FlexPod Datacenter with Cisco UCSM M6, VMware vSphere 8, and NetApp ONTAP 9.12.1 Design Guide

Published: March 2023

In partnership with:

About the Cisco Validated Design Program

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

Executive Summary

The FlexPod Datacenter solution is a validated approach for deploying Cisco and NetApp technologies and products to build shared private and public cloud infrastructure. Cisco and NetApp have partnered to deliver a series of FlexPod solutions that enable strategic data-center platforms. The success of the FlexPod solution is driven through its ability to evolve and incorporate both technology and product innovations in the areas of management, compute, storage, and networking. This document explains the design details of incorporating the Cisco UCS B and C-Series M6 servers in UCS Managed Mode, NetApp ONTAP 9.12.1, and VMware vSphere 8.0 into FlexPod Datacenter and the ability to monitor and manage FlexPod components from the cloud using Cisco Intersight. Some of the key advantages of integrating these components into the FlexPod infrastructure are:

- **Simpler and programmable infrastructure:** infrastructure as code delivered through a single partner integrable open API.
- **Latest Software Feature Support with Cisco UCS M6:** demonstrating the latest features of NetApp ONTAP 9.12.1 and VMware vSphere 8.0 with the Cisco UCS B and C-Series servers.
- **Introduction of the Cisco UCS 5th Generation Virtual Interface Cards (VICs) with Cisco UCS M6:** validating the Cisco UCS VIC 15411 in the Cisco UCS B200 M6, and the Cisco UCS VIC 15238 and 15428 in the Cisco UCS C220 M6.
- **Extend the capabilities of current investments:** get more capabilities out of existing Cisco UCS and NetApp hardware with the support of VMware vSphere 8.0 and NetApp ONTAP 9.12.1.

In addition to the compute-specific hardware and software innovations, the integration of the Cisco Intersight cloud platform with VMware vCenter and NetApp Active IQ Unified Manager delivers monitoring, orchestration, and workload optimization capabilities for different layers (virtualization and storage) of the FlexPod infrastructure. The modular nature of the Cisco Intersight platform also provides an easy upgrade path to additional services, such as workload optimization.

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

Solution Overview

This chapter contains the following:

Introduction

The Cisco Unified Computing System (Cisco UCS) with UCS Managed Mode is a modular compute system, configured and managed from Cisco UCS Manager on Cisco UCS fabric Interconnects. It is designed to meet the needs of modern applications and to improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco UCS Manager platform delivers simplified configuration, deployment, maintenance, and support.

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 UCS managed FlexPod Datacenter infrastructure. The document introduces various design elements and explains various considerations and best practices for a successful deployment. The document also highlights the design and product requirements for integrating virtualization and storage systems to Cisco Intersight to deliver a true cloud-based integrated approach to infrastructure management.

What's New in this Release?

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

- Integration of the Cisco UCS 5th Generation VIC 15411 in Cisco UCS B200 M6 and the Cisco UCS 5th Generation VIC 15238 and 15428 in Cisco UCS C220 M6 into FlexPod Datacenter
- An integrated, more complete end-to-end Infrastructure as Code (IaC) Day 0 configuration of the FlexPod Infrastructure utilizing Ansible Scripts
- VMware vSphere 8.0
- Continued Integration with Cisco Intersight
- Design Guide Information and Deployment Guidance on Sustainability and Security

Solution Summary

The FlexPod Datacenter solution with Cisco UCS M6, VMware 8.0, and NetApp ONTAP 9.12.1 offers the following key customer benefits:

- Simplified cloud-based monitoring of 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
- Easy to deploy, consume, and manage architecture, which saves time and resources required to research, procure, and integrate off-the-shelf components
- Support for component monitoring, solution automation and orchestration, and workload optimization

Like all other FlexPod solution designs, FlexPod Datacenter with Cisco UCS M6 is configurable according to demand and usage. Customers can purchase exactly the infrastructure they need for their current application requirements and can then scale up by adding more resources to the FlexPod system or scale out by adding more FlexPod instances. All components can be scaled as needed, allowing the FlexPod to meet the exact customer needs.

Technology Overview

This chapter contains the following:

FlexPod Datacenter

FlexPod Datacenter architecture is built using the following infrastructure components for compute, network, and storage:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus and Cisco MDS switches
- NetApp All Flash FAS (AFF), FAS, and All SAN Array (ASA) storage systems

Figure 1. FlexPod Datacenter Components

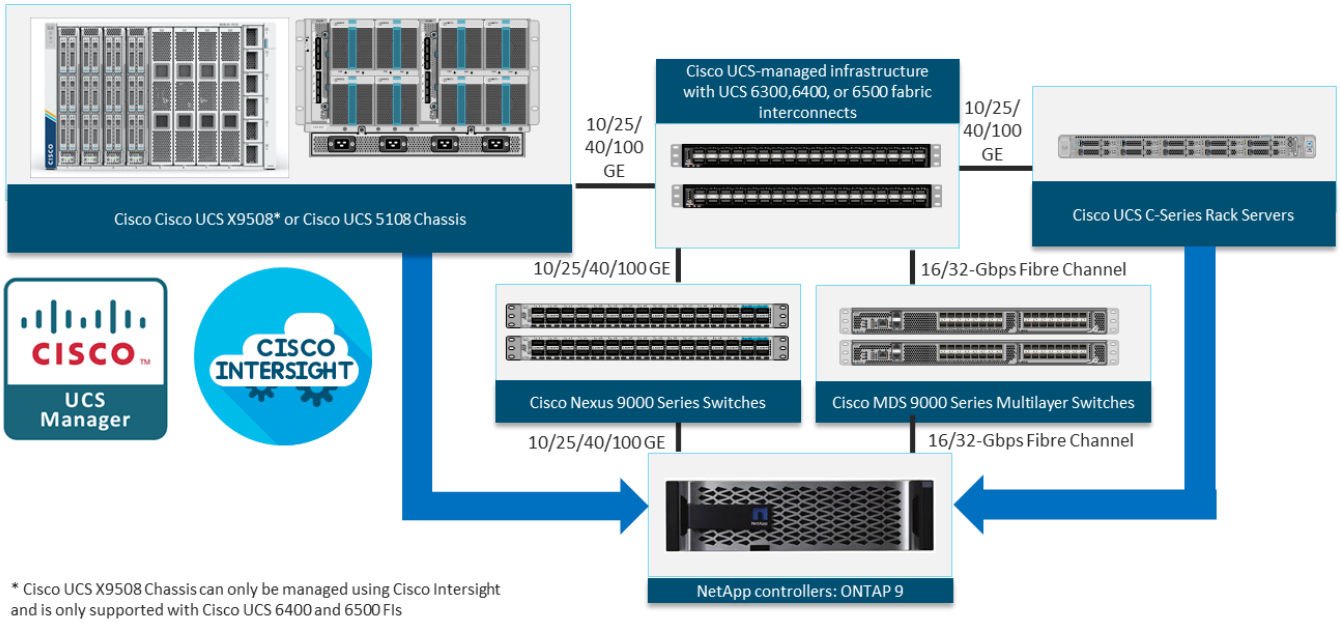

All the FlexPod components have been integrated so that customers can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlexPod is its ability to maintain consistency at scale. Each of the component families shown in Figure 1 (Cisco UCS, Cisco Nexus, Cisco MDS, and NetApp controllers) offers platform and resource options to scale up or scale out the infrastructure while supporting the same features.

The FlexPod Datacenter solution with Cisco UCS M6 is built using the following hardware components:

- Cisco UCS 5108 Chassis with Cisco UCS 2408 Fabric Extenders and up to eight Cisco UCS B200 M6 Compute Nodes
- Fourth-generation Cisco UCS 6454 Fabric Interconnects to support 10/25/40/100GbE and 16/32GbFC connectivity from various components
- Cisco UCS C220 M6 servers connected directly to the Cisco UCS 6454 Fabric Interconnects with either 25GbE (Cisco UCS VIC 15428 or 1455) or 100GbE (Cisco UCS VIC 15238)
- High-speed Cisco NX-OS-based Cisco Nexus 93180YC-FX switching design to support up to 100GE connectivity and optional 32G FC connectivity
- Cisco MDS 9132T switches to support 32G FC connectivity
- NetApp AFF A400 end-to-end NVMe storage with up to 100GE connectivity and 32G FC connectivity

The software components of the solution consist of:

- Cisco UCS Manager to deploy the Cisco UCS components, and maintain and support the FlexPod components

- Cisco Intersight Assist Virtual Appliance to help connect NetApp AIQUM, Cisco Nexus Switches, and VMware vCenter to Cisco Intersight
- NetApp Active IQ Unified Manager to monitor and manage the storage and for NetApp ONTAP integration with Cisco Intersight
- VMware vCenter to set up and manage the virtual infrastructure as well as Cisco Intersight integration
- NetApp ONTAP to set up and manage the NetApp AFF A400 storage

Infrastructure as Code with Ansible

This FlexPod solution provides a fully automated solution deployment that explains all sections of the infrastructure and application layer. The configuration of the NetApp ONTAP Storage, Cisco Network and Compute, and VMware layers are automated by leveraging Ansible playbooks that have been developed to setup the components as per the solution best practices that were identified during the testing and validation.

Note: There are two modes to configure Cisco UCS, one is UCS Managed, and the other is Intersight Managed Mode (IMM). Here, the Ansible scripts will configure Cisco UCS Managed mode.

The automated deployment using Ansible provides a well-defined sequence of execution across the different constituents of this solution. Certain phases of the deployment also involve the exchange of parameters or attributes between compute, network, storage, and virtualization and also involve some manual intervention. All phases have been clearly demarcated and the implementation with automation is split into equivalent phases via Ansible playbooks with a tag-based execution of a specific section of the component's configuration.

Figure 2. Infrastructure as Code with Ansible

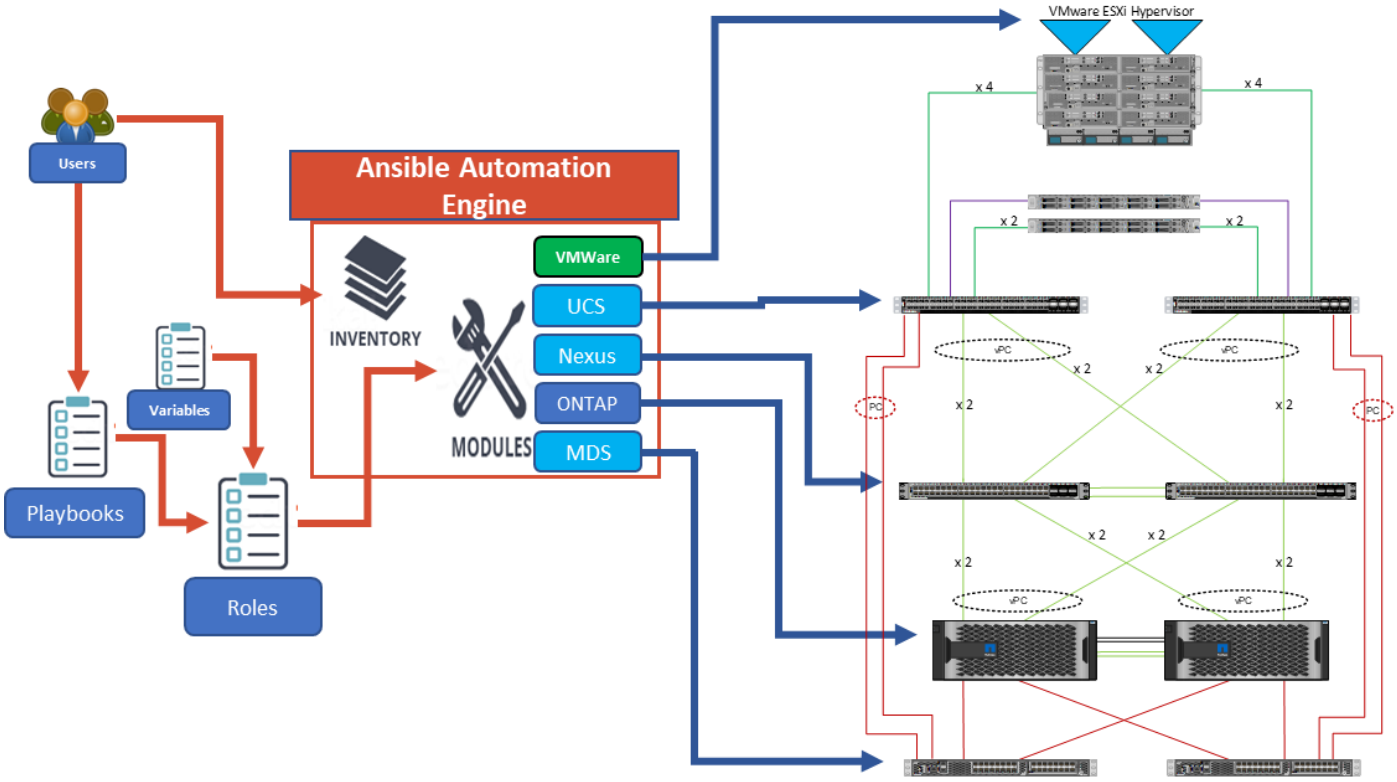

As illustrated in Figure 2, the Ansible playbooks to configure the different sections of the solution invoke a set of Roles and consume the associated variables that are required to setup the solution. The variables needed for this solution can be split into two categories – user input and defaults/best practices. Based on the installation environment customers can choose to modify the variables to suit their requirements and proceed with the automated installation.

Note: The automation for ONTAP is scalable in nature that can configure anywhere from a single HA pair to a fully scaled 24 node ONTAP cluster.

After the base infrastructure is setup with NetApp ONTAP, Cisco Network and Compute, and VMware, customers can also deploy the FlexPod Management Tools like SnapCenter Plug-in for VMware vSphere, and Active IQ Unified Manager in an automated fashion.

Cisco Unified Computing System

Cisco UCS B200 M6 Blade Servers

The Cisco UCS B200 M6 server, shown in Figure 3, is a half-width blade upgrade from the Cisco UCS B200 M5.

Figure 3. Cisco UCS B200 M6 Blade Server

Key features:

- 3rd Gen Intel Xeon Scalable processors with up to 40 cores per socket, 2-socket
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane DC Persistent Memory
- Up to 80 Gbps of I/O throughput with Cisco UCS 6332, 6454, or 6536 FI
- Two optional, hot-pluggable, Solid-State Drives (SSDs), or Non-Volatile Memory Express (NVMe) 2.5-inch drives with a choice of enterprise-class Redundant Array of Independent Disks (RAIDs) or pass-through controllers or 4 M.2 SATA drives for flexible boot and local storage capabilities

For more information about the Cisco UCS B200 M6 Blade Servers, see Cisco UCS B200 M6 Blade Server Data Sheet - Cisco.

Cisco UCS C220 M6 Rack Servers

The Cisco UCS C220 M6 rack server shown in Figure 4, is a high-density 2-socket rack server that is an upgrade from the Cisco UCS C220 M5.

Figure 4. Cisco UCS C220 M6 Rack Server

It features:

- 3rd Gen Intel Xeon Scalable processors with up to 40 cores per socket, 2-socket
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane DC Persistent Memory
- Up to 3 PCIe 4.0 slots plus a modular LAN on Motherboard (mLOM) slot
- Up to 10 SAS/SATA or NVMe disk drives
- Up to two GPUs supported
- Up to 100 Gbps of I/O throughput with Cisco UCS 6454 or 6536 FI

For more information about the Cisco UCS C220 M6 Rack Servers, see: Cisco UCS C220 M6 Rack Server Data Sheet - Cisco.

Intel Optane DC Persistent Memory in Cisco UCS B200 M6 and Cisco UCS C220 M6 Servers

Intel 200 Series Optane DC Persistent Memory was validated in Cisco UCS B200 and C220 M6 servers in two ways. The first validation was running a VMware ESXi 8.0 Host in VMware supported Memory Mode as specified in vSphere Support for Intel's Optane Persistent Memory (PMEM) (67645). In this validation, the 2-socket server platform with 2TB available memory was setup with the Balanced Profile BIOS setting, and no issues were seen with the ESXi host. The second validations were in App Direct Mode. The first validation in App Direct Mode was to configure Intel Optane in App Direct Mode, not use it, and show that it had no effect on running applications. Testing was also done in assigning NVDIMMs to Windows 11 VMs and using them as direct access (DAX) fast disks. Intel 200 Series Optane memory is 3200MHz along with the standard DIMMs.

Note: Intel Optane DC Persistent Memory has been discontinued by Intel and will have no further development, but the 200 series PMEM is still available, and its capabilities can be used with Cisco UCS B200 and C220 M6 servers.

Cisco UCS 6400 Series Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide a single point for connectivity and management for the entire Cisco Unified Computing System. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly available management domain controlled by Cisco UCS Manager.

The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in deterministic I/O latency regardless of a server or virtual machine's topological location in the system.

The Cisco UCS Fabric Interconnect provides both network connectivity and management capabilities for Cisco Unified Computing System. IOM modules in the blade chassis support power supply, along with fan and blade management. They also support port channeling and, thus, better use of bandwidth. The IOMs support virtualization-aware networking in conjunction with the Fabric Interconnects and Cisco Virtual Interface Cards (VIC).

The Cisco UCS 6400 Series Fabric Interconnect is a core part of Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6400 Series offers line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and 32 Gigabit Fibre Channel functions. The 100-Gbps ports on the Cisco UCS 6400 Series Fabric Interconnects can now be configured as Server Ports, allowing Cisco UCS C-Series servers with 100-Gbps VIC cards to be directly connected to the Fabric Interconnect at 100-Gbps.

Figure 5. Cisco UCS 6454 Fabric Interconnect

The Cisco UCS 6454 54-Port Fabric Interconnect is a One-Rack-Unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE and Fibre Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has 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-Gbps Ethernet ports or 8/16/32-Gbps Fibre Channel ports. All Ethernet ports are capable of supporting FCoE.

The Cisco UCS 64108 Fabric Interconnect (FI) is a 2-RU top-of-rack switch that mounts in a standard 19-inch rack such as the Cisco R Series rack. The 64108 is a 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 7.42 Tbps throughput and up to 108 ports. The switch has 16 unified ports (port numbers 1-16) that can support 10/25-Gbps SFP28 Ethernet ports or 8/16/32-Gbps Fibre Channel ports, 72 10/25-Gbps Ethernet SFP28 ports (port numbers 17-88), 8 1/10/25-Gbps Ethernet SFP28 ports (port numbers 89-96), and 12 40/100-Gbps Ethernet QSFP28 uplink ports (port numbers 97-108). All Ethernet ports are capable of supporting FCoE. The Cisco UCS 64108 FI is supported in the FlexPod solution but was not validated in this project.

Figure 6. Cisco UCS 64108 Fabric Interconnect

For more information on the Cisco UCS 6400 Series Fabric Interconnects, see the Cisco UCS 6400 Series Fabric Interconnects Data Sheet.

Cisco UCS 2408 Fabric Extender

The Cisco UCS 2408 connects the I/O fabric between the Cisco UCS 6454 Fabric Interconnect and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic converged fabric to connect all blades and chassis together. Since the fabric extender is similar to a distributed line card, it does not perform any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity, and enabling Cisco UCS to scale to many chassis without multiplying the number of switches needed, reducing TCO, and allowing all chassis to be managed as a single, highly available management domain.

The Cisco UCS 2408 Fabric Extender has eight 25-Gigabit Ethernet, FCoE-capable, Small Form-Factor Pluggable (SFP28) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2408 provides 10-Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis, giving it a total 32 10G interfaces to Cisco UCS blades. Typically configured in pairs for redundancy, two fabric extenders provide up to 400 Gbps of I/O from FI 6400's to 5108 chassis.

Cisco UCS 6536 Fabric Interconnects

The Cisco UCS 6536 is a 36-port fabric interconnect. This single RU device includes up to 36 10/25/40/100 Gbps Ethernet ports, and 16 8/16/32-Gbps Fibre Channel ports via 4 128 Gbps to 4x32 Gbps breakouts on ports 33-36. All 36 ports support breakout cables or QSA interfaces. The Cisco UCS 6536 Fabric Interconnect is supported but was not validated in this project. The Cisco UCS 6536 supports Cisco UCS Manager and connects to the Cisco UCS 2408 in the Cisco UCS 5100 Series Blade Server Chassis using 100 Gbps to 4x25 Gbps breakouts.

Figure 7. Cisco UCS 6536 Fabric Interconnect

Cisco UCS 5th Generation Virtual Interface Cards (VICs)

This section describes the Cisco UCS B200 M6 servers that support the following 5th Generation Cisco VIC cards.

Cisco VIC 15411

Cisco VIC 15411 fits in the mLOM slot in the Cisco UCS B200 M6 and when coupled with the optional Port Expander in the mezzanine slot enables up to 40 Gbps of unified fabric connectivity to each of the chassis 2408 IOMs for a total of 80 Gbps of connectivity per server. Because the 10 GBASE-KR lanes are fixed length in the chassis, the 4x 10GBASE-KR lanes can be port channeled into a true 40 Gbps signal. It is important to note that since the links between the Cisco UCS 2408 IOM and the 6454 FIs are 25 Gbps and port channeled up to an 8x25 Gbps port channel, that a single session (TCP or FCoE) is limited to 25 Gbps and that 40 Gbps is attained by multiplexing multiple sessions on the link. Cisco VIC 15411 supports 512 virtual interfaces (both FCoE and Ethernet) along with the latest networking innovations such as NVMeoF over FC or TCP, VxLAN/NVGRE offload, and so forth.

