FlexPod Datacenter with Red Hat OpenShift Container Platform and NetApp Astra

Design Guide

Published Date: December 2022

CISCO VALIDATED DESIGN

FlexPod®

In partnership with:

NetApp

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 by its ability to evolve and incorporate both technology and product innovations in the areas of management, computing, storage, and networking. To help organizations with their digital transformation and application modernization practices, Cisco and NetApp have partnered to produce this Cisco Validated Design (CVD) for the FlexPod™ Datacenter for Red Hat OpenShift Container Platform solution. As the hybrid-cloud operation is the new new de-facto default for many companies, the network connection to the public cloud, the Kubernetes cluster management, and the workload management across on-premises and public clouds are covered as part of this solution design.

FlexPod delivers an integrated architecture that incorporates compute, storage, and network design best practices, thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance, and support that can be used in various stages (planning, designing, and implementation) of a deployment. FlexPod delivered as IaC further eliminates error-prone manual tasks, allowing quicker and more consistent solution deployments.

Red Hat® OpenShift® Container Platform (OCP) is an enterprise-ready Kubernetes container platform with full-stack automated operations to manage hybrid cloud and multi-cloud deployments. Red Hat OCP is optimized to improve developer productivity and promote innovation. The Red Hat OCP gives developers a self-service platform on which to build and run containerized applications. With Red Hat OCP, you can quickly start creating new cloud-native applications or cloud-enabling existing applications and spawning an environment for a new microservice in minutes.

Combining Red Hat OCP with the FlexPod solution can simplify the deployment and management of the container infrastructure. The Red Hat Ansible integration with the FlexPod solution automates the deployment of the FlexPod infrastructure along with Red Hat OCP installation enabling customers to take advantage of programming and automating the infrastructure at scale with agility, extending the benefits of automation to the entire stack.

Some of the key advantages of integrating Cisco FlexPod Datacenter as a workload domain into Red Hat OCP 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 physical or 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 solution includes integration of the Cisco Intersight with NetApp Active IQ Unified Manager and, if required, VMware vCenter 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: https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html
.

Solution Overview – FlexPod Datacenter with Red Hat OpenShift Container Platform

This chapter contains the following:

Introduction

The featured FlexPod Datacenter for RedHat OpenShift Container platform solution is a pre-designed, integrated, and validated architecture for the data center that combines Cisco UCS servers, the Cisco Nexus/MDS family of switches, and NetApp AFF A-series storage into a single, flexible architecture. FlexPod is designed for high availability (HA), with no single point of failure, while maintaining cost-effectiveness and flexibility in the design to support a wide variety of workloads. The Red Hat OCP software provides multiple ways of deployment: Assisted Installer, Installer Provisioned Infrastructure (IPI), and User Provisioned Infrastructure (UPI). The FlexPod solution with Red Hat OCP was tested, validated with different options, and captured in the related CVD deployment guides and white papers. A list of available deployment documents is available in a later section of this design document.

Integration between the OpenShift Container Platform and the storage and data management services occurs at several levels, all of which are captured in this document. The main storage integration is based on Container Storage Interface (CSI) Astra Trident for Kubernetes Driver for NetApp storage systems, which enables container orchestrators such as Kubernetes to manage the life cycle of persistent storage.

The main focus of the FlexPod Datacenter with Red Hat OpenShift Container Platform solution is on a bare metal cluster with the OCP nodes running Red Hat Enterprise Linux CoreOS (RHCOS) on Cisco UCS servers. To better support smaller deployments and Test/Dev installations, the deployment of virtualized OCP cluster and single node deployments is validated as well.

The following design and deployment aspects of the FlexPod for OCP solution are explained in this document:

- FlexPod converged infrastructure
- Red Hat OpenShift Container Platform 4.x
- Red Hat OpenShift Virtualization
- CSI Astra Trident for Kubernetes
- Astra Control Center

The document also covers key design aspects based on the validated environment and best practices.

Audience

The intended audience of this document includes but is not limited to IT architects, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides design guidance around incorporating the Cisco Intersight—managed Cisco UCS platform within FlexPod Datacenter infrastructure to run Red Hat OpenShift Container Platform (OCP). The document introduces various design elements and covers various considerations and best practices for a successful deployment. The document also highlights the design and product requirements for integrating virtualization and storage systems to Cisco Intersight to deliver a true cloud-based integrated approach to infrastructure management.

What's New in this Release?

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

- Cisco Intersight Infrastructure Manager
- NetApp ONTAP 9.11.1
- Cisco UCS X-Series Modular
- Red Hat OpenShift Container Platform 4.10
- Integration with the FlexPod Integrated System in Cisco Intersight
- NetApp Astra Control Center

Solution Summary

These components are integrated and validated, and – where possible - the entire stack is automated 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 ground up.

The FlexPod Datacenter for Red Hat OpenShift Container Platform 4 solution offers the following key customer benefits:

- Integrated solution that supports the entire Red Hat 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 that aligns with Cisco, NetApp, and Red Hat best practices and compatibility requirements
- Support for component monitoring, solution automation and orchestration, and workload optimization

Figure 1. **FlexPod Datacenter for Red Hat OpenShift Container Platform Solution Stack**

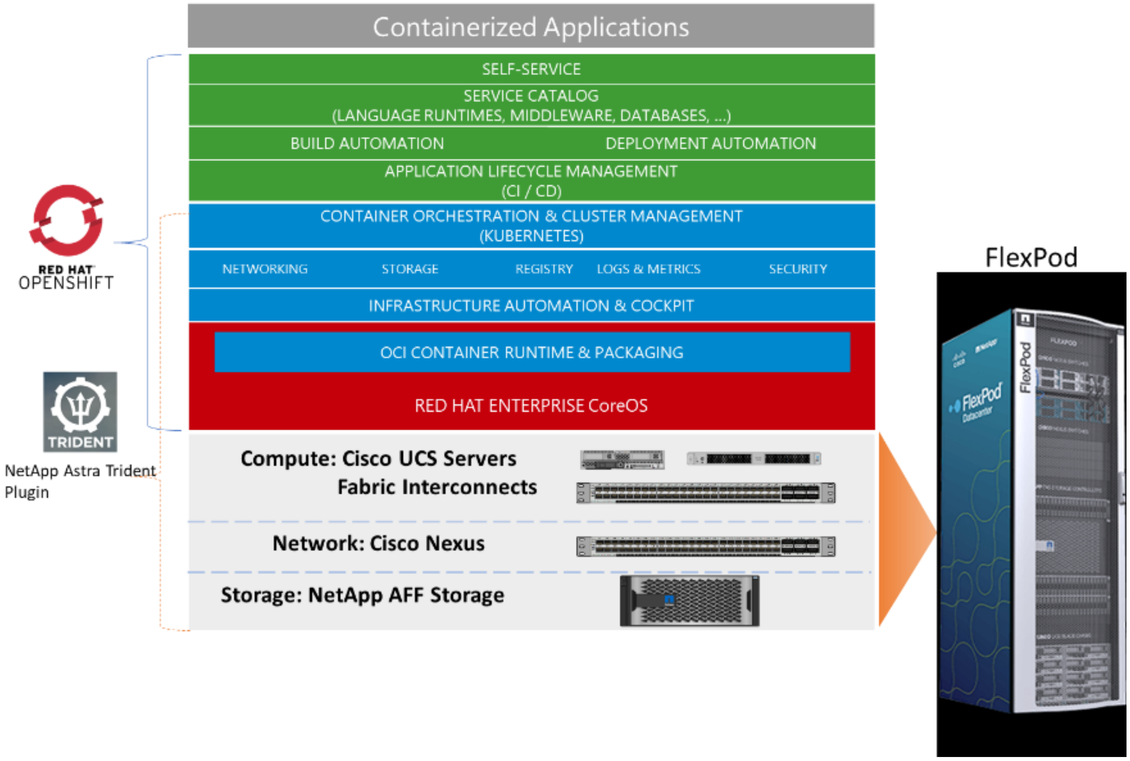
Containerized Applications
SELF-SERVICE
SERVICE CATALOG
(LANGUAGE RUNTIMES, MIDDLEWARE, DATABASES, …)
BUILD AUTOMATION
DEPLOYMENT AUTOMATION
APPLICATION LIFECYCLE MANAGEMENT
(CI / CD)
CONTAINER ORCHESTRATION & CLUSTER MANAGEMENT
(KUBERNETES)
NETWORKING
STORAGE
REGISTRY
LOGS & METRICS
SECURITY
INFRASTRUCTURE AUTOMATION & COCKPIT
OCI CONTAINER RUNTIME & PACKAGING
RED HAT ENTERPRISE CoreOS
RED HAT OPENSHIFT
TRIDENT
NetApp Astra Trident Plugin
Compute: Cisco UCS Servers
Fabric Interconnects
Network: Cisco Nexus
Storage: NetApp AFF Storage
FlexPod

Like all other FlexPod solution designs, FlexPod Datacenter for Red Hat OpenShift Container Platform 4 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. The FlexPod solution can also be deployed as a multi-site solution (stretched cluster or metro cluster) to address extended high-availability options and connected to the public cloud using NetApp Cloud Storage options for simplified data management across multiple locations.

Technology Overview

This chapter contains the following:

FlexPod Datacenter

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

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus® and Cisco MDS switches
- NetApp All Flash FAS (AFF) storage systems

Figure 2. FlexPod Datacenter Components

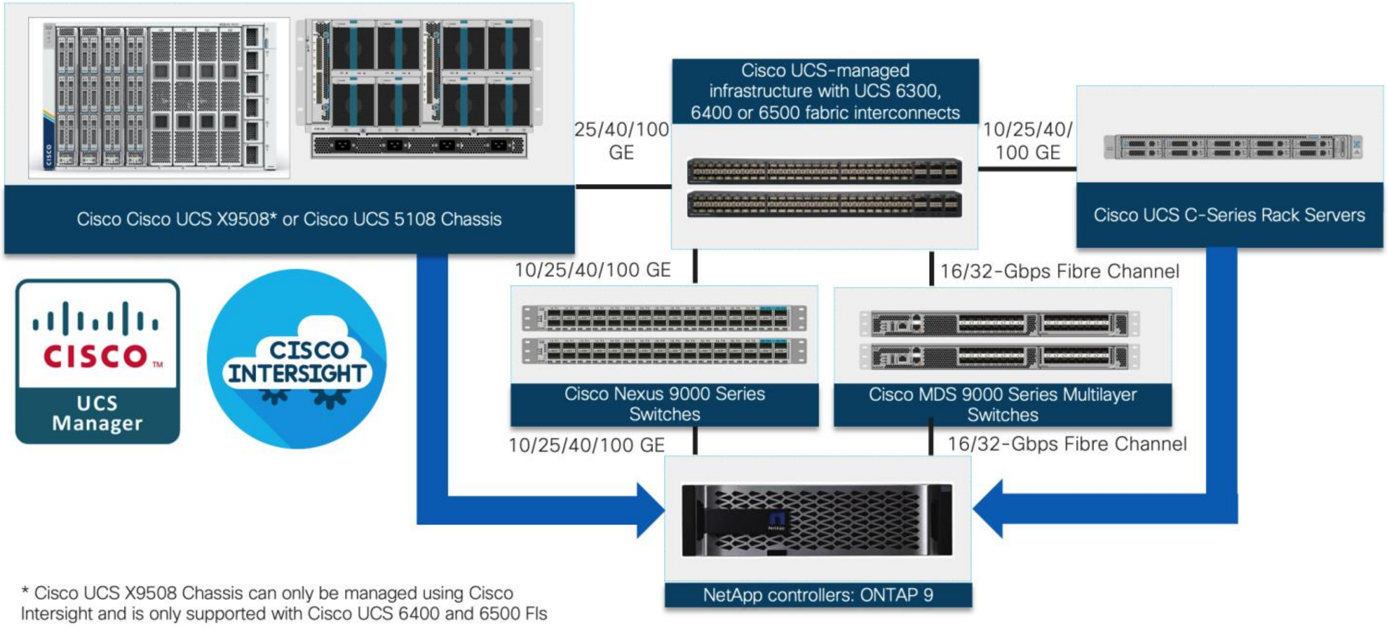
Cisco UCS-managed infrastructure with UCS 6300, 6400 or 6500 fabric interconnects
25/40/100 GE
10/25/40/ 100 GE
Cisco Cisco UCS X9508* or Cisco UCS 5108 Chassis
Cisco UCS C-Series Rack Servers
10/25/40/100 GE
16/32-Gbps Fibre Channel
CISCO UCS Manager
CISCO INTERSIGHT
Cisco Nexus 9000 Series Switches
Cisco MDS 9000 Series Multilayer Switches
10/25/40/100 GE
16/32-Gbps Fibre Channel
NetApp controllers: ONTAP 9

* Cisco UCS X9508 Chassis can only be managed using Cisco Intersight and is only supported with Cisco UCS 6400 and 6500 FIs

These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed), or it can scale out for environments that require multiple consistent deployments (such as rolling out of additional FlexPod stacks).

One of the key benefits of FlexPod is its ability to maintain consistency during scale. Each of the component families shown (Cisco UCS, Cisco Nexus, Cisco MDS, and NetApp AFF) offers platform and resource options to scale the infrastructure up or down while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod. The key features and highlights of the FlexPod components are explained below.

Cisco Unified Computing System

Cisco Unified Computing System™ (Cisco UCS®) is an integrated computing infrastructure with intent-based management to automate and accelerate deployment of all your applications, including virtualization and cloud computing, scale-out and bare-metal workloads, and in-memory analytics, as well as edge computing that supports remote and branch locations and massive amounts of data from the Internet of Things (IoT). The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

Cisco UCS Management

While Cisco UCS is a stateless, programmable infrastructure, the Cisco UCS unified API is how management tools program it. This enables the tools to help guarantee consistent, error-free, policy-based alignment of server personalities with workloads. Through automation, transforming the server and networking components of your infrastructure into a complete solution is fast and error-free because programmability eliminates the error-prone manual configuration of servers and integration into solutions. Server, network, and storage administrators are now free to focus on strategic initiatives rather than spending their time performing tedious tasks.

With Cisco Unified Computing System, Cisco introduced the Cisco UCS Manager to manage the system. Over the last few years, the way companies run and operate their data center and operate it has changed so has the management option for Cisco UCS.   The next-generation Cisco UCS management is called Cisco Intersight and is available in an as-a-Service model.

Cisco UCS Manager is only valid to manage the Cisco UCS B-Series blade server and C-Series rack server up to M6 generation. Cisco Intersight is the go-to manager for Cisco UCS X-Series blades and for Cisco UCS B-Series Blade server and C-Series Rack server from M5 generation forward.

Cisco UCS Manager

Cisco UCS® Manager (UCSM) provides unified, integrated management for all software and hardware components in Cisco UCS manages a single domain through an intuitive HTML 5-based GUI. The Cisco UCS Manager software is embedded in each fabric interconnect. Running in a redundant, high-availability configuration, it creates a single, self-aware, self-integrating unified system that recognizes and integrates components as they are added to the system. It quickly and accurately configures computing, network, storage, and storage-access resources to reduce the chance of errors that can cause downtime. Its role and policy-based approach help organizations more easily align policies and configurations with workloads. While Cisco UCS Manager requires an "always on" connection, our other tools are evolving to manage systems to which they are not continuously connected.

For more information about the Cisco UCS Manager, see https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/data_sheet_c78-520522.html.

Figure 3. **Cisco UCS System Manager - Server Overview**

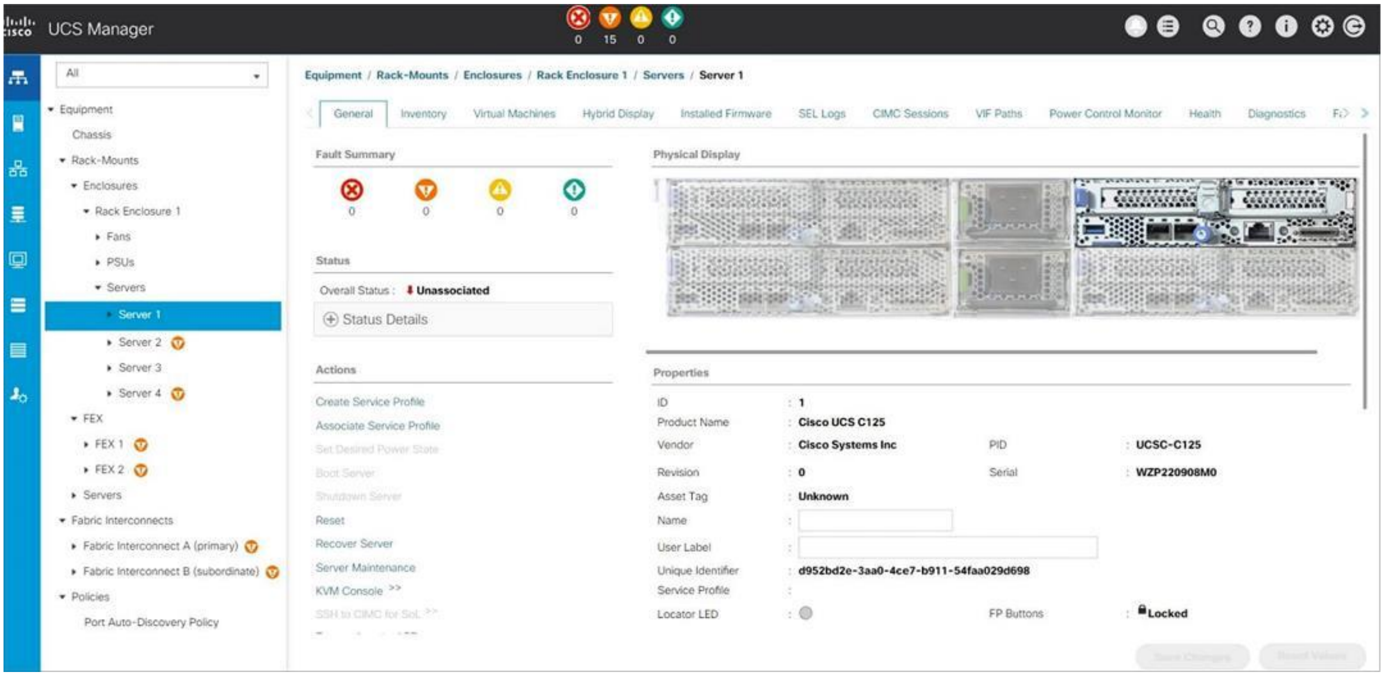

Overall Status : Unassociated

ID : 1
Product Name : Cisco UCS C125
Vendor : Cisco Systems Inc   PID : UCSC-C125
Revision : 0   Serial : WZP220908M0
Asset Tag : Unknown
Name :
User Label :
Unique Identifier : d952bd2e-3aa0-4ce7-b911-54faa029d698
Service Profile :
Locator LED :   FP Buttons : Locked

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 4. Cisco Intersight Overview

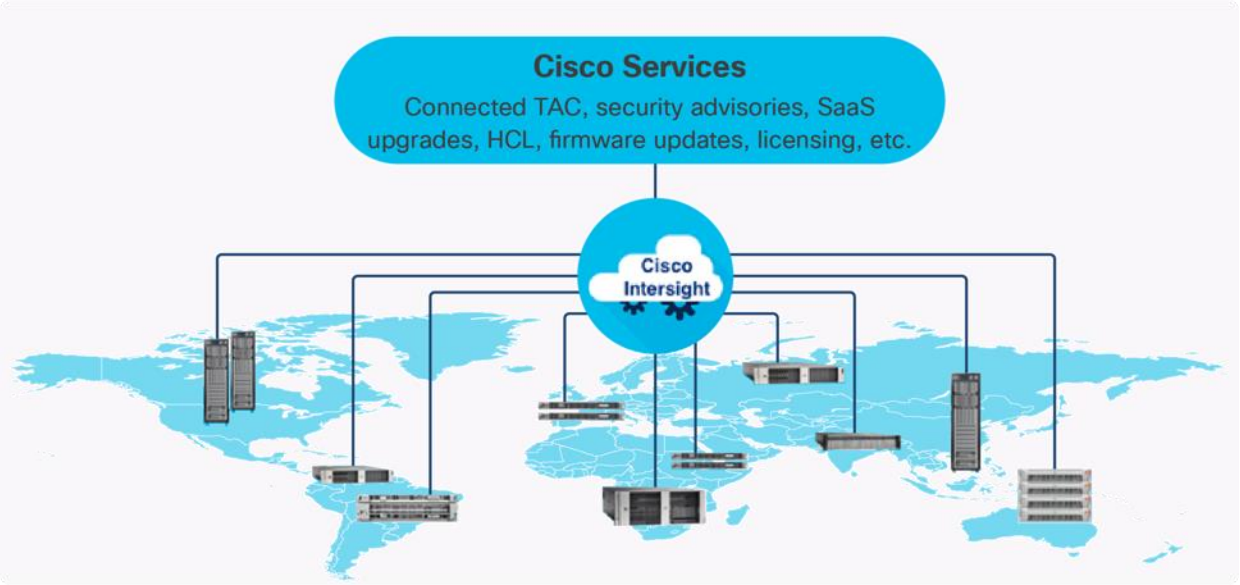
Cisco Services
Connected TAC, security advisories, SaaS upgrades, HCL, firmware updates, licensing, etc.
Cisco Intersight

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.

