



# FlexPod Datacenter using Cisco UCS X-Series with VMware Horizon and VMware vSphere 7 for up to 2600 Seats

Deployment Guide for Virtual Desktop Infrastructure built on Cisco UCS 210C M6X-Series with 3<sup>rd</sup> Generation Intel Xeon Processors, Cisco Intersight 5.0(1b), NetApp Storage for VMware Horizon Remote Desktop Server Hosted Sessions, Windows 10 Desktops, and VMware vSphere 7.0 U3 Hypervisor

---

Published: December 2022



In partnership with:



---

## About the Cisco Validated Design Program

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

---

## Executive Summary

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

The solution explains the deployment of a predesigned, best-practice datacenter architecture with VMware Horizon Remote Desktop Server Hosted (RDSH) sessions and Windows 10 Virtual desktops and VMware vSphere built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and NetApp Storage AFF A400 all flash array supporting Fibre Channel storage access.

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

When deployed, the architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Desktop Infrastructure (VDI).

This document provides a Reference Architecture for a virtual desktop and application design using VMware Remote Desktop Server Hosted (RDSH) and VMware Windows 10 Virtual Desktops built on Cisco UCS with a NetApp All Flash FAS (AFF) A400 storage and the VMware vSphere ESXi 7.0.U3 hypervisor platform.

This document explains the deployment details of incorporating the Cisco X-Series modular platform into the FlexPod Datacenter and the ability to manage and orchestrate FlexPod components from the cloud using Cisco Intersight. Some of the key advantages of integrating Cisco UCS X-Series into the FlexPod infrastructure are:

- **Simpler and programmable infrastructure:** infrastructure as a code delivered through a single partner integrable open API
- **Power and cooling innovations:** higher power headroom and lower energy loss due to a 54V DC power delivery to the chassis
- **Better airflow:** midplane-free design with fewer barriers, therefore lower impedance
- **Fabric innovations:** PCIe/Compute Express Link (CXL) topology for heterogeneous compute and memory composability
- **Innovative cloud operations:** continuous feature delivery and no need for maintaining on-premises virtual machines supporting management functions
- **Built for investment protections:** design ready for future technologies such as liquid cooling and high-Wattage CPUs; CXL-ready

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

The landscape of desktop and application virtualization is changing constantly. The high-performance Cisco UCS B series blade servers and Cisco UCS unified fabric combined as part of the FlexPod Proven Infrastructure with the latest generation NetApp AFF storage result in a more compact, more powerful, more reliable, and more efficient platform.

---

This document provides the architecture and design of a virtual desktop infrastructure for up to 2600 compute end users. The solution virtualized on Cisco UCS X-Series blade server, booting VMware vSphere 7.0 Update 3 through FC SAN from the AFF A400 storage array. The virtual desktops are powered using VMware Remote Desktop Server Hosted (RDSH) Sessions and VMware Win 10 Virtual Desktops, with a mix of RDS hosted shared desktops (2600), pooled/non-persistent hosted virtual Windows 10 Instant Clones desktops (1800) and persistent Full clone virtual Windows 10 desktops.

The solution provides outstanding virtual desktop end-user experience as measured by the Login VSI 4.1.40 Knowledge Worker workload running in benchmark mode.

The 2600 solution provides a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

---

## Solution Overview

- [Introduction](#)
- [Audience](#)
- [Purpose of this Document](#)
- [What's New in this Release](#)
- [FlexPod Cisco Validated Design Advantages for VDI](#)
- [Cisco Desktop Virtualization Solutions: Datacenter](#)
- [Physical Topology](#)
- [Configuration Guidelines](#)
- [What is FlexPod?](#)

### Introduction

The current industry trend in datacenter design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility, and reducing costs. Cisco, NetApp storage, and VMware have partnered to deliver this Cisco Validated Design, which uses best of breed storage, server, and network components to serve for the foundation for desktop virtualization workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

The Cisco UCS X-Series Blade Server delivers performance, flexibility, and optimization for deployments in datacenters, cloud, and remote sites. This enterprise-class server offers market-leading versatility, and density without compromise for workloads, including web infrastructure, distributed databases, Virtual Desktop Infrastructure (VDI), converged infrastructure, and enterprise applications such as SAP HANA and Oracle. The Cisco UCS X-Series Blade Server can quickly deploy stateless physical and virtual workloads through a programmable, easy-to-use Cisco UCS Manager and Cisco Intersight and simplified server access through Cisco SingleConnect technology.

### Audience

The intended audience for this document includes, but is not limited to IT architects, sales engineers, field consultants, professional services, IT managers, IT engineers, partners, and customers who are interested in learning about and deploying the Virtual Desktop Infrastructure (VDI)

### Purpose of this document

This document provides a step-by-step design, configuration, and implementation guide for the Cisco Validated Design for a large-scale VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions and Windows 10 Virtual Desktops with NetApp AFF A400, NS224 NVMe Disk Shelf, Cisco UCS X210c M6 compute nodes, Cisco Nexus 9000 Series Ethernet Switches and Cisco MDS 9000 Series Multilayer Fibre Channel Switches.

### What's New in this Release?

This version of the FlexPod VDI Design based on the latest Cisco FlexPod Virtual Server Infrastructure and introduces the Cisco UCS 6<sup>th</sup> generation servers and a NetApp AFF A-Series system.

---

This is the first VMware Horizon Remote Desktop Server Hosted (RDSH) sessions and VMware Horizon Win 10 virtual desktops virtualization Cisco Validated Design with Cisco UCS 6<sup>th</sup> generation servers and a NetApp AFF A-Series system.

It incorporates the following features:

- Integration of Cisco UCS X-Series into FlexPod Datacenter
- Deploying and managing Cisco UCS X9508 chassis equipped with Cisco UCS X210c M6 compute nodes from the cloud using Cisco Intersight
- Integration of Cisco Intersight with NetApp Active IQ Unified Manager for storage monitoring and orchestration
- Integration of Cisco Intersight with VMware vCenter for interacting with, monitoring, and orchestrating the virtual environment
- Support for Cisco UCS Cisco UCS X210c M6 compute nodes with Intel Xeon Scalable Family processors and 3200 MHz memory
- Validation of Cisco Nexus 9000 with NetApp AFF A400 system
- Validation of Cisco MDS 9000 with NetApp AFF A400 system
- Support for the Cisco UCS 5.0(1b) release and Cisco X-Series Compute nodes
- VMware vSphere 7.0.U3 Hypervisor
- VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions for Server 2019 RDS hosted shared virtual desktops
- VMware Horizon non-persistent instant clones virtual Windows 10 desktops
- VMware Horizon persistent full clones virtual Windows 10 desktops
- Support for NetApp Storage AFF A400 with ONTAP version 9.10.1P1
- VMware Horizon 2209 Remote Desktop Sever Hosted (RDSH) Sessions
- VMware Horizon 2209 Horizon instant clone virtual machines
- VMware Horizon 2209 Horizon persistent full desktops
- Support for VMware vSphere 7.0.Update 3
- Fully automated solution deployment covering FlexPod infrastructure and vSphere virtualization

## **FlexPod Cisco Validated Design Advantages for VDI**

The datacenter market segment is shifting toward heavily virtualized private, hybrid and public cloud computing models running on industry-standard systems. These environments require uniform design points that can be repeated for ease of management and scalability.

These factors have led to the need predesigned computing, networking and storage building blocks optimized to lower the initial design cost, simplify management, and enable horizontal scalability and high levels of utilization. The use cases include:

- Enterprise Datacenter (small failure domains)
- Service Provider Datacenter (small failure domains)

- 
- Commercial Datacenter
  - Remote Office/Branch Office
  - SMB Standalone Deployments
  - Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both Horizon Microsoft Windows 10 virtual desktops and Horizon RDSH server desktop sessions based on Microsoft Server 2019. The mixed workload solution includes Cisco UCS hardware and Data Platform software, Cisco Nexus® switches, the Cisco Unified Computing System (Cisco UCS), VMware Horizon and VMware vSphere software in a single package. The design is efficient such that the networking, computing, and storage components occupy 18-rack units footprint in an industry standard 42U rack. Port density on the Cisco Nexus switches and Cisco UCS Fabric Interconnects enables the networking components to accommodate multiple UCS clusters in a single Cisco UCS domain.

A key benefit of the Cisco Validated Design architecture is the ability to customize the environment to suit a customer's requirements. A Cisco Validated Design scales easily as requirements and demand change. The unit can be scaled both up (adding resources to a Cisco Validated Design unit) and out (adding more Cisco Validated Design units).

The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a hyper-converged desktop virtualization solution. A solution capable of consuming multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

The combination of technologies from Cisco Systems, Inc. and VMware Inc. produced a highly efficient, robust, and affordable desktop virtualization solution for a virtual desktop, hosted shared desktop or mixed deployment supporting different use cases. Key components of the solution include the following:

- **More power, same size.** Cisco UCS X210c M6 compute nodes with dual 28-core 2.6 GHz Intel Xeon Gold (6348) Scalable Family processors with 1 TB of 3200 Mhz memory with VMware Horizon support more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel Xeon Gold 6348, 28-core scalable family processors used in this study provided a balance between increased per-server capacity and cost.
- **Fault-tolerance with high availability built into the design.** The various designs are based on multiple Cisco UCS X-Series blade servers for virtual desktop and infrastructure workloads. The design provides N+1 server fault tolerance for every payload type tested.
- **Stress-tested to the limits during aggressive boot scenario.** The 2600 user Remote Desktop Hosted (RDSH) Sessions and 1800 Win 10 Virtual Desktops environment booted and registered with the Horizon 8 in under 10 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.
- **Stress-tested to the limits during simulated login storms.** The 2600 user Remote Desktop Hosted (RDSH) Sessions and 1800 Win 10 Virtual Desktops environment ready state in 48-minutes without overwhelming the processors, exhausting memory, or exhausting the storage subsystems, providing customers with a desktop virtualization system that can easily handle the most demanding login and startup storms.
- **Ultra-condensed computing for the datacenter.** The rack space required to support the initial 1800 user system is 8 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects. Incremental seat Cisco converged solutions clusters can be added one at a time to a total of 32 nodes.



- **100 percent virtualized** This CVD presents a validated design that is 100 percent virtualized on VMware ESXi 7.0. U3 All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, VMware Horizon Connection Server components, Horizon VDI virtual desktops and RDSH servers were hosted as virtual machines.
- **Cisco datacenter management:** Cisco maintains industry leadership with the new Cisco UCS Manager 5.0(1b) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director ensure that customer environments are consistent locally, across Cisco UCS Domains and across the globe. Cisco UCS software suite offers increasingly simplified operational and deployment management, and it continues to widen the span of control for customer organizations' subject matter experts in compute, storage, and network.
- **Our 25G unified fabric story** gets additional validation on Cisco UCS 6400 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- **NetApp AFF A400** array provides industry-leading storage solutions that efficiently handle the most demanding I/O bursts (for example, login storms), profile management, and user data management, deliver simple and flexible business continuance, and help reduce storage cost per desktop.
- **NetApp AFF A400** array provides a simple to understand storage architecture for hosting all user data components (VMs, profiles, user data) on the same storage array.
- **NetApp clustered Data ONTAP software** enables to seamlessly add, upgrade, or remove storage from the infrastructure to meet the needs of the virtual desktops.
- **VMware Horizon 8 advantage:** VMware Horizon 8 follows a new unified product architecture that supports both Virtual Desktops and Remote Desktop Server Hosted server sessions. This new Horizon release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of user increase. In addition, PCoIP and Blast extreme enhancements help to optimize performance and improve the user experience across a variety of endpoint device types, from workstations to mobile devices including laptops, tablets, and smartphones.
- **Optimized for performance and scale.** For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the Horizon 8 RDSH virtual machines did not exceed the number of hyper-threaded (logical) cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for the virtual machines running virtualized RDS systems.
- **Provisioning desktop machines made easy:** VMware Horizon 8 provisions Remote Desktop Hosted Sessions (RDSH) virtual desktops as well as hosted shared desktop virtual machines for this solution using a single method for both, the "Automated floating assignment desktop pool." "Dedicated user assigned desktop pool" for persistent desktops was provisioned in the same Horizon 8 administrative console. The new method of Instant Clone greatly reduces the amount of life-cycle spend and the maintenance windows for the guest OS.

## Cisco Desktop Virtualization Solutions: Datacenter

### The Evolving Workplace

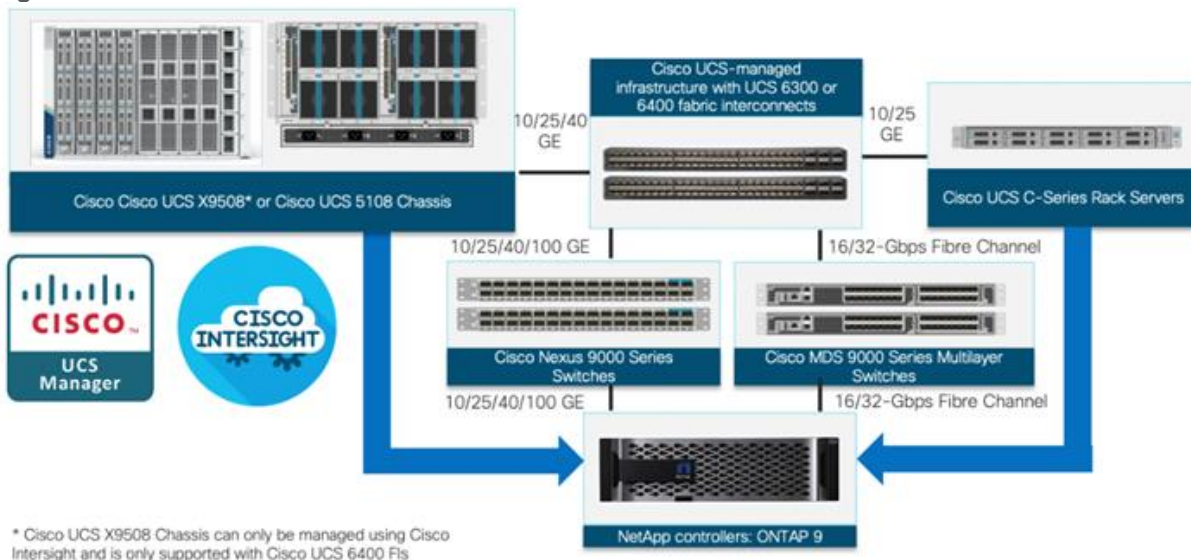
Today's IT departments are facing a rapidly evolving workplace environment. The workforce is becoming increasingly diverse and geographically dispersed, including offshore contractors, distributed call center

operations, knowledge and task workers, partners, consultants, and executives connecting from locations around the world at all times.

This workforce is also increasingly mobile, conducting business in traditional offices, conference rooms across the enterprise campus, home offices, on the road, in hotels, and at the local coffee shop. This workforce wants to use a growing array of client computing and mobile devices that they can choose based on personal preference. These trends are increasing pressure on IT to ensure protection of corporate data and prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 1).

These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 10 and productivity tools, namely Microsoft Office 2016.

**Figure 1. Cisco Datacenter Partner Collaboration**



Some of the key drivers for desktop virtualization are increased data security, the ability to expand and contract capacity and reduced TCO through increased control and reduced management costs.

### Cisco Desktop Virtualization Focus

Cisco focuses on three key elements to deliver the best desktop virtualization datacenter infrastructure: simplification, security, and scalability. The software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform.

#### Simplified

Cisco UCS and NetApp provide a radical new approach to industry-standard computing and provides the core of the datacenter infrastructure for desktop virtualization. Among the many features and benefits of Cisco UCS are the drastic reduction in the number of servers needed, in the number of cables used per server and the capability to rapidly deploy or re-provision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified. Thousands of desktops can be provisioned in minutes with Cisco UCS Manager service profiles and Cisco storage partners' storage-based cloning. This approach accelerates the time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

Cisco UCS Manager automates many mundane, error-prone datacenter operations such as configuration and provisioning of server, network, and storage access infrastructure. In addition, Cisco UCS B-Series Blade

---

Servers, C-Series and HX-Series Rack Servers with large memory footprints enable high desktop density that helps reduce server infrastructure requirements.

Simplification also leads to more successful desktop virtualization implementation. Cisco and its technology partners like VMware have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlexPod. Cisco Desktop Virtualization Solutions have been tested with VMware vSphere.

## **Secure**

Although virtual desktops are inherently more secure than their physical predecessors, they introduce new security challenges. Mission-critical web and application servers using a common infrastructure such as virtual desktops are now at a higher risk for security threats. Inter-virtual machine traffic now poses an important security consideration that IT managers need to address, especially in dynamic environments in which virtual machines, using VMware vMotion, move across the server infrastructure.

Desktop virtualization, therefore, significantly increases the need for virtual machine-level awareness of policy and security, especially given the dynamic and fluid nature of virtual machine mobility across an extended computing infrastructure. The ease with which new virtual desktops can proliferate magnifies the importance of a virtualization-aware network and security infrastructure. Cisco datacenter infrastructure (Cisco UCS and Cisco Nexus Family solutions) for desktop virtualization provides strong datacenter, network, and desktop security, with comprehensive security from the desktop to the hypervisor. Security is enhanced with segmentation of virtual desktops, virtual machine-aware policies and administration, and network security across the LAN and WAN infrastructure.

## **Scalable**

The growth of a desktop virtualization solution is all but inevitable, so a solution must be able to scale, and scale predictably, with that growth. The Cisco Desktop Virtualization Solutions built on FlexPod Datacenter infrastructure supports high virtual-desktop density (desktops per server), and additional servers and storage scale with near-linear performance. FlexPod Datacenter provides a flexible platform for growth and improves business agility. Cisco UCS Manager service profiles allow on-demand desktop provisioning and make it just as easy to deploy dozens of desktops as it is to deploy thousands of desktops.

Cisco UCS servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 1 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco storage partners NetApp help maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs for End User Computing based on FlexPod solutions have demonstrated scalability and performance, with up to 2600 desktops up and running in less than 15 minutes.

FlexPod Datacenter provides an excellent platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization, datacenter applications, and cloud computing.

Cisco UCS and Cisco Nexus datacenter infrastructure provides an excellent platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization, datacenter applications, and cloud computing.

---

## Savings and Success

The simplified, secure, scalable Cisco datacenter infrastructure for desktop virtualization solutions saves time and money compared to alternative approaches. Cisco UCS enables faster payback and ongoing savings (better ROI and lower TCO) and provides the industry's greatest virtual desktop density per server, reducing both capital expenditures (CapEx) and operating expenses (OpEx). The Cisco UCS architecture and Cisco Unified Fabric also enables much lower network infrastructure costs, with fewer cables per server and fewer ports required. In addition, storage tiering and deduplication technologies decrease storage costs, reducing desktop storage needs by up to 50 percent.

The simplified deployment of Cisco NetApp FlexPod solution for desktop virtualization accelerates the time to productivity and enhances business agility. IT staff and end users are more productive more quickly, and the business can respond to new opportunities quickly by deploying virtual desktops whenever and wherever they are needed. The high-performance Cisco Systems and network deliver a near-native end-user experience, allowing users to be productive anytime and anywhere.

The key measure of desktop virtualization for any organization is its efficiency and effectiveness in both the near term and the long term. The Cisco Desktop Virtualization Solutions are very efficient, allowing rapid deployment, requiring fewer devices and cables, and reducing costs. The solutions are also extremely effective, providing the services that end users need on their devices of choice while improving IT operations, control, and data security. Success is bolstered through Cisco's best-in-class partnerships with leaders in virtualization and through tested and validated designs and services to help customers throughout the solution lifecycle. Long-term success is enabled through the use of Cisco's scalable, flexible, and secure architecture as the platform for desktop virtualization.

The ultimate measure of desktop virtualization for any end-user is a great experience. Cisco NetApp deliver class-leading performance with sub-second base line response times and index average response times at full load of just under one second.

## Use Cases


The following are some typical use cases:

- Healthcare: Mobility between desktops and terminals, compliance, and cost
- Federal government: Teleworking initiatives, business continuance, continuity of operations (COOP), and training centers
- Financial: Retail banks reducing IT costs, insurance agents, compliance, and privacy
- Education: K-12 student access, higher education, and remote learning
- State and local governments: IT and service consolidation across agencies and interagency security
- Retail: Branch-office IT cost reduction and remote vendors
- Manufacturing: Task and knowledge workers and offshore contractors
- Microsoft Windows 10 migration
- Graphic intense applications
- Security and compliance initiatives
- Opening of remote and branch offices or offshore facilities
- Mergers and acquisitions

## Physical Topology

Figure 2 illustrates the physical architecture.

Figure 2. Physical Architecture



The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-FX switches
- Two Cisco MDS 9132T 32GB Fibre Channel switches
- Two Cisco UCS 6454 Fabric Interconnects
- Two Cisco UCS X-Series Chassis
- Eight Cisco X-Series 210c Compute Nodes (for VDI workload)
- Infrastructure VMs for VDI were housed on an external cluster
- One NetApp AFF A400 Storage System (HA Pair)
- Two NetApp NS224 Disk Shelves

For desktop virtualization, the deployment includes VMware Horizon Remote Desktop Session Hosts (RDSH) Sessions and Win 10 virtual desktops running on VMware vSphere 7.0.U3.

The design is intended to provide a large-scale building block for VMware Horizon Remote Desktop Session Hosted (RDSH) Sessions workloads consisting of Remote Desktops Server Hosted (RDSH) Sessions with Windows Server 2019 hosted shared desktop sessions and Windows 10 non-persistent and persistent hosted desktops in the following:

- 2600 VMware Horizon Remote Desktop Server Hosted (RDSH) Random Hosted Shared Server 2019 user sessions with Microsoft Office 2016 (Instant Clone RDS deployment)

- 
- 1800 VMWare Horizon Pooled Windows 10 Desktops with Microsoft Office 2016 (Instant Clone deployed virtual machines)
  - 1800 VMWare Horizon Persistent Full Clone Windows 10 Desktops with Microsoft Office 2016 (Full Clone deployed virtual machines)

The data provided in this document will allow our customers to adjust the mix of Remote Desktop Server Hosted (RDSH) Sessions and Win 10 Virtual Desktops to suit their environment. For example, additional blade servers and chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the detailed steps for deploying the base architecture. This procedure covers everything from physical cabling to network, compute, and storage device configurations.

## Configuration Guidelines

This Cisco Validated Design provides details for deploying a fully redundant, highly available 2600 seats workload virtual sessions /desktop solution with VMware on a FlexPod Datacenter architecture. Configuration guidelines are provided that refer the reader to which redundant component is being configured with each step. For example, storage controller 01 and storage controller 02 are used to identify the two AFF A400 storage controllers that are provisioned with this document, Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured, and Cisco MDS A or Cisco MDS B identifies the pair of Cisco MDS switches that are configured.

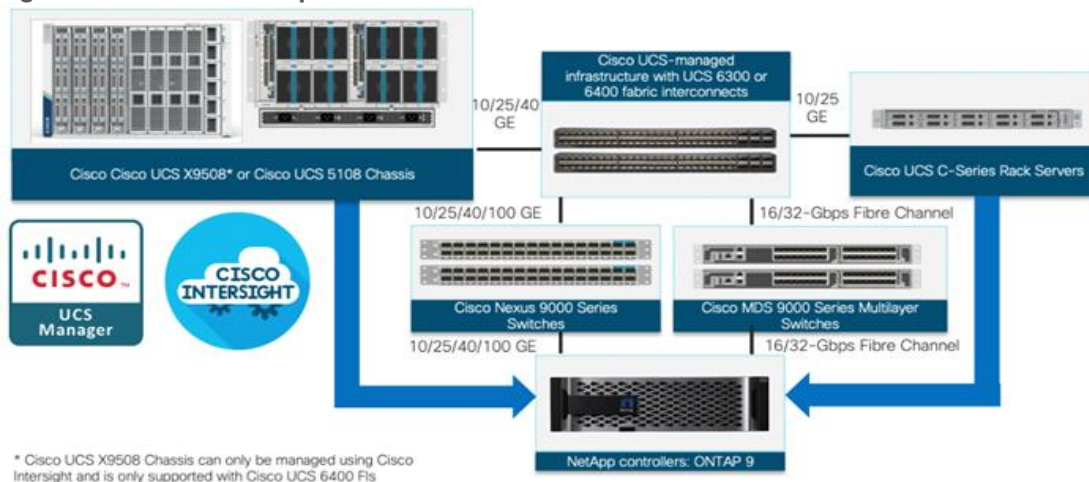
The Cisco UCS 6454 Fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, VM-Host-RDSH-01, VM-Host-VDI-01 and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure.

## What is FlexPod?

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere® built on FlexPod includes NetApp AFF storage, Cisco Nexus® networking, Cisco MDS storage networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one datacenter rack or be deployed according to a customer's datacenter design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a Fibre Channel and IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

**Figure 3. FlexPod Component Families**



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). The reference architecture covered in this document leverages Cisco Nexus 9000 for the network switching element and pulls in the Cisco MDS 9000 for the SAN switching component.

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, 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 following lists the benefits of FlexPod:

- Consistent Performance and Scalability
  - Consistent sub-millisecond latency with 100% flash storage
  - Consolidate 100's of enterprise-class applications in a single rack
  - Scales easily, without disruption
  - Continuous growth through multiple FlexPod CI deployments
- Operational Simplicity
  - Fully tested, validated, and documented for rapid deployment
  - Reduced management complexity
  - Auto-aligned 512B architecture removes storage alignment issues
  - No storage tuning or tiers necessary
- Lowest TCO
  - Dramatic savings in power, cooling, and space with 100 percent Flash
  - Industry leading data reduction
- Enterprise-Grade Resiliency
  - Highly available architecture with no single point of failure
  - Nondisruptive operations with no downtime
  - Upgrade and expand without downtime or performance loss

- 
- Native data protection: snapshots and replication
  - Suitable for even large resource-intensive workloads such as real-time analytics or heavy transactional databases



---

## Solution Components


This chapter contains the following:

- [Cisco Unified Compute System X-Series](#)
- [Cisco Intersight](#)
- [Cisco Switches](#)
- [Cisco Intersight Cloud-Based Management](#)
- [VMware Horizon Remote Desktop Server Hosted Sessions and Windows 10 Desktops](#)
- [NetApp A-Series All Flash FAS](#)
- [NetApp ONTAP 9.10.1P1](#)
- [VMware vSphere 7.0](#)
- [Cisco Intersight Assist Device Connector for VMware vCenter and NetApp ONTAP](#)
- [NetApp ONTAP Tools for VMware vSphere](#)
- [NetApp NFS Plug-in for VMware VAAI](#)
- [NetApp SnapCenter Plug-In for VMware vSphere](#)
- [NetApp Active IQ Unified Manager 9.10P1](#)
- [NetApp XCP File Analytics](#)

### Cisco Unified Computing System X-Series

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

**Figure 4. Cisco UCS X9508 Chassis**



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

### Cisco UCS X9508 Chassis

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

**Figure 5. Cisco UCS X9508 Chassis - Midplane Free Design**



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

## Cisco UCSX 9108-25G Intelligent Fabric Modules

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

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




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

## Cisco UCS X210c M6 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to 8 Cisco UCS X210c M6 Compute Nodes. The hardware details of the Cisco UCS X210c M6 Compute Nodes are shown in [Figure 7](#):

**Figure 7. Cisco UCS X210c M6 Compute Node**



The Cisco UCS X210c M6 features:

- **CPU:** Up to 2x 3rd Gen Intel Xeon Scalable Processors with up to 40 cores per processor and 1.5 MB Level 3 cache per core.
- **Memory:** Up to 32 x 256 GB DDR4-3200 DIMMs for a maximum of 8 TB of main memory. The Compute Node can also be configured for up to 16 x 512-GB Intel Optane persistent memory DIMMs for a maximum of 12 TB of memory.

- **Disk storage:** Up to 6 SAS or SATA drives can be configured with an internal RAID controller, or customers can configure up to 6 NVMe drives. 2 M.2 memory cards can be added to the Compute Node with RAID 1 mirroring.
- **Virtual Interface Card (VIC):** Up to 2 VICs including an mLOM Cisco VIC 14425 and a mezzanine Cisco VIC card 14825 can be installed in a Compute Node.
- **Security:** The server supports an optional Trusted Platform Module (TPM). Additional security features include a secure boot FPGA and ACT2 anticounterfeit provisions.

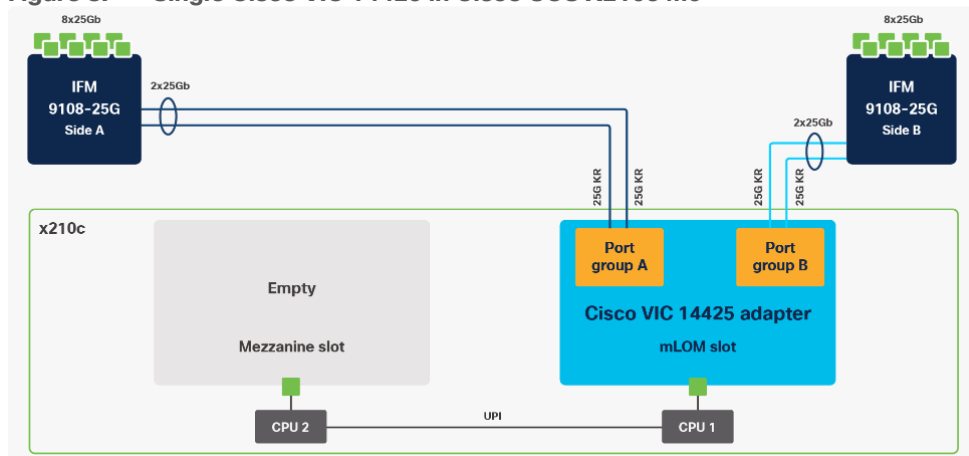
## Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS X210c M6 Compute Nodes support the following two Cisco fourth-generation VIC cards:

### Cisco VIC 14425

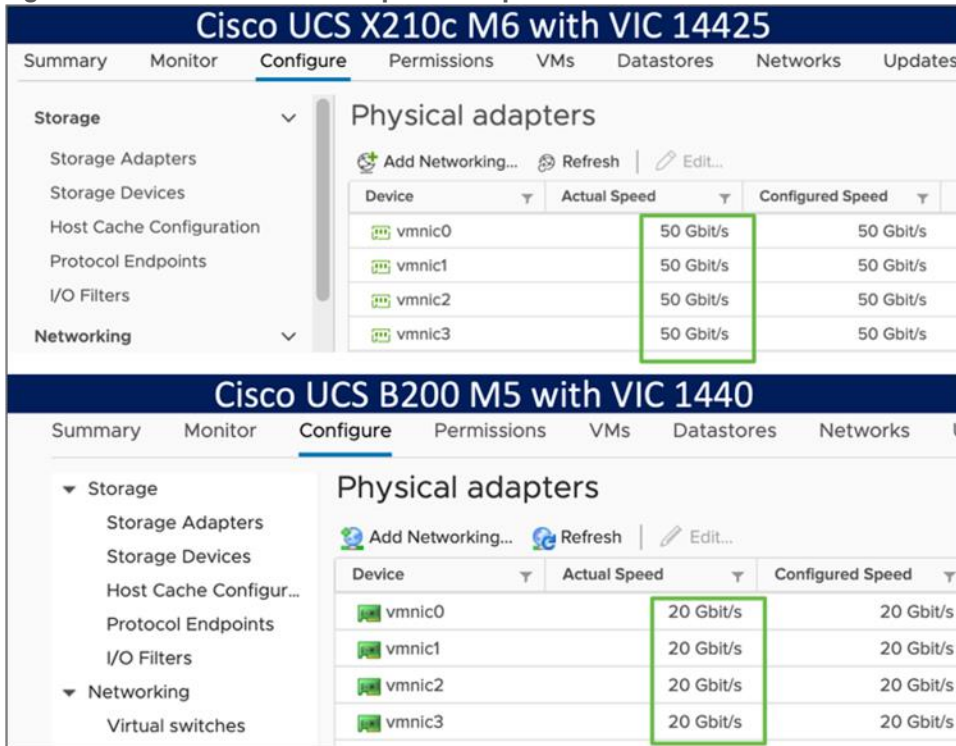
Cisco VIC 14425 fits the mLOM slot in the Cisco X210c Compute Node and enables up to 50 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 100 Gbps of connectivity per server. Cisco VIC 14425 connectivity to the IFM and up to the fabric interconnects is delivered through 4x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

**Figure 8. Single Cisco VIC 14425 in Cisco UCS X210c M6**



The connections between the 4<sup>th</sup> generation Cisco VIC (Cisco UCS VIC 1440) in the Cisco UCS B200 blades and the I/O modules in the Cisco UCS 5108 chassis comprise of multiple 10Gbps KR lanes. The same connections between Cisco VIC 14425 and IFMs in Cisco UCS X-Series comprise of multiple 25Gbps KR lanes resulting in 2.5x better connectivity in Cisco UCS X210c M6 Compute Nodes. The network interface speed comparison between VMware ESXi installed on Cisco UCS B200 M6 with VIC 1440 and Cisco UCS X210c M6 with VIC 14425 is shown in [Figure 9](#).

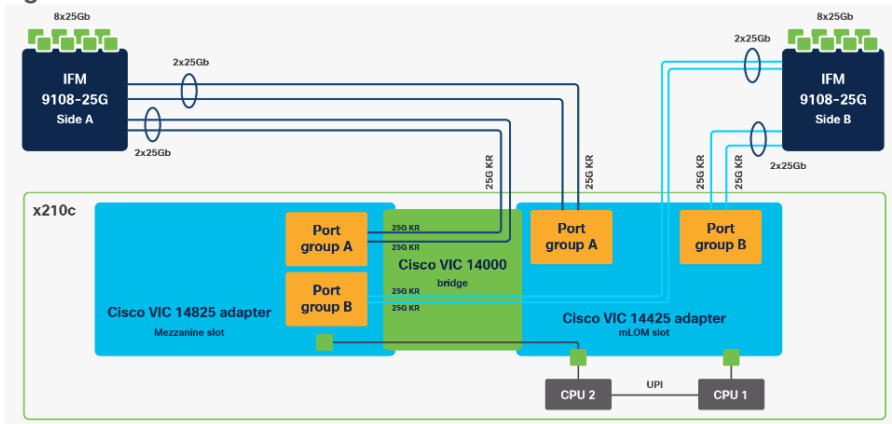
Figure 9. Network Interface Speed Comparison



### Cisco VIC 14825

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

Figure 10. Cisco VIC 14425 and 14825 in Cisco UCS X210c M6



### Cisco UCS 6400 Series Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide a single point 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 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 11. Cisco UCS 6454 Fabric Interconnect**



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

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

## Cisco Intersight

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

**Figure 12. Cisco Intersight Overview**



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

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

## Cisco Intersight Virtual Appliance and Private Virtual Appliance

In addition to the SaaS deployment model running on Intersight.com, on-premises options can be purchased separately. The Cisco Intersight Virtual Appliance and Cisco Intersight Private Virtual Appliance are available for

---

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

## Cisco Intersight Assist

Cisco Intersight Assist helps customers add endpoint devices to Cisco Intersight. A datacenter 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 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 configuration is covered in later sections.

## Licensing Requirements

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

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

Servers in the Cisco Intersight managed mode require at least the Essentials license. For more information about the features provided in the various licensing tiers, see

[https://intersight.com/help/getting\\_started#licensing\\_requirements](https://intersight.com/help/getting_started#licensing_requirements).

## Cisco Switches

### Cisco Nexus 93180YC-FX Switches

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

- Architectural Flexibility


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

A Cisco 40 Gbe [bidirectional transceiver](#) allows reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet Support for 1 Gbe and 10 Gbe access connectivity for datacenters migrating access switching infrastructure to faster speed. The following is supported:

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



**Figure 13. Cisco Nexus 93180YC-FX Switch**



### Cisco MDS 9132T 32-Gb Fiber Channel Switch

The next-generation Cisco® MDS 9132T 32-Gb 32-Port Fibre Channel Switch ([Figure 14](#)) provides high-speed Fibre Channel connectivity from the server rack to the SAN core. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the dual benefits of greater bandwidth and consolidation.

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

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

Additionally, investing in this switch for the lower-speed (4- or 8- or 16-Gb) server rack gives you the option to upgrade to 32-Gb server connectivity in the future using the 32-Gb Host Bus Adapter (HBA) that are available today. The Cisco® MDS 9132T 32-Gb 32-Port Fibre Channel switch also provides unmatched flexibility through a unique port expansion module ([Figure 9](#)) that provides a robust cost-effective, field swappable, port upgrade option.

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

**Figure 14. Cisco MDS 9132T 32-Gb 32-Port Fabric Channel Switch**



**Figure 15. Cisco MDS 9132T 32-Gb 16-Port Fibre Channel Port Expansion Module**



- Features

- High performance: MDS 9132T architecture, with chip-integrated nonblocking arbitration, provides consistent 32-Gb low-latency performance across all traffic conditions for every Fibre Channel port on the switch.
- Capital Expenditure (CapEx) savings: The 32-Gb ports allow users to deploy them on existing 16- or 8-Gb transceivers, reducing initial CapEx with an option to upgrade to 32-Gb transceivers and adapters in the future.
- High availability: MDS 9132T switches continue to provide the same outstanding availability and reliability as the previous-generation Cisco MDS 9000 Family switches by providing optional redundancy on all major components such as the power supply and fan. Dual power supplies also facilitate redundant power grids.
- Pay-as-you-grow: The MDS 9132T Fibre Channel switch provides an option to deploy as few as eight 32-Gb Fibre Channel ports in the entry-level variant, which can grow by 8 ports to 16 ports, and thereafter with a port expansion module with sixteen 32-Gb ports, to up to 32 ports. This approach results in lower initial investment and power consumption for entry-level configurations of up to 16 ports compared to a fully loaded switch. Upgrading through an expansion module also reduces the overhead of managing multiple instances of port activation licenses on the switch. This unique combination of port upgrade options allow four possible configurations of 8 ports, 16 ports, 24 ports and 32 ports.
- Next-generation Application-Specific Integrated Circuit (ASIC): The MDS 9132T Fibre Channel switch is powered by the same high-performance 32-Gb Cisco ASIC with an integrated network processor that powers the Cisco MDS 9700 48-Port 32-Gb Fibre Channel Switching Module. Among all the advanced features that this ASIC enables, one of the most notable is inspection of Fibre Channel and Small Computer System Interface (SCSI) headers at wire speed on every flow in the smallest form-factor Fibre Channel switch without the need for any external taps or appliances. The recorded flows can be analyzed on the switch and also exported using a dedicated 10/100/1000BASE-T port for telemetry and analytics purposes.
- Intelligent network services: Slow-drain detection and isolation, VSAN technology, Access Control Lists (ACLs) for hardware-based intelligent frame processing, smart zoning, and fabric wide Quality of Service (QoS) enable migration from SAN islands to enterprise-wide storage networks. Traffic encryption is optionally available to meet stringent security requirements.
- Sophisticated diagnostics: The MDS 9132T provides intelligent diagnostics tools such as Inter-Switch Link (ISL) diagnostics, read diagnostic parameters, protocol decoding, network analysis tools, and integrated Cisco Call Home capability for greater reliability, faster problem resolution, and reduced service costs.
- Virtual machine awareness: The MDS 9132T provides visibility into all virtual machines logged into the fabric. This feature is available through HBAs capable of priority tagging the Virtual Machine Identifier (VMID) on every FC frame. Virtual machine awareness can be extended to intelligent fabric services such as analytics[1] to visualize performance of every flow originating from each virtual machine in the fabric.
- Programmable fabric: The MDS 9132T provides powerful Representational State Transfer (REST) and Cisco NX-API capabilities to enable flexible and rapid programming of utilities for the SAN as well as polling point-in-time telemetry data from any external tool.
- Single-pane management: The MDS 9132T can be provisioned, managed, monitored, and troubleshoot using Cisco Datacenter Network Manager (DCNM), which currently manages the entire suite of Cisco datacenter products.

- Self-contained advanced anticounterfeiting technology: The MDS 9132T uses on-board hardware that protects the entire system from malicious attacks by securing access to critical components such as the bootloader, system image loader and Joint Test Action Group (JTAG) interface.
- 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 Nexus switches are added with the appropriate credentials and licensing, monitoring of the SAN and Ethernet fabrics can begin. Additionally, VSANs, device aliases, zones, and zone sets can be added, modified, and deleted using the DCNM point-and-click interface. Device Manager can also be used to configure the Cisco MDS switches. SAN Analytics can be added to Cisco MDS switches to provide insights into the fabric by allowing customers to monitor, analyze, identify, and troubleshoot performance issues.
- Cisco DCNM integration with Cisco Intersight
  - The Cisco Network Insights Base (Cisco NI Base) application provides several TAC assist functionalities which are useful when working with Cisco TAC. 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.

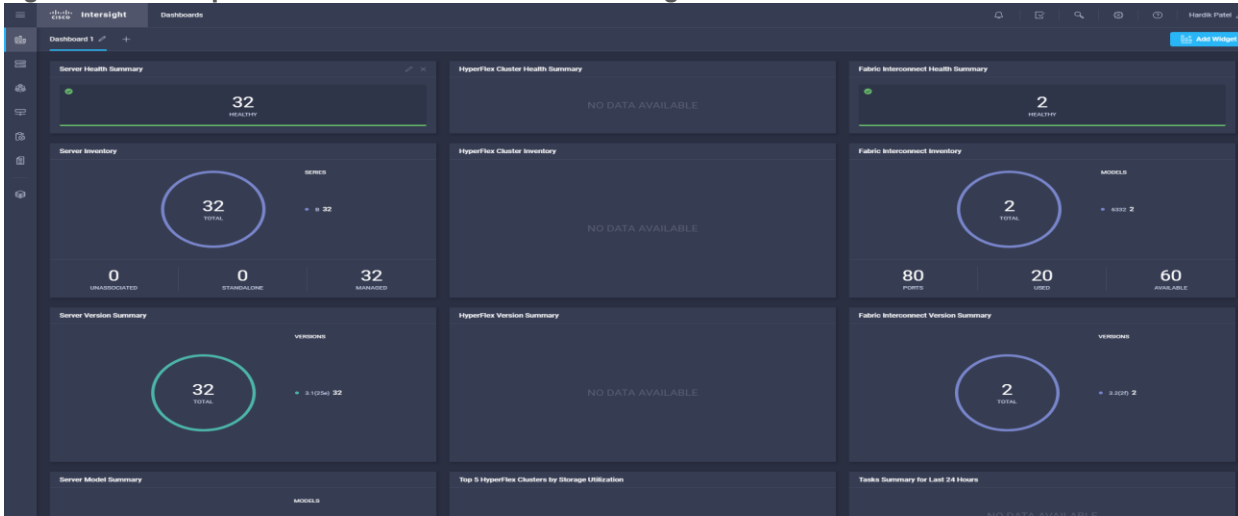
## Cisco Intersight Cloud-Based Management

[Cisco Intersight](#) is Cisco's new systems management platform that delivers intuitive computing through cloud-powered intelligence. This platform offers a more intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in ways that were not possible with prior generations of tools. This capability empowers organizations to achieve significant savings in Total Cost of Ownership (TCO) and to deliver applications faster, so they can support new business initiatives. The advantages of the model-based management of the Cisco UCS platform plus Cisco Intersight are extended to Cisco UCS servers.

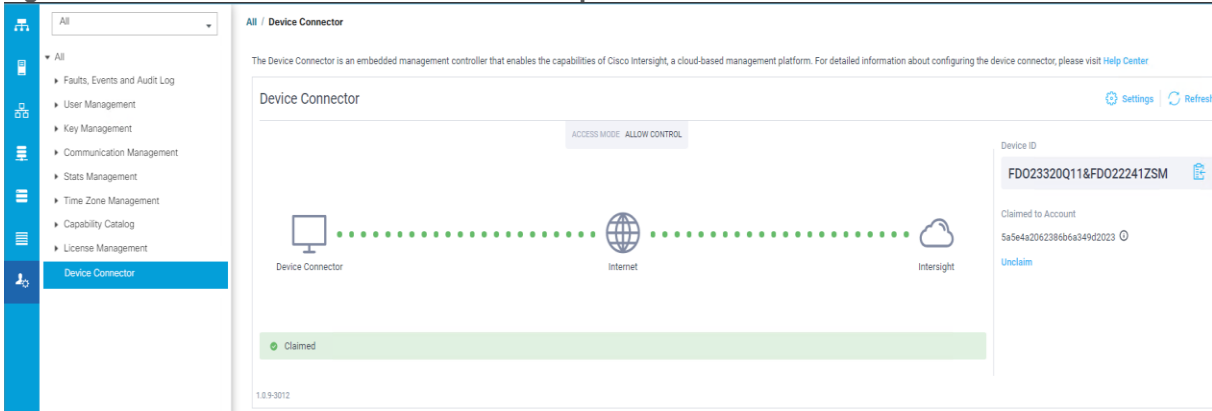
The Cisco UCS platform uses model-based management to provision servers and the associated storage and fabric automatically, regardless of form factor. Cisco Intersight works in conjunction with Cisco UCS Manager and the Cisco® Integrated Management Controller (IMC). By simply associating a model-based configuration with a resource through service profiles, your IT staff can consistently align policy, server personality, and workloads. These policies can be created once and used by IT staff with minimal effort to deploy servers. The result is improved productivity and compliance and lower risk of failures due to inconsistent configuration.

Cisco Intersight will be integrated with datacenter, hybrid cloud platforms, and services to securely deploy and manage infrastructure resources across datacenter and edge environments. In addition, Cisco will provide future integrations to third-party operations tools to allow customers to use their existing solutions more effectively.

**Figure 16. Example of User-Customizable Cisco Intersight Dashboard for FlexPod UCS Domain**



**Figure 17. Cisco UCSM Device Connector Example**



## VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions and Windows 10 Desktops

The virtual app and desktop solution is designed for an exceptional experience.

Today's employees spend more time than ever working remotely, causing companies to rethink how IT services should be delivered. To modernize infrastructure and maximize efficiency, many are turning to desktop as a service (DaaS) to enhance their physical desktop strategy, or they are updating on-premises virtual desktop infrastructure (VDI) deployments. Managed in the cloud, these deployments are high-performance virtual instances of desktops and apps that can be delivered from any datacenter or public cloud provider.

DaaS and VDI capabilities provide corporate data protection as well as an easily accessible hybrid work solution for employees. Because all data is stored securely in the cloud or datacenter, rather than on devices, end-users can work securely from anywhere, on any device, and over any network—all with a fully IT-provided experience. IT also gains the benefit of centralized management, so they can scale their environments quickly and easily. By separating endpoints and corporate data, resources stay protected even if the devices are compromised.

As a leading VDI and DaaS provider, VMware provides the capabilities organizations need for deploying virtual apps and desktops to reduce downtime, increase security, and alleviate the many challenges associated with traditional desktop management.

For more information, go to:

## NetApp A-Series All Flash FAS

Powered by [NetApp® ONTAP® data management software](#), [NetApp® AFF A-Series systems](#) (NetApp AFF) deliver the industry’s highest performance, superior flexibility, and best-in-class data services and cloud integration to help you accelerate, manage, and protect business-critical data across your hybrid cloud. It is a robust scale-out platform built for virtualized environments, combining low-latency performance with best-in-class data management, built-in efficiencies, integrated data protection, multiprotocol support, and nondisruptive operations. Deploy as a stand-alone system or as a high-performance tier in a NetApp ONTAP® configuration.

A wide range of organizations, from enterprise to midsize businesses, rely on NetApp AFF A-Series to:

- Simplify operations with seamless data management, on the premises and in the cloud.
- Accelerate traditional and new-generation applications.
- Keep business-critical data available, protected, and secure.
- Accelerates applications and future-proofs your infrastructure

In the modern datacenter, IT is charged with driving maximum performance for business-critical workloads, scaling without disruption as the business grows, and enabling the business to take on new data-driven initiatives. NetApp AFF A-Series systems handle all of it with ease.

The NetApp AFF A-Series lineup includes the A250, A400, A700, A800 and A900. These controllers and their technical specifications are listed in [Table 1](#). For more information about the A-Series AFF controllers, see:

<http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>

<https://hww.netapp.com/Controller/Index?platformTypeld=5265148>

**Table 1. NetApp AFF Technical Specifications**

Specifications	AFF A250	AFF A400	AFF A800	AFF A900
Maximum scale-out	2-24 nodes (12 HA pair)	2-24 nodes (12 HA pair)	2-24 nodes (12 HA pair)	2-24 nodes (12 HA pair)
Maximum SSDs	576	5760	2880	5760
Max effective capacity	35PB	702.7PB	316.3PB	702.7PB
Controller form factor	2U	4U	4U with 48 SSD slots	8U
PCIe expansion slots	4	10	8	20
FC target ports (32Gb autoranging)	16	24	32	64
FC target ports (16Gb autoranging)	n/a	32(with FC mezzanine card)	32	64

Specifications	AFF A250	AFF A400	AFF A800	AFF A900
FCoE target ports, UTA2	n/a	n/a	n/a	64
100GbE ports (40GbE autoranging)	4	16	20	32
25GbE ports (10GbE autoranging)	20	16	16	64
10GbE ports	n/a	32	32	64
12Gb/6Gb SAS ports	8	32	n/a	64
Storage networking supported	NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3	NVMe/RDMA, NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3	NVMe/RDMA, NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3	NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3
OS version	ONTAP 9.8 RC1 or later	ONTAP 9.7 RC1 or later	ONTAP 9.7 RC1 or later	ONTAP 9.10.1 RC2 or later

Below are few advantages of NetApp AFF:

- Maximum performance for your most demanding applications

NetApp AFF A-Series systems deliver industry-leading performance proven by SPC-1 and SPEC SFS industry benchmarks, making them ideal for demanding, highly transactional applications such as Oracle, Microsoft SQL Server, MongoDB databases, VDI, and server virtualization.

With the power of front-end NVMe/FC and NVMe/TCP host connectivity and back-end NVMe-attached SSDs, our high-end AFF A900 systems deliver latency as low as 100µs. Based on a high-resiliency design, the A900 also delivers high RAS and enables non-disruptive in-chassis upgrade from its predecessor A700. The A800 delivers high performance in a compact form factor and is especially suited for EDA and Media & Entertainment workloads. The midrange, most versatile NetApp AFF A400 system features hardware acceleration technology that significantly enhances performance and storage efficiency. And our entry-level, budget-friendly NetApp AFF A250, provides 40% more performance and 33% more efficiency at no extra cost compared with its predecessor.

- NetApp AFF A-Series also lets you:
  - Drive mission-critical SAN workloads with symmetric active-active host connectivity for continuous availability and instant failover.
  - Consolidate workloads to deliver up to 14.4 million IOPS at 1ms latency in a cluster with a truly unified scale-out architecture. Built-in adaptive quality of service (QoS) safeguards SLAs in multi-workload and multitenant environments.
  - Manage massively scalable NAS containers of up to 20PB and 400 billion files with a single namespace.

- Improve the speed and productivity of collaboration across multiple locations and increase data throughput for read-intensive applications with NetApp FlexCache® software.
- Modernize with advanced connectivity
- NetApp AFF A-Series all-flash systems deliver industry-leading performance, density, scalability, security, and network connectivity. As the first enterprise-grade storage systems to support both NVMe/TCP and NVMe/FC, NetApp AFF A-Series systems boost performance with modern network connectivity. With NVMe/TCP, which uses the commonly available Ethernet infrastructure, you don't have to invest in new hardware to take advantage of the faster host connectivity. With NVMe/FC, you can get twice the IOPS and cut application response time in half compared with traditional FC. These systems support a range of ecosystems, including VMware, Microsoft Windows 10, and Linux, with storage path failover. For most customers, integrating NVMe/FC into an existing SAN is a simple, nondisruptive software upgrade.

- Scale without disruption

With NetApp AFF A-Series, you can integrate new technologies and private or public cloud into your infrastructure nondisruptively. NetApp AFF A-Series is the only all-flash array that enables you to combine different controllers, SSD sizes, and new technologies so that your investment is protected. The NVMe-based AFF systems also support SAS SSDs, maximizing the flexibility and cost effectiveness of your upgrade:

- Best balance between price, technology, features, and performance.
- Increase operational efficiency
- IT departments are striving to make budgets go further and to allow IT staff to focus on new value-added projects rather than on day-to-day IT management. NetApp AFF systems simplify IT operations, which therefore reduces datacenter cost. In particular, our entry-level system, the NetApp AFF A250, delivers best-in-class performance and efficiency to mid-size business customers so they can consolidate more workloads and eliminate silos.

- Provision storage in minutes

NetApp AFF systems offer broad application ecosystem support and deep integration for enterprise applications, virtual desktop infrastructure (VDI), database, and server virtualization, supporting Oracle, Microsoft SQL Server, VMware, SAP, MySQL, and more. You can quickly provision storage in less than 10 minutes with NetApp ONTAP System Manager. In addition, infrastructure management tools simplify and automate common storage tasks so you can:

- Easily provision and rebalance workloads by monitoring clusters and nodes.
- Use one-click automation and self-service for provisioning and data protection.
- Upgrade OS and firmware with a single-click
- Import LUNs from third-party storage arrays directly into an AFF system to seamlessly migrate data.

Additionally, the NetApp® Active IQ® Digital Advisor engine enables you to optimize your NetApp systems with predictive analytics and proactive support. Fueled by the massive NetApp user base, AI and machine learning create actionable insights that help you prevent problems, optimize your configuration, save time, and make smarter decisions.

- Achieve outstanding storage savings

NetApp employs various capabilities to promote optimal capacity savings and to drive down your TCO. AFF A-Series system's support for solid-state drives (SSDs) with multistream write technology, combined with advanced SSD partitioning, provides maximum usable capacity, regardless of the type of data that

---

you store. Thin provisioning; NetApp Snapshot™ copies; and inline data reduction features, such as deduplication, compression, and compaction, provide substantial additional space savings—without affecting performance—enabling you to purchase the least amount of storage capacity possible.

- Build your hybrid cloud with ease

Your data fabric built by NetApp helps you simplify and integrate data management across cloud and on-premises environments to meet business demands and gain a competitive edge. With AFF A-Series, you can connect to more clouds for more data services, data tiering, caching, and disaster recovery. You can also:

- Maximize performance and reduce overall storage costs by automatically tiering cold data to the cloud with FabricPool.
- Instantly deliver data to support efficient collaboration across your hybrid cloud
- Protect your data by taking advantage of Amazon Simple Storage Service (Amazon S3) cloud resources—on premises and in the public cloud.
- Accelerate read performance for data that is shared widely throughout your organization and across hybrid cloud deployments.
- Keep data available, protected, and secure

As organizations become more data driven, the business impact of data loss can be increasingly dramatic—and costly. IT must protect data from both internal and external threats, ensure data availability, eliminate maintenance disruptions, and quickly recover from failures.

- Integrated data protection

AFF A-Series systems come with a full suite of acclaimed NetApp integrated and application-consistent data protection software. Key capabilities include:

- Native space efficiency with cloning and NetApp Snapshot copies reduce storage costs and minimize performance impact. Up to 1,023 copies are supported.
- [NetApp® SnapCenter®](#) software provides application-consistent data protection and clone management to simplify application management.
- [NetApp® SnapMirror®](#) technology replicates to any NetApp FAS or AFF system on the premises or in the cloud, reducing overall system costs.
- Business continuity and fast disaster recovery

With AFF, you can maintain constant data availability with zero data loss and zero downtime. NetApp MetroCluster™ software provides synchronous replication to protect your entire system, and NetApp SnapMirror Business Continuity provides a more flexible, cost-effective business continuity to even with more granular replication of selected critical data.

- Security everywhere

Flexible encryption and key management help guard your sensitive data on the premises, in the cloud, and in transit. The market-leading anti-ransomware protection for both preemption and post-attack recovery safeguards your critical data from ransomware attacks and can prevent catastrophic financial consequences. With the simple and efficient security solutions, you can:

- Achieve FIPS 140-2 compliance (Level 1 and Level 2) with self-encrypting drives and use any type of drives with software-based encryption.



- Meet governance, risk, and compliance requirements with security features such as secure purge; logging and auditing monitors; and write once, read many (WORM) file locking.
- Protect against threats with multifactor authentication, role-based access control, secure multitenancy, and storage-level file security.

## 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.

Figure 18. NetApp AFF A400 Front View



Figure 19. NetApp AFF A400 Rear View



**Note:** We used 4 port 32Gb FC HBA on slot 1 (1a,1b, other two ports unused) for front-end FC SAN connection, 4x25Gb Ethernet NICs on slot 0 (e0e, e0f, e0g, e0h) for NAS connectivity, 2x100Gb ethernet ports on slot 3 (e3a, e3b) used for cluster interconnect, 2x25Gb ethernet on slot 0 (e0a, e0b) used for Node HA interconnect, 2x100Gb ethernet on slot 0 (e0c, e0d) and 2x100Gb ethernet on slot 5 (e5a, e5b) are used for backend NVMe storage connectivity.

