FlexPod as a Workload Domain for VMware Cloud Foundation

Design Guide

Published: December 2022

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 covers the design details of incorporating FlexPod Datacenter as a workload domain for VMware Cloud Foundation. VMware Cloud Foundation provides a complete set of software defined services to run enterprise apps, both traditional and containerized, in private or public cloud environments. VMware Cloud Foundation simplifies the private cloud deployment and provides a streamlined path to the hybrid cloud by delivering a single integrated solution that is easy to deploy, operate and manage.

VMware Cloud Foundation (VCF) provides following benefits in a data center environment:

- **Integrated Stack**: VCF is an engineered solution that integrates the entire VMware software-defined stack with guaranteed interoperability.
- **Standardized Architecture**: VCF is built upon standard VMware Validated Design architecture and therefore ensures quick, repeatable deployments while eliminating risk of misconfigurations.
- **Lifecycle Management**: VCF includes lifecycle management services that automate day 0 to day 2 operations, resources provisioning and patching/upgrades.

Some of the key advantages of integrating Cisco FlexPod Datacenter as a workload domain for VMware Cloud Foundation are:

- **Simpler and programmable infrastructure:** FlexPod infrastructure delivered as infrastructure-as-a-code through a single partner integrable open API.
- **Latest hardware and software compute innovations:** policy-based configurations, delivered using Cisco Intersight, to deploy and manage the latest processor, memory, network, and power/cooling improvements.
- **Storage Modernization**: deliver high-speed, consistent, low latency, multi-tenant storage using a range of NetApp all-flash arrays.
- **Innovative cloud operations:** continuous feature delivery and no need for maintaining on-premises virtual machines supporting management functions.
- **Built for investment protections:** design ready for future technologies such as liquid cooling and high-Wattage CPUs; CXL-ready.

The FlexPod workload domain includes integration of the Cisco Intersight with VMware vCenter and NetApp Active IQ Unified Manager to deliver 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 Intersight Workload Optimization and Intersight Cloud Orchestrator.

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:

VMware Cloud Foundation enables data center administrators to provision an application environment in a quick, repeatable, and automated manner. VMware Cloud Foundation consists of workload domains which represent an application-ready infrastructure. A workload domain represents a logical unit that groups ESXi hosts managed by a vCenter Server instance with specific characteristics according to VMware best practices.

To deploy and manage the workload domains, VMware Cloud Foundation introduces VMware Cloud Builder and VMware Cloud Foundation Software Defined Data Center (SDDC) Manager. VMware Cloud Builder automates the deployment of the software defined stack, creating the first software defined unit known as the management domain. After the management domain is successfully setup, using the newly deployed SDDC Manager, virtual infrastructure administrator or cloud administrator provisions FlexPod Datacenter as a new workload domain to manage life cycle and other operational activities.

Workload domain definition requires administrators to configure network, compute and storage as well as install VMware vSphere ESXi software on the hosts that become part of workload domains (including the management domain). To automate the infrastructure setup, Cisco Intersight (or Cisco UCS Manager for Non-UCS-X-Series systems), NetApp ONTAP and Cisco NxOS configurations are (optionally) automated using RedHat Ansible framework for an easy on-boarding experience.

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 FlexPod Datacenter as a workload domain for VMware Cloud Foundation single site deployment. The document introduces various design elements and covers various considerations and best practices for a successful deployment. While VMware Cloud Foundation can be utilized in public cloud such as VMware Cloud on AWS as well as hybrid cloud solutions, the discussion in this document focuses solely on the on-prem data center design and requirements.

What's New in this Release?

Following design elements distinguish this FlexPod Datacenter Cisco Validated Design from previous designs:

- VMware Cloud Foundation deployment on vSAN ready nodes.
- Integration of FlexPod Datacenter as a workload domain in VMware Cloud Foundation.
- Automated configuration of the ESXi hosts for both the VMware Cloud Foundation management and workload domains using Cisco Intersight.
- Integration of compute and storage infrastructure management tools to administer the FlexPod workload domain.

- Integration of the FlexPod workload domain virtual environment (VMware vCenter) and NetApp storage (Active IQ Unified Manager) with Cisco Intersight for visibility and orchestration support.

Solution Summary

The FlexPod as a workload domain for VMware Cloud Foundation solution offers the following key customer benefits:

- Integrated solution that supports entire VMware software defined stack
- Standardized architecture for quick, repeatable, error free deployments of FlexPod based workload domains
- Automated life cycle management to keep all the system components up to date
- Simplified cloud-based management of various FlexPod components
- Hybrid-cloud-ready, policy-driven modular design
- Highly available, flexible, and scalable FlexPod architecture
- Cooperative support model and Cisco Solution Support
- Easy to deploy, consume, and manage design which aligns with Cisco, NetApp and VMware best practices and compatibility requirements
- Support for component monitoring, solution automation and orchestration, and workload optimization

Like all other FlexPod solution designs, FlexPod as a workload domain for VMware Cloud Foundation solution 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. By offloading the workload domain management to VMware Cloud Foundation and moving the infrastructure management into the cloud, the solution can respond to the speed and scale of customer deployments swiftly at cloud-scale.

Technology Overview

This chapter contains the following:

- FlexPod Datacenter
- Cisco Unified Compute System X-Series
- Cisco Nexus 93180YC-FX3 Ethernet Switch
- Cisco MDS 9132T 32G Multilayer Fabric Switch
- NetApp AFF A-Series Storage
- Cisco Intersight
- Infrastructure as Code with Ansible to setup FlexPod and VCF Management Domain
- VMware Cloud Foundation

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) storage systems

Figure 1. FlexPod Datacenter Components

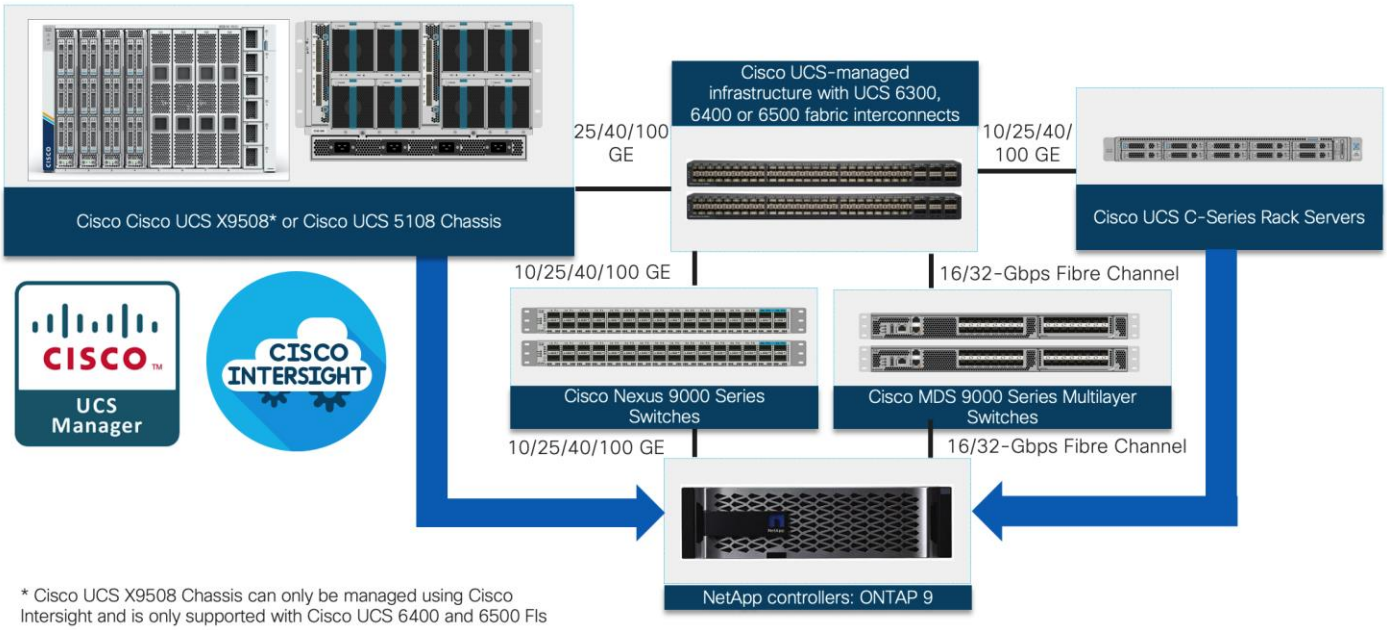

All the FlexPod components have been integrated so that customers can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlexPod is its ability 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 as a workload domain for VMware Cloud Foundation is built using the following hardware components:

- Cisco UCS X9508 Chassis with three Cisco UCS X210c M6 Compute Nodes
- Fourth-generation Cisco UCS 6454 Fabric Interconnects to support 25Gb and 100Gb ethernet connectivity and 32Gb FC connectivity from various components
- High-speed Cisco NX-OS-based Nexus 93180YC-FX3 switching design to support up to 100GE connectivity
- NetApp AFF A400 end-to-end NVMe storage with high-speed Ethernet (up to 100Gb) and 32 Gbps Fibre Channel connectivity
- Cisco MDS 9132T switches to support 32G FC connectivity

The software components to manage the FlexPod infrastructure consist of:

- Cisco Intersight platform 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 vSphere ESXi software installed on the Cisco UCS X210c compute nodes using Cisco custom ESXi software that includes up to date drivers and software

Note: VMware vCenter installation and configuration is performed by VMware Cloud Foundation and therefore VMware vCenter is covered in the overview section of VMware Cloud Foundation

Cisco Unified Computing System X-Series

The Cisco UCS X-Series Modular System is designed to take the current generation of the Cisco UCS platform to the next level with its future-ready design and cloud-based management. Decoupling and moving the platform management to the cloud allows Cisco UCS to respond to customer feature and scalability requirements in a much faster and efficient manner. Cisco UCS X-Series state of the art hardware simplifies the data center design by providing flexible server options. A single server type, supporting a broader range of workloads, results in fewer different datacenter products to manage and maintain.

Figure 2. Cisco UCS X9508 Chassis

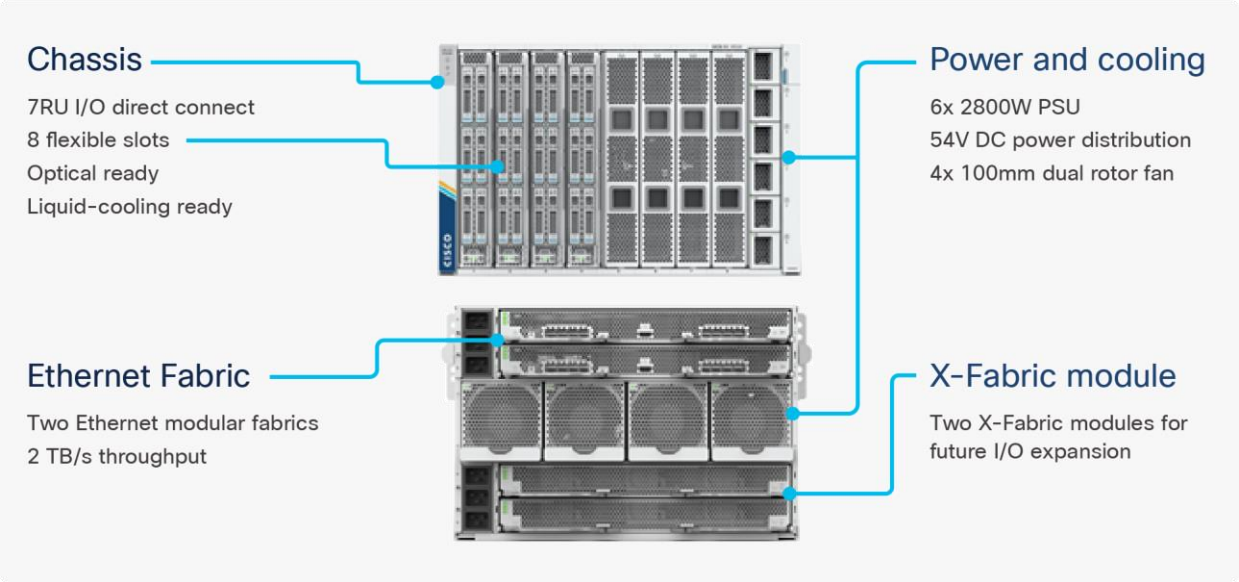

The various components of the Cisco UCS X-Series are described in the following sections.

Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As seen in Figure 3, Cisco UCS X9508 chassis has only a power-distribution midplane. This innovative design provides fewer obstructions for better airflow. For I/O connectivity, vertically oriented compute nodes intersect with horizontally oriented fabric modules, allowing the chassis to support future fabric innovations. Cisco UCS X9508 Chassis' superior packaging enables larger compute nodes, thereby providing more space for actual compute components, such as memory, GPU, drives, and accelerators. Improved airflow through the chassis enables support for higher power components, and more space allows for future thermal solutions (such as liquid cooling) without limitations.

Figure 3. Cisco UCS X9508 Chassis – Innovative Design

The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes, GPU accelerators and a pool of future I/O resources that may include disk storage, and memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) that connect the chassis to upstream Cisco UCS 6400 or 6500 Series Fabric Interconnects. At the bottom rear of the chassis are slots ready to house future X-Fabric modules that can flexibly connect the compute nodes with I/O devices. Six 2800W Power Supply Units (PSUs) provide 54V power to the chassis with N, N+1, and N+N redundancy. A higher voltage allows efficient power delivery with less copper and reduced power loss. Efficient, 100mm, dual

counter-rotating fans deliver industry-leading airflow and power efficiency, and optimized thermal algorithms enable different cooling modes to best support the customer's environment.

Cisco UCSX 9108-25G Intelligent Fabric Modules

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

Figure 4. Cisco UCSX 9108-25G Intelligent Fabric Module

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

Note: The current design was validated with Cisco UCSX 9108-25G IFMs.

Cisco UCS 9108-100G Intelligent Fabric Modules (for 100Gbps connectivity support)

In the end-to-end 100Gbps Ethernet design, for the Cisco UCS X9508 Chassis, the network connectivity is provided by a pair of Cisco UCS 9108-100G Intelligent Fabric Modules (IFMs). Like the fabric extenders used in the Cisco UCS 5108 Blade Server Chassis, these modules carry all network traffic to a pair of Cisco UCS 6536 Fabric Interconnects (FIs). IFMs also host the Chassis Management Controller (CMC) for chassis management.

Figure 5. Cisco UCS 9108-100G Intelligent Fabric Module

Each IFM supports eight 100Gb uplink ports for connecting the Cisco UCS X9508 Chassis to the Cisco UCS fifth generation 6556 FIs and 8 100Gb server ports for the eight compute nodes. IFM server ports can provide up to 200 Gbps of unified fabric connectivity per compute node across the two IFMs. The uplink ports connect the chassis to the Cisco UCS FIs, providing up to 1600Gbps connectivity across the two IFMs.

Cisco UCS X210c M6 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to 8 Cisco UCS X210c M6 Compute Nodes. The hardware details of the Cisco UCS X210c M6 Compute Nodes are shown in Figure 6:

Figure 6. Cisco UCS X210c M6 Compute Node

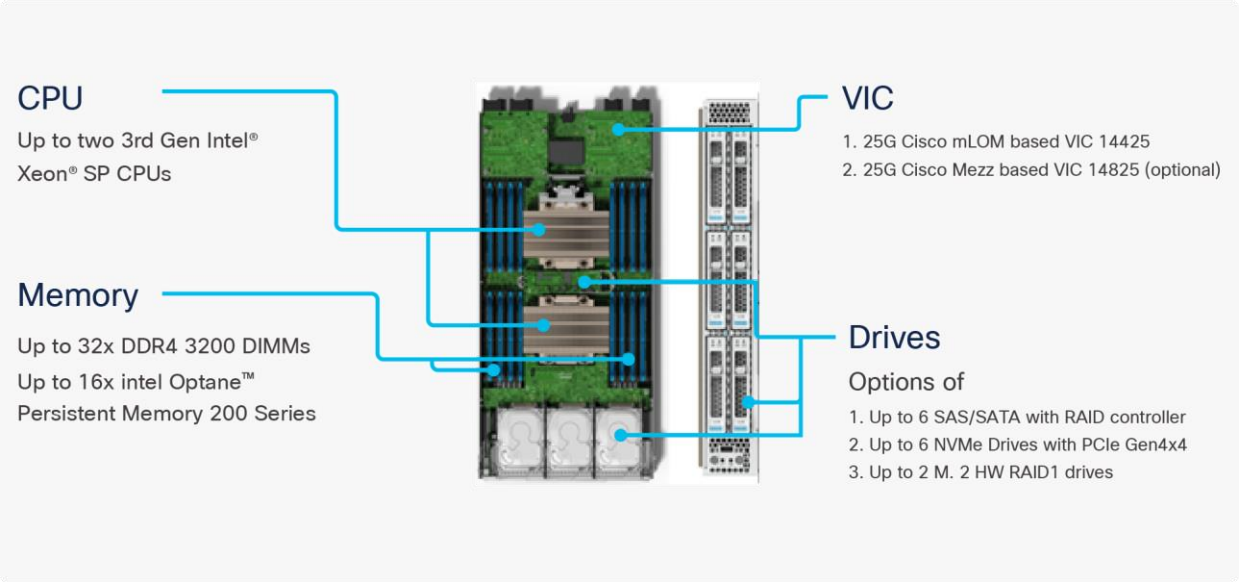

The Cisco UCS X210c M6 features:

- **CPU:** Up to 2 x 3rd Gen Intel Xeon Scalable Processors with up to 40 cores per processor and 1.5 MB Level 3 cache per core.
- **Memory:** Up to 32 x 256 GB DDR4-3200 DIMMs for a maximum of 8 TB of main memory. The Compute Node can also be configured for up to 16 x 512-GB Intel Optane persistent memory DIMMs for a maximum of 12 TB of memory.
- **Disk storage:** Up to 6 SAS or SATA drives can be configured with an internal RAID controller, or customers can configure up to 6 NVMe drives. 2 M.2 memory cards can be added to the Compute Node with RAID 1 mirroring.
- **Virtual Interface Card (VIC):** Up to 2 VICs including an mLOM Cisco VIC 14425* and a mezzanine Cisco VIC card 14825 can be installed in a Compute Node. Cisco VIC 15231 can also be used when setting up 100 Gbps compute connectivity.
- **Security:** The server supports an optional Trusted Platform Module (TPM). Additional security features include a secure boot FPGA and ACT2 anticounterfeit provisions.

Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS X210c M6 Compute Nodes support the following two 25Gbps Cisco fourth-generation VIC cards and one 100Gbps fifth generation VIC card:

Cisco VIC 14425

Cisco VIC 14425 fits the mLOM slot in the Cisco X210c Compute Node and enables up to 50 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 100 Gbps of connectivity per server. Cisco VIC 14425 connectivity to the IFM and up to the fabric interconnects is delivered through 4x 25-Gbps connections, which are configured automatically as 2x 50-Gbps port channels. Cisco VIC 14425 supports 256 virtual interfaces (both Fibre Channel and Ethernet) along with the latest networking innovations such as NVMeoF over RDMA (ROCEv2) and VxLAN/NVGRE offload.

