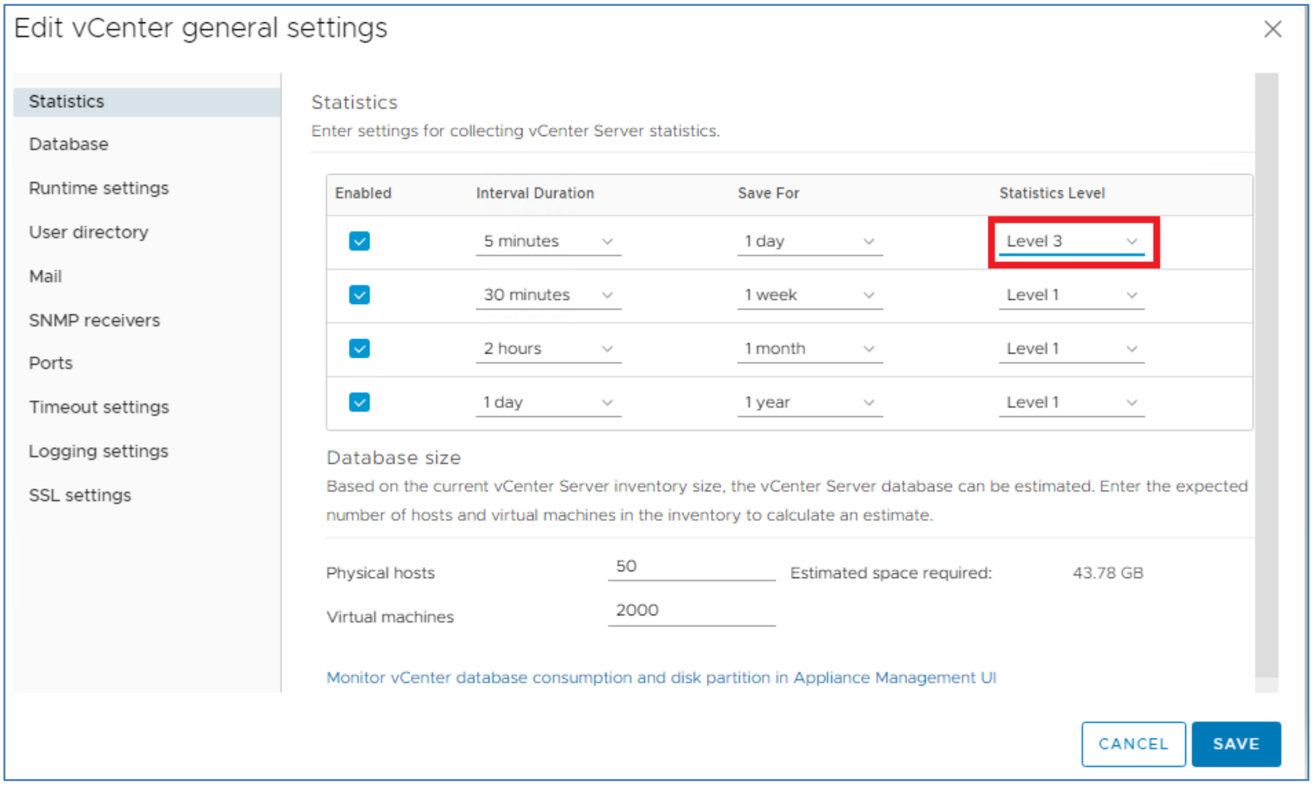

Statistics

Enter settings for collecting vCenter Server statistics.

Enabled	Interval Duration	Save For	Statistics Level
☑	5 minutes	1 day	Level 3
☑	30 minutes	1 week	Level 1
☑	2 hours	1 month	Level 1
☑	1 day	1 year	Level 1

Database size

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

Physical hosts 50 Estimated space required: 43.78 GB

Virtual machines 2000

Monitor vCenter database consumption and disk partition in Appliance Management UI

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

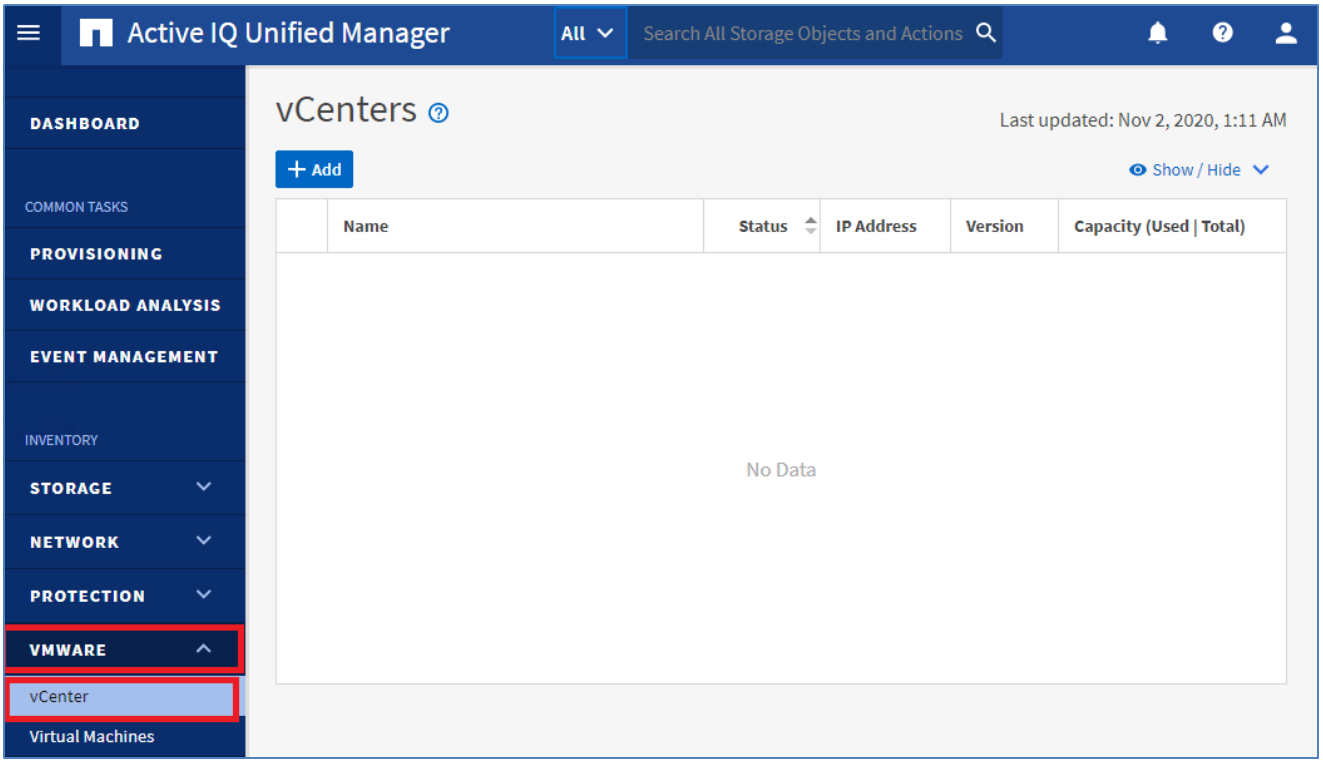

vCenters

Last updated: Nov 2, 2020, 1:11 AM

Name	Status	IP Address	Version	Capacity (Used | Total)

No Data

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

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

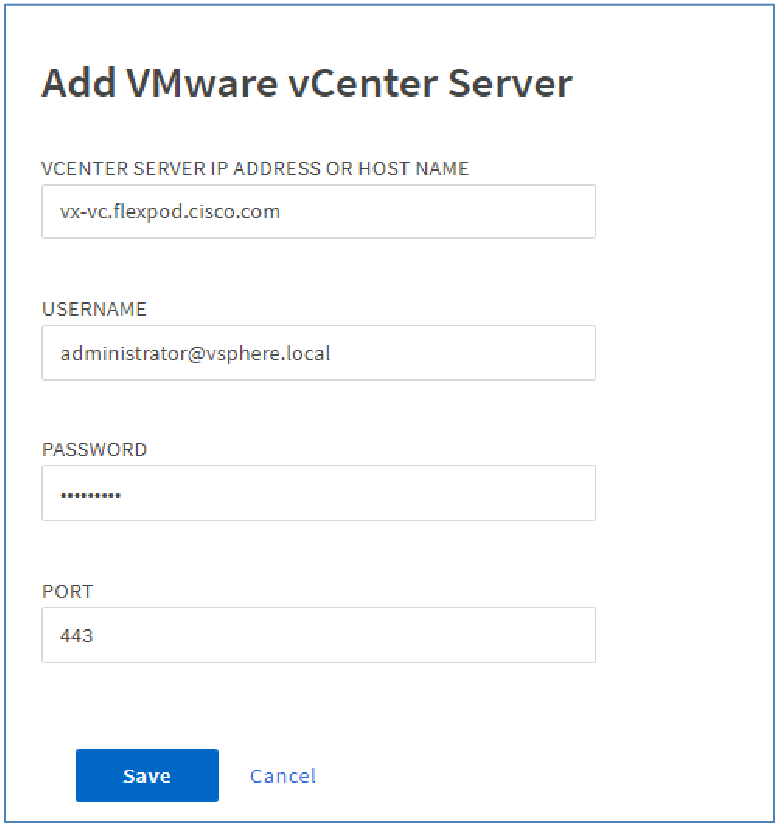

Add VMware vCenter Server

VCENTER SERVER IP ADDRESS OR HOST NAME: vx-vc.flexpod.cisco.com

USERNAME: administrator@vsphere.local

PASSWORD: •••••••••

PORT: 443

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

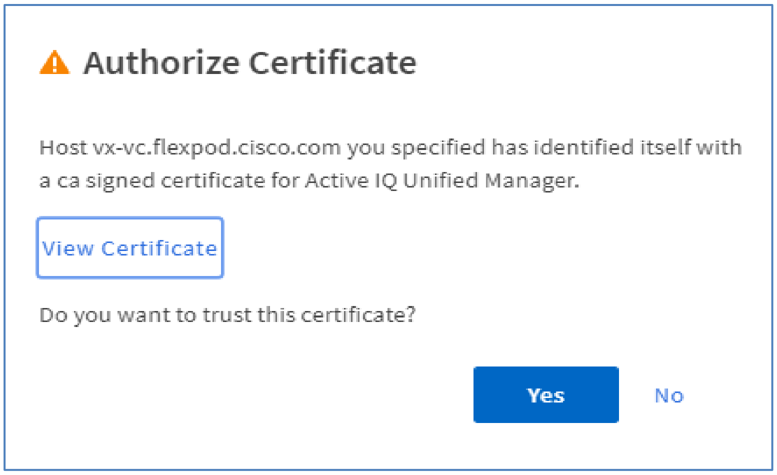

Authorize Certificate

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

Do you want to trust this certificate?

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

View Virtual Machine Inventory

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

To review the virtual machine topology and statics, follow these steps:

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

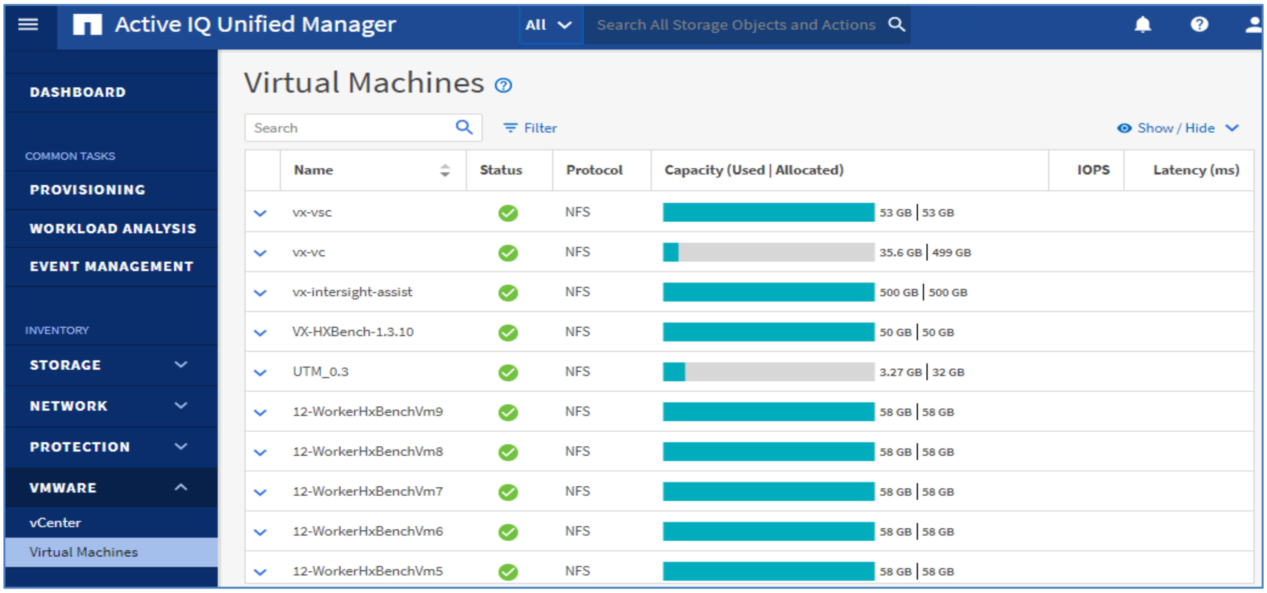

Virtual Machines

	Name	Status	Protocol	Capacity (Used | Allocated)	IOPS	Latency (ms)
⌄	vx-vsc	✓	NFS	53 GB | 53 GB		
⌄	vx-vc	✓	NFS	35.6 GB | 499 GB		
⌄	vx-intersight-assist	✓	NFS	500 GB | 500 GB		
⌄	VX-HXBench-1.3.10	✓	NFS	50 GB | 50 GB		
⌄	UTM_0.3	✓	NFS	3.27 GB | 32 GB		
⌄	12-WorkerHxBenchVm9	✓	NFS	58 GB | 58 GB		
⌄	12-WorkerHxBenchVm8	✓	NFS	58 GB | 58 GB		
⌄	12-WorkerHxBenchVm7	✓	NFS	58 GB | 58 GB		
⌄	12-WorkerHxBenchVm6	✓	NFS	58 GB | 58 GB		
⌄	12-WorkerHxBenchVm5	✓	NFS	58 GB | 58 GB		