Figure 8. Cisco VIC 15411 in Cisco UCS B200 M6

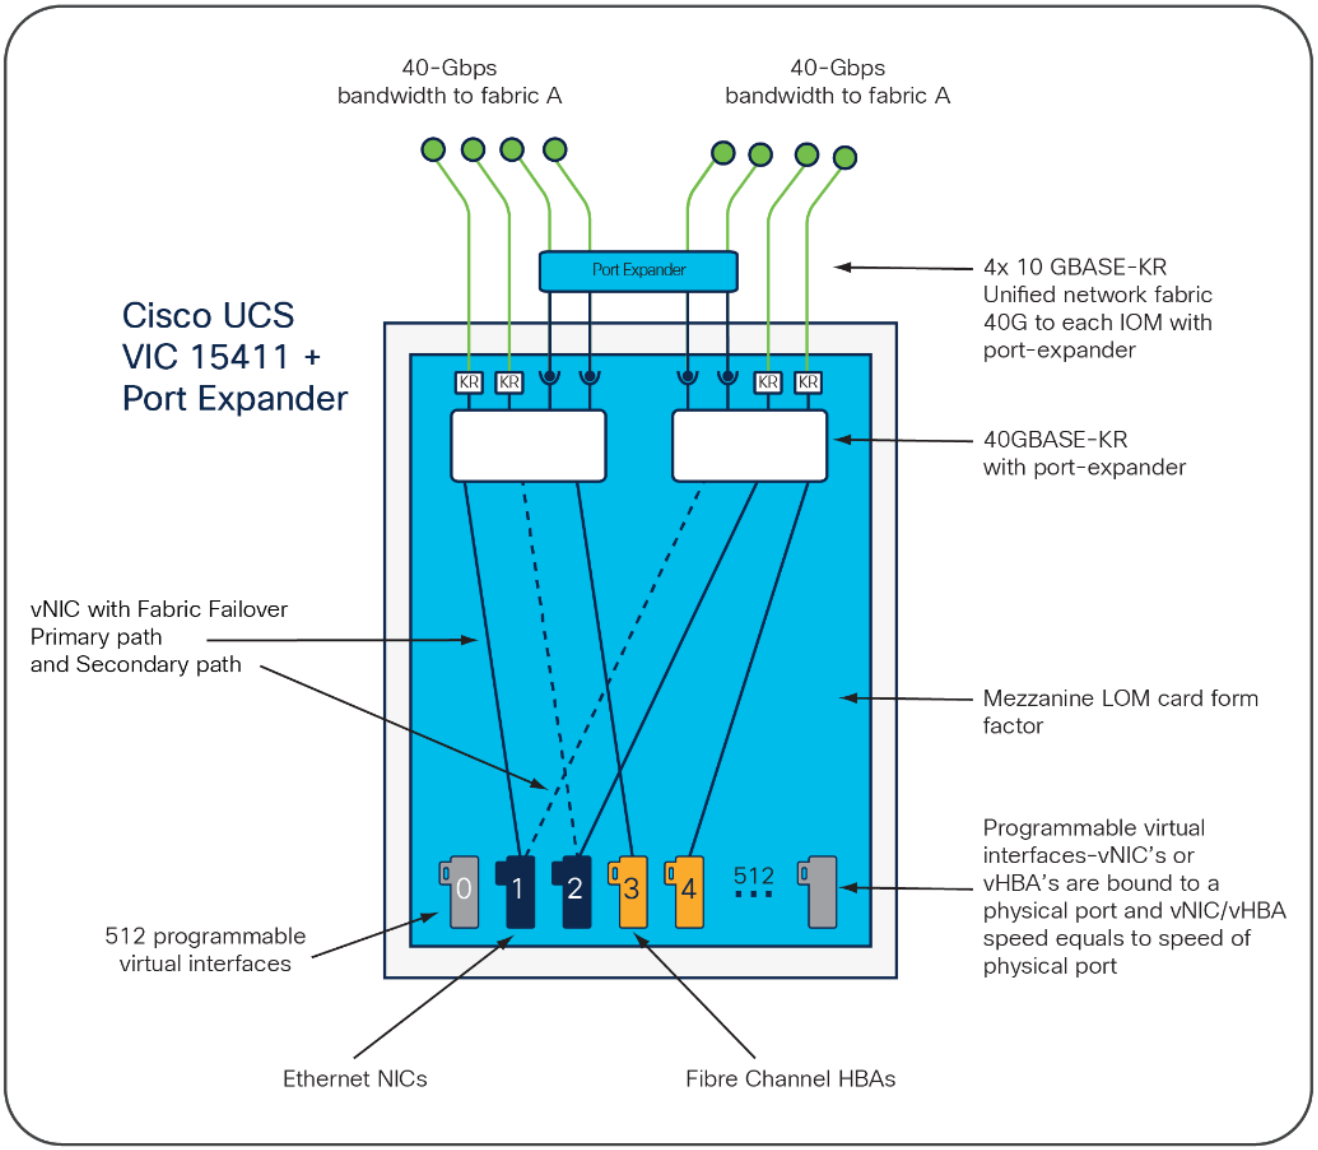

Cisco VIC 15238

Cisco VIC 15238 fits in the mLOM slot in the Cisco UCS C220 M6 (and C225 and C245) server and enables up to 100 Gbps of unified fabric connectivity to each fabric interconnect for a total of 200 Gbps of connectivity per server. Cisco VIC 15238 connectivity to the fabric interconnects is delivered through 2x 100-Gbps links. Cisco VIC 15238 supports 512 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 9. Single Cisco VIC 15238 in Cisco UCS C220 M6

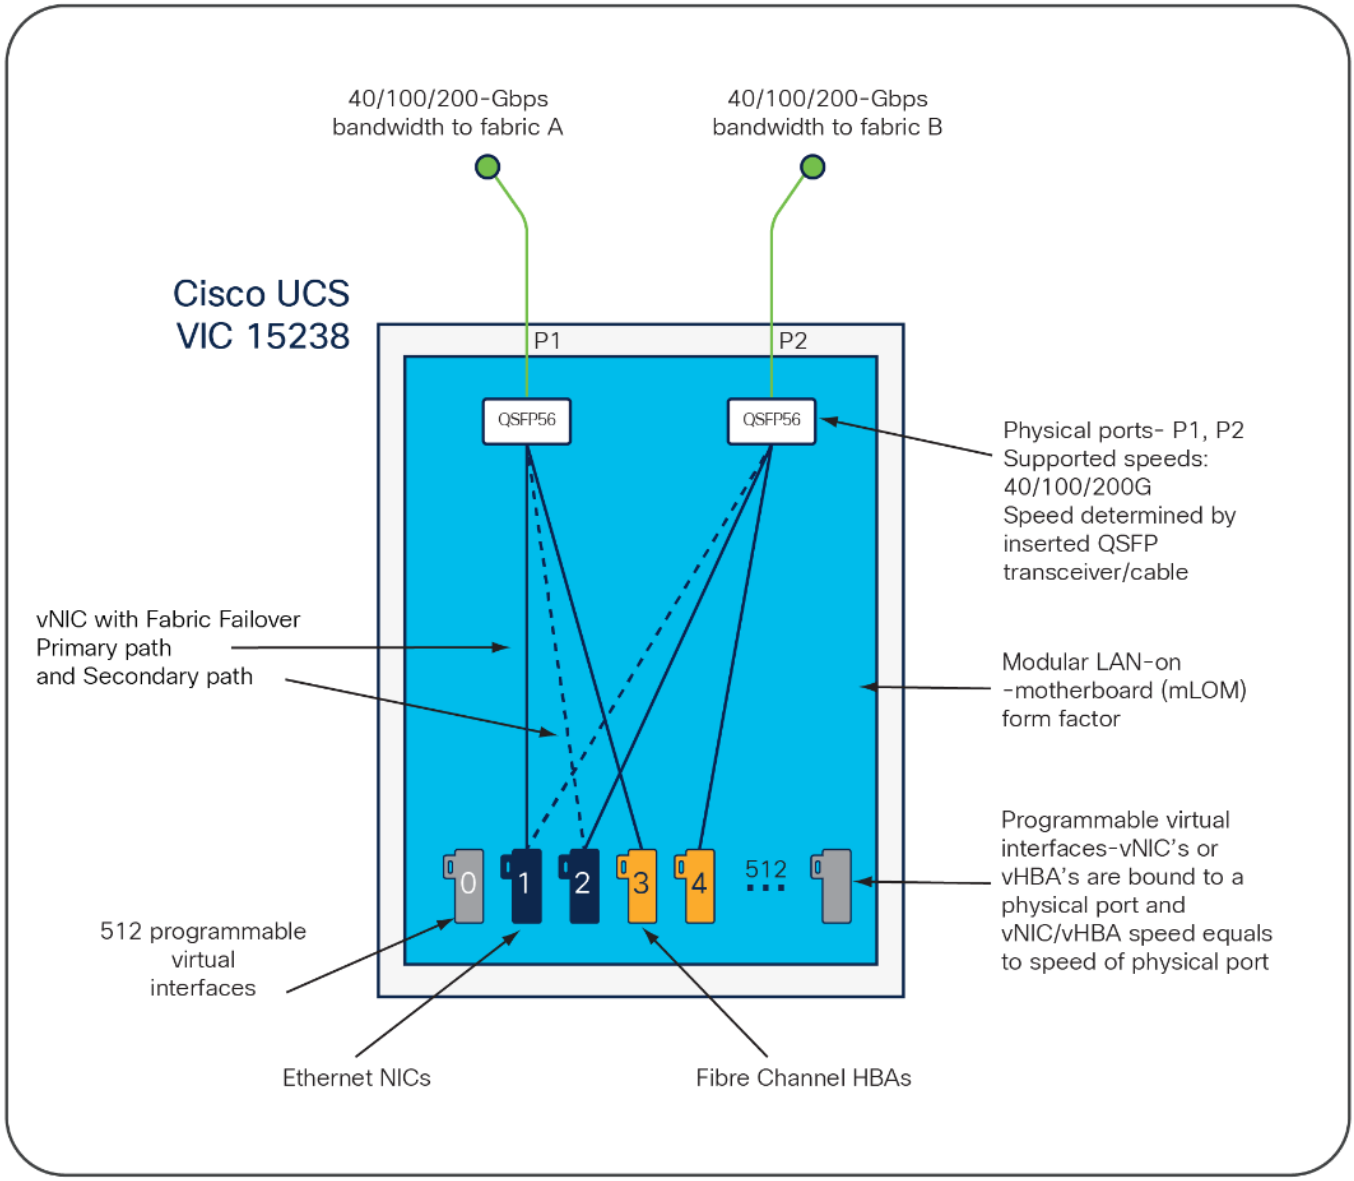

Cisco VIC 15428

Cisco VIC 15428 fits in the mLOM slot in the Cisco UCS C220 M6 (and C225 and C245) server and enables up to 50 Gbps of unified fabric connectivity to fabric interconnects for a total of 100 Gbps of connectivity per server. Cisco VIC 15428 connectivity to the fabric interconnects is delivered through 4x 25-Gbps links, with 2 port

channel links to each fabric interconnect. Cisco VIC 15428 supports 512 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 10. Cisco VIC 15238 in Cisco UCS C220 M6

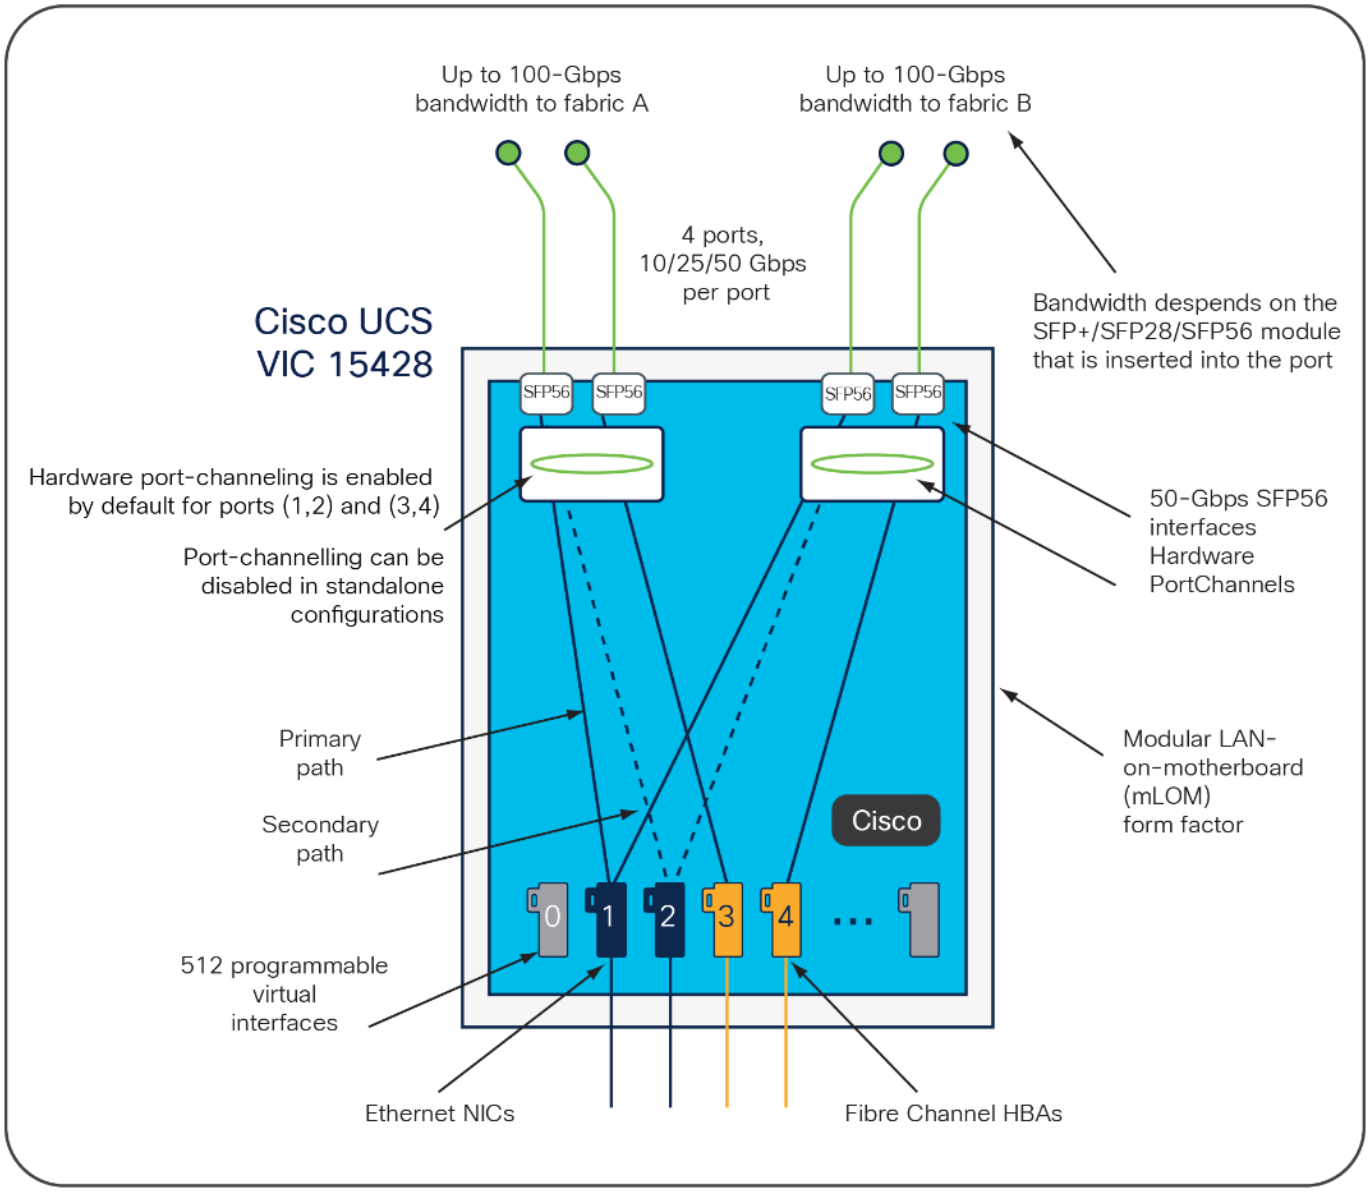

Various 4th generation VICs are supported with the Cisco UCS B- and C-Series M6 servers and have been explained in previous CVD Design and Deployment Guides.

Cisco UCS C225 M6 and Cisco UCS C245 M6 Rack Servers

The Cisco UCS C225 M6 and UCS C245 Rack Servers extend the capabilities of Cisco's UCS portfolio with the addition of the AMD EPYC CPUs as well as 16 DIMM slots per CPU for 3200-MHz DDR4 DIMMs with individual DIMM capacity points up to 256 GB. The maximum memory capacity for 2 CPUs is 8 TB (for 32 x 256 GB DDR4

DIMMs1). The Cisco UCS C225M6 has a 1-Rack-Unit (RU) form factor while the Cisco UCS C245 has a 2-RU form factor and can hold more GPUs than the Cisco UCS C225. These servers can connect directly to the Cisco UCS Fabric Interconnects at 2x100Gbps with 5th Generation Cisco UCS VIC 15238 and 4th Generation Cisco UCS VIC 1477 and 1495. These servers can also connect directly to the Cisco UCS Fabric Interconnects with the 5th Generation Cisco VIC 15428 and 4th Generation Cisco VICs 1467 and 1455. The Cisco UCS C-series servers can also connect to the FI using the Cisco Nexus 93180YC-FX3 in FEX-mode. These servers are supported in FlexPod but were not validated in this project.

Figure 11. Cisco UCS C225 M6 Rack Server

Figure 12. Cisco UCS C245 M6 Rack Server

Cisco UCS Sustainability and Security

From a Sustainability perspective, Cisco UCS server BIOS token settings are configured in this solution to provide power conservation while minimally affecting performance. Also, the Intelligent Platform Management Interface (IPMI) is configured on each server to allow interaction with VMware Distributed Power Management (DPM), described later in this document. Finally, Cisco UCS Manager provides power policies at both the Cisco UCS Chassis and Server level, which allow customers to adjust the balance of power consumption and performance to meet application needs.

From a Security perspective, all Cisco UCS user interfaces are hardened with the latest security ciphers and protocols including redirection of http to https, password and password expiry policies, integration with secure authentication systems, etc. Additionally, Cisco UCS servers support confidential computing (both Intel SGX and AMD based), although confidential computing is not addressed in this CVD. Finally, almost all Cisco UCS servers now sold come with Trusted Platform Modules (TPMs), that in VMware allows attestation of Unified Extended Firmware Interface Forum (UEFI) secure boot, which allows only securely signed code to be loaded. Many of the latest available operating systems, such as Microsoft Windows 11 require a TPM. The latest versions of VMware allow the assignment of a virtual TPM to VMs running operating systems that require a TPM.

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 13. Cisco Intersight Overview

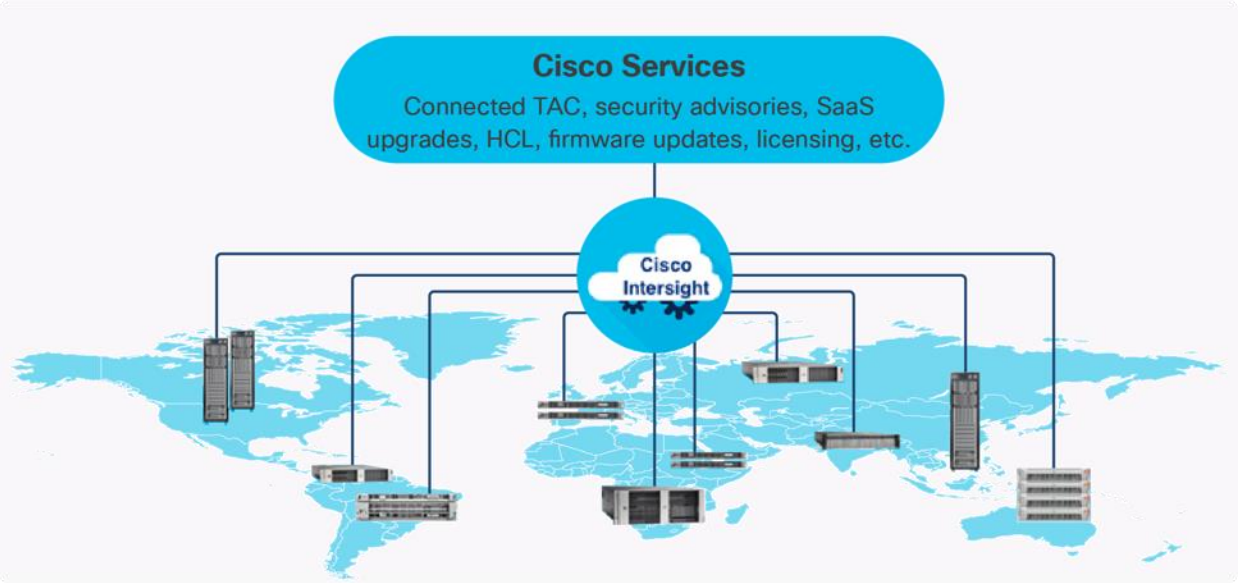

The main benefits of Cisco Intersight infrastructure services are as follows:

- 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 other 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, but some features are reduced

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 directly with it, will need a connection mechanism. Cisco Intersight Assist provides that connection mechanism. In FlexPod, VMware vCenter, Cisco Nexus Switches, Cisco MDS Switches, and NetApp Active IQ Unified Manager connect to Intersight with the help of the Intersight Assist VM.

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 explained 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 9300 Series Switches offer both modular and fixed 1/10/25/40/100/400 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 14. Cisco Nexus 933180YC-FX Switch

The Cisco Nexus 9000 series switch featured in this design is the Cisco Nexus 93180YC-FX 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-FX Switch is a 1RU switch that supports 3.6 Tbps of bandwidth and 1.2 bpps. The 48 downlink ports on the 93180YC-FX2 can support 1-, 10-, or 25-Gbps Ethernet or 16- or 32-Gbps Fibre Channel ports, offering deployment flexibility and investment protection. The 6 uplink ports can be configured as 40- or 100-Gbps Ethernet, offering flexible migration options. If more scale or 100 Gbps ports are needed, the Cisco Nexus 93360YC-FX2 or 9336C-FX2-E switches can be used.

The Cisco Nexus 93180YC-FX, 93360YC-FX2, and 9336C-FX2-E switches now support SAN switching, allowing both Ethernet and Fibre Channel SAN switching in a single switch. In addition to 16- or 32-Gbps Fibre Channel, these switches also support 100-Gbps FCoE, allowing port-channeled up to 100-Gbps FCoE uplinks from the Cisco UCS 6454 Fabric Interconnects to Cisco Nexus switches in SAN switching mode.

