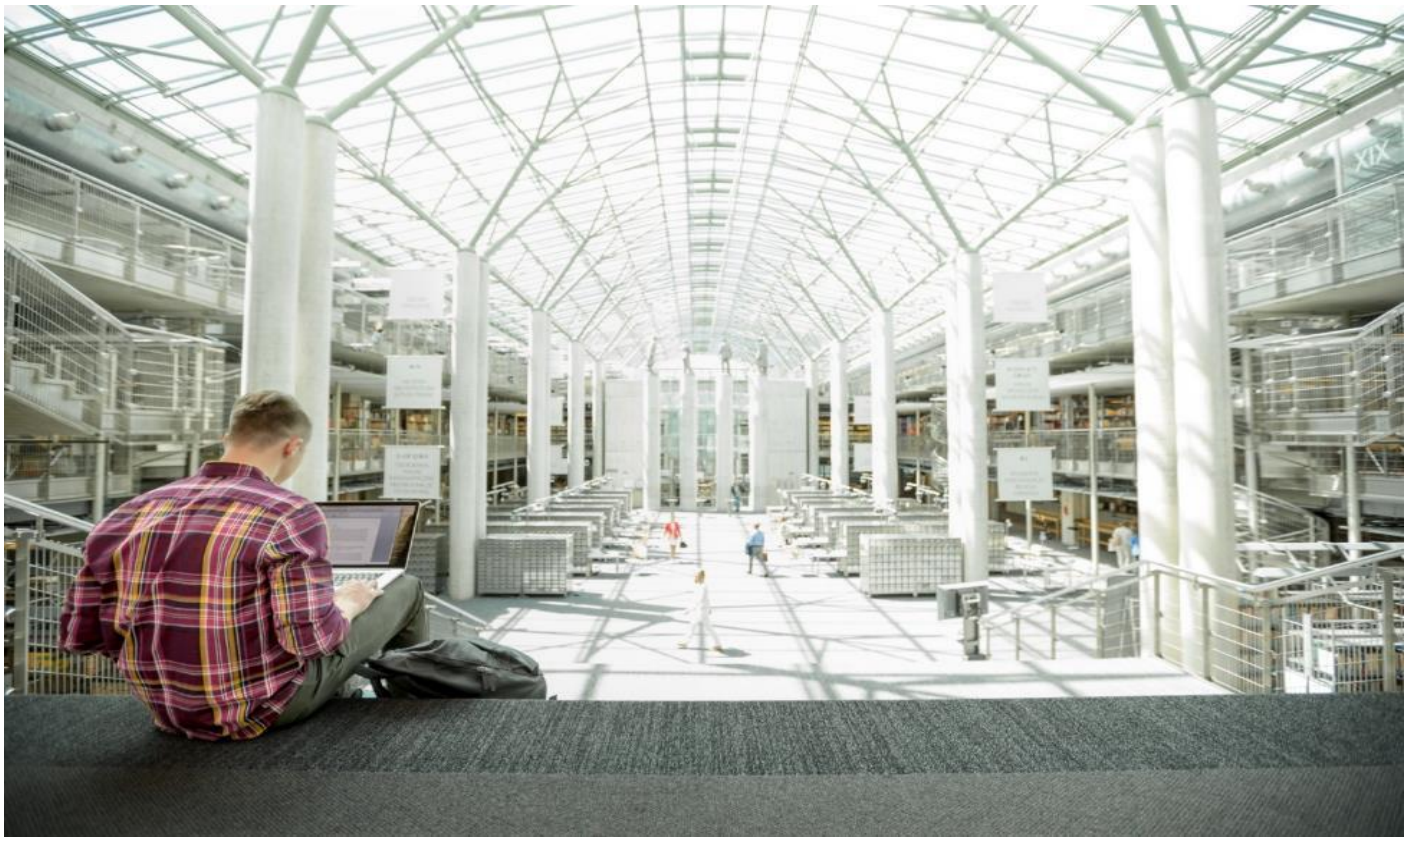

FlashStack Virtual Server Infrastructure with Cisco UCS X-Series and VMware 7.0 U2 Design Guide

Published: November 2021

CISCO VALIDATED DESIGN

FlashStack

In partnership with:

PURESTORAGE

About the Cisco Validated Design Program

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

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

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series. Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_P2)

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2021 Cisco Systems, Inc. All rights reserved.

Contents

Executive Summary 4
Solution Overview 5
Technology Overview 7
Solution Design 22
Deployment Hardware and Software 57
Validation 59
Summary 60
Appendix 61
About the Authors 62
Feedback 63

Executive Summary

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

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

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

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

Solution Overview

Introduction

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

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

Audience

The intended audience of this document includes but is not limited to 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 around incorporating the Cisco Intersight software-managed Cisco UCS X-Series platform within the FlashStack solution. The document introduces various design elements and addresses various considerations and best practices for a successful deployment. It also highlights the design and product requirements for integrating virtualization and storage systems with the Cisco Intersight platform to deliver a true cloud-based integrated approach to infrastructure management.

What's New in this Release?

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

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

Solution Summary

The FlashStack VSI with Cisco UCS X-Series and VMware 7.0 U2 offers the following key customer benefits:

- Simplified cloud-based management of the solution components
- Hybrid cloud-ready, policy-driven modular design
- Highly available and scalable platform with flexible architecture that supports various deployment models

- Cooperative support model and Cisco® solution support
- Architecture that is easy to deploy, consume, and manage, saving time and resources required to research, procure, and integrate off-the-shelf components
- Support for component monitoring, solution orchestration, and workload optimization

Like all other FlashStack solution designs, FlashStack VSI with Cisco UCS X-Series and VMware 7.0 U2 is configurable according to the demand and usage. Customers can purchase exactly the infrastructure they need for their current application requirements and then can scale up by adding more resources to the FlashStack system or scale out by adding more FlashStack instances. By moving the management from the fabric interconnects into the cloud, the solution can respond to speed and scale of customer deployments with a constant stream of new capabilities delivered from the Cisco Intersight SaaS model at cloud scale.

Technology Overview

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

FlashStack Components

FlashStack architecture is built using the following infrastructure components for compute, network, and storage (Figure 1):

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus® switches
- Cisco MDS 9000 switches
- Pure Storage FlashArray

Figure 1. FlashStack Components

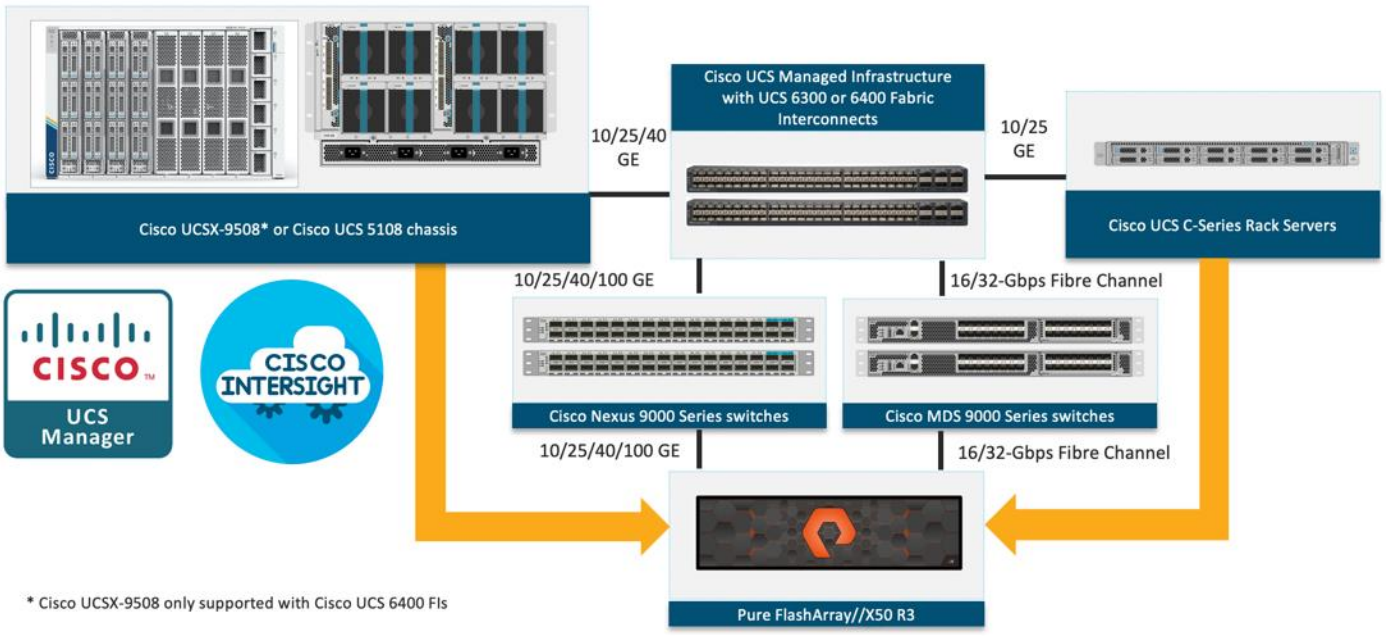

All the FlashStack components are integrated, so customers can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlashStack is its ability 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.

The FlashStack solution with Cisco UCS X-Series uses following hardware components:

- Cisco UCS X9508 chassis with any number of Cisco UCS X210c M6 compute nodes

- Cisco fourth-generation 6454 fabric interconnects to support 25- and 100-GE connectivity from various components
- High-speed Cisco NXOS-based Nexus 93180YC-FX3 switching design to support up to 100-GE connectivity
- Pure Storage FlashArray//X50 R3 with high-speed Ethernet or Fibre Channel connectivity
- Pure FlashArray//X50 R3 storage with 25GbE connectivity to Cisco Nexus switching fabric and 32Gb FC connectivity to Cisco MDS switching fabric.

The software components consist of:

- 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
- VMware vCenter 7.0 U2 to set up and manage the virtual infrastructure as well as integration of the virtual environment with Cisco Intersight software

The next section outlines these critical product highlights and features.

Cisco Unified Compute 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 design that will support future innovations and management in the cloud (Figure 2). Decoupling and moving platform management to the cloud allows the Cisco UCS platform to respond to features and scalability requirements much faster and more efficiently. Cisco UCS X-Series state-of-the-art hardware simplifies the datacenter design by providing flexible server options. A single server type that supports a broader range of workloads results in fewer different datacenter products to manage and maintain. The Cisco Intersight cloud management platform manages the Cisco UCS X-Series as well as integrates with third-party devices. These devices include VMware vCenter and Pure Storage to provide visibility, optimization, and orchestration from a single platform, thereby enhancing agility and deployment consistency.

Figure 2. Cisco UCS X9508 Chassis

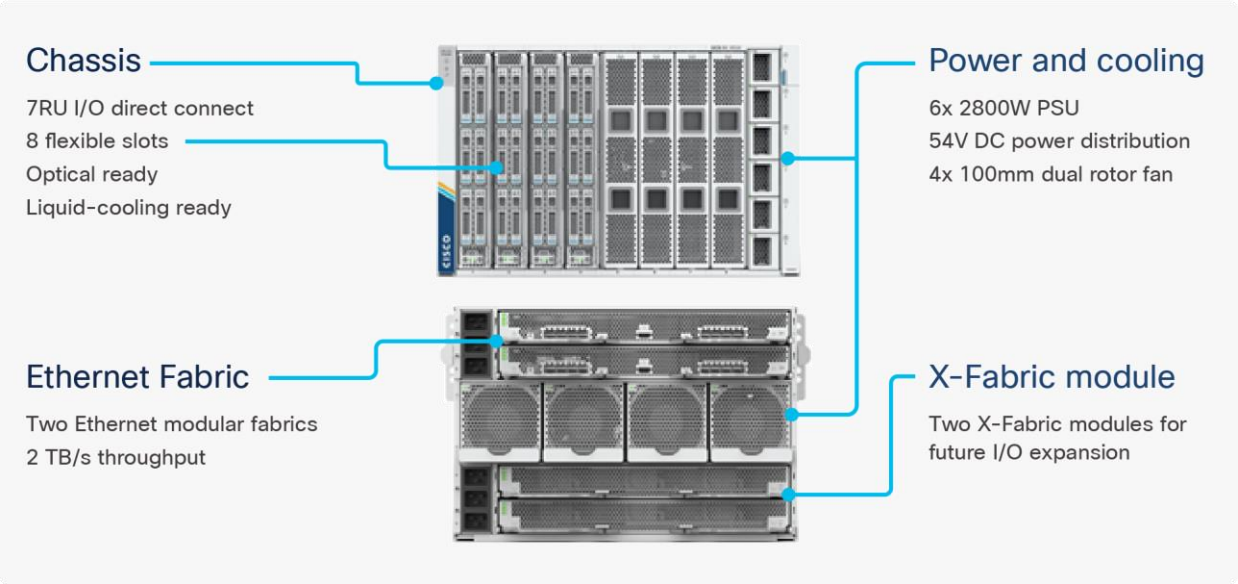

The various components of the Cisco UCS X-Series are described in the following sections.

Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As seen in Figure 3, Cisco UCS X9508 chassis has only a power-distribution midplane. This innovative 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 3. Cisco UCS X9508 Chassis – Innovative Design

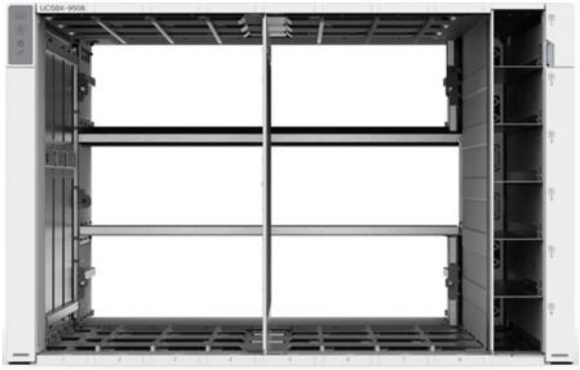

The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes and a pool of future I/O resources that may include GPU accelerators, disk storage, and non-volatile memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) 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 X9508s 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 4. 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 M6 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to 8 Cisco UCS X210c M6 Compute Nodes. The hardware details of the Cisco UCS X210c M6 Compute Nodes are shown in Figure 5:

Figure 5. Cisco UCS X210c M6 Compute Node

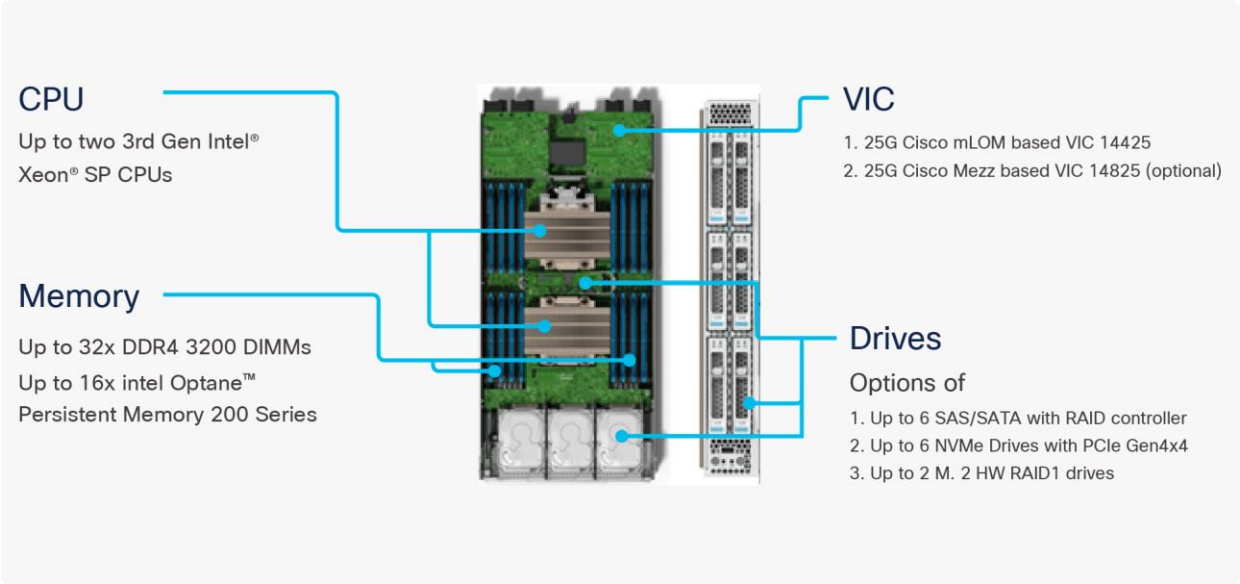

The Cisco UCS X210c M6 features:

- **CPU:** Up to 2x 3rd Gen Intel Xeon Scalable Processors with up to 40 cores per processor and 1.5 MB Level 3 cache per core
- **Memory:** Up to 32 x 256 GB DDR4-3200 DIMMs for a maximum of 8 TB of main memory. The Compute Node can also be configured for up to 16 x 512-GB Intel Optane persistent memory DIMMs for a maximum of 12 TB of memory
- **Disk storage:** Up to 6 SAS or SATA drives can be configured with an internal RAID controller, or customers can configure up to 6 NVMe drives. 2 M.2 memory cards can be added to the Compute Node with RAID 1 mirroring.
- **Virtual Interface Card (VIC):** Up to 2 VICs including an mLOM Cisco VIC 14425 and a mezzanine Cisco VIC card 14825 can be installed in a Compute 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 M6 Compute Nodes support the following two Cisco fourth-generation VIC cards:

Cisco VIC 14425

Cisco VIC 14425 fits the mLOM slot in the Cisco X210c Compute Node and enables up to 50 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 100 Gbps of connectivity per server. Cisco VIC 14425 connectivity to the IFM and up to the fabric interconnects is delivered through 4x 25-Gbps connections, which are configured automatically as 2x 50-Gbps port channels. Cisco VIC 14425 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.

Figure 6. Single Cisco VIC 14425 in Cisco UCS X210c M6

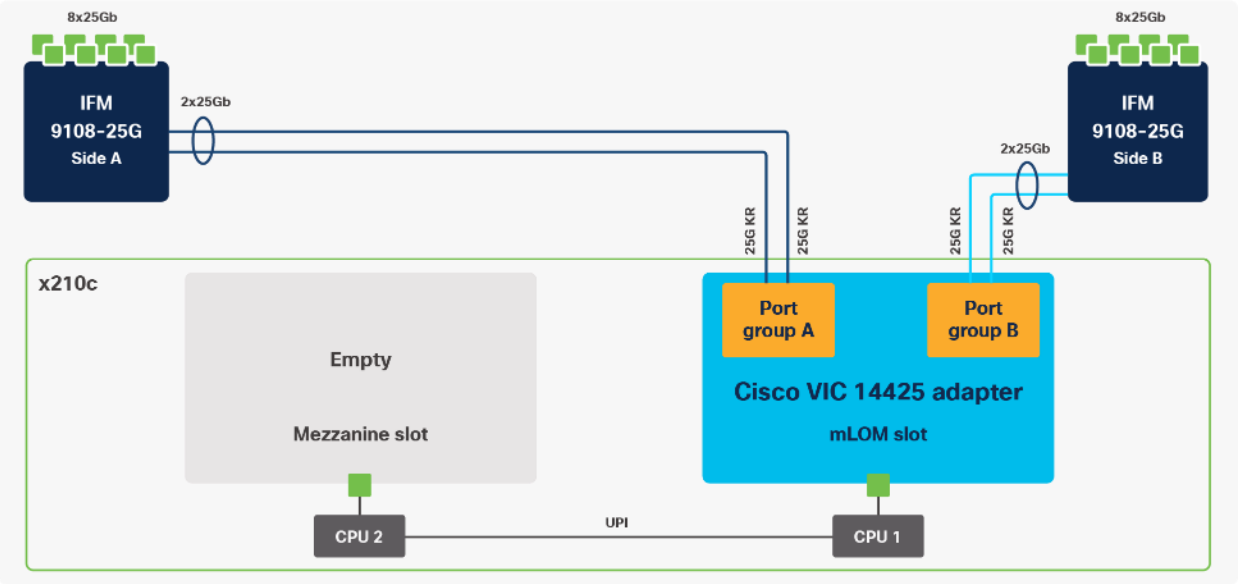

The connections between the 4th generation Cisco VIC (Cisco UCS VIC 1440) in the Cisco UCS B200 blades and the I/O modules in the Cisco UCS 5108 chassis comprise of multiple 10Gbps KR lanes. The same connections between Cisco VIC 14425 and IFMs in Cisco UCS X-Series comprise of multiple 25Gbps KR lanes resulting in 2.5x better connectivity in Cisco UCS X210c M6 Compute Nodes. The network interface speed comparison between VMware ESXi installed on Cisco UCS B200 M5 with VIC 1440 and Cisco UCS X210c M6 with VIC 14425 is shown in Figure 7.