## NetApp ONTAP 9

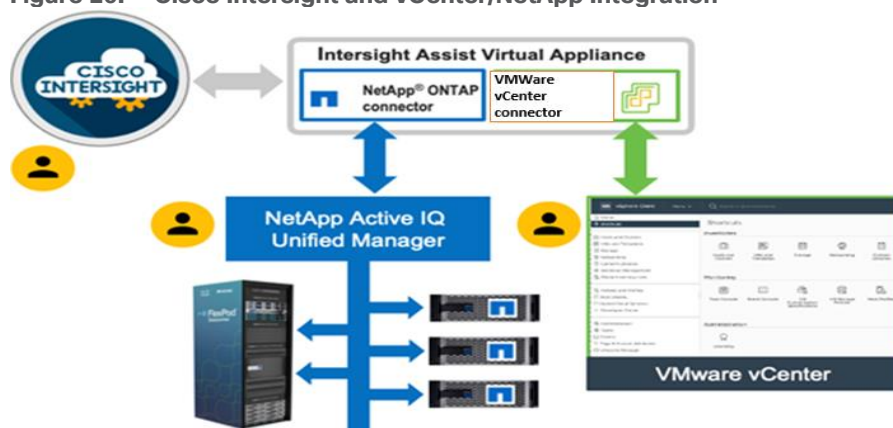
NetApp storage systems harness the power of ONTAP to simplify the data infrastructure from edge, core, and cloud with a common set of data services and 99.9999 percent availability. NetApp ONTAP 9 data management software from NetApp enables customers to modernize their infrastructure and transition to a cloud-ready

datacenter. 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 AFF, FAS or ASA series arrays and in private, public, or hybrid clouds (NetApp Private Storage and NetApp Cloud Volumes ONTAP). Specialized implementations offer best-in-class converged infrastructure, featured here as part of the FlexPod Datacenter solution or with access to third-party storage arrays (NetApp FlexArray virtualization). Together these implementations form the basic framework of the NetApp Data Fabric, with a common software-defined approach to data management, and fast efficient replication across systems. FlexPod and ONTAP architectures can serve as the foundation for both hybrid cloud and private cloud designs.

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

Figure 20. Cisco Intersight and vCenter/NetApp Integration



## NetApp ONTAP 9.10.1P1

### NetApp ONTAP Features for VDI

The following are the ONTAP features for VDI:

- Secure Multi-Tenancy—Tenants can be in overlapping subnet or can use identical IP subnet range.
- Multi-Protocol—Same storage system can be used for Block/File/Object storage demands.
- FlexGroup Volumes—High performance and massive capacity (~20PB and ~40 billion files) for file shares and for hosting VDI pools.
- FlexCache—Enables Single Global Namespace can be consumed around the clouds or multi-site.
- File System Analytics—Fast query to file metadata on the SMB file share.
- Ease of management with vCenter Plugins—Best practices are validated and implemented while provisioning. Supports VAAI and VASA for fast provisioning and storage capability awareness.
- SnapCenter integration with vCenter—Space efficient data protection with snapshots and FlexClones.
- Automation support—Supports RESTapi, has modules for Ansible, PowerShell, and so on.

- Storage Efficiency—Supports inline dedupe, compression, thin provisioning, etc. Guaranteed dedupe of 8:1 for VDI.
- Adaptive QoS—Adjusts QoS setting based on space consumption.
- ActiveIQ Unified Manager—Application based storage provisioning, Performance Monitoring, End-End storage visibility diagrams.

## Storage Efficiency

Storage efficiency has always been a primary architectural design point of ONTAP. A wide array of features allows businesses to store more data using less space. In addition to deduplication and compression, businesses can store their data more efficiently by using features such as unified storage, multi-tenancy, thin provisioning, and NetApp Snapshot® technology.

Starting with ONTAP 9, NetApp guarantees that the use of NetApp storage efficiency technologies on AFF systems reduce the total logical capacity used to store customer data by 75 percent, a data reduction ratio of 4:1. This space reduction is a combination of several different technologies, such as deduplication, compression, and compaction, which provide additional reduction to the basic features provided by ONTAP.

Compaction, which is introduced in ONTAP 9, is the latest patented storage efficiency technology released by NetApp. In the NetApp WAFL® file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on the disk-to-save space. This process is illustrated in [Figure 21](#).

## Storage Efficiency Features

The storage efficiency features are as follows:

- Deduplication

Deduplication reduces the amount of physical storage required for a volume (or all the volumes in an AFF aggregate) by discarding duplicate blocks and replacing them with references to a single shared block. Reads of deduplicated data typically incur no performance charge. Writes incur a negligible charge except on overloaded nodes.

As data is written during normal use, WAFL uses a batch process to create a catalog of block signatures. After deduplication starts, ONTAP compares the signatures in the catalog to identify duplicate blocks. If a match exists, a byte-by-byte comparison is done to verify that the candidate blocks have not changed since the catalog was created. Only if all the bytes match is the duplicate block discarded and its disk space reclaimed.

- Compression

Compression reduces the amount of physical storage required for a volume by combining data blocks in compression groups, each of which is stored as a single block. Reads of compressed data are faster than in traditional compression methods because ONTAP decompresses only the compression groups that contain the requested data, not an entire file or LUN.

You can perform inline or postprocess compression, separately or in combination:

- Inline compression compresses data in memory before it is written to disk, significantly reducing the amount of write I/O to a volume, but potentially degrading write performance. Performance-intensive operations are deferred until the next postprocess compression operation, if any.
- Postprocess compression compresses data after it is written to disk, on the same schedule as deduplication.

- Compaction
- Small files or I/O padded with zeros are stored in a 4 KB block whether or not they require 4 KB of physical storage. Inline data compaction combines data chunks that would ordinarily consume multiple 4 KB blocks into a single 4 KB block on disk. Compaction takes place while data is still in memory, so it is best suited to faster controllers

Figure 21. Storage Efficiency Features

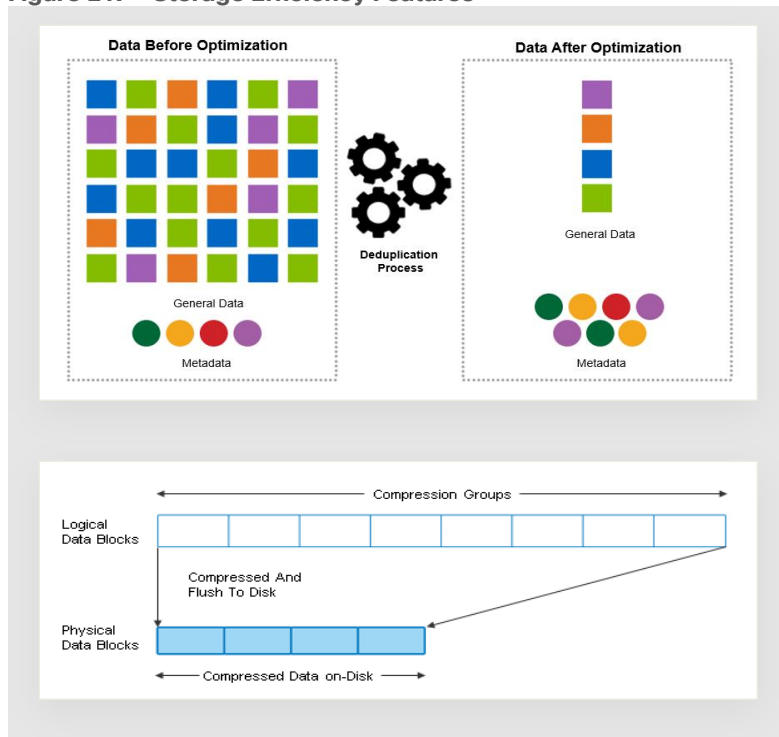
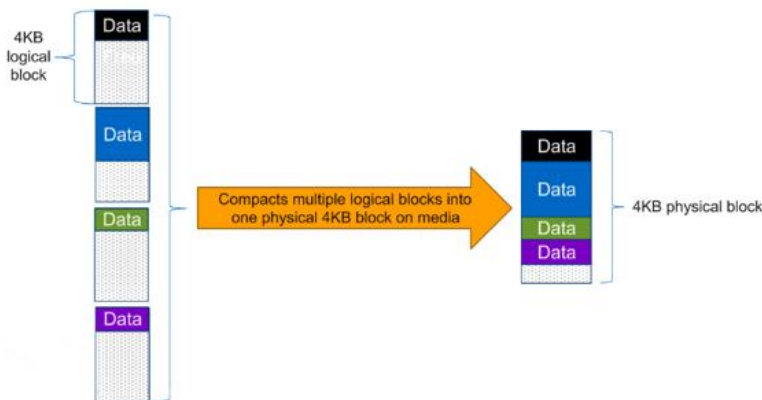


Figure 22. Storage Efficiency



**Note:** Some applications such as Oracle and SQL have unique headers in each of their data blocks that prevent the blocks to be identified as duplicates. So, for such applications, enabling deduplication does not result in significant savings. So, deduplication is not recommended to be enabled for databases. However, NetApp data compression works very well with databases, and we strongly recommend enabling compression for databases. [Table 2](#) lists some guidelines where compression, deduplication and/or inline Zero block deduplication can be used. These are guidelines, not rules; environment may have different performance requirements and specific use cases.

**Table 2. Compression and Deduplication Guidelines**

Workload	Storage Efficiency Guidelines		
	All Flash FAS (AFF)	Flash Pool (Sized as per Flash Pool Best Practice)	Hard Disk Drives
<b>Database (Oracle, SQL)</b>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Inline zero-block deduplication</li> <li>Inline deduplication (Data ONTAP 8.3.2 and above)</li> </ul>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Inline zero-block deduplication</li> <li>Inline deduplication (Data ONTAP 8.3.2 and above)</li> </ul>	For primary workloads, use: <ul style="list-style-type: none"> <li>Inline zero-block deduplication</li> </ul> For secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Adaptive background compression</li> <li>Inline zero-block deduplication</li> </ul>
<b>VDI and SVI</b>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> <li>Inline deduplication (Data ONTAP 8.3.2 and above)</li> </ul>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> <li>Inline deduplication (Data ONTAP 8.3.2 and above)</li> </ul>	For primary workloads, use: <ul style="list-style-type: none"> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul> For secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Adaptive background compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>
<b>Exchange</b>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Set schedule to off peak hours</li> <li>Inline zero-block</li> </ul>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Inline secondary compression</li> <li>Background secondary compression</li> <li>Deduplication</li> </ul>

Workload	Storage Efficiency Guidelines		
	All Flash FAS (AFF)	Flash Pool (Sized as per Flash Pool Best Practice)	Hard Disk Drives
<b>File Services</b>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>	deduplication <ul style="list-style-type: none"> <li>Inline zero-block deduplication</li> </ul> For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Adaptive background compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>
<b>Mixed Workload</b>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>	For primary and secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>	For primary workloads, use: <ul style="list-style-type: none"> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul> For secondary workloads, use: <ul style="list-style-type: none"> <li>Adaptive inline compression</li> <li>Adaptive background compression</li> <li>Deduplication</li> <li>Inline zero-block deduplication</li> </ul>

**Space Savings**

[Table 3](#) lists the storage efficiency data reduction ratio ranges for different applications. A combination of synthetic datasets and real-world datasets has been used to determine the typical savings ratio range. The savings ratio range mentioned is only indicative.

**Table 3. Typical Savings Ratios with ONTAP 9—Sample Savings Achieved with Internal and Customer Testing**

Typical Savings Ratios with ONTAP 9	
Workload [with deduplication, data compaction, adaptive compression and	Ratio Range


FlexClone volumes (where applicable) technologies]	
Home directories	1.5:1.-.2:1
Software development	2:1 - 10:1
VDI VMware Horizon full clone desktops (persistent)	6:1 - 10:1
VDI VMware Horizon Instant clone desktops (nonpersistent)	5:1 - 7:1
VDI VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions Instant clone desktops (nonpersistent)	6:1 - 10:1
Virtual Servers (OS and Applications)	2:1.-.4:1
Oracle databases (with no database compression)	2.1 - 4:1
SQL 2014 databases (with no database compression)	2.1 - 4:1
Microsoft Exchange	1.6:1
Mongo DB	1.3:1 - 1.5:1
Recompressed data (such as video and image files, audio files, pdfs, and so on)	No Savings

### NetApp Storage Virtual Machine (SVM)

An SVM is a logical abstraction that represents the set of physical resources of the cluster. This adds extra security and peace of mind to your VDI environment, giving you another place besides vCenter to apply HA, High Availability. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and may reside on any node in the cluster to which the SVM has been given access. An SVM may own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node to another. For example, a flexible volume can be non-disruptively moved to a new node and aggregate, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware, and it is not tied to any specific physical hardware.

An SVM can support multiple data protocols concurrently. Volumes within the SVM can be joined together to form a single NAS namespace, which makes all of an SVM's data available through a single share or mount point to create a VMware NFS datastore for your VDI desktop folders. SVMs also support block-based protocols, and LUNs can be created and exported by using iSCSI, FC, or FCoE. Any or all of these data protocols can be configured for use within a given SVM to support your VDI needs.

**Figure 23. NetApp Storage Virtual Machine**



*Service providers use SVMs in multitenant environments to isolate tenant data and simplify chargeback.*


## FlexClones

FlexClone technology references Snapshot metadata to create writable, point-in-time copies of a volume. Copies share data blocks with their parents, consuming no storage except what is required for metadata until changes are written to the copy. FlexClone files and FlexClone LUNs use identical technology, except that a backing Snapshot copy is not required.

Where traditional copies can take minutes or even hours to create, FlexClone software lets you copy even the largest datasets almost instantaneously. That makes it ideal for situations in which you need multiple copies of identical datasets (a virtual desktop deployment, for example) or temporary copies of a dataset (testing an application against a production dataset).

You can clone an existing FlexClone volume, clone a volume containing LUN clones, or clone mirror and vault data. You can split a FlexClone volume from its parent, in which case the copy is allocated its own storage.

**Figure 24. Traditional Copy vs FlexClone Copy**



*FlexClone copies share data blocks with their parents, consuming no storage except what is required for metadata.*

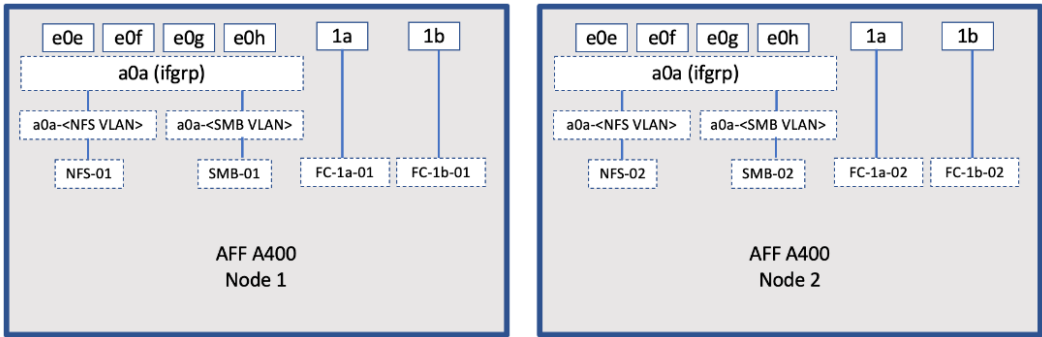
## SAN Boot

NetApp recommends implementing SAN boot for Cisco UCS servers in the FlexPod Datacenter solution. Doing so enables the ESXI host to be safely secured by the NetApp All Flash FAS storage system, providing better performance. In this design, FC SAN boot is validated.

In FC SAN boot, each Cisco UCS server boots by connecting the NetApp All Flash FAS storage to the Cisco MDS switch. The 32G FC storage ports, in this example 0g and 0h, are connected to Cisco MDS switch. The FC LIFs are created on the physical ports and each FC LIF is uniquely identified by its target WWPN. The storage

system target WWPNs can be zoned with the server initiator WWPNs in the Cisco MDS switches. The FC boot LUN is exposed to the servers through the FC LIF using the MDS switch; this enables only the authorized server to have access to the boot LUN.

**Figure 25. FC - SVM ports and LIF layout**




Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead if a network interface becomes unavailable, the ESXI host chooses a new optimized path to an available network interface. ALUA is a standard supported by NetApp used to provide information about SCSI targets, which allows a host to identify the best path to the storage.

**FlexGroups**

ONTAP 9.3 brought an innovation in scale-out NAS file systems: NetApp FlexGroup volumes, which plays a major role to give ONTAP the ability to be scaled nondisruptively out to 24 storage nodes while not degrading the performance of the VDI infrastructure.

With FlexGroup volumes, a storage administrator can easily provision a massive single namespace in a matter of seconds. FlexGroup volumes have virtually no capacity or file count constraints outside of the physical limits of hardware or the total volume limits of ONTAP. Limits are determined by the overall number of constituent member volumes that work in collaboration to dynamically balance load and space allocation evenly across all members. There is no required maintenance or management overhead with a FlexGroup volume. You simply create the FlexGroup volume and share it with your NAS clients. ONTAP does the rest.

**Figure 26. NetApp FlexGroups**



**Storage QoS**

Storage QoS (Quality of Service) can help you manage risks around meeting your performance objectives. You use Storage QoS to limit the throughput to workloads and to monitor workload performance. You can reactively limit workloads to address performance problems and you can pro-actively limit workloads to prevent performance problems.

A workload represents the input/output (I/O) operations to one of the following kinds of storage objects:




- FlexVol volumes
- LUNs

You assign a storage object to a policy group to control and monitor a workload. You can monitor workloads without controlling them.

[Figure 27](#) illustrates an example environment before and after using Storage QoS. On the left, workloads compete for cluster resources to transmit I/O. These workloads get "best effort" performance, which means you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right are the same workloads assigned to policy groups. The policy groups enforce a maximum throughput limit.

**Figure 27. Before and After using Storage QoS**



NetApp storage quality of service (QoS) works with both SAN and NAS storage, and it runs across the entire NetApp product line from entry to enterprise. Storage QoS offers significant benefits for all types of VDI environments. It lets you:

- Achieve greater levels of consolidation
- Set maximum and minimum limits on multiple VDI workloads that require separate service level agreements (SLAs)
- Add additional workloads with less risk of interference
- Make sure your customers get what they pay for, but not more

## Adaptive QoS

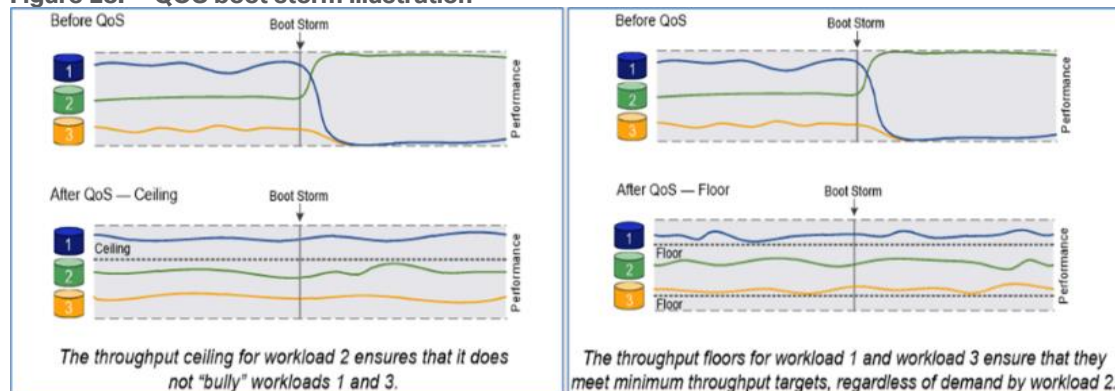
Adaptive QoS automatically scales the policy group (A *policy group* defines the throughput ceiling for one or more workloads) value to workload (A *workload* represents the I/O operations for a storage object: a volume, file, qtree or LUN, or all the volumes, files, qtrees, or LUNs in an SVM) size, for the size of the workload changes. That is a significant advantage when you are managing hundreds or thousands of workloads in a VDI deployment. With Adaptive QoS, Ceiling and Floor limit can be set using allocated or used space. The QoS also address HA and Scaling as it will assist in both efforts to produce a non-disruptive change during VDI growth by maintaining the ratio of IOPS to TBs/GBs. To assist in managing your QoS, Active IQ unified manager will provide QoS suggestions based on historical performance and usage.

Three default adaptive QoS policy groups are available, as shown in [Table 4](#). You can apply these policy groups directly to a volume.

**Table 4. Available Default Adaptive QoS Policy Groups**

Default Policy Group	Expected IOPS/TB	Peak IOPS/TB	Absolute Min IOPS
extreme	6,144	12,288	1000
performance	2,048	4,096	500
Value	128	512	75

**Figure 28. QoS boot storm illustration**



## Security and Data Protection

### Vscan

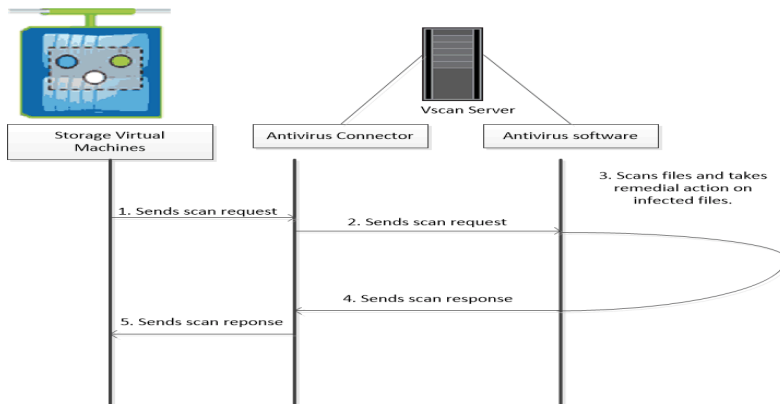
With Vscan you can use integrated antivirus functionality on NetApp storage systems to protect data from being compromised by viruses or other malicious code. NetApp virus scanning, called Vscan, combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when.

Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. The ONTAP Antivirus Connector, provided by NetApp and installed on the external server, handles communication between the storage system and the antivirus software.

You can use *on-access scanning* to check for viruses when clients open, read, rename, or close files over CIFS. File operation is suspended until the external server reports the scan status of the file. If the file has already been scanned, ONTAP allows the file operation. Otherwise, it requests a scan from the server.

You can use *on-demand scanning* to check files for viruses immediately or on a schedule. You might want to run scans only in off-peak hours, for example. The external server updates the scan status of the checked files, so that file-access latency for those files (assuming they have not been modified) is typically reduced when they are next accessed over CIFS. You can use on-demand scanning for any path in the SVM namespace, even for volumes that are exported only through NFS.

Typically, you enable both scanning modes on an SVM. In either mode, the antivirus software takes remedial action on infected files based on your settings in the software.

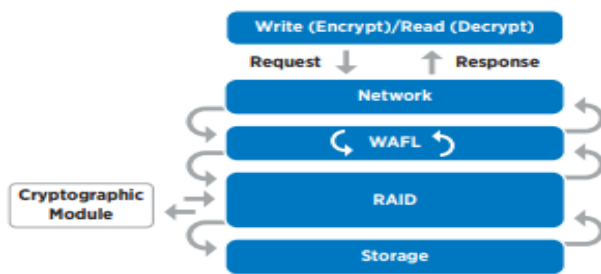


## NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE)

NetApp Volume Encryption is a software-based, data-at-rest encryption solution that is FIPS 140-2 compliant. NVE allows ONTAP to encrypt data for each volume for granularity. NAE, is an outgrowth of NVE; it allows ONTAP to encrypt data for each volume, and the volumes can share keys across the aggregate. NVE and NAE enable you to use storage efficiency features that would be lost with encryption at the application layer. For greater storage efficiency, you can use aggregate deduplication with NAE.

Here's how the process works: The data leaves the disk encrypted, is sent to RAID, is decrypted by the CryptoMod, and is then sent up the rest of the stack. This process is illustrated in [Figure 29](#).

**Figure 29. NVE and NAE Process**



To view the latest security features for ONTAP 9, go to: [Security Features in ONTAP 9 | NetApp](#).

## ONTAP Rest API

ONTAP Rest API enables you to automate the deployment and administration of your ONTAP storage systems using one of several available options. The ONTAP REST API provides the foundation for all the various ONTAP automation technologies.

Beginning with ONTAP 9.6, ONTAP includes an expansive workflow-driven REST API that you can use to automate deployment and management of your storage. In addition, NetApp provides a Python client library, which makes it easier to write robust code, as well as support for ONTAP automation based on Ansible.

## AutoSupport and Active IQ Digital Advisor

ONTAP offers artificial intelligence-driven system monitoring and reporting through a web portal and through a mobile app. The AutoSupport component of ONTAP sends telemetry that is analyzed by Active IQ Digital Advisor. Active IQ enables you to optimize your data infrastructure across your global hybrid cloud by delivering actionable predictive analytics and proactive support through a cloud-based portal and mobile app. Data-driven

---

insights and recommendations from Active IQ are available to all NetApp customers with an active SupportEdge contract (features vary by product and support tier).

The following are some things you can do with Active IQ:

- Plan upgrades. Active IQ identifies issues in your environment that can be resolved by upgrading to a newer version of ONTAP and the Upgrade Advisor component helps you plan for a successful upgrade.
- View system wellness. Your Active IQ dashboard reports any issues with wellness and helps you correct those issues. Monitor system capacity to make sure you never run out of storage space.
- Manage performance. Active IQ shows system performance over a longer period than you can see in ONTAP System Manager. Identify configuration and system issues that are impacting your performance.
- Maximize efficiency. View storage efficiency metrics and identify ways to store more data in less space.
- View inventory and configuration. Active IQ displays complete inventory and software and hardware configuration information. View when service contracts are expiring to ensure you remain covered.

## VMware vSphere 7.0

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

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

- Fully featured vSphere Client (HTML5) client. (The flash-based vSphere Web Client has been deprecated and is no longer available.)
- Improved Distributed Resource Scheduler (DRS) – a very different approach that results in a much more granular optimization of resources
- Assignable hardware – a new framework that was developed to extend support for vSphere features when customers utilize hardware accelerators
- vSphere Lifecycle Manager – a replacement for VMware Update Manager, bringing a suite of capabilities to make lifecycle operations better
- Refactored vMotion – improved to support today’s workloads

For more information about VMware vSphere and its components, see:

<https://www.vmware.com/products/vsphere.html>.

## VMware vSphere vCenter

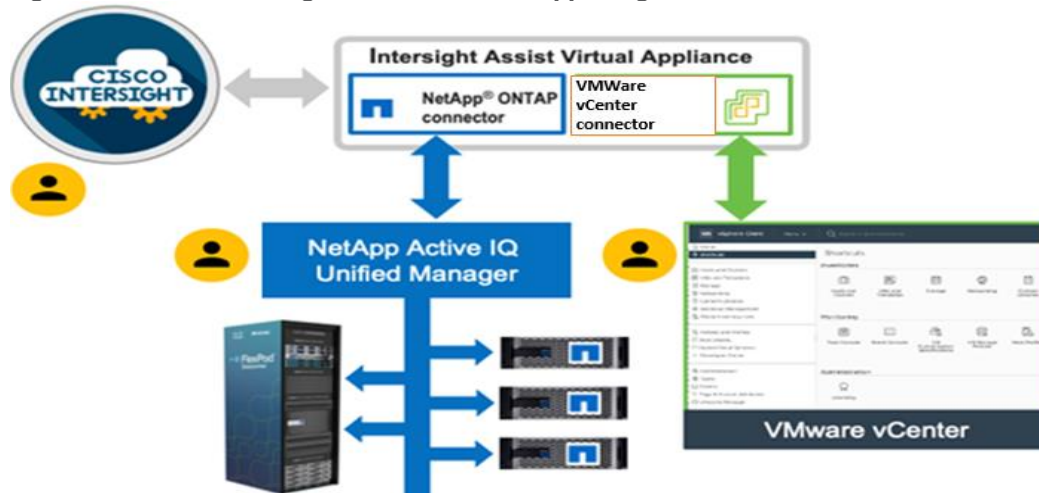
VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest 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 and NetApp ONTAP

Cisco Intersight integrates with VMware vCenter and NetApp storage 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.

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 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 and NetApp Active IQ Unified Manager 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.

## NetApp ONTAP Tools for VMware vSphere

NetApp ONTAP tools for VMware vSphere is a unified appliance that includes vSphere Storage Console (VSC), VASA Provider and SRA Provider. This vCenter web client plug-in that provides Context sensitive menu to provision traditional datastores & Virtual Volume (vVol) datastore.

ONTAP tools provides visibility into the NetApp storage environment from within the vSphere web client. VMware administrators can easily perform tasks that improve both server and storage efficiency while still using role-based access control to define the operations that administrators can perform. It includes enhanced REST APIs that provide vVols metrics for SAN storage systems using ONTAP 9.7 and later. So, NetApp OnCommand API Services is no longer required to get metrics for ONTAP systems 9.7 and later.

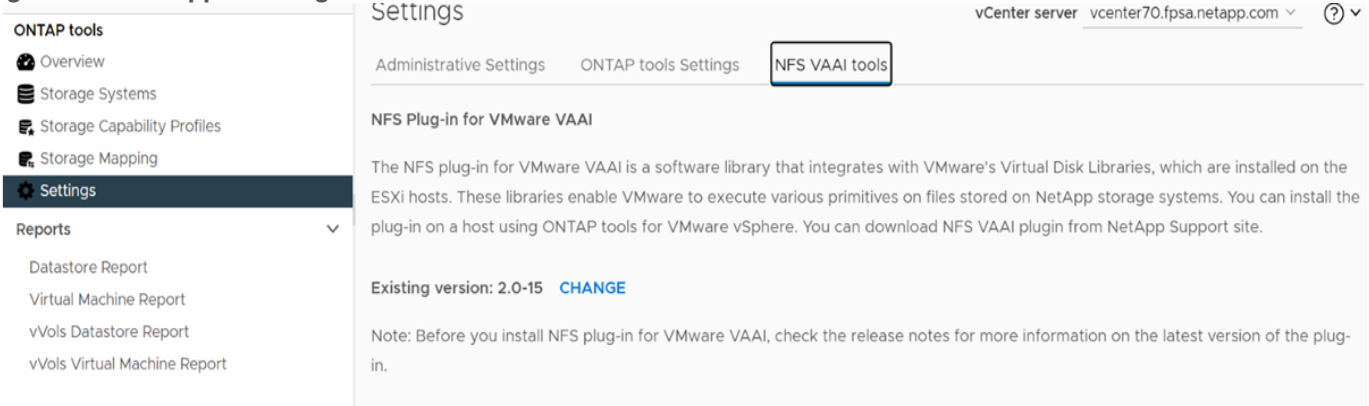
To download ontap tools for vmware vsphere, go to:

<https://mysupport.netapp.com/site/products/all/details/otv/downloads-tab>.

## NetApp NFS Plug-in for VMware VAAI

The NetApp NFS Plug-in for VMware vStorage APIs - Array Integration (VAAI) is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. The VMware VAAI package enables the offloading of certain tasks from the physical hosts to the storage array. Performing those tasks at the array level can reduce the workload on the ESXi hosts.

Figure 31. NetApp NFS Plug-in for VMware VAAI



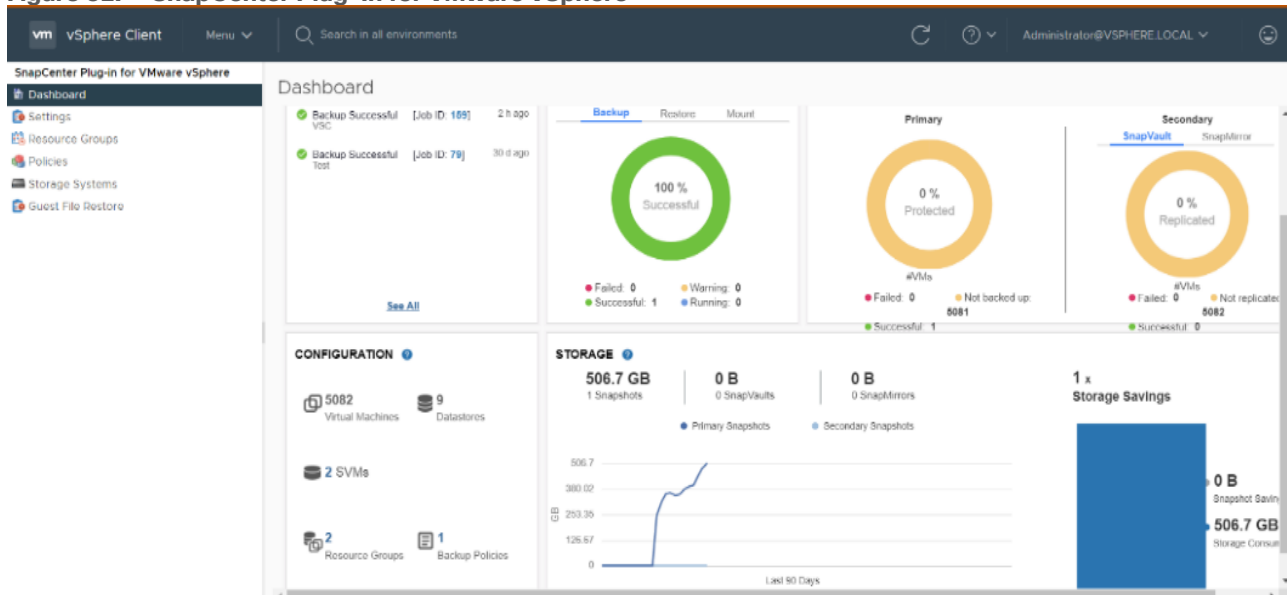
The copy offload feature and space reservation feature improve the performance of VSC operations. The NetApp NFS Plug-in for VAAI is not shipped with VSC, but you can install it by using VSC. You can download the plug-in installation package and obtain the instructions for installing the plug-in from the NetApp Support Site.

For more information about the NetApp VSC for VMware vSphere, see the [NetApp Virtual Infrastructure Management Product Page](#).

## NetApp SnapCenter Plug-In for VMware vSphere

NetApp SnapCenter Plug-in for VMware vSphere enables VM-consistent and crash-consistent backup and restore operations for VMs and datastores from the vCenter server. The SnapCenter plug-in is deployed as a virtual appliance, and it integrates with the vCenter server web client GUI.

Figure 32. SnapCenter Plug-In for VMware vSphere



---

Here are some of the functionalities provided by the SnapCenter plug-in to help protect your VMs and datastores:

- Backup VMs, virtual machine disks (VMDKs), and datastores
  - You can back up VMs, underlying VMDKs, and datastores. When you back up a datastore, you back up all the VMs in that datastore.
  - You can create mirror copies of backups on another volume that has a SnapMirror relationship to the primary backup or perform a disk-to-disk backup replication on another volume that has a NetApp SnapVault® relationship to the primary backup volume.
  - Backup operations are performed on all the resources defined in a resource group. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.
- Restore VMs and VMDKs from backups
  - You can restore VMs from either a primary or secondary backup to the same ESXi server. When you restore a VM, you overwrite the existing content with the backup copy that you select.
  - You can restore one or more VMDKs on a VM to the same datastore. You can restore existing
- VMDKs, or deleted or detached VMDKs from either a primary or a secondary backup
  - You can attach one or more VMDKs from a primary or secondary backup to the parent VM (the same VM that the VMDK was originally associated with) or an alternate VM. You can detach the VMDK after you have restored the files you need.
  - You can restore a deleted VM from a datastore primary or secondary backup to an ESXi host that you select.

**Note:** For application-consistent backup and restore operations, the NetApp SnapCenter Server software is required.

**Note:** For additional information, requirements, licensing, and limitations of the NetApp SnapCenter Plug-In for VMware vSphere, see the [NetApp Product Documentation](#).

## NetApp Active IQ Unified Manager 9.10.1P1

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

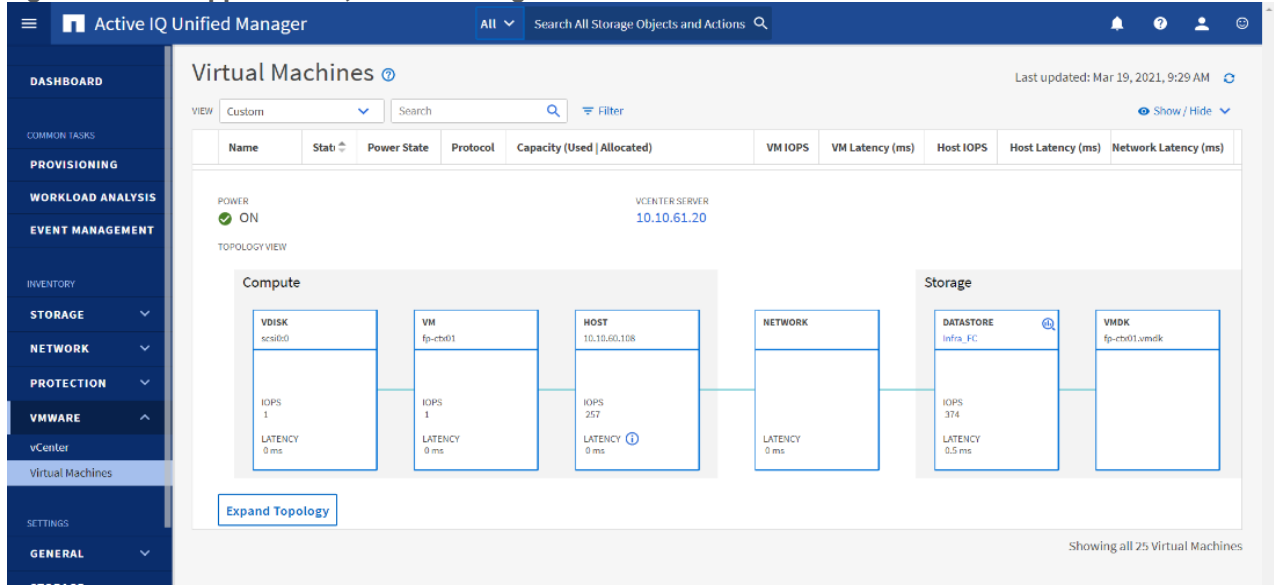
Active IQ Unified Manager enables monitoring your ONTAP storage clusters, VMware vCenter server and VMs from a single redesigned, intuitive interface that delivers intelligence from community wisdom and AI analytics. It provides comprehensive operational, performance, and proactive insights into the storage environment and the VMs running on it. When an issue occurs on the storage or virtual infrastructure, Active IQ Unified Manager can notify you about the details of the issue to help with identifying the root cause.

Unified Manager enables to manage storage objects in your environment by associating them with annotations. You can create custom annotations and dynamically associate clusters, SVMs, and volumes with the annotations through rules.

The VM dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the vSphere host down through the network and finally to the storage. Some events also

provide remedial actions that can be taken to rectify the issue. You can also configure custom alerts for events so that when issues occur, you are notified through email and SNMP traps.

**Figure 33. NetApp Active IQ Unified Manager Virtual Machine Dashboard**



## NetApp XCP File Analytics

NetApp XCP file analytics is host-based software to scan the file shares, collect and analyzes the data and provide insights into the file system. NetApp XCP file analytics works for both NetApp and non-NetApp systems and runs on Linux or Windows host. For more info, go to: <http://docs.netapp.com/us-en/xcp/index.html>



---

## Architecture and Design Considerations for Desktop Virtualization

This chapter contains the following:

- [Understanding Applications and Data](#)
- [Project Planning and Solution Sizing Sample Questions](#)
- [Hypervisor Selection](#)
- [Desktop Virtualization Design Fundamentals](#)
- [Storage Considerations](#)

There are many reasons to consider a virtual desktop solution such as an ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Device (BYOD) to work programs. The first step in designing a virtual desktop solution is to understand the user community and the type of tasks that are required to successfully execute their role. The following user classifications are provided:

- Knowledge Workers today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- External Contractors are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.
- Task Workers perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.
- Mobile Workers need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- Shared Workstation users are often found in state-of-the-art University and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications for the needs of the organization change, tops the list.

After the user classifications have been identified and the business requirements for each user classification have been defined, it becomes essential to evaluate the types of virtual desktops that are needed based on user requirements. There are essentially five potential desktops environments for each user:

- Traditional PC: A traditional PC is what typically constitutes a desktop environment: a physical device with a locally installed operating system.
- Hosted Shared Desktop: A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2016/2019, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Remoted Desktop Server Hosted Server sessions: A hosted virtual desktop is a virtual desktop running on a virtualization layer (ESX). The user does not work with and sit in front of the desktop, but instead, the user interacts through a delivery protocol.

- **Published Applications:** Published applications run entirely on the VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions RDS server virtual machines and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft Office, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.
- **Streamed Applications:** Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly, but the resources may only be available while they are connected to the network.
- **Local Virtual Desktop:** A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a primary type and is synchronized with the datacenter when the device is connected to the network.

For the purposes of the validation represented in this document, both VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions and Win 10 Virtual Desktops sessions were validated. Each of the sections provides some fundamental design decisions for this environment.

## Understanding Applications and Data

When the desktop user groups and sub-groups have been identified, the next task is to catalog group application and data requirements. This can be one of the most time-consuming processes in the VDI planning exercise but is essential for the VDI project's success. If the applications and data are not identified and co-located, performance will be negatively affected.

The process of analyzing the variety of application and data pairs for an organization will likely be complicated by the inclusion of cloud applications, for example, Salesforce.com. This application and data analysis is beyond the scope of this Cisco Validated Design but should not be omitted from the planning process. There are a variety of third-party tools available to assist organizations with this crucial exercise.

## Project Planning and Solution Sizing Sample Questions

Now that user groups, their applications, and their data requirements are understood, some key project and solution sizing questions may be considered.

General project questions should be addressed at the outset, including:

- Has a VDI pilot plan been created based on the business analysis of the desktop groups, applications, and data?
- Is there infrastructure and budget in place to run the pilot program?
- Are the required skill sets to execute the VDI project available? Can we hire or contract for them?
- Do we have end user experience performance metrics identified for each desktop sub-group?
- How will we measure success or failure?
- What is the future implication of success or failure?

Below is a short, non-exhaustive list of sizing questions that should be addressed for each user sub-group:

- What is the desktop OS planned? Windows 10 or Windows 11?
- 32 bit or 64 bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 10?
- How much memory per target desktop group desktop?

- 
- Are there any rich media, Flash, or graphics-intensive workloads?
  - Are there any applications installed? What application delivery methods will be used, Installed, Streamed, Layered, Hosted, or Local?
  - What is the OS planned for RDS Server Roles? Windows Server 2019 or Server 2022?
  - What is the hypervisor for the solution?
  - What is the storage configuration in the existing environment?
  - Are there sufficient IOPS available for the write-intensive VDI workload?
  - Will there be storage dedicated and tuned for VDI service?
  - Is there a voice component to the desktop?
  - Is anti-virus a part of the image?
  - What is the SQL server version for the database? SQL server 2017 or 2019?
  - Is user profile management (for example, non-roaming profile based) part of the solution?
  - What is the fault tolerance, failover, disaster recovery plan?
  - Are there additional desktop sub-group specific questions?

## Hypervisor Selection

VMware vSphere has been identified for the hypervisor for both VMware Horizon Remoted Server Desktop Hosted (RDSH) Sessions and Win 10 Virtual Desktops.

VMware vSphere: VMware vSphere comprises the management infrastructure or virtual center server software and the hypervisor software that virtualizes the hardware resources on the servers. It offers features like Distributed Resource Scheduler, vMotion, high availability, Storage vMotion, VMFS, and a multi-pathing storage layer. More information on vSphere can be obtained at the VMware website:

<http://www.vmware.com/products/datacentervirtualization/vsphere/overview.html>.

**Note:** For this CVD, the hypervisor used was VMware ESXi 7.0. Update 3.

## Desktop Virtualization Design Fundamentals

An ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Device (BYOD) to work programs are prime reasons for moving to a virtual desktop solution.

## VMware Horizon Design Fundamentals

VMware Horizon 8 integrates Remote Desktop Server Hosted sessions users and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, mixed users and manageable solution for delivering Windows applications and desktops as a service.

Users can select applications from an easy-to-use “store” that is accessible from tablets, smartphones, PCs, Macs, and thin clients. VMware Horizon delivers a native touch-optimized experience via PCoIP or Blast Extreme high-definition performance, even over mobile networks.

## Horizon Remote Desktop Server Hosted (RDSH) Sessions and Win10 VDI Virtual Desktop Pools

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Desktop Pool. In this CVD, VM provisioning relies on VMware View Instant Cloning aligning with VMware Horizon View Connection Server and vCenter Server components. Machines in these Pools are configured to run either a Windows Server 2019 OS (for RDSH Remote Desktop Hosted Server Sessions) or a Windows 10 Desktop OS (for, instant clone and persistent full clone VDI Windows 10 desktops).

**Note:** Server OS and Desktop OS Machines were configured in this CVD to support Remote Desktop Server Hosted (RDSH) Sessions hosted shared desktops and a variety of Win 10 VDI Virtual desktops.

Figure 34. VMware Horizon Design Overview

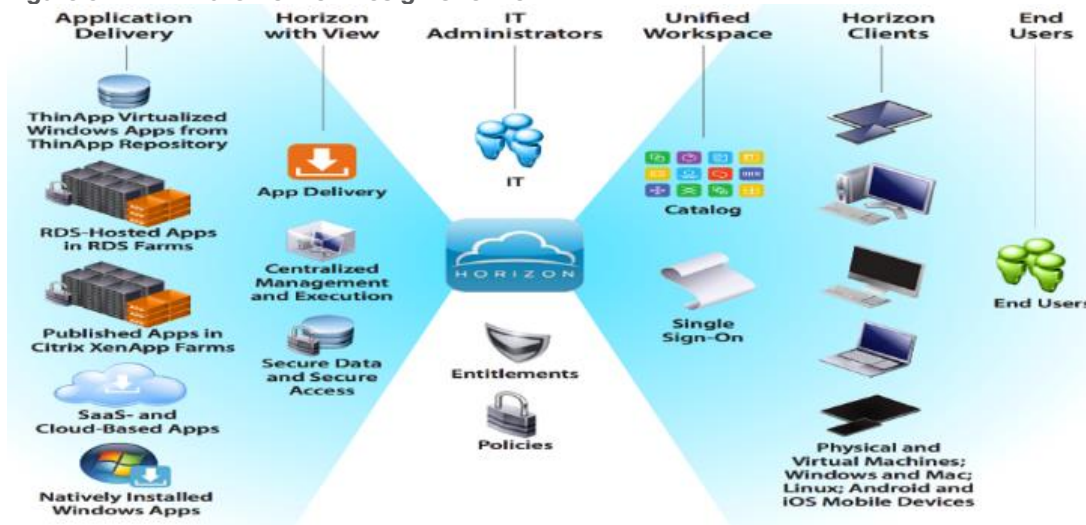
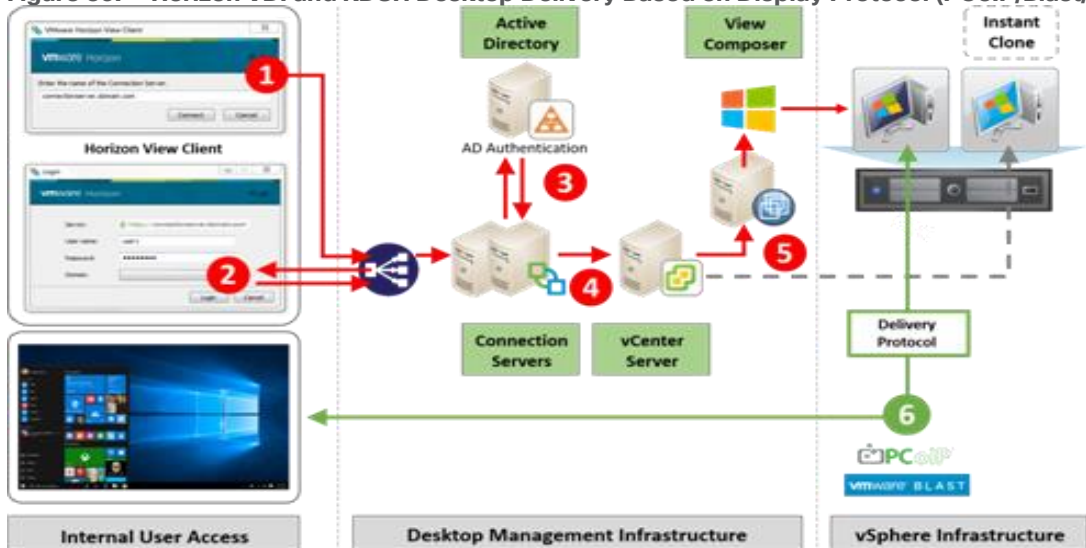


Figure 35. Horizon VDI and RDSH Desktop Delivery Based on Display Protocol (PCoIP/Blast/RDP)

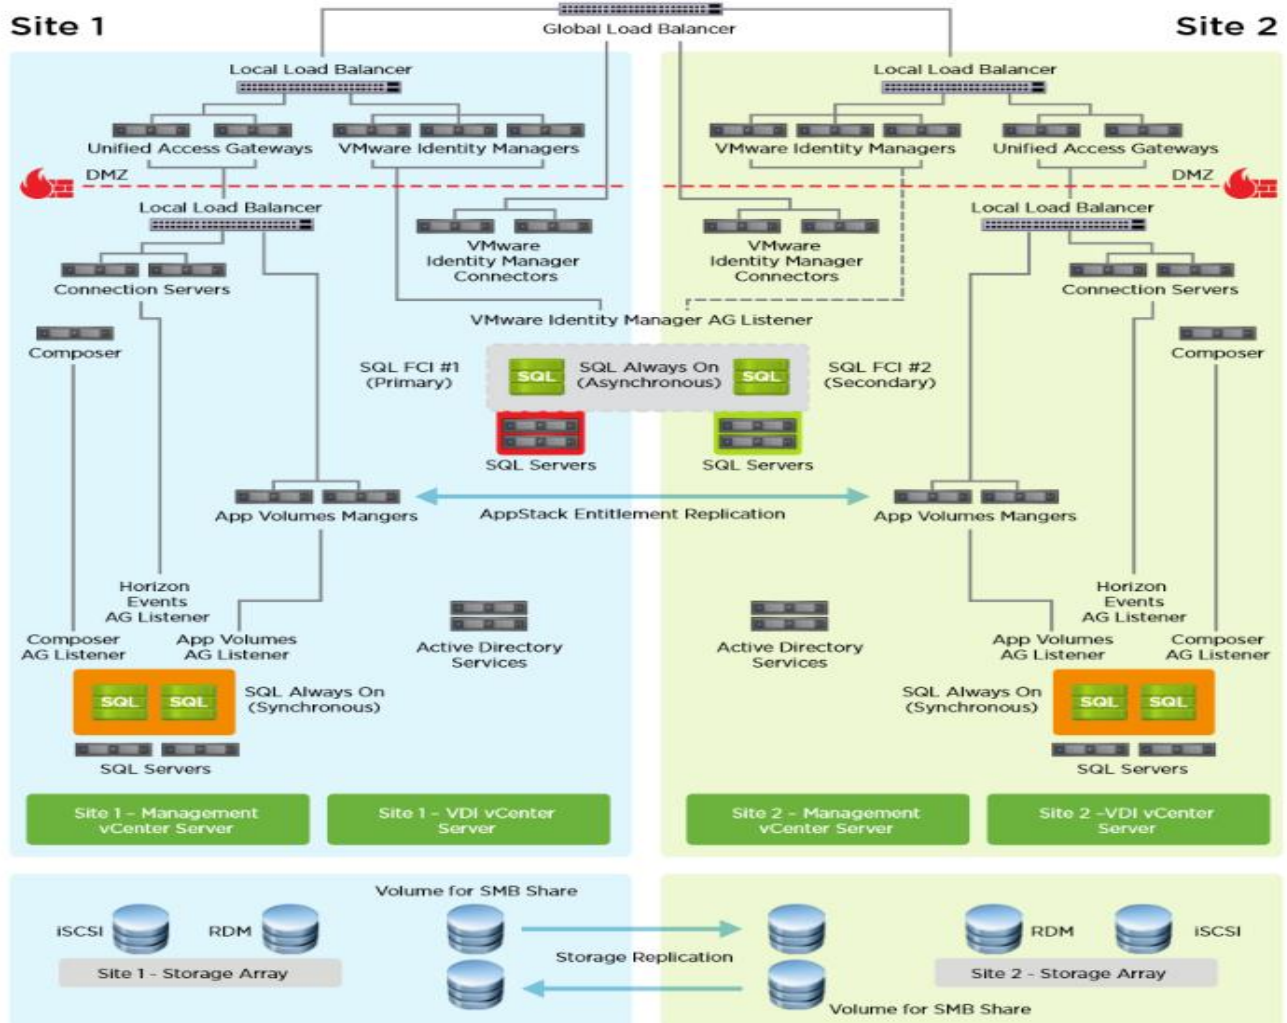


## Multiple-Site Configuration

If you have multiple regional sites, you can use any of the Load Balances Tools to direct the user connections to the most appropriate site to deliver the desktops and application to users.

Figure 36 illustrating sites, shows a site created in two datacenters. Having two sites globally, rather than just one, minimizes the amount of unnecessary WAN traffic. Two Cisco blade servers host the required infrastructure services (Domain Controllers, DNS, DHCP, Profile, SQL, VMware Horizon View Connection Servers, View Composer server and web servers).

Figure 36. Multisite Configuration Overview



Based on the requirement and no of datacenters or remote location, we can choose any of the available Load balancing software or tools accelerates the application performance, load balances servers, increases security, and optimizes the user experience.

**Note:** Multi-Site configuration is shown as the example. Not used as part of this CVD testing

### Designing a VMware Horizon Environment for Various Workload Types

With VMware Horizon 8, the method you choose to provide applications or desktops to users depends on the types of applications and desktops you are hosting and available system resources, as well as the types of users and user experience you want to provide.

Machine	User Type and Experience
Server OS machines	<p>You want: Inexpensive server-based deliver to minimize the cost of delivering applications to a large number of users while providing a secure, high-definition user experience.</p> <p>Your users: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations.</p> <p>Application types: Any application.</p>
Desktop OS machines	<p>You want: A client-base application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition.</p> <p>Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require offline access to hosted applications.</p> <p>Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines. Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.</p>
Remote PC Access	<p>You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the datacenter.</p> <p>Your users: Employees or contractors that have the option to work from home but need access to specific software or data on their corporate desktops to perform their jobs remotely.</p> <p>Host: The same as Desktop OS machines.</p> <p>Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.</p>

For the Cisco Validated Design described in this document, following designs are included:

- Multi-session OS Solution
  - VMware Remote Desktop Servers Hosted (RDSH) Sessions: 2600 Windows Server 2019 random pooled desktops were configured and tested
- Single-session OS Solution:
  - VMware Horizon Instant Clones non-persistent virtual machines: 1800 Windows 10 Virtual desktops random pooled were configured and tested
  - VMware Horizon Full Clone persistent virtual machines: 1800 Windows 10 Virtual desktops random pooled were configured and tested

For the Cisco Validated Design described in this document, individual configuration of Remote Desktop Server Hosted sessions (RDSH) using RDS-based Server OS machines and Virtual Desktops using Desktop OS machines via Instant-clone, and Full clone persistent desktops were configured and tested. The following sections discuss design decisions relative to the VMware Horizon deployment, including this CVD test environment.

---

## Storage Considerations

### Boot from SAN

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

### NetApp AFF Storage Considerations

**Note:** Make sure each NetApp AFF Controller is connected to BOTH storage fabrics (A/B).

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

### Port Connectivity

10/25/40/100 Gbe connectivity support – while both 10 and 25 Gbe is provided through 2 onboard NICs on each AFF controller, if more interfaces are required or if 40Gbe connectivity is also required, then make sure to provision for additional NICs have been included in the original AFF BOM.

16/32Gb Fiber Channel supports the NetApp Storage up to 32Gb FC support on the latest AFF A400 series arrays. Always make sure the correct number of HBAs and the speed of SFPs are included in the original AFFBOM.

### Overprovision

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

### Topology

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

### VMware Virtual Volumes Considerations

vCenters that are in Enhanced Linked Mode will each be able to communicate with the same AFF, however vCenters that are not in Enhanced Linked Mode must use CA-Signed Certificates using the same AFF. If multiple vCenters need to use the same AFF for vVols, they should be configured in Enhanced Linked Mode.

There are some AFF limits on Volume Connections per Host, Volume Count, and Snapshot Count. For more information about NetApp AFF limits review the following: <https://hww.NetApp.com/Controller/Index>

When a Storage Policy is applied to a vVol VM, the volumes associated with that VM are added to the designated protection group when applying the policy to the VM. If replication is part of the policy, be mindful of the amount of VMs using that storage policy and replication group. A large amount of VMs with a high change

---

rate could cause replication to miss its schedule due to increased replication bandwidth and time needed to complete the scheduled snapshot. NetApp Storage recommends vVol VMs that have Storage Policies applied be balanced between protection groups.



