The bridge to possible

Deployment Guide
Cisco Public

FlashStack for Healthcare with Epic on Cisco UCS M7 X-Series, Pure Storage FlashArray//X R3 Series, and VMware vSphere 8.0

Deployment Guide using Cisco UCS X210c M7 Servers with 4th Generation Intel Xeon Scalable Processors running on Cisco Intersight with Pure Storage FlashArray//X R3 Series

Published: July 2023

In partnership with:

About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to: https://www.cisco.com/c/en/us/solutions/design-zone.html

Executive Summary

Cisco Validated Designs (CVDs) consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.

This document details the design of the FlashStack with Epic for VMware vSphere 8.0 Design Guide, which describes a validated Converged Infrastructure (CI) jointly developed by Cisco and Pure Storage.

The solution explains the deployment of a predesigned, best-practice data center architecture with:

- Epic for Healthcare
- VMware vSphere
- Cisco Unified Computing System (Cisco UCS) incorporating the Cisco X-Series modular platform
- Cisco Nexus 9000 family of switches
- Cisco MDS 9000 family of Fibre Channel switches
- Pure Storage FlashArray//X R3 All Flash Array supporting Fibre Channel storage protocol

Additionally, this FlashStack solution is delivered as Infrastructure as Code (IaC) to eliminate error-prone manual tasks, allowing quicker and more consistent solution deployments. The Cisco Intersight cloud platform delivers monitoring, orchestration, workload optimization and lifecycle management capabilities for the FlashStack solution.

Customers interested in understanding the 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, here: Data Center Design Guides - FlashStack Platforms

Solution Overview

This chapter contains the following:

- Introduction
- Audience
- Purpose of this Document
- What’s New in this Release?

Introduction

The Cisco Unified Computing System (Cisco UCS) X-Series is a brand-new modular compute system, configured and managed from the cloud. It is designed to meet the needs of modern applications and to improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco Intersight platform is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support. FlashStack systems deployed to host Epic HyperSpace, InterSystems Iris database, Cogito Clarity analytics and reporting suite, and services servers hosting the Epic application layer provide an integrated platform for a dependable, high-performance infrastructure that can be deployed rapidly.

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 data-center management but also allows the infrastructure to adapt to the unpredictable needs of modern applications as well as traditional workloads. With the Cisco Intersight platform, customers get all the benefits of SaaS delivery and the full lifecycle management of Intersight-connected distributed servers and integrated Pure Storage systems across data centers, remote sites, branch offices, and edge environments.

Audience

The intended audience of this document includes but is not limited to IT architects, sales engineers, field consultants, professional services, IT managers, partner engineering, 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 design guidance for incorporating the Cisco Intersight-managed Cisco UCS X-Series platform within the FlashStack Datacenter infrastructure for deploying Epic Electronic Health Records (EHR). This document introduces various design elements and covers various considerations and best practices for a successful deployment. This document also highlights the design and product requirements for integrating virtualization and storage systems to Cisco Intersight to deliver a true cloud-based integrated approach to infrastructure management.

What’s New in this Release?

This version of the FlashStack design is based on the latest FlashStack Virtual Server Infrastructure and introduces the Cisco UCS X-Series modular platform.

Highlights for this design include:

- Support for Cisco UCS X9508 chassis with Cisco UCS X210c M7 compute nodes
- Support for Pure Storage FlashArray//X R3 with Purity version 6.3.7

- Support for VMware vSphere 8.0
- Support for Cisco Intersight platform to deploy, maintain, and support the FlashStack components
- Support for Cisco Intersight Assist virtual appliance to help connect the Pure Storage FlashArray and VMware vCenter with the Cisco Intersight platform

The use cases include:

- Enterprise Data Center
- Service Provider Data Center
- Large Commercial Data Center

Technology Overview

This chapter contains the following:

- FlashStack
- Cisco Unified Computing System
- Cisco UCS Fabric Interconnect
- Cisco Unified Computing System X-Series
- Cisco UCS Virtual Interface Cards (VICs)
- Cisco Switching
- VMware vSphere 8.0
- Red Hat Ansible
- Cisco Intersight Assist Device Connector for VMware vCenter and Pure Storage FlashArray
- Purity for FlashArray
- Pure1
- InterSystems Corporation
- Epic Systems

Cisco and Pure Storage have partnered to deliver several Cisco Validated Designs. These designs use best-in-class storage, server, and network components to serve as the foundation for virtualized workloads such as Epic, enabling efficient architectural designs that you can deploy quickly and confidently.

FlashStack

The FlashStack architecture was jointly developed by Cisco and Pure Storage. All FlashStack components are integrated, allowing customers can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlashStack is its ability to maintain consistency at scale. Each of the component families shown in Figure 1 (Cisco UCS, Cisco Nexus, Cisco MDS, and Pure Storage FlashArray systems) offers platform and resource options to scale up or scale out the infrastructure while supporting the same features and functions.

Figure 1. FlashStack components

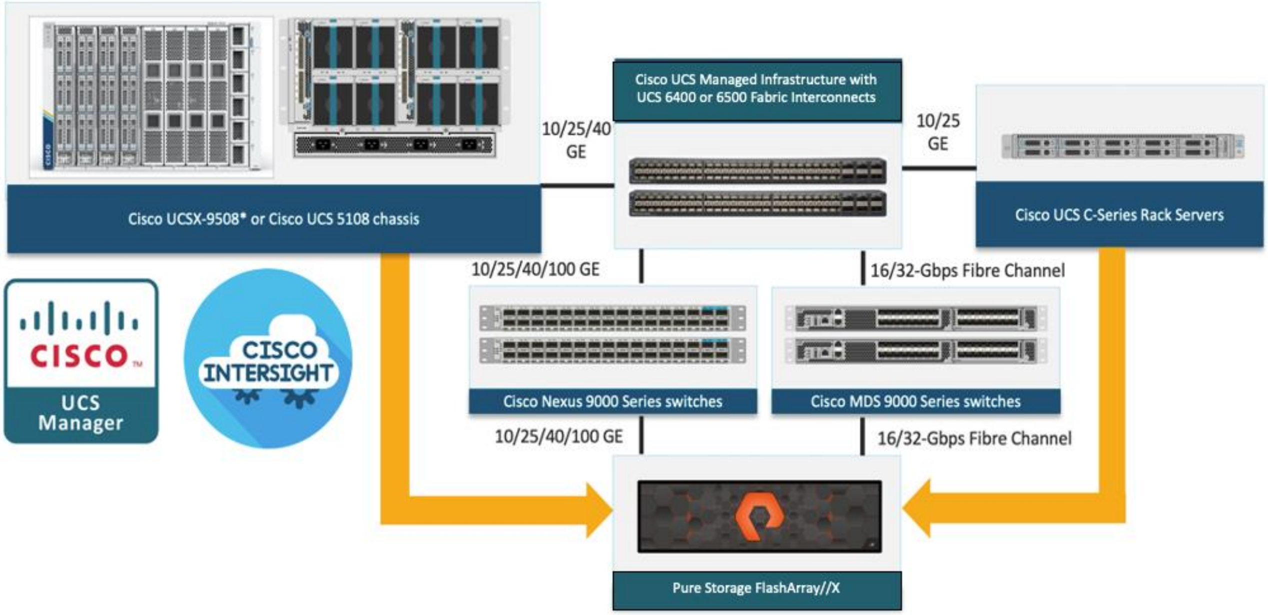
Cisco UCSX-9508* or Cisco UCS 5108 chassis
10/25/40 GE
Cisco UCS Managed Infrastructure with UCS 6400 or 6500 Fabric Interconnects
10/25 GE
Cisco UCS C-Series Rack Servers
10/25/40/100 GE
16/32-Gbps Fibre Channel
CISCO UCS Manager
CISCO INTERSIGHT
Cisco Nexus 9000 Series switches
Cisco MDS 9000 Series switches
10/25/40/100 GE
16/32-Gbps Fibre Channel
Pure Storage FlashArray//X

Cisco Unified Computing System

Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10-100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform with a unified management domain for managing all resources.

Cisco Unified Computing System consists of the following subsystems:

- Compute: The compute piece of the system incorporates servers based on the third-generation Intel Xeon Scalable processors. Servers are available in blade and rack form factor, managed by Cisco UCS Manager.
- Network: The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lowers costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.
- Virtualization: The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtual environments to support evolving business needs.
- Storage access: Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.

- Management: The system uniquely integrates compute, network, and storage access subsystems, enabling it to be managed as a single entity through Cisco Intersight. Cisco Intersight increases IT staff productivity by enabling storage, network, and server administrators to collaborate on Service Profiles that define the desired physical configurations and infrastructure policies for applications. Service Profiles increase business agility by enabling IT to automate and provision re-sources in minutes instead of days.

Cisco UCS Differentiators

Cisco Unified Computing System is revolutionizing the way servers are managed in the datacenter. The following are the unique differentiators of Cisco Unified Computing System and Cisco UCS Manager:

- Embedded Management: In Cisco UCS, the servers are managed by the embedded firmware in the fabric interconnects, eliminating the need for any external physical or virtual devices to manage the servers.
- Unified Fabric: In Cisco UCS, from blade server chassis or rack servers to FI, there is a single ethernet cable used for LAN, SAN, and management traffic. This converged I/O results in reduced cables, SFPs and adapters – reducing capital and operational expenses of the overall solution.
- Auto Discovery: By simply inserting the blade server in the chassis or connecting the rack server to the fabric interconnect, discovery and inventory of compute resources occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.
- Policy Based Resource Classification: Once a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD showcases the policy-based resource classification of Cisco UCS Manager.
- Combined Rack and Blade Server Management: Cisco UCS Manager can manage Cisco UCS B-series blade servers and Cisco UCS C-series rack servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.
- Model based Management Architecture: The Cisco UCS Manager architecture and management database is model based, and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.
- Policies, Pools, Templates: The management approach in Cisco UCS Manager is based on defining policies, pools, and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network, and storage resources.
- Loose Referential Integrity: In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each other. This provides great flexibility where different experts from different domains, such as network, storage, security, server, and virtualization work together to accomplish a complex task.
- Policy Resolution: In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organizational hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then the special policy named "default" is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.
- Service Profiles and Stateless Computing: A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as

far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.

- Built-in Multi-Tenancy Support: The combination of policies, pools and templates, loose referential integrity, policy resolution in the organizational hierarchy and a service profiles-based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environments typically observed in private and public clouds.
- Simplified QoS: Even though Fibre Channel and Ethernet are converged in the Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

Cisco Intersight

Cisco Intersight is a lifecycle management platform for your infrastructure, regardless of where it resides. In your enterprise data center, at the edge, in remote and branch offices, at retail and industrial sites—all these locations present unique management challenges and have typically required separate tools. Cisco Intersight Software as a Service (SaaS) unifies and simplifies your experience of the Cisco Unified Computing System (Cisco UCS).

Cisco Intersight software delivers a new level of cloud-powered intelligence that supports lifecycle management with continuous improvement. It is tightly integrated with the Cisco Technical Assistance Center (TAC). Expertise and information flow seamlessly between Cisco Intersight and IT teams, providing global management of Cisco infrastructure, anywhere. Remediation and problem resolution are supported with automated upload of error logs for rapid root-cause analysis.

Figure 2. Cisco Intersight

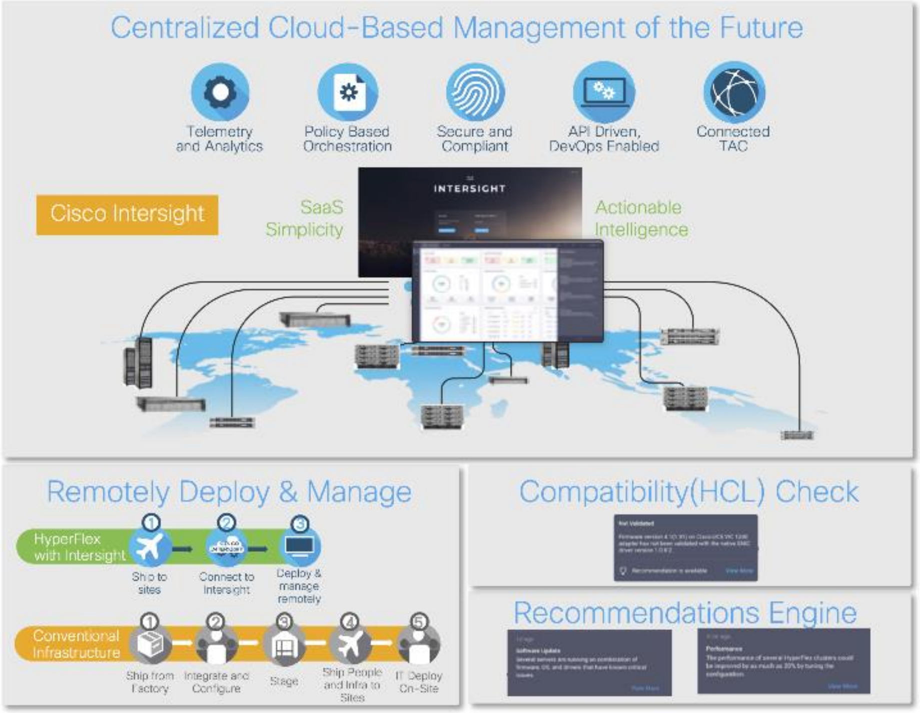

- Automate your infrastructure

  Cisco has a strong track record for management solutions that deliver policy-based automation to daily operations. Intersight SaaS is a natural evolution of our strategies. Cisco designed Cisco UCS to be 100 percent programmable. Cisco Intersight simply moves the control plane from the network into the cloud. Now you can manage your Cisco UCS and infrastructure wherever it resides through a single interface.

- Deploy your way

  If you need to control how your management data is handled, comply with data locality regulations, or consolidate the number of outbound connections from servers, you can use the Cisco Intersight Virtual Appliance for an on-premises experience. Cisco Intersight Virtual Appliance is continuously updated just like the SaaS version, so regardless of which approach you implement, you never have to worry about whether your management software is up to date.

- DevOps ready

  If you are implementing DevOps practices, you can use the Cisco Intersight API with either the cloud-based or virtual appliance offering. Through the API you can configure and manage infrastructure as code—you are not merely configuring an abstraction layer; you are managing the real thing. Through the API and support of cloud-based RESTful API, Terraform providers, Microsoft PowerShell scripts, or Python software, you can automate the deployment of settings and software for both physical and virtual layers. Using the API, you can simplify infrastructure lifecycle operations and increase the speed of continuous application delivery.

- Pervasive simplicity

  Simplify the user experience by managing your infrastructure regardless of where it is installed.

- Actionable intelligence
- Use best practices to enable faster, proactive IT operations
- Gain actionable insight for ongoing improvement and problem avoidance
- Manage anywhere
- Deploy in the data center and at the edge with massive scale
- Get visibility into the health and inventory detail for your Intersight Managed environment on-the-go with the Cisco Inter-sight Mobile App

For more information about Cisco Intersight and the different deployment options, go to: Cisco Intersight – Manage your systems anywhere.

Cisco UCS Fabric Interconnect

The Cisco UCS Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. Depending on the model chosen, the Cisco UCS Fabric Interconnect offers line-rate, low-latency, lossless 10 Gigabit, 25 Gigabit, 40 Gigabit, or 100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel connectivity. Cisco UCS Fabric Interconnects provide the management and communication backbone for the Cisco UCS X-Series and Cisco UCS 9000 Series Server Chassis. All servers and chassis attached to the Cisco UCS Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabrics, the Cisco UCS Fabric Interconnects provide both the LAN and SAN connectivity for all servers within its domain.

For networking performance, the Cisco UCS 6454 Series uses a cut-through architecture, supporting deterministic, low latency, line rate 10/25/40/100 Gigabit Ethernet ports, 3.82 Tbps of switching capacity, and 320 Gbps bandwidth per Cisco 5108 blade chassis when connected through the IOM 2208 model. The product family supports Cisco low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet net-works. The Fabric Interconnect supports multiple traffic classes over the Ethernet fabric from the servers to the uplinks. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6454 Fabric Interconnect

The Cisco UCS 6454 Fabric Interconnect is a one-rack-unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has eight (8) 10/25-Gbps fixed Ethernet ports, which optionally can be configured as 8/16/32-Gbps FC ports (ports 1 to 8), thirty-six (36) 10/25-Gbps fixed Ethernet ports (ports 9 to 44), four (4) 1/10/25-Gbps Ethernet ports (ports 45 to 48), and finally six (6) 40/100-Gbps Ethernet uplink ports (ports 49 to 54). For more information, refer to the Cisco UCS 6454 Fabric Interconnect spec sheet: https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/6400-specsheet.pdf

Figure 3. Cisco UCS 6454 Fabric Interconnect

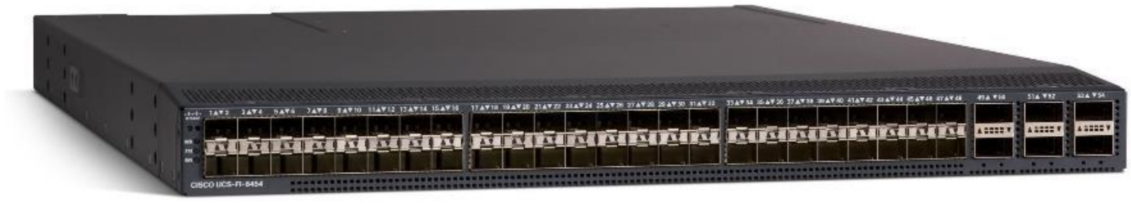

Cisco Unified Computing System X-Series

The Cisco UCS X-Series Modular System is designed to take the current generation of the Cisco UCS platform to the next level with its future-ready design and cloud-based management. Decoupling and moving the platform management to the cloud allows Cisco UCS to respond to customer feature and scalability requirements in a much faster and efficient manner. Cisco UCS X-Series state of the art hardware simplifies the data-center design by providing flexible server options. A single server type, supporting a broader range of workloads, results in fewer different data-center products to manage and maintain. The Cisco Intersight cloud-management platform manages Cisco UCS X-Series as well as integrating with third-party devices, including VMware vCenter and Pure Storage, to provide visibility, optimization, and orchestration from a single platform, thereby driving agility and deployment consistency.

Figure 4. Cisco UCS X9508 Chassis

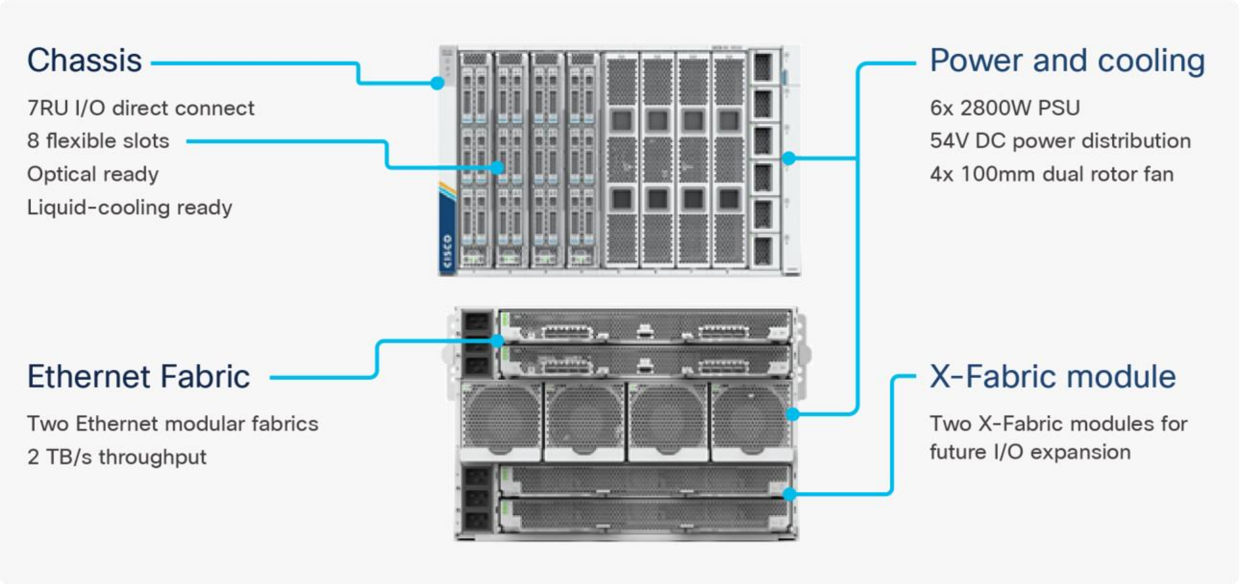

Chassis

7RU I/O direct connect

8 flexible slots

Optical ready

Liquid-cooling ready

Power and cooling

6x 2800W PSU

54V DC power distribution

4x 100mm dual rotor fan

Ethernet Fabric

Two Ethernet modular fabrics

2 TB/s throughput

X-Fabric module

Two X-Fabric modules for future I/O expansion

The various components of the Cisco UCS X-Series are described in the following sections.

Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As shown in Figure 5, Cisco UCS X9508 chassis has only a power-distribution midplane. This midplane-free design provides fewer obstructions for better airflow. For I/O connectivity, vertically oriented compute nodes intersect with horizontally oriented fabric modules, allowing the chassis to support future fabric innovations. Cisco UCS X9508 Chassis' superior packaging enables larger compute nodes, thereby providing more space for actual compute components, such as memory, GPU, drives, and accelerators. Improved airflow through the chassis enables support for higher power components, and more space allows for future thermal solutions (such as liquid cooling) without limitations.

Figure 5. Cisco UCS X9508 Chassis - Midplane Free Design

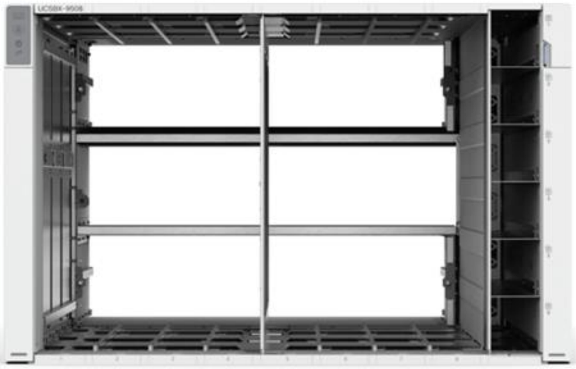

The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes and a pool of future I/O resources that may include GPU accelerators, disk storage, and nonvolatile memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) that connect the chassis to upstream Cisco UCS 6400 Series Fabric Interconnects. At the bottom rear of the chassis are slots ready to house future X-Fabric modules that can flexibly connect the compute nodes with I/O devices. Six 2800W Power Supply Units (PSUs) provide 54V power to the chassis with N, N+1, and N+N redundancy. A higher voltage allows efficient power delivery with less copper and reduced power loss. Efficient, 100mm, dual counter-rotating fans deliver industry-leading airflow and power efficiency, and optimized thermal algorithms enable different cooling modes to best support the customer's environment.

Cisco UCSX 9108-25G Intelligent Fabric Modules

For the Cisco UCS X9508 Chassis, the network connectivity is provided by a pair of Cisco UCSX 9108-25G Intelligent Fabric Modules (IFMs). Like the fabric extenders used in the Cisco UCS 5108 Blade Server Chassis, these modules carry all network traffic to a pair of Cisco UCS 6400 Series Fabric Interconnects (FIs). IFMs also host the Chassis Management Controller (CMC) for chassis management. In contrast to systems with fixed networking components, Cisco UCS X9508's midplane-free design enables easy upgrades to new networking technologies as they emerge making it straightforward to accommodate new network speeds or technologies in the future.

Figure 6. Cisco UCSX 9108-25G Intelligent Fabric Module

Each IFM supports eight 25Gb uplink ports for connecting the Cisco UCS X9508 Chassis to the FIs and 32 25Gb server ports for the eight compute nodes. IFM server ports can provide up to 200 Gbps of unified fabric connectivity per compute node across the two IFMs. The uplink ports connect the chassis to the Cisco UCS FIs, providing up to 400Gbps connectivity across the two IFMs. The unified fabric carries management, VM, and Fibre Channel over Ethernet (FCoE) traffic to the FIs, where management traffic is routed to the Cisco Intersight cloud operations platform, FCoE traffic is forwarded to the native Fibre Channel interfaces through unified ports on the FI (to Cisco MDS switches), and data Ethernet traffic is forwarded upstream to the data center network via Cisco Nexus switches.

Cisco UCS X210c M7 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to eight Cisco UCS X210c M7 Compute Nodes. The hardware details of the Cisco UCS X210c M7 Compute Nodes are shown in Figure 7:

Figure 7. Cisco UCS X210c M7 Compute Node

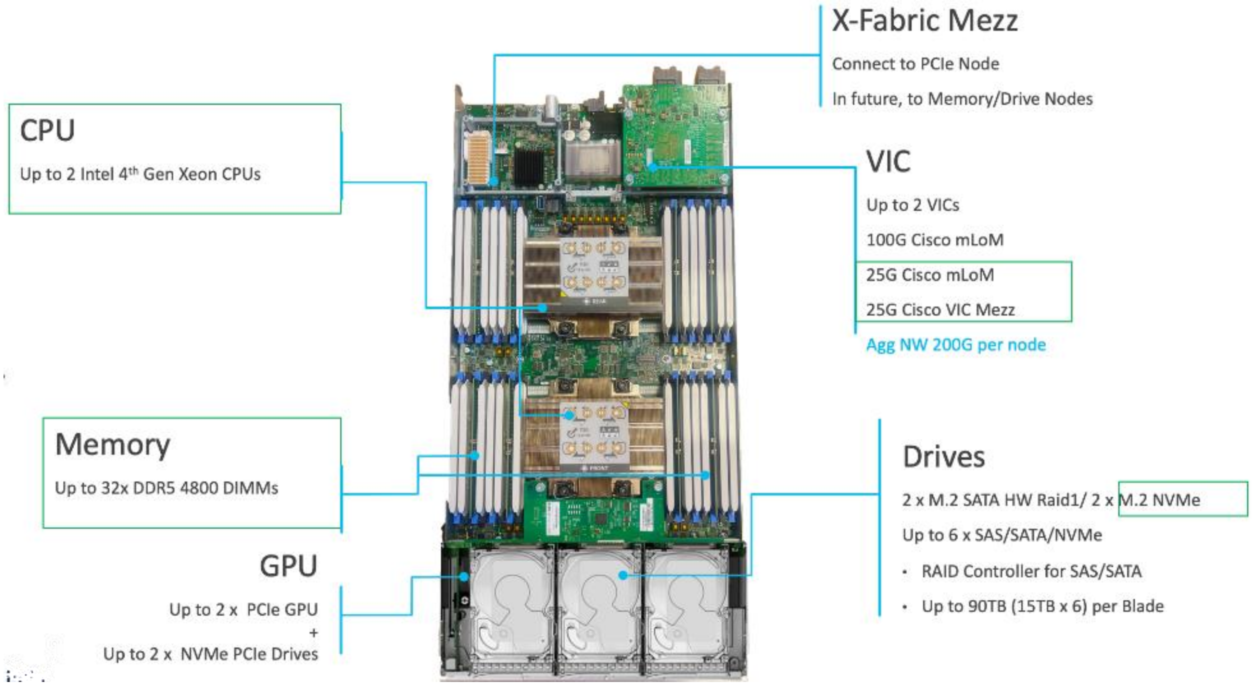
X-Fabric Mezz
Connect to PCIe Node
In future, to Memory/Drive Nodes
CPU
Up to 2 Intel 4th Gen Xeon CPUs
VIC
Up to 2 VICs
100G Cisco mLoM
25G Cisco mLoM
25G Cisco VIC Mezz
Agg NW 200G per node
Memory
Up to 32x DDR5 4800 DIMMs
Drives
2 x M.2 SATA HW Raid1/ 2 x M.2 NVMe
Up to 6 x SAS/SATA/NVMe
• RAID Controller for SAS/SATA
• Up to 90TB (15TB x 6) per Blade
GPU
Up to 2 x PCIe GPU
+
Up to 2 x NVMe PCIe Drives

The Cisco UCS X210c M7 features:

- CPU: Up to 2x 4th Gen Intel Xeon Scalable Processors with up to 60 cores per processor and 2.625 MB Level 3 cache per core and up to 112.5 MB per CPU.
- Memory: Up to 8TB of main memory with 32x 256 GB DDR5-4800 DIMMS.
- Disk storage: Up to six hot-pluggable, solid-state drives (SSDs), or non-volatile memory express (NVMe) 2.5-inch drives with a choice of enterprise-class redundant array of independent disks (RAIDs) or passthrough controllers, up to two M.2 SATA drives with optional hardware RAID.
- Optional front mezzanine GPU module: The Cisco UCS front mezzanine GPU module is a passive PCIe Gen 4.0 mezzanine option with support for up to two U.2 NVMe drives and two HHHL GPUs.
- mLOM virtual interface cards:
  - Cisco UCS Virtual Interface Card (VIC) 15420 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 50 Gbps of unified fabric connectivity to each of the chassis intelligent fabric modules (IFMs) for 100 Gbps connectivity per server.

Cisco UCS Virtual Interface Card (VIC) 15231 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis intelligent fabric modules (IFMs) for 100 Gbps connectivity per server.

- Optional mezzanine card:
  - Cisco UCS 5th Gen Virtual Interface Card (VIC) 15422 can occupy the server's mezzanine slot at the bottom rear of the chassis. The card's I/O connectors link the Cisco UCS X-Fabric technology. An included bridge card extends this VIC's 2x 50 Gbps of network connections through IFM connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per server).
  - Cisco UCS PCI Mezz card for X-Fabric can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric modules and enable connectivity to the Cisco UCS X440p PCIe Node.
  - All VIC mezzanine cards also provide I/O connections from the X210C M7 compute node to the X440p PCIe Node.
- Security: The server supports an optional Trusted Platform Module (TPM). Additional security features include a secure boot FPGA and ACT2 anticounterfeit provisions.

Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS X210c M7 Compute Nodes support the following Cisco fifth-generation VIC cards:

Cisco VIC 15420

Cisco VIC 15420 fits the mLOM slot in the Cisco X210c Compute Node and enables up to 25 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 100 Gbps of connectivity per server. Cisco VIC 15420 connectivity to the IFM and up to the fabric interconnects is delivered through 4x 25-Gbps connections, which are configured automatically as 4x 25-Gbps port channels. Cisco VIC 15420 supports 256 virtual interfaces (both Fibre Channel and Ethernet) along with the latest networking innovations such as NVMeoF over RDMA (ROCEv2), VxLAN/NVGRE offload, and so on.

Single Cisco VIC 15420 in Cisco UCS X210c M7

The connections between the 5th generation Cisco UCS VIC (Cisco UCS VIC 15420) in the Cisco UCS B200 blades and the I/O mod-ules in the Cisco UCS 5108 chassis comprise of multiple 10Gbps KR lanes. The same connections between Cisco VIC 15420 and IFMs in Cisco UCS X-Series comprise of multiple 25Gbps KR lanes resulting in 2.5x better connectivity in Cisco UCS X210c M7 Compute Nodes. The network interface speed comparison between VMware ESXi installed on Cisco UCS X210C M7 with Cisco UCS VIC 15420 and Cisco UCS X210c M7 with Cisco UCS VIC 15420.

Cisco Switching

Cisco Nexus 93180YC-FX Switches

The Cisco Nexus 93180YC-FX Switch provides a flexible line-rate Layer 2 and Layer 3 feature set in a compact form factor. Designed with Cisco Cloud Scale technology, it supports highly scalable cloud architectures. With the option to operate in Cisco NX-OS or Application Centric Infrastructure (ACI) mode, it can be deployed across enterprise, service provider, and Web 2.0 data centers.

- Architectural Flexibility
  - Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
  - Leaf node support for Cisco ACI architecture is provided in the roadmap

    - Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support
- Feature Rich
    - Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
    - ACI-ready infrastructure helps users take advantage of automated policy-based systems management
    - Virtual Extensible LAN (VXLAN) routing provides network services
    - Rich traffic flow telemetry with line-rate data collection
    - Real-time buffer utilization per port and per queue, for monitoring traffic micro-bursts and application traffic patterns
- Highly Available and Efficient Design
    - High-density, non-blocking architecture
    - Easily deployed into either a hot-aisle and cold-aisle configuration
    - Redundant, hot-swappable power supplies and fan trays
- Simplified Operations
    - Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
    - An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infra-structure
    - Python Scripting for programmatic access to the switch command-line interface (CLI)
    - Hot and cold patching, and online diagnostics
- Investment Protection

A Cisco 40 Gbe bidirectional transceiver allows reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Giga-bit Ethernet Support for 1 Gbe and 10 Gbe access connectivity for data centers migrating access switching infrastructure to faster speed. The following is supported:

- 1.8 Tbps of bandwidth in a 1 RU form factor
- 48 fixed 1/10/25-Gbe SFP+ ports
- 6 fixed 40/100-Gbe QSFP+ for uplink connectivity
- Latency of less than 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies
- Hot swappable 3+1 redundant fan trays

Figure 8. Cisco Nexus 93180YC-EX Switch

Cisco MDS 9132T 32-Gb Fiber Channel Switch

The next-generation Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switch (Figure 9) provides high-speed Fibre Channel connectivity from the server rack to the SAN core. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the dual benefits of greater bandwidth and consolidation.

Small-scale SAN architectures can be built from the foundation using this low-cost, low-power, non-blocking, line-rate, and low-latency, bi-directional airflow capable, fixed standalone SAN switch connecting both storage and host ports.

Medium-size to large-scale SAN architectures built with SAN core directors can expand 32-Gb connectivity to the server rack using these switches either in switch mode or Network Port Virtualization (NPV) mode.

Additionally, investing in this switch for the lower-speed (4- or 8- or 16-Gb) server rack gives you the option to upgrade to 32-Gb server connectivity in the future using the 32-Gb Host Bus Adapter (HBA) that are available today. The Cisco MDS 9132T 32-Gb 32-Port Fibre Channel switch also provides unmatched flexibility through a unique port expansion module (Figure 15) that provides a robust cost-effective, field swappable, port upgrade option.

This switch also offers state-of-the-art SAN analytics and telemetry capabilities that have been built into this next-generation hardware platform. This new state-of-the-art technology couples the next-generation port ASIC with a fully dedicated Network Processing Unit designed to complete analytics calculations in real time. The telemetry data extracted from the inspection of the frame headers are calculated on board (within the switch) and, using an industry-leading open format, can be streamed to any analytics-visualization platform. This switch also includes a dedicated 10/100/1000BASE-T telemetry port to maximize data delivery to any telemetry receiver including Cisco Data Center Network Manager.

Figure 9. Cisco MDS 9132T 32-Gb Fibre Channel Switch

Figure 10. Cisco MDS 9132T 32-Gb 16-Port Fibre Channel Port Expansion Module

- Features
    - High performance: Cisco MDS 9132T architecture, with chip-integrated nonblocking arbitration, provides consistent 32-Gb low-latency performance across all traffic conditions for every Fibre Channel port on the switch.
    - Capital Expenditure (CapEx) savings: The 32-Gb ports allow users to deploy them on existing 16- or 8-Gb transceivers, reducing initial CapEx with an option to upgrade to 32-Gb transceivers and adapters in the future.
    - High availability: Cisco MDS 9132T switches continue to provide the same outstanding availability and reliability as the previous-generation Cisco MDS 9000 Family switches by providing optional redundancy on all major components such as the power supply and fan. Dual power supplies also facilitate redundant power grids.

- ◦ Pay-as-you-grow: The Cisco MDS 9132T Fibre Channel switch provides an option to deploy as few as eight 32-Gb Fibre Channel ports in the entry-level variant, which can grow by 8 ports to 16 ports, and thereafter with a port expansion module with sixteen 32-Gb ports, to up to 32 ports. This approach results in lower initial investment and power consumption for entry-level configurations of up to 16 ports compared to a fully loaded switch. Upgrading through an expansion module also reduces the overhead of managing multiple instances of port activation licenses on the switch. This unique combination of port upgrade options allow four possible configurations of 8 ports, 16 ports, 24 ports and 32 ports.
- ◦ Next-generation Application-Specific Integrated Circuit (ASIC): The Cisco MDS 9132T Fibre Channel switch is powered by the same high-performance 32-Gb Cisco ASIC with an integrated network processor that powers the Cisco MDS 9700 48-Port 32-Gb Fibre Channel Switching Module. Among all the advanced features that this ASIC enables, one of the most notable is inspection of Fibre Channel and Small Computer System Interface (SCSI) headers at wire speed on every flow in the smallest form-factor Fibre Channel switch without the need for any external taps or appliances. The recorded flows can be analyzed on the switch and also exported using a dedicated 10/100/1000BASE-T port for telemetry and analytics purposes.
- ◦ Intelligent network services: Slow-drain detection and isolation, VSAN technology, Access Control Lists (ACLs) for hardware-based intelligent frame processing, smartzoning and fabric wide Quality of Service (QoS) enable migration from SAN islands to enterprise-wide storage networks. Traffic encryption is optionally available to meet stringent security requirements.
- ◦ Sophisticated diagnostics: The Cisco MDS 9132T provides intelligent diagnostics tools such as Inter-Switch Link (ISL) diagnostics, read diagnostic parameters, protocol decoding, network analysis tools, and integrated Cisco Call Home capability for greater reliability, faster problem resolution, and reduced service costs.
- ◦ Virtual machine awareness: The Cisco MDS 9132T provides visibility into all virtual machines logged into the fabric. This feature is available through HBAs capable of priority tagging the Virtual Machine Identifier (VMID) on every FC frame. Virtual machine awareness can be extended to intelligent fabric services such as analytics[1] to visualize performance of every flow originating from each virtual machine in the fabric.
- ◦ Programmable fabric: The Cisco MDS 9132T provides powerful Representational State Transfer (REST) and Cisco NX-API capabilities to enable flexible and rapid programming of utilities for the SAN as well as polling point-in-time telemetry data from any external tool.
- ◦ Single-pane management: The Cisco MDS 9132T can be provisioned, managed, monitored, and troubleshot using Cisco Data Center Network Manager (DCNM), which currently manages the entire suite of Cisco data center products.
- ◦ Self-contained advanced anticounterfeiting technology: The Cisco MDS 9132T uses on-board hardware that protects the entire system from malicious attacks by securing access to critical components such as the bootloader, system image loader and Joint Test Action Group (JTAG) interface.

VMware vSphere 8.0

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructures (resources including CPUs, storage, and networking) as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

VMware vSphere 8.0 has several improvements and simplifications including, but not limited to:

- vSphere Memory Monitoring and Remediation, and support for snapshots of PMem VMs: vSphere Memory Monitoring and Remediation collects data and provides visibility of performance statistics to help you determine if your application workload is regressed due to Memory Mode. VMware vSphere 8.0 also adds support for snapshots of PMem VMs.

- Improved interoperability between vCenter Server and ESXi versions: Starting with vSphere 7.0 Update 3, vCenter Server can manage ESXi hosts from the previous two major releases and any ESXi host from version 7.0 and 7.0 updates. For example, vCenter Server 8.0 can manage ESXi hosts of versions 6.5, 6.7 and 7.0, all 7.0 update releases, including later than Update 3, and a mixture of hosts between major and update versions.
- New VMNIC tag for NVMe-over-RDMA (NVME/RoCEv2) storage traffic: ESXi 7.0 Update 3 adds a new VMNIC tag for NVMe-over-RDMA (NVMe/RoCEv2) storage traffic. This VMkernel port setting enables NVMe-over-RDMA traffic to be routed over the tagged interface. You can also use the ESXCLI command esxcli network ip interface tag add -i <interface name> -t NVMeRDMA to enable the NVMeRDMA VMNIC tag.
- NVMe over TCP support: vSphere 8.0 8 the NVMe-oF suite with the NVMe over TCP storage protocol to enable high performance and parallelism of NVMe devices over a wide deployment of TCP/IP networks.
- Micro-second level time accuracy for workloads: ESXi 7.0 Update 3 adds the hardware timestamp Precision Time Protocol (PTP) to enable micro-second level time accuracy. For more information, see Use PTP for Time and Date Synchronization of a Host.

For more information about VMware vSphere and its components, see: https://www.vmware.com/products/vsphere.html.

VMware vSphere vCenter

VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest OS, and other critical components of a virtual infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

Red Hat Ansible

Ansible is simple and powerful, allowing users to easily manage various physical devices within FlashStack including the provisioning of Cisco UCS servers, Cisco Nexus switches, Pure Storage FlashArray and VMware vSphere. Using Ansible’s playbook-based automation is easy and integrates into your current provisioning infrastructure.

Cisco Intersight Assist Device Connector for VMware vCenter and Pure Storage FlashArray

Cisco Intersight integrates with VMware vCenter and Pure Storage FlashArray as follows:

- Cisco Intersight uses the device connector running within Cisco Intersight Assist virtual appliance to communicate with the VMware vCenter.
- Cisco Intersight uses the device connector running within a Cisco Intersight Assist virtual appliance to integrate with all Pure Storage FlashArray /models. The newest version 1.1 of Pure Storage integration to Cisco Intersight introduces support for REST API 2.x for FlashArray products (running Purity//FA 6.0.3 or later), along with User Agent support (for telemetry). Intersight Cloud Orchestrator now has new storage tasks for adding/removing a Pure Storage snapshot and copy a Pure Storage volume from snapshot.

Figure 11. **Cisco Intersight and vCenter and Pure Storage Integration**

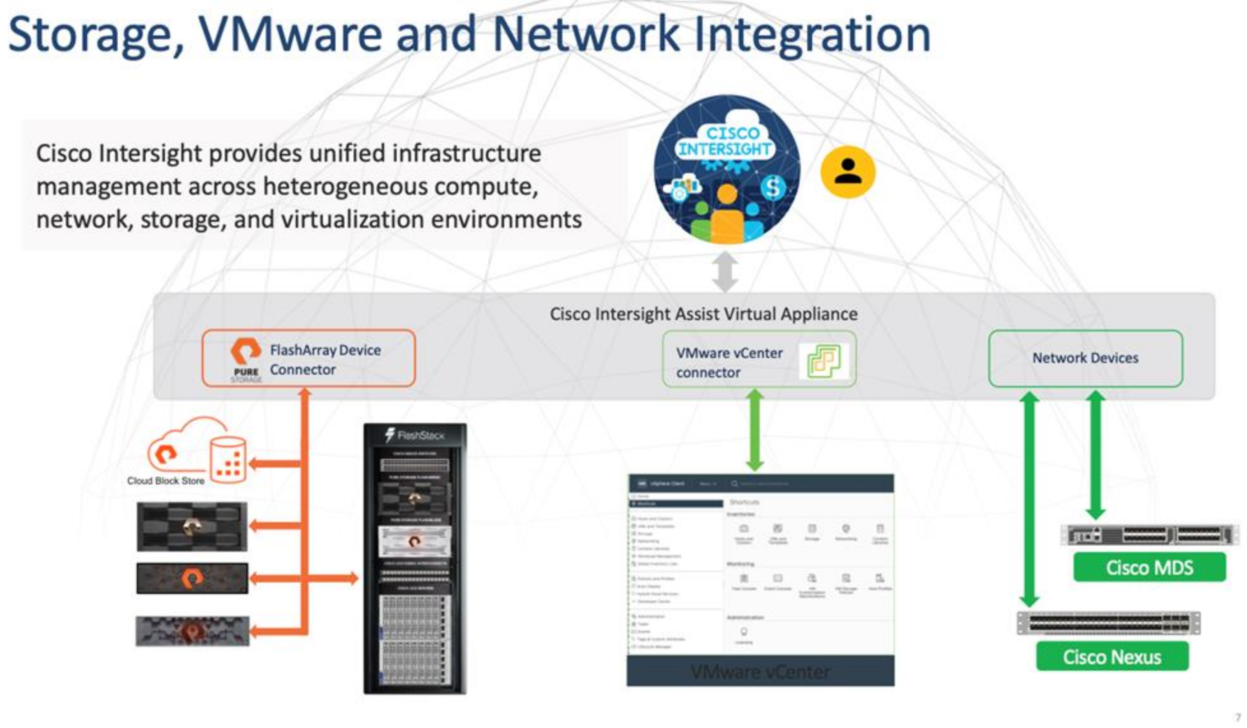

Storage, VMware and Network Integration

Cisco Intersight provides unified infrastructure management across heterogeneous compute, network, storage, and virtualization environments

The device connector provides a safe way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure Internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware hypervisor and FlashArray storage environments. The integration architecture enables FlashStack customers to use new management capabilities with no compromise in their existing VMware or FlashArray operations. IT users will be able to manage heterogeneous infrastructure from a centralized Cisco Intersight portal. At the same time, the IT staff can continue to use VMware vCenter and the Pure Storage dashboard for comprehensive analysis, diagnostics, and reporting of virtual and storage environments. The next section addresses the functions that this integration provides.

Purity for FlashArray

The essential element of every FlashArray is the Purity Operating Environment software. Purity implements advanced data reduction, storage management, and flash management features, enabling organizations to enjoy Tier 1 data services for all workloads, proven 99.9999% availability over multiple years (inclusive of maintenance and generational upgrades), completely non-disruptive operations, 2X better data reduction versus alternative all-flash solutions, and the power and efficiency of DirectFlash.

Figure 12. **Pure Storage FlashArray Family**

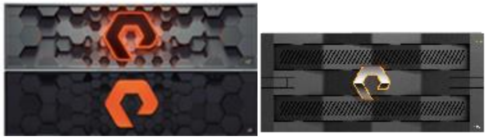

Moreover, Purity includes enterprise-grade data security, modern data protection options, and complete business continuity and global disaster recovery through ActiveCluster multi-site stretch cluster and ActiveDR for continuous replication with near zero RPO. All these features are included with every array.

For healthcare customers, Pure Storage takes more time to allow a new Purity release to mature in the field before it's recommended for use. The criteria for adoption is as follows:

- Purity releases must be in a General Availability (GA) status for at least 45 days.
- Purity releases must have at least 10,000 total array days.
- Purity releases should have a minimum number of events per array days. There should be no more than 1 event seen per 3000 arrays days in the field for high acceptance.
- Arrays that are not running active production or non-production workloads, are not bound by this extended Purity release cycle. It's always best to consult with your local Pure Storage Systems Engineer for any additional information related to code releases.

Healthcare customers can request from Pure Storage the latest available code recommendations by sending an email with the subject "HC Code Request" or "HC Purity Request", as shown below. Pure Storage has created an autoresponder that will send out the latest Purity Release levels for Epic-specific storage devices based on the above criteria.

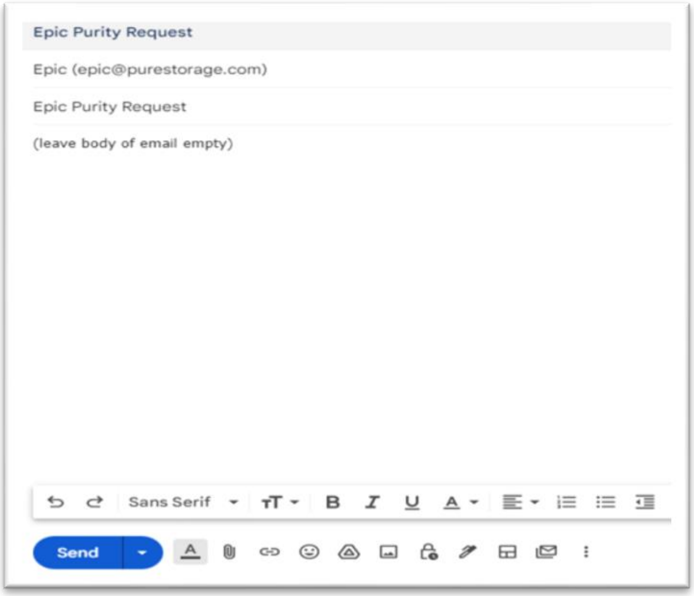

Epic Purity Request

Epic (epic@purestorage.com)

Epic Purity Request

(leave body of email empty)

IMPORTANT! There may be situations where a Purity release will require immediate consideration based on engineering needs. If this happens, Pure Storage will send out an alert with the updated Purity release level along with a justification statement. There can be a situation where there will be two Purity releases supported for Epic customers. Typically, this is due to varying models of arrays. For example, currently the FlashArray //X R3 platform supports 6.3.7+. But the FlashArray //XL platform ships with Purity release 6.4.x. In this situation, expect the purity releases to stay with their designated array models unless otherwise noted.

FlashArray File Services

Pure Storage acquired Compuverde last year, and they've been busy at work integrating this technology into the Purity//FA operating system. They emphasize the "integrating," because they didn't just take the existing product, drop it onto a FlashArray system, and run it on top of Purity. Instead, they incorporated key parts of it into Purity to give you the advantages of native files alongside blocks.

The SMB and NFS protocols bring consolidated storage to the Purity//FA operating system, complementing its block capabilities, while the file system offers features like directory snapshots and directory-level performance and space monitoring. For the purposes of this reference architecture, we will be focusing on using File Services for User Profile management.

Figure 13. **FlashArray//X Specifications**

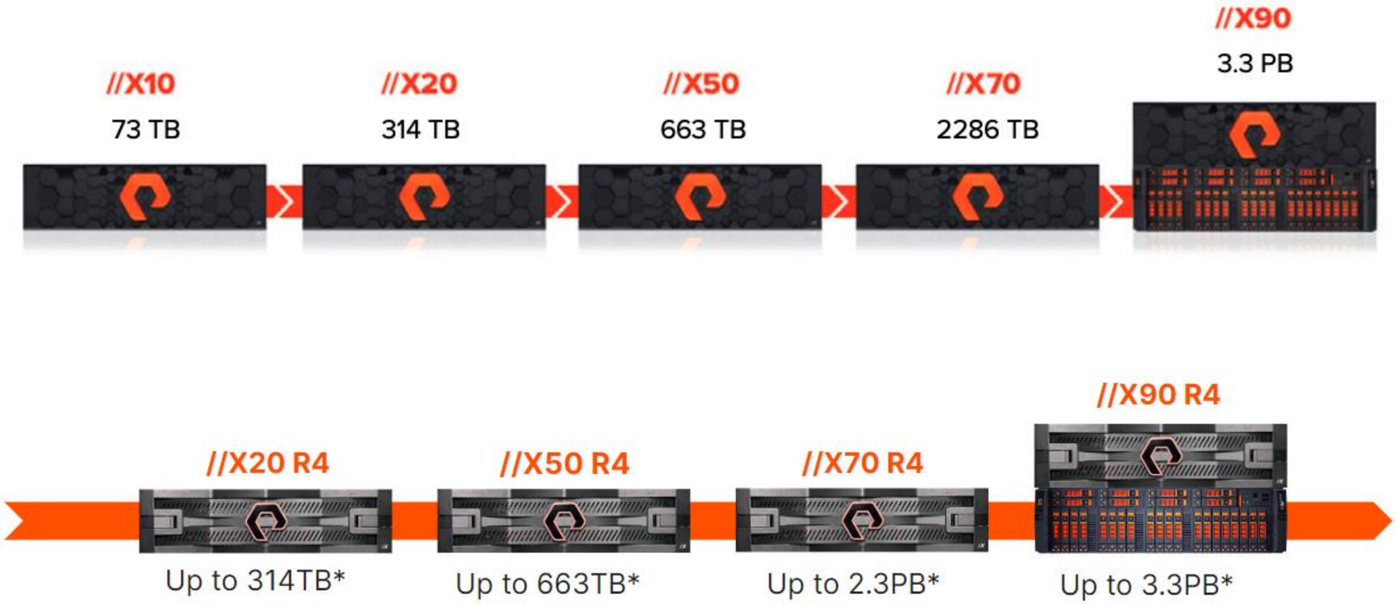

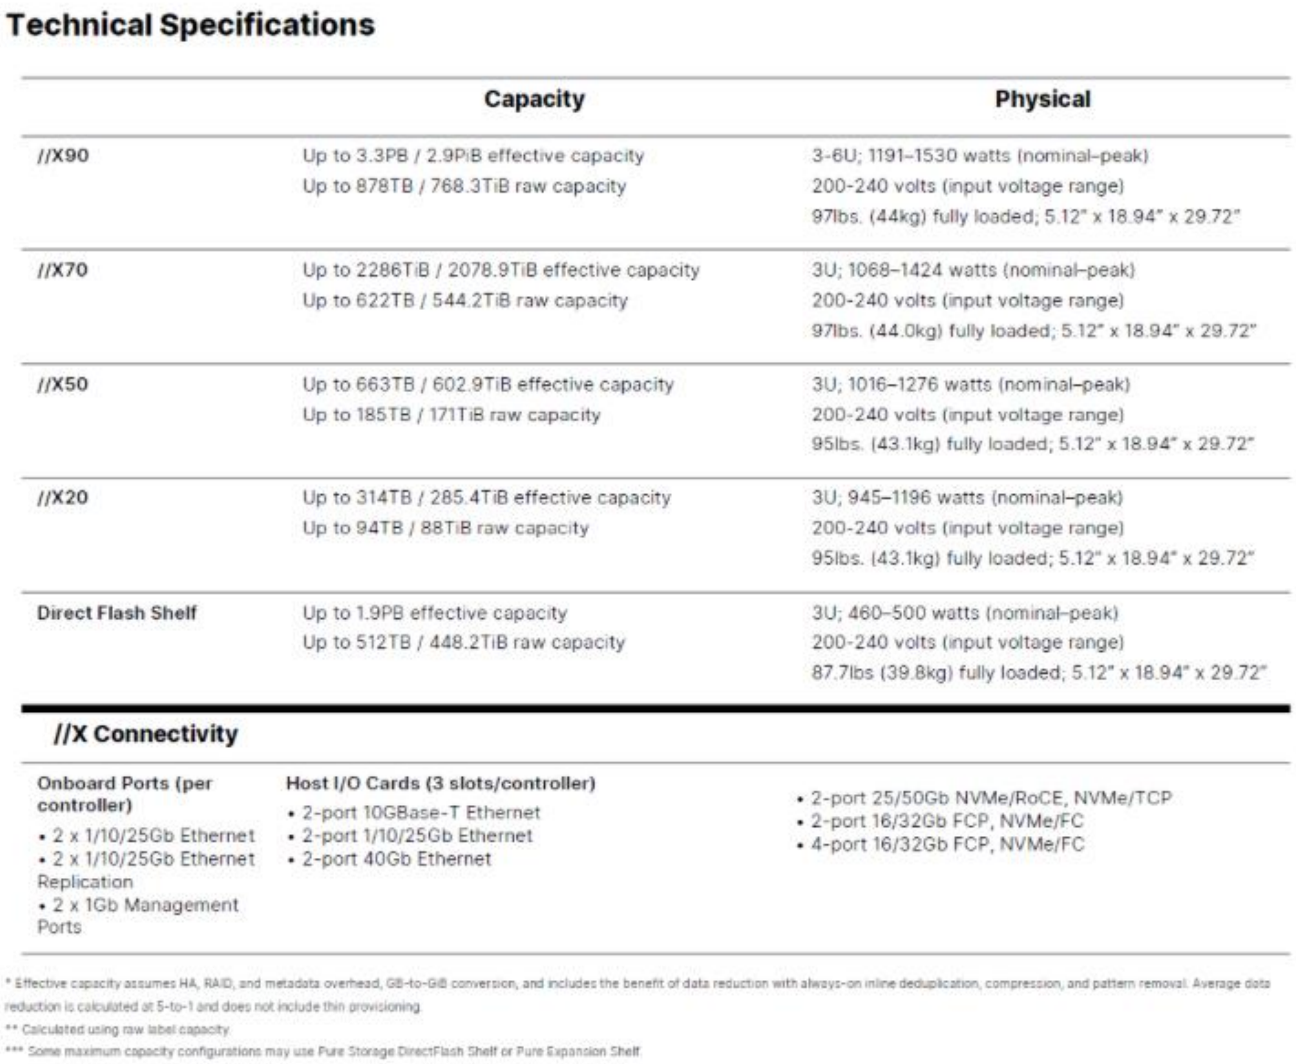

Technical Specifications

	Capacity	Physical
//X90	Up to 3.3PB / 2.9PiB effective capacity Up to 878TB / 768.3TiB raw capacity	3-6U; 1191–1530 watts (nominal–peak) 200-240 volts (input voltage range) 97lbs. (44kg) fully loaded; 5.12" x 18.94" x 29.72"
//X70	Up to 2286TiB / 2078.9TiB effective capacity Up to 622TB / 544.2TiB raw capacity	3U; 1068–1424 watts (nominal–peak) 200-240 volts (input voltage range) 97lbs. (44.0kg) fully loaded; 5.12" x 18.94" x 29.72"
//X50	Up to 663TB / 602.9TiB effective capacity Up to 185TB / 171TiB raw capacity	3U; 1016–1276 watts (nominal–peak) 200-240 volts (input voltage range) 95lbs. (43.1kg) fully loaded; 5.12" x 18.94" x 29.72"
//X20	Up to 314TB / 285.4TiB effective capacity Up to 94TB / 88TiB raw capacity	3U; 945–1196 watts (nominal–peak) 200-240 volts (input voltage range) 95lbs. (43.1kg) fully loaded; 5.12" x 18.94" x 29.72"
Direct Flash Shelf	Up to 1.9PB effective capacity Up to 512TB / 448.2TiB raw capacity	3U; 460–500 watts (nominal–peak) 200-240 volts (input voltage range) 87.7lbs (39.8kg) fully loaded; 5.12" x 18.94" x 29.72"

//X Connectivity

Onboard Ports (per controller)	Host I/O Cards (3 slots/controller)	
• 2 x 1/10/25Gb Ethernet • 2 x 1/10/25Gb Ethernet Replication • 2 x 1Gb Management Ports	• 2-port 10GBase-T Ethernet • 2-port 1/10/25Gb Ethernet • 2-port 40Gb Ethernet	• 2-port 25/50Gb NVMe/RoCE, NVMe/TCP • 2-port 16/32Gb FCP, NVMe/FC • 4-port 16/32Gb FCP, NVMe/FC

* Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1 and does not include thin provisioning.

** Calculated using raw label capacity.

*** Some maximum capacity configurations may use Pure Storage DirectFlash Shelf or Pure Expansion Shelf.

** Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1 and does not include thin provisioning or snapshots.

† Array accepts Pure Storage DirectFlash Shelf and/or Pure Storage SAS-based expansion shelf.

Evergreen Storage

You can deploy storage once and enjoy a subscription to continuous innovation through Pure Storage's Evergreen Storage ownership model: expand and improve performance, capacity, density, and/or features for 10 years or more – all without downtime, performance impact, or data migrations. Pure Storage has disrupted the industry's 3-5-year rip-and-replace cycle by engineering compatibility for future technologies right into its products, notably nondisruptive capability to upgrade from //M to //X with NVMe, DirectMemory, and NVMe-oF capability.

Pure1

Pure1, our cloud-based management, analytics, and support platform, expands the self-managing, plug-n-play design of Pure Storage all-flash arrays with the machine learning predictive analytics and continuous scanning of Pure1 Meta to enable an effortless, worry-free data platform.

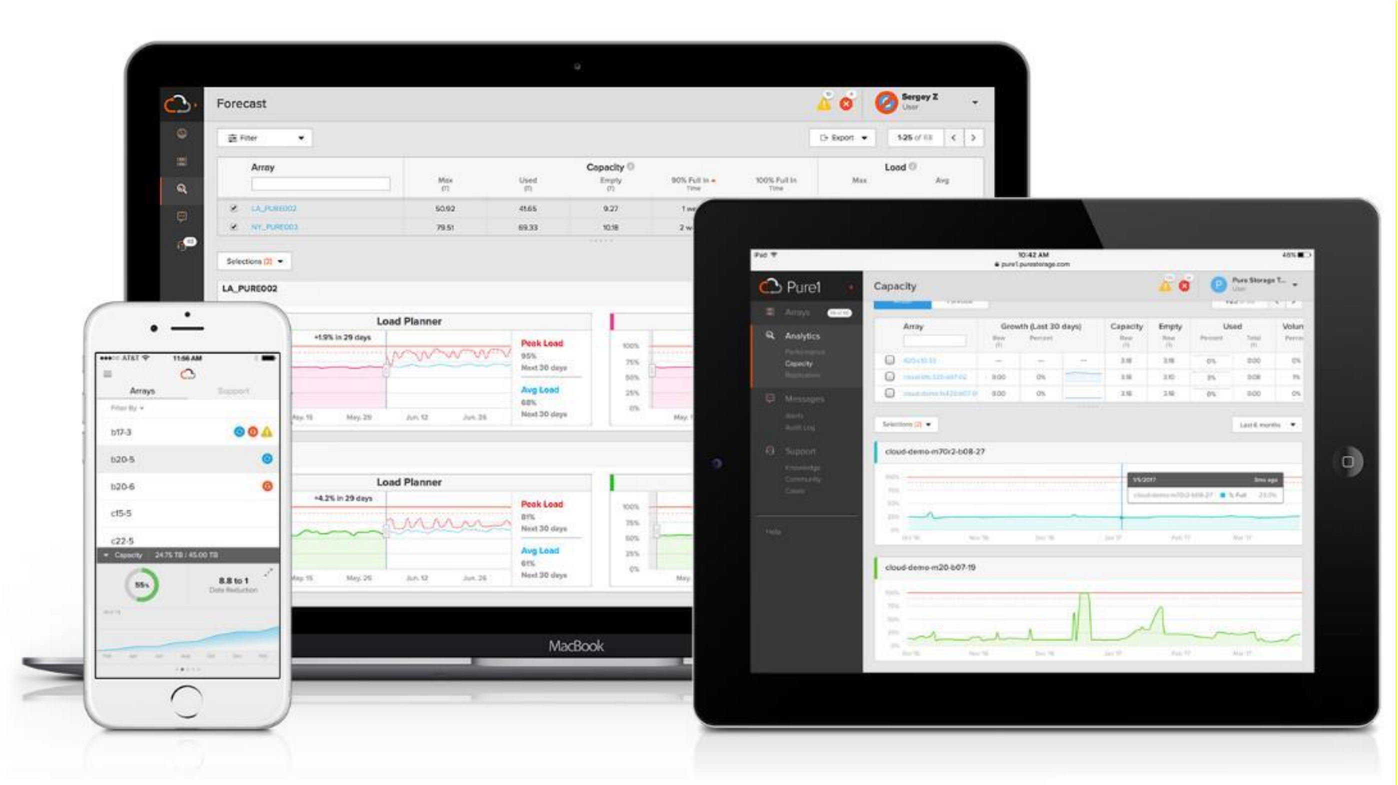

Pure1 Manage

In the Cloud IT operating model, installing, and deploying management software is an oxymoron: you simply login. Pure1 Manage is SaaS-based, allowing you to manage your array from any browser or from the Pure1 Mobile App – with nothing extra to purchase, deploy, or maintain. From a single dashboard you can manage all your arrays, with full visibility on the health and performance of your storage.

Pure1 Analyze

Pure1 Analyze delivers true performance forecasting – giving customers complete visibility into the performance and capacity needs of their arrays – now and in the future. Performance forecasting enables intelligent consolidation and unprecedented workload optimization.

Pure1 Capacity/Planning/Reporting

This feature of Pure1 is a set of tools that Pure Storage customers can leverage to understand how their environment is performing, how their storage environment is growing, and how to present capacity and performance planning materials to leadership.

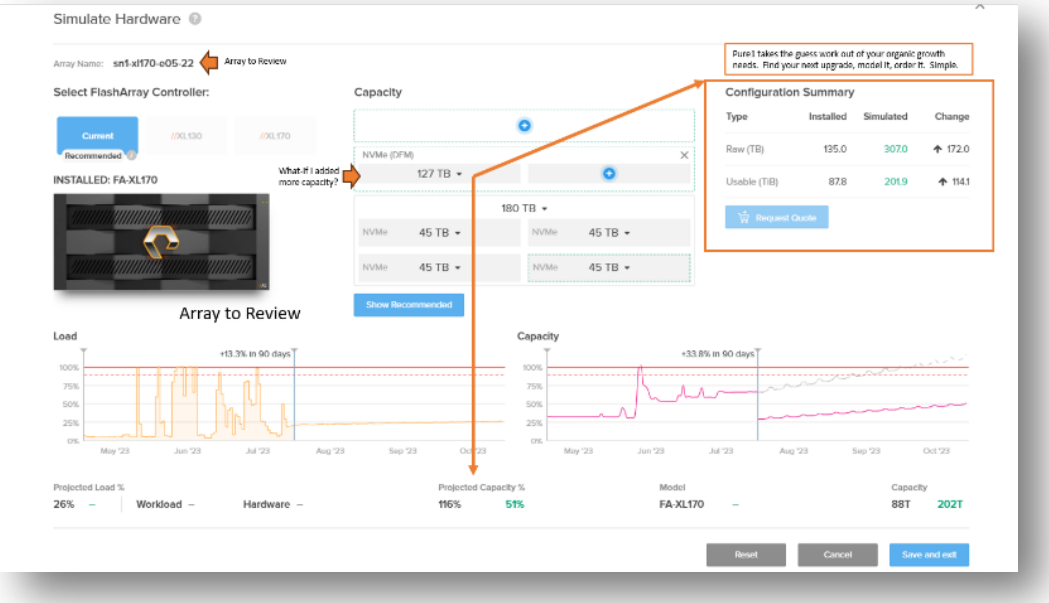

Simulate Hardware

Array Name: sn1-xl170-e05-22 ← Array to Review

Pure1 takes the guess work out of your organic growth needs. Find your next upgrade, model it, order it. Simple.

Select FlashArray Controller:

INSTALLED: FA-XL170

Array to Review

What-If I added more capacity?

Capacity

NVMe (DFM) 127 TB

180 TB

NVMe 45 TB | NVMe 45 TB
NVMe 45 TB | NVMe 45 TB

Configuration Summary

Type	Installed	Simulated	Change
Raw (TB)	135.0	307.0	↑ 172.0
Usable (TiB)	87.8	201.9	↑ 114.1

Projected Load %: 26% – | Workload – | Hardware –

Projected Capacity %: 116% 51%

Model: FA-XL170 –

Capacity: 88T 202T

Pure1 takes the sizing of future upgrades for capacity and gives the healthcare customer the power to understand their own organic scaling needs.

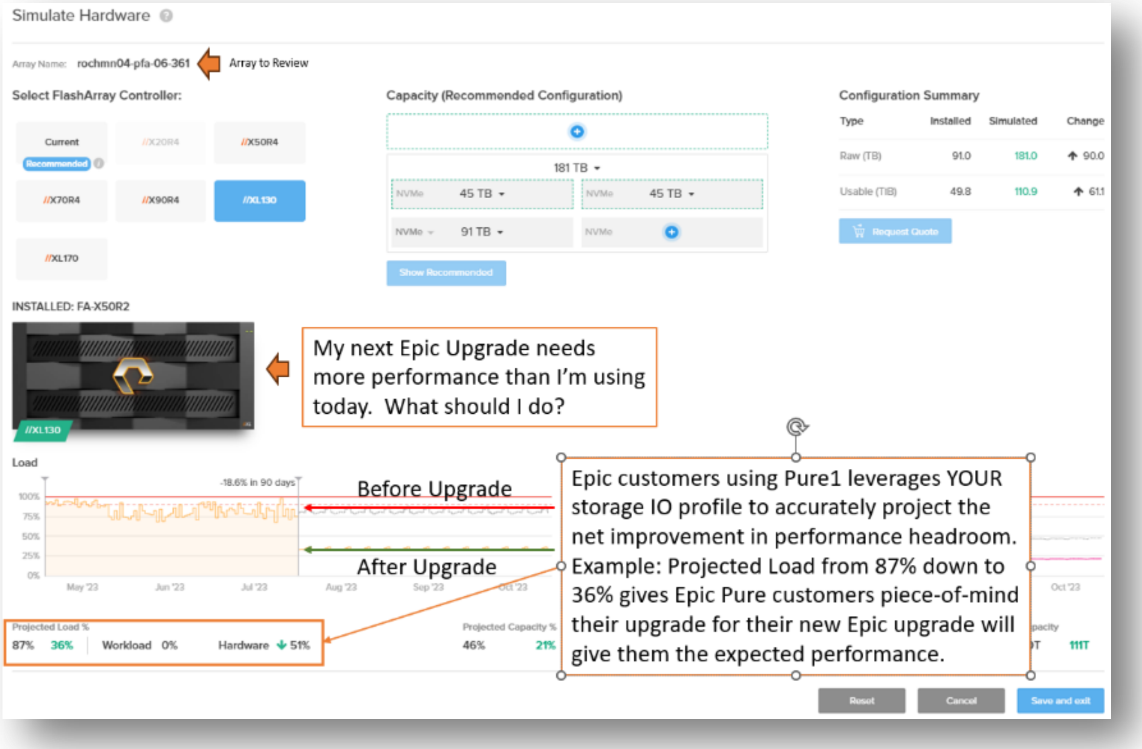

Simulate Hardware

Array Name: rochmn04-pfa-06-361 ← Array to Review

Select FlashArray Controller:

INSTALLED: FA-X50R2

Capacity (Recommended Configuration)

181 TB

NVMe 45 TB | NVMe 45 TB
NVMe 91 TB | NVMe

Configuration Summary

Type	Installed	Simulated	Change
Raw (TB)	91.0	181.0	↑ 90.0
Usable (TiB)	49.8	110.9	↑ 61.1

My next Epic Upgrade needs more performance than I'm using today. What should I do?

Before Upgrade

After Upgrade

Epic customers using Pure1 leverages YOUR storage IO profile to accurately project the net improvement in performance headroom. Example: Projected Load from 87% down to 36% gives Epic Pure customers piece-of-mind their upgrade for their new Epic upgrade will give them the expected performance.

Projected Load %: 87% 36% | Workload 0% | Hardware ↓ 51%

Projected Capacity %: 46% 21%

Capacity: T 111T

Pure Storage customers using Pure1 can determine using their own real workload performance characteristics to model their next healthcare application upgrade to ensure the storage platform model they require will do the job. Easy. No Surprises.

Pure1 Support

Pure Storage combines an ultra-proactive support team with the predictive intelligence of Pure1 Meta to deliver unrivaled support that's a key component in our proven FlashArray 99.9999 percent availability. Customers are often surprised and delighted when we fix issues they did not even know existed.

Pure1 META

The foundation of Pure1 services, Pure1 Meta is global intelligence built from a massive collection of storage array health and performance data. By continuously scanning call-home telemetry from Pure Storage's installed base, Pure1 Meta uses machine learning predictive analytics to help resolve potential issues and optimize workloads. The result is both a white glove customer support experience and breakthrough capabilities like accurate performance forecasting.

Pure1 Meta is always expanding and refining what it knows about array performance and health, moving the Data Platform toward a future of self-driving storage. This gives Epic's customers a storage product that is always evolving and improving using a global sampling of real-world telemetry.

Pure1 VM Analytics

Pure1 helps you narrow down the troubleshooting steps in your virtualized environment. VM Analytics provides you with a visual representation of the IO path from the VM all the way through to the FlashArray. Other tools and features guide you through identifying where an issue might be occurring in order to help eliminate potential candidates for a problem.

VM Analytics doesn't only help when there's a problem. The visualization allows you to identify which volumes and arrays particular applications are running on. This brings the whole environment into a more manageable domain.

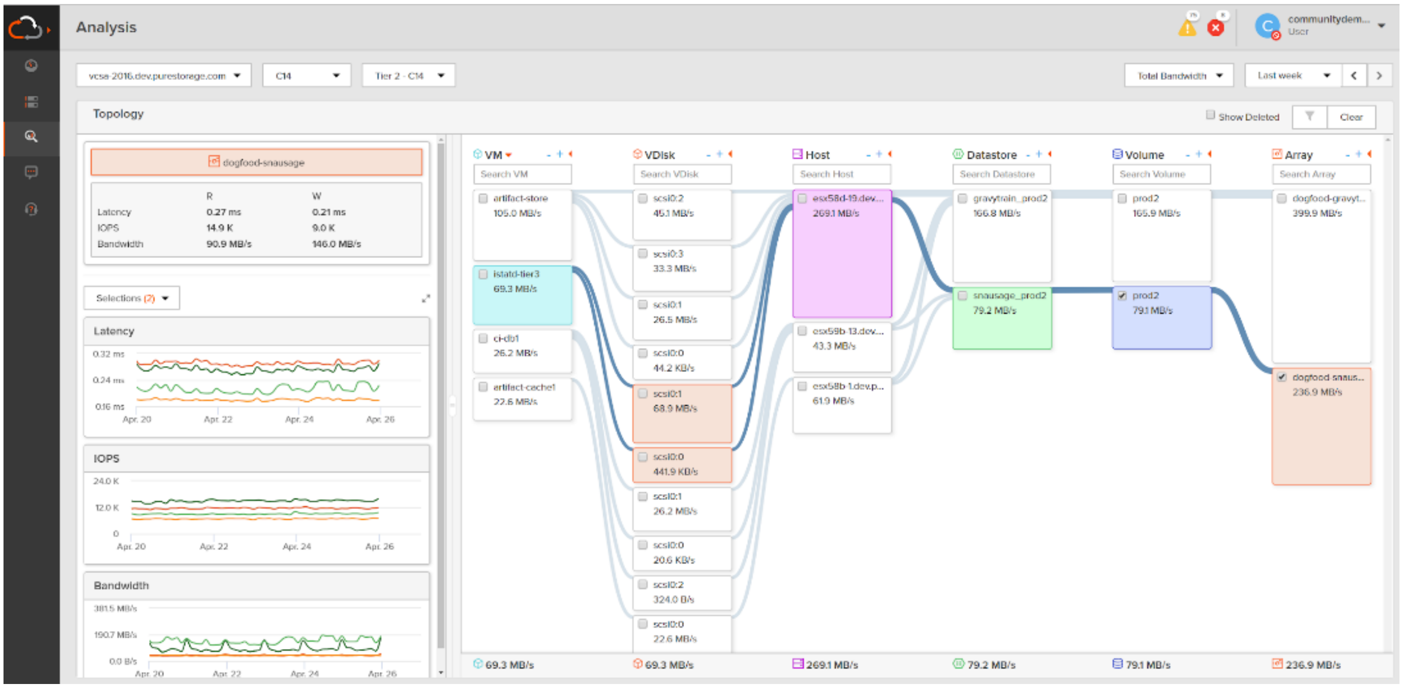

Analysis

vcsa-2016.dev.purestorage.com ▾ | C14 ▾ | Tier 2 - C14 ▾ | Total Bandwidth ▾ | Last week ▾

Topology

dogfood-snausage

	R	W
Latency	0.27 ms	0.21 ms
IOPS	14.9 K	9.0 K
Bandwidth	90.9 MB/s	146.0 MB/s

Pure1 and VM Analytics complement Epic's System Pulse tool to give customers the insight needed to manage their healthcare implementation effectively.

CloudSnap

Pure Storage portable snapshots provide simple, built-in, local and cloud protection for Pure Storage FlashArrays. Purity Snapshots enable free movement of space-efficient copies between FlashArrays, to FlashBlade, to 3rd party NFS servers, and to the cloud. Pure Storage portable snapshot technology encapsulates metadata along with data into the snapshot, making the snapshot portable, so it can be offloaded from a Pure Storage FlashArray to the cloud in a format that is recoverable to any FlashArray.

Benefits

CloudSnap is a self-backup technology built into FlashArray. It does not require the purchase of additional backup software or hardware, nor is there a need to learn and use an additional management interface. CloudSnap is natively managed via Pure Storage FlashArray's GUI, CLI, and REST interfaces and is integrated with the Pure1 Snapshot Catalog. Since FlashArray connects to AWS via https, data is encrypted in transit and stored in an encrypted format in the S3 bucket using server side encryption. Since CloudSnap was built from scratch for FlashArray, it is deeply integrated with the Purity Operating Environment, resulting in highly efficient operation. The following are a few examples of the efficiency of CloudSnap:

- CloudSnap preserves data compression on the wire, and in the S3 bucket, saving network bandwidth and increasing storage space efficiency.
- CloudSnap preserves data reduction across snapshots of a volume. After offloading the initial baseline snapshot of a volume, it only sends delta changes for subsequent snaps of the same volume. The snapshot differencing engine runs within the Purity Operating Environment in FlashArray and uses a local copy of the previous snapshot to compute the delta changes. Therefore, there is no back and forth network traffic between FlashArray and the cloud to compute deltas between snapshots, further reducing network congestion and data access costs in the cloud.
- CloudSnap knows which data blocks already exist on FlashArray, so during restores it only pulls back missing data blocks to rebuild the complete snapshot on FlashArray. In addition, CloudSnap uses dedupe preserving restores, so when data is restored from the offload target to FlashArray, it is deduped to save space on FlashArray.

The highly efficient operation of CloudSnap provides the following benefits:

- Less space is consumed in the S3 bucket
- Network utilization is minimized
- Backup windows are much smaller
- Data retrieval costs from the S3 bucket are lower

Note: Healthcare customers who would like to leverage Pure Storage's CloudSnap tool for their environments should reach out to their local presales engineering account team for more information.

InterSystems Corporation

Organizations in every industry are looking to exploit the strategic and operational benefits of shortening and eliminating the delay between event, insight, and action. They also strive to embed data-driven intelligence into their real time business processes. When successful, turning these goals into reality offers myriad benefits, including:

- Delivering new and innovative business services
- Increasing revenues Improving customer experiences
- Streamlining operations

- Identifying and decreasing risk
- Complying with new and ever-changing industry regulations
- Reducing costs

InterSystems leverages it’s highly scalable InterSystems IRIS Data Platform to healthcare customers to manage the extreme scaling needs of today’s healthcare_customers, big or small. InterSystems has a strong presence in both financial and healthcare markets including their own healthcare software.

Epic Systems

Epic is a healthcare-focused software company that develops and offers products to improve the health of people. The company’s products are used at community hospitals, academic medical centers, children's organizations, safety net providers, retail clinics, multispecialty groups, integrated delivery networks, rehab centers, and patients’ homes as well as in many other aspects. Epic Systems Corporation was established in 1979 and is based in Verona, Wisconsin. Today, Epic supports over 250 million patients across the globe for their healthcare needs.

Solution Design

This chapter contains the following:

- Storage Considerations
- Production

Storage Considerations

Boot from SAN

When utilizing Cisco UCS Server technology, it is recommended to configure Boot from SAN and store the boot partitions on remote storage, this enabled architects and administrators to take full advantage of the stateless nature of service profiles for hardware flexibility across lifecycle management of server hardware generational changes, Operating Systems/Hypervisors, and overall portability of server identity. Boot from SAN also removes the need to populate local server storage creating more administrative overhead.

Pure Storage FlashArray Considerations

Make sure Each FlashArray Controller is connected to BOTH storage fabrics (A/B).

Within Purity, it's best practice to map Hosts to Host Groups and then Host Groups to Volumes, this ensures the Volume is presented on the same LUN ID to all hosts and allows for simplified management of ESXi Clusters across multiple nodes.

How big should a Volume be? With the Purity Operating Environment, we remove the complexities of aggregates, RAID groups, and so on. When managing storage, you just create a volume based on the size required, availability and performance are taken care of through RAID-HD and DirectFlash Software. As an administrator you can create 1 10TB volume or 10 1TB Volumes and their performance/availability will be the same, so instead of creating volumes for availability or performance you can think about recoverability, manageability, and administrative considerations. For example, what data do I want to present to this application or what data do I want to store together so I can replicate it to another site/system/cloud, and so on.

IRIS volumes are required to use VMware's Raw Device Mapping (RDM) for Linux regardless of the Linux builds. Currently, only LINUX_RHEL and Ubuntu are supported for use for ODB operating systems.

Port Connectivity

10/25/40Gbe connectivity support – while both 10 and 25 Gbe is provided through 2 onboard NICs on each FlashArray controller, if more interfaces are required or if 40Gbe connectivity is also required, then make sure to provision for additional NICs have been included in the original FlashArray BOM.

32/16Gb Fiber Channel support (N-2 support) – Pure Storage offers up to 32Gb FC support on the latest FlashArray//X-Series arrays (and 64Gb FC support for FlashArray//XL). Always make sure the correct number of HBAs and the speed of SFPs are included in the original FlashArray BOM. It is preferred for implementation that all available front-end ports be for the entire arrays use. Do not fence or isolate front-end FC ports for portions of the workload this is popular in older generation storage devices.

