Deployment Guide
Cisco Public

FlashStack Virtual Server Infrastructure with Cisco UCS X-Series and VMware 7.0 U2

Deployment Guide for FlashStack with VMware vSphere 7.0 U2, Cisco UCS X9508 Chassis with Cisco UCS X210c M6 Compute Nodes, and Pure Storage FlashArray//X R3 Series

Published: February 2022

In partnership with:

Contents

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.

Executive Summary

The FlashStack solution is a validated, converged infrastructure developed jointly by Cisco and Pure Storage. The solution offers a predesigned data center architecture that incorporates computing, storage, and network design best practices to reduce IT risk by validating the architecture and helping ensure compatibility among the components. The FlashStack solution is successful because of its ability to evolve and incorporate both technology and product innovations in the areas of management, compute, storage, and networking. This document covers the design details of incorporating the Cisco Unified Computing System™ (Cisco UCS®) X-Series modular platform into the FlashStack Virtual Server Infrastructure (VSI) and its ability to manage and orchestrate FlashStack components from the cloud using the Cisco Intersight™. Some of the most important advantages of integrating the Cisco UCS X-Series into the FlashStack infrastructure include:

- Simpler and programmable infrastructure: Infrastructure as a code delivered through an open application programming interface (API)
- Power and cooling innovations: Higher-power headroom and lower energy loss because of a 54V DC power delivery to the chassis
- Better airflow: Midplane free design with fewer barriers, thus lower impedance
- Fabric innovations: PCIe/Compute Express Link (CXL) topology for heterogeneous compute and memory composability
- Innovative cloud operations: Continuous feature delivery and no need for managing virtual machines
- Built for investment protections: Design-ready for future technologies such as liquid-cooling and high-wattage CPUs; CXL-ready

In addition to the compute-specific hardware and software innovations, integration of the Cisco Intersight cloud platform with VMware vCenter and Pure Storage FlashArray delivers monitoring, orchestration, and workload optimization capabilities for different layers (virtualization and storage) of the FlashStack solution. The modular nature of the Cisco Intersight platform also provides an easy upgrade path to additional services such as workload optimization and Kubernetes.

Customers interested in understanding FlashStack design and deployment details, including the configuration of various elements of design and associated best practices, should refer to Cisco Validated Designs for FlashStack at: https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/data-center-design-guides-all.html - FlashStack.

Solution Overview

Introduction

The Cisco UCS X-Series is a new modular compute system configured and managed from the cloud. It is designed to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco Intersight platform software-as-a-service (SaaS) infrastructure lifecycle management platform delivers simplified configuration, deployment, maintenance, and support.

Powered by the Cisco Intersight cloud operations platform, the Cisco UCS X-Series enables the next-generation cloud-operated FlashStack infrastructure that not only simplifies the datacenter management but also allows the infrastructure to adapt to unpredictable needs of the modern applications as well as traditional workloads. With the Cisco Intersight platform, you get all the benefits of SaaS delivery and the full lifecycle management of Cisco Intersight connected, distributed servers and integrated Pure Storage FlashArray across data centers, remote sites, branch offices, and edge environments.

Audience

The intended audience of this document includes but is not limited to data scientists, IT architects, sales engineers, field consultants, professional services, IT managers, partner engineering, DevOps, and Site Reliability Engineers (SREs), 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 guidance around incorporating the Cisco Intersight software-managed Cisco UCS X-Series platform within the FlashStack solution. The document introduces various design elements and addresses various considerations and best practices for a successful deployment. It also highlights the design and product requirements for integrating virtualization and storage systems with the Cisco Intersight platform to deliver a true cloud-based integrated approach to infrastructure management.

What’s New in this Release?

The following design elements distinguish this version of the FlashStack VSI solution from previous models:

- Integration of the Cisco UCS X-Series into FlashStack
- Management of the Cisco UCS X-Series from the cloud using the Cisco Intersight platform
- Integration of the Cisco Intersight platform with Pure Storage FlashArray for storage monitoring and orchestration
- Integration of the Cisco Intersight software with VMware vCenter for interacting with, monitoring, and orchestrating the virtual environment

Deployment Hardware and Software

Architecture

The FlashStack VSI with Cisco UCS X-Series and VMware vSphere 7.0 U2 delivers a cloud-managed infrastructure solution on the latest Cisco UCS hardware. VMware vSphere 7.0 U2 hypervisor is installed on the Cisco UCS X210c M6 Compute Nodes configured for stateless compute design using boot from SAN. Pure Storage FlashArray//X50 R3 provides the storage infrastructure required for setting up the VMware environment. The Cisco Intersight cloud-management platform is utilized to configure and manage the infrastructure. The solution requirements and design details are covered in this section.

Requirements

The FlashStack VSI with Cisco UCS X-Series meets the following general design requirements:

- Resilient design across all layers of the infrastructure with no single point of failure
- Scalable design with the flexibility to add compute and storage capacity or network bandwidth as needed
- Modular design that can be replicated to expand and grow as the needs of the business grow
- Flexible design that can support different models of various components with ease
- Simplified design with the ability to integrate and automate with external automation tools
- Cloud-enabled design which can be configured, managed and orchestrated from the cloud using GUI or APIs

Physical Topology

FlashStack with Cisco UCS X-Series supports both IP and Fibre Channel (FC) based storage access design. For the IP-based solution, iSCSI configuration on Cisco UCS and Pure Storage FlashArray is utilized to set up storage access including boot from SAN configuration for the compute nodes. For the Fibre Channel designs, Pure Storage FlashArray and Cisco UCS X-Series are connected using Cisco MDS 9132T switches and storage access, including boot from SAN, is provided over the Fibre Channel network. The physical connectivity details for both IP and FC designs are explained below.

IP-based Storage Access

The physical topology for the IP-based FlashStack is shown in Figure 1.

Figure 1. FlashStack - physical topology for IP connectivity

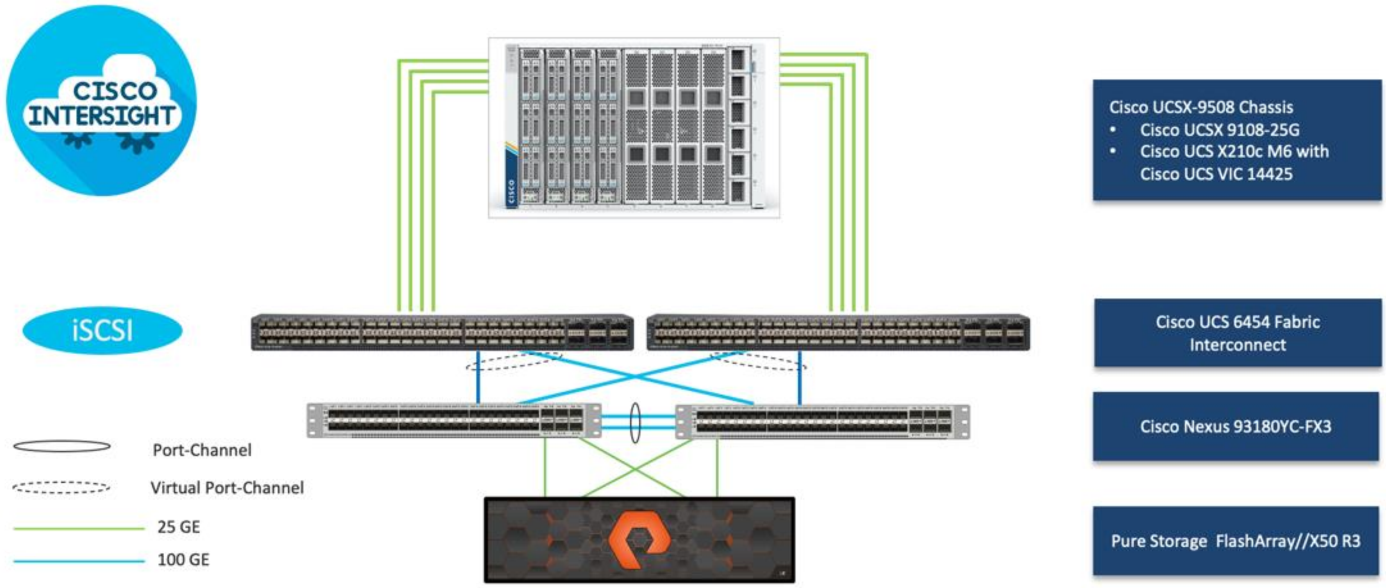

To validate the IP-based storage access in a FlashStack configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCSX 9108-25G intelligent fabric modules (IFMs), where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate FI. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCSX-210c M6 Compute Nodes contain fourth-generation Cisco 14425 virtual interface cards.
- Cisco Nexus 93180YC-FX3 Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 100-Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX3 Switches in a Virtual Port Channel (vPC) configuration.
- The Pure Storage FlashArray//X50 R3 connects to the Cisco Nexus 93180YC-FX3 switches using four 25-GE ports.
- VMware 7.0 U2 ESXi software is installed on Cisco UCSX-210c M6 Compute Nodes to validate the infrastructure.

FC-based Storage Access

Figure 2 illustrates the FlashStack physical topology for FC connectivity.

Figure 2. FlashStack- physical topology for FC connectivity

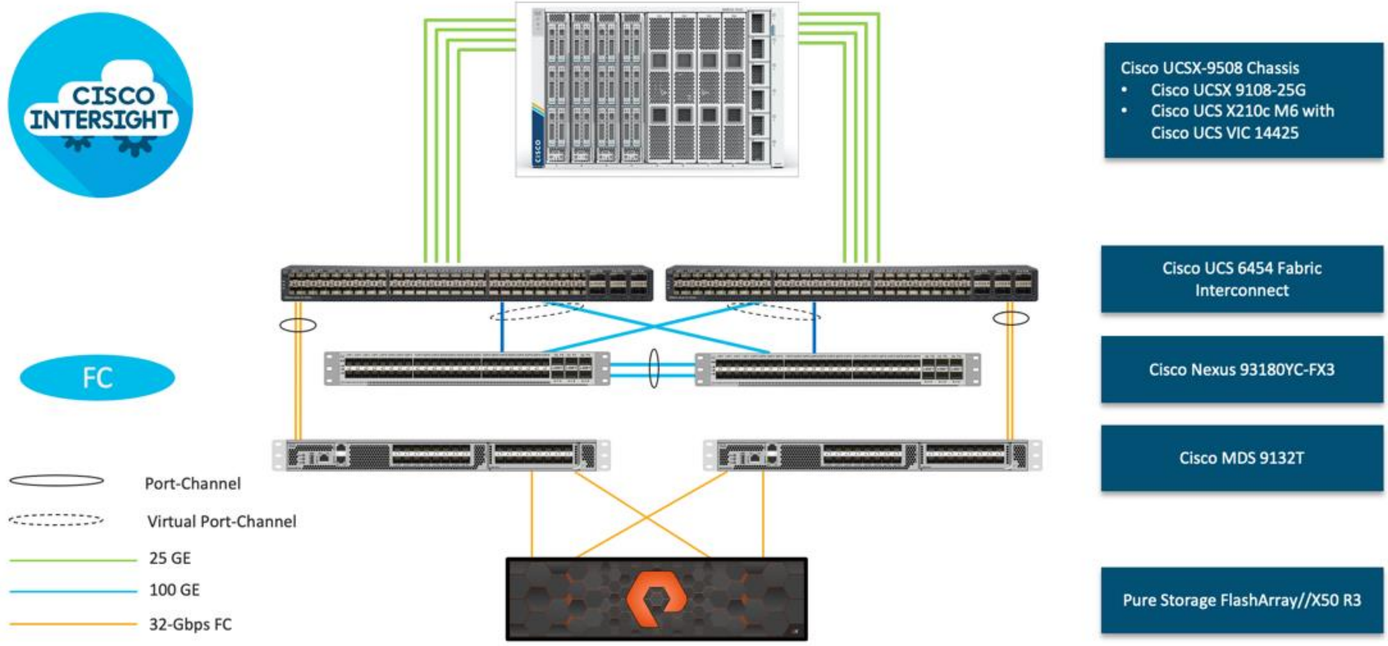

To validate the FC-based storage access in a FlashStack configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCSX 9108-25G Intelligent Fabric Modules (IFMs), where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate FI.
- Cisco UCS X210c M6 Compute Nodes contain fourth-generation Cisco 14425 virtual interface cards.
- Cisco Nexus 93180YC-FX3 Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 100 Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX3 Switches in a vPC configuration.
- Cisco UCS 6454 Fabric Interconnects are connected to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections configured as a single port channel for SAN connectivity.
- The Pure Storage FlashArray//X50 R3 connects to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections for SAN connectivity.
- VMware 7.0 U2 ESXi software is installed on Cisco UCS X210c M6 Compute Nodes to validate the infrastructure.

Deployment Hardware and Software

Table 1 lists the hardware and software versions used during solution validation. It is important to note that the validated FlashStack solution explained in this document adheres to Cisco, Pure Storage, and VMware interoperability matrix to determine support for various software and driver versions. Customers should use the same interoperability matrix to determine support for components that are different from the current validated design.

Click the following links for more information:

- Cisco UCS Hardware and Software Interoperability Tool: http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html
- Pure Storage Interoperability (note, this interoperability list will require a support login form Pure): https://support.purestorage.com/FlashArray/Getting_Started/Compatibility_Matrix
- Pure Storage FlashStack Compatibility Matrix (note, this interoperability list will require a support login from Pure): https://support.purestorage.com/FlashStack/Product_Information/FlashStack_Compatibility_Matrix
- VMware Compatibility Guide: http://www.vmware.com/resources/compatibility/search.php

Additionally, it is also strongly suggested to align FlashStack deployments with the recommended release for the Cisco Nexus 9000 switches used in the architecture:

- Cisco Nexus: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended_release/b_Minimum_and_Recommended_Cisco_NX-OS_Releases_for_Cisco_Nexus_9000_Series_Switches.html
- MDS: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b_MDS_NX-OS_Recommended_Releases.html

Table 1. Hardware and Software Revisions

Component		Software
Network	Cisco Nexus 93180YC-FX3	9.3(7)
	Cisco MDS 9132T	8.4(2c)
Compute	Cisco UCS Fabric Interconnect 6454 and UCSX 9108-25G IFM	4.2(1h)
	Cisco UCS X210C with VIC 14425	5.0(1b)
	VMware ESXi	7.0 U2a
	Cisco VIC ENIC Driver for ESXi	1.0.35.0
	Cisco VIC FNIC Driver for ESXi	5.0.0.15
	VMware vCenter Appliance	7.0 U2b
	Cisco Intersight Assist Virtual Appliance	1.0.9-342
Storage	Pure Storage FlashArray//X50 R3	6.1.11
	Pure Storage VASA Provider	3.5
	Pure Storage Plugin	5.0.0

Configuration Guidelines

This document details the step-by-step configuration of a fully redundant and highly available Virtual Server Infrastructure built on Cisco and Pure Storage components. References are made to which component is being configured with each step, either 01 or 02 or A and B. For example, controller-1 and controller-2 are used to identify the two controllers within the Pure Storage FlashArray//X 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-FCP-01, VM-Host-Infra-FCP-02 to represent Fibre Channel booted infrastructure and production hosts deployed to the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in each step, <<text>> appears as part of the command structure. The following is an example of a configuration step for both Cisco Nexus switches:

```
BB08-93180YC-FX-A (config)# ntp server <<var_oob_ntp>> use-vrf management
```

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, and Table 3 lists the external dependencies necessary for deployment as outlined in this guide.

Table 2. Necessary VLANs

VLAN ID	Name	Usage
2	Native-VLAN	Use VLAN 2 as Native VLAN instead of default VLAN (1)
15	OOB-MGMT-VLAN	Out-of-Band Management VLAN to connect the management ports for various devices
115	IB-MGMT-VLAN	In-Band Management VLAN utilized for all in-band management connectivity for example, ESXi hosts, VM management, and so on.
1101	VM-Traffic-VLAN	VM data traffic VLAN.
1130	vMotion-VLAN	VMware vMotion traffic.
901*	iSCSI-A-VLAN	iSCSI-A path for supporting boot-from-san for both Cisco UCS B-Series and Cisco UCS C-Series servers
902*	iSCSI-B-VLAN	iSCSI-B path for supporting boot-from-san for both Cisco UCS B-Series and Cisco UCS C-Series servers

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

Table 3. Virtual Machines

Virtual Machine Description	Host Name	IP Address
vCenter Server		
Cisco Data Center Network Manager (DCNM)		
Cisco Intersight Assist		

Table 4. Configuration Variables

Variable Name	Variable Description	Customer Variable Name
<<var_nexus_A_hostname>>	Cisco Nexus switch A Host name (Example: BB08-91380YX-FX-A)	
<<var_nexus_A_mgmt_ip>>	Out-of-band management IP for Cisco Nexus switch A (Example: 10.1.164.61)	
<<var_oob_mgmt_mask>>	Out-of-band network mask (Example: 255.255.255.0)	
<<var_oob_gateway>>	Out-of-band network gateway (Example: 10.1.164.254)	
<<var_oob_ntp>>	Out-of-band management network NTP Server (Example: 10.1.164.254)	
<<var_nexus_B_hostname>>	Cisco Nexus switch B Host name (Example: BB08-91380YX-FX-B)	
<<var_nexus_B_mgmt_ip>>	Out-of-band management IP for Nexus switch B (Example: 10.1.164.62)	
<<var_flasharray_hostname>>	Array Hostname set during setup (Example: BB08-FlashArrayR3)	
<<var_flasharray_vip>>	Virtual IP that will answer for the active management controller (Example: 10.2.164.100)	
<<var_contoller-1_mgmt_ip>>	Out-of-band management IP for FlashArray controller-1 (Example:10.2.164.101)	
<<var_contoller-1_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)	
<<var_contoller-1_mgmt_gateway>>	Out-of-band management network default gateway (Example: 10.2.164.254)	
<<var_contoller-2_mgmt_ip>>	Out-of-band management IP for FlashArray controller-2 (Example:10.2.164.102)	
<<var_contoller-2_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)	
<<var_ contoller-2_mgmt_gateway>>	Out-of-band management network default gateway (Example: 10.2.165.254)	
<<var_password>>	Administrative password (Example: Fl@shSt4x)	
<<var_dns_domain_name>>	DNS domain name (Example: flashstack.cisco.com)	
<<var_nameserver_ip>>	DNS server IP(s) (Example: 10.1.164.125)	
<<var_smtp_ip>>	Email Relay Server IP Address or FQDN (Example: smtp.flashstack.cisco.com)	
<<var_smtp_domain_name>>	Email Domain Name (Example: flashstack.cisco.com)	
<<var_timezone>>	FlashStack time zone (Example: America/New_York)	

Variable Name	Variable Description	Customer Variable Name
<<var_oob_mgmt_vlan_id>>	Out-of-band management network VLAN ID (Example: 15)	
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID (Example: 215)	
<<var_ib_mgmt_vlan_netmask_length>>	Length of IB-MGMT-VLAN Netmask (Example: /24)	
<<var_ib_gateway_ip>>	In-band management network VLAN ID (Example: 10.2.164.254)	
<<var_vmotion_vlan_id>>	vMotion network VLAN ID (Example: 1130)	
<<var_vmotion_vlan_netmask_length>>	Length of vMotion VLAN Netmask (Example: /24)	
<<var_native_vlan_id>>	Native network VLAN ID (Example: 2)	
<<var_app_vlan_id>>	Example Application network VLAN ID (Example: 1101)	
<<var_snmp_contact>>	Administrator e-mail address (Example: admin@flashstack.cisco.com)	
<<var_snmp_location>>	Cluster location string (Example: RTP1-AA19)	
<<var_mds_A_mgmt_ip>>	Cisco MDS Management IP address (Example: 10.1.164.63)	
<<var_mds_A_hostname>>	Cisco MDS hostname (Example: BB08-MDS-9132T-A)	
<<var_mds_B_mgmt_ip>>	Cisco MDS Management IP address (Example: 10.1.164.64)	
<<var_mds_B_hostname>>	Cisco MDS hostname (Example: BB08-MDS-9132T-B)	
<<var_vsan_a_id>>	VSAN used for the A Fabric between the FlashArray/MDS/FI (Example: 100)	
<<var_vsan_b_id>>	VSAN used for the B Fabric between the FlashArray/MDS/FI (Example: 200)	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name (Example: BB08-FI-6454)	
<<var_ucs_a_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address (Example: 10.1.164.51)	
<<var_ucs_mgmt_vip>>	Cisco UCS fabric interconnect (FI) Cluster out-of-band management IP address (Example: 10.1.164.50)	
<<var_ucs b_mgmt_ip>>	Cisco UCS fabric interconnect (FI) Cluster out-of-band management IP address (Example: 10.1.164.52)	
<<var_vm_host_fc_01_ip>>	VMware ESXi host 01 in-band management IP (Example:10.1.164.111)	

Variable Name	Variable Description	Customer Variable Name
<<var_vm_host_fc_vmotion_01_ip>>	VMware ESXi host 01 vMotion IP (Example: 192.168.130.101)	
<<var_vm_host_fc_02_ip>>	VMware ESXi host 02 in-band management IP (Example:10.1.164.112)	
<<var_vm_host_fc_vmotion_02_ip>>	VMware ESXi host 02 vMotion IP (Example: 192.168.130.102)	
<<var_vmotion_subnet_mask>>	vMotion subnet mask (Example: 255.255.255.0)	
<<var_vcenter_server_ip>>	IP address of the vCenter Server (Example: 10.1.164.110)	

Physical Infrastructure

FlashStack Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlashStack 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 Pure FlashArray//X R3 running Purity 6.1.11.

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.

Note: Make sure to use the cabling directions in this section as a guide.

Figure 3 details the cable connections used in the validation lab for FlashStack topology based on the Cisco UCS 6454 fabric interconnect. Four 32Gb uplinks connect as port-channels to each Cisco UCS Fabric Interconnect from the MDS switches, and a total of eight 32Gb links connect the MDS switches to the Pure FlashArray//X R3 controllers, four of these have been used for scsi-fc and the other four to support nvme-fc. The 100Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the Pure FlashArray//X R3 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 FlashStack infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each FlashArray controller has a connection to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

Note: Although the following diagram includes the Cisco UCS 5108 chassis with Cisco UCS B-Series M6 servers, this document describes the configuration of only Cisco UCS X210c M6 servers in the Cisco UCS X9508 chassis. However, Cisco UCS X9508 chassis with X210c M6 servers and Cisco UCS 5108 chassis with Cisco UCS B200 M6 servers can be connected to the same set of fabric interconnects with common management using Cisco Intersight.

Figure 3. **FlashStack Cabling with Cisco UCS 6454 Fabric Interconnect**

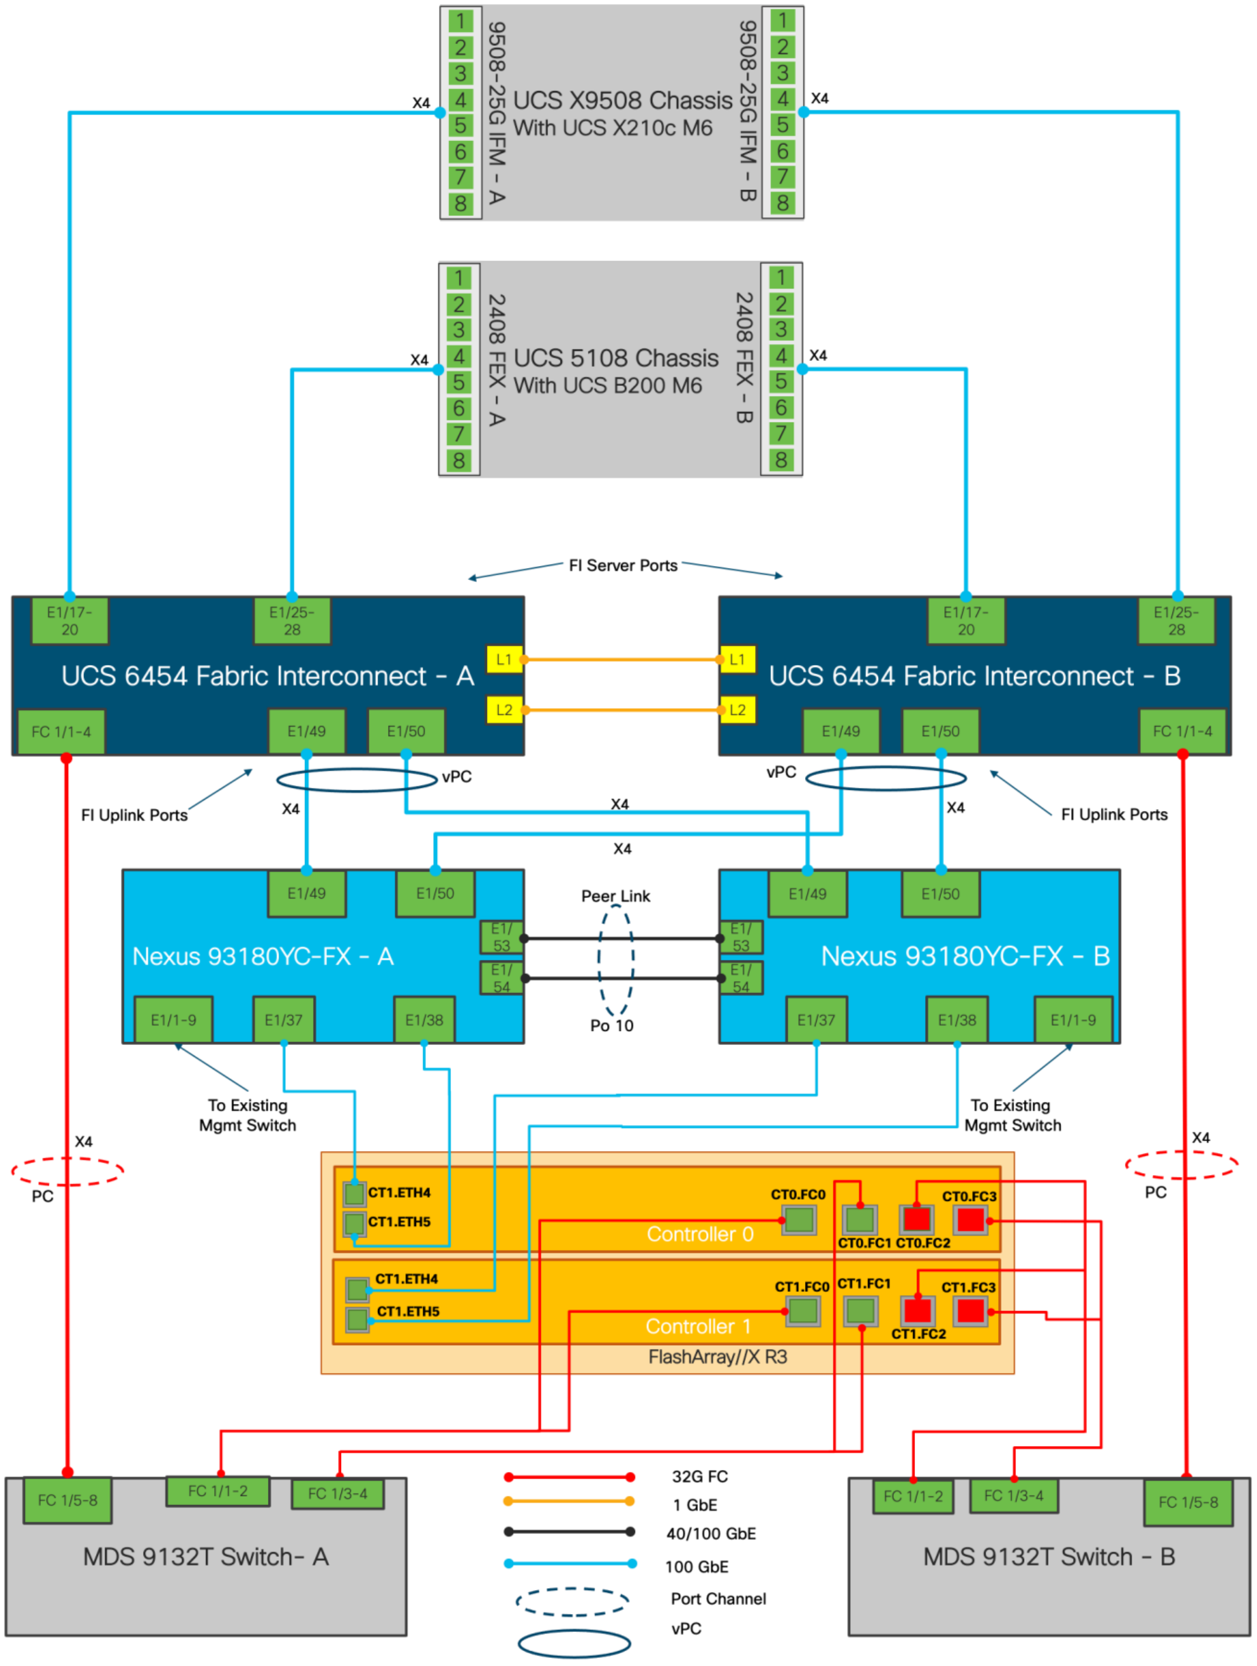

Note: Cisco UCS Fabric Interconnect's to the Cisco Nexus 93180YC-FX switches connectivity can be done using the 100Gbe or 25Gbe ports based on the bandwidth requirements, this document includes the usage of 100Gbe ports with aggregate bandwidth of 200Gbe per port channel from the Cisco UCS FI to the Cisco Nexus switches.

Note: * iSCSI connectivity is not required if iSCSI storage access is not being implemented.

Table 5. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco Nexus 93180YC-FX-A	Eth 1/49	100Gbe	Cisco UCS 6454-A	Eth 1/49
	Eth 1/50	100Gbe	Cisco UCS 6454-B	Eth 1/49
	Eth 1/53	100Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/53
	Eth 1/54	100Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/54
	Eth 1/9	10Gbe or 25 Gbe	Upstream Network Switch	Any
	Mgmt0	Gbe	Gbe Management Switch	Any
	Eth 1/37 *	25Gbe	FlashArray//X50 R3 Controller 1	CT0.ETH4
	Eth 1/38 *	25Gbe	FlashArray//X50 R3 Controller 2	CT1.ETH4

Table 6. Cisco Nexus 93180YC-FX-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco Nexus 93180YC-FX-B	Eth 1/49	100Gbe	Cisco UCS 6454-A	Eth 1/50
	Eth 1/50	100Gbe	Cisco UCS 6454-B	Eth 1/50
	Eth 1/9	10Gbe or 25 Gbe	Upstream Network Switch	Any
	Mgmt0	Gbe	Gbe Management Switch	Any
	Eth 1/37 *	25Gbe	FlashArray//X50 R3 Controller 1	CT0.ETH5
	Eth 1/38 *	25Gbe	FlashArray//X50 R3 Controller 2	CT1.ETH5

Table 7. Cisco UCS-6545-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco UCS-6454-A	Eth 1/49	100Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/49
	Eth 1/50	100Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/49
	Eth 1/17	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G A	IFM 1/1
	Eth 1/18	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G A	IFM 1/2
	Eth 1/19	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G A	IFM 1/3
	Eth 1/20	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G A	IFM 1/4
	FC1/1	32G FC	Cisco MDS 9132T-A	FC1/1
	FC1/2	32G FC	Cisco MDS 9132T-A	FC1/2
	FC1/3	32G FC	Cisco MDS 9132T-A	FC1/3
	FC1/4	32G FC	Cisco MDS 9132T-A	FC1/4
	Mgmt0	Gbe	Gbe Management Switch	Any

Table 8. Cisco UCS-6545-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco UCS-6454-B	Eth 1/49	100Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/50
	Eth 1/50	100Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/50
	Eth 1/17	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G B	IFM 1/1
	Eth 1/18	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G B	IFM 1/2
	Eth 1/19	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G B	IFM 1/3
	Eth 1/20	25Gbe	Cisco UCS Chassis 1 IFM 9108-25G B	IFM 1/4
	FC1/1	32G FC	Cisco MDS 9132T-B	FC1/1
	FC1/2	32G FC	Cisco MDS 9132T-B	FC1/2
	FC1/3	32G FC	Cisco MDS 9132T-B	FC1/3
	FC1/4	32G FC	Cisco MDS 9132T-B	FC1/4

Local Device	Local Port	Connection	Remote Device	Remote port
	Mgmt0	Gbe	Gbe Management Switch	Any

Table 9. Cisco MDS-9132T-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco MDS-9132T-A	FC1/5	32Gb FC	Cisco UCS 6454-A	FC1/1
	FC1/6	32Gb FC	Cisco UCS 6454-A	FC1/2
	FC 1/7	32Gb FC	Cisco UCS 6454-A	FC1/3
	FC 1/8	32Gb FC	Cisco UCS 6454-A	FC1/4
	FC1/1	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC0 (scsi-fc)
	FC1/2	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC0 (scsi-fc)
	FC1/3	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC1 (nvme-fc)
	FC1/4	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC1 (nvme-fc)
	Mgmt0	Gbe	Gbe Management Switch	Any

Note: This design uses SCSI-FCP for boot and datastore storage access and Port numbers 0 and 2 on each Pure FlashArray Controller have been used for the fibre channel connectivity, the ports 1 and 3 are used for FC-NVMe datastore access. All the four ports can be used for SCSI-FCP or FC-NVMe as needed but each port can only function as an SCSI-FCP or FC-NVMe port.

Table 10. Cisco MDS-9132T-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco MDS-9132T-B	FC1/5	32Gb FC	Cisco UCS 6454-B	FC1/1
	FC1/6	32Gb FC	Cisco UCS 6454-B	FC1/2
	FC 1/7	32Gb FC	Cisco UCS 6454-B	FC1/3
	FC 1/8	32Gb FC	Cisco UCS 6454-B	FC1/4
	FC1/1	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC2 (scsi-fc)
	FC1/2	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC2 (scsi-fc)
	FC1/3	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC3 (nvme-fc)
	FC1/4	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC3 (nvme-fc)
	Mgmt0	Gbe	Gbe Management Switch	Any

Note: This design uses SCSI-FCP for boot and datastore storage access and Port numbers 0 and 2 on each Pure FlashArray Controller have been used for the fibre channel connectivity, the ports 1 and 3 are used for FC-NVMe datastore access. All the four ports can be used for SCSI-FCP or FC-NVMe as needed but each port can only function as an SCSI-FCP or FC-NVMe port.

Table 11. Pure Storage FlashArray//X50 R3 Controller 1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
FlashArray//X50 R3 Controller 1	CT0.FC0 (scsi-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/1
	CT0.FC2 (scsi-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/1
	CT0.FC1 (nvme-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/3
	CT0.FC3 (nvme-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/3
	CT0.ETH4 *	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/37
	CT0.ETH5 *	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/37

Note: * Required only if iSCSI storage access is implemented.

Table 12. Pure Storage FlashArray//X50 R3 Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
FlashArray//X50 R3 Controller 2	CT1.FC0 (scsi-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/2
	CT1.FC2 (scsi-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/2
	CT1.FC1 (nvme-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/4
	CT1.FC3 (nvme-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/4
	CT1.ETH4 *	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/38
	CT1.ETH5 *	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/38

Note: * Required only if iSCSI storage access is implemented.

Network Switch Configuration

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlashStack environment. This procedure assumes the use of Cisco Nexus 93180YC-FX switches running NX-OS 9.3(8). Configuring on a differing model of Cisco Nexus 9000 series switches should be comparable but may differ slightly with model and changes in NX-OS release. The Cisco Nexus 93180YC-FX switch and the NX-OS 9.3(8) release were used in validating this FlashStack solution, so the steps will reflect this model and release.

Figure 4. Network Configuration workflow

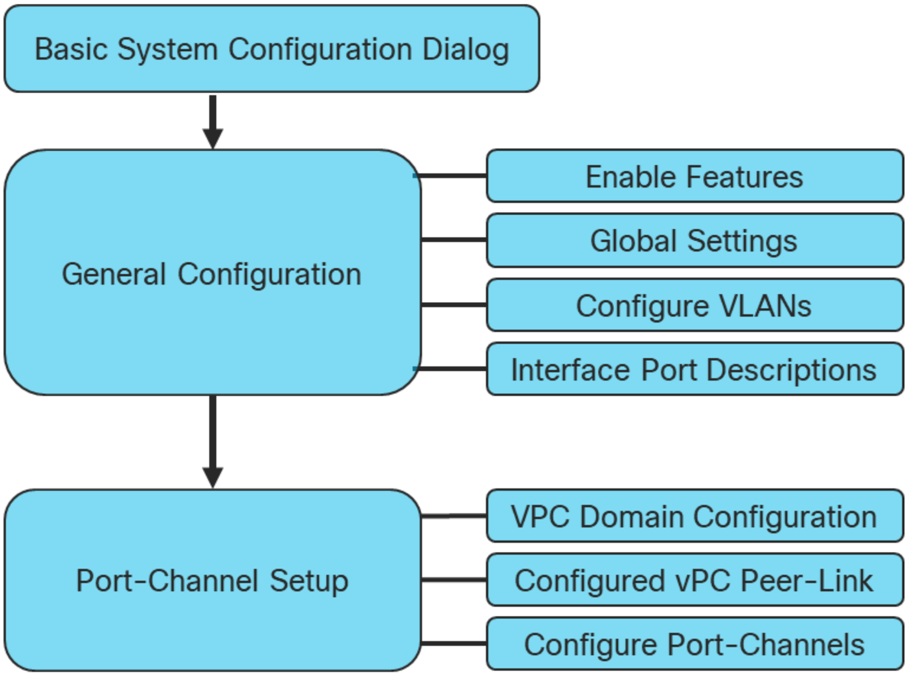

Physical Connectivity

Physical cabling should be completed by following the diagram and table references in section FlashStack Cabling.

FlashStack Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus 93180YC-FX switches for use in a base FlashStack environment. This procedure assumes the use of Cisco Nexus 9000 9.3(8), the Cisco suggested Nexus switch release at the time of this validation.

Note: The following procedure includes the setup of NTP distribution on both the mgmt0 port and the in-band management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes that the default VRF is used to route the in-band management VLAN.

Cisco Nexus A

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

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no)[no]: yes
Disabling POAP.......Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

         ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: n
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus B

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

1. Configure the switch.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no)[no]: yes
Disabling POAP.......Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

         ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-B-hostname>
```

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

FlashStack Cisco Nexus Switch Configuration

Enable Features

Cisco Nexus A and Cisco Nexus B

To enable the appropriate features on the Cisco Nexus switches, follow these steps:

1. Log in as admin.

2. Run the following commands:

```
config t
feature udld
feature interface-vlan
feature lacp
feature vpc
feature lldp
feature nxapi
```

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

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

1. Run the following commands to set global configurations:

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

Note: It is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summer time, please see Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3(x). Sample clock commands for the United States Eastern timezone are:
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60

Create VLANs

Cisco Nexus A and Cisco Nexus B

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

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

```
vlan <oob-mgmt-vlan-id>
name OOB-MGMT
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-Vlan
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus A

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

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

Cisco Nexus B

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

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

Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, follow these steps:

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

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

```
interface Eth1/49
description <ucs-clustername>-A:1/49
udld enable
interface Eth1/50
description <ucs-clustername>-B:1/49
udld enable
```

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

```
interface Ethernet1/53
description Peer Link <<nexus-B-hostname>>:Eth1/53
interface Ethernet1/54
description Peer Link <<nexus-B-hostname>>:Eth1/54
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B and to enable aggressive UDLD on copper interfaces connected to Cisco UCS systems, follow this step:

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

```
interface Eth1/49
description <ucs-clustername>-A:1/50
udld enable
interface Eth1/50
description <ucs-clustername>-B:1/50
udld enable
```

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

```
interface Ethernet1/53
description Peer Link <<nexus-A-hostname>>:Eth1/53
interface Ethernet1/54
description Peer Link <<nexus-A-hostname>>:Eth1/54
```

Create Port Channels

Cisco Nexus A and Cisco Nexus B

To create the necessary port channels between devices, follow this step on both switches:

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

```
interface Po10
description vPC peer-link
interface Eth1/53-54
channel-group 10 mode active
no shutdown
interface Po121
description <ucs-clustername>-A
interface Eth1/49
channel-group 121 mode active
no shutdown
interface Po123
description <ucs-clustername>-B
interface Eth1/50
channel-group 123 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, follow this step on both switches:

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

```
interface Po10
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>, <oob-mgmt-vlan-id>
spanning-tree port type network
speed 100000
duplex full
state enabled

interface Po121
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>, <oob-mgmt-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled


interface Po123
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>, <oob-mgmt-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

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

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

```
vpc domain <nexus-vpc-domain-id>
```

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

Cisco Nexus B

To configure vPCs for switch B, follow this step:

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

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po121
vpc 121
interface Po123
vpc 123
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlashStack environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlashStack environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Switch Testing Commands

The following commands can be used to check for correct switch configuration:

Note: Some of these commands need to run after further configuration of the FlashStack components are complete to see complete results.

```
show run
show vpc
show port-channel summary
show ntp peer-status
show cdp neighbors
show lldp neighbors
show run int
show int
show udld neighbors
show int status
```

Storage Configuration

Pure Storage FlashArray//X50 R3 Initial Configuration

FlashArray Initial Configuration

The following information should be gathered to enable the installation and configuration of the FlashArray. An official representative of Pure Storage will help rack and configure the new installation of the FlashArray.

Array Settings	Variable Name
Array Name (Hostname for Pure Array):	<<var_flasharray_hostname>>
Virtual IP Address for Management:	<<var_flasharray_vip>>
Physical IP Address for Management on Controller 0 (CT0):	<<var_contoller-1_mgmt_ip >>
Physical IP Address for Management on Controller 1 (CT1):	<<var_contoller-2_mgmt_ip>>
Netmask:	<<var_contoller-1_mgmt_mask>>
Gateway IP Address:	<<var_contoller-1_mgmt_gateway>>
DNS Server IP Address(es):	<<var_nameserver_ip>>
DNS Domain Suffix: (Optional)	<<var_dns_domain_name>>
NTP Server IP Address or FQDN:	<<var_oob_ntp>>
Email Relay Server (SMTP Gateway IP address or FQDN): (Optional)	<<var_smtp_ip>>
Email Domain Name:	<<var_smtp_domain_name>>
Alert Email Recipients Address(es): (Optional)	
HTTP Proxy Server ad Port (For Pure1): (Optional)	
Time Zone:	<<var_timezone>>

When the FlashArray has completed initial configuration, it is important to configure the Cloud Assist phone-home connection to provide the best pro-active support experience possible. Furthermore, this will enable the analytics functionalities provided by Pure1.

Add an Alert Recipient

The Alerts sub-view is used to manage the list of addresses to which Purity delivers alert notifications, and the attributes of alert message delivery. You can designate up to 19 alert recipients. The Alert Recipients section displays a list of email addresses that are designated to receive Purity alert messages. Up to 20 alert recipients can be designated.

Note: The list includes the built-in flasharray-alerts@purestorage.com address, which cannot be deleted.

The email address that Purity uses to send alert messages includes the sender domain name and is comprised of the following components:

<Array_Name>-<Controller_Name>@<Sender_Domain_Name>.com

To add an alert recipient, follow these steps:

1. Select Settings.

2. In the Alert Watchers section, enter the email address of the alert recipient and click the + icon.

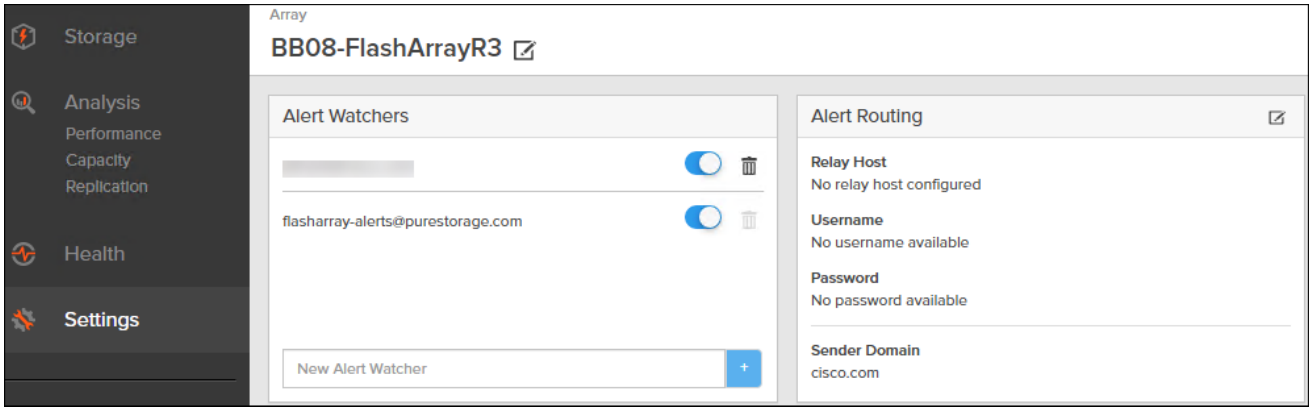

Array

BB08-FlashArrayR3

Alert Watchers

[redacted]

flasharray-alerts@purestorage.com

New Alert Watcher

Alert Routing

Relay Host
No relay host configured

Username
No username available

Password
No password available

Sender Domain
cisco.com

The Relay Host section displays the hostname or IP address of an SMTP relay host, if one is configured for the array. If you specify a relay host, Purity routes the email messages via the relay (mail forwarding) address rather than sending them directly to the alert recipient addresses.

In the Sender Domain section, the sender domain determines how Purity logs are parsed and treated by Pure Storage Support and Escalations. By default, the sender domain is set to the domain name please-configure.me.

It is crucial that you set the sender domain to the correct domain name. If the array is not a Pure Storage test array, set the sender domain to the actual customer domain name. For example, mycompany.com.

Configure Pure1 Support

The Pure1 Support section manages settings for Phone Home, Remote Assist, and Support Logs.

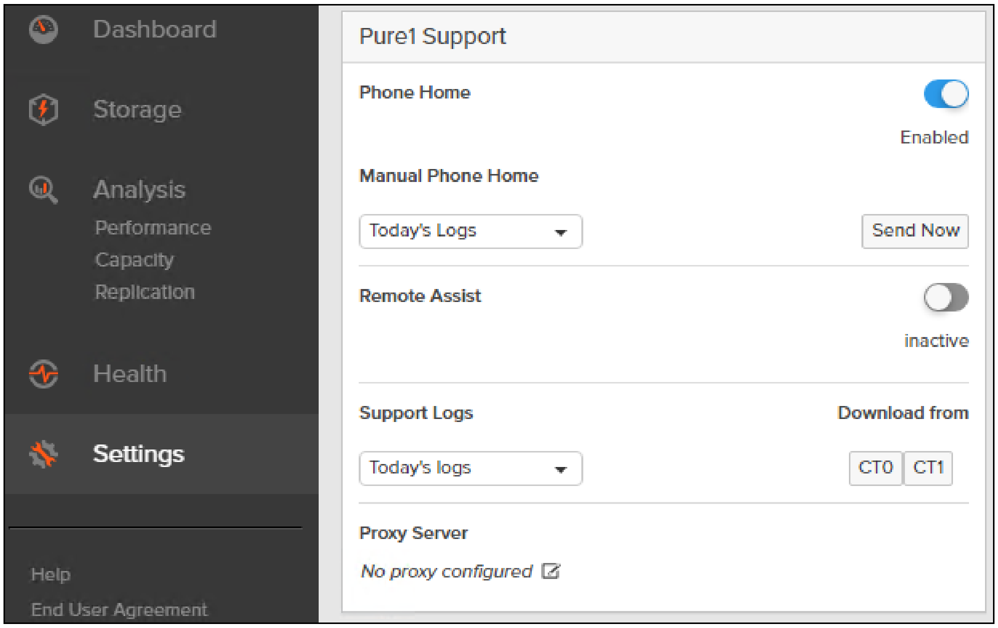

Pure1 Support

Phone Home: Enabled

Manual Phone Home: Today's Logs

Remote Assist: inactive

Support Logs: Today's logs

Download from

Proxy Server: *No proxy configured*

- The phone home facility provides a secure direct link between the array and the Pure Storage Technical Support web site. The link is used to transmit log contents and alert messages to the Pure Storage Support team so that when diagnosis or remedial action is required, complete recent history about array performance and significant events is available. By default, the phone home facility is enabled. If the phone home facility is enabled to send information automatically, Purity transmits log and alert information directly to Pure Storage Support via a secure network connection. Log contents are transmitted hourly and stored at the support web site, enabling detection of array performance and error rate trends. Alerts are reported immediately when they occur so that timely action can be taken.
- Phone home logs can also be sent to Pure Storage Technical support on demand, with options including Today's Logs, Yesterday's Logs, or All Log History.

  The Remote Assist section displays the remote assist status as " Connected" or " Disconnected" . By default, remote assist is disconnected. A connected remote assist status means that a remote assist session has been opened, allowing Pure Storage Support to connect to the array. Disconnect the remote assist session to close the session.
- The Support Logs section allows you to download the Purity log contents of the specified controller to the current administrative workstation. Purity continuously logs a variety of array activities, including performance summaries, hardware and operating status reports, and administrative actions.

Configure DNS Server IP Addresses

To configure the DNS server IP addresses, follow these steps:

1. Select Settings > Network.

2. In the DNS section, hover over the domain name and click the pencil icon. The Edit DNS dialog box appears.

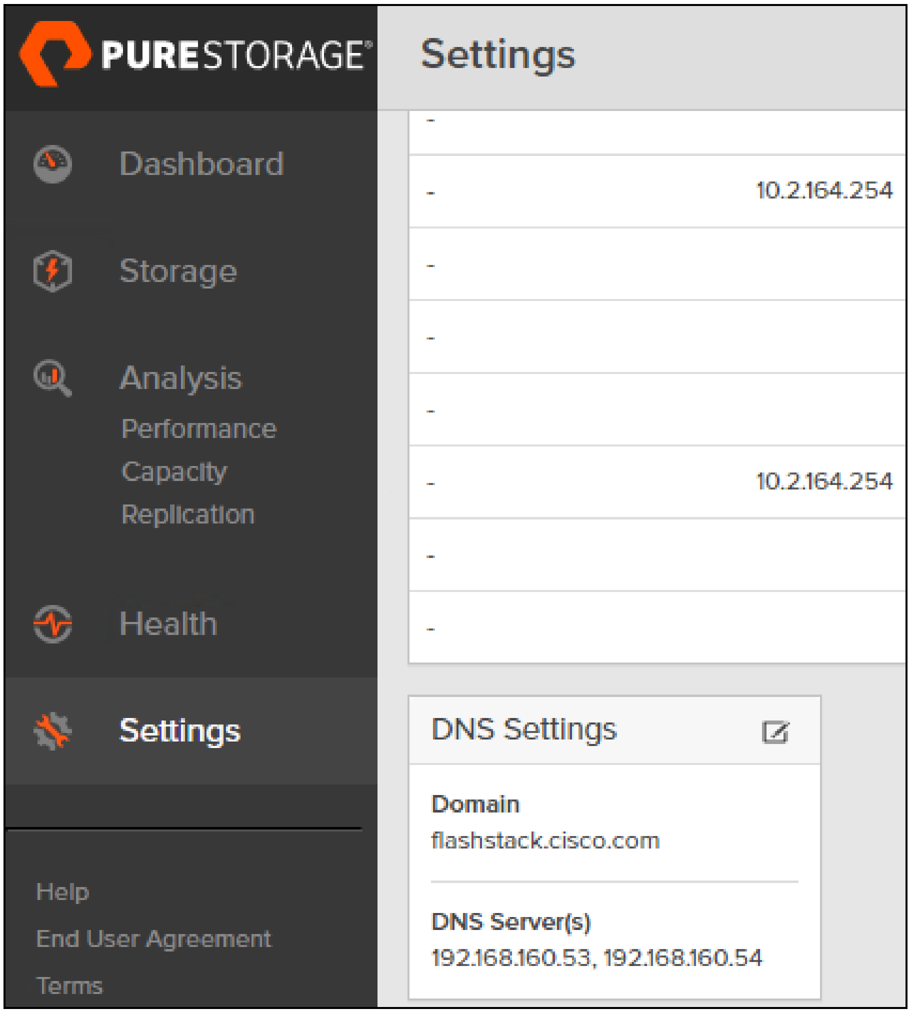

DNS Settings

Domain
flashstack.cisco.com

DNS Server(s)
192.168.160.53, 192.168.160.54

3. Complete the following fields:
   a. Domain: Specify the domain suffix to be appended by the array when doing DNS lookups.
   b. NS#: Specify up to three DNS server IP addresses for Purity to use to resolve hostnames to IP addresses. Enter one IP address in each DNS# field. Purity queries the DNS servers in the order that the IP addresses are listed.
4. Click Save.

Directory Service

The Directory Service manages the integration of FlashArray with an existing directory service. When the Directory Service sub-view is configured and enabled, the FlashArray leverages a directory service to perform user account and permission level searches. Configuring directory services is OPTIONAL.

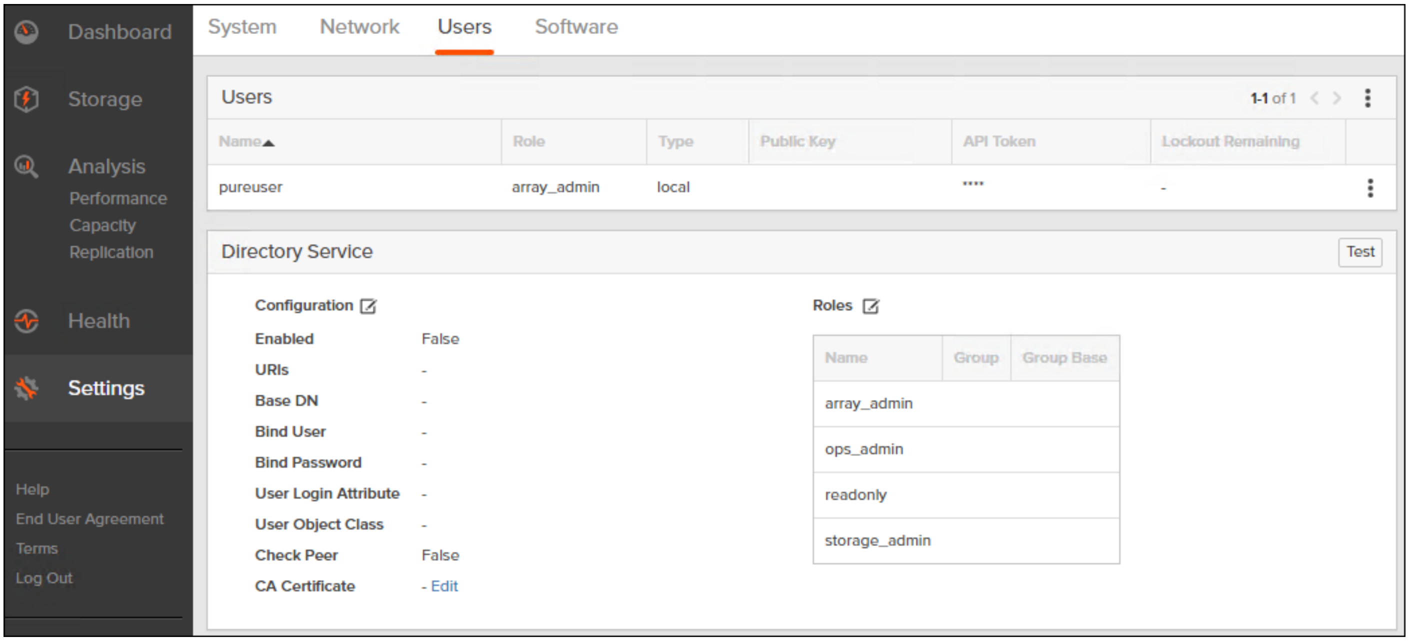

Users

Name▲	Role	Type	Public Key	API Token	Lockout Remaining
pureuser	array_admin	local		****	-

Directory Service

Configuration

Enabled	False
URIs	-
Base DN	-
Bind User	-
Bind Password	-
User Login Attribute	-
User Object Class	-
Check Peer	False
CA Certificate	- Edit

Roles

Name	Group	Group Base
array_admin		
ops_admin		
readonly		
storage_admin		

The FlashArray is delivered with a single local user, named pureuser, with array-wide (Array Admin) permissions.

To support multiple FlashArray users, integrate the array with a directory service, such as Microsoft Active Directory or OpenLDAP.

Role-based access control is achieved by configuring groups in the directory that correspond to the following permission groups (roles) on the array:

- Read Only Group. Read Only users have read-only privilege to run commands that convey the state of the array. Read Only uses cannot alter the state of the array.
- Storage Admin Group. Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering volumes, hosts, and host groups. Storage Admin users cannot perform operations that deal with global and system configurations.
- Array Admin Group. Array Admin users have all the privileges of Storage Admin users, plus the ability to perform array-wide changes. In other words, Array Admin users can perform all FlashArray operations.

To configure the Directory Service, follow these steps:

1. Select Settings > Access > Users.

2. Select the ✎ icon in the Directory Services panel:

   - Enabled: Select the check box to leverage the directory service to perform user account and permission level searches.
   - URI: Enter the comma-separated list of up to 30 URIs of the directory servers. The URI must include a URL scheme (ldap, or ldaps for LDAP over SSL), the hostname, and the domain. You can optionally specify a port. For example, ldap://ad.company.com configures the directory

service with the hostname " ad" in the domain " company.com" while specifying the unencrypted LDAP protocol.

- Base DN: Enter the base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist only of domain components (DCs). For example, for ldap://ad.storage.company.com, the Base DN would be: “DC=storage,DC=company,DC=com”
- Bind User: Username used to bind to and query the directory. For Active Directory, enter the username - often referred to as sAMAccountName or User Logon Name - of the account that is used to perform directory lookups. The username cannot contain the characters " [ ] : ; | = + * ? < > / \ and cannot exceed 20 characters in length. For OpenLDAP, enter the full DN of the user. For example, " CN=John,OU=Users,DC=example,DC=com" .
- Bind Password: Enter the password for the bind user account.
- Group Base: Enter the organizational unit (OU) to the configured groups in the directory tree. The Group Base consists of OUs that, when combined with the base DN attribute and the configured group CNs, complete the full Distinguished Name of each groups. The group base should specify " OU=" for each OU and multiple OUs should be separated by commas. The order of OUs should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: " OU=PureGroups,OU=SANManagers" .
- Array Admin Group: Common Name (CN) of the directory service group containing administrators with full privileges to manage the FlashArray. Array Admin Group administrators have the same privileges as pureuser. The name should be the Common Name of the group without the " CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, " pureadmins,OU=PureStorage" , where pureadmins is the common name of the directory service group.
- Storage Admin Group: Common Name (CN) of the configured directory service group containing administrators with storage related privileges on the FlashArray. The name should be the Common Name of the group without the " CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, " pureusers,OU=PureStorage" , where pureusers is the common name of the directory service group.
- Read Only Group: Common Name (CN) of the configured directory service group containing users with read-only privileges on the FlashArray. The name should be the Common Name of the group without the " CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, " purereadonly,OU=PureStorage" , where purereadonly is the common name of the directory service group.
- Check Peer: Select the check box to validate the authenticity of the directory servers using the CA Certificate. If you enable Check Peer, you must provide a CA Certificate.
- CA Certificate: Enter the certificate of the issuing certificate authority. Only one certificate can be configured at a time, so the same certificate authority should be the issuer of all directory server certificates. The certificate must be PEM formatted (Base64 encoded) and include the " -----BEGIN CERTIFICATE-----" and " -----END CERTIFICATE-----" lines. The certificate cannot exceed 3000 characters in total length.

3. Click Save.

4. Click Test to test the configuration settings. The LDAP Test Results pop-up window appears. Green squares represent successful checks. Red squares represent failed checks.

SSL Certificate

Self-Signed Certificate

Purity creates a self-signed certificate and private key when you start the system for the first time. The SSL Certificate sub-view allows you to view and change certificate attributes, create a new self-signed certificate, construct certificate signing requests, import certificates and private keys, and export certificates.

Creating a self-signed certificate replaces the current certificate. When you create a self-signed certificate, include any attribute changes, specify the validity period of the new certificate, and optionally generate a new private key.

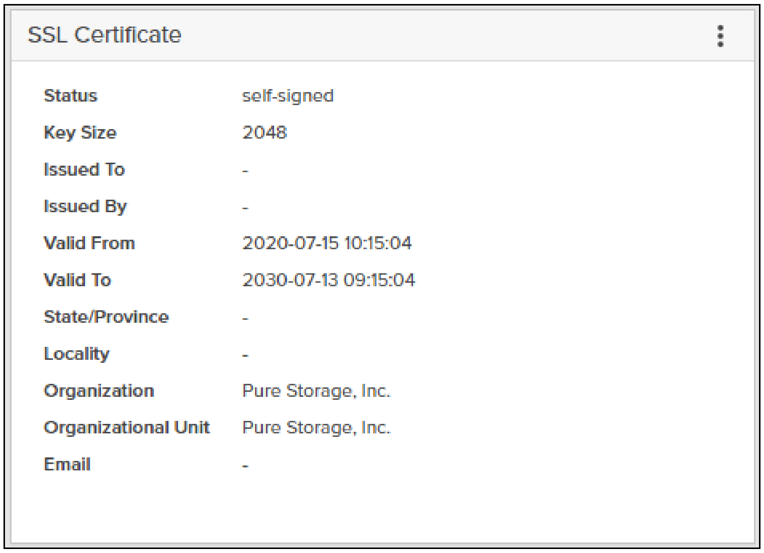

SSL Certificate

Status	self-signed
Key Size	2048
Issued To	-
Issued By	-
Valid From	2020-07-15 10:15:04
Valid To	2030-07-13 09:15:04
State/Province	-
Locality	-
Organization	Pure Storage, Inc.
Organizational Unit	Pure Storage, Inc.
Email	-

When you create the self-signed certificate, you can generate a private key and specify a different key size. If you do not generate a private key, the new certificate uses the existing key.

You can change the validity period of the new self-signed certificate. By default, self-signed certificates are valid for 3650 days

CA-Signed Certificate

Certificate authorities (CA) are third party entities outside the organization that issue certificates. To obtain a CA certificate, you must first construct a certificate signing request (CSR) on the array.

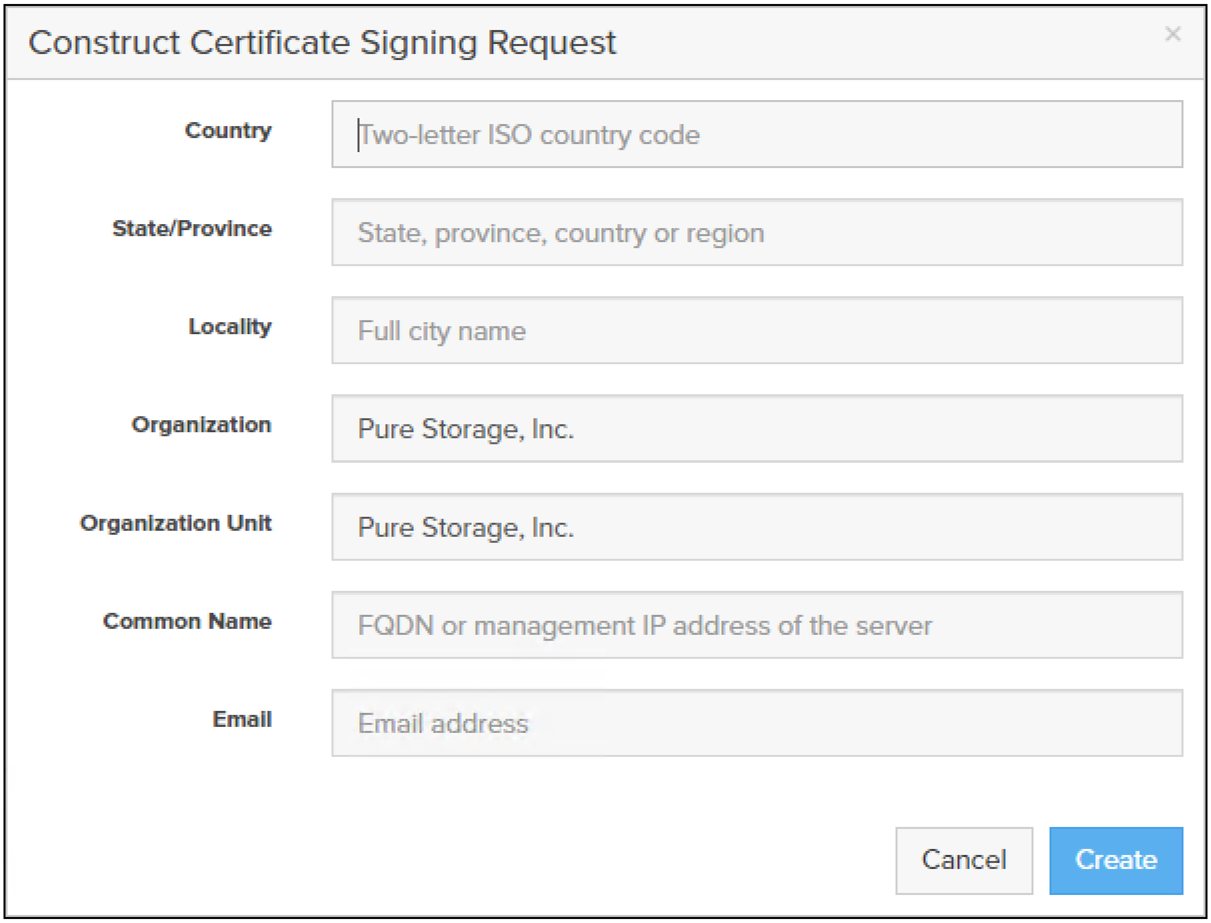

Construct Certificate Signing Request

Country	Two-letter ISO country code
State/Province	State, province, country or region
Locality	Full city name
Organization	Pure Storage, Inc.
Organization Unit	Pure Storage, Inc.
Common Name	FQDN or management IP address of the server
Email	Email address

The CSR represents a block of encrypted data specific to your organization. You can change the certificate attributes when you construct the CSR; otherwise, Purity will reuse the attributes of the current certificate (self-signed or imported) to construct the new one. Note that the certificate attribute changes will only be visible after you import the signed certificate from the CA.

Send the CSR to a certificate authority for signing. The certificate authority returns the SSL certificate for you to import. Verify that the signed certificate is PEM formatted (Base64 encoded), includes the " -----BEGIN CERTIFICATE-----" and " -----END CERTIFICATE-----" lines, and does not exceed 3000 characters in total length. When you import the certificate, also import the intermediate certificate if it is not bundled with the CA certificate.

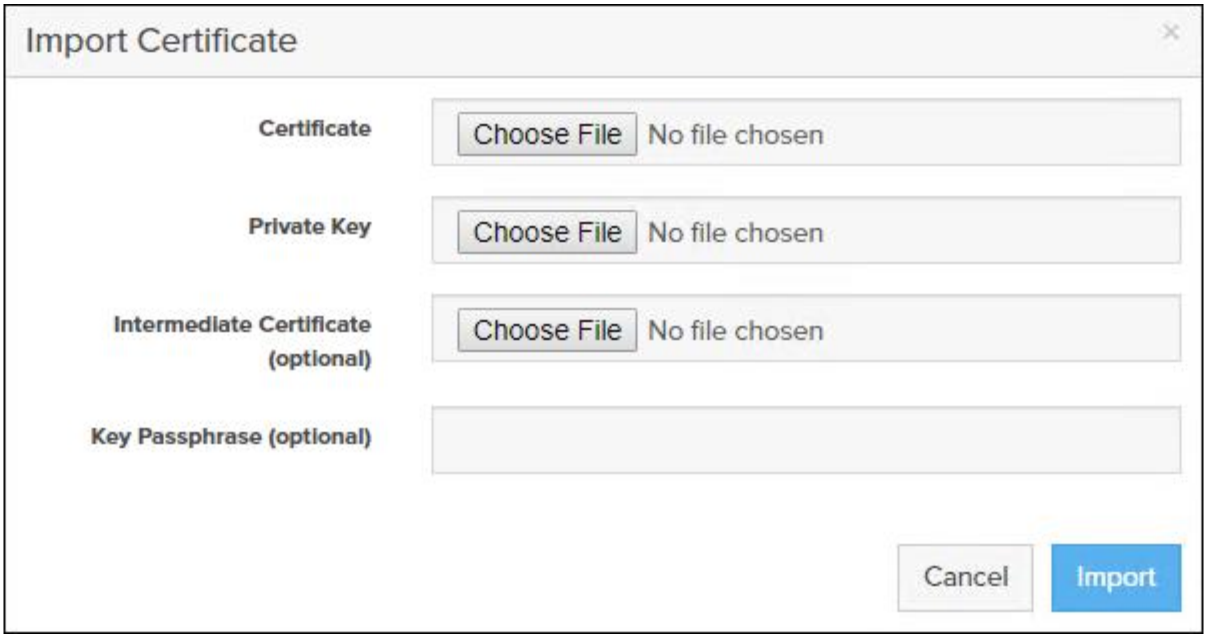

Import Certificate

Certificate: Choose File No file chosen

Private Key: Choose File No file chosen

Intermediate Certificate (optional): Choose File No file chosen

Key Passphrase (optional):

If the certificate is signed with the CSR that was constructed on the current array and you did not change the private key, you do not need to import the key. However, if the CSR was not constructed on the current array or if the private key has changed since you constructed the CSR, you must import the private key. If the private key is encrypted, also specify the passphrase.

Note: If FC-NVMe is being implemented, the FC ports personality on the FlashArray need to be converted to nvme-fc from the default sccsi-fc. In this design we have used two scsi-fc and two nvme-fc ports to support both SCSI and NVMe over Fibre Channel. The ports can be converted to nvme-fc with the help of Pure support.

Cisco UCS Configuration

The following procedures describe how to configure the Cisco UCS domain for use in a base FlashStack environment. This procedure assumes you're using Cisco UCS Fabric Interconnects running in Intersight managed mode.

Figure 5. Cisco UCS Configuration Workflow

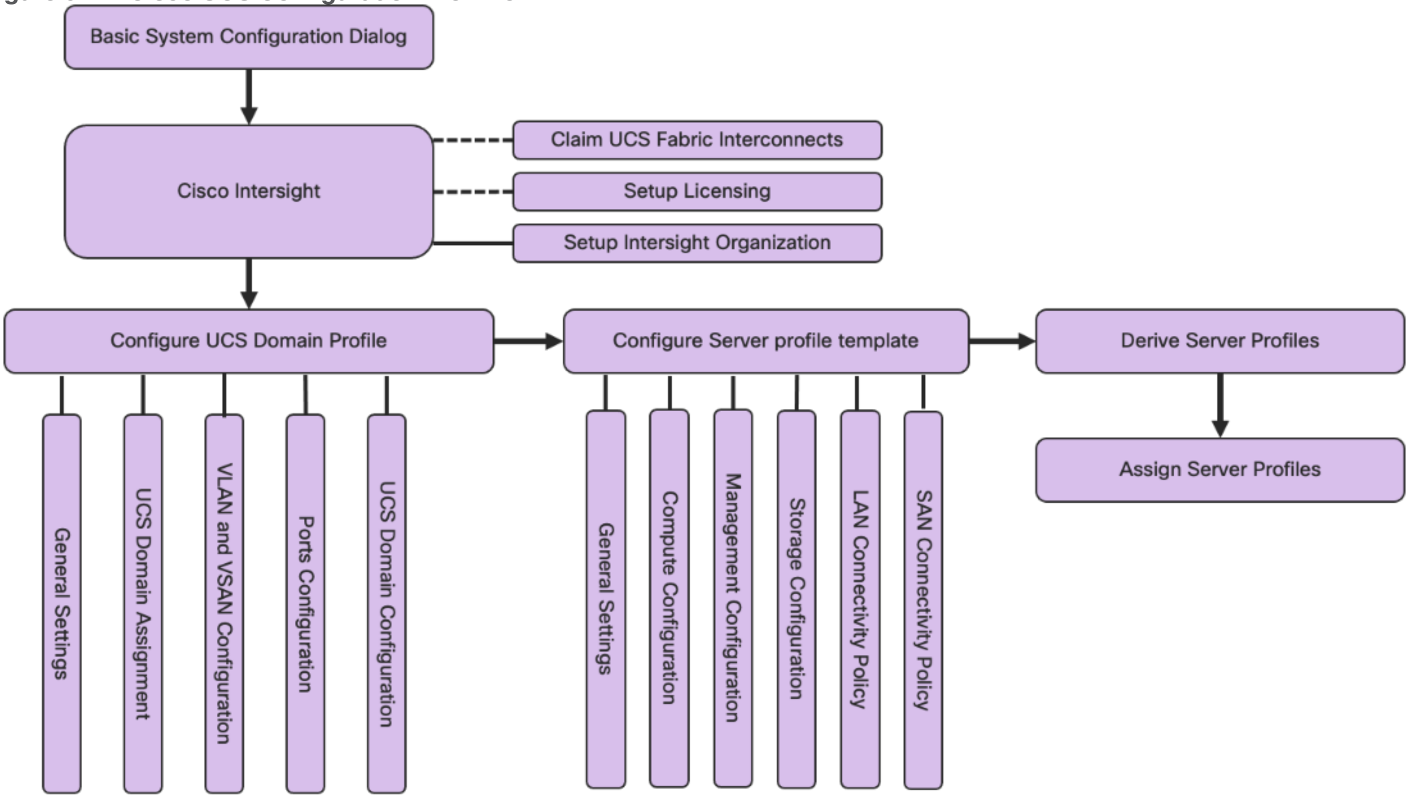

Physical Connectivity

Physical cabling should be completed by following the diagram and table references in section FlashStack Cabling.

Cisco Intersight Managed Mode Configuration

The Cisco Intersight™ platform is a management solution delivered as a service with embedded analytics for Cisco® and third-party IT infrastructures. The Cisco Intersight managed mode (also referred to as Cisco IMM or Intersight managed mode) is a new architecture that manages Cisco Unified Computing System™ (Cisco UCS®) fabric interconnect-attached systems through a Redfish-based standard model. Cisco Intersight managed mode standardizes both policy and operation management for Cisco UCSX X210c M6 compute nodes used in this deployment guide.

Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

The Cisco UCS fabric interconnects need to be set up to support Cisco Intersight managed mode. When converting an existing pair of Cisco UCS fabric interconnects from Cisco UCS Manager (UCSM) mode to Intersight Mange Mode (IMM), first erase the configuration and reboot your system.

Note: Converting fabric interconnects to Cisco Intersight managed mode is a disruptive process, and configuration information will be lost. Customers are encouraged to make a backup of their existing configuration. If a UCS software version that supports Intersight Managed Mode (4.1(3) or later) is already installed on Cisco UCS Fabric Interconnects, do not upgrade the software to a recommended recent release using Cisco UCS Manager. The software upgrade will be performed using Cisco Intersight to make sure Cisco UCS X-series firmware is part of the software upgrade.

This section provides the detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlashStack environment. These steps are necessary to provision the Cisco UCS Compute nodes and should be followed precisely to avoid improper configuration.

Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlashStack environment in Intersight managed mode, follow these steps:

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

2. Power on the Fabric Interconnect.

3. Power-on self-test messages will be displayed as the Fabric Interconnect boots.

4. When the unconfigured system boots, it prompts you for the setup method to be used. Enter console to continue the initial setup using the console CLI.

5. Enter the "intersight" as the management mode for the Fabric Interconnect:

   - Intersight to manage the Fabric Interconnect through Cisco Intersight.
   - ucsm to manage the Fabric Interconnect through Cisco UCS Manager.

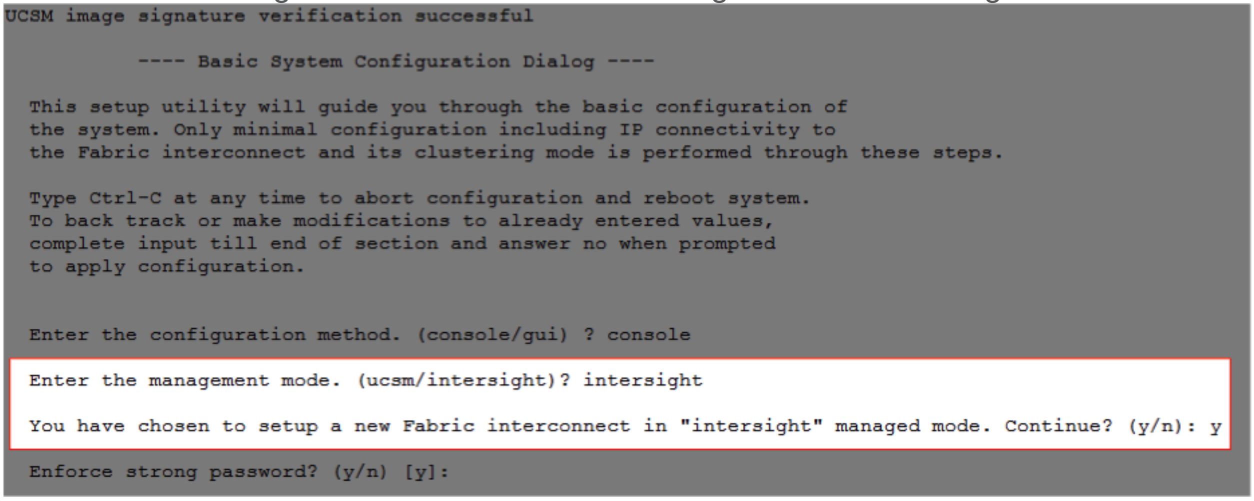

```
UCSM image signature verification successful

          ---- Basic System Configuration Dialog ----

  This setup utility will guide you through the basic configuration of
  the system. Only minimal configuration including IP connectivity to
  the Fabric interconnect and its clustering mode is performed through these steps.

  Type Ctrl-C at any time to abort configuration and reboot system.
  To back track or make modifications to already entered values,
  complete input till end of section and answer no when prompted
  to apply configuration.


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

  Enter the management mode. (ucsm/intersight)? intersight

  You have chosen to setup a new Fabric interconnect in "intersight" managed mode. Continue? (y/n): y

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

6. Enter y to confirm that you want to continue the initial setup.

7. To use a strong password, enter y.

8. Enter the password for the admin account.

9. To confirm, re-enter the password for the admin account.

10. Enter yes to continue the initial setup for a cluster configuration.

11. Enter the Fabric Interconnect fabric (either A or B).

12. Enter the system name.

13. Enter the IPv4 or IPv6 address for the management port of the Fabric Interconnect.

Note: If you enter an IPv4 address, you will be prompted to enter an IPv4 subnet mask. If you enter an IPv6 address, you will be prompted to enter an IPv6 network prefix.

14. Enter the respective IPv4 subnet mask or IPv6 network prefix, then press Enter.

Note: You are prompted for an IPv4 or IPv6 address for the default gateway, depending on the address type you entered for the management port of the Fabric Interconnect.

15. Enter either of the following:

- IPv4 address of the default gateway
- IPv6 address of the default gateway

16. Enter the IPv4 or IPv6 address for the DNS server.

Note: The address type must be the same as the address type of the management port of the Fabric Interconnect.

17. Enter yes if you want to specify the default Domain name, or no if you do not.

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

Enter the management mode. (ucsm/intersight)? intersight

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

You have chosen to setup a new Fabric interconnect in “ucsm” managed mode. Continue? (y/n): y

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

Enter the password for "admin": <password>
Confirm the password for "admin": <password>

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: y

Enter the switch fabric (A/B) []: A

Enter the system name:  <ucs-cluster-name>

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

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

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

Cluster IPv4 address : <ucs-cluster-ip>

Configure the DNS Server IP address? (yes/no) [n]: y

  DNS IP address : <dns-server-1-ip>

Configure the default domain name? (yes/no) [n]: y
```

```
    Default domain name : <ad-dns-domain-name>

  Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

18. 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 FlashStack environment, follow these steps:

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

2. Power up the Fabric Interconnect.

3. When the unconfigured system boots, it prompts you for the setup method to be used. Enter console to continue the initial setup using the console CLI.

```
  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
```

4. Wait for the login prompt for UCS Fabric Interconnect B before proceeding to the next section.

Set up Cisco Intersight Account

In this step, using the unique device information for the Cisco UCS, you set up a new Cisco Intersight account. Customers also can choose to add the Cisco UCS devices set up for Cisco Intersight managed mode to an existing Cisco Intersight account; however, that procedure is not covered in this document.

Claim a device

After completing the initial configuration for the fabric interconnects, log into Fabric Interconnect A using your web browser to capture the Cisco Intersight connectivity information. To claim a device, follow these steps:

1. Use the management IP address of Fabric Interconnect A to access the device from a web browser and the previously configured admin password to log into the device.

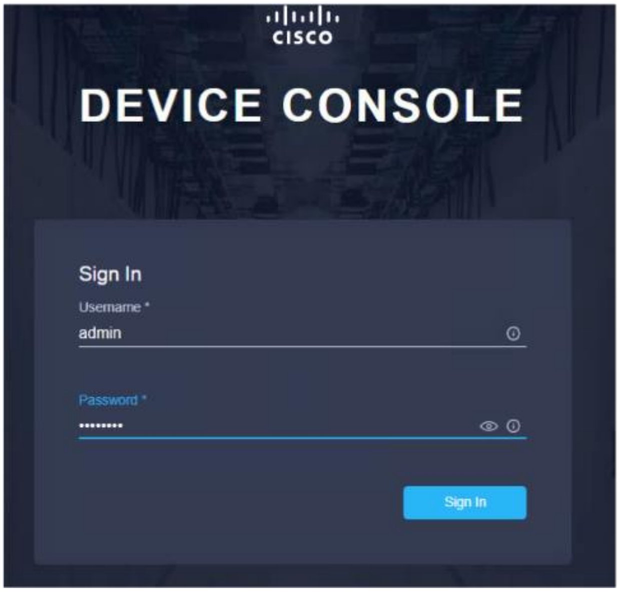

Username *
admin

Password *
••••••••

2. Under DEVICE CONNECTOR, you should see the current device status as “Not claimed.” Note, or copy, the Device ID and Claim Code information to use to set up a new Cisco Intersight account.

Note: The Device ID and Claim Code information can also be used to claim the Cisco UCS devices set up with Cisco Intersight managed mode in an existing Cisco Intersight account.

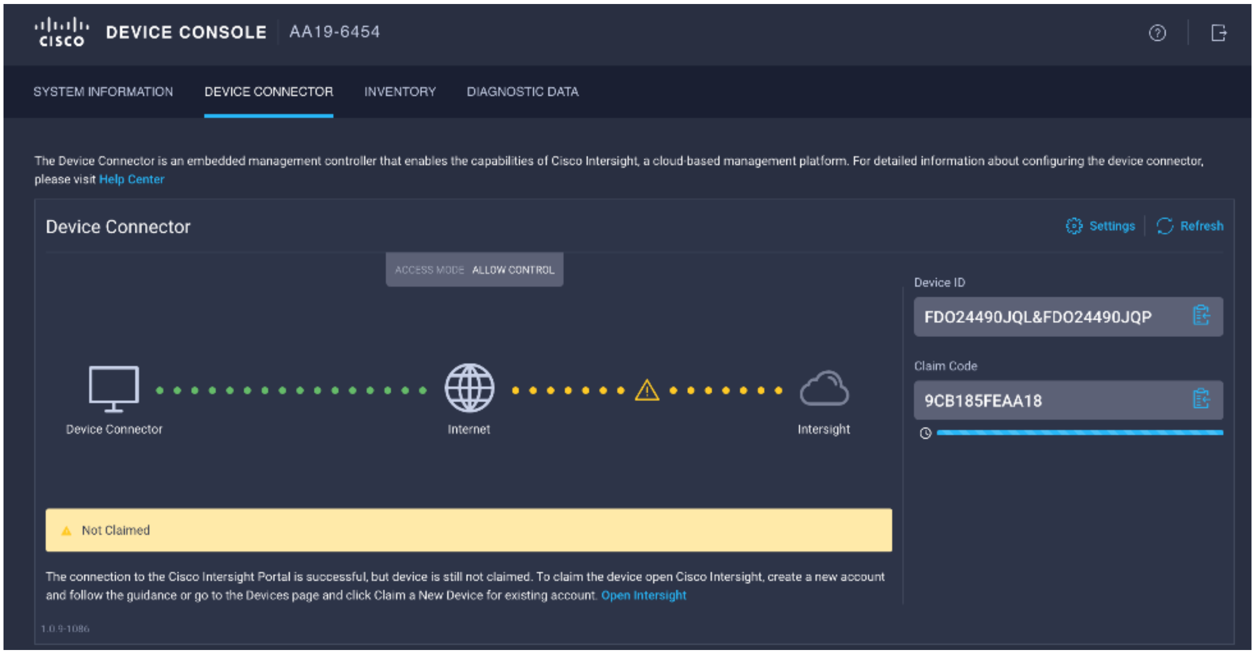

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

Device ID
FDO24490JQL&FDO24490JQP

Claim Code
9CB185FEAA18

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

Create a new Cisco Intersight account

To create a new Cisco Intersight account, follow these steps:

1. Go to https://www.intersight.com and click “Don't have an Intersight Account? Create an account.”

2. Accept the end User license Agreement.

3. Click Next.

4. Click Create.

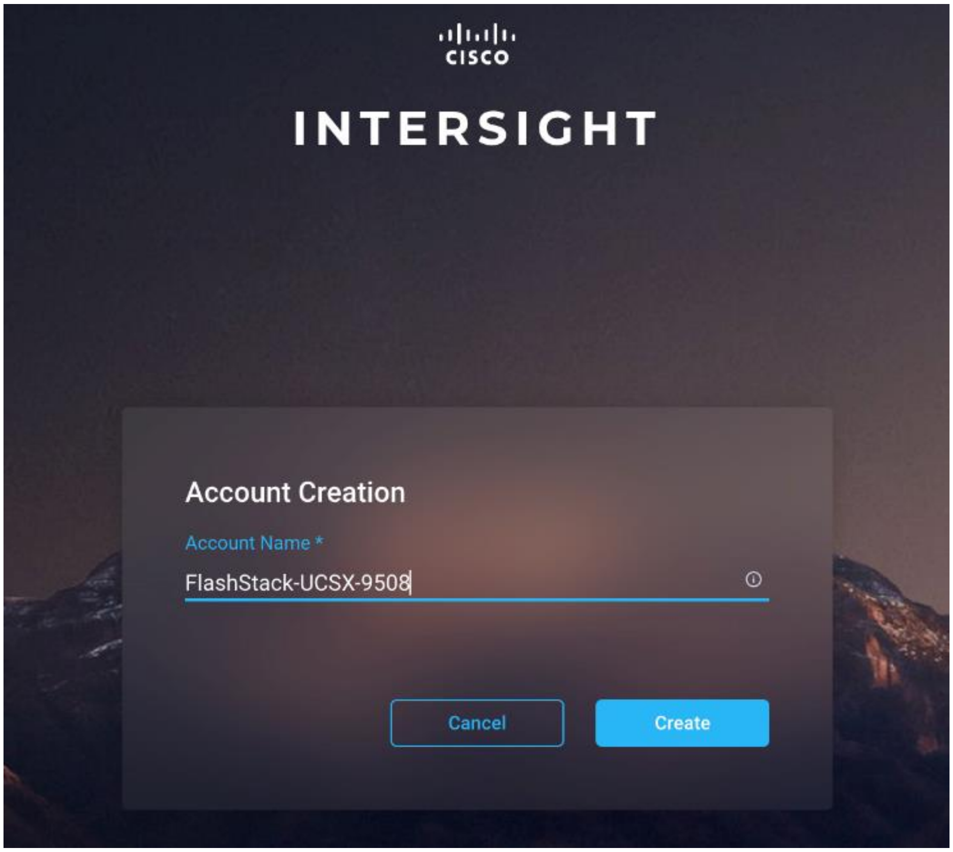

INTERSIGHT

Account Creation

Account Name *

FlashStack-UCSX-9508

5. After the account has been created successfully, click “Go To Intersight.”

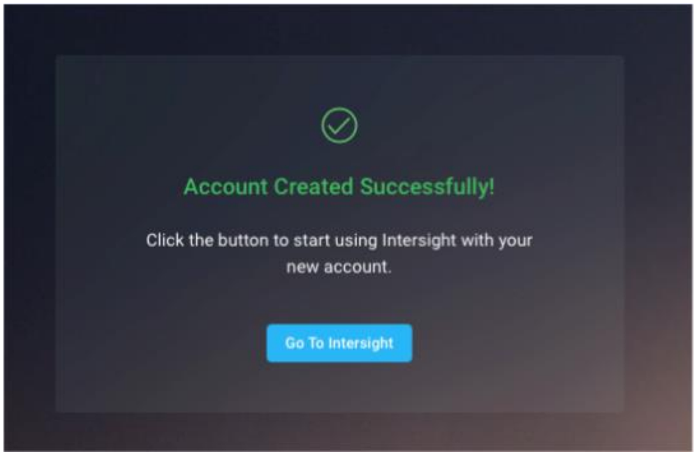

Account Created Successfully!

Click the button to start using Intersight with your new account.

You will see a screen with your Cisco Intersight account as shown below:

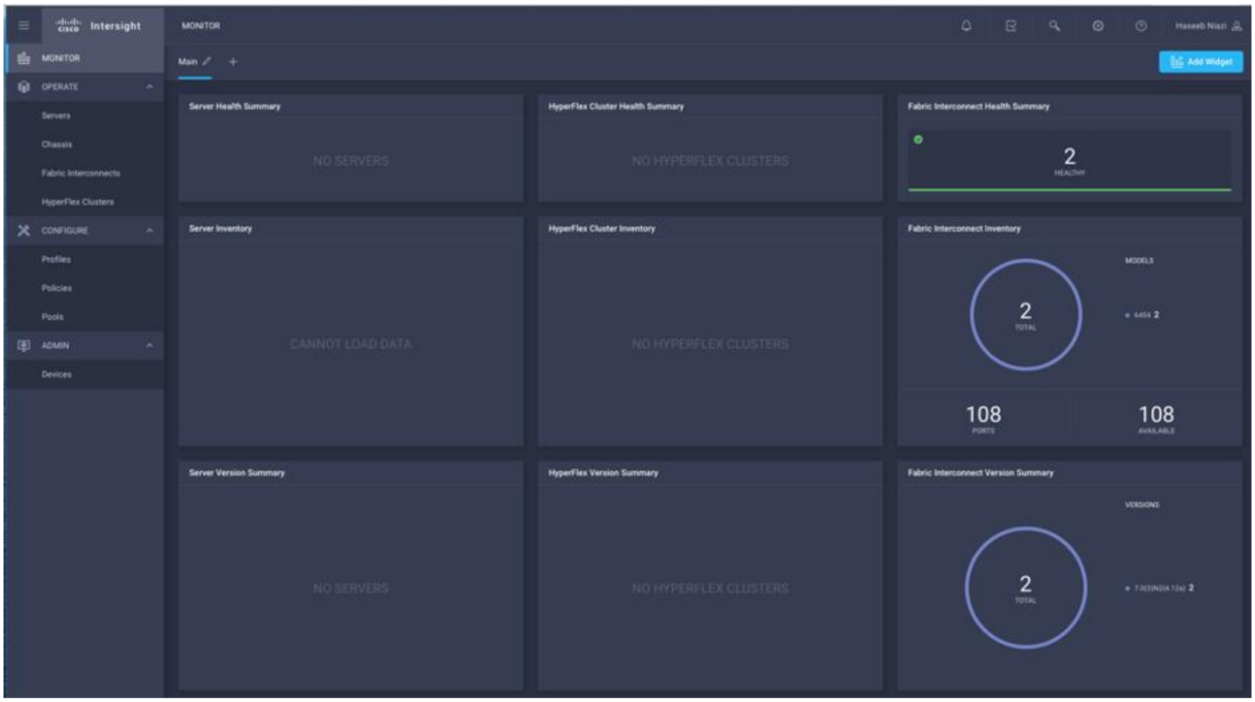

Claim UCS Fabric Interconnects to Cisco Intersight

To claim a new target, follow these steps:

1. Log into Intersight with the Account Administrator, Device Administrator, or Device Technician privileges.

2. Navigate to ADMIN > Targets > Claim a New Target.

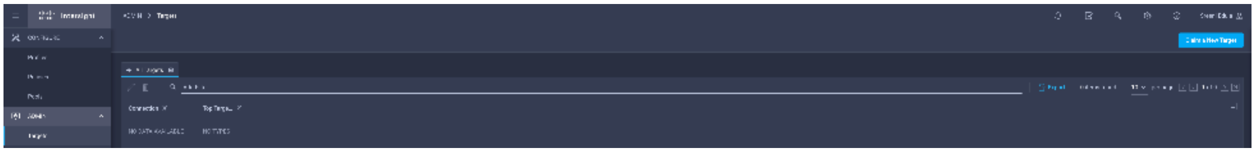

3. Choose Available for Claiming and select the Cisco UCS Domain (Intersight managed) as target type you want to claim. Click Start.

4. Enter the Device ID and Claim Code details captured from Device Connector tab earlier and click Claim to complete the claiming process.

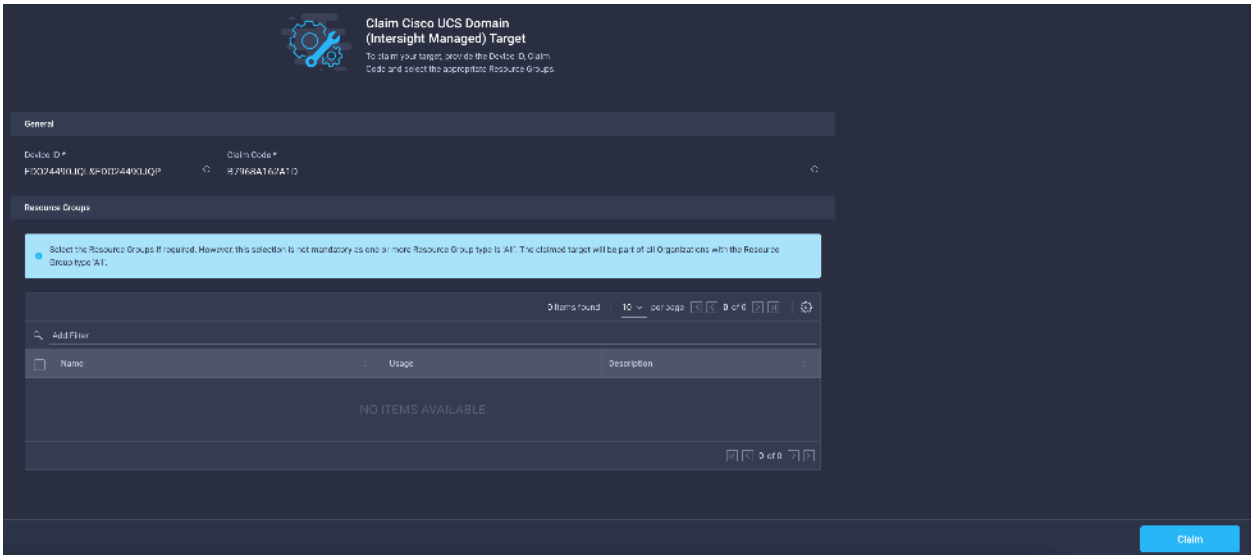

Claim Cisco UCS Domain (Intersight Managed) Target

To claim your target, provide the Device ID, Claim Code and select the appropriate Resource Groups.

General

Device ID * FDO24490JQL&FDO24490JQP

Claim Code * B7968A162A1D

Resource Groups

Select the Resource Groups if required. However, this selection is not mandatory as one or more Resource Group type is 'All'. The claimed target will be part of all Organizations with the Resource Group type 'All'.

0 items found 10 per page 0 of 0

□ Name	Usage	Description

NO ITEMS AVAILABLE

Verify addition of Cisco UCS Fabric Interconnects to Cisco Intersight

To verify that the Cisco UCS fabric interconnects are added to your account in Cisco Intersight, follow these steps:

1. Go to the web GUI of the Cisco UCS fabric interconnect and click Refresh.

The fabric interconnect status should now be set to Claimed, as shown below:

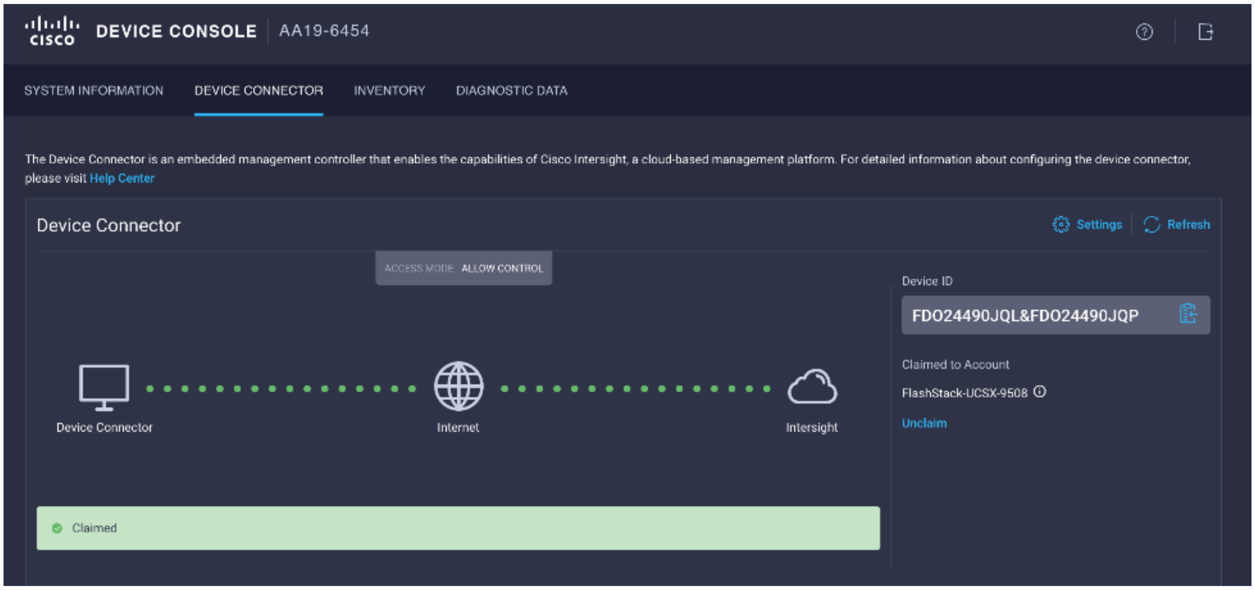

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

Device Connector

ACCESS MODE ALLOW CONTROL

Device ID FDO24490JQL&FDO24490JQP

Claimed to Account FlashStack-UCSX-9508

Claimed

The fabric interconnect now is listed as a Claimed Target, as shown below:

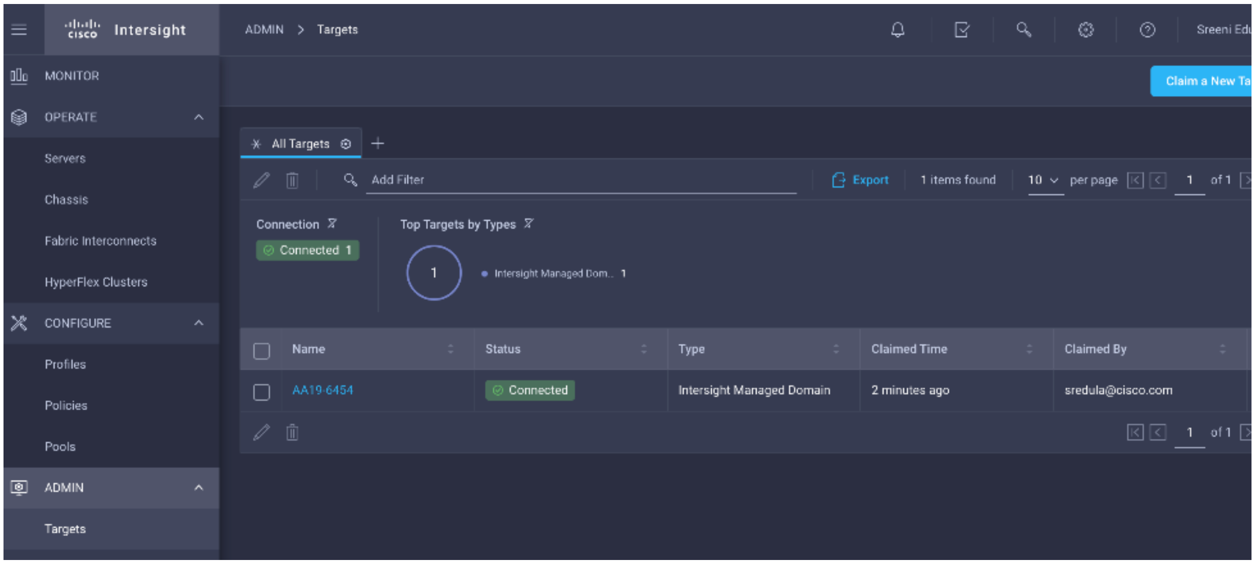

Name	Status	Type	Claimed Time	Claimed By
AA19-6454	Connected	Intersight Managed Domain	2 minutes ago	sredula@cisco.com

Set up licensing

When setting up a new Cisco Intersight account (as discussed in this document), the account needs to be enabled for Cisco Smart Software Licensing. To set up licensing, follow these steps:

1. Associate the Cisco Intersight account with Cisco Smart Licensing by following these steps:

   a. Log into the Cisco Smart Licensing portal: https://software.cisco.com/software/csws/ws/platform/home?locale=en_US#module/SmartLicensing.

   b. Select the correct virtual account.

   c. Under Inventory > General, generate a new token for product registration.

   d. Copy this newly created token.

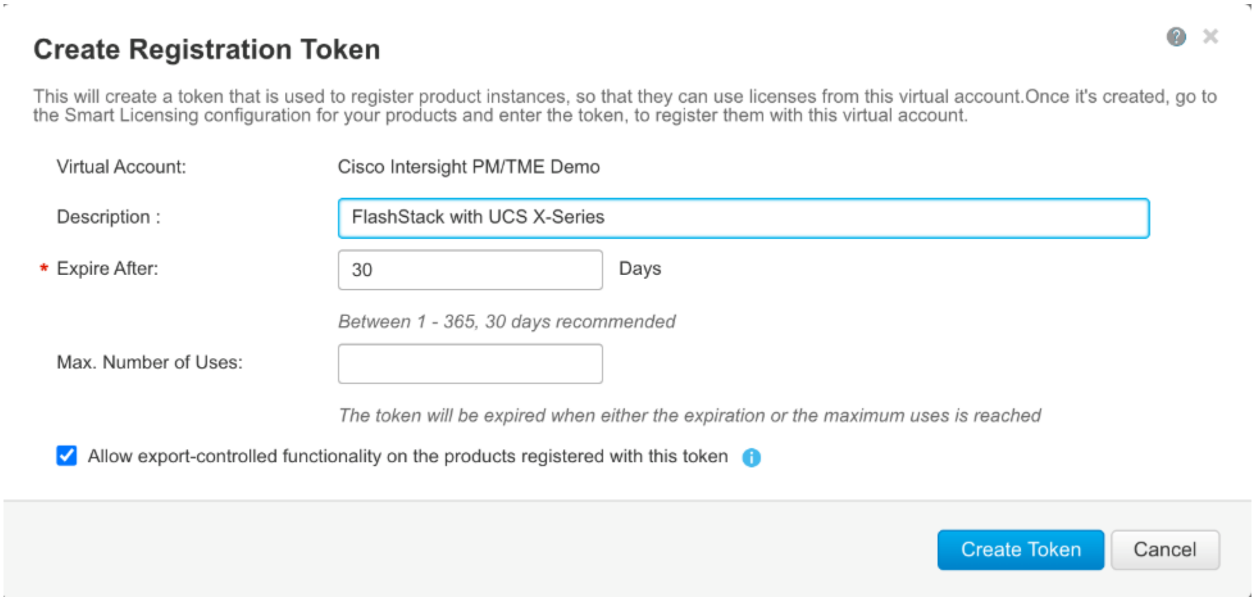

Create Registration Token

This will create a token that is used to register product instances, so that they can use licenses from this virtual account.Once it's created, go to the Smart Licensing configuration for your products and enter the token, to register them with this virtual account.

Virtual Account: Cisco Intersight PM/TME Demo

Description : FlashStack with UCS X-Series

* Expire After: 30 Days

Between 1 - 365, 30 days recommended

Max. Number of Uses:

The token will be expired when either the expiration or the maximum uses is reached

☑ Allow export-controlled functionality on the products registered with this token

2. With the Cisco Intersight account associated with Cisco Smart Licensing, log into the Cisco Intersight portal and click Settings (the gear icon) in the top-right corner. Choose Licensing.

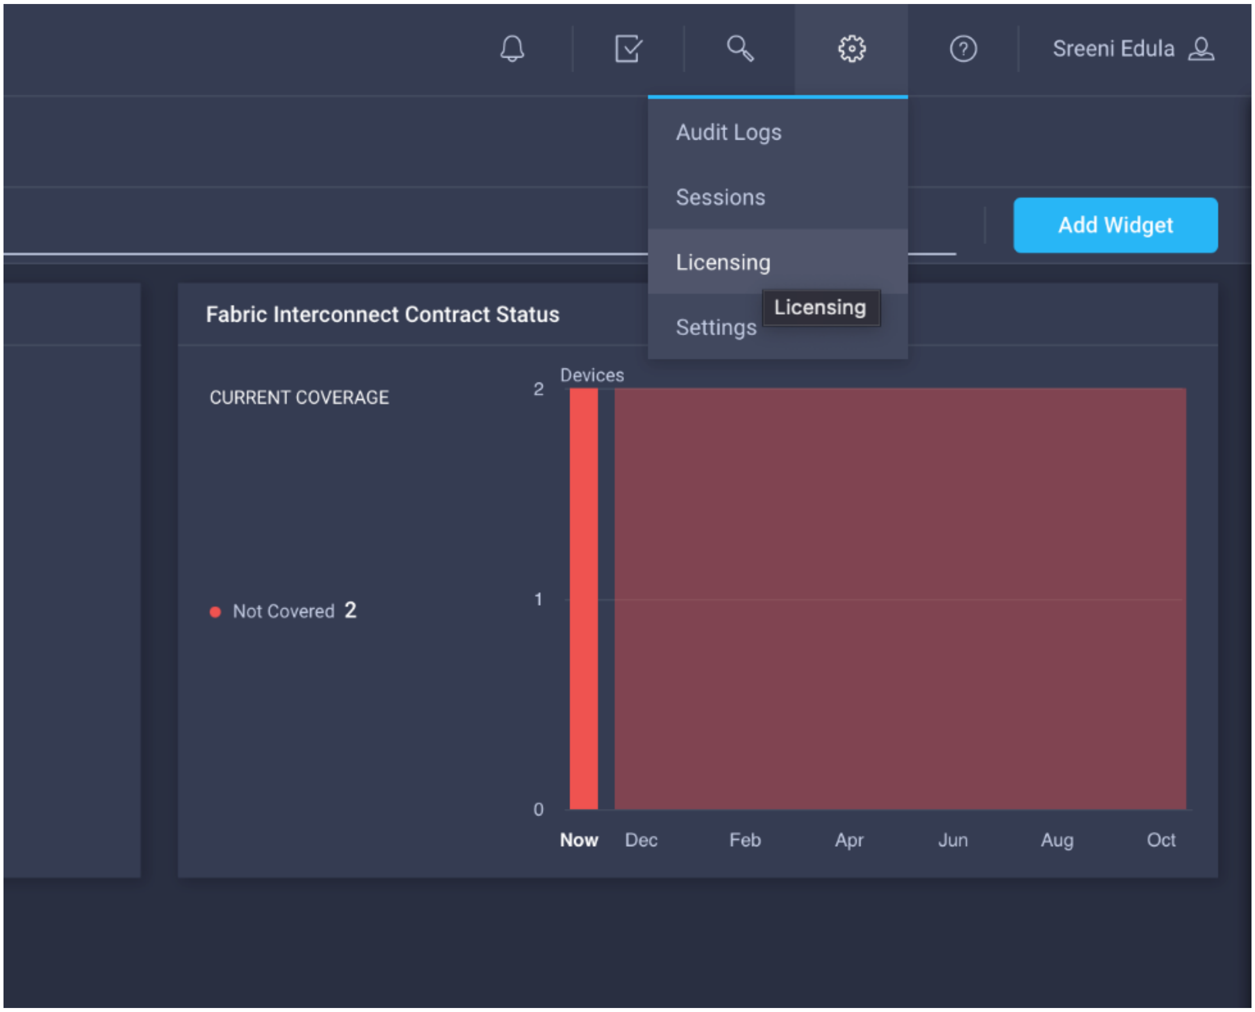

3. Under Cisco Intersight > Licensing, click Register.

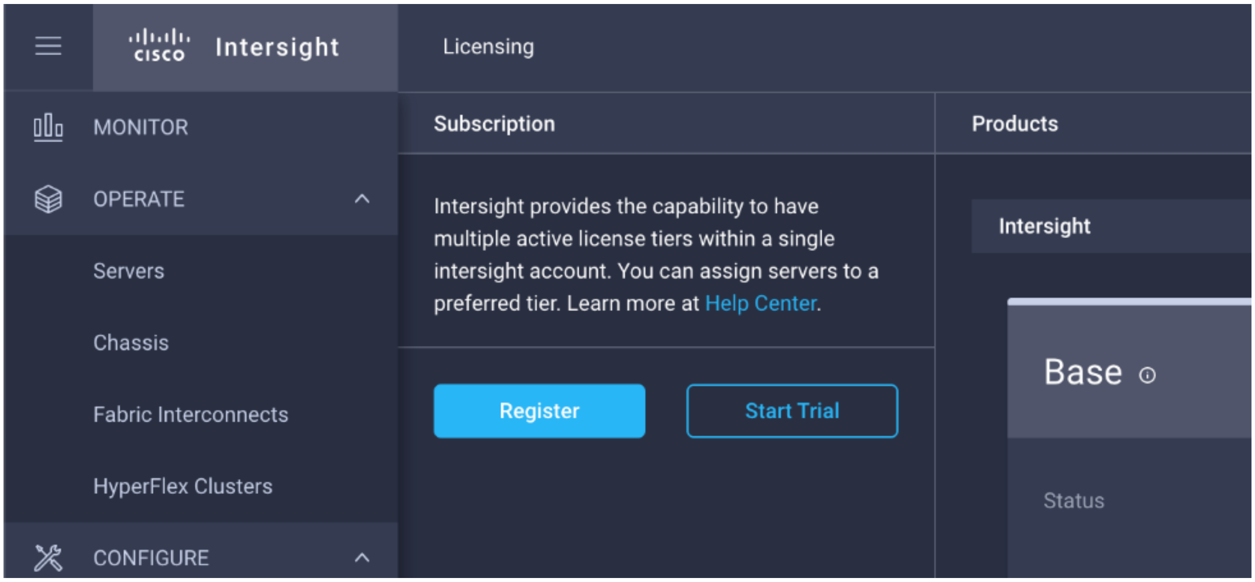

Intersight provides the capability to have multiple active license tiers within a single intersight account. You can assign servers to a preferred tier. Learn more at Help Center.

4. Enter the copied token from the Cisco Smart Licensing portal.

5. Click Next.

6. Select the appropriate default Tier.

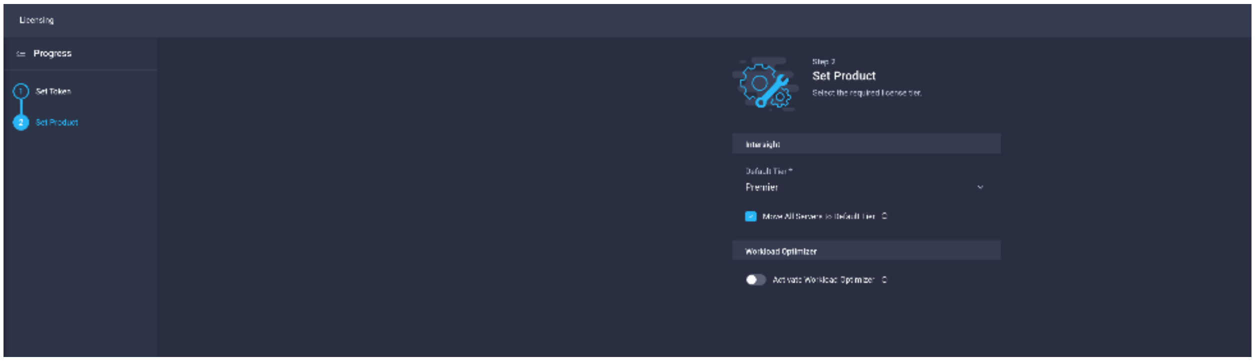

Step 2

Set Product

Select the required license tier.

Intersight

Default Tier *

Premier

☑ Move All Servers to Default Tier

Workload Optimizer

○ Activate Workload Optimizer

7. Click Register and wait for registration to process.

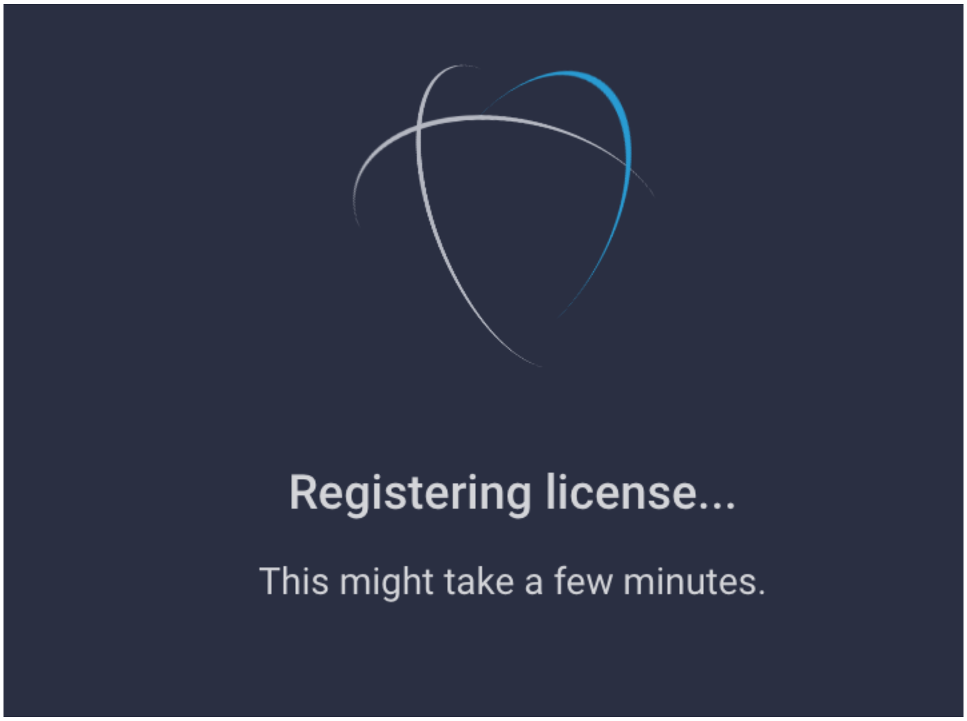

Registering license...

This might take a few minutes.

When the registration is successful, the information about the associated Cisco Smart account is displayed, as shown below:

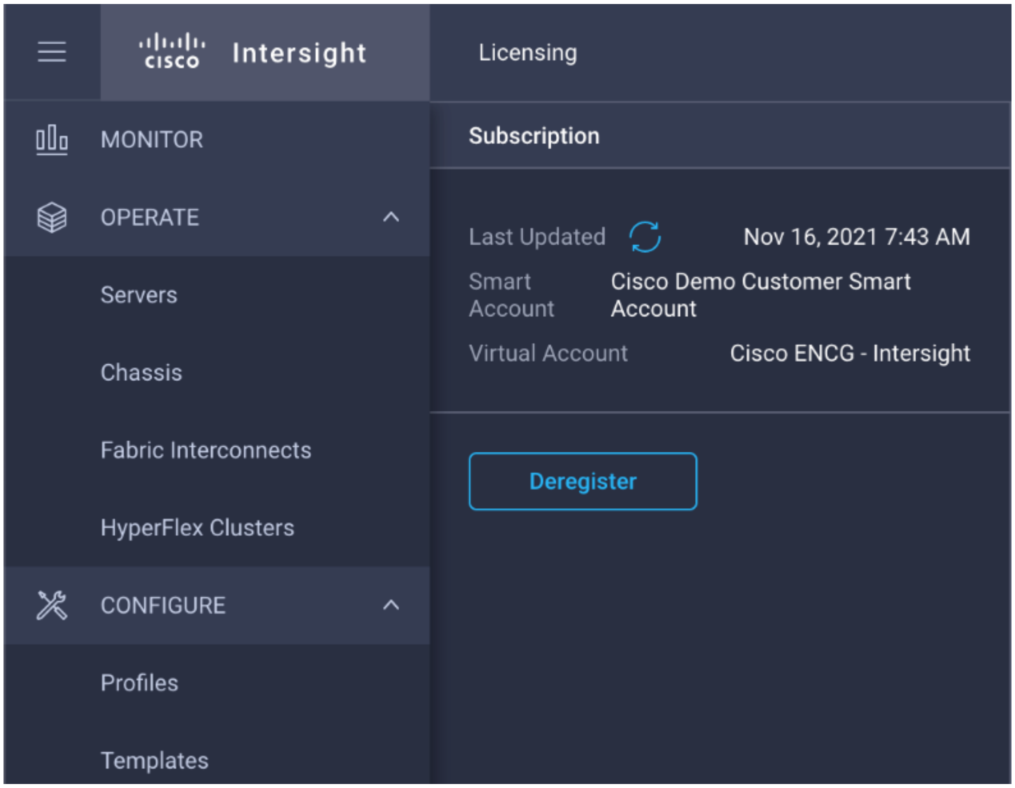

Subscription

Last Updated Nov 16, 2021 7:43 AM

Smart Account Cisco Demo Customer Smart Account

Virtual Account Cisco ENCG - Intersight

8. If default licensing tier is not set to desired level, change it by clicking Set Default Tier on right side of the page. For Cisco Intersight managed mode, the default tier needs to be changed to Essential or a higher tier.

Note: The Default Tier was set to Premier during registration in the earlier step, however describing the procedure to set it up if it's missed during license registration.

9. Select the tier supported by your Smart License.

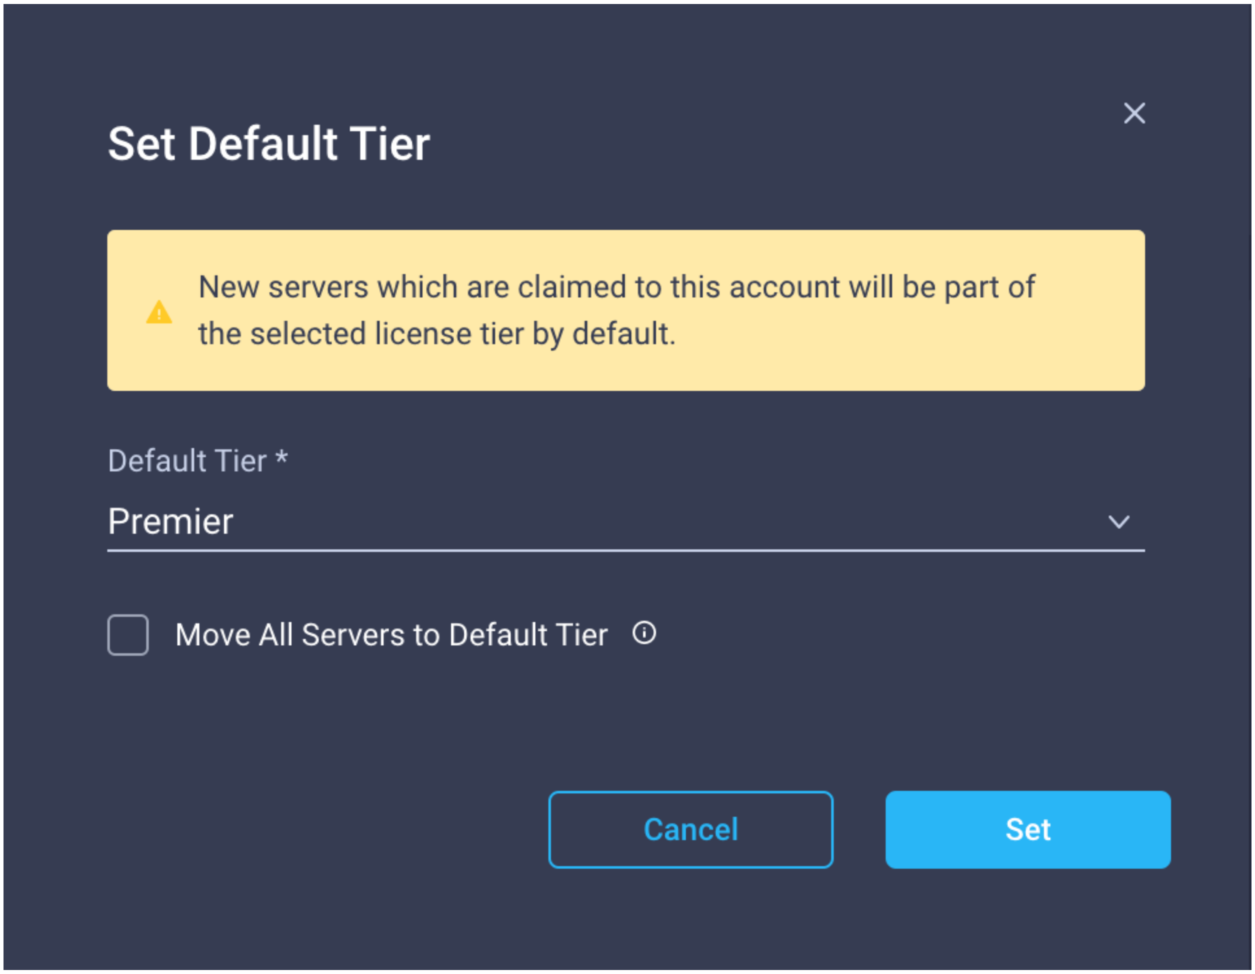

Set Default Tier

⚠ New servers which are claimed to this account will be part of the selected license tier by default.

Default Tier *

Premier

□ Move All Servers to Default Tier

Note: In this deployment, the default license tier is set to Premier.

Setup Intersight Organization

You need to define all Cisco Intersight managed mode configurations for Cisco UCS, including policies, under an organization. To define a new organization, follow these steps:

Create Resource Group

Note: Optionally, create a resource group to place the claimed targets. The Default Resource Group automatically groups all the resources of a User Account. A user with Account Administrator privilege can remove resources from the Default Resource Group when required.

1. Log into the Cisco Intersight portal.

2. Click Settings (the gear icon).

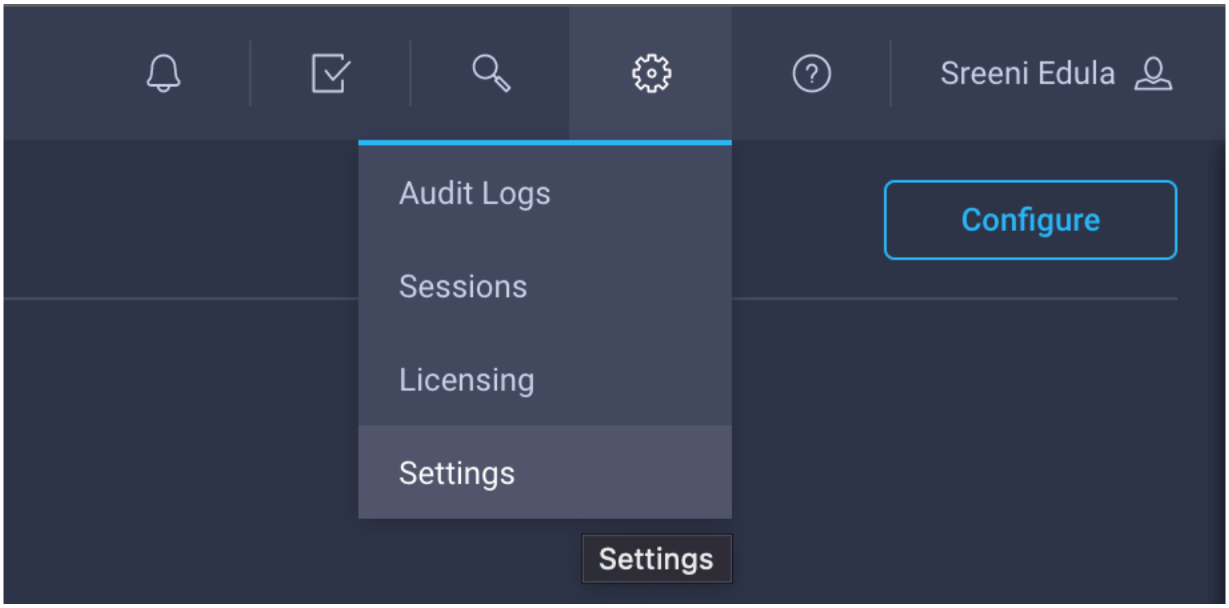

3. Click Resource Groups.

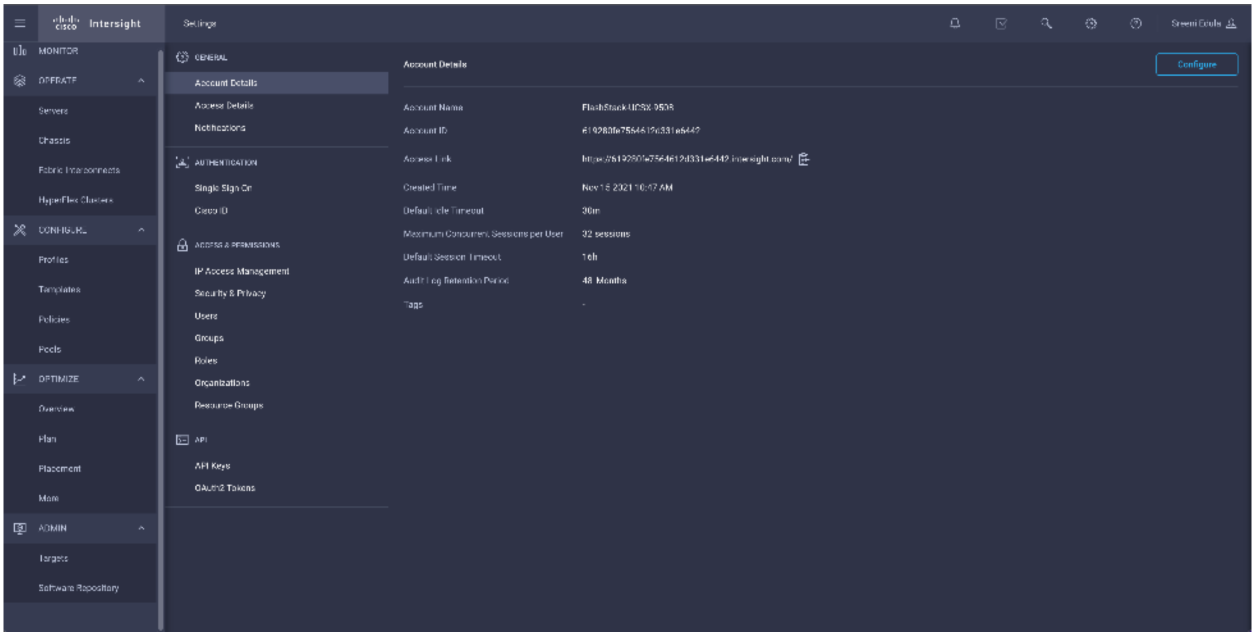

Account Details

Account Name FlashStack-UCSX-9508

Account ID 619280fe7564612d331e6442

Access Link https://619280fe7564612d331e6442.intersight.com/

Created Time Nov 15 2021 10:47 AM

Default Idle Timeout 30m

Maximum Concurrent Sessions per User 32 sessions

Default Session Timeout 16h

Audit Log Retention Period 48 Months

Tags -

4. Click Create Resource Group.

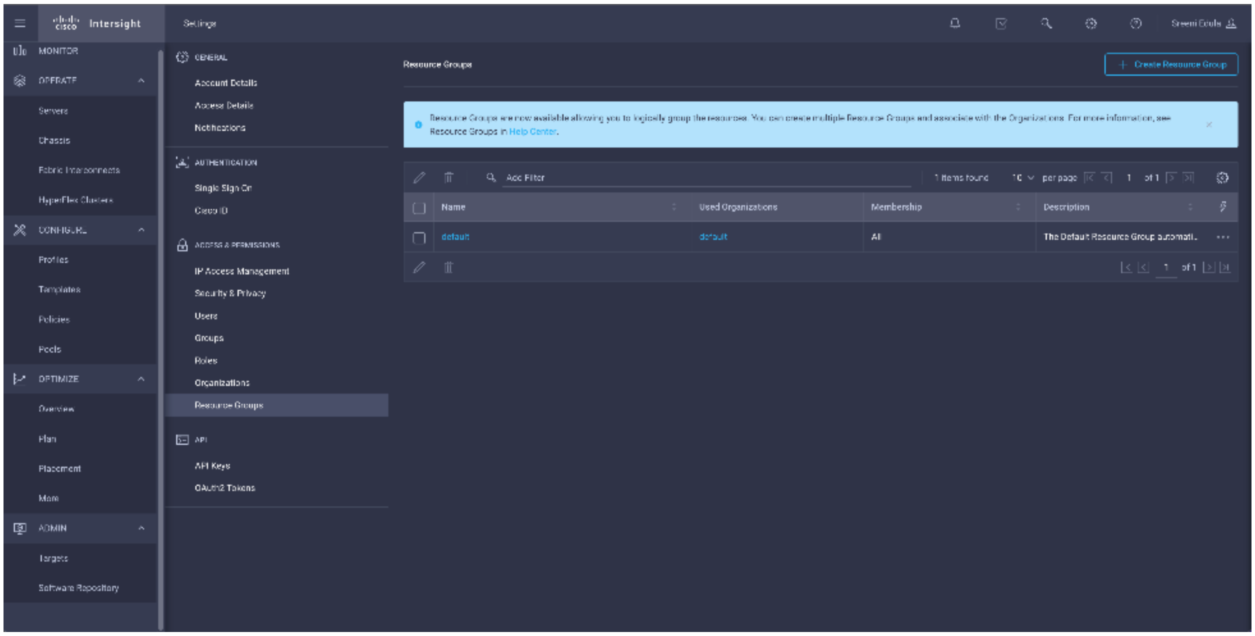

Resource Groups

Resource Groups are now available allowing you to logically group the resources. You can create multiple Resource Groups and associate with the Organizations. For more information, see Resource Groups in Help Center.

	Name	Used Organizations	Membership	Description
□	default	default	All	The Default Resource Group automati...

5. Provide a name for the Resource Group (for example **FSV-RG**).

6. Under Memberships, select Custom.

7. Select the recently added Cisco UCS device for this Resource Group.

8. Click Create.

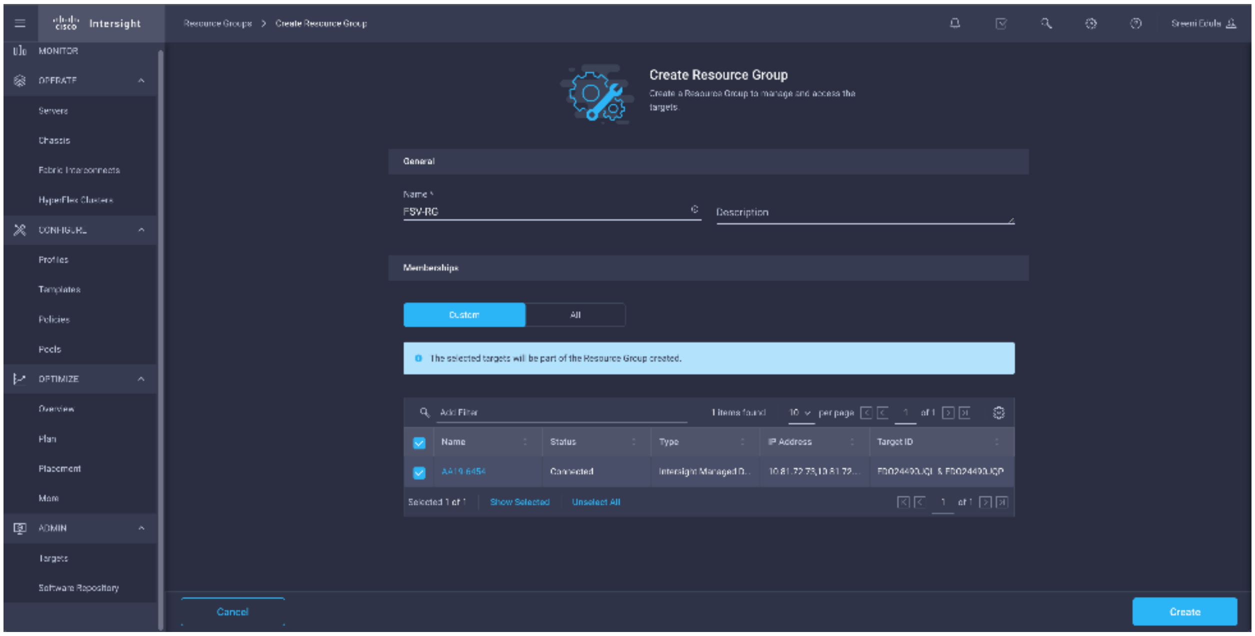

Create Resource Group

Create a Resource Group to manage and access the targets.

General

Name * FSV-RG

Description

Memberships

Custom | All

The selected targets will be part of the Resource Group created.

☑	Name	Status	Type	IP Address	Target ID
☑	AA19-6454	Connected	Intersight Managed D...	10.81.72.73,10.81.72....	FDO24490JQL & FDO24490JQP

Selected 1 of 1

9. Click Organizations.

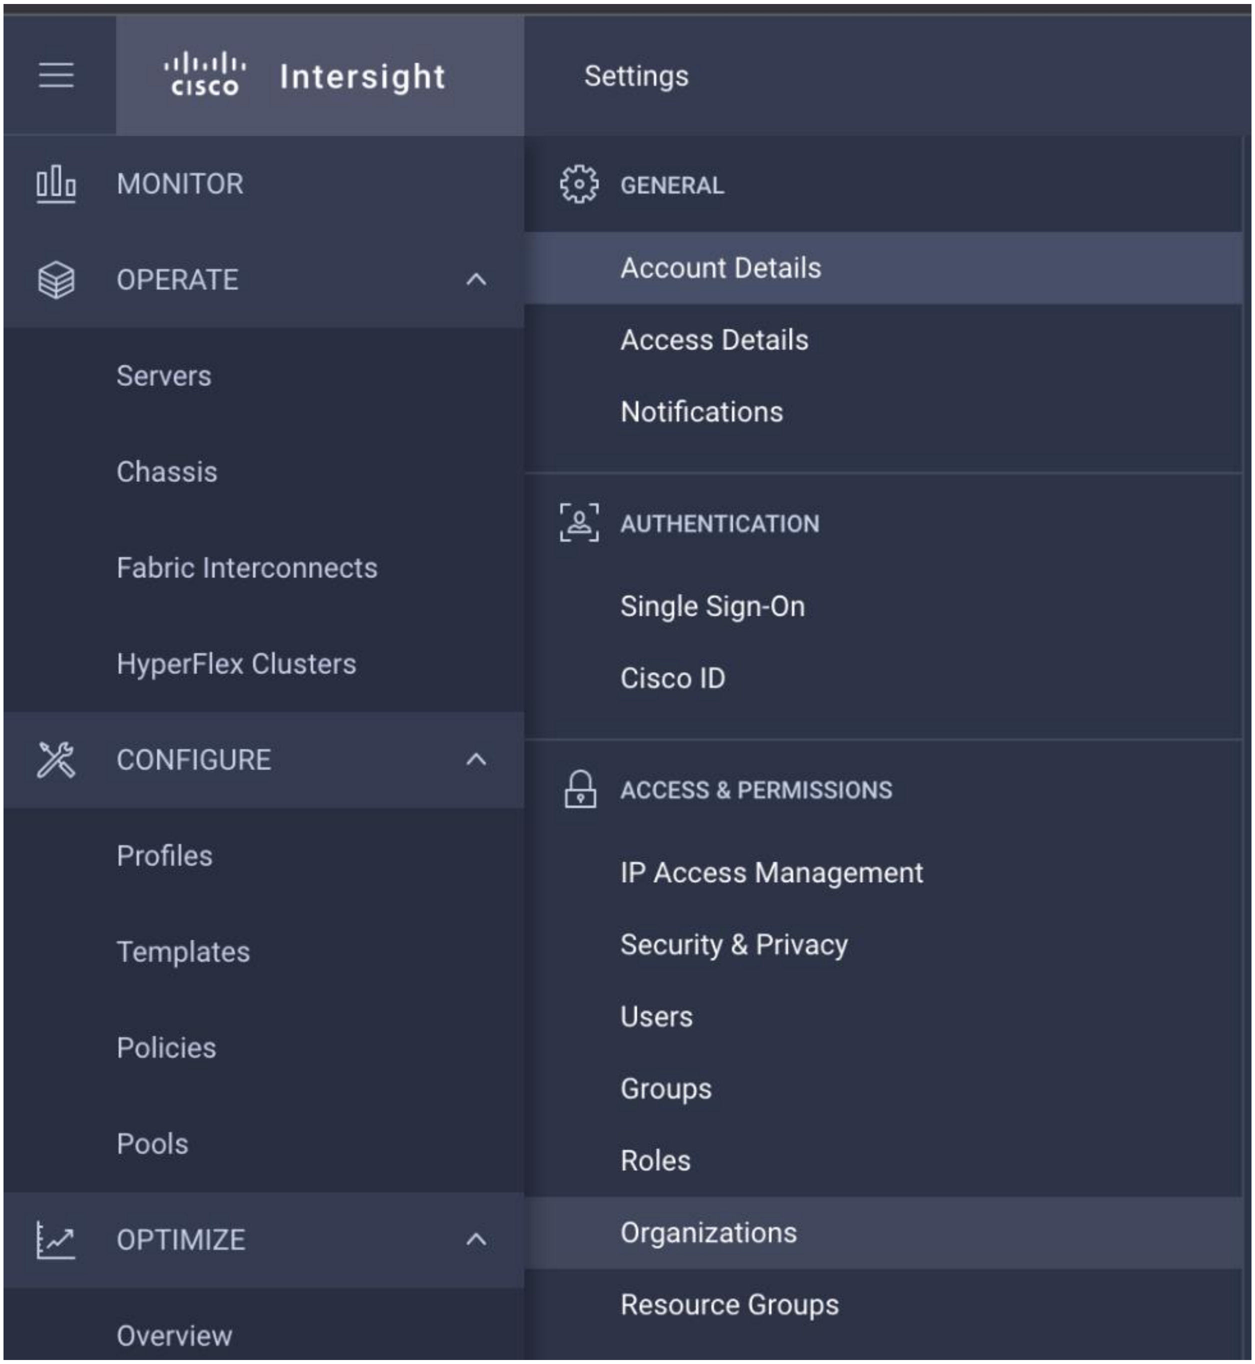

10. Click Create Organization.

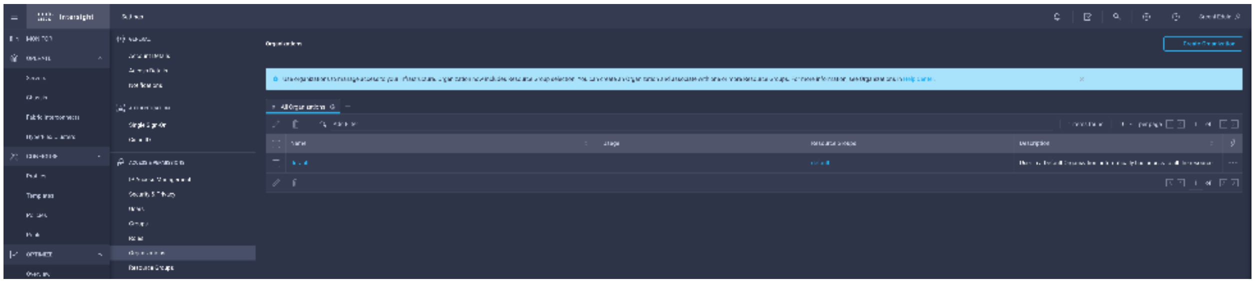

11. Provide a name for the organization (for example FSV) and select the previously created resource group (**FSV-RG**).

12. Click Create.

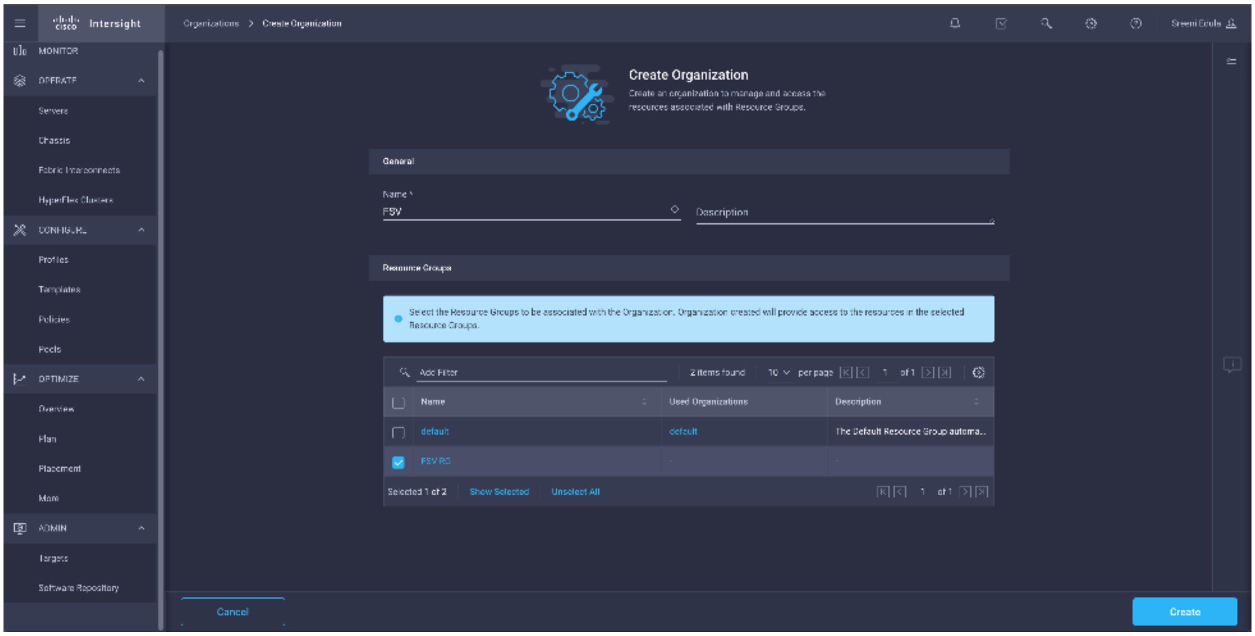

Create Organization

Create an organization to manage and access the resources associated with Resource Groups.

General

Name * FSV

Description

Resource Groups

Select the Resource Groups to be associated with the Organization. Organization created will provide access to the resources in the selected Resource Groups.

Add Filter — 2 items found — 10 per page — 1 of 1

	Name	Used Organizations	Description
□	default	default	The Default Resource Group automa..
☑	FSV-RG	-	-

Selected 1 of 2 — Show Selected — Unselect All — 1 of 1

Upgrade Fabric Interconnect Firmware using Cisco Intersight

Cisco UCS Manager does not support Cisco UCS X-Series, therefore upgrading Fabric Interconnect software using Cisco UCS Manager does not contain the firmware for Cisco UCS X-series. Before setting up a UCS domain profile and discovering the chassis, upgrade the Fabric Interconnect firmware to the latest recommended release using Cisco Intersight.

Note: If Cisco UCS Fabric Interconnects were upgraded to the latest recommended software using Cisco UCS Manager, this upgrade process through Intersight will still work and will copy the X-Series firmware to the Fabric Interconnects.

To perform the software upgrade, follow these steps:

1. Log in to the Cisco Intersight portal.

2. Click to expand OPERATE in the left pane and select Fabric Interconnects.

3. Click the ellipses "…" at the end of the row for either of the Fabric Interconnects and select Upgrade Firmware.

4. Click Start.

5. Verify the Fabric Interconnect information and click Next.

6. Enable Advanced Mode using the toggle switch and uncheck Fabric Interconnect Traffic Evacuation.

7. Select the recommended release from the list and click Next.

8. Verify the information and click Upgrade to start the upgrade process.

9. Keep an eye on the Request panel of the main Intersight screen as the system will ask for user permission before upgrading each FI. Click on the Circle with Arrow and follow the prompts on screen to grant permission.

10. Wait for both the FIs to successfully upgrade.

Configure a Cisco UCS Domain Profile

A Cisco UCS domain profile configures a fabric interconnect pair through reusable policies, allows configuration of the ports and port channels, and configures the VLANs and VSANs in the network. It defines the characteristics of and configures ports on fabric interconnects. The domain-related policies can be attached to the profile either at the time of creation or later. One Cisco UCS domain profile can be assigned to one fabric interconnect domain.

To create a Cisco UCS domain profile, follow these steps:

1. Log into the Cisco Intersight portal

2. Click to expand CONFIGURE in the left pane and select Profiles.

3. In the main window, select UCS Domain Profiles and click Create UCS Domain Profile.

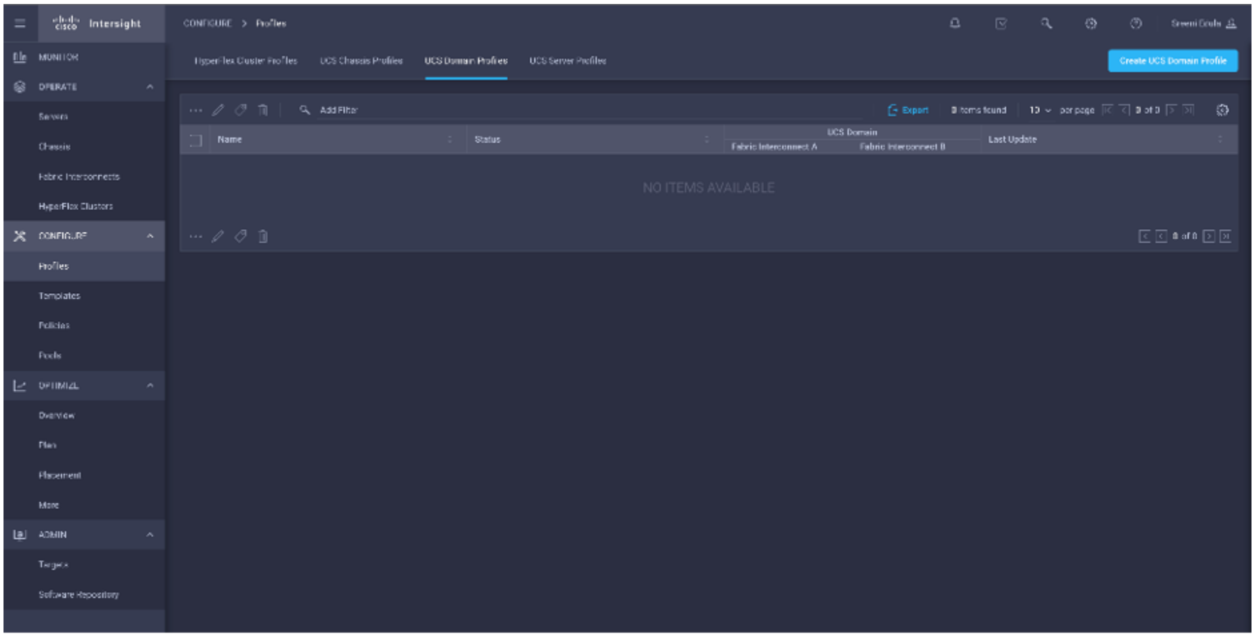

4. In the “Create UCS Domain Profile” screen, click Start.

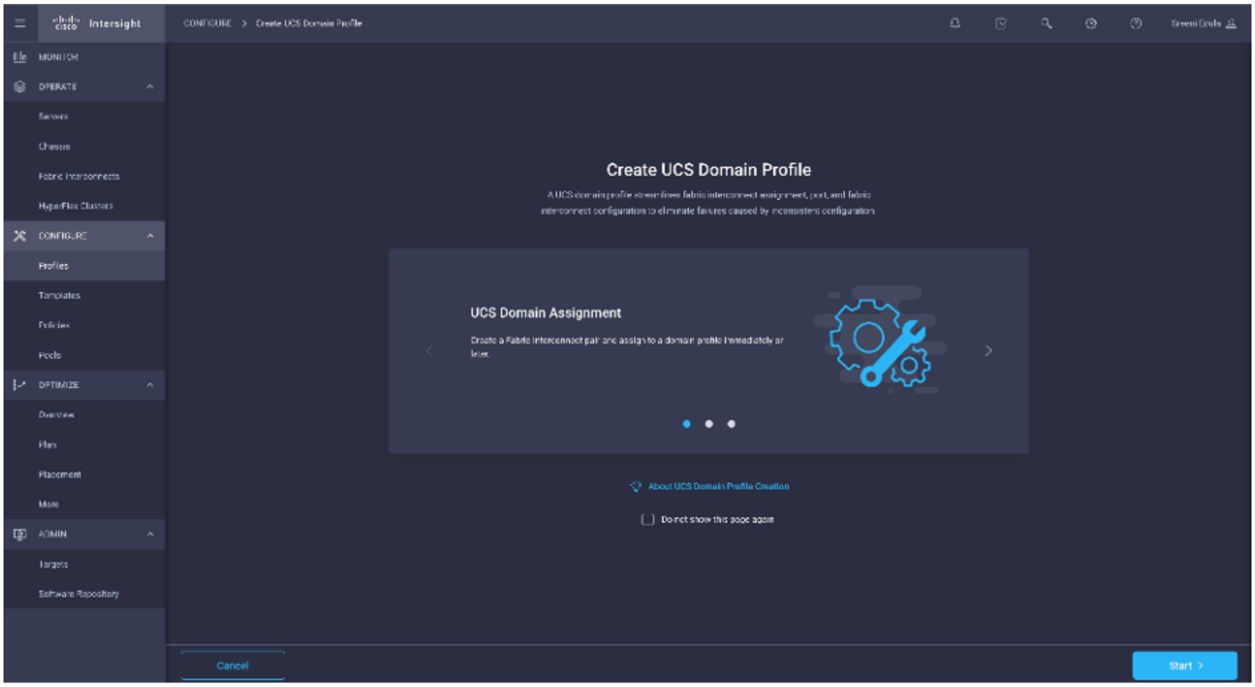

Create UCS Domain Profile

A UCS domain profile streamlines fabric interconnect assignment, port, and fabric interconnect configuration to eliminate failures caused by inconsistent configuration.

UCS Domain Assignment

Create a Fabric Interconnect pair and assign to a domain profile immediately or later.

About UCS Domain Profile Creation

□ Do not show this page again

Step 1 – General

1. Choose the organization from the drop-down list (for example, FSV).

2. Provide a name for the domain profile (for example AA19-Domain-Profile).

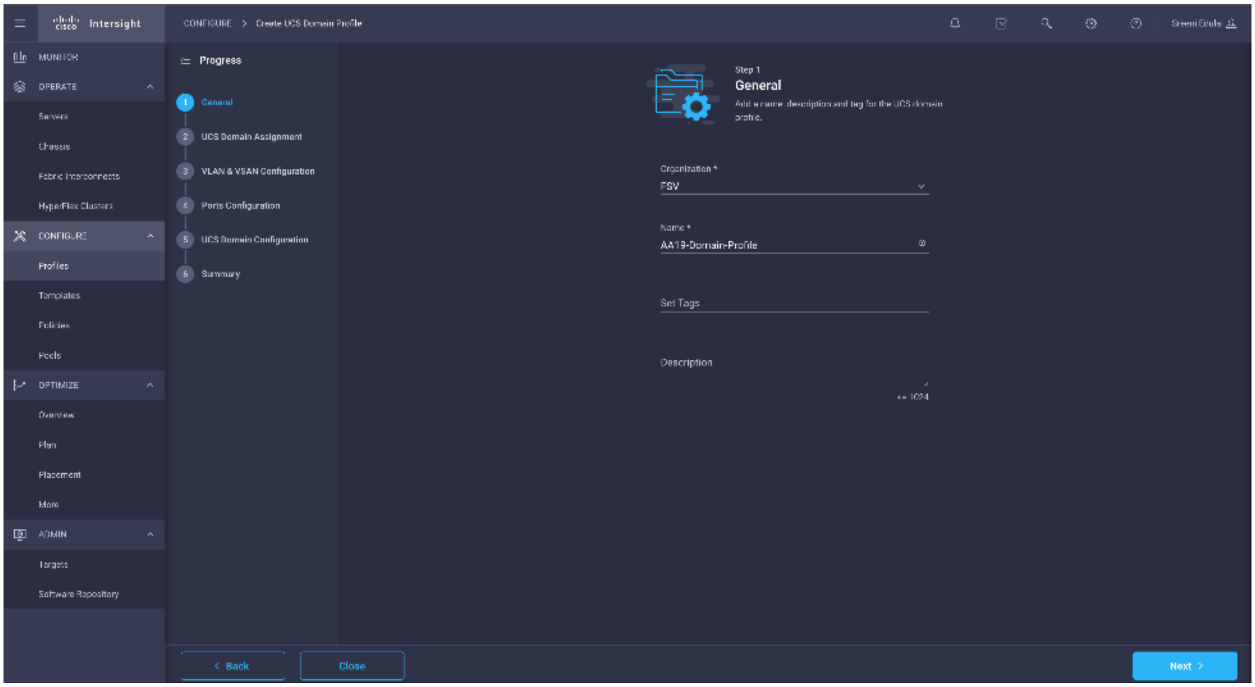

Step 1

General

Add a name, description and tag for the UCS domain profile.

Organization *
FSV

Name *
AA19-Domain-Profile

Set Tags

Description

<= 1024

3. Click Next.

Step 2 – UCS Domain Assignment

To create the Cisco UCS domain assignment, follow these steps:

1. Assign the Cisco UCS domain to this new domain profile by clicking Assign Now and selecting the previously added Cisco UCS domain (AA19-6454).

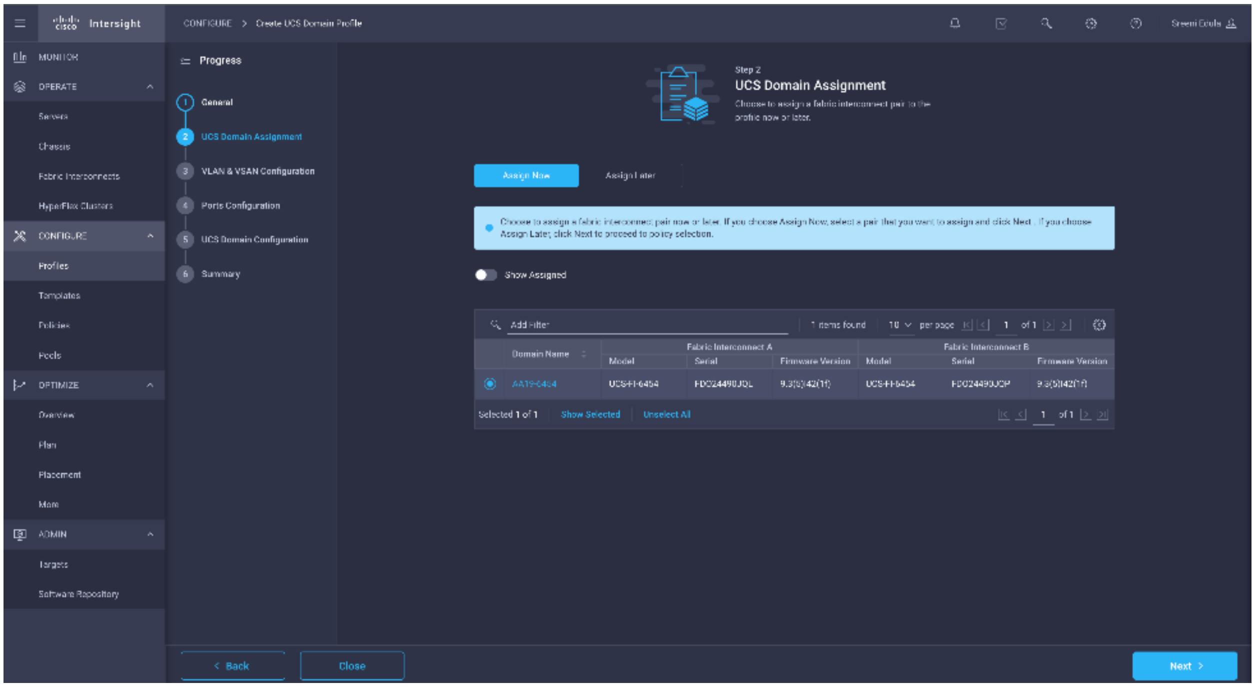

Step 2

UCS Domain Assignment

Choose to assign a fabric interconnect pair to the profile now or later.

Choose to assign a fabric interconnect pair now or later. If you choose Assign Now, select a pair that you want to assign and click Next . If you choose Assign Later, click Next to proceed to policy selection.

1 items found 10 per page 1 of 1

	Fabric Interconnect A			Fabric Interconnect B		
Domain Name	Model	Serial	Firmware Version	Model	Serial	Firmware Version
AA19-6454	UCS-FI-6454	FDO24490JQL	9.3(5)I42(1f)	UCS-FI-6454	FDO24490JCP	9.3(5)I42(1f)

Selected 1 of 1 Show Selected Unselect All

2. Click Next.

Step 3 – VLAN and VSAN Configuration

In this step, a single VLAN policy will be created for both FIs, but individual policies will be created for the VSANs as the VSAN IDs are unique for each FI.

Create and apply the VLAN Policy

1. Click Select Policy next to VLAN Configuration under Fabric Interconnect A and in the pane on the right, click Create New.