Figure 5. Cisco Intersight Dashboard

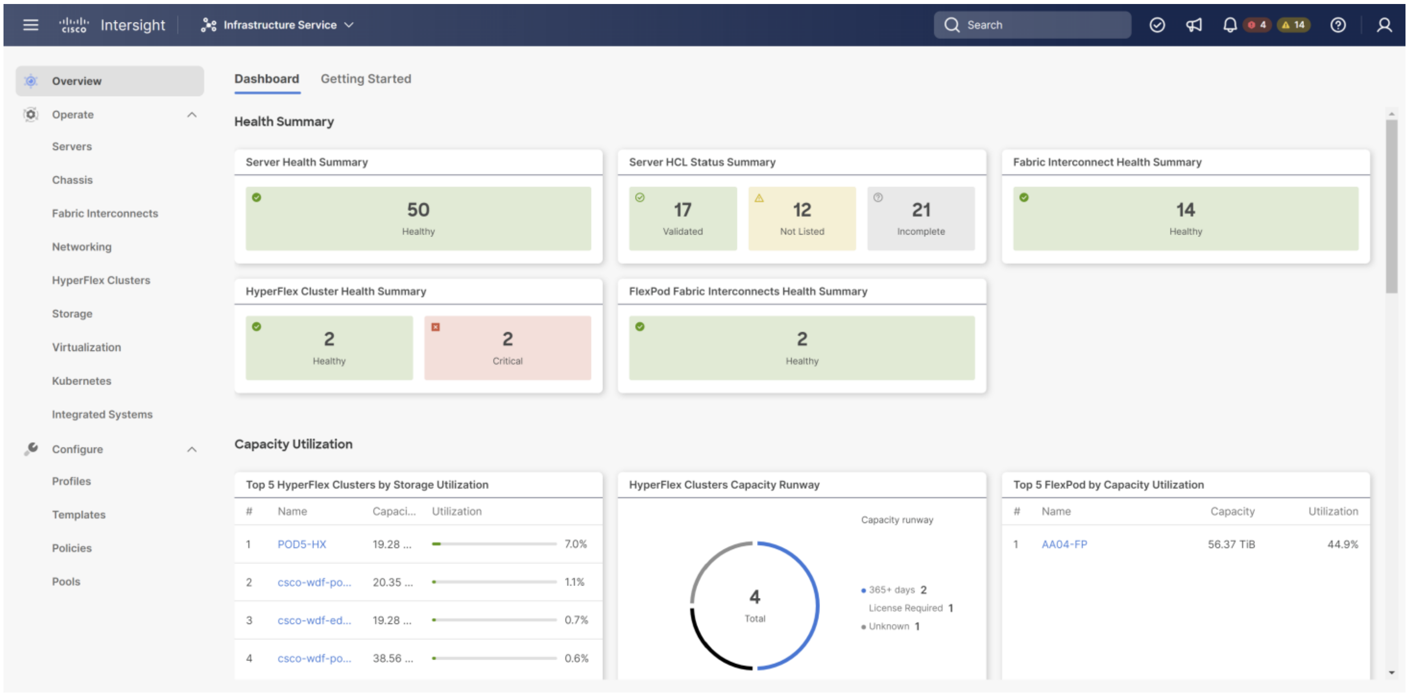

Top 5 HyperFlex Clusters by Storage Utilization

#	Name	Capaci...	Utilization
1	POD5-HX	19.28 ...	7.0%
2	csco-wdf-po...	20.35 ...	1.1%
3	csco-wdf-ed...	19.28 ...	0.7%
4	csco-wdf-po...	38.56 ...	0.6%

Top 5 FlexPod by Capacity Utilization

#	Name	Capacity	Utilization
1	AA04-FP	56.37 TiB	44.9%

Cisco Intersight Integration with NetApp ONTAP Storage

Using NetApp Active IQ Unified Manager (AIQUM) and Cisco Intersight Assist, NetApp ONTAP storage controllers can now be shown in Cisco Intersight with general and inventory information. NetApp AQIUM is an OVA based VMware virtual machine that can monitor multiple NetApp ONTAP storage clusters and also provides an API gateway to those storage clusters where the individual storage cluster credentials are managed by AIQUM, and all authentications can be handled with just AIQUM's credentials. When AIQUM is claimed by Cisco Intersight through the Intersight Assist appliance, all NetApp ONTAP storage clusters configured in AIQUM are pulled into Intersight. If you have installed the Intersight Advantage or Premier license, you can view this general and target inventory information, such as nodes, storage virtual machines, aggregates, disks, volumes, LUNs, initiator groups, network ports and network interfaces. With Premier License you can also execute NetApp ONTAP Storage tasks as workflows. The Virtualization and NetApp Storage tasks can be combined and executed as a single workflow.

DevOps and Tool Support

The Cisco UCS unified API is of great benefit to developers and administrators who want to treat physical infrastructure the way they treat other application services, using processes that automatically provision or change IT resources. Similarly, your IT staff needs to provision, configure, and monitor physical and virtual resources; automate routine activities; and rapidly isolate and resolve problems. The Cisco UCS unified API integrates with DevOps management tools and processes and enables you to easily adopt DevOps methodologies.

Cisco UCS 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 do-main controlled by the Cisco UCS Manager. 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.

The Cisco UCS Fabric Interconnects provides the management and communication backbone for the Cisco UCS B-Series Blade Servers in the Cisco UCS 5108 B-Series Server Chassis, the Cisco UCS X-Series Blade Servers in the Cisco UCS 9508 X-Series Server Chassis, and Cisco UCS Managed C-Series Rack Servers. All servers attached to the Cisco UCS Fabric Interconnect become part of a single, highly available management domain. In addition, by supporting a unified fabric the Cisco UCS Fabric Interconnect provides both LAN and SAN connectivity for all servers within its domain.

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 6. Cisco UCS 6454 Fabric Interconnect

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

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

For more information about the Cisco UCS 6400 series Fabric Interconnect see: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html.

Cisco UCS 6500 series Fabric Interconnects

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 7. Cisco UCS 6536 Fabric Interconnect

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

For more information about the Cisco UCS 6500 series Fabric Interconnect see: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs6536-fabric-interconnect-ds.html.

Cisco UCS Virtual Interface Cards (VICs)

The Cisco UCS Virtual Interface Card (VIC) extends the network fabric directly to both servers and virtual machines so that a single connectivity mechanism can be used to connect both physical and virtual servers with the same level of visibility and control. Cisco VICs provide complete programmability of the Cisco UCS I/O infrastructure, with the number and type of I/O interfaces configurable on demand with a zero-touch model.

Cisco VICs support Cisco SingleConnect technology, which provides an easy, intelligent, and efficient way to connect and manage computing in your data center. Cisco SingleConnect unifies LAN, SAN, and systems management into one simplified link for rack servers and blade servers. This technology reduces the number of network adapters, cables, and switches needed and radically simplifies the network, reducing complexity.

Cisco UCS 1400/14000 Series Virtual Interface Cards (VICs)

The Cisco UCS VIC 1400/14000 series is designed for Cisco UCS B-Series and Cisco UCS X-Series M5 and M6 Blade Servers, C-Series M5 and M6 Rack Servers, and S-Series M5 Storage Servers. The adapters are capable of supporting 10/25/40/100-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). It incorporates Cisco's next-generation Converged Network Adapter (CNA) technology and offers a comprehensive feature set, providing investment protection for future feature software releases. In addition, the VIC supports Cisco's Data Center Virtual Machine Fabric Extender (VM-FEX) technology. This technology extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

Cisco VIC1400/14000 can support 256 Express (PCIe) virtual devices, either virtual Network Interface Cards (vNICs) or virtual Host Bus Adapters (vHBAs), with a high rate of I/O Operations Per Second (IOPS), support for lossless Ethernet, and 10/25/40/100-Gbps connection to servers. The PCIe Generation 3 x16 interface helps ensure optimal bandwidth to the host for network-intensive applications, with a redundant path to the fabric interconnects. Cisco VICs support NIC teaming with fabric failover for increased reliability and availability. In addition, it provides a policy-based, stateless, agile server infrastructure for your data center.

Cisco UCS VIC 1480

The Cisco UCS VIC 1480 is a single-port 40-Gbps or 4x10-Gbps Ethernet/FCoE capable mezzanine card (Mezz) designed exclusively for the M5/M6 generation of Cisco UCS B-Series Blade Servers. The card enables a policy-based, stateless, agile server infrastructure that can present PCIe standards-compliant interfaces to the host that can be dynamically configured as either NICs or HBAs.

Cisco UCS VIC 14425

The Cisco UCS VIC 14425 is a 4x25-Gbps Ethernet/FCoE capable modular L AN On Motherboard (mLOM) designed exclusively for Cisco UCS X210 Compute Node. The Cisco UCS VIC 14425 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.

Figure 8. Cisco UCS VIC High-Level Configuration option with VIC14225 or VIC14825

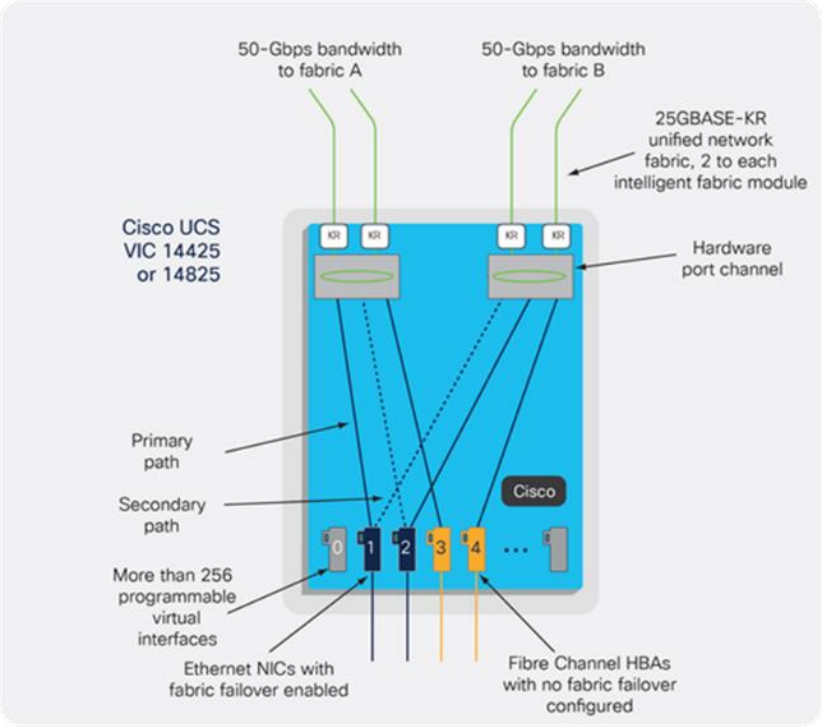
50-Gbps bandwidth to fabric A
50-Gbps bandwidth to fabric B
25GBASE-KR unified network fabric, 2 to each intelligent fabric module
Cisco UCS VIC 14425 or 14825
KR
KR
KR
KR
Hardware port channel
Primary path
Secondary path
Cisco
0
1
2
3
4
...
More than 256 programmable virtual interfaces
Ethernet NICs with fabric failover enabled
Fibre Channel HBAs with no fabric failover configured

Figure 9. Single Cisco UCS VIC 14425 in Cisco UCS X210c M6

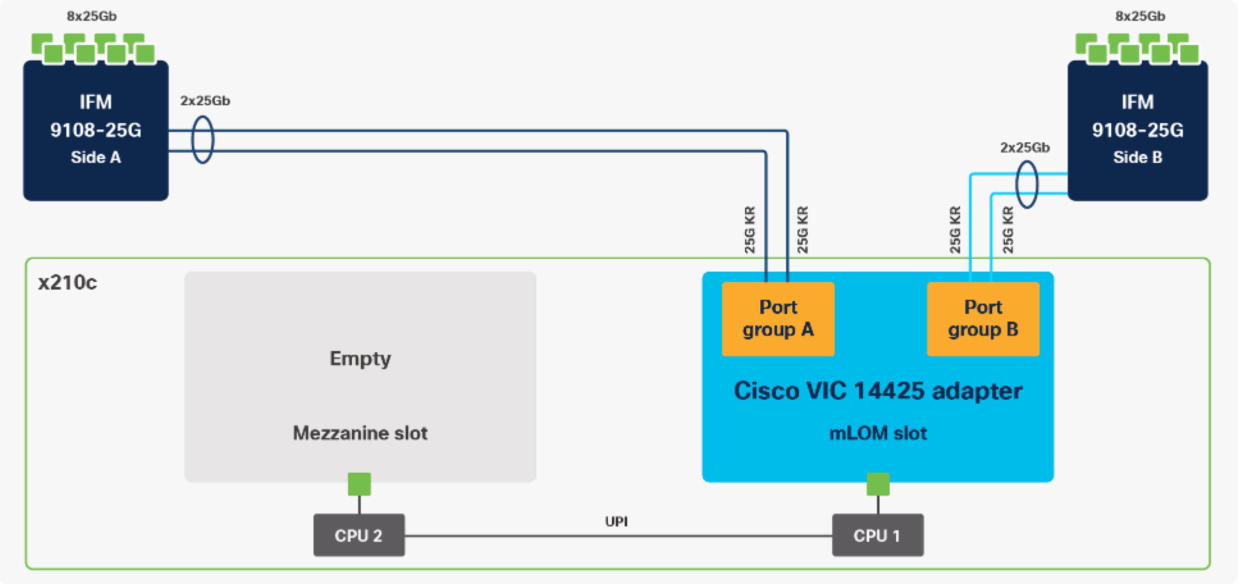

Cisco UCS VIC 14825

The optional Cisco UCS 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 10. Cisco VIC 14425 and 14825 in Cisco UCS X210c M6

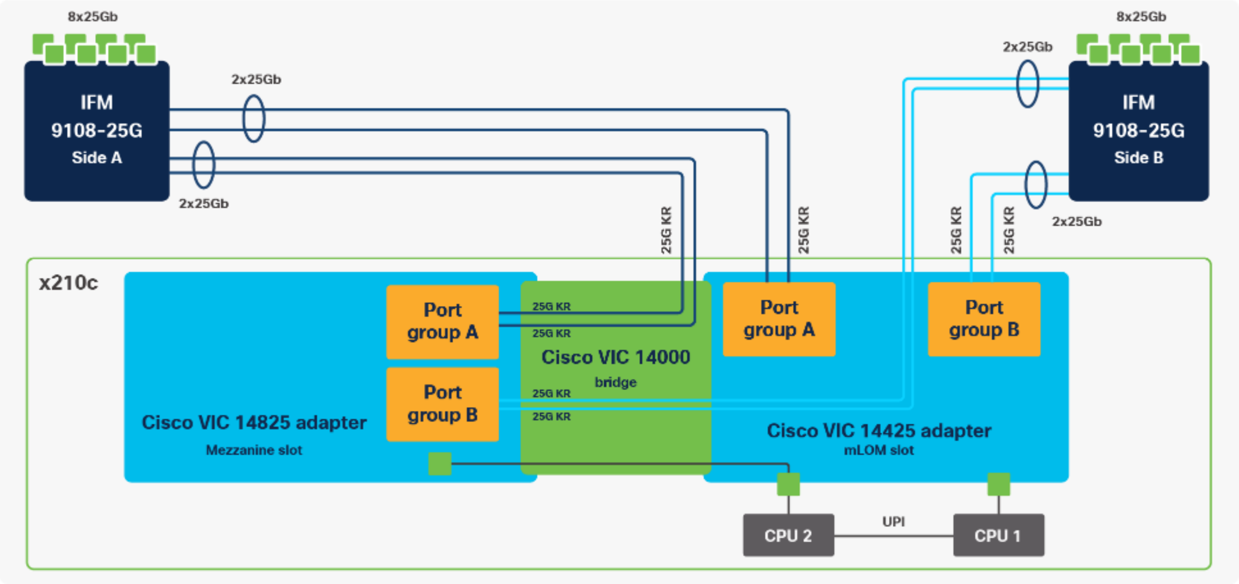

Cisco UCS VIC 1467

The Cisco UCS VIC 1467 is a quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M6 generation of Cisco UCS C-Series Rack Servers. The card supports 10/25-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBA

Cisco UCS VIC 1477

The Cisco UCS VIC 1477 is a dual-port Quad Small Form-Factor (QSFP28) mLOM card designed for the M6 generation of Cisco UCS C-Series Rack Servers. The card supports 40/100-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as NICs or HBAs.

Cisco UCS 15000 Series Virtual Interface Cards (VICs)

The Cisco UCS VIC 15000 series is designed for Cisco UCS X-Series M6 Blade Servers and Cisco UCS C-Series M6 Rack Servers. The adapters are capable of supporting 10/25/50/100/200-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). They incorporate Cisco's next-generation Converged Network Adapter (CNA) technology and offer a comprehensive feature set, providing investment protection for future feature software releases.

The Cisco VIC 15000 series can support 512 PCI Express (PCIe) virtual devices, either virtual network interface cards (vNICs) or virtual Host Bus Adapters (vHBAs), with a high rate of I/O operations per second (IOPS), support for lossless Ethernet, and 10/25/50/100/200-Gbps connection to servers. The PCIe Generation 4 x16 interface helps ensure optimal bandwidth to the host for network-intensive applications, with a redundant path to the fabric interconnects. Cisco VICs support NIC teaming with fabric failover for increased reliability and availability. In addition, it provides a policy-based, stateless, agile server infrastructure for your data center.

Cisco UCS VIC 15231

The Cisco UCS VIC 15231 is a 2x100-Gbps Ethernet/FCoE capable modular LAN on motherboard (mLOM) designed exclusively for the Cisco UCS X210 Compute Node. The Cisco UCS VIC 15231 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.

Figure 11. **Single Cisco VIC 15231 in Cisco UCS X210c M6**

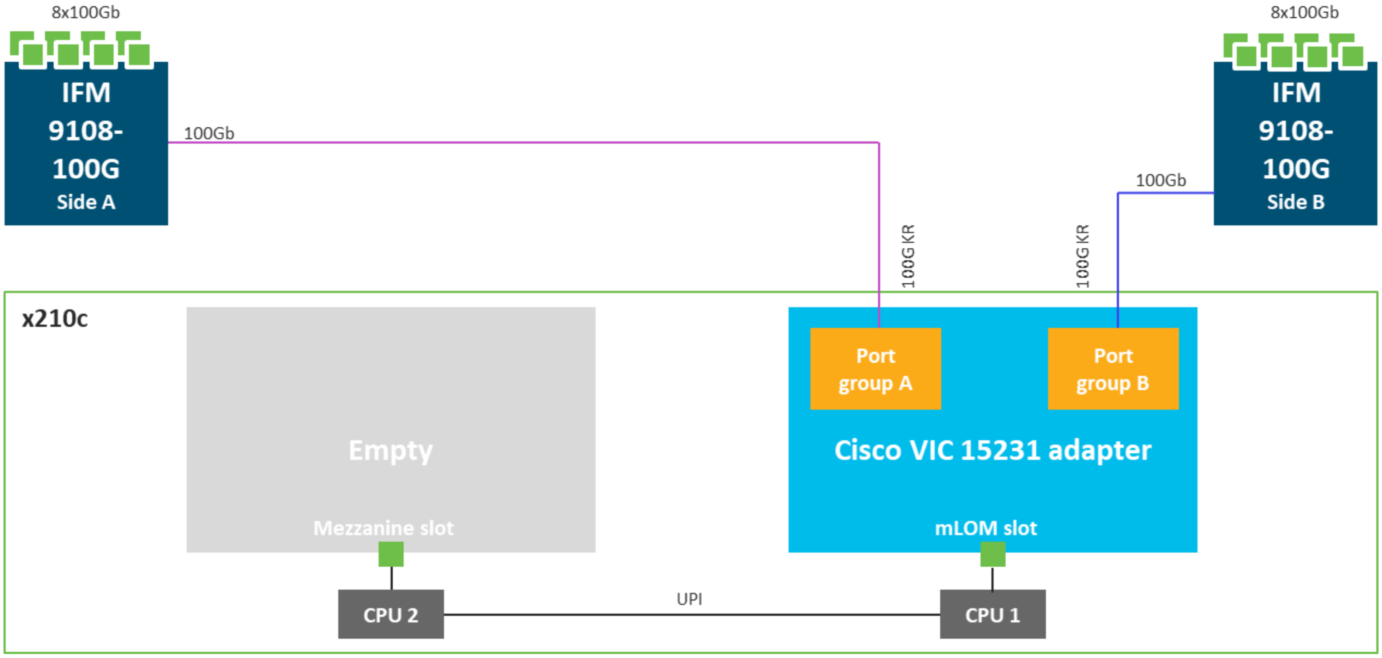

The Cisco UCS VIC 15000 series is capable to work with 25Gbps and 100Gbps Intelligent Fabric Modules as shown in Figure 12.