Oversubscription

To reduce the impact of an outage or maintenance scheduled downtime it Is good practice when designing fabrics to provide oversubscription of bandwidth, this enables a similar performance profile during component failure and protects workloads from being impacted by a reduced number of paths during a component failure or maintenance event. Oversubscription can be achieved by increasing the number of physically cabled connections between storage and compute. These connections can then be utilized to deliver performance and reduced latency to the underlying workloads running on the solution.

Topology

When configuring your SAN, it’s important to remember that the more hops you have, the more latency you will see. For best performance, the ideal topology is a “Flat Fabric” where the FlashArray is only one hop away from any applications being hosted on it.

Pure Storage FlashArray Best Practices for VMware vSphere 8.0

The following Pure Storage best practices for VMware vSphere should be followed as part of a design:

- FlashArray Volumes are automatically presented to VMware vSphere using the Round Robin Path Selection Policy (PSP) and appropriate vendor Storage Array Type Plugin (SATP) for VMware vSphere 8.0.
- The Pure Storage-VMware multipathing personality is critical to have in place **prior** to associating Pure Storage volumes to VMware as RDM’s
- Multipathing is one of the most important areas to ensure that storage for Linux is optimized for the storage platform. In order to set this up, Pure Storage has a “personality” file that a simple copy-and-paste into the multipath.conf file which multipathing reads so any volumes from Pure Storage have the best-practices for storage devices presented to Linux.
- Before you configure the multipathing for Linux, backup the multipath.conf file first:

```
#cp /etc/multipath.conf /etc/multipath.conf.orig
```

- Edit the /etc/multipath.conf to contain the following personality file:

```
defaults {
        polling_interval       10
}


devices {
    device {
        vendor                      "NVME"
        product                     "Pure Storage FlashArray"
        path_selector               "queue-length 0"
        path_grouping_policy        group_by_prio
        prio                        ana
        failback                    immediate
        fast_io_fail_tmo            10
        user_friendly_names         no
        no_path_retry               0
        features                    0
        dev_loss_tmo                60
    }
    device {
        vendor                   "PURE"
        product                  "FlashArray"
        path_selector            "round-robin 0"
        hardware_handler         "1 alua"
        path_grouping_policy     group_by_prio
        prio                     alua
        failback                 immediate
        path_checker             tur
        fast_io_fail_tmo         10
        user_friendly_names      no
        no_path_retry            0
        features                 0
        dev_loss_tmo             600
    }
}
```

- An example of a proper multipathing output:

```
#multipath -ll
3624a93705a9272924e4b4d7e00011cb5 dm-16 PURE,FlashArray
size=5.0T features='0' hwhandler='1 alua' wp=rw
-+- policy='round-robin 0' prio=50 status=active
|- 12:0:17:23 sdy     65:128   active ready running
|- 13:0:17:23 sdgy    132:224  active ready running
|- 14:0:0:23  sdse    135:288  active ready running
|- 15:0:6:23  sdaop   67:1168  active ready running
|- 12:0:18:23 sdbo    68:32    active ready running
|- 13:0:18:23 sdkd    66:272   active ready running

This output will result back the settings for Multipathing.  Policy='round-robin 0' and hwhandler='1 alua'
status=active, and PURE, FlashArray all indicate that the multipathing is properly set up.

The sdy, sdgy, and more all referred to as the "backing" devices.  This are logical devices created by the
multipathing subsystem to allow IO to flow across each device to ensure even IO flow control and high-
availability at the LVM layer.
```

- **This is a critical step:** The number of volumes you want to present from the Pure Storage FlashArray to ESX and the VM guest follows general guidelines. The guideline to follow is 2 volumes for APP files, 2 volumes for Journaling, and 8-16 volumes for the ODB IRIS.dat files. See Appendix D – example ODB LVM creation steps.
- The logical volume creation step is a critical performance-impacting configuration. This is described in detail in the Epic ODB guide and should be reviewed. Appendix D outlines these steps in detail. It is important to ensure you are properly creating the striped logical volume based on the number of volumes for the logical volumes you are creating.
- VMware vSphere 8.0 also uses the Latency SATP that was introduced in VMware vSphere 6.7U1 (this replaces the I/O Operations Limit of 1 SATP, which was the default from VMware vSphere 6.5U1).
- When using iSCSI connected Pure Storage FlashArray volumes, it is recommended to set DelayedAck to false (disabled) and LoginTimeout to 30 seconds. Jumbo Frames are optional when using iSCSI. Using iSCSI requires a rigorous battery of test scenarios that focuses on resiliency and performance.

Note: iSCSI is not a supported storage connectivity option – Fiber Channel is best suited for this.

- For VMFS-6, keep automatic UNMAP enabled.
- DataMover.HardwareAcceleratedMove, DataMover.HardwareAcceleratedInit, and VMFS3.HardwareAcceleratedLocking should all be enabled.
- Ensure all ESXi hosts are connected to both Pure Storage FlashArray controllers. A minimum of two paths to each. Aim for total redundancy.
- Install VMware tools or Open VM tools whenever possible.
- The General volume layout for both VMware and the LINUX guest breakdown is below. Its very important that Raw Device Mapping (RDM) is the best-practice for the use-cases.

Production

Our best practices for storage layout still apply (see the Epic Hardware Configuration Guide and the Storage Area Network Considerations whitepaper, which is available on the Epic UserWeb for details). For specific LVM configuration, see the Epic Storage Configuration Quick-Reference document.

Based on considerations for backups, use the following layout for production:

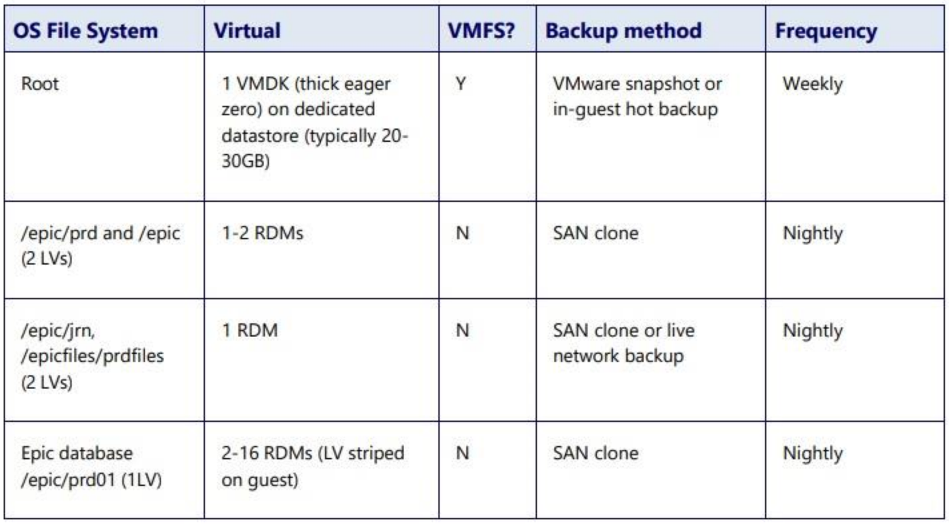

OS File System	Virtual	VMFS?	Backup method	Frequency
Root	1 VMDK (thick eager zero) on dedicated datastore (typically 20-30GB)	Y	VMware snapshot or in-guest hot backup	Weekly
/epic/prd and /epic (2 LVs)	1-2 RDMs	N	SAN clone	Nightly
/epic/jrn, /epicfiles/prdfiles (2 LVs)	1 RDM	N	SAN clone or live network backup	Nightly
Epic database /epic/prd01 (1LV)	2-16 RDMs (LV striped on guest)	N	SAN clone	Nightly

- RDMs should be physical or set to independent persistent to avoid being a snapshot with the rest of the VM.
- RDMs are used rather than VMFS/VMDKs so backups can still be taken with storage level clones. VMware snapshots are not feasible for this due to the change rate of the production database storage.
- RDMs should have VMDK pointer files stored on the datastore with the VM.
- Queue depths should be left at the default. Changing queue depths on the ESXi host is a tweak and should only be examined if a performance problem (high latency) is observed. A value of 32 is a good starting point. With more paths and/or more bandwidth per port, this number becomes less critical but 32 is a good number to start with.
- When mounting snapshots, use the ESXi resignature option and avoid force-mounting.
- Configure Host Groups on the FlashArray identically to clusters in VMware vSphere. For example, if a cluster has four hosts in it, create a corresponding Host Group on the relevant Pure Storage FlashArray with exactly those four hosts—no more, no less.
- When possible, use Paravirtual SCSI adapters for virtual machines. For Epic, they are requesting pSCSI.

SCSI Device Configuration

The following is the SCSI device configuration:

- Use pvscsi devices
- Utilize all four SCSI devices for optimal performance

  Suggested layout:

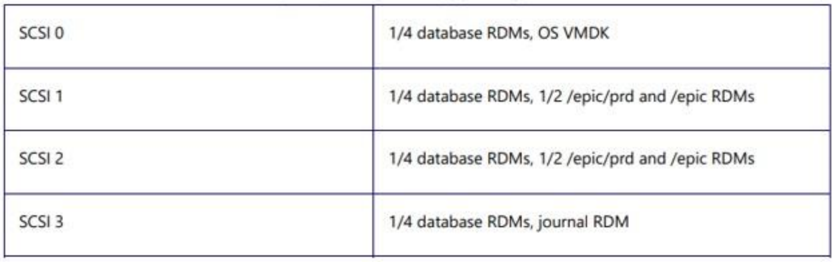

SCSI 0	1/4 database RDMs, OS VMDK
SCSI 1	1/4 database RDMs, 1/2 /epic/prd and /epic RDMs
SCSI 2	1/4 database RDMs, 1/2 /epic/prd and /epic RDMs
SCSI 3	1/4 database RDMs, journal RDM

- Atomic Test and Set (ATS) is required on all Pure Storage volumes. This is a default configuration, and no changes should normally be needed.

Deployment Hardware and Software

This chapter contains the following:

- Architecture
- Products Deployed
- Physical Topology
- Configuration Guidelines

Architecture

This FlashStack architecture delivers an healthcare infrastructure that is redundant and uses the best practices of Cisco and Pure Storage.

It includes:

- VMware vSphere 8.0 hypervisor installed on the Cisco UCS X210c M7 compute nodes configured for stateless compute design using boot from SAN.
- Pure Storage FlashArray//X R3 provides the storage infrastructure required for VMware vSphere hypervisors and the workload.
- Cisco Intersight provides UCS infrastructure management with lifecycle management capabilities.

The architecture deployed is highly modular. While each customer's environment might vary in its exact configuration, the reference architecture contained in this document once built, can easily be scaled as requirements, and demands change. This includes scaling both up (adding additional resources within a Cisco UCS Domain) and out (adding additional Cisco UCS Domains and Pure Storage).

Pure Storage takes a different approach to providing storage products for healthcare customers. Because of application dynamics, Pure Storage offers three options for storage designs as described in the following sections.

Products Deployed

- VMware vSphere ESXi 8.0
- VMware vCenter 8.0 to set up and manage the virtual infrastructure as well as integration of the virtual environment with Cisco Intersight software
- RHEL 8.7
- Cisco Intersight platform to deploy, maintain, and support the FlashStack components
- Cisco Intersight Assist virtual appliance to help connect the Pure Storage FlashArray and VMware vCenter with the Cisco Intersight platform

Physical Topology

FlashStack with Cisco UCS X-Series Modular System is a Fibre Channel (FC) based storage access design. Pure Storage FlashArray and Cisco UCS are connected through Cisco MDS 9132T switches and storage access utilizes the FC network. For IP-based file share storage access Pure Storage FlashArray and Cisco UCS are connected through Cisco Nexus C93180YC-FX switches.

Figure 14 details the physical hardware and cabling deployed to enable this solution:

- Two Cisco Nexus 93180YC-FX Switches in NX-OS Mode.

- Two Cisco MDS 9132T 32-Gb Fibre Channel Switches.
- One Cisco UCS X9508 Chassis with two Cisco UCSX 9108 25G IF Modules.
- One Cisco UCS X210c M7 Compute Nodes with Intel(R) Xeon(R) Platinum 8452Y CPU 2.0GHz 36-core processors, 256GB 4800MHz RAM, and one Cisco UCS VIC 15420 mezzanine card, providing N+1 server fault tolerance.
- Pure Storage FlashArray//X R3 with dual redundant controllers, with 10 1.92TB DirectFlash NVMe drives.

Figure 14. **FlashStack – Physical Topology for FC**

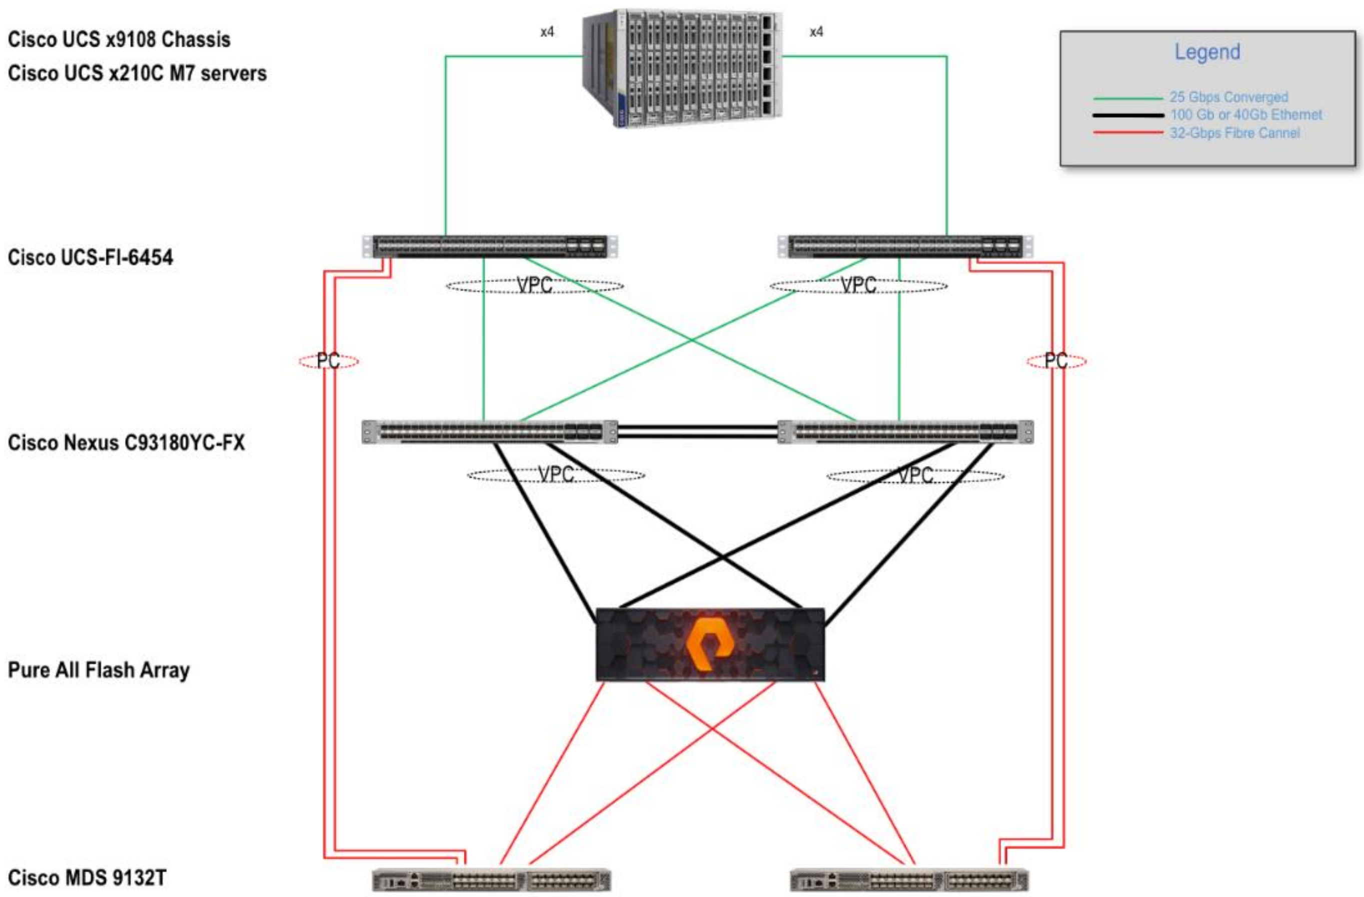

[Table 1](#) lists the software versions of the primary products installed in the environment.

Table 1. Software and Firmware Versions

Vendor	Product/Component	Version/Build/Code
Cisco	UCS Component Firmware	5.1
Cisco	UCS x210c Compute Node	5.0(2e)
Cisco	VIC 15420	5.2(2e)
Cisco	Cisco Nexus 93180YC-FX	9.3(3)

Vendor	Product/Component	Version/Build/Code
Cisco	Cisco MDS 9132T	8.4(2d)
Pure Storage	FlashArray//X R3	Purity//FA 6.3.7
VMware	vCenter Server Appliance	8.0.0 Build:20395099
VMware	vSphere 8.0	8.0.0.20513097
Cisco	Intersight Assist	1.0.11-759
VMware	Tools	12325

Configuration Guidelines

The EHR solution described in this document provides details for configuring a fully redundant, highly-available configuration. Configuration guidelines are provided that refer to which redundant component is being configured with each step, whether that be A or B. For example, Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are configured similarly.

Note: This document is intended to allow the reader to configure the customer environment as a stand-alone solution.

VLANs

The VLAN configuration recommended for the environment includes a total of six VLANs as listed in Table 2.

Table 2. VLANs Configured in this study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
FS-InBand-Mgmt_40	40	In-Band management interfaces
FS-Infra-Mgmt_41	41	Infrastructure Virtual Machines
FS-_42	42	VM Data
FS-vMotion_43	43	VMware vMotion
OOB-Mgmt	164	Out of Band management interfaces

VSANs

Table 3 lists the two virtual SANs that were configured for communications and fault tolerance in this design.

Table 3. VSANs Configured in this study

VSAN Name	VSAN ID	VSAN Purpose
VSAN 700	700	VSAN for Primary SAN communication
VSAN 701	701	VSAN for Secondary SAN communication

Solution Configuration

This chapter contains the following:

- Solution Cabling

Solution Cabling

This section details the physical connectivity configuration of the FlashStack VMware environment.

The information provided in this section is a reference for cabling the physical equipment in this Cisco Validated Design environment.

Note: 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: Be sure to follow the cabling directions in this section. Failure to do so will result in problems with your deployment.

Figure 15 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 Storage FlashArray//X R3 controllers, four of these have been used for SCSI-fc and the other four to support NVMe-fc. Also, 25Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the Pure Storage 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.

Figure 15. **FlashStack solution cabling diagram**

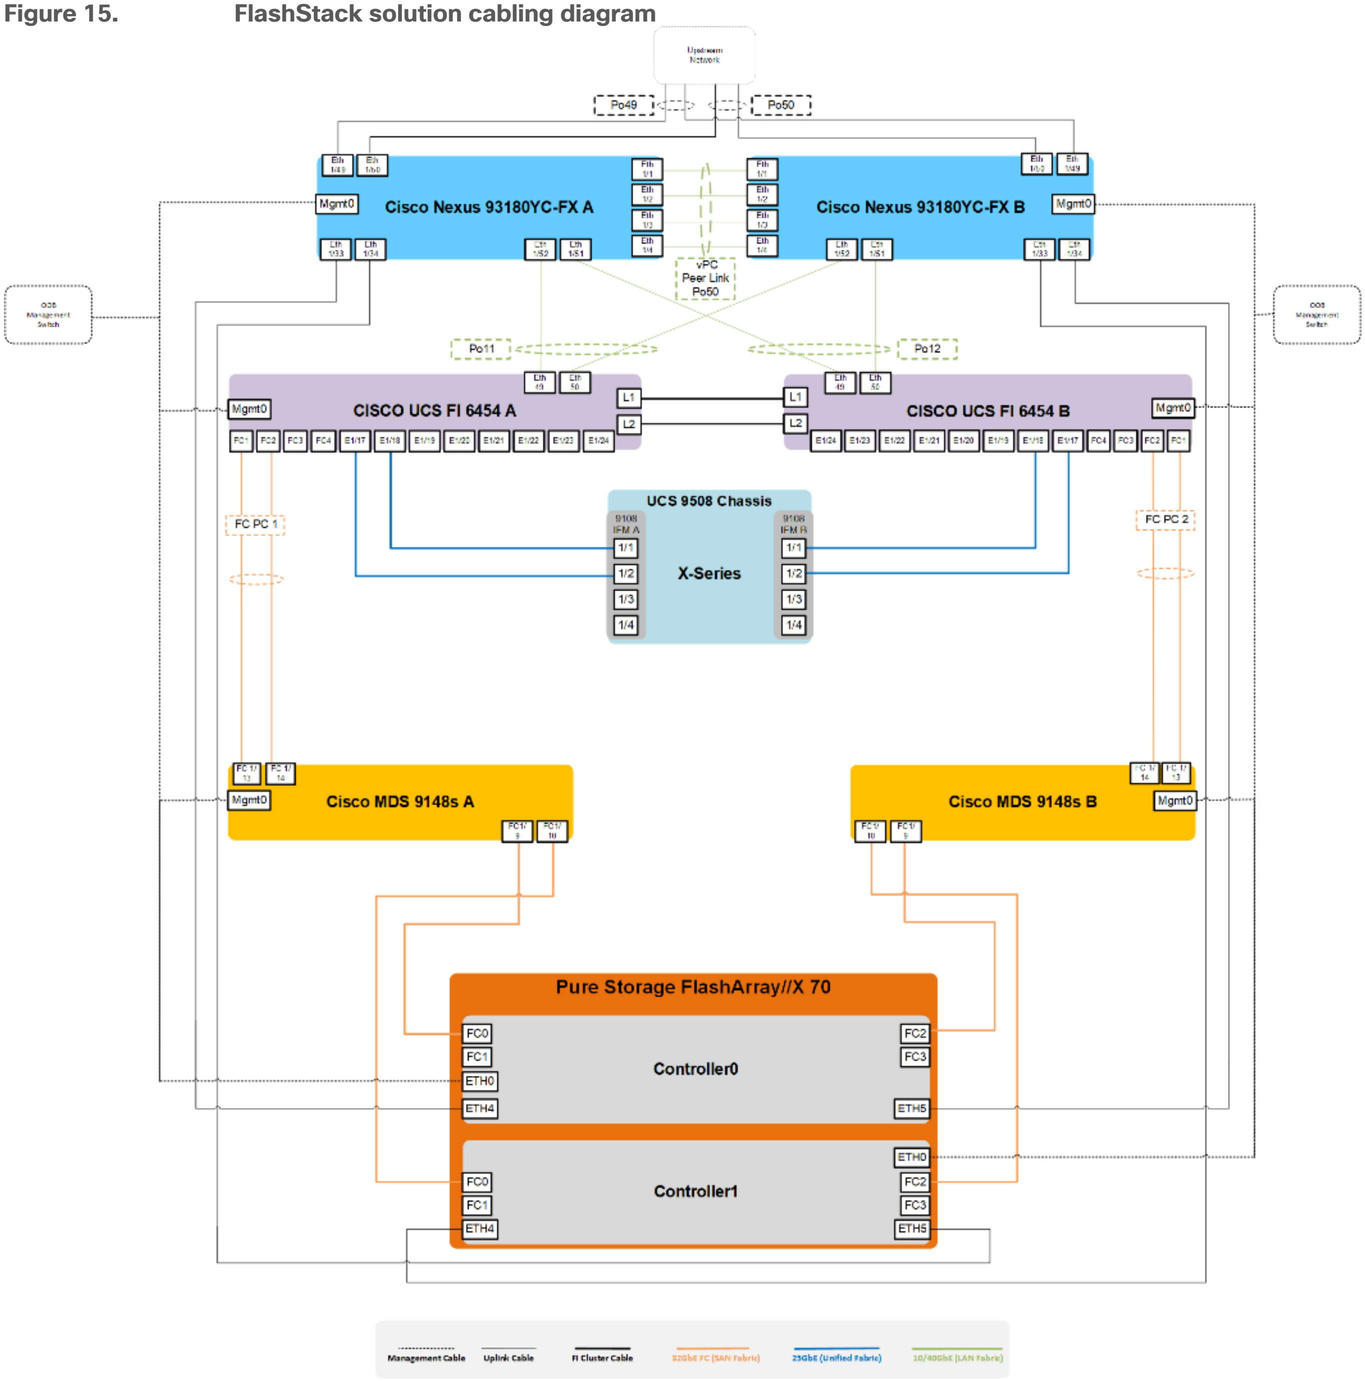

Configuration and Installation

This chapter contains the following:

- FlashStack Automated Deployment with Ansible
- FlashStack Manual Deployment
- Cisco UCS X-Series Configuration – Intersight Managed Mode (IMM)
- Cisco MDS 9132T 32-Gb FC Switch Configuration
- Pure Storage FlashArray//X R3 to MDS SAN Fabric Connectivity
- Configure Pure Storage FlashArray//X R3
- Install and Configure VMware ESXi 8.0
- Cisco Intersight Orchestration
- Supported LINUX Hosts for Operations Database (ODB)
- Safe Mode Considerations
- Purity Code Updating for Healthcare Customers

FlashStack Automated Deployment with Ansible

This repository contains Ansible playbooks to configure all the components of FlashStack including:

- Cisco UCS in Intersight Managed Mode (IMM)
- Cisco Nexus and MDS Switches
- Pure Storage FlashArray
- VMware ESXi and VMware vCenter

FlashStack Manual Deployment

Cisco Intersight Managed Mode standardizes policy and operation management for Cisco UCS X-Series. The compute nodes in Cisco UCS X-Series are configured using server profiles defined in Cisco Intersight. These server profiles derive all the server characteristics from various policies and templates. At a high level, configuring Cisco UCS using Cisco Intersight Managed Mode consists of the steps shown in Figure 16.

Figure 16. **Configuration Steps for Cisco Intersight Managed Mode**

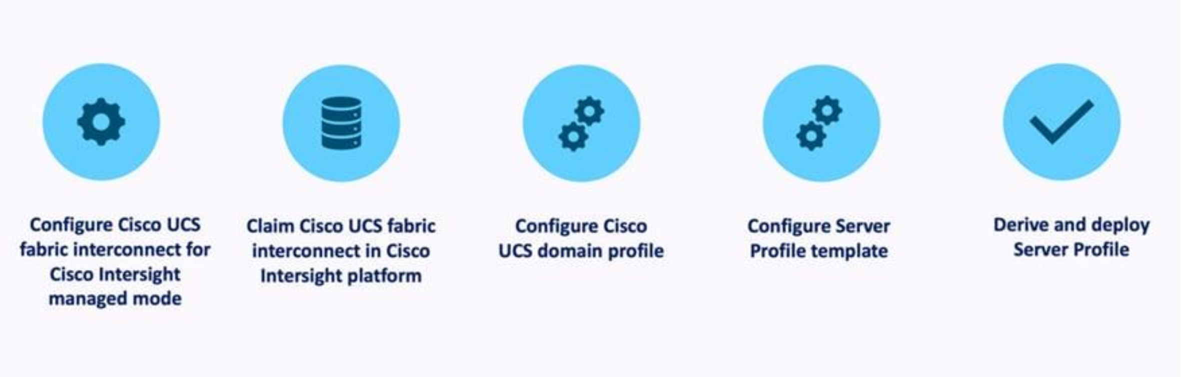

Configure Cisco UCS fabric interconnect for Cisco Intersight managed mode
Claim Cisco UCS fabric interconnect in Cisco Intersight platform
Configure Cisco UCS domain profile
Configure Server Profile template
Derive and deploy Server Profile

Cisco UCS X-Series Configuration - Intersight Managed Mode (IMM)

Procedure 1. Configure Cisco UCS Fabric Interconnects for IMM

Step 1. Verify the following physical connections on the fabric interconnect:

- The management Ethernet port (mgmt0) is connected to an external hub, switch, or router
- The L1 ports on both fabric interconnects are directly connected to each other
- The L2 ports on both fabric interconnects are directly connected to each other

Step 2. Connect to the console port on the first Fabric Interconnect.

Step 3. Configure Fabric Interconnect A (FI-A). On the Basic System Configuration Dialog screen, set the management mode to Intersight. All the remaining settings are similar to those for the Cisco UCS Manager managed mode (UCSM-Managed).

Cisco UCS Fabric Interconnect A

Procedure 1. Configure the Cisco UCS for use in Intersight Managed Mode

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

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

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

  You have chosen to setup a new Fabric interconnect in “intersight” 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>

  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>

  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>
<SNIP>

  Verify and save the configuration.
```

Step 2. After applying the settings, make sure you can ping the fabric interconnect management IP address. When Fabric Interconnect A is correctly set up and is available, Fabric Interconnect B will automatically discover Fabric Interconnect A during its setup process as shown in the next step.

Step 3. Configure Fabric Interconnect B (FI-B). For the configuration method, choose console. Fabric Interconnect B will detect the presence of Fabric Interconnect A and will prompt you to enter the admin password for Fabric Interconnect A. Provide the management IP address for Fabric Interconnect B and apply the configuration.

```
Cisco UCS Fabric Interconnect B
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>