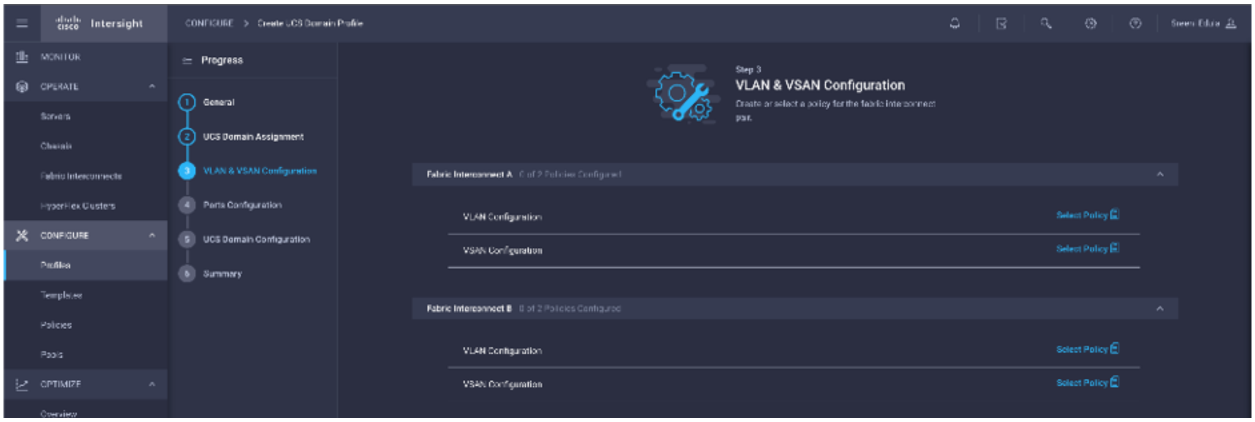

Step 3

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 0 of 2 Policies Configured

VLAN Configuration

VSAN Configuration

Fabric Interconnect B 0 of 2 Policies Configured

VLAN Configuration

VSAN Configuration

2. Verify the correct Organization is selected (for example, FSV).

3. Provide a name for the policy (for example, AA19-VLAN-Pol).

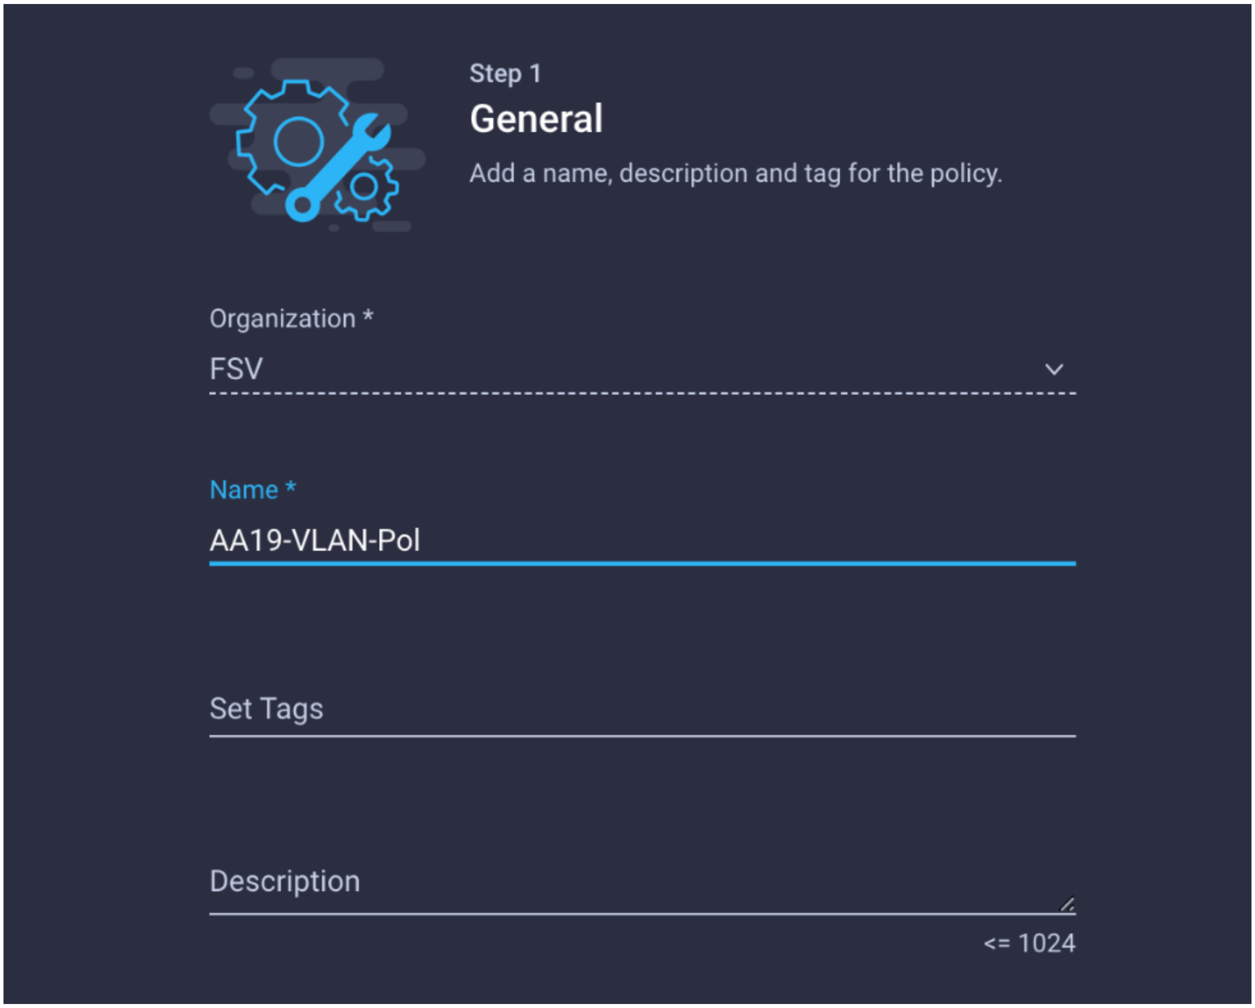
Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-VLAN-Pol

Set Tags

Description

<= 1024

4. Click Next.

5. Click Add VLANs.

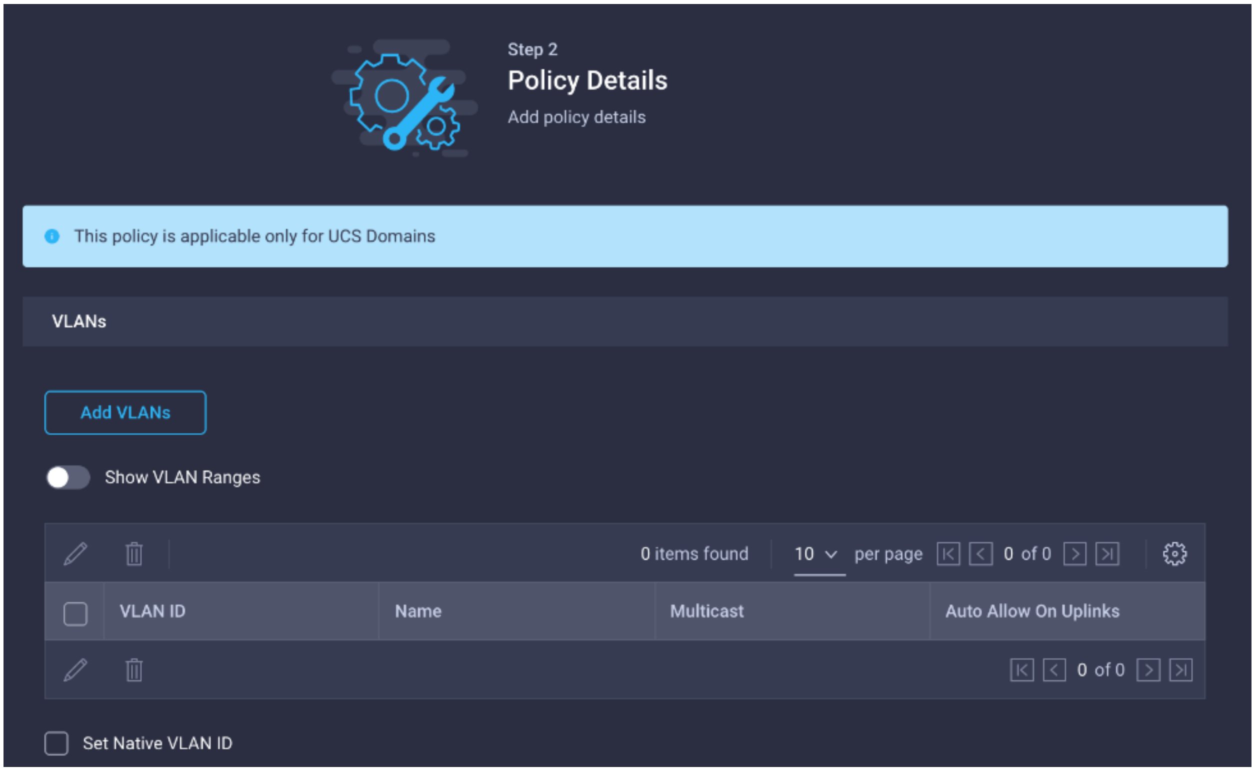

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

○ Show VLAN Ranges

0 items found | 10 per page | 0 of 0

□	VLAN ID	Name	Multicast	Auto Allow On Uplinks

0 of 0

□ Set Native VLAN ID

6. Provide the Name and VLAN ID for the native VLAN (for example, Native-VLAN: 2).

7. Make sure Auto Allow On Uplinks is enabled.

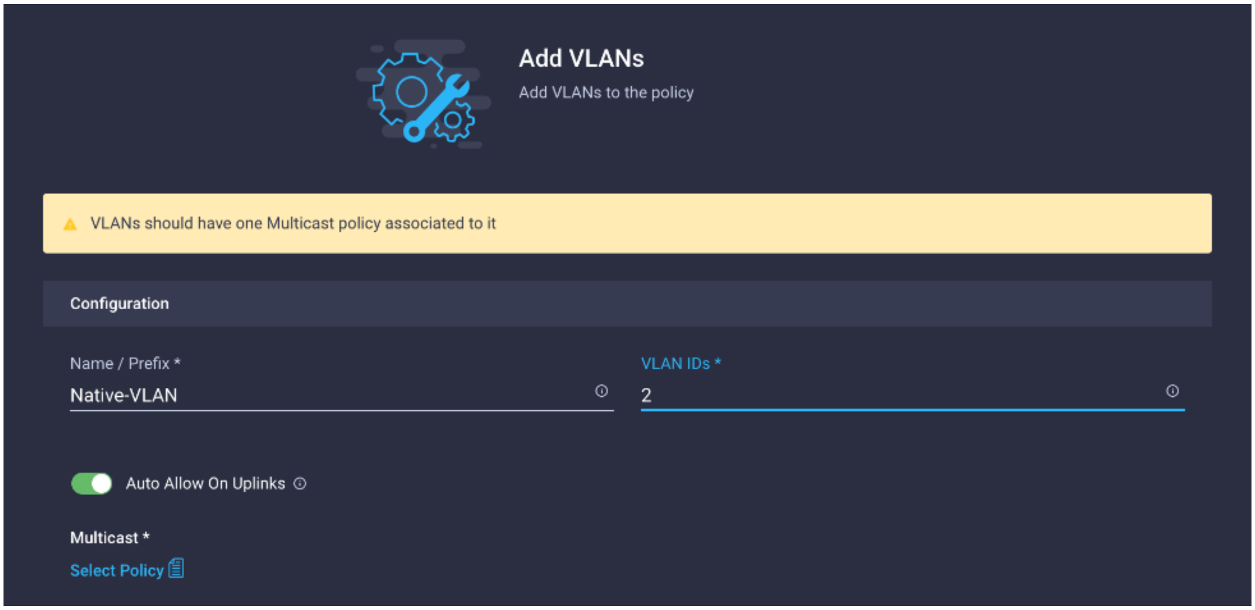

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix *: Native-VLAN

VLAN IDs *: 2

● Auto Allow On Uplinks

Multicast *

8. Click Select Policy under Multicast in the pane on the left, click Create New.

9. Verify the correct Organization is selected (for example, FSV).

10. Provide a name for the policy (for example, AA19-MCAST-Pol).

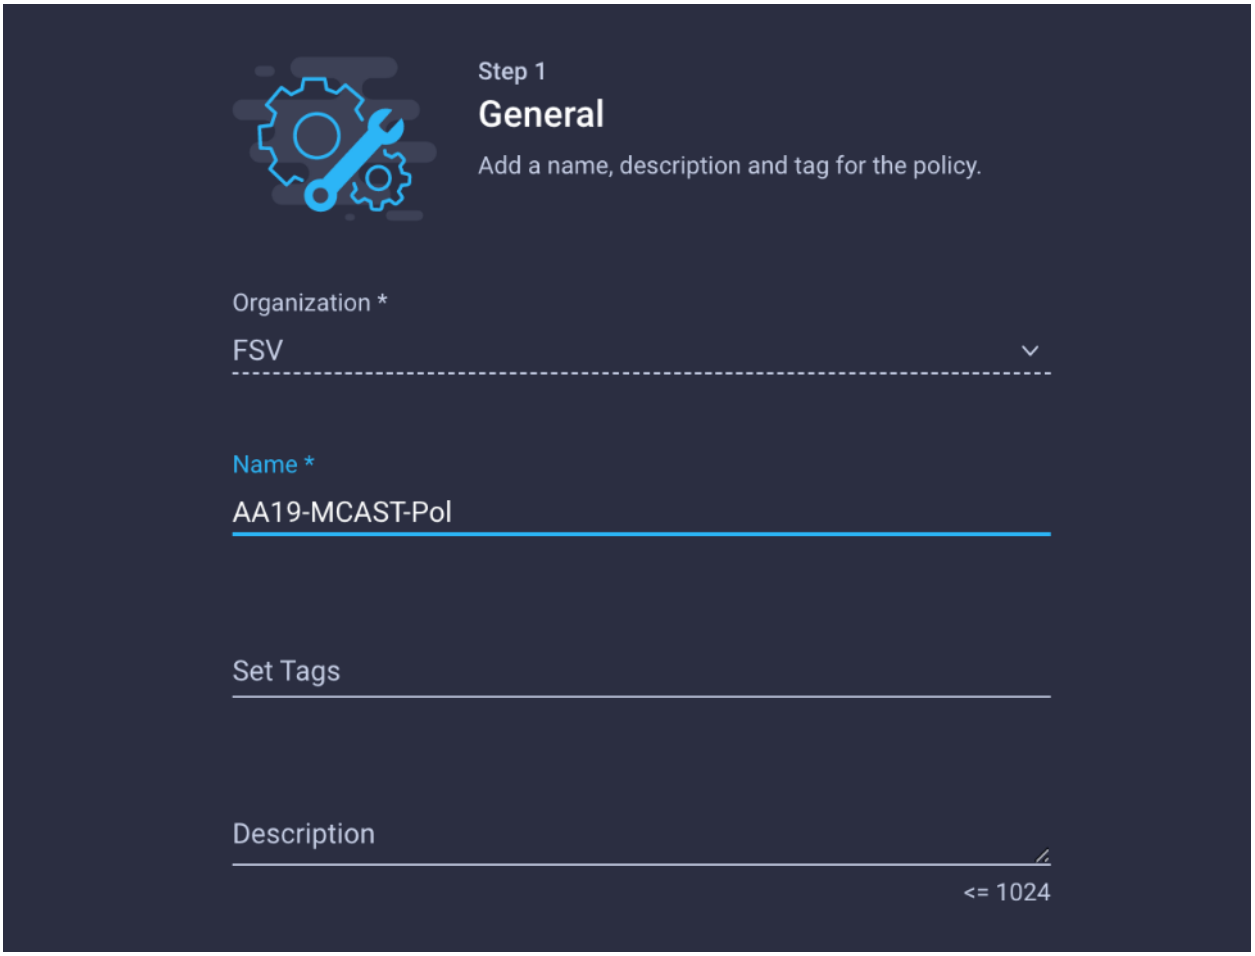

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-MCAST-Pol

Set Tags

Description

<= 1024

11. Click Next.

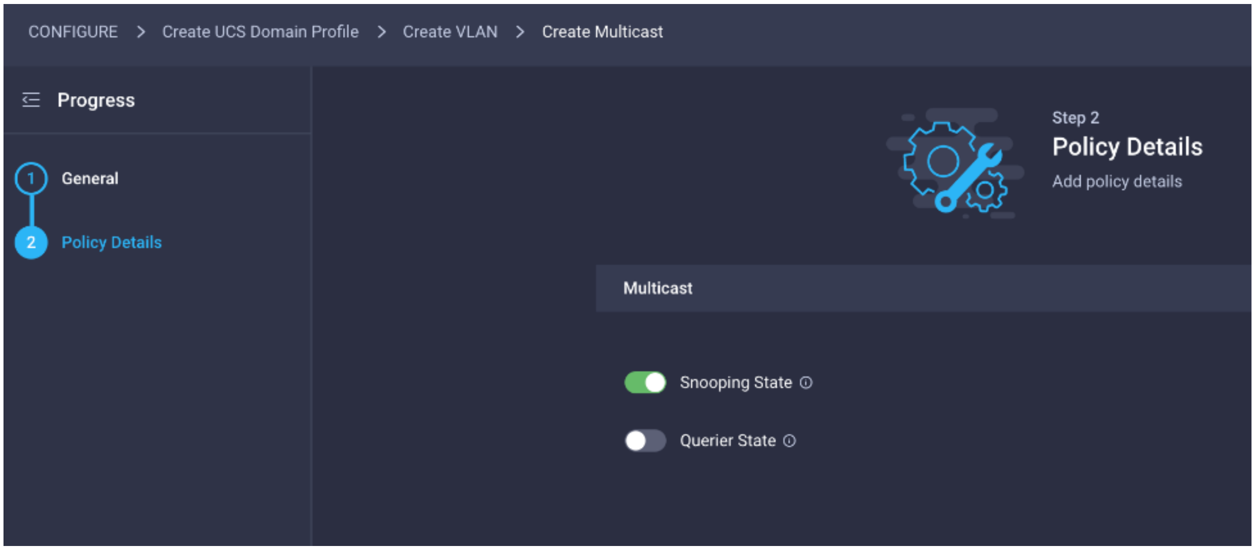

CONFIGURE > Create UCS Domain Profile > Create VLAN > Create Multicast

Step 2

Policy Details

Add policy details

Multicast

Snooping State (on)

Querier State (off)

12. Click Create.

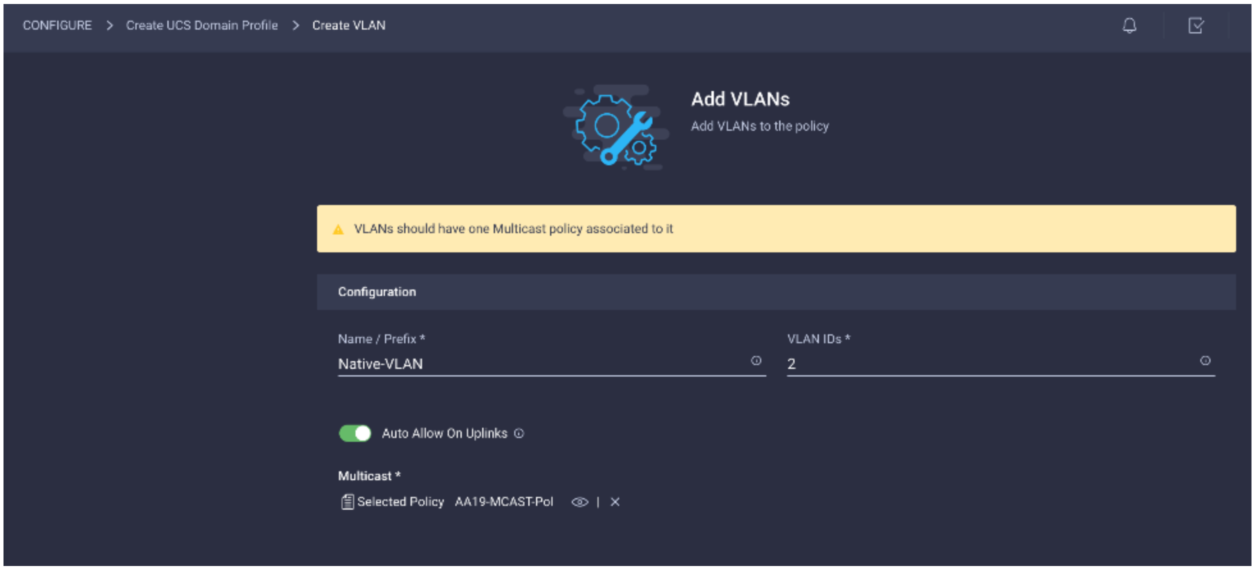
CONFIGURE > Create UCS Domain Profile > Create VLAN

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix * Native-VLAN

VLAN IDs * 2

Auto Allow On Uplinks

Multicast *

Selected Policy AA19-MCAST-Pol

13. Click Add to add the VLAN.

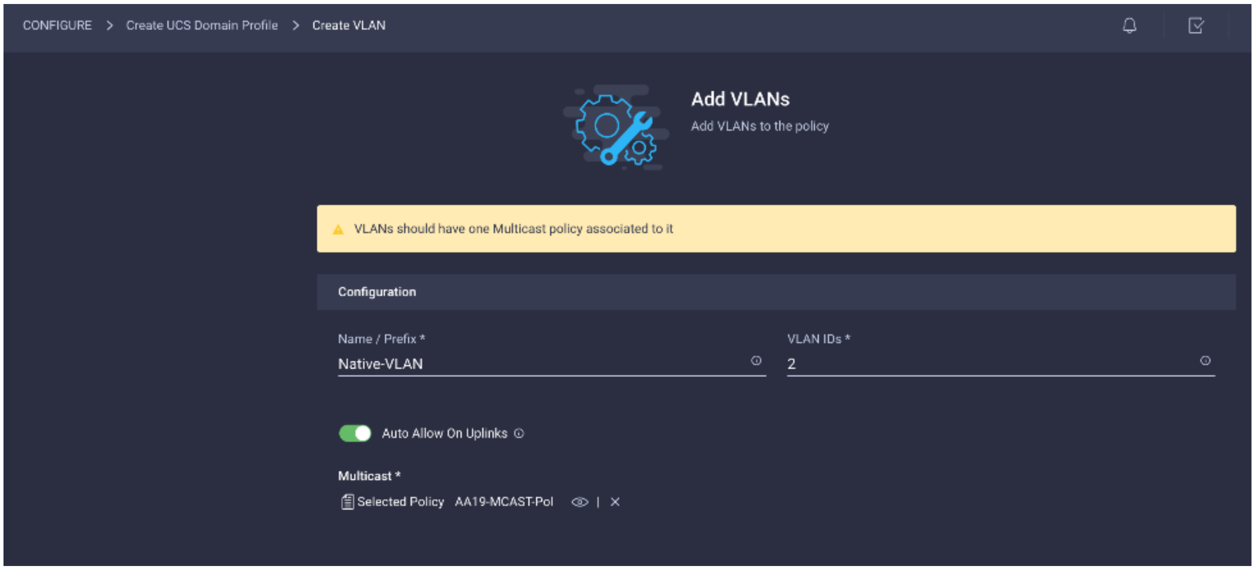
CONFIGURE > Create UCS Domain Profile > Create VLAN

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix * Native-VLAN

VLAN IDs * 2

Auto Allow On Uplinks

Multicast *

Selected Policy AA19-MCAST-Pol

14. Select Set Native VLAN ID and enter VLAN number (for example, 2) under the VLAN ID.

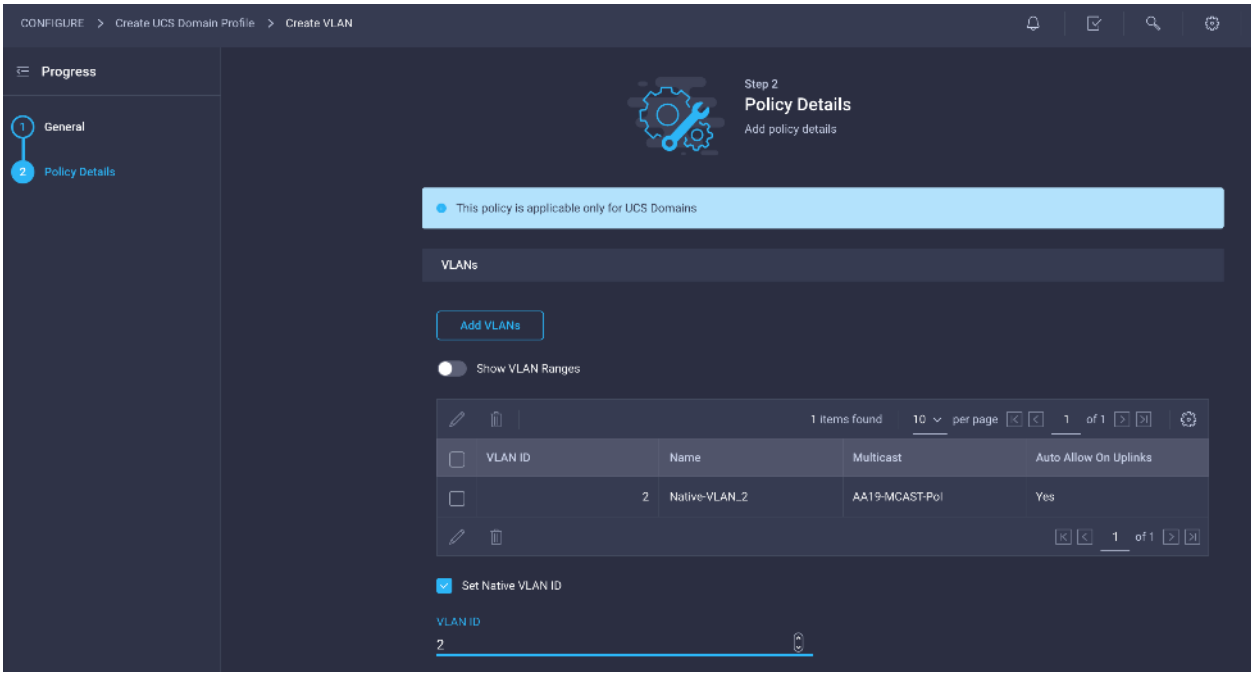

CONFIGURE > Create UCS Domain Profile > Create VLAN

Progress

1 General

2 Policy Details

Step 2
Policy Details
Add policy details

This policy is applicable only for UCS Domains

VLANs

Add VLANs

Show VLAN Ranges

1 items found 10 per page 1 of 1

☐	VLAN ID	Name	Multicast	Auto Allow On Uplinks
☐	2	Native-VLAN_2	AA19-MCAST-Pol	Yes

☑ Set Native VLAN ID

VLAN ID
2

15. Add remaining VLANs for FlashStack by clicking Add VLANs and entering the VLANs one by one while selecting the Multicast policy (AA19-IB-MGMT-VLAN) created in earlier steps. The VLANs used during this validation for FC based storage are shown below.

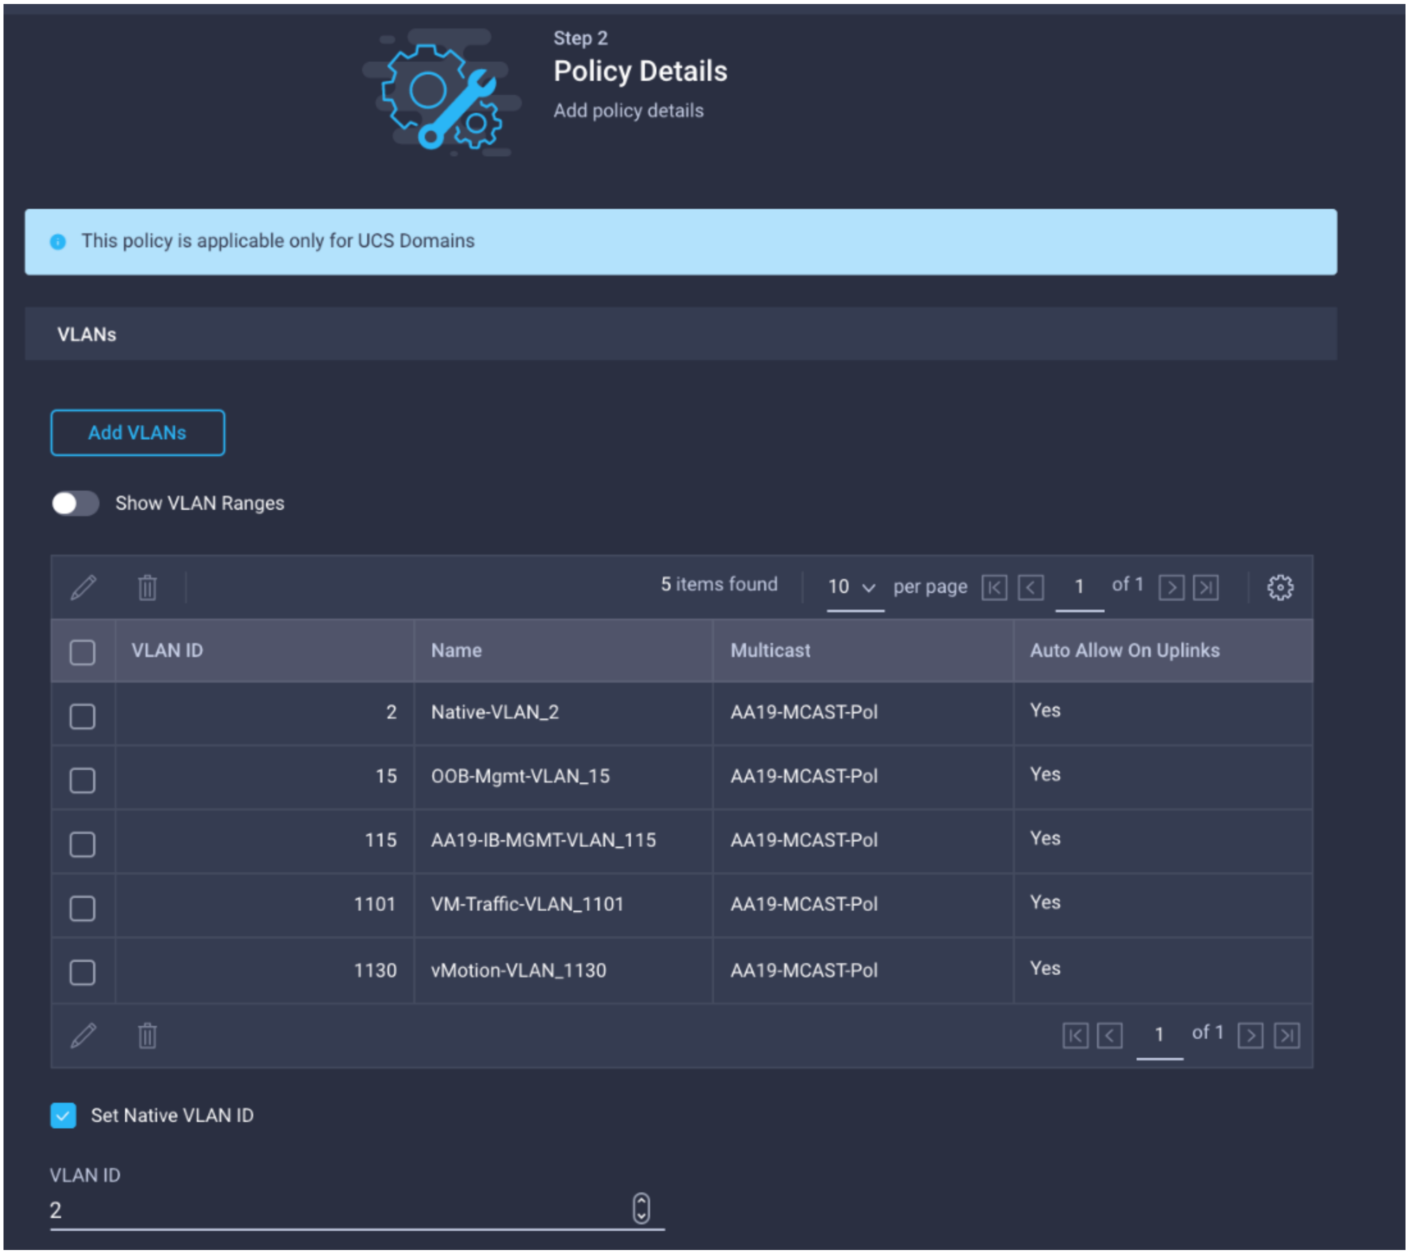
Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

Show VLAN Ranges

5 items found 10 per page 1 of 1

☐	VLAN ID	Name	Multicast	Auto Allow On Uplinks
☐	2	Native-VLAN_2	AA19-MCAST-Pol	Yes
☐	15	OOB-Mgmt-VLAN_15	AA19-MCAST-Pol	Yes
☐	115	AA19-IB-MGMT-VLAN_115	AA19-MCAST-Pol	Yes
☐	1101	VM-Traffic-VLAN_1101	AA19-MCAST-Pol	Yes
☐	1130	vMotion-VLAN_1130	AA19-MCAST-Pol	Yes

1 of 1

☑ Set Native VLAN ID

VLAN ID

2

Note: If implementing iSCSI storage, add the iSCSI-A and iSCSI-B VLANs.

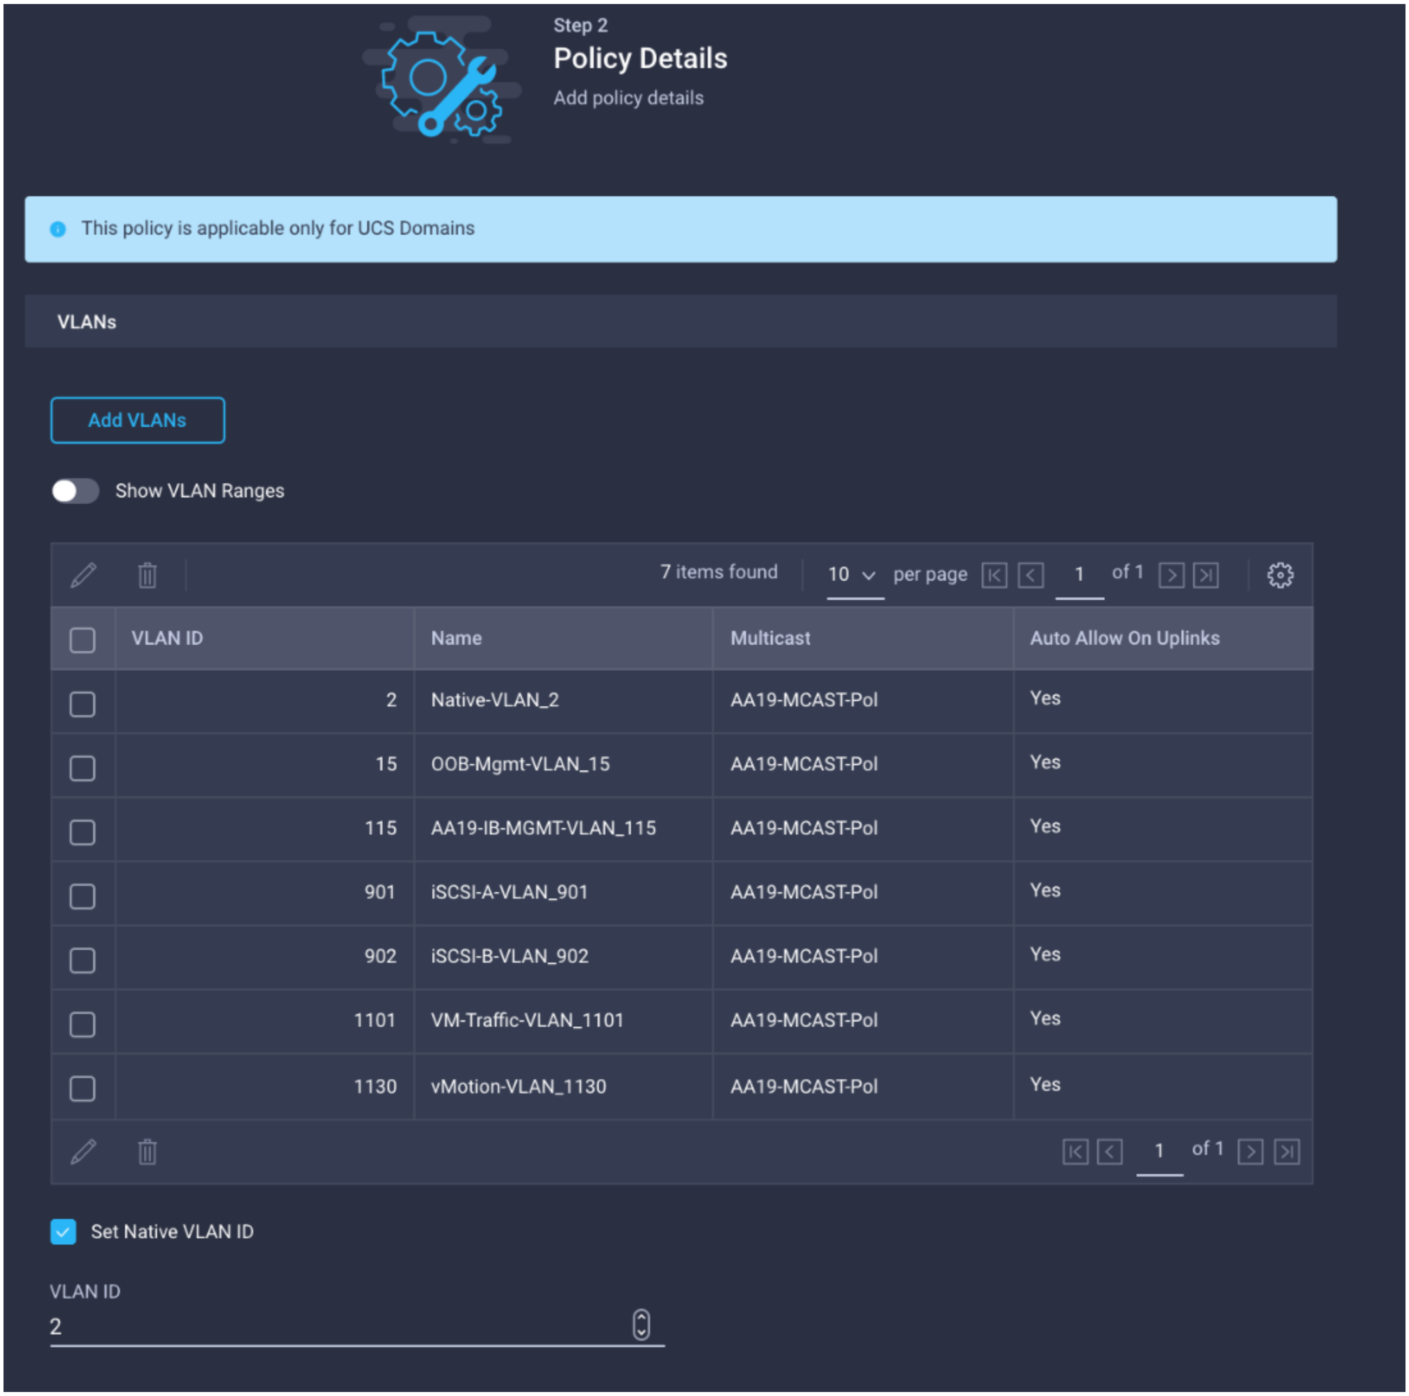

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

☐	VLAN ID	Name	Multicast	Auto Allow On Uplinks
☐	2	Native-VLAN_2	AA19-MCAST-Pol	Yes
☐	15	OOB-Mgmt-VLAN_15	AA19-MCAST-Pol	Yes
☐	115	AA19-IB-MGMT-VLAN_115	AA19-MCAST-Pol	Yes
☐	901	iSCSI-A-VLAN_901	AA19-MCAST-Pol	Yes
☐	902	iSCSI-B-VLAN_902	AA19-MCAST-Pol	Yes
☐	1101	VM-Traffic-VLAN_1101	AA19-MCAST-Pol	Yes
☐	1130	vMotion-VLAN_1130	AA19-MCAST-Pol	Yes

☑ Set Native VLAN ID

VLAN ID
2

16. Click Create to create all the VLANs.

17. Click Select Policy next to VLAN Configuration for FI-B and select the same VALN policy that was created in the last step.

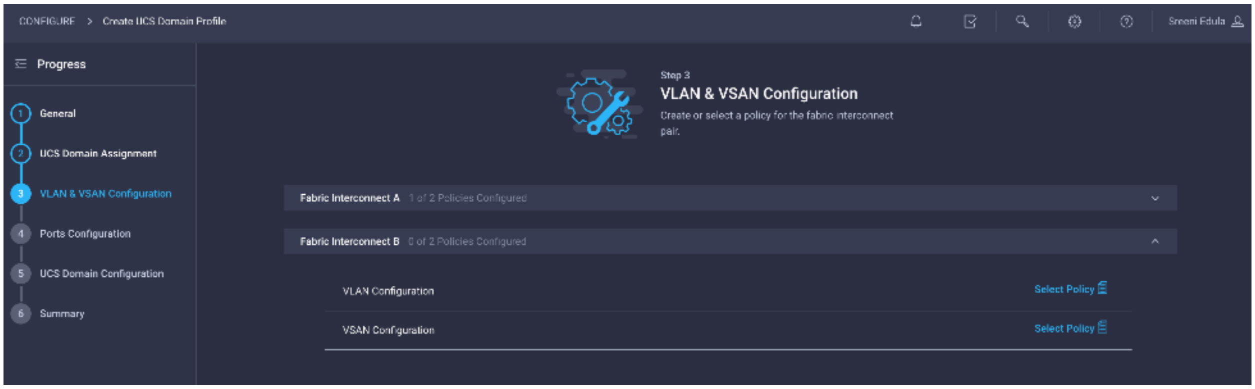

Step 3

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 1 of 2 Policies Configured

Fabric Interconnect B 0 of 2 Policies Configured

VLAN Configuration

VSAN Configuration

Create and apply the VSAN Policies (FC configuration Only)

To create and apply the VSAN policy, follow these steps:

Note: A VSAN policy is only needed when configuring Fibre Channel and can be skipped when configuring IP-only storage access.

1. Click Select Policy next to VSAN Configuration under Fabric Interconnect A and in the pane on the right, click Create New.

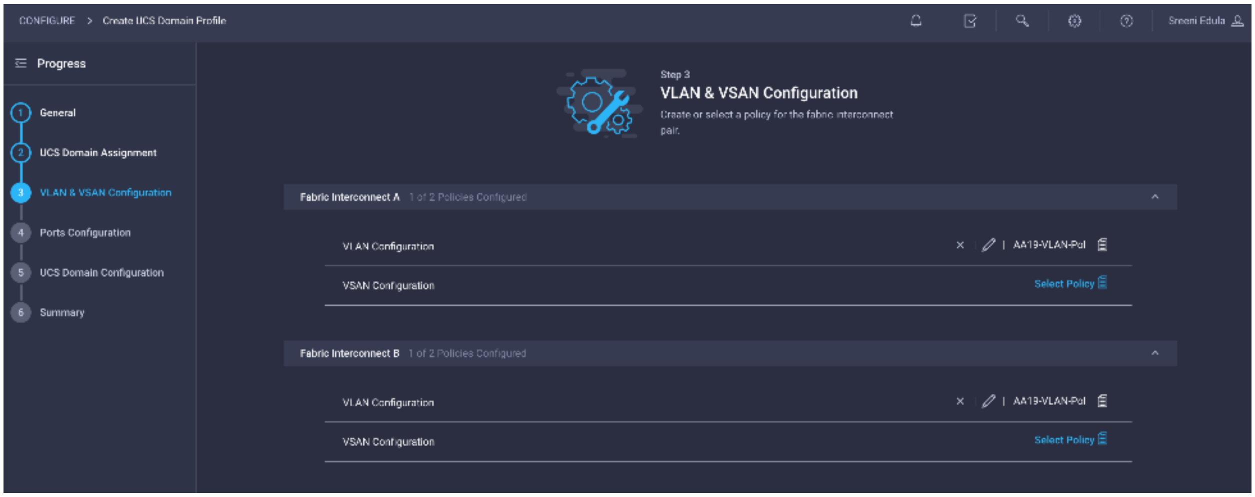

Step 3

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 1 of 2 Policies Configured

VLAN Configuration | AA19-VLAN-Pol

VSAN Configuration

Fabric Interconnect B 1 of 2 Policies Configured

VLAN Configuration | AA19-VLAN-Pol

VSAN Configuration

2. Verify the correct Organization is selected (for example, FSV).

3. Provide a name for the policy (for example, AA19-VSAN-Pol-A).

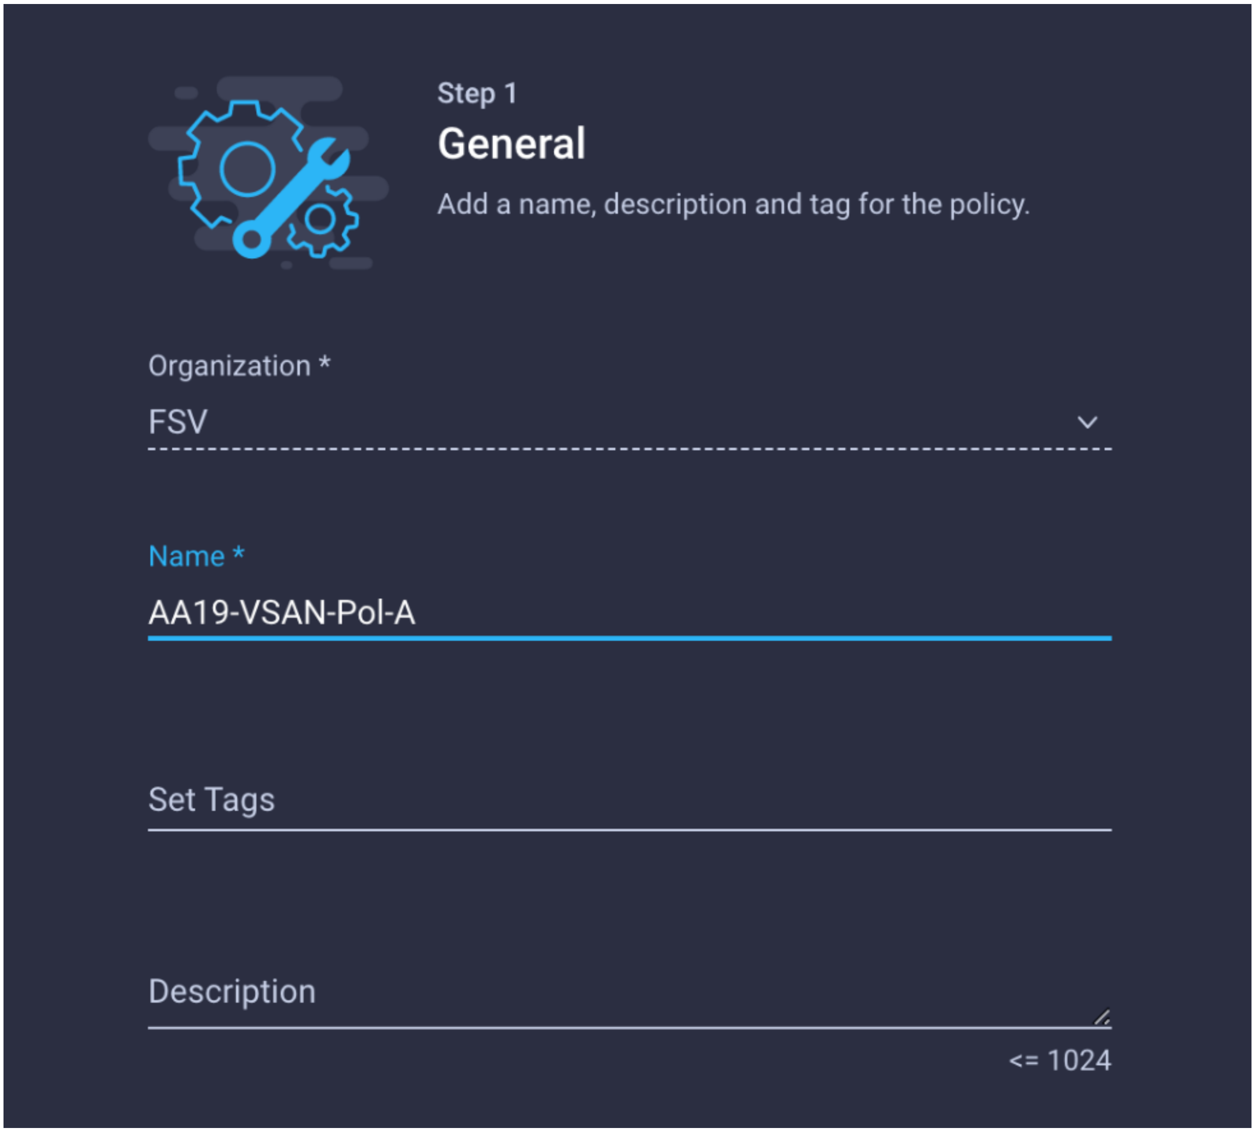

Step 1

General

Add a name, description and tag for the policy.

Organization * FSV

Name * AA19-VSAN-Pol-A

4. Click Next.

5. Click Add VSAN and provide the Name (for example, VSAN-A), VSAN ID (for example, 101) and associated FCoE VLAN ID (for example, 101) for SAN-A.

6. Click Add.

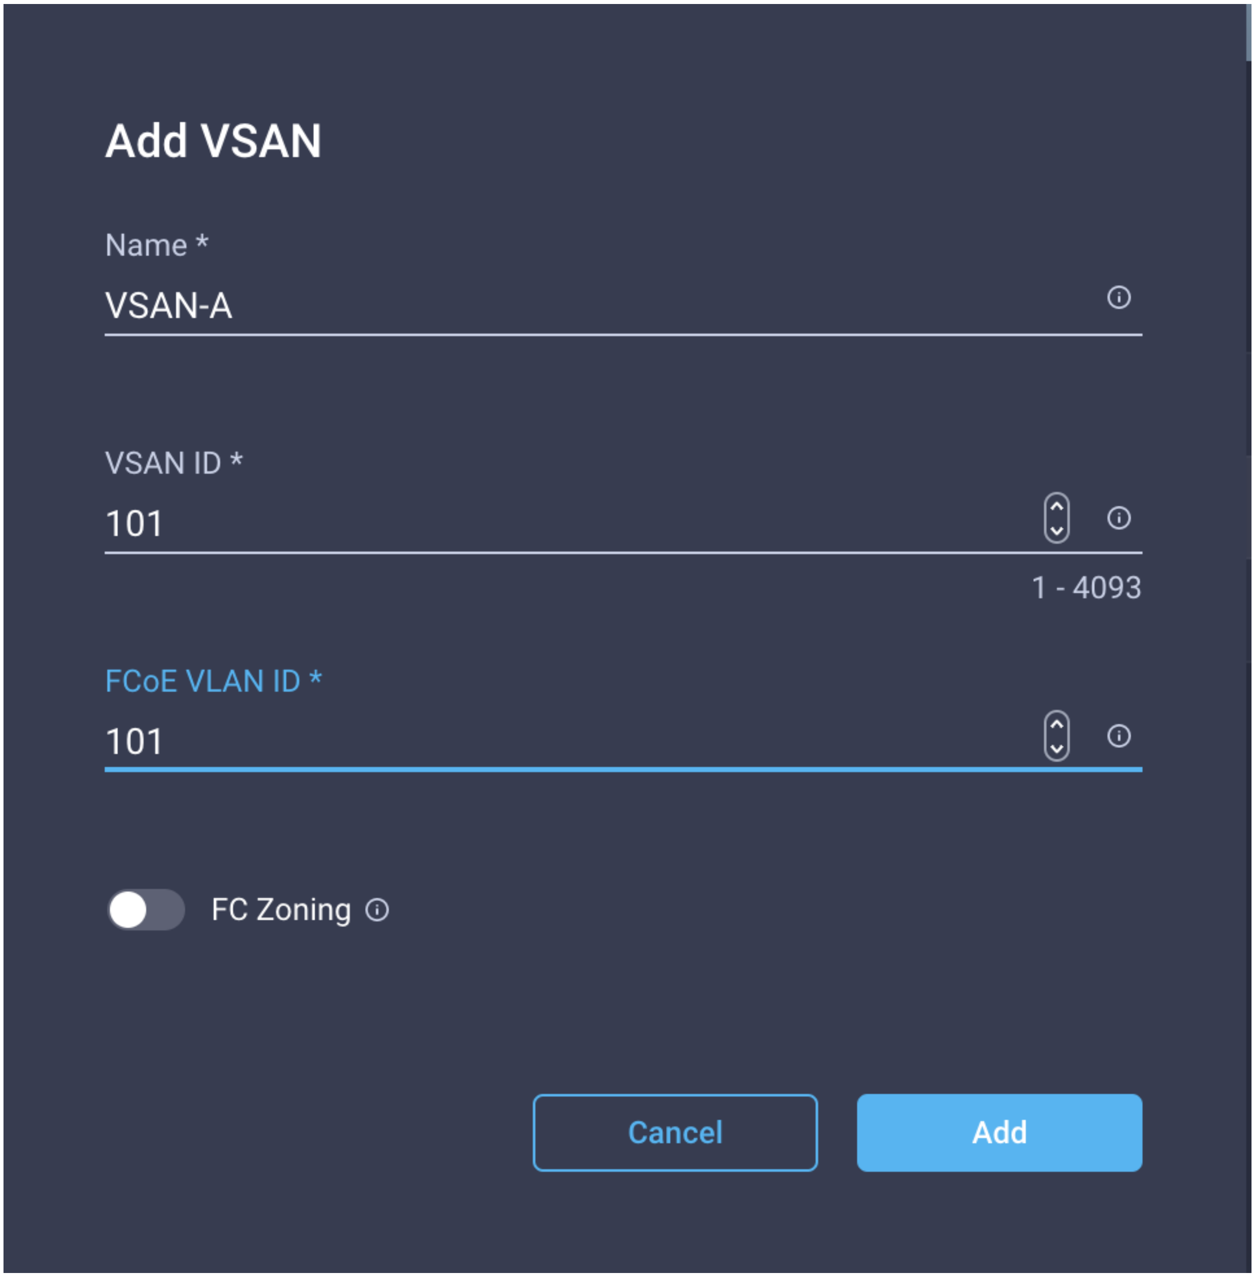

Add VSAN

Name *

VSAN-A

VSAN ID *

101

1 - 4093

FCoE VLAN ID *

101

FC Zoning

7. Enable Uplink Trunking for this VSAN.

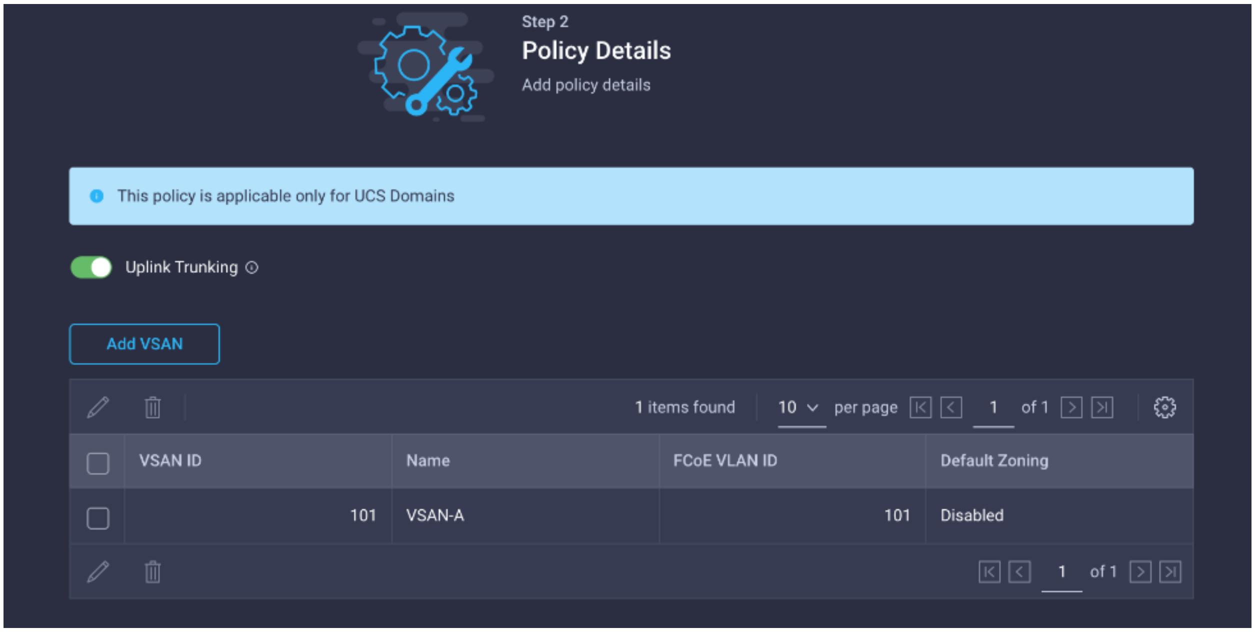

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Uplink Trunking: on

VSAN ID	Name	FCoE VLAN ID	Default Zoning
101	VSAN-A	101	Disabled

8. Click Create.

9. Repeat steps 1 – 8 to create a new VSAN policy for SAN-B.

10. Click Select Policy next to VSAN Configuration under Fabric Interconnect B and in the pane on the right, click Create New.

11. Verify the correct Organization is selected (for example, FSV).

12. Provide a name for the policy (for example, AA19-VSAN-Pol-B).

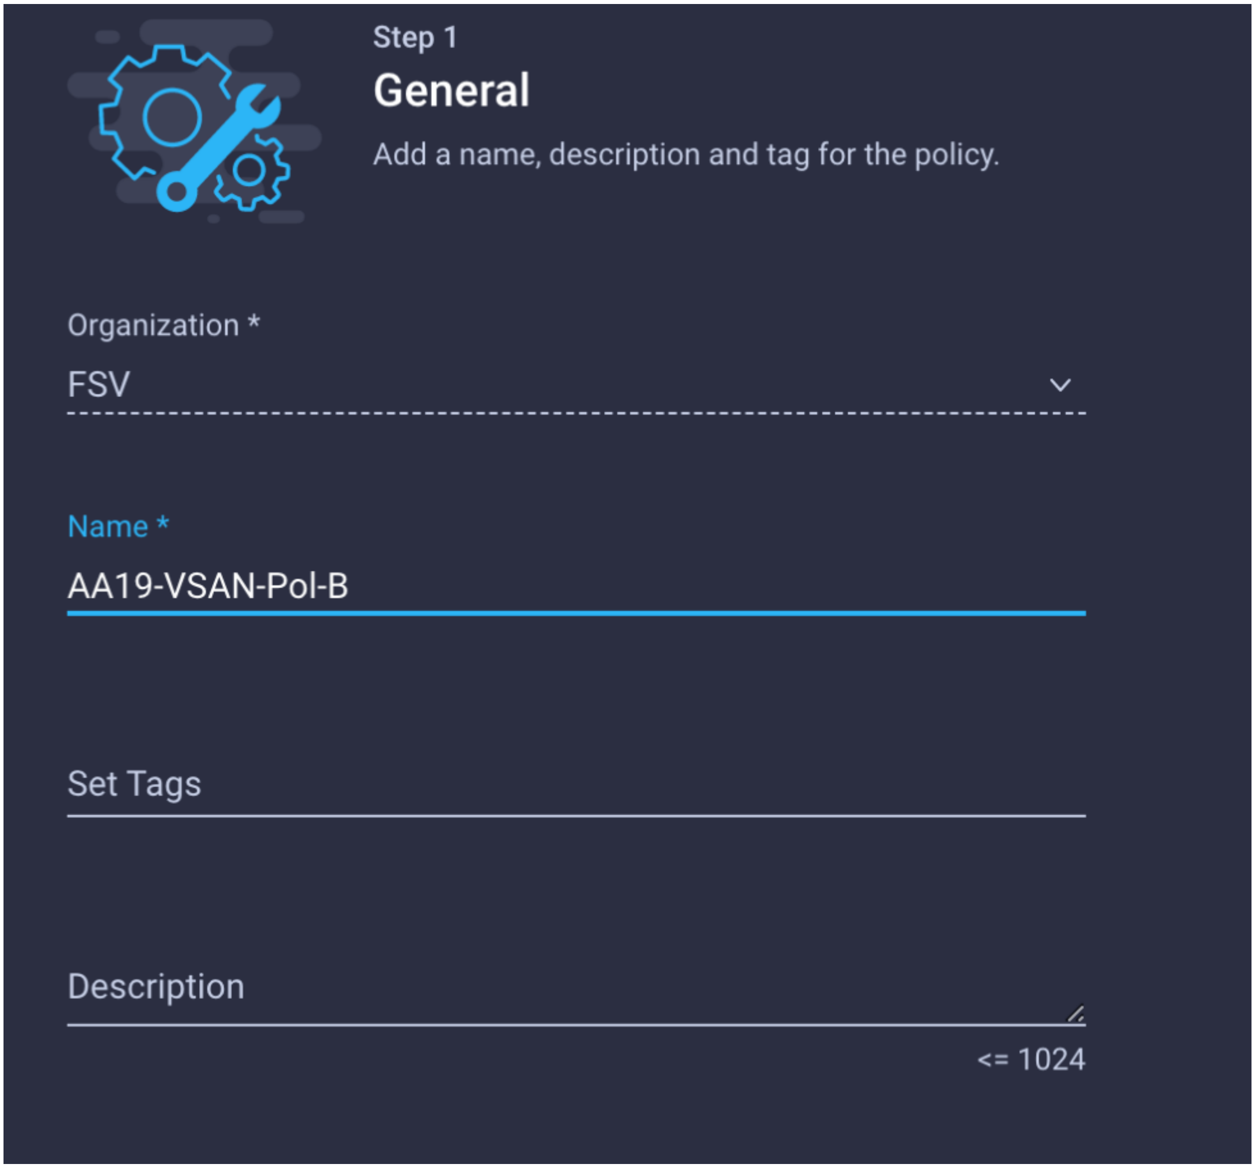

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-VSAN-Pol-B

Set Tags

Description

<= 1024

13. Click Next.

14. Click Add VSAN and provide the Name (for example, VSAN-B), VSAN ID (for example, 102) and associated FCoE VLAN ID (for example, 102) for SAN-B.

15. Click Add.

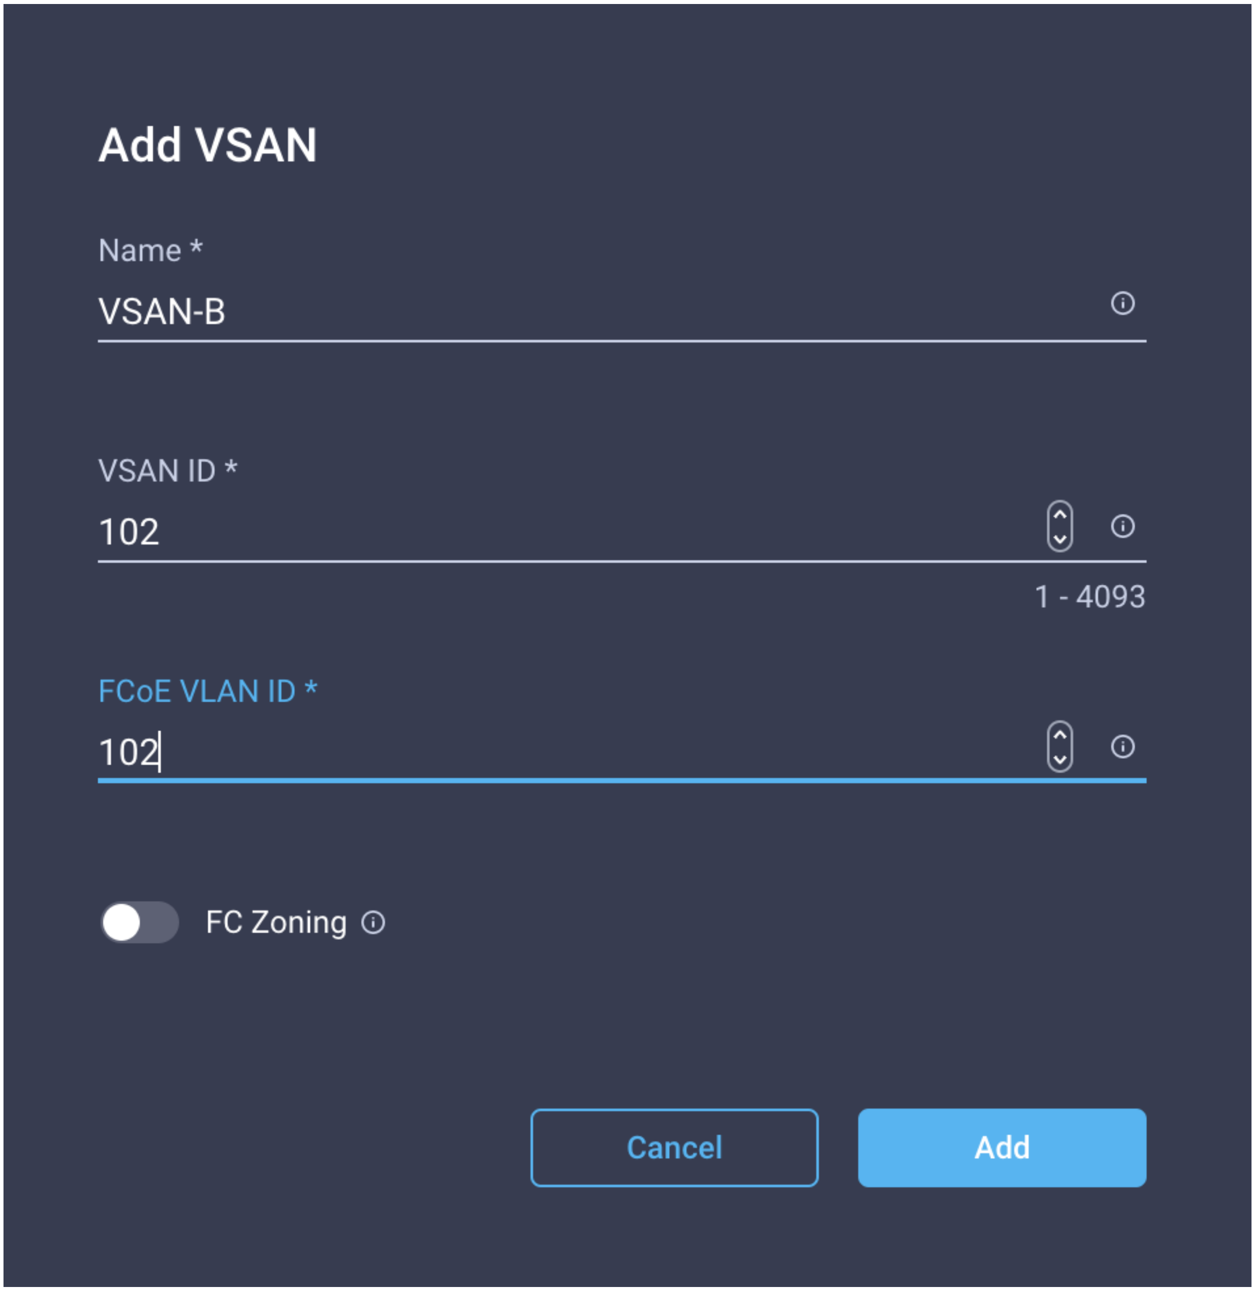

Add VSAN

Name *

VSAN-B

VSAN ID *

102

1 - 4093

FCoE VLAN ID *

102

FC Zoning

16. Enable Uplink Trunking for this VSAN.

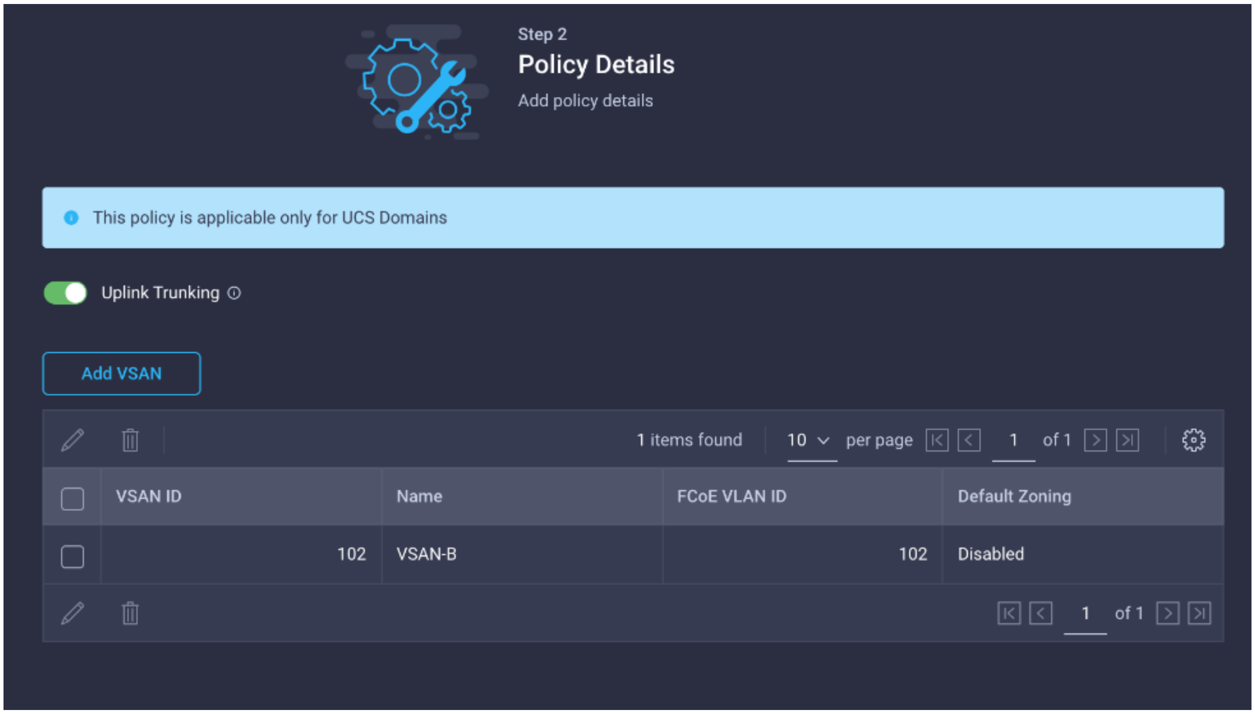

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Uplink Trunking: on

1 items found | 10 per page | 1 of 1

	VSAN ID	Name	FCoE VLAN ID	Default Zoning
□	102	VSAN-B	102	Disabled

1 of 1

17. Click Create.

18. Verify that a common VLAN policy and two unique VSAN policies are associated with the two FIs.

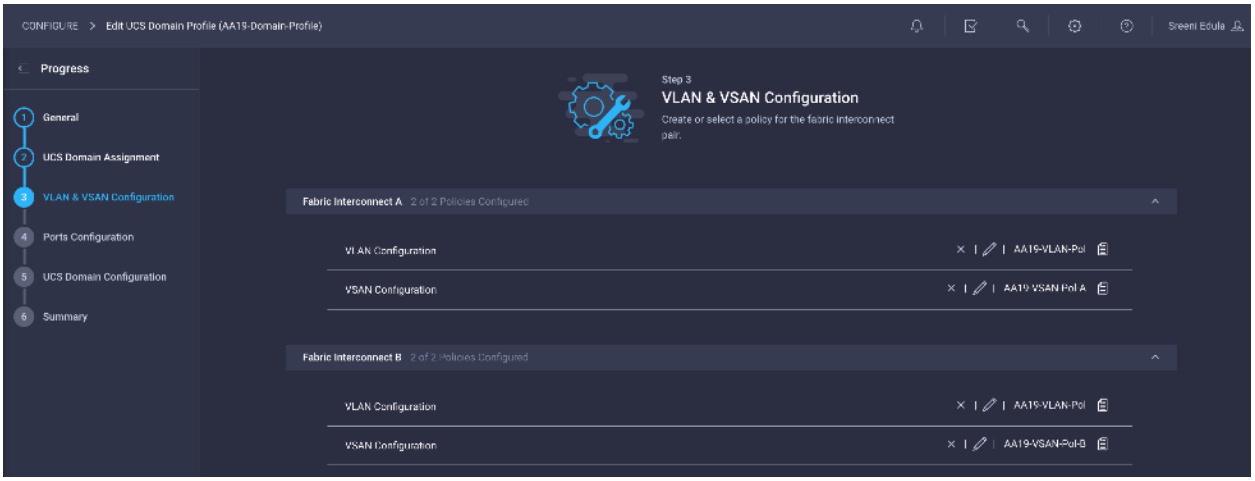

CONFIGURE > Edit UCS Domain Profile (AA19-Domain-Profile)

Step 3

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 2 of 2 Policies Configured

VLAN Configuration: AA19-VLAN-Pol

VSAN Configuration: AA19 VSAN Pol A

Fabric Interconnect B 2 of 2 Policies Configured

VLAN Configuration: AA19-VLAN-Pol

VSAN Configuration: AA19-VSAN-Pol-B

19. Click Next.

Step 4 – Ports Configuration

To configure the ports on Fabric Interconnects, follow these steps:

1. Click Select Policy for Fabric Interconnect A.

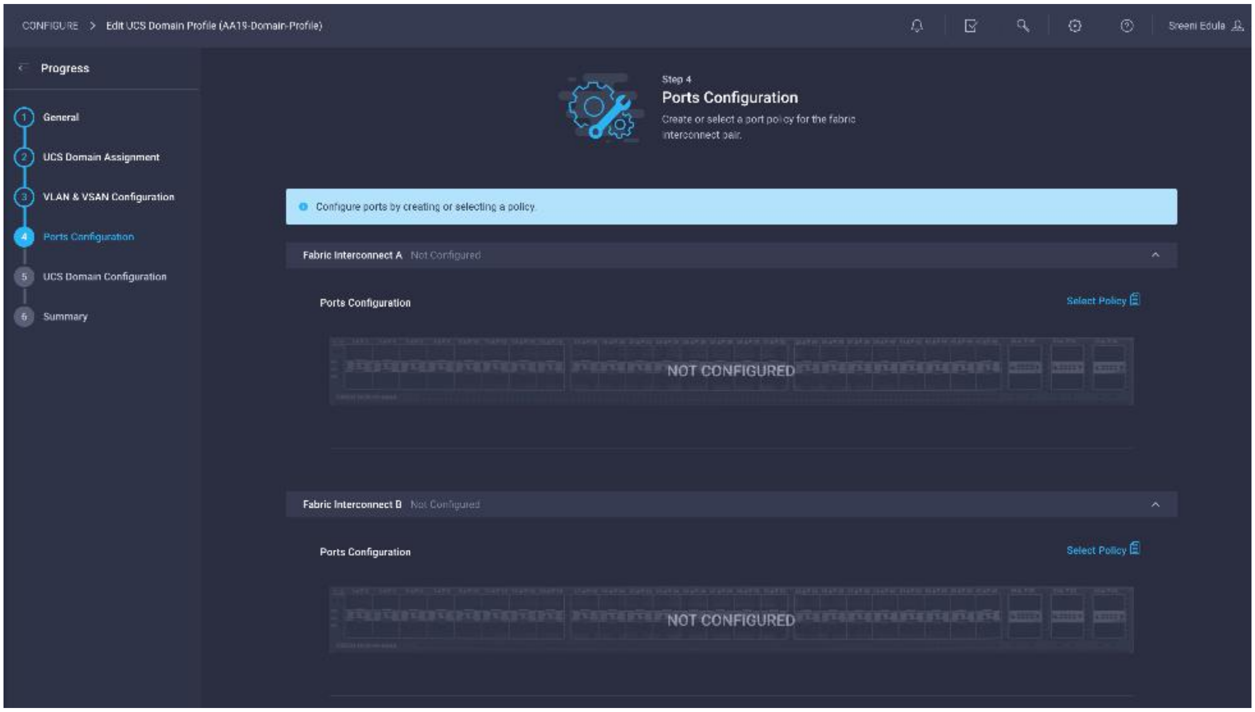

CONFIGURE > Edit UCS Domain Profile (AA19-Domain-Profile)

Sreeni Edula

Progress

1 General

2 UCS Domain Assignment

3 VLAN & VSAN Configuration

4 Ports Configuration

5 UCS Domain Configuration

6 Summary

Step 4

Ports Configuration

Create or select a port policy for the fabric interconnect pair.

Configure ports by creating or selecting a policy.

Fabric Interconnect A Not Configured

Ports Configuration

Select Policy

NOT CONFIGURED

Fabric Interconnect B Not Configured

Ports Configuration

Select Policy

NOT CONFIGURED

2. Click Create New in the right-hand pane to define new port configuration policy.

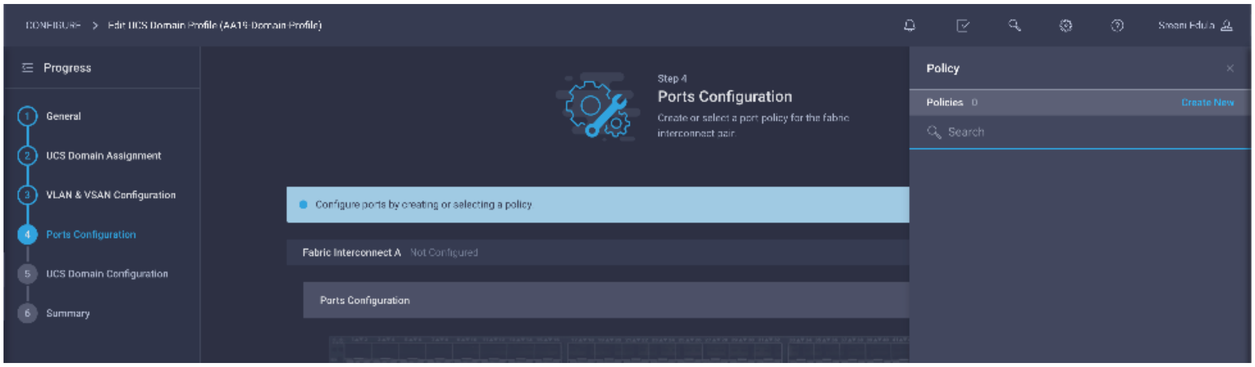

CONFIGURE > Edit UCS Domain Profile (AA19-Domain-Profile)

Sreeni Edula

Progress

1 General

2 UCS Domain Assignment

3 VLAN & VSAN Configuration

4 Ports Configuration

5 UCS Domain Configuration

6 Summary

Step 4

Ports Configuration

Create or select a port policy for the fabric interconnect pair.

Configure ports by creating or selecting a policy.

Fabric Interconnect A Not Configured

Ports Configuration

Policy

Policies 0

Create New

Search

Note: Use two separate port policies for the fabric interconnects. Using separate policies provide flexibility when port configuration (port numbers or speed) differs between the two FIs. When configuring Fibre Channel, two policies are required because each fabric interconnect uses a unique Fibre Channel and VSAN.

3. Choose the organization from the drop-down list.

4. Provide a name for the policy (for example, AA19-Port-Pol-A)

5. From the drop-down list, select the correct FI model under the Switch Model (for example, UCS-FI-6454).

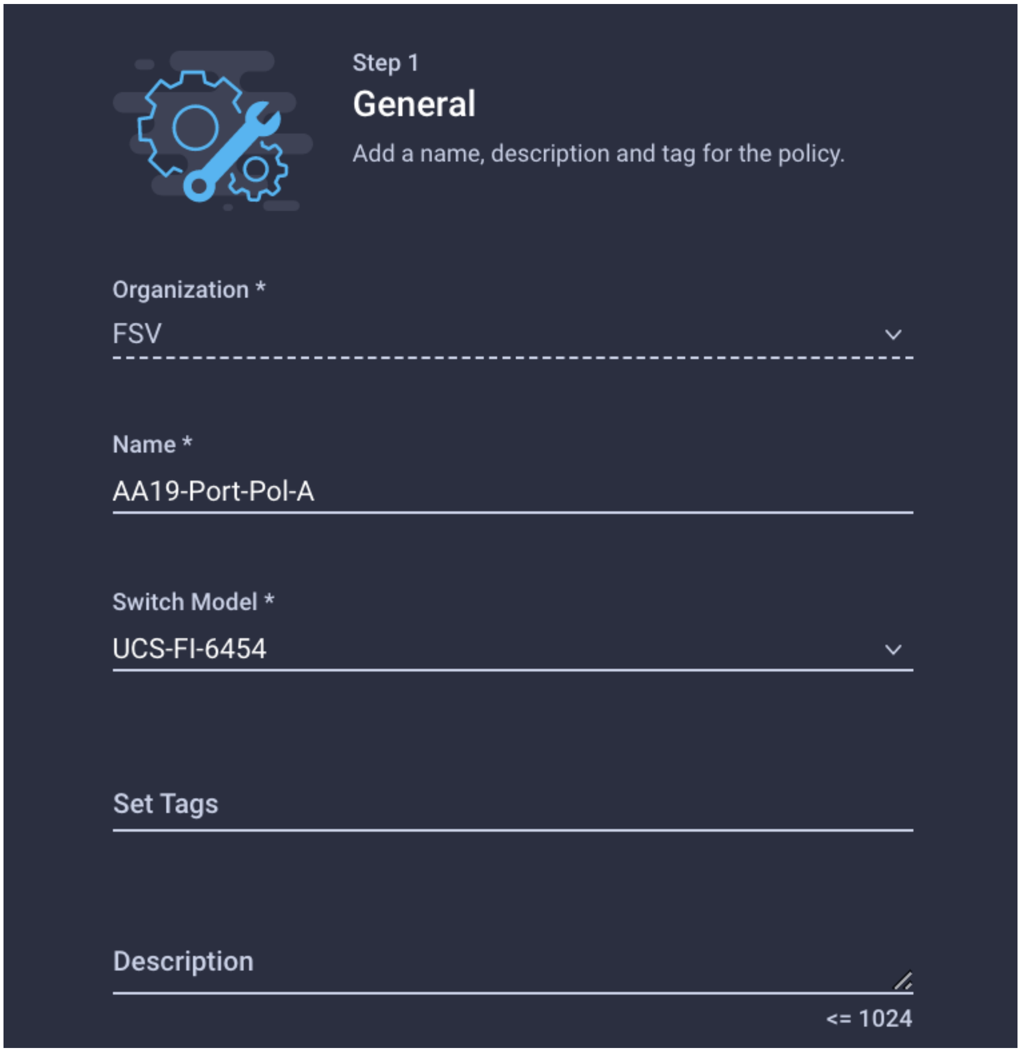

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-Port-Pol-A

Switch Model *

UCS-FI-6454

Set Tags

Description

<= 1024

6. Click Next.

7. Move the slider to set up unified ports. In this example, the first four ports were selected as Fibre Channel ports. Click Next.

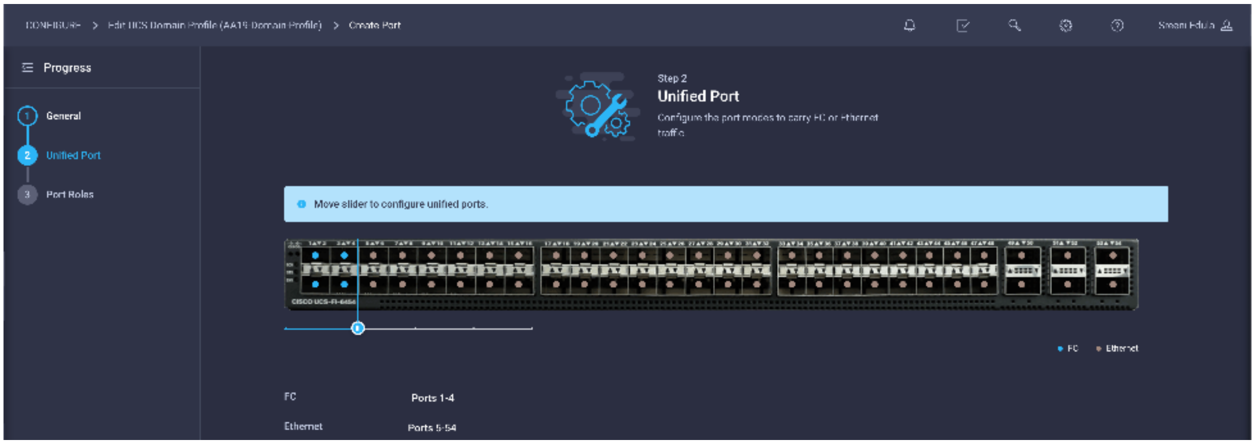

Step 2

Unified Port

Configure the port modes to carry FC or Ethernet traffic.

Move slider to configure unified ports.

FC Ports 1-4

Ethernet Ports 5-54

8. Verify the ports 1-4 are indeed configured as FC ports.

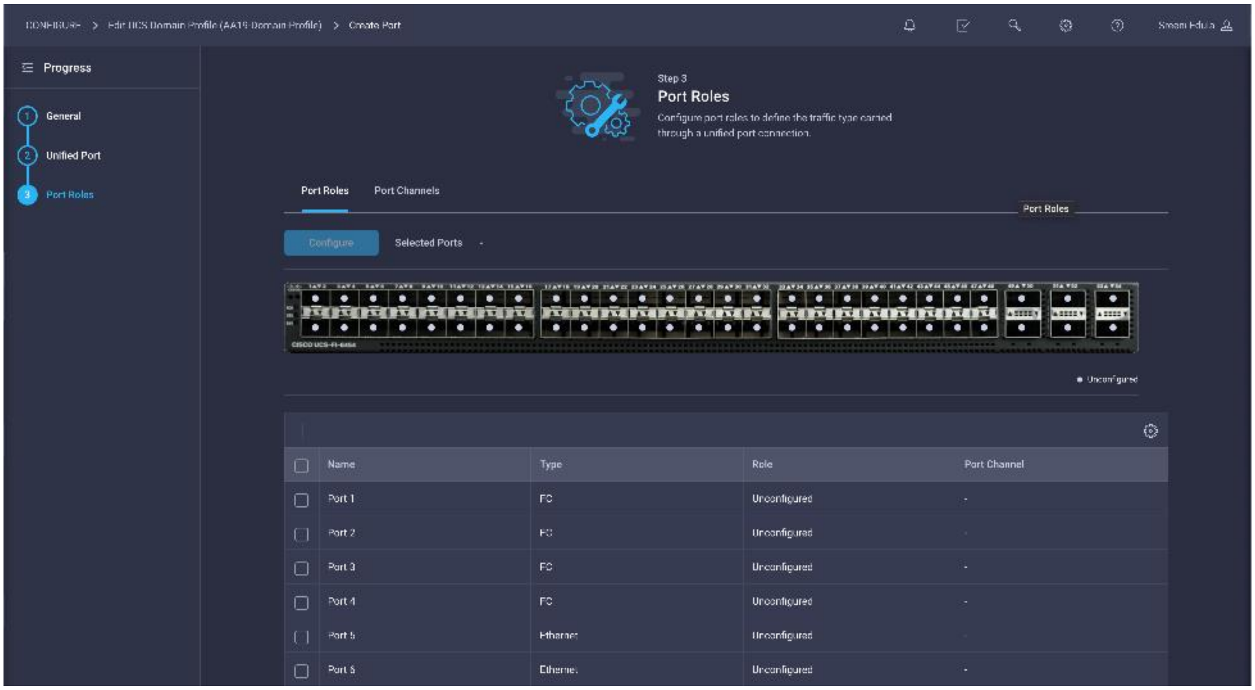
CONFIGURE > Edit UCS Domain Profile (AA19 Domain Profile) > Create Port
Sreeni Edula
Progress
1 General
2 Unified Port
3 Port Roles
Step 3
Port Roles
Configure port roles to define the traffic type carried through a unified port connection.
Port Roles
Port Channels
Port Roles
Configure
Selected Ports -
CISCO UCS-FI-6454
● Unconfigured

☐	Name	Type	Role	Port Channel
☐	Port 1	FC	Unconfigured	-
☐	Port 2	FC	Unconfigured	-
☐	Port 3	FC	Unconfigured	-
☐	Port 4	FC	Unconfigured	-
☐	Port 5	Ethernet	Unconfigured	-
☐	Port 6	Ethernet	Unconfigured	-

9. Select all the ports that need to be configured as server ports by clicking the ports in the graphics (or from the list below the graphic). When all ports are selected, click Configure.

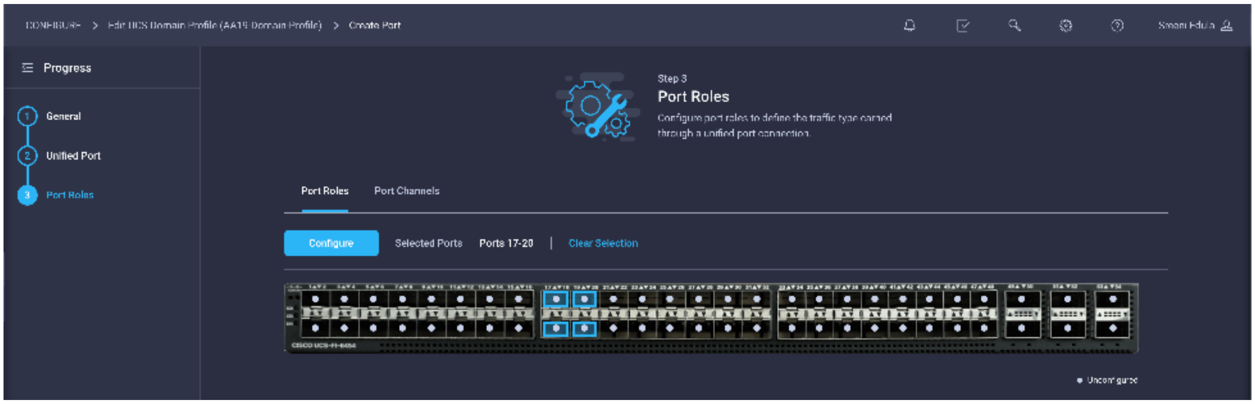
CONFIGURE > Edit UCS Domain Profile (AA19 Domain Profile) > Create Port
Sreeni Edula
Progress
1 General
2 Unified Port
3 Port Roles
Step 3
Port Roles
Configure port roles to define the traffic type carried through a unified port connection.
Port Roles
Port Channels
Configure
Selected Ports Ports 17-20 | Clear Selection
CISCO UCS-FI-6454
● Unconfigured

10. From the drop-down list, select Server as the Role. Click Save.

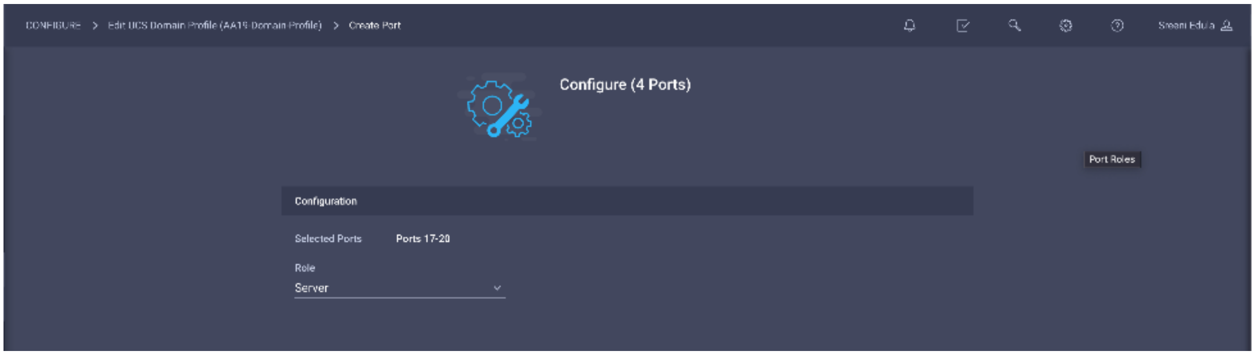
CONFIGURE > Edit UCS Domain Profile (AA19 Domain Profile) > Create Port
Sreeni Edula
Configure (4 Ports)
Port Roles
Configuration

Selected Ports Ports 17-20

Role
Server

11. Configure Uplink Ethernet Port-Channel by selecting the Port Channel in the main pane and then clicking Create Port Channel.

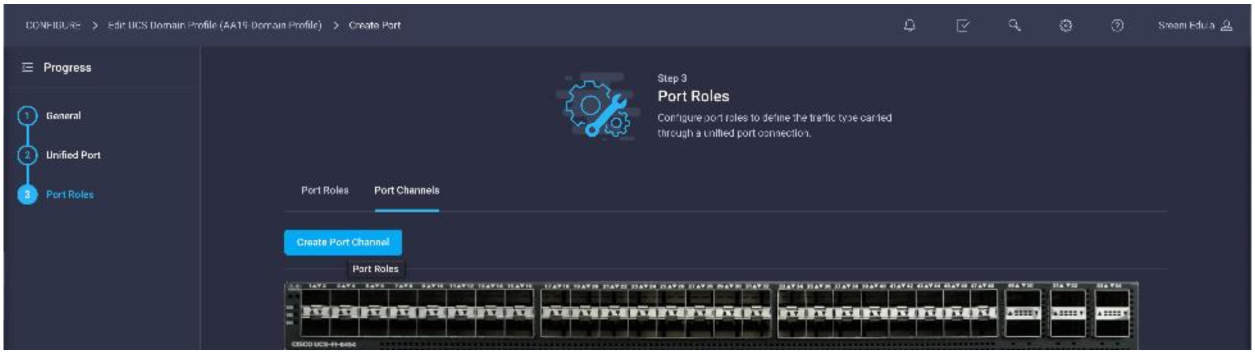

Step 3

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

12. Select Ethernet Uplink Port Channel as the Role, provide a Port-Channel ID (for example, 11) and the Admin Speed (for example, Auto).

Note: Customers can create the Ethernet Network Group, Flow Control, Ling Aggregation, or Link control policy for defining disjoint Layer-2 domain or fine-tune port-channel parameters. These policies were not configured for this deployment.

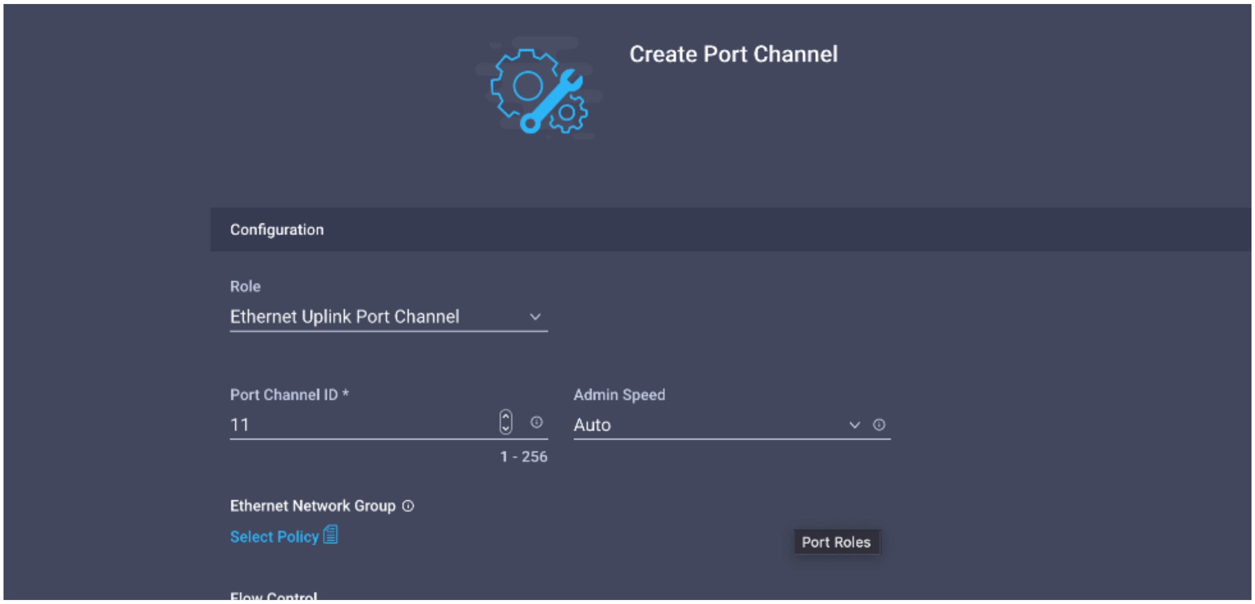

Role: Ethernet Uplink Port Channel

Port Channel ID *: 11

Admin Speed: Auto

13. Click Select Policy under Flow Control Policy in the pane on the left, click Create New.

14. Verify the correct Organization is selected (for example, FSV).

15. Provide a name for the policy (for example, AA19-FlowControl-Pol).

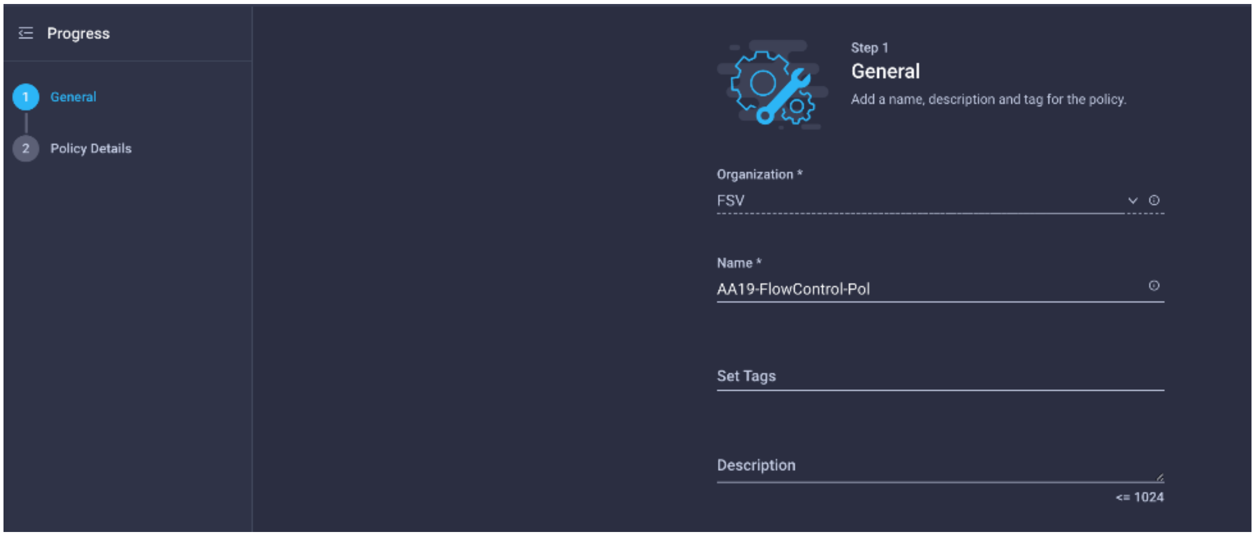

Step 1

General

Add a name, description and tag for the policy.

Organization * FSV

Name * AA19-FlowControl-Pol

Set Tags

Description

<= 1024

16. Click Next.

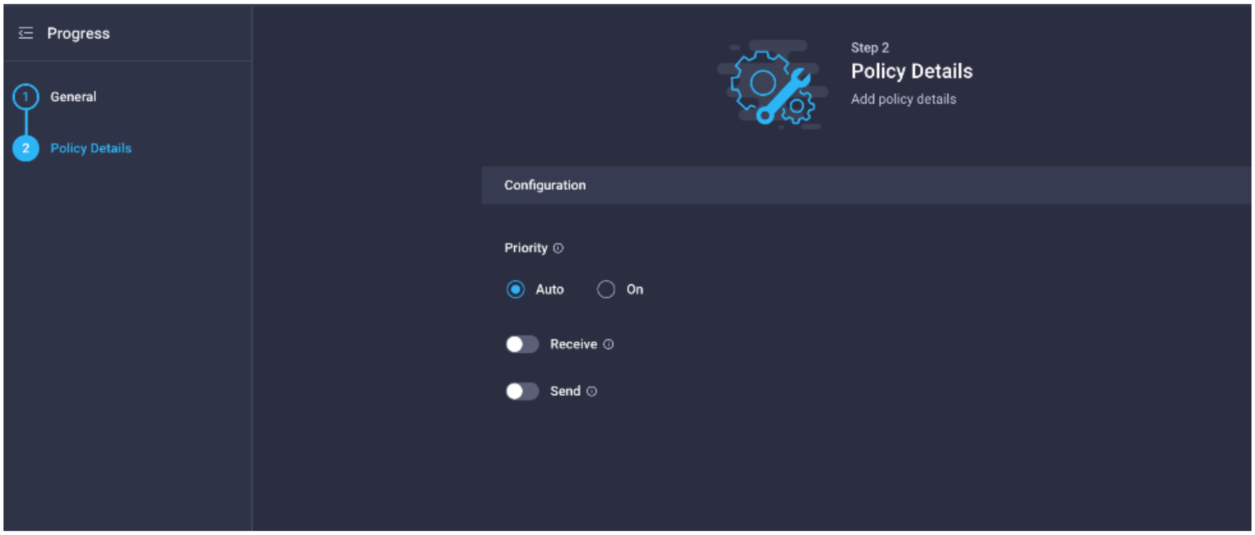

Step 2

Policy Details

Add policy details

Configuration

Priority

● Auto ○ On

Receive

Send

17. Click Select Policy under Link Aggregation Policy in the pane on the left, click Create New.

18. Verify the correct Organization is selected (for example, FSV).

19. Provide a name for the policy (for example, AA19-LinkAgg-Pol).

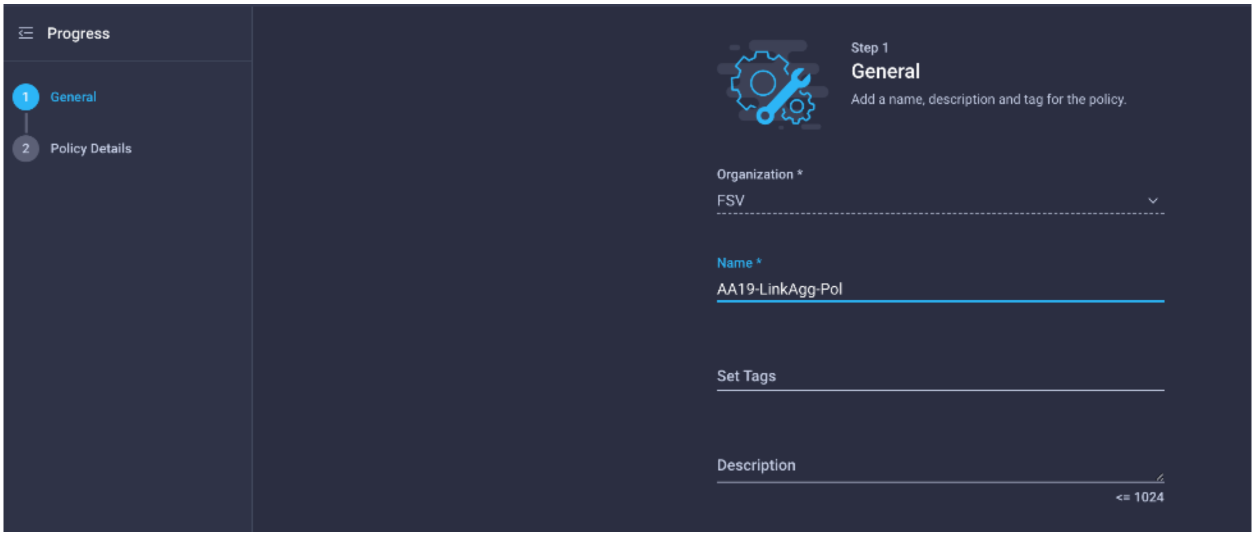

Step 1

General

Add a name, description and tag for the policy.

Organization * FSV

Name * AA19-LinkAgg-Pol

Set Tags

Description

<= 1024

20. Click Next.

21. Make sure you have the following enabled.

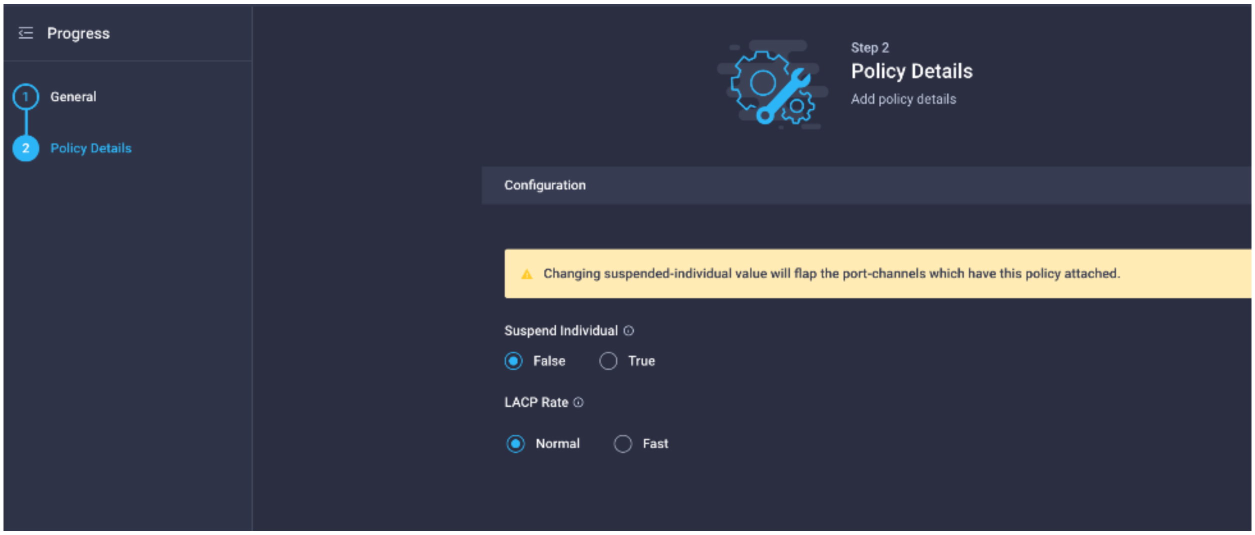

Step 2

Policy Details

Add policy details

Configuration

⚠ Changing suspended-individual value will flap the port-channels which have this policy attached.

Suspend Individual

● False ○ True

LACP Rate

● Normal ○ Fast

22. Click Select Policy under Link Control Policy in the pane on the left, click Create New.

23. Verify the correct Organization is selected (for example, FSV).

24. Provide a name for the policy (for example, AA19-LinkControl-Pol).

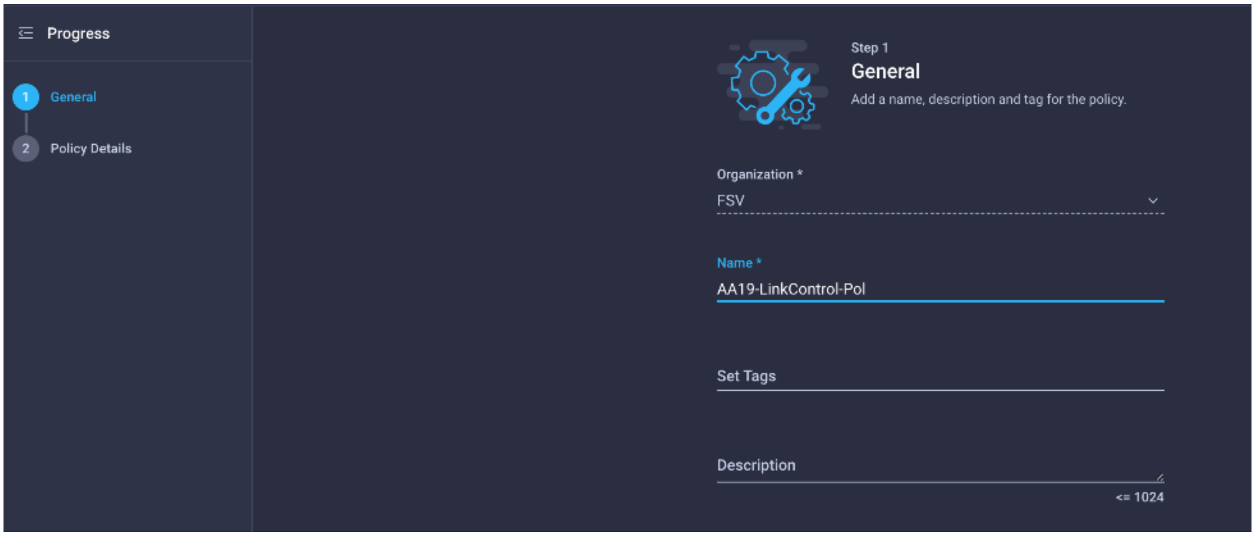

Progress

1 General

2 Policy Details

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-LinkControl-Pol

Set Tags

Description

<= 1024

25. Click Next.

26. Keep Admin State enabled and Mode as Normal.

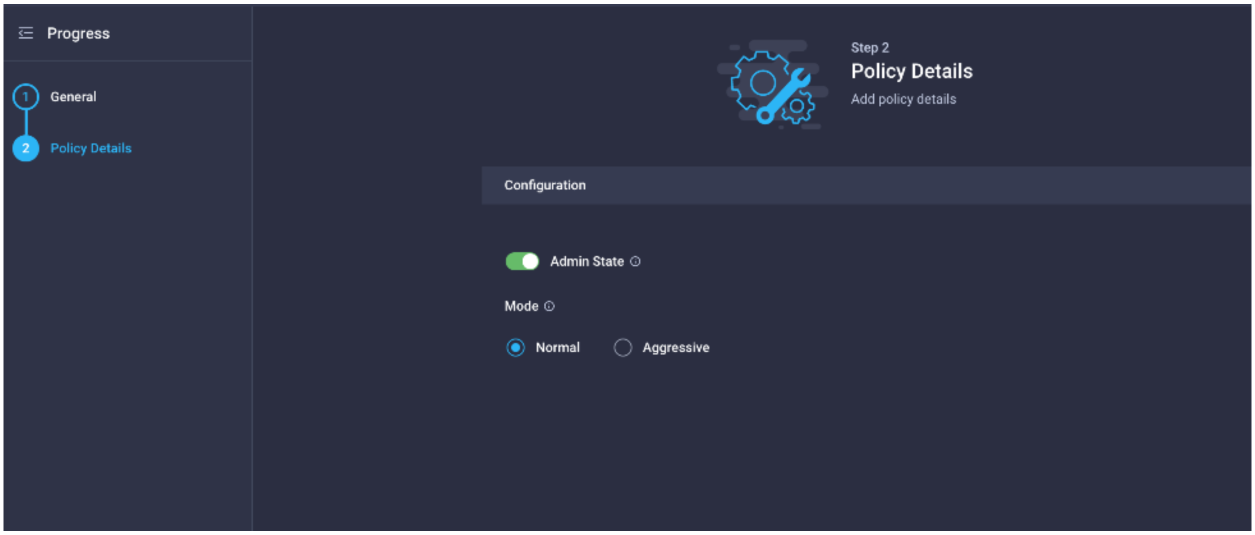

Progress

1 General

2 Policy Details

Step 2

Policy Details

Add policy details

Configuration

Admin State: enabled

Mode

● Normal ○ Aggressive

27. Select ports that are going to be the members of Port Channel under Select Ports dialogue.

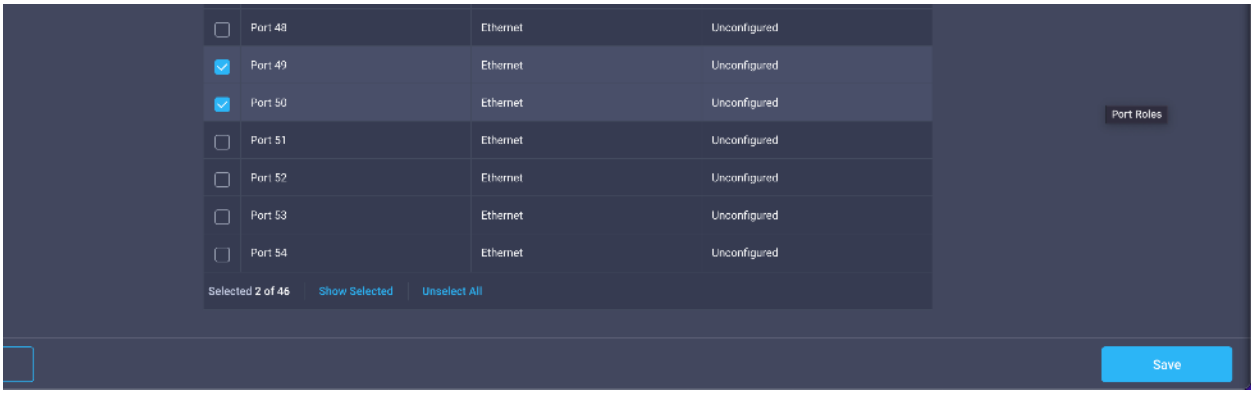

□	Port 48	Ethernet	Unconfigured
☑	Port 49	Ethernet	Unconfigured
☑	Port 50	Ethernet	Unconfigured
□	Port 51	Ethernet	Unconfigured
□	Port 52	Ethernet	Unconfigured
□	Port 53	Ethernet	Unconfigured
□	Port 54	Ethernet	Unconfigured

Selected 2 of 46

28. Click Save.

Configure FC Port Channel (FC configuration only)

FC uplink port channel is only needed when configuring SAN and can be skipped when configuring IP-only storage access. To configure FC port-channel, follow these steps:

- Name of the port policy: AA19-Port-Pol-A
- Ethernet port-Channel ID: 11
- FC port-channel ID: 101
- FC VSAN ID: 101

1. Configure Fibre Channel Port-Channel from the main pane again by clicking Create Port Channel.

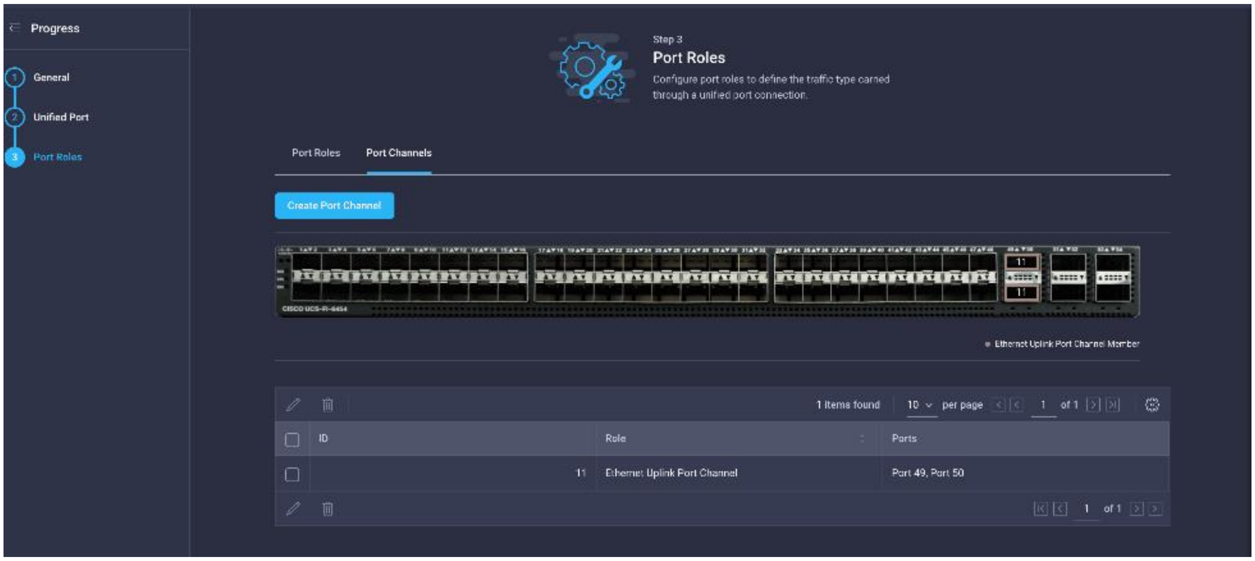

Step 3

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

1 items found 10 per page 1 of 1

☐	ID	Role	Ports
☐	11	Ethernet Uplink Port Channel	Port 49, Port 50

1 of 1

2. In the drop-down list under Role, choose FC Uplink Port Channel.

3. Provide a port-channel ID (for example, 101), choose a value for Admin Speed (for example, 32Gbps), and provide a VSAN ID (for example, 101).

4. Select the Port Channel member Ports.

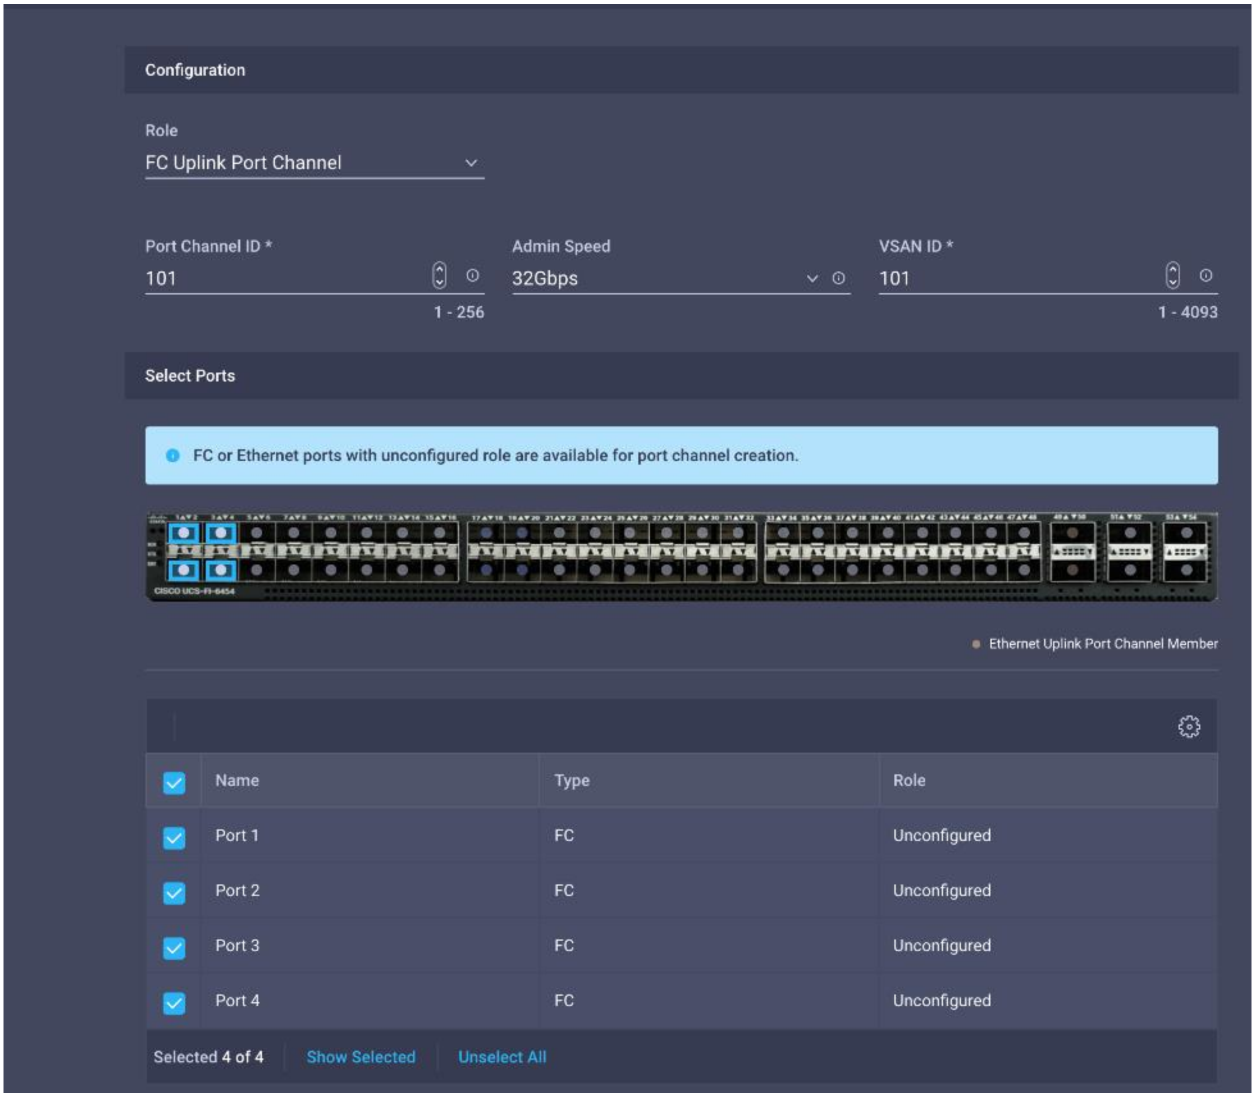

Configuration

Role
FC Uplink Port Channel

Port Channel ID *
101
1 - 256

Admin Speed
32Gbps

VSAN ID *
101
1 - 4093

Select Ports

FC or Ethernet ports with unconfigured role are available for port channel creation.

● Ethernet Uplink Port Channel Member

☑	Name	Type	Role
☑	Port 1	FC	Unconfigured
☑	Port 2	FC	Unconfigured
☑	Port 3	FC	Unconfigured
☑	Port 4	FC	Unconfigured

Selected 4 of 4 Show Selected Unselect All

5. Click Save.

6. Verify the port-channel IDs and ports after both the Ethernet uplink port channel and the Fibre Channel uplink port channel have been created.

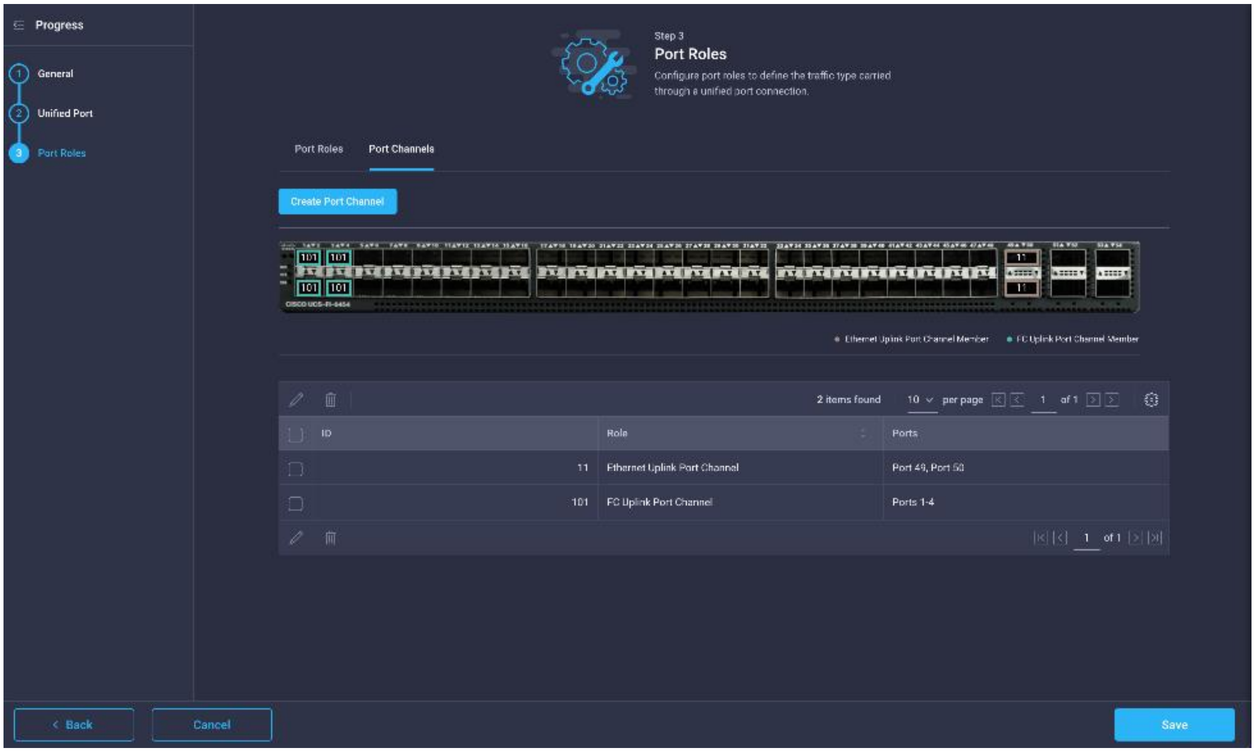

ID	Role	Ports
11	Ethernet Uplink Port Channel	Port 49, Port 50
101	FC Uplink Port Channel	Ports 1-4

7. Click Save to create the port policy for FI-A. The summary screen, shown below, can be used to verify ports were selected and configured correctly.

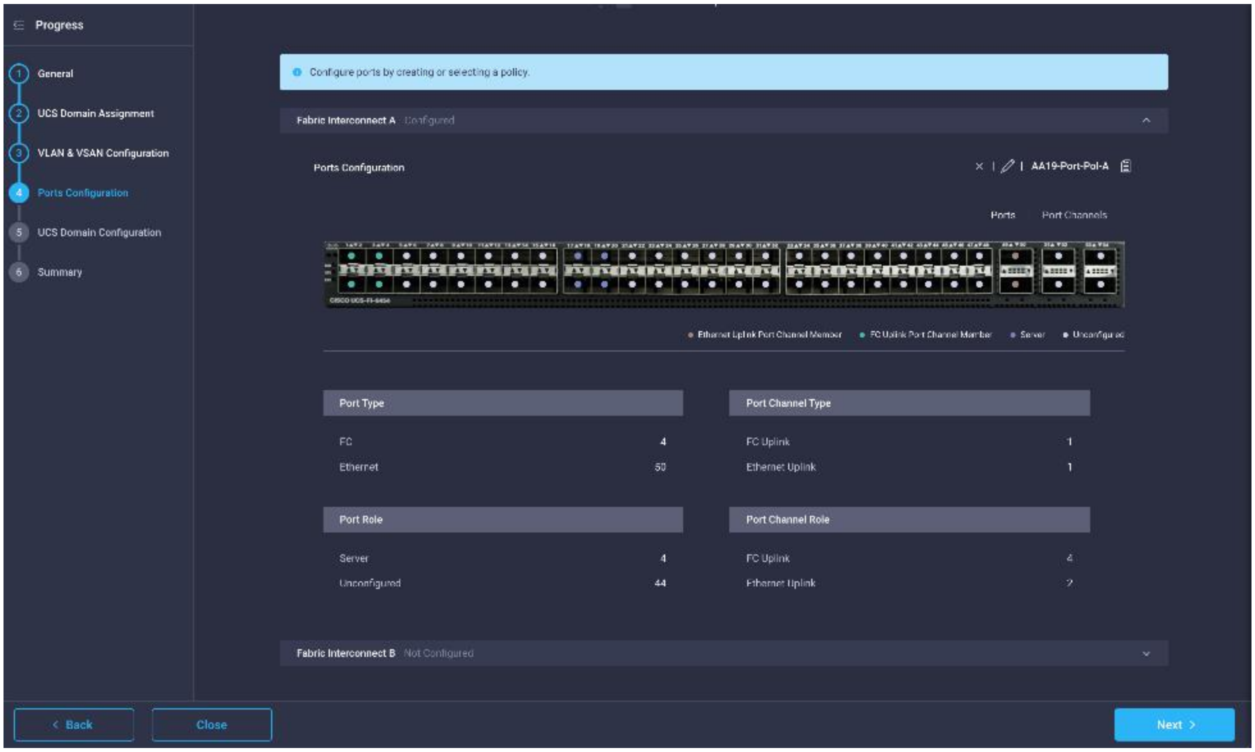
Progress
1 General
2 UCS Domain Assignment
3 VLAN & VSAN Configuration
4 Ports Configuration
5 UCS Domain Configuration
6 Summary
Configure ports by creating or selecting a policy.
Fabric Interconnect A Configured
Ports Configuration
AA19-Port-Pol-A
Ports
Port Channels
CISCO UCS-FI-6454
Ethernet Uplink Port Channel Member
FC Uplink Port Channel Member
Server
Unconfigured
Fabric Interconnect B Not Configured
< Back
Close
Next >

Port Type	
FC	4
Ethernet	50

Port Channel Type	
FC Uplink	1
Ethernet Uplink	1

Port Role	
Server	4
Unconfigured	44

Port Channel Role	
FC Uplink	4
Ethernet Uplink	2

Port configuration for Fabric Interconnect B

Repeat the steps from section Configure FC Port Channel (FC configuration only) to create the port policy for Fabric Interconnect B including the Ethernet port-channel and the FC port-channel (if configuring SAN). Use the following values for various parameters:

- Name of the port policy: AA19-Port-Pol-B
- Ethernet port-Channel ID: 12
- FC port-channel ID: 102
- FC VSAN ID: 102

1. Click Select Policy for Fabric Interconnect B, and in the pane at the right, click Create New.
2. Verify the Organization from the drop-down list (for example, FSV).
3. Provide a name for the policy (for example, AA19-Port-Pol-B).
4. Select the correct UCS FI Model under the Switch Model (for example, UCS-FI-6454).
5. Click Next.
6. Repeat the steps you used for Fabric Interconnect A to configure Fibre Channel ports, server ports, and Ethernet and Fibre Channel port channels with appropriate IDs.

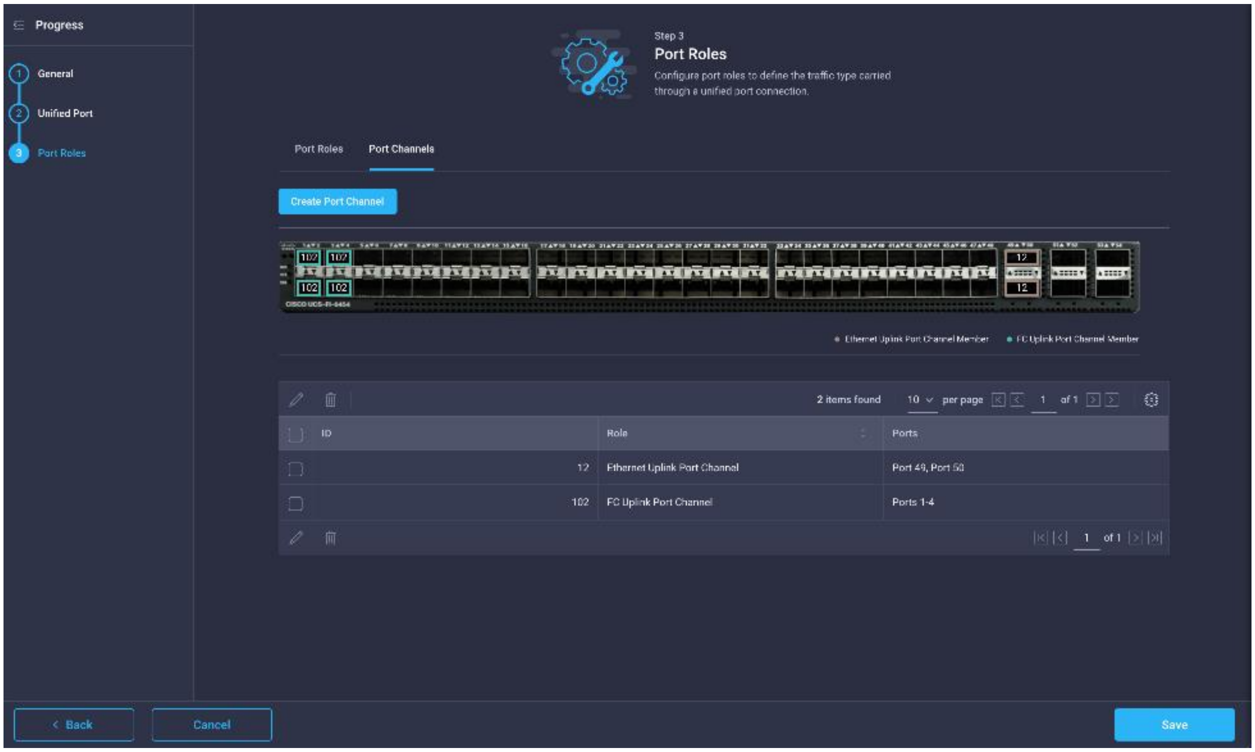

Step 3

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

ID	Role	Ports
12	Ethernet Uplink Port Channel	Port 49, Port 50
102	FC Uplink Port Channel	Ports 1-4

7. Click Save.
8. Use the summary screen shown here to verify that the ports were selected and configured correctly for Fabric Interconnect B.

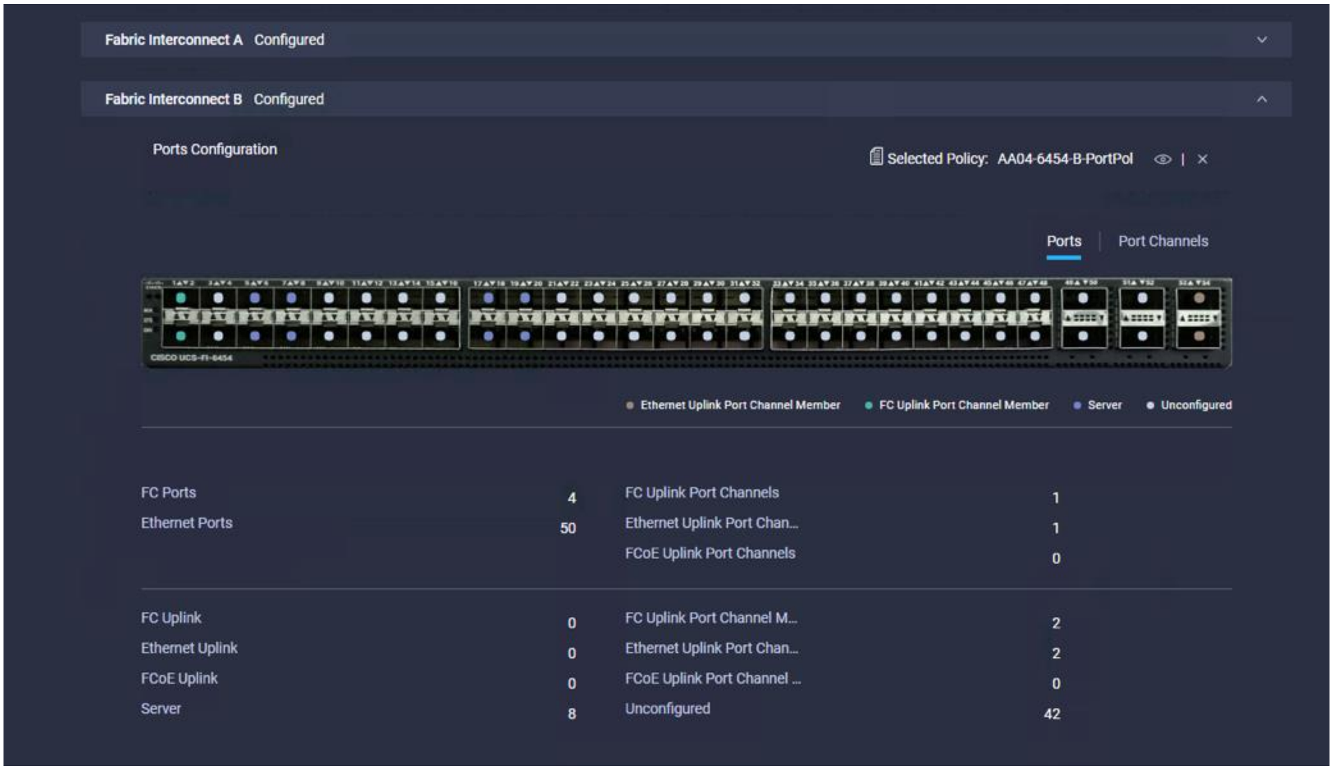

Fabric Interconnect A Configured

Fabric Interconnect B Configured

Ports Configuration

Selected Policy: AA04-6454-B-PortPol

● Ethernet Uplink Port Channel Member ● FC Uplink Port Channel Member ● Server ● Unconfigured

FC Ports	4	FC Uplink Port Channels	1
Ethernet Ports	50	Ethernet Uplink Port Chan...	1
		FCoE Uplink Port Channels	0
FC Uplink	0	FC Uplink Port Channel M...	2
Ethernet Uplink	0	Ethernet Uplink Port Chan...	2
FCoE Uplink	0	FCoE Uplink Port Channel ...	0
Server	8	Unconfigured	42

9. When the port configuration for both FIs is complete and looks good, click Next.

Step 5 – UCS Domain Configuration

Some additional policies such as NTP, Network Connectivity and System QoS need to be defined for the UCS Domain Configuration.

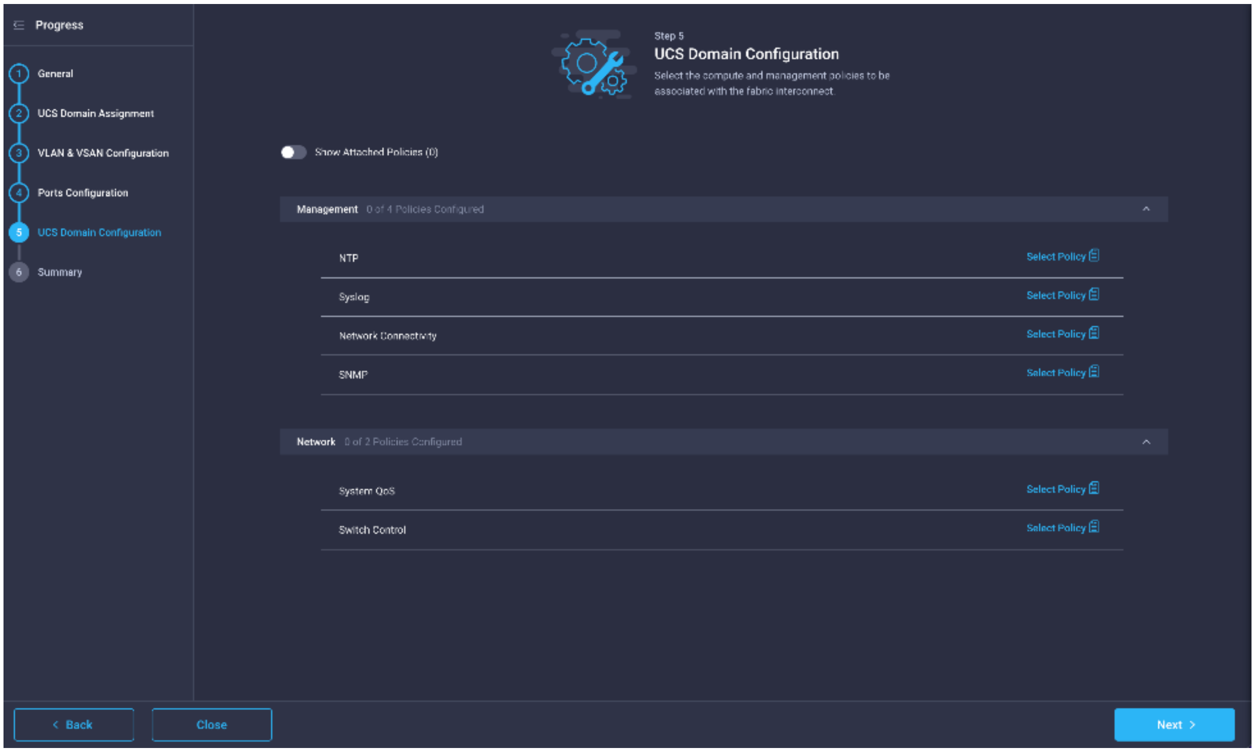

Step 5

UCS Domain Configuration

Select the compute and management policies to be associated with the fabric interconnect

Show Attached Policies (0)

Management 0 of 4 Policies Configured

NTP	Select Policy
Syslog	Select Policy
Network Connectivity	Select Policy
SNMP	Select Policy

Network 0 of 2 Policies Configured

System QoS	Select Policy
Switch Control	Select Policy

Configure NTP Policy

To define an NTP server for the UCS Domain, an NTP policy must be configured by following these steps:

1. Click Select Policy next to NTP and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-NTP-Pol).

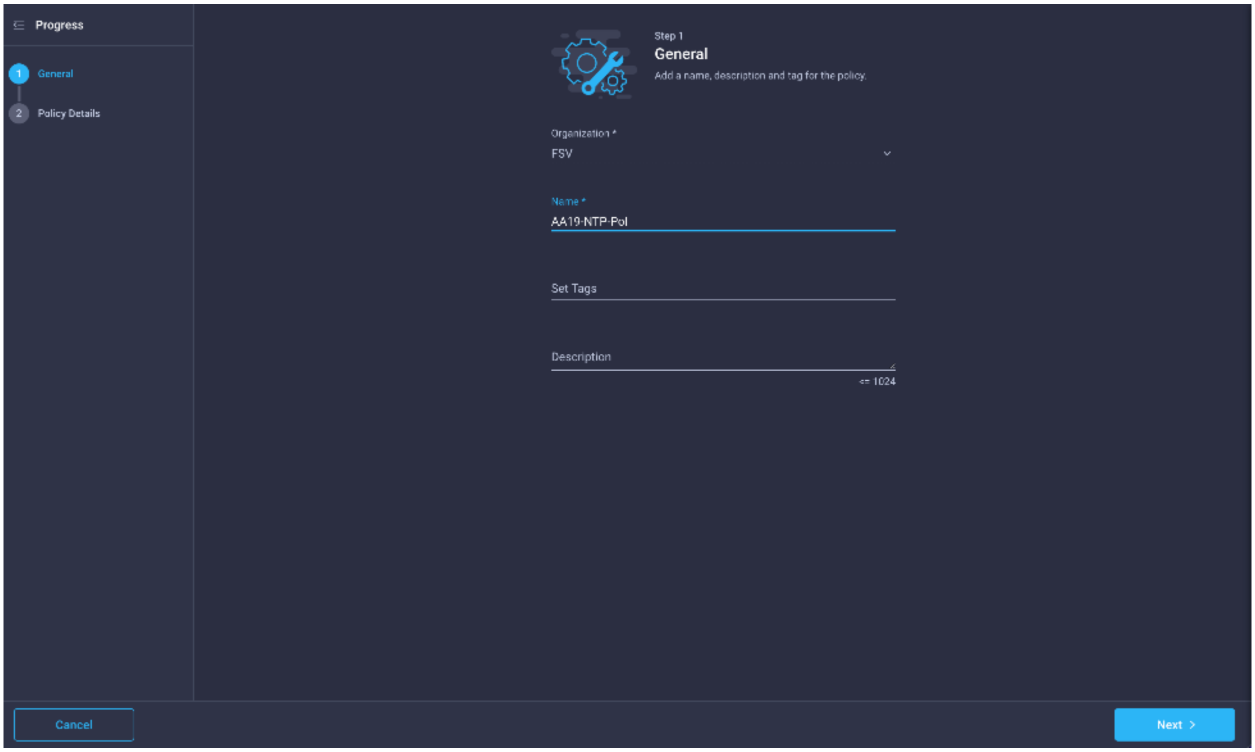

Organization * FSV

Name * AA19-NTP-Pol

Set Tags

Description <= 1024

3. Click Next.

4. Enable NTP, provide the NTP server IP address(s) (for example, 10.81.72.18, 10.81.72.19) and select the Timezone from the drop-down list (for example, America/New_York).

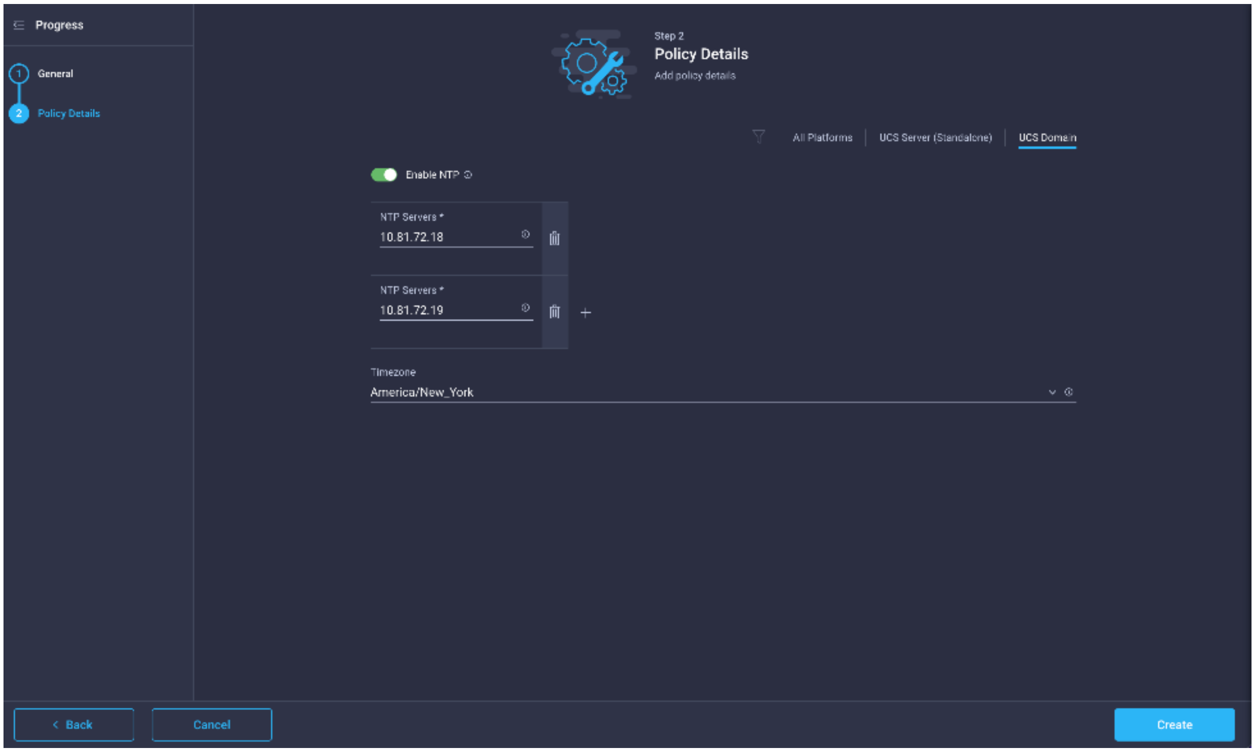

Step 2
Policy Details
Add policy details

Enable NTP

NTP Servers *
10.81.72.18

NTP Servers *
10.81.72.19

Timezone
America/New_York

5. Click Create.

Configure Network Connectivity Policy

To define the DNS servers for the UCS, Network Connectivity Policy must be configured by following these steps:

1. Click Select Policy next to Network Connectivity under Management and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-NetConn-Pol).

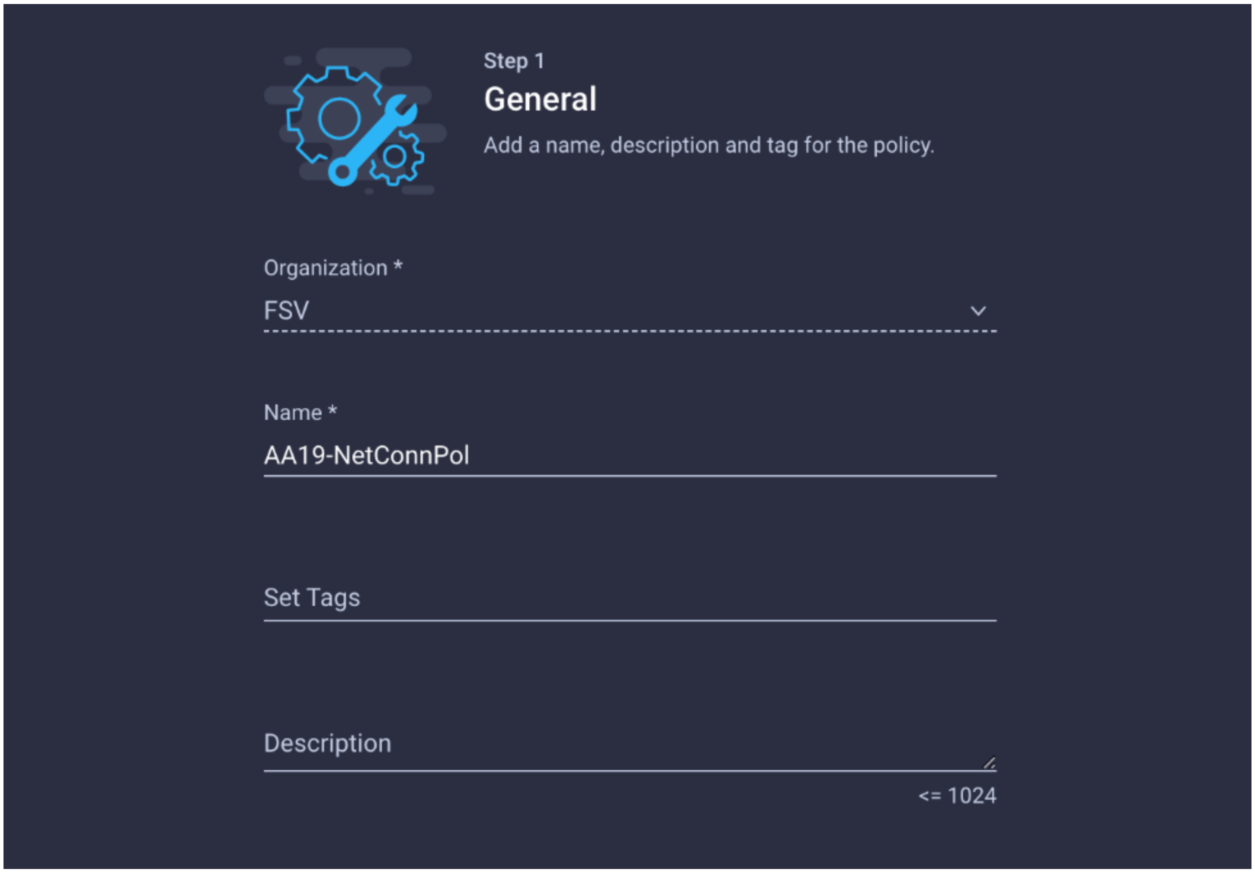

Step 1

General

Add a name, description and tag for the policy.

Organization * FSV

Name * AA19-NetConnPol

Set Tags

Description

<= 1024

3. Provide DNS server IP addresses for the UCS (for example, 10.81.72.40 and 10.81.72.41).

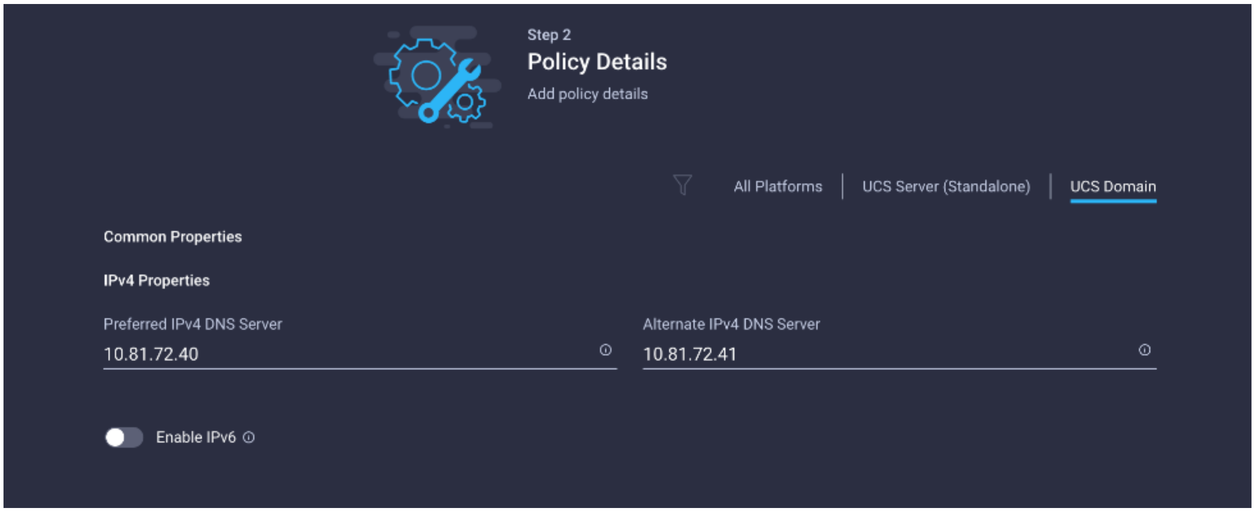

Step 2

Policy Details

Add policy details

Common Properties

IPv4 Properties

Preferred IPv4 DNS Server 10.81.72.40

Alternate IPv4 DNS Server 10.81.72.41

Enable IPv6

4. Click Create.

Configure System QoS Policy

To define the QoS settings for the UCS, QoS System QoS Policy must be configured by following these steps:

1. Click Select Policy next to System QoS under Network and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-QoS-Pol).

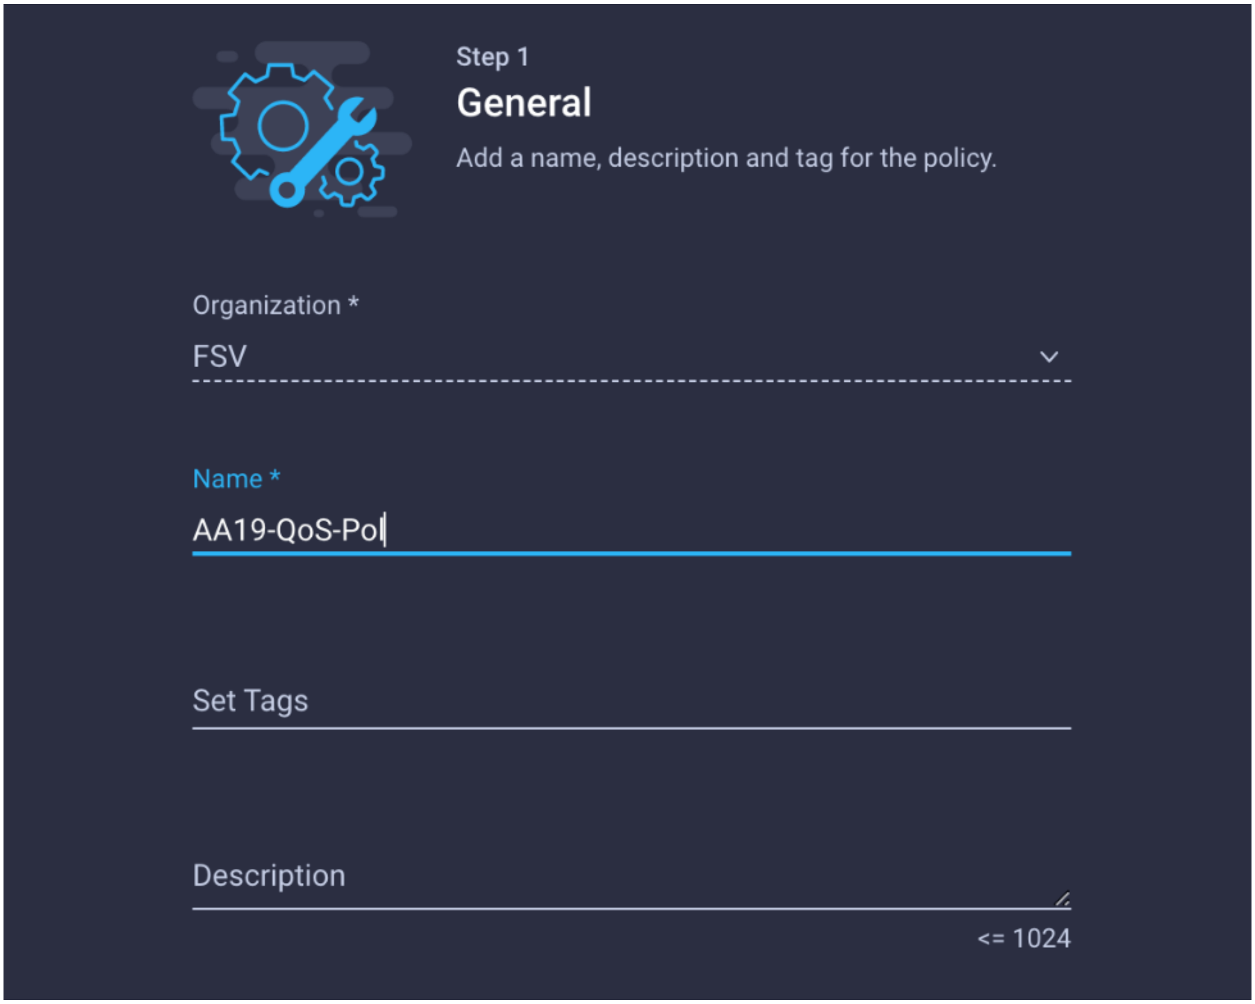

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-QoS-Pol

Set Tags

Description

<= 1024

3. Change the MTU value to 9000 for the Best Effort class and leave the rest to defaults.

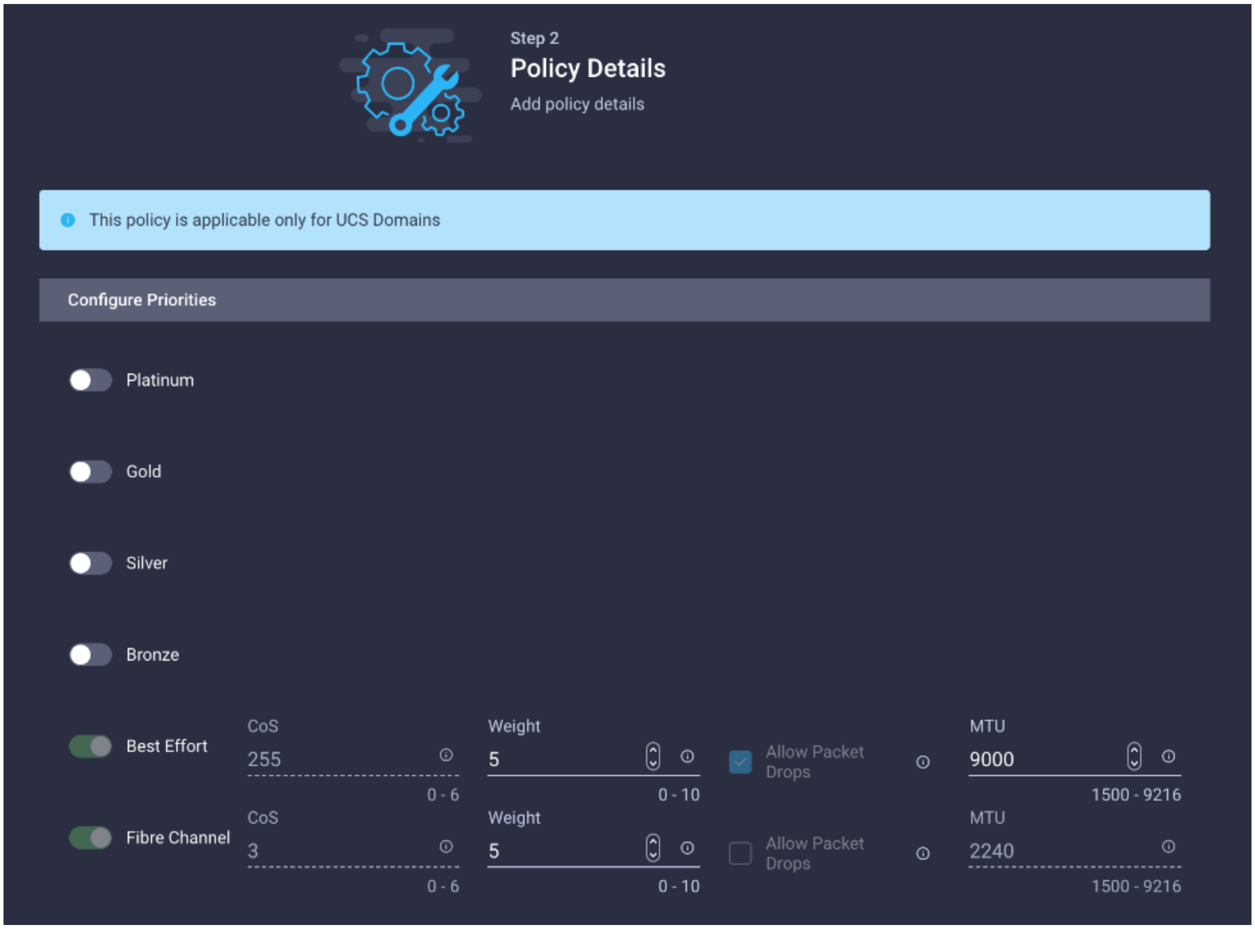

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Configure Priorities

○ Platinum

○ Gold

○ Silver

○ Bronze

	CoS	Weight		MTU
● Best Effort	255 0 - 6	5 0 - 10	☑ Allow Packet Drops	9000 1500 - 9216
● Fibre Channel	3 0 - 6	5 0 - 10	□ Allow Packet Drops	2240 1500 - 9216

4. Click Create.

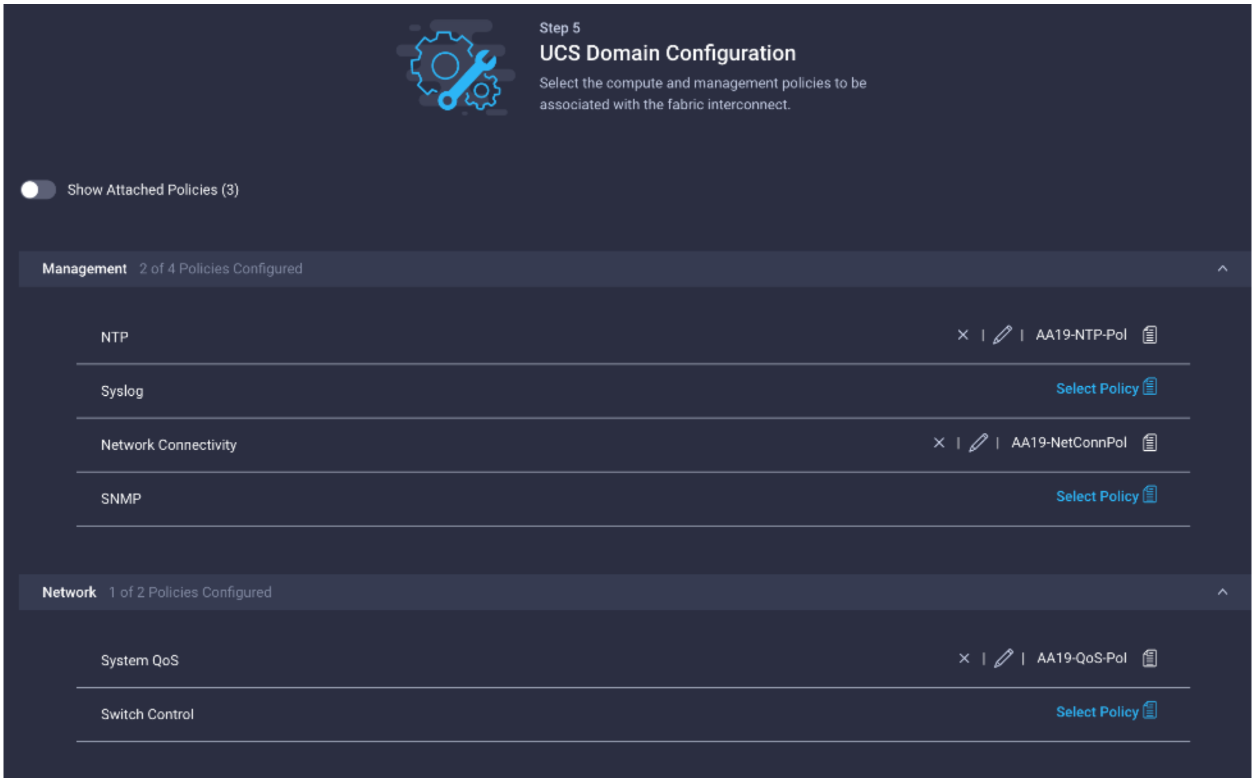

Step 5

UCS Domain Configuration

Select the compute and management policies to be associated with the fabric interconnect.