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

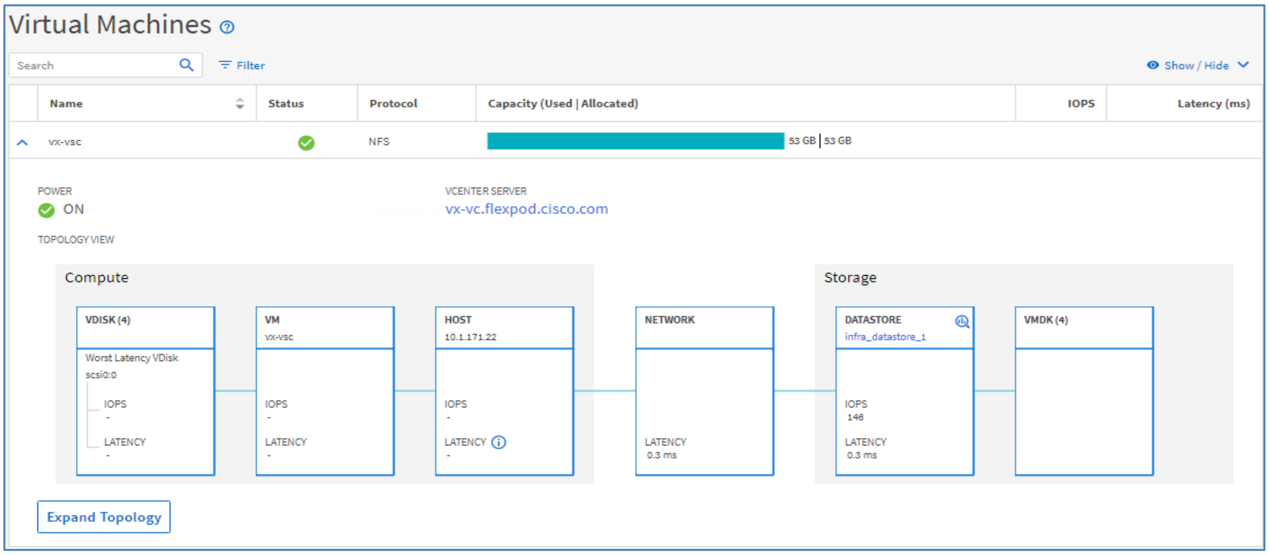

Virtual Machines

	Name	Status	Protocol	Capacity (Used | Allocated)	IOPS	Latency (ms)
^	vx-vsc	✓	NFS	53 GB | 53 GB		

POWER
✓ ON

VCENTER SERVER
vx-vc.flexpod.cisco.com

TOPOLOGY VIEW

VDISK (4)	VM vx-vsc	HOST 10.1.171.22	NETWORK	DATASTORE infra_datastore_1	VMDK (4)
Worst Latency VDisk scsi0:0 IOPS - LATENCY -	IOPS - LATENCY -	IOPS - LATENCY -	LATENCY 0.3 ms	IOPS 146 LATENCY 0.3 ms	

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

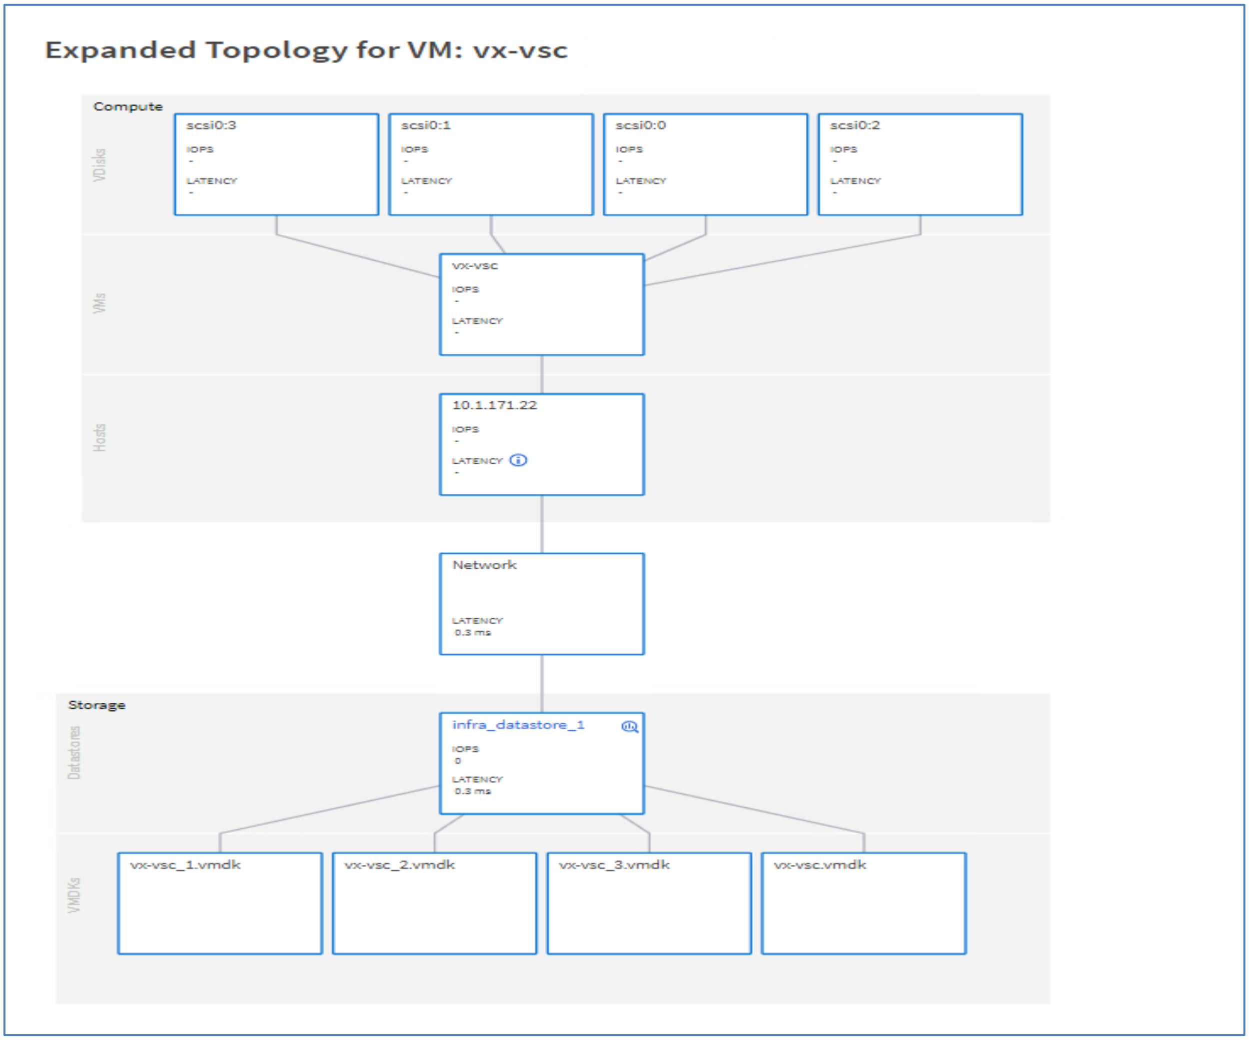

Review Security Compliance with Active IQ Unified Manager

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

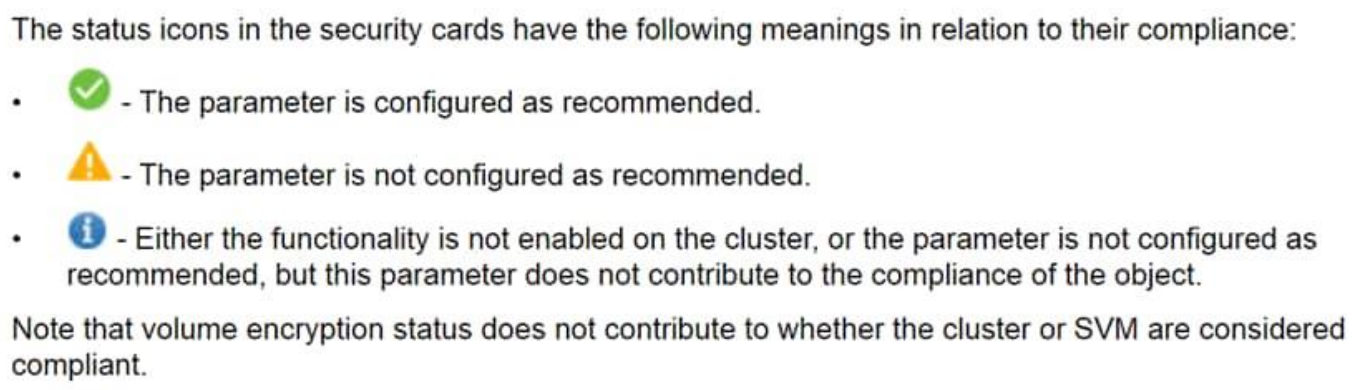

The status icons in the security cards have the following meanings in relation to their compliance:

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

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

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

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

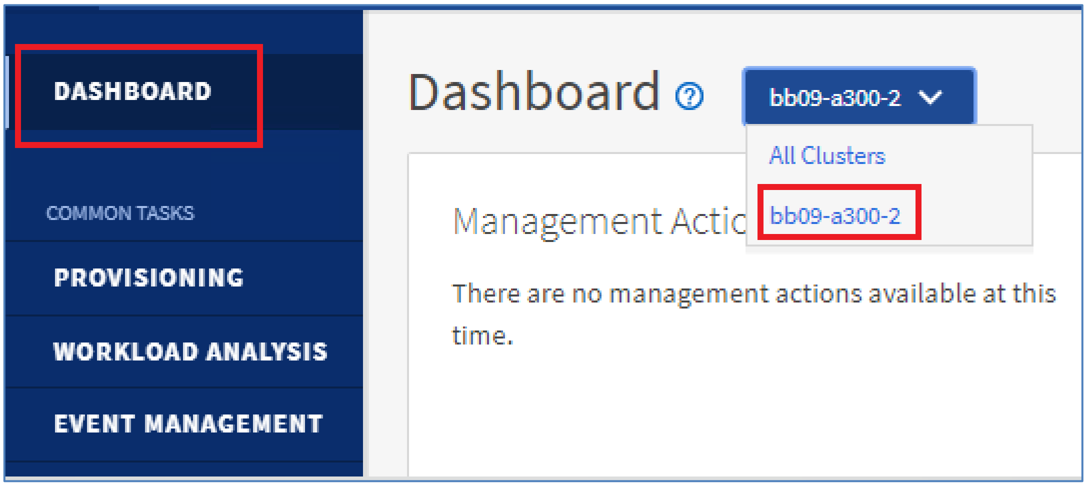

There are no management actions available at this time.

4. Locate the Security card and note the compliance level of the cluster and SVM. Click the blue arrow to expand the findings.