Figure 7. Single Cisco VIC 14425 in Cisco UCS X210c M6

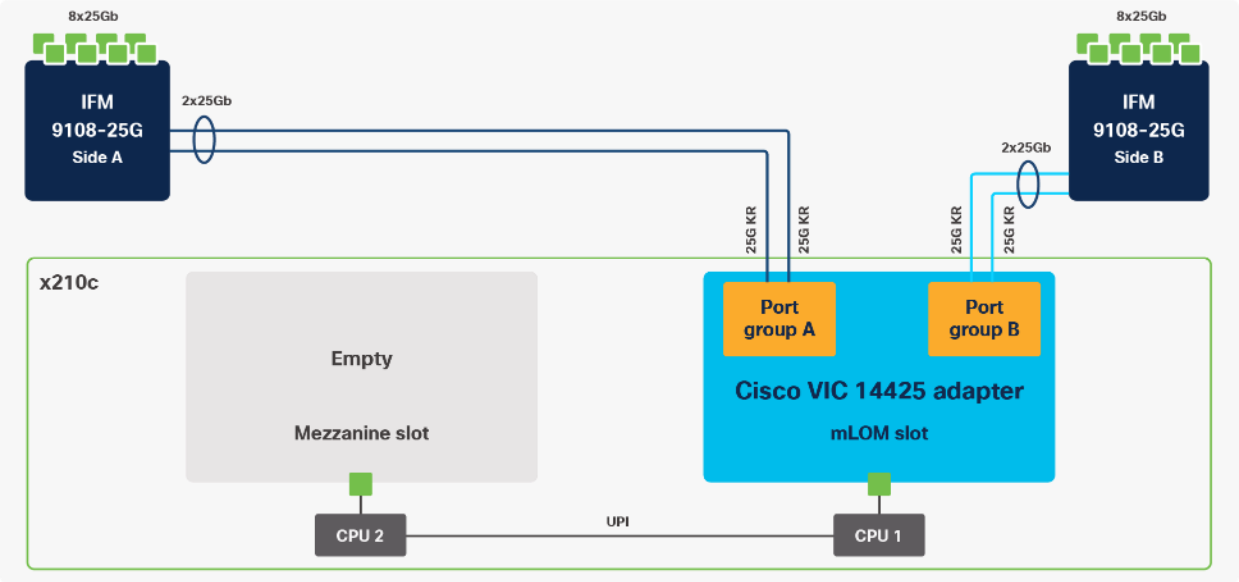

Note: In the current design, each compute node was installed with a single Cisco VIC 14425.

Cisco VIC 14825

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

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

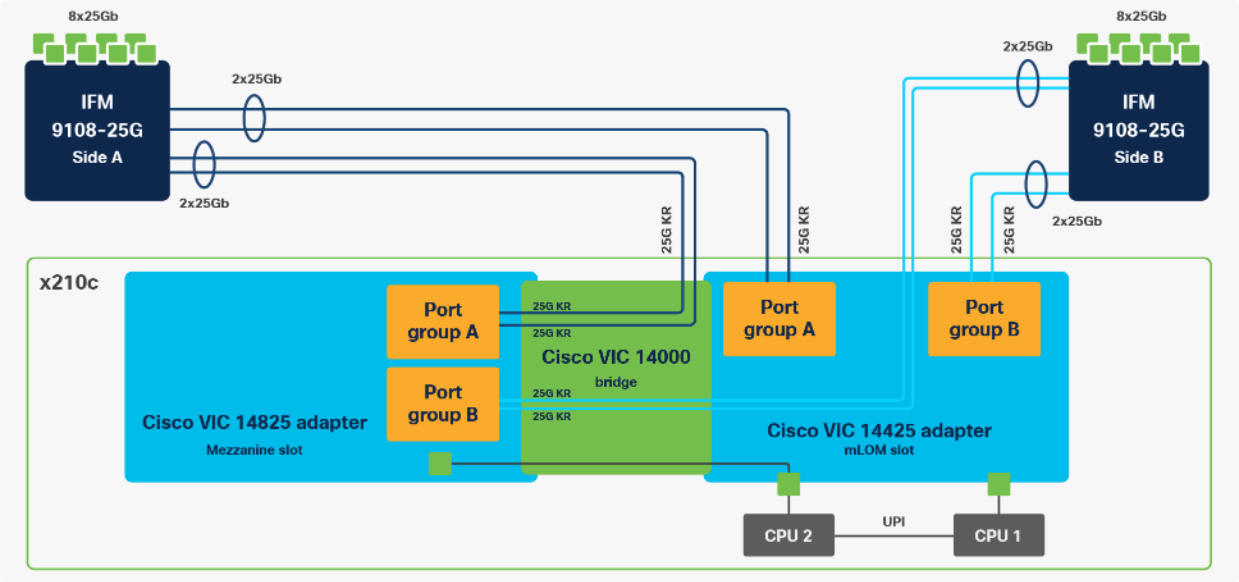

Cisco VIC 15231 (for 100Gbps connectivity support)

Cisco VIC 15231 fits the mLOM slot in the Cisco X210c Compute Node and enables up to 100 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 200 Gbps of connectivity per server. Cisco VIC 15231 connectivity to the IFM and up to the fabric interconnects is delivered through 100Gbps connections. Cisco VIC 15231 supports 512 virtual interfaces (both FCoE and Ethernet) along with the latest networking innovations such as NVMeoF over FC or TCP and VxLAN/NVGRE offload.

Figure 9. Single Cisco VIC 15231 in Cisco UCS X210c M6

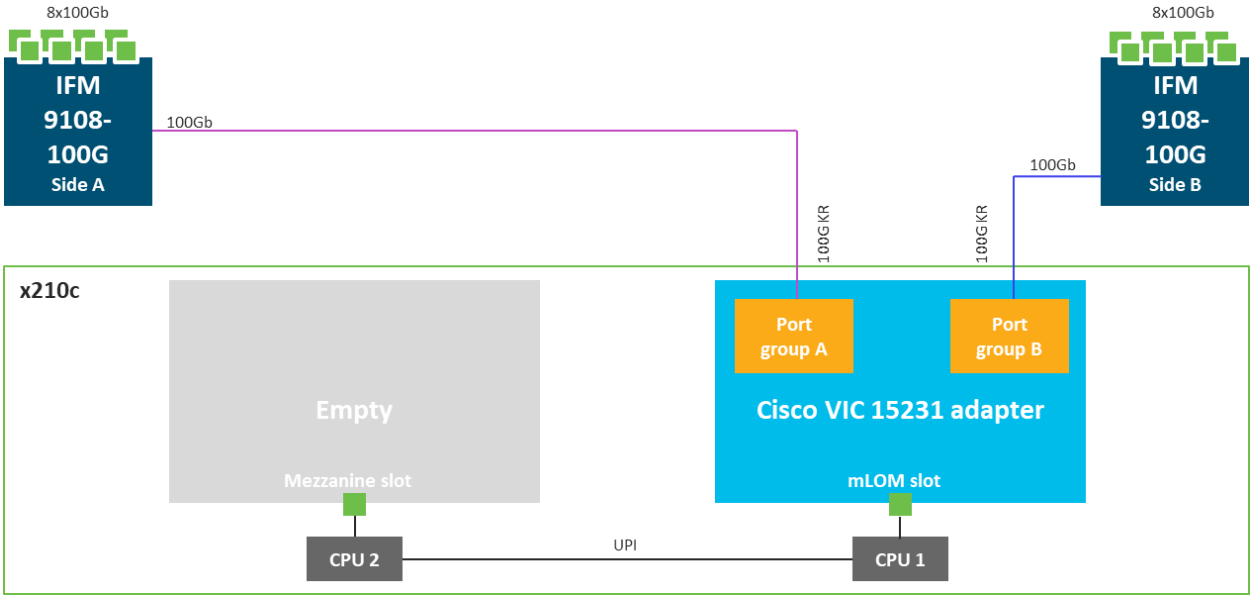

Cisco UCS 6400 Series Fabric Interconnects

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

Figure 10. Cisco UCS 6454 Fabric Interconnect

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

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

Cisco UCS 6536 Fabric Interconnects (for 100Gbps connectivity support)

The Cisco UCS fifth generation FI 6536 is a 36-port Fabric Interconnect. This single RU device includes up to 36 10/25/40/100 Gbps Ethernet ports and 32-Gbps Fibre Channel ports via 4 128 Gbps to 4x32 Gbps breakouts on ports 33-36. All 36 ports support ethernet breakout cables or QSA interfaces.

Figure 11. Cisco UCS 6536 Fabric Interconnect

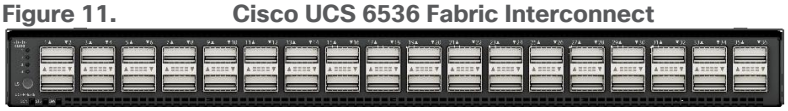

Note: The Cisco UCS 6536 FI currently only supports Intersight Managed Mode (IMM).

Cisco Nexus 93180YC-FX3 Ethernet Switch

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

Figure 12. Cisco Nexus 93180YC-FX3 Switch

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

Note: Cisco Nexus 93180YC-FX, 93360YC-FX2 and 9336C-FX2-E support SAN switching on the Cisco Nexus switch. SAN switching on Cisco Nexus (instead of MDS) allows customers using Cisco UCS 5th generation FI 6536 to utilize the 100G FCoE uplinks on FI to carry both ethernet and FC traffic eliminating the need for 128G to 4x32G FC breakout cables. Currently the SAN analytics feature is not supported when SAN switching is configured on Cisco Nexus.

Cisco MDS 9132T 32G Multilayer Fabric Switch

The Cisco MDS 9132T 32G Multilayer Fabric Switch is a highly reliable, flexible, and low-cost Cisco MDS 9100 Series switch. 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 13. Cisco MDS 9132T 32G Multilayer Fabric Switch

The Cisco MDS 9132T delivers advanced storage networking features and functions with ease of management and compatibility for reliable end-to-end connectivity. This switch also offers SAN analytics and telemetry capabilities built into this next-generation hardware platform. The telemetry 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 is 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.

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 customers accelerate, manage, and protect business-critical data on-prem and across hybrid clouds. As the first enterprise-grade storage systems to support both NVMe over Fibre Channel (NVMe/FC) and NVMe over TCP (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 applications. With a simple software upgrade to the modern NVMe/FC or

NVMe/TCP SAN infrastructure, customers can run more workloads, with faster response times, and without disruption or data migration.

NetApp offers a wide range of AFF-A series controllers to meet varying demands of the field. This solution design covers midrange, most versatile NetApp AFF A400 system featuring hardware acceleration technology that significantly enhances performance and storage efficiency.

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

Download technical specifications of the AFF A-series controllers here: https://www.netapp.com/us/media/ds-3582.pdf

NetApp AFF A400

The NetApp AFF A400 offers full end-to-end NVMe support. The frontend NVMe/FC 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. 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 25GbE or 100GbE, as well as 32Gb/FC and NVMe/FC network connectivity. 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.

Note: Cisco UCS X-Series, like Cisco UCS 5108, is supported with all NetApp AFF systems running NetApp ONTAP 9 release.

Figure 14. NetApp AFF A400 Front View

Figure 15. NetApp AFF A400 Rear View

NetApp ONTAP 9.11.1

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. 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 all-flash storage system 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 FAS or AFF 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.

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

See the ONTAP 9 release notes for more details on specific features and what's new: ONTAP® 9 Release Notes (netapp.com)

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 the storage systems and virtual infrastructure. The Unified Manager can be deployed on a Linux server, a Windows server, or as a virtual appliance on a VMware host.

Active IQ Unified Manager enables monitoring 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 on the storage infrastructure, Unified Manager can notify storage admins about the details of the issue to help identify the root cause. The virtual machine dashboard provides performance statistics for the VM so that users 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. Custom alerts can be configured for events so that when issues occur, notifications are sent via email or using SNMP Traps. Active IQ Unified Manager enables planning for the storage requirements by forecasting capacity and usage trends to proactively act before issues arise.

For more information on NetApp Active IQ Unified Manager, refer to the following link:

https://docs.netapp.com/us-en/active-iq-unified-manager/

NetApp ONTAP Tools for VMware vSphere

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

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

Virtual Storage Console (VSC)

NetApp Virtual Storage Console enables customers to perform the following tasks:

- Add storage controllers, assign credentials, and set up permissions for storage controllers

- Provision datastores
- Monitor the performance of the datastores and virtual machines in the vCenter Server environment
- View and update the host settings of the ESXi hosts that are connected to NetApp storage
- Control administrator access to the vCenter Server objects by using role-based access control (RBAC)

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 an Open API design that natively integrates with third-party platforms and tools.

Figure 16. Cisco Intersight Overview

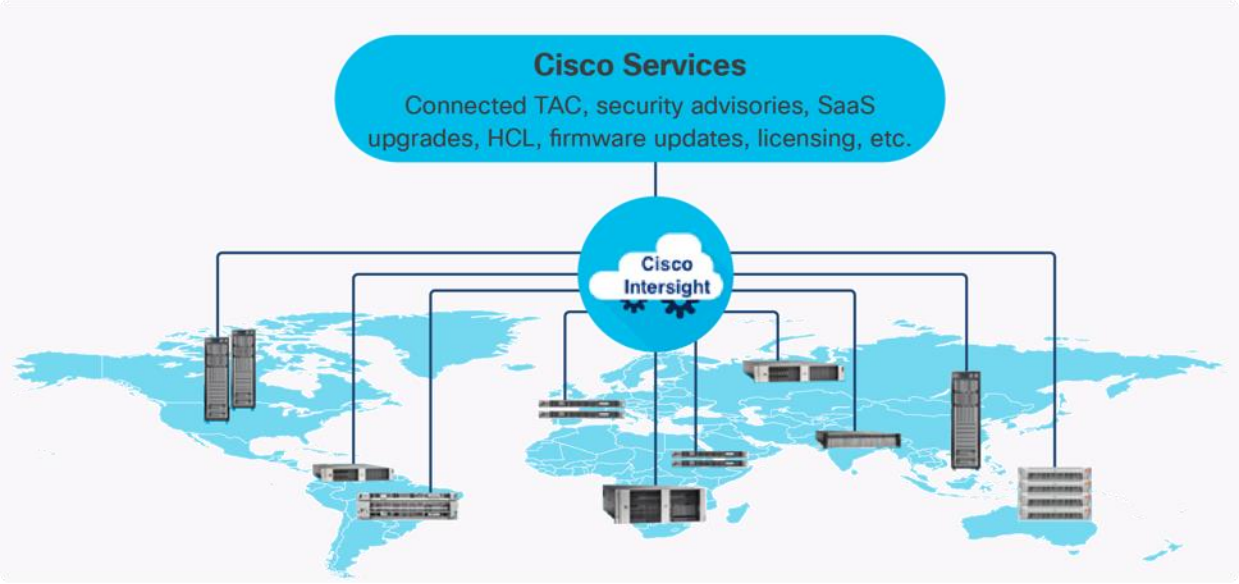

The main benefits of Cisco Intersight infrastructure services are as follows:

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

Cisco Intersight Virtual Appliance and Private Virtual Appliance

In addition to the SaaS deployment model running on Intersight.com, on-premises options can be purchased separately. The Cisco Intersight Virtual Appliance and Cisco Intersight Private Virtual Appliance are available for organizations that have additional data locality or security requirements for managing systems. The Cisco

Intersight Virtual Appliance delivers the management features of the Cisco Intersight platform in an easy-to-deploy VMware Open Virtualization Appliance (OVA) or Microsoft Hyper-V Server virtual machine that allows customers to control the system details that leave their premises. The Cisco Intersight Private Virtual Appliance is provided in a form factor specifically designed for users who operate in disconnected (air gap) environments. The Private Virtual Appliance requires no connection to public networks or back to Cisco to operate.

Cisco Intersight Assist

Cisco Intersight Assist helps customers add endpoint devices to Cisco Intersight. A data center could have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight, but does not connect directly with it, will need a connection mechanism. Cisco Intersight Assist provides that connection mechanism. In FlexPod, VMware vCenter and NetApp Active IQ Unified Manager, Cisco Nexus and Cisco MDS switches connect to Intersight with the help of 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 is covered in later sections.

Licensing Requirements

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

- **Cisco Intersight Essentials:** Essentials includes all the functions of the Base license plus additional features, including Cisco UCS Central Software and Cisco Integrated Management Controller (IMC) supervisor entitlement, policy-based configuration with server profiles, firmware management, and evaluation of compatibility with the Cisco Hardware Compatibility List (HCL).
- **Cisco Intersight Advantage:** Advantage offers all the features and functions of the Base and Essentials tiers. It includes storage widgets and cross-domain inventory correlation across compute, storage, and virtual environments. It also includes OS installation for supported Cisco UCS platforms.
- **Cisco Intersight Premier:** In addition to all 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 detailed information about the features provided in the various licensing tiers, see https://intersight.com/help/getting_started#licensing_requirements.

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

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

- Cisco Intersight uses the device connector running within Cisco Intersight Assist virtual appliance to communicate with the VMware vCenter.
- Cisco Intersight uses the device connector running within a Cisco Intersight Assist virtual appliance to integrate with NetApp Active IQ Unified Manager. The NetApp AFF A400 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 17. Cisco Intersight and vCenter/NetApp Integration

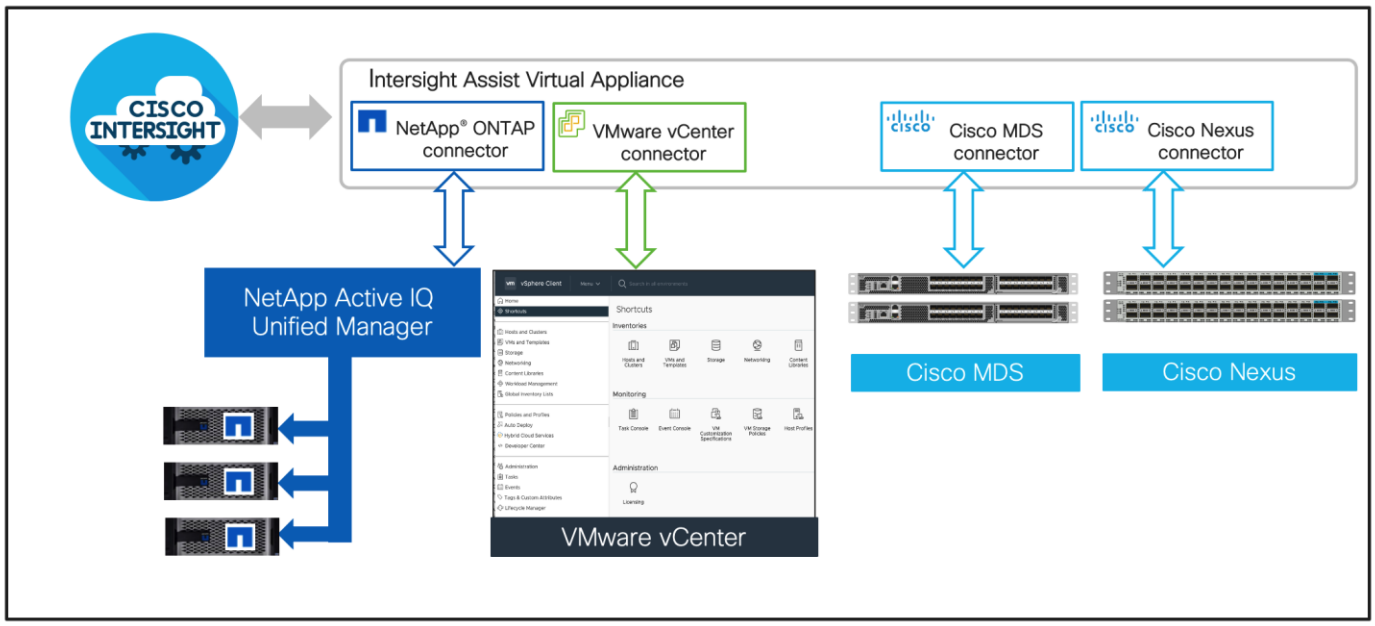

The device connector provides a secure way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware 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 Cisco switching, VMware or ONTAP 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 NX-OS command line interface (CLI) for comprehensive analysis, diagnostics, and reporting of virtual and storage environments. The functionality provided through this integration is covered in the upcoming solution design section.