○ Show Attached Policies (3)

Management 2 of 4 Policies Configured

NTP	AA19-NTP-Pol
Syslog	Select Policy
Network Connectivity	AA19-NetConnPol
SNMP	Select Policy

Network 1 of 2 Policies Configured

System QoS	AA19-QoS-Pol
Switch Control	Select Policy

5. Click Next.

Step 6 – Summary

To verify all the settings (including expanding the Fabric Interconnect settings) and make sure the configuration is correct, follow these steps:

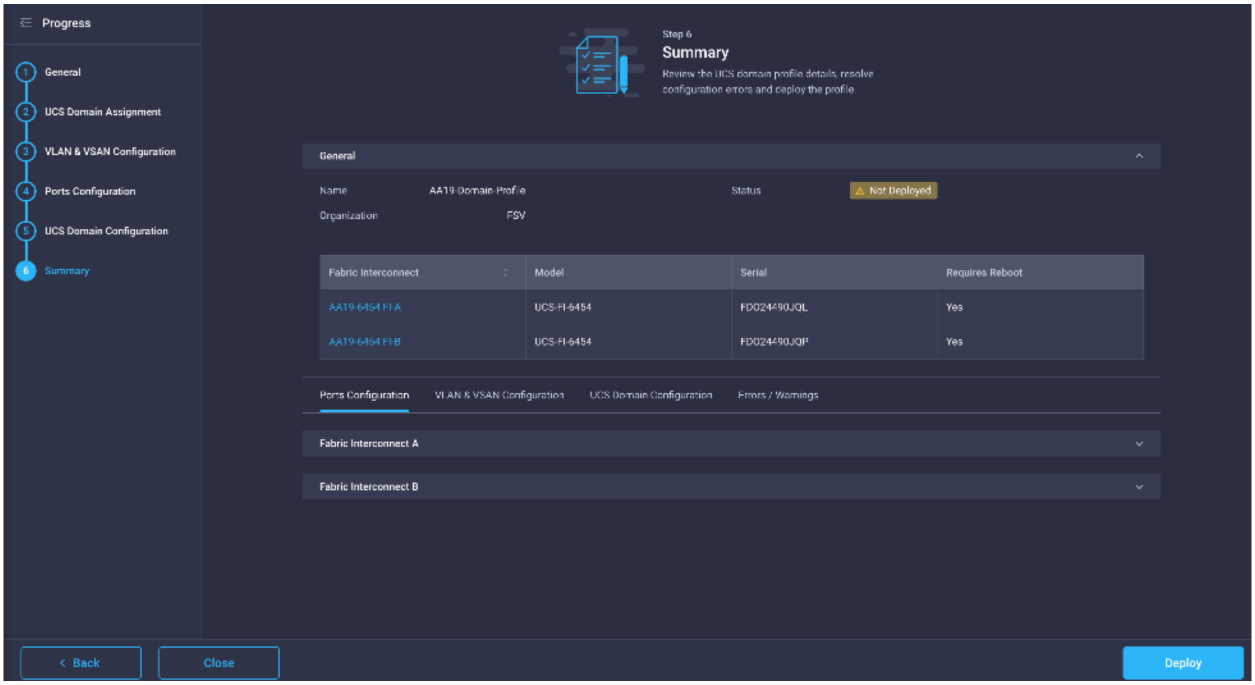

Name AA19-Domain-Profile

Status ⚠ Not Deployed

Organization FSV

Fabric Interconnect	Model	Serial	Requires Reboot
AA19-6454 FI-A	UCS-FI-6454	FDO24490JQL	Yes
AA19-6454 FI-B	UCS-FI-6454	FDO24490JQP	Yes

Deploy the Cisco UCS Domain Profile

After verifying the domain profile configuration, deploy the Cisco UCS profile by following these steps:

6. From the UCS domain profile Summary view, click Deploy.

7. Acknowledge any warnings and click Deploy again.

The system will take some time to validate and configure the settings on the fabric interconnects. Log into the console servers to see when the Cisco UCS fabric interconnects have finished configuration and are successfully rebooted.

Verify Cisco UCS Domain Profile Deployment

When the Cisco UCS domain profile has been successfully deployed, the Cisco UCS chassis and the blades should be successfully discovered.

Note: It takes a while to discover the blades for the first time. Keep an eye on the number of outstanding tasks in Intersight:

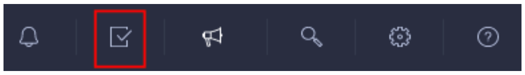

8. Log into the Cisco Intersight. Under CONFIGURE > Profiles > UCS Domain Profiles, verify that the domain profile has been successfully deployed.

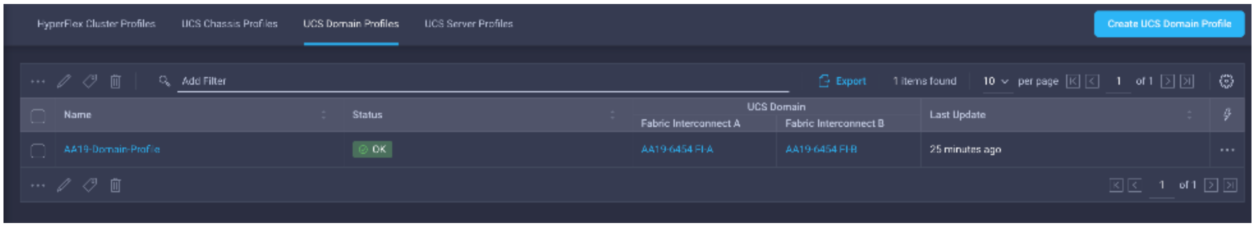

Name	Status	UCS Domain		Last Update
		Fabric Interconnect A	Fabric Interconnect B	
AA19-Domain-Profile	OK	AA19-6454 FI-A	AA19-6454 FI-B	25 minutes ago

9. At this time, the chassis should be discovered and visible under OPERATE->Chassis:

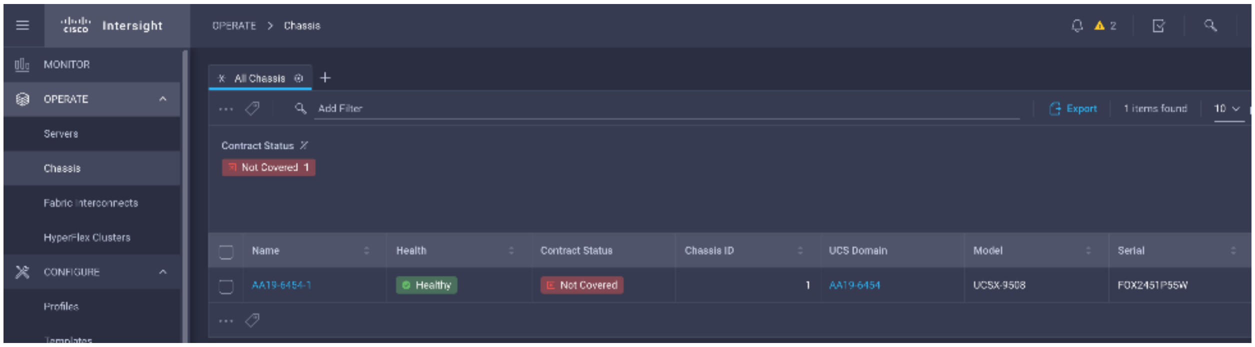

Name	Health	Contract Status	Chassis ID	UCS Domain	Model	Serial
AA19-6454-1	Healthy	Not Covered	1	AA19-6454	UCSX-9508	FOX2451P55W

10. Servers should have been successfully discovered and visible under OPERATE -> Servers.

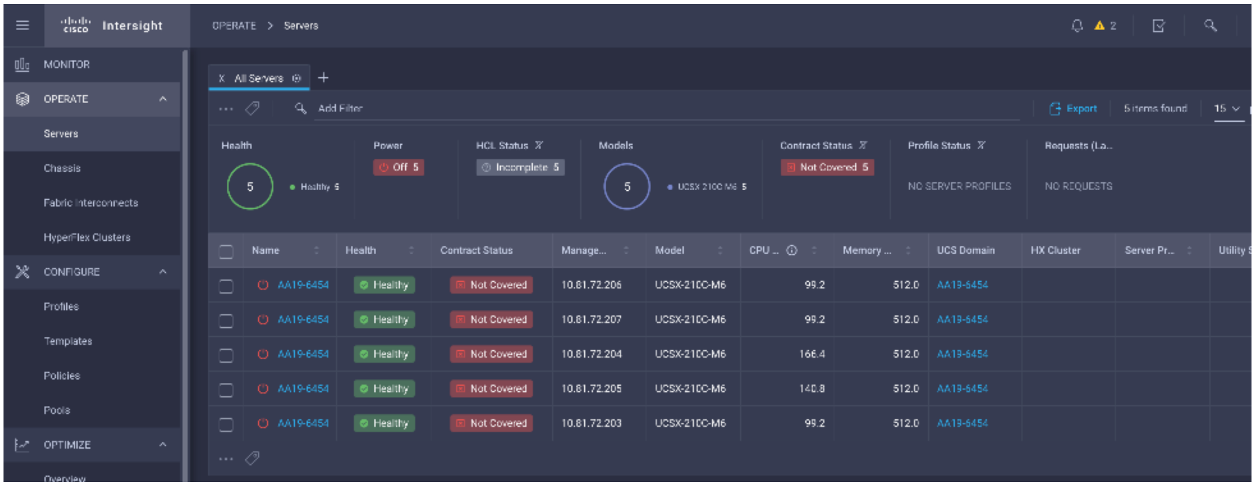

Name	Health	Contract Status	Manage...	Model	CPU ...	Memory ...	UCS Domain	HX Cluster	Server Pr...	Utility
AA19-6454	Healthy	Not Covered	10.81.72.206	UCSX-210C-M6	99.2	512.0	AA19-6454			
AA19-6454	Healthy	Not Covered	10.81.72.207	UCSX-210C-M6	99.2	512.0	AA19-6454			
AA19-6454	Healthy	Not Covered	10.81.72.204	UCSX-210C-M6	166.4	512.0	AA19-6454			
AA19-6454	Healthy	Not Covered	10.81.72.205	UCSX-210C-M6	140.8	512.0	AA19-6454			
AA19-6454	Healthy	Not Covered	10.81.72.203	UCSX-210C-M6	99.2	512.0	AA19-6454			

Configure Cisco UCS Chassis Profile

Cisco UCS Chassis profile in Cisco Intersight allow customers to configure various parameters for chassis, including:

- IMC Access Policy: IP configuration for the in-band chassis connectivity. This setting is independent of Server IP connectivity and only applies to communication to and from chassis.
- SNMP Policy, and SNMP trap settings.
- Power Policy to enable power management and power supply redundancy mode.
- Thermal Policy to control the speed of FANs (only applicable to Cisco UCS 5108)

A chassis policy can be assigned to any number of chassis profiles to provide a configuration baseline for a chassis. In this deployment, no chassis profile was created or attached to the chassis, but customers can configure policies to configure SNMP or Power parameters and attach them to the chassis.

Configure Server Profile template (FCP)

In the Cisco Intersight platform, a server profile enables resource management by simplifying policy alignment and server configuration. The server profiles are derived from a server profile template. The server profile template and its associated policies can be created using the server profile template wizard. After creating a server profile template, customers can derive multiple consistent server profiles from the template.

Note: If setting up a system with iSCSI boot, the sections with (FCP) in the heading can be skipped and then complete the Cisco UCS iSCSI Configuration section in the Appendix.

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot within the FSV Organization. To create the service profile template, follow these steps:

Note: The server profile captured in this deployment guide supports both Cisco UCS B200 M6 blades and Cisco UCSX X210c M6 compute nodes.

vNIC and vHBA Placement for Server Profile Template

The vNIC and vHBA layout is explained below for FC connected storage. In this deployment, four vHBAs were created to support scsi-fc and fc-nvme storage access.

- Four vNICs and four vHBAs are configured to support FC boot from SAN. Two vHBAs (vHBA-A and vHBA-B) are used for boot from SAN connectivity and the remaining two vHBAs are used to support FC-NVMe. These devices are manually placed as follows:

Table 13. vHBA and vNIC placement for FC with FC-NVMe connected storage

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-A	MLOM	A	0
vHBA-B	MLOM	B	1
vHBA-NVMe-A	MLOM	A	2
vHBA-NVMe-B	MLOM	B	3
00-vSwitch0-A	MLOM	A	4
01-vSwitch0-B	MLOM	B	5
02-VDS0-A	MLOM	A	6
03-VDS0-B	MLOM	B	7

Note: If FC-NVMe connectivity is not required, please use the following vNIC and vHBA layout.

Note: Four vNICs and two vHBAs are configured to support FC boot from SAN. These devices are manually placed as follows:

Table 14. vHBA and vNIC placement for FC connected storage

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-A	MLOM	A	0

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-B	MLOM	B	1
00-vSwitch0-A	MLOM	A	2
01-vSwitch0-B	MLOM	B	3
02-VDS0-A	MLOM	A	4
03-VDS0-B	MLOM	B	5

Server Profile Template Creation

1. Log into Cisco Intersight portal.
2. Go to Configure -> Templates and from the main window, select Create UCS Server Profile Template.

Step 1 – General

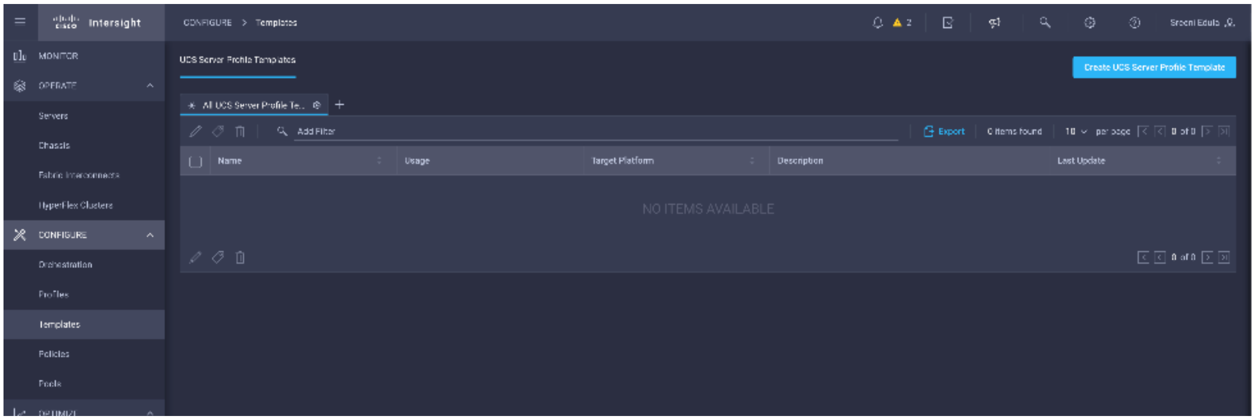

3. Select the Organization from the drop-down list (for example, FSV).
4. Provide a name for the Server Profile Template (for example, VM-Host-Infra-FCP).
5. Select UCS Server (FI-Attached).

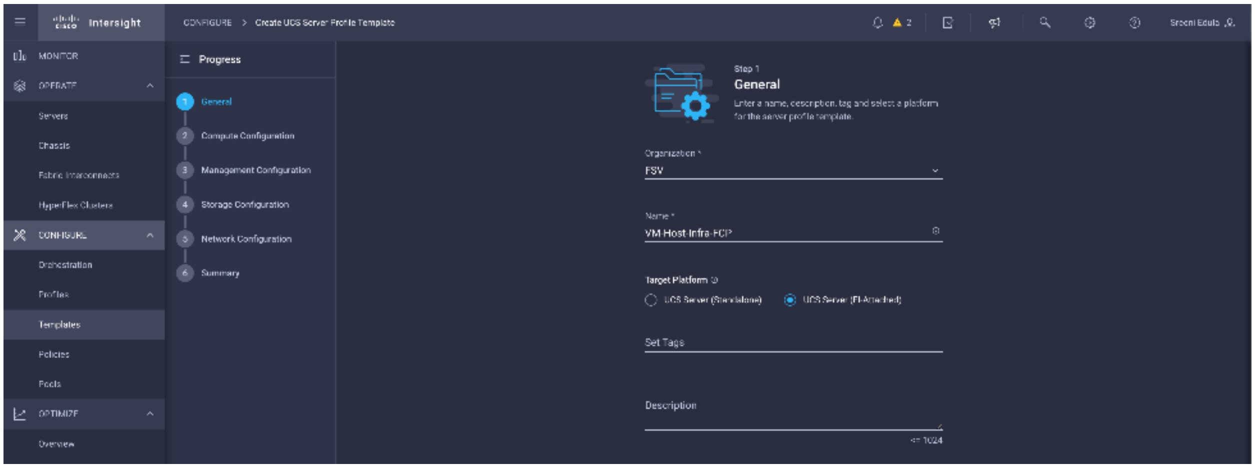

Step 1

General

Enter a name, description, tag and select a platform for the server profile template.

Organization * FSV

Name * VM-Host-Infra-FCP

Target Platform ○ UCS Server (Standalone) ● UCS Server (FI-Attached)

Set Tags

Description

<= 1024

6. Click Next.

Step 2 - Compute Configuration

Configure UUID Pool

1. Click Select Pool next to UUID Pool and in the pane on the right, click Create New.

2. Provide a name for the pool (for example, AA19-UUID-Pool).

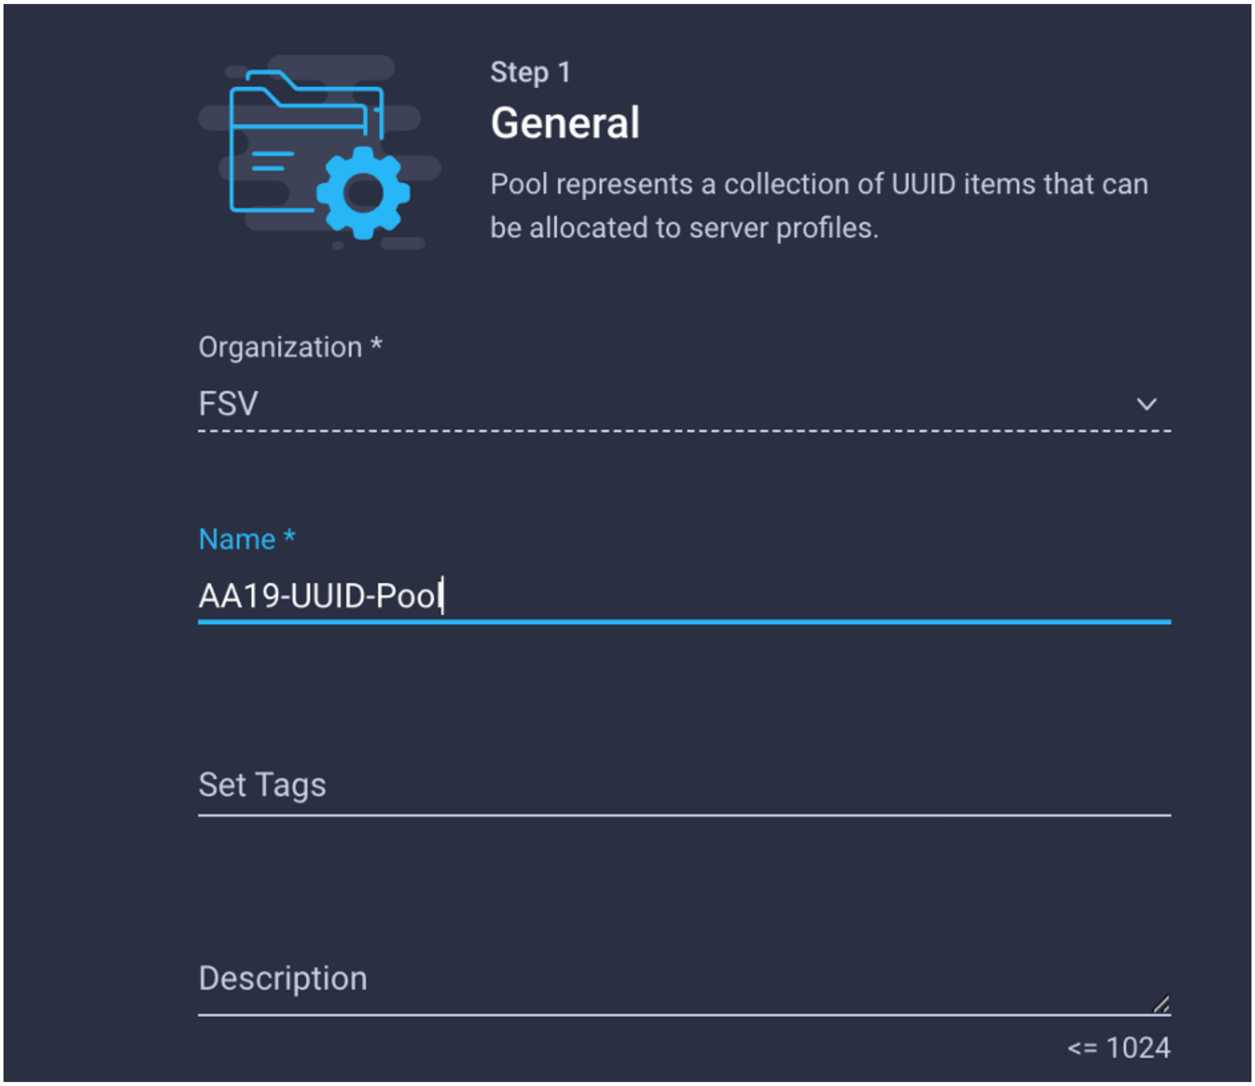

Step 1

General

Pool represents a collection of UUID items that can be allocated to server profiles.

Organization *

FSV

Name *

AA19-UUID-Pool

Set Tags

Description

<= 1024

3. Click Next. The Pool Details page appears.

4. In the Configuration section, add the UUID Prefix number in a hexadecimal format. Example, 1728E89A-7B43-47DE.

5. In the UUID Blocks section, add the following configuration details:

   a. From—Indicates the UUID suffix of the block in a hexadecimal format. Example, 0000-0000A1900001

   b. Size—Indicates the number of UUID identifiers in the block. The size ranges from 1 to 1000.

6. Click Create.

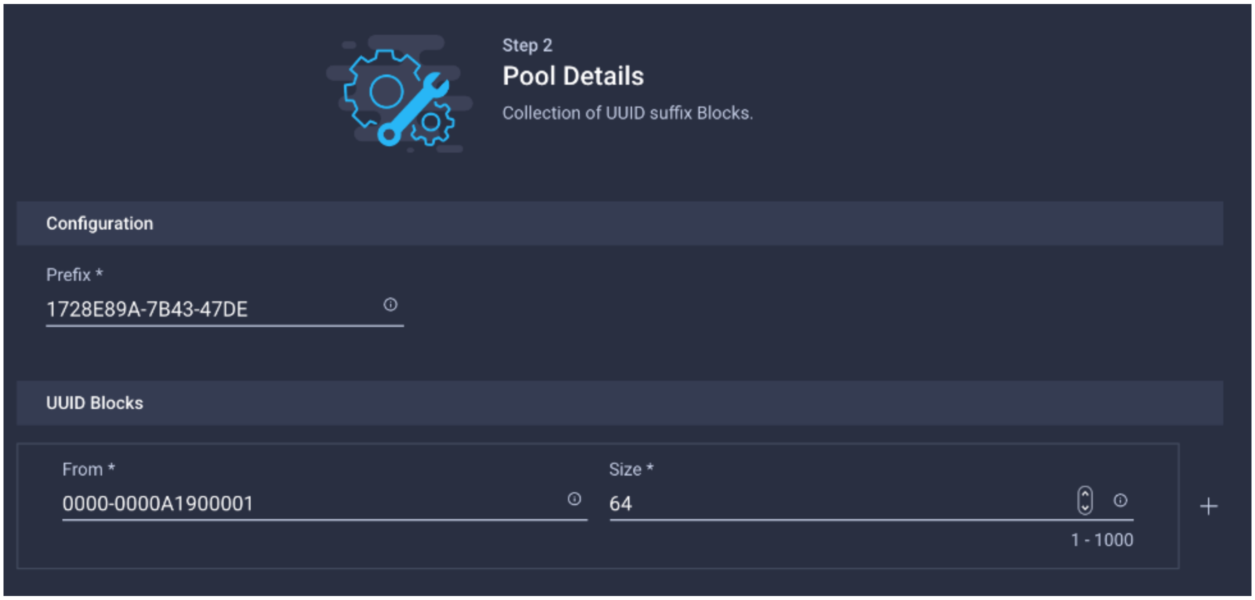

Step 2

Pool Details

Collection of UUID suffix Blocks.

Configuration

Prefix *

1728E89A-7B43-47DE

UUID Blocks

From *

0000-0000A1900001

Size *

64

1 - 1000

Configure BIOS Policy

1. Click Select Policy next to BIOS Configuration and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-BIOS-Pol).

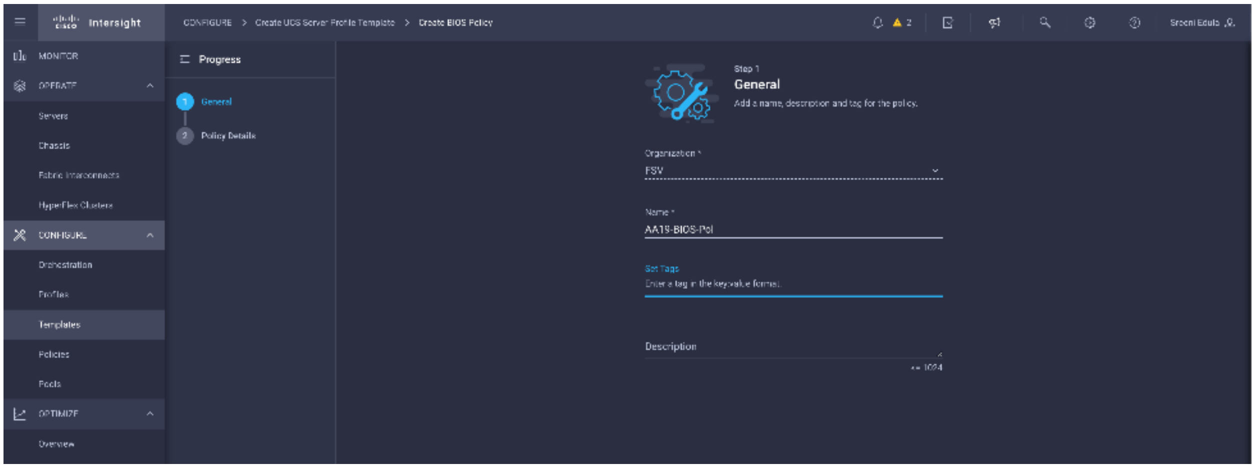

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-BIOS-Pol

Set Tags

Enter a tag in the key:value format.

Description

<= 1024

3. Click Next.

4. On the Policy Details screen, select appropriate values for the BIOS settings. In this deployment, the BIOS values were selected based on recommendations in the performance tuning guide for Cisco UCS M6 BIOS: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.html.

5. Select the appropriate values on the following screen:

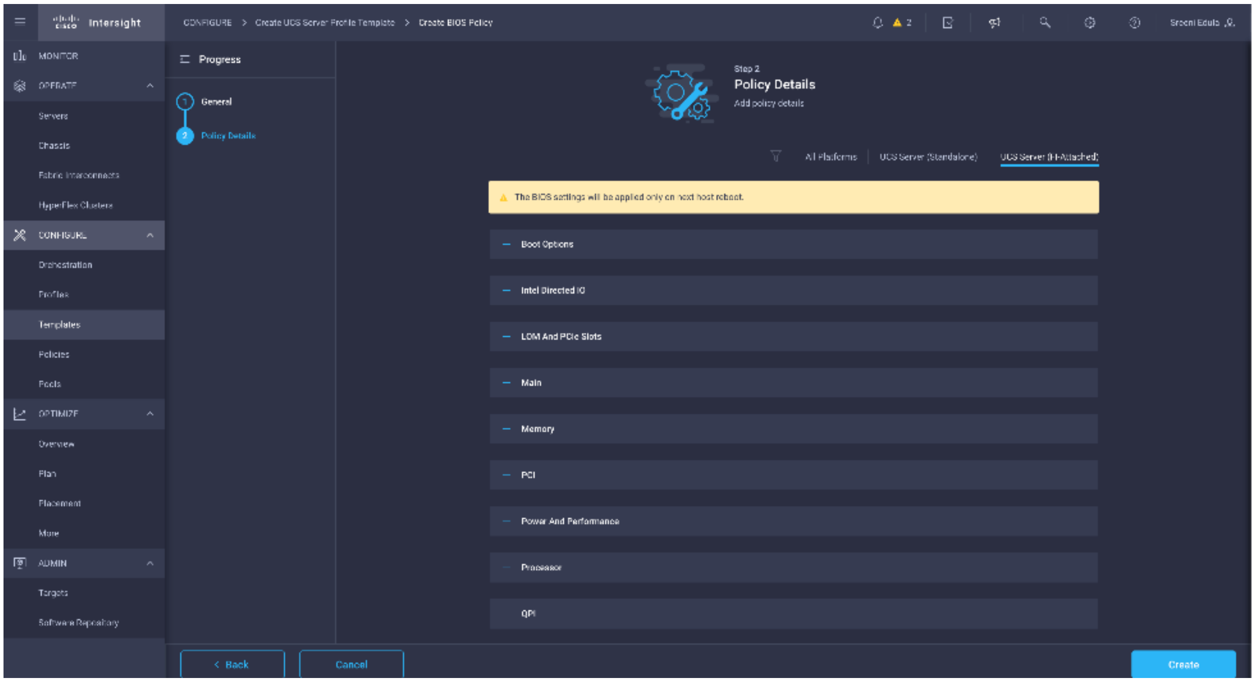

The BIOS settings will be applied only on next host reboot.

- LOM and PCIe Slot -> CDN Support for LOM: **Enabled**
- Processor -> Enhanced CPU performance: **Auto**
- Memory -> NVM Performance Setting: **Balanced Profile**

6. Click Create.

Configure Boot Order Policy

1. Click Select Policy next to Boot Configuration and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-FS-BootOrder-Pol).

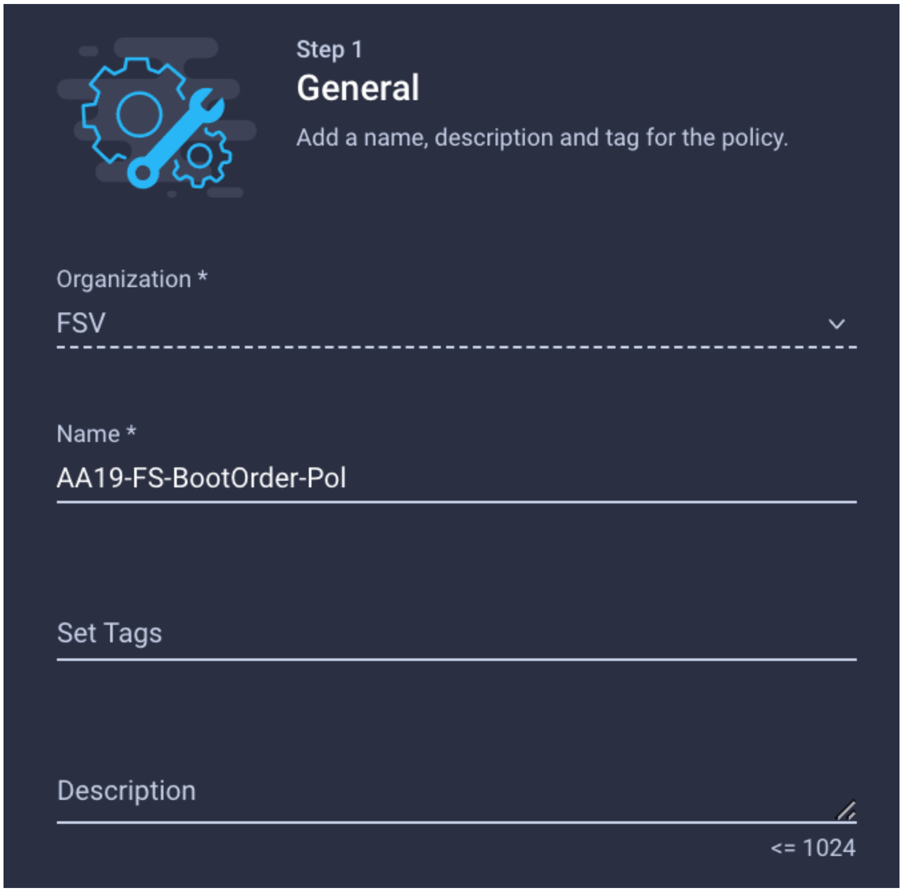

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-FS-BootOrder-Pol

Set Tags

Description

<= 1024

3. Click Next.
4. From Configure Boot Mode, select Unified Extensible Firmware Interface (UEFI).
5. Click Enable Secure Boot.

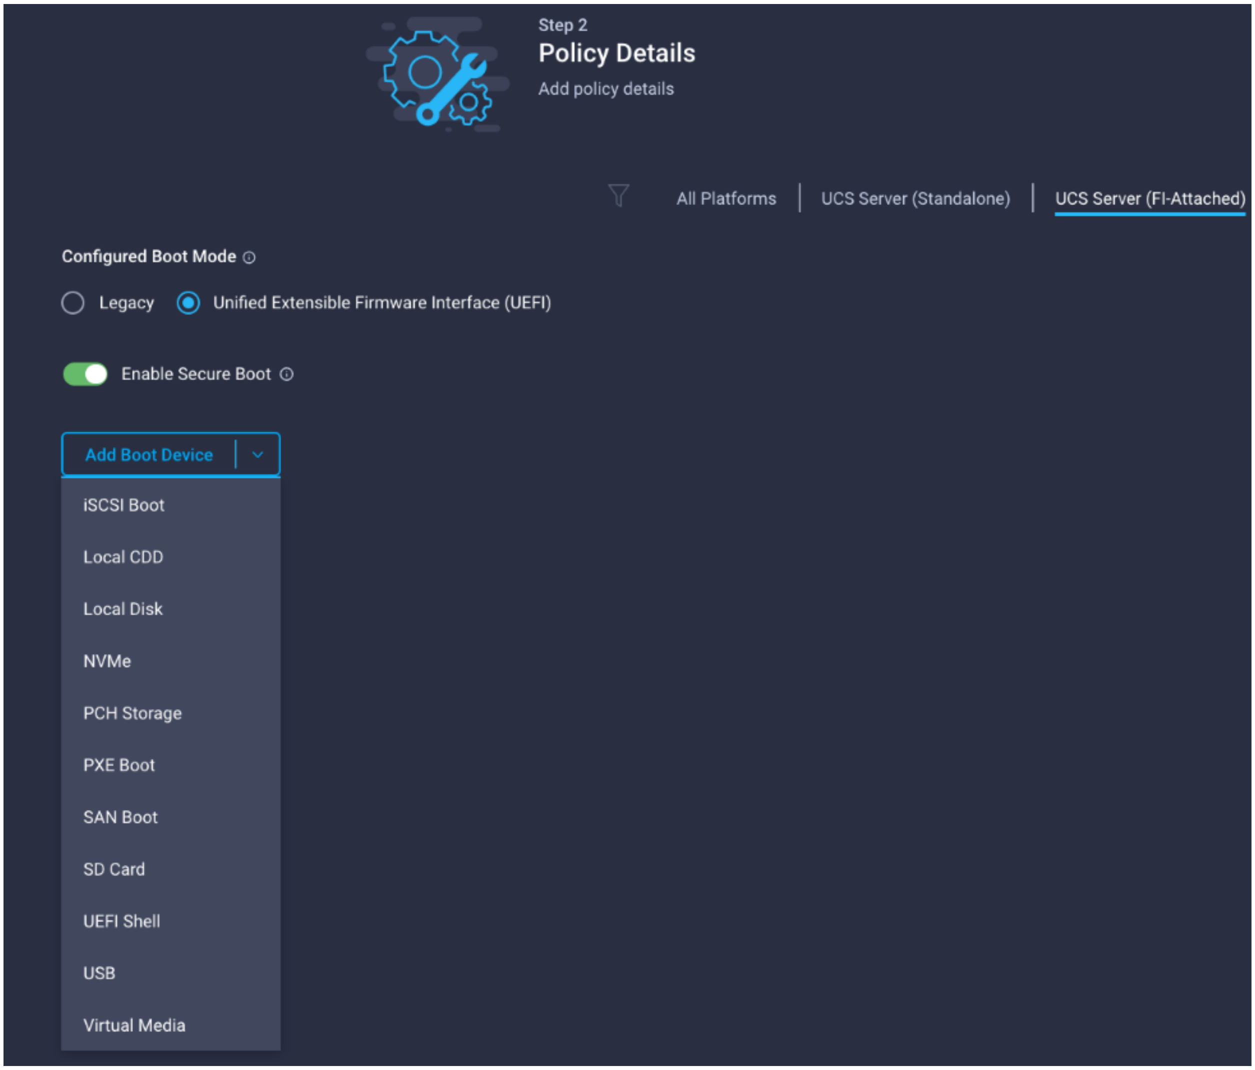

Step 2

Policy Details

Add policy details

Configured Boot Mode

○ Legacy ● Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot: On

6. From the Add Boot Device drop-down list, select Virtual Media.

7. Provide a Device Name (for example, ISO) and Sub-Type KVM Mapped DVD.

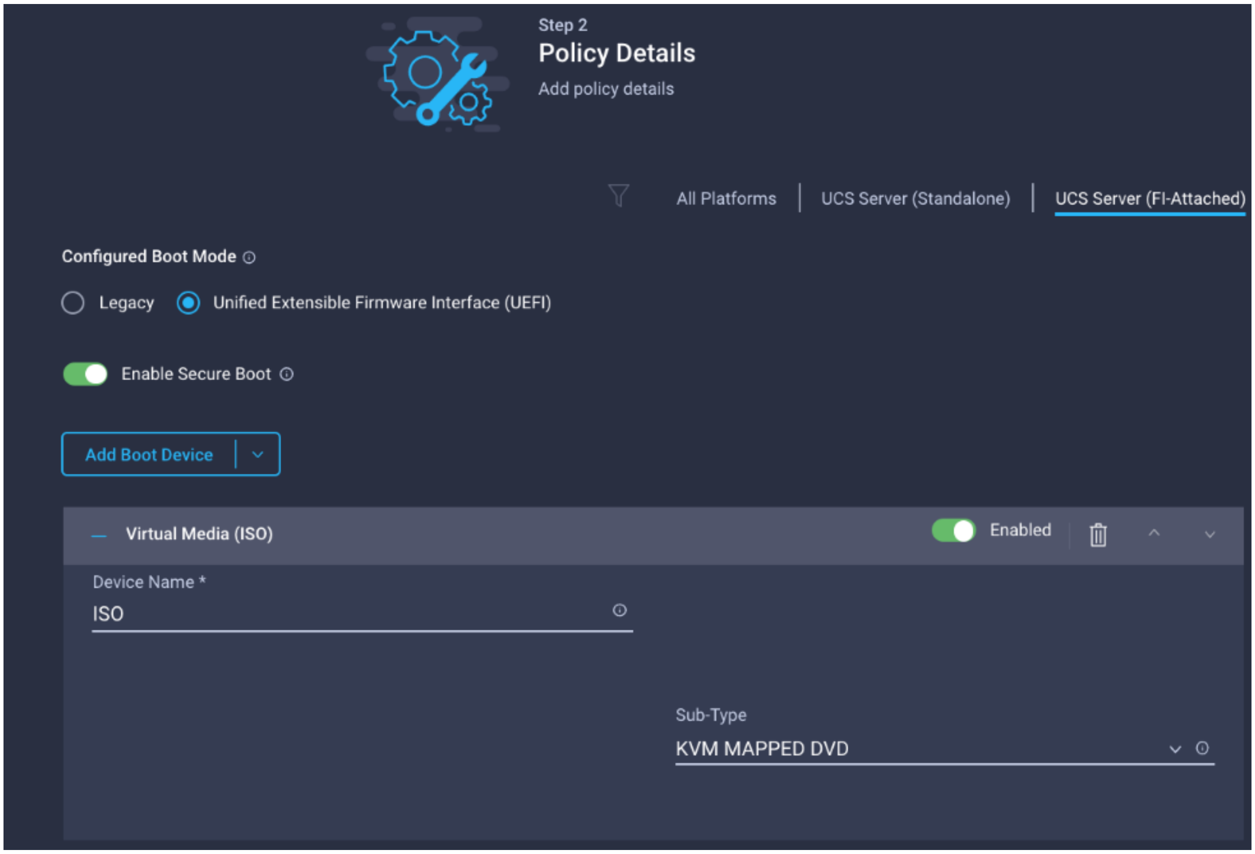

Step 2

Policy Details

Add policy details

Configured Boot Mode

○ Legacy ● Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot

Virtual Media (ISO) Enabled

Device Name * ISO

Sub-Type KVM MAPPED DVD

For Fibre Channel SAN boot, add all four Pure Storage scsi-fc interfaces as boot options. The four interfaces are named as follows:

- **FlashArray-CT0FC0**: FlashArray Controller 0, FC0 (SAN-A)
- **FlashArray-CT1FC0**: FlashArray Controller 1, FC0 (SAN-A)
- **FlashArray-CT0FC2**: FlashArray Controller 0, FC1 (SAN-B)
- **FlashArray-CT1FC2**: FlashArray Controller 1, FC1 (SAN-B)

8. From the Add Boot Device drop-down list, select SAN Boot.

9. Provide the Device Name (for example, FlashArray-CT0FC0) and LUN value (for example, 1).

10. Provide an Interface Name (for example, vHBA-A) and note this name to be used for vHBA definition later. This value is important and should match the vHBA name.

Note: vHBA-A is used to access **CT0FC0** and **CT1FC0** and vHBA-B is used to access **CT0FC2** and **CT1FC2**.

11. Add the appropriate WWPN value in the Target WWPN. This value can be obtained from Pure Storage FlashArray using “pureport list” command using the FlashArray//X CLI or from the Connections tab under the Health section of the FlashArray//X with the Web GUI.

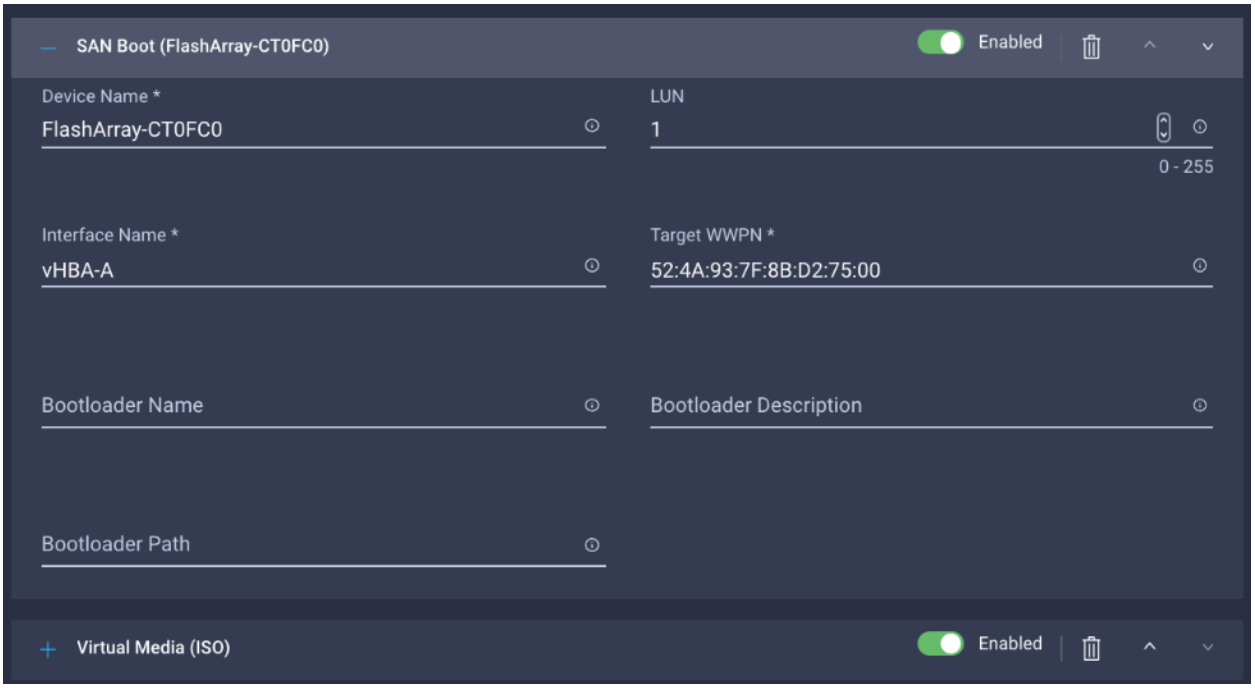

— SAN Boot (FlashArray-CT0FC0) Enabled

Device Name *: FlashArray-CT0FC0

LUN: 1 (0 - 255)

Interface Name *: vHBA-A

Target WWPN *: 52:4A:93:7F:8B:D2:75:00

Bootloader Name

Bootloader Description

Bootloader Path

+ Virtual Media (ISO) Enabled

12. Click SAN Boot again to add the second Pure FlashArray target on the Fabric A side.

13. Repeat steps 1 – 12 three more times to add all the FlashArray Interfaces.

14. Verify the order of the boot policies and adjust the boot order as necessary using the arrows.

15. After adding all the boot devices, the list should look like as shown below:

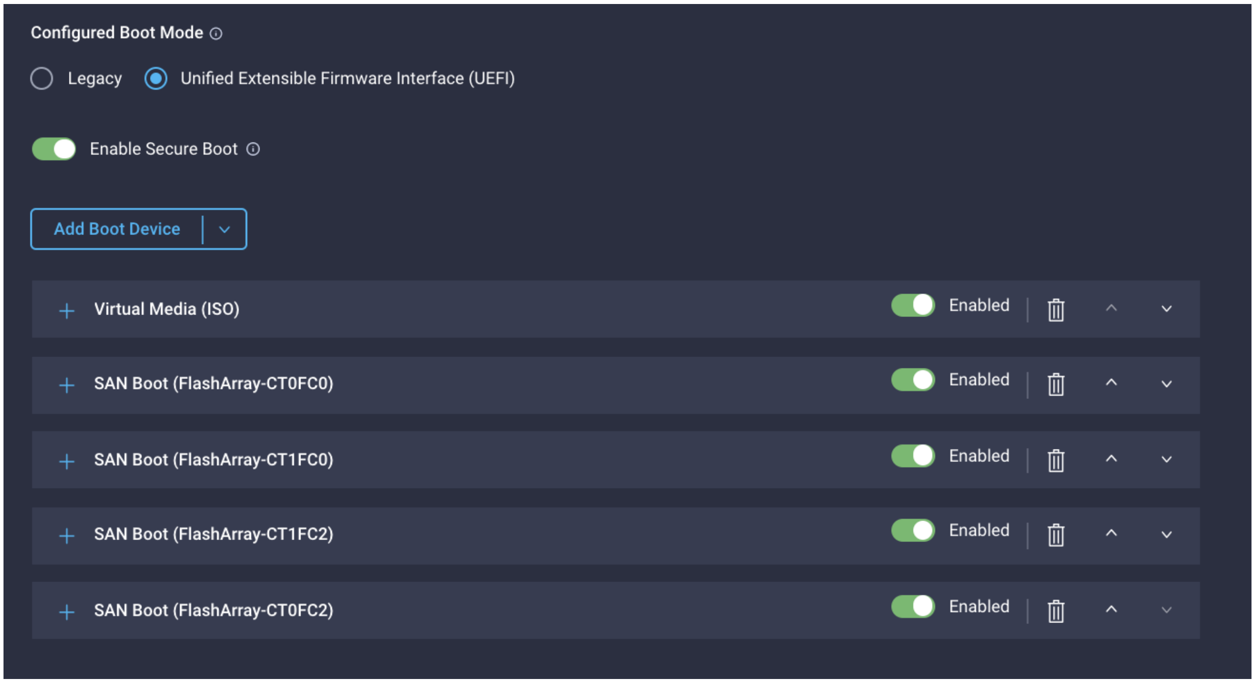

Configured Boot Mode

○ Legacy ● Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot

+ Virtual Media (ISO) Enabled

+ SAN Boot (FlashArray-CT0FC0) Enabled

+ SAN Boot (FlashArray-CT1FC0) Enabled

+ SAN Boot (FlashArray-CT1FC2) Enabled

+ SAN Boot (FlashArray-CT0FC2) Enabled

16. Click Next.

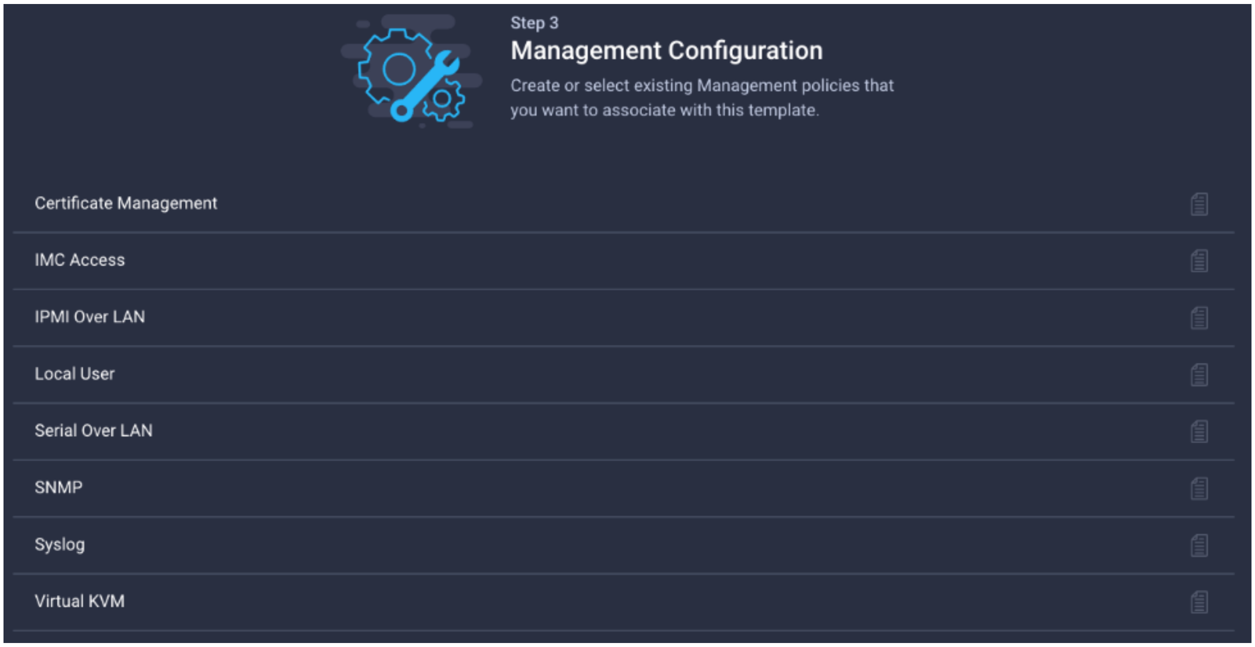

Step 3

Management Configuration

Create or select existing Management policies that you want to associate with this template.

Certificate Management

IMC Access

IPMI Over LAN

Local User

Serial Over LAN

SNMP

Syslog

Virtual KVM

17. Click Create.

18. Click Next to move to Management Configuration.

Step 3 – Management Configuration

Configure the management policy. These policies will be added to the management configuration:

- Certificate Management Policy (Optional) to use external certificates
- IMC Access to define the pool of IP addresses for compute node KVM access
- IPMI Over LAN to allow Intersight to manage IPMI messages
- Local User to provide local administrator to access KVM

Configure Certificate Management Policy (Optional)

The Certificate Management policy allows you to specify the certificate and private key-pair details for an external certificate and attach the policy to servers for IMC access.

1. Click Select Policy next to Certificate Management and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-Cert-Pol).

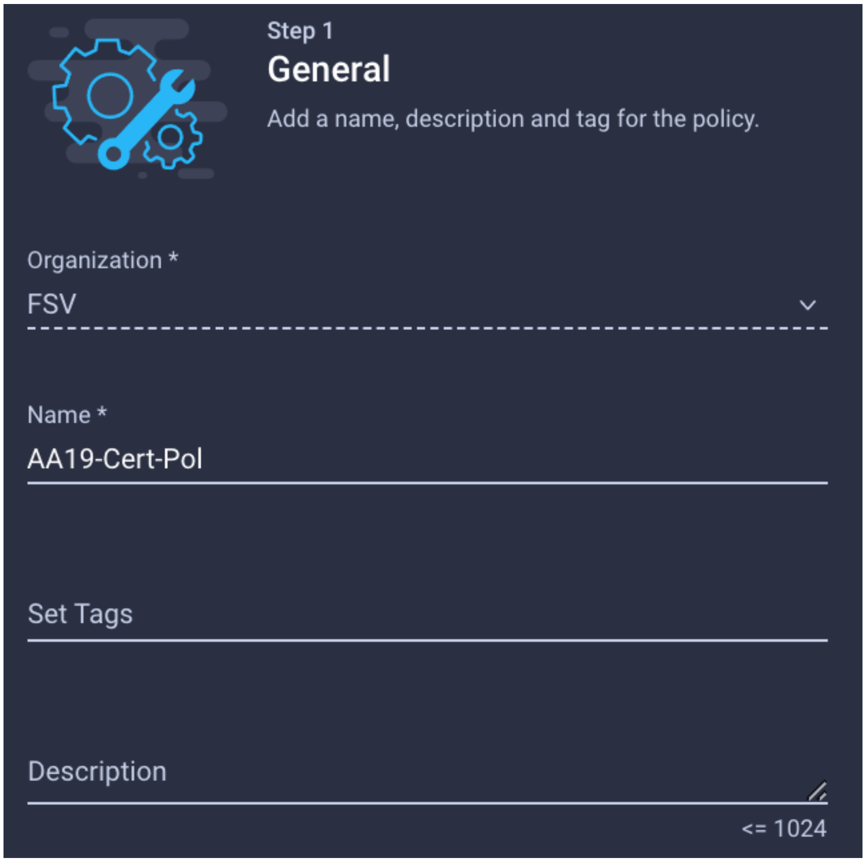
Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-Cert-Pol

Set Tags

Description

<= 1024

3. Click Next.
4. Enter the certificate details and Private Key.
5. Click add.

Configure IMC Access Policy

1. Click Select Policy next to IMC Access and in the pane on the right, click Create New.
2. Provide a name for the policy (for example, AA19-IMC-Access).

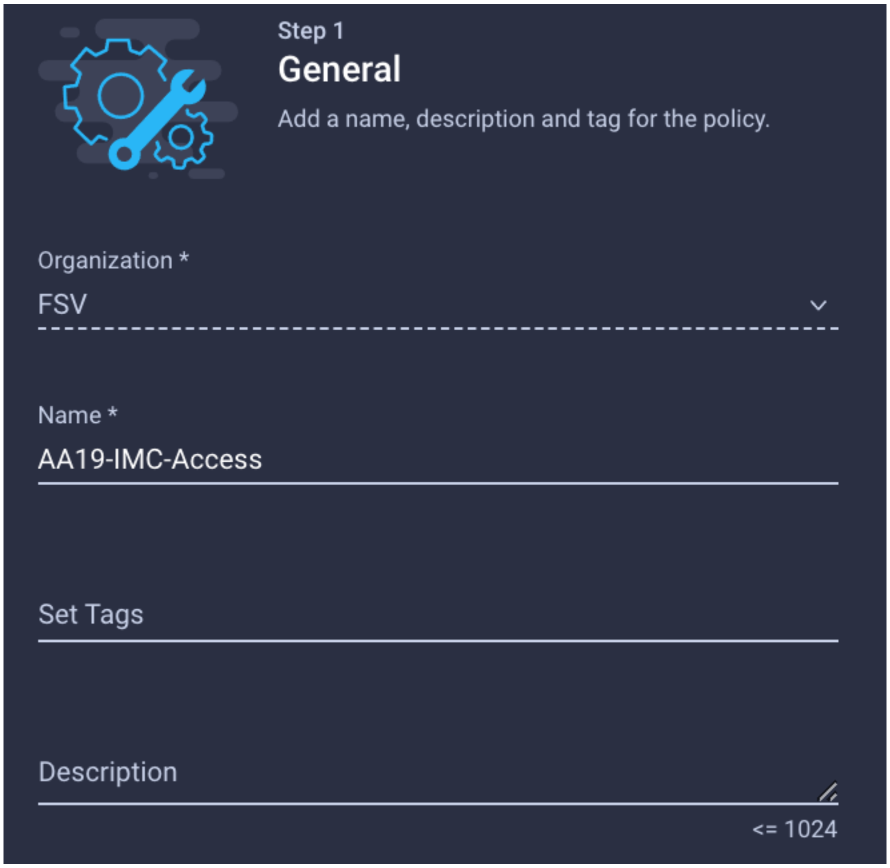

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-IMC-Access

Set Tags

Description

<= 1024

3. Click Next.

4. Provide the in-band (or out-of-band) management VLAN ID (for example, 115)

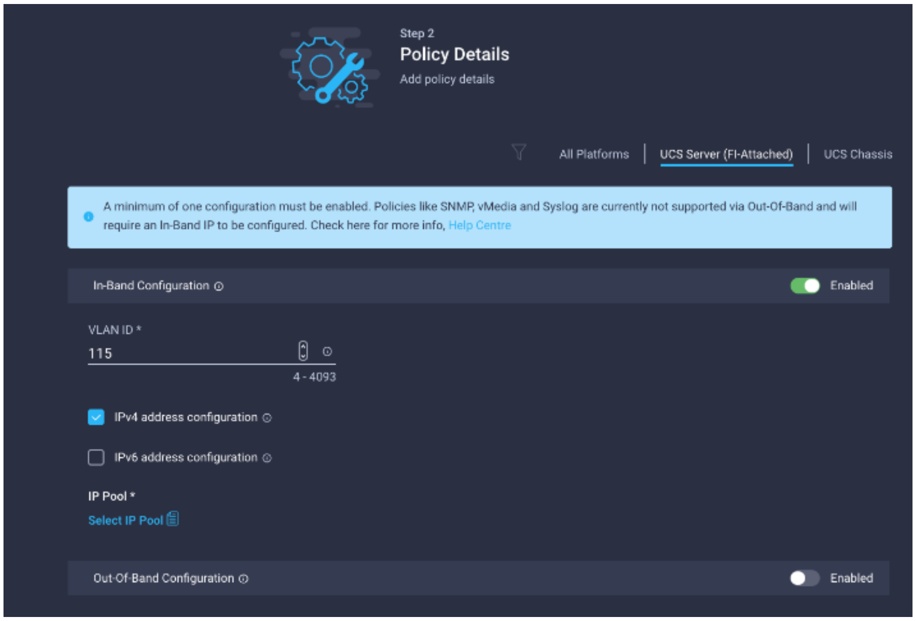

Step 2

Policy Details

Add policy details

A minimum of one configuration must be enabled. Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Check here for more info, Help Centre

In-Band Configuration — Enabled

VLAN ID *

115

4 - 4093

☑ IPv4 address configuration

□ IPv6 address configuration

IP Pool *

Select IP Pool

Out-Of-Band Configuration — Enabled

5. Select Configure IPv4 address configuration and click Select IP Pool for defining a KVM IP address assignment pool.

6. Click Create New in the menu on the right.

7. Provide a name for the policy (for example, AA19-IMC-Pool).

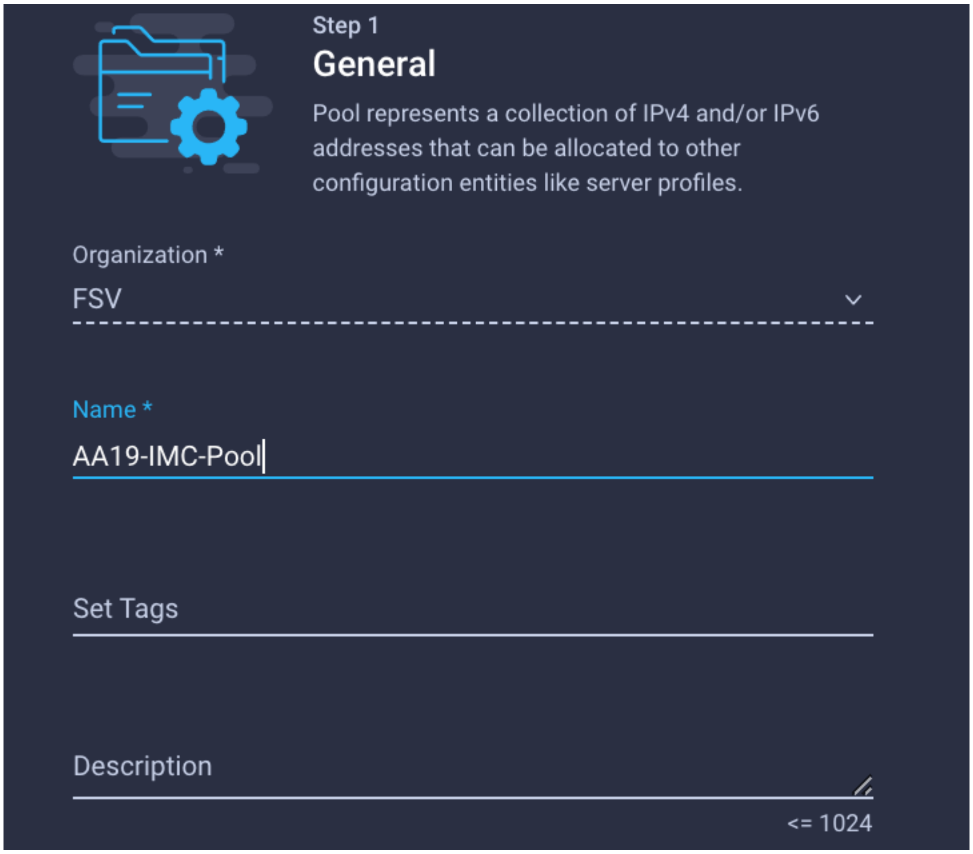

Step 1

General

Pool represents a collection of IPv4 and/or IPv6 addresses that can be allocated to other configuration entities like server profiles.

Organization *

FSV

Name *

AA19-IMC-Pool

Set Tags

Description

<= 1024

8. Select Configure IPv4 Pool and provide the information to define a pool for KVM IP address assignment.

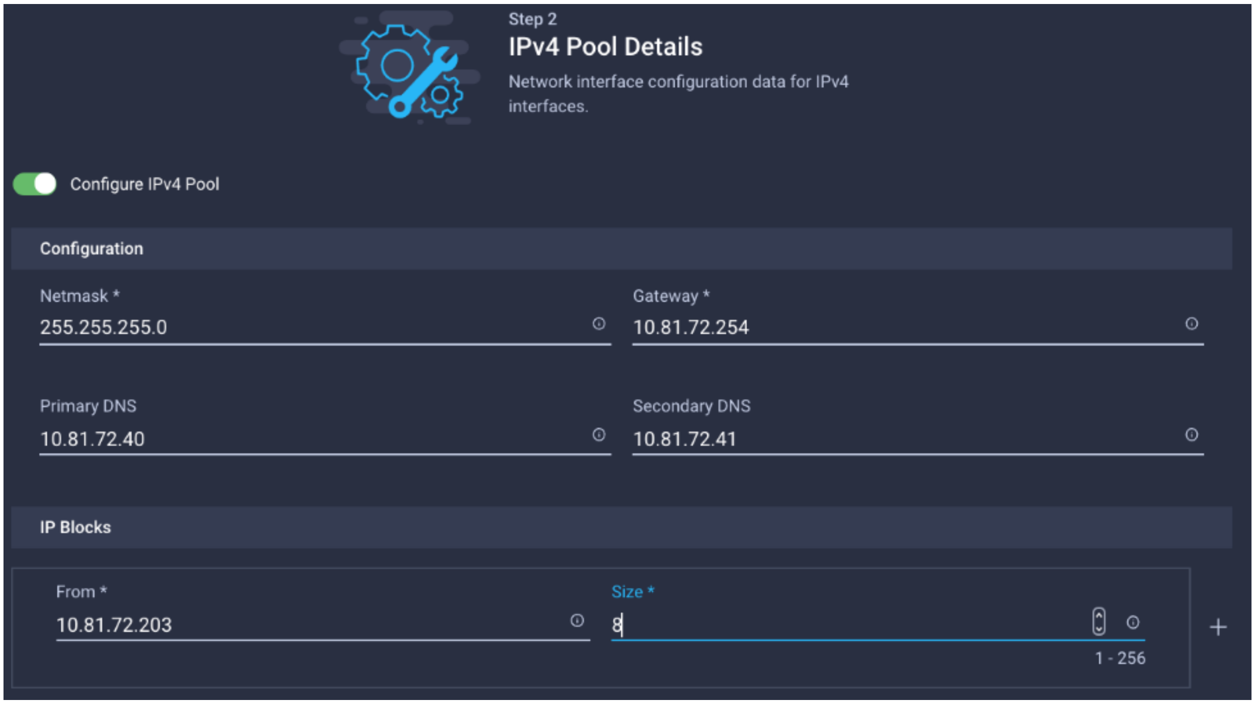

Step 2

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Configuration

Netmask * 255.255.255.0

Gateway * 10.81.72.254

Primary DNS 10.81.72.40

Secondary DNS 10.81.72.41

IP Blocks

From * 10.81.72.203

Size * 8

1 - 256

Note: The management IP pool subnet should be accessible from the host trying to open the KVM connection. In the example above, the hosts trying to open a KVM connection would need to be able to route to 192.168.160.0 subnet.

9. Click Next.

10. Unselect Configure IPv6 Pool.

11. Click Create to finish configuring the IP pool.

12. Click Create to finish configuring the IMC Access Policy.

Configure IPMI Over LAN Policy

To configure IPMI Over LAN policy, follow these steps:

1. Click Select Policy next to IPMI Over LAN and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, AA19) and provide a name for the policy (for example, Enable-IPMIoLAN).

3. Turn on Enable IPMI Over LAN.

4. From the Privilege Level drop-down list, select admin.

5. Click Create.

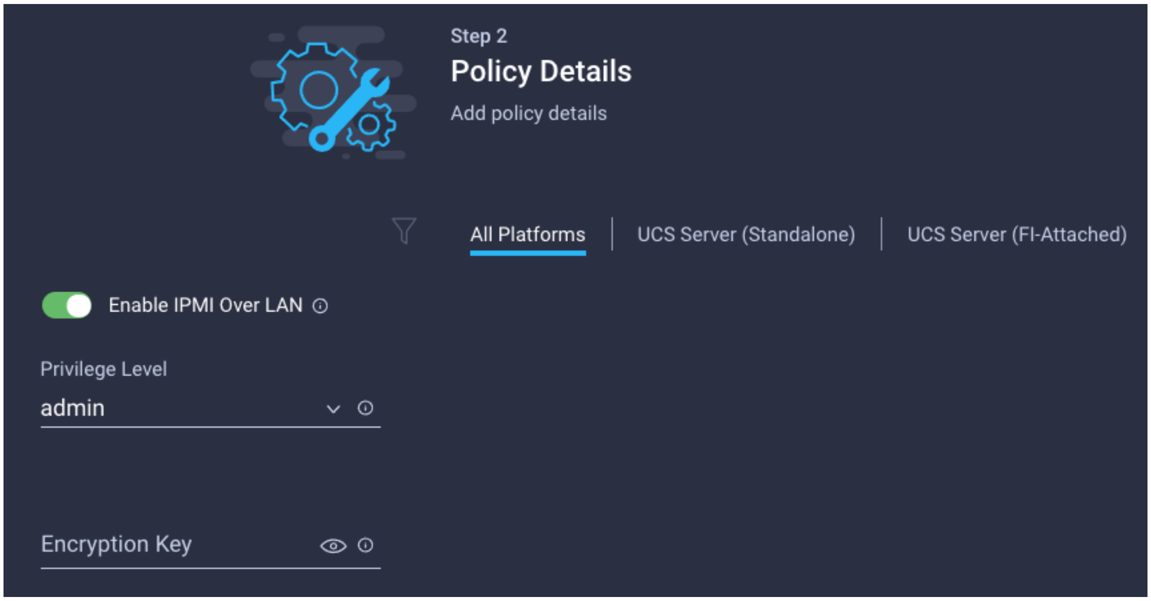

Step 2

Policy Details

Add policy details

Enable IPMI Over LAN

Privilege Level: admin

Encryption Key

Configure Local User Policy

To configure local user policy, follow these steps:

1. Click Select Policy next to Local User and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-LocalUser-Pol).

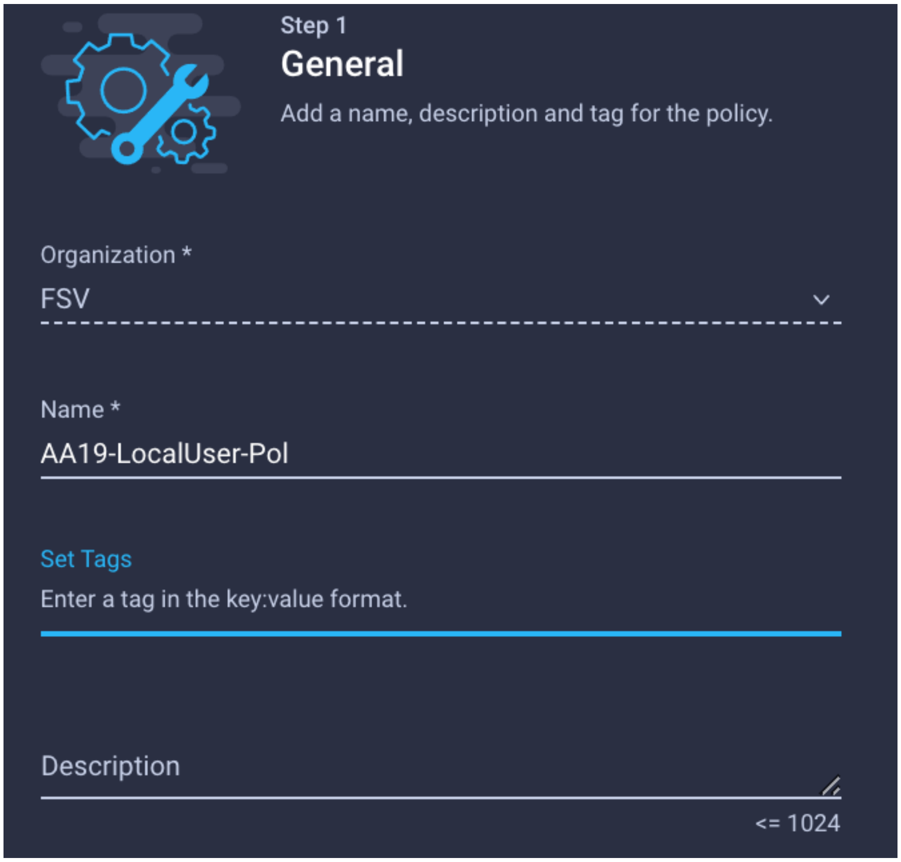

Step 1

General

Add a name, description and tag for the policy.

Organization *
FSV

Name *
AA19-LocalUser-Pol

Set Tags
Enter a tag in the key:value format.

Description

<= 1024

3. Verify UCS Server (FI-Attached) is selected.

4. Verify Enforce Strong Password is selected.

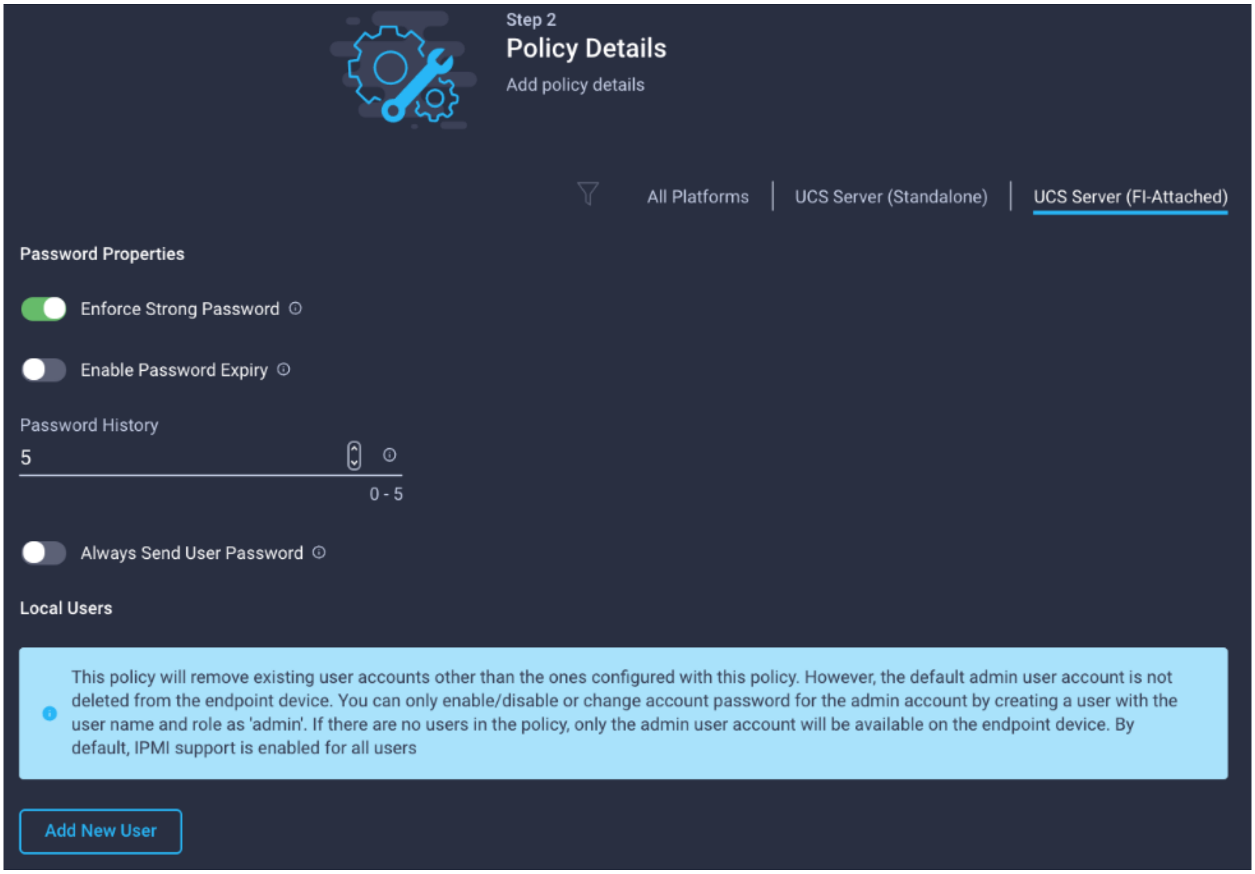
Step 2

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

Password Properties

Enforce Strong Password: on

Enable Password Expiry: off

Password History: 5 (0 - 5)

Always Send User Password: off

Local Users

This policy will remove existing user accounts other than the ones configured with this policy. However, the default admin user account is not deleted from the endpoint device. You can only enable/disable or change account password for the admin account by creating a user with the user name and role as 'admin'. If there are no users in the policy, only the admin user account will be available on the endpoint device. By default, IPMI support is enabled for all users

5. Click Add New User.

6. Provide the Username (for example, flashadmin), Role (for example, admin) and the password.

Note: The username and password combination defined here will be used to log into KVMs. The typical UCS admin/password combination cannot be used for KVM access.

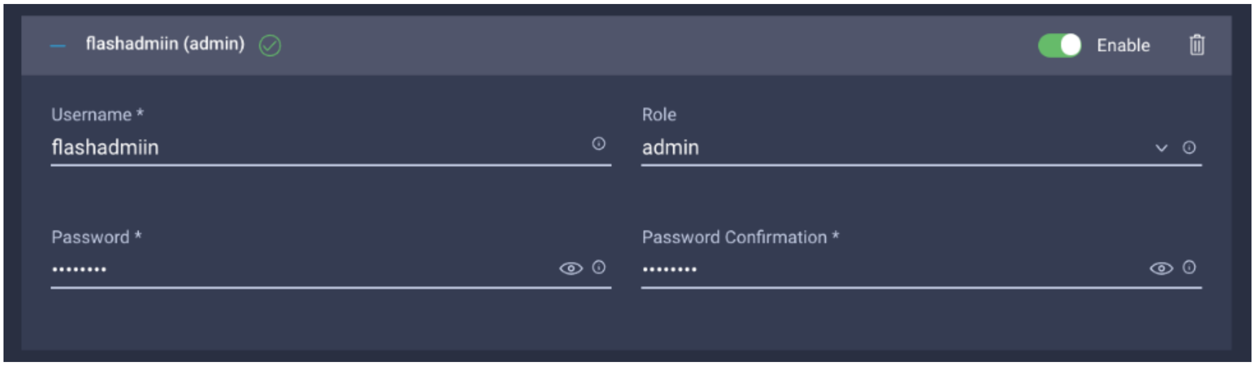
flashadmiin (admin)

Enable: on

Username *: flashadmiin

Role: admin

Password *: ••••••••

Password Confirmation *: ••••••••

7. Click Create to finish configuring the user.

8. Click Create to finish configuring Local User Policy.

9. Click Next.

Step 4 – Storage Configuration

Click Next on the Storage Configuration screen. No configuration is needed in the local storage system.

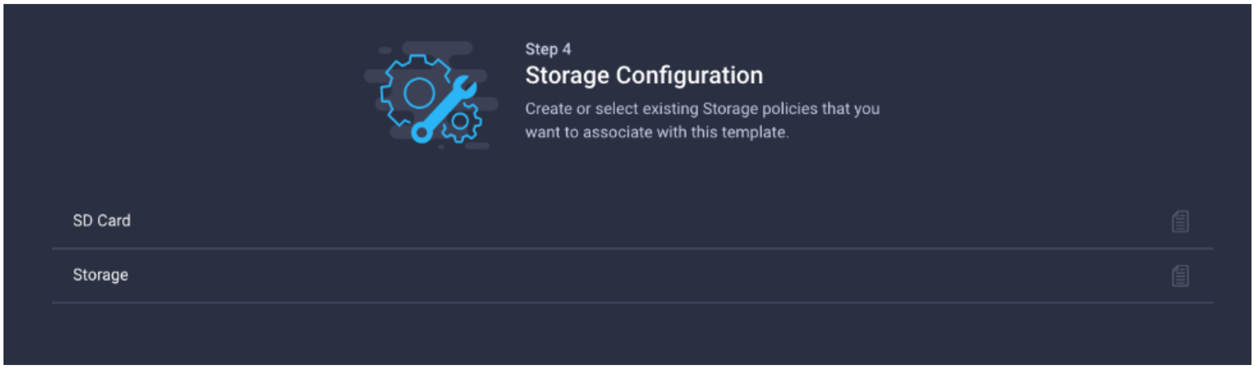

Step 4

Storage Configuration

Create or select existing Storage policies that you want to associate with this template.

SD Card

Storage

Step 5a – Network Configuration – LAN Connectivity Policy

LAN connectivity policy defines the connections and network communication resources between the server and the LAN. This policy uses pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network.

For consistent vNIC and vHBA placement, manual vHBA/vNIC placement is utilized. FC host LAN connectivity policy is explained separately in this section.

Table 15. vNICs for FC LAN Connectivity

vNIC/vHBA Name	Slot ID	Switch ID	PCI Order	VLANs
00-vSwitch0-A	MLOM	A	4	IB-MGMT
01-vSwitch0-B	MLOM	B	5	IB-MGMT
02-VDS0-A	MLOM	A	6	VM Traffic, vMotion
03-VDS0-B	MLOM	B	7	VM Traffic, vMotion

Note: If fc-nvme access is not required, the PCI Order needs to be adjusted accordingly as shown in Table 14.

1. Click Select Policy next to LAN Connectivity and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-FC-ESXi-LANConn-Manual).

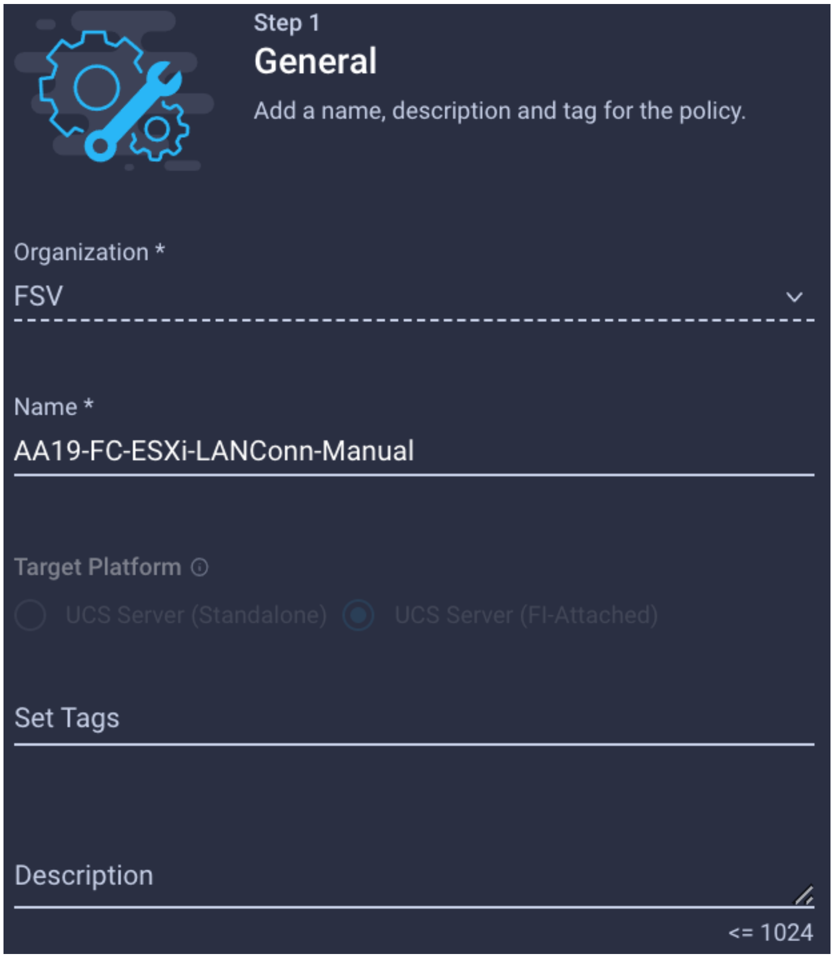

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-FC-ESXi-LANConn-Manual

Target Platform

○ UCS Server (Standalone) ● UCS Server (FI-Attached)

Set Tags

Description

<= 1024

3. To manually select the vNIC placement, select Manual vNIC Placement.

4. Click Add vNIC.

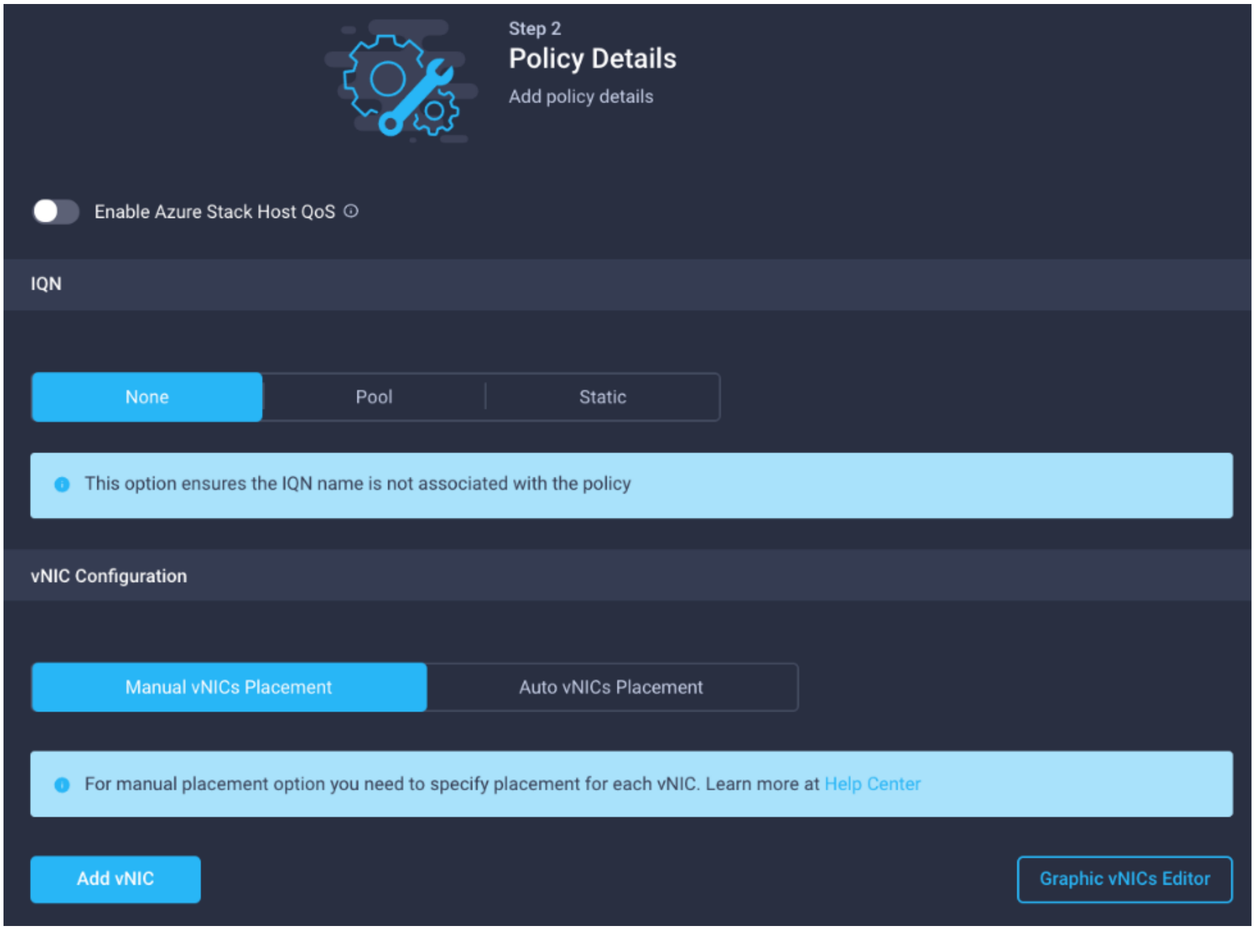

This option ensures the IQN name is not associated with the policy

For manual placement option you need to specify placement for each vNIC. Learn more at Help Center

5. Configure the following parameters:

   a. Provide the name of vNIC (for example, 00-vSwitch0-A).

   b. Enter PCIe slod ID as MLOM where the VIC adapter is installed.

   c. Enter the PCI Link as 0

   d. Select the Switch ID as A, the fabric port to which the vNICs are associated.

   e. Enter PCI Order a 4, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter

   f. From the drop-down list, select vNIC Name for Consistent Device Naming (CDN).

   g. Leave the Failover disabled. The failover will be supported by attaching multiple NICs to VMware vSwitch and VDS.

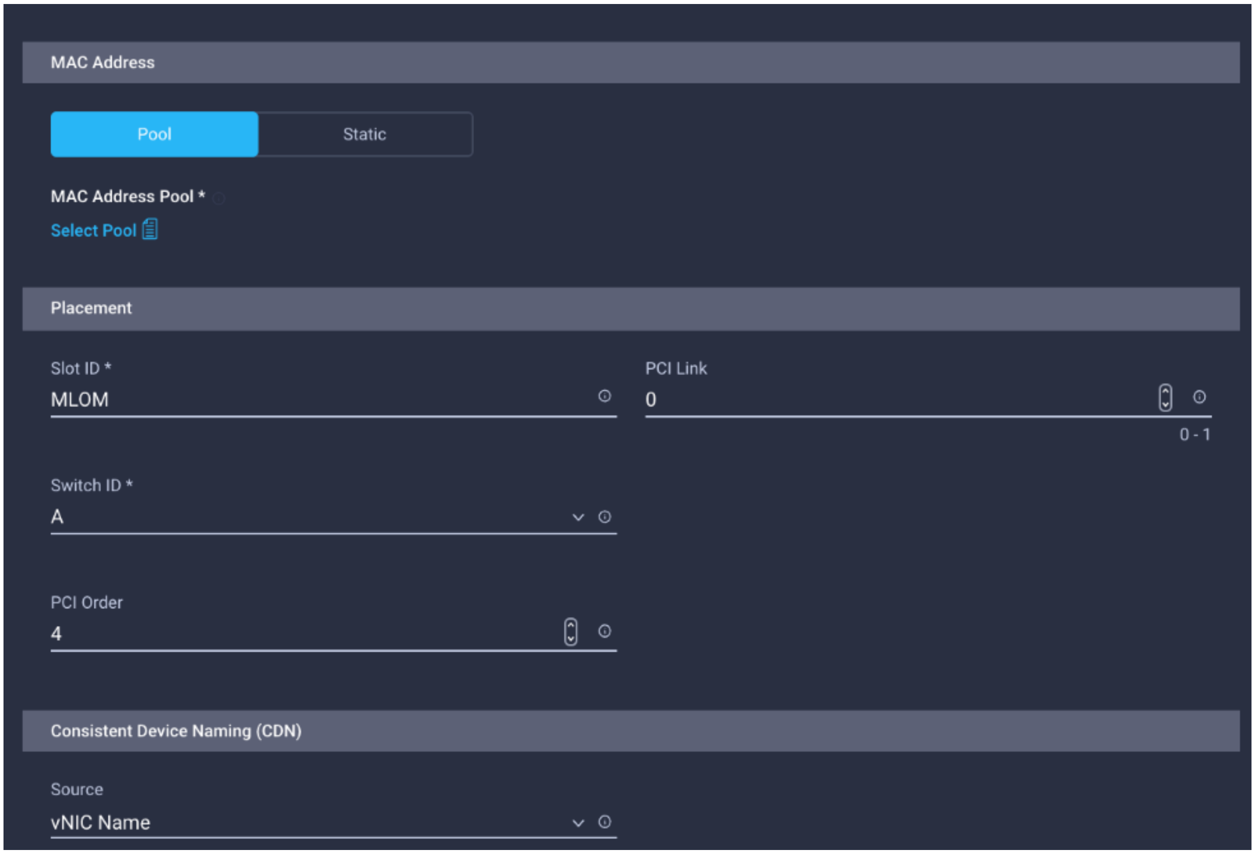

MAC Address

MAC Address Pool *

Placement

Slot ID *: MLOM

PCI Link: 0 (0 - 1)

Switch ID *: A

PCI Order: 4

Consistent Device Naming (CDN)

Source: vNIC Name

Since the MAC Address Pool has not been defined yet, we will create a new MAC address Pool for Fabric-A. This pool will be re-used for the all Fabric-A vNICs.

Create MAC Address Pool

When creating the first vNIC, the MAC address pool has not been defined yet therefore a new MAC address pool will need to be created. Two separate MAC address pools are configured for each Fabric. MAC-Pool-A will be reused for all Fabric-A vNICs, and MAC-Pool-B will be reused for all Fabric-B vNICs.

Table 16. MAC Address Pools

Pool Name	Starting MAC Address	Size	vNICs	Pool Name
MAC-Pool-A	00:25:B5:19:0A:00	128*	01-vSwitch0-A, 03-VDS0-A	MAC-Pool-A
MAC-Pool-B	00:25:B5:19:0B:00	128*	02-vSwitch0-B, 04-VDS0-B	MAC-Pool-B

Note: Each server requires 2 MAC addresses from the pool. Adjust the size of the pool according to your requirements.

1. Click Select Pool under MAC Address Pool and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-MAC-Pool-A).

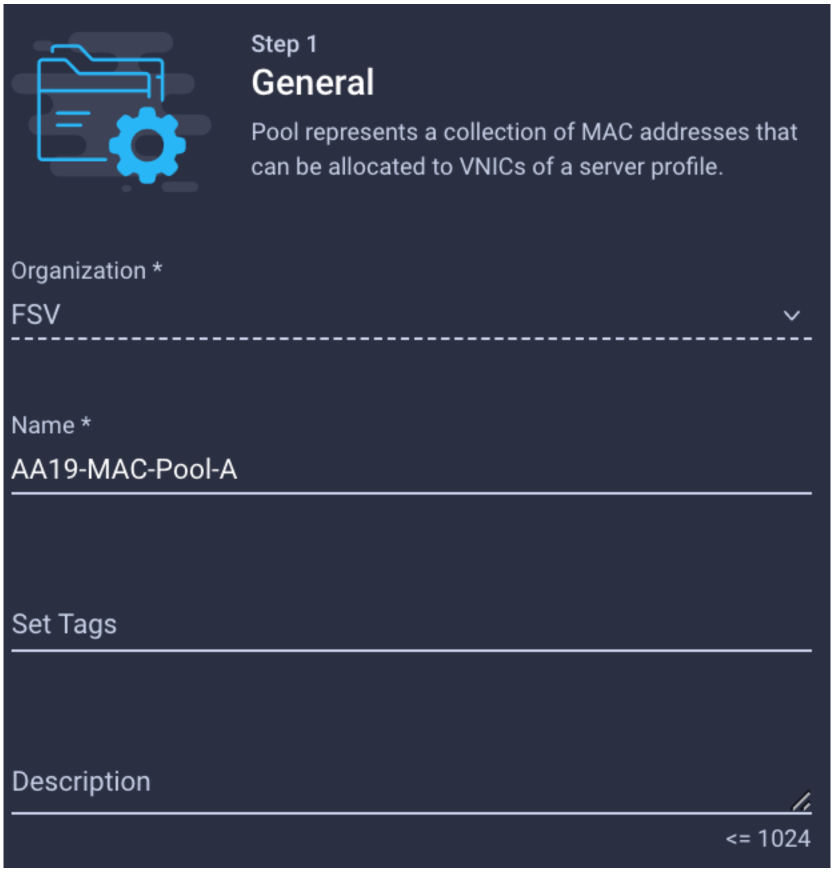

Step 1

General

Pool represents a collection of MAC addresses that can be allocated to VNICs of a server profile.

Organization *

FSV

Name *

AA19-MAC-Pool-A

Set Tags

Description

<= 1024

3. Click Next.

4. Provide the start MAC address. Recommended prefix for MAC addresses is 00:25:B5:xx:xx:xx. As a best practice, some extra information is always coded into the MAC address pool for ease of troubleshooting. For example, in the figure below, 19 represents the Rack ID and 0A represents Fabric A.

5. Provide the size of the MAC Pool (for example, 128).

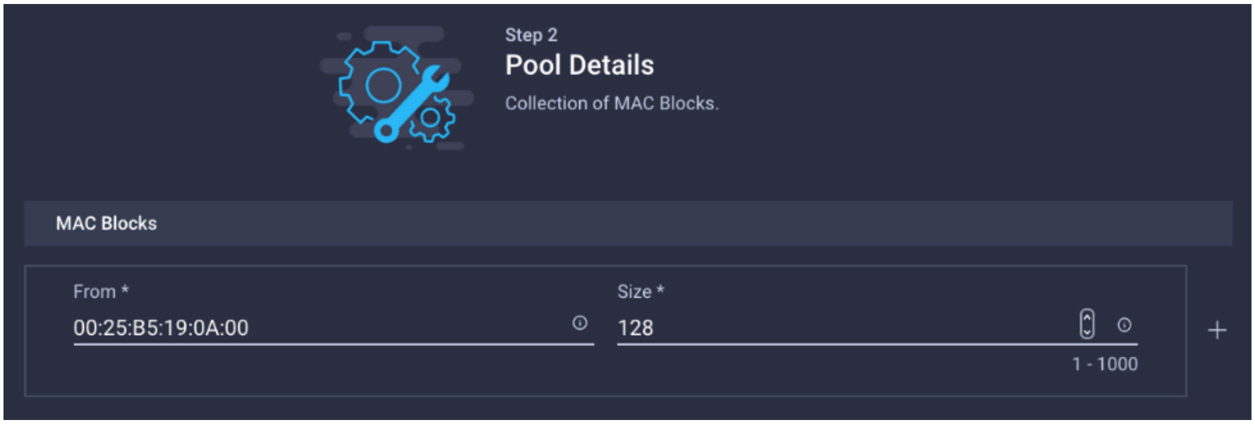

Step 2

Pool Details

Collection of MAC Blocks.

MAC Blocks

From *

00:25:B5:19:0A:00

Size *

128

1 - 1000

6. Click Create to finish creating the MAC Address Pool.

7. Back on the Add vNIC window, from the drop-down list, select “A” as the Switch ID.

Since the Ethernet policies have not been created yet, these policies will be created at this time. These policies will be re-used when defining additional vNICs.

Ethernet Network Group Policy

Ethernet Network Group policies will be created and reused on applicable vNICs as covered below. Ethernet network group policy defines the VLANs allowed for a particular vNIC therefore multiple network group policies will be defined for this deployment as follows:

Table 17. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs	VLANs	Group Policy Name
AA19-vSwitch0-NetGrp	Native-VLAN (2)	00-vSwitch0-A, 01-vSwitch0-B	IB-MGMT, OOB-Mgmt	AA19-vSwitch0-NetGrp
AA19-VDS0-NetGrp	Native-VLAN (2)	02-VDS0-A, 03-VDS0-B	VM Traffic, vMotion	AA19-VDS0-NetGrp

To define Ethernet Group Policy for a vNIC, follow these steps:

1. Click Select Policy under Ethernet Network Group Policy and in the pane on the right, click Create New.

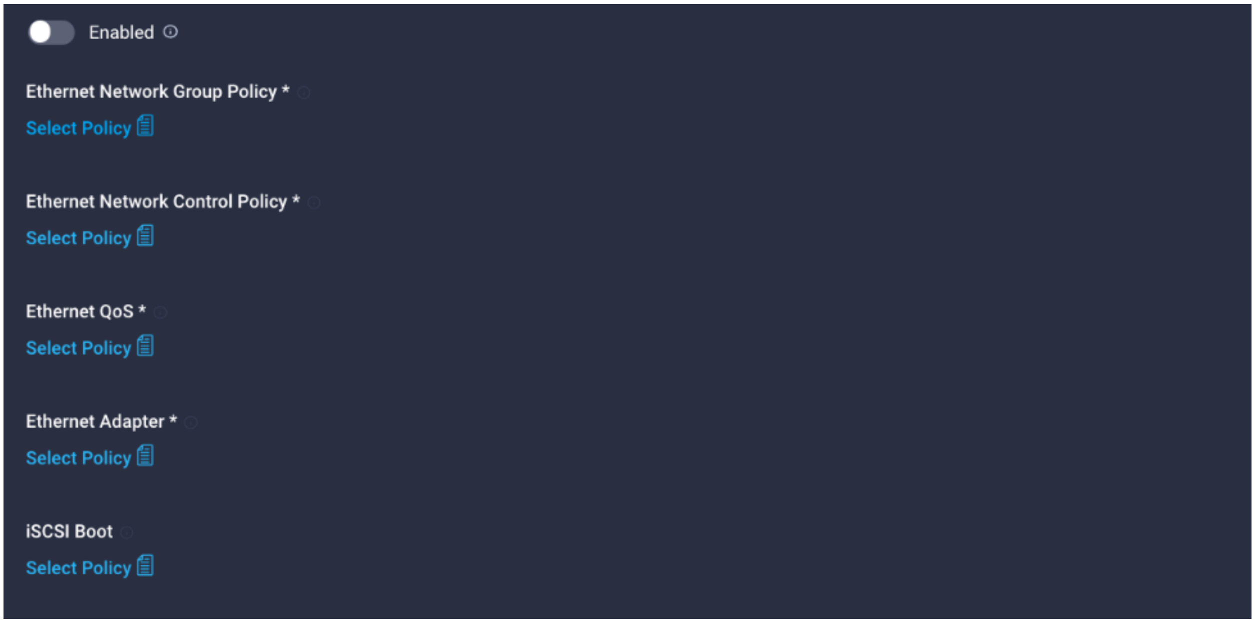

2. Provide a name for the policy (for example, AA19-vSwitch0-NetGrp).

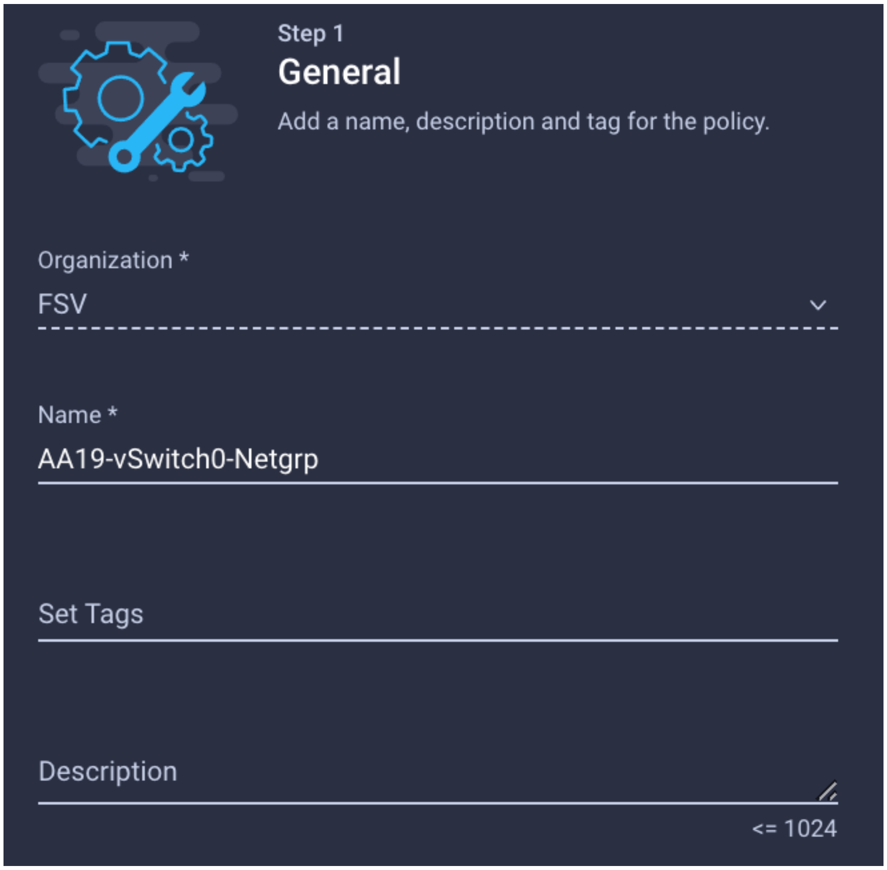

Step 1

General

Add a name, description and tag for the policy.

Organization * FSV

Name * AA19-vSwitch0-Netgrp

Set Tags

Description

<= 1024

3. Click Next.
4. Enter Allowed VLANs from Table 2 (for example, 15,115) and Native VLAN id (for example, 2).

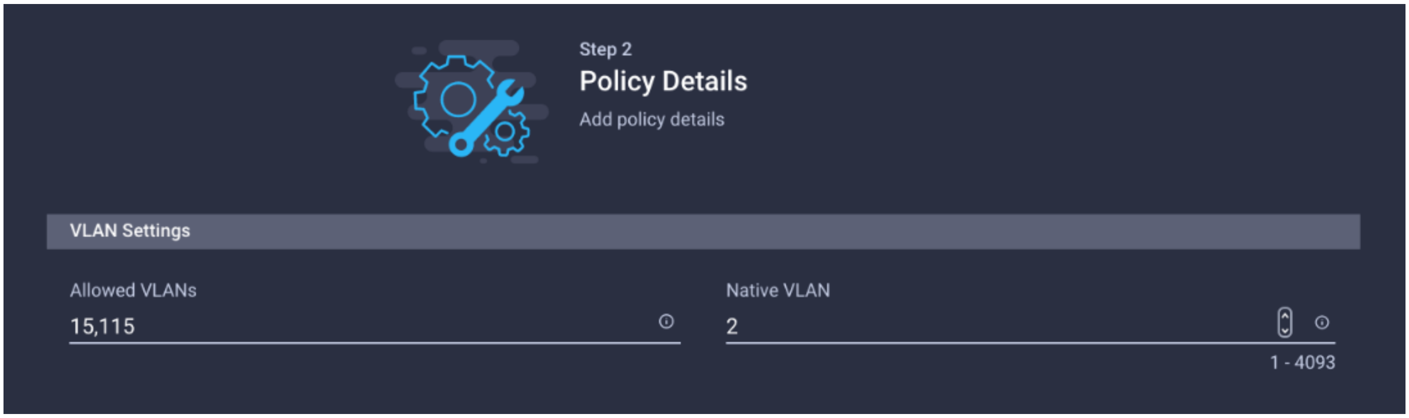

Step 2

Policy Details

Add policy details

VLAN Settings

Allowed VLANs 15,115

Native VLAN 2

1 - 4093

5. Click Create to finish configuring the “Ethernet Network Group Policy.”

Ethernet Network Control Policy

Ethernet Network Control Policy is used to enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) for the vNICs. A single policy will be created here and reused for all the vNICs.

To create the Ethernet Network Control Policy, follow these steps:

1. Click Select Policy under Ethernet Network Control Policy and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-Enable-CDP-LLDP).

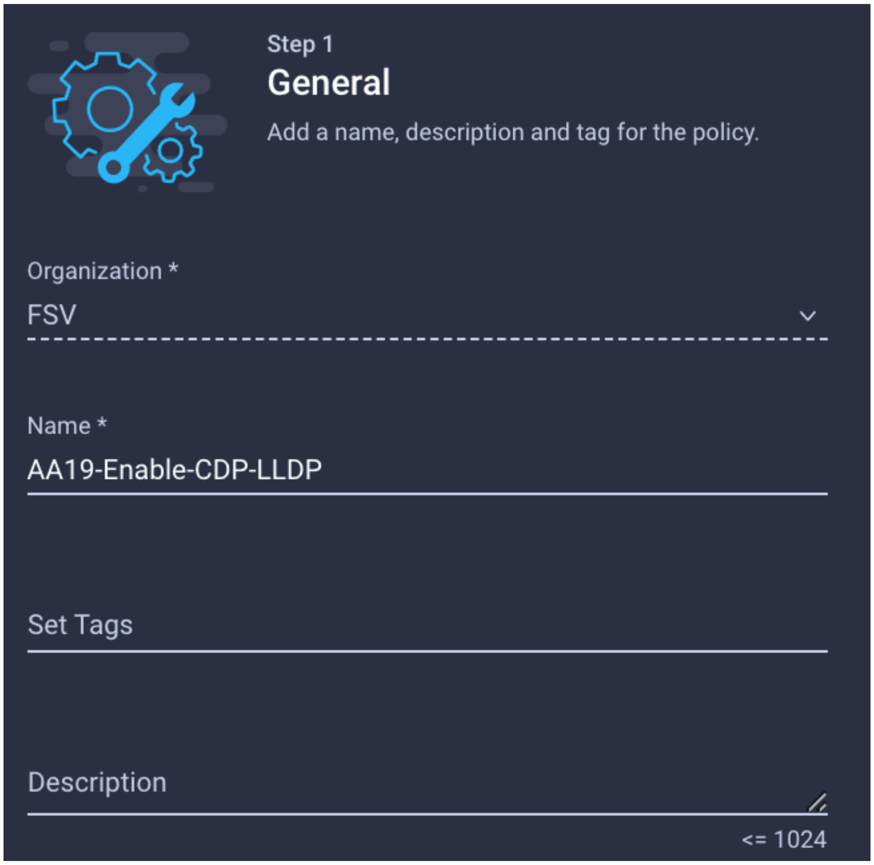

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-Enable-CDP-LLDP

Set Tags

Description

<= 1024

3. Click Next.

4. Enable CDP and both Transmit and Receive under LLDP.

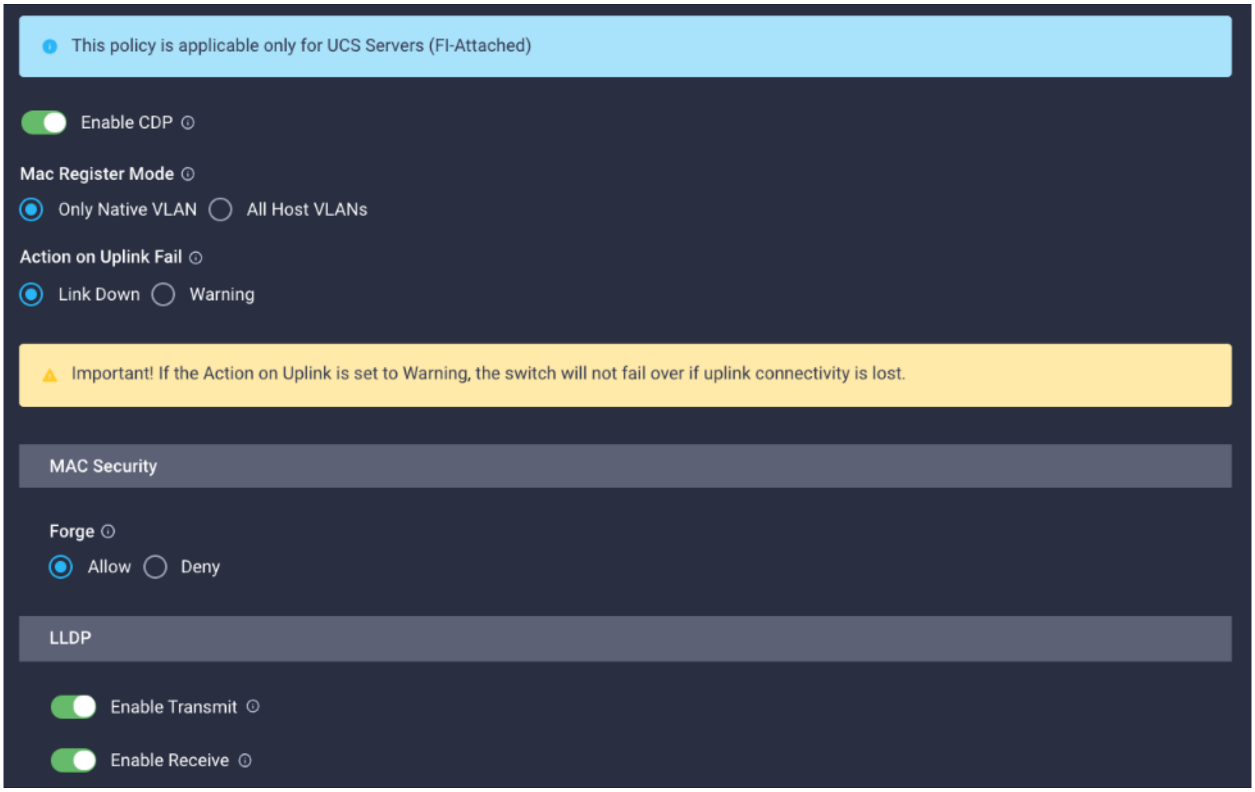

This policy is applicable only for UCS Servers (FI-Attached)

Enable CDP (on)

Mac Register Mode

● Only Native VLAN ○ All Host VLANs

Action on Uplink Fail

● Link Down ○ Warning

Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost.

MAC Security

Forge

● Allow ○ Deny

LLDP

Enable Transmit (on)

Enable Receive (on)

5. Click Create to finish creating Ethernet Network Control Policy.

Create Ethernet QoS Policy

Ethernet QoS policy is used to enable jumbo maximum transmission units (MTUs) for all the vNICs. A single policy will be created and reused for all the vNICs. To create the ethernet QoS policy, follow these steps:

1. Click Select Policy under Ethernet QoS and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-EthQoS-Pol).

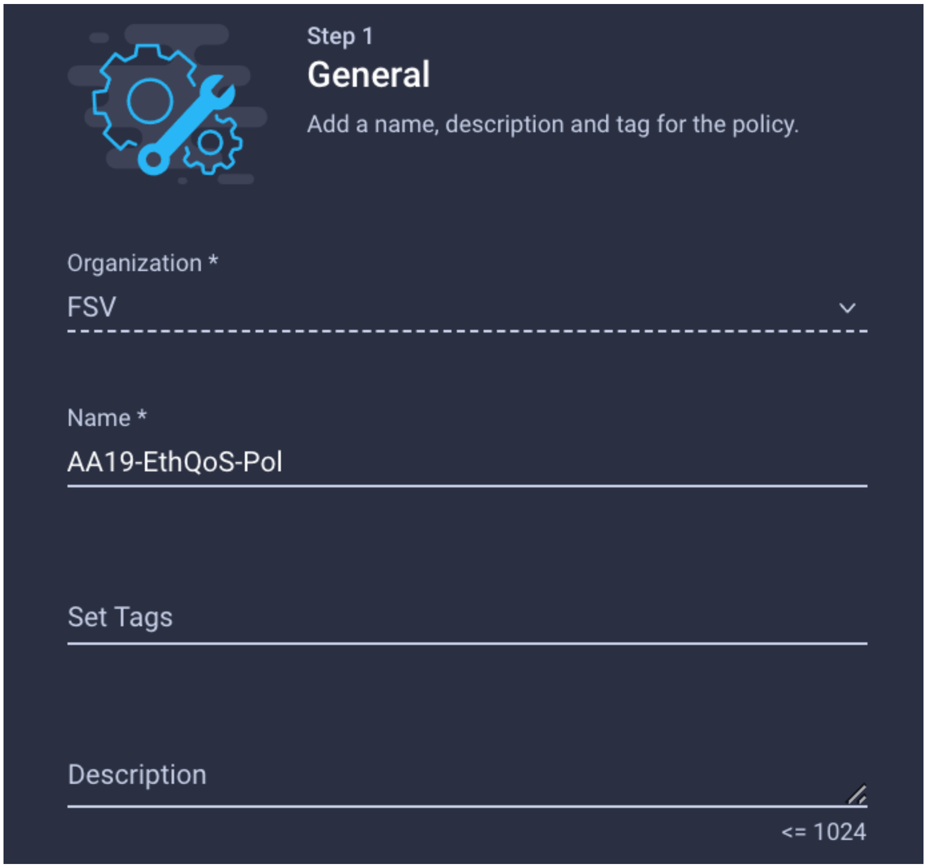

Step 1

General

Add a name, description and tag for the policy.

Organization *
FSV

Name *
AA19-EthQoS-Pol

Set Tags

Description

<= 1024

3. Click Next.

4. Change the MTU, Bytes to 9000.

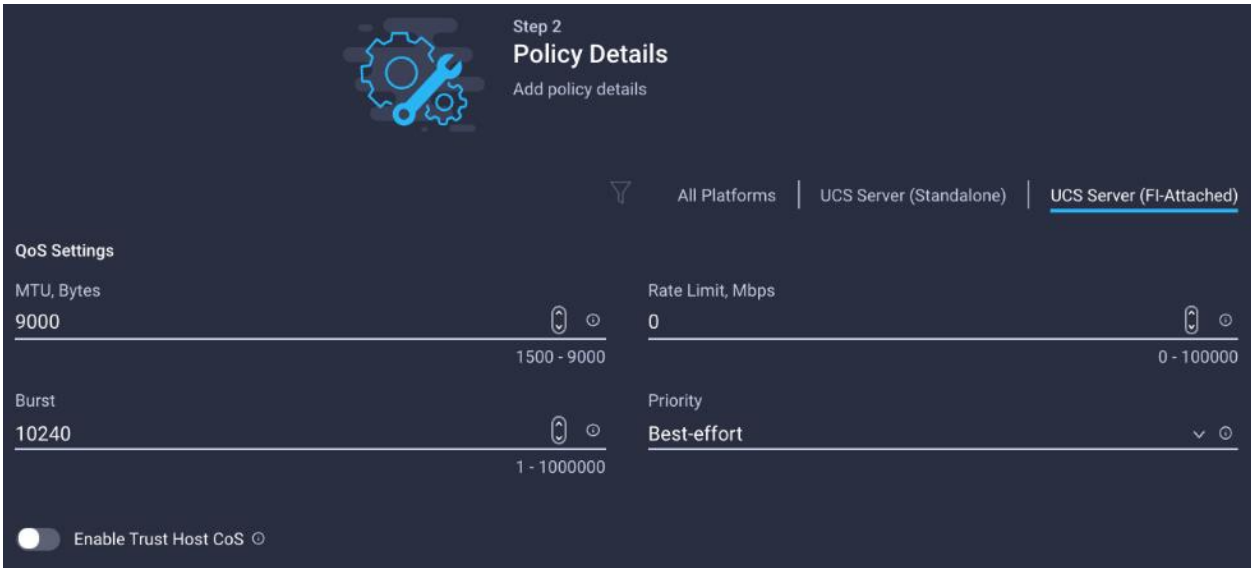

Step 2

Policy Details

Add policy details

QoS Settings

MTU, Bytes
9000
1500 - 9000

Rate Limit, Mbps
0
0 - 100000

Burst
10240
1 - 1000000

Priority
Best-effort

Enable Trust Host CoS

5. Click Create to finish setting up the Ethernet QoS Policy.

Create Ethernet Adapter Policy

Ethernet adapter policy is used to set the interrupts and the send and receive queues. The values are set according to the best-practices guidance for the operating system in use. Cisco Intersight provides default VMware Ethernet Adapter policy for typical VMware deployments.

Customers can optionally configure a tweaked ethernet adapter policy for additional hardware receive queues handled by multiple CPUs in scenarios where there is a lot of vMotion traffic and multiple flows. In this deployment, a modified ethernet adapter policy, AA19-VMware-High-Traffic, is created and attached to the 02-VDS0-A and 03-VDS0-B interfaces which handle vMotion.

Table 18. Ethernet Adapter Policy association to vNICs

Policy Name	vNICs	Policy Name
AA19-EthAdapter-VMware	00-vSwitch0-A, 01-vSwitch0-B	AA19-EthAdapter-VMware
AA19-VMware-High-Traffic	02-VDS0-A, 03-VDS0-B,	AA19-VMware-High-Traffic

1. Click Select Policy under Ethernet Adapter and in the pane on the right, click Create New.
2. Provide a name for the policy (for example, AA19-EthAdapter-VMware).

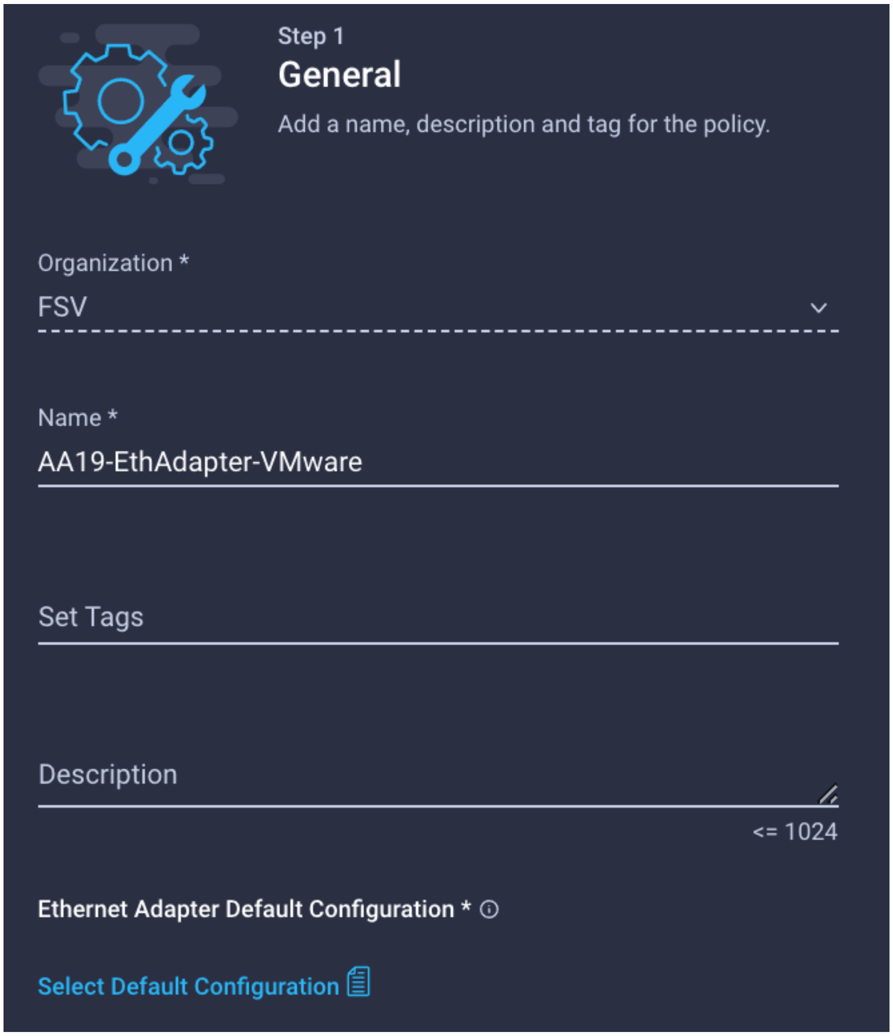
Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-EthAdapter-VMware

Set Tags

Description

<= 1024

Ethernet Adapter Default Configuration *

Select Default Configuration

3. Select VMWare under Ethernet Adapter Default Configuration:

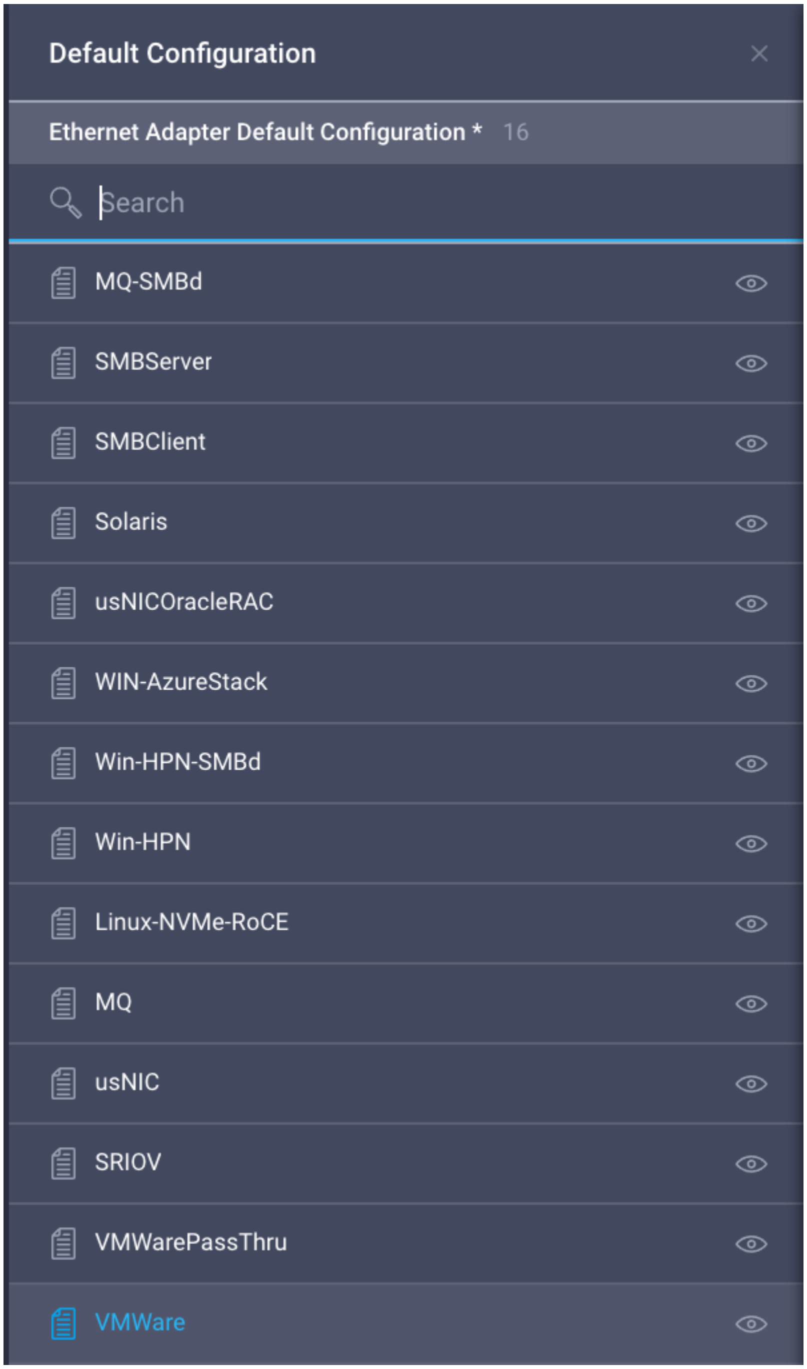

Default Configuration

Ethernet Adapter Default Configuration * 16

- MQ-SMBd
- SMBServer
- SMBClient
- Solaris
- usNICOracleRAC
- WIN-AzureStack
- Win-HPN-SMBd
- Win-HPN
- Linux-NVMe-RoCE
- MQ
- usNIC
- SRIOV
- VMWarePassThru
- VMWare

4. Verify all the Policies are assigned to the vNIC 00-vSwitch0-A.

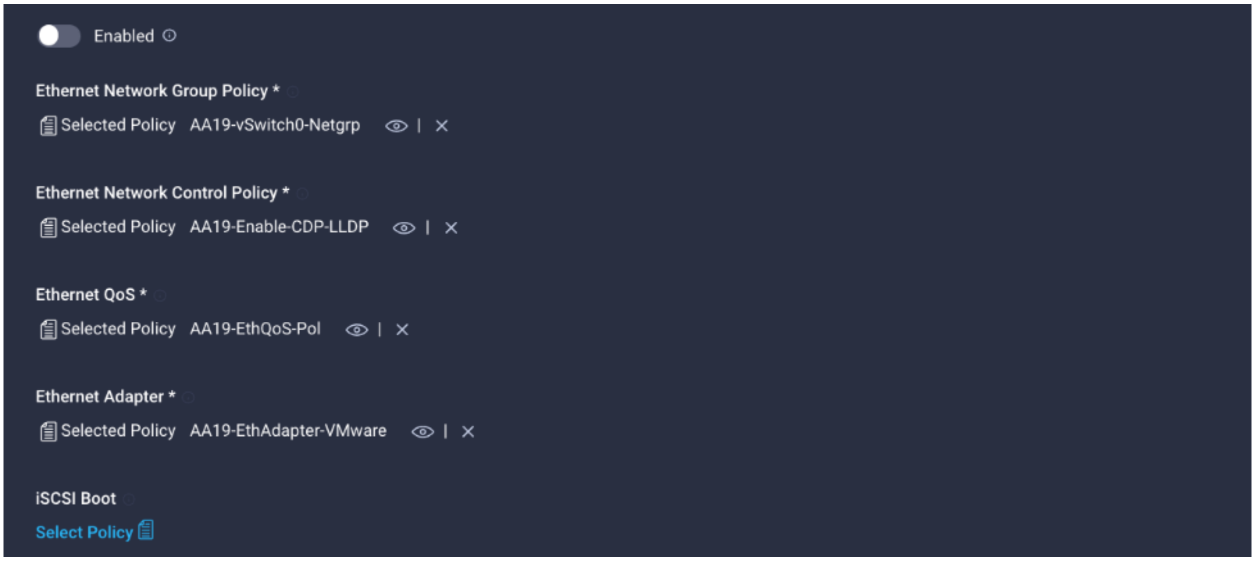

Enabled

Ethernet Network Group Policy *

Selected Policy AA19-vSwitch0-Netgrp

Ethernet Network Control Policy *

Selected Policy AA19-Enable-CDP-LLDP

Ethernet QoS *

Selected Policy AA19-EthQoS-Pol

Ethernet Adapter *

Selected Policy AA19-EthAdapter-VMware

iSCSI Boot

Select Policy

5. Click Add to add the vNIC.

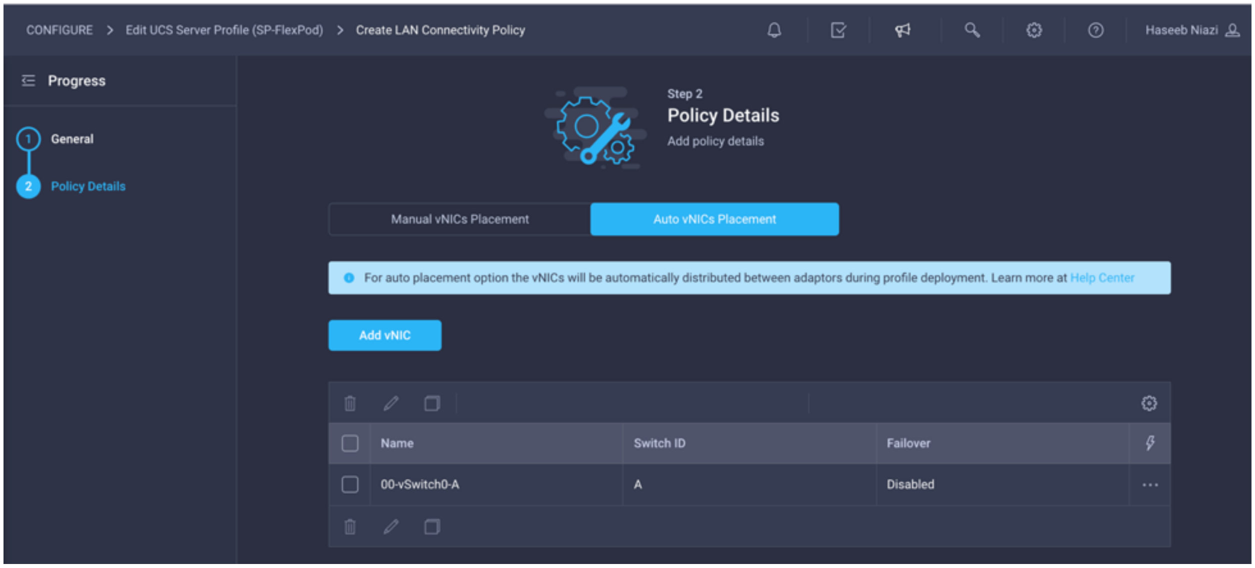

For auto placement option the vNICs will be automatically distributed between adaptors during profile deployment. Learn more at Help Center

Name	Switch ID	Failover
00-vSwitch0-A	A	Disabled

Note: Repeat the steps under Step 5 – Network Configuration to create additional vNICs. Most of the policies created for the first vNIC will be re-used for the remaining vNICs. MAC-Pool-B and VDS0-NetGrp-Policy used for subsequent vNICs are explained below.

6. The MAC-Pool-B is used by vNICs 01-vSwitch0-B and 03-VDS0-B. When adding the vNIC 01-vSwitch0-B, click Select Pool under MAC Address Pool and click Create New in the pane on the right.

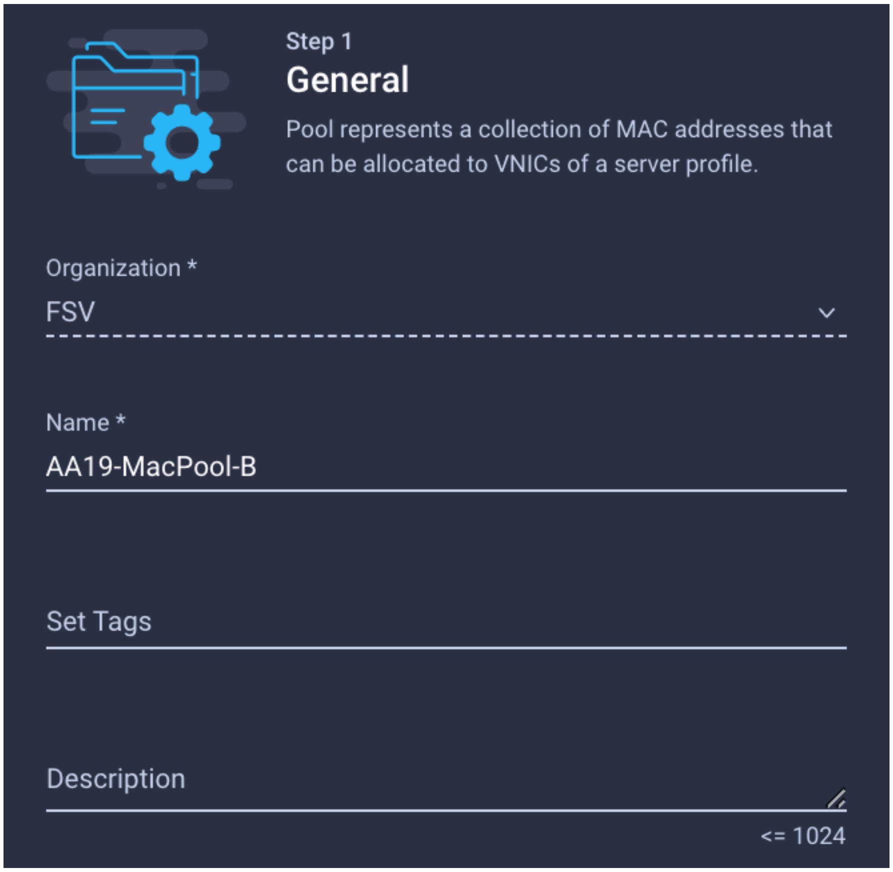

Step 1

General

Pool represents a collection of MAC addresses that can be allocated to VNICs of a server profile.

Organization *

FSV

Name *

AA19-MacPool-B

Set Tags

Description

<= 1024

7. While the same prefix 00:25:B5:19 is used for MAC Pool B, 0B in the second to last octet signifies these MAC addresses are assigned to vNICs associated with Fabric-B.

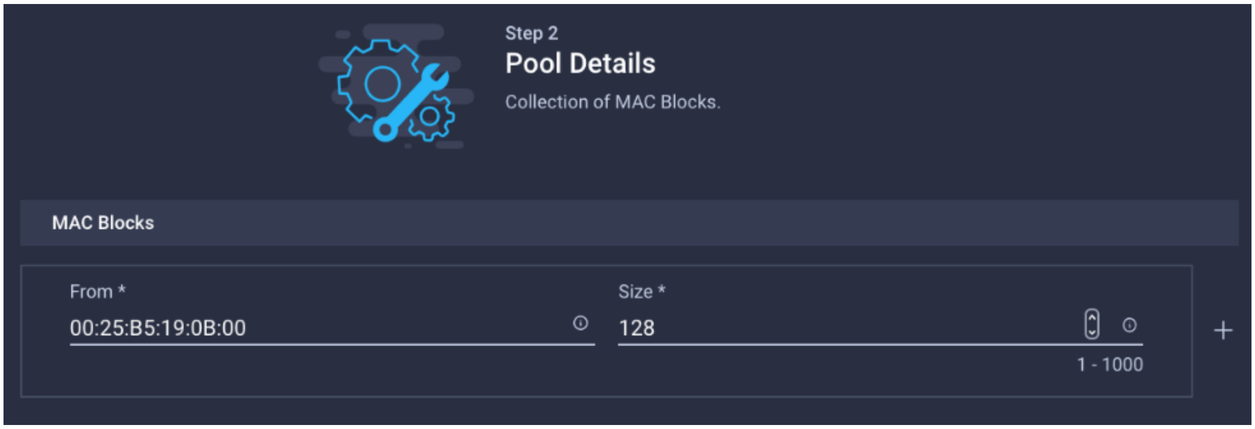

Step 2

Pool Details

Collection of MAC Blocks.

MAC Blocks

From *

00:25:B5:19:0B:00

Size *

128

1 - 1000

8. Configure the following parameters:

    a. Provide the name of vNIC (for example, 01-vSwitch0-B).

    b. Enter PCIe slod ID as MLOM where the VIC adapter is installed.

    c. Enter the value of PCI Link as 0

    d. Select the Switch ID as B, the fabric port to which the vNICs are associated.

e. Enter PCI Order a 5, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter

f. From the drop-down list, select vNIC Name for Consistent Device Naming (CDN).

g. Leave the Failover disabled. The failover will be supported by attaching multiple NICs to VMware vSwitch and VDS.

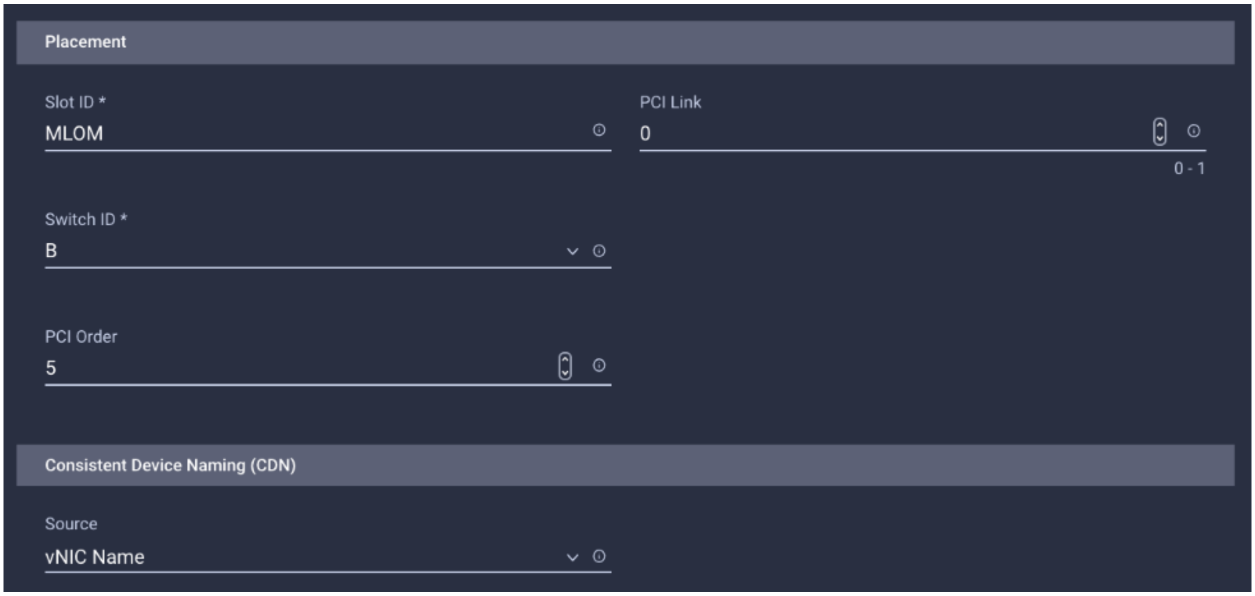

Placement

Slot ID * MLOM

PCI Link 0

0 - 1

Switch ID * B

PCI Order 5

Consistent Device Naming (CDN)

Source vNIC Name

9. The figure below shows various settings associate with 01-vSwitch0-B:

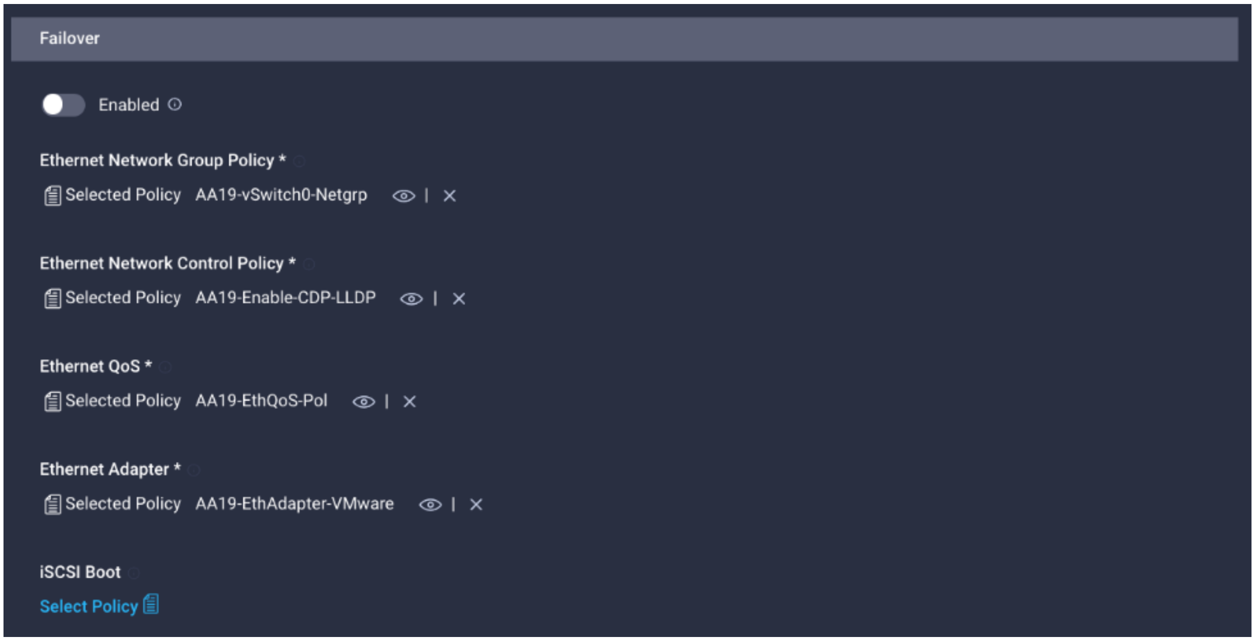

Failover

Enabled

Ethernet Network Group Policy *

Selected Policy AA19-vSwitch0-Netgrp

Ethernet Network Control Policy *

Selected Policy AA19-Enable-CDP-LLDP

Ethernet QoS *

Selected Policy AA19-EthQoS-Pol

Ethernet Adapter *

Selected Policy AA19-EthAdapter-VMware

iSCSI Boot

Select Policy

10. Click Select Policy under Ethernet Network Control Policy and in the pane on the right, click Create New.

11. Provide a name for the policy (for example, Enable-CDP-LLDP).

Network Group Policy for VDS0

The Network Group Policy for vNICs 02-VDS0-A and 03-VDS0-B is different because the VLANs used for the VDS are different, also the Network adapter policy is different for high throughput. The Network Group Policy and the Network Adapter Policy for VDS0 can be defined when adding the 02-VDS0-A vNIC. To configure the policy for VDS0, follow these steps:

1. To manually select the vNIC placement, select Manual vNIC Placement.

2. Click Add vNIC.

3. Configure the following parameters

   a. Provide the name of vNIC (for example, 02-VDS0-A).

   b. Enter PCIe slot ID as MLOM where the VIC adapter is installed.

   c. Enter the value of PCI Link as 0

   d. Select the Switch ID as A, the fabric port to which the vNICs are associated.

   e. Enter PCI Order a 6, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter

   f. From the drop-down list, select vNIC Name for Consistent Device Naming (CDN).

   g. Leave the Failover disabled. The failover will be supported by attaching multiple NICs to VMware vSwitch and VDS.

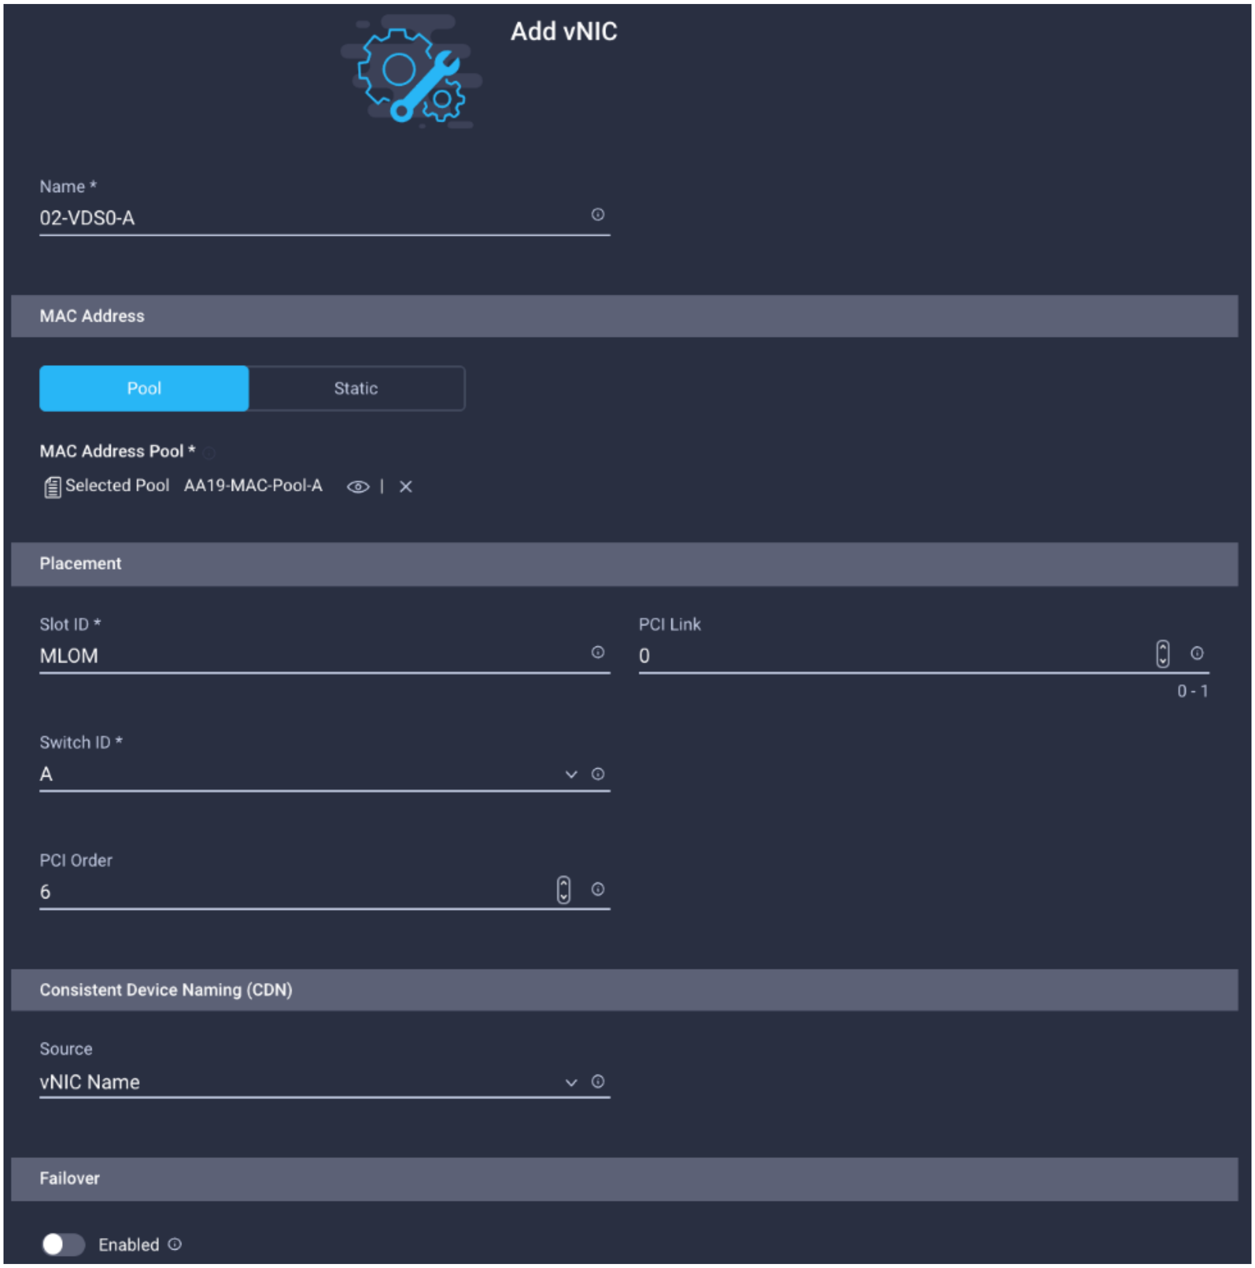

Add vNIC

Name *
02-VDS0-A

MAC Address

Pool | Static

MAC Address Pool *
Selected Pool AA19-MAC-Pool-A

Placement

Slot ID *
MLOM

PCI Link
0
0 - 1

Switch ID *
A

PCI Order
6

Consistent Device Naming (CDN)

Source
vNIC Name

Failover

Enabled

4. Instead of selecting the pre-exiting Network Group Policy, select Create New in the pane on the right.

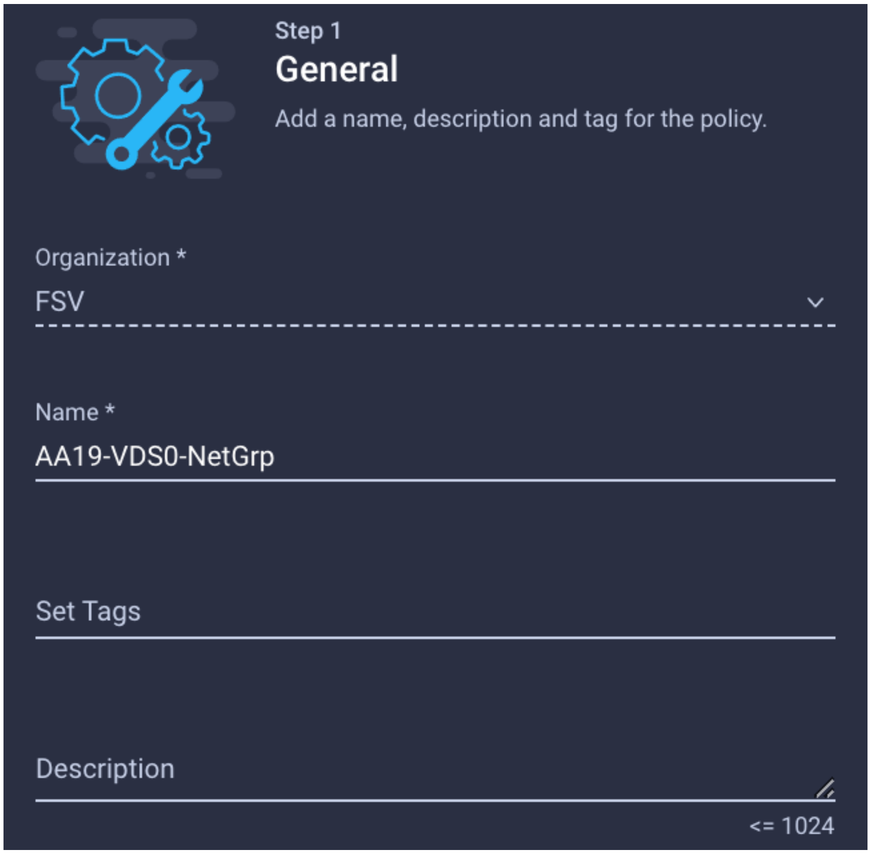

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-VDS0-NetGrp

Set Tags

Description

<= 1024

5. Define the correct VLANs associated with VDS (application traffic and vMotion traffic).

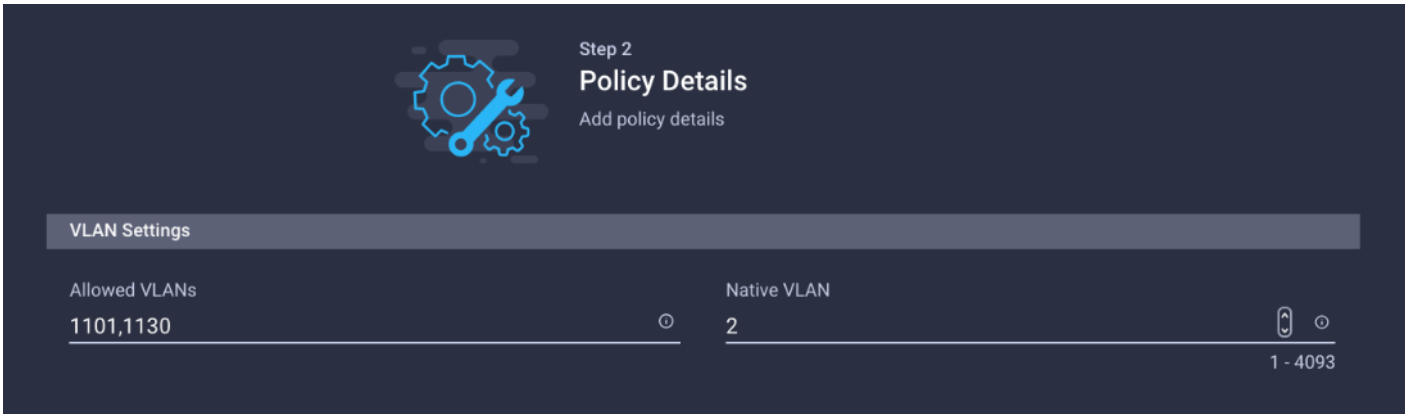

Step 2

Policy Details

Add policy details

VLAN Settings

Allowed VLANs

1101,1130

Native VLAN

2

1 - 4093

Create Ethernet Adapter Policy

To create the Ethernet Adapter Policy for 02-VDS0-A and 03-VDS-B interfaces which handle vMotion, follow these steps:

1. Click Select Policy under Ethernet Adapter and in the pane on the right, click Create New.
2. Provide a name for the policy (for example, AA19-VMware-High-Traffic).