Figure 12. Cisco UCS VIC Configuration option with VIC15231

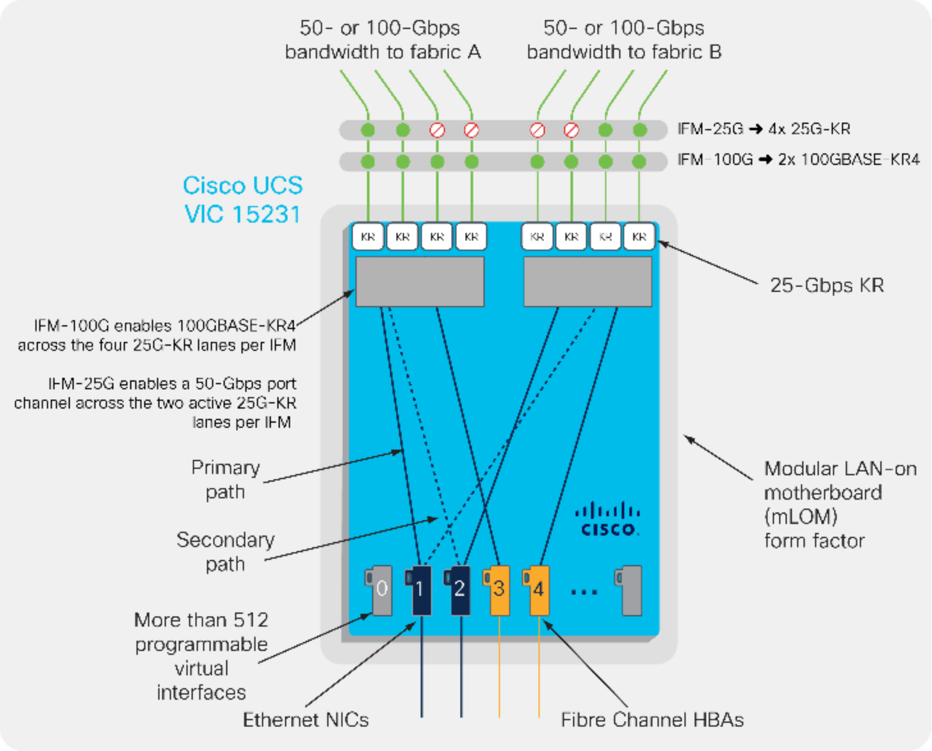
50- or 100-Gbps bandwidth to fabric A
50- or 100-Gbps bandwidth to fabric B
IFM-25G ➜ 4x 25G-KR
IFM-100G ➜ 2x 100GBASE-KR4
Cisco UCS VIC 15231
KR
25-Gbps KR
IFM-100G enables 100GBASE-KR4 across the four 25G-KR lanes per IFM
IFM-25G enables a 50-Gbps port channel across the two active 25G-KR lanes per IFM
Primary path
Secondary path
Modular LAN-on motherboard (mLOM) form factor
cisco
0
1
2
3
4
...
More than 512 programmable virtual interfaces
Ethernet NICs
Fibre Channel HBAs

Cisco UCS VIC 15428

The Cisco UCS VIC 15428 is a quad-port Small Form-Factor Pluggable (SFP+/SFP28/SFP56) mLOM card designed for the M6 generation of Cisco UCS C-series rack servers. The card supports 10/25/50-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

Cisco Unified Compute System X-Series

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

Figure 13. **Cisco UCS X9508 Chassis**

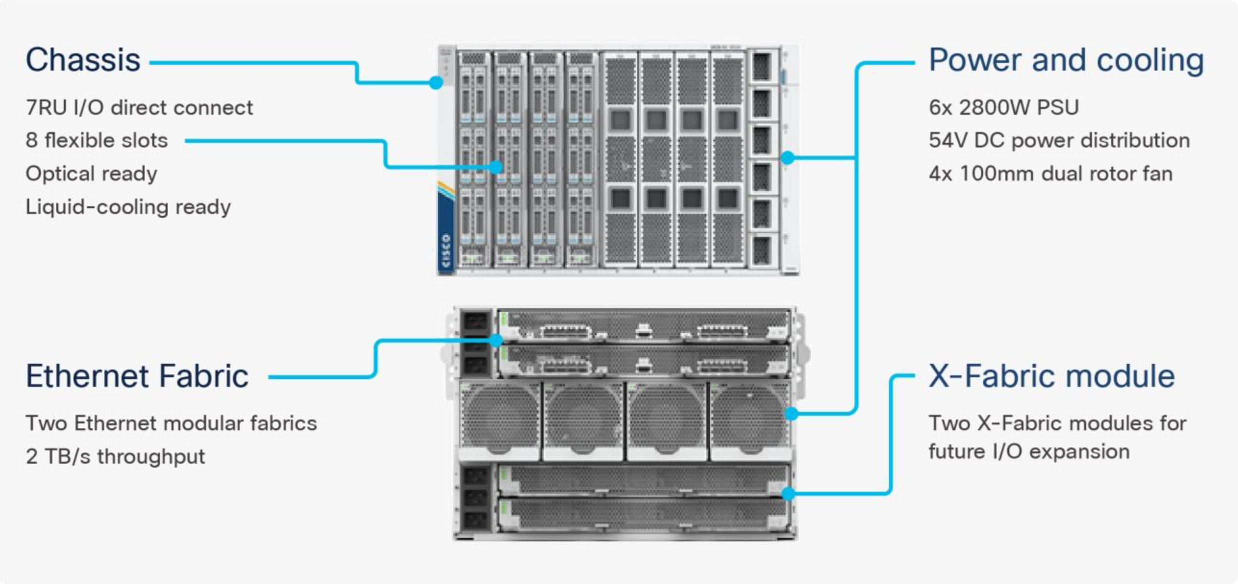

Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As seen in Figure 14 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 14. **Cisco UCS X9508 Chassis – Innovative Design**

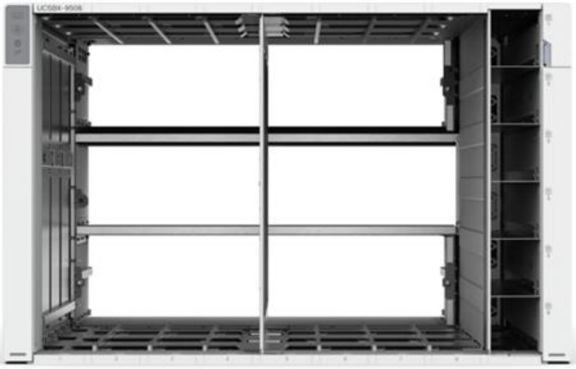

The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes, GPU accelerators and a pool of future I/O resources that may include disk storage, and memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) that connect 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 15. Cisco UCSX 9108-25G Intelligent Fabric Module

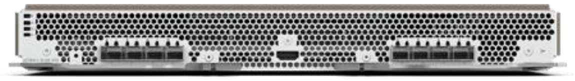

Each IFM supports eight 25Gb uplink ports for connecting the Cisco UCS X9508 Chassis to the FIs and 32 25Gb server ports for the eight compute nodes. IFM server ports can provide up to 200 Gbps of unified fabric connectivity per compute node across the two IFMs. The uplink ports connect the chassis to 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).

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 16. Cisco UCS 9108-100G Intelligent Fabric Module

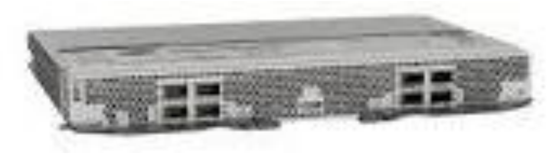

Each IFM supports eight 100Gb uplink ports for connecting the Cisco UCS X9508 Chassis to the Cisco UCS fifth generation 6556 FIs and 8 100Gb server ports for the eight compute nodes. IFM server ports can provide up to 200 Gbps of unified fabric connectivity per compute node across the two IFMs. The uplink ports connect the chassis to the 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 17:

Figure 17. **Cisco UCS X210c M6 Compute Node**

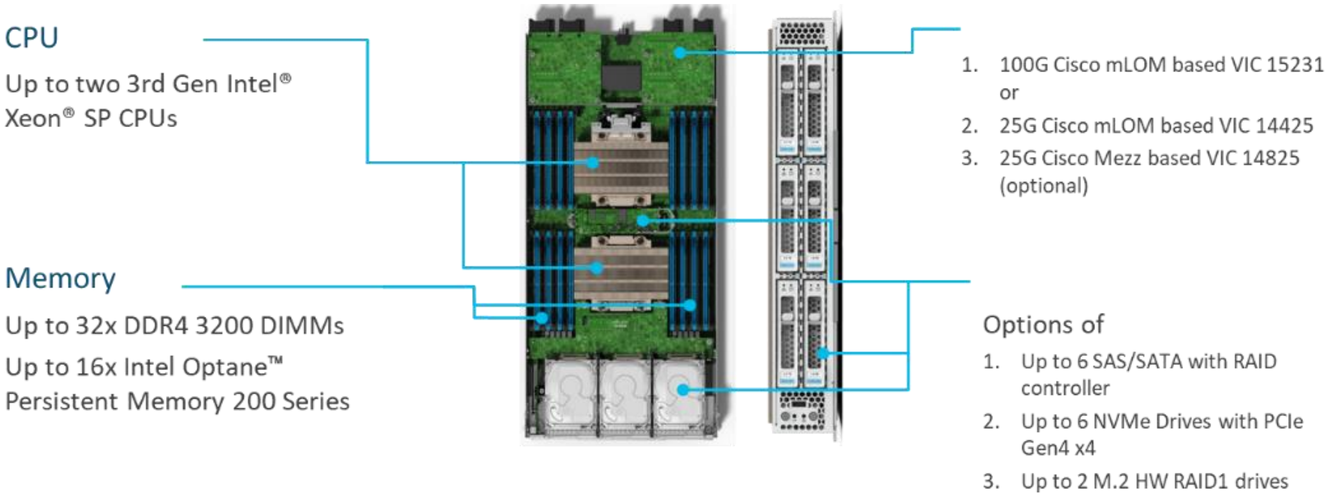
CPU
Up to two 3rd Gen Intel® Xeon® SP CPUs
Memory
Up to 32x DDR4 3200 DIMMs
Up to 16x Intel Optane™ Persistent Memory 200 Series
1. 100G Cisco mLOM based VIC 15231 or
2. 25G Cisco mLOM based VIC 14425
3. 25G Cisco Mezz based VIC 14825 (optional)
Options of
1. Up to 6 SAS/SATA with RAID controller
2. Up to 6 NVMe Drives with PCIe Gen4 x4
3. Up to 2 M.2 HW RAID1 drives

The Cisco UCS X210c M6 features:

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

Cisco UCS B-Series Blade System

This section describes the Cisco UCS B-Series Blade System.

Cisco UCS 5108 Server Chassis

The Cisco UCS 5108 server chassis revolutionizes the use and deployment of blade-based systems. By incorporating unified fabric, integrated, embedded management, and fabric extender technology, the Cisco Unified Computing System enables the chassis to have fewer physical components, no independent management, and to be more energy efficient than traditional blade server chassis.

This simplicity eliminates the need for dedicated chassis management and blade switches, reduces cabling, and enables the Cisco Unified Computing System to scale to 40 chassis without adding complexity. The Cisco UCS 5108 server chassis is a critical component in delivering the Cisco Unified Computing System benefits of data center simplicity and IT responsiveness.

For more information about the UCS 5108 server chassis see: https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-5100-series-blade-server-chassis/index.html.

Cisco UCS 2408 Fabric Extender

The Cisco UCS Fabric Extender connects the I/O fabric between the Cisco UCS Fabric Interconnect and the Cisco UCS Blade Server Chassis, enabling a lossless and deterministic converged fabric to connect all blades and chassis together.

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

Figure 18. Cisco UCS 2408 Fabric Extender

For more information about the UCS 2408 Fabric Extender see: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-742624.pdf.

Cisco UCS B200 M6 Blade Servers

The Cisco UCS B200 M6 server shown in Figure 20, is the last half-width blade supported in the Cisco UCS 5108 chassis.

Figure 19. Cisco UCS B200 M6 Blade Server

It features the following:

- 3rd Gen Intel® Xeon® Scalable and Intel® Xeon® Scalable processors with up to 40 cores per socket.
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane™ DC Persistent Memory.
- Up to two GPUs.
- Two Small-Form-Factor (SFF) drive slots.
- Up to four M.2 SATA drives.
- Up to 80 Gbps of I/O throughput with Cisco UCS 6454 FI.

For more information about the Cisco UCS B200 M6 Blade Servers see: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/datasheet-c78-2368888.html.

Cisco UCS C-Series Rack Servers

Cisco UCS C-Series Rack Servers deliver unified computing in an industry-standard form factor to reduce TCO and increase agility. Each server addresses varying workload challenges through a balance of processing, memory, I/O, and internal storage resources.

Cisco UCS C220 M6 Rack Servers

The Cisco UCS C220 M6 rack server shown in Figure 20, is a high-density 2-socket rack server that is an upgrade from the Cisco UCS C220 M5.

Figure 20. Cisco UCS C220 M6 Rack Server

It features the following:

- 3rd Gen Intel® Xeon® Scalable and Intel® Xeon® Scalable processors, 2-socket.
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane™ DC Persistent Memory.
- Up to 10 Small-Form-Factor (SFF) 2.5-inch drives or 4 Large-Form-Factor (LFF) 3.5-inch drives.

- Support for 12-Gbps SAS modular RAID controller in a dedicated slot, leaving the remaining PCIe Generation 4.0 slots available for other expansion cards.
- Modular LAN-On-Motherboard (mLOM) slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot.
- Dual embedded Intel x550 10GBASE-T LAN-On-Motherboard (LOM) ports.
- Up to 100 Gbps of I/O throughput with Cisco UCS 6454 FI.

For more information about the Cisco UCS C220 M6 Blade Servers see: https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/ucs-c220-m6-rack-server-ds.html.

Cisco UCS C240 M6 Rack Servers

The Cisco UCS C220 M6 rack server shown in Figure 22, is a high-density 2-socket rack server that is an upgrade from the Cisco UCS C220 M5.

Figure 21. Cisco UCS C240 M6 Rack Server

It features the following:

- 3rd Gen Intel® Xeon® Scalable and Intel® Xeon® Scalable processors, 2-socket.
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane™ DC Persistent Memory.
- Up to 28 Small-Form-Factor (SFF) 2.5-inch drives or 16 Large-Form-Factor (LFF) 3.5-inch drives.
- Support for 12-Gbps SAS modular RAID controller in a dedicated slot, leaving the remaining PCIe Generation 4.0 slots available for other expansion cards.
- Modular LAN-On-Motherboard (mLOM) slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot.
- Dual embedded Intel x550 10GBASE-T LAN-On-Motherboard (LOM) ports.
- Up to 100 Gbps of I/O throughput with Cisco UCS 6454 FI.

For more information about the Cisco UCS C240 M6 Blade Servers see: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/ucs-c240-m6-rack-server-ds.html.

Cisco UCS Scalability in FlexPod

The scale of Cisco UCS in the context of FlexPod is possible in scale-up and scale-out, based on performance, high availability, and locality requirements. A single UCS Domain can scale-up up to 160 server nodes with any mixture of blade and rack server options. The key KPI limiting the scale-up option is the required network bandwidth from/to the servers in the UCS domain. The more ports are used to connect servers, the fewer ports are available to connect the UCS Fabric Interconnects to the next-hop network switches. The use of multiple UCS Domains within one FlexPod is not only limited to bandwidth requirements in case multiple rooms in the data center are used to deploy FlexPod i.e., to provide better high availability or limited impact to the FlexPod in case of environmental issues.

In addition to local scale options, FlexPod is also supported as dual-site and multi-site deployments with multiple NetApp storages and Cisco UCS Domains. Some of the deployment options will be shown in the Deployment Option Chapter in this document.

Cisco Nexus Switching Fabric

Cisco Nexus series switches provide an Ethernet switching fabric for communications between the Cisco UCS, NetApp storage controllers, and the rest of a customer's network. There are many factors to consider when choosing the main data switch in this type of architecture to support both the scale and the protocols required for the resulting applications. All Cisco Nexus switch models of the Cisco Nexus 9000 series are supported in the FlexPod design. However, be aware that there may be slight differences in setup and configuration based on the switch used. The validation for this deployment leverages the Cisco Nexus 9000 series fixed switch configuration, which deliver high-performance 10/25/40/50/100/400GbE ports, density, low latency, and exceptional power efficiency in a broad range of compact form factors.

Many of the most recent FlexPod designs also use this switch due to the advanced feature set and the ability to support Application Centric Infrastructure (ACI) mode. When leveraging ACI fabric mode, the Cisco Nexus 9000 series switches are deployed in a spine-leaf architecture. Although the reference architecture covered in this design does not leverage ACI, it lays the foundation for customer migration to ACI in the future and fully supports ACI today if required.

For more information, go
to: http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html.

This FlexPod design deploys a single pair of Cisco Nexus 9000 series switches like the Cisco Nexus 93240YC-FX2 top-of-rack switch (Figure 22) within each placement, using the traditional standalone mode running NX-OS.

Figure 22. **Cisco Nexus 93240YC-FX2**

For larger-scale deployment with a single pair of Cisco Nexus 9000 switches, the Cisco Nexus 93360YC-FX2 provides more ports and throughput.

Figure 23. Cisco Nexus 933360YC-FX2 Switch

The Cisco Nexus 93360YC-FX2 Switch is a 2RU switch that supports 7.2 Tbps of bandwidth and 2.4 bpps. The 96 downlink ports on the 93360YC-FX2 can support 1-, 10-, or 25-Gbps Ethernet or 16- or 32-Gbps Fibre Channel ports, offering deployment flexibility and investment protection. The 12 uplink ports can be configured as 40- or 100-Gbps Ethernet, offering flexible migration options. This switch was chosen for this solution because of the extra flexibility and scaling the 12 40- or 100-Gbps uplink ports offer.

The Cisco Nexus 93180YC-FX, 93360YC-FX2, and 9336C-FX2-E switches now support SAN switching, allowing both Ethernet and Fibre Channel SAN switching in a single switch. In addition to 16- or 32-Gbps Fibre Channel, these switches also support 100-Gbps FCoE, allowing port-channeled 100-Gbps FCoE uplinks from the Cisco UCS 6536 Fabric Interconnects to Cisco Nexus switches in SAN switching mode.

Cisco Nexus Dashboard

Cisco Nexus Dashboard Fabric Controller (NDFC) is an application deployed on top of a Cisco Nexus Dashboard installation and can monitor, configure, and analyze Cisco Nexus 9000 switching fabrics. Cisco Nexus Dashboard is deployed as a virtual appliance from an OVA, as a physical appliance, or on top of a Linux operating system, and is managed through a web browser. Once the Cisco Nexus switches are added with the appropriate credentials and licensing, monitoring the Ethernet fabric can begin. The NDFC application introduces the Spine-Leaf based Easy Fabric which configures a VXLAN/BGP EVPN network fabric and supports the connection to remote sites and public cloud providers through Cisco Nexus Dashboard Orchestrator.

Cisco Nexus Dashboard Orchestrator (NDO) is an application deployed on top of Cisco Nexus Dashboard. It configures multi-fabric, multi-site, and hybrid-cloud connections for Cisco Nexus Dashboard and ACI based networks.

Cisco Nexus Dashboard Insights (NDI) in an application deployed on top of Cisco Nexus Dashboard.

Cisco Nexus Dashboard integration with Cisco Intersight

The Cisco Network Dashboard Insights (NDI) application provides several TAC assist functionalities that are useful when working with Cisco TAC. The Cisco NDI app collects the CPU, device name, device product id, serial number, version, memory, device type, and disk usage information for the nodes in the fabric. The Cisco NDI application is connected to the Cisco Intersight cloud portal through a device connector which is embedded in the management controller of the Cisco Nexus Dashboard platform. The device connector provides a secure way for the connected Cisco Nexus Dashboard to send and receive information from the Cisco Intersight portal, using a secure Internet connection.

Cisco MDS 9132T 32G Multilayer Fabric Switch

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

Figure 24. Cisco MDS 9132T 32G Multilayer Fabric Switch

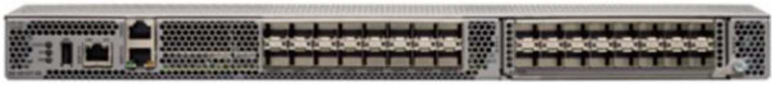

The Cisco MDS 9132T delivers advanced storage networking features and functions with ease of management and compatibility with the entire Cisco MDS 9000 family portfolio for reliable end-to-end connectivity. This switch also offers state-of-the-art SAN analytics and telemetry capabilities that have been built into this next-generation hardware platform. This new state-of-the-art technology couples the next-generation port ASIC with a fully dedicated network processing unit designed to complete analytics calculations in real time. The telemetry data extracted from the inspection of the frame headers are calculated on board (within the switch) and, using an industry-leading open format, can be streamed to any analytics-visualization platform. This switch also includes a dedicated 10/100/1000BASE-T telemetry port to maximize data delivery to any telemetry receiver, including Cisco Data Center Network Manager.

Cisco DCNM-SAN

Cisco DCNM-SAN can be used to monitor, configure, and analyze Cisco 32Gbps Fibre Channel fabrics and show information about the Cisco Nexus switching fabric. Cisco DCNM-SAN is deployed as a virtual appliance from an OVA and is managed through a web browser. Once the Cisco MDS and Cisco Nexus switches are added with the appropriate credentials and licensing, monitoring of the SAN and Ethernet fabrics can begin. Additionally, VSANs, device aliases, zones, and zone sets can be added, modified, and deleted using the DCNM point-and-click interface. Device Manager can also be used to configure the Cisco MDS switches. SAN Analytics can be added to Cisco MDS switches to provide insights into the fabric by allowing customers to monitor, analyze, identify, and troubleshoot performance issues. Cisco DCNM-SAN has been renamed Cisco Nexus Dashboard Fabric Controller (NDFC), which is a Cisco Nexus Dashboard App. However, as of the publishing date of this CVD, Cisco DCNM 11.5(4) was the suggested release and was what was used in this CVD. The next FlexPod CVD will use NDFC as part of Cisco Nexus Dashboard.

Cisco DCNM integration with Cisco Intersight

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

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 A800 and A400 have been chosen for solution validation although any other AFF series could be used instead.

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 25. NetApp AFF A400 Front View

Figure 26. NetApp AFF A400 Rear View

NetApp AFF A800

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

Figure 27. NetApp AFF A800 Front View

Figure 28. NetApp AFF A800 Rear View

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.

You can view or download more technical specifications of the AFF A-series controllers here: https://www.netapp.com/us/media/ds-3582.pdf

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

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

NetApp ONTAP 9.11.1

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 below 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, go to: https://docs.netapp.com/us-en/active-iq-unified-manager/

NetApp SnapCenter

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

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

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

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

Data protection is supported for Microsoft Exchange Server, Microsoft SQL Server, Oracle Databases on Linux or AIX, SAP HANA database, and Windows Host Filesystems running on ONTAP systems. It is also supported for other standard or custom applications and databases by providing a framework to create user-defined SnapCenter plug-ins. You may install only the plug-ins that are appropriate for the data that you want to protect.

For more information on SnapCenter 4.7, go to the SnapCenter software documentation:
https://docs.netapp.com/us-en/snapcenter/index.html

NetApp Astra Trident CSI Plugin

Astra Trident is an open-source, fully supported storage orchestrator for containers created by NetApp. It has been designed from the ground up to help you meet your containerized applications' persistence demands using industry-standard interfaces, such as the Container Storage Interface (CSI). With Astra Trident, microservices and containerized applications can take advantage of enterprise-class storage services provided by the full NetApp portfolio of storage systems. In a FlexPod environment, Astra Trident is utilized to allow end users to dynamically provision and manage persistent volumes for containers backed by FlexVols and LUNs hosted on ONTAP-based products such as NetApp AFF and FAS systems.