FlexPod XCS

FlexPod XCS is Cisco Intersight Integrated System that groups the FlexPod components (Cisco UCS, NetApp ONTAP storage, Cisco switches, and VMware vCenter) into an Integrated System. This grouping enhances full-stack visibility and provides FlexPod-level dashboards and widgets within the stack. For current information on FlexPod XCS, see https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/flexpod-xcs-solution-with-intersight-wp.html.

Figure 18. FlexPod XCS in Cisco Intersight

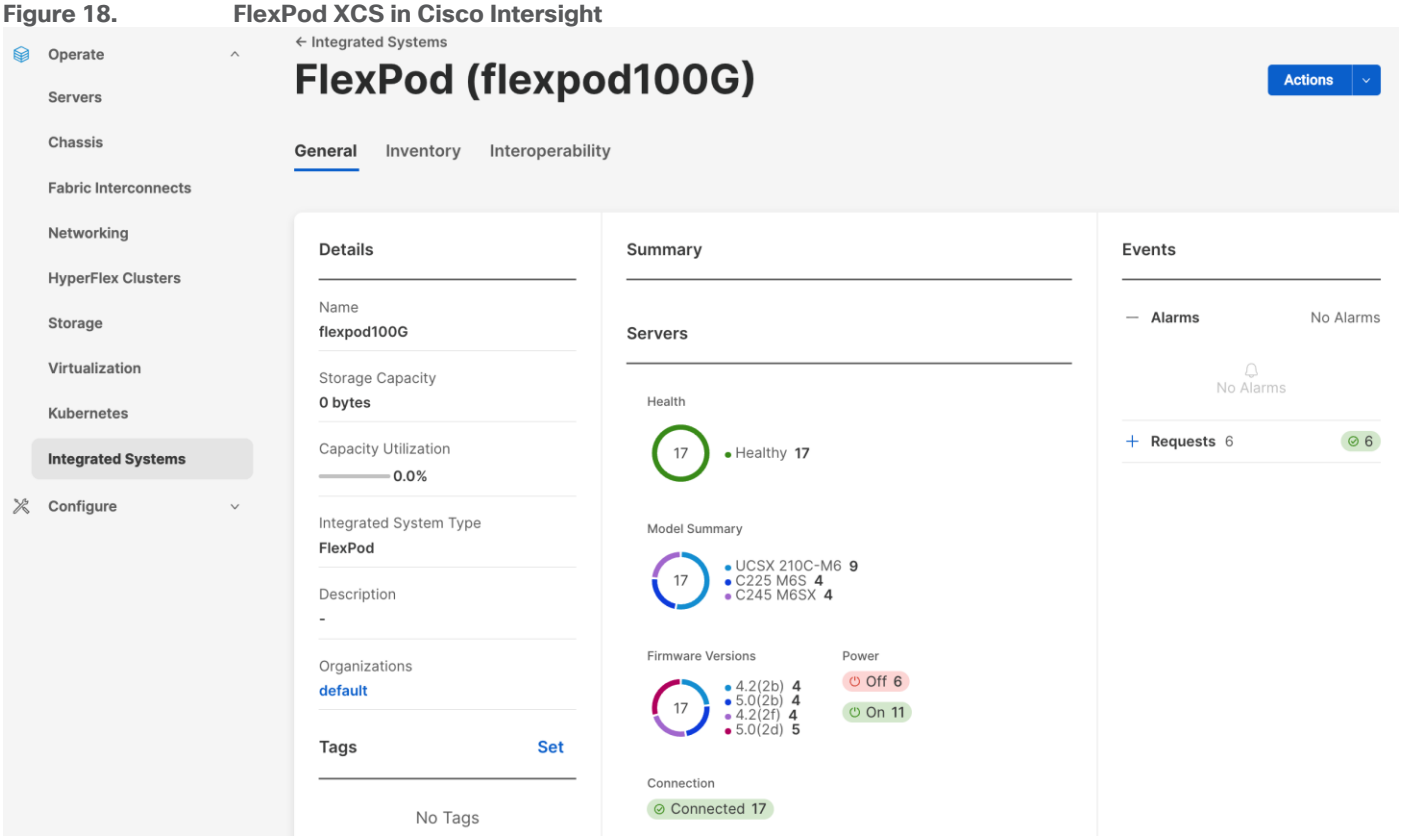

Infrastructure as Code with Ansible to Setup FlexPod and VCF Management Domain

This FlexPod solution provides a fully automated solution deployment that covers all components of the infrastructure. The configuration of the Cisco Network and Compute, NetApp ONTAP Storage, and VMware vSphere are automated by leveraging Ansible playbooks that have been developed to setup the components according to the solution best practices. Customers can use Ansible automation to configure the management domain servers as well as FlexPod Virtual Infrastructure (VI) domain servers, install vSphere ESXi on these servers, setup various required parameters (such as setting up NTP, enabling SSH, and so on) and present the servers for commissioning through VMware Cloud Foundation.

The automated deployment using Ansible provides a well-defined sequence of steps across the different elements of this solution. The automated deployment involves exchange of parameters or attributes between compute, network, storage, and virtualization and require some level of manual intervention. The workflow is clearly defined and documented for the customers. The Ansible playbooks to configure the different sections of the solution invoke a set of Roles which consume several user configurable variables. Based on the installation environment, customers can choose to modify the variables to suit their needs 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 FlexPod VI workload domain is onboarded, NetApp Management Tools such as ONTAP Tools for VMware vSphere (formerly Virtual Storage Console), SnapCenter Plug-in for VMware vSphere, and Active IQ Unified Manager can also be deployed in an automated fashion.

VMware Cloud Foundation

Based on a proven and comprehensive software-defined stack including VMware vSphere with VMware Tanzu, VMware vSAN, VMware NSX-T Data Center, and VMware vRealize Suite, VMware Cloud Foundation provides a complete set of software-defined services for compute, storage, network, container, and cloud management. The result is agile, reliable, efficient cloud infrastructure that offers consistent operations across private and public clouds. VMware Cloud Foundation enables data center cloud administrators to provision an application environment in a rapid, repeatable, automated way versus the traditional manual process.

Figure 19. **VMware Cloud Foundation overview**

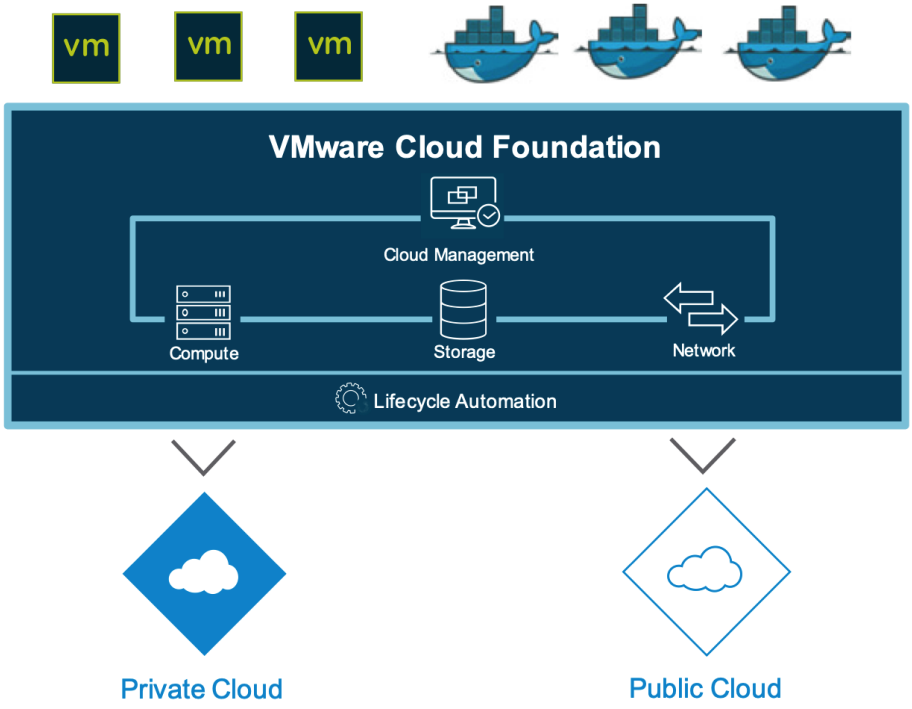

Management Domain

The VMware Cloud Foundation management domain contains the components for managing the virtual infrastructure of the workload domains in the SDDC. The management domain infrastructure requires vSAN datastore deployment to host all the infrastructure VMs for various workload domains. The management domain is created during the bring-up process by VMware Cloud Builder and contains the VMware Cloud Foundation management components as follows:

- Minimum four ESXi hosts
- An instance of vCenter Server
- A three-node NSX Manager cluster
- SDDC Manager
- vSAN datastore

Workload Domains in VMware Cloud Foundation

VMware Cloud Foundation consists of workload domains which represent an application-ready infrastructure. A workload domain represents a logical unit that groups ESXi hosts managed by a vCenter Server instance with

specific characteristics according to VMware best practices. A workload domain can consist of one or more vSphere clusters, provisioned automatically by SDDC Manager.

Each workload domain contains the following components:

- ESXi hosts
- One VMware vCenter Server™ instance
- At least one vSphere cluster with vSphere HA and vSphere DRS enabled.
- At least one vSphere Distributed Switch per cluster for system traffic and NSX segments for workloads.
- One NSX Manager cluster for configuring and implementing software-defined networking.
- One* NSX Edge cluster, added after you create the workload domain, which connects the workloads in the workload domain for logical switching, logical dynamic routing, and load balancing.
- One or more shared storage allocations.

Note: * Customers can still use the traditional VLAN based application networking where the gateways are configured on the upstream hardware switches (for example, Cisco Nexus 9000) and slow migrate their applications to NSX-T edge cluster as needed.

VI Workload Domains

For each VI workload domain, customers can choose multiple storage options - vSAN, NFS, vVols, or VMFS on FC. A VI workload domain consists of one or more vSphere clusters. Each cluster starts with a minimum of three hosts and can scale up to the vSphere maximum of 64 hosts. SDDC Manager automates the creation of the VI workload domain and the underlying vSphere clusters.

For the first VI workload domain, SDDC Manager deploys a vCenter Server instance and a three-node NSX Manager cluster in the management domain. For each subsequent VI workload domain, SDDC Manager deploys an additional vCenter Server instance. New VI workload domains have an option to share the same NSX Manager cluster with an existing VI workload domain, or a new NSX Manager cluster can be deployed. VI workload domains cannot use the NSX Manager cluster created for the management domain.

Note: FlexPod is added as a VI workload domain to VMware Cloud Foundation.

VMware Cloud Builder

VMware Cloud Builder automates the deployment of the software-defined stack, creating the first software-defined unit known as the management domain.

VMware Cloud Foundation SDDC Manager

SDDC Manager allows customers to perform administration tasks on the VMware Cloud Foundation instance. The SDDC Manager UI provides an integrated view of the physical and virtual infrastructure and centralized access to manage the physical and logical resources. From the SDDC Manager interface, the virtual infrastructure administrator or cloud administrator can provision new private cloud resources, monitor changes to the logical infrastructure, and manage life cycle and other operational activities.

vSphere

vSphere uses virtualization to transform individual data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. VMware vSphere manages these infrastructures as a unified operating environment and provides customers with the tools to administer the data centers that participate in

that environment. The two core components of VMware vSphere are ESXi and vCenter Server. ESXi is the virtualization platform which allows customers to create and run virtual machines and virtual appliances. vCenter Server is the service through which virtualization administrators manage multiple hosts.

VMware vSphere 7.0 U3

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

VMware vSphere 7.0 U3 has several improvements and simplifications including, but not limited to:

- Support for the NVMe-TCP storage protocol with VMFS6 datastores
- Improvements to vSphere Cluster Services (vCLS), including the ability to designate a datastore to store vCLS virtual machines
- Improved Maintenance Mode Reliability and Workload Placement
- Enhanced Performance Statistics for Memory
- vSphere Lifecycle Management (vLCM) with Hardware Support Manager (HSM) Integration with Cisco Intersight
- A VMware-Recommended 128GB SAN boot LUN for VMware ESXi

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

VMware vSphere vCenter

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

vSAN

vSAN aggregates local or direct-attached data storage devices to create a single storage pool that is shared across all hosts in the vSAN cluster. Using vSAN simplifies storage configuration and virtual machine provisioning. Built-in policies allow for flexibility in data availability. VMware Cloud Foundation requires vSAN to be configured as storage on the management domain ESXi hosts. The VI workload domains, e.g., FlexPod, support additional common storage types including NFS, VMFS on FC, and vVol.

Note: For the FlexPod based VI workload domain, NFS storage is utilized.

NSX-T Data Center

NSX-T Data Center is focused on providing networking, security, automation, and operational simplicity for emerging application frameworks and architectures that have heterogeneous endpoint environments and technology stacks. NSX-T Data Center supports cloud-native applications, bare-metal workloads, multi-hypervisor environments, public clouds, and multiple clouds.

Note: The FlexPod based VI workload domain is configured for NSX-T networking. The workload domain(s) also support the traditional VLAN based application deployment.

vRealize Suite Lifecycle Manager

VMware Cloud Foundation supports automated deployment of vRealize Suite Lifecycle Manager. Customer can then deploy and manage the life cycle of various components and deploy the additional VMware vRealize Suite products (vRealize Log Insight, vRealize Automation, and vRealize Operations Manager) by using vRealize Suite Lifecycle Manager.

VMware Cloud Foundation Deployment Architecture

VMware Cloud Foundation supports two architecture models, standard and consolidated.

Standard Architecture Model

With the standard architecture model, management workloads run on a dedicated management domain and customer workloads are deployed in separate virtual infrastructure (VI) workload domains. Each workload domain is managed by a separate vCenter Server instance which provides for scalability and allows for autonomous licensing and life cycle management.

Figure 20. VMware Cloud Foundation Standard Architecture

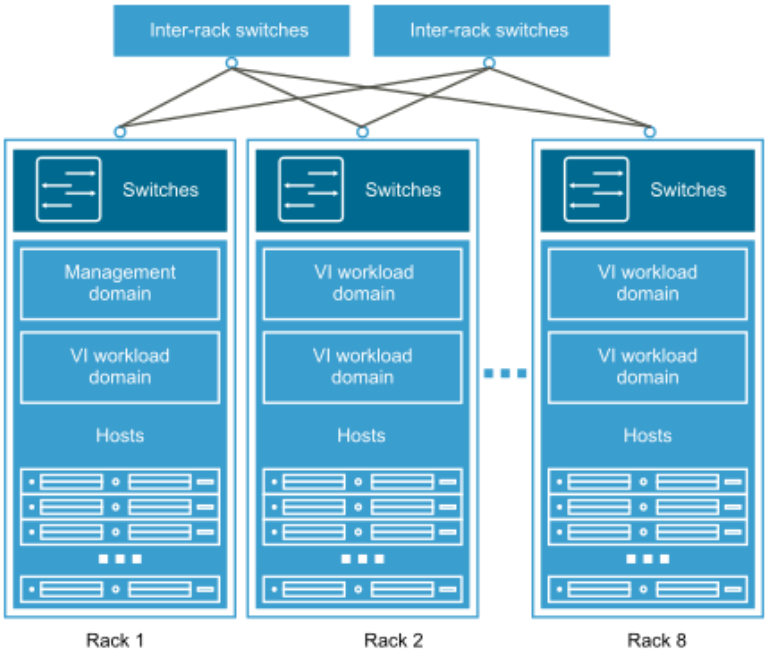

Standard architecture is the recommended model because it aligns with the VMware best practice of separating management workloads from customer workloads and provides better long-term flexibility and expansion options.

Consolidated Architecture Model

In the consolidated model, both the management and customer workloads run together on a shared management domain. The environment is managed from a single vCenter Server and vSphere resource pools provide isolation between management and customer workloads. Resource pools must be properly configured as the domain is shared by the management and compute workloads.

Figure 21. VMware Cloud Foundation Consolidated Architecture

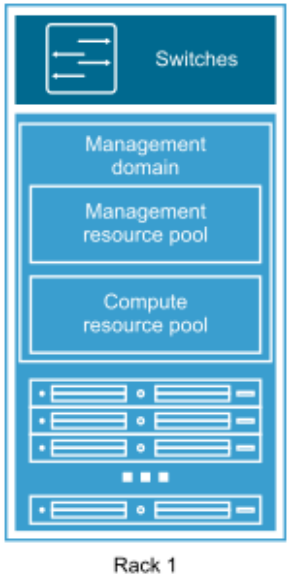

When additional hosts are added to a VMware Cloud Foundation system deployed on a consolidated architecture, customers can migrate to the standard architecture by creating a VI workload domain and moving the customer workload VMs from the compute resource pool to the newly created VI workload domain.

Note: In the current FlexPod as a workload domain for VMware Cloud Foundation design, VMware Cloud Foundation standard architecture was deployed with FlexPod added as a VI workload domain.

Solution Design

This chapter contains the following:

The FlexPod as a workload domain for VMware Cloud Foundation delivers a VMware Cloud Foundation VI workload domain solution built on Cisco UCS X-Series based FlexPod infrastructure. VMware vSphere 7.0 U3 hypervisor is installed on the Cisco UCS X210c M6 Compute Nodes configured for stateless compute design using boot from SAN. NetApp AFF A400 provides the storage infrastructure required for setting up the VMware environment. The Cisco Intersight cloud-management platform is utilized to configure and manage the infrastructure. The solution requirements and design details are explained in this section.

For setting up the VMware Cloud Foundation management domain, 4 vSAN ready or vSAN certified Cisco UCS C-Series compute nodes are also needed in addition the FlexPod infrastructure.

Requirements

The FlexPod as a workload domain for VMware Cloud Foundation 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 VMware Cloud Foundation and other external automation tools
- Cloud-enabled design which can be configured, managed, and orchestrated from the cloud using GUI or APIs

Physical Topology

FlexPod as a workload domain for VMware Cloud Foundation supports both Fibre Channel (FC) and IP based storage access design.

FlexPod Datacenter with Fibre Channel Design

For the FC designs, NetApp AFF A400 and Cisco UCS X-Series are connected through Cisco MDS 9132T Fibre Channel Switches and boot from SAN for stateless compute uses the FC network. When adding FlexPod as VI workload domain, both NFS and VMFS on FC storage is supported.