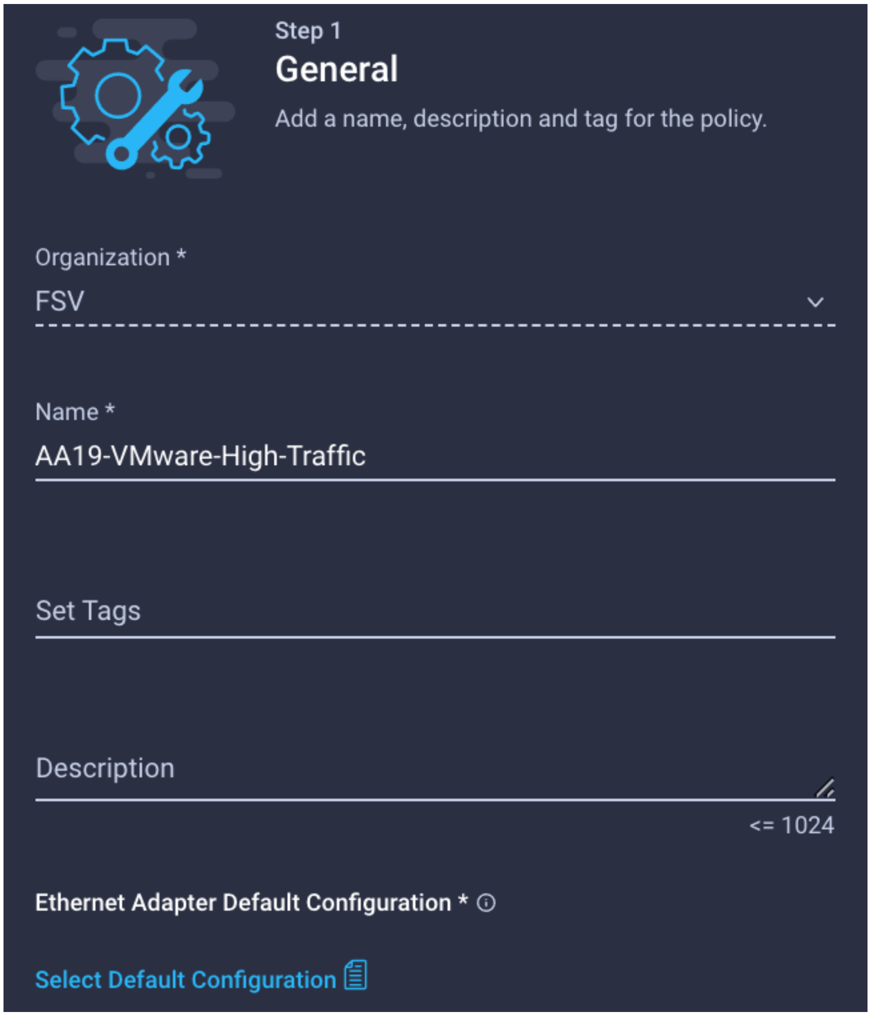

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-VMware-High-Traffic

Set Tags

Description

<= 1024

Ethernet Adapter Default Configuration *

3. Select VMWare under Ethernet Adapter Default Configuration

4. Click Next.

5. For the AA19-EthAdapter-VMware policy, click Create and skip the rest of the steps in this “Create Ethernet Adapter Policy” section.

6. For the optional AA19-VMware-High-Traffic policy (for VDS interfaces), make the following modifications to the policy:

   - Increase Interrupts to 11
   - Increase Receive Queue Count to 8
   - Increase Completion Queue Count to 9
   - Enable Receive Side Scaling

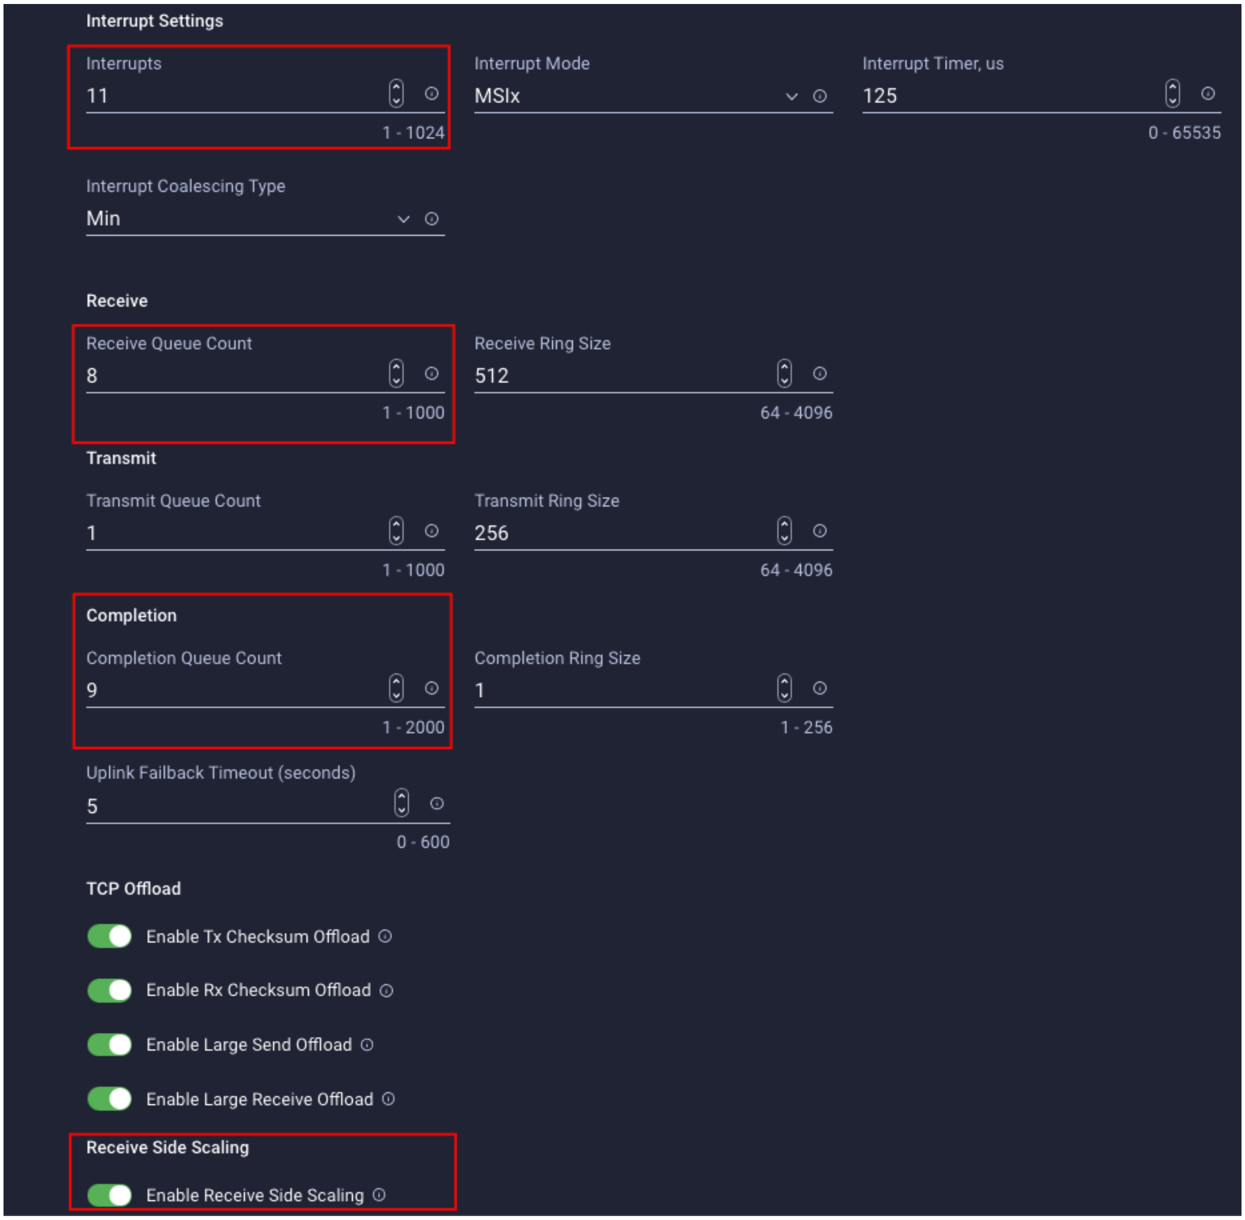

Interrupt Settings

Interrupts: 11 (1 - 1024)

Interrupt Mode: MSIx

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type: Min

Receive

Receive Queue Count: 8 (1 - 1000)

Receive Ring Size: 512 (64 - 4096)

Transmit

Transmit Queue Count: 1 (1 - 1000)

Transmit Ring Size: 256 (64 - 4096)

Completion

Completion Queue Count: 9 (1 - 2000)

Completion Ring Size: 1 (1 - 256)

Uplink Failback Timeout (seconds): 5 (0 - 600)

TCP Offload

Enable Tx Checksum Offload

Enable Rx Checksum Offload

Enable Large Send Offload

Enable Large Receive Offload

Receive Side Scaling

Enable Receive Side Scaling

7. Keep Receive Side Scaling and IPv4 Hash Enabled:

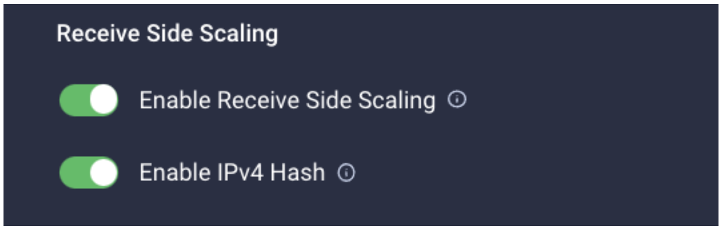

Receive Side Scaling

Enable Receive Side Scaling

Enable IPv4 Hash

8. Click Create.

9. Verify all the Policies are assigned to the vNIC 02-VDS0-A.

vNIC 02-VDS0-A Policy Mappings

The figure below shows various settings associated with 02-VDS0-A:

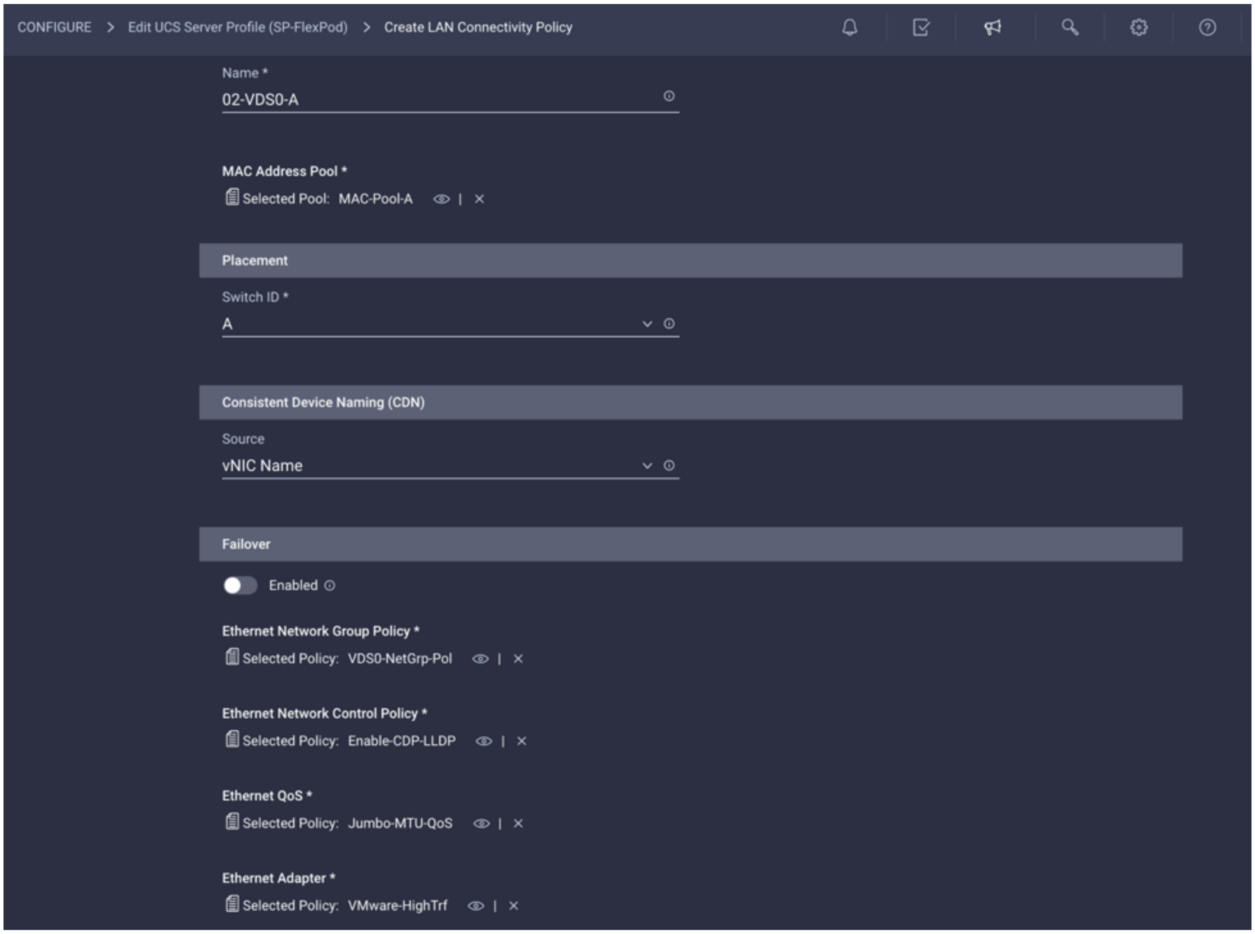

CONFIGURE > Edit UCS Server Profile (SP-FlexPod) > Create LAN Connectivity Policy

Name *
02-VDS0-A

MAC Address Pool *
Selected Pool: MAC-Pool-A

Placement

Switch ID *
A

Consistent Device Naming (CDN)

Source
vNIC Name

Failover

Enabled

Ethernet Network Group Policy *
Selected Policy: VDS0-NetGrp-Pol

Ethernet Network Control Policy *
Selected Policy: Enable-CDP-LLDP

Ethernet QoS *
Selected Policy: Jumbo-MTU-QoS

Ethernet Adapter *
Selected Policy: VMware-HighTrf

1. Click Add vNIC.

2. To manually select the vNIC placement, select Manual vNIC Placement.

3. Configure the following parameters:

   a. Provide the name of vNIC (for example, 02-VDS0-B).
   b. Enter the value of PCI Link as 0
   c. Select the Switch ID as B, the fabric port to which the vNICs are associated.
   d. Enter PCI Order a 7, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter
   e. From the drop-down list, select vNIC Name for Consistent Device Naming (CDN).
   f. Leave the Failover disabled. The failover will be supported by attaching multiple NICs to VMware vSwitch and VDS.

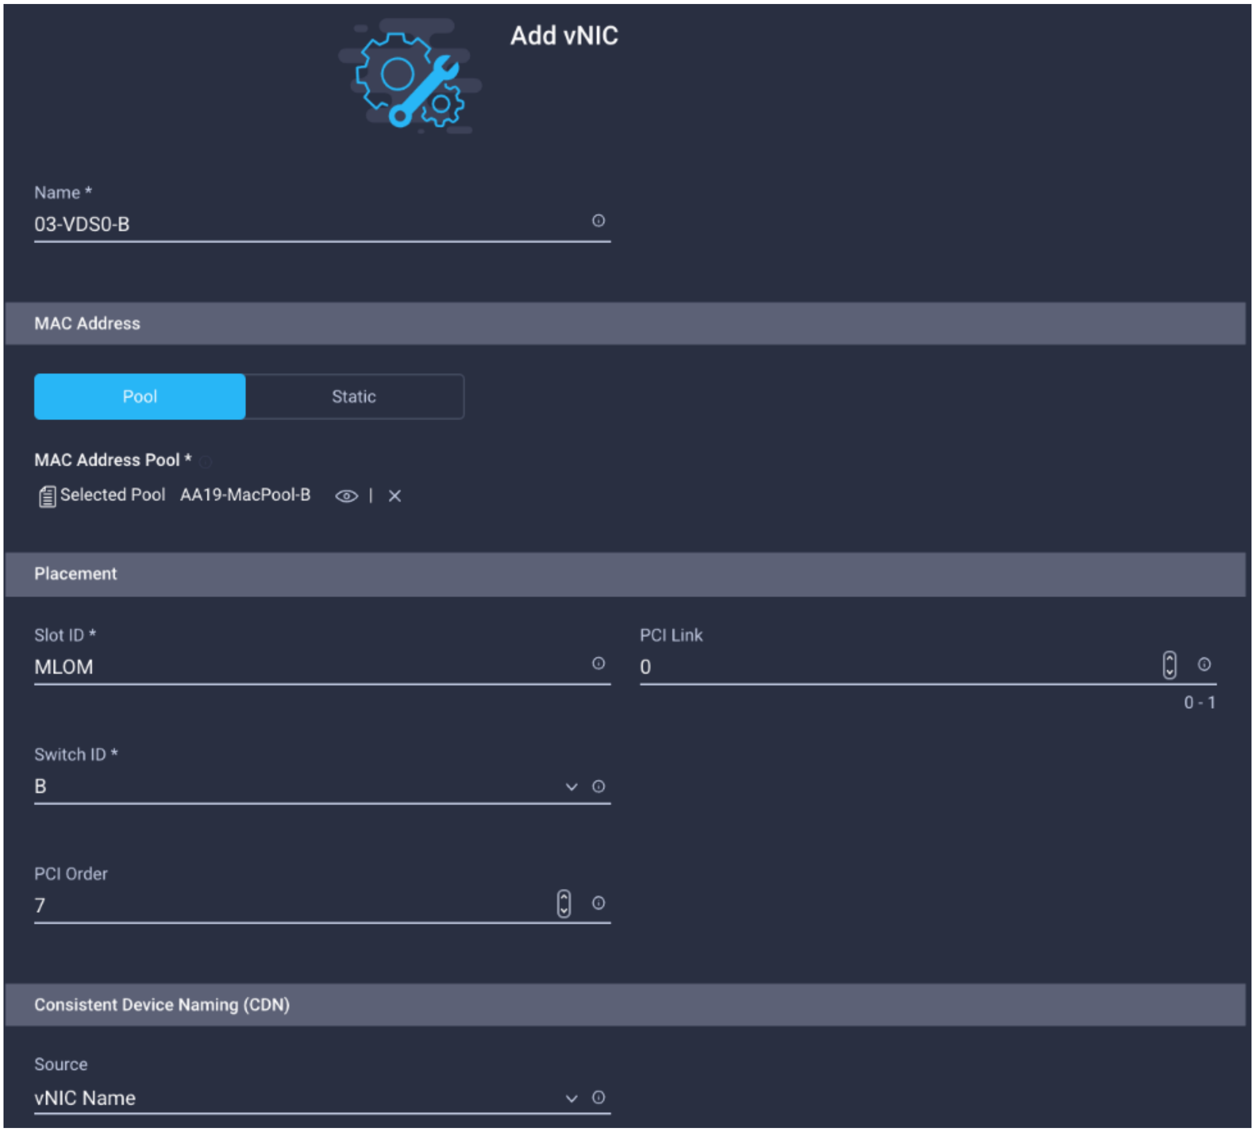

Add vNIC

Name * 03-VDS0-B

MAC Address

MAC Address Pool * Selected Pool AA19-MacPool-B

Placement

Slot ID * MLOM

PCI Link 0

0 - 1

Switch ID * B

PCI Order 7

Consistent Device Naming (CDN)

Source vNIC Name

The figure below shows various settings associated with 03-VDS0-B:

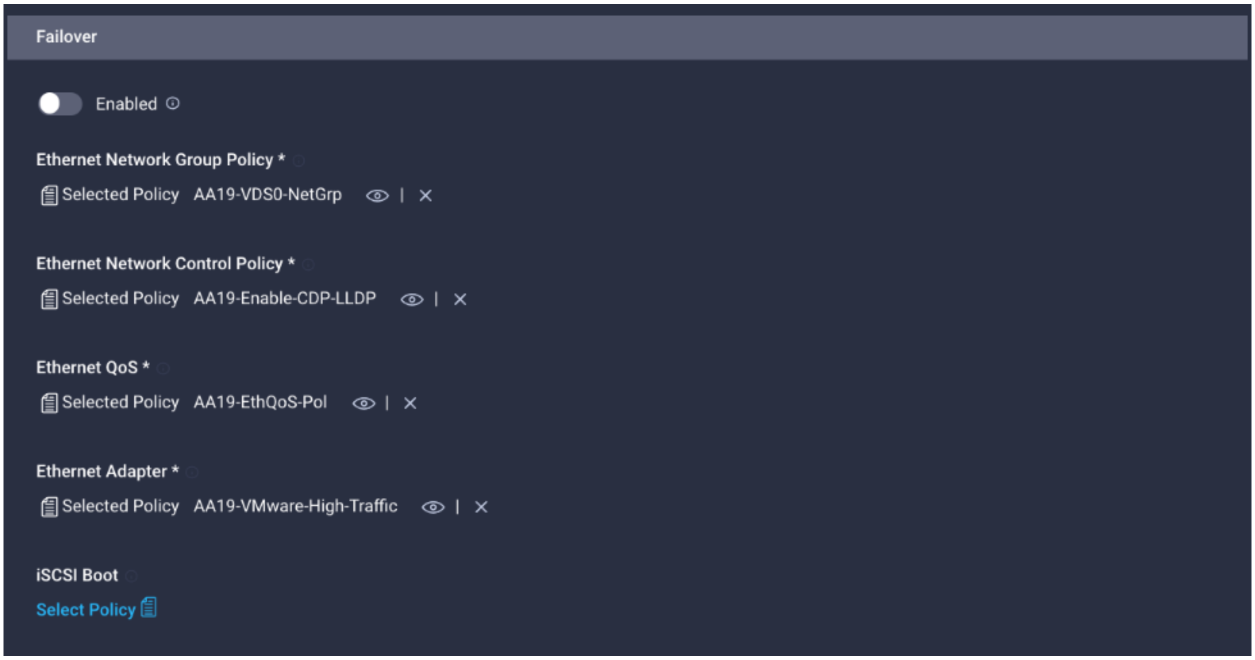

Failover

Enabled

Ethernet Network Group Policy *

Selected Policy AA19-VDS0-NetGrp

Ethernet Network Control Policy *

Selected Policy AA19-Enable-CDP-LLDP

Ethernet QoS *

Selected Policy AA19-EthQoS-Pol

Ethernet Adapter *

Selected Policy AA19-VMware-High-Traffic

iSCSI Boot

Select Policy

4. Verify all the vNICs are added successfully before moving on to SAN connectivity policies.

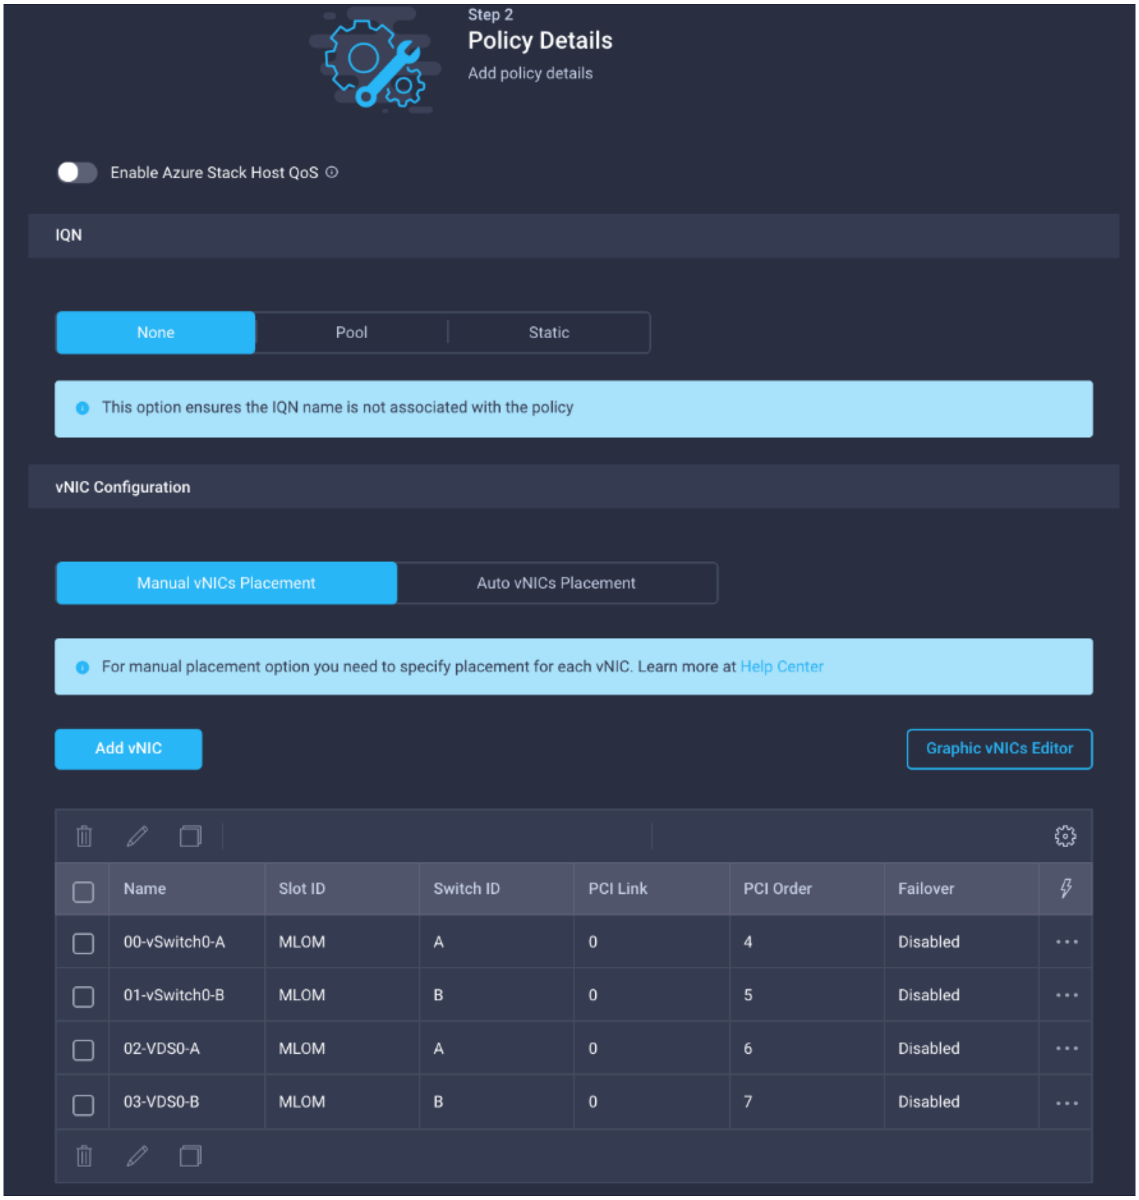

Step 2

Policy Details

Add policy details

Enable Azure Stack Host QoS

IQN

This option ensures the IQN name is not associated with the policy

vNIC Configuration

For manual placement option you need to specify placement for each vNIC. Learn more at Help Center

Name	Slot ID	Switch ID	PCI Link	PCI Order	Failover
00-vSwitch0-A	MLOM	A	0	4	Disabled
01-vSwitch0-B	MLOM	B	0	5	Disabled
02-VDS0-A	MLOM	A	0	6	Disabled
03-VDS0-B	MLOM	B	0	7	Disabled

5. Click Create.

Step 5b – Network Connectivity – SAN Connectivity Policy

A SAN connectivity policy determines the network storage resources and the connections between the server and the storage device on the network. This policy enables customers to configure the vHBAs that the servers use to communicate with the SAN.

Table 19 lists the details of two vHBAs that are used to provide FC connectivity and boot from SAN functionality. The other two vHBAs are used to provide fc-nvme connectivity.

Table 19. vHBA for boot from FC SAN

vNIC/vHBA Name	Slot	Switch ID
vHBA-A	MLOM	A
vHBA-B	MLOM	B

vNIC/vHBA Name	Slot	Switch ID
vHBA-NVMe-A	MLOM	A
vHBA-NVMe-B	MLOM	B
vNIC/vHBA Name	Slot	Switch ID

1. Click Select Policy next to SAN Connectivity and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-FC-ESXi-SANConn-Pol).

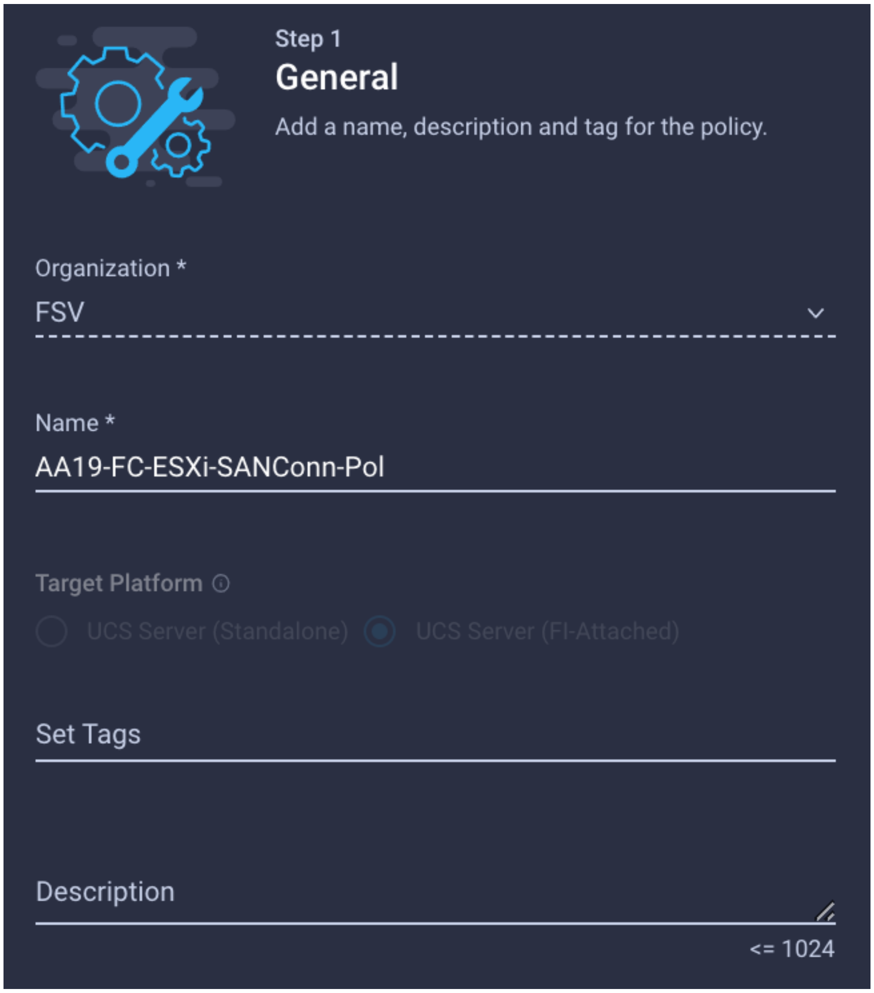

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-FC-ESXi-SANConn-Pol

Target Platform

○ UCS Server (Standalone) ● UCS Server (FI-Attached)

Set Tags

Description

<= 1024

Create vHBA - SAN A

Ethernet Network Control Policy is used to enable CDP and LLDP for the vNICs. A single policy will be created and reused for all the vNICs. To create the vHBA, follow these steps:

1. To manually select the vHBA placement, select Manual vHBAs Placement.

2. Click Add vHBA.

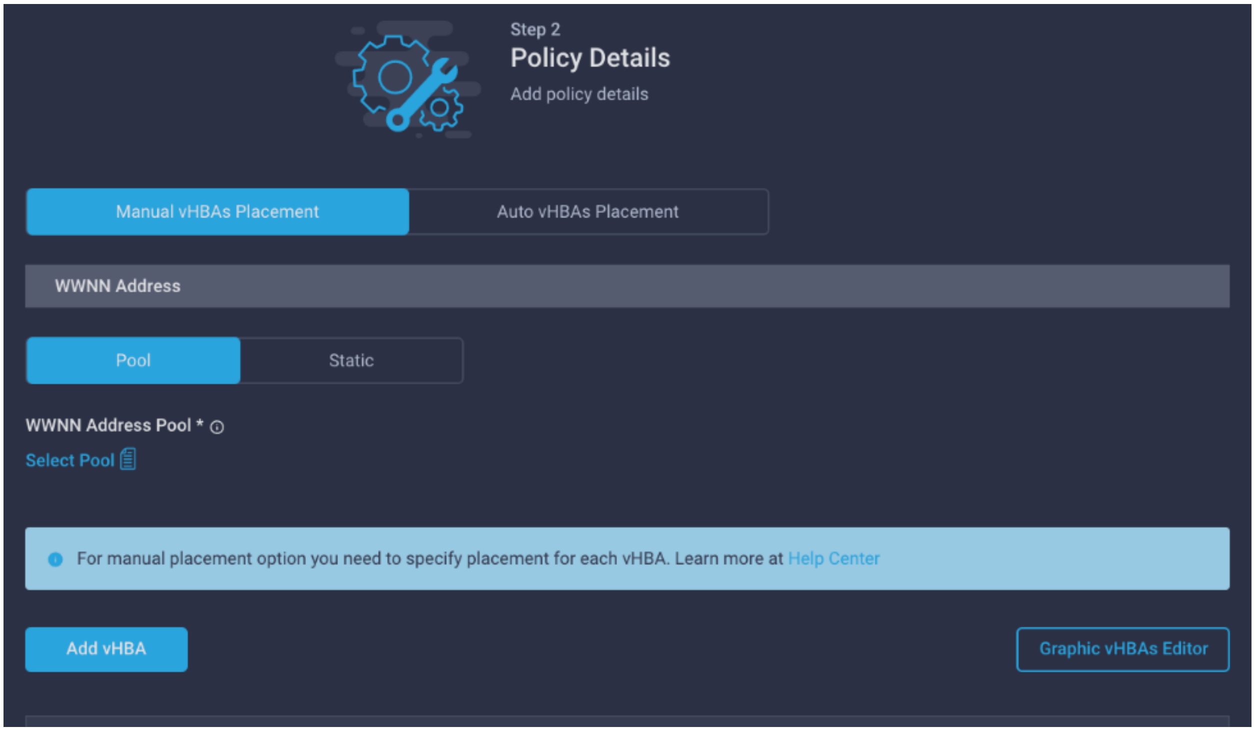

Step 2

Policy Details

Add policy details

WWNN Address

WWNN Address Pool * ⓘ

Select Pool

For manual placement option you need to specify placement for each vHBA. Learn more at Help Center

3. Configure the following parameters

   a. Provide the name of vNIC (for example, vHBA-A).

   b. Enter PCIe slod ID as MLOM where the VIC adapter is installed.

   c. Enter the PCI Link as 0

   d. Select the Switch ID as A, the fabric port to which the vHBAs are associated.

   e. Enter PCI Order a 0, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter

   f. In the vHBA Type, select fc-initiator from the drop-down list.

   g. Select Switch ID A from the drop-down list.

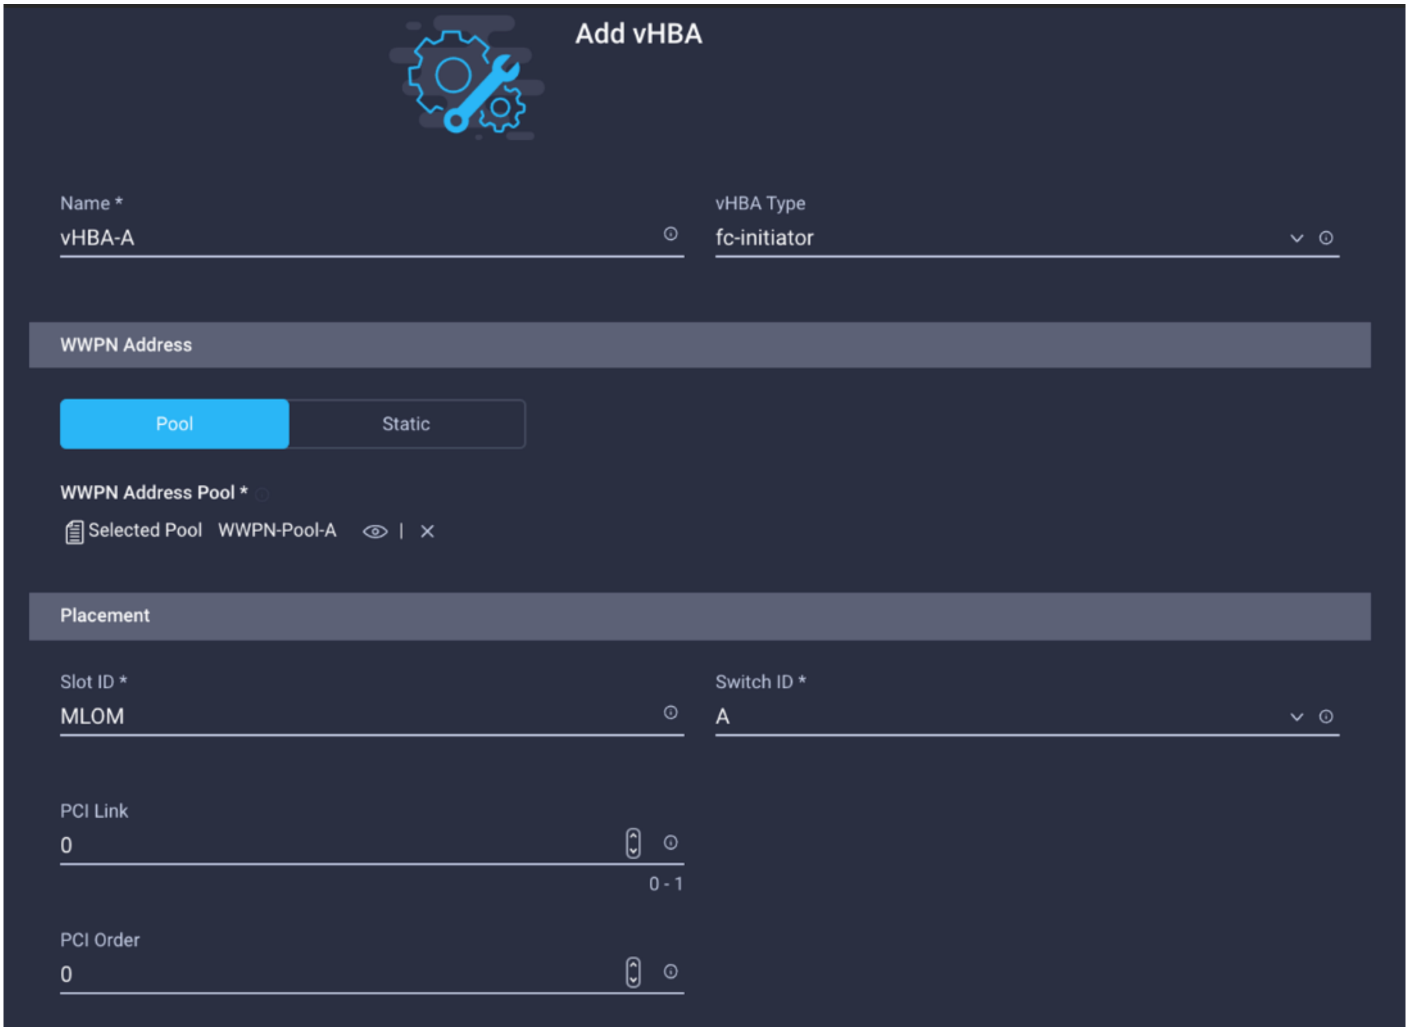

Add vHBA

Name * vHBA-A

vHBA Type fc-initiator

WWPN Address

WWPN Address Pool *

Selected Pool WWPN-Pool-A

Placement

Slot ID * MLOM

Switch ID * A

PCI Link 0

0 - 1

PCI Order 0

Since the WWNN Address Pools have not been defined yet, you will create a new WWNN address Pool first.

Create WWNN Address Pool

1. Click Select Pool under WWNN Address Pool and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, WWNN-Pool).

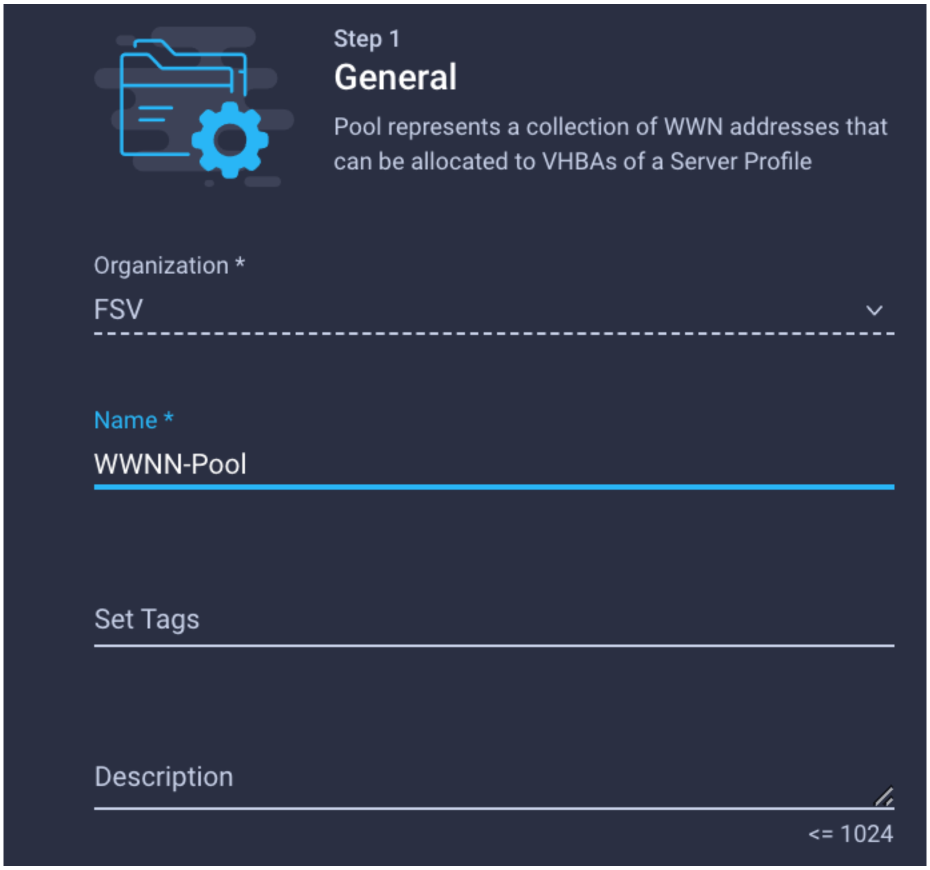

Step 1

General

Pool represents a collection of WWN addresses that can be allocated to VHBAs of a Server Profile

Organization * FSV

Name * WWNN-Pool

Set Tags

Description

<= 1024

3. Click Next.

4. Provide the start WWNN Block Address. Recommended prefix for WWNN addresses is 20:00:00:25:B5:xx:xx:xx. As a best practice, some extra information is always coded into the WWNN address pool for ease of troubleshooting. For example, in the figure below, 19 is the Rack ID.

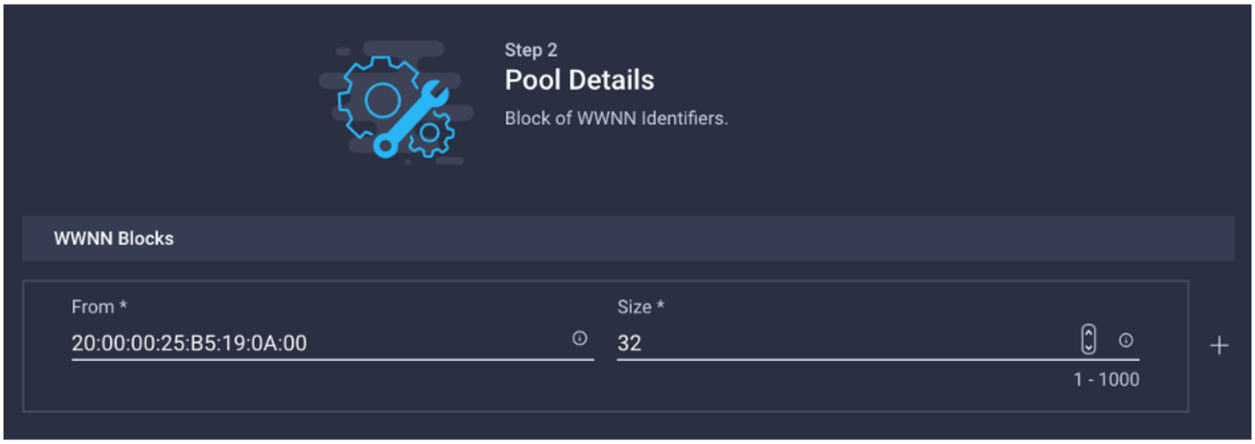

Step 2

Pool Details

Block of WWNN Identifiers.

WWNN Blocks

From * 20:00:00:25:B5:19:0A:00

Size * 32

1 - 1000

5. Click Create to finish creating the WWNN address pool.

Note: Since the WWPN Address Pool has not been defined yet, you will create a WWPN address Pool for Fabric-A first.

Create WWPN Pool – SAN A

1. Click Select Pool under WWPN Address Pool and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, WWPN-Pool-A).

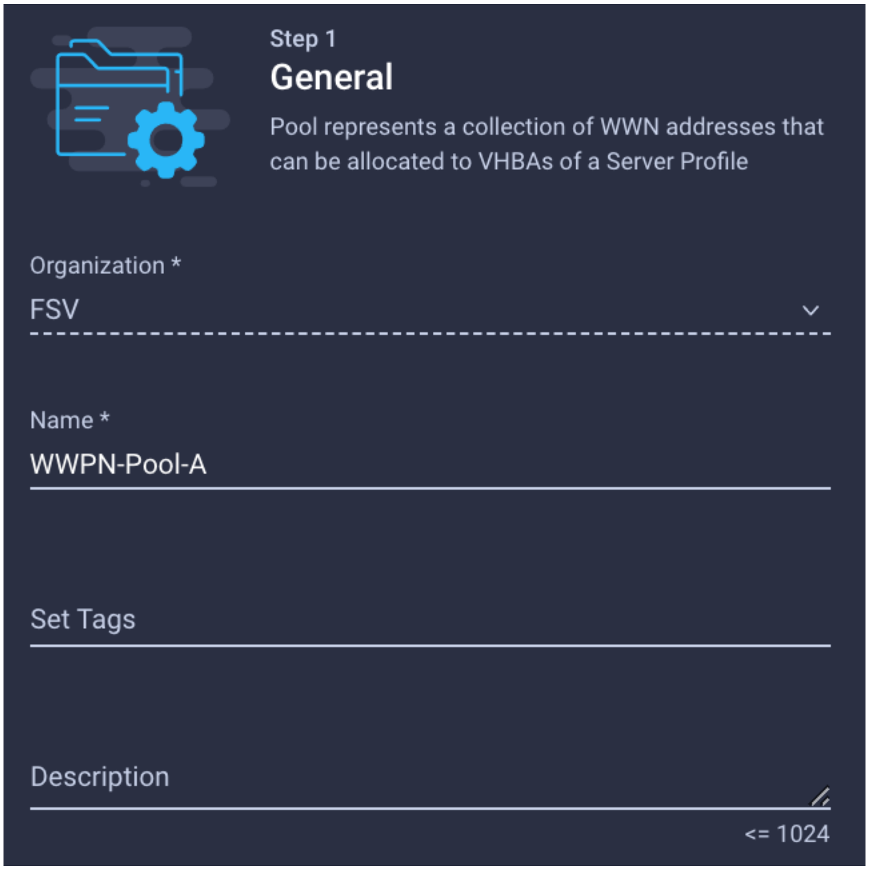

Step 1

General

Pool represents a collection of WWN addresses that can be allocated to VHBAs of a Server Profile

Organization *
FSV

Name *
WWPN-Pool-A

Set Tags

Description

<= 1024

3. Provide the start WWPN Block Address for SAN A. Recommended prefix for WWPN addresses is 20:00:00:25:B5:xx:xx:xx. As a best practice, in FlashStack some extra information is always coded into the WWPN address pool for ease of troubleshooting. For example, in the figure below, 19 is the Rack ID and 0A signifies SAN A.

4. Provide the size of the pool (for example, 32).

5. Click Create.

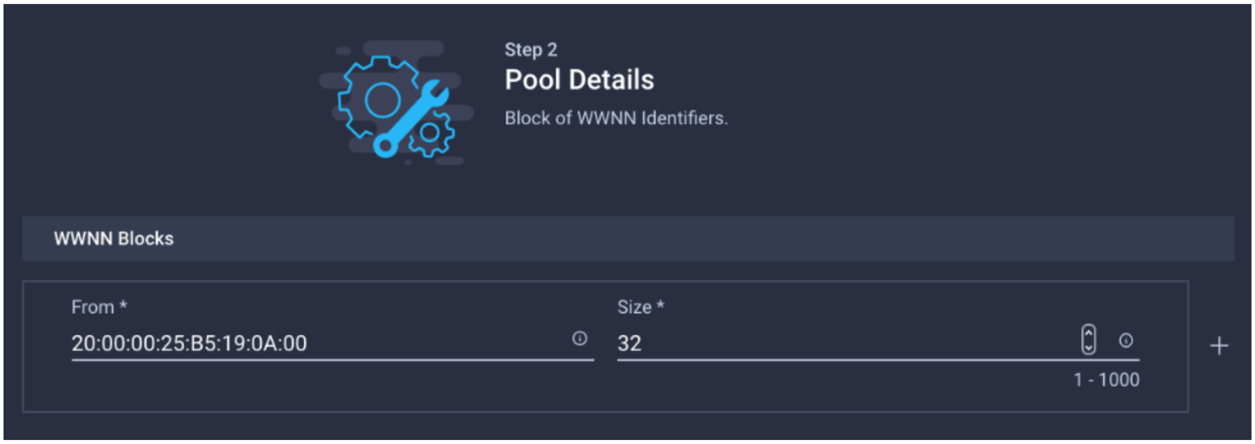
Step 2

Pool Details

Block of WWNN Identifiers.

WWNN Blocks

From *
20:00:00:25:B5:19:0A:00

Size *
32

1 - 1000

6. Click Create.

Fibre Channel Network Policy - SAN A

A Fibre Channel Network policy governs the Virtual Storage Area Network (VSAN) configuration for the virtual interfaces. VSAN 101 will be used for vHBA-A while VSAN 102 will be used for vHBA-B. To create the policy, follow these steps:

1. Click Select Policy under Fibre Channel Network and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, SAN-A-Network).

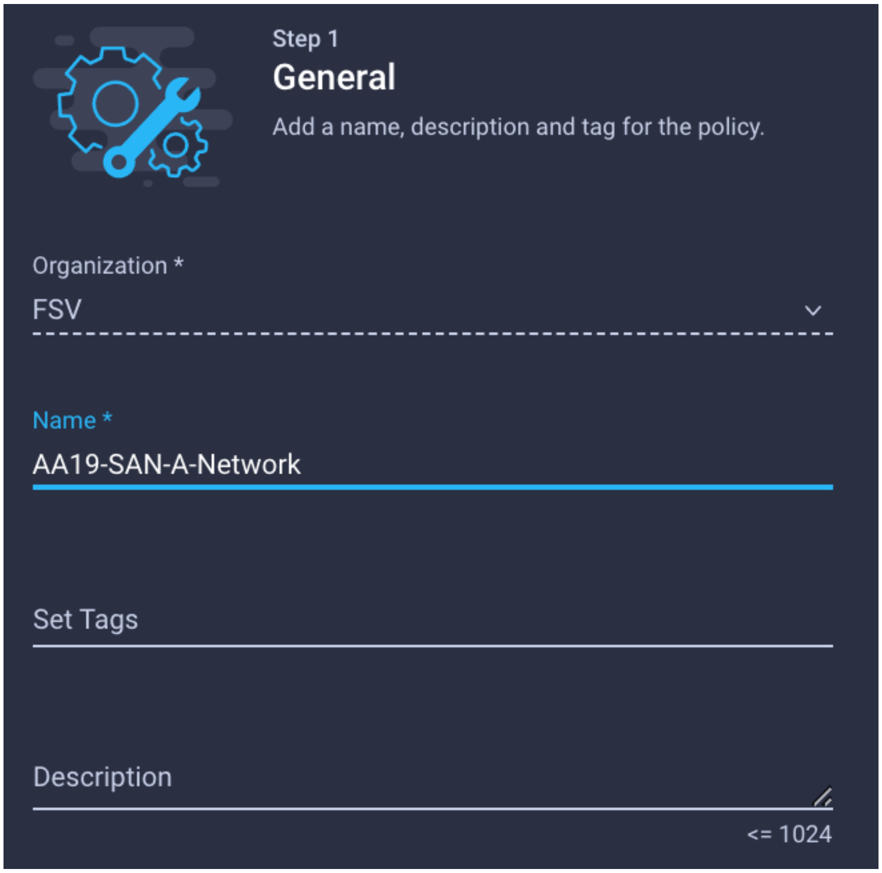
Step 1

General

Add a name, description and tag for the policy.

Organization *
FSV

Name *
AA19-SAN-A-Network

Set Tags

Description

<= 1024

3. For the scope, select UCS Server (FI-Attached).

4. Enter the VSAN information (for example, 101) under Default VLAN.

Note: The current GUI shows **Default VLAN** instead of **Default VSAN**. Enter the VLAN associated with VSAN-A.

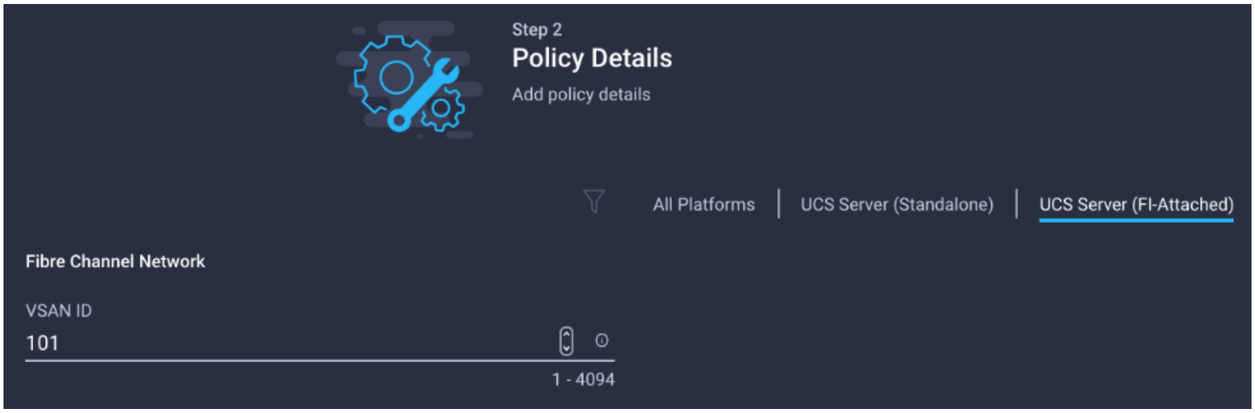

Step 2

Policy Details

Add policy details

Fibre Channel Network

VSAN ID

101

1 - 4094

5. Click Create to finish creating the “FC Network Policy.”

Fibre Channel QoS Policy

The Fibre Channel QoS policy assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. The Fibre Channel QoS policy used in this deployment will use default values and will be shared by both vHBA-A and vHBA-B. To create the policy, follow these steps:

1. Click Select Policy under Fibre Channel QoS and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, FC-QoS).

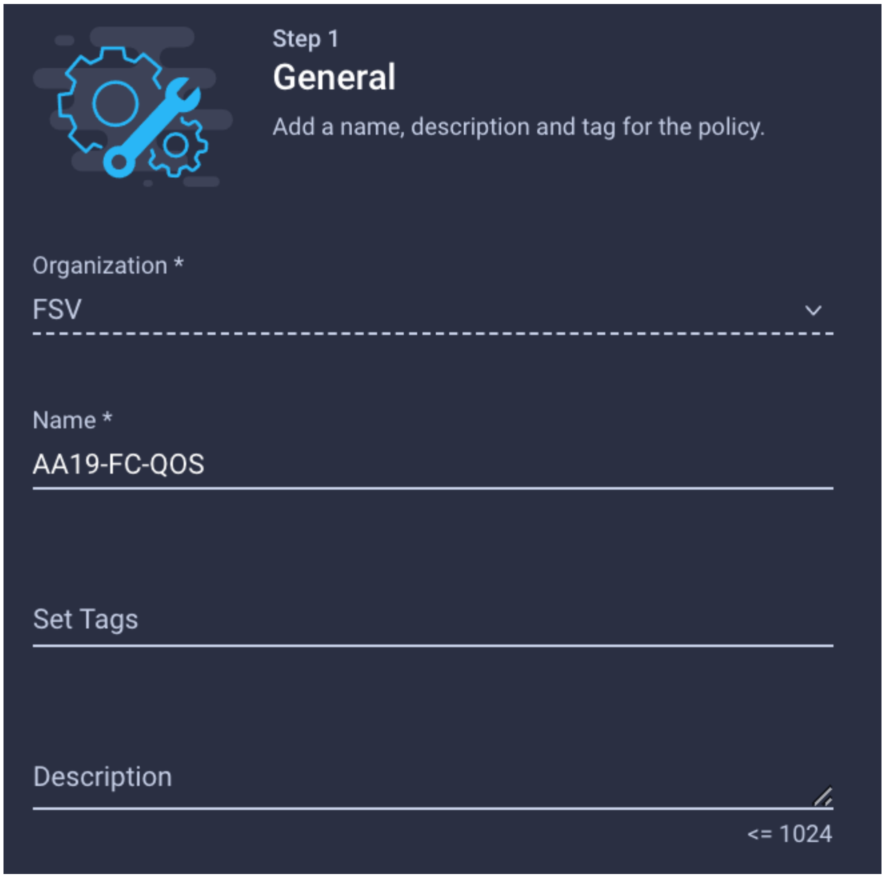

Step 1

General

Add a name, description and tag for the policy.

Organization *
FSV

Name *
AA19-FC-QOS

Set Tags

Description

<= 1024

3. For the scope, select UCS Server (FI-Attached).

4. Do not change the default values.

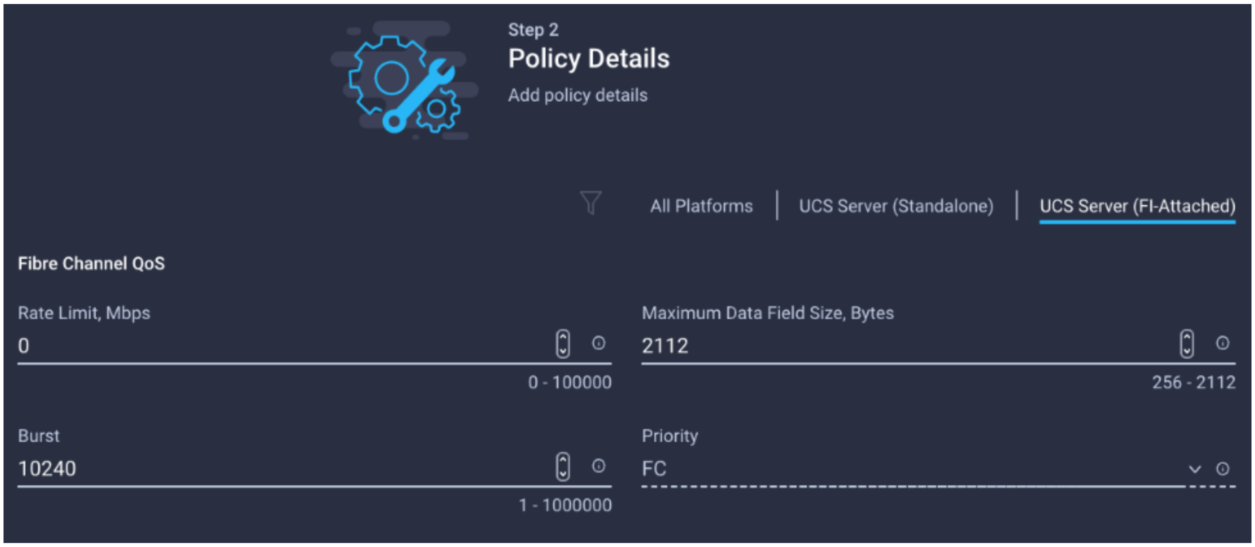

Step 2

Policy Details

Add policy details

Fibre Channel QoS

Rate Limit, Mbps
0
0 - 100000

Maximum Data Field Size, Bytes
2112
256 - 2112

Burst
10240
1 - 1000000

Priority
FC

5. Click Create to finish creating the FC QoS Policy.

Fibre Channel Adapter Policy

A Fibre Channel adapter policy governs the host-side behavior of the adapter, including how the adapter handles traffic. In this validation, we will use the default values for the adapter policy and the policy will be shared by both vHBA-A and vHBA-B. To create the policy, follow these steps:

1. Click Select Policy under Fibre Channel Adapter and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, FC-Adapter).

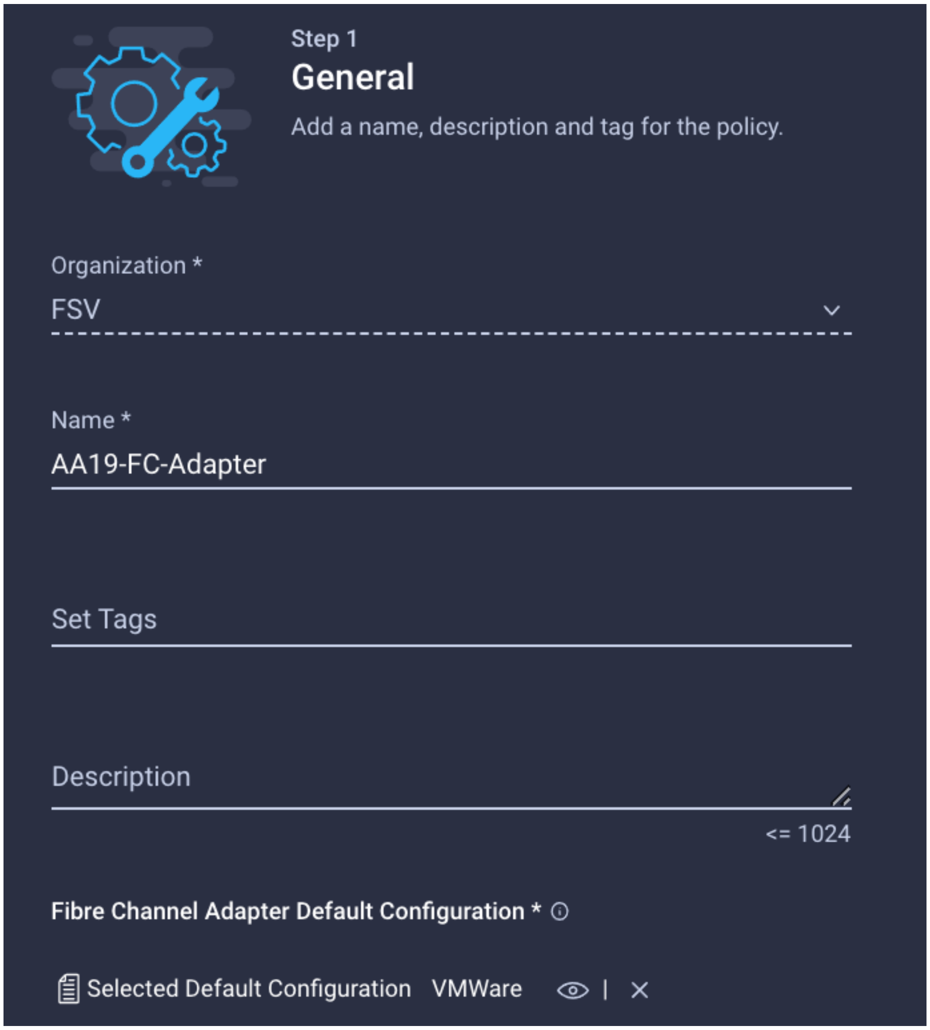

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-FC-Adapter

Set Tags

Description

<= 1024

Fibre Channel Adapter Default Configuration *

Selected Default Configuration VMWare

3. For the scope, select UCS Server (FI-Attached).

Note: Do not change the default values.

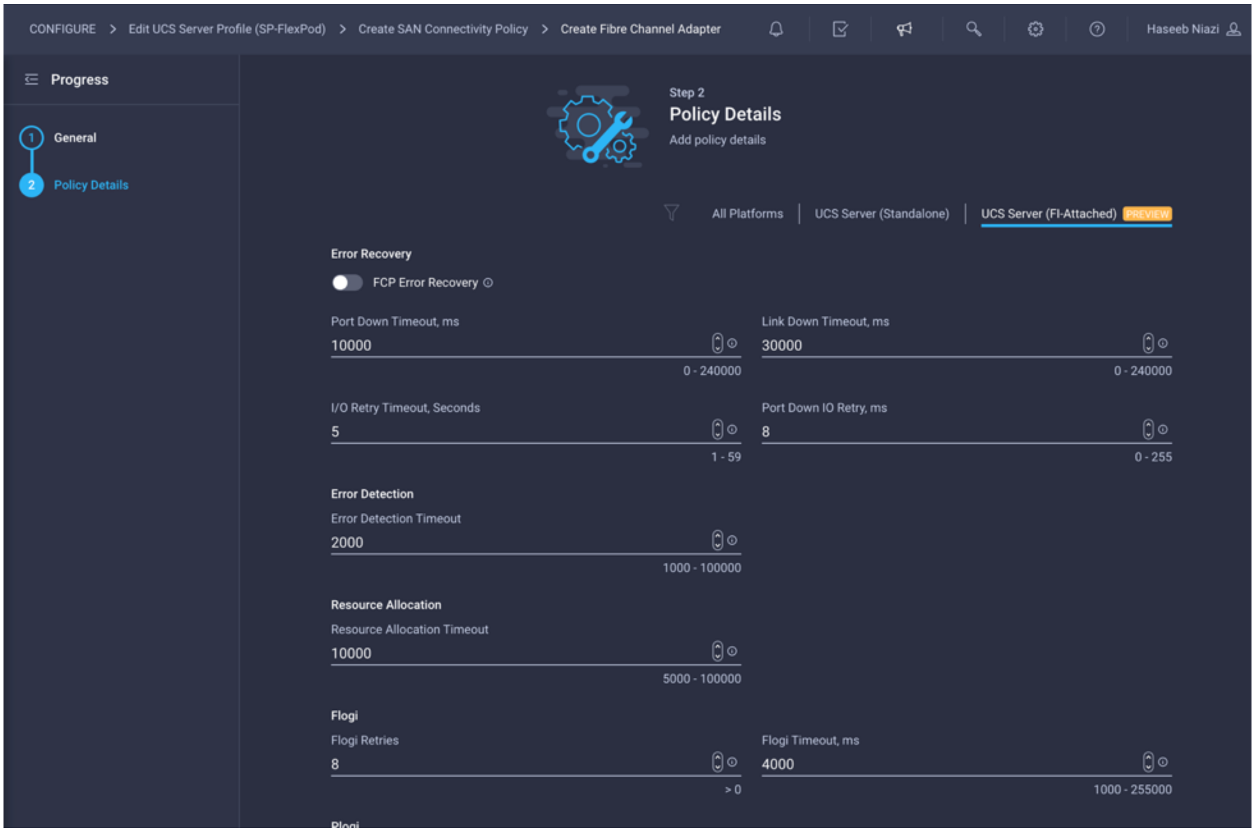

Error Recovery

FCP Error Recovery

Port Down Timeout, ms: 10000 (0 - 240000)

Link Down Timeout, ms: 30000 (0 - 240000)

I/O Retry Timeout, Seconds: 5 (1 - 59)

Port Down IO Retry, ms: 8 (0 - 255)

Error Detection

Error Detection Timeout: 2000 (1000 - 100000)

Resource Allocation

Resource Allocation Timeout: 10000 (5000 - 100000)

Flogi

Flogi Retries: 8 (> 0)

Flogi Timeout, ms: 4000 (1000 - 255000)

Plogi

4. Click Create to finish creating the FC Adapter Policy.

5. Click Add to create the vHBA-A.

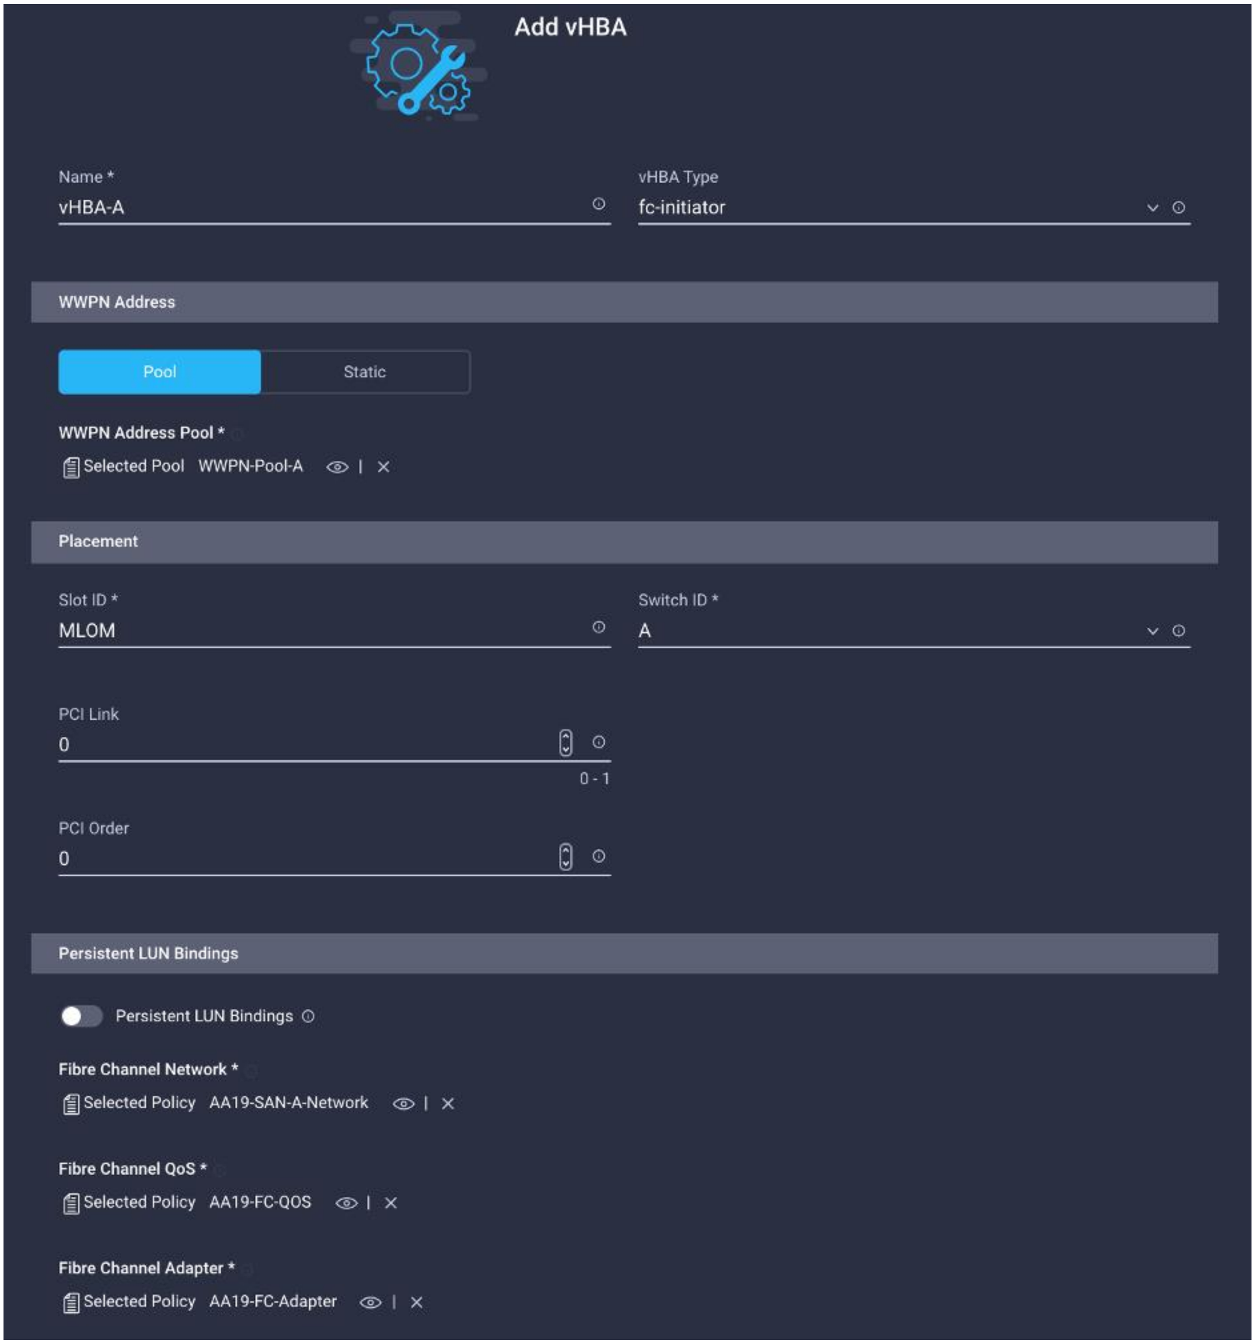

Add vHBA

Name * vHBA-A

vHBA Type fc-initiator

WWPN Address

WWPN Address Pool *

Selected Pool WWPN-Pool-A

Placement

Slot ID * MLOM

Switch ID * A

PCI Link 0 (0 - 1)

PCI Order 0

Persistent LUN Bindings

Fibre Channel Network *

Selected Policy AA19-SAN-A-Network

Fibre Channel QoS *

Selected Policy AA19-FC-QOS

Fibre Channel Adapter *

Selected Policy AA19-FC-Adapter

Create vHBA - SAN B

Repeat the steps in section Fibre Channel Adapter Policy to add vHBA-B for SAN-B. Select Switch ID **B** for this vHBA. The WWPN Pool and Fibre Channel Network policy (VSAN) for this vHBA are unique while the Fibre Channel QoS and Fibre Channel Adapter policies defined above for vHBA-A will be re-used.

1. To manually select the vHBA placement, select Manual vHBAs Placement.
2. Click Add vHBA.

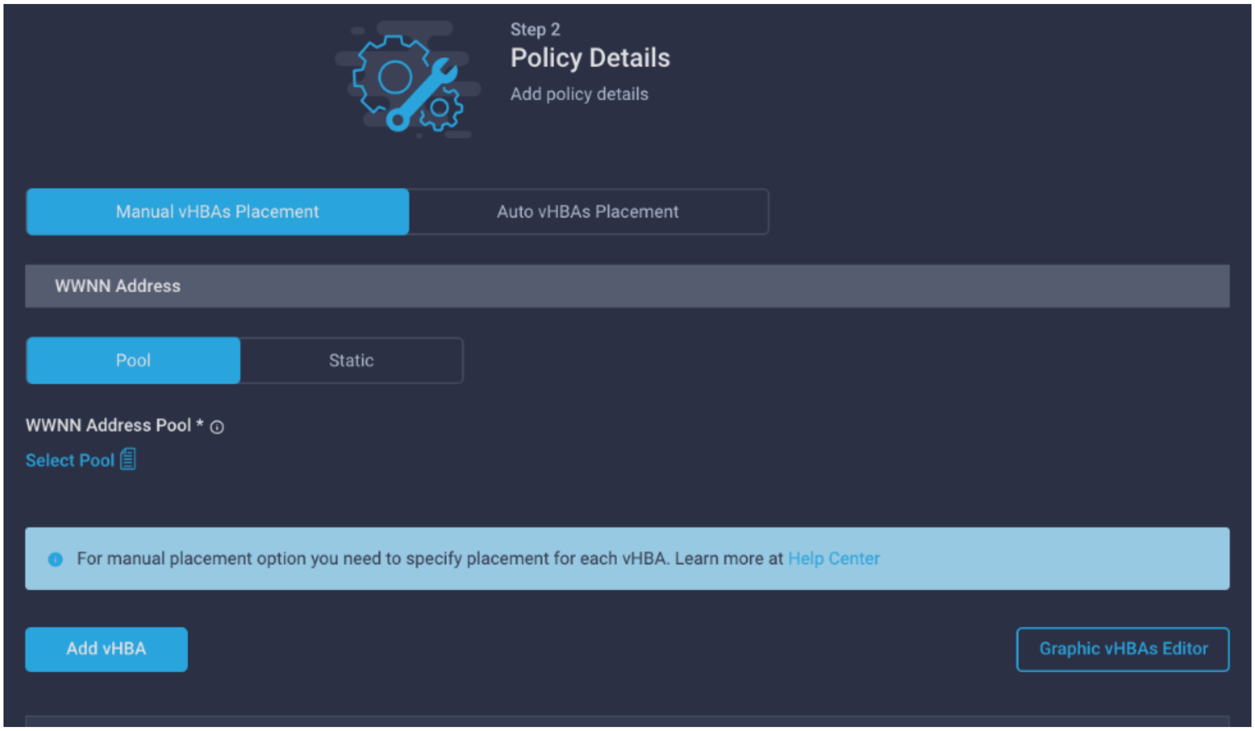

Step 2

Policy Details

Add policy details

WWNN Address

WWNN Address Pool *

Select Pool

For manual placement option you need to specify placement for each vHBA. Learn more at Help Center

3. Configure the following parameters:
   a. Provide the name of vNIC (for example, vHBA-B).
   b. Enter PCIe slod ID as MLOM where the VIC adapter is installed.
   c. Enter the PCI Link as 0
   d. Select the Switch ID as A, the fabric port to which the vHBAs are associated.
   e. Enter PCI Order a 1, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter
   f. In the vHBA Type, select fc-initiator from the drop-down list.
   g. Select Switch ID B from the drop-down list.

The WWPN and Fibre Channel Network information used in this validation is shown here for your reference:

Figure 6. **WWPN-Pool-B**

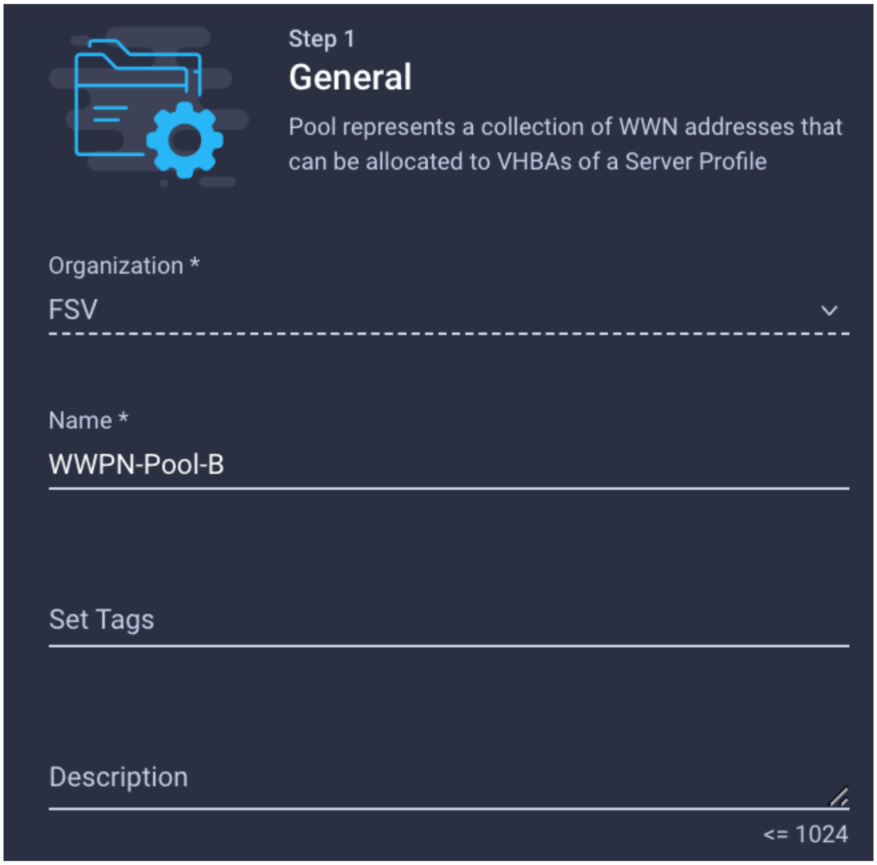

Step 1

General

Pool represents a collection of WWN addresses that can be allocated to VHBAs of a Server Profile

Organization *

FSV

Name *

WWPN-Pool-B

Set Tags

Description

<= 1024

4. Recommended prefix for WWPN addresses is 20:00:00:25:B5:xx:xx:xx. As a best practice, in FlashStack some extra information is always coded into the WWPN address pool for ease of troubleshooting. For example, in the figure below, 19 is the Rack ID and 0B signifies SAN B.

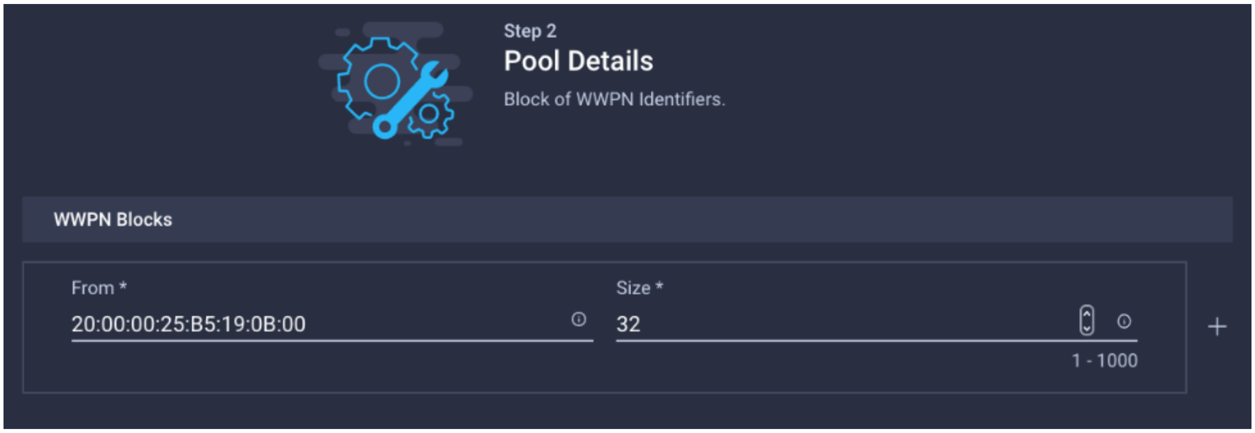

Step 2

Pool Details

Block of WWPN Identifiers.

WWPN Blocks

From *

20:00:00:25:B5:19:0B:00

Size *

32

1 - 1000

Fibre Channel Network Policy – SAN B

1. For the Fibre Channel Network policy, VSAN 102 will be used for vHBA-B.

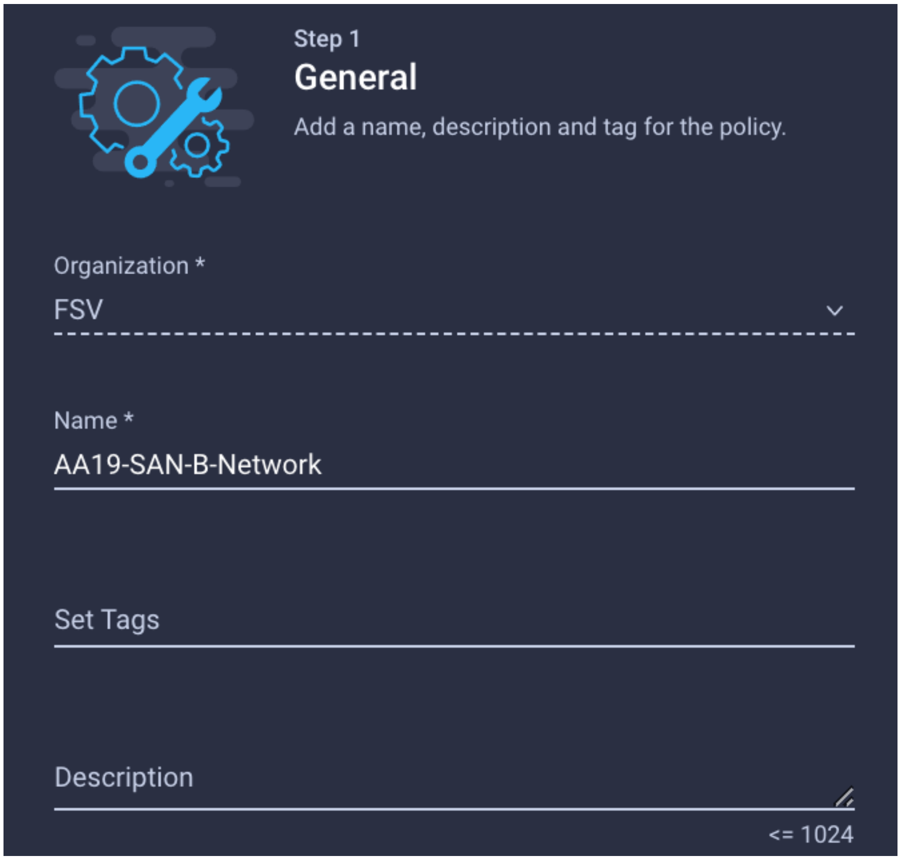

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-SAN-B-Network

Set Tags

Description

<= 1024

2. For the scope, select UCS Server (FI-Attached) and enter the VSAN information (for example, 102) under Default VLAN.

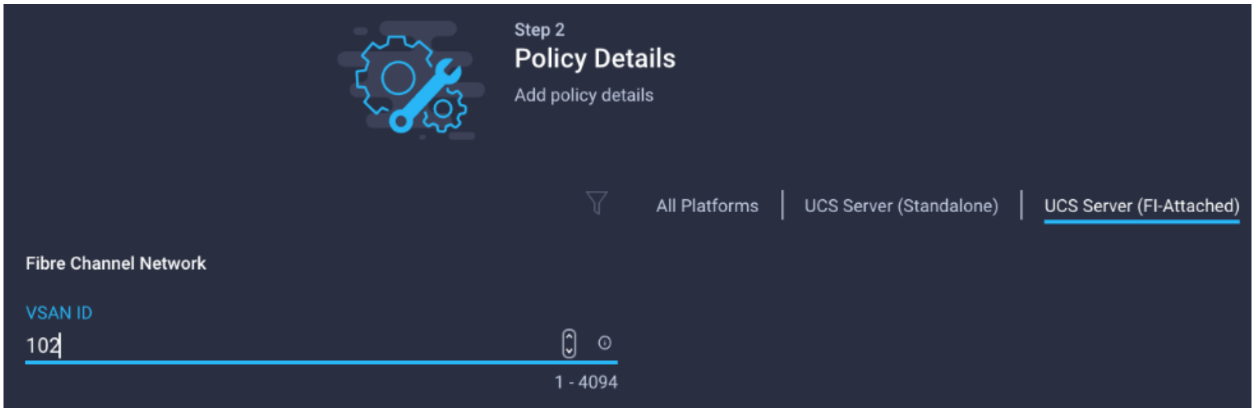

Step 2

Policy Details

Add policy details

Fibre Channel Network

VSAN ID

102

1 - 4094

3. Click Add to create the vHBA-B.

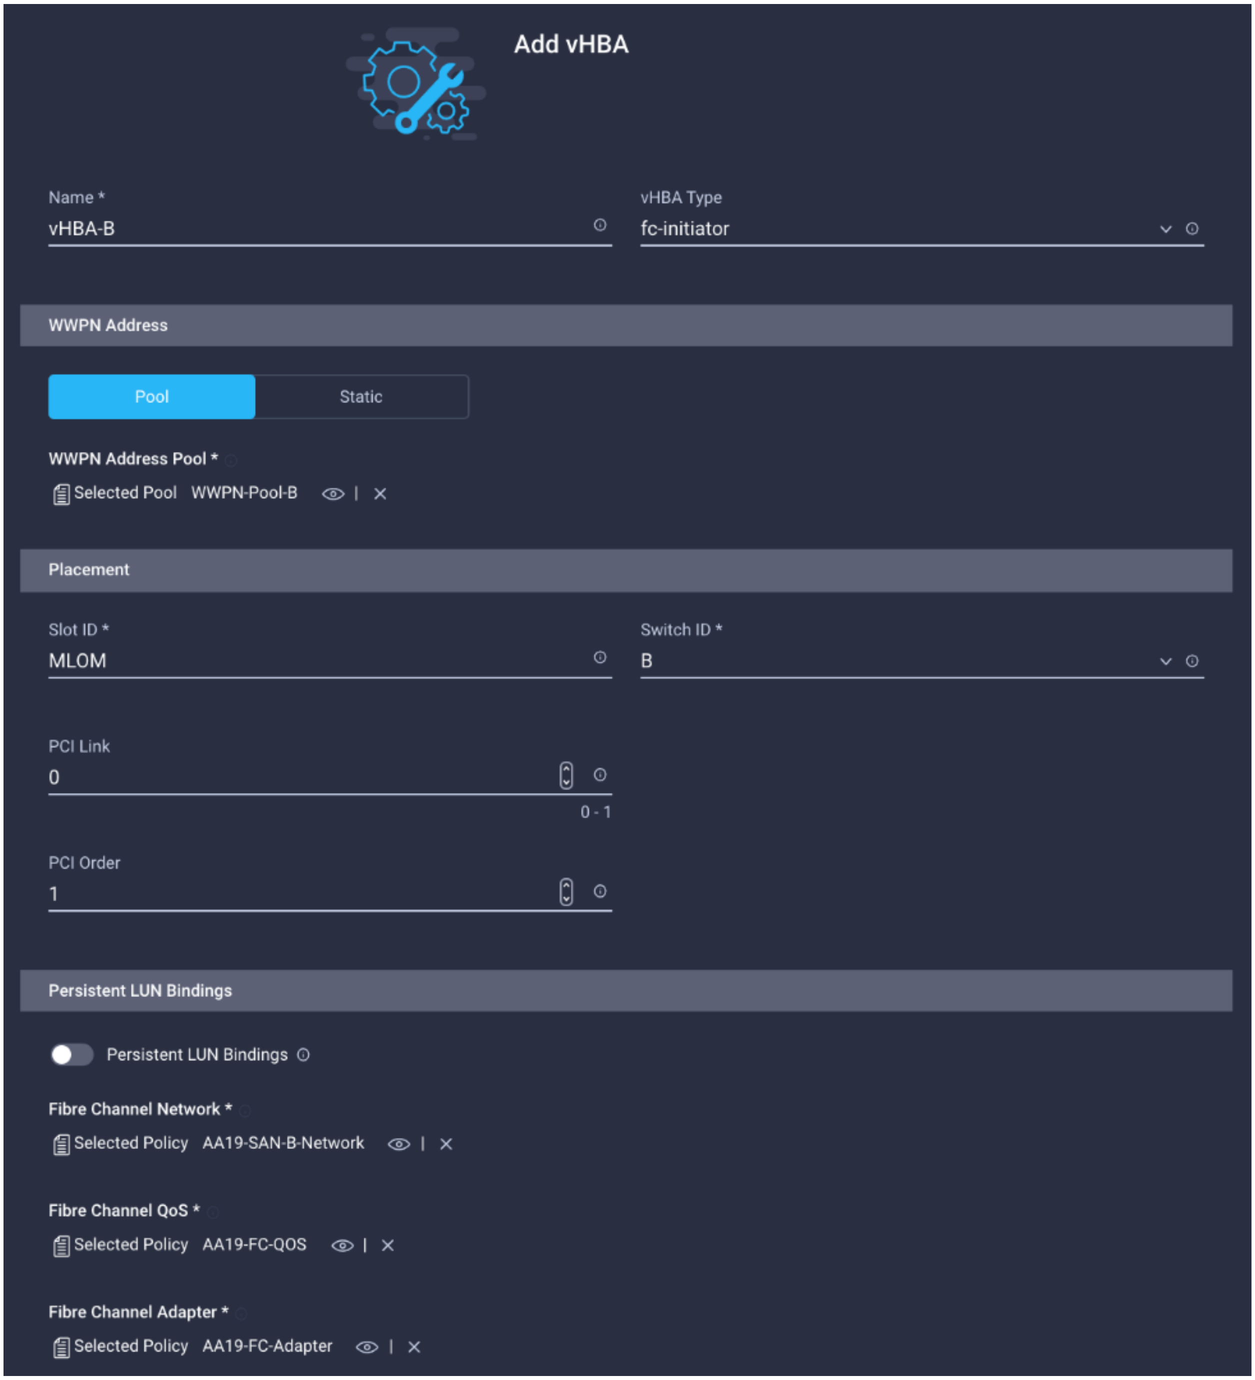

Add vHBA

Name * vHBA-B

vHBA Type fc-initiator

WWPN Address

Pool | Static

WWPN Address Pool *

Selected Pool WWPN-Pool-B

Placement

Slot ID * MLOM

Switch ID * B

PCI Link 0

0 - 1

PCI Order 1

Persistent LUN Bindings

Persistent LUN Bindings

Fibre Channel Network *

Selected Policy AA19-SAN-B-Network

Fibre Channel QoS *

Selected Policy AA19-FC-QOS

Fibre Channel Adapter *

Selected Policy AA19-FC-Adapter

Create vHBA (fc-nvme) - SAN A

Repeat the steps in section Fibre Channel Adapter Policy to add FC-NVMe vHBA-A for SAN-A. Select Switch ID **A** for this vHBA.

Note: Skip creating the FC-NVMe initiators if FC-NVMe storage connectivity is not required.

1. Click Add vHBA.

2. To keep the vHBA placement manual, select Manual vHBAs Placement.

3. Configure the following parameters

a. Provide the name of vNIC (for example, vHBA-NVMe-A).
b. Enter PCIe slot ID as MLOM where the VIC adapter is installed.
c. Enter the PCI Link as 0
d. Select the Switch ID as A, the fabric port to which the vHBAs are associated.
e. Enter PCI Order a 2, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter
f. In the vHBA Type, select fc-nvme-initiator from the drop-down list.
g. Select Switch ID A from the drop-down list.
h. Select the previously create WWPN pool.
i. Select the Fibre Channel Network, Fibre Channel Adapter, and the Fibre Channel QoS policies that have been created earlier for vHBA A.

4. Click Create.

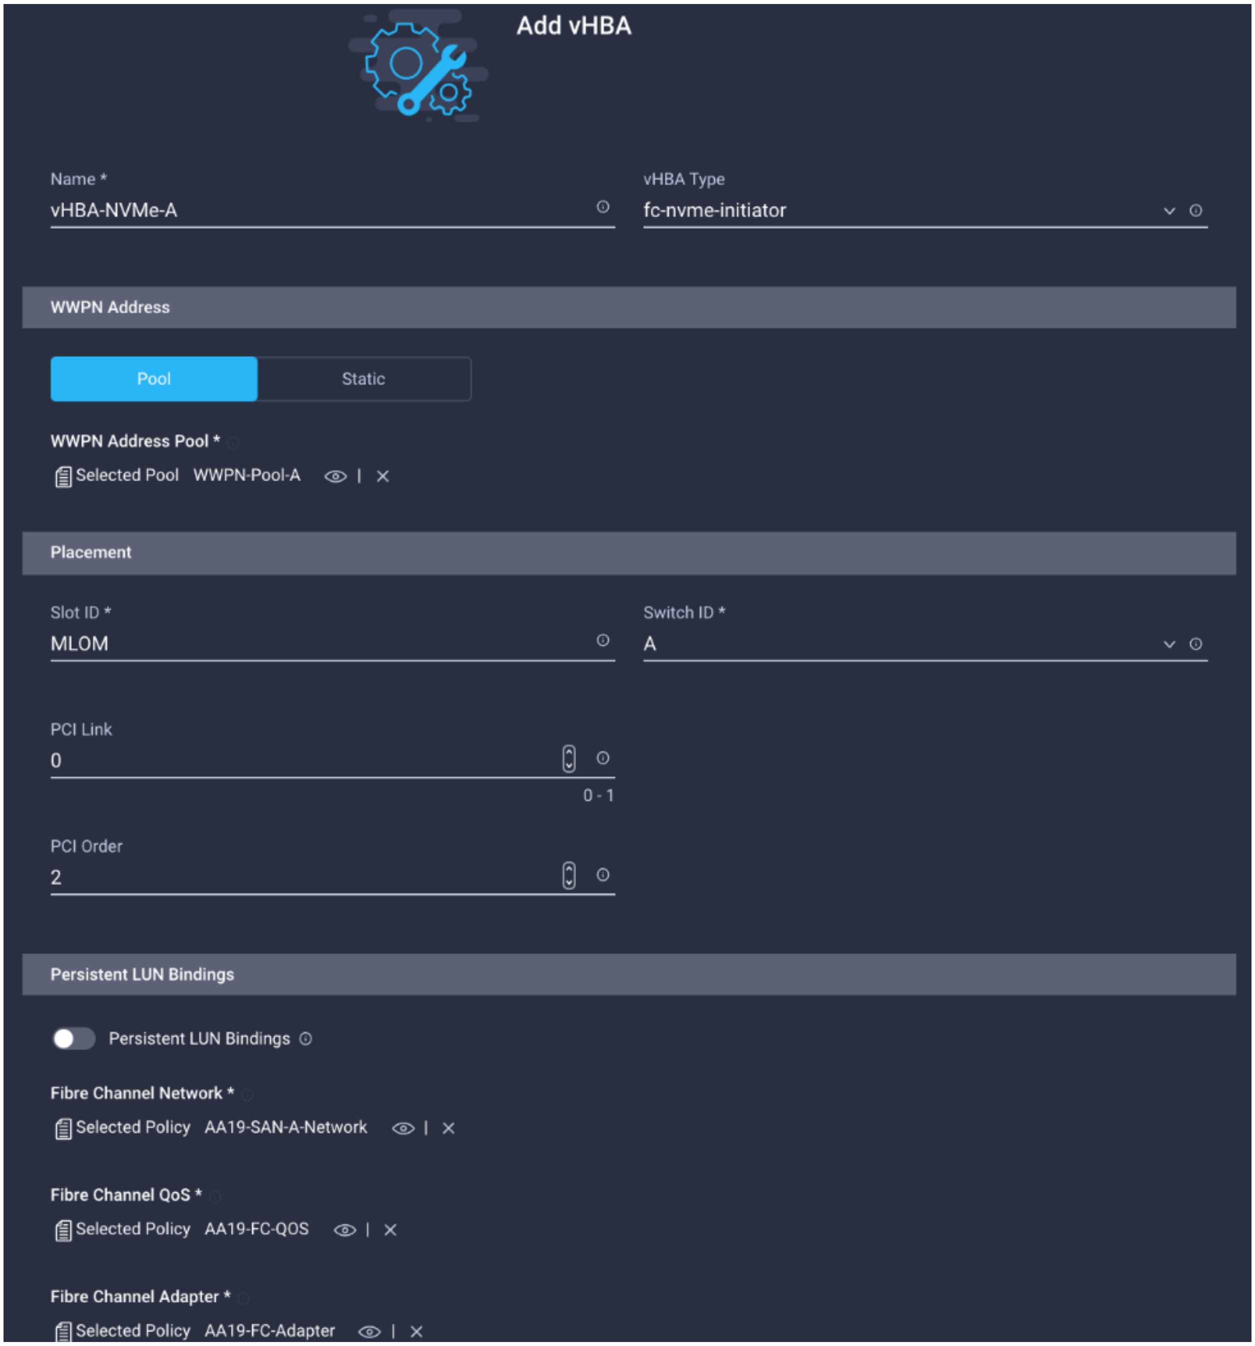

Add vHBA

Name * vHBA-NVMe-A

vHBA Type fc-nvme-initiator

WWPN Address

WWPN Address Pool *

Selected Pool WWPN-Pool-A

Placement

Slot ID * MLOM

Switch ID * A

PCI Link 0

PCI Order 2

Persistent LUN Bindings

Fibre Channel Network *

Selected Policy AA19-SAN-A-Network

Fibre Channel QoS *

Selected Policy AA19-FC-QOS

Fibre Channel Adapter *

Selected Policy AA19-FC-Adapter

Create vHBA (fc-nvme) - SAN B

Repeat the steps in section Fibre Channel Adapter Policy to add FC-NVMe vHBA-A for SAN-A. Select Switch ID B for this vHBA.

1. Click Add vHBA.

2. To keep the vHBA placement manual, select Manual vHBAs Placement.

3. Configure the following parameters

   a. Provide the name of vNIC (for example, vHBA-NVMe-B).

b. Enter PCIe slot ID as MLOM where the VIC adapter is installed.

c. Enter the PCI Link as 0

d. Select the Switch ID as A, the fabric port to which the vHBAs are associated.

e. Enter PCI Order a 3, this is the order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter

f. In the vHBA Type, select fc-nvme-initiator from the drop-down list.

g. Select Switch ID **B** from the drop-down list.

h. Select the previously create WWPN pool.

i. Select the Fibre Channel Network, Fibre Channel Adapter, and the Fibre Channel QoS policies that have been created earlier for vHBA A.

4. Click Create.

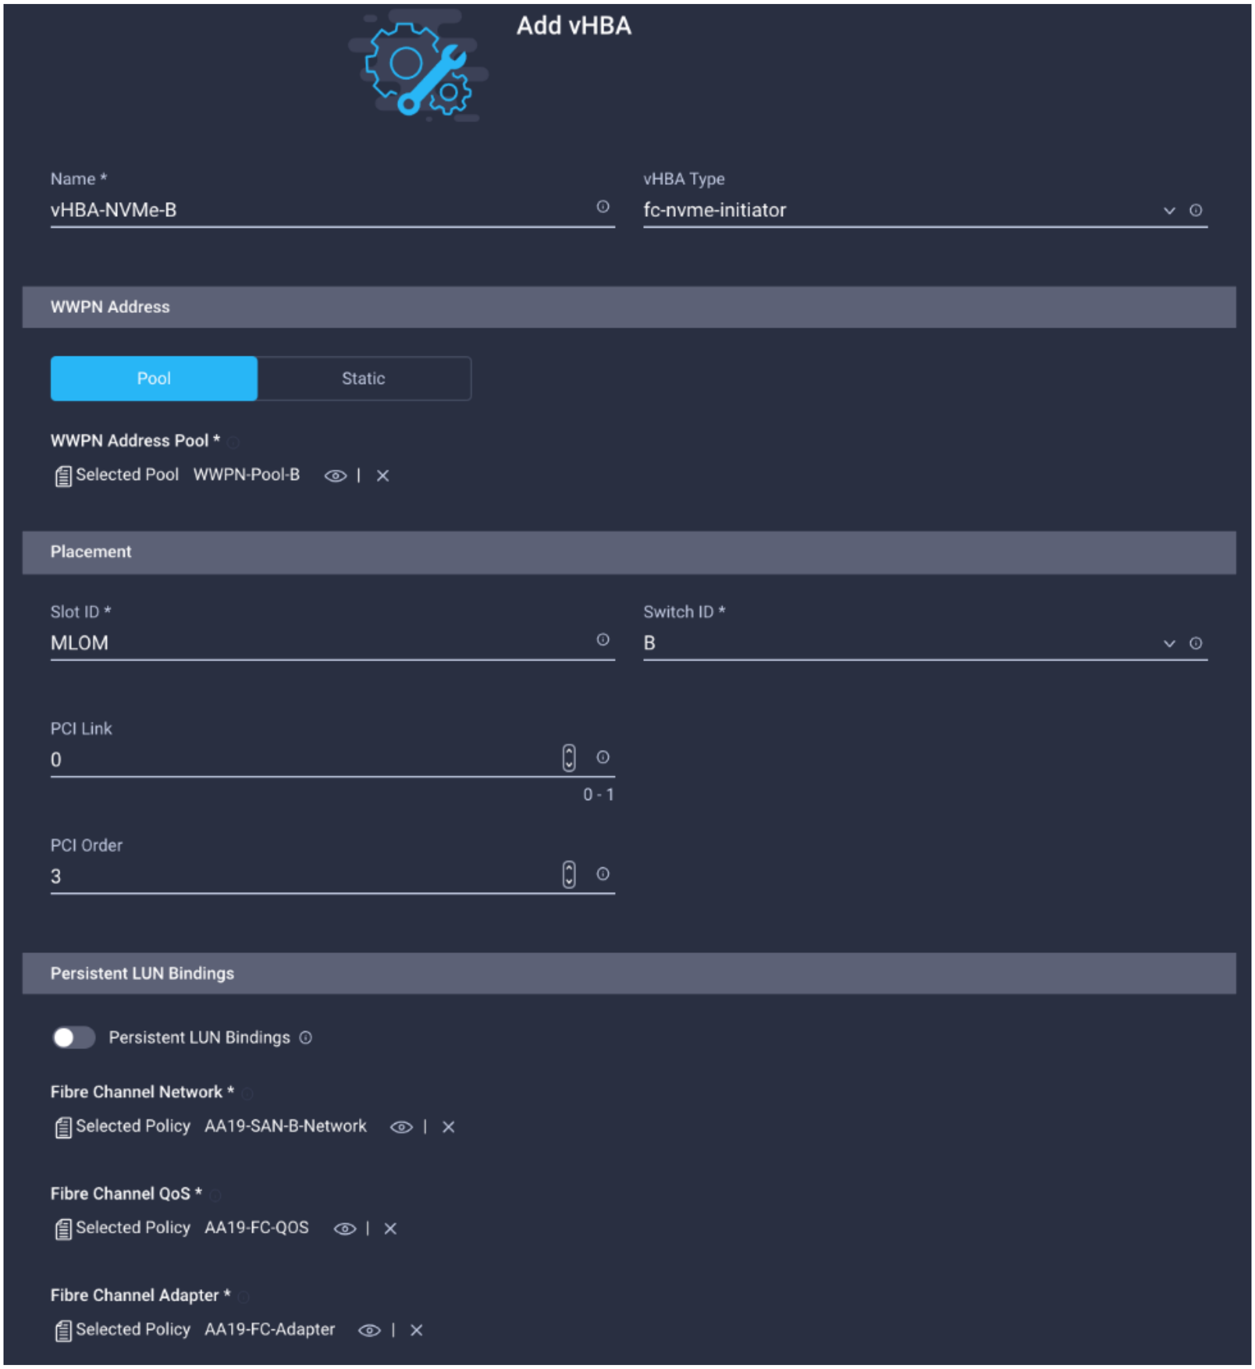

Add vHBA

Name * vHBA-NVMe-B

vHBA Type fc-nvme-initiator

WWPN Address

Pool | Static

WWPN Address Pool *

Selected Pool WWPN-Pool-B

Placement

Slot ID * MLOM

Switch ID * B

PCI Link 0

0 - 1

PCI Order 3

Persistent LUN Bindings

Persistent LUN Bindings

Fibre Channel Network *

Selected Policy AA19-SAN-B-Network

Fibre Channel QoS *

Selected Policy AA19-FC-QOS

Fibre Channel Adapter *

Selected Policy AA19-FC-Adapter

5. Verify all the vHBAs are added and in intended order.

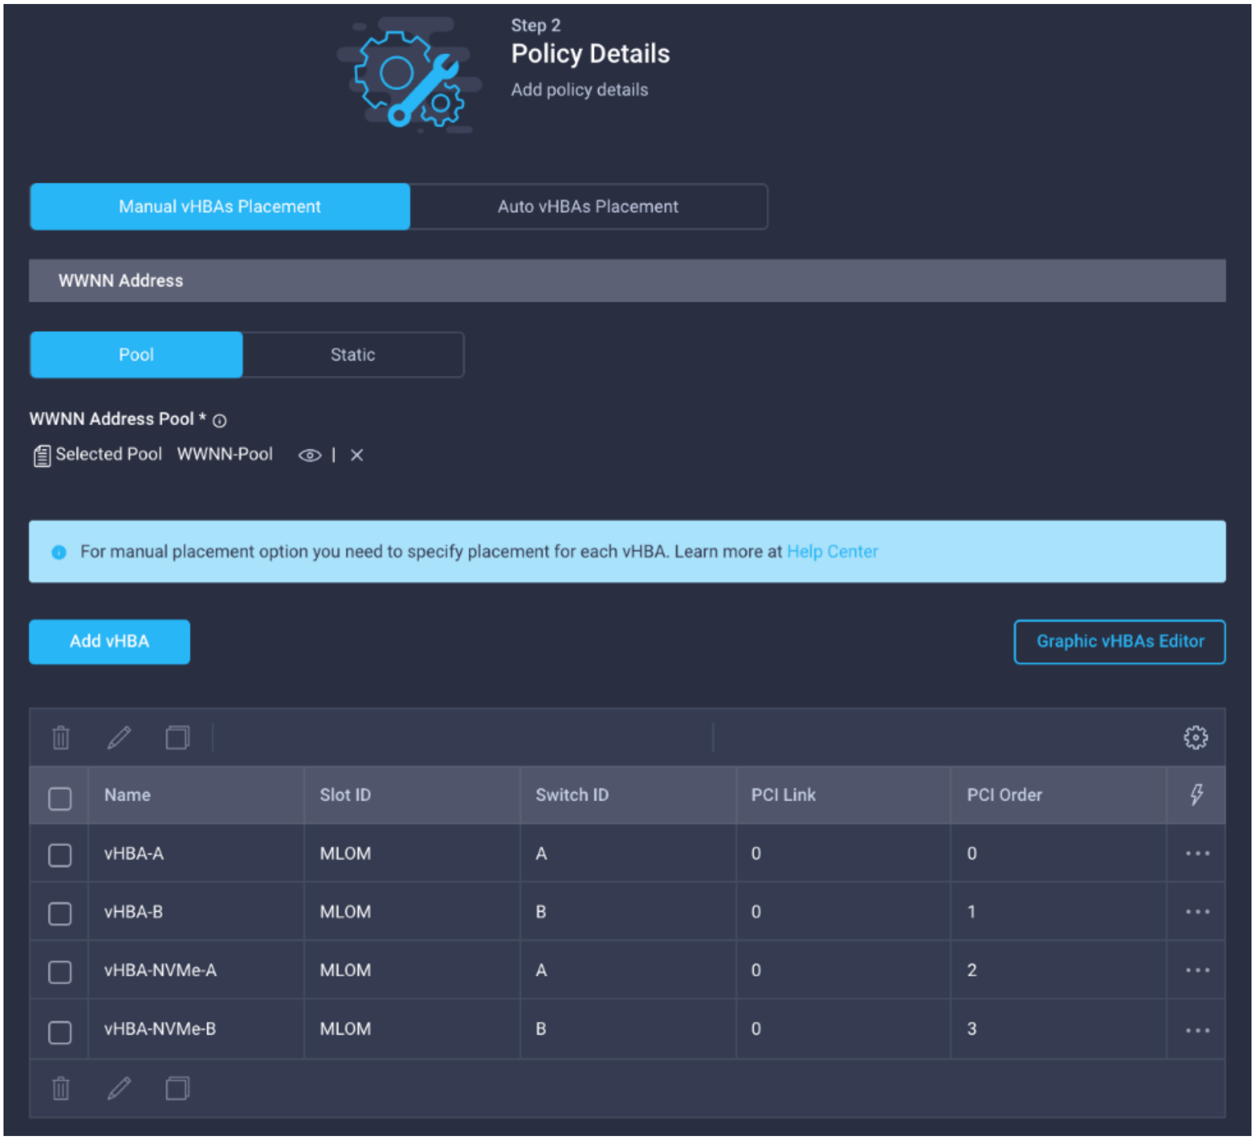
Step 2

Policy Details

Add policy details

WWNN Address

WWNN Address Pool *

Selected Pool WWNN-Pool

For manual placement option you need to specify placement for each vHBA. Learn more at Help Center

□	Name	Slot ID	Switch ID	PCI Link	PCI Order
□	vHBA-A	MLOM	A	0	0
□	vHBA-B	MLOM	B	0	1
□	vHBA-NVMe-A	MLOM	A	0	2
□	vHBA-NVMe-B	MLOM	B	0	3

6. Click Next.

Step 6 – Summary

On the summary screen, you can verify what policies are mapped to various settings and the status of the Server Profile Template.

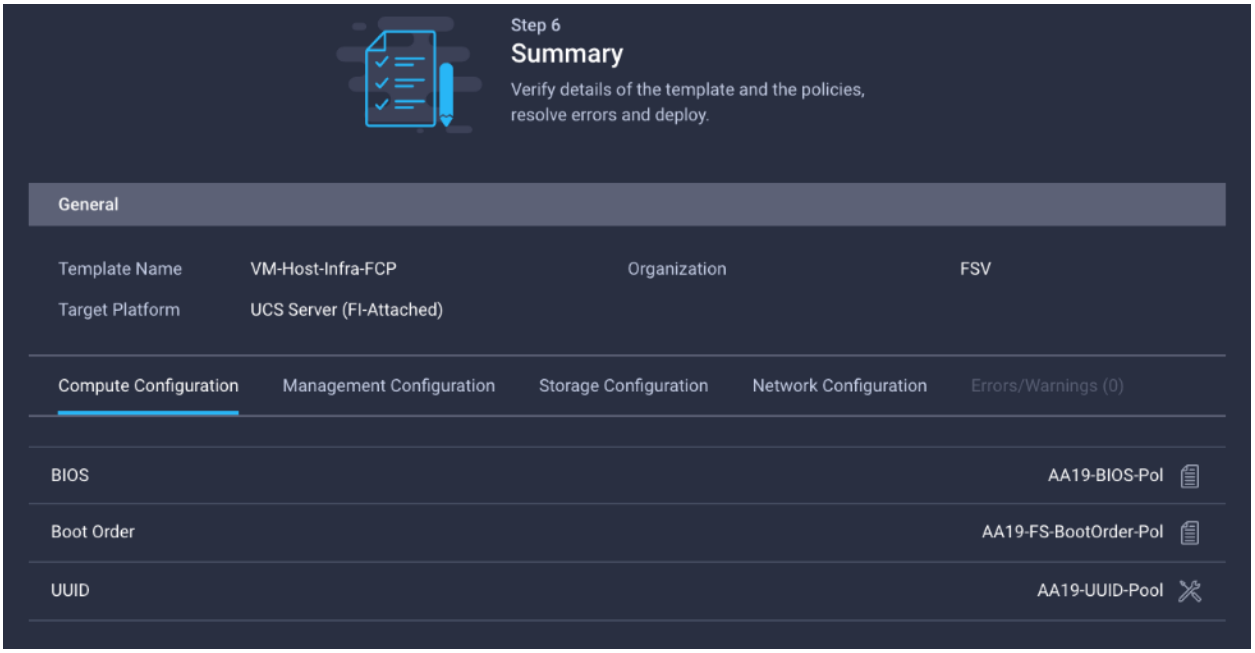
Step 6

Summary

Verify details of the template and the policies, resolve errors and deploy.

General

Template Name	VM-Host-Infra-FCP	Organization	FSV
Target Platform	UCS Server (FI-Attached)		

BIOS	AA19-BIOS-Pol
Boot Order	AA19-FS-BootOrder-Pol
UUID	AA19-UUID-Pool

Derive Server Profiles

To derive one or many server profiles from the configured template, follow these steps:

1. From the Server profile template Summary screen, click Derive Profiles.

Note: This action can also be performed later by navigating to **Templates**, clicking **"..."** next to the template name and selecting **Derive Profiles**.

2. Under the Server Assignment, select Assign Now and pick Cisco UCS X210c M6 servers. Customers can select one or more servers depending on the number of profiles to be deployed.

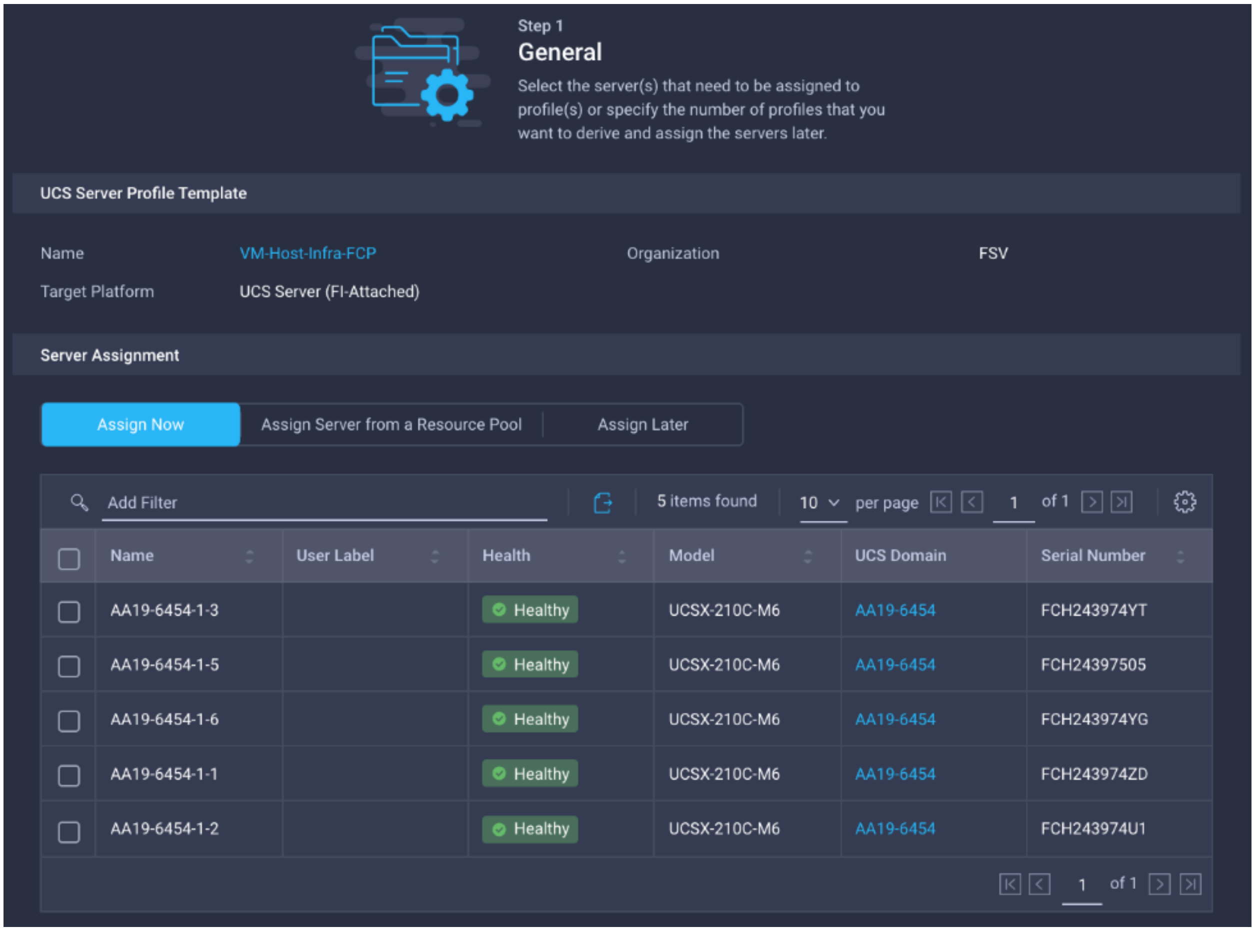

Step 1

General

Select the server(s) that need to be assigned to profile(s) or specify the number of profiles that you want to derive and assign the servers later.

UCS Server Profile Template

Name	VM-Host-Infra-FCP	Organization	FSV
Target Platform	UCS Server (FI-Attached)		

Server Assignment

□	Name	User Label	Health	Model	UCS Domain	Serial Number
□	AA19-6454-1-3		Healthy	UCSX-210C-M6	AA19-6454	FCH243974YT
□	AA19-6454-1-5		Healthy	UCSX-210C-M6	AA19-6454	FCH24397505
□	AA19-6454-1-6		Healthy	UCSX-210C-M6	AA19-6454	FCH243974YG
□	AA19-6454-1-1		Healthy	UCSX-210C-M6	AA19-6454	FCH243974ZD
□	AA19-6454-1-2		Healthy	UCSX-210C-M6	AA19-6454	FCH243974U1

3. Select the servers that you want the server profiles to be assigned to which will be used as FC nodes in the ESXi cluster.

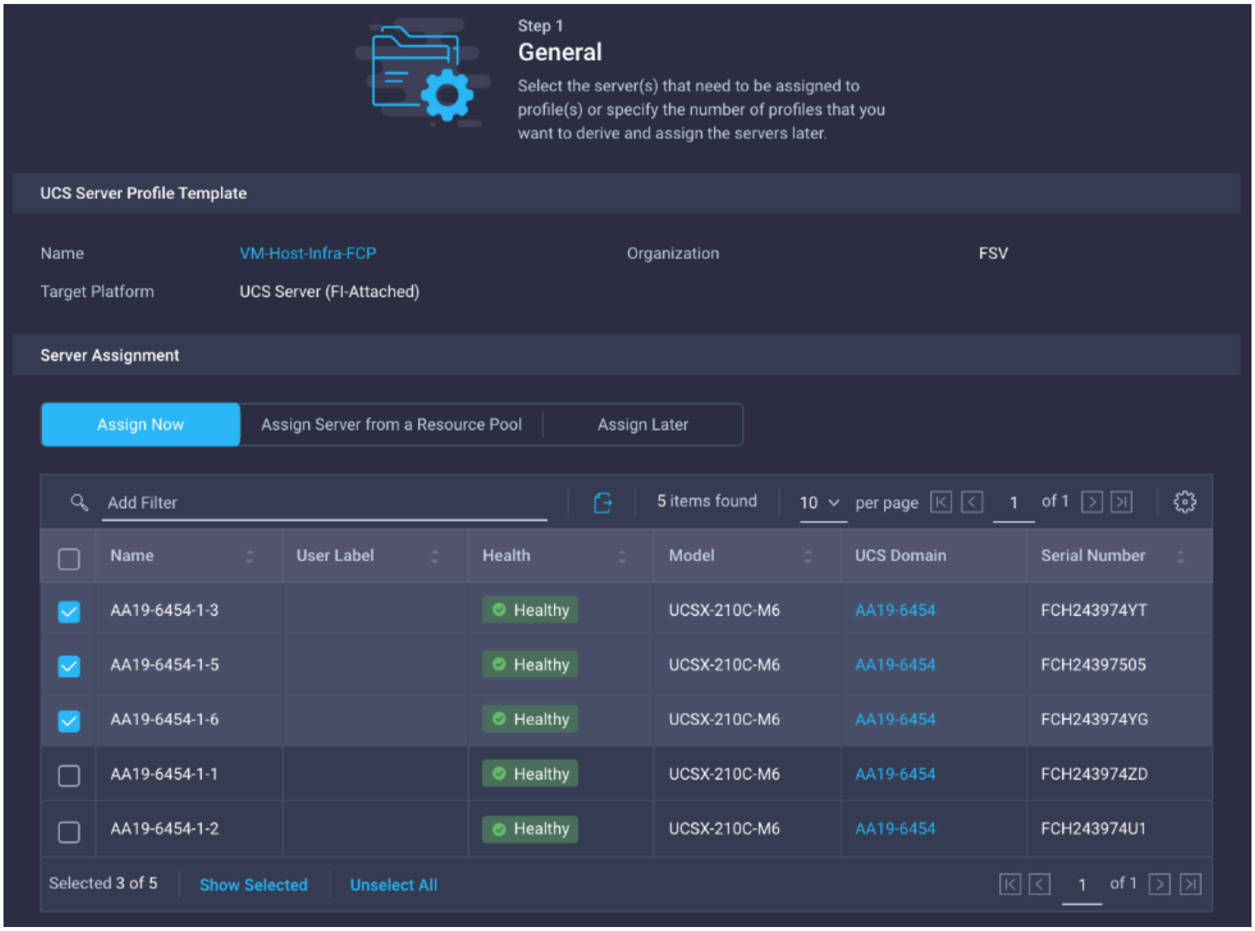

Step 1

General

Select the server(s) that need to be assigned to profile(s) or specify the number of profiles that you want to derive and assign the servers later.

UCS Server Profile Template

Name	VM-Host-Infra-FCP	Organization	FSV
Target Platform	UCS Server (FI-Attached)		

Server Assignment

□	Name	User Label	Health	Model	UCS Domain	Serial Number
☑	AA19-6454-1-3		Healthy	UCSX-210C-M6	AA19-6454	FCH243974YT
☑	AA19-6454-1-5		Healthy	UCSX-210C-M6	AA19-6454	FCH24397505
☑	AA19-6454-1-6		Healthy	UCSX-210C-M6	AA19-6454	FCH243974YG
□	AA19-6454-1-1		Healthy	UCSX-210C-M6	AA19-6454	FCH243974ZD
□	AA19-6454-1-2		Healthy	UCSX-210C-M6	AA19-6454	FCH243974U1

Selected 3 of 5

4. Click Next.

5. Intersight will fill in default information for the number of servers selected (3 in this case).

6. Change the server profile names (for ex: VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03), prefix and number can be changed as needed.

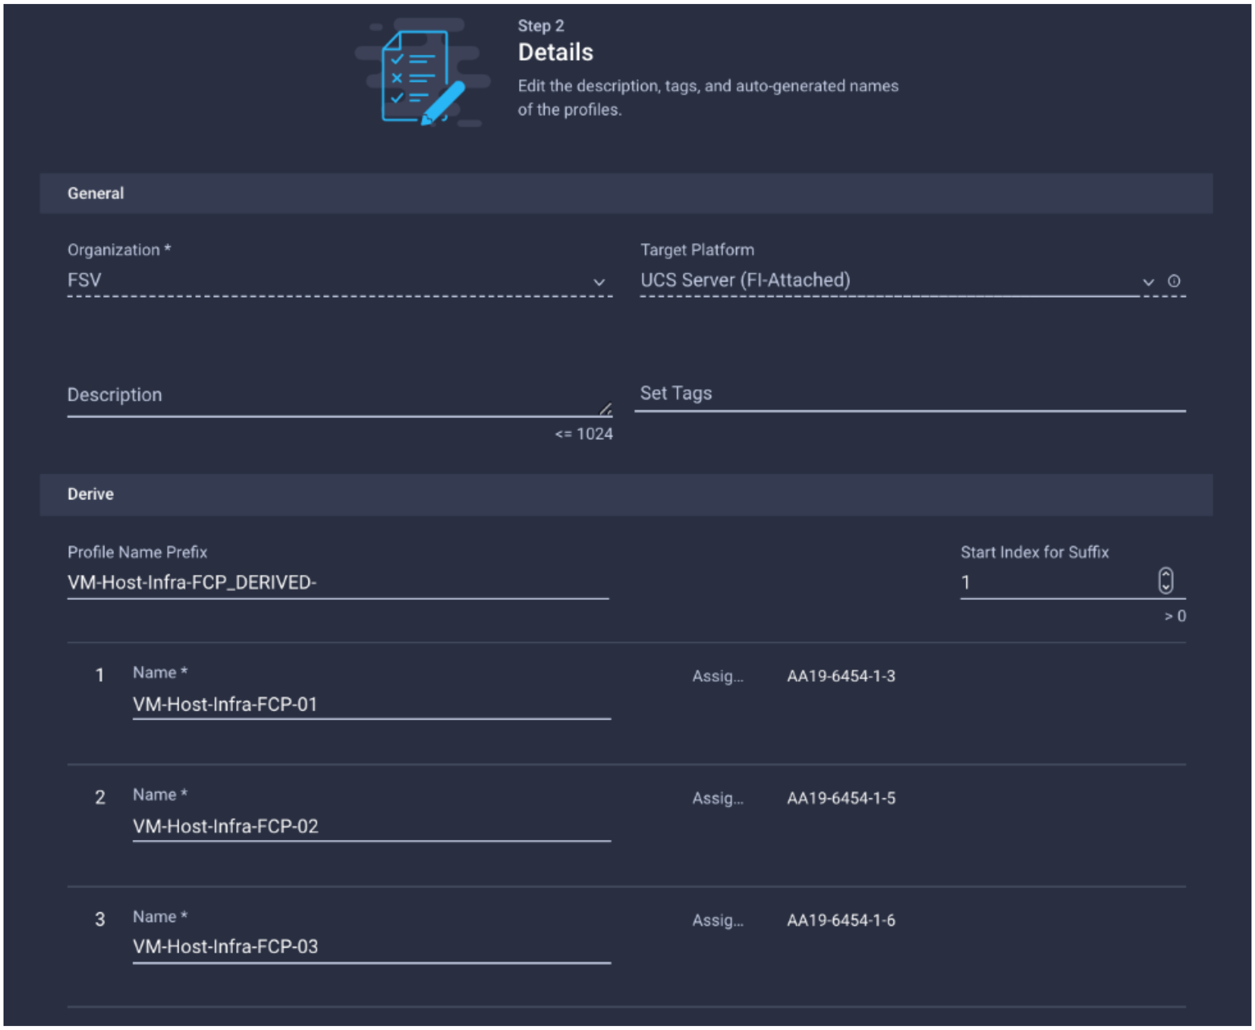

Step 2

Details

Edit the description, tags, and auto-generated names of the profiles.

General

Organization * FSV

Target Platform UCS Server (FI-Attached)

Description

<= 1024

Set Tags

Derive

Profile Name Prefix VM-Host-Infra-FCP_DERIVED-

Start Index for Suffix 1

> 0

1	Name * VM-Host-Infra-FCP-01	Assig...	AA19-6454-1-3
2	Name * VM-Host-Infra-FCP-02	Assig...	AA19-6454-1-5
3	Name * VM-Host-Infra-FCP-03	Assig...	AA19-6454-1-6

7. On the summary screen, you can verify what profiles are assigned to servers.

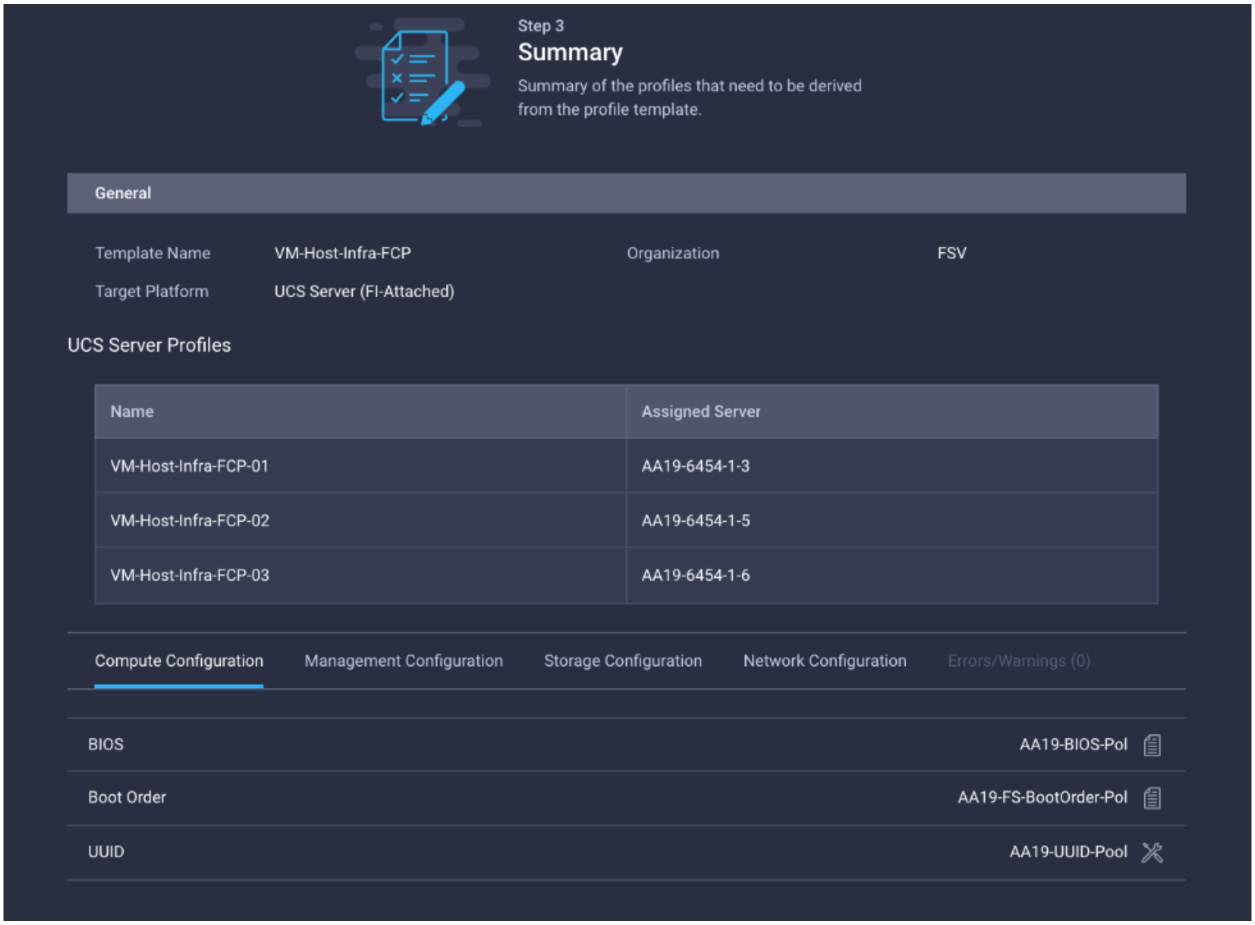

Step 3

Summary

Summary of the profiles that need to be derived from the profile template.

General

Template Name	VM-Host-Infra-FCP	Organization	FSV
Target Platform	UCS Server (FI-Attached)		

UCS Server Profiles

Name	Assigned Server
VM-Host-Infra-FCP-01	AA19-6454-1-3
VM-Host-Infra-FCP-02	AA19-6454-1-5
VM-Host-Infra-FCP-03	AA19-6454-1-6

BIOS	AA19-BIOS-Pol
Boot Order	AA19-FS-BootOrder-Pol
UUID	AA19-UUID-Pool

8. Click Next.

9. Verify the information and click Derive to create the Server Profiles.

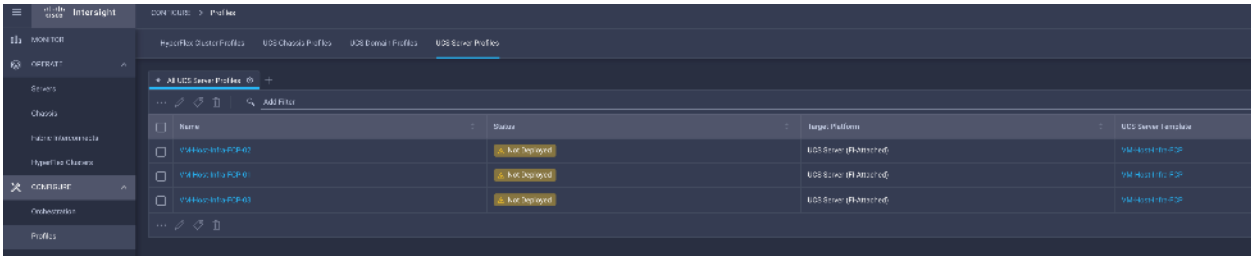

☐	Name	Status	Target Platform	UCS Server Template
☐	VM-Host-Infra-FCP-02	Not Deployed	UCS Server (FI-Attached)	VM-Host-Infra-FCP
☐	VM-Host-Infra-FCP-01	Not Deployed	UCS Server (FI-Attached)	VM-Host-Infra-FCP
☐	VM-Host-Infra-FCP-03	Not Deployed	UCS Server (FI-Attached)	VM-Host-Infra-FCP

10. Click ... on the right-hand side of the derived server profiles and click deploy.

11. When the Server Profiles are deployed successfully, they will appear under the Server Profiles with the status of OK.

Gather Necessary Information

After the Cisco UCS server profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlashStack deployment, specific information must be gathered from each Cisco UCS server and from the Pure FlashArray controllers.

Table 20. WWPNs from Pure FlashArray//X R3 Storage

FlashArray	Adapter	MDS Switch	Target: WWPN
BB08-FlashArray//X-R3	CT0.FC0	Fabric A	<CT0.FC0-wwpn>
	CT0.FC2	Fabric B	<CT0.FC2-wwpn>
	CT1.FC0	Fabric A	<CT1.FC0-wwpn>
	CT1.FC2	Fabric B	<CT1.FC2-wwpn>

Table 21. WWPNs for Cisco UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
VM-Host-Infra-FCP-01	Fabric A	<vm-host-infra-fcp-01-wwpna>
	Fabric B	<vm-host-infra-fcp-01-wwpnb>
VM-Host-Infra-FCP-02	Fabric A	<vm-host-infra-fcp-02-wwpna>
	Fabric B	<vm-host-infra-fcp-02-wwpnb>
VM-Host-Infra-FCP-03	Fabric A	<vm-host-infra-fcp-03-wwpna>
	Fabric B	<vm-host-infra-fcp-03-wwpnb>

Note: To obtain the FC vHBA WWPN information in Cisco Intersight. Log into Intersight Portal. Go to CONFIGURE -> Profiles and select the Server Profiles just deployed.

SAN Switch Configuration

This section explains how to configure the Cisco MDS 9000s for use in a FlashStack environment.

Note: Follow the steps precisely because failure to do so could result in an improper configuration.

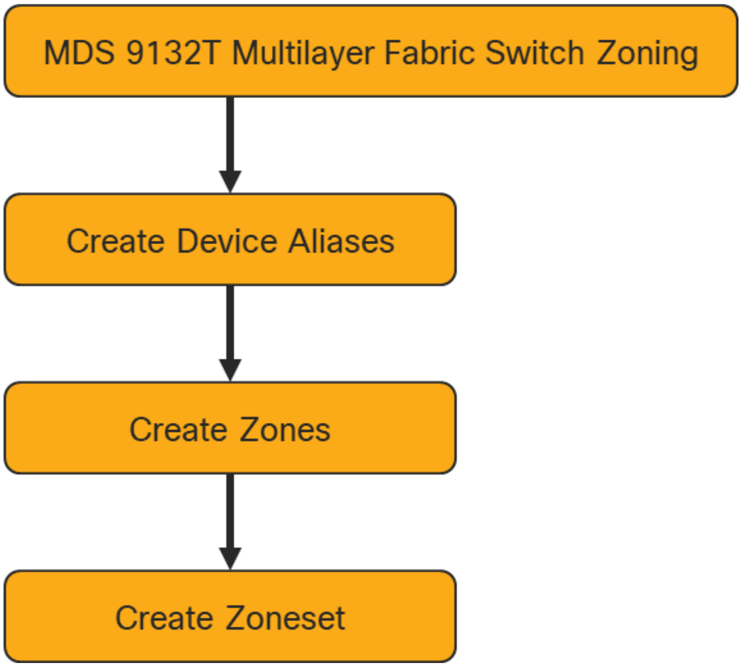

Note: If directly connecting storage to the Cisco UCS fabric interconnects, skip this section.

Physical Connectivity

Follow the physical connectivity guidelines for FlashStack as explained in section FlashStack Cabling.

FlashStack Cisco MDS Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlashStack environment. This procedure assumes you are using the Cisco MDS 9132T with NX-OS 8.4(2c).

Cisco MDS 9132T A

To set up the initial configuration for the Cisco MDS A switch, <mds-A-hostname>, follow these steps:

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

```
         ---- System Admin Account Setup ----


Do you want to enforce secure password standard (yes/no) [y]: Enter

Enter the password for "admin": <password>
Confirm the password for "admin": <password>
```

```
Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name : <mds-A-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address : <mds-A-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-A-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway : <mds-A-mgmt0-gw>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Configure congestion/no_credit drop for fc interfaces? (yes/no)     [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter

Configure timezone? (yes/no) [n]: Enter

Configure summertime? (yes/no) [n]: Enter

Configure the ntp server? (yes/no) [n]: yes

NTP server IPv4 address : <nexus-A-mgmt0-ip>

Configure default switchport interface state (shut/noshut) [shut]: Enter

Configure default switchport trunk mode (on/off/auto) [on]: auto

Configure default switchport port mode F (yes/no) [n]: yes

Configure default zone policy (permit/deny) [deny]: Enter

Enable full zoneset distribution? (yes/no) [n]: Enter

Configure default zone mode (basic/enhanced) [basic]: Enter
```

2. Review the configuration.

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

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

Cisco MDS 9132T B

To set up the initial configuration for the Cisco MDS B switch, <mds-B-hostname>, follow these steps:

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

```
          ---- System Admin Account Setup ----


Do you want to enforce secure password standard (yes/no) [y]: Enter

Enter the password for "admin": <password>
Confirm the password for "admin": <password>

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name : <mds-B-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address : <mds-B-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-B-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway : <mds-B-mgmt0-gw>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Configure congestion/no_credit drop for fc interfaces? (yes/no)     [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter

Configure timezone? (yes/no) [n]: Enter

Configure summertime? (yes/no) [n]: Enter

Configure the ntp server? (yes/no) [n]: yes

NTP server IPv4 address : <nexus-A-mgmt0-ip>

Configure default switchport interface state (shut/noshut) [shut]: Enter

Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes

Configure default zone policy (permit/deny) [deny]: Enter

Enable full zoneset distribution? (yes/no) [n]: Enter

Configure default zone mode (basic/enhanced) [basic]: Enter
```

2. Review the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter
```

FlashStack Cisco MDS Switch Configuration

Enable Licenses

Cisco MDS 9132T A and Cisco MDS 9132T B

To enable the correct features on the Cisco MDS switches, follow these steps:

1. Log in as admin.

2. Run the following commands:

```
configure terminal
feature npiv
feature fport-channel-trunk
```

Add Second NTP Server and Local Time Configuration

Cisco MDS 9132T A and Cisco MDS 9132T B

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

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

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

Note: It is important to configure the local time so that logging time alignment, any backup schedules, and SAN Analytics forwarding are correct. For more information on configuring the timezone and daylight savings time or summer time, please see Cisco MDS 9000 Series Fundamentals Configuration Guide, Release 8.x. Sample clock commands for the United States Eastern timezone are:
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60

Configure Individual Ports

Cisco MDS 9132T A

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

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

```
interface fc1/1
switchport description BB08-X50R3-ct0fc0
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description BB08-X50R3-ct1fc0
switchport speed 32000
switchport trunk mode off
no shutdown
```

```
exit

interface fc1/1
switchport description BB08-X50R3-ct0fc1
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description BB08-X50R3-ct1fc1
switchport speed 32000
switchport trunk mode off
no shutdown
exit


interface fc1/5
  switchport description BB08-6454-A:fc1/1
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown

interface fc1/6
  switchport description BB08-6454-A:fc1/2
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown

interface fc1/7
  switchport description BB08-6454-A:fc1/3
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown

interface fc1/8
  switchport description BB08-6454-A:fc1/4
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown


interface port-channel15
  switchport mode F
  switchport trunk allowed vsan 100
  switchport description BB08-6454-A
  switchport speed 32000
  switchport rate-mode dedicated
  switchport trunk mode auto
```

If VSAN trunking is not being used between the Cisco UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-a-id>” for interface port-channel15. Note also that the default setting of switchport trunk mode auto is being used for the port channel.

Cisco MDS 9132T B

To configure individual ports and port-channels for switch B, follow these steps:

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

```
interface fc1/1
switchport description BB08-X50R3-ct0fc2
switchport speed 32000
switchport trunk mode off
no shutdown
exit
```

```
interface fc1/2
switchport description BB08-X50R3-ct1fc2
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/3
switchport description BB08-X50R3-ct0fc3
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/4
switchport description BB08-X50R3-ct1fc3
switchport speed 32000
switchport trunk mode off
no shutdown
exit


interface fc1/5
  switchport description BB08-6454-B:fc1/1
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown

interface fc1/6
  switchport description BB08-6454-B:fc1/2
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown

interface fc1/7
  switchport description BB08-6454-B:fc1/3
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown

interface fc1/8
  switchport description BB08-6454-B:fc1/4
  switchport trunk mode auto
  port-license acquire
  channel-group 15 force
  no shutdown
```

If VSAN trunking is not being used between the Cisco UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-b-id>” for interface port-channel15. Note also that the default setting of the switchport trunk mode auto is being used for the port channel.

Create VSANs

Cisco MDS 9132T A

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

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

```
vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fc1/1
vsan <vsan-a-id> interface fc1/2
vsan <vsan-a-id> interface fc1/3
vsan <vsan-a-id> interface fc1/4
vsan <vsan-a-id> interface port-channel15
exit
```

Cisco MDS 9132T B

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

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

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fc1/1
vsan <vsan-b-id> interface fc1/2
vsan <vsan-b-id> interface fc1/3
vsan <vsan-b-id> interface fc1/4
vsan <vsan-b-id> interface port-channel15
exit
```

At this point, it may be necessary to go into Cisco UCS Manager and disable and enable the FC port-channel interfaces to get the port-channels to come up.

Create Device Aliases

Cisco MDS 9132T A

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

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

```
device-alias mode enhanced
device-alias database
device-alias name FlashArray-CT0FC0 pwwn 52:4a:93:77:de:d7:21:00
device-alias name FlashArray-CT1FC0 pwwn 52:4a:93:77:de:d7:21:10
device-alias name FlashArray-CT0FC1 pwwn 52:4a:93:77:de:d7:21:01
device-alias name FlashArray-CT1FC1 pwwn 52:4a:93:77:de:d7:21:11
device-alias name VM-Host-Infra-FCP-01-A pwwn 20:00:00:25:b5:a4:0a:00
device-alias name VM-Host-Infra-FCP-02-A pwwn 20:00:00:25:b5:a4:0a:01
device-alias name VM-Host-Infra-FCP-03-A pwwn 20:00:00:25:b5:a4:0a:02
device-alias name VM-Host-Infra-FC-NVMe-01-A pwwn 20:00:00:25:b5:a4:0a:03
device-alias name VM-Host-Infra-FC-NVMe-02-A pwwn 20:00:00:25:b5:a4:0a:04
device-alias name VM-Host-Infra-FC-NVMe-03-A pwwn 20:00:00:25:b5:a4:0a:05
device-alias commit
```

Cisco MDS 9132T B

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

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

```
device-alias mode enhanced
device-alias database
device-alias name FlashArray-CT0FC2 pwwn 52:4a:93:77:de:d7:21:02
device-alias name FlashArray-CT1FC2 pwwn 52:4a:93:77:de:d7:21:12
device-alias name FlashArray-CT0FC2 pwwn 52:4a:93:77:de:d7:21:03
device-alias name FlashArray-CT1FC2 pwwn 52:4a:93:77:de:d7:21:13
device-alias name VM-Host-Infra-FCP-01-B pwwn 20:00:00:25:b5:a4:0b:00
device-alias name VM-Host-Infra-FCP-02-B pwwn 20:00:00:25:b5:a4:0b:01
device-alias name VM-Host-Infra-FCP-03-B pwwn 20:00:00:25:b5:a4:0b:02
device-alias name VM-Host-Infra-FC-NVMe-01-B pwwn 20:00:00:25:b5:a4:0b:03
```

```
device-alias name VM-Host-Infra-FC-NVMe-02-B pwwn 20:00:00:25:b5:a4:0b:04
device-alias name VM-Host-Infra-FC-NVMe-03-B pwwn 20:00:00:25:b5:a4:0b:05
device-alias commit
```

Create Zones and Zoneset

Cisco MDS 9132T A

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

```
configure terminal
zone name Infra-VSI-Fabric-A vsan <vsan-a-id>
member device-alias FlashArray-CT0FC0 target
member device-alias FlashArray-CT1FC0 target
member device-alias Infra-Host-FCP-01-A init
member device-alias Infra-Host-FCP-02-A init
member device-alias Infra-Host-FCP-03-A init
exit
zone name Infra-VSI-NVMe-Fabric-A vsan <vsan-a-id>
member device-alias FlashArray-CT0FC1 target
member device-alias FlashArray-CT1FC1 target
member device-alias Infra-Host-FC-NVMe-01-A init
member device-alias Infra-Host-FC-NVMe-02-A init
member device-alias Infra-Host-FC-NVMe-03-A init
exit
zoneset name Fabric-A vsan <vsan-a-id>
member Infra-VSI-Fabric-A
member Infra-VSI-NVMe-Fabric-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
show zoneset active
copy r s
```

Note: Since Smart Zoning is enabled, a single zone for each storage protocol (FCP and FC-NVMe) is created with all host boot initiators and boot targets for the FlashArray//X R3 instead of creating a separate zone for each host with the host initiator and boot targets. If a new host is added, its boot initiator can simply be added to the single zone in each MDS switch and then the zoneset reactivated. If another FlashArray is added to the FlashStack with FC targets, a new zone can be added for that FlashArray.

Cisco MDS 9132T B

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

```
configure terminal

zone name Infra-VSI-Fabric-B vsan <vsan-b-id>
member device-alias FlashArray-CT0FC2 target
member device-alias FlashArray-CT1FC2 target
member device-alias Infra-Host-FCP-01-B init
member device-alias Infra-Host-FCP-02-B init
member device-alias Infra-Host-FCP-03-B init
exit
zone name Infra-VSI-NVMe-Fabric-B vsan <vsan-b-id>
member device-alias FlashArray-CT0FC3 target
member device-alias FlashArray-CT1FC3 target
member device-alias Infra-Host-FC-NVMe-01-B init
member device-alias Infra-Host-FC-NVMe-02-B init
member device-alias Infra-Host-FC-NVMe-03-B init
exit
zoneset name Fabric-B vsan <vsan-b-id>
member Infra-VSI-Fabric-B
member Infra-VSI-NVMe-Fabric-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
```

```
show zoneset active
copy r s
```

Storage Configuration

FlashArray Storage Deployment

The Pure Storage FlashArray//X is accessible to the FlashStack, but no storage has been deployed at this point. The storage to be deployed will include:

- ESXi FC Boot LUNs
- VMFS Datastores
- FC-NVMe Data stores

The FC Boot LUNs will need to be setup from the Pure Storage Web Portal, and the VMFS datastores can be provisioned from the Pure Storage Web Portal or can be directly provisioned from the vSphere Web Client after the Pure Storage vSphere Web Client Plugin has later been registered with the vCenter.

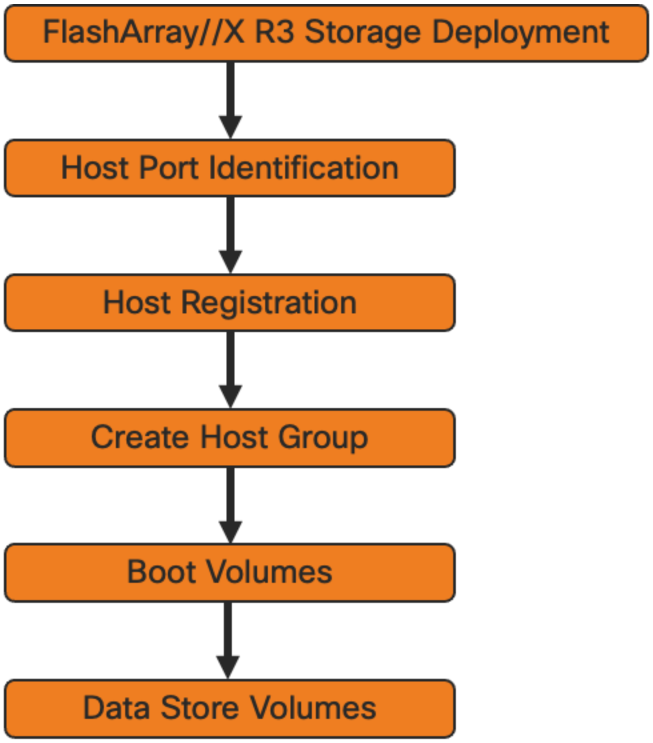

Host Port Identification

FC Boot LUNs will be mapped by the FlashArray//X using the assigned Initiator PWWN to the provisioned server profiles. This information can be found within the server profile located within the Cisco Intersight > Configure > UCS Server Profiles:

Host Registration

To register the Host, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.

2. Select the + icon in the Hosts Panel.

3. After clicking the Create Host (+) option, a pop-up will appear to create an individual host entry on the FlashArray.

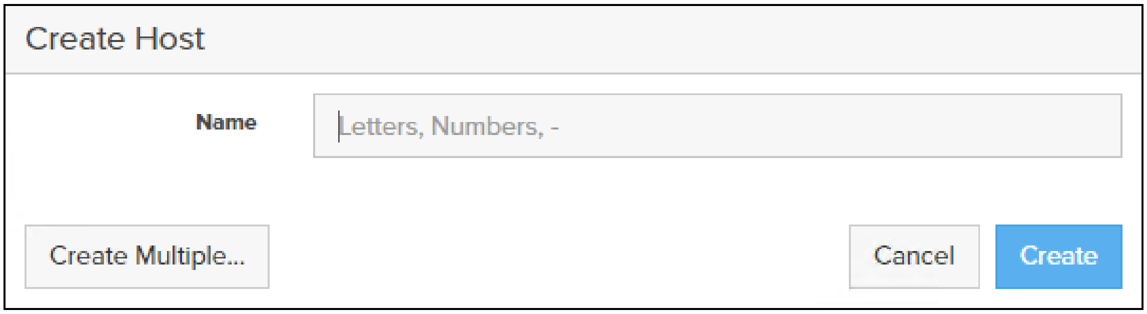

Create Host

Name: Letters, Numbers, -

4. To create more than one host entry, click the Create Multiple... option, filling in the Name, Start Number, Count, Personality as ESXi and Number of Digits, with a "#" appearing in the name where an iterating number will appear:

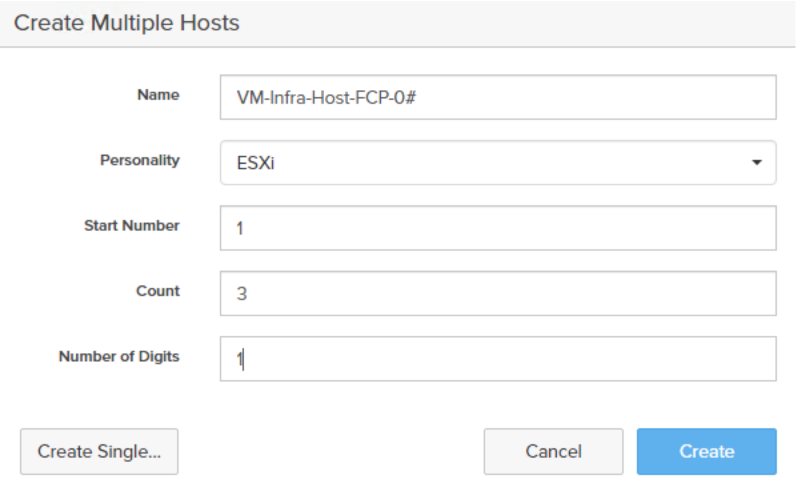

Create Multiple Hosts

Name: VM-Infra-Host-FCP-0#

Personality: ESXi

Start Number: 1

Count: 3

Number of Digits: 1

5. Click Create to add the hosts.

6. For each host created, select the host.

7. In the Host view, select 'Configure WWNs...' from the Host Ports menu.

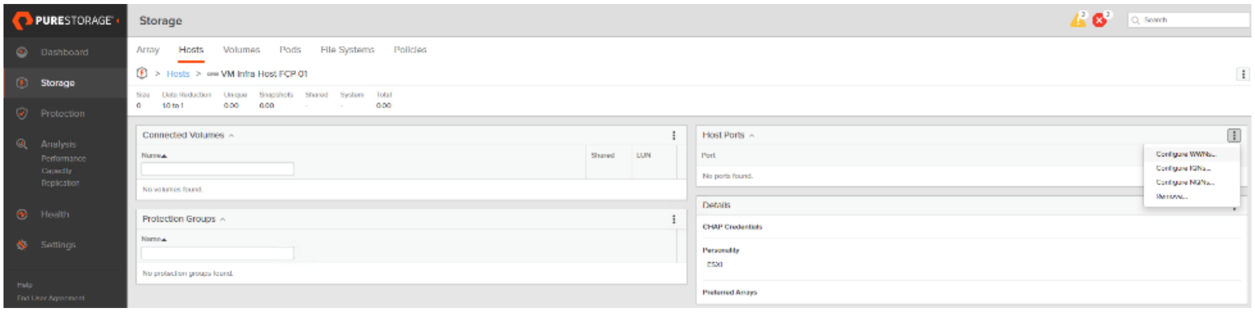

> Hosts > VM Infra Host FCP 01

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	10 to 1	0.00	0.00	-	-	0.00

Connected Volumes

Name▲	Shared	LUN

No volumes found.

Protection Groups

Name▲

No protection groups found.

Host Ports

Port

No ports found.

Details

CHAP Credentials

Personality
ESXi

Preferred Arrays

8. A pop-up will appear for Configure Fibre Channel WWNs <host being configured>. Within this pop-up, select the appropriate Existing WWNs from the discovered list.