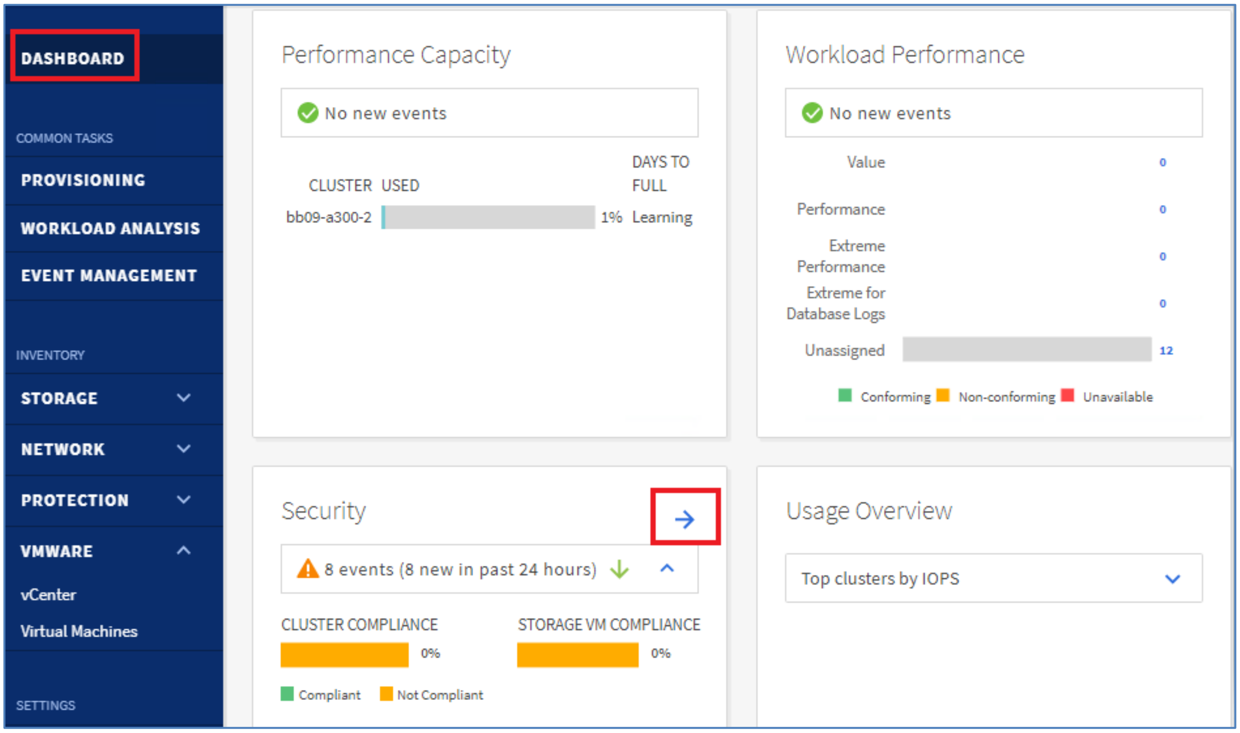

5. From the drop-down list choose View All.

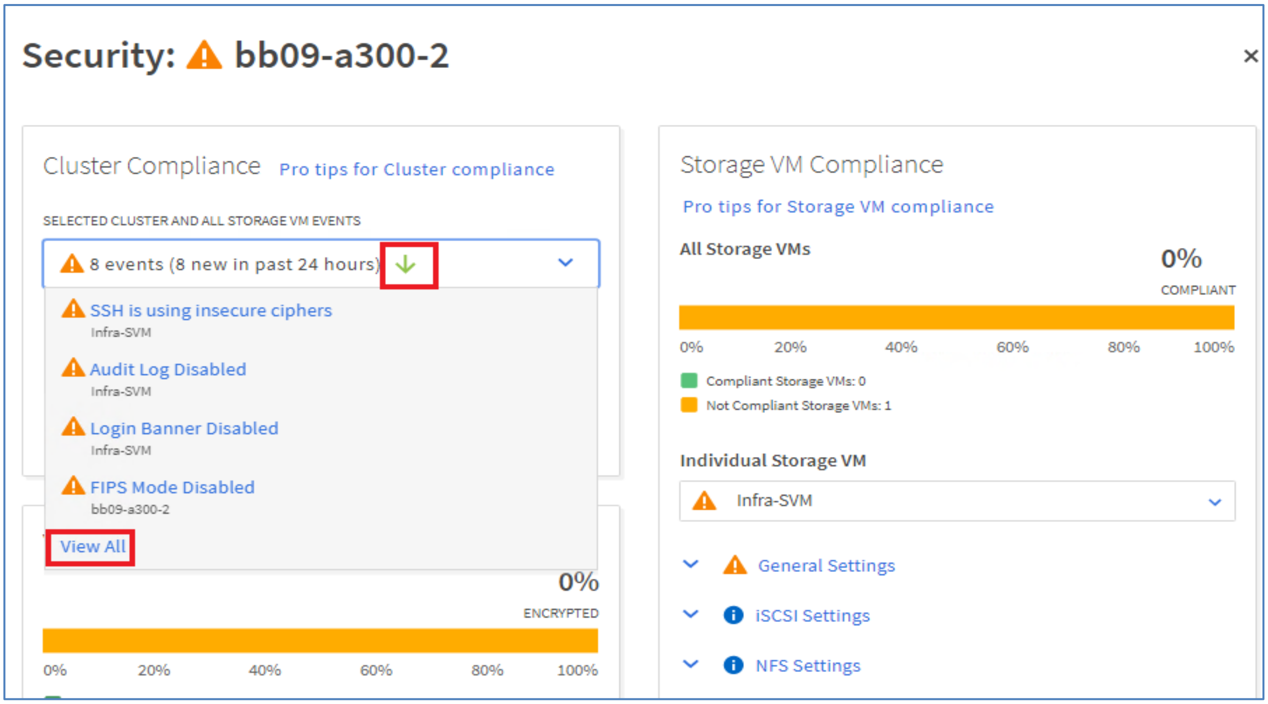

6. Choose an event from the list and click the name of the event to view the remediation steps.

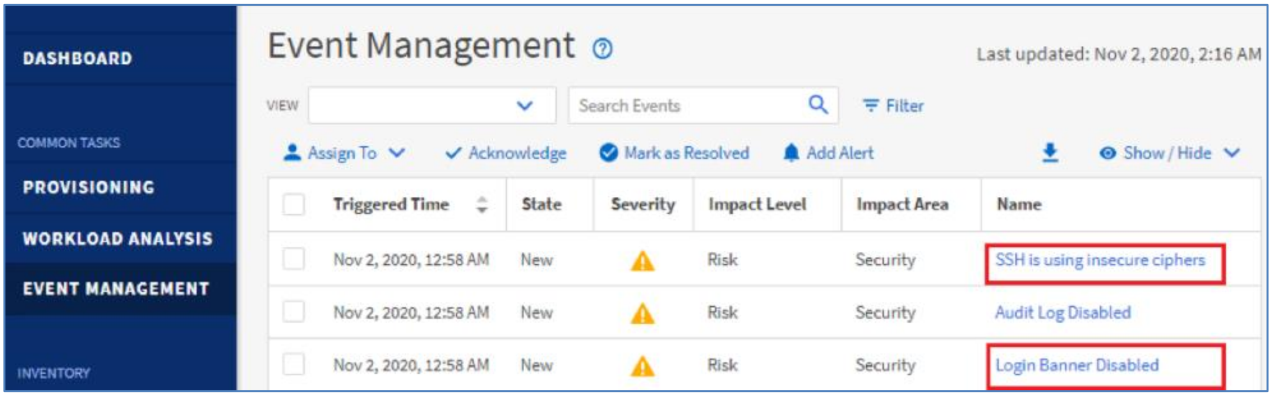

Event Management

Last updated: Nov 2, 2020, 2:16 AM

□	Triggered Time	State	Severity	Impact Level	Impact Area	Name
□	Nov 2, 2020, 12:58 AM	New	⚠	Risk	Security	SSH is using insecure ciphers
□	Nov 2, 2020, 12:58 AM	New	⚠	Risk	Security	Audit Log Disabled
□	Nov 2, 2020, 12:58 AM	New	⚠	Risk	Security	Login Banner Disabled

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

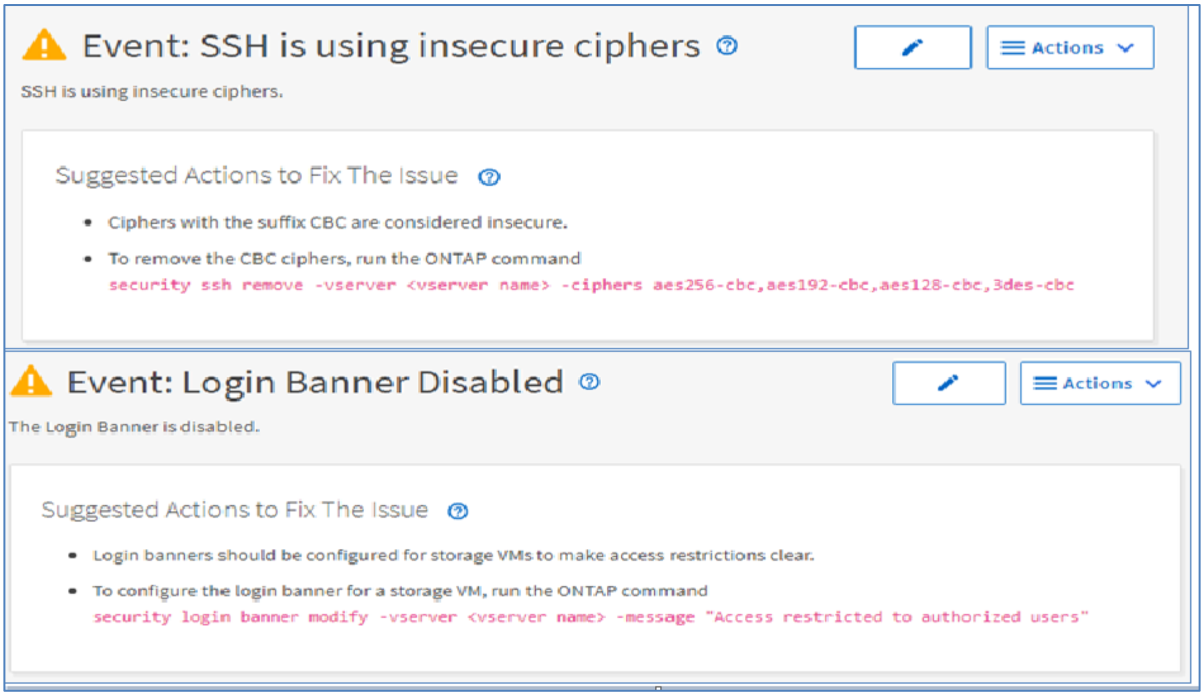

Event: SSH is using insecure ciphers

SSH is using insecure ciphers.

Suggested Actions to Fix The Issue

- Ciphers with the suffix CBC are considered insecure.
- To remove the CBC ciphers, run the ONTAP command `security ssh remove -vserver <vserver name> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc`

Event: Login Banner Disabled

The Login Banner is disabled.

Suggested Actions to Fix The Issue

- Login banners should be configured for storage VMs to make access restrictions clear.
- To configure the login banner for a storage VM, run the ONTAP command `security login banner modify -vserver <vserver name> -message "Access restricted to authorized users"`

Remediate Security Compliance Findings

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

Correct Cluster Risks

To correct cluster risks, follow these steps:

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

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

2. Enable the login banner on the cluster:

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

Correct Infrastructure Storage VM Risks

To correct infrastructure storage VM risks, follow these steps:

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

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

2. Enable the login banner on the cluster:

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

NetApp Active IQ

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

Additional Active IQ documentation is available on the Active IQ Documentation Resources web page.

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

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

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

2. Navigate to the Active IQ portal at https://activeiq.netapp.com/

3. Login with you NetApp support account ID.

4. At the welcome screen enter the cluster name or one of controller serial numbers in the search box. Active IQ will automatically begin searching for the cluster and display results below.

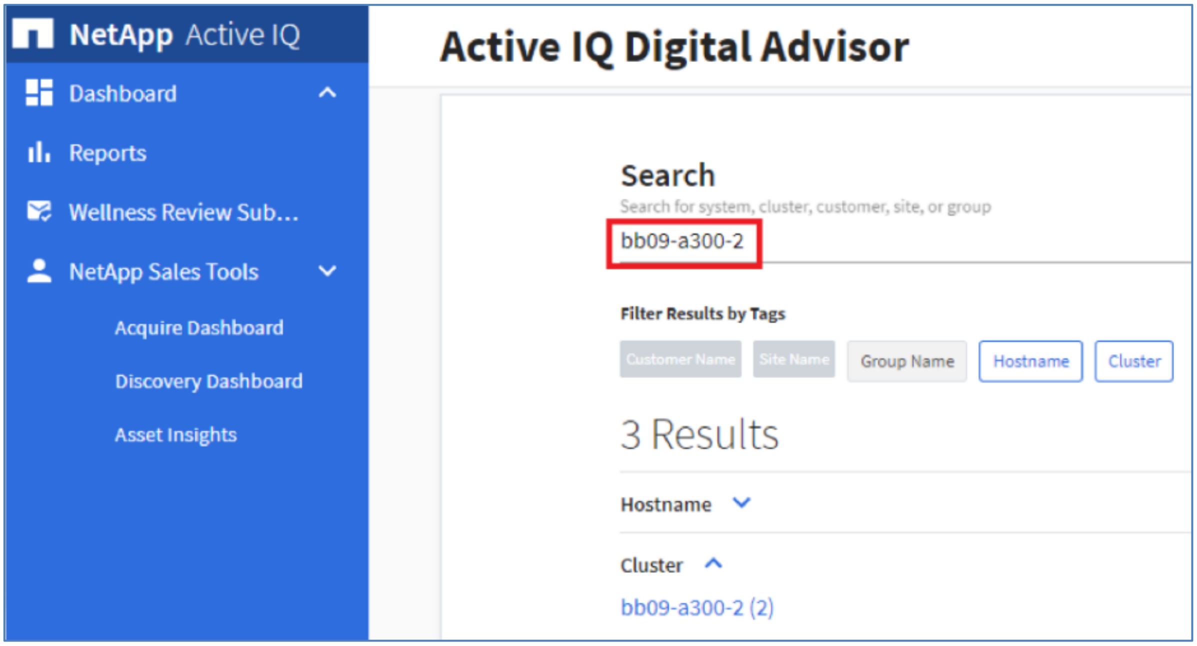

5. Choose the cluster name to launch the main dashboard.

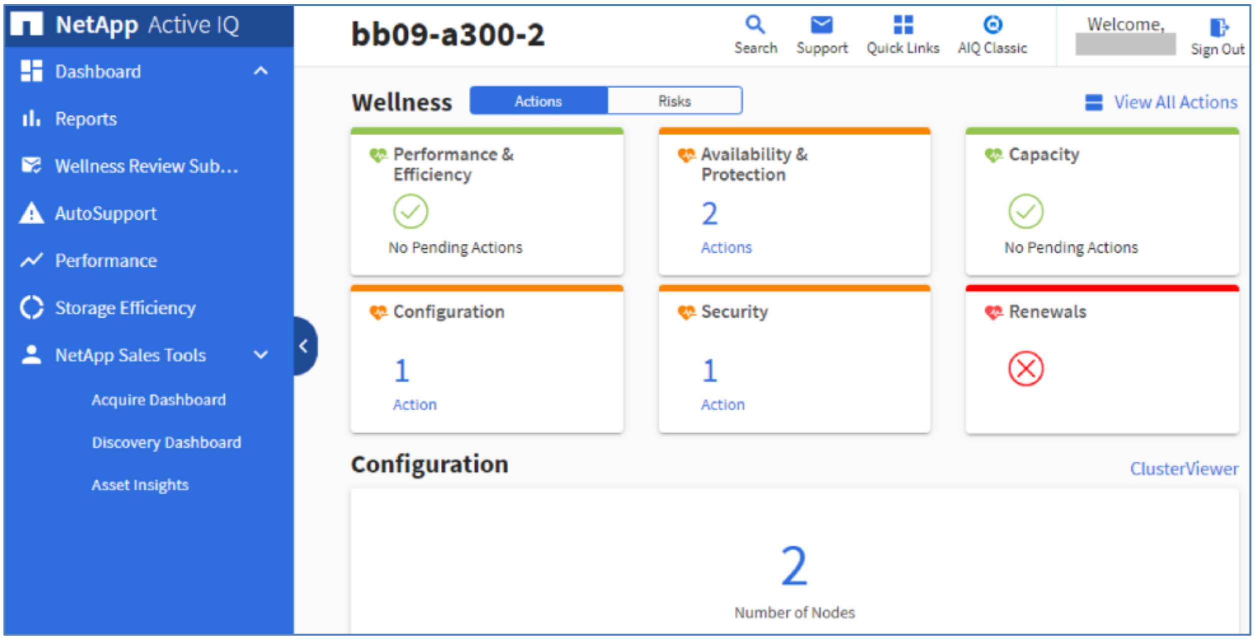

Create Active IQ Digital Advisor Dashboard

The system level dashboard is the default view for systems in Active IQ. To create a dashboard, follow these steps:

1. On the Create Dashboard page, click here to create a dashboard.

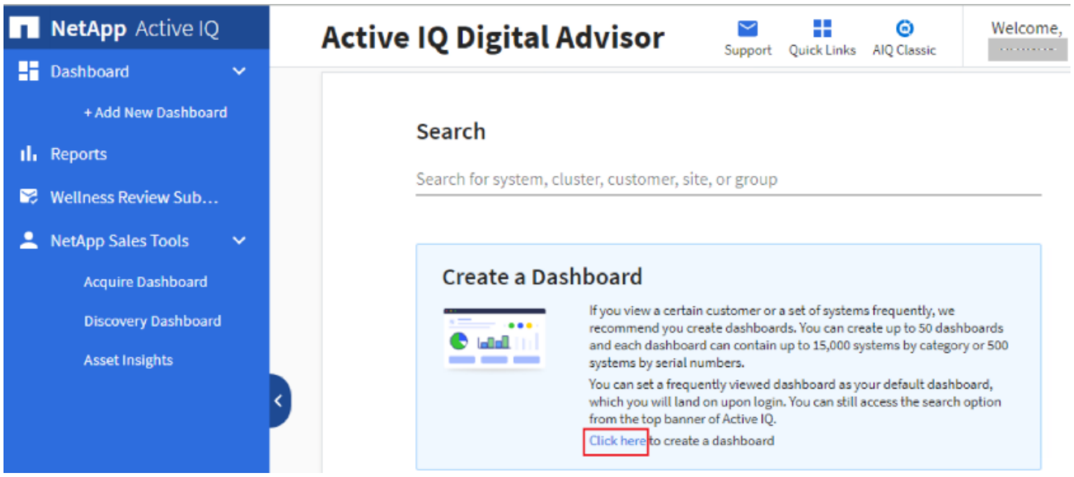

[redacted]

Search

Search for system, cluster, customer, site, or group

Create a Dashboard

If you view a certain customer or a set of systems frequently, we recommend you create dashboards. You can create up to 50 dashboards and each dashboard can contain up to 15,000 systems by category or 500 systems by serial numbers.

You can set a frequently viewed dashboard as your default dashboard, which you will land on upon login. You can still access the search option from the top banner of Active IQ.