Note: In the current FlexPod as a workload domain for VMware Cloud Foundation design, only NFS storage was verified for FC based FlexPod design.

FlexPod Datacenter with IP-only Design

For the IP-only solution, ESXi is installed on local M2 drives of Cisco UCS X210c compute nodes. When adding FlexPod as VI workload domain, only NFS storage is supported. iSCSI based datastores can be added as supplementary storage after the ESXi hosts are commissioned and workload domain is defined in VMware cloud foundation.

Note: The IP based FlexPod design is supported but was not validated as part of the current FlexPod as a workload domain for VMware Cloud Foundation solution.

VMware Cloud Foundation management domain with FlexPod VI workload domain

The physical topology for FlexPod as a workload domain for VMware Cloud Foundation is shown in Figure 22.

Figure 22. VMware Cloud Foundation with FlexPod

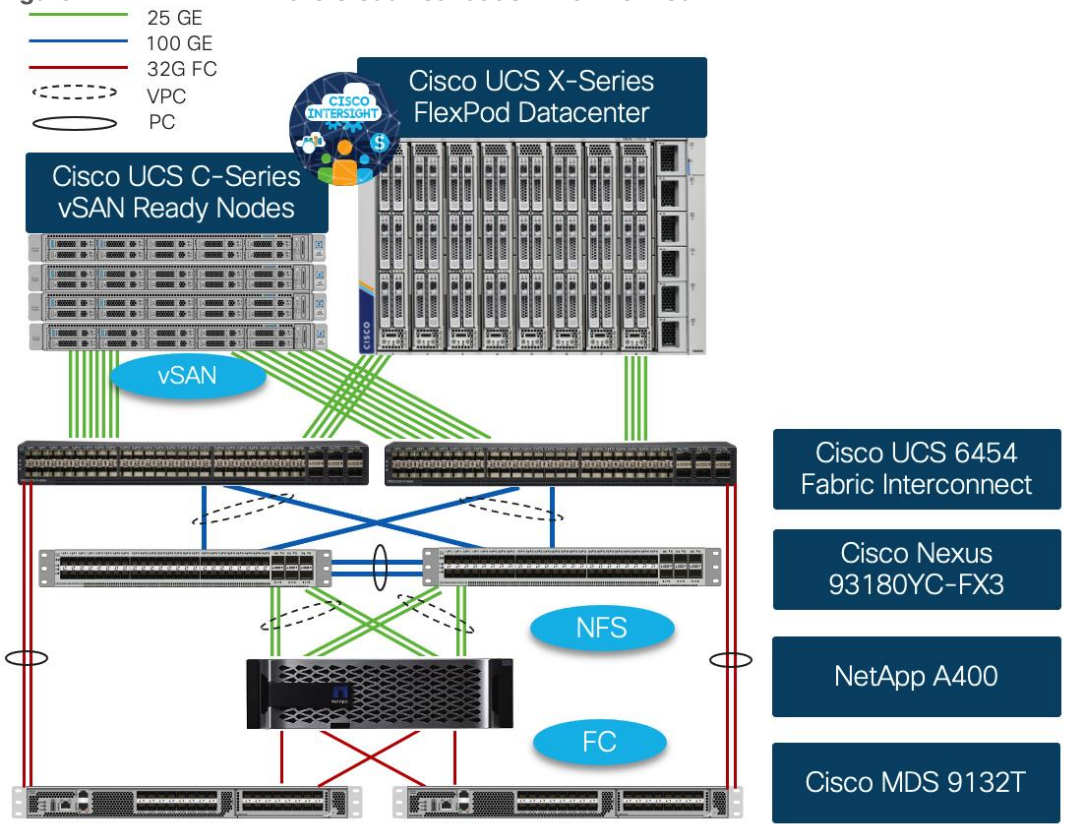

To validate the FlexPod as a workload domain for VMware Cloud Foundation configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the rack server and chassis to network connectivity.

- 4 Cisco UCS C-Series* vSAN ready nodes are connected to fabric interconnects (FI) and are managed using Cisco Intersight. Two 25 Gigabit Ethernet ports from each Cisco UCS C-Series server are connected to each FI.
- The Cisco UCS X9508 Chassis connects to FIs using Cisco UCSX 9108-25G Intelligent Fabric Modules (IFMs), where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate FI. Remaining 4 ports from each IFM can be connected FIs if additional bandwidth is required.
- Cisco Nexus 93180YC-FX3 Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 100 Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX3 Switches in a vPC configuration.
- The NetApp AFF A400 controller connects to the Cisco Nexus 93180YC-FX3 Switches using four 25 GE ports from each controller configured as a vPC for NFS traffic.
- For compute SAN connectivity, Cisco UCS 6454 Fabric Interconnects connect to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections configured as a single port channel.
- For storage SAN connectivity, each NetApp AFF A400 controller connects to both Cisco MDS 9132T switches using 32-Gbps Fibre Channel.
- VMware 7.0 U3 ESXi software is installed on Cisco UCS C-Series servers to deploy the VMware Cloud Foundation management cluster.
- VMware 7.0 U3 ESXi software is also installed on Cisco UCS X210c M6 compute nodes to provide FlexPod based VI workload domain.

Note: * Since Cisco UCS C-series is being managed and configured by Cisco Intersight Managed Mode, the vSAN ready nodes must satisfy the software and hardware requirements outlined here: https://intersight.com/help/saas/supported_systems

VLAN Configuration

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

Table 1. VLAN Usage

VLAN ID	Name	Description	Subnet
2	Native-VLAN	Use VLAN 2 as native VLAN instead of default VLAN (1)	
1010	Mgmt	Existing management VLAN where all the management interfaces for various devices will be connected	10.101.0.0/24
1011	IB-Mgmt	FlexPod In-band management VLAN is utilized for all in-band management connectivity such as ESXi hosts, VM management, and VCF components (Cloud Builder, SDDC Manager, all NSX managers, all vCenters)	10.101.1.0/24
1012	VM-Traffic	Application VLAN (one of many) where application VMs will be deployed. Adjust the name and add more VLANs as needed.	10.101.2.0/24
1017	NFS	VLAN for ESXi NFS datastore access in FlexPod VI workload domain	10.101.7.0/24
3001	Mgmt-vSAN	vSAN VLAN for the management domain	192.168.1.0/24
3002	Mgmt-Host-	NSX-T Host Overlay Network VLAN for the management domain	192.168.2.0/24

VLAN ID	Name	Description	Subnet
	Overlay		
3003	WD-Host-Overlay	NSX-T Host Overlay Network VLAN for the FlexPod VI workload domain	192.168.3.0/24
3030	vMotion	A common vMotion VLAN for both management and VI workload domains	192.168.30.0/24

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

- VLAN 1010 is the management VLAN where out of band management interfaces of all the physical devices are connected.
- VLAN 1011 is used for in-band management of VMs, ESXi hosts, and other infrastructure services in the FlexPod environment. This VLAN is also used for deploying VMware Cloud Foundation components.
- VLAN 1012 is used as application traffic VLAN. This VLAN will be mapped to a port-group on VDS02 of the VI workload domain ESXi hosts. Customers can modify the name or add additional VLANs as needed.
- VLAN 1017 provides FlexPod VI workload domain ESXi hosts access to the NSF datastores hosted on the NetApp Controllers. NFS storage is used as primary storage for VI domain.
- VLAN 3001 is used for VMware Cloud Foundation management domain vSAN configuration.
- VLANs 3002 and 3003 are separate NSX-T host overlay VLANs for VMware Cloud Foundation management and FlexPod VI workload domains. Depending on the customer requirements, a single VLAN can be used.
- VLAN 3030 is the common VM vMotion VLAN for VMware Cloud Foundation management and FlexPod VI workload domains. Depending on the customer requirements, separate VLANs can be used for vMotion traffic isolation.

Physical Components

Table 2 lists the required hardware components used to build the validated solution. Customers are encouraged to review their requirements and adjust the size or quantity of various components as needed.

Table 2. FlexPod as a Workload Domain for VMware Cloud Foundation Hardware Components

Component	Hardware	Comments
Cisco Nexus Switches	Two Cisco Nexus 93180YC-FX3 switches	
Cisco MDS Switches	Two Cisco MDS 9132T switches	
NetApp A400	A NetApp AFF A400 with appropriate storage and network connectivity	Customer requirements will determine the amount of storage. The NetApp A400 should support both 25Gbps (or 100 Gbps) ethernet and 32Gbps (or 16 Gbps) FC connectivity
Fabric Interconnects	Two Cisco UCS 6454 Fabric Interconnects	These fabric interconnects will be shared between the management and the workload domain
Management Domain Compute		

Component	Hardware	Comments
Cisco UCS Servers	A minimum of four Cisco UCS C-Series vSAN ready (or vSAN compatible) nodes	vSAN ready nodes are recommended for ease of deployment however, customers can also utilize existing Cisco UCS C-Series servers with vSAN supported components
FlexPod VI Workload Domain Compute		
Cisco UCS Chassis	A minimum of one UCS X9508 chassis.	Single chassis can host up to 8 Cisco UCS X210c compute nodes
Cisco UCS Compute Nodes	A minimum of three Cisco UCS X210c compute nodes	Four compute nodes are recommended but three compute nodes will work.

Software Components

Table 3 lists various software releases used in the solution. The exact versions of the components listed in Table 3 and additional drivers and software tool (for example, various NetApp software tools, Cisco Intersight Assist, and so on) versions will be covered in the deployment guide.

Table 3. Software Components and Versions

Component	Version
Cisco Nexus 93180YC-FX3	9.3(10)
Cisco MDS 9132T	8.4(2d)
Cisco UCS Fabric Interconnects	4.2(2c)
Cisco UCS C-Series vSAN ready nodes	4.2(2a)
Cisco UCS X210c compute nodes	5.0(2b)
NetApp A400 - ONTAP	9.11.1
VMware Cloud Foundation	
Cloud Builder VM	4.4.1
SDDC Manager	4.4.1
VMware vCenter	7.0 Update 3d
VMware ESXi	7.0 Update 3d
VMware NSX-T	3.1.3.7.4

Logical Design

VMware Cloud Foundation (VCF) deployment is fully automated. Before starting the VCF deployment process for either the management domain or the VI workload domain, Cisco UCS C-Series vSAN ready nodes as well as Cisco UCS X210c compute nodes need to be configured with appropriate compute policies (BIOS etc.), Network Interface Card (NIC) and VLAN configuration. All Cisco UCS servers are equipped with a Cisco Virtual Interface Card (VIC) configured for multiple virtual Network Interfaces (vNICs). The server design as well as

connectivity including VLAN/VSAN usage between the server profile for an ESXi host and the storage configuration on NetApp AFF A400 controllers is captured in the following subsections.

Management Domain - vSAN Ready Nodes

Each management domain server profile is deployed on Cisco UCS C-Series vSAN ready nodes. VMware Cloud foundation supports three different network profiles for vSphere Distribute Switch (VDS) deployment. In this deployment, the ESXi hosts are configured where two VDSs are deployed using 4 NICs. First (primary) VDS carries management, vMotion, and vSAN traffic while the second VDS carries Host Overlay network to be used by NSX-T.

Note: Application specific VLANs (if required) are configured on the second VDS.

Figure 23 illustrates the ESXi host design and network connectivity for the VMware Cloud Foundation management domain.

Figure 23. Management Domain ESXi Host

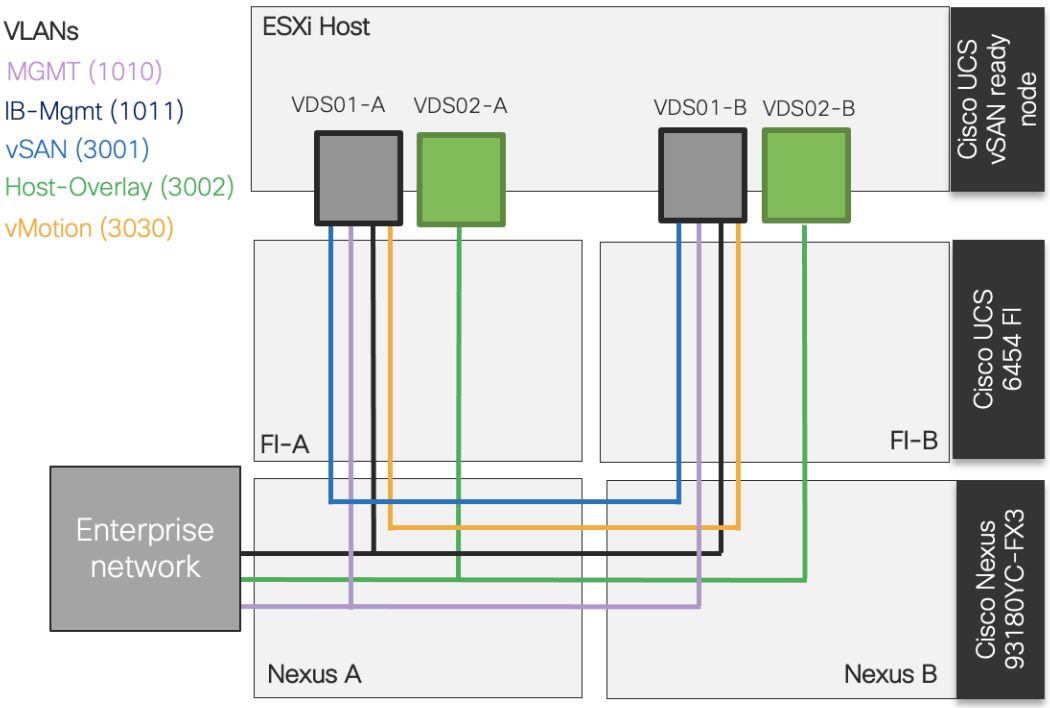

ESXi host design highlights:

- The four vNICs in each ESXi host are configured as follows:
  - Two redundant vNICs (VDS01-A and VDS01-B) are configured for the primary VDS to carry management, VMware vSAN, and vMotion traffic. Jumbo MTU (9000) is enabled on these vNICs.
  - Two redundant vNICs (VDS02-A and VDS02-B) are configured for the second VDS and carry NSX-T host overlay network traffic. Jumbo MTU (9000) is enabled on these vNICs.
- VMware vSphere ESXi is installed on local disk of the host interactively using Cisco custom ISO.
- A DNS entry is created for the ESXi host.
- Management IP address is assigned to the host and host and domain name is configured.

- SSH access is enabled on the host.
- vSwitch0 (default virtual switch) MTU is set to 9000.
- Default VM Network port group is updated with the management VLAN because this port group is used to deploy the management vCenter during installation.
- NTP server is configured on the host.
- Self-Signed certificates are regenerated on the host.
- The disks to be used by VMware vSAN are presented to the ESXi host in JBOD configuration

VI Workload Domain - FlexPod Hosts

Each VI workload domain server profile is deployed on Cisco UCS X210c compute nodes in the FlexPod. In this deployment, the ESXi hosts are configured with two VDSs using 4 NICs. First VDS carries management and NFS traffic while the second VDS carries, vMotion, NSX-T Host Overlay network, and application traffic VLANs.

Figure 24 illustrates the ESXi host design and connectivity for the VMware Cloud Foundation VI workload domain.

Figure 24. VI Workload Domain ESXi Host

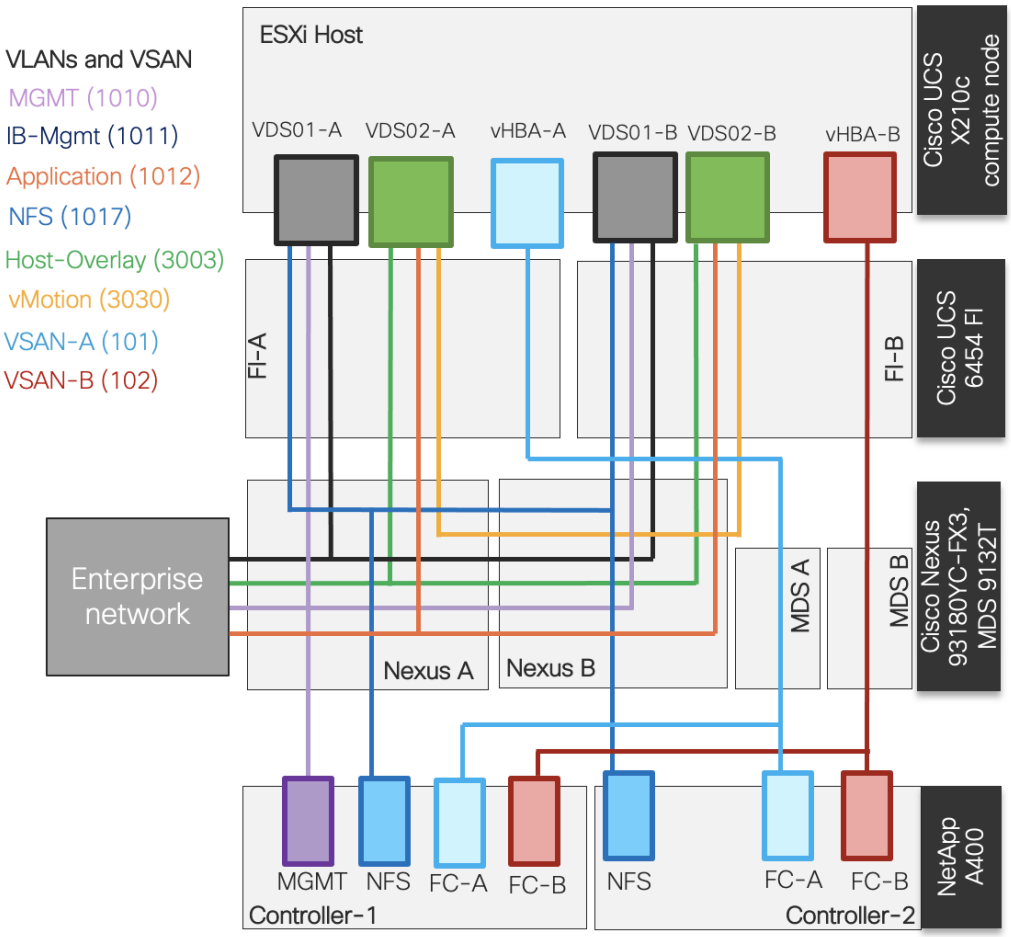

ESXi host design highlights:

- The four vNICs in each ESXi host are configured as follows:

- Two redundant vNICs (VDS01-A and VDS01-B) are configured for the first VDS to carry management, NFS, and vMotion traffic. Jumbo MTU (9000) is enabled on these vNICs.
- Two redundant vNICs (VDS02-A and VDS02-B) are configured for the second VDS and carry vMotion, NSX-T host overlay network, and application traffic. Jumbo MTU (9000) is enabled on these vNICs.

- Stateless boot from SAN using FC is configured for the host and boot LUN is configured on NetApp A400.
- vHBA-A is defined on Fabric A to provide access to SAN-A path for accessing NetApp A400.
- vHBA-B defined on Fabric B to provide access to SAN-B path for accessing NetApp A400.
- VMware vSphere ESXi is installed using the boot LUN on each host interactively using Cisco custom ISO.
- A DNS entry is created for the ESXi host.
- Management IP address is assigned to the host and host and domain name is configured.
- SSH access is enabled on the host.
- vSwitch0 (default virtual switch) MTU is set to 9000.
- NTP server is configured on the host.
- Self-Signed certificates are regenerated on the host.