Astra Trident has a rapid development cycle, and just like Kubernetes, is released four times a year. Starting with the v21.04.0 release, the setup of Astra Trident is performed by the Trident operator using a Helm chart which makes large scale deployments easier and provides additional support including self-healing for the pods that are deployed as a part of the Astra Trident install.

NetApp Astra Control Center

NetApp Astra Control Center offers a rich set of storage and application-aware data management services for stateful Kubernetes workloads deployed in an on-premises environment and hybrid cloud environment powered by NetApp data protection technology.

Astra Control offers critical capabilities for Kubernetes application data lifecycle management:

- Automatically manage persistent storage
- Create application-aware, on-demand snapshots and backups
- Automate policy-driven snapshot and backup operations
- Replicate application to a remote system using NetApp SnapMirror technology
- Migrate applications and data from one Kubernetes cluster to another
- Easily clone an application from production to staging
- Visualize application health and protection status
- Use a user interface or an API to implement your backup and migration workflows

Red Hat OpenShift Container Platform

The RedHat OpenShift Container Platform (OCP) is a container application platform that brings together CRI-0 and Kubernetes and provides an API and web interface to manage these services. CRI-O is an implementation of the Kubernetes CRI (Container Runtime Interface) to enable using Open Container Initiative (OCI) compatible runtimes. It is a lightweight alternative to using Docker as the runtime for Kubernetes.

OCP allows customers to create and manage containers. Containers are standalone processes that run within their own environment, independent of operating system and the underlying infrastructure. OCP helps developing, deploying, and managing container-based applications. It provides a self-service platform to create, modify, and deploy applications on demand, thus enabling faster development and release life cycles. OCP has a micro-services-based architecture of smaller, decoupled units that work together. It runs on top of a Kubernetes cluster, with data about the objects stored in etcd, a reliable clustered key-value store.

Figure 29. OpenShift Container Platform Overview

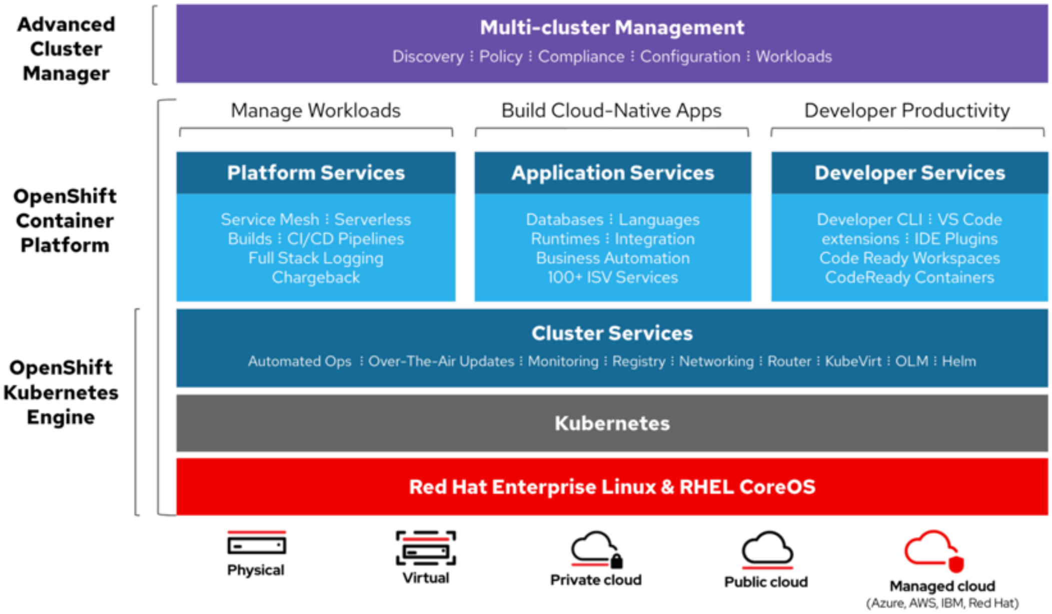

Kubernetes Infrastructure

Within OpenShift Container Platform, Kubernetes manages containerized applications across a set of CRI-O runtime hosts and provides mechanisms for deployment, maintenance, and application-scaling. The CRI-O service packages, instantiates, and runs containerized applications.

A Kubernetes cluster consists of one or more masters and a set of worker nodes. This solution design includes HA functionality at the hardware as well as the software stack. A Kubernetes cluster is designed to run in HA mode with 3 master nodes and a minimum of 2 worker nodes to help ensure that the cluster has no single point of failure.

Red Hat Core OS

OpenShift Container Platform uses Red Hat Enterprise Linux CoreOS (RHCOS), a container-oriented operating system that combines some of the best features and functions of the CoreOS and Red Hat Atomic Host operating systems. RHCOS is specifically designed for running containerized applications from OpenShift Container Platform and works with new tools to provide fast installation, Operator-based management, and simplified upgrades.

RHCOS includes the following:

- Ignition, which OpenShift Container Platform uses as a first boot system configuration for initially bringing up and configuring machines.

- CRI-O, a Kubernetes native container runtime implementation that integrates closely with the operating system to deliver an efficient and optimized Kubernetes experience. CRI-O provides facilities for running, stopping, and restarting containers. It fully replaces the Docker Container Engine, which was used in OpenShift Container Platform 3.
- Kubelet, the primary node agent for Kubernetes that is responsible for launching and monitoring containers.

Note: RHCOS was used on all control planes and worker nodes to support the automated OCP 4 deployment.

Red Hat Ansible

Red Hat Ansible Automation helps Red Hat OpenShift Container Platform users create and run reusable infrastructure code and automate provisioning tasks for infrastructure components.

Ansible is simple and powerful, allowing users to easily manage various physical devices within FlexPod- including the provisioning of Cisco UCS bare metal servers, Cisco Nexus switches, and NetApp AFF storage. Using Ansible's Playbook-based automation is easy and integrates into your current provisioning infrastructure.

Finally, Ansible also provides robust container and native Kubernetes management, expanding to Red Hat OpenShift Container Platform and other container technologies.

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.

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.

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

Cisco Intersight integrates with VMware vCenter, NetApp Storage, 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 30. Cisco Intersight and vCenter/NetApp Integration

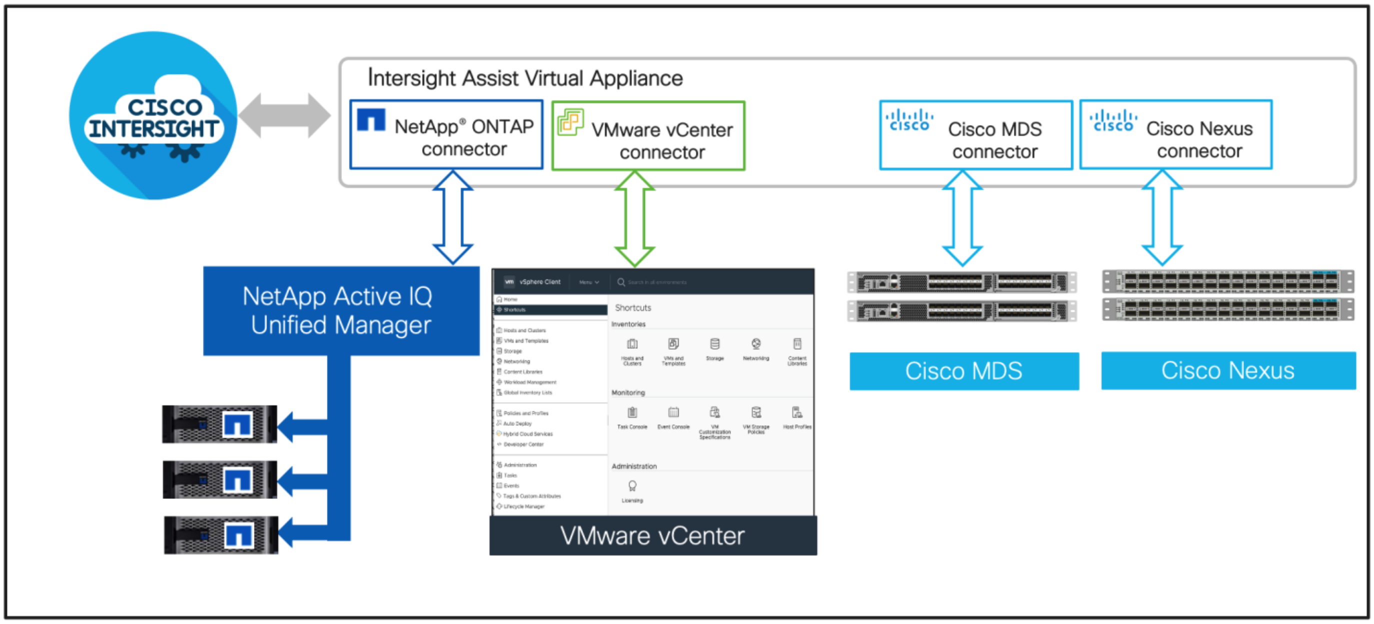

The device connector provides a secure way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware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

FlexPod is available as 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, 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 31. FlexPod in Cisco Intersight

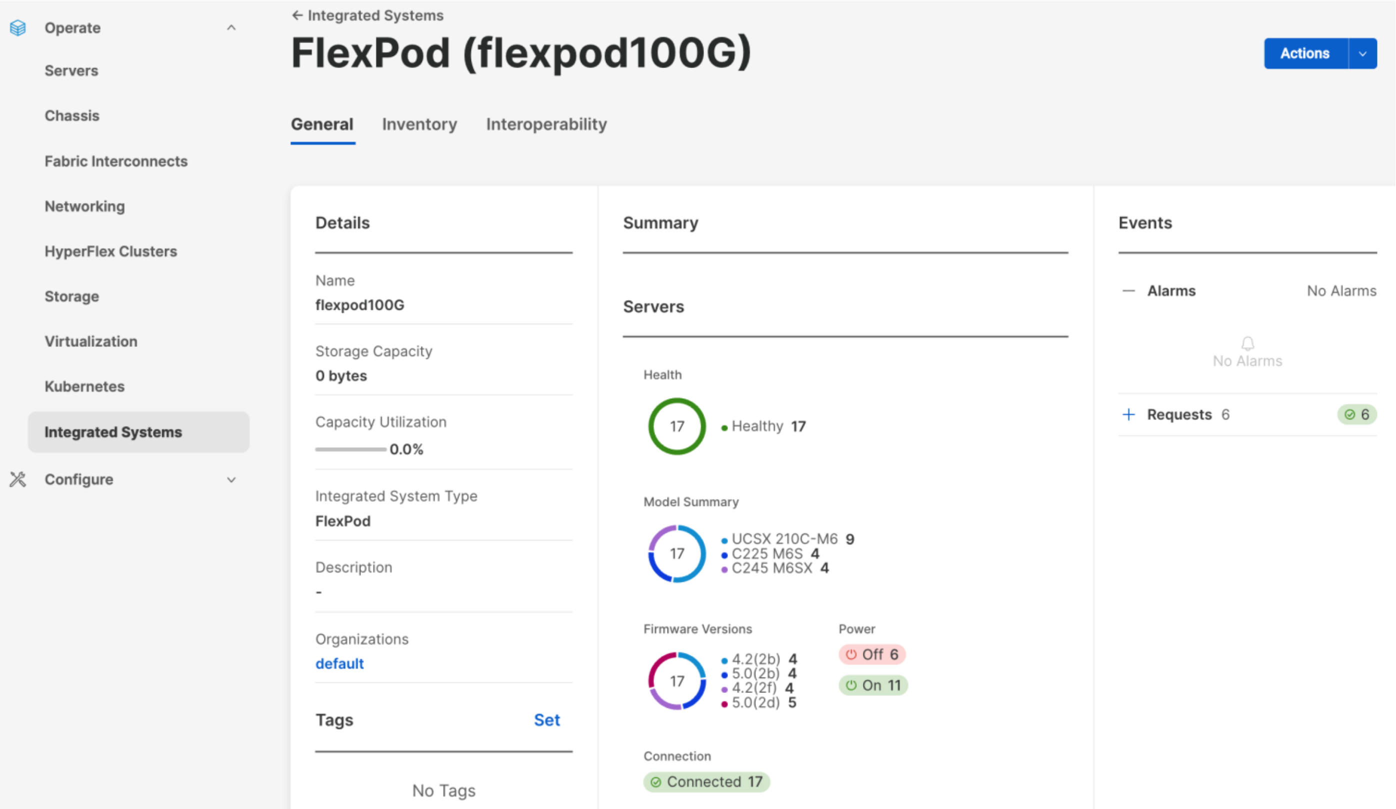

FlexPod (flexpod100G)

Details

Name
flexpod100G

Storage Capacity
0 bytes

Capacity Utilization
0.0%

Integrated System Type
FlexPod

Description
-

Organizations
default

Tags Set

No Tags

Summary

Servers

Health
17 ● Healthy 17

Model Summary
17
● UCSX 210C-M6 9
● C225 M6S 4
● C245 M6SX 4

Firmware Versions
17
● 4.2(2b) 4
● 5.0(2b) 4
● 4.2(2f) 4
● 5.0(2d) 5

Power
Off 6
On 11

Connection
Connected 17

Events

— Alarms No Alarms

No Alarms

+ Requests 6 ✓ 6

Infrastructure as Code with Ansible

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

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

Figure 32. Infrastructure as Code with Ansible

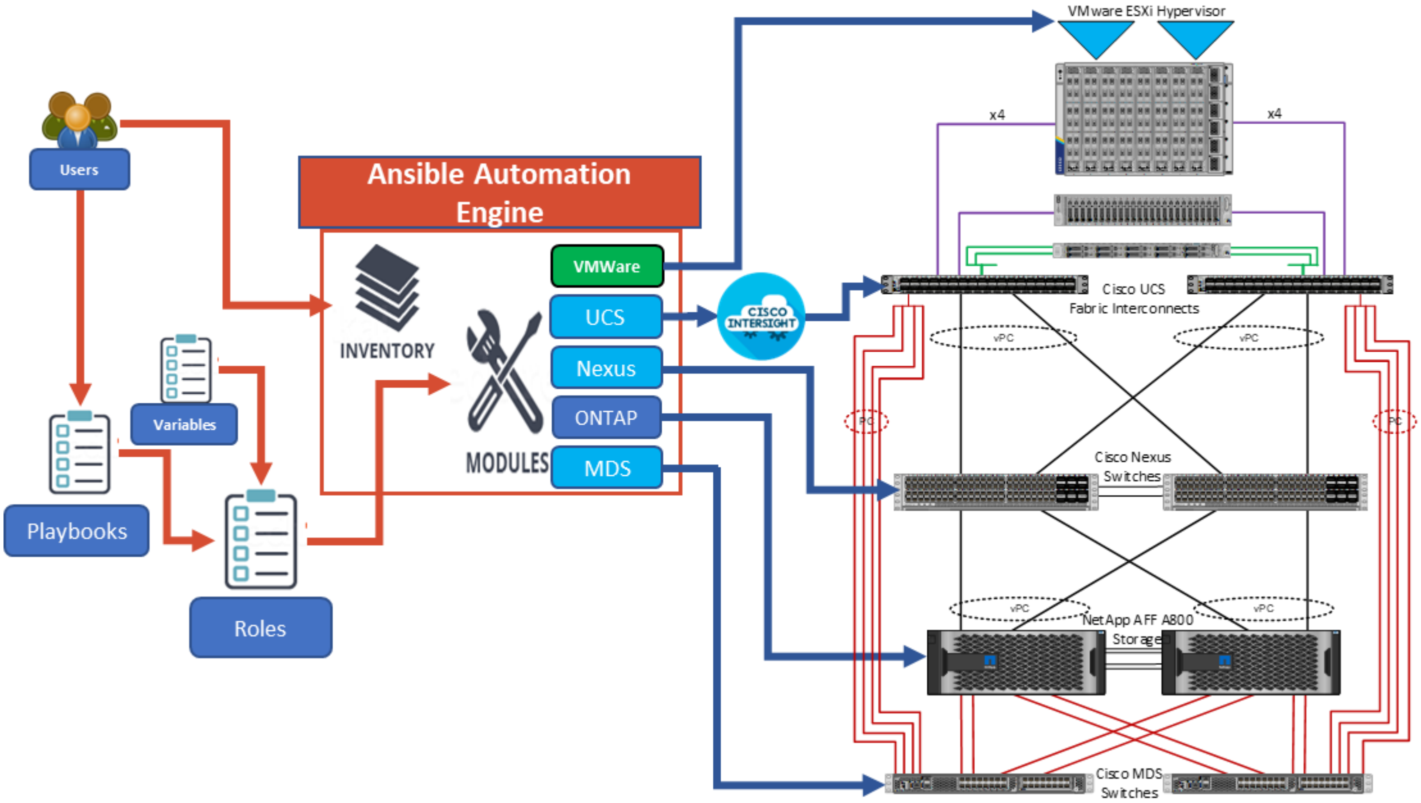

As illustrated in Figure 32, the Ansible playbooks to configure the different sections of the solution invoke a set of Roles and consume the associated variables that are required to setup the solution. The variables needed for this solution can be split into two categories – user input and defaults/ best practices. Based on the installation environment customers can choose to modify the variables to suit their requirements and proceed with the automated installation.

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

Solution Design

This chapter contains the following:

The FlexPod Datacenter for Red Hat OpenShift Container Platform provides an end-to-end architecture with Cisco and NetApp technologies that demonstrate support for OCP workloads with high availability and server redundancy. The architecture supports the deployment of OCP as bare metal cluster or virtualized on top of supported virtual machine monitors like Vmware vSphere or KVM on Cisco UCS servers within FlexPod infrastructure, with the Cisco UCS servers and NetApp storage attached to the Cisco Nexus switches in NXOS mode.

Figure 33 illustrates a sample design with the required management components, like Intersight Assist or Active IQ Unified Manager (AIQUM), installed on the FlexPod stack as virtual machines. Each of the components can be scaled easily to support specific business requirements. For example, additional OCP nodes can be deployed to scale the OCP environment to increase compute capacity, additional storage controllers or disk shelves can be deployed to improve I/O capability and throughput.

Note: The solution was validated using Cisco UCS X210 M6, UCS B200 M6, UCS B200 M5, and UCS C240 M5 servers to show the versatility of the Cisco UCS platform. Customers can choose to deploy OCP on any Cisco UCS servers depending on their requirements.

Figure 33. FlexPod Datacenter with Red Hat OpenShift Container Platform 4 Architecture View

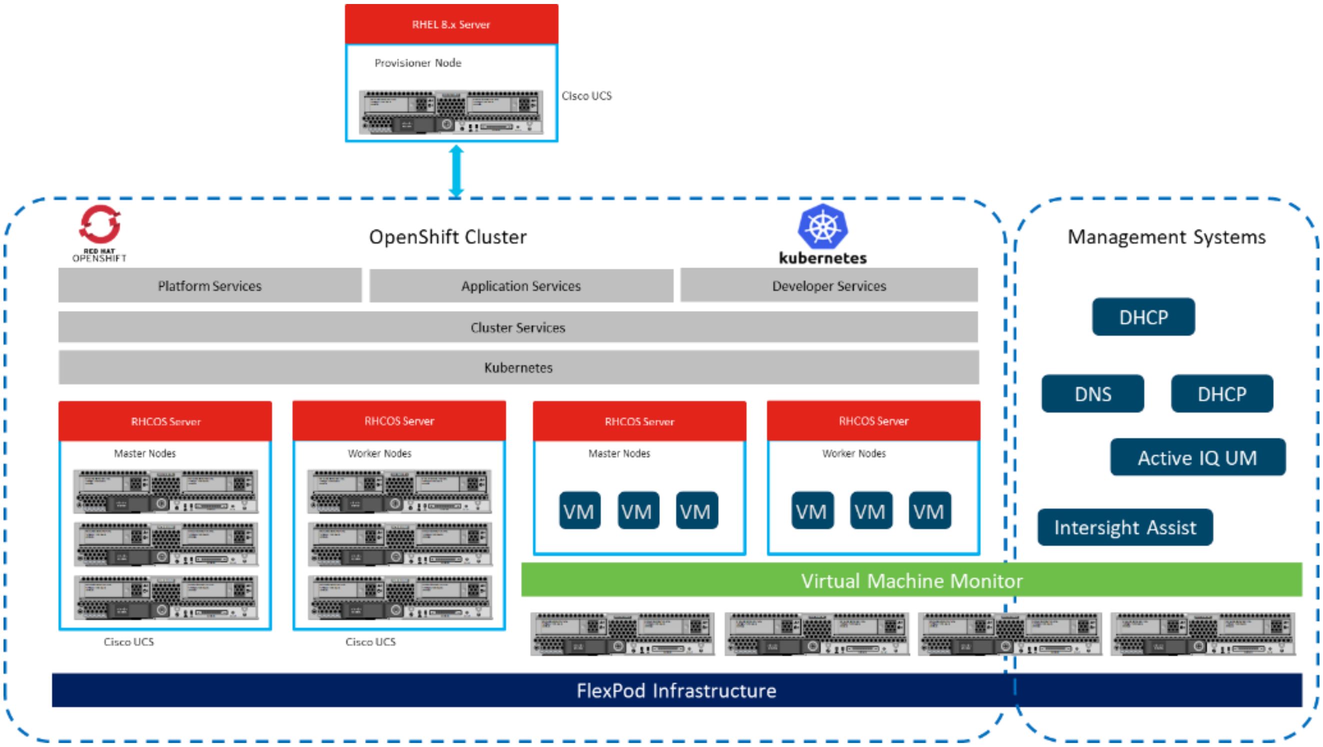

OpenShift Virtualization is an add-on to OpenShift Container Platform that allows you to run and manage virtual machine workloads alongside container workloads. The OpenShift Virtualization feature has been validated within this solution to deploy traditional VMs into OpenShift where they run side by side with containers on the same OCP cluster deployed on the FlexPod infrastructure.