Figure 7. Network Interface Speed Comparison

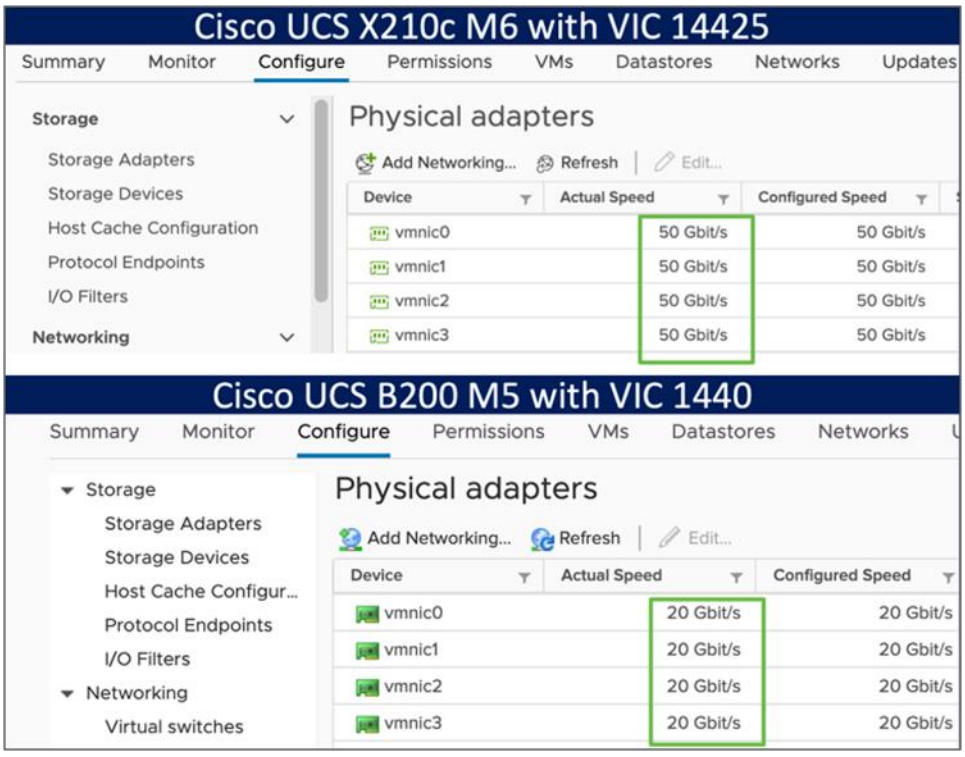

Cisco VIC 14825

The optional Cisco VIC 14825 fits the mezzanine slot on the server. A bridge card (UCSX-V4-BRIDGE) extends this VIC's 2x 50 Gbps of network connections up to the mLOM slot and out through the mLOM's IFM connectors, bringing the total bandwidth to 100 Gbps per fabric for a total bandwidth of 200 Gbps per server.

Figure 8. Cisco VIC 14425 and 14825 in Cisco UCS X210c M6

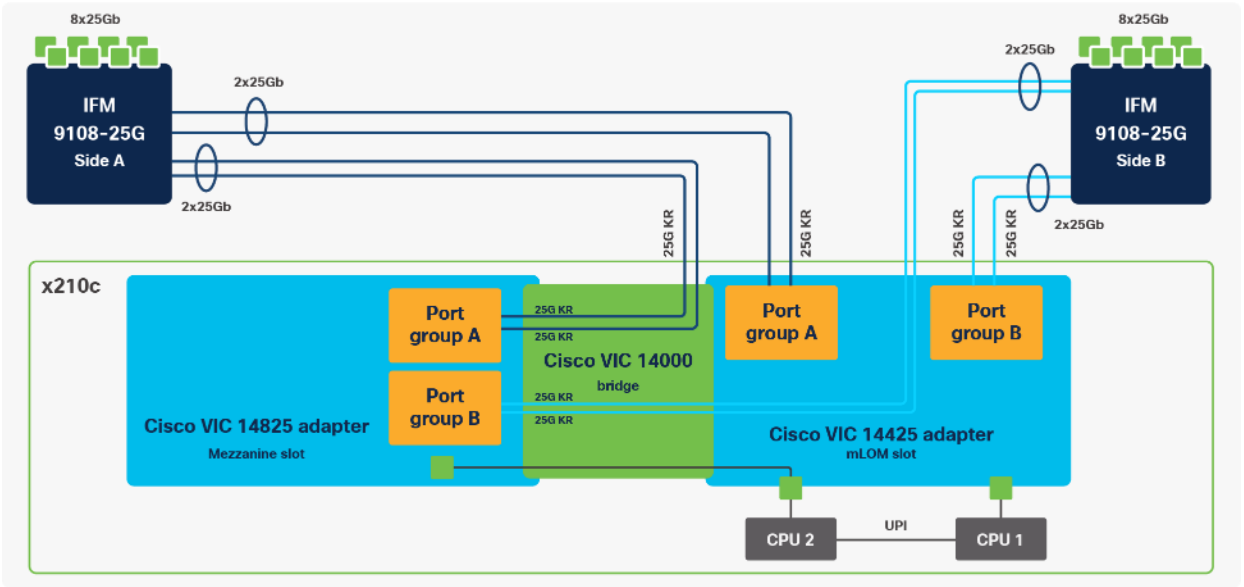

Cisco UCS 6400 Series Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide a single point of connectivity and management for the entire Cisco UCS system. Typically deployed as an active/active pair, the system's FIs integrate all components into a single, highly available management domain controlled by the Cisco UCS Manager or Cisco Intersight. Cisco UCS FIs provide a single unified fabric for the system, with low-latency, lossless, cut-through switching that supports LAN, SAN, and management traffic using a single set of cables.

Figure 9. Cisco UCS 6454 Fabric Interconnect

Cisco UCS 6454 utilized in the current design is a 54-port Fabric Interconnect. This single RU device includes 28 10/25 Gbps Ethernet ports, 4 1/10/25-Gbps Ethernet ports, 6 40/100-Gbps Ethernet uplink ports, and 16 unified ports that can support 10/25 Gigabit Ethernet or 8/16/32-Gbps Fibre Channel, depending on the SFP.

For supporting the Cisco UCS X-Series, the fabric interconnects must be configured in Intersight Managed Mode (IMM). This option replaces the local management with Cisco Intersight cloud or appliance-based management.

Cisco Intersight

The Cisco Intersight platform is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support. The Cisco Intersight platform is designed to be modular, so customers can adopt services based on their individual requirements. The platform significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses a unified Open API design that natively integrates with third-party platforms and tools.

Figure 10. Cisco Intersight Overview

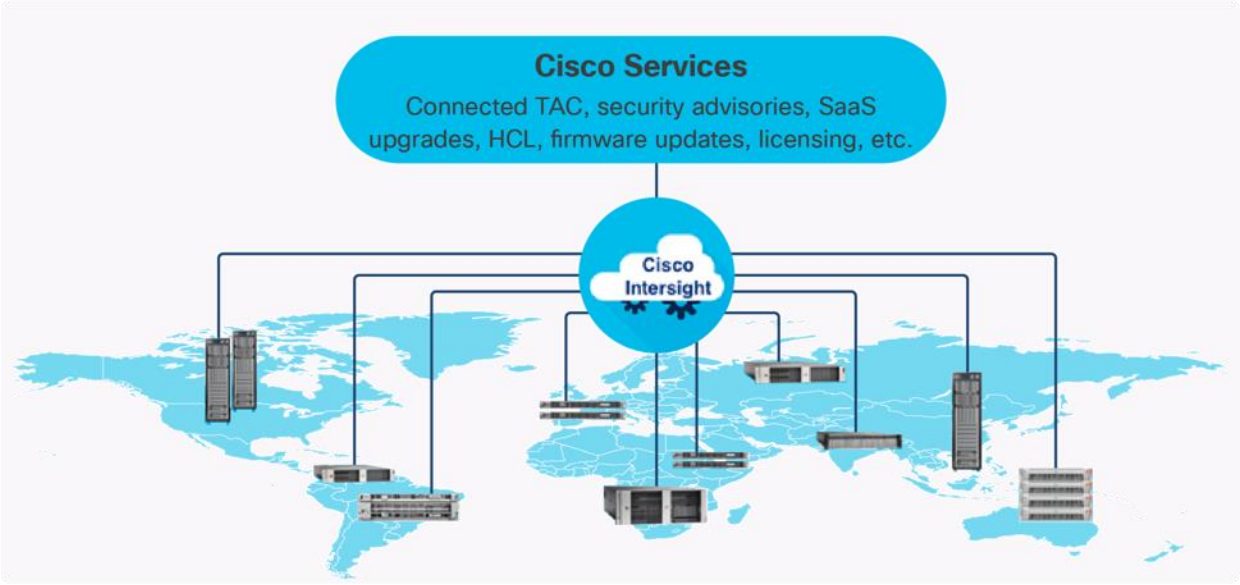

The main benefits of Cisco Intersight infrastructure services are as follows:

- Simplify daily operations by automating many daily manual tasks
- Combine the convenience of a SaaS platform with the capability to connect from anywhere and manage infrastructure through a browser or mobile app
- Stay ahead of problems and accelerate trouble resolution through advanced support capabilities
- Gain global visibility of infrastructure health and status along with advanced management and support capabilities
- Upgrade to add workload optimization and Kubernetes services when needed

Cisco Intersight Virtual Appliance and Private Virtual Appliance

In addition to the SaaS deployment model running on Intersight.com, on-premises options can be purchased separately. The Cisco Intersight Virtual Appliance and Cisco Intersight Private Virtual Appliance are available for organizations that have additional data locality or security requirements for managing systems. The Cisco Intersight Virtual Appliance delivers the management features of the Cisco Intersight platform in an easy-to-deploy VMware Open Virtualization Appliance (OVA) or Microsoft Hyper-V Server virtual machine that allows you to control the system details that leave your premises. The Cisco Intersight Private Virtual Appliance is provided in a form factor specifically designed for users who operate in disconnected (air gap) environments. The Private Virtual Appliance requires no connection to public networks or back to Cisco to operate.

Cisco Intersight Assist

Cisco Intersight Assist helps customers add endpoint devices to Cisco Intersight. A data center could have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight but does not connect to Intersight directly needs Cisco Intersight Assist to provide the necessary connectivity. In FlashStack, VMware vCenter and Pure Storage FlashArray connect to Intersight with the help of Intersight Assist appliance.

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. More details about the Cisco Intersight Assist VM deployment configuration is covered in later sections.

Licensing Requirements

The Cisco Intersight platform uses a subscription-based license with multiple tiers. Customers can purchase a subscription duration of one, three, or five years and choose the required Cisco UCS server volume tier for the selected subscription duration. Each Cisco endpoint automatically includes a Cisco Intersight Base license at no additional cost when customers access the Cisco Intersight portal and claim a device. Customers can purchase any of the following higher-tier Cisco Intersight licenses using the Cisco ordering tool:

- **Cisco Intersight Essentials:** Essentials includes all the functions of the Base license plus additional features, including Cisco UCS Central Software and Cisco Integrated Management Controller (IMC) supervisor entitlement, policy-based configuration with server profiles, firmware management, and evaluation of compatibility with the Cisco Hardware Compatibility List (HCL).
- **Cisco Intersight Advantage:** Advantage offers all the features and functions of the Base and Essentials tiers. It includes storage widgets and cross-domain inventory correlation across compute, storage, and virtual environments (VMWare ESXi). It also includes OS installation for supported Cisco UCS platforms.

- **Cisco Intersight Premier:** In addition to all of the functions provided in the Advantage tier, Premier includes full subscription entitlement for Intersight Orchestrator, which provides orchestration across Cisco UCS and third-party systems.

Servers in the Cisco Intersight managed mode require at least the Essentials license. For more information about the features provided in the various licensing tiers, see https://intersight.com/help/getting_started#licensing_requirements.

Cisco Nexus Switching Fabric

The Cisco Nexus 9000 Series Switches offer both modular and fixed 1/10/25/40/100 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of nonblocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.

Figure 11. Cisco Nexus 93180YC-FX3 Switch

The Cisco Nexus 9000 series switch featured in this design is the Cisco Nexus 93180YC-FX3 configured in NX-OS standalone mode. NX-OS is a purpose-built data-center operating system designed for performance, resiliency, scalability, manageability, and programmability at its foundation. It provides a robust and comprehensive feature set that meets the demanding requirements of virtualization and automation.

The Cisco Nexus 93180YC-FX3 Switch is a 1RU switch that supports 3.6 Tbps of bandwidth and 1.2 bpps. The 48 downlink ports on the 93180YC-FX3 can support 1-, 10-, or 25-Gbps Ethernet, offering deployment flexibility and investment protection. The six uplink ports can be configured as 40- or 100-Gbps Ethernet, offering flexible migration options.

Cisco MDS 9132T 32G Multilayer Fabric Switch

The Cisco MDS 9132T 32G Multilayer Fabric Switch is the next generation of the highly reliable, flexible, and low-cost Cisco MDS 9100 Series switches. It combines high performance with exceptional flexibility and cost effectiveness. This powerful, compact one Rack-Unit (1RU) switch scales from 8 to 32 line-rate 32 Gbps Fibre Channel ports.

Figure 12. Cisco MDS 9132T 32G Multilayer Fabric Switch

The Cisco MDS 9132T delivers advanced storage networking features and functions with ease of management and compatibility with the entire Cisco MDS 9000 family portfolio for reliable end-to-end connectivity. 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.

Cisco Data Center Network Manager (DCNM)-SAN

Cisco DCNM-SAN can be used to monitor, configure, and analyze Cisco 32Gbps Fibre Channel fabrics and show information about the Cisco Nexus switching fabric. Cisco DCNM-SAN is deployed as a virtual appliance from an OVA and is managed through a web browser. Once the Cisco MDS and Nexus switches are added with the appropriate credentials and licensing, monitoring of the SAN and Ethernet fabrics can begin. Additionally, VSANs, device aliases, zones, and zone sets can be added, modified, and deleted using the DCNM point-and-click interface. Device Manager can also be used to configure the Cisco MDS switches. SAN Analytics can be added to Cisco MDS switches to provide insights into the fabric by allowing customers to monitor, analyze, identify, and troubleshoot performance issues.

Cisco DCNM Integration with Cisco Intersight

The Cisco Network Insights Base (Cisco NI Base) application provides several TAC assist functionalities which are useful when working with Cisco TAC. Cisco NI base provides a way for Cisco customers to collect technical support information across multiple devices and upload them to Cisco Cloud. The Cisco NI Base app collects the CPU, device name, device product id, serial number, version, memory, device type, and disk usage information for the nodes in the fabric. Cisco NI Base application is connected to the Cisco Intersight cloud portal through a device connector which is embedded in the management controller of the Cisco DCNM platform. The device connector provides a safe way for connected Cisco DCNM to send and receive information from the Cisco Intersight portal, using a secure Internet connection.

Pure Storage FlashArray//X

The Pure Storage FlashArray Family delivers software-defined all-flash power and reliability for businesses of every size. FlashArray is all-flash enterprise storage that is up to 10X faster, space and power efficient, reliable, and far simpler than other available solutions. Compared to traditional performance disk arrays, FlashArray costs less with total cost of ownership (TCO) savings of up to 50%. At the top of the FlashArray line is FlashArray//X—the first mainstream, 100-percent NVMe, enterprise-class all-flash array. //X represents a higher performance tier for mission-critical databases, top-of-rack flash deployments, and tier 1 application consolidation. //X, at 1PB in 3RU, hundred-microsecond range latency, and GBs of bandwidth, delivers unparalleled performance density and consolidation. FlashArray//X is ideal for cost-effective consolidation of everything on flash, including accelerating a single database, scaling virtual desktop environments, or powering an all-flash cloud.

Purity for FlashArray (Purity//FA 6)

Every FlashArray is driven by Purity Operating Environment software. Purity//FA6 implements advanced data reduction, storage management, and flash management features, enabling customers to enjoy tier 1 data services for all workloads. Purity software provides proven 99.9999-percent availability over 2 years, completely nondisruptive operations, 2X better data reduction, and the power and efficiency of DirectFlash™. Purity also includes enterprise-grade data security, comprehensive data-protection options, and complete business continuity with an ActiveCluster multi-site stretch cluster. All these features are included with every Pure Storage array.

FlashArray//X R3 specification

Table 1 lists both capacity and physical aspects of various FlashArray systems.

Table 1. FlashArray//X R3 Specifications

//X10	Up to 73 TB/66.2 TiB (tebibyte) effective	3RU; 640–845 watts (nominal – peak)

	capacity** Up to 22 TB/19.2 TiB raw capacity	95 lb. (43.1 kg) fully loaded; 5.12 x 18.94 x 29.72 in.
//X20	Up to 314 TB/285.4 TiB (tebibyte) effective capacity** Up to 94 TB/88 TiB raw capacity†	3RU; 741–973 watts (nominal – peak) 95 lb. (43.1 kg) fully loaded; 5.12 x 18.94 x 29.72 in.
//X50	Up to 663 TB/602.9 TiB effective capacity** Up to 185 TB/171 TiB raw capacity†	3RU; 868–1114 watts (nominal – peak) 95 lb. (43.1 kg) fully loaded; 5.12 x 18.94 x 29.72 in.
//X70	Up to 2286 TB/2078.9 TiB effective capacity** Up to 622 TB/544.2 TiB raw capacity†	3RU; 1084–1344 watts (nominal – peak) 97 lb. (44.0 kg) fully loaded; 5.12 x 18.94 x 29.72 in.
//X90	Up to 3.3 PB/3003.1 TiB effective capacity** Up to 878 TB/768.3 TiB raw capacity†	3–6RU; 1160–1446 watts (nominal – peak) 97 lb. (44 kg) fully loaded; 5.12 x 18.94 x 29.72 in.
DirectFlash Shelf	Up to 1.9 PB effective capacity** Up to 512 TB/448.2 TiB raw capacity	3RU; 460–500 watts (nominal – peak) 87.7 lb. (39.8kg) fully loaded; 5.12 x 18.94 x 29.72 in.

** Effective capacity assumes high availability, 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.

† Array accepts Pure Storage DirectFlash Shelf and/or Pure.

Table 2 lists the various connectivity options using both onboard and host I/O cards.

Table 2. FlashArray //X Connectivity

Two 1-/10-/25-GE	2-port 10GBASE-T Ethernet	2-port 25-/50 or 100-Gb NVMe/RoCE
Two 1-/10-/25-GE replication	2-port 1/10/25 GE	2-port 16-/32-Gb Fibre Channel (NVMe-oF Ready)
Two 1-Gb management ports	2-port 40 GE	4-port 16-/32-Gb Fibre Channel (NVMe-oF Ready)

Pure1

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

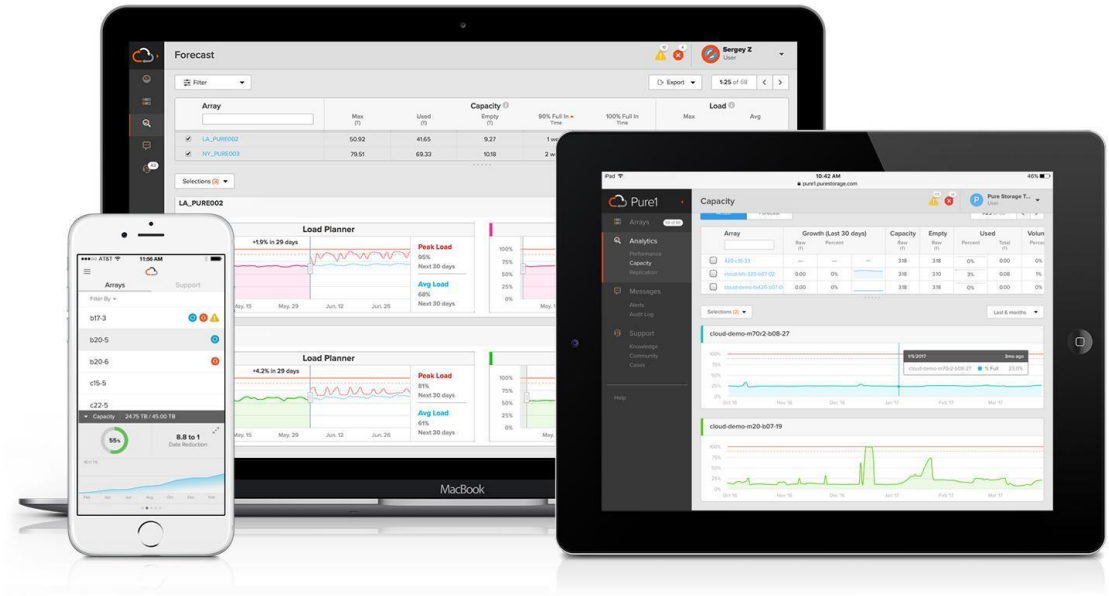

Pure1 Manage

Pure1 Manage is a SaaS-based offering that allows customers to manage their array from any browser or from the Pure1 Mobile App with nothing extra to purchase, deploy, or maintain. From a single dashboard, customers can manage all their arrays and have full storage health and performance visibility.