Note: An NFS based volume is configured on NetApp A400 and each ESXi host will mount the NFS datastores from NetApp AFF A400 over the NFS VLAN.

Compute System Connectivity

vSAN Ready Nodes Connectivity

The Cisco UCS C-Series vSAN Ready Nodes are equipped with the Cisco 4th generation VIC (Cisco VIC 1457 or 1467). Each C-series server connects to each Cisco UCS 6454 FI using two 25GE ports, as shown in Figure 25.

Figure 25. Cisco UCS C-Series vSAN Ready Nodes Connectivity to Cisco UCS Fabric Interconnects

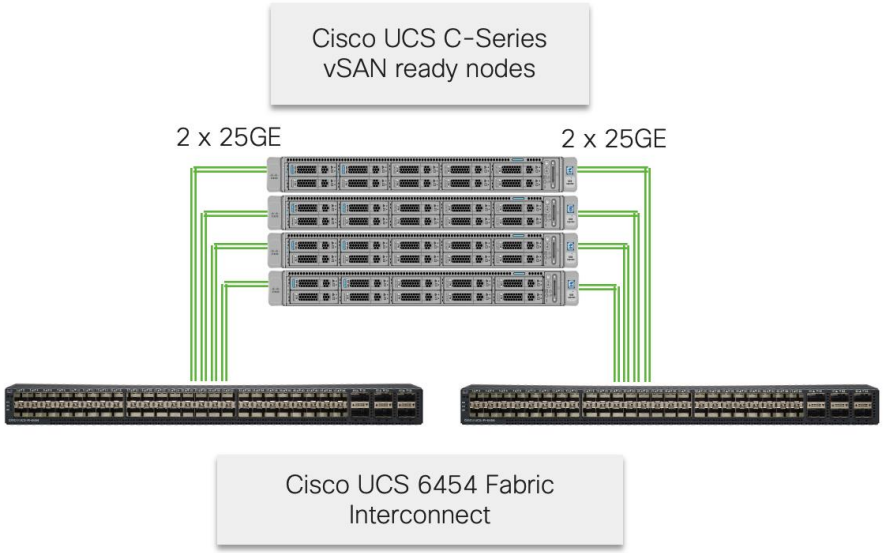

Cisco UCS X-Series Connectivity

The Cisco UCS X9508 Chassis is equipped with the Cisco UCSX 9108-25G intelligent fabric modules (IFMs) and connects to each Cisco UCS 6454 FI using four 25GE ports, as shown in Figure 26. If the customers require more bandwidth, all eight ports on the IFMs can be connected to each FI.

Figure 26. Cisco UCS X9508 Chassis connectivity to Cisco UCS Fabric Interconnects

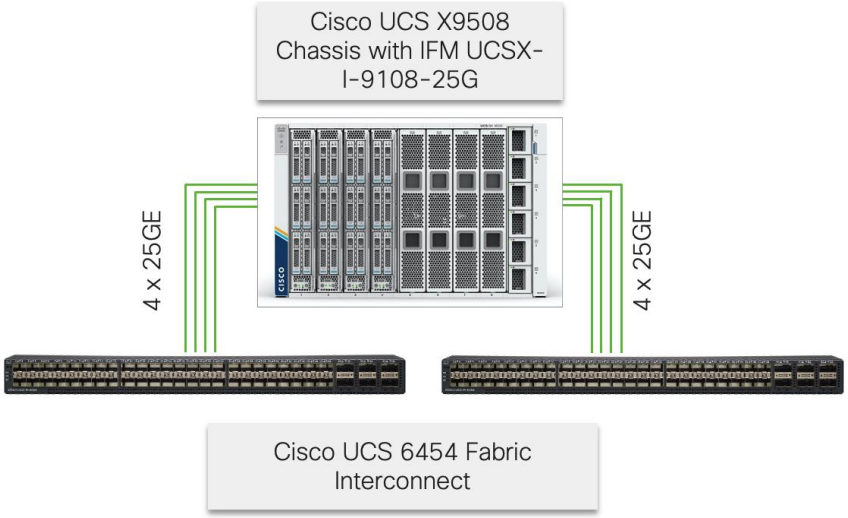

Cisco Nexus Ethernet Connectivity

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

- Feature interface-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 to Cisco Nexus 93180YC-FX3 switches using 100GE connections configured as virtual port channels. Both FIs are connected to both Cisco Nexus switches using a 100G connection. Additional links can easily be added to the port channel to increase the network bandwidth as needed. Figure 27 illustrates the physical connectivity details.

Figure 27. Cisco UCS 6454 FI Ethernet Connectivity

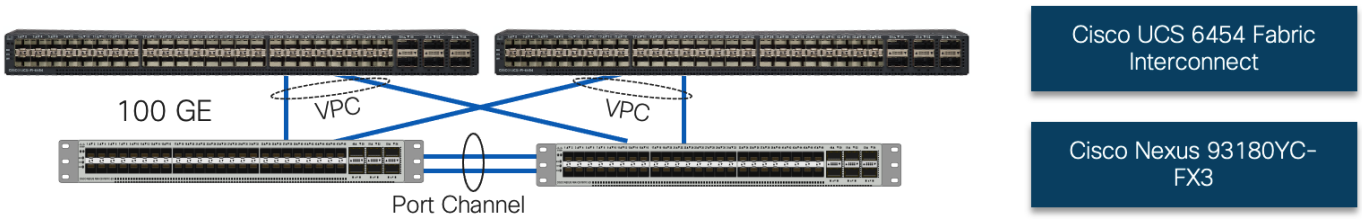

NetApp AFF A400 Ethernet Connectivity

NetApp AFF A400 controllers are connected to Cisco Nexus 93180YC-FX3 switches using 25GE connections configured as virtual port channels. The storage controllers are deployed in a switchless cluster configuration and are connected to each other using the 100GE ports e3a and e3b. Figure 28 illustrates the physical connectivity details.

Note: In Figure 28, the two storage controllers in the high-availability pair are drawn separately for clarity. Physically, the two controllers exist within a single chassis.

Figure 28. NetApp AFF A400 Ethernet Connectivity

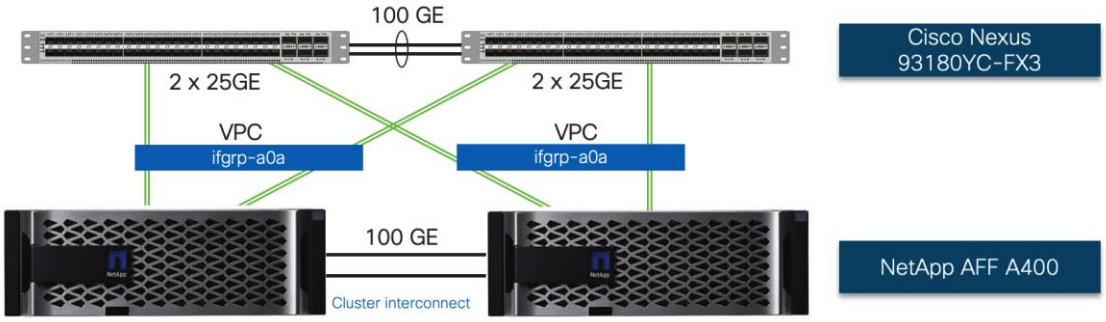

Cisco MDS SAN Connectivity

The Cisco MDS 9132T brings together the 32Gbps Fibre Channel (FC) capabilities to the FlexPod design. A redundant 32 Gbps FC 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 by identifying the initiators and targets in the environment.

Cisco UCS Fabric Interconnect 6454 SAN Connectivity

For SAN connectivity, each Cisco UCS 6454 Fabric Interconnect is connected to a Cisco MDS 9132T SAN switch using 2 x 32G Fibre Channel port-channel connection, as shown in Figure 29.

Figure 29. Cisco UCS 6454 Fabric Interconnect SAN Connectivity

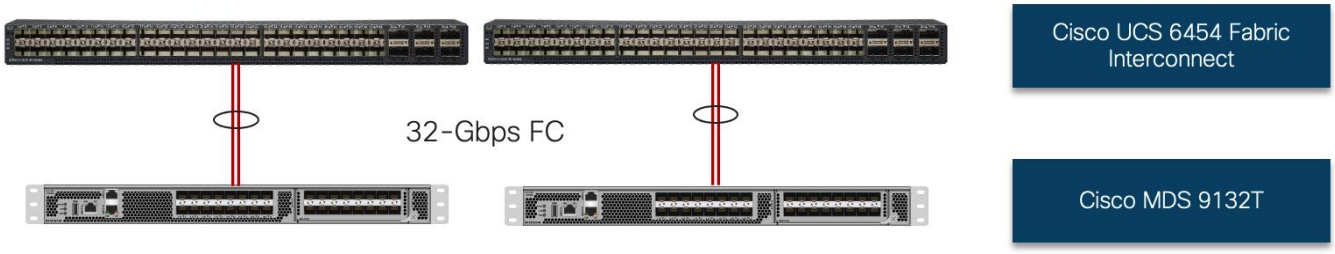

NetApp AFF A400 SAN Connectivity

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 30.

Note: In Figure 30. , the two storage controllers in the high-availability pair are drawn separately for clarity. Physically, the two controllers exist within a single chassis.

Figure 30. NetApp AFF A400 SAN Connectivity

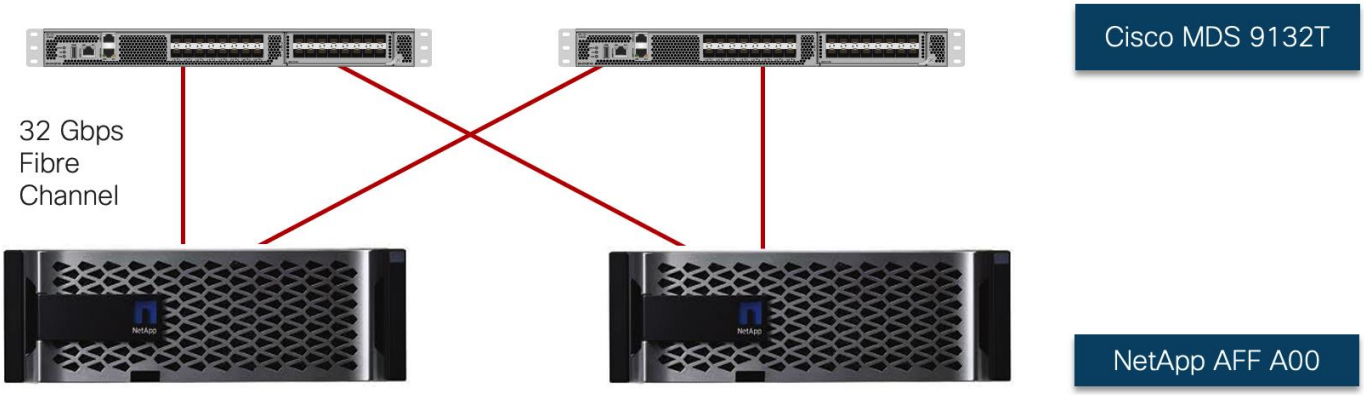

Cisco UCS Configuration - Cisco Intersight Managed Mode

Cisco Intersight Managed Mode standardizes policy and operation management for Cisco UCS hardware used in this CVD. The Cisco UCS vSAN ready nodes and the Cisco UCS X210c compute nodes are configured using separate server profiles in Cisco Intersight. These server profiles derive all the server characteristics from various associated policies and templates. Some of the design highlights of the Cisco UCS configuration using Intersight Managed Mode are explained below.

Set up Cisco UCS Fabric Interconnect for Cisco Intersight Managed Mode

During the initial setup of the FIs, the Intersight Managed Mode (IMM) must be selected. Cisco UCS X-Series can only be managed when FIs are configured to support IMM.

Note: Customers can switch the management mode for the fabric interconnects between Cisco Intersight and Cisco UCS Manager at any time however this is a disruptive process.

Claim a Cisco UCS Fabric Interconnect in the Cisco Intersight Platform

After setting up the Cisco UCS fabric interconnect for Cisco Intersight Managed Mode, FIs must be claimed to a new or an existing Cisco Intersight account. When a Cisco UCS fabric interconnect is successfully added to the Cisco Intersight platform (Figure 31), all future configuration steps are completed in the Cisco Intersight web based graphical user interface (GUI).

Figure 31. Cisco UCS Fabric Interconnect as an IMM Target in Cisco Intersight

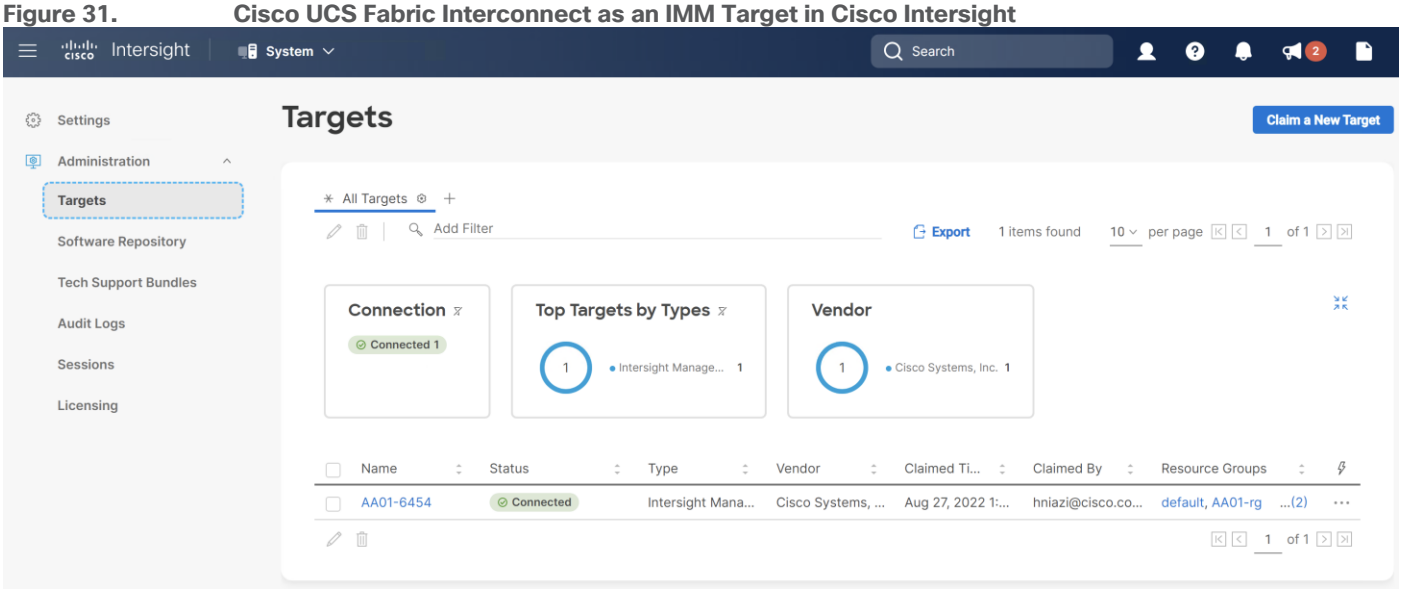

Cisco UCS Chassis Profile (Optional)

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

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

Cisco UCS Domain Profile

A Cisco UCS domain profile configures a fabric interconnect pair through reusable policies, allows configuration of the ports and port channels, and configures the VLANs and VSANs to be used in the network. It defines the characteristics of and configures the ports on the fabric interconnects. One Cisco UCS domain profile can be assigned to one fabric interconnect domain, and the Cisco Intersight platform supports the attachment of one port policy per Cisco UCS domain profile.

Some of the key characteristics of the Cisco UCS domain profile for setting up FlexPod as a workload domain for VMware Cloud Foundation are:

- A single domain profile is created for the pair of Cisco UCS fabric interconnects.
- Unique port policies are defined for the two fabric interconnects.
- All the VLANs used in the VMware Cloud Foundation management and VI workload domain are defined in the VLAN policy.
- The VLAN configuration policy is common to the fabric interconnect pair because both fabric interconnects are configured for the same set of VLANs (Figure 32).

- The VSAN configuration policies are unique for the two fabric interconnects because the VSANs are unique (Figure 32). All the server ports including the ports connected to Cisco C-Series vSAN ready nodes and the Cisco UCS X-Series chassis are enabled in the port policies.
- The Network Time Protocol (NTP), network connectivity, Link Control (UDLD), and system Quality-of-Service (QoS) policies are common to the fabric interconnect pair.

Figure 32. Domain Profile showing VLAN and VSAN Policies

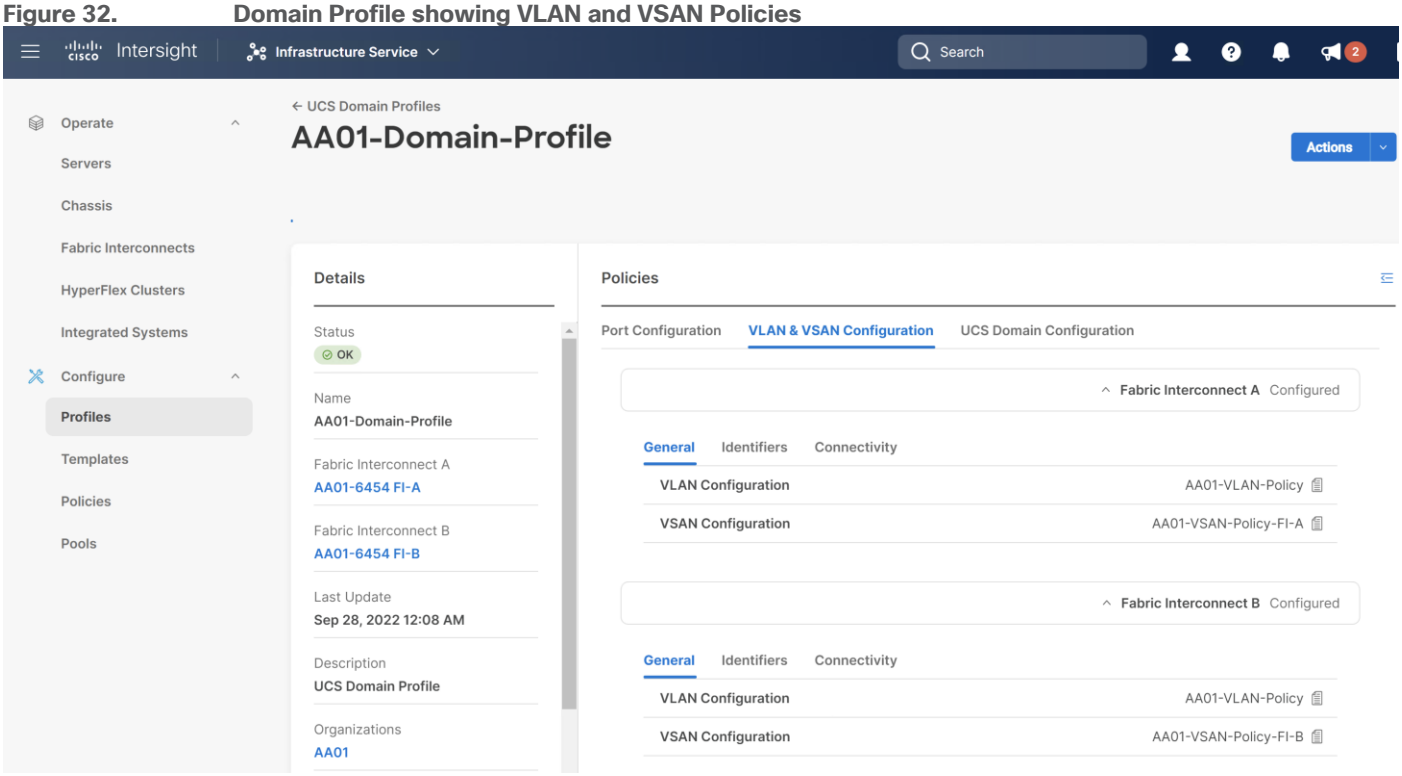

After the Cisco UCS domain profile has been successfully created and deployed, the policies including the port policies are pushed to the Cisco UCS fabric interconnects. Cisco UCS vSAN ready nodes and the Cisco UCS X9508 Chassis and Cisco UCS X210c M6 Compute Nodes are automatically discovered when the ports are successfully configured using the domain profile.