Figure 34 illustrates a high-level overview of the FlexPod for OCP cluster architecture.

Figure 34. OpenShift Virtualization

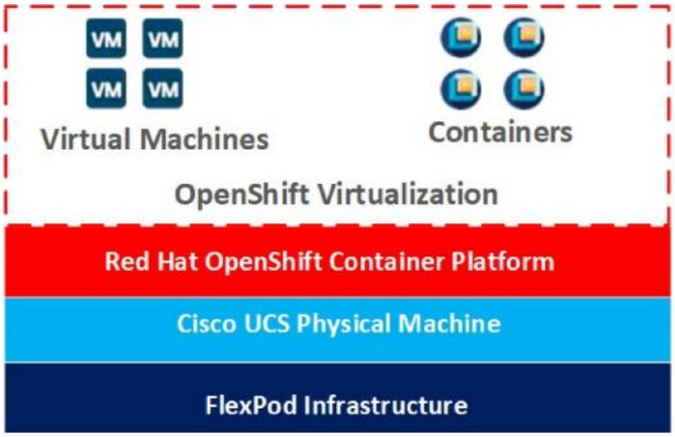

Requirements

This section explains the key design requirement and various prerequisites for delivering this new solution.

The FlexPod Datacenter with Red Hat OpenShift Container Platform solution closely aligns with all FlexPod CVDs and 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 components beyond what is validated and documented in this guide.
- Simplified design with ability to automate and integrate with external automation and orchestration tools.

For Red Hat OCP 4 integration into a traditional FlexPod solution, the following specific design considerations are also observed:

- Deployment option for one or three master nodes, where the option with one master is only recommended for non-productive installations.
- A minimum of 2 worker nodes with ability to increase the nodes as the load requirements increase.
- Automating the FlexPod infrastructure deployment and OCP installation by utilizing Ansible Playbooks to simplify the installation and reduce the deployment time.
- Present persistent storage (volumes) to the containerized applications by utilizing the NetApp Astra Trident CSI framework.
- Dedicated Cisco UCS vNICs for different traffic needs with UCS Fabric Failover for high availability.

Physical Topology

This FlexPod design utilizes Cisco UCS servers connected and managed through Cisco UCS Fabric Interconnects and the Intersight Infrastructure Manager (IMM) to manage the servers. These high-performance servers are configured as compute nodes where Red Hat Core OS (RHCOS) is loaded using SAN boot leveraging FC LUNs from the NetApp AFF storage, optional the use of local disk for boot is documented. The persistent storage volumes for containers are provisioned on the NetApp AFF A400 using NFS NAS storage and iSCSI or FC block storage.

IP-based Storage Access: iSCSI and NFS

A typical topology for the iSCSI FlexPod Datacenter is shown in Figure 35.

Figure 35. FlexPod Physical Topology

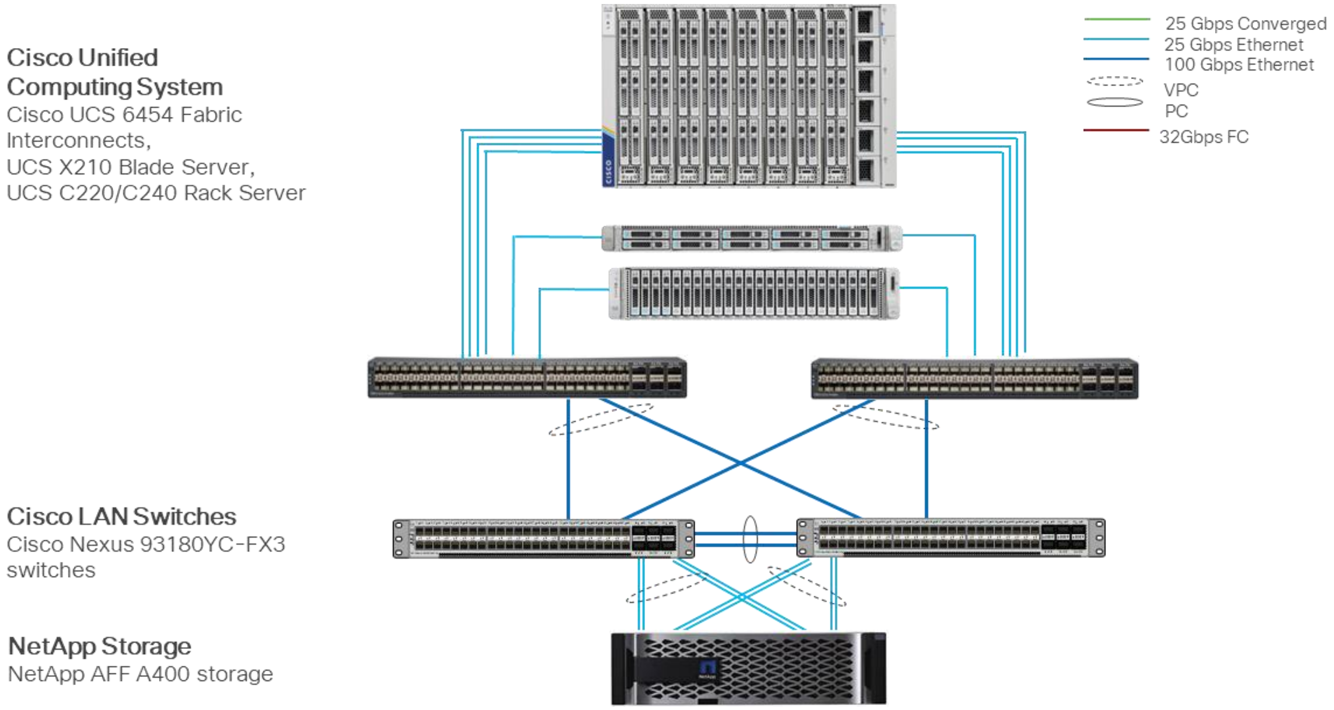

To build an IP only storage access in a FlexPod configuration, the component set up is as follows:

- Cisco UCS Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Chassis connects to fabric interconnects using Cisco UCS 2408 Fabric Extender, where four 25 Gigabit Ethernet ports are used on each IOM to connect to the appropriate Fabric Interconnect. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCSB-200 M6 Compute Nodes contain fourth-generation Cisco 1467 virtual interface cards (VICs).
- The Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCS X9108 Intelligent Fabric Module, where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate Fabric Interconnect. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCSX-210 M6 Compute Nodes contain fourth-generation Cisco 14425 virtual interface cards (VICs).
- Cisco UCS C220 or C240 Servers with fourth-generation VICs 1457 VICs connect to the fabric interconnects with 25GE.
- Cisco Nexus 93180YC-FX Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 25-Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX Switches in a Virtual Port Channel (vPC) configuration.
- The NetApp AFF A400 controllers connect to the Cisco Nexus 93180YC-FX Switches using two 100 GE ports from each controller configured as a vPC.
- Red Hat OpenShift software is installed on Cisco UCS Compute Nodes with local disks to validate the infrastructure.

Note: Red Hat OpenShift Container Platform 4.x does not support the RHCOS installation on iSCSI disks, to install the OS local disks or FC LUNs are required.

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

A typical topology for the FC-booted FlexPod Datacenter is shown in Figure 36.

Figure 36. FlexPod Datacenter Physical Topology for FC, FC-NVMe, and NFS

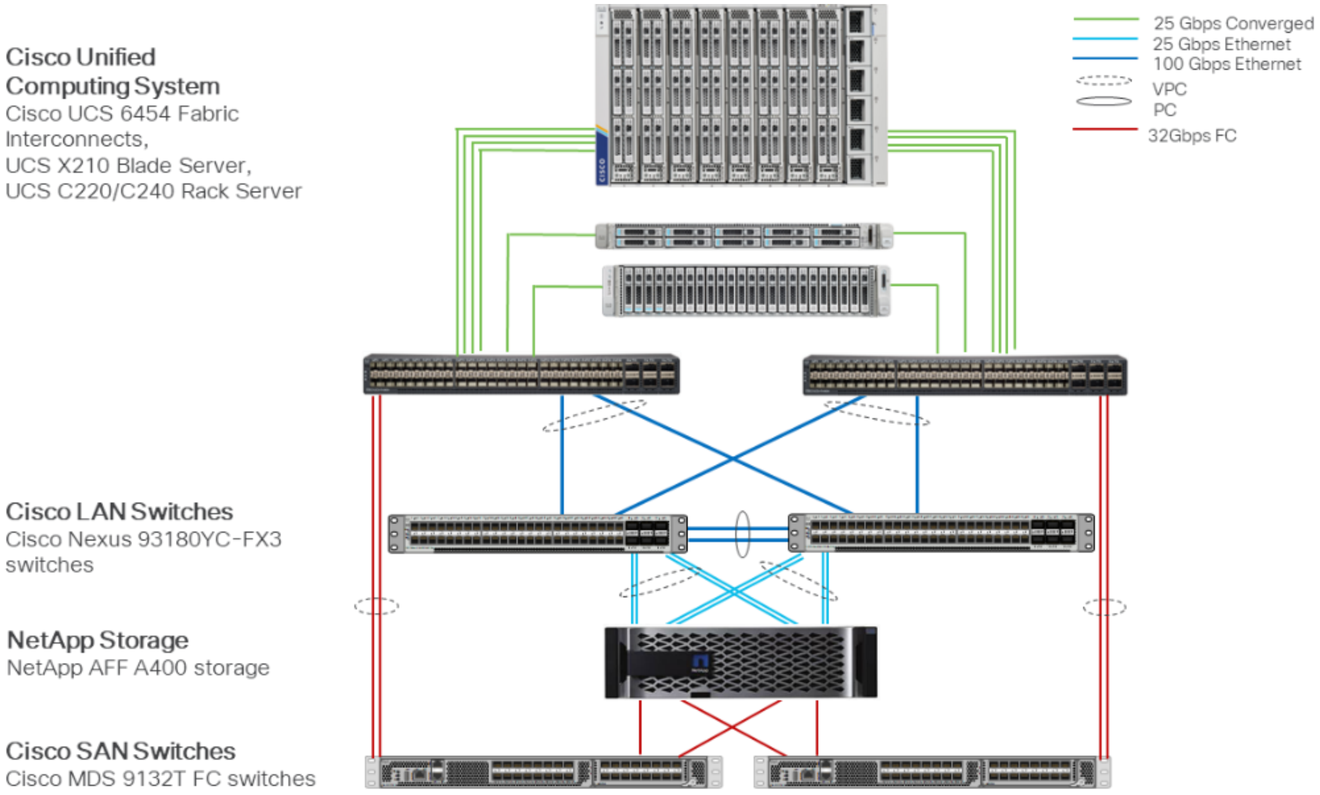

To build an FC-based storage access in a FlexPod configuration, the component set up is as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Chassis connects to fabric interconnects using Cisco UCS 2408 Fabric Extender, where four 25 Gigabit Ethernet ports are used on each FEX to connect to the appropriate FI. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCSB-200 M6 Compute Nodes contain fourth-generation Cisco 1467 virtual interface cards (VICs).
- The Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCS X9108 Intelligent Fabric Module, where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate Fabric Interconnect. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCSX-210 M6 Compute Nodes contain fourth-generation Cisco 14425 virtual interface cards (VICs).
- Cisco UCS C22x or C24x Servers with either fourth-generation VICs or fifth-generation 15428 VICs connect to the fabric interconnects 25GE.

- Cisco Nexus 93240YC-FX2 Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 25-Gigabit Ethernet uplink ports connect to Cisco Nexus 93240YC-FX2 Switches in a Virtual Port Channel (vPC) configuration.
- The NetApp AFF A400 controllers connect to the Cisco Nexus 93240YC-FX2 Switches using two 40 GE ports from each controller configured as a vPC for NFS storage access.
- Cisco UCS 6454 Fabric Interconnects are connected to the Cisco MDS 9132T switches using multiple 32-Gbps Fibre Channel connections configured as a single port channel for SAN connectivity.
- The NetApp AFF controllers connect to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections for SAN connectivity.
- Red Hat OpenShift software is installed on Cisco UCS Compute Nodes with FC LUN to validate the infrastructure.
- Persistent Storage for containerized applications are provides by NetApp Trident CSI via NFS protocol.

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

The typical topology for the FC-boot FlexPod Datacenter with Cisco Nexus SAN Switching is shown in Figure 37.

Figure 37. FlexPod Datacenter Physical Topology for FC, FC-NVMe, and NFS

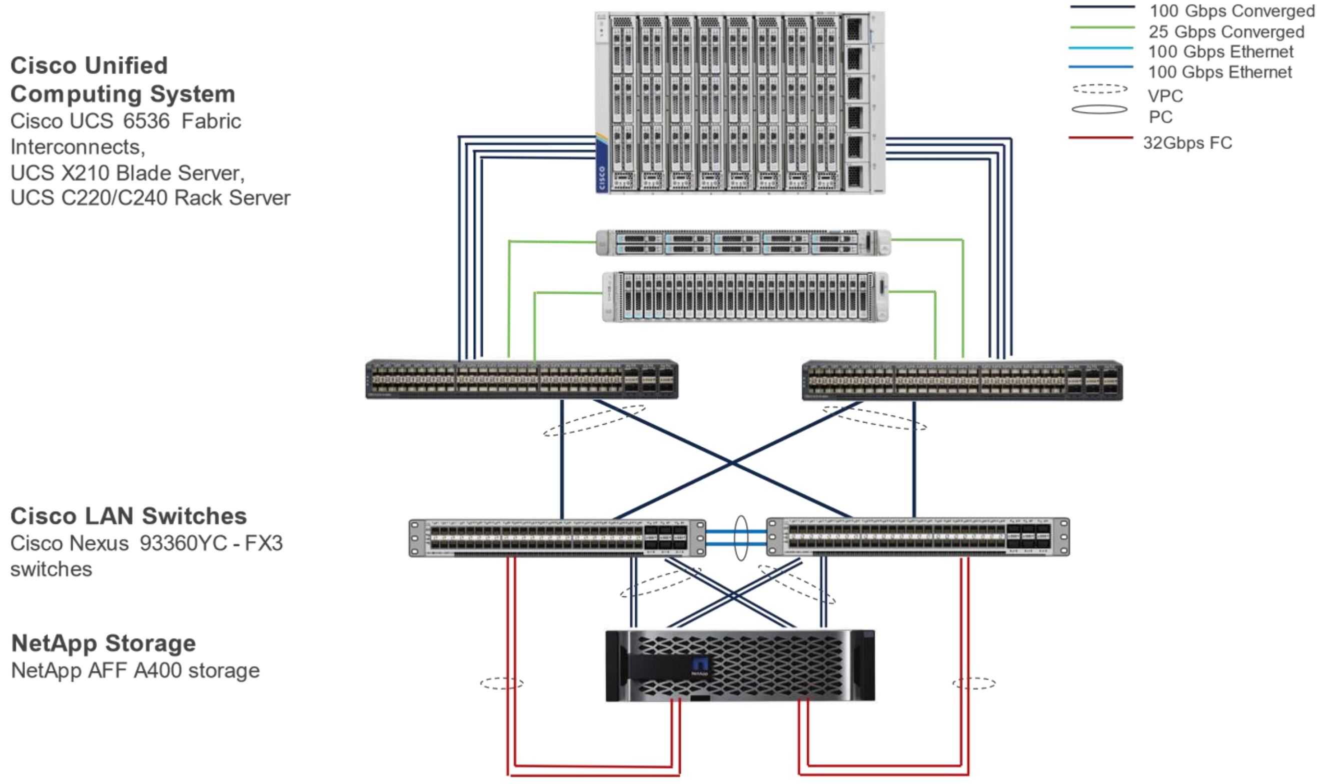

To build a FC-based storage access in a FlexPod configuration with Cisco Nexus SAN switching, the component set up is as follows:

- Cisco UCS 6536 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS X9508 Chassis connects to fabric interconnects using Cisco UCS X9108 Intelligent Fabric Module, where four 25 Gigabit Ethernet ports are used on each IFM to connect to the appropriate Fabric Interconnect. If additional bandwidth is required, all eight 25G ports can be utilized.
- Cisco UCSX-210 M6 Compute Nodes contain fourth-generation Cisco 14425 virtual interface cards (VICs).
- Cisco UCS C22x or C24x Servers with either fourth-generation VICs or fifth-generation 15428 VICs connect to the fabric interconnects 25GE.
- Cisco Nexus 93360YC-FX3 Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6536 Fabric Interconnect 100-Gigabit Ethernet uplink ports connect to Cisco Nexus 93360YC-FX3 Switches in a Virtual Port Channel (vPC) configuration.
- The NetApp AFF A400 controllers connect to the Cisco Nexus 93360YC-FX3 Switches using two 100 GE ports from each controller configured as a vPC for NFS storage access.
- Cisco UCS 6536 Fabric Interconnects are connected to the Cisco Nexus 93360YC-FX3 switches using multiple 32-Gbps Fibre Channel connections configured as a single port channel for SAN connectivity.
- The NetApp AFF controllers connect to the Cisco Nexus 93360YC-FX3 switches using 32-Gbps Fibre Channel connections for SAN connectivity.
- Red Hat OpenShift software is installed on Cisco UCS Compute Nodes with FC LUN to validate the infrastructure.
- Persistent Storage for containerized applications are provides by NetApp Trident CSI via NFS protocol.

VLAN Configuration

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

Table 1. VLAN Usage

VLAN ID	Name	Usage
2	Native-VLAN	Use VLAN 2 as native VLAN instead of default VLAN (1)
1020	OOB-MGMT-VLAN	Out-of-band management VLAN to connect management ports for various devices
1021	IB-MGMT-VLAN	In-band management VLAN utilized for all in-band management connectivity - for example, ESXi hosts, VM management, etc.
1022	OCP-Provisioning	VLAN for OpenShift Provisioning Network
1023	OCP1-Traffic	Data traffic VLAN from/to RH OCP

VLAN ID	Name	Usage
		cluster 1
1024	OCP-VM	Data traffic VLAN from/to RH OCP Virtual Machines
3022	OCP1-NFS	NFS storage traffic VLAN for OCP cluster 1
3050	NFS-VLAN	NFS VLAN for Infrastructure components
3010*	iSCSI-A	iSCSI-A path for boot-from-san traffic
3020*	iSCSI-B	iSCSI-B path for boot-from-san traffic
3030*	NVMe-TCP-A	NVMe-TCP-A path for NVMe datastores
3040*	NVMe-TCP-B	NVMe-TCP-B path for NVMe datastores
3000	vMotion	VMware vMotion traffic

* iSCSI, NVMe-TCP, and vMotion VLANs are optional only.

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

- VLAN 1020 allows customers to manage and access out-of-band management interfaces of various devices and is brought into the infrastructure to allow CIMC access to the UCS servers and is also available to infrastructure virtual machines (VMs). Interfaces in this VLAN are configured with MTU 1500.
- VLAN 1021 is used for in-band management of VMs, hosts, and other infrastructure services. Interfaces in this VLAN are configured with MTU 1500.
- VLAN 3050 provides ESXi hosts, management systems and other infrastructure services access to the NFS storage hosted on the NetApp Controllers. Interfaces in this VLAN are configured with MTU 9000.
- A pair of iSCSI VLANs (3010 and 3020) is configured to provide access to boot LUNs for ESXi hosts and iSCSI datastores. These VLANs are not needed when configuring Fibre Channel connectivity. Interfaces in these VLANs are configured with MTU 9000.
- A pair of NVMe-TCP VLANs (3030 and 3040) is configured to provide access to NVMe datastores. These VLANs are not needed when configuring Fibre Channel connectivity. Interfaces in these VLANs are configured with MTU 9000.
- VLAN 1022 is used as provisioning network for OCP cluster 1. Interfaces in this VLAN are configured with MTU 1500.

- VLAN 1023 is used as access network for OCP cluster 1 to access all OCP hosts, and services deployed on top. Interfaces in this VLAN are configured with MTU 1500.
- VLAN 1024 is used as access network for OCP Virtual machines and services deployed on top. Interfaces in this VLAN are configured with MTU 1500.
- VLAN 3022 provides services deployed on top of OCP cluster 1access to the NFS storage hosted on the NetApp Controllers managed by NetApp Astra Trident CSI. Interfaces in this VLAN are configured with MTU 9000.

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 Datacenter with Red Hat OCP 4 hardware components

Component	Hardware	Comments
Cisco Nexus Switches	Two Cisco Nexus 93000 series switches, such as Cisco Nexus 93240YC-FX	The switch model is dependent on the number of ports and the port speed required for the planned installation.
Cisco MDS Switches	Two Cisco MDS 9100 series switches, i.e. MDS 9132T	The supported port speed of the selected MDS switch must match the port speed of the Fabric Interconnect and the NetApp storage.
NetApp AFF Storage	A NetApp AFF series storage with appropriate capacity and network connectivity, i.e. NetApp AFF A400	Customer requirements will determine the amount of storage. The NetApp AFF storage should support both 25Gbps or 100 Gbps ethernet and 32Gbps or 16 Gbps FC connectivity
Fabric Interconnects	Two Cisco UCS Fabric Interconnects, such as. Cisco UCS 6454 FI	
Management Cluster Compute		
Cisco UCS Servers	A minimum of two Cisco UCS servers to host management components like Intersight Assist and NetApp Active IQ Unified Manager systems.	To reduce the number of physical servers the use of a supported virtualization software like Vmware ESXi is recommended.
Red Hat OCP 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	Six compute nodes are recommended to build an OCP cluster with three control and two worker nodes, but three compute nodes

Component	Hardware	Comments
		will work to build a three node cluster with combined control and worker function.

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 explained in the deployment guide.