---

## Deployment Hardware and Software

This chapter contains the following:

- [Architecture](#)
- [Products Deployed](#)
- [Physical Topology](#)
- [Logical Architecture](#)
- [Configuration Guidelines](#)

### Architecture

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

The FlexPod Datacenter solution includes Cisco networking, Cisco UCS and NetApp AFF A400, which efficiently fit into a single datacenter rack, including the access layer network switches.

### Products Deployed

This CVD details the deployment of up to 2600 Multi-session OS, 1800 Single-session Windows 10 OS VDI users featuring the following software:

- VMware vSphere ESXi 7 Update 3 Hypervisor
- Microsoft SQL Server 2019 SP1
- Microsoft Windows Server 2019 and Windows 10 64-bit virtual machine Operating Systems
- VMware Horizon 2209 Remote Desktop Server Hosted (RDSH) Sessions provisioned as Instant Clones and stored on the NFS storage
- VMware Horizon 2209 Non-Persistent Win10 Virtual Desktops (VDI) provisioned as Instant Clones and stored on NFS storage
- VMware Horizon 2209 Persistent Win10 Virtual Desktops (VDI) provisioned as Full Clones and stored on NFS storage
- NetApp ONTAP Tools for VMware vSphere 9.10P1
- FSlogix for Profile Management
- VMware vSphere ESXi 7.0.3 Hypervisor
- Microsoft Windows Server 2019 and Windows 10 (build 2004) 64-bit virtual machine Operating Systems  
NetApp ONTAP 9.10.1P1


FlexPod with Cisco UCS M6 servers, VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions and Windows 10 virtual desktops on vSphere 7.0 U3 delivers a Virtual Desktop Infrastructure that is redundant, using the best practices of Cisco and NetApp Storage. The solution includes VMware vSphere 7.0 U3 hypervisor installed on the Cisco UCS M6 X Series server configured for stateless compute design using boot from SAN. NetApp Storage AFF A400 provides the storage infrastructure required for setting up the VDI workload. Cisco

UCS Manager is utilized to configure and manage the UCS infrastructure with Cisco Intersight providing lifecycle management capabilities. The solution requirements and design details are covered in this section.

## Physical Topology

FlexPod VDI with Cisco UCS M6 servers is a Fibre Channel (FC) based storage access design. NetApp Storage AFF and Cisco UCS are connected through Cisco MDS 9132T switches and storage access utilizes the FC network. For VDI IP based file share storage access NetApp Storage AFF Cisco UCS are connected through Cisco Nexus 93180YC-FX switches. The physical connectivity details are explained below.

**Figure 37. FlexPod VDI – Physical Topology**



**Figure 37** details the physical hardware and cabling deployed to enable this solution:

- 2 Cisco Nexus 93180YC-FX Switches in NX-OS Mode.
- 2 Cisco MDS 9132T 32-Gb Fibre Channel Switches.
- One Cisco UCS X9508 chassis with two built-in UCS 9108-25G IO Modules
- 8 Cisco UCS X210C M6 Blade Servers with Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM, and one Cisco VIC14425 mezzanine card, providing N+1 server fault tolerance.
- NetApp AFF A400 Storage System with dual redundant controllers, 2x disk shelves, and 48 x 1.75 TB solid-state NVMe drives providing storage and NVME/FC/NFS/CIFS connectivity.

**Note:** The common services and LoginVSI Test infrastructure are not a part of the physical topology of this solution.

Table 5 lists the software versions of the primary products installed in the environment.

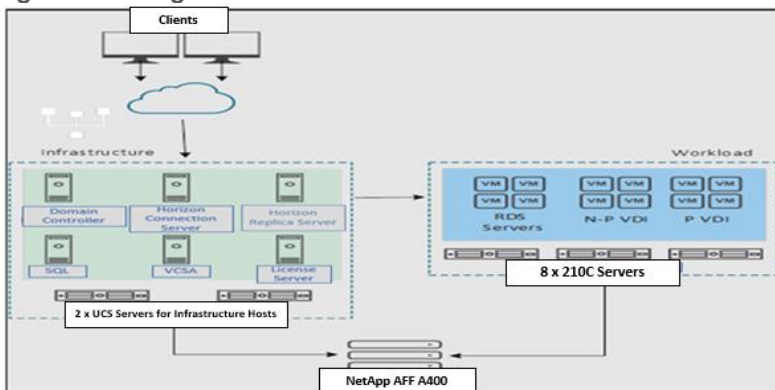
**Table 5. Software and Firmware Versions**

Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	5.0(1b) bundle release
Cisco	UCS Manager	5.0(1b) bundle release
Cisco	UCS X-Series Blades	5.0(1b) bundle release
Cisco	VIC 14425	5.0(1b) bundle release
Cisco	Cisco Nexus 93180YC-FX	9.3(7a)
Cisco	Cisco MDS 9132T	8.5(1a)
NetApp	AFF A400	ONTAP 9.10.1P1
NetApp	ONTAP Tools for VMWare vSphere	9.10
NetApp	NetApp NFS Plug-in for VMWare VAAI	2.0
NetApp	Active IQ Unified Manager	9.10P1
NetApp	SnapCenter Plug-In for VMWare vSphere	4.6
VMware	vCenter Server Appliance	7.0.3.20150588
VMware	vSphere 7.0U3	7.0.3.00700
VMware	Horizon Connection Server	8.7.0.20649599
VMware	Horizon Agent	8.7.0.20606795
VMware	Tools	11.3.5.18557794

## Logical Architecture

The logical architecture of the validated solution which is designed to support up to 2600 users on a single chassis containing eight Cisco UCS X-Series blade servers, with physical redundancy for the blade servers for each workload type is illustrated in Figure 38.

**Figure 38. Logical Architecture Overview**



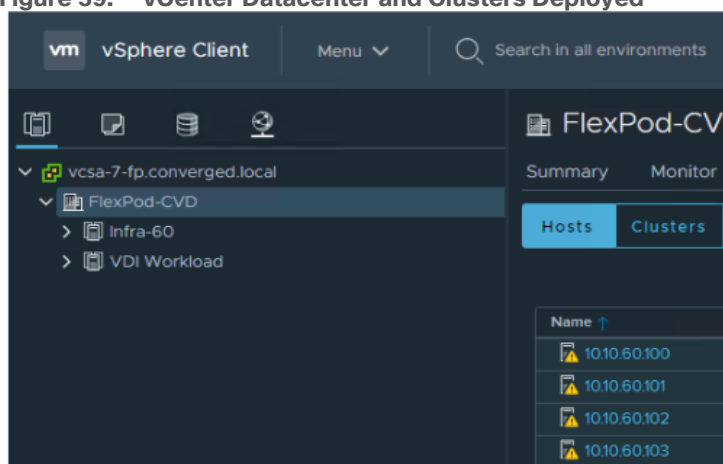
## VMware Clusters

Two VMware Clusters in one vCenter datacenter were utilized to support the solution and testing environment:

- VDI Cluster FlexPod Datacenter with Cisco UCS
  - Infrastructure: Infra VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, VMware & Desktops Controllers, Provisioning Servers, and NetApp ONTAP Tools for VMware vSphere, ActiveIQ Unified Manager, VSMS, and so on)
  - VDI Workload VMs (Windows Server 2019 provisioned with VMware Horizon, Windows 10 provisioned with VMware Horizon Windows 10 Instant Clones and Persistent desktops)
- VSI Launchers and Launcher Cluster

For Example, the cluster(s) configured for running LoginVSI workload for measuring VDI End User Experience is LVS-Launcher-CLSTR: (The Login VSI infrastructure cluster consists of Login VSI data shares, LVSI Web Servers and LVSI Management Control VMs etc. were connected using the same set of switches and vCenter instance but was hosted on separate storage. LVS-Launcher-CLSTR configured and used for the purpose of testing LoginVSI End User Experience for VDI multi session users and VDI Win 10 users.

**Figure 39. vCenter Datacenter and Clusters Deployed**



## Configuration Guidelines

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

This document is intended to allow the reader to configure the VMware Horizon 2209 customer environment as a stand-alone solution.

## VLANs

The VLAN configuration recommended for the environment includes a total of six VLANs as listed in [Table 6](#).

**Table 6. VLANs Configured in this Study**

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN

VLAN Name	VLAN ID	VLAN Purpose
In-Band-Mgmt	60	In-Band management interfaces
Infra-Mgmt	61	Infrastructure Virtual Machines
NFS- VLAN	62	VLAN for Infrastructure NFS traffic
CIFS-VLAN	63	CIFS Storage access
VCC/VM-Network	64	RDSH, VDI Persistent and Non-Persistent
vMotion	66	VMware vMotion
OOB-Mgmt	173	Out-of-Band management interfaces

## VSANs

Two virtual SANs configured for communications and fault tolerance in this design as outlined in [Table 7](#).

**Table 7. VSANs Configured in this Study**

VSAN Name	VSAN ID	VSAN Purpose
VSAN 400	400	VSAN for Primary SAN communication
VSAN 401	401	VSAN for Secondary SAN communication

---

## Solution Configuration

This chapter contains the following:

- [Solution Cabling](#)
- [Network Switch Configuration](#)
- [FlexPod Cisco Nexus Switch Configuration](#)

### Solution Cabling

The following sections detail the physical connectivity configuration of the FlexPod VMware VDI environment.

The information provided in this section is a reference for cabling the physical equipment in this Cisco Validated Design environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain the details for the prescribed and supported configuration of the NetApp Storage AFF A400 storage array to the Cisco 6454 Fabric Interconnects through Cisco MDS 9132T 32-Gb FC switches.

**Note:** This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

**IMPORTANT! Be sure to follow the cabling directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.**

**Note:** Be sure to use the cabling directions in this section as a guide.

[Figure 40](#) details the cable connections used in the validation lab for FlexPod topology based on the Cisco UCS 6454 fabric interconnect. Four 32Gb uplinks connect as port-channels to each Cisco UCS Fabric Interconnect from the MDS switches, and a total of eight 32Gb links connect the MDS switches to the NetApp AFF A400 controllers, four of these have been used for scsi-fc and the other four to support nvme-fc. Also, 40Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the NetApp AFF A400 controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each All Flash Array controller has a connection to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

The architecture is divided into three distinct layers:

1. Cisco UCS Compute Platform
2. Network Access layer and LAN
3. Storage Access to the NetApp AFF400

Figure 40. FlexPod Solution Cabling Diagram

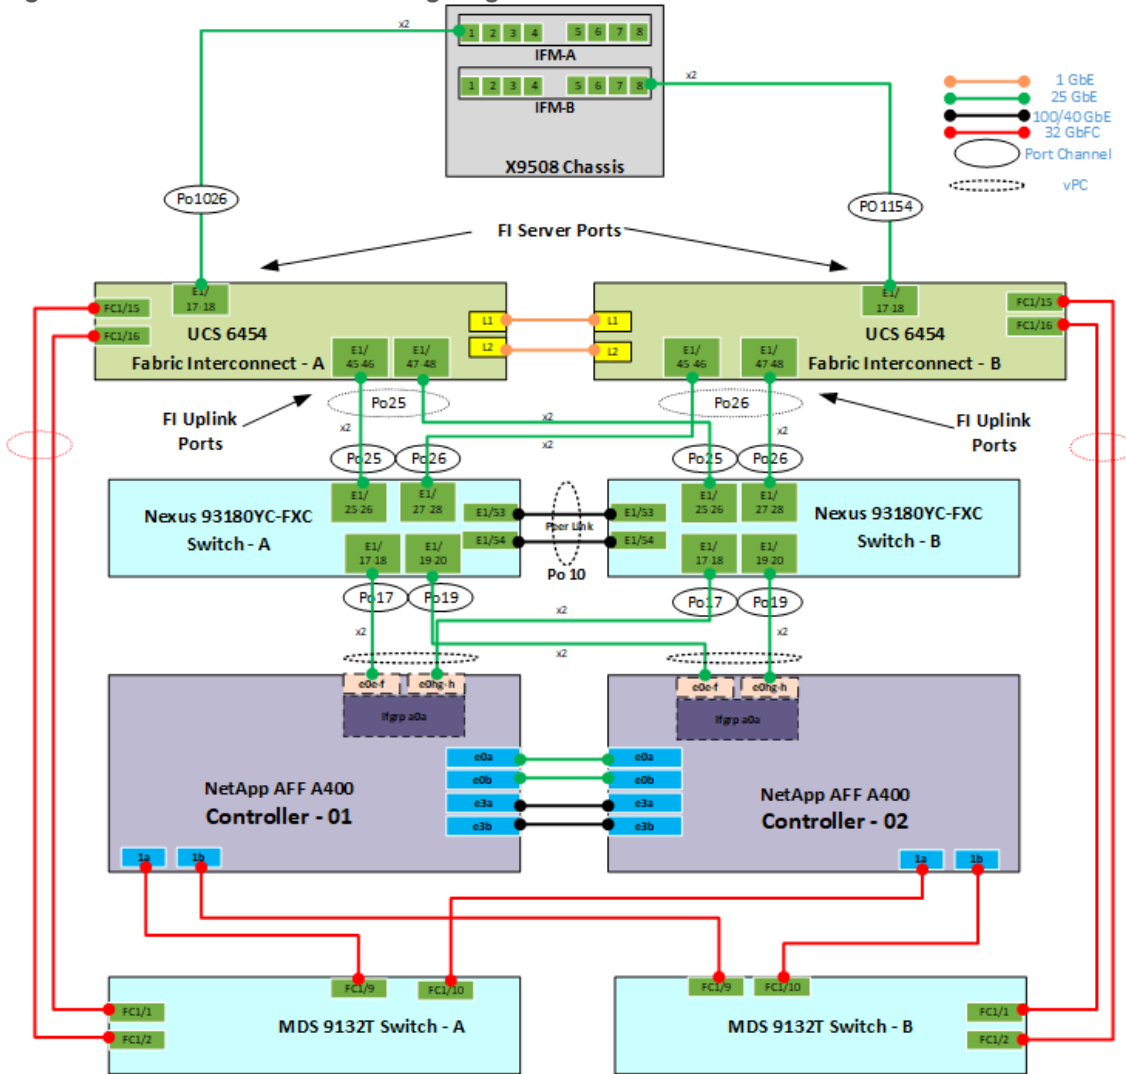


Table 8. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/17	25GbE	NetApp Controller 1	e0e
	Eth1/18	25GbE	NetApp Controller 1	e0h
	Eth1/19	25GbE	NetApp Controller 2	e0e
	Eth1/20	25GbE	NetApp Controller 2	e0h
	Eth1/25	25GbE	Cisco UCS fabric interconnect A	Eth2/1
	Eth1/26	25GbE	Cisco UCS fabric interconnect A	Eth2/2
	Eth1/27	25GbE	Cisco UCS fabric interconnect B	Eth2/3
	Eth1/28	25GbE	Cisco UCS fabric interconnect B	Eth2/4
	Eth1/53	40GbE	Cisco Nexus 93180YC-FX B	Eth1/53

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/54	40GbE	Cisco Nexus 93180YC-FX B	Eth1/54
	MGMT0	GbE	GbE management switch	Any

**Note:** For devices requiring GbE connectivity, use the GbE Copper SFP+s (GLC-T=).

**Table 9. Cisco Nexus 93180YC-FX-B Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/17	25GbE	NetApp Controller 1	e0e
	Eth1/18	25GbE	NetApp Controller 1	e0f
	Eth1/19	25GbE	NetApp Controller 1	e0f
	Eth1/20	25GbE	NetApp Controller 2	e0g
	Eth1/25	25GbE	Cisco UCS fabric interconnect A	eth1/1
	Eth1/26	25GbE	Cisco UCS fabric interconnect A	eth1/2
	Eth1/27	25GbE	Cisco UCS fabric interconnect B	eth1/3
	Eth1/28	25GbE	Cisco UCS fabric interconnect B	eth1/4
	Eth1/53	40GbE	Cisco Nexus 93180YC-FX A	Eth1/53
	Eth1/54	40GbE	Cisco Nexus 93180YC-FX A	Eth1/54
	MGMT0	GbE	GbE management switch	Any

**Table 10. NetApp Controller-1 Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp AFF400 Node 1	e0M	1GbE	1GbE management switch	Any
	e0s	GbE	GbE management switch	Any
	e0a	25GbE	NetApp Controller 2	e0a
	e0b	25GbE	NetApp Controller 2	e0b
	e0c	100GbE	NS224-1	e0a
	e0d	100GbE	NS224-2	e0b
	e0e	25GbE	Cisco Nexus 93180YC-FX B	Eth1/17
	e0f	25GbE	Cisco Nexus 93180YC-FX B	Eth1/18
	e0g	25GbE	Cisco Nexus 93180YC-FX A	Eth1/18



Local Device	Local Port	Connection	Remote Device	Remote Port
	e0h	25GbE	Cisco Nexus 93180YC-FX A	Eth1/17
	e3a	100GbE	NetApp Controller 2	e3a
	e3b	100GbE	NetApp Controller 2	e3b
	e1a	100GbE	NS224-2	e0a
	e1b	100GbE	NS224-1	e0b

**Note:** When the term e0M is used, the physical Ethernet port to which the table is referring is the port indicated by a wrench icon on the rear of the chassis.

**Table 11. NetApp Controller 2 Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp AFF400 Node 2	e0M	100E	100MbE management switch	Any
	e0s	GbE	GbE management switch	Any
	e0a	25GbE	NetApp Controller 1	e0a
	e0b	25GbE	NetApp Controller 1	e0b
	e0c	100GbE	NS224-1	e0a
	e0d	100GbE	NS224-2	e0b
	e0e	40GbE	Cisco Nexus 93180YC-FX A	Eth1/19
	e0f	40GbE	Cisco Nexus 93180YC-FX B	Eth1/19
	e0g	40GbE	Cisco Nexus 93180YC-FX B	Eth1/20
	e0h	40GbE	Cisco Nexus 93180YC-FX A	Eth1/20
	e3a	100GbE	NetApp Controller 1	e3a
	e3b	100GbE	NetApp Controller 1	e3b
	e1a	100GbE	NS224-2	e0a
	e1b	100GbE	NS224-1	e0b

**Table 12. Cisco UCS Fabric Interconnect A Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect	eth2/1	25GbE	Cisco Nexus 93180YC-FX A	Eth1/45
	eth2/2	25GbE	Cisco Nexus 93180YC-FX A	Eth1/46

Local Device	Local Port	Connection	Remote Device	Remote Port
A	eth2/3	25GbE	Cisco Nexus 93180YC-FX B	Eth1/47
	eth2/4	25GbE	Cisco Nexus 93180YC-FX B	Eth1/48
	Eth1/1	25GbE	Cisco UCS Chassis1 FEX A	IOM 1/1
	Eth1/2	25GbE	Cisco UCS Chassis1 FEX A	IOM 1/2
	Eth1/3	25GbE	Cisco UCS Chassis1 FEX A	IOM 1/3
	Eth1/4	25GbE	Cisco UCS Chassis1 FEX A	IOM 1/4
	Eth1/5	25GbE	Cisco UCS Chassis2 FEX A	IOM 1/1
	Eth1/6	25GbE	Cisco UCS Chassis2 FEX A	IOM 1/2
	Eth1/7	25GbE	Cisco UCS Chassis2 FEX A	IOM 1/3
	Eth1/8	25GbE	Cisco UCS Chassis2 FEX A	IOM 1/4
	Eth1/9	25GbE	Cisco UCS Chassis3 FEX A	IOM 1/1
	Eth1/10	25GbE	Cisco UCS Chassis3 FEX A	IOM 1/2
	Eth1/11	25GbE	Cisco UCS Chassis3 FEX A	IOM 1/3
	Eth1/12	25GbE	Cisco UCS Chassis3 FEX A	IOM 1/4
	Eth1/13	25GbE	Cisco UCS Chassis4 FEX A	IOM 1/1
	Eth1/14	25GbE	Cisco UCS Chassis4 FEX A	IOM 1/2
	Eth1/15	25GbE	Cisco UCS Chassis4 FEX A	IOM 1/3
	Eth1/16	25GbE	Cisco UCS Chassis4 FEX A	IOM 1/4
	FC1/1	32GbE	Cisco MDS 9132T A	IOM 1/15
	FC1/2	32GbE	Cisco MDS 9132T A	IOM 1/16
	MGMT0	GbE	GbE Management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

**Table 13. Cisco UCS Fabric Interconnect B Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth2/1	25GbE	Cisco Nexus 93108 A	Eth1/45
	Eth2/2	25GbE	Cisco Nexus 93108 A	Eth1/46
	Eth2/3	25GbE	Cisco Nexus 93108 B	Eth1/47

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth2/4	25GbE	Cisco Nexus 93108 B	Eth1/48
	Eth1/1	25GbE	Cisco UCS Chassis1 FEX B	IOM 2/1
	Eth1/2	25GbE	Cisco UCS Chassis1 FEX B	IOM 2/2
	Eth1/3	25GbE	Cisco UCS Chassis1 FEX B	IOM 2/3
	Eth1/4	25GbE	Cisco UCS Chassis1 FEX B	IOM 2/4
	Eth1/5	25GbE	Cisco UCS Chassis2 FEX B	IOM 2/1
	Eth1/6	25GbE	Cisco UCS Chassis2 FEX B	IOM 2/2
	Eth1/7	25GbE	Cisco UCS Chassis2 FEX B	IOM 2/3
	Eth1/8	25GbE	Cisco UCS Chassis2 FEX B	IOM 2/4
	Eth1/9	25GbE	Cisco UCS Chassis3 FEX B	IOM 2/1
	Eth1/10	25GbE	Cisco UCS Chassis3 FEX B	IOM 2/2
	Eth1/11	25GbE	Cisco UCS Chassis3 FEX B	IOM 2/3
	Eth1/12	25GbE	Cisco UCS Chassis3 FEX B	IOM 2/4
	Eth1/13	25GbE	Cisco UCS Chassis4 FEX B	IOM 2/1
	Eth1/14	25GbE	Cisco UCS Chassis4 FEX B	IOM 2/2
	Eth1/15	25GbE	Cisco UCS Chassis4 FEX B	IOM 2/3
	Eth1/16	25GbE	Cisco UCS Chassis4 FEX B	IOM 2/4
	FC1/1	32GbE	Cisco MDS 9132T B	IOM 1/15
	FC1/2	32GbE	Cisco MDS 9132T B	IOM 1/16
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect A	L1
	L2	GbE	Cisco UCS fabric interconnect A	L2

## Network Switch Configuration

This subject contains the following procedures:

- [Set Up Initial Configuration on Cisco Nexus A](#)
- [Set Up Initial Configuration on Cisco Nexus B](#)

This section provides a detailed procedure for configuring the Cisco Nexus 93180YC-FX switches for use in a FlexPod environment. The Cisco Nexus 93180YC-FX will be used LAN switching in this solution.

**IMPORTANT!** Follow these steps precisely because failure to do so could result in an improper configuration.

## Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in section [Solution Cabling](#).

### FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 9.3(7a), the Cisco suggested Nexus switch release at the time of this validation.

The following procedure includes the setup of NTP distribution on both the mgmt0 port and the in-band management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes that the default VRF is used to route the in-band management VLAN.

**Note:** In this validation, port speed and duplex are hard set at both ends of every 100GE connection.

### Procedure 1. Set Up Initial Configuration on Cisco Nexus A

Set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>

**Step 1.** On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass
password and basic configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
```

```
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: yes
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

**Step 2. Review the configuration summary before enabling the configuration.**

```
Use this configuration and save it? (yes/no) [y]: Enter
```

## Procedure 2. Set Up Initial Configuration on Cisco Nexus B

Set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>.

**Step 1. On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.**

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass
password and basic configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: yes
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

**Step 2. Review the configuration summary before enabling the configuration.**

```
Use this configuration and save it? (yes/no) [y]: Enter
```

## FlexPod Cisco Nexus Switch Configuration

This subject contains the following procedures:

- [Enable Features on Cisco Nexus A and Cisco Nexus B](#)
- [Set Global Configurations on Cisco Nexus A and Cisco Nexus B](#)
- [Create VLANs on Cisco Nexus A and Cisco Nexus B](#)
- [Add NTP Distribution Interface on Cisco Nexus A](#)
- [Add NTP Distribution Interface on Cisco Nexus B](#)
- [Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus A](#)
- [Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus B](#)
- [Create Port Channels on Cisco Nexus A and Cisco Nexus B](#)
- [Configure Port Channel Parameters on Cisco Nexus A and Cisco Nexus B](#)
- [Configure Virtual Port Channels on Cisco Nexus A](#)
- [Configure Virtual Port Channels on Cisco Nexus B](#)

### Procedure 1. Enable Features on Cisco Nexus A and Cisco Nexus B

SAN switching requires both the SAN\_ENTERPRISE\_PKG and FC\_PORT\_ACTIVATION\_PKG licenses. Please ensure these licenses are installed on each Cisco Nexus 93180YC-FX switch.

**Step 1.** Log in as admin.

**Step 2.** Since basic FC configurations were entered in the setup script, feature-set fcoe has been automatically installed and enabled. Run the following commands:

```
config t
feature udld
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

### Procedure 2. Set Global Configurations on Cisco Nexus A and Cisco Nexus B

**Step 1.** Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day> <end-month> <end-time> <offset-minutes>
```

```
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

#### Tech tip

It is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summer time, please see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3\(x\)](#). Sample clock commands for the United States Eastern timezone are:

```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60.
```

### Procedure 3. Create VLANs on Cisco Nexus A and Cisco Nexus B

**Step 1.** From the global configuration mode, run the following commands:

```
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
vlan <infra-nfs-vlan-id>
name Infra-NFS-VLAN
vlan <infra-CIFS-vlan-id>
name Infra-CIFS-VLAN
exit
```

### Procedure 4. Add NTP Distribution Interface on Cisco Nexus A

**Step 1.** From the global configuration mode, run the following commands:

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-b-ntp-ip> use-vrf default
```

### Procedure 5. Add NTP Distribution Interface on Cisco Nexus B

**Step 1.** From the global configuration mode, run the following commands:

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-a-ntp-ip> use-vrf default
```

### Procedure 6. Add Port Profiles on Cisco Nexus A and Cisco Nexus B

This version of the FlexPod solution uses port profiles for virtual port channel (vPC) connections to NetApp Storage, Cisco UCS, and the vPC peer link.

**Step 1.** From the global configuration mode, run the following commands:

```
port-profile type port-channel FP-ONTAP-Storage
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <infra-CIFS-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled

port-profile type port-channel FP-UCS
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <infra-CIFS-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled

port-profile type port-channel vPC-Peer-Link
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <infra-CIFS-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>
spanning-tree port type network
speed 100000
duplex full
state enabled
```

### Procedure 7. Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus A

**Note:** In this step and in the following sections, configure the AFF nodename <st-node> and Cisco UCS 6454 fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

**Step 1.** From the global configuration mode, run the following commands:

```
interface Eth1/21
description <ucs-clustername>-a:1/45
udld enable
interface Eth1/22
description <ucs-clustername>-a:1/46
udld enable
interface Eth1/23
description <ucs-clustername>-b:1/45
udld enable
interface Eth1/24
```



```
description <ucs-clustername>-b:1/46
udld enable
```

**Step 2.** For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected. If you have fibre optic connections, do not enter the `udld enable` command.

```
interface Eth1/17
description <st-clustername>-01:e0e
interface Eth1/18
description <st-clustername>-01:e0f
interface Eth1/19
description <st-clustername>-02:e0e
interface Eth1/20
description <st-clustername>-02:e0f
interface Eth1/53
description <nexus-b-hostname>:1/53
interface Eth1/54
description <nexus-b-hostname>:1/54
exit
```

### Procedure 8. Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus B

**Step 1.** From the global configuration mode, run the following commands:

```
interface Eth1/21
description <ucs-clustername>-a:1/47
udld enable
interface Eth1/22
description <ucs-clustername>-a:1/48
udld enable
interface Eth1/23
description <ucs-clustername>-b:1/47
udld enable
interface Eth1/24
description <ucs-clustername>-b:1/48
udld enable
```

**Step 2.** For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected.

```
interface Eth1/17
description <st-clustername>-01:e0g
interface Eth1/18
description <st-clustername>-01:e0h
interface Eth1/19
description <st-clustername>-02:e0g
interface Eth1/20
```

```
description <st-clustername>-02:e0h
interface Eth1/53
description <nexus-a-hostname>:1/53
interface Eth1/54
description <nexus-a-hostname>:1/54
exit
```

## Procedure 9. Create Port Channels on Cisco Nexus A and Cisco Nexus B

**Step 1.** From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
interface Eth1/49-50
channel-group 10 mode active
no shutdown
interface Po117
description <st-clustername>-01
interface Eth1/17-18
channel-group 117 mode active
no shutdown
interface Po119
description <st-clustername>-02
interface Eth1/19-20
channel-group 119 mode active
no shutdown
interface Po141
description <ucs-clustername>-a
interface Eth1/21-22
channel-group 121 mode active
no shutdown
interface Po123
description <ucs-clustername>-b
interface Eth1/23-24
channel-group 123 mode active
no shutdown
exit
copy run start
```

## Procedure 10. Configure Port Channel Parameters on Cisco Nexus A and Cisco Nexus B

**Step 1.** From the global configuration mode, run the following commands:

```
interface Po10
inherit port-profile vPC-Peer-Link
interface Po117
```

```
inherit port-profile FP-ONTAP-Storage
interface Po119

inherit port-profile FP-ONTAP-Storage
interface Po121

inherit port-profile FP-UCS
interface Po123

inherit port-profile FP-UCS
exit

copy run start
```

## Procedure 11. Configure Virtual Port Channels on Cisco Nexus A

**Step 1.** From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po121
vpc 121
interface Po123
vpc 123
exit

copy run start
```

## Procedure 12. Configure Virtual Port Channels on Cisco Nexus B

**Step 1.** From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
```

```
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po121
vpc 121
interface Po123
vpc 123
exit
copy run start
```

## Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after the configuration is completed.

## Switch Testing Commands

The following commands can be used to check for correct switch configuration:

**Note:** Some of these commands need to run after further configuration of the FlexPod components are complete to see complete results.

```
show run
show vpc
show port-channel summary
show ntp peer-status
show cdp neighbors
show lldp neighbors
show run int
show int
show udld neighbors
show int status
```

---

## Storage Configuration

This chapter contains the following:

- [NetApp Hardware Universe](#)
- [NetApp ONTAP 9.10.1P1](#)

### NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It also provides configuration information for all the NetApp storage appliances currently supported by ONTAP software and a table of component compatibilities.

To confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install, follow these steps at the [NetApp Support](#) site.

1. Access the [HWU application](#) to view the System Configuration guides. Click the Platforms menu to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

### Controllers

Follow the physical installation procedures for the controllers found here: <https://docs.netapp.com/us-en/ontap-systems/index.html>.

### Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported by the AFF A400 and AFF A800 is available at the [NetApp Support](#) site.

When using SAS disk shelves with NetApp storage controllers, refer to: <https://docs.netapp.com/us-en/ontap-systems/sas3/index.html> for proper cabling guidelines.

When using NVMe drive shelves with NetApp storage controllers, refer to: <https://docs.netapp.com/us-en/ontap-systems/ns224/index.html> for installation and servicing guidelines.

### NetApp ONTAP 9.10.1P1

This subject contains the following procedures:

- [Configure Node 01](#)
- [Configure Node 02](#)
- [Set Up Node](#)
- [Log into the Cluster](#)
- [Verify Storage Failover](#)
- [Set Auto-Revert on Cluster Management](#)
- [Zero All Spare Disks](#)
- [Set Up Service Processor Network Interface](#)
- [Create Manual Provisioned Aggregates \(Optional\)](#)

- 
- [Remove Default Broadcast Domains](#)
  - [Disable Flow Control on 25/100GbE Data Ports](#)
  - [Disable Auto-Negotiate on Fibre Channel Ports \(Required only for FC configuration\)](#)
  - [Enable Cisco Discovery Protocol](#)
  - [Enable Link-layer Discovery Protocol on all Ethernet Ports](#)
  - [Create Management Broadcast Domain](#)
  - [Create NFS Broadcast Domain](#)
  - [Create CIFS Broadcast Domain](#)
  - [Create iSCSI Broadcast Domains \(Required only for iSCSI configuration\)](#)
  - [Create Interface Groups](#)
  - [Change MTU on Interface Groups](#)
  - [Create VLANs](#)
  - [Configure Time Synchronization on the Cluster](#)
  - [Configure Simple Network Management Protocol \(SNMP\)](#)
  - [Configure SNMPv3 Access](#)
  - [Create an infrastructure SVM](#)
  - [Configure CIFS Servers](#)
  - [Modify Storage Virtual Machine Option](#)
  - [Create Load-Sharing Mirrors of a SVM Root Volume](#)
  - [Create FC Block Protocol Service \(required only for FC configuration\)](#)
  - [Create iSCSI Block Protocol Service \(required only for iSCSI configuration\)](#)
  - [Vserver Protocol Verification](#)
  - [Configure HTTPS Access to the Storage Controller](#)
  - [Configure NFSv3 and NFSv4.1](#)
  - [Create CIFS Export Policy](#)
  - [Create a NetApp FlexVol Volume](#)
  - [Create a NetApp FlexGroup Volume](#)
  - [Modify Volume Efficiency](#)
  - [Create CIFS Shares](#)
  - [Create NFS LIFs](#)
  - [Create CIFS LIFs](#)
  - [Create FC LIFs \(required only for FC configuration\)](#)
  - [Create iSCSI LIFs \(required only for iSCSI configuration\)](#)

- [Configure FC-NVMe Datastore for vSphere 7U2 on existing SVM \(Infra-FC\) for FC-NVMe configuration only](#)
- [Add Infrastructure SVM Administrator and SVM Administration LIF to In-band Management Network](#)
- [Configure and Test AutoSupport](#)

## Complete Configuration Worksheet

Before running the setup script, complete the [Cluster setup worksheet](#) in the ONTAP 9 Documentation Center. You must have access to the [NetApp Support](#) site to open the cluster setup worksheet.

## Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [Software setup](#) section of the [ONTAP 9 Documentation Center](#) to learn about configuring ONTAP. [Table 14](#) lists the information needed to configure two ONTAP nodes. Customize the cluster-detail values with the information applicable to your deployment.

**Table 14. ONTAP Software Installation Prerequisites**

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
ONTAP 9.10.1P1 URL (http server hosting ONTAP software)	<url-boot-software>

### Procedure 1. Configure Node 01

**Step 1.** Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when the following message displays: Starting AUTOBOOT press Ctrl-C to abort...

**Step 2.** Allow the system to boot up.

```
autoboot
```

**Step 3.** Press Ctrl-C when prompted.

**Note:** If ONTAP 9.10.1P1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.10.1P1 is the version being booted, select option 8 and `y` to reboot the node. Then continue with section [Set Up Node](#).

**Step 4.** To install new software, select option 7 from the menu.

**Step 5.** Enter `y` to continue the installation.

**Step 6.** Select `e0M` for the network port for the download.

**Step 7.** Enter `n` to skip the reboot.

**Step 8.** Select option 7 from the menu: `Install new software first`

**Step 9.** Enter `y` to continue the installation.

**Step 10.** Enter the IP address, netmask, and default gateway for `e0M`.

**Step 11.** Enter the IP address for port `e0M`: `<node01-mgmt-ip>`

**Step 12.** Enter the netmask for port `e0M`: `<node01-mgmt-mask>`

**Step 13.** Enter the IP address of the default gateway: `<node01-mgmt-gateway>`

**Step 14.** Enter the URL where the software can be found.

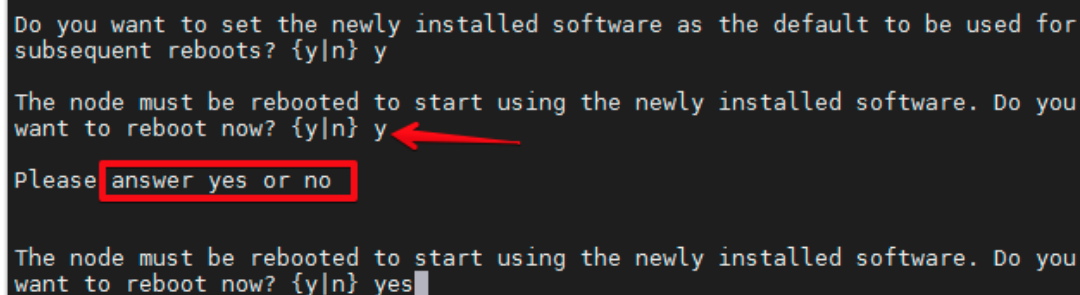
**Step 15.** The `e0M` interface should be connected to management network and the web server must be reachable (using ping) from node 01.

```
<url-boot-software>
```

**Step 16.** Press Enter for the user name, indicating no user name.

**Step 17.** Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

**Step 18.** Enter `yes` to reboot the node.



```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y

Please answer yes or no

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} yes
```

**Note:** When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

**Note:** During the ONTAP installation, a prompt to reboot the node requests a Y/N response. The prompt requires the entire Yes or No response to reboot the node and continue the installation.

**Step 19.** Press `Ctrl-C` when the following message displays:

```
Press Ctrl-C for Boot Menu
```

**Step 20.** Select option 4 for Clean Configuration and Initialize All Disks.

**Step 21.** Enter `y` to zero disks, reset config, and install a new file system.

**Step 22.** Enter `yes` to erase all the data on the disks.

**Note:** The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the configuration of node 02 while the disks for node 01 are zeroing.

## Procedure 2. Configure Node 02

**Step 1.** Connect to the storage system console port. You should see a Loader-B prompt. However, if the storage system is in a reboot loop, press `Ctrl-C` to exit the autoboot loop when the following message displays: Starting AUTOBOOT press `Ctrl-C` to abort...



**Step 2.** Allow the system to boot up.

```
autoboot
```

**Step 3.** Press Ctrl-C when prompted.

**Note:** If ONTAP 9.10.1P1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.10.1P1 is the version being booted, select option 8 and `y` to reboot the node, then continue with section [Set Up Node](#).

**Step 4.** To install new software, select option 7.

**Step 5.** Enter `y` to continue the installation.

**Step 6.** Select e0M for the network port you want to use for the download.

**Step 7.** Enter `n` to skip the reboot.

**Step 8.** Select option 7: Install new software first

**Step 9.** Enter `y` to continue the installation

**Step 10.** Enter the IP address, netmask, and default gateway for e0M.

**Step 11.** Enter the IP address for port e0M: <node02-mgmt-ip>

**Step 12.** Enter the netmask for port e0M: <node02-mgmt-mask>

**Step 13.** Enter the IP address of the default gateway: <node02-mgmt-gateway>

**Step 14.** Enter the URL where the software can be found.

**Step 15.** The web server must be reachable (ping) from node 02.

```
<url-boot-software>
```

**Step 16.** Press `Enter` for the username, indicating no user name.

**Step 17.** Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

**Step 18.** Enter `yes` to reboot the node.

```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} y
```

```
The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y
```

```
Please answer yes or no
```

```
The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} yes
```

**Note:** When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

**Note:** During the ONTAP installation a prompt to reboot the node requests a Y/N response. The prompt requires the entire `Yes` or `No` response to reboot the node and continue the installation.

**Step 19.** Press Ctrl-C when you see this message: Press Ctrl-C for Boot Menu.

**Step 20.** Select option 4 for Clean Configuration and Initialize All Disks.

**Step 21.** Enter `y` to zero disks, reset config, and install a new file system.

**Step 22.** Enter `yes` to erase all the data on the disks.

**Note:** The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

### Procedure 3. Set Up Node

**Step 1.** From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.10.1P1 boots on the node for the first time.

**Step 2.** Follow the prompts to set up node 01.

**Step 3.** Welcome to node setup.

- You can enter the following commands at any time:
  - "help" or "?" - if you want to have a question clarified,
  - "back" - if you want to change previously answered questions, and
  - "exit" or "quit" - if you want to quit the setup wizard.

#### Tech tip

Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup."

To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical Support.

To disable this feature, enter "autosupport modify -support disable" within 24 hours.

Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.

For further information on AutoSupport, see: <http://support.netapp.com/autosupport/>

**Step 4.** Type `yes` to confirm and continue {yes}: `yes`

**Step 5.** Enter the node management interface port [e0M]: `Enter`

**Step 6.** Enter the node management interface IP address: `<node01-mgmt-ip>`

**Step 7.** Enter the node management interface netmask: `<node01-mgmt-mask>`

**Step 8.** Enter the node management interface default gateway: `<node01-mgmt-gateway>`

**Step 9.** A node management interface on port e0M with IP address `<node01-mgmt-ip>` has been created

**Step 10.** Use your web browser to complete cluster setup by accessing <https://<node01-mgmt-ip>>. Otherwise press `Enter` to complete cluster setup using the command line interface.

**Step 11.** To complete cluster setup, open a web browser and navigate to <https://<node01-mgmt-ip>>.

**Table 15. Cluster Create in ONTAP Prerequisites**

Cluster Detail	Cluster Detail Value
Cluster name	<code>&lt;clustername&gt;</code>
Cluster Admin SVM	<code>&lt;cluster-adm-svm&gt;</code>
Infrastructure Data SVM	<code>&lt;infra-data-svm&gt;</code>

Cluster Detail	Cluster Detail Value
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-sp-ip>
Node 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>
NTP server B IP address	<switch-b-ntp-ip>
SNMPv3 User	<snmp-v3-usr>
SNMPv3 Authentication Protocol	<snmp-v3-auth-proto>
SNMPv3 Privacy Protocol	<snmpv3-priv-proto>

**Note:** Cluster setup can also be performed using the CLI. This document describes the cluster setup using the NetApp ONTAP System Manager guided setup.

**Step 12.** Complete the required information on the Initialize Storage System screen:

**Step 13.** In the Cluster screen, enter the cluster name and administrator password.

**Step 14.** Complete the Networking information for the cluster and each node.

**Tech tip**

The nodes should be discovered automatically; if they are not, Refresh the browser page. By default, the cluster interfaces are created on all the new factory shipping storage controllers.

If all the nodes are not discovered, then configure the cluster using the command line.

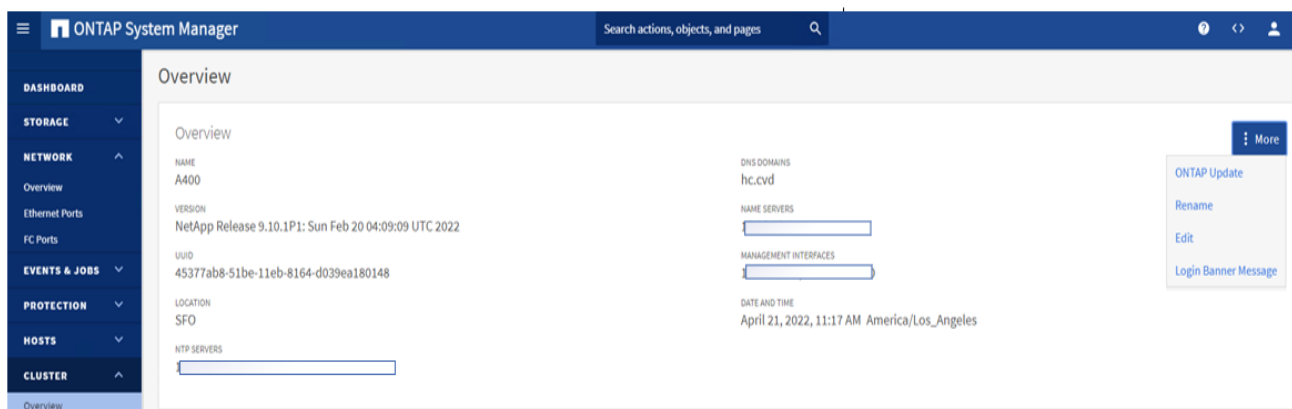
The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

**Step 15.** Click Submit.

**Step 16.** A few minutes will pass while the cluster is configured. When prompted, login to ONTAP System Manager to continue the cluster configuration.

**Step 17.** From the Dashboard click the Cluster menu and click Overview.

**Step 18.** Click the More ellipsis button in the Overview pane and click Edit.

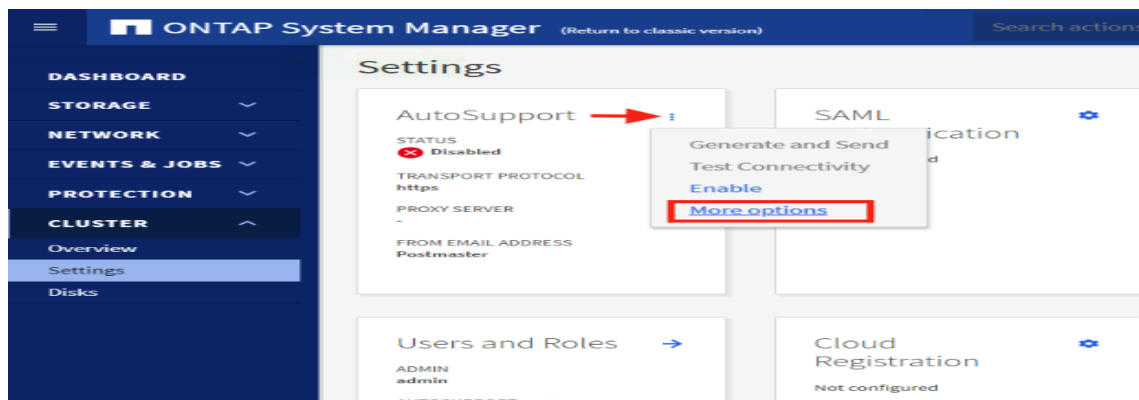


**Step 19.** Add additional cluster configuration details and click Save to make the changes persistent:

- Cluster location
- DNS domain name
- DNS server IP addresses
- DNS server IP addresses can be added individually or with a comma separated list on a single line.

**Step 20.** Click Save to make the changes persistent.

**Step 21.** Select the Settings menu under the Cluster menu.



**Step 22.** If AutoSupport was not configured during the initial setup, click the ellipsis in the AutoSupport tile and select More options.

**Step 23.** To enable AutoSupport click the slider.

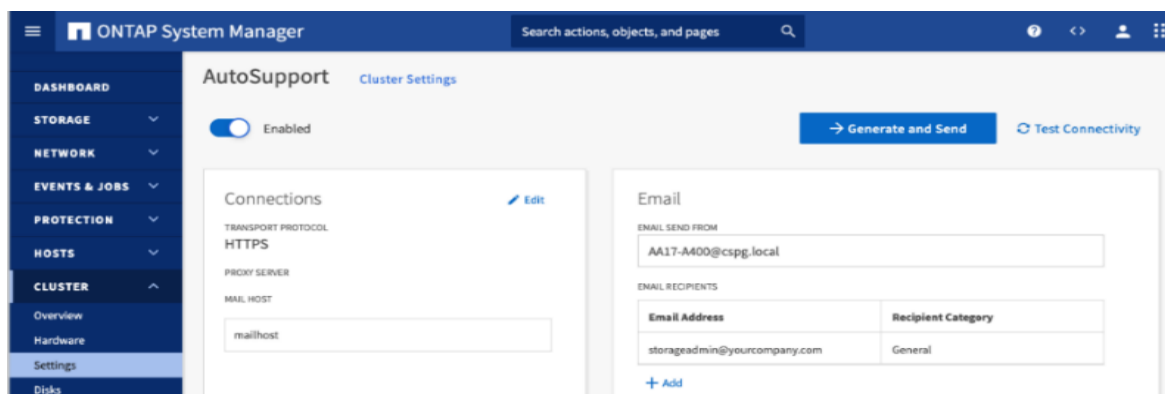
**Step 24.** Click Edit to change the transport protocol, add a proxy server address and a mail host as needed.

**Step 25.** Click Save to enable the changes.

**Step 26.** In the Email tile to the right, click Edit and enter the desired email information:

- Email send from address
- Email recipient addresses
- Recipient Category

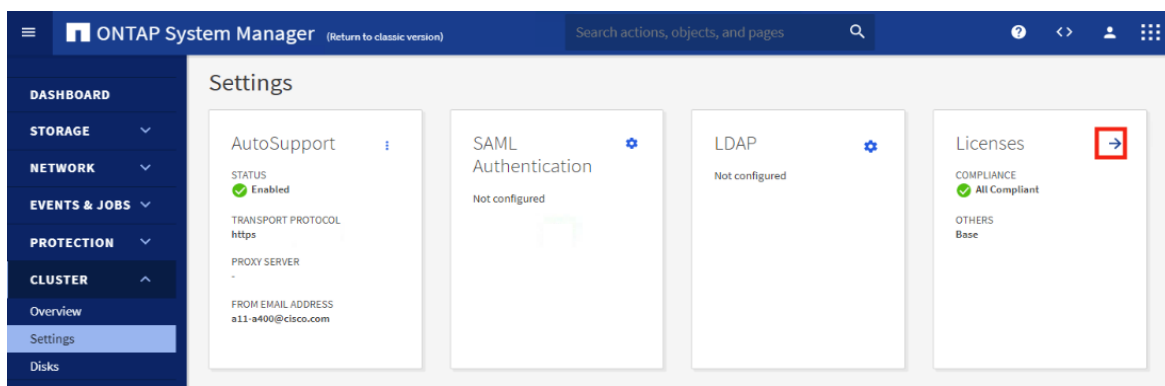
**Step 27.** Click Save when complete.



**Step 28.** Select CLUSTER > Settings at the top left of the page to return to the cluster settings page.

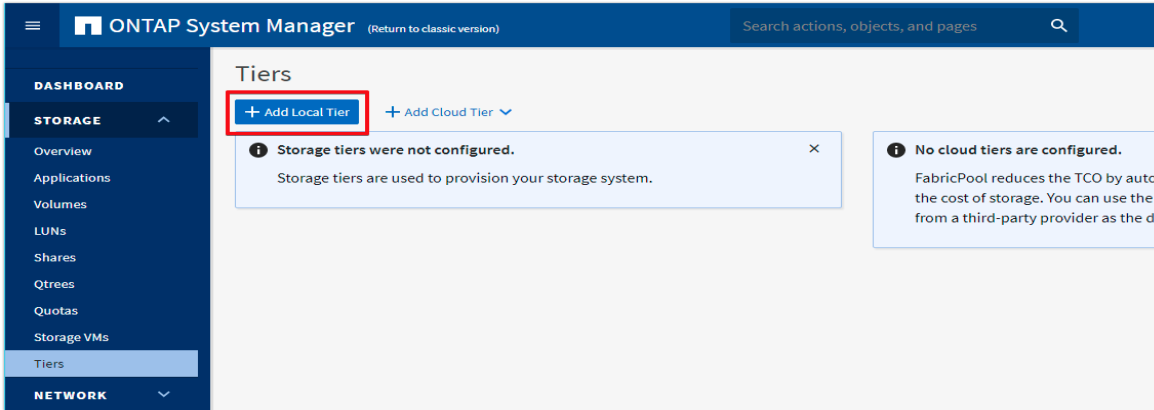
**Step 29.** Locate the Licenses tile on the right and click the detail arrow.

**Step 30.** Add the desired licenses to the cluster by clicking Add and entering the license keys in a comma separated list.



**Step 31.** Configure storage aggregates by selecting the Storage menu on the left and selecting Tiers.

**Step 32.** Click Add Local Tier and allow ONTAP System Manager to recommend a storage aggregate configuration.



**Step 33.** ONTAP will use best practices to recommend an aggregate layout. Click the Recommended details link to view the aggregate information.

**Step 34.** Optionally, enable NetApp Aggregate Encryption (NAE) by checking the box for Configure Onboard Key Manager for encryption.

**Step 35.** Enter and confirm the passphrase and save it in a secure location for future use.

**Step 36.** Click Save to make the configuration persistent.

### Add Local Tier ✕

**Storage Recommendation**

**32.6 TB**  
USABLE

2 local tiers can be added on nodes aa16-a400-02 and aa16-a400-01.

^ Recommendation details ←

Node Name	Local Tier	Usable Size	Type
aa16-a400-02	aa16_a400_02_NVME_...	16.3 TB	SSD
aa16-a400-01	aa16_a400_01_NVME_...	16.3 TB	SSD

---

**Encryption** Considerations

Configure Onboard Key Manager for encryption

✕ 👁

i Save the passphrase for future use. You will need the passphrase if the system needs to be recovered.

Cancel
Save

**Note:** Aggregate encryption may not be supported for all deployments. Please review the [NetApp Encryption Power Guide](#) and the [Security Hardening Guide for NetApp ONTAP 9 \(TR-4569\)](#) to help determine if aggregate encryption is right for your environment.

## Procedure 4. Log into the Cluster

**Step 1.** Open an SSH connection to either the cluster IP or the host name.

**Step 2.** Log into the admin user with the password you provided earlier.

## Procedure 5. Verify Storage Failover

**Step 1. Verify the status of the storage failover:**

```
A400::> storage failover show

Node                Partner      Takeover
Possible State Description
-----
A400-01  A400-02  true    Connected to A400-02
A400-02  A400-01  true    Connected to A400-01
2 entries were displayed.
```

**Note:** Both <st-node01> and <st-node02> must be capable of performing a takeover. Continue with step 2 if the nodes can perform a takeover.

**Step 2. Enable failover on one of the two nodes if it was not completed during the installation:**

```
storage failover modify -node <st-node01> -enabled true
```

**Note:** Enabling failover on one node enables it for both nodes.

**Step 3. Verify the HA status for a two-node cluster:**

**Note:** This step is not applicable for clusters with more than two nodes.

```
A400::> cluster ha show
High-Availability Configured: true
```

**Note:** If HA is not configured use the following commands. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

**Step 4. Verify that hardware assist is correctly configured:**

```
A400::> storage failover hwassist show
Node
-----
A400-01
Partner: A400-02
Hwassist Enabled: true
Hwassist IP: 192.0.2.84
Hwassist Port: 162
Monitor Status: active
Inactive Reason: -
Corrective Action: -
Keep-Alive Status: healthy

A400-02
Partner: A400-01
Hwassist Enabled: true
Hwassist IP: 192.0.2.85
Hwassist Port: 162
```

```
Monitor Status: active
Inactive Reason: -
Corrective Action: -
Keep-Alive Status: healthy
2 entries were displayed.
```

**Step 5.** If hwassist storage failover is not enabled, enable using the following commands:

```
storage failover modify -hwassist-partner-ip <node02-mgmt-ip> -node <st-node01>
storage failover modify -hwassist-partner-ip <node01-mgmt-ip> -node <st-node02>
```

## Procedure 6. Set Auto-Revert on Cluster Management

**Step 1.** Set the `auto-revert` parameter on the cluster management interface:

**Note:** A storage virtual machine (SVM) is referred to as a Vserver or `vserver` in the GUI and CLI.

```
net interface modify -vserver <clustername> -lif cluster_mgmt_lif_1 -auto-revert true
```

## Procedure 7. Zero All Spare Disks

**Step 1.** Zero all spare disks in the cluster by running the following command:

```
disk zerospares
```

### Tech tip

Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an AFF configuration. Disk autoassign should have assigned one data partition to each node in an HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

## Procedure 8. Set Up Service Processor Network Interface

**Step 1.** Assign a static IPv4 address to the Service Processor on each node by running the following commands:

```
system service-processor network modify -node <st-node01> -address-family IPv4 -enable true -dhcp none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>
system service-processor network modify -node <st-node02> -address-family IPv4 -enable true -dhcp none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```

**Note:** The Service Processor IP addresses should be in the same subnet as the node management IP addresses.

## Procedure 9. Create Manual Provisioned Aggregates - Optional

**Note:** An aggregate containing the root volume is created during the ONTAP setup process. To manually create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it should contain.

**Step 1.** Create new aggregates by running the following commands:



```

storage aggregate create -aggregate <aggr1_node01> -node <st-node01> -diskcount <num-
disks> -disktype SSD-NVM

storage aggregate create -aggregate <aggr1_node02> -node <st-node02> -diskcount <num-
disks> -disktype SSD-NVM

```

**Note:** You should have the minimum number of hot spare disks for the recommended hot spare disk partitions for their aggregate.

**Note:** For all-flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.

#### Tech tip

In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all, but one remaining disk (spare) assigned to the controller.

**Step 2.** The aggregate cannot be created until disk zeroing completes. Run the `storage aggregate show` command to display the aggregate creation status. Do not proceed until both `aggr1_node1` and `aggr1_node2` are online.

### Procedure 10. Remove Default Broadcast Domains

**Note:** By default, all network ports are included in separate default broadcast domain. Network ports used for data services (for example, `e0e`, `e0f`, and so on) should be removed from their default broadcast domain and that broadcast domain should be deleted.

**Step 1.** Run the following commands:

```

network port broadcast-domain delete -broadcast-domain <Default-N> -ipspace Default
network port broadcast-domain show

```

**Step 2.** Delete the Default broadcast domains with Network ports (Default-1, Default-2, and so on).

### Procedure 11. Disable Flow Control on 25/100GbE Data Ports

**Step 1.** Disable the flow control on 25 and 100GbE data ports by running the following command to configure the ports on node 01:

```

network port modify -node <st-node01> -port e3a,e3b -flowcontrol-admin none
network port modify -node <st-node01> -port e0e,e0f,e0g,e0h -flowcontrol-admin none

```

**Step 2.** Run the following command to configure the ports on node 02:

```

network port modify -node <st-node02> -port e3a,e3b -flowcontrol-admin none
network port modify -node <st-node02> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
A400::> net port show -node * -port e0e,e0f,e0g,e0h -fields speed-admin,duplex-
admin,flowcontrol-admin

```

```

(network port show)
node    port duplex-admin speed-admin flowcontrol-admin
-----
A400-01 e0e  auto          auto          none
A400-01 e0f  auto          auto          none
A400-01 e0g  auto          auto          none

```

```

A400-01 e0h auto auto none
A400-02 e0e auto auto none
A400-02 e0f auto auto none
A400-02 e0g auto auto none
A400-02 e0h auto auto none

```

8 entries were displayed.