Cisco MDS 9132T 32G Multilayer Fabric Switch

The Cisco MDS 9132T 32G Multilayer Fabric Switch is the current 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 15. 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 MDS 9124V 64G 24-Port Fibre Channel Switch

The next-generation Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switch (Figure 16) supports 64, 32, and 16 Gbps Fibre Channel ports and provides high-speed Fibre Channel connectivity for all-flash arrays and high-performance hosts. This switch offers state-of-the-art analytics and telemetry capabilities built into its next-generation Application-Specific Integrated Circuit (ASIC) chipset. This switch allows seamless transition to Fibre Channel Non-Volatile Memory Express (NVMe/FC) workloads whenever available without any hardware upgrade in the SAN. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the benefits of greater bandwidth, scale, and consolidation. This switch is now orderable from Cisco, is supported in FlexPod, but was not validated in this design.

Figure 16. **Cisco MDS 9124V 64G 24-Port Fibre Channel Switch**

The Cisco MDS 9124V 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 Cisco 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. The Cisco MDS 9148V 48-Port Fibre Channel Switch is also available when more ports are needed.

Cisco Nexus Dashboard Fabric Controller (NDFC) SAN

NDFC SAN can be used to monitor, configure, and analyze Cisco 32Gbps Fibre Channel fabrics. Cisco NDFC SAN is deployed as an app on Cisco Nexus Dashboard. A single-server instance of the virtualized Nexus Dashboard with NDFC SAN and SAN Analytics is supported. Once the Cisco MDS switches and Cisco UCS Fabric Interconnects are added with the appropriate credentials and licensing, monitoring of the SAN fabrics can begin. Additionally, VSANs, device aliases, zones, and zone sets can be added, modified, and deleted using the NDFC 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.

NetApp AFF A-Series Storage

NetApp AFF A-Series controller lineup provides industry leading performance while continuing to provide a full suite of enterprise-grade data services for a shared environment across on-premises data centers and the cloud. Powered by NetApp ONTAP data management software, NetApp AFF A-Series systems deliver the industry's highest performance, superior flexibility, and best-in-class data services and cloud integration to help you accelerate, manage, and protect business-critical data across your hybrid clouds. As the first enterprise-grade storage systems to support both FC-NVMe and NVMe-TCP, AFF A-Series systems boost performance with modern network connectivity. These systems deliver the industry's lowest latency for an enterprise all-flash array, making them a superior choice for running the most demanding workloads and AI/DL applications. With a simple software upgrade to the modern FC-NVMe or NVMe-TCP SAN infrastructure, you can run more workloads with faster response times, without disruption or data migration.

NetApp offers a wide range of AFF-A series controllers to meet varying demands of the field. The high-end NetApp AFF A900 systems have a highly resilient design that enables non-disruptive in-chassis upgrades. It de-

livers latency as low as 100µs with FC-NVMe technology. The A800 delivers high performance in a compact form factor and is especially suited for EDA and Media & Entertainment workloads. The midrange, most versatile AFF A400 system features hardware acceleration technology that significantly enhances performance and storage efficiency. The budget friendly, the NetApp AFF A150 is an excellent entry-level performance flash option for customers.

NetApp AFF A150

The NetApp AFF A150 entry-level performance all-flash array provides 40% more performance compared with its predecessor. In addition to the performance enhancement, the NetApp AFF A150 system also supports more expansion options than its predecessor. It supports 24 internal 960GB, 3.8TB, and 7.6TB SAS SSD drives and up to two external expansion shelves for a maximum of 72 SAS SSDs per HA pair. The flexible AFF A150 system can be tailored to meet various solution requirements, including starting very small with 8 x 960GB SSD drives.

The NetApp AFF A150 offers 10GbE ports for IP based transport or Unified Target Adapter 2 (UTA2) ports for either 10GbE Ethernet connectivity for IP-based traffic or 16Gb FC connectivity for FC and FC-NVMe traffic. The two miniSAS ports can be used to connect up to two expansion shelves per HA pair. Additional AFF A150, or compatible AFF / FAS HA pairs, can be added to the cluster to scale out the solution to meet the performance requirements, subject to the platform mixing rules and support limits. Customers can start protecting their business with AFF A150 by taking advantage of the ONTAP data protection features to create instantaneous Snapshots, set up SnapMirror data replication, and deploy MetroCluster IP or SnapMirror Business Continuity solutions for disaster recovery and to ensure business continuity.

Figure 17. **NetApp AFF A150 Front View**

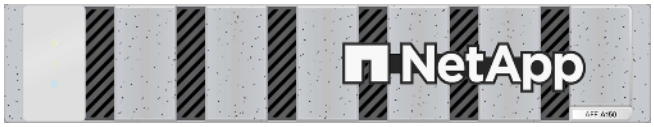

Figure 18. **NetApp AFF A150 Rear View (with UTA2)**

NetApp AFF A400

The NetApp AFF A400 offers full end-to-end NVMe support. The frontend FC-NVMe connectivity makes it possible to achieve optimal performance from an all-flash array for workloads that include artificial intelligence, machine learning, and real-time analytics as well as business-critical databases. The frontend NVMe-TCP connectivity enables customers to take advantage of NVMe technology over existing ethernet infrastructure for faster host connectivity. On the back end, the A400 supports both serial-attached SCSI (SAS) and NVMe-attached SSDs, offering the versatility for current customers to move up from their legacy A-Series systems and satisfying the increasing interest that all customers have in NVMe-based storage.

The NetApp AFF A400 offers greater port availability, network connectivity, and expandability. The NetApp AFF A400 has 10 PCIe Gen3 slots per high availability pair. The NetApp AFF A400 offers 10GbE, 25GbE and 100GbE ports for IP based transport, and 16/32Gb ports for FC and FC-NVMe traffic. This model was created to keep up with changing business needs and performance and workload requirements by merging the latest technology for data acceleration and ultra-low latency in an end-to-end NVMe storage system.

Figure 19. NetApp AFF A400 Front View

Figure 20. NetApp AFF A400 Rear View

NetApp AFF A800

The NetApp AFF A800 is a higher end model that offers superior performance and higher port count (both 32G FC and 100G Ethernet) than NetApp AFF A400. NetApp AFF A800 single chassis HA Pair supports 48 internal SSD drives and up to 8 external NS224 shelves allowing up to 240 NVMe SSD drives. It offers ultra-low latency of 100us and up to 300 GB/s throughput enabling it to be an ultimate choice to power data hungry applications such as artificial intelligence, deep learning, and big data analytics.

Figure 21. NetApp AFF A800 Front View

Figure 22. NetApp AFF A800 Rear View

For more information about the NetApp AFF A-series controllers, see the NetApp AFF product page: https://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx.

You can view or download more technical specifications of the NetApp AFF A-series controllers here: https://www.netapp.com/pdf.html?item=/media/7828-DS-3582-AFF-A-Series.pdf

NetApp AFF C-Series Storage

NetApp AFF C-Series storage systems help move more data to flash with the latest high-density NVMe QLC capacity flash technology. These systems are suited for large-capacity deployment with a small footprint as an affordable way to modernize data center to all flash and also connect to the cloud. Powered by NetApp ONTAP data management software, NetApp AFF C-Series systems deliver industry-leading efficiency, superior flexibility, and best-in-class data services and cloud integration to help scale IT infrastructure, simplify data management, reduce storage cost, rack space usage, power consumption, and improve sustainability significantly.

NetApp offers several AFF-C series controllers to meet varying demands of the field. The high-end NetApp AFF C800 systems offer superior performance. The midrange NetApp AFF C4004 delivers high performance and good expansion capability. The entry-level NetApp AFF C250 system balanced performance, connectivity, and expansion options for a small footprint deployment.

NetApp AFF C250

The NetApp AFF C250 is an entry-level small form-factor capacity flash model. The 2U dual-controller system supports 24 internal drives for space efficient deployment. The NetApp AFF C250 offers scale-out performance, storage expansion, flexibility in network connectivity, and a rich set of data management and data protection capabilities powered by NetApp ONTAP software.

The NetApp AFF C250 offers both 25 GbE and 100 GbE Ethernet connectivity as well as 32Gb FC connectivity for deploying reliable Ethernet and FC solutions. By adding external NVMe expansion shelves for additional NVMe QLC SSD, the platform is capable of meeting the substantial capacity needs of the data centers.

Figure 23. NetApp AFF C250 Front View

Figure 24. NetApp AFF C250 Rear View

NetApp AFF C400

The NetApp AFF C400 is a midrange model which offers full end-to-end NVMe support. The frontend FC-NVMe and NVMe-TCP connectivity enables customers to take advantage of NVMe technology over existing FC and Ethernet infrastructure for faster host connectivity. On the back end, the NetApp AFF C400 supports both serial-attached SCSI (SAS) and NVMe-attached SSDs, offering the versatility for current customers to move up from their existing systems and adopt NVMe-based storage.

Compared to the entry-level NetApp AFF C250 model, the NetApp AFF C400 offers greater port availability, network connectivity, and expandability. The NetApp AFF C400 has 10 PCIe Gen3 slots per high availability pair.

The NetApp AFF C400 offers 10GbE, 25GbE and 100GbE ports for IP based transport, and 16/32Gb ports for FC and FC-NVMe traffic.

Figure 25. NetApp AFF C400 Front View

Figure 26. NetApp AFF C400 Rear View

NetApp AFF C800

The NetApp AFF C800 is a higher end model that offers superior performance and higher port count (both 32G FC and 100G Ethernet) than NetApp AFF C400. The NetApp AFF C800 single chassis HA Pair supports 48 internal SSD drives and up to 8 external NS224 shelves allowing up to 240 NVMe SSD drives. It is an ultimate choice to power data hungry applications such as artificial intelligence, deep learning, and big data analytics.

Figure 27. NetApp AFF C800 Front View

Figure 28. NetApp AFF C800 Rear View

For more information about the NetApp AFF C-series controllers, see the NetApp AFF C-Series product page: https://www.netapp.com/data-storage/aff-c-series/

You can view or download more technical specifications of the AFF C-Series controllers here: https://www.netapp.com/media/81583-da-4240-aff-c-series.pdf

You can look up the detailed NetApp storage product configurations and limits here: https://hwu.netapp.com/

Note: FlexPod CVDs provide reference configurations and there are many more supported IMT configurations that can be used for FlexPod deployments, including NetApp hybrid storage arrays.

NetApp ONTAP 9.12.1

NetApp storage systems harness the power of ONTAP to simplify the data infrastructure from edge, core, and cloud with a common set of data services and 99.9999 percent availability. NetApp ONTAP 9 data management software from NetApp enables customers to modernize their infrastructure and transition to a cloud-ready data center. NetApp ONTAP 9 has a host of features to simplify deployment and data management, accelerate and protect critical data, and make infrastructure future-ready across hybrid-cloud architectures.

NetApp ONTAP 9 is the data management software that is used with the NetApp AFF A400 and A800 all-flash storage systems in this solution design. ONTAP software offers secure unified storage for applications that read and write data over block- or file-access protocol storage configurations. These storage configurations range from high-speed flash to lower-priced spinning media or cloud-based object storage. ONTAP implementations can run on NetApp engineered AFF, FAS, or ASA series arrays and in private, public, or hybrid clouds (NetApp Private Storage and NetApp Cloud Volumes ONTAP). Specialized implementations offer best-in-class converged infrastructure, featured here as part of the FlexPod Datacenter solution or with access to third-party storage arrays (NetApp FlexArray virtualization). Together these implementations form the basic framework of the NetApp Data Fabric, with a common software-defined approach to data management, and fast efficient replication across systems. FlexPod and ONTAP architectures can serve as the foundation for both hybrid cloud and private cloud designs.

The NetApp AFF C-Series family of capacity flash storage system comes with NetApp ONTAP One, the all-in-one software license for on-premises operations. At any time, the AFF C-Series customers can start using and taking advantage of the rich ONTAP data management capabilities.

Read more about the capabilities of ONTAP data management software here: https://www.netapp.com/us/products/data-management-software/ontap.aspx.

For more information on new features and functionality in latest NetApp ONTAP software, refer to the NetApp ONTAP release notes: NetApp ONTAP 9 Release Notes (netapp.com)

Note: The support for the AFF A150 and the AFF C-Series platforms was introduced with ONTAP 9.12.1P1.

NetApp Storage Sustainability

Data centers consume a significant amount of electricity and contribute to global greenhouse gas emissions. NetApp is providing lifetime carbon footprint estimates to help customers better understand the environmental impacts of NetApp storage systems.

NetApp uses Product Attribute to Impact Algorithm (PAIA) to calculate the carbon emissions associated with a product through its lifecycle, including acquisition of raw materials, manufacturing, distribution, product use, and

final disposition. PAIA is a streamlined lifecycle assessment (LCA) methodology for assessing environmental impacts associated with the entire lifecycle of a product. The PAIA model was developed by the Materials Systems Laboratory at the Massachusetts Institute of Technology (MIT) and is a leading and globally accepted methodology for streamlining the product carbon footprint process.

You can use ONTAP REST API to access environment data from the ONTAP storage system for sustainability assessments. You can also utilize NetApp Harvest tool, which is an open-metrics endpoint for ONTAP and StorageGRID, to collect performance, capacity, hardware, and environmental metrics and display them in Grafana dashboards to gain sustainability insights.

Note: For more information on NetApp storage system environmental certifications and product carbon footprint report, ONTAP REST API, and NetApp Harvest, refer to the following respective references:

- https://www.netapp.com/company/environmental-certifications/
- https://docs.netapp.com/us-en/ontap-automation/reference/api_reference.html
- https://github.com/NetApp/harvest

NetApp Storage Security and Ransomware Protection

NetApp storage administrators use local or remote login accounts to authenticate themselves to the cluster and storage VM. Role-Based Access Control (RBAC) determines the commands to which an administrator has access. In addition to RBAC, NetApp ONTAP supports multi-factor authentication (MFA) and multi-admin verification (MAV) to enhance the security of the storage system.

With NetApp ONTAP, you can use the security login create command to enhance security by requiring that administrators log in to an admin or data SVM with both an SSH public key and a user password. Beginning with NetApp ONTAP 9.12.1, you can use Yubikey hardware authentication devices for SSH client MFA using the FIDO2 (Fast IDentity Online) or Personal Identity Verification (PIV) authentication standards.

With NetApp ONTAP 9.11.1, you can use multi-admin verification (MAV) to ensure that certain operations, such as deleting volumes or Snapshot copies, can be executed only after approvals from designated administrators. This prevents compromised, malicious, or inexperienced administrators from making undesirable changes or deleting data.

Also with NetApp ONTAP 9.10.1, the Autonomous Ransomware Protection (ARP) feature uses workload analysis in NAS (NFS and SMB) environments to proactively detect and warn about abnormal activity that might indicate a ransomware attack.

While NetApp ONTAP includes features like FPolicy, Snapshot copies, SnapLock, and Active IQ Digital Advisor to help protect from ransomware, ARP utilizes machine-learning and simplifies the detection of and the recovery from a ransomware attack.

ARP can detect the spread of most ransomware attacks after only a small number of files are encrypted, take action automatically to protect data, and alert you that a suspected attack is happening.

When an attack is suspected, the system takes a volume Snapshot copy at that point in time and locks that copy. If the attack is confirmed later, the volume can be restored to this proactively taken snapshot to minimize the data loss.

For more information on MFA, MAV, and ransomware protection, refer to the following:

- https://docs.netapp.com/us-en/ontap/authentication/setup-ssh-multifactor-authentication-task.html
- https://docs.netapp.com/us-en/ontap/multi-admin-verify/
- https://www.netapp.com/pdf.html?item=/media/17055-tr4647pdf.pdf
- https://docs.netapp.com/us-en/ontap/anti-ransomware/

NetApp Active IQ Unified Manager

NetApp Active IQ Unified Manager is a comprehensive monitoring and proactive management tool for NetApp ONTAP systems to help manage the availability, capacity, protection, and performance risks of your storage systems and virtual infrastructure. The Unified Manager can be deployed on a Linux server, on a Windows server, or as a virtual appliance on a VMware host.

Active IQ Unified Manager enables monitoring your ONTAP storage clusters from a single redesigned, intuitive interface that delivers intelligence from community wisdom and AI analytics. It provides comprehensive operational, performance, and proactive insights into the storage environment and the virtual machines running on it. When an issue occurs with the storage infrastructure, Unified Manager can notify you about the details of the issue to help with identifying the root cause. The virtual machine dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the vSphere host down through the network and finally to the storage. Some events also provide remedial actions that can be taken to rectify the issue. You can configure custom alerts for events so that when issues occur, you are notified through email and SNMP Traps. Active IQ Unified Manager enables planning for the storage requirements of your users by forecasting capacity and usage trends to proactively act before issues arise, preventing reactive short-term decisions that can lead to additional problems in the long term.

For more information on NetApp Active IQ Unified Manager, go to:
https://docs.netapp.com/us-en/active-iq-unified-manager/

NetApp ONTAP Tools for VMware vSphere

The NetApp ONTAP tools for VMware vSphere provides end-to-end life cycle management for virtual machines in VMware environments that use NetApp storage systems. It simplifies storage and data management for VMware environment by enabling administrators to directly manage storage within the vCenter Server.

Note: Each component in ONTAP tools provides capabilities to help manage your storage more efficiently.

Virtual Storage Console (VSC)

VSC enables you to perform the following tasks:

- Add storage controllers, assign credentials, and set up permissions for storage controllers of VSC, that both SRA and VASA Provider can leverage
- Provision datastores
  - Monitor the performance of the datastores and virtual machines in your vCenter Server environment
  - View and update the host settings of the ESXi hosts that are connected to NetApp storage

- Manage access to the vCenter Server objects and ONTAP objects by using the vCenter Server role-based access control (RBAC) and ONTAP RBAC

VASA Provider

VASA Provider for ONTAP uses VMware vSphere APIs for Storage Awareness (VASA) to send information about storage used by VMware vSphere to the vCenter Server. ONTAP tools has VASA Provider integrated with VSC. VASA Provider enables you to perform the following tasks:

- Provision VMware Virtual Volumes (vVols) datastores
  - Create and use storage capability profiles that define different storage service level objectives (SLOs) for your environment
  - Verify for compliance between the datastores and the storage capability profiles
  - Set alarms to warn you when volumes and aggregates are approaching the threshold limits
  - Monitor the performance of virtual machine disks (VMDKs) and the virtual machines that are created on vVols datastores

Storage Replication Adapter (SRA)

SRA enables you to use array-based replication (ABR) for protected sites and recovery sites for disaster recovery in the event of a failure. When SRA is enabled and used in conjunction with VMware Site Recovery Manager (SRM), you can recover the vCenter Server datastores and virtual machines in the event of a failure.

Note: The ONTAP tools for VMware vSphere 9.12 release supports and interoperates with VMware vSphere 8.0. It also supports NVMe-oF vVols introduced with vSphere 8.0 in conjunction with Storage Policy Based Management for performance and availability requirement configurations. For more information on ONTAP tools for VMware vSphere, go to: https://docs.netapp.com/us-en/ontap-tools-vmware-vsphere/index.html

NetApp SnapCenter

SnapCenter Software is a simple, centralized, scalable platform that provides application consistent data protection for applications, databases, host file systems, and VMs running on ONTAP systems anywhere on premise or in the Hybrid Cloud.

SnapCenter leverages NetApp Snapshot, SnapRestore, FlexClone, SnapMirror, and SnapVault technologies to provide:

- Fast, space-efficient, application-consistent, disk-based backups
- Rapid, granular restore, and application-consistent recovery
- Quick, space-efficient cloning

SnapCenter includes both SnapCenter Server and individual lightweight plug-ins. You can automate deployment of plug-ins to remote application hosts, schedule backup, verification, and clone operations, and monitor all data protection operations.

Data protection is supported for Microsoft Exchange Server, Microsoft SQL Server, Oracle Databases on Linux or AIX, SAP HANA database, and Windows Host Filesystems running on ONTAP systems. It is also supported for

other standard or custom applications and databases by providing a framework to create user-defined SnapCenter plug-ins. You may install only the plug-ins that are appropriate for the data that you want to protect.

Note: For more information on SnapCenter 4.8, refer to the SnapCenter software documentation: https://docs.netapp.com/us-en/snapcenter/index.html

NetApp BlueXP

NetApp BlueXP is a unified control plane that provides a hybrid multicloud experience for storage and data services across on-premises and cloud environments. NetApp BlueXP is an evolution of Cloud Manager and enables the management of your NetApp storage and data assets from a single interface.

You can use BlueXP to move, protect, and analyze data, and to control on-prem storage devices like ONTAP, E-Series, and StorgeGRID, and to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files).

The BlueXP backup and recovery service provides efficient, secure, and cost-effective data protection for NetApp ONTAP data, Kubernetes persistent volumes, databases, and virtual machines, both on premises and in the cloud. Backups are automatically generated and stored in an object store in your public or private cloud account.

BlueXP ransomware protection provides a single point of visibility and control to manage and to refine data security across various working environments and infrastructure layers to better respond to threats as they occur.

Note: For more information on BlueXP, refer to the BlueXP documentation: https://docs.netapp.com/us-en/cloud-manager-family/

VMware vSphere 8.0

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

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

- Limits with vSphere 8 have been increased including number of GPU devices is increased to 8, the number of ESXi hosts that can be managed by Lifecycle Manager is increased from 400 to 1000, the maximum number of VMs per cluster is increased from 8,000 to 10,000, and the number of VM DirectPath I/O devices per host is increased from 8 to 32.
- Security improvements including adding an SSH timeout on ESXi hosts, a TPM Provisioning policy allowing a vTPM to be replaced when cloning VMs, and TLS 1.2 as the minimum supported TLS version.
- Implementation of VMware vMotion Unified Data Transport (UDT) to significantly reduce the time to storage migrate powered off virtual machines.
- Lifecycle Management improvements including VMware vSphere Configuration Profiles as a new alternative to VMware Host Profiles, staging cluster images and remediating up to 10 ESXi hosts in parallel instead of one at a time.