Server Profile Template

A server profile template enables resource management by simplifying policy alignment and server configuration. A server profile template is created using the server profile template wizard. Figure 33 shows various policies that should be defined for creating a server profile template.

Figure 33. Cisco Intersight – Server Profile Template

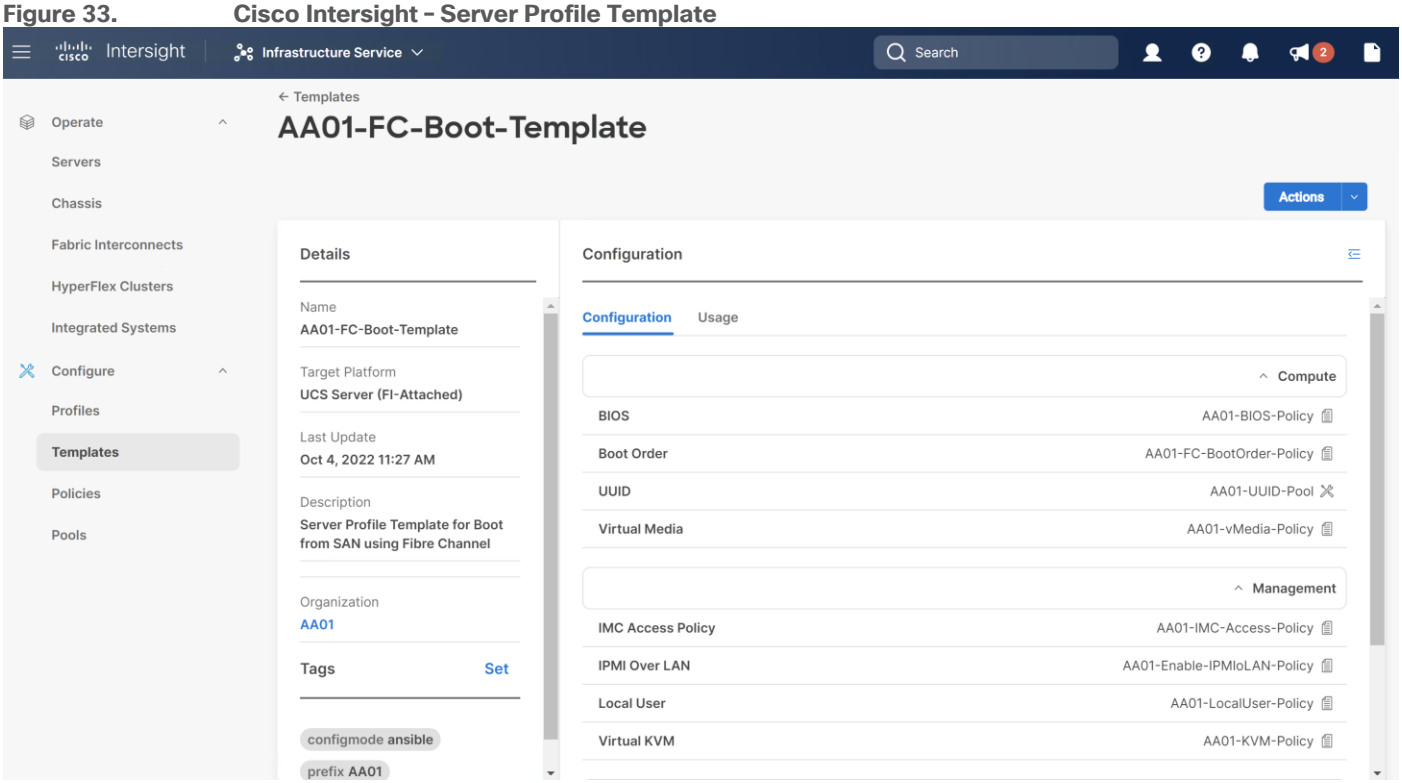

VMware Cloud Foundation Management Domain Server Profile Template

Some of the characteristics of the server profile template for FlexPod are as follows:

- BIOS policy is created to specify various server parameters in accordance with Cisco UCS Performance Tuning Guides.
- Boot order policy defines virtual media (KVM mapped DVD) and local disk as ESXi installation options.
- IMC access policy defines the management IP address pool for KVM access.
- Local user policy is used to enable KVM-based user access.
- LAN connectivity policy is used to create four virtual network interface cards (vNICs); two for primary vSphere Distributed Switch (VDS01) and two for second Virtual Distributed Switch (VDS02). Various policies and pools are also defined during the vNIC configuration.
- Appropriate VLANs are enabled on each of the vNIC in the LAN connectivity policy as covered in logical design.

VMware Cloud Foundation VI Workload Domain Server Profile Template

Some of the characteristics of the server profile template for FlexPod are as follows:

- BIOS policy is created to specify various server parameters in accordance with FlexPod best practices and Cisco UCS Performance Tuning Guides.
- Boot order policy defines virtual media (KVM mapped DVD) and all SAN paths for NetApp Fibre Channel logical interfaces (LIFs).
- IMC access policy defines the management IP address pool for KVM access.
- Local user policy is used to enable KVM-based user access.

- LAN connectivity policy is used to create four virtual network interface cards (vNICs); two for management /NFS vSphere Distributed Switch (VDS01) and two for vMotion, application traffic and NSX-T vSphere Distributed Switch (VDS02); along with various policies and pools.
- SAN connectivity policy is used to create two virtual host bus adapters (vHBAs) one each for SAN A and for SAN B; along with various policies and pools.
- Appropriate VLANs are enabled on each of the vNIC in the LAN connectivity policy as covered in logical design.

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

The Cisco Intersight server profile allows server configurations to be deployed directly on the compute nodes based on polices defined in the server profile template. After the management and VI workload server profile templates have been successfully created, server profiles can be derived from the template and associated with the appropriate Cisco UCS Compute Nodes. Figure 34 shows multiple (six) server profiles, associated with six UCS compute nodes, derived form a single server profile template.

Figure 34. Multiple Server Profiles derived from a single Server Profile Template

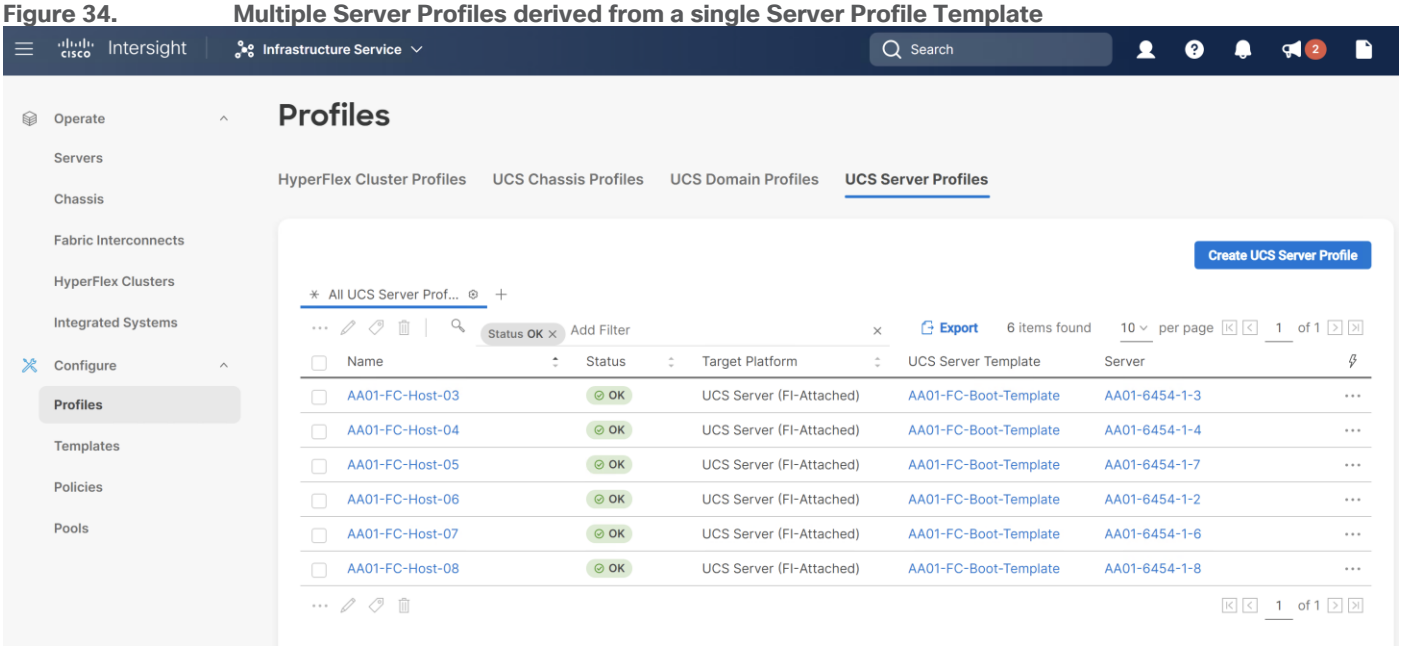

NetApp AFF A400 – Storage Virtual Machine (SVM) Design

To provide the necessary data segregation and management, a dedicated SVM is created for each FlexPod VI workload domain. The SVM contains the following volumes and logical interfaces (LIFs):

- Volumes
  - ESXi boot LUNs used to enable ESXi host boot from SAN functionality.
  - NFS datastore that acts as primary storage for the VMware Cloud Foundation VI workload domain.
  - NFS based VM swap file datastore and any additional/optional Infrastructure datastores for the VI workload domain to be added after workload domain provisioning is complete in VCF.
- Logical interfaces (LIFs):
  - NFS LIFs to mount NFS datastores in the vSphere environment.
  - FC LIFs for supporting FC SAN traffic.

Details of volumes, VLANs, and logical interfaces (LIFs) are shown in Figure 35.

Figure 35. NetApp AFF A400 – Infra-SVM for FC Boot

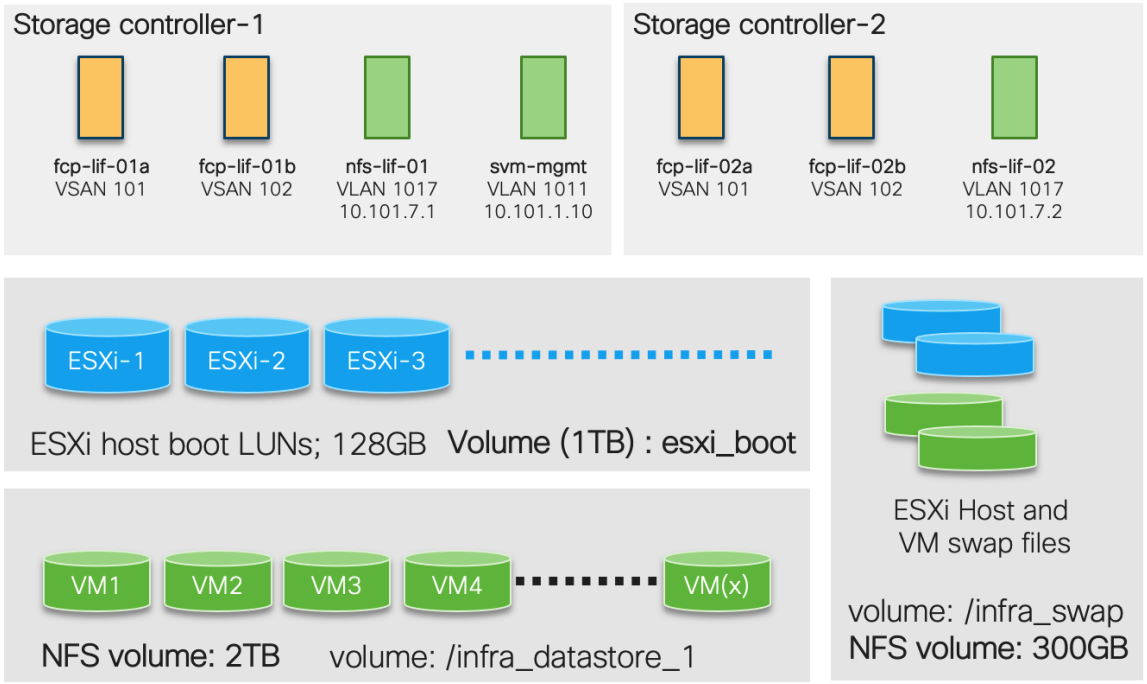

VMware Cloud Foundation Design

The deployment of VMware Cloud Foundation is automated. To manage the infrastructure, VMware Cloud Foundation augments the VMware virtualization and management components with VMware Cloud Builder and VMware Cloud Foundation SDDC Manager. VMware Cloud Builder is used to deploy the management domain, SDDC Manager is used to deploy Virtual Infrastructure (VI) workload domains for customer workloads, and VMware vRealize Suite Lifecycle Manager in VMware Cloud Foundation mode is used to deploy optional vRealize Suite products.

VMware Cloud Foundation deployment can be divided into three unique steps:

- Preparing the existing infrastructure*
- Deploying the management domain*
- Deploying the VI workload domain

Note: *For customers who only need to onboard FlexPod VI workload domain in an existing VMware Cloud Foundation setup, the management domain setup and related infrastructure preparation steps can be skipped. This design guide assumes customers are setting up a new VMware Cloud Foundation from the beginning.

Prepare the Existing Infrastructure

Before starting the automated deployment of the management domain using VMware Cloud Builder, the environment must meet target prerequisites and be in a specific starting state. Figure 36 shows the prerequisite components.

Figure 36. VCF Existing Infrastructure Requirements

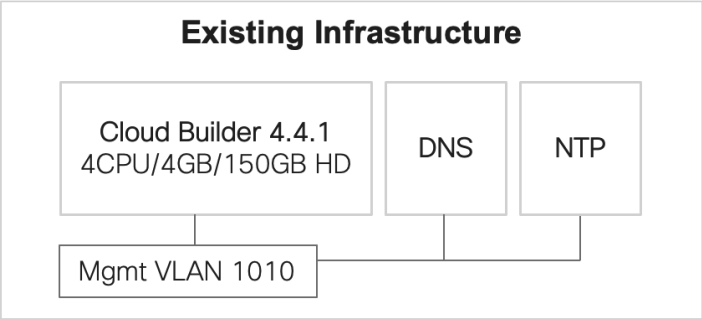

- The deployment environment should contain an NTP server for the ESXi hosts
- The deployment environment should have a DNS infrastructure and all following VM hostnames should be programmed in the DNS server (both forwards and reverse lookups):
  - VMware Cloud Builder VM
  - VMware SDDC Manager
  - VMware vCenter for management and VI domains
  - VMware NSX-T manager VMs and cluster VIPs for management and VI domains
  - All the ESXi hosts for management and VI domains
- The cloud builder VM is deployed in the customer existing environment. Customers should have a vSphere environment available to deploy the cloud builder VM OVF file.
- The management network where the VCF components are being deployed should be accessible within the enterprise

Deploying the Management Domain on vSAN Ready Nodes

VMware Cloud Foundation management domain is deployed using VMware cloud builder appliance. VMware Cloud Builder is a virtual appliance that is used to deploy and configure the first cluster of the management domain and transfer inventory and control to SDDC Manager. The management domain consists of four ESXi hosts at a minimum. During the deployment process, the VMware Cloud Builder appliance validates the provided network information in the deployment parameter workbook such as DNS, network (VLANS, IPs, MTUs), and credentials.

During the VMware Cloud Foundation setup definition process, deployment information specific to the environment such as networks, hosts, license keys, and other information is added to the deployment parameter workbook and uploaded to the VMware Cloud Builder appliance. During bring-up, the management domain is created on the ESXi hosts specified in the deployment parameter workbook. The VMware Cloud Foundation software components are automatically deployed, configured, and licensed using the information provided in the workbook.

Figure 37 shows various virtual machines and their deployment configuration in the management domain. Depending on the size of the deployment chosen in the deployment worksheet, the virtual machine size could be different. The workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for high availability.

Note: The deployment parameter workbook is downloaded from the VMware Cloud Builder appliance and the completed workbook is uploaded back to the VM. The specifics of the information contained and updated in the worksheet will be explained in the deployment guide.

Figure 37. VMware Cloud Foundation Management Workload Domain

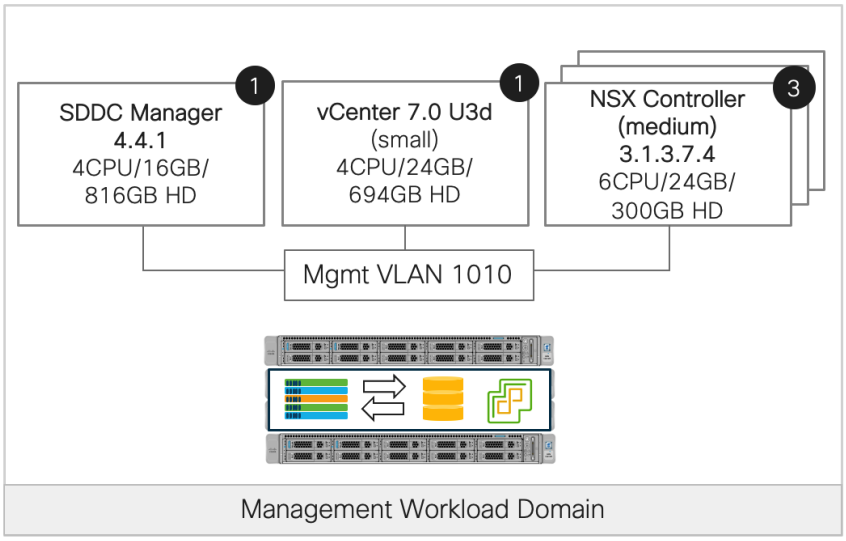

For more details on deployment of the management cluster using VMware cloud builder appliance, see: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.4/vcf-deploy/GUID-78EEF782-CF21-4228-97E0-37B8D2165B81.html

Management Domain ESXi Host Design

The vNICs created for the ESXi hosts using the Cisco Intersight server profile are assigned to specific distributed switches as covered below. Before VMware cloud builder starts configuring the ESXi host for management domain cluster, the ESXi host should only contain a single vSwitch using a single vNIC. The vNICs are assigned as follows:

- Two vNICs* (one on each fabric) are assigned to vSphere Virtual Distributed Switch VDS01 to support core services such as management, vSAN and vMotion traffic.
- Two vNICs (one on each fabric) for vSphere Virtual Distributed Switch (VDS02) to support NSX-T host overlay traffic.