```

A400::> net port show -node * -port e3a,e3b -fields speed-admin,duplex-
admin,flowcontrol-admin

```

(network port show)

```

node    port duplex-admin speed-admin flowcontrol-admin
-----
A400-01 e3a auto auto none
A400-01 e3b auto auto none
A400-02 e3a auto auto none
A400-02 e3b auto auto none

```

4 entries were displayed.

### Procedure 12. Disable Auto-Negotiate on Fibre Channel Ports - Required only for FC configuration

In accordance with the best practices for FC host ports, to disable auto-negotiate on each FCP adapter in each controller node, follow these steps:

**Step 1.** Disable each FC adapter in the controllers with the `fc adapter modify` command:

```

fc adapter modify -node <st-node01> -adapter 1a -status-admin down
fc adapter modify -node <st-node01> -adapter 1b -status-admin down
fc adapter modify -node <st-node02> -adapter 1a -status-admin down
fc adapter modify -node <st-node02> -adapter 1b -status-admin down

```

**Step 2.** Set the desired speed on the adapter and return it to the online state:

```

fc adapter modify -node <st-node01> -adapter 1a -speed 32 -status-admin up
fc adapter modify -node <st-node01> -adapter 1b -speed 32 -status-admin up
fc adapter modify -node <st-node02> -adapter 1a -speed 32 -status-admin up
fc adapter modify -node <st-node02> -adapter 1b -speed 32 -status-admin up

```

### Procedure 13. Enable Cisco Discovery Protocol

**Step 1.** Enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers by running the following command to enable CDP in ONTAP:

```

node run -node * options cdpd.enable on

```

### Procedure 14. Enable Link-layer Discovery Protocol on all Ethernet Ports

**Step 1.** Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches, on all ports, of all nodes in the cluster, by running the following command:

```

node run * options lldp.enable on

```

### Procedure 15. Create Management Broadcast Domain

**Step 1.** If the management interfaces are required to be on a separate VLAN, create a new broadcast domain for those interfaces by running the following command:

```
network port broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500
```

### Procedure 16. Create NFS Broadcast Domain

**Step 1.** To create a NFS, data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands to create a broadcast domain for NFS in ONTAP:

```
network port broadcast-domain create -broadcast-domain Infra-NFS -mtu 9000
```

### Procedure 17. Create CIFS Broadcast Domain

**Step 1.** To create a CIFS data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands to create a broadcast domain for CIFS in ONTAP:

```
network port broadcast-domain create -broadcast-domain Infra-CIFS -mtu 9000
```

### Procedure 18. Create ISCSI Broadcast Domains - Required only for iSCSI configuration

**Step 1.** To create an ISCSI-A and ISCSI-B data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands to create a broadcast domain for NFS in ONTAP:

```
network port broadcast-domain create -broadcast-domain Infra-ISCSI-A -mtu 9000
network port broadcast-domain create -broadcast-domain Infra-ISCSI-B -mtu 9000
```

### Procedure 19. Create Interface Groups

**Step 1.** To create the LACP interface groups for the 25GbE data interfaces, run the following commands:

```
network port ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0e
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0f
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0h
network port ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0e
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0f
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0h
```

**Step 2.** To verify, run the following:

```
A400::> network port ifgrp show
      Port      Distribution      Active
Node   IfGrp      Function      MAC Address      Ports   Ports
-----
A400-01
      a0a      port      d2:39:ea:18:01:48  full   e0e, e0f, e0g, e0h
A400-02
      a0a      port      d2:39:ea:17:e8:78  full   e0e, e0f, e0g, e0h
2 entries were displayed.
```

### Procedure 20. Change MTU on Interface Groups

**Step 1.** To change the MTU size on the base interface-group ports before creating the VLAN ports, run the following commands:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

## Procedure 21. Create VLANs

**Step 1.** Create the management VLAN ports and add them to the management broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<ib-mgmt-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<ib-mgmt-vlan-id>

network port broadcast-domain add-ports -broadcast-domain IB-MGMT -ports <st-
node01>:a0a-<ib-mgmt-vlan-id>,<st-node02>:a0a-<ib-mgmt-vlan-id>
```

**Step 2.** To verify, run the following command:

```
A400::> network port vlan show
```

```

                Network Network
Node  VLAN Name Port  VLAN ID  MAC Address
-----
A400-01
    a0a-31  a0a    31      d2:39:ea:18:01:48
    a0a-33  a0a    33      d2:39:ea:18:01:48
    a0a-41  a0a    41      d2:39:ea:18:01:48
    a0a-42  a0a    42      d2:39:ea:18:01:48
    a0a-45  a0a    45      d2:39:ea:18:01:48
    a0a-60  a0a    60      d2:39:ea:18:01:48
    a0a-61  a0a    61      d2:39:ea:18:01:48
    a0a-62  a0a    62      d2:39:ea:18:01:48
    a0a-63  a0a    63      d2:39:ea:18:01:48
A400-02
    a0a-31  a0a    31      d2:39:ea:17:e8:78
    a0a-33  a0a    33      d2:39:ea:17:e8:78
    a0a-41  a0a    41      d2:39:ea:17:e8:78
    a0a-42  a0a    42      d2:39:ea:17:e8:78
    a0a-45  a0a    45      d2:39:ea:17:e8:78
    a0a-60  a0a    60      d2:39:ea:17:e8:78
    a0a-61  a0a    61      d2:39:ea:17:e8:78
    a0a-62  a0a    62      d2:39:ea:17:e8:78
    a0a-63  a0a    63      d2:39:ea:17:e8:78
```

18 entries were displayed.

**Step 3.** Create the NFS VLAN ports and add them to the `Infra-NFS` broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>
```

```
network port broadcast-domain add-ports -broadcast-domain Infra-NFS -ports <st-
node01>:a0a-<infra-nfs-vlan-id>,<st-node02>:a0a-<infra-nfs-vlan-id>
```

**Step 4.** Create the CIFS VLAN ports and add them to the Infra-CIFS broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-cifs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-cifs-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-CIFS -ports <st-
node01>:a0a-<infra-cifs-vlan-id>,<st-node02>:a0a-<infra-cifs-vlan-id>
```

**Step 5.** If configuring iSCSI, create VLAN ports for the iSCSI LIFs on each storage controller and add them to the broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-b-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-a-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-b-vlan-id>

network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-
node01>:a0a-<infra-iscsi-a-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-
node01>:a0a-<infra-iscsi-b-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-A -ports <st-
node02>:a0a-<infra-iscsi-a-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-iSCSI-B -ports <st-
node02>:a0a-<infra-iscsi-b-vlan-id>
```

## Procedure 22. Configure Time Synchronization on the Cluster

**Step 1.** Set the time zone for the cluster:

```
timezone -timezone <timezone>
```

**Note:** For example, in the eastern United States, the time zone is America/New\_York.

## Procedure 23. Configure Simple Network Management Protocol - SNMP

**Step 1.** Configure basic SNMP information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP:

```
snmp contact <snmp-contact>
snmp location "<snmp-location>"
snmp init 1
options snmp.enable on
```

**Step 2.** Configure SNMP traps to send to remote hosts, such as an Active IQ Unified Manager server or another fault management system:

```
snmp traphost add <oncommand-um-server-fqdn>
```

## Procedure 24. Configure SNMPv3 Access

**Note:** SNMPv3 offers advanced security by using encryption and passphrases. The SNMPv3 user can run SNMP utilities from the traphost using the authentication and privacy settings that you specify.

**Step 1.** Configure the SNMPv3 access by running the following command:

```
security login create -user-or-group-name <<snmp-v3-usr>> -application snmp -
authentication-method usm
```

## Step 2. Enter the authoritative entity's EngineID [local EngineID]:

```
Which authentication protocol do you want to choose (none, md5, sha, sha2-256) [none]:
<<snmp-v3-auth-proto>>
Enter the authentication protocol password (minimum 8 characters long):
Enter the authentication protocol password again:
Which privacy protocol do you want to choose (none, des, aes128) [none]: <<snmpv3-priv-
proto>>
Enter privacy protocol password (minimum 8 characters long):
Enter privacy protocol password again:
```

Refer to the [SNMP Configuration Express Guide](#) for additional information when configuring SNMPv3 security users.

## Procedure 25. Create an Infrastructure SVM

### Step 1. Run the `vserver create` command:

```
vserver create -vserver Infra-FC -rootvolume infra_FC_root -aggregate aggr1_node01 -
rootvolume-security-style unix
```

**Note:** It is recommended to remove iSCSI or FCP protocols if the protocol is not in use.

### Step 2. Add the two data aggregates to the Infra-FC aggregate list for the NetApp ONTAP Tools:

```
vserver modify -vserver Infra-FC -aggr-list <aggr1_node01>,<aggr1_node02>
```

### Step 3. Enable and run the NFS protocol in the Infra-FC:

```
vserver nfs create -vserver Infra-FC -udp disabled -v3 enabled -v4.1 enabled -vstorage
enabled
```

**Note:** If the NFS license was not installed during the cluster configuration, make sure to install the license before starting the NFS service.

**Note:** Verify the NFS `vstorage` parameter for the NetApp NFS VAAI plug-in was enabled:

```
A400::> vserver nfs show -fields vstorage
vserver      vstorage
-----
Infra-FC      enabled
epic-svm      enabled
hc-svm-fc     enabled
hc-svm-nvme   enabled
hc-svm-nvme2  enabled
test-SVM      enabled
7 entries were displayed.
```

## Procedure 26. Configure CIFS Servers

**Note:** You can enable and configure CIFS servers on storage virtual machines (SVMs) with NetApp FlexVol<sup>®</sup> volumes to let SMB clients access files on your cluster. Each data SVM in the cluster can be bound to exactly one Active Directory domain. However, the data SVMs do not need to be bound to the same domain. Each data SVM can be bound to a unique Active Directory domain.

### Step 1. Configure the DNS for your SVM.

```
dns create -vserver Infra-FC -domains <domain_name> -name-servers <dns_server_ip>
```

**Note:** The node management network interfaces should be able to route to the Active Directory domain controller to which you want to join the CIFS server. Alternatively, a data network interface must exist on the SVM that can route to the Active Directory domain controller.

### Step 2. Create a network interface on the IB-MGMT VLAN:

```
network interface create -vserver Infra-FC -lif <<svm_mgmt_lif_name>> -role data -data-protocol none -home-node <<st-node-01>> -home-port a0a-<IB-MGMT-VLAN> -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```

### Step 3. Create the CIFS service:

```
vserver cifs create -vserver Infra-FC -cifs-server FS01 -domain <domain.com>
```

In order to create an Active Directory machine account for the CIFS server, you must supply the name and password of a Windows account with sufficient privileges to add computers to the "CN=Computers" container within the "DOMAIN.COM" domain.

Enter the user name: Administrator@active diectory.local

Enter the password:

## Procedure 27. Modify Storage Virtual Machine Option

**Note:** NetApp ONTAP can use automatic node referrals to increase SMB client performance on SVMs with FlexVol volumes. This feature allows the SVM to automatically redirect a client request to a network interface on the node where the FlexVol volume resides.

### Step 1. Run the following command to enable automatic node referrals on your SVM:

```
set -privilege advanced
vserver cifs options modify -vserver Infra-FC -is-referral-enabled true
```

## Procedure 28. Create Load-Sharing Mirrors of a SVM Root Volume

### Step 1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node:

```
volume create -vserver Infra-FC -volume Infra_FC_root_m01 -aggregate <aggr1_node01> -size 1GB -type DP
volume create -vserver Infra-FC -volume Infra_FC_root_m01 -aggregate <aggr1_node02> -size 1GB -type DP
```

### Step 2. Create a job schedule to update the root volume mirror relationships every 15 minutes:

```
job schedule interval create -name 15min -minutes 15
```

### Step 3. Create the mirroring relationships:

```
snapmirror create -source-path Infra-FC:infra_FC_root -destination-path Infra-FC:infra_FC_root_m01 -type LS -schedule 15min
snapmirror create -source-path Infra-FC:infra_FC_root -destination-path Infra-FC:infra_FC_root_m02 -type LS -schedule 15min
```

Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path Infra-FC:infra_FC_root
```

### Step 4. To verify, run the following:

```
A400::> snapmirror show -type ls
```

Source Path	Destination Type	Mirror Path	Relationship State	Relationship Status	Total Progress	Healthy	Last Updated
A400://Infra-FC/Infra_FC_root	LS	A400://Infra-FC/Infra_FC_root_m01	Snapmirrored	Idle	-	true	-
		A400://Infra-FC/Infra_FC_root_m02	Snapmirrored	Idle	-	true	-
A400://hc-svm-fc/hc_svm_fc_root	LS	A400://hc-svm-fc/hc_svm_fc_root_m01	Snapmirrored	Idle	-	true	-
		A400://hc-svm-fc/hc_svm_fc_root_m02	Snapmirrored	Idle	-	true	-
A400://hc-svm-nvme/hc_svm_nvme_root	LS	A400://hc-svm-nvme/hc_svm_nvme_root_m01	Snapmirrored	Idle	-	true	-
		A400://hc-svm-nvme/hc_svm_nvme_root_m02	Snapmirrored	Idle	-	true	-
A400://hc-svm-nvme2/hc_svm_nvme2_root	LS	A400://hc-svm-nvme2/hc_svm_nvme2_root_m01	Snapmirrored	Idle	-	true	-
		A400://hc-svm-nvme2/hc_svm_nvme2_root_m02	Snapmirrored	Idle	-	true	-

8 entries were displayed.

### Procedure 29. Create FC Block Protocol Service -required only for FC configuration

**Step 1.** Run the following command to create the FCP service. This command also starts the FCP service and sets the worldwide name (WWN) for the SVM:

```
vserver fcp create -vserver Infra-FC -status-admin up
```

**Step 2.** To verify, run the following:

```
A400::> vserver fcp show
```

Vserver	Target Name	Status Admin
---------	-------------	--------------



```

-----
CIT-VDI-FC      20:32:d0:39:ea:18:01:47      up
Infra-FC        20:0a:d0:39:ea:18:01:47      up
epic-svm        20:01:d0:39:ea:18:01:47      up
hc-svm-fc       20:0f:d0:39:ea:18:01:47      up
hc-svm-nvme     20:14:d0:39:ea:18:01:47      up
hc-svm-nvme2    20:1e:d0:39:ea:18:01:47      up

```

6 entries were displayed.

If the FC license was not installed during the cluster configuration, make sure to install the license before creating the FC service.

### Procedure 30. Vserver Protocol Verification

**Step 1.** Verify the protocols are added to the Infra vsver by running the following:

```

A400::> vsver show-protocols -vsver Infra-FC
Vserver: Infra-FC
Protocols: nfs, cifs, fcp

```

**Step 2.** If a protocol is not present, use the following command to add the protocol to the vsver:

```

vsver add-protocols -vsver <infra-data-svm> -protocols < iscsi or fcp >

```

### Procedure 31. Configure HTTPS Access to the Storage Controller

**Step 1.** Increase the privilege level to access the certificate commands:

```

set -privilege diag
Do you want to continue? {y|n}: y

```

**Step 2.** Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, the <serial-number>) by running the following command:

```

security certificate show

```

**Step 3.** For each SVM shown, the certificate common name should match the DNS fully qualified domain name (FQDN) of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```

security certificate delete -vsver Infra-FC -common-name Infra-FC -ca Infra-FC -type
server -serial <serial-number>

```

**Step 4.** Deleting expired certificates before creating new certificates is a best practice. Run the `security certificate delete` command to delete the expired certificates. In the following command, use TAB completion to select and delete each default certificate.

**Step 5.** To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Infra-FC and the cluster SVM. Use TAB completion to aid in the completion of these commands:

```

security certificate create -common-name <cert-common-name> -type server -size 2048 -
country <cert-country> -state <cert-state> -locality <cert-locality> -organization
<cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol
SSL -hash-function SHA256 -vsver Infra-FC

```

**Step 6.** To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the `security certificate show` command.

**Step 7.** Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use TAB completion to aid in the completion of these commands:

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -
ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

**Step 8.** Disable HTTP cluster management access:

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```

**Note:** It is normal for some of these commands to return an error message stating that the entry does not exist.

**Step 9.** Return to the normal admin privilege level and verify that the system logs are available in a web browser:

```
set -privilege admin
https://<node01-mgmt-ip>/spi
https://<node02-mgmt-ip>/spi
```

### Procedure 32. Configure NFSv3 and NFSv4.1

**Step 1.** Create a new rule for the infrastructure NFS subnet in the default export policy:

```
vserver export-policy rule create -vserver Infra-FC -policyname default -ruleindex 1 -
protocol nfs -clientmatch <infra-nfs-subnet-cidr> -rorule sys -rwrule sys -superuser sys
-allow-suid true
```

**Step 2.** Assign the FlexPod export policy to the infrastructure SVM root volume:

```
volume modify -vserver Infra-FC -volume infra_FC_root -policy default
```

### Procedure 33. Create CIFS Export Policy

**Note:** Optionally, you can use export policies to restrict CIFS access to files and folders on CIFS volumes. You can use export policies in combination with share level and file level permissions to determine effective access rights.

**Step 1.** Run the following command to create an export policy that limits access to devices in the domain:

```
export-policy create -vserver Infra-FC -policyname cifs
export-policy rule create -vserver Infra-FC -policyname cifs -clientmatch <domain_name>
-rorule
krb5i,krb5p -rwrule krb5i,krb5p
```

### Procedure 34. Create a NetApp FlexVol Volume

The following information is required to create a NetApp FlexVol® volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

**Step 1.** Run the following commands:

```
volume create -vserver Infra-FC -volume infra_datastore_1 -aggregate <aggr1_node01> -
size 1TB -state online -policy default -junction-path /infra_datastore_01 -space-
guarantee none -percent-snapshot-space 0
volume create -vserver Infra-FC -volume infra_datastore_2 -aggregate <aggr1_node02> -
size 1TB -state online -policy default -junction-path /infra_datastore_02 -space-
guarantee none -percent-snapshot-space 0
```

```
volume create -vserver Infra-FC -volume infra_swap -aggregate <aggr1_node01> -size 100GB
-state online -policy default -junction-path /infra_swap -space-guarantee none -percent-
snapshot-space 0 -snapshot-policy none.
```

```
volume create -vserver Infra-FC -volume esxi_boot -aggregate <aggr1_node01> -size 320GB
-state online -policy default -space-guarantee none -percent-snapshot-space 0
```

```
snapmirror update-ls-set -source-path Infra-FC:infra_FC_root
```

**Step 2.** If you are going to setup and use SnapCenter to backup the infra\_datastore volume, add “-snapshot-policy none” to the end of the volume create command for the infra\_datastore volume.

### Procedure 35. Create a NetApp FlexGroup Volume

#### Tech tip

A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. A FlexGroup Volume contains several constituents that automatically and transparently share the traffic. A FlexGroup volume is a single namespace container that can be managed in a similar way as FlexVol volumes.

**Step 1.** Run the following commands to create FlexGroup volumes:

```
volume create -vserver Infra-FC -volume cifs_vol_01 -aggr-list
aggr01_node01,aggr01_node02-aggr-list-multiplier4-state online -policy cifs_policy -size
800GB -junction-path /cifs_vol_01 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-FC -volume cifs_vol_02 -aggr-list
aggr01_node01,aggr01_node02-aggr-list-multiplier4-state online -policy cifs_policy -size
800GB -junction-path /cifs_vol_02 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-FC -volume cifs_vol_03 -aggr-list
aggr01_node01,aggr01_node02-aggr-list-multiplier4-state online -policy cifs_policy -size
800GB -junction-path /cifs_vol_03 -space-guarantee none -percent-snapshot-space 0
```

### Procedure 36. Modify Volume Efficiency

**Step 1.** On NetApp AFF systems, deduplication is enabled by default. To disable the efficiency policy on the infra\_swap volume, run the following command:

```
volume efficiency off -vserver Infra-FC -volume infra_swap
```

### Procedure 37. Create CIFS Shares

**Note:** A CIFS share is a named access point in a volume that enables CIFS clients to view, browse, and manipulate files on a file server.

**Step 1.** Run the following commands to create CIFS shares:

```
cifs share create -vserver Infra-FC -share-name <CIFS_share_1> -path /infra_datastore_01
-share properties oplocks,browsable,continuously-available,showsnapshot
cifs share create -vserver Infra-FC -share-name <CIFS_share_2> -path /infra_datastore_02
-share properties oplocks,browsable,continuously-available,showsnapshot
cifs share create -vserver Infra-FC -share-name <CIFS_share_3> -path /cifs_vol_03 -
share properties oplocks,browsable,continuously-available,showsnapshot
```

### Procedure 38. Create NFS LIFs

**Step 1.** Run the following commands to create NFS LIFs:

```
network interface create -vserver Infra-FC -lif NFS-1-A400-01 -role data -data-protocol
nfs -home-node <st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs-lif-
```

```
01-ip> -netmask <node01-nfs-lif-01-mask> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-FC -lif NFS-2-A400-02 -role data -data-protocol
nfs -home-node <st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs-lif-
02-ip> -netmask <node02-nfs-lif-02-mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
```

**Step 2. Run the following commands to verify:**

```
A400::> network interface show -vserver Infra-FC -data-protocol nfs
```

	Logical	Status	Network	Current	Current	Is
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port	Home
-----						
Infra-FC						
	NFS-1-A400-01	up/up	10.10.63.10/24	A400-01	a0a-63	true
	NFS-2-A400-02	up/up	10.10.63.11/24	A400-02	a0a-63	true

2 entries were displayed.

**Procedure 39. Create CIFS LIFs**

**Step 1. Run the following commands to create CIFS LIFs:**

```
network interface create -vserver Infra-FC -lif cifs_lif01 -role data -data-protocol
cifs -home-node <st-node01> -home-port a0a-<infra-cifs-vlan-id> -address <node01-
cifs_lif01-ip> -netmask <node01-cifs_lif01-mask> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-FC -lif cifs_lif02 -role data -data-protocol
cifs -home-node <st-node02> -home-port a0a-<infra-cifs-vlan-id> -address <node02-
cifs_lif02-ip> -netmask <node02-cifs_lif02-mask>> -status-admin up -failover-policy
broadcast-domain-wide -firewall-policy data -auto-revert true
```

**Procedure 40. Create FC LIFs - required only for FC configuration**

**Step 1. Run the following commands to create four FC LIFs (two on each node):**

```
network interface create -vserver Infra-FC -lif fcp-lif-01a -role data -data-protocol
fcp -home-node <st-node01> -home-port 1a -status-admin up

network interface create -vserver Infra-FC -lif fcp-lif-01b -role data -data-protocol
fcp -home-node <st-node01> -home-port 1b -status-admin up

network interface create -vserver Infra-FC -lif fcp-lif-02a -role data -data-protocol
fcp -home-node <st-node02> -home-port 1a -status-admin up

network interface create -vserver Infra-FC -lif fcp-lif-02b -role data -data-protocol
fcp -home-node <st-node02> -home-port 1b -status-admin up
```

**Step 2. Run the following commands to verify:**

```
A400::> network interface show -vserver Infra-FC -data-protocol fcp
```

	Logical	Status	Network	Current	Current	Is
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port	Home
-----						
Infra-FC						
	fcp-lif-01a	up/up	20:01:d0:39:ea:29:ce:d4	A400-01	1a	true
	fcp-lif-01b	up/up	20:02:d0:39:ea:29:ce:d4	A400-01	1b	true

```
fcp-lif-02a up/up 20:03:d0:39:ea:29:ce:d4 A400-02 1a true
fcp-lif-02b up/up 20:04:d0:39:ea:29:ce:d4 A400-02 1b true
4 entries were displayed.
```

## Procedure 41. Add Infrastructure SVM Administrator and SVM Administration LIF to In-band Management Network

### Step 1. Run the following commands:

```
network interface create -vserver Infra-FC -lif svm-mgmt -role data -data-protocol none
-home-node <st-node02> -home-port a0a-<ib-mgmt-vlan-id> -address <svm-mgmt-ip> -netmask
<svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy
mgmt -auto-revert true
```

### Step 2. Create a default route that enables the SVM management interface to reach the outside world:

```
network route create -vserver Infra-FC -destination 0.0.0.0/0 -gateway <svm-mgmt-
gateway>
```

### Step 3. To verify, run the following:

```
A400::> network route show -vserver Infra-FC
Vserver          Destination      Gateway          Metric
-----
Infra-FC         0.0.0.0/0       10.10.61.1      20
```

### Step 4. Set a password for the SVM vsadmin user and unlock the user:

```
security login password -username vsadmin -vserver Infra-FC
Enter a new password: <password>
Enter it again: <password>
security login unlock -username vsadmin -vserver Infra-FC
```

A cluster serves data through at least one and possibly several SVMs. By completing these steps, you have created a single data SVM. You can create additional SVMs depending on their requirement.

## Procedure 42. Configure and Test AutoSupport

NetApp AutoSupport sends support summary information to NetApp through HTTPS.

### Step 1. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <mailhost> -transport
https -support enable -noteto <storage-admin-email>
```

### Step 2. Test the AutoSupport configuration by sending a message from all nodes of the cluster:

```
autosupport invoke -node * -type all -message "FlexPod storage configuration completed"
```

The following is the configuration information that was modified from the platform guide to validate this solution:

- 32 Gbps HBA on slot 1 which was used for boot from SAN using FC. It can also be used for NVMe when required. By default, it stays in initiator type. You will need to change the type to target for the fcp adapter to be listed under network ports:

```
system node hardware unified-connect modify -node * -adapter <adapter-port>
```

- 3 FlexVol volumes are created for hosting virtual desktops, PVS share, and SMB share:

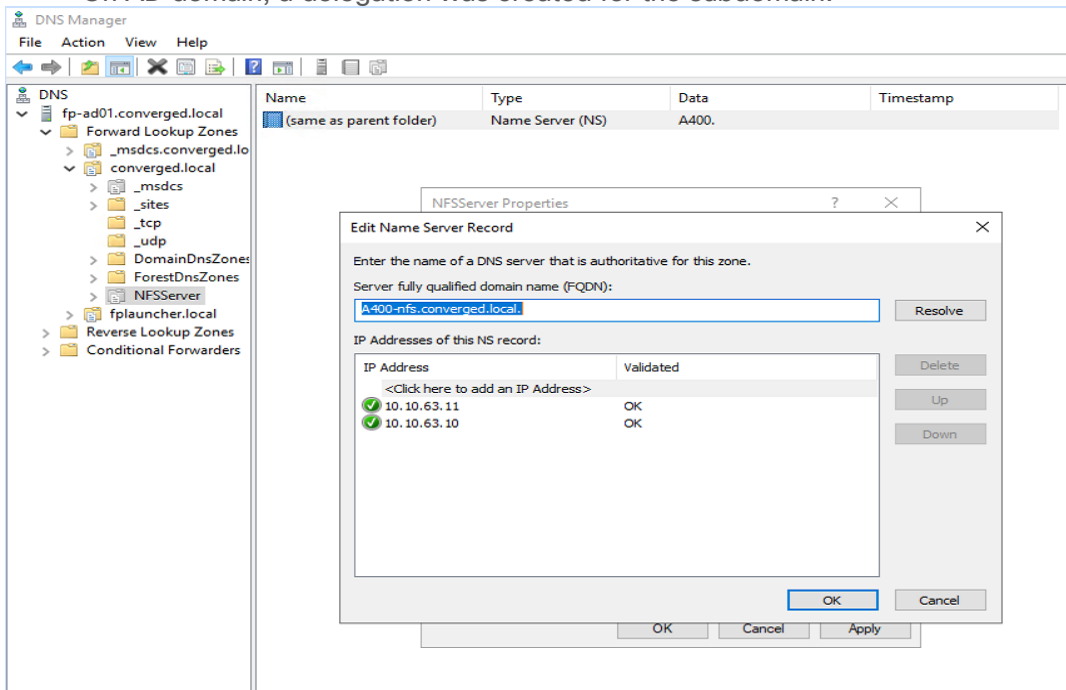
```
volume create -server <vserver> -volume <volumename> -aggr-list <aggr-node-01>,<aggr-node-
02> -aggr-list-multiplier <number_of_member_volume/aggr> -size <allocation_size> -security-
style <unix/ntfs> -qos-adaptive-policy-group <aqos_policy>
```

Name	Number of Members	Size	Adaptive QoS Policy	Expected IOPS (2048 * Allocated Space)	Peak IOPS (4096 * Used Space)
VDI	8	30TB (12% used)	performance	61440	14745.6
Data	8	10TB (25% used)	performance	20480	10240

For NFS, the DNS Load balancing feature was used and is available on ONTAP. (physical, interface groups, and VLANs). With DNS load balancing, LIFs are associated with the load balancing zone of an SVM. A site-wide DNS server is configured to forward all DNS requests and return the least-loaded LIF based on the network traffic and the availability of the port resources (CPU usage, throughput, open connections, and so on). DNS load balancing provides the following benefits:

- New client connections balanced across available resources.
- No manual intervention required for deciding which LIFs to use when mounting a particular SVM.
- DNS load balancing supports NFSv3, NFSv4, NFSv4.1, CIFS, SMB 2.0, SMB 2.1, and SMB 3.0.
- network interface modify -vserver <vserver\_name> -lif <lif\_name> -dns-zone <zone\_name>  
for example, network interface modify -vserver Infra-FC -lif NFS-1-A400-01 -dns-zone nfsserver.converged.local

On AD domain, a delegation was created for the subdomain.



## Cisco Intersight Managed Mode Configuration

This chapter contains the following:

- [Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects](#)
- [Configure a Cisco UCS Domain Profile](#)
- [Cisco UCS Domain Configuration](#)
- [Configure Cisco UCS Chassis Profile](#)
- [Configure Server Profile Template](#)
- [Management Configuration](#)
- [SAN Switch Configuration](#)
- [FlexPod Cisco MDS Switch Configuration](#)

The Cisco Intersight platform is a management solution delivered as a service with embedded analytics for Cisco and third-party IT infrastructures. The Cisco Intersight managed mode (also referred to as Cisco IMM or Intersight managed mode) is a new architecture that manages Cisco Unified Computing System (Cisco UCS) fabric interconnect-attached systems through a Redfish-based standard model. Cisco Intersight managed mode standardizes both policy and operation management Cisco UCS X210c M6 compute nodes used in this deployment guide.

### Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

This subject contains the following procedures:

- [Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects](#)
- [Set up a new Cisco Intersight account](#)
- [Set up Cisco Intersight Account and Associate it with Cisco Smart Licensing](#)
- [Set up Cisco Intersight Resource Group](#)
- [Set up Cisco Intersight Organization](#)
- [Claim Cisco UCS Fabric Interconnects in Cisco Intersight](#)
- [Verify the addition of Cisco UCS Fabric Interconnects to Cisco Intersight](#)

**Note:** Cisco UCS C-Series M6 servers, connected and managed through Cisco UCS FIs, are also supported by IMM. For a complete list of supported platforms, visit:

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/Intersight/b\\_Intersight\\_Managed\\_Mode\\_Configuration\\_Guide/b\\_intersight\\_managed\\_mode\\_guide\\_chapter\\_01010.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide/b_intersight_managed_mode_guide_chapter_01010.html)

#### Procedure 1. Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

**Note:** The Cisco UCS fabric interconnects need to be set up to support Cisco Intersight managed mode. When converting an existing pair of Cisco UCS fabric interconnects from Cisco UCS Manager mode to Intersight Managed Mode (IMM), first erase the configuration and reboot your system.

**WARNING! Converting fabric interconnects to Cisco Intersight managed mode is a disruptive process, and configuration information will be lost. You are encouraged to make a backup of their existing configuration.**

**Step 1.** Configure Fabric Interconnect A (FI-A). On the Basic System Configuration Dialog screen, set the management mode to Intersight. All the remaining settings are similar to those for the Cisco UCS Manager managed mode (UCSM-Managed).

```
Cisco UCS Fabric Interconnect A
To configure the Cisco UCS for use in a FlexPod environment in intersight managed mode, follow these steps:
```

**Step 2.** Connect to the console port on the first Cisco UCS fabric interconnect.

```
Enter the configuration method. (console/gui) ? console
Enter the management mode. (ucsm/intersight)? intersight
You have chosen to setup a new Fabric interconnect in "intersight" managed mode. Continue? (y/n): y
Enforce strong password? (y/n) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Enter the switch fabric (A/B) []: A
Enter the system name: <ucs-cluster-name>
Physical Switch Mgmt0 IP address : <ucsa-mgmt-ip>
Physical Switch Mgmt0 IPv4 netmask : <ucsa-mgmt-mask>
IPv4 address of the default gateway : <ucsa-mgmt-gateway>
Configure the DNS Server IP address? (yes/no) [n]: y
    DNS IP address : <dns-server-1-ip>
Configure the default domain name? (yes/no) [n]: y
    Default domain name : <ad-dns-domain-name>
<SNIP>
Verify and save the configuration.
```

**Step 3.** After applying the settings, make sure you can ping the fabric interconnect management IP address. When Fabric Interconnect A is correctly set up and is available, Fabric Interconnect B will automatically discover Fabric Interconnect A during its setup process as shown in the next step.

**Step 4.** Configure Fabric Interconnect B (FI-B). For the configuration method, choose console. Fabric Interconnect B will detect the presence of Fabric Interconnect A and will prompt you to enter the admin password for Fabric Interconnect A. Provide the management IP address for Fabric Interconnect B and apply the configuration.

```
Cisco UCS Fabric Interconnect A
Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added
to the cluster. Continue (y/n) ? y

Enter the admin password of the peer Fabric interconnect: <password>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucsa-mgmt-mask>

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address

Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>

Local fabric interconnect model(UCS-FI-6454)
Peer fabric interconnect is compatible with the local fabric interconnect. Continuing with the installer...

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```



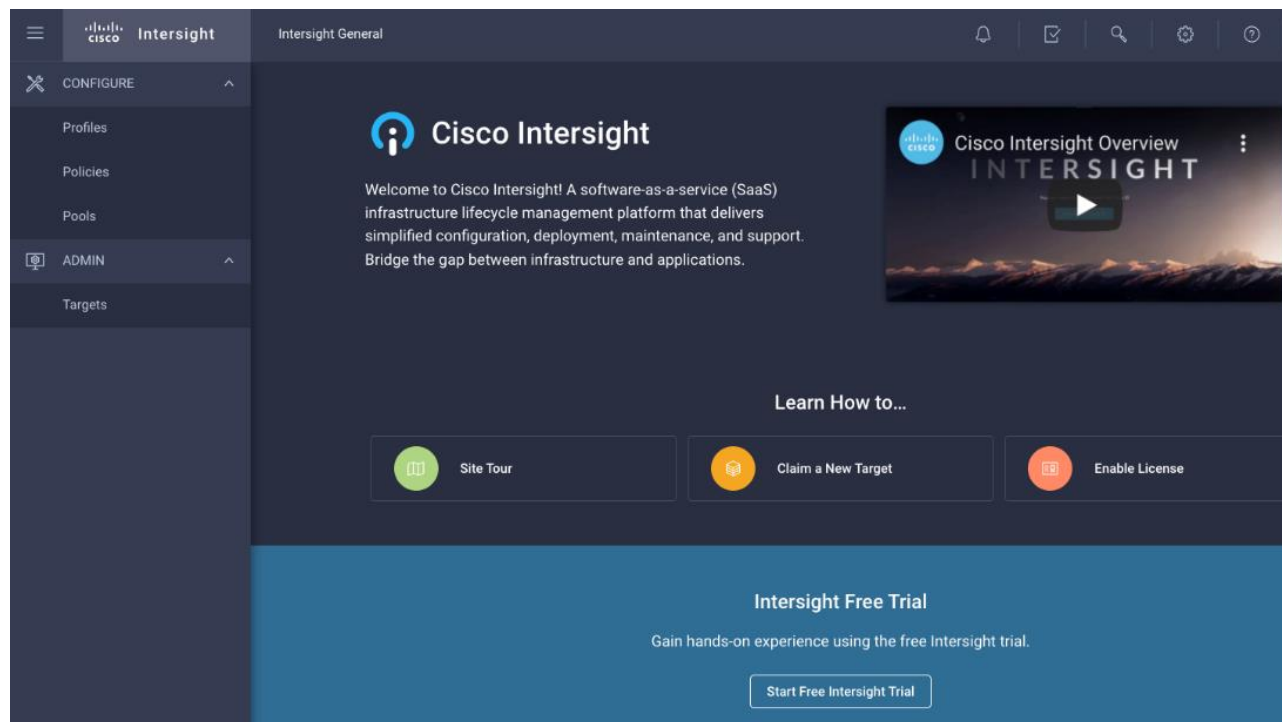
## Procedure 2. Set up a new Cisco Intersight account

**Step 1.** Go to <https://intersight.com> and click **Create an account**.

**Step 2.** Read and accept the license agreement. Click **Next**.

**Step 3.** Provide an Account Name and click **Create**.

On successful creation of the Intersight account, following page will be displayed:



**Note:** You can also choose to add the Cisco UCS FIs to an existing Cisco Intersight account.

## Procedure 3. Set up Cisco Intersight account and associate it with Cisco Smart Licensing

**Note:** When setting up a new Cisco Intersight account (as described in this document), the account needs to be enabled for Cisco Smart Software Licensing.

**Step 1.** Log into the Cisco Smart Licensing portal:

[https://software.cisco.com/software/cs/ws/platform/home?locale=en\\_US#module/SmartLicensing](https://software.cisco.com/software/cs/ws/platform/home?locale=en_US#module/SmartLicensing).


**Step 2.** Verify that the correct virtual account is selected.

**Step 3.** Under **Inventory > General**, generate a new token for product registration.

**Step 4.** Copy this newly created token.

## Create Registration Token

This will create a token that is used to register product instances, so that they can use licenses from this virtual account. Once it's created, go to the Smart Licensing configuration for your products and enter the token, to register them with this virtual account.


Virtual Account: Cisco  Intersight

Description :

\* Expire After:  Days  
Between 1 - 365, 30 days recommended


Max. Number of Uses:

The token will be expired when either the expiration or the maximum uses is reached

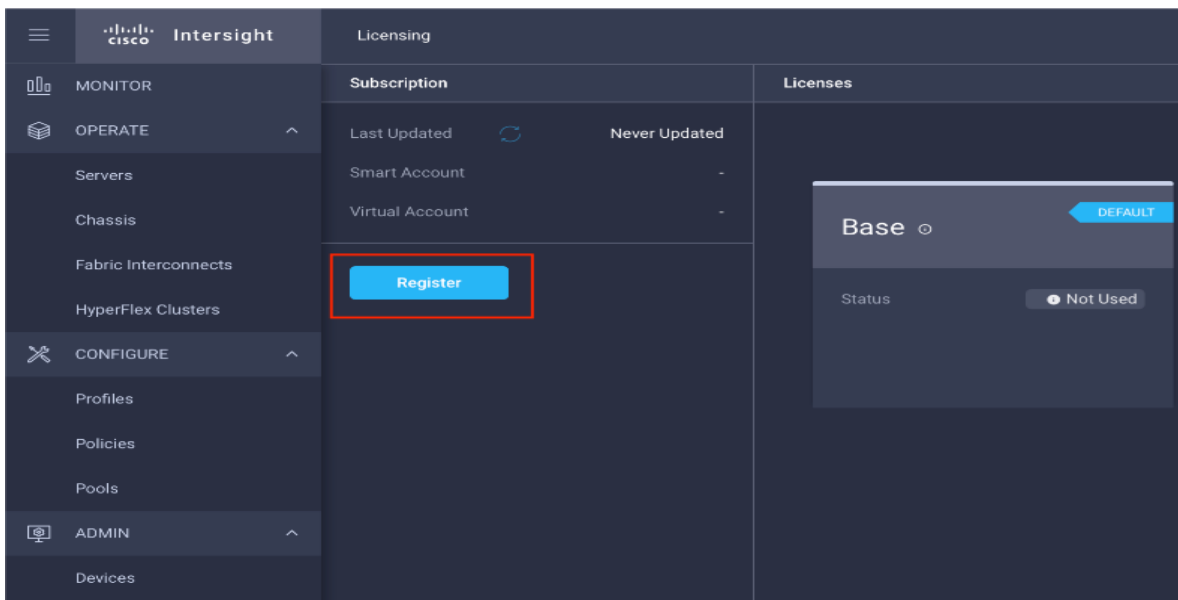
Allow export-controlled functionality on the products registered with this token 

Create Token Cancel

**Step 5.** Log into the Cisco Intersight portal and click **Settings** (the gear icon) in the top-right corner. Click **Licensing**.



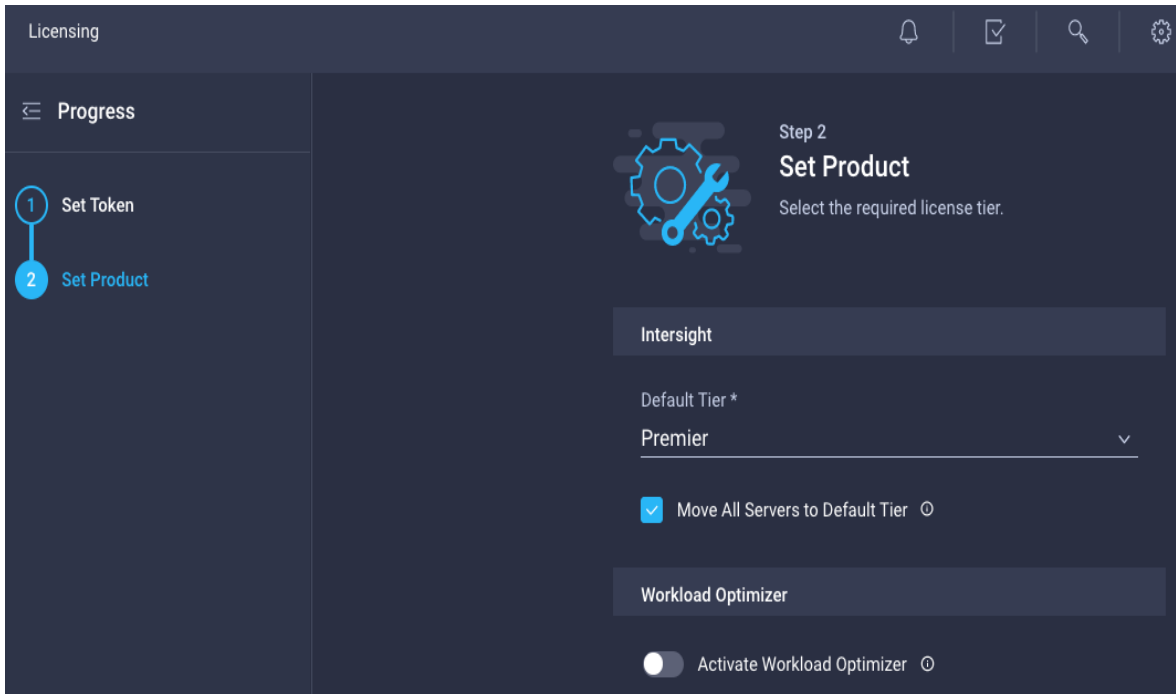
**Step 6.** Under **Cisco Intersight > Licensing**, click **Register**.



**Step 7.** Enter the copied token from the Cisco Smart Licensing portal. Click **Next**.

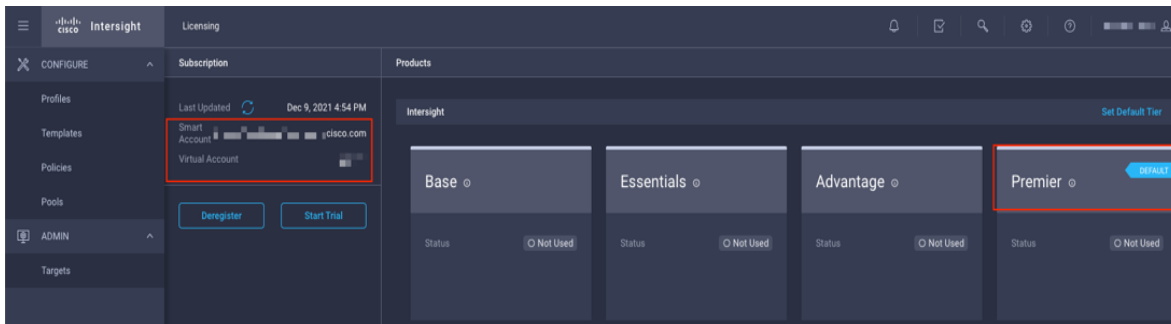
**Step 8.** From the drop-down list, select the pre-selected Default Tier \* and select the license type (for example, Premier).

**Step 9.** Select **Move All Servers to Default Tier**.



**Step 10. Click Register.**

When the registration is successful (takes a few minutes), the information about the associated Cisco Smart account and default licensing tier selected in the last step is displayed.

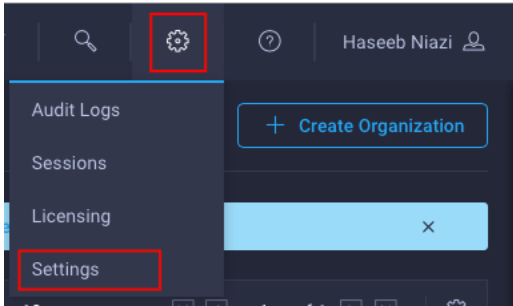


**Procedure 4. Set up Cisco Intersight Resource Group**

**Note:** In this step, a Cisco Intersight resource group is created where resources such as targets will be logically grouped. In this deployment, a single resource group is created to host all the resources, but customers can choose to create multiple resource groups for granular control of the resources.

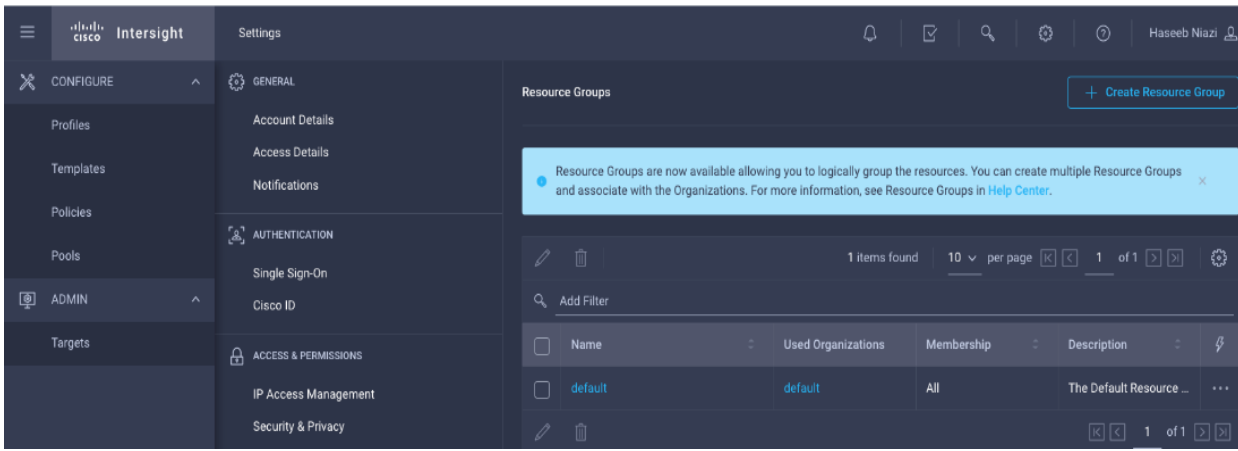
**Step 1. Log into Cisco Intersight.**

**Step 2. Click Settings (the gear icon) and click Settings.**

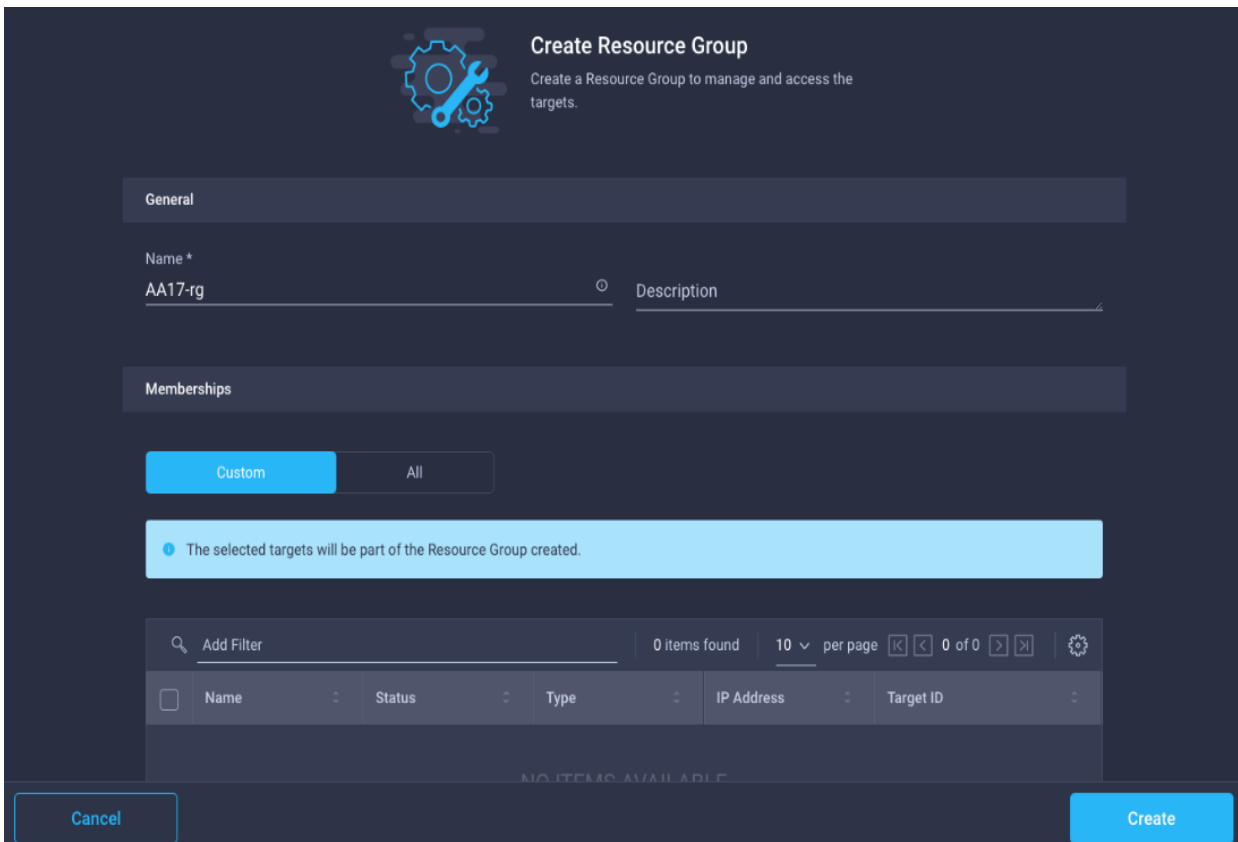


**Step 3.** Click **Resource Groups** in the middle panel.

**Step 4.** Click **+ Create Resource Group** in the top-right corner.



**Step 5.** Provide a name for the Resource Group (for example, rg).



**Step 6.** Under Memberships, click **Custom**.

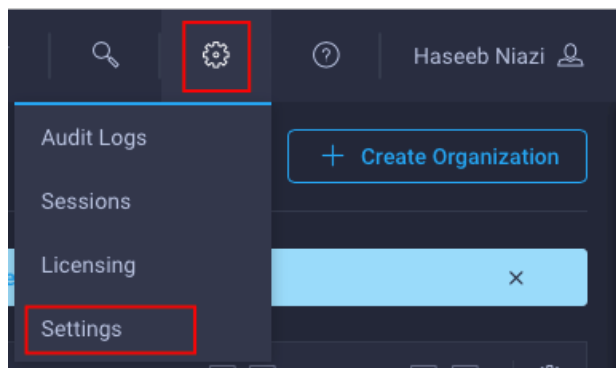
**Step 7.** Click **Create**.

### Procedure 5. Set up Cisco Intersight Organization

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

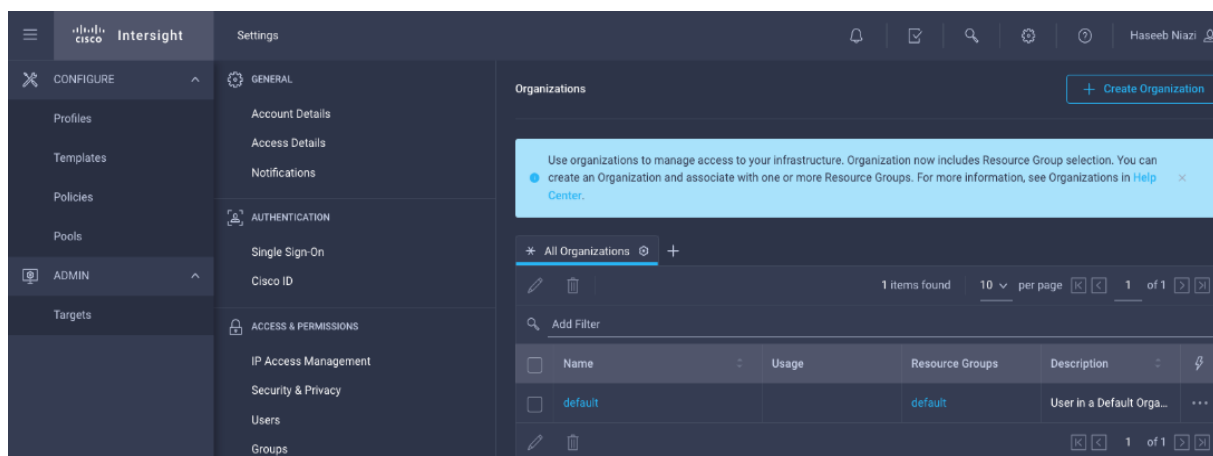
**Step 1.** Log into the Cisco Intersight portal.

**Step 2.** Click **Settings** (the gear icon) and select **Settings**.



**Step 3.** Click **Organizations** in the middle panel.

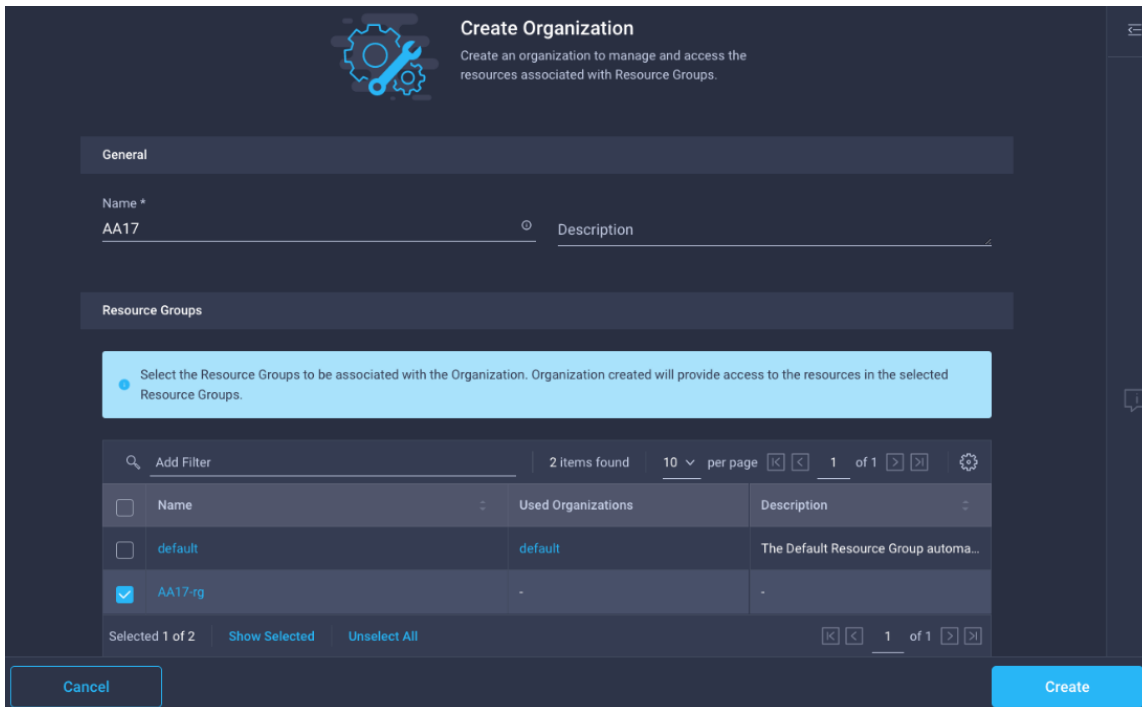
**Step 4.** Click **+ Create Organization** in the top-right corner.