- New Virtual Hardware in VM hardware version 20 supporting the latest guest operating systems, including Windows 11.
- Distributed Resource Scheduler and vMotion improvements.
- Implementation of the VMware Balanced Power Management Policy on each server, which reduces energy consumption with minimal performance compromise.
- Implementation of VMware Distributed Power Management, which along with configuration of the Intelligent Platform Management Interface (IPMI) on each Cisco UCS server allows a VMware host cluster to reduce its power consumption by powering hosts on and off based on cluster resource utilization.

For more information about VMware vSphere and its components, go to: 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.

Cisco Intersight Assist Device Connector for VMware vCenter, NetApp ONTAP, and Cisco Nexus and MDS Switches

Cisco Intersight integrates with VMware vCenter, NetApp storage, and Cisco Nexus switches as follows:

- Cisco Intersight uses the device connector running within the 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 NetApp Active IQ Unified Manager. The NetApp AFF A400/A800 should be added to NetApp Active IQ Unified Manager.
- Cisco Intersight uses the device connector running within the Cisco Intersight Assist virtual appliance to communicate with Cisco Nexus 9000 and MDS switches.

Figure 29. Cisco Intersight and vCenter/NetApp/Cisco Switch Integration

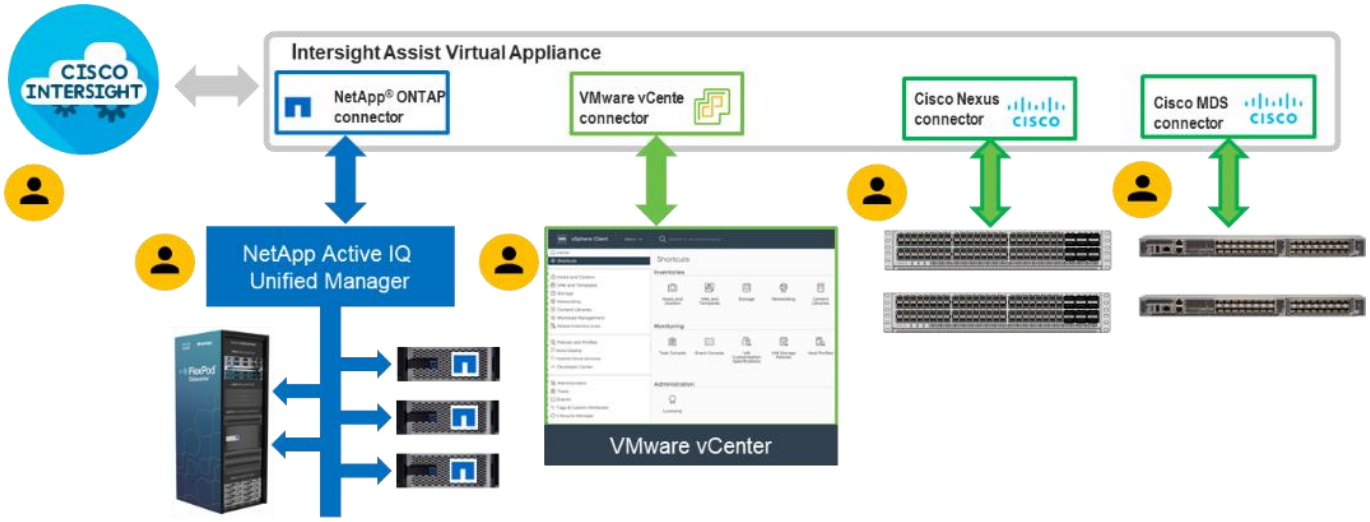

The device connector provides a secure way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware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

Enterprise SAN and NAS workloads can benefit equally from the integrated management solution. The integration architecture enables FlexPod customers to use new management capabilities with no compromise in their existing VMware, ONTAP, or switch 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, NetApp Active IQ Unified Manager, and Cisco Switch Interfaces for comprehensive analysis, diagnostics, and reporting of virtual, storage, and switching environments. The functionality provided through this integration is explained in the upcoming solution design section.

Solution Design

This chapter contains the following:

The FlexPod Datacenter with Cisco UCS M6 servers, NetApp ONTAP 9.12.1, and VMware vSphere 8.0 solution extends the customer investment in Cisco UCS M6 hardware with the latest Cisco UCS, NetApp, and VMware software. The VMware ESXi 8.0 hypervisor is installed on the Cisco UCS B200 M6 and Cisco UCS C220 M6 Compute Nodes configured for stateless compute design using boot from SAN. The NetApp AFF A400 provides the storage infrastructure required for setting up the VMware environment. The Cisco Intersight platform is utilized to monitor the infrastructure.

Requirements

The FlexPod Datacenter with Cisco UCS M6 servers, NetApp ONTAP 9.12.1, and VMware vSphere 8.0 design 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 capacity, storage, 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

FlexPod Datacenter with Cisco UCS M6 servers, NetApp ONTAP 9.12.1, and VMware vSphere 8.0 supports both IP and Fibre Channel (FC)–based storage access design. For the IP-based solution, iSCSI configuration on Cisco UCS and NetApp AFF A400 is utilized to set up boot from SAN for the Compute Nodes. For the FC designs, NetApp AFF A400 and Cisco UCS Servers are connected through Cisco MDS 9132T Fibre Channel Switches and

boot from SAN uses the FC network. In both these designs, VMware ESXi hosts access the VM datastore volumes on NetApp using NFS. The physical connectivity details for both IP and FC designs are explained below.

IP-based Storage Access: iSCSI, NFS, and NVMe-TCP

The physical topologies for the IP-based FlexPod Datacenter are shown in Figure 30 and Figure 31.

Figure 30. FlexPod Datacenter Physical Topology for iSCSI, NFS, and NVMe-TCP with 25 Gbps Core Networking

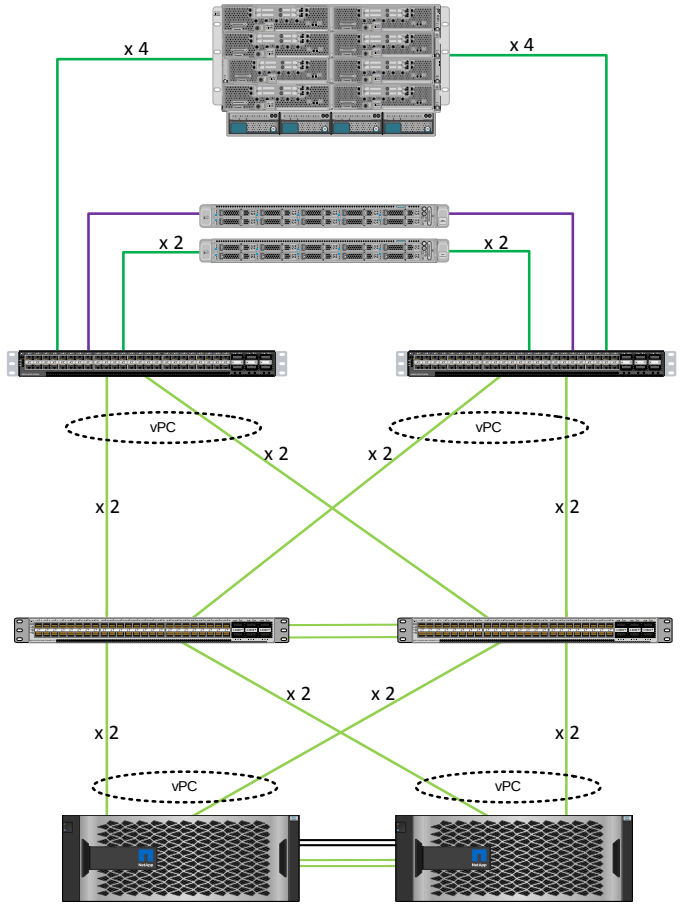

Figure 31. FlexPod Datacenter Physical Topology for iSCSI, NFS, and NVMe-TCP with 100 Gbps Core Networking

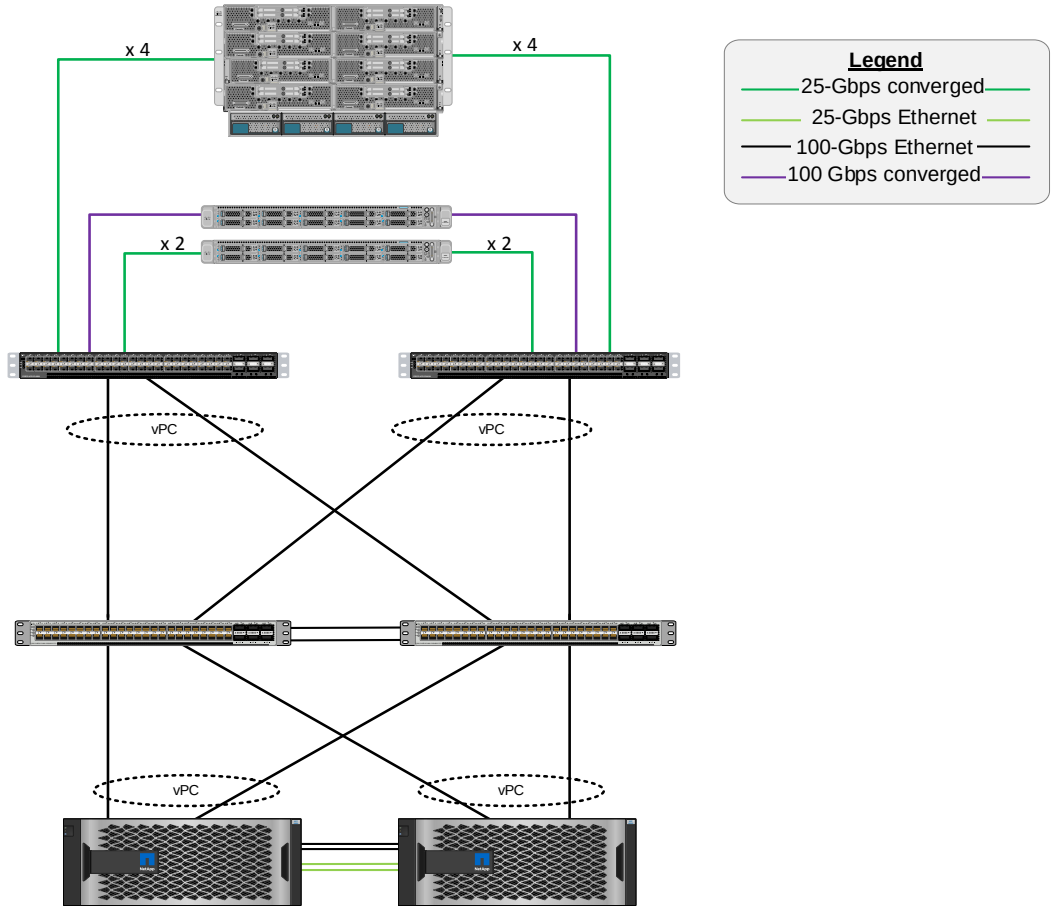

To validate the IP-based storage access in a FlexPod configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Chassis connects to fabric interconnects using Cisco UCS 2408 Fabric Extender modules, where four 25 Gigabit Ethernet ports are used on each IOM or Fabric Extender to connect to the appropriate FI. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCS B200 M6 Compute Nodes contain fifth-generation Cisco 15411 virtual interface cards (VICs).
- Cisco UCS C220 M6 with either fifth-generation Cisco UCS 15238 or 15428 VICs, or fourth-generation Cisco UCS VICs connect to the fabric interconnects with either 100GE or 25GE.
- Cisco Nexus 93180YC-FX Switches in Cisco NX-OS mode provide the switching fabric.
- In the 25GE design, 2-4 25GE links can be used for the vPC peer link.

- Cisco UCS 6454 Fabric Interconnect 25-Gigabit or 100-Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX Switches in a Virtual Port Channel (vPC) configuration.
- The NetApp AFF A400 controllers connect to the Cisco Nexus 93360YC-FX2 Switches using either for 25 GE ports or two 100 GE ports from each controller configured as a vPC.
- VMware 8.0 ESXi software is installed on all M6 Compute Nodes to validate the infrastructure.

FC-based Storage Access: FC, FC-NVMe, and NFS

The physical topologies for the FC-booted FlexPod Datacenter are shown in Figure 32 and Figure 33.

Figure 32. FlexPod Datacenter Physical Topology for FC, FC-NVMe, and NFS with 25 Gpbs Core Networking

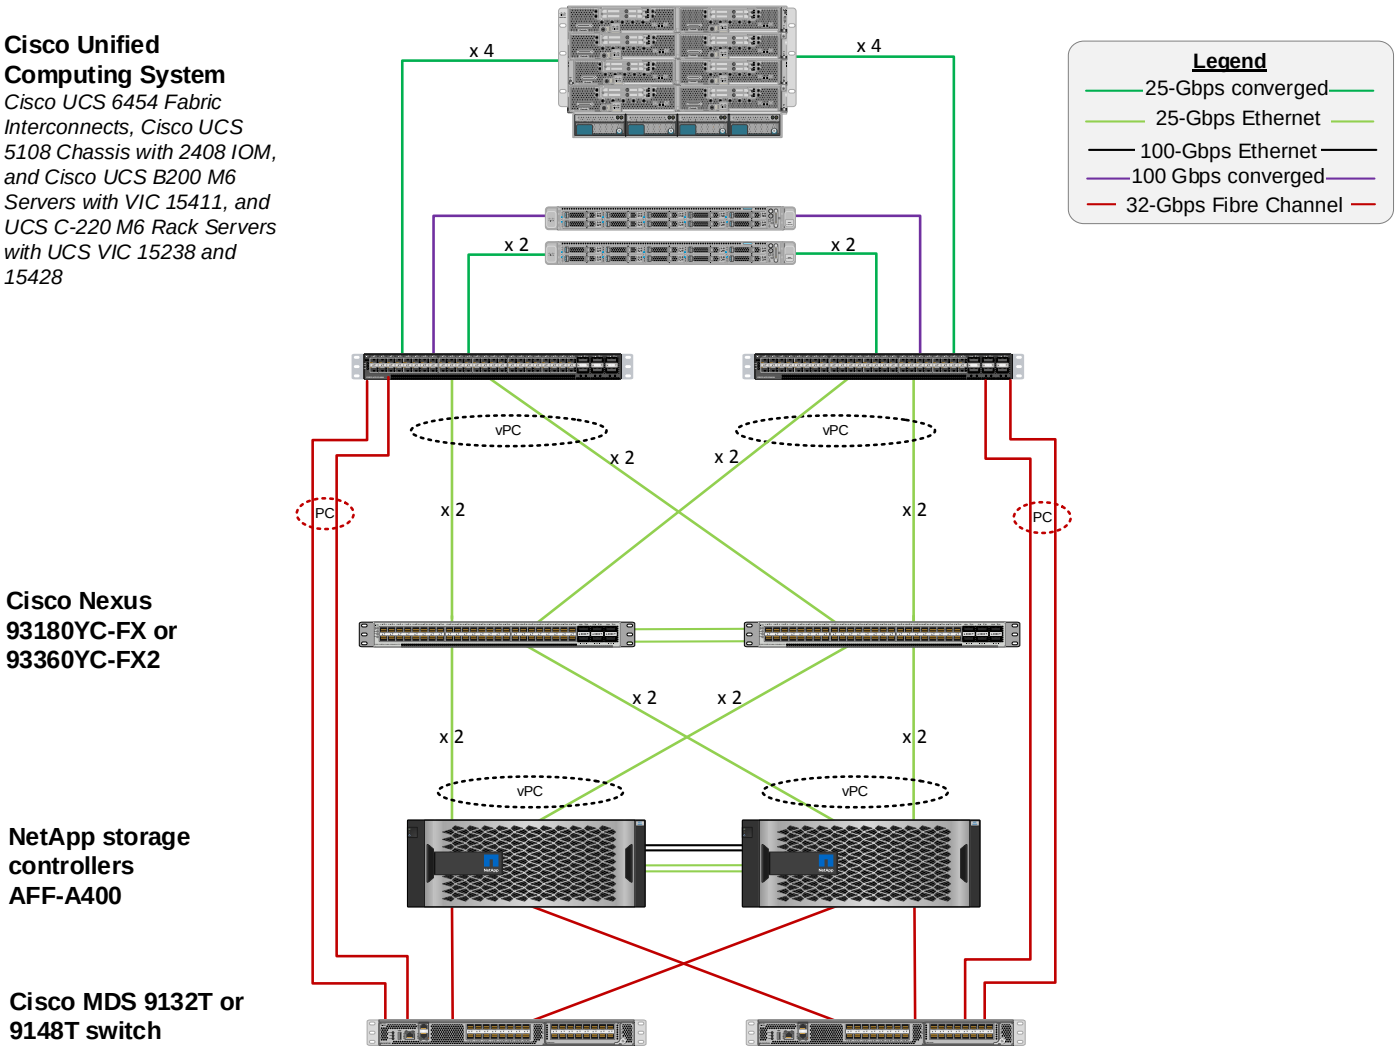

Figure 33. FlexPod Datacenter Physical Topology for FC, FC-NVMe, and NFS with 100 Gpbs Core Networking

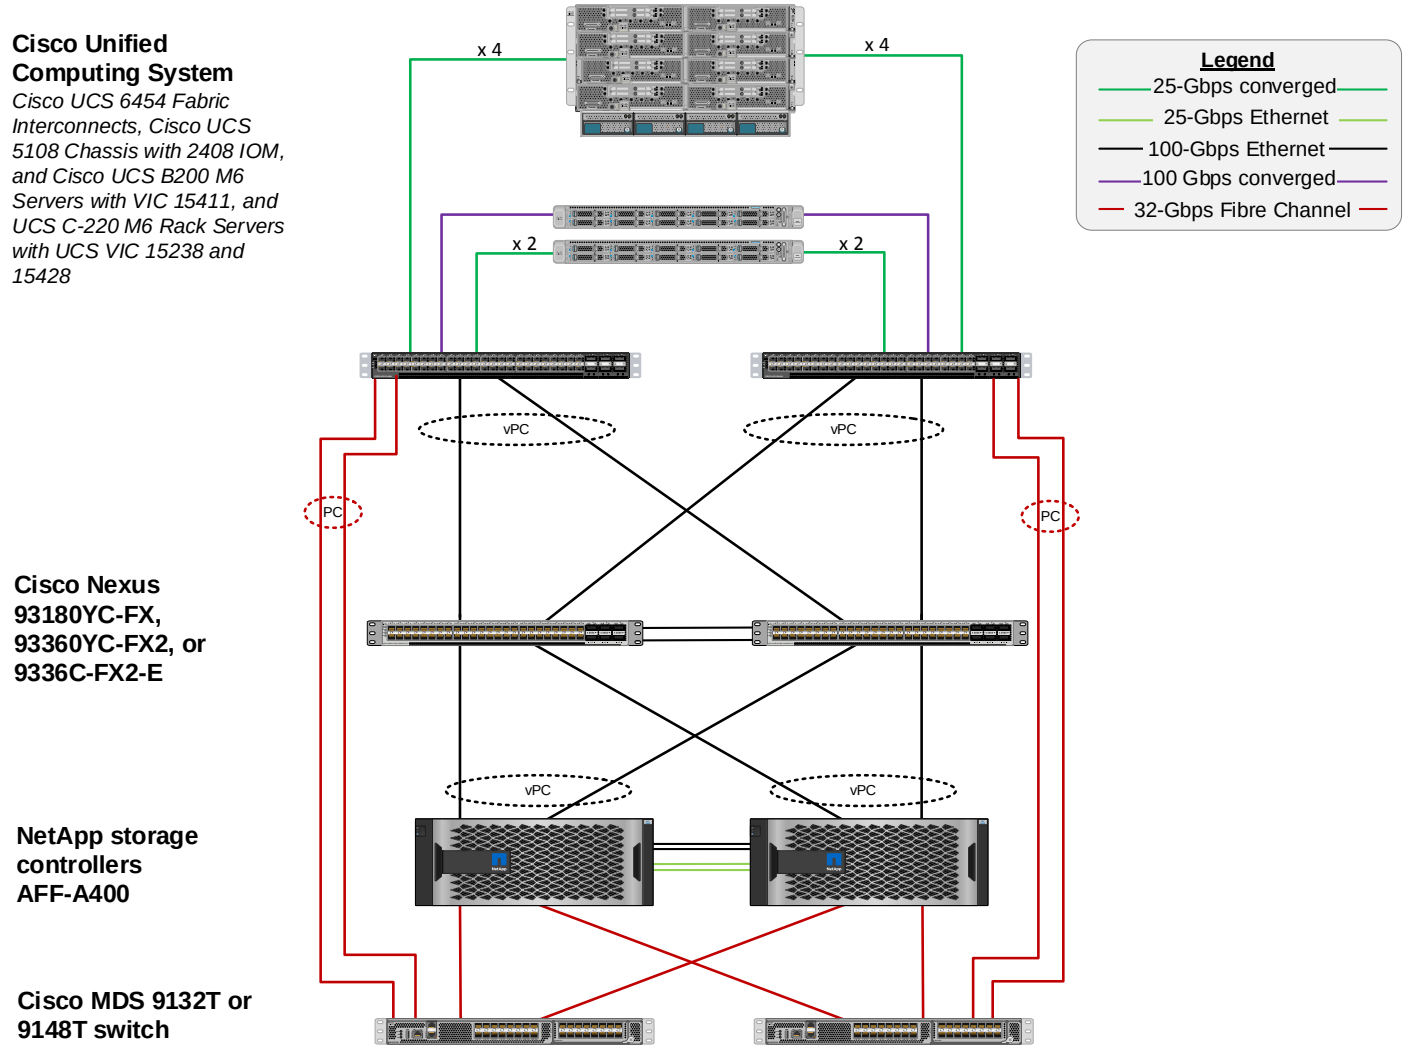

To validate the FC-based storage access in a FlexPod configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Chassis connects to the fabric interconnects using Cisco UCS 2408 Fabric Extenders, where four 25 Gigabit Ethernet ports are used on each Fabric Extender or IOM to connect to the appropriate FI.
- Cisco UCS B200 M6 Compute Nodes contain fifth-generation Cisco UCS 15411 VICs.
- Cisco UCS C220 M6 with either fifth-generation Cisco UCS 15238 or 15428 VICs or fourth-generation Cisco UCS VICs connect to the fabric interconnects with either 100GE or 25GE.
- Cisco Nexus 93180YC-FX Switches in Cisco NX-OS mode provide the switching fabric.
- In the 25GE design, 2-4 25GE links can be used for the vPC peer link.