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 workload optimization.

Pure1 Support

Pure Storage support team with the predictive intelligence of Pure1 Meta delivers unrivaled support that's a key component in FlashArray 99.9999% availability. Some of the customer issues are identified and fixed without any customer intervention.

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's installed base, Pure1 Meta uses machine learning predictive analytics to help resolve potential issues, optimize workloads, and provide accurate forecasting. Meta is always expanding and refining what it knows about array performance and health.

VMware vSphere 7.0 U2

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 7.0 U2 has several improvements and simplifications including, but not limited to:

- VMware vSphere Virtual Volumes statistics for better debugging - track performance statistics for vSphere Virtual Volumes to quickly identify issues such as latency in third-party VASA provider responses. By using a set of commands, you can get statistics for all VASA providers in your system, or for a specified namespace or entity in the given namespace or enable statistics tracking for the complete namespace.
- vSphere Native Key Provider - Enables the use of vTPMs, vSphere Virtual Machine Encryption, and vSAN Data at Rest Encryption, when you do not require or want an external key server.
- vSphere HA support for Persistent Memory (PMEM) workloads - to deliver DRS initial placement and vSphere High Availability support for workloads that use non-volatile, persistent memory technologies.
- vSphere Lifecycle Manager fast upgrades – a replacement for VMware Update Manager, bringing a suite of capabilities to make lifecycle operations better.
- vMotion Auto Scaling - enables vSphere to automatically tune vMotion for best performance on modern, high-speed 25, 40, and 100 Gbps Ethernet networks.

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 bare metal servers, Cisco Nexus switches, Pure FlashArray storage and VMware vSphere. Using Ansible's Playbook-based automation is easy and integrates into your current provisioning infrastructure. This solution offers Ansible Playbooks that are made available from a GitHub repository that customers can access to automate the FlashStack deployment.

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 to integrate with Pure Storage FlashArray//X50 R3.

Figure 13. Cisco Intersight and vCenter and Pure Storage Integration

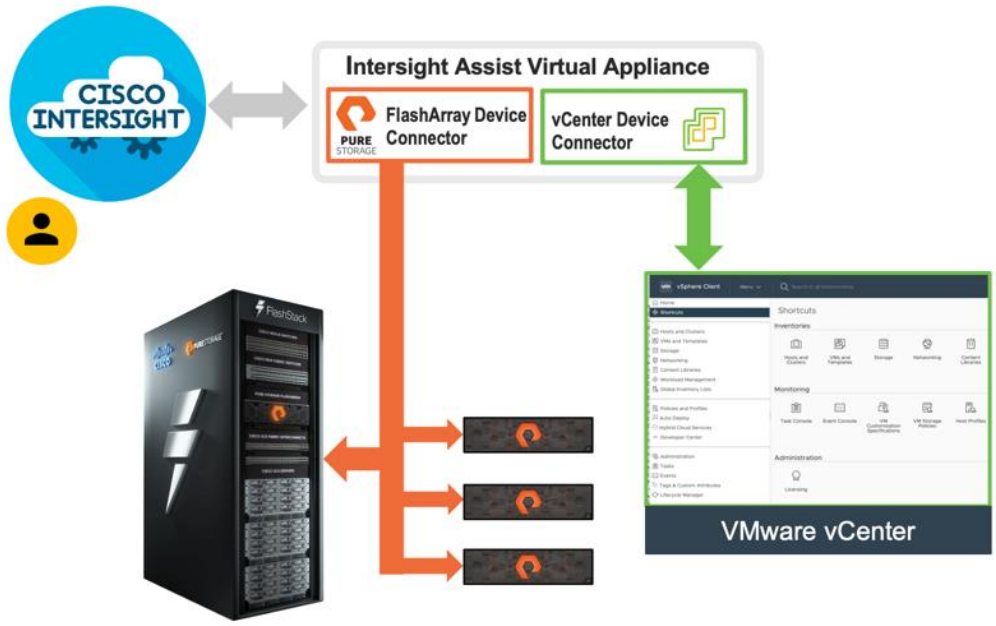

The device connector provides a safe way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure Internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware 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.

Solution Design

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

Requirements

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

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

Physical Topology

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

IP-based Storage Access

The physical topology for the IP-based FlashStack is shown in Figure 14.

Figure 14. FlashStack - Physical Topology for IP Connectivity

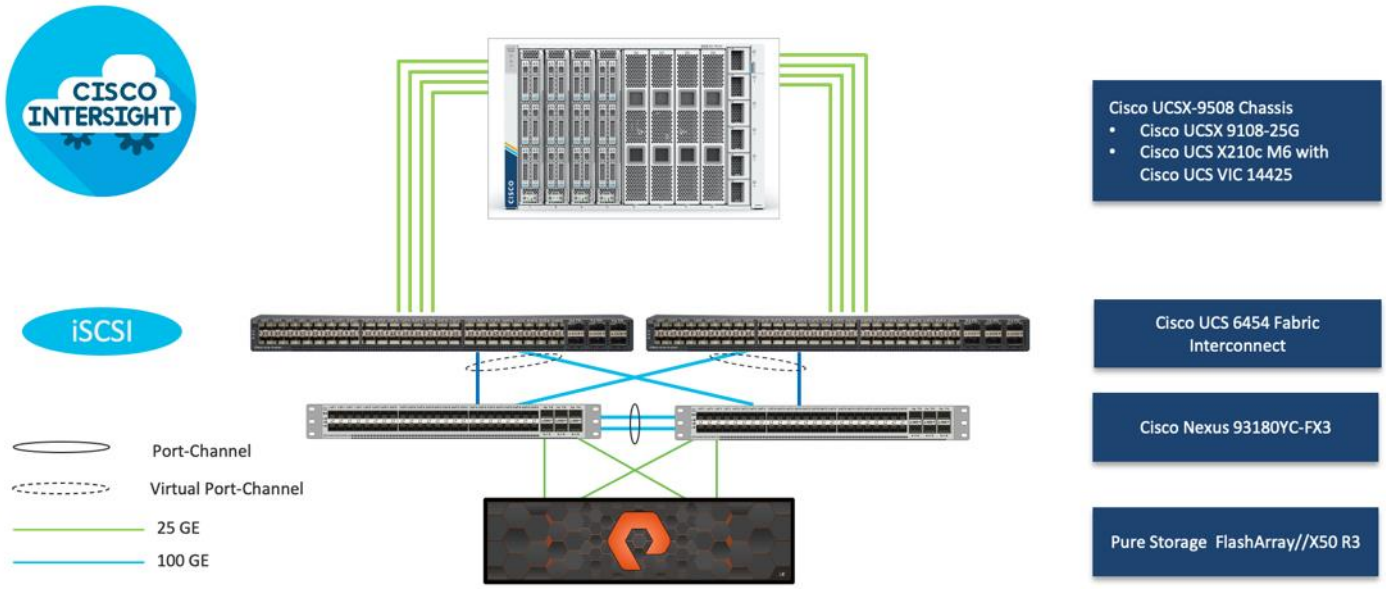

To validate the IP-based storage access in a FlashStack configuration, the components are set up as follows:

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

FC-based Storage Access

The physical topology of the FlashStack for FC connectivity is shown in Figure 15.

Figure 15. FlashStack - Physical Topology for FC Connectivity

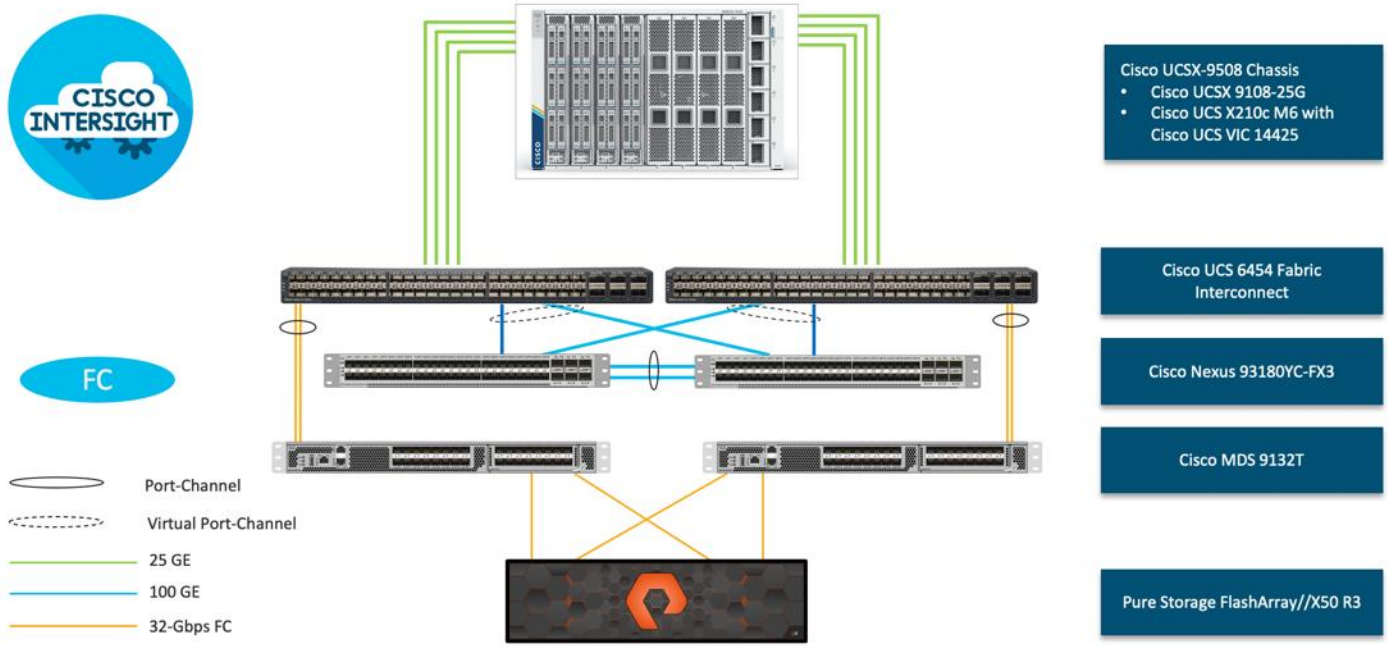

To validate the FC-based storage access in a FlashStack configuration, the components are set up as follows:

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

VLAN Configuration

Table 3 lists VLANs configured for setting up the FlashStack environment.

Table 3. VLAN Usage

2	Native-VLAN	Use VLAN 2 as native VLAN instead of default VLAN (1).
3072	OOB-MGMT-VLAN	Out-of-band management VLAN to connect management ports for various devices
19	IB-MGMT-VLAN	In-band management VLAN utilized for all in-band management connectivity - for example, ESXi hosts, VM management, and so on.
172	VM-Traffic	VM data traffic VLAN
3119*	iSCSI-A	iSCSI-A path for storage traffic including boot-from-san traffic
3219*	iSCSI-B	iSCSI-B path for storage traffic including boot-from-san traffic
3319	vMotion	VMware vMotion traffic

* iSCSI VLANs are not required if using FC storage connectivity.

Some of the key highlights of VLAN usage are as follows:

- VLAN 3072 allows customers to manage and access out-of-band management interfaces of various devices.
- VLAN 19 is used for in-band management of VMs, ESXi hosts, and other infrastructure services
- A pair of iSCSI VLANs (3119 and 3219) is configured to provide storage access including access to boot LUNs for ESXi hosts. These VLANs are not needed when configuring Fibre Channel connectivity.
- VLAN 3319 is used for VM vMotion

Logical Topology

In FlashStack deployments, each Cisco UCS server equipped with a Cisco Virtual Interface Card (VIC) is configured for multiple virtual Network Interfaces (vNICs), which appear as standards-compliant PCIe endpoints to the OS. The end-to-end logical connectivity including VLAN/VSAN usage between the server profile for an ESXi host and the storage configuration on Pure Storage FlashArray is captured in the following subsections.

Logical Topology for IP-based Storage Access

Figure 16 illustrates the end-to-end connectivity design for IP-based storage access.

Figure 16. Logical End-to-End Connectivity for iSCSI Design

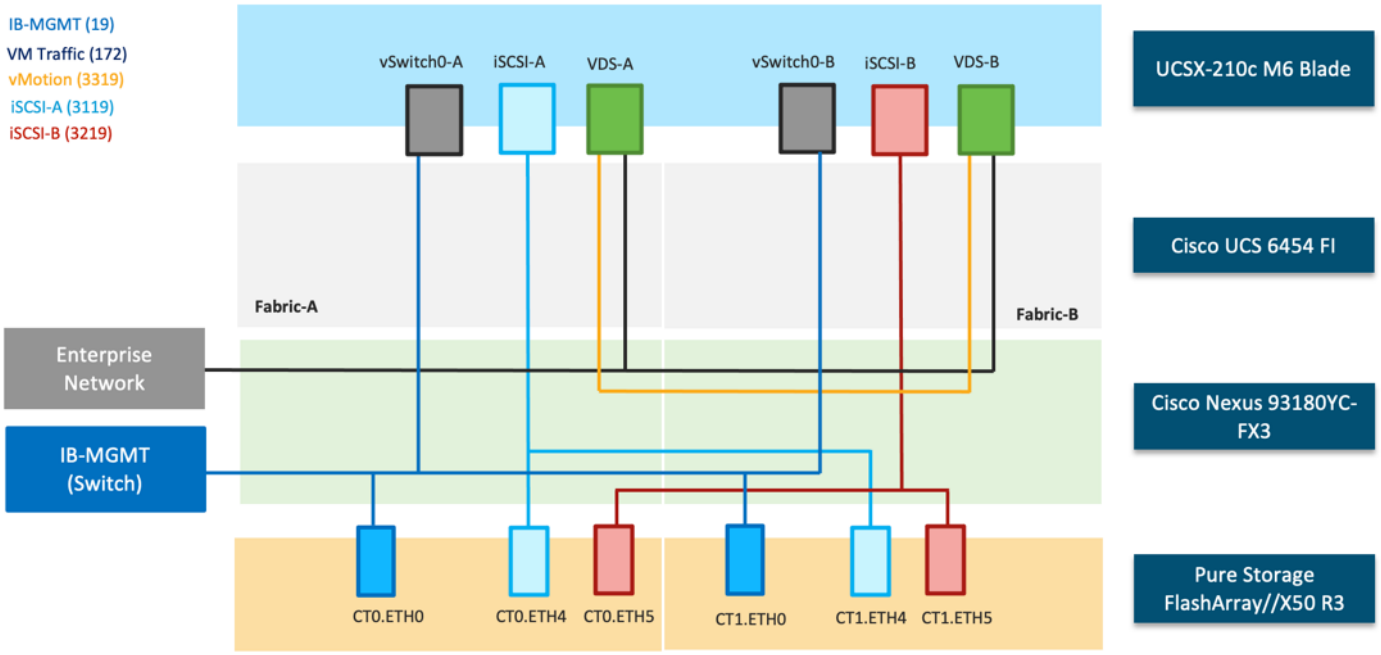

Each ESXi server profile supports:

- Managing the ESXi hosts using a common management segment
- Diskless SAN boot using iSCSI with persistent operating system installation for true stateless computing
- Six vNICs where:
  - Two redundant vNICs (vSwitch0-A and vSwitch0-B) carry management traffic. The maximum transmission unit (MTU) value for these vNICs is set as a Jumbo MTU (9000).
  - The vSphere distributed switch uses two redundant vNICs (VDS-A and VDS-B) to carry VMware vMotion traffic and customer application data traffic. The MTU for the vNICs is set to Jumbo MTU (9000).
  - The iSCSI-A vSwitch uses one iSCSI-A vNIC to provide access to the iSCSI-A path. The MTU value for the vNIC is set to Jumbo MTU (9000).
  - The iSCSI-B vSwitch uses one iSCSI-B vNIC to provide access to the iSCSI-B path. The MTU value for this vNIC is set to Jumbo MTU (9000).
- Each ESXi host (compute node) accesses datastores from Pure Storage FlashArray using iSCSI to deploy virtual machines.

Logical Topology for FC-based Storage Access

Figure 17 illustrates the end-to-end connectivity design for FC-based storage access.

Figure 17. Logical End-to-End Connectivity for FC Design

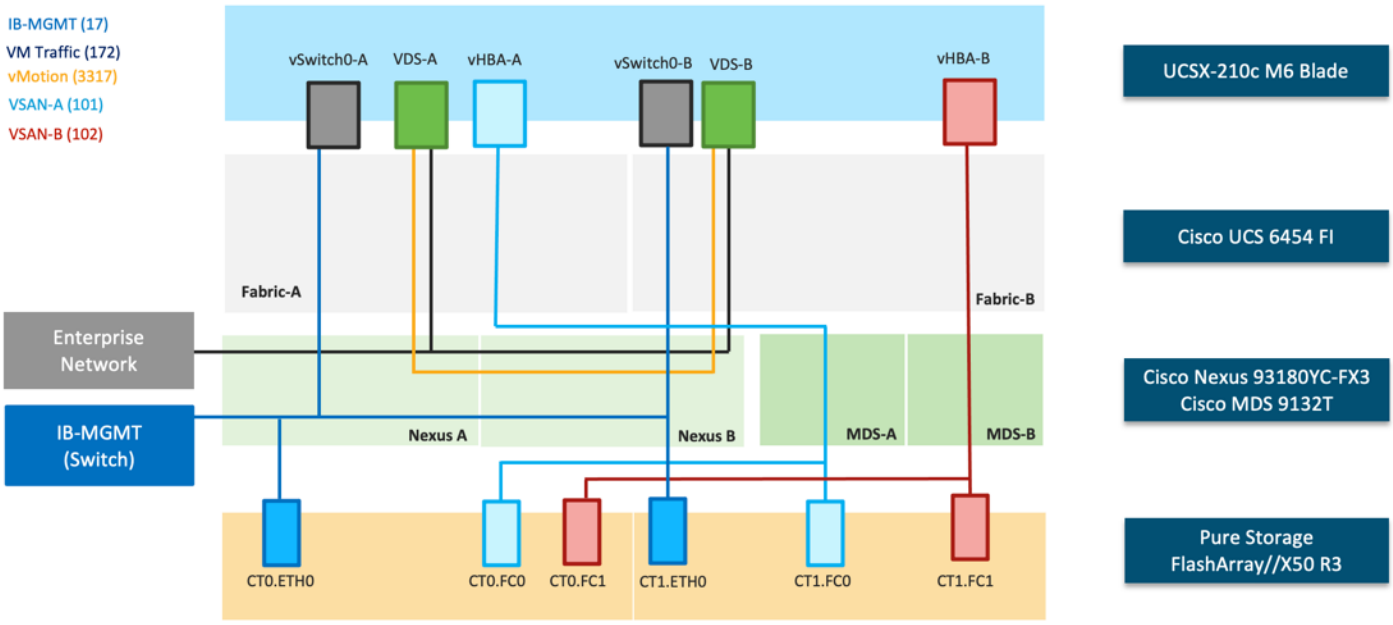

Each ESXi server profile supports:

- Managing the ESXi hosts using a common management segment
- Diskless SAN boot using Fibre Channel with persistent operating system installation for true stateless computing
- Four vNICs where:
  - Two redundant vNICs (vSwitch0-A and vSwitch0-B) carry management traffic. The MTU value for these vNICs is set as a Jumbo MTU (9000).
  - The vSphere Distributed switch uses two redundant vNICs (VDS-A and VDS-B) to carry VMware vMotion traffic and customer application data traffic. The MTU for the vNICs is set to Jumbo MTU (9000).
- Two vHBAs where:
  - One vHBA defined on Fabric A provides access to the SAN-A path.
  - One vHBA defined on Fabric B provides access to the SAN-B path.
- Each ESXi host (compute node) accesses datastores from Pure Storage FlashArray using Fibre Channel to deploy virtual machines.

Compute System Connectivity

The Cisco UCS X9508 Chassis is equipped with the Cisco UCSX 9108-25G intelligent fabric modules (IFMs). The Cisco UCS X9508 Chassis connects to each Cisco UCS 6454 FI using four 25GE ports, as shown in Figure 18. If the customers require more bandwidth, all eight ports on the IFMs can be connected to each FI.