**Step 5.** Provide a name for the organization (for example, AA17).

**Step 6.** Select the Resource Group created in the last step (for example, rg).

**Step 7.** Click **Create**.

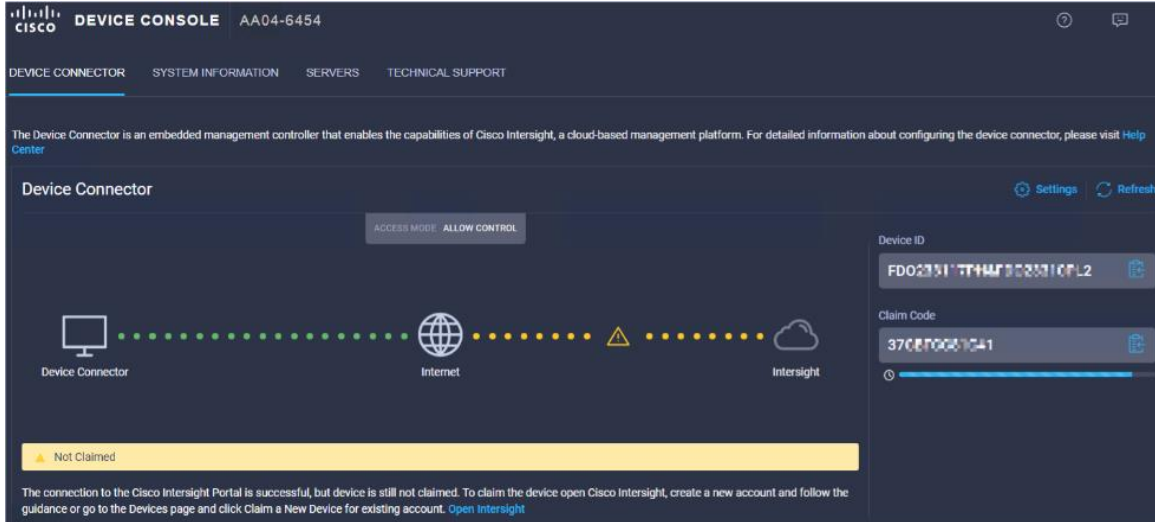


**Procedure 6. Claim Cisco UCS Fabric Interconnects in Cisco Intersight**

**Note:** Make sure the initial configuration for the fabric interconnects has been completed. Log into Fabric Interconnect A using a web browser to capture the Cisco Intersight connectivity information.

**Step 1.** Use the management IP address of Fabric Interconnect A to access the device from a web browser and the previously configured admin password to Log into the device.

**Step 2.** Under DEVICE CONNECTOR, the current device status will show “Not claimed.” Note, or copy, the Device ID, and Claim Code information for claiming the device in Cisco Intersight.

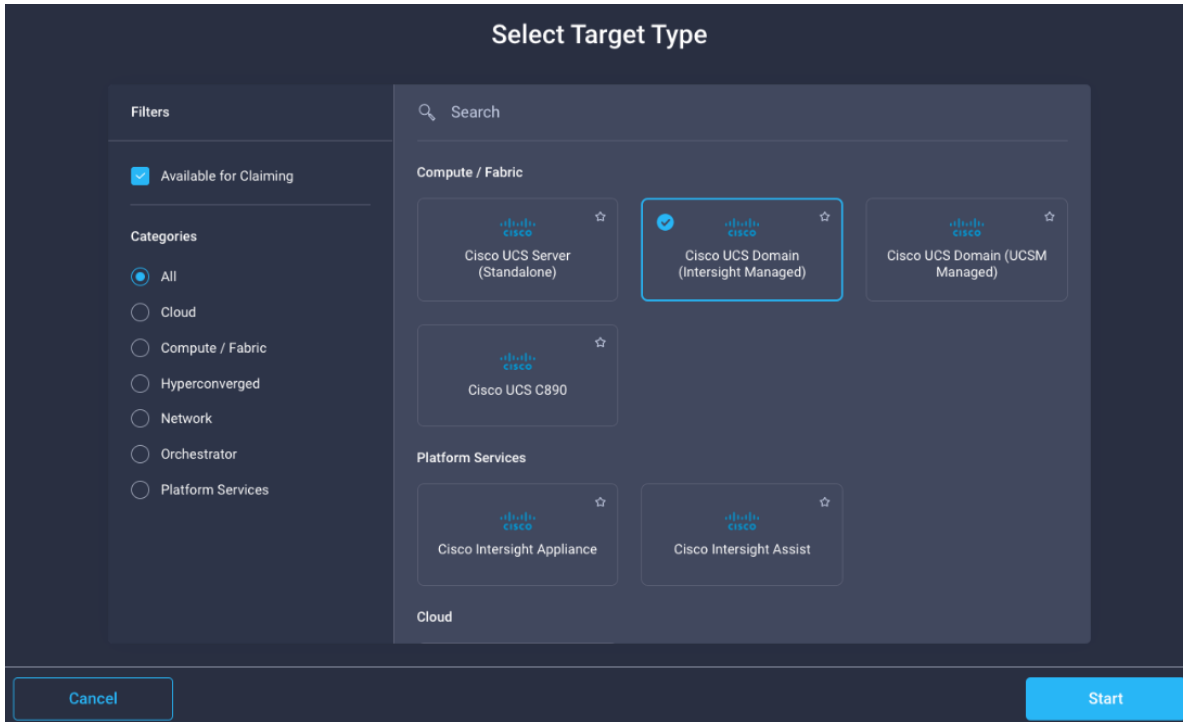


**Step 3.** Log into Cisco Intersight.

**Step 4.** Click **Targets** from the left menu.

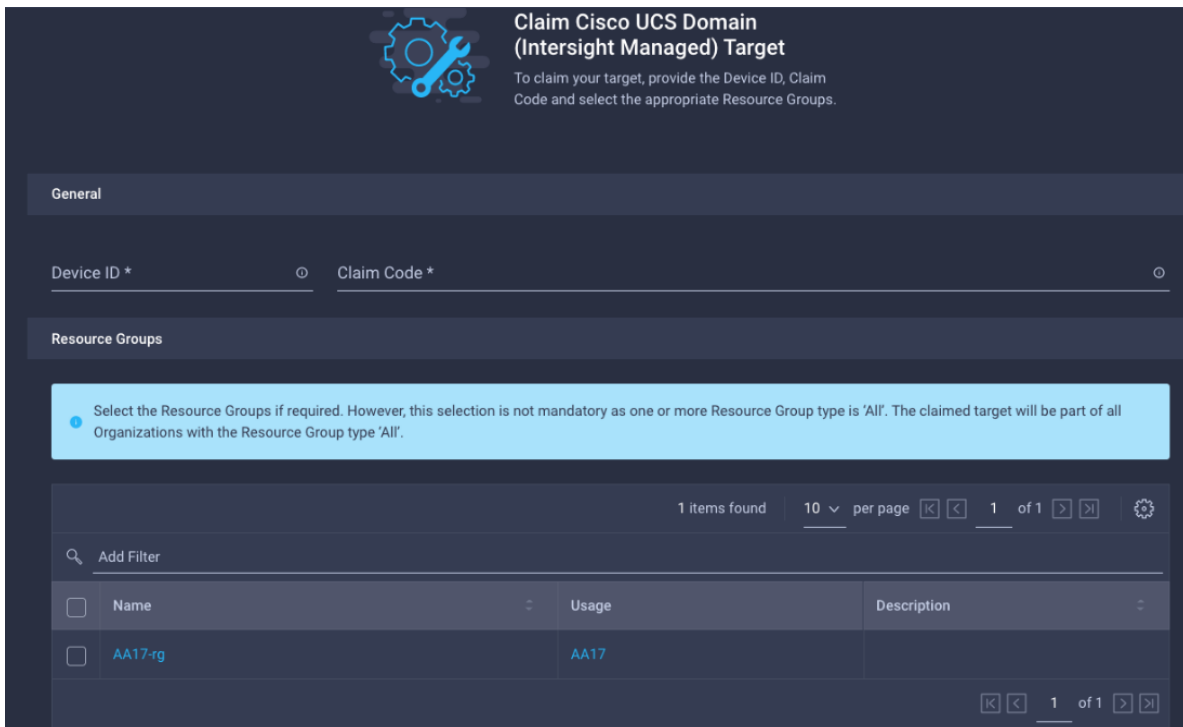
**Step 5.** Click **Claim New Target**.

**Step 6.** Select **Cisco UCS Domain (Intersight Managed)** and click **Start**.




**Step 7.** Enter the Device ID and Claim Code captured from the Cisco UCS FI.

**Step 8.** Select the previously created Resource Group and click **Claim**.



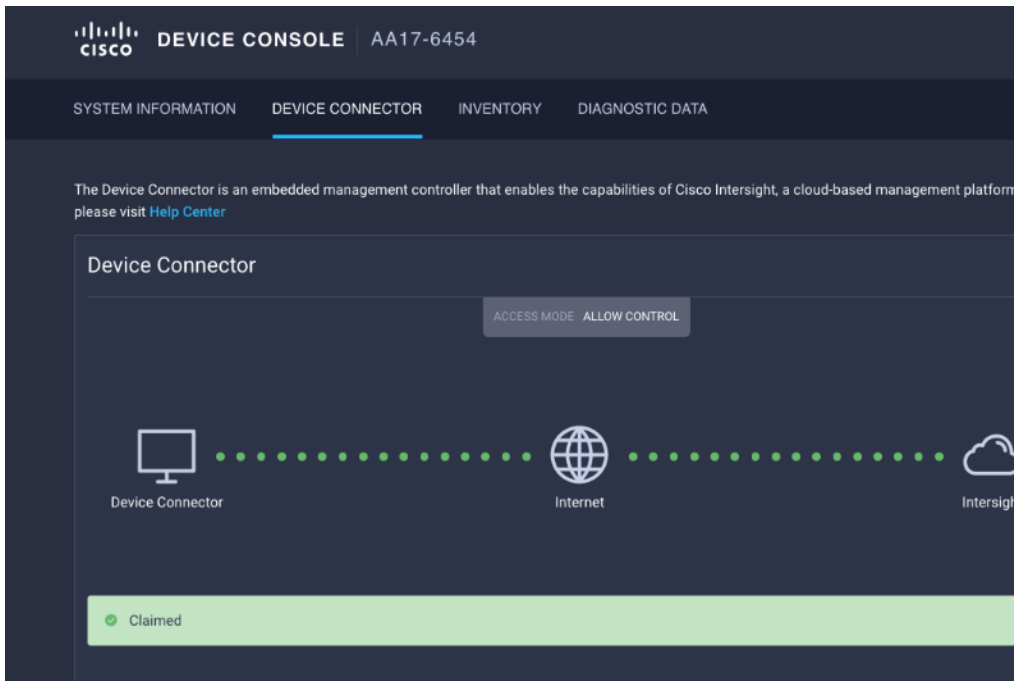
On successfully device claim, Cisco UCS FI should appear as a target in Cisco Intersight.



**Procedure 7. Verify the addition of Cisco UCS Fabric Interconnects to Cisco Intersight**

**Step 1.** Log back into the web GUI of the Cisco UCS fabric interconnect and click **Refresh**.

The fabric interconnect status should now be set to **Claimed**.



**Configure a Cisco UCS Domain Profile**

This subject contains the following procedures:

- [Create a Cisco UCS Domain Profile](#)
- [General Configuration](#)
- [Cisco UCS Domain Assignment](#)
- [Create and apply the VLAN Policy](#)
- [Create and apply VSAN policy \(FC configuration only\)](#)
- [Configure the Ports on the Fabric Interconnects](#)
- [Configure FC Port Channel \(FC configuration only\)](#)
- [Port Configuration for Fabric Interconnect B](#)

A Cisco UCS domain profile configures a fabric interconnect pair through reusable policies, allows configuration of the ports and port channels, and configures the VLANs and VSANs in the network. It defines the characteristics of and configured ports on fabric interconnects. The domain-related policies can be attached to the profile either at the time of creation or later. One Cisco UCS domain profile can be assigned to one fabric interconnect domain.

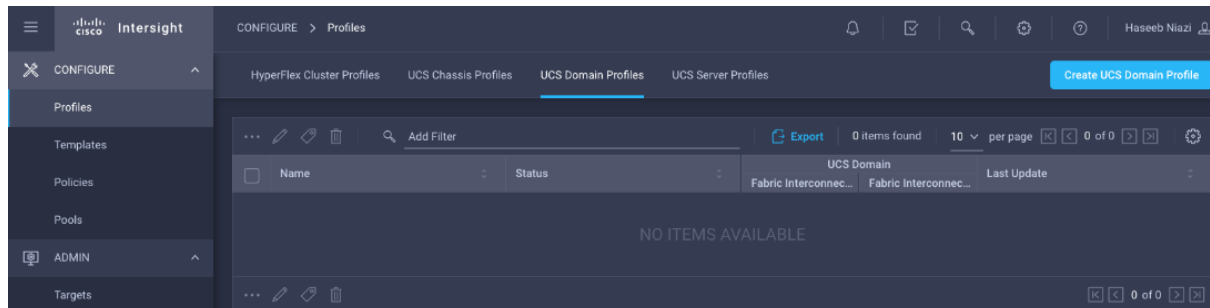


## Procedure 1. Create a Cisco UCS Domain Profile

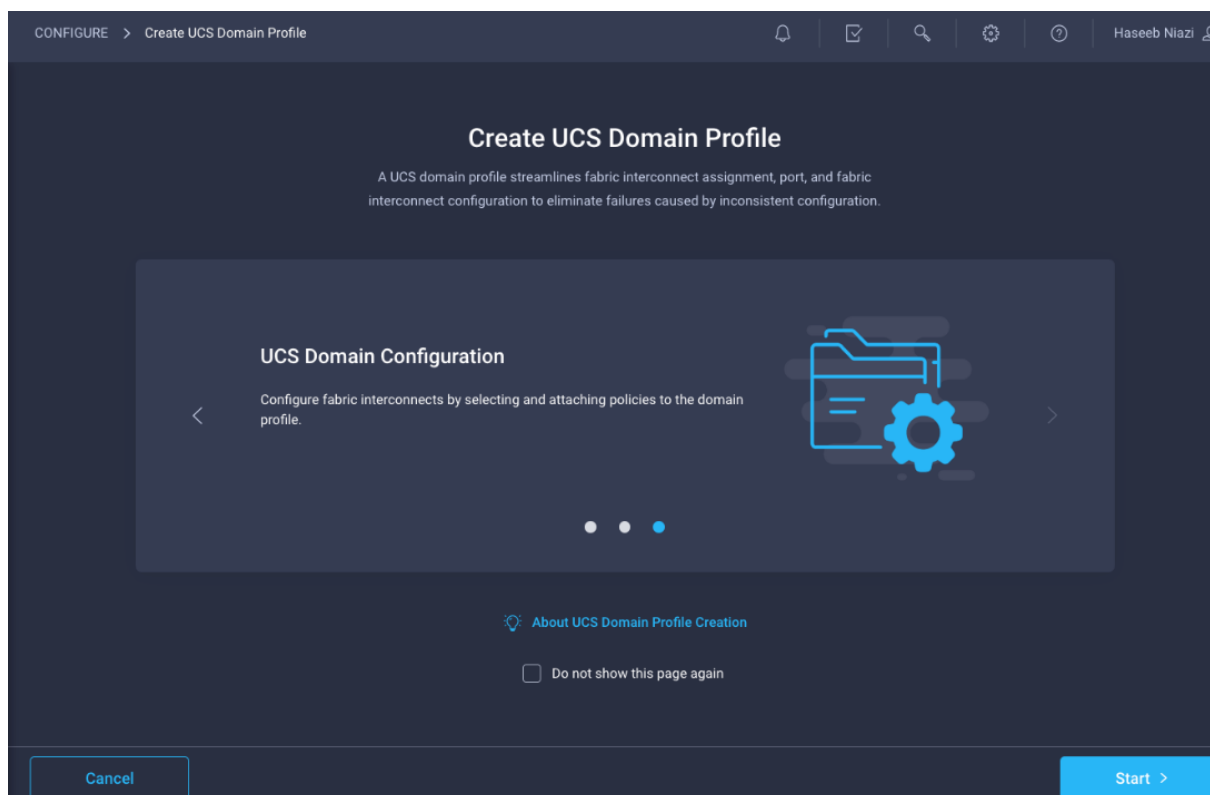
**Step 1.** Log into the **Cisco Intersight** portal.

**Step 2.** Click to expand **CONFIGURE** in the left pane and click **Profiles**.

**Step 3.** In the main window, click **UCS Domain Profiles** and click **Create UCS Domain Profile**.



**Step 4.** On the Create UCS Domain Profile screen, click **Start**.



## Procedure 2. General Configuration

**Step 1.** Select the organization from the drop-down list (for example, AA17).

**Step 2.** Provide a name for the domain profile (for example, Domain-Profile).

**Step 3.** Provide an optional Description.

Step 1  
**General**  
Add a name, description and tag for the UCS domain profile.

Organization \*  
AA17

Name \*  
AA17-Domain-Profile

Set Tags

Description

<= 1024

**Step 4.** Click **Next**.

### Procedure 3. Cisco UCS Domain Assignment

**Step 1.** Assign the Cisco UCS domain to this new domain profile by clicking **Assign Now** and selecting the previously added Cisco UCS domain (for example, 6454).

Progress

- 1 General
- 2 **UCS Domain Assignment**
- 3 VLAN & VSAN Configuration
- 4 Ports Configuration
- 5 UCS Domain Configuration
- 6 Summary

Step 2  
**UCS Domain Assignment**  
Choose to assign a fabric interconnect pair to the profile now or later.

**Assign Now** **Assign Later**

Choose to assign a fabric interconnect pair now or later. If you choose **Assign Now**, select a pair that you want to assign and click **Next**. If you choose **Assign Later**, click **Next** to proceed to policy selection.

Show Assigned

Add Filter

Domain Name	Fabric Interconnect A			Fabric Interconnect B		
	Model	Serial	Firmware Ver...	Model	Serial	Firmware Ver...
AA17-6454	UCS-FI-6454	FDO24350M...	9.3(5)42(1f)	UCS-FI-6454	FDO24350G32	9.3(5)42(1f)

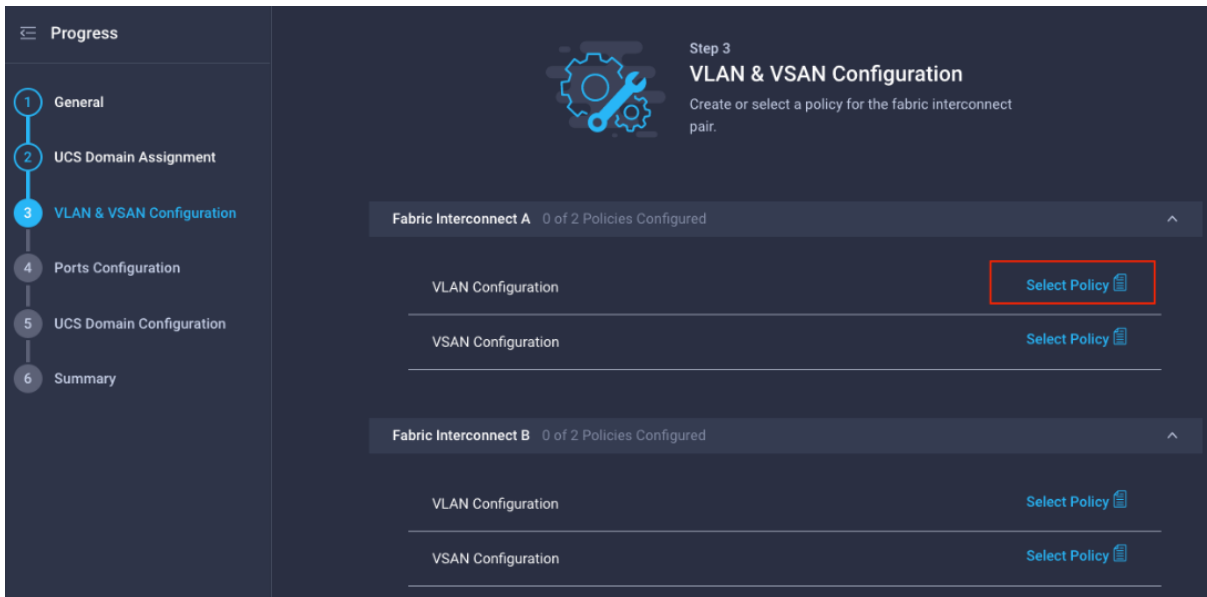
< Back Close **Next** >

**Step 2.** Click **Next**.

### Procedure 4. Create and apply the VLAN Policy

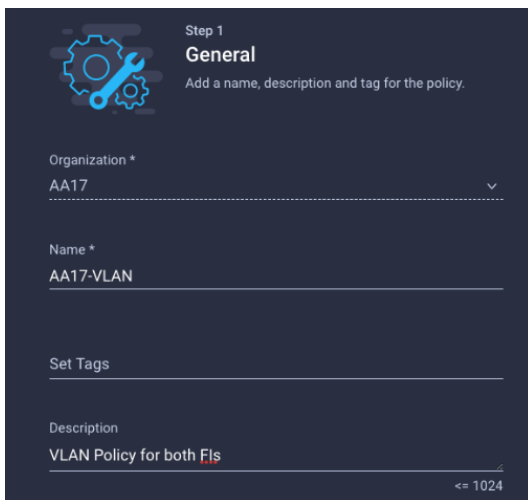
**Note:** In this step, a single VLAN policy is created for both fabric interconnects and two individual VSAN policies are created because the VSAN IDs are unique for each fabric interconnect.

**Step 1.** Click **Select Policy** next to VLAN Configuration under Fabric Interconnect A.



**Step 2.** In the pane on the right, click **Create New**.

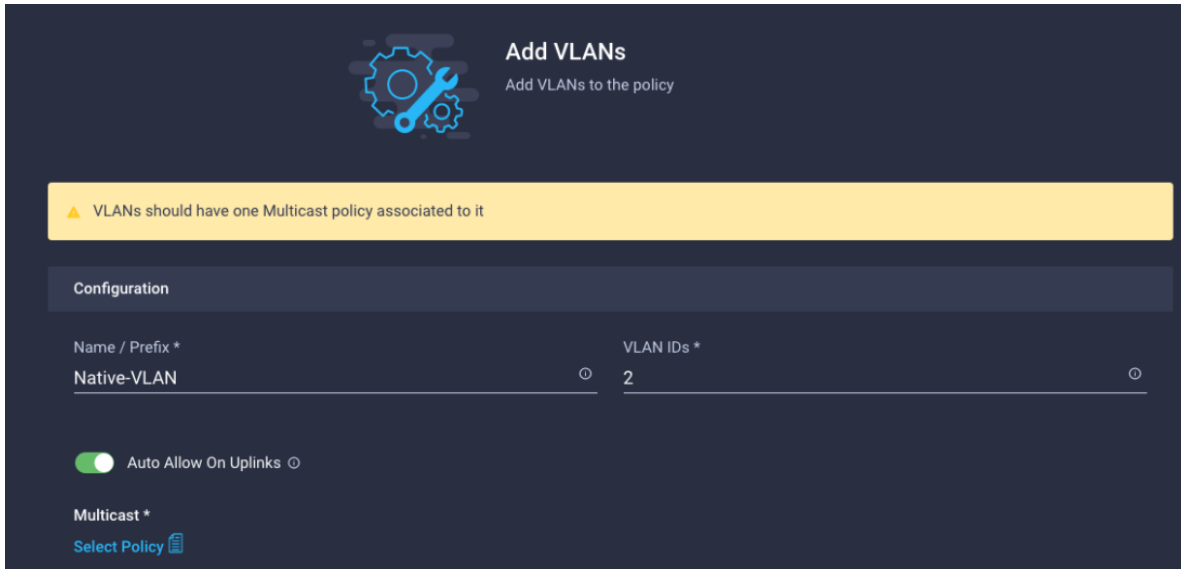
**Step 3.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, VLAN).



**Step 4.** Click **Next**.

**Step 5.** Click **Add VLANs**.

**Step 6.** Provide a name and VLAN ID for the native VLAN.



**Step 7.** Make sure **Auto Allow On Uplinks** is enabled.

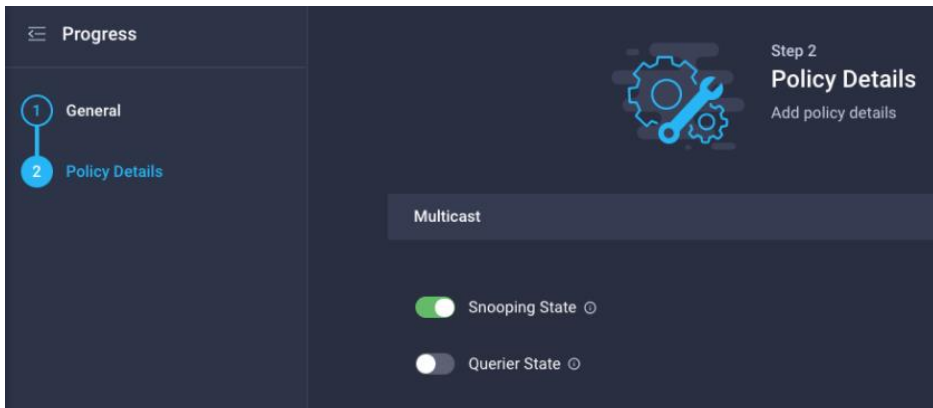
**Step 8.** To create the required Multicast policy, click **Select Policy** under Multicast\*.

**Step 9.** In the window on the right, click **Create New** to create a new Multicast Policy.

**Step 10.** Provide a Name for the Multicast Policy (for example, MCAST-Pol).

**Step 11.** Provide optional Description and click **Next**.

**Step 12.** Leave the Snooping State selected and click **Create**.



**Step 13.** Click **Add** to add the VLAN.

**Step 14.** Select **Set Native VLAN ID** and enter the VLAN number (for example, 2) under VLAN ID.

This policy applicable only for UCS Domain

VLANs

Add VLANs

Show VLAN Ranges

1 items found | 10 per page | 1 of 1

<input type="checkbox"/>	VLAN ID	Name
<input type="checkbox"/>	2	Native-Vlan_2

Selected 1 of 1 | Show Selected | Unselect All | 1 of 1

Set Native VLAN ID

VLAN ID  
2

**Step 15.** Add the remaining VLANs for FlexPod by clicking Add VLANs and entering the VLANs one by one. Reuse the previously created multicast policy for all the VLANs.

The VLANs created during this validation are shown below:

8 items found | 15 per page | 1 of 1

<input type="checkbox"/>	VLAN ID	Name	Multicast	Auto Allow On Uplinks
<input type="checkbox"/>	2	Native_2	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	17	AA17-IB-Mgmt_17	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	172	vm-traffic_172	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	3017	nfs_3017	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	3072	OOB-Mgmt_3072	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	3117	iscsi-a_3117	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	3217	iscsi-b_3217	AA17-MCAST-Pol	Yes
<input type="checkbox"/>	3317	vmotion_3317	AA17-MCAST-Pol	Yes

1 of 1

Set Native VLAN ID

VLAN ID  
2

**Note:** The iSCSI VLANs shown in the screen image above are only needed when iSCSI is configured in the environment.

**Step 16.** Click **Create** to finish creating the VLAN policy and associated VLANs.

**Step 17.** Click **Select Policy** next to VLAN Configuration for Fabric Interconnect B and select the same VLAN policy.

### Procedure 5. Create and apply VSAN policy (FC configuration only)

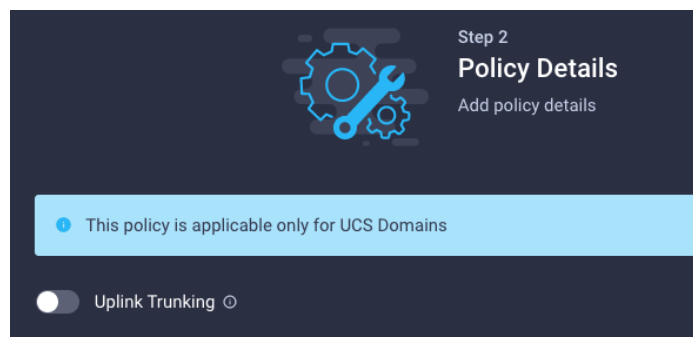
**Note:** A VSAN policy is only needed when configuring Fibre Channel and can be skipped when configuring IP-only storage access.

**Step 1.** Click **Select Policy** next to VSAN Configuration under Fabric Interconnect A. Click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, VSAN-Pol-A).

**Step 3.** Click **Next**.

**Step 4.** Enable **Uplink Trunking**.



**Step 5.** Click **Add VSAN** and provide a name (for example, VSAN-A), VSAN ID (for example, 400), and associated Fibre Channel over Ethernet (FCoE) VLAN ID (for example, 400) for SAN A.

**Step 6.** Set VLAN Scope as **Uplink**.

#### vdi-VSAN-FIA

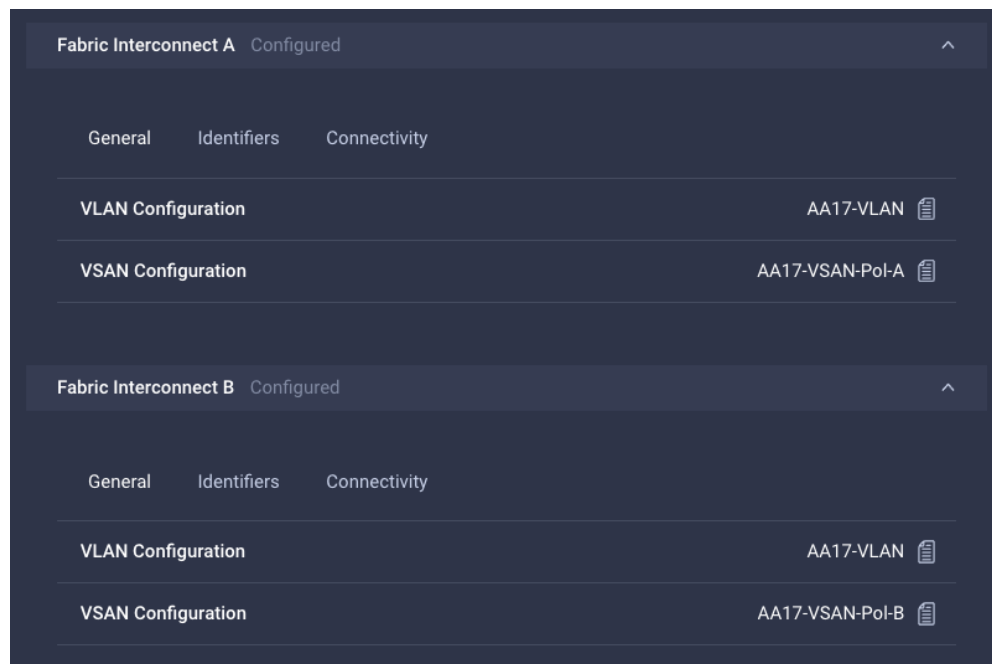
Details	Usage	Configuration												
<p>Name <b>vdi-VSAN-FIA</b></p> <p>Description -</p> <p>Type <b>VSAN</b></p> <p>Usage 1</p> <p>Last Update Aug 12, 2022 3:04 PM</p> <p>Organization <b>X-Series</b></p> <p>Tags <a href="#">Set</a></p>	<p>1 items found 4 per page 1 of 1</p> <p>Add Filter</p> <table border="1"><thead><tr><th>Name</th><th>Status</th><th>Platform Type</th><th>Type</th><th>Device Name</th><th>Last Up...</th></tr></thead><tbody><tr><td>vdi-dom-prfl-7</td><td>OK</td><td>UCS Domain</td><td>Profile</td><td>vdi-tme FI-A</td><td>Aug 12, 2022</td></tr></tbody></table>	Name	Status	Platform Type	Type	Device Name	Last Up...	vdi-dom-prfl-7	OK	UCS Domain	Profile	vdi-tme FI-A	Aug 12, 2022	<p>Uplink Trunking <b>Off</b></p> <p>VSAN ID 400</p> <p>Name <b>VSAN400</b></p> <p>VSAN Scope <b>Uplink</b></p> <p>FCoE VLAN ID <b>410</b></p>
Name	Status	Platform Type	Type	Device Name	Last Up...									
vdi-dom-prfl-7	OK	UCS Domain	Profile	vdi-tme FI-A	Aug 12, 2022									

**Step 7.** Click **Add**.

**Step 8.** Click **Create** to finish creating VSAN policy for fabric A.

**Step 9.** Repeat steps 1 – 8 to create a new VSAN policy for SAN-B. Name the policy to identify the SAN-B configuration (for example, VSAN-Pol-B) and use appropriate VSAN and FCoE VLAN (for example, 400).

**Step 10.** Verify that a common VLAN policy and two unique VSAN policies are associated with the two fabric interconnects.



**Step 11.** Click **Next**.

## Procedure 6. Configure the Ports on the Fabric Interconnects

**Step 1.** Click **Select Policy** for Fabric Interconnect A.

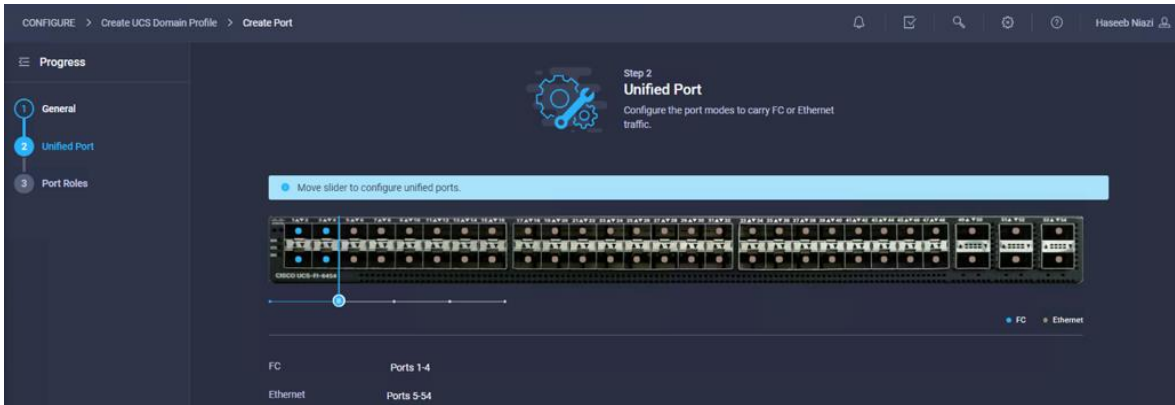
**Step 2.** Click **Create New** in the pane on the right to define a new port configuration policy.

**Note:** Use two separate port policies for the fabric interconnects. Using separate policies provide flexibility when port configuration (port numbers or speed) differs between the two FIs. When configuring Fibre Channel, two port policies are required because each fabric interconnect uses unique Fibre Channel VSAN ID.

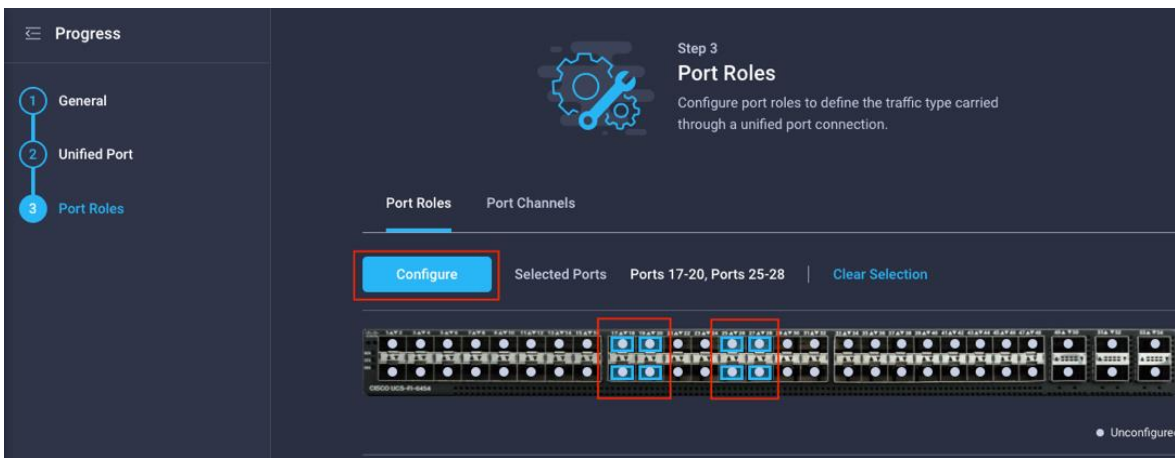
**Step 3.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, PortPol-A).

**Step 4.** Click **Next**.


**Step 5.** Move the slider to set up unified ports. In this deployment, the first four ports were selected as Fibre Channel ports. Click **Next**.



**Step 6.** Select the ports that need to be configured as server ports by clicking the ports in the graphics (or select from the list below the graphic). When all ports are selected, click **Configure**.

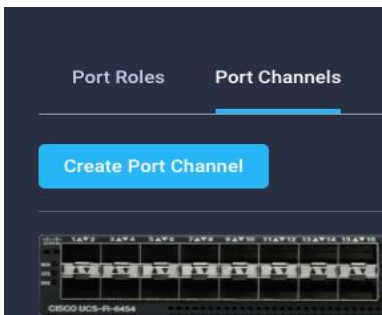


**Step 7.** From the drop-down list, select **Server** as the role.



**Step 8.** Click **Save**.

**Step 9.** Configure the Ethernet uplink port channel by selecting the Port Channel in the main pane and then clicking **Create Port Channel**.





**Step 10.** Select **Ethernet Uplink Port Channel** as the role, provide a port-channel ID (for example, 11), and select a value for Admin Speed from drop down menu (for example, Auto).

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

**Step 11.** Scroll down and select uplink ports from the list of available ports (for example, port 49 and 50)

**Step 12.** Click **Save**.

### Procedure 7. Configure FC Port Channel (FC configuration only)

**Note:** FC uplink port channel is only needed when configuring FC SAN and can be skipped for IP-only (iSCSI) storage access.

**Step 1.** Configure a Fibre Channel Port Channel by selecting the **Port Channel** in the main pane again and clicking **Create Port Channel**.

**Step 2.** In the drop-down list under Role, choose **FC Uplink Port Channel**.

**Step 3.** Provide a port-channel ID (for example, 1), select a value for Admin Speed (for example, 32Gbps), and provide a VSAN ID (for example, 400).

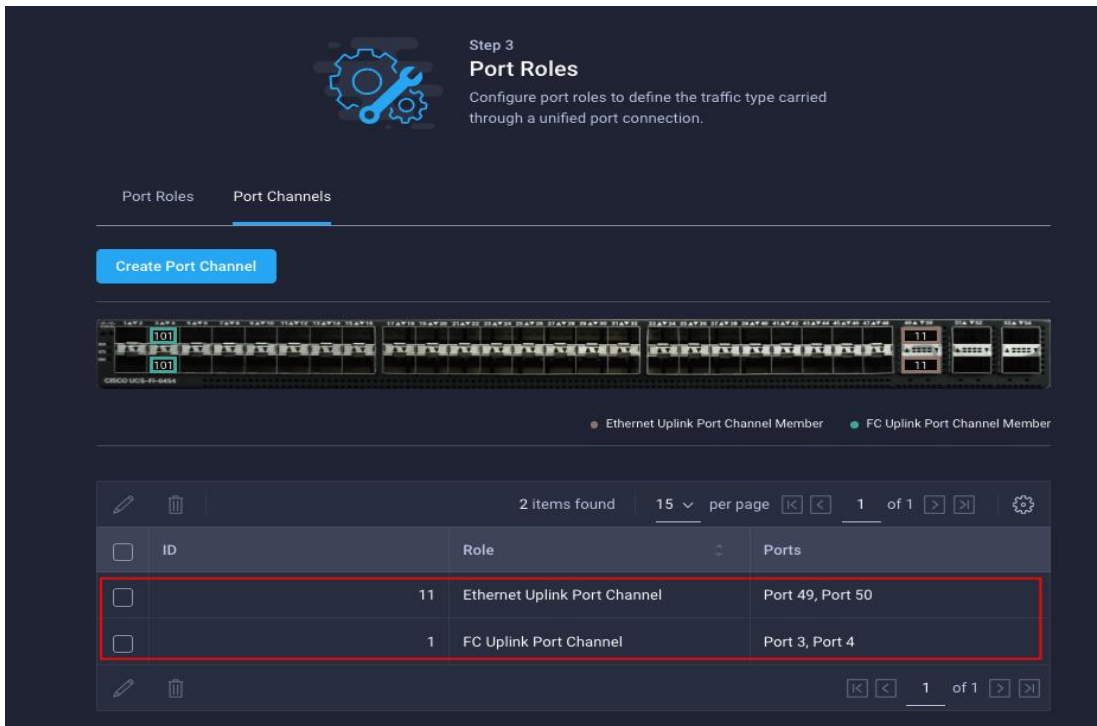
**vdi-FCN-A**

Details	Usage	Configuration
Name <b>vdi-FCN-A</b>	• Usage applies to Policies that are directly associated with a Profile or Template type.	Fibre Channel Network
Description -		Default VLAN <b>400</b>
Type <b>Fibre Channel Network</b>		VSAN ID <b>400</b>
Usage <b>N/A</b> ⓘ		
Last Update <b>Aug 12, 2022 3:04 PM</b>		
Organization <b>X-Series</b>		

**Step 4.** Select ports (for example, 3 and 4).

**Step 5.** Click **Save**.

**Step 6.** Verify the port-channel IDs and ports after both the Ethernet uplink port channel and the Fibre Channel uplink port channel have been created.



**Step 7.** Click **Save** to create the port policy for Fabric Interconnect A.

**Note:** Use the summary screen to verify that the ports were selected and configured correctly.

### Procedure 8. Port Configuration for Fabric Interconnect B

**Step 1.** Repeat the steps from [Procedure 7. Configure FC Port Channel \(FC configuration only\)](#) to create the port policy for Fabric Interconnect B including the Ethernet port-channel and the FC port-channel (if configuring SAN). Use the following values for various parameters:

- Name of the port policy: PortPol-B
- Ethernet port-Channel ID: 12
- FC port-channel ID: 2
- FC VSAN ID: 401

**Step 2.** When the port configuration for both fabric interconnects is complete and looks good, click **Next**.

## Cisco UCS Domain Configuration

This subject contains the following procedures:

- [Configure NTP Policy for the Cisco UCS Domain](#)
- [Configure Network Connectivity Policy](#)
- [Configure System QoS Policy](#)
- [Verify Settings](#)
- [Deploy the Cisco UCS Domain Profile](#)
- [Verify Cisco UCS Domain Profile Deployment](#)

**Note:** Under Cisco UCS domain configuration, additional policies can be configured to setup NTP, Syslog, DNS settings, SNMP, QoS and UCS operating mode (end host or switch mode). For this deployment, three policies (NTP, Network Connectivity and System QoS) will be configured.

### Procedure 1. Configure NTP Policy for the Cisco UCS Domain

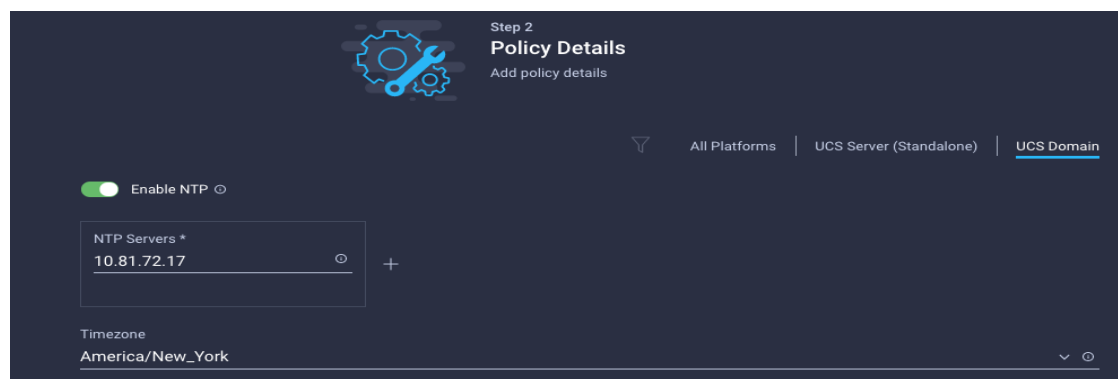
**Step 1.** Click **Select Policy** next to NTP and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, NTPPol).

**Step 3.** Click **Next**.

**Step 4.** **Enable NTP**, provide the NTP server IP addresses, and select the **Timezone** from the drop-down list.

**Step 5.** If required, add a second NTP server by clicking **+** next to the first NTP server IP address.



**Step 6.** Click **Create**.

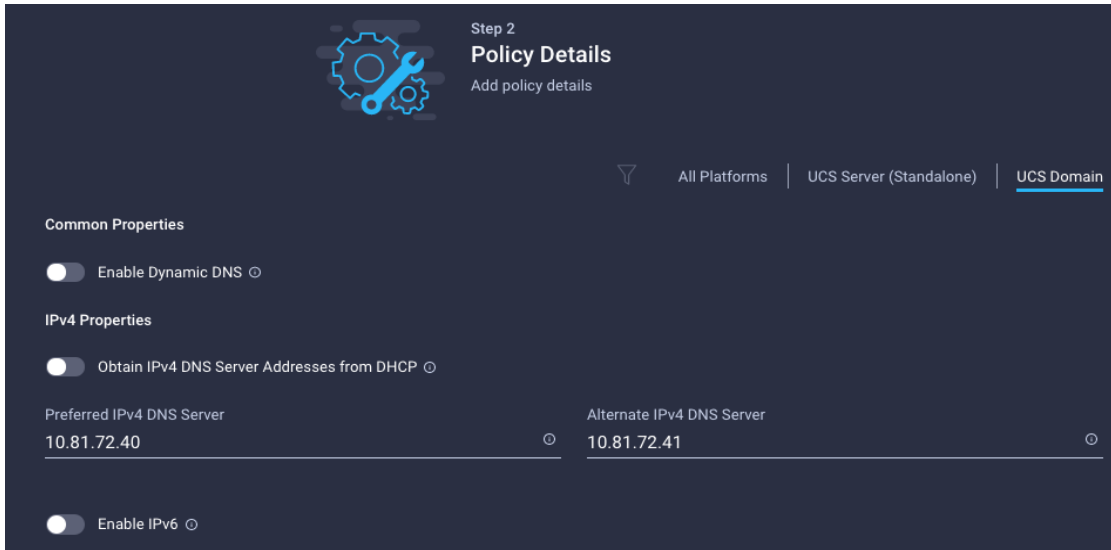
### Procedure 2. Configure Network Connectivity Policy

**Note:** To define the Domain Name Service (DNS) servers for Cisco UCS, configure the network connectivity policy.

**Step 1.** Click **Select Policy** next to Network Connectivity and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, NetConn-Pol).

**Step 3.** Provide DNS server IP addresses for Cisco UCS (for example, 10.81.72.40 and 10.81.72.41).



**Step 4.** Click **Create**.

### Procedure 3. Configure System QoS Policy

**Step 1.** Click **Select Policy** next to System QoS\* and click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, QoSPol).

**Step 3.** Click **Next**.

**Step 4.** Change the MTU for Best Effort class to **9216**.

**Step 5.** Keep the default selections or change the parameters if necessary.



**Step 6.** Click **Create**.

**Step 7.** Click **Next**.

#### Procedure 4. Verify Settings

**Step 1.** Verify all the settings including the fabric interconnect settings, by expanding the settings and make sure that the configuration is correct.

The screenshot shows the 'Step 6 Summary' view for a UCS domain profile. On the left, a 'Progress' sidebar lists steps 1 through 6, with 'Summary' selected. The main content area shows the 'General' tab for the profile 'AA17-Domain-Profile' under organization 'AA17'. Below this is a table of fabric interconnects:

Fabric Interconnect	Model	Serial	Requires Reboot
AA17-6454 FI-A	UCS-FI-6454	FD# ■■■■	No
AA17-6454 FI-B	UCS-FI-6454	FD# ■■■■	No

Below the table are tabs for 'Ports Configuration', 'VLAN & VSAN Configuration', 'UCS Domain Configuration', and 'Errors / Warnings'. Under 'Ports Configuration', there are expandable sections for 'Fabric Interconnect A' and 'Fabric Interconnect B'.

#### Procedure 5. Deploy the Cisco UCS Domain Profile

**Note:** After verifying the domain profile configuration, deploy the Cisco UCS profile.

**Step 1.** From the UCS domain profile Summary view, Click **Deploy**.

**Step 2.** Acknowledge any warnings and click **Deploy** again.

**Step 3.** The system will take some time to validate and configure the settings on the fabric interconnects. Log into the console servers to see when the Cisco UCS fabric interconnects have finished configuration and are successfully rebooted.

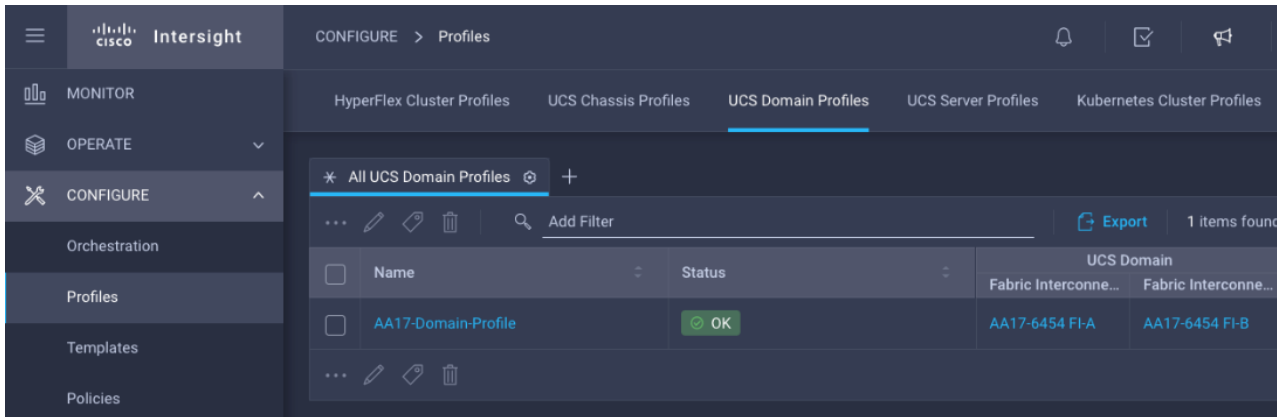
#### Procedure 6. Verify Cisco UCS Domain Profile Deployment

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

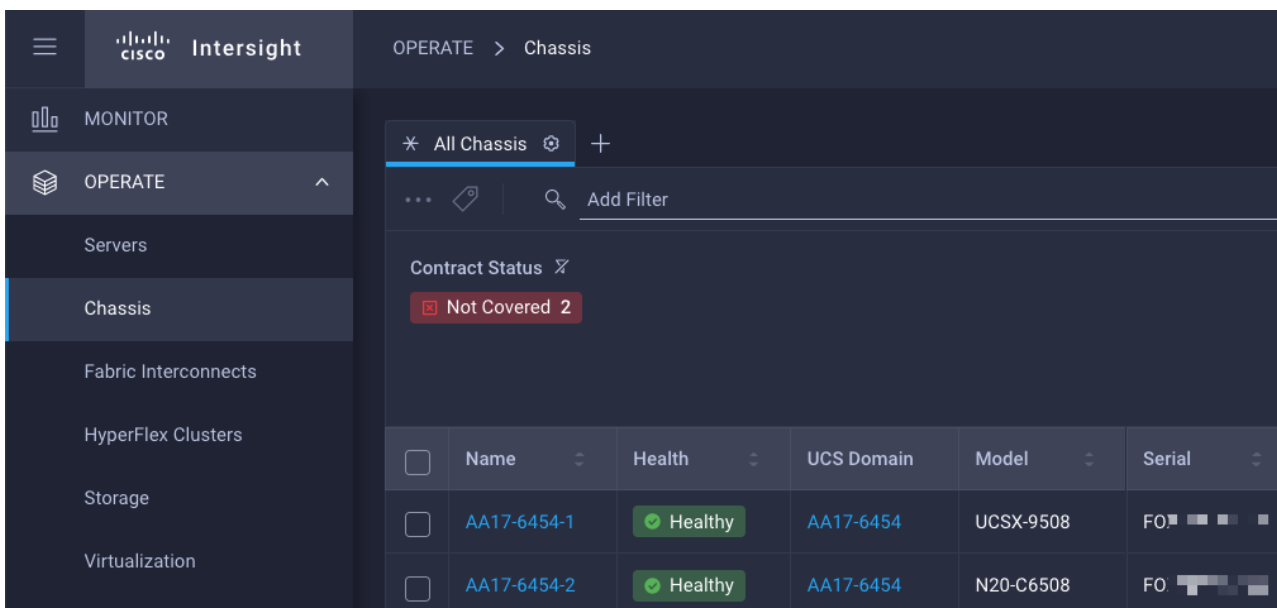
**Note:** It takes a while to discover the blades for the first time. Watch the number of outstanding tasks in Cisco Intersight:



**Step 1.** Log into **Cisco Intersight**. Under **CONFIGURE > Profiles > UCS Domain Profiles**, verify that the domain profile has been successfully deployed.



**Step 2.** Verify that the chassis has been discovered and is visible under **OPERATE > Chassis**.



**Step 3.** Verify that the servers have been successfully discovered and are visible under **OPERATE > Servers**.

<input type="checkbox"/>	AA17-6454-1-1	UCSX-210C-M6	140.8	5.0(1b)
<input type="checkbox"/>	AA17-6454-1-2	UCSX-210C-M6	140.8	5.0(1b)
<input type="checkbox"/>	AA17-6454-1-3	UCSX-210C-M6	140.8	5.0(1b)
<input type="checkbox"/>	AA17-6454-1-5	UCSX-210C-M6	166.4	5.0(1b)
<input type="checkbox"/>	AA17-6454-1-6	UCSX-210C-M6	166.4	5.0(1b)
<input type="checkbox"/>	AA17-6454-1-7	UCSX-210C-M6	166.4	5.0(1b)

## Configure Cisco UCS Chassis Profile

Cisco UCS Chassis profile in Cisco Intersight allows you to configure various parameters for the chassis, including:

- IMC Access Policy: IP configuration for the in-band chassis connectivity. This setting is independent of Server IP connectivity and only applies to communication to and from chassis.

- SNMP Policy, and SNMP trap settings.
- Power Policy to enable power management and power supply redundancy mode.
- Thermal Policy to control the speed of FANs (only applicable to Cisco UCS 5108)

A chassis policy can be assigned to any number of chassis profiles to provide a configuration baseline for a chassis. In this deployment, no chassis profile was created or attached to the chassis, but you can configure policies to configure SNMP or Power parameters and attach them to the chassis.

## Configure Server Profile Template

This subject contains the following procedures:

- [Configure a Server Profile Template](#)
- [Configure UUID Pool](#)
- [Configure BIOS Policy](#)
- [Configure Boot Order Policy for FC Hosts](#)

In the Cisco Intersight platform, a server profile enables resource management by simplifying policy alignment and server configuration. The server profiles are derived from a server profile template. The server profile template and its associated policies can be created using the server profile template wizard. After creating server profile template, you can derive multiple consistent server profiles from the template.

**Note:** The server profile captured in this deployment guide supports both Cisco UCS X-Series blade servers and Cisco UCS X210c M6 compute nodes.

### vNIC and vHBA Placement for Server Profile Template

In this deployment, separate server profile templates are created for iSCSI connected storage and for FC connected storage. The vNIC and vHBA layout is explained below. While most of the policies are common across various templates, the LAN connectivity and SAN connectivity policies are unique and will use the information in the tables below.