Table 3. Software components and versions

Component	Version
Cisco Nexus 93240YC-FX	9.3(10)
Cisco MDS 9132T	8.4(2d)
Cisco UCS Fabric Interconnects	4.2(2c)
Cisco UCS B-Series blade server	4.2(2a)
Cisco UCS X-Series blade server	5.0(2b)
Cisco UCS C-Series rack server	4.2.2(a)
NetApp A400 - ONTAP	9.11.1
NetApp Astra Trident CSI	22.10
Red Hat OpenShift Container Platform	
OCP	4.10
RHCOS	4.10
Red Hat Enterprise Linux	8.6

Logical Design

The Red Hat OpenShift deployment is fully automated. Before starting the deployment process, Cisco UCS compute nodes need to be configured with appropriate compute policies (BIOS, for example), 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 OCP host and the storage configuration on NetApp AFF A400 controllers is captured in the following subsections.

Figure 38 illustrates the FlexPod Datacenter with Red Hat OpenShift Container Platform logical topology with OCP components utilizing compute, network, and storage resources on FlexPod. The storage and network connectivity to Red Hat OpenShift Container Platform Cluster nodes running on Cisco UCS servers is enabled by the Cisco Nexus 9000 series switches within FlexPod.

Persistent storage is a critical part of running stateful containers, and Red Hat OpenShift Container Platform with Kubernetes simplifies storage management by abstracting details of how storage is provisioned and how it is consumed. Persistent volumes for containers can be static or dynamically provisioned, in this case it is dynamic with FlexPod and is enabled by the NetApp Astra Trident CSI Driver. Dynamic volume provisioning allows storage volumes to be created on-demand, NetApp Astra Trident CSI eliminates the need to pre-provision storage for containers and allows persistent storage provisioning during the container deployment. This solution used NFS and iSCSI storage for dynamic storage provisioning.

OpenShift Container Platform uses a software-defined networking (SDN) approach to provide a unified cluster network that enables communication between pods across the OpenShift Container Platform cluster. This pod network is established and maintained by the OpenShift SDN, which configures an overlay network using Open vSwitch (OVS). The default OpenShift SDN solution is built on top of Open vSwitch (OVS). With OpenShift, the cluster admin can choose to deploy with one of the OpenShift native SDN plug-ins or they can opt to deploy the cluster using a third-party SDN from the supported ecosystem such as Cisco ACI. For this solution, we have used the OpenShift native SDN plug-in (OVN-Kubernetes).

Figure 38. **FlexPod Datacenter for Red Hat OCP 4 Bare Metal Logical Topology**

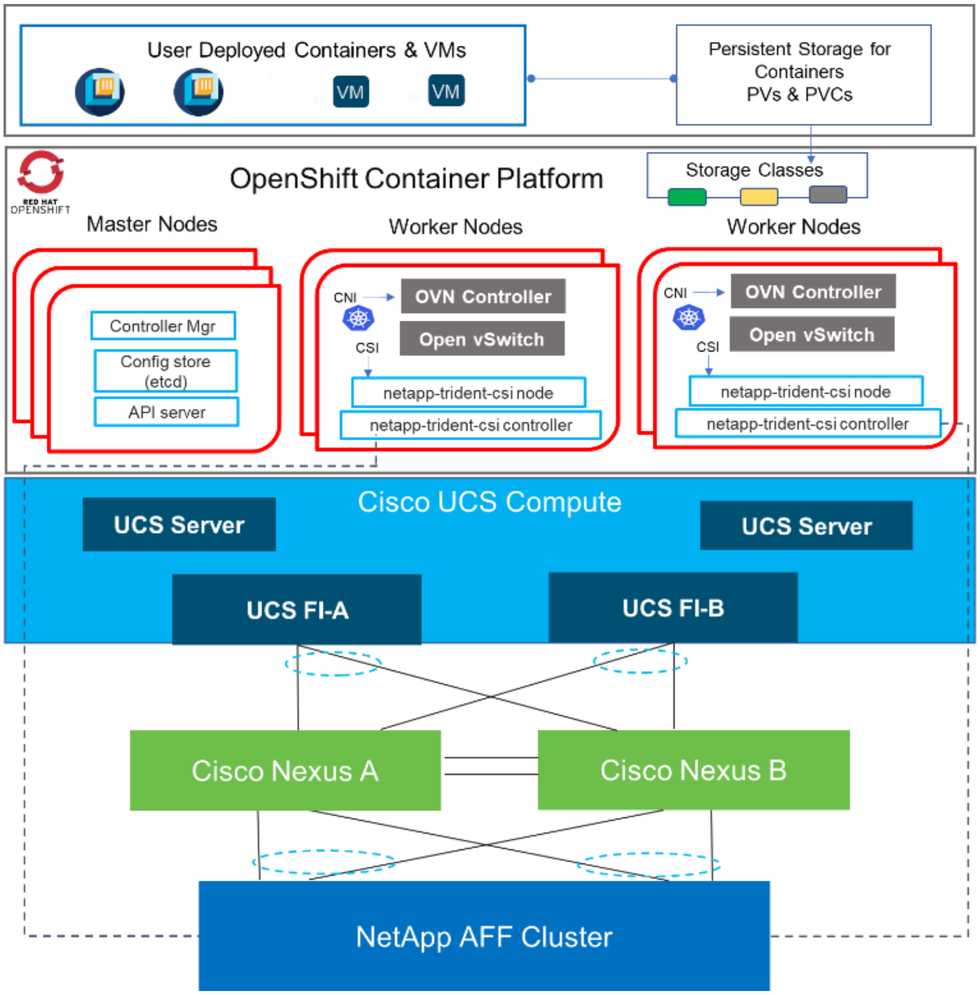

FlexPod Network Connectivity and Design

The Layer 2 network connection to each Fabric Interconnect is implemented as Virtual Port Channels (vPC) from the upstream Cisco Nexus Switches as shown in Figure 39. In the switching environment, the vPC provides the following benefits:

- Allows a single device to use a Port Channel across two upstream devices
- Eliminates Spanning Tree Protocol blocked ports and use all available uplink bandwidth
- Provides a loop-free topology
- Provides fast convergence if either one of the physical links or a device fails
- Helps ensure high availability of the network

The upstream network switches can connect to the Cisco UCS Fabric Interconnects using 10G, 25G, 40G, or 100G port speeds. Virtual port channels were also configured between the Cisco Nexus switches and NetApp AFF storage to also transport the storage traffic between the Cisco UCS servers and the NetApp AFF storage.

Figure 39. Network Connectivity – vPC Enabled Connections

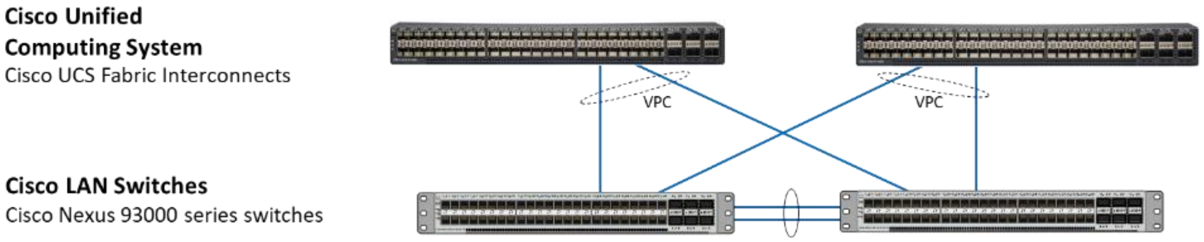
Cisco Unified
Computing System
Cisco UCS Fabric Interconnects
VPC
VPC
Cisco LAN Switches
Cisco Nexus 93000 series switches

NetApp AFF Storage Ethernet Connectivity

NetApp AFF storage controllers are connected with port channels (NetApp Interface Groups) to Cisco Nexus 93000 series switches using 25GE or 100GE connections configured as virtual port channels. The storage controllers are deployed in a switchless cluster interconnect configuration and are connected to each other using the 100GE ports e0a and e1a. Figure 40 illustrates the physical connectivity details.

In Figure 40, the two storage controllers in the high-availability pair are drawn separately for clarity. Physically, the two controllers exist within a single chassis.

Figure 40. NetApp AFF A800 Ethernet Connectivity

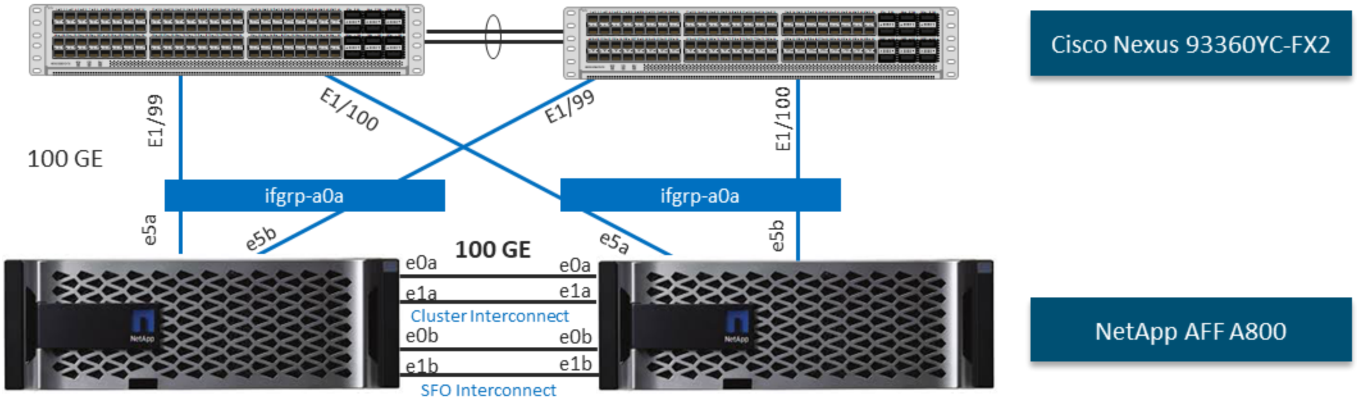

Cisco MDS SAN Connectivity – Fibre Channel Design Only

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

- Feature NPIV—N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port.
- Feature fport-channel-trunk—F-port-channel-trunks allow for the fabric logins from the NPV switch to be virtualized over the port channel. This provides nondisruptive redundancy should individual member links fail.
- Enhanced Device Alias – a feature that allows device aliases (a name for a WWPN) to be used in zones instead of WWPNs, making zones more readable. Also, if a WWPN for a vHBA or NetApp FC LIF changes, the

device alias can be changed, and this change will carry over into all zones that use the device alias instead of changing WWPNs in all zones.

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

Cisco UCS Fabric Interconnect 6536 SAN Connectivity

For SAN connectivity, each Cisco UCS 6536 Fabric Interconnect is connected to a Cisco MDS 9132T SAN switch using a breakout on ports 33-36 to a 4 x 32G Fibre Channel port-channel connection, as shown in Figure 41.

Figure 41. Cisco UCS 6454 FI SAN Connectivity

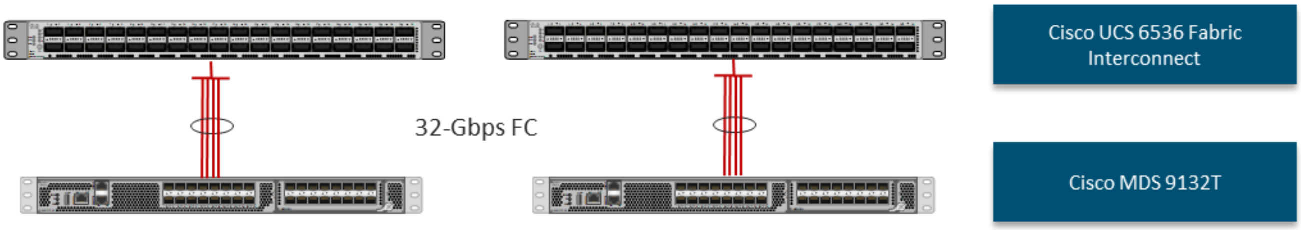

NetApp AFF SAN Connectivity

For SAN connectivity, each NetApp AFF storage controller is connected to both of Cisco MDS 9132T SAN switches using 32G Fibre Channel connections, as shown in Figure 42. FC-NVMe LIFs can be put on the same FC ports on the NetApp storage controllers as FC LIFs.

In Figure 42, the two storage controllers in the high-availability pair are drawn separately for clarity. Physically, the two controllers exist within a single chassis.

Figure 42. NetApp AFF A800 SAN Connectivity

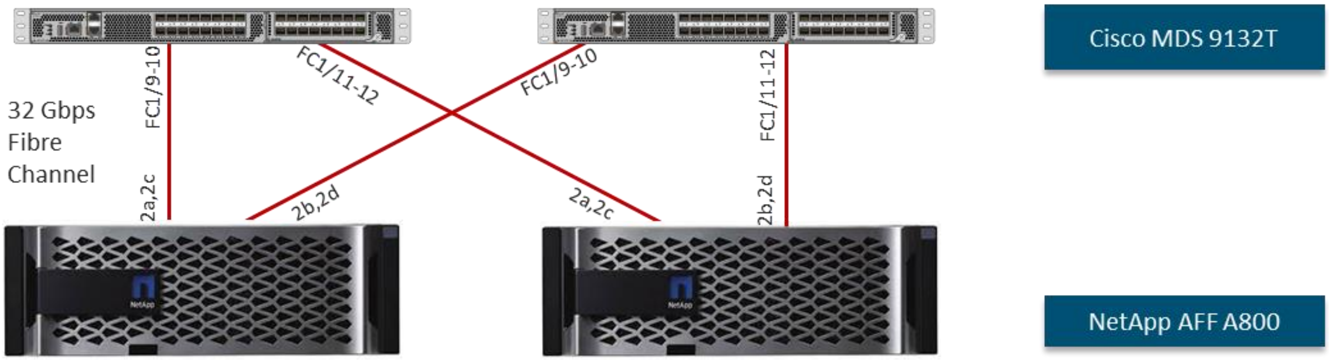

FlexPod Compute Connectivity

The FlexPod compute design supports Cisco UCS X-Series, B-Series, and C-Series. Cisco UCS supports the Red Hat OpenShift environment by providing robust, highly available, and integrated compute resources centrally managed from Cisco UCS Manager in the Enterprise or from Cisco Intersight Software as a Service (SaaS) in the cloud. In this validation effort, the Cisco UCS servers are booted from local SSDs, these drives are configured in Raid 1 using a Cisco Raid Controller. Or the Cisco UCS servers are booted form a FC LUN provided from the NetApp AFF storage for deployments with. The servers have access to NFS and iSCSI storage for persistent storage volumes presented from the NetApp AFF storage cluster.

Figure 43. Compute Connectivity

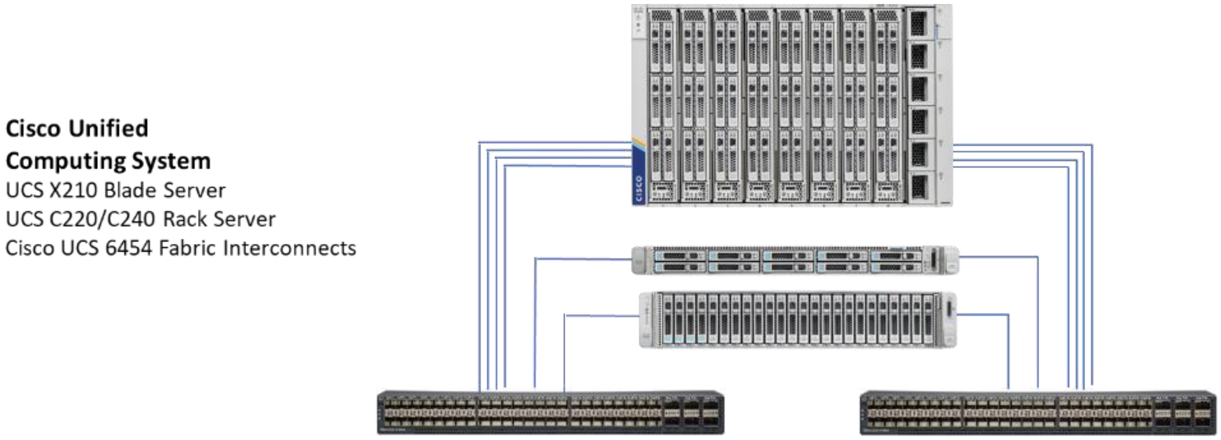
Cisco Unified
Computing System
UCS X210 Blade Server
UCS C220/C240 Rack Server
Cisco UCS 6454 Fabric Interconnects

The Cisco UCS chassis in the design are populated with Cisco UCS X210 M6 blade servers and each of these blade servers contain one physical network adapter (VIC) that passes converged fibre channel over Ethernet (FCoE) and Ethernet traffic to the X9108-25 IFM. The IFMs are redundantly connected to the fabric interconnects using 4X25Gbps ports per IFM to deliver an aggregate bandwidth of 200Gbps to the chassis. Full population of each X9108-25 IFM can support 8x25Gbps ports, providing an aggregate bandwidth of 400Gbps to the chassis. The connections from the Cisco UCS Fabric Interconnects to the IFMs are automatically configured as port channels.

The Cisco UCS C-Series nodes are equipped with Cisco UCS VIC 1457 or Cisco UCS PCIe VIC 1455. Cisco UCS VIC 1455/1457 has four 25GbE ports which are connected to the Cisco UCS Fabric Interconnect in pairs such that ports 1 and 2 are connected to the Cisco UCS FI-A and the ports 3 and 4 are connected to the FI-B as shown in Figure 43. Optionally, only ports 1 and 3 from each VIC 1455/57 and be connected with network bandwidth reduced from 50 Gbps to 25 Gbps to each FI.

Fabric Failover for Ethernet: Architecture for High Availability

Cisco UCS is designed for high availability, with no single point of failure in its network infrastructure. Each adapter in Cisco UCS connects to both fabrics (A and B) and the fabric interconnects are designed to work in an active-active model, with automated failover of network in the event of a failure. The system is designed so that if either fabric A or fabric B fails, the remaining fabric can take on the traffic from the failed fabric. Cisco VICs support fabric failover by moving traffic from one fabric to the other according to failover policies established on a per-vNIC basis. This eliminates complicated operating system NIC teaming configurations. Fabric failover makes the failure of a fabric transparent to the operating system.

Figure 44 illustrates the Cisco UCS Fabric failover mechanism, in this example one vNIC connects to fabric A but fails over to fabric B in the event of a fabric failover.

Cisco UCS fabric failover is an important feature because it reduces the complexity of defining NIC teaming software for failover on the host. It does this transparently in the fabric based on the network property that is defined in the service profile. With Cisco UCS fabric failover, NIC teaming is not necessary on the OCP nodes, and the high availability is managed at the UCS level more efficiently.

Figure 44. Cisco UCS Fabric Failover

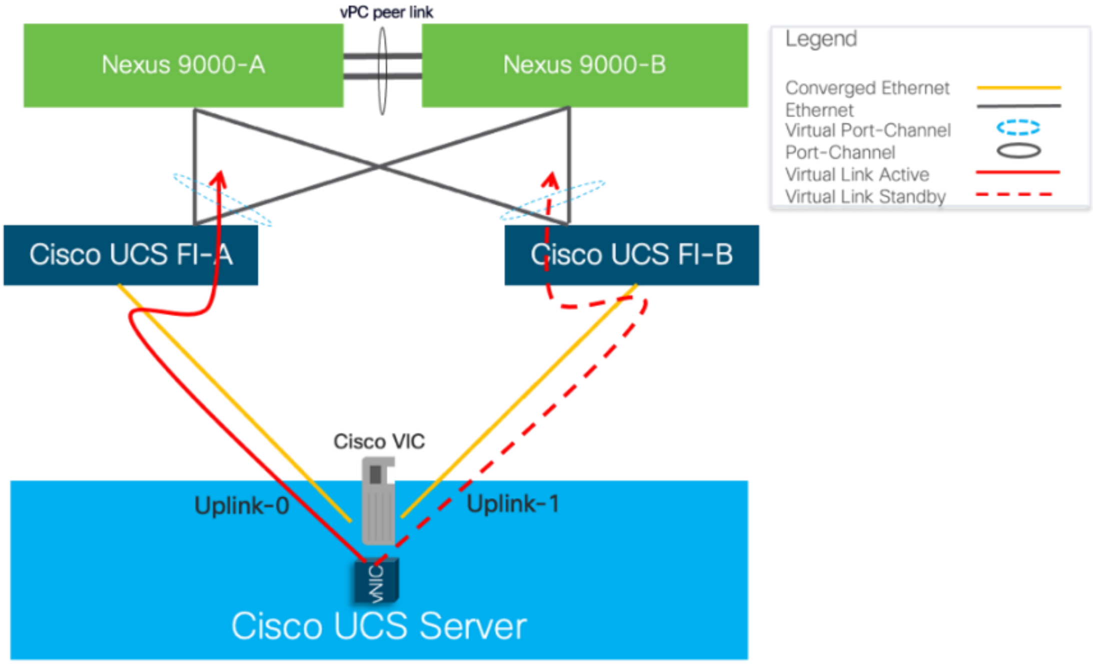

NIC Bonding versus Cisco UCS Fabric Failover

OpenShift Container Platform network requirements in this design are standard Ethernet only, while OCP 4 deployment can work with two network interfaces in bonded mode for each traffic type (bare metal public VLAN and VM network VLAN), it is recommended to use a single network interface for each traffic type and enable Cisco UCS Fabric Failover for resiliency versus NIC bonding in the operating system. With Cisco UCS Fabric Failover the management and operation of failover and link aggregation is handled in the networking fabric. The Fabric Failover is enabled in the vNIC policy within the LAN Connectivity Policy which makes it easy to implement NIC resiliency across any number of servers managed by Cisco UCS, this eliminates the need to configure every server individually.

NIC teaming is often implemented to aggregate lower-speed NICs in order to gain throughput. Since OCP design with Cisco UCS leverages 25/100GbE connections, aggregation is generally not required.

Intersight Server Profile for OCP Hosts

In FlexPod deployments, each Cisco UCS server (X-Series, B-Series, or C-Series), equipped with a Cisco Virtual Interface Card (VIC), is configured for multiple virtual interfaces (vNICs) which appear as standards-compliant PCIe devices to the OS. The server profile configuration for an OCP host is shown in Table 4, Figure 45, and Figure 46 for OCP Worker and Master nodes, respectively.