Figure 38 illustrates the ESXi vNIC configurations in detail.

Figure 38. Management Domain ESXi Host Networking Configuration

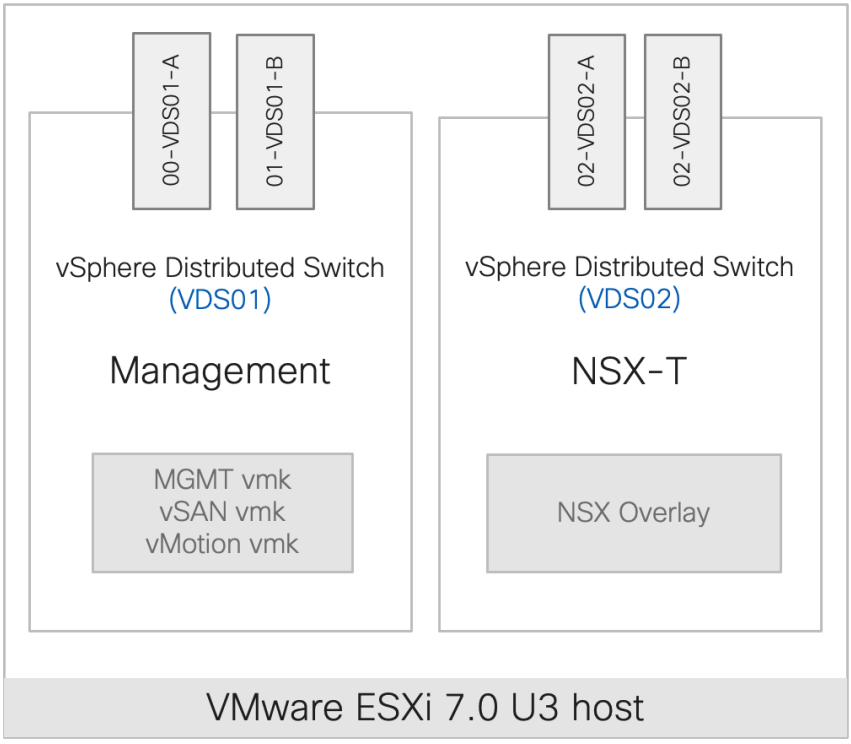

VMware SDDC Manager – Management Domain

On successful deployment of the management workload domain, customers can log into the VMware SDDC manager and find various deployment parameters, perform lifecycle management, and gather information about the vCenter and NSX manager. A VMware SDDC Manager GUI showing a VCF management domain is detailed in Figure 39, where customers can find an overview of the resource utilization and access details for various components in the management cluster.

Figure 39. SDDC Manager showing Management Workload Domain

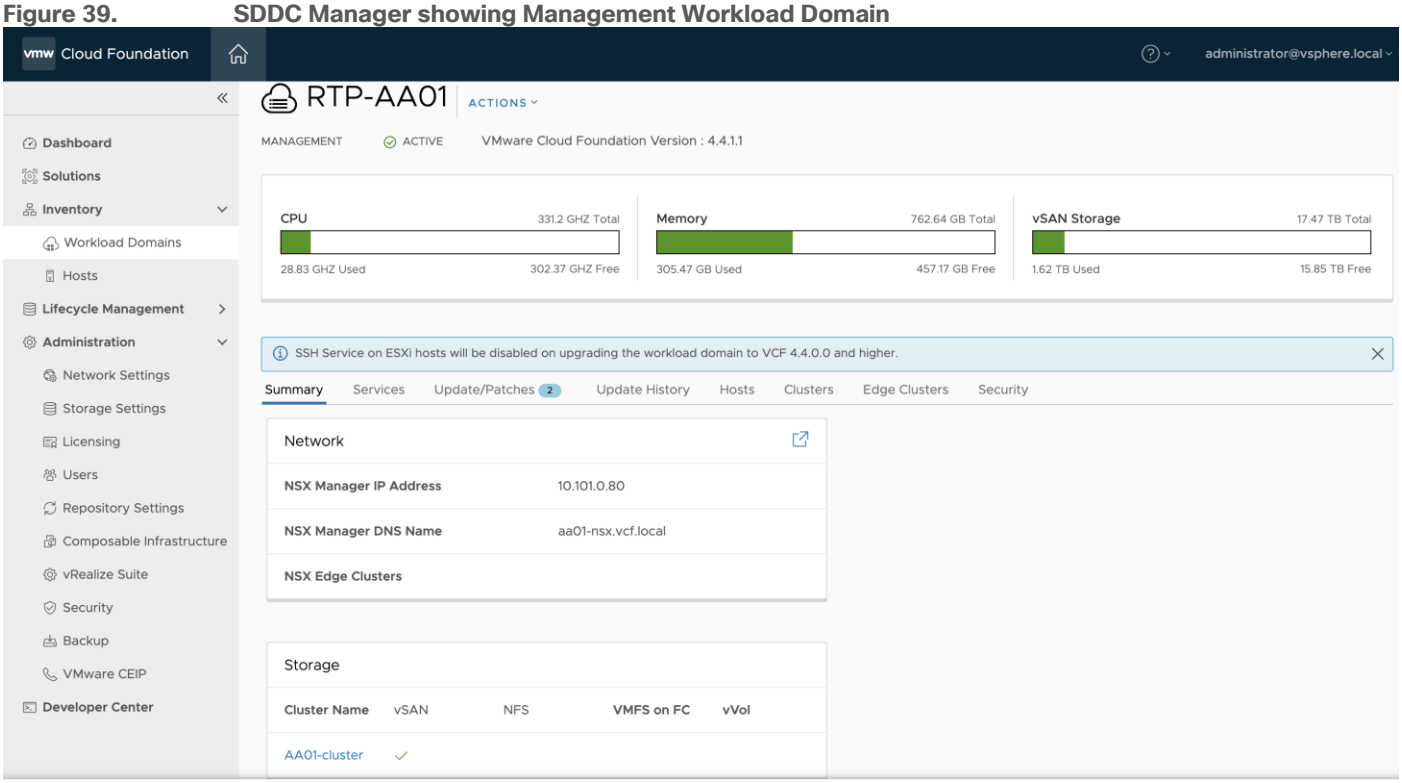

Deploy the VI Workload Domain on FlexPod Datacenter

VMware Cloud Foundation VI workload domain is deployed using VMware SDDC manager. When deploying a VI workload domain, the storage type, compute, and networking details are provided to the SDDC manager. Based on the selected storage (NFS for FlexPod), NFS share details are also provided at the time of deployment. This storage becomes the primary storage for the VI workload domain. The VI workload domain consists of three ESXi hosts at a minimum. As part of VI workload onboarding, VMware SDDC manager automatically:

- Deploys a vCenter Server Appliance for the new VI workload domain within the management domain. By using a separate vCenter Server instance per VI workload domain, software updates can be applied without impacting other VI workload domains. It also allows for each VI workload domain to have additional isolation as needed.
- Connects the specified ESXi servers to this vCenter Server instance and groups them into a cluster. Each host is configured with the port groups applicable for the VI workload domain.
- Configures networking on each host.
- Configures NFS storage on the ESXi hosts.

The VI workload onboarding comprises of two steps:

- Commissioning the FlexPod ESXi hosts in VMware SDDC Manager
- Creating a new VI workload domain using the newly commissioned hosts

Figure 40 shows various virtual machines and their deployment configuration in the VI workload domain. The VMware SDDC Manager workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for high availability.

Note: By default, VI workload domains do not include any NSX Edge clusters and are isolated. To provide north-south routing and network services, utilize the traditional VLAN based application deployment or add one or more NSX Edge clusters to a VI workload domain.

Figure 40. VMware Cloud Foundation VI Workload Domain

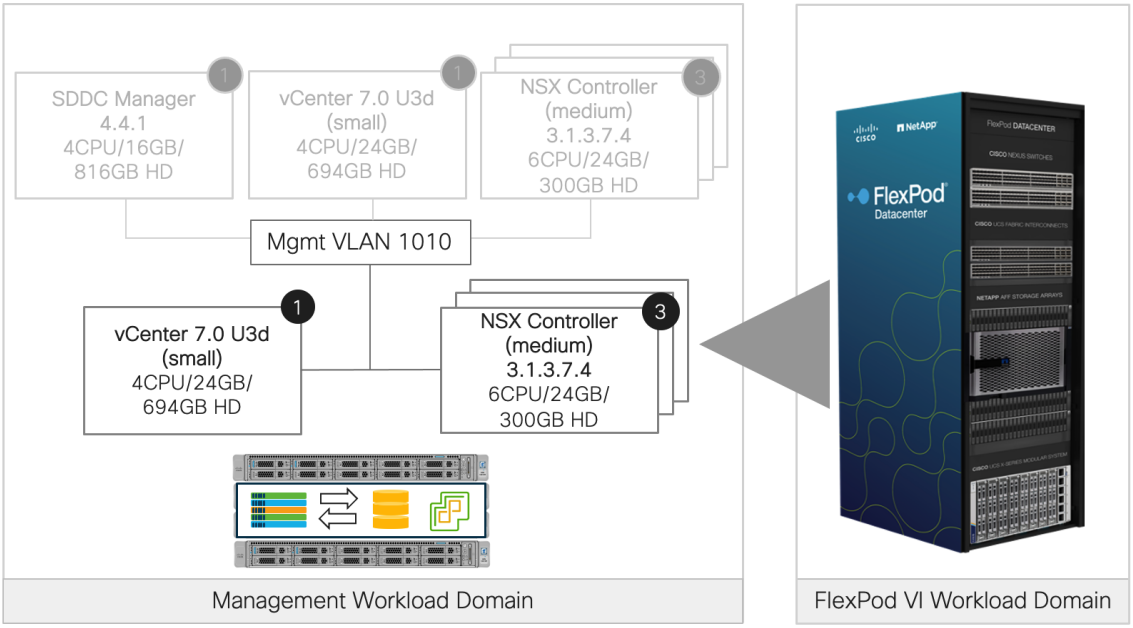

For the first VI workload domain, the workflow deploys a cluster of three NSX Managers in the management domain and configures a virtual IP (VIP) address for the NSX Manager cluster. The workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for high availability. Subsequent VI workload domains can share an existing NSX Manager cluster or deploy a new one.

Note: To share an NSX Manager cluster, the VI workload domains must use the same update method. The VI workload domains must both use vSphere Lifecycle Manager (vLCM) images, or they must both use vLCM baselines.

VI Workload Domain Creation using VMware Cloud Foundation API

VMware SDDC manager allows customers to create a new workload domain using the SDDC Manager web graphical user interface (GUI) or by creating a description file using JSON and using VMware Cloud Foundation API. The VI workload domain deployment using GUI is simpler however the GUI only supports creation of a single VDS in the ESXi host. The FlexPod ESXi hosts contain at least four vNICs and require creation of 2 VDSs so traffic can be segregated and controlled on the vNIC basis. This multi-VDS configuration is completed using VMware Cloud Foundation API. A JSON file is created with appropriate network parameters and definitions and pushed to VMware Cloud Foundation API. Using these parameters, VMware Cloud Foundation deploys two VDSs in the VI workload domain.

FlexPod VI Workload Domain ESXi Host Design

The vNICs and vHBAs created for the ESXi hosts using the Cisco Intersight server profile are assigned to specific distributed switches as covered below. Before VMware SDDC Manager starts commissioning the ESXi host for VI workload domain, the ESXi host should only contain a single vSwitch using a single vNIC. The vNICs are assigned as follows:

- Two vNICs* (one on each fabric) are assigned to vSphere Virtual Distributed Switch VDS01 to support core services such as management, vSAN and vMotion traffic.
- Two vNICs (one on each fabric) for vSphere Virtual Distributed Switch (VDS02) to support NSX-T host overlay traffic.
- One vHBA each for Fabric-A and Fabric-B for FC stateless boot.

Figure 41 shows the ESXi vNIC configurations in detail. vHBAs are not shown in this figure.

Figure 41. FlexPod VI Workload Domain ESXi Host Networking Configuration

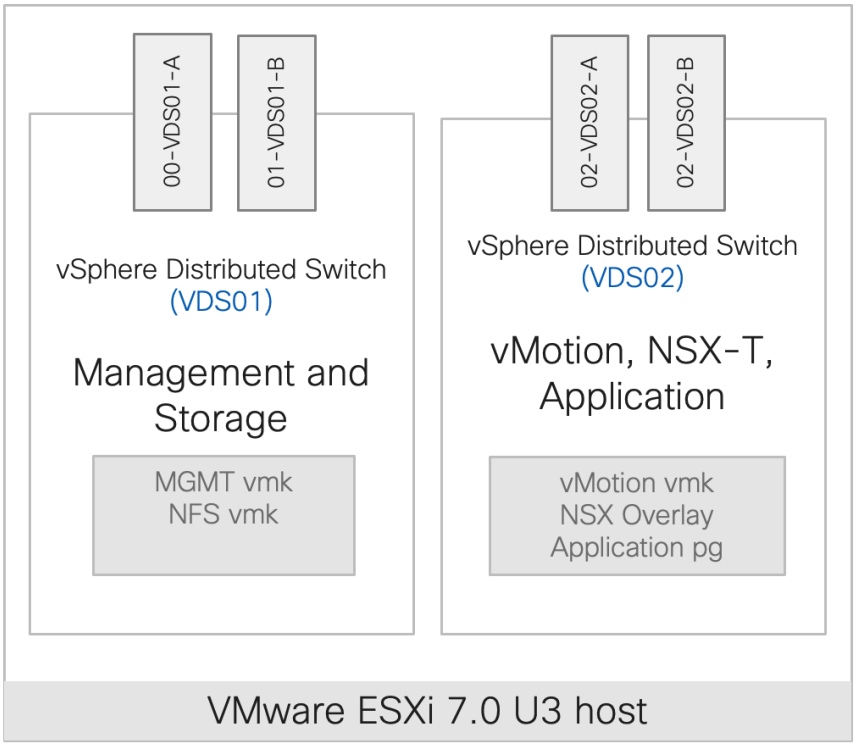

VMware SDDC Manager – FlexPod VI Workload Domain

On successful deployment of the VI workload domain, customers can log into the VMware SDDC manager and find various deployment parameters, perform lifecycle management, and gather information about the vCenter and NSX manager assigned to the newly created VI workload domain. A VMware SDDC Manager GUI showing a FlexPod VI workload domain is captured in Figure 42, where customers can find an overview of the resource utilization and access details for various components in the VI workload domain cluster.

Figure 42. SDDC Manager showing FlexPod VI Workload Domain

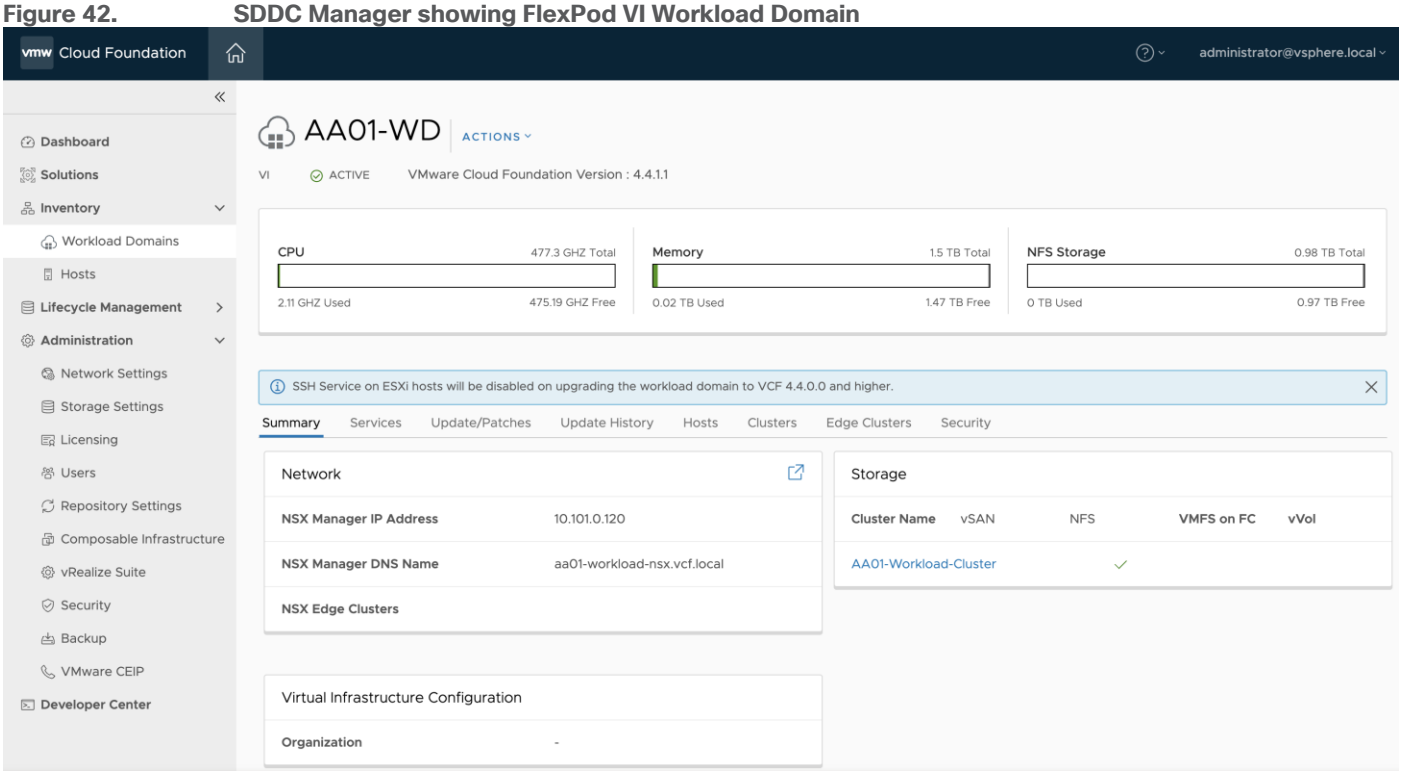

Design Considerations

Some of the key design considerations for the FlexPod Datacenter with Cisco UCS X-Series are explained in this section.

Cisco UCS C-Series vSAN Ready Node Management Options

Cisco UCS C-Series vSAN ready nodes that host the VMware Cloud Foundation management domain can be deployed and managed in one of the following two configurations:

Cisco UCS C-Series vSAN Ready Nodes Supported by IMM

For vSAN ready nodes that contain Intersight Managed Mode (IMM) supported components are connected to the same Cisco UCS FIs where FlexPod Cisco UCS X-Series chassis is connected as shown in Figure 43.