Click here to create a dashboard

2. Click Create Watchlist and enter a name for the watchlist.

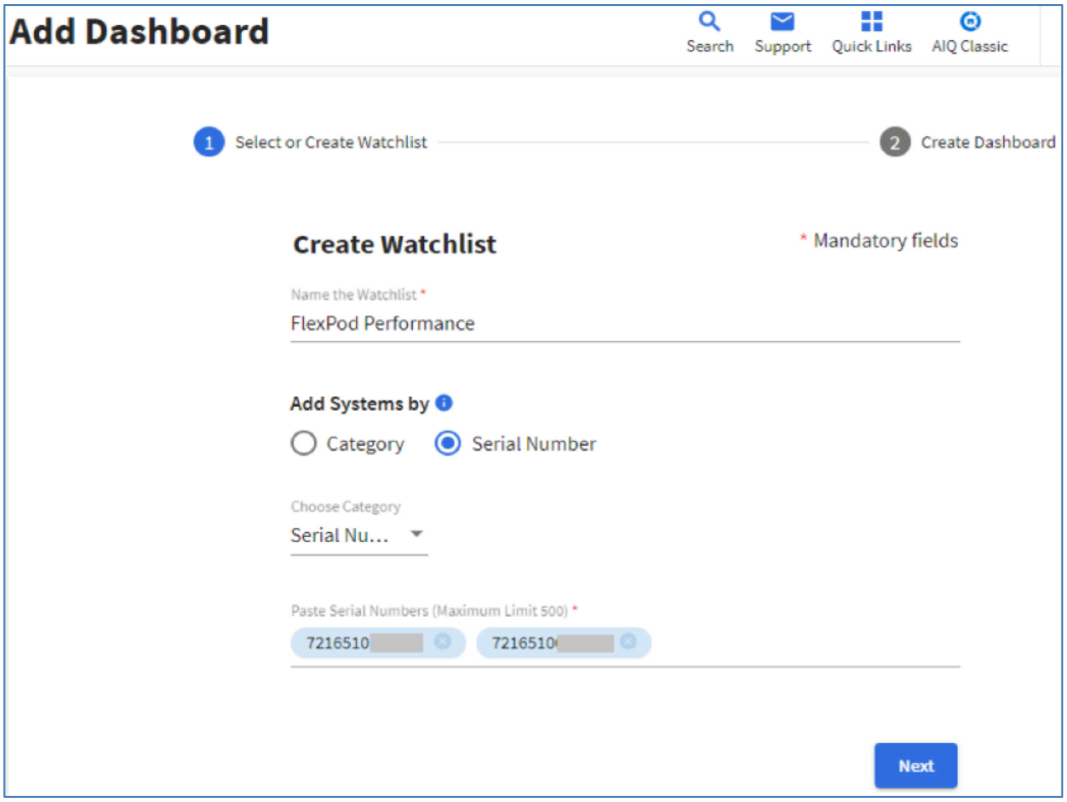

Create Watchlist

* Mandatory fields

Name the Watchlist *

FlexPod Performance

Add Systems by

○ Category ● Serial Number

Choose Category

Serial Nu…

Paste Serial Numbers (Maximum Limit 500) *

7216510[redacted] 7216510[redacted]

3. Choose the radio button to add systems by serial number and enter the cluster serial numbers to the watchlist.

4. Check the box for Make this my default dashboard and click Create.

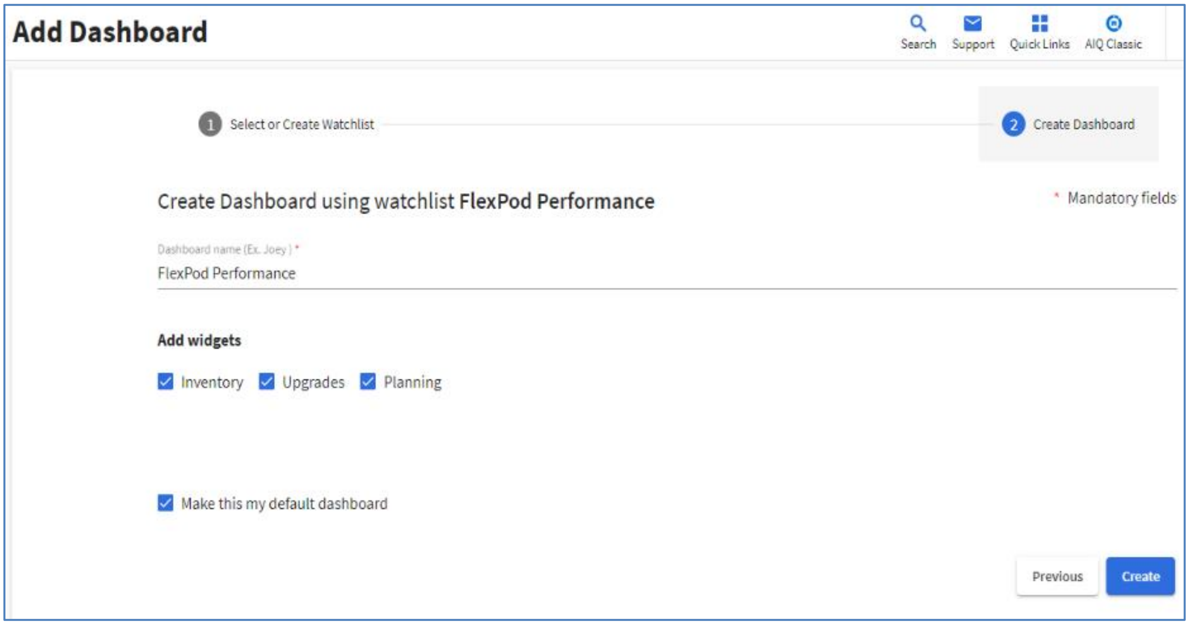

Create Dashboard using watchlist **FlexPod Performance**

* Mandatory fields

Dashboard name (Ex. Joey) *

FlexPod Performance

Add widgets

☑ Inventory ☑ Upgrades ☑ Planning

☑ Make this my default dashboard

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

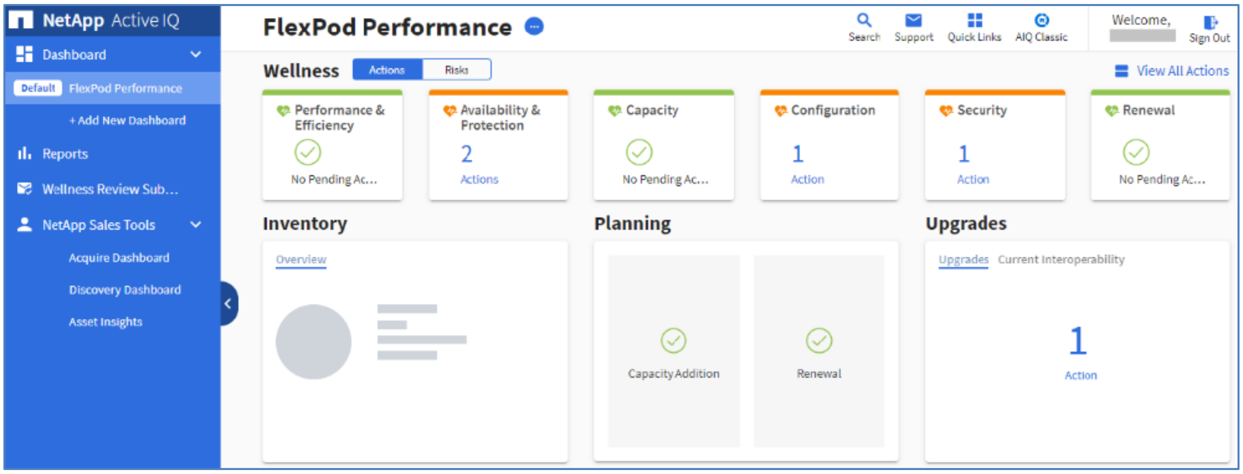

6. Switch between the Actions and Risks tabs to view the risks broken down by category or a list of all risks with their impact and links to corrective actions.

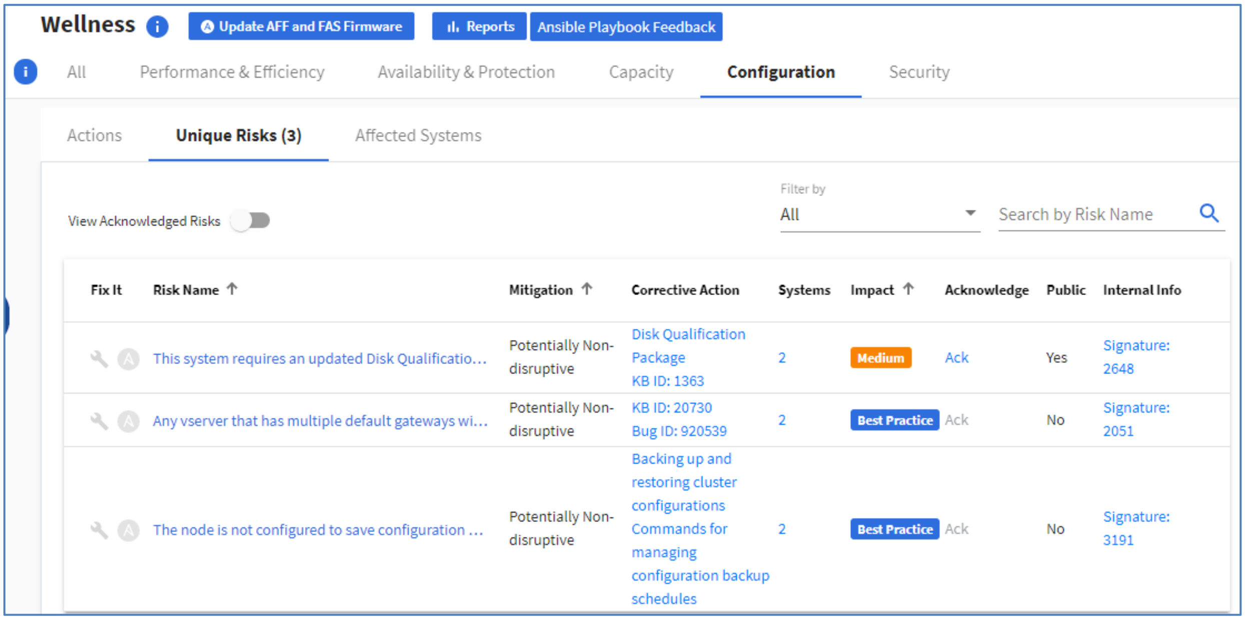

Fix It	Risk Name	Mitigation	Corrective Action	Systems	Impact	Acknowledge	Public	Internal Info
	This system requires an updated Disk Qualificatio…	Potentially Non-disruptive	Disk Qualification Package KB ID: 1363	2	Medium	Ack	Yes	Signature: 2648
	Any vserver that has multiple default gateways wi…	Potentially Non-disruptive	KB ID: 20730 Bug ID: 920539	2	Best Practice	Ack	No	Signature: 2051
	The node is not configured to save configuration …	Potentially Non-disruptive	Backing up and restoring cluster configurations Commands for managing configuration backup schedules	2	Best Practice	Ack	No	Signature: 3191

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

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

Cisco Intersight

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

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

To configure Cisco Intersight, follow these steps:

1. If you do not already have a Cisco Intersight account, to claim your Cisco UCS system into a new account on Cisco Intersight, connect to https://Intersight.com.

2. Click Create an account.

3. Click Continue. Complete the Sign in process with your Cisco ID.

4. Read the Offer Description carefully and accept it. Click Next.

5. Enter an Account Name, Device ID, and Claim Code. The Device ID and Claim Code can be obtained by connecting to Cisco UCS Manager and selecting Admin > All > Device Connector. The Device ID and Claim Code are on the right. Click Claim.

6. Click Create. After the account has been created, click Log me in to log into Cisco Intersight.

7. To claim your Cisco UCS system into an existing Intersight account, log into the account at https://Intersight.com. Choose Administration > Devices. Click Claim a New Device. Under Direct Claim, fill in the Device ID and Claim Code. The Device ID and Claim Code can be obtained by connecting to Cisco UCS Manager and selecting Admin > All > Device Connector. The Device ID and Claim Code are on the right.

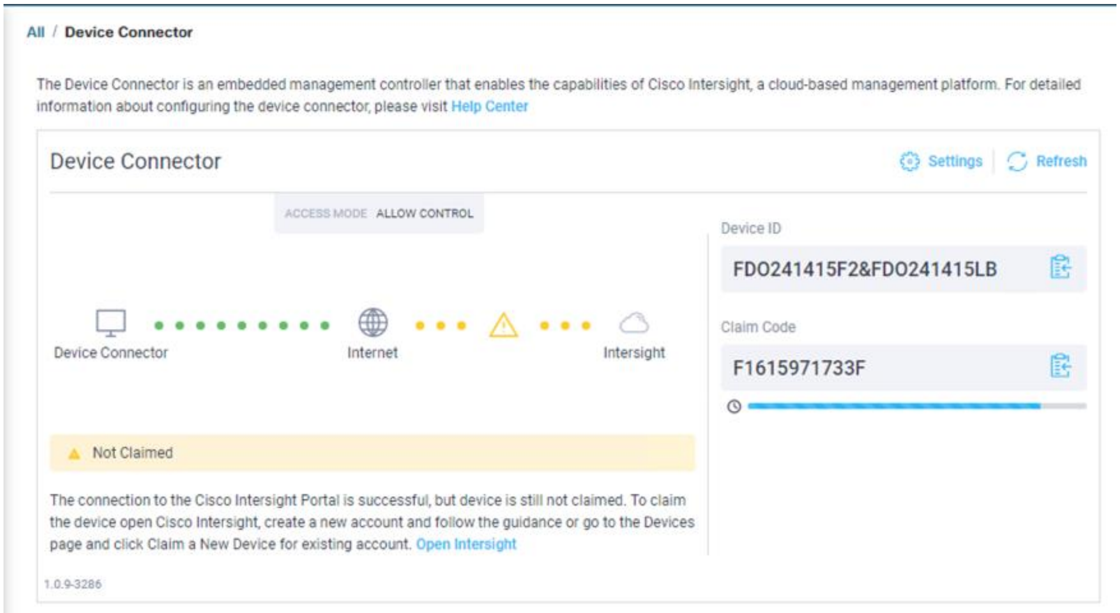

All / Device Connector

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

Device Connector

Device ID

FDO241415F2&FDO241415LB

Claim Code

F1615971733F

Not Claimed

The connection to the Cisco Intersight Portal is successful, but device is still not claimed. To claim the device open Cisco Intersight, create a new account and follow the guidance or go to the Devices page and click Claim a New Device for existing account. Open Intersight

1.0.9-3286

8. From the Cisco Intersight window, click ⚙ and then click Licensing. If this is a new account, all servers connected to the UCS Domain will appear under the Base license Tier. If you have purchased Cisco Intersight licenses and have them in your Cisco Smart Account, click Register and follow the prompts to register this Cisco Intersight account to your Cisco Smart Account. Cisco Intersight also offers a one-time 90-day trial of Premier licensing for new accounts. Click Start Trial and then Start to begin this evaluation. The remainder of this section will assume Premier licensing.

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

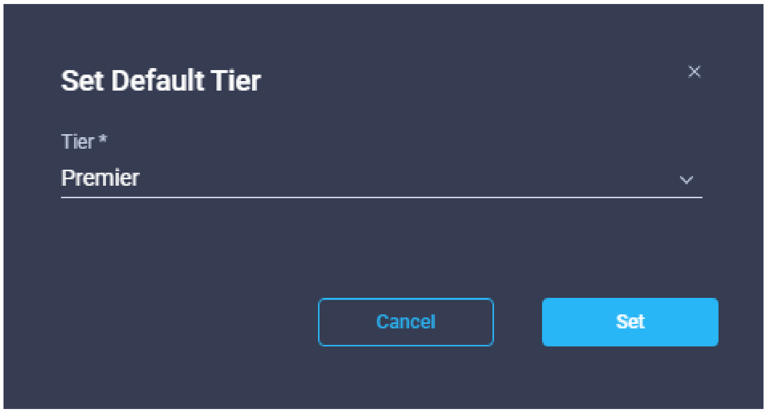

Set Default Tier

Tier *
Premier

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

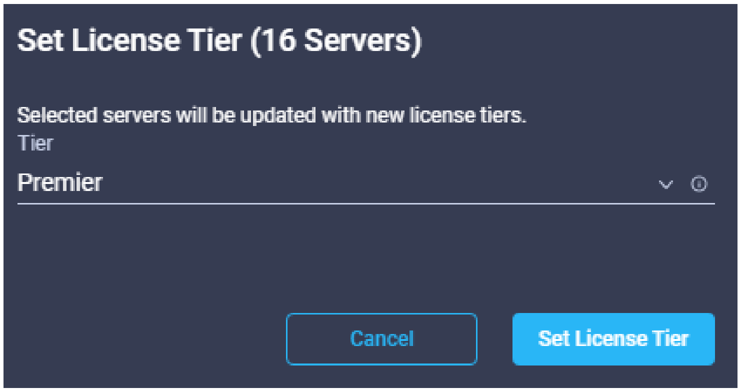

Set License Tier (16 Servers)

Selected servers will be updated with new license tiers.