Each OCP host server profile supports:

- Managing the OCP hosts using a common management segment
- OS local boot using the mirrored onboard disk drives on the Cisco UCS servers or boot from FC LUN (boot from iSCSI was not supported with OCP nodes at the time of this validation).

- Six vNICs (NFS and iSCSI vNICs are only required on the Worker Nodes) used as follows in the same order specified below:
- One vNIC for provisioning traffic to support OCP installer provisioned infrastructure. The MTU value for this interface is set as a Jumbo MTU (9000). This should be the first vNIC on the UCS servers and this supports provisioning network which is non-routable network used for provisioning the underlying operating system (RHCOS) on each node that is part of the OpenShift Container Platform Cluster.
- One vNIC for OCP Public Bare Metal Network traffic. The bare metal network is a routable network. The second vNIC on the UCS servers is used to support the bare metal network.
- One vNIC for OpenShift Virtualization VM Network traffic.
- One NFS vNIC for NFS storage traffic. The MTU value for this interface is set as a Jumbo MTU (9000).
- One iSCSI-A vNIC utilizes iSCSI-A VLAN (defined only on Fabric A) to provide access to iSCSI-A path. The MTU value for this interface is set as a Jumbo MTU (9000).
- One iSCSI-B vNIC utilizes iSCSI-B VLAN (defined only on Fabric B) to provide access to iSCSI-B path. The MTU value for this interface is set as a Jumbo MTU (9000).

Table 4. OCP Host Service Profile

Machine	Provisioning Traffic	BareMetal Public Traffic	OpenShift Virtualization	NFS	iSCSI-A	iSCSI-B
Provisioner Node	vNIC1 Fabric-A failover to Fabric-B	vNIC2 Fabric-B failover to Fabric-A	vNIC3 Fabric-A failover to Fabric-B	vNIC4 Fabric-B failover to Fabric-A	vNIC5 Fabric-A only	vNIC6 Fabric-B only
Master Node	vNIC1 Fabric-A failover to Fabric-B	vNIC2 Fabric-B failover to Fabric-A	N/A	N/A	N/A	N/A
Worker Node	vNIC1 Fabric-A failover to Fabric-B	vNIC2 Fabric-B failover to Fabric-A	vNIC3 Fabric-B failover to Fabric-A	vNIC4 Fabric-B failover to Fabric-A	vNIC5 Fabric-A only	vNIC6 Fabric-B only

Figure 45. Cisco UCS – Network Interface Design for OCP Worker Nodes

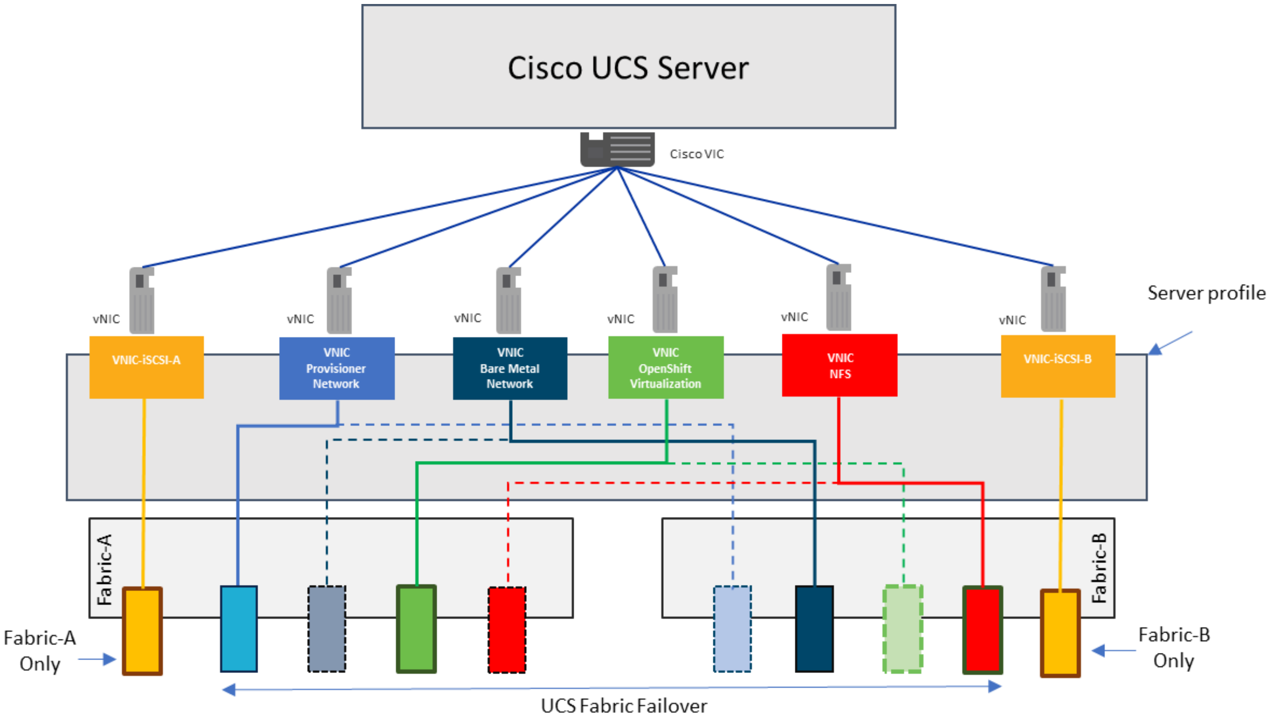

Figure 46. Cisco UCS – Network Interface Design for OCP Master Nodes

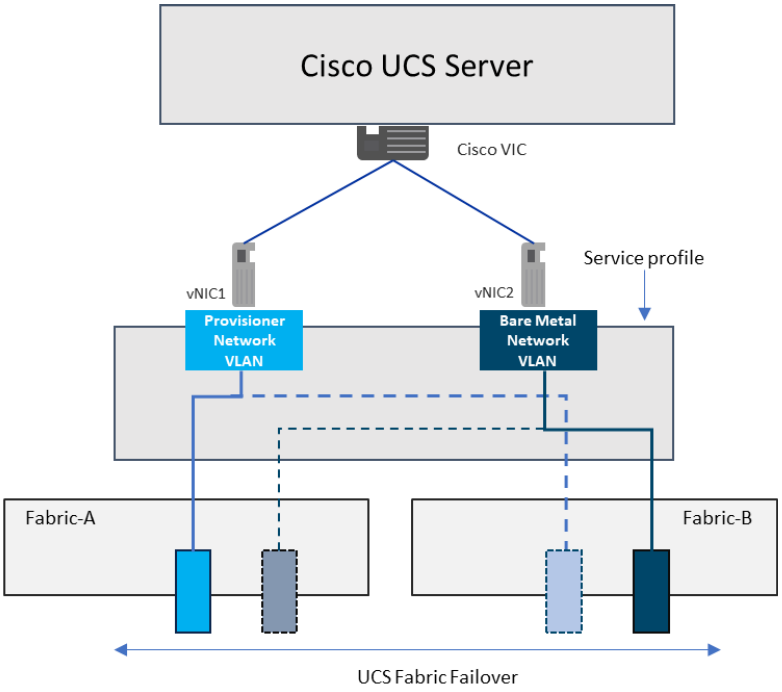

FlexPod Storage Design for OCP

The FlexPod Datacenter for Red Hat OpenShift Container Platform uses the NetApp Astra Trident CSI driver that is an add-on component that needs to be installed on the OpenShift Container Platform cluster. Astra Trident enables the integration between the storage and OCP cluster.

The NetApp AFF series storage array supports both NFS and iSCSI protocols. For the purpose of this validated design, both NFS and iSCSI were used for dynamic persistent storage for Containers and VMs.

Dynamic Storage Provisioning

Note: OpenShift provides dynamic provisioning of storage for applications by utilizing the StorageClass resource. Using dynamic storage, you can select different types of back-end storage. The back-end storage is segregated into different tiers depending on the needs of your application. When requesting storage, you can specify a PersistentVolumeClaim with an annotation that specifies the value of the StorageClass they prefer.

To order the storage, you must create a PVC. The PVC determines the specification for the storage that you want to provision. After the PVC is created, the storage device and the PV are automatically created for you.

Figure 47 illustrates how block storage is dynamically provisioned in a cluster. This sample flow works similarly with other storage types, such as file storage.

Figure 47. Dynamic Storage Provisioning Workflow

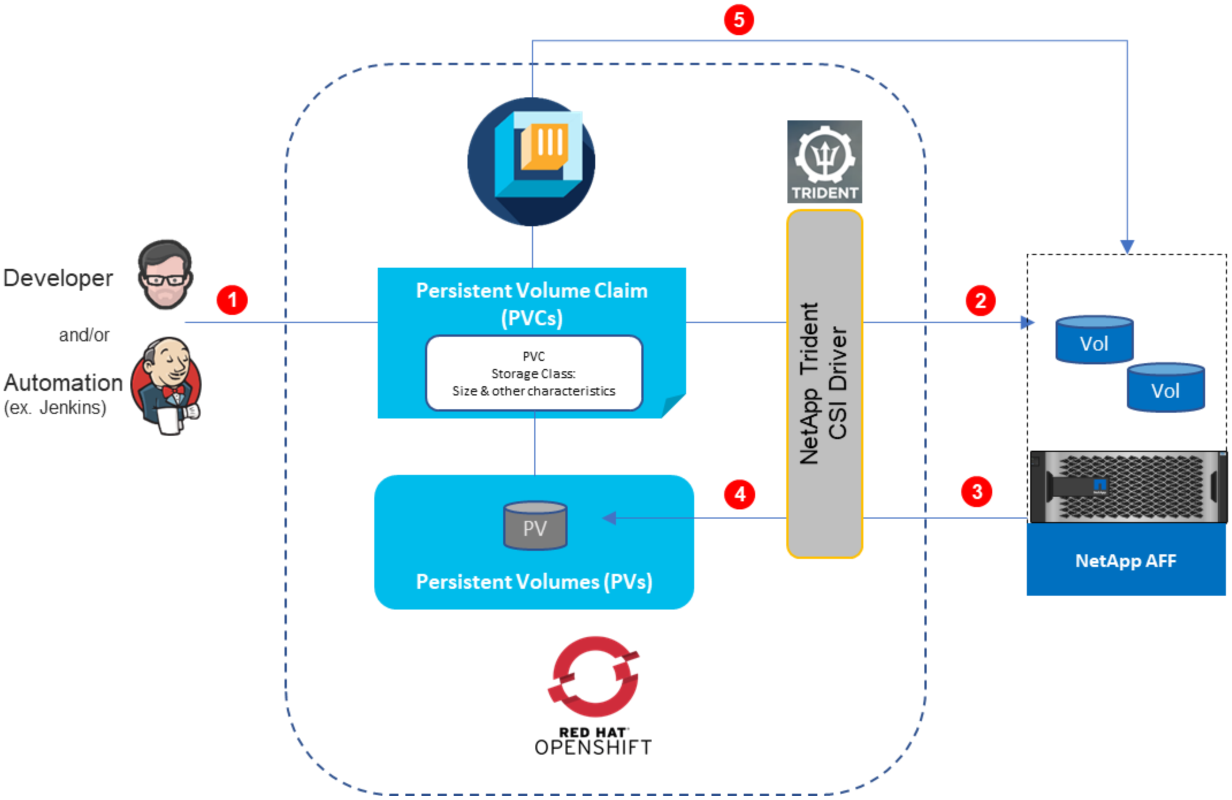

Developer/Automation submits storage requirements in the form of standard Persistent Volume Claims that specifies the storage type, storage class, size, and so on.

NetApp Astra Trident CSI Plugin listens to, and intercepts Persistent Volume Claims based on Storage Class. Creating a PVC in a cluster automatically triggers the storage plug-in for the requested type of storage to provision storage with the given specification.

Storage provisioning API call sent to NetApp AFF, and storage is provisioned.

The storage plug-in automatically creates a persistent volume (PV) in the cluster, a virtual storage device that points to the actual storage device on your NetApp AFF.

The PVC and PV are automatically connected to each other. The status of the PVC and the PV changes to Bound and the PVC is used to mount persistent storage to your app. If you delete the PVC, the PV and related storage instance are also deleted.

OCP Virtual Switching Architecture

The OpenShift Container Platform cluster uses a virtualized network for pod and service networks. The OVN-Kubernetes Container Network Interface (CNI) plug-in is a network provider for the default cluster network. A cluster that uses the OVN-Kubernetes network provider also runs Open vSwitch (OVS) on each node. OVN configures OVS on each node to implement the declared network configuration.

The OVN-Kubernetes default Container Network Interface (CNI) network provider implements the following features:

- Uses OVN (Open Virtual Network) to manage network traffic flows. OVN is a community developed, vendor agnostic network virtualization solution.
- Implements Kubernetes network policy support, including ingress and egress rules.
- Uses the Geneve (Generic Network Virtualization Encapsulation) protocol rather than VXLAN to create an overlay network between nodes.

Figure 48 shows the distribution of network interfaces on each OCP worker node with one Open vSwitch bridge, for Pod-to-Pod communication and the other dedicated Linux bridge created for VM external network when VMs are deployed leveraging the OpenShift Virtualization feature. Each bridge has one NIC within the OS providing access to external networks. The other four network interfaces on the OCP nodes are used for communication via the provisioning network and access to NFS, iSCSI-A, and iSCSI-B traffic via dedicated interfaces. Appropriate VLANs are enabled at the UCS level to support different traffic types.

Figure 48. Virtual Switching and Connectivity Diagram for a Cisco UCS OCP Host

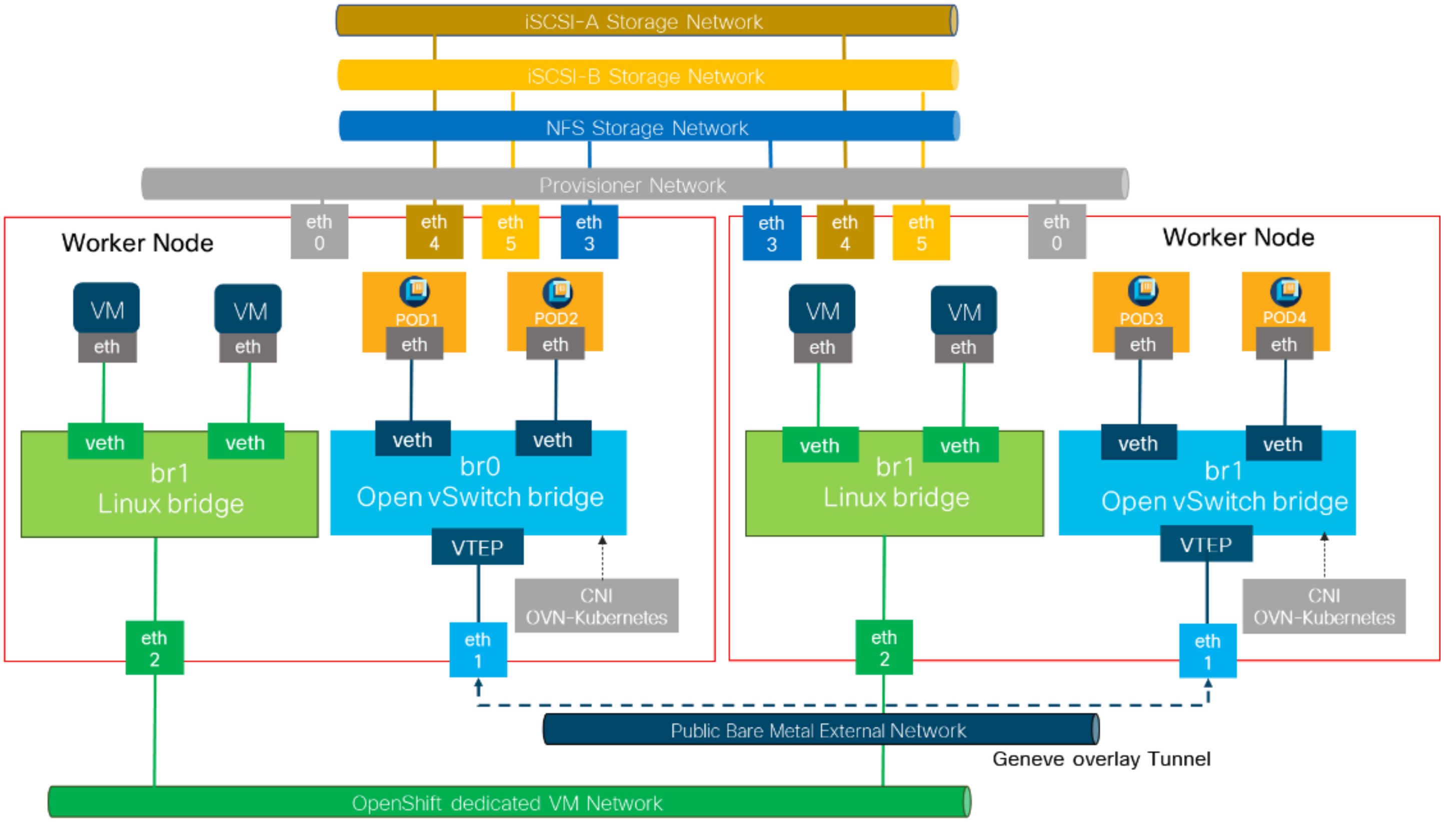

With OpenShift Virtualization, each VM deployed is controlled via a virt-launcher pod that is created with each VM. The default networking type for OpenShift Virtualization VMs is Masquerade. The VM will be assigned a non-routable IP and you can access the VM using the IP of the virt-launcher pod that was deployed alongside it. This makes the VM accessible in the same way that containers are accessed.

Alternatively, you can connect the VM to the host network by creating a bridge interface on the OCP nodes using Nmstate. The Nmstate operator is installed with OpenShift Virtualization and provides you with the Node Network Configuration Policy (NNCP) object to update the host network settings. Figure 48 has a sample config bridge called br1 created from an interface called eth2 (the interface name differs based on how the host views it) on the OCP nodes.

Bare Metal Compute Options and Sizing

The validation of FlexPod Datacenter with Red Hat OpenShift Container Platform was done with various UCS server models and is not limited to a specific set of server types. It is important to check if the used operating system release is supported on the server and firmware version available in the Cisco UCS Hardware Compatibility List (HCL), and if one of the validated boot options (local disk or FC LUN) are possible.

Table 5 lists some Cisco UCS compute options that are tested for OCP installation.

Table 5. Cisco UCS Server Node Configuration Options

Server Node	CPU	Boot storage

Server Node	CPU	Boot storage
Cisco UCS B200 M5	2x 2nd Gen Intel® Xeon® Scalable	RAID1 on local SSD, FC LUN
Cisco UCS C220 M5	2x 2nd Gen Intel® Xeon® Scalable	RAID1 on local SSD, FC LUN
Cisco UCS C225 M5	2x 3rd Gen AMD EPYC	RAID1 on local SSD, FC LUN
Cisco UCS B200 M6	2x 3rd Gen Intel® Xeon® Scalable	RAID1 on local SSD, FC LUN
Cisco UCS X210 M6	2x 3rd Gen Intel® Xeon® Scalable	RAID1 on local SSD, FC LUN

Sizing

This is a general recommendation and not specific to a customer environment. It is important to properly size the solution with all of its components by a qualified Engineer or Architect per the specific requirements of the customer. There is no one size fits all approach, hence specific sizing and performance testing were excluded from the validation process.

For example, at the Cisco UCS level, customers have the option to include servers with different processors and core counts, and with the combination of the right amount of memory the servers can be optimized for the right cost-performance configuration. The same strategy is applicable across all the layers of FlexPod including network and storage.

It is important to size the servers to meet the minimal requirements of the OCP platform, to account for failures of servers and by that to make sure that OCP HA related rules can be followed upon server failure with enough resources available for OCP to redistribute the workloads from the failing host or when performing upgrades and other maintenance tasks.

Example Sizing Guidelines (Worker Nodes)

Determine how many nodes and pods are required for your OpenShift Container Platform cluster. Cluster scalability correlates to the number of pods in a cluster environment. That number influences the other numbers in your setup. See Cluster Limits for the latest limits for objects in OpenShift Container Platform.

Environment sizing can be done according to tested cluster maximums or according to your application requirements. While planning your environment, determine how many pods are expected to fit per node:

```
Required Pods per Cluster / Pods per Node = Total Number of Nodes Needed
```

If you want to scope your cluster at 3000 pods, assuming the 500 maximum pods per node, you will need at least ten nodes:

```
3000 / 500 = 6
```

If you increase the number of nodes to 8, the pods per node distribution changes to 375 pods per node.

The current maximum number of pods per node is 500. However, the number of pods that fit on a node is dependent on the application itself. Consider the application’s memory, CPU, and storage requirements.

Table 6 lists components you might consider for a sample application environment.

Table 6. Environment Components

Pod type	Pod quantity	Max memory	CPU cores	Persistent storage
apache	100	500 MB	0.5	1 GB
node.js	200	1 GB	1	1 GB
postgresql	100	1 GB	2	10 GB

The overall resource requirements for this application are: `450 CPU cores, 350GB RAM, and 1.3TB storage` plus overhead required for OCP operations.

Secure Multi-Tenancy / Multi-Tenancy

Secure Multi-Tenancy defines the deployment option to share the infrastructure resource across multiple tenants/customer/user and separate them so that tenant 1 do have no access to any resources of tenant 2. Such as, tenant 1 is unable to access any storage resources or files of any other tenant stored on the same physical storage. The FlexPod architecture provides all required functions to enable network and storage separation on a logical level. In the example shown in Figure 49, each OCP cluster do have dedicated networks – meaning VLAN IDs and IP ranges, and a dedicated virtual storage server (vserver). All storge networks (iSCSI and NFS) are not routed, meaning not accessible from outside of the configured VLAN ID. The external network and VM network are either routed and controlled by access control lists or firewall rules or mapped to transfer network to the public cloud or remote sites to enable hybrid-cloud operations. To allows central management of all components the OCP nodes and storages are connected to the IB-Management network. In case of management dedication – each tenant is maintaining the OCP nodes – the IB-Management network can be replaced by dedicated tenant management networks.