    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
```

Procedure 2. Claim a Cisco UCS Fabric Interconnect in the Cisco Intersight Platform

If you do not already have a Cisco Intersight account, you need to set up a new account in which to claim your Cisco UCS deployment. Start by connecting to https://intersight.com.

All information about Cisco Intersight features, configurations can be accessed in the Cisco Intersight Help Center.

Step 1. Click Create an account.

Step 2. Sign in with your Cisco ID.

Step 3. Read, scroll through, and accept the end-user license agreement. Click Next.

Step 4. Enter an account name and click Create.

If you have an existing Cisco Intersight account, connect to https://intersight.com and sign in with your Cisco ID, select the appropriate account.

Note: In this step, a Cisco Intersight organization is created where all Cisco Intersight managed mode configurations including policies are defined.

Step 5. Log into the Cisco Intersight portal as a user with account administrator role.

Step 6. From the Service Selector drop-down list, select System.

Step 7. Navigate to Settings > General > Resource Groups.

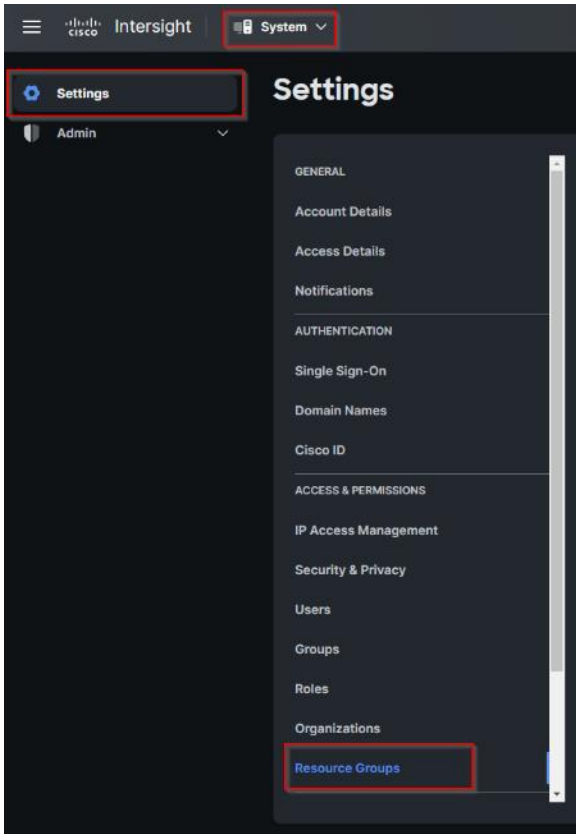

Step 8. On Resource Groups panel click + Create Resource Group.

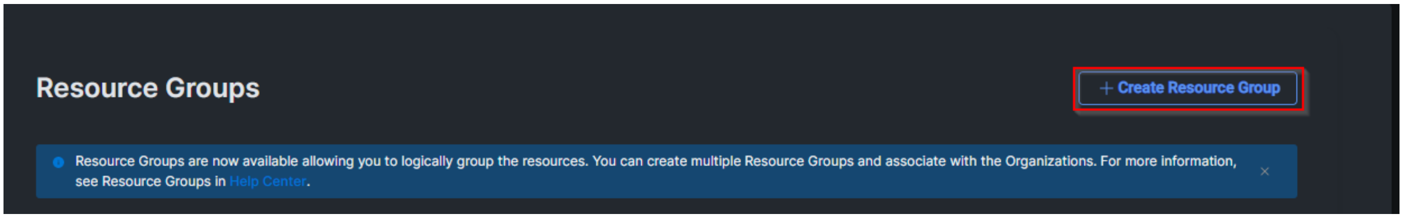

Resource Groups are now available allowing you to logically group the resources. You can create multiple Resource Groups and associate with the Organizations. For more information, see Resource Groups in Help Center.

Step 9. Provide a name for the Resource Group (for example, FlashStack-L151-DMZ).

Step 10. Click Create.

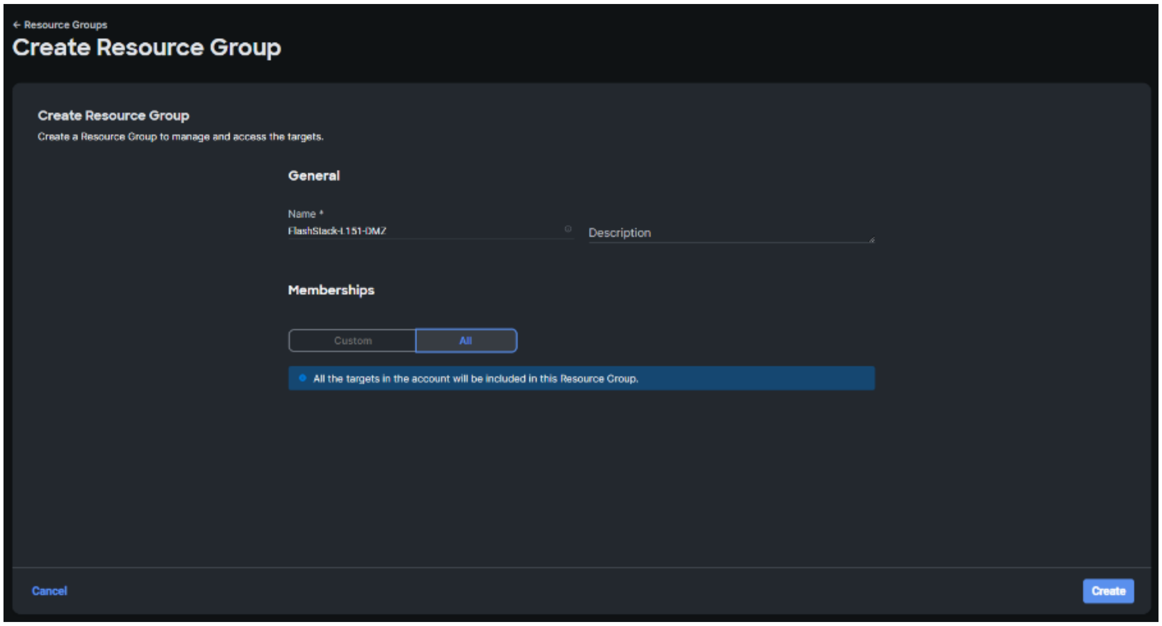

Create a Resource Group to manage and access the targets.

Name * FlashStack-L151-DMZ

Description

Memberships: All

All the targets in the account will be included in this Resource Group.

Step 11. Navigate to Settings > General > Organizations.

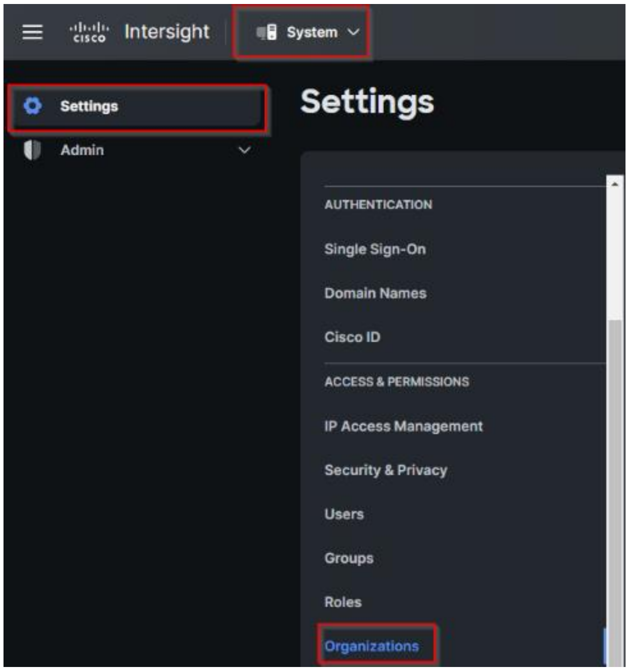

Step 12. On Organizations panel click + Create Organization.

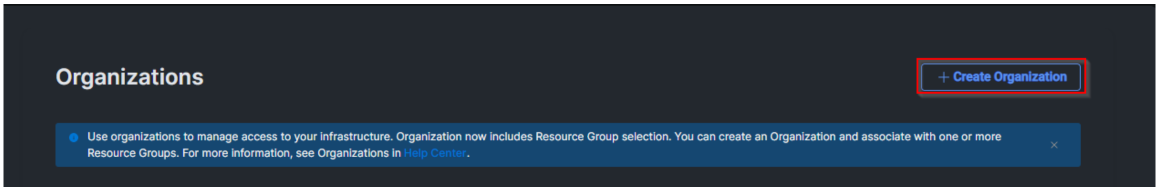

Use organizations to manage access to your infrastructure. Organization now includes Resource Group selection. You can create an Organization and associate with one or more Resource Groups. For more information, see Organizations in Help Center.

Step 13. Provide a name for the organization (FlashStack).

Step 14. Select the Resource Group created in the last step (for example, FlashStack-L151-DMZ).

Step 15. Click Create.

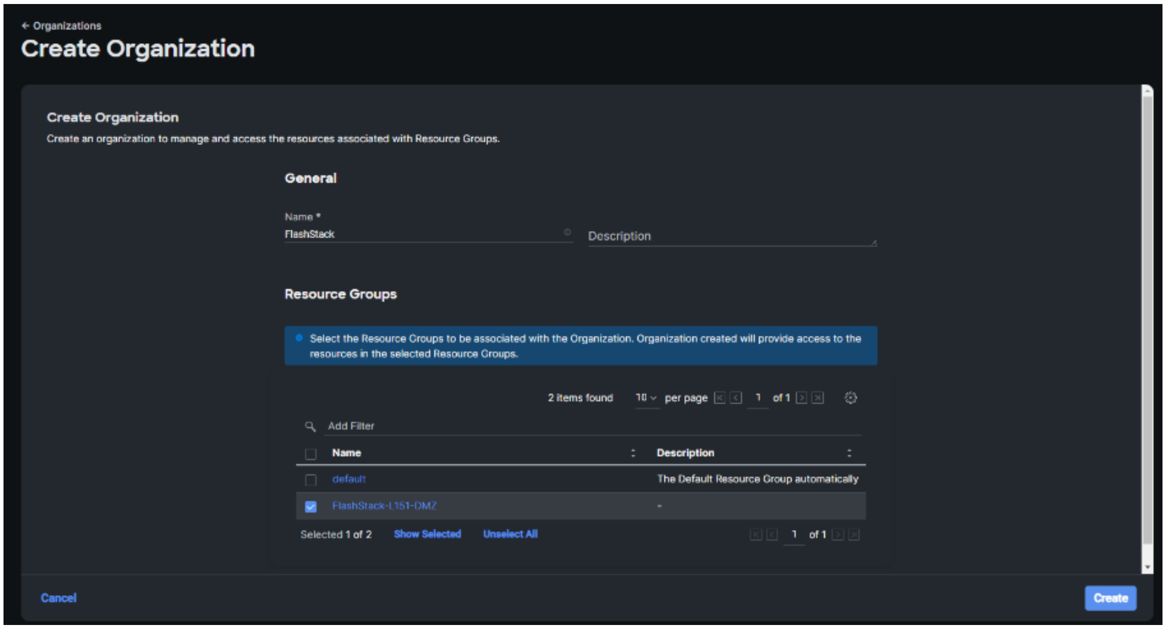

← Organizations

Create Organization

Create Organization

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

General

Name * FlashStack

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.

2 items found 10 per page 1 of 1

Add Filter

☐	Name	Description
☐	default	The Default Resource Group automatically
☑	FlashStack-L151-DMZ	-

Selected 1 of 2 Show Selected Unselect All 1 of 1

Step 16. 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.

Step 17. Under DEVICE CONNECTOR, the current device status will show “Not claimed.” Note, or copy, the Device ID, and Claim Code information for claiming the device in Cisco Intersight.

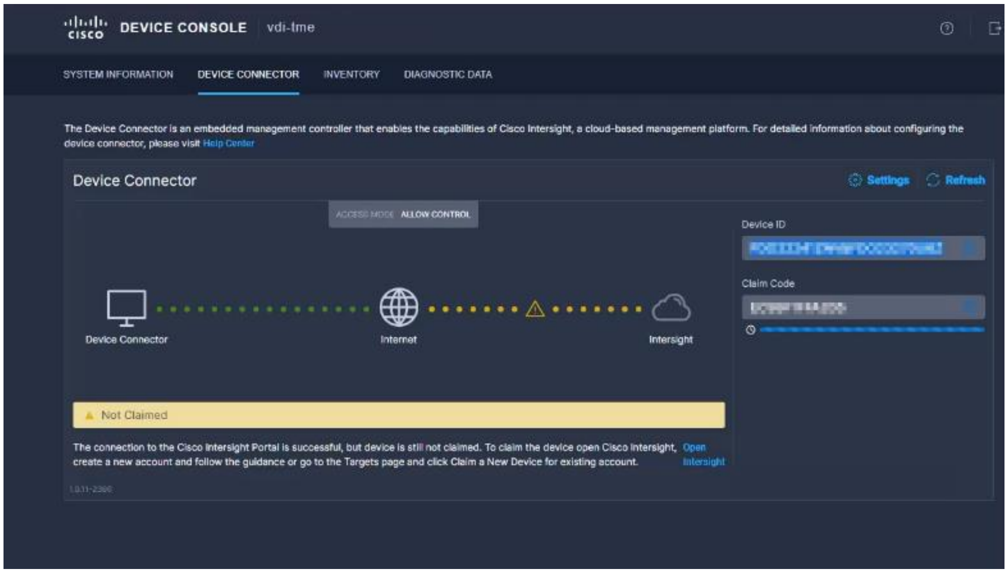

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

Device Connector

ACCESS MODE ALLOW CONTROL

Device ID [redacted]

Claim Code [redacted]

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 Targets page and click Claim a New Device for existing account. Open Intersight

1.0.11-2300

Step 18. Navigate to Admin > General > Targets.

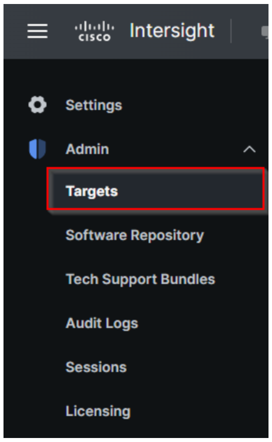

Step 19. On the Targets panel, click Claim a New Target.

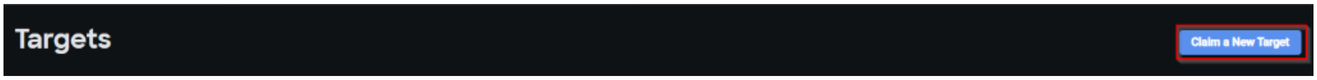

Step 20. Select Cisco UCS Domain (Intersight Managed) and click Start.

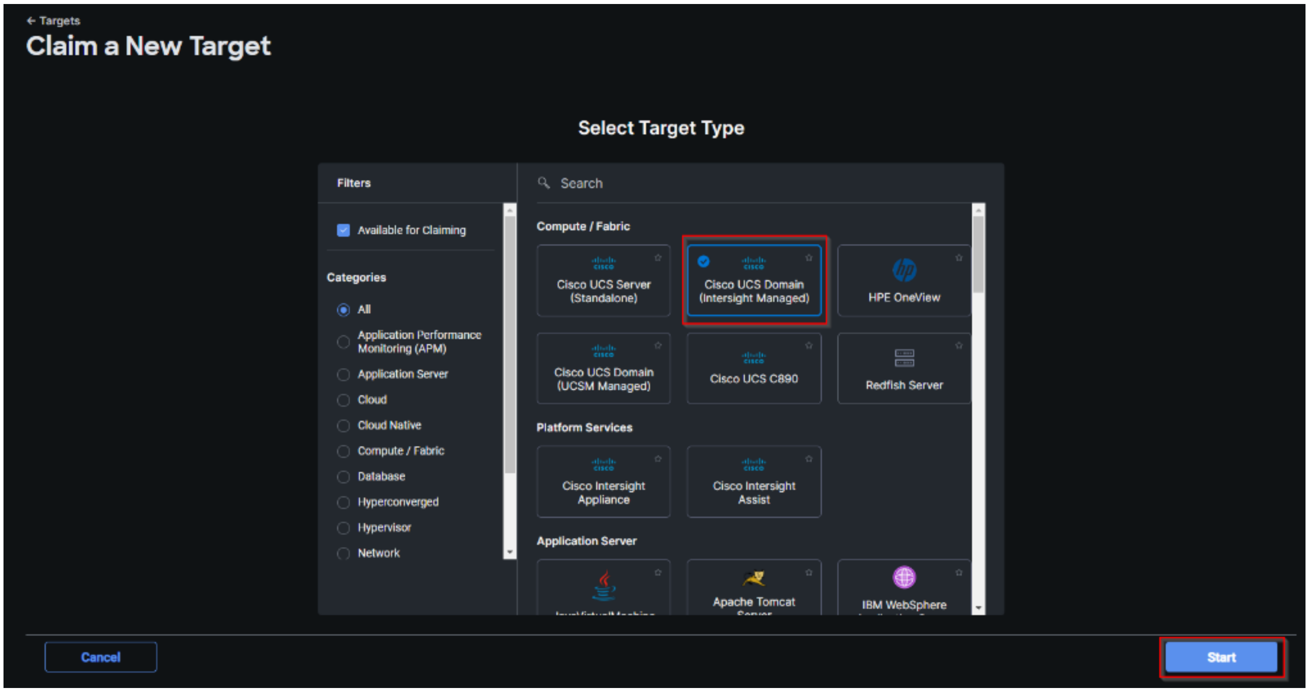

☑ Available for Claiming

● All

Step 21. Enter the Device ID and Claim Code captured from the Cisco UCS FI.

Step 22. Select the previously created Resource Group and click Claim.

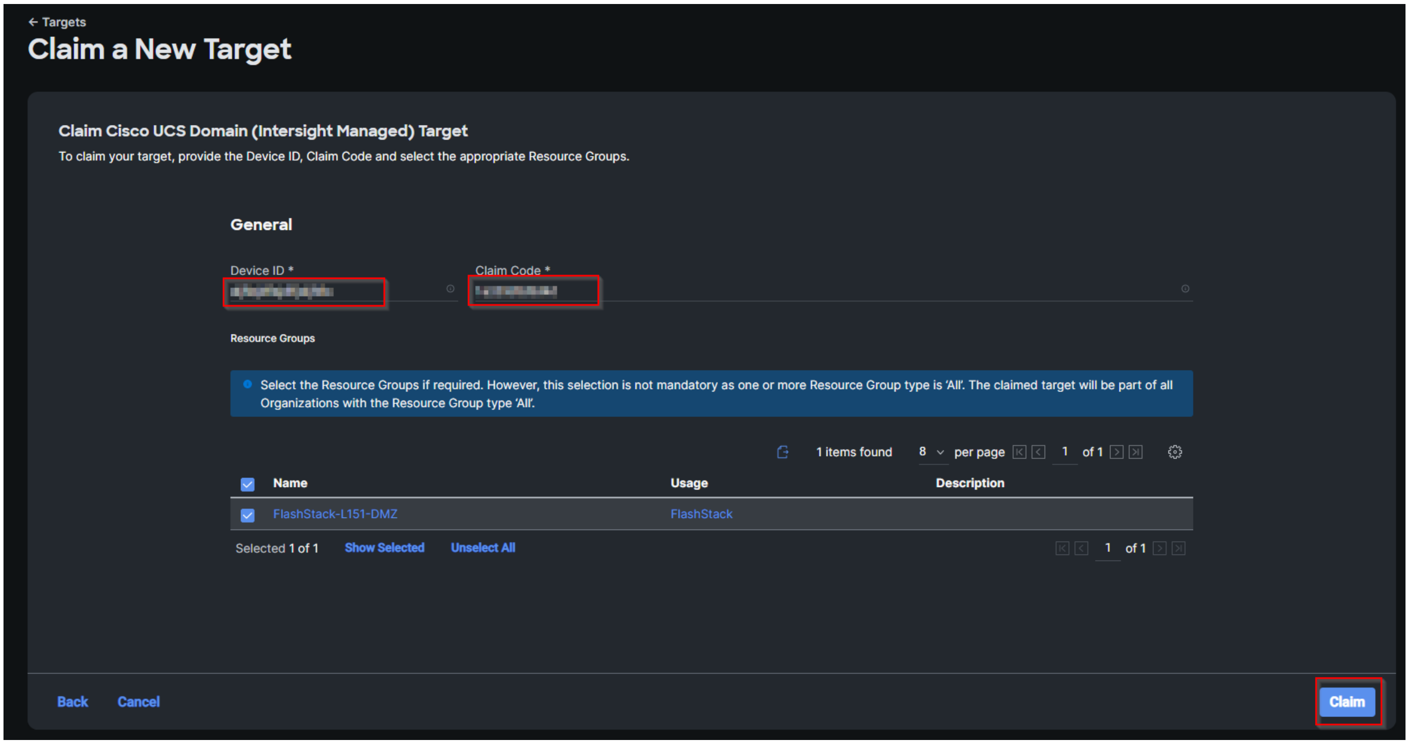
← Targets

Claim a New Target

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 * [redacted]

Claim Code * [redacted]

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'.

1 items found 8 per page 1 of 1

☑	Name	Usage	Description
☑	FlashStack-L151-DMZ	FlashStack	

Selected 1 of 1 Show Selected Unselect All

Step 23. On successfully device claim, Cisco UCS FI should appear as a target in Cisco Intersight.

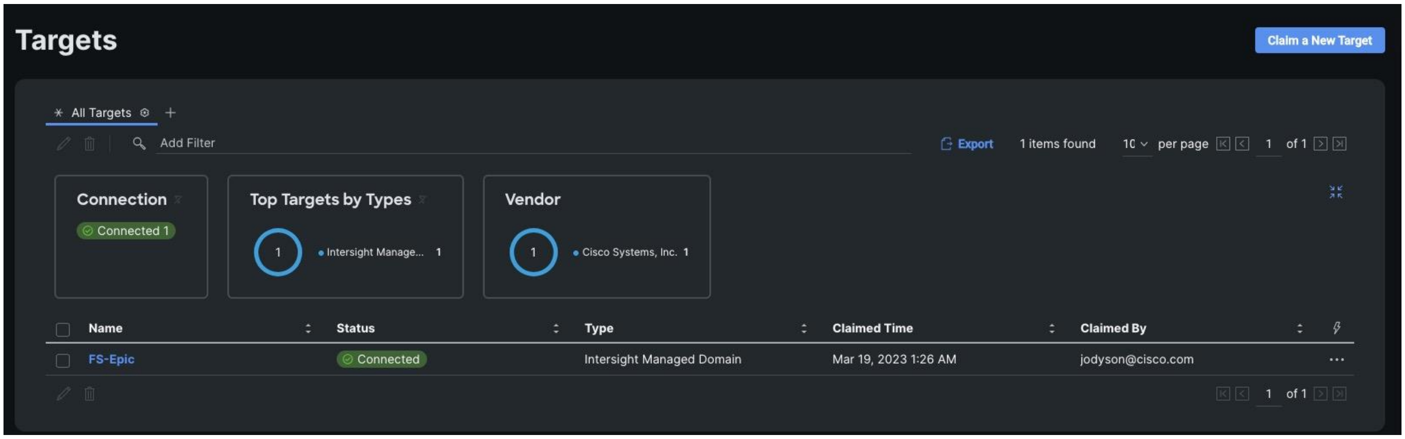
Targets

1 items found 10 per page 1 of 1

Connection: Connected 1

Top Targets by Types: 1 — Intersight Manage... 1

Vendor: 1 — Cisco Systems, Inc. 1

☐	Name	Status	Type	Claimed Time	Claimed By
☐	FS-Epic	Connected	Intersight Managed Domain	Mar 19, 2023 1:26 AM	jodyson@cisco.com

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 configured 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.

After the Cisco UCS domain profile has been successfully created and deployed, the policies including the port policies are pushed to Cisco UCS Fabric Interconnects. Cisco UCS domain profile can easily be cloned to install additional Cisco UCS systems. When cloning the UCS domain profile, the new UCS domains utilize the existing policies for consistent deployment of additional Cisco UCS systems at scale.

Procedure 1. Create a Domain Profile

Step 1. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Configure > Profiles, to launch the Profiles Table view.

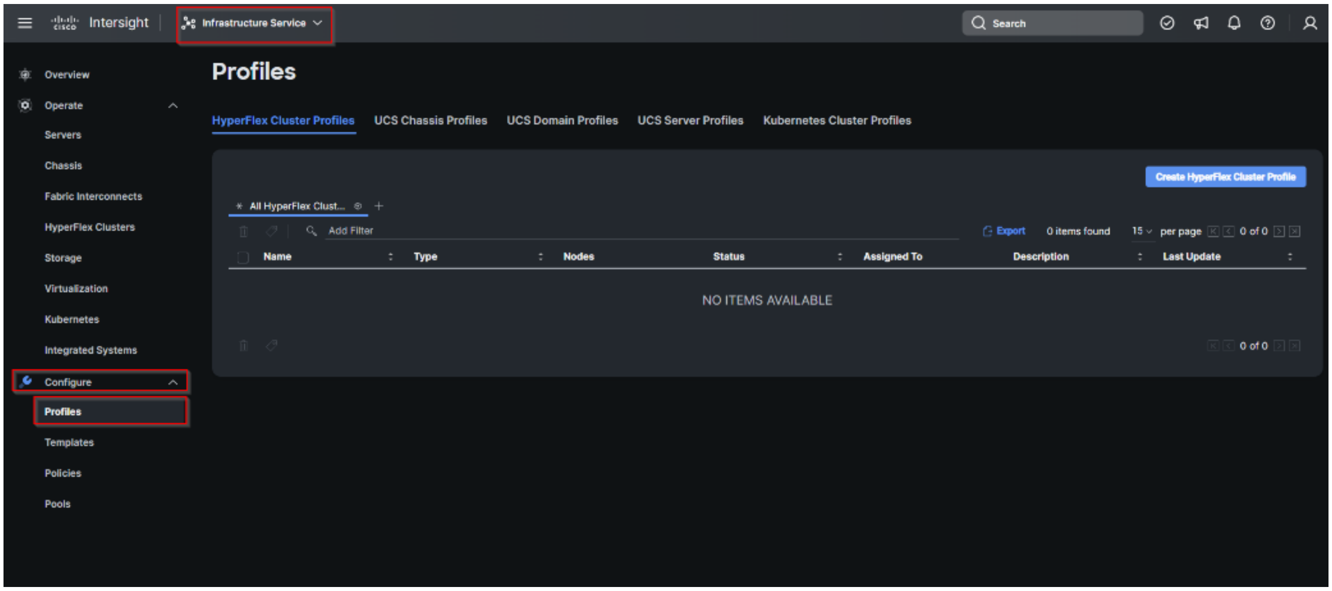

Step 2. Navigate UCS Domain Profiles tab and click Create UCS Domain Profile.

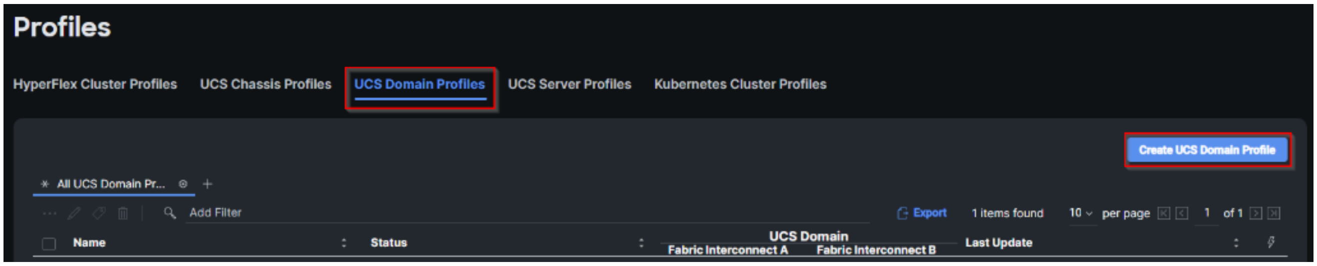

Step 3. On the Create UCS Domain Profile screen, click Start.

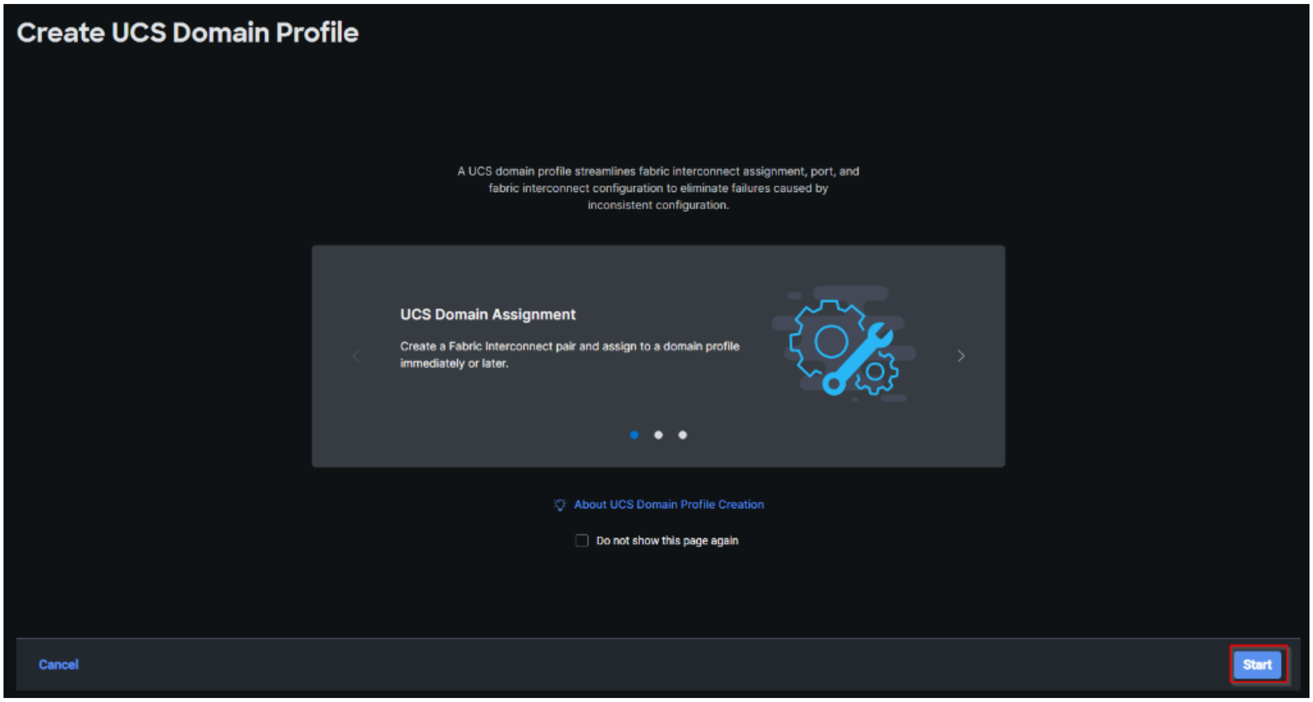

Create UCS Domain Profile

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
Cancel
Start

Step 4. On the General page, select the organization created before and enter a name for your profile (for example, FS-L152-DMZ-K4). Optionally, include a short description and tag information to help identify the profile. Tags must be in the key:value format. For example, Org: IT or Site: APJ. Click Next.

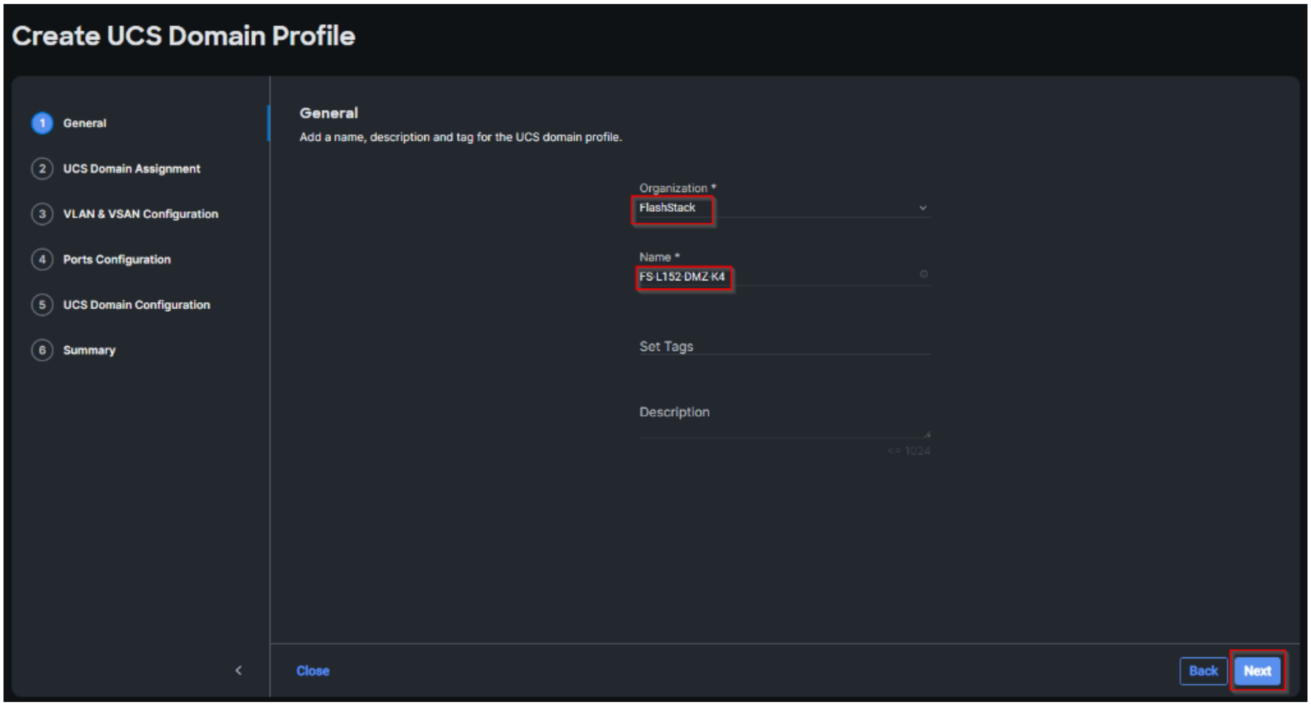

Create UCS Domain Profile

General

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

Organization * FlashStack

Name * FS-L152-DMZ-K4

Set Tags

Description

Step 5. On the Domain Assignment page, assign a switch pair to the Domain profile. Click Next.

Note: You can also click Assign Later and assign a switch pair to the Domain profile at a later time.

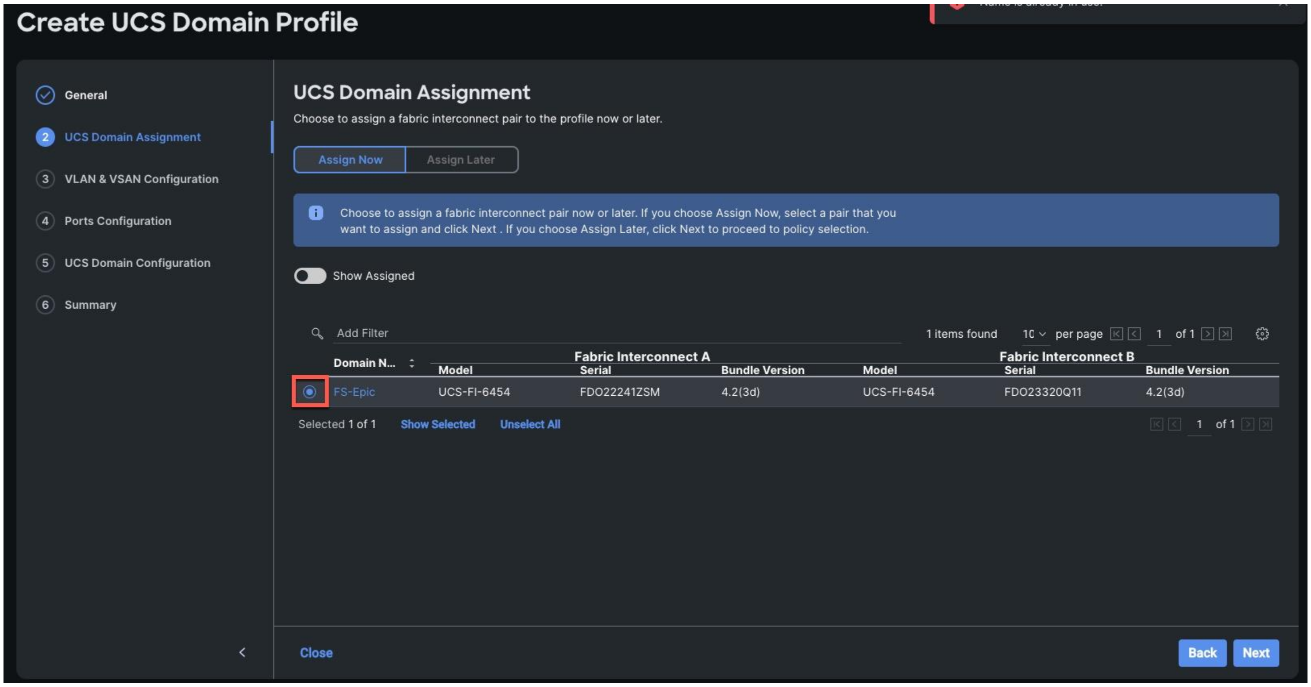

Create UCS Domain Profile

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.

		Fabric Interconnect A			Fabric Interconnect B	
Domain N...	Model	Serial	Bundle Version	Model	Serial	Bundle Version
● FS-Epic	UCS-FI-6454	FDO22241ZSM	4.2(3d)	UCS-FI-6454	FDO23320Q11	4.2(3d)

Step 6. On the VLAN & VSAN Configuration page, attach VLAN and VSAN policies for each switch to the UCS Domain Profile.

Note: In this step, a single VLAN policy is created for both fabric interconnects and two individual VSAN policies are created because the VSAN IDs are unique for each fabric interconnect.

Step 7. Click Select Policy next to VLAN Configuration under Fabric Interconnect A.

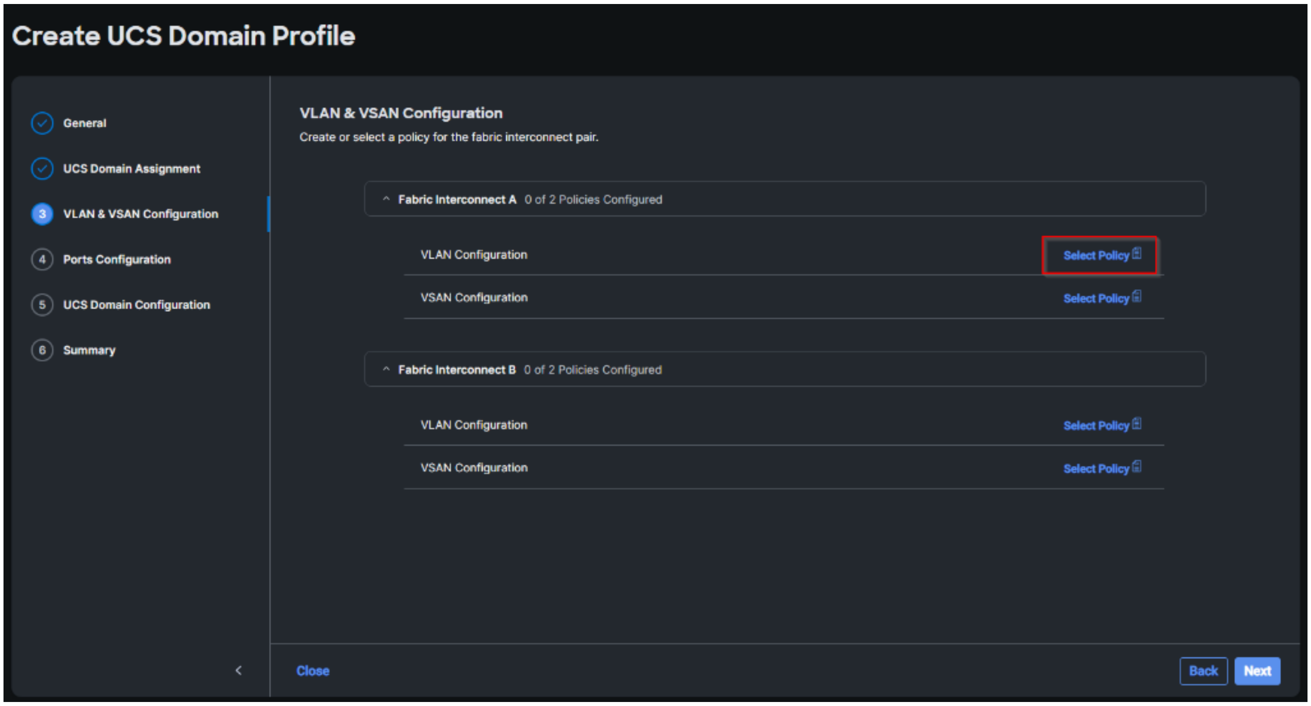

Create UCS Domain Profile

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 0 of 2 Policies Configured

- VLAN Configuration — Select Policy
- VSAN Configuration — Select Policy

Fabric Interconnect B 0 of 2 Policies Configured

- VLAN Configuration — Select Policy
- VSAN Configuration — Select Policy

Step 8. In the pane on the right, click Create New.

Step 9. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-VLAN). Click Next.

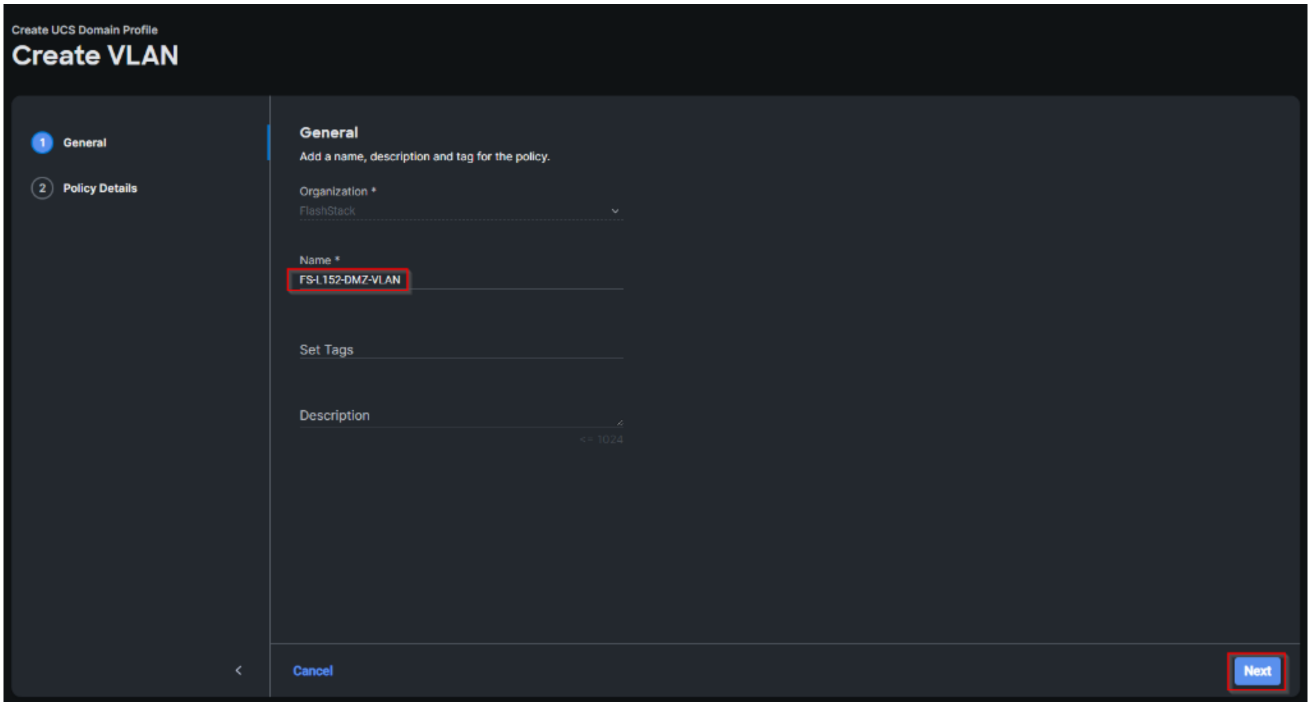

Create UCS Domain Profile

Create VLAN

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-VLAN

Set Tags

Description

<= 1024

Step 10. Click Add VLANs.

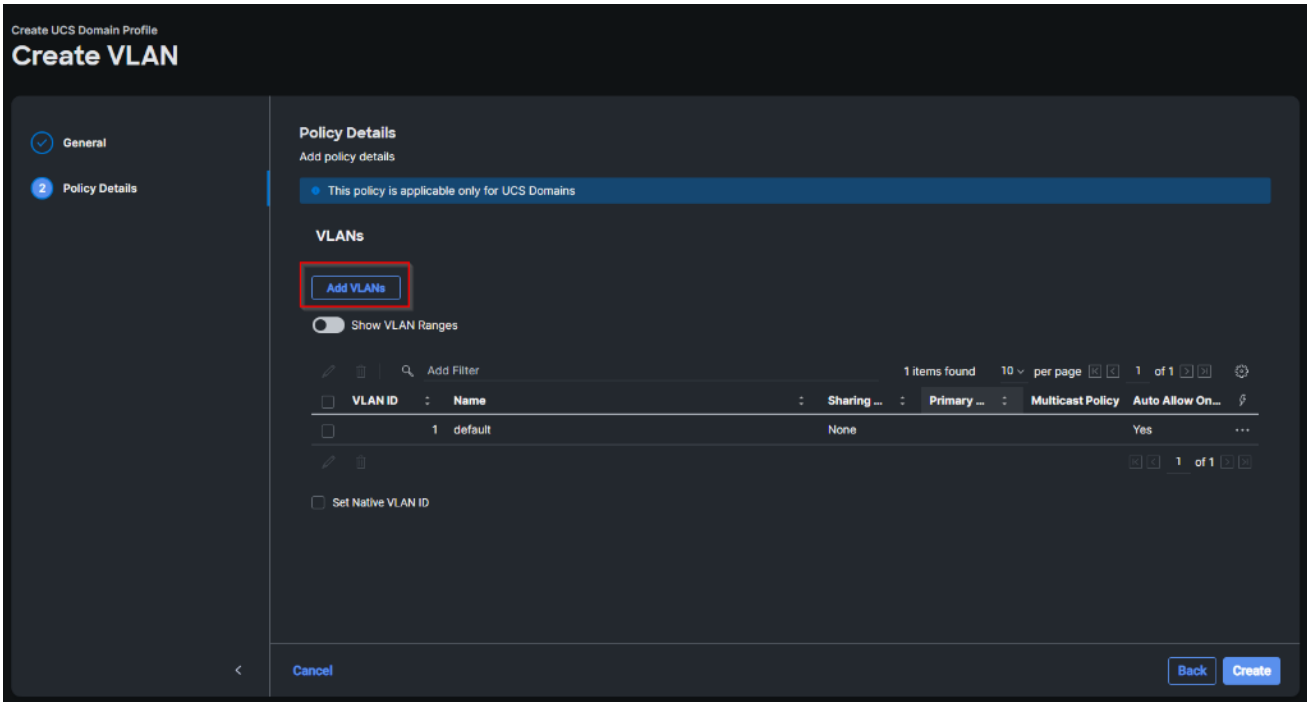

Create UCS Domain Profile

Create VLAN

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

Show VLAN Ranges

Add Filter 1 items found 10 per page 1 of 1

	VLAN ID	Name	Sharing ...	Primary ...	Multicast Policy	Auto Allow On...
□	1	default	None			Yes

1 of 1

□ Set Native VLAN ID

Step 11. Provide a name and VLAN ID for the VLAN from you list (for example, 40, 41, 42,43). Enable Auto Allow On Uplinks. To create the required Multicast policy, click Select Policy under Multicast*.

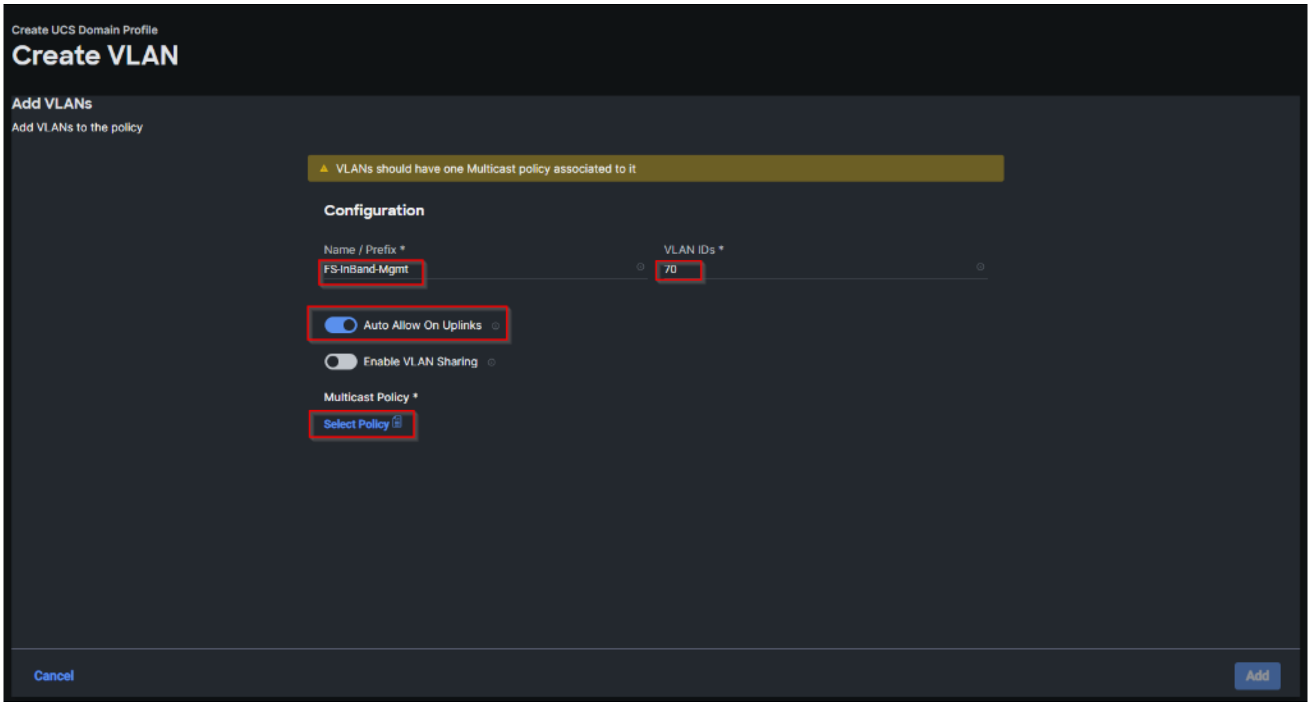
Create UCS Domain Profile

Create VLAN

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix * FS-InBand-Mgmt

VLAN IDs * 70

Auto Allow On Uplinks (enabled)

Enable VLAN Sharing

Multicast Policy *

Step 12. In the window on the right, click Create New to create a new Multicast Policy.

Step 13. Provide a Name for the Multicast Policy (for example, FS-L152-DMZ-McastPol). Provide optional Description and click Next.

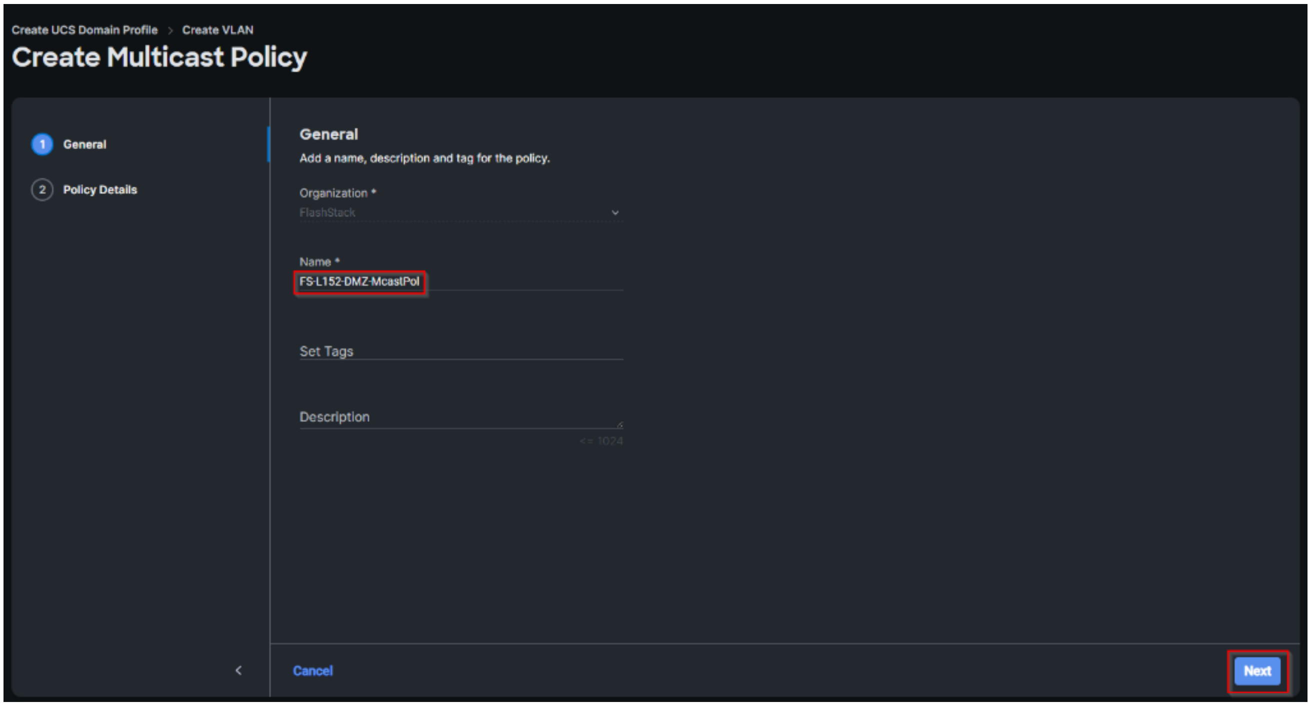
Create UCS Domain Profile > Create VLAN

Create Multicast Policy

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-McastPol

Set Tags

Description

Step 14. Leave defaults selected and click Create.

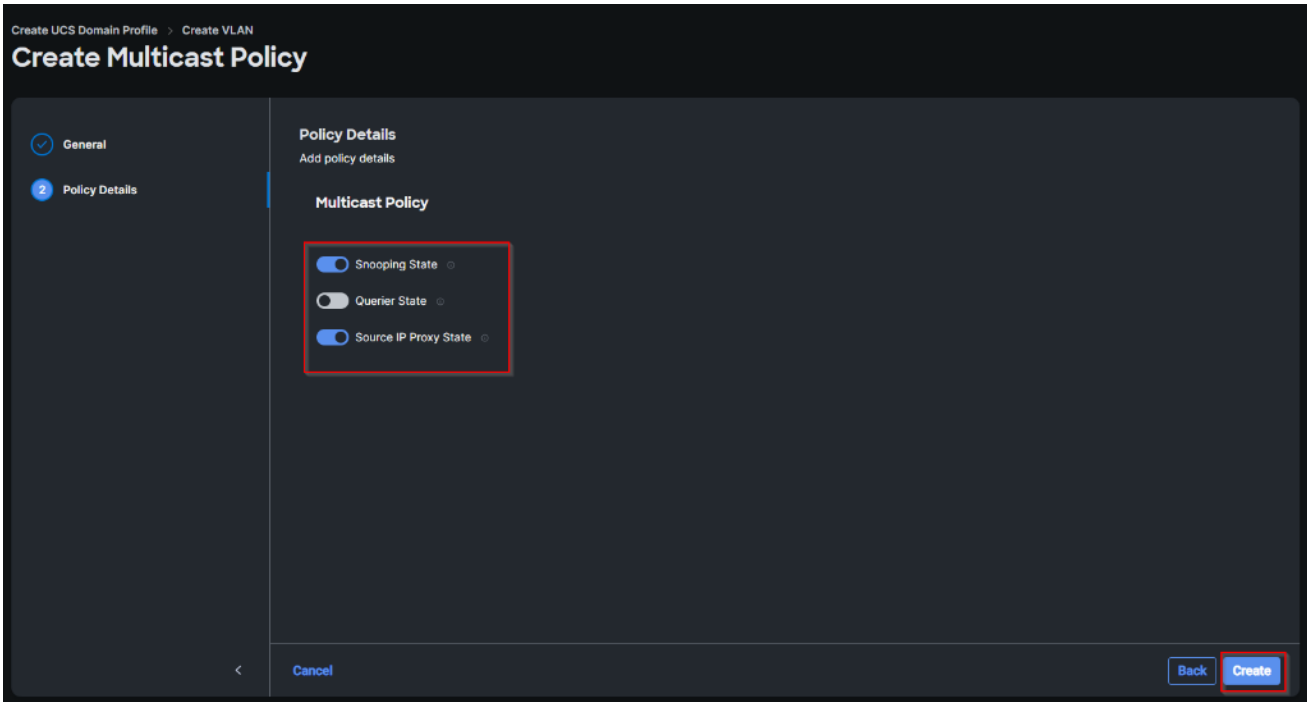
Create UCS Domain Profile > Create VLAN

Create Multicast Policy

✓ General

2 Policy Details

Policy Details

Add policy details

Multicast Policy

● Snooping State

○ Querier State

● Source IP Proxy State

Step 15. Click Add to add the VLAN.

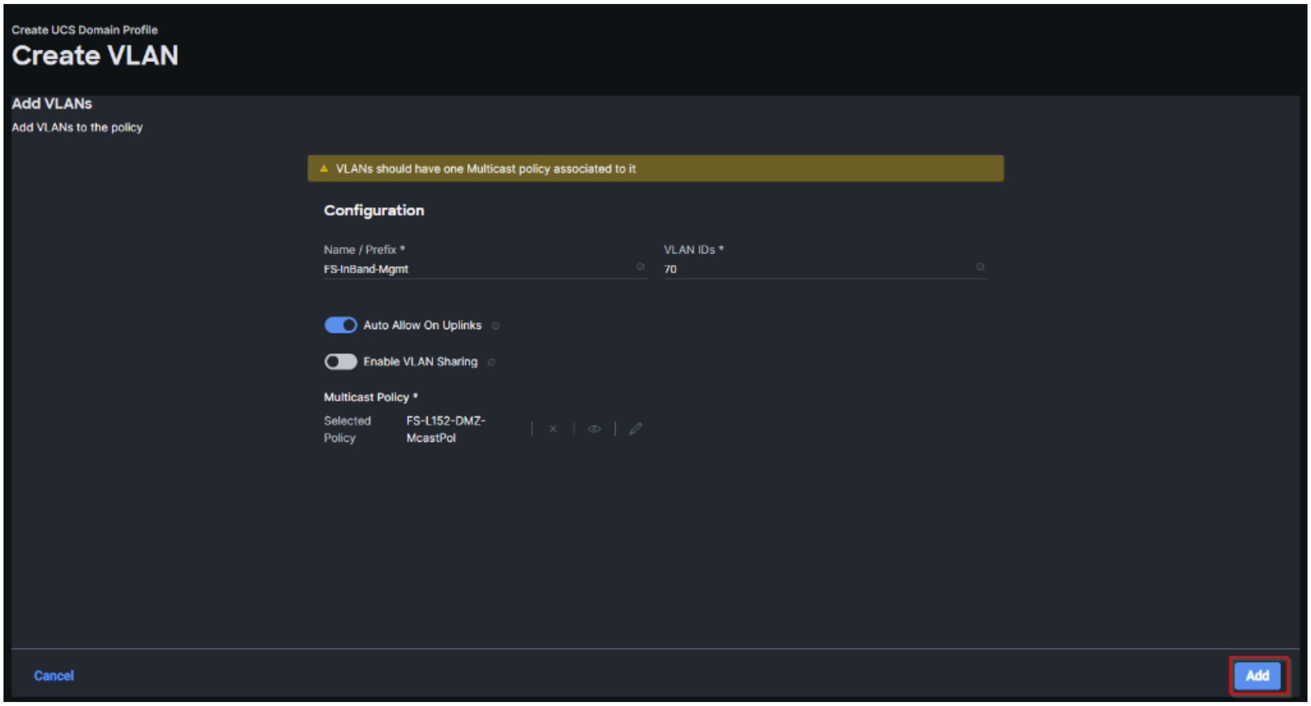
Create UCS Domain Profile

Create VLAN

Add VLANs

Add VLANs to the policy

⚠ VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix * FS-InBand-Mgmt

VLAN IDs * 70

● Auto Allow On Uplinks

○ Enable VLAN Sharing

Multicast Policy *

Selected Policy FS-L152-DMZ-McastPol

Step 16. Add the remaining VLANs from you list by clicking Add VLANs and entering the VLANs one by one. Reuse the previously created multicast policy for all the VLANs.

The VLANs created during this validation are shown below:

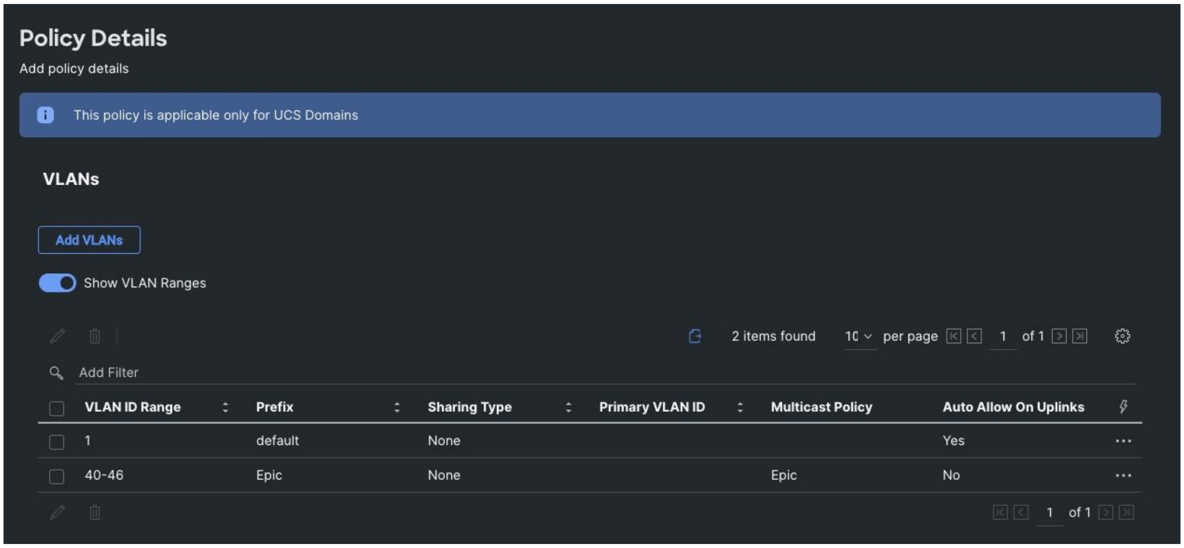

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

VLAN ID Range	Prefix	Sharing Type	Primary VLAN ID	Multicast Policy	Auto Allow On Uplinks
1	default	None			Yes
40-46	Epic	None		Epic	No

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

Step 17. Click Select Policy next to VSAN Configuration under Fabric Interconnect A. Click Create New.

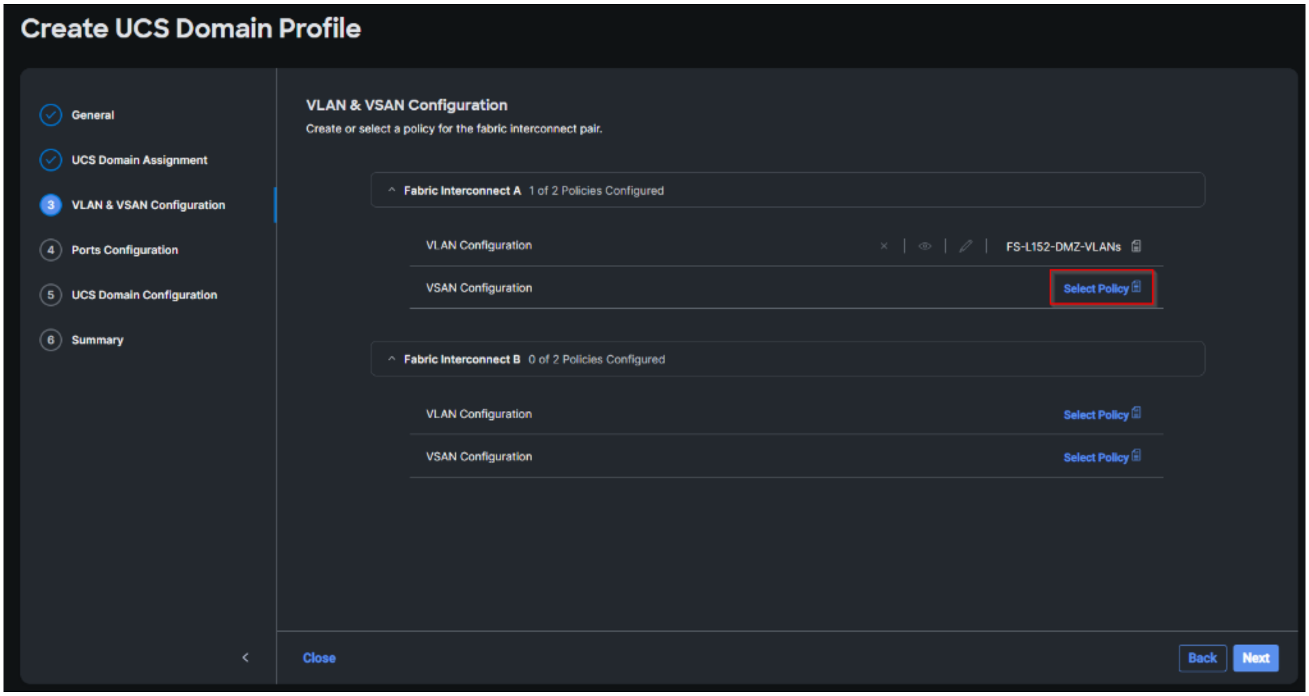

Create UCS Domain Profile

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 1 of 2 Policies Configured

VLAN Configuration FS-L152-DMZ-VLANs

VSAN Configuration Select Policy

Fabric Interconnect B 0 of 2 Policies Configured

VLAN Configuration Select Policy

VSAN Configuration Select Policy

Step 18. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-VSAN-A). Click Next.

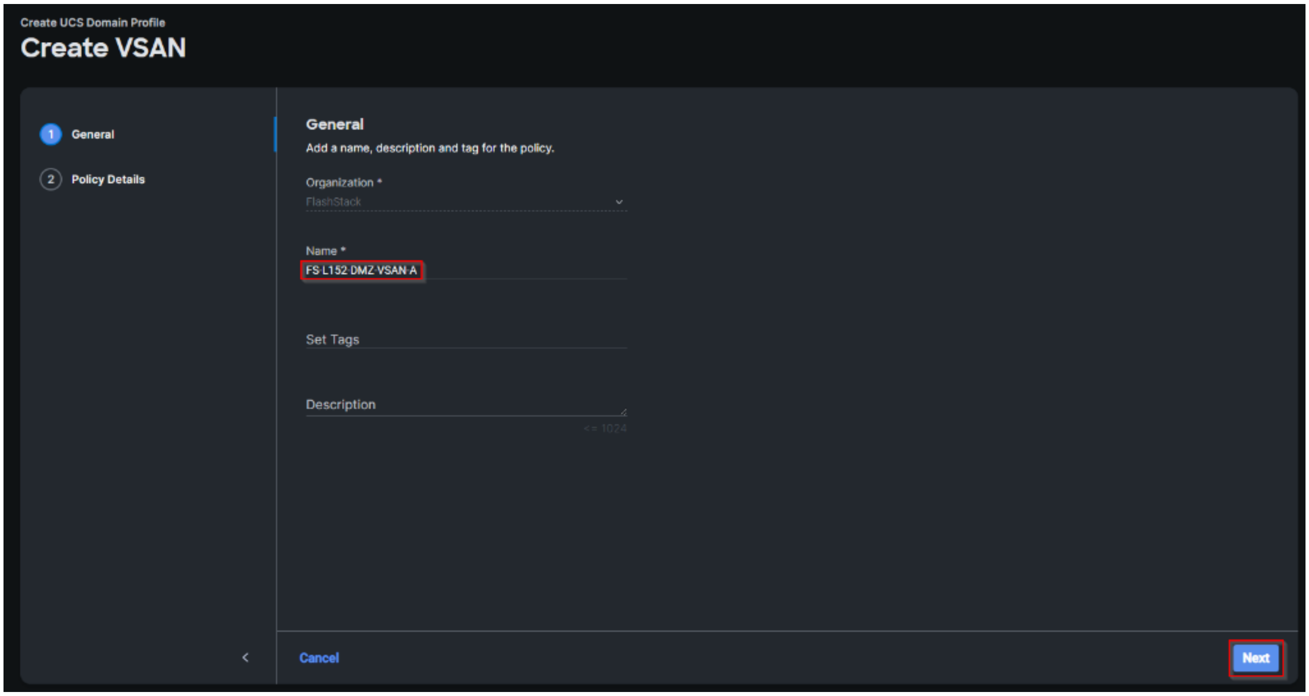
Create UCS Domain Profile

Create VSAN

General

Add a name, description and tag for the policy.

Organization *: FlashStack

Name *: FS-L152-DMZ-VSAN-A

Step 19. Click Add VSAN.

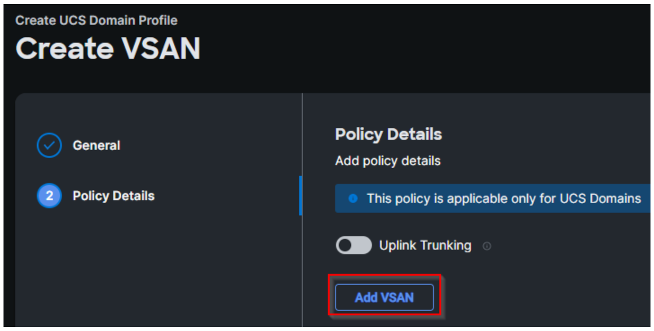
Create UCS Domain Profile

Create VSAN

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Uplink Trunking: off

Step 20. Provide a name (for example, VSAN-A), VSAN ID (for example, 700), and associated Fibre Channel over Ethernet (FCoE) VLAN ID (for example, 700) for VSAN A.

Step 21. Set VLAN Scope as Uplink.

Step 22. Click Add.

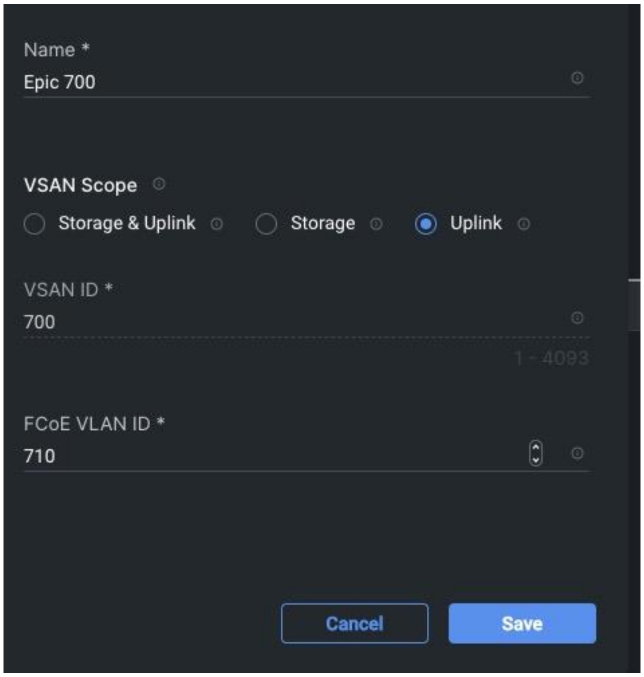

Name *
Epic 700

VSAN Scope

○ Storage & Uplink ○ Storage ● Uplink

VSAN ID *
700
1 - 4093

FCoE VLAN ID *
710

Step 23. Click Create to finish creating VSAN policy for fabric A.

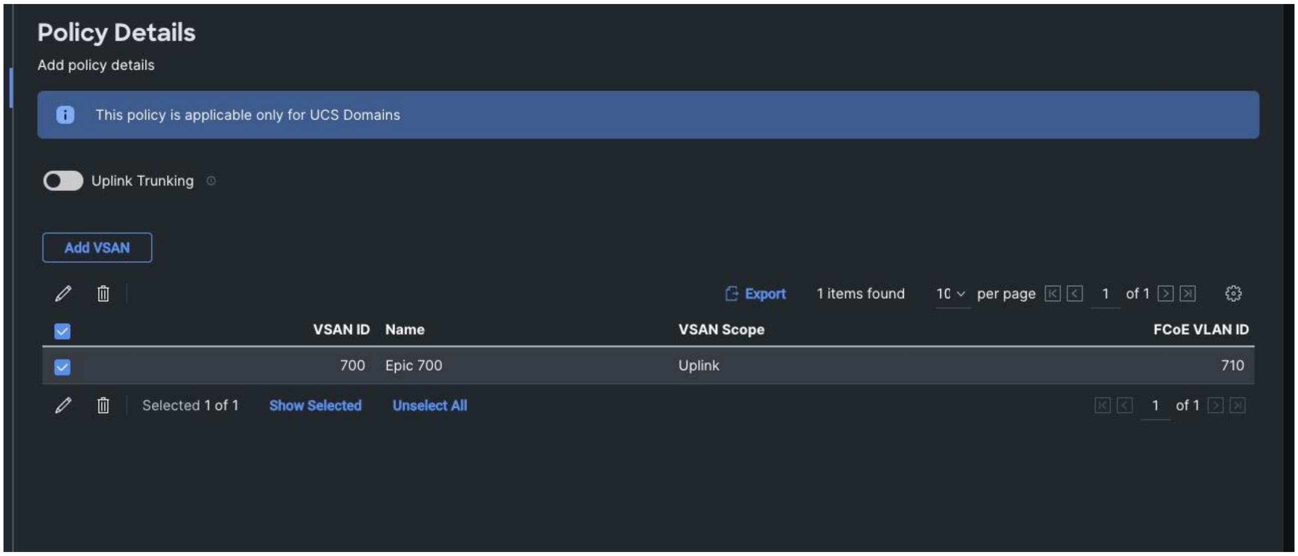

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Uplink Trunking

☑	VSAN ID	Name	VSAN Scope	FCoE VLAN ID
☑	700	Epic 700	Uplink	710

Step 24. Repeat steps 13 - 19 for fabric interconnect B assigning the VLAN policy created previously and creating a new VSAN policy for VSAN-B. Name the policy to identify the SAN-B configuration (for example, FS-L152-DMZ-VSAN-B) and use appropriate VSAN and FCoE VLAN (for example, 701).

Step 25. Verify that a common VLAN policy and two unique VSAN policies are associated with the two fabric interconnects. Click Next.

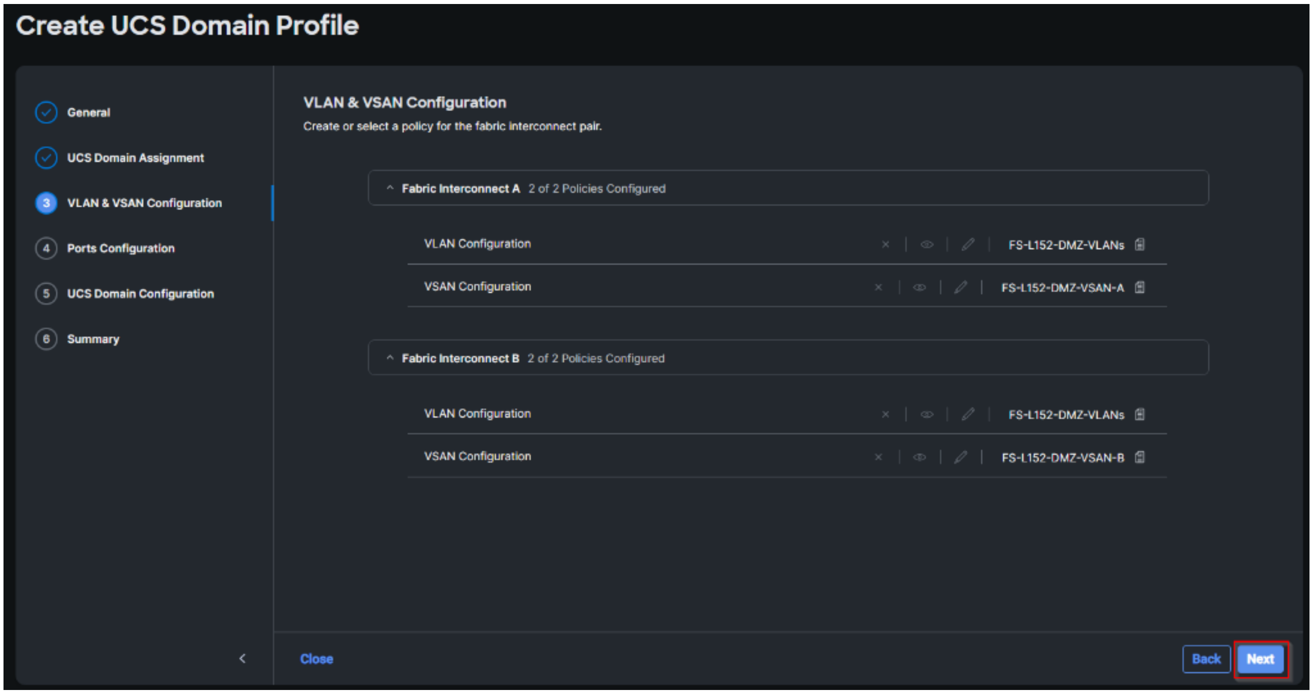
Create UCS Domain Profile

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 2 of 2 Policies Configured

VLAN Configuration: FS-L152-DMZ-VLANs

VSAN Configuration: FS-L152-DMZ-VSAN-A

Fabric Interconnect B 2 of 2 Policies Configured

VLAN Configuration: FS-L152-DMZ-VLANs

VSAN Configuration: FS-L152-DMZ-VSAN-B

Step 26. On the Ports Configuration page, attach port policies for each switch to the UCS Domain Profile.

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 port policies are required because each fabric interconnect uses unique Fibre Channel VSAN ID.

Step 27. Click Select Policy for Fabric Interconnect A.

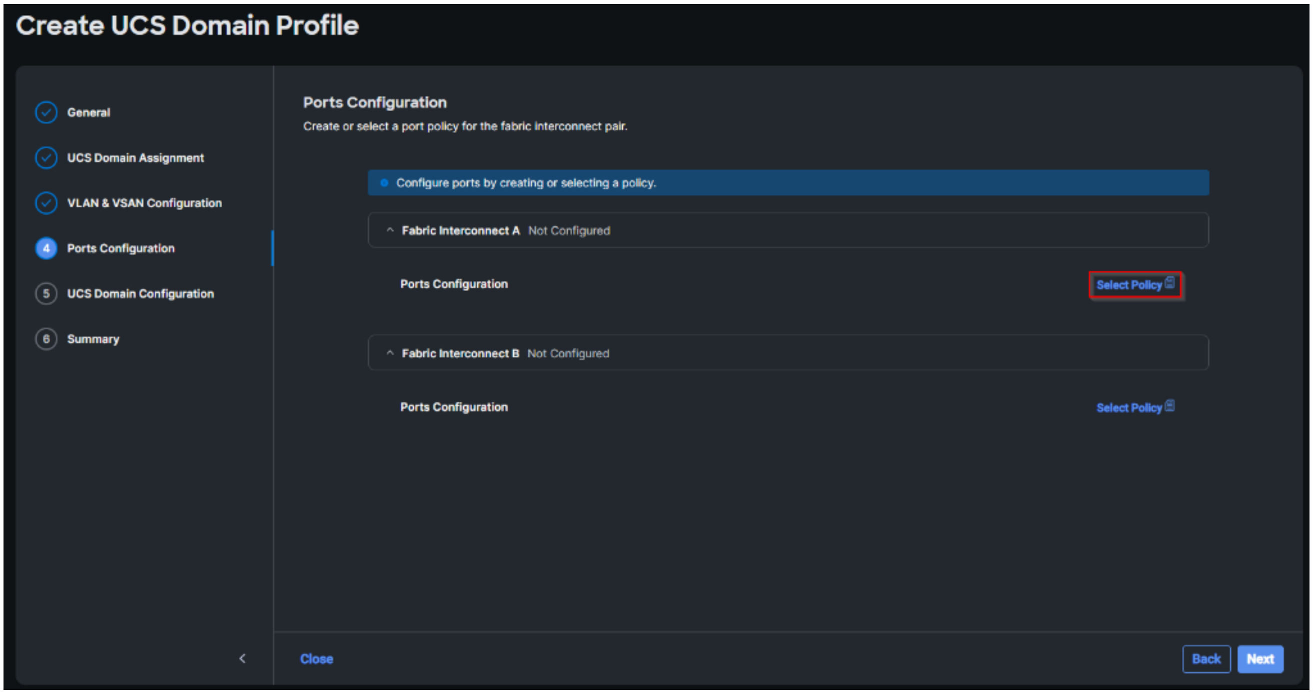
Create UCS Domain Profile

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

Fabric Interconnect B Not Configured

Ports Configuration: Select Policy

Step 28. Click Create New.

Step 29. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-K4-FI-A). Click Next.

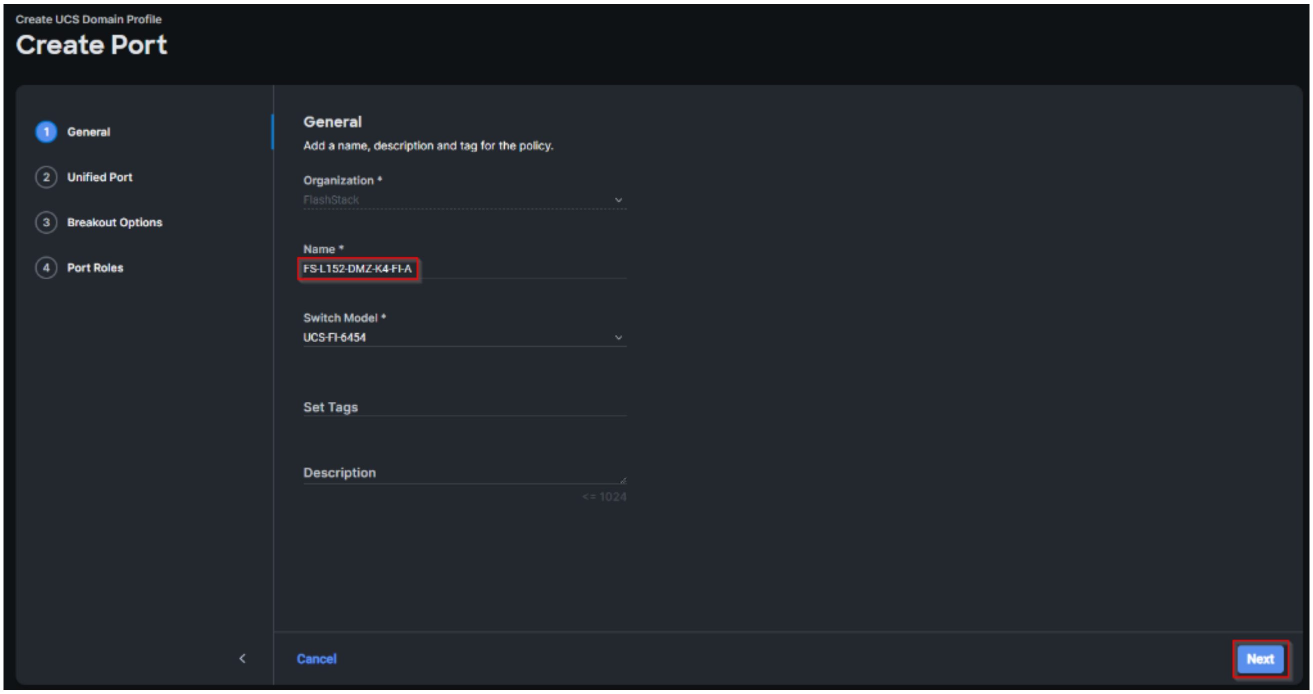

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-K4-FI-A

Switch Model * UCS-FI-6454

Set Tags

Description

<= 1024

Step 30. Move the slider to set up unified ports. In this deployment, the first four ports were selected as Fibre Channel ports. Click Next.

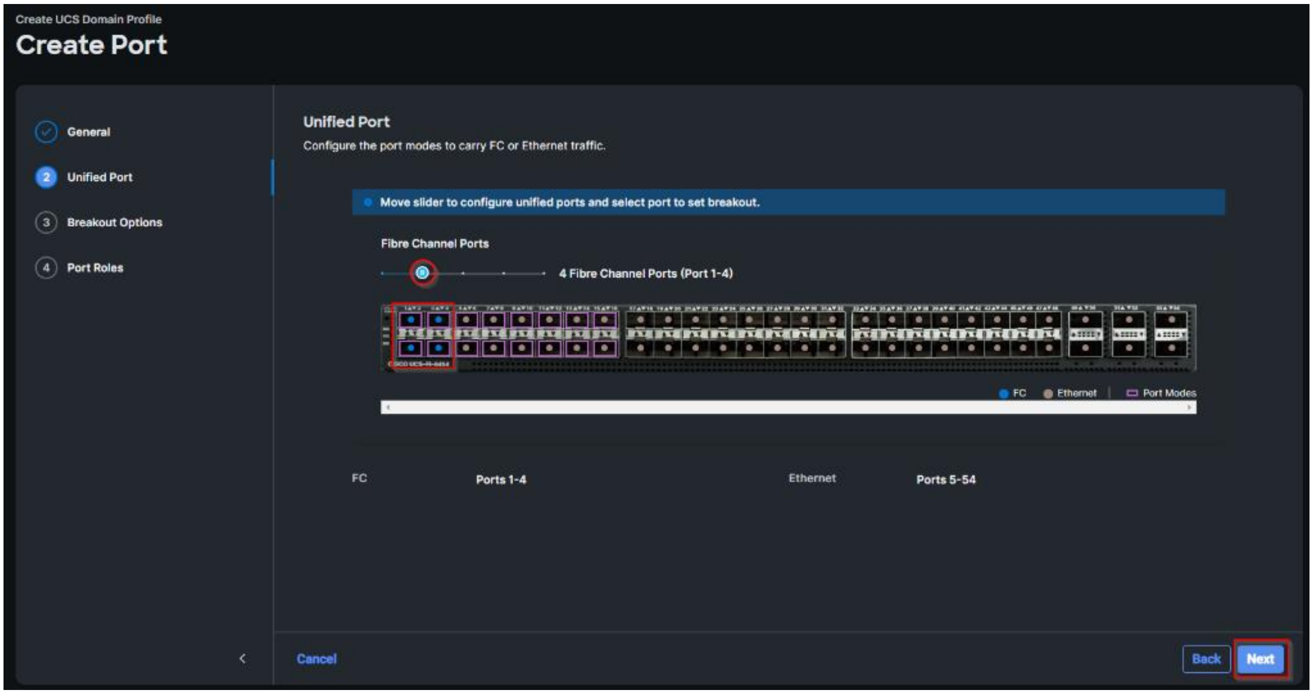

Unified Port

Configure the port modes to carry FC or Ethernet traffic.

Move slider to configure unified ports and select port to set breakout.

Fibre Channel Ports

4 Fibre Channel Ports (Port 1-4)

FC Ports 1-4

Ethernet Ports 5-54

Step 31. On the breakout Options page click Next.

Note: No Ethernet/Fibre Channel breakouts were used in this validation.

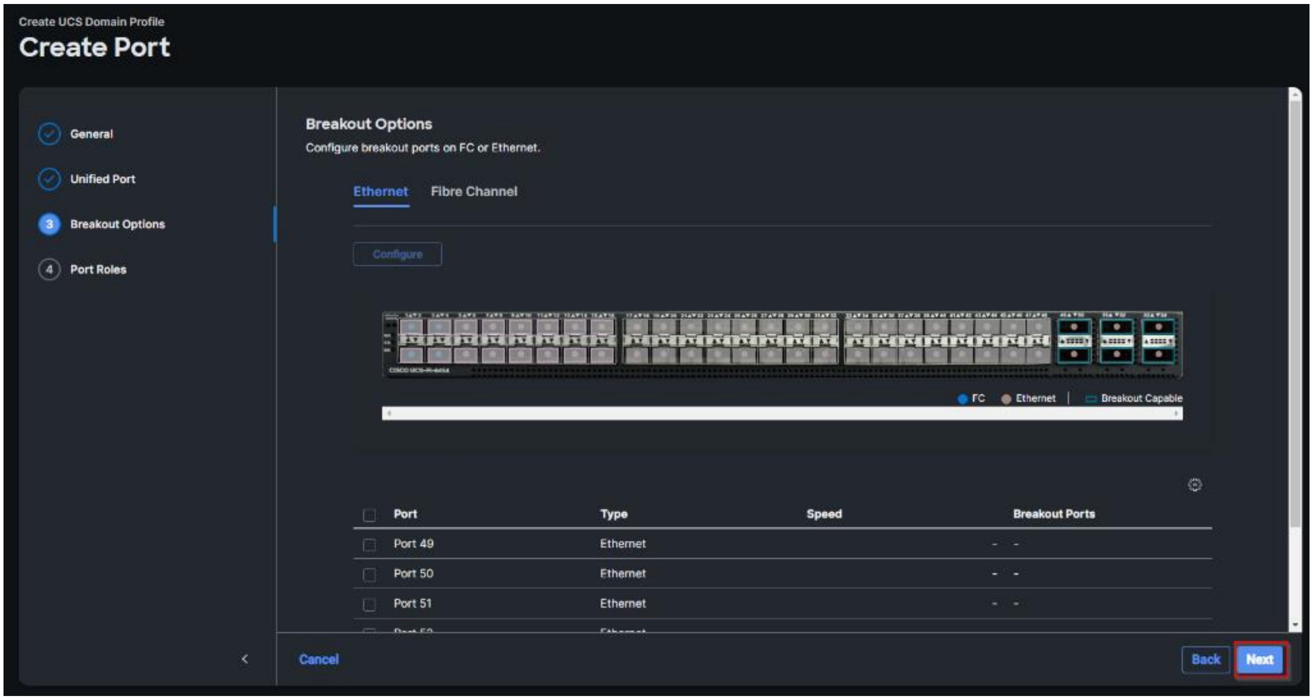

Create UCS Domain Profile

Create Port

- General
- Unified Port
- 3 Breakout Options
- 4 Port Roles

Breakout Options

Configure breakout ports on FC or Ethernet.

□	Port	Type	Speed	Breakout Ports
□	Port 49	Ethernet	-	-
□	Port 50	Ethernet	-	-
□	Port 51	Ethernet	-	-

Step 32. Select the ports that need to be configured as server ports by clicking the ports in the graphics (or select from the list below the graphic). When all ports are selected, click Configure.

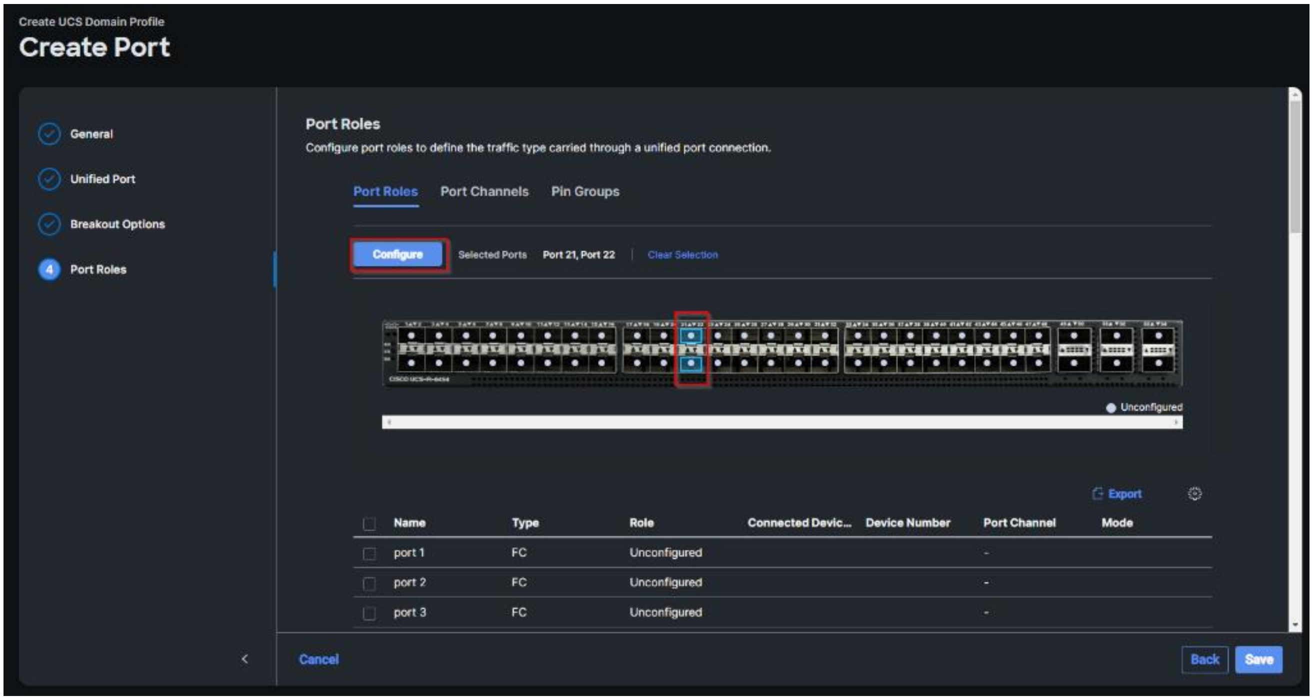

Create UCS Domain Profile

Create Port

- General
- Unified Port
- Breakout Options
- 4 Port Roles

Port Roles

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

Selected Ports Port 21, Port 22 | Clear Selection

□	Name	Type	Role	Connected Devic...	Device Number	Port Channel	Mode
□	port 1	FC	Unconfigured			-	
□	port 2	FC	Unconfigured			-	
□	port 3	FC	Unconfigured			-	

Step 33. From the drop-down list, select Server as the role. Click Save.

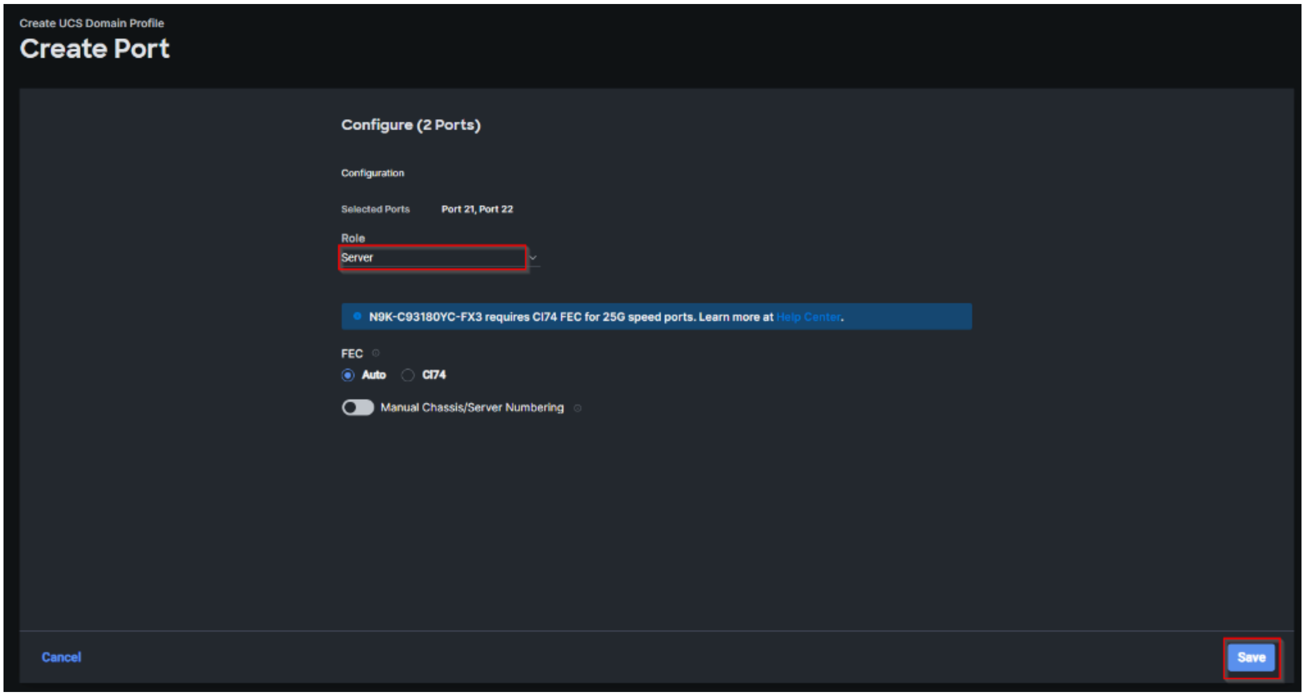

Create UCS Domain Profile

Create Port

Configure (2 Ports)

Configuration

Selected Ports Port 21, Port 22

Role: Server

N9K-C93180YC-FX3 requires Cl74 FEC for 25G speed ports. Learn more at Help Center.

FEC: ● Auto ○ Cl74

Manual Chassis/Server Numbering

Step 34. Configure the Ethernet uplink port channel by selecting the Port Channel in the main pane and then clicking Create Port Channel.

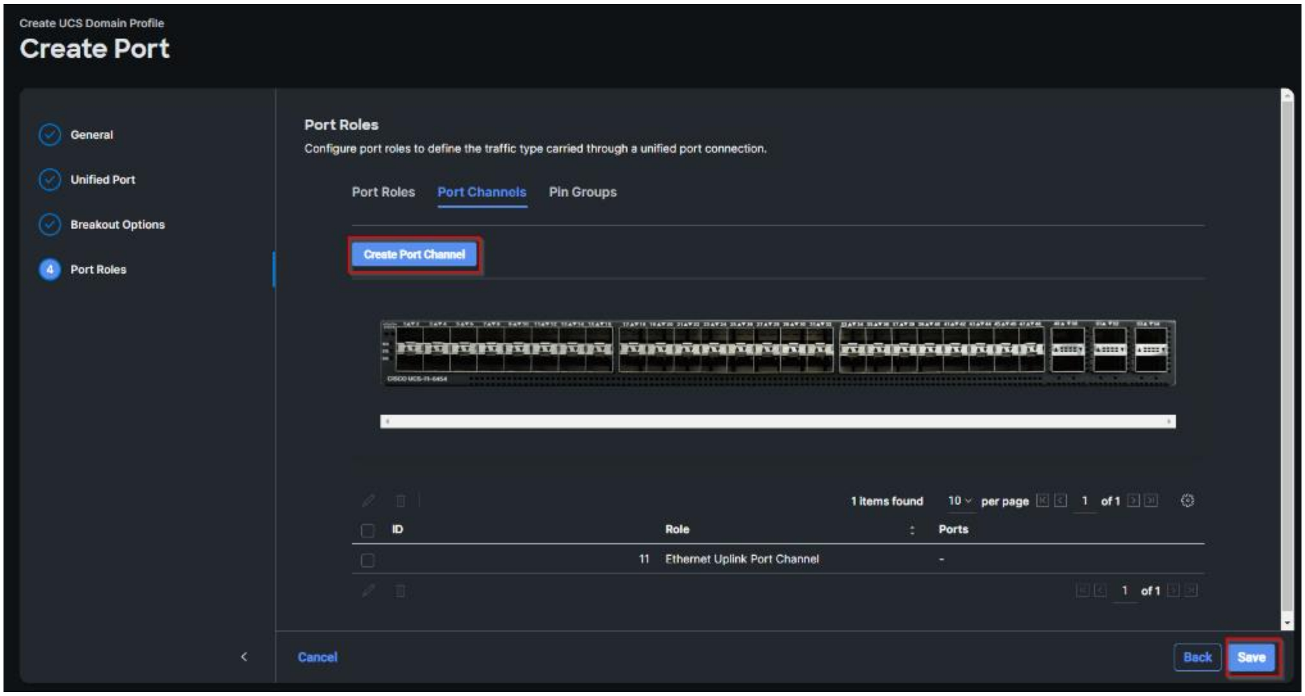

Create UCS Domain Profile

Create Port

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	-

Step 35. Select Ethernet Uplink Port Channel as the role, provide a port-channel ID (for example, 11).

Note: You can create the Ethernet Network Group, Flow Control, Link Aggregation or Link control policy for defining disjoint Layer-2 domain or fine tune port-channel parameters. These policies were not used in this deployment and system default values were utilized.

Step 36. Scroll down and select uplink ports from the list of available ports (for example, port 49 and 50).

Step 37. Click Save.

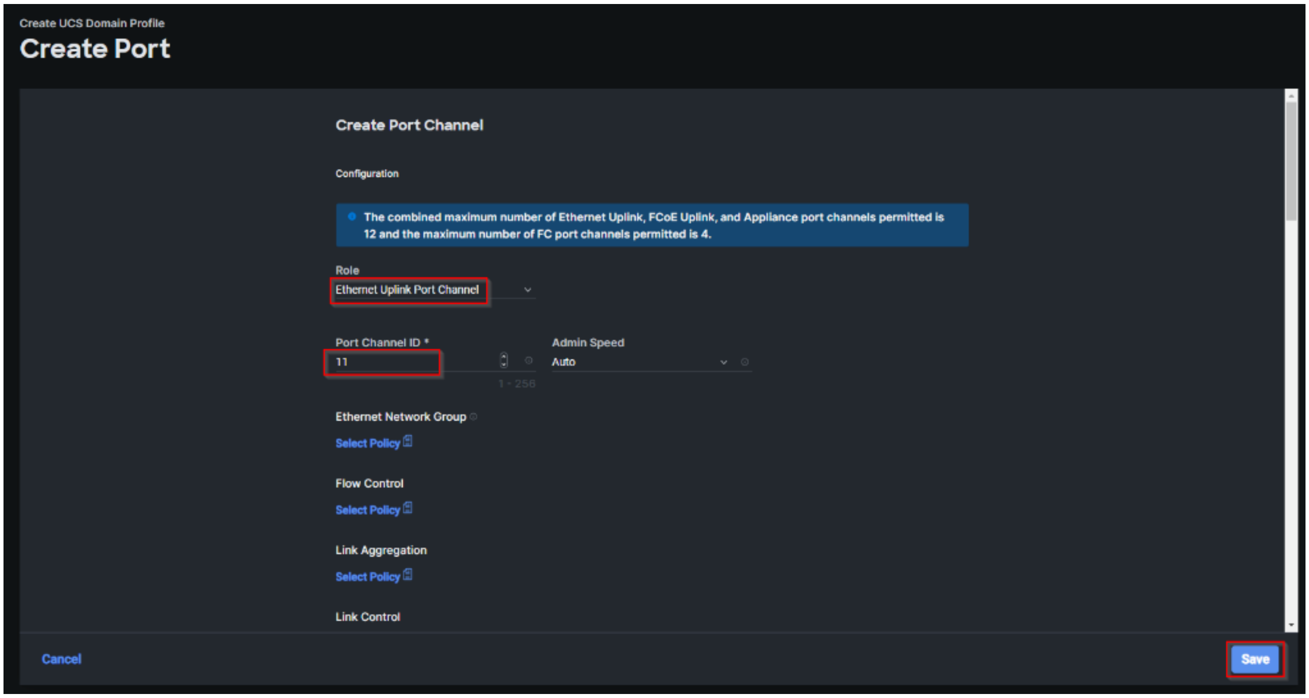

Create UCS Domain Profile

Create Port

Create Port Channel

Configuration

The combined maximum number of Ethernet Uplink, FCoE Uplink, and Appliance port channels permitted is 12 and the maximum number of FC port channels permitted is 4.

Role: Ethernet Uplink Port Channel

Port Channel ID *: 11 (1 - 256)

Admin Speed: Auto

Step 38. Repeat steps 54-63 to create the port policy for Fabric Interconnect B. Use the following values for various parameters:

- Name of the port policy: FS-L152-DMZ-K4-FI-B
- Ethernet port-Channel ID: 12

Step 39. When the port configuration for both fabric interconnects is complete and looks good, click Next.

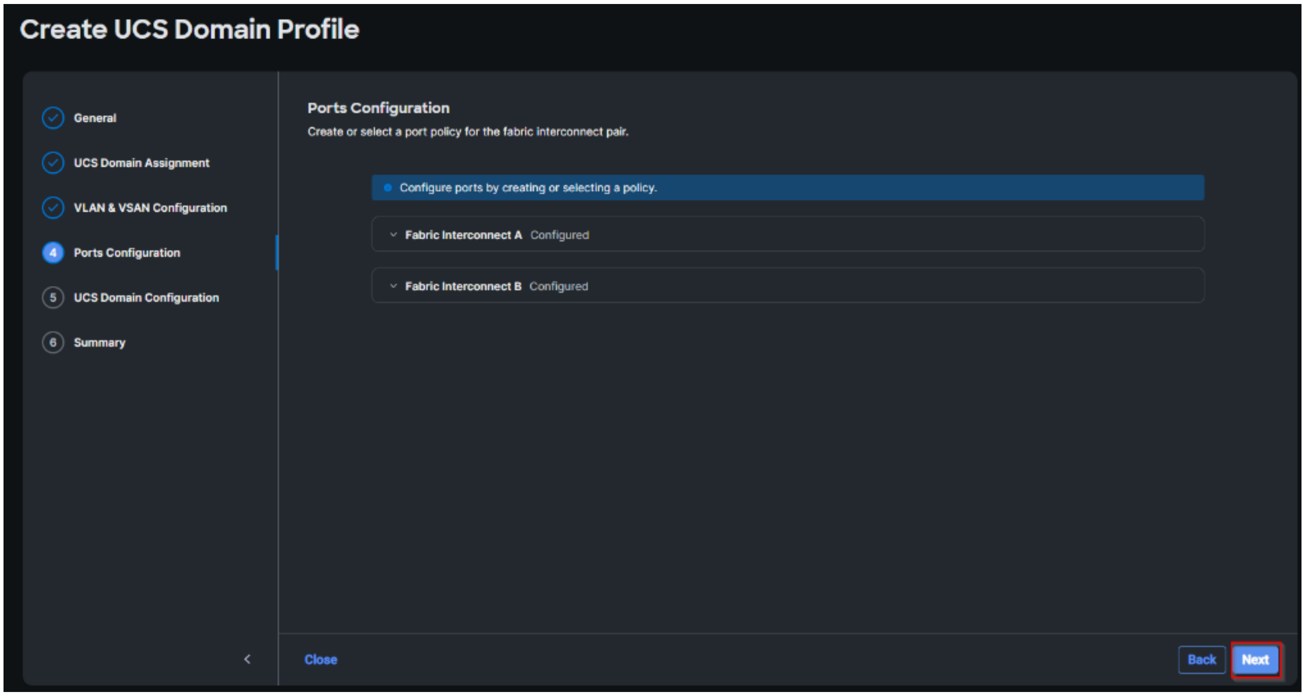

Create UCS Domain Profile

Ports Configuration

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

Configure ports by creating or selecting a policy.

Fabric Interconnect A Configured

Fabric Interconnect B Configured

Step 40. Under UCS domain configuration, additional policies can be configured to setup NTP, Syslog, DNS settings, SNMP, QoS and UCS operating mode (end host or switch mode). For this deployment, System QoS will be configured.

Step 41. Click Select Policy next to System QoS* and click Create New to define the System QOS policy.

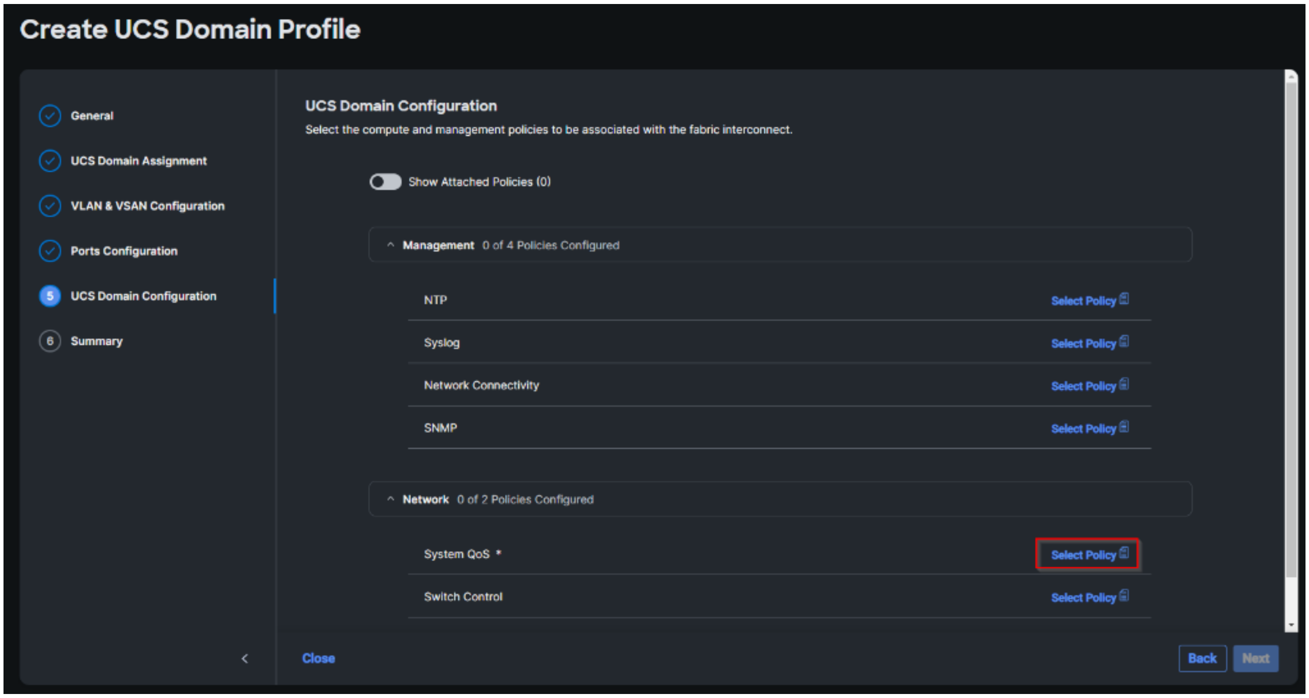

Create UCS Domain Profile

- General
- UCS Domain Assignment
- VLAN & VSAN Configuration
- Ports Configuration
- 5 UCS Domain Configuration
- 6 Summary

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

Step 42. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-QosPol). Click Next.

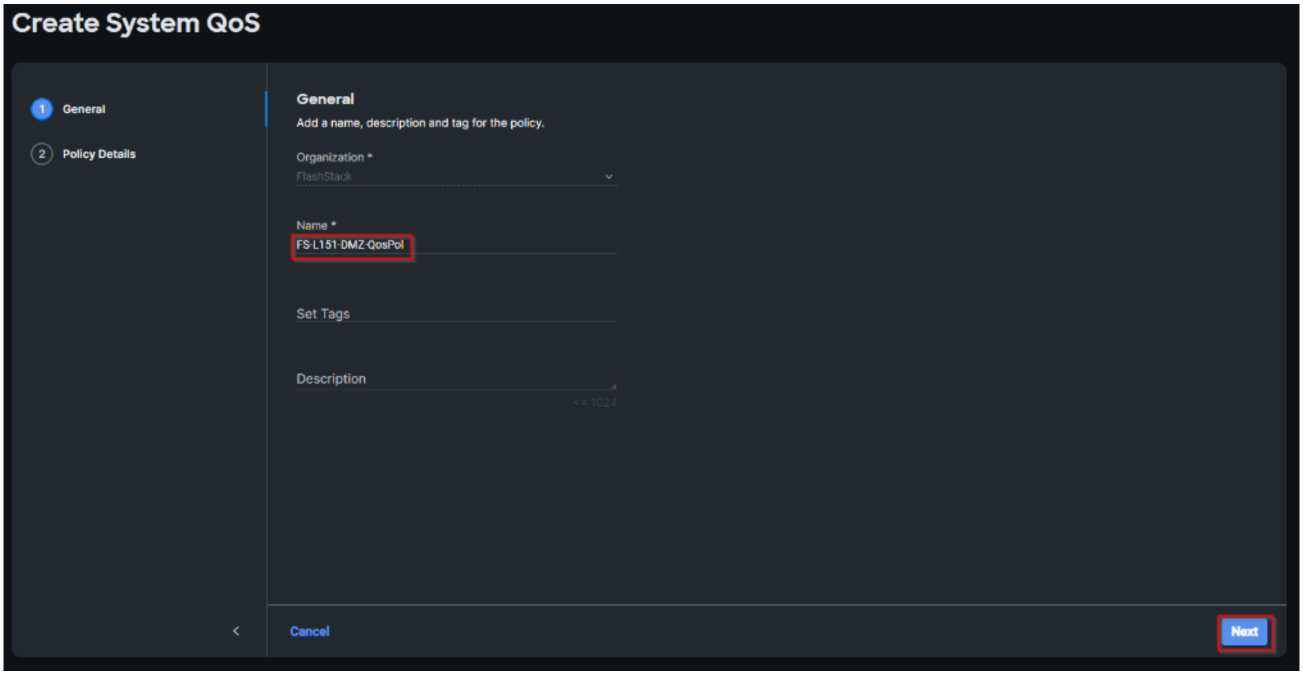

Create System QoS

- 1 General
- 2 Policy Details

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L151-DMZ-QosPol

Set Tags

Description

<= 1024

Step 43. Change the MTU for Best Effort class to 9216. Keep the rest default selections. Click Create.

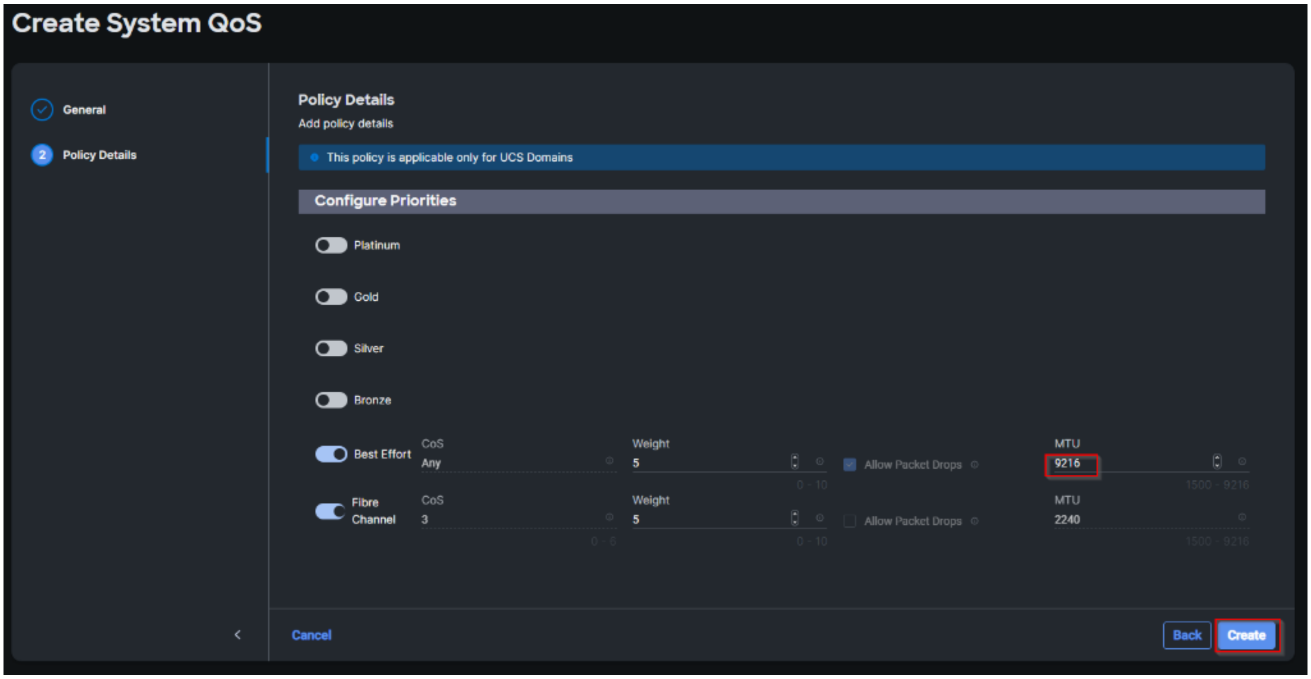
Create System QoS

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Configure Priorities

- Platinum
- Gold
- Silver
- Bronze

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

Step 44. Click Next.

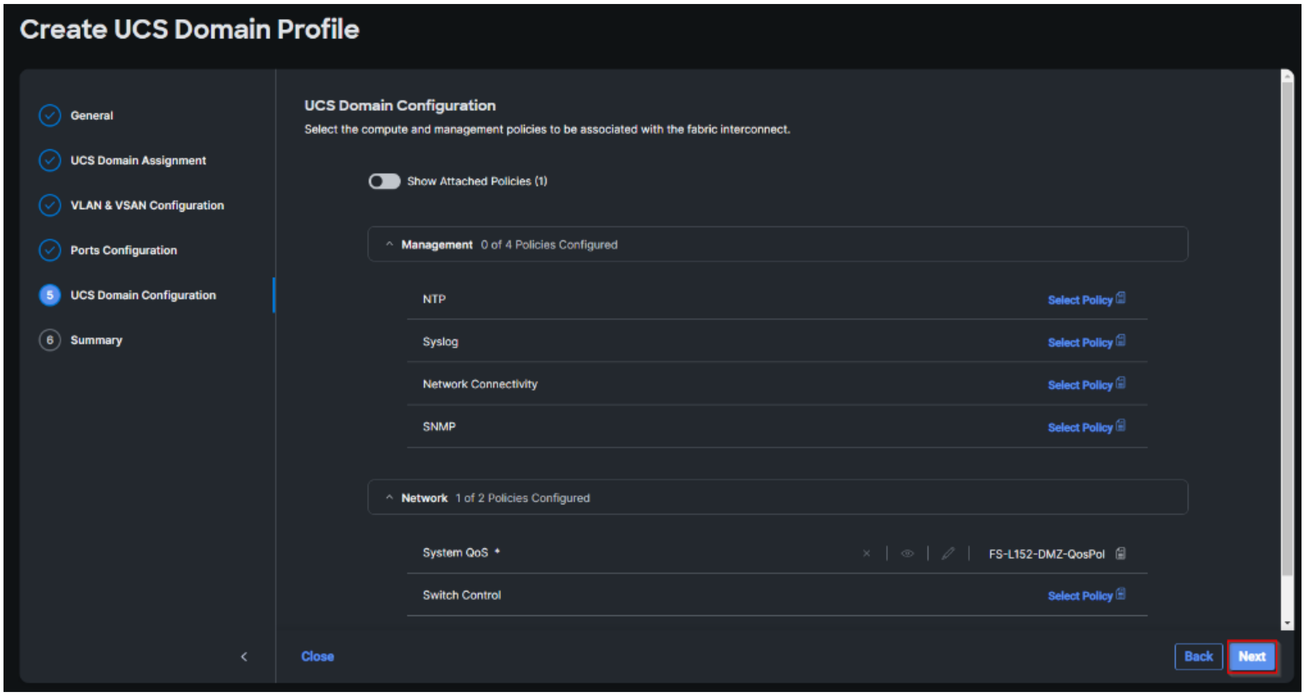
Create UCS Domain Profile

UCS Domain Configuration

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

Show Attached Policies (1)

Management 0 of 4 Policies Configured

- NTP — Select Policy
- Syslog — Select Policy
- Network Connectivity — Select Policy
- SNMP — Select Policy

Network 1 of 2 Policies Configured

- System QoS * — FS-L152-DMZ-QosPol
- Switch Control — Select Policy

Step 45. From the UCS domain profile Summary view, verify all the settings including the fabric interconnect settings, by expanding the settings and make sure that the configuration is correct. Click Deploy.

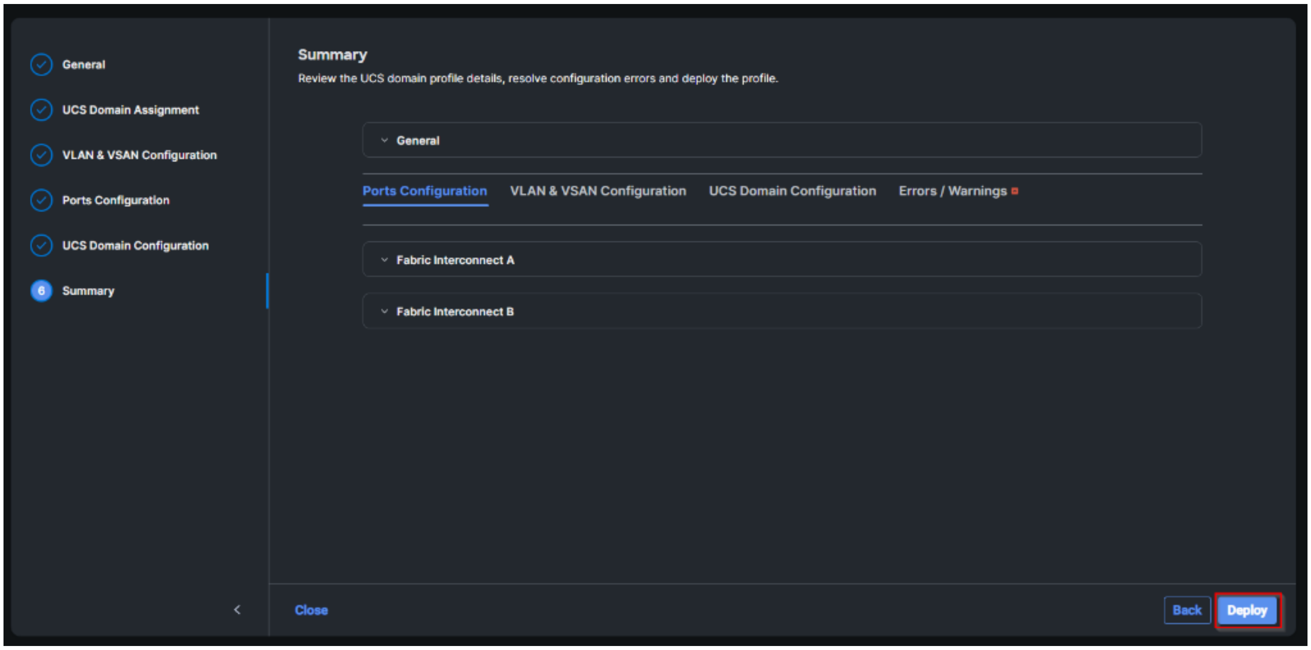

Summary

Review the UCS domain profile details, resolve configuration errors and deploy the profile.

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.

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

It takes a while to discover the blades for the first time. Cisco Intersight provides an ability to view the progress in the Requests page:

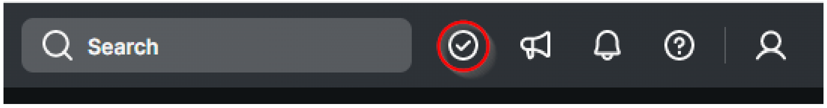

Step 46. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Configure > Profiles, select UCS Domain Profiles, verify that the domain profile has been successfully deployed.

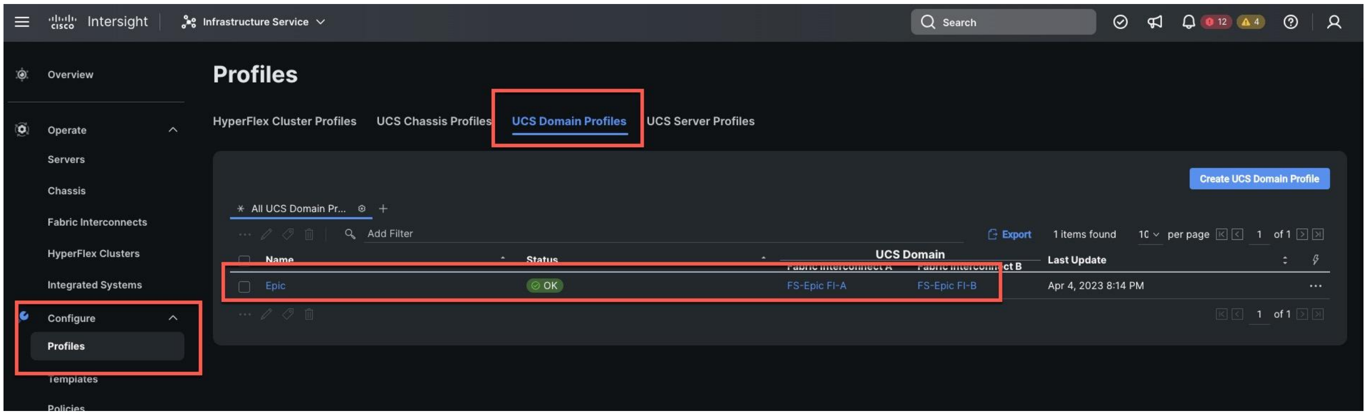

Profiles

Name	Status	UCS Domain Fabric Interconnect A	 Fabric Interconnect B	Last Update
Epic	OK	FS-Epic FI-A	FS-Epic FI-B	Apr 4, 2023 8:14 PM

Step 47. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Chassis, verify that the chassis has been discovered.

Step 48. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Servers, verify that the servers have been successfully discovered.

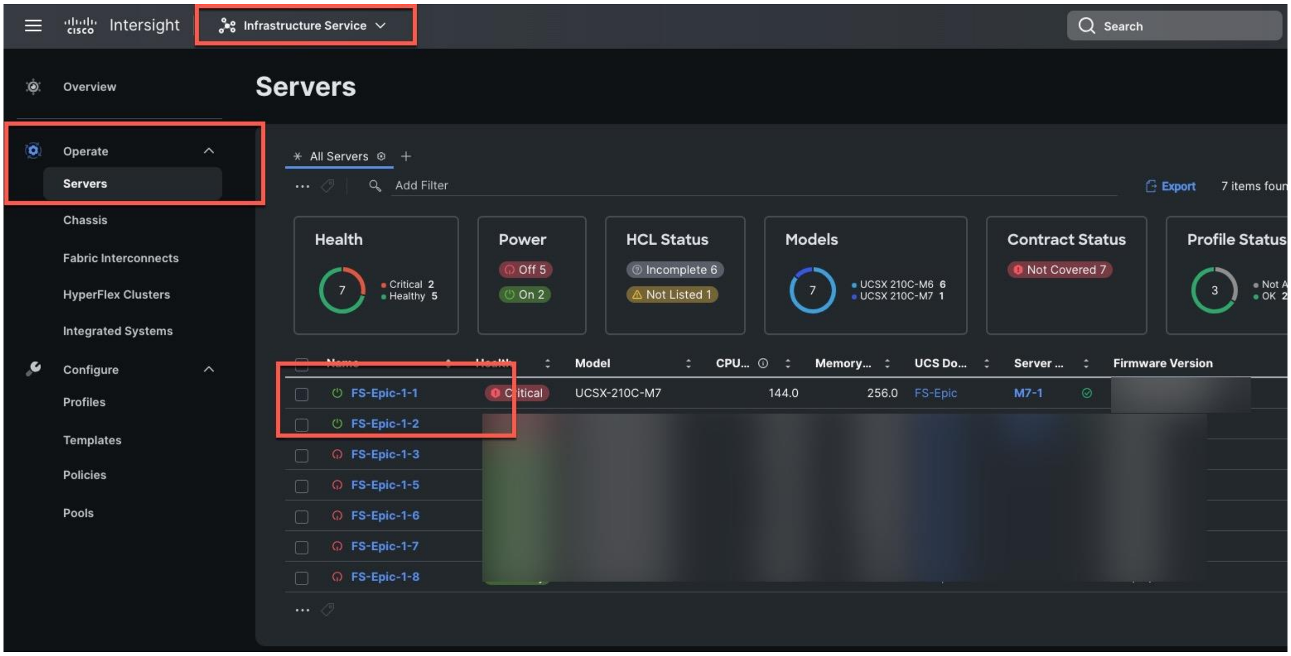

Servers

Health	Power	HCL Status	Models	Contract Status	Profile Status
7 Critical 2 Healthy 5	Off 5 On 2	Incomplete 6 Not Listed 1	7 UCSX 210C-M6 6 UCSX 210C-M7 1	Not Covered 7	3 Not A OK 2

Name	Health	Model	CPU...	Memory...	UCS Do...	Server ...	Firmware Version
FS-Epic-1-1	Critical	UCSX-210C-M7	144.0	256.0	FS-Epic	M7-1 ✓	[redacted]
FS-Epic-1-2	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
FS-Epic-1-3	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
FS-Epic-1-5	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
FS-Epic-1-6	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
FS-Epic-1-7	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
FS-Epic-1-8	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

Configure Cisco UCS Chassis Profile

Cisco UCS Chassis profile in Cisco Intersight allows you to configure various parameters for the 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, chassis profile was created and attached to the chassis with following settings shown in Figure 17.

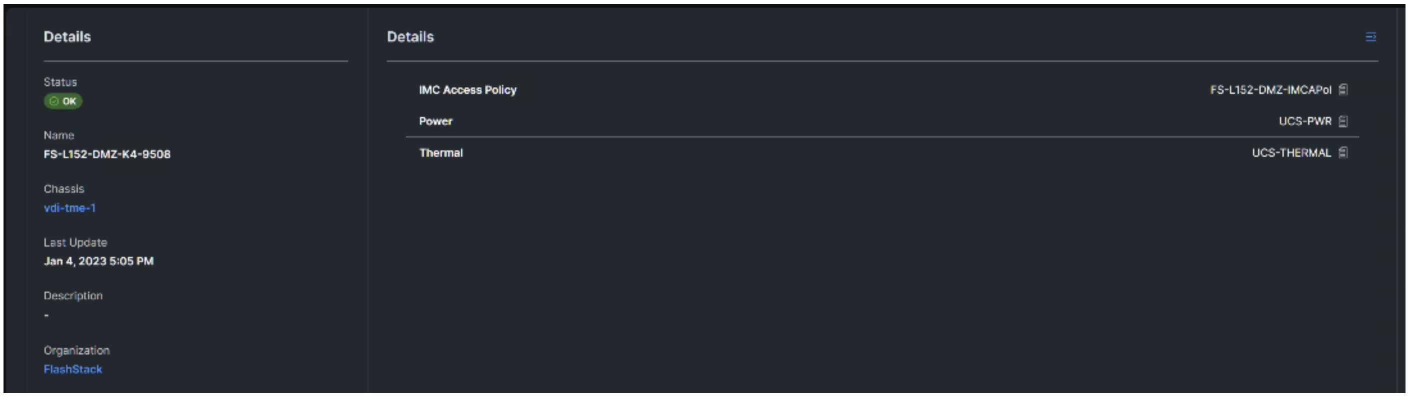

Details

Status
OK

Name
FS-L152-DMZ-K4-9508

Chassis
vdi-tme-1

Last Update
Jan 4, 2023 5:05 PM

Description
-

Organization
FlashStack

Details

IMC Access Policy	FS-L152-DMZ-IMCAPol
Power	UCS-PWR
Thermal	UCS-THERMAL

Figure 17. Chassis policy detail

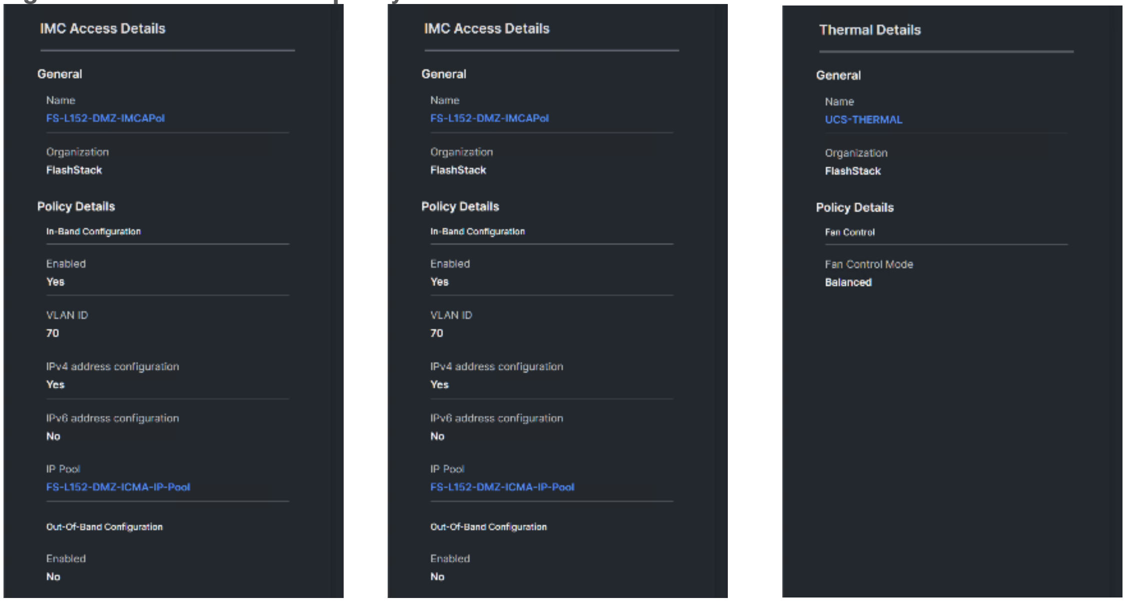

IMC Access Details

General

Name
FS-L152-DMZ-IMCAPol

Organization
FlashStack

Policy Details

In-Band Configuration

Enabled
Yes

VLAN ID
70

IPv4 address configuration
Yes

IPv6 address configuration
No

IP Pool
FS-L152-DMZ-ICMA-IP-Pool

Out-Of-Band Configuration

Enabled
No

IMC Access Details

General

Name
FS-L152-DMZ-IMCAPol

Organization
FlashStack

Policy Details

In-Band Configuration

Enabled
Yes

VLAN ID
70

IPv4 address configuration
Yes

IPv6 address configuration
No

IP Pool
FS-L152-DMZ-ICMA-IP-Pool

Out-Of-Band Configuration

Enabled
No

Thermal Details

General

Name
UCS-THERMAL

Organization
FlashStack

Policy Details

Fan Control

Fan Control Mode
Balanced

Configure Server Profiles

Configure Server Profile Template

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 server profile template, you can derive multiple consistent server profiles from the template.

Note: The server profile captured in this deployment guide supports both Cisco UCS X-Series blade servers and Cisco UCS X210c M7 compute nodes.

Procedure 1. Create vNIC and vHBA Placement for the Server Profile Template

In this deployment, four vNICs and two vHBAs are configured. These devices are manually placed as listed in Table 4.

Table 4. vHBA and vNIC placement for FC connected storage

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

Note: Two vHBAs (vHBA-A and vHBA-B) are configured to support FC boot from SAN.

Step 1. Log into the Cisco Intersight portal as a user with account administrator role.

Step 2. Navigate to Configure > Templates and click Create UCS Server Profile Template.

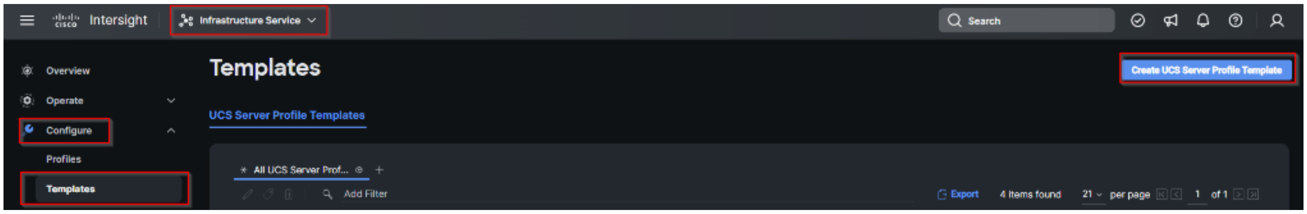

Templates

Step 3. Select the organization from the drop-down list. Provide a name for the server profile template (for example, FS-L151-DMZ-K4-X210CM7) for FI-Attached UCS Server. Click Next.

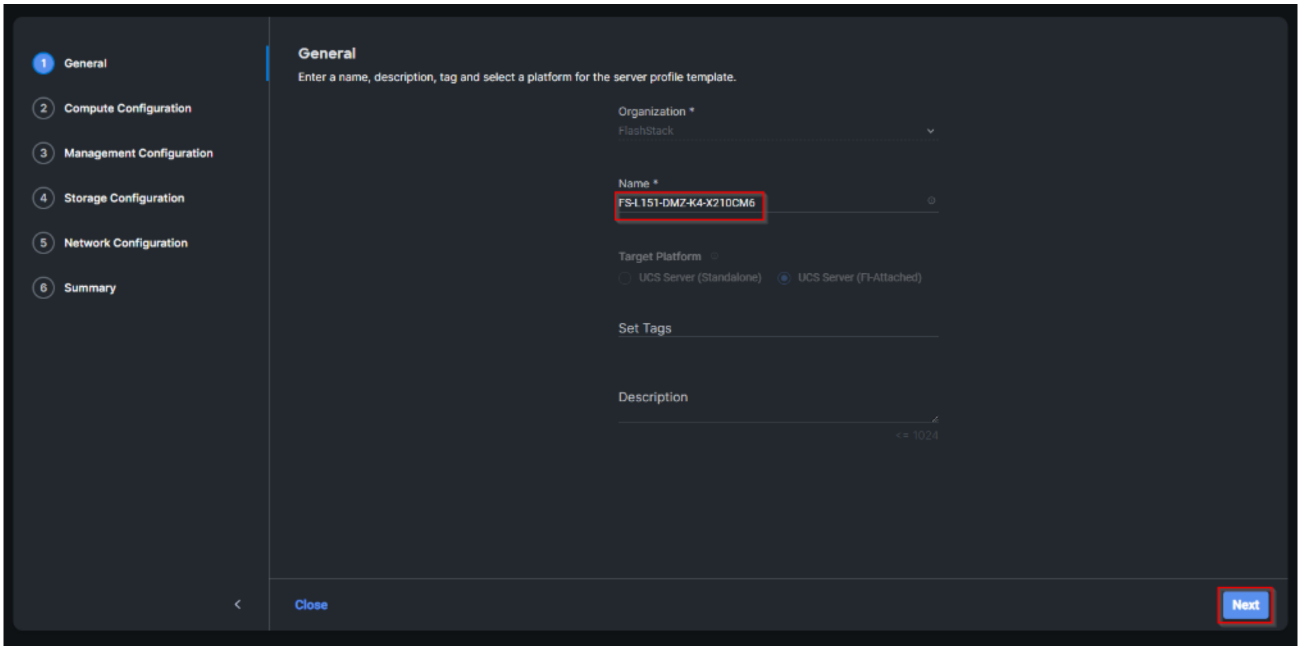

General

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

Organization *
FlashStack

Name *
FS-L151-DMZ-K4-X210CM6

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

Set Tags

Description

<= 1024

Step 4. Click Select Pool under UUID Pool and then click Create New.

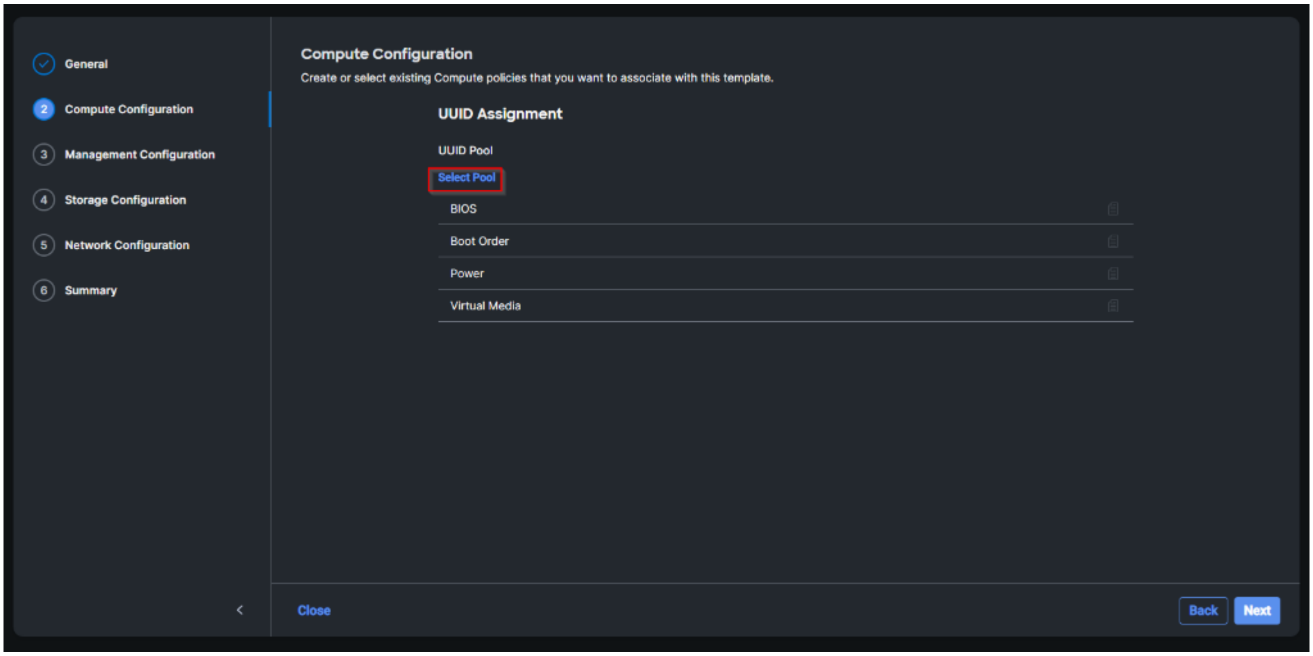

Compute Configuration

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

UUID Assignment

UUID Pool

BIOS

Boot Order

Power

Virtual Media

Step 5. Verify correct organization is selected from the drop-down list and provide a name for the UUID Pool (for example, FS-L151-DMZ-UUID-Pool). Provide an optional Description and click Next.

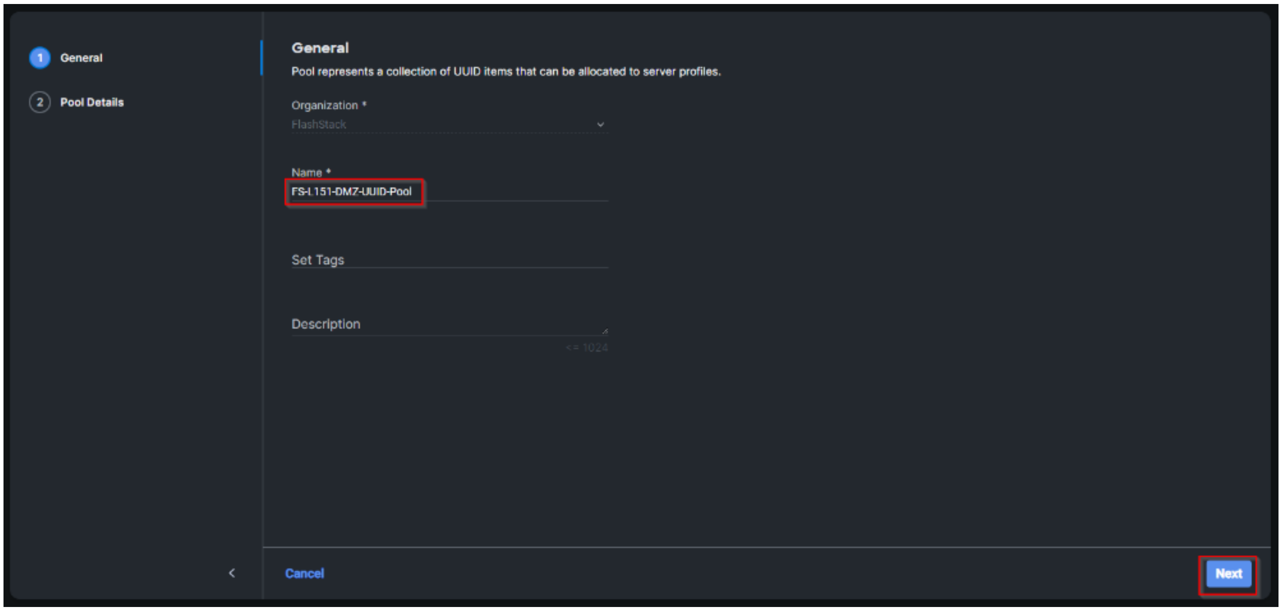

General

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

Organization *
FlashStack

Name *
FS-L151-DMZ-UUID-Pool

Set Tags

Description
<= 1024

Step 6. Provide a UUID Prefix (for example, a random prefix of A11A14B6-B193-49C7 was used). Add a UUID block of appropriate size. Click Create.

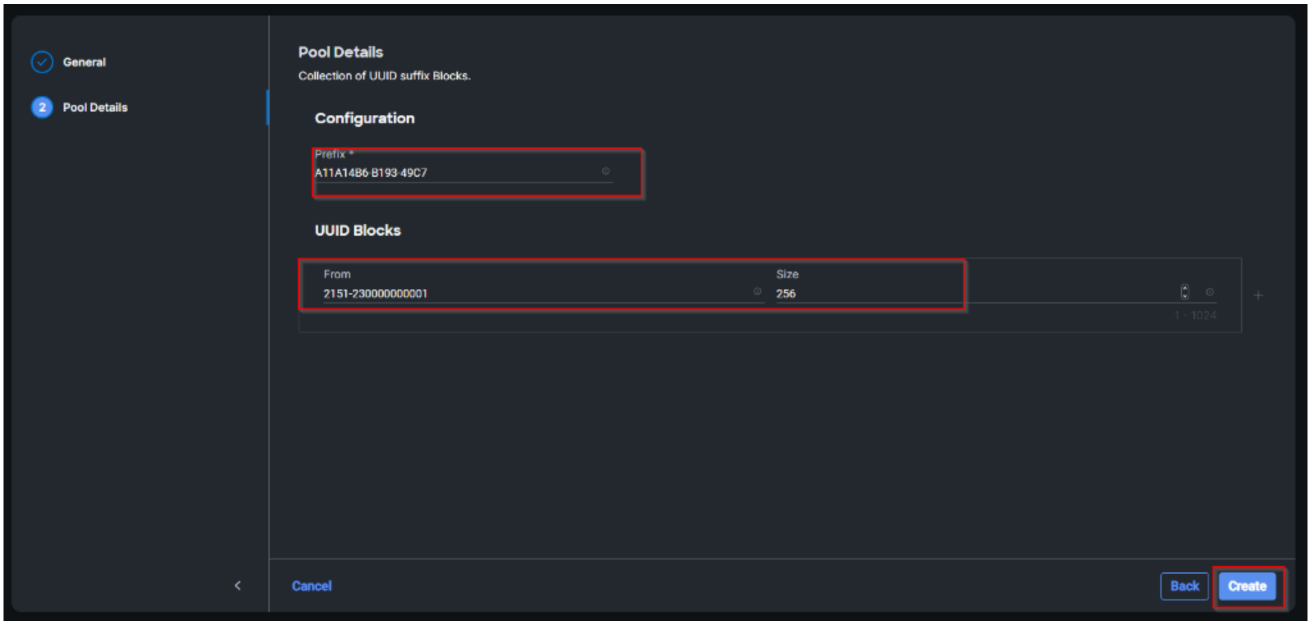

Pool Details

Collection of UUID suffix Blocks.

Configuration

Prefix *
A11A14B6-B193-49C7

UUID Blocks

From
2151-230000000001

Size
256
1 - 1024

Step 7. Click Select Policy next to BIOS and click Create New.

Step 8. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-M7-BIOS-Perf).

Step 9. Click Next.

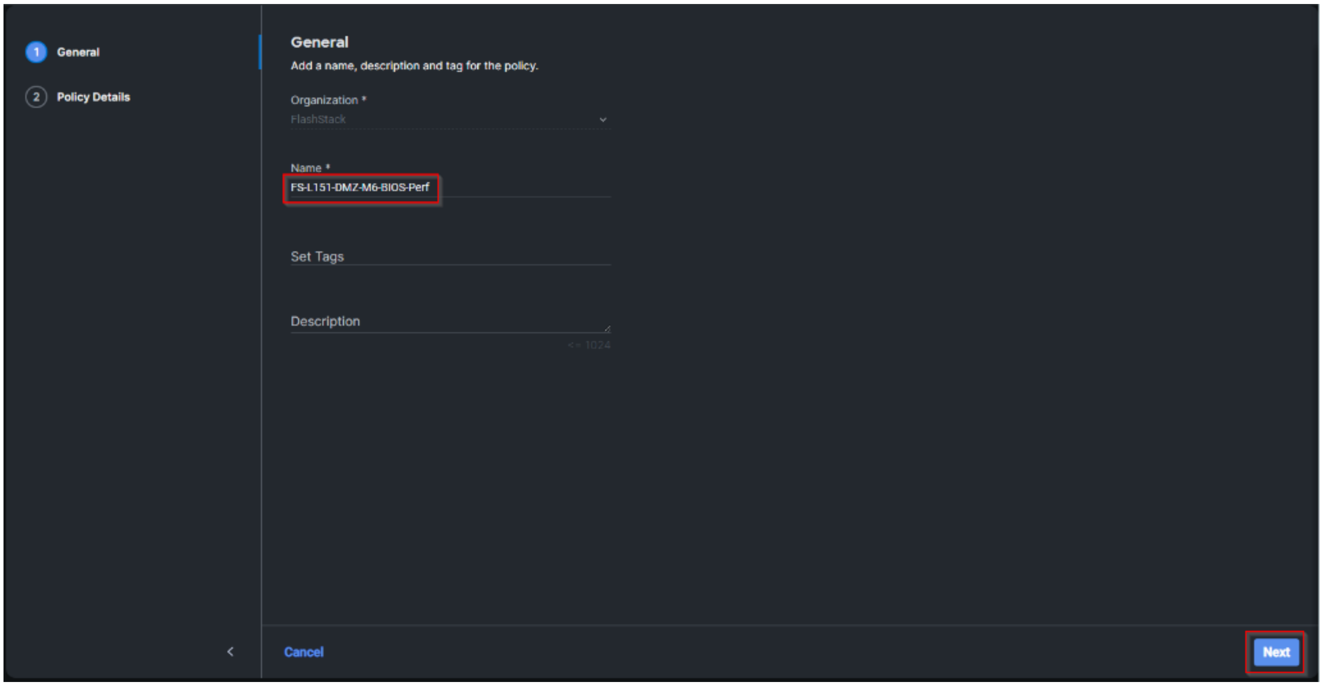

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L151-DMZ-M6-BIOS-Perf

Set Tags

Description

<= 1024

Step 10. On the Policy Details screen, select appropriate values for the BIOS settings. Click Create.

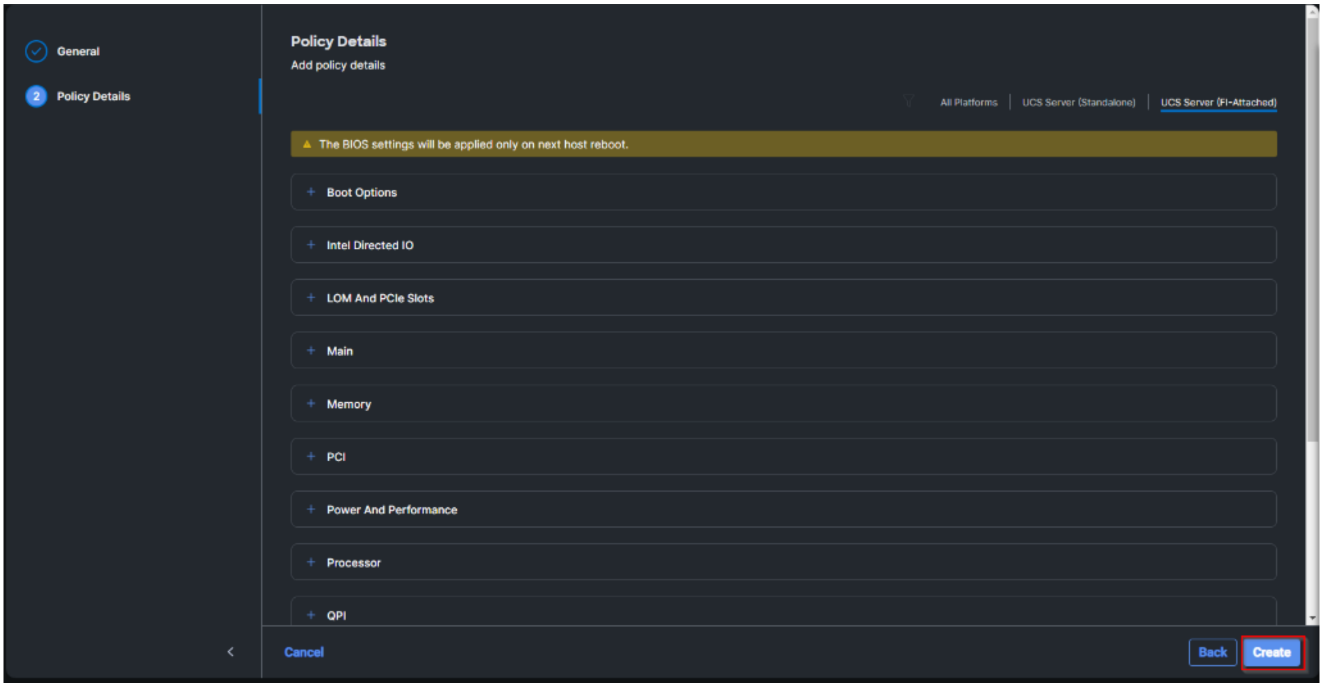

Policy Details

Add policy details

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

Note: 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.

Table 5. FS-L151-DMZ-M7-BIOS-Perf token values

BIOS Token	Value
Intel Directed IO	

BIOS Token	Value
Intel VT for Directed IO	enabled
Memory	
Memory RAS Configuration	maximum-performance
Power And Performance	
Core Performance Boost	Auto
Enhanced CPU Performance	Auto
LLC Dead Line	disabled
UPI Link Enablement	1
UPI Power Management	enabled
Processor	
Altitude	auto
Boot Performance Mode	Max Performance
Core Multi Processing	all
CPU Performance	enterprise
Power Technology	performance
Direct Cache Access Support	enabled
DRAM Clock Throttling	Performance
Enhanced Intel Speedstep(R) Technology	enabled
Execute Disable Bit	enabled
IMC Interleaving	1-way Interleave
Intel HyperThreading Tech	Enabled
Intel Turbo Boost Tech	enabled
Intel(R) VT	enabled
DCU IP Prefetcher	enabled
Processor C1E	disabled
Processor C3 Report	disabled
Processor C6 Report	disabled
CPU C State	disabled

BIOS Token	Value
Sub Numa Clustering	enabled
DCU Streamer Prefetch	enabled

Step 11. Click Select Policy next to Boot Order and then click Create New.

Step 12. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-BootPol). Click Next.

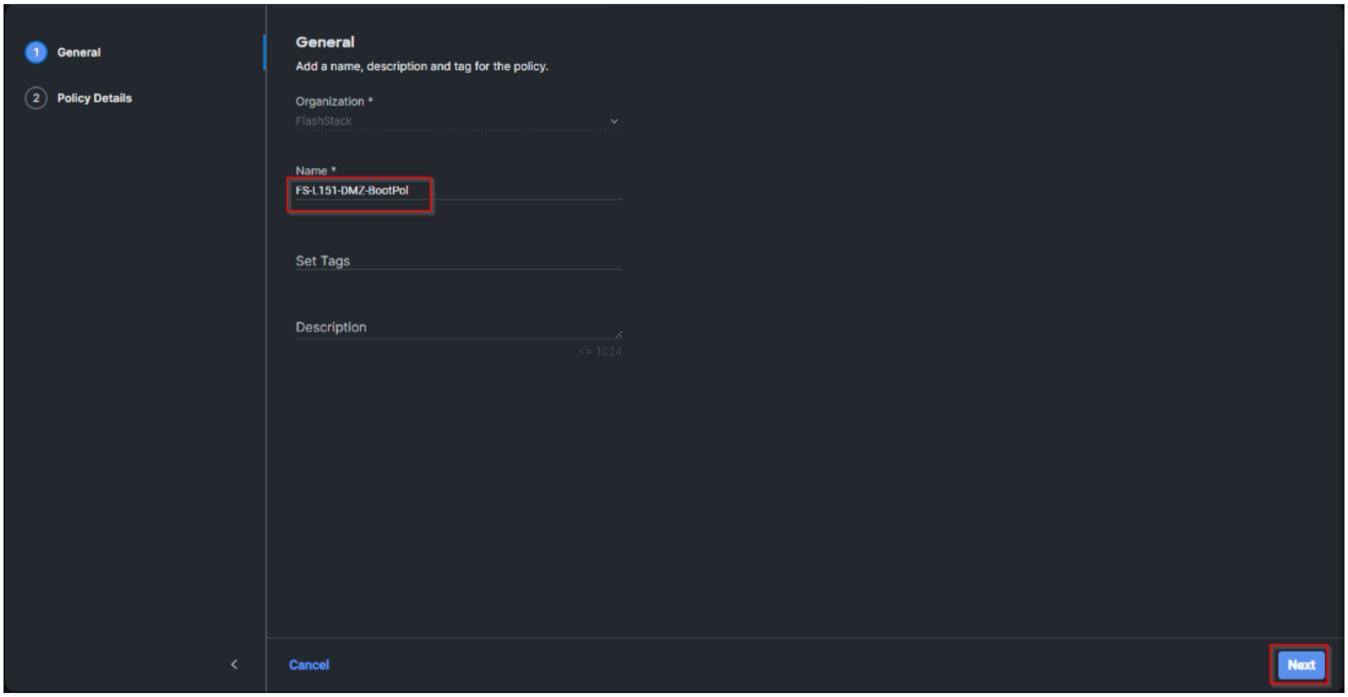

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L151-DMZ-BootPol

Set Tags