Tier
Premier

11. Click Refresh to refresh the Intersight window with Premier, Advantage, and Essentials features added.

12. Click ? in the Intersight window and click Take a Site Tour. Follow the prompts for a tour of Cisco Intersight.

13. The Essentials tier of Cisco Intersight includes a Cisco driver check against the Cisco Hardware Compatibility List (HCL). In the Servers list, choose one of the servers in your VMware FlexPod-Management cluster by clicking the server name. Review the detailed General and Inventory information for the server. Click the HCL tab. Review the server information, the version of VMware ESXi, and the Cisco VIC driver versions and RAID card if present.

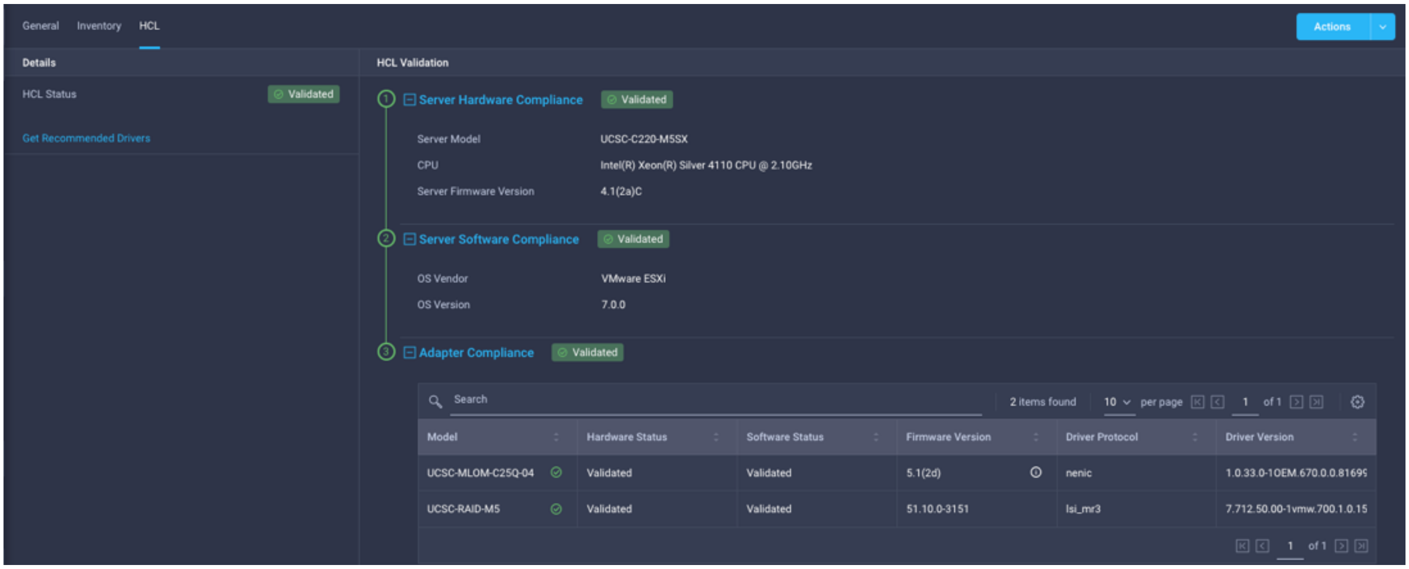

Details

HCL Status ✓ Validated

Get Recommended Drivers

HCL Validation

1 Server Hardware Compliance ✓ Validated

Server Model	UCSC-C220-M5SX
CPU	Intel(R) Xeon(R) Silver 4110 CPU @ 2.10GHz
Server Firmware Version	4.1(2a)C

2 Server Software Compliance ✓ Validated

OS Vendor	VMware ESXi
OS Version	7.0.0

3 Adapter Compliance ✓ Validated

Search 2 items found 10 per page 1 of 1

Model	Hardware Status	Software Status	Firmware Version	Driver Protocol	Driver Version
UCSC-MLOM-C25Q-04 ✓	Validated	Validated	5.1(2d)	nenic	1.0.33.0-1OEM.670.0.0.81699
UCSC-RAID-M5 ✓	Validated	Validated	51.10.0-3151	lsi_mr3	7.712.50.00-1vmw.700.1.0.15

1 of 1

14. Using the Intersight Assist personality of the Cisco Intersight Virtual Appliance VMware vCenter currently can be monitored (Advantage Licensing Tier) and configured (Premier Licensing Tier Tech Preview not to be used in production environments). To install Intersight Assist from an Open Virtual Appliance (OVA) in your VMware FlexPod-Management Cluster, first download the latest release of the OVA from https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-148.

15. Refer to https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/cisco-intersight-assist-getting-started-guide/m-installing-cisco-intersight-assist.html and set up the DNS entries for the Intersight Assist hostname as specified under Before you begin.

16. From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the FlexPod-Management cluster and click Deploy OVF Template.

17. Specify a URL or either browse to the intersight-virtual-appliance-1.0.9-148.ova or latest release file. Click NEXT.

18. Name the Intersight Assist VM and choose the location. Click NEXT.

19. Choose the FlexPod-Management cluster and click NEXT.

20. Review details and click NEXT.

21. Choose a deployment configuration (Tiny recommended) and click NEXT.

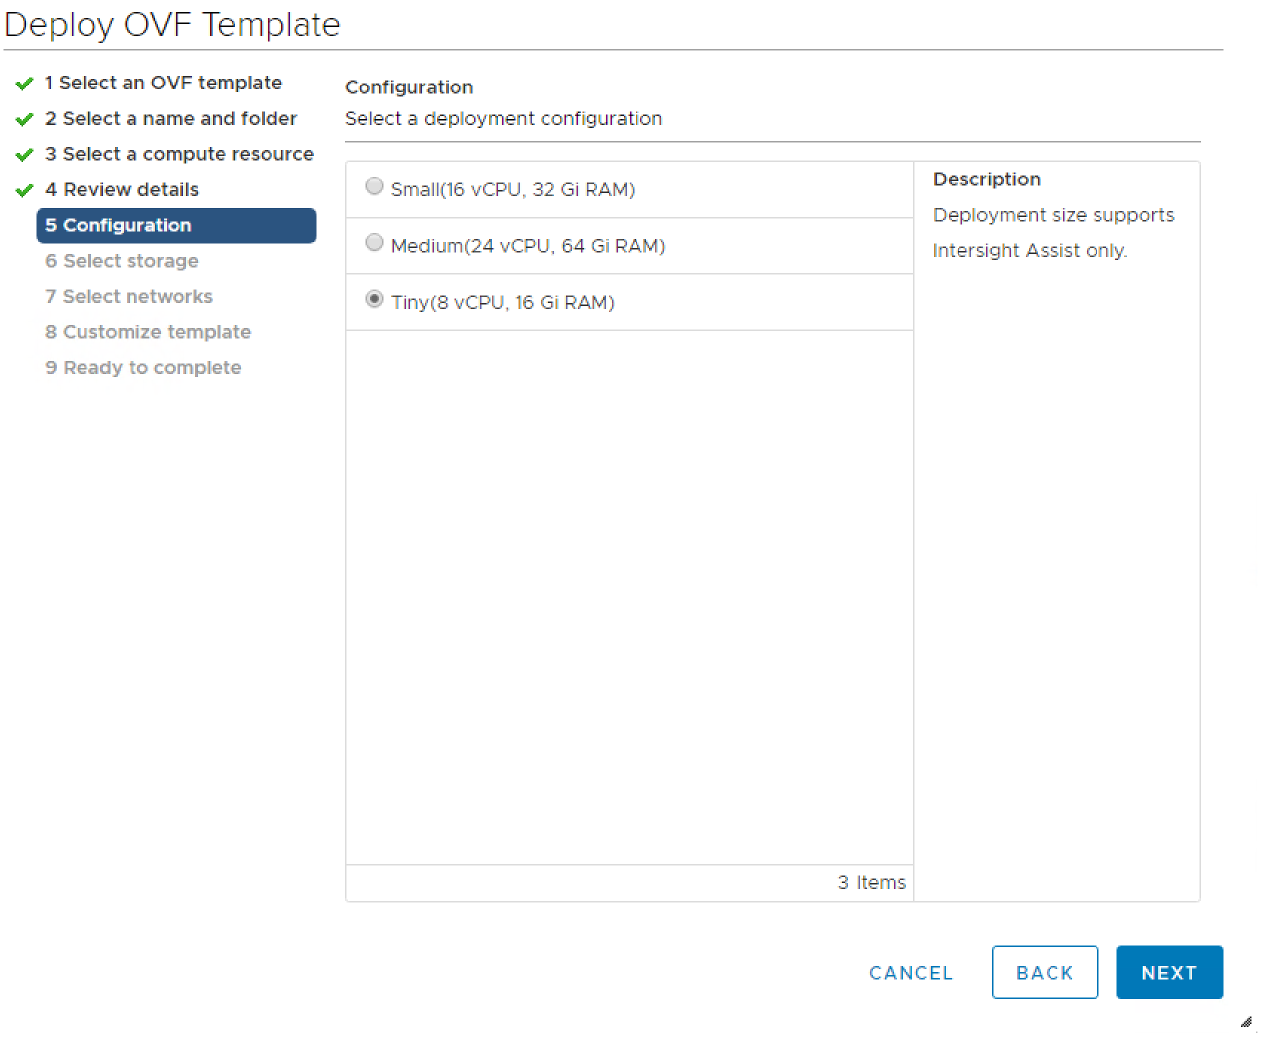

Deploy OVF Template

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

Configuration

Select a deployment configuration

○ Small(16 vCPU, 32 Gi RAM)

○ Medium(24 vCPU, 64 Gi RAM)

● Tiny(8 vCPU, 16 Gi RAM)

3 Items

Description

Deployment size supports Intersight Assist only.

22. Choose infra_datastore_1 for storage and choose the Thin Provision virtual disk format. Click NEXT.

23. Choose Common-Services for the VM Network. Click NEXT.

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

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

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

27. Expand CPU and adjust the Cores per Socket so that the number of Sockets matches your server CPU configuration. In this example 2 Sockets are shown. Click OK.

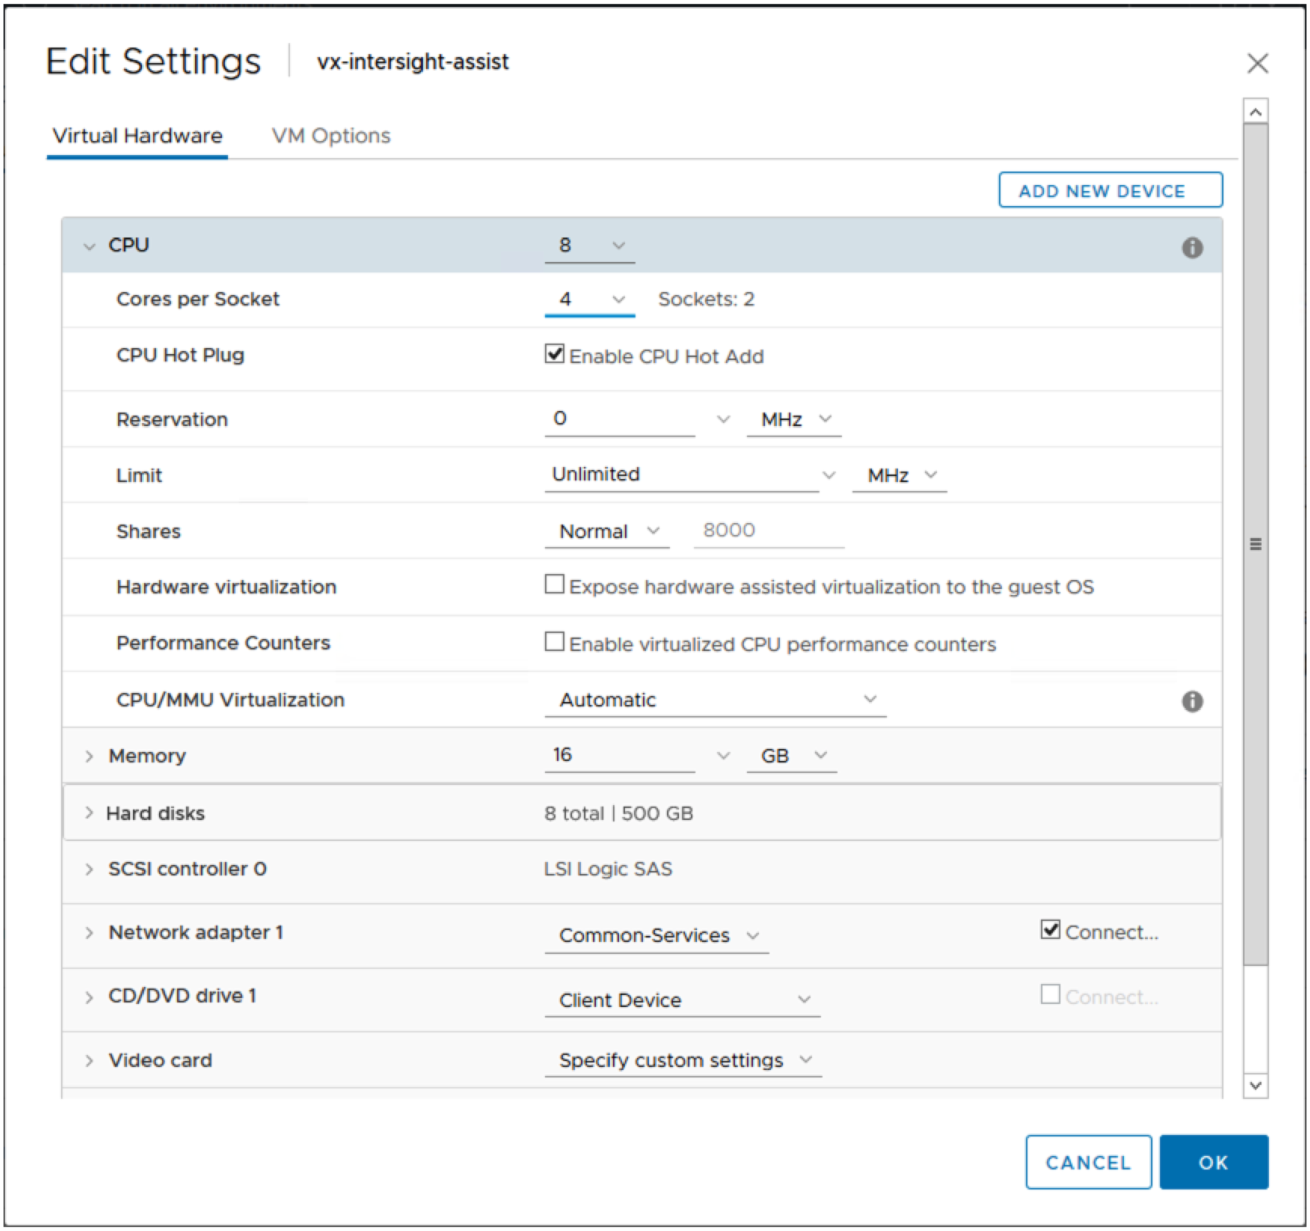

Edit Settings | vx-intersight-assist

Virtual Hardware | VM Options

ADD NEW DEVICE

CPU	8	
Cores per Socket	4	Sockets: 2
CPU Hot Plug	☑ Enable CPU Hot Add	
Reservation	0	MHz
Limit	Unlimited	MHz
Shares	Normal	8000
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS	
Performance Counters	□ Enable virtualized CPU performance counters	
CPU/MMU Virtualization	Automatic	
Memory	16	GB
Hard disks	8 total | 500 GB	
SCSI controller 0	LSI Logic SAS	
Network adapter 1	Common-Services	☑ Connect...
CD/DVD drive 1	Client Device	□ Connect...
Video card	Specify custom settings	