Figure 18. Cisco UCS X9508 Chassis Connectivity to Cisco UCS Fabric Interconnects

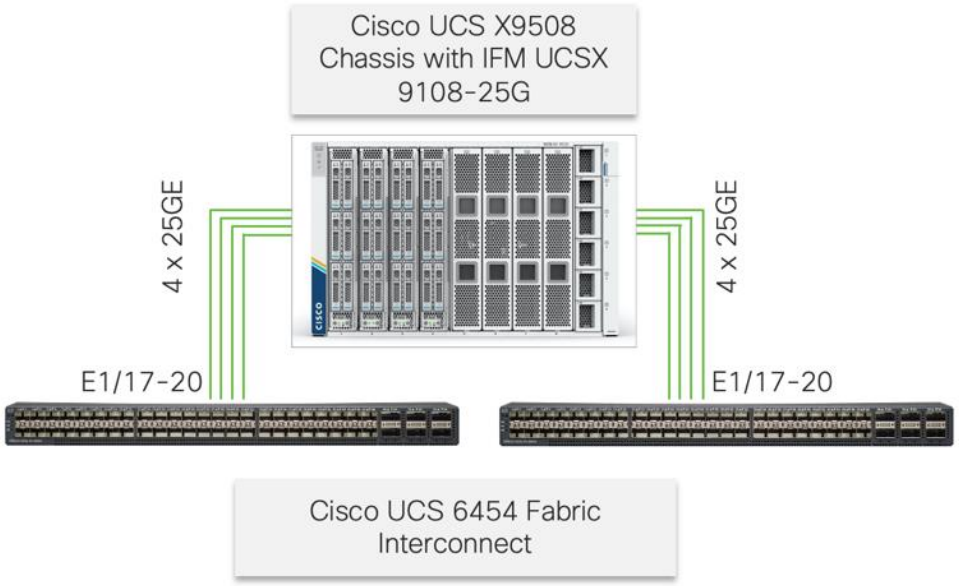

Cisco Nexus Ethernet Connectivity

The Cisco Nexus 93180YC-FX3 device configuration covers the core networking requirements for Layer 2 and Layer 3 communication. Some of the key NX-OS features implemented within the design are:

- Feature interface-vlan – Allows for VLAN IP interfaces to be configured within the switch as gateways.
- Feature HSRP – Allows for Hot Standby Routing Protocol configuration for high availability.
- Feature LACP – Allows for the utilization of Link Aggregation Control Protocol (802.3ad) by the port channels configured on the switch.
- Feature vPC – Virtual Port-Channel (vPC) presents the two Nexus switches as a single “logical” port channel to the connecting upstream or downstream device.
- Feature LLDP - Link Layer Discovery Protocol (LLDP), a vendor-neutral device discovery protocol, allows the discovery of both Cisco devices and devices from other sources.
- Feature NX-API – NX-API improves the accessibility of CLI by making it available outside of the switch by using HTTP/HTTPS. This feature helps with configuring the Cisco Nexus switch remotely using the automation framework.
- Feature UDLD – Enables unidirectional link detection for various interfaces.

Cisco UCS Fabric Interconnect 6454 Ethernet Connectivity

Cisco UCS 6454 FIs are connected to Cisco Nexus 93180YC-FX3 switches using 100GE connections configured as virtual port channels. Each FI is connected to both Cisco Nexus switches using a 100G connection; additional links can easily be added to the port channel to increase the bandwidth as needed. Figure 19 illustrates the physical connectivity details.

Figure 19. Cisco UCS 6454 FI Ethernet Connectivity

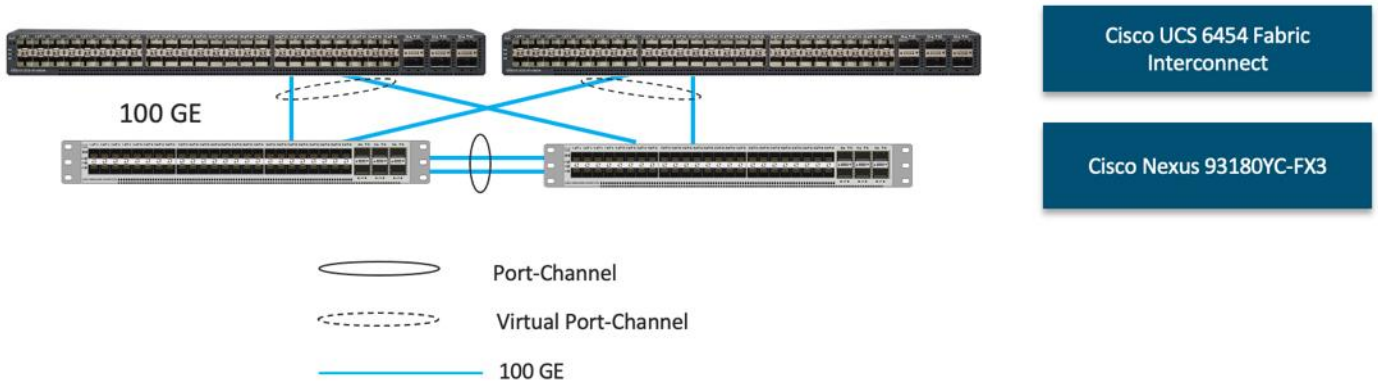

Pure Storage FlashArray//X50 R3 Ethernet Connectivity

Pure Storage FlashArray controllers are connected to Cisco Nexus 93180YC-FX3 switches using redundant 25-GE. Figure 20 illustrates the physical connectivity details.

Figure 20. Pure Storage FlashArray//X50 R3 Ethernet Connectivity

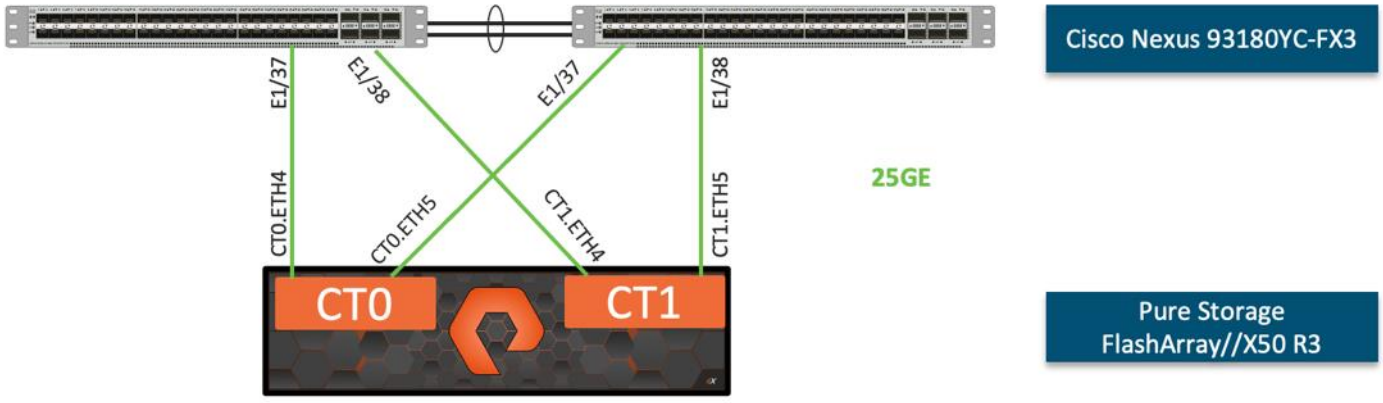

Cisco MDS SAN Connectivity – Fibre Channel Design

The Cisco MDS 9132T is the key design component bringing together the 32Gbps Fibre Channel (FC) capabilities to the FlashStack design. A redundant 32 Gbps Fibre Channel SAN configuration is deployed utilizing two MDS 9132Ts switches. Some of the key MDS features implemented within the design are:

- Feature NPIV – N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port.
- Feature fport-channel-trunk – F-port-channel-trunks allow for the fabric logins from the NPV switch to be virtualized over the port channel. This provides nondisruptive redundancy should individual member links fail.
- Smart-Zoning – a feature that reduces the number of TCAM entries by identifying the initiators and targets in the environment.

Cisco UCS Fabric Interconnect 6454 SAN Connectivity

For SAN connectivity, each Cisco UCS 6454 Fabric Interconnect is connected to a Cisco MDS 9132T SAN switch using 2 x 32G Fibre Channel port-channel connection, as shown in Figure 21.

Figure 21. Cisco UCS 6454 FI FC Connectivity

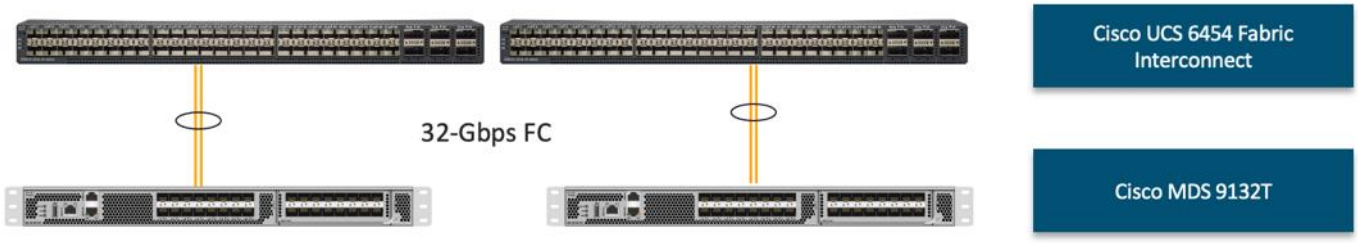

Pure Storage FlashArray//X50 R3 SAN Connectivity

For SAN connectivity, each Pure FlashArray controller is connected to both of Cisco MDS 9132T SAN switches using 32G Fibre Channel connections, as shown in Figure 22.

Figure 22. Pure Storage FlashArray FC Connectivity

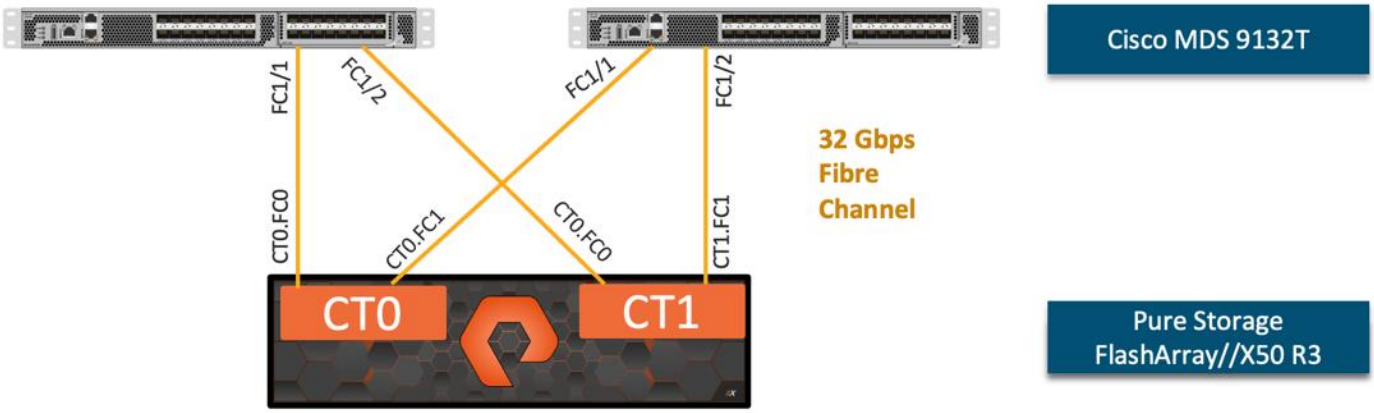

Cisco UCS X-Series Configuration - Cisco Intersight Managed Mode

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 Intersight Managed Mode consists of the steps shown in Figure 23.

Figure 23. Configuration Steps for Cisco Intersight Managed Mode

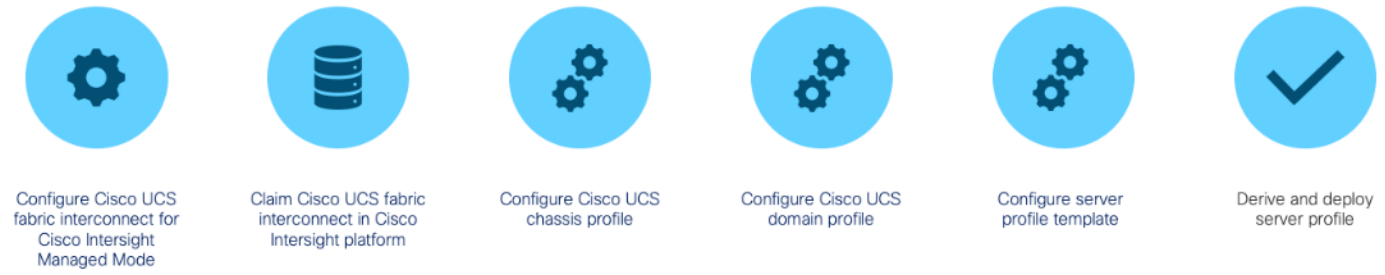

Set Up Cisco UCS Fabric Interconnect for Cisco Intersight Managed Mode

During the initial configuration, for the management mode the configuration wizard enables customers to choose whether to manage the fabric interconnect through Cisco UCS Manager or the Cisco Intersight platform. Customers can switch the management mode for the fabric interconnects between Cisco Intersight and Cisco UCS Manager at any time; however, Cisco UCS FIs must be set up in Intersight Managed Mode (IMM) for configuring the Cisco UCS X-Series system. Figure 24 shows the dialog during initial configuration of Cisco UCS FIs for setting up IMM.