Description

<= 1024

Step 13. For Configured Boot Mode, select Unified Extensible Firmware Interface (UEFI).

Step 14. Turn on Enable Secure Boot.

Step 15. Click Add Boot Device drop-down list and select Virtual Media.

Step 16. Provide a device name (for example, vKVM-DVD) and then, for the subtype, select KVM Mapped DVD.

For Fibre Channel SAN boot, four connected FC ports on Pure Storage FlashArray//X R3 controllers will be added as boot options. The four FC ports are as follows:

- CT0.FC0, CT1.FC0 are connected to SAN-A
- CT1.FC2, CT0.FC2 are connected to SAN-B

Figure 18. Pure Storage FlashArray//X R3 Array Ports

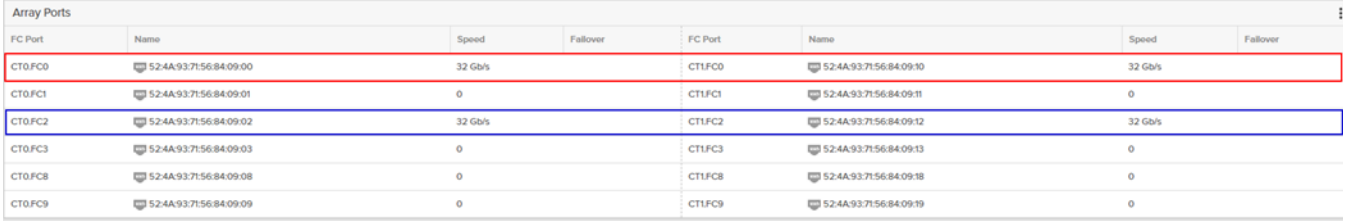

Array Ports

FC Port	Name	Speed	Failover	FC Port	Name	Speed	Failover
CT0.FC0	52:4A:93:71:56:84:09:00	32 Gb/s		CT1.FC0	52:4A:93:71:56:84:09:10	32 Gb/s	
CT0.FC1	52:4A:93:71:56:84:09:01	0		CT1.FC1	52:4A:93:71:56:84:09:11	0	
CT0.FC2	52:4A:93:71:56:84:09:02	32 Gb/s		CT1.FC2	52:4A:93:71:56:84:09:12	32 Gb/s	
CT0.FC3	52:4A:93:71:56:84:09:03	0		CT1.FC3	52:4A:93:71:56:84:09:13	0	
CT0.FC8	52:4A:93:71:56:84:09:08	0		CT1.FC8	52:4A:93:71:56:84:09:18	0	
CT0.FC9	52:4A:93:71:56:84:09:09	0		CT1.FC9	52:4A:93:71:56:84:09:19	0	

Step 17. From the Add Boot Device drop-down list, select SAN Boot (Repeat steps for all 4 FC ports)

Step 18. Provide the Device Name: CT0FC0 and the Logical Unit Number (LUN) value (for example, 1).

Step 19. Provide an interface name vHBA-A. This value is important and should match the vHBA name.

Note: vHBA-A is used to access CT0.FC0, CT1.FC0 and vHBA-B is used to access CT1.FC2, CT0.FC2.

Step 20. Add the appropriate World Wide Port Name (WWPN) as the Target WWPN (for example, 52:4A:93:71:56:84:09:00).

Step 21. Provide bootloader name as BOOTX64.EFI.

Step 22. Provide bootloader name as \EFI\BOOT.

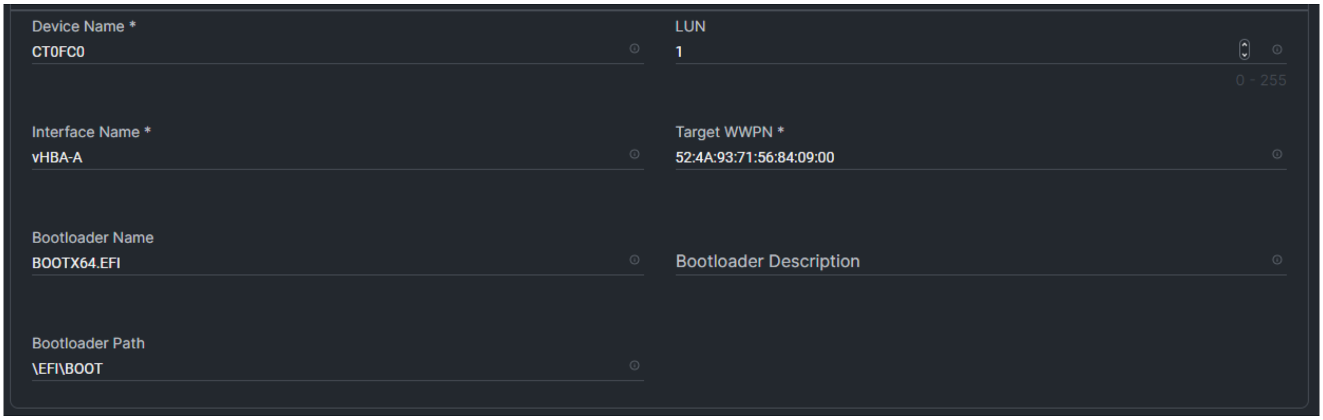

Device Name * CT0FC0

LUN 1 (0 - 255)

Interface Name * vHBA-A

Target WWPN * 52:4A:93:71:56:84:09:00

Bootloader Name BOOTX64.EFI

Bootloader Description

Bootloader Path \EFI\BOOT

Step 23. Verify the order of the boot policies and adjust the boot order as necessary using the arrows next to delete icon. Click Create.

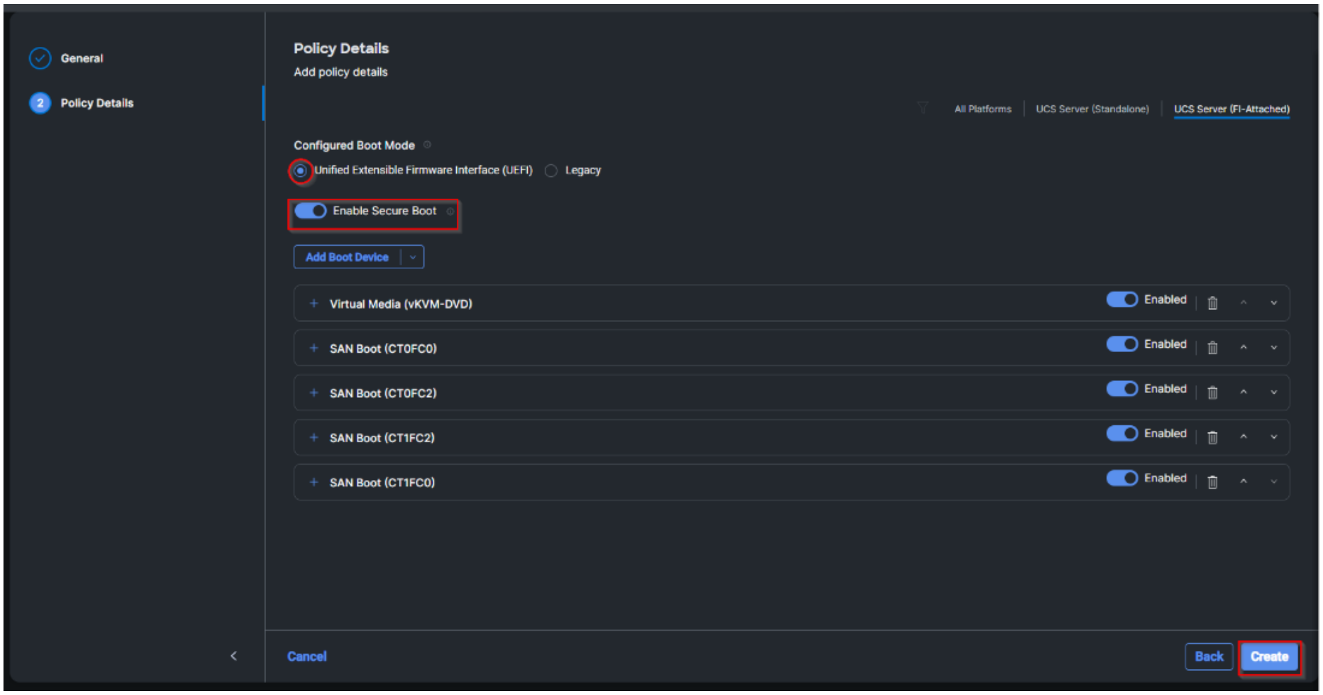

Policy Details

Add policy details

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

Configured Boot Mode

● Unified Extensible Firmware Interface (UEFI) ○ Legacy

Enable Secure Boot (enabled)

- Virtual Media (vKVM-DVD) — Enabled
- SAN Boot (CT0FC0) — Enabled
- SAN Boot (CT0FC2) — Enabled
- SAN Boot (CT1FC2) — Enabled
- SAN Boot (CT1FC0) — Enabled

Step 24. Click Select Policy next to Power and click Create New.

Step 25. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, UCS-PWR). Click Next.

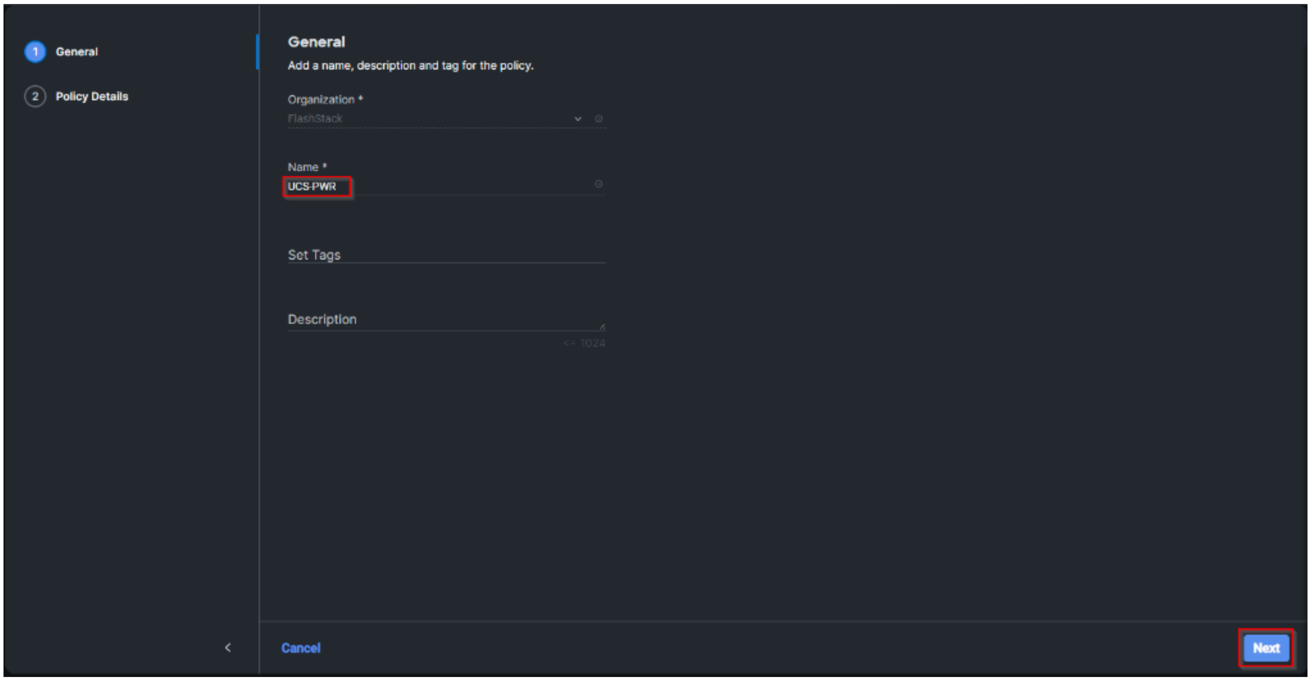

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * UCS-PWR

Set Tags

Description

<= 1024

Step 26. Enable Power Profiling and select High from the Power Priority drop-down list. Click Create.

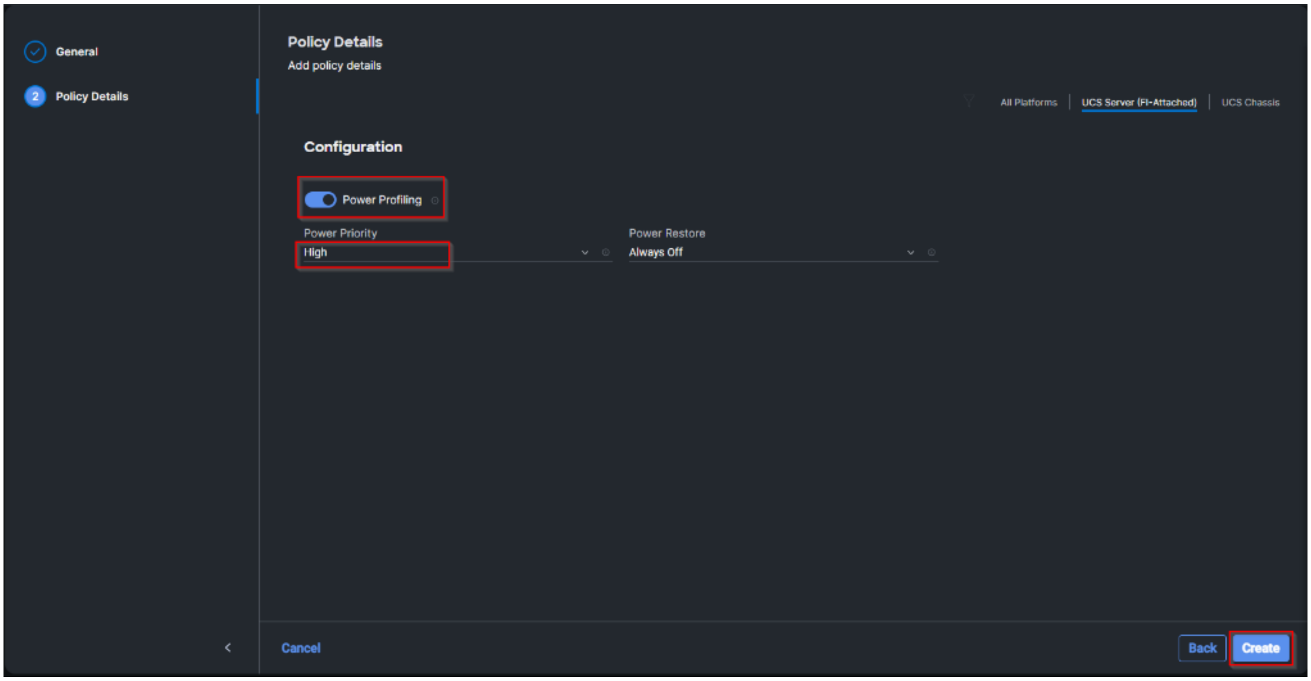

Policy Details

Add policy details

Configuration

Power Profiling

Power Priority: High

Power Restore: Always Off

Step 27. Click Select Policy next to Virtual Media and click Create New (Optional).

Step 28. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-Vmedia-Pol). Click Next.

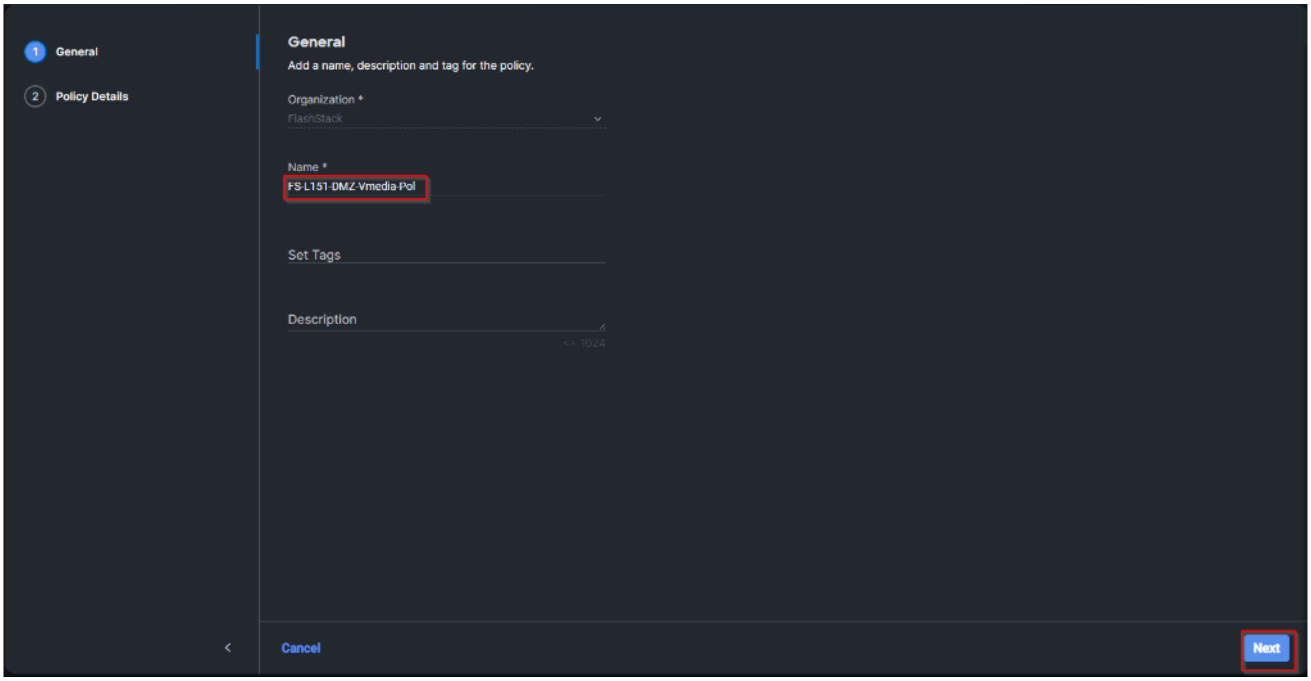
1 General

2 Policy Details

General

Add a name, description and tag for the policy.

Organization *
FlashStack

Name *
FS-L151-DMZ-Vmedia-Pol

Set Tags

Description
<= 1024

Cancel

Next

Step 29. Disable Lower Power USB and click Create.

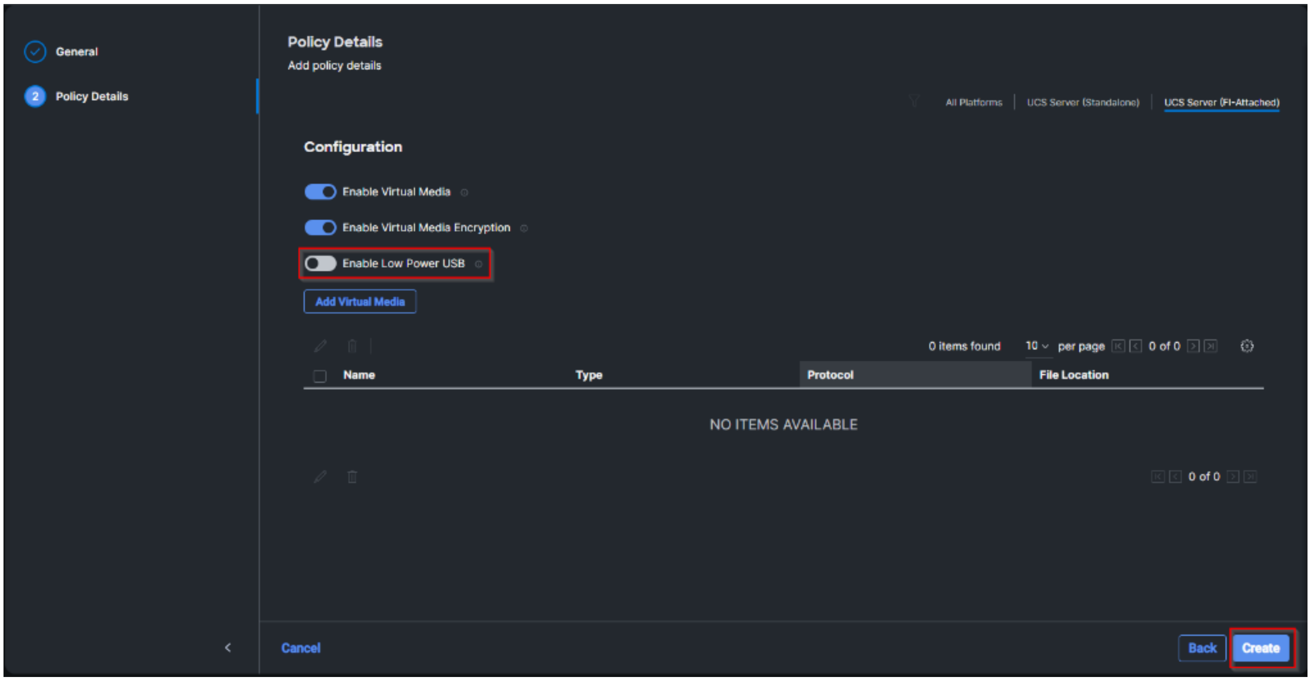
General

2 Policy Details

Policy Details

Add policy details

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

Configuration

Enable Virtual Media

Enable Virtual Media Encryption

Enable Low Power USB

Add Virtual Media

0 items found 10 per page 0 of 0

Name	Type	Protocol	File Location

NO ITEMS AVAILABLE

0 of 0

Cancel

Back

Create

Step 30. Click Next to go to Management Configuration.

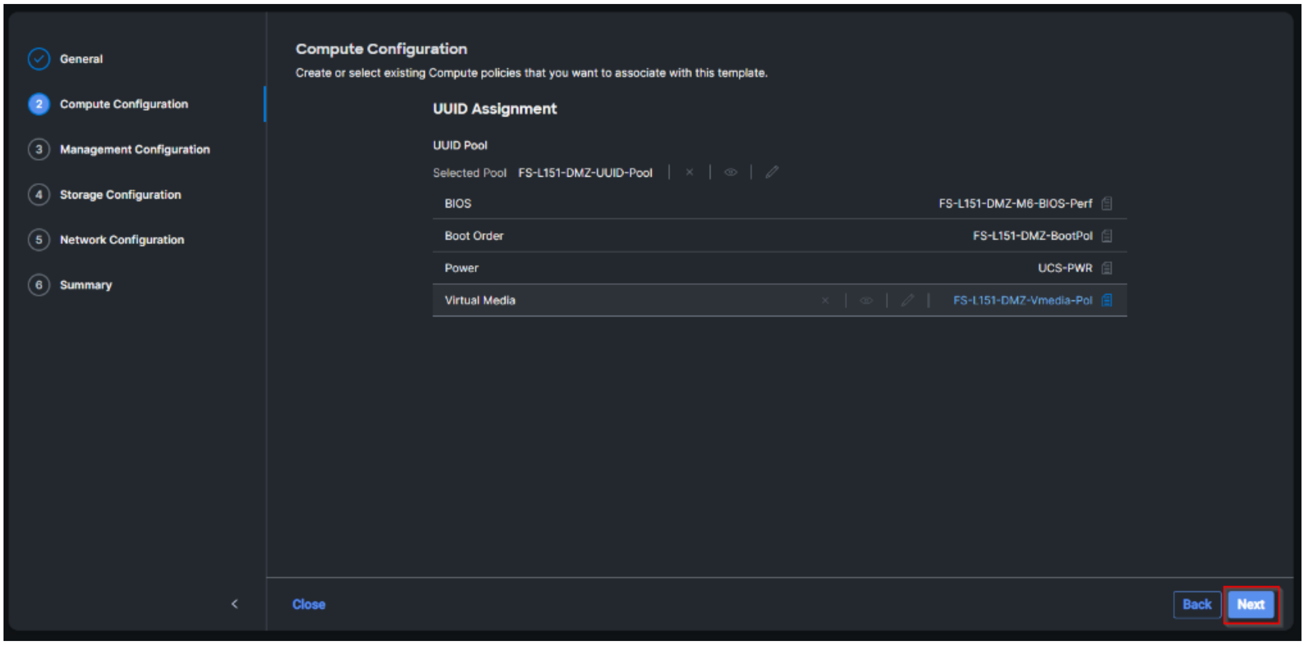

Compute Configuration

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

UUID Assignment

UUID Pool

Selected Pool FS-L151-DMZ-UUID-Pool

BIOS	FS-L151-DMZ-M6-BIOS-Perf
Boot Order	FS-L151-DMZ-BootPol
Power	UCS-PWR
Virtual Media	FS-L151-DMZ-Vmedia-Pol

Step 31. Click Select Policy next to IMC Access and then click Create New.

Step 32. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-IMCAPol). Click Next.

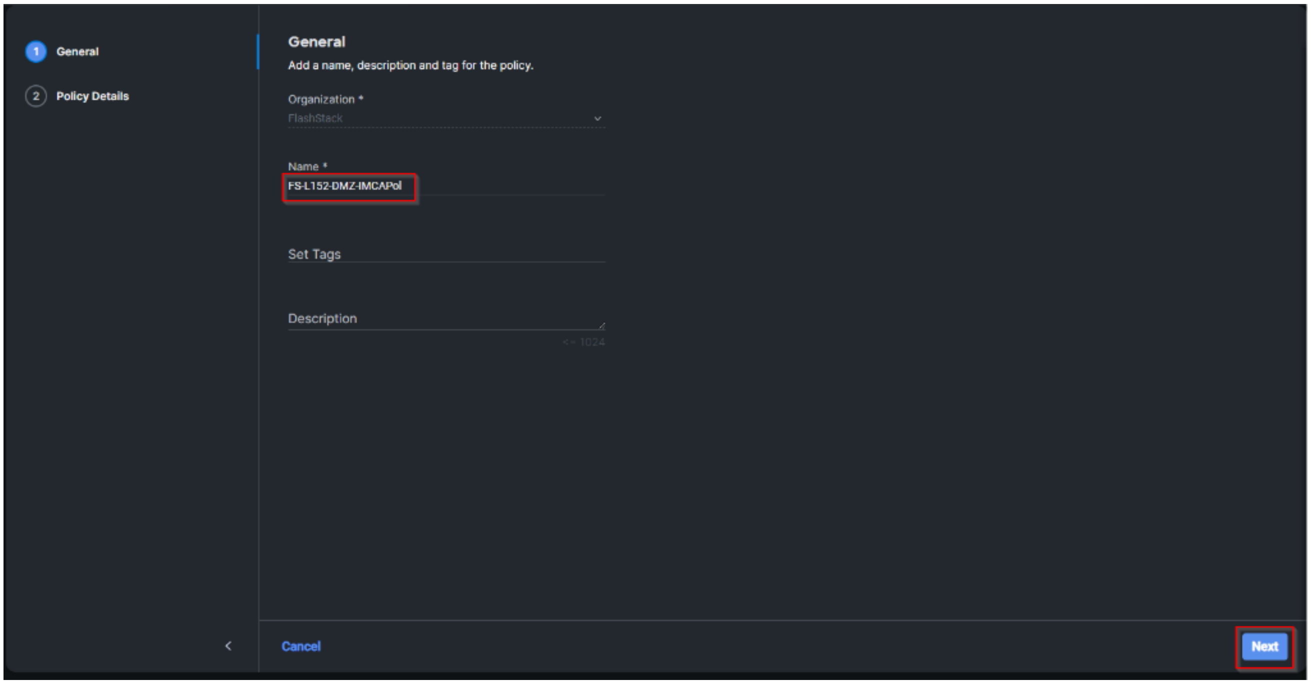

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-IMCAPol

Set Tags

Description

<= 1024

Note: You can select in-band management access to the compute node using an in-band management VLAN (for example, VLAN 70) 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.

Step 33. Click UCS Server (FI-Attached). Enable In-Band Configuration and type VLAN Id designated for the In-Band management (for example, 70).

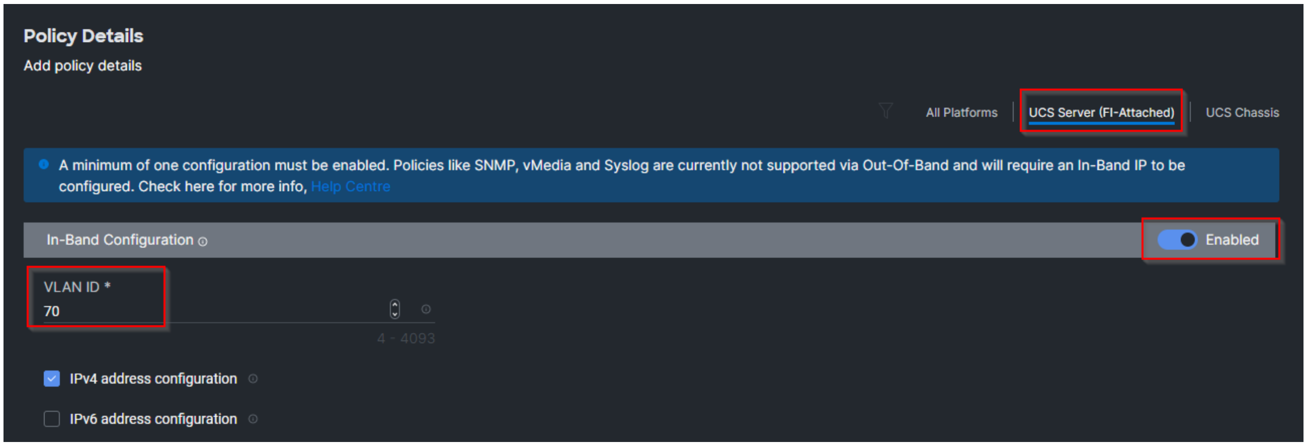

Policy Details

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 * 70 (4 - 4093)

☑ IPv4 address configuration

☐ IPv6 address configuration

Step 34. Under IP Pool, click Select IP Pool and then click Create New.

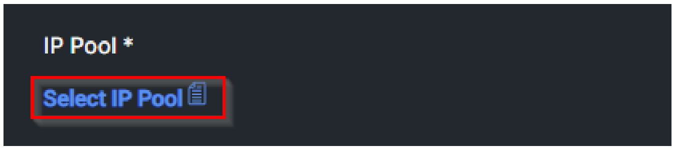

IP Pool *

Step 35. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-ICMA-IP-Pool). Click Next.

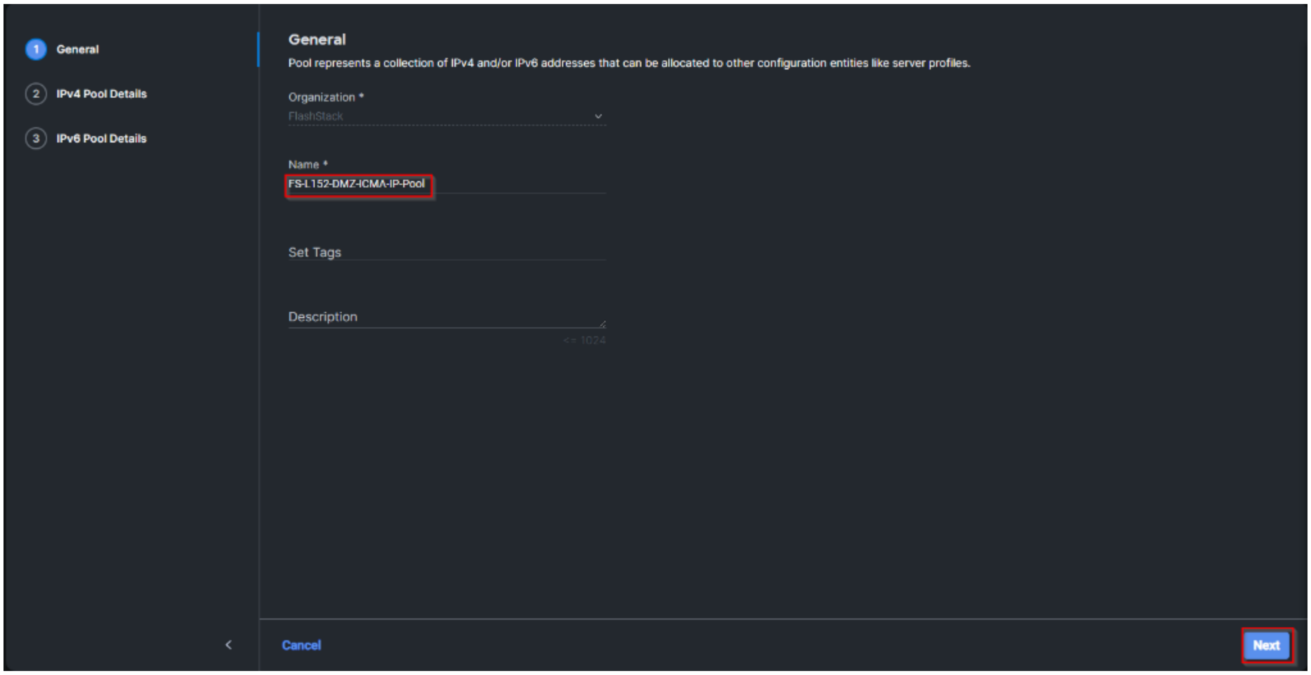

General

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

Organization * FlashStack

Name * FS-L152-DMZ-ICMA-IP-Pool

Set Tags

Description

<= 1024

Step 36. Select Configure IPv4 Pool and provide the information to define a pool for KVM IP address assignment including an IP Block.

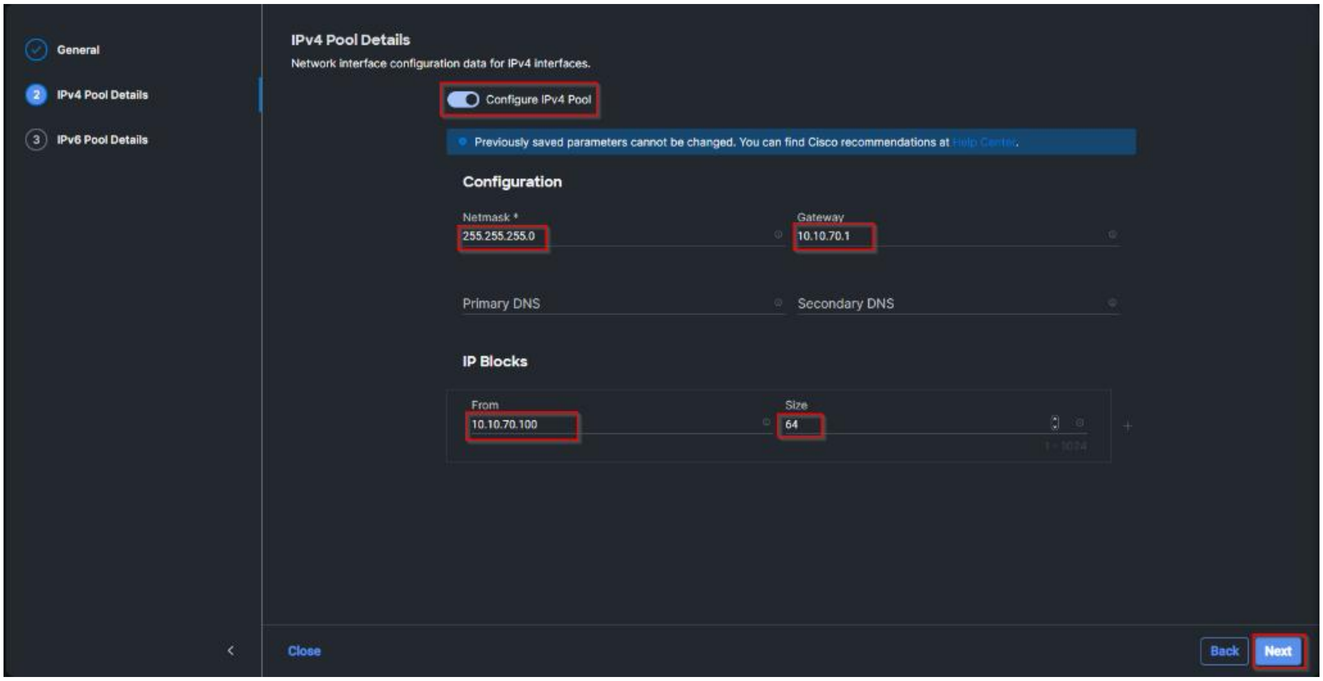

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Previously saved parameters cannot be changed. You can find Cisco recommendations at Help Center.

Configuration

Netmask *: 255.255.255.0

Gateway: 10.10.70.1

Primary DNS

Secondary DNS

IP Blocks

From: 10.10.70.100

Size: 64

1 - 1024

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 10.10.70.0/24 subnet.

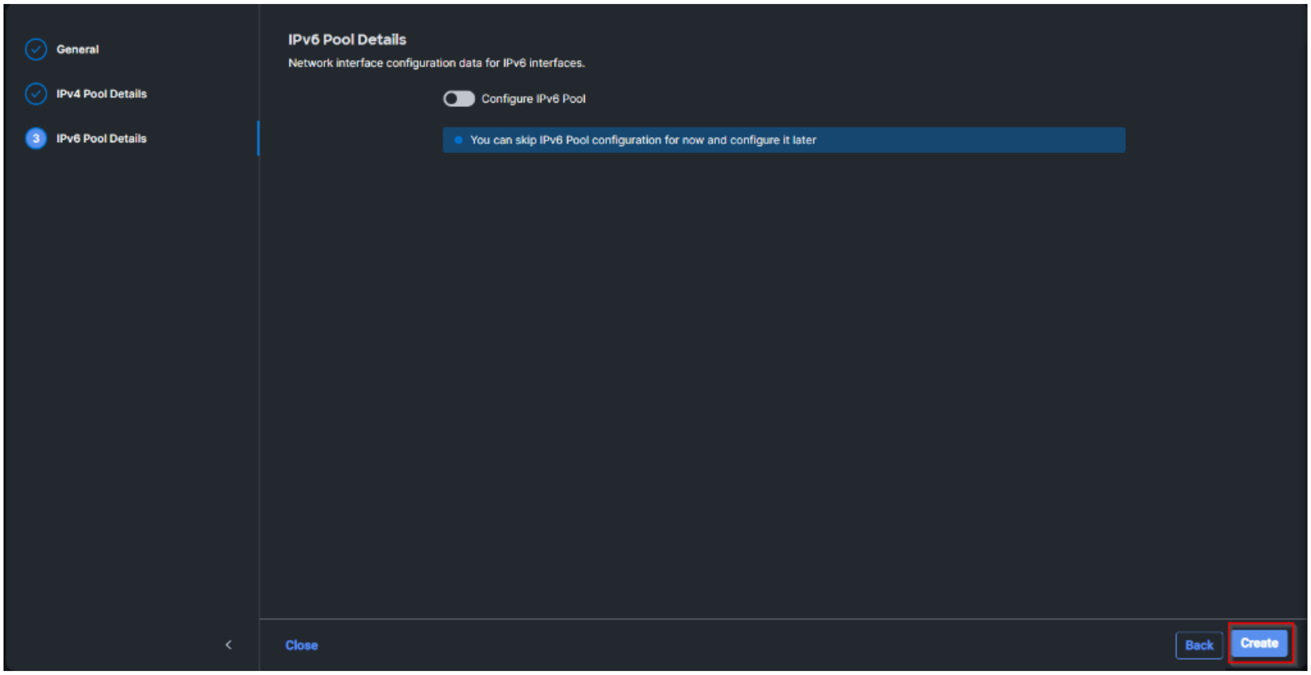

IPv6 Pool Details

Network interface configuration data for IPv6 interfaces.

Configure IPv6 Pool

You can skip IPv6 Pool configuration for now and configure it later

Step 37. Click Select Policy next to IPMI Over LAN and then click Create New.

Step 38. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, Enable-IPMIoLAN). Click Next.

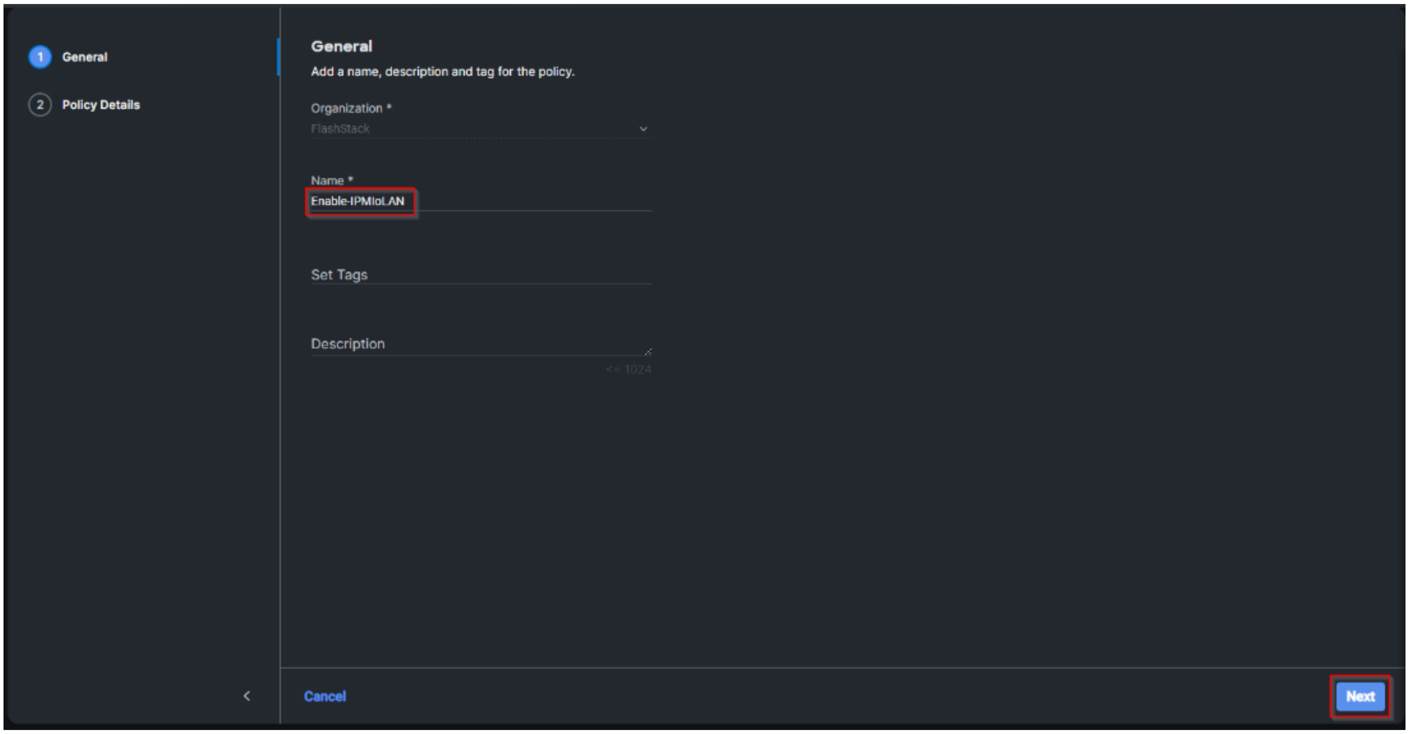

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * Enable-IPMIoLAN

Set Tags

Description

<= 1024

Step 39. Turn on Enable IPMI Over LAN.

Step 40. From the Privilege Level drop-down list, select admin.

Step 41. Click Create.

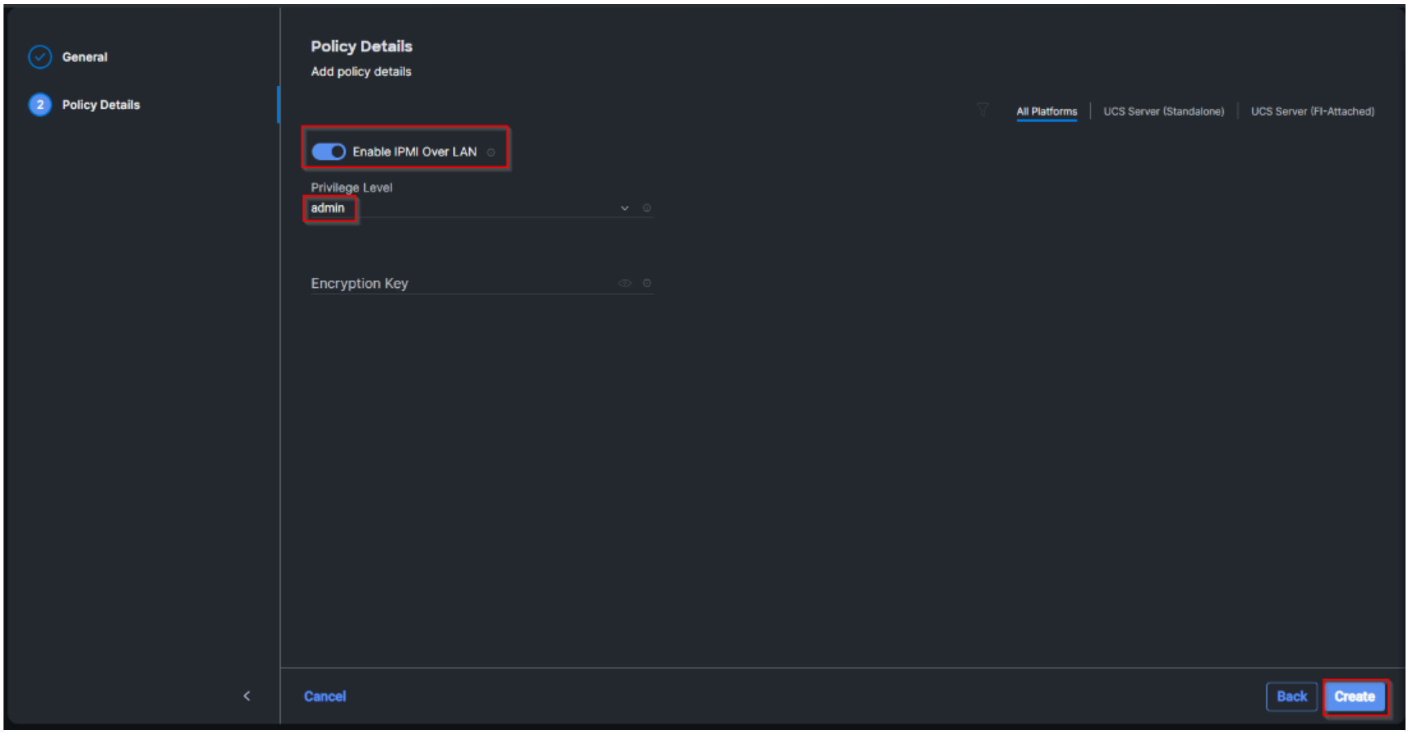

Policy Details

Add policy details

Enable IPMI Over LAN (on)

Privilege Level admin

Encryption Key

Step 42. Click Select Policy next to Local User and click Create New.

Step 43. Verify the correct organization is selected from the drop-down list and provide a name for the policy.

Step 44. Verify that UCS Server (FI-Attached) is selected.

Step 45. Verify that Enforce Strong Password is selected.

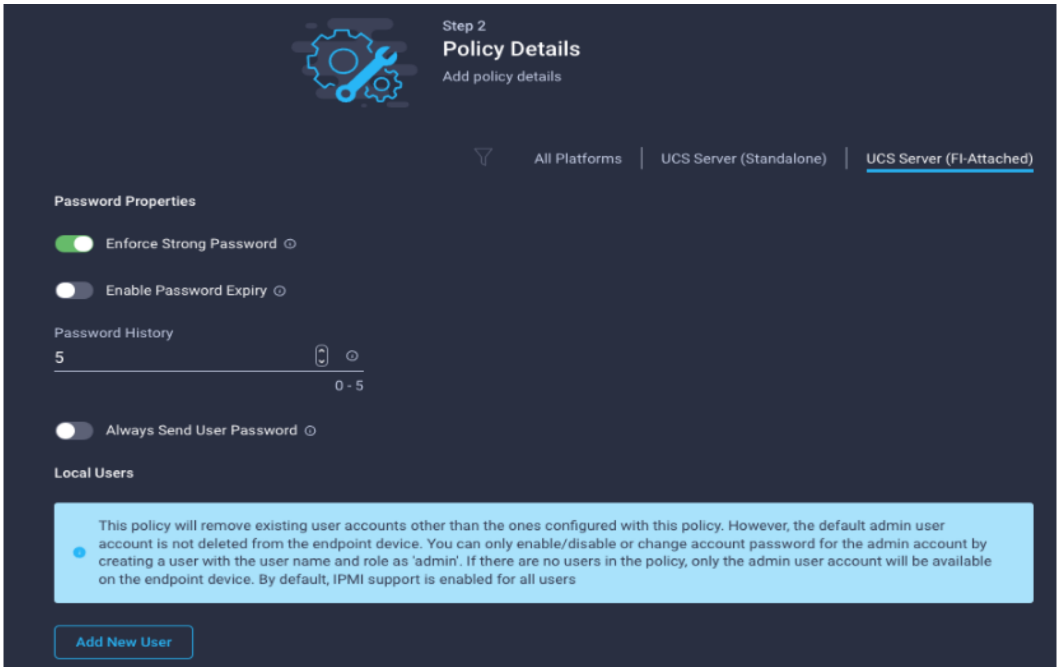

Step 2
Policy Details
Add policy details

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

Step 46. Click Add New User and then click + next to the New User.

Step 47. Provide the username (for example, fsadmin), choose a role for example, admin), and provide a password.

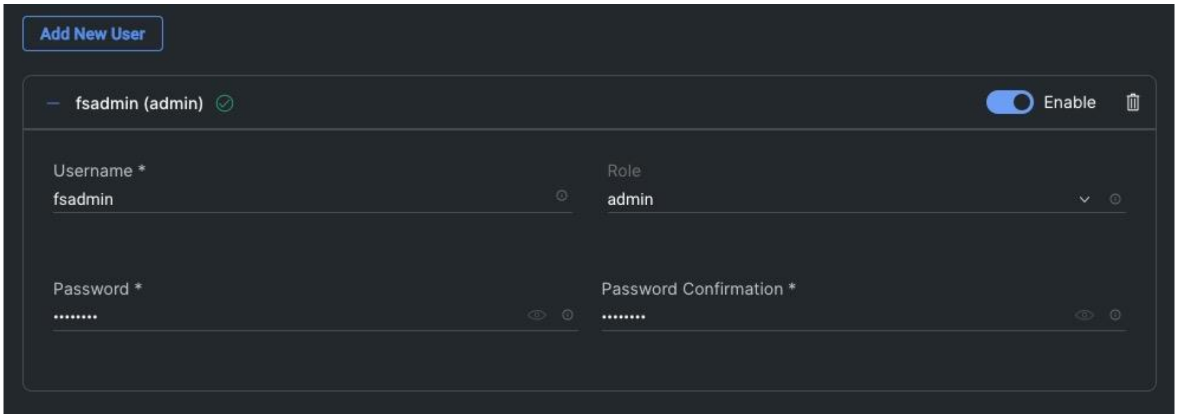

fsadmin (admin)

Username *: fsadmin

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.

Step 48. Click Create to finish configuring the user.

Step 49. Click Create to finish configuring local user policy.

Step 50. Click Next to move to Storage Configuration.

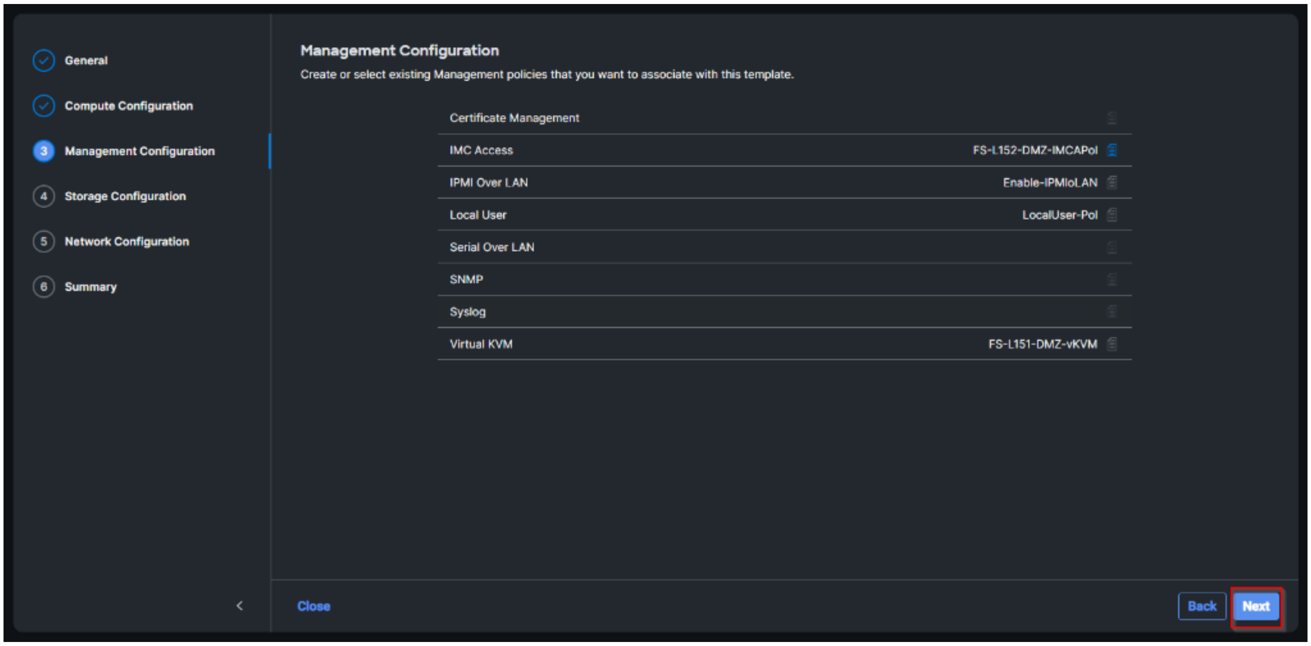

Management Configuration

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

Certificate Management	
IMC Access	FS-L152-DMZ-IMCAPol
IPMI Over LAN	Enable-IPMIoLAN
Local User	LocalUser-Pol
Serial Over LAN	
SNMP	
Syslog	
Virtual KVM	FS-L151-DMZ-vKVM

Step 51. Click Next on the Storage Configuration screen. No configuration is needed in the local storage system.

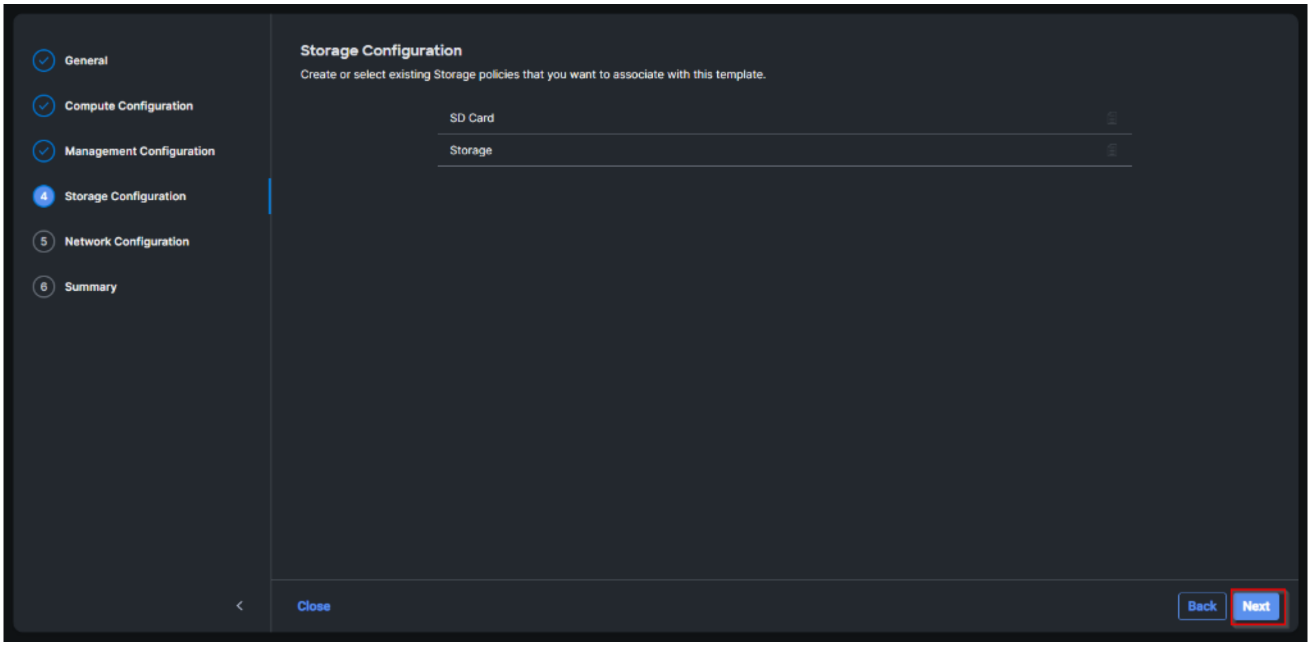

Storage Configuration

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

SD Card	
Storage	

Step 52. Click Select Policy next to LAN Connectivity and then click Create New.

Note: 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 placement, manual vNIC placement is utilized.

The FC boot from SAN hosts that uses four vNICs configured as list in Table 6.

Table 6. vNICs for LAN Connectivity

vNIC	Slot ID	Switch ID	PCI Order	VLANs
vSwitch0-A	MLOM	A	2	FS-InBand-Mgmt_40
vSwitch0-B	MLOM	B	3	FS-InBand-Mgmt_40
VDS0-A	MLOM	A	4	FS-_42, FS-vMotion_43
VDS0-B	MLOM	B	5	FS-_42, FS-vMotion_43

Note: The PCI order 0 and 1 will be used in the SAN Connectivity policy to create vHBA-A and vHBA-B.

Step 53. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-LAN-Conn-Pol). Click Next.

Step 54. Under vNIC Configuration, select Manual vNICs Placement.

Step 55. Click Add vNIC.

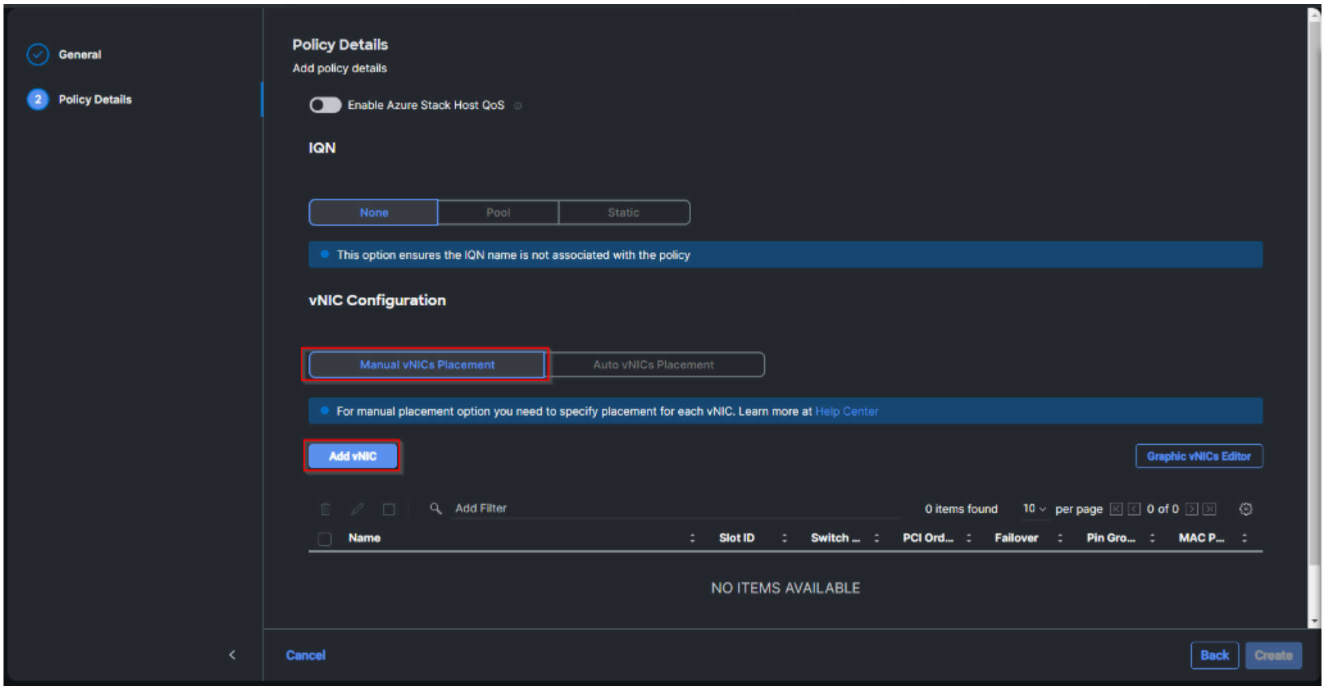

Policy Details

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

0 items found 10 per page 0 of 0

Name	Slot ID	Switch ...	PCI Ord...	Failover	Pin Gro...	MAC P...

NO ITEMS AVAILABLE

Step 56. Click Select Pool under MAC Address Pool and then click Create New.

Note: 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 7. MAC Address Pools

Pool Name	Starting MAC Address	Size	vNICs
FS-L151-DMZ-MAC-Pool-A	00:25:B5:04:0A:00	256*	vSwitch0-A, VDS0-A
FS-L151-DMZ-MAC-Pool-B	00:25:B5:04:0B:00	256*	vSwitch0-B, VDS0-B

Step 57. Verify the correct organization is selected from the drop-down list and provide a name for the pool from Table 7 depending on the vNIC being created (for example, FS-L151-DMZ-MAC-Pool-A for Fabric A).

Step 58. Click Next.

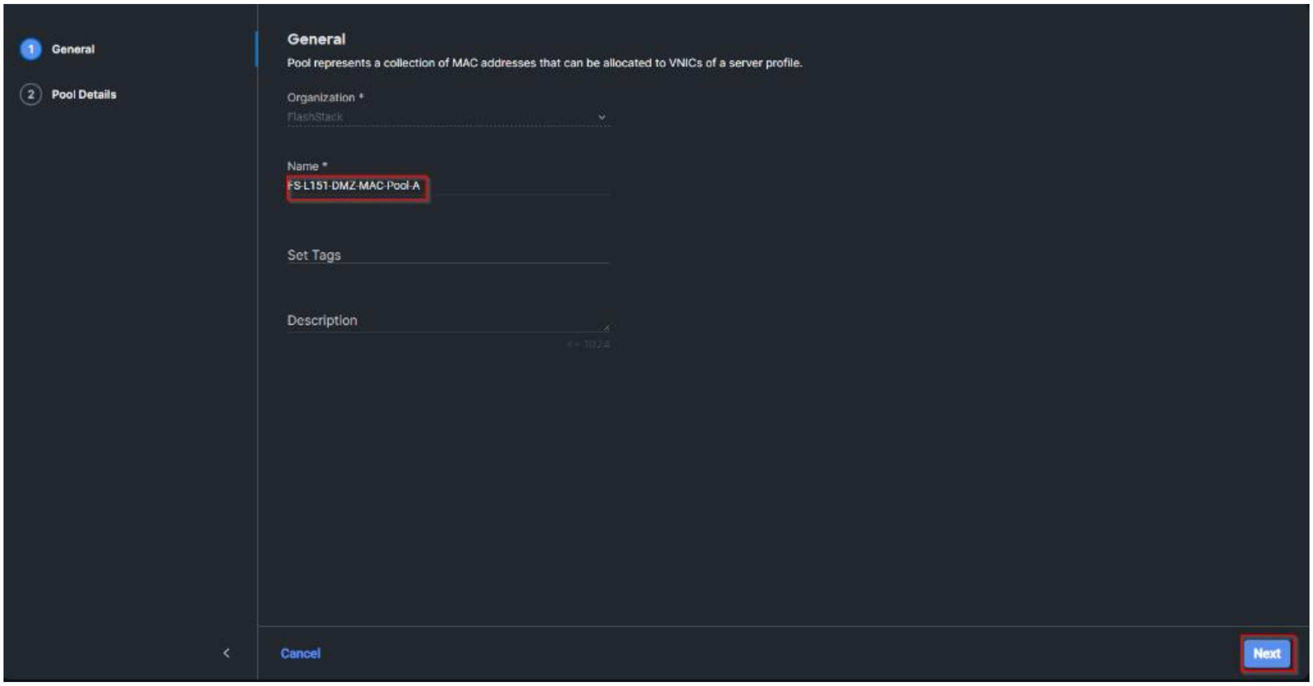

General

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

Organization *: FlashStack

Name *: FS-L151-DMZ-MAC-Pool-A

Set Tags

Description

<= 1024

Step 59. Provide the starting MAC address from Table 7 (for example, 00:25:B5:04:0A:00) and the size of the MAC address pool (for example, 256). Click Create to finish creating the MAC address pool.

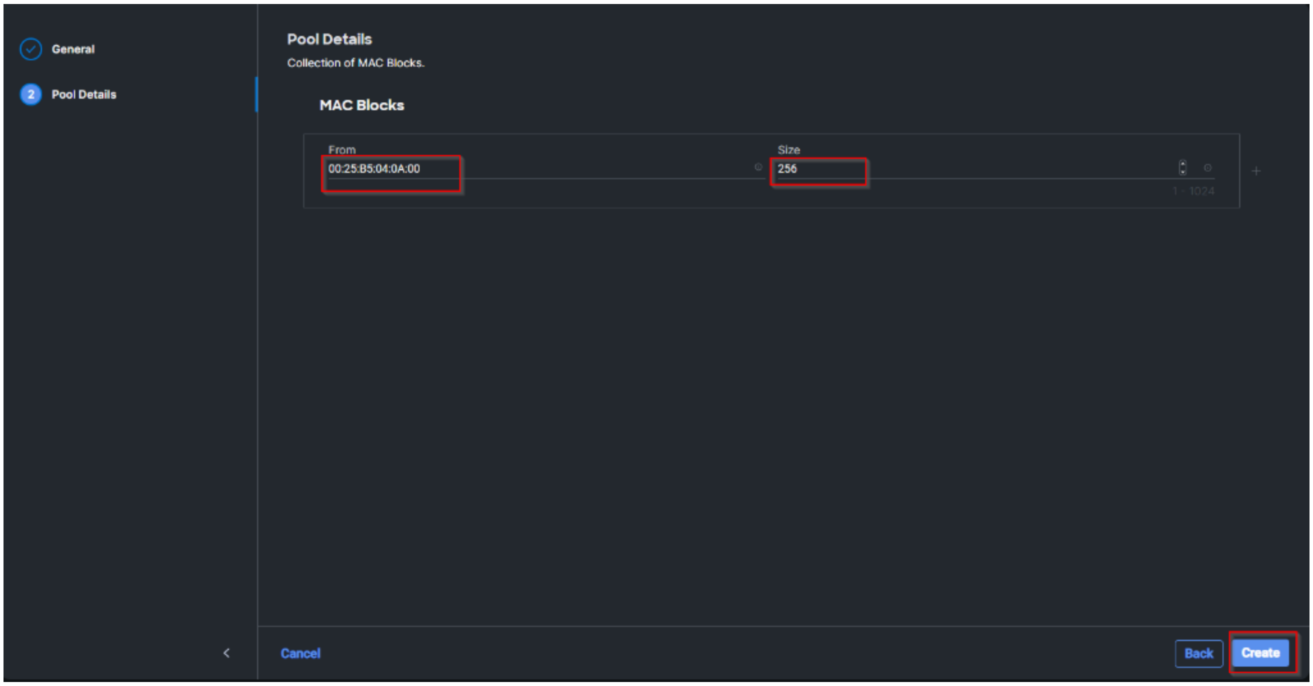

Pool Details

Collection of MAC Blocks.

MAC Blocks

From: 00:25:B5:04:0A:00

Size: 256

1 - 1024

Step 60. From the Add vNIC window, provide vNIC Name, Slot ID, Switch ID, and PCI Order information from Table 7.

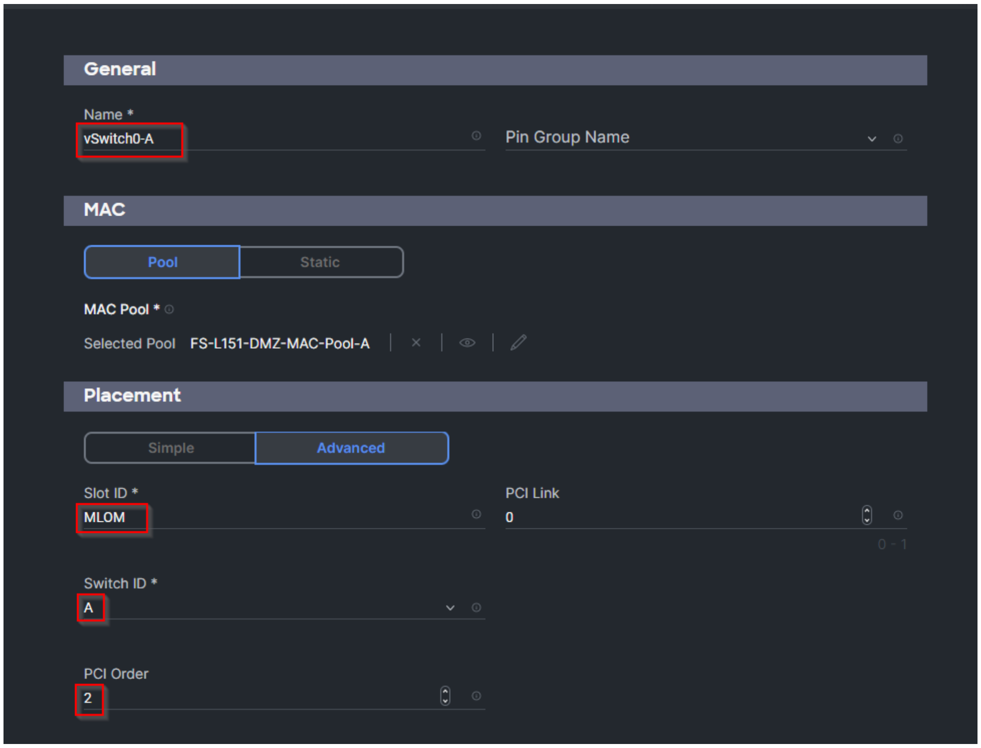

General

Name * vSwitch0-A

Pin Group Name

MAC

MAC Pool *

Selected Pool FS-L151-DMZ-MAC-Pool-A

Placement

Slot ID * MLOM

PCI Link 0

Switch ID * A

PCI Order 2

Step 61. For Consistent Device Naming (CDN), from the drop-down list, select vNIC Name.

Step 62. Verify that Failover is disabled because the failover will be provided by attaching multiple vNICs to the VMware vSwitch and VDS.

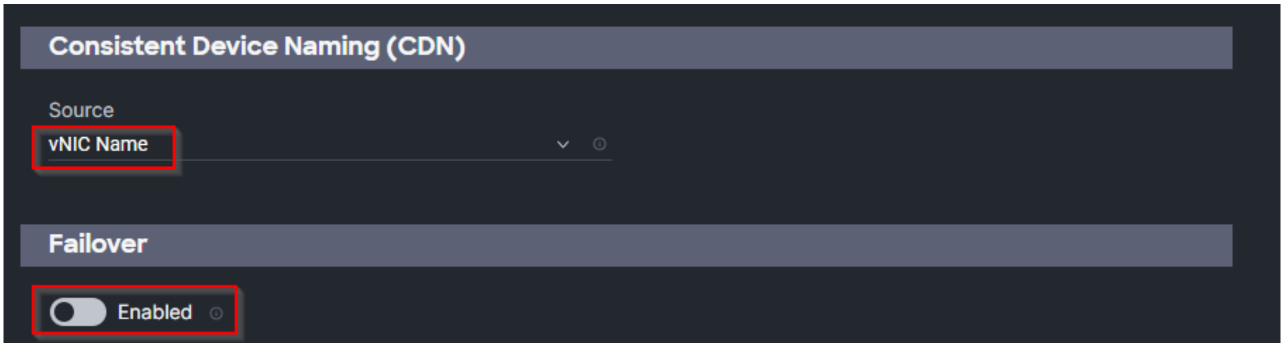

Consistent Device Naming (CDN)

Source vNIC Name

Failover

Enabled

Step 63. Click Select Policy under Ethernet Network Group Policy and then click Create New.

Note: The Ethernet Network Group policies will be created and reused on applicable vNICs as explained below. The Ethernet Network Group policy defines the VLANs allowed for a particular vNIC, therefore multiple network group policies will be defined for this deployment as listed in Table 8.

Table 8. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs	VLANs
FS-L151-DMZ-vSwitch0-NetGrp-Pol	Native-VLAN (1)	vSwitch0-A, vSwitch0-B	FS-InBand-Mgmt_40
FS-L151-DMZ-vSwitch1-NetGrp-Pol	Native-VLAN (1)	VDS0-A, VDS0-B	FS-_42, FS-vMotion_43

Step 64. Verify the correct organization is selected from the drop-down list and provide a name for the policy from Table 8 (for example, FS-L151-DMZ-vSwitch0-NetGrp-Pol). Click Next.

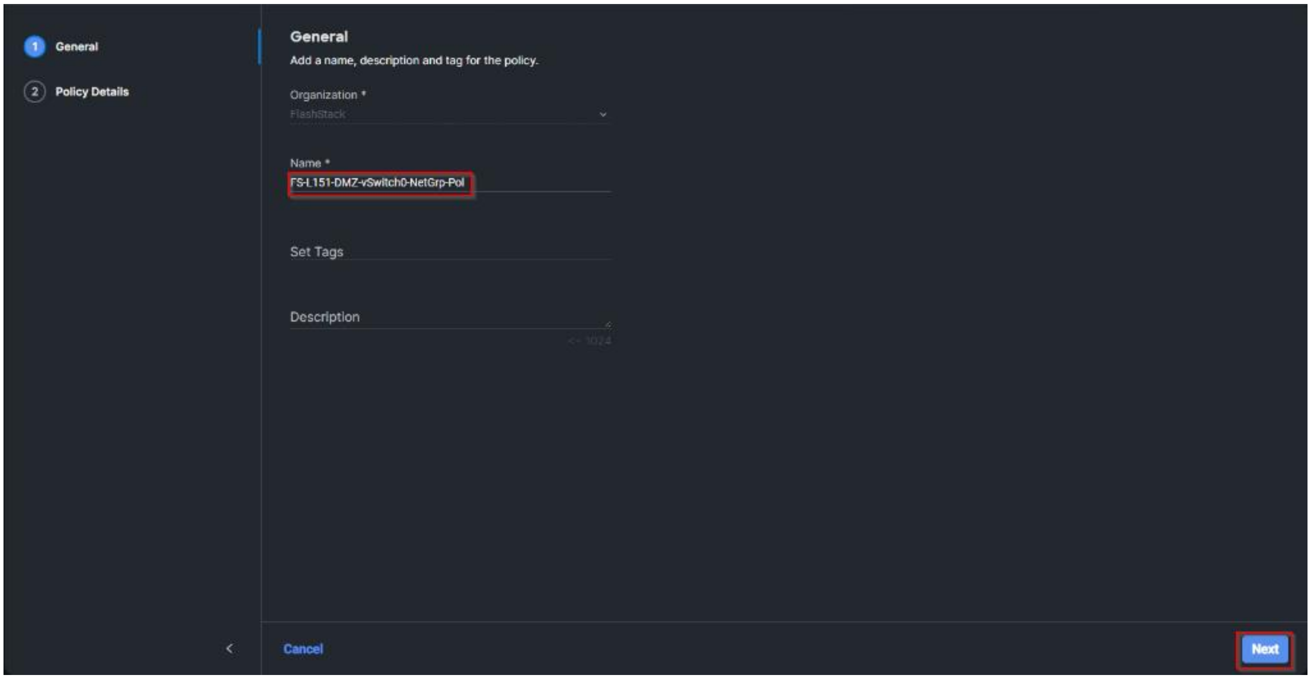

Organization *
FlashStack

Name *
FS-L151-DMZ-vSwitch0-NetGrp-Pol

Step 65. Enter the allowed VLANs from Table 7 (for example, 70) and the native VLAN ID from Table 8 (for example, 1). Click Create.

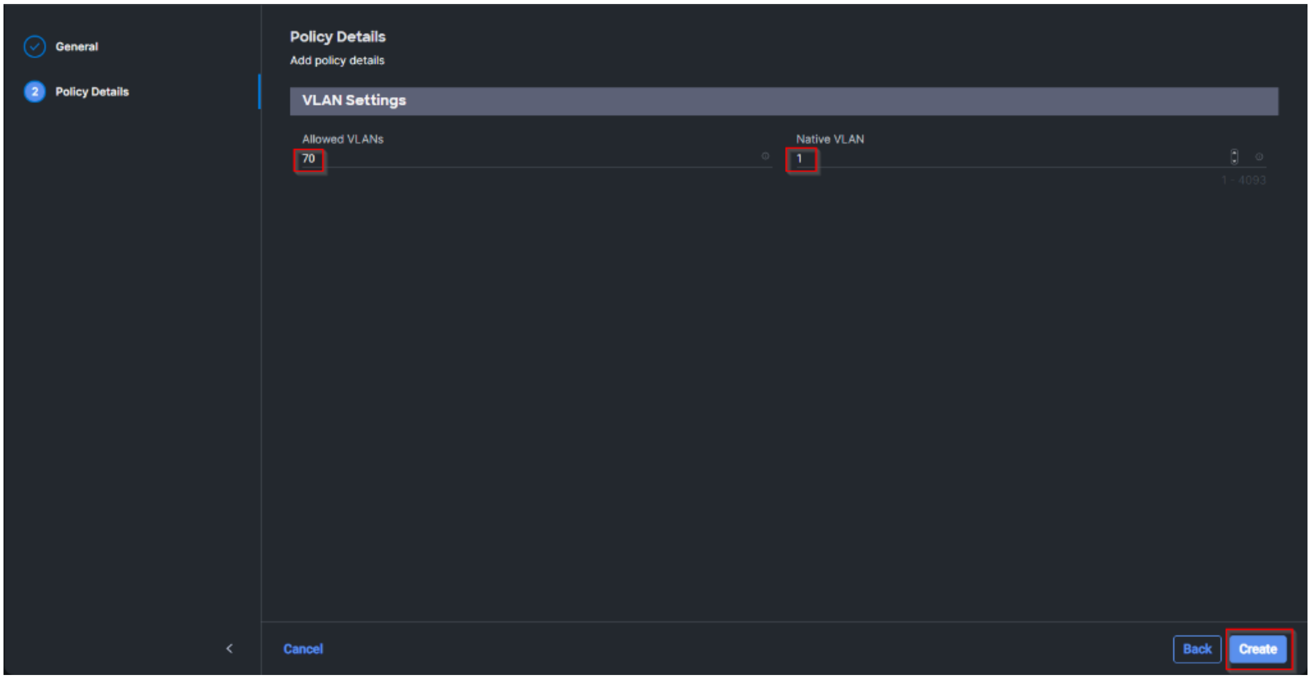

Allowed VLANs
70

Native VLAN
1

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, click Select Policy and pick the previously defined ethernet group policy from the list on the right.

Step 66. Click Select Policy under Ethernet Network Control Policy and then click Create New.

Note: The 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 and reused for all the vNICs.

Step 67. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-NetCtrl-Pol).

Step 68. Click Next.

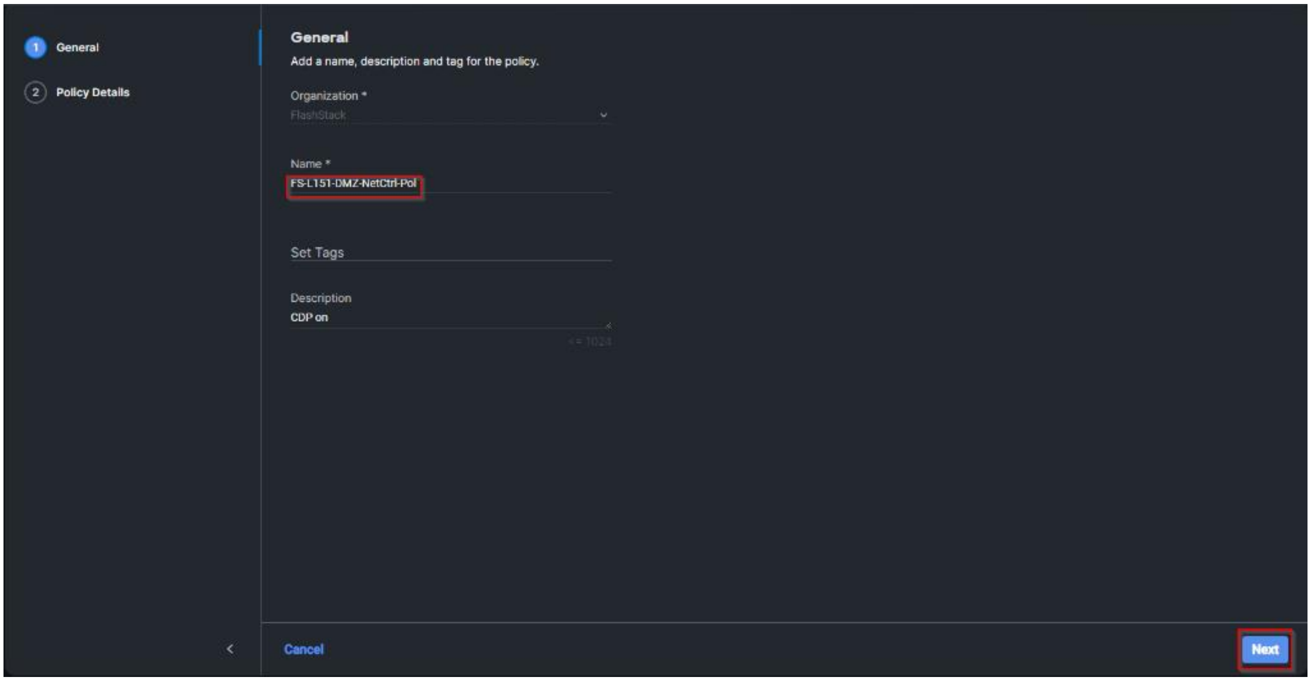

1 General

2 Policy Details

General

Add a name, description and tag for the policy.

Organization *: FlashStack

Name *: FS-L151-DMZ-NetCtrl-Pol

Set Tags

Description: CDP on

<= 1024

Cancel

Next

Step 69. Enable Cisco Discovery Protocol and both Enable Transmit and Enable Receive under LLDP. Click Create.

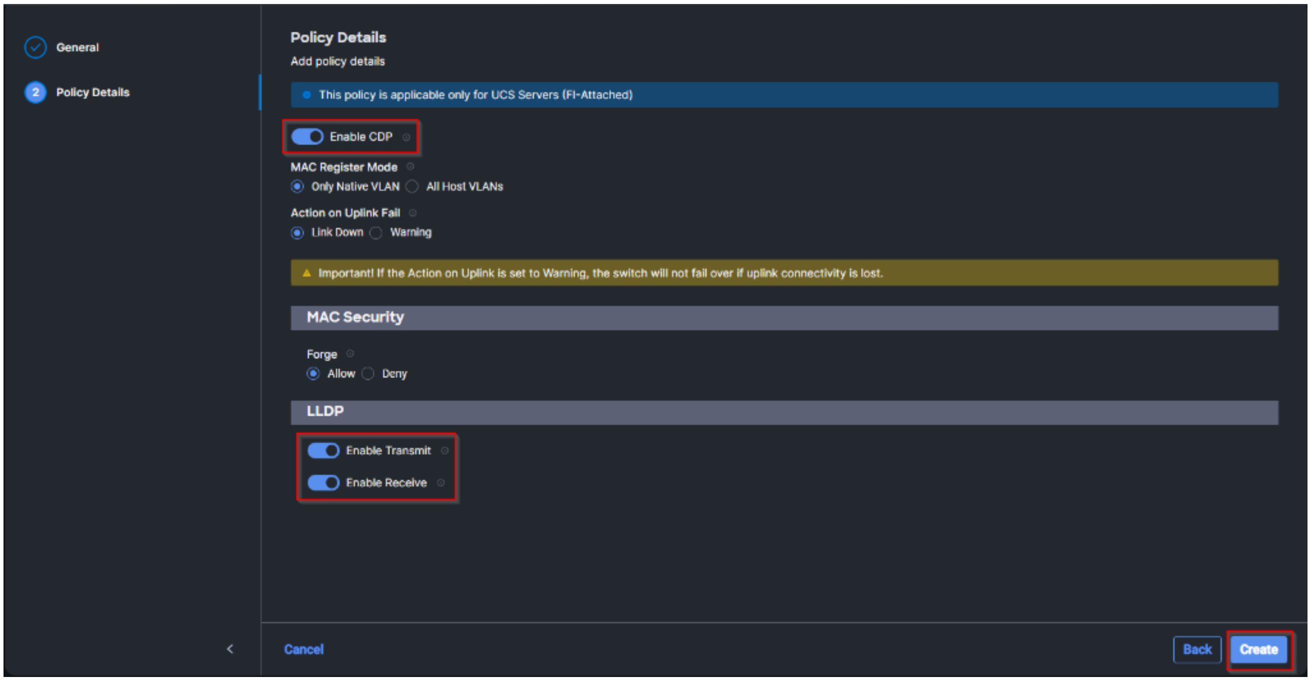

General

2 Policy Details

Policy Details

Add policy details

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

Cancel

Back

Create

Step 70. Click Select Policy under Ethernet QoS and click Create New.

Note: The 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.

Step 71. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-QOS).

Step 72. Click Next.

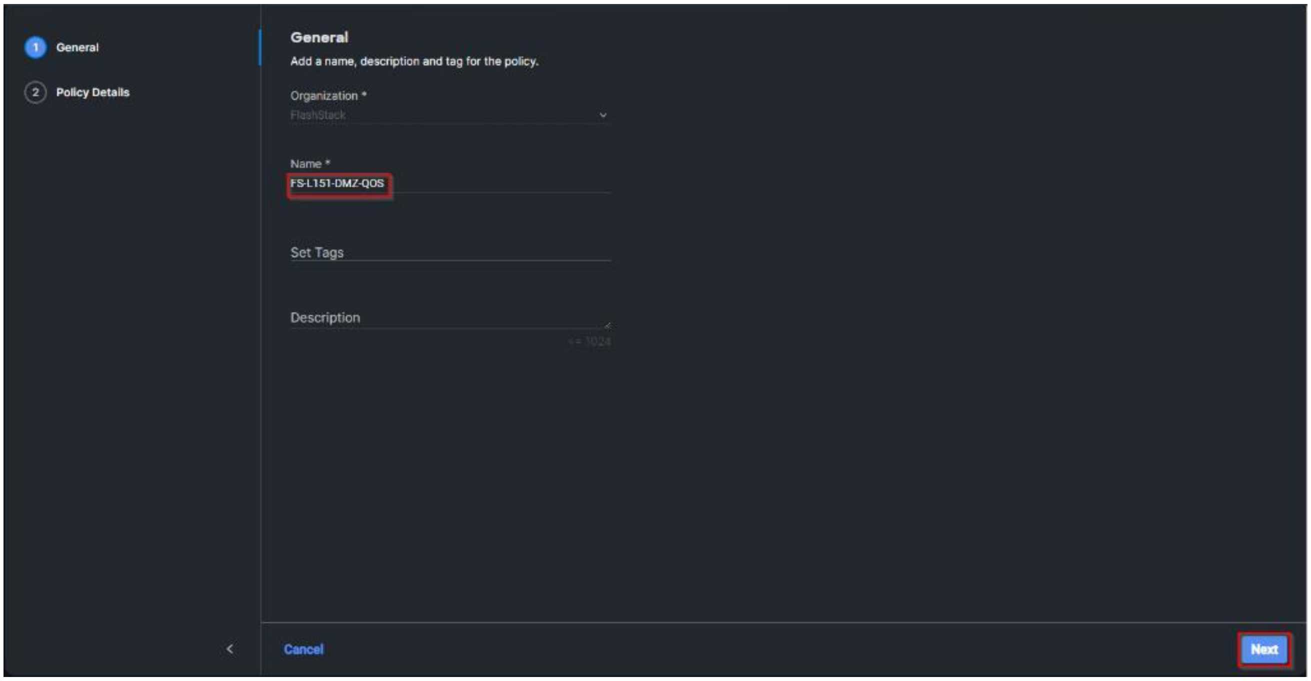

General

Add a name, description and tag for the policy.

Organization *: FlashStack

Name *: FS-L151-DMZ-QOS

Set Tags

Description

<= 1024

Step 73. Change the MTU Bytes value to 9000. Click Create.

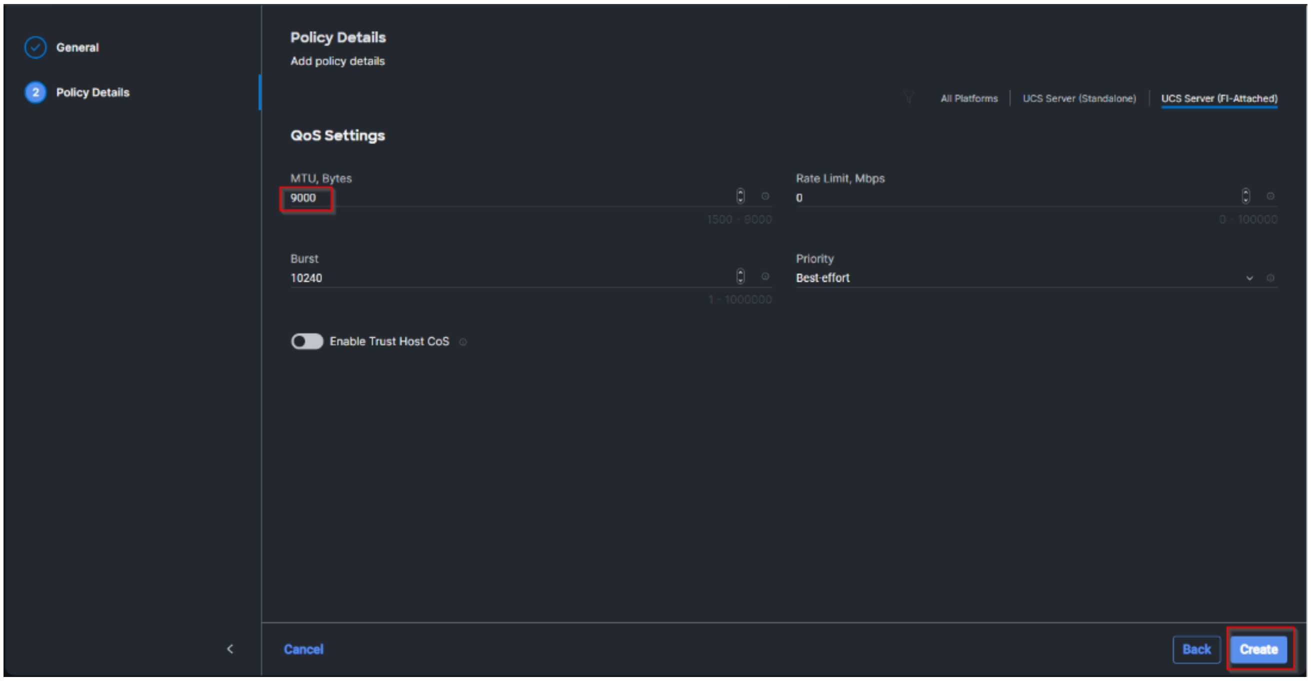

Policy Details

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

Step 74. Click Select Policy under Ethernet Adapter and then click Create New.

Note: The 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. Optionally, you can 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, FS-L151-DMZ-EthAdapt-VMware-HiTraffic, is created and attached to the VDS0-A and VDS0-B interfaces which handle vMotion.

Table 9. Ethernet Adapter Policy association to vNICs

Policy Name	vNICS
FS-L151-DMZ-EthAdapt-VMware	vSwitch0-A, vSwitch0-B
FS-L151-DMZ-EthAdapt-VMware-HiTraffic	VDS0-A, VDS0-B,

Step 75. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-EthAdapt-VMware).

Step 76. Click Select Default Configuration under Ethernet Adapter Default Configuration.

Step 77. From the list, select VMware. Click Next.

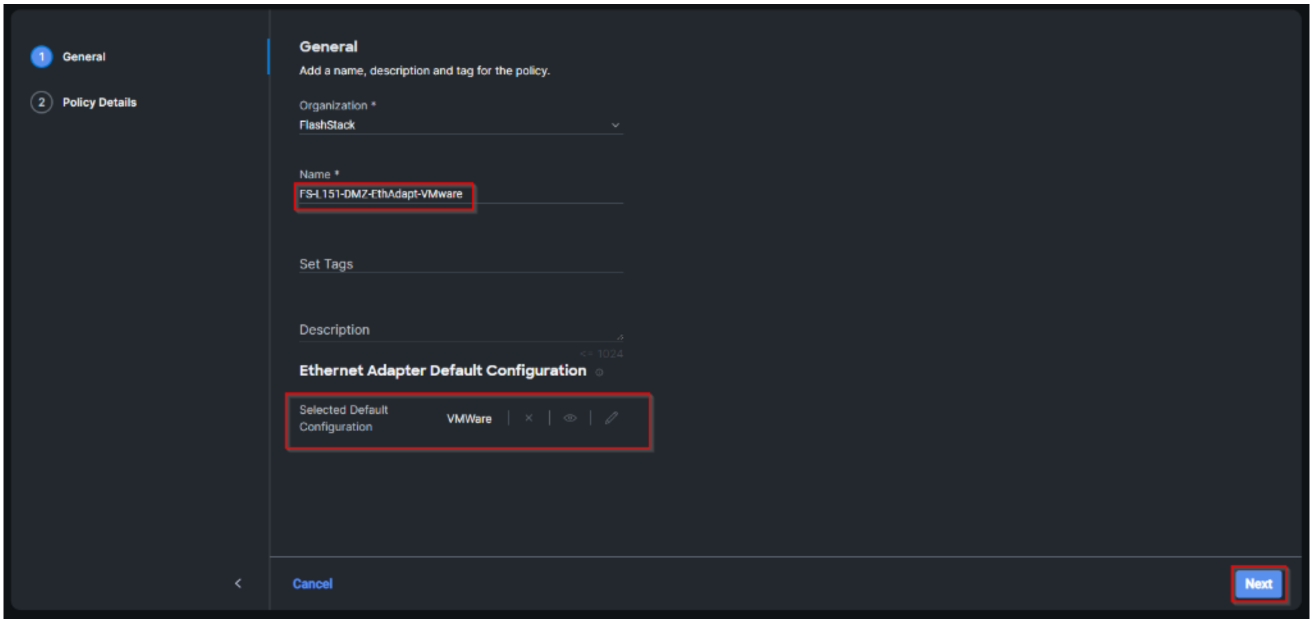

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L151-DMZ-EthAdapt-VMware

Set Tags

Description

<= 1024

Ethernet Adapter Default Configuration

Selected Default Configuration VMWare

Step 78. For the FS-L151-DMZ-EthAdapt-VMware policy, click Create and skip the rest of the steps in this section.

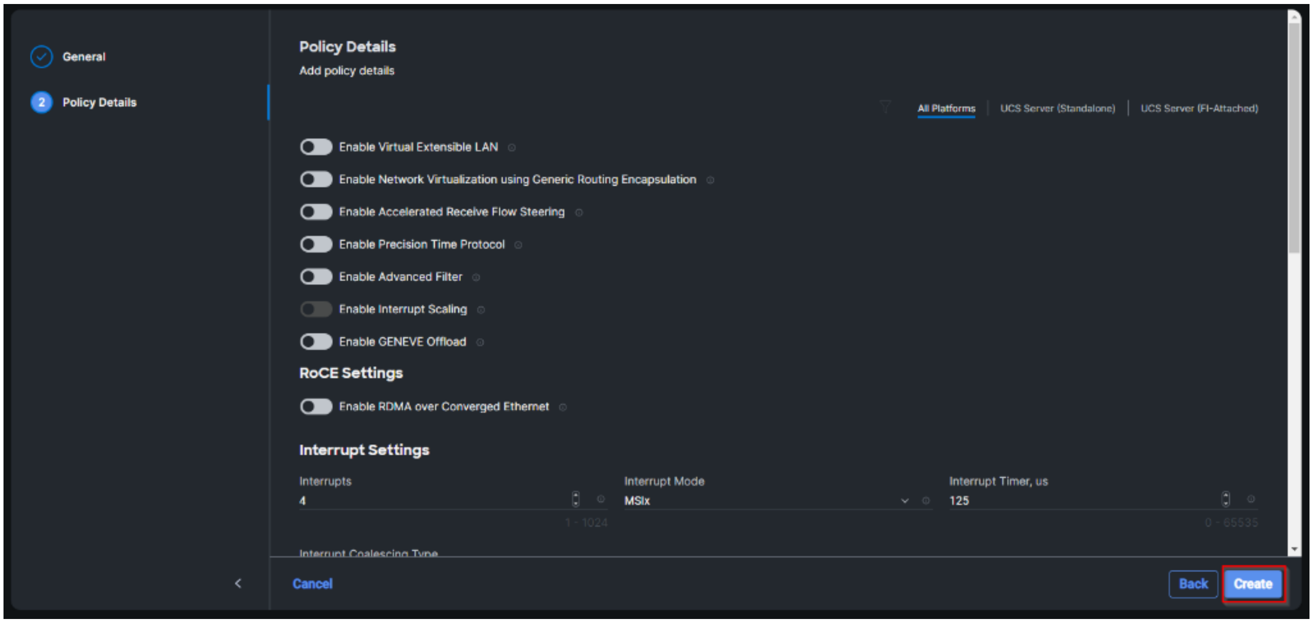

Policy Details

Add policy details

- Enable Virtual Extensible LAN
- Enable Network Virtualization using Generic Routing Encapsulation
- Enable Accelerated Receive Flow Steering
- Enable Precision Time Protocol
- Enable Advanced Filter
- Enable Interrupt Scaling
- Enable GENEVE Offload

RoCE Settings

- Enable RDMA over Converged Ethernet

Interrupt Settings

Interrupts	Interrupt Mode	Interrupt Timer, us
4 (1 - 1024)	MSIx	125 (0 - 65535)

Interrupt Coalescing Type

Step 79. For the optional FS-L151-DMZ-EthAdapt-VMware-HiTraffic policy used 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

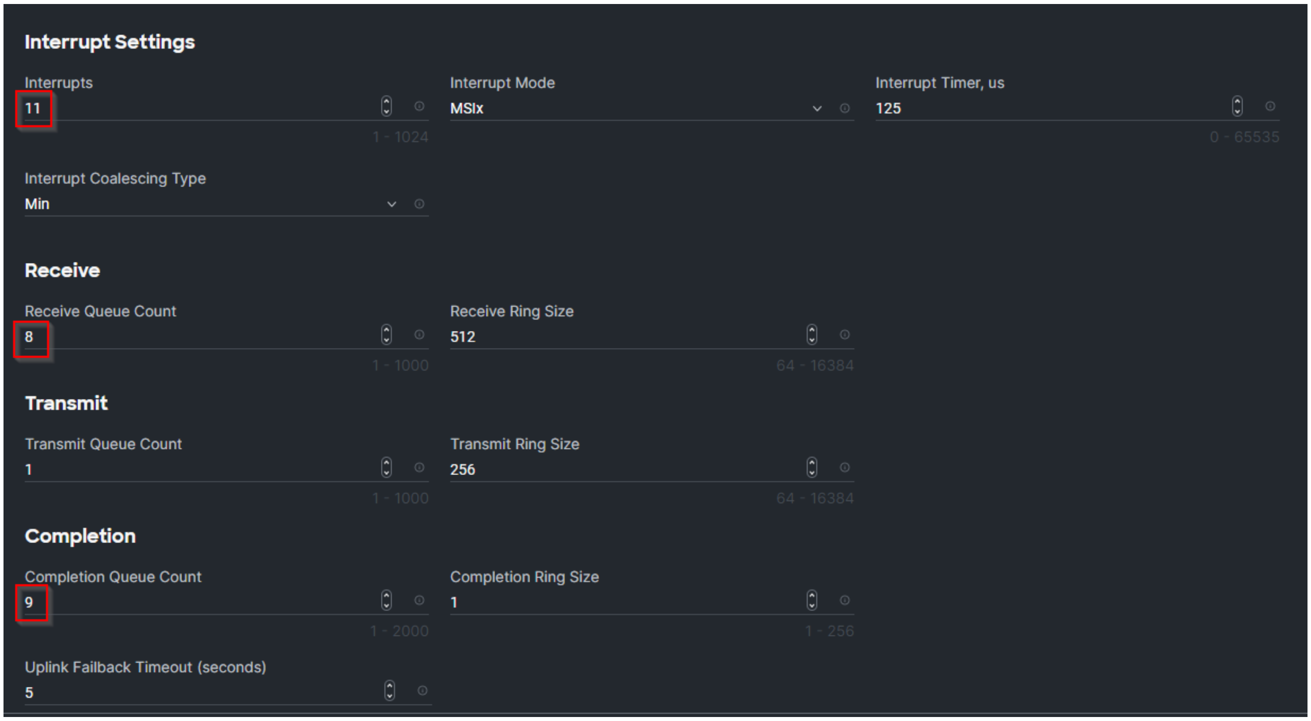

Interrupt Settings

Interrupts	Interrupt Mode	Interrupt Timer, us
11 (1 - 1024)	MSIx	125 (0 - 65535)

Interrupt Coalescing Type: Min

Receive

Receive Queue Count	Receive Ring Size
8 (1 - 1000)	512 (64 - 16384)

Transmit

Transmit Queue Count	Transmit Ring Size
1 (1 - 1000)	256 (64 - 16384)

Completion

Completion Queue Count	Completion Ring Size
9 (1 - 2000)	1 (1 - 256)

Uplink Failback Timeout (seconds): 5

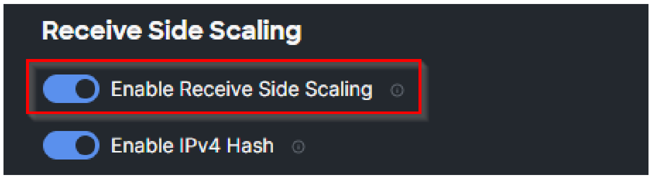

Receive Side Scaling

Enable Receive Side Scaling: on

Enable IPv4 Hash: on

Step 80. Click Create.

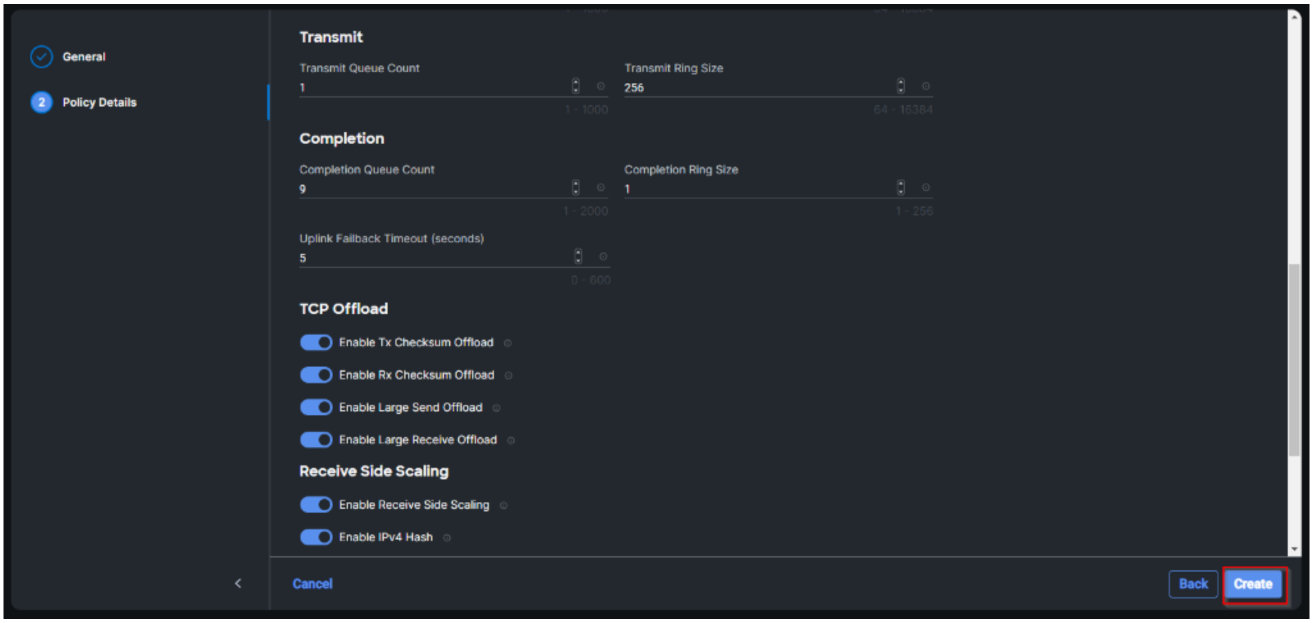

Transmit

Transmit Queue Count: 1 (1 - 1000)

Transmit Ring Size: 256 (64 - 16384)

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: on

Enable Rx Checksum Offload: on

Enable Large Send Offload: on

Enable Large Receive Offload: on

Receive Side Scaling

Enable Receive Side Scaling: on

Enable IPv4 Hash: on

Step 81. Click Create to finish creating the vNIC.

Step 82. Repeat the vNIC creation steps for the rest of vNICs. Verify all four vNICs were successfully created. Click Create.

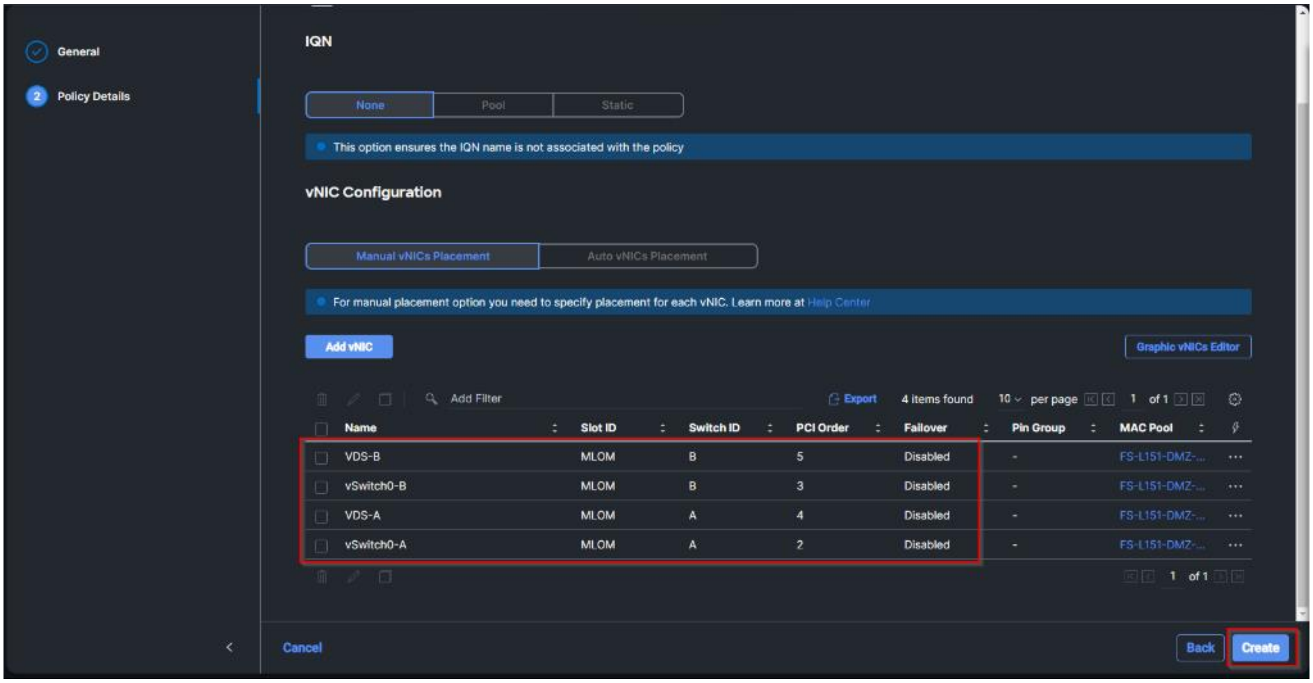

IQN

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

4 items found 10 per page 1 of 1

Name	Slot ID	Switch ID	PCI Order	Failover	Pin Group	MAC Pool
VDS-B	MLOM	B	5	Disabled	-	FS-L151-DMZ-...
vSwitch0-B	MLOM	B	3	Disabled	-	FS-L151-DMZ-...
VDS-A	MLOM	A	4	Disabled	-	FS-L151-DMZ-...
vSwitch0-A	MLOM	A	2	Disabled	-	FS-L151-DMZ-...

Step 83. Click Select Policy next to SAN Connectivity and then click Create New.

Note: 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 10. vHBA for boot from FC SAN

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-A	MLOM	A	0
vHBA-B	MLOM	B	1

Step 84. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-FC-SAN-Conn-Pol).

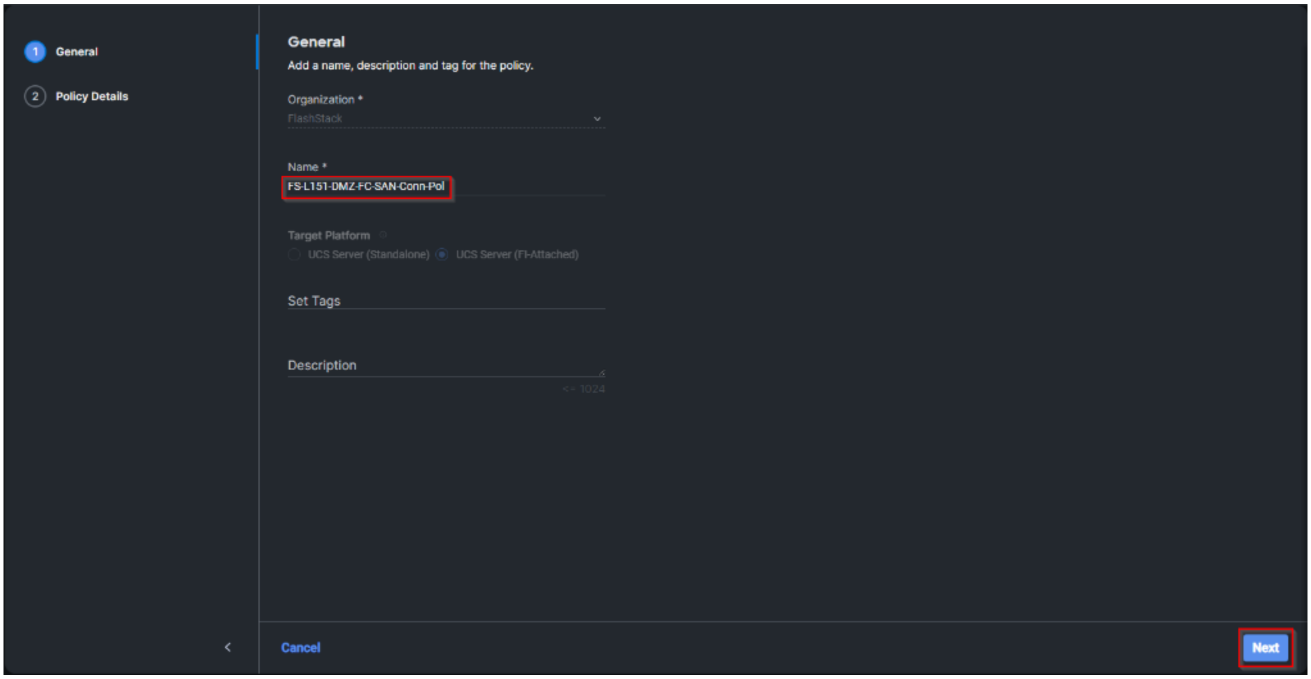
1 General
2 Policy Details
General
Add a name, description and tag for the policy.
Target Platform
○ UCS Server (Standalone) ● UCS Server (FI-Attached)
Set Tags
Description
<= 1024
Cancel
Next

Organization *: FlashStack

Name *: FS-L151-DMZ-FC-SAN-Conn-Pol

Step 85. Select Manual vHBAs Placement.

Step 86. Select Pool under WWNN.

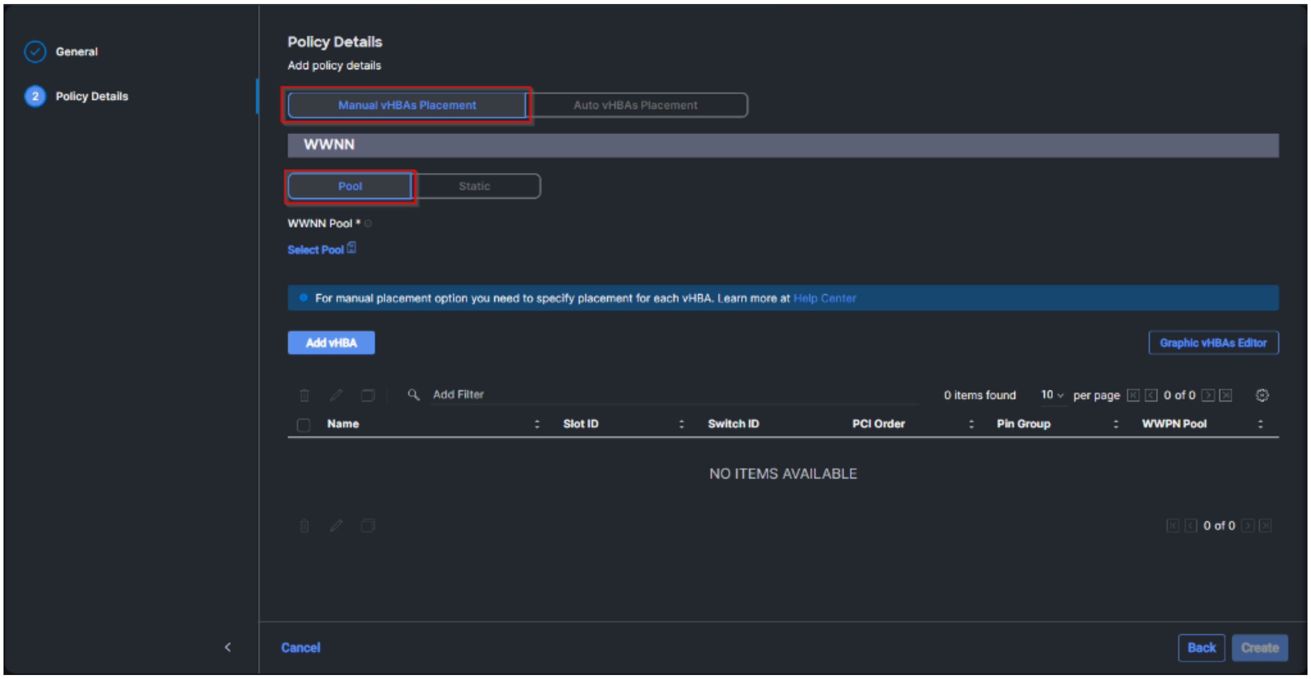

Policy Details

Add policy details

WWNN

WWNN Pool *

Select Pool

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

0 items found 10 per page 0 of 0

Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN Pool

NO ITEMS AVAILABLE

0 of 0

Note: The WWNN address pools have not been defined yet therefore a new WWNN address pool has to be defined.

Step 87. Click Select Pool under WWNN Pool and then click Create New.

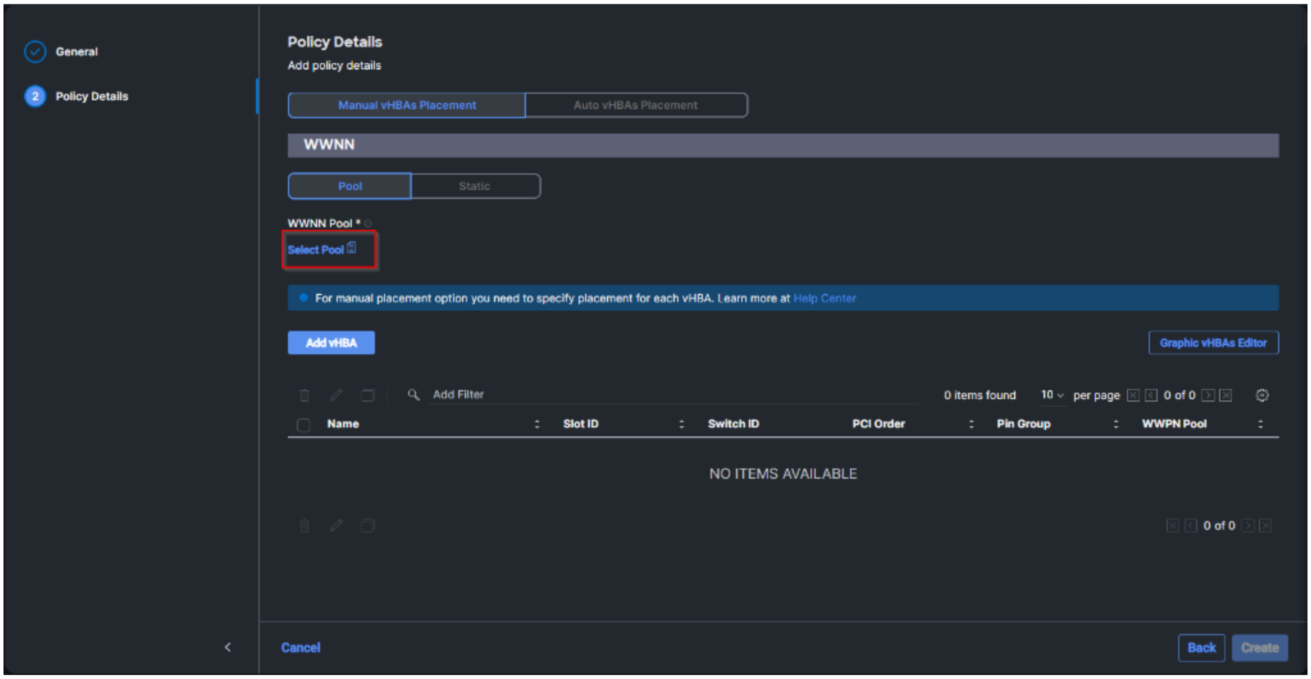

Policy Details

Add policy details

WWNN

WWNN Pool *

Select Pool

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

0 items found 10 per page 0 of 0

Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN Pool

NO ITEMS AVAILABLE

0 of 0

Step 88. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-WWN-Pool).

Step 89. Click Next.

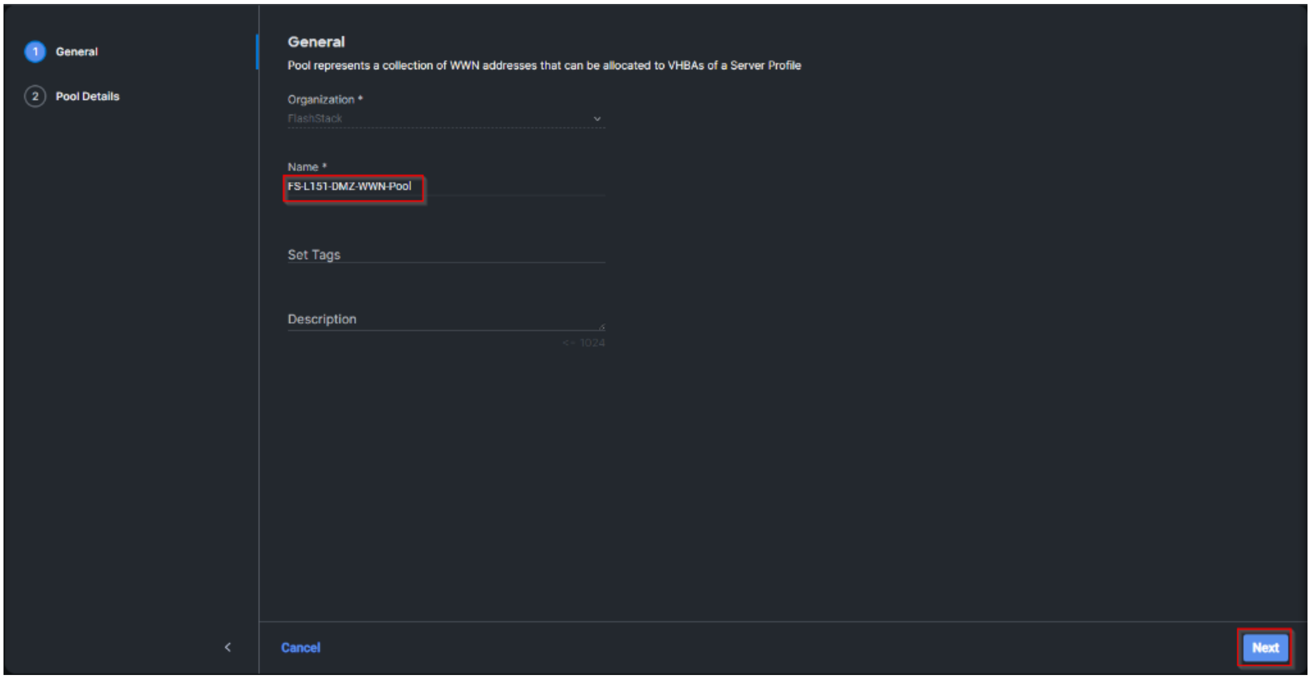

General

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

Organization *: FlashStack

Name *: FS-L151-DMZ-WWN-Pool

Set Tags

Description

<= 1024

Step 90. Provide the starting WWNN block address and the size of the pool. Click Create.

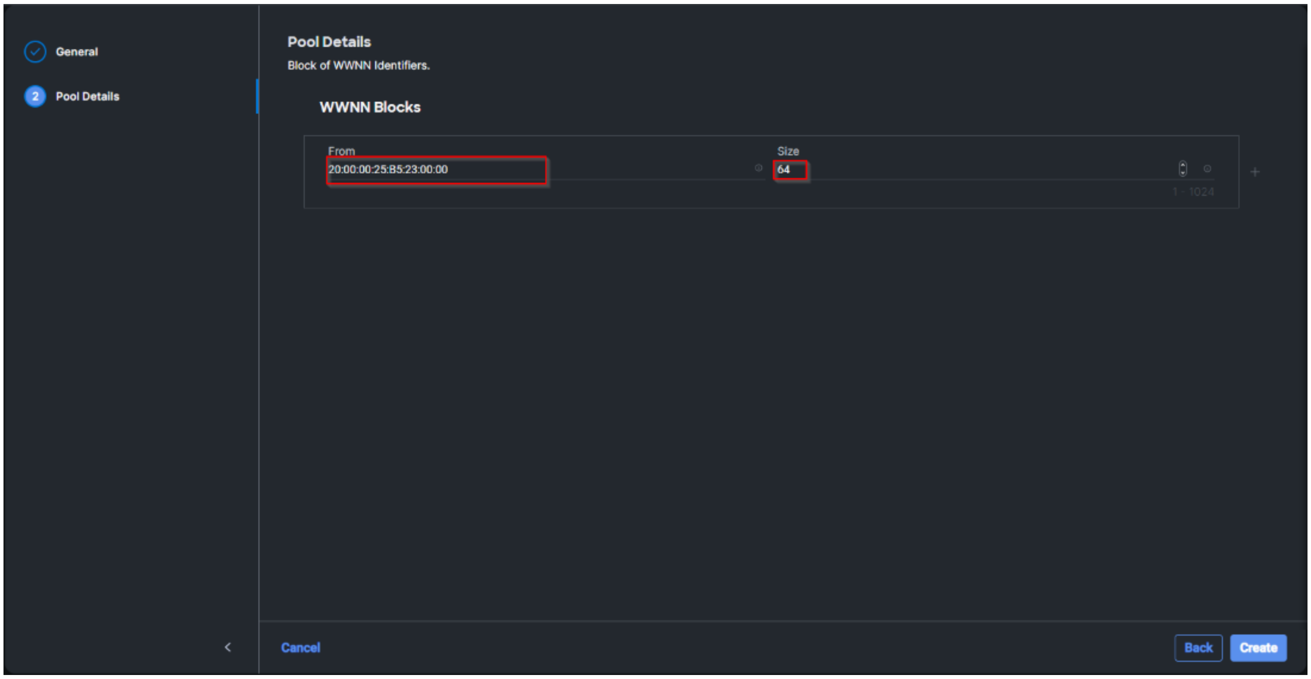

Pool Details

Block of WWNN Identifiers.

WWNN Blocks

From: 20:00:00:25:B5:23:00:00

Size: 64

1 - 1024

Note: As a best practice, additional information should always be coded into the WWNN address pool for troubleshooting. For example, in the address 20:00:00:25:B5:23:00:00, 23 is the rack ID.

Step 91. Click Add vHBA.

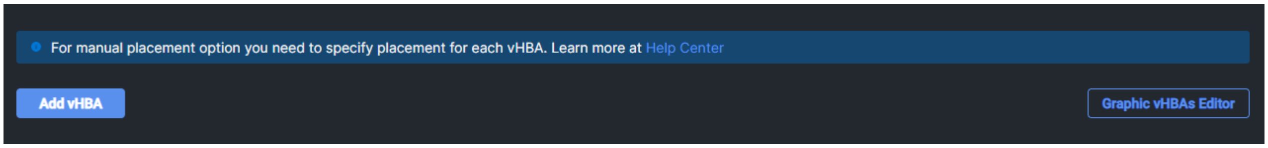
For manual placement option you need to specify placement for each vHBA. Learn more at Help Center

Step 92. Enter vHBA-A for the Name and select fc-initiator from the drop-down list.

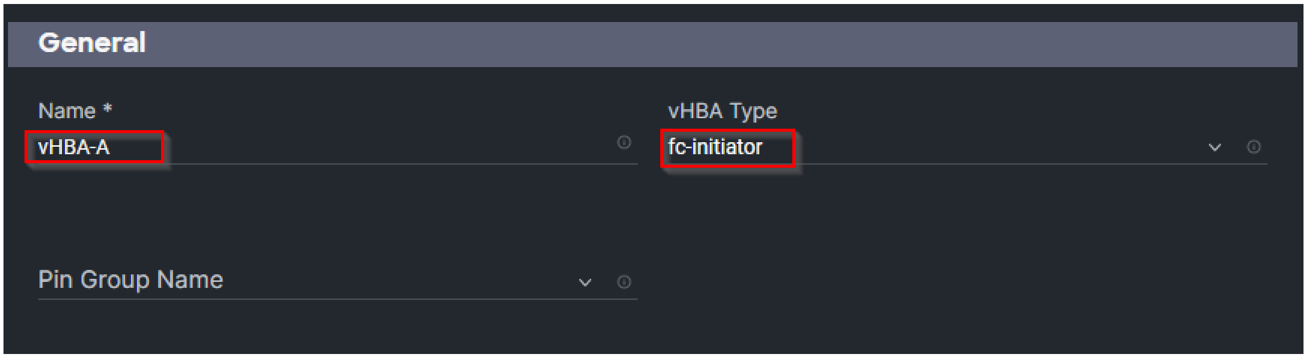

General

Name * vHBA-A

vHBA Type fc-initiator

Note: The WWPN address pool has not been defined yet therefore a WWPN address pool for Fabric A will be defined.

Step 93. Click Select Pool under WWPN Address Pool and then click Create New.

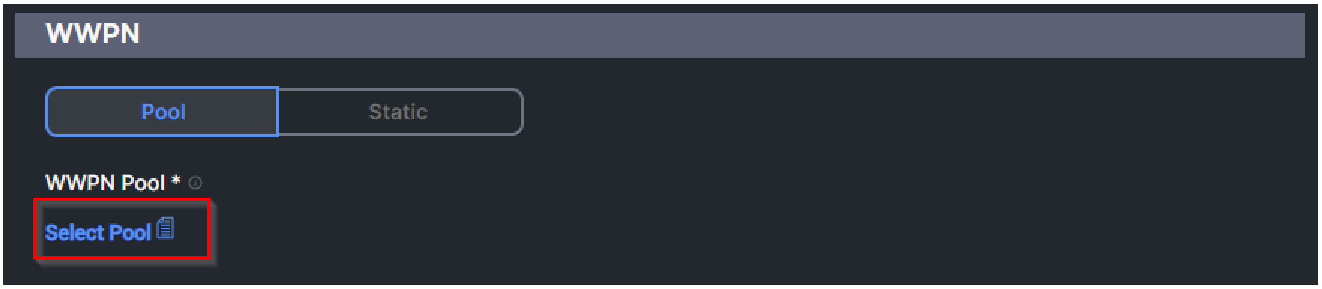

WWPN

Step 94. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-WWPN-Pool-A).

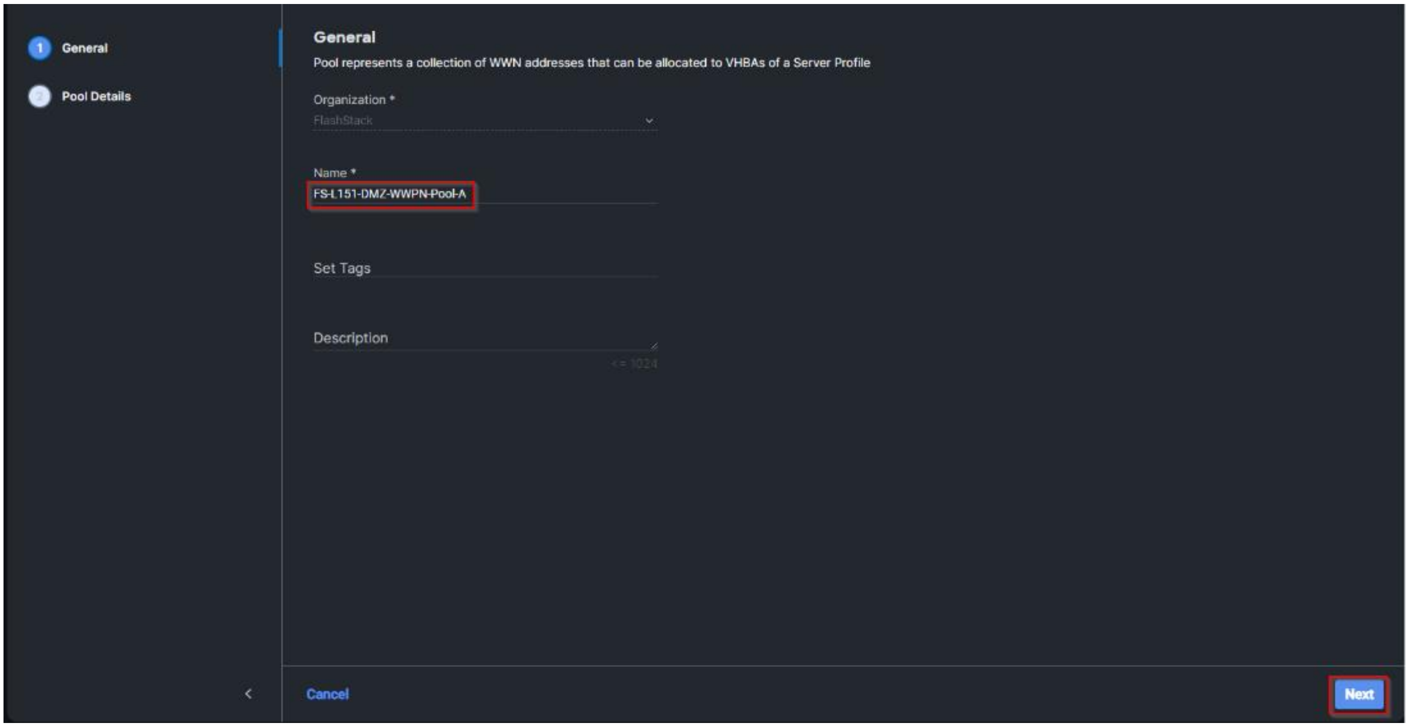

General

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

Organization * FlashStack

Name * FS-L151-DMZ-WWPN-Pool-A

Step 95. Provide the starting WWPN block address for SAN A and the size. Click Create.

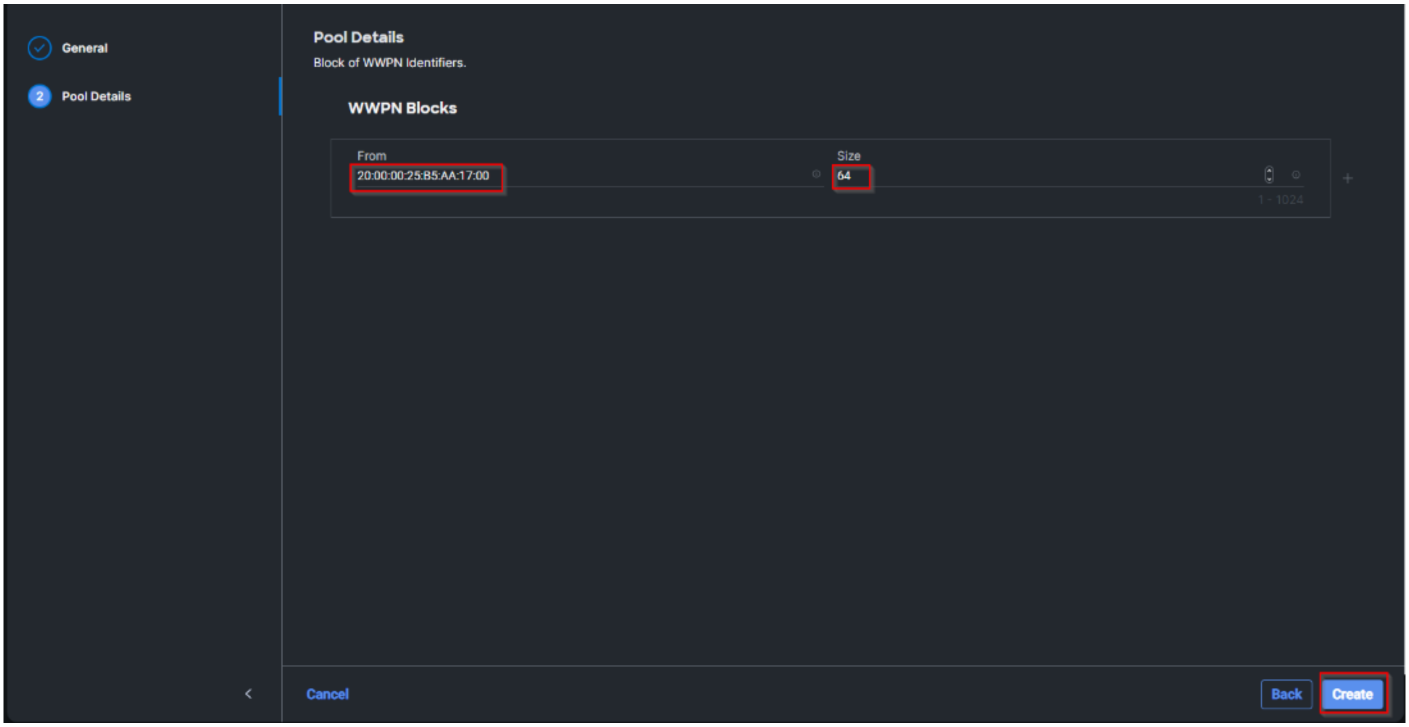

Pool Details

Block of WWPN Identifiers.

WWPN Blocks

From: 20:00:00:25:B5:AA:17:00

Size: 64 (1 - 1024)

Step 96. Provide the Switch ID (for example, A) and PCI Order (for example, 0) from Table 10.

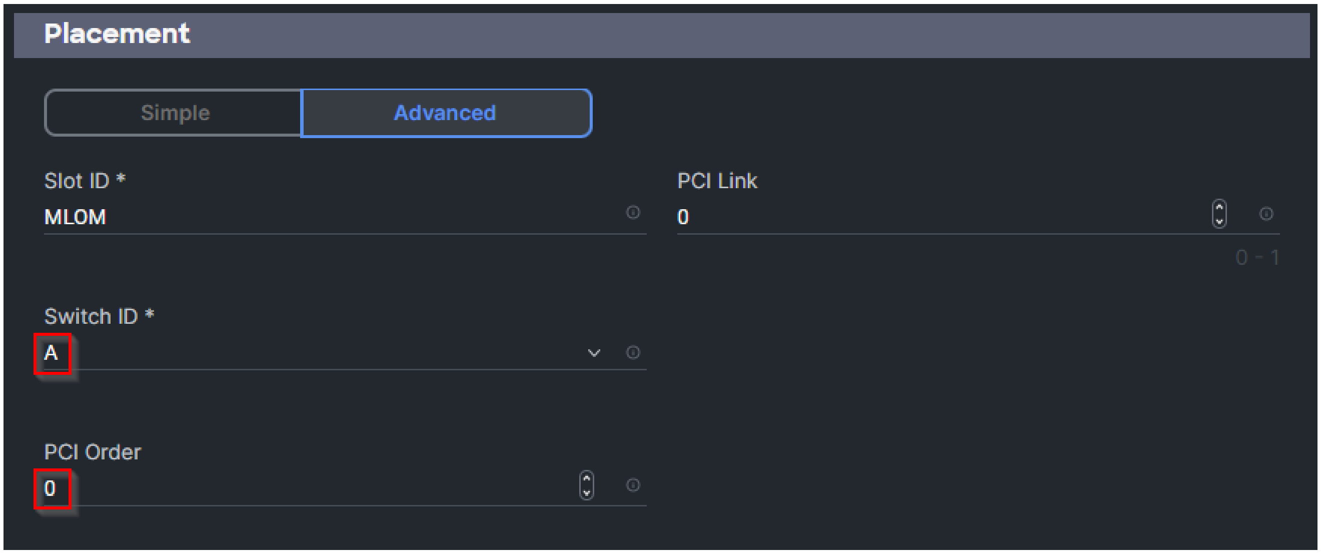

Placement

Slot ID *: MLOM

PCI Link: 0 (0 - 1)

Switch ID *: A

PCI Order: 0

Step 97. Click Select Policy under Fibre Channel Network and then click Create New.

Note: A Fibre Channel network policy governs the VSAN configuration for the virtual interfaces. In this deployment, VSAN 700 will be used for vHBA-A and VSAN 701 will be used for vHBA-B.

Step 98. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-K4-FCN-A). Click Next.

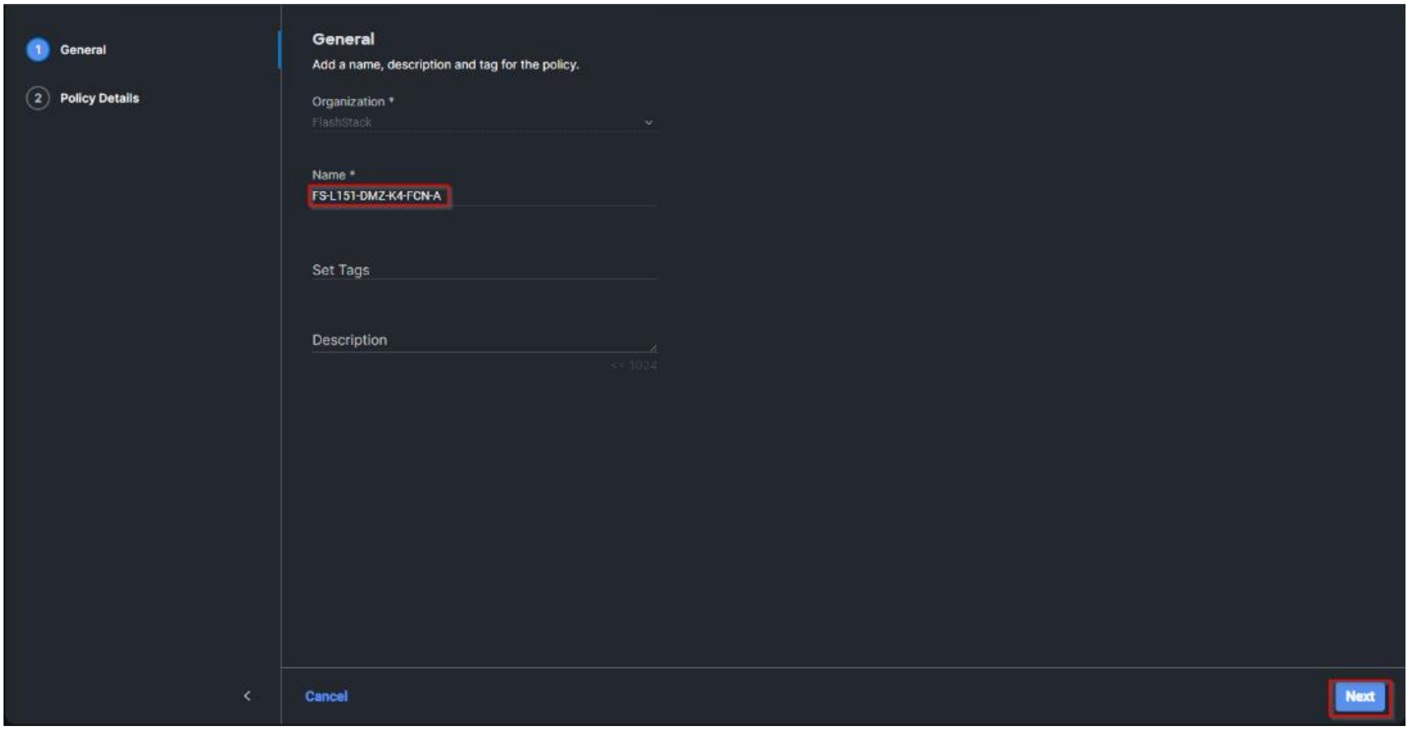

Organization *
FlashStack

Name *
FS-L151-DMZ-K4-FCN-A

Step 99. For the scope, select UCS Server (FI-Attached).

Step 100. Under VSAN ID, provide the VSAN information (for example, 700).

Step 101. Click Create.

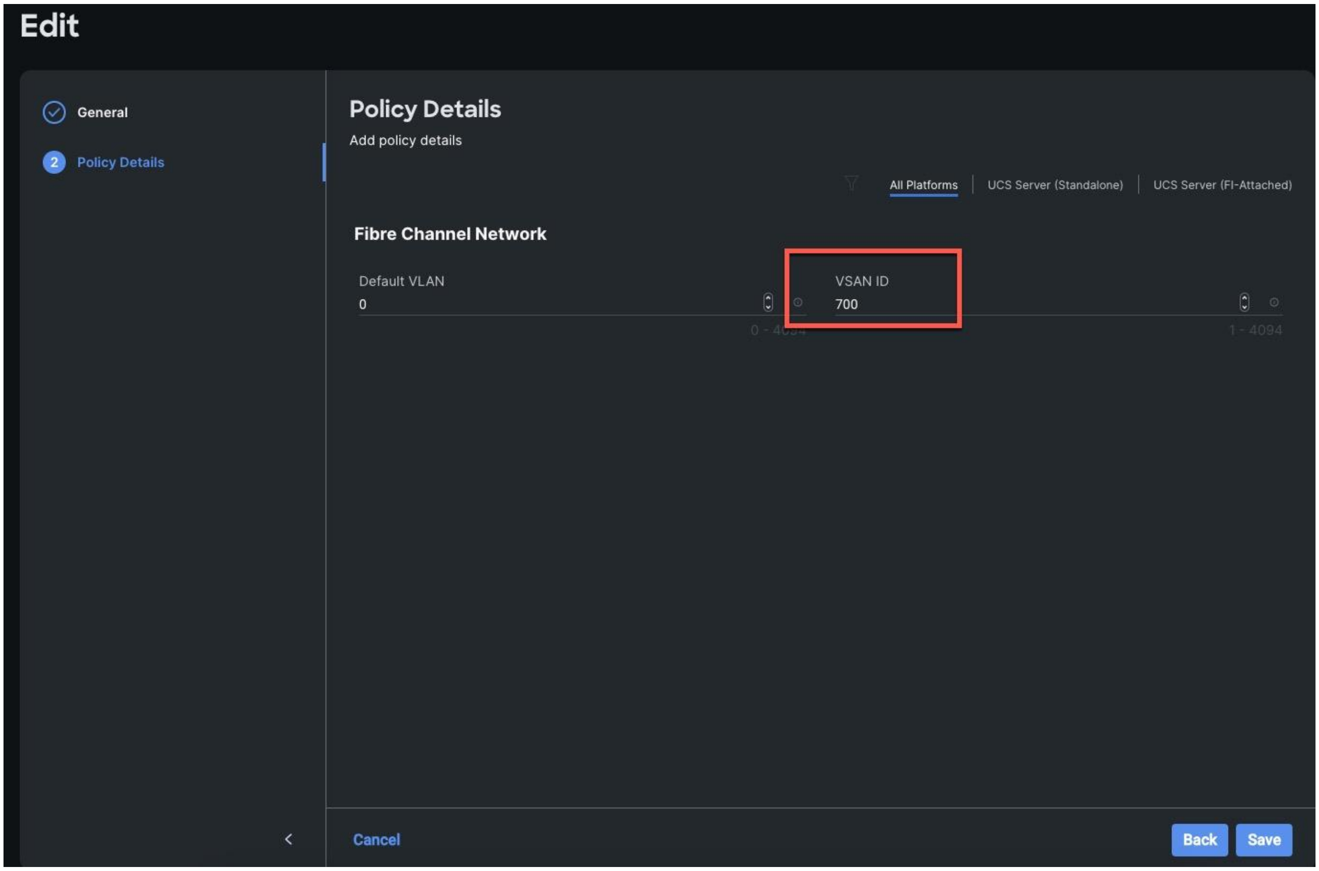

Default VLAN
0
0 - 4094

VSAN ID
700
1 - 4094

Step 102. Click Select Policy under Fibre Channel QoS and then click Create New.

Note: 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 uses default values and will be shared by all vHBAs.

Step 103. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-FCQOS-Pol). Click Next.

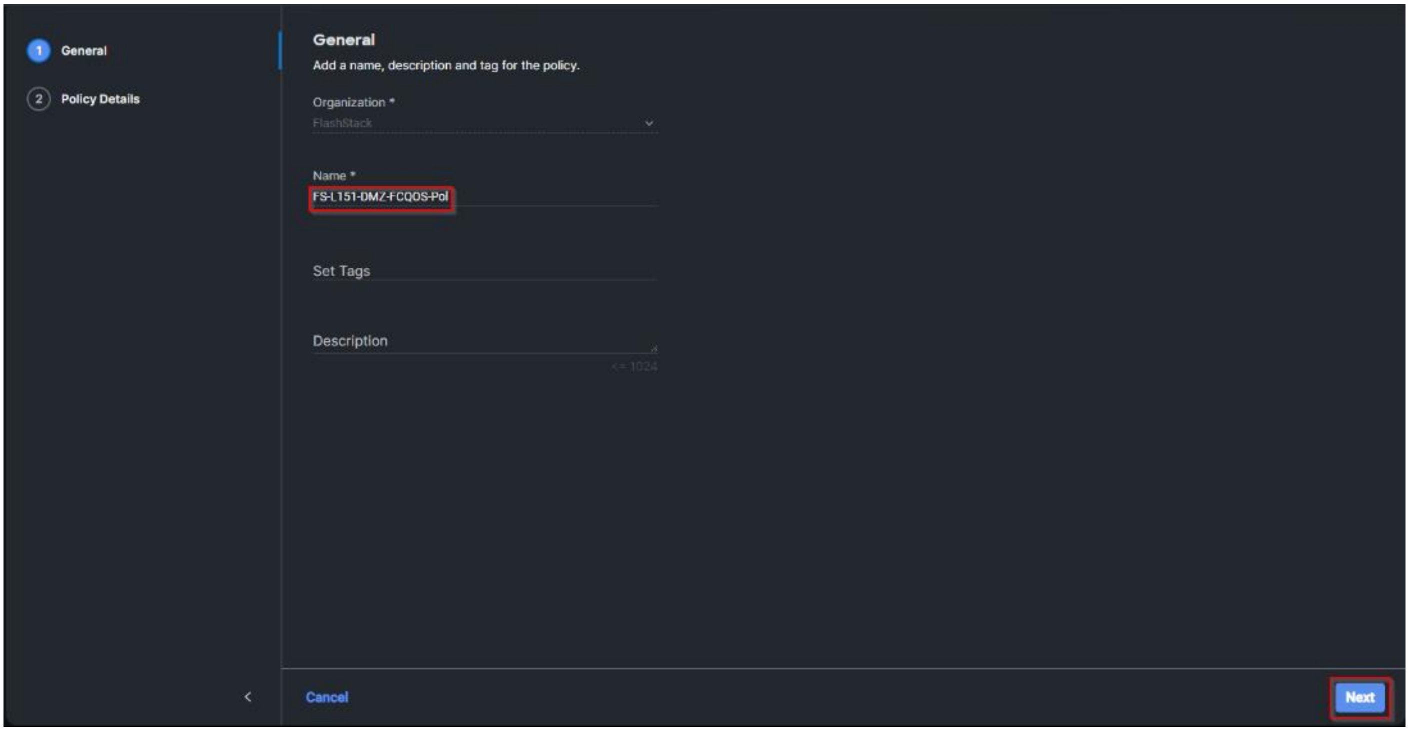

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L151-DMZ-FCQOS-Pol

Step 104. For the scope, select UCS Server (FI-Attached).

Note: Do not change the default values on the Policy Details screen.

Step 105. Click Create.

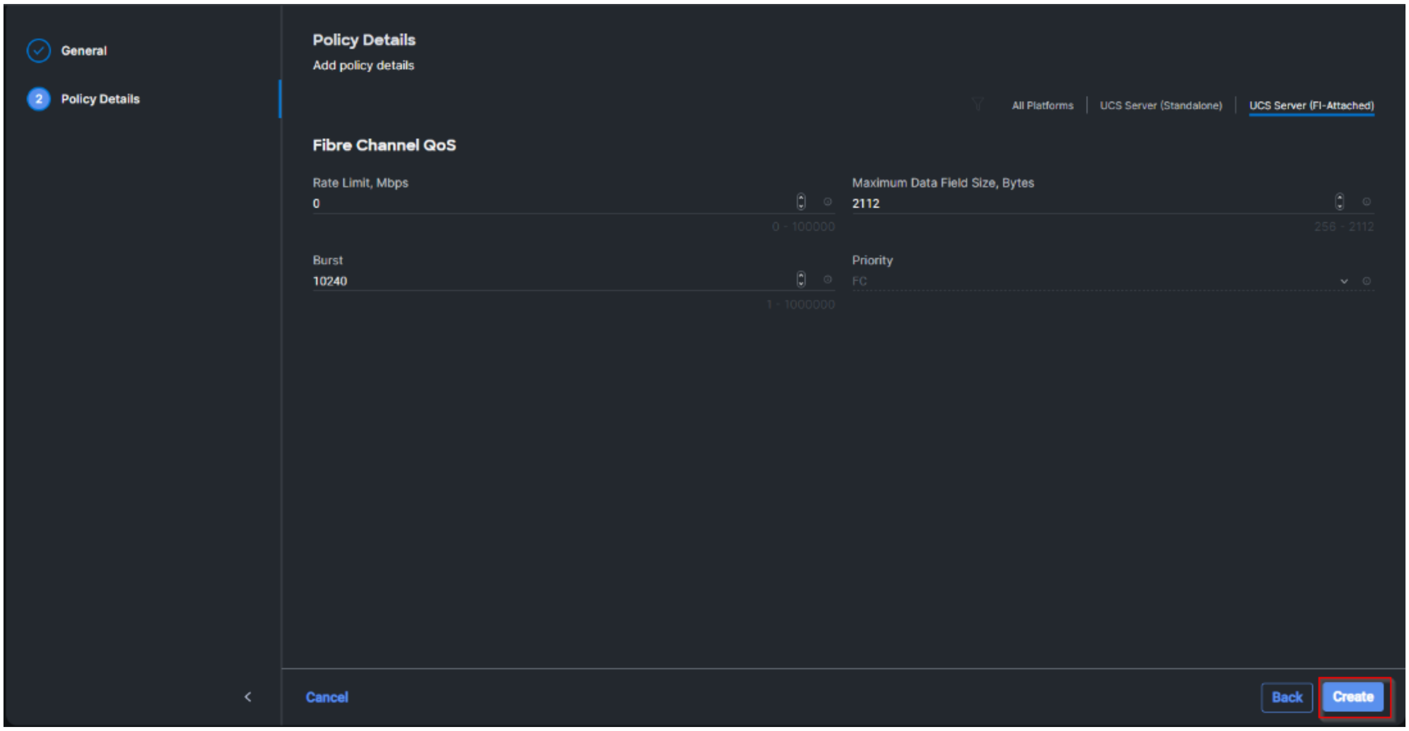

Fibre Channel QoS

Rate Limit, Mbps: 0 (0 - 100000)

Maximum Data Field Size, Bytes: 2112 (256 - 2112)

Burst: 10240 (1 - 1000000)

Priority: FC

Step 106. Click Select Policy under Fibre Channel Adapter and then click Create New.

Note: A Fibre Channel adapter policy governs the host-side behavior of the adapter, including the way that the adapter handles traffic. This validation uses the default values for the adapter policy, and the policy will be shared by all the vHBAs.

Step 107. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-FC-Adapter-Pol).

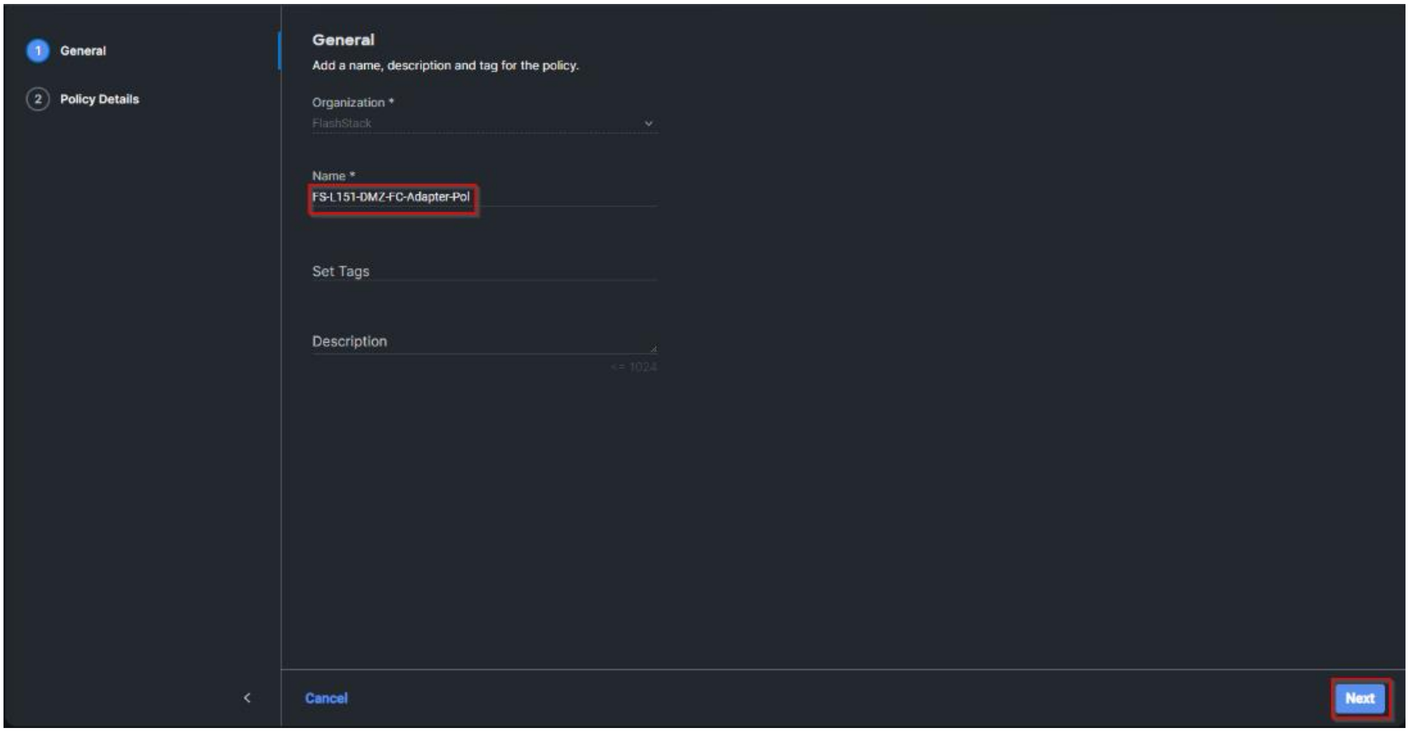

Organization *: FlashStack

Name *: FS-L151-DMZ-FC-Adapter-Pol

Set Tags

Description (<= 1024)

Step 108. For the scope, select UCS Server (FI-Attached).

Note: Do not change the default values on the Policy Details screen.

Step 109. Click Create.

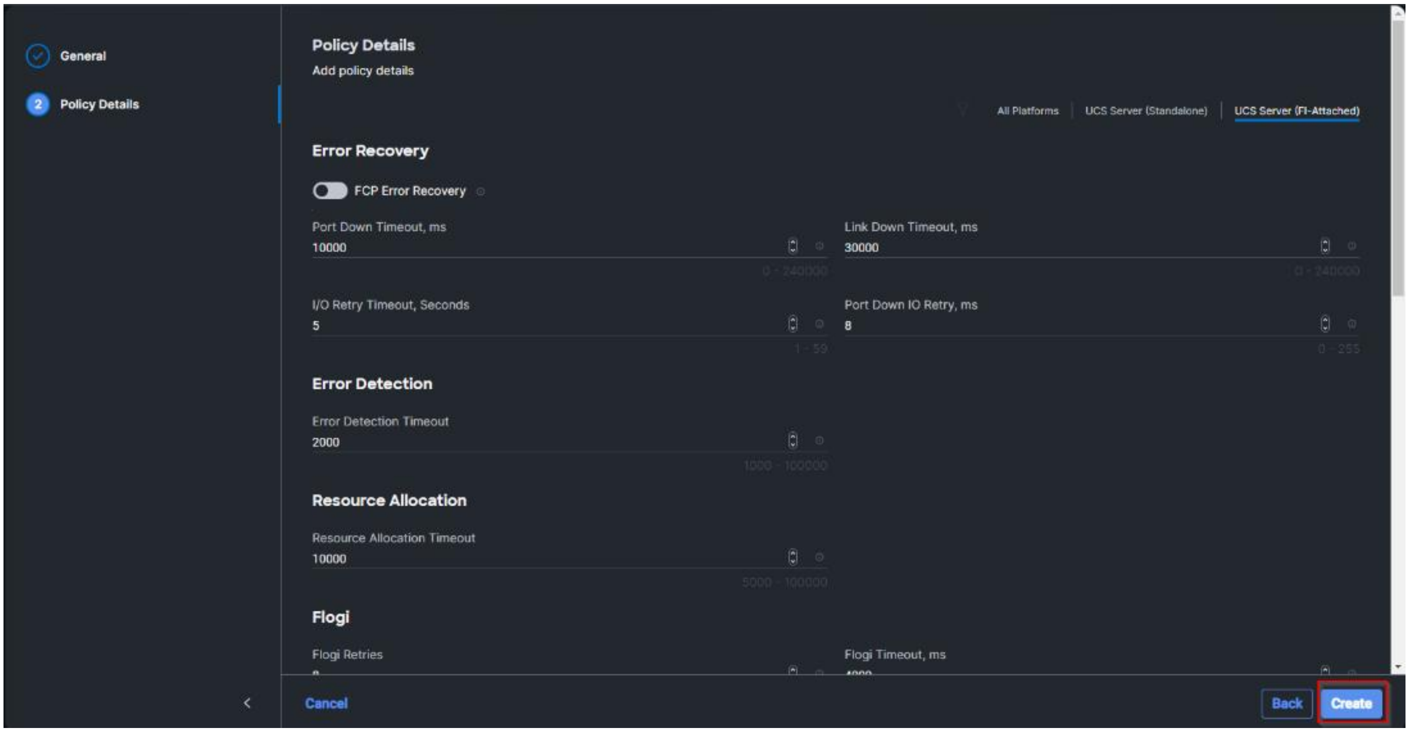

Policy Details
Add policy details

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

Flogi Timeout, ms

Step 110. Click Add to create vHBA-A.

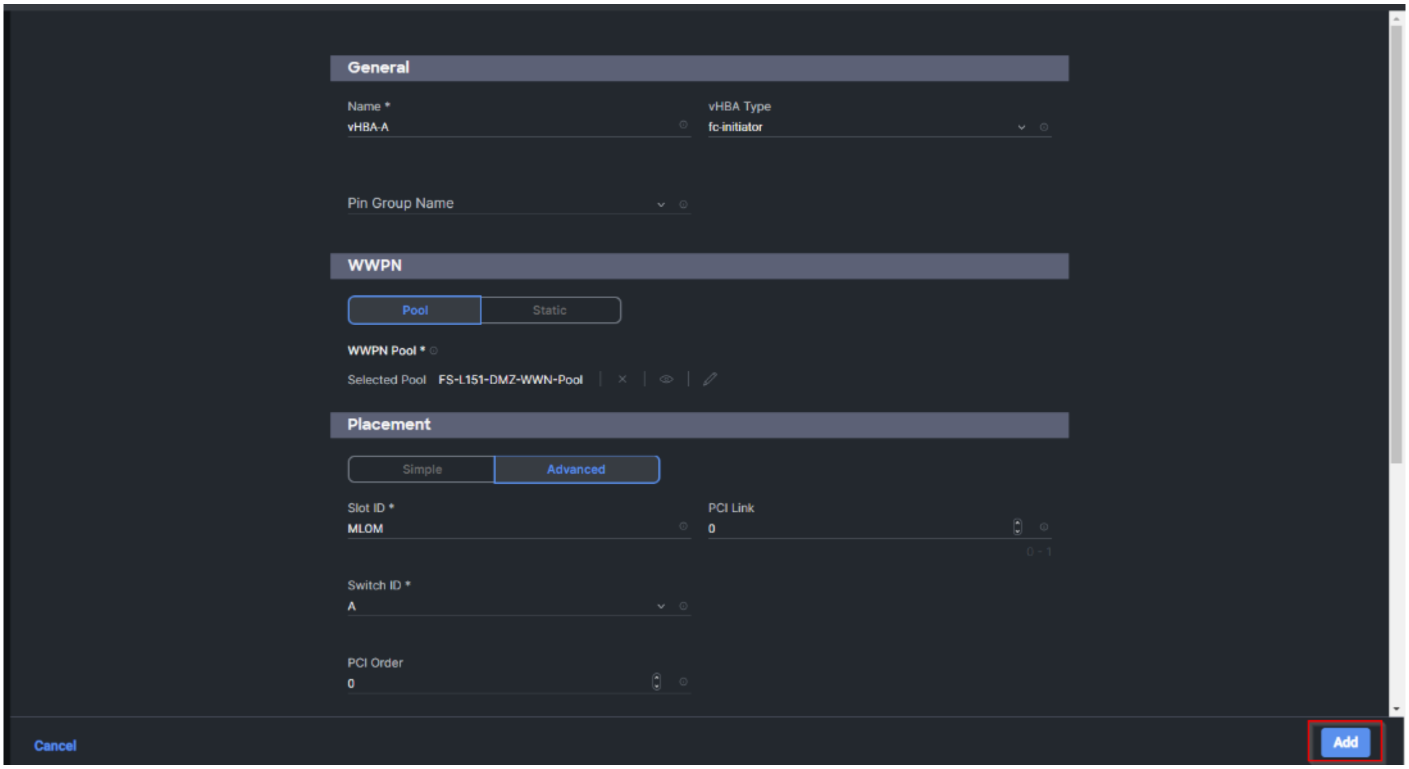

General

Name *: vHBA-A

vHBA Type: fc-initiator

Pin Group Name

WWPN

WWPN Pool *

Selected Pool FS-L151-DMZ-WWN-Pool

Placement

Slot ID *: MLOM

PCI Link: 0 (0 - 1)

Switch ID *: A

PCI Order: 0

Step 111. Create the vHBA-B using the same steps using pools and Fibre Channel Network policy for SAN-B.

Step 112. Verify both vHBAs are added to the SAN connectivity policy.

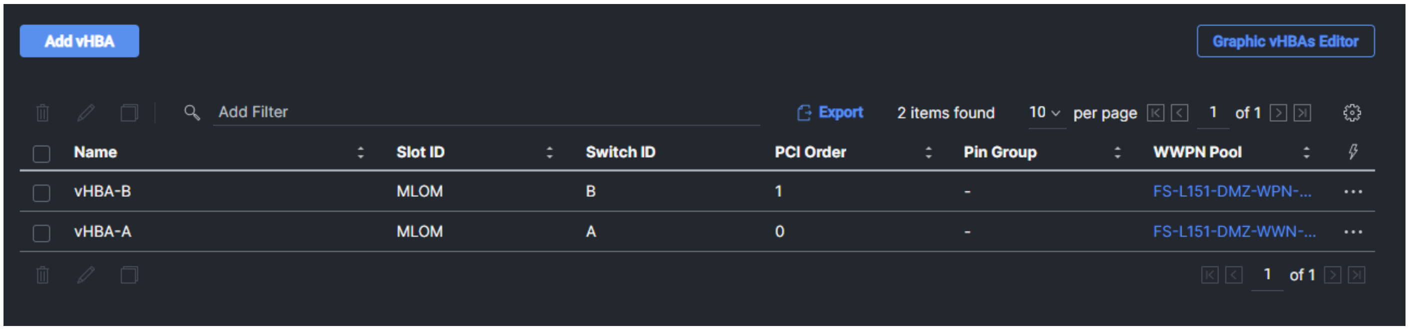

Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN Pool
vHBA-B	MLOM	B	1	-	FS-L151-DMZ-WPN-...
vHBA-A	MLOM	A	0	-	FS-L151-DMZ-WWN-...

Step 113. When the LAN connectivity policy and SAN connectivity policy are created and assigned, click Next to move to the Summary screen.

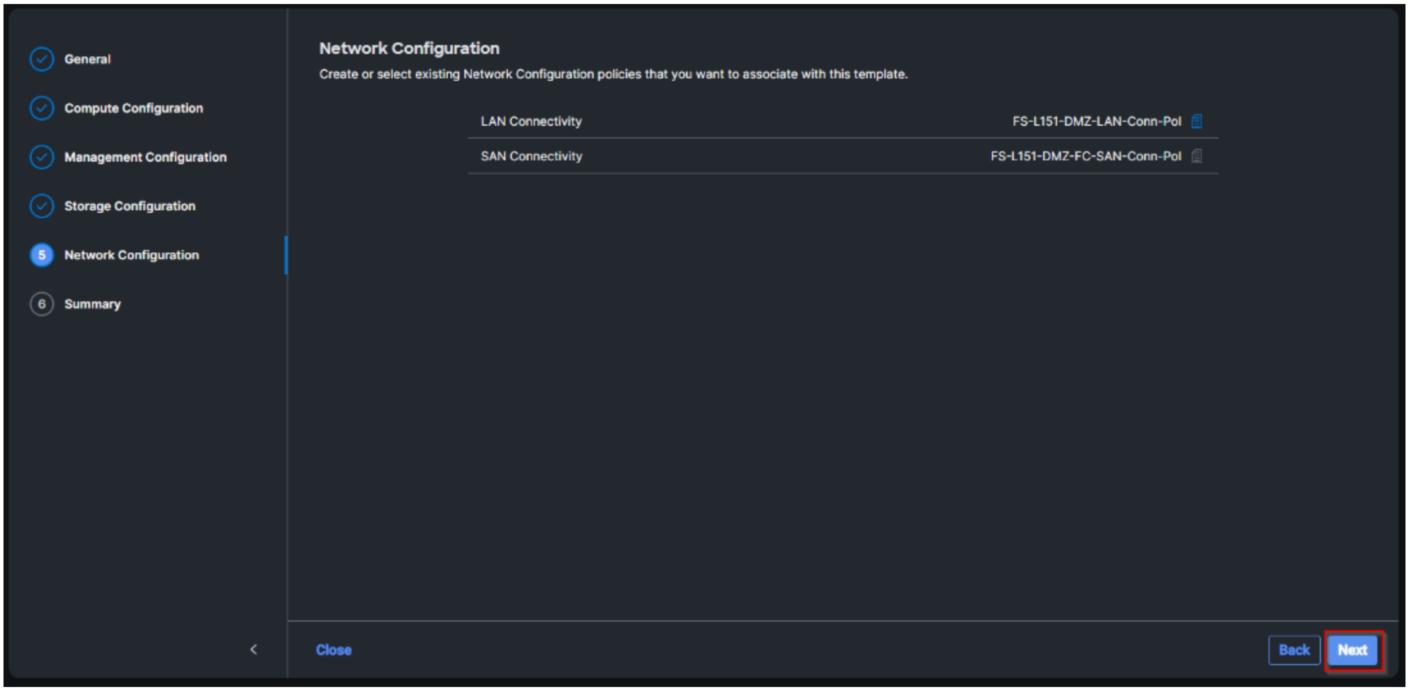

Network Configuration

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

LAN Connectivity	FS-L151-DMZ-LAN-Conn-Pol
SAN Connectivity	FS-L151-DMZ-FC-SAN-Conn-Pol

Step 114. 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.

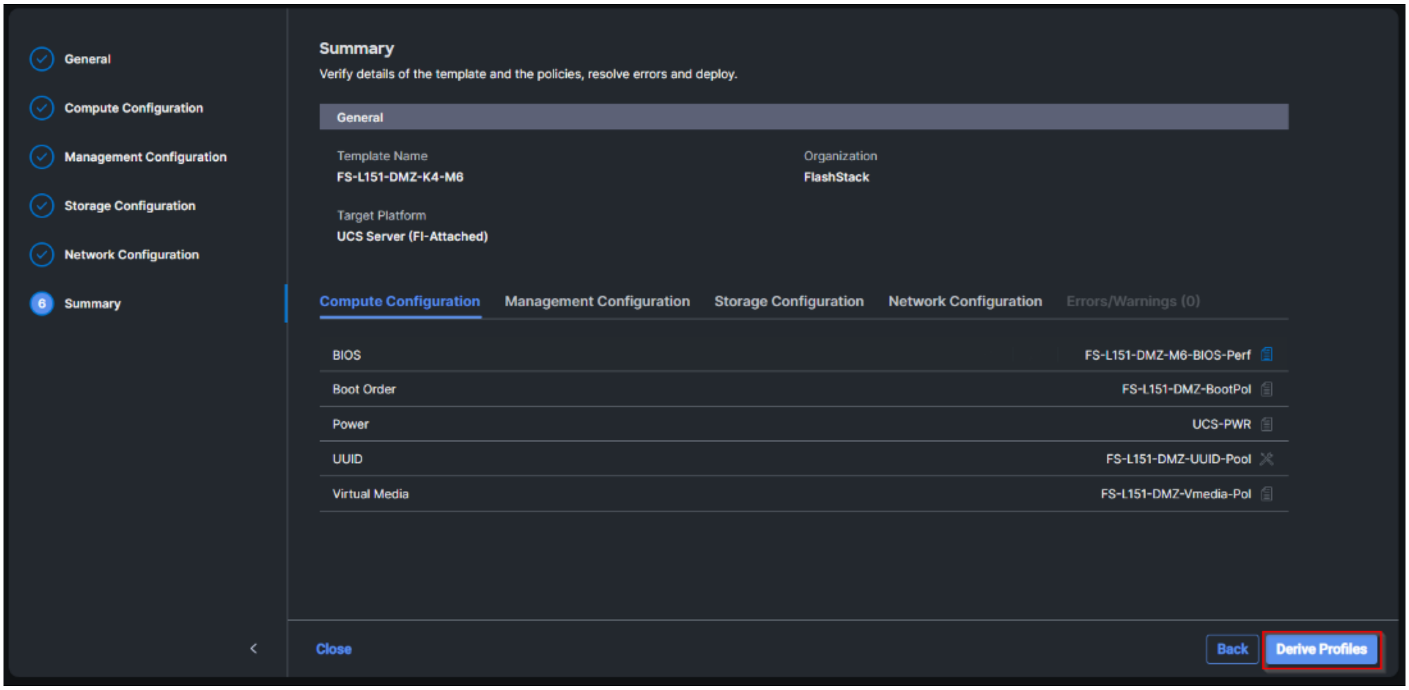

Summary

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

General

Template Name
FS-L151-DMZ-K4-M6

Organization
FlashStack

Target Platform
UCS Server (FI-Attached)

BIOS	FS-L151-DMZ-M6-BIOS-Perf
Boot Order	FS-L151-DMZ-BootPol
Power	UCS-PWR
UUID	FS-L151-DMZ-UUID-Pool
Virtual Media	FS-L151-DMZ-Vmedia-Pol

Step 115. Under the Server Assignment, select Assign Now and select Cisco UCS X210c M7 Nodes. You can select one or more servers depending on the number of profiles to be deployed. Click Next.

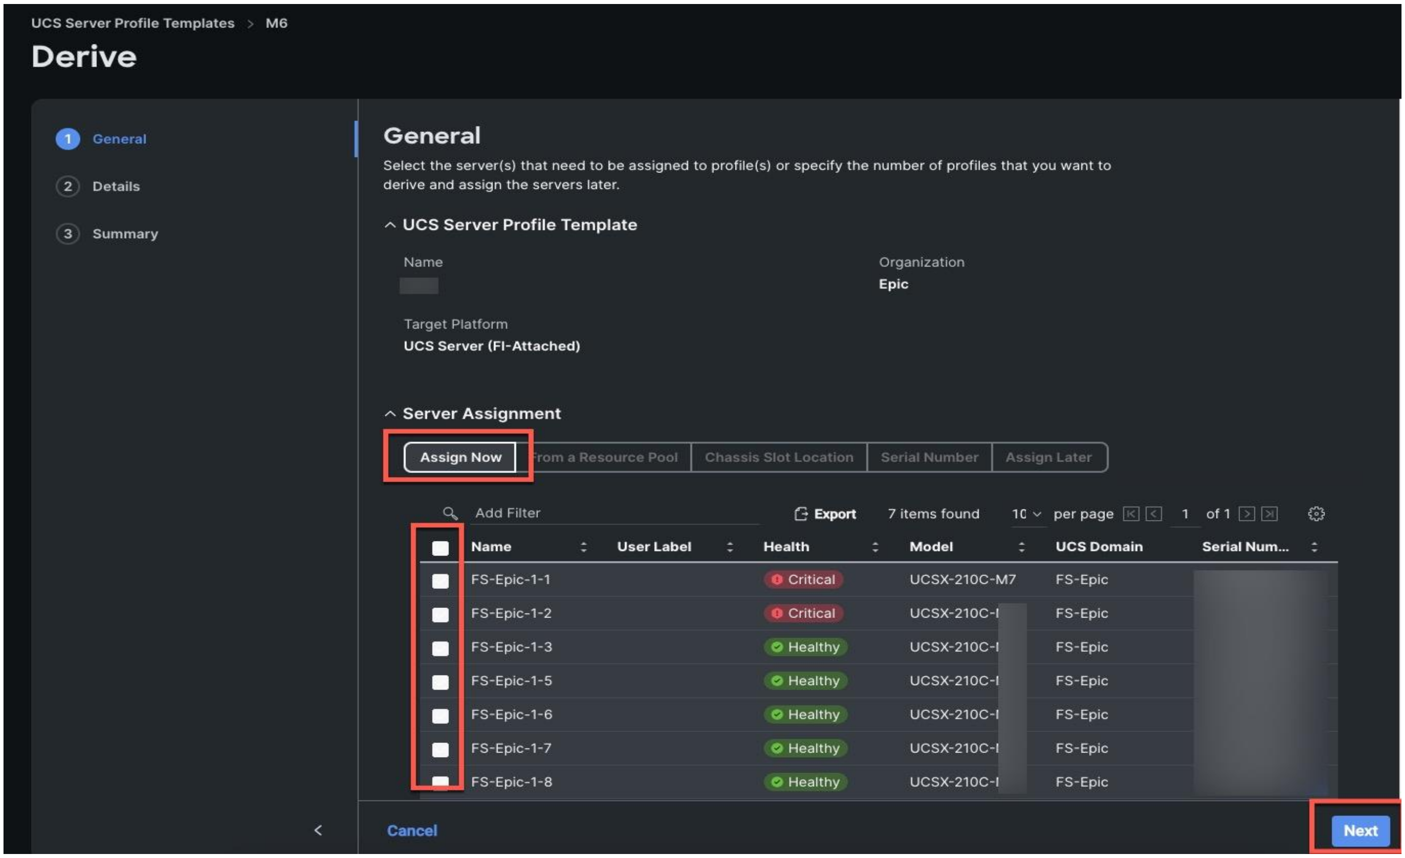

Derive

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
[redacted]

Organization
Epic

Target Platform
UCS Server (FI-Attached)

Server Assignment

7 items found 10 per page 1 of 1

	Name	User Label	Health	Model	UCS Domain	Serial Num...
□	FS-Epic-1-1		Critical	UCSX-210C-M7	FS-Epic	[redacted]
□	FS-Epic-1-2		Critical	UCSX-210C-[redacted]	FS-Epic	[redacted]
□	FS-Epic-1-3		Healthy	UCSX-210C-[redacted]	FS-Epic	[redacted]
□	FS-Epic-1-5		Healthy	UCSX-210C-[redacted]	FS-Epic	[redacted]
□	FS-Epic-1-6		Healthy	UCSX-210C-[redacted]	FS-Epic	[redacted]
□	FS-Epic-1-7		Healthy	UCSX-210C-[redacted]	FS-Epic	[redacted]
□	FS-Epic-1-8		Healthy	UCSX-210C-[redacted]	FS-Epic	[redacted]

Cisco Intersight will fill the default information for the number of servers selected.

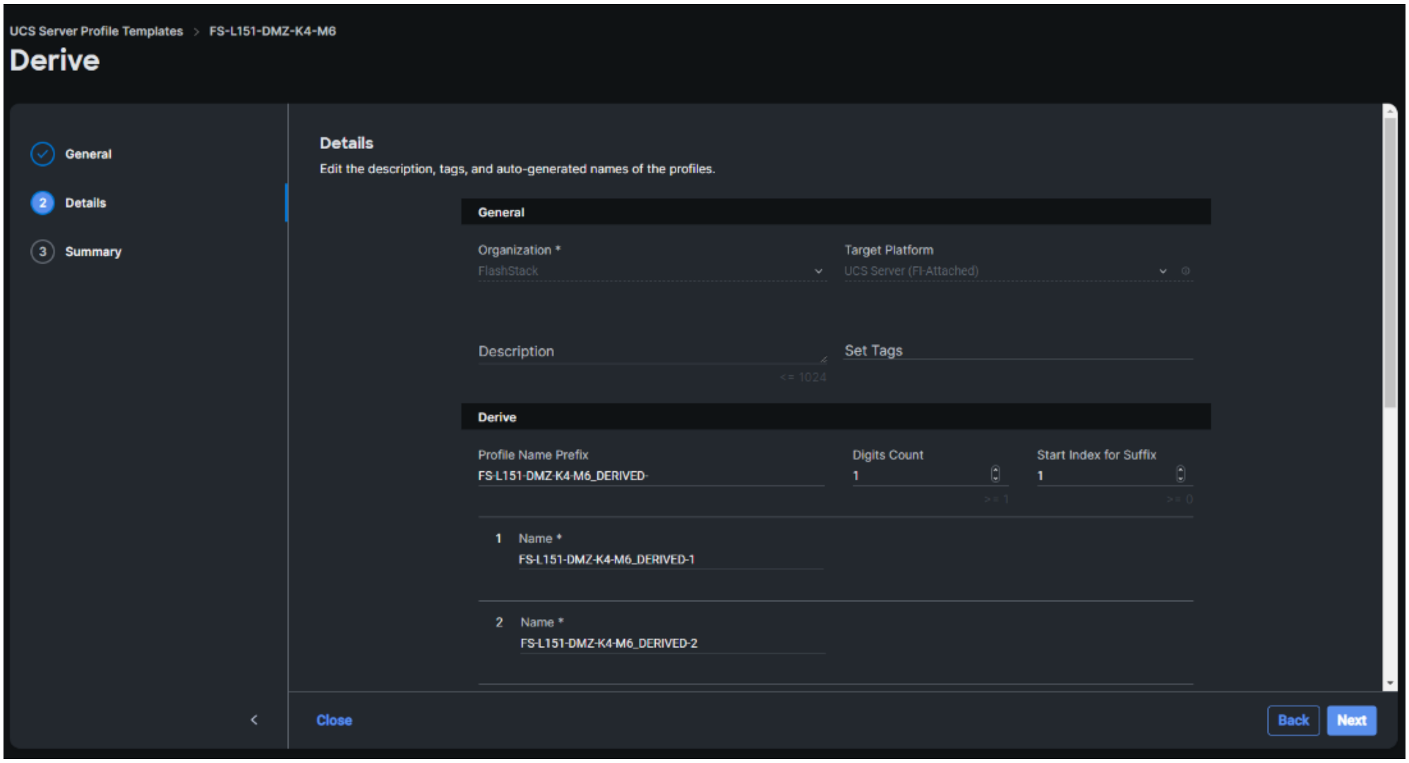

UCS Server Profile Templates > FS-L151-DMZ-K4-M6

Derive

Details

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

General

Organization * FlashStack

Target Platform UCS Server (FI-Attached)

Description <= 1024

Set Tags

Derive

Profile Name Prefix FS-L151-DMZ-K4-M6_DERIVED-

Digits Count 1 >= 1

Start Index for Suffix 1 >= 0

1 Name * FS-L151-DMZ-K4-M6_DERIVED-1

2 Name * FS-L151-DMZ-K4-M6_DERIVED-2

Step 116. Adjust the Prefix and number as needed. Click Next.

Step 117. Verify the information and click Derive to create the Server Profiles.

Configure Cisco Nexus 93180YC-FX Switches

This section details the steps for the Cisco Nexus 93180YC-FX switch configuration.

Procedure 1. Configure Global Settings for Cisco Nexus A and Cisco Nexus B

Step 1. Log in as admin user into the Cisco Nexus Switch A and run the following commands to set global configurations and jumbo frames in QoS:

```
conf terminal
policy-map type network-qos jumbo
class type network-qos class-default
mtu 9216
exit
class type network-qos class-fcoe
pause no-drop
mtu 2158
exit
exit
system qos
service-policy type network-qos jumbo
exit
copy running-config startup-config
```

Step 2. Log in as admin user into the Cisco Nexus Switch B and run the same commands (above) to set global configurations and jumbo frames in QoS.

Procedure 2. Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

Note: For this solution, we created VLAN 40, 41, 42, 43 and 46.

Step 1. Log in as admin user into the Cisco Nexus Switch A.

Step 2. Create VLAN 40:

```
config terminal
VLAN 40
name InBand-Mgmt
no shutdown
exit
copy running-config startup-config
```

Step 3. Log in as admin user into the Nexus Switch B and create VLANs.

Virtual Port Channel (vPC) Summary for Data and Storage Network

In the Cisco Nexus 93180YC-FX switch topology, a single vPC feature is enabled to provide HA, faster convergence in the event of a failure, and greater throughput. Cisco Nexus 93180YC-FX vPC configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers is listed in Table 11.

Table 11. vPC Summary

vPC Domain	vPC Name	vPC ID
50	Peer-Link	1
50	vPC Port-Channel to FI-A	11
50	vPC Port-Channel to FI-B	12

As listed in Table 11, a single vPC domain with Domain ID 50 is created across two Cisco Nexus 93180YC-FX member switches to define vPC members to carry specific VLAN network traffic. In this topology, a total number of 3 vPCs were defined:

- vPC ID 1 is defined as Peer link communication between two Nexus switches in Fabric A and B.
- vPC IDs 11 and 12 are defined for traffic from Cisco UCS fabric interconnects.

Cisco Nexus 93180YC-FX Switch Cabling Details

The following tables list the cabling information.

Table 12. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch A	Eth1/51	40Gbe	Cisco UCS fabric interconnect B	Eth1/49
	Eth1/52	40Gbe	Cisco UCS fabric interconnect A	Eth1/49
	Eth1/1	25Gbe	Cisco Nexus 93180YC-FX B	Eth1/1
	Eth1/2	25Gbe	Cisco Nexus 93180YC-FX B	Eth1/2
	Eth1/3	25Gbe	Cisco Nexus 93180YC-FX B	Eth1/3
	Eth1/4	25Gbe	Cisco Nexus 93180YC-FX B	Eth1/4
	MGMT0	1Gbe	Gbe management switch	Any

Table 13. Cisco Nexus 93180YC-FX-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch B	Eth1/51	40Gbe	Cisco UCS fabric interconnect B	Eth1/50
	Eth1/52	40Gbe	Cisco UCS fabric interconnect A	Eth1/50
	Eth1/1	25Gbe	Cisco Nexus 93180YC-FX A	Eth1/1
	Eth1/2	25Gbe	Cisco Nexus 93180YC-FX A	Eth1/2
	Eth1/3	25Gbe	Cisco Nexus 93180YC-FX A	Eth1/3
	Eth1/4	25Gbe	Cisco Nexus 93180YC-FX A	Eth1/4
	MGMT0	1Gbe	Gbe management switch	Any

Cisco UCS Fabric Interconnect 6454 Cabling

The following tables list the Cisco UCS FI 6454 cabling information.

Table 14. Cisco UCS Fabric Interconnect (FI) A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-A	FC 1/1	32G FC	Cisco MDS 9132T 32-Gb-A	FC 1/13
	FC 1/2	32G FC	Cisco MDS 9132T 32-Gb-A	FC 1/14
	Eth1/17-	25Gbe	UCS X9508 Chassis IFM-A Chassis 1	Intelligent Fabric Module 1 Port1-2
	Eth1/49	40Gbe	Cisco Nexus 93180YC-FX Switch A	Eth1/52
	Eth1/50	40Gbe	Cisco Nexus 93180YC-FX Switch B	Eth1/52
	Mgmt 0	1Gbe	Management Switch	Any
	L1	1Gbe	Cisco UCS FI - A	L1
	L2	1Gbe	Cisco UCS FI - B	L2

Table 15. Cisco UCS Fabric Interconnect (FI) B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-B	FC 1/1	32Gb FC	Cisco MDS 9132T 32-Gb-B	FC 1/13
	FC 1/2	32Gb FC	Cisco MDS 9132T 32-Gb-B	FC 1/14
	Eth1/17-	25Gbe	Cisco UCS X9508 Chassis IFM-	Intelligent Fabric

Local Device	Local Port	Connection	Remote Device	Remote Port
	18		B Chassis 1	Module 1 Port1-2
	Eth1/49	40Gbe	Cisco Nexus 93180YC-FX Switch A	Eth1/51
	Eth1/50	40Gbe	Cisco Nexus 93180YC-FX Switch B	Eth1/51
	Mgmt 0	1Gbe	Management Switch	Any
	L1	1Gbe	Cisco UCS FI - A	L1
	L2	1Gbe	Cisco UCS FI - B	L2

Procedure 1. Create vPC Peer-Link Between the Two Cisco Nexus Switches

Step 1. Log in as “admin” user into the Cisco Nexus Switch A.

Note: For vPC 1 as Peer-link, we used interfaces 53-54 for Peer-Link. You may choose the appropriate number of ports for your needs.

Step 2. Create the necessary port channels between devices by running these commands on both Cisco Nexus switches:

```
config terminal
feature vpc
feature lacp
vpc domain 50
peer-keepalive destination 173.37.52.104 source 173.37.52.103
exit
interface port-channel 10
description VPC peer-link
switchport mode trunk
switchport trunk allowed VLAN 1,40-46
spanning-tree port type network
vpc peer-link
interface Ethernet1/1
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,40-46,132
channel-group 10 mode active
no shutdown
exit

interface Ethernet1/2
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,40-46,132
channel-group 10 mode active
no shutdown
exit

interface Ethernet1/3
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,40-46,132
channel-group 10 mode active
no shutdown
exit
```

```
interface Ethernet1/4
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,40-46,132
channel-group 10 mode active
no shutdown
exit
copy running-config startup-config
```

Step 3. Log in as admin user into the Nexus Switch B and repeat the above steps to configure second Cisco Nexus switch.

Step 4. Make sure to change the peer-keepalive destination and source IP address appropriately for Cisco Nexus Switch B.

Procedure 2. Create vPC Configuration Between Cisco Nexus 93180YC-FX and Cisco Fabric Interconnects

Create and configure vPC 11 and 12 for the data network between the Cisco Nexus switches and fabric interconnects.

Note: Create the necessary port channels between devices, by running the following commands on both Cisco Nexus switches.

Step 1. Log in as admin user into Cisco Nexus Switch A and enter the following:

```
config terminal
interface port-channel11
description FI-A-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,40-46
spanning-tree port type edge trunk
vpc 11
no shutdown
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,40-46
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,40-46
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,40-46
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

Step 2. Log in as admin user into the Nexus Switch B and complete the following for the second switch configuration:

```
config Terminal
interface port-channel11
description FI-A-Uplink
```

```
switchport mode trunk
switchport trunk allowed VLAN 1,40-46
spanning-tree port type edge trunk
vpc 11
no shutdown
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,40-46
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,40-46
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,40-46
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

Verify all vPC Status is up on both Cisco Nexus Switches

Figure 19 shows the verification of the vPC status on both Cisco Nexus Switches.

Figure 19. vPC Description for Cisco Nexus Switch A and B

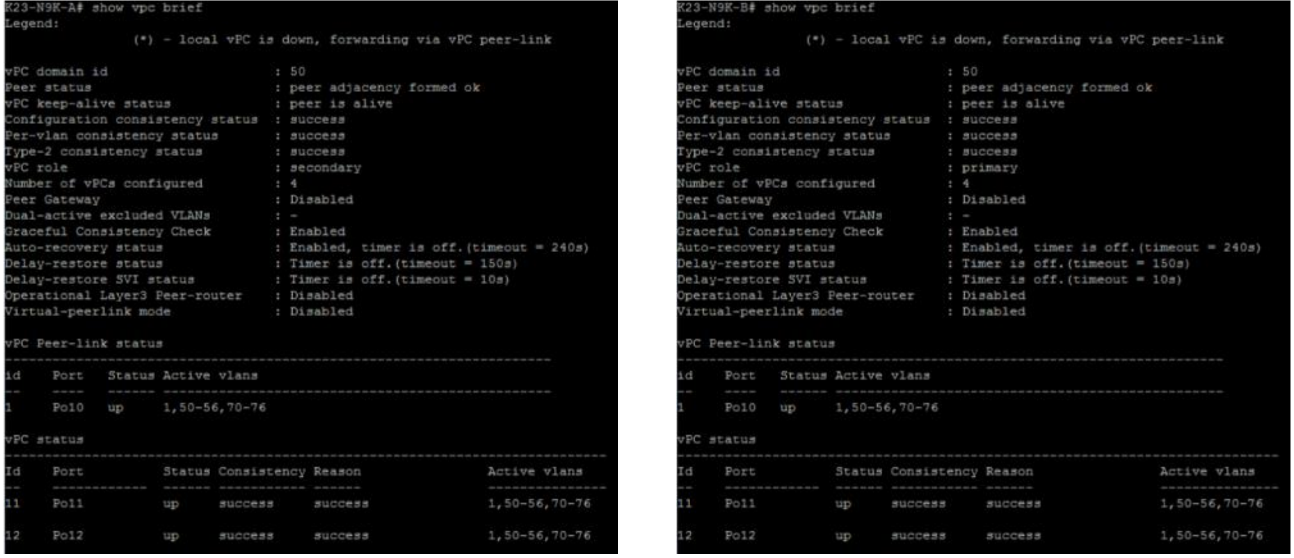

```
K23-N9K-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 50
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 150s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
Virtual-peerlink mode             : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ -------------------------------------------------
1     Po10   up     1,50-56,70-76

vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
11    Po11          up     success     success               1,50-56,70-76

12    Po12          up     success     success               1,50-56,70-76
```

```
K23-N9K-B# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 50
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 150s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
Virtual-peerlink mode             : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ -------------------------------------------------
1     Po10   up     1,50-56,70-76

vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
11    Po11          up     success     success               1,50-56,70-76

12    Po12          up     success     success               1,50-56,70-76
```

Cisco MDS 9132T 32-Gb FC Switch Configuration

Figure 15 illustrates the cable connectivity between the Cisco MDS 9132T 32-Gb switch and the Cisco 6454 Fabric Interconnects and Pure Storage FlashArray//X R3 storage.

Note: We used two 32Gb FC connections from each fabric interconnect to each MDS switch and two 32Gb FC connections from each Pure Storage FlashArray//X R3 array controller to each MDS switch.

Table 16. Cisco MDS 9132T-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-A	FC1/9	32Gb FC	Pure Storage FlashArray//X R3 Controller 0	CT0.FC0
	FC1/10	32Gb FC	Pure Storage FlashArray//X R3 Controller 1	CT1.FC0
	FC1/13	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/1
	FC1/14	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/2

Table 17. Cisco MDS 9132T-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-B	FC1/9	32Gb FC	Pure Storage FlashArray//X R3 Controller 0	CT0.FC2
	FC1/10	32Gb FC	Pure Storage FlashArray//X R3 Controller 1	CT1.FC2
	FC1/13	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/1
	FC1/14	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/2

Pure Storage FlashArray//X R3 to MDS SAN Fabric Connectivity

Pure Storage FlashArray//X R3 to MDS A and B Switches using VSAN 700 for Fabric A and VSAN 101 Configured for Fabric B

In this solution, two ports (ports FC1/9 and FC1/10) of MDS Switch A and two ports (ports FC1/9 and FC1/10) of MDS Switch B are connected to Pure Storage FlashArray as listed in Table 18. All ports connected to the Pure Storage FlashArray carry 32 Gb/s FC Traffic.

Table 18. MDS 9132T 32-Gb switch Port Connection to Pure Storage System

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-A	FC1/9	32Gb FC	Pure Storage FlashArray//X R3 Controller 0	CT0.FC0
	FC1/10	32Gb FC	Pure Storage FlashArray//X R3 Controller 1	CT1.FC0
Cisco MDS 9132T-B	FC1/9	32Gb FC	Pure Storage FlashArray//X R3 Controller 0	CT0.FC2
	FC1/10	32Gb FC	Pure Storage FlashArray//X R3 Controller 1	CT1.FC2

Procedure 1. Configure Features and Name for MDS Switch A and MDS Switch B

Follow these steps on both MDS switches.

Step 1. Log in as admin user into MDS Switch A:

```
config terminal
feature npiv
feature telnet
```

```
switchname FlashStack-MDS-A
copy running-config startup-config
```

Step 2. Log in as admin user into MDS Switch B. Repeat step 1 on MDS Switch B.

Procedure 2. Configure VSANs for MDS Switch A and MDS Switch B

Step 1. Log in as admin user into MDS Switch A. Create VSAN 700 for Storage Traffic:

```
config terminal
VSAN database
VSAN 700
exit
zone smart-zoning enable vsan 700
vsan database
VSAN 700 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 700
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

Step 2. Log in as admin user into MDS Switch B. Create VSAN 701 for Storage Traffic:

```
config terminal
VSAN database
vsan 701
exit
zone smart-zoning enable vsan 701
vsan database
vsan 701 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 701
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

Procedure 3. Create and Configure Fiber Channel Zoning

This procedure sets up the Fibre Channel connections between the Cisco MDS 9132T 32-Gb switches, the Cisco UCS Fabric Interconnects, and the Pure Storage FlashArray systems.

Note: Before you configure the zoning details, decide how many paths are needed for each LUN and extract the WWPN numbers for each of the HBAs from each server. We used 2 HBAs for each Server. One of the HBAs (HBA-A) is connected to MDS Switch-A and other HBAs (HBA-B) is connected to MDS Switch-B.

Step 1. Log into the Cisco Intersight portal as a user with account administrator role.

Step 2. From the Service Selector drop-down list, choose Infrastructure Service.

Step 3. Navigate to Configure > Pools. Filter WWPN type pools.

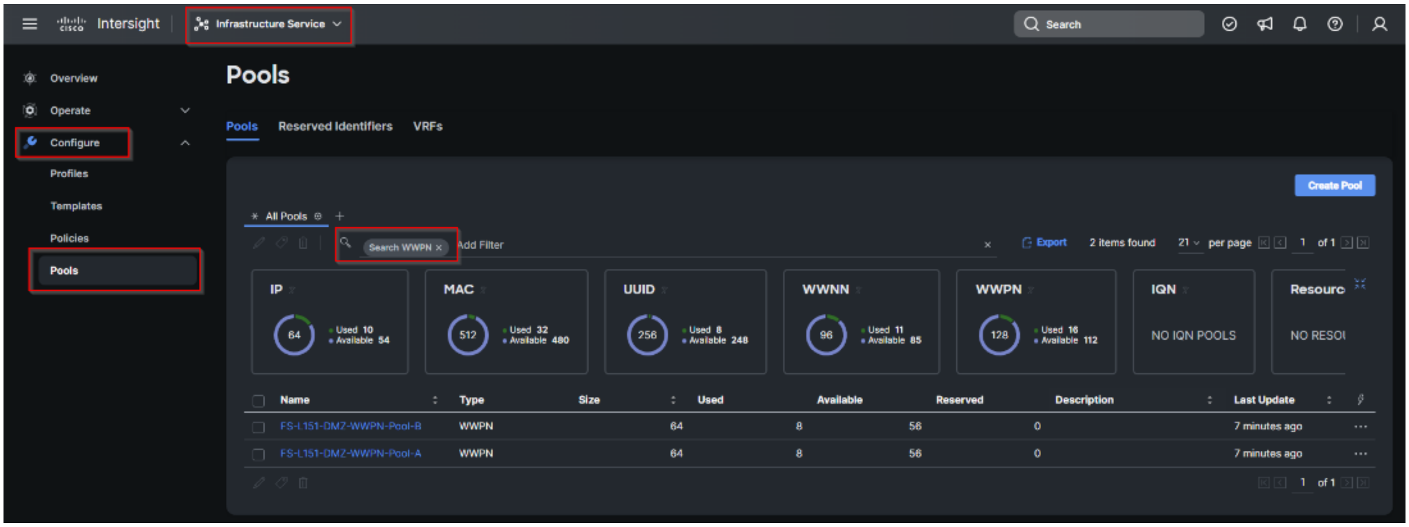

Pools

IP	MAC	UUID	WWNN	WWPN	IQN	Resourc
64 Used 10 Available 54	512 Used 32 Available 480	256 Used 8 Available 248	96 Used 11 Available 85	128 Used 16 Available 112	NO IQN POOLS	NO RESOU

Name	Type	Size	Used	Available	Reserved	Description	Last Update
FS-L151-DMZ-WWPN-Pool-B	WWPN	64	8	56	0		7 minutes ago
FS-L151-DMZ-WWPN-Pool-A	WWPN	64	8	56	0		7 minutes ago

Step 4. Select Usage tab and collect the WWPNs and profiles to which they are assigned.

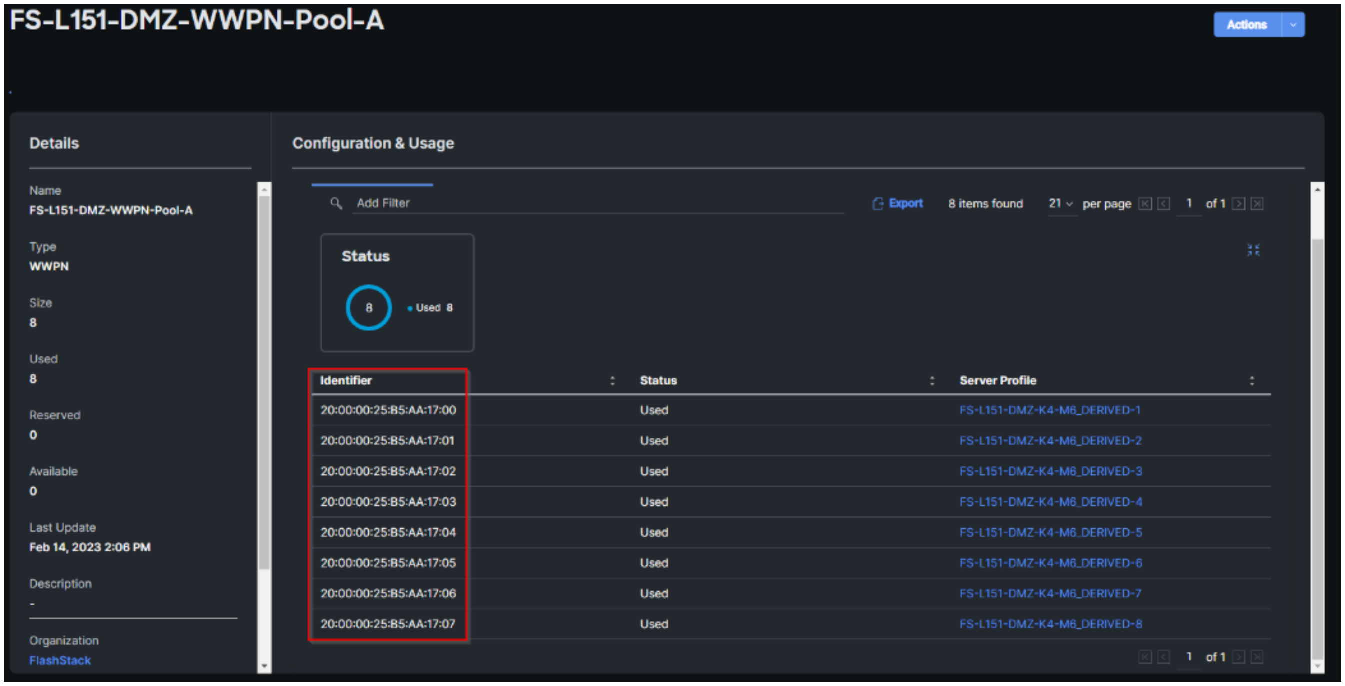

FS-L151-DMZ-WWPN-Pool-A

Name
FS-L151-DMZ-WWPN-Pool-A

Type
WWPN

Size
8

Used
8

Reserved
0

Available
0

Last Update
Feb 14, 2023 2:06 PM

Description
-

Organization
FlashStack

Identifier	Status	Server Profile
20:00:00:25:B5:AA:17:00	Used	FS-L151-DMZ-K4-M6_DERIVED-1
20:00:00:25:B5:AA:17:01	Used	FS-L151-DMZ-K4-M6_DERIVED-2
20:00:00:25:B5:AA:17:02	Used	FS-L151-DMZ-K4-M6_DERIVED-3
20:00:00:25:B5:AA:17:03	Used	FS-L151-DMZ-K4-M6_DERIVED-4
20:00:00:25:B5:AA:17:04	Used	FS-L151-DMZ-K4-M6_DERIVED-5
20:00:00:25:B5:AA:17:05	Used	FS-L151-DMZ-K4-M6_DERIVED-6
20:00:00:25:B5:AA:17:06	Used	FS-L151-DMZ-K4-M6_DERIVED-7
20:00:00:25:B5:AA:17:07	Used	FS-L151-DMZ-K4-M6_DERIVED-8

Step 5. Connect to the Pure Storage System Health and go to the Connections tab and extract the WWPN of FC Ports connected to the Cisco MDS Switches from Array Ports section.

Note: We connected 4 FC ports from Pure Storage System to Cisco MDS Switches. FC ports CT0.FC0, CT1.FC0 are connected to MDS Switch-A and similarly FC ports CT1.FC2, CT0.FC2 are connected to MDS Switch-B.

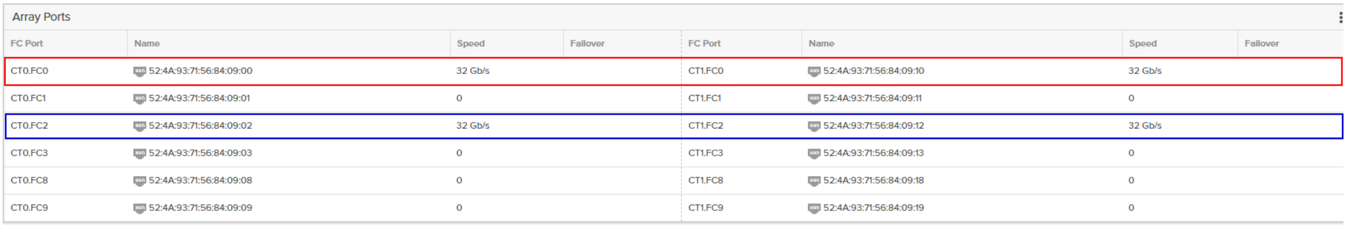
Array Ports

FC Port	Name	Speed	Failover	FC Port	Name	Speed	Failover
CT0.FC0	52:4A:93:71:56:84:09:00	32 Gb/s		CT1.FC0	52:4A:93:71:56:84:09:10	32 Gb/s	
CT0.FC1	52:4A:93:71:56:84:09:01	0		CT1.FC1	52:4A:93:71:56:84:09:11	0	
CT0.FC2	52:4A:93:71:56:84:09:02	32 Gb/s		CT1.FC2	52:4A:93:71:56:84:09:12	32 Gb/s	
CT0.FC3	52:4A:93:71:56:84:09:03	0		CT1.FC3	52:4A:93:71:56:84:09:13	0	
CT0.FC8	52:4A:93:71:56:84:09:08	0		CT1.FC8	52:4A:93:71:56:84:09:18	0	
CT0.FC9	52:4A:93:71:56:84:09:09	0		CT1.FC9	52:4A:93:71:56:84:09:19	0	

Procedure 4. Create Device Aliases for Fiber Channel Zoning for SAN Boot Paths and Datapaths on Cisco MDS Switch A

Step 1. Log in as admin user and run the following commands from the global configuration mode:

```
configure terminal
device-alias mode enhanced
device-alias database
device-alias name Host-FCP-1-HBA0 pwwn 20:00:00:25:B5:AA:17:00
device-alias name X20R3-CT0-FC0 pwwn 52:4A:93:71:56:84:09:00
device-alias name X20R3-CT1-FC0 pwwn 52:4A:93:71:56:84:09:10
exit
device-alias commit
```

Procedure 5. Create Device Aliases for Fiber Channel Zoning for SAN Boot Paths and Datapaths on Cisco MDS Switch B

Step 1. Log in as admin user and run the following commands from the global configuration mode:

```
configure terminal
device-alias mode enhanced
device-alias database
device-alias name Host-FCP-1-HBA1 pwwn 20:00:00:25:b5:bb:17:00
device-alias name X20R3-CT0-FC2 pwwn 52:4A:93:71:56:84:09:02
device-alias name X20R3-CT1-FC2 pwwn 52:4A:93:71:56:84:09:12
exit
device-alias commit
```

Procedure 6. Create Fiber Channel Zoning for Cisco MDS Switch A for each Service Profile

Step 1. Log in as admin user and create the zone:

```
configure terminal
zone name FlashStack-Fabric-A VSAN 700
    member device-alias X20R3-CT0-FC0 target
    member device-alias X20R3-CT1-FC0 target
    member device-alias Host-FCP-1-HBA0 init
```

Step 2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the created zones as members:

```
configure terminal
zoneset name -Fabric-A VSAN 700
    member FlashStack-Fabric-A
```

Step 3. Activate the zone set by running following commands:

```
zoneset activate name -Fabric-A VSAN 700
exit
copy running-config startup-config
```

Procedure 7. Create Fiber Channel Zoning for Cisco MDS Switch B for each Service Profile

Step 1. Log in as admin user and create the zone as shown below:

```
configure terminal zone name FlashStack-Fabric-B vsan 701
    member device-alias X20R3-CT0-FC2 target
    member device-alias X20R3-CT1-FC2 target
    member device-alias Host-FCP-1-HBA1 init
```

Step 2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members:

```
zoneset name -Fabric-B vsan 701
    member FlashStack-Fabric-B
```

Step 3. Activate the zone set by running following commands:

```
zoneset activate name -Fabric-B vsan 701
exit
copy running-config startup-config
```

Configure Pure Storage FlashArray//X R3

The design goal of the reference architecture is to best represent a real-world environment as closely as possible. The approach included the features of Cisco UCS to rapidly deploy stateless servers and use Pure Storage FlashArray's boot LUNs to provision the ESXi on top of Cisco UCS. Zoning was performed on the Cisco MDS 9132T 32-Gb switches to enable the initiators discover the targets during boot process.

A Service Profile was created within Cisco UCS Manager to deploy the thirty-two servers quickly with a standard configuration. SAN boot volumes for these servers were hosted on the same Pure Storage FlashArray//X R3. Once the stateless servers were provisioned, following process was performed to enable rapid deployment of thirty-two blade servers.

Each Blade Server has dedicated single LUN to install operating system and all the thirty-two blade servers configured to boot from SAN. For this solution, we installed the vSphere ESXi 8.0 Cisco Custom ISO on this LUNs to create solution.

Using logical servers that are disassociated from the physical hardware removes many limiting constraints around how servers are provisioned. Cisco UCS Service Profiles contain values for a server's property settings, including virtual network interface cards (vNICs), MAC addresses, boot policies, firmware policies, fabric connectivity, external management, and HA information. The service profiles represent all the attributes of a logical server in Cisco UCS model. By abstracting these settings from the physical server into a Cisco Service Profile, the Service Profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, Service Profiles can, at any time, be migrated from one physical server to another. Furthermore, Cisco is the only hardware provider to offer a truly unified management platform, with Cisco UCS Service Profiles and hardware abstraction capabilities extending to both blade and rack servers.

In addition to the service profiles, the use of Pure Storage's FlashArray's with SAN boot policy provides the following benefits:

- Scalability - Rapid deployment of new servers to the environment in a very few steps.
- Manageability - Enables seamless hardware maintenance and upgrades without any restrictions. This is a significant benefit in comparison to another appliance model like Exadata.
- Flexibility - Easy to repurpose physical servers for different applications and services as needed.
- Availability - Hardware failures are not impactful and critical. In rare case of a server failure, it is easier to associate the logical service profile to another healthy physical server to reduce the impact.

Configure Host, WWNs, and Volume Connectivity with FlashArray Management Tools

Procedure 1. Configure Host

Note: Before using a boot volume (LUN) by a Cisco UCS Blade Server, a host representing this blade server must be defined on Pure Storage FlashArray.

Step 1. Log into Pure Storage FlashArray Management interface.

Step 2. Click the Storage tab.

Step 3. Click the + sign in the Hosts section and select Create Host.

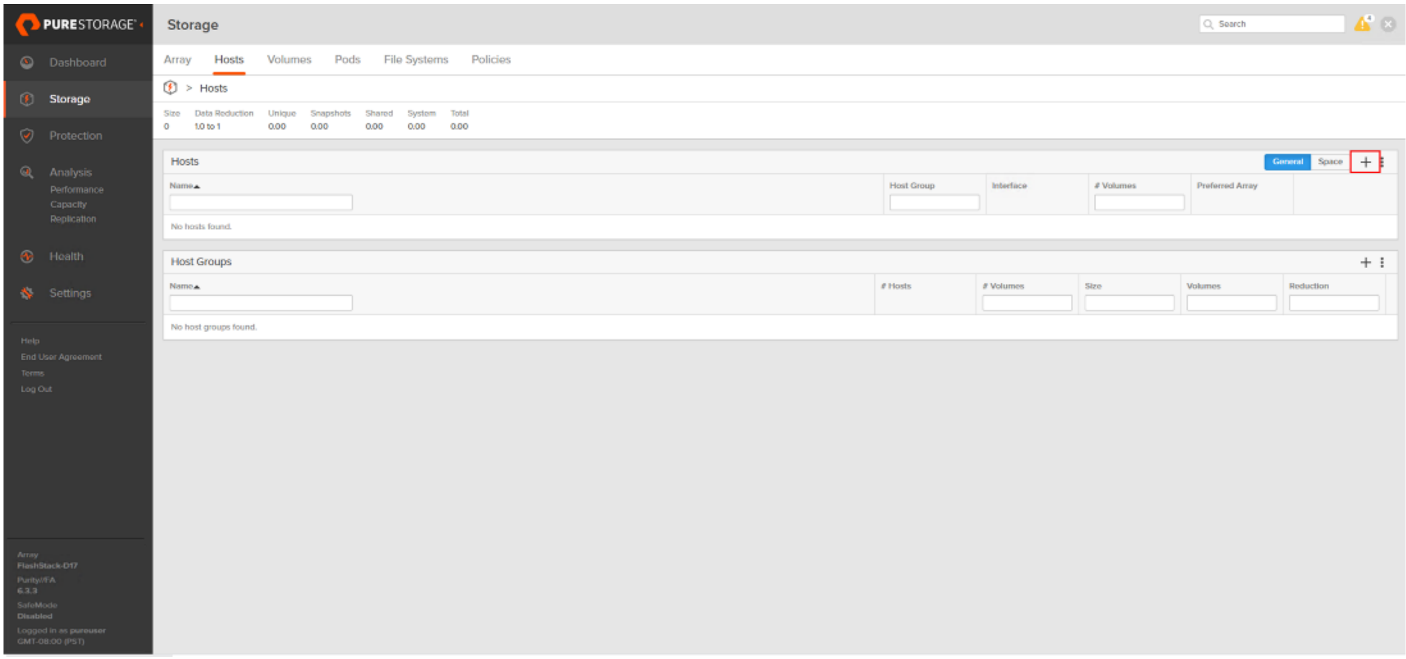

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	0.00	0.00	0.00

No hosts found.

No host groups found.

Array FlashStack-D17
Purity//FA 6.3.3
SafeMode Disabled
Logged in as pureuser
GMT-08:00 (PST)

Step 4. Click Create Multiple to create Host entries under the Hosts category.

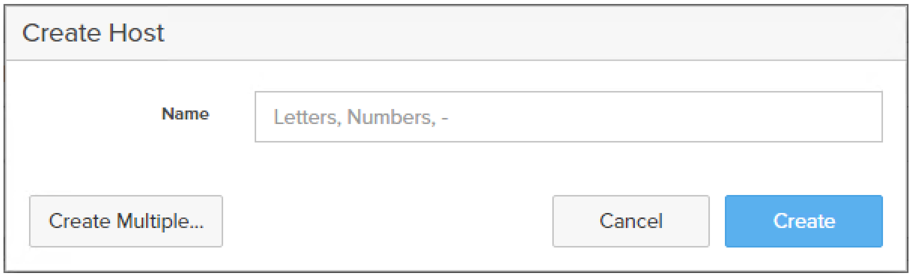

Create Host

Name: Letters, Numbers, -

Step 5. Enter the required information and click Create.

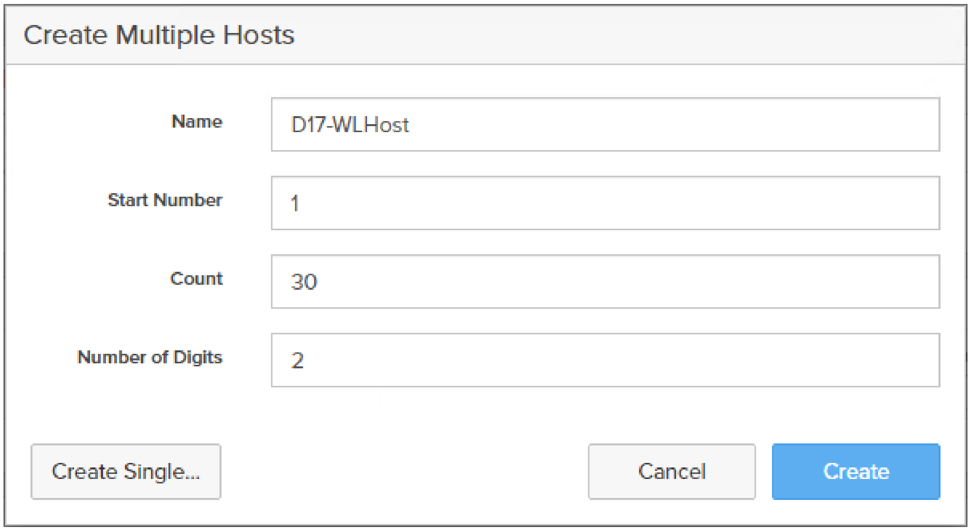

Create Multiple Hosts

Name: D17-WLHost

Start Number: 1

Count: 30

Number of Digits: 2

Step 6. Select one of the newly created hosts, in Host Ports section from the drop-down list select Configure WWNs.

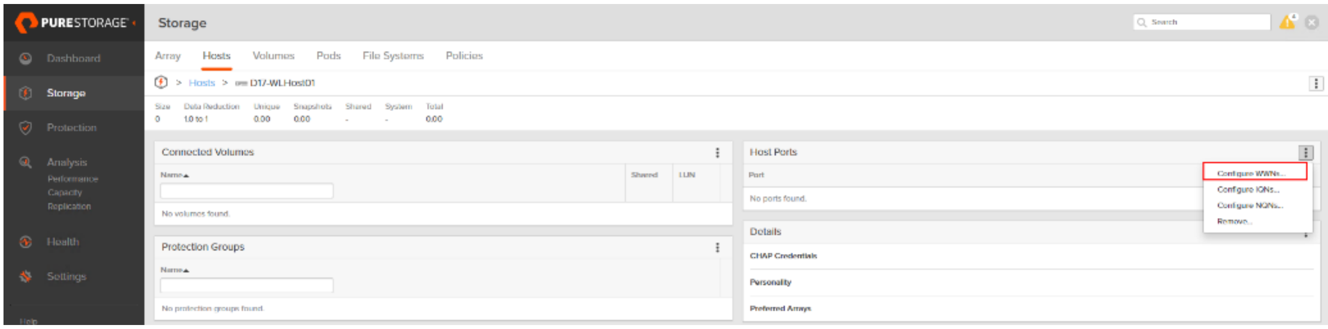

Step 7. Select the list of WWNs that belongs to the host in the next window and click Add.

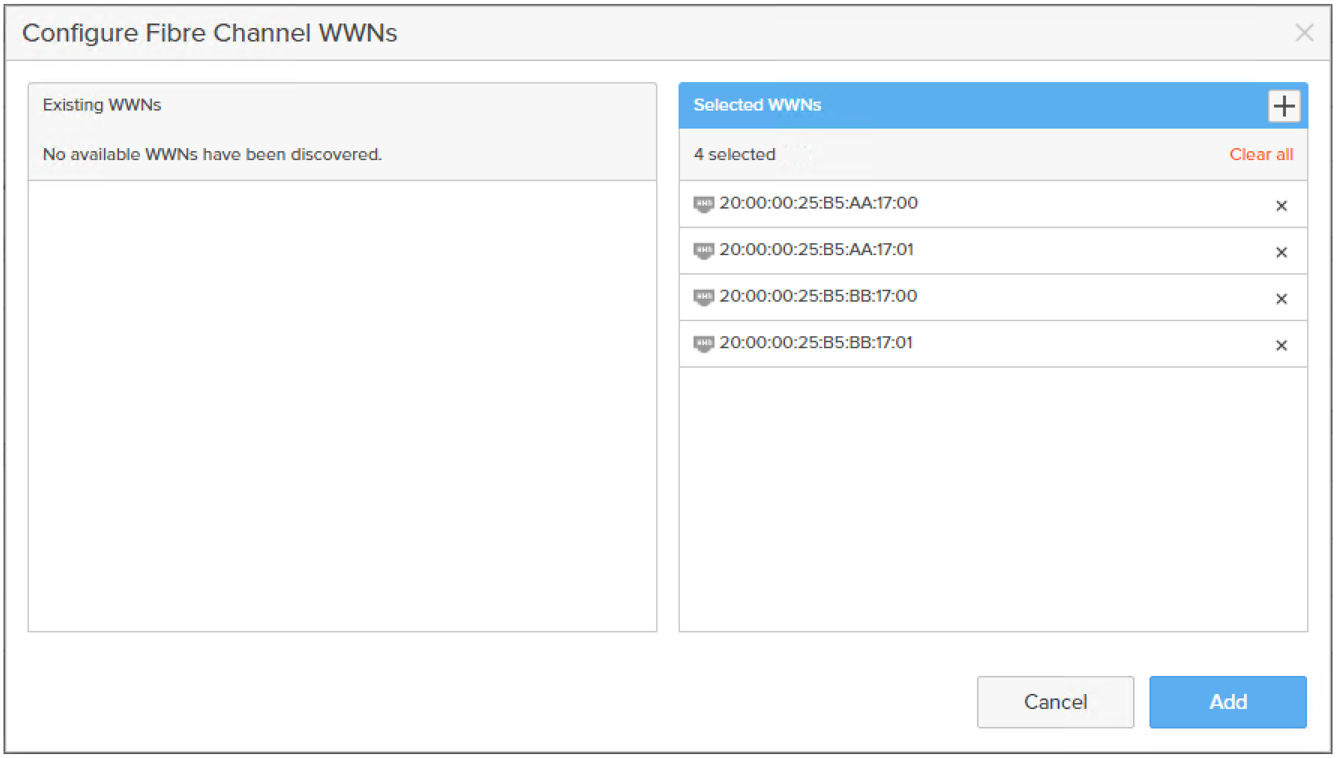

Configure Fibre Channel WWNs

Existing WWNs

No available WWNs have been discovered.

Selected WWNs

4 selected

- 20:00:00:25:B5:AA:17:00
- 20:00:00:25:B5:AA:17:01
- 20:00:00:25:B5:BB:17:00
- 20:00:00:25:B5:BB:17:01

Note: Make sure the zoning has been setup to include the WWNs details of the initiators along with the target, without which the SAN boot will not work.

Note: WWNs will appear only if the appropriate FC connections were made, and the zones were setup on the underlying FC switch, and the servers were powered on.

Note: Alternatively, the WWN can be added manually by clicking the + in the Selected WWNs section and manually inputting the blade’s WWNs.

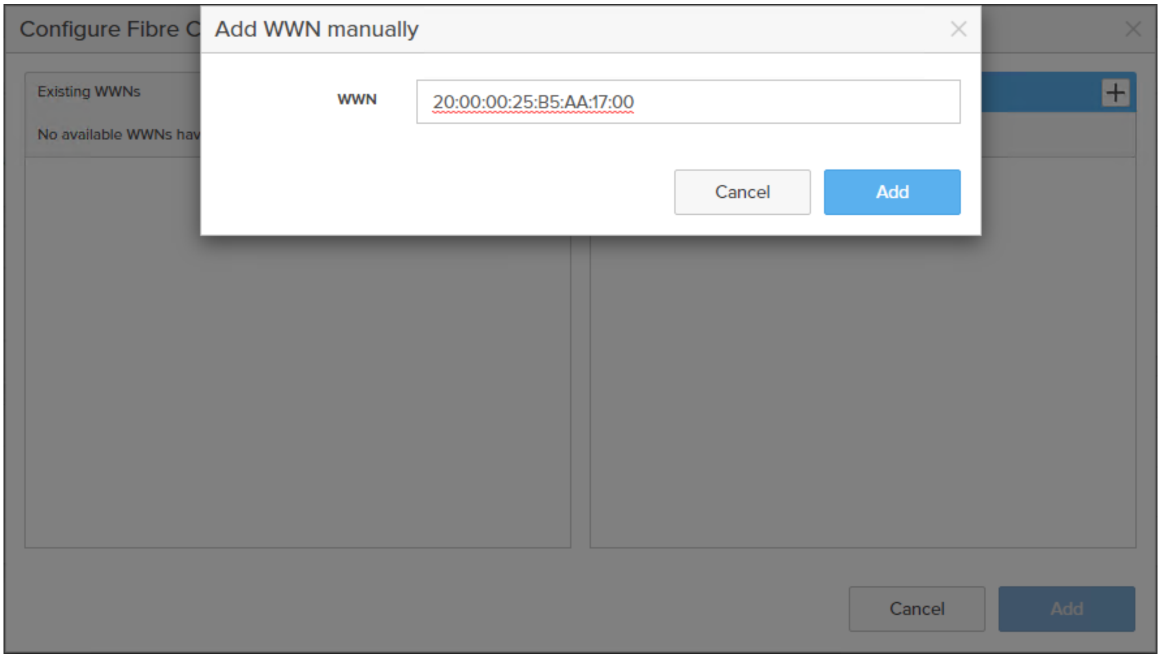

Add WWN manually

WWN: 20:00:00:25:B5:AA:17:00

Procedure 2. Configure Volume Connectivity

Step 1. Click the Storage tab.

Step 2. Click the + sign in the Volumes section and click Create Volume.

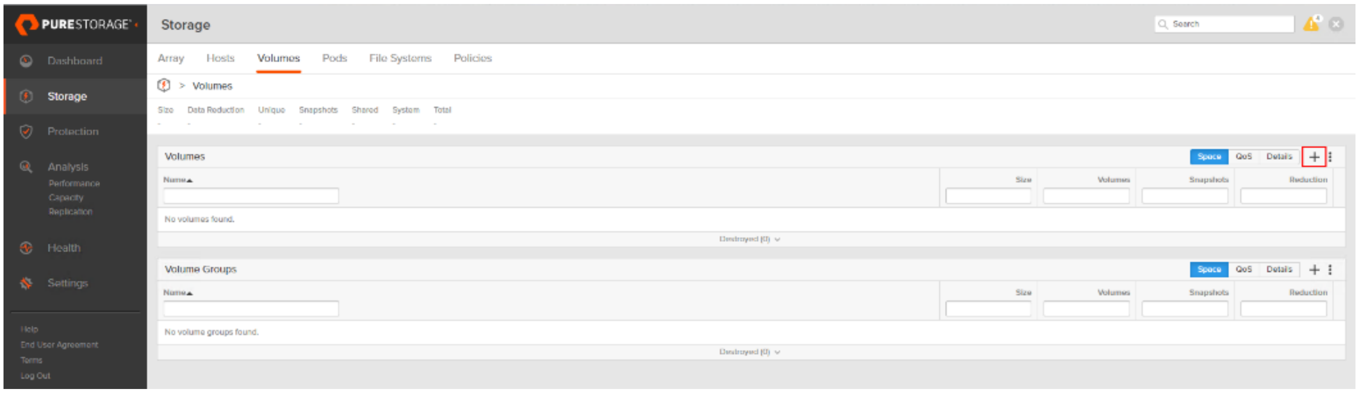

No volumes found.

No volume groups found.

Step 3. Click Create Multiple to open Create Multiple Volumes wizard.

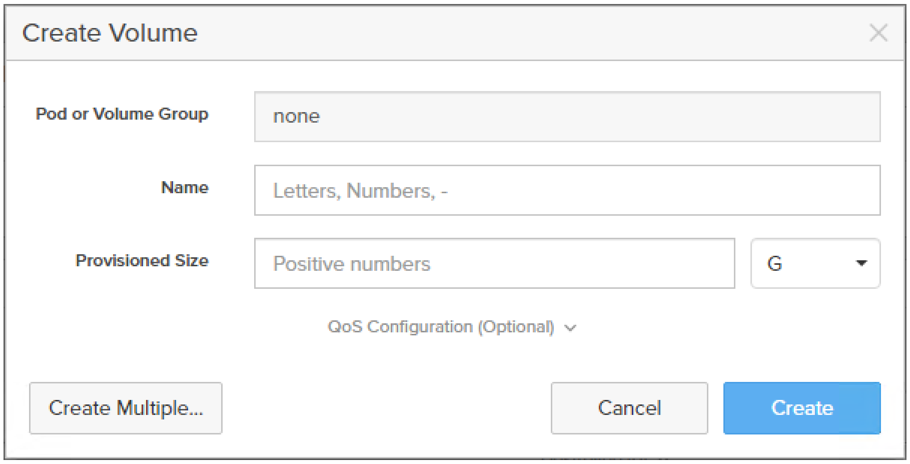

Create Volume

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

Step 4. Provide the common name of the volume, size, choose the size type (KB, MB, GB, TB, PB) and click Create to create volumes.

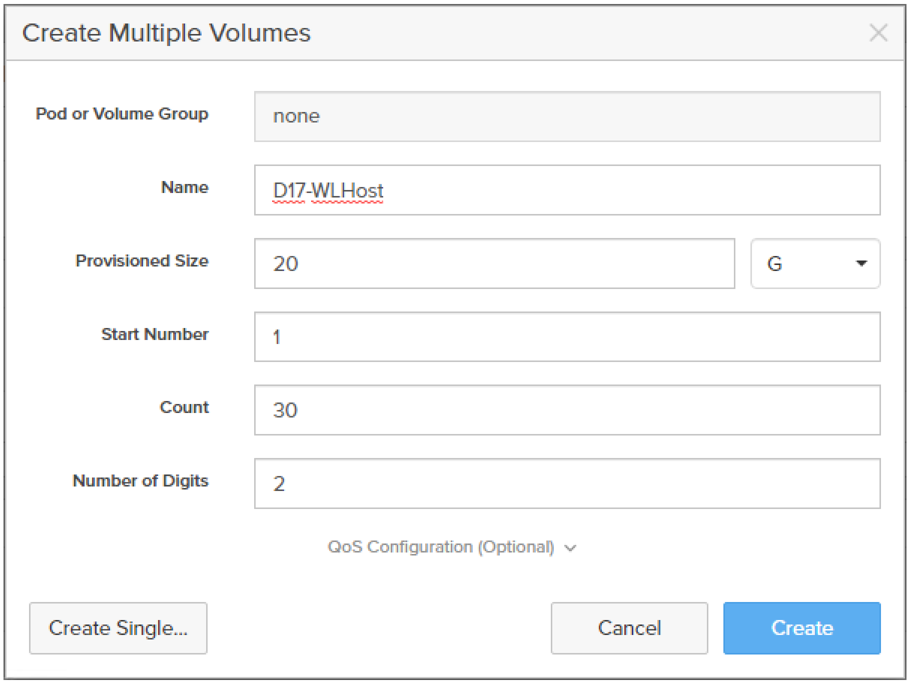

Create Multiple Volumes

Pod or Volume Group	none
Name	D17-WLHost
Provisioned Size	20 G
Start Number	1
Count	30
Number of Digits	2

Step 5. Select one of the hosts and in Connected Volumes section from the drop-down list select Connect.

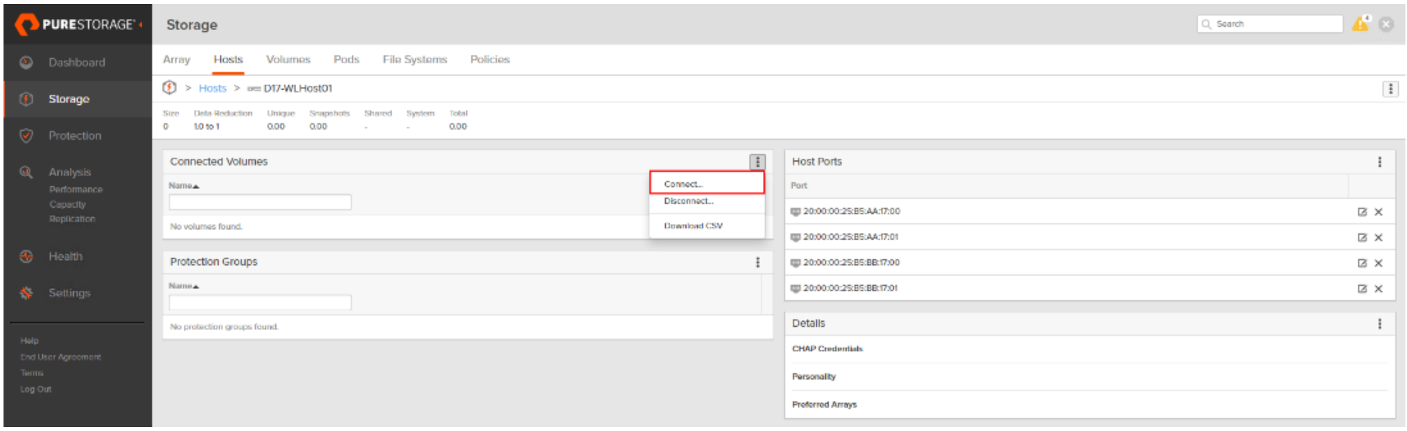

> Hosts > D17-WLHost01

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:AA:17:00
20:00:00:25:B5:AA:17:01
20:00:00:25:B5:BB:17:00
20:00:00:25:B5:BB:17:01

Details

CHAP Credentials

Personality

Preferred Arrays

Step 6. In the Connect Volumes to Host wizard select the volume configured for ESXi installation, click Connect.

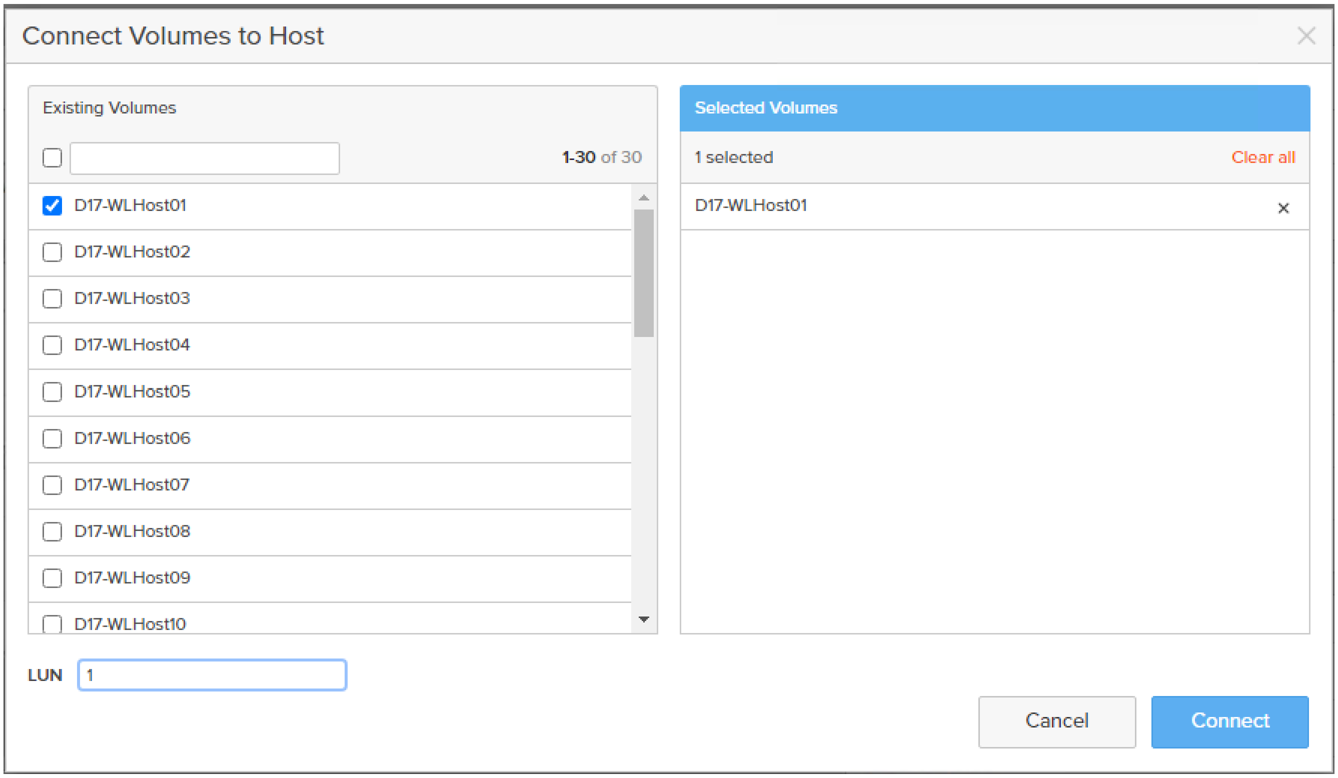

Connect Volumes to Host

Existing Volumes

□ 1-30 of 30

- ☑ D17-WLHost01
- □ D17-WLHost02
- □ D17-WLHost03
- □ D17-WLHost04
- □ D17-WLHost05
- □ D17-WLHost06
- □ D17-WLHost07
- □ D17-WLHost08
- □ D17-WLHost09
- □ D17-WLHost10

Selected Volumes

1 selected

D17-WLHost01

LUN: 1

Note: Make sure the SAN Boot Volumes has the LUN ID "1" since this is important while configuring Boot from SAN. You will also configure the LUN ID as "1" when configuring Boot from SAN policy in Cisco UCS Manager.

Note: More LUNs can be connected by adding a connection to existing or new volume(s) to an existing node.

Configure File Services

FA File services can be activated by Pure Storage Technical Services (Support). Please refer to FA File Services Support Matrix to verify that your hardware offers support for running File Services.

Currently, all FA File services activations require Pure Storage support approval. Customers can work with their local account representatives to obtain approval to activate File Services.

For additional information on FA File Services setup and configuration see:

- FA File Services Quick Start Guide
- FA File Services Best Practices

Procedure 1. Create Virtual Interface(s)

The VIF provides high-availability network access across 2 physical Ethernet ports per array controller. Each VIF requires 2 physical ports per controller. Any physical ethernet port can be used with the restriction that any port that is in use by management services, a bond, or subnet configuration cannot be part of a VIF. For the maximum number of VIFs supported, please see the FA File Services Limits KB.

Note: VIFs are created by CLI over SSH, configured and enabled using the Management Console. An account with administrator privileges is required.

Step 1. Connect to the array via SSH.

Step 2. Run the following syntax to create the VIF on the array:

```
purenetwork create vif -subinterfacelist ct0.ethX,ct1.ethX,ct0.ethY,ct1.ethY <name of interface>
```

Procedure 2. Configure and Enable the Virtual Interface for File Services

Step 1. Connect to the array GUI.

Step 2. Navigate to Settings > Network.

Step 3. Locate the File VIF in the interface list and click the edit icon.

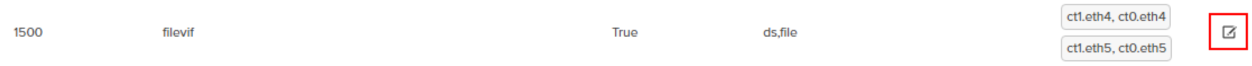

1500 filevif True ds,file ct1.eth4, ct0.eth4
ct1.eth5, ct0.eth5

Step 4. In the Edit Interface dialog turn on the Enabled option, provide the IP Address, Netmask, and Gateway used by the interface. Click Save.

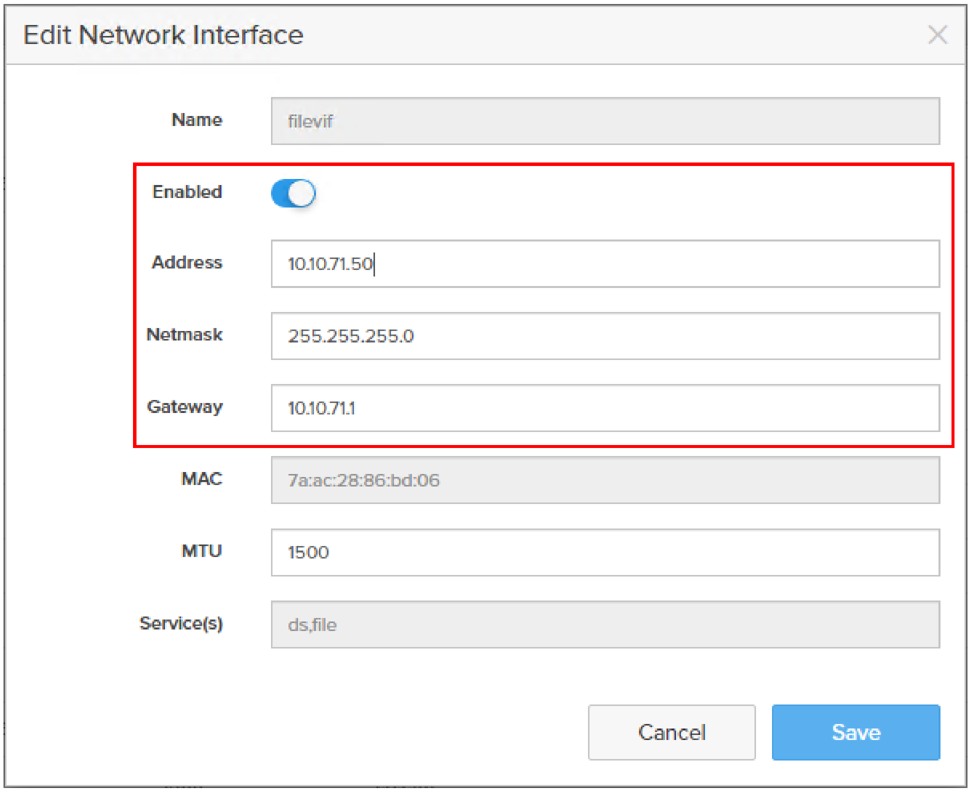

Edit Network Interface

Name: filevif

Enabled: ●

Address: 10.10.71.50

Netmask: 255.255.255.0

Gateway: 10.10.71.1

MAC: 7a:ac:28:86:bd:06

MTU: 1500

Service(s): ds,file

Step 5. Scroll to the bottom of the Network tab and click the edit icon for DNS Settings.

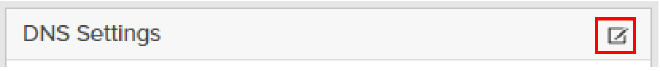

Step 6. In the Edit DNS Settings dialog, enter desired values for Domain and DNS server Ips. Click Save.

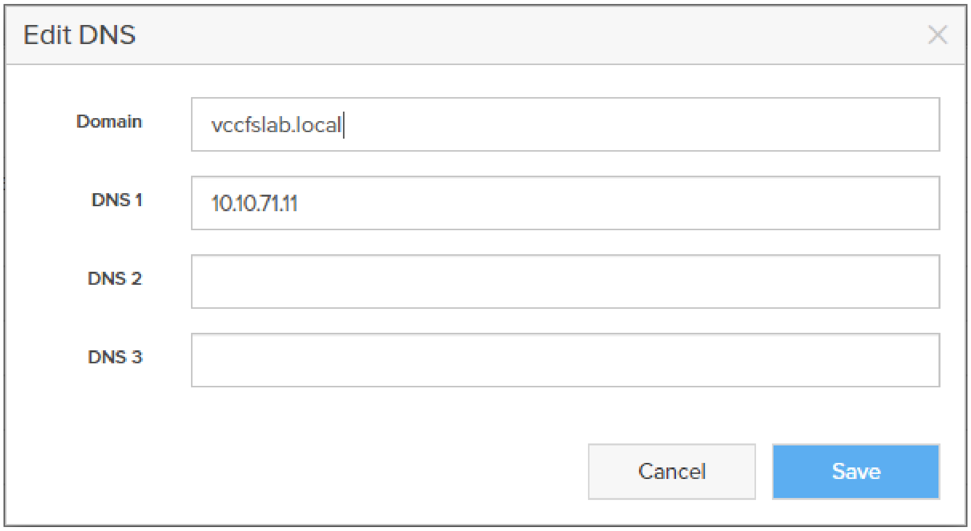

Domain: vccfslab.local

DNS 1: 10.10.71.11

DNS 2:

DNS 3:

Note: More than one DNS server can be configured with the caveat that all DNS servers must have a record for Directory Service servers such as LDAP or Microsoft Active Directory.

Procedure 3. Create Active Directory Account for the Array

Step 1. Navigate to Settings > Access > Active Directory Accounts.

Step 2. To open the Create Dialog, click the + icon.

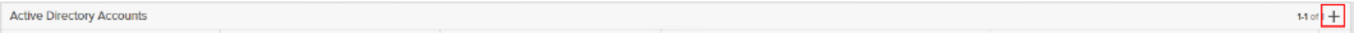

Step 3. Enter the following information:

- Name = Array management name for this AD account
- Domain = AD domain name
- Computer Name = Computer Object name within AD
- User = Domain user that can create computer objects and join to the domain.
- Password = Users password for the above domain user

Step 4. Click Create to finalize AD account creation.

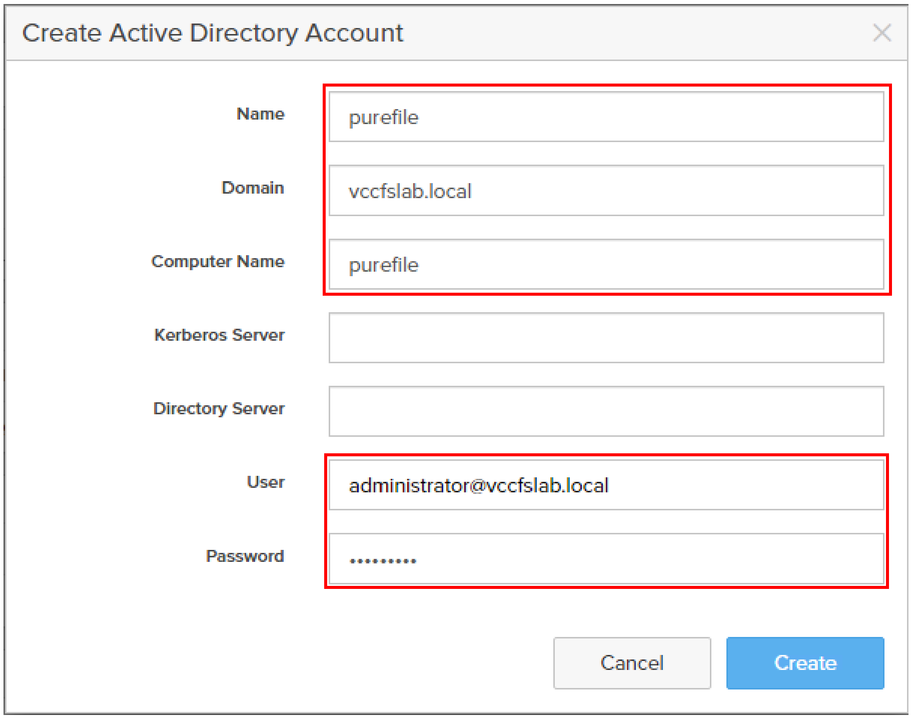
Create Active Directory Account

Name: purefile

Domain: vccfslab.local

Computer Name: purefile

Kerberos Server:

Directory Server:

User: administrator@vccfslab.local

Password: •••••••••

Procedure 4. Create a File System and Shared Directory

Step 1. Navigate to Storage > File Systems.

Step 2. Click the + icon.

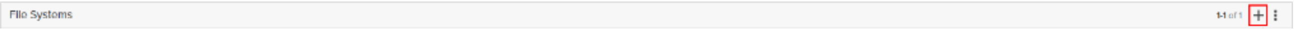

Step 3. In Create File System enter a file system name and click Create.

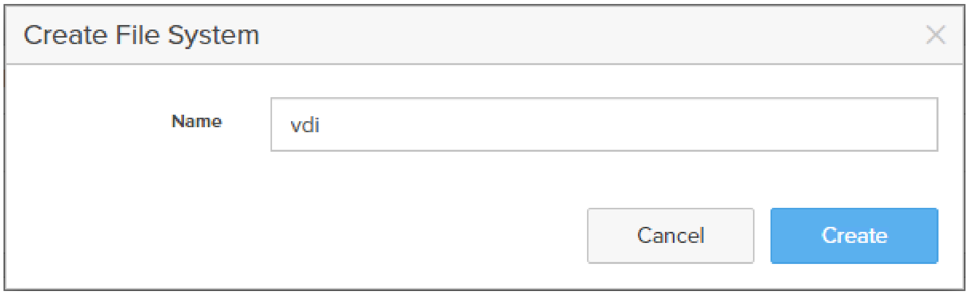
Create File System

Name: vdi

Step 4. Navigate to Storage > File Systems > Directories.

Step 5. Click the + icon.

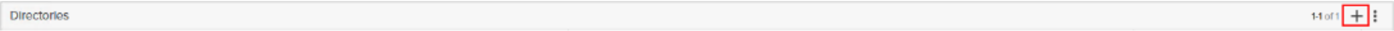

Step 6. In Create Directory, enter Select a file system from the drop-down list, enter the desired management name of the directory, and enter the directory path in the file system. (for example, dir or /dir, for sub-level directories /dir/subdir or /dir/subdir/subdir1 can be used). Click Create.

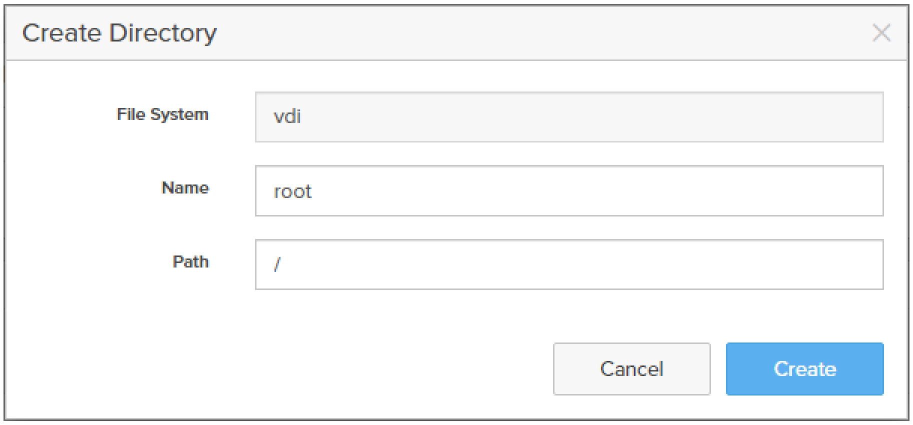

Create Directory

File System: vdi

Name: root

Path: /

Note: Polices for exports/shares/snapshots can only be attached to managed directories at the file system root or 1 level deep (/ and /dir in the example above). Space and performance metrics can be seen at all levels of managed directories.

Step 7. Navigate to Storage > Policies.

Step 8. Click the + icon.

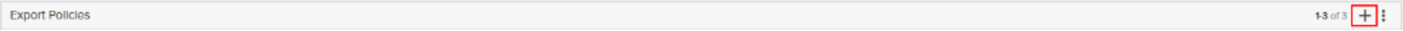

Step 9. In the Create Export Policy pop-up choose SMB from the Type drop-down list and enter a name for the policy. Click Create.

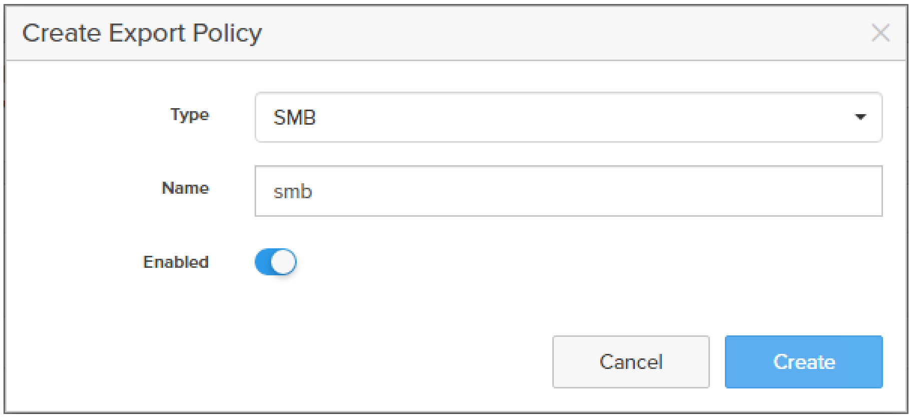

Create Export Policy

Type: SMB

Name: smb

Enabled: ●

Step 10. Click Created Policy and click the + icon.

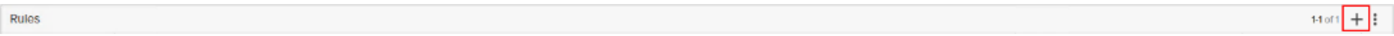

Step 11. Complete the Client filter for read-write access and click Add to complete the rule creation.

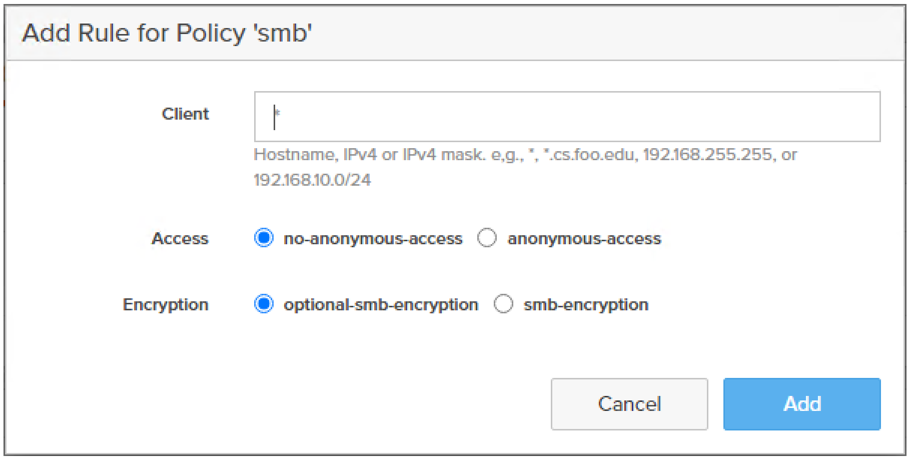

Add Rule for Policy 'smb'

Client: *

Hostname, IPv4 or IPv4 mask. e,g., *, *.cs.foo.edu, 192.168.255.255, or 192.168.10.0/24

Access: ● no-anonymous-access ○ anonymous-access

Encryption: ● optional-smb-encryption ○ smb-encryption

Step 12. Attach the export policy(s) to a managed directory. Click the + icon.

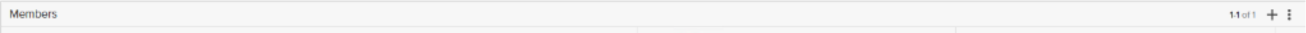

Step 13. Select a managed directory from the drop-down list, enter a share/export name, and click Create.

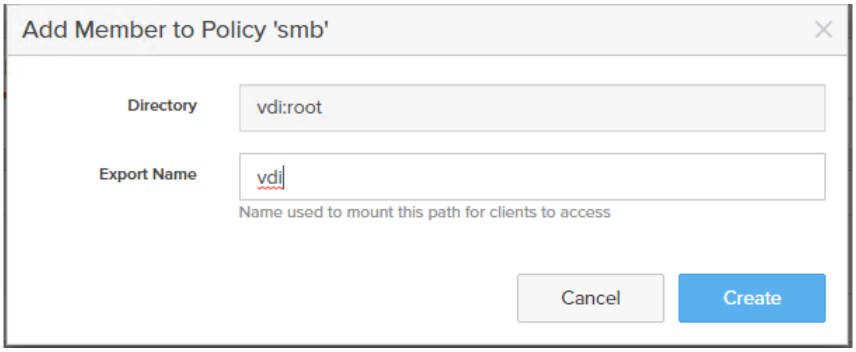

Add Member to Policy 'smb'

Directory: vdi:root

Export Name: vdi

Name used to mount this path for clients to access

Step 14. Verify access to the created share from the Windows client.

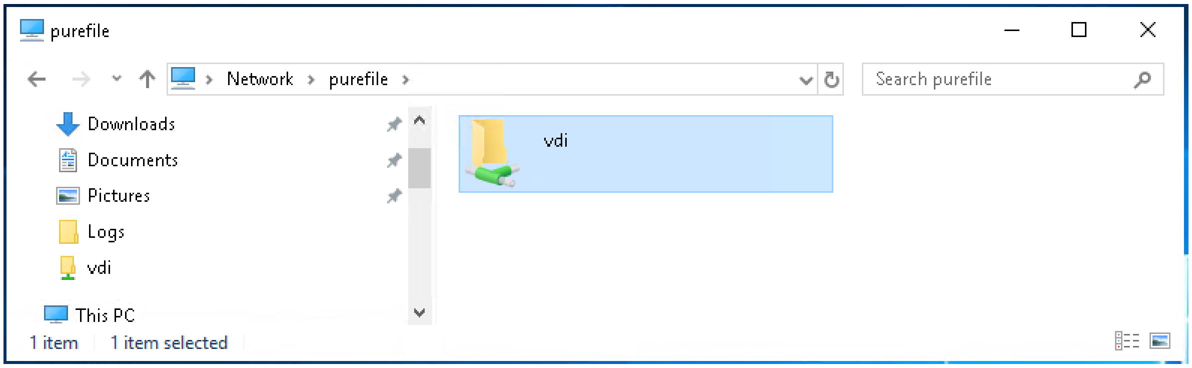

Install and Configure VMware ESXi 8.0

This section explains how to install VMware ESXi 8.0 in an environment.

There are several methods to install 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 install ESXi on boot logical unit number (LUN). Upon completion of steps outlined here, ESXi hosts will be booted from their corresponding SAN Boot LUNs.

Procedure 1. Download Cisco Custom Image for VMware vSphere ESXi 8.0

Step 1. To download the Cisco Custom Image for VMware ESXi 7.0 Update 3d, from the VMware vSphere Hypervisor 7.0 U3d page click the Custom ISOs tab.

Procedure 2. Install VMware vSphere ESXi 7.0 U3d

Step 1. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Servers.

Step 2. Right-click on ... icon for the server being access and select Launch vKVM.

Step 3. Click Boot Device and then select vKVM Mapped vDVD.

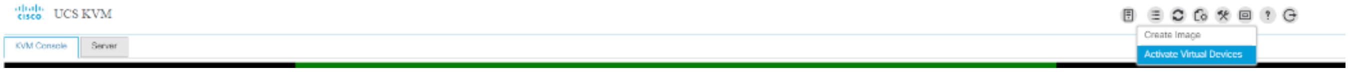

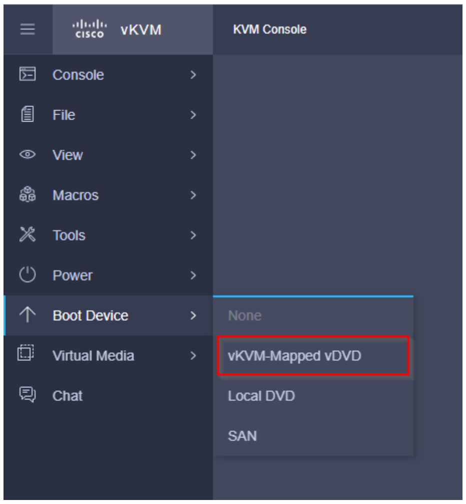

Step 4. Browse to the ESXi iso image file. Click Map Drive to mount the ESXi ISO image.

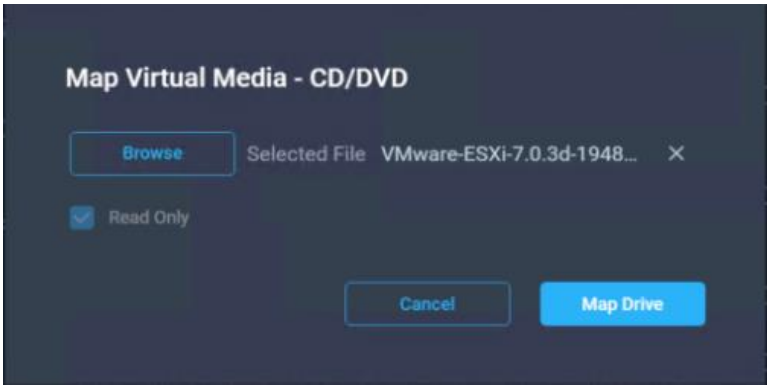
Map Virtual Media - CD/DVD

Selected File VMware-ESXi-7.0.3d-1948...

Step 5. Boot into ESXi installer and follow the prompts to complete installing VMware vSphere ESXi hypervisor.

Step 6. When selecting a storage device to install ESXi, select Remote LUN provisioned through Pure Storage Administrative console and access through FC connection.

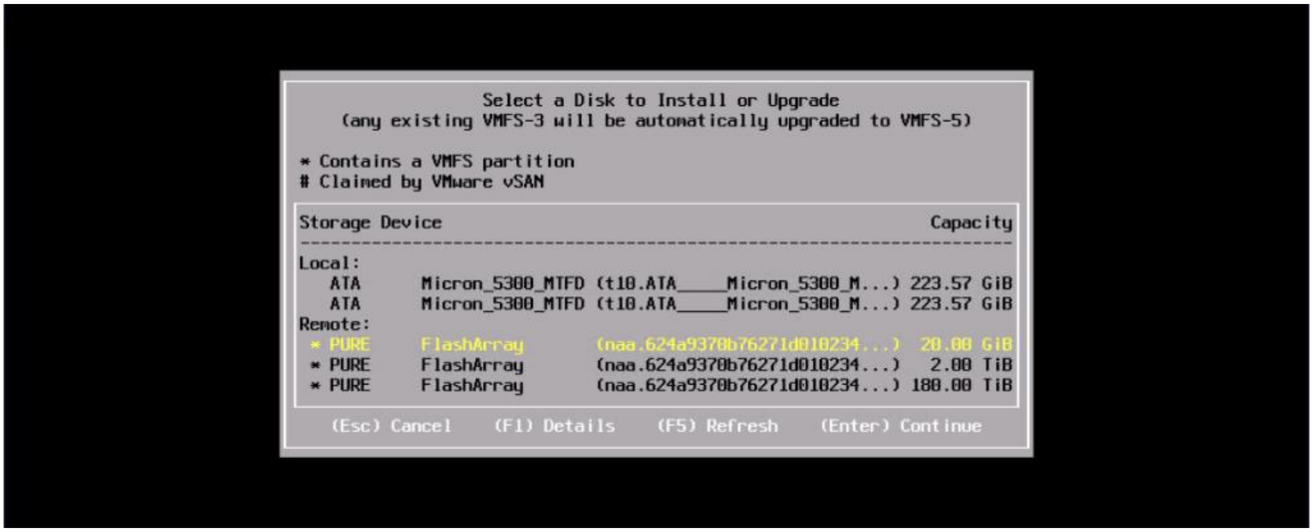

```
             Select a Disk to Install or Upgrade
   (any existing VMFS-3 will be automatically upgraded to VMFS-5)

* Contains a VMFS partition
# Claimed by VMware vSAN

Storage Device                                                  Capacity
------------------------------------------------------------------------
Local:
   ATA      Micron_5300_MTFD (t10.ATA_____Micron_5300_M...) 223.57 GiB
   ATA      Micron_5300_MTFD (t10.ATA_____Micron_5300_M...) 223.57 GiB
Remote:
 * PURE     FlashArray       (naa.624a9370b76271d010234...)  20.00 GiB
 * PURE     FlashArray       (naa.624a9370b76271d010234...)   2.00 TiB
 * PURE     FlashArray       (naa.624a9370b76271d010234...) 180.00 TiB

   (Esc) Cancel     (F1) Details     (F5) Refresh     (Enter) Continue
```

Procedure 3. Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host and connection to vCenter Server. Select the IP address that can communicate with existing or new vCenter Server.

Step 1. After the server has finished rebooting, press F2 to enter in to configuration wizard for ESXi Hypervisor.

Step 2. Log in as root and enter the corresponding password.

Step 3. Select the Configure the Management Network option and press Enter.

Step 4. Select the VLAN (Optional) option and press Enter. Enter the VLAN In-Band management ID and press Enter.

Step 5. From the Configure Management Network menu, select IP Configuration and press Enter.

Step 6. Select the Set Static IP Address and Network Configuration option by using the space bar. Enter the IP address to manage the first ESXi host. Enter the subnet mask for the first ESXi host. Enter the default gateway for the first ESXi host. Press Enter to accept the changes to the IP configuration.

Note: Ipv6 Configuration is set to automatic.

Step 7. Select the DNS Configuration option and press Enter.

Step 8. Enter the IP address of the primary and secondary DNS server. Enter Hostname

Step 9. Enter DNS Suffixes.

Note: Since the IP address is assigned manually, the DNS information must also be entered manually.

Note: The steps provided vary based on the configuration. Please make the necessary changes according to your configuration.

Figure 20. Sample ESXi Configure Management Network

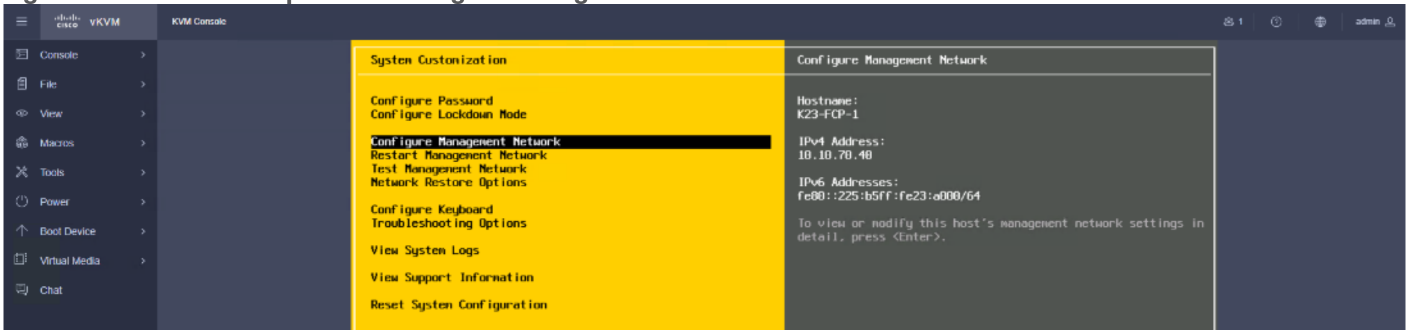

```
System Customization                        Configure Management Network

Configure Password                          Hostname:
Configure Lockdown Mode                     K23-FCP-1

Configure Management Network                IPv4 Address:
Restart Management Network                  10.10.70.40
Test Management Network
Network Restore Options                     IPv6 Addresses:
                                            fe80::225:b5ff:fe23:a000/64
Configure Keyboard
Troubleshooting Options                     To view or modify this host's management network settings in
                                            detail, press <Enter>.
View System Logs