CANCEL | OK

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

29. Click ▶ to power on the VM.

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

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

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

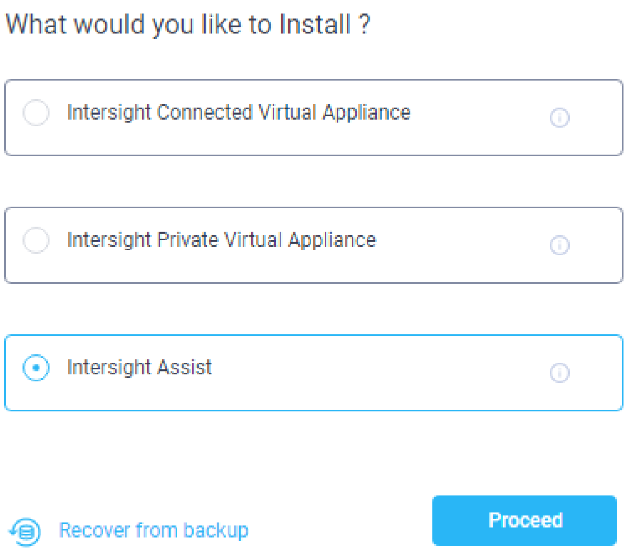

What would you like to Install ?

○ Intersight Connected Virtual Appliance

○ Intersight Private Virtual Appliance

● Intersight Assist

32. From Cisco Intersight, click ADMIN > Devices. Click Claim a New Device. Copy and paste the Device ID and Claim Code shown in the Intersight Assist web interface to the Cisco Intersight Device Claim Direct Claim window. In Cisco Intersight, click Claim. Intersight Assist will now appear as a claimed device.

33. In the Intersight Assist web interface, reload if necessary to reflect the connection, then click Continue.

34. The Intersight Assist software will now be downloaded and installed into the Intersight Assist VM. This can take up to an hour to complete.

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

35. When the software download is complete, an Intersight Assist login screen will appear. Log into Intersight Assist with the admin@local user and the password supplied in the OVA installation. Check the Intersight Assist status and log out of Intersight Assist.

36. To claim the vCenter, from Cisco Intersight, click ADMIN > Devices. Click Claim a New Device. In the Device Claim window, choose Claim Through Intersight Assist. Fill in the vCenter information and click Claim.

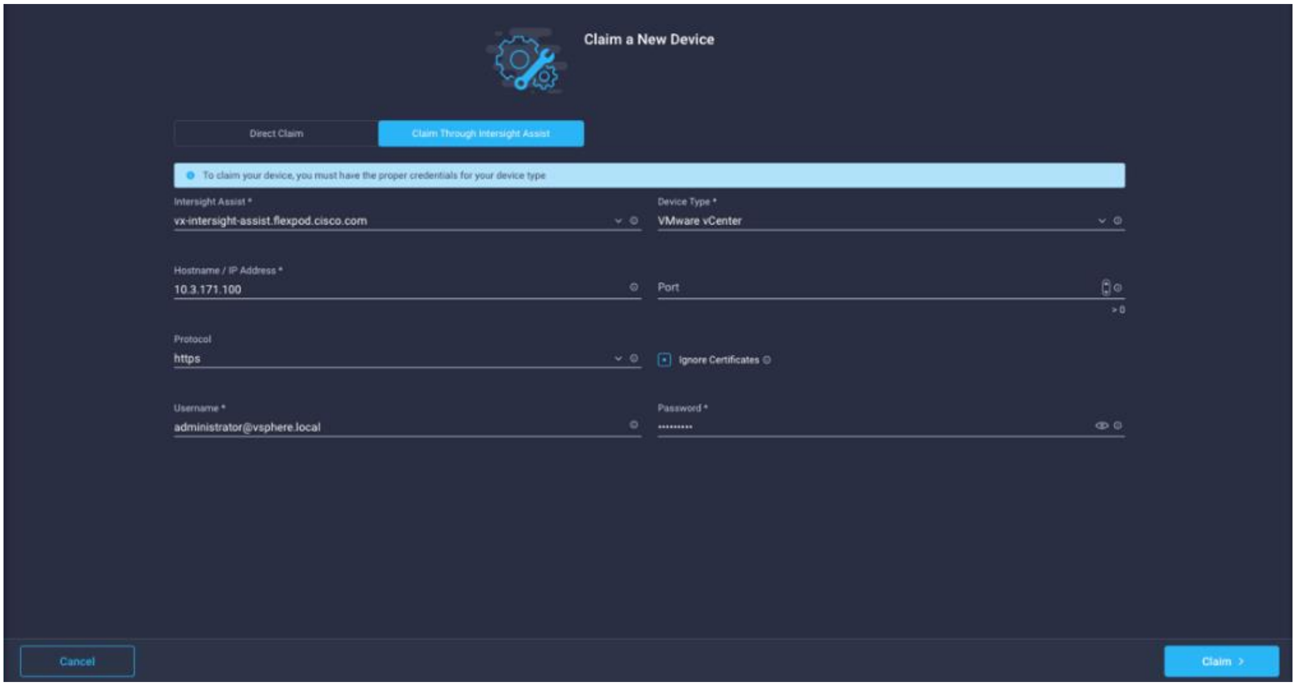

Claim a New Device

To claim your device, you must have the proper credentials for your device type

Intersight Assist *: vx-intersight-assist.flexpod.cisco.com

Device Type *: VMware vCenter

Hostname / IP Address *: 10.3.171.100

Port:

Protocol: https

☑ Ignore Certificates

Username *: administrator@vsphere.local

Password *: •••••••••

37. After a few minutes, the VMware vCenter will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

38. Detailed information obtained from the vCenter can now be viewed by clicking Virtualization from the menu.

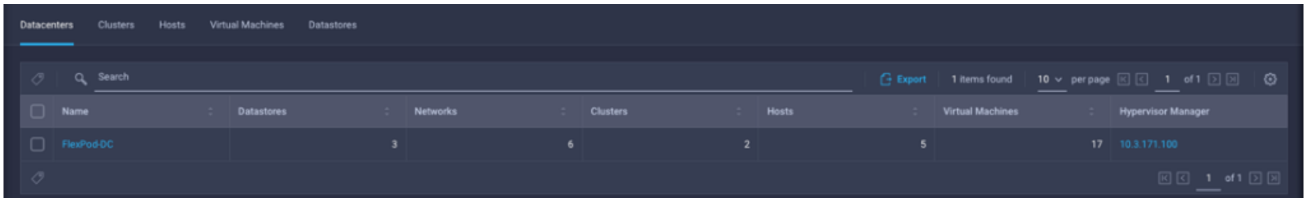

Name	Datastores	Networks	Clusters	Hosts	Virtual Machines	Hypervisor Manager
FlexPod-DC	3	6	2	5	17	10.3.171.100

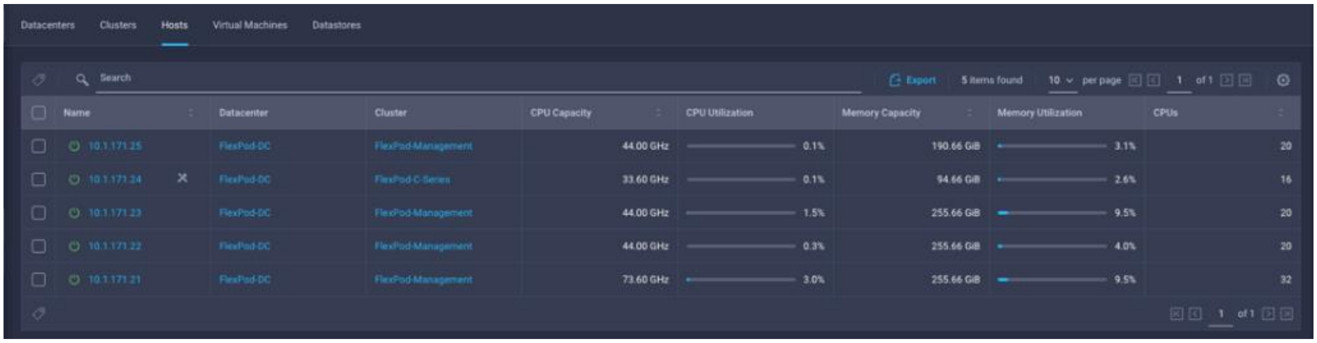

Name	Datacenter	Cluster	CPU Capacity	CPU Utilization	Memory Capacity	Memory Utilization	CPUs
10.1.171.25	FlexPod-DC	FlexPod-Management	44.00 GHz	0.1%	190.66 GiB	3.1%	20
10.1.171.24	FlexPod-DC	FlexPod-C-Series	33.60 GHz	0.1%	94.66 GiB	2.6%	16
10.1.171.23	FlexPod-DC	FlexPod-Management	44.00 GHz	1.5%	255.66 GiB	9.5%	20
10.1.171.22	FlexPod-DC	FlexPod-Management	44.00 GHz	0.3%	255.66 GiB	4.0%	20
10.1.171.21	FlexPod-DC	FlexPod-Management	73.60 GHz	3.0%	255.66 GiB	9.5%	32

Solution Deployment - Sample Tenant Provisioning

Provision a Sample Application Tenant

This section describes a sample procedure for provisioning an application tenant. The procedure here refers to previous sections of this document and can be used as a guide and modified as needed when provisioning an application tenant.

Plan your application tenant and determine what storage protocols will be provided in the tenant. In the architecture covered in this document, NFS, iSCSI, and CIFS/SMB can be provided to the tenant. Also, plan what network VLANs the tenant will use. It is recommended to have a VLAN for SVM management traffic. The tenant application VLAN can be used for SVM management. One or two VLANs (iSCSI needs two if VMware RDM LUNs or iSCSI datastores will be provisioned) are also needed for each storage protocol used. If the infrastructure NFS VLAN will be used in the tenant, consider migrating the infrastructure NFS VMkernel port on each host to the vDS to take advantage of Ethernet adapter policy queuing.

In the DCNM, create the necessary application and application storage VLANs and enable them on the access-layer vPC connections to Cisco UCS and NetApp storage. Also, for some networks, the fabric can be used for strictly Layer 2 forwarding (for example, storage networks) and for others the VXLAN fabric can serve as the default gateway to reach other networks connected to the same shared fabric or for connectivity outside the fabric. See Solution Deployment – Network section for more details on provisioning the network fabric to support the Application Tenant and associated networks.

Once the fabric is provisioned, configure the VLANs on Cisco UCS and NetApp storage systems and enable them on the uplinks to the VXLAN fabric.

In the storage cluster:

- Create a broadcast domain with MTU 1500 for the tenant SVM management interface. Create a broadcast domain with MTU 9000 for each tenant storage protocol.
- Create VLAN interface ports on the node interface group on each node for tenant SVM management (VM VLAN) and for the VLAN for each storage protocol. Add these VLAN ports to the appropriate broadcast domains.
- Create the tenant SVM and follow all procedures in that section.
- Create Load-Sharing Mirrors for the tenant SVM.
- Create the iSCSI service for the tenant SVM if iSCSI is being deployed in this tenant.
- Optionally, create a self-signed security certificate for the tenant SVM.
- Configure NFSv3 for the tenant SVM.
- Create a VM datastore volume in the tenant SVM.
- For iSCSI configure four iSCSI LIFs in the tenant SVM on the iSCSI VLAN interfaces.
- Create one NFS LIF in the tenant SVM on each storage node.
- Create a boot LUN in the esxi_boot volume in the Infra-SVM for each tenant VMware ESXi host.

- Add the tenant SVM Administrator, SVM management LIF on the SVM management VLAN port, and default route for the SVM.

In Cisco UCS, one method of tenant setup is to dedicate a VMware ESXi cluster and set of UCS servers to each tenant. Service profiles will be generated for at least two tenant ESXi hosts. These hosts can boot from LUNs from the esxi_boot volume in the Infra-SVM but will also have access to iSCSI storage in the tenant SVM.

- Create a Server Pool for the tenant ESXi host servers.
- Create all tenant VLANs in the LAN Cloud.
- Add the tenant VLANs to the vDS vNIC templates.
- Generate service profiles from the service profile template with the vMedia policy for the tenant ESXi hosts. Remember to bind these service profiles to the service profile template without the vMedia policy after VMware ESXi installation.

In the storage cluster:

- Create igroups for the tenant ESXi hosts in both the Infra-SVM and tenant SVM. Also, create an igroup in the tenant SVM that includes the IQNs for all tenant ESXi hosts to support shared storage from the tenant SVM.
- In Infra-SVM, map the boot LUNs created earlier to the tenant ESXi hosts. Tenant iSCSI storage can be created later using NetApp VSC.
- Install and configure VMware ESXi on all tenant host servers. Optionally, if needed then infra_datastore_1/_2 datastores can be mapped to the tenant hosts.
- In VMware vCenter, create a cluster for the tenant ESXi hosts. Add the hosts to the cluster.
- Using the vCenter HTML5 Client, add the tenant hosts to vDS0 or create a tenant vDS (using the vDS0 vNICs) and add the hosts to it. In vDS0 or the tenant vDS, add port-profiles for the tenant VLANs.
- Back in vCenter, add in any necessary VMkernel ports for storage interfaces remembering to set the MTU correctly on these interfaces. Mount the tenant NFS datastore on the tenant cluster if one was created. Tenant iSCSI VMkernel ports can be created on the vDS with the port groups pinned to the appropriate fabric. Add the tenant iSCSI LIF IP addresses as Dynamic Targets on the VMware ESXi hosts in the vCenter HTML5 Client.
- Using the NetApp VSC plugin to the vCenter HTML5 Client, set recommended values for all tenant ESXi hosts. Ensure the NetApp NFS Plug-in for VMware VAAI is installed on all tenant hosts and reboot each host.

You can now begin provisioning virtual machines on the tenant cluster. The NetApp VSC plugin can be used to provision iSCSI and NFS datastores. Optionally, use NetApp SnapCenter to provision backups of tenant virtual machines.

Appendix

The leaf and spine configurations deployed in the environment by DCNM are provided below. The FlexPod compute and storage infrastructure connect to the Leaf switches included in this section. The configuration for the optional Border Leaf switches are not included here.