Figure 49. Secure Multi-Tenancy Option – Multiple Red Hat OCP clusters

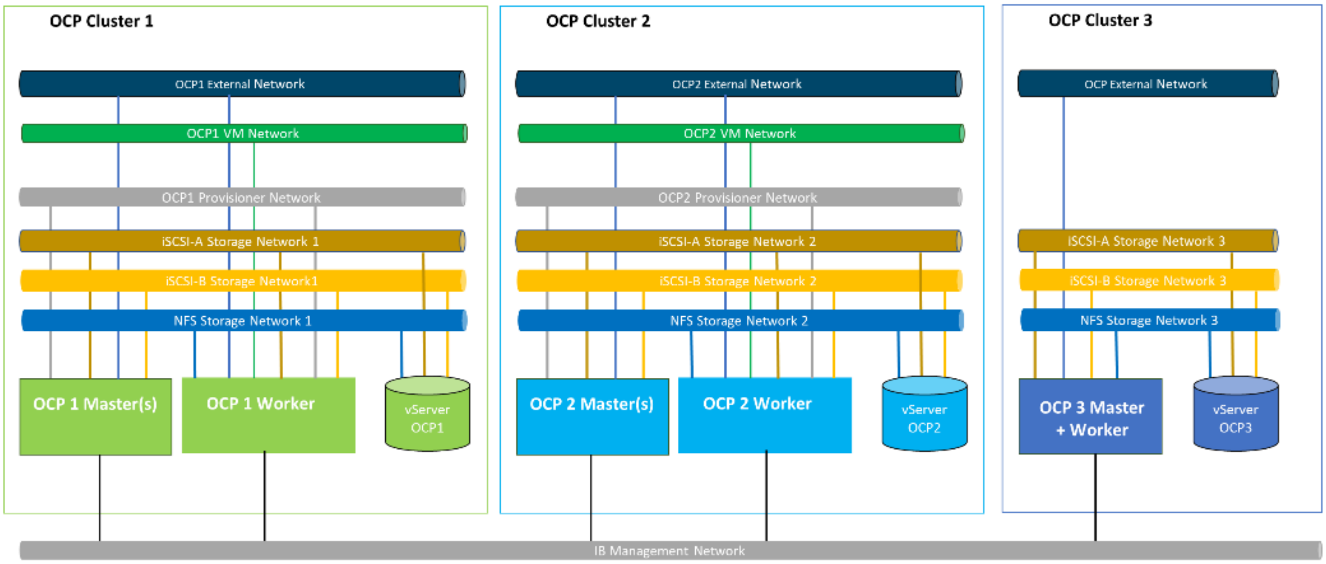

Another option to provide multi-tenancy is combining network and storage separation with RedHat OCP configurations as shown in Figure 50. This option can use a single set of OCP Master nodes to control multiple sets of worker nodes and manage the separation by labels and tags.

Note: As there is only one set of OCP master nodes, all set of worker nodes must match the OCP version requirements of this master.

Figure 50. Secure Mutli-Tenancy Option – Single Red Hat OCP Cluster with Worker Node Sets

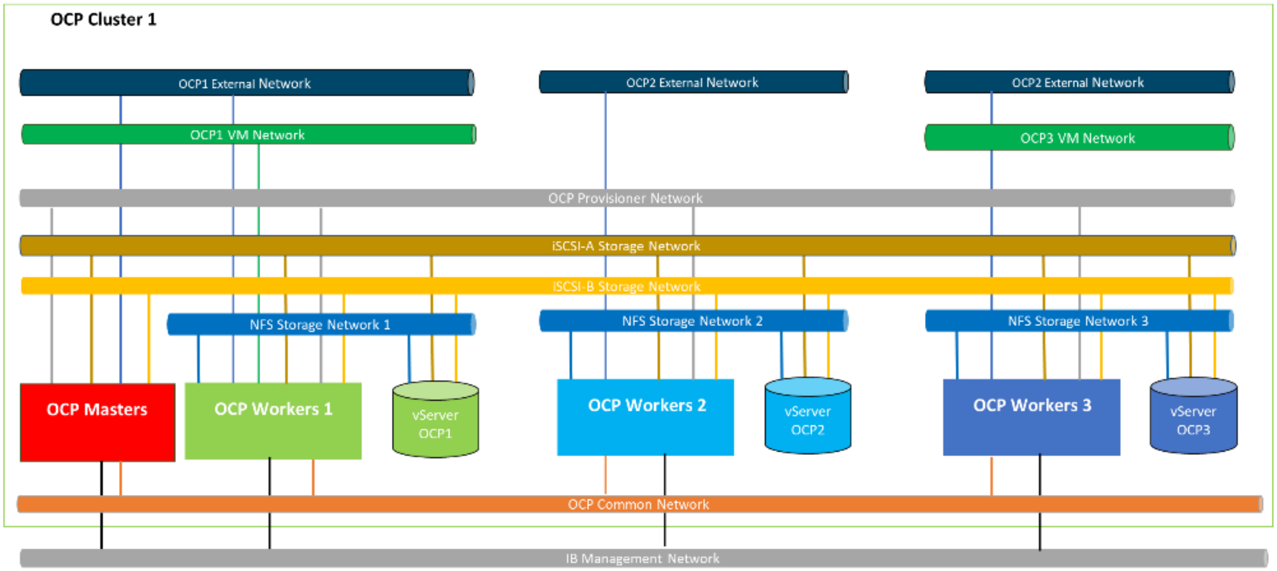

Hybrid-Cloud

The terminology hybrid-cloud can be used for more than one type of solution. One of them is managing hardware deployed on-premises together with resources from the public cloud from a single as a Service “aaS” solution. Or managing workloads deployed at on-premises stack and in the public cloud like production system in the local data center and test or development systems in the public cloud. The FlexPod Datacenter with Red Hat OpenShift Container Platform solution works with both options. Cisco Intersight can manage the FlexPod hardware in local data centers and remote locations together with resources from the public cloud like AWS EC2 instances and Red Hat Hybrid Console is the single place to manage OCP deployments on bare metal infrastructure, virtualized or in the public cloud.

The key aspect of managing distributed workloads or multiple application instances across on-premises and public clouds is the network in between. Deploying and managing the virtual private network (VPN) connection between the locations and control access is an important aspect of every hybrid cloud deployment. This FlexPod Datacenter with RedHat OpenShift solution is tested with two key networking solutions from Cisco. The first is Cisco Application Centric Infrastructure (ACI) combined with Cisco Cloud APIC – shown in Figure 51, and the second is Cisco Nexus Dashboard Fabric Controller (NDFC) together with Cisco Nexus Dashboard Orchestrator (NDO) and the Cisco Cloud APIC to deploy and manage workload dependent network connections. With those two options a End-2-end tested, and documented solution is available. There are other networking options available which will work but they will not be tested and documented as part of this solution.

Figure 51. Sample Hybrid Cloud setup - ACI Managed Network Connections

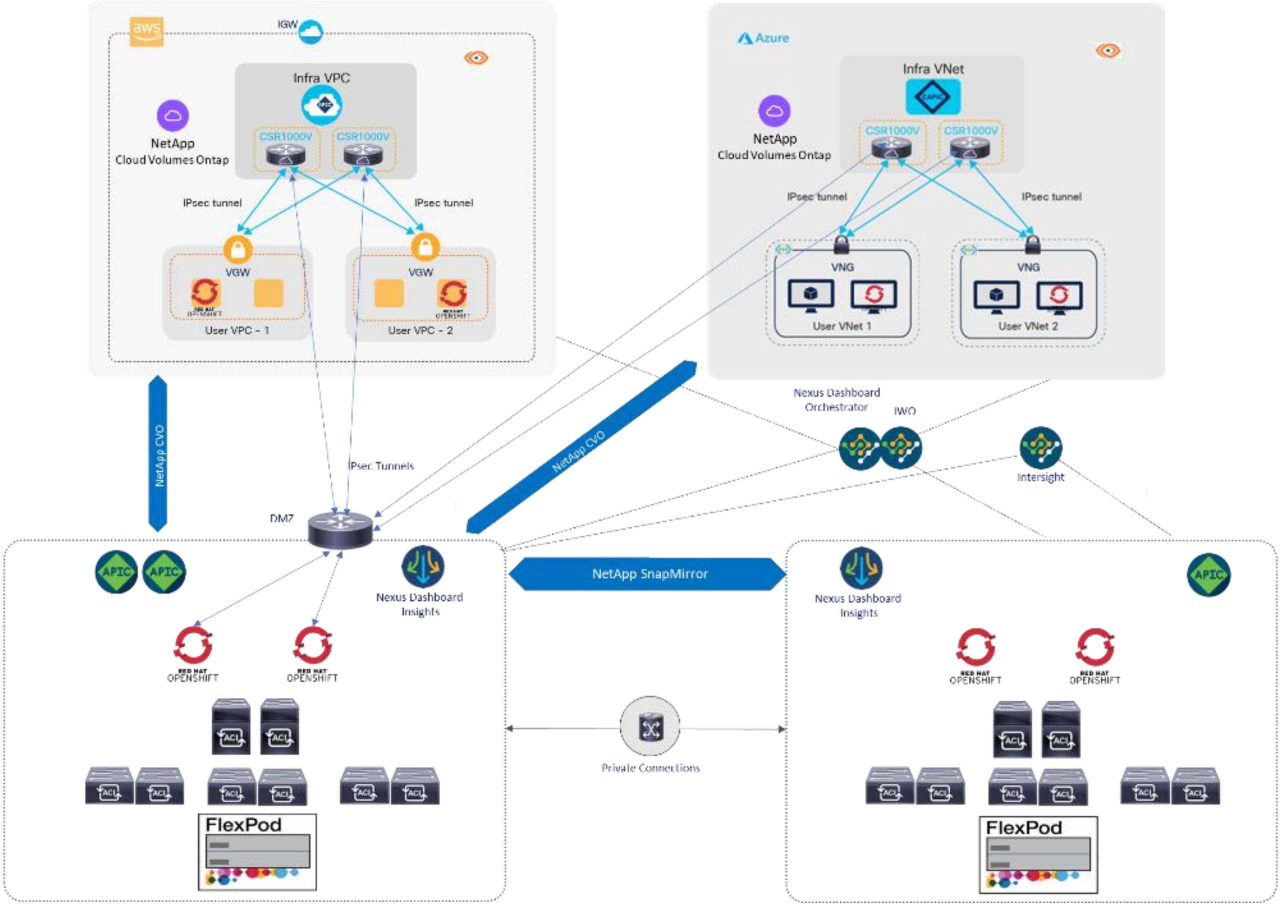

Management Design Considerations

Out-of-band Management Network

The management interface of every physical device in FlexPod is connected to a dedicated out-of-band management switch which can be part of the existing management infrastructure in a customer's environment. The out of band management network provides management access to all the devices in the FlexPod environment for initial and on-going configuration changes. The routing and switching configuration for this network is independent of FlexPod deployment and therefore changes in FlexPod configurations do not impact management access to the devices. In this design, the out-of-band management network is connected to the Cisco Nexus uplinks to allow Cisco UCS CIMC connectivity and to provide the out-of-band management network to management virtual machines (Cisco DCNM) when necessary.

In-band Management Network

The in-band management VLAN configuration is part of FlexPod design. The in-band VLAN is configured on Cisco Nexus switches and Cisco UCS within the FlexPod solution to provide management connectivity for management components deployed on this solution or on a dedicated management system. The changes to FlexPod configuration can impact the in-band management network and misconfigurations can cause loss of access to the management components hosted by FlexPod. It is also required that the out-of-band management network have

Layer 3 access to the in-band management network so that management machines with only in-band management interfaces can manage FlexPod hardware devices.

Jumbo Frames

An MTU of 9216 is configured at all network levels to allow jumbo frames as needed by the guest OS and application layer. This allows the network at every point to negotiate an MTU up to 9000 with the end point. For VLANs that leave the FlexPod via the Cisco Nexus switch uplinks (OOB-MGMT, IB-MGMT, OCP-VM), all endpoints should have MTU 1500. For Storage and replication VLANs that stay within the FlexPod, MTU 9000 should be used on all endpoints for higher performance. It is important that all endpoints within a VLAN have the same MTU setting. It is important to remember that most virtual machine network interfaces have MTU 1500 set by default and that it may be difficult to change this setting to 9000, especially on a large number of virtual machines. This difficulty should be considered when implementing storage protocols such as NFS or SMB. Note that a VLAN tagged trunk can contain both VLANs with MTU 1500 and VLANs with MTU 9000 interfaces.

NTP

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

Boot From SAN

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

UEFI Secure Boot

This validation of FlexPod uses Unified Extensible Firmware Interface (UEFI) Secure Boot. UEFI is a specification that defines a software interface between an operating system and platform firmware. With UEFI secure boot enabled, all executables, such as boot loaders and adapter drivers, are authenticated as properly signed by the BIOS before they can be loaded. Additionally, a Trusted Platform Module (TPM) is also installed in the Cisco UCS compute nodes. Red Hat Core OS supports UEFI Secure Boot and UEFI Secure Boot Attestation between the TPM module and the OS, validating that UEFI Secure Boot has properly taken place.

NetApp Astra Trident

Astra Trident is an open-source and fully supported storage orchestrator for containers and Kubernetes distributions, including Red Hat OpenShift. Trident works with the entire NetApp storage portfolio, including the NetApp ONTAP storage systems, and it also supports NFS and iSCSI connections. Trident accelerates the DevOps workflow by allowing end users to provision and manage storage from their NetApp storage systems without requiring intervention from a storage administrator.

An administrator can configure a number of storage backends based on project needs and storage system models that enable advanced storage features, including compression, specific disk types, or QoS levels that guarantee a certain level of performance. After they are defined, these backends can be used by developers in their projects to create persistent volume claims (PVCs) and to attach persistent storage to their containers on demand

Deployment Documents

Upcoming Publications

FlexPod Datacenter for Red Hat OpenShift Container Platform with NetApp Astra

Cisco UCS M5 and M6, RedHat OCP 4.10, NetApp Ontap 9.11, NetApp Astra

Cisco Validated Design describing the deployment of FlexPod Datacenter with Red Hat OpenShift based on Cisco Intersight with the option of automation with Red Hat Ansible and network connection to the public cloud for Hybrid-cloud operation. In addition the use of NetApp Astra Control Center to protect the containerized workloads is covered in the document.

FlexPod Datacenter for Red Hat OpenShift Container Platform with NetApp Astra – Cloud Connection to Azure with NDFC and NDO

Whitepaper describing the network setup between the FlexPod solution and the Azure public cloud with Cisco Nexus Dashboard Fabric Controller (NDFC) and Cisco Nexus Dashboard Orchestrator (NDO). This will include data protection and data management with NetApp Astra Control Center to protect all containerized workloads and support workload migration between the locations.

FlexPod Datacenter for Red Hat OpenShift Container Platform with NetApp Astra – Cloud connection to AWS with NDFC and NDO

Whitepaper describing the network setup between the FlexPod solution and the AWS public cloud with Cisco Nexus Dashboard Fabric Controller (NDFC) and Nexus Dashboard Orchestrator (NDO). This will include data protection and data management with NetApp Astra Control Center to protect all containerized workloads and support workload migration between the locations.

FlexPod Datacenter for Red Hat OpenShift Container Platform with NetApp Astra – Cloud connection to Google Cloud with NDFC and NDO

Whitepaper describing the network setup between the FlexPod solution and the Google public cloud with Cisco Nexus Dashboard Fabric Controller (NDFC) and Cisco Nexus Dashboard Orchestrator (NDO). This will include data protection and data management with NetApp Astra Control Center to protect all containerized workloads and support workload migration between the locations.

FlexPod Datacenter for Red Hat OpenShift Container Platform with NetApp Astra – Cloud connection to AWS with ACI

Whitepaper describing the network setup between the FlexPod solution and the Azure public cloud with Cisco Application Centric Infrastructure (ACI) and Cisco Nexus Dashboard Orchestrator (NDO). This will include data protection and data management with NetApp Astra Control Center to protect all containerized workloads and support workload migration between the locations.

Available Publications

FlexPod Infrastructure as Code (IaC) for Red Hat OpenShift Container Platform 4.7 Bare Metal

Cisco UCS M5, Red Hat OCP 4.7, NetApp ONTAP 9.8

Cisco Validated Design describing the deployment of the FlexPod Datacenter stack and the Red Hat OpenShift software automated with Red Hat Ansible.

FlexPod Datacenter for OpenShift Container Platform 4 - Cisco

Cisco UCS M5, Red Hat OCP 4.4, NetApp ONTAP 9.7

Cisco Validated Design describing the deployment of FlexPod Datacenter with Red Hat OpenShift based on Cisco UCS Manager.

FlexPod for Hybrid Cloud using Cisco Intersight Service and Cloud Volumes ONTAP Replication - Cisco

Cisco UCS M5, NetApp ONTAP 9.9, Cloud Volumes ONTAP 9.10

Cisco Validated Design describing in detail the design and configuration steps to replicate date from NetApp AFF storage system to NetApp Cloud Volumes ONTAP running at AWS.

Appendix

This appendix is organized into the following:

- Compute
- Network
- Storage
- Container Platform
- Interoperability Matrix
- Glossary of Acronyms
- Glossary of Terms

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 6536 Fabric Interconnects:
https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs6536-fabric-interconnect-ds.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://community.netapp.com/t5/Tech-ONTAP-Blogs/Introducing-NetApp-Active-IQ-Unified-Manager-9-11/ba-p/435519

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

Container Platform

Red Hat OpenShift Container Platform: http://www.redhat.com/ocp

Red Hat Core OS: https://www.redhat.com/coreos

Interoperability Matrix

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

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

Glossary of Acronyms

AAA–Authentication, Authorization, and Accounting

ACP–Access-Control Policy

ACI–Cisco Application Centric Infrastructure

ACK–Acknowledge or Acknowledgement

ACL–Access-Control List

AD–Microsoft Active Directory

AFI–Address Family Identifier

AMP–Cisco Advanced Malware Protection

AP–Access Point

API–Application Programming Interface

APIC– Cisco Application Policy Infrastructure Controller (ACI)

ASA–Cisco Adaptative Security Appliance

ASM–Any-Source Multicast (PIM)

ASR–Aggregation Services Router

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

AVC–Application Visibility and Control

BFD–Bidirectional Forwarding Detection

BGP–Border Gateway Protocol

BMS–Building Management System

BSR–Bootstrap Router (multicast)

BYOD–Bring Your Own Device

CAPWAP–Control and Provisioning of Wireless Access Points Protocol

CDP–Cisco Discovery Protocol

CEF–Cisco Express Forwarding

CMD–Cisco Meta Data

CPU–Central Processing Unit

CSR–Cloud Services Routers

CTA–Cognitive Threat Analytics

CUWN–Cisco Unified Wireless Network

CVD–Cisco Validated Design

CYOD–Choose Your Own Device

DC–Data Center

DHCP–Dynamic Host Configuration Protocol

DM–Dense-Mode (multicast)

DMVPN–Dynamic Multipoint Virtual Private Network

DMZ–Demilitarized Zone (firewall/networking construct)

DNA–Cisco Digital Network Architecture

DNS–Domain Name System

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

DWDM–Dense Wavelength Division Multiplexing

ECMP–Equal Cost Multi Path

EID–Endpoint Identifier

EIGRP–Enhanced Interior Gateway Routing Protocol

EMI–Electromagnetic Interference

ETR–Egress Tunnel Router (LISP)

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

FHR–First-Hop Router (multicast)

FHRP–First-Hop Redundancy Protocol

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

OCP—OpenShift Container Platform

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)

RHCOS—Red Hat Core OS

RHEL–Red Hat Enterprise Linux

RIB–Routing Information Base

RLOC–Routing Locator (LISP)

RP–Rendezvous Point (multicast)

RP–Redundancy Port (WLC)

RP–Route Processer

RPF–Reverse Path Forwarding

RR–Route Reflector (BGP)

RTT–Round-Trip Time

SA–Source Active (multicast)

SAFI–Subsequent Address Family Identifiers (BGP)

SD–Software-Defined

SDA–Cisco Software Defined-Access

SDN–Software-Defined Networking

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

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

SGACL–Security-Group ACL

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

SM–Spare-mode (multicast)

SNMP–Simple Network Management Protocol

SSID–Service Set Identifier (wireless)

SSM–Source-Specific Multicast (PIM)

SSO–Stateful Switchover

STP–Spanning-tree protocol

SVI–Switched Virtual Interface

SVL–Cisco StackWise Virtual

SWIM–Software Image Management

SXP–Scalable Group Tag Exchange Protocol

Syslog–System Logging Protocol

TACACS+–Terminal Access Controller Access-Control System Plus

TCP–Transmission Control Protocol (OSI Layer 4)

UCS– Cisco Unified Computing System

UDP–User Datagram Protocol (OSI Layer 4)

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

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

URL–Uniform Resource Locator

VLAN–Virtual Local Area Network

VM—Virtual Machine

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

VNI–Virtual Network Identifier (VXLAN)

vPC–virtual Port Channel (Cisco Nexus)

VPLS–Virtual Private LAN Service

VPN–Virtual Private Network

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

VPWS–Virtual Private Wire Service

VRF–Virtual Routing and Forwarding

VSL–Virtual Switch Link (Cisco VSS component)

VSS–Cisco Virtual Switching System

VXLAN–Virtual Extensible LAN

WAN–Wide-Area Network

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

WoL–Wake-on-LAN

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

Glossary of Terms

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

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

	connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity.” https://en.wikipedia.org/wiki/Colocation_centre

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

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

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

Feedback

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

CVD Program

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series. Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cis-co MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trade-marks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_1)

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

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**.