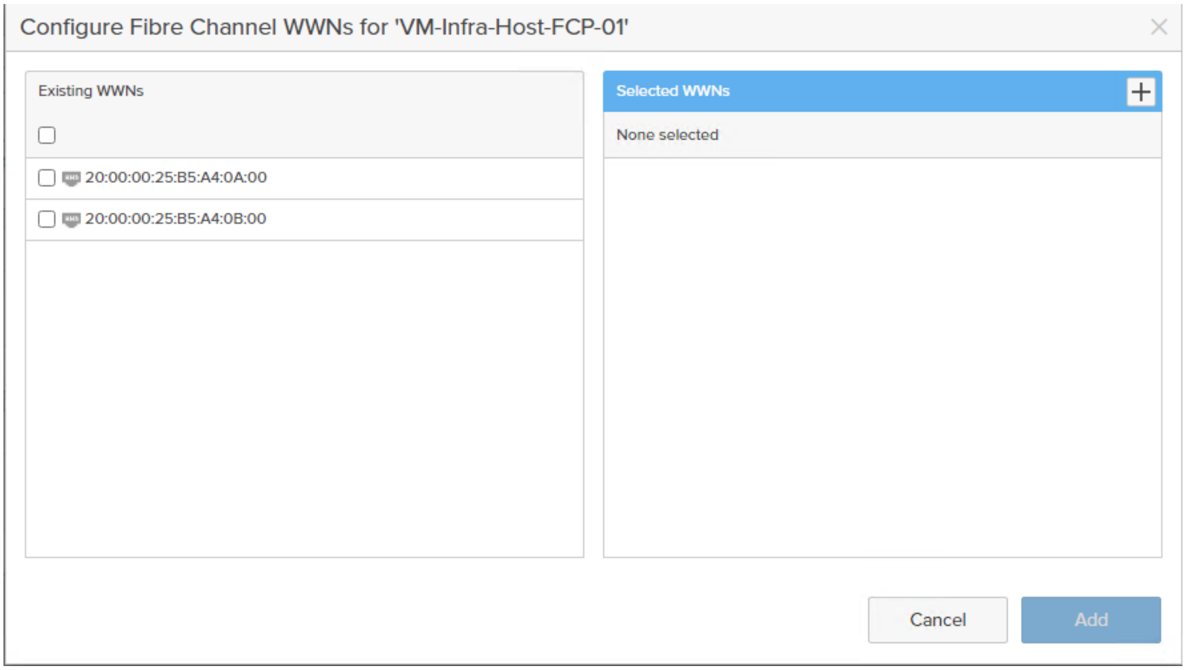

Configure Fibre Channel WWNs for 'VM-Infra-Host-FCP-01'

Existing WWNs

- □ 20:00:00:25:B5:A4:0A:00
- □ 20:00:00:25:B5:A4:0B:00

Selected WWNs

None selected

9. Or you may enter the WWN manually by Selecting the +.

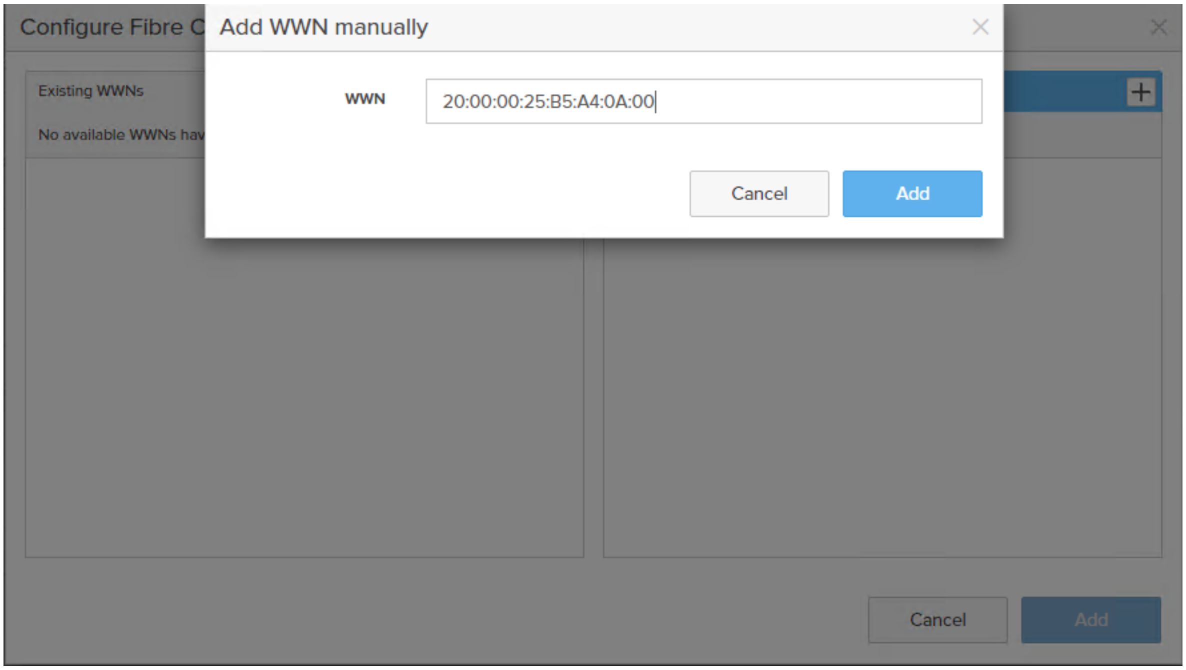

Add WWN manually

WWN: 20:00:00:25:B5:A4:0A:00

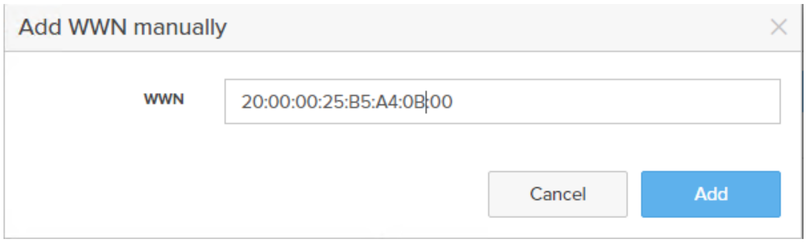

Add WWN manually

WWN: 20:00:00:25:B5:A4:0B:00

10. After entering the PWWN/WWPN, click Add to add the Host Ports.

11. Repeat steps 1-10 for each host created.

Create Host Group

Host Groups allow the Administrator to map Volumes to a group of hosts at once with the same LUN ID. To create a Host Group, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.

2. Select the + icon in the Host Groups Panel.

3. A pop-up will appear to create a host group on the FlashArray.

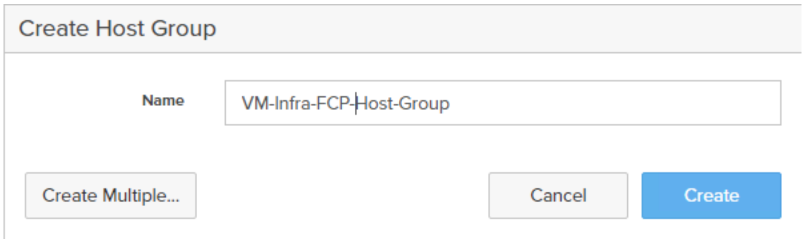

Create Host Group

Name: VM-Infra-FCP-Host-Group

4. Provide a name for the group and click Create.

5. Select the group in the Host Groups Panel.

6. In the Host Group view, select 'Add...' from the Member Hosts menu.

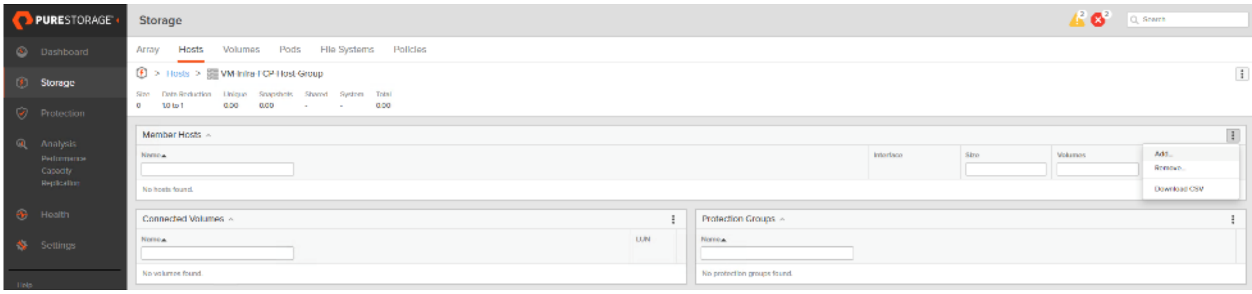

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

Member Hosts

Name ▲	Interface	Size	Volumes

No hosts found.

Connected Volumes

Name ▲	LUN

No volumes found.

Protection Groups

Name ▲

No protection groups found.

7. Select the host to be part of the host group.

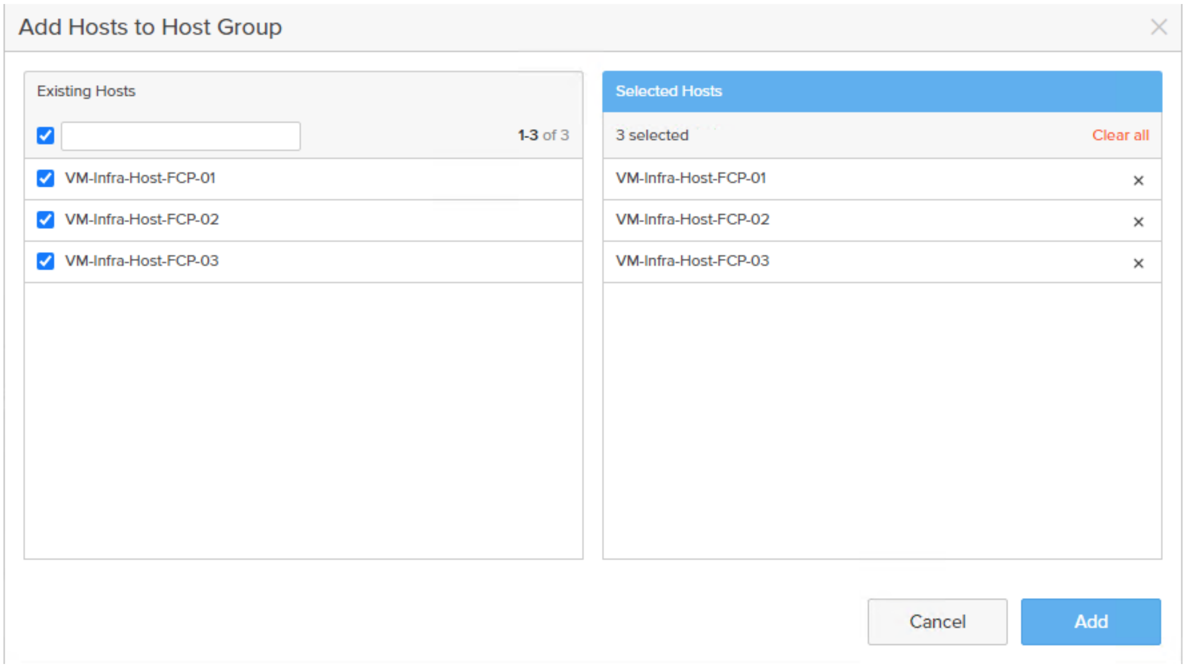

Add Hosts to Host Group

Existing Hosts	Selected Hosts
☑ 1-3 of 3	3 selected Clear all
☑ VM-Infra-Host-FCP-01	VM-Infra-Host-FCP-01 ×
☑ VM-Infra-Host-FCP-02	VM-Infra-Host-FCP-02 ×
☑ VM-Infra-Host-FCP-03	VM-Infra-Host-FCP-03 ×

8. Click Add.

Private Boot Volumes for each ESXi Host

To create private boot volumes for each ESXi Host, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.

2. Select the + icon in the Volumes Panel.

3. A pop-up will appear to create a volume on the FlashArray.

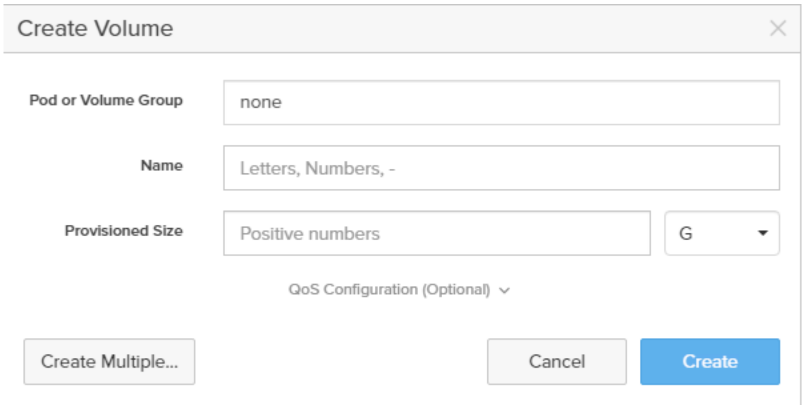

Create Volume

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

4. To create more than one volume, click the Create Multiple... option, filling in the Name, Provisioned Size, Staring Number, Count, and Number of Digits, with a “#” appearing in the name where an iterating number will appear.

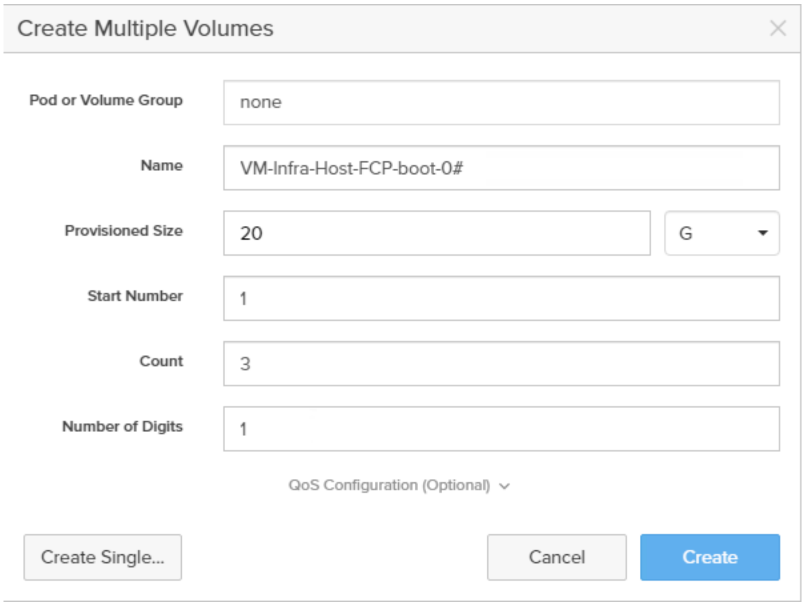
Create Multiple Volumes

Pod or Volume Group: none

Name: VM-Infra-Host-FCP-boot-0#

Provisioned Size: 20 G

Start Number: 1

Count: 3

Number of Digits: 1

5. Click Create to provision the volumes to be used as FC boot LUNs.

6. Go back to the Hosts section under the Storage tab. Click one of the hosts and select the gear icon drop-down within the Connected Volumes tab within that host.

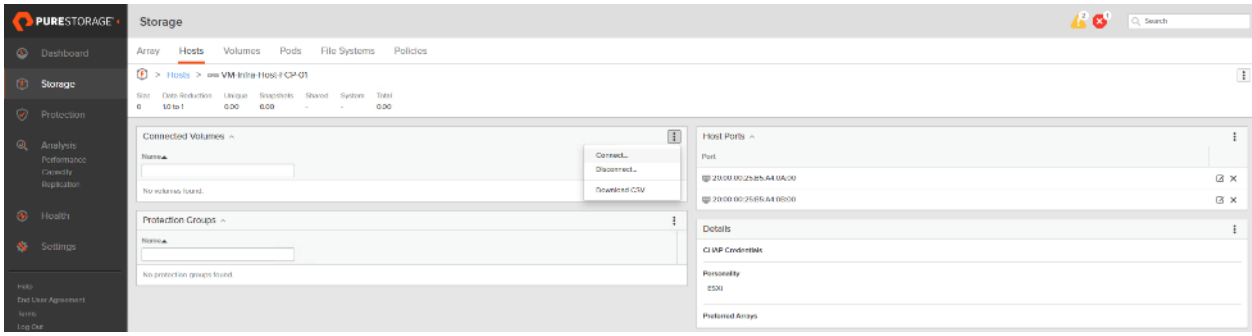

Hosts > VM-Infra-Host-FCP-01

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

Connected Volumes

Name▲

No volumes found.

Protection Groups

Name▲

No protection groups found.

Host Ports

Port

20:00:00:25:B5:A4:0A:00

20:00:00:25:B5:A4:0B:00

Details

CHAP Credentials

Personality

ESXi

Preferred Arrays

7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

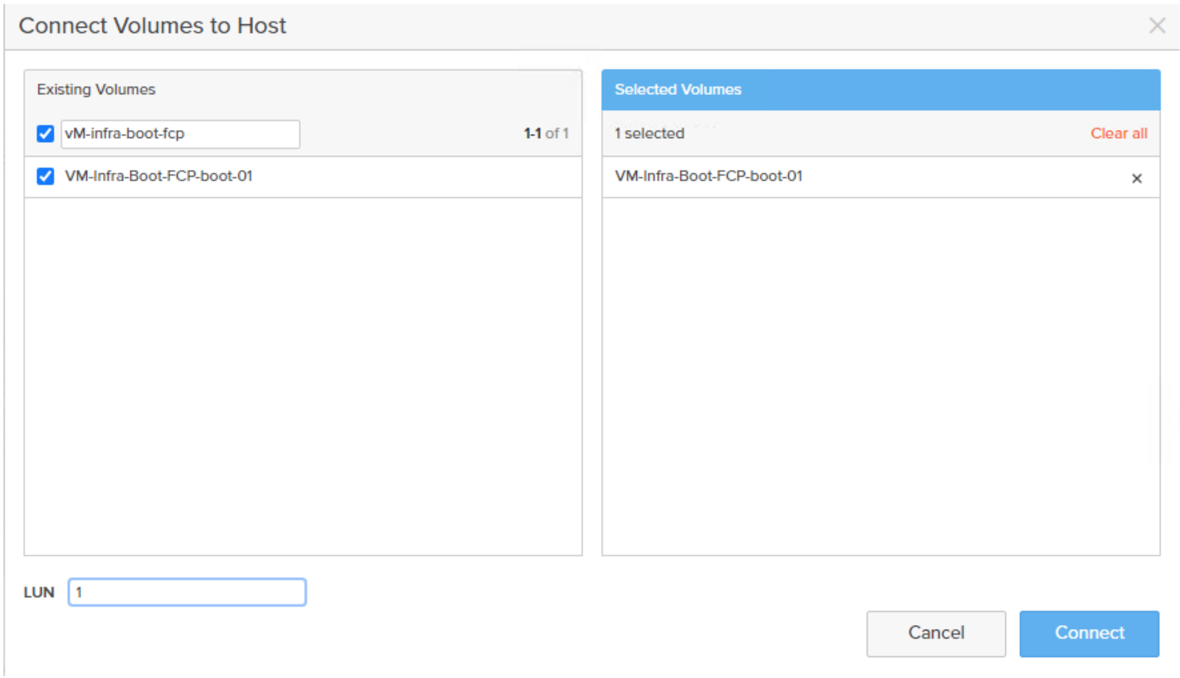

Connect Volumes to Host

Existing Volumes

☑ vM-infra-boot-fcp 1-1 of 1

☑ VM-Infra-Boot-FCP-boot-01

Selected Volumes

1 selected

VM-Infra-Boot-FCP-boot-01

LUN 1

Note: LUN ID 1 should be used for the boot.

8. Select the volume that has been provisioned for the host, set the LUN ID for the volume, click the + next to the volume, and select Confirm to proceed. Repeat the steps for connecting volumes for each of the host/volume pairs configured.

Create Infra Datastores

To create datastore volumes for the ESXi Cluster, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.
2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

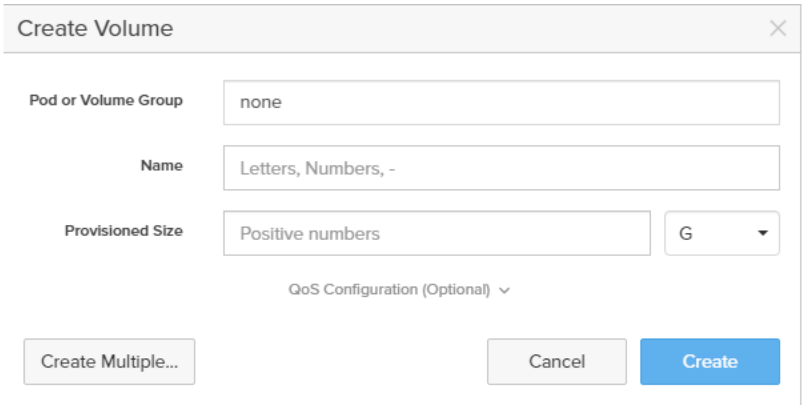

Create Volume

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

4. Fill in the Name and Provisioned Size.
5. Click Create to provision the volumes to be used as Infra datastore LUN.

6. Go back to the Hosts section under the Storage tab. Click ESXi cluster host group created earlier and select the gear icon drop-down within the Connected Volumes tab within that host group.

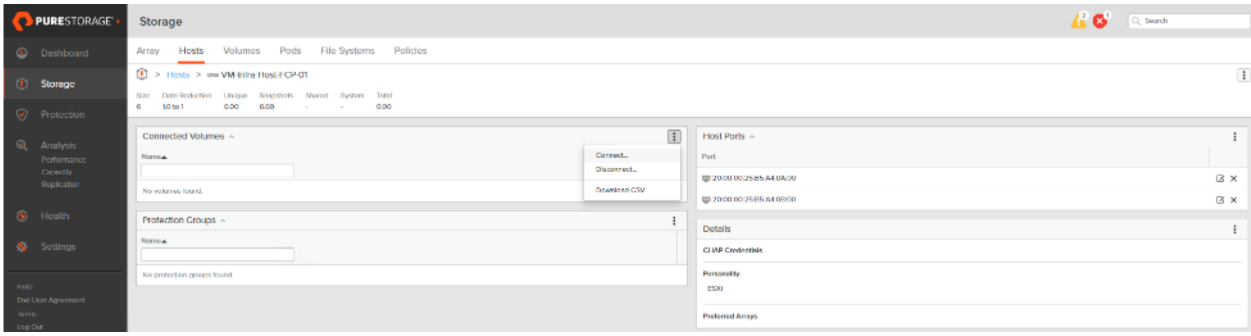

7. Within the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

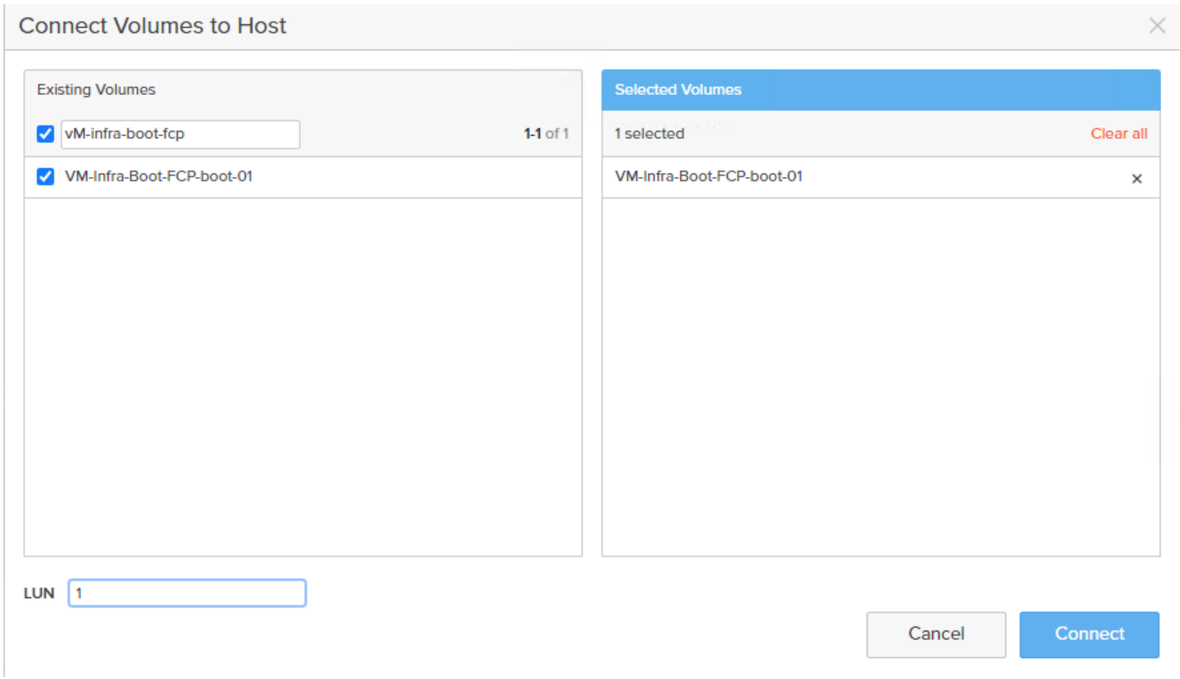
Connect Volumes to Host

Existing Volumes

☑ vM-infra-boot-fcp 1-1 of 1

☑ VM-Infra-Boot-FCP-boot-01

Selected Volumes

1 selected Clear all

VM-Infra-Boot-FCP-boot-01 ×

LUN 1

8. Select the Infra datastore volume that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.

Configure Storage Policy Based Management

vSphere can communicate to the array via VASA provider to find out what features it supports and allow the vSphere administrator to assign, change, or remove functionality on a VVol on demand and via policies. Below is an example of how to configure a Protection group that will provide hourly snapshots that will be retained for 1 day, with 4 snapshots per day retained for 7 days. These policies should be configured based on application snapshot need.

To configure Storage Policy Based Management, follow these steps

1. From the Pure Storage Web Portal, Select Protection > Protection Groups > select the + icon in the Source Protection Groups.

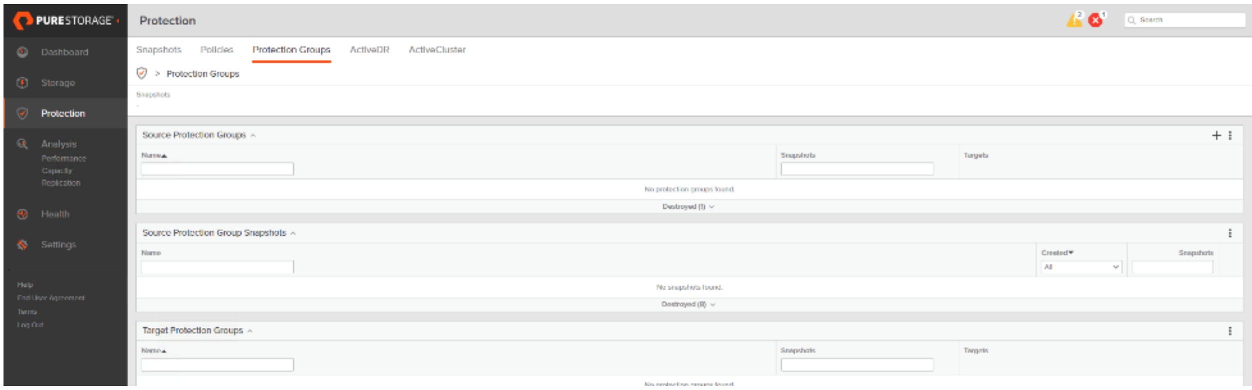

No protection groups found.

No snapshots found.

No protection groups found.

2. Enter a name.

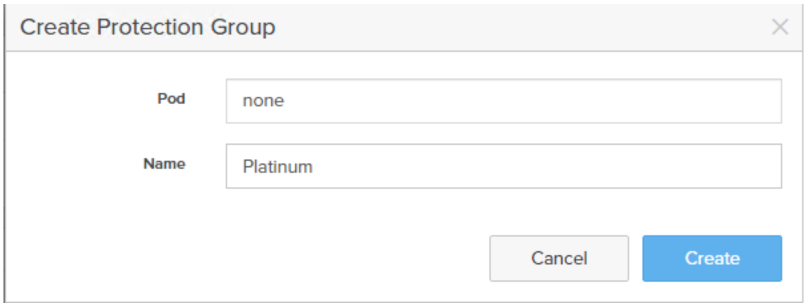

Create Protection Group

Pod: none

Name: Platinum

3. Select the protection group.

4. Edit the Snapshot Schedule based on your operational requirements.

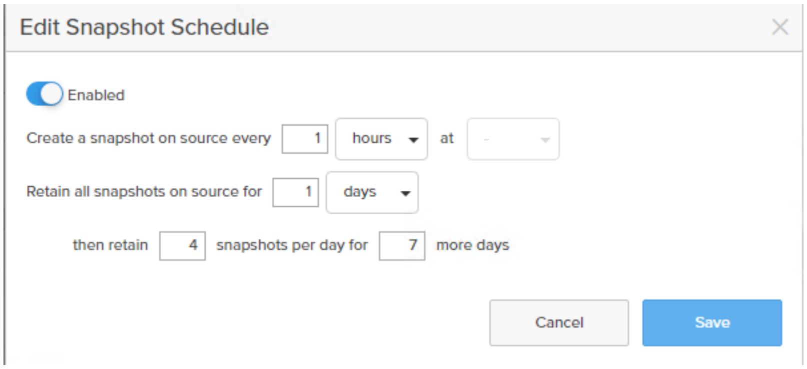

Edit Snapshot Schedule

Enabled

Create a snapshot on source every 1 hours at -

Retain all snapshots on source for 1 days

then retain 4 snapshots per day for 7 more days

5. Click Save.

VMware vSphere 7.0 U2 Setup

VMware ESXi 7.0 U2

This section provides detailed instructions for installing VMware ESXi 7.0 U2 in a FlashStack 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 U2 from VMware

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

1. Click this link: https://customerconnect.vmware.com/downloads/details?downloadGroup=OEM-ESXI70U2-CISCO&productId=974

Note: You will need a user id and password on vmware.com to download this software.

2. Download the .iso file.

Log into Cisco Intersight

Cisco Intersight

When a server profile is successfully deployed, to install operating system follow these steps:

To log into the Cisco UCS environment, follow these steps:

1. Go to Servers and click **...** next to the server and select Launch the vKVM.

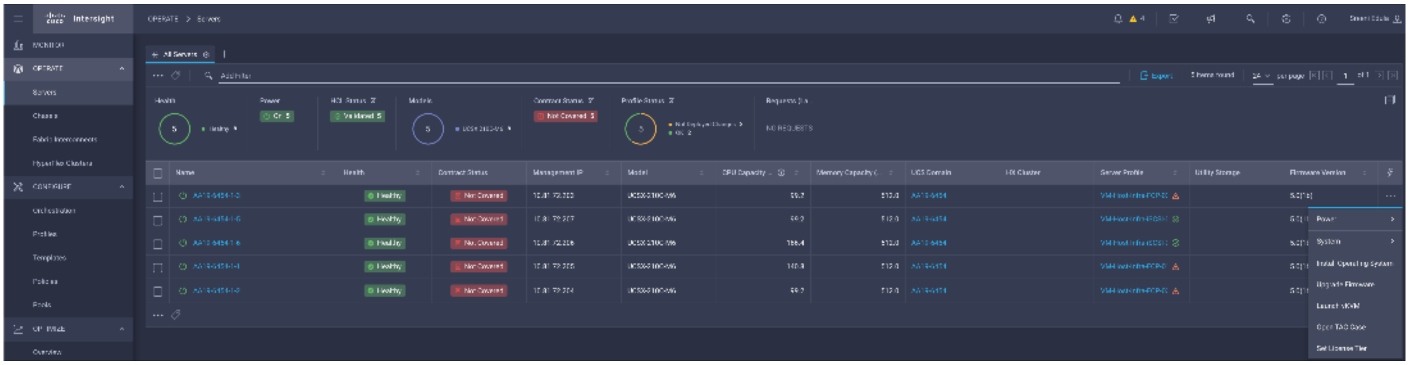

Note: Make sure the VM/host trying to access the KVM can route to the management IP address pool.

2. Login using the username (for example, flashadmin) and password previously defined in the Local User Policy.

3. In this new KVM tab on the browser, click Virtual Media, click Activate Virtual Devices.

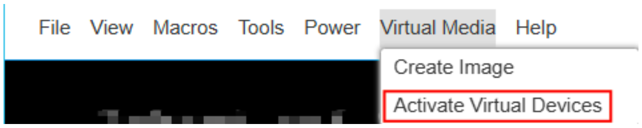

4. Click Virtual Media again and click Map CD/DVD...

5. Browse to ESXi ISO and click Map Drive.

6. In the Intersight portal:

   a. Power Cycle the Server by clicking ... next to the server in and selecting Power Cycle.

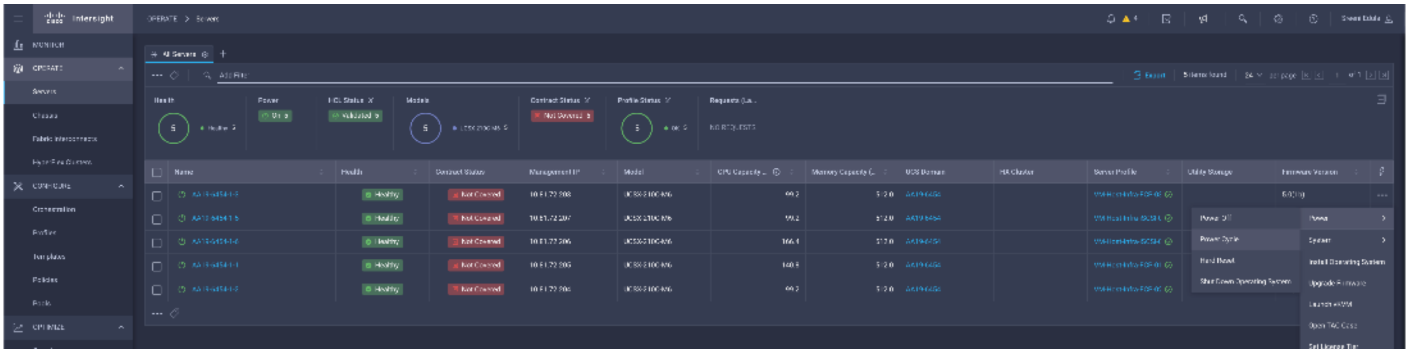

   b. Select Set One Time Boot Device and select ISO (label previously created for CD/DVD) from the drop-down list.

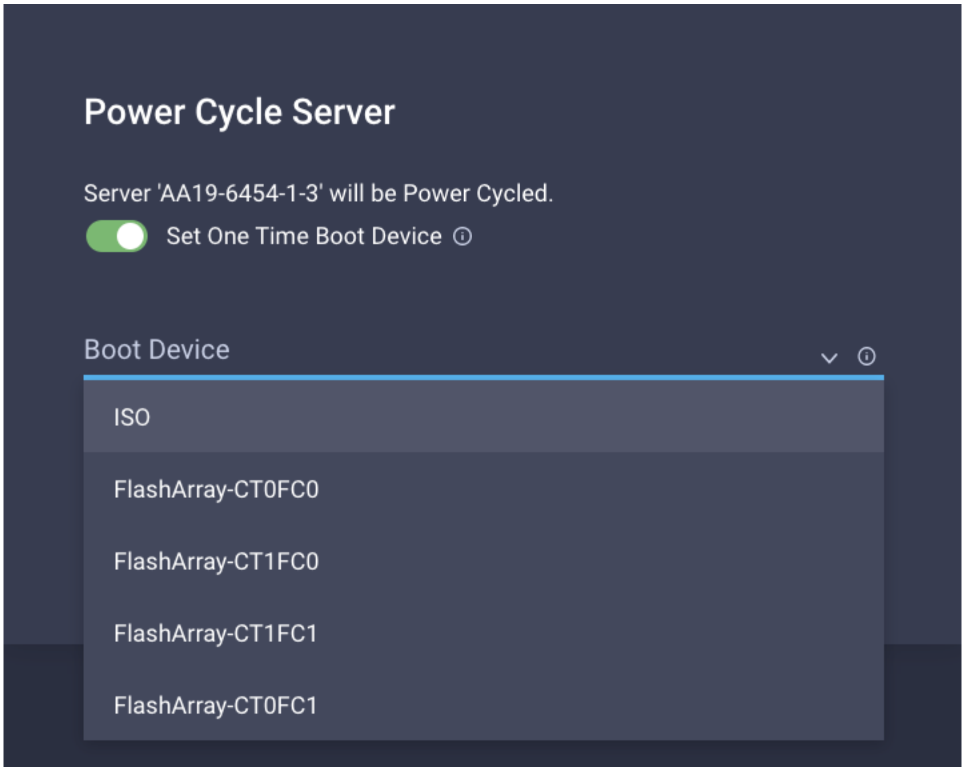

Power Cycle Server

Server 'AA19-6454-1-3' will be Power Cycled.

Set One Time Boot Device ⓘ

Boot Device ⓘ

- ISO
- FlashArray-CT0FC0
- FlashArray-CT1FC0
- FlashArray-CT1FC1
- FlashArray-CT0FC1

7. Click Power Cycle.

8. In the KVM Tab:

- In the KVM window, you should see the server being power cycled. If the zoning and boot LUNs were configured correctly, you will see the server has successfully discovered Boot LUN over all four paths.
- After ESXi installed is loaded, navigate through the ESXi installer instructions. The installer should discover the Pure Storage Boot LUN as an installation location.

9. Proceed with the ESXi installation and when installation is complete, un-map the installer using Deactivate Virtual Devices menu option under Virtual Media and reboot the server.

Note: You may follow the FlashStack CVD for setting up vCenter and other management tools

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To prepare the server for the OS installation, follow these steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Choose Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and choose Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM Console tab to monitor the server boot.

Install ESXi

ESXi Hosts VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To install VMware ESXi to the bootable LUN of the hosts, follow these steps on each host:

1. Boot the server by selecting Boot Server in the KVM and click OK, then click OK again.
2. On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.

Note: 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.

Note: 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.

Note: The ESXi installation image will be automatically unmapped in the KVM when Enter is pressed.

Set Up Management Networking for ESXi Hosts

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

ESXi Host VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To configure each ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, in the UCS KVM console, press F2 to customize VMware ESXi.

2. Log in as root, enter the corresponding password, and press Enter to log in.

3. Use the down arrow key to choose Troubleshooting Options and press Enter.

4. Choose Enable ESXi Shell and press Enter.

5. Choose Enable SSH and press Enter.

6. Press Esc to exit the Troubleshooting Options menu.

7. Choose the Configure Management Network option and press Enter.

8. Choose Network Adapters and press Enter.

9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field. If the numbers do not match, note the mapping of vmnic ports to vNIC ports for later use.

10. Using the spacebar, choose vmnic1.

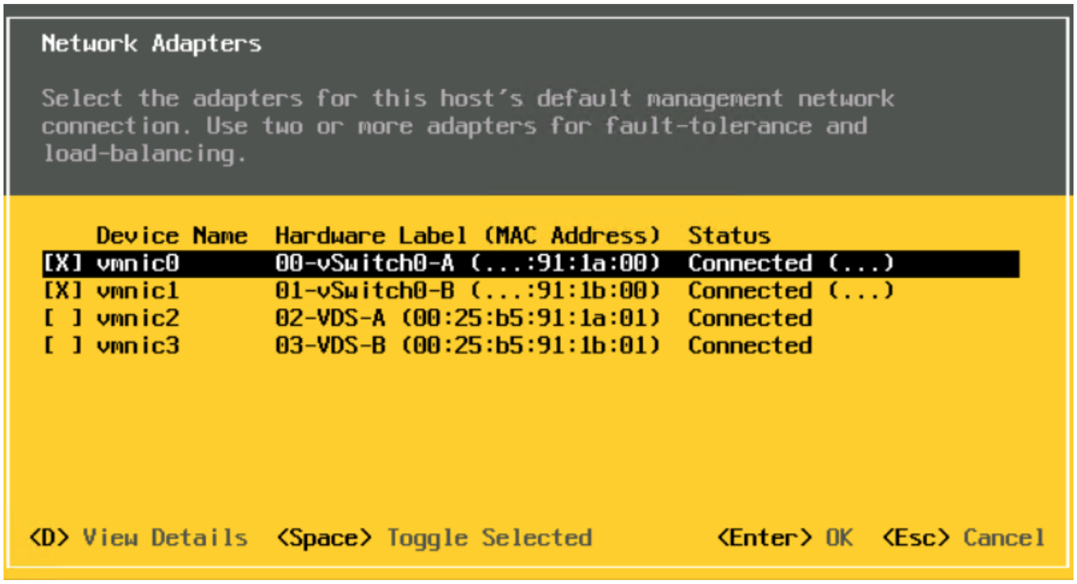

```
Network Adapters

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 (...:91:1a:00)  Connected (...)
[X] vmnic1       01-vSwitch0-B (...:91:1b:00)  Connected (...)
[ ] vmnic2       02-VDS-A (00:25:b5:91:1a:01)  Connected
[ ] vmnic3       03-VDS-B (00:25:b5:91:1b:01)  Connected

<D> View Details  <Space> Toggle Selected        <Enter> OK  <Esc> Cancel
```

Note: 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.

Note: Since the IP address is assigned manually, the DNS information must also be entered manually.

23. Using the spacebar, choose “Use the following DNS server addresses and hostname:”

24. Move to the Primary DNS Server field and enter the IP address of the primary DNS server.

25. Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.

26. Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.

27. Press Enter to accept the changes to the DNS configuration.

28. Press Esc to exit the Configure Management Network submenu.

29. Press Y to confirm the changes and reboot the ESXi host.

Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional)

ESXi VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port it is placed on. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will exist because vmk0 will retain the assigned MAC address unless the ESXi System Configuration is reset. To reset the MAC address of vmk0 to a random VMware-assigned MAC address, follow these steps:

1. From the ESXi console menu main screen, type Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of Static Macros.

2. Log in as root.

3. Type esxcfg-vmknic –l to get a detailed listing of interface vmk0. vmk0 should be a part of the "Management Network" port group. Note the IP address and netmask of vmk0.

4. To remove vmk0, type esxcfg-vmknic –d "Management Network".

5. To re-add vmk0 with a random MAC address, type esxcfg-vmknic –a –i <vmk0-ip> -n <vmk0-netmask> "Management Network".

6. Verify vmk0 has been re-added with a random MAC address by typing esxcfg-vmknic –l.

7. Tag vmk0 as the management interface by typing esxcli network ip interface tag add -i vmk0 -t Management.

8. When vmk0 was re-added, if a message popped up saying vmk1 was marked as the management interface, type esxcli network ip interface tag remove -i vmk1 -t Management.

9. If this VMware ESXi host is iSCSI booted, the vmk1, iScsiBootPG-A interface's MAC address can also be reset to a random, VMware-assigned MAC address.

   a. Type esxcfg-vmknic –l to get a detailed listing of interface vmk1. vmk1 should be a part of the "iScsiBootPG-A" port group and should have a MAC address from the UCS MAC Pool. Note the IP address and netmask of vmk1.

   a. To remove vmk1, type esxcfg-vmknic –d "iScsiBootPG-A".

b. To re-add vmk1 with a random MAC address, type esxcfg-vmknic –a –i <vmk1-ip> -n <vmk1-netmask> -m 9000 “iScsiBootPG-A”.

c. Verify vmk1 has been re-added with a random MAC address by typing esxcfg-vmknic –l.

d. 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 VMware VIC Driver to the Management workstation:

UCS Tools Component for ESXi 7.0 1.2.1 (ucs-tool-esxi_1.2.1-1OEM.zip)

VMware ESXi 7.0 nfnic 5.0.0.15 Driver for Cisco VIC Adapters (Cisco-nfnic_5.0.0.15-1OEM.700.1.0.15843807_18697950.zip)

nenic Driver version 1.0.35.0 (nenic driver is included with the Cisco ESXi installation ISO).

This document is using the driver versions shown above. These were the versions validated and supported at the time this document was published. This document can be used as a guide for configuring future versions of software. Consult the Cisco UCS Hardware Compatibility List and the Pure Interoperability Matrix Tool to determine supported combinations

ESXi Hosts VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To install UCS Tools on the ESXi host ESXi VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03, follow these steps:

Note: The latest nenic driver is already included with the ESXi install ISO and is not required to be updated if the Cisco Custom ISO for ESXi 7.0 U2 is used.

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/Cisco-nfnic_5.0.0.15-1OEM.700.1.0.15843807_18697950.zip