Leaf A

```
version 9.3(5) Bios:version 05.42
switchname AA01-9336C-FX2-1
vdc AA01-9336C-FX2-1 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource vni_bd minimum 4096 maximum 4096

feature nxapi
cfs ipv4 distribute
nv overlay evpn
feature ospf
feature bgp
feature pim
feature interface-vlan
feature vn-segment-vlan-based
feature lacp
feature dhcp
feature vpc
feature lldp
feature nv overlay
feature ngoam

username admin password 5 $5$LL7Z.BYy$U6AAWL6OvWLdvJ6SB5Hul.UCZ2I5EskdSKoubMKKvm
1  role network-admin
ip domain-lookup
copp profile strict
configure profile FPV-Foundation_VRF
  vlan 3500
    name FPV_Foundation_VRF_VLAN
    vn-segment 30000
  interface Vlan3500
    description FPV_Foundation_VRF_Interface
    vrf member fpv-foundation_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-foundation_vrf
    description FPV_Foundation_VRF
    vni 30000
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  router bgp 65001
    vrf fpv-foundation_vrf
      address-family ipv4 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
      address-family ipv6 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
  interface nve1
    member vni 30000 associate-vrf
configure terminal
configure profile FPV-iSCSI-A_Network
  vlan 3010
    vn-segment 20000
    name FPV-iSCSI-A_VLAN
  interface nve1
    member vni 20000
      mcast-group 239.1.1.0
  evpn
    vni 20000 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-iSCSI-B_Network
  vlan 3020
    vn-segment 20001
    name FPV-iSCSI-B_VLAN
  interface nve1
    member vni 20001
      mcast-group 239.1.1.0
  evpn
    vni 20001 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-InfraNFS_Network
  vlan 3050
    vn-segment 20002
    name FPV-InfraNFS_VLAN
  interface nve1
    member vni 20002
      mcast-group 239.1.1.0
  evpn
    vni 20002 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-InBand-SiteA_Network
  vlan 122
    vn-segment 20003
    name FPV-InBand-SiteA_VLAN
  interface Vlan122
    description FPV-InBand-SiteA_Interface
    vrf member fpv-foundation_vrf
    no ip redirects
    no ipv6 redirects
    ip address 10.1.171.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 20003
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 20003 l2
      rd auto
      route-target import auto
      route-target export auto
```

```
configure terminal
configure profile FPV-vMotion_Network
  vlan 3000
    vn-segment 20004
    name FPV-vMotion_VLAN
  interface nve1
    member vni 20004
      mcast-group 239.1.1.0
  evpn
    vni 20004 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-CommonServices_Network
  vlan 322
    vn-segment 20005
    name FPV-CommonServices_VLAN
  interface Vlan322
    description FPV-CommonServices_Interface
    vrf member fpv-foundation_vrf
    no ip redirects
    no ipv6 redirects
    ip address 10.3.171.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 20005
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 20005 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-Application_VRF
  vlan 3501
    name FPV-Application_VRF_VLAN
    vn-segment 30001
  interface Vlan3501
    description FPV-Application_VRF_Interface
    vrf member fpv-application_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-application_vrf
    description FPV-Application_VRF
    vni 30001
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  router bgp 65001
    vrf fpv-application_vrf
      address-family ipv4 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
      address-family ipv6 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
  interface nve1
    member vni 30001 associate-vrf
configure terminal
configure profile FPV-App-1_Network
  vlan 1001
    vn-segment 21001
    name FPV-App-1_VLAN
  interface Vlan1001
    description FPV-App-1_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.1.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21001
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21001 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-App-2_Network
  vlan 1002
    vn-segment 21002
    name FPV-App-2_VLAN
  interface Vlan1002
    description FPV-App-2_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.2.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21002
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21002 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-App-3_Network
  vlan 1003
    vn-segment 21003
    name FPV-App-3_VLAN
  interface Vlan1003
    description FPV-App-3_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.3.254/24 tag 12345
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21003
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21003 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
snmp-server user admin network-admin auth md5 0xba043903263cde0a9b23f130b0aabfa0
 priv 0xba043903263cde0a9b23f130b0aabfa0 localizedkey
snmp-server host 172.26.163.142 traps version 2c public udp-port 2162
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 172.26.163.254 use-vrf management
ntp server 172.26.164.254 use-vrf management

fabric forwarding anycast-gateway-mac 2020.0000.00aa
ipv6 switch-packets lla
ip pim rp-address 10.11.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
vlan 1,122,322,1001-1003,3000,3010,3020,3050,3500-3501

route-map fabric-rmap-redist-subnet permit 10
  match tag 12345
service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 172.26.163.254
hardware access-list tcam region ing-racl 1792
hardware access-list tcam region ing-flow-redirect 512
vpc domain 1
  peer-switch
  peer-keepalive destination 172.26.163.224 source 172.26.163.223
  virtual peer-link destination 10.11.0.4 source 10.11.0.3 dscp 56
  delay restore 150
  peer-gateway
  auto-recovery reload-delay 360
  ipv6 nd synchronize
  ip arp synchronize
ngoam install acl

nxapi http port 80


interface Vlan1
  no ip redirects
  no ipv6 redirects

interface port-channel1
  description To FXV-AA01-UCS6454FI-A: e1/53
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,1001-1003,3000,3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 1

interface port-channel2
  description To FXV-AA01-UCS6454FI-B: e1/53
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,1001-1003,3000,3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 2
```

```
interface port-channel3
  description To FXV-BB09-A300-2-01: e2a
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,1001-1003,3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 3

interface port-channel4
  description To FXV-BB09-A300-2-02: e2a
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,1001-1003,3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 4

interface port-channel500
  description "vpc-peer-link"
  switchport
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link

interface nve1
  no shutdown
  host-reachability protocol bgp
  advertise virtual-rmac
  source-interface loopback1

interface Ethernet1/1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,1001-1003,3000,3010,3020,3050
  mtu 9216
  channel-group 1 mode active
  no shutdown

interface Ethernet1/2
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,1001-1003,3000,3010,3020,3050
  mtu 9216
  channel-group 2 mode active
  no shutdown

interface Ethernet1/3
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/4
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/5
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,1001-1003,3010,3020,3050
  mtu 9216
  channel-group 3 mode active
  no shutdown

interface Ethernet1/6
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,1001-1003,3010,3020,3050
  mtu 9216
  channel-group 4 mode active
  no shutdown

interface Ethernet1/7
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/8
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/9
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/10
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/13
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/14
```

```
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/15
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/16
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/17
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/18
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/19
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/20
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/21
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/22
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/23
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/24
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/25
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/26
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/27
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/28
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/29
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/30
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/31
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/32
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/33
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/34
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/35
  description connected-to-AA01-9364C-1-Ethernet1/3
  mtu 9216
  port-type fabric
  ip address 10.11.0.25/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/36
  description connected-to-AA01-9364C-2-Ethernet1/3
  mtu 9216
  port-type fabric
  ip address 10.11.0.33/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface mgmt0
  vrf member management
  ip address 172.26.163.223/24

interface loopback0
  description Routing loopback interface
  ip address 10.11.0.3/32
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode

interface loopback1
  description VTEP loopback interface
  ip address 10.11.1.3/32
  ip address 10.11.1.5/32 secondary
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
icam monitor scale

line console
line vty
```

```
boot nxos bootflash:/nxos.9.3.5.bin
router ospf Site-A_UNDERLAY
  router-id 10.11.0.3
router bgp 65001
  router-id 10.11.0.3
  address-family l2vpn evpn
    advertise-pip
  neighbor 10.11.0.1
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
  neighbor 10.11.0.2
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended

apply profile FPV-Foundation_VRF
apply profile FPV-iSCSI-A_Network
apply profile FPV-iSCSI-B_Network
apply profile FPV-InfraNFS_Network
apply profile FPV-InBand-SiteA_Network
apply profile FPV-vMotion_Network
apply profile FPV-CommonServices_Network
apply profile FPV-Application_VRF
apply profile FPV-App-1_Network
apply profile FPV-App-2_Network
apply profile FPV-App-3_Network
```

Leaf B

```
version 9.3(5) Bios:version 05.42
switchname AA01-9336C-FX2-2
vdc AA01-9336C-FX2-2 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource vni_bd minimum 4096 maximum 4096

feature nxapi
cfs ipv4 distribute
nv overlay evpn
feature ospf
feature bgp
feature pim
feature interface-vlan
feature vn-segment-vlan-based
feature lacp
feature dhcp
feature vpc
feature lldp
feature nv overlay
feature ngoam

username admin password 5 $5$kHQwAaVM$Ehn5K/FQeT68soExah6QdS9T7F71QpEEqmdmf7pdiU
C  role network-admin
ip domain-lookup
copp profile strict
configure profile FPV-Foundation_VRF
  vlan 3500
    name FPV_Foundation_VRF_VLAN
    vn-segment 30000
  interface Vlan3500
    description FPV_Foundation_VRF_Interface
    vrf member fpv-foundation_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-foundation_vrf
    description FPV_Foundation_VRF
    vni 30000
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  router bgp 65001
    vrf fpv-foundation_vrf
      address-family ipv4 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
      address-family ipv6 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
  interface nve1
    member vni 30000 associate-vrf
configure terminal
configure profile FPV-iSCSI-A_Network
  vlan 3010
    vn-segment 20000
    name FPV-iSCSI-A_VLAN
  interface nve1
    member vni 20000
      mcast-group 239.1.1.0
  evpn
    vni 20000 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-iSCSI-B_Network
  vlan 3020
    vn-segment 20001
    name FPV-iSCSI-B_VLAN
  interface nve1
    member vni 20001
      mcast-group 239.1.1.0
  evpn
    vni 20001 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-InfraNFS_Network
  vlan 3050
    vn-segment 20002
    name FPV-InfraNFS_VLAN
  interface nve1
    member vni 20002
      mcast-group 239.1.1.0
  evpn
    vni 20002 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-InBand-SiteA_Network
  vlan 122
    vn-segment 20003
    name FPV-InBand-SiteA_VLAN
  interface Vlan122
    description FPV-InBand-SiteA_Interface
    vrf member fpv-foundation_vrf
    no ip redirects
    no ipv6 redirects
    ip address 10.1.171.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 20003
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 20003 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-vMotion_Network
  vlan 3000
    vn-segment 20004
    name FPV-vMotion_VLAN
  interface nve1
    member vni 20004
      mcast-group 239.1.1.0
  evpn
    vni 20004 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-CommonServices_Network
  vlan 322
    vn-segment 20005
    name FPV-CommonServices_VLAN
  interface Vlan322
    description FPV-CommonServices_Interface
    vrf member fpv-foundation_vrf
    no ip redirects
    no ipv6 redirects
    ip address 10.3.171.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 20005
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 20005 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-Application_VRF
  vlan 3501
    name FPV-Application_VRF_VLAN
    vn-segment 30001
  interface Vlan3501
    description FPV-Application_VRF_Interface
    vrf member fpv-application_vrf
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context fpv-application_vrf
    description FPV-Application_VRF
    vni 30001
    rd auto
    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  router bgp 65001
    vrf fpv-application_vrf
      address-family ipv4 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
      address-family ipv6 unicast
        advertise l2vpn evpn
        redistribute direct route-map fabric-rmap-redist-subnet
        maximum-paths ibgp 2
  interface nve1
    member vni 30001 associate-vrf
configure terminal
configure profile FPV-App-1_Network
  vlan 1001
    vn-segment 21001
    name FPV-App-1_VLAN
  interface Vlan1001
    description FPV-App-1_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.1.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21001
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21001 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-App-2_Network
  vlan 1002
    vn-segment 21002
    name FPV-App-2_VLAN
  interface Vlan1002
    description FPV-App-2_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.2.254/24 tag 12345
    mtu 9216
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21002
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21002 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
configure profile FPV-App-3_Network
  vlan 1003
    vn-segment 21003
    name FPV-App-3_VLAN
  interface Vlan1003
    description FPV-App-3_Interface
    vrf member fpv-application_vrf
    no ip redirects
    no ipv6 redirects
    ip address 172.22.3.254/24 tag 12345
    fabric forwarding mode anycast-gateway
    no shutdown
  interface nve1
    member vni 21003
      mcast-group 239.1.1.0
      suppress-arp
  evpn
    vni 21003 l2
      rd auto
      route-target import auto
      route-target export auto
configure terminal
snmp-server user admin network-admin auth md5 0xe476741bdd531efcbdb71ba42173560d
 priv 0xe476741bdd531efcbdb71ba42173560d localizedkey
snmp-server host 172.26.163.142 traps version 2c public udp-port 2162
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 172.26.163.254 use-vrf management
ntp server 172.26.164.254 use-vrf management

fabric forwarding anycast-gateway-mac 2020.0000.00aa
ipv6 switch-packets lla
ip pim rp-address 10.11.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
vlan 1,122,322,1001-1003,3000,3010,3020,3050,3500-3501

route-map fabric-rmap-redist-subnet permit 10
  match tag 12345
service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 172.26.163.254
hardware access-list tcam region ing-racl 1792
hardware access-list tcam region ing-flow-redirect 512
vpc domain 1
  peer-switch
  peer-keepalive destination 172.26.163.223 source 172.26.163.224
  virtual peer-link destination 10.11.0.3 source 10.11.0.4 dscp 56
  delay restore 150
  peer-gateway
  auto-recovery reload-delay 360
  ipv6 nd synchronize
  ip arp synchronize
ngoam install acl

nxapi http port 80
```

```
interface Vlan1
  no ip redirects
  no ipv6 redirects

interface port-channel1
  description To FXV-AA01-UCS6454FI-A: e1/54
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,3000,3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 1

interface port-channel2
  description To FXV-AA01-UCS6454FI-B: e1/54
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,3000,3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 2

interface port-channel3
  description To FXV-BB09-A300-2-01: e2e
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 3

interface port-channel4
  description To FXV-BB09-A300-2-02: e2e
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3010,3020,3050
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  mtu 9216
  vpc 4

interface port-channel500
  description "vpc-peer-link"
  switchport
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link

interface nve1
  no shutdown
  host-reachability protocol bgp
  advertise virtual-rmac
  source-interface loopback1

interface Ethernet1/1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,1001-1003,3000,3010,3020,3050
  mtu 9216
  channel-group 1 mode active
  no shutdown

interface Ethernet1/2
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,322,1001-1003,3000,3010,3020,3050
  mtu 9216
  channel-group 2 mode active
  no shutdown

interface Ethernet1/3
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/4
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/5
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,1001-1003,3010,3020,3050
  mtu 9216
  channel-group 3 mode active
  no shutdown

interface Ethernet1/6
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 122,1001-1003,3010,3020,3050
  mtu 9216
  channel-group 4 mode active
  no shutdown

interface Ethernet1/7
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/8
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/9
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/10
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown
```

```
interface Ethernet1/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/13
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/14
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/15
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/16
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/17
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/18
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/19
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/20
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/21
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/22
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/23
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/24
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/25
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/26
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/27
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/28
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/29
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/30
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/31
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/32
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/33
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/34
  switchport
  switchport mode trunk
  switchport trunk allowed vlan none
  spanning-tree port type edge trunk
  mtu 9216
  no shutdown

interface Ethernet1/35
  description connected-to-AA01-9364C-1-Ethernet1/4
  mtu 9216
  port-type fabric
  ip address 10.11.0.37/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/36
  description connected-to-AA01-9364C-2-Ethernet1/4
  mtu 9216
  port-type fabric
  ip address 10.11.0.29/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface mgmt0
  vrf member management
  ip address 172.26.163.224/24

interface loopback0
  description Routing loopback interface
  ip address 10.11.0.4/32
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode

interface loopback1
  description VTEP loopback interface
  ip address 10.11.1.1/32
  ip address 10.11.1.5/32 secondary
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
line console
line vty
boot nxos bootflash:/nxos.9.3.5.bin
router ospf Site-A_UNDERLAY
  router-id 10.11.0.4
router bgp 65001
  router-id 10.11.0.4
  address-family l2vpn evpn
    advertise-pip
  neighbor 10.11.0.1
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
  neighbor 10.11.0.2
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended

apply profile FPV-Foundation_VRF
apply profile FPV-iSCSI-A_Network
apply profile FPV-iSCSI-B_Network
apply profile FPV-InfraNFS_Network
apply profile FPV-InBand-SiteA_Network
apply profile FPV-vMotion_Network
apply profile FPV-CommonServices_Network
apply profile FPV-Application_VRF
apply profile FPV-App-1_Network
apply profile FPV-App-2_Network
apply profile FPV-App-3_Network
```