Four vNICs and two vHBAs are configured to support FC boot from SAN. These devices are manually placed as follows:

**Table 16. vHBA and vNIC placement for FC connected storage**

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

Four vNICs and four vHBAs are configured to support FC boot from SAN. Two vHBAs (vHBA-A and vHBA-B) are used for boot from SAN connectivity and the remaining two vHBAs are used to support NVMe-o-FC. These devices are manually placed as follows:

**Table 17. vHBA and vNIC placement for FC with NVMe-o-FC connected storage**

vNIC/vHBA Name	Slot	Switch ID	PCI Order	Comment
vHBA-A	MLOM	A	0	Used for boot from SAN
vHBA-B	MLOM	B	1	Used for boot from SAN
01-vSwitch0-A	MLOM	A	2	
02-vSwitch0-B	MLOM	B	3	
03-VDS0-A	MLOM	A	4	
04-VDS0-B	MLOM	B	5	
vHBA-NVMe-A	MLOM	A	6	Used for NVMe-o-FC
vHBA-NVMe-B	MLOM	B	7	Used for NVMe-o-FC

**Procedure 1. Configure a Server Profile Template**

**Step 1.** Log into **Cisco Intersight**.

**Step 2.** Go to **CONFIGURE > Templates** and in the main window click **Create UCS Server Profile Template**.

**Step 3.** Select the organization from the drop-down list (for example, AA17).

**Step 4.** Provide a name for the server profile template. The names used in this deployment are:

- FC-Boot-Template (FC boot from SAN)
- FC-Boot-NVME-Template (FC boot from SAN with support for NVMe-o-FC).

**Step 5.** Click **UCS Server (FI-Attached)**.

**Step 6.** Provide an optional description.



**Step 1**  
**General**  
Enter a name, description, tag and select a platform for the server profile.

Organization \*  
X-Series

Name \*  
FC-Boot-vdi

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

Set Tags

Description  
Supports FC

<= 1024

**Step 7.** Click **Next**.

## Procedure 2. Configure UUID Pool

**Step 1.** Click **Select Pool** under UUID Pool and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the UUID Pool (for example, UUID-Pool).

**Step 3.** Provide an optional Description and click **Next**.

**Step 4.** Provide a UUID Prefix (for example, a random prefix of 33FB3F9C-BF35-4BDE was used).

**Step 5.** Add a UUID block.

Configuration

Prefix \*  
33FB3F9C-BF35-4BDE

UUID Blocks

From *	Size *
0000-0000A1700001	64

1 - 1000

**Step 6.** Click **Create**.

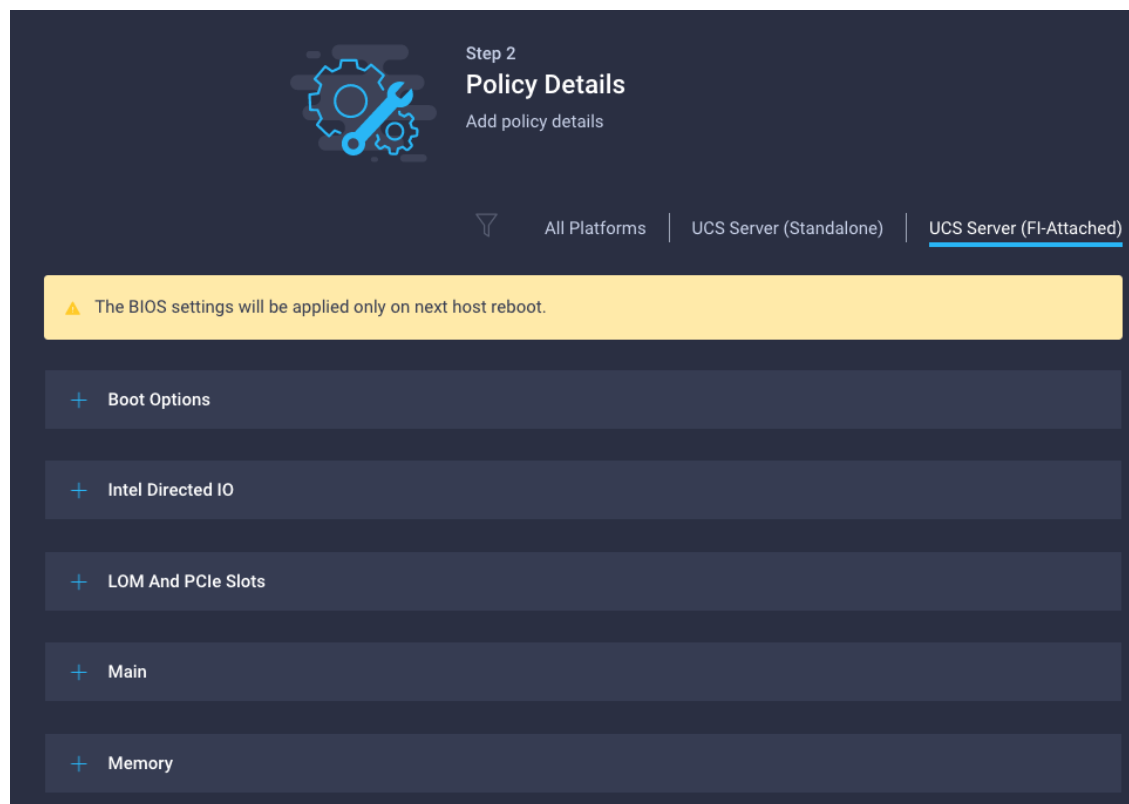
## Procedure 3. Configure BIOS Policy

**Step 1.** Click **Select Policy** next to BIOS and in the pane on the right, click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, NTPPol).

**Step 3.** Click **Next**.

**Step 4.** On the Policy Details screen, select appropriate values for the BIOS settings. In this deployment, the BIOS values were selected based on recommendations in the performance tuning guide for Cisco UCS M6 BIOS: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.html>.



- LOM and PCIe Slot > CDN Support for LOM: Enabled
- Processor > Enhanced CPU performance: Auto
- Memory > NVM Performance Setting: Balanced Profile

**Step 5.** Click **Create**.

#### **Procedure 4.** Configure Boot Order Policy for FC Hosts

**Note:** The FC boot order policy applies to all FC hosts including hosts that support NVMe-o-FC storage access.

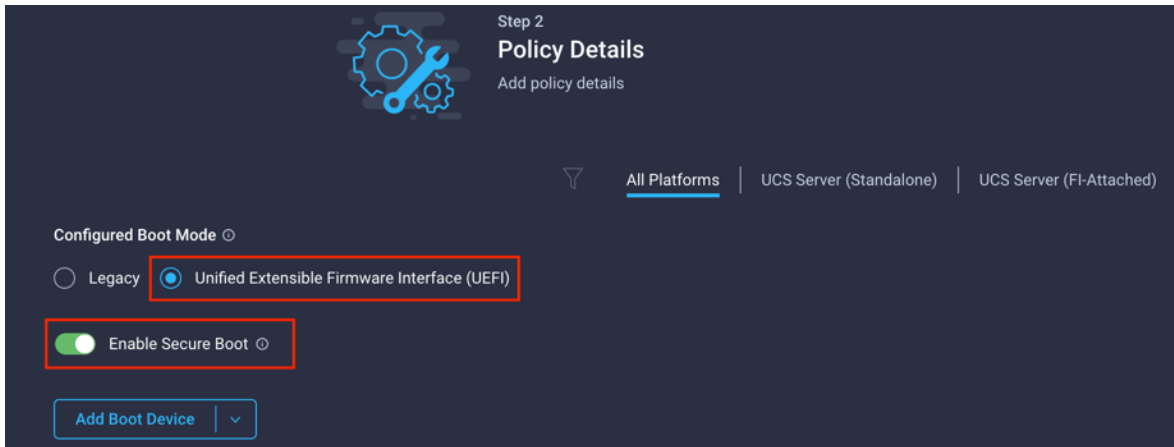
**Step 1.** Click **Select Policy** next to BIOS Configuration and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, FC-BootOrder-Pol).

**Step 3.** Click **Next**.


**Step 4.** For Configured Boot Mode, select **Unified Extensible Firmware Interface (UEFI)**.

**Step 5.** Turn on **Enable Secure Boot**.



**Step 6.** Click **Add Boot Device** drop-down list and select **Virtual Media**.

**Step 7.** Provide a device name (for example, ISO) and then, for the subtype, select **KVM Mapped DVD**.



For Fibre Channel SAN boot, all four NetApp controller LIFs will be added as boot options. The four LIFs are as follows:

- **FCP-LIF01a:** NetApp Controller 1, LIF for Fibre Channel SAN A
- **FCP-LIF01b:** NetApp Controller 1, LIF for Fibre Channel SAN B
- **FCP-LIF02a:** NetApp Controller 2, LIF for Fibre Channel SAN A
- **FCP-LIF02b:** NetApp Controller 2, LIF for Fibre Channel SAN B

**Step 8.** From the **Add Boot Device** drop-down list, select **SAN Boot**.




**Step 9.** Provide the Device Name: FCP-LIF01a and the Logical Unit Number (LUN) value (for example, 0).

**Step 10.** Provide an interface name vHBA-A. This value is important and should match the vHBA name.

**Note:** vHBA-A is used to access FCP-LIF01a and FCP-LIF02a and vHBA-B is used to access FCP-LIF01b and FCP-LIF02b.

**Step 11.** Add the appropriate World Wide Port Name (WWPN) as the Target WWPN.

**Note:** To obtain the WWPN values, log into NetApp controller using SSH and enter the following command: **network interface show -vserver Infra-FC -data-protocol fcp**.




SAN Boot (FCP-LIF01a) Enabled   

Device Name *	LUN
FCP-LIF01a	0
	0 - 255
Slot	Interface Name *
	vHBA-A
Target WWPN *	
20:01:d0:39:ea:29:ce:d4	
Bootloader Name	Bootloader Description

**Step 12.** Repeat steps 8-11 three more times to add all the NetApp LIFs.

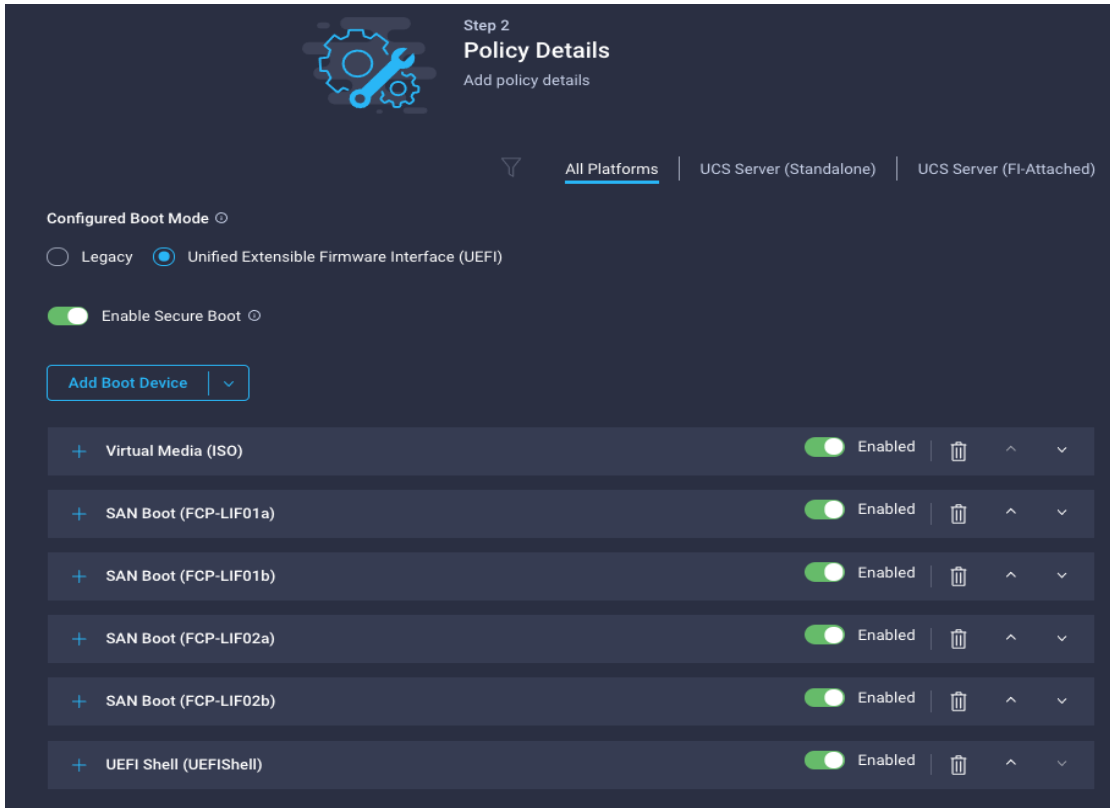
**Step 13.** From the **Add Boot Device** drop-down list, select **UEFI Shell**.

**Step 14.** Add Device Name **UEFIshell**.

UEFI Shell (UEFIshell) Enabled   

Device Name *
UEFIshell

**Step 15.** Verify the order of the boot policies and adjust the boot order as necessary using arrows next to delete button.



**Step 16.** Click **Create**.

**Step 17.** Click **Next** to move to Management Configuration.

## Management Configuration

This subject contains the following procedures:

- [Configure Cisco IMC Access Policy](#)
- [Configure IPMI Over LAN Policy](#)
- [Configure Local User Policy](#)
- [Storage Configuration](#)
- [Create LAN Connectivity Policy for FC Hosts](#)
- [Create MAC Address Pool for Fabric A and B](#)
- [Create Ethernet Network Group Policy for a vNIC](#)
- [Create Ethernet Network Control Policy](#)
- [Create Ethernet QoS Policy](#)
- [Create Ethernet Adapter Policy](#)
- [Create the SAN Connectivity Policy](#)
- [Create the WWNN Address Pool](#)
- [Create the vHBA-A for SAN A](#)
- [Create the WWPN Pool for SAN A](#)

- [Create Fibre Channel Network Policy for SAN A](#)
- [Create Fibre Channel QoS Policy](#)
- [Create Fibre Channel Adapter Policy](#)
- [Create the vHBA for SAN B](#)
- [Create the WWPN Pool for SAN B](#)
- [Create Fibre Channel Network Policy for SAN B](#)
- [Configure vHBA-NVMe-A and vHBA-NVMe-B](#)
- [Configure vHBA-NVMe-A](#)
- [Configure vHBA-NVMe-B](#)
- [Verify Summary](#)
- [Derive Server Profile](#)

Three policies will be added to the management configuration:

- IMC Access to define the pool of IP addresses for compute node KVM access
- IPMI Over LAN to allow Intersight to manage IPMI messages
- Local User to provide local administrator to access KVM

#### Procedure 1. Configure Cisco IMC Access Policy

**Step 1.** Click **Select Policy** next to IMC Access and then click **Create New**.

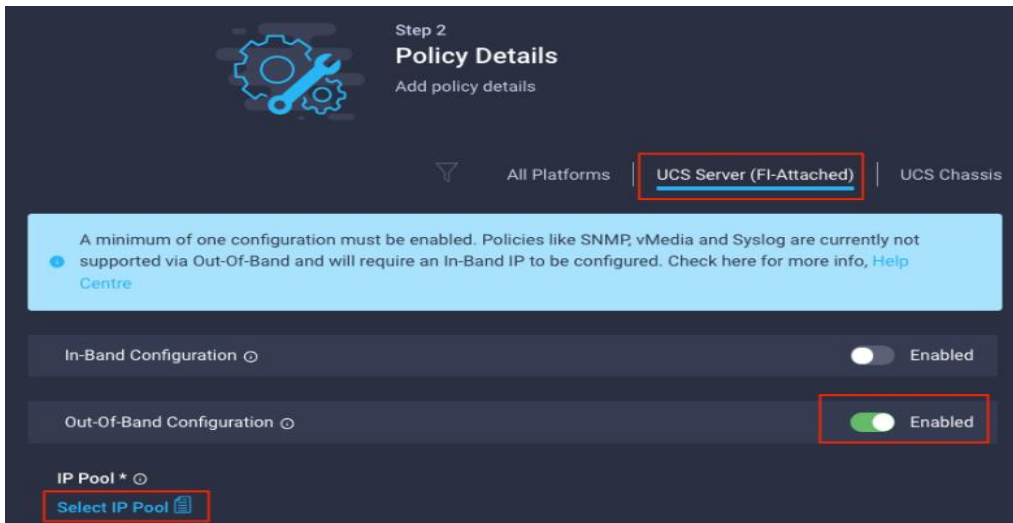
**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, IMC-Access).

**Step 3.** Click **Next**.

**Note:** You can select in-band management access to the compute node using an in-band management VLAN (for example, VLAN 17) or out-of-band management access via the Mgmt0 interfaces of the FIs. KVM Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Since these policies were not configured in this deployment, out-of-band management access was configured so that KVM access to compute nodes is not impacted by any potential switching issues in the fabric.

**Step 4.** Click **UCS Server (FI-Attached)**.

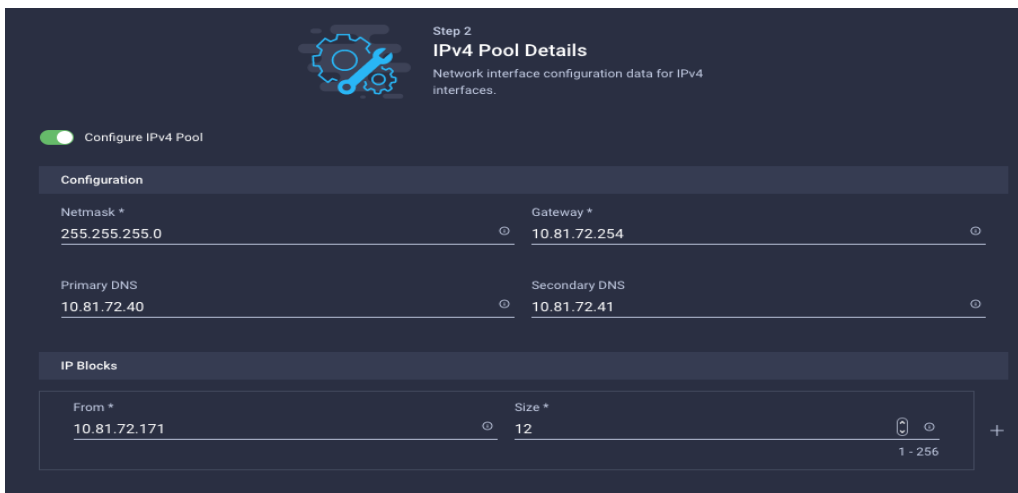
**Step 5.** **Enable** Out-Of-Band Configuration.



**Step 6.** Under IP Pool, click **Select IP Pool** and then click **Create New**.

**Step 7.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, IMC-OOB-Pool).

**Step 8.** Select **Configure IPv4 Pool** and provide the information to define a pool for KVM IP address assignment including an IP Block.



**Note:** The management IP pool subnet should be accessible from the host that is trying to open the KVM connection. In the example shown here, the hosts trying to open a KVM connection would need to be able to route to 10.81.72.0/24 subnet.

**Step 9.** Click **Next**.

**Step 10.** Deselect **Configure IPv6 Pool**.

**Step 11.** Click **Create** to finish configuring the IP address pool.

**Step 12.** Click **Create** to finish configuring the IMC access policy.

## Procedure 2. Configure IPMI Over LAN Policy

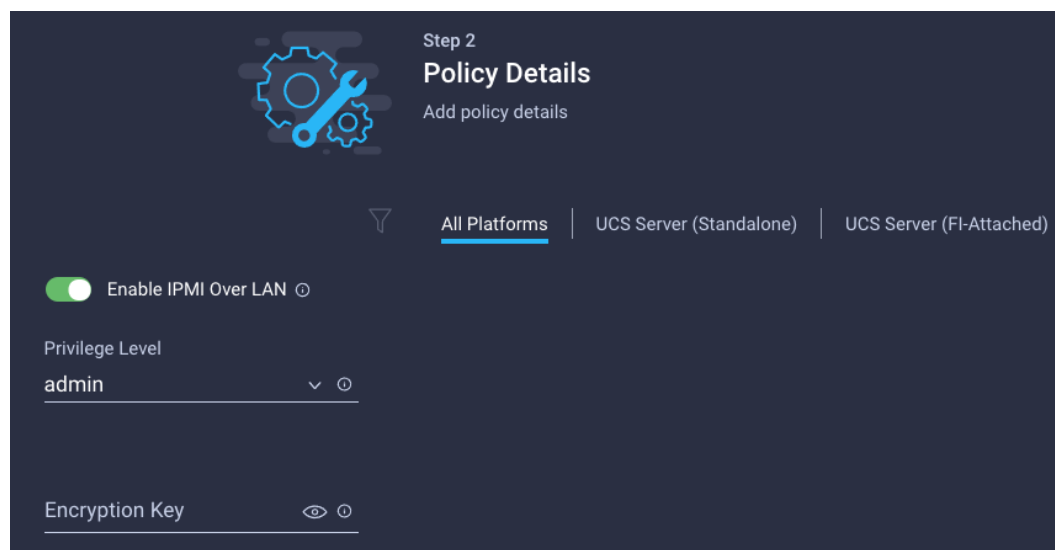
**Step 1.** Click **Select Policy** next to IPMI Over LAN and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, Enable-IPMIoLAN).

**Step 3.** Turn on **Enable IPMI Over LAN**.

**Step 4.** From the **Privilege Level** drop-down list, select **admin**.

**Step 5.** Click **Create**.



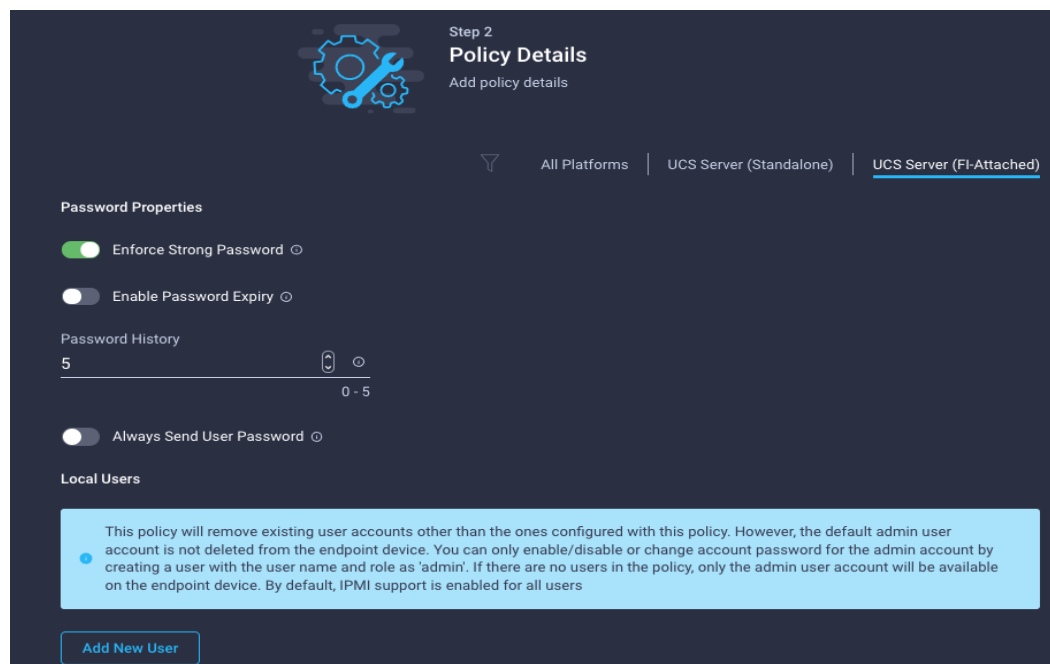
### Procedure 3. Configure Local User Policy

**Step 1.** Click **Select Policy** next to Local User and then, in the pane on the right, click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, LocalUser-Pol).

**Step 3.** Verify that **UCS Server (FI-Attached)** is selected.

**Step 4.** Verify that **Enforce Strong Password** is selected.



**Step 5.** Click **Add New User** and then click **+** next to the New User



**Step 6.** Provide the username (for example, fpadmin), choose a role for example, admin), and provide a password.

The screenshot shows a user configuration interface. At the top, there is a button labeled 'Add New User'. Below it, a user entry for 'fpadmin (admin)' is shown with a green checkmark and an 'Enable' toggle switch. The configuration fields are as follows:

- Username \***: fpadmin
- Role**: admin
- Password \***: [masked]
- Password Confirmation \***: [masked]

**Note:** The username and password combination defined here will be used to log into KVMs. The typical Cisco UCS admin username and password combination cannot be used for KVM access.

**Step 7.** Click **Create** to finish configuring the user.

**Step 8.** Click **Create** to finish configuring local user policy.

**Step 9.** Click **Next** to move to Storage Configuration.

#### Procedure 4. Storage Configuration

**Step 1.** Click **Next** on the Storage Configuration screen. No configuration is needed in the local storage system.

#### Step 2. Network Configuration > LAN Connectivity

LAN connectivity policy defines the connections and network communication resources between the server and the LAN. This policy uses pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For iSCSI hosts, this policy also defined an IQN address pool.

For consistent vNIC and vHBA placement, manual vHBA/vNIC placement is utilized. iSCSI boot from SAN hosts and FC boot from SAN hosts require different number of vNICs/vHBAs and different placement order therefore the iSCSI host and the FC host LAN connectivity policies are explained separately in this section.

#### Procedure 5. Create LAN Connectivity Policy for FC Hosts

The FC boot from SAN hosts uses 4 vNICs configured as follows:

**Table 18. vNICs for FC LAN Connectivity**

vNIC/vHBA Name	Slot ID	Switch ID	PCI Order	VLANs
01-vSwitch0-A	MLOM	A	2	IB-MGMT, NFS
02-vSwitch0-B	MLOM	B	3	IB-MGMT, NFS
03-VDS0-A	MLOM	A	4	VM Traffic, vMotion

vNIC/vHBA Name	Slot ID	Switch ID	PCI Order	VLANs
04-VDS0-B	MLOM	B	5	VM Traffic, vMotion

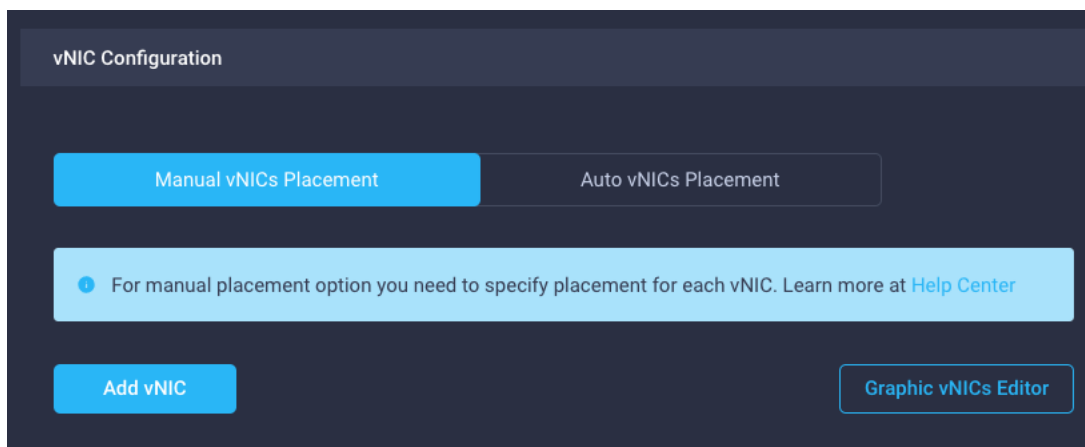
**Note:** The PCI order 0 and 1 will be used in the SAN Connectivity policy to create vHBA-A and vHBA-B.

**Step 1.** Click **Select Policy** next to LAN Connectivity and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, FC-ESXi-LanConn). Click **Next**.

**Step 3.** Under vNIC Configuration, select **Manual vNICs Placement**.

**Step 4.** Click **Add vNIC**.



### Procedure 6. Create MAC Address Pool for Fabric A and B

When creating the first vNIC, the MAC address pool has not been defined yet, therefore a new MAC address pool will need to be created. Two separate MAC address pools are configured for each Fabric. MAC-Pool-A will be reused for all Fabric-A vNICs, and MAC-Pool-B will be reused for all Fabric-B vNICs.

**Table 19. MAC Address Pools**

Pool Name	Starting MAC Address	Size	vNICs
MAC-Pool-A	00:25:B5:17:0A:00	64*	01-vSwitch0-A, 03-VDS0-A
MAC-Pool-B	00:25:B5:17:0B:00	64*	02-vSwitch0-B, 04-VDS0-B

**Note:** Each server requires 2 MAC addresses from the pool. Adjust the size of the pool according to your requirements.

**Step 1.** Click **Select Pool** under MAC Address Pool and then click **Create New**.

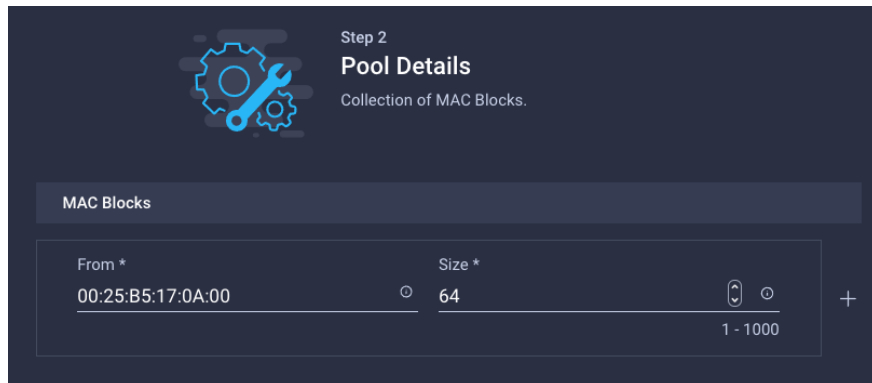
**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the pool from Table 19 depending on the vNIC being created (for example, MAC-Pool-A for Fabric A).

**Step 3.** Click **Next**.

**Step 4.** Provide the starting MAC address from [Table 19](#) (for example, 00:25:B5:17:0A:00)

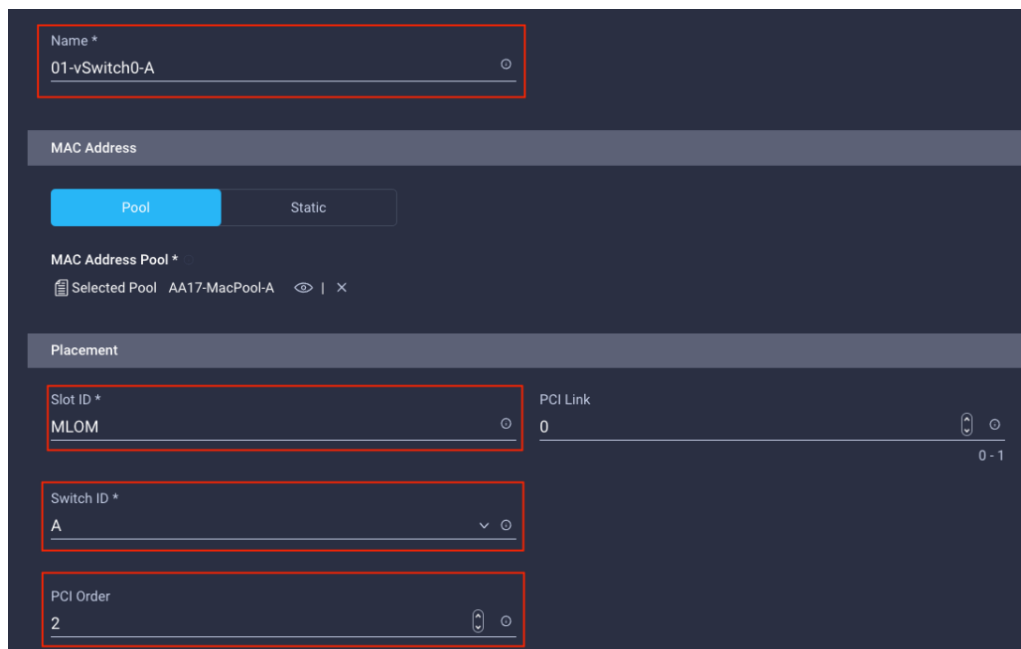
**Note:** For troubleshooting FlexPod, some additional information is always coded into the MAC address pool. For example, in the starting address 00:25:B5:17:0A:00, 17 is the rack ID and 0A indicates Fabric A.

**Step 5.** Provide the size of the MAC address pool from [Table 19](#) (for example, 64).




**Step 6.** Click **Create** to finish creating the MAC address pool.

**Step 7.** From the Add vNIC window, provide vNIC Name, Slot ID, Switch ID, and PCI Order information from [Table 18](#).



**Step 8.** For Consistent Device Naming (CDN), from the drop-down list, select **vNIC Name**.

**Step 9.** Verify that Failover is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and VDS.



**Procedure 7. Create Ethernet Network Group Policy for a vNIC**

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

**Table 20. Ethernet Group Policy Values**

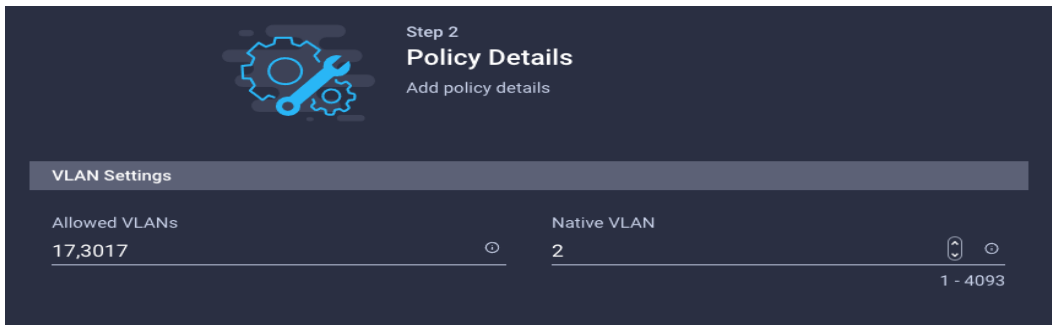
Group Policy Name	Native VLAN	Apply to vNICs	VLANs
vSwitch0-NetGrp	Native-VLAN (2)	01-vSwitch0-A, 02-vSwitch0-B	IB-MGMT, NFS
VDS0-NetGrp	Native-VLAN (2)	03-VDS0-A, 04-VDS0-B	VM Traffic, vMotion

**Step 10.** Click **Select Policy** under Ethernet Network Group Policy and then click **Create New**.

**Step 11.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy from [Table 20](#) (for example, vSwitch0-NetGrp).

**Step 12.** Click **Next**.

**Step 13.** Enter the allowed VLANs from [Table 20](#) (for example, 17,3017) and the native VLAN ID from [Table 20](#) (for example, 2).



**Step 14.** Click **Create** to finish configuring the Ethernet network group policy.

**Note:** When ethernet group policies are shared between two vNICs, the ethernet group policy only needs to be defined for the first vNIC. For subsequent vNIC policy mapping, just click **Select Policy** and pick the previously defined ethernet group policy from the list on the right.

**Procedure 8. Create Ethernet Network Control Policy**

The Ethernet Network Control Policy is used to enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) for the vNICs. A single policy will be created and reused for all the vNICs.

**Step 1.** Click **Select Policy** under Ethernet Network Control Policy and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, Enable-CDP-LLDP).

**Step 3.** Click **Next**.

**Step 4.** **Enable Cisco Discovery Protocol** and both **Enable Transmit** and **Enable Receive** under LLDP.

The screenshot displays a configuration page for a network policy. At the top, a light blue banner states: "This policy is applicable only for UCS Servers (FI-Attached)". Below this, the "Enable CDP" toggle is turned on and highlighted with a red box. Under "Mac Register Mode", "Only Native VLAN" is selected. Under "Action on Uplink Fail", "Link Down" is selected. A yellow warning box contains the text: "Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost." The "MAC Security" section shows "Forge" set to "Allow". The "LLDP" section at the bottom has "Enable Transmit" and "Enable Receive" toggles both turned on, with a red box highlighting both.

**Step 5.** Click **Create** to finish creating Ethernet network control policy.

## Procedure 9. Create Ethernet QoS Policy

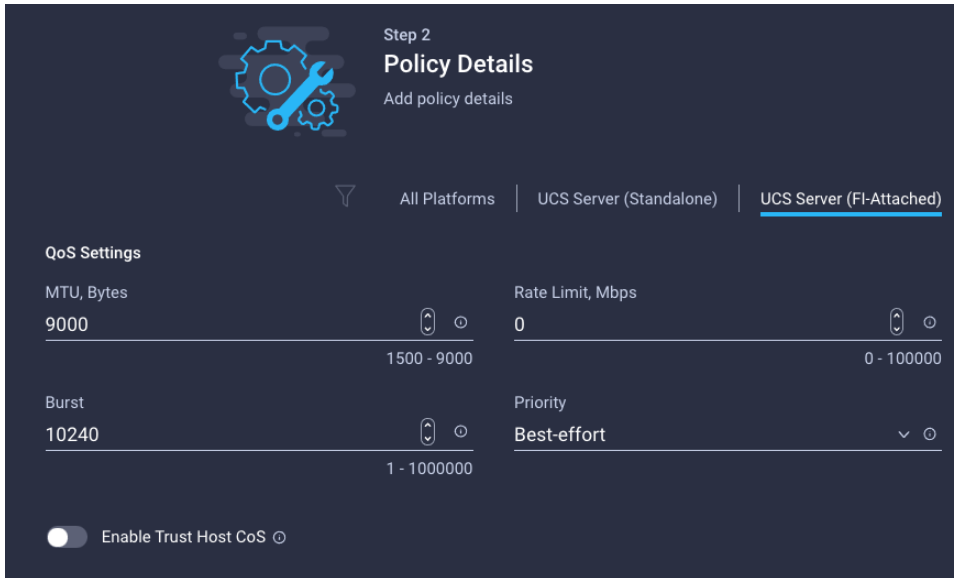
Ethernet QoS policy is used to enable jumbo maximum transmission units (MTUs) for all the vNICs. A single policy will be created and reused for all the vNICs.

**Step 1.** Click **Select Policy** under Ethernet QoS and click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, EthQos-Pol).

**Step 3.** Click **Next**.

**Step 4.** Change the MTU, Bytes value to **9000**.



**Step 5.** Click **Create** to finish setting up the Ethernet QoS policy.

## Procedure 10. Create Ethernet Adapter Policy

Ethernet adapter policy is used to set the interrupts and the send and receive queues. The values are set according to the best-practices guidance for the operating system in use. Cisco Intersight provides default VMware Ethernet Adapter policy for typical VMware deployments.

Optionally, you can configure a tweaked ethernet adapter policy for additional hardware receive queues handled by multiple CPUs in scenarios where there is a lot of vMotion traffic and multiple flows. In this deployment, a modified ethernet adapter policy, VMware-High-Traffic, is created and attached to the 03-VDS0-A and 04-VDS0-B interfaces which handle vMotion.

**Table 21. Ethernet Adapter Policy association to vNICs**

Policy Name	vNICs
EthAdapter-VMware	01-vSwitch0-A, 02-vSwitch0-B
VMware-High-Traffic	03-VDS0-A, 04-VDS0-B,

**Step 1.** Click **Select Policy** under Ethernet Adapter and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, EthAdapter-VMware).

**Step 3.** Click **Select Default Configuration** under Ethernet Adapter Default Configuration.

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
AA17

Name \*  
AA17-EthAdapter-VMware

Set Tags

Description  
<= 1024

Ethernet Adapter Default Configuration \*

Select Default Configuration

**Step 4.** From the list, select **VMware**.

**Step 5.** Click **Next**.

**Step 6.** For the EthAdapter-VMware policy, click **Create** and skip the rest of the steps in this section.

**Step 7.** For the optional VMware-High-Traffic policy (for VDS interfaces), make the following modifications to the policy:

- Increase Interrupts to 11
- Increase Receive Queue Count to 8
- Increase Completion Queue Count to 9
- Enable Receive Side Scaling

**Interrupt Settings**

Interrupts: 11 (1 - 1024)

Interrupt Mode: MSIx

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type: Min

**Receive**

Receive Queue Count: 8 (1 - 1000)

Receive Ring Size: 512 (64 - 4096)

**Transmit**

Transmit Queue Count: 1 (1 - 1000)

Transmit Ring Size: 256 (64 - 4096)

**Completion**

Completion Queue Count: 9 (1 - 2000)

Completion Ring Size: 1 (1 - 256)

Uplink Failback Timeout (seconds): 5 (0 - 600)

**TCP Offload**

- Enable Tx Checksum Offload
- Enable Rx Checksum Offload
- Enable Large Send Offload
- Enable Large Receive Offload

**Receive Side Scaling**

- Enable Receive Side Scaling

**Step 8.** Click **Create**.

**Step 9.** Click **Create** to finish creating the vNIC.

**Step 10.** Go back to [Step 1](#) and repeat vNIC creation for all four vNICs.

**Step 11.** Verify all four vNICs were successfully created.



<input type="checkbox"/>	Name	Slot ID	Switch ID	PCI Link	PCI Order	Failover	
<input type="checkbox"/>	01-vSwitch0-A	MLOM	A	0	2	Disabled	...
<input type="checkbox"/>	03-VDS0-A	MLOM	A	0	4	Disabled	...
<input type="checkbox"/>	02-vSwitch0-B	MLOM	B	0	3	Disabled	...
<input type="checkbox"/>	04-VDS0-B	MLOM	B	0	5	Disabled	...

**Step 12.** Click **Create** to finish creating the LAN Connectivity policy for FC hosts.

### Procedure 11. Create the SAN Connectivity Policy

A SAN connectivity policy determines the network storage resources and the connections between the server and the storage device on the network. This policy enables customers to configure the vHBAs that the servers use to communicate with the SAN.

**Note:** SAN Connectivity policy is not needed for iSCSI boot from SAN hosts and can be skipped.

[Table 22](#) lists the details of two vHBAs that are used to provide FC connectivity and boot from SAN functionality.

**Table 22. vHBA for boot from FC SAN**

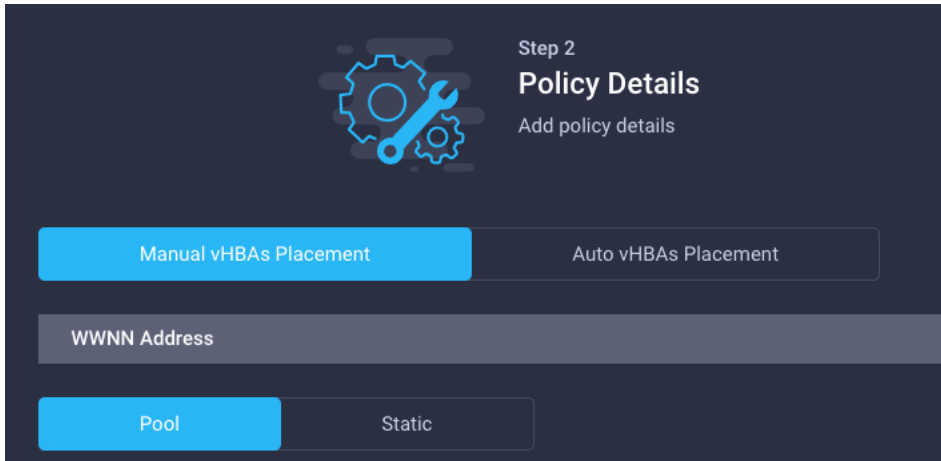
vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-A	MLOM	A	0
vHBA-B	MLOM	B	1

**Step 1.** Click **Select Policy** next to SAN Connectivity and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, SanConn-Pol).

**Step 3.** Select **Manual vHBAs Placement**.

**Step 4.** Select **Pool** under WWNN Address.



**Procedure 12. Create the WWNN Address Pool**

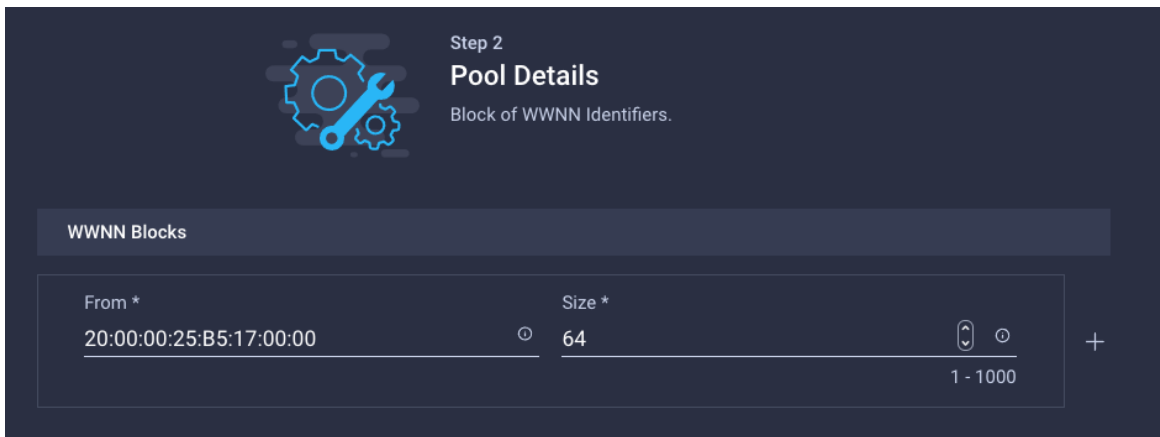
The WWNN address pools have not been defined yet therefore a new WWNN address pool has to be defined.

**Step 1.** Click **Select Pool** under WWNN Address Pool and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, WWNN-Pool).

**Step 3.** Click **Next**.

**Step 4.** Provide the starting WWNN block address and the size of the pool.



**Note:** As a best practice, in FlexPod some additional information is always coded into the WWNN address pool for troubleshooting. For example, in the address 20:00:00:25:B5:17:00:00, 17 is the rack ID.

**Step 5.** Click **Create** to finish creating the WWNN address pool.

**Procedure 13. Create the vHBA-A for SAN A**

**Step 1.** Click **Add vHBA**.

**Step 2.** For vHBA Type, select **fc-initiator** from the drop-down list.

**Procedure 14. Create the WWPN Pool for SAN A**

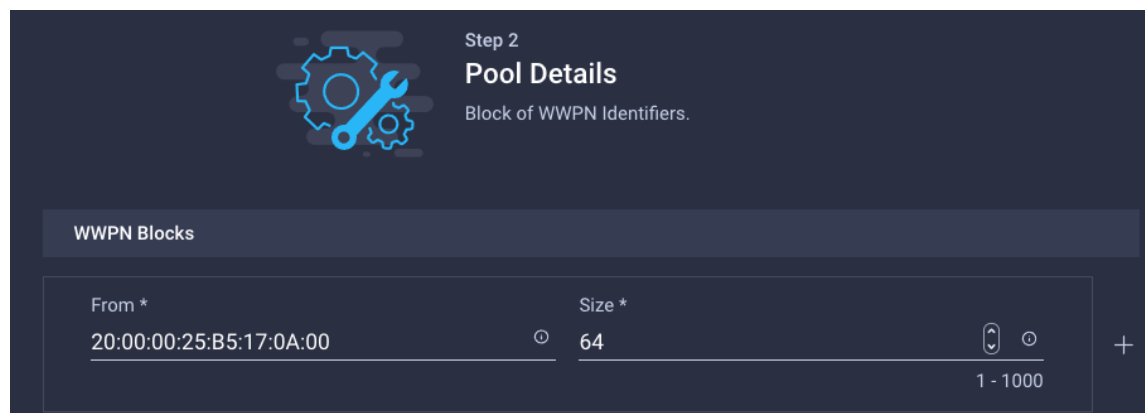
The WWPN address pool has not been defined yet therefore a WWPN address pool for Fabric A will be defined. This pool will also be used for the NVMe-o-FC vHBAs if the vHBAs are defined.

**Step 1.** Click **Select Pool** under WWPN Address Pool and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, WWPN-Pool-A).

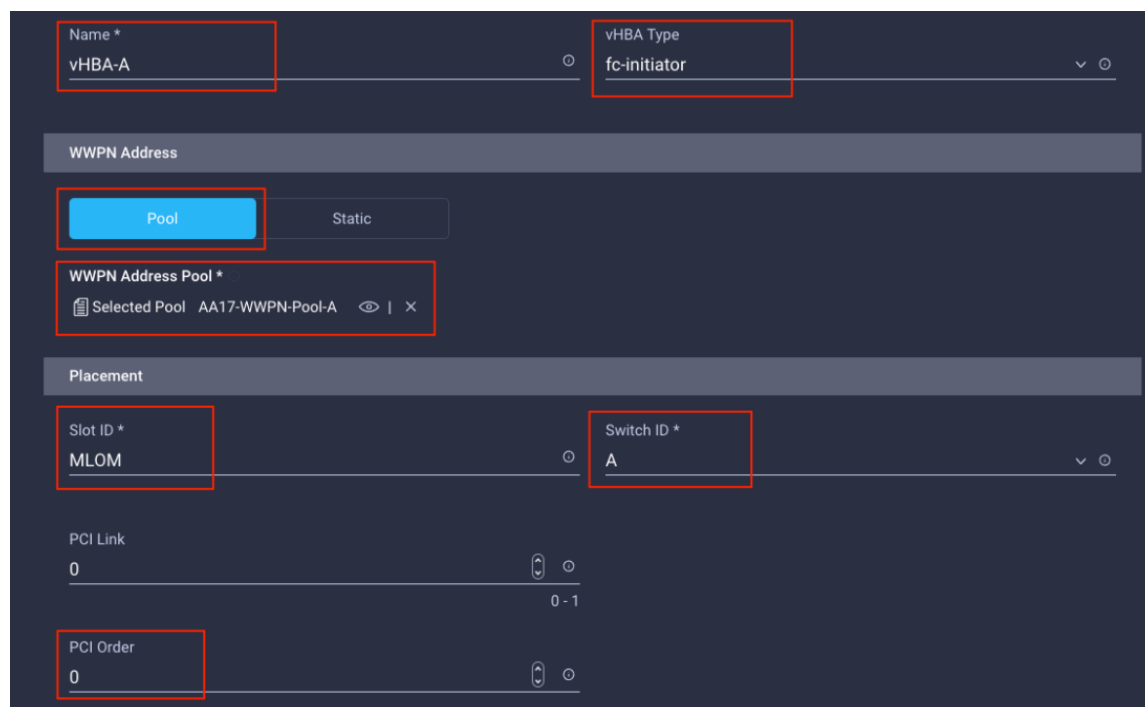
**Step 3.** Provide the starting WWPN block address for SAN A and the size.

**Note:** As a best practice, in FlexPod some additional information is always coded into the WWPN address pool for troubleshooting. For example, in the address 20:00:00:25:B5:17:0A:00, 17 is the rack ID and 0A signifies SAN A.



**Step 4.** Click **Create** to finish creating the WWPN pool.

**Step 5.** Back in the Create vHBA window, provide the Name (for example, vHBA-A), Switch ID (for example, A) and PCI Order from [Table 22](#).



### Procedure 15. Create Fibre Channel Network Policy for SAN A

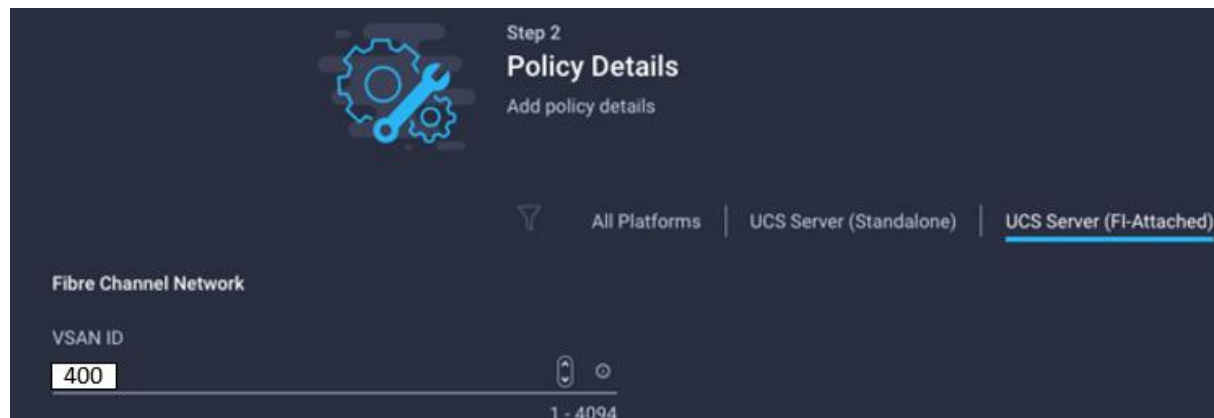
A Fibre Channel network policy governs the VSAN configuration for the virtual interfaces. In this deployment, VSAN 400 will be used for vHBA-A.

**Step 1.** Click **Select Policy** under Fibre Channel Network and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, SAN-A-Network).

**Step 3.** For the scope, select **UCS Server (FI-Attached)**.

**Step 4.** Under VSAN ID, provide the VSAN information (for example, 400).



**Step 5.** Click **Create** to finish creating the Fibre Channel network policy.

### Procedure 16. Create Fibre Channel QoS Policy

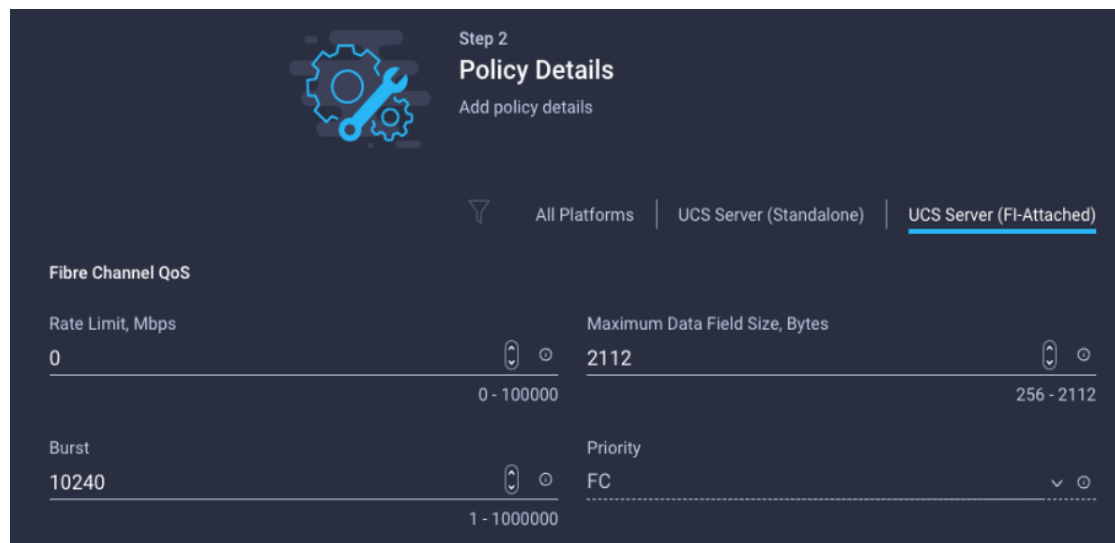
The Fibre Channel QoS policy assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. The Fibre Channel QoS policy used in this deployment uses default values and will be shared by all vHBAs.

**Step 1.** Click **Select Policy** under Fibre Channel QoS and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, FC-QoS).

**Step 3.** For the scope, select **UCS Server (FI-Attached)**.

**Step 4.** Do not change the default values on the Policy Details screen.



**Step 5.** Click **Create** to finish creating the Fibre Channel QoS policy.

**Step 6.** Click **Create** to finish creating the Fibre Channel QoS policy.

### Procedure 17. Create Fibre Channel Adapter Policy

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

**Step 1.** Click **Select Policy** under Fibre Channel Adapter and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, FC-Adapter).

**Step 3.** For the scope, select **UCS Server (FI-Attached)**.

**Step 4.** Do not change the default values on the Policy Details screen.

Step 2  
**Policy Details**  
Add policy details

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

**Error Recovery**

FCP Error Recovery

Port Down Timeout, ms: 10000 (0 - 240000)

Link Down Timeout, ms: 30000 (0 - 240000)

I/O Retry Timeout, Seconds: 5 (1 - 59)

Port Down IO Retry, ms: 30 (0 - 255)

**Error Detection**

Error Detection Timeout: 2000 (1000 - 100000)

**Step 5.** Click **Create** to finish creating the Fibre Channel adapter policy.

**Step 6.** Click **Add** to create vHBA-A.

### Procedure 18. Create the vHBA for SAN B

**Step 1.** Click **Add vHBA**.

**Step 2.** For vHBA Type, select **fc-initiator** from the drop-down list.

### Procedure 19. Create the WWPN Pool for SAN B

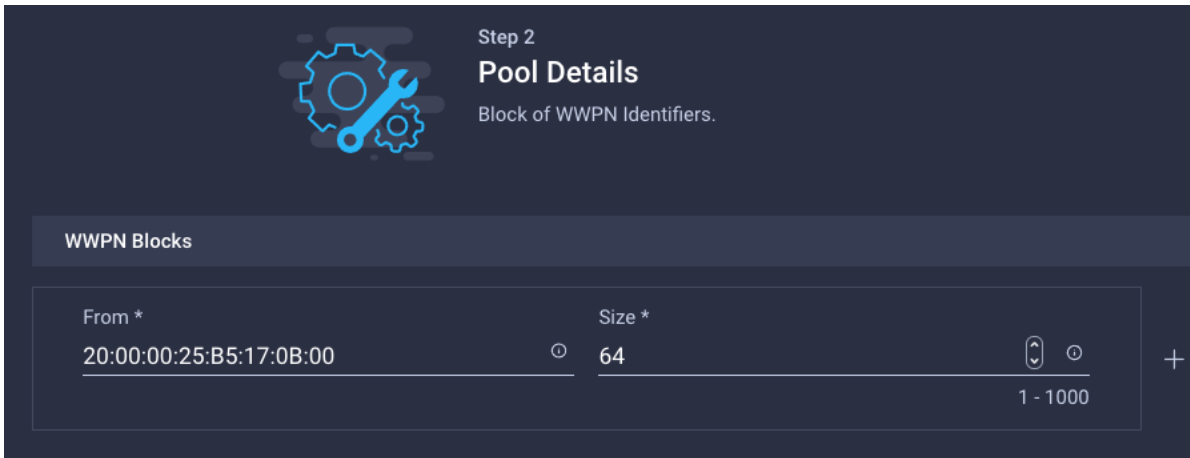
The WWPN address pool has not been defined yet therefore a WWPN address pool for Fabric B will be defined. This pool will also be used for the NVMe-o-FC vHBAs if the vHBAs are defined.