- Cisco UCS 6454 Fabric Interconnect 25 Gigabit or 100 Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX Switches in a vPC configuration.
- The NetApp AFF A400 controller connects to the Cisco Nexus 93180YC-FX Switches using either for 25 GE ports or two 100 GE ports from each controller configured as a vPC for NFS traffic.
- The Cisco UCS 6454 Fabric Interconnects are connected to the Cisco MDS 9132T switches using multiple 32-Gbps Fibre Channel connections configured as a single port channel for SAN connectivity.
- The NetApp AFF A400 controllers connect to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections for SAN connectivity.
- VMware 8.0 ESXi software is installed on all Cisco UCS M6 Compute Nodes to validate the infrastructure.

FC-based Storage Access: FC, FC-NVMe, and NFS Utilizing Nexus SAN Switching

The physical topology for the FC-boot FlexPod Datacenter with Nexus SAN Switching is shown in Figure 34.

Figure 34. FlexPod Datacenter Physical Topology for FC, FC-NVMe, and NFS Utilizing Nexus SAN Switching

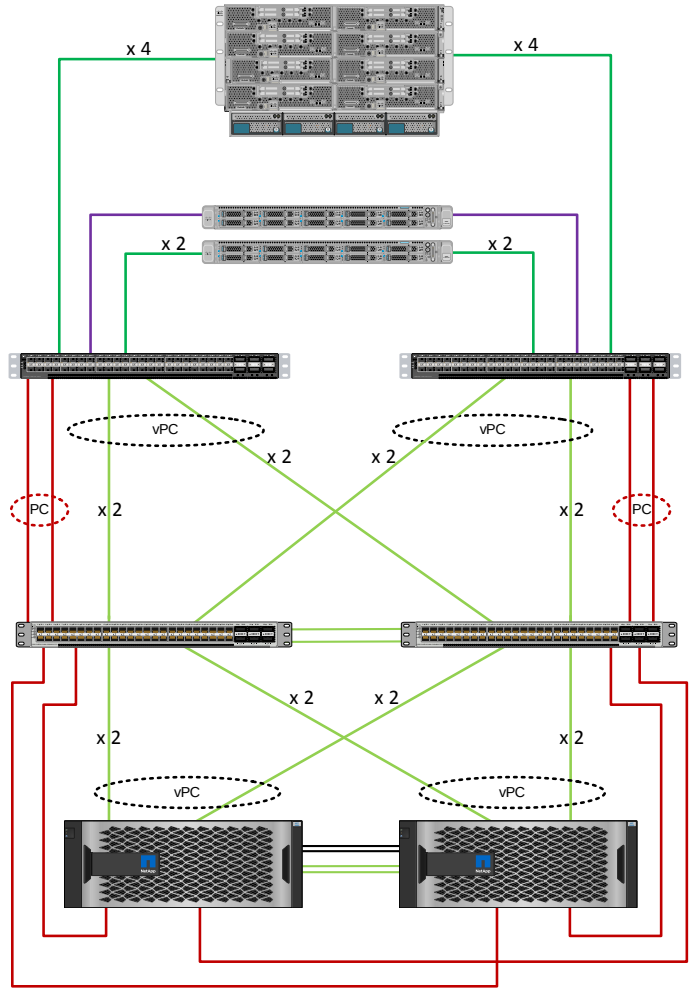

To validate the FC-based storage access in a FlexPod configuration with Nexus SAN switching, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Chassis connects to fabric interconnects using Cisco UCS 2408 Fabric Extenders, where four 25 Gigabit Ethernet ports are used on each Port Extender or IOM to connect to the appropriate FI.
- Cisco UCS B200 M6 Compute Nodes contain fifth-generation Cisco 15411 VICs.
- Cisco UCS C220 M6 with either fifth-generation Cisco UCS 15238 or 15428 VICs or fourth-generation Cisco UCS VICs connect to the fabric inter-connects with either 100GE or 25GE.
- Cisco Nexus 93180YC-FX Switches in Cisco NX-OS mode provide both the switching fabric and the SAN fabric.
- In the 25GE design, 2-4 25GE links can be used for the vPC peer link.
- Cisco UCS 6454 Fabric Interconnect 25 or 100 Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX Switches in a vPC configuration.
- The NetApp AFF A400 controller connects to the Cisco Nexus 93180YC-FX switches using either four 25 GE ports or two 100 GE ports from each controller configured as a vPC for NFS traffic.
- Cisco UCS 6454 Fabric Interconnects are connected to the Cisco Nexus 93180YC-FX switches using multiple 32-Gbps FC uplinks configured as a single FC port channel. Multiple 25 GE or 100 GE FCoE uplinks configured as a single Ethernet port channel can also be used.
- The NetApp AFF controllers connect to the Cisco Nexus 93180YC-FX switches using 32-Gbps Fibre Channel connections for SAN connectivity.
- VMware 8.0 ESXi software is installed on Cisco M6 Compute Nodes to validate the infrastructure.
- From a sustainability perspective, this configuration does not have the power draw of two MDS switches.

VLAN Configuration

Table 1 lists VLANs configured for setting up the FlexPod environment along with their usage.

Table 1. VLAN Usage

VLAN ID	Name	Usage
2	Native-VLAN	Use VLAN 2 as native VLAN instead of default VLAN (1)
13	OOB-MGMT-VLAN	Out-of-band management VLAN to connect management ports for various devices
113	IB-MGMT-VLAN	In-band management VLAN utilized for all in-band management connectivity - for example, ESXi hosts, VM management, etc.
800	VM-Traffic	VM data traffic VLAN

VLAN ID	Name	Usage
3050	NFS-VLAN	NFS VLAN for mounting datastores in ESXi servers for VMs
3010*	iSCSI-A	iSCSI-A path for boot-from-san traffic
3020*	iSCSI-B	iSCSI-B path for boot-from-san traffic
3030*	NVMe-TCP-A	NVMe-TCP-A path for NVMe datastores
3040*	NVMe-TCP-B	NVMe-TCP-B path for NVMe datastores
3000	vMotion	VMware vMotion traffic

* iSCSI and NVMe-TCP VLANs are not required if using FC storage access.

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

- VLAN 13 allows customers to manage and access out-of-band management interfaces of various devices and is brought into the infrastructure to allow CIMC access to the Cisco UCS servers and is also available to infrastructure virtual machines (VMs). Interfaces in this VLAN are configured with MTU 1500.
- VLAN 113 is used for in-band management of VMs, ESXi hosts, and other infrastructure services. Interfaces in this VLAN are configured with MTU 1500.
- VLAN 3050 provides ESXi hosts access to the NFS datastores hosted on the NetApp Controllers for deploying VMs. Interfaces in this VLAN are configured with MTU 9000.
- A pair of iSCSI VLANs (3010 and 3020) is configured to provide access to boot LUNs for ESXi hosts and iSCSI datastores. These VLANs are not needed when configuring Fibre Channel connectivity. Interfaces in these VLANs are configured with MTU 9000.
- A pair of NVMe-TCP VLANs (3030 and 3040) is configured to provide access to NVMe datastores. These VLANs are not needed when configuring Fibre Channel connectivity. Interfaces in these VLANs are configured with MTU 9000.

Logical Topology

In FlexPod Datacenter 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 service profile for an ESXi host and the storage configuration on NetApp AFF A400 controllers is described below.

Logical Topology for IP-based Storage Access

Figure 35 illustrates the end-to-end connectivity design for IP-based storage access.

Figure 35. Logical End-to-End Connectivity for iSCSI Design

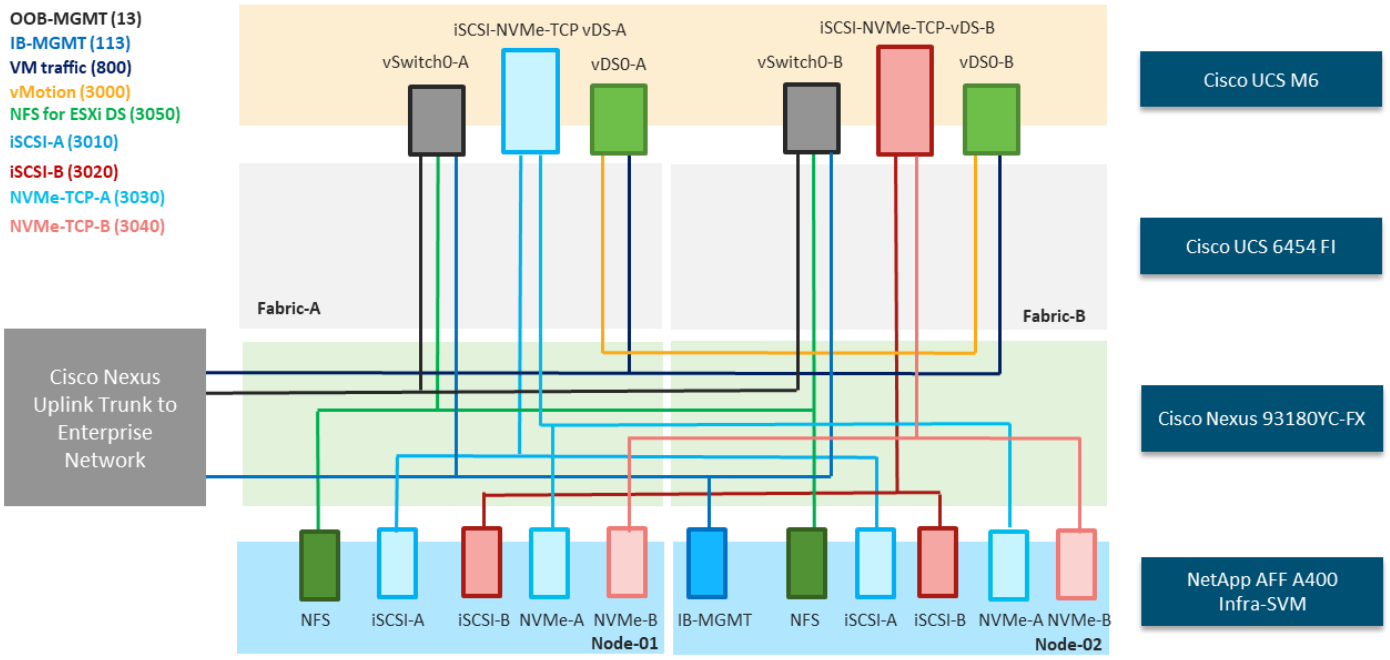

Each ESXi service profile supports:

- 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 and infrastructure NFS traffic. The MTU value for these vNICs is set as a Jumbo MTU (9000), but management interfaces with MTU 1500 can be placed on these vNICs.
  - Two redundant vNICs (vDS0-A and vDS0-B) are used by the first vSphere Distributed switch (vDS) and carry VMware vMotion traffic and customer application data traffic. The MTU for the vNICs is set to Jumbo MTU (9000), but interfaces that require MTU 1500 can be placed on these vNICs.
  - Two vNICs (iSCSi/NVMe-TCP-A and iSCSi/NVMe-TCP-B) are used by the iSCSI-NVMe-TCP vDS. The iSCSI VLANs are set as native on the corresponding vNICs, and the NVMe-TCP VLANs are set as tagged VLANs on the corresponding vNICs. The MTU value for the vNICs and all interfaces on the vDS is set to Jumbo MTU (9000). The initial VMware ESXi setup utilizes two vSwitches, but the vNICs and VMkernel ports are migrated to the second vDS.
- Each ESXi host (compute node) mounts VM datastores from NetApp AFF A400 controllers using NFS for deploying virtual machines.

Logical Topology for FC-based Storage Access

Figure 36 illustrates the end-to-end connectivity design for FC-based storage access.

Figure 36. Logical End-to-End Connectivity for FC Design

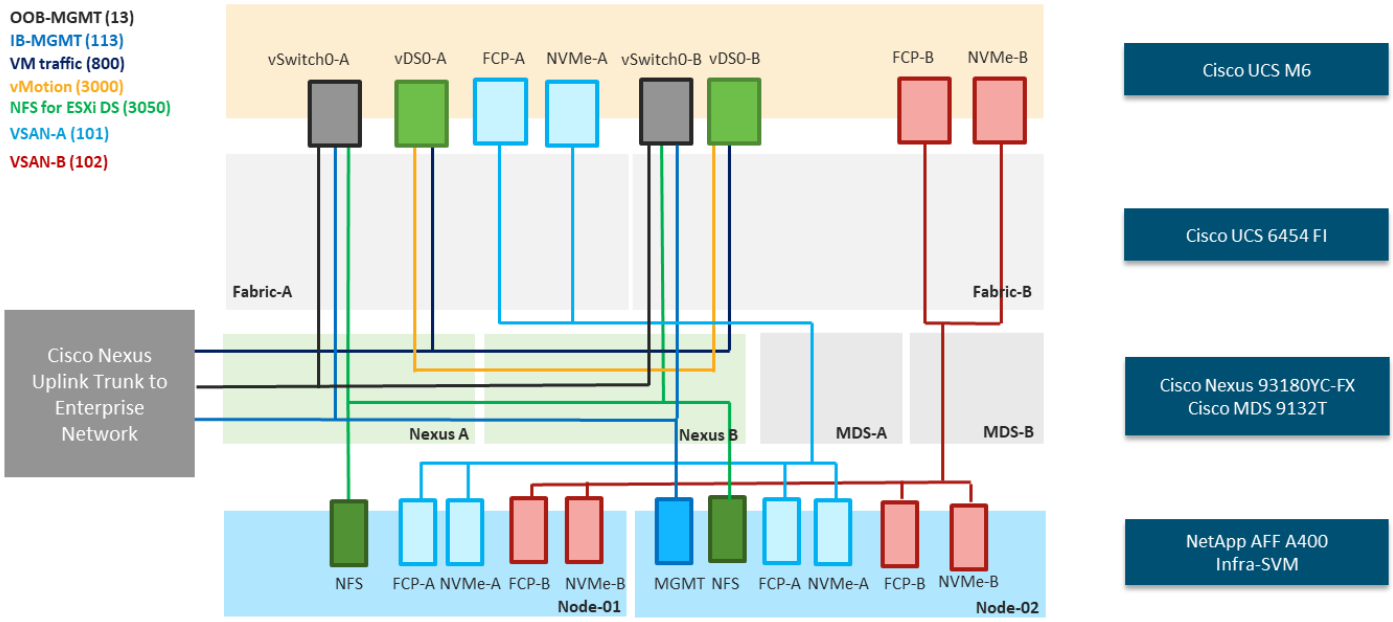

Each ESXi service profile supports:

- Managing the ESXi hosts using a common management segment.
- Diskless SAN boot using FC with persistent operating system installation for true stateless computing.
- Four vNICs where:
  - Two redundant vNICs (vSwitch0-A and vSwitch0-B) carry management and Infrastructure NFS VLANs. The MTU value for these vNICs is set as a Jumbo MTU (9000), but management interfaces with MTU 1500 can be placed on these vNICs.
  - Two redundant vNICs (vDS0-A and vDS0-B) are used by vDS0 and carry VMware vMotion traffic and customer application data traffic. The MTU for the vNICs is set to Jumbo MTU (9000), but interfaces that require MTU 1500 can be placed on these vNICs.
  - Two vHBAs (one for FC and one for FC-NVMe) defined on Fabric A to provide access to SAN-A path.
  - Two vHBAs (one for FC and one for FC-NVMe) defined on Fabric B to provide access to SAN-B path.
- Each ESXi host (compute node) mounts VM datastores from NetApp AFF A400 controllers using NFS for deploying virtual machines.

Compute System Connectivity

The Cisco UCS 5108 Chassis is equipped with the Cisco UCS 2408 Fabric Extenders or Input/Output Modules (IOMs). The Cisco UCS 5108 Chassis connects to each Cisco UCS 6454 FI using four 25GE ports, as shown in Figure 37. If the customers require more bandwidth, all eight ports on the IOMs can be connected to each FI.

Figure 37. Cisco UCS 5108 Chassis Connectivity to Cisco UCS Fabric Interconnects

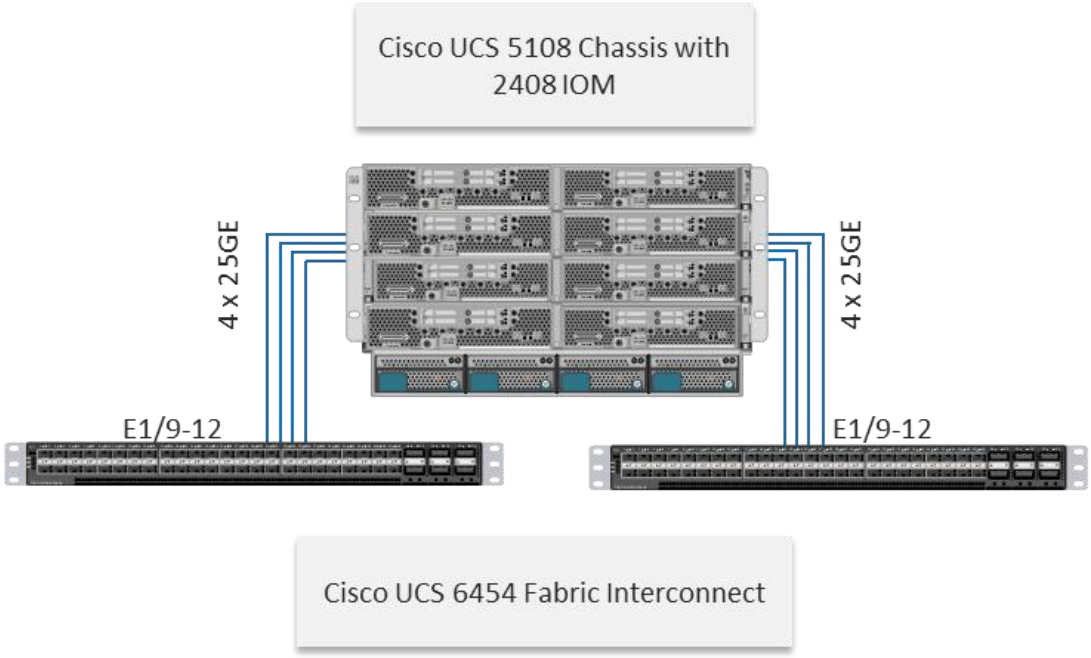

Cisco Nexus Ethernet Connectivity

The Cisco Nexus 93180YC-FX device configuration explains 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-vans—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 with port channels to Cisco Nexus 931800YC-FX switches using either 25GE (shown) or 100GE connections configured as virtual port channels. Each FI is connected to both Cisco Nexus switches using two 25G connections; additional links can easily be added to the port channel to increase the bandwidth as needed. Figure 38 illustrates the physical connectivity details.

Figure 38. Cisco UCS 6454 FI Ethernet Connectivity

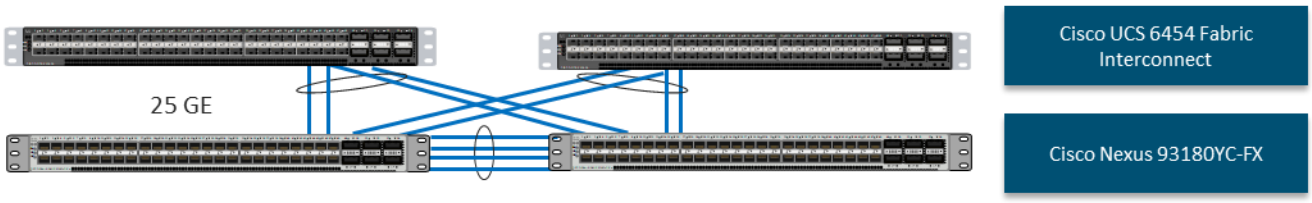

NetApp AFF A400 Ethernet Connectivity

NetApp AFF A400 controllers are connected with port channels (NetApp Interface Groups) to Cisco Nexus 93180YC-FX switches using either 25GE (shown) or 100GE connections configured as virtual port channels. The storage controllers are deployed in a switchless cluster interconnect configuration and are connected to each other using the 100GE ports e0a and e1a. Figure 26 illustrates the physical connectivity details.

In Figure 39, the two storage controllers in the high-availability pair are drawn separately for clarity. Physically, the two controllers exist within a single chassis.

Figure 39. NetApp AFF A400 Ethernet Connectivity

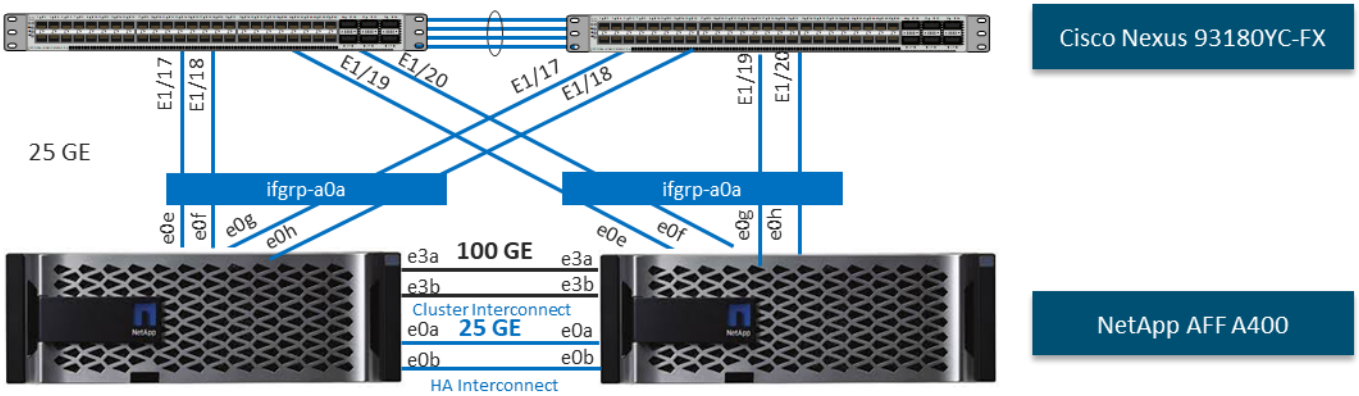

Cisco MDS SAN Connectivity – Fibre Channel Design Only

The Cisco MDS 9132T is the key design component bringing together the 32Gbps Fibre Channel (FC) capabilities to the FlexPod 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.
- Enhanced Device Alias – a feature that allows device aliases (a name for a WWPN) to be used in zones instead of WWPNs, making zones more readable. Also, if a WWPN for a vHBA or NetApp FC LIF changes, the device alias can be changed, and this change will carry over into all zones that use the device alias instead of changing WWPNs in all zones.