esxcli software component apply -d /tmp/ucs-tool-esxi_1.2.1-1OEM.zip

reboot
```

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

```
esxcli software vib list | grep nenic

esxcli software component list | grep nfnic
esxcli software component list | grep ucs
```

Log into the First VMware ESXi Host by Using VMware Host Client

ESXi Host VM-Host-Infra-FCP-01

To log into the VM-Host-Infra-FCP-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-FCP-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-FCP-01

To set up the VMkernel ports and the virtual switches on the first ESXi host, follow these steps:

Note: In this procedure, you’re only setting up the first ESXi host. The second and third hosts will be added to vCenter and setup from the vCenter HTML5 Interface.

1. From the Host Client Navigator, choose Networking.
2. In the center pane, choose the Virtual switches tab.
3. Highlight the vSwitch0 line.
4. Choose Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, choose vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.
9. Click Save.
10. Choose Networking, then choose the Port groups tab.
11. In the center pane, right-click VM Network and choose Edit settings.
12. Name the port group IB-MGMT Network and enter <ib-mgmt-vlan-id> in the VLAN ID field.

13. Click Save to finalize the edits for the IB-MGMT Network.

14. Click Add port group.

15. Name the port group OOB-MGMT Network and enter the <OOB-MGMT-vlan-id> for the VLAN ID.

16. Click Add to finalize the edits for the OOB-MGMT port group.

17. At the top, choose the VMkernel NICs tab.

18. Click VMkernel NICs tab.

19. Click Add VMkernel NIC.

20. For New port group, enter VMkernel-vMotion.

21. For Virtual switch, choose vSwitch0.

22. Enter <vmotion-vlan-id> for the VLAN ID.

23. Change the MTU to 9000.

24. Choose Static IPv4 settings and expand IPv4 settings.

25. Enter the ESXi host vMotion IP address and netmask.

26. Choose the vMotion stack for TCP/IP stack.

27. Click Create.

28. Choose the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be like the following example:

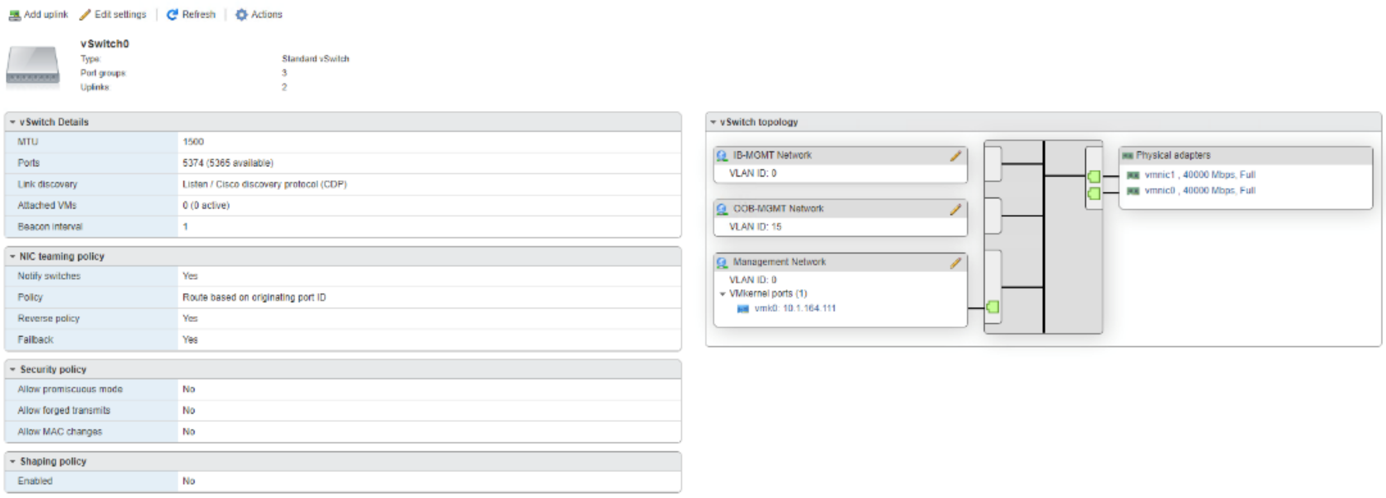

vSwitch0

Type:	Standard vSwitch
Port groups:	3
Uplinks:	2

vSwitch Details

MTU	1500
Ports	5374 (5365 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

NIC teaming policy

Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Failback	Yes

Security policy

Allow promiscuous mode	No
Allow forged transmits	No
Allow MAC changes	No

Shaping policy

Enabled	No

vSwitch topology

IB-MGMT Network
VLAN ID: 0

OOB-MGMT Network
VLAN ID: 15

Management Network
VLAN ID: 0
VMkernel ports (1)
vmk0: 10.1.164.111

Physical adapters
vmnic1 , 40000 Mbps, Full
vmnic0 , 40000 Mbps, Full

29. Choose Networking and the VMkernel NICs tab to confirm configured virtual adapter. The adapter listed should be like the following example:

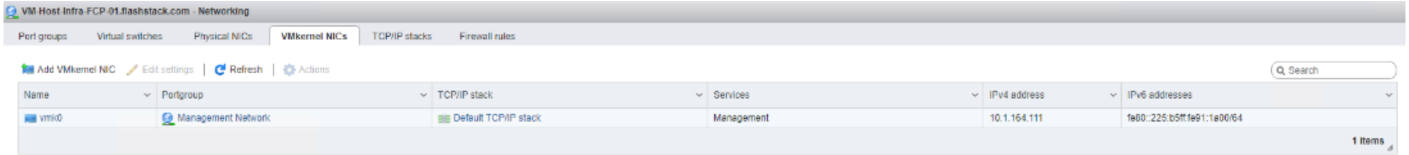

VM-Host-Infra-FCP-01.flashstack.com - Networking

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.164.111	fe80::225:b5ff:fe91:1a00/64

1 items

Mount Required Datastores

ESXi Host VM-Host-Infra-FCP-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. Click New datastore to add a new datastore.

4. In the New datastore, choose Create new VMFS datastore and click Next.

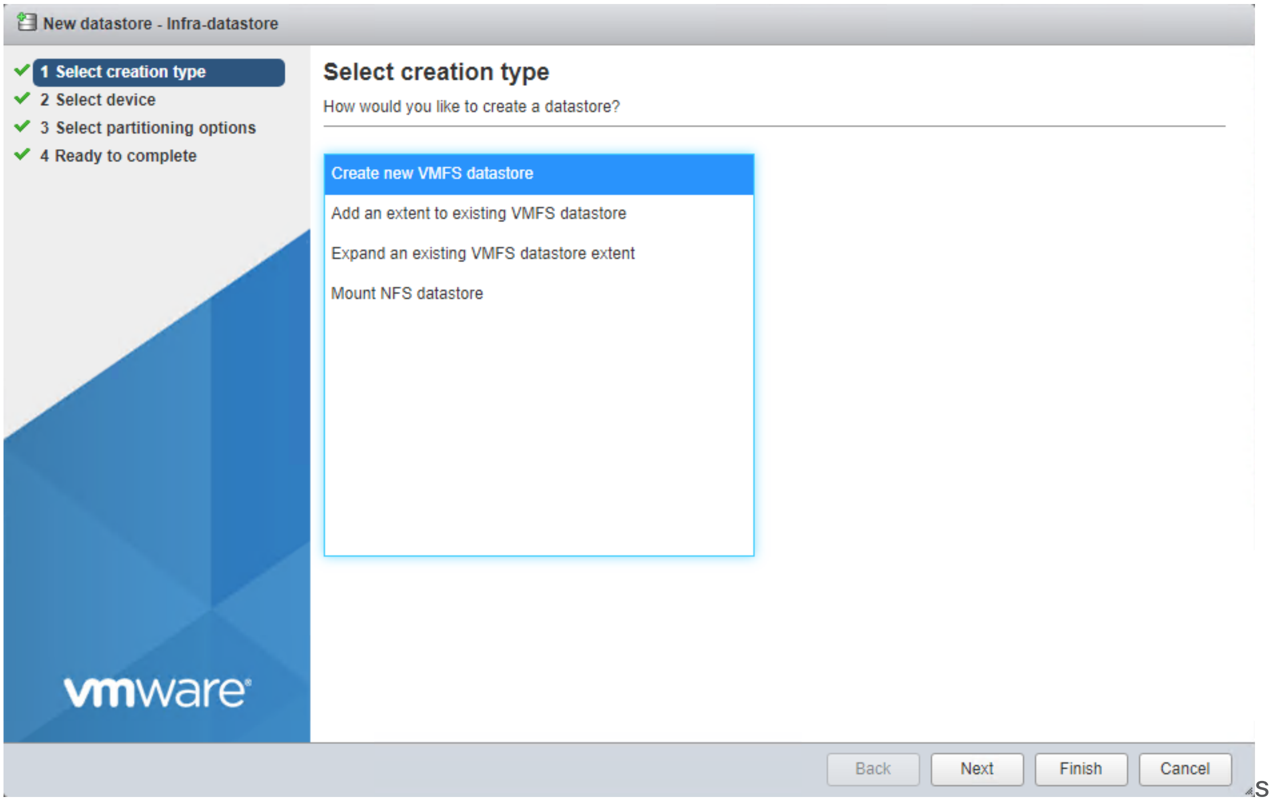

New datastore - Infra-datastore

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

S

5. Input Infra-Datastore1 for the datastore name.

6. Select the Pure LUN that will be used for the data store.

7. Click Next.

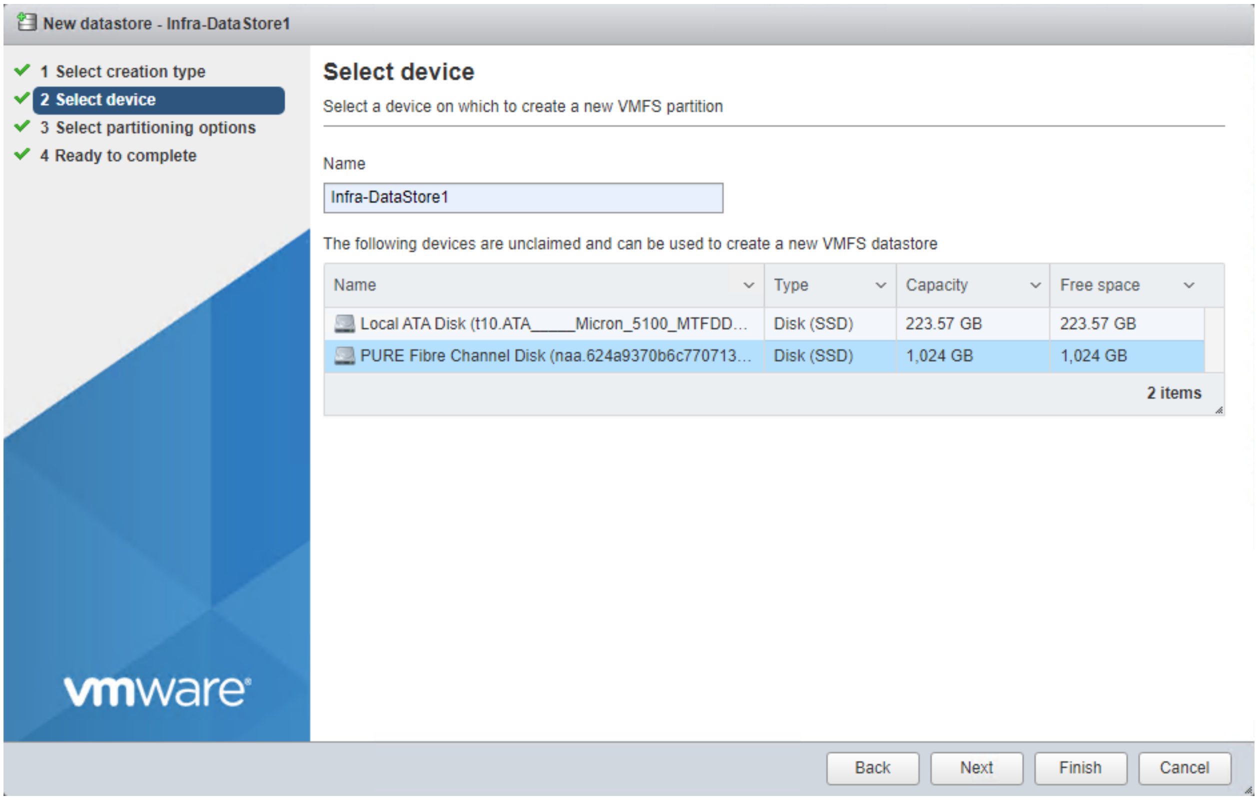

Select device

Select a device on which to create a new VMFS partition

Name: Infra-DataStore1

The following devices are unclaimed and can be used to create a new VMFS datastore

Name	Type	Capacity	Free space
Local ATA Disk (t10.ATA_____Micron_5100_MTFDD...	Disk (SSD)	223.57 GB	223.57 GB
PURE Fibre Channel Disk (naa.624a9370b6c770713...	Disk (SSD)	1,024 GB	1,024 GB

2 items

8. Click Next.

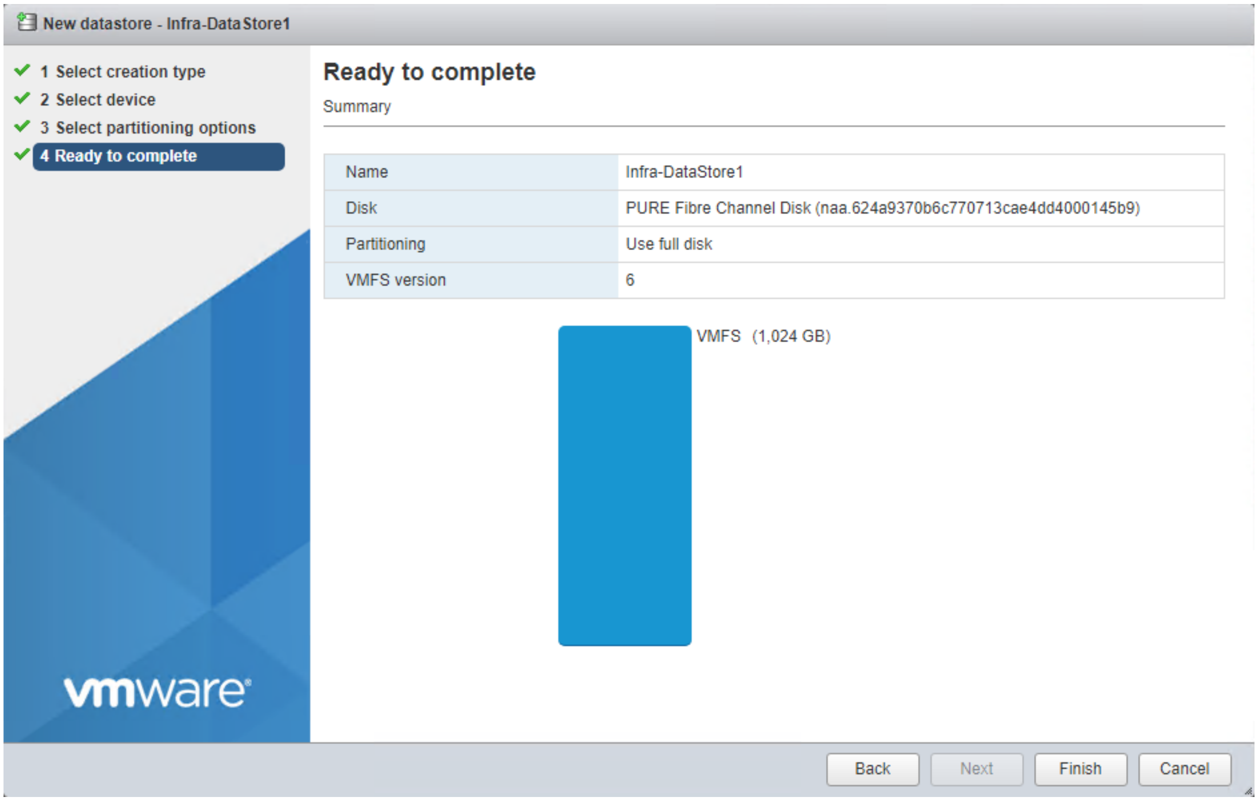

Ready to complete

Summary

Name	Infra-DataStore1
Disk	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd4000145b9)
Partitioning	Use full disk
VMFS version	6

9. Click Finish. The datastore should now appear in the datastore list.

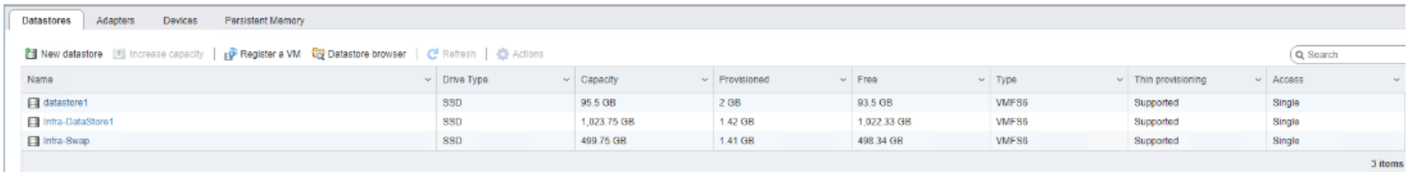

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
datastore1	SSD	95.5 GB	2 GB	93.5 GB	VMFS6	Supported	Single
Infra-DataStore1	SSD	1,023.75 GB	1.42 GB	1,022.33 GB	VMFS6	Supported	Single
Infra-Swap	SSD	499.75 GB	1.41 GB	498.34 GB	VMFS6	Supported	Single

3 items

Configure NTP on First ESXi Host

ESXi Host VM-Host-Infra-FCP-01

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

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

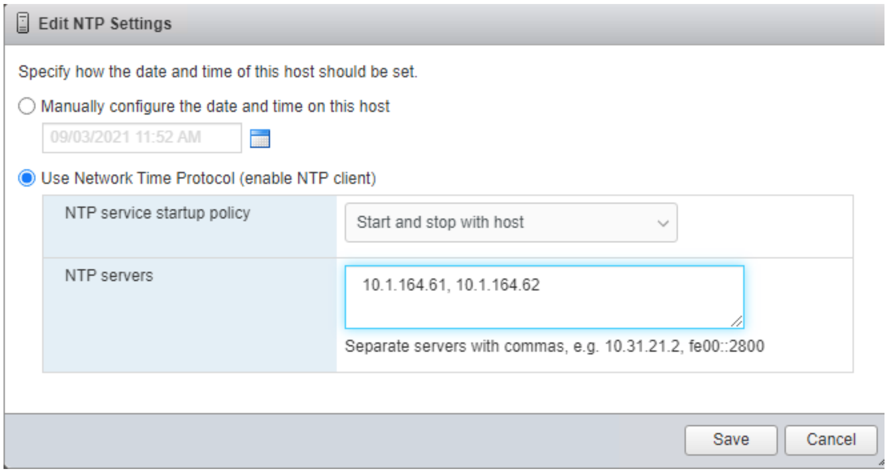

Specify how the date and time of this host should be set.

○ Manually configure the date and time on this host

09/03/2021 11:52 AM

● Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 10.1.164.61, 10.1.164.62

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

8. Click Save to save the configuration changes.

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

Configure Host Power Policy

ESXi Host VM-Host-Infra-FCP-01

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

Note: Implementation of this policy is recommended in Performance Tuning for Cisco UCS M6 Servers for maximum performance. If your organization has specific power policies, please set this policy accordingly.

1. From the Host Client, choose Manage.

2. In the center pane, choose Hardware > Power Management.

3. Choose Change policy.

4. Choose High performance and click OK.

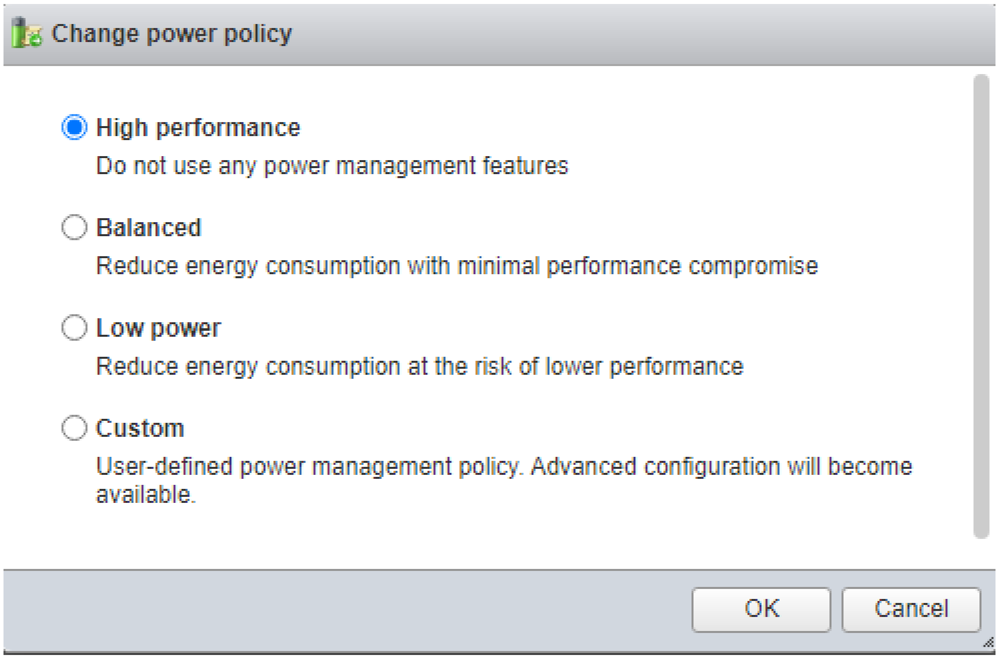

Change power policy

● 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.

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

VMware vCenter 7.0 U2B (Optional)

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 7.0U2B Server Appliance in a FlashStack 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-7.0.2-17958471.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 7.0 U2 vCenter Server Appliance.

Note: It is important to use at minimum VMware vCenter release 7.0U2 to ensure access to all needed features.

2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012 and above).

3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appears.

4. Click Install to start the vCenter Server Appliance deployment wizard.

5. Click NEXT in the Introduction section.

6. Read and accept the license agreement and click NEXT.

7. In the “vCenter Server deployment target” window, enter the host name or IP address of the first ESXi host, User name (root), and Password. Click NEXT.

8. Click YES to accept the certificate.

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

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

11. Choose Infra-DataStore1 for storage. Click NEXT.

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

    a. Choose a Network: IB-MGMT Network.

Note: It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it isn’t moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, and it is attempted to bring up vCenter on a different host than the one it was running on before the shutdown, vCenter will not have a functional network connection. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. 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-server1>,<dns-server2>

13. Click NEXT.

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

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

15. Click CONTINUE to proceed with stage 2 configuration.

16. Click NEXT.

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

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

18. Click NEXT.

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

20. Click NEXT.

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

22. Click NEXT.

23. Review the configuration and click FINISH.

24. Click OK.

Note: 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-FCP-01 management IP address.

2. Enter root for the user name.

3. Enter the root password.

4. Click Login to connect.

5. On the left, choose Virtual Machines.

6. In the center pane, right-click the vCenter VM and choose Edit settings.

7. In the Edit settings window, expand CPU and check the value of Sockets.

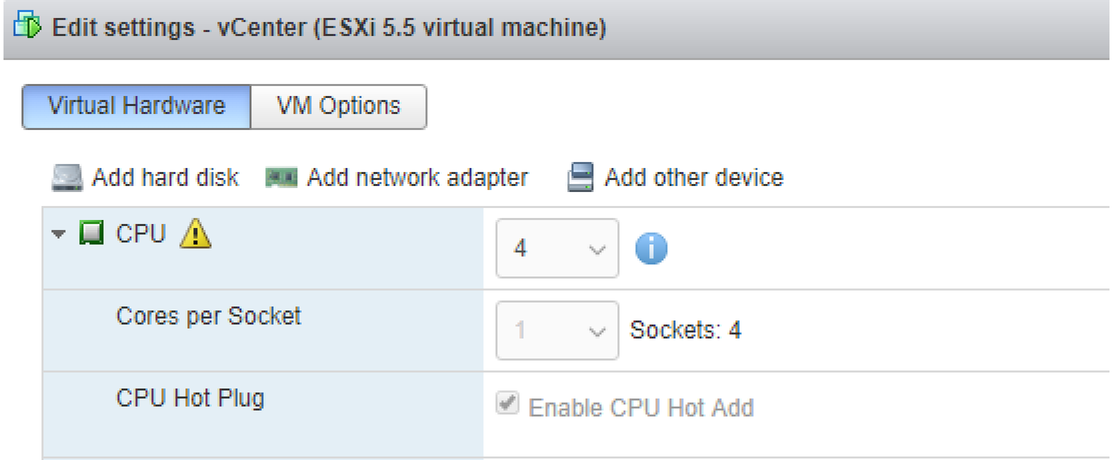

Edit settings - vCenter (ESXi 5.5 virtual machine)

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 (normally 2).

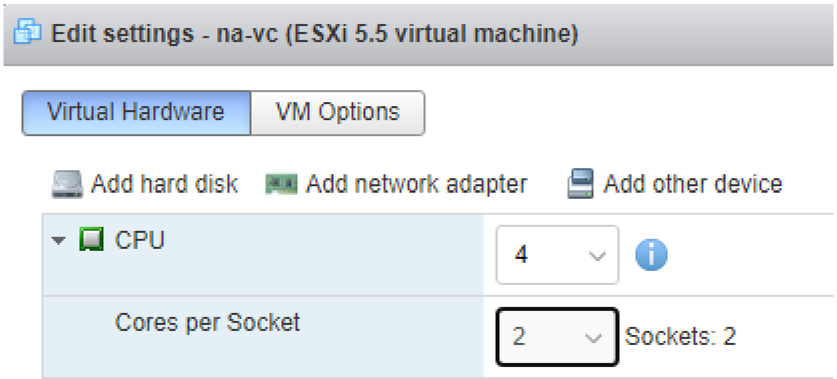

Edit settings - na-vc (ESXi 5.5 virtual machine)

CPU	4
Cores per Socket	2 Sockets: 2

   d. Click Save.

   e. Right-click the vCenter VM and choose Power > Power on. Wait approximately 10 minutes for vCenter to come up.

Setup VMware vCenter Server

To setup the VMware vCenter Server, follow these steps:

1. Using a web browser, navigate to https://<vcenter-ip-address>:5480. You will need to navigate security screens.

2. Log into the 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 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 7.0.2.00200 was installed.

10. In the upper right-hand corner of the screen, choose root > Logout to logout of the Appliance Management interface.

11. Using a web browser, navigate to https://<vcenter-fqdn>. You will need to navigate security screens.

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

12. Choose LAUNCH VSPHERE CLIENT (HTML5).

Note: Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is the only option starting with 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 currently.

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

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

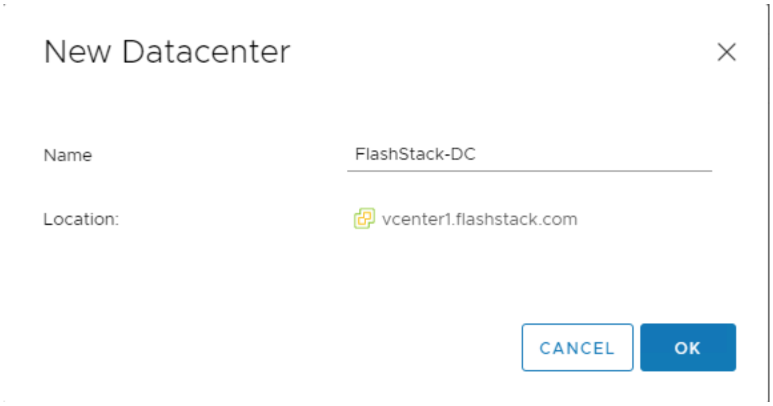

New Datacenter

Name: FlashStack-DC

Location: vcenter1.flashstack.com

16. Click OK.

17. Expand the vCenter on the left.

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

19. Name the cluster FlashStack-Management.

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

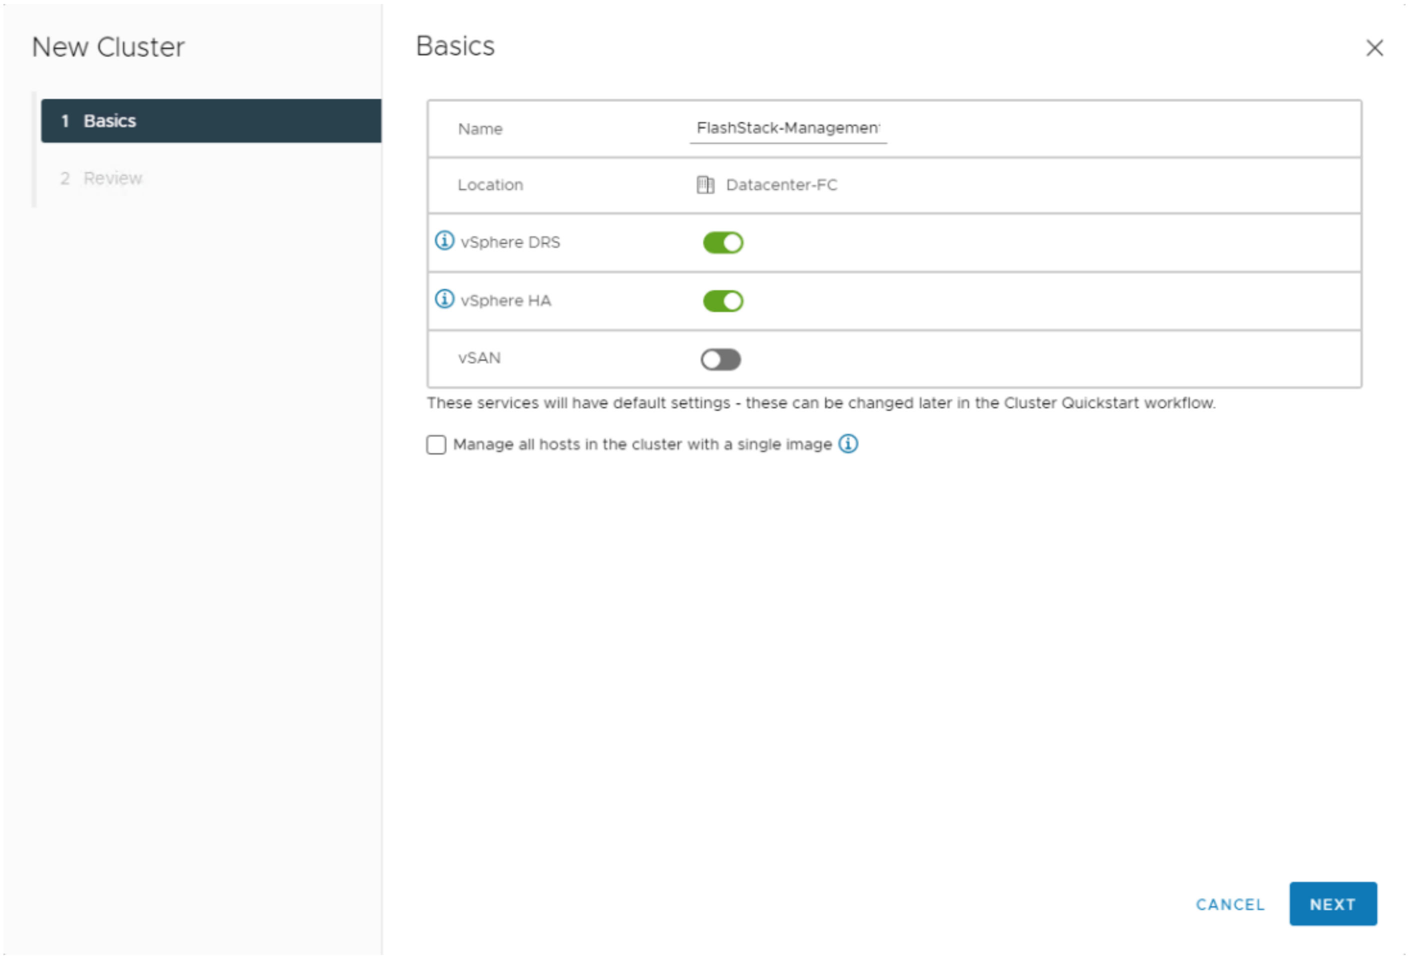

Name: FlashStack-Managemen

Location: Datacenter-FC

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 “FlashStack-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.

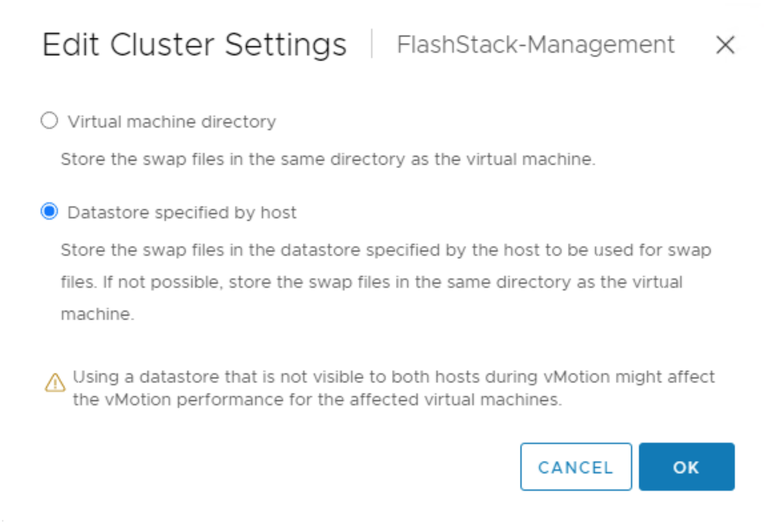

Edit Cluster Settings | FlashStack-Management

○ 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 “FlashStack-Management” and click Add Hosts.

26. In the IP address or FQDN field, enter either the IP address or the FQDN of the first VMware ESXi host. Enter root as the Username and the root password. Click NEXT.

27. In the Security Alert window, choose the host and click OK.

28. Verify the Host summary information and click NEXT.

29. Ignore warnings about the host being moved to Maintenance Mode and click FINISH to complete adding the host to the cluster.

30. The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.

31. In the list, right-click the added ESXi host and choose Settings.

32. In the center pane under Virtual Machines, choose Swap File location.

33. On the right, click EDIT.

34. Choose the Infra-Swap datastore and click OK.

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.

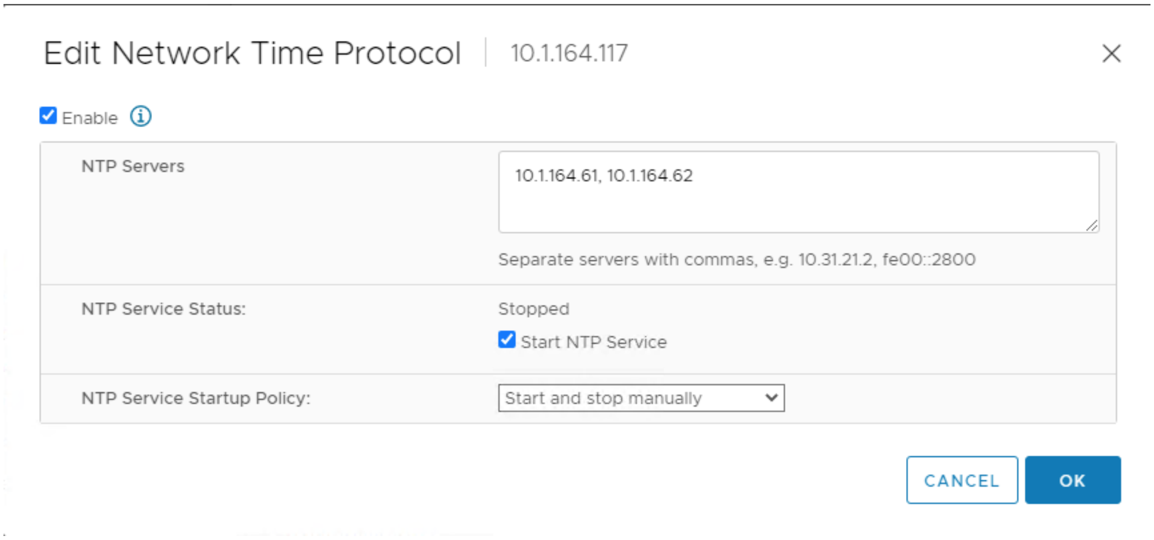

Edit Network Time Protocol | 10.1.164.117

☑ Enable

NTP Servers	10.1.164.61, 10.1.164.62 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 manually

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 Pure Fibre Channel Disk LUN 1 or Pure iSCSI Disk LUN 1 is selected.

41. Choose the Paths tab.

42. Ensure that 4 paths appear, which should have the status Active (I/O).

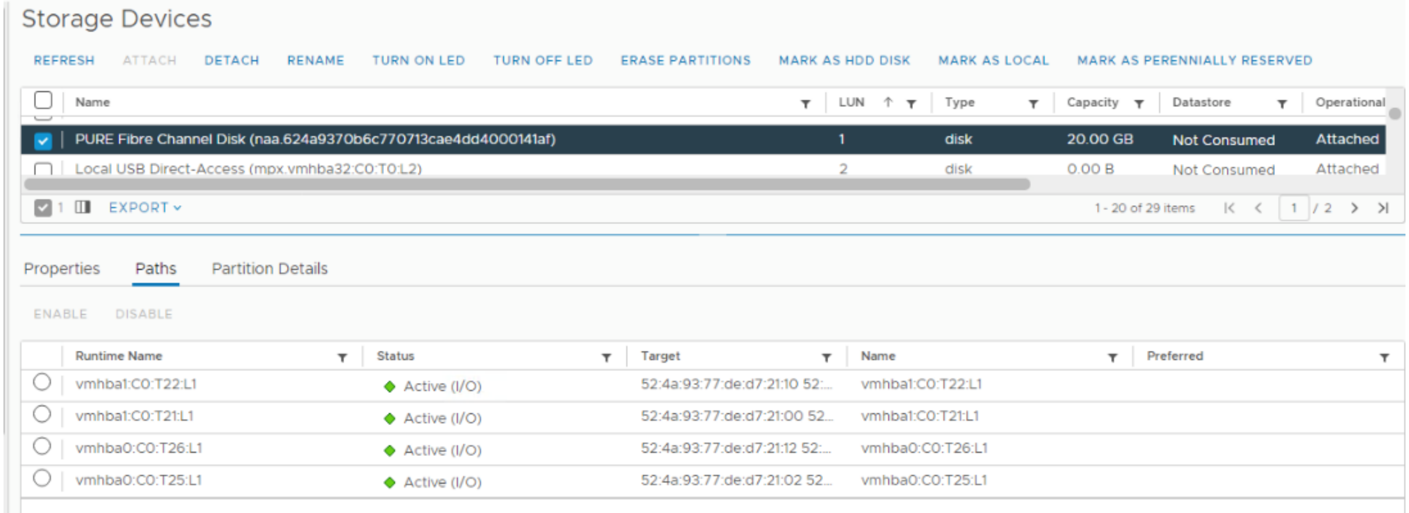

Storage Devices

	Name	LUN	Type	Capacity	Datastore	Operational
☑	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd4000141af)	1	disk	20.00 GB	Not Consumed	Attached
□	Local USB Direct-Access (mpx.vmhba32:C0:T0:L2)	2	disk	0.00 B	Not Consumed	Attached

	Runtime Name	Status	Target	Name	Preferred
○	vmhba1:C0:T22:L1	Active (I/O)	52:4a:93:77:de:d7:21:10 52:...	vmhba1:C0:T22:L1	
○	vmhba1:C0:T21:L1	Active (I/O)	52:4a:93:77:de:d7:21:00 52...	vmhba1:C0:T21:L1	
○	vmhba0:C0:T26:L1	Active (I/O)	52:4a:93:77:de:d7:21:12 52:...	vmhba0:C0:T26:L1	
○	vmhba0:C0:T25:L1	Active (I/O)	52:4a:93:77:de:d7:21:02 52...	vmhba0:C0:T25:L1	

Add AD User Authentication to vCenter (Optional)

If an AD Infrastructure is set up in this FlashStack environment, you can set up in AD and authenticate from vCenter.

To add an AD user authentication to the vCenter, follow these steps:

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flashadmin (FlashStack 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 FlashStack Admin username or the Domain Admins group. Leave the Role set to Administrator. Choose the Propagate to children checkbox.

Note: The FlashStack Admin user was created in the Domain Admins group. The selection here depends on whether the FlashStack Admin user will be the only user used in this FlashStack 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 FlashStack Admin user. You will need to add the domain name to the user, for example, flashadmin@domain.

FlashStack VMware vSphere Distributed Switch (vDS)

This section provides detailed procedures for installing the VMware vDS in vCenter and on the first FlashStack ESXi Management Host.

In the Cisco UCS setup section of this document two sets of vNICs were setup. The vmnic ports associated with the vDS0-A and B vNICs will be placed on the VMware vDS in this procedure. The vMotion VMkernel port(s) will be placed on the vDS.

A vMotion, and a VM-Traffic port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would need to have the corresponding VLANs added to the Cisco UCS LAN cloud, to the Cisco UCS vDS0-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC peer-link interfaces on the switches.

In this document, the infrastructure ESXi management VMkernel ports, the In-Band management interfaces including the vCenter management interface are left on vSwitch0 to facilitate bringing the virtual environment back up in the event it needs to be completely shut down. The vMotion VMkernel ports are moved to the vDS to allow QoS marking of vMotion to be done at the VLAN level in the vDS if vMotion needs to have QoS policies applied in the future. The vMotion port group is also pinned to Cisco UCS fabric B. Pinning should be done in a vDS to ensure consistency across all ESXi hosts.

Configure the VMware vDS in vCenter for the VMware vSphere Web Client

To configure the vDS, follow these steps:

1. After logging into the VMware vSphere HTML5 Client, choose Networking under Menu.

2. Right-click the FlashStack-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 7.0.2 – ESXi 7.0.2 and later is selected and click NEXT.

5. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter VM-Traffic for the Port group name. Click NEXT.

6. Review the information and click FINISH to complete creating the vDS.

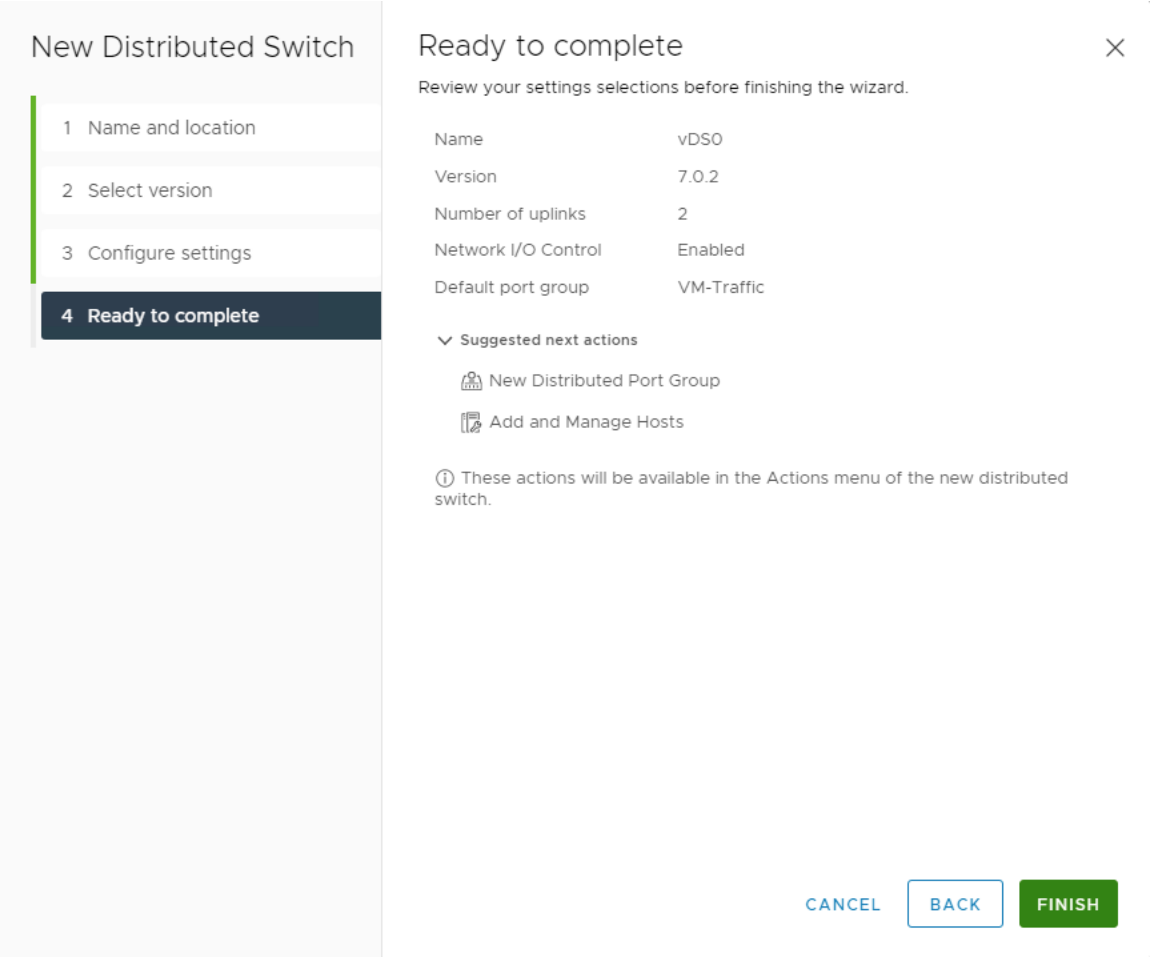

Ready to complete

Review your settings selections before finishing the wizard.

Name	vDS0
Version	7.0.2
Number of uplinks	2
Network I/O Control	Enabled
Default port group	VM-Traffic

These actions will be available in the Actions menu of the new distributed switch.

7. Expand the FlashStack-DC datacenter and the newly created vDS. Choose the newly created vDS.

8. Right-click the VM-Traffic port group and choose Edit Settings.

9. Choose VLAN.

10. Choose VLAN for VLAN type and enter the VM-Traffic VLAN ID. Click OK.

11. Right-click the vDS and choose Settings > Edit Settings.

12. In the Edit Settings window, choose Advanced.

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.

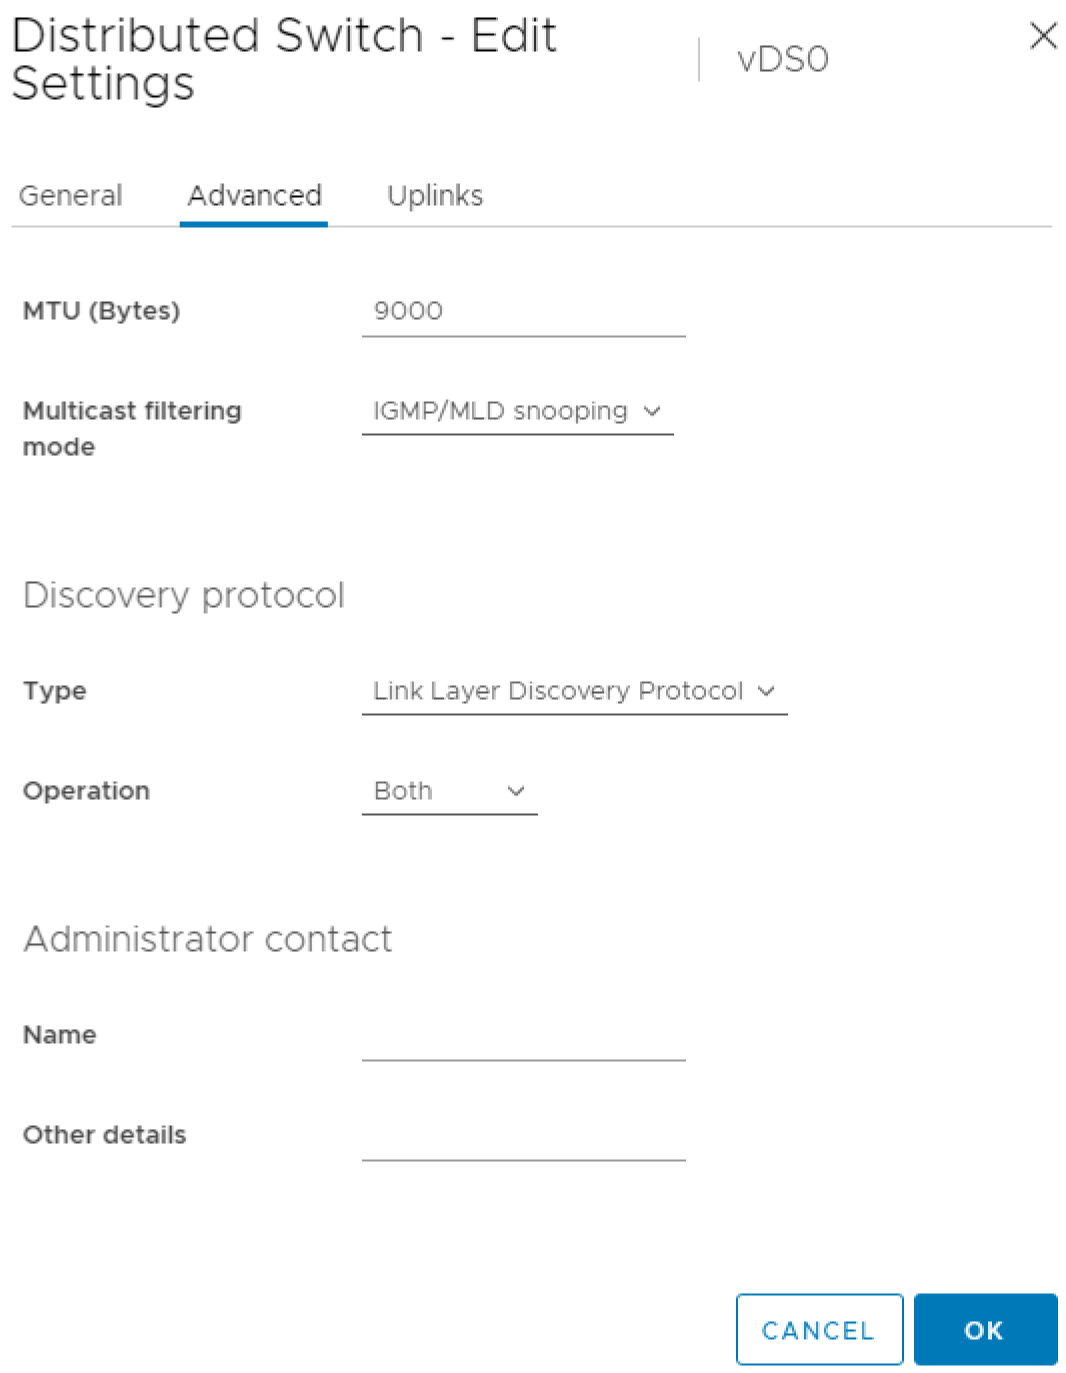

Distributed Switch - Edit Settings | vDS0

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 VMkernel-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 move down tab 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.

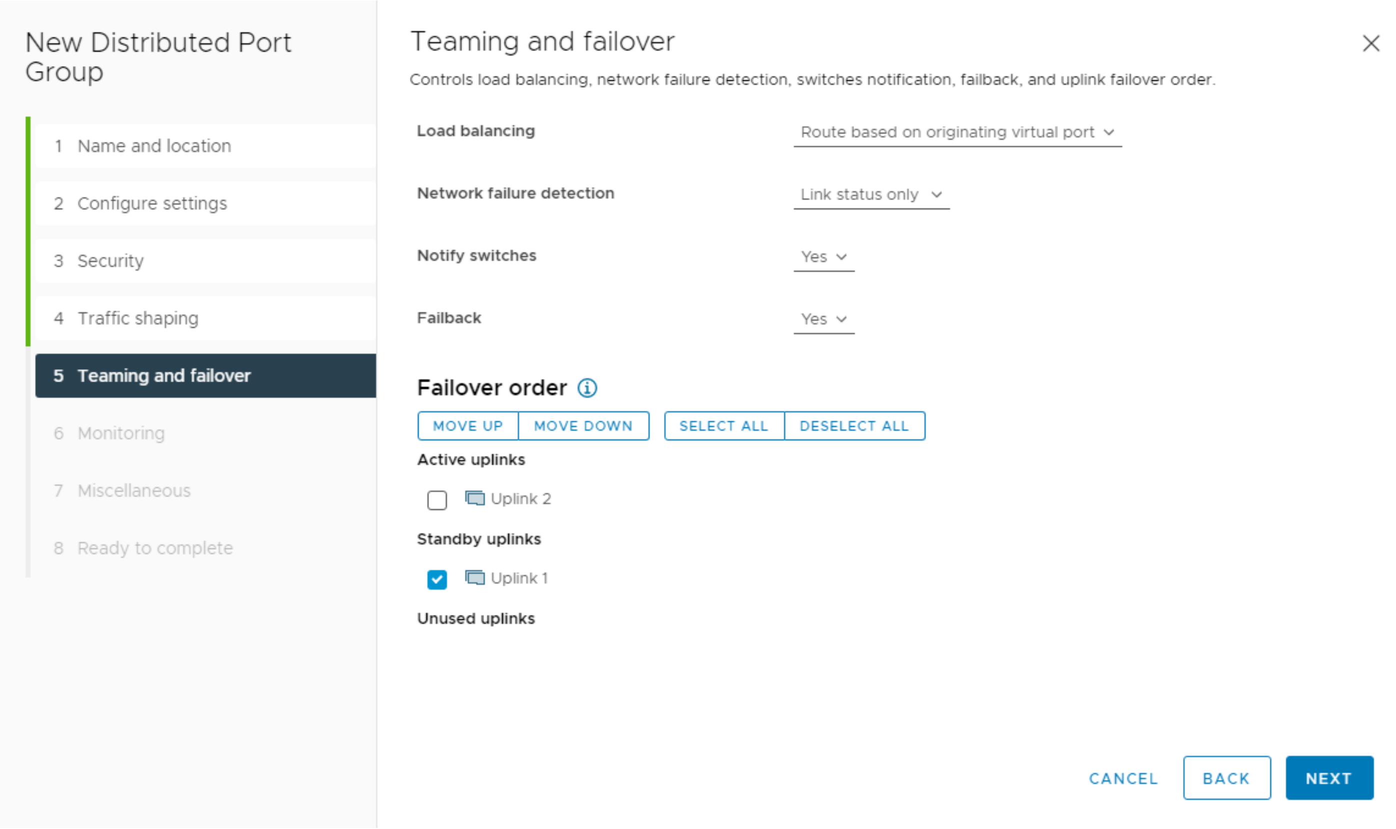

New Distributed Port Group

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. Right-click the vDS and choose Add and Manage Hosts.

25. Make sure Add hosts is selected and click NEXT.

26. Click the + sign to add New hosts. Choose the FlashStack ESXi hosts and click OK. Click NEXT.

27. Choose vmnic2 and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 and click Assign uplink. Choose Uplink 2 and click OK. If more than one host is being connected to the vDS, use the Apply this uplink assignment to the rest of the hosts checkbox.

Note: It is important to assign the uplinks as shown below. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.

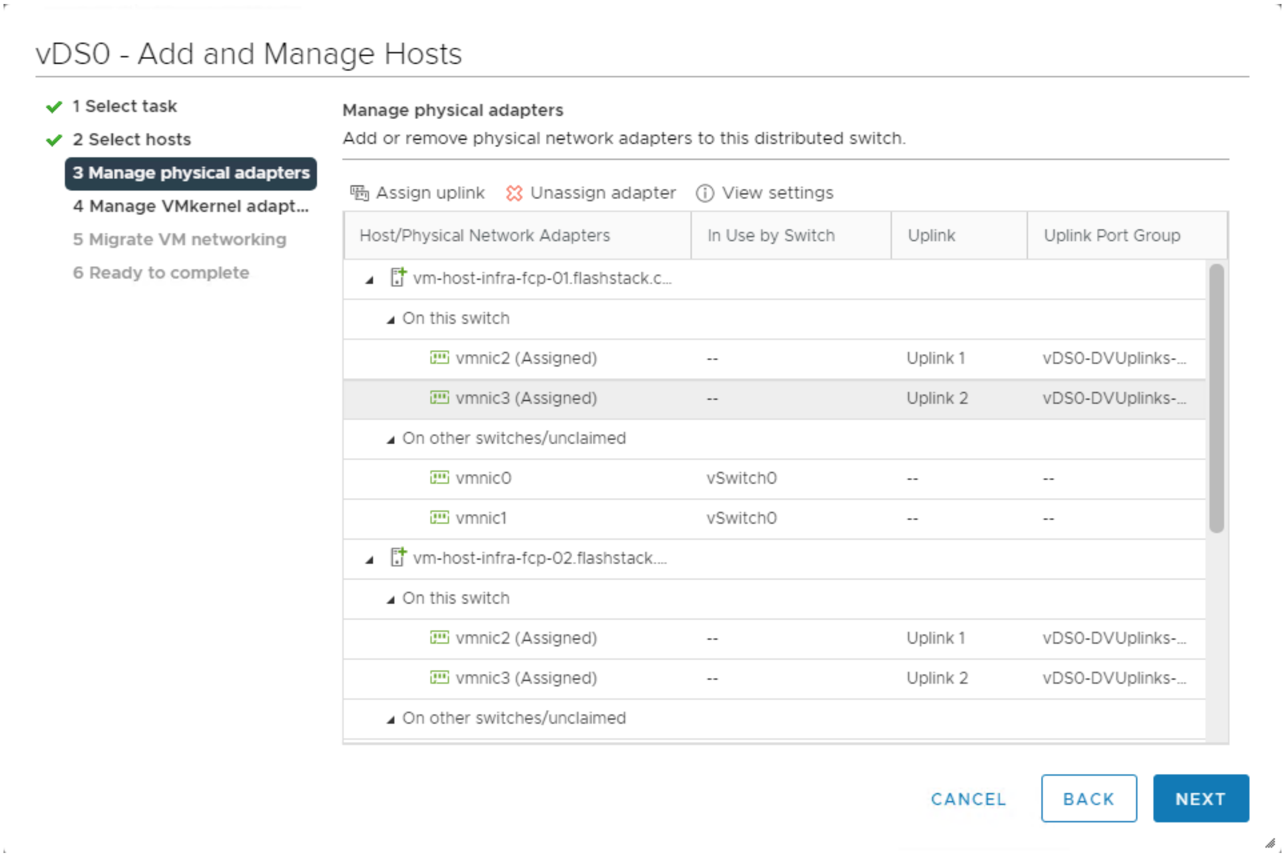

vDS0 - Add and Manage Hosts

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
vm-host-infra-fcp-01.flashstack.c...			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vm-host-infra-fcp-02.flashstack....			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			

28. Click NEXT.

29. Do not migrate any VMkernel ports and click NEXT.

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

31. Click FINISH to complete adding the ESXi host(s) to the vDS.

Add the vMotion VMkernel Port(s) to the ESXi Host

ESXi Host VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To add the vMotion VMkernel Port to the ESXi host(s) on the VMware vDS, follow these steps on the host:

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, click the Configure tab.

3. In the list under Networking, choose VMkernel adapters.

4. Choose Add Networking to Add host networking.

5. Make sure VMkernel Network Adapter is selected and click NEXT.

6. Choose BROWSE to the right of Select an existing network.

7. Choose vMotion on the vDS and click OK.

8. Click NEXT.

9. Make sure the Network label is vMotion with the vDS in parenthesis. From the drop-down list, select Custom for MTU and make sure the MTU is set to 9000. Choose the vMotion TCP/IP stack and click NEXT.

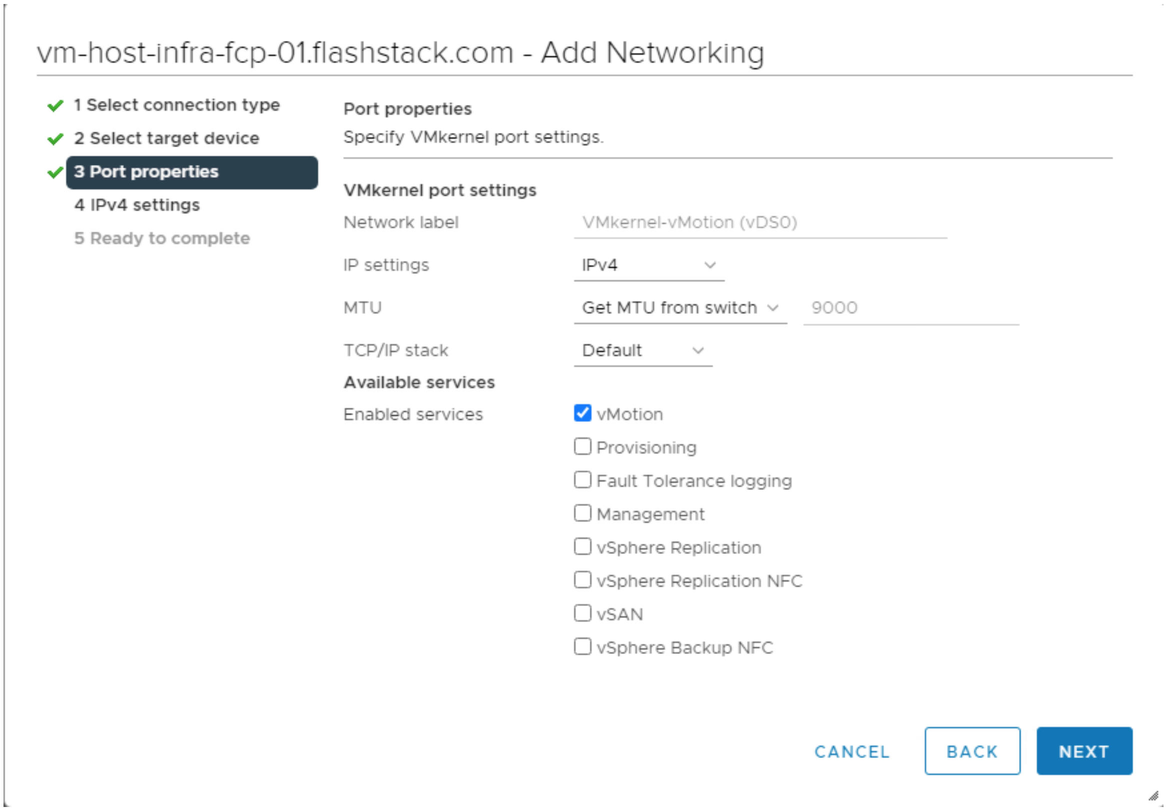

vm-host-infra-fcp-01.flashstack.com - Add Networking

Port properties

Specify VMkernel port settings.

VMkernel port settings

Network label	VMkernel-vMotion (vDS0)
IP settings	IPv4
MTU	Get MTU from switch 9000
TCP/IP stack	Default

Available services

Enabled services

☑ vMotion
□ Provisioning
□ Fault Tolerance logging
□ Management
□ vSphere Replication
□ vSphere Replication NFC
□ vSAN
□ vSphere Backup NFC

10. Choose Use static IPv4 settings and input the host’s vMotion IPv4 address and Subnet mask.

11. Click NEXT.

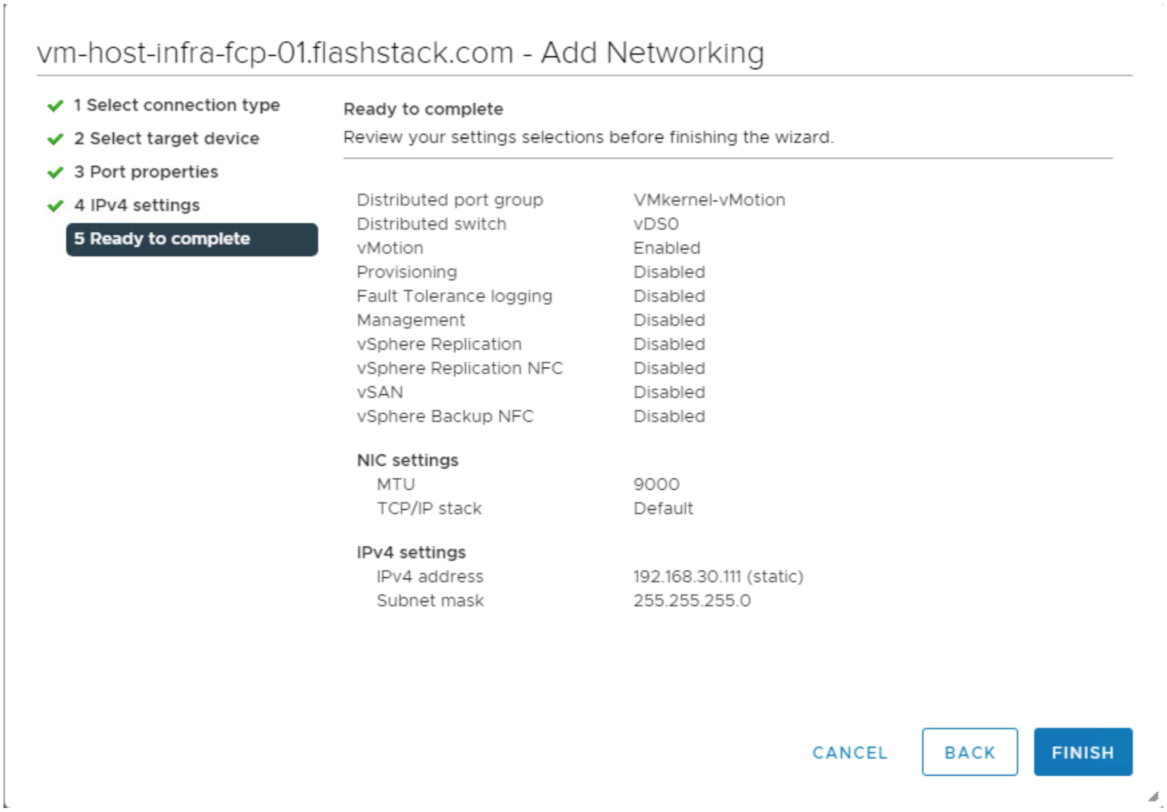

vm-host-infra-fcp-01.flashstack.com - Add Networking

Ready to complete

Review your settings selections before finishing the wizard.

Distributed port group	VMkernel-vMotion
Distributed switch	vDS0
vMotion	Enabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled
vSphere Backup NFC	Disabled
NIC settings	
MTU	9000
TCP/IP stack	Default
IPv4 settings	
IPv4 address	192.168.30.111 (static)
Subnet mask	255.255.255.0

12. Review the parameters and click FINISH to add the vMotion VMkernel port.

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 7.0 U2 installed, the management IP address set, the nfnic driver updated and the Cisco UCS Tool 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 Hosts to vCenter

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 “FlashStack-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).

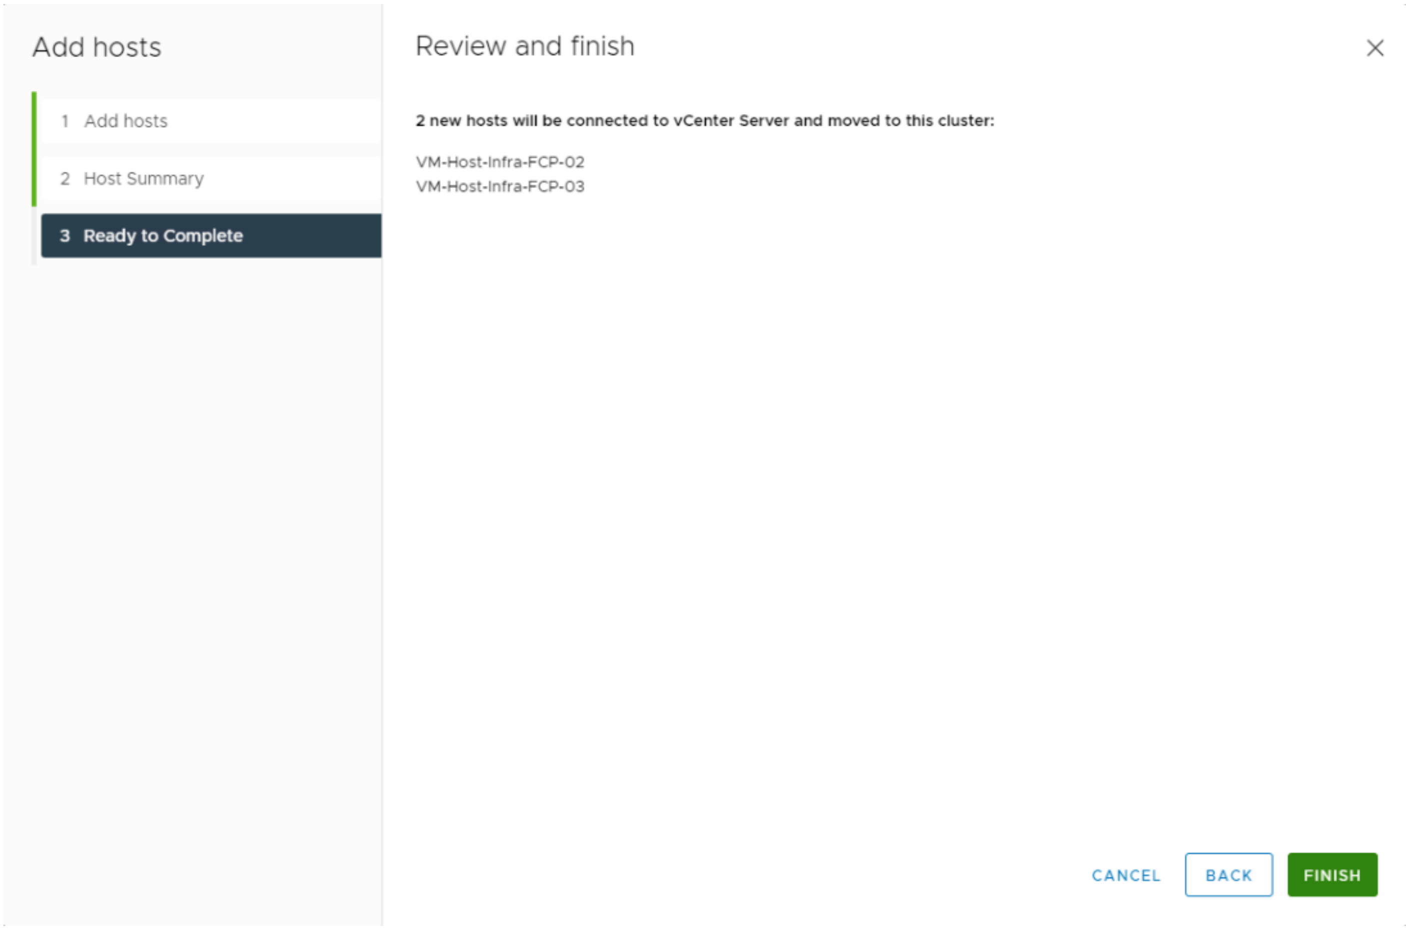

2 new hosts will be connected to vCenter Server and moved to this cluster:

VM-Host-Infra-FCP-02
VM-Host-Infra-FCP-03

The added ESXi host(s) will be placed in Maintenance Mode and will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.

Set Up VMkernel Ports and Virtual Switch for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-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.

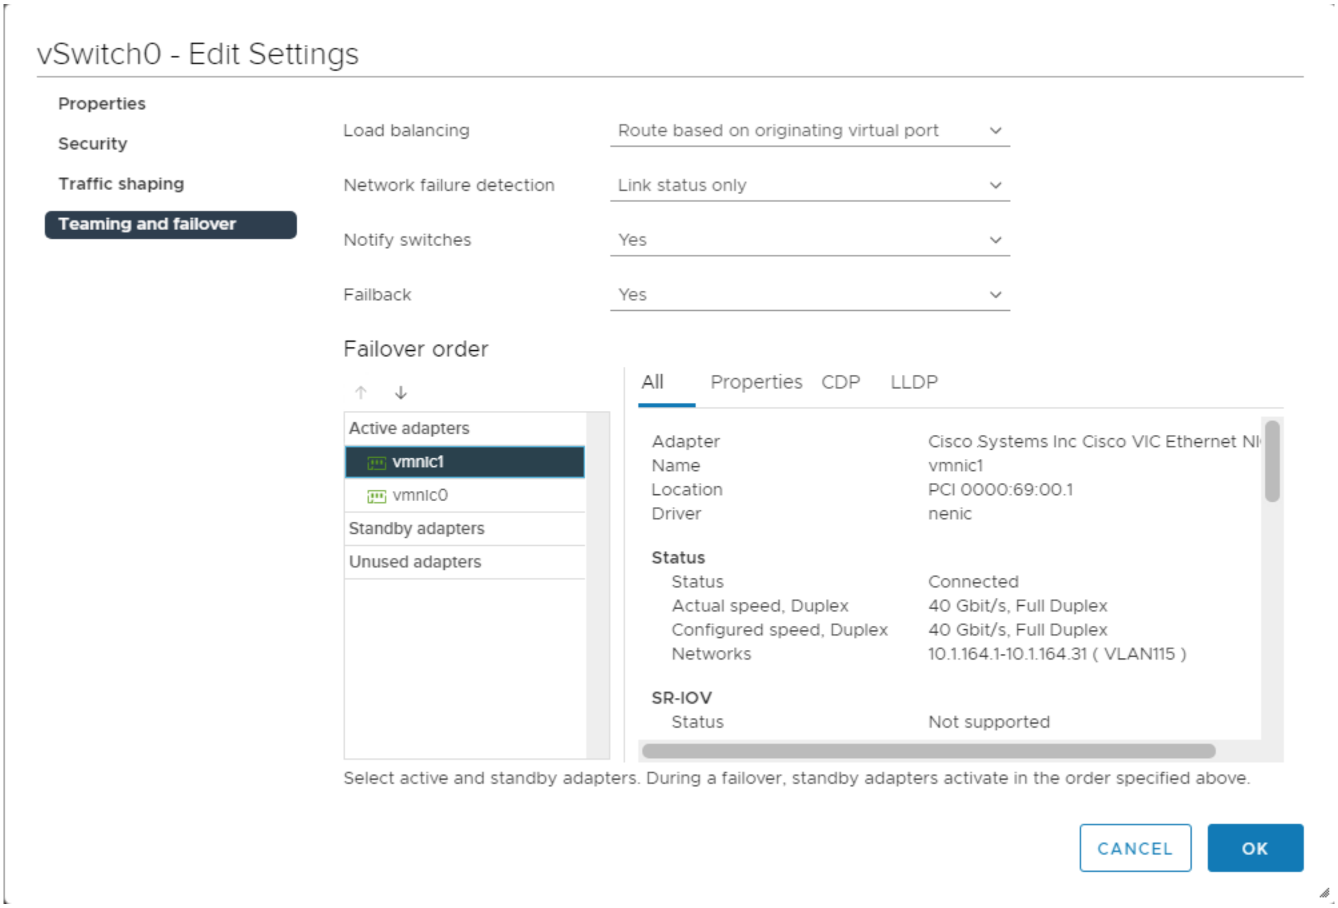

vSwitch0 - Edit Settings

Load balancing: Route based on originating virtual port

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active adapters
- vmnic1
- vmnic0

Standby adapters

Unused adapters

Adapter	Cisco Systems Inc Cisco VIC Ethernet NI
Name	vmnic1
Location	PCI 0000:69:00.1
Driver	nenic
Status	
Status	Connected
Actual speed, Duplex	40 Gbit/s, Full Duplex
Configured speed, Duplex	40 Gbit/s, Full Duplex
Networks	10.1.164.1-10.1.164.31 (VLAN115)
SR-IOV	
Status	Not supported

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

9. Click OK.

10. In the center pane, to the right of VM Network click ... > Remove to remove the port group. Click YES on the confirmation.

11. Click ADD NETWORKING to add a new VM port group.

12. Choose Virtual Machine Port Group for a Standard Switch and click NEXT.

13. Ensure vSwitch0 is shown for Select an existing standard switch and click NEXT.

14. Name the port group “IB-MGMT Network” and leave the VLAN ID field set to None (0). Click NEXT.

Note: In the Cisco UCS section of this document, the IB-MGMT VLAN was set as the native VLAN for the vSwitch0 vNIC templates, allowing DHCP to be used on ESXi vmk0 without putting in a VLAN ID for this port. Since this port group is in the same VLAN, the port group’s VLAN ID should also be set to 0.

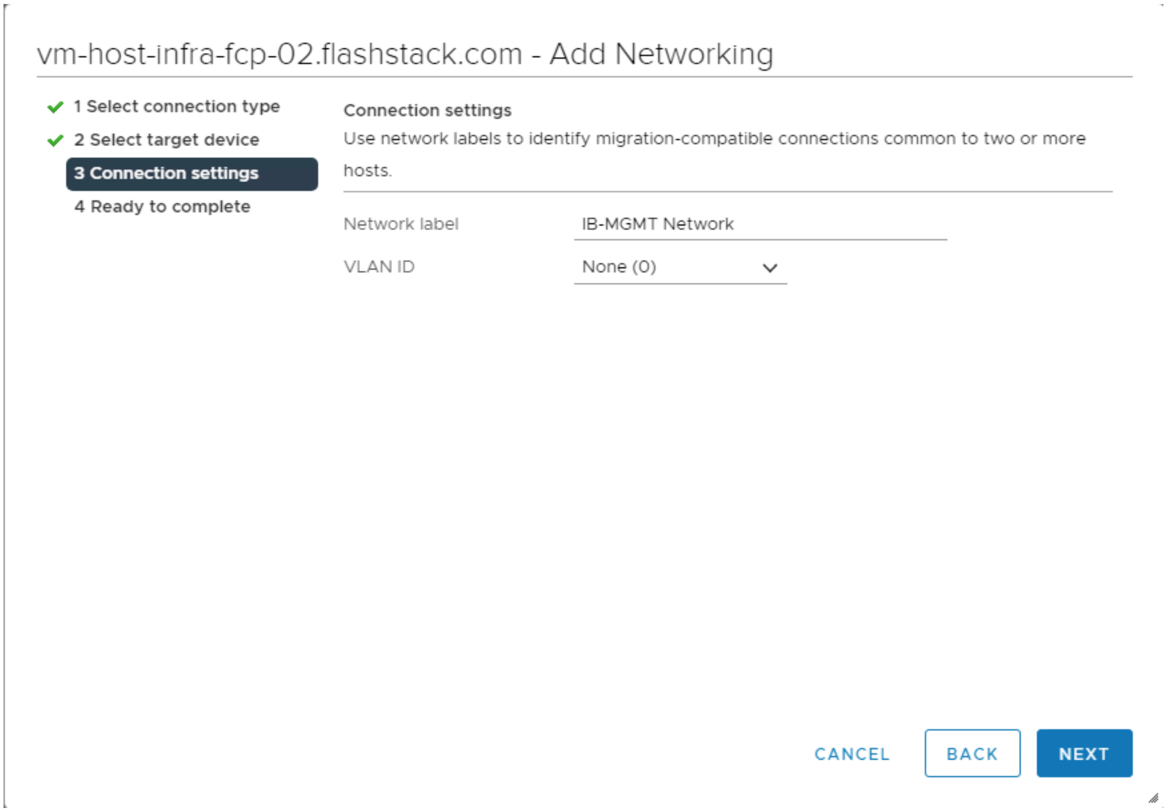

vm-host-infra-fcp-02.flashstack.com - Add Networking

Connection settings

Use network labels to identify migration-compatible connections common to two or more hosts.

Network label: IB-MGMT Network

VLAN ID: None (0)

15. Click FINISH to complete adding the IB-MGMT Network VM port group.

16. Click ADD NETWORKING to add a new VM port group.

17. Choose Virtual Machine Port Group for a Standard Switch and click NEXT.

18. Ensure vSwitch0 is shown for Select an existing standard switch and click NEXT.

19. Name the port group “OOB-MGMT Network” and input <OOB-MGMT-vlan-id> for the VLAN ID field. Click NEXT.

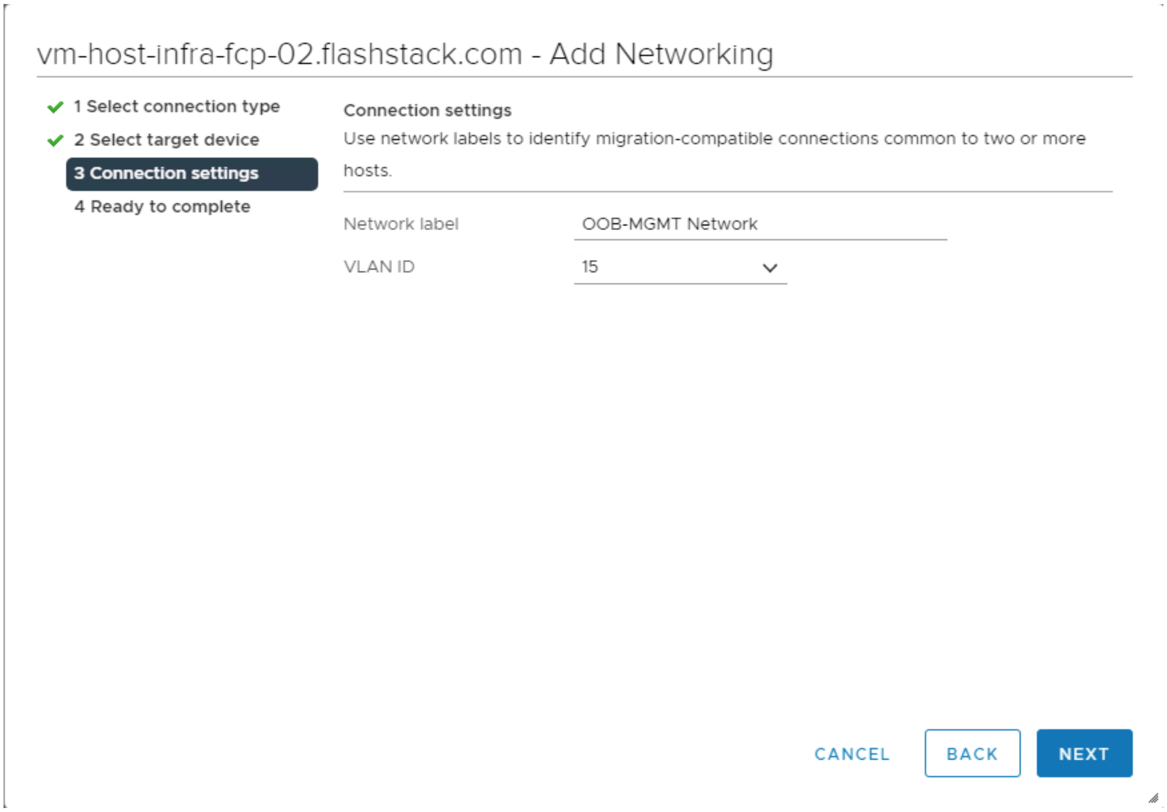

vm-host-infra-fcp-02.flashstack.com - Add Networking

Connection settings

Use network labels to identify migration-compatible connections common to two or more hosts.

Network label	OOB-MGMT Network
VLAN ID	15

20. Click FINISH to complete adding the OOB-MGMT Network VM port group.

21. Under Networking, choose Virtual switches. Expand vSwitch0. The properties for vSwitch0 should be like the following example:

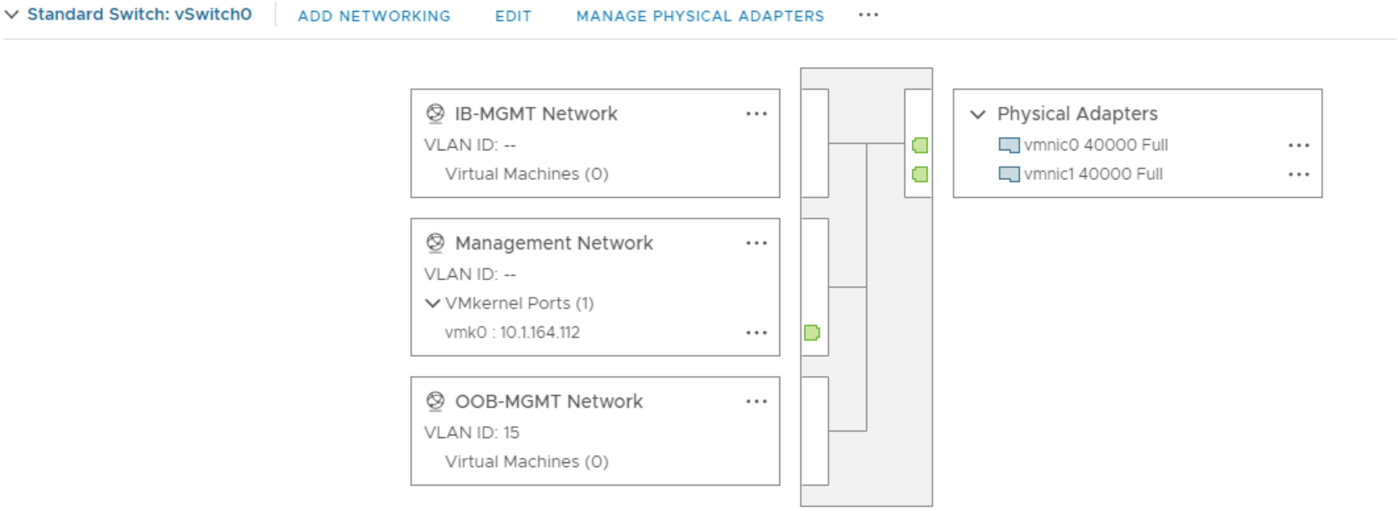

IB-MGMT Network
VLAN ID: --
Virtual Machines (0)

Management Network
VLAN ID: --
VMkernel Ports (1)
vmk0 : 10.1.164.112

OOB-MGMT Network
VLAN ID: 15
Virtual Machines (0)

Physical Adapters
vmnic0 40000 Full
vmnic1 40000 Full

22. Repeat steps 1-21 for all hosts being added.

Mount Required Datastores for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To mount the required datastores, follow these steps on the ESXi host(s):

1. From the vCenter Home screen, choose Menu > Storage.

2. Expand FlashStack-DC.

3. Located on the left, right-click Infra-DataStore1 and choose Mount Datastore to Additional Hosts.

4. Choose the ESXi host(s) and click OK.

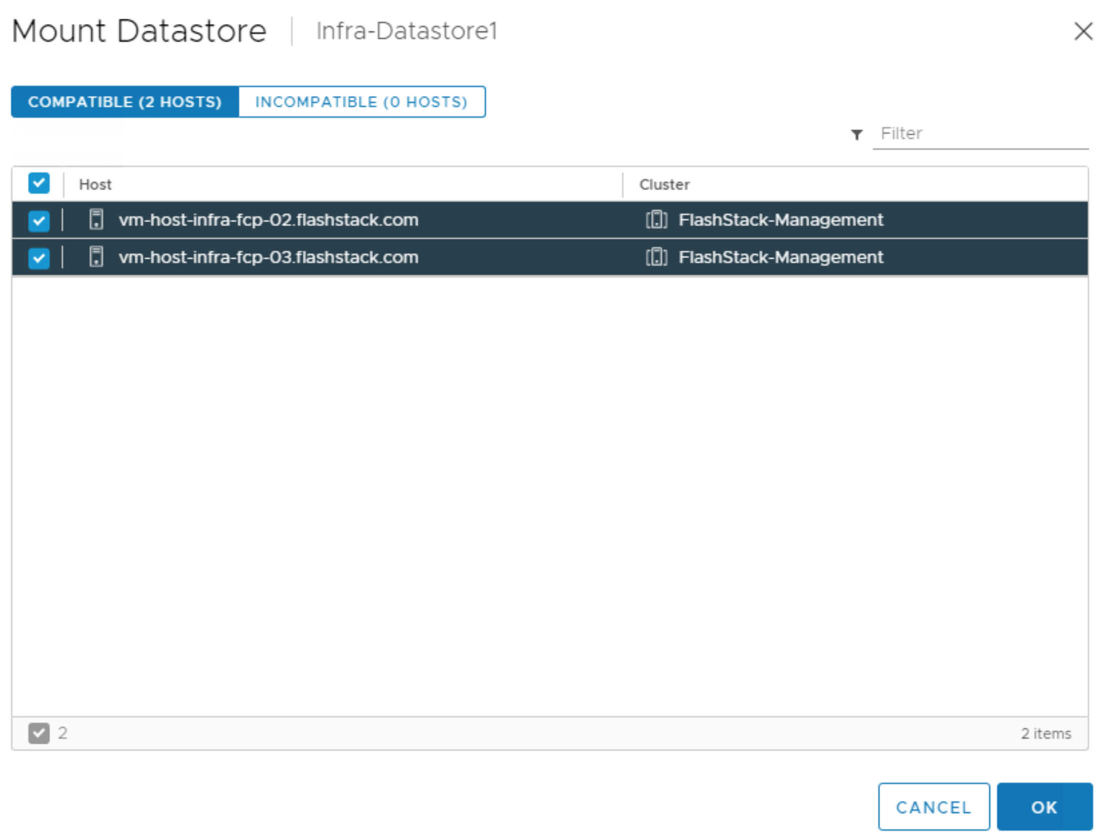

☑	Host	Cluster
☑	vm-host-infra-fcp-02.flashstack.com	FlashStack-Management
☑	vm-host-infra-fcp-03.flashstack.com	FlashStack-Management

☑ 2 2 items

5. Repeat steps 1-4 to mount the Infra-Swap datastore to the ESXi host(s).

6. Choose Infra-DataStore1. In the center pane, choose Hosts. Verify the ESXi host(s) now has the datastore mounted. Repeat this process to also verify that Infra-Swap is also mounted.

Configure NTP on ESXi Host for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

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

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, choose the Configure tab.

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

4. To the right of Manual Time Configuration, click EDIT.

5. Set the correct local time and click OK.

6. To the right of Network Time Protocol, click EDIT.

7. Choose the Enable checkbox.

8. Enter the two Nexus switch NTP IP addresses in the NTP servers box separated by a comma.

9. Click the Start NTP Service checkbox.

10. Use the drop-down list to choose Start and stop with host.

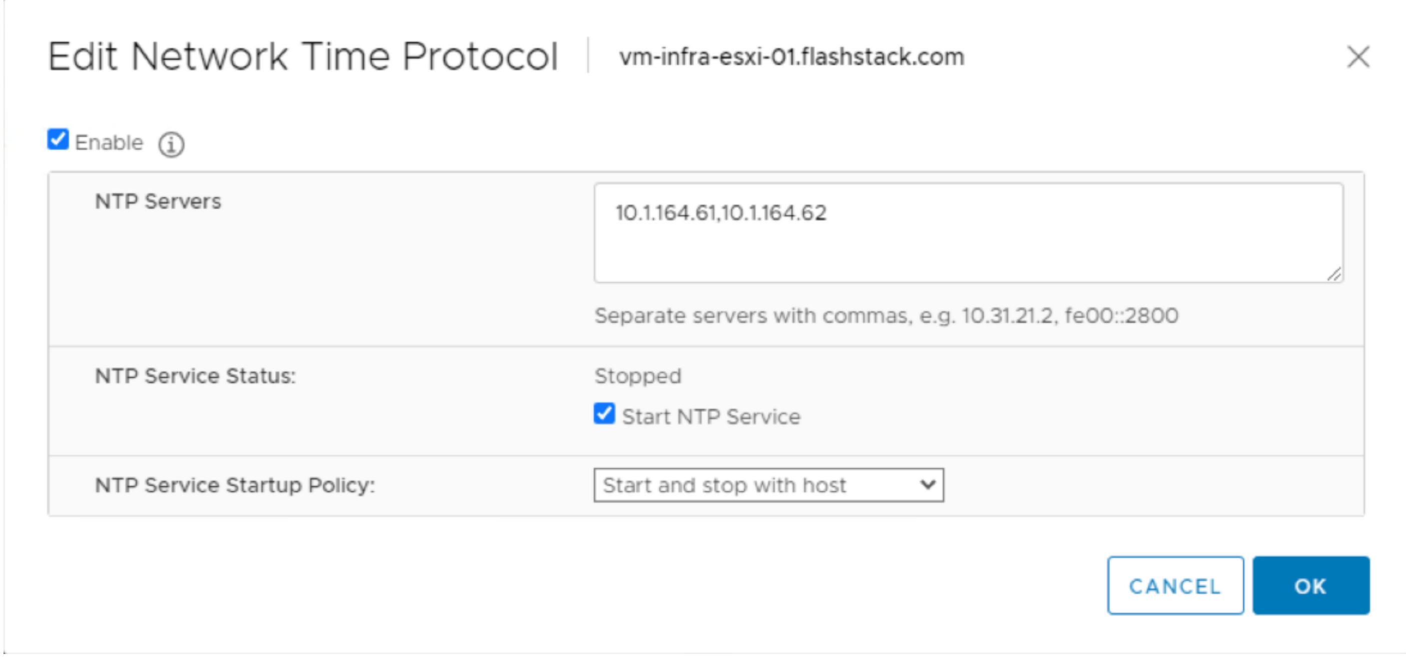

Edit Network Time Protocol | vm-infra-esxi-01.flashstack.com

☑ Enable

NTP Servers: 10.1.164.61,10.1.164.62

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.

Change ESXi Power Management Policy for Cisco UCS M6 Hosts for the ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To change the ESXi power management policy for the Cisco UCS M6 hosts, follow these steps:

Note: Implementation of this policy is recommended in Performance Tuning for Cisco UCS M6 Server with Intel 3rd Gen Processors for maximum performance. If your organization has specific power policies, please set this policy accordingly.

1. In the list under Hardware, choose Overview. Scroll to the bottom and to the right of Power Management, choose EDIT POWER POLICY.

2. Choose High performance and click OK.

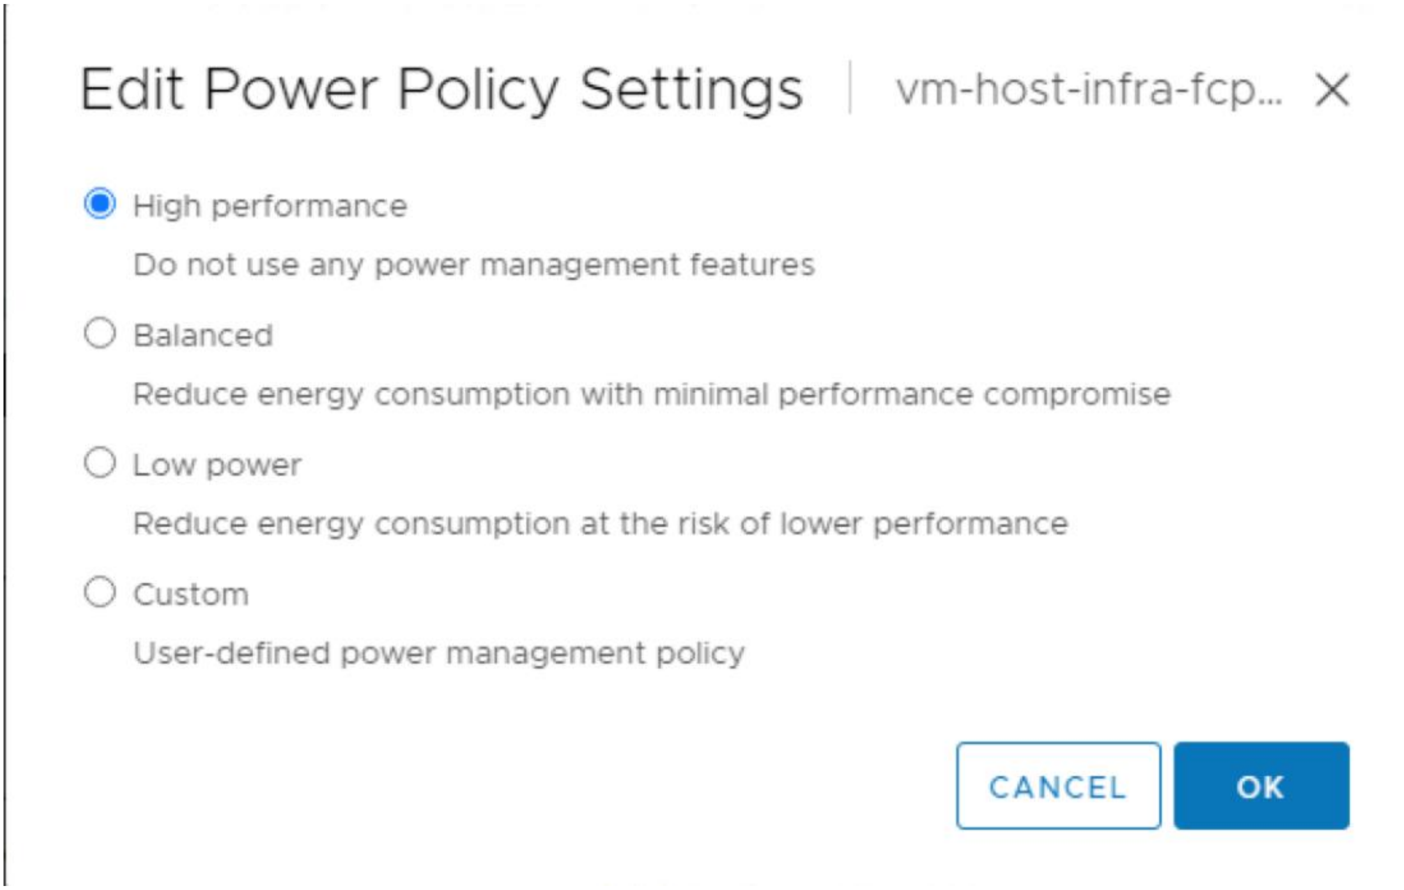

Edit Power Policy Settings | vm-host-infra-fcp...

● 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

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

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

1. In the list under Storage, choose Storage Devices. Make sure the Pure FlashArray Fibre Channel Disk is selected.

2. Choose the Paths tab.

3. Ensure that 4 fibre channel paths appear, all four should have the status Active (I/O).

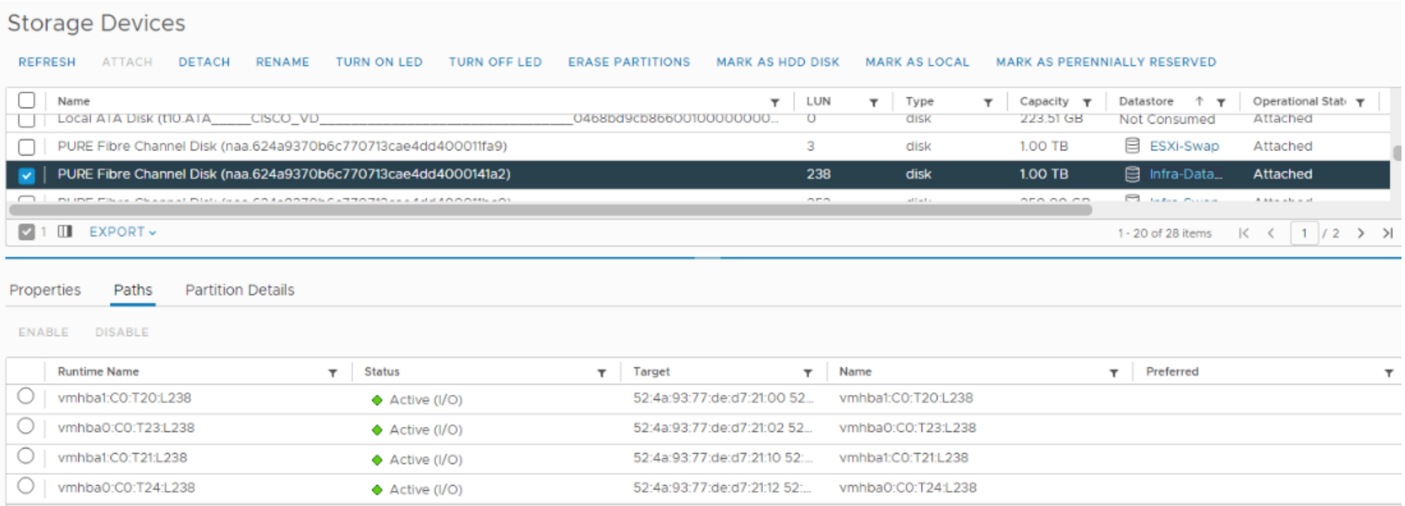

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS HDD DISK MARK AS LOCAL MARK AS PERENNIALLY RESERVED

☐	Name	LUN	Type	Capacity	Datastore	Operational Stat
☐	Local ATA Disk (t10.ATA_____CISCO_VD______________________________0468bd9cb86600100000000...	0	disk	223.51 GB	Not Consumed	Attached
☐	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd400011fa9)	3	disk	1.00 TB	ESXi-Swap	Attached
☑	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd4000141a2)	238	disk	1.00 TB	Infra-Data...	Attached

☑ 1 EXPORT 1 - 20 of 28 items 1 / 2

Properties Paths Partition Details

ENABLE DISABLE

	Runtime Name	Status	Target	Name	Preferred
○	vmhba1:C0:T20:L238	Active (I/O)	52:4a:93:77:de:d7:21:00 52...	vmhba1:C0:T20:L238	
○	vmhba0:C0:T23:L238	Active (I/O)	52:4a:93:77:de:d7:21:02 52...	vmhba0:C0:T23:L238	
○	vmhba1:C0:T21:L238	Active (I/O)	52:4a:93:77:de:d7:21:10 52:...	vmhba1:C0:T21:L238	
○	vmhba0:C0:T24:L238	Active (I/O)	52:4a:93:77:de:d7:21:12 52:...	vmhba0:C0:T24:L238	

Add the ESXi Host(s) to the VMware Virtual Distributed Switch to the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

Follow this procedure if there are hosts to be added to vDS, skip if already added earlier. 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 FlashStack 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.

Note: It is important to assign the uplinks as shown below. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.

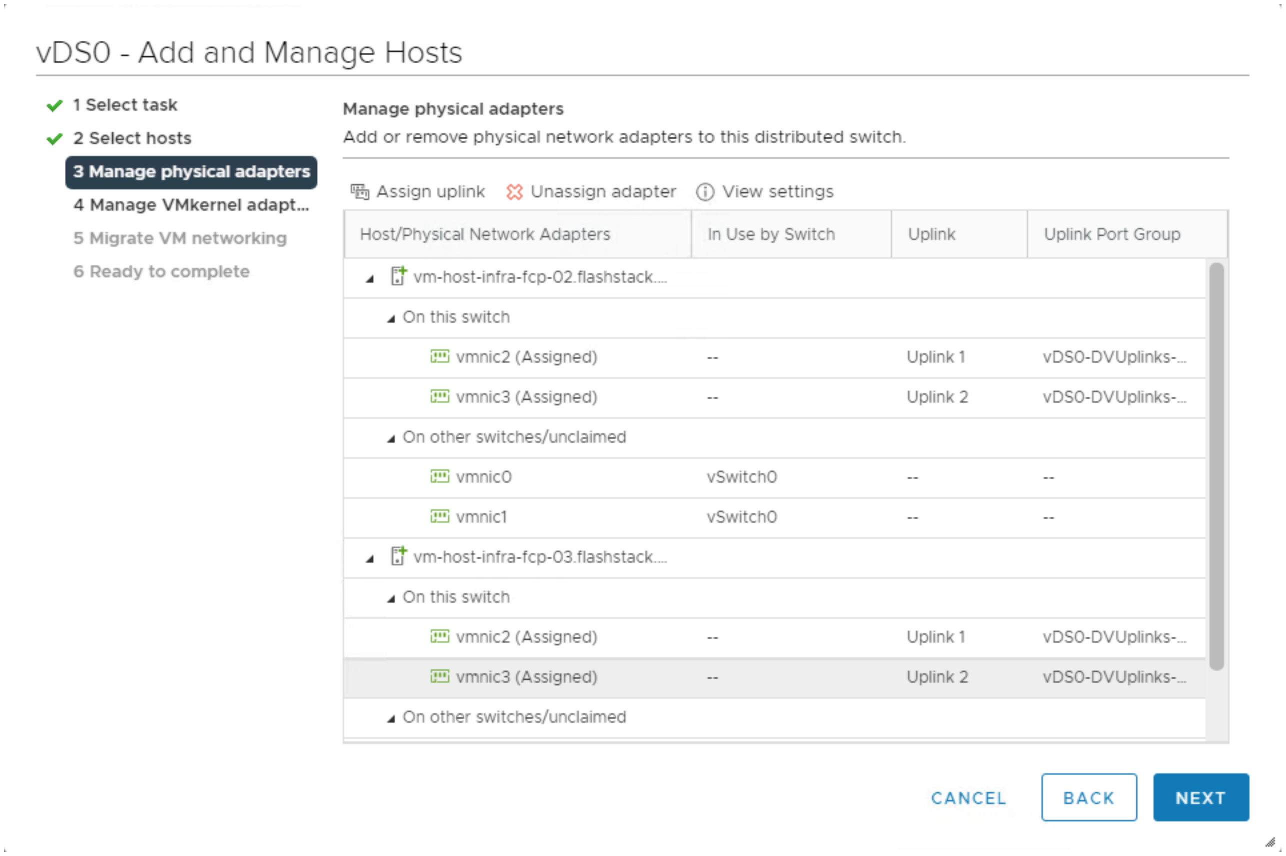

vDS0 - Add and Manage Hosts

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
vm-host-infra-fcp-02.flashstack....			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vm-host-infra-fcp-03.flashstack....			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			

6. Click NEXT.

7. Do not migrate any VMkernel ports and click NEXT.

8. Do not migrate any VM ports and click NEXT.

9. Click FINISH to complete adding the ESXi host(s) to the vDS.

Add the vMotion VMkernel Port(s) to the ESXi Host to the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To add the vMotion VMkernel Port to the ESXi host(s) on the VMware vDS, follow these steps on the host:

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.

2. In the center pane, click the Configure tab.

3. In the list under Networking, choose VMkernel adapters.

4. Choose Add Networking to Add host networking.

5. Make sure VMkernel Network Adapter is selected and click NEXT.

6. Choose BROWSE to the right of Select an existing network.

7. Choose vMotion on the vDS and click OK.

8. Click NEXT.

9. Make sure the Network label is vMotion with the vDS in parenthesis. From the drop-down list, select Custom for MTU and make sure the MTU is set to 9000. Choose the vMotion TCP/IP stack and click NEXT.

10. Choose Use static IPv4 settings and input the host's vMotion IPv4 address and Subnet mask.

11. Click NEXT.

12. Review the parameters and click FINISH to add the vMotion VMkernel port.

13. If this is an iSCSI-booted host, execute the instructions in the Appendix for an iSCSI-booted host being added in vCenter.

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

VMware ESXi 7.0 U2 TPM Attestation

If your Cisco UCS servers have Trusted Platform Module (TPM) 2.0 modules installed, the TPM can provide assurance that ESXi has booted with UEFI Secure Boot enabled and using only digitally signed code. In the Cisco UCS section of this document, UEFI secure boot was enabled in the boot policy. A server can boot with UEFI Secure Boot with or without a TPM 2.0 module. If it has a TPM, VMware vCenter can attest that the server booted with UEFI Secure Boot. Follow these steps:

1. If your Cisco UCS servers have TPM 2.0 modules installed, TPM Attestation can be verified in the vSphere HTML5 Client. To get to the HTML5 client from the Web Client, click "Launch vSphere Client (HTML5) in the upper center portion of the Web Client window.

2. From the Hosts and Clusters window in the vSphere Client, click the FlashStack-Management cluster. In the center pane, click Monitor > Security. The Attestation status will appear as shown below, where 2 of the 3 hosts have TPM 2.0 modules installed:

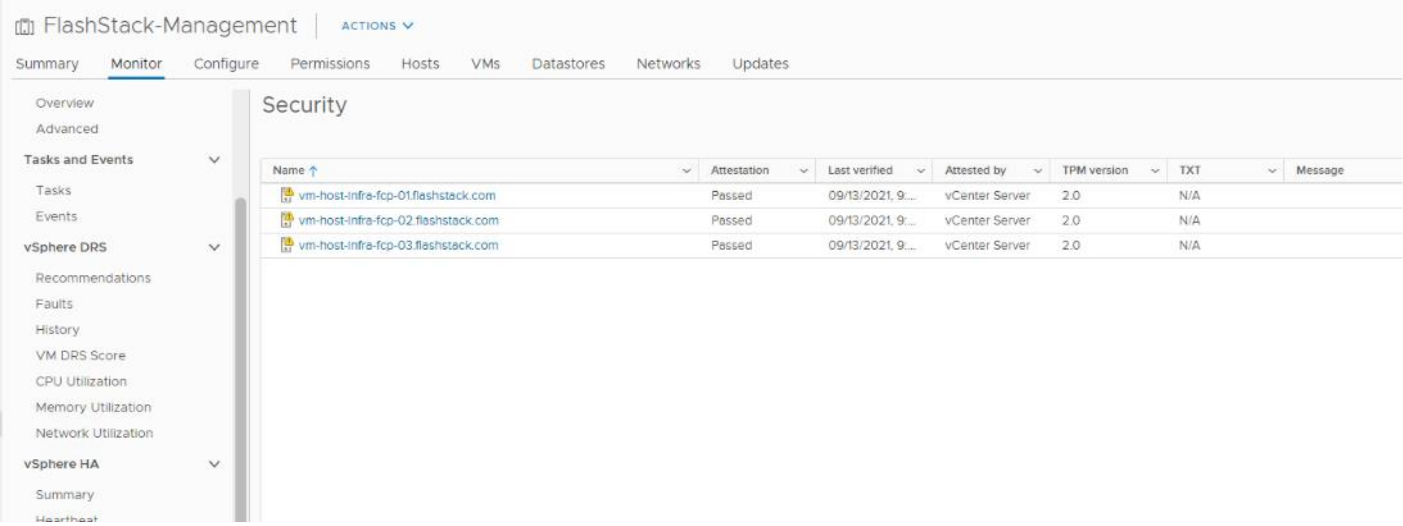

Security

Name ↑	Attestation	Last verified	Attested by	TPM version	TXT	Message
vm-host-infra-fcp-01.flashstack.com	Passed	09/13/2021, 9:...	vCenter Server	2.0	N/A	
vm-host-infra-fcp-02.flashstack.com	Passed	09/13/2021, 9:...	vCenter Server	2.0	N/A	
vm-host-infra-fcp-03.flashstack.com	Passed	09/13/2021, 9:...	vCenter Server	2.0	N/A	

Note: It may be necessary to disconnect and reconnect a host from vCenter to get it to pass attestation the first time. Also, in this example, only the second host had a TPM module installed.

FlashStack Management Tools Setup

Cisco Data Center Network Manager (DCNM)-SAN

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

Prerequisites

The following prerequisites need to be configured:

1. Licensing. Cisco DCNM-SAN includes a 60-day server-based trial license that can be used to monitor and configure Cisco MDS Fibre Channel switches and monitor Cisco Nexus switches. Both DCNM server-based and switch-based licenses can be purchased. Additionally, SAN Insights and SAN Analytics requires an additional switch-based license on each switch. Cisco MDS 32Gbps Fibre Channel switches provide a 120-day grace period to trial SAN Analytics.

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

2. Passwords. Cisco DCNM-SAN passwords should adhere to the following password requirements:

    a. It must be at least eight characters long and contain at least one alphabet and one numeral.

    b. It can contain a combination of alphabets, numerals, and special characters.

    c. Do not use any of these special characters in the DCNM password for all platforms: <SPACE> " & $ % ' ^ = < > ; : ` \ | / , .*

3. DCNM SNMPv3 user on switches. Each switch (both Cisco MDS and Nexus) needs an SNMPv3 user added for DCNM to use to query and configure the switch. On each switch, enter the following command in configure terminal mode (in the example, the userid is snmpuser): snmp-server user snmpadmin network-admin auth sha <password> priv aes-128 <privacy-password>

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

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

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

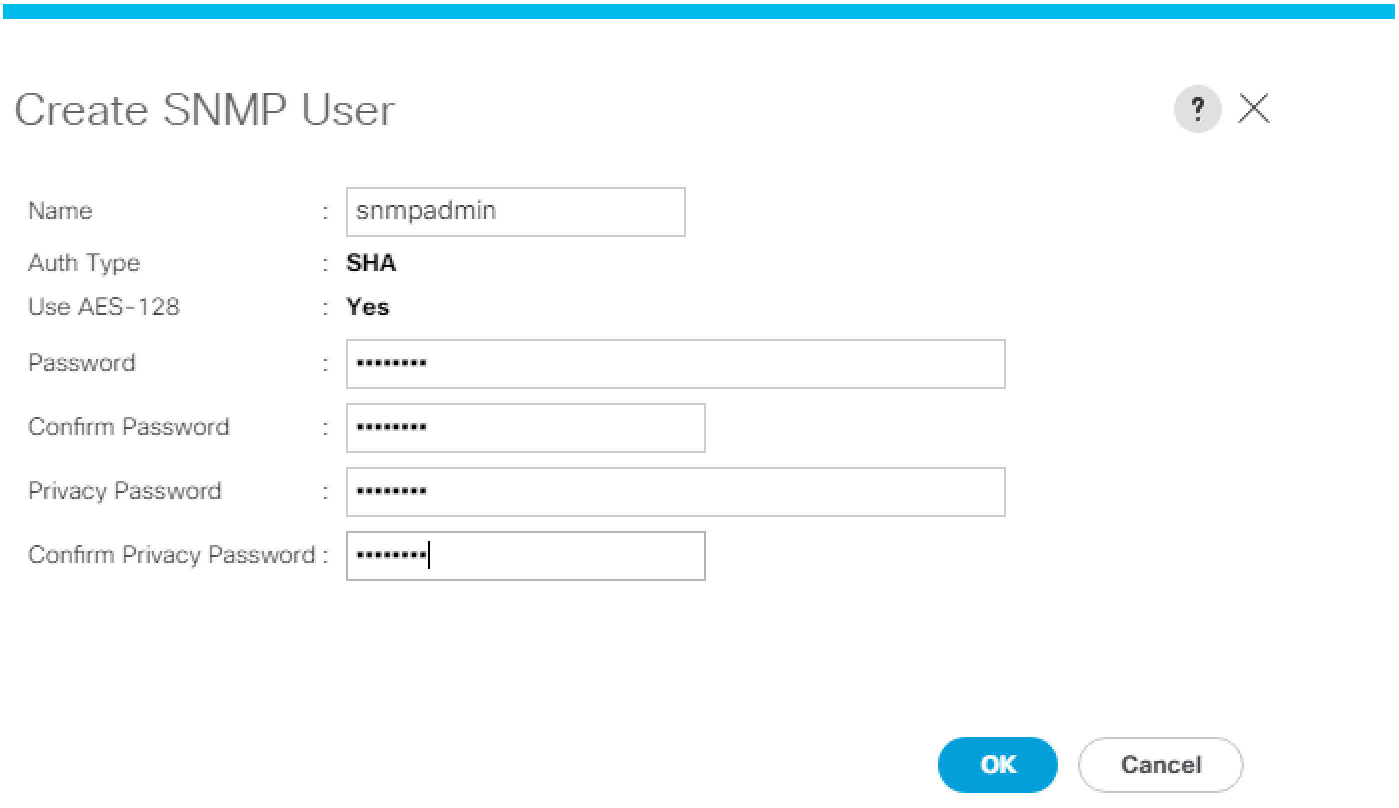

Create SNMP User

Name :	snmpadmin
Auth Type :	SHA
Use AES-128 :	Yes
Password :	•••••••••
Confirm Password :	•••••••••
Privacy Password :	•••••••••
Confirm Privacy Password :	•••••••••

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

Deploy the Cisco DCNM-SAN OVA

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

1. Download the Cisco DCNM 11.5.1 Open Virtual Appliance for VMware from https://software.cisco.com/download/home/281722751/type/282088134/release/11.5(1). Extract dcnm-va.11.5.1.ova from the ZIP file.

2. In the VMware vCenter HTML5 interface, click Menu > Hosts and Clusters.

3. Right-click the FlashStack-Management cluster and select Deploy OVF Template.

4. Choose Local file then click UPLOAD FILES. Navigate to choose dcnm-va.11.5.1.ova and click Open. Click NEXT.

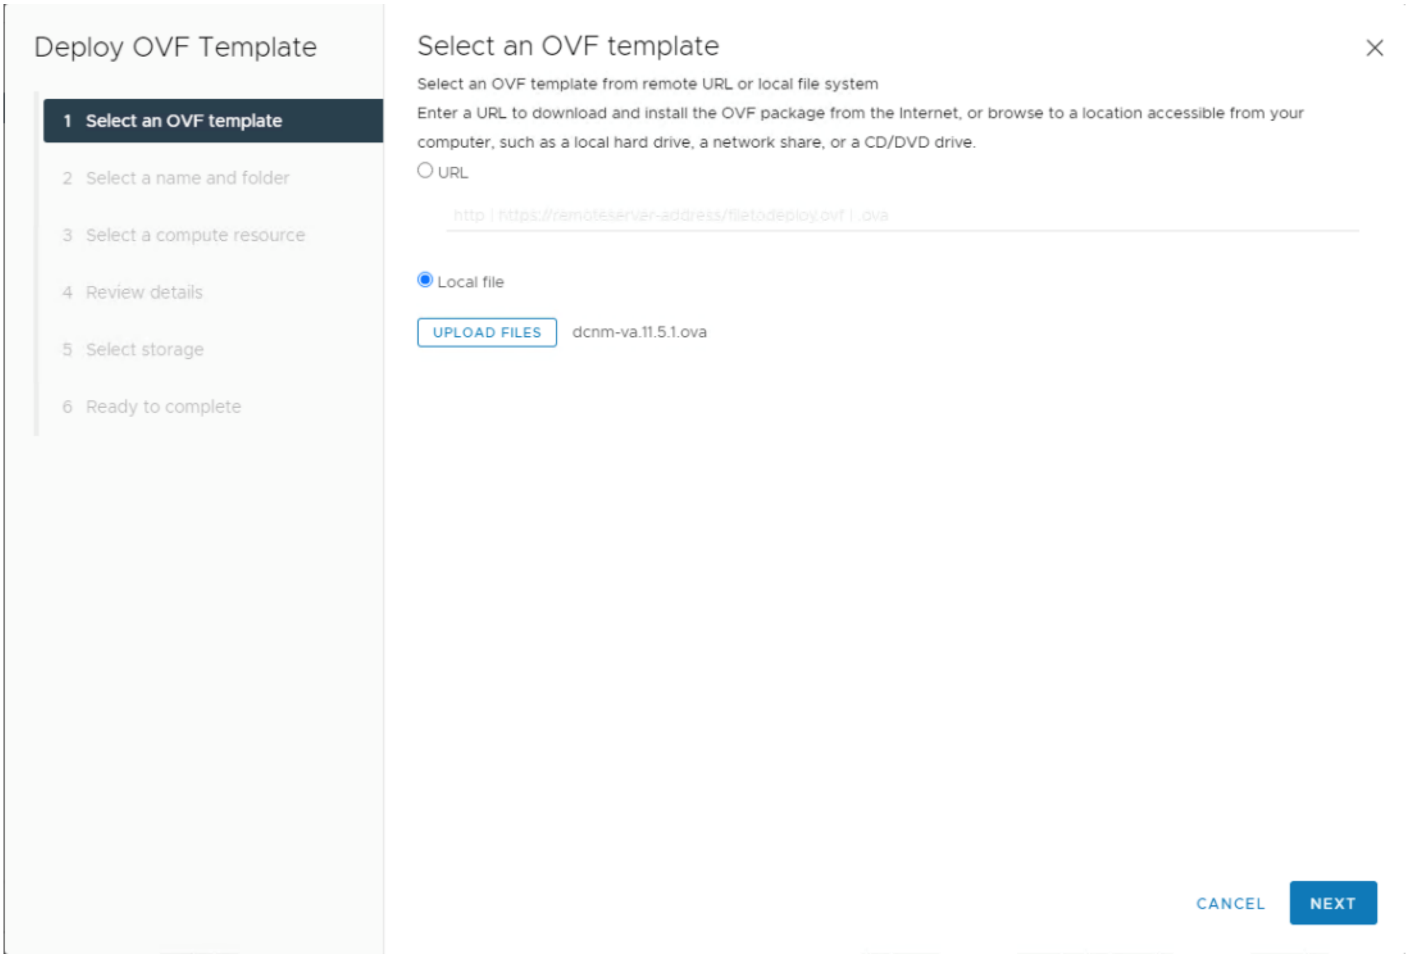

Deploy OVF Template

Select an OVF template

Select an OVF template from remote URL or local file system

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

○ URL

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

● Local file

UPLOAD FILES dcnm-va.11.5.1.ova

5. Name the virtual machine and choose the FlashStack-DC datacenter. Click NEXT.

6. Choose the FlashStack-Management cluster and click NEXT.

7. Review the details and click NEXT.

8. Scroll through and accept the license agreements. Click NEXT.

9. Choose the appropriate deployment configuration size and click NEXT.

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

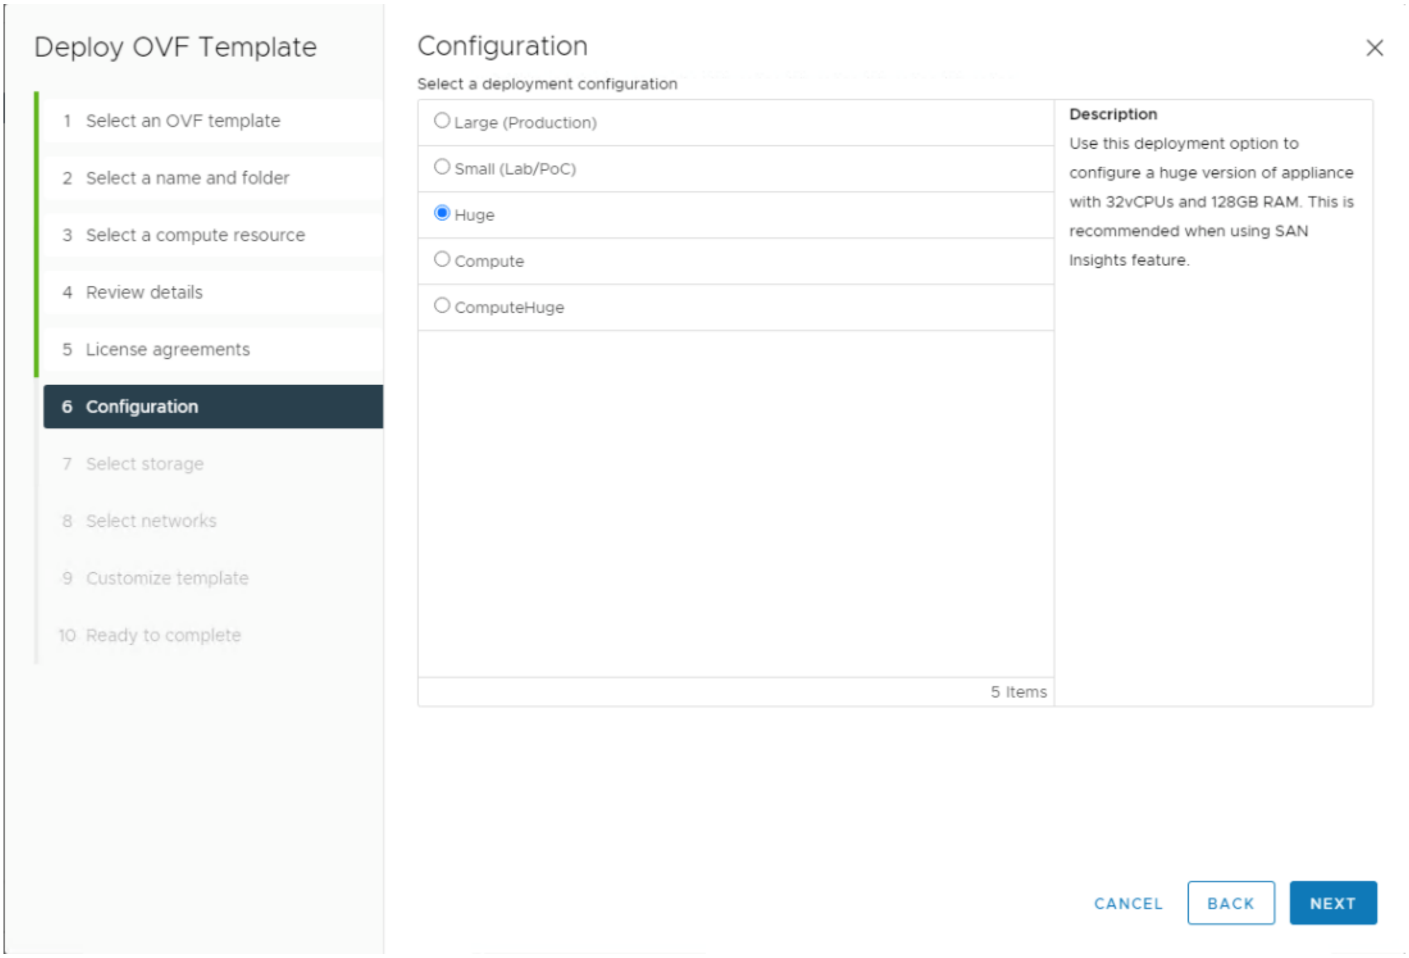

Configuration

Select a deployment configuration

○ Large (Production)
○ Small (Lab/PoC)
● Huge
○ Compute
○ ComputeHuge

5 Items

Description

Use this deployment option to configure a huge version of appliance with 32vCPUs and 128GB RAM. This is recommended when using SAN Insights feature.

10. Choose Infra-DataStore1 and the Thin Provision virtual disk format. Click NEXT.

11. Choose IB-MGMT Network for all three Source Networks. Click NEXT.

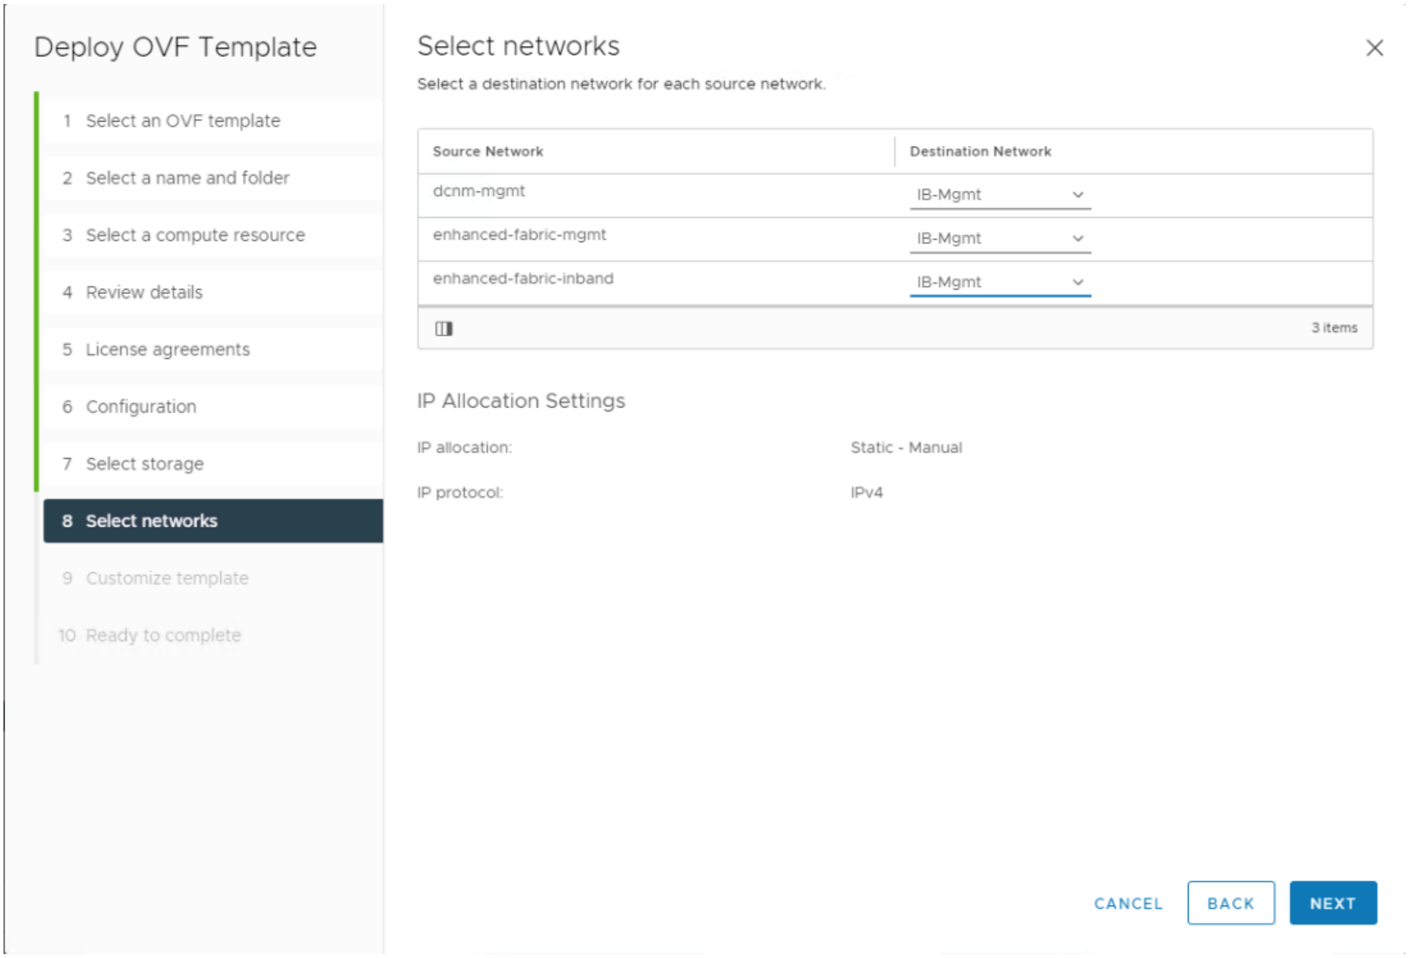

Deploy OVF Template

Select networks

Select a destination network for each source network.

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

3 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

12. Fill-in the management IP address, subnet mask, and gateway. Set the Extra Disk Size according to how many Cisco MDS switches you will be monitoring with this DCNM. If you are only monitoring the two Cisco MDS switches in this FlashStack deployment, set this field to 32. Click NEXT.

13. Review the settings and click FINISH to deploy the OVA.

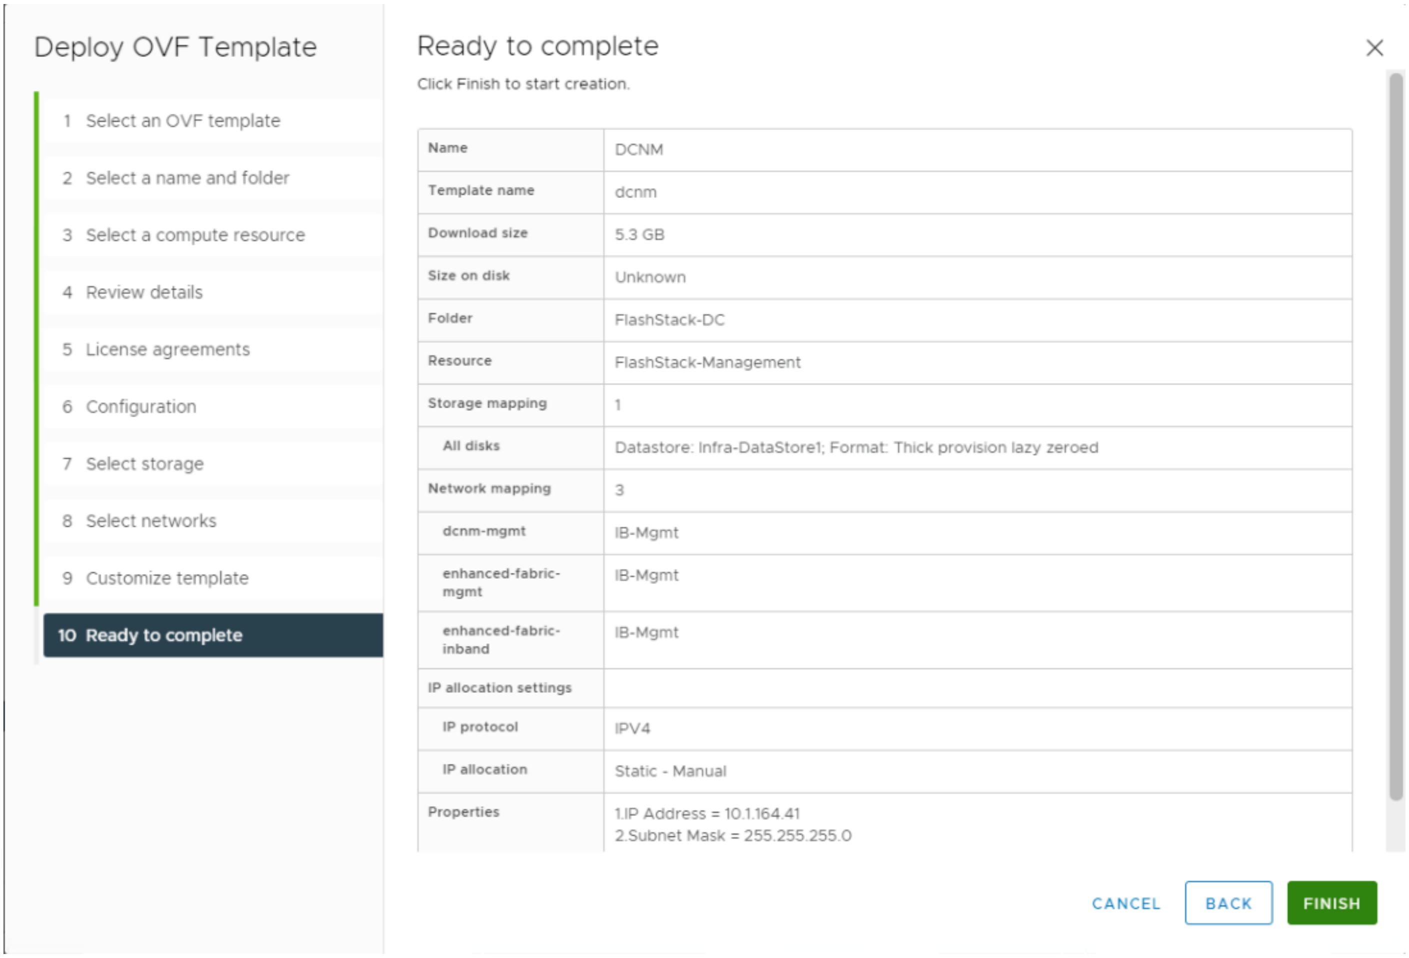

Deploy OVF Template

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

Ready to complete

Click Finish to start creation.

Name	DCNM
Template name	dcnm
Download size	5.3 GB
Size on disk	Unknown
Folder	FlashStack-DC
Resource	FlashStack-Management
Storage mapping	1
All disks	Datastore: Infra-DataStore1; Format: Thick provision lazy zeroed
Network mapping	3
dcnm-mgmt	IB-Mgmt
enhanced-fabric-mgmt	IB-Mgmt
enhanced-fabric-inband	IB-Mgmt
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual
Properties	1.IP Address = 10.1.164.41 2.Subnet Mask = 255.255.255.0

14. After deployment is complete, right-click the newly deployed DCNM VM and click Edit Settings. Expand CPU and adjust the Cores per Socket setting until the number of Sockets is set to match the number of CPUs in the UCS servers used in this deployment. The following example shows 2 sockets.

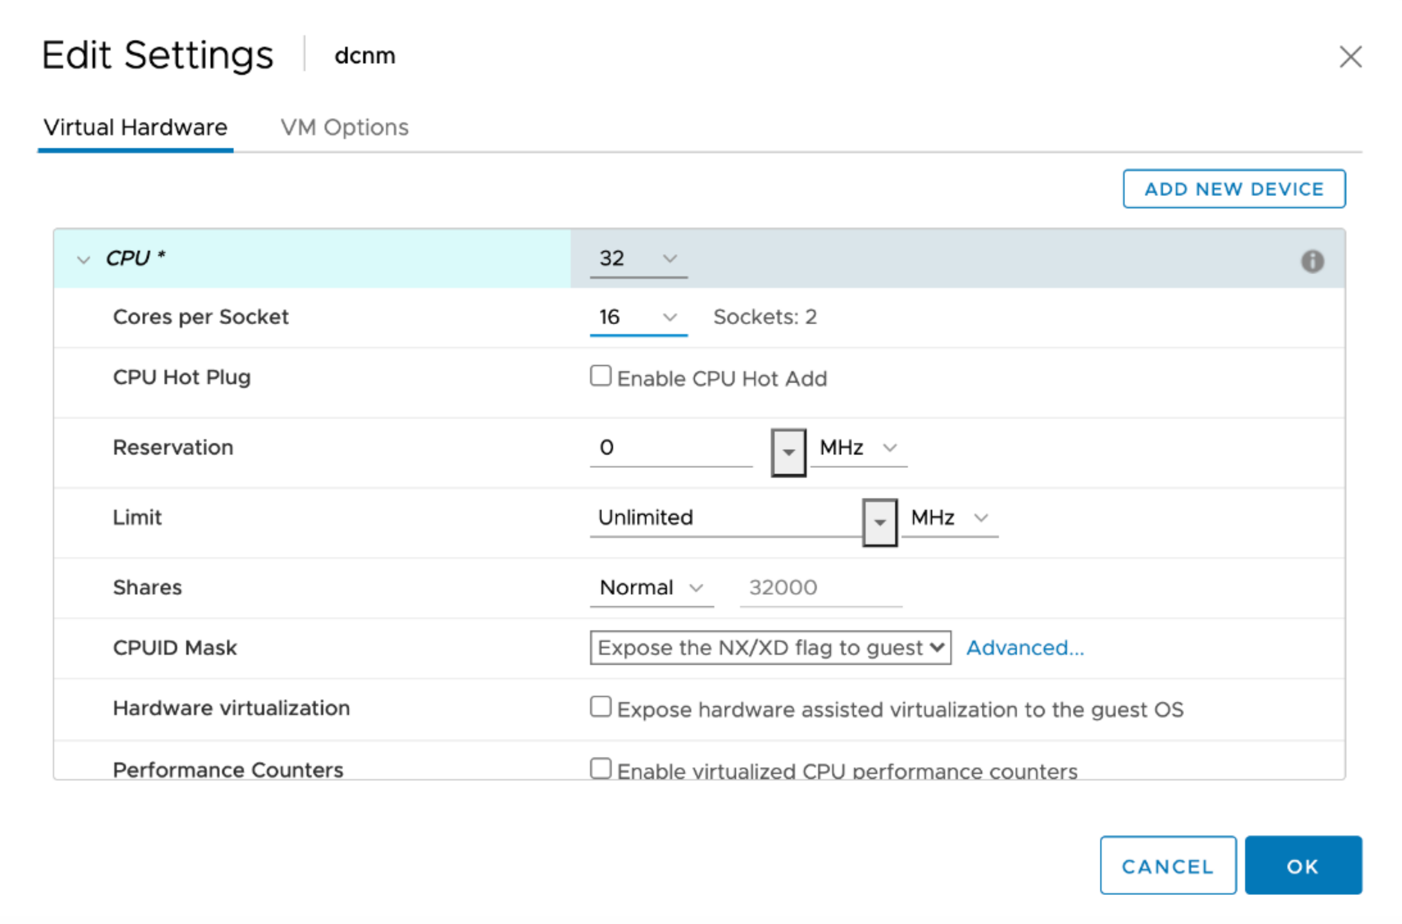

Edit Settings | dcnm

CPU *	32
Cores per Socket	16 Sockets: 2
CPU Hot Plug	□ Enable CPU Hot Add
Reservation	0 MHz
Limit	Unlimited MHz
Shares	Normal 32000
CPUID Mask	Expose the NX/XD flag to guest Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS
Performance Counters	□ Enable virtualized CPU performance counters

15. Click OK to complete the change.

16. Right-click the newly deployed DCNM VM and click Open Remote Console. Once the console is up, click ▶ to power on the VM. Once the VM has powered up, point a web browser to the URL displayed on the console.

17. Navigate the security prompts and click Get started.

18. Make sure Fresh installation – Standalone is selected and click Continue.

19. Choose SAN only for the Installation mode and leave Cisco Systems, Inc. for the OEM vendor and click Next.

20. Enter and repeat the administrator and database passwords and click Next.

21. Enter the DCNM FQDN, a comma-separated list of DNS servers, a comma-separated list of NTP servers, and select the appropriate time zone. Click Next.

22. The Management Network settings should be filled in. For Out-of-Band Network, a different IP address in the same subnet as the management address should be used. Only input the IPV4 address with prefix. Do not put in the Gateway IPv4 Address. Scroll down and click Next.

23. Leave Internal Application Services Network set at the default setting and click Next.

24. Review the Summary details and click Start installation.

25. When the Installation status is complete, click Continue.

26. In the vCenter HTML5 client under Hosts and Clusters, choose the DCNM VM and click the Summary Tab. If an alert is present that states “A newer version of VMware Tools is available for this virtual machine.”, click Upgrade VMware Tools. Choose Automatic Upgrade and click UPGRADE. Wait for the VMware Tools upgrade to complete.

Configure DCNM-SAN

To configure the DCNM-SAN, follow these steps:

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

1. Log in as admin with the password entered above.

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

Note: If you have purchased DCNM server-based or switch-based licenses, follow the instructions that came with the licenses to install them. A new DCNM installation also has a 60-day trial license.

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

4. Click + to add LAN switches. In the Add LAN Devices window, enter the mgmt0 IP address of Nexus switch A in the Seed Switch box. Enter the snmpadmin user name and password set up in the Prerequisites section above. Set Auth-Privacy to SHA_AES. Click Next.

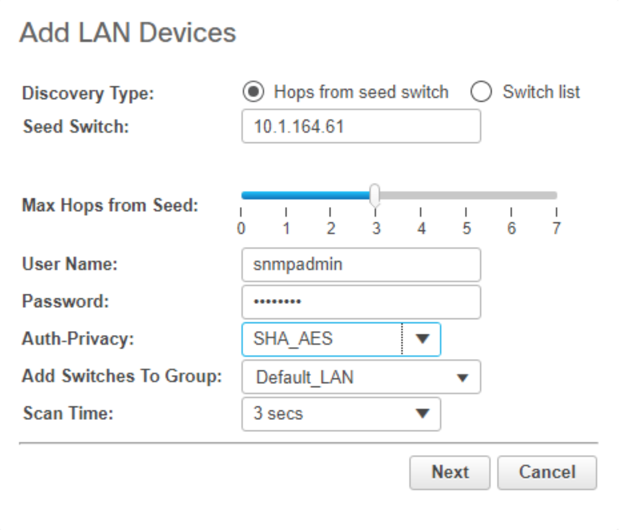

Add LAN Devices

Discovery Type: ● Hops from seed switch ○ Switch list

Seed Switch: 10.1.164.61

Max Hops from Seed: 3

User Name: snmpadmin

Password: ••••••••

Auth-Privacy: SHA_AES

Add Switches To Group: Default_LAN

Scan Time: 3 secs

5. LAN switch discovery will take a few minutes. In the LAN Discovery list that appears, the two Nexus switches and two Fabric Interconnects that are part of this FlashStack should appear with a status of “manageable”. Using the checkboxes on the left, choose the two Nexus switches and two Fabric Interconnects that are part of this FlashStack. Click Add.

6. After a few minutes (click the Refresh icon in the upper right-hand corner), the two Nexus switches and two Fabric Interconnects that are part of this FlashStack will appear with detailed information. The SSH warning under SNMP Status can be ignored since only SNMP can be used to monitor Fabric Interconnects.

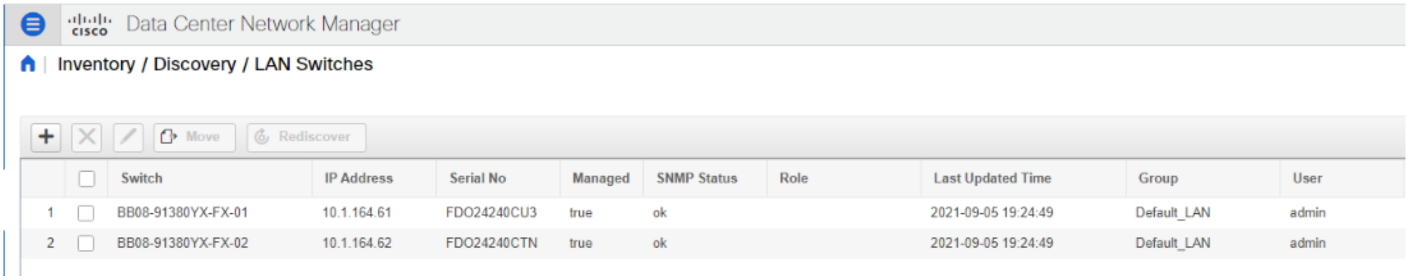
cisco Data Center Network Manager
Inventory / Discovery / LAN Switches
+
Move
Rediscover

	□	Switch	IP Address	Serial No	Managed	SNMP Status	Role	Last Updated Time	Group	User
1	□	BB08-91380YX-FX-01	10.1.164.61	FDO24240CU3	true	ok		2021-09-05 19:24:49	Default_LAN	admin
2	□	BB08-91380YX-FX-02	10.1.164.62	FDO24240CTN	true	ok		2021-09-05 19:24:49	Default_LAN	admin

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

8. Click + to add a switching fabric.

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

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

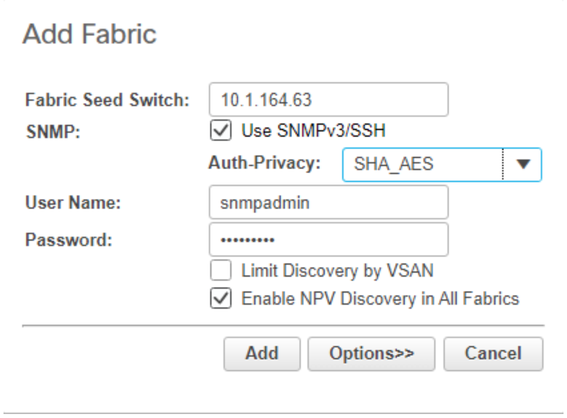

Fabric Seed Switch: 10.1.164.63
SNMP: ☑ Use SNMPv3/SSH
Auth-Privacy: SHA_AES
User Name: snmpadmin
Password: •••••••••
□ Limit Discovery by VSAN
☑ Enable NPV Discovery in All Fabrics

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

The two SAN fabrics appear in the Inventory.

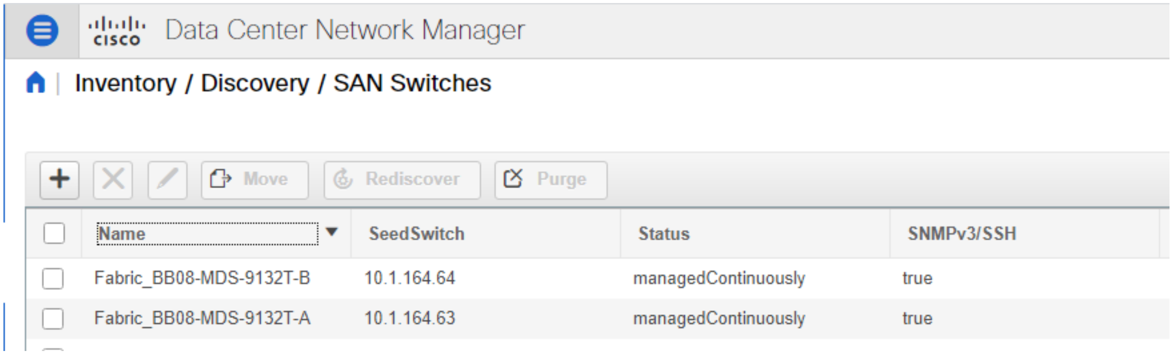

Data Center Network Manager
Inventory / Discovery / SAN Switches
+
Move
Rediscover
Purge

□	Name	SeedSwitch	Status	SNMPv3/SSH
□	Fabric_BB08-MDS-9132T-B	10.1.164.64	managedContinuously	true
□	Fabric_BB08-MDS-9132T-A	10.1.164.63	managedContinuously	true

11. Choose Inventory > Discovery > Virtual Machine Manager.

12. Click + to add the vCenter.

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

14. The vCenter should now appear in the inventory.

15. Choose Administration > Performance Setup > LAN Collections.

16. Choose the Default_LAN group and all information you would like to collect. Click Apply. Click Yes to restart the Performance Collector.

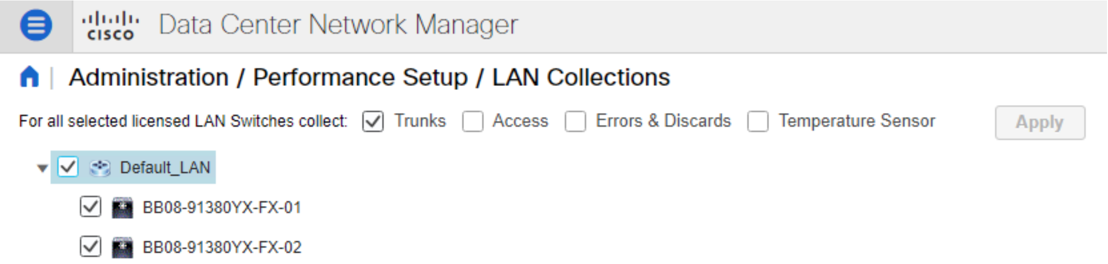

Data Center Network Manager
Administration / Performance Setup / LAN Collections

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

- ☑ Default_LAN
  - ☑ BB08-91380YX-FX-01
  - ☑ BB08-91380YX-FX-02

17. Choose Administration > Performance Setup > SAN Collections.

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

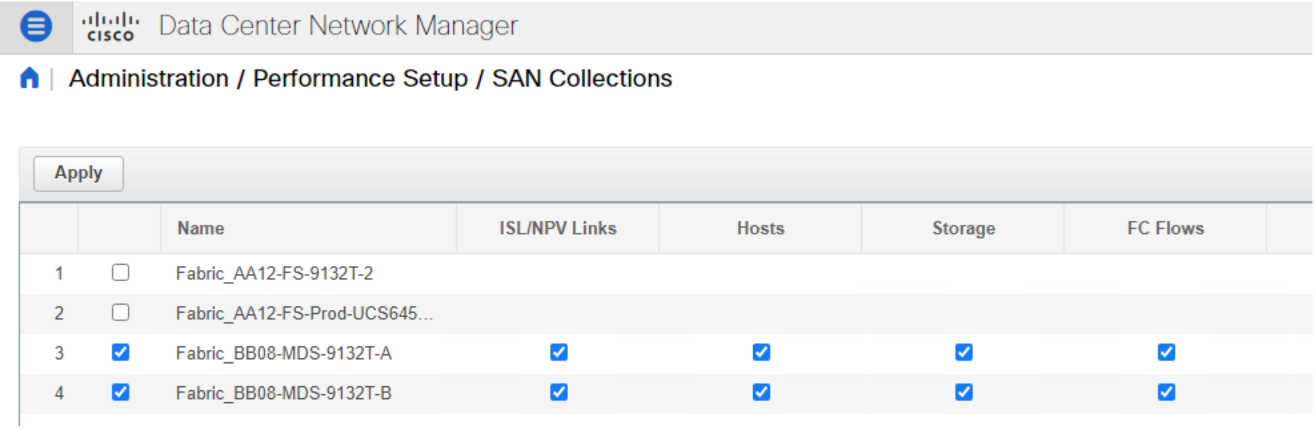

Data Center Network Manager
Administration / Performance Setup / SAN Collections
Apply

		Name	ISL/NPV Links	Hosts	Storage	FC Flows
1	☐	Fabric_AA12-FS-9132T-2				
2	☐	Fabric_AA12-FS-Prod-UCS645...				
3	☑	Fabric_BB08-MDS-9132T-A	☑	☑	☑	☑
4	☑	Fabric_BB08-MDS-9132T-B	☑	☑	☑	☑

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

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

You can now explore all of the different options and information provided by DCNM SAN. See Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.5(1).

Configure SAN Insights in DCNM SAN

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

- Ensure that the time configurations set above, including daylight savings settings, are consistent across the MDS switches and Cisco DCNM.
- SAN Insights requires the installation of a switch-based SAN Analytics license on each switch. To trial the feature, each switch includes a one-time 120-day grace period for SAN Analytics from the time the feature is first enabled.
- SAN Insights supports current Fibre Channel Protocol (SCSI) and NVMe over Fibre Channel (NVMe).
- SAN Insights works by enabling SAN Analytics and Telemetry Streaming on each switch. The switches then stream the SAN Analytics data to DCNM, which collects, correlates, and displays statistics. All configurations can be done from DCNM.
- Only Cisco MDS switches support SAN Analytics. Nexus 93180YC-FX switches do not support SAN Analytics.
- For more information on SAN Insights, see the SAN Insights sections: Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.5(1).

- For more information on SAN Analytics, see: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/san_analytics/cisco-mds9000-san-analytics-telemetry-streaming-config-guide-8x.html.

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

1. Click Configure > SAN > SAN Insights. Click Continue.

2. Choose Fabric A. Click Continue.

3. Choose the Fabric A Cisco MDS switch. Under Install Query click None and from the drop-down list click Storage. Under Subscriptions, choose SCSI & NVMe. Optionally, under Receiver, choose the second IP address in the In-Band Management subnet configured for DCNM. Click Save, then click Continue.

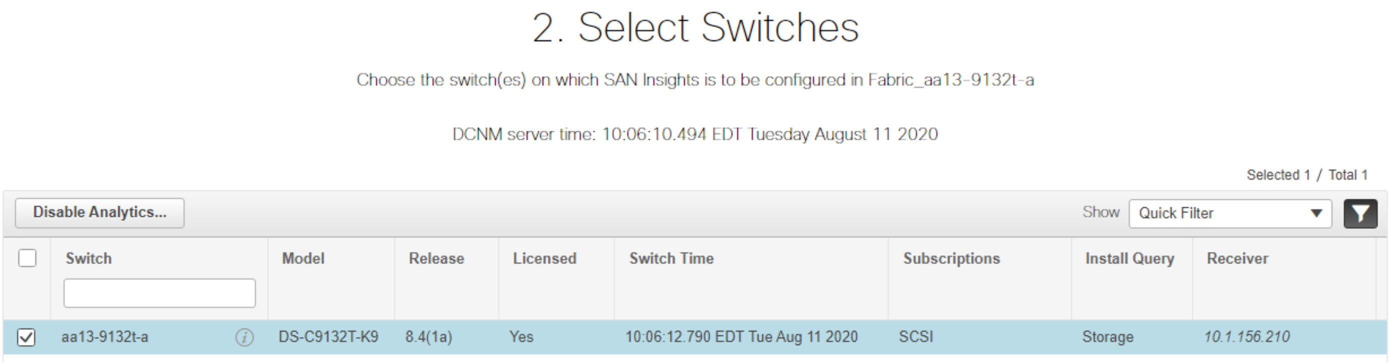

2. Select Switches

Choose the switch(es) on which SAN Insights is to be configured in Fabric_aa13-9132t-a

DCNM server time: 10:06:10.494 EDT Tuesday August 11 2020

Selected 1 / Total 1

□	Switch	Model	Release	Licensed	Switch Time	Subscriptions	Install Query	Receiver
☑	aa13-9132t-a	DS-C9132T-K9	8.4(1a)	Yes	10:06:12.790 EDT Tue Aug 11 2020	SCSI	Storage	10.1.156.210

4. Review the information and click Continue.

5. Expand the switch and then the module. Under Enable / Disable SCSI Telemetry, click the left icon to enable telemetry on the ports connected to the FlashArray//X R3 Click Continue.

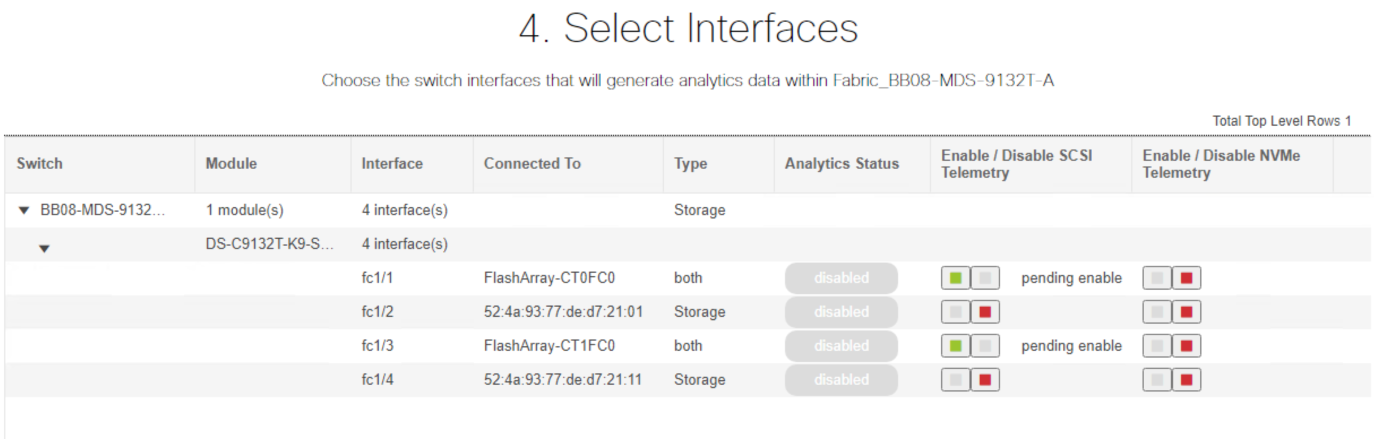

4. Select Interfaces

Choose the switch interfaces that will generate analytics data within Fabric_BB08-MDS-9132T-A

Total Top Level Rows 1

Switch	Module	Interface	Connected To	Type	Analytics Status	Enable / Disable SCSI Telemetry	Enable / Disable NVMe Telemetry
▼ BB08-MDS-9132...	1 module(s)	4 interface(s)		Storage			
▼	DS-C9132T-K9-S...	4 interface(s)					
		fc1/1	FlashArray-CT0FC0	both	disabled	🟩 □ pending enable	□ 🟥
		fc1/2	52:4a:93:77:de:d7:21:01	Storage	disabled	□ 🟥	□ 🟥
		fc1/3	FlashArray-CT1FC0	both	disabled	🟩 □ pending enable	□ 🟥
		fc1/4	52:4a:93:77:de:d7:21:11	Storage	disabled	□ 🟥	□ 🟥

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

7. Ensure that the two operations were successful and click Close.

8. Repeat steps 1-7 to install SAN Analytics and Telemetry on the Fabric B switch.

9. After approximately two hours, you can view SAN Analytics data under the Dashboard and Monitor.

Cisco Intersight Orchestration

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. FlashStack environment includes multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight but does not connect directly with it, will need a connection mechanism. Cisco Intersight Assist provides that connection mechanism and helps you add devices into Cisco Intersight.

Cisco Intersight Assist is available within the Cisco Intersight Virtual Appliance, which is distributed as a deployable virtual machine contained within an Open Virtual Appliance (OVA) file format. You can install the appliance on an ESXi server. For more information, see the Cisco Intersight Virtual Appliance Getting Started Guide.

After claiming Intersight Assist into Intersight, you can claim endpoint devices using the Claim Through Intersight Assist option.

The following sections describe some of the sample FlashStack Orchestration and lifecycle management tasks that can be performed using Cisco Intersight.

To configure Cisco Intersight Assist Virtual Appliance, follow these steps:

1. To install Intersight Assist from an Open Virtual Appliance (OVA) in your VMware FlashStack-Management Cluster, first download the latest release of the OVA from https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-230.

2. 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.

3. From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the FlashStack-Management cluster and click Deploy OVF Template.

4. Specify a URL or browse to the intersight-virtual-appliance-1.0.9-230.ova file. Click NEXT.

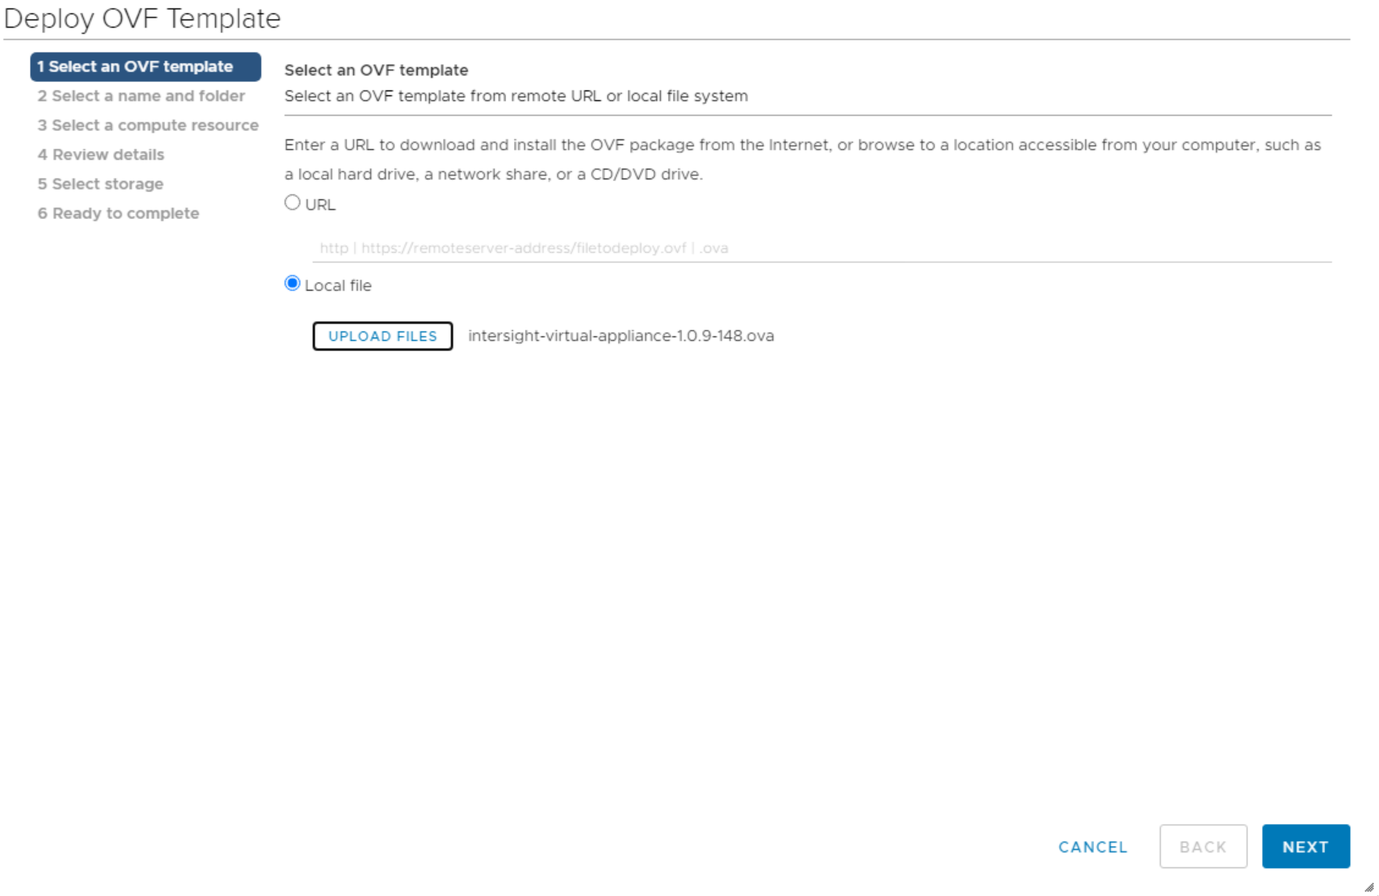

Deploy OVF Template

Select an OVF template

Select an OVF template from remote URL or local file system

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

○ URL

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

● Local file

UPLOAD FILES intersight-virtual-appliance-1.0.9-148.ova

5. Name the Intersight Assist VM and choose the location. Click NEXT.

6. Choose the FlashStack-Management cluster and click NEXT.

7. Review details and click NEXT.

8. Choose a deployment configuration (Tiny recommended) and click NEXT.

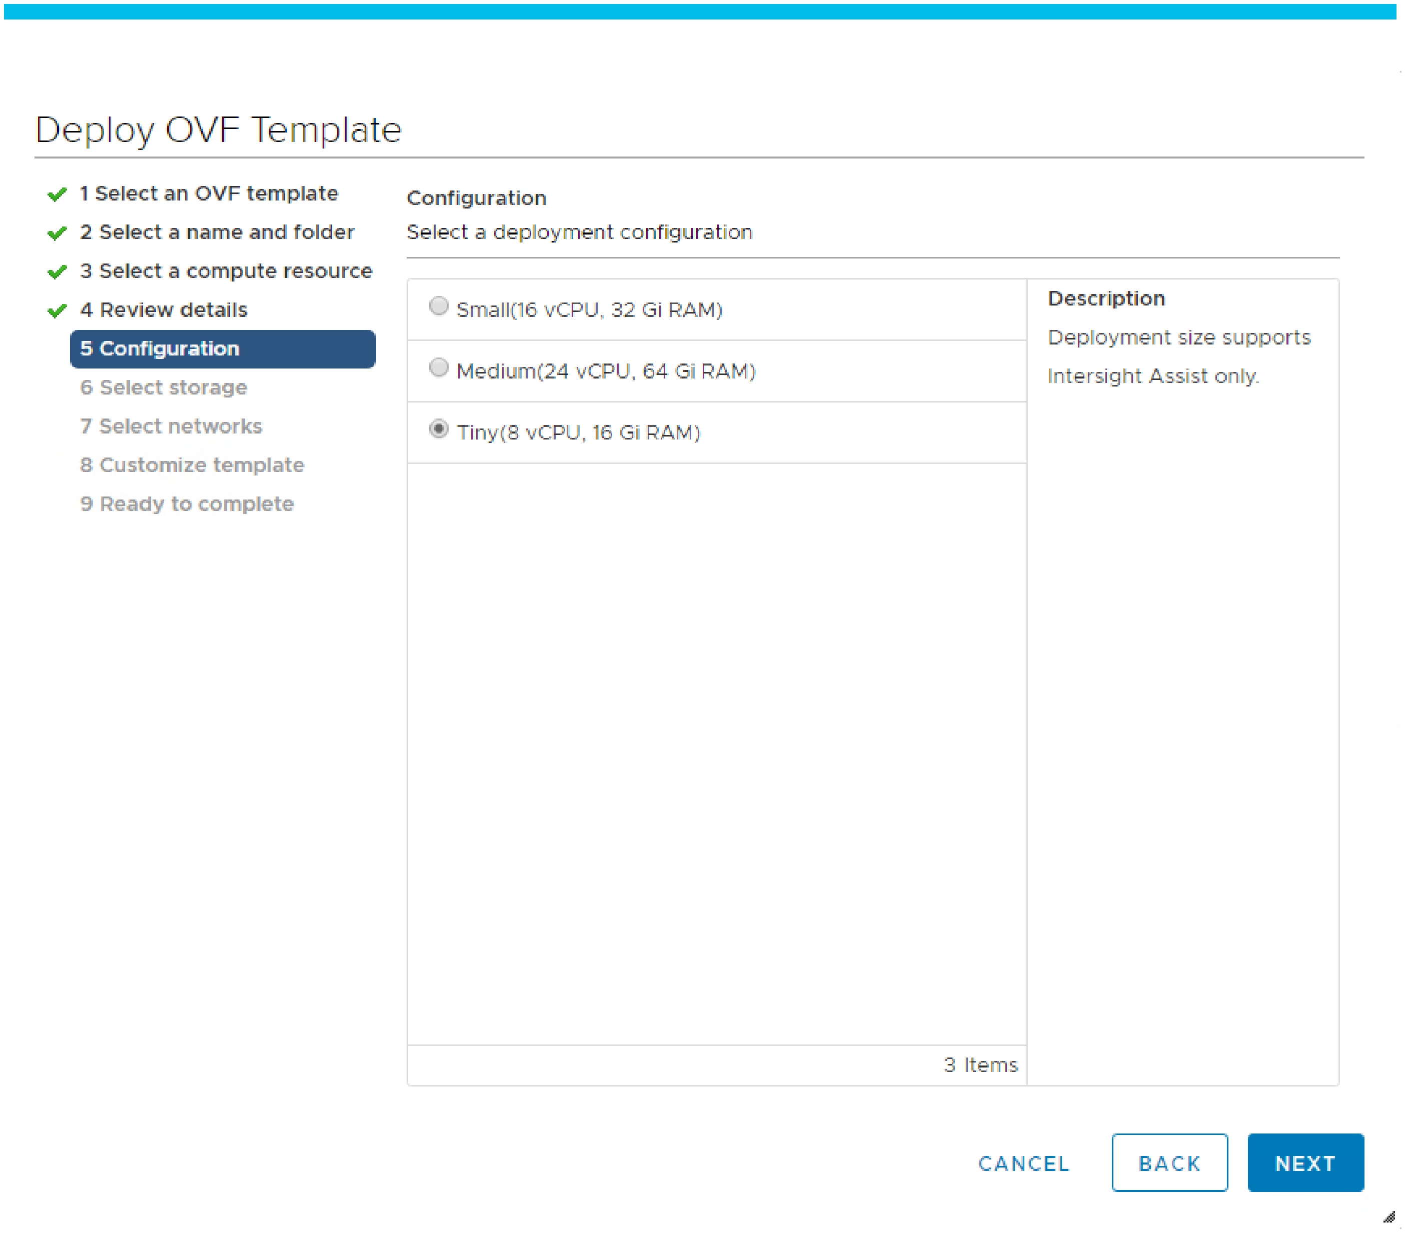

Deploy OVF Template

- ✓ 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.

9. Choose Infra-DataStore1 for storage and choose the Thin Provision virtual disk format. Click NEXT.

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

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

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

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

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

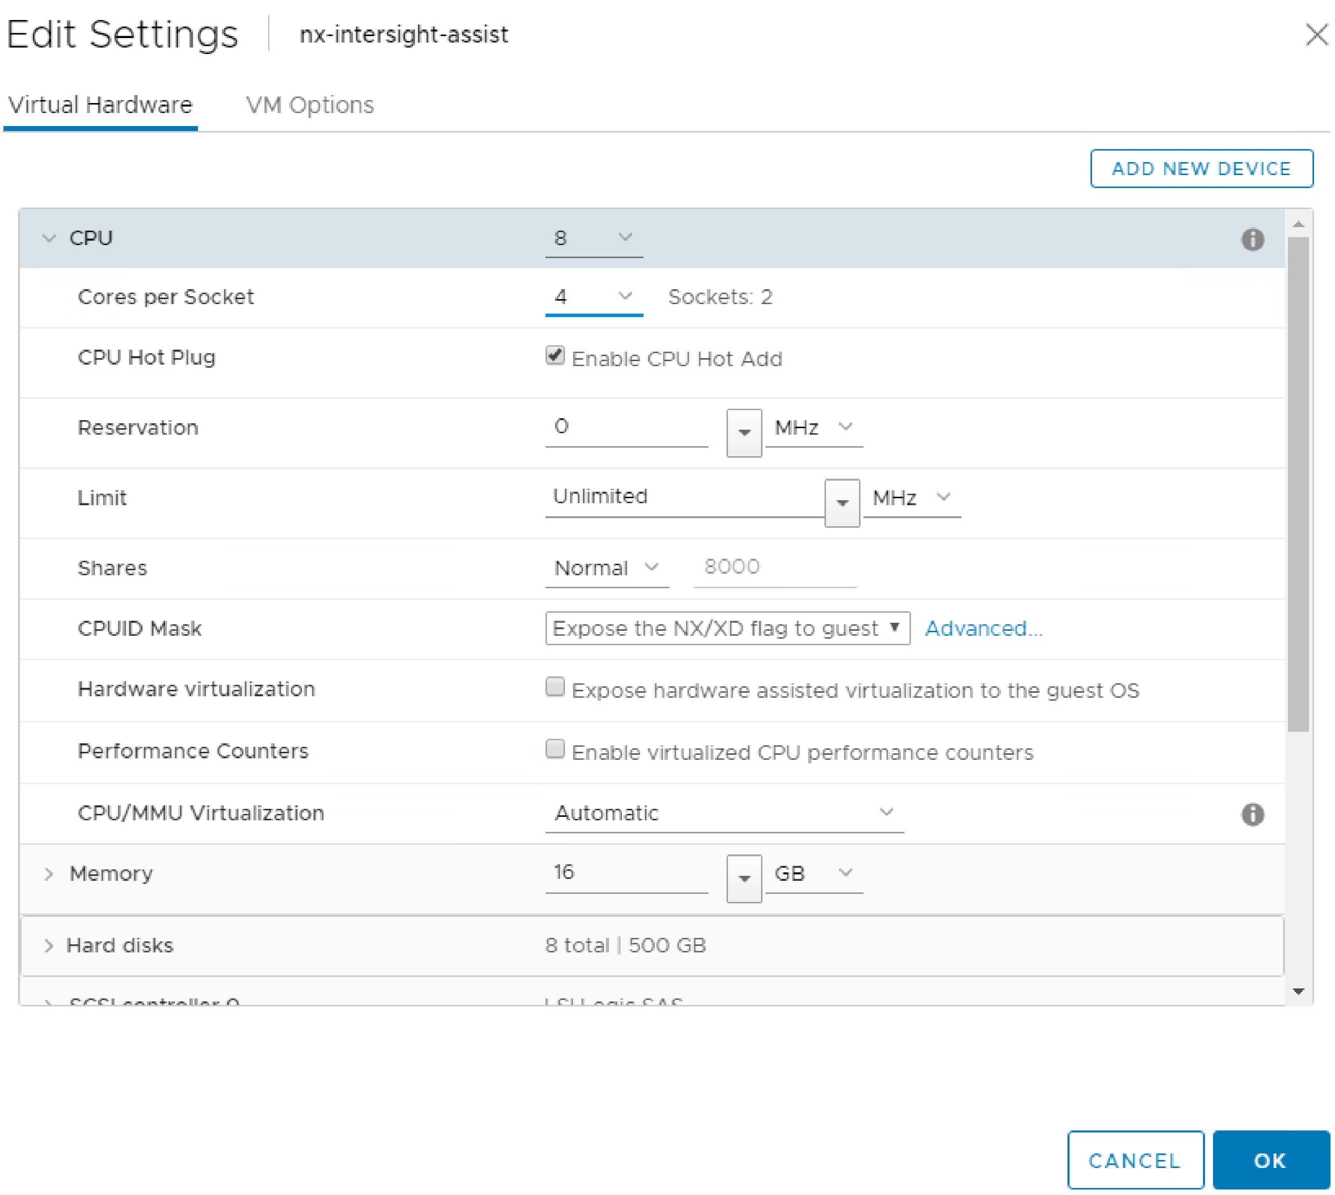

Edit Settings | nx-intersight-assist

Virtual Hardware | VM Options

ADD NEW DEVICE

CPU	8
Cores per Socket	4 Sockets: 2
CPU Hot Plug	☑ Enable CPU Hot Add
Reservation	0 MHz
Limit	Unlimited MHz
Shares	Normal 8000
CPUID Mask	Expose the NX/XD flag to guest Advanced...
Hardware virtualization	□ 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

CANCEL OK

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

16. Click ▶ to power on the VM.

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

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

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

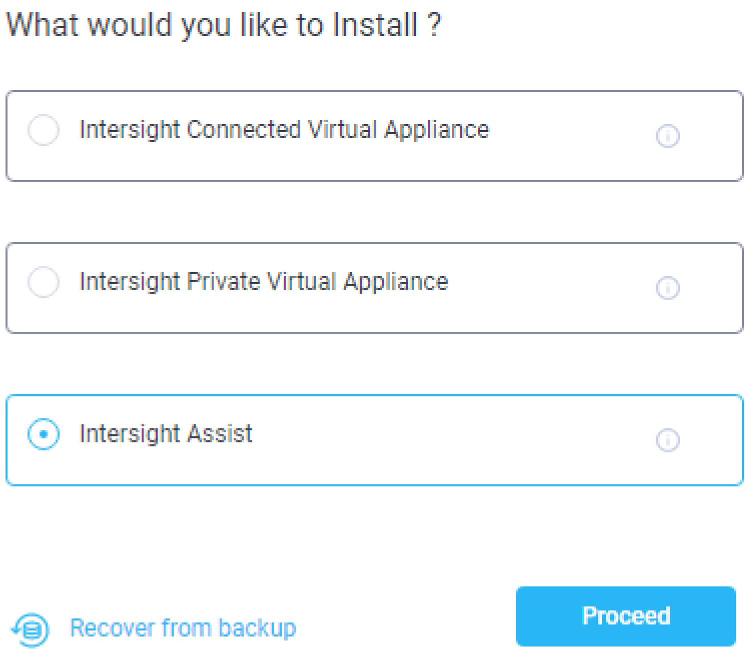

What would you like to Install ?

○ Intersight Connected Virtual Appliance

○ Intersight Private Virtual Appliance

● Intersight Assist

19. 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.

20. In the Intersight Assist web interface, click Continue.

21. The Intersight Assist software will now be downloaded and installed into the Intersight Assist VM. This can take up to an hour to complete.

Note: 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.

22. When the software download is complete, navigate the security prompts and an Intersight Assist login screen will appear. Log into Intersight Assist with the admin user and the password supplied in the OVA installation. Check the Intersight Assist status and log out of Intersight Assist.

23. To claim the vCenter, from Cisco Intersight, click ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select VMware vCenter under Hypervisor and click Start. In the VMware vCenter window, make sure the Intersight Assist is correctly selected, fill in the vCenter information, and click Claim.

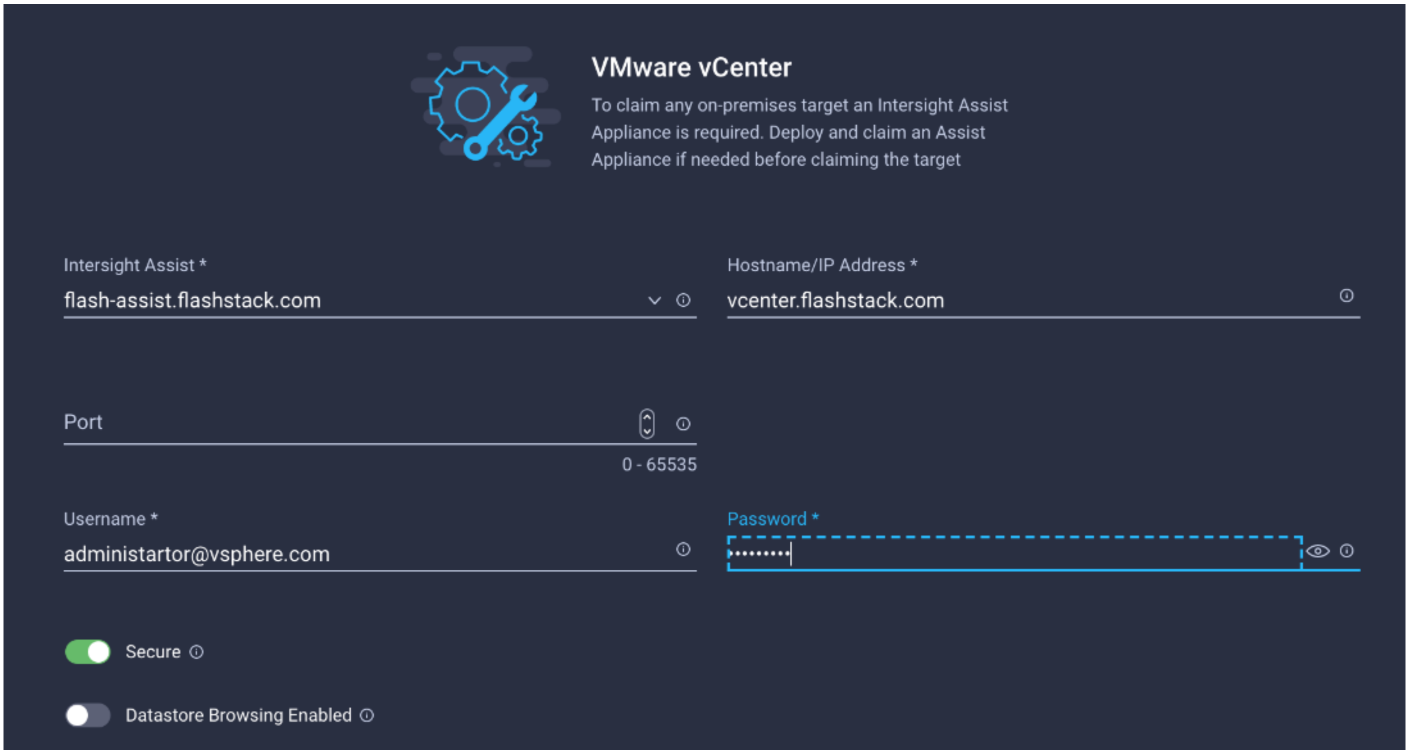

VMware vCenter

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

Intersight Assist * flash-assist.flashstack.com

Hostname/IP Address * vcenter.flashstack.com

Port

0 - 65535

Username * administartor@vsphere.com

Password * ••••••••

Secure

Datastore Browsing Enabled

24. 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.

25. Detailed information obtained from the vCenter can now be viewed by clicking Virtualization from the menu.

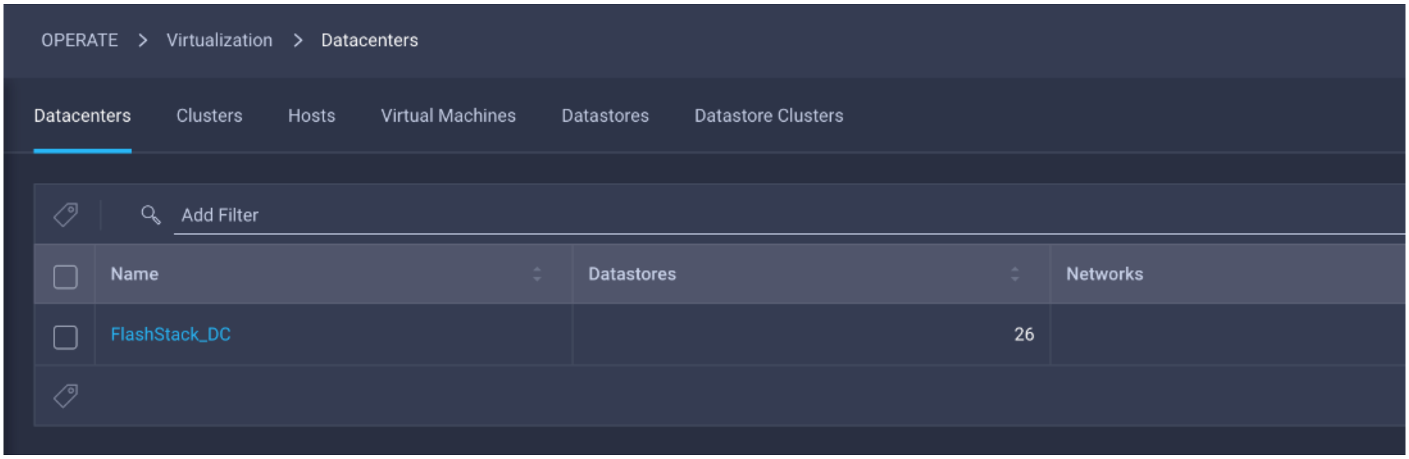

OPERATE > Virtualization > Datacenters

□	Name	Datastores	Networks
□	FlashStack_DC	26	

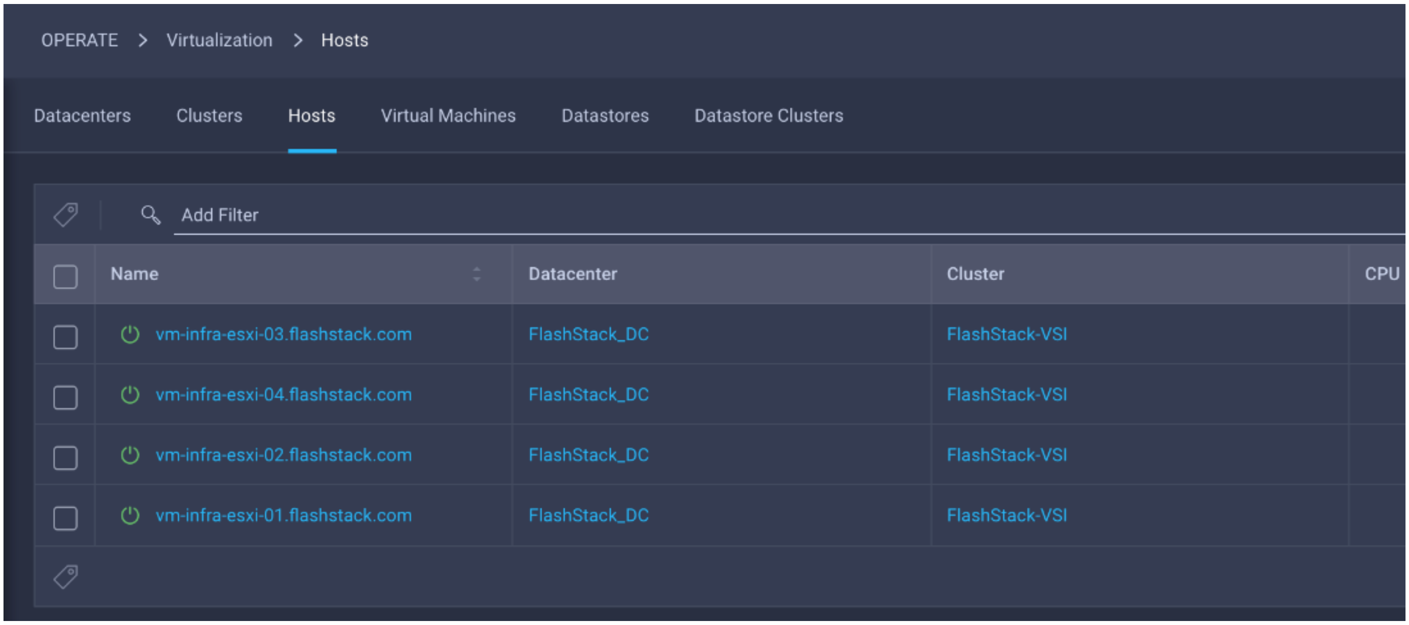

Name	Datacenter	Cluster	CPU
vm-infra-esxi-03.flashstack.com	FlashStack_DC	FlashStack-VSI	
vm-infra-esxi-04.flashstack.com	FlashStack_DC	FlashStack-VSI	
vm-infra-esxi-02.flashstack.com	FlashStack_DC	FlashStack-VSI	
vm-infra-esxi-01.flashstack.com	FlashStack_DC	FlashStack-VSI	

Claim FlashArray//X in Cisco Intersight

Claiming a Pure Storage FlashArray also requires the use of an Intersight Assist virtual machine. Deploy an Intersight assist appliance using the above-described procedure if one doesn't exist. To claim FlashArray//X in Cisco Intersight, follow these steps:

1. Open a browser to Cisco Intersight, https://intersight.com, and log into your Intersight account.

2. Click Admin > Devices.

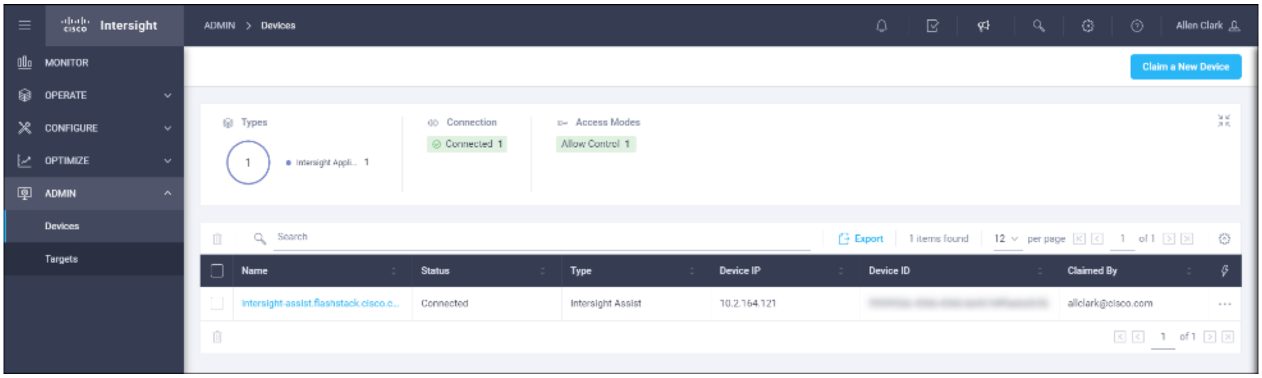

Name	Status	Type	Device IP	Device ID	Claimed By
intersight-assist.flashstack.cisco.c...	Connected	Intersight Assist	10.2.164.121	[redacted]	allclark@cisco.com

3. Click Claim a New Device and select Claim Though Intersight Assist.

4. Set Type to Pure Storage FlashArray.

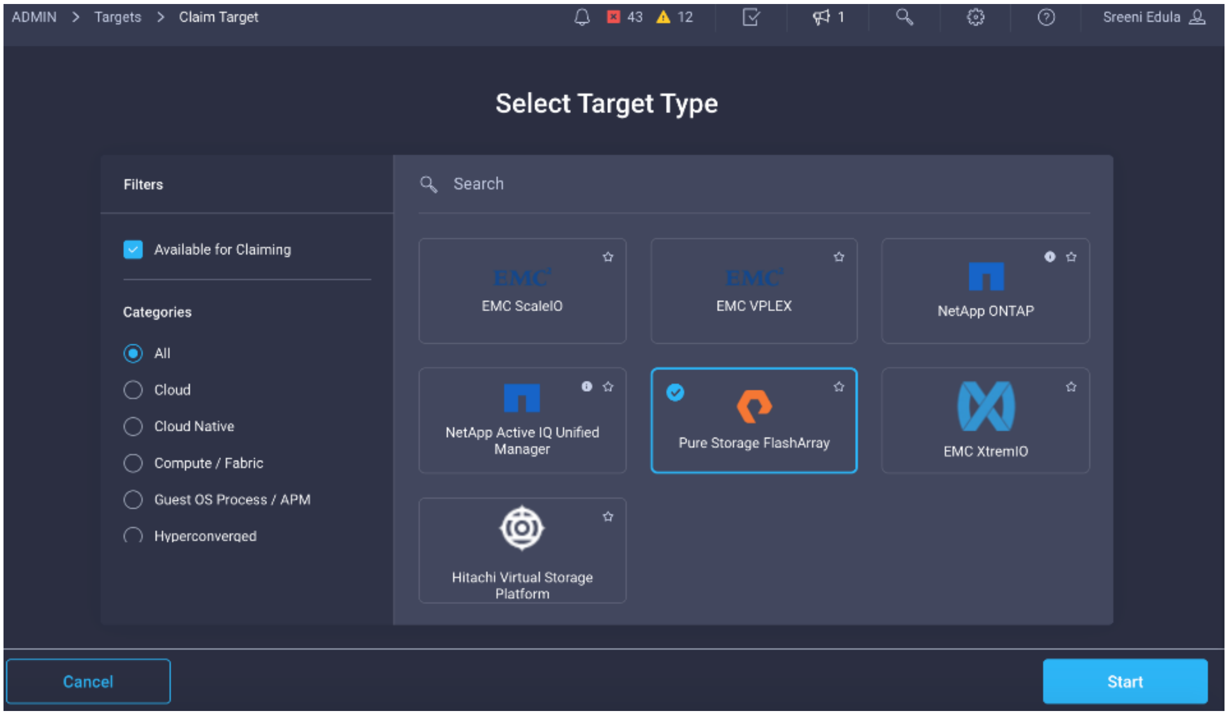

Select Target Type

☑ Available for Claiming

● All

☑ Pure Storage FlashArray

5. Click Start.

6. Enter FlashArray Hostname/ IP address and credentials.

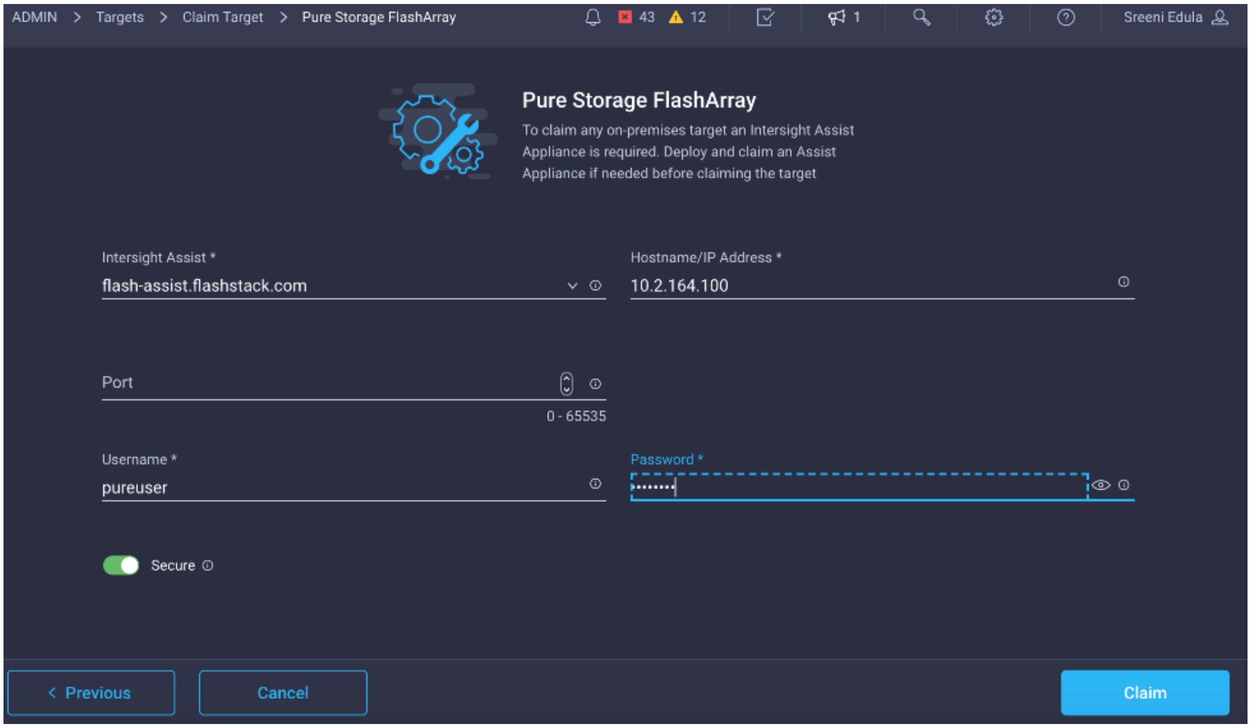

Pure Storage FlashArray

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

Intersight Assist * flash-assist.flashstack.com

Hostname/IP Address * 10.2.164.100

Port

0 - 65535

Username * pureuser

Password * ••••••••

Secure ●

7. Click Claim.

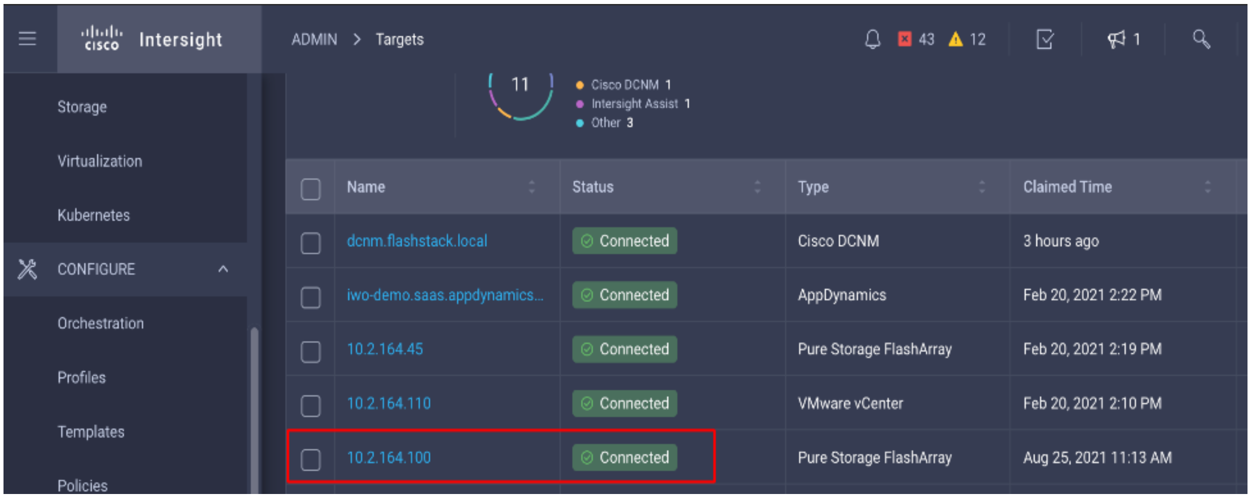

☐	Name	Status	Type	Claimed Time
☐	dcnm.flashstack.local	Connected	Cisco DCNM	3 hours ago
☐	iwo-demo.saas.appdynamics...	Connected	AppDynamics	Feb 20, 2021 2:22 PM
☐	10.2.164.45	Connected	Pure Storage FlashArray	Feb 20, 2021 2:19 PM
☐	10.2.164.110	Connected	VMware vCenter	Feb 20, 2021 2:10 PM
☐	10.2.164.100	Connected	Pure Storage FlashArray	Aug 25, 2021 11:13 AM

FC Host Registration using Cisco Intersight

To register the FC host using Cisco Intersight, follow these steps:

1. Click Configure > Orchestration.
2. Click New Storage Host.

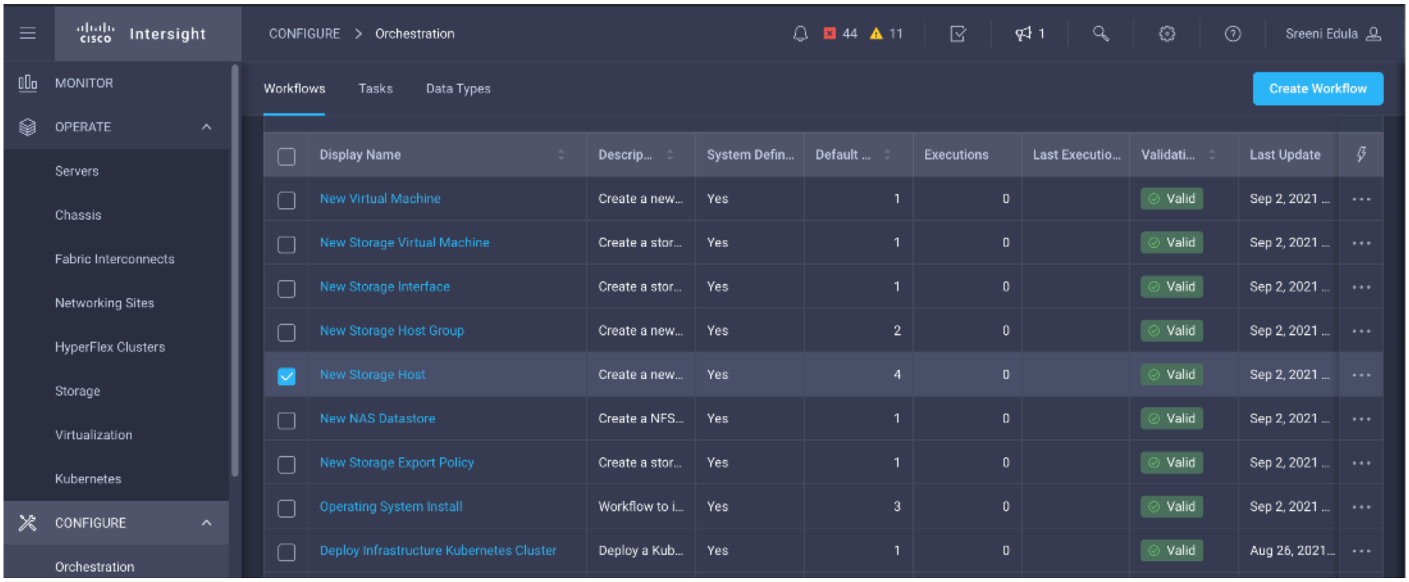

☐	Display Name	Descrip...	System Defin...	Default ...	Executions	Last Executio...	Validati...	Last Update
☐	New Virtual Machine	Create a new...	Yes	1	0		Valid	Sep 2, 2021 ...
☐	New Storage Virtual Machine	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...
☐	New Storage Interface	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...
☐	New Storage Host Group	Create a new...	Yes	2	0		Valid	Sep 2, 2021 ...
☑	New Storage Host	Create a new...	Yes	4	0		Valid	Sep 2, 2021 ...
☐	New NAS Datastore	Create a NFS...	Yes	1	0		Valid	Sep 2, 2021 ...
☐	New Storage Export Policy	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...
☐	Operating System Install	Workflow to i...	Yes	3	0		Valid	Sep 2, 2021 ...
☐	Deploy Infrastructure Kubernetes Cluster	Deploy a Kub...	Yes	1	0		Valid	Aug 26, 2021...

3. Click Execute.

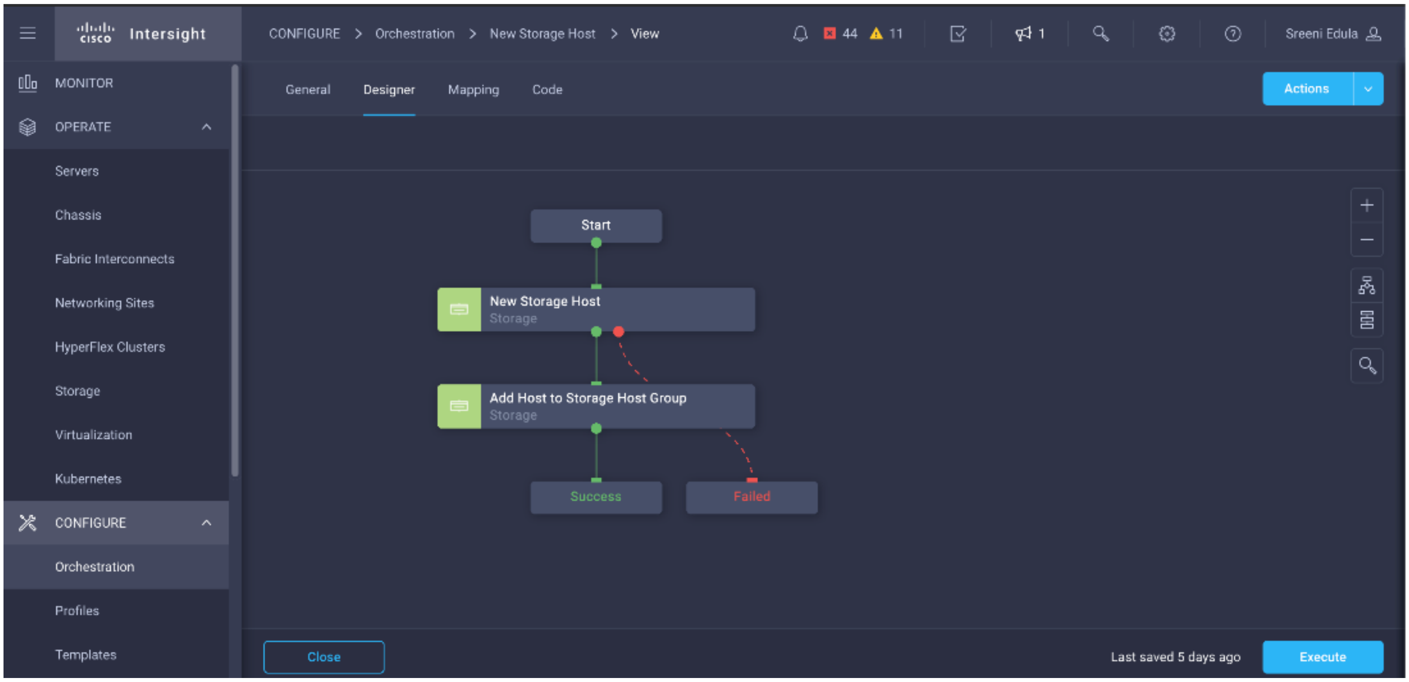

4. Select the appropriate Organization (default by default).

5. Select the appropriate Pure Storage device.

6. Enter the name of the Host name and WWNs for host VM-Host-Infra-FCP-01.

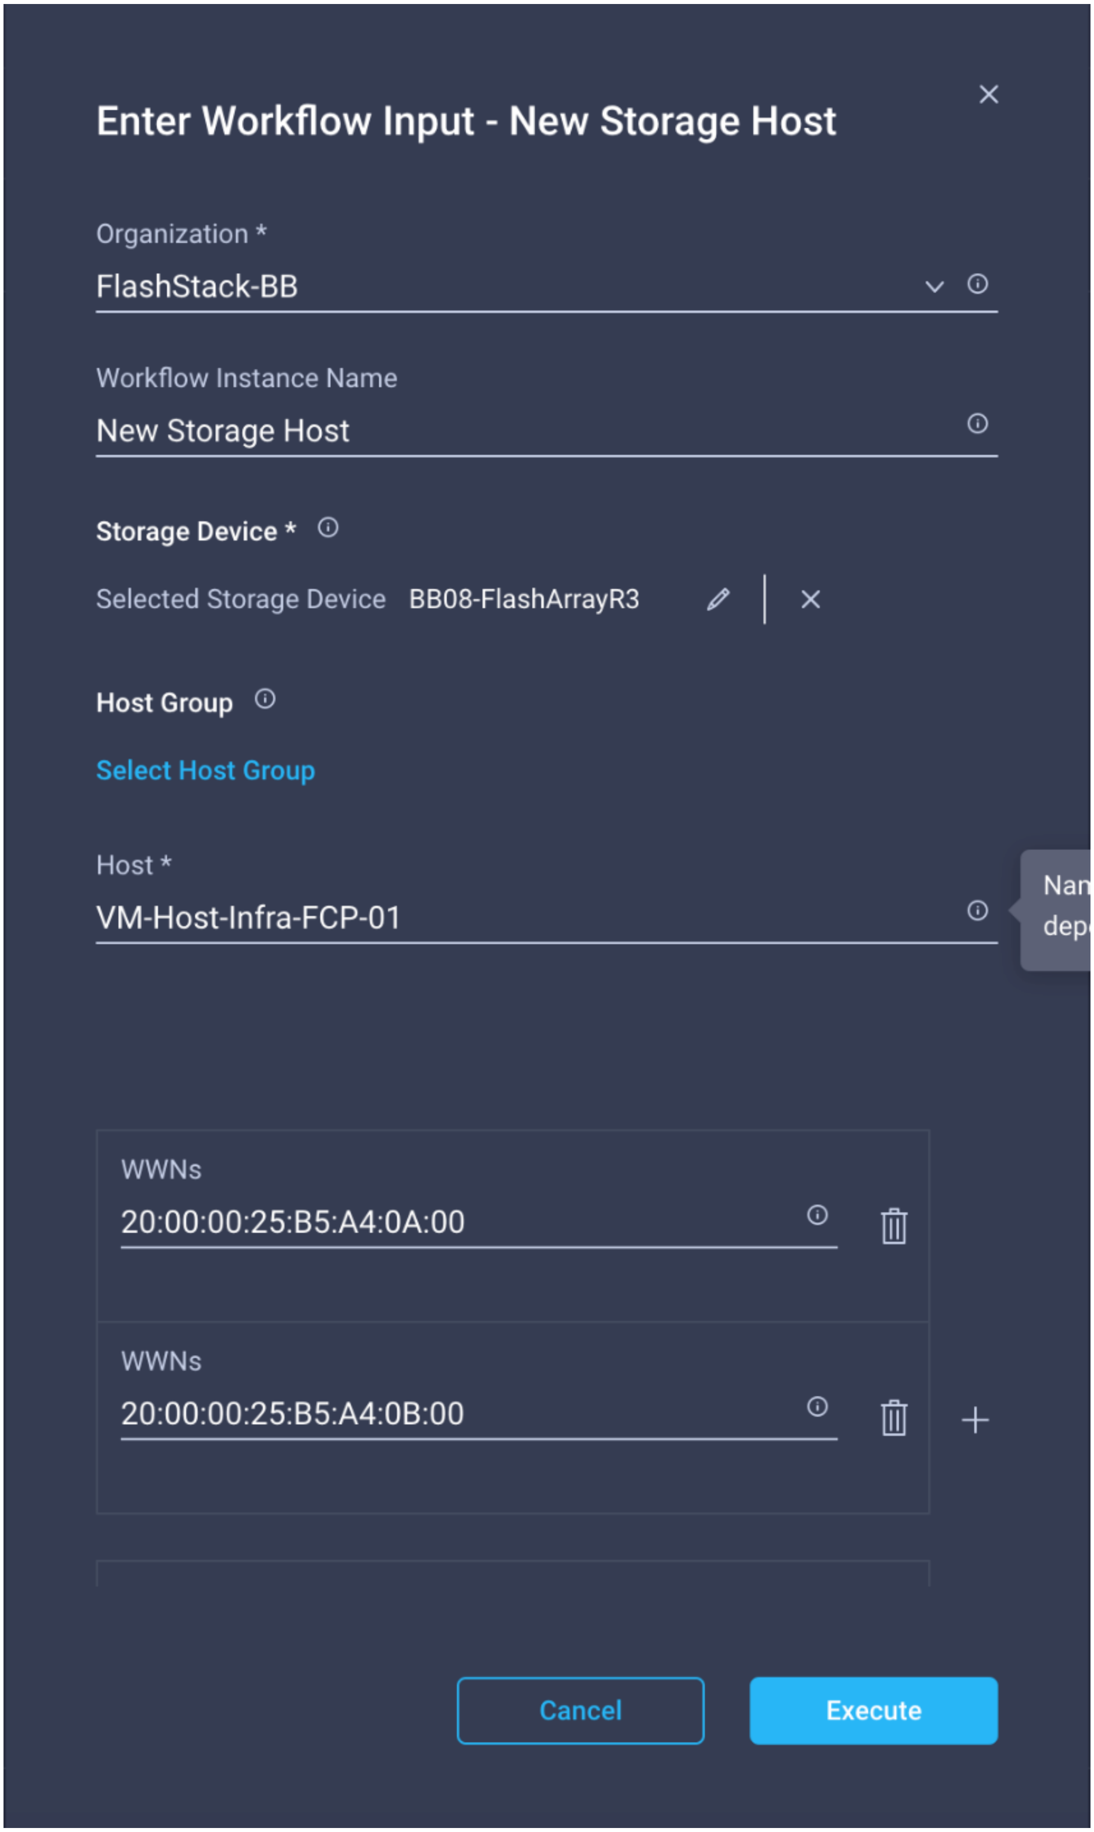

Enter Workflow Input - New Storage Host

Organization *

FlashStack-BB

Workflow Instance Name

New Storage Host

Storage Device *

Selected Storage Device BB08-FlashArrayR3

Host Group

Select Host Group

Host *

VM-Host-Infra-FCP-01

WWNs

20:00:00:25:B5:A4:0A:00

WWNs

20:00:00:25:B5:A4:0B:00

7. Click Execute.

8. Repeat Steps 2-7 for all hosts.

Create FC Host Group using Cisco Intersight

To create an FC host group using Cisco Intersight, follow these steps:

1. Click Configure > Orchestration.

2. Click New Storage Host Group.

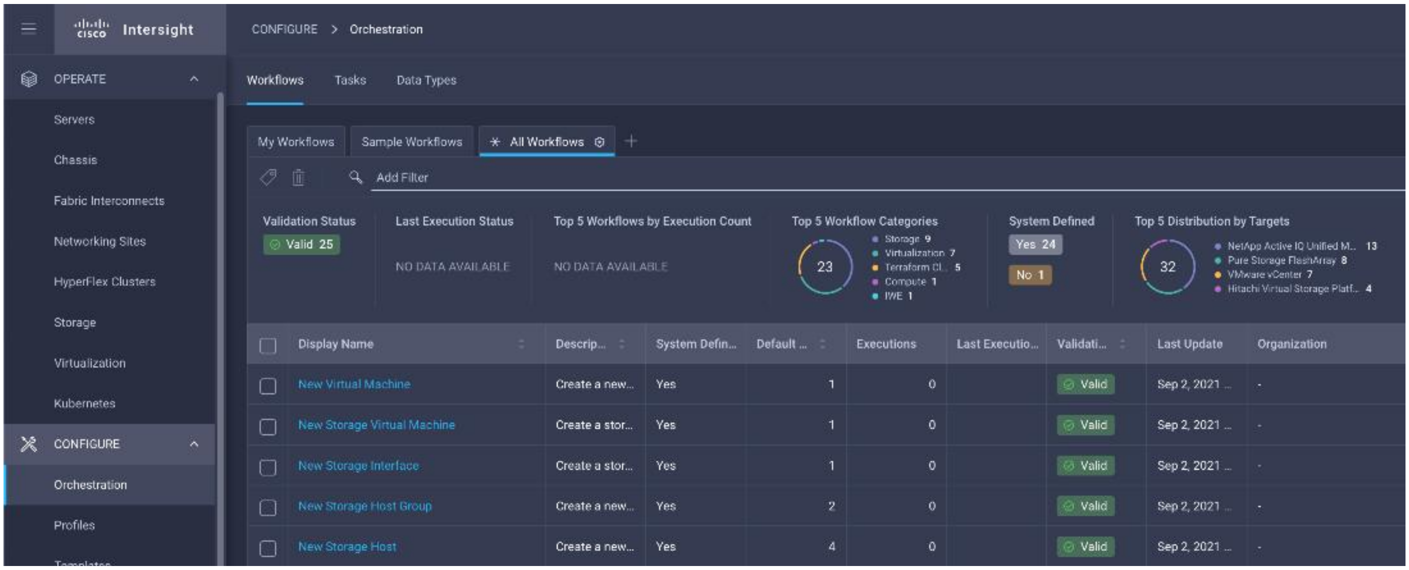

Display Name	Descrip...	System Defin...	Default ...	Executions	Last Executio...	Validati...	Last Update	Organization
New Virtual Machine	Create a new...	Yes	1	0		Valid	Sep 2, 2021 ...	-
New Storage Virtual Machine	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
New Storage Interface	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
New Storage Host Group	Create a new...	Yes	2	0		Valid	Sep 2, 2021 ...	-
New Storage Host	Create a new...	Yes	4	0		Valid	Sep 2, 2021 ...	-

3. Click Execute.

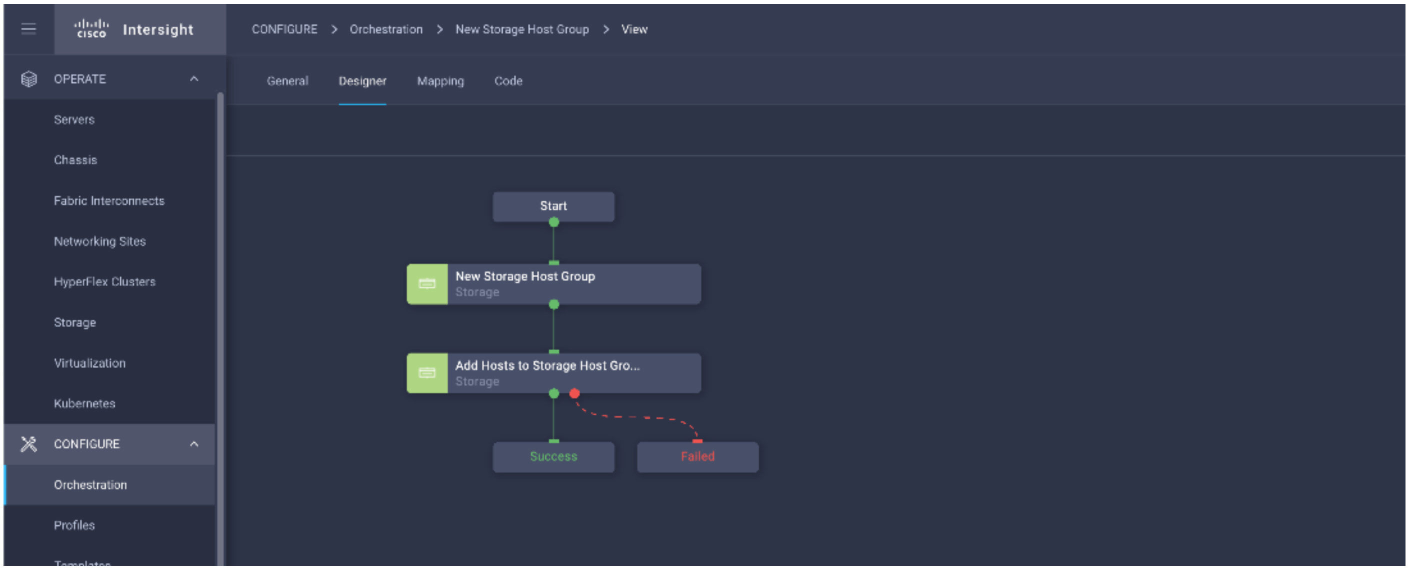

4. Select the appropriate Organization (default by default).

5. Select the appropriate Pure Storage device.

6. Enter the name of the Host Group and of the Hosts created during Host Registration. VM-Infra-Host-FCP-01, VM-Infra-Host-FC-02 and VM-Infra-Host-FCP-03 are the hosts used in this deployment.

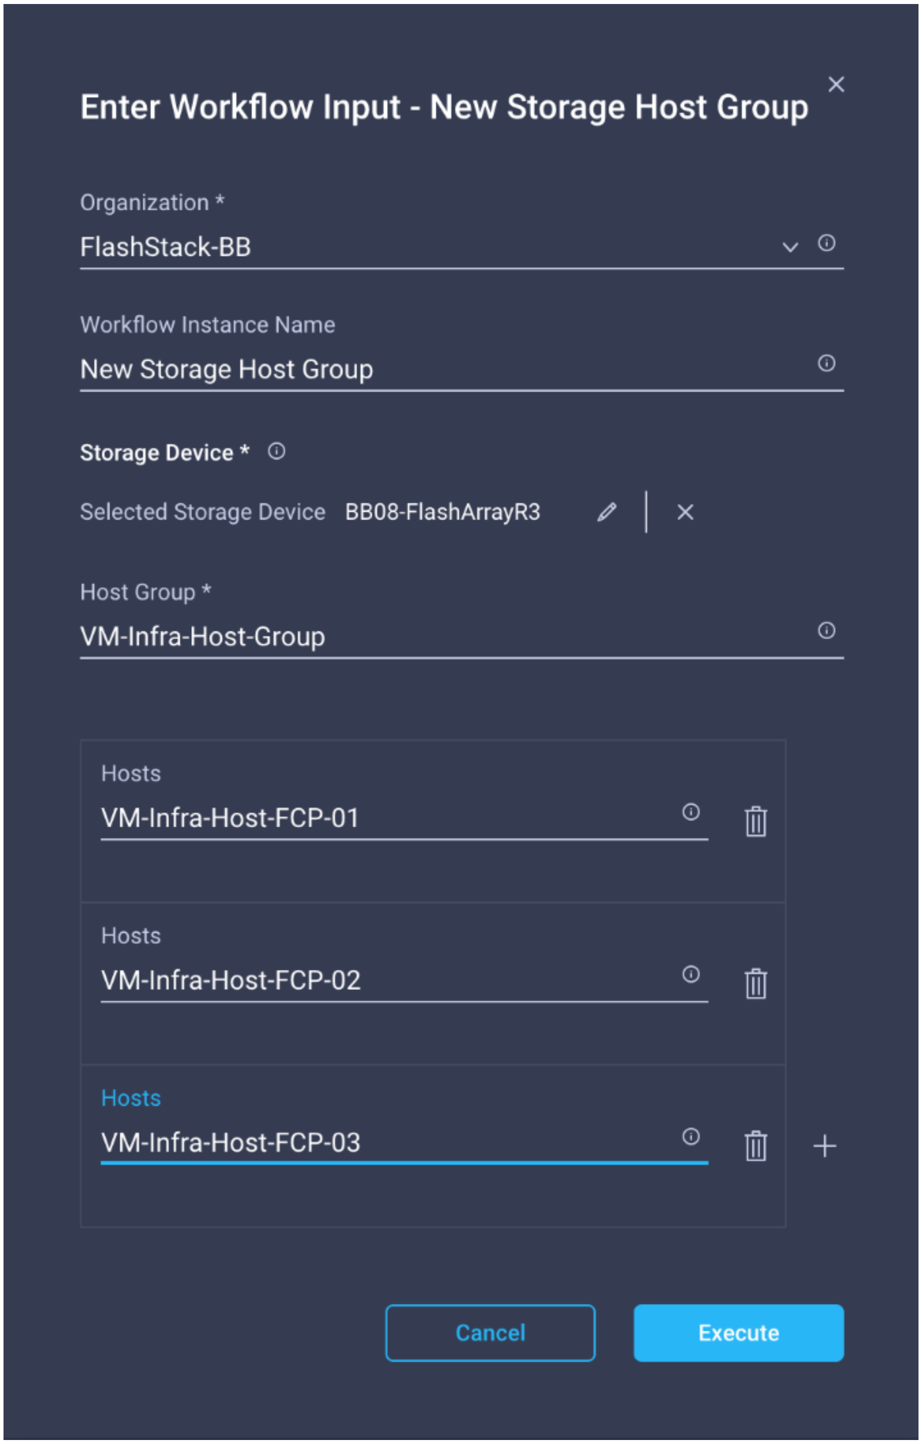

Enter Workflow Input - New Storage Host Group

Organization *
FlashStack-BB

Workflow Instance Name
New Storage Host Group

Storage Device *

Selected Storage Device BB08-FlashArrayR3

Host Group *
VM-Infra-Host-Group

Hosts
VM-Infra-Host-FCP-01

Hosts
VM-Infra-Host-FCP-02

Hosts
VM-Infra-Host-FCP-03

7. Click Execute.

iSCSI Host Registration using Cisco Intersight

To register the iSCSI Host using Cisco Intersight, follow these steps:

1. Click Configure > Orchestration.

2. Click New Storage Host.

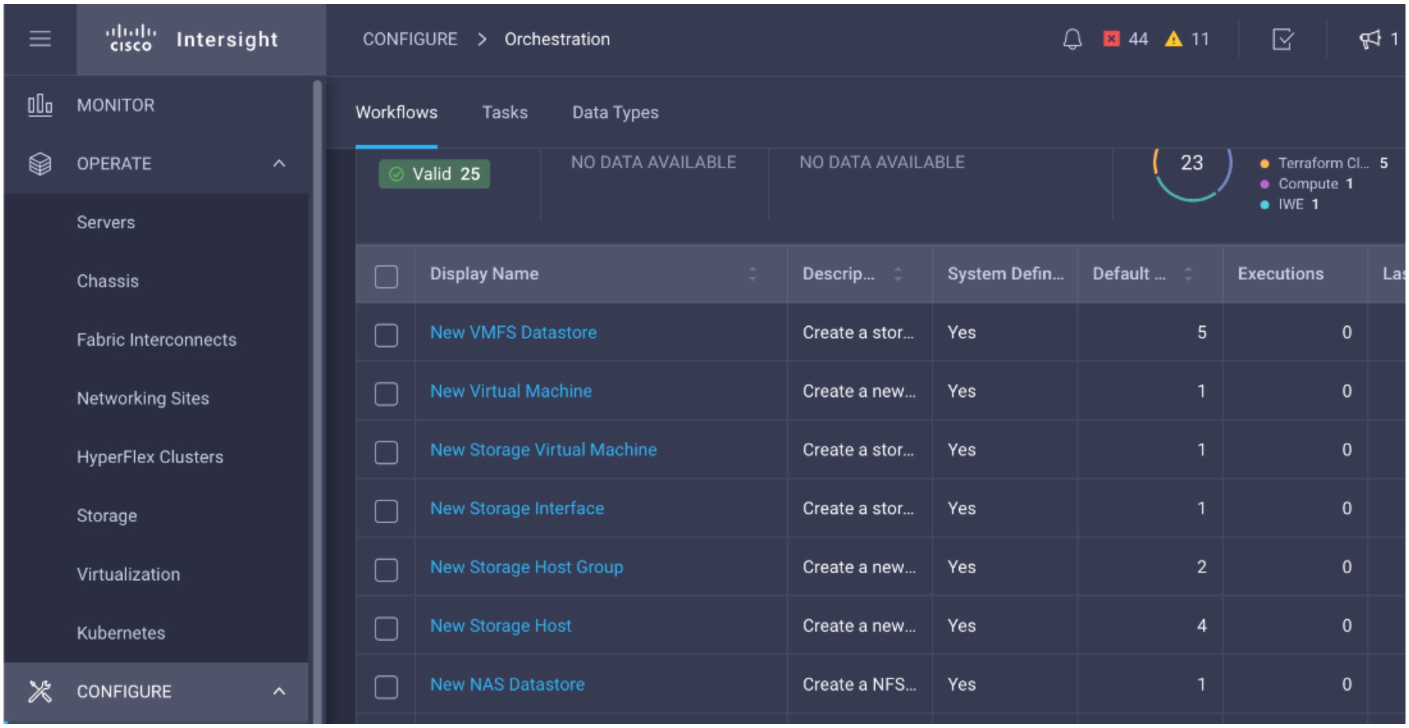

☐	Display Name	Descrip...	System Defin...	Default ...	Executions	Las
☐	New VMFS Datastore	Create a stor...	Yes	5	0	
☐	New Virtual Machine	Create a new...	Yes	1	0	
☐	New Storage Virtual Machine	Create a stor...	Yes	1	0	
☐	New Storage Interface	Create a stor...	Yes	1	0	
☐	New Storage Host Group	Create a new...	Yes	2	0	
☐	New Storage Host	Create a new...	Yes	4	0	
☐	New NAS Datastore	Create a NFS...	Yes	1	0	

3. Click Execute.

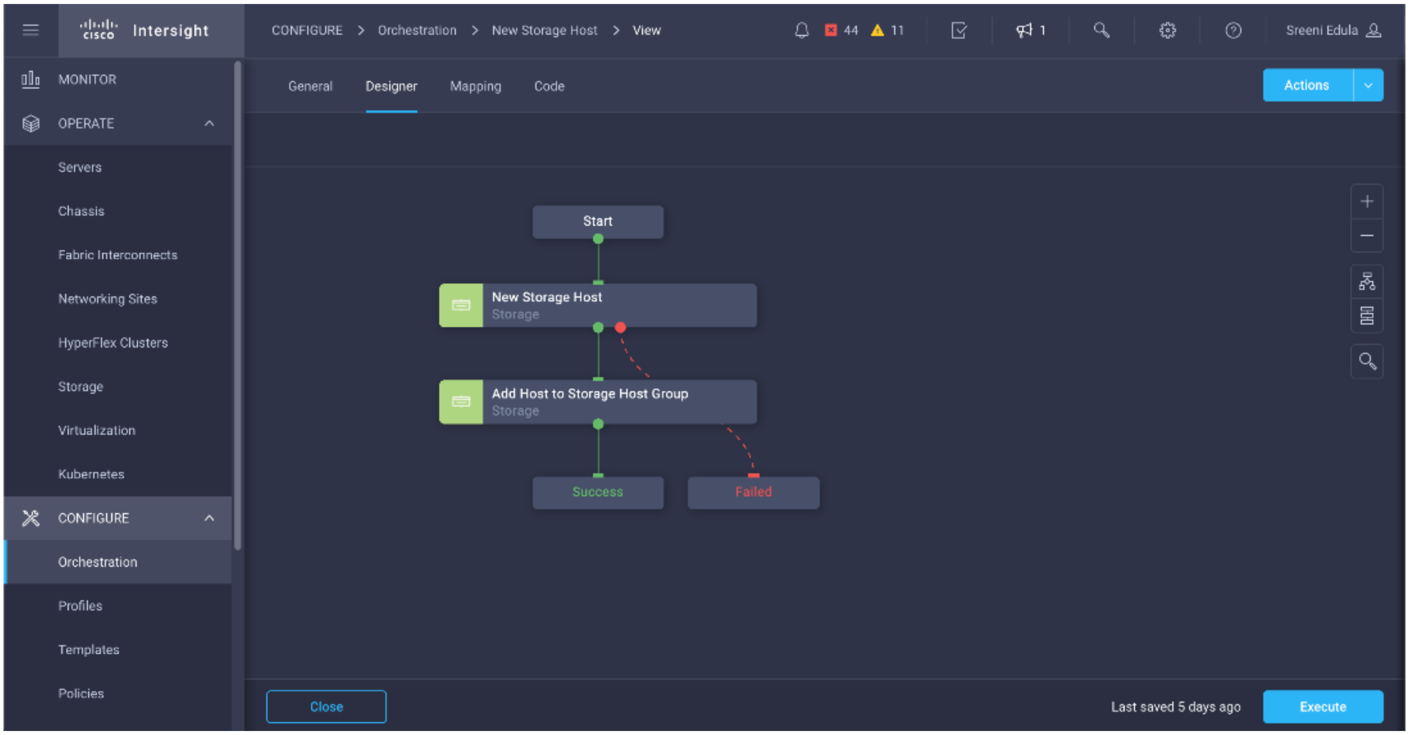

4. Select the appropriate Organization (default by default).

5. Select the appropriate Pure Storage device.

6. Enter the name of the Host name and IQN for host VM-Host-Infra-iSCSI-01.

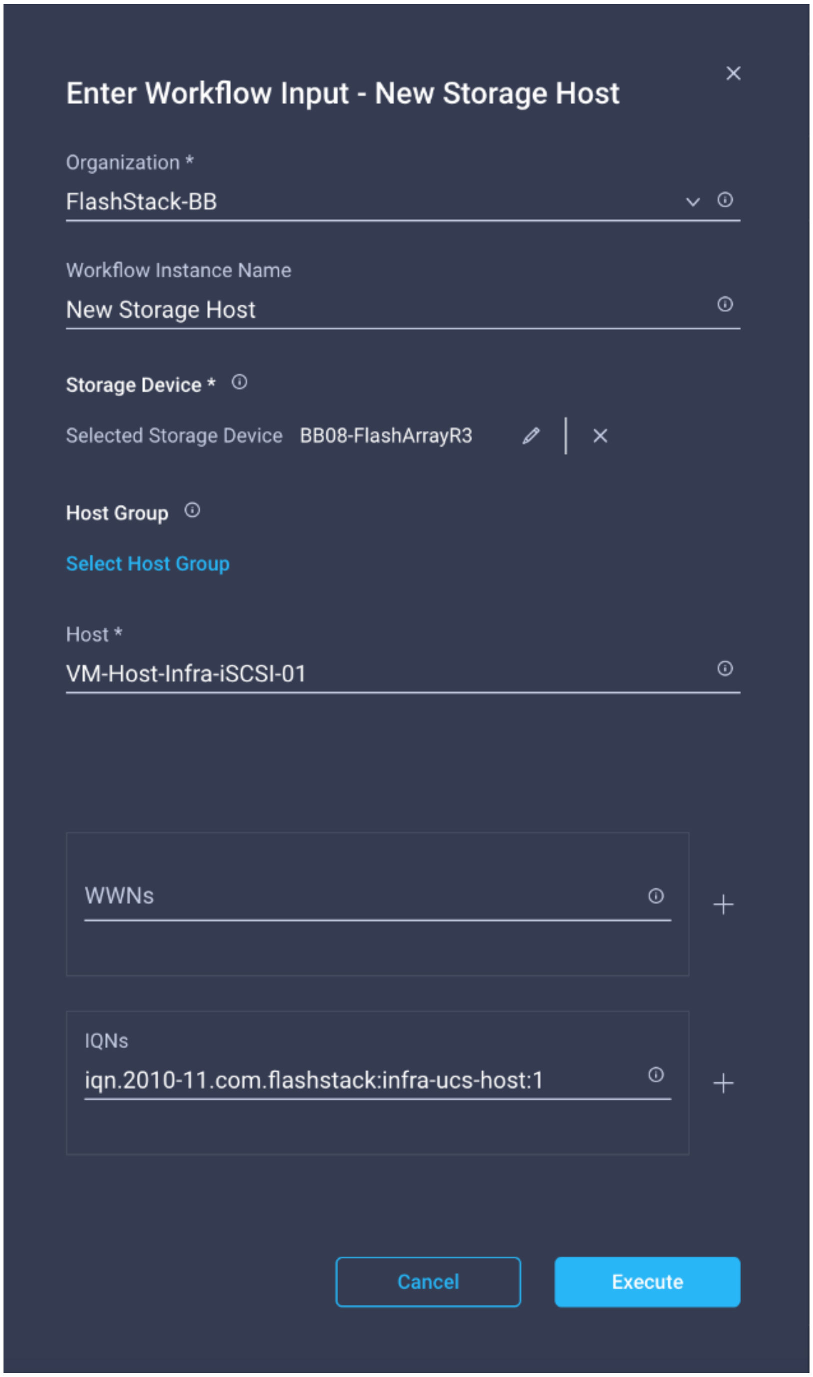

Enter Workflow Input - New Storage Host

Organization * FlashStack-BB

Workflow Instance Name: New Storage Host

Storage Device *

Selected Storage Device BB08-FlashArrayR3

Host Group

Host * VM-Host-Infra-iSCSI-01

IQNs: iqn.2010-11.com.flashstack:infra-ucs-host:1

7. Click Execute.

8. Repeat Steps 2-7 for all host.

Create Host Group using Cisco Intersight

To create a Host group using Cisco Intersight, follow these steps:

1. Click Configure > Orchestration.

2. Click New Storage Host Group.

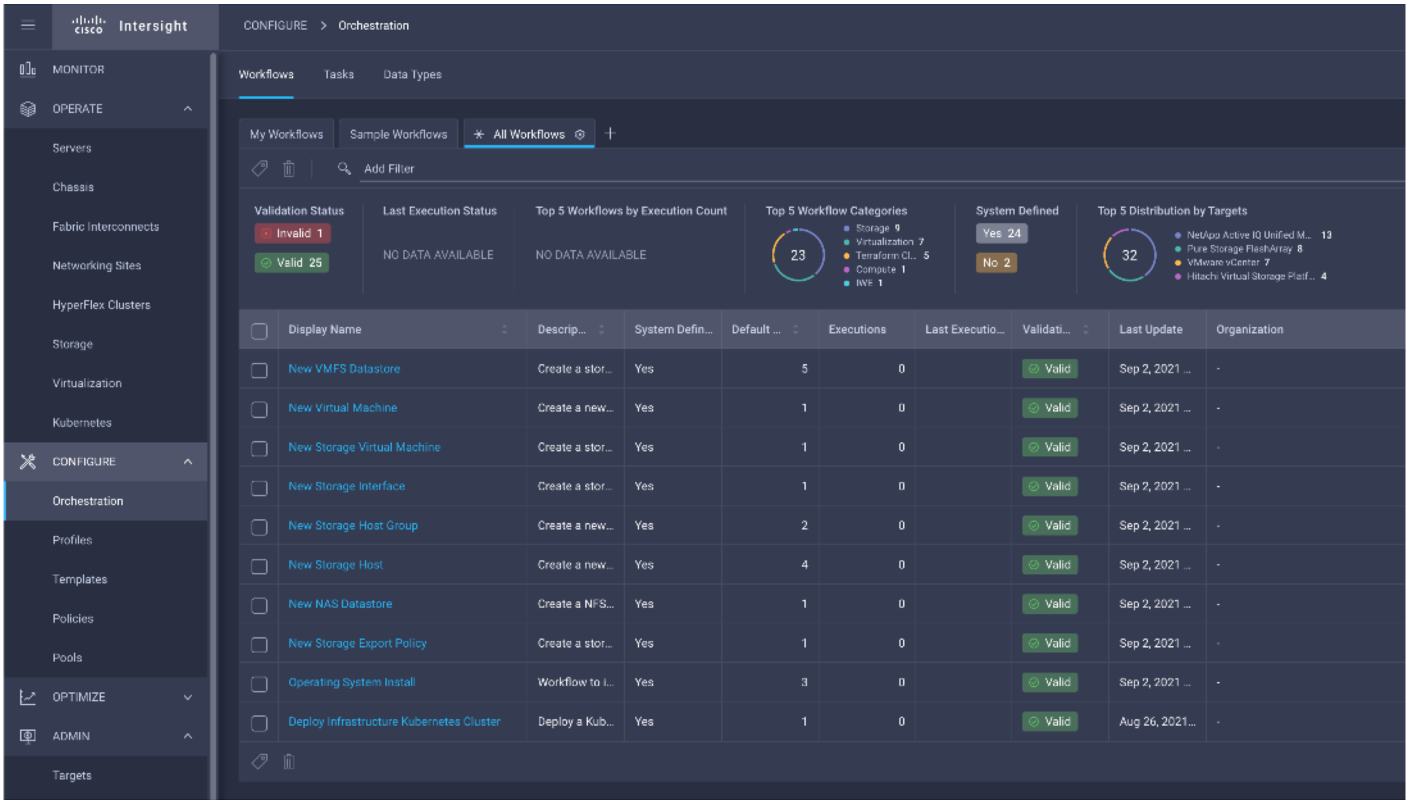

□	Display Name	Descrip...	System Defin...	Default ...	Executions	Last Executio...	Validati...	Last Update	Organization
□	New VMFS Datastore	Create a stor...	Yes	5	0		Valid	Sep 2, 2021 ...	-
□	New Virtual Machine	Create a new...	Yes	1	0		Valid	Sep 2, 2021 ...	-
□	New Storage Virtual Machine	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
□	New Storage Interface	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
□	New Storage Host Group	Create a new...	Yes	2	0		Valid	Sep 2, 2021 ...	-
□	New Storage Host	Create a new...	Yes	4	0		Valid	Sep 2, 2021 ...	-
□	New NAS Datastore	Create a NFS...	Yes	1	0		Valid	Sep 2, 2021 ...	-
□	New Storage Export Policy	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
□	Operating System Install	Workflow to i...	Yes	3	0		Valid	Sep 2, 2021 ...	-
□	Deploy Infrastructure Kubernetes Cluster	Deploy a Kub...	Yes	1	0		Valid	Aug 26, 2021...	-

3. Click Execute.

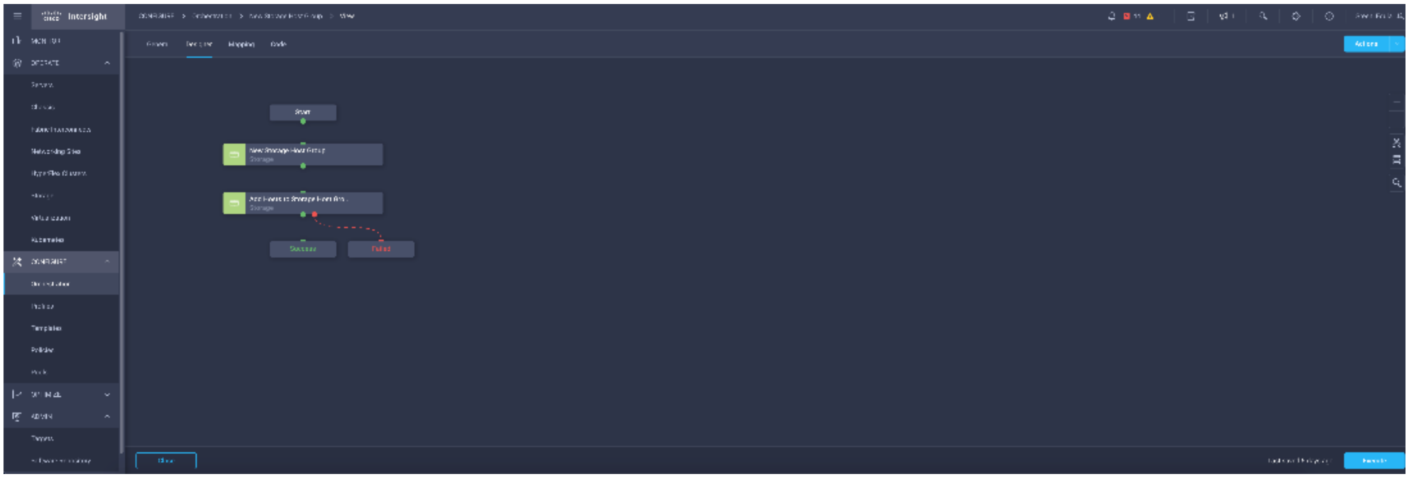

4. Select the appropriate Organization (default by default).

5. Select the appropriate Pure Storage device.

6. Enter the name of the Host Group and of the Hosts created during Host Registration. VM-Infra-Host-iSCSI-01, VM-Infra-Host-iSCSI-02 and VM-Infra-Host-iSCSI-03 are the hosts used in this deployment.

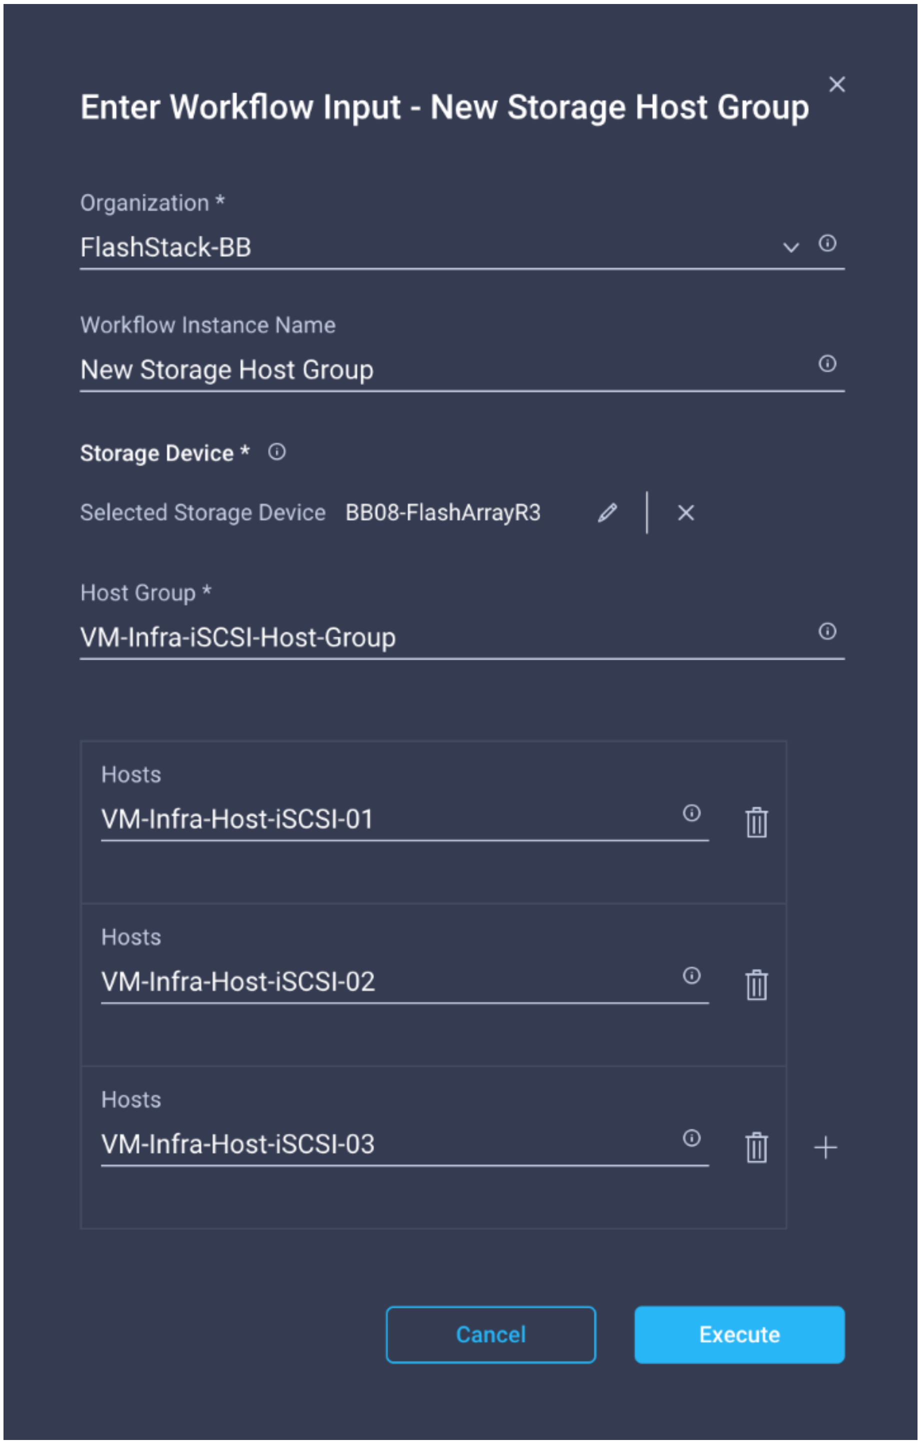

Enter Workflow Input - New Storage Host Group

Organization *
FlashStack-BB

Workflow Instance Name
New Storage Host Group

Storage Device *

Selected Storage Device BB08-FlashArrayR3

Host Group *
VM-Infra-iSCSI-Host-Group

Hosts
VM-Infra-Host-iSCSI-01

Hosts
VM-Infra-Host-iSCSI-02

Hosts
VM-Infra-Host-iSCSI-03

7. Click Execute.

Cisco Infrastructure Firmware Upgrade (Fabric Interconnects) using Cisco Intersight

To upgrade Cisco UCS Fabric Interconnects using Cisco Intersight, follow these steps in Intersight SaaS Portal:

1. From the left navigation pane, click Fabric Interconnects, select a Fabric Interconnect, and perform an Upgrade Firmware action on it.

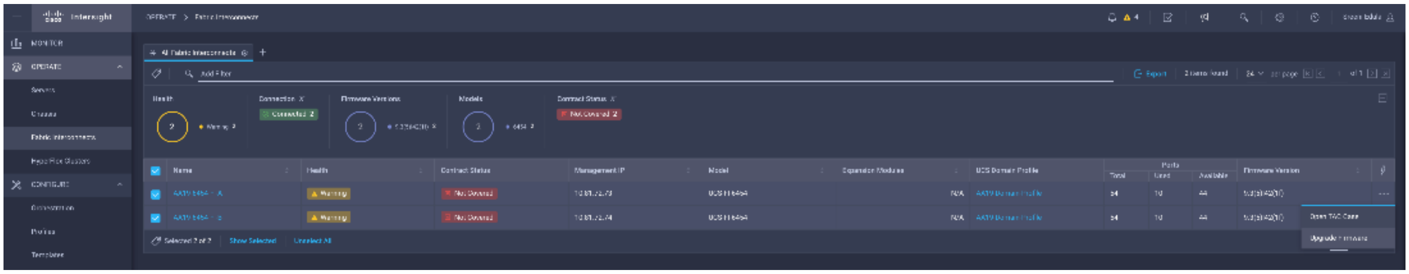

OPERATE > Fabric Interconnects

All Fabric Interconnects

Health: 2 ● Warning 2 | Connection: Connected 2 | Firmware Versions: 2 ● 9.3(5)I42(1f) 2 | Models: 2 ● 6454 2 | Contract Status: Not Covered 2

☑	Name	Health	Contract Status	Management IP	Model	Expansion Modules	UCS Domain Profile	Ports Total	Ports Used	Ports Available	Firmware Version
☑	AA19-6454 - A	Warning	Not Covered	10.81.72.73	UCS-FI-6454	N/A	AA19 Domain Profile	54	10	44	9.3(5)I42(1f)
☑	AA19-6454 - B	Warning	Not Covered	10.81.72.74	UCS-FI-6454	N/A	AA19 Domain Profile	54	10	44	9.3(5)I42(1f)

Selected 2 of 2 Show Selected Unselect All

2. On the Upgrade Firmware page, click Start.

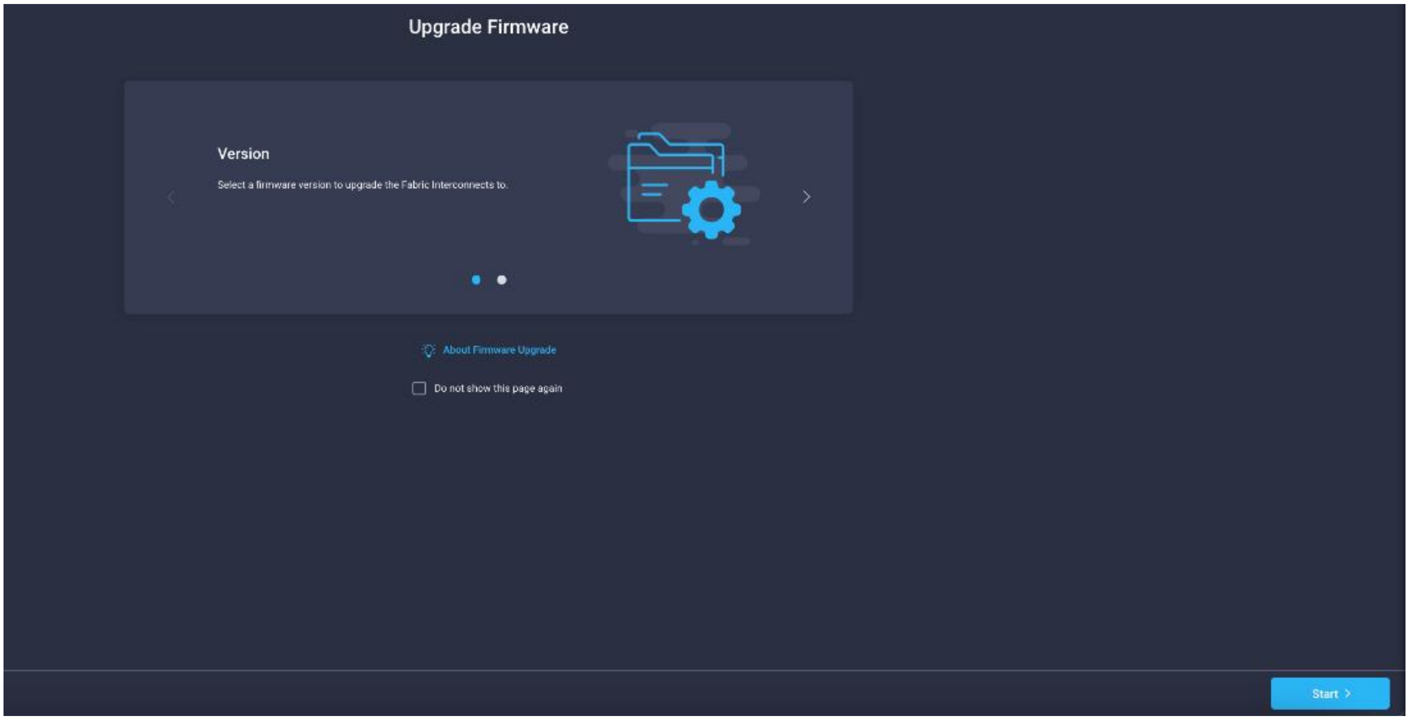

Upgrade Firmware

Version

Select a firmware version to upgrade the Fabric Interconnects to.

About Firmware Upgrade

□ Do not show this page again

3. On the General page, confirm selection of the switch Domain and click Next.

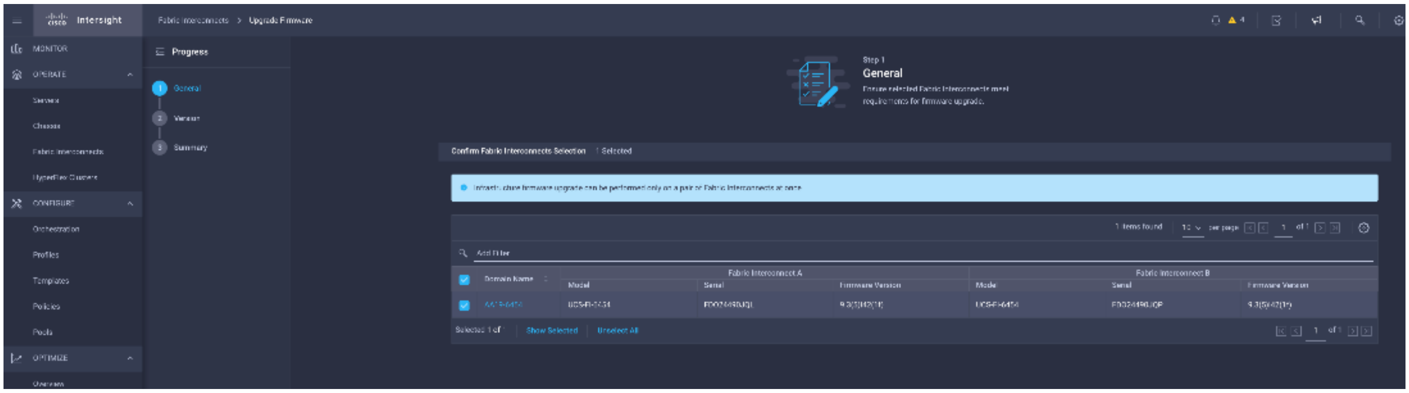

Fabric Interconnects > Upgrade Firmware

Step 1

General

Ensure selected Fabric Interconnects meet requirements for firmware upgrade.

Confirm Fabric Interconnects Selection 1 Selected

Infrastructure firmware upgrade can be performed only on a pair of Fabric Interconnects at once.

1 items found 10 per page 1 of 1

☑	Domain Name	Fabric Interconnect A			Fabric Interconnect B		
		Model	Serial	Firmware Version	Model	Serial	Firmware Version
☑	AA19-6454	UCS-FI-6454	FDO24490JQL	9.3(5)I42(1f)	UCS-FI-6454	FDO24490JQP	9.3(5)I42(1f)

Selected 1 of 1 Show Selected Unselect All

4. On the Version page, select the fabric firmware bundle to which the Fabric Interconnects need to be upgraded, and click Next.

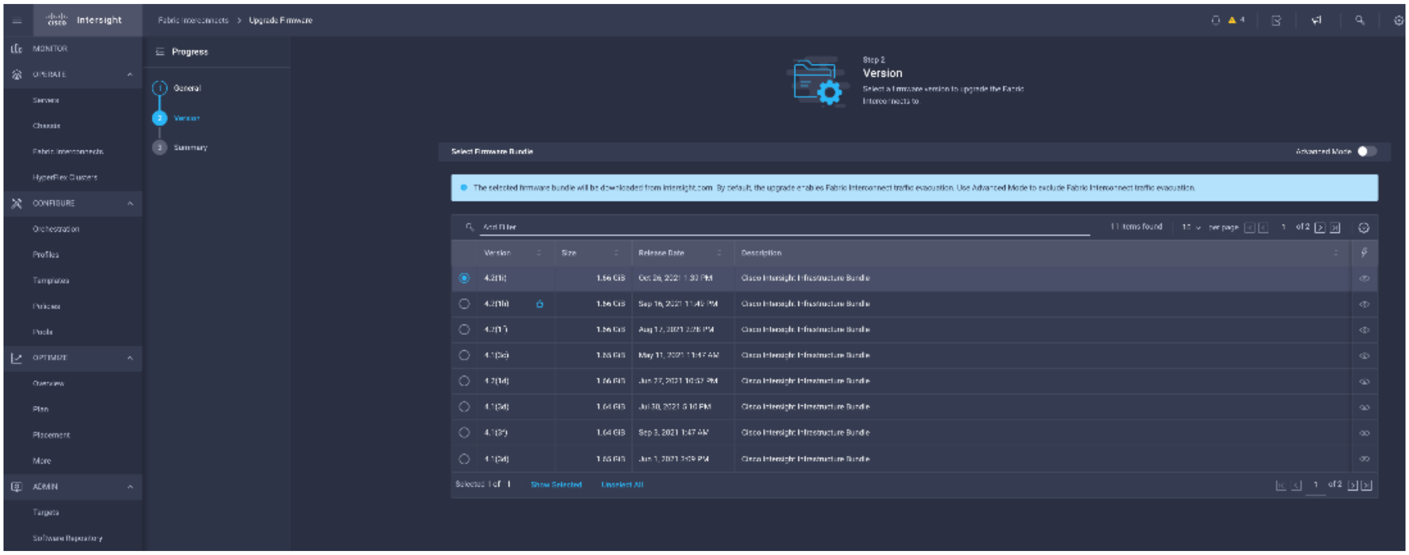

Step 2

Version

Select a firmware version to upgrade the Fabric Interconnects to.

Select Firmware Bundle

Advanced Mode

The selected firmware bundle will be downloaded from intersight.com. By default, the upgrade enables Fabric Interconnect traffic evacuation. Use Advanced Mode to exclude Fabric Interconnect traffic evacuation.

Add Filter — 11 items found — 10 per page — 1 of 2

	Version	Size	Release Date	Description
●	4.2(1i)	1.66 GiB	Oct 26, 2021 1:30 PM	Cisco Intersight Infrastructure Bundle
○	4.2(1h)	1.56 GiB	Sep 16, 2021 11:49 PM	Cisco Intersight Infrastructure Bundle
○	4.2(1f)	1.56 GiB	Aug 17, 2021 3:26 PM	Cisco Intersight Infrastructure Bundle
○	4.1(3e)	1.65 GiB	May 11, 2021 11:47 AM	Cisco Intersight Infrastructure Bundle
○	4.2(1d)	1.66 GiB	Jun 27, 2021 10:57 PM	Cisco Intersight Infrastructure Bundle
○	4.1(3d)	1.64 GiB	Jul 30, 2021 5:10 PM	Cisco Intersight Infrastructure Bundle
○	4.1(3f)	1.64 GiB	Sep 3, 2021 1:47 AM	Cisco Intersight Infrastructure Bundle
○	4.1(3d)	1.65 GiB	Jun 1, 2021 2:09 PM	Cisco Intersight Infrastructure Bundle

Selected 1 of 1 Show Selected Unselect All — 1 of 2

5. On the Summary screen, verify the summary of the selected switches, the firmware version running on them, and the firmware version to which they will be upgraded, and click Upgrade.

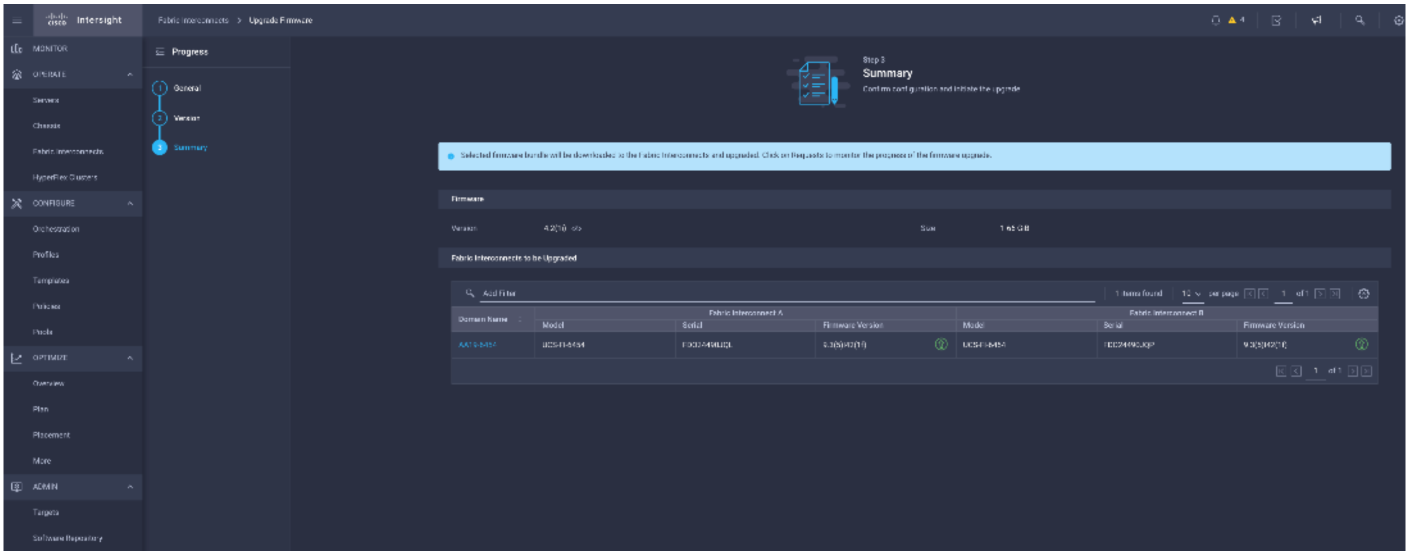

Step 3

Summary

Confirm configuration and initiate the upgrade

Selected firmware bundle will be downloaded to the Fabric Interconnects and upgraded. Click on Requests to monitor the progress of the firmware upgrade.

Firmware

Version 4.2(1i) Size 1.65 GiB

Fabric Interconnects to be Upgraded

Add Filter — 1 items found — 10 per page — 1 of 1

Domain Name	Fabric Interconnect A			Fabric Interconnect B		
	Model	Serial	Firmware Version	Model	Serial	Firmware Version
AA19-6454	UCS-FI-6454	FDO2449LLQL	9.3(5)I42(1f)	UCS-FI-6454	FDO24490JQP	9.3(5)I42(1f)

6. Confirm the upgrade request.

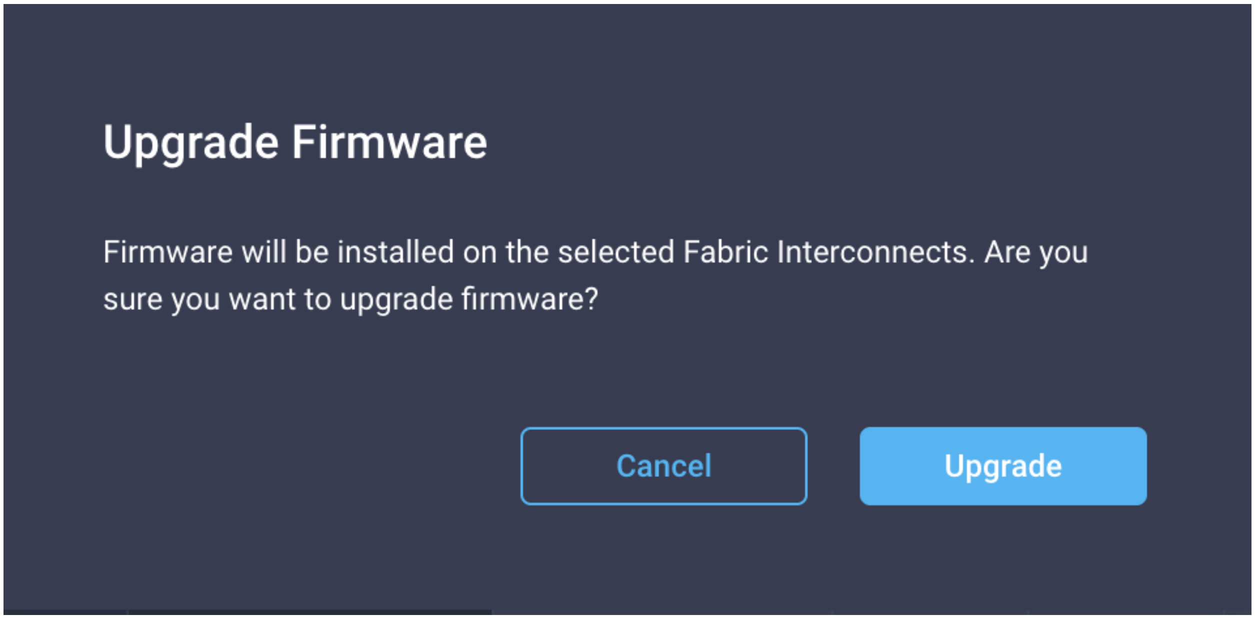

Upgrade Firmware

Firmware will be installed on the selected Fabric Interconnects. Are you sure you want to upgrade firmware?

The firmware upgrade workflow begins.

7. You can check the status of the upgrade workflow in the Execution Flow pane. Acknowledge any messages in the Execution Flow pane and click Continue to proceed with the upgrade.

8. Click Continue.

9. Verify if the upgrade is successful.

Cisco UCS Server Upgrades

To upgrade the Cisco UCS Servers using Intersight, follow these steps in Intersight SaaS Portal:

Note: Only servers in associated state can be upgraded.

Note: Servers associated with server profiles bound to updating templates cannot be upgraded.

Note: Servers associated with global server profiles cannot be upgraded.

1. From the left navigation pane, click Servers, select a server, and perform an Upgrade Firmware action on it.

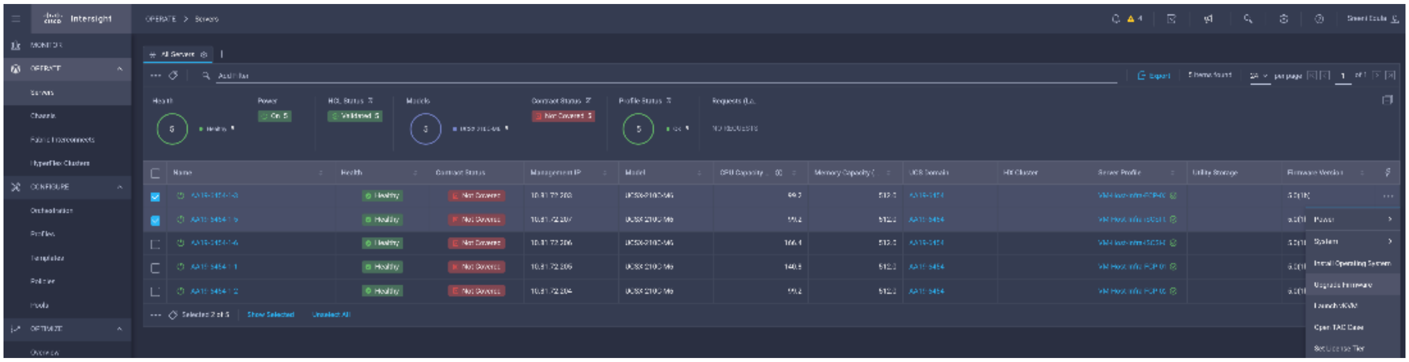

2. On the Upgrade Firmware page, click Start.

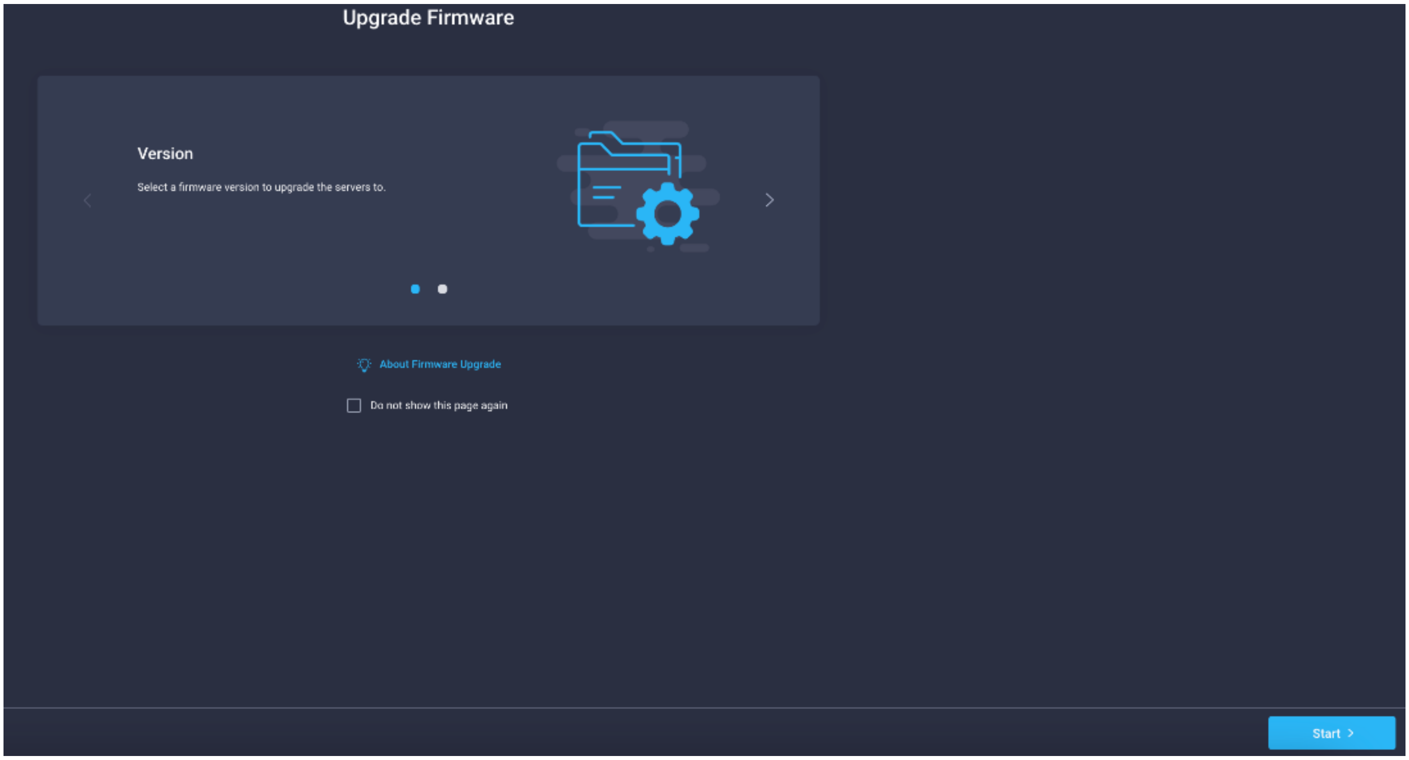

Upgrade Firmware

Version

Select a firmware version to upgrade the servers to.

About Firmware Upgrade

□ Do not show this page again

3. On the General page, confirm selection of the server and click Next.

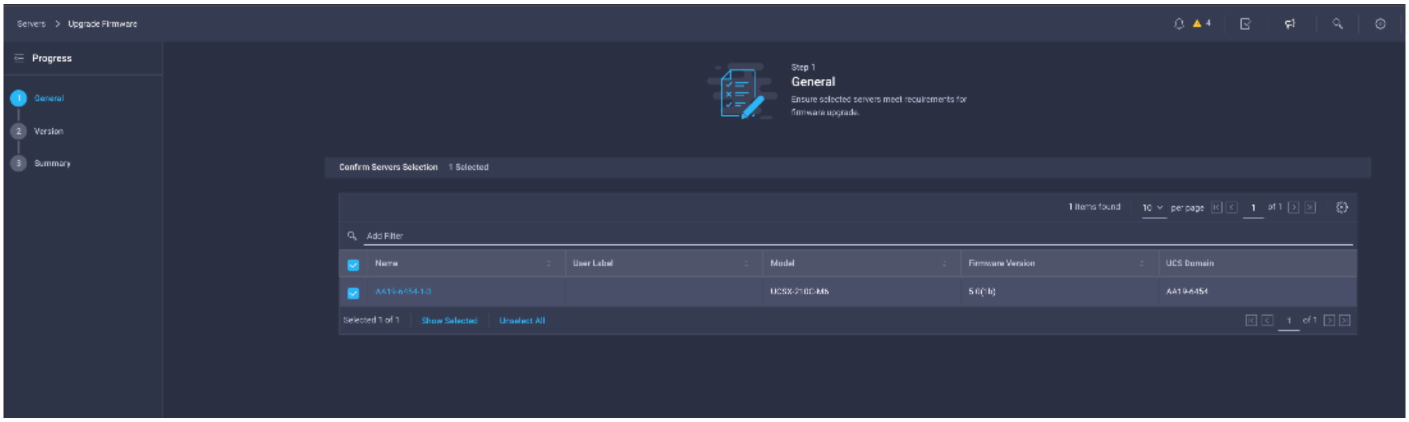

Servers > Upgrade Firmware

Progress

1 General
2 Version
3 Summary

Step 1

General

Ensure selected servers meet requirements for firmware upgrade.

Confirm Servers Selection 1 Selected

1 items found 10 per page 1 of 1

Add Filter

☑	Name	User Label	Model	Firmware Version	UCS Domain
☑	AA19-6454-1-3		UCSX-210C-M6	5.0(1b)	AA19-6454

Selected 1 of 1 Show Selected Unselect All

4. On the Version page, select the fabric firmware bundle to which the Fabric Interconnects need to be upgraded, and click Next.

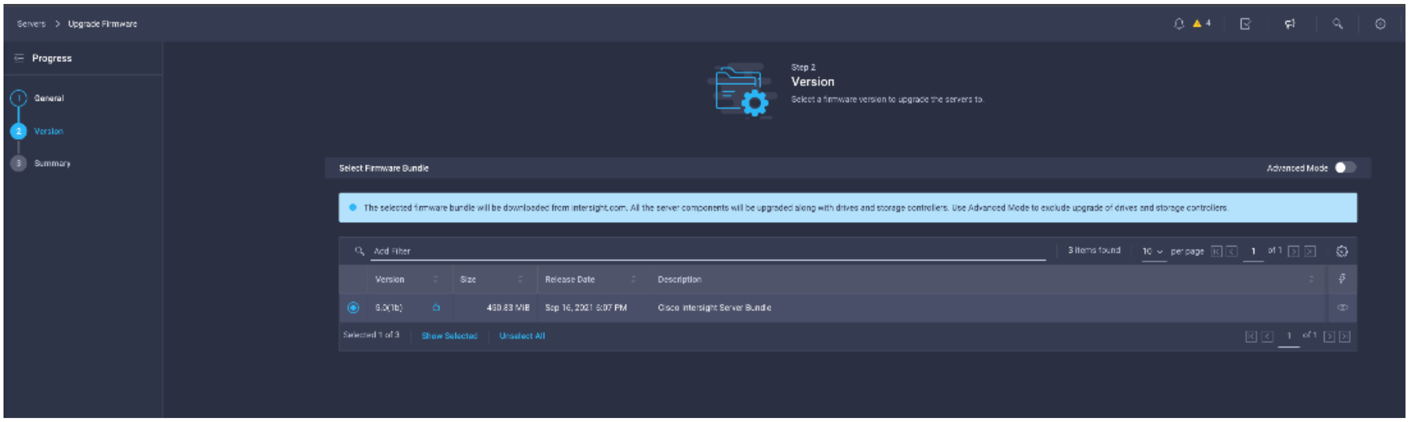

Step 2
Version
Select a firmware version to upgrade the servers to.

Select Firmware Bundle — Advanced Mode

The selected firmware bundle will be downloaded from intersight.com. All the server components will be upgraded along with drives and storage controllers. Use Advanced Mode to exclude upgrade of drives and storage controllers.

Add Filter — 3 items found — 10 per page — 1 of 1

	Version	Size	Release Date	Description
●	5.0(1b)	450.83 MiB	Sep 16, 2021 6:07 PM	Cisco Intersight Server Bundle

Selected 1 of 3 — Show Selected — Unselect All

5. On the Summary screen, verify the summary of the selected switches, the firmware version running on them, and the firmware version to which they will be upgraded, and click Upgrade.

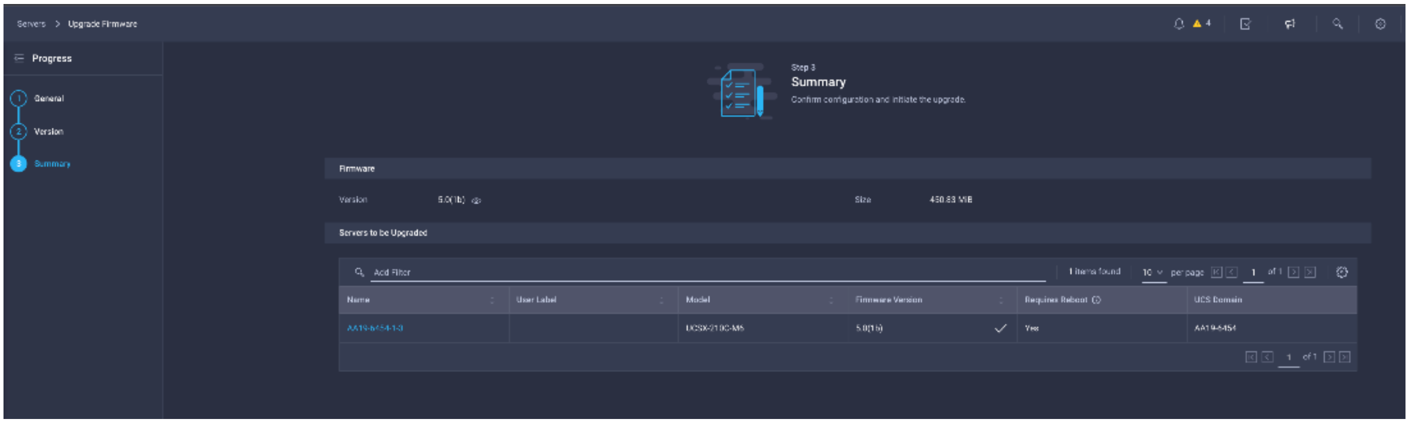

Step 3
Summary
Confirm configuration and initiate the upgrade.

Firmware

Version 5.0(1b) — Size 450.83 MiB

Servers to be Upgraded

Add Filter — 1 items found — 10 per page — 1 of 1

Name	User Label	Model	Firmware Version	Requires Reboot	UCS Domain
AA19-6454-1-3		UCSX-210C-M5	5.0(1b) ✓	Yes	AA19-6454

6. Select Reboot Immediately to Begin Upgrade and Confirm the upgrade request and monitor the process for successful upgrade.

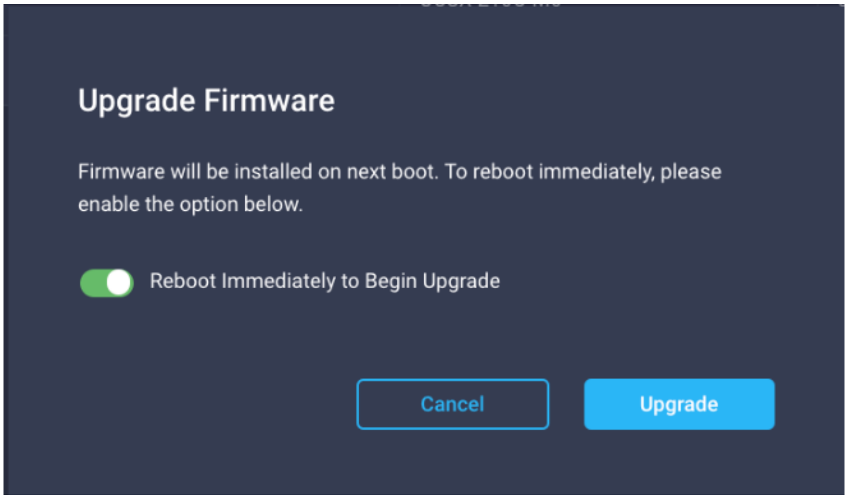

Upgrade Firmware

Firmware will be installed on next boot. To reboot immediately, please enable the option below.

● Reboot Immediately to Begin Upgrade

Pure Storage vSphere Client Plugin

The Pure Storage Plugin for the vSphere Client provides the ability to VMware users to have insight into and control of their Pure Storage FlashArray environment while directly logged into the vSphere Client. The Pure Storage plugin extends the vSphere Client interface to include environmental statistics and objects that underpin the VMware objects in use and to provision new resources as needed.

The Pure Storage vSphere Client Plugin will be accessible through the vSphere Client after registration through the Pure Storage Web Portal.

To access the Pure Storage vSphere Client Plugin, follow these steps:

1. Go to Settings > Software.

2. Click the edit icon in the vSphere Plugin panel.

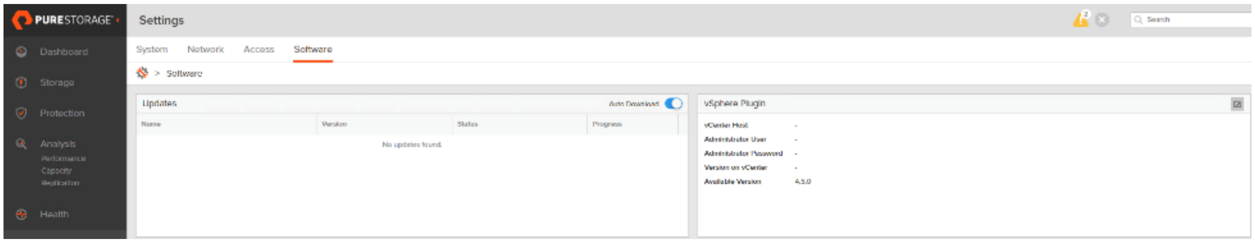

No updates found.

vSphere Plugin

vCenter Host	-
Administrator User	-
Administrator Password	-
Version on vCenter	-
Available Version	4.5.0

3. Enter the vCenter information in the pop-up window and click Save.

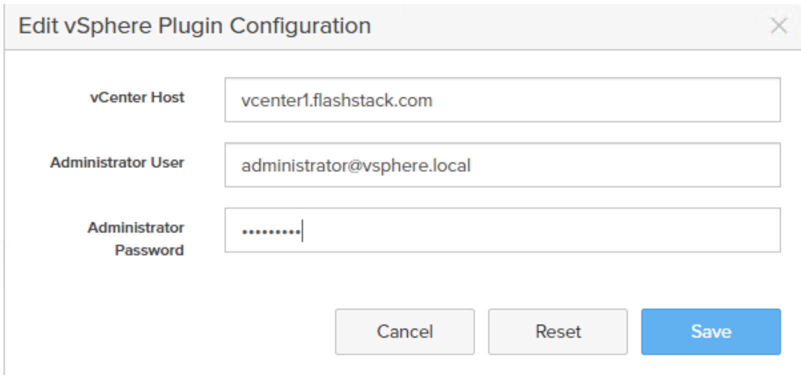

Edit vSphere Plugin Configuration

vCenter Host	vcenter1.flashstack.com
Administrator User	administrator@vsphere.local
Administrator Password	•••••••••

4. After the discovery completes, click Install.

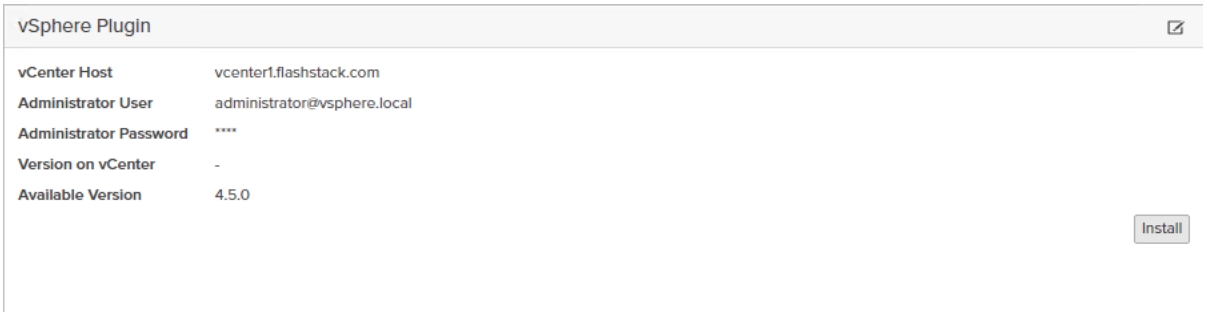

vSphere Plugin

vCenter Host	vcenter1.flashstack.com
Administrator User	administrator@vsphere.local
Administrator Password	****
Version on vCenter	-
Available Version	4.5.0

5. In vCenter, select Pure Storage from the Menu.

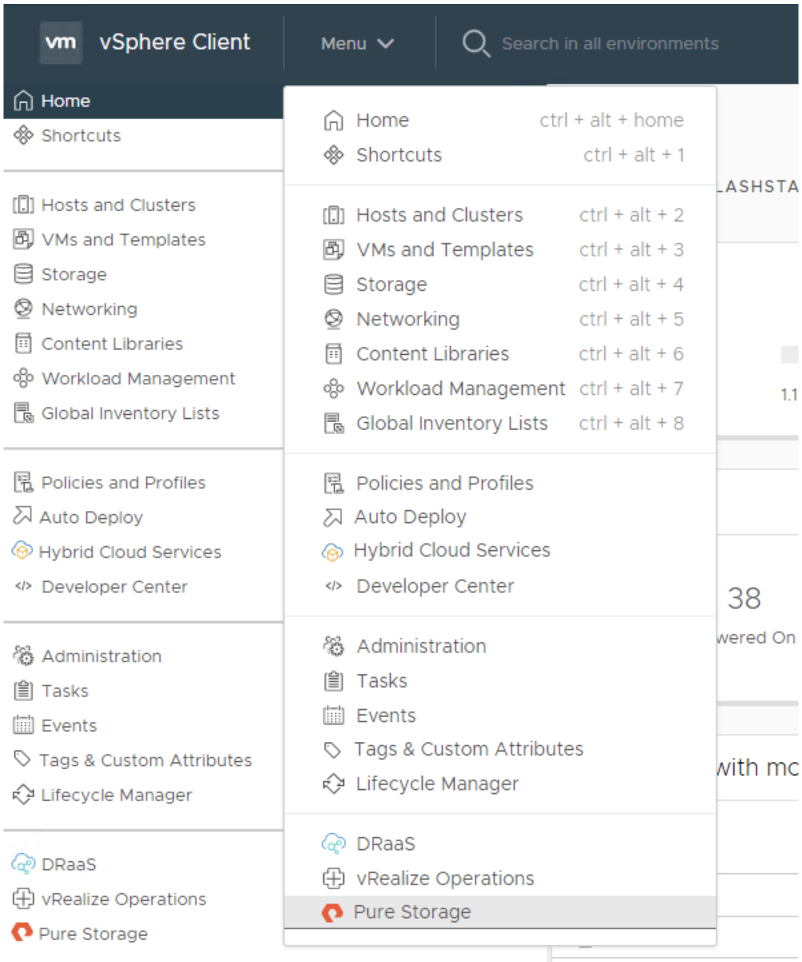

6. Click Authenticate with Pure1.

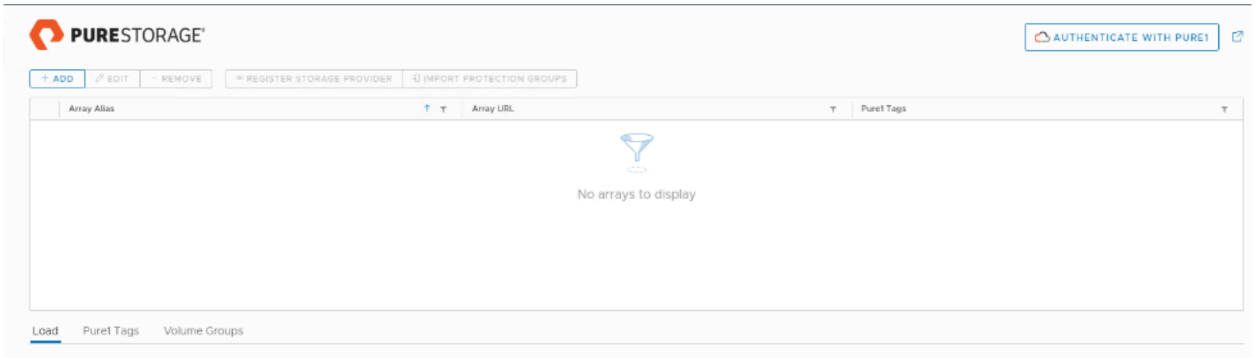

7. Input your Pure1 JWT (link).

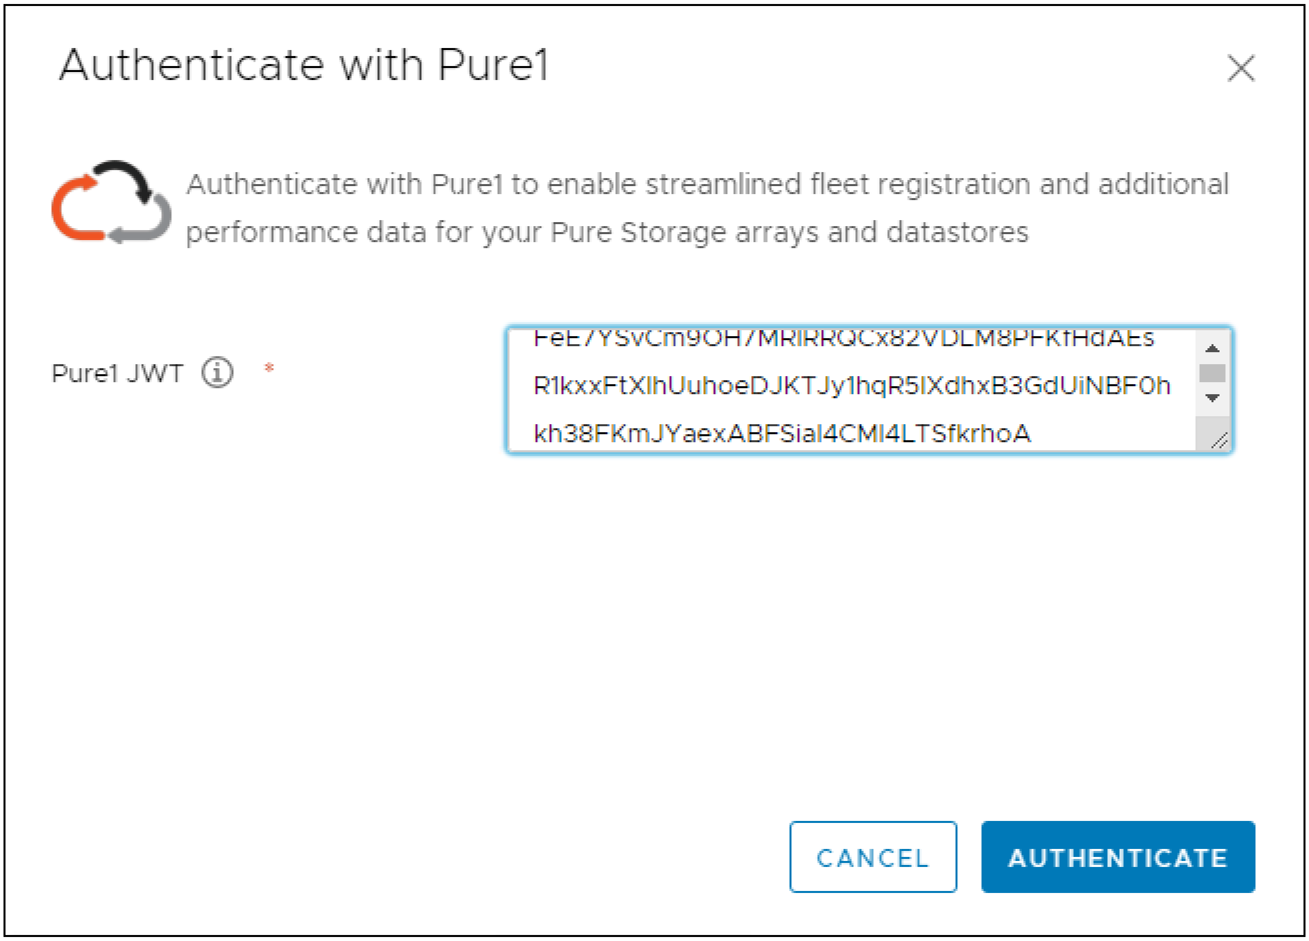

Authenticate with Pure1

Authenticate with Pure1 to enable streamlined fleet registration and additional performance data for your Pure Storage arrays and datastores

Pure1 JWT *

FeE7YSvCm9OH7MRlRRQCx82VDLM8PFKfHdAEs
R1kxxFtXlhUuhoeDJKTJy1hqR5lXdhxB3GdUiNBF0h
kh38FKmJYaexABFSial4CMl4LTSfkrhoA

8. Click Authenticate.

9. Click Add.

10. Click Import Arrays from Pure1 and input the Username and Password.

11. Click Import Arrays from Pure1 and input the Username and Password.

12. Click Done.

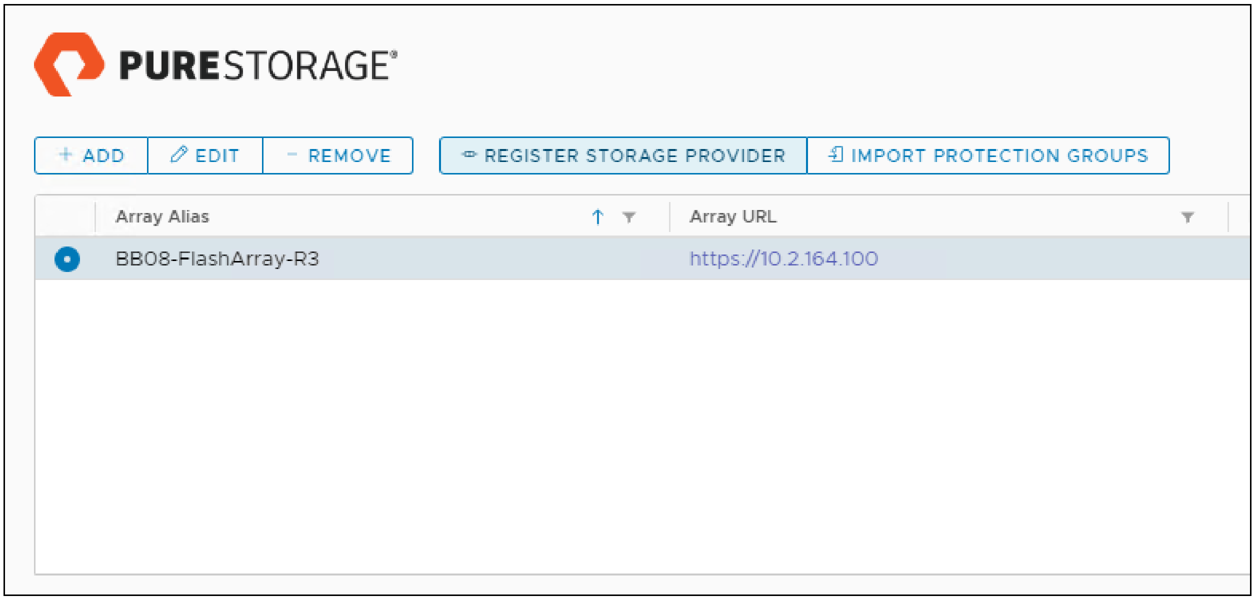

	Array Alias	Array URL
●	BB08-FlashArray-R3	https://10.2.164.100

13. Alternatively, provide array details in the Add a Single Array tab to add the Array manually.

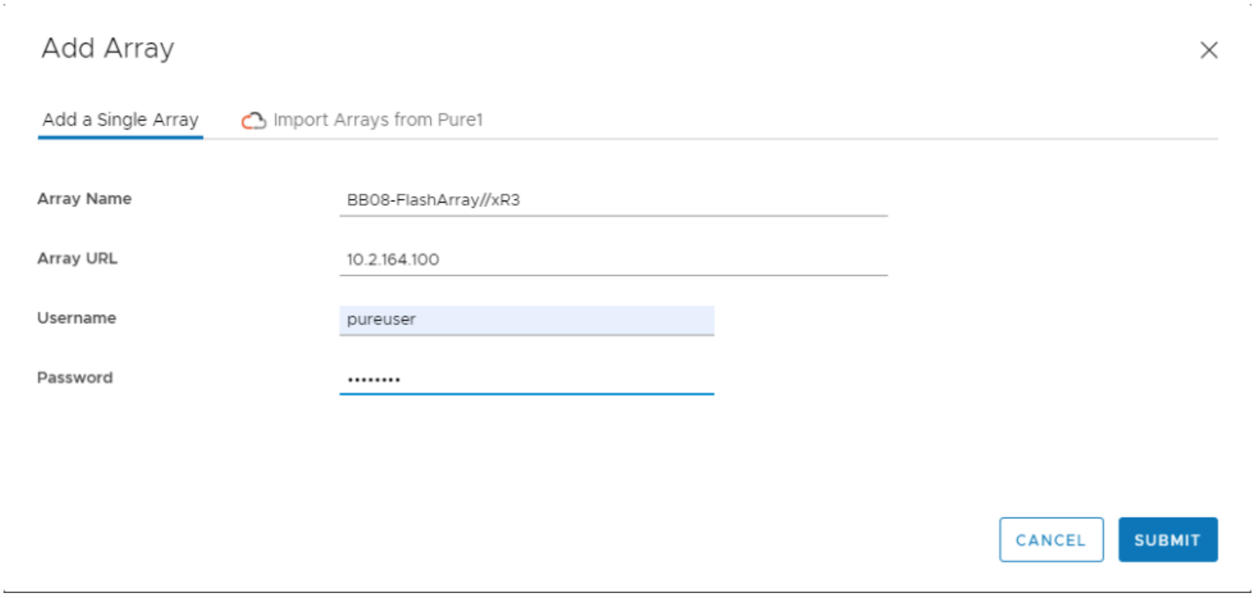

Array Name	BB08-FlashArray//xR3
Array URL	10.2.164.100
Username	pureuser
Password	••••••••

14. Select the newly added array.

15. Click Register Storage Provider.

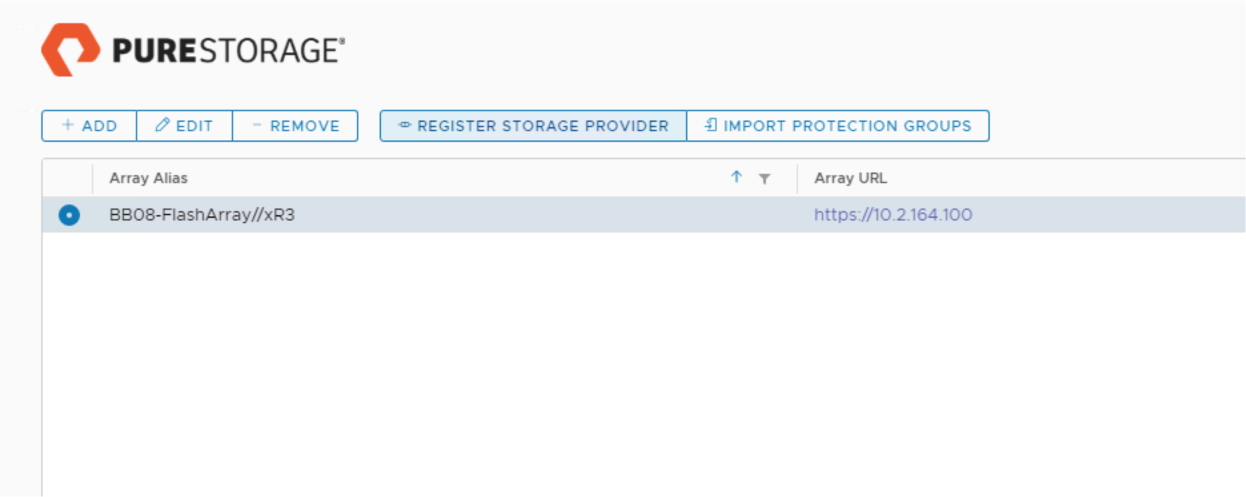

Array Alias	Array URL
BB08-FlashArray//xR3	https://10.2.164.100

16. Enter Username and Password.

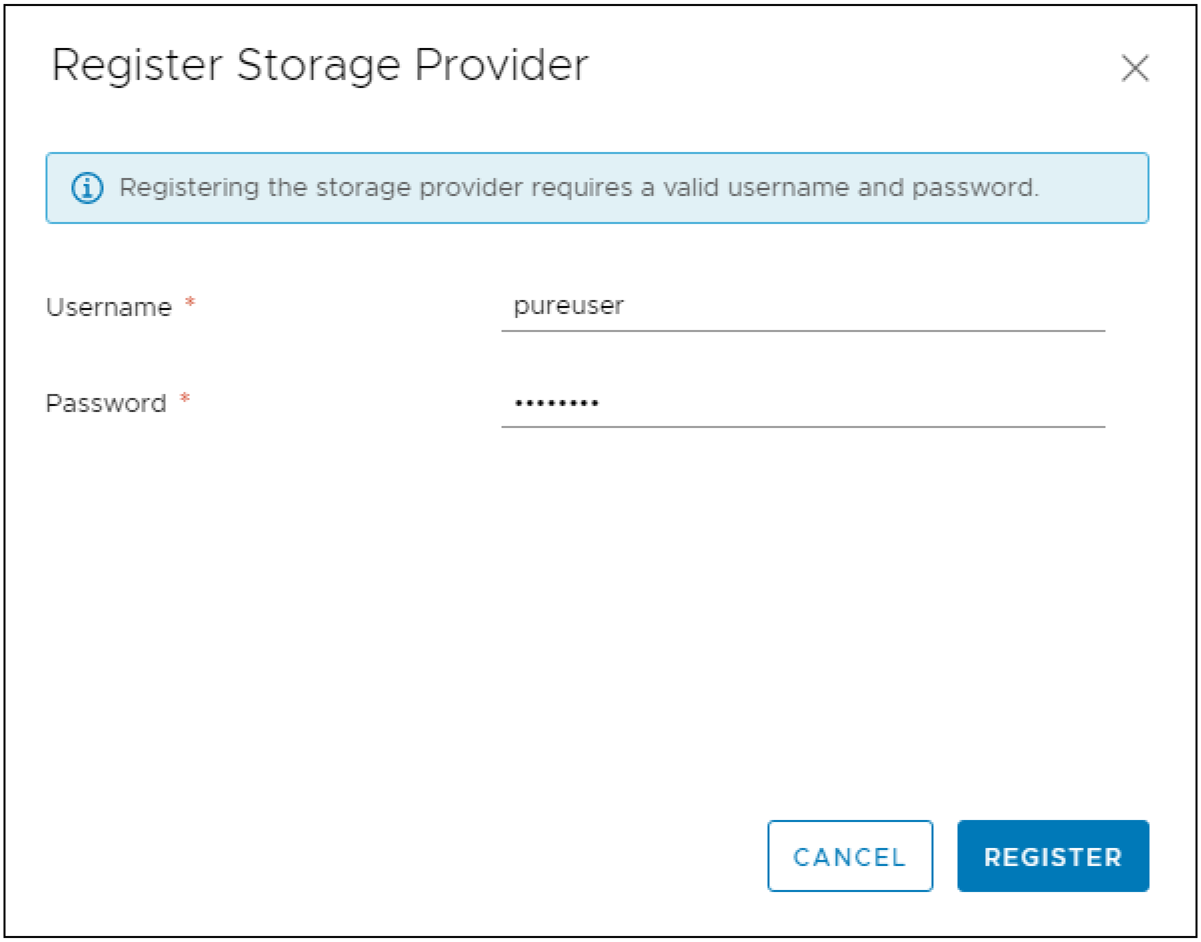

Register Storage Provider

Registering the storage provider requires a valid username and password.

Username * pureuser

Password * •••••••••

17. Click Register.

Create VMDS Datastore using Pure vSphere Plugin

To create VMDS datastore using the Pure vSphere plugin, follow these steps:

1. In vCenter, click Host and Clusters.

2. Right-click the FlashStack Cluster and Select Pure Storage > Create Datastore.

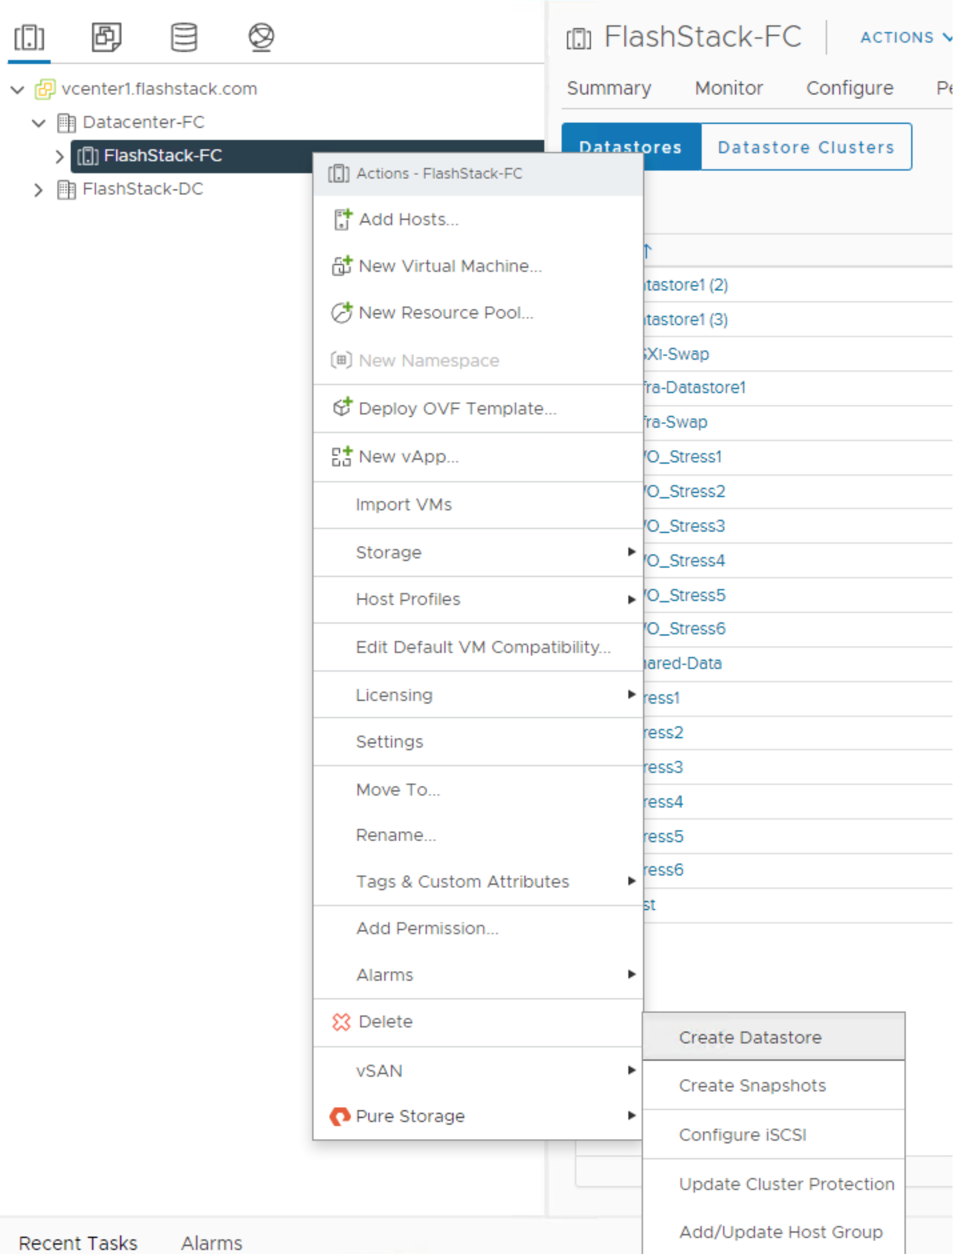

3. Click VMFS.

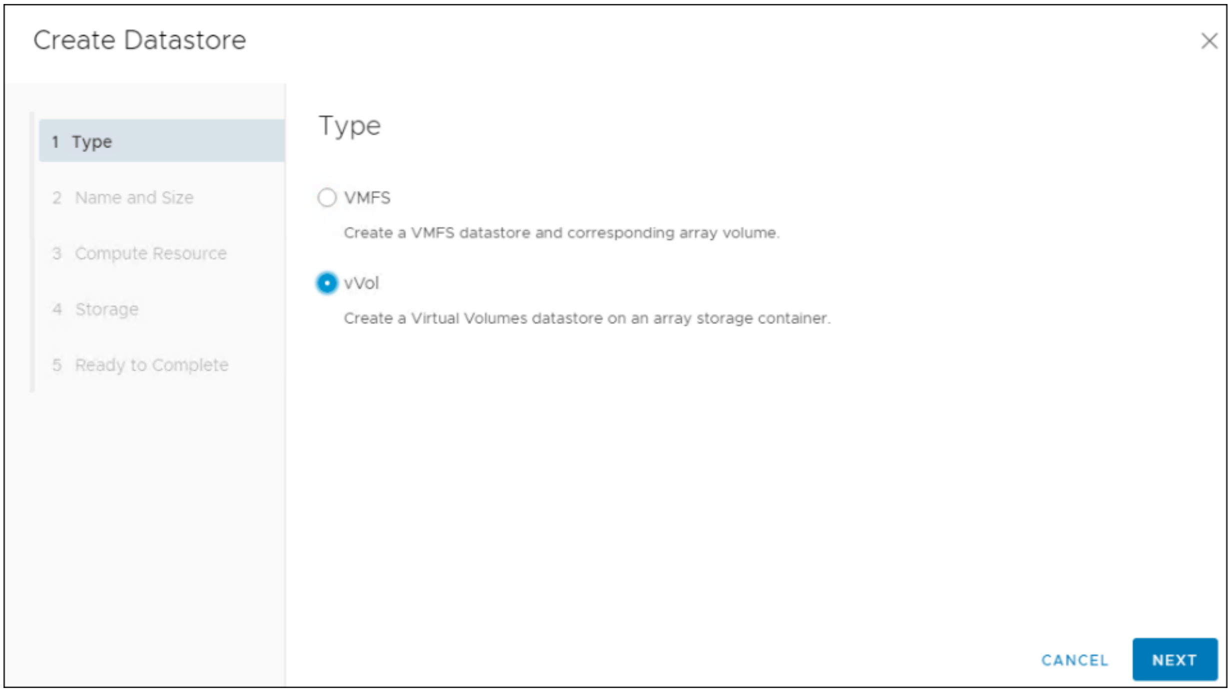
Create Datastore
1 Type
2 Name and Size
3 Compute Resource
4 Storage
5 Ready to Complete
Type
○ VMFS
Create a VMFS datastore and corresponding array volume.
● vVol
Create a Virtual Volumes datastore on an array storage container.
CANCEL
NEXT

4. Click Next.

5. Keep VMFS 6 selected.

6. Click Next.

7. Enter a Datastore Name and Datastore Size.

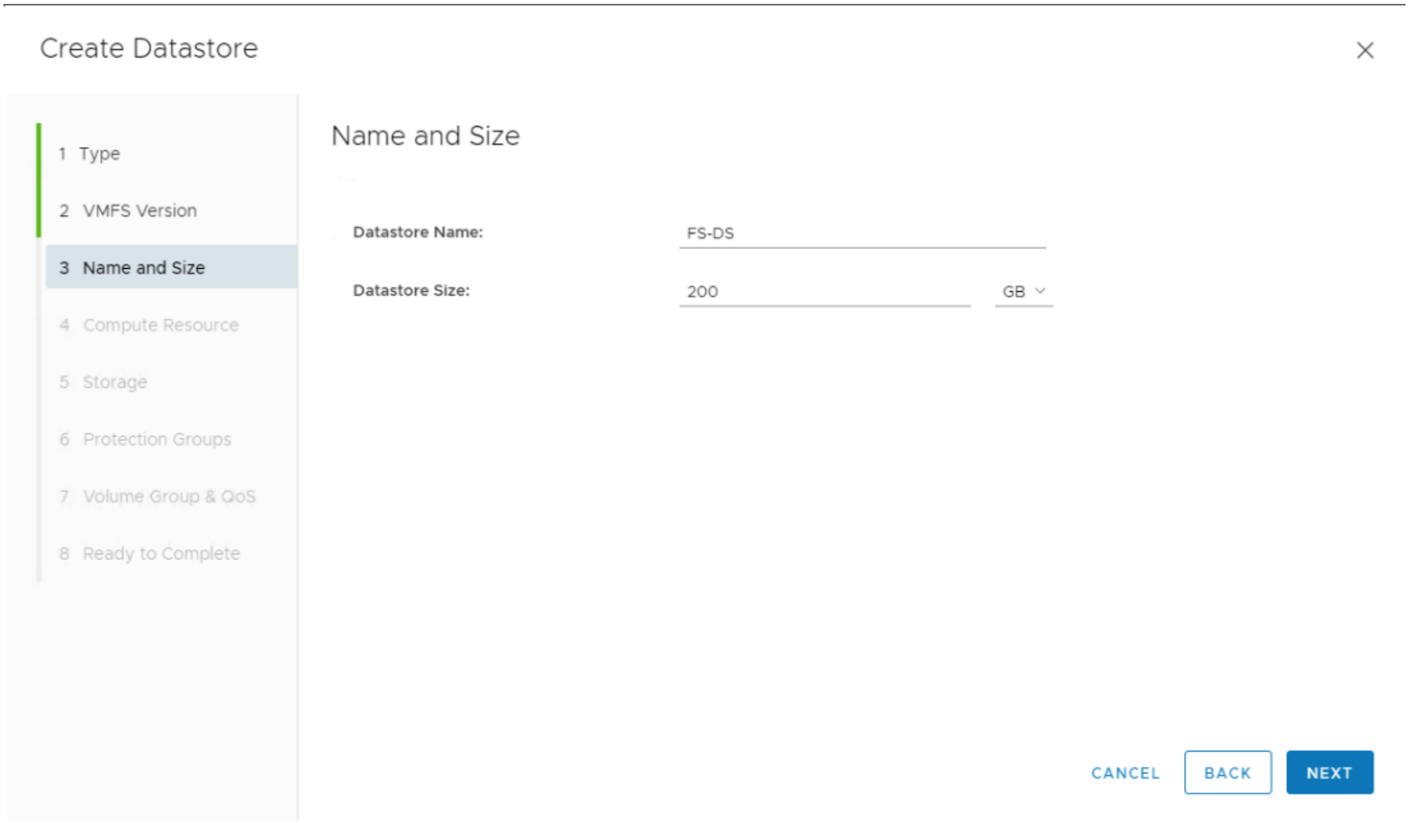

Datastore Name: FS-DS

Datastore Size: 200 GB

8. Click Next.

9. Select the Cluster under Compute Resources.

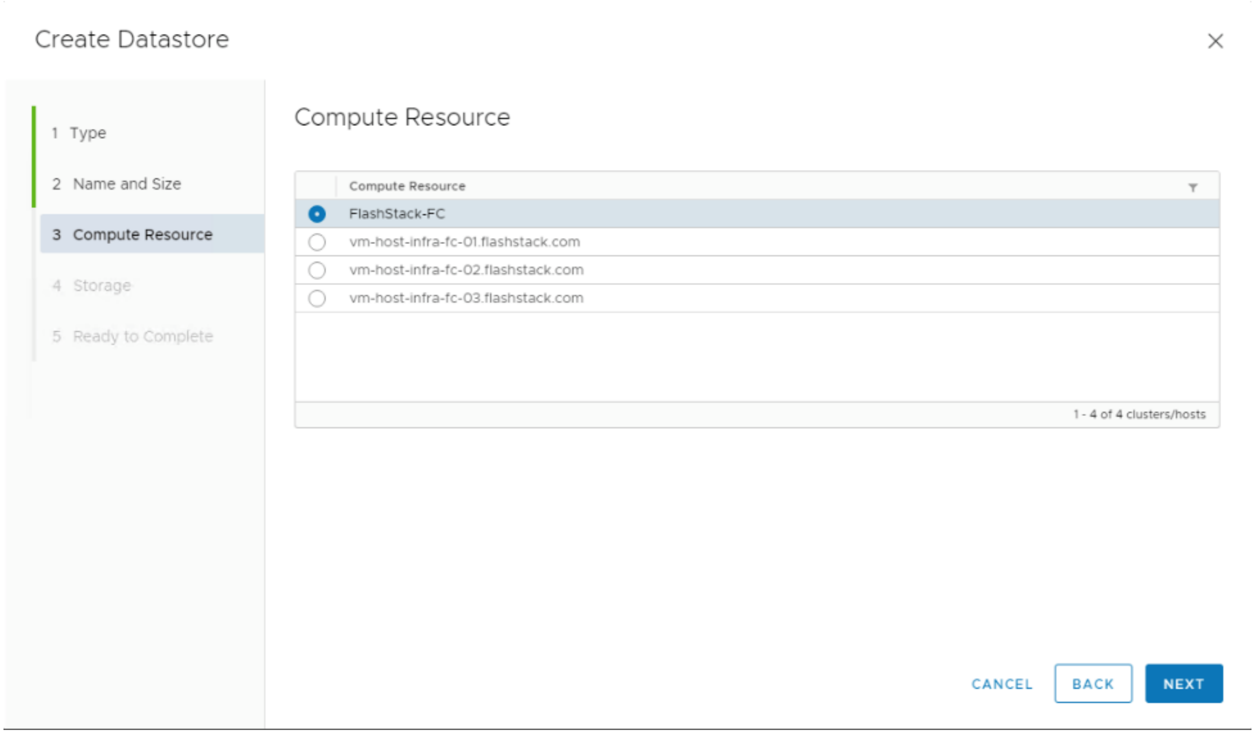

Create Datastore

Compute Resource

	Compute Resource
●	FlashStack-FC
○	vm-host-infra-fc-01.flashstack.com
○	vm-host-infra-fc-02.flashstack.com
○	vm-host-infra-fc-03.flashstack.com

1 - 4 of 4 clusters/hosts

10. Click Next.

11. Select the Registered FlashArray.

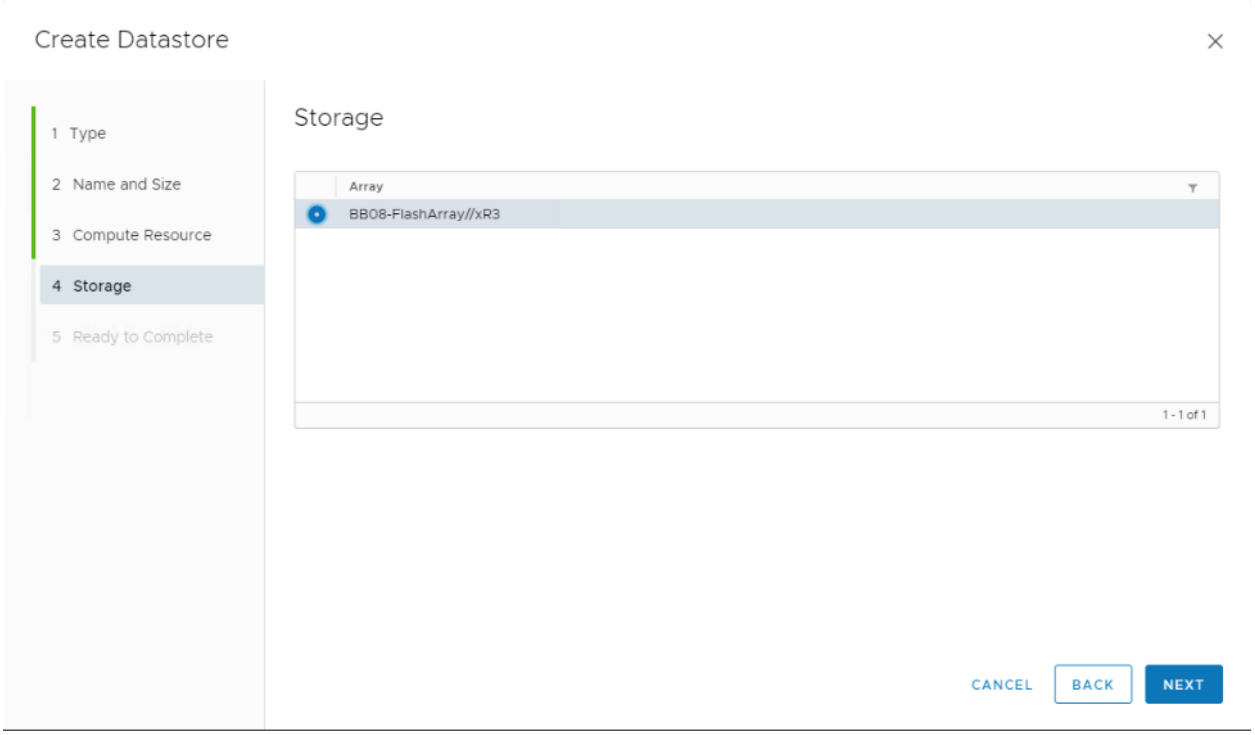

Create Datastore

Storage

	Array
●	BB08-FlashArray//xR3

1 - 1 of 1

12. Optionally, add to the protection group created earlier and click Next.

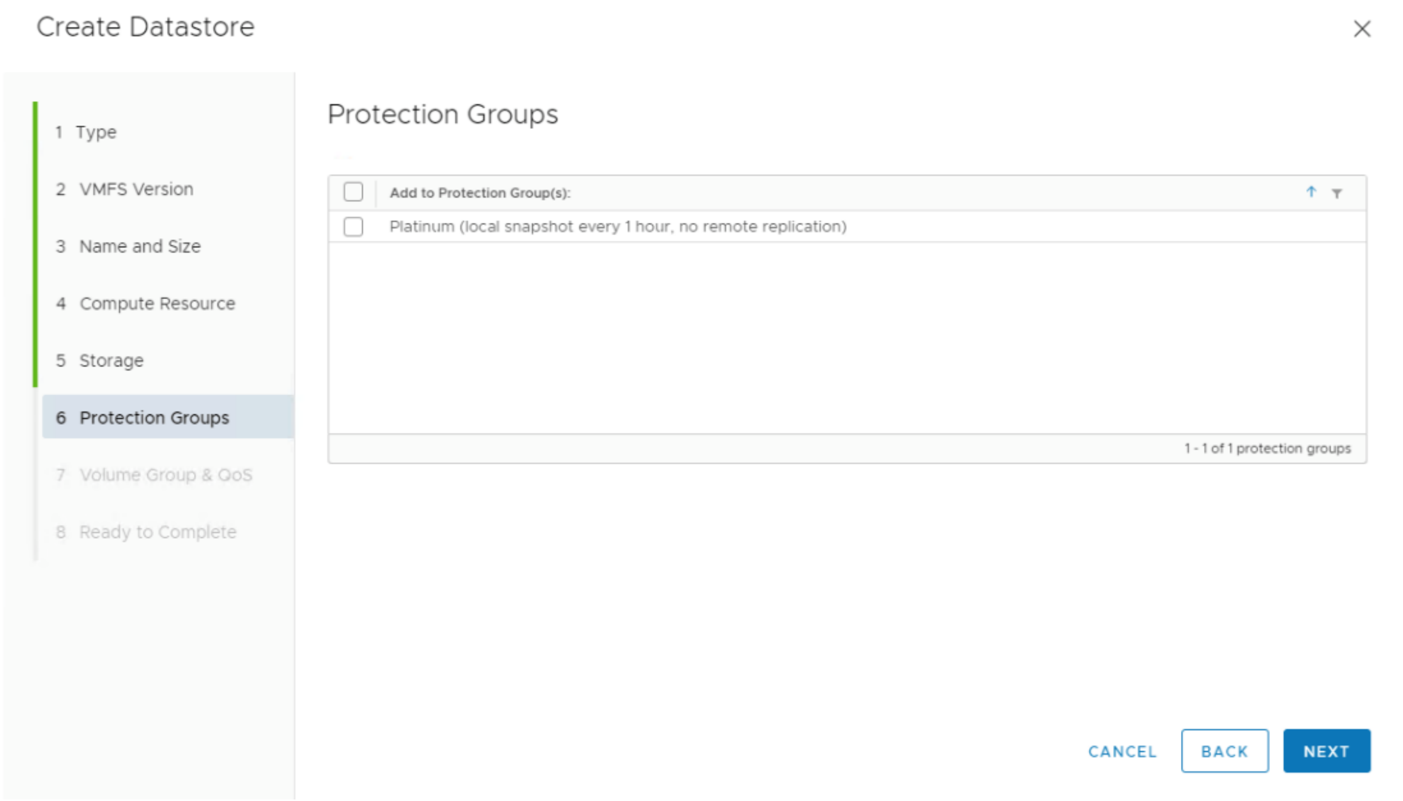

□	Add to Protection Group(s):
□	Platinum (local snapshot every 1 hour, no remote replication)

1 - 1 of 1 protection groups

13. Click Next on the Volume Group & QoS page.

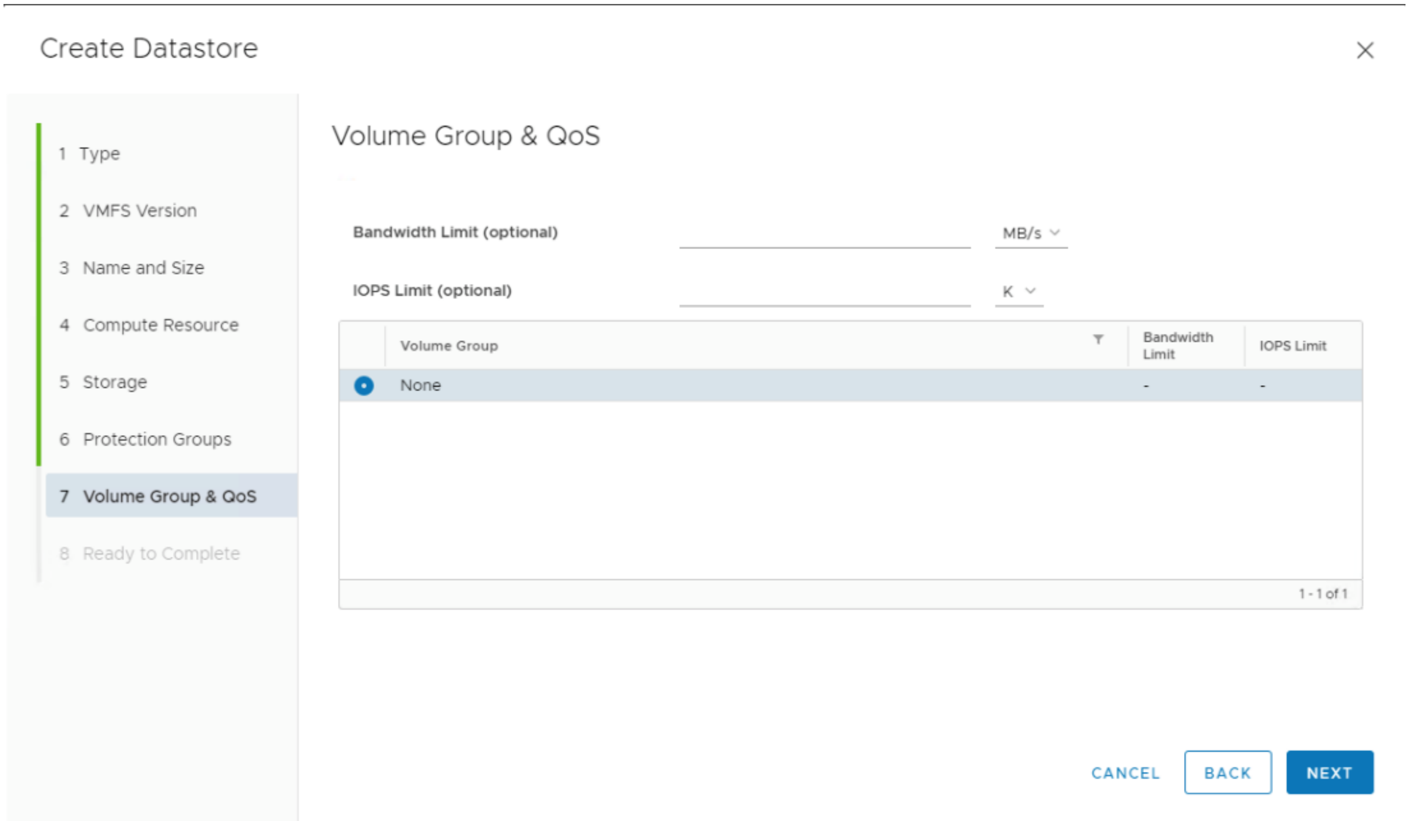

Volume Group & QoS

Bandwidth Limit (optional) MB/s

IOPS Limit (optional) K

	Volume Group	Bandwidth Limit	IOPS Limit
●	None	-	-

1 - 1 of 1

14. Review the information and click Finish.

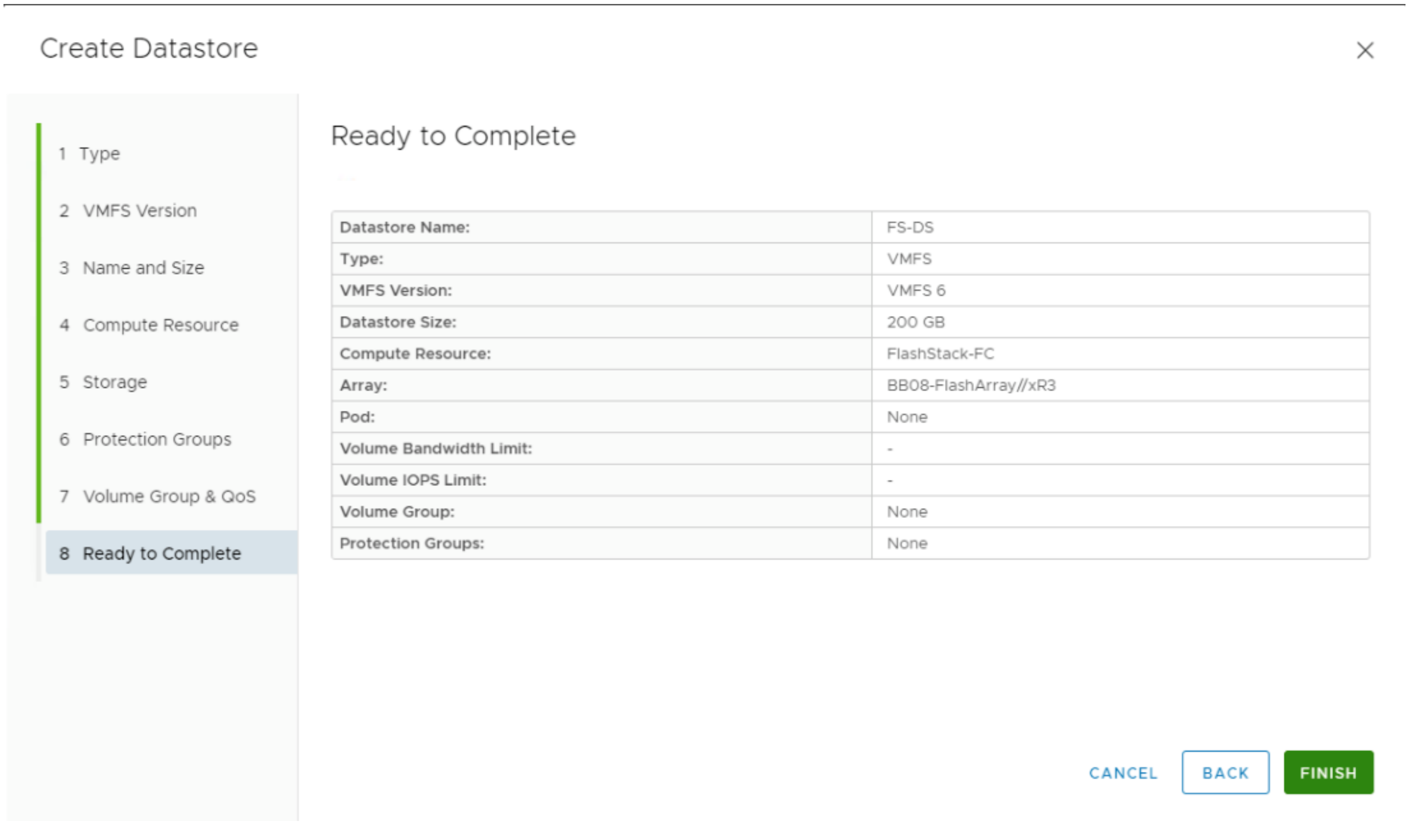

Ready to Complete

Datastore Name:	FS-DS
Type:	VMFS
VMFS Version:	VMFS 6
Datastore Size:	200 GB
Compute Resource:	FlashStack-FC
Array:	BB08-FlashArray//xR3
Pod:	None
Volume Bandwidth Limit:	-
Volume IOPS Limit:	-
Volume Group:	None
Protection Groups:	None

Create vVol Datastore

1. In vCenter, Select Host and Clusters.

2. Right-click the FlashStack Cluster and Select Pure Storage > Create Datastore.

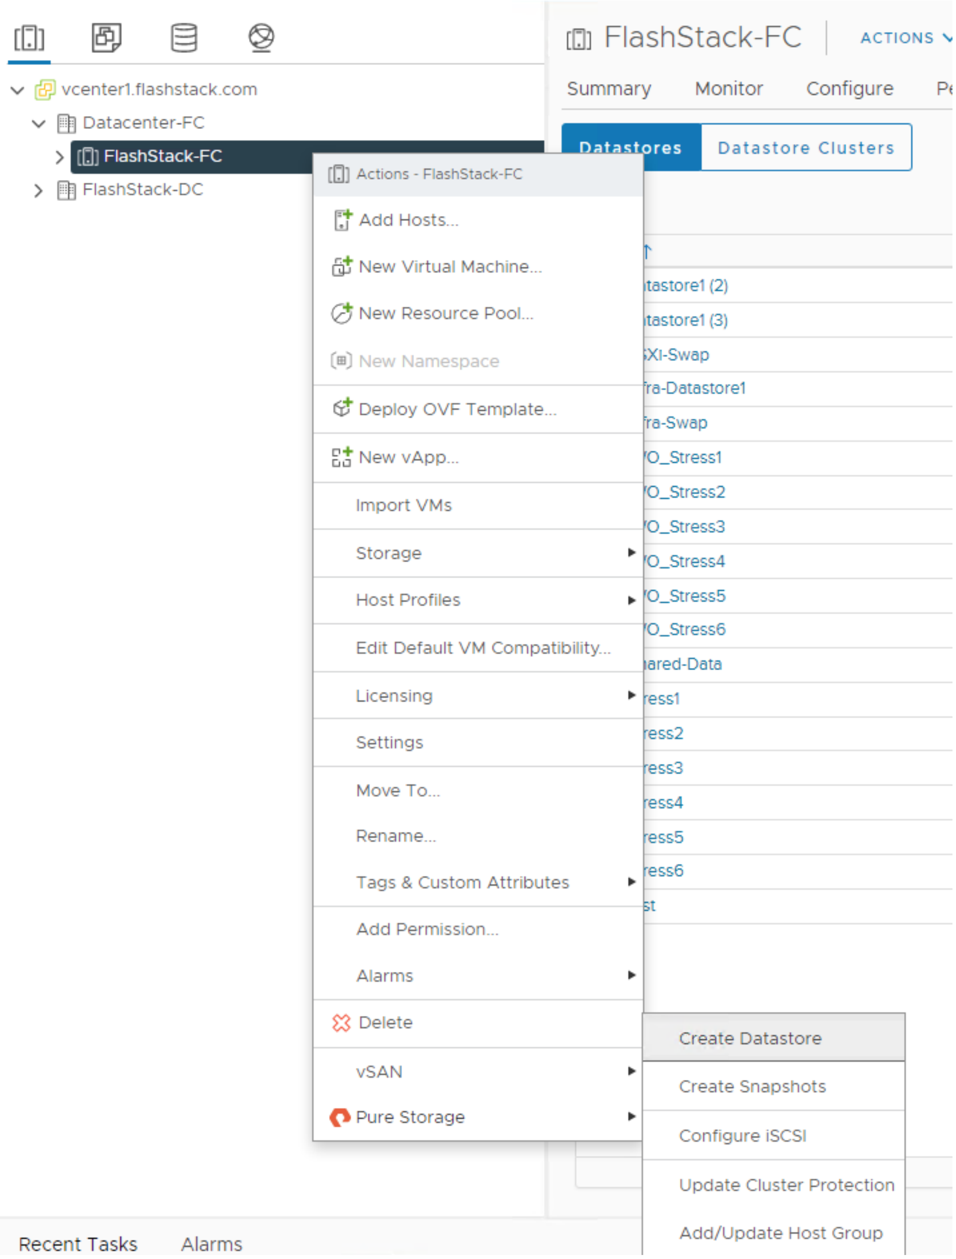

3. Click vVol.

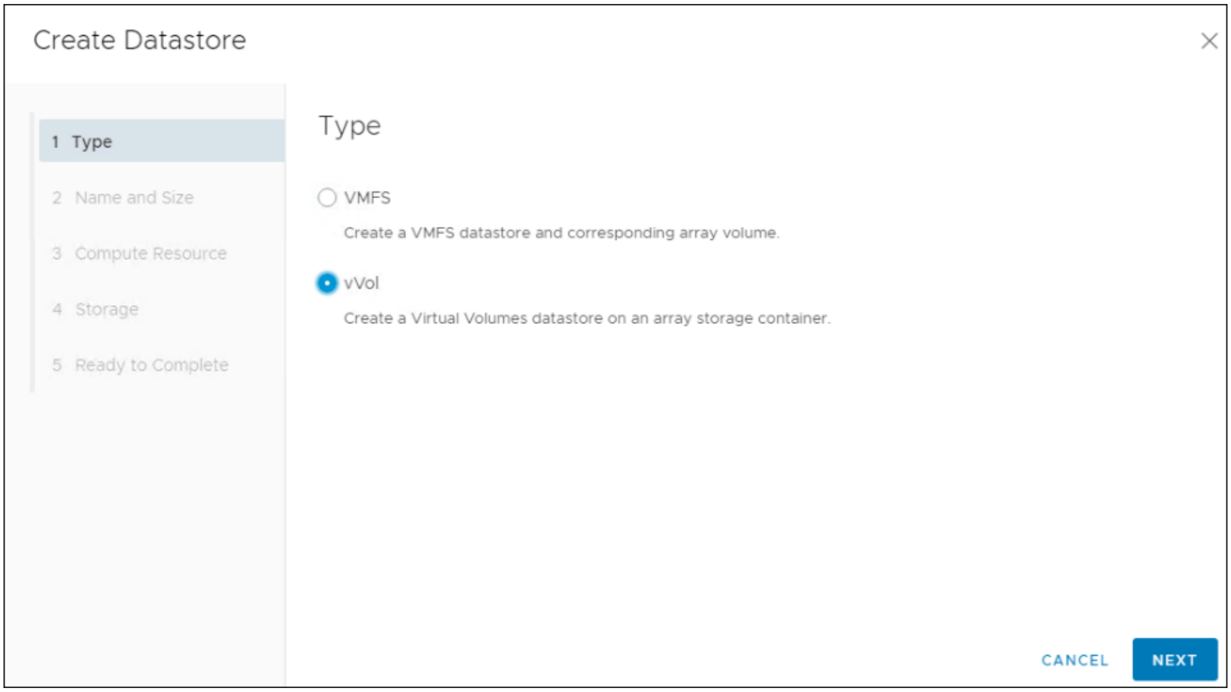

Create Datastore

Type

○ VMFS

Create a VMFS datastore and corresponding array volume.

● vVol

Create a Virtual Volumes datastore on an array storage container.

4. Click Next.

5. Enter a Datastore Name.

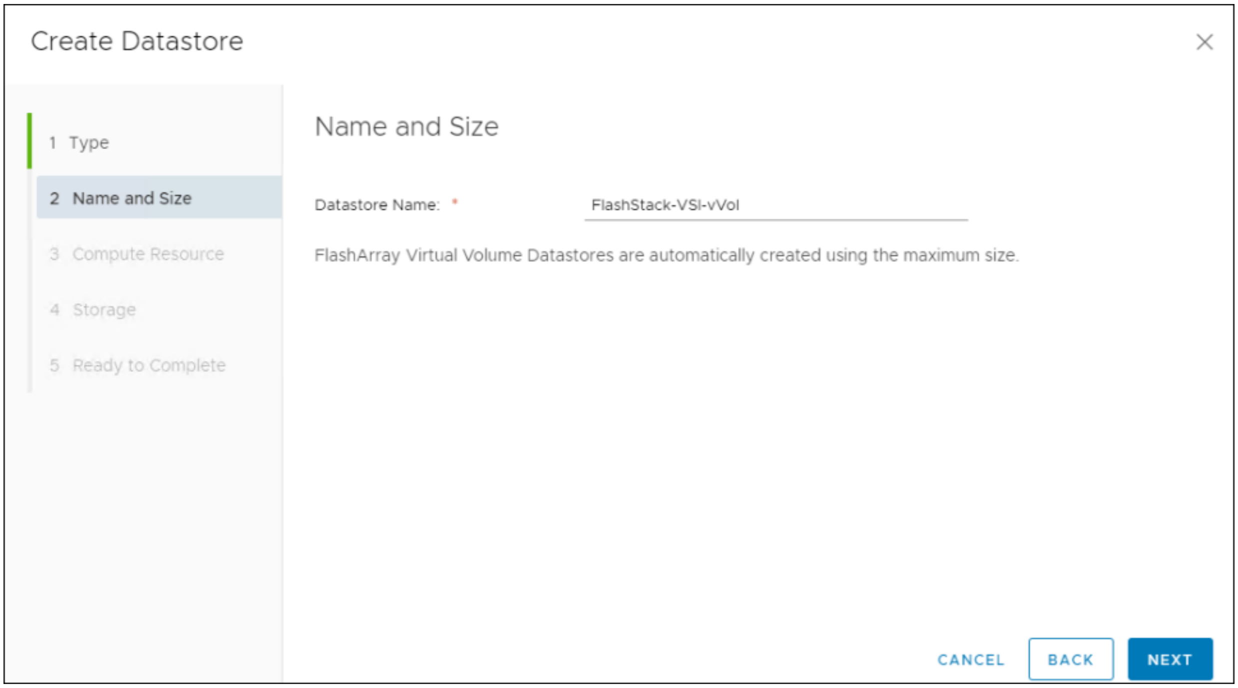

Create Datastore

Name and Size

Datastore Name: * FlashStack-VSI-vVol

FlashArray Virtual Volume Datastores are automatically created using the maximum size.

6. Click Next.

7. Click the Cluster under Compute Resources.

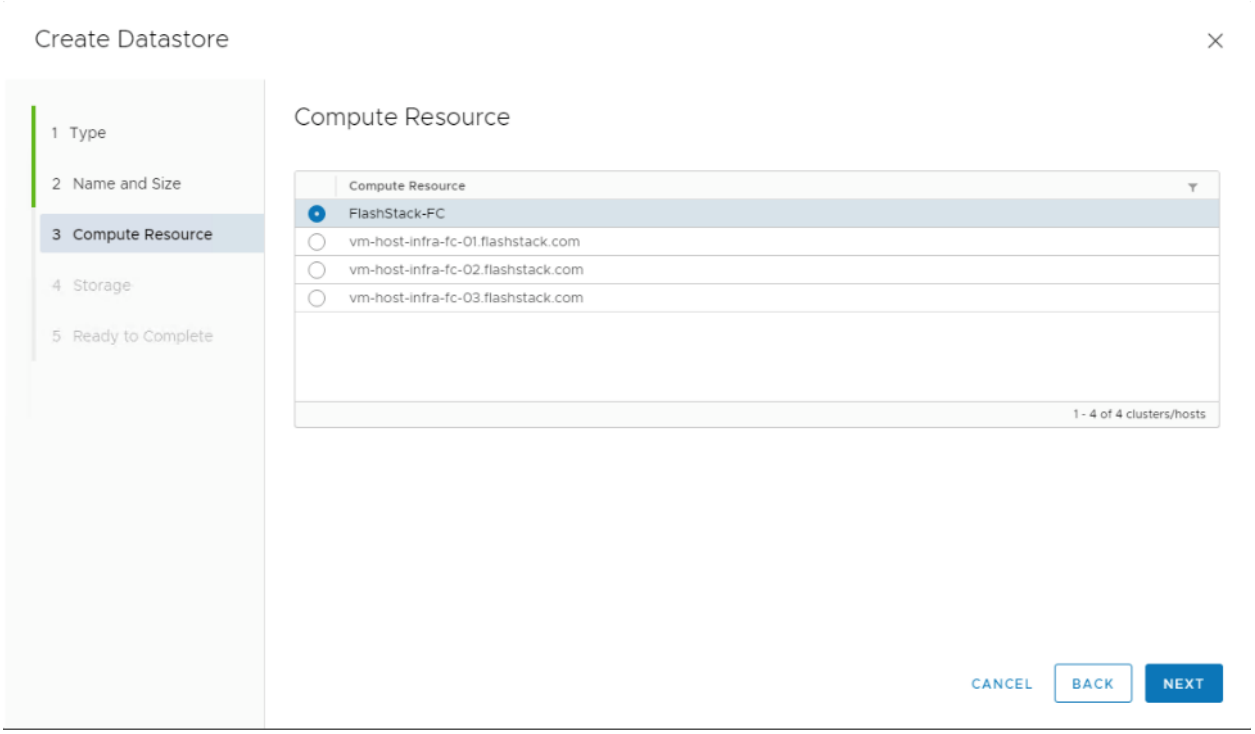

Create Datastore

1 Type
2 Name and Size
3 Compute Resource
4 Storage
5 Ready to Complete

Compute Resource

	Compute Resource
●	FlashStack-FC
○	vm-host-infra-fc-01.flashstack.com
○	vm-host-infra-fc-02.flashstack.com
○	vm-host-infra-fc-03.flashstack.com

1 - 4 of 4 clusters/hosts

8. Click Next.

9. Click the Registered FlashArray.

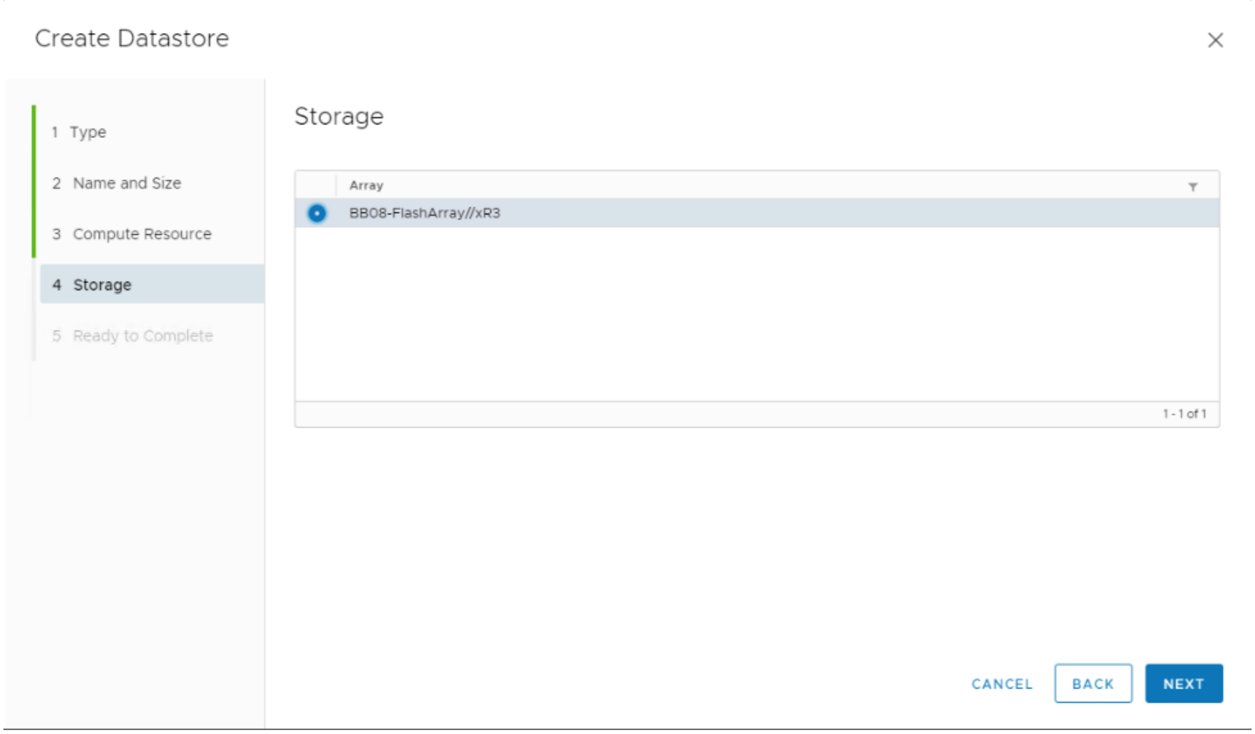

Create Datastore

1 Type
2 Name and Size
3 Compute Resource
4 Storage
5 Ready to Complete

Storage

	Array
●	BB08-FlashArray//xR3

1 - 1 of 1

10. Click Next.

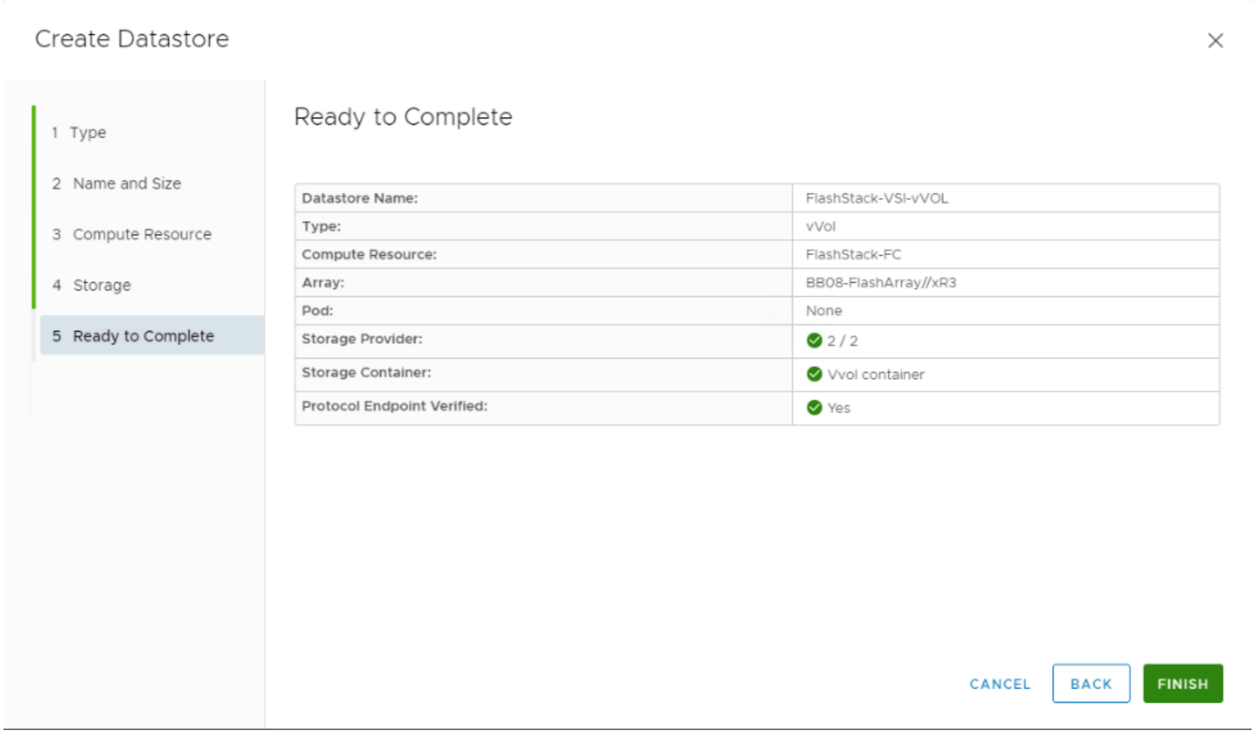

Ready to Complete

Datastore Name:	FlashStack-VSI-vVOL
Type:	vVol
Compute Resource:	FlashStack-FC
Array:	BB08-FlashArray//xR3
Pod:	None
Storage Provider:	✓ 2 / 2
Storage Container:	✓ Vvol container
Protocol Endpoint Verified:	✓ Yes

11. Review the information and click Finish.

Configure NVMe over FC on ESXi Host

To configure the NVMe over FC on the ESXi host, follow these steps:

1. Log into vCenter and on the ESXi host verify the storage adapter information, there will be four adapters listed, two among them being the FC-NVMe initiators.

2. Once you click on one, you will see more information appear in the Details panel:

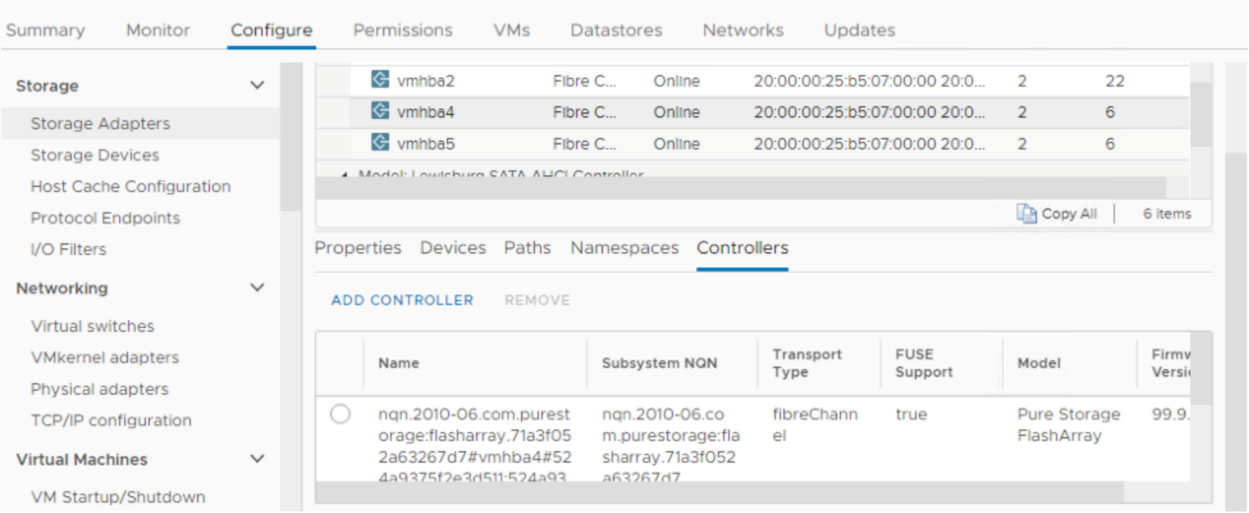

vmhba2	Fibre C...	Online	20:00:00:25:b5:07:00:00 20:0...	2	22
vmhba4	Fibre C...	Online	20:00:00:25:b5:07:00:00 20:0...	2	6
vmhba5	Fibre C...	Online	20:00:00:25:b5:07:00:00 20:0...	2	6

Model: Lewisburg SATA AHCI Controller

	Name	Subsystem NQN	Transport Type	FUSE Support	Model	Firmv Versi
○	nqn.2010-06.com.purestorage:flasharray.71a3f052a63267d7#vmhba4#524a9375f2e3d511:524a93	nqn.2010-06.com.purestorage:flasharray.71a3f052a63267d7	fibreChannel	true	Pure Storage FlashArray	99.9.

Note: If the zoning is complete at this point no additional steps are required.

3. The next step is to create the host and host group objects on the FlashArray. In NVMe-oF, initiators use something called an **N**VMe **Q**ualified **N**ame (NQN).

Note: The initiator has one and so does the target (the FlashArray). With NVMe-oF/FC, NQNs do not **replace** FC WWNs-they both exist.

Note: The WWN of each side is what is advertised on the FC layer to enable physical connectivity and zoning. The NQN is what enables the NVMe layer to communicate to the correct endpoints on the FC fabric. You can look at it in a similar way as networking in IP (MAC addresses and IPs).

4. For each ESXi host, you need to create a host object on the FlashArray, then add the NQN to it. So where do you get the NQN? However, not from the vSphere Client. For now, you need to use esxcli.

5. SSH back into the ESXi host and run:

```
esxcli nvme info get
```

6. Copy the NQN.

7. Log into the FlashArray.

Host Registration

For Host registration, follow these steps in the Pure Storage Web Portal:

1. Click Storage > Hosts.

2. Click the + icon in the Hosts Panel.

3. After clicking the Create Host (+) option, a pop-up will appear to create an individual host entry on the FlashArray.

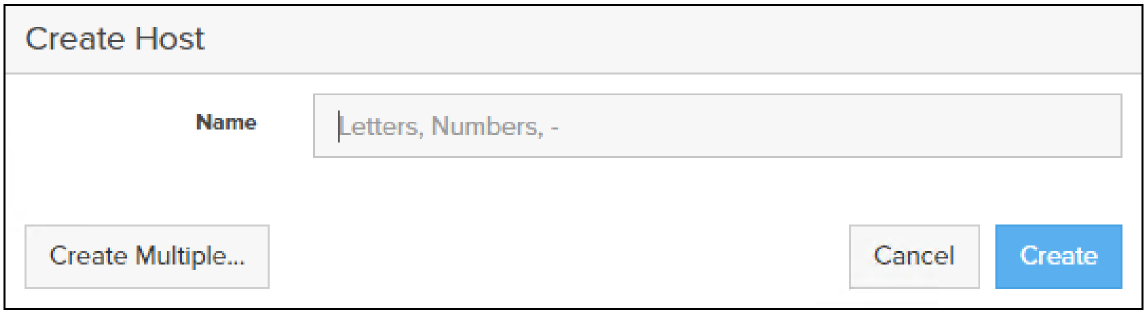
Create Host

Name: Letters, Numbers, -

4. To create more than one host entry, click the Create Multiple... option, filling in the Name, Start Number, Count, Personality as ESXi and Number of Digits, with a “#” appearing in the name where an iterating number will appear:

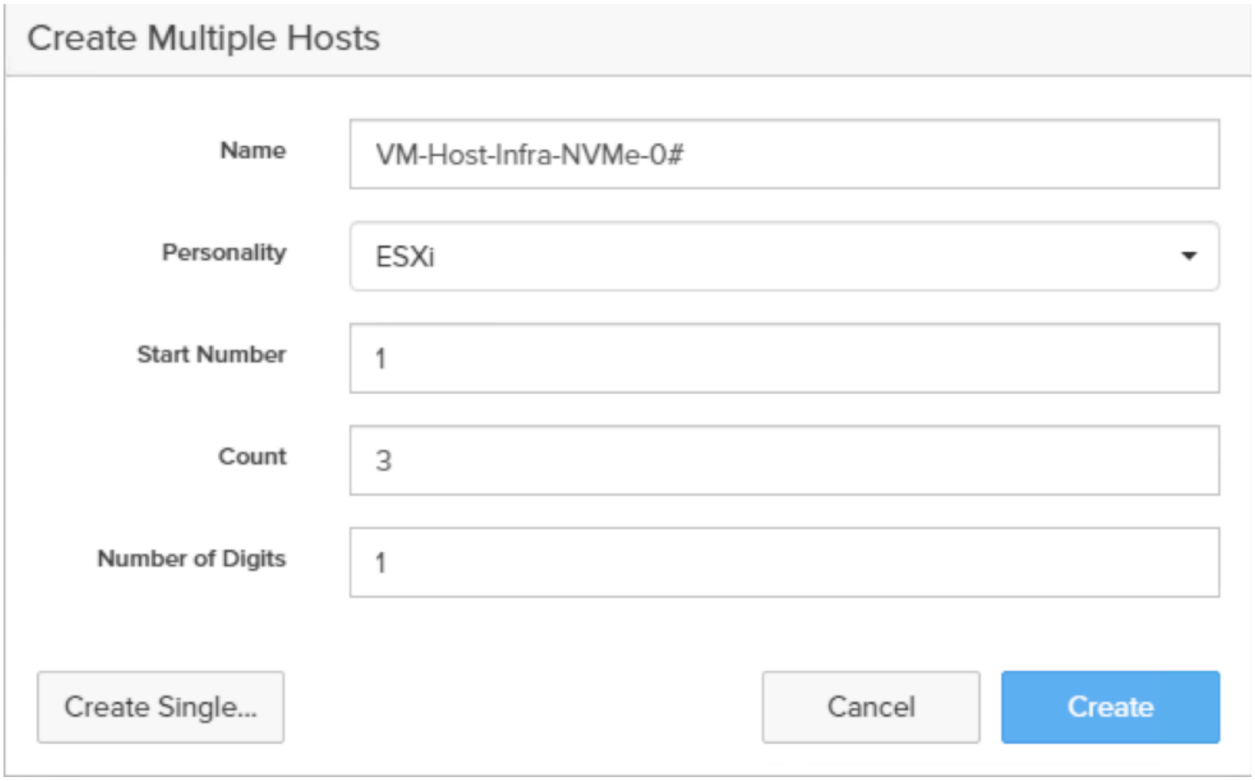

Create Multiple Hosts

Name: VM-Host-Infra-NVMe-0#

Personality: ESXi

Start Number: 1

Count: 3

Number of Digits: 1

5. Click Create to add the hosts.

6. For each host created, select the host.

7. In the Host view, select ‘Configure NQNs...’ from the Host Ports menu.

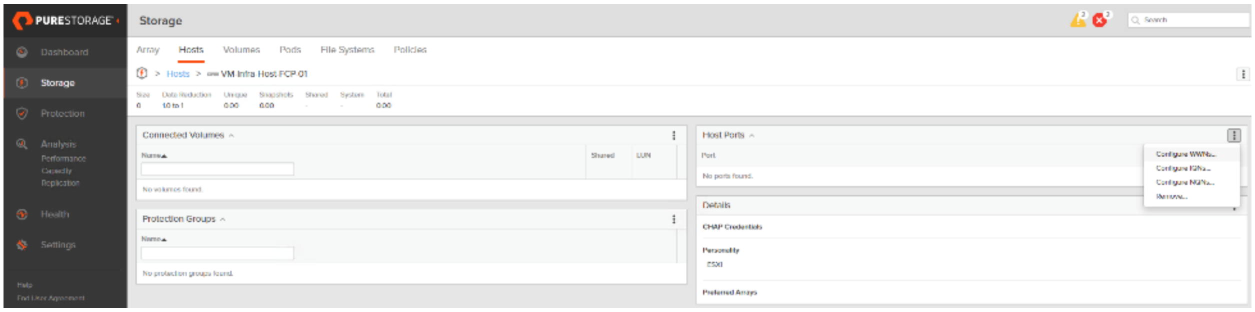

8. A pop-up will appear for Configure NVMe-oF NQNs for <Host> Within this pop-up, enter the appropriate NQN of this specific host.

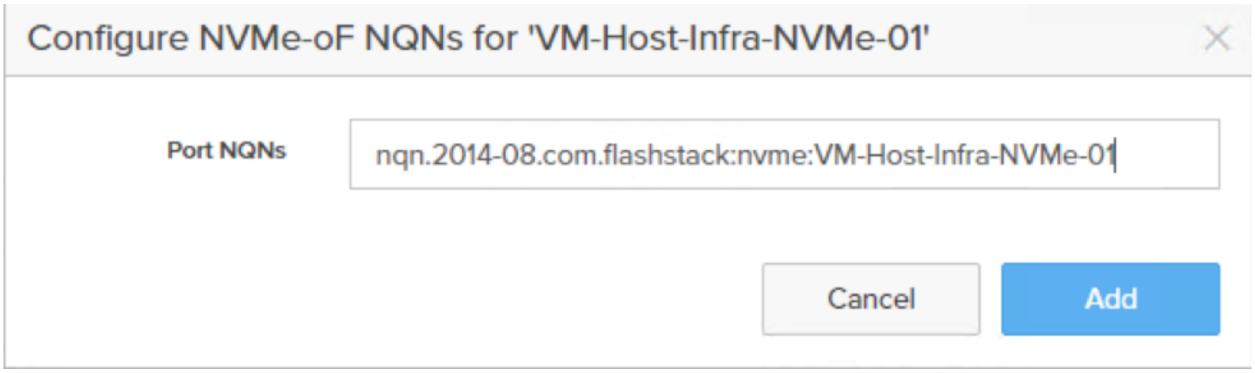

Configure NVMe-oF NQNs for 'VM-Host-Infra-NVMe-01'

Port NQNs: nqn.2014-08.com.flashstack:nvme:VM-Host-Infra-NVMe-01

9. Click Add.

10. Repeat steps 1-9 for each host created.

Create NVMe Host Group

Host Groups allow the Administrator to map Volumes to a group of hosts at once with the same LUN ID. To create a Host Group, follow these steps in the Pure Storage Web Portal:

1. Click Storage > Hosts.

2. Click the + icon in the Host Groups Panel.

3. A pop-up will appear to create a host group on the FlashArray.

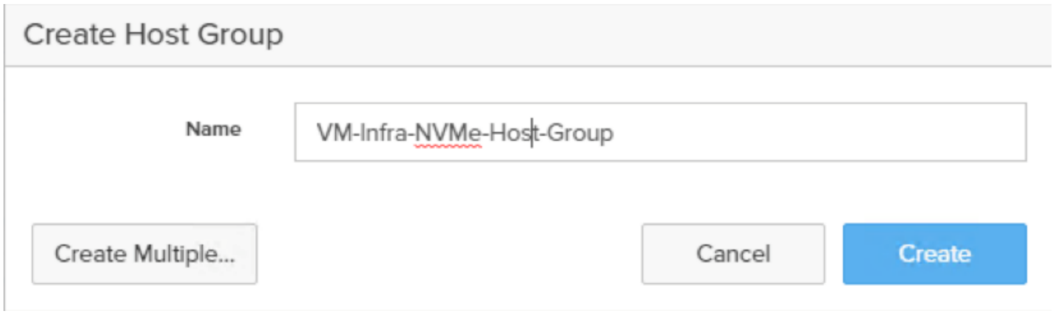

Create Host Group

Name: VM-Infra-NVMe-Host-Group

4. Provide a name for the group and click Create.

5. Select the group in the Host Groups Panel.

6. In the Host Group view, select 'Add...' from the Member Hosts menu.

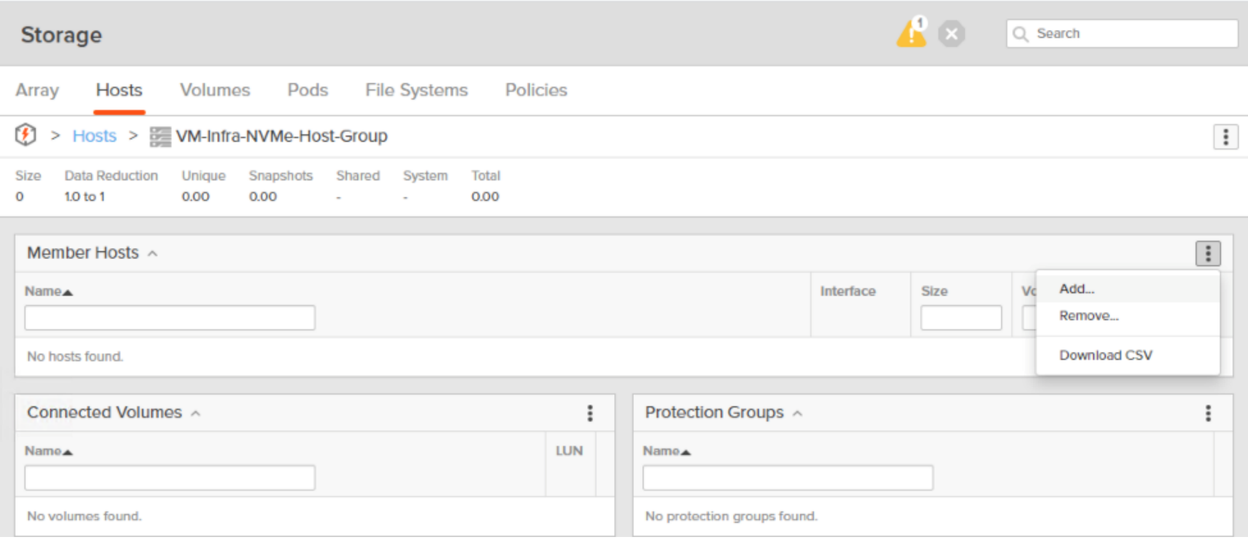

Storage

Hosts > VM-Infra-NVMe-Host-Group

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

Member Hosts: No hosts found.

Connected Volumes: No volumes found.

Protection Groups: No protection groups found.

7. Select the host to be part of the host group.

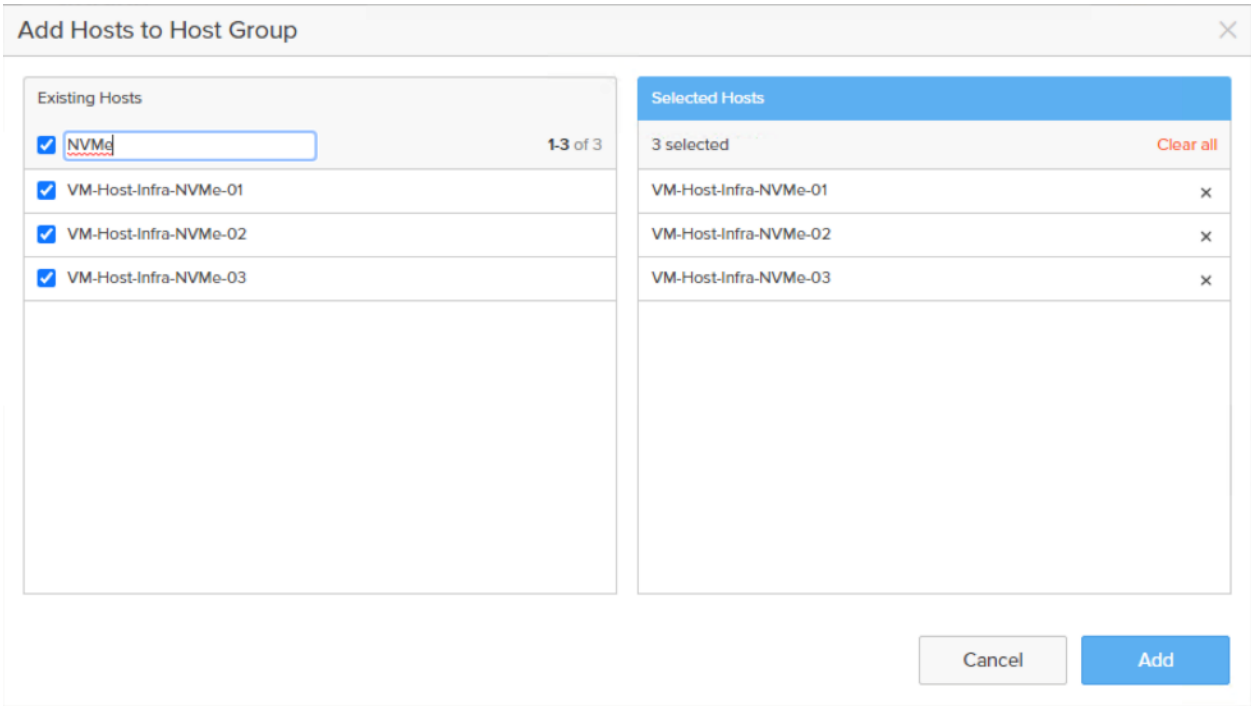

8. Click Add.

Create NVMe datastores

To create datastore volumes for the ESXi Cluster, follow these steps in the Pure Storage Web Portal:

1. Click Storage > Volumes.
2. Click the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

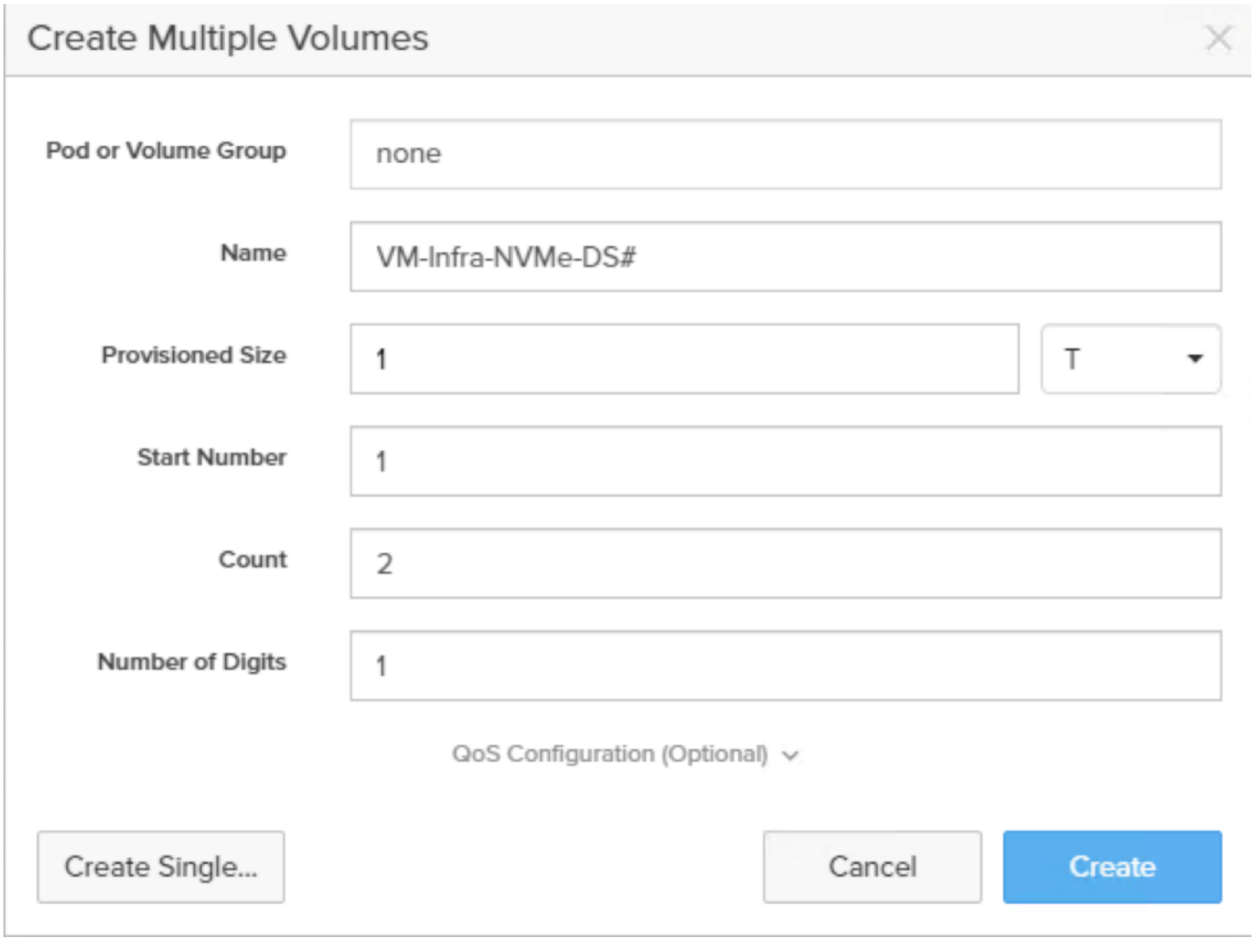

Create Multiple Volumes

Pod or Volume Group: none

Name: VM-Infra-NVMe-DS#

Provisioned Size: 1 T

Start Number: 1

Count: 2

Number of Digits: 1

4. Fill in the Name and Provisioned Size.

5. Click Create to provision the volumes to be used as Infra datastore LUN.

6. Go back to the Hosts section under the Storage tab. Click ESXi cluster NVMe host group created earlier and select the gear icon drop-down within the Connected Volumes tab within that host group.

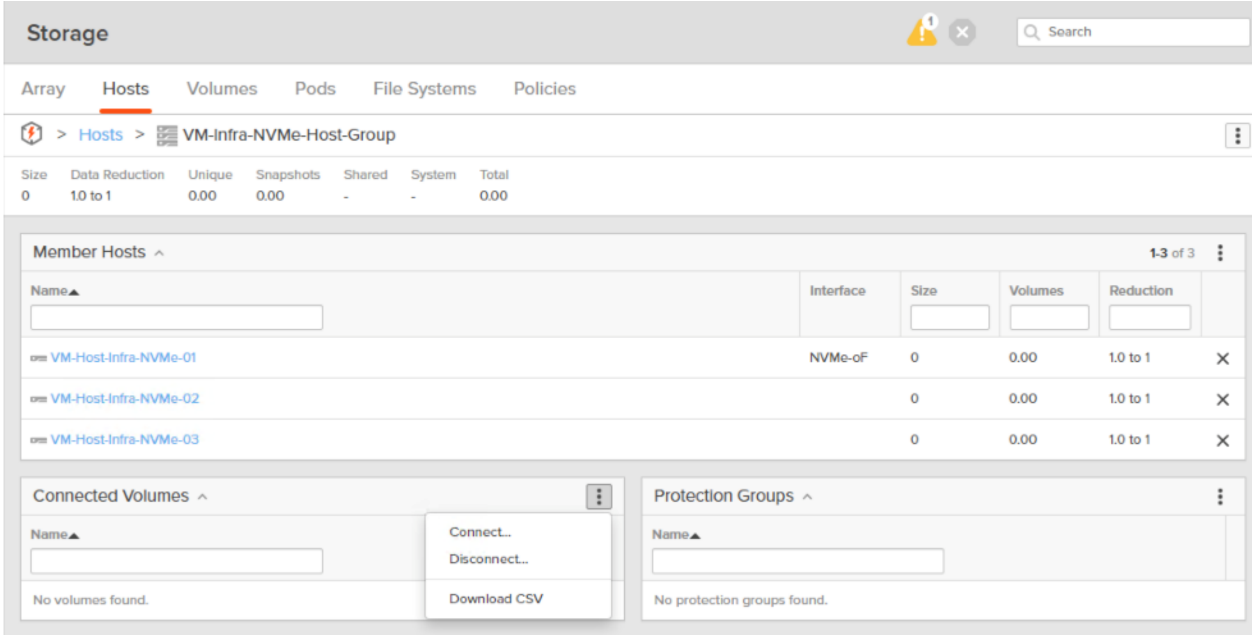
Storage

Array | Hosts | Volumes | Pods | File Systems | Policies

> Hosts > VM-Infra-NVMe-Host-Group

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

Member Hosts 1-3 of 3

Name	Interface	Size	Volumes	Reduction
VM-Host-Infra-NVMe-01	NVMe-oF	0	0.00	1.0 to 1
VM-Host-Infra-NVMe-02		0	0.00	1.0 to 1
VM-Host-Infra-NVMe-03		0	0.00	1.0 to 1

Connected Volumes

Name

No volumes found.

Protection Groups

Name

No protection groups found.

7. Within the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

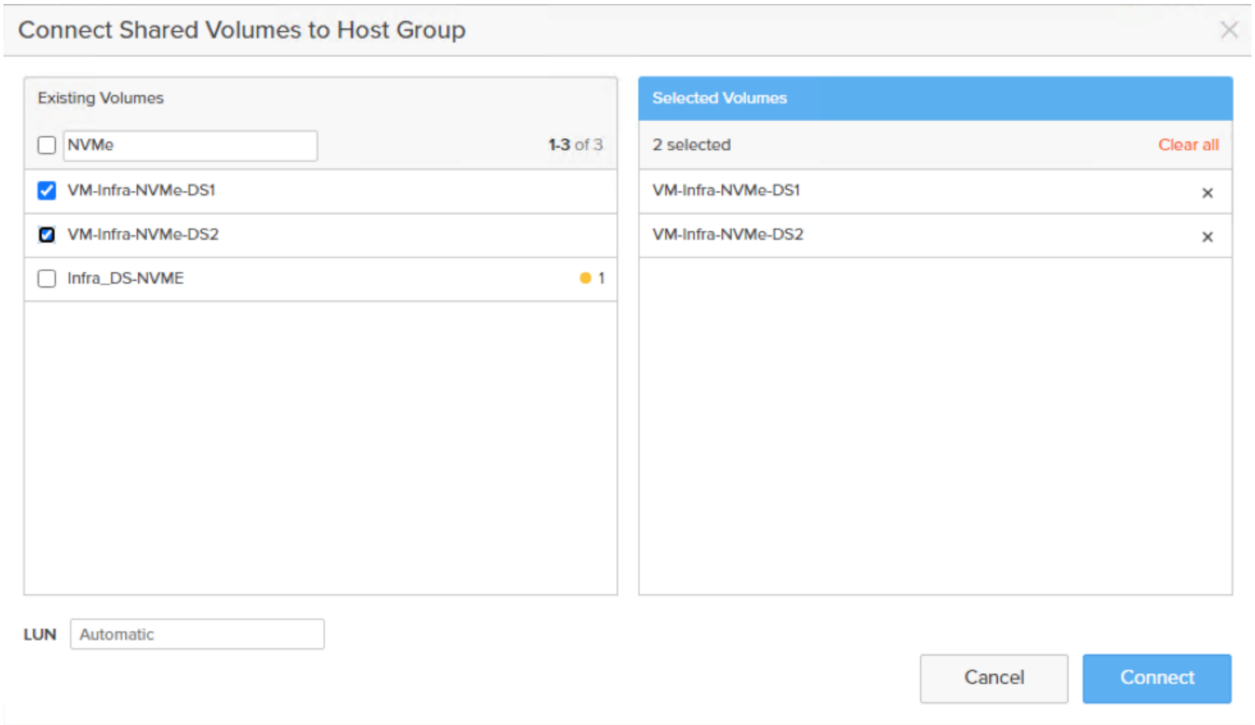
Connect Shared Volumes to Host Group

Existing Volumes

□ NVMe 1-3 of 3

☑ VM-Infra-NVMe-DS1

☑ VM-Infra-NVMe-DS2

□ Infra_DS-NVME ● 1

Selected Volumes

2 selected Clear all

VM-Infra-NVMe-DS1

VM-Infra-NVMe-DS2

LUN Automatic

8. Select the Infra datastore NVMe volumes that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.

ESXi Host NVMe over FC Datastore Configuration

To configure the ESXi host NVMe over FC datastore, follow these steps:

1. The remaining steps in the VMware vSphere Client are manual steps that should be completed whether Ansible configuration or manual configuration is being done. Verify that the NVMe Fibre Channel Disk is mounted on each ESXi host. Under Hosts and Clusters select the ESXi host. In the center pane, select Configure > Storage > Storage Devices. The NVMe Fibre Channel Disk should be listed under Storage Devices. Select the NVMe Fibre Channel Disk, then select Paths underneath. Verify 4 paths have a status of Active (I/O). Repeat this for all 3 hosts.

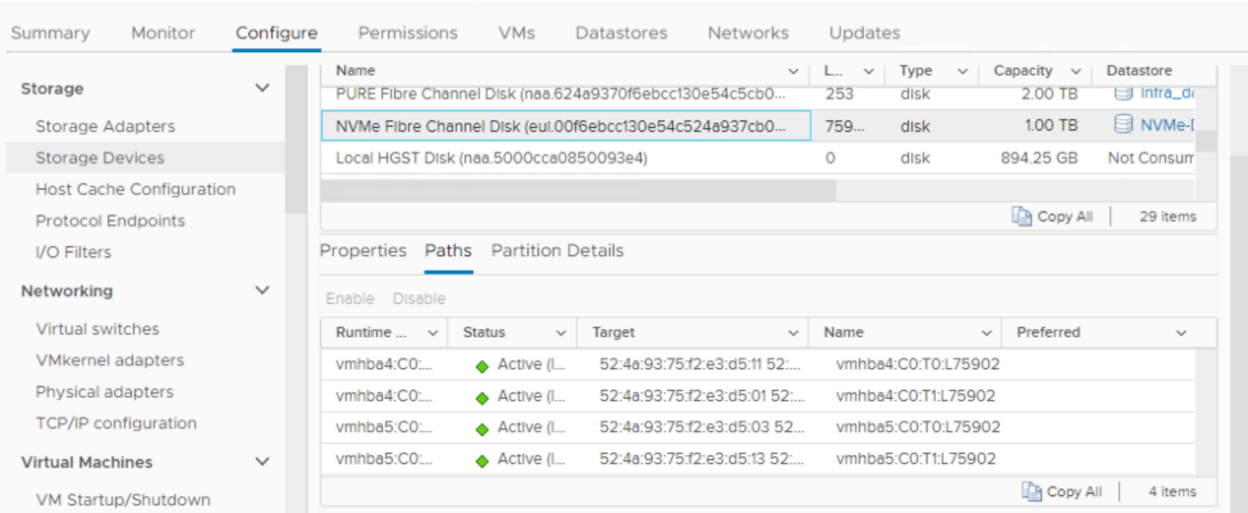

Name	L...	Type	Capacity	Datastore
PURE Fibre Channel Disk (naa.624a9370f6ebcc130e54c5cb0...	253	disk	2.00 TB	Infra_d
NVMe Fibre Channel Disk (eui.00f6ebcc130e54c524a937cb0...	759...	disk	1.00 TB	NVMe-I
Local HGST Disk (naa.5000cca0850093e4)	0	disk	894.25 GB	Not Consum

Copy All | 29 items

Runtime ...	Status	Target	Name	Preferred
vmhba4:C0:...	Active (I...	52:4a:93:75:f2:e3:d5:11 52:...	vmhba4:C0:T0:L75902	
vmhba4:C0:...	Active (I...	52:4a:93:75:f2:e3:d5:01 52:...	vmhba4:C0:T1:L75902	
vmhba5:C0:...	Active (I...	52:4a:93:75:f2:e3:d5:03 52...	vmhba5:C0:T0:L75902	
vmhba5:C0:...	Active (I...	52:4a:93:75:f2:e3:d5:13 52:...	vmhba5:C0:T1:L75902	

Copy All | 4 items

2. For any of the three hosts, right-click the host under Hosts and Clusters and select Storage > New Datastore. Leave VMFS selected and click NEXT.

3. Name the datastore and select the NVMe Fibre Channel Disk. Click NEXT.

4. Leave VMFS 6 selected and click NEXT.

5. Leave all Partition configuration values at the default values and click NEXT.

6. Review the information and click FINISH.

7. Click Storage and select the just-created NVMe datastore. In the center pane, select Hosts. Ensure all three hosts have the datastore mounted.

ESXi Host Multipathing Configuration

To configure the ESXi host multipathing, follow these steps:

1. From the vCenter management GUI.

2. Go to Hosts and Clusters view.

3. Click a Host.

4. Click the Configure tab.

5. Click Storage Devices.

6. Click an NVMe device.

7. Click Edit Multipathing.

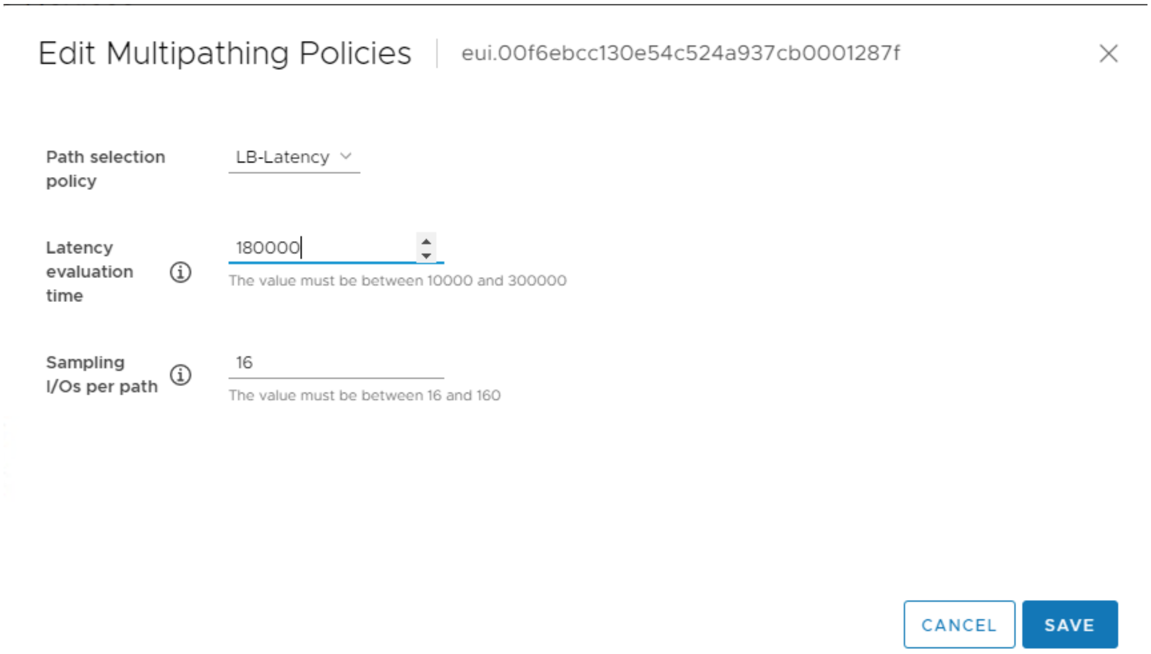

Edit Multipathing Policies | eui.00f6ebcc130e54c524a937cb0001287f

Path selection policy: LB-Latency

Latency evaluation time: 180000

The value must be between 10000 and 300000

Sampling I/Os per path: 16

The value must be between 16 and 160

Appendix

FlashStack iSCSI Addition

Cisco Nexus Switch Configuration

This section is a delta section for adding infrastructure iSCSI to the Cisco Nexus switches. This section should be executed after the Cisco Nexus Switch Configuration section in the main document is completed.

Create Infrastructure iSCSI VLANs on Cisco Nexus A and Cisco Nexus B

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

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

```
config t
vlan <infra-iscsi-a-vlan-id>
name Infra-iSCSI-A-VLAN
vlan <infra-iscsi-b-vlan-id>
name Infra-iSCSI-B-VLAN
exit
```

Add iSCSI Individual Port Descriptions for Troubleshooting and Enable UDLD for Pure iSCSI Interfaces

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A connected to Cisco Pure FlashArray//X R3, follow this step, follow this step:

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

```
config t
interface Ethernet1/37
description <<var_flasharray_hostname>>-CT0.ETH4
interface Ethernet1/38
description <<var_flasharray_hostname>>-CT1.ETH4
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B connected to Cisco Pure FlashArray//X R3, follow this step:

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

```
config t
interface Ethernet1/37
description <<var_flasharray_hostname>>-CT0.ETH5
interface Ethernet1/38
description <<var_flasharray_hostname>>-CT1.ETH5
```

Configure iSCSI interfaces for Cisco Nexus 93180YC-FX-A

To configure iSCSI interfaces for this deployment, run the following commands on Cisco Nexus 93180YC-FX-A:

```
config t
interface Ethernet1/37
switchport
switchport access valn <<var-iscsi-a-vlan-id>>
mtu 9216
no negoriate auto
no shut
interface Ethernet1/38
switchport
switchport access valn <<var-iscsi-a-vlan-id>>
mtu 9216
no negoriate auto
no shut
```

Configure iSCSI interfaces for Cisco Nexus 93180YC-FX-B

To configure iSCSI interfaces for this deployment, run the following commands on Cisco Nexus 93180YC-FX-B:

```
config t
interface Ethernet1/37
switchport
switchport access valn <<var-iscsi-b-vlan-id>>
mtu 9216
no negoriate auto
no shut
interface Ethernet1/38
switchport
switchport access valn <<var-iscsi-b-vlan-id>>
mtu 9216
no negoriate auto
no shut
```

Add Infrastructure iSCSI VLANs to Port-Channels on Cisco Nexus A and Cisco Nexus B

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

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

```
interface Po10
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
interface Po121
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
interface Po123
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
copy run start
```

FlashArray //X R3 iSCSI Interface Configuration

The iSCSI traffic will be carried on two VLANs, A (901) and B (902) that are configured in our example with the following values:

Table 22.iSCSI A FlashArray//X50 R3 Interface Configuration Settings

FlashArray Controller	iSCSI Port	IP Address	Subnet Mask
FlashArray//X R3 Controller 0	CT0.ETH4	192.168.101.146	255.255.255.0
FlashArray//X R3 Controller 1	CT0.ETH4	192.168.101.147	255.255.255.0

Table 23.iSCSI B FlashArray//X50 R3 Interface Configuration Settings

FlashArray Controller	iSCSI Port	IP Address	Subnet Mask
FlashArray//X R3 Controller 0	CT0.ETH5	192.168.102.146	255.255.255.0
FlashArray//X R3 Controller 1	CT0.ETH5	192.168.102.147	255.255.255.0

To configure iSCSI interfaces for environments deploying iSCSI boot LUNs and/or datastores, follow these steps from Pure FlashArray Web Portal:

1. Click Settings > Network
2. Click Edit for interface CT0.eth4.
3. Click Enable and add the IP information from the above tables and set the MTU to 9000.

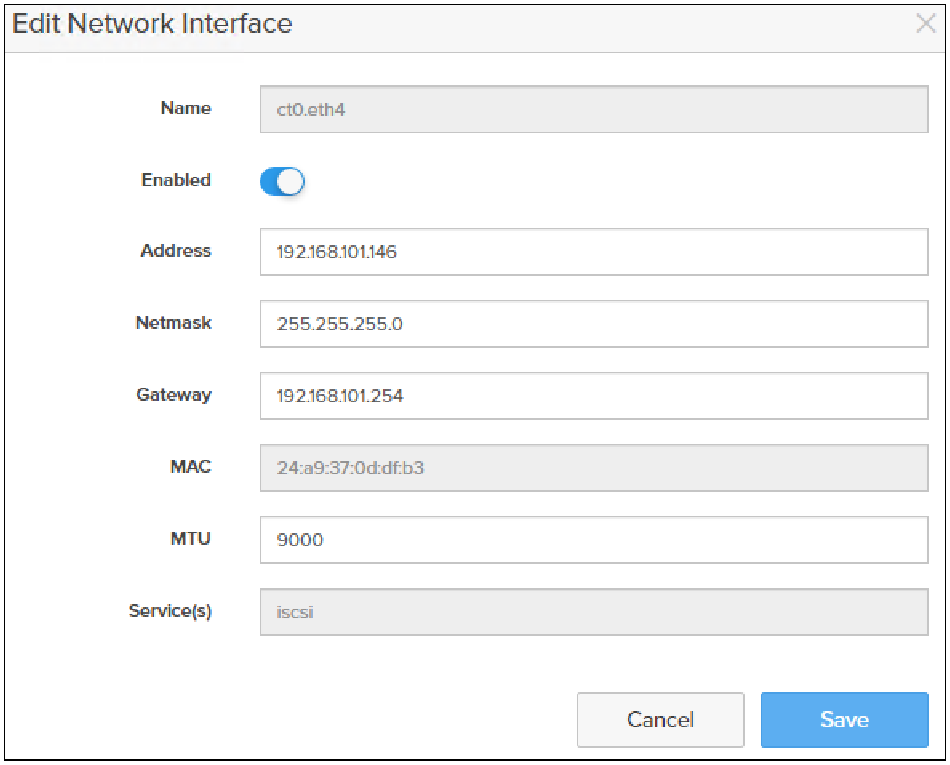

Edit Network Interface

Name	ct0.eth4
Enabled	●
Address	192.168.101.146
Netmask	255.255.255.0
Gateway	192.168.101.254
MAC	24:a9:37:0d:df:b3
MTU	9000
Service(s)	iscsi

4. Click Save.
5. Repeat steps 1-4 for CT0.eth5, CT1.eth4, and CT1.eth5.

Cisco UCS iSCSI Configuration

The following subsections can be completed to add infrastructure iSCSI to the Cisco UCS. In this procedure, one service profile template for Infrastructure ESXi hosts within the FSV Organization is created for Fabric A boot. This procedure needs to be completed after creating and deploying the domain profile to the fabric interconnects following the procedure described in the section Configure a Cisco UCS Domain Profile.

Create Server Profile Template

1. Log into the Cisco Intersight portal.

2. Go to Configure -> Templates and from the main window, select Create UCS Server Profile Template.

Step 1 - General

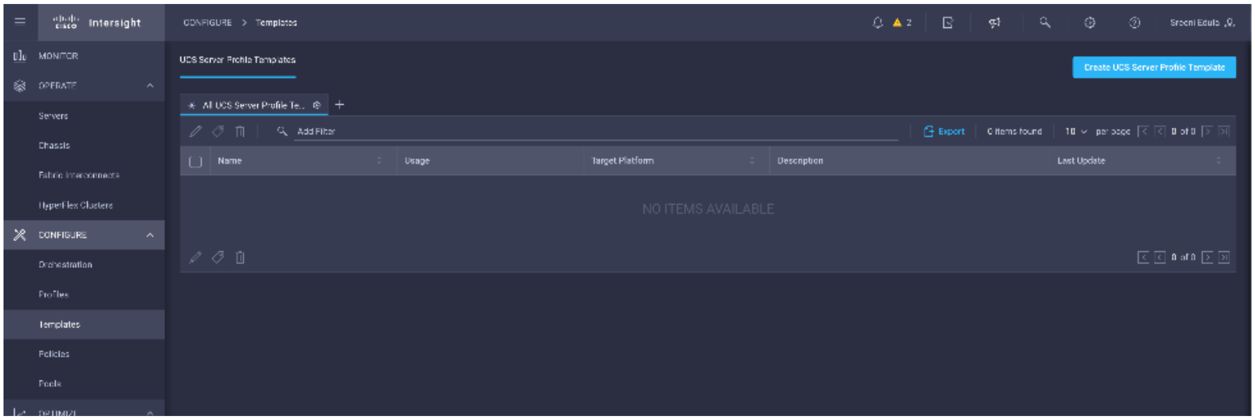

Name	Usage	Target Platform	Description	Last Update

NO ITEMS AVAILABLE

3. Select the Organization from the drop-down list (for example, FSV).

4. Provide a name for the Server Profile Template (for example, VM-Host-Infra-iSCSI).

5. Click UCS Server (FI-Attached).

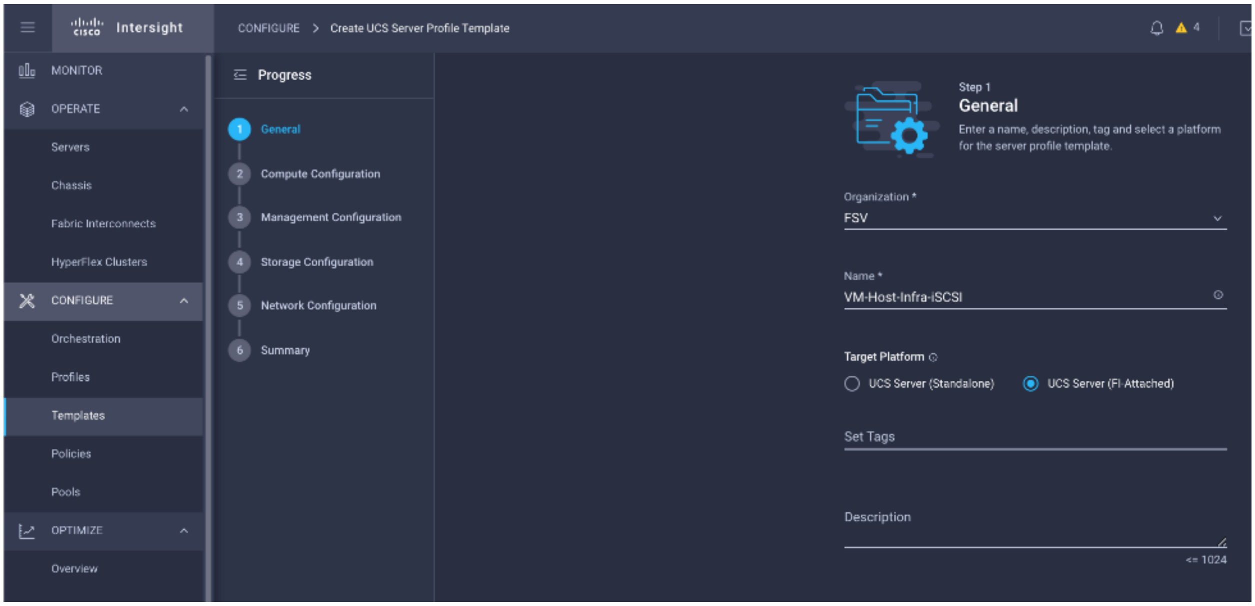

Step 1

General

Enter a name, description, tag and select a platform for the server profile template.

Organization *: FSV

Name *: VM-Host-Infra-iSCSI

Target Platform: ○ UCS Server (Standalone) ● UCS Server (FI-Attached)

Set Tags

Description

<= 1024

6. Click Next.

Step 2 - Compute Configuration

Configure UUID Pool

1. Click Select Pool next to UUID Pool and in the pane on the right, click Create New.

2. Provide a name for the pool (for example, AA19-UUID-Pool).

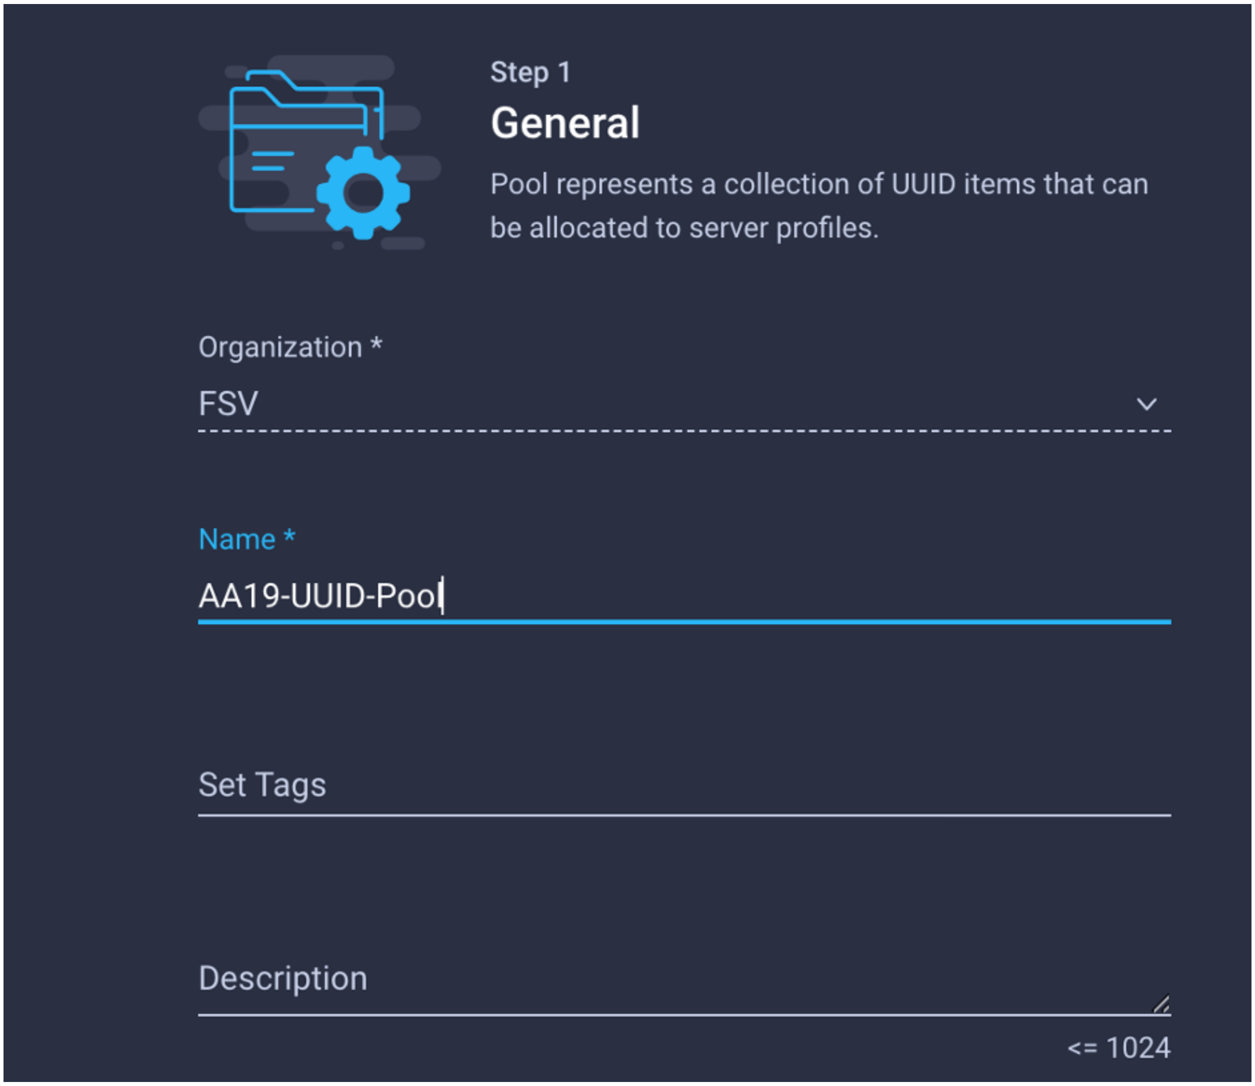

Step 1

General

Pool represents a collection of UUID items that can be allocated to server profiles.

Organization *

FSV

Name *

AA19-UUID-Pool

Set Tags

Description

<= 1024

3. Click Next. The Pool Details page appears.

4. In the Configuration section, add the UUID Prefix number in a hexadecimal format. Example, 1728E89A-7B43-47DE.

5. In the UUID Blocks section, add the following configuration details:

    a. From—Indicates the UUID suffix of the block in a hexadecimal format. Example, 0000-0000A1900001

    b. Size—Indicates the number of UUID identifiers in the block. The size ranges from 1 to 1000.

6. Click Create.

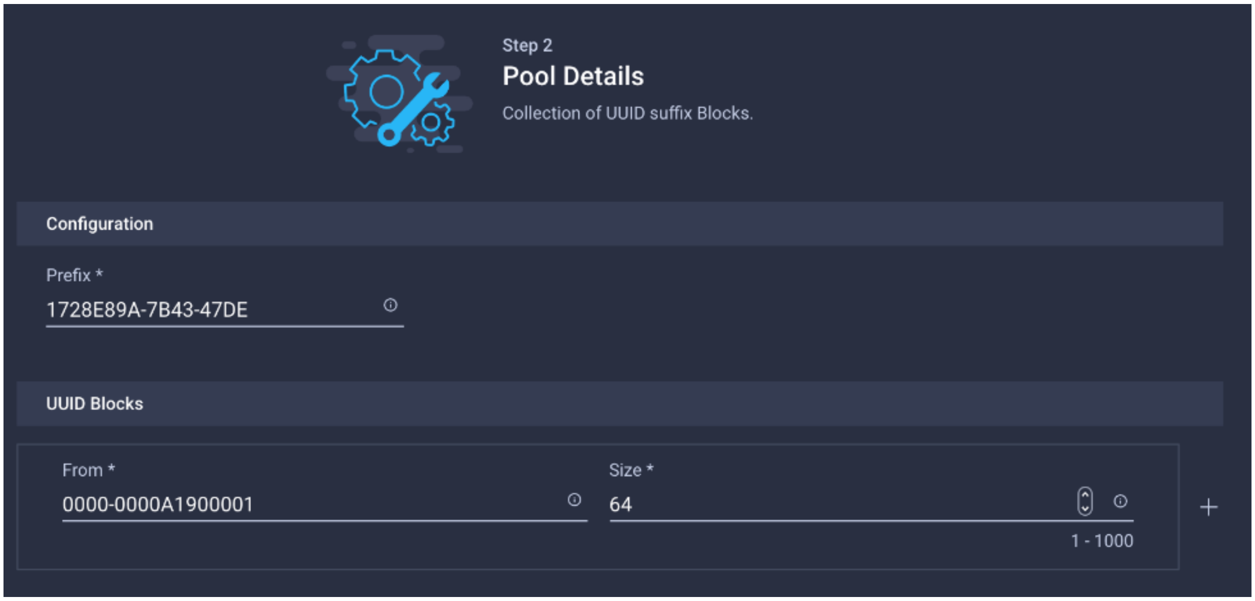

Step 2

Pool Details

Collection of UUID suffix Blocks.

Configuration

Prefix *

1728E89A-7B43-47DE

UUID Blocks

From *

0000-0000A1900001

Size *

64

1 - 1000

Configure BIOS Policy

1. Click Select Policy next to BIOS Configuration and in the pane on the right, click Create New.

2. Provide a name for the policy (for example, AA19-BIOS-Pol).

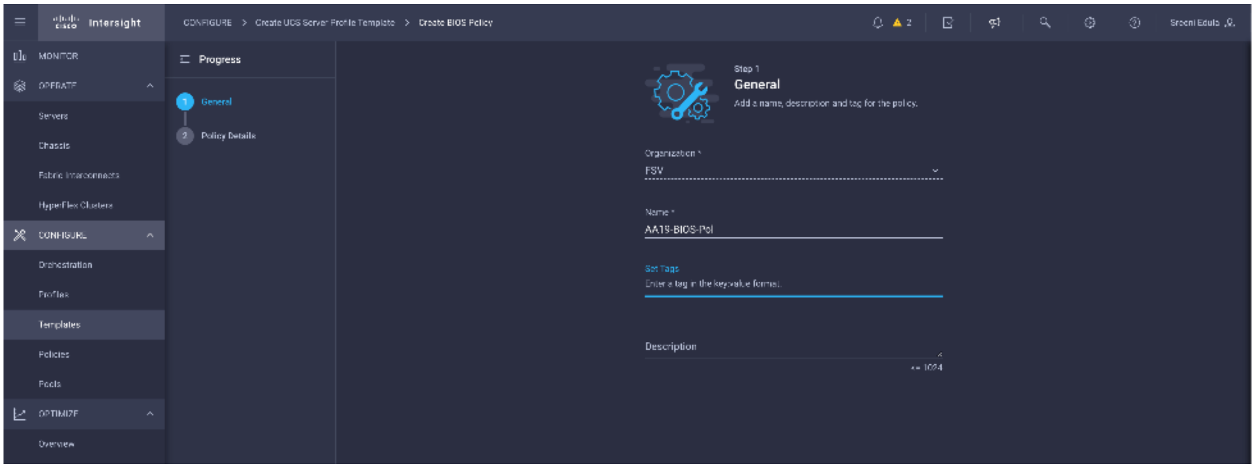

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-BIOS-Pol

Set Tags

Enter a tag in the key:value format.

Description

<= 1024

3. Click Next.

4. On the Policy Details screen, select appropriate values for the BIOS settings. In this deployment, the BIOS values were selected based on recommendations in the performance tuning guide for Cisco UCS M6 BIOS: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.html.

5. Select the appropriate values on the following screen:

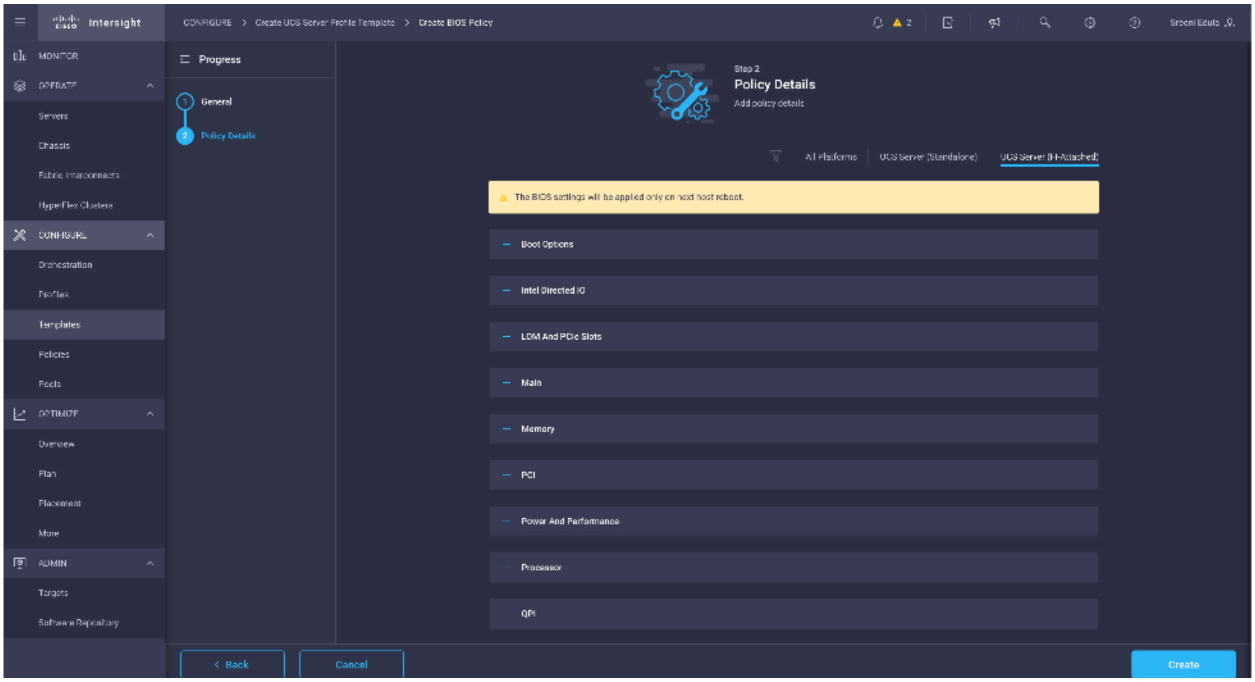

The BIOS settings will be applied only on next host reboot.

- LOM and PCIe Slot -> CDN Support for LOM: **Enabled**
- Processor -> Enhanced CPU performance: **Auto**
- Memory -> NVM Performance Setting: **Balanced Profile**

6. Click Create.

Configure Boot Order Policy for iSCSI Hosts

To configure Boot Order policy for iSCSI hosts, follow these steps. The FC boot order policy is different from iSCSI boot policy and is explained in the following section.

1. Click Select Policy next to BIOS Configuration and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-BootOrder-Pol).

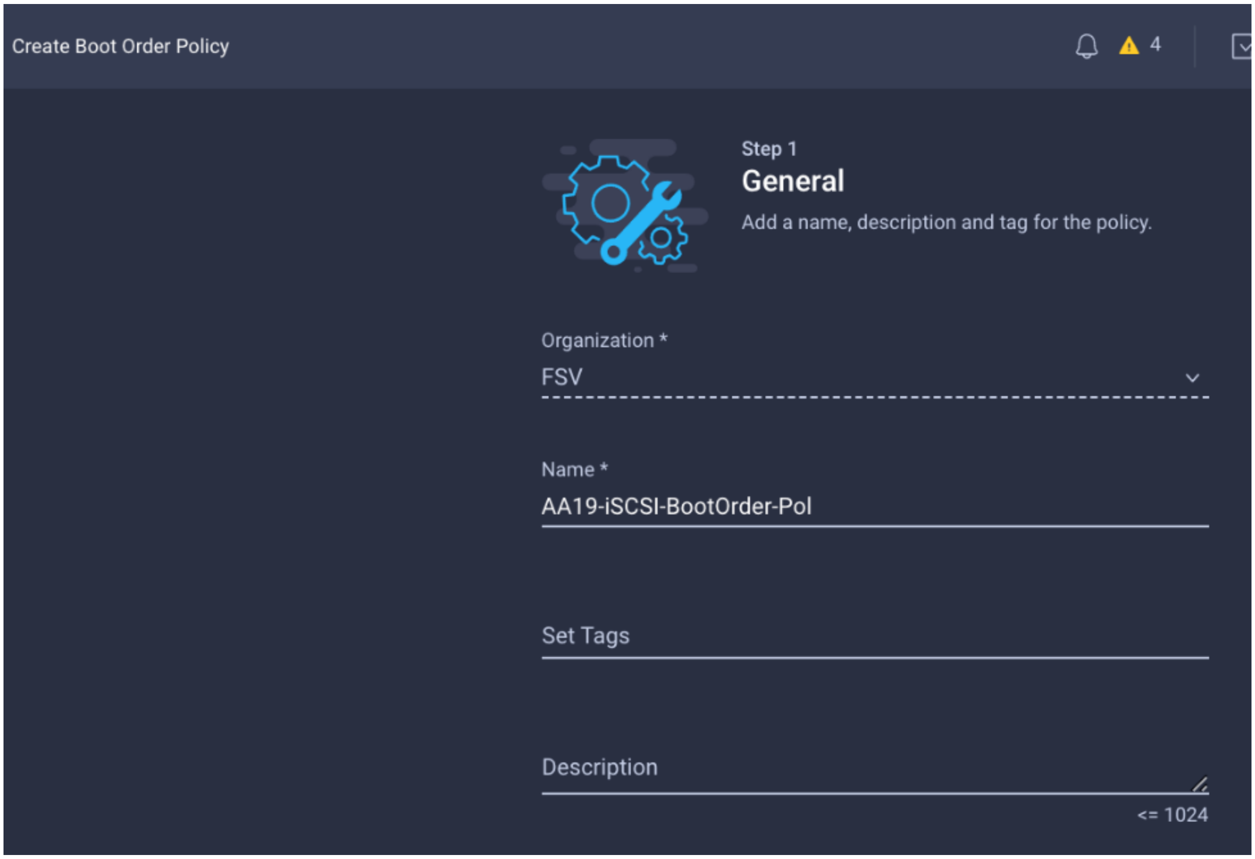

Create Boot Order Policy

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-iSCSI-BootOrder-Pol

Set Tags

Description

<= 1024

3. Click Next.

4. For Configured Boot Mode, select Unified Extensible Firmware Interface (UEFI).

5. Turn on Enable Secure Boot.

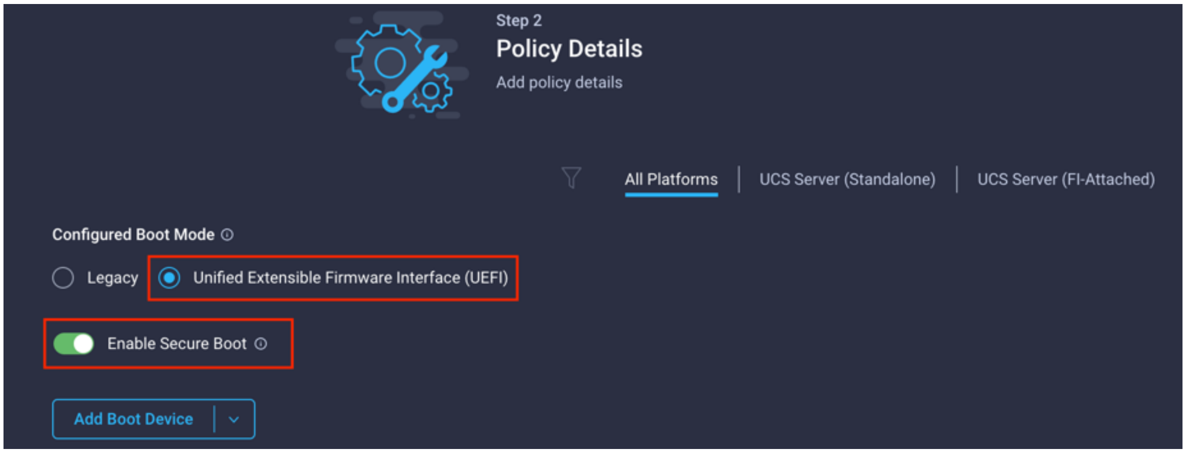

Step 2

Policy Details

Add policy details

Configured Boot Mode

○ Legacy ● Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot: on

6. Click Add Boot Device drop-down list and select Virtual Media.

7. Provide a device name (for example, ISO) and then, for the subtype, select KVM Mapped DVD.

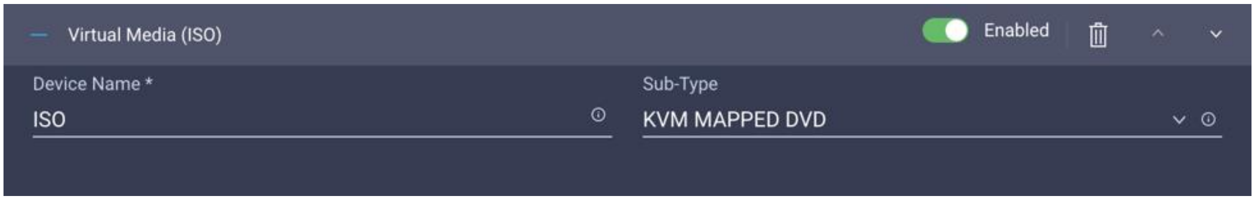

Device Name * ISO

Sub-Type KVM MAPPED DVD

8. From the Add Boot Device drop-down list, select iSCSI Boot.

9. Provide the Device Name: iSCSI-A-Boot and the exact name of the interface used for iSCSI boot under Interface Name: 04-iSCSI-A.

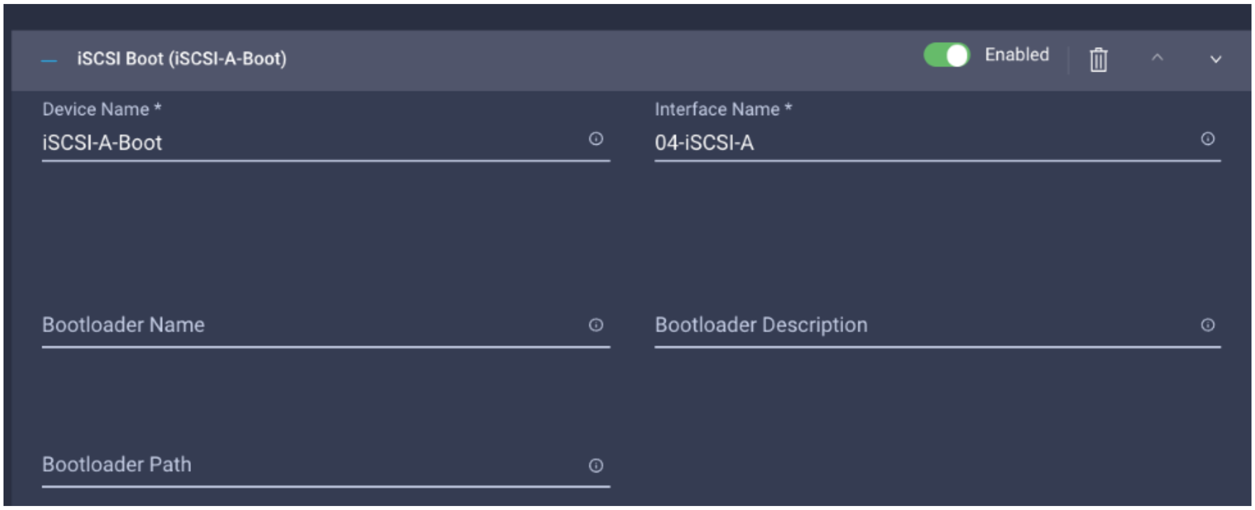

Device Name * iSCSI-A-Boot

Interface Name * 04-iSCSI-A

Bootloader Name

Bootloader Description

Bootloader Path

Note: The device names (iSCSI-A-Boot and iSCSI-B-Boot) are being defined here and will be used in the later steps of the ISCSI configuration.

10. From the Add Boot Device drop-down list, select iSCSI Boot.

11. Provide the Device Name: iSCSI-B-Boot and the exact name of the interface used for iSCSI boot under Interface Name: 05-iSCSI-B.

12. From the Add Boot Device drop-down list, select UEFI Shell.

13. Add Device Name UEFIShell.

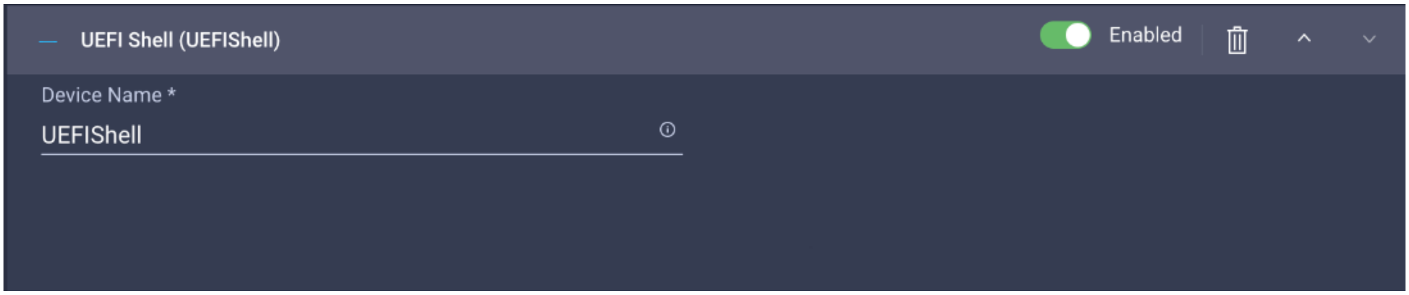

Device Name * UEFIShell

14. Verify the order of the boot policies and adjust the boot order as necessary using arrows next to delete button.

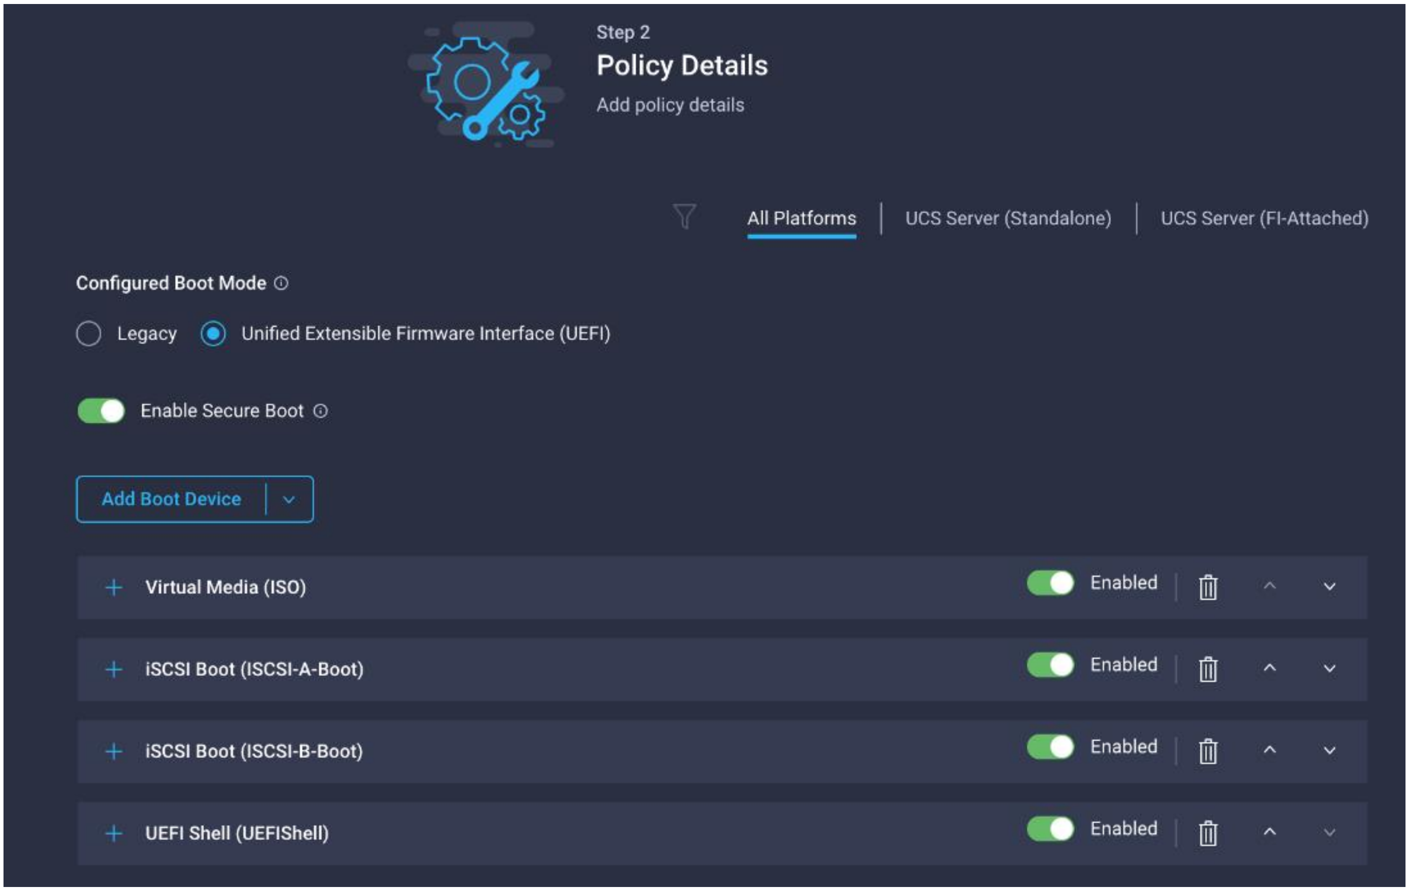

Configured Boot Mode

○ Legacy ● Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot: on

Boot Device	State
Virtual Media (ISO)	Enabled
iSCSI Boot (ISCSI-A-Boot)	Enabled
iSCSI Boot (ISCSI-B-Boot)	Enabled
UEFI Shell (UEFIShell)	Enabled

15. Click Create.

16. Click Next to move to Management Configuration.

Step 4: Management Configuration

Next, configure management policy. There policies will be added to the management configuration

- IMC Access to define the pool of IP addresses for compute node KVM access
- IPMI Over LAN to allow Intersight to manage IPMI messages
- Local User to provide local administrator to access KVM

Configure Cisco IMC access policy

To configure Cisco IMC access policy, follow these steps:

1. Click Select Policy next to IMC Access and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-IMC-Access).

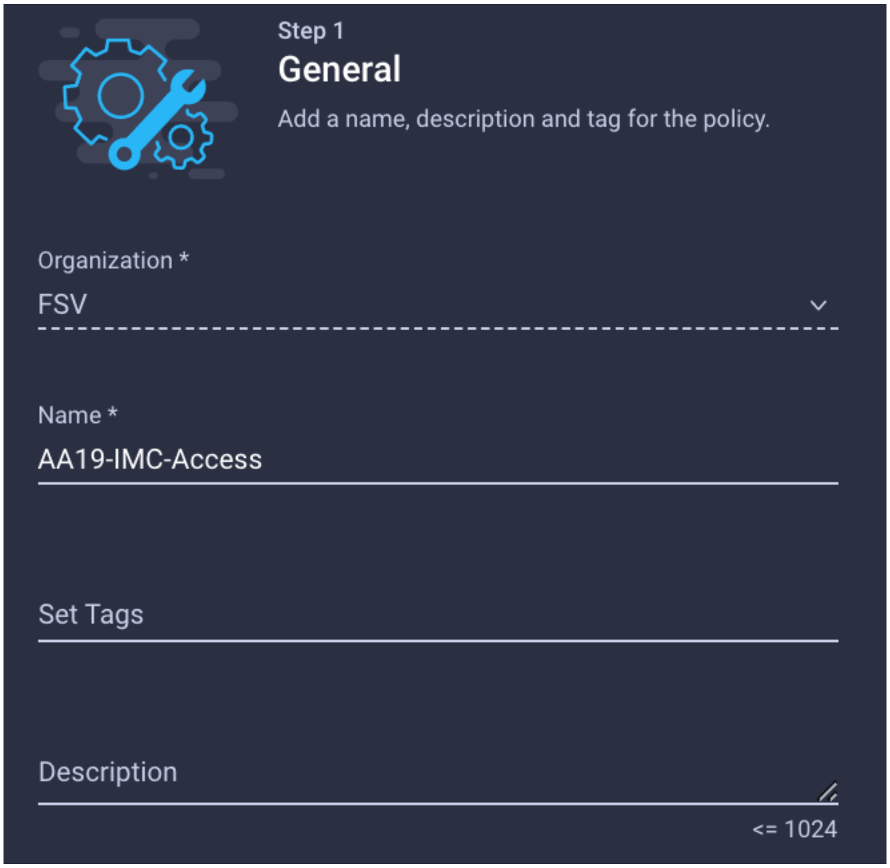

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-IMC-Access

Set Tags

Description

<= 1024

3. Click **Next**.

Note: Customers can select in-band management access to the compute node using an in-band management VLAN (for example, VLAN 115) or out-of-band management access via the Mgmt0 interfaces of the FIs. KVM Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Since these policies were not configured in this deployment, out-of-band management access was configured so that KVM access to compute nodes is not impacted by any potential switching issues in the fabric.

4. Click UCS Server (FI-Attached).

5. Enable Out-Of-Band Configuration.

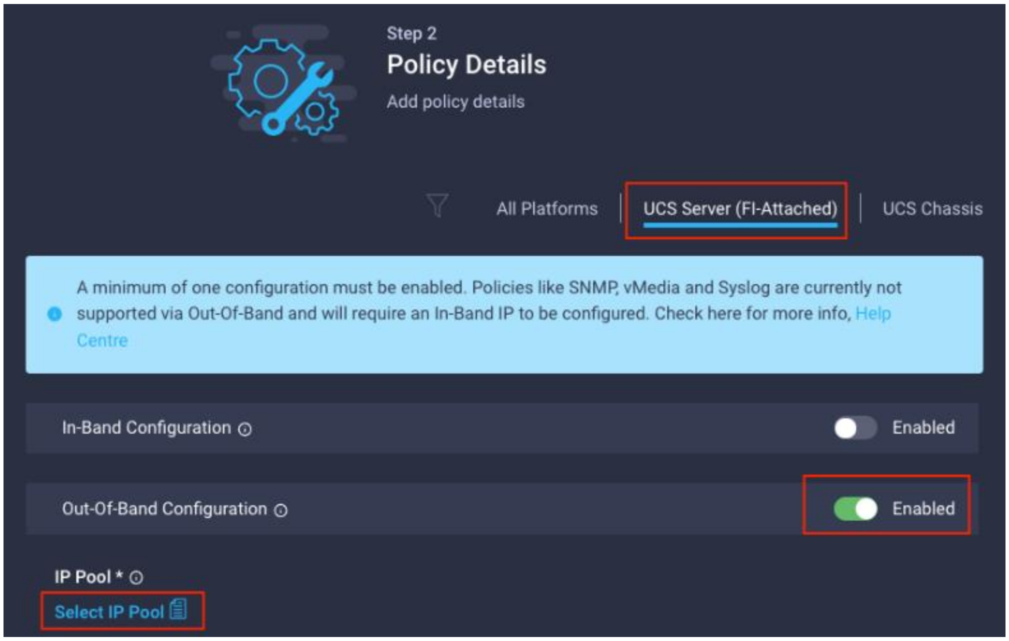

Step 2

Policy Details

Add policy details

A minimum of one configuration must be enabled. Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Check here for more info, Help Centre

In-Band Configuration ○ Enabled

Out-Of-Band Configuration ● Enabled

IP Pool *

Select IP Pool

6. Under IP Pool, click Select IP Pool and then, in the pane on the right, click Create New.

7. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-IMC-Pool).

8. Select Configure IPv4 Pool and provide the information to define a pool for KVM IP address assignment including an IP Block.