Figure 24. Fabric Interconnect Setup for Cisco Intersight Managed Mode

```
UCSM image signature verification successful

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

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

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


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

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

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

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

Claim a Cisco UCS Fabric Interconnect in the Cisco Intersight Platform

After setting up the Cisco UCS Fabric Interconnect for Cisco Intersight Managed Mode, FIs can be claimed to a new or an existing Cisco Intersight account. When a Cisco UCS Fabric Interconnect is successfully added to Cisco Intersight, all future configuration steps are completed in the Cisco Intersight portal.

Figure 25. Cisco Intersight: Adding Fabric Interconnects

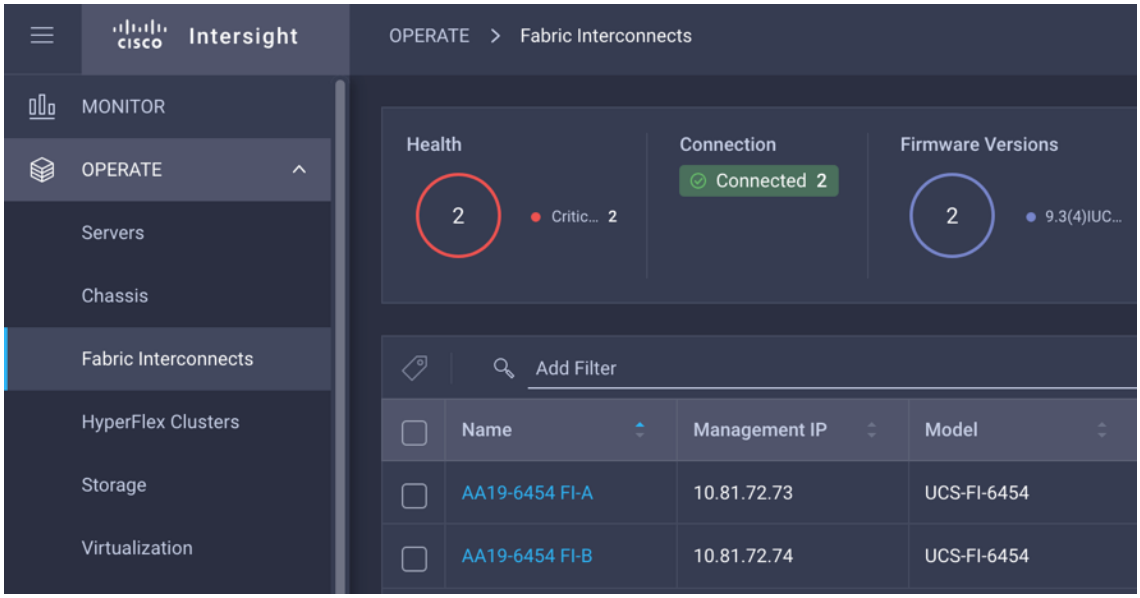

Customers can verify whether a Cisco UCS Fabric Interconnect is in Cisco UCS Manager managed mode or Cisco Intersight Managed Mode by clicking on the fabric interconnect name and looking at the detailed information screen for the FI, as shown in Figure 26.

Figure 26. Cisco UCS FI in Intersight Managed Mode

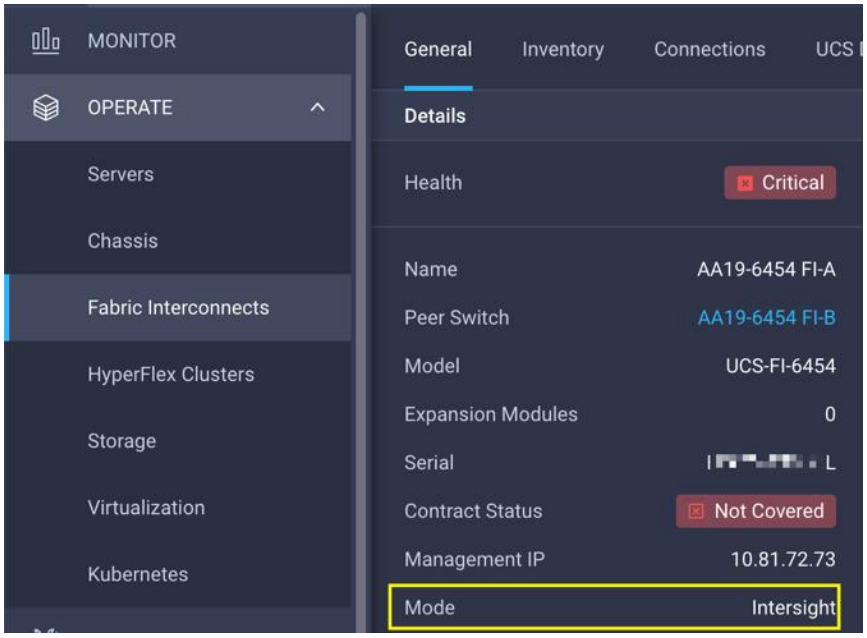

Cisco UCS Chassis Profile

A Cisco UCS Chassis profile configures and associates the chassis policy to a Cisco UCS chassis. The chassis profile feature is available in Intersight only if customers have installed the Intersight Essentials License. The chassis-related policies can be attached to the profile either at the time of creation or later.

The chassis profile in a FlashStack is used to set the power policy for the chassis. By default, UCS X-Series power supplies are configured in GRID mode, but power policy can be utilized to set the power supplies in non-redundant or N+1/N+2 redundant modes.

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 to be used in the network. It defines the characteristics of and configures the ports on the fabric interconnects. One Cisco UCS domain profile can be assigned to one fabric interconnect domain.

Some of the characteristics of the Cisco UCS domain profile in the FlashStack environment are:

- A single domain profile is created for the pair of Cisco UCS fabric interconnects.
- Unique port policies are defined for the two fabric interconnects.
- The VLAN configuration policy is common to the fabric interconnect pair because both fabric interconnects are configured for the same set of VLANs.
- The VSAN configuration policies (FC connectivity option) are unique for the two fabric interconnects because the VSANs are unique.
- The Network Time Protocol (NTP), network connectivity, and system Quality-of-Service (QoS) policies are common to the fabric interconnect pair.

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.

Figure 27. Cisco UCS Domain Profile

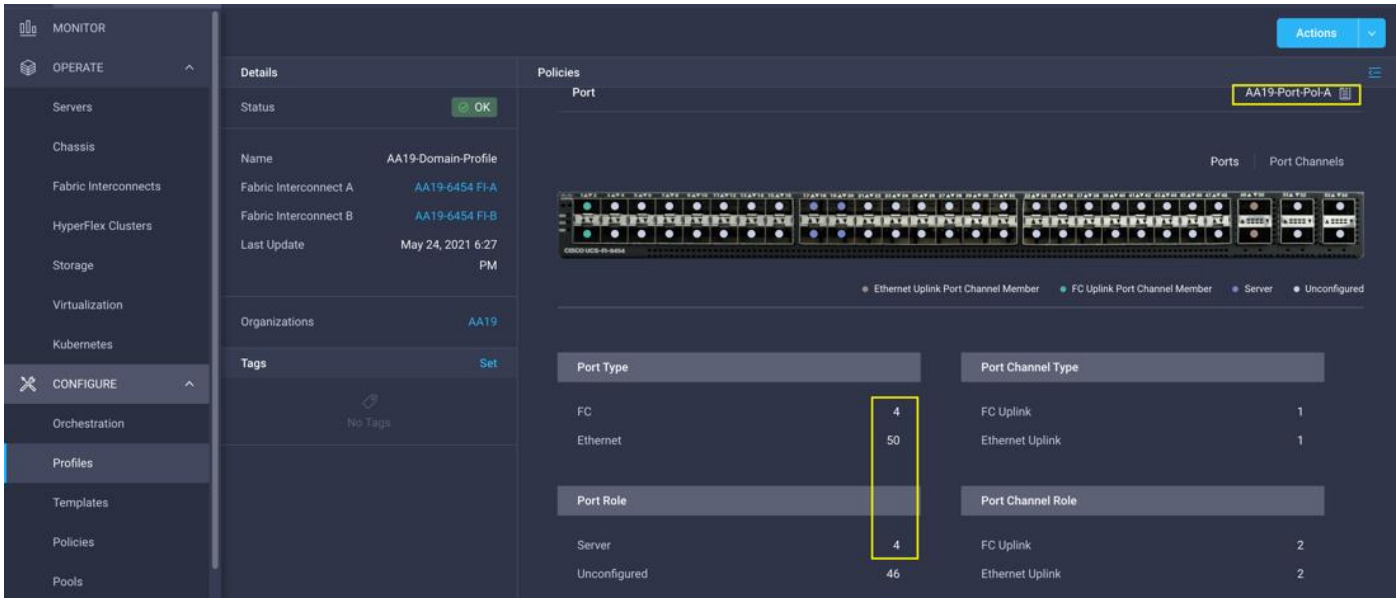

The Cisco UCS X9508 Chassis and Cisco UCS X210c M6 Compute Nodes are automatically discovered when the ports are successfully configured using the domain profile as shown in Figure 28, Figure 29, and Figure 30.

Figure 28. Cisco UCS X9508 Chassis Front View

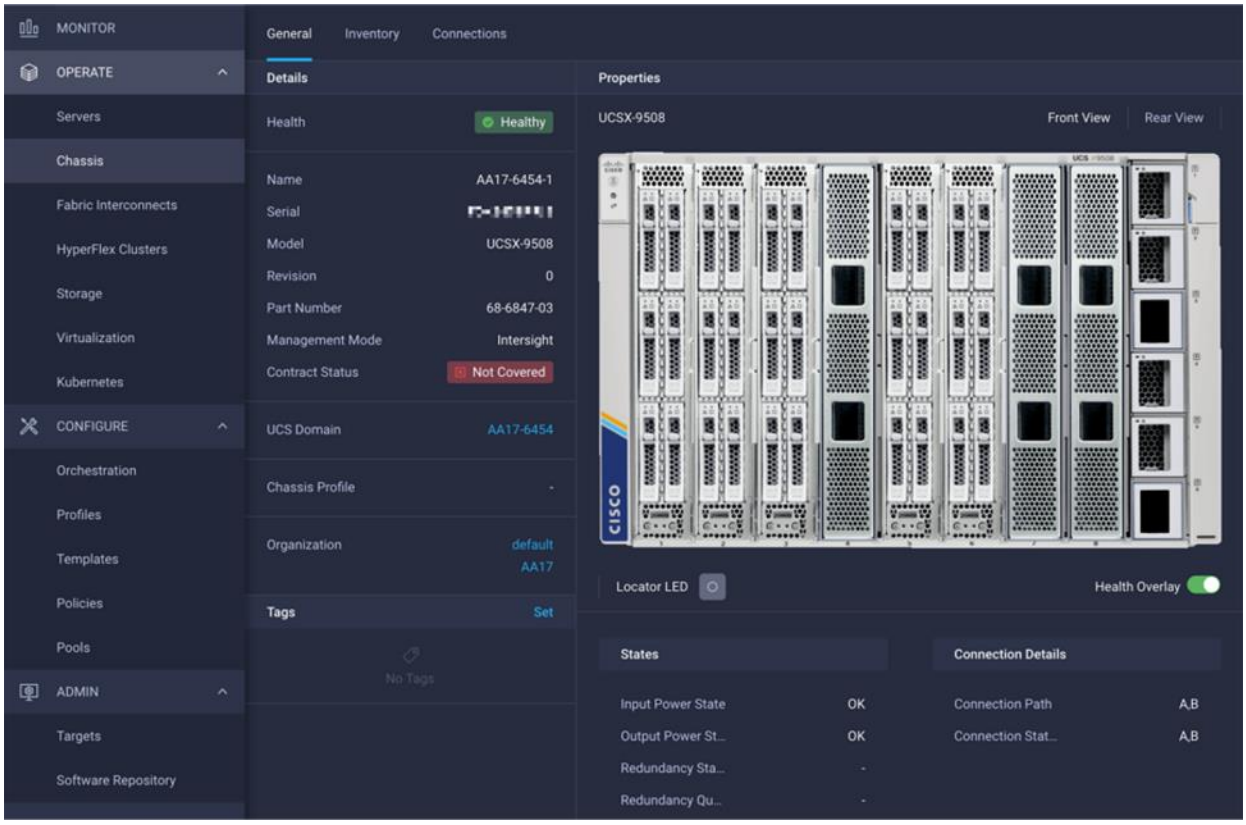

Figure 29. Cisco UCS X9508 Chassis Rear View

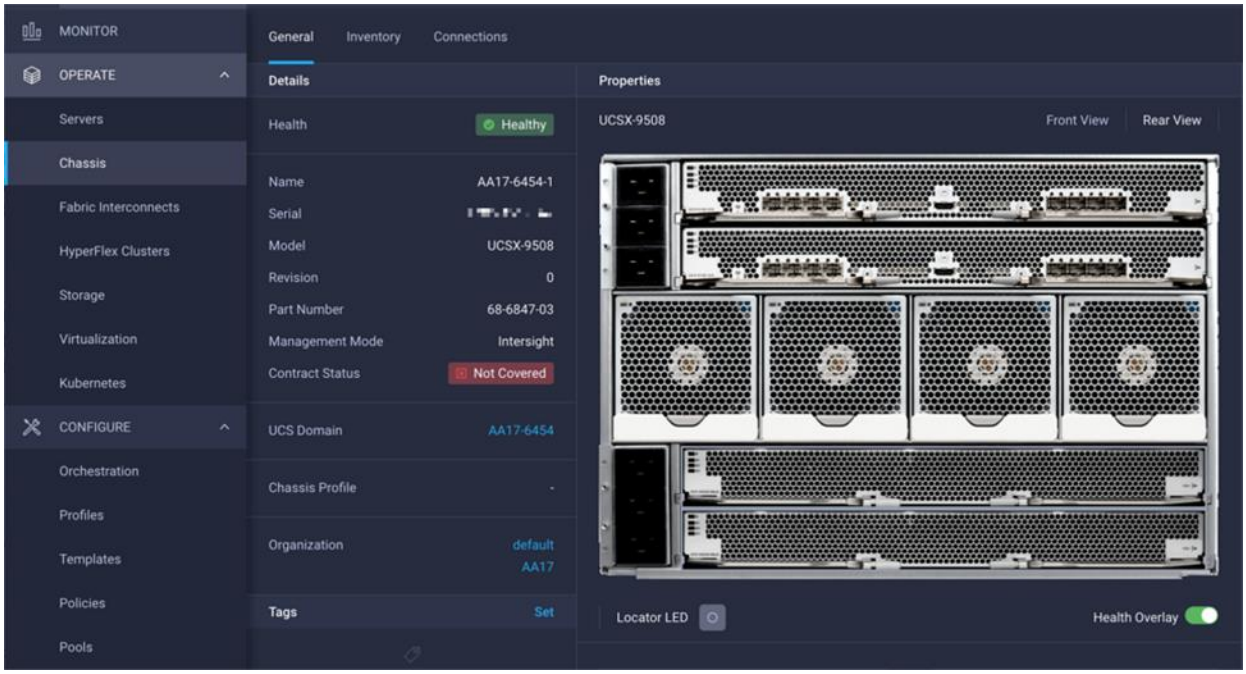

Figure 30. Cisco UCS X210c M6 Compute Nodes

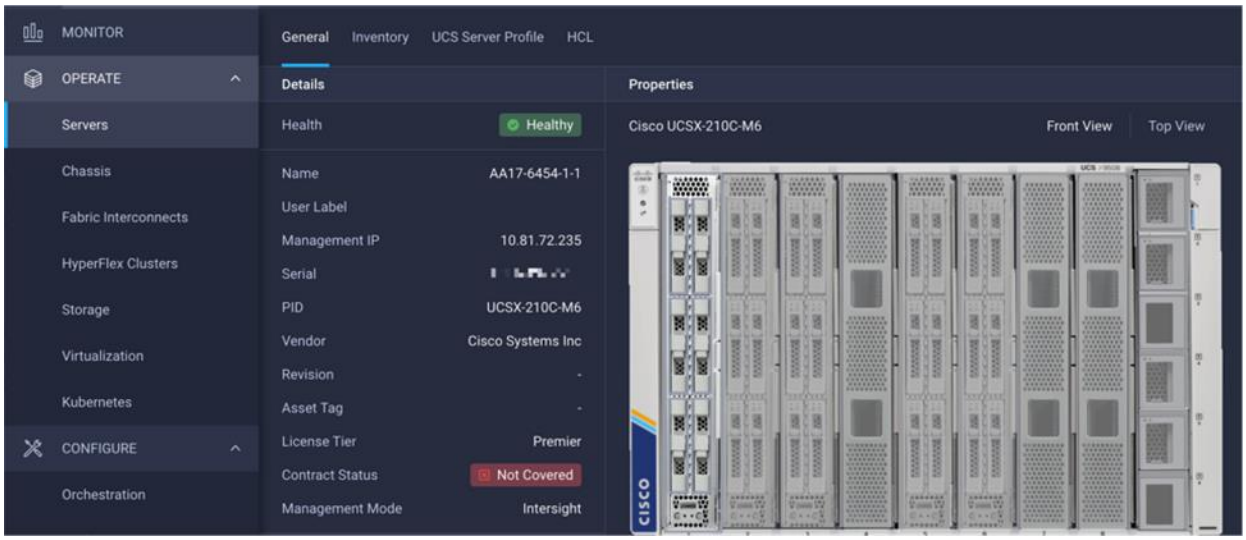

Server Profile Template

A server profile template enables resource management by simplifying policy alignment and server configuration. A server profile template is created using the server profile template wizard. The server profile template wizard groups the server policies into the following four categories to provide a quick summary view of the policies that are attached to a profile:

- Compute policies: BIOS, boot order, and virtual media policies
- Network policies: adapter configuration, LAN connectivity, and SAN connectivity policies

    - The LAN connectivity policy requires you to create Ethernet network policy, Ethernet adapter policy, and Ethernet QoS policy.
    - The SAN connectivity policy requires you to create Fibre Channel (FC) network policy, Fibre Channel adapter policy, and Fibre Channel QoS policy. SAN connectivity policy is only required for the FC connectivity option.
- Storage policies configure local storage and are not used in FlashStack
- Management policies: device connector, Intelligent Platform Management Interface (IPMI) over LAN, Lightweight Directory Access Protocol (LDAP), local user, network connectivity, Simple Mail Transfer Protocol (SMTP), Simple Network Management Protocol (SNMP), Secure Shell (SSH), Serial over LAN (SOL), syslog, and virtual Keyboard, Video, and Mouse (KVM) policies

Some of the characteristics of the server profile template for FlashStack are:

- BIOS policy is created to specify various server parameters in accordance with FlashStack best practices.
- Boot order policy defines virtual media (KVM mapper DVD), all SAN paths for Pure Storage FlashArray (iSCSI or Fibre Channel interfaces), and UEFI Shell.
- IMC access policy defines the management IP address pool for KVM access.
- Local user policy is used to enable KVM-based user access.
- For the iSCSI boot from SAN configuration, LAN connectivity policy is used to create six virtual network interface cards (vNICs) – two for management virtual switch (vSwitch0), two for application Virtual Distributed Switch (VDS), and one each for iSCSI A/B vSwitches. Various policies and pools are also created for the vNIC configuration.
- For the FC boot from SAN configuration, LAN connectivity policy is used to create four vNICs – two for management virtual switches (vSwitch0) and two for application VDS – along with various policies and pools.
- For the FC connectivity option, SAN connectivity policy is used to create two virtual host bus adapters (vHBAs) – one for SAN A and one for SAN B – along with various policies and pools. The SAN connectivity policy is not required for iSCSI setup.

Figure 31 shows various policies associated with the server profile template.

Figure 31. Server Profile Template for iSCSI Boot from SAN

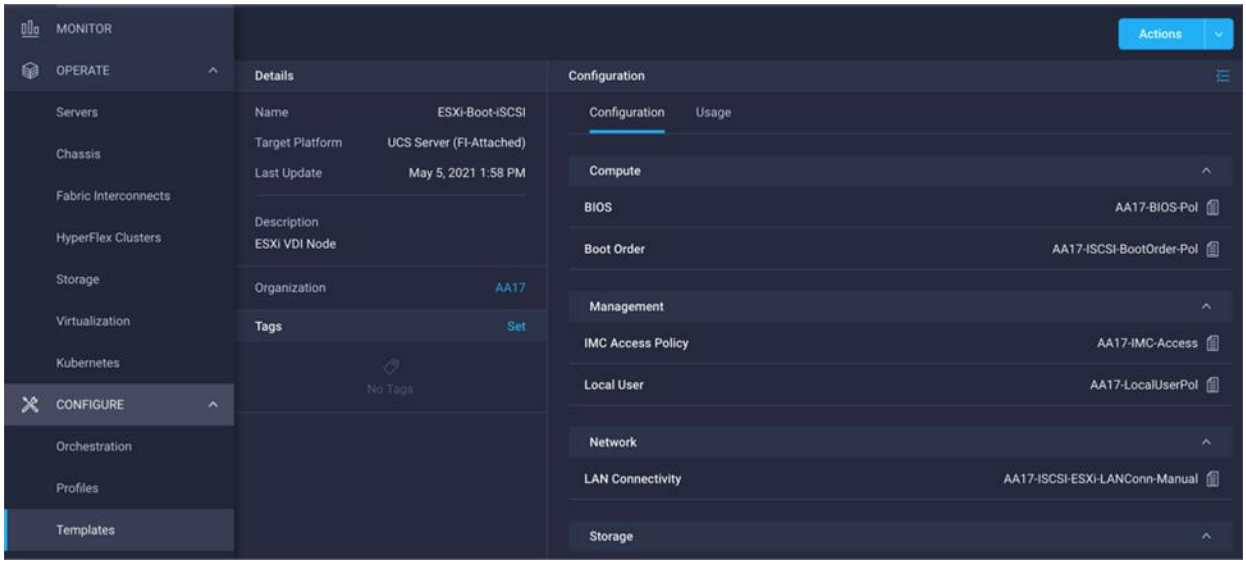

Derive and Deploy Server Profiles from the Cisco Intersight Server Profile Template

The Cisco Intersight server profile allows server configurations to be deployed directly on the compute nodes based on polices defined in the server profile template. After a server profile template has been successfully created, server profiles can be derived from the template and associated with the Cisco UCS X210c M6 Compute Nodes, as shown in Figure 32.

Figure 32. Deriving a Server Profile from Templates

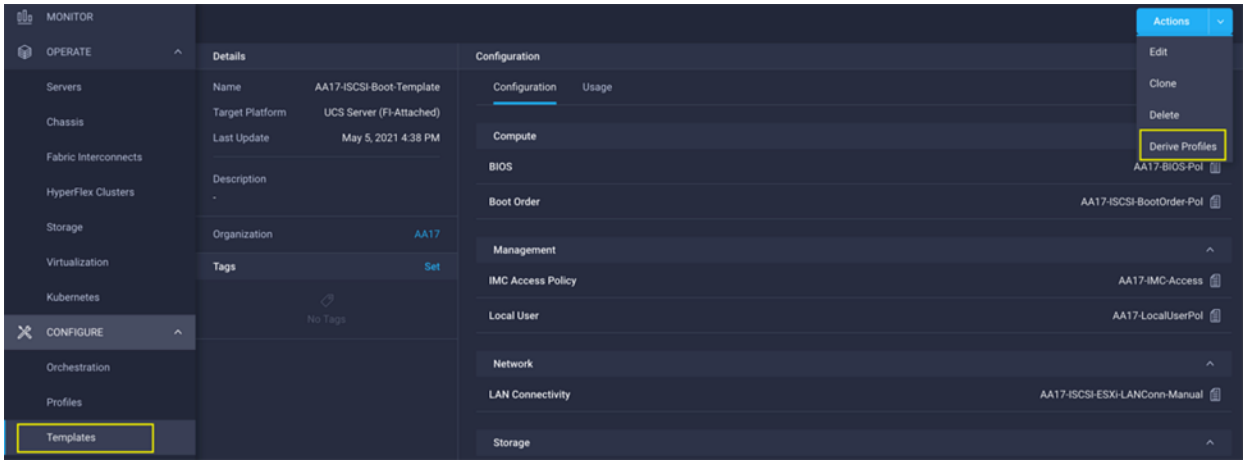

On successful deployment of the server profile, the Cisco UCS X210c M6 Compute Nodes are configured with parameters defined in the server profile and can boot from the storage LUN hosted on Pure Storage FlashArray.

Pure Storage FlashArray – Storage Design

To set up Pure Storage FlashArray customers must configure the following items:

- Volumes

- ESXi boot LUNs: These LUNs enable ESXi host boot from SAN functionality using iSCSI or Fibre Channel.
- The vSphere environment: vSphere uses the infrastructure datastore(s) to store the virtual machines.

- Hosts
  - All FlashArray ESXi hosts are defined.
  - Add every active initiator for a given ESXi host.
- Host groups
  - All ESXi hosts in a VMware cluster are part of the host group.
  - Host groups are used to mount VM infrastructure datastores in the VMware environment.

The volumes, interfaces, and VLAN/VSAN details are shown in Figure 33 and Figure 34 for iSCSI and Fibre Channel connectivity, respectively.

Figure 33. Pure Storage FlashArray Volumes and Interfaces – iSCSI Configuration

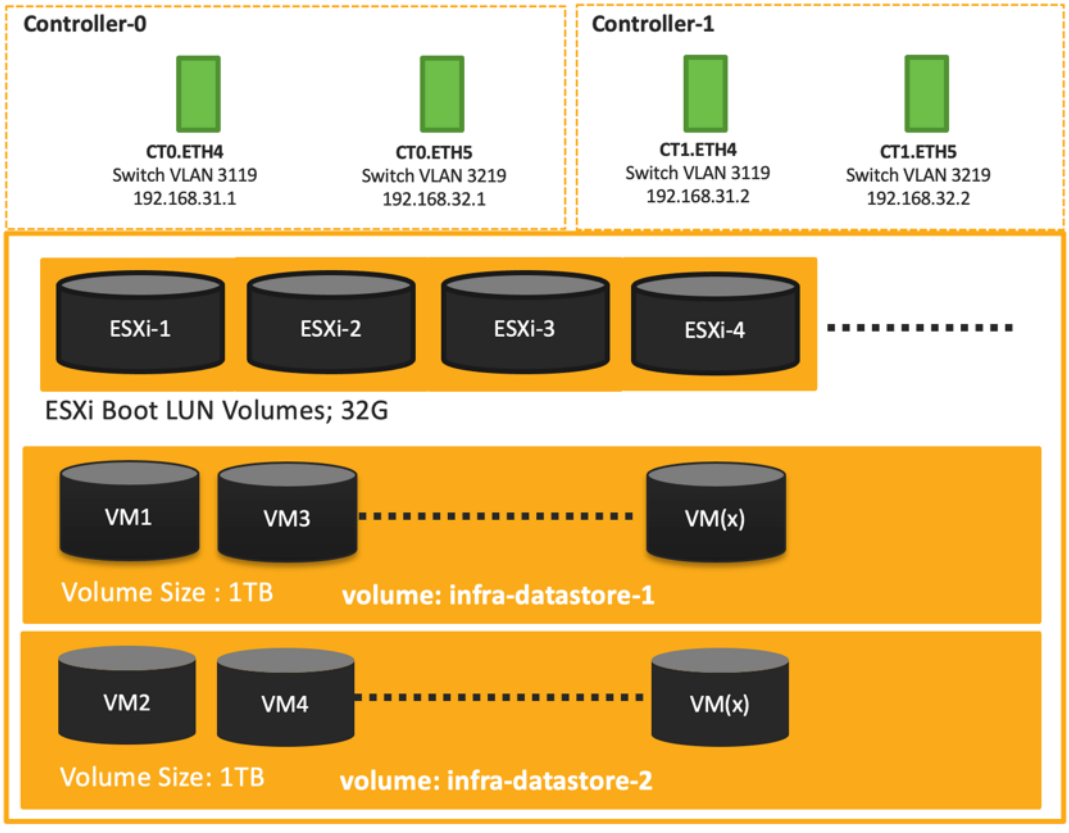

Figure 34. Pure Storage FlashArray Volumes and Interfaces - Fibre Channel Configuration

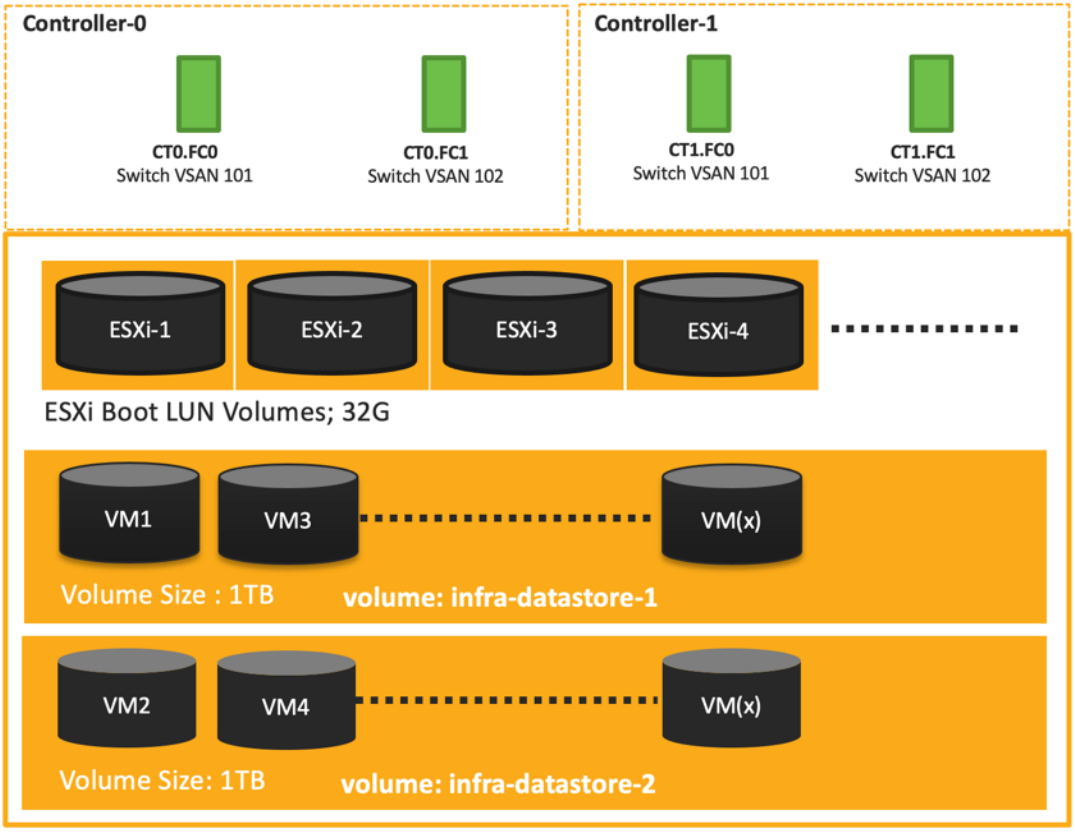

VMware vSphere - ESXi Design

Multiple vNICs (and vHBAs) are created for the ESXi hosts using the Cisco Intersight server profile and are then assigned to specific virtual and distributed switches. The vNIC and (optional) vHBA distribution for the ESXi hosts is as follows:

- Two vNICs (one on each fabric) for vSwitch0 to support core services such as management traffic.
- Two vNICs (one on each fabric) for vSphere Virtual Distributed Switch (VDS) to support customer data traffic and vMotion traffic.
- One vNIC each for Fabric-A and Fabric-B for iSCSI stateless boot. These vNICs are only required when iSCSI boot from SAN configuration is desired.
- One vHBA each for Fabric-A and Fabric-B for FC stateless boot. These vHBAs are only required when FC connectivity is desired.

Typically, customers will either have iSCSI vNICs for IP based storage access or the FC vHBAs for Fibre Channel SAN connectivity.

Figure 35 and Figure 36 show the ESXi vNIC configurations in detail.

Figure 35. VMware vSphere – ESXi Host Networking for iSCSI Boot from SAN

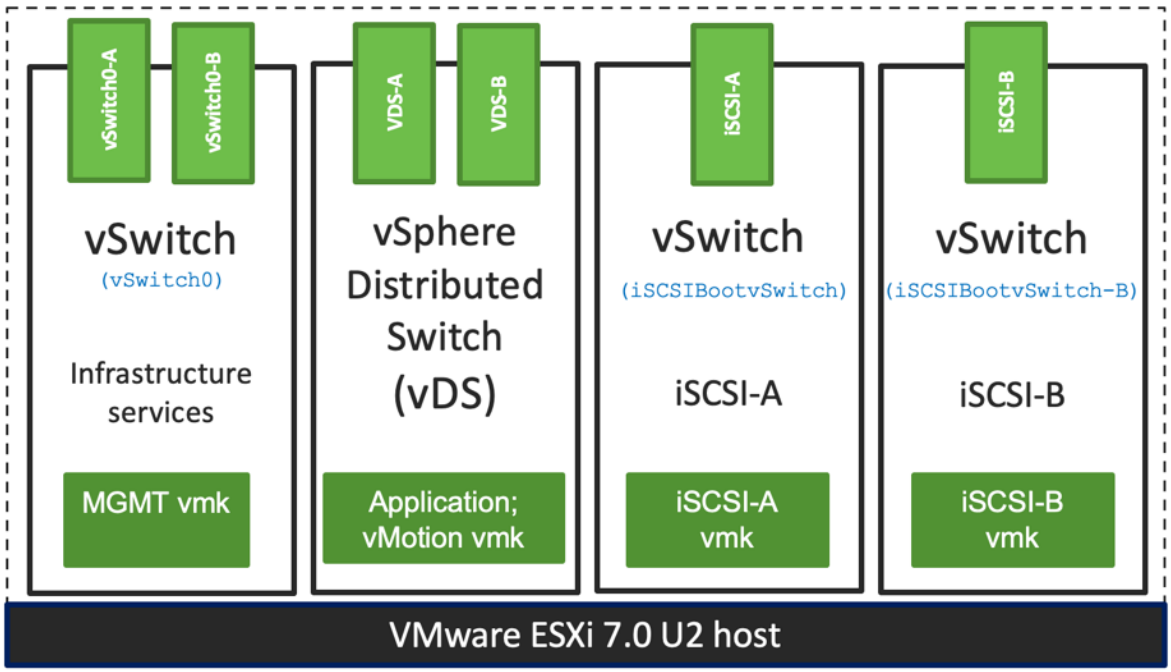

Figure 36. VMware vSphere – ESXi Host Networking for FC Boot from SAN

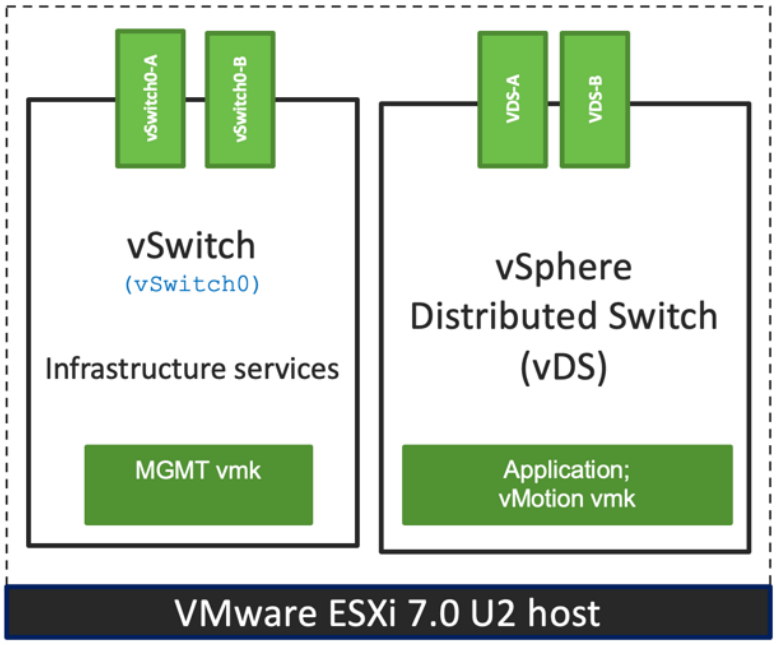

Cisco Intersight Integration with VMware vCenter and Pure Storage FlashArray

Cisco Intersight works with Pure Storage FlashArray and VMware vCenter using third-party device connectors. Since third-party infrastructure does not contain any built-in Intersight device connector, Cisco Intersight Assist virtual appliance enables Cisco Intersight to communicate with these non-Cisco devices.

A single Cisco Intersight Assist virtual appliance can support both Pure Storage FlashArray and VMware vCenter.

Cisco Intersight integration with VMware vCenter and Pure Storage FlashArray enables customers to perform following tasks right from the Intersight dashboard:

- Monitor the virtualization and storage environment.
- Add various dashboard widgets to obtain useful at-a-glance information.
- Perform common Virtual Machine tasks such as power on/off, remote console and so on.
- Orchestrate virtual and storage environment to perform common configuration tasks.

The following sections explain the details of these operations. Since Cisco Intersight is a SaaS platform, the monitoring and orchestration capabilities are constantly being added and delivered seamlessly from the cloud.

The monitoring capabilities and orchestration tasks and workflows listed below provide an in-time snapshot for your reference. For the most up to date list of capabilities and features, customers should use the help and search capabilities in Cisco Intersight.

Figure 37. Managing Pure Storage FlashArray and VMware vCenter through Cisco Intersight using Intersight Assist

Licensing Requirement