- Smart-Zoning—a feature that reduces the number of TCAM entries and administrative overhead by identifying the initiators and targets in the environment.

Cisco UCS Fabric Interconnect 6454 FC SAN Connectivity

For SAN connectivity, each Cisco UCS 6454 Fabric Interconnect is connected to a Cisco MDS 9132T SAN switch using at least two 32G Fibre Channel ports in a port-channel connection, as shown in Figure 40.

Figure 40. Cisco UCS 6454 FI FC SAN Connectivity

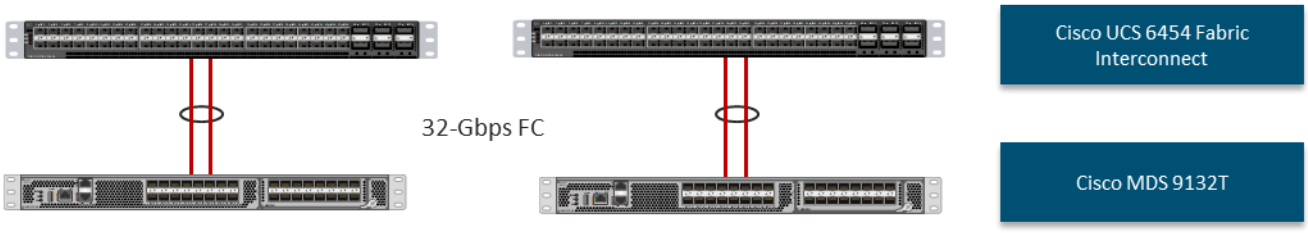

NetApp AFF A400 FC SAN Connectivity

For SAN connectivity, each NetApp AFF A400 controller is connected to both of Cisco MDS 9132T SAN switches using 32G Fibre Channel connections, as shown in Figure 41. FC-NVMe LIFs can be put on the same FC ports on the NetApp storage controllers as FC LIFs.

In Figure 41, the two storage controllers in the high-availability pair are drawn separately for clarity. Physically, the two controllers exist within a single chassis.

Figure 41. NetApp AFF A400 FC SAN Connectivity

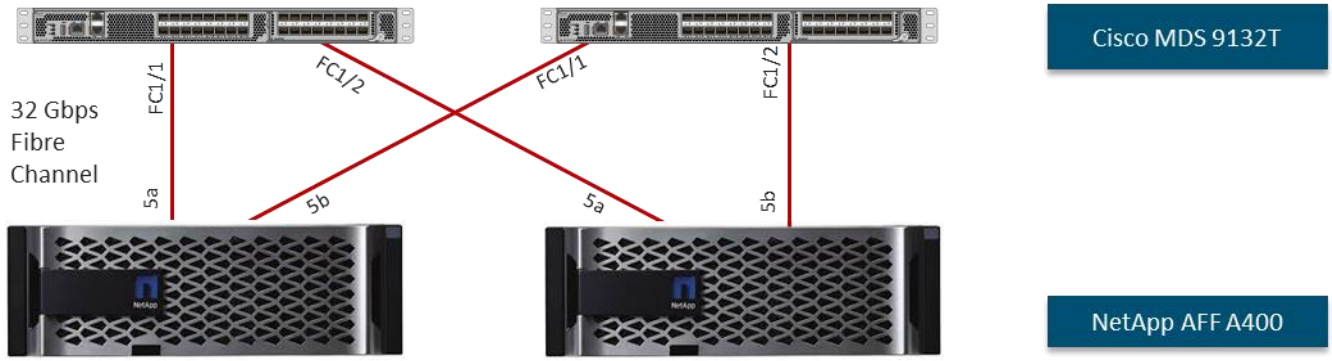

Cisco UCS Configuration – Cisco UCS Managed Mode

Cisco USC Managed Mode standardizes policy and operation management for Cisco UCS B-Series and the remaining Cisco UCS hardware used in this CVD. The Cisco UCS compute nodes are configured using service profiles defined in Cisco UCS Manager. These service profiles derive all the server characteristics from various policies and templates. At a high level, configuring Cisco UCS using Cisco UCS Manager consists of the steps shown in Figure 42.

Figure 42. Configuration Steps for Cisco UCS Managed Mode

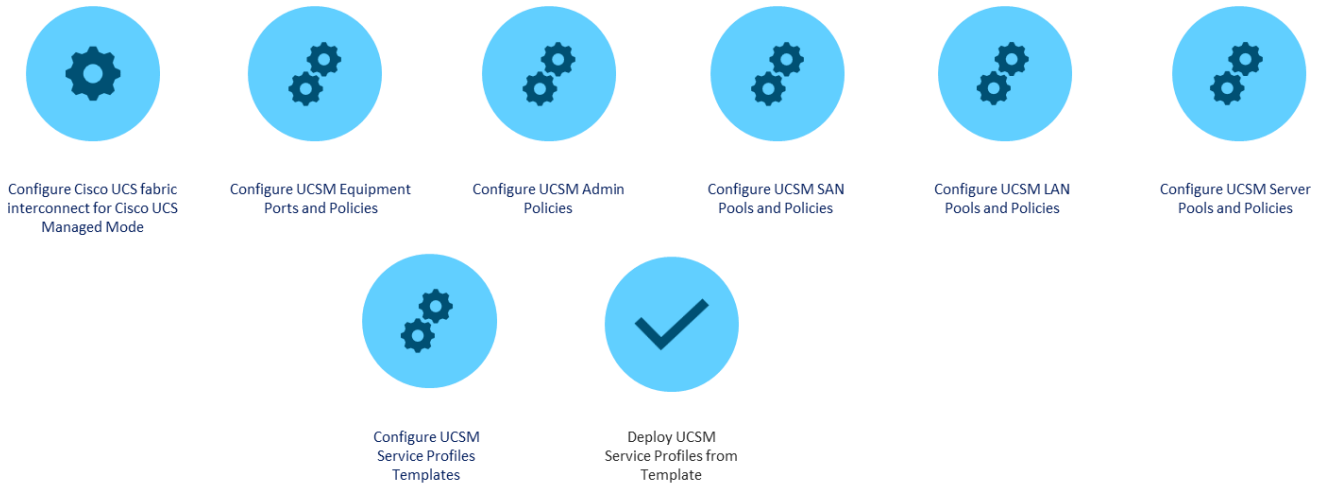

Cisco UCS Ethernet Adapter Policies

One point of optimization with Cisco UCS in FlexPod is use of Cisco UCS Ethernet Adapter policies to optimize network traffic into multiple receive (RX) queues and maximize the use of multiple CPU cores in servicing these queues resulting in higher network throughput on up to 100Gbps interfaces. Cisco Managed Mode adapter policies allow the number of transmit (TX) and RX queues and the queue ring size (buffer size) to be adjusted, and features such as Receive Side Scaling (RSS) to be enabled. RSS allows multiple RX queues to each be assigned to a different CPU core, allowing parallel processing of incoming Ethernet traffic. VMware ESXi 8.0 supports RSS, a single TX queue, and up to 16 RX queues. This CVD utilizes the fifth-generation Cisco VICs which support a ring size up to 16K (16,384), where the previous fourth-generation VICs support a ring size up to 4K (4096). Increasing the ring size can result in increased latency, but with the higher speed interfaces used in this CVD, the data moves through the buffers in less time, minimizing the latency increase. In this CVD, up to four Ethernet Adapter policies are defined in Table 2.

Table 2. Ethernet Adapter Policies

Policy Name	TX Queues	TX Ring Size	RX Queues	RX Ring Size	RSS
VMware-Default	1	256	1	512	Disabled
VMware-High Traffic	1	4096	8	4096	Enabled
VMware-4G-16RXQs	1	4096	16	4096	Enabled
VMware-5G-16RXQs	1	16384	16	16384	Enabled

Figure 43 shows part of the VMware-5G-16RXQs Ethernet Adapter policy in Cisco Intersight. Notice that not only the fields in the above table have been modified, but also Completion Queue Count (TX Queues + RX Queues) and Interrupts (Completion Queue Count + 2) have also been modified. For more information on configuring Ethernet Adapter polices, go to: https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/white-paper-c11-744754.html.

Figure 43. VMware-5G-16RxQs Ethernet Adapter Policy

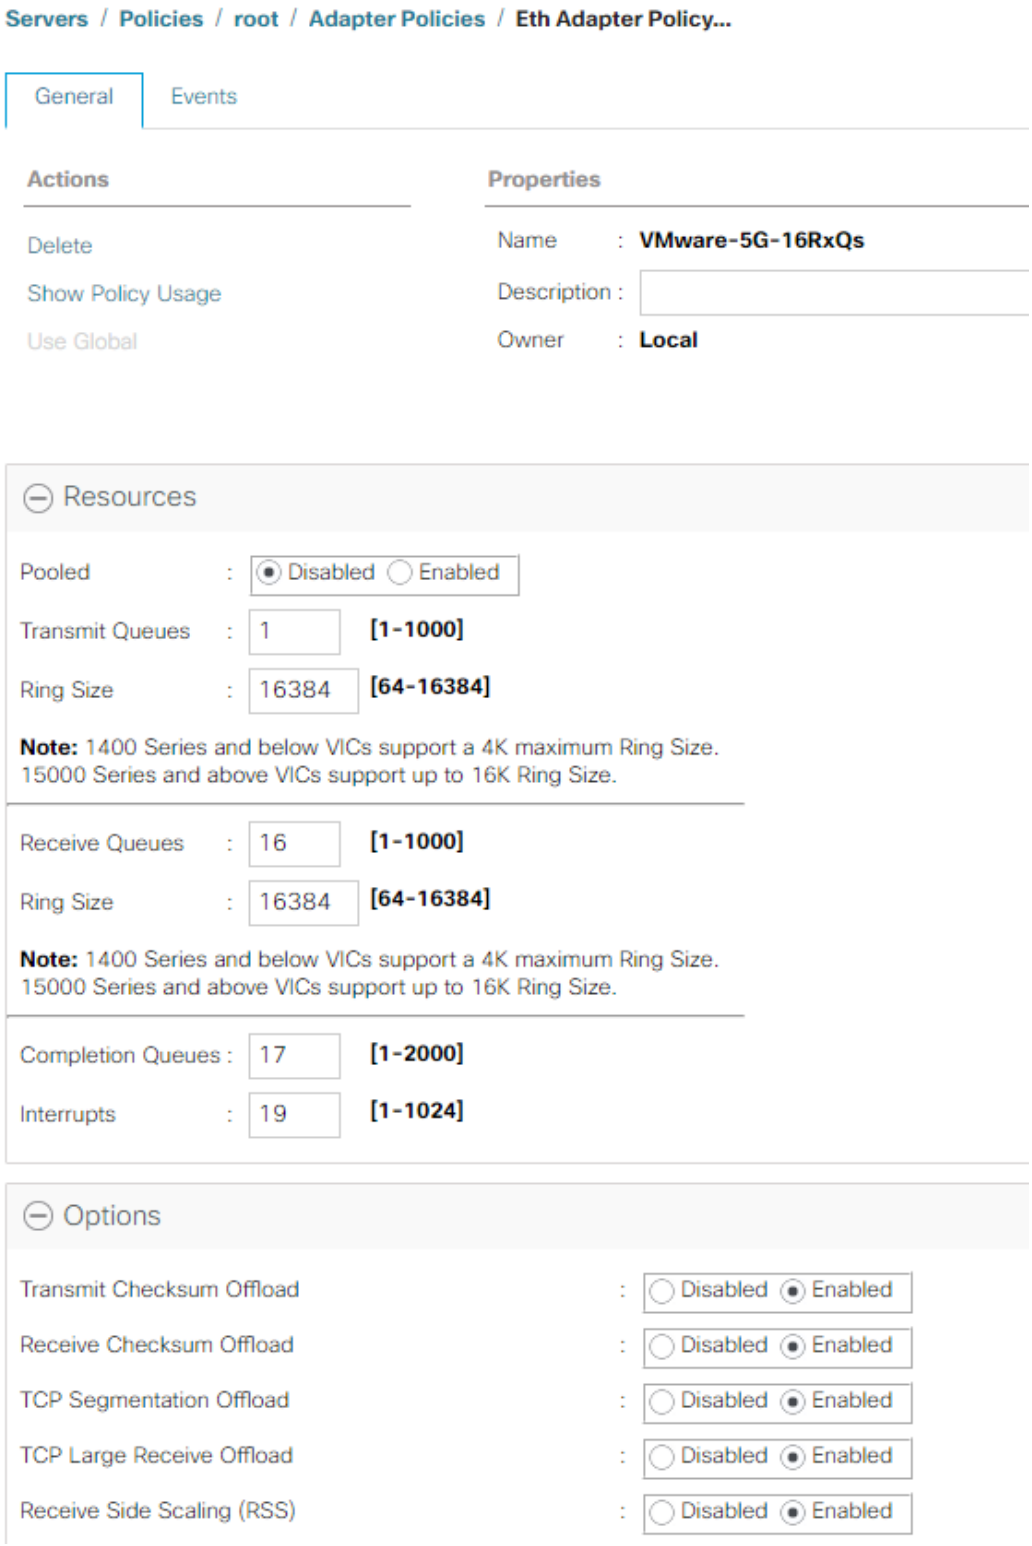

NetApp AFF – Storage Virtual Machine (SVM) Design

To provide the necessary data segregation and management, a dedicated SVM (Infra-SVM) is created for hosting the VMware environment. The SVM contains the following volumes and logical interfaces (LIFs):

- Volumes
  - ESXi boot volume (esxi_boot) that consists of ESXi boot LUNs, used to enable ESXi host boot using iSCSI or FC boot from SAN. The boot LUNs are 128GB in size and thin provisioned as per VMware recommendation.
  - Infrastructure datastores used by the vSphere environment to store the VMs and swap files. Separate datastores to be configured for NFS volume and NVMe namespace. The datastore configured for NVMe may be used for NVMe-TCP or FC-NVMe.
  - Datastore used by the vSphere environment to host vSphere Cluster Services (vCLS) VMs. By default, the datastore placement logic chooses an available datastore hence it is recommended to create a dedicated datastore for vCLS VMs.

Note: It is a NetApp best practice to create Load sharing mirror for each SVM root volume that serves NAS data in the cluster. For more information on LSM, go to: https://docs.netapp.com/us-en/ontap/data-protection/manage-snapmirror-root-volume-replication-concept.html

- Logical interfaces (LIFs)
  - NFS LIFs to mount NFS datastores in the VMware vSphere environment
  - NVMe-TCP LIFs to connect to the NVMe namespace from VMs using NVMe over TCP
  - iSCSI A/B LIFs or FC LIFs to connect to ESXi boot LUNs or application data using iSCSI and FC Protocol respectively
  - FC-NVMe LIFs for VMs to connect to NVMe datastores using NVMe over FC traffic

Each LIF belongs to specific VLANs or VSANs assigned for that traffic, as described earlier in this document. For IP based LIFs, IP addresses are assigned from subnets assigned to the respective VLAN. The IP based LIFs configured for IP SAN storage (iSCSI, , NVMe-TCP) require 2 IP addresses per controller to allow all 4 paths between the end host and storage. LIFs configured for NFS require one IP address per controller. For FC based LIFs, WWPNs are automatically assigned by storage. Both the FCP and FC-NVMe LIFs require 2 WWPNs per controller to allow all 4 paths for each protocol between the end host and storage.

A visual representation of volumes and logical interfaces (LIFs) are shown in Figure 44 and Figure 45, for iSCSI and FC boot, respectively.

Figure 44. NetApp AFF A400 - Infra-SVM for iSCSI Boot

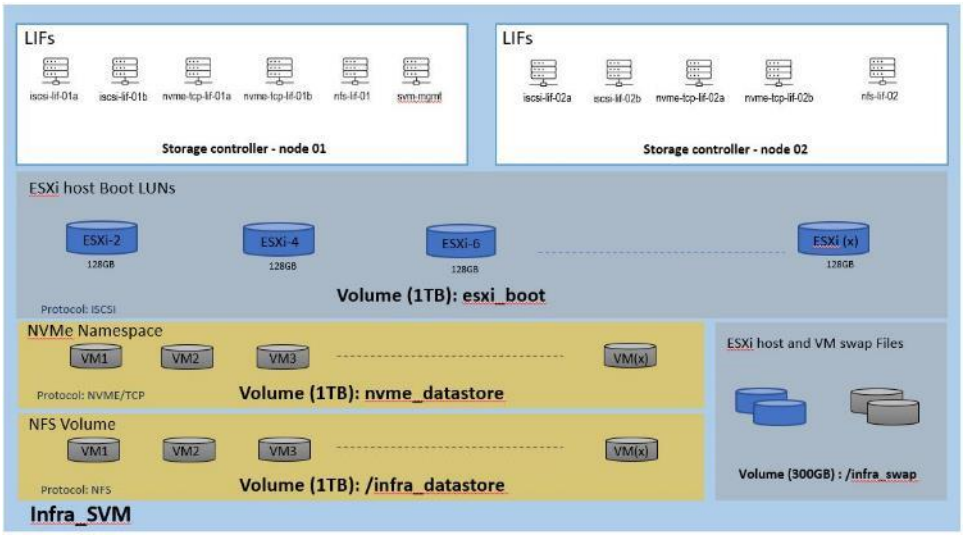

Figure 45. NetApp AFF A400 - Infra-SVM for FC Boot

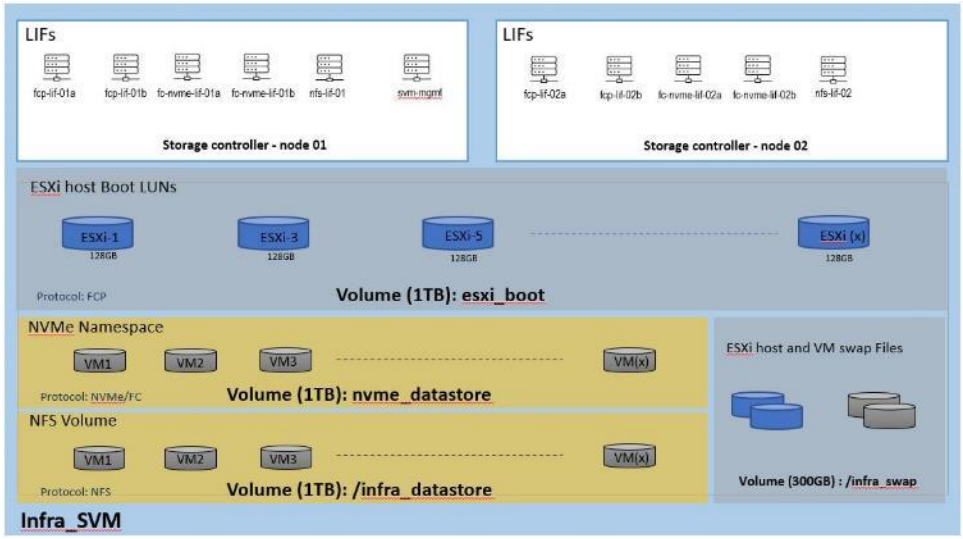

VMware vSphere - ESXi Design

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 and NFS traffic. The standard VMware-Default Cisco UCS Ethernet adapter policy is assigned to these vNICs.
- Two vNICs (one on each fabric) for vSphere Virtual Distributed Switch (vDS0) to support customer data traffic and vMotion traffic. In this vDS, vMotion is pinned to Cisco UCS Fabric B so that vMotion is switched in the B-side fabric interconnect. The higher performance VMware-HighTraffic Cisco UCS Ethernet adapter policy utilizing receive side scaling (RSS) is assigned to these vNICs. If higher performance for infrastructure

NFS is desired, the NFS VMkernel ports can be migrated to this vDS, provided the NFS VLAN is configured in the Ethernet network group policy for the vNICs on this vDS.

- Two vNICs (one on each fabric) for the iSCSI-NVMe-TCP vSphere Virtual Distributed Switch (iSCSI-NVMe-TCP-vDS) to support iSCSI (including boot) and NVMe-TCP traffic. In this vDS, both the iSCSI and NVMe-TCP VMkernel ports are pinned to the appropriate fabric. A maximum performance VMware-5G-16RXQs or VMware-4G-16RXQs Cisco UCS Ethernet adapter policy, utilizing receive side scaling (RSS) and maximum buffer size is assigned to these vNICs.

Note: Typically, you will either have iSCSI vNICs or the FC vHBAs configured for stateless boot from SAN of the ESXi servers.

Figure 46 and Figure 47 show the ESXi vNIC configurations in detail.