Figure 43. vSAN Ready Nodes with IMM Supported Components

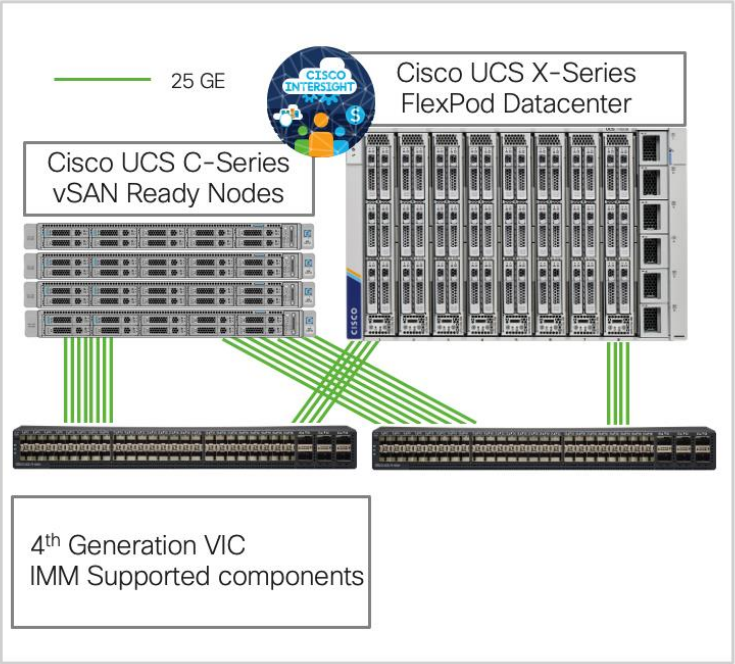

This design uses topology shown in Figure 43 to connect the vSAN ready nodes to shared Cisco UCS FIs.

Note: The IMM supported components for a server can be found here: https://intersight.com/help/saas/supported_systems

Cisco UCS C-Series vSAN Ready Nodes not Supported by IMM

Customers might own certain vSAN ready nodes or Cisco UCS C-Series systems that use VMware vSAN certified components (disks, raid controllers, and so on), that are not supported in IMM. One of the most common examples of the non-IMM supported configuration is the 3rd generation VIC. These C-Series servers cannot be connected to the same Cisco UCS FIs where FlexPod Cisco UCS X-Series chassis is connected. Customers can connect the servers to a separate pair of FIs which will be configured and managed using Cisco UCS Manager as shown in Figure 44.

Figure 44. vSAN Ready Nodes Without IMM Supported Components

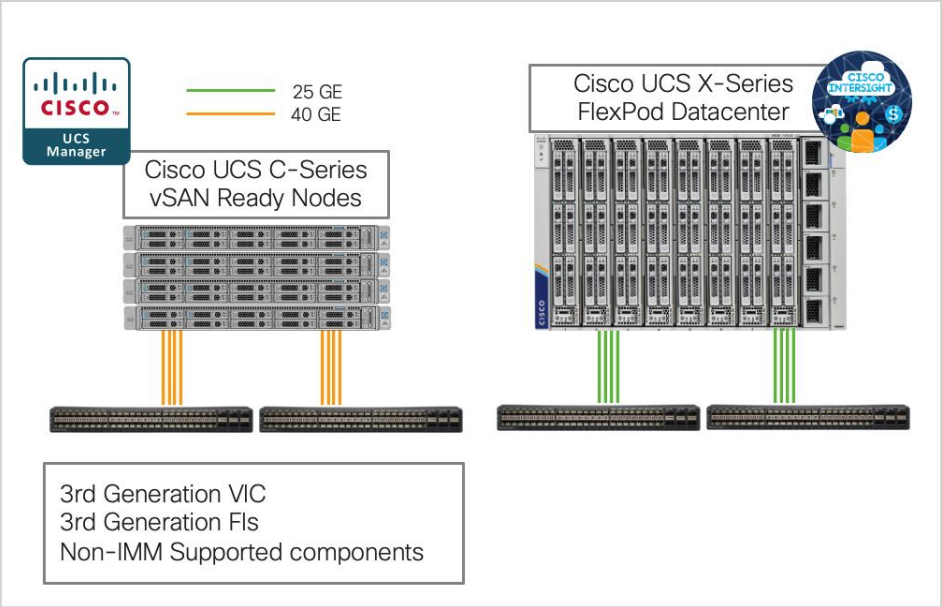

Note: The non-IMM supported Cisco UCS C-series vSAN ready nodes can also be connected directly to the Cisco Nexus switches and configured using CIMC. This option was not explored during the validation.

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-band Management Network

The in-band management VLAN configuration is part of FlexPod design. The in-band VLAN is configured on 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.

For the VMware Cloud Foundation component deployment, customers can use either:

- An existing enterprise management network or
- The FlexPod dedicated In-Band network

During this validation, VMware cloud foundation elements were deployed on an existing management segment (VLAN 1010; 10.101.0.0/24) which was different from the FlexPod In-Band management segment (VLAN 1011; 10.101.1.0/24) and routing was configured between the two networks. Customers can choose to use a single management VLAN/network for all the VMware Cloud Foundation components as well as ESXi hosts, and management tools deployed in the FlexPod environment.

Jumbo Frames

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

Boot From SAN

In FlexPod infrastructure, 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 X-Series Server Profiles for hardware flexibility across the server hardware and overall portability of server identity. Boot from SAN also removes the need to populate local server storage thereby reducing cost and administrative overhead.

UEFI Secure Boot

This validation of 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 by the BIOS before they can be loaded. Cisco UCS X210C compute nodes also contain a Trusted Platform Module (TPM). VMware ESXi 7.0 U3 supports UEFI Secure Boot and VMware vCenter 7.0 U3 supports UEFI Secure Boot Attestation between the TPM module and ESXi, validating that UEFI Secure Boot has properly taken place.

Solution Automation

In addition to command line interface (CLI) and graphical user interface (GUI) configurations, explained in the deployment guide, all FlexPod components support configurations through automation frameworks such as Ansible and Terraform etc. The FlexPod solution validation team develops automation modules to configure Cisco Nexus, Cisco UCS, Cisco MDS, NetApp ONTAP, NetApp ONTAP Tools for VMware, Active IQ Unified Manager and VMware ESXi (initial configuration). This community-supported GitHub repository is meant to expedite customer adoption of automation by providing them sample configuration playbooks that can be easily developed or integrated into existing customer automation frameworks. Another key benefit of the automation package is the reusability of the code and roles to help customers execute repeatable tasks within their environment.

To setup VMware Cloud Foundation management hosts using Ansible playbooks, following two GitHub repositories are available:

- Setup Cisco UCS C-Series servers configured in Intersight Managed Mode: https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/IMM-VCF-MgmtDomain
- Setup Cisco UCS C-Series servers configured in UCSM Managed Mode: https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/UCSM-VCF-MgmtDomain

Note: A comprehensive GitHub repository is being developed to setup all the components of the solution including switching and storage and will be added to the document when available.

Validate

A high-level overview of the FlexPod as a workload domain for VMware Cloud Foundation design validation is provided in this section. Solution validation covers various aspects of the converged infrastructure including compute, virtualization, network, and storage. The test scenarios are divided into four broad categories:

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

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

- Verification of correct and repeatable deployment of VMware Cloud Foundation VI workload domain
- Verification of features configured on various FlexPod components
- Powering off and rebooting redundant devices and removing redundant links to verify high availability
- Path MTU verification including both storage and virtual machine traffic
- Failure and recovery of vCenter and ESXi hosts in a cluster
- Failure and recovery of storage access paths across AFF nodes, MDS and Nexus switches, and fabric interconnects
- Load generations using IOMeter VMs hosted on FlexPod components and path verification

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

Conclusion

The FlexPod Datacenter solution is a validated approach for deploying Cisco and NetApp technologies and products for building shared private and public cloud infrastructure. VMware Cloud Foundation enables data center administrators to provision an application environment in a quick, repeatable, and automated manner. FlexPod as a workload domain for VMware Cloud Foundation provides following benefits in any data center environment:

- Integrated solution that supports entire VMware software defined stack
- Standardized architecture for quick, repeatable, error free deployments of FlexPod based workload domains
- Automated life cycle management to keep all the system components up to date
- Simplified cloud-based management of various FlexPod components
- Hybrid-cloud-ready, policy-driven modular design
- Highly available, flexible, and scalable FlexPod architecture
- Cooperative support model and Cisco Solution Support
- Easy to deploy, consume, and manage design which aligns with Cisco, NetApp and VMware best practices and compatibility requirements
- Support for component monitoring, solution automation and orchestration, and workload optimization

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 and this document highlights the design details of incorporating FlexPod as a workload domain for VMware Cloud Foundation.

About the Authors

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

Haseeb Niazi has over 23 years of experience at Cisco in the Datacenter, Enterprise and Service Provider Solutions and Technologies. As a member of various solution teams and Advanced Services, Haseeb has helped many enterprise and service provider customers evaluate and deploy a wide range of Cisco solutions. As a technical marking engineer at Cisco UCS Solutions group, Haseeb focuses on network, compute, virtualization, storage, and orchestration aspects of various Compute Stacks. Haseeb holds a master's degree in Computer Engineering from the University of Southern California and is a Cisco Certified Internetwork Expert (CCIE 7848).

Ruchika Lahoti, Technical Marketing Engineer, NetApp

Ruchika has more than five years of experience in the IT industry. She focuses on FlexPod hybrid cloud infrastructure solution design, implementation, validation, and automation. Ruchika holds a bachelor's degree in Computer Science.

Acknowledgements

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

- John George, Technical Marketing Engineer, Cisco Systems, Inc.
- Chris Dunk, Principal Engineer, Cisco Systems, Inc.
- Abhinav Singh, Senior Technical Marketing Engineer, NetApp

Appendix

This appendix is organized as follows:

Appendix A – References Used in Guide

Compute

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

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

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

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

Network

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

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

Storage

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

NetApp Active IQ Unified Manager: https://docs.netapp.com/ocum-98/index.jsp?topic=%2Fcom.netapp.doc.onc-um-isg-lin%2FGUID-FA7D1835-F32A-4A84-BD5A-993F7EE6BBAE.html

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

Virtualization

VMware Cloud Foundation 4.4 release notes: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.4.1/rn/vmware-cloud-foundation-441-release-notes/index.html

VMware Cloud Foundation 4.4 Deployment Guide: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.4/vcf-deploy/GUID-F2DCF1B2-4EF6-444E-80BA-8F529A6D0725.html

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/

Appendix B – Terms Glossary

This glossary addresses some terms used in this document, for the purposes of aiding understanding. This is not a complete list of all multicloud terminology. Some Cisco product links are supplied here also, where considered useful for the purposes of clarity, but this is by no means intended to be a complete list of all applicable Cisco products.

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

	power, cooling, and physical security for the server, storage, and networking equipment of other firms and also connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity.” https://en.wikipedia.org/wiki/Colocation_centre

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

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

Appendix C - Acronym Glossary

AAA—Authentication, Authorization, and Accounting

ACP—Access-Control Policy

ACI—Cisco Application Centric Infrastructure

ACK—Acknowledge or Acknowledgement

ACL—Access-Control List

AD—Microsoft Active Directory

AFI—Address Family Identifier

AMP—Cisco Advanced Malware Protection

AP—Access Point

API—Application Programming Interface

APIC— Cisco Application Policy Infrastructure Controller (ACI)

ASA—Cisco Adaptative Security Appliance

ASM—Any-Source Multicast (PIM)

ASR—Aggregation Services Router

Auto-RP—Cisco Automatic Rendezvous Point protocol (multicast)

AVC—Application Visibility and Control

BFD—Bidirectional Forwarding Detection

BGP—Border Gateway Protocol

BMS—Building Management System

BSR—Bootstrap Router (multicast)

BYOD—Bring Your Own Device

CAPWAP—Control and Provisioning of Wireless Access Points Protocol

CDP—Cisco Discovery Protocol

CEF—Cisco Express Forwarding

CMD—Cisco Meta Data

CPU—Central Processing Unit

CSR—Cloud Services Routers

CTA—Cognitive Threat Analytics

CUWN—Cisco Unified Wireless Network

CVD—Cisco Validated Design

CYOD—Choose Your Own Device

DC—Data Center

DHCP—Dynamic Host Configuration Protocol

DM—Dense-Mode (multicast)

DMVPN—Dynamic Multipoint Virtual Private Network

DMZ–Demilitarized Zone (firewall/networking construct)

DNA–Cisco Digital Network Architecture

DNS–Domain Name System

DORA–Discover, Offer, Request, ACK (DHCP Process)

DWDM–Dense Wavelength Division Multiplexing

ECMP–Equal Cost Multi Path

EID–Endpoint Identifier

EIGRP–Enhanced Interior Gateway Routing Protocol

EMI–Electromagnetic Interference

ETR–Egress Tunnel Router (LISP)

EVPN–Ethernet Virtual Private Network (BGP EVPN with VXLAN data plane)

FHR–First-Hop Router (multicast)

FHRP–First-Hop Redundancy Protocol

FMC–Cisco Firepower Management Center

FTD–Cisco Firepower Threat Defense

GBAC–Group-Based Access Control

GbE–Gigabit Ethernet

Gbit/s–Gigabits Per Second (interface/port speed reference)

GRE–Generic Routing Encapsulation

GRT–Global Routing Table

HA–High-Availability

HQ–Headquarters

HSRP–Cisco Hot-Standby Routing Protocol

HTDB–Host-tracking Database (SD-Access control plane node construct)

IBNS–Identity-Based Networking Services (IBNS 2.0 is the current version)

ICMP– Internet Control Message Protocol

IDF–Intermediate Distribution Frame; essentially a wiring closet.

IEEE–Institute of Electrical and Electronics Engineers

IETF–Internet Engineering Task Force

IGP–Interior Gateway Protocol

IID–Instance-ID (LISP)

IOE–Internet of Everything

IoT—Internet of Things

IP—Internet Protocol

IPAM—IP Address Management

IPS—Intrusion Prevention System

IPSec—Internet Protocol Security

ISE—Cisco Identity Services Engine

ISR—Integrated Services Router

IS-IS—Intermediate System to Intermediate System routing protocol

ITR—Ingress Tunnel Router (LISP)

LACP—Link Aggregation Control Protocol

LAG—Link Aggregation Group

LAN—Local Area Network

L2 VNI—Layer 2 Virtual Network Identifier; as used in SD-Access Fabric, a VLAN.

L3 VNI— Layer 3 Virtual Network Identifier; as used in SD-Access Fabric, a VRF.

LHR—Last-Hop Router (multicast)

LISP—Location Identifier Separation Protocol

MAC—Media Access Control Address (OSI Layer 2 Address)

MAN—Metro Area Network

MEC—Multichassis EtherChannel, sometimes referenced as ***MCEC***

MDF—Main Distribution Frame; essentially the central wiring point of the network.

MnT—Monitoring and Troubleshooting Node (Cisco ISE persona)

MOH—Music on Hold

MPLS—Multiprotocol Label Switching

MR—Map-resolver (LISP)

MS—Map-server (LISP)

MSDP—Multicast Source Discovery Protocol (multicast)

MTU—Maximum Transmission Unit

NAC—Network Access Control

NAD—Network Access Device

NAT—Network Address Translation

NBAR—Cisco Network-Based Application Recognition (NBAR2 is the current version).

NFV—Network Functions Virtualization

NSF–Non-Stop Forwarding

OSI–Open Systems Interconnection model

OSPF–Open Shortest Path First routing protocol

OT–Operational Technology

PAgP–Port Aggregation Protocol

PAN–Primary Administration Node (Cisco ISE persona)

PCI DSS–Payment Card Industry Data Security Standard

PD–Powered Devices (PoE)

PETR–Proxy-Egress Tunnel Router (LISP)

PIM–Protocol-Independent Multicast

PITR–Proxy-Ingress Tunnel Router (LISP)

PnP–Plug-n-Play

PoE–Power over Ethernet (Generic term, may also refer to IEEE 802.3af, 15.4W at PSE)

PoE+–Power over Ethernet Plus (IEEE 802.3at, 30W at PSE)

PSE–Power Sourcing Equipment (PoE)

PSN–Policy Service Node (Cisco ISE persona)

pxGrid–Platform Exchange Grid (Cisco ISE persona and publisher/subscriber service)

PxTR–Proxy-Tunnel Router (LISP – device operating as both a PETR and PITR)

QoS–Quality of Service

RADIUS–Remote Authentication Dial-In User Service

REST–Representational State Transfer

RFC–Request for Comments Document (IETF)

RIB–Routing Information Base

RLOC–Routing Locator (LISP)

RP–Rendezvous Point (multicast)

RP–Redundancy Port (WLC)

RP–Route Processer

RPF–Reverse Path Forwarding

RR–Route Reflector (BGP)

RTT–Round-Trip Time

SA–Source Active (multicast)

SAFI–Subsequent Address Family Identifiers (BGP)

SD–Software-Defined

SDA–Cisco Software Defined-Access

SDN–Software-Defined Networking

SFP–Small Form-Factor Pluggable (1 GbE transceiver)

SFP+– Small Form-Factor Pluggable (10 GbE transceiver)

SGACL–Security-Group ACL

SGT–Scalable Group Tag, sometimes reference as Security Group Tag

SM–Spare-mode (multicast)

SNMP–Simple Network Management Protocol

SSID–Service Set Identifier (wireless)

SSM–Source-Specific Multicast (PIM)

SSO–Stateful Switchover

STP–Spanning-tree protocol

SVI–Switched Virtual Interface

SVL–Cisco StackWise Virtual

SWIM–Software Image Management

SXP–Scalable Group Tag Exchange Protocol

Syslog–System Logging Protocol

TACACS+–Terminal Access Controller Access-Control System Plus

TCP–Transmission Control Protocol (OSI Layer 4)

UCS– Cisco Unified Computing System

UDP–User Datagram Protocol (OSI Layer 4)

UPoE–Cisco Universal Power Over Ethernet (60W at PSE)

UPoE+– Cisco Universal Power Over Ethernet Plus (90W at PSE)

URL–Uniform Resource Locator

VCF–VMware Cloud Foundation

vHBA–virtual Host Bus Adapter

VLAN–Virtual Local Area Network

VM—Virtual Machine

VN–Virtual Network, analogous to a VRF in SD-Access

VNI–Virtual Network Identifier (VXLAN)

vNIC–virtual Network Interface Card

vPC—virtual Port Channel (Cisco Nexus)

VPLS—Virtual Private LAN Service

VPN—Virtual Private Network

VPNv4—BGP address family that consists of a Route-Distinguisher (RD) prepended to an IPv4 prefix

VPWS—Virtual Private Wire Service

VRF—Virtual Routing and Forwarding

VSL—Virtual Switch Link (Cisco VSS component)

VSS—Cisco Virtual Switching System

VXLAN—Virtual Extensible LAN

WAN—Wide-Area Network

WLAN—Wireless Local Area Network (generally synonymous with IEEE 802.11-based networks)

WoL—Wake-on-LAN

xTR—Tunnel Router (LISP – device operating as both an ETR and ITR)

Appendix D – Recommended for You

FlexPod Datacenter with Cisco UCS X-Series Design Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_esxi7u2_design.html

FlexPod Datacenter with UCS X-Series Design Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_vmware_7u2.html

FlexPod Datacenter with End-to-End 100G Design Guide:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_ucs_xseries_e2e_ontap_design.html

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**.