To integrate and view various Pure Storage FlashArray and VMware vCenter parameters from Cisco Intersight, a Cisco Intersight Advantage license is required. To use Cisco Intersight orchestration and workflows to provision the storage and virtual environments, an Intersight Premier license is needed.

Integrate Cisco Intersight with Pure Storage FlashArray

To integrate Pure Storage FlashArray with the Cisco Intersight platform, you must deploy a Cisco Intersight Assist virtual appliance and claim Pure Storage FlashArray as a target in the Cisco Intersight application, as shown in Figure 38.

Figure 38. Claiming Pure Storage FlashArray as a Target in Cisco Intersight

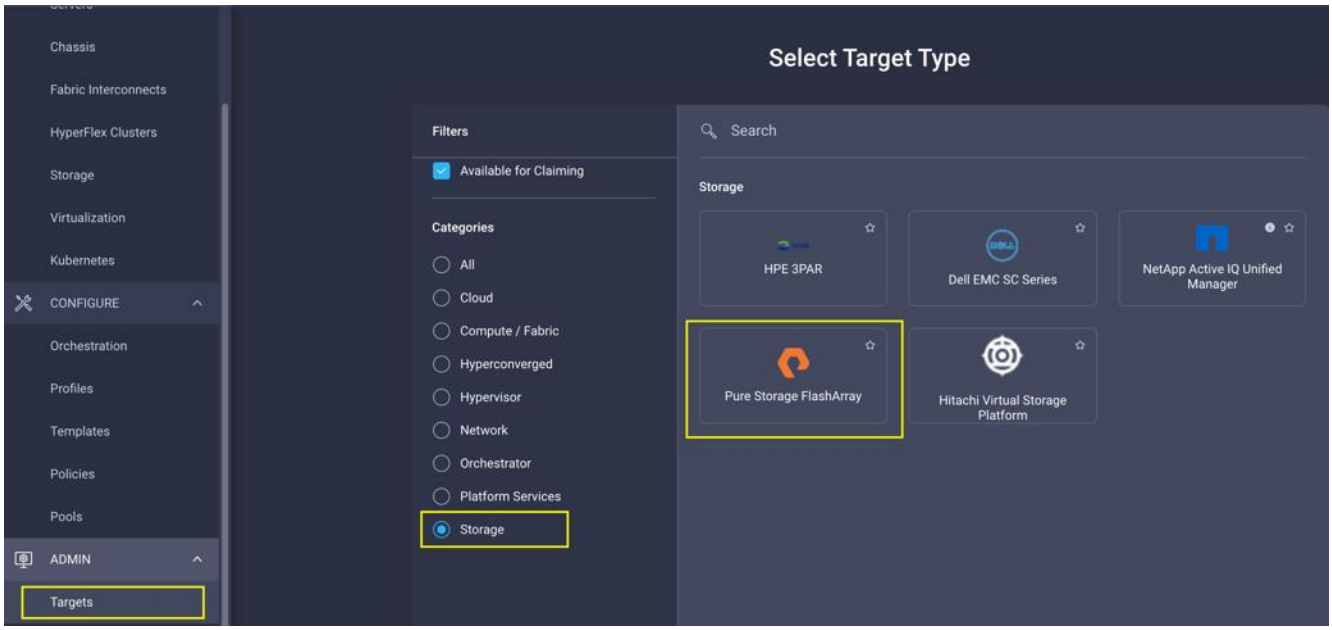

Obtain Storage-level Information

After successfully claiming Pure Storage FlashArray as a target, customers can view storage-level information in Cisco Intersight.

Figure 39. Pure Storage FlashArray Information in Cisco Intersight

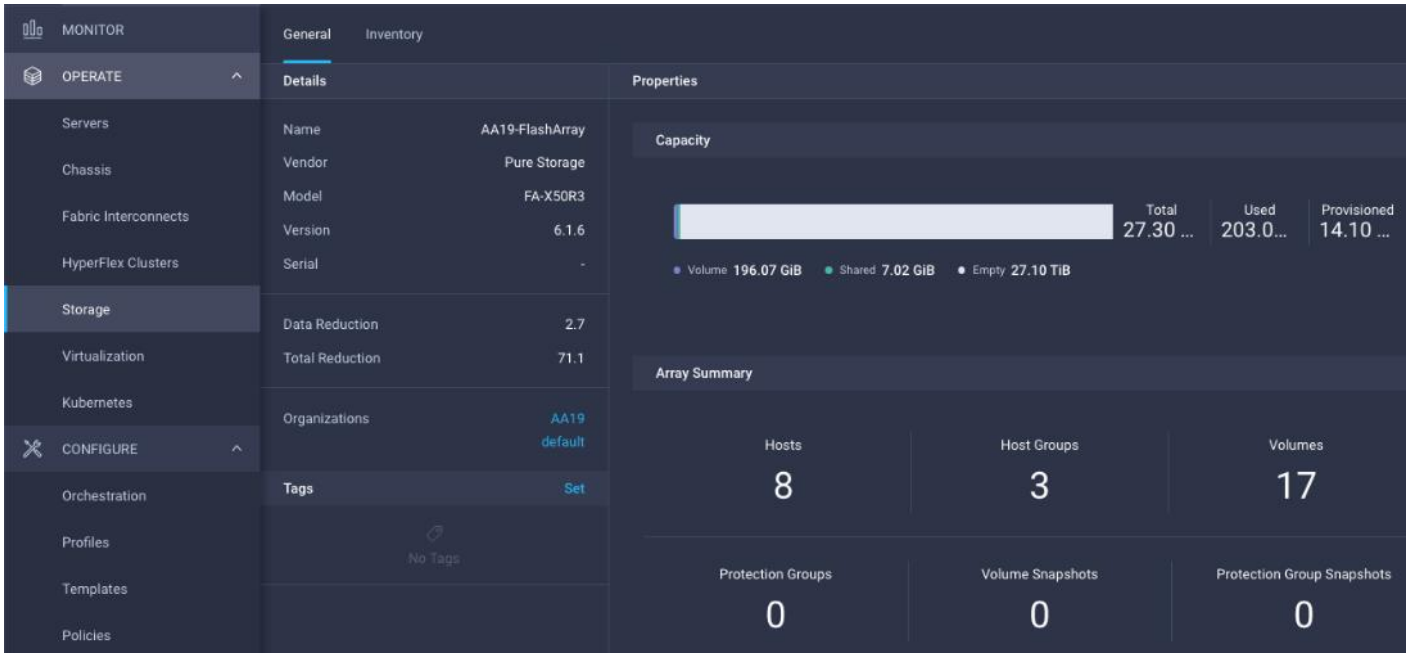

Table 4 lists some of the Pure Storage FlashArray information presented through Cisco Intersight.

Table 4. Pure Storage FlashArray Information in Cisco Intersight

General	Name	Name of the controller
	Vendor	Pure Storage
	Model	Pure Storage FlashArray model information (for example, FA-X50R3)
	Version	Software version
	Serial	Serial number
	Data reduction	Storage efficiency
	Total reduction	Storage efficiency
Monitoring	Capacity	Total, used, and provisioned system capacity
	Array	Summary of hosts, host groups, volumes, and so on, in the system
Inventory	Hosts	Hosts defined in the system and associated ports, volumes, and protection of group information
	Host groups	Host groups defined in the system and associated hosts, volumes, and protection of groups in the system
	Volumes	Configured volumes and volume-specific information such as capacity, data reduction, and so on.

	Protection group	Protection groups defined in the system and associated targets, members, etc.
	Controllers	FlashArray controllers and their state, version, and model information
	Drives	Storage drive-related information, including type and capacity information
	Ports	Information related to physical ports, including World Wide Port Name (WWPN) and iSCSI Qualified Name (IQN) information

Storage Widget in the Dashboard

Customers can also add the storage dashboard widgets to Cisco Intersight for viewing Pure Storage FlashArray at a glance information on the Cisco Intersight dashboard, as shown in Figure 40.

Figure 40. Storage Widgets in Cisco Intersight Dashboard

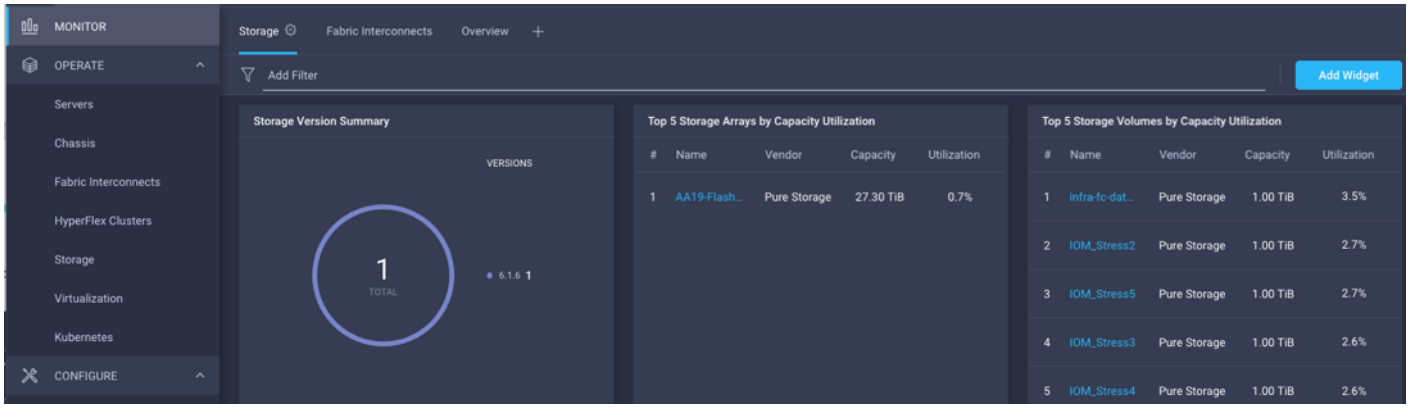

These storage widgets provide useful information such as:

- Storage versions summary, providing information about the software version and the number of storage systems running that version
- Storage arrays and capacity utilization
- Top-five storage volumes by capacity utilization

Cisco Intersight Orchestrator – Pure Storage FlashArray

Cisco Intersight Orchestrator provides various workflows that can be used to automate storage provisioning. Some of the sample storage workflows available for Pure Storage FlashArray are listed in Table 5.

Table 5. Pure Storage FlashArray Workflows in Cisco Intersight Orchestrator

New storage host	Create a new storage host; if a host group is provided as input, then the host is added to the host group.
New storage host group	Create a new storage host group; if hosts are provided as inputs, the workflow will add the hosts to the host group.

New VMFS Datastore	Create a storage volume and build a VMFS datastore on the volume.
Remove storage host	Remove a storage host. If a host group name is provided as input, the workflow will also remove the host from the host group.
Remove storage host group	Remove a storage host group. If hosts are provided as input, the workflow will remove the hosts from the host group.
Remove VMFS datastore	Remove a VMFS datastore and remove the backing volume from the storage device.
Update storage host	Update the storage host details. If the inputs for a task are provided, then the task is run; otherwise, it is skipped.
Update VMFS datastore	Expand a datastore on the hypervisor manager by extending the backing storage volume to specified capacity, and then expanding the data store to use the additional capacity.

In addition to the above workflows, Cisco Intersight Orchestrator also provides many storage and virtualization tasks for customers to create custom workflow based on their specific needs. A sample subset of these tasks is highlighted in Figure 41.

Figure 41. Storage Tasks for Pure Storage FlashArray

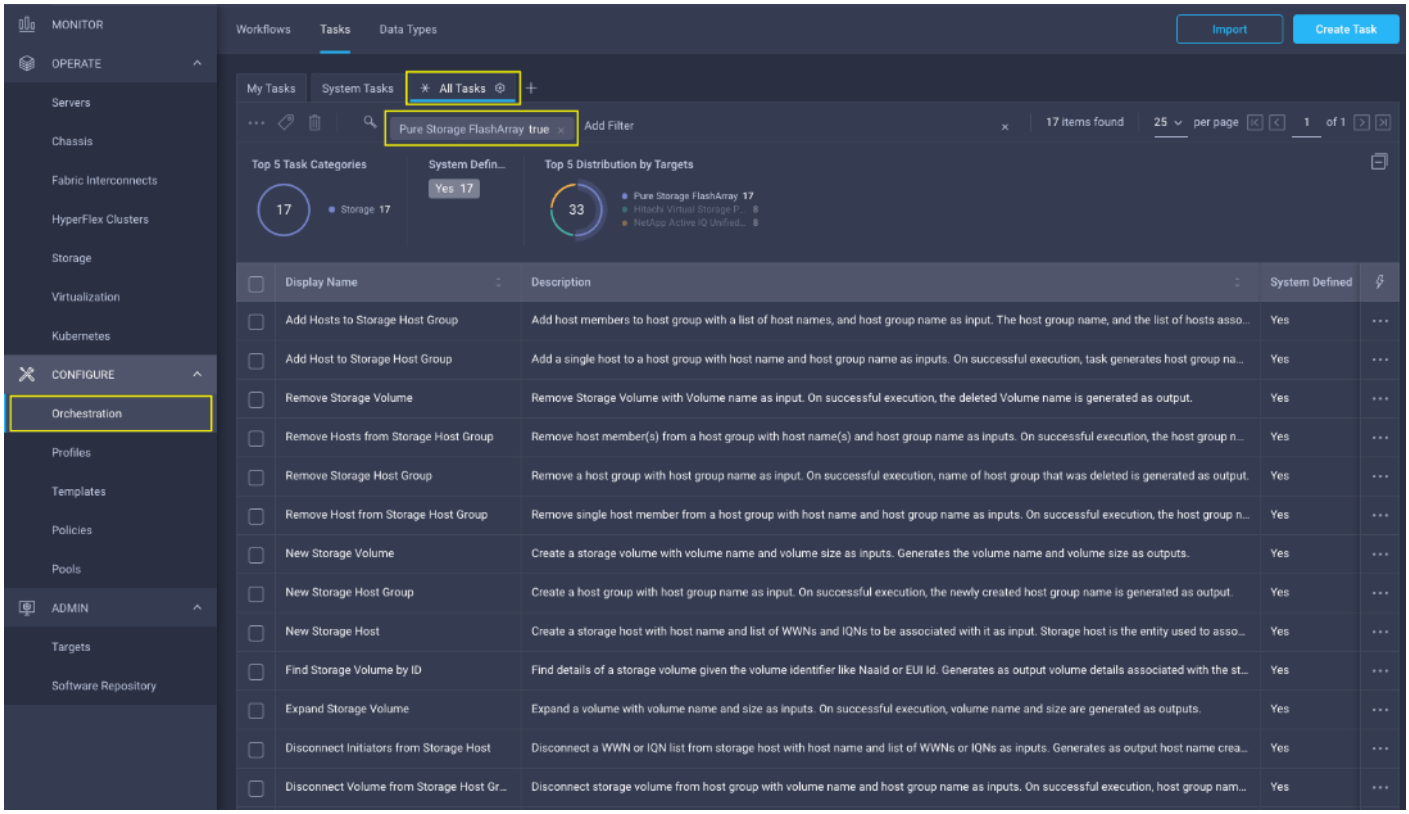

Integrate Cisco Intersight with VMware vCenter

To integrate VMware vCenter with Cisco Intersight, VMware vCenter can be claimed as a target using Cisco Intersight Assist Virtual Appliance, as shown in Figure 42.

Figure 42. Claim VMware vCenter in Cisco Intersight as a Target

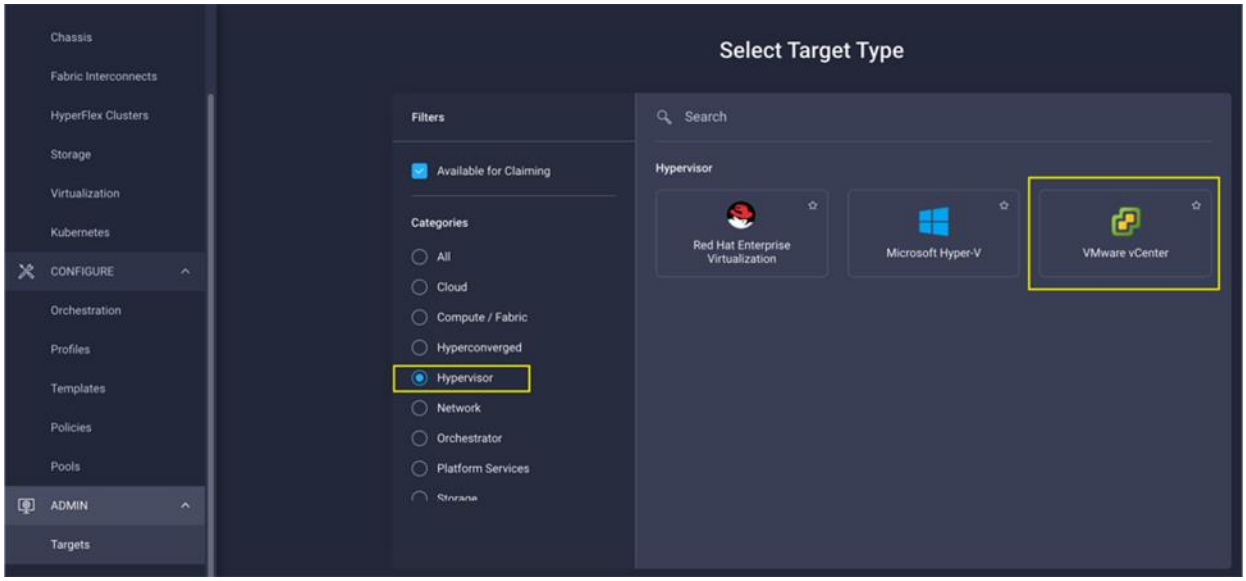

Obtain Hypervisor-level Information

After successfully claiming the VMware vCenter as a target, customers can view hypervisor-level information in Cisco Intersight including hosts, VMs, clusters, datastores, and so on.

Figure 43. VMware vCenter Information in Cisco Intersight

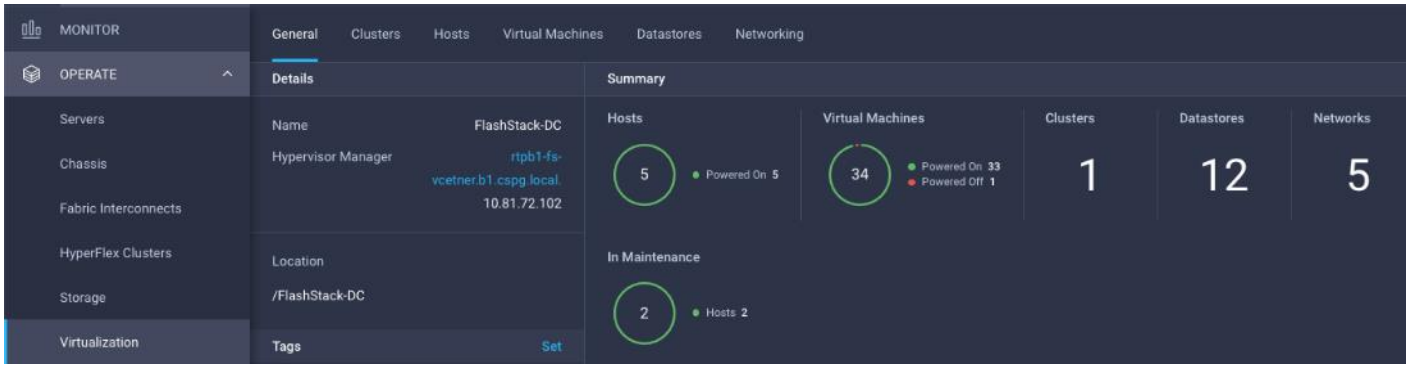

Table 6 lists some of the main virtualization properties presented in Cisco Intersight.

Table 6. Virtualization (VMware vCenter) Information in Cisco Intersight

General	Name	Name of the data center
	Hypervisor manager	Host name or IP address of the vCenter
Clusters	Name	Name of the cluster
	Data center	Name of the data center

	Hypervisor type	ESXi
	Hypervisor manager	vCenter IP address or the host name
	CPU capacity	CPU capacity in the cluster (GHz)
	CPU consumed	CPU cycles consumed by workloads (percentage and GHz)
	Memory capacity	Total memory in the cluster (GB)
	Memory consumed	Memory consumed by workloads (percentage and GB)
	Total cores	All the CPU cores across the CPUs in the cluster
	VMware cluster information allows customers to access additional details about hosts and virtual machines associated with the cluster.	
Hosts	Name	Host name or IP address
	Server	Server profile associated with the ESXi host
	Cluster	Cluster information if the host is part of a cluster
	Data center	VMware data center
	Hypervisor type	ESXi
	Hypervisor manager	vCenter IP address of host name
	Uptime	Host uptime
	Virtual Machines	Number and state of VMs running on a host
	CPU Information	CPU cores, sockets, vendor, speed, capacity, consumption, and other CPU related information
	Memory Information	Memory capacity and consumption information
	Hardware Information	Compute node hardware information such as serial number, model etc.
	Host information allows customers to access additional details about clusters, VMs, datastores, and networking related to the current ESXi host.	
Virtual machines	Name	Name of the VM
	Guest OS	Operating system, for example, RHEL, CentOS, etc.
	Hypervisor type	ESXi
	Host	ESXi host information for the VM
	Cluster	VMware cluster name
	Data center	VMware data center name
	IP address	IP address(s) assigned to the VM
	Hypervisor manager	IP address of host name of the vCenter

	Resource Information	CPU, memory, disk, and network information
	Guest Information	Hostname, IP address and operating system information
	VM information allows customers to access additional details about clusters, hosts, datastores, networking, and virtual disks related to the current VM.	
Datastores	Name	Name of the datastore in VMware vCenter
	Type	VMFS or NFS etc.
	Accessible	Yes, if datastore is accessible; No, if datastore is inaccessible
	Thin provisioning	Yes, if thin provisioning is allowed; No, if thin provisioning is not allowed
	Multiple host access	Yes, if multiple hosts can mount the datastore; No, if the datastore only allows a single host
	Storage capacity	Space in GB or TB
	Storage consumes	Percentage and GB
	Data center	Name of VMware vCenter data center
	Hypervisor manager	vCenter hostname or IP address
	Datastore Cluster	Datastore cluster information if datastore cluster is configured
	Hosts and Virtual Machines	Number if hosts connected to a datastore and number of VM hosted on the datastore
	Datastore information allows customers to access additional details about hosts and VMs associated with the datastore.	

Interact with Virtual Machines

VMware vCenter integration with Cisco Intersight allows customers to directly interact with the virtual machines (VMs) from the Cisco Intersight dashboard. In addition to obtaining in-depth information about a VM, including the operating system, CPU, memory, host name, and IP addresses assigned to the virtual machines, customers can use Intersight to perform following actions on the virtual machines (Figure 44).

- Launch VM console
- Power off
- Reset
- Shutdown guest OS
- Restart guest OS
- Suspend

Figure 44. Virtual Machine Actions in Cisco Intersight

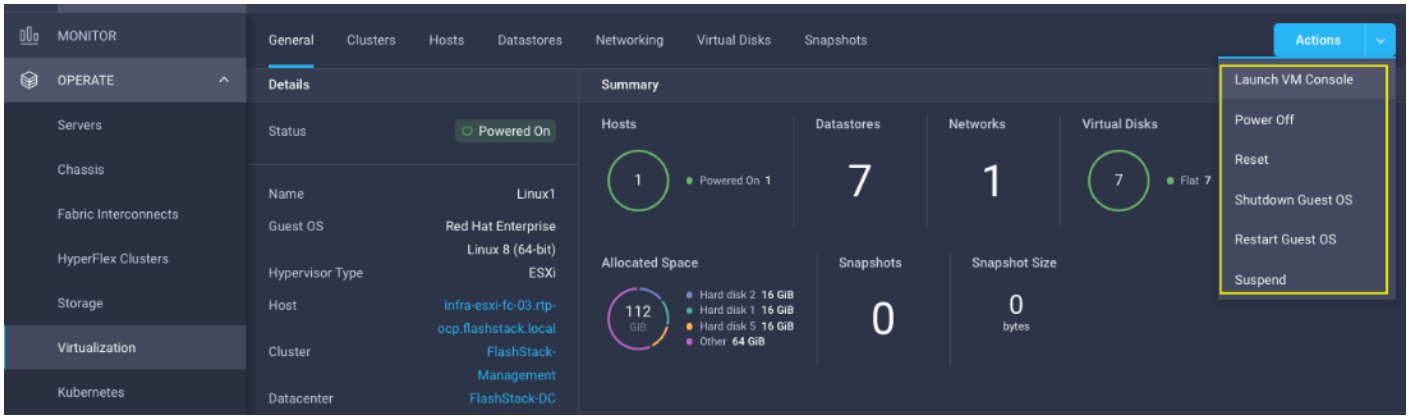

Cisco Intersight Orchestrator – VMware vCenter

Cisco Intersight Orchestrator provides various workflows that can be used for the VM and hypervisor provisioning. Some of the sample workflows available for VMware vCenter are captured in Table 7.

Table 7. VMware vCenter Workflows in Cisco Intersight Orchestrator

New VMFS Datastore	Create a storage volume and build VMFS datastore on the volume.
New Virtual Machine	Create a new virtual machine on the hypervisor from an OVA or OVF file. Datastore, Host/Cluster, and Image URL fields are mandatory. All other inputs are optional.
Remove VMFS Datastore	Remove VMFS datastore and remove the backing volume from the storage device.
Update VMFS Datastore	Expand a datastore on hypervisor manager by extending the backing storage volume to specified capacity, and then grow the datastore to utilize the additional capacity.

In addition to the above workflows, Cisco Intersight Orchestrator provides many tasks for customers to create custom workflows depending on their specific requirements. A sample subset of these tasks is highlighted in Figure 45.

Figure 45. VMware vCenter Tasks in Cisco Intersight Orchestrator

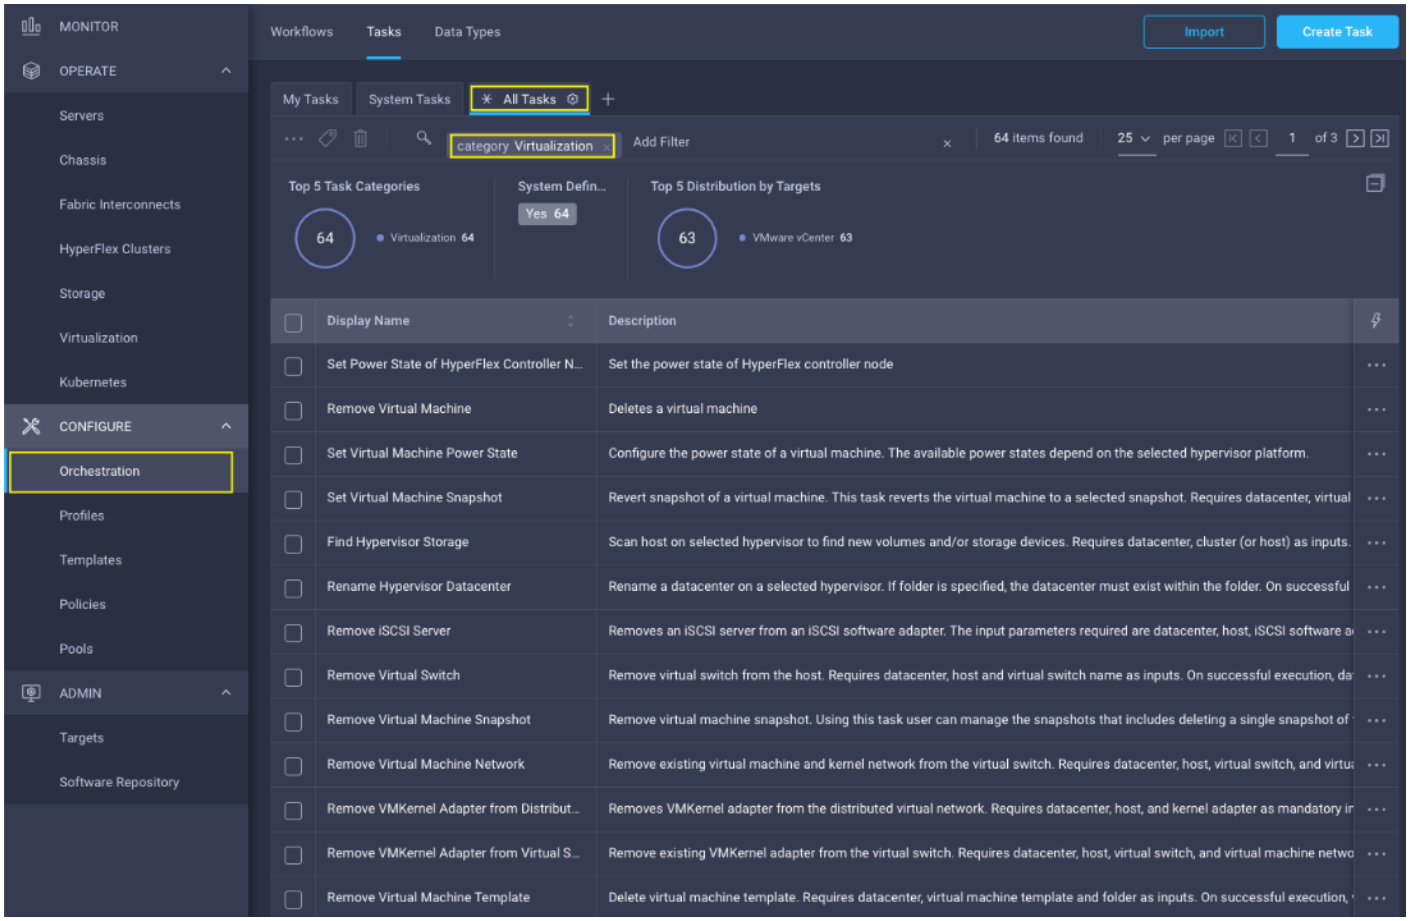

Design Considerations

Some of the key design considerations for the FlashStack with Cisco UCS X-Series and VMware 7.0 U2 are explained in this section.

Management Design Considerations

Out-of-band Management Network

The management interface of every physical device in FlashStack is connected to a dedicated out-of-band management switch, which can be part of the existing management infrastructure in a customer's environment. The out-of-band management network provides management access to all the devices in FlashStack for initial and on-going configuration changes. The routing and switching configuration for this network is independent of FlashStack deployment and therefore changes in FlashStack configurations do not impact management access to the devices.