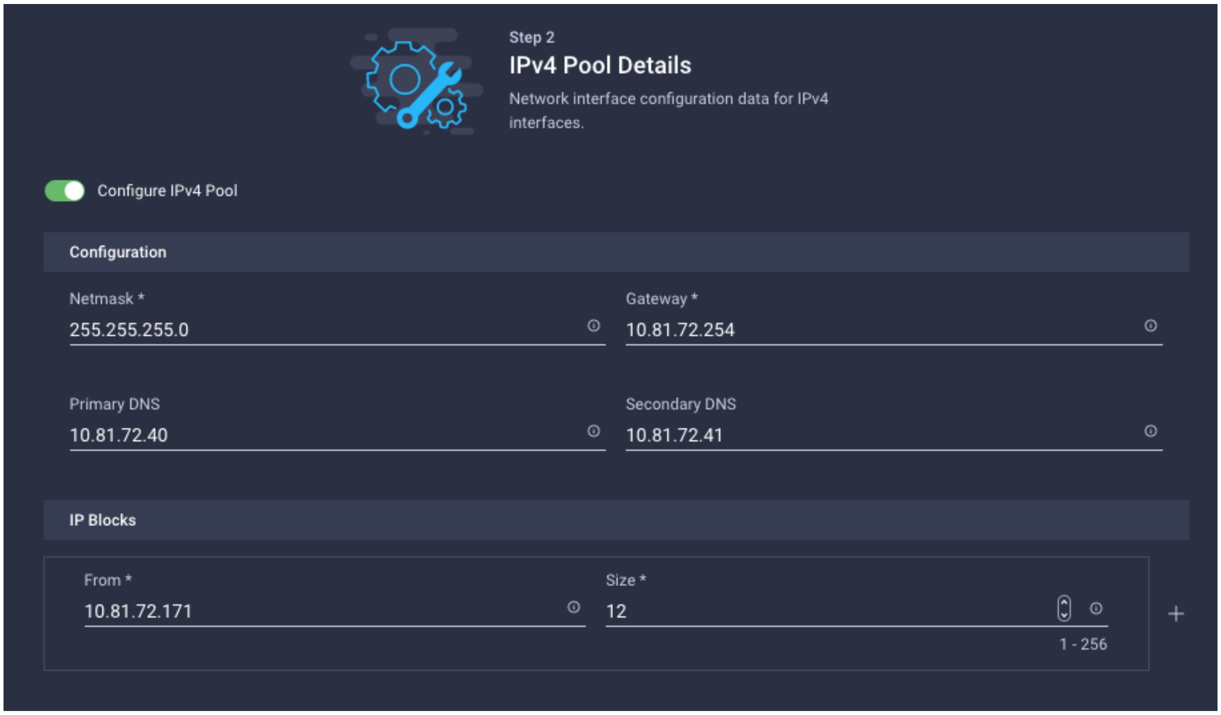

Step 2

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

● Configure IPv4 Pool

Configuration

Netmask * 255.255.255.0

Gateway * 10.81.72.254

Primary DNS 10.81.72.40

Secondary DNS 10.81.72.41

IP Blocks

From * 10.81.72.171

Size * 12

Note: The management IP pool subnet should be accessible from the host that is trying to open the KVM connection. In the example shown here, the hosts trying to open a KVM connection would need to be able to route to 10.81.72.0/24 subnet.

9. Click Next.

10. Unselect Configure IPv6 Pool.

11. Click Create to finish configuring the IP address pool.

Configure IPMI Over LAN policy

To configure IPMI Over LAN policy, follow these steps:

1. Click Select Policy next to IPMI Over LAN and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, Enable-IPMIoLAN).

3. Turn on Enable IPMI Over LAN.

4. From the Privilege Level drop-down list, select admin.

5. Click Create.

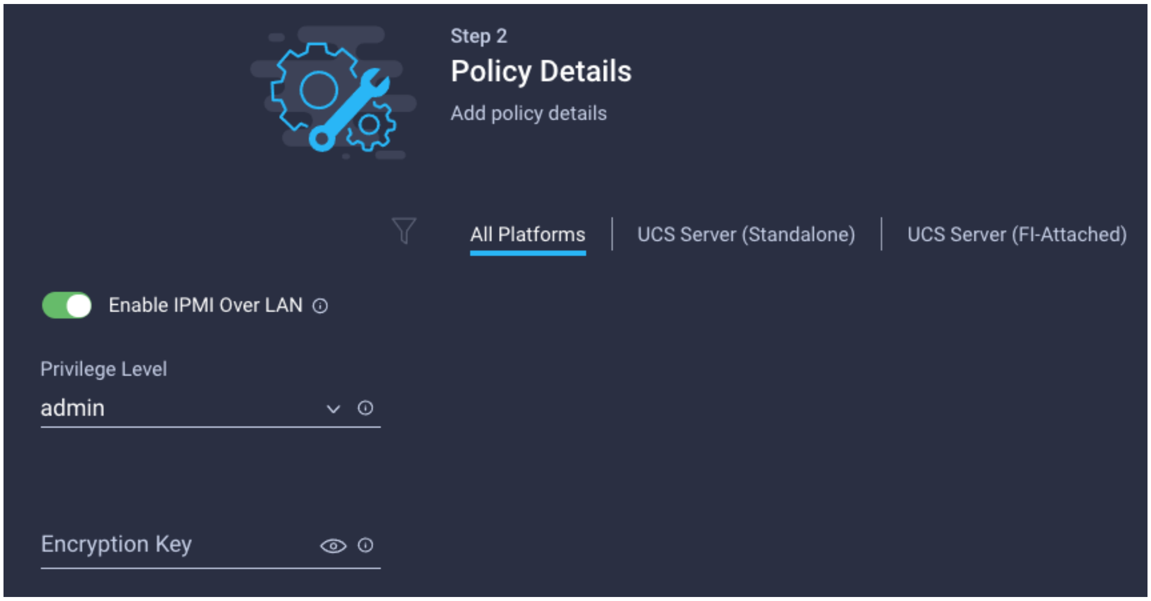

Step 2

Policy Details

Add policy details

Enable IPMI Over LAN: on

Privilege Level: admin

Encryption Key

Configure local user policy

To configure local user policy, follow these steps:

1. Click Select Policy next to Local User and the in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-LocalUser-Pol).

3. Verify that UCS Server (FI-Attached) is selected.

4. Verify that Enforce Strong Password is selected.

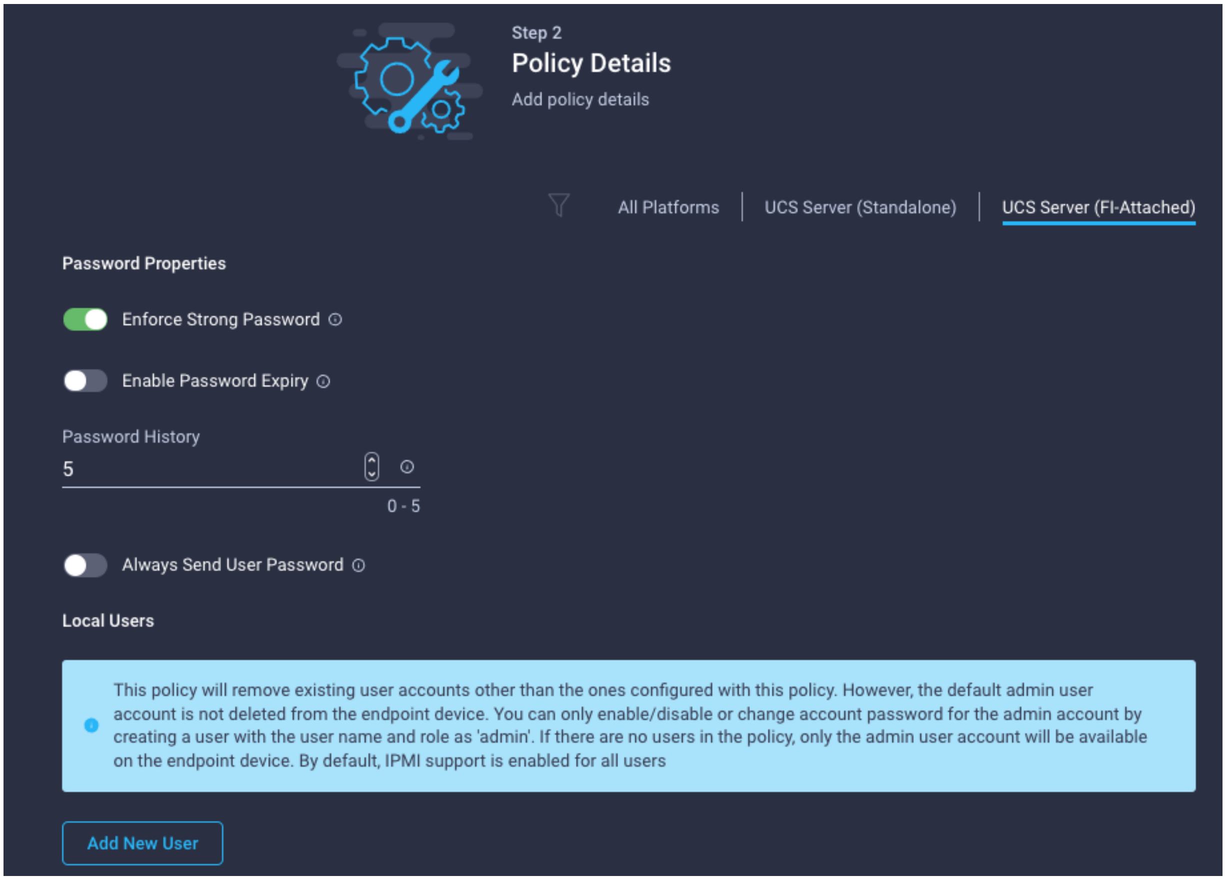

Step 2

Policy Details

Add policy details

Password Properties

Enforce Strong Password: enabled

Enable Password Expiry: disabled

Password History: 5 (0 - 5)

Always Send User Password: disabled

Local Users

This policy will remove existing user accounts other than the ones configured with this policy. However, the default admin user account is not deleted from the endpoint device. You can only enable/disable or change account password for the admin account by creating a user with the user name and role as 'admin'. If there are no users in the policy, only the admin user account will be available on the endpoint device. By default, IPMI support is enabled for all users

5. Click Add New User and then click + next to the New User

6. Provide the username (for example, flashadmin), choose a role for example, admin), and provide a password.

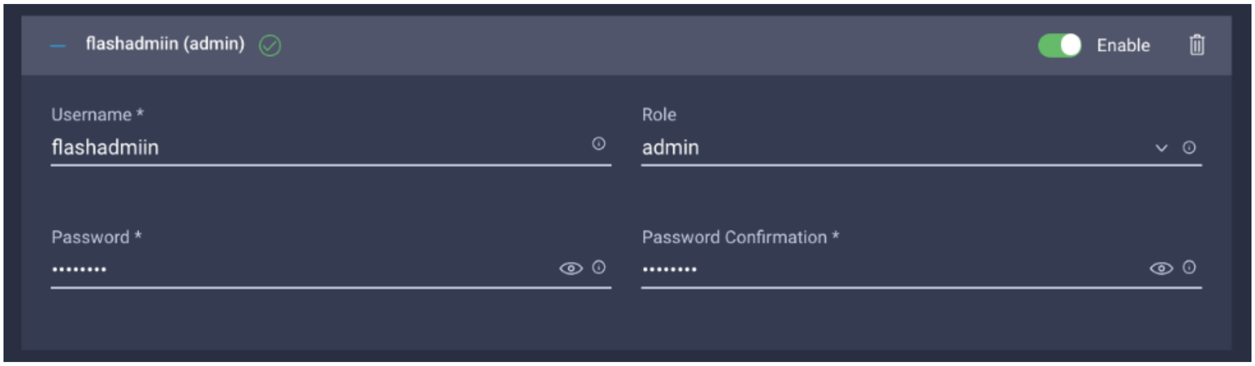

flashadmiin (admin)

Enable

Username * flashadmiin

Role admin

Password * ••••••••

Password Confirmation * ••••••••

Note: The username and password combination defined here will be used to log into KVMs. The typical Cisco UCS admin username and password combination cannot be used for KVM access.

7. Click Create to finish configuring the user.

8. Click Create to finish configuring local user policy.

9. Click Next to move to Storage Configuration.

Step 5: Storage Configuration

Click Next on the Storage Configuration screen. No configuration is needed in the local storage system.

Step 6a: Network Configuration > LAN Connectivity

LAN connectivity policy defines the connections and network communication resources between the server and the LAN. This policy uses pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For iSCSI hosts, this policy also defined an IQN address pool.

For consistent vNIC placement, manual vHBA/vNIC placement is utilized. iSCSI boot from SAN hosts and FC boot from SAN hosts require different number of vNICs/vHBAs and different placement order therefore the iSCSI host and the FC host LAN connectivity policies are covered separately in this document.

LAN Connectivity Policy for iSCSI Hosts

The iSCSI boot from SAN hosts uses 6 vNICs configured as follows:

Table 24. vNICs for iSCSI LAN Connectivity

vNIC/vHBA Name	Slot ID	Switch ID	PCI Order	VLANs
00-vSwitch0-A	MLOM	A	0	IB-MGMT, OOB-MGMT
01-vSwitch0-B	MLOM	B	1	IB-MGMT, OOB-MGMT
02-VDS0-A	MLOM	A	2	VM Traffic, vMotion
03-VDS0-B	MLOM	B	3	VM Traffic, vMotion
04-iSCSI-A	MLOM	A	4	iSCSI-A-VLAN
05-iSCSI-B	MLOM	B	5	iSCSI-B-VLAN

To create LAN connectivity policy for iSCSI hosts, follow these steps:

1. Click Select Policy next to LAN Connectivity and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-ESXi-LanConn-Manual). Click Next.

3. Under IQN, select Pool.

4. Click Select Pool under IQN Pool and then, in the pane on the right, click Create New.

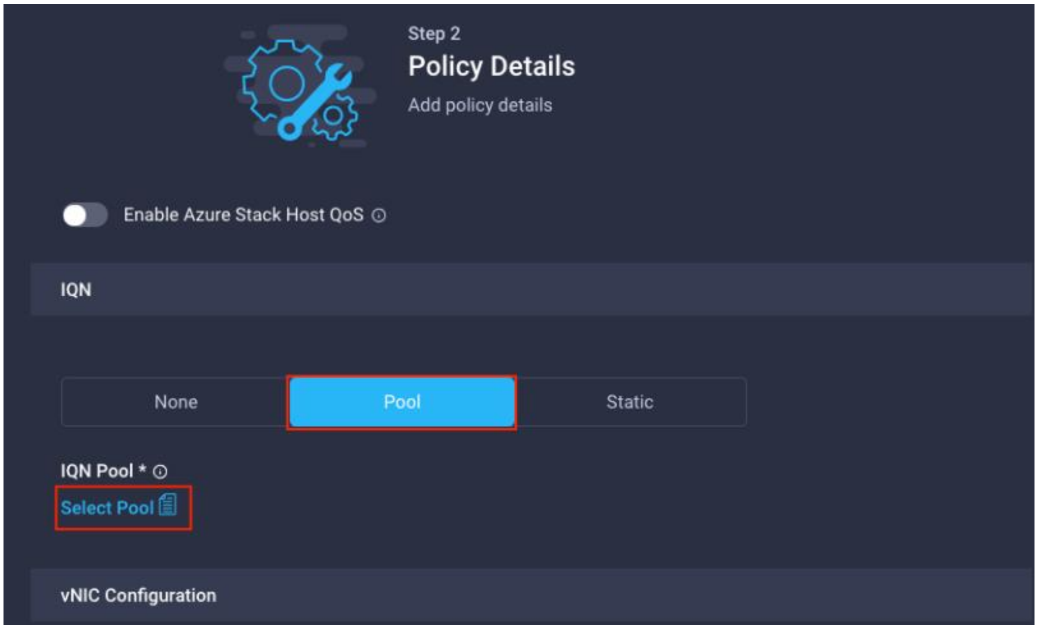

Step 2

Policy Details

Add policy details

Enable Azure Stack Host QoS

IQN

○ None ● Pool ○ Static

IQN Pool *

Select Pool

vNIC Configuration

5. Verify correct organization is selected from the drop-down list (for example, AA19) and provide a name for the IQN Pool (for example, AA19-IQN Pool).

6. Click Next.

7. Provide the values for Prefix and IQN Block to create the IQN pool.

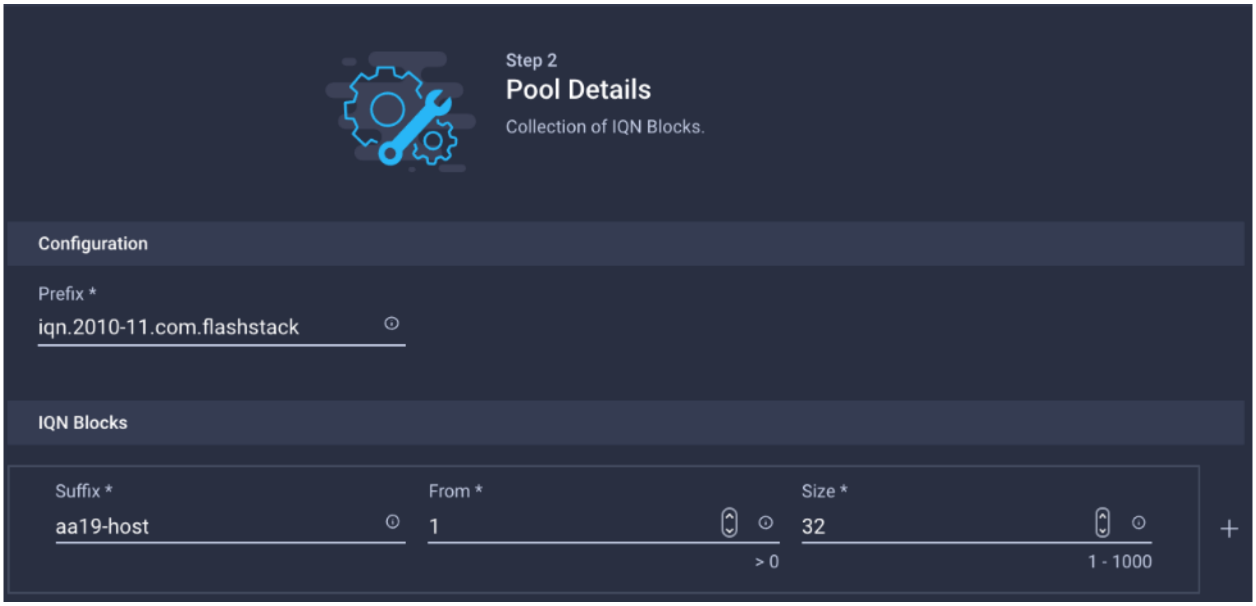

Step 2

Pool Details

Collection of IQN Blocks.

Configuration

Prefix *: iqn.2010-11.com.flashstack

IQN Blocks

Suffix *: aa19-host

From *: 1 (> 0)

Size *: 32 (1 - 1000)

8. Click Create.

9. Under vNIC Configuration, select Manual vNICs Placement.

10. Click Add vNIC.

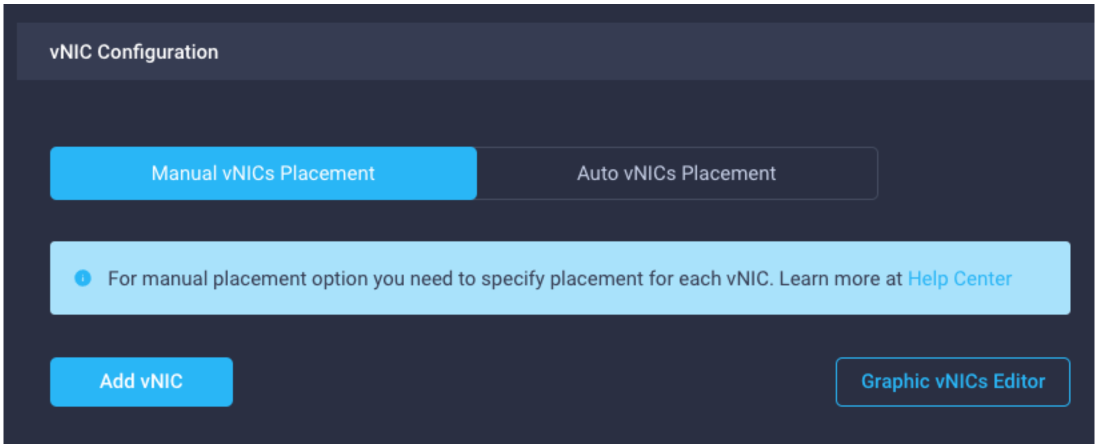
vNIC Configuration

For manual placement option you need to specify placement for each vNIC. Learn more at Help Center

Create MAC address pool

When creating the first vNIC, the MAC address pool has not been defined yet therefore a new MAC address pool will need to be created. Two separate MAC address pools are configured for each Fabric. MAC-Pool-A will be reused for all Fabric-A vNICs, and MAC-Pool-B will be reused for all Fabric-B vNICs.

Table 25. MAC Address Pools

Pool Name	Starting MAC Address	Size
MAC-Pool-A	00:25:B5:19:0A:00	128*
MAC-Pool-B	00:25:B5:19:0B:00	128*

Note: Each server requires 3 MAC addresses from the pool. Adjust the size of the pool according to your requirements.

To define the MAC pool for Fabric A/B, follow these steps:

1. Click Select Pool under MAC Address Pool and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the pool from Table 25 depending on the vNIC being created (for example, MAC-Pool-A for Fabric A).

3. Click Next.

4. Provide the starting MAC address from Table 25 (for example, 00:25:B5:19:0A:00)

Note: For ease of troubleshooting FlexPod, some additional information is always coded into the MAC address pool. For example, in the starting address 00:25:B5:19:0A:00, 19 is the rack ID and 0A indicates Fabric A.

5. Provide the size of the MAC address pool from Table 25 (for example, 128).

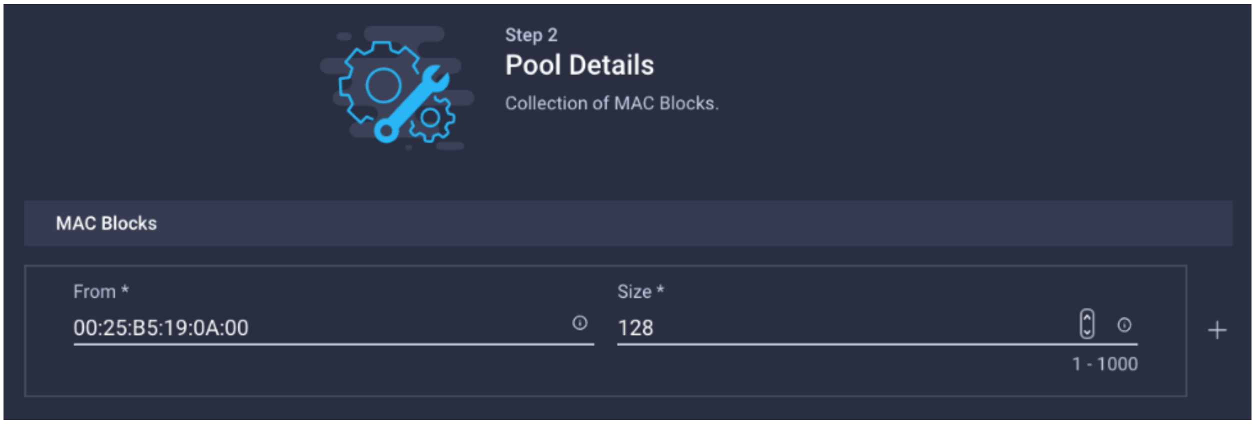

Step 2

Pool Details

Collection of MAC Blocks.

MAC Blocks

From * 00:25:B5:19:0A:00

Size * 128

1 - 1000

6. Click Create to finish creating the MAC address pool.

7. From the Add vNIC window:

   a. Provide vNIC Name, Slot ID, Switch ID and PCI Order information from Table 24.

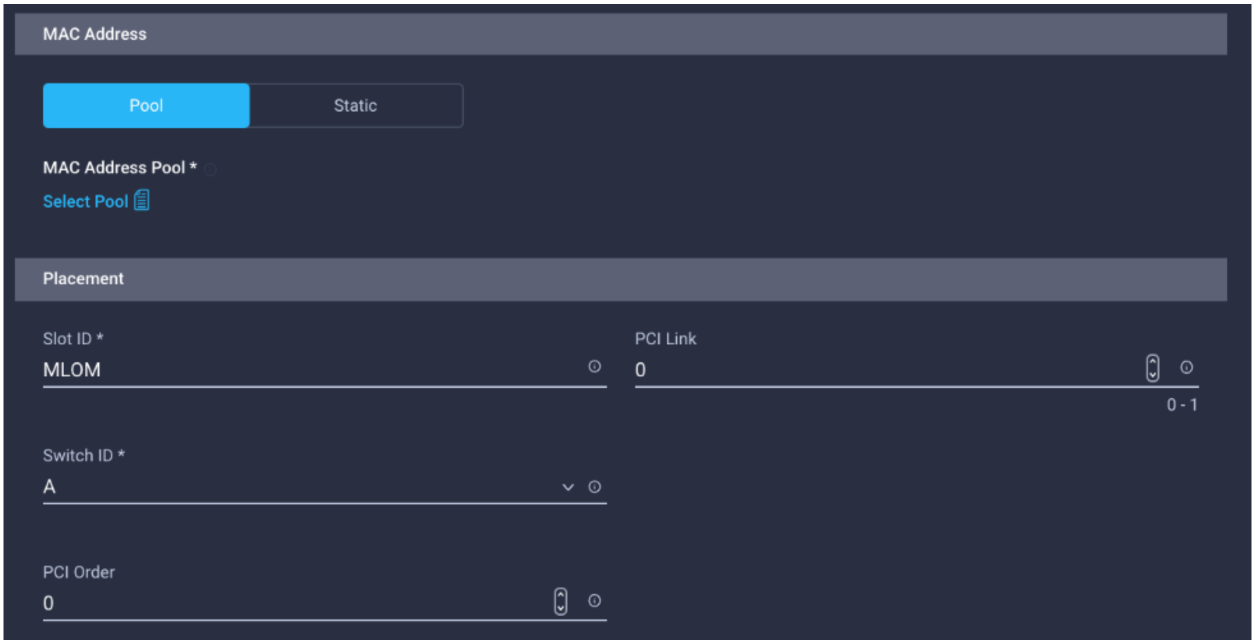

MAC Address

MAC Address Pool *

Placement

Slot ID * MLOM

PCI Link 0

0 - 1

Switch ID * A

PCI Order 0

8. For Consistent Device Naming (CDN), from the drop-down list, select **vNIC Name**.

9. Verify that Failover is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and VDS.

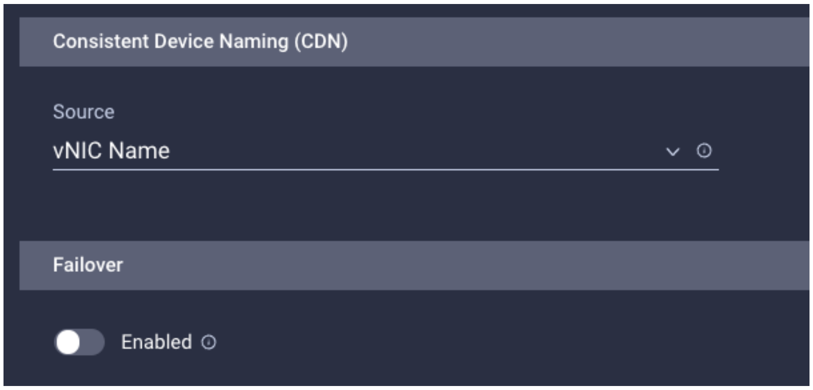
Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled: off

Create Ethernet Network Group Policy

Ethernet Network Group policies will be created and reused on applicable vNICs as covered below. Ethernet network group policy defines the VLANs allowed for a particular vNIC therefore multiple network group policies will be defined for this deployment as follows:

Table 26. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs
AA19-vSwitch0-NetGrp	Native-VLAN (2)	00-vSwitch0-A, 01-vSwitch0-B
AA19-VDS0-NetGrp	Native-VLAN (2)	02-VDS0-A, 03-VDS0-B
AA19-iSCSI-A-NetGrp	iSCSI-A-VLAN	05-iSCSI-A
AA19- iSCSI-B-NetGrp	iSCSI-B-VLAN	06-iSCSI-B

To define Ethernet Group Policy for a vNIC, follow these steps:

1. Click Select Policy under Ethernet Network Group Policy and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy from the Table 26 (for example, AA19-vSwitch0-NetGrp).

3. Click Next.

4. Enter the allowed VLANs from the Table 2 (for example, 15,115) and the native VLAN ID from Table 2 (for example, 2).

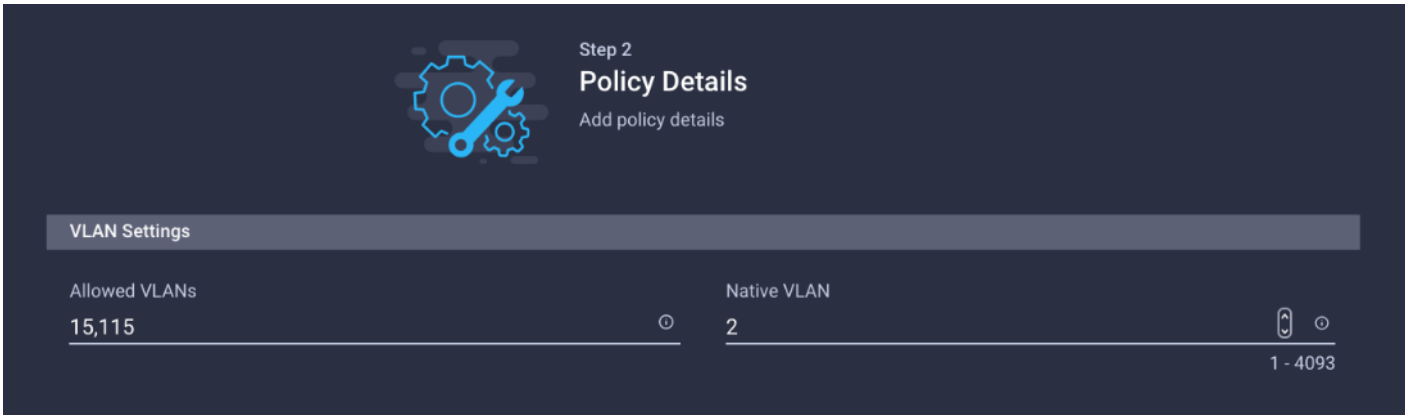

VLAN Settings

Allowed VLANs 15,115

Native VLAN 2

5. Click Create to finish configuring the Ethernet network group policy.

6. Repeat steps 1 – 5 to create AA19-iSCSI-A-NetGrp policy.

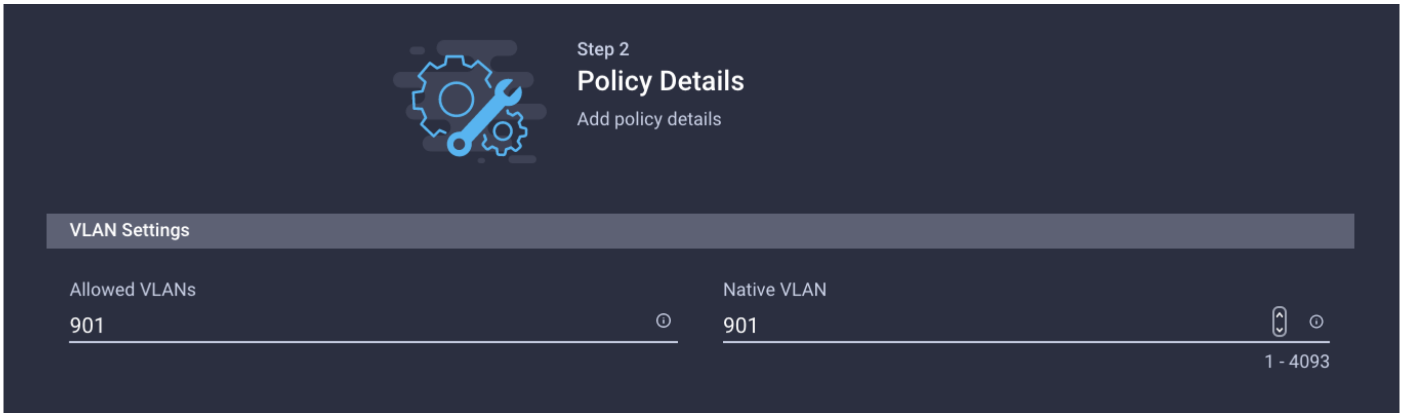

VLAN Settings

Allowed VLANs 901

Native VLAN 901

7. Repeat steps 1 – 5 again to create AA19-iSCSI-B-NetGrp Policy.

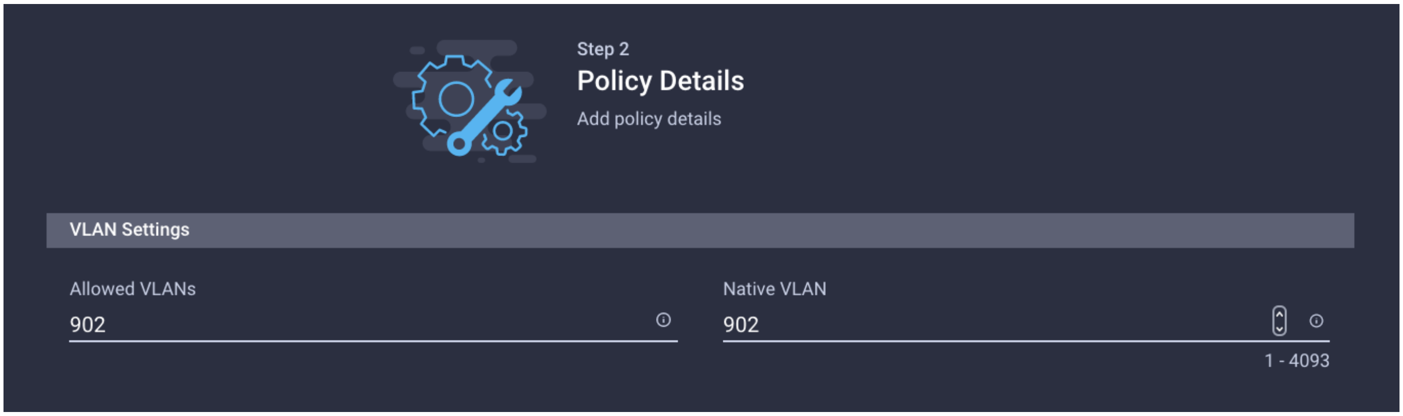

VLAN Settings

Allowed VLANs 902

Native VLAN 902

Note: When ethernet group policies are shared between two vNICs, the ethernet group policy only needs to be defined for the first vNIC. For subsequent vNIC policy mapping, just click **Select Policy** and pick the previously defined ethernet group policy from the list on the right.

Create Ethernet Network Control Policy

Ethernet Network Control Policy is used to enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) for the vNICs. A single policy will be created here and reused for all the vNICs.

To create the Ethernet Network Control Policy, follow these steps:

1. Click Select Policy under Ethernet Network Control Policy and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, AA19) and provide a name for the policy (for example, AA19-Enable-CDP-LLDP).

3. Click Next.

4. Enable Cisco Discovery Protocol and both Enable Transmit and Enable Receive under LLDP.

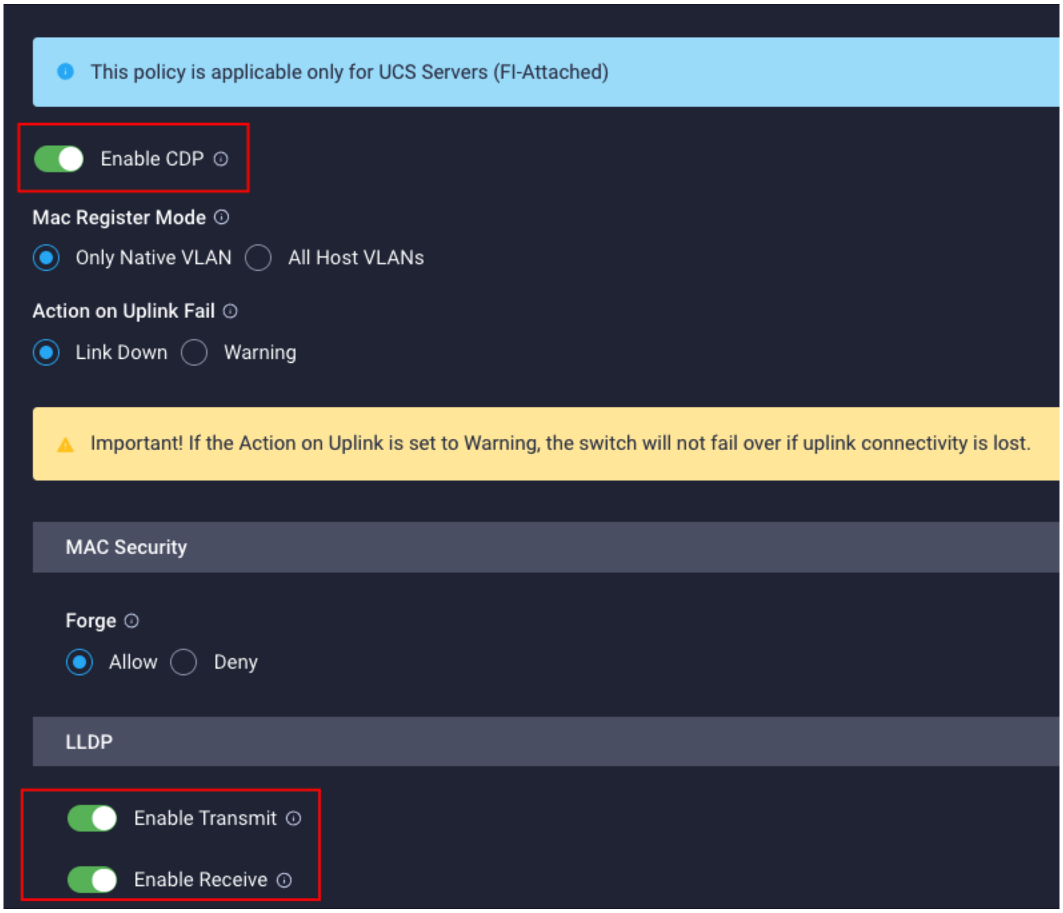

This policy is applicable only for UCS Servers (FI-Attached)

Enable CDP

Mac Register Mode

● Only Native VLAN ○ All Host VLANs

Action on Uplink Fail

● Link Down ○ Warning

Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost.

MAC Security

Forge

● Allow ○ Deny

LLDP

Enable Transmit

Enable Receive

5. Click Create to finish creating Ethernet network control policy.

Create Ethernet QoS policy

Ethernet QoS policy is used to enable jumbo maximum transmission units (MTUs) for all the vNICs. A single policy will be created and reused for all the vNICs. To create the ethernet QoS policy, follow these steps:

1. Click Select Policy under Ethernet QoS and in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-EthQos-Pol).

3. Click Next.

4. Change the MTU, Bytes value to 9000.

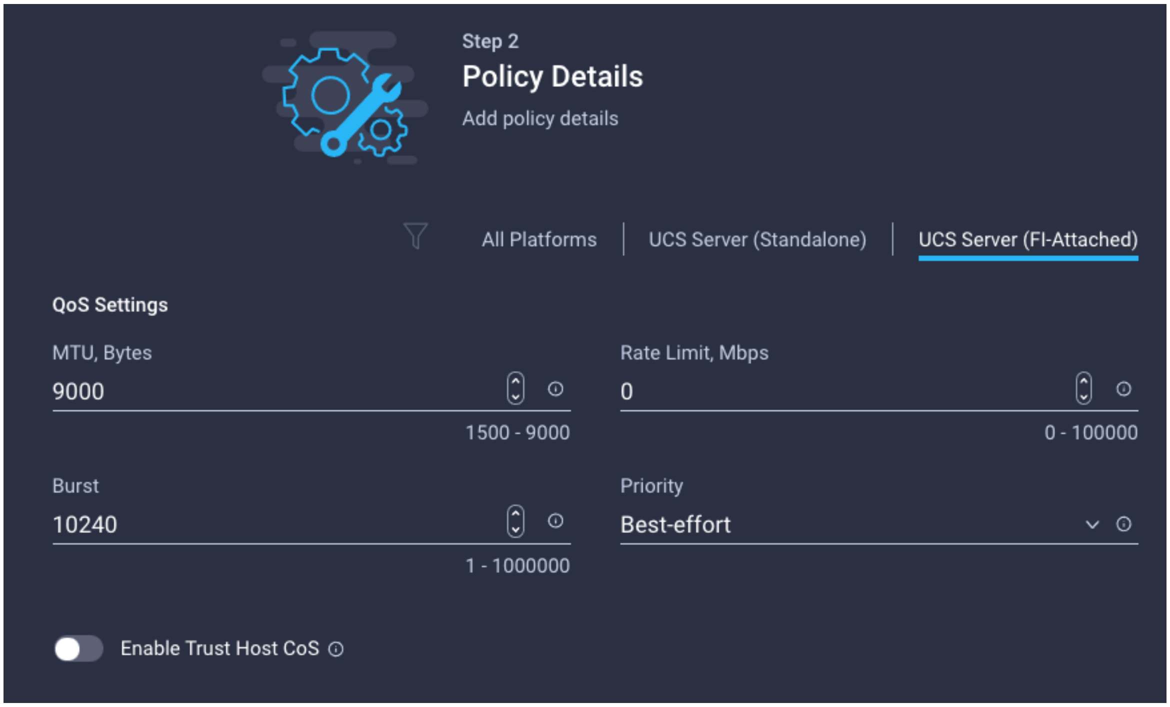

Step 2

Policy Details

Add policy details

QoS Settings

MTU, Bytes
9000
1500 - 9000

Rate Limit, Mbps
0
0 - 100000

Burst
10240
1 - 1000000

Priority
Best-effort

○ Enable Trust Host CoS

5. Click Create to finish setting up the Ethernet QoS policy.

Create Ethernet adapter policy

Ethernet adapter policy is used to set the interrupts and the send and receive queues. The values are set according to the best-practices guidance for the operating system in use. Cisco Intersight provides default VMware Ethernet Adapter policy for typical VMware deployments.

Customers can optionally configure a tweaked ethernet adapter policy for additional hardware receive queues handled by multiple CPUs in scenarios where there is a lot of vMotion traffic and multiple flows. In this deployment, a modified ethernet adapter policy, AA19-VMware-High-Traffic, is created and attached to the 02-VDS0-A and 03-VDS0-B interfaces which handle vMotion.

Table 27. Ethernet Adapter Policy association to vNICs

Policy Name	vNICs
AA19-EthAdapter-VMware	00-vSwitch0-A, 01-vSwitch0-B, 04-iSCSI-A, 05-iSCSI-B
AA19-VMware-High-Traffic	02-VDS0-A, 03-VDS0-B,

1. Click Select Policy under Ethernet Adapter and then, in the pane on the right, click Create New.

2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-EthAdapter-VMware).

3. Click Select Default Configuration under Ethernet Adapter Default Configuration.

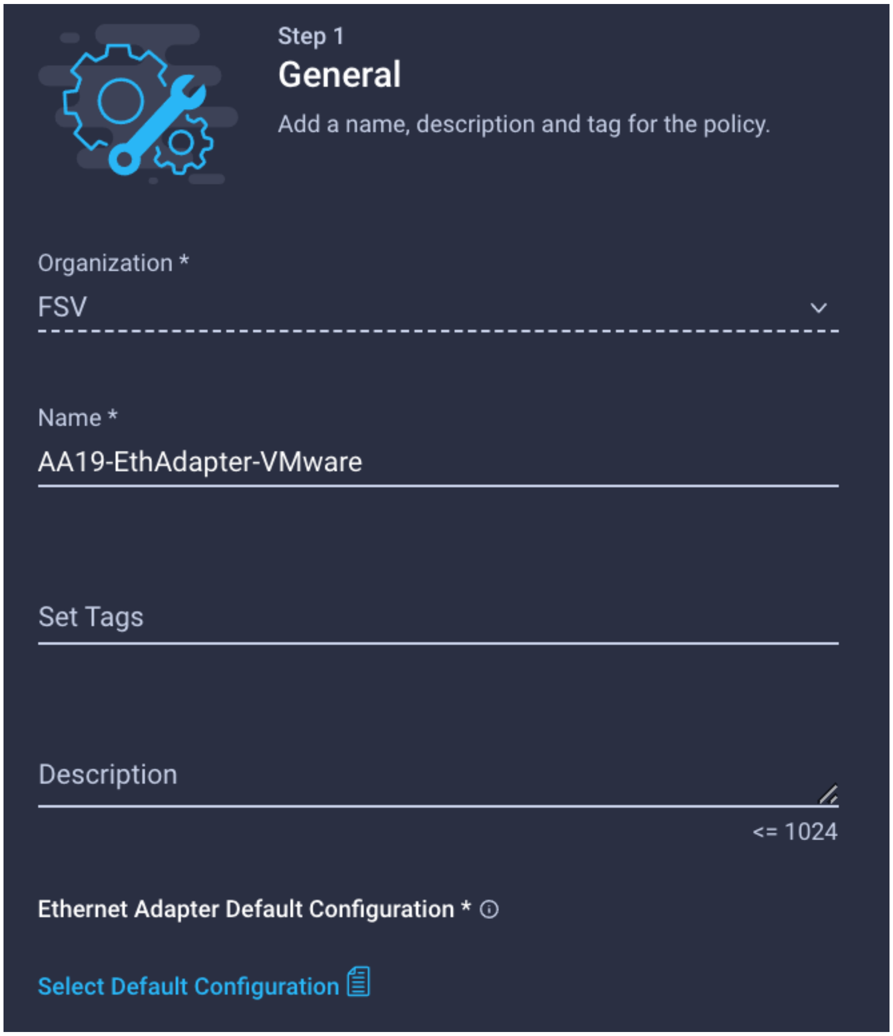

Step 1

General

Add a name, description and tag for the policy.

Organization *

FSV

Name *

AA19-EthAdapter-VMware

Set Tags

Description

<= 1024

Ethernet Adapter Default Configuration *

Select Default Configuration

4. From the list, select VMware.

5. Click Next.

6. For the AA19-EthAdapter-VMware policy, click Create and skip the rest of the steps in this “Create Ethernet Adapter Policy” section.

7. For the optional AA19-VMware-High-Traffic policy (for VDS interfaces), make the following modifications to the policy:

   - Increase Interrupts to 11
   - Increase Receive Queue Count to 8
   - Increase Completion Queue Count to 9
   - Enable Receive Side Scaling

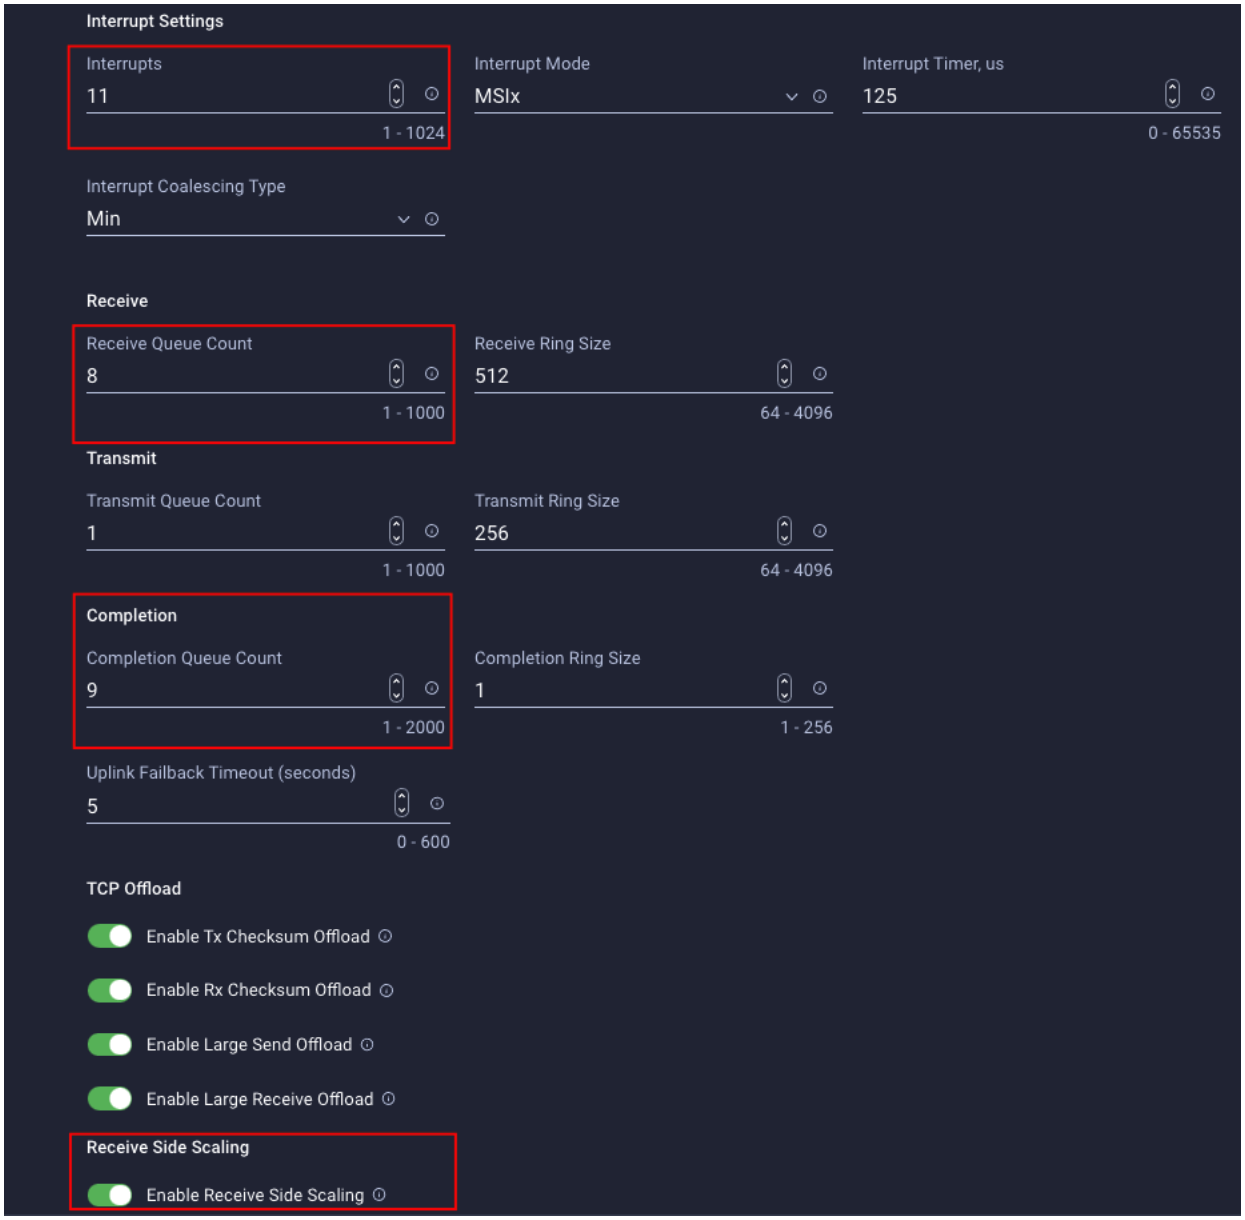

Interrupt Settings

Interrupts: 11 (1 - 1024)

Interrupt Mode: MSIx

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type: Min

Receive

Receive Queue Count: 8 (1 - 1000)

Receive Ring Size: 512 (64 - 4096)

Transmit

Transmit Queue Count: 1 (1 - 1000)

Transmit Ring Size: 256 (64 - 4096)

Completion

Completion Queue Count: 9 (1 - 2000)

Completion Ring Size: 1 (1 - 256)

Uplink Failback Timeout (seconds): 5 (0 - 600)

TCP Offload

Enable Tx Checksum Offload

Enable Rx Checksum Offload

Enable Large Send Offload

Enable Large Receive Offload

Receive Side Scaling

Enable Receive Side Scaling

8. Click Create.

Create iSCSI-A Policy

iSCSI-A policy only applied to vNICs 04-ISCSI-A and should not be created for data vNICs (vSwitch0 and VDS). The iSCSI-B policy creation is explained in the following section.

To create the iSCSI boot policy, follow these steps:

1. Click Select Policy under iSCSI Boot and then, in the pane on the right, click Create New.
2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-Boot-A).
3. Click Next.
4. Select Static under Configuration.

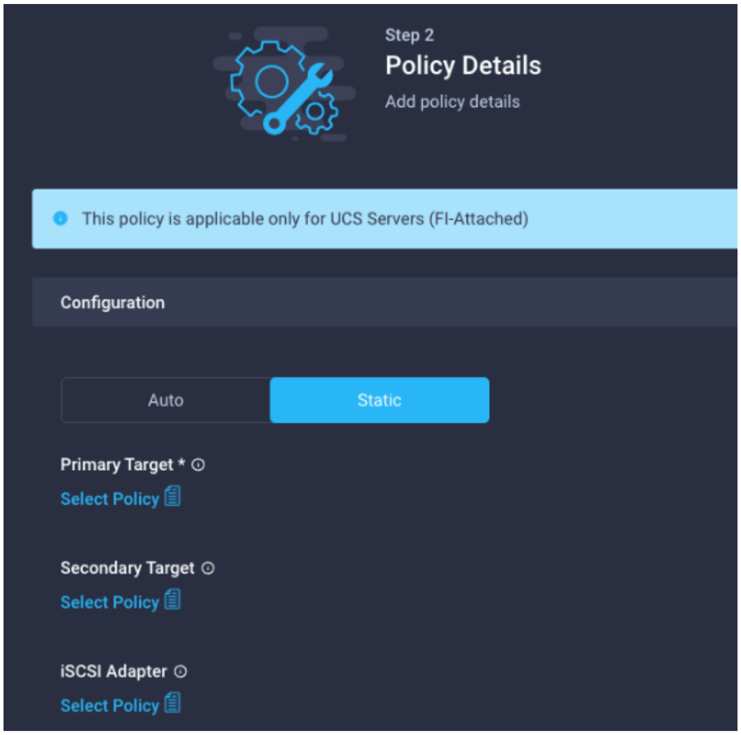

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Configuration

Primary Target *

Select Policy

Secondary Target

Select Policy

iSCSI Adapter

Select Policy

5. Click Select Policy under Primary Target and then, in the pane on the right, click Create New.

6. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-A-Primary-Target).

7. Click Next.

8. Provide the Target Name captured from Pure FlashArray, IP Address of ct0.eth4, Port 3260 and Lun ID of 1.

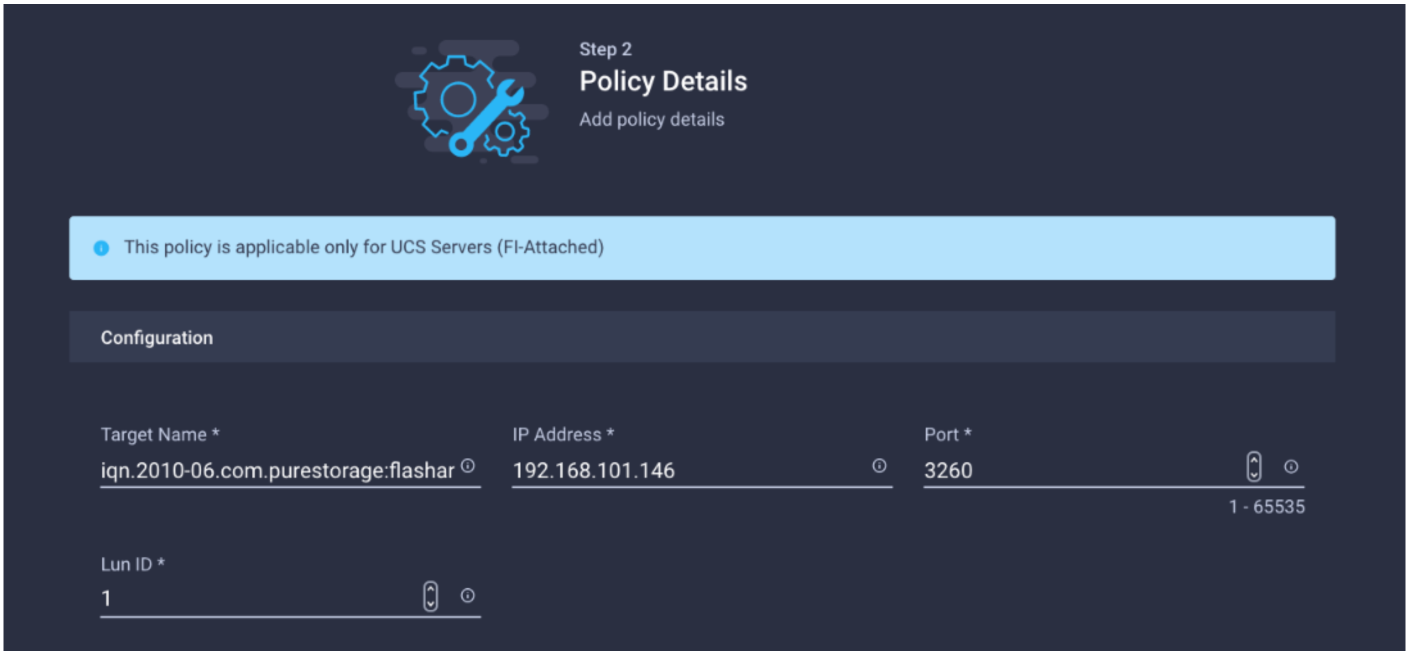

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Configuration

Target Name *: iqn.2010-06.com.purestorage:flashar

IP Address *: 192.168.101.146

Port *: 3260

1 - 65535

Lun ID *: 1

9. Click Create.

10. Click Select Policy under Secondary Target and then, in the pane on the right, click Create New.

11. Verify correct organization is selected from the drop-down list (for example, AA19) and provide a name for the policy (for example, AA19-iSCSI-A-Secondary-Target).

12. Click Next.

13. Provide the Target Name captured from Pure FlashArray, IP Address of ct1.eth4, Port 3260 and Lun ID of 1

14. Click Create.

15. Click Select Policy under iSCSI Adapter and then, in the pane on the right, click Create New.

16. Verify correct organization is selected from the drop-down list (for example, AA19) and provide a name for the policy (for example, AA19-iSCSI-AdapterPol).

17. Click Next.

18. Accept the default policies. Customers can adjust the timers if necessary.

19. Click Create.

20. Scroll down to Initiator IP Source and make sure Pool is selected.

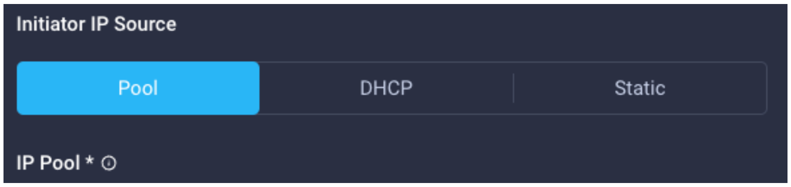

21. Click Select Pool under IP Pool and then, in the pane on the right, click Create New.

22. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the pool (for example, AA19-iSCSI-A-Pool).

23. Click Next.

24. Make sure Configure IPv4 Pool is selected. Enter the IP pool information for iSCSI-A subnet.

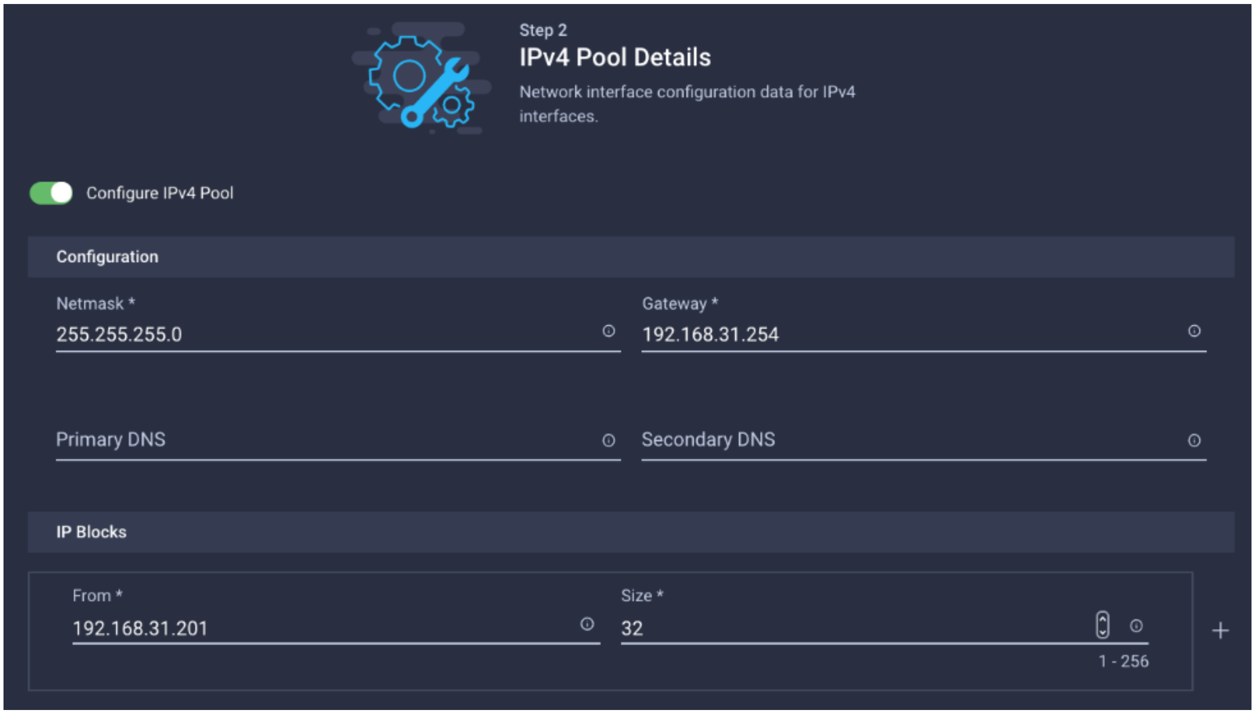

Step 2

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Configuration

Netmask * 255.255.255.0

Gateway * 192.168.31.254

Primary DNS

Secondary DNS

IP Blocks

From * 192.168.31.201

Size * 32

1 - 256

Note: Since the iSCSI network is not routable and all the VMkernel ports and LIFs are layer-2 adjacent, there is no need to define a gateway or DNS.

25. Click Next.

26. Disable Configure IPv6 Pool.

27. Click Create.

28. Verify all the policies and pools are correctly mapped for the iSCSI-A policy.

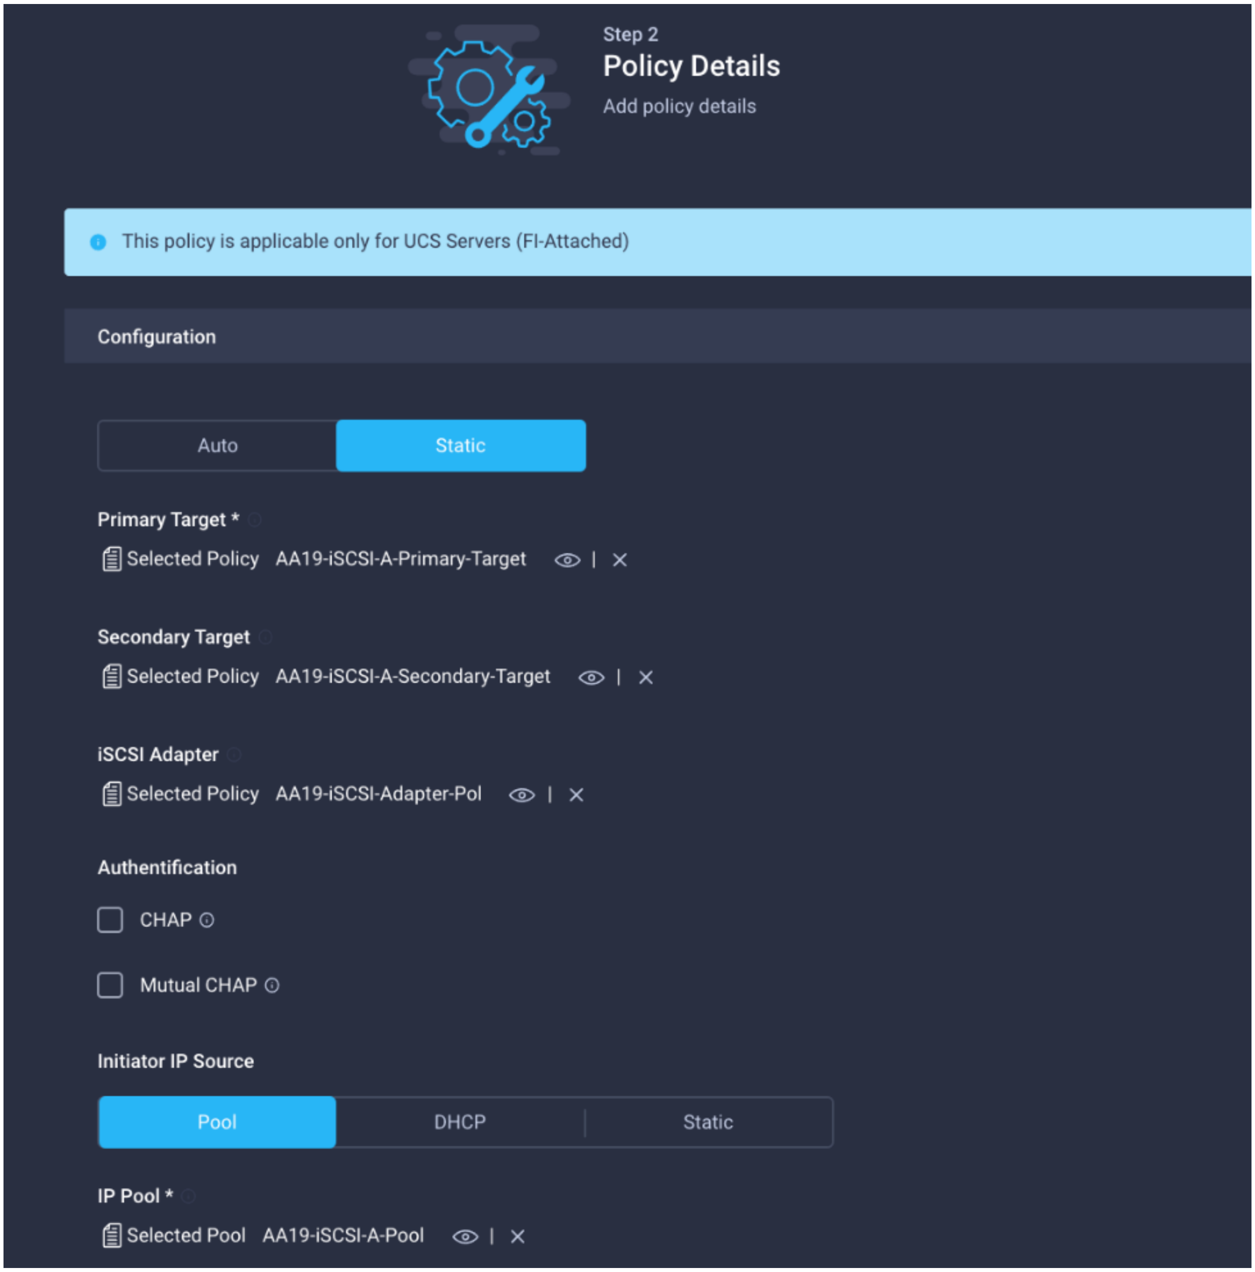

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Configuration

Primary Target *

Selected Policy AA19-iSCSI-A-Primary-Target

Secondary Target

Selected Policy AA19-iSCSI-A-Secondary-Target

iSCSI Adapter

Selected Policy AA19-iSCSI-Adapter-Pol

Authentification

☐ CHAP

☐ Mutual CHAP

Initiator IP Source

IP Pool *

Selected Pool AA19-iSCSI-A-Pool

29. Click Create.

Create iSCSI-B Policy

Note: The iSCSI-B policy is only applied to vNICs 05-ISCSI-B and should not be created for data vNICs (vSwitch0 and VDS).

To create the iSCSI boot policy, follow these steps:

1. Click Select Policy under iSCSI Boot and then, in the pane on the right, click Create New.
2. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-Boot-B).
3. Click Next.

4. Select Static under Configuration.

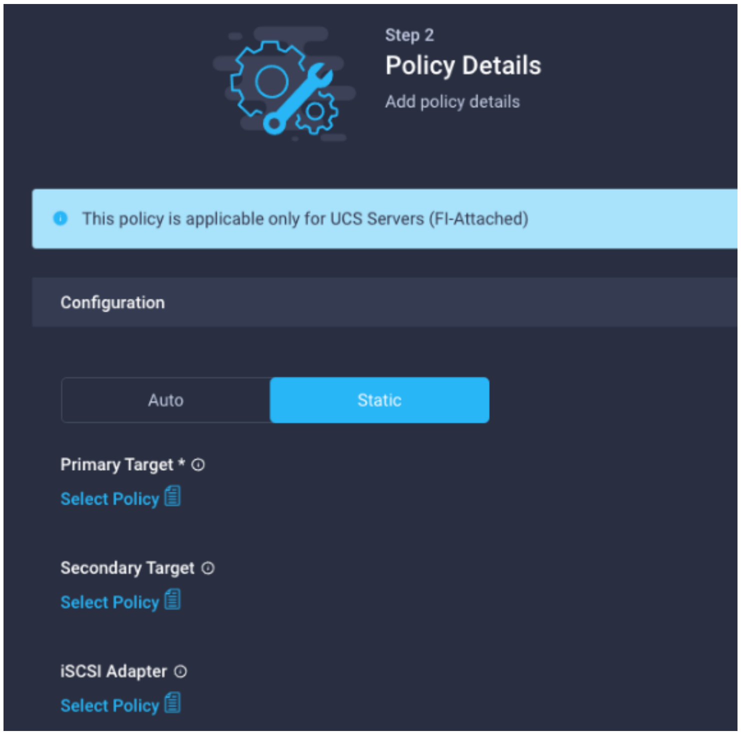

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Configuration

Primary Target *

Select Policy

Secondary Target

Select Policy

iSCSI Adapter

Select Policy

5. Click Select Policy under Primary Target and then, in the pane on the right, click Create New.

6. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-B-Primary-Target).

7. Click Next.

8. Provide the Target Name captured from Pure FlashArray, IP Address of ct0.eth5, Port 3260 and Lun ID of 1.

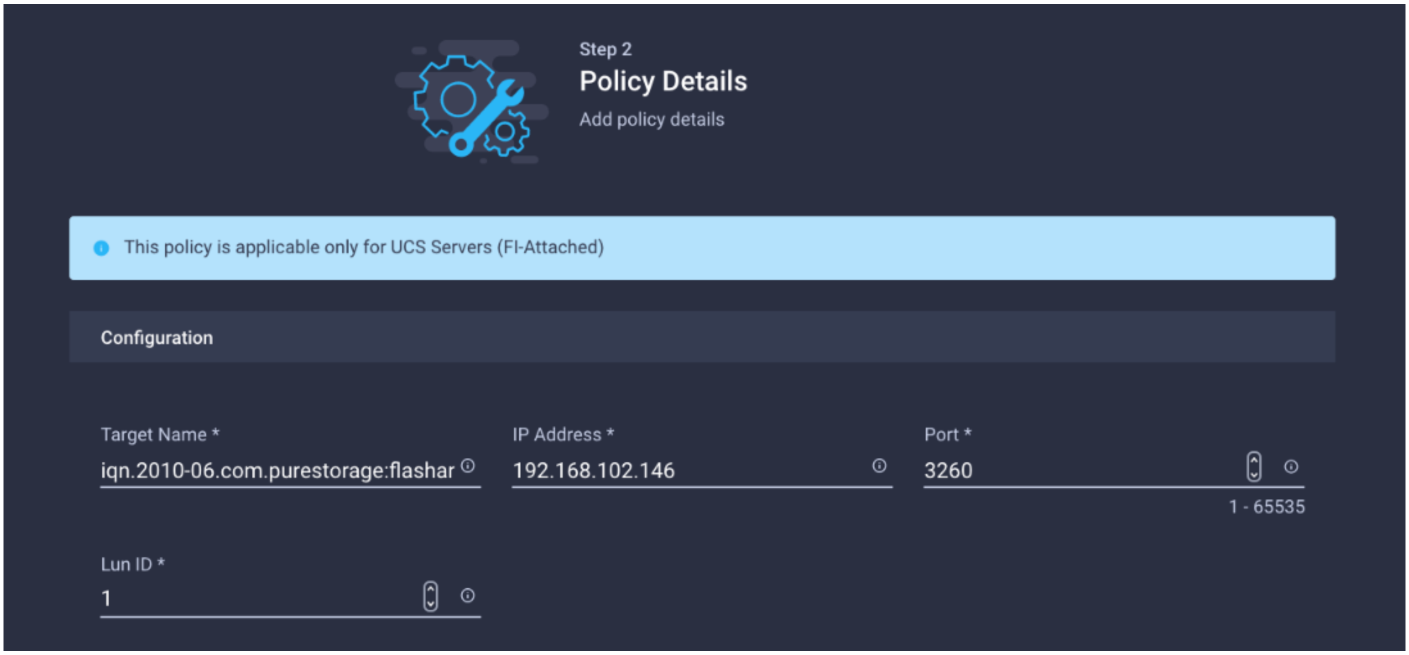
Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Configuration

Target Name * iqn.2010-06.com.purestorage:flashar

IP Address * 192.168.102.146

Port * 3260

1 - 65535

Lun ID * 1

9. Click Create.

10. Click Select Policy under Secondary Target and then, in the pane on the right, click Create New.

11. Verify correct organization is selected from the drop-down list (for example, FSV) and provide a name for the policy (for example, AA19-iSCSI-B-Secondary-Target).

12. Click Next.

13. Provide the Target Name captured from Pure FlashArray, IP Address of ct0.eth5, Port 3260 and Lun ID of 1

14. Click Create.

15. Click Select Policy under iSCSI Adapter and then, in the pane on the right, select the previously configured adapter policy AA19-iSCSI-AdapterPol).

16. Scroll down to Initiator IP Source and make sure Pool is selected.

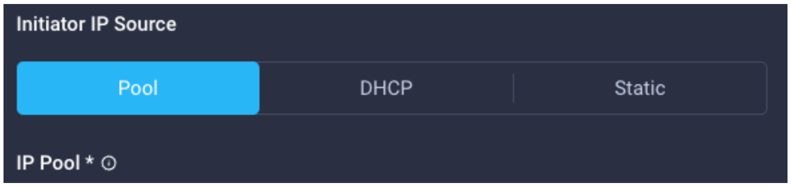
Initiator IP Source

Pool (selected) | DHCP | Static

IP Pool *

17. Click Select Pool under IP Pool and then, in the pane on the right, click Create New.

18. Verify correct organization is selected from the drop-down list (for example, AA19) and provide a name for the pool (for example, AA19-iSCSI-B-Pool).

19. Click Next.

20. Make sure Configure IPv4 Pool is selected. Enter the IP pool information for iSCSI-B subnet.

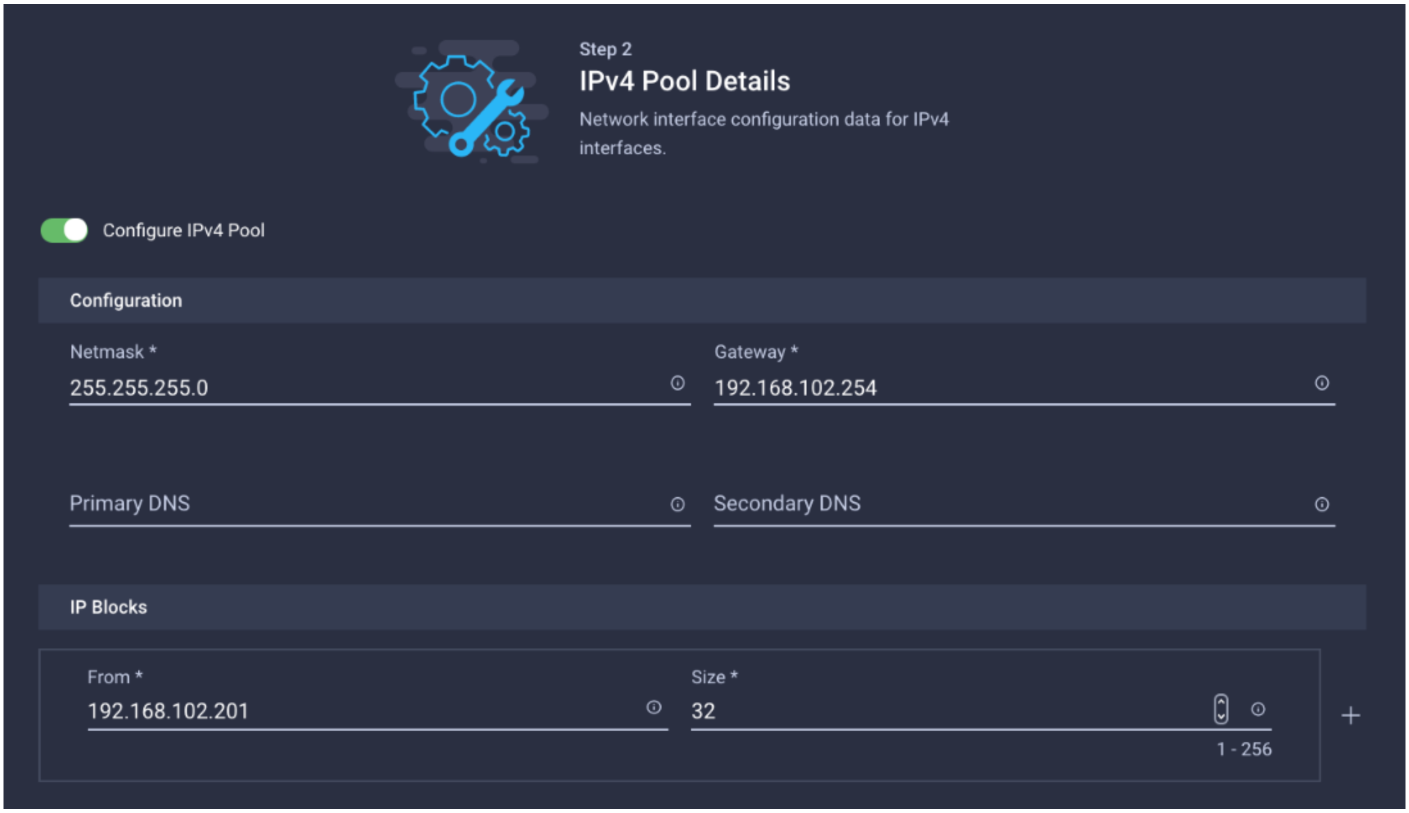

Step 2

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Configuration

Netmask * 255.255.255.0

Gateway * 192.168.102.254

Primary DNS

Secondary DNS

IP Blocks

From * 192.168.102.201

Size * 32

1 - 256

Note: Since the iSCSI network is not routable and all the VMkernel ports and LIFs are layer-2 adjacent, there is no need to define a gateway or DNS.

21. Click Next.

22. Disable Configure IPv6 Pool.

23. Click Create.

24. Verify all the policies and pools are correctly mapped for the iSCSI-B policy.

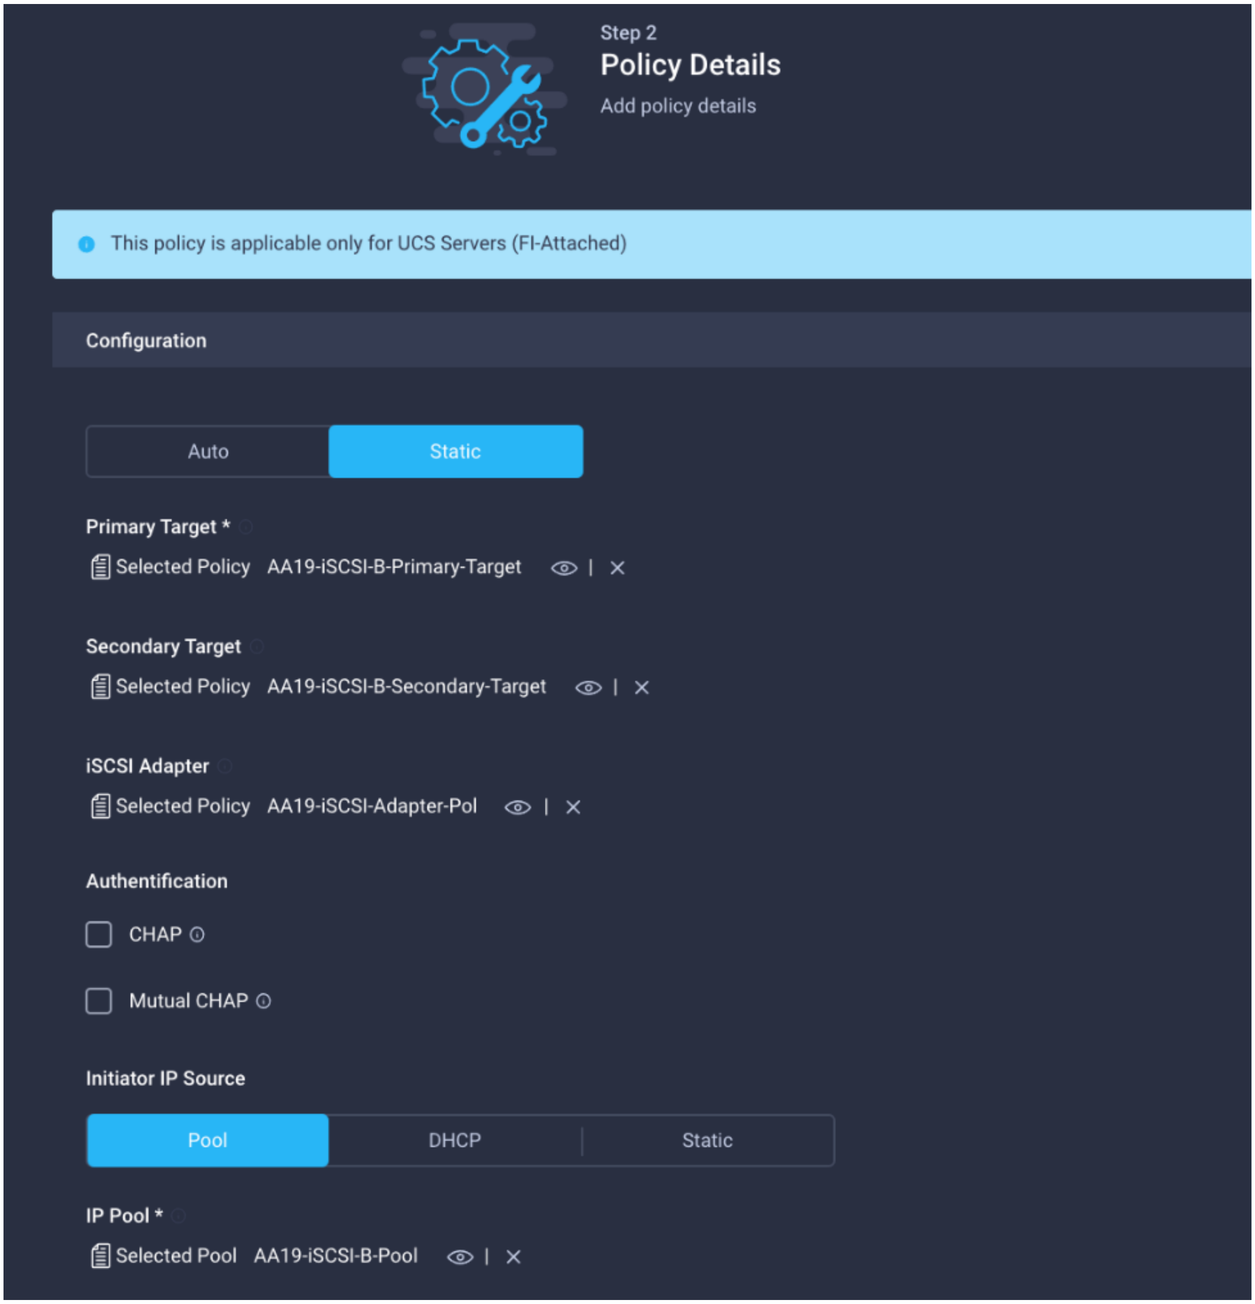

Step 2

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Configuration

Auto | Static (selected)

Primary Target *

Selected Policy AA19-iSCSI-B-Primary-Target

Secondary Target

Selected Policy AA19-iSCSI-B-Secondary-Target

iSCSI Adapter

Selected Policy AA19-iSCSI-Adapter-Pol

Authentification

□ CHAP

□ Mutual CHAP

Initiator IP Source

Pool (selected) | DHCP | Static

IP Pool *

Selected Pool AA19-iSCSI-B-Pool

25. Click Create.

26. Click Create to finish creating the vNIC.

27. Jump back to step 10 Add vNIC and repeat vNIC creation for all six vNICs.

28. Verify all six vNICs were successfully created.

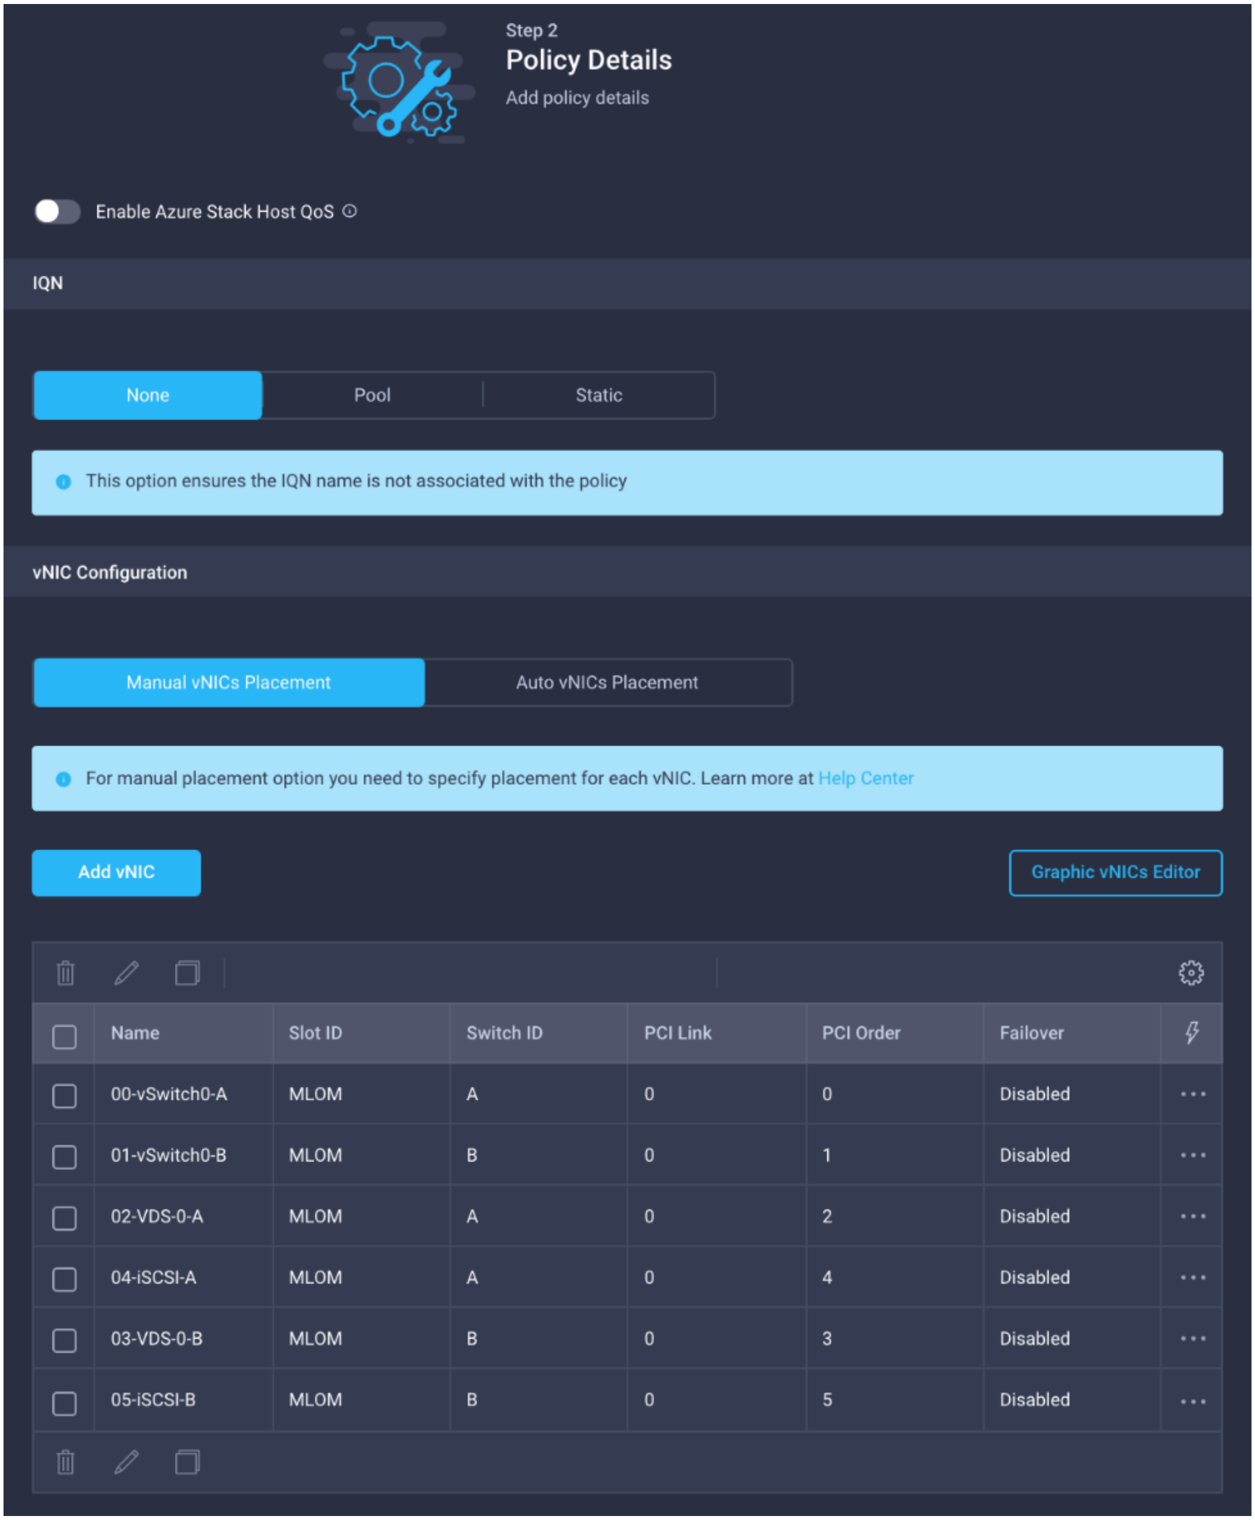

Step 2

Policy Details

Add policy details

Enable Azure Stack Host QoS

IQN

[✓ None] Pool Static

This option ensures the IQN name is not associated with the policy

vNIC Configuration

[✓ Manual vNICs Placement] Auto vNICs Placement

For manual placement option you need to specify placement for each vNIC. Learn more at Help Center

□	Name	Slot ID	Switch ID	PCI Link	PCI Order	Failover
□	00-vSwitch0-A	MLOM	A	0	0	Disabled
□	01-vSwitch0-B	MLOM	B	0	1	Disabled
□	02-VDS-0-A	MLOM	A	0	2	Disabled
□	04-iSCSI-A	MLOM	A	0	4	Disabled
□	03-VDS-0-B	MLOM	B	0	3	Disabled
□	05-iSCSI-B	MLOM	B	0	5	Disabled

29. Click Create to finish creating the LAN Connectivity policy for iSCSI hosts.

30. Click Next.

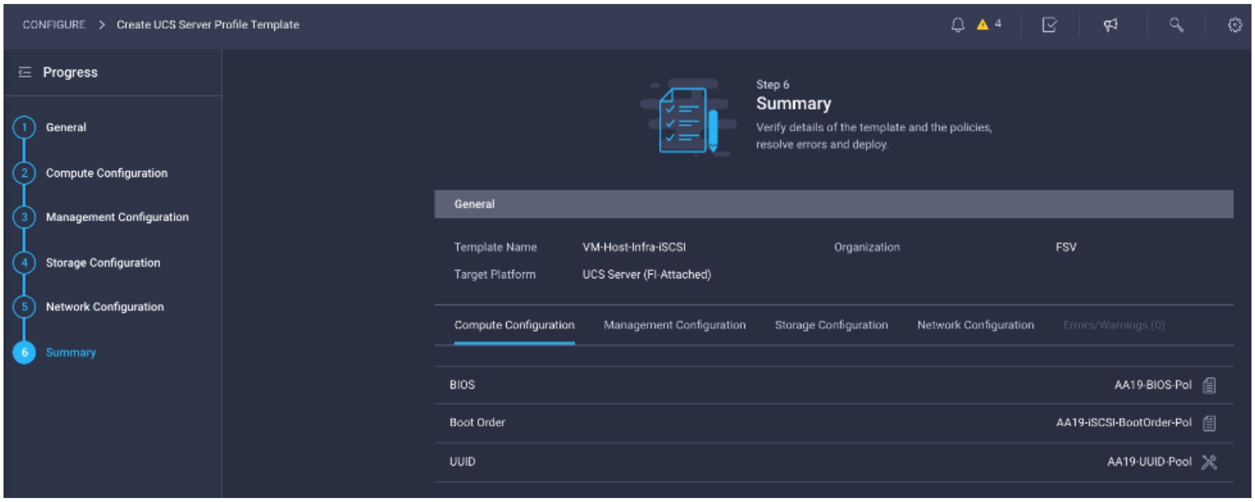

Step 6
Summary
Verify details of the template and the policies, resolve errors and deploy.

General

Template Name	VM-Host-Infra-iSCSI	Organization	FSV
Target Platform	UCS Server (FI-Attached)		

BIOS	AA19-BIOS-Pol
Boot Order	AA19-iSCSI-BootOrder-Pol
UUID	AA19-UUID-Pool

Derive Server Profile

To derive one or many server profiles from the configured template, follow these steps:

1. From the Server profile template Summary screen, click Derive Profiles.

Note: This action can also be performed later by navigating to **Templates**, clicking **"…"** next to the template name and selecting **Derive Profiles**.

2. Under the Server Assignment, select Assign Now and pick Cisco UCS X210c M6 servers. Customers can select one or more servers depending on the number of profiles to be deployed.

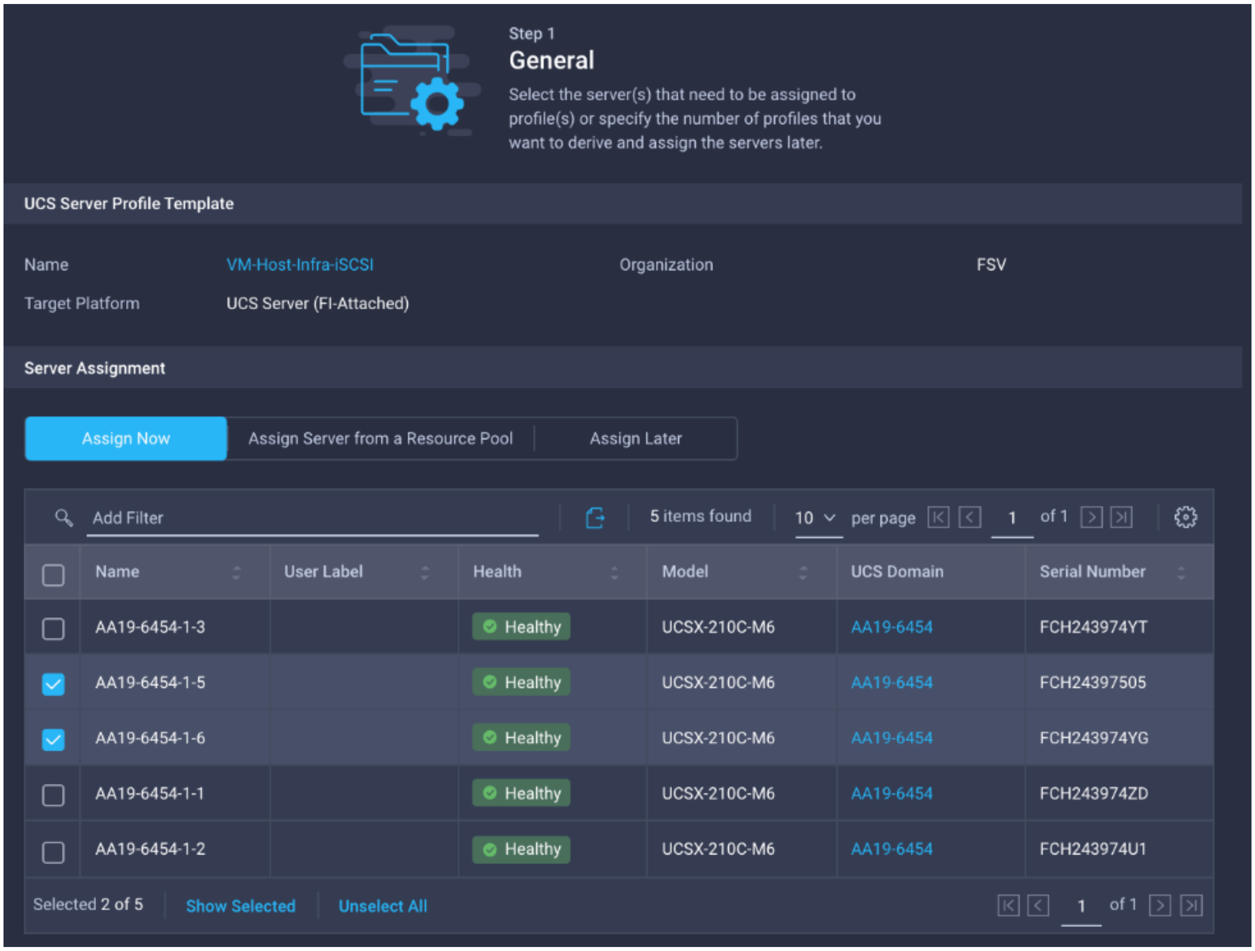

Select the server(s) that need to be assigned to profile(s) or specify the number of profiles that you want to derive and assign the servers later.

UCS Server Profile Template

Name	VM-Host-Infra-iSCSI	Organization	FSV
Target Platform	UCS Server (FI-Attached)		

Server Assignment

□	Name	User Label	Health	Model	UCS Domain	Serial Number
□	AA19-6454-1-3		Healthy	UCSX-210C-M6	AA19-6454	FCH243974YT
☑	AA19-6454-1-5		Healthy	UCSX-210C-M6	AA19-6454	FCH24397505
☑	AA19-6454-1-6		Healthy	UCSX-210C-M6	AA19-6454	FCH243974YG
□	AA19-6454-1-1		Healthy	UCSX-210C-M6	AA19-6454	FCH243974ZD
□	AA19-6454-1-2		Healthy	UCSX-210C-M6	AA19-6454	FCH243974U1

Selected 2 of 5

Note: The server profile template and policies in this document apply to both Cisco UCS X210x M6 and Cisco UCS B200 M6 servers.

3. Click Next.

4. Intersight will fill in default information for the number of servers selected (2 in this case), modify the names if needed.

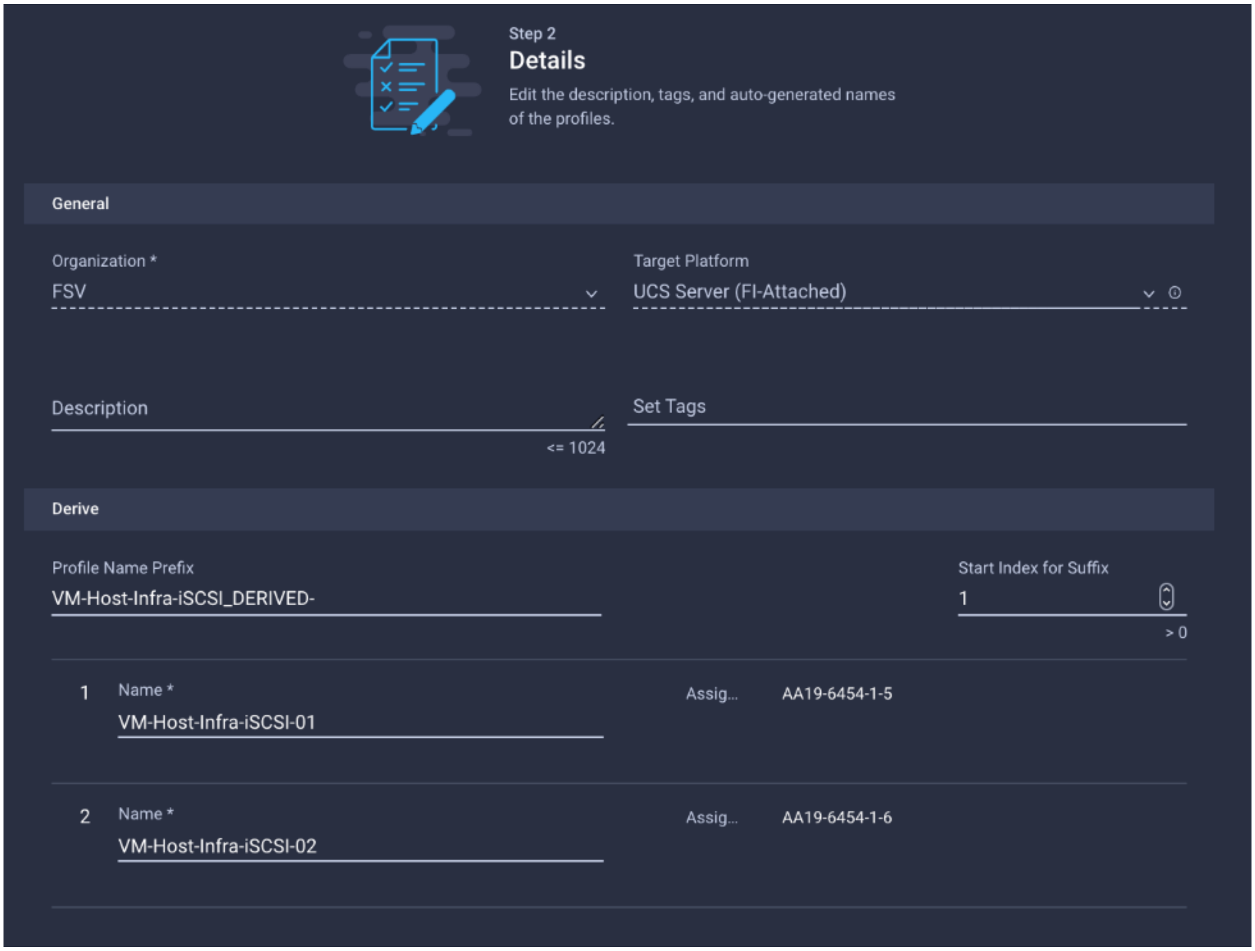

Step 2

Details

Edit the description, tags, and auto-generated names of the profiles.

General

Organization *: FSV

Target Platform: UCS Server (FI-Attached)

Description

<= 1024

Set Tags

Derive

Profile Name Prefix: VM-Host-Infra-iSCSI_DERIVED-

Start Index for Suffix: 1

> 0

	Name *	Assig...	
1	VM-Host-Infra-iSCSI-01		AA19-6454-1-5
2	VM-Host-Infra-iSCSI-02		AA19-6454-1-6

5. Click Next.

6. Verify the information and click Derive to create the Server Profiles.

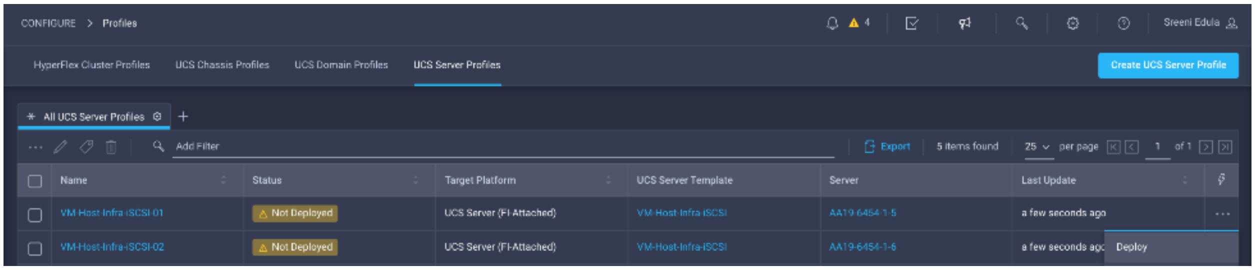

Name	Status	Target Platform	UCS Server Template	Server	Last Update
VM-Host-Infra-iSCSI-01	Not Deployed	UCS Server (FI-Attached)	VM-Host-Infra-iSCSI	AA19-6454-1-5	a few seconds ago
VM-Host-Infra-iSCSI-02	Not Deployed	UCS Server (FI-Attached)	VM-Host-Infra-iSCSI	AA19-6454-1-6	a few seconds ago

7. Click ... on the right-hand side of the derived server profiles and click deploy.

8. When the Server Profiles are deployed successfully, they will appear under the Server Profiles with the status of OK.

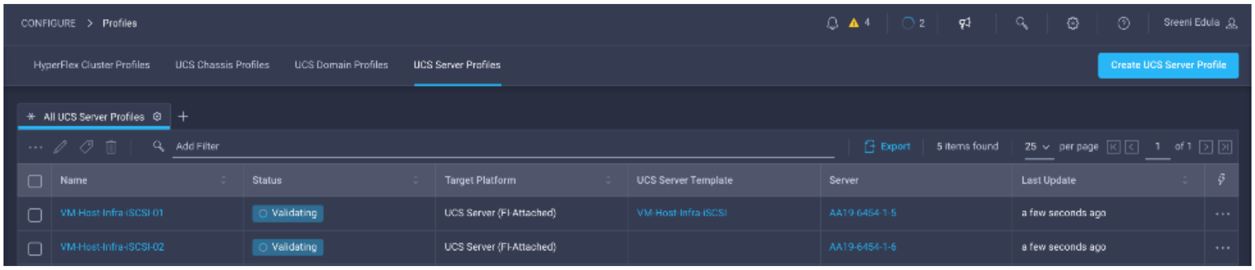

	Name	Status	Target Platform	UCS Server Template	Server	Last Update	
□	VM-Host-Infra-iSCSI-01	○ Validating	UCS Server (FI-Attached)	VM-Host-Infra-iSCSI	AA19-6454-1-5	a few seconds ago	...
□	VM-Host-Infra-iSCSI-02	○ Validating	UCS Server (FI-Attached)		AA19-6454-1-6	a few seconds ago	...

FlashArray Storage Deployment

The Pure Storage FlashArray//X is accessible to the FlashStack, but no storage has been deployed at this point. The storage to be deployed will include:

- ESXi iSCSI Boot LUNs
- VMFS Datastores
- vVOL Data Stores

The iSCSI Boot LUNs will need to be setup from the Pure Storage Web Portal, and the VMFS datastores can be directly provisioned from the vSphere Web Client after the Pure Storage vSphere Web Client Plugin has later been registered with the vCenter.

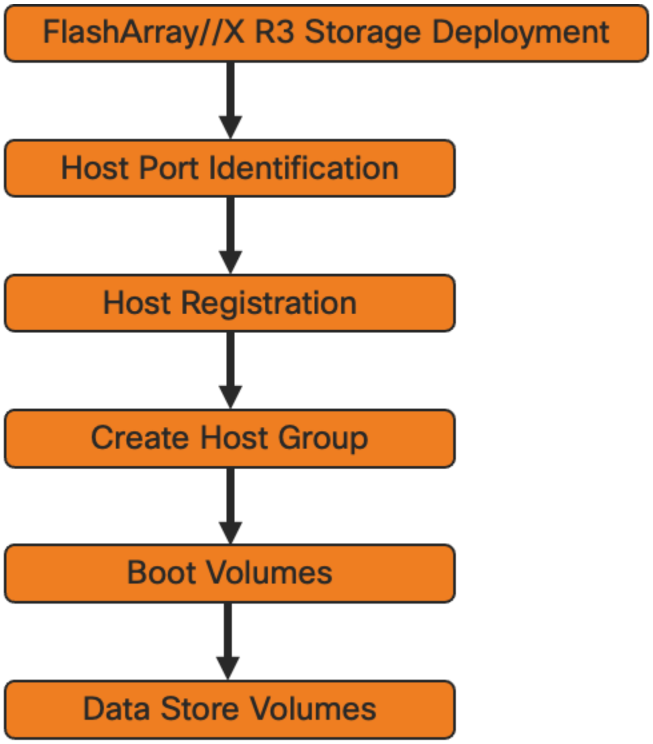

Host Port Identification

iSCSI Boot LUNs will be mapped by the FlashArray//X R3 using the assigned Initiator IQN to the provisioned service profiles. This information can be found within the service profile located within the iSCSI vNIC tab:

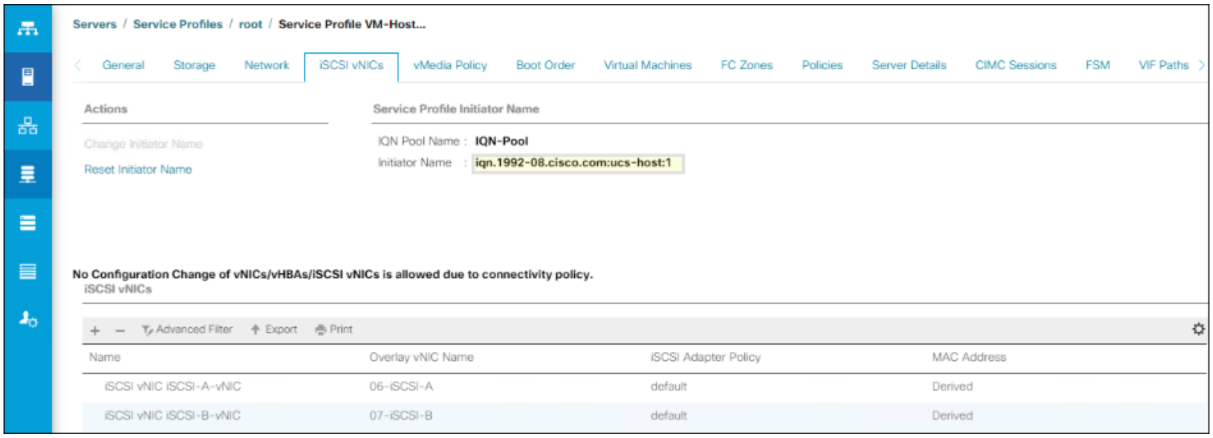

IQN Pool Name : IQN-Pool
Initiator Name : iqn.1992-08.cisco.com:ucs-host:1

No Configuration Change of vNICs/vHBAs/iSCSI vNICs is allowed due to connectivity policy.

iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-A-vNIC	06-iSCSI-A	default	Derived
iSCSI vNIC iSCSI-B-vNIC	07-iSCSI-B	default	Derived

Host Registration

To register the Host, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.
2. Select the + icon in the Hosts Panel.
3. After clicking the Create Host (+) option, a pop-up will appear to create an individual host entry on the FlashArray.

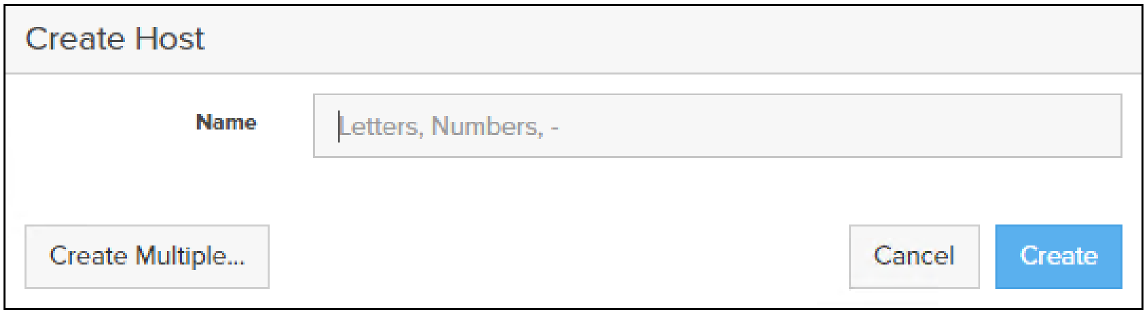

4. To create more than one host entry, click the Create Multiple... option, filling in the Name, Start Number, Count, Personality as ESXi and Number of Digits, with a “#” appearing in the name where an iterating number will appear:

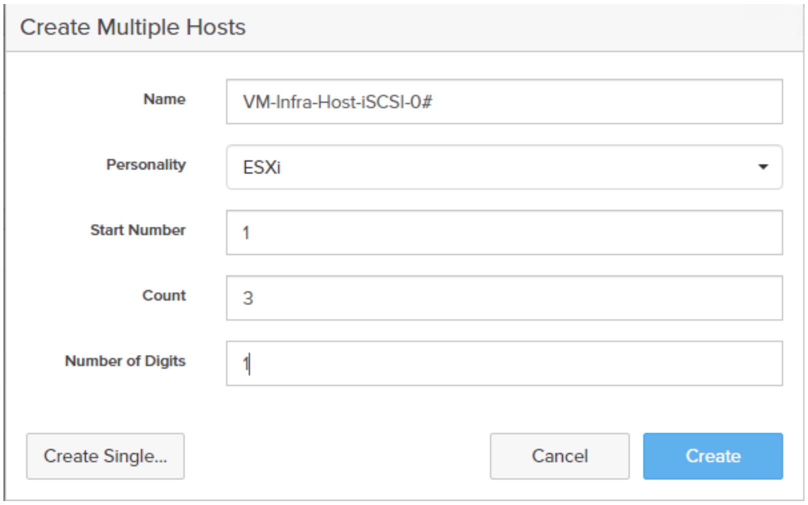

Name: VM-Infra-Host-iSCSI-0#
Personality: ESXi
Start Number: 1
Count: 3
Number of Digits: 1

5. Click Create to add the hosts.

6. For each host created, select the host.

7. In the Host view, select 'Configure IQNs…' from the Host Ports menu.

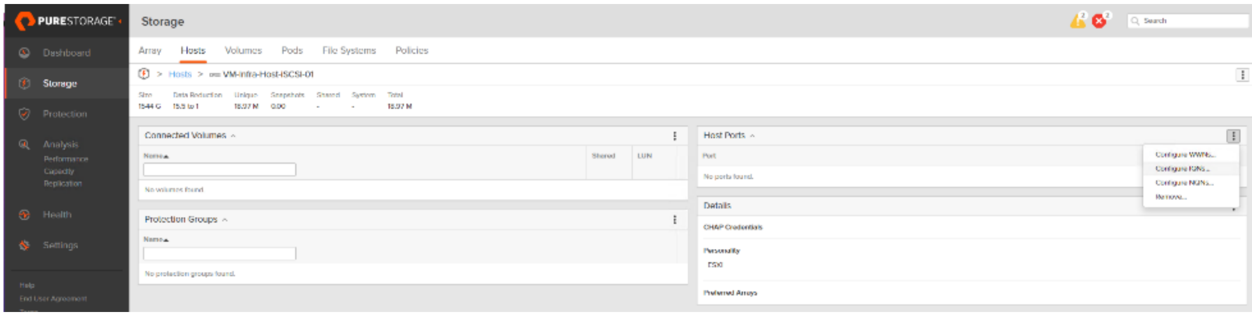

8. A pop-up will appear for Configure iSCSI IQNs for Host <host being configured>. Within this pop-up, enter the IQN Initiator Name found within the service profile for the host being configured:

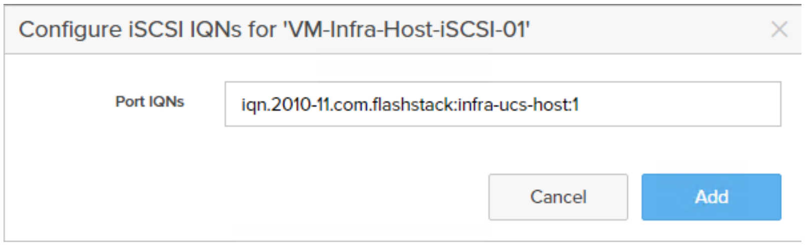
Configure iSCSI IQNs for 'VM-Infra-Host-iSCSI-01'

Port IQNs: iqn.2010-11.com.flashstack:infra-ucs-host:1

9. After entering the IQN, click Add to add the Host Ports.

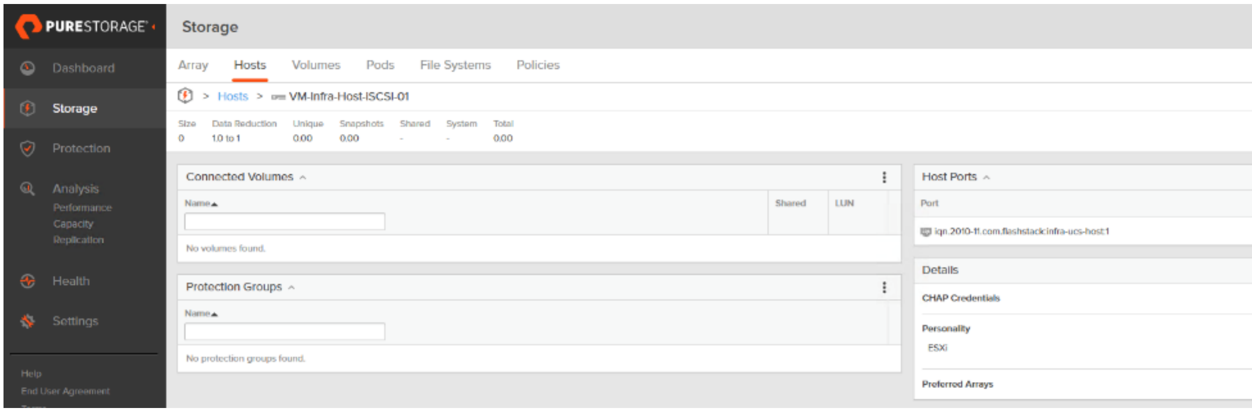

Host Ports – Port: iqn.2010-11.com.flashstack:infra-ucs-host:1

10. Repeat steps 1-9 for each host created.

Create Host Group

Host Groups allow the Administrator to map Volumes to a group of hosts at once with the same LUN ID. To create a Host Group, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.

2. Select the + icon in the Host Groups Panel.

3. A pop-up will appear to create a host group on the FlashArray.

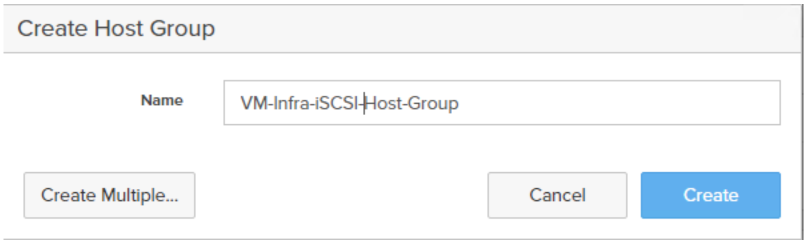

Create Host Group

Name: VM-Infra-iSCSI-Host-Group

4. Provide a name for the group and click Create.

5. Select the group in the Host Groups Panel.

6. In the Host Group view, select ‘Add...’ from the Member Hosts menu.

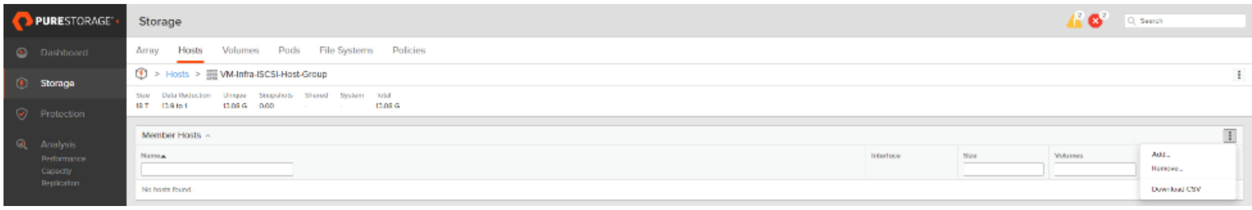

Hosts > VM-Infra-iSCSI-Host-Group

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
18 T	13.9 to 1	13.08 G	0.00	-	-	13.08 G

Member Hosts

Name▲	Interface	Size	Volumes
No hosts found.			

7. Select the host to be part of the host group.

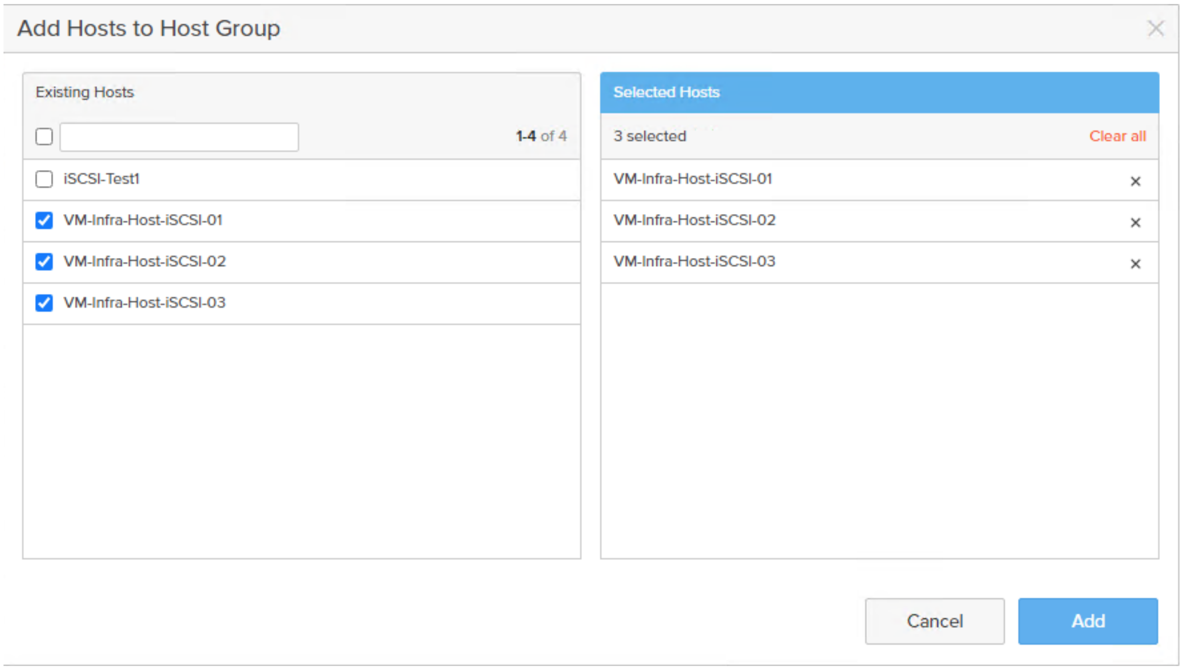

Add Hosts to Host Group

Existing Hosts — 1-4 of 4

- ☐ iSCSI-Test1
- ☑ VM-Infra-Host-iSCSI-01
- ☑ VM-Infra-Host-iSCSI-02
- ☑ VM-Infra-Host-iSCSI-03

Selected Hosts — 3 selected — Clear all

- VM-Infra-Host-iSCSI-01 ×
- VM-Infra-Host-iSCSI-02 ×
- VM-Infra-Host-iSCSI-03 ×

8. Click Add.

Private Boot Volumes for each ESXi Host

To create private boot volumes for each ESXi Host, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.

2. Select the + icon in the Volumes Panel.

3. A pop-up will appear to create a volume on the FlashArray.

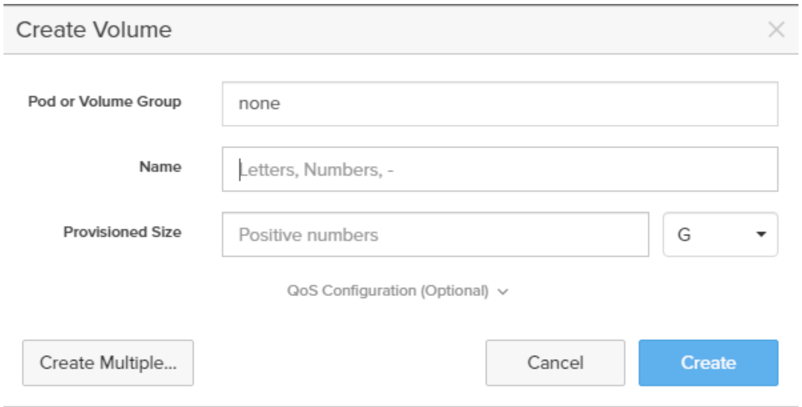

Create Volume

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

4. To create more than one volume, click the Create Multiple... option, filling in the Name, Provisioned Size, Staring Number, Count, and Number of Digits, with a "#" appearing in the name where an iterating number will appear.

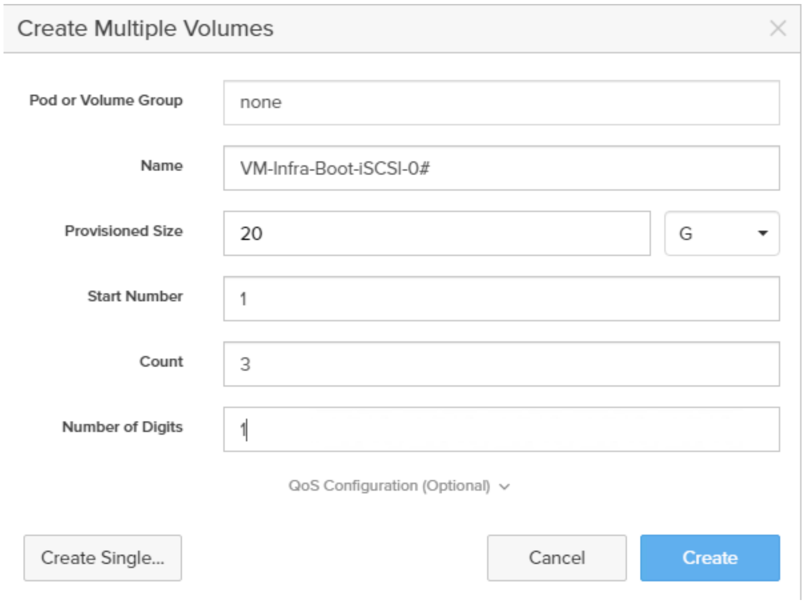

Create Multiple Volumes

Pod or Volume Group: none

Name: VM-Infra-Boot-iSCSI-0#

Provisioned Size: 20 G

Start Number: 1

Count: 3

Number of Digits: 1

5. Click Create to provision the volumes to be used as iSCSI boot LUNs.

6. Go back to the Hosts section under the Storage tab. Click one of the hosts and select the gear icon drop-down within the Connected Volumes tab within that host.

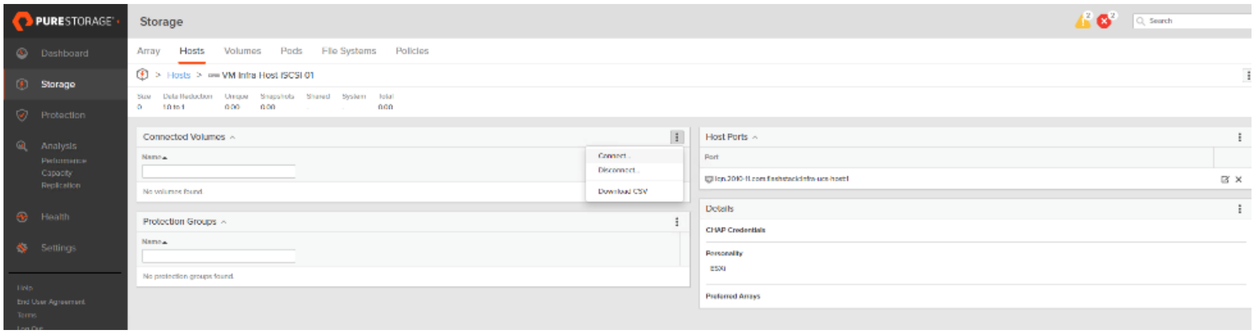

7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

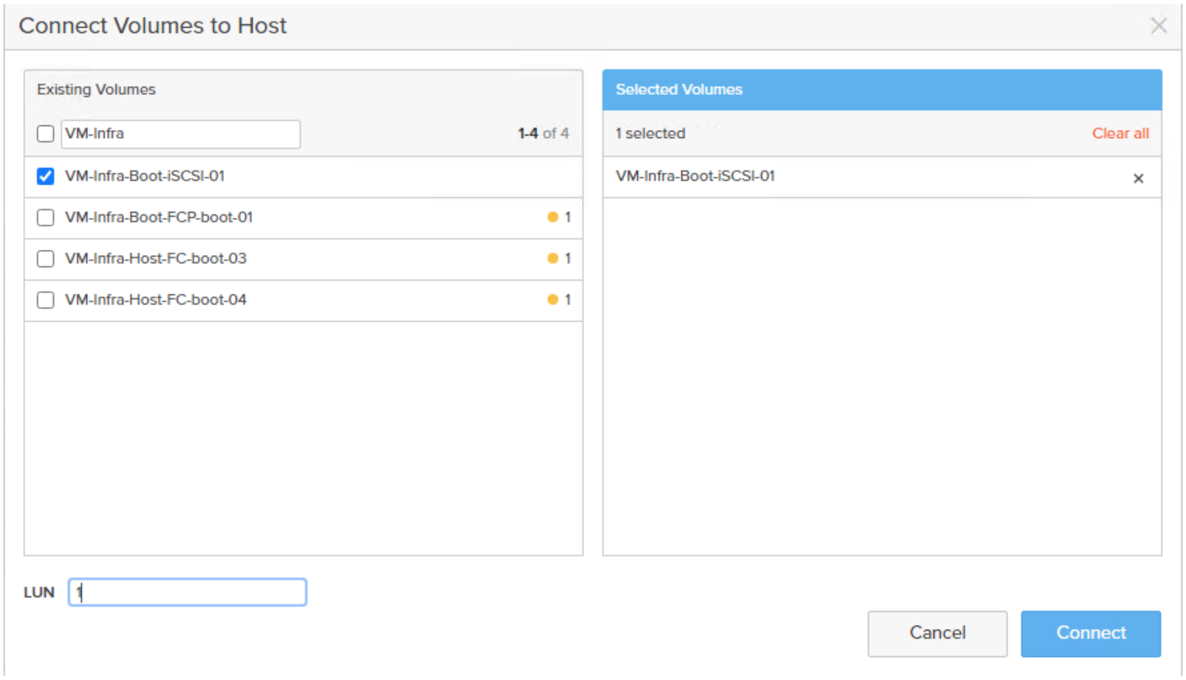

Note: LUN ID 1 should be used for the boot .

8. Select the volume that has been provisioned for the host, set the LUN ID for the volume, click the + next to the volume, and select Confirm to proceed. Repeat the steps for connecting volumes for each of the host/volume pairs configured.

Create Infra Datastores

To create datastore volumes for the ESXi Cluster, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.

2. Select the + icon in the Volumes Panel.

3. A pop-up will appear to create a volume on the FlashArray.

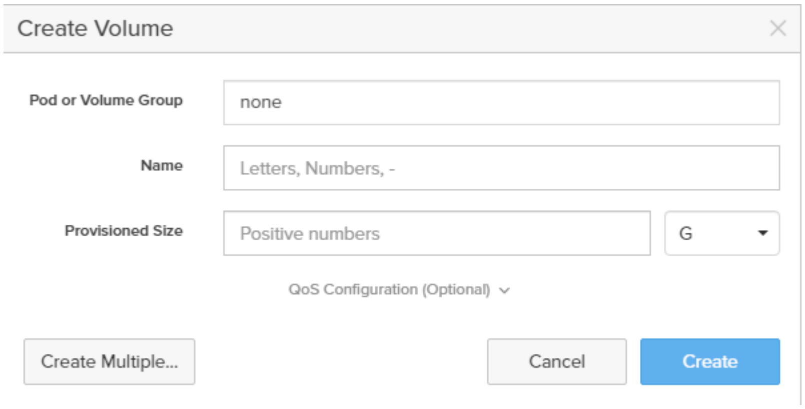

4. Fill in the Name and Provisioned Size.

5. Click Create to provision the volumes to be used as Infra datastore LUN.

6. Go back to the Hosts section under the Storage tab. Click ESXi cluster host group created earlier and select the gear icon pull-down within the Connected Volumes tab within that host group.

7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

8. Select the Infra datastore volume that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.

VMware vSphere Configuration

Set Up VMkernel Ports and Virtual Switch on ESXi Host VM-Host-Infra-iSCSI-01

To add the iSCSI networking configuration on the first ESXi host, follow the steps at the end of section Set Up VMkernel Ports and Virtual Switch. In this section, a single iSCSI Boot vSwitch is configured with two uplinks, one to UCS fabric A and the other to fabric B. The first VMkernel port will be mapped only to the fabric A uplink and the second one will be mapped to the fabric B uplink.

To setup VMkernel ports and virtual switches on ESXi hosts on VM-Host-Infra-iSCSI-01, follow these steps:

1. From the Host Client Navigator, click Networking.

2. In the center pane, choose the Virtual switches tab.

3. Highlight the iScsiBootvSwitch line.

4. Choose Edit settings.

5. Change the MTU to 9000.

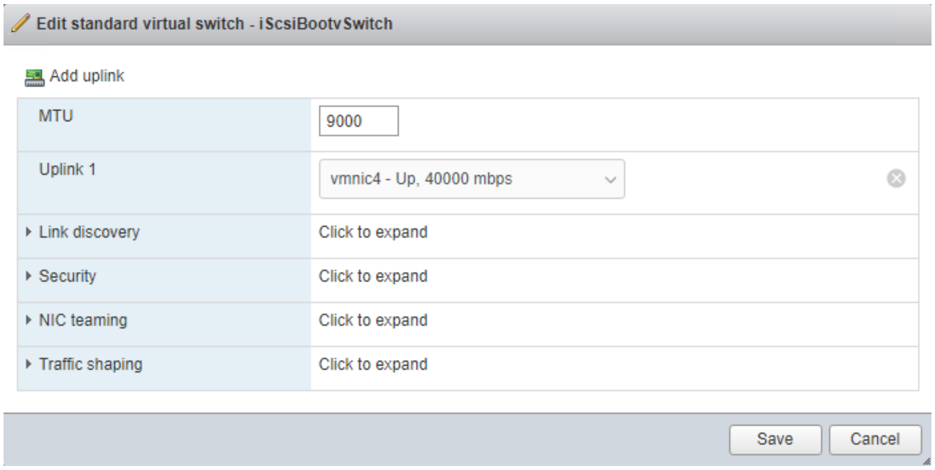

Edit standard virtual switch - iScsiBootvSwitch

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

6. Click Save to save the changes to iScsiBootvSwitch.

7. Click vmk1 entry.

8. Click Edit Settings.

9. From Port properties update the MTU value to 9000.

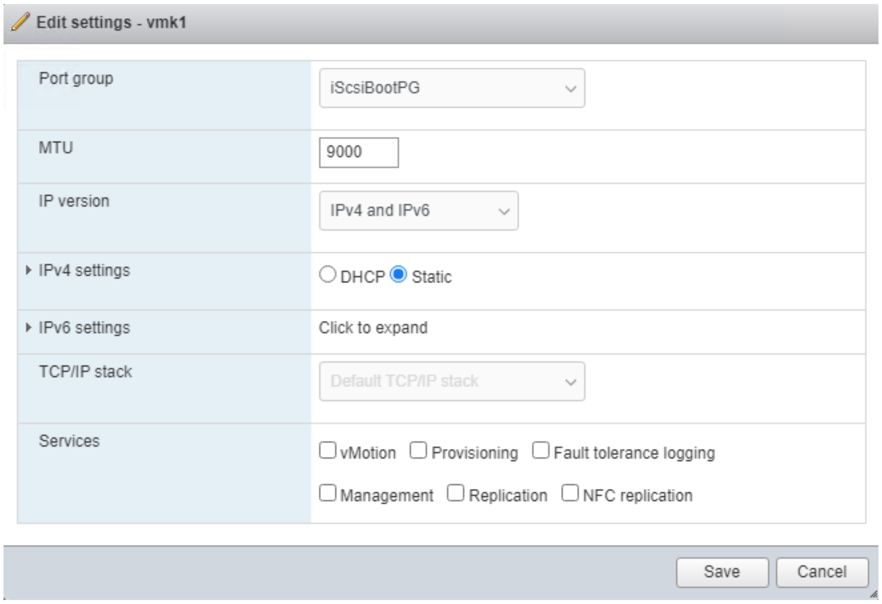

Edit settings - vmk1

Port group	iScsiBootPG
MTU	9000
IP version	IPv4 and IPv6
IPv4 settings	○ DHCP ● Static
IPv6 settings	Click to expand
TCP/IP stack	Default TCP/IP stack
Services	□ vMotion □ Provisioning □ Fault tolerance logging □ Management □ Replication □ NFC replication

10. Click the IPv4 Settings.

11. Change the IPv4 settings from the Cisco UCS Manager iSCSI-A-Pool assigned IP to one that is not in the IP block.

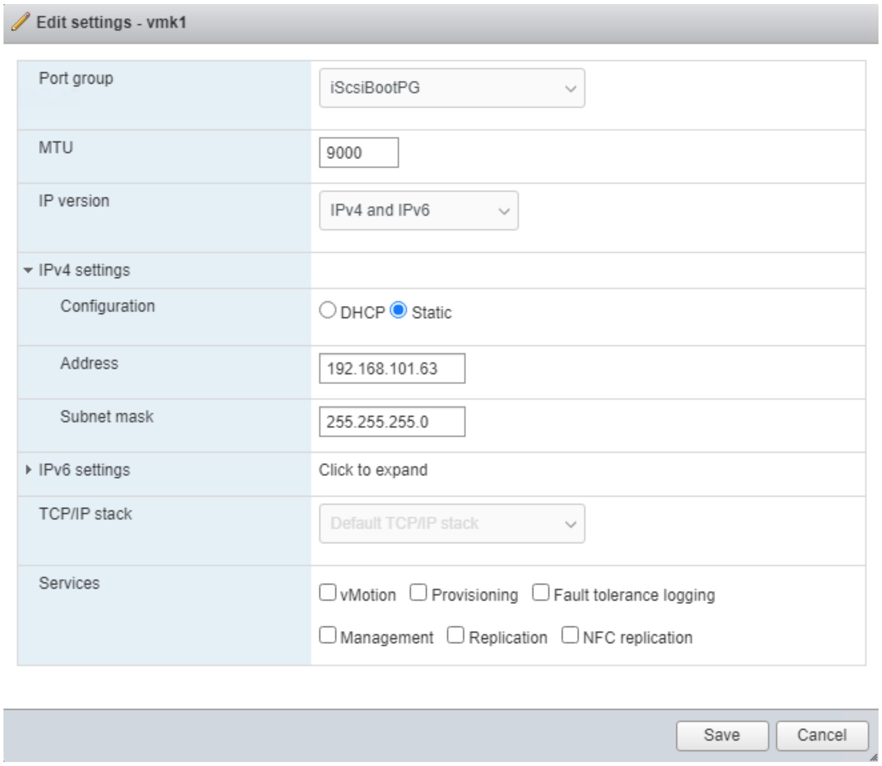

Edit settings - vmk1

Port group	iScsiBootPG
MTU	9000
IP version	IPv4 and IPv6
IPv4 settings	
Configuration	○ DHCP ● Static
Address	192.168.101.63
Subnet mask	255.255.255.0
IPv6 settings	Click to expand
TCP/IP stack	Default TCP/IP stack
Services	□ vMotion □ Provisioning □ Fault tolerance logging □ Management □ Replication □ NFC replication

12. Click OK to apply the changes.

Configure iSCSI B vSwitch and VMkernel

To configure the iSCSI vSwitch and VMkernel, follow these steps:

1. From the Host Client Navigator, click Networking.

2. In the center pane, choose the Virtual switches tab.

3. Click add standard virtual switch.

4. Name the switch iScsiBootvSwitch-B.

5. Change the MTU to 9000.

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

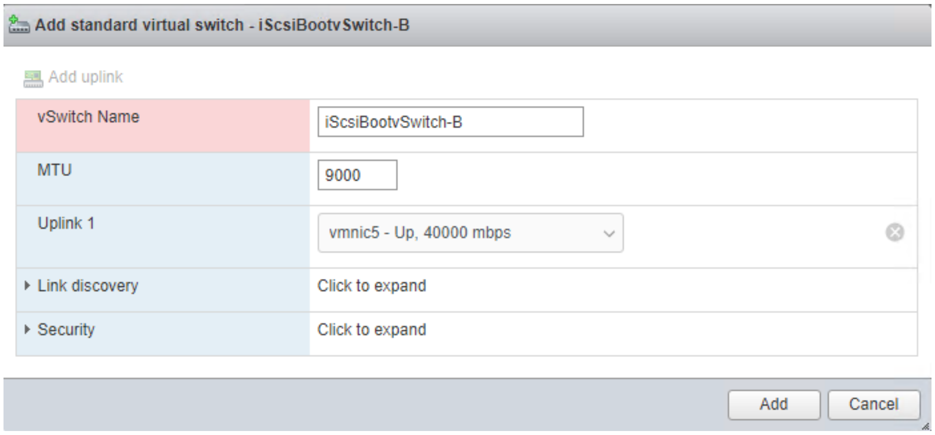

Add standard virtual switch - iScsiBootvSwitch-B

Add uplink

vSwitch Name	iScsiBootvSwitch-B
MTU	9000
Uplink 1	vmnic5 - Up, 40000 mbps
▸ Link discovery	Click to expand
▸ Security	Click to expand

7. Choose Add to add iScsiBootvSwitch-B.

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

9. Choose Add VMkernel NIC.

10. For New port group, enter iScsiBootPG-B.

11. For Virtual switch, use the pull-down to choose vSwitch1.

12. Change the MTU to 9000.

13. For IPv4 settings, choose Static.

14. 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.

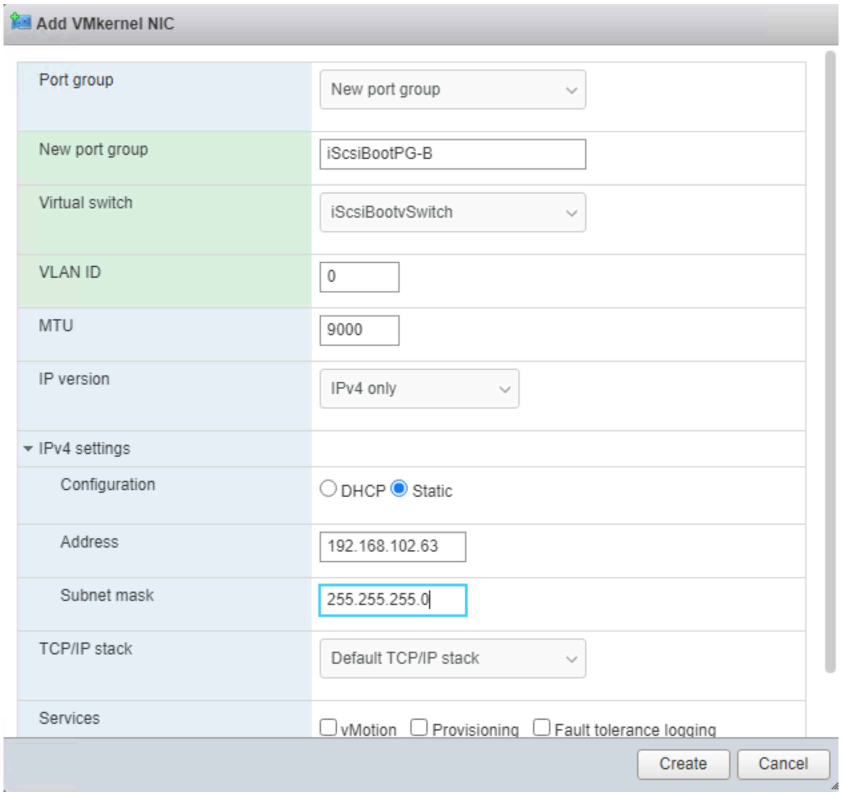

Add VMkernel NIC

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	iScsiBootvSwitch
VLAN ID	0
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	○ DHCP ● Static
Address	192.168.102.63
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	□ vMotion □ Provisioning □ Fault tolerance logging

15. Click Create to complete creating the VMkernel NIC.

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

17. Highlight the iScsiBootPG line.

18. Choose Edit settings.

19. Change the Name to iScsiBootPG-A.

20. Click Save to complete editing the port group name.

21. Click Storage, then in the center pane choose the Adapters tab.

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

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

24. Choose to add address and enter the IP address of ct0.eth4 from Pure FlashArray//X R3. Press Return.

25. Repeat above steps to add the IP addresses for ct0.eth5, ct1.eth4 and ct1.eth5.

26. Click Save configuration.

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

28. Verify that four static targets and four dynamic targets are listed for the host.

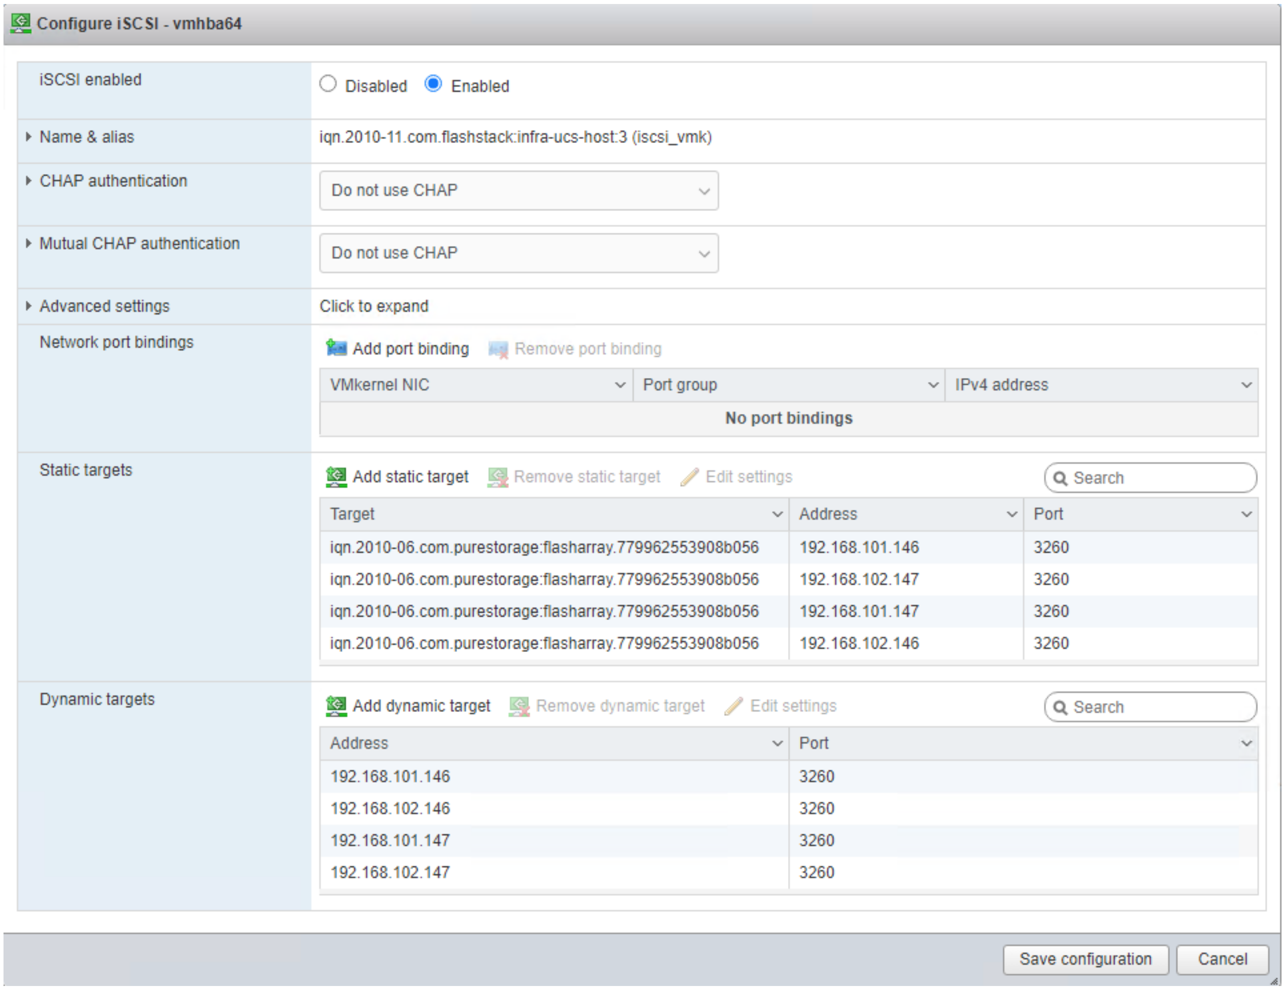
Configure iSCSI - vmhba64

iSCSI enabled	○ Disabled ● Enabled
Name & alias	iqn.2010-11.com.flashstack:infra-ucs-host:3 (iscsi_vmk)
CHAP authentication	Do not use CHAP
Mutual CHAP authentication	Do not use CHAP
Advanced settings	Click to expand
Network port bindings	No port bindings

Static targets

Target	Address	Port
iqn.2010-06.com.purestorage:flasharray.779962553908b056	192.168.101.146	3260
iqn.2010-06.com.purestorage:flasharray.779962553908b056	192.168.102.147	3260
iqn.2010-06.com.purestorage:flasharray.779962553908b056	192.168.101.147	3260
iqn.2010-06.com.purestorage:flasharray.779962553908b056	192.168.102.146	3260

Dynamic targets

Address	Port
192.168.101.146	3260
192.168.102.146	3260
192.168.101.147	3260
192.168.102.147	3260

29. Click Cancel to close the window.

Note: If the host shows an alarm stating that connectivity with the boot disk was lost, place the host in Maintenance Mode and reboot the host.

Add iSCSI Configuration to a VMware ESXi Host Added in vCenter

This section details the steps to add iSCSI configuration to an ESXi host added and configured in vCenter. This section assumes the host has been added to vCenter and the basic networking completed, and the time configuration and swap files added.

To add an iSCSI configuration to an ESXi host, follow these steps:

1. In the vSphere HTML5 Client, under Hosts and Clusters, choose the ESXi host.

2. In the center pane, click Configure. In the list under Networking, select Virtual switches.

3. In the center pane, expand iScsiBootvSwitch. Click EDIT to edit settings for the vSwitch.

4. Change the MTU to 9000 and click OK.

5. Choose ... > Edit Settings to the right of iScsiBootPG. Change the Network label to iScsiBootPG-A and click OK.

6. Choose ... > Edit Settings to the right of the VMkernel Port IP address. Change the MTU to 9000.

7. Click IPv4 settings on the left. Change the IP address to a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.

Note: 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 pure FlashArray storage controller's ct0.eth4and click OK.

20. Repeat steps 1- 19 to add the IPs for ct0.eth5, ct1.eth4, and ct1.eth5.

21. Under Storage Adapters, click Rescan Adapter to rescan the iSCSI Software Adapter.

Under Static Discovery, four static targets are listed.

22. Under Paths, four paths should now be listed with two of the paths having the "Active (I/O)" Status.

FlashStack Backups

Cisco Intersight SaaS Platform

Cisco Intersight SaaS platform maintains customer configurations online. No separate backup was created for the UCSX configuration.

VMware VCSA Backup

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

1. Connect to the VCSA Console here: https://<VCSA IP>:5480 as root.
2. Click Backup in the list to open up the Backup Appliance Dialogue.
3. To the right of the Backup Schedule, click CONFIGURE.
4. Specify the following:
   a. The Backup location with the protocol to use [FTPS, HTTPS, SFTP, FTP, NFS, SMB, HTTP]
   b. The User name and password.
   c. The Number of backups to retain.

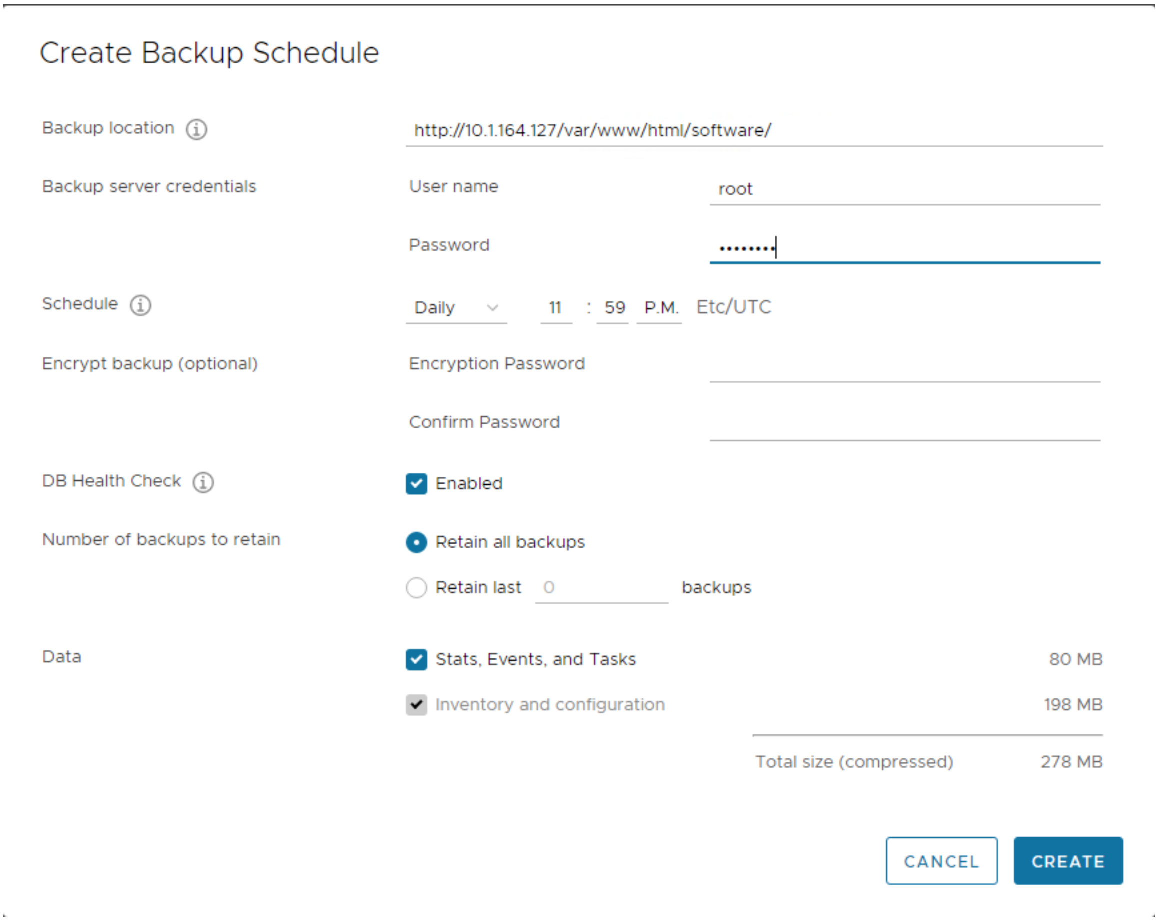

Create Backup Schedule

Backup location ⓘ	http://10.1.164.127/var/www/html/software/	
Backup server credentials	User name	root
	Password	••••••••
Schedule ⓘ	Daily : 11 : 59 P.M. Etc/UTC	
Encrypt backup (optional)	Encryption Password	
	Confirm Password	
DB Health Check ⓘ	☑ Enabled	
Number of backups to retain	● Retain all backups	
	○ Retain last 0 backups	
Data	☑ Stats, Events, and Tasks	80 MB
	☑ Inventory and configuration	198 MB
	Total size (compressed)	278 MB

5. Click CREATE.

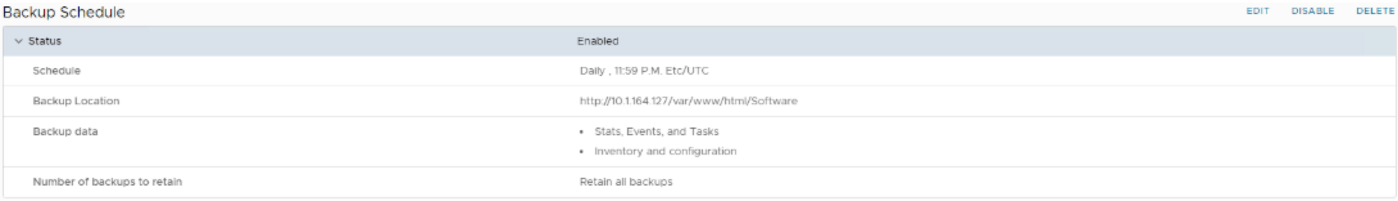

Backup Schedule

Status	Enabled
Schedule	Daily , 11:59 P.M. Etc/UTC
Backup Location	http://10.1.164.127/var/www/html/Software
Backup data	• Stats, Events, and Tasks • Inventory and configuration
Number of backups to retain	Retain all backups

6. The Backup Schedule should now show a Status of Enabled.
7. Restoration can be initiated with the backed-up files using the Restore function of the VCSA 7.0 U2 Installer.

About the Authors

Sreenivasa Edula, Technical Marketing Engineer, Cisco UCS Data Center Solutions Engineering, Cisco Systems, Inc.

Sreeni is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Engineering team focusing on converged and hyper-converged infrastructure solutions, prior to that he worked as a Solutions Architect at EMC Corporation. He has experience in Information Systems with expertise across the Cisco Data Center technology portfolio, including DC architecture design, virtualization, compute, network, storage, and cloud computing.

Joe Houghes, Senior Solutions Architect, Pure Storage, Inc.

Joe is a Senior Solutions Architect in the Portfolio Solutions team within Pure Storage, focused on solutions on the FlashStack platform along with automation and integration. He has experience from over 15 years in Information Technology across various customer/vendor organizations with architecture and operations expertise covering compute, networking, storage, virtualization, business continuity, and disaster recovery, along with cloud computing technologies, plus automation and integration across many applications and vendor platforms.

Acknowledgments

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- John George, Technical Marketing Engineer, Cisco Systems, Inc.
- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Craig Waters, Technical Director, Pure Storage, 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**.