Spine A

```
version 9.3(5) Bios:version 05.42
switchname AA01-9364C-1
vdc AA01-9364C-1 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource vni_bd minimum 4096 maximum 4096
```

```
feature nxapi
nv overlay evpn
feature ospf
feature bgp
feature pim
feature lldp
feature nv overlay
feature ngoam

username admin password 5 $5$HNHHFM$IbBvjXFU4PJAO38qyINz5gtai/hlc.gqRANcOeYU/44
 role network-admin
ip domain-lookup
copp profile strict
snmp-server user admin network-admin auth md5 0x7f6f76fdde44e67a9a1c0a1ab74af65d
 priv 0x7f6f76fdde44e67a9a1c0a1ab74af65d localizedkey
snmp-server host 172.26.163.142 traps version 2c public udp-port 2162
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 172.26.163.254 use-vrf management

ipv6 switch-packets lla
ip pim rp-address 10.11.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ip pim anycast-rp 10.11.254.1 10.11.0.1
ip pim anycast-rp 10.11.254.1 10.11.0.2
vlan 1

vrf context management
  ip route 0.0.0.0/0 172.26.163.254
ngoam install acl

nxapi http port 80


interface Ethernet1/1
  description connected-to-AA01-93180LC-EX-1-Ethernet1/31
  mtu 9216
  ip address 10.11.0.9/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/2
  description connected-to-AA01-93180LC-EX-2-Ethernet1/31
  mtu 9216
  ip address 10.11.0.18/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/3
  description connected-to-AA01-9336C-FX2-1-Ethernet1/35
  mtu 9216
  ip address 10.11.0.26/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/4
  description connected-to-AA01-9336C-FX2-2-Ethernet1/35
  mtu 9216
  ip address 10.11.0.38/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/5
  mtu 9216
  no shutdown

interface Ethernet1/6
  mtu 9216
  no shutdown

interface Ethernet1/7
  mtu 9216
  no shutdown

interface Ethernet1/8
  mtu 9216
  no shutdown

interface Ethernet1/9
  mtu 9216
  no shutdown

interface Ethernet1/10
  mtu 9216
  no shutdown

interface Ethernet1/11
  mtu 9216
  no shutdown

interface Ethernet1/12
  mtu 9216
  no shutdown

interface Ethernet1/13
  mtu 9216
  no shutdown

interface Ethernet1/14
  mtu 9216
  no shutdown

interface Ethernet1/15
  mtu 9216
  no shutdown

interface Ethernet1/16
  mtu 9216
  no shutdown

interface Ethernet1/17
  mtu 9216
  no shutdown

interface Ethernet1/18
  mtu 9216
  no shutdown

interface Ethernet1/19
  mtu 9216
  no shutdown

interface Ethernet1/20
  mtu 9216
  no shutdown
```

```
interface Ethernet1/21
  mtu 9216
  no shutdown

interface Ethernet1/22
  mtu 9216
  no shutdown

interface Ethernet1/23
  mtu 9216
  no shutdown

interface Ethernet1/24
  mtu 9216
  no shutdown

interface Ethernet1/25
  mtu 9216
  no shutdown

interface Ethernet1/26
  mtu 9216
  no shutdown

interface Ethernet1/27
  mtu 9216
  no shutdown

interface Ethernet1/28
  mtu 9216
  no shutdown

interface Ethernet1/29
  mtu 9216
  no shutdown

interface Ethernet1/30
  mtu 9216
  no shutdown

interface Ethernet1/31
  mtu 9216
  no shutdown

interface Ethernet1/32
  mtu 9216
  no shutdown

interface Ethernet1/33
  mtu 9216
  no shutdown

interface Ethernet1/34
  mtu 9216
  no shutdown

interface Ethernet1/35
  mtu 9216
  no shutdown

interface Ethernet1/36
  mtu 9216
  no shutdown

interface Ethernet1/37
  mtu 9216
  no shutdown

interface Ethernet1/38
```

```
  mtu 9216
  no shutdown

interface Ethernet1/39
  mtu 9216
  no shutdown

interface Ethernet1/40
  mtu 9216
  no shutdown

interface Ethernet1/41
  mtu 9216
  no shutdown

interface Ethernet1/42
  mtu 9216
  no shutdown

interface Ethernet1/43
  mtu 9216
  no shutdown

interface Ethernet1/44
  mtu 9216
  no shutdown

interface Ethernet1/45
  mtu 9216
  no shutdown

interface Ethernet1/46
  mtu 9216
  no shutdown

interface Ethernet1/47
  mtu 9216
  no shutdown

interface Ethernet1/48
  mtu 9216
  no shutdown

interface Ethernet1/49
  mtu 9216
  no shutdown

interface Ethernet1/50
  mtu 9216
  no shutdown

interface Ethernet1/51
  mtu 9216
  no shutdown

interface Ethernet1/52
  mtu 9216
  no shutdown

interface Ethernet1/53
  mtu 9216
  no shutdown

interface Ethernet1/54
  mtu 9216
  no shutdown

interface Ethernet1/55
  mtu 9216
  no shutdown

interface Ethernet1/56
  mtu 9216
  no shutdown

interface Ethernet1/57
  mtu 9216
  no shutdown

interface Ethernet1/58
  mtu 9216
  no shutdown

interface Ethernet1/59
  mtu 9216
  no shutdown

interface Ethernet1/60
  mtu 9216
  no shutdown

interface Ethernet1/61
  mtu 9216
  no shutdown

interface Ethernet1/62
  mtu 9216
  no shutdown

interface Ethernet1/63
  mtu 9216
  no shutdown

interface Ethernet1/64
  mtu 9216
  no shutdown

interface Ethernet1/65
  mtu 9216
  no shutdown

interface Ethernet1/66
  mtu 9216
  no shutdown

interface mgmt0
  vrf member management
  ip address 172.26.163.231/24

interface loopback0
  description Routing loopback interface
  ip address 10.11.0.1/32
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode

interface loopback254
  description RP loopback interface
  ip address 10.11.254.1/32
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
line console
line vty
boot nxos bootflash:/nxos.9.3.5.bin
router ospf Site-A_UNDERLAY
  router-id 10.11.0.1
router bgp 65001
  router-id 10.11.0.1
  neighbor 10.11.0.3
```

```
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
  neighbor 10.11.0.4
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
  neighbor 10.11.0.5
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
  neighbor 10.11.0.6
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
```

Spine B

```
version 9.3(5) Bios:version 05.42
switchname AA01-9364C-2
vdc AA01-9364C-2 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource vni_bd minimum 4096 maximum 4096

feature nxapi
nv overlay evpn
feature ospf
feature bgp
feature pim
feature lldp
feature nv overlay
feature ngoam

username admin password 5 $5$JFLHLI$F14yAUx3SBPY9mG3JQkqDJ7R58UVW167CIjP6ElRO0C
 role network-admin
ip domain-lookup
copp profile strict
snmp-server user admin network-admin auth md5 0x380cfbb28eafe6263adbb5e5ce18a630
 priv 0x380cfbb28eafe6263adbb5e5ce18a630 localizedkey
snmp-server host 172.26.163.142 traps version 2c public udp-port 2162
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 172.26.163.254 use-vrf management

ipv6 switch-packets lla
ip pim rp-address 10.11.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ip pim anycast-rp 10.11.254.1 10.11.0.1
ip pim anycast-rp 10.11.254.1 10.11.0.2
vlan 1

vrf context management
  ip route 0.0.0.0/0 172.26.163.254
ngoam install acl

nxapi http port 80


interface Ethernet1/1
  description connected-to-AA01-93180LC-EX-1-Ethernet1/32
  mtu 9216
  ip address 10.11.0.13/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/2
  description connected-to-AA01-93180LC-EX-2-Ethernet1/32
  mtu 9216
  ip address 10.11.0.22/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/3
  description connected-to-AA01-9336C-FX2-1-Ethernet1/36
  mtu 9216
  ip address 10.11.0.34/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/4
  description connected-to-AA01-9336C-FX2-2-Ethernet1/36
  mtu 9216
  ip address 10.11.0.30/30
  ip ospf network point-to-point
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet1/5
  mtu 9216
  no shutdown

interface Ethernet1/6
  mtu 9216
  no shutdown

interface Ethernet1/7
  mtu 9216
  no shutdown

interface Ethernet1/8
  mtu 9216
  no shutdown

interface Ethernet1/9
  mtu 9216
  no shutdown

interface Ethernet1/10
  mtu 9216
  no shutdown

interface Ethernet1/11
  mtu 9216
  no shutdown

interface Ethernet1/12
  mtu 9216
  no shutdown

interface Ethernet1/13
  mtu 9216
  no shutdown

interface Ethernet1/14
  mtu 9216
  no shutdown

interface Ethernet1/15
  mtu 9216
  no shutdown

interface Ethernet1/16
  mtu 9216
  no shutdown

interface Ethernet1/17
  mtu 9216
  no shutdown

interface Ethernet1/18
  mtu 9216
  no shutdown

interface Ethernet1/19
  mtu 9216
  no shutdown

interface Ethernet1/20
  mtu 9216
  no shutdown

interface Ethernet1/21
  mtu 9216
  no shutdown

interface Ethernet1/22
  mtu 9216
  no shutdown

interface Ethernet1/23
  mtu 9216
  no shutdown

interface Ethernet1/24
  mtu 9216
  no shutdown

interface Ethernet1/25
  mtu 9216
  no shutdown

interface Ethernet1/26
  mtu 9216
  no shutdown

interface Ethernet1/27
  mtu 9216
  no shutdown
```

```
interface Ethernet1/28
  mtu 9216
  no shutdown

interface Ethernet1/29
  mtu 9216
  no shutdown

interface Ethernet1/30
  mtu 9216
  no shutdown

interface Ethernet1/31
  mtu 9216
  no shutdown

interface Ethernet1/32
  mtu 9216
  no shutdown

interface Ethernet1/33
  mtu 9216
  no shutdown

interface Ethernet1/34
  mtu 9216
  no shutdown

interface Ethernet1/35
  mtu 9216
  no shutdown

interface Ethernet1/36
  mtu 9216
  no shutdown

interface Ethernet1/37
  mtu 9216
  no shutdown

interface Ethernet1/38
  mtu 9216
  no shutdown

interface Ethernet1/39
  mtu 9216
  no shutdown

interface Ethernet1/40
  mtu 9216
  no shutdown

interface Ethernet1/41
  mtu 9216
  no shutdown

interface Ethernet1/42
  mtu 9216
  no shutdown

interface Ethernet1/43
  mtu 9216
  no shutdown

interface Ethernet1/44
  mtu 9216
  no shutdown
```

```
interface Ethernet1/45
  mtu 9216
  no shutdown

interface Ethernet1/46
  mtu 9216
  no shutdown

interface Ethernet1/47
  mtu 9216
  no shutdown

interface Ethernet1/48
  mtu 9216
  no shutdown

interface Ethernet1/49
  mtu 9216
  no shutdown

interface Ethernet1/50
  mtu 9216
  no shutdown

interface Ethernet1/51
  mtu 9216
  no shutdown

interface Ethernet1/52
  mtu 9216
  no shutdown

interface Ethernet1/53
  mtu 9216
  no shutdown

interface Ethernet1/54
  mtu 9216
  no shutdown

interface Ethernet1/55
  mtu 9216
  no shutdown

interface Ethernet1/56
  mtu 9216
  no shutdown

interface Ethernet1/57
  mtu 9216
  no shutdown

interface Ethernet1/58
  mtu 9216
  no shutdown

interface Ethernet1/59
  mtu 9216
  no shutdown

interface Ethernet1/60
  mtu 9216
  no shutdown

interface Ethernet1/61
  mtu 9216
  no shutdown

interface Ethernet1/62
```

```
  mtu 9216
  no shutdown

interface Ethernet1/63
  mtu 9216
  no shutdown

interface Ethernet1/64
  mtu 9216
  no shutdown

interface Ethernet1/65
  mtu 9216
  no shutdown

interface Ethernet1/66
  mtu 9216
  no shutdown

interface mgmt0
  vrf member management
  ip address 172.26.163.232/24

interface loopback0
  description Routing loopback interface
  ip address 10.11.0.2/32
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode

interface loopback254
  description RP loopback interface
  ip address 10.11.254.1/32
  ip router ospf Site-A_UNDERLAY area 0.0.0.0
  ip pim sparse-mode
line console
line vty
boot nxos bootflash:/nxos.9.3.5.bin
router ospf Site-A_UNDERLAY
  router-id 10.11.0.2
router bgp 65001
  router-id 10.11.0.2
  neighbor 10.11.0.3
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
  neighbor 10.11.0.4
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
  neighbor 10.11.0.5
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
  neighbor 10.11.0.6
    remote-as 65001
    update-source loopback0
    address-family l2vpn evpn
      send-community
      send-community extended
      route-reflector-client
```

About the Authors

Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.

Ramesh Isaac is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Group. Ramesh has worked in the data center and mixed-use lab settings since 1995. He started in information technology supporting UNIX environments and focused on designing and implementing multi-tenant virtualization solutions in Cisco labs before entering Technical Marketing where he has supported converged infrastructure and virtual services as part of solution offerings as Cisco. Ramesh has certifications from Cisco, VMware, and Red Hat.

Abhinav Singh, Technical Marketing Engineer, Hybrid Cloud Infrastructures, NetApp

Abhinav is a Technical Marketing Engineer in the Hybrid Cloud Infrastructures Engineering team at NetApp. He has more than 11 years of experience in data center infrastructure solutions which includes On-prem and Hybrid cloud space. He focuses on the validating, supporting, implementing cloud infrastructure solutions that include NetApp products. Prior to joining the Hybrid Cloud Infrastructure Engineering team at NetApp, he was with Cisco Systems as Technical Consulting Engineer working on Cisco Application Centric Infrastructure (ACI). Abhinav holds multiple certifications like Cisco Certified Network Professional (R&S), Double VCP (DCV,NV) and VMware Certified Implementation Expert for Network Virtualization (VCIX-NV). Abhinav holds a bachelor's degree in Electrical & Electronics.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Archana Sharma, Technical Marketing Engineer, Cisco Systems, Inc.
- John George, Technical Marketing Engineer, Cisco Systems, Inc.

Feedback

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

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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