In-band Management Network

The in-band management VLAN configuration is part of FlashStack design. The in-band VLAN is configured on Nexus switches and Cisco UCS within the FlashStack solution to provide management connectivity for vCenter, ESXi and other management components. The changes to FlashStack configuration can impact the in-band

management network and misconfigurations can cause loss of access to the management components hosted on the FlashStack.

Cisco Nexus 9000 Series vPC Best Practices

The following Cisco Nexus 9000 vPC design best practices and recommendations were used in this design:

- vPC peer keepalive link should not be routed over a vPC peer-link.
- The out-of-band management network is used as the vPC peer keepalive link in this design.
- Only vPC VLANs are allowed on the vPC peer-links. For deployments that require non-vPC VLAN traffic to be carried across vPC peer switches, a separate Layer 2 link should be deployed.

QoS Considerations

When using iSCSI for storage traffic, it may be necessary to prioritize the storage traffic over vMotion traffic if network bandwidth is constrained. FlashStack design allows customers to easily increase network bandwidth by adding additional links. Configuring QoS to prioritize storage traffic is also supported but QoS should always include a comprehensive plan for the individual customer environment.

Cisco UCS Fabric Interconnect (FI) Best Practices

Cisco UCS Fabric Interconnect is configured in default end-host mode. In this mode, the FIs will only learn MAC addresses from devices connected on Server and Appliance ports and FIs do not run spanning-tree. Loops avoidance is achieved using a combination of Deja-Vu check and Reverse Path Forwarding (RFP).

Oversubscription

To reduce the impact of an outage or scheduled downtime, it is a good practice to overprovision link bandwidth to enable a sustainable performance profile during component failure. Appropriately sized oversubscription protects workloads from being impacted by a reduced number of paths during a failure or maintenance event. Oversubscription can be achieved by increasing the number of physically cabled connections between storage and compute.

SAN Topology

For best performance, the ideal Fibre Channel SAN topology is a “Flat Fabric” where the FlashArray is only one hop away from any applications accessing the storage because additional hops add additional latency. Similarly, for iSCSI-based SAN design, it is recommended to reduce the number of network hops and not enable routing for the iSCSI storage LAN.

Pure Storage FlashArray Considerations

Connectivity

- Each FlashArray Controller should be connected to BOTH storage fabrics (A/B).
- Both 10 and 25 Gbps ports are provided via 2 onboard NICs on each FlashArray controller, if additional interfaces or 40 and 100 GE connectivity is required, additional NICs can be included in the original FlashArray BOM.

- Pure Storage offers up to 32Gb FC support on the latest FlashArray//X series arrays. Always make sure the correct number of HBAs and SFPs (with appropriate speed) are included in the original FlashArray BOM.

Host Groups and Volumes

It is a best practice to map Hosts to Host Groups and the Host Groups to Volumes in Purity. 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.

Size of the Volume

Purity removes the complexities of aggregates and RAID groups. When managing storage, a volume should be created based on the size required and purity takes care of availability and performance via RAID-HD and DirectFlash software. Customers can create 1 10-TB volume or 10 1-TB volumes and the performance and availability for these volumes will always be consistent. This feature allows customers to focus on recoverability, manageability, and administrative considerations of volumes instead of dwelling on availability or performance.

vCenter Deployment Consideration

While hosting the vCenter on the same ESXi hosts that the vCenter will manage is supported, it is a best practice to deploy the vCenter on a separate management infrastructure. The ESXi hosts in this new FlashStack with Cisco UCS X-Series environment can also be added to an existing customer vCenter. The in-band management VLAN will provide connectivity between the vCenter and the ESXi hosts deployed in the new FlashStack environment.

Jumbo Frames

An MTU of 9216 is configured at all network levels to allow jumbo frames as needed by the guest OS and application layer. The MTU value of 9000 is used on all the vSwitches and vSphere Distributed Switches (VDS) in the VMware environment.

Boot From SAN

When utilizing Cisco UCS Server technology with shared storage, it is recommended to configure boot from SAN and store the boot LUNs on remote storage. This enables architects and administrators to take full advantage of the stateless nature of Cisco UCS X-Series Server Profiles for hardware flexibility across the server hardware and overall portability of server identity. Boot from SAN also removes the need to populate local server storage thereby reducing cost and administrative overhead.

UEFI Secure Boot

This validation of FlashStack uses Unified Extensible Firmware Interface (UEFI) Secure Boot. UEFI is a specification that defines a software interface between an operating system and platform firmware. With UEFI secure boot enabled, all executables, such as boot loaders and adapter drivers, are authenticated by the BIOS before they can be loaded. Cisco UCS X210C compute nodes also contain a Trusted Platform Module (TPM). VMware ESXi 7.0 U2 supports UEFI Secure Boot and VMware vCenter 7.0 U2 supports UEFI Secure Boot Attestation between the TPM module and ESXi, validating that UEFI Secure Boot has properly taken place.

Pure Storage FlashArray considerations for VMware vSphere 7.0

The following Pure Storage design considerations and best practices for VMware vSphere were followed in this FlashStack 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 vSphere 7.0.
- vSphere 7.0 uses the Latency SATP that was introduced in vSphere 6.7U1. This replaces the I/O operations limit of 1 SATP, which was the default from vSphere 6.5U1. It is recommended to set samplingCycles - 16 and latencyEvalTime - 180000 ms.
- DataMover.HardwareAcceleratedMove, DataMover.HardwareAcceleratedInit, and VMFS3.HardwareAcceleratedLocking should all be enabled.
- When using iSCSI connected FlashArray volumes, it is recommended to set TCP DelayedAck to false (disabled) and LoginTimeout to 30 seconds.
- 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.
- Install VMware tools or Open VM tools whenever possible.
- When mounting snapshots, use the ESXi resignature option and avoid force-mounting.
- Ensure all ESXi hosts are connected to both FlashArray controllers and at a minimum, ensure two physical paths to each controller to achieve complete redundancy.
- Configure Host Groups on the FlashArray identical to clusters in vSphere. For example, if a cluster has four hosts in it, create a corresponding Host Group on the relevant FlashArray with exactly those four hosts.
- Use Paravirtual SCSI adapters for virtual machines whenever possible.
- Atomic Test and Set (ATS) is required on all Pure Storage volumes. This is a default configuration, and no configuration changes are needed.

For further details about the VMware vSphere Pure Storage FlashArray Best Practices please refer to:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/User_Guides_for_VMware_Solutions/FlashArray_VMware_Best_Practices_User_Guide/Quick_Reference%3A_Best_Practice_Settings

VMware Virtual Volumes

This validation of FlashStack supports VMware Virtual Volumes (vVols) for customers looking for more granular control of their SAN environment. vVol is a storage technology that provides policy-based, granular storage configuration and control of VMs. Through API-based interaction with an underlying array, VMware administrators can maintain storage configuration compliance using only native VMware interfaces. The Pure Storage FlashArray Plugin for the vSphere Web Client makes it possible to create, manage, and use vVols from within the Web Client.

Figure 46. vSphere Virtual Volumes Architecture

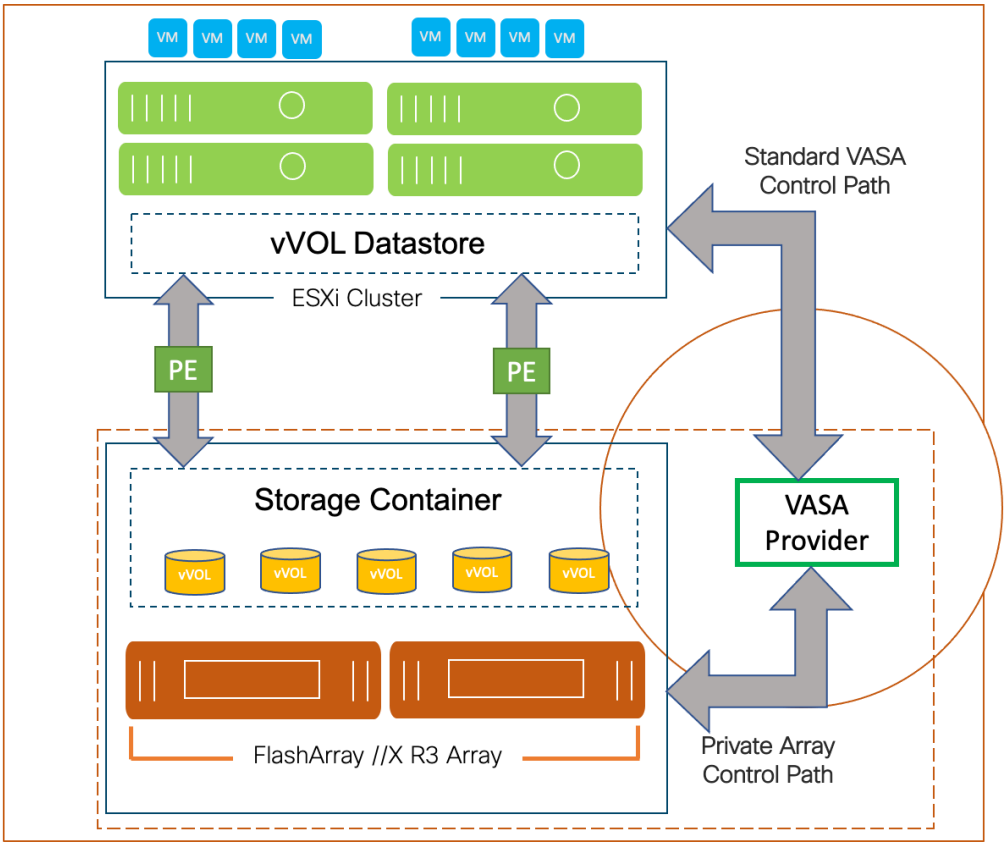

To start using vVOLs with the Pure Storage FlashArray, the FlashArray storage providers must be registered in vCenter Server. The Protocol Endpoint (PE) is then connected to the hostgroup and the vVol datastore is created.

FlashArray Virtual Volumes Considerations

VMware vCenters in Enhanced Linked Mode will each be able to communicate with the same FlashArray. However, vCenters that are not in Enhanced Linked Mode must use CA-Signed Certificates to use the same FlashArray. To support multiple VMware vCenters accessing the same FlashArray for vVols, the vCenters should be configured in Enhanced Linked Mode.

A VM's Config vVol stores the files required to build and manage the VM. Ensure that the Config vVol is part of an existing FlashArray Protection Group. Alternately, if customers are using storage policy that include snapshot or if customers prefer manual snapshots, Config vVol should be part of these snapshots. This will help with the VM recovery process if the VM is deleted.

When a Storage Policy is applied to a vVol VM, the volumes associated with that VM are added to the designated protection group. If replication is part of Storage Policy, the number of VMs using the storage policy as well as the replication groups becomes an important consideration. A large number of VMs with high change rate could cause replication to miss its schedule due to increased replication bandwidth and time needed to complete the scheduled snapshot. Pure Storage recommends vVol VMs with Storage Policies applied to be balanced between protection groups. To understand FlashArray limits on volume connections per host, volume count and snapshot Count, review the following document:

https://support.purestorage.com/FlashArray/PurityFA/General_Troubleshooting/Pure_Storage_FlashArray_Limits.

Pure Storage FlashArray Best Practices for vVols

Along with the above Pure Storage vVol considerations, following best practices should be considered during implementation of vVols:

- Create a Local FlashArray Array-Admin user to register the storage provider instead of using the local "pure user" account.
- Use the Round Robin pathing policy (default) for the Protocol Endpoint.
- Use the Pure Storage Plugin for the vSphere Client to register the FlashArray storage provider and mount the vVols datastore.
- When registering the storage providers manually, register both VASA providers with CT0.ETH0 and CT1.ETH0. The support for ETH1 interfaces are supported if a custom certificate is used.