Figure 46. VMware vSphere - ESXi Host Networking for iSCSI Boot from SAN

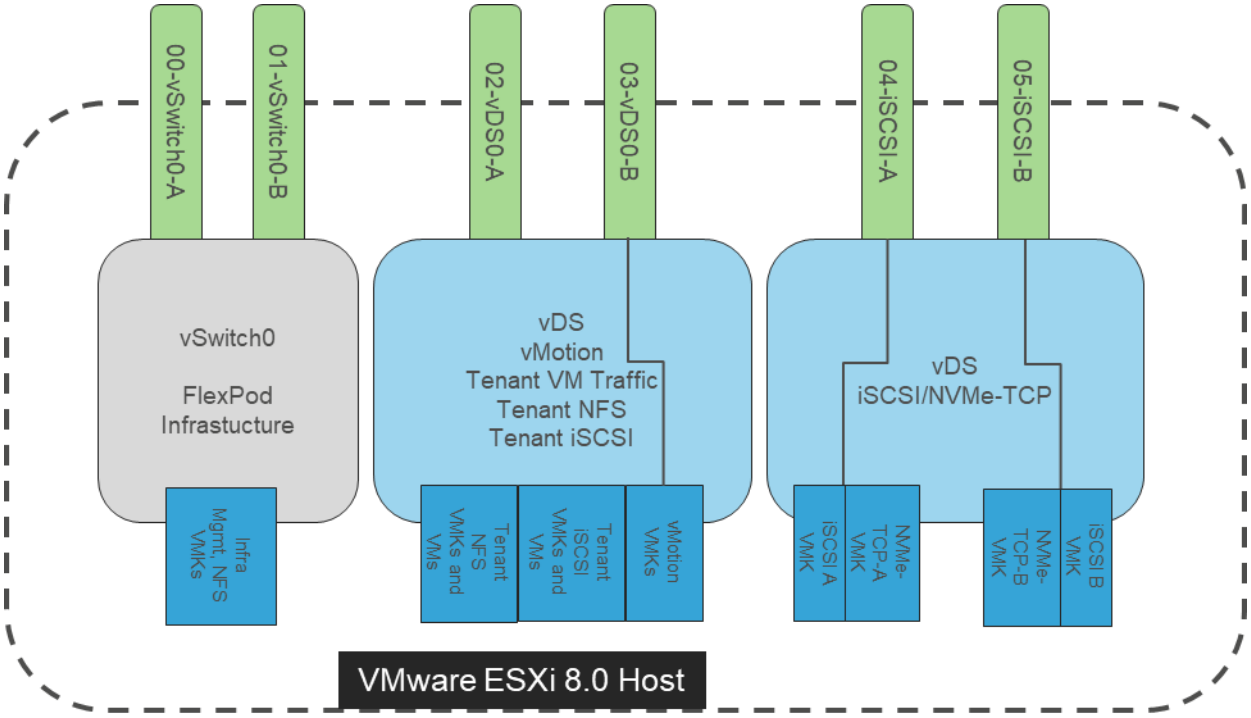

Figure 47. VMware vSphere - ESXi Host Networking for FC Boot from SAN

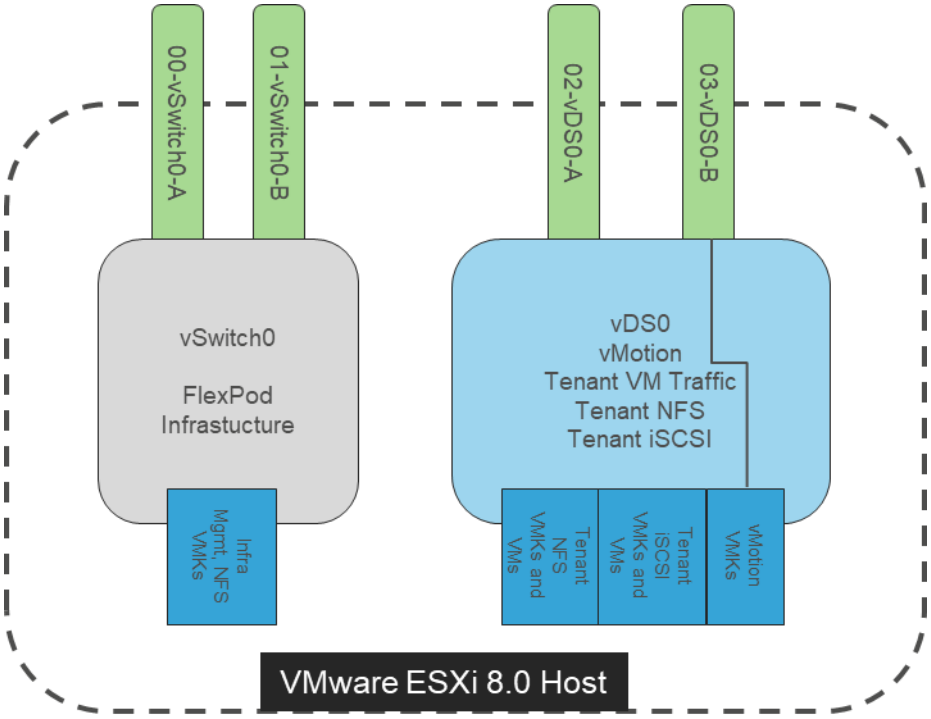

Cisco Intersight Integration with VMware vCenter, NetApp ONTAP Storage, and Cisco Switches

Cisco Intersight works with NetApp's ONTAP storage and VMware vCenter using third-party device connectors, and Cisco Nexus and MDS switches using a Cisco device connector. Since third-party infrastructure does not contain any built-in Intersight device connector, Cisco Intersight Assist virtual appliance enables Cisco Intersight to communicate with both non-Cisco devices and supported Cisco switches.

Note: A single Cisco Intersight Assist virtual appliance can support NetApp ONTAP storage, VMware vCenter, and Cisco switches.

Cisco Intersight integration with VMware vCenter, NetApp ONTAP, and Cisco switches enables customers to perform the following tasks right from the Intersight dashboard:

- Monitor the virtualization, storage, and switching 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, storage, and switching, 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.

Figure 48. Managing NetApp and VMware vCenter through Cisco Intersight using Intersight Assist

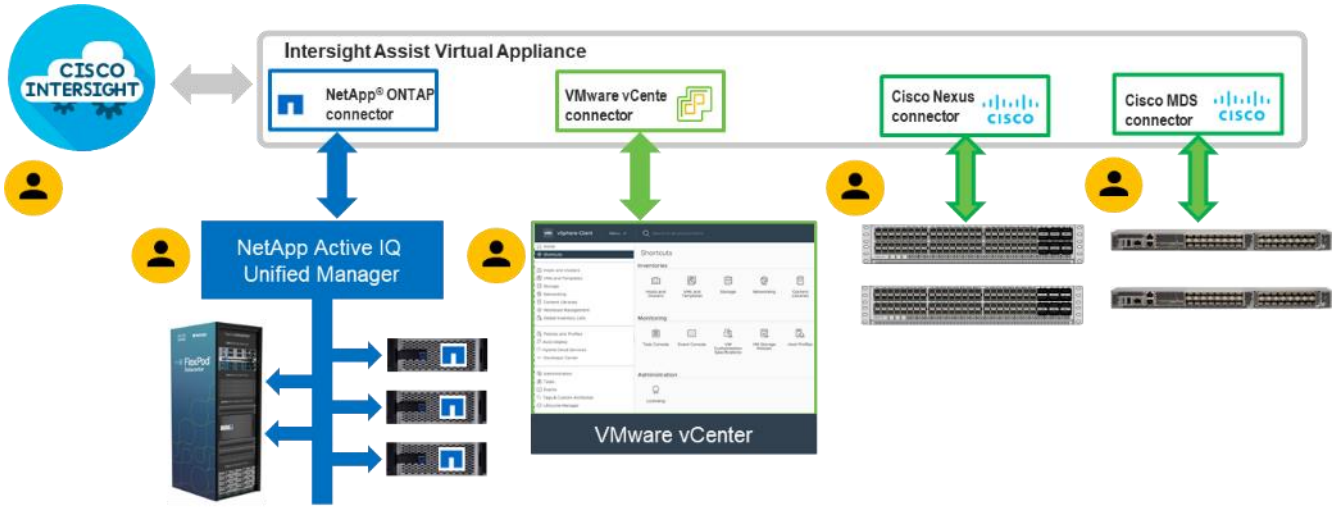

Licensing Requirement

To integrate and view various NetApp storage, VMware vCenter, and Cisco switch 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 required.

Integrate Cisco Intersight with NetApp ONTAP Storage

To integrate NetApp AFF A400 with Cisco Intersight, you need to deploy and configure:

- Cisco Intersight Assist virtual appliance
- NetApp Active IQ Unified Manager virtual appliance

Using the Cisco Intersight Assist, NetApp Active IQ Unified Manager (AIQUM) is claimed as a target in Cisco Intersight. Once NetApp AIQUM is claimed, the NetApp storage clusters configured in AIQUM will appear in Intersight and can be monitored and orchestrated.

Obtain Storage-level Information

After successfully claiming the NetApp Active IQ Unified Manager as a target, customers can view storage-level information in Cisco Intersight if they have already added NetApp AFF A400 to the NetApp Active IQ Unified Manager.

Figure 49. NetApp AFF A400 Information in Cisco Intersight

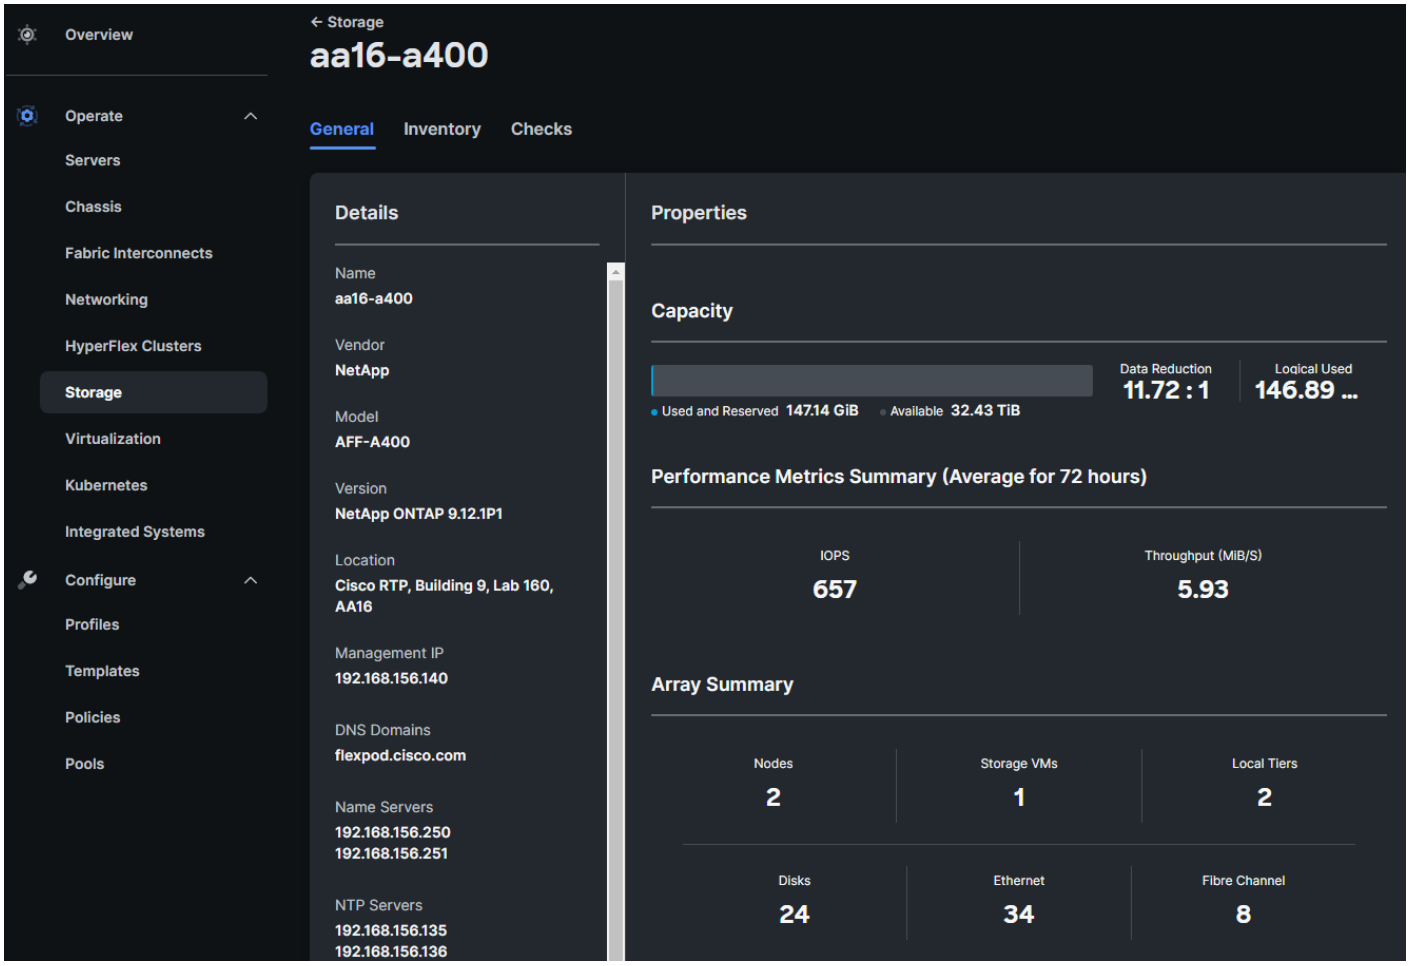

Integrate Cisco Intersight with VMware vCenter, Cisco Nexus Switches, and Cisco MDS Switches

To integrate VMware vCenter and supported Cisco switches with Cisco Intersight, you need use the deployed Cisco Intersight Assist virtual appliance. Using the Cisco Intersight Assist, VMware vCenter and supported Cisco switches are claimed as targets in Cisco Intersight.

Obtain VMware vCenter and Cisco Switch Information

After successfully claiming the VMware vCenter and supported Cisco switches as targets, you can view information on these products in Cisco Intersight.

Figure 50. VMware vCenter Information in Cisco Intersight

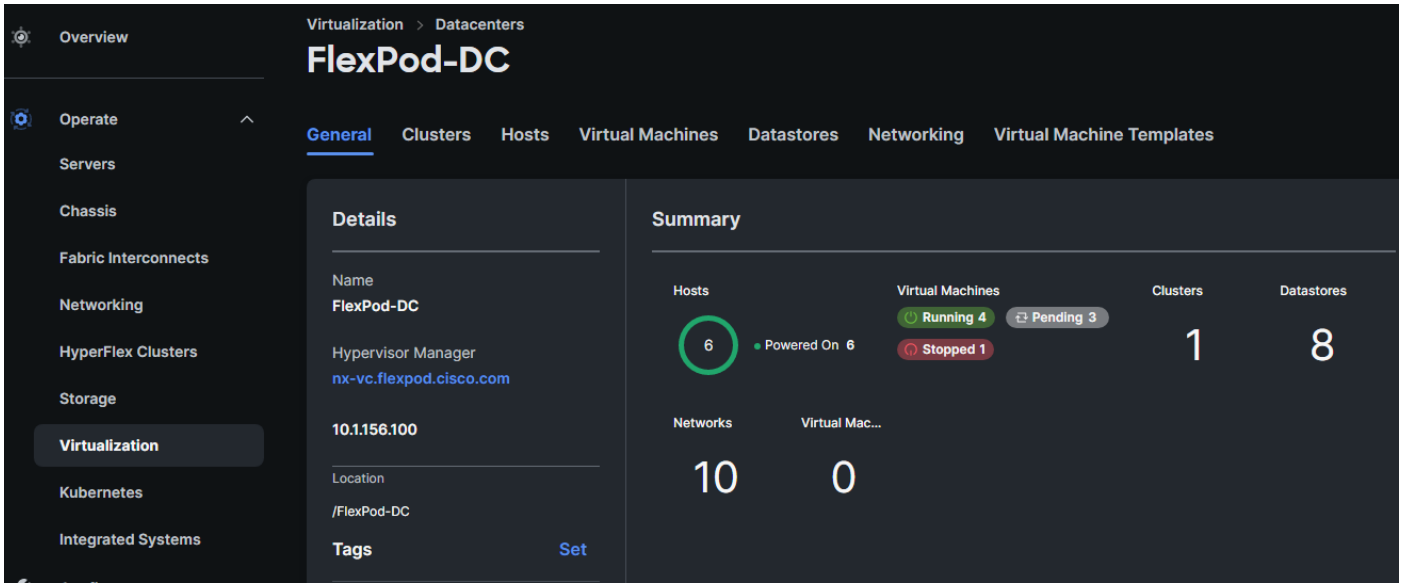

Figure 51. Cisco Nexus Information in Cisco Intersight

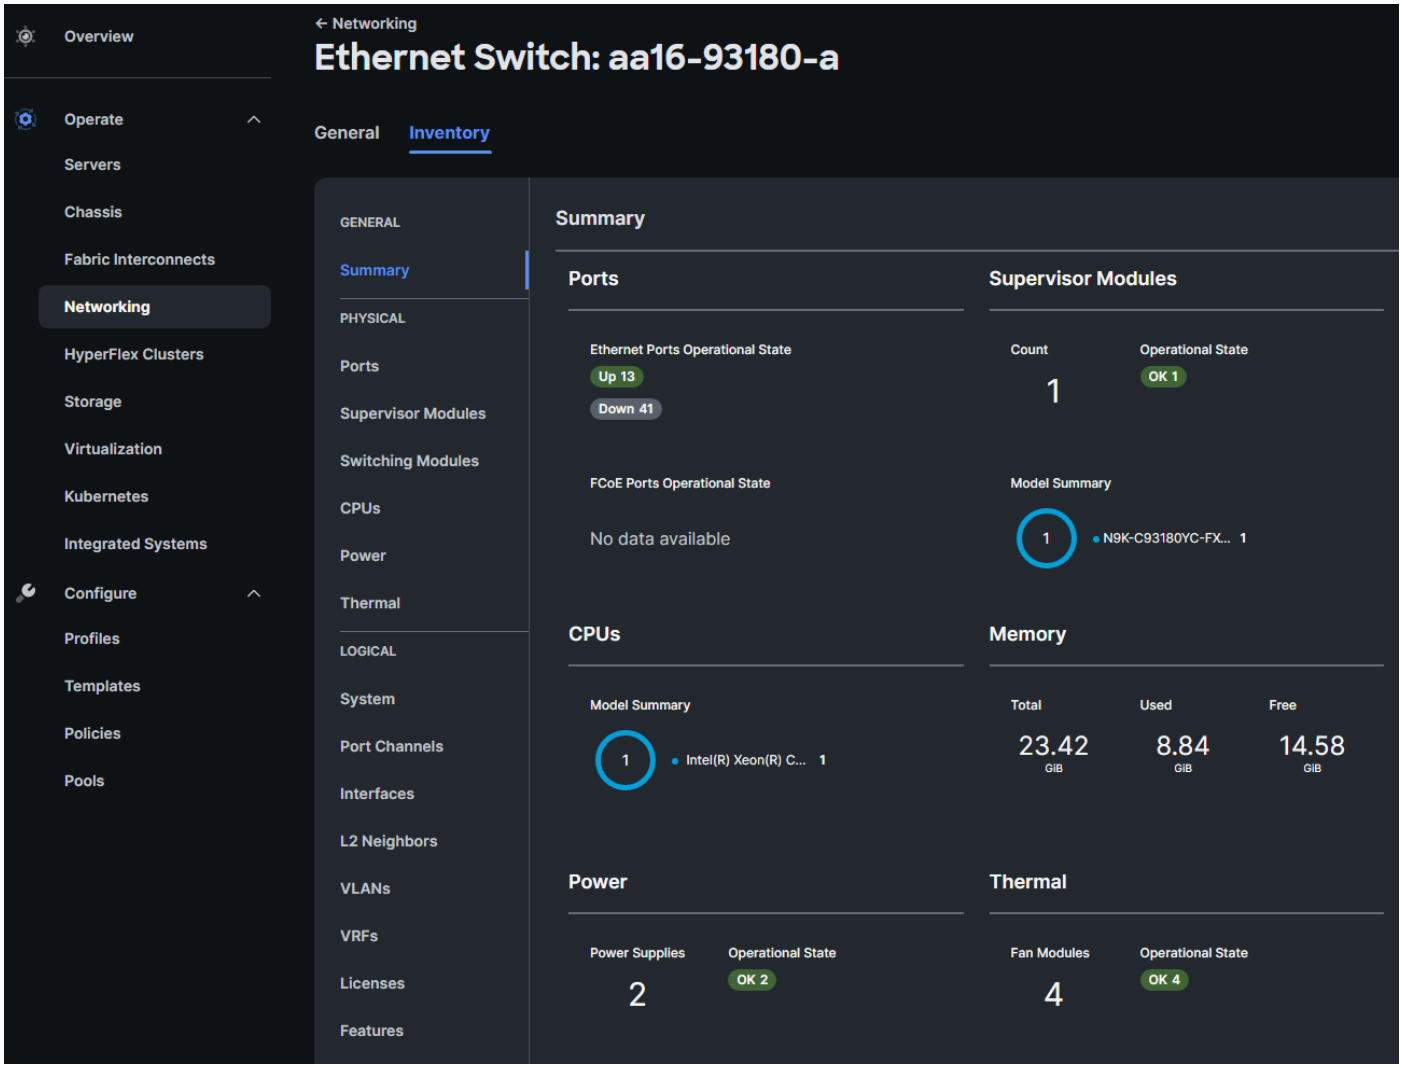

Design Considerations

Some of the key design considerations for the FlexPod Datacenter with end-to-end 100Gbps Ethernet are explained in this section.

Management Design Considerations

Out-of-band Management Network

The management interface of every physical device in FlexPod 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 the FlexPod environment for initial and on-going configuration changes. The routing and switching configuration for this network is independent of FlexPod deployment and therefore changes in FlexPod configurations do not impact management access to the devices. In this CVD, the out-of-band management network is connected to the Cisco Nexus uplinks to allow Cisco UCS CIMC connectivity and to provide the out-of-band management network to management virtual machines (Cisco DCNM) when necessary.

In-band Management Network

The in-band management VLAN configuration is part of FlexPod design. The in-band VLAN is configured on Cisco Nexus switches and Cisco UCS within the FlexPod solution to provide management connectivity for vCenter, ESXi and other management components. The changes to FlexPod configuration can impact the in-band management network and misconfigurations can cause loss of access to the management components hosted by FlexPod. It is also required that the out-of-band management network have Layer 3 access to the in-band management network so that management virtual machines with only in-band management interfaces can manage FlexPod hardware devices.

vCenter Deployment Consideration

While hosting the vCenter on the same ESXi hosts that the vCenter is managing is supported, it is a best practice to deploy the vCenter on a separate management infrastructure. Similarly, the ESXi hosts in this new FlexPod 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 FlexPod environment. In this CVD Deployment Guide, the steps for installing vCenter on FlexPod environment are included, but the vCenter can also be installed in another environment with L3 reachability to the ESXi hosts in the FlexPod.

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. This allows the network at every point to negotiate an MTU up to 9000 with the end point. For VLANs that leave the FlexPod via the Nexus switch uplinks (OOB-MGMT, IB-MGMT, VM-Traffic), all endpoints should have MTU 1500. For Storage and vMotion VLANs that stay within the FlexPod, MTU 9000 should be used on all endpoints for higher performance. It is important that all endpoints within a VLAN have the same MTU setting. It is important to remember that most virtual machine network interfaces have MTU 1500 set by default and that it may be difficult to change this setting to 9000, especially on a large number of virtual machines. This difficulty should be considered when implementing storage protocols such as CIFS or SMB. Note that a VLAN tagged trunk can contain both VLANs with MTU 1500 and VLANs with MTU 9000 interfaces.

NTP

For many reasons, including authentication and log correlation, it is critical within a FlexPod environment that all components are properly synchronized to a time-of-day clock. In order to support this synchronization, all components of FlexPod support network time protocol (NTP). In the FlexPod setup, the two Cisco Nexus switches are synchronized via NTP to at least two external NTP sources. Cisco Nexus NTP distribution is then set up and all the other components of the FlexPod can use the IP of any of the switches' L3 interfaces, including mgmt0 as an NTP source. If a customer already has NTP distribution in place, that can used instead of Cisco Nexus switch NTP distribution.

Boot From SAN

When utilizing Cisco UCS Server technology with shared storage, it is recommended to configure boot from SAN and store the boot partitions on remote storage. This enables architects and administrators to take full advantage of the stateless nature of Cisco UCS 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 FlexPod 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 as properly signed by the BIOS before they can be loaded. Additionally, a Trusted Platform Module (TPM) is also installed in the Cisco UCS compute nodes. VMware ESXi 8.0 supports UEFI Secure Boot and VMware vCenter 8.0 supports UEFI Secure Boot Attestation between the TPM module and ESXi, validating that UEFI Secure Boot has properly taken place.

NVMe over Fibre Channel

This validation of FlexPod supports NVMe over Fibre Channel (FC-NVMe) to provide the high-performance and low-latency benefits of NVMe across fabrics connecting servers and storage. FC-NVMe is implemented through the Fibre Channel over NVMe (FC-NVMe) standard which is designed to enable NVMe based message commands to transfer data and status information between a host computer and a target storage subsystem over a Fibre Channel network fabric. FC-NVMe simplifies the NVMe command sets into basic FCP instructions.

FC-NVMe requires the creation of additional Fibre Channel interfaces on Cisco UCS Compute nodes and NetApp controllers. Appropriate zoning configurations are also required on Cisco MDS switches.

NVMe over TCP

This validation of FlexPod supports NVMe over TCP (NVMe-TCP) that provides excellent performance scalability for large scale deployments and longer distances. NVMe-TCP has almost all the benefits of FC-NVMe while radically simplifying the networking requirements, including operating over routed networks. The NVMe-TCP targets are connected to the network through a standard TCP infrastructure using Ethernet switches and host-side adapters. NVMe-TCP target is supported beginning with ONTAP 9.10.1 release.

NVMe-TCP requires configuration of 2 additional LIFs per controller. Similarly, 2 additional VMkernel ports tagged for NVME-TCP are required on the ESXi hosts.

Deployment Hardware and Software

This chapter contains the following:

Hardware and Software Revisions

Table 3 lists the hardware and software used in this solution.

Table 3. Hardware and Software Revisions

Component		Software
Network	Cisco Nexus 93180YC-FX	10.2(5)M
	Cisco MDS 9132T	9.2(2)
Compute	Cisco UCS Fabric Interconnect 6536 and UCS 9108-100G IFM	4.2(3d)
	Cisco UCS B200 M6	4.2(3d)
	VMware ESXi	8.0
	Cisco VIC ENIC Driver for ESXi	1.0.45.0
	Cisco VIC FNIC Driver for ESXi	5.0.0.37
	VMware vCenter Appliance	8.0 or latest
	Cisco Intersight Assist Virtual Appliance	1.0.9-499 (automatically upgrades to current release)
Storage	NetApp AFF A400	ONTAP 9.12.1 latest P release
	NetApp Active IQ Unified Manager	9.12
	NetApp SnapCenter Plug-in for VMware vSphere	4.8
	NetApp ONTAP Tools for VMware vSphere	9.12

About the Authors

John George, Technical Marketing Engineer, Cisco Systems, Inc.

John has been involved in designing, developing, validating, and supporting the FlexPod Converged Infrastructure since it was developed over 12 years ago. Before his roles with FlexPod, he supported and administered a large worldwide training network and VPN infrastructure. John holds a master's degree in Computer Engineering from Clemson University.

Jyh-shing Chen, Senior Technical Marketing Engineer, Hybrid Cloud Infra & OEM Solutions, NetApp Inc.

Jyh-shing Chen is a Senior Technical Marketing engineer at NetApp. His current focus is on FlexPod Converged Infrastructure solution enablement, validation, deployment and management simplification, and solution integration with Cisco Intersight. Jyh-shing joined NetApp in 2006 and had worked on storage interoperability and integration projects with Solaris and VMware vSphere operating systems, and qualifications of ONTAP MetroCluster solutions and Cloud Volumes data services. Before joining NetApp, Jyh-shing's engineering experiences include software and firmware development on cardiology health imaging system, mass spectrometer system, Fibre Channel virtual tape library, and the research and development of microfluidic devices. Jyh-shing earned his BS and MS degrees from National Taiwan University, MBA degree from Meredith College, and PhD degree from MIT.

Appendix

This appendix contains following:

Compute

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

Cisco UCS 6536 Fabric Interconnects:
https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs6536-fabric-interconnect-ds.html

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

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

NetApp ONTAP: https://docs.netapp.com/ontap-9/index.jsp

NetApp Active IQ Unified Manager: https://docs.netapp.com/us-en/active-iq-unified-manager/

ONTAP Storage Connector for Cisco Intersight:
https://www.netapp.com/pdf.html?item=/media/25001-tr-4883.pdf

ONTAP tools for VMware vSphere: https://docs.netapp.com/us-en/ontap-tools-vmware-vsphere/index.html

NetApp SnapCenter: https://docs.netapp.com/us-en/snapcenter/index.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

NetApp Interoperability Matrix Tool: http://support.netapp.com/matrix/

Glossary of Acronyms

- **AAA**—Authentication, Authorization, and Accounting
- **ACP**—Access-Control Policy
- **ACI**—Cisco Application Centric Infrastructure
- **ACK**—Acknowledge or Acknowledgement
- **ACL**—Access-Control List
- **AD**—Microsoft Active Directory
- **AFI**—Address Family Identifier
- **AMP**—Cisco Advanced Malware Protection
- **AP**—Access Point
- **API**—Application Programming Interface
- **APIC**— Cisco Application Policy Infrastructure Controller (ACI)
- **ASA**—Cisco Adaptative Security Appliance
- **ASM**—Any-Source Multicast (PIM)
- **ASR**—Aggregation Services Router
- **Auto-RP**—Cisco Automatic Rendezvous Point protocol (multicast)
- **AVC**—Application Visibility and Control
- **BFD**—Bidirectional Forwarding Detection
- **BGP**—Border Gateway Protocol
- **BMS**—Building Management System
- **BSR**—Bootstrap Router (multicast)
- **BYOD**—Bring Your Own Device

- **CAPWAP**—Control and Provisioning of Wireless Access Points Protocol
- **CDP**—Cisco Discovery Protocol
- **CEF**—Cisco Express Forwarding
- **CMD**—Cisco Meta Data
- **CPU**—Central Processing Unit
- **CSR**—Cloud Services Routers
- **CTA**—Cognitive Threat Analytics
- **CUWN**—Cisco Unified Wireless Network
- **CVD**—Cisco Validated Design
- **CYOD**—Choose Your Own Device
- **DC**—Data Center
- **DHCP**—Dynamic Host Configuration Protocol
- **DM**—Dense-Mode (multicast)
- **DMVPN**—Dynamic Multipoint Virtual Private Network
- **DMZ**—Demilitarized Zone (firewall/networking construct)
- **DNA**—Cisco Digital Network Architecture
- **DNS**—Domain Name System
- **DORA**—Discover, Offer, Request, ACK (DHCP Process)
- **DWDM**—Dense Wavelength Division Multiplexing
- **ECMP**—Equal Cost Multi Path
- **EID**—Endpoint Identifier
- **EIGRP**—Enhanced Interior Gateway Routing Protocol
- **EMI**—Electromagnetic Interference
- **ETR**—Egress Tunnel Router (LISP)
- **EVPN**—Ethernet Virtual Private Network (BGP EVPN with VXLAN data plane)
- **FHR**—First-Hop Router (multicast)
- **FHRP**—First-Hop Redundancy Protocol
- **FI**—Fabric Interconnect
- **FMC**—Cisco Firepower Management Center
- **FTD**—Cisco Firepower Threat Defense
- **GBAC**—Group-Based Access Control

- **GbE**—Gigabit Ethernet
- **Gbit/s**—Gigabits Per Second (interface/port speed reference)
- **GRE**—Generic Routing Encapsulation
- **GRT**—Global Routing Table
- **HA**—High-Availability
- **HQ**—Headquarters
- **HSRP**—Cisco Hot-Standby Routing Protocol
- **HTDB**—Host-tracking Database (SD-Access control plane node construct)
- **IBNS**—Identity-Based Networking Services (IBNS 2.0 is the current version)
- **ICMP**— Internet Control Message Protocol
- **IDF**—Intermediate Distribution Frame; essentially a wiring closet.
- **IEEE**—Institute of Electrical and Electronics Engineers
- **IETF**—Internet Engineering Task Force
- **IGP**—Interior Gateway Protocol
- **IID**—Instance-ID (LISP)
- **IOE**—Internet of Everything
- **IoT**—Internet of Things
- **IP**—Internet Protocol
- **IPAM**—IP Address Management
- **IPS**—Intrusion Prevention System
- **IPSec**—Internet Protocol Security
- **ISE**—Cisco Identity Services Engine
- **ISR**—Integrated Services Router
- **IS-IS**—Intermediate System to Intermediate System routing protocol
- **ITR**—Ingress Tunnel Router (LISP)
- **LACP**—Link Aggregation Control Protocol
- **LAG**—Link Aggregation Group
- **LAN**—Local Area Network
- **L2 VNI**—Layer 2 Virtual Network Identifier; as used in SD-Access Fabric, a VLAN.
- **L3 VNI**— Layer 3 Virtual Network Identifier; as used in SD-Access Fabric, a VRF.
- **LHR**—Last-Hop Router (multicast)

- **LISP**—Location Identifier Separation Protocol
- **MAC**—Media Access Control Address (OSI Layer 2 Address)
- **MAN**—Metro Area Network
- **MEC**—Multichassis EtherChannel, sometimes referenced as ***MCEC***
- **MDF**—Main Distribution Frame; essentially the central wiring point of the network.
- **MnT**—Monitoring and Troubleshooting Node (Cisco ISE persona)
- **MOH**—Music on Hold
- **MPLS**—Multiprotocol Label Switching
- **MR**—Map-resolver (LISP)
- **MS**—Map-server (LISP)
- **MSDP**—Multicast Source Discovery Protocol (multicast)
- **MTU**—Maximum Transmission Unit
- **NAC**—Network Access Control
- **NAD**—Network Access Device
- **NAT**—Network Address Translation
- **NBAR**—Cisco Network-Based Application Recognition (NBAR2 is the current version).
- **NFV**—Network Functions Virtualization
- **NSF**—Non-Stop Forwarding
- **OSI**—Open Systems Interconnection model
- **OSPF**—Open Shortest Path First routing protocol
- **OT**—Operational Technology
- **PAgP**—Port Aggregation Protocol
- **PAN**—Primary Administration Node (Cisco ISE persona)
- **PCI DSS**—Payment Card Industry Data Security Standard
- **PD**—Powered Devices (PoE)
- **PETR**—Proxy-Egress Tunnel Router (LISP)
- **PIM**—Protocol-Independent Multicast
- **PITR**—Proxy-Ingress Tunnel Router (LISP)
- **PnP**—Plug-n-Play
- **PoE**—Power over Ethernet (Generic term, may also refer to IEEE 802.3af, 15.4W at PSE)
- **PoE+**—Power over Ethernet Plus (IEEE 802.3at, 30W at PSE)

- **PSE**—Power Sourcing Equipment (PoE)
- **PSN**—Policy Service Node (Cisco ISE persona)
- **pxGrid**—Platform Exchange Grid (Cisco ISE persona and publisher/subscriber service)
- **PxTR**—Proxy-Tunnel Router (LISP – device operating as both a PETR and PITR)
- **QoS**—Quality of Service
- **RADIUS**—Remote Authentication Dial-In User Service
- **REST**—Representational State Transfer
- **RFC**—Request for Comments Document (IETF)
- **RIB**—Routing Information Base
- **RLOC**—Routing Locator (LISP)
- **RP**—Rendezvous Point (multicast)
- **RP**—Redundancy Port (WLC)
- **RP**—Route Processer
- **RPF**—Reverse Path Forwarding
- **RR**—Route Reflector (BGP)
- **RTT**—Round-Trip Time
- **SA**—Source Active (multicast)
- **SAFI**—Subsequent Address Family Identifiers (BGP)
- **SD**—Software-Defined
- **SDA**—Cisco Software Defined-Access
- **SDN**—Software-Defined Networking
- **SFP**—Small Form-Factor Pluggable (1 GbE transceiver)
- **SFP+**— Small Form-Factor Pluggable (10 GbE transceiver)
- **SGACL**—Security-Group ACL
- **SGT–**Scalable Group Tag, sometimes reference as Security Group Tag
- **SM**—Spare-mode (multicast)
- **SNMP**—Simple Network Management Protocol
- **SSID**—Service Set Identifier (wireless)
- **SSM**—Source-Specific Multicast (PIM)
- **SSO**—Stateful Switchover
- **STP**—Spanning-tree protocol

- **SVI**—Switched Virtual Interface
- **SVL**—Cisco StackWise Virtual
- **SWIM**—Software Image Management
- **SXP**—Scalable Group Tag Exchange Protocol
- **Syslog**—System Logging Protocol
- **TACACS+**—Terminal Access Controller Access-Control System Plus
- **TCP**—Transmission Control Protocol (OSI Layer 4)
- **UCS**— Cisco Unified Computing System
- **UDP**—User Datagram Protocol (OSI Layer 4)
- **UPoE**—Cisco Universal Power Over Ethernet (60W at PSE)
- **UPoE+**— Cisco Universal Power Over Ethernet Plus (90W at PSE)
- **URL**—Uniform Resource Locator
- **VLAN**—Virtual Local Area Network
- **VM**—Virtual Machine
- **VN**—Virtual Network, analogous to a VRF in SD-Access
- **VNI**—Virtual Network Identifier (VXLAN)
- **vPC**—virtual Port Channel (Cisco Nexus)
- **VPLS**—Virtual Private LAN Service
- **VPN**—Virtual Private Network
- **VPNv4**—BGP address family that consists of a Route-Distinguisher (RD) prepended to an IPv4 prefix
- **VPWS**—Virtual Private Wire Service
- **VRF**—Virtual Routing and Forwarding
- **VSL**—Virtual Switch Link (Cisco VSS component)
- **VSS**—Cisco Virtual Switching System
- **VXLAN**—Virtual Extensible LAN
- **WAN**—Wide-Area Network
- **WLAN**—Wireless Local Area Network (generally synonymous with IEEE 802.11-based networks)
- **WoL**—Wake-on-LAN
- **xTR**—Tunnel Router (LISP - device operating as both an ETR and ITR)

Glossary of Terms

This glossary addresses some terms used in this document, for the purposes of aiding understanding. This is not a complete list of all multicloud terminology. Some Cisco product links are supplied here also, where considered

useful for the purposes of clarity, but this is by no means intended to be a complete list of all applicable Cisco products.

aaS/XaaS **(IT capability provided as a Service)**	Some IT capability, X, provided as a service (XaaS). Some benefits are: • The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it. • There are very low barriers to entry, so that services can be quickly adopted and dropped in response to business demand, without the penalty of inefficiently utilized CapEx. • The service charge is an IT OpEx cost (pay-as-you-go), whereas the CapEx and the service infrastructure is the responsibility of the provider. • Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes. Such services are typically implemented as "microservices," which are accessed via REST APIs. This architectural style supports composition of service components into systems. Access to and management of aaS assets is via a web GUI and/or APIs, such that Infrastructure-as-code (IaC) techniques can be used for automation, for example, Ansible and Terraform. The provider can be any entity capable of implementing an aaS "cloud-native" architecture. The cloud-native architecture concept is well-documented and supported by open-source software and a rich ecosystem of services such as training and consultancy. The provider can be an internal IT department or any of many third-party companies using and supporting the same open-source platforms. Service access control, integrated with corporate IAM, can be mapped to specific users and business activities, enabling consistent policy controls across services, wherever they are delivered from.
Ansible	An infrastructure automation tool, used to implement processes for instantiating and configuring IT service components, such as VMs on an IaaS platform. Supports the consistent execution of processes defined in YAML "playbooks" at scale, across multiple targets. Because the Ansible artefacts (playbooks) are text-based, they can be stored in a Source Code Management (SCM) system, such as GitHub. This allows for software development like processes to be applied to infrastructure automation, such as, Infrastructure-as-code (see IaC below). https://www.ansible.com
AWS **(Amazon Web Services)**	Provider of IaaS and PaaS. https://aws.amazon.com
Azure	Microsoft IaaS and PaaS. https://azure.microsoft.com/en-gb/
Co-located data center	"A colocation center (CoLo)...is a type of data center where equipment, space, and bandwidth are available for rental to retail customers. Colocation facilities provide space, power, cooling, and physical security for the server, storage, and networking equipment of other firms and also connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity." https://en.wikipedia.org/wiki/Colocation_centre

Containers **(Docker)**	A (Docker) container is a means to create a package of code for an application and its dependencies, such that the application can run on different platforms which support the Docker environment. In the context of aaS, microservices are typically packaged within Linux containers orchestrated by Kubernetes (K8s). https://www.docker.com https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html
DevOps	The underlying principle of DevOps is that the application development and operations teams should work closely together, ideally within the context of a toolchain that automates the stages of development, test, deployment, monitoring, and issue handling. DevOps is closely aligned with IaC, continuous integration and deployment (CI/CD), and Agile software development practices. https://en.wikipedia.org/wiki/DevOps https://en.wikipedia.org/wiki/CI/CD
Edge compute	Edge compute is the idea that it can be more efficient to process data at the edge of a network, close to the endpoints that originate that data, or to provide virtualized access services, such as at the network edge. This could be for reasons related to low latency response, reduction of the amount of unprocessed data being transported, efficiency of resource utilization, and so on. The generic label for this is Multi-access Edge Computing (MEC), or Mobile Edge Computing for mobile networks specifically. From an application experience perspective, it is important to be able to utilize, at the edge, the same operations model, processes, and tools used for any other compute node in the system. https://en.wikipedia.org/wiki/Mobile_edge_computing
IaaS **(Infrastructure as-a-Service)**	Infrastructure components provided aaS, located in data centers operated by a provider, typically accessed over the public Internet. IaaS provides a base platform for the deployment of workloads, typically with containers and Kubernetes (K8s).
IaC **(Infrastructure as-Code)**	Given the ability to automate aaS via APIs, the implementation of the automation is typically via Python code, Ansible playbooks, and similar. These automation artefacts are programming code that define how the services are consumed. As such, they can be subject to the same code management and software development regimes as any other body of code. This means that infrastructure automation can be subject to all of the quality and consistency benefits, CI/CD, traceability, automated testing, compliance checking, and so on, that could be applied to any coding project. https://en.wikipedia.org/wiki/Infrastructure_as_code
IAM **(Identity and Access Management)**	IAM is the means to control access to IT resources so that only those explicitly authorized to access given resources can do so. IAM is an essential foundation to a secure multicloud environment. https://en.wikipedia.org/wiki/Identity_management_
IBM **(Cloud)**	IBM IaaS and PaaS. https://www.ibm.com/cloud

Intersight	Cisco Intersight is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support. https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html
GCP **(Google Cloud Platform)**	Google IaaS and PaaS. https://cloud.google.com/gcp
Kubernetes **(K8s)**	Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. https://kubernetes.io
Microservices	A microservices architecture is characterized by processes implementing fine-grained services, typically exposed via REST APIs and which can be composed into systems. The processes are often container-based, and the instantiation of the services often managed with Kubernetes. Microservices managed in this way are intrinsically well suited for deployment into IaaS environments, and as such, are the basis of a cloud native architecture. https://en.wikipedia.org/wiki/Microservices
PaaS **(Platform-as-a-Service)**	PaaS is a layer of value-add services, typically for application development, deployment, monitoring, and general lifecycle management. The use of IaC with IaaS and PaaS is very closely associated with DevOps practices.
Private on-premises data center	A data center infrastructure housed within an environment owned by a given enterprise is distinguished from other forms of data center, with the implication that the private data center is more secure, given that access is restricted to those authorized by the enterprise. Thus, circumstances can arise where very sensitive IT assets are only deployed in a private data center, in contrast to using public IaaS. For many intents and purposes, the underlying technology can be identical, allowing for hybrid deployments where some IT assets are privately deployed but also accessible to other assets in public IaaS. IAM, VPNs, firewalls, and similar are key technologies needed to underpin the security of such an arrangement.
REST API	Representational State Transfer (REST) APIs is a generic term for APIs accessed over HTTP(S), typically transporting data encoded in JSON or XML. REST APIs have the advantage that they support distributed systems, communicating over HTTP, which is a well-understood protocol from a security management perspective. REST APIs are another element of a cloud-native applications architecture, alongside microservices. https://en.wikipedia.org/wiki/Representational_state_transfer
SaaS **(Software-as-a-Service)**	End-user applications provided “aaS” over the public Internet, with the underlying software systems and infrastructure owned and managed by the provider.
SAML **(Security Assertion**	Used in the context of Single-Sign-On (SSO) for exchanging authentication and authorization data between an identity provider, typically an IAM system, and a service provider (some form of SaaS). The SAML protocol exchanges XML documents that contain security assertions used by

Markup Language)	the aaS for access control decisions. https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language
Terraform	An open-source IaC software tool for cloud services, based on declarative configuration files. https://www.terraform.io

Feedback

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

CVD Program

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**.