**Step 1.** Click **Select Pool** under WWPN Address Pool and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, WWPN-Pool-B).

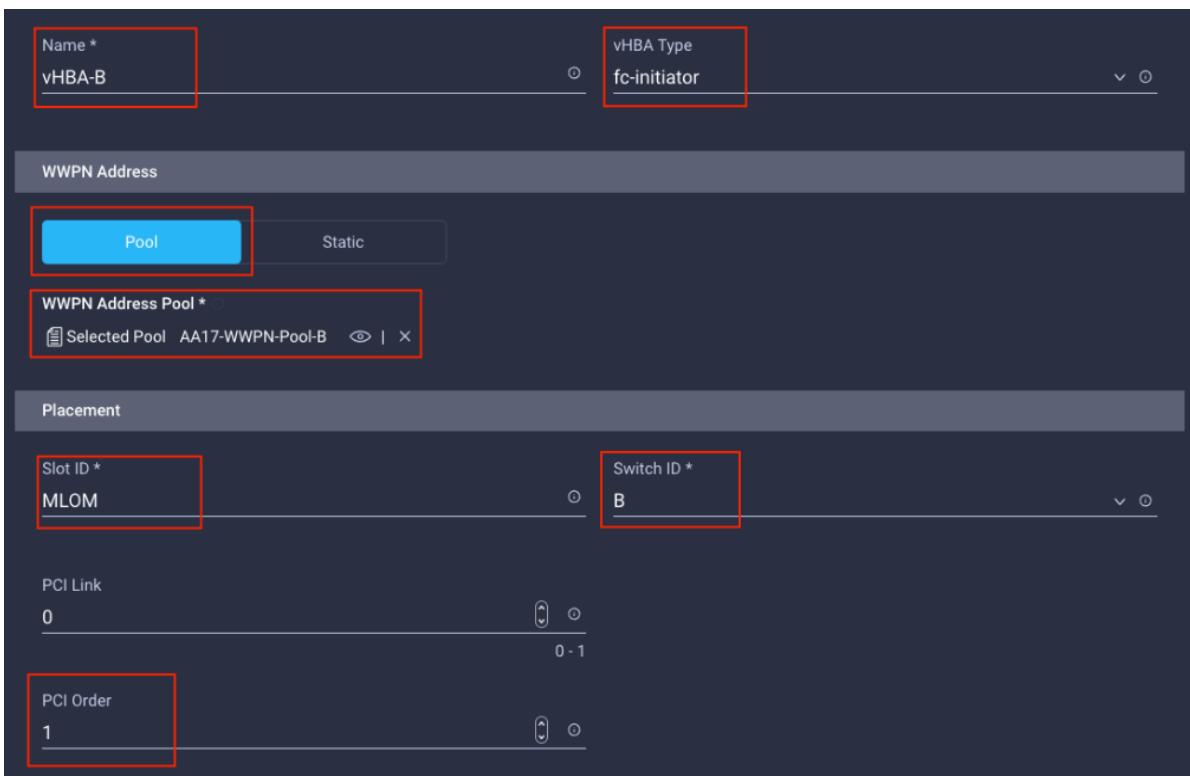
**Step 3.** Provide the starting WWPN block address for SAN B and the size.

**Note:** As a best practice, in FlexPod some additional information is always coded into the WWPN address pool for troubleshooting. For example, in the address 20:00:00:25:B5:17:0B:00, 17 is the rack ID and 0B signifies SAN B.



**Step 4.** Click **Create** to finish creating the WWPN pool.

**Step 5.** Back in the Create vHBA window, provide the Name (for example, vHBA-B), Switch ID (for example, B) and PCI Order from [Table 22](#).



## Procedure 20. Create Fibre Channel Network Policy for SAN B

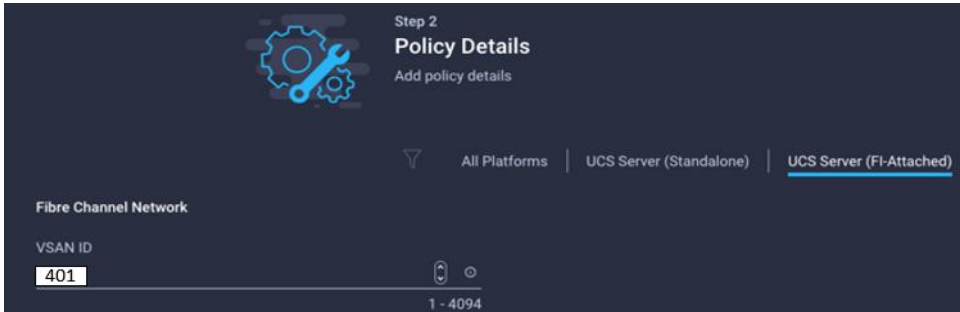
**Note:** In this deployment, VSAN 401 will be used for vHBA-B.

**Step 1.** Click **Select Policy** under Fibre Channel Network and then click **Create New**.

**Step 2.** Verify correct organization is selected from the drop-down list (for example, AA17) and provide a name for the policy (for example, SAN-B-Network).

**Step 3.** For the scope, select UCS Server (FI-Attached).

**Step 4.** Under VSAN ID, provide the VSAN information (for example, 401).



**Step 5.** Click **Create**.

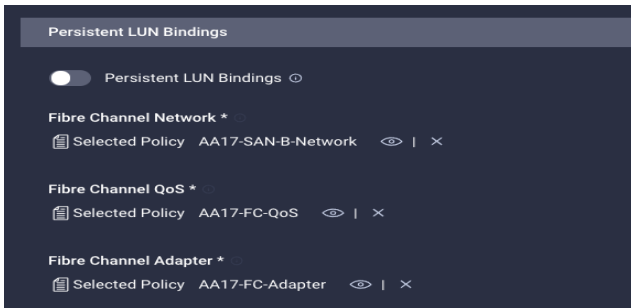
**Step 6.** Select Fibre Channel QoS policy for SAN B

**Step 7.** Click **Select Policy** under Fibre Channel QoS and then select the previously created QoS policy FC-QoS.

**Step 8.** Select Fibre Channel Adapter policy for SAN B

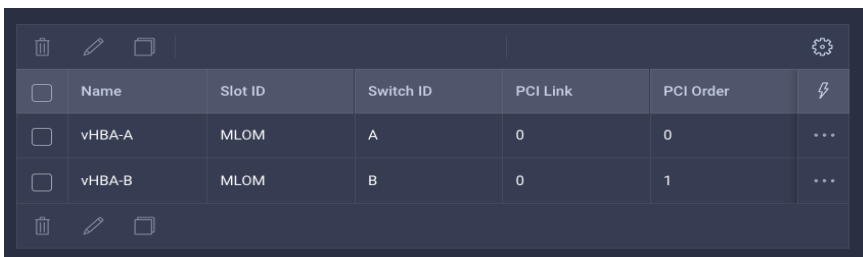
**Step 9.** Click **Select Policy** under Fibre Channel Adapter and then select the previously created Adapter policy FC-Adapter.

**Step 10.** Verify all the vHBA policies are mapped.



**Step 11.** Click **Add** to add the vHBA-B.

**Step 12.** Verify both the vHBAs are added to the SAN connectivity policy.



	Name	Slot ID	Switch ID	PCI Link	PCI Order	
<input type="checkbox"/>	vHBA-A	MLOM	A	0	0	...
<input type="checkbox"/>	vHBA-B	MLOM	B	0	1	...

**Note:** If the customers don't need the NVMe-o-FC connectivity, skip the following sections for creating NVMe vHBAs.

### Procedure 21. Configure vHBA-NVMe-A and vHBA-NVMe-B

To configure (optional) NVMe-o-FC, two vHBAs, one for each fabric, needs to be added to the server profile template. These vHBAs are in addition to the FC boot from SAN vHBAs, vHBA-A and vHBA-b.

**Table 23. vHBA placement for NVMe-o-FC**

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-NVMe-A	MLOM	A	6
vHBA-NVMe-B	MLOM	B	7

**Procedure 22. Configure vHBA-NVMe-A**

**Step 1.** Click **Add vHBA**.

**Step 2.** For vHBA Type, select **fc-nvme-initiator** from the drop-down list.

**Step 3.** Click **Select Pool** under WWPN Address Pool and then select the previously created pool WWPN-Pool-A.

**Step 4.** Provide the Name (for example, vHBA-NVMe-A), Switch ID (for example, A) and PCI Order from [Table 23](#).

The screenshot shows a configuration form for a vHBA. The fields are as follows:

- Name \***: vHBA-NVMe-A
- vHBA Type**: fc-nvme-initiator
- WWPN Address**: Pool (selected), Static
- WWPN Address Pool \***: Selected Pool AA17-WWPN-Pool-A
- Placement**:
  - Slot ID \***: MLOM
  - Switch ID \***: A
  - PCI Link**: 0 (range 0-1)
  - PCI Order**: 6 (range 0-1)

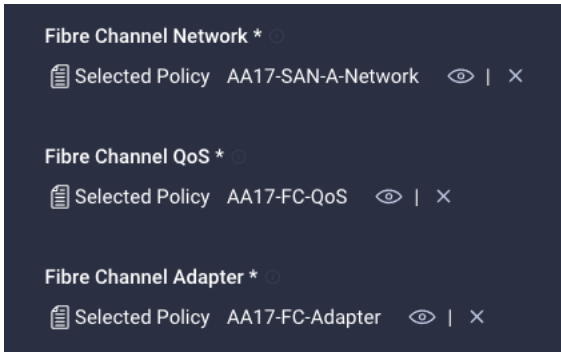
**Step 5.** Click **Select Policy** under Fibre Channel Network and then select the previously created policy for SAN A, SAN-A-Network.

**Step 6.** Click **Select Policy** under Fibre Channel QoS and then select the previously created QoS policy FC-QoS.

**Step 7.** Click **Select Policy** under Fibre Channel Adapter and then select the previously created Adapter policy FC-Adapter.

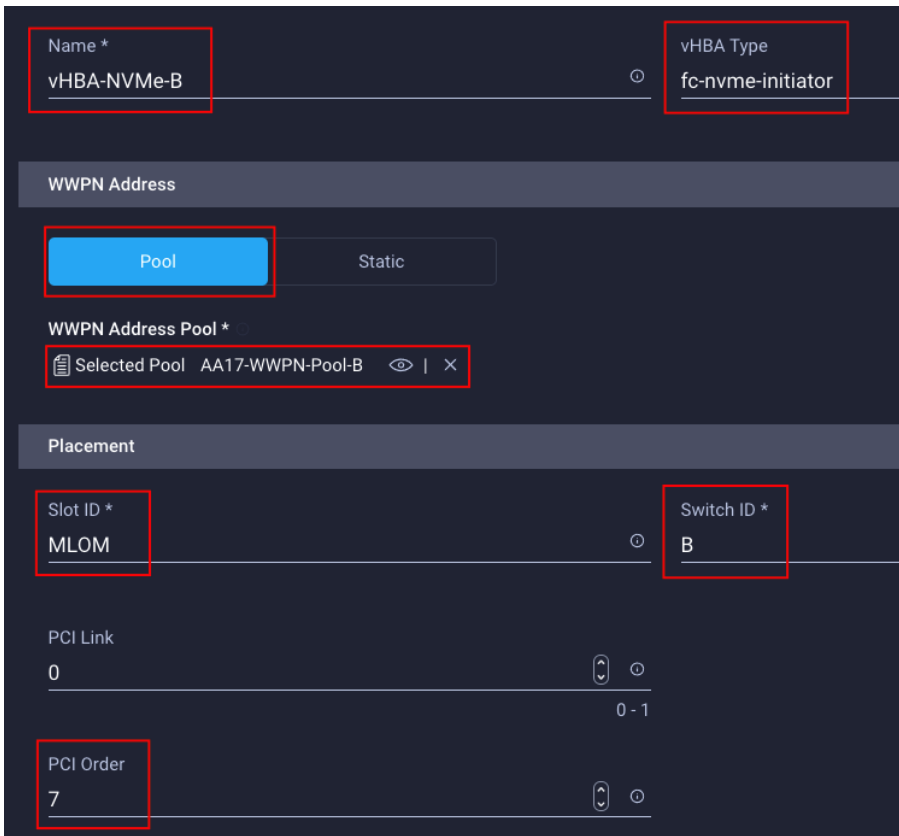
**Step 8.** Verify all the vHBA policies are mapped.





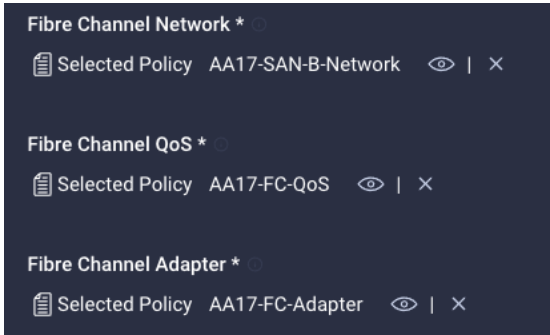
**Procedure 23. Configure vHBA-NVMe-B**

- Step 1.** Click **Add vHBA**.
- Step 2.** For vHBA Type, select **fc-nvme-initiator** from the drop-down list.
- Step 3.** Click **Select Pool** under WWPN Address Pool and then select the previously created pool WWPN-Pool-B.
- Step 4.** Provide the Name (for example, vHBA-NVMe-B), Switch ID (for example, B) and PCI Order from [Table 23](#).



- Step 5.** Click **Select Policy** under Fibre Channel Network and then select the previously created policy for SAN B, SAN-B-Network.
- Step 6.** Click **Select Policy** under Fibre Channel QoS and then select the previously created QoS policy FC-QoS.
- Step 7.** Click **Select Policy** under Fibre Channel Adapter and then select the previously created Adapter policy FC-Adapter.

**Step 8.** Verify all the vHBA policies are mapped correctly.



**Step 9.** Verify all four vHBAs are added to the SAN connectivity policy.

Name	Slot ID	Switch ID	PCI Link	PCI Order
vHBA-NVMe-A	MLOM	A	0	6
vHBA-B	MLOM	B	0	1
vHBA-A	MLOM	A	0	0
vHBA-NVMe-B	MLOM	B	0	7

**Step 10.** Click **Create** to create the SAN connectivity policy with NVMe-o-FC support.

## Procedure 24. Verify Summary

**Step 1.** When the LAN connectivity policy and SAN connectivity policy (for FC) is created, click **Next** to move to the Summary screen.

**Step 2.** On the summary screen, verify policies mapped to various settings. The screenshots below provide summary view for a FC boot from SAN server profile template.

**Step 6 Summary**  
Verify details of the template and the policies, resolve errors and deploy.

---

**General**

Template Name	FC-Boot-Template	Organization	AA17
Target Platform	UCS Server (FI-Attached)		

Description  
FC Boot

<u>Compute Configuration</u>	Management Configuration	Storage Configuration	Network Configuration	Errors/Warnings (0)
------------------------------	--------------------------	-----------------------	-----------------------	---------------------

---

BIOS	AA17-BIOS-Pol
Boot Order	AA17-FC-BootOrder-Pol
UUID	AA17-UUID-Pool

Description FC Boot				
Compute Configuration	Management Configuration	Storage Configuration	Network Configuration	Errors/Warnings (0)
IMC Access			AA17-IMC-Access	
IPMI Over LAN			Enable-IPMIoLAN	
Local User			AA17-LocalUserPol	

Description FC Boot				
Compute Configuration	Management Configuration	Storage Configuration	Network Configuration	Errors/Warnings (0)
LAN Connectivity			AA17-FC-ESXi-LANConn-Manual	
SAN Connectivity			AA17-SanConn-Pol	
Adapter #MLOM (vNICs)	4	Adapter #MLOM (vHBAs)	2	

## Procedure 25. Derive Server Profile

**Step 1.** From the Server profile template Summary screen, click **Derive Profiles**.

**Note:** This action can also be performed later by navigating to **Templates**, clicking “...” next to the template name and selecting **Derive Profiles**.

**Step 2.** Under the Server Assignment, select **Assign Now** and click **Cisco UCS X210c M6**. You can select one or more servers depending on the number of profiles to be deployed.

**Server Assignment**

Assign Now    Assign Server from a Resource Pool    Assign Later

Add Filter    14 items found    10

<input checked="" type="checkbox"/>	Name	User Label	Health	Model
<input checked="" type="checkbox"/>	vdi-tme-2-7		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-3		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-4		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-8		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-1-5		Healthy	UCSBX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-1		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-1-8		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-2		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-6		Healthy	UCSX-210C-M6
<input checked="" type="checkbox"/>	vdi-tme-2-5		Healthy	UCSX-210C-M6

**Note:** The server profile template and policies in this document apply to both Cisco UCS X210x M6 and Cisco UCS X-Series servers.

**Step 3.** Click **Next**.

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

**Step 2**  
**Details**  
Edit the description, tags, and auto-generated names of the profiles.

**General**

Organization \*  
X-Series

Target Platform  
UCS Server (FI-Attached)

Description  
VDI\_FC\_Boot

Set Tags

**Derive**

Profile Name Prefix  
vdi-FC-Boot\_DERIVED-

Start Index for Suffix  
1

1 Name \*  
vdi-FC-Boot\_DERIVED-1

**Step 4.** Adjust the Prefix and number if needed.

**Step 2 Details**  
Edit the description, tags, and auto-generated names of the profiles.

**General**

Organization \*  
X-Series

Target Platform  
UCS Server (FI-Attached)

Description  
VDI\_FC\_Boot

Set Tags

<= 1024

**Derive**

Profile Name Prefix  
vdi-tme

Start Index for Suffix  
1


> 0

1	Name *
	vdi-tme-1

**Step 5.** Click **Next**.

**Step 6.** Verify the information and click **Derive** to create the Server Profiles.

**Step 7.** Cisco Intersight will start configuring the server profiles and will take some time to apply all the policies. Use the Requests tab to see the progress.



**Step 8.** When the Server Profiles are deployed successfully, they will appear under the Server Profiles with the status of OK.

CONFIGURE > Profiles

HyperFlex Cluster Profiles UCS Chassis Profiles UCS Domain Profiles **UCS Server Profiles**

\* All UCS Server Profiles

Add Filter

<input type="checkbox"/>	Name	Status	Target Platform	UCS Server Template
<input type="checkbox"/>	vdi-09	OK	UCS Server (FI-Attached)	vdi-SvcPrfl
<input type="checkbox"/>	vdi-08	OK	UCS Server (FI-Attached)	vdi-SvcPrfl
<input type="checkbox"/>	vdi-07	OK	UCS Server (FI-Attached)	vdi-SvcPrfl
<input type="checkbox"/>	vdi-06	OK	UCS Server (FI-Attached)	vdi-SvcPrfl
<input type="checkbox"/>	vdi-05	OK	UCS Server (FI-Attached)	vdi-SvcPrfl
<input type="checkbox"/>	vdi-04	OK	UCS Server (FI-Attached)	vdi-SvcPrfl
<input type="checkbox"/>	vdi-03	OK	UCS Server (FI-Attached)	vdi-SvcPrfl

## SAN Switch Configuration

This subject explains how to configure the Cisco MDS 9000s for use in a FlexPod environment.

**IMPORTANT! Follow the steps precisely because failure to do so could result in an improper configuration.**

**Note:** If you're directly connecting storage to the Cisco UCS fabric interconnects, skip this section.

## Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained here: [FlexPod Cabling](#).

## FlexPod Cisco MDS Switch Configuration

This subject has the following procedures:

- [Configure Cisco MDS 9132T A Switch](#)
- [Configure Cisco MDS 9132T B Switch](#)
- [Enable Features on Cisco MDS 9132T A and Cisco MDS 9132T B](#)
- [Configure the Second NTP Server and Add Local Time](#)
- [Configure Individual Ports for Cisco MDS 9132T A](#)
- [Configure Individual Ports for Cisco MDS 9132T B](#)
- [Create VSANs for Cisco MDS 9132T A](#)
- [Create VSANs for Cisco MDS 9132T B](#)
- [Create Device Aliases for Cisco MDS 9132T A](#)
- [Create Device Aliases for Cisco MDS 9132T B](#)
- [Create Zones and Zoneset for Cisco MDS 9132T A](#)
- [Create Zones and Zoneset for Cisco MDS 9132T B](#)

### Procedure 1. Configure Cisco MDS 9132T A Switch

**Step 1.** On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning. Enter y to get to the System Admin Account Setup.

**Step 2.** Configure the switch using the command line:

```
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name : <mds-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address : <mds-A-mgmt0-ip>
Mgmt0 IPv4 netmask : <mds-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <mds-A-mgmt0-gw>
```

```

Configure advanced IP options? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter
Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter
Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter
Enable the http-server? (yes/no) [y]: Enter
Configure clock? (yes/no) [n]: Enter
Configure timezone? (yes/no) [n]: Enter
Configure summertime? (yes/no) [n]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <nexus-A-mgmt0-ip>
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: Enter
Configure default zone mode (basic/enhanced) [basic]: Enter

```

**Step 3. Run the following commands to review the configuration:**

```

Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter

```

## Procedure 2. Configure Cisco MDS 9132T B Switch

**Step 1.** On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning. Enter y to get to the System Admin Account Setup.

**Step 2.** Configure the switch using the command line:

```

---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name : <mds-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address : <mds-B-mgmt0-ip>
Mgmt0 IPv4 netmask : <mds-B-mgmt0-netmask>

```

```

Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <mds-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter
Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter
Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter
Enable the http-server? (yes/no) [y]: Enter
Configure clock? (yes/no) [n]: Enter
Configure timezone? (yes/no) [n]: Enter
Configure summertime? (yes/no) [n]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <nexus-A-mgmt0-ip>
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: Enter
Configure default zone mode (basic/enhanced) [basic]: Enter

```

**Step 3. Run the following commands to review the configuration:**

```

Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter

```

**Procedure 3. Enable Features on Cisco MDS 9132T A and Cisco MDS 9132T B**

**Step 1.** Log in as admin.

**Step 2.** Run the following commands:

```

configure terminal
feature npiv
feature fport-channel-trunk

```

**Procedure 4. Configure the Second NTP Server and Add Local Time**

**Step 1.** From the global configuration mode, run the following command:

```

ntp server <nexus-B-mgmt0-ip>
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-
week> <end-day> <end-month> <end-time> <offset-minutes>

```

**Note:** It is important to configure the local time so that logging time alignment, any backup schedules, and SAN Analytics forwarding are correct. For more information on configuring the timezone and daylight



savings time or summer time, go to: [Cisco MDS 9000 Series Fundamentals Configuration Guide, Release 8.x](#). Sample clock commands for the United States Eastern timezone are:  
clock timezone EST -5 0  
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60

## Procedure 5. Configure Individual Ports for Cisco MDS 9132T A

**Step 1.** From the global configuration mode, run the following commands:

```
interface fc1/3
switchport description <st-clustername>-1:1a
switchport speed 32000
switchport trunk mode off
no shutdown
exit
interface fc1/4
switchport description <st-clustername>-2:1a
switchport speed 32000
switchport trunk mode off
no shutdown
exit
interface fc1/1
switchport description <ucs-clustername>-a:1/1
channel-group 15
no shutdown
exit
interface fc1/2
switchport description <ucs-clustername>-a:1/2
channel-group 15
no shutdown
exit
interface port-channel15
channel mode active
switchport trunk allowed vsan <vsan-a-id>
switchport description <ucs-clustername>-a
switchport speed 32000
no shutdown
exit
```

**Note:** If VSAN trunking is not being used between the Cisco UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-a-id>” for interface port-channel15. The default setting of switchport trunk mode auto is being used for the port channel.

## Procedure 6. Configure Individual Ports for Cisco MDS 9132T B

**Step 1.** From the global configuration mode, run the following commands:

```
interface fc1/5
```

```

switchport description <st-clustername>-1:1b
switchport speed 32000
switchport trunk mode off
no shutdown
exit
interface fc1/6
switchport description <st-clustername>-2:1b
switchport speed 32000
switchport trunk mode off
no shutdown
exit
interface fc1/1
switchport description <ucs-clustername>-b:1/1
channel-group 15
no shutdown
exit
interface fc1/2
switchport description <ucs-clustername>-b:1/2
channel-group 15
no shutdown
exit
interface port-channel15
channel mode active
switchport trunk allowed vsan <vsan-b-id>
switchport description <ucs-clustername>-b
switchport speed 32000
no shutdown
exit

```

**Note:** If VSAN trunking is not being used between the Cisco UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-b-id>” for interface port-channel15. The default setting of switchport trunk mode auto is being used for the port channel.

## Procedure 7. Create VSANs for Cisco MDS 9132T A

**Step 1.** From the global configuration mode, run the following commands:

```

vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fc1/9
vsan <vsan-a-id> interface fc1/10

```

```
vsan <vsan-a-id> interface port-channel15
exit
```

## Procedure 8. Create VSANs for Cisco MDS 9132T B

**Step 1.** From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fc1/9
vsan <vsan-b-id> interface fc1/10
vsan <vsan-b-id> interface port-channel15
exit
```

**Step 2.** At this point, it may be necessary to go into Cisco UCS Manager and disable and then enable the FC port-channel interfaces to get the port-channels to come up.

## Procedure 9. Create Device Aliases for Cisco MDS 9132T A

**Note:** Device aliases for Fabric A will be used to create zones.

**Step 1.** From the global configuration mode, run the following commands:

```
device-alias mode enhanced
device-alias database
device-alias name Infra-FC-fcp-lif-01a pwnn <fcp-lif-01a-wwpn>
device-alias name Infra-FC-fcp-lif-02a pwnn <fcp-lif-02a-wwpn>
device-alias name VM-Host-Infra-01-A pwnn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwnn <vm-host-infra-02-wwpna>
device-alias name VM-Host-Infra-03-A pwnn <vm-host-infra-03-wwpna>
device-alias commit
```

## Procedure 10. Create Device Aliases for Cisco MDS 9132T B

**Step 1.** From the global configuration mode, run the following commands:

```
device-alias mode enhanced
device-alias database
device-alias name Infra-FC-fcp-lif-01b pwnn <fcp-lif-01b-wwpn>
device-alias name Infra-FC-fcp-lif-02b pwnn <fcp-lif-02b-wwpn>
device-alias name VM-Host-Infra-01-B pwnn <vm-host-infra-01-wwpnb>
device-alias name VM-Host-Infra-02-B pwnn <vm-host-infra-02-wwpnb>
device-alias name VM-Host-Infra-03-B pwnn <vm-host-infra-03-wwpnb>
device-alias commit
```

## Procedure 11. Create Zones and Zoneset for Cisco MDS 9132T A

**Step 1.** To create the required zones and zoneset on Fabric A, run the following commands:

```
configure terminal
zone name Infra-FC-Fabric-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-01-A init
member device-alias VM-Host-Infra-02-A init
member device-alias VM-Host-Infra-03-A init
member device-alias Infra-FC-fcp-lif-01a target
member device-alias Infra-FC-fcp-lif-02a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member Infra-FC-Fabric-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
show zoneset active
copy r s
```

**Note:** Since Smart Zoning is enabled, a single zone is created with all host boot initiators and boot targets for the Infra-FC instead of creating a separate zone for each host with the host initiator and boot targets. If a new host is added, its boot initiator can simply be added to the single zone in each MDS switch and then the zoneset reactivated. If another SVM is added to the FlexPod with FC targets, a new zone can be added for that SVM.

## Procedure 12. Create Zones and Zoneset for Cisco MDS 9132T B

**Step 1.** To create the required zones and zoneset on Fabric B, run the following commands:

```
configure terminal
zone name Infra-FC-Fabric-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-01-B init
member device-alias VM-Host-Infra-02-B init
member device-alias VM-Host-Infra-03-B init
member device-alias Infra-FC-fcp-lif-01b target
member device-alias Infra-FC-fcp-lif-02b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member Infra-FC-Fabric-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active
copy r s
```

## Storage Configuration – Boot LUNs

This chapter contains the following:

- [NetApp ONTAP Boot Storage Setup](#)
- [Install VMware ESXi 7.0](#)
- [VMware vCenter 7.0](#)

### NetApp ONTAP Boot Storage Setup

This subject contains the following procedures:

- [Create Boot LUNs](#)
- [Create igroups](#)
- [Map Boot LUNs to igroups](#)

#### Procedure 1. Create Boot LUNs

**Step 1.** Run the following commands to create three boot LUNs,:

```
lun create -vserver Infra-FC -path /vol/esxi_boot/VM-Host-Infra-01 -size 32GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-FC -path /vol/esxi_boot/VM-Host-Infra-02 -size 32GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-FC -path /vol/esxi_boot/VM-Host-Infra-03 -size 32GB -ostype vmware -space-reserve disabled
```

#### Procedure 2. Create igroups

**Step 1.** Create initiator groups (igroups) by entering the following commands from the storage cluster management node Secure Shell (SSH) connection:

```
lun igroup create -vserver Infra-FC -igroup VM-Host-Infra-01 -protocol fcp -ostype vmware -initiator <vm-host-infra-01-wwpna>, <vm-host-infra-01-wwpnb>
lun igroup create -vserver Infra-FC -igroup VM-Host-Infra-02 -protocol fcp -ostype vmware -initiator <vm-host-infra-02-wwpna>, <vm-host-infra-02-wwpnb>
lun igroup create -vserver Infra-FC -igroup VM-Host-Infra-03 -protocol fcp -ostype vmware -initiator <vm-host-infra-03-wwpna>, <vm-host-infra-03-wwpnb>
lun igroup create -vserver Infra-FC -igroup MGMT-Hosts -protocol fcp -ostype vmware -initiator <vm-host-infra-01-wwpna>, <vm-host-infra-01-wwpnb>, <vm-host-infra-02-wwpna>, <vm-host-infra-02-wwpnb>, <vm-host-infra-03-wwpna>, <vm-host-infra-03-wwpnb>
```

**Step 2.** To view the three igroups just created, use the command `lun igroup show`:

```
lun igroup show -protocol fcp
```

#### Procedure 3. Map Boot LUNs to igroups

**Step 1.** From the storage cluster management SSH connection, enter the following commands:

```
lun mapping create -vserver Infra-FC -path /vol/esxi_boot/VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-id 0
lun mapping create -vserver Infra-FC -path /vol/esxi_boot/VM-Host-Infra-02 -igroup VM-Host-Infra-02 -lun-id 0
lun mapping create -vserver Infra-FC -path /vol/esxi_boot/VM-Host-Infra-03 -igroup VM-Host-Infra-03 -lun-id 0
```

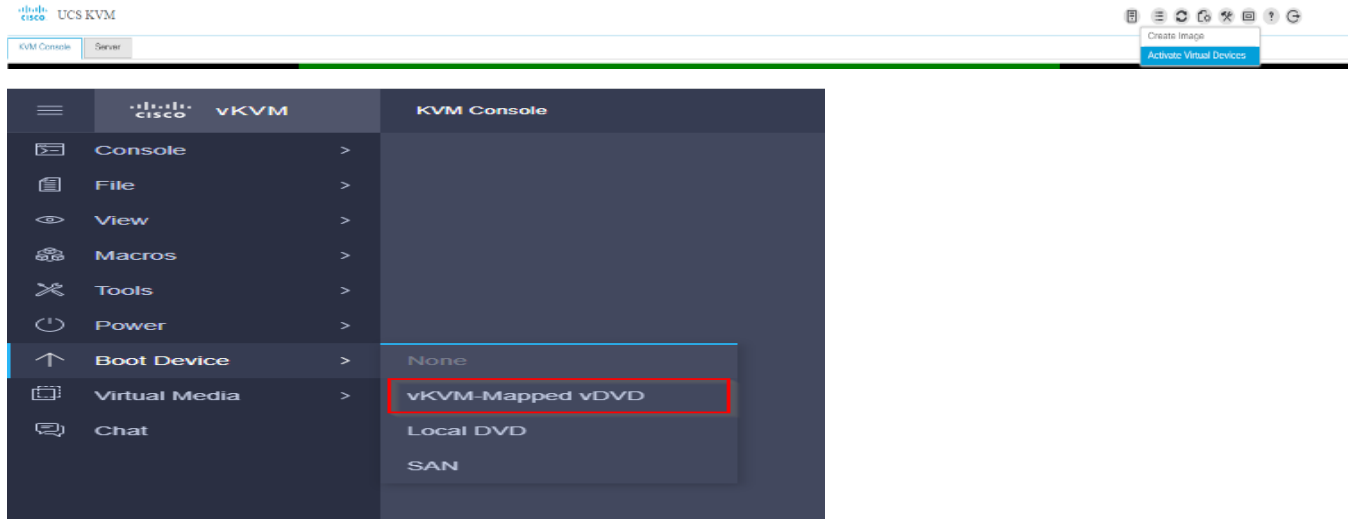
**Step 2.** Download the Cisco Custom Image for VMware ESXi 7.0 Update 3: <https://customerconnect.vmware.com/en/downloads/details?downloadGroup=OEM-ESXi70U3-CISCO&productId=974&download=true&fileId=9f1e4a251c08193f303ad9a5912b6540&uuld=1a0d093d-17e2-4760-9dbb-4917ba7fdc21>. From the page click the “Custom ISOs” tab.

**Step 3.** In the Cisco UCS Manager navigation pane, click the Equipment tab.


**Step 4.** Under Servers > Service Profiles> VDI-Host1

**Step 5.** Right-click on VDI-Host1 and select KVM Console.

**Step 6.** Click Boot Device and then select CD/DVD.

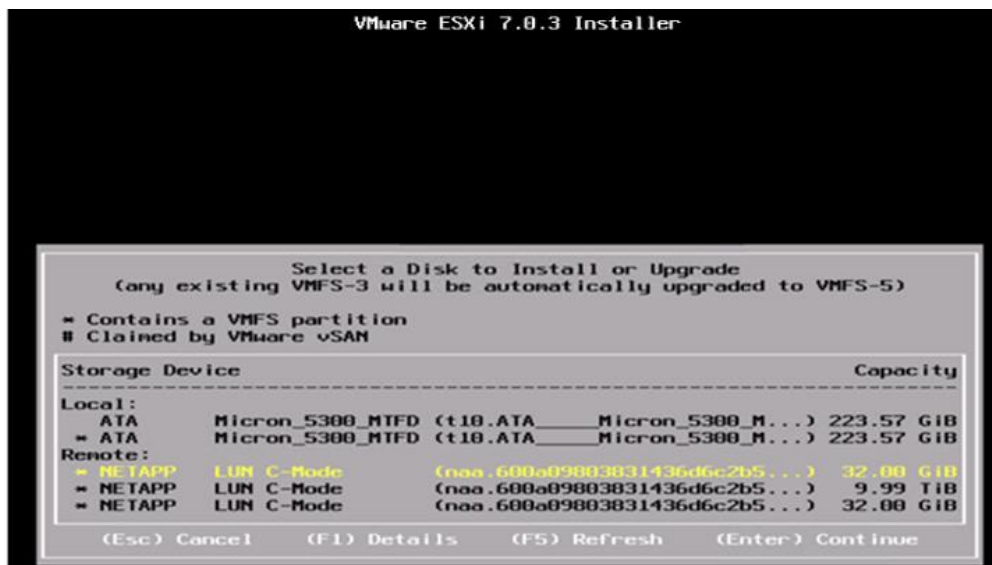


**Step 7.** Click Virtual Media and Mount the ESXi ISO image.



**Step 8.** Boot into ESXi installer and follow the prompts to complete installing VMware vSphere ESXi hypervisor.

**Step 9.** When selecting a storage device to install ESXi, select Remote LUN provisioned through NetApp Storage Administrative console and access through FC connection.



**Note:** Adding a management network for each VMware host is necessary for managing the host and connection to vCenter Server. Please select the IP address that can communicate with an existing or a new vCenter Server.

**Step 10.** After the server has finished rebooting, press F2 to enter into configuration wizard for ESXi Hypervisor.

**Step 11.** Log in as root and enter the corresponding password.

**Step 12.** Select the Configure the Management Network option and press Enter.

**Step 13.** Select the VLAN (Optional) option and press Enter. Enter the VLAN In-Band management ID and press Enter.

**Step 14.** From the Configure Management Network menu, select “IP Configuration” and press Enter.

**Step 15.** Select “Set Static IP Address and Network Configuration” option by using the space bar. Enter the IP address to manage the first ESXi host. Enter the subnet mask for the first ESXi host. Enter the default gateway for the first ESXi host. Press Enter to accept the changes to the IP configuration.

**Step 16.** IPv6 Configuration is set to automatic.

**Step 17.** Select the DNS Configuration option and press Enter.

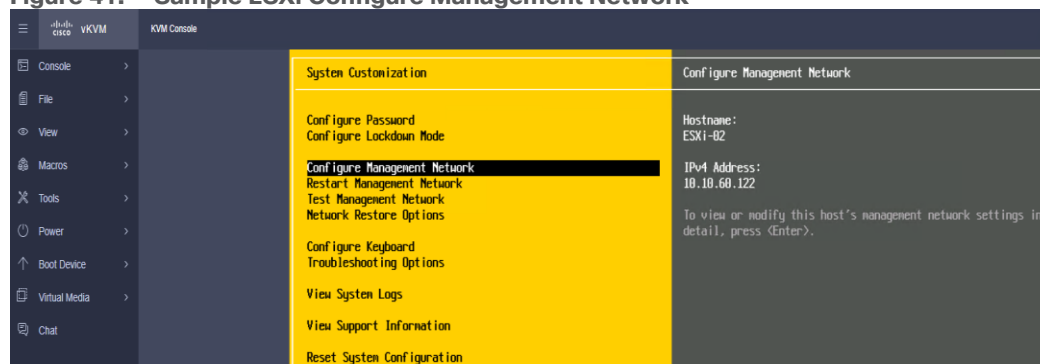
**Step 18.** Enter the IP address of the primary and secondary DNS server. Enter Hostname

**Step 19.** Enter DNS Suffixes.

**Step 20.** Since the IP address is assigned manually, the DNS information must also be entered manually.

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

**Figure 41. Sample ESXi Configure Management Network**



## Install VMware ESXi 7.0

This subject contains the following procedures:

- [Download ESXi 7.0 from VMware](#)
- [Log into the Cisco UCS Environment using Cisco UCS Manager](#)
- [Prepare the Server for the OS Installation](#)
- [Install VMware ESXi to the Bootable LUN of the Hosts](#)
- [Set Up Management Networking for ESXi Hosts](#)
- [Reset VMware ESXi Host VMkernel Port vmk0 MAC Address \(Optional\)](#)
- [Install VMware and Cisco VIC Drivers for the ESXi Host](#)
- [Install VMware VIC Drivers and the NetApp NFS Plug-in for VMware VAAI on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02](#)
- [Log into the First VMware ESXi Host by Using VMware Host Client](#)
- [Set Up VMkernel Ports and Virtual Switch for ESXi Host VM-Host-Infra-01](#)
- [Mount Required Datastores on ESXi Host VM-Host-Infra-01](#)
- [Configure NTP on First ESXi Host on ESXi Host VM-Host-Infra-01](#)
- [Configure ESXi Host Swap on ESXi Host VM-Host-Infra-01](#)
- [Configure Host Power Policy on ESXi Host VM-Host-Infra-01](#)

This section provides detailed instructions for installing VMware ESXi 7.0 in a FlexPod environment. After the procedures are completed, three booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

### Procedure 1. Download ESXi 7.0 from VMware

**Step 1.** Click this following link: [Cisco Custom ISO for UCS 4.1.2a](#). You will need a user id and password on vmware.com to download this software:

<https://customerconnect.vmware.com/downloads/details?downloadGroup=OEM-ESXI70U3-CISCO&productId=974>



---

**Note:** The Cisco Custom ISO for UCS 4.1.2a should also be used for Cisco UCS software release 5.0(1b) and VMware vSphere 7.0.

**Step 2.** Download the .iso file.

### **Procedure 2. Log into the Cisco UCS Environment using Cisco UCS Manager**

The Cisco UCS IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log into the Cisco UCS environment to run the IP KVM.

**Step 1.** Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.

**Step 2.** Click the Launch UCS Manager link to launch the HTML 5 UCS Manager GUI.

**Step 3.** If prompted to accept security certificates, accept, as necessary.

**Step 4.** When prompted, enter admin for the username and enter the administrative password.

**Step 5.** To log into Cisco UCS Manager, click Login.

**Step 6.** From the main menu, click Servers.

**Step 7.** Click Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-01.

**Step 8.** In the Actions pane, click KVM Console.

**Step 9.** Follow the prompts to launch the HTML5 KVM console.

**Step 10.** Click Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-02.

**Step 11.** In the Actions pane, click KVM Console.

**Step 12.** Follow the prompts to launch the HTML5 KVM console.

**Step 13.** Go to Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-03.

**Step 14.** In the Actions pane, click KVM Console.

**Step 15.** Follow the prompts to launch the HTML5 KVM console.

### **Procedure 3. Prepare the Server for the OS Installation**

**Note:** Skip this section if you're using vMedia policies; the ISO file will already be connected to KVM.

**Step 1.** In the KVM window, click Virtual Media.

**Step 2.** Select Activate Virtual Devices.

**Step 3.** If prompted to accept an Unencrypted KVM session, accept, as necessary.

**Step 4.** Click Virtual Media and select Map CD/DVD.

**Step 5.** Browse to the ESXi installer ISO image file and click Open.

**Step 6.** Click Map Device.

**Step 7.** Click the KVM Console tab to monitor the server boot.

### **Procedure 4. Install VMware ESXi to the Bootable LUN of the Hosts**

**Step 1.** Boot the server by selecting Boot Server in the KVM and click OK, then click OK again.


**Step 2.** On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.

- 
- Step 3.** If the ESXi installer fails to load because the software certificates cannot be validated, reset the server, and when prompted, press F2 to go into BIOS and set the system time and date to current. Now the ESXi installer should load properly.
- Step 4.** After the installer is finished loading, press Enter to continue with the installation.
- Step 5.** Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
- Step 6.** It may be necessary to map function keys as User Defined Macros under the Macros menu in the Cisco UCS KVM console.
- Step 7.** Select the LUN that was previously set up for the installation disk for ESXi and press Enter to continue with the installation.
- Step 8.** Select the appropriate keyboard layout and press Enter.
- Step 9.** Enter and confirm the root password and press Enter.
- Step 10.** The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
- Step 11.** After the installation is complete, press Enter to reboot the server.
- Step 12.** The ESXi installation image will be automatically unmapped in the KVM when Enter is pressed.
- Step 13.** In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

#### **Procedure 5. Set Up Management Networking for ESXi Hosts**

**Note:** Adding a management network for each VMware host is necessary for managing the host.

- Step 1.** After the server has finished rebooting, in the UCS KVM console, press F2 to customize VMware ESXi.
- Step 2.** Log in as root, enter the corresponding password, and press Enter to log in.
- Step 3.** Use the down arrow key to select Troubleshooting Options and press Enter.
- Step 4.** Select Enable ESXi Shell and press Enter.
- Step 5.** Select Enable SSH and press Enter.
- Step 6.** Press Esc to exit the Troubleshooting Options menu.
- Step 7.** Select the Configure Management Network option and press Enter.
- Step 8.** Select Network Adapters and press Enter.
- Step 9.** Verify that the numbers in the Hardware Label field match the numbers in the Device Name field. If the numbers do not match, note the mapping of vmnic ports to vNIC ports for later use.
- Step 10.** Using the spacebar, select vmnic1.



**Note:** In lab testing, examples were seen where the vmnic and device ordering do not match. In this case, use the Consistent Device Naming (CDN) to note which vmnics are mapped to which vNICs and adjust the upcoming procedure accordingly.

**Step 11.** Press Enter.

**Step 12.** Select the VLAN (Optional) option and press Enter.

**Step 13.** Enter the <ib-mgmt-vlan-id> and press Enter.

**Step 14.** Select IPv4 Configuration and press Enter.

**Step 15.** Select the “Set static IPv4 address and network configuration” option by using the arrow keys and space bar.

**Step 16.** Move to the IPv4 Address field and enter the IP address for managing the ESXi host.

**Step 17.** Move to the Subnet Mask field and enter the subnet mask for the ESXi host.

**Step 18.** Move to the Default Gateway field and enter the default gateway for the ESXi host.

**Step 19.** Press Enter to accept the changes to the IP configuration.

**Step 20.** Select the IPv6 Configuration option and press Enter.

**Step 21.** Using the spacebar, select Disable IPv6 (restart required) and press Enter.

**Step 22.** Select the DNS Configuration option and press Enter.

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

**Step 23.** Using the spacebar, select “Use the following DNS server addresses and hostname:”

**Step 24.** Move to the Primary DNS Server field and enter the IP address of the primary DNS server.

**Step 25.** Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.

**Step 26.** Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.

**Step 27.** Press Enter to accept the changes to the DNS configuration.

**Step 28.** Press Esc to exit the Configure Management Network submenu.

**Step 29.** Press Y to confirm the changes and reboot the ESXi host.

## Procedure 6. Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional)

**Note:** By default, the MAC address of the management VMkernel port vmk0 is the same for the MAC address of the Ethernet port it is placed on. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will exist because vmk0 will retain the assigned MAC address unless the ESXi System Configuration is reset.

**Step 1.** From the ESXi console menu main screen, type Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of Static Macros.

**Step 2.** Log in as root.

**Step 3.** Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk0. vmk0 should be a part of the "Management Network" port group. Note the IP address and netmask of vmk0.

**Step 4.** To remove vmk0, type `esxcfg-vmknic -d "Management Network"`.

**Step 5.** To add vmk0 with a random MAC address, type `esxcfg-vmknic -a -i <vmk0-ip> -n <vmk0-netmask> "Management Network."`

**Step 6.** Verify vmk0 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.

**Step 7.** Tag vmk0 for the management interface by typing `esxcli network ip interface tag add -i vmk0 -t Management`.

**Step 8.** When vmk0 was added, if a message popped up saying vmk1 was marked for the management interface, type `esxcli network ip interface tag remove -i vmk1 -t Management`.

**Step 9.** If this VMware ESXi host is iSCSI booted, the vmk1, iScsiBootPG-A interface's MAC address can also be reset to a random, VMware-assigned MAC address.

**Step 10.** Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk1. vmk1 should be a part of the "iScsiBootPG-A" port group and should have a MAC address from the UCS MAC Pool. Note the IP address and netmask of vmk1.

**Step 11.** To remove vmk1, type `esxcfg-vmknic -d "iScsiBootPG-A"`.

**Step 12.** To re-add vmk1 with a random MAC address, type `esxcfg-vmknic -a -i <vmk1-ip> -n <vmk1-netmask> -m 9000 "iScsiBootPG-A"`.

**Step 13.** Verify vmk1 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.

**Step 14.** Type `exit` to log out of the command line interface.

**Step 15.** Type Ctrl-Alt-F2 to return to the ESXi console menu interface.

## **Procedure 7. Install VMware and Cisco VIC Drivers for the ESXi Host**

**Step 1.** Download the offline bundle for the Cisco UCS Tools Component and the NetApp NFS Plug-in for VMware VAAI to the Management workstation:

[Cisco UCS Tools Component for ESXi 7.0 1.1.5](#) (ucs-tool-esxi\_1.1.5-1OEM.zip)

(NetAppNasPluginV2.0.zip )

**Note:** This document describes using the driver versions shown above along with Cisco VIC nenic version 1.0.33.0 and nfnic version 4.0.0.56 along with VMware vSphere version 7.0. U3, Cisco UCS version 4.1(2a), and the latest patch NetApp ONTAP 9.10.1P1. These were the versions validated and supported at the time this document was published. This document can be used as a guide for configuring future versions of software. Consult the [Cisco UCS Hardware Compatibility List](#) and the [NetApp Interoperability Matrix Tool](#) to determine supported combinations of firmware and software.

## **Procedure 8. Install VMware VIC Drivers and the NetApp NFS Plug-in for VMware VAAI on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02**

**Step 1.** Using an SCP program such as WinSCP, copy the two offline bundles referenced above to the /tmp directory on each ESXi host.

**Step 2.** Using a ssh tool such as PuTTY, ssh to each VMware ESXi host. Log in as root with the root password.

**Step 3.** Type `cd /tmp`.

**Step 4.** Run the following commands on each host:

```
esxcli software component apply -d /tmp/ucs-tool-esxi_1.1.5-1OEM.zip
esxcli software vib install -d /tmp/NetAppNasPlugin.v23.zip
reboot
```

**Step 5.** After reboot, log back into each host and run the following commands and ensure the correct version is installed:

```
esxcli software component list | grep ucs
esxcli software vib list | grep NetApp
```

### Procedure 9. Log into the First VMware ESXi Host by Using VMware Host Client

**Step 1.** Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.

**Step 2.** Enter root for the User name.

**Step 3.** Enter the root password.

**Step 4.** Click Login to connect.

**Step 5.** Decide whether to join the VMware Customer Experience Improvement Program and click OK.

### Procedure 10. Set Up VMkernel Ports and Virtual Switch for ESXi Host VM-Host-Infra-01

**Note:** In this procedure, you're only setting up the first ESXi host. The second and third hosts will be added to vCenter and setup from the vCenter HTML5 Interface.

**Step 1.** From the Host Client Navigator, click Networking.

**Step 2.** In the center pane, click the Virtual switches tab.

**Step 3.** Highlight the vSwitch0 line.

**Step 4.** Click Edit settings.

**Step 5.** Change the MTU to 9000.

**Step 6.** Expand NIC teaming.

**Step 7.** In the Failover order section, click vmnic1 and click Mark active.

**Step 8.** Verify that vmnic1 now has a status of Active.

**Step 9.** Click Save.

**Step 10.** Click Networking, then click the Port groups tab.

**Step 11.** In the center pane, right-click VM Network and click Edit settings.

**Step 12.** Name the port group IB-MGMT Network and enter <ib-mgmt-vlan-id> in the VLAN ID field.

**Step 13.** Click Save to finalize the edits for the IB-MGMT Network.

**Step 14.** At the top, click the VMkernel NICs tab.

**Step 15.** Click Add VMkernel NIC.

**Step 16.** For New port group, enter VMkernel-Infra-NFS.

- 
- Step 17.** For Virtual switch, click vSwitch0.
- Step 18.** Enter <infra-nfs-vlan-id> for the VLAN ID.
- Step 19.** Change the MTU to 9000.
- Step 20.** Click Static IPv4 settings and expand IPv4 settings.
- Step 21.** Enter the ESXi host Infrastructure NFS IP address and netmask.
- Step 22.** Leave TCP/IP stack set at Default TCP/IP stack and do not choose any of the Services.
- Step 23.** Click Create.
- Step 24.** Click Add VMkernel NIC.
- Step 25.** For New port group, enter VMkernel-vMotion.
- Step 26.** For Virtual switch, click vSwitch0.
- Step 27.** Enter <vmotion-vlan-id> for the VLAN ID.
- Step 28.** Change the MTU to 9000.
- Step 29.** Click Static IPv4 settings and expand IPv4 settings.
- Step 30.** Enter the ESXi host vMotion IP address and netmask.
- Step 31.** Click the vMotion stack for TCP/IP stack.
- Step 32.** Click Create.
- Step 33.** Optionally, create two more vMotion VMkernel NICs to increase the speed of multiple simultaneous vMotion on this solution's 40 and 50GE vNICs:
- a. Click Add VMkernel NIC.
  - b. For New port group, enter VMkernel-vMotion1.
  - c. For Virtual switch, click vSwitch0.
  - d. Enter <vmotion-vlan-id> for the VLAN ID.
  - e. Change the MTU to 9000.
  - f. Click Static IPv4 settings and expand IPv4 settings.
  - g. Enter the ESXi host's second vMotion IP address and netmask.
  - h. Click the vMotion stack for TCP/IP stack.
  - i. Click Create.
  - j. Click Add VMkernel NIC.
  - k. For New port group, enter VMkernel-vMotion2.
  - l. For Virtual switch, click vSwitch0.
  - m. Enter <vmotion-vlan-id> for the VLAN ID.
  - n. Change the MTU to 9000.
  - o. Click Static IPv4 settings and expand IPv4 settings.
  - p. Enter the ESXi host's third vMotion IP address and netmask.
  - q. Click the vMotion stack for TCP/IP stack.
  - r. Click Create.

**Step 34.** Click the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:

**vSwitch0**

Type: Standard vSwitch

Port groups: 6

Uplinks: 2

**vSwitch Details**

MTU	9000
Ports	11776 (11763 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

**NIC teaming policy**

Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Failback	Yes

**Security policy**

Allow promiscuous mode	No
Allow forged transmits	No
Allow MAC changes	No

**Shaping policy**

Enabled	No
---------	----

**vSwitch topology**

- IB-MGMT Network (VLAN ID: 113)
- VMkernel-vMotion2 (VLAN ID: 3000)
  - VMkernel ports (1)
    - vmk4: 192.168.100.211
- VMkernel-vMotion1 (VLAN ID: 3000)
  - VMkernel ports (1)
    - vmk3: 192.168.100.201
- VMkernel-vMotion (VLAN ID: 3000)
  - VMkernel ports (1)
    - vmk2: 192.168.100.191
- VMkernel-Infra-NFS (VLAN ID: 3050)
  - VMkernel ports (1)
    - vmk1: 192.168.50.191
- Management Network (VLAN ID: 113)
  - VMkernel ports (1)
    - vmk0: 10.1.156.191

**Physical adapters**

- vnic1, 50000 Mbps, Full
- vnic0, 50000 Mbps, Full

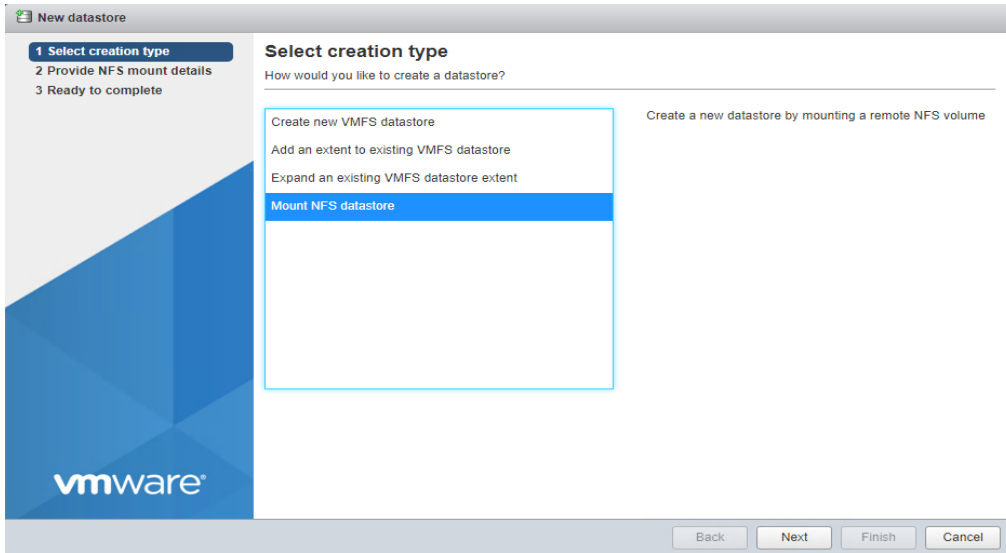
**Step 35.** Click Networking and the VMkernel NICs tab to confirm configured virtual adapters. The adapters listed should be similar to the following example:

Port groups Virtual switches Physical NICs <b>VMkernel NICs</b> TCP/IP stacks Firewall rules						
Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses	
vmk0	Management Network	Default TCP/IP stack	Management	10.1.156.191	None	
vmk1	VMkernel-Infra-NFS	Default TCP/IP stack		192.168.50.191	None	
vmk2	VMkernel-vMotion	vMotion stack	vMotion	192.168.100.191	None	
vmk3	VMkernel-vMotion1	vMotion stack	vMotion	192.168.100.201	None	
vmk4	VMkernel-vMotion2	vMotion stack	vMotion	192.168.100.211	None	

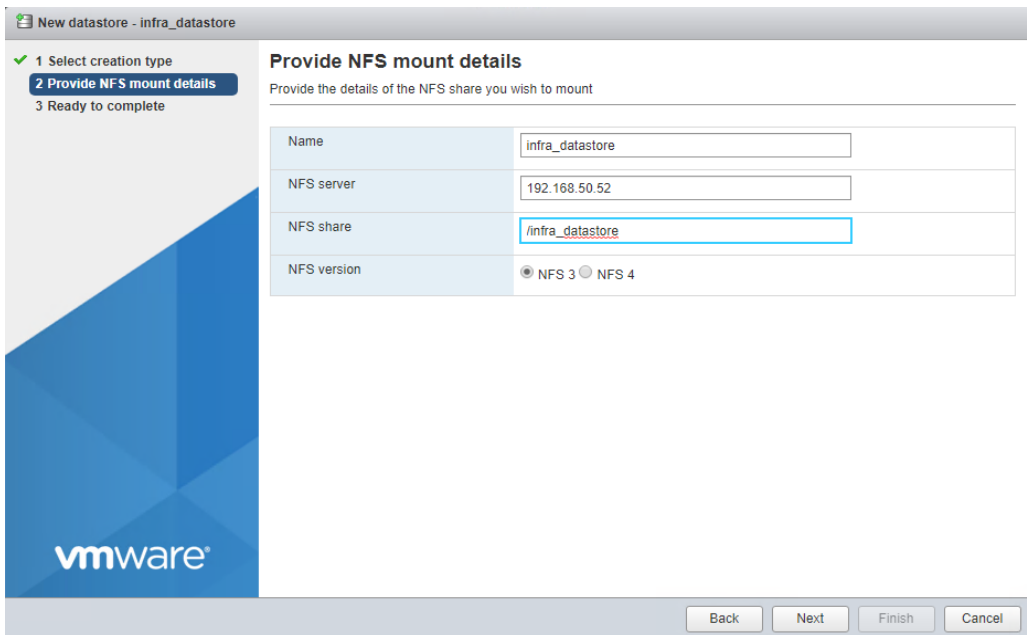
5 items

## Procedure 11. Mount Required Datastores on ESXi Host VM-Host-Infra-01

- Step 1.** From the Host Client, click Storage.
- Step 2.** In the center pane, click the Datastores tab.
- Step 3.** In the center pane, click New Datastore to add a new datastore.
- Step 4.** In the New datastore popup, click Mount NFS datastore and click Next.



**Step 5.** Input `infra_datastore` for the datastore name. Input the IP address for the `nfs-lif-02` LIF for the NFS server. Input `/infra_datastore` for the NFS share. Leave the NFS version set at NFS 3. Click Next.



**Step 6.** Click Finish. The datastore should now appear in the datastore list.

**Step 7.** In the center pane, click New Datastore to add a new datastore.

**Step 8.** In the New datastore popup, click Mount NFS datastore and click Next.

**Step 9.** Input `infra_swap` for the datastore name. Input the IP address for the `nfs-lif-01` LIF for the NFS server. Input `/infra_swap` for the NFS share. Leave the NFS version set at NFS 3. Click Next.

**Step 10.** Click Finish. The datastore should now appear in the datastore list.



Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
infra_datastore	Unknown	1,024 GB	3.85 MB	1,024 GB	NFS	Supported	Single
infra_swap	Unknown	100 GB	364 KB	100 GB	NFS	Supported	Single

## Procedure 12. Configure NTP on First ESXi Host on ESXi Host VM-Host-Infra-01

- Step 1.** From the Host Client, click Manage.
- Step 2.** In the center pane, click System > Time & date.
- Step 3.** Click Edit NTP settings.
- Step 4.** Make sure “Manually configure the date and time on this host and enter the approximate date and time.
- Step 5.** Select Use Network Time Protocol (enable NTP client).
- Step 6.** Use the drop-down list to click Start and stop with host.
- Step 7.** Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.

**Edit time configuration**

Specify how the date and time of this host should be set.

Manually configure the date and time on this host

07/22/2020 6:56 PM

Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 10.1.156.11,10.1.156.12

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

Save Cancel

- Step 8.** Click Save to save the configuration changes.

**Note:** Currently, it isn't possible to start NTP from the ESXi Host Client. NTP will be started from vCenter. The NTP server time may vary slightly from the host time.

## Procedure 13. Configure ESXi Host Swap on ESXi Host VM-Host-Infra-01

- Step 1.** From the Host Client, click Manage.
- Step 2.** In the center pane, click System > Swap.
- Step 3.** Click Edit settings.
- Step 4.** Use the drop-down list to click infra\_swap. Leave all other settings unchanged.

**Step 5.** Click Save to save the configuration changes.

#### Procedure 14. Configure Host Power Policy on ESXi Host VM-Host-Infra-01

**Note:** Implementing this policy is recommended in [Performance Tuning Guide for Cisco UCS M5 Servers](#) for maximum VMware ESXi performance. If your organization has specific power policies, please set this policy accordingly.

**Step 1.** From the Host Client, click Manage.

**Step 2.** Go to Hardware > Power Management.

**Step 3.** Click Change policy.

**Step 4.** Click High performance and click OK.

## VMware vCenter 7.0

This subject contains the following:

- [Build the VMware vCenter Server Appliance](#)
- [Adjust vCenter CPU Settings](#)
- [Set up VMware vCenter Server](#)

#### Procedure 1. Build the VMware vCenter Server Appliance

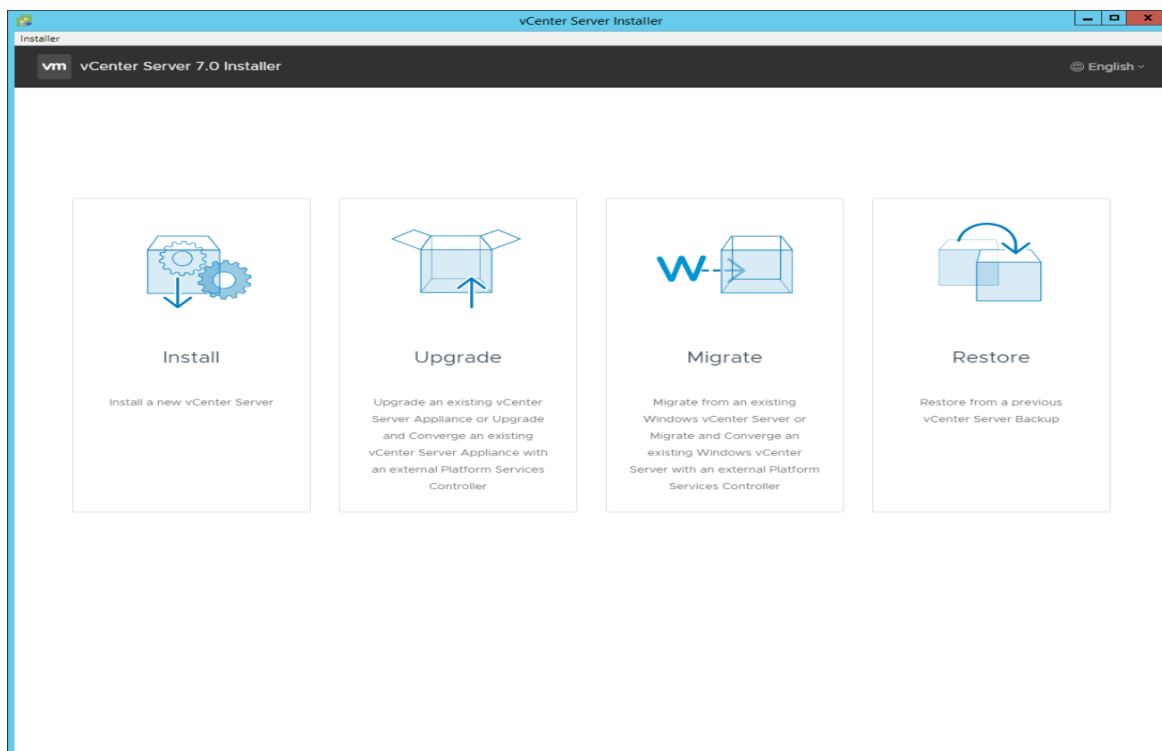
**Note:** The VCSA deployment consists of 2 stages: install and configuration.

**Step 1.** Locate and copy the VMware-VCSA-all-7.0.U3-.20150558 iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 7.0.U3 vCenter Server Appliance.

**Note:** It is important to use at minimum VMware vCenter release 7.0B to ensure access to all needed features.

**Step 2.** Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012 and above).

**Step 3.** In the mounted disk directory, navigate to the `vcasa-ui-installer > win32` directory and double-click `installer.exe`. The vCenter Server Appliance Installer wizard appears.

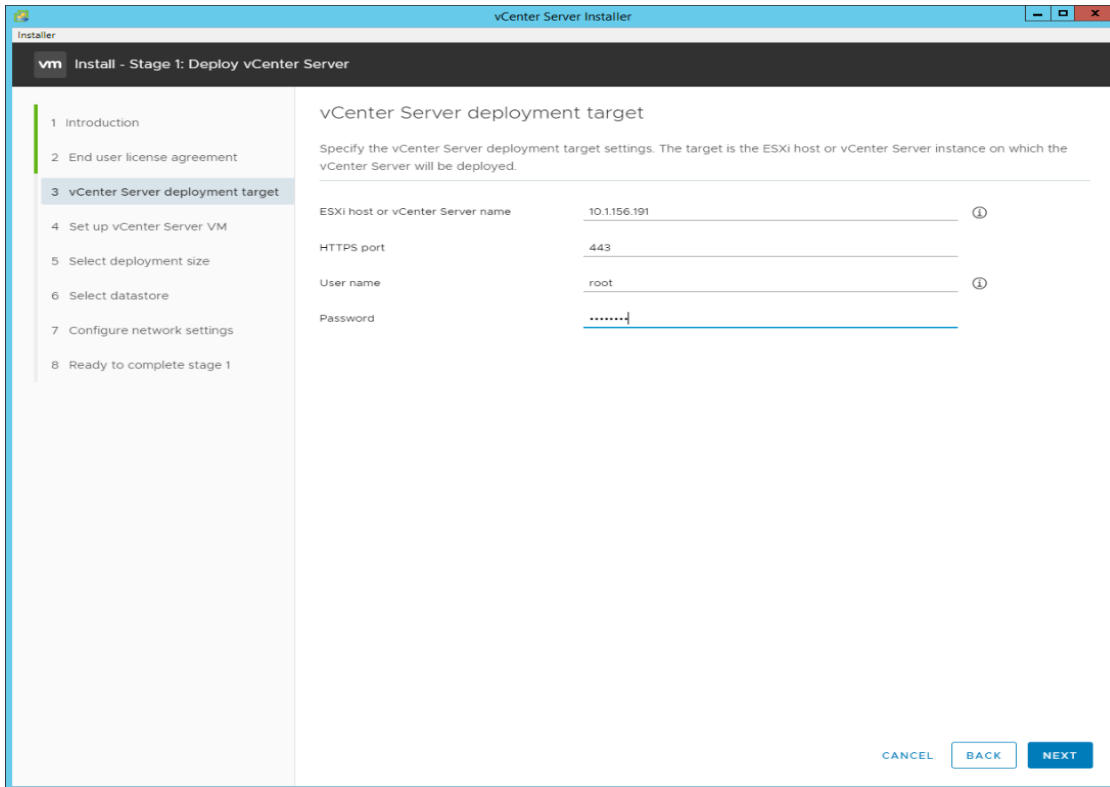


**Step 4.** Click Install to start the vCenter Server Appliance deployment wizard.

**Step 5.** Click NEXT in the Introduction section.

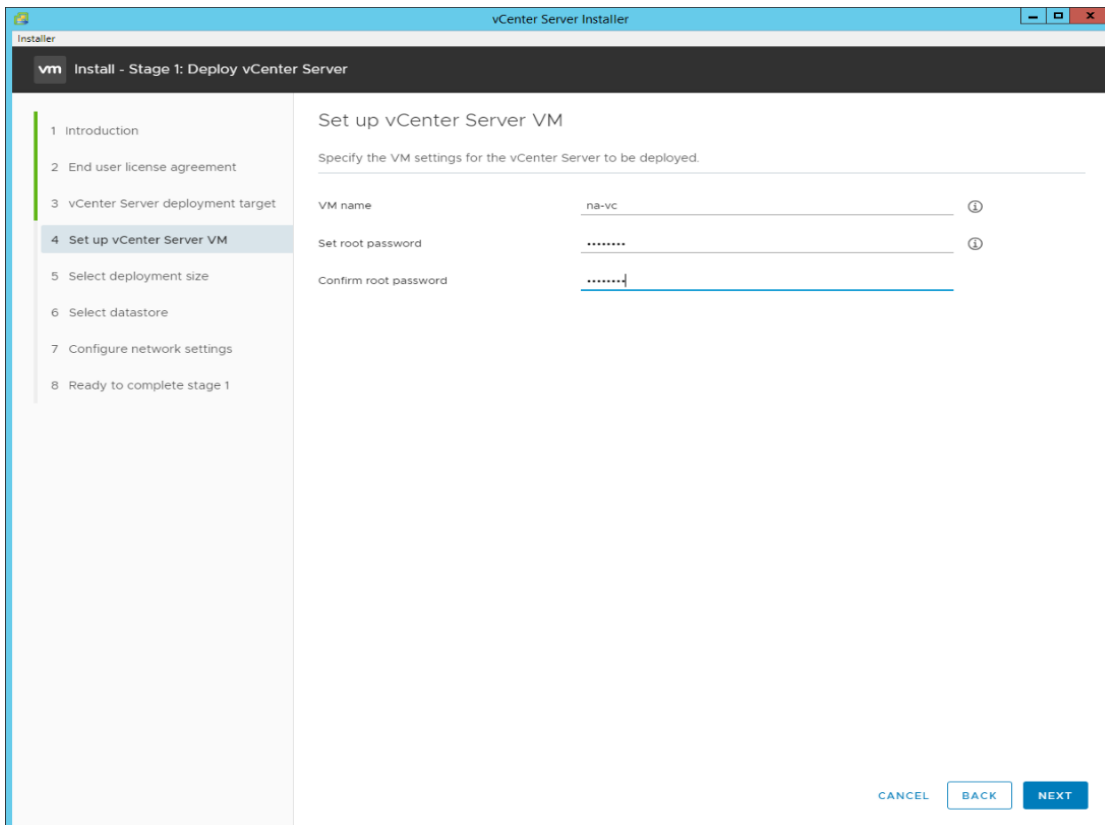
**Step 6.** Read and accept the license agreement and click NEXT.

**Step 7.** In the “vCenter Server deployment target” window, enter the host name or IP address of the first ESXi host, Username (root) and Password. Click NEXT.

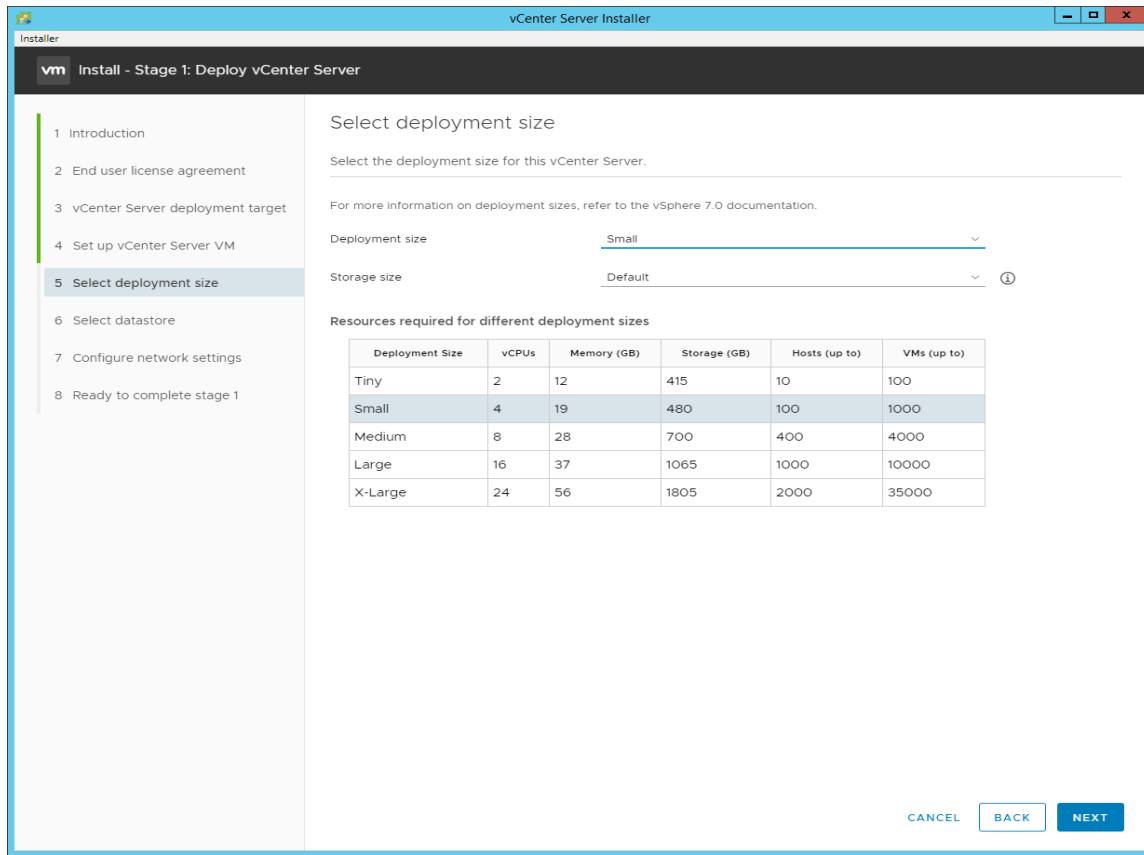


**Step 8.** Click YES to accept the certificate.

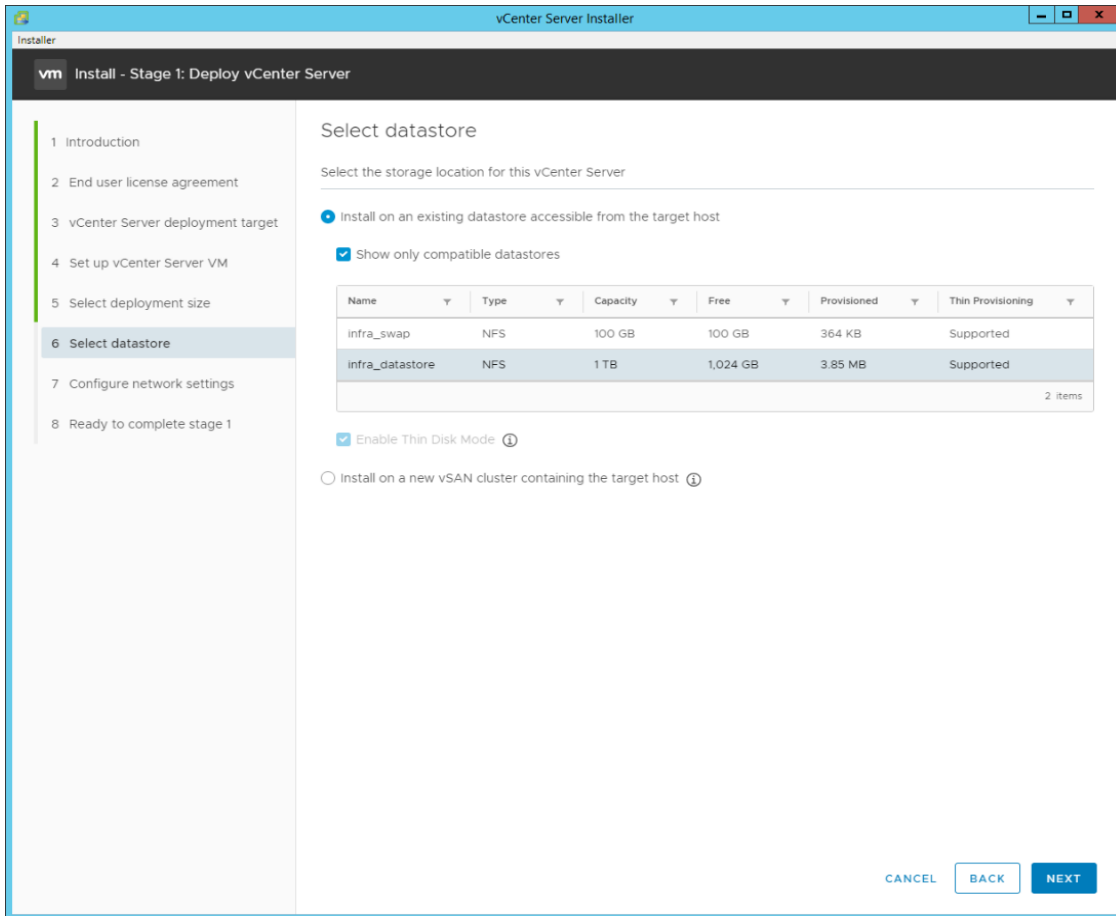
**Step 9.** Enter the Appliance VM name and password details in the “Set up vCenter Server VM” section. Click NEXT.



**Step 10.** In the “Select deployment size” section, click the Deployment size and Storage size. For example, click “Small” and “Default.” Click NEXT.



**Step 11.** Click infra\_datastore for storage. Click NEXT.

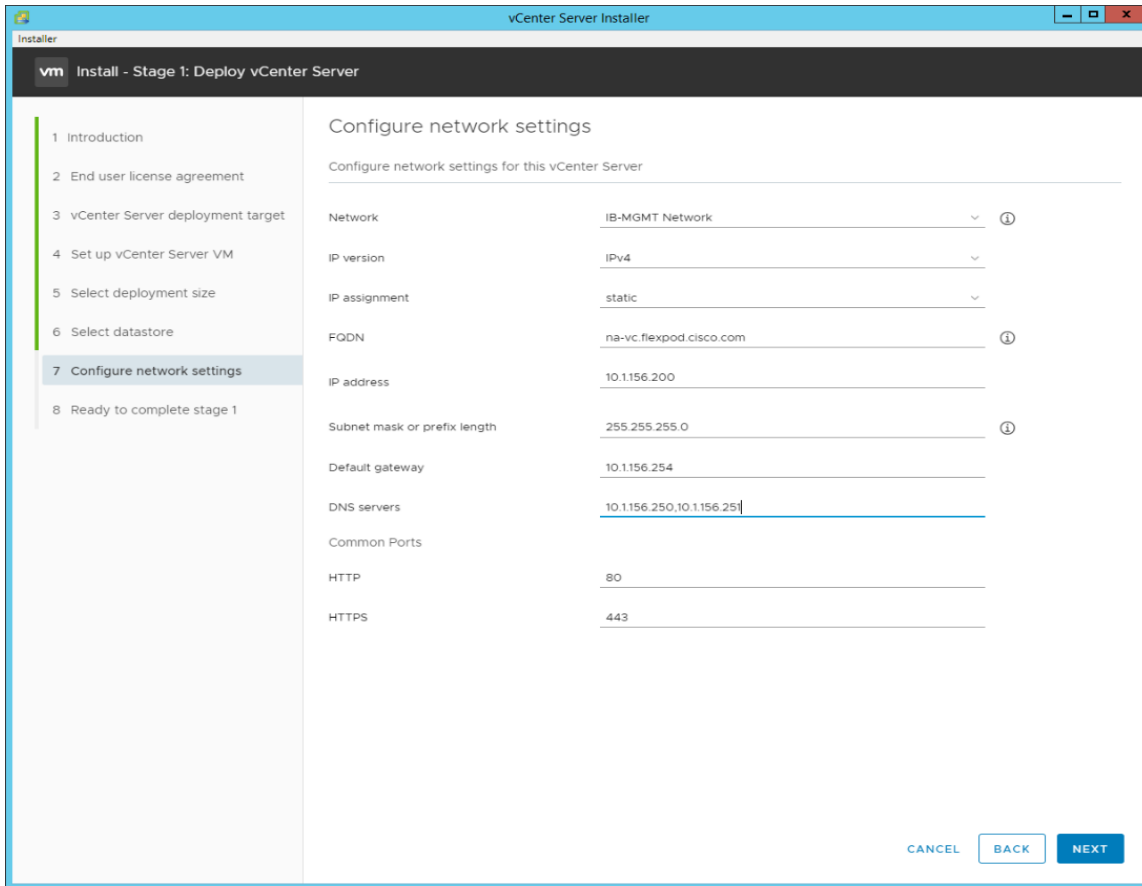


**Step 12.** In the “Network Settings” section, configure the following settings:

- a. Click a Network: IB-MGMT Network.

**Note:** It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it not get moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, and it is attempted to bring up vCenter on a different host than the one it was running on before the shutdown, vCenter will not have a functional network connection. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. Moving vCenter to a different host on vSwitch0 to be brought up always occurs correctly without requiring vCenter to already be up and running.

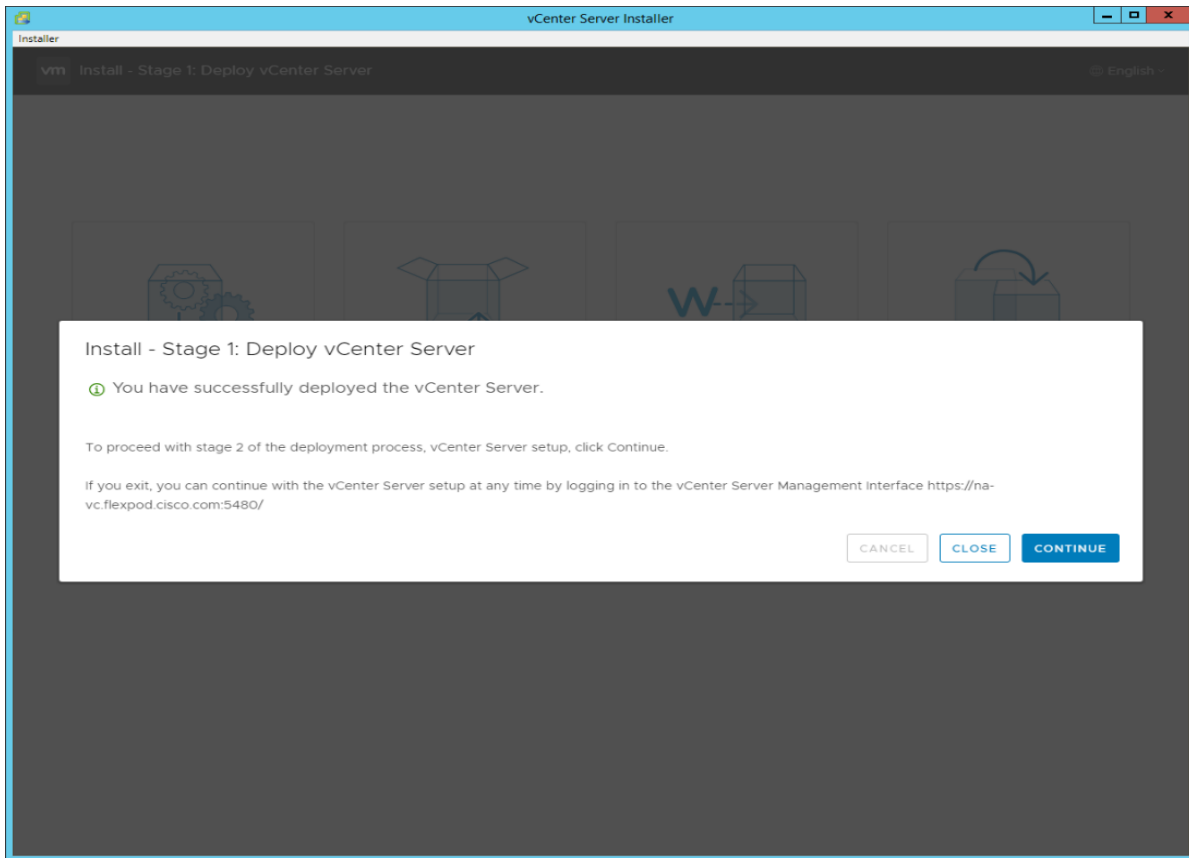
- b. IP version: IPV4
- c. IP assignment: static
- d. FQDN: <vcenter-fqdn>
- e. IP address: <vcenter-ip>
- f. Subnet mask or prefix length: <vcenter-subnet-mask>
- g. Default gateway: <vcenter-gateway>
- h. DNS Servers: <dns-server1>,<dns-server2>



**Step 13.** Click NEXT.

**Step 14.** Review all values and click FINISH to complete the installation.

**Note:** The vCenter Server appliance installation will take a few minutes to complete.



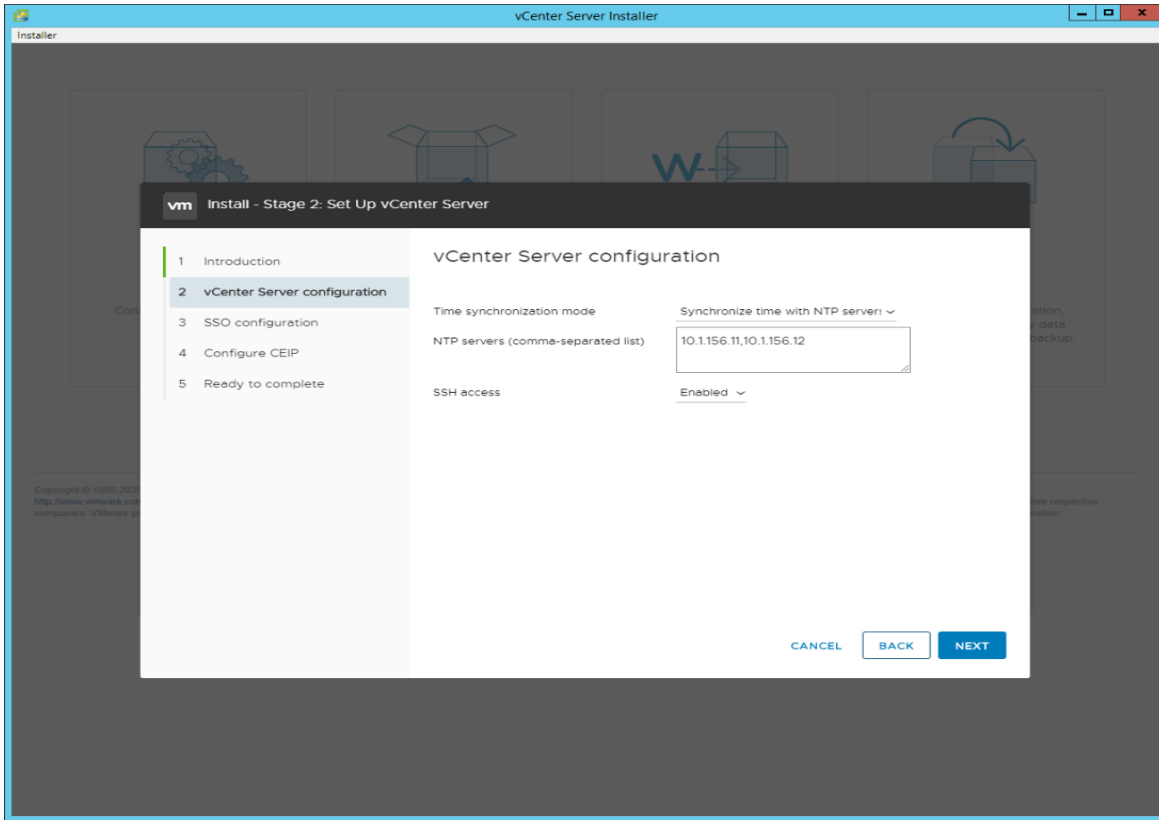
**Step 15.** Click CONTINUE to proceed with stage 2 configuration.

**Step 16.** Click NEXT.

**Step 17.** In the vCenter Server configuration window, configure these settings:

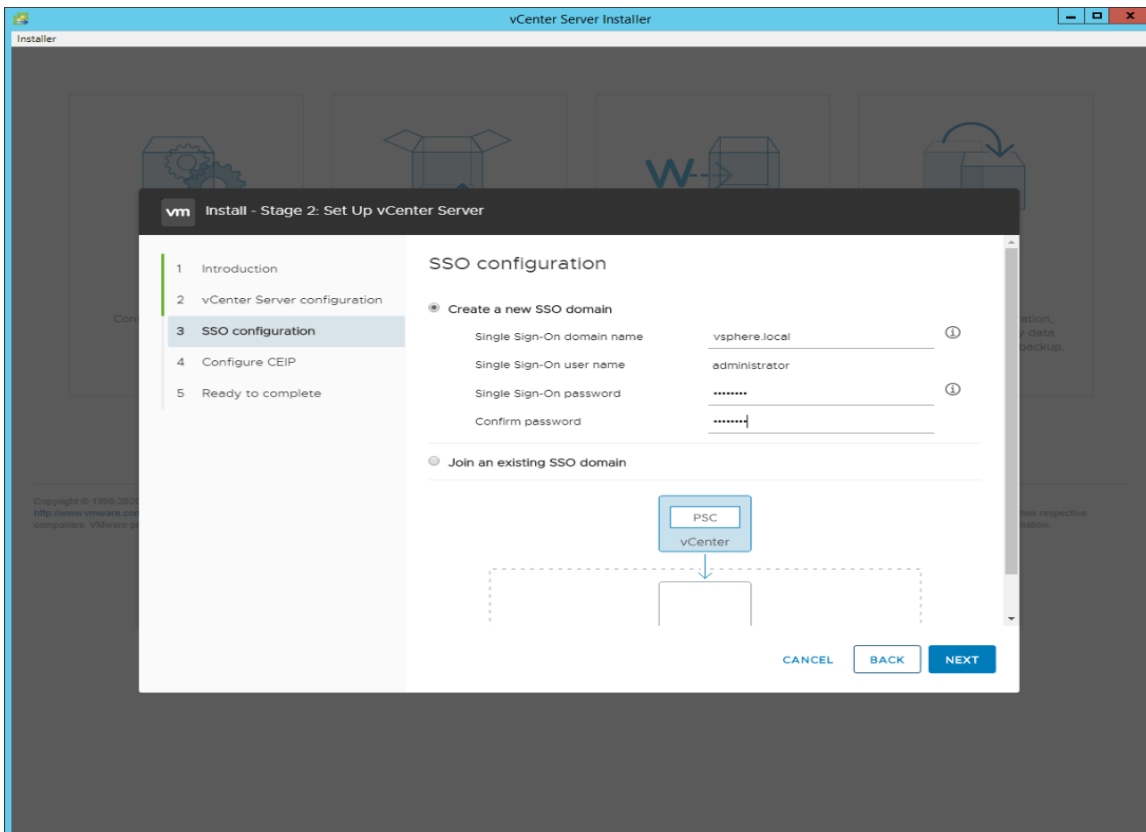
- a. Time Synchronization Mode: Synchronize time with NTP servers.
- b. NTP Servers: <nexus-a-ntp-ip>,<nexus-b-ntp-ip>
- c. SSH access: Enabled.





**Step 18.** Click NEXT.

**Step 19.** Complete the SSO configuration as shown below or according to your organization's security policies:



**Step 20.** Click NEXT.

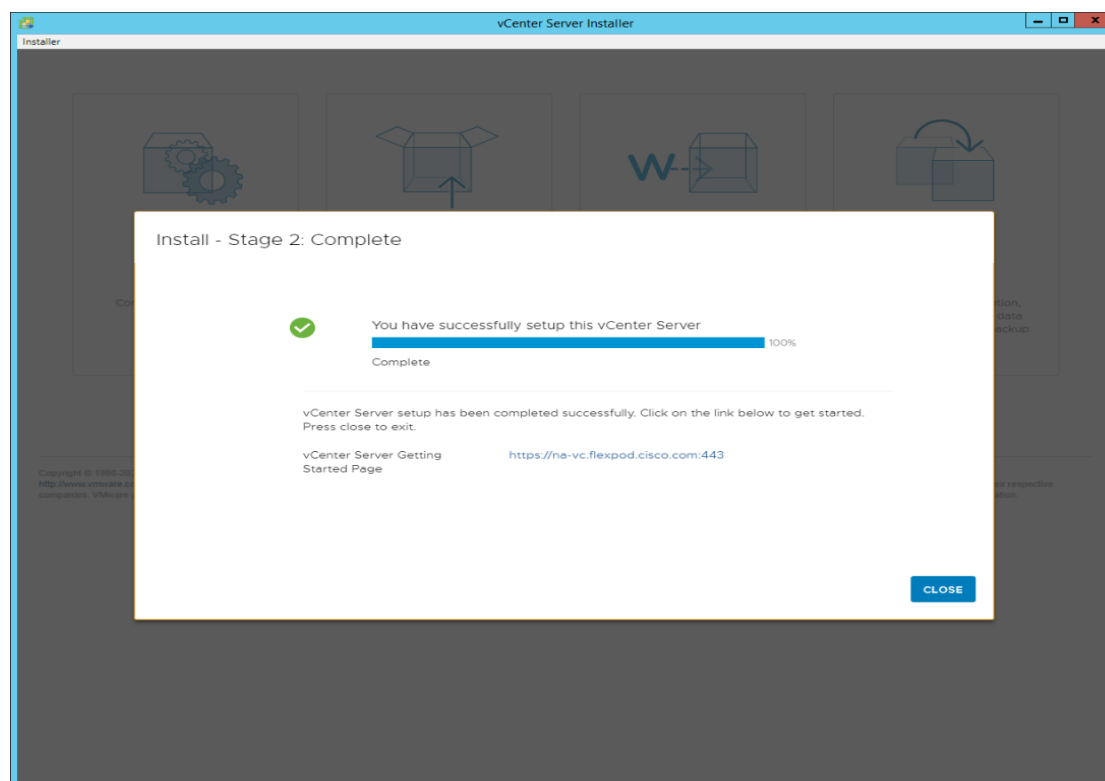
**Step 21.** Decide whether to join VMware's Customer Experience Improvement Program (CEIP).

**Step 22.** Click NEXT.

**Step 23.** Review the configuration and click FINISH.

**Step 24.** Click OK.

**Note:** The Server setup will take a few minutes to complete.



**Step 25.** Click CLOSE. Eject or unmount the VCSA installer ISO.

## Procedure 2. Adjust vCenter CPU Settings

**Note:** If a vCenter deployment size Small or Larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration. Cisco UCS B and C-Series servers are normally 2-socket servers. In this validation, the Small deployment size was selected and vCenter was setup for a 4-socket server. This setup will cause issues in the VMware ESXi cluster Admission Control. To resolve the Admission Control issue, follow these steps:

**Step 1.** Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.

**Step 2.** Enter root for the user name.


**Step 3.** Enter the root password.

**Step 4.** Click Login to connect.

**Step 5.** On the left, click Virtual Machines.

**Step 6.** In the center pane, right-click the vCenter VM and click Edit settings.

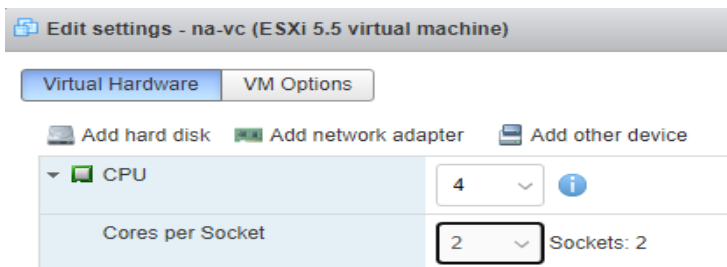
**Step 7.** In the Edit settings window, expand CPU and check the value of Sockets.



**Step 8.** If the number of Sockets does not match your server configuration, it will need to be adjusted. Click Cancel.

**Step 9.** If the number of Sockets needs to be adjusted:

- a. Right-click the vCenter VM and click Guest OS > Shut down. Click Yes on the confirmation.
- b. Once vCenter is shut down, right-click the vCenter VM and click Edit settings.
- c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to your server configuration (normally 2).



- d. Click Save.
- e. Right-click the vCenter VM and click Power > Power on. Wait approximately 10 minutes for vCenter to come up.

### Procedure 3. Set up VMware vCenter Server

**Step 1.** Using a web browser, navigate to <https://<vcenter-ip-address>:5480>.

**Step 2.** Log into the VMware vCenter Server Management interface as root with the root password set in the vCenter installation.

**Step 3.** In the menu on the left, click Time.

**Step 4.** Click EDIT.

**Step 5.** Select the appropriate Time zone and click SAVE.

**Step 6.** In the menu on the left click Administration.

**Step 7.** According to your Security Policy, adjust the settings for the root user and password.

**Step 8.** Click Update.

**Step 9.** Follow the prompts to STAGE AND INSTALL any available vCenter updates. In this validation, vCenter version 7.0.3.00700 was installed.

**Step 10.** Go to root > Logout to logout of the Appliance Management interface.

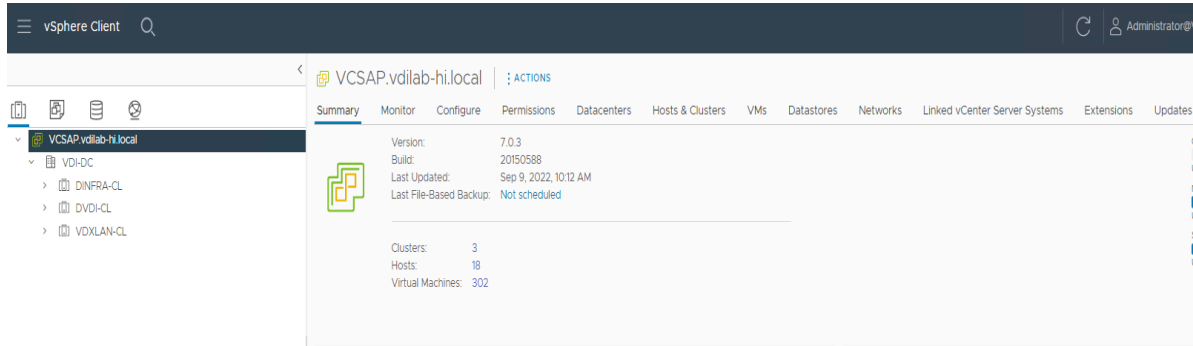
**Step 11.** Using a web browser, navigate to <https://<vcenter-fqdn>>. You will need to navigate security screens.

**Note:** With VMware vCenter 7.0.U3 the use of the vCenter FQDN is required.

**Step 12.** Click LAUNCH VSPHERE CLIENT (HTML5).

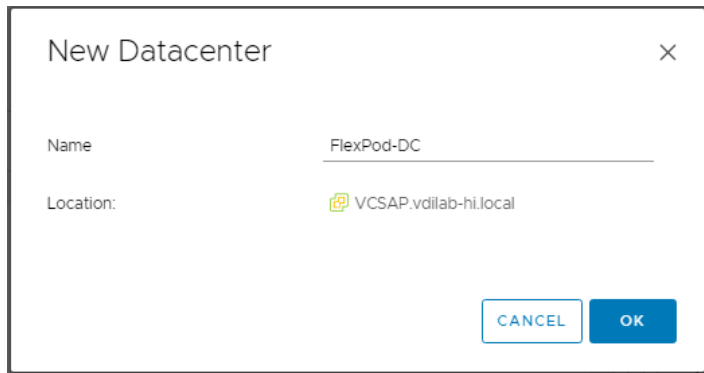
**Note:** Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is the only option in vSphere 7 and will be used going forward.

**Step 13.** Log in using the Single Sign-On username ([administrator@vsphere.local](mailto:administrator@vsphere.local)) and password created during the vCenter installation. Dismiss the Licensing warning at this time.



**Step 14.** Click ACTIONS > New Datacenter.

**Step 15.** Type “FlexPod-DC” in the Datacenter name field.



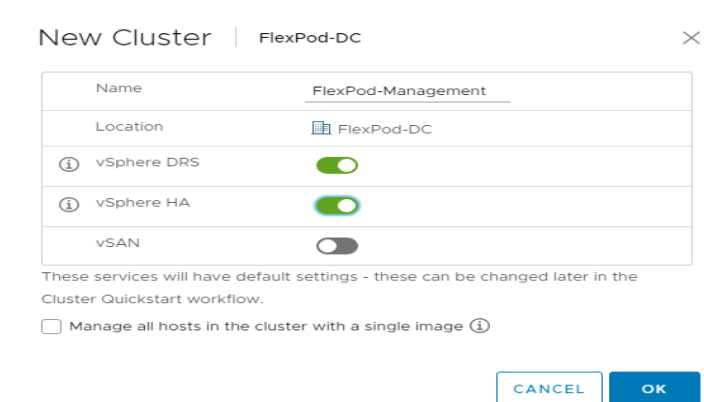
**Step 16.** Click OK.

**Step 17.** Expand the vCenter.

**Step 18.** Right-click the datacenter FlexPod-DC in the list in the left pane. Click New Cluster.

**Step 19.** Name the cluster FlexPod-Management.

**Step 20.** Turn on DRS and vSphere HA. Do not turn on vSAN.

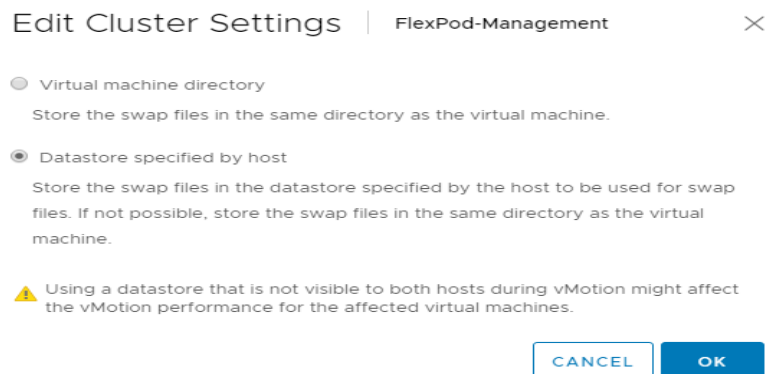


**Step 21.** Click OK to create the new cluster.

**Step 22.** Right-click “FlexPod-Management” and click Settings.

**Step 23.** Click Configuration > General in the list located on the left and click EDIT located on the right of General.

**Step 24.** Click Datastore specified by host and click OK.



**Step 25.** Right-click “FlexPod-Management” and click Add Hosts.

**Step 26.** In the IP address or FQDN field, enter either the IP address or the FQDN of the first VMware ESXi host. Enter root for the Username and the root password. Click NEXT.

**Step 27.** In the Security Alert window, click the host and click OK.

**Step 28.** Verify the Host summary information and click NEXT.

**Step 29.** Ignore warnings about the host being moved to Maintenance Mode and click FINISH to complete adding the host to the cluster.

**Step 30.** The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.

**Step 31.** In the list, right-click the added ESXi host and click Settings.

**Step 32.** In the center pane under Virtual Machines, click Swap File location.

**Step 33.** Click EDIT.

**Step 34.** Click the infra\_swap datastore and click OK.

## Edit Swap File Location | na-esxi-1.flexpod.cisco.com



Select a location to store the swap files.

Virtual machine directory

Store the swap files in the same directory as the virtual machine.

Use a specific datastore

Store the swap files in the specified datastore. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

Name	Capacity	Provisioned	Free Space	Type	Thin Provisioned
Infra_datastore	1.00 TB	504.92 GB	1,011.55 GB	NFS	Supported
Infra_swap	100.00 GB	8.42 MB	99.99 GB	NFS	Supported

2 items

CANCEL

OK

**Step 35.** In the list under System, click Time Configuration.

**Step 36.** Click EDIT to the right of Manual Time Configuration. Set the time and date to the correct local time and click OK.

**Step 37.** Click EDIT to the right of Network Time Protocol.

**Step 38.** In the Edit Network Time Protocol window, select Enable and then select Start NTP Service. Ensure the other fields are filled in correctly and click OK.

## Edit Network Time Protocol | na-esxi-1.flexpod.cisco.com



Enable

NTP Servers	10.1.156.11,10.1.156.12
Separate servers with commas, e.g. 10.31.21.2, fe00::2800	
NTP Service Status:	Stopped
	<input checked="" type="checkbox"/> Start NTP Service
NTP Service Startup Policy:	Start and stop with host

CANCEL

OK

**Step 39.** In the list under Hardware, click Overview. Scroll to the bottom and ensure the Power Management Active policy is High Performance. If the Power Management Active policy is not High Performance, to the right of Power Management, click EDIT POWER POLICY. Click High performance and click OK.

**Step 40.** In the list under Storage, click Storage Devices. Make sure the NETAPP Fibre Channel Disk LUN 0 or NETAPP iSCSI Disk LUN 0 is selected.

**Step 41.** Click the Paths tab.

**Step 42.** Ensure that 4 paths appear, two of which should have the status Active (I/O).

## Storage Devices

Refresh | 
 Attach | 
 Detach | 
 Rename... | 
 Turn On LED | 
 Turn Off LED | 
 Erase Partitions... | 
 Mark as HDD Disk | 
 Mark as Local | 
 Mark as Perennially Reserved

Name	LUN	Type	Capacity	Datastore	Operational St...	Hardware Accelerat...	Drive Ty...	Transport
Local ATA Disk (t10.ATA____Micron_5100_MTF...	0	disk	223.57 GB	Not Consu...	Attached	Not supported	Flash	Block Adapter
NETAPP Fibre Channel Disk (naa.600a09803831...	0	disk	32.00 GB	Not Consu...	Attached	Supported	Flash	Fibre Channel
Local ATA Disk (t10.ATA____Micron_5100_MTF...	0	disk	223.57 GB	Not Consu...	Attached	Not supported	Flash	Block Adapter

Copy All | 3 items

[Properties](#) | 
 [Paths](#) | 
 [Partition Details](#)

Enable | 
  Disable

Runtime Name	Status	Target	Name	Preferred
vmhba0:C0:T1:L0	◆ Active (I/O)	20:00:d0:39:ea:16:6b:8b 20:01:d0:39:ea:16:6b:8b	vmhba0:C0:T1:L0	
vmhba1:C0:T2:L0	◆ Active	20:00:d0:39:ea:16:6b:8b 20:04:d0:39:ea:16:6b:...	vmhba1:C0:T2:L0	
vmhba1:C0:T1:L0	◆ Active (I/O)	20:00:d0:39:ea:16:6b:8b 20:02:d0:39:ea:16:6b:...	vmhba1:C0:T1:L0	
vmhba0:C0:T2:L0	◆ Active	20:00:d0:39:ea:16:6b:8b 20:03:d0:39:ea:16:6b:...	vmhba0:C0:T2:L0	

## Configuration and Installation

This chapter contains the following:

- [FlexPod Automated Deployment with Ansible](#)
- [Prerequisites](#)
- [Software Infrastructure Configuration](#)
- [Horizon 8 Infrastructure Components Installation](#)

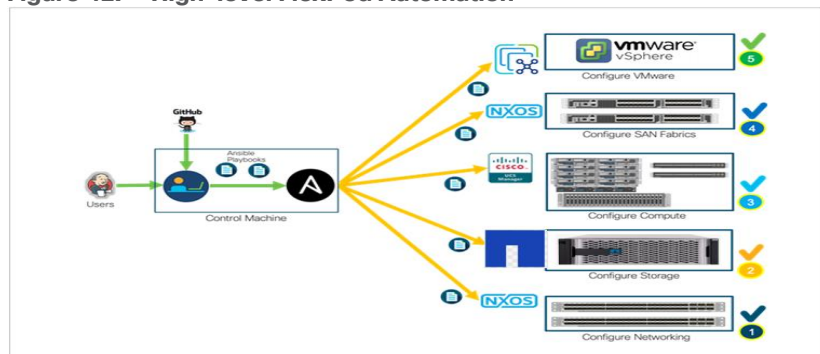
### FlexPod Automated Deployment with Ansible

If using the published Ansible playbooks to configure the FlexPod infrastructure, follow the procedures detailed in this section.

#### Ansible Automation Workflow and Solution Deployment

This FlexPod with vSphere 7.0 U3 and Cisco UCS M6 solution uses a management workstation (control machine) to run Ansible playbooks to configure Cisco Nexus, Cisco UCS, NetApp Storage, and Install VMware Cluster.

Figure 42. High-level FlexPod Automation



### Prerequisites

This subject contains the following procedure:

- [Prepare Management Workstation \(Control Machine\)](#)
- [Update Cisco VIC Drivers for ESXi](#)

Setting up the solution begins with a management workstation that has access to the internet and has a working installation of Ansible. The management workstation runs a variant of Linux or MacOS for ease of use with these command-line-based tools. Instructions for installing the workstation are not included in this document, but the basic installation and configuration of Ansible is explained. The following is a list of prerequisites:

- To use the Ansible playbooks demonstrated in this document ([Getting Started with Red Hat Ansible](#)), the management workstation must also have a working installation of Git and access to the Cisco DevNet public GitHub repository. The Ansible playbooks used in this document are cloned from the public repositories, located at the following links:
  - Cisco DevNet: <https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/Flexpod-laC-UCSM6>



◦ GitHub repository for FlexPod infrastructure setup: [GitHub - ucs-compute-solutions/FlexPod-UCSM-M6: Ansible configuration of FlexPod with UCSM 4.2\(1f\), NetApp ONTAP 9.9.1, and VMware vSphere 7.0U2](https://github.com/ucs-compute-solutions/FlexPod-UCSM-M6)

- The Cisco Nexus Switches, NetApp Storage and Cisco UCS must be physically racked, cabled, powered, and configured with the management IP addresses before the Ansible-based installation procedure can begin as shown in the cabling diagram (Figure 42). If necessary, upgrade the Nexus Switches to release 9.3(7) and the Cisco UCS System to 4.2(1f) with the default firmware packages for both blades and rack servers set to 4.2(1f).
- Before running each Ansible Playbook to setup the Network, Storage, Cisco UCS and VMware, various variables must be updated based on the customers environment and specific implementation with values such as the VLANs, pools and ports on Cisco UCS, IP addresses for iSCSI interfaces and values needed for the ESXi installation and configuration.

**Note:** Day 2 Configuration tasks such as adding datastores or ESXi servers have been performed manually or with Cisco Intersight Cloud Orchestrator (ICO) and the information has been provided in the respective sections of this document.

## Procedure 1. Prepare Management Workstation (Control Machine)

**Note:** In this section, the installation steps are performed on the CentOS management host to prepare the host for solution deployment to support the automation of Cisco UCS, Cisco Nexus, NetApp Storage and VMware installation using Ansible Playbooks.

**Step 1.** Install the EPEL repository on the management host:

```
[root@FSV-Automation ~]# yum install epel-release
```

**Step 2.** Install Ansible engine.

```
[root@FSV-Automation ~]# yum install ansible
```

**Step 3.** Verify the Ansible version to make sure it's at least release 2.9:

```
[root@FS-Automation tasks]# ansible --version
ansible 2.10.7
  config file = None
  configured module search path = ['/root/.ansible/plugins/modules', '/usr/share/ansible/plugins/modules']
  ansible python module location = /usr/local/lib/python3.6/site-packages/ansible
  executable location = /usr/local/bin/ansible
  python version = 3.6.8 (default, Aug 24 2020, 17:57:11) [GCC 8.3.1 20191121 (Red Hat 8.3.1-5)]
```

**Step 4.** Install pip the package installer for Python:

```
[root@FSV-Automation ~]# yum install python-pip
```

**Step 5.** Install the UCS SDK:

```
[root@FSV-Automation ~]# pip3 install ucmsdk
```

**Step 6.** Install the paramiko package for Cisco Nexus automation:

```
[root@FSV-Automation ~]# pip3 install paramiko
```

**Step 7.** SSH into each of the Cisco Nexus and Cisco MDS switches using Ansible so that the SSH keys are cached:

```
[root@FSV-Automation ~]# ssh admin@10.1.164.61
The authenticity of host '10.1.164.61 (10.1.164.61)' can't be established.
RSA key fingerprint is SHA256:mtomJluZVkcITgSLhVygocSnojlyPPDPmcJLQX2dfu4.
RSA key fingerprint is MD5:b4:e3:86:97:99:58:df:0d:5d:20:b2:5b:d5:69:aa:23.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.1.164.61' (RSA) to the list of known hosts.
User Access Verification
```

Password:

**Step 8.** Install the NetApp specific python module:

```
[root@FSV-Automation ~]# pip3 install netapp-lib
```

**Step 9.** Install ansible-galaxy collections for Cisco UCS, Cisco Nexus/MDS switches and NetApp Storage Array as follows:

```
[root@FSV-Automation ~]# ansible-galaxy collection install cisco.nxos
[root@FSV-Automation ~]# ansible-galaxy collection install cisco.ucs
[root@FSV-Automation ~]# ansible-galaxy collection install netapp.ontap
[root@FSV-Automation ~]# ansible-galaxy collection install community.vmware
```

**Note:** We validated the Ansible automation with both python 2.7.5 and python 3.6 as the python interpreter for Ansible.

## Procedure 2. Update Cisco VIC Drivers for ESXi

**Note:** When ESXi is installed from Cisco Custom ISO, you might have to update the Cisco VIC drivers for VMware ESXi Hypervisor to match the current [Cisco Hardware and Software Interoperability Matrix](#).

In this Validated Design the following drivers were used:

- Cisco-nenic- 1.0.33.0
- Cisco-nfnic- 4.0.0.56

**Step 1.** Log into your VMware Account to download required drivers for FNIC and NENIC as