- Manually mounting the vVol datastore requires users to connect the protocol endpoint (PE).
- A single PE utilizing the default device queue depth is sufficient in the design.
- VM Templates associated with the vVol VMs should also be kept on vVols.
- VMDK resizing of VMs that resides on a vVol should be completed from vSphere Client and not from FlashArray GUI.
- ESXi Hosts, vCenter Server and FlashArray should synchronize time to the same NTP Server.
- TCP port 8084 must be open and accessible from vCenter Servers and ESXi hosts to the FlashArray that will be used for vVol.
- vCenter Server should not reside on vVols.
- The FlashArray Protocol Endpoint object 'pure-protocol-endpoint' must exist. The FlashArray admin must not rename, delete or otherwise edit the default FlashArray Protocol Endpoint.

For more information on vVols best practices refer to the following summary:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/User_Guides_for_VMware_Solutions/Virtual_Volumes_User_Guide/vVols_User_Guide%3A_Best_Practice_Summary.

NVMe over Fabrics

NVMe over Fabrics (NVMe-oF) is an extension of the NVMe network protocol to Ethernet and Fibre Channel delivering faster and more efficient connectivity between storage and servers as well as a reduction in CPU utilization of application host servers. This validation of FlashStack supports NVMe over Fibre Channel (NVMe/FC) to provide the high-performance and low-latency benefits of NVMe across fabrics. In this solution, NVMe initiators consisting of Cisco UCS X210c compute nodes access Pure FlashArray NVMe targets over Fibre Channel.

Figure 47. End-to-End NVMe over Fibre Channel Connectivity

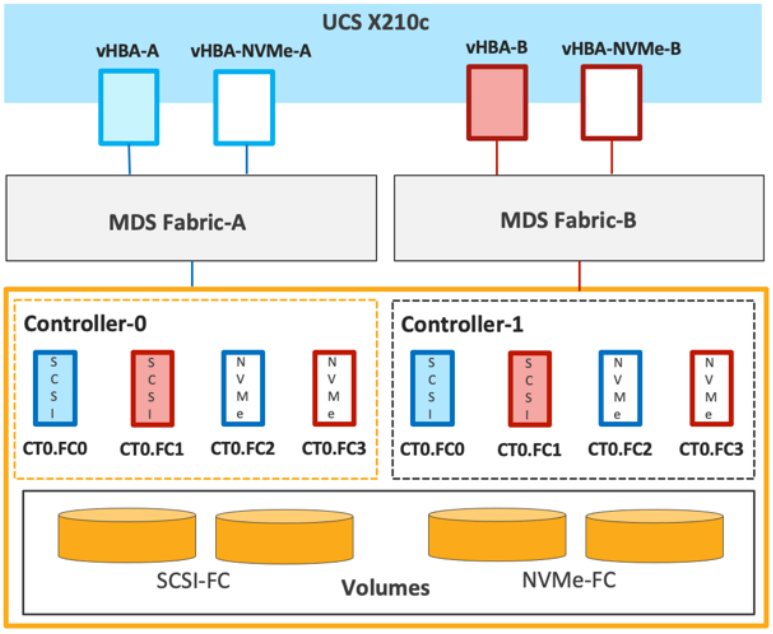

Each port on the Pure FlashArray can be configured as traditional scsi-fc port or as a nvme-fc port to support NVMe end-to-end via fibre channel from the host to storage array. Two ports on each Pure Storage FlashArray controller are configured as SCSI ports and two ports are configured as NVMe ports as shown in Figure 47.

A given FC port on Pure Storage FlashArray can either be configured as FC-SCSI or FC-NVMe port.

In a Cisco UCS server profile, both standard Fibre Channel and FC-NVMe vHBAs can be created. A default Fibre Channel adapter policy named fc-nvme-initiator is preconfigured in Cisco Intersight. This policy contains recommended adapter settings for FC-NVMe. Both Fibre Channel and FC-NVMe vHBAs can exist in a Cisco UCS server profile on a single server.

To support NVMe over Fabric, four vHBAs, two FC-NVME initiators and two Fibre Channel initiators (one on each Fibre Channel fabric), are created for each server profile. Cisco MDS 9132T switches are configured with appropriate zoning to connect the FC-NVMe and Fibre Channel vHBAs to appropriate storage targets. Single-initiator, multiple-target zones are used for both FCP and FC-NVMe. VMware ESXi automatically connects to Pure FlashArray NVMe subsystem and discovers all shared NVMe storage devices that it can reach once the SAN zoning on MDS switches, and the configuration of host/host groups and volumes is completed on the Pure FlashArray.

Solution Automation

In addition to command line interface (CLI) and graphical user interface (GUI) configurations, explained in the deployment guide, all FlashStack components support configurations through automation frameworks such as Ansible and Terraform etc. The FlashStack solution validation team will share automation modules to configure Cisco Nexus, Cisco UCS, Cisco MDS, Pure Storage FlashArray, VMware ESXi, and VMware vCenter. This community-supported GitHub repository is meant to expedite customer adoption of automation by providing them sample configuration playbooks that can be easily developed or integrated into existing customer automation frameworks.

Deployment Hardware and Software

Table 8 lists the hardware and software versions used during solution validation. It is important to note that the validated FlashStack solution explained in this document adheres to Cisco, Pure Storage, and VMware interoperability matrix to determine support for various software and driver versions. Customers should use the same interoperability matrix to determine support for components that are different from the current validated design.

Click the following links for more information:

- Pure Storage Interoperability Matrix. Note, this interoperability list will require a support login form Pure:

  https://support.purestorage.com/FlashArray/Getting_Started/Compatibility_Matrix

- Pure Storage FlashStack Compatibility Matrix. Note, this interoperability list will require a support login from Pure:

  https://support.purestorage.com/FlashStack/Product_Information/FlashStack_Compatibility_Matrix

- Cisco UCS Hardware and Software Interoperability Tool:

  http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html

- VMware Compatibility Guide:

  http://www.vmware.com/resources/compatibility/search.php

Table 8. Hardware and Software Revisions

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

	Pure Storage VASA Provider	3.5
	Pure Storage Plugin	5.0.0

Validation

A high-level overview of the FlashStack design validation is provided in this section. Solution validation covers various aspects of the converged infrastructure including compute, virtualization, network, and storage. The test scenarios are divided into four broad categories:

- Functional validation – physical and logical setup validation
- Feature verification – feature verification for FlashStack design
- Availability testing – link and device redundancy and high availability testing
- Infrastructure as a code validation – verify automation and orchestration of solution components

The goal of solution validation is to test functional aspects of the design and unless explicitly called out, the performance and scalability is not covered during solution validation. However, limited load is always generated using tools such as IOMeter and/or iPerf to help verify test setup. Some of the examples of the types of tests executed include:

- Verification of features configured on various FlashStack components
- Powering off and rebooting redundant devices and removing redundant links to verify high availability
- Path MTU verification including both storage and virtual machine traffic
- Failure and recovery of vCenter and ESXi hosts in a cluster
- Failure and recovery of storage access paths across FlashArray controllers, MDS and Nexus switches, and fabric interconnects
- Server Profile migration between compute nodes
- Load generation using IOMeter VMs hosted on FlashStack components and path validation

As part of the validation effort, solution validation team identifies the problems, works with the appropriate development teams to fix the problem and provides work arounds as necessary.

Summary

The FlashStack solution is a validated approach for deploying Cisco and Pure Storage technologies and products for building shared private and public cloud infrastructure. With the introduction of Cisco X-Series modular platform to FlashStack, customers can now manage and orchestrate the next-generation Cisco UCS platform from the cloud using Cisco Intersight. Some of the key advantages of integrating Cisco UCS X-Series and Cisco Intersight into the FlashStack infrastructure are:

- Simpler and programmable infrastructure
- Power and cooling innovations and better airflow
- Fabric innovations for heterogeneous compute and memory composability
- Innovative cloud operations providing continuous feature delivery
- Future-ready design built for investment protection

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

Appendix

This section includes links to various product pages.

Compute

Cisco Intersight: https://www.intersight.com

Cisco Intersight Managed Mode: https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide.html

Cisco Unified Computing System: http://www.cisco.com/en/US/products/ps10265/index.html

Cisco UCS 6400 Series Fabric Interconnects: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html

Network

Cisco Nexus 9000 Series Switches: http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html

Cisco MDS 9132T Switches: https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9100-series-multilayer-fabric-switches/datasheet-c78-739613.html

Storage

Pure Storage FlashArray//X: https://www.purestorage.com/products/flasharray-x.html

Virtualization

VMware vCenter Server: http://www.vmware.com/products/vcenter-server/overview.html

VMware vSphere: https://www.vmware.com/products/vsphere

Interoperability Matrix

Cisco UCS Hardware Compatibility Matrix: https://ucshcltool.cloudapps.cisco.com/public/

VMware and Cisco Unified Computing System: http://www.vmware.com/resources/compatibility

Pure Storage Interoperability Matrix. Note, this interoperability list will require a support login form Pure: https://support.purestorage.com/FlashArray/Getting_Started/Compatibility_Matrix

Pure Storage FlashStack Compatibility Matrix. Note, this interoperability list will require a support login from Pure: https://support.purestorage.com/FlashStack/Product_Information/FlashStack_Compatibility_Matrix

About the Authors

Haseeb Niazi, Principal Technical Marketing Engineer, Cisco Systems, Inc.

Haseeb Niazi has over 22 years of experience at Cisco in the Datacenter, Enterprise and Service Provider Solutions and Technologies. As a member of various solution teams and Advanced Services, Haseeb has helped many enterprise and service provider customers evaluate and deploy a wide range of Cisco solutions. As a technical marking engineer at Cisco UCS Solutions group, Haseeb focuses on network, compute, virtualization, storage, and orchestration aspects of various Compute Stacks. Haseeb holds a master's degree in Computer Engineering from the University of Southern California and is a Cisco Certified Internetwork Expert (CCIE 7848).

Joe Houghes, Senior Solutions Architect, Pure Storage, Inc.

Joe is a Senior Solutions Architect in the Portfolio Solutions team within Pure Storage, focused on solutions on the FlashStack platform along with automation and integration. He has experience from over 15 years in Information Technology across various customer/vendor organizations with architecture and operations expertise covering compute, networking, storage, virtualization, business continuity and disaster recovery, and cloud computing technologies, plus automation and integration across many applications & vendor platforms.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the author would like to thank:

- Sreenivasa Edula, Technical Marketing Engineer, Cisco Systems, Inc.
- Craig Waters, Technical Director, Pure Storage, Inc.

Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community at https://cs.co/en-cvds.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)