View Support Information

Reset System Configuration
```

Update Cisco VIC Drivers for ESXi

When ESXi is installed from Cisco Custom ISO, you might have to update the Cisco VIC drivers for VMware ESXi Hypervisor to match the current Cisco Hardware and Software Interoperability Matrix. Additionally, Cisco Intersight incorporates an HCL check.

Figure 21. Servers HCL Status in Cisco Intersight Infrastructure Services

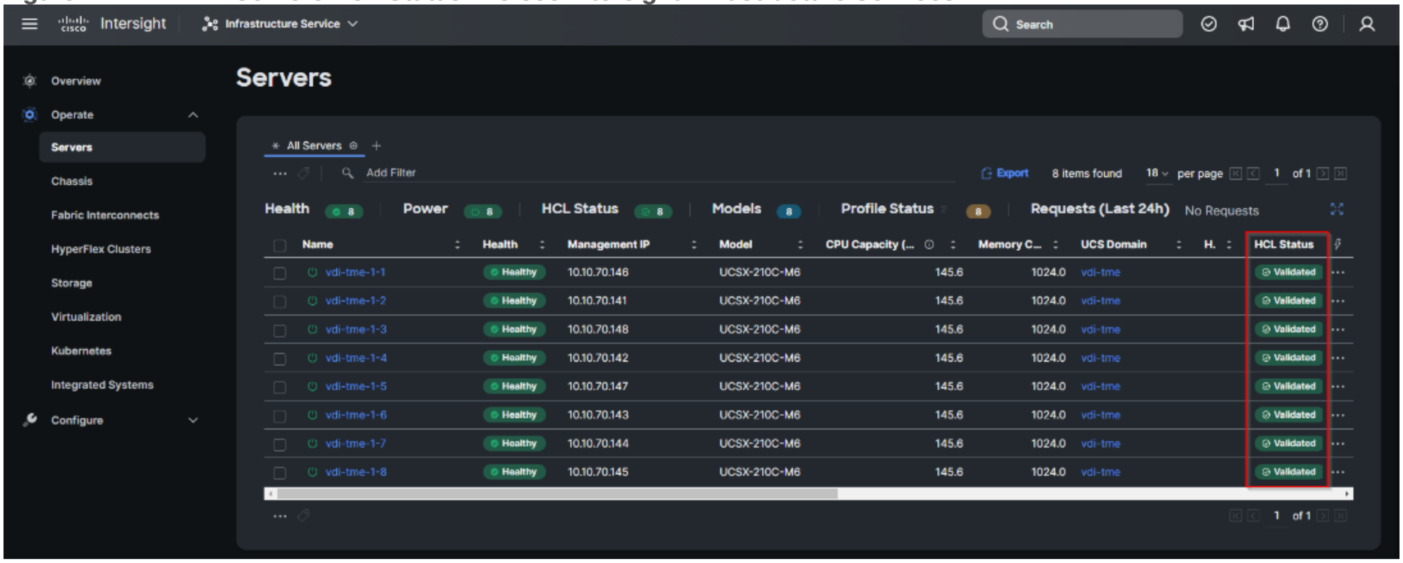

Servers

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status
vdi-tme-1-1	Healthy	10.10.70.146	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-2	Healthy	10.10.70.141	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-3	Healthy	10.10.70.148	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-4	Healthy	10.10.70.142	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-5	Healthy	10.10.70.147	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-6	Healthy	10.10.70.143	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-7	Healthy	10.10.70.144	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-8	Healthy	10.10.70.145	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated

In this Cisco Validated Design, the following drivers were used (VMware-ESXi-7.0.3d-19482537-Custom-Cisco-4.2.2-a):

- Cisco-nenic- 1.0.42.0-1OEM.670.0.0.8169
- Cisco-nfnic- 4.0.0.87-1OEM.670.0.0

Note: For additional information on how to update Cisco VIC drivers on ESXi, refer to the Cisco UCS Virtual Interface Card Drivers for ESX Installation Guide.

VMware Clusters

The VMware vSphere Client was configured to support the solution and testing environment as follows:

- Datacenter: FlashStack – Pure Storage FlashArray//X R3 with Cisco UCS
- Cluster: FlashStack- - EHR_workloads

Figure 22. VMware vSphere WebUI Reporting Cluster Configuration for this Cisco Validated Design

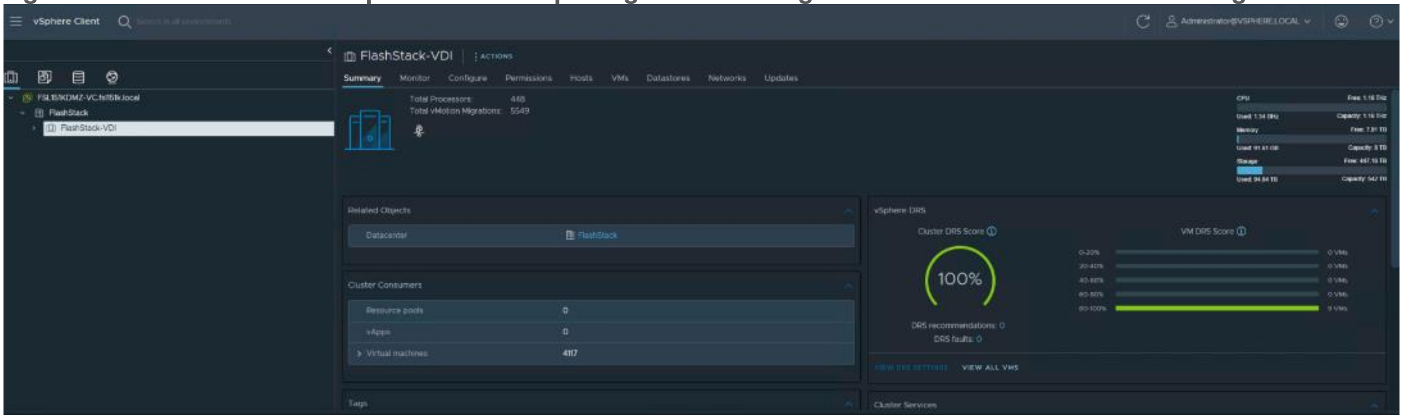

FlashStack-VDI

Total Processors: 448

Total vMotion Migrations: 5549

CPU Free: 1.16 THz; Used: 1.34 GHz; Capacity: 1.16 THz

Memory Free: 7.91 TB; Used: 91.61 GB; Capacity: 8 TB

Storage Free: 447.16 TB; Used: 94.84 TB; Capacity: 542 TB

Datacenter: FlashStack

Resource pools: 0

vApps: 0

Virtual machines: 4117

DRS recommendations: 0

DRS faults: 0

VM DRS Score	
0-20%	0 VMs
20-40%	0 VMs
40-60%	0 VMs
60-80%	0 VMs
80-100%	9 VMs

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 Cisco Intersight Assist into Cisco Intersight, you can claim endpoint devices using the Claim Through Intersight Assist option.

Procedure 1. Configure Cisco Intersight Assist Virtual Appliance

Step 1. To install Cisco 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/1.0.9-342.

Step 2. To set up the DNS entries for the Cisco Intersight Assist hostname as specified under Before you Begin, go 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.

Step 3. From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the FlashStack-Management cluster and click Deploy OVF Template.

Step 4. Specify a URL or browse to the intersight-appliance-installer-vsphere-1.0.9-342.ova file. Click NEXT.

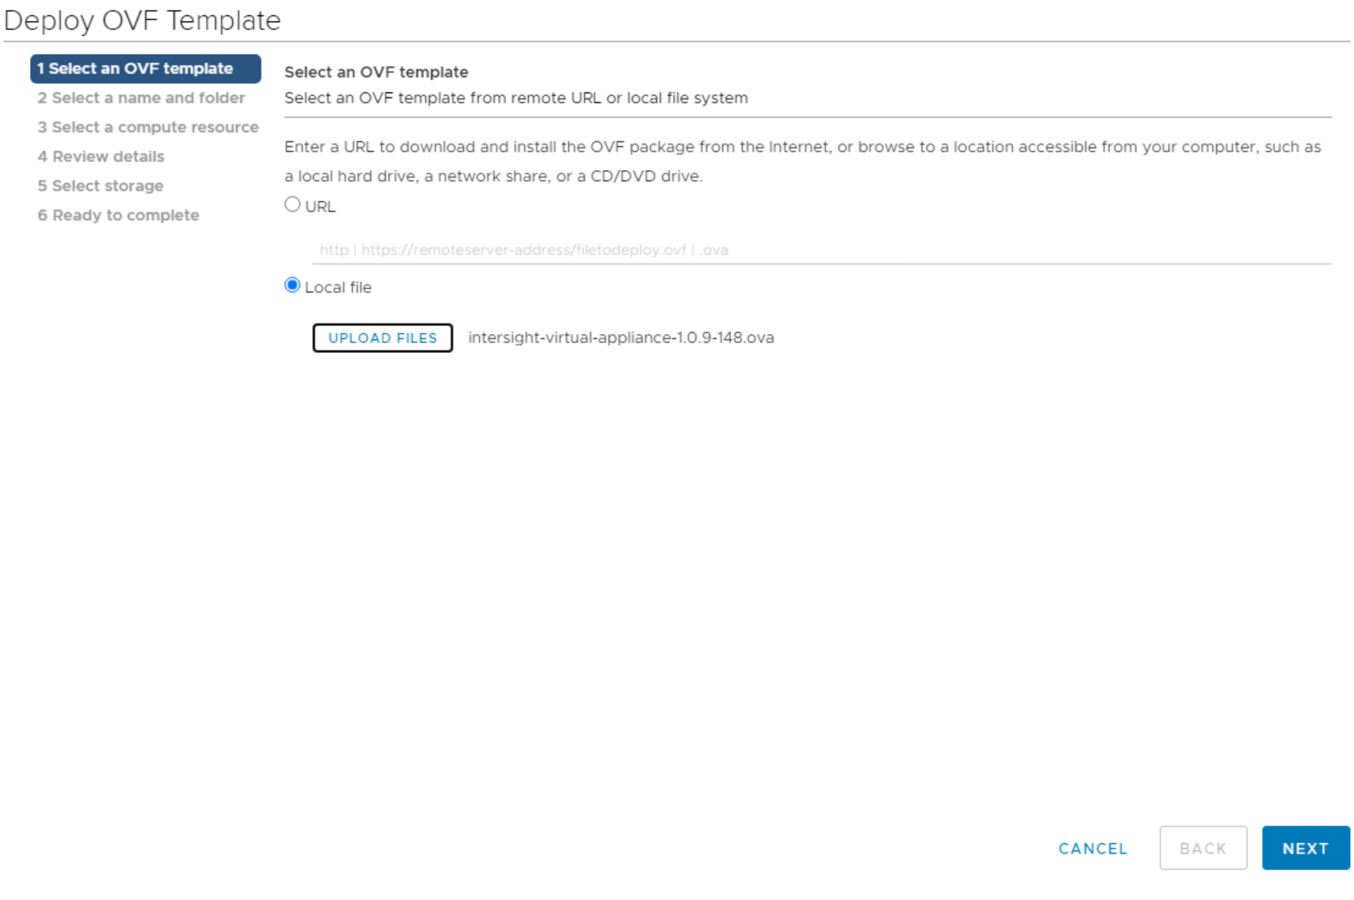

Deploy OVF Template

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

Step 5. Name the Cisco Intersight Assist VM and choose the location. Click NEXT.

Step 6. Select the FlashStack-Management cluster and click NEXT.

Step 7. Review details and click NEXT.

Step 8. Select a deployment configuration (Tiny recommended) and click NEXT.

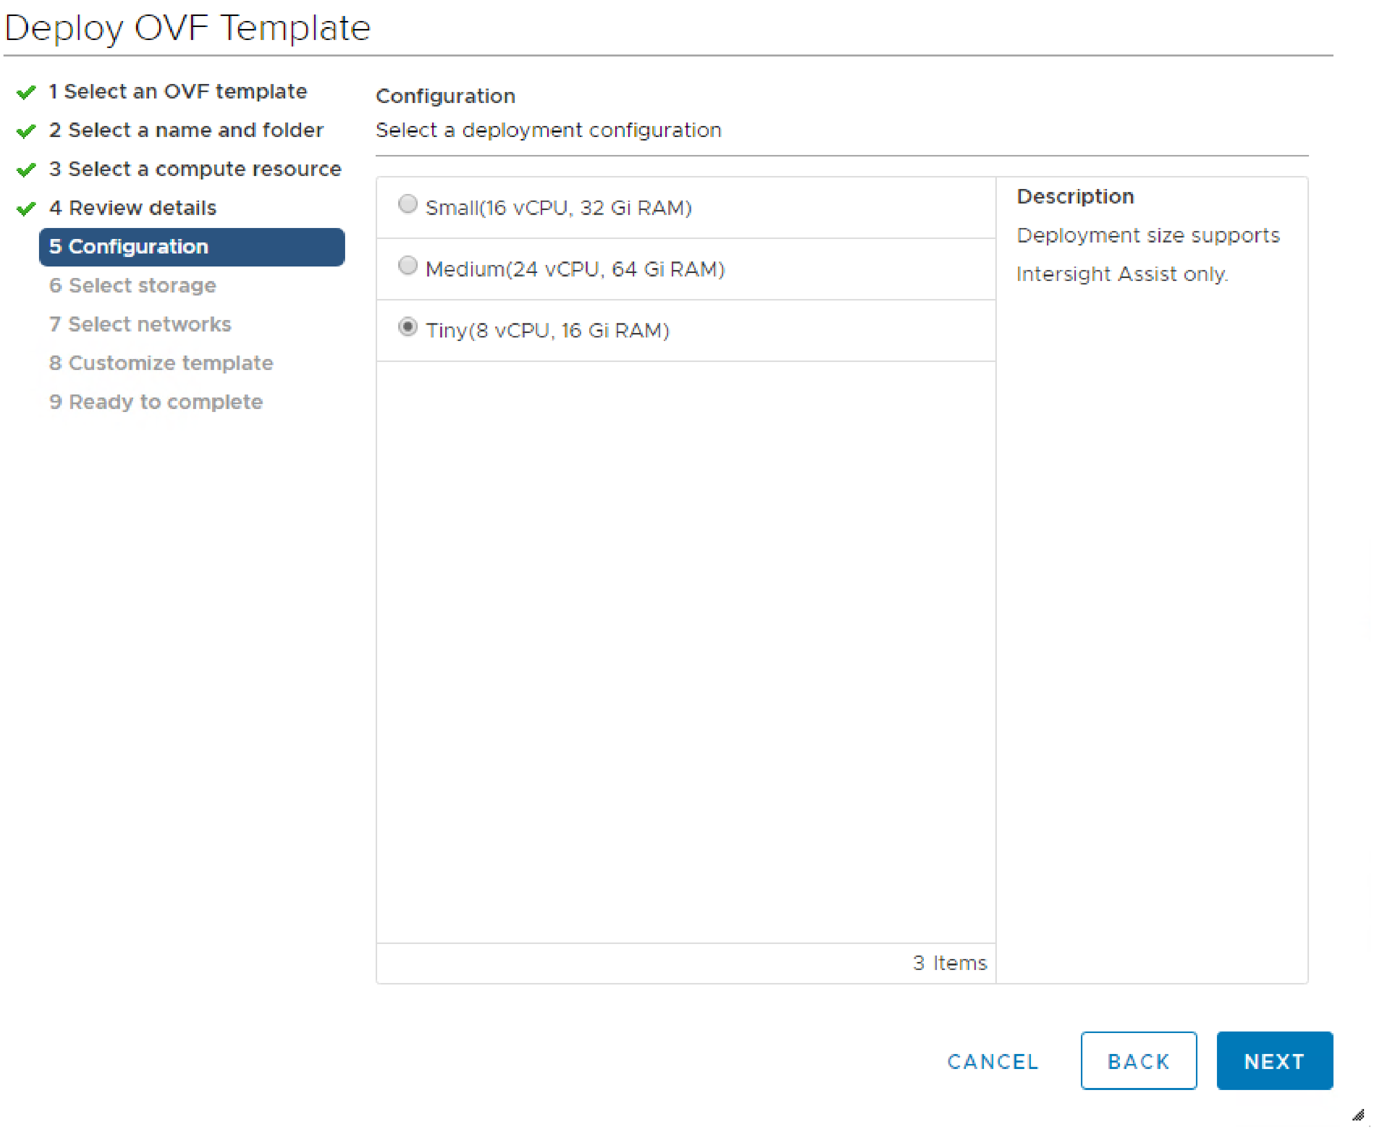

Deploy OVF Template

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.

Step 9. Select the appropriate datastore for storage and select the Thin Provision virtual disk format. Click NEXT.

Step 10. Select IB-MGMT Network for the VM Network. Click NEXT.

Step 11. Fill in all values to customize the template. Click NEXT.

Step 12. Review the deployment information and click FINISH to deploy the appliance.

Step 13. Once the OVA deployment is complete, right-click the Cisco Intersight Assist VM and click Edit Settings.

Step 14. Expand CPU and adjust the Cores per Socket so that 2 Sockets are shown. Click OK.

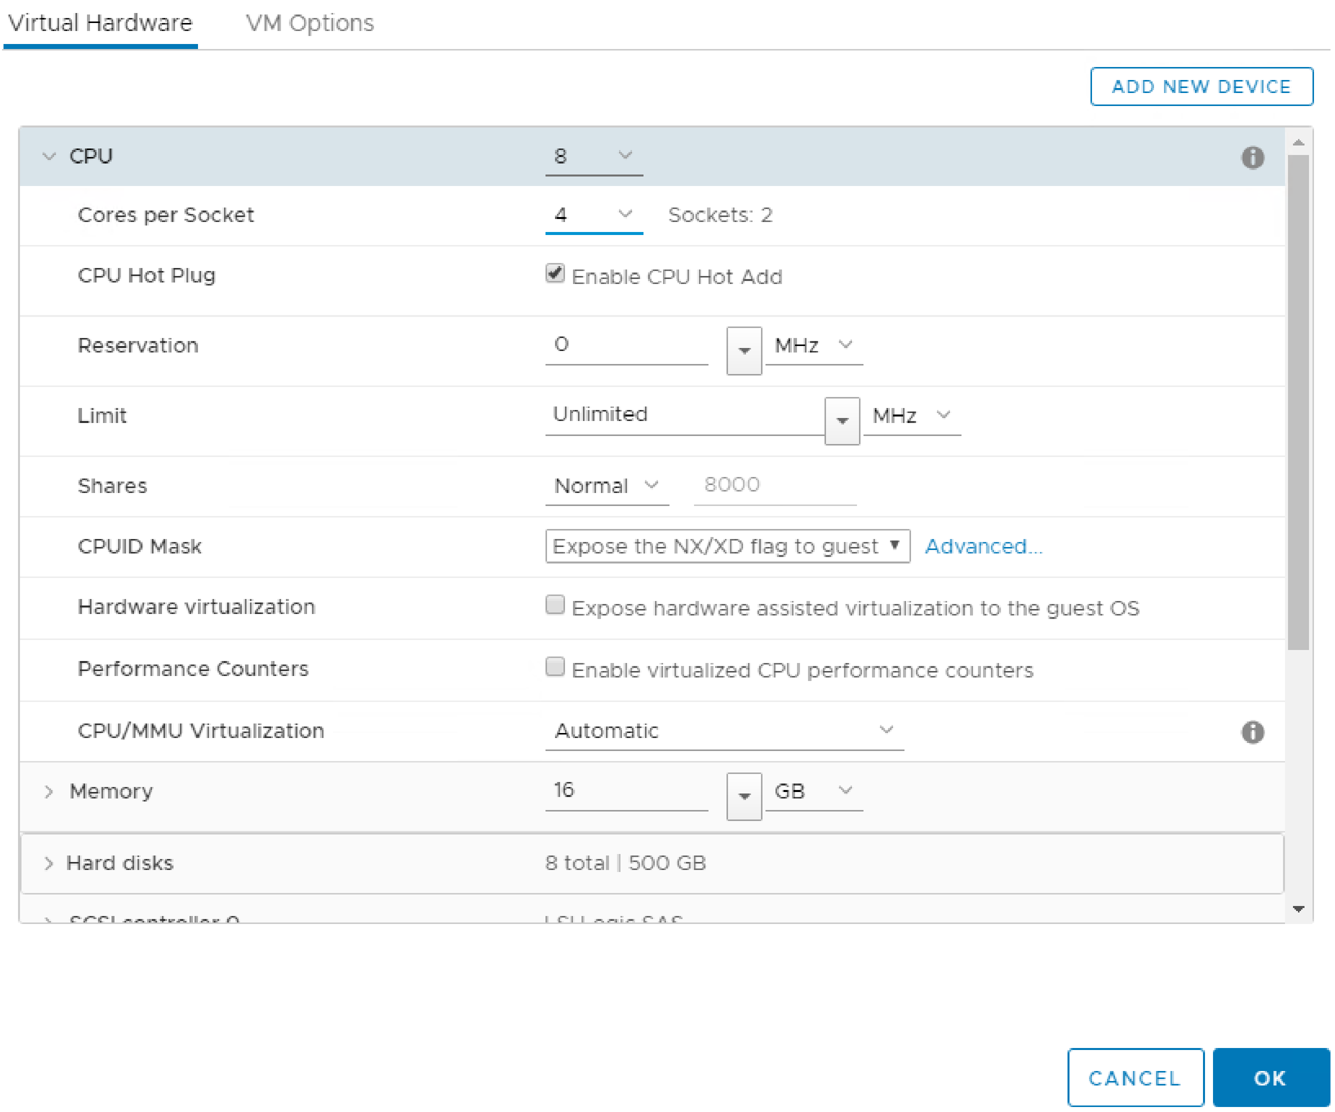

CPU	8
Cores per Socket	4 Sockets: 2
CPU Hot Plug	☑ Enable CPU Hot Add
Reservation	0 MHz
Limit	Unlimited MHz
Shares	Normal 8000
CPUID Mask	Expose the NX/XD flag to guest Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS
Performance Counters	□ Enable virtualized CPU performance counters
CPU/MMU Virtualization	Automatic
Memory	16 GB
Hard disks	8 total | 500 GB
SCSI controller 0	LSI Logic SAS

Step 15. Right-click the Cisco Intersight Assist VM and select Open Remote Console.

Step 16. Power on the VM.

Step 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.

Step 18. Navigate the security prompts and select Intersight Assist. Click Proceed.

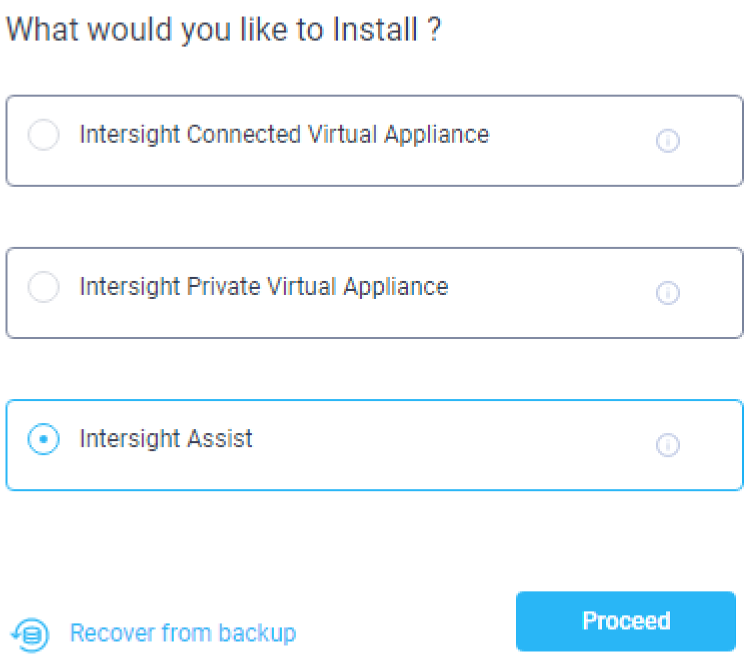

What would you like to Install ?

○ Intersight Connected Virtual Appliance

○ Intersight Private Virtual Appliance

● Intersight Assist

Step 19. From Cisco Intersight, click ADMIN > Devices. Click Claim a New Device. Copy and paste the Device ID and Claim Code shown in the Cisco Intersight Assist web interface to the Cisco Intersight Device Claim Direct Claim window. In Cisco Intersight, click Claim.

Step 20. In the Cisco Intersight Assist web interface, click Continue.

Note: The Cisco Intersight Assist software will now be downloaded and installed into the Cisco Intersight Assist VM. This can take up to an hour to complete.

Note: The Cisco 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.

Step 21. When the software download is complete, navigate the security prompts and a Cisco Intersight Assist login screen will appear. Log into Cisco Intersight Assist with the admin user and the password supplied in the OVA installation. Check the Cisco Intersight Assist status and log out of Intersight Assist.

Procedure 2. Claim Intersight Assist into Cisco Intersight

Step 1. To claim the Intersight assist appliance, from the Service Selector drop-down list, select System.

Step 2. From Cisco Intersight, click ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select Cisco Intersight Assist under Platform Services and click Start.

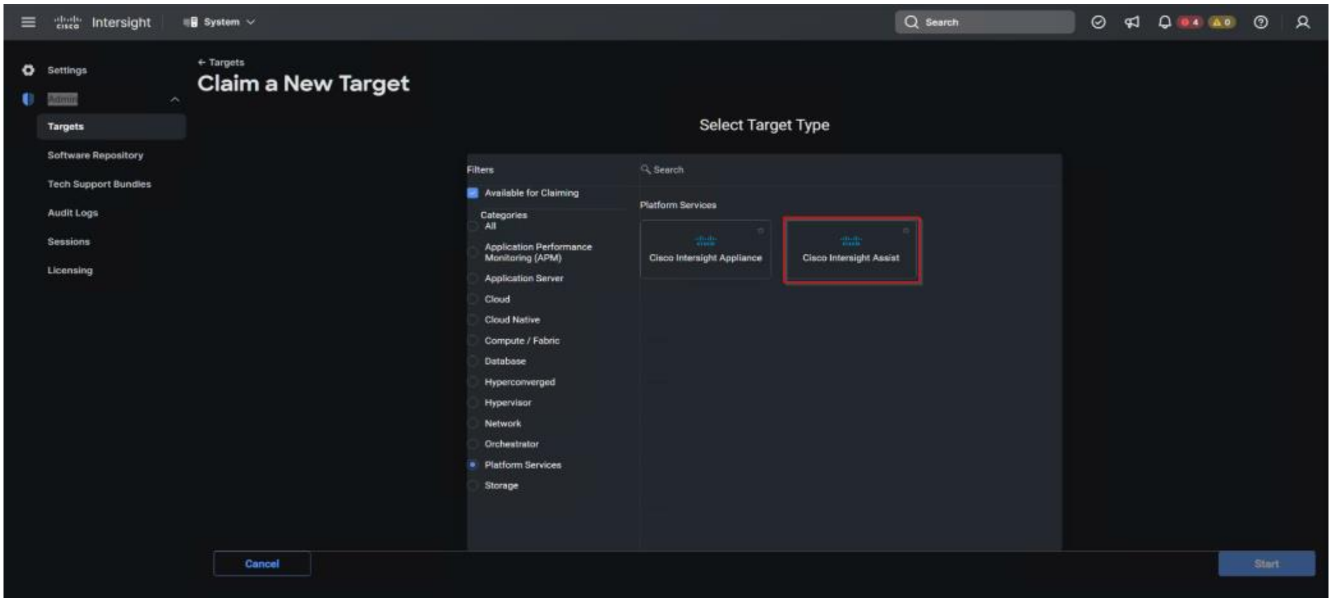

Step 3. Fill in the Intersight Assist information and click Claim.

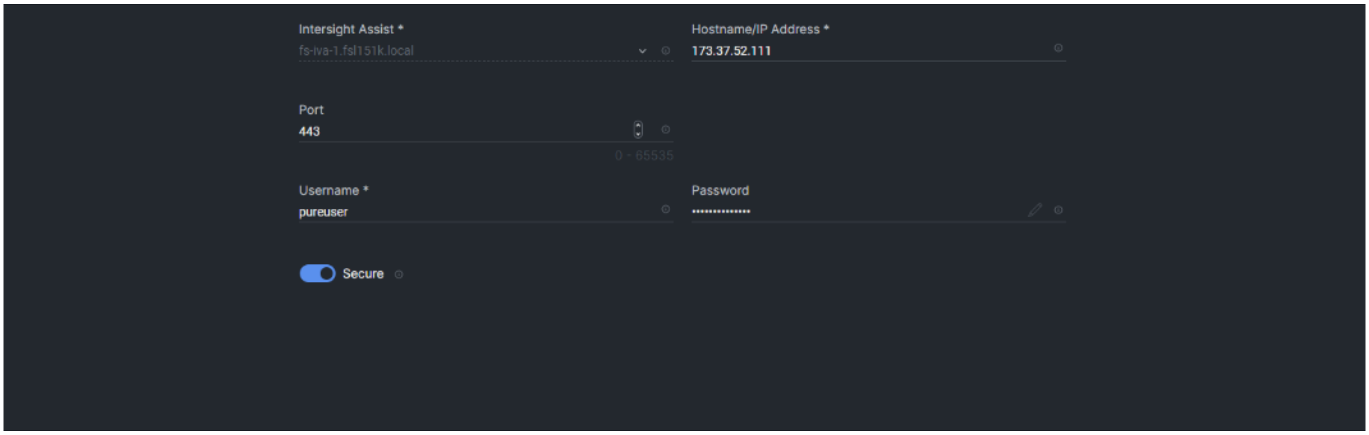

Intersight Assist *
fs-iva-1.fsl151k.local

Hostname/IP Address *
173.37.52.111

Port
443
0 - 65535

Username *
pureuser

Password
••••••••••••••

Secure

After a few minutes, Cisco Intersight Assist will appear in the Targets list.

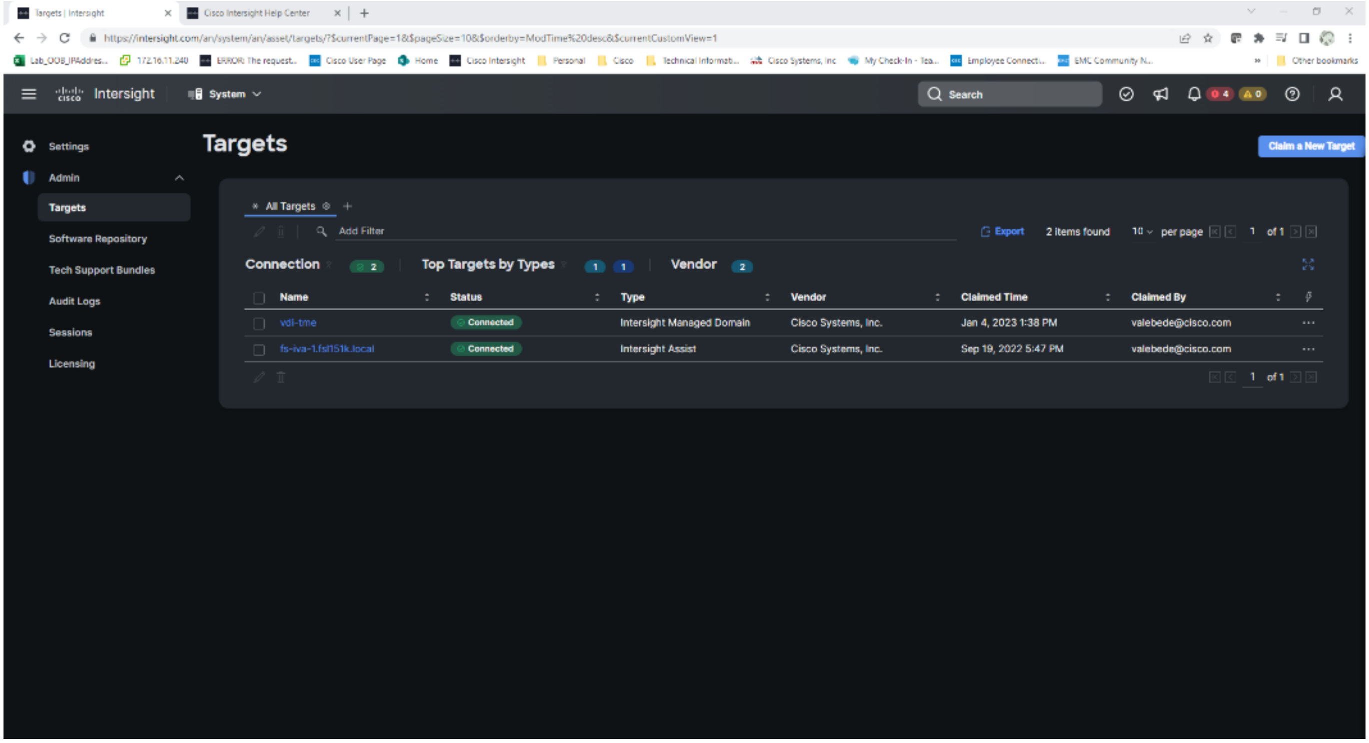

Name	Status	Type	Vendor	Claimed Time	Claimed By
vdi-tme	Connected	Intersight Managed Domain	Cisco Systems, Inc.	Jan 4, 2023 1:38 PM	valebede@cisco.com
fs-iva-1.fsl151k.local	Connected	Intersight Assist	Cisco Systems, Inc.	Sep 19, 2022 5:47 PM	valebede@cisco.com

Procedure 3. Claim vCenter in Cisco Intersight

Step 1. 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.

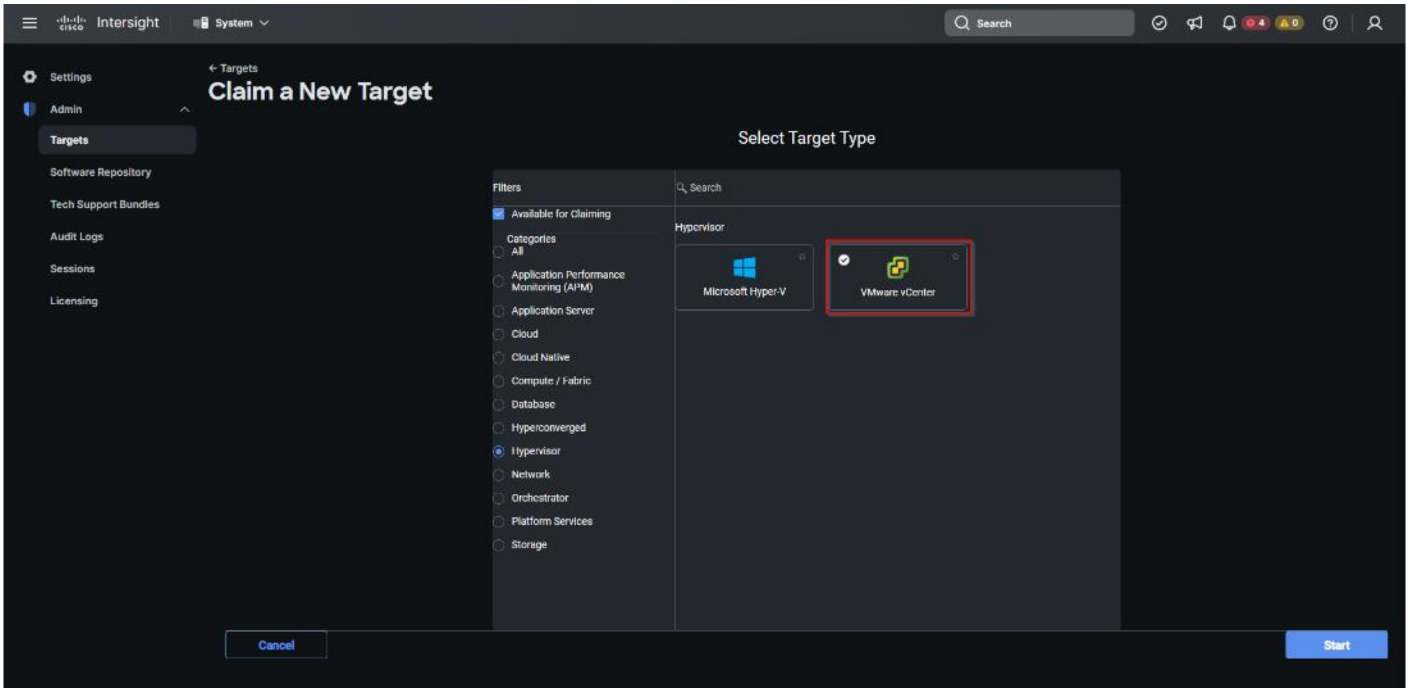

Step 2. In the VMware vCenter window, make sure the Intersight Assist is correctly selected, fill in the vCenter information, and click Claim.

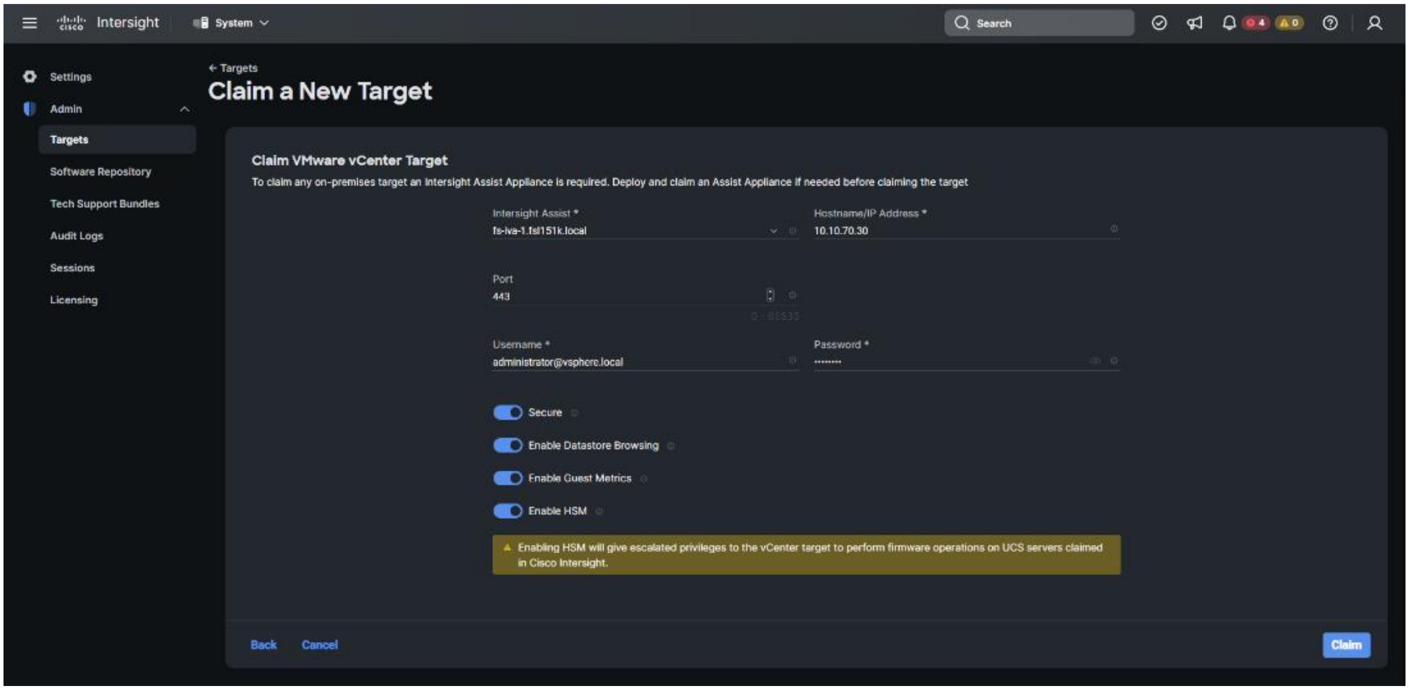

Claim a New Target

Claim VMware vCenter Target

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 * fs-iva-1.fsl151k.local

Hostname/IP Address * 10.10.70.30

Port 443

Username * administrator@vsphere.local

Password * ••••••••

Secure

Enable Datastore Browsing

Enable Guest Metrics

Enable HSM

Enabling HSM will give escalated privileges to the vCenter target to perform firmware operations on UCS servers claimed in Cisco Intersight.

Step 3. 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.

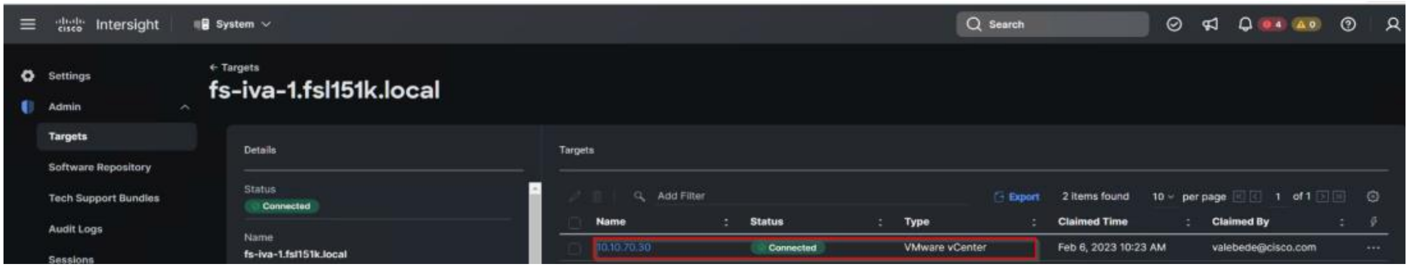

fs-iva-1.fsl151k.local

Details

Status Connected

Name fs-iva-1.fsl151k.local

Targets

2 items found 10 per page 1 of 1

Name	Status	Type	Claimed Time	Claimed By
10.10.70.30	Connected	VMware vCenter	Feb 6, 2023 10:23 AM	valebede@cisco.com

Step 4. Detailed information obtained from the vCenter can now be viewed by clicking Virtualization from the Infrastructure service > Operate menu.

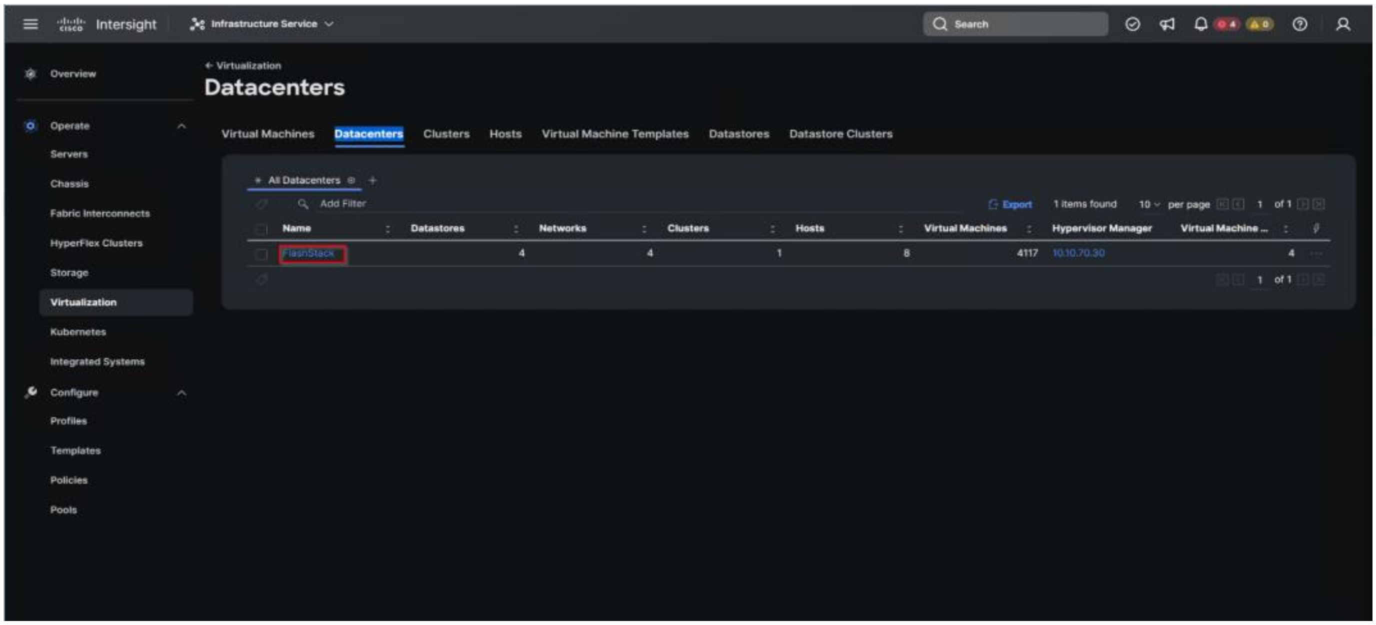

Datacenters

1 items found 10 per page 1 of 1

Name	Datastores	Networks	Clusters	Hosts	Virtual Machines	Hypervisor Manager	Virtual Machine ...
FlashStack	4	4	1	8	4117	10.10.70.30	4

Procedure 4. Claim FlashArray//X in Cisco Intersight

Note: Claiming a Pure Storage FlashArray also requires the use of an Intersight Assist virtual machine.

Step 1. To claim the vCenter, from Cisco Intersight, click ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select Pure Storage FlashArray under Storage and click Start.

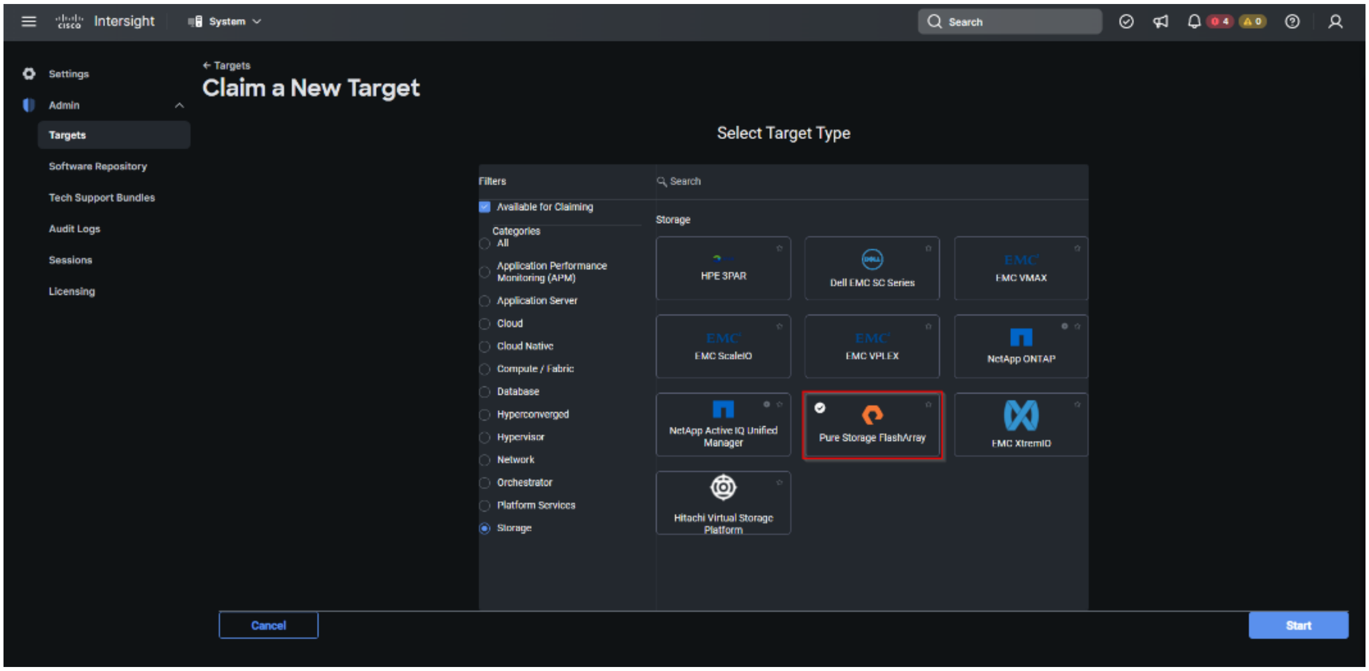

Claim a New Target

Select Target Type

Filters

☑ Available for Claiming

Categories

● Storage

Storage: ✓ Pure Storage FlashArray

Step 2. Enter the Pure Storage FlashArray Hostname/IP address and credentials. Click Claim.

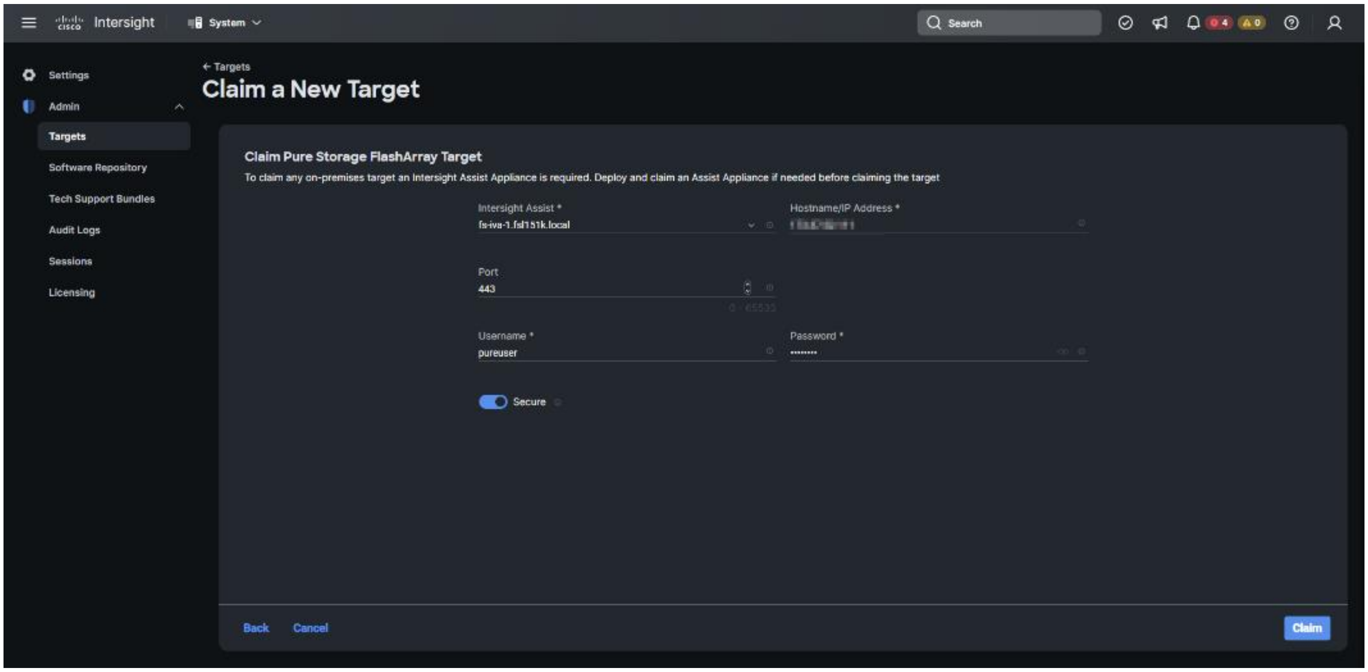

Claim a New Target

Claim Pure Storage FlashArray Target

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 * fs-iva-1.fsl151k.local

Hostname/IP Address * [redacted]

Port 443

0 - 65535

Username * pureuser

Password * ••••••••

Secure: on

Supported LINUX Hosts for Operations Database (ODB)

Currently, there are only two supported LINUX Operating Systems for use as the ODB hosts; RedHat Enterprise and Ubuntu.

Data Protections Options

The following data protection options apply to this solution:

- Pure Storage provides customers at no additional cost, an easy yet powerful set of data protection tools to copy and move data between arrays for several data protection and utility situations. In nearly all cases, an asynchronous replication methodology is preferred over synchronous to ensure data integrity and movement performance is not hampered by latency challenges.
- The best practice to leverage Pure Protection Groups (PG) to act as a logical container for volumes related to the database (cache.dat files), the journals/WIJ, and application files specific to the instance. Using this container provides the value of requiring only a single pure CLI or REST call enabling a single action to be performed on multiple volumes exactly at the same time simplifying your scripting process. This will give us a crash-consistent copy of the database instance. By wrapping this process around the IRIS freeze and thaw scripts, this will give you an application-consistent copy of the database instance.
- One of the most important processes to understand to ensure you are working with an application-consistent copy of data is how to integrate instance Freeze (instfreeze) and Thaw (instthaw) within a snapshotting workflow. Epic provides these scripts to customers at no cost and are critical in creating an application consistent copy of data for an Epic instance. Because of this there's a general workflow that is universal for data preparation for Epic.

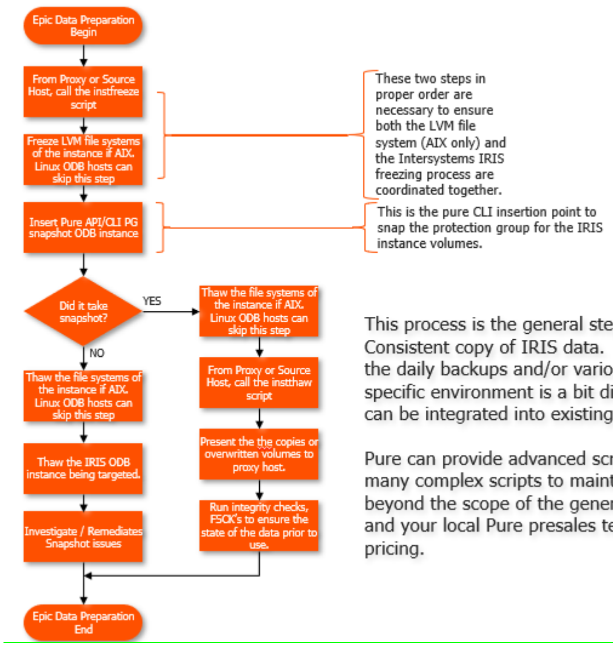
Epic Data Preparation Begin
From Proxy or Source Host, call the instfreeze script
Freeze LVM file systems of the instance if AIX. Linux ODB hosts can skip this step
Insert Pure API/CLI PG snapshot ODB instance
Did it take snapshot?
YES
NO
Thaw the file systems of the instance if AIX. Linux ODB hosts can skip this step
From Proxy or Source Host, call the instthaw script
Present the the copies or overwritten volumes to proxy host.
Run integrity checks, FSCK's to ensure the state of the data prior to use.
Thaw the file systems of the instance if AIX. Linux ODB hosts can skip this step
Thaw the IRIS ODB instance being targeted.
Investigate / Remediates Snapshot issues
Epic Data Preparation End

These two steps in proper order are necessary to ensure both the LVM file system (AIX only) and the Intersystems IRIS freezing process are coordinated together.

This is the pure CLI insertion point to snap the protection group for the IRIS instance volumes.

This process is the general steps needed to create an Application Consistent copy of IRIS data. This workflow is general used for making the daily backups and/or various refresh points. Every Epic customer's specific environment is a bit different but following this general workflow can be integrated into existing workflows easily.

Pure can provide advanced scripting services in the event there are many complex scripts to maintain, or the scripting needs are advanced beyond the scope of the general workflow. This is a paid engagement, and your local Pure presales team can assist you with scoping and pricing.

- As Epic customers, you also will have tools provided by Epic and InterSystems to help provide data movement capabilities using IRIS mirroring and shadowing as well as instance refreshing scripts. Epic requires customers to utilize IRIS mirroring typically for the production (PRD) instance to its DR counterpart (DR-PRD) typically in another data center. Along with PRD mirroring, it's common to see an IRIS mirror from the production (PRD) instance to the IRIS reporting (RPT) instance as well. It's a best practice to separate the PRD instance from the RPT instance for HA and potential performance reasons. IRIS mirroring is typically configured to run asynchronously to prevent any potential lag in replication from busy network connections.
- Pure Storage customers can use any of Pure Storage's data protection technologies alongside Epic's recommended tools. It's critical for customers to follow their guidance related to business. Beyond that, Pure Storage customers can also leverage Pure Storage tools to provide added replication capability based on specific customer needs. This would be in the areas of the services-tier of the application for servers like Interconnect Services, WBS, Web Servers for MyChart, EPS, BCA, and other proxy hosts.

While these servers tend to be smaller, their protection should be considered as critical as any other server in the infrastructure. Work with both your Epic Server Systems team along with your local Pure Storage account team to ensure you are getting the most out of these tools.

- Customers who are using Pure Storage's SafeMode feature to protect their array resources from ransomware and malicious deletions and encryption often will leverage on-array snapshots as a first point of recoverability.

Safe Mode Considerations

The following Pure Storage Safe Mode considerations apply to this solution:

- SafeMode is Pure Storage's answer to assist customers who want to gain an extra layer of ransomware and malicious attack mitigation.
- With SafeMode, the workflow for staging copies of data using the instfreeze and instthaw does not change. Also, the best practice of leveraging Protection Groups that contain IRIS database files, Journals, and app files is still valid when using SafeMode.
- A key element of the introduction of SafeMode for an healthcare customer is to engage Epic support once the volumes have been returned to an online/attached state on the host(s). It's critical that customers contact Epic prior to a restart of the IRIS database engine as there may be additional forensics needed to be run. This would include IRIS database integrity checks and other specific tools deemed necessary based on the situation.
- Pure Storage recommended no fewer than two trusted customer names be identified when a SafeMode call to Pure Support is made. The best practice related to trusted SafeMode users should be up to six people identified. These people should be a mix of technical and leadership people. For example, one or two storage admins, the Director/Lead for the Technology Team, a Corporate Compliance/Security Officer, and finally a DBA and/or Computer Operator. This level of diversity should provide the necessary trusted access when needed. Each of these individuals will receive their own unique PIN code generated by Pure Storage and the code should not be shared with anyone.

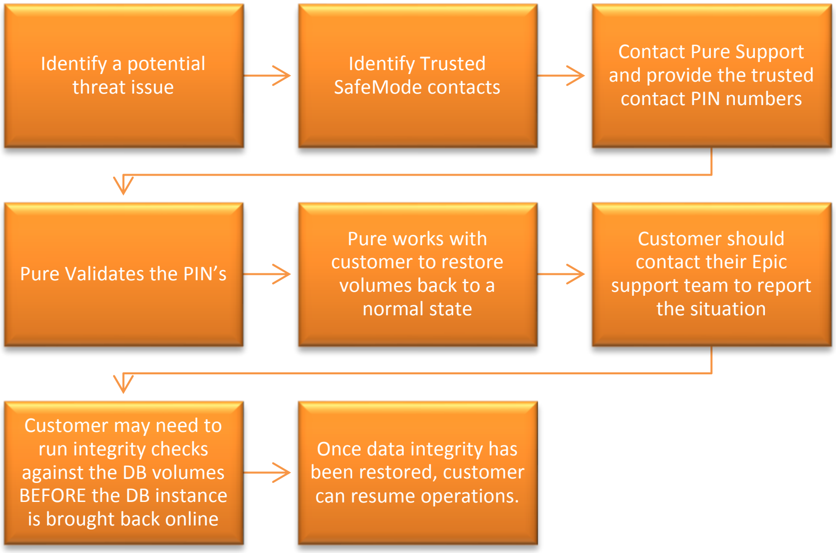
Identify a potential threat issue
Identify Trusted SafeMode contacts
Contact Pure Support and provide the trusted contact PIN numbers
Pure Validates the PIN's
Pure works with customer to restore volumes back to a normal state
Customer should contact their Epic support team to report the situation
Customer may need to run integrity checks against the DB volumes BEFORE the DB instance is brought back online
Once data integrity has been restored, customer can resume operations.

- When establishing the SafeMode eradication policy on initial setup, it's important to think through how snapshots are eradicated today. *The initial best-practice is five to seven days*. A good rule-of-thumb is to use the retention timeframe of your IRIS journaling reclamation. Customers than modify this value to a higher number once it's initially established but will not be able to change the SafeMode policy that's lower than the configured amount of time. If it’s necessary to change the value lower, the customer should contact Pure Support with their PIN authentication to assist with the modification.

- You need to be mindful that the longer the duration of their SafeMode policy, can impact their effective storage capacity. It’s best to start out with a lower SafeMode policy retention value to observe the impact. If capacity is not an issue, the customer can reach out to Pure Storage support for adjusting the value. In upcoming releases, this will be something you will be able to change within the Purity interface.

Purity Code Updating for Healthcare Customers

For healthcare customers, Pure Storage requires a longer code maturity timeframe than the normal Purity release cycle. This is done to assure there’s been ample run-time for new code releases with the goal of assuring stability. Pure Storage manages code releases eligible for healthcare customers to implement within the engineering group. The criteria we use for maturing the Purity code for healthcare customers is the following:

- Purity Releases must be released for at a minimum of 45 days.
- Purity Releases must have at least 10,000 total array days.
- Purity releases should have a minimum number of events per array day. These should be no more than 1 event seen per 3000 array days in the field for high confidence.
- Code eligibility questions should be directed to your local Pure Storage account team for the proper upgrade guidance. In the event there is a significant issue or feature release, Pure Storage may allow for a more aggressive upgrade release timeline than the standard criteria above.
- If there’s a code level this is not reflected in the approved release table, please reach out to your local account team for clarification and potential release information before any action is taken to apply new releases.

For future Purity code upgrades, Pure Storage has an auto-responder that will push the latest qualified Purity code release.

How to get the latest Purity code release:

1. Open an email and direct it to epic@purestorage.com.
2. In the Subject line, enter “HC Code Request” or “HC Purity Request.”
3. You do not need to put anything in the body of the email.
4. Send the email and in a few minutes, you’ll get a response. – below is an example of the message:

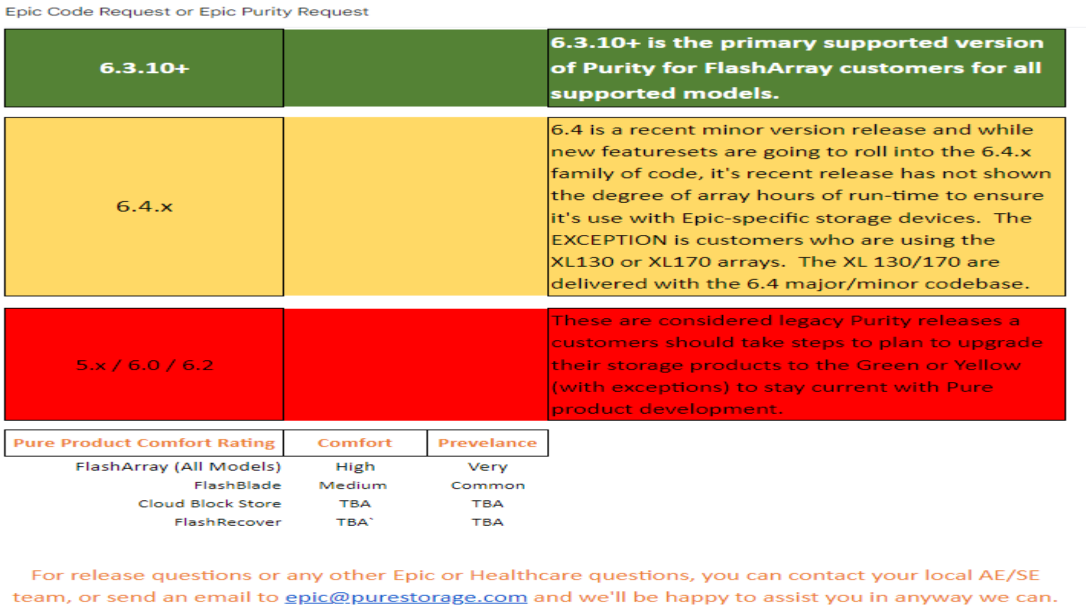

Epic Code Request or Epic Purity Request

6.3.10+		6.3.10+ is the primary supported version of Purity for FlashArray customers for all supported models.
6.4.x		6.4 is a recent minor version release and while new featuresets are going to roll into the 6.4.x family of code, it's recent release has not shown the degree of array hours of run-time to ensure it's use with Epic-specific storage devices. The EXCEPTION is customers who are using the XL130 or XL170 arrays. The XL 130/170 are delivered with the 6.4 major/minor codebase.
5.x / 6.0 / 6.2		These are considered legacy Purity releases a customers should take steps to plan to upgrade their storage products to the Green or Yellow (with exceptions) to stay current with Pure product development.

Pure Product Comfort Rating	Comfort	Prevelance
FlashArray (All Models)	High	Very
FlashBlade	Medium	Common
Cloud Block Store	TBA	TBA
FlashRecover	TBA`	TBA

For release questions or any other Epic or Healthcare questions, you can contact your local AE/SE team, or send an email to epic@purestorage.com and we'll be happy to assist you in anyway we can.

Generate IO or more commonly referred to as "GENIO" is an Epic-developed performance assessment tool that helps customers who are needing to test a storage solution to determine it's performance. The test is made up of two primary executables: RampRun.pl and dgen.pl. The tool uses generated IRIS dummy files that are created to mimic the files that you would normally have in a real Epic implementation. The following will give you the direct needed to setup the tool in a customer's environment based on the specific GENIO requirements and then executed.

IMPORTANT! DO NOT RUN dgen.pl AGAINST A REAL IRIS DATABASE! IT WILL DESTROY THE IRIS INSTANCE.

The next sections will break down the configuration requirements for both dgen.pl and RampRun.pl, as well as how to properly setup the ODB host. GENIO is a tool created and distributed by Epic and managed through the Server Systems team. Before you begin to setup the GENIO test, it's a good idea to check with your Epic Server Systems Team to ensure you have the latest version of the testing software. The documentation packaged within the GenerateIO bundle will always supersede this documentation.

To run GENIO, you need a host that's running either LINUX or AIX. For this section, the LINUX option is what will be reflected here but the same general guidelines will apply for AIX as well.

It's important to know about the GENIO test is the success in using it is to also understand the resources it needs in order to provide the best results. This includes a VMware-based VM running on the FlashStack compute running LINUX, SAN connectivity from the storage array to the host, and finally presenting resources from FlashArray to execute the test. Lastly, we will assume the LINUX or AIX ODB hosts are setup with the optimum storage design with multipathing, failover, and so on. Refer to the Epic documentation to create an ODB Linux Server guidelines and well as AIX. Pure Storage also has some additional materials for consideration.

Now you can FTP the GenerateIO tar file to the ODB host. It's a good idea to put in a place off the root volume group. You do not want to run it from the shared storage you are using for the test itself.

Once you have the host created and all the various pathing and communications all operating correctly, you can move in to create your fake IRIS ODB.

Epic's dgen.pl is the first step to perform the GENIO testing. Before you can run dgen.pl you must create a fake LVM structure for the dgen.pl to operate against. While this is a test, it's a good idea to use this opportunity to lock in your LVM settings. This will help make sure the test runs correctly but also can help troubleshoot any Host / Storage needs. There are two ways to create a file system layout for dgen.pl to operate against. A twelve directory build or a four directory. Why? In today's NVMe storage systems, the days of creating hundreds of devices to get optimal performance are done. The twelve directory option is a more explicit way of carving up the files, but going with the four directory option saves a lot of time with the same results.

The best option for a customer is up to the storage admin's needs but it's important to contrast the differences. In today's storage performance, the three directories work just as well as larger options.

Datagen or dgen.pl is a tool that takes inputs around the LVM structure above and uses it to create the files. The following is the execution parameters to kick off the dgen.pl.

As root, you enter the following command string to launch the dgen.pl test. You will notice that there are no two-digit numbers at the end of the /epic/prd/. When you build your file system, you include the 01-08 based on your option. The dgen.pl tool will assume there will be a two-digit number after /epic/prd[01]. IF this is not in the file system, the script will error out.

Once completed, in the installed location of the GenerateIO tool, a 'files' file will be created and is used as input for RampRun.pl.

Now that the faux Epic database is completed, you can begin to run the actual test using the RampRun.pl. RampRun.pl is a wrapper script of several "forked" scripts that execute in an order of operation that will simulate the working of an InterSystems IRIS database. This script, unlike dgen.pl, has a large number of switches that allow the script to run in a variety of ways and in turn, can influence how RampRun.pl achieves its results. The focus will be only on the relevant switches based on Pure Storage and Epic testing.

RampRun.pl has a very distinct IO profile that is easy to see when you monitor the storage array as it runs.

The Generate IO Performance Profile

A combination of 8K reads / 8K writes that steps or "Ramps" incremental performance based on the parameters passed to the command line.

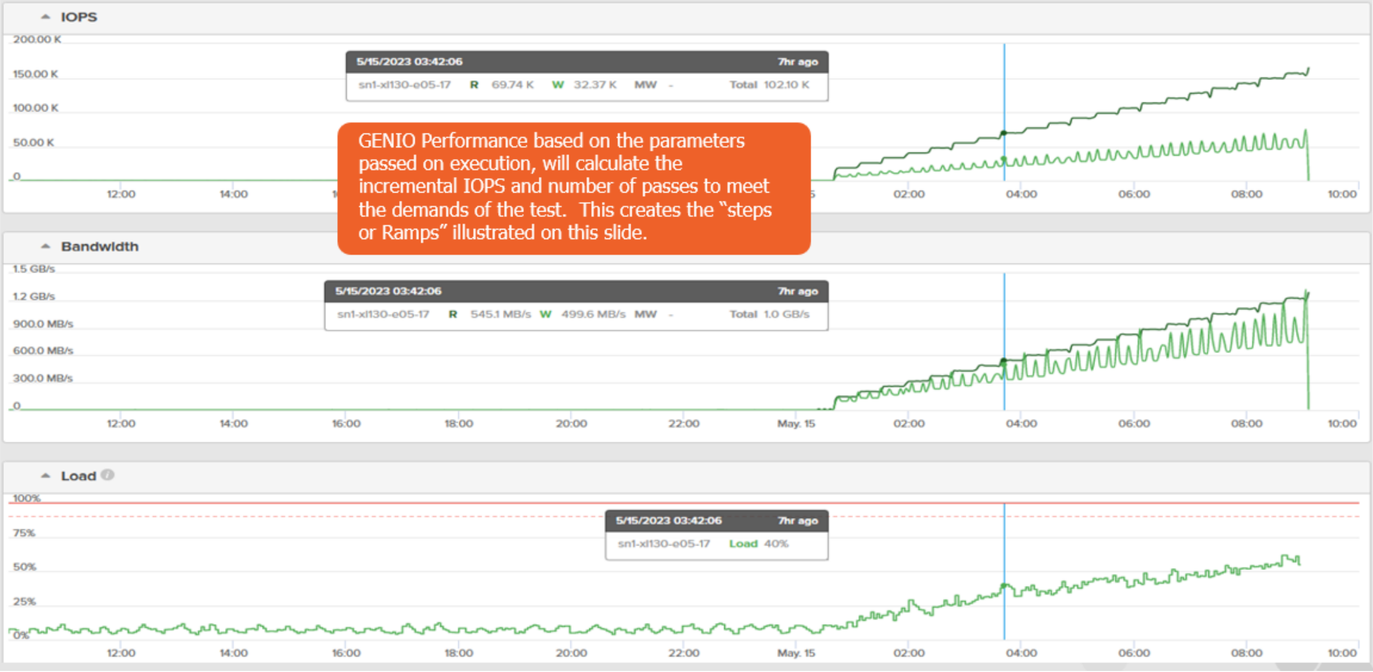

5/15/2023 03:42:06 7hr ago
sn1-xl130-e05-17 R 69.74 K W 32.37 K MW - Total 102.10 K

GENIO Performance based on the parameters passed on execution, will calculate the incremental IOPS and number of passes to meet the demands of the test. This creates the "steps or Ramps" illustrated on this slide.

5/15/2023 03:42:06 7hr ago
sn1-xl130-e05-17 R 545.1 MB/s W 499.6 MB/s MW - Total 1.0 GB/s

5/15/2023 03:42:06 7hr ago
sn1-xl130-e05-17 Load 40%

Remember that RampRun.pl is not just a storage performance tool, but also is influenced by the server running RampRun.pl as well.

The key outcome measure is the <45 second write flush. When a test is completed, a directory is created that has the name based on the name you supplied in the –comment switch. Inside that directory, you are going to look for the summary file in it's raw .csv format. This file is easily imported into excel or any tool that can read a .csv file. For the examples shown here, MS Excel is used.

Example of the output and its interpretation:

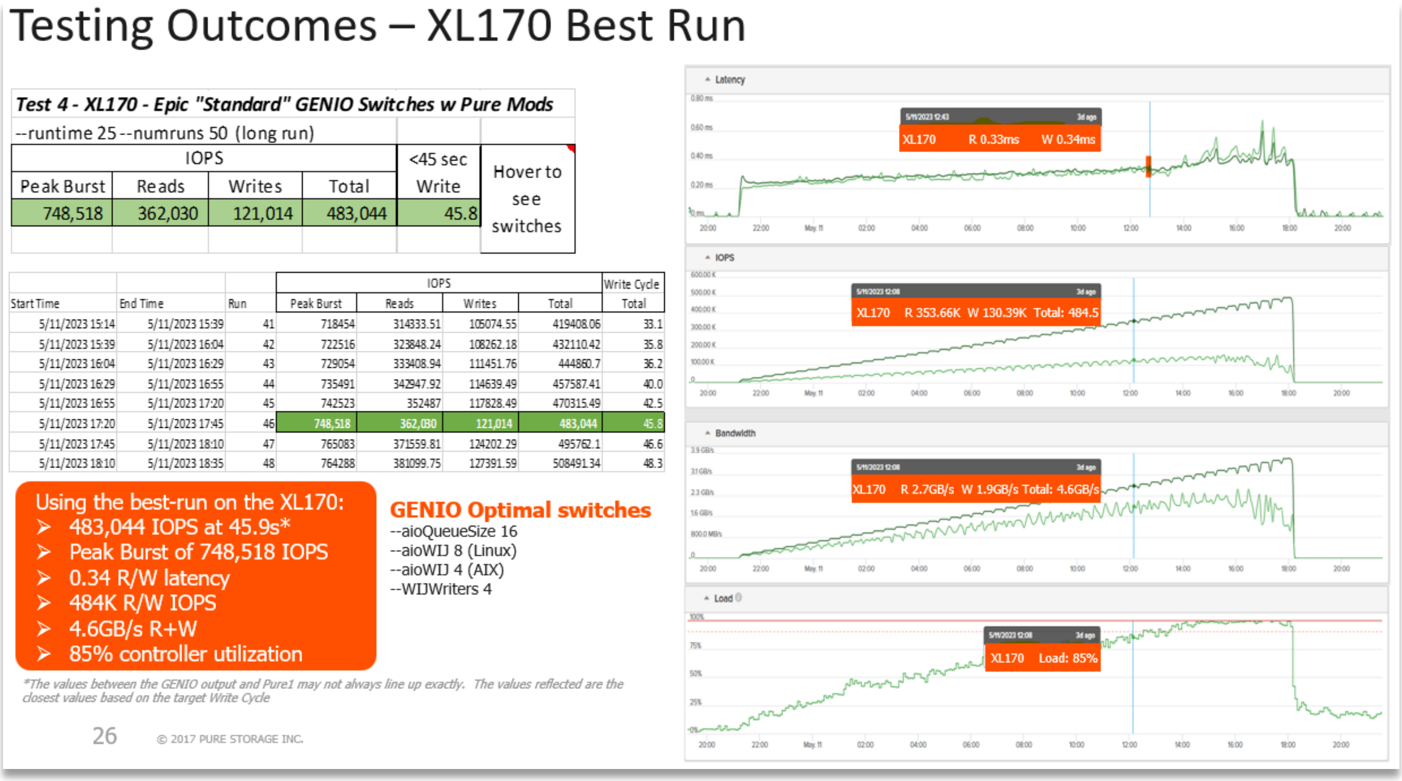

Testing Outcomes – XL170 Best Run

Test 4 - XL170 - Epic "Standard" GENIO Switches w Pure Mods					
--runtime 25 --numruns 50 (long run)					
IOPS				<45 sec	Hover to see switches
Peak Burst	Reads	Writes	Total	Write	
748,518	362,030	121,014	483,044	45.8	

			IOPS				Write Cycle
Start Time	End Time	Run	Peak Burst	Reads	Writes	Total	Total
5/11/2023 15:14	5/11/2023 15:39	41	718454	314333.51	105074.55	419408.06	33.1
5/11/2023 15:39	5/11/2023 16:04	42	722516	323848.24	108262.18	432110.42	35.8
5/11/2023 16:04	5/11/2023 16:29	43	729054	333408.94	111451.76	444860.7	36.2
5/11/2023 16:29	5/11/2023 16:55	44	735491	342947.92	114639.49	457587.41	40.0
5/11/2023 16:55	5/11/2023 17:20	45	742523	352487	117828.49	470315.49	42.5
5/11/2023 17:20	5/11/2023 17:45	46	748,518	362,030	121,014	483,044	45.8
5/11/2023 17:45	5/11/2023 18:10	47	765083	371559.81	124202.29	495762.1	46.6
5/11/2023 18:10	5/11/2023 18:35	48	764288	381099.75	127391.59	508491.34	48.3

Using the best-run on the XL170:

- 483,044 IOPS at 45.9s*
- Peak Burst of 748,518 IOPS
- 0.34 R/W latency
- 484K R/W IOPS
- 4.6GB/s R+W
- 85% controller utilization

GENIO Optimal switches

```
--aioQueueSize 16
--aioWIJ 8 (Linux)
--aioWIJ 4 (AIX)
--WIJWriters 4
```

**The values between the GENIO output and Pure1 may not always line up exactly. The values reflected are the closest values based on the target Write Cycle*

26 © 2017 PURE STORAGE INC.

The key area is the Write Cycle Total column. A valid benchmark is a value of < 45.9 seconds. At this point is where you assess the performance captured at that moment in the test. The charts on the right are visual graphs from the Pure1 for the array being tested. By using the results from the GENIO output and using the timelines on the graph, you can see how the array performed up to the 45 second point.

The evaluation of the above data is that when GENIO hit the 45 second benchmark, the array's overall Reads vs Writes were 362K Reads and 121K Writes or a total of 483K IOPS. GENIO will also predict a Peak Burst in IOPS of 748K based on the run outcomes. The array's controller pressure at this point was roughly 85 percent. GENIO will continue to test the array past the 45 second benchmark to satisfy the –runtime and –numruns values used when you submitted the run.

Based on the run, the Pure Storage FlashArray measured is capable of executing workloads of up to 483K IOPS at 45 seconds with a .34ms latency with controller utilization at 85 percent. It's important to remember that GENIO is also running against a server and if there's server resource concerns, it will affect GENIO's outcomes. For advanced troubleshooting, Cisco , Pure Storage, and Epic can assist with determining results.

Pure Storage FlashArray Product Family GENIO Outcomes

The testing of GENIO against the XL platform showed a dramatic uplift in IOPS and throughput while demonstrating low latency and reduced controller saturation at the Epic <45sec benchmark

When considering our deployment approach (see below) we feel the FlashArray product family gives customers more upgrade options and a clear path from Xr3 and Xr4 to the XL family. All supporting NDU controller upgrades and NVMe

These results are based on our internal GENIO testing and while every customer's environment is different, these results are a good representation of what Pure can deliver to customers.

Flash Array Platform GENIO-Tested Performance Benchmarks

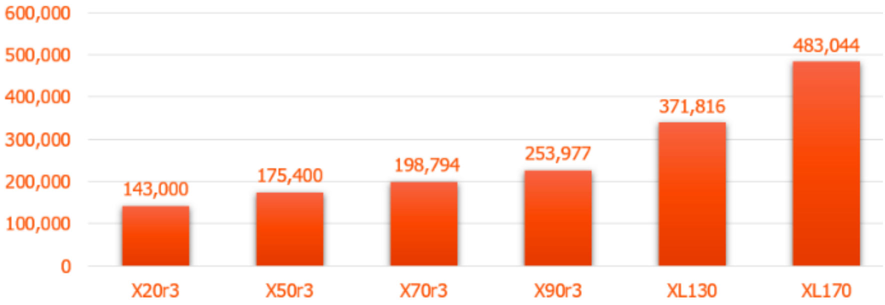

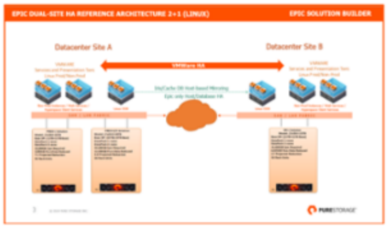

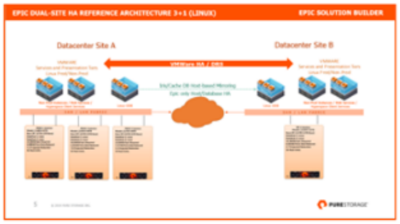

Summary

FlashStack is a powerful and reliable platform that has been specifically developed for Epic Healthcare computing deployments and cloud data centers. It utilizes a range of innovative technologies, including Cisco UCS Blade Servers and Cisco UCS Rack Servers, Cisco UCS Fabric Interconnects, Cisco Nexus 9000 Switches, Cisco MDS 9100 Fibre Channel switches, and Pure Storage FlashArray//X R3 Storage Array, to provide customers with a comprehensive solution that is designed and validated using best practices for compute, network, and storage.

With the introduction of Cisco UCS X-Series modular platform and Cisco Intersight, FlashStack now offers even more benefits to its users. These new technologies enhance the ability to provide complete visibility and orchestration across all elements of the FlashStack datacenter, enabling users to modernize their infrastructure and operations. This means that users can achieve higher levels of efficiency, scalability, and flexibility while also reducing deployment time, project risk, and IT costs.

FlashStack has been validated using industry-standard benchmarks to ensure that it meets the highest standards of performance, management, scalability, and resilience. This makes it the ideal choice for customers who are looking to deploy enterprise-class and other IT initiatives. With its powerful combination of hardware and software, FlashStack is capable of meeting the demands of the most complex and demanding IT environments, ensuring that users can focus on their core business objectives without having to worry about the underlying infrastructure.

Get More Business Value with Services

Whether you are planning your next-generation environment, need specialized know-how for a major deployment, or want to get the most from your current storage, Cisco Advanced Services, Pure Storage FlashArray//X R3 storage and our certified partners can help. We collaborate with you to enhance your IT capabilities through a full portfolio of services for your IT lifecycle with:

- Strategy services to align IT with your business goals:
    - Design services to architect your best storage environment
    - Deploy and transition services to implement validated architectures and prepare your storage environment
    - Operations services to deliver continuous operations while driving operational excellence and efficiency

Cisco Advanced Services and Pure Storage Support provide in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Author

Jeff Nichols, Technical Marketing Manager, Cisco Systems, Inc.

Jeff Nichols is a member of the Cisco's Computing Systems Product Group team focusing on design, testing, solutions validation, technical content creation, and performance testing/benchmarking. He has years of experience in Virtual Desktop Infrastructure (VDI), Server and Desktop Virtualization using Microsoft and VMware products.

Jeff is a subject matter expert on FlashStack, Cisco X-Series, Desktop/Server virtualization, Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA/AMD Graphics.

Tom Whalen, Senior Principal Solutions Architect, Pure Storage, Inc.

Tom Whalen has 30 years in the IT field and is a Solutions Architect and Evangelist for the global Healthcare team for Pure Storage. Prior to Pure Storage, Tom worked for six years at EMC, then Dell EMC, and 15+ years in IT at Aspirus Health Care in Wausau, WI. This included running applications like Epic, Hyland OnBase, Sunquest Laboratory, Sectra, FUJI, GE Medical, and Philips PACS/VNA technologies, and other medical and back-office application suites.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the following for their contribution and expertise that resulted in developing this document:

- Christopher Dudkiewicz, Senior Product Manager, Cisco Systems, Inc.
- Joel Dyson, Technical Marketing Engineer, Computing, Cisco Systems, Inc
- Joe Houghes, Senior Solutions Architect, Pure Storage, Inc.
- Craig Waters, Technical Director, Pure Storage, Inc.

Appendices

This appendix contains the following:

- Appendix A – References used in this guide
- Appendix B – Glossary
- Appendix C – Acronyms
- Appendix D – LVM creation examples

Appendix A – References used in this guide

This section provides links to additional information for each partner's solution component of this document.

- Cisco UCS X-Series Modular System

  https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-x-series-modular-system/series.html

  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/solution-overview-c22-2432175.html

  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/cisco-ucs-x9508-chassis-aag.html

  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x210c-m6-compute-node-aag.html

- Cisco UCS Manager Configuration Guides

  http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html

  https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-release-notes-list.html

- Cisco UCS Virtual Interface Cards

  https://www.cisco.com/c/en/us/products/interfaces-modules/unified-computing-system-adapters/index.html

- Cisco Nexus Switching References

  http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html

  https://www.cisco.com/c/en/us/products/switches/nexus-93180yc-fx-switch/index.html

- Cisco MDS 9000 Service Switch References

  http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html

  http://www.cisco.com/c/en/us/products/storage-networking/product-listing.html

  https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9100-series-multilayer-fabric-switches/datasheet-c78-739613.html

- Cisco Intersight References

  https://www.cisco.com/c/en/us/products/cloud-systems-management/intersight/index.html https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/intersight-ds.html

- FlashStack Cisco Design Guides

  https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/data-center-design-guides-all.html#FlashStack

- VMware References

  https://docs.vmware.com/en/VMware-vSphere/index.html

- Pure Storage Reference Documents

  https://www.flashstack.com/

  https://www.purestorage.com/content/dam/purestorage/pdf/datasheets/ps_ds_flasharray_03.pdf

  https://www.purestorage.com

  https://www.purestorage.com/products/evergreen-subscriptions.html

  https://www.purestorage.com/solutions/infrastructure/.html

  https://www.purestorage.com/solutions/infrastructure/-calculator.html

  https://support.purestorage.com/FlashArray/PurityFA/FlashArray_File_Services/001_Getting_Started/001_FA_File_Services_Quick_Start_Guide

  https://support.purestorage.com/FlashArray/PurityFA/FlashArray_File_Services/001_Getting_Started/002_FA_File_Services_Requirements_and_Best_Practices

- Epic Systems

  Documentation from Epic are provided via Epic's Galaxy customer-facing documentation portal. Traditionally technology partners do not have direct access to these materials. Epic will assist in providing documentation when requested. The following are two of the main documents needed to assure success:

  - ODB on Linux on VMware Architecture
  - Storage Configuration Quick Reference

Appendix B – Glossary

This glossary addresses some terms used in this document, for the purposes of aiding understanding. This is not a complete list of all multicloud terminology. Some Cisco product links are supplied here also, where considered useful for the purposes of clarity, but this is by no means intended to be a complete list of all applicable Cisco products.

aaS/XaaS **(IT capability provided as a Service)**	Some IT capability, X, provided as a service (XaaS). Some benefits are: • The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it. • There are very low barriers to entry, so that services can be quickly adopted and dropped in response to business demand, without the penalty of inefficiently utilized CapEx. • The service charge is an IT OpEx cost (pay-as-you-go), whereas the CapEx and the service infrastructure is the responsibility of the provider. • Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes. Such services are typically implemented as "microservices," which are accessed via REST APIs. This architectural style supports composition of service components into systems. Access to and management of aaS assets is via a web GUI and/or APIs, such that

	Infrastructure-as-code (IaC) techniques can be used for automation, for example, Ansible and Terraform. The provider can be any entity capable of implementing an aaS "cloud-native" architecture. The cloud-native architecture concept is well-documented and supported by open-source software and a rich ecosystem of services such as training and consultancy. The provider can be an internal IT department or any of many third-party companies using and supporting the same open-source platforms. Service access control, integrated with corporate IAM, can be mapped to specific users and business activities, enabling consistent policy controls across services, wherever they are delivered from.
Ansible	An infrastructure automation tool, used to implement processes for instantiating and configuring IT service components, such as VMs on an IaaS platform. Supports the consistent execution of processes defined in YAML "playbooks" at scale, across multiple targets. Because the Ansible artefacts (playbooks) are text-based, they can be stored in a Source Code Management (SCM) system, such as GitHub. This allows for software development like processes to be applied to infrastructure automation, such as, Infrastructure-as-code (see IaC below). https://www.ansible.com
AWS **(Amazon Web Services)**	Provider of IaaS and PaaS. https://aws.amazon.com
Azure	Microsoft IaaS and PaaS. https://azure.microsoft.com/en-gb/
Co-located data center	"A colocation center (CoLo)...is a type of data center where equipment, space, and bandwidth are available for rental to retail customers. Colocation facilities provide space, power, cooling, and physical security for the server, storage, and networking equipment of other firms and also connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity." https://en.wikipedia.org/wiki/Colocation_centre

Containers **(Docker)**	A (Docker) container is a means to create a package of code for an application and its dependencies, such that the application can run on different platforms which support the Docker environment. In the context of aaS, microservices are typically packaged within Linux containers orchestrated by Kubernetes (K8s). https://www.docker.com https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html
DevOps	The underlying principle of DevOps is that the application development and operations teams should work closely together, ideally within the context of a toolchain that automates the stages of development, test, deployment, monitoring, and issue handling. DevOps is closely aligned with IaC, continuous integration and deployment (CI/CD), and Agile software development practices. https://en.wikipedia.org/wiki/DevOps https://en.wikipedia.org/wiki/CI/CD
Edge compute	Edge compute is the idea that it can be more efficient to process data at the edge of a network, close to the endpoints that originate that data, or to provide virtualized access services, such as at the network edge. This could be for reasons related to low latency response, reduction of the amount of unprocessed data being transported, efficiency of resource utilization, and so on. The generic label for this is Multi-access Edge Computing (MEC), or Mobile Edge Computing for mobile networks specifically. From an application experience perspective, it is important to be able to utilize, at the edge, the same operations model, processes, and tools used for any other compute node in the system. https://en.wikipedia.org/wiki/Mobile_edge_computing
IaaS **(Infrastructure as-a-Service)**	Infrastructure components provided aaS, located in data centers operated by a provider, typically accessed over the public Internet. IaaS provides a base platform for the deployment of workloads, typically with containers and Kubernetes (K8s).
IaC **(Infrastructure as-Code)**	Given the ability to automate aaS via APIs, the implementation of the automation is typically via Python code, Ansible playbooks, and similar. These automation artefacts are programming code that define how the services are consumed. As such, they can be subject to the same code management and software development regimes as any other body of code. This means that infrastructure automation can be subject to all of the quality and consistency benefits, CI/CD, traceability, automated testing, compliance checking, and so on, that could be applied to any coding project. https://en.wikipedia.org/wiki/Infrastructure_as_code
IAM **(Identity and Access Management)**	IAM is the means to control access to IT resources so that only those explicitly authorized to access given resources can do so. IAM is an essential foundation to a secure multicloud environment. https://en.wikipedia.org/wiki/Identity_management
IBM **(Cloud)**	IBM IaaS and PaaS. https://www.ibm.com/cloud
Intersight	Cisco Intersight is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.

	https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html
GCP **(Google Cloud Platform)**	Google IaaS and PaaS. https://cloud.google.com/gcp
Kubernetes **(K8s)**	Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. https://kubernetes.io
Microservices	A microservices architecture is characterized by processes implementing fine-grained services, typically exposed via REST APIs and which can be composed into systems. The processes are often container-based, and the instantiation of the services often managed with Kubernetes. Microservices managed in this way are intrinsically well suited for deployment into IaaS environments, and as such, are the basis of a cloud native architecture. https://en.wikipedia.org/wiki/Microservices
PaaS **(Platform-as-a-Service)**	PaaS is a layer of value-add services, typically for application development, deployment, monitoring, and general lifecycle management. The use of IaC with IaaS and PaaS is very closely associated with DevOps practices.
Private on-premises data center	A data center infrastructure housed within an environment owned by a given enterprise is distinguished from other forms of data center, with the implication that the private data center is more secure, given that access is restricted to those authorized by the enterprise. Thus, circumstances can arise where very sensitive IT assets are only deployed in a private data center, in contrast to using public IaaS. For many intents and purposes, the underlying technology can be identical, allowing for hybrid deployments where some IT assets are privately deployed but also accessible to other assets in public IaaS. IAM, VPNs, firewalls, and similar are key technologies needed to underpin the security of such an arrangement.
REST API	Representational State Transfer (REST) APIs is a generic term for APIs accessed over HTTP(S), typically transporting data encoded in JSON or XML. REST APIs have the advantage that they support distributed systems, communicating over HTTP, which is a well-understood protocol from a security management perspective. REST APIs are another element of a cloud-native applications architecture, alongside microservices. https://en.wikipedia.org/wiki/Representational_state_transfer
SaaS **(Software-as-a-Service)**	End-user applications provided “aaS” over the public Internet, with the underlying software systems and infrastructure owned and managed by the provider.
SAML **(Security Assertion Markup Language)**	Used in the context of Single-Sign-On (SSO) for exchanging authentication and authorization data between an identity provider, typically an IAM system, and a service provider (some form of SaaS). The SAML protocol exchanges XML documents that contain security assertions used by the aaS for access control decisions. https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language
Terraform	An open-source IaC software tool for cloud services, based on declarative configuration files. https://www.terraform.io

Appendix C – Acronyms

AAA–Authentication, Authorization, and Accounting

ACP–Access-Control Policy

ACI–Cisco Application Centric Infrastructure

ACK–Acknowledge or Acknowledgement

ACL–Access-Control List

AD–Microsoft Active Directory

AFI–Address Family Identifier

AMP–Cisco Advanced Malware Protection

AP–Access Point

API–Application Programming Interface

APIC– Cisco Application Policy Infrastructure Controller (ACI)

ASA–Cisco Adaptative Security Appliance

ASM–Any-Source Multicast (PIM)

ASR–Aggregation Services Router

Auto-RP–Cisco Automatic Rendezvous Point protocol (multicast)

AVC–Application Visibility and Control

BFD–Bidirectional Forwarding Detection

BGP–Border Gateway Protocol

BMS–Building Management System

BSR–Bootstrap Router (multicast)

BYOD–Bring Your Own Device

CAPWAP–Control and Provisioning of Wireless Access Points Protocol

CDP–Cisco Discovery Protocol

CEF–Cisco Express Forwarding

CMD–Cisco Meta Data

CPU–Central Processing Unit

CSR–Cloud Services Routers

CTA–Cognitive Threat Analytics

CUWN–Cisco Unified Wireless Network

CVD–Cisco Validated Design

CYOD–Choose Your Own Device

DC–Data Center

DHCP–Dynamic Host Configuration Protocol

DM–Dense-Mode (multicast)

DMVPN–Dynamic Multipoint Virtual Private Network

DMZ–Demilitarized Zone (firewall/networking construct)

DNA–Cisco Digital Network Architecture

DNS–Domain Name System

DORA–Discover, Offer, Request, ACK (DHCP Process)

DWDM–Dense Wavelength Division Multiplexing

ECMP–Equal Cost Multi Path

EID–Endpoint Identifier

EIGRP–Enhanced Interior Gateway Routing Protocol

EMI–Electromagnetic Interference

ETR–Egress Tunnel Router (LISP)

EVPN–Ethernet Virtual Private Network (BGP EVPN with VXLAN data plane)

FHR–First-Hop Router (multicast)

FHRP–First-Hop Redundancy Protocol

FMC–Cisco Firepower Management Center

FTD–Cisco Firepower Threat Defense

GBAC–Group-Based Access Control

GbE–Gigabit Ethernet

Gbit/s–Gigabits Per Second (interface/port speed reference)

GRE–Generic Routing Encapsulation

GRT–Global Routing Table

HA–High-Availability

HQ–Headquarters

HSRP–Cisco Hot-Standby Routing Protocol

HTDB–Host-tracking Database (SD-Access control plane node construct)

IBNS–Identity-Based Networking Services (IBNS 2.0 is the current version)

ICMP– Internet Control Message Protocol

IDF–Intermediate Distribution Frame; essentially a wiring closet.

IEEE–Institute of Electrical and Electronics Engineers

IETF–Internet Engineering Task Force

IGP–Interior Gateway Protocol

IID–Instance-ID (LISP)

IOE–Internet of Everything

IoT–Internet of Things

IP–Internet Protocol

IPAM–IP Address Management

IPS–Intrusion Prevention System

IPSec–Internet Protocol Security

ISE–Cisco Identity Services Engine

ISR–Integrated Services Router

IS-IS–Intermediate System to Intermediate System routing protocol

ITR–Ingress Tunnel Router (LISP)

LACP–Link Aggregation Control Protocol

LAG–Link Aggregation Group

LAN–Local Area Network

L2 VNI–Layer 2 Virtual Network Identifier; as used in SD-Access Fabric, a VLAN.

L3 VNI– Layer 3 Virtual Network Identifier; as used in SD-Access Fabric, a VRF.

LHR–Last-Hop Router (multicast)

LISP–Location Identifier Separation Protocol

MAC–Media Access Control Address (OSI Layer 2 Address)

MAN–Metro Area Network

MEC–Multichassis EtherChannel, sometimes referenced as ***MCEC***

MDF–Main Distribution Frame; essentially the central wiring point of the network.

MnT–Monitoring and Troubleshooting Node (Cisco ISE persona)

MOH–Music on Hold

MPLS–Multiprotocol Label Switching

MR–Map-resolver (LISP)

MS–Map-server (LISP)

MSDP–Multicast Source Discovery Protocol (multicast)

MTU–Maximum Transmission Unit

NAC–Network Access Control

NAD–Network Access Device

NAT–Network Address Translation

NBAR–Cisco Network-Based Application Recognition (NBAR2 is the current version).

NFV–Network Functions Virtualization

NSF–Non-Stop Forwarding

OSI–Open Systems Interconnection model

OSPF–Open Shortest Path First routing protocol

OT–Operational Technology

PAgP–Port Aggregation Protocol

PAN–Primary Administration Node (Cisco ISE persona)

PCI DSS–Payment Card Industry Data Security Standard

PD–Powered Devices (PoE)

PETR–Proxy-Egress Tunnel Router (LISP)

PIM–Protocol-Independent Multicast

PITR–Proxy-Ingress Tunnel Router (LISP)

PnP–Plug-n-Play

PoE–Power over Ethernet (Generic term, may also refer to IEEE 802.3af, 15.4W at PSE)

PoE+–Power over Ethernet Plus (IEEE 802.3at, 30W at PSE)

PSE–Power Sourcing Equipment (PoE)

PSN–Policy Service Node (Cisco ISE persona)

pxGrid–Platform Exchange Grid (Cisco ISE persona and publisher/subscriber service)

PxTR–Proxy-Tunnel Router (LISP – device operating as both a PETR and PITR)

QoS–Quality of Service

RADIUS–Remote Authentication Dial-In User Service

REST–Representational State Transfer

RFC–Request for Comments Document (IETF)

RIB–Routing Information Base

RLOC–Routing Locator (LISP)

RP–Rendezvous Point (multicast)

RP–Redundancy Port (WLC)

RP–Route Processer

RPF–Reverse Path Forwarding

RR–Route Reflector (BGP)

RTT–Round-Trip Time

SA–Source Active (multicast)

SAFI–Subsequent Address Family Identifiers (BGP)

SD–Software-Defined

SDA–Cisco Software Defined-Access

SDN–Software-Defined Networking

SFP–Small Form-Factor Pluggable (1 GbE transceiver)

SFP+– Small Form-Factor Pluggable (10 GbE transceiver)

SGACL–Security-Group ACL

SGT–Scalable Group Tag, sometimes reference as Security Group Tag

SM–Spare-mode (multicast)

SNMP–Simple Network Management Protocol

SSID–Service Set Identifier (wireless)

SSM–Source-Specific Multicast (PIM)

SSO–Stateful Switchover

STP–Spanning-tree protocol

SVI–Switched Virtual Interface

SVL–Cisco StackWise Virtual

SWIM–Software Image Management

SXP–Scalable Group Tag Exchange Protocol

Syslog–System Logging Protocol

TACACS+–Terminal Access Controller Access-Control System Plus

TCP–Transmission Control Protocol (OSI Layer 4)

UCS– Cisco Unified Computing System

UDP–User Datagram Protocol (OSI Layer 4)

UPoE–Cisco Universal Power Over Ethernet (60W at PSE)

UPoE+– Cisco Universal Power Over Ethernet Plus (90W at PSE)

URL–Uniform Resource Locator

VLAN–Virtual Local Area Network

VM—Virtual Machine

VN–Virtual Network, analogous to a VRF in SD-Access

VNI–Virtual Network Identifier (VXLAN)

vPC–virtual Port Channel (Cisco Nexus)

VPLS–Virtual Private LAN Service

VPN–Virtual Private Network

VPNv4–BGP address family that consists of a Route-Distinguisher (RD) prepended to an IPv4 prefix

VPWS–Virtual Private Wire Service

VRF–Virtual Routing and Forwarding

VSL–Virtual Switch Link (Cisco VSS component)

VSS–Cisco Virtual Switching System

VXLAN–Virtual Extensible LAN

WAN–Wide-Area Network

WLAN–Wireless Local Area Network (generally synonymous with IEEE 802.11-based networks)

WoL–Wake-on-LAN

xTR–Tunnel Router (LISP - device operating as both an ETR and ITR)

Appendix D - LVM creation examples

This appendix is supported by both the Pure Storage and Epic best-practice guides. Before volumes are masked to VMware from the Pure Storage FlashArray, multipathing should be set-up to ensure that when volumes are presented, they inherit the attributes for Pure Storage and settings for Epic as directed by the guides. If you have any questions, please contact to your Epic Server Systems liaison and/or your local Pure Storage System Engineer.

```
Step 1 - create lvm physical volumes
          Before you do this, make sure have the UUID's on the host meaning the host
                  imported them.  In some instances, they may not all appear.  This is normal
                  but is a pain.  The best way to handle this is to reboot the host and it will
                  force rediscovery of any FC devices

          Using the new HAID (Host & Array ID), create the physical volumes, enter the cmd:

          pvcreate /dev/mapper/HAID
                  example: /dev/mapper/ABC123000222333444555662345320

          -Do this for each of the volumes you created on the host.  The LVM will confirm
                  that the physical devices have been created.  If this fails, double check the
                  numbers for a typo.

                  PRO-TIP: After your rescan of deviceas to bring in the devices from the LUN's
provisioned,
                           run the multipath -ll command to view the settings of the devices.  You
                                     are looking for the policy: round-robin for every device mapper device
(dm-#)


                                     example:
```

```
                                     3624a93705a9272924e4b4d7e00011cb5 dm-16 PURE,FlashArray
                                           size=5.0T features='0' hwhandler='1 alua' wp=rw
                                           -+- policy='round-robin 0' prio=50 status=active
                                             |- 12:0:17:23 sdy     65:128   active ready running
                                             |- 13:0:17:23 sdgy    132:224  active ready running
                                             |- 14:0:0:23  sdse    135:288  active ready running
                                             |- 15:0:6:23  sdaop   67:1168  active ready running
                                             |- 12:0:18:23 sdbo    68:32    active ready running
                                             |- 13:0:18:23 sdkd    66:272   active ready running
                                             |- 14:0:1:23  sdyl    129:592  active ready running
                                             |- 15:0:0:23  sdnv    128:272  active ready running
                                             |- 12:0:19:23 sdde    70:192   active ready running
                                             |- 13:0:19:23 sdoa    128:352  active ready running
                                             |- 14:0:2:23  sdadi   65:832   active ready running
                                             |- 15:0:1:23  sduq    67:544   active ready running
                                             |- 12:0:20:23 sdew    129:128  active ready running
                                             |- 13:0:20:23 sdur    67:560   active ready running
                                             |- 14:0:3:23  sdakd   132:848  active ready running
                                             `- 15:0:2:23  sdzk    130:736  active ready running

                                  The "policy='round-robin 0'" is the indicator that the devices have
                used this option when the devices were imported.


Step 2 - Create the LINUX Volume Group
         -This is NOT the same as the array Volume Group construct - but using a similar name to the array
                    may help with tracking down errors and stuff ***

                  vgcreate vgLINUXLVMVolumeGroup [/dev/mapper/ABC1230002223334445556623453 2C0]
                  example: vgcreate vgLINUX01 /dev/mapper/ABC12300022233344455566234532C0
/dev/mapper/ABC12300022233344455566234532C0
             -for more than one physical volume to add it to the volume group, include a SINGLE SPACE
          bewteen mapper commands like shown above

          If no errors, the LVM should indicate the VG has been created

Step 3 - Create the Logical Volumes
         Epic's best practiceto use striped, I'll show you that command.
         ** Each logical volume defined is assumed to be a mount-point but it doesn't have to be.
         Each logical volume can be a mount point.  Below is a high-level breakdown - see the manpage for
further definition

          lvcreate -n [lvLogicalVolNameXX] -L {size of LV] -i {# of disks to stripe] -L (Size of Stripe)
[LINUX Volume Group Name]
          example: lvcreate -n lvLogicalVolume01 -L 1000G -i 8 -I 4M vgVolumeGroup

          -This will create a logical volume(s) using the capacity from the vgVolumeGroup of 1000G
          across 8 or 16 disks with a stripe size of 4MB
          **IMPORTANT - You can only allocate capacity from the total size of the Volume Group.
            so if you want to have 4 logical volumes, you do some simple math; 100G VG capacity
            and you need 4 Logical Volumes, the size of the LV's wouild be 100/4 or 25G each

Step 4 - Make the File Systems
         This will actually format the volumes with the LVM type you want to use.  IMO 'xfs or xf2' are
                  the best with my personal preference of 'xfs2'  Do this for each logical volume you want
to
                  create. Tip - make sure you know the correct format (xfs,xfs2, etc.) before you 'mkfs'
                  lvchangre
                  mkfs -t xfs[2] -K /dev/[vgVolumeGroupName]/[lvLogicalVolNameXX]
                  example: mkfs -t xfs2 /dev/vgVolumeGroup/lvLogicalVolume00
                           mkfs -t xfs2 /dev/vgVolumeGroup/lvLogicalVolume01  (and so on...)
                  -if all goes correctly, Linux should respond with a logical volumes created.

Step 5 - Create the file systems in LINUX
         The Home stretch - This will create the file systems needed to allow you access
                  to the allocation of space you carved out of the array for the host to consume.
                  -This is simply making file systems.  IMPORTANT - if you are creating a nested file
systems
                  you want to make sure you build them in the right order to ensure the FS tree is what you
want.
                  example: If I need "/epic" and then "/epic/prd01", you create the /epic file system first,
then
```

```
                  the "/epic/prd01" file system next.
                  -Also remember, before you create the file systems, make sure you are in the root-level
                  of the file systems.  enter 'pwd' and the response should be "/".  If not, then simplt
                  type "cd /" and another 'pwd' and you should at the root file system level.

                  mkdir /[filesystem name]
                  examples:   mkdir /filesys01
                              mkdir /filesys01/filesysA
                              mkdir /filesys01/filesysA/filesys01A

Step 6 - Mounting the FS's
         This step is where you will be able to mount the file systems for use.  This is done in two
                  ways to ensure you are mounting them correctly.

                  mount [/dev/vgVolumeGroup/lvLogicalVolume] [/file system]
                  example:
                          mount -o noatime /dev/vgVolumeGroup/lvLogicalVolume00 /filesys00
                          mount -o noatime /dev/vgVolumeGroup/lvLogicalVolume01 /filesys01

                  IMPORTANT - Linux may not return a message back other than giving you back the cursor.
                              You can simply do a 'df -H' and you should see your mounted file systems
appear.
                  Pro-Tip: When mounting the fs's with their respective LV's you will sometimes run into an
error
                           that will indicate that a mount action can't be done.  In most cases, the reason
this
                           happenes is based in the order of the mount operations.  What happess is the LVM
tags
                           the LV with a status of 'NOT Available'.
                          example:
                          #lvdisplay
                          --- Logical volume ---
                          LV Path                /dev/vgVolumeGroup/lvELogicalVolume
                          LV Name                lvEpicFA13020WIJ01
                          VG Name                vgEpicFA13020WIJ
                          LV UUID                d2mMMl-hk7A-0FL7-2eVh-GECU-fkTP-UTeTEJ
                          LV Write Access        read/write
                          LV Creation host, time sn1-intel-e11-33, 2022-11-30 08:40:13 +0000
                          --->  LV Status              NOT available <---
                          LV Size                7.32 TiB
                          Current LE             1920000
                          Segments               1
                          Allocation             inherit
                          Read ahead sectors     auto
                            Corrective Action:
                              #lvchange [volume group name] --activate y
                               example: lvchange vgEpicFA13020WIJ --activate y
                               check the LV status - it should now be 'avaialble'
                               #lvdisplay
                                 --- Logical volume ---
                                 LV Path                /dev/vgEpicFA13020WIJ/lvEpicFA13020WIJ01
                                   LV Name                lvEpicFA13020WIJ01
                                   VG Name                vgEpicFA13020WIJ
                                   LV UUID                d2mMMl-hk7A-0FL7-2eVh-GECU-fkTP-UTeTEJ
                                   LV Write Access        read/write
                                   LV Creation host, time sn1-intel-e11-33, 2022-11-30 08:40:13 +0000
                                --->  LV Status              available  <---
                                   LV Size                7.32 TiB
                                   Current LE             1920000
                                   Segments               1
                                   Allocation             inherit
                                   Read ahead sectors     auto

Step 7 - Dealing with the fstab file
         The fstab file is file that is used to mount the file systems you want when the system is
         up and running.  This is a step that you do have to do right away....
         its advised to not initially add the entries into the fstab until you have tested your file
         mounts and they work like they should. Also the fstab file has to be done correctly or when
         you attempt to boot the server, it will hang at a file system mount step and you may not know this
         right away.  Let me know if you want them added and I can do that for you.
         If you want to UNmount the filesystems, you simple type:
```

```
umount /filesys00
umount /filesys01
...
Pro-Tip: unless the directory is a individual mount point, you will want to umount
         is the opposite order from how they were mounted.
                     mount:  mount /epic
                                  mount /epic/prd01
                    umount:  umount /epic/prd01
                             umount /epic
```

Feedback

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

CVD Program

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 IMPLE-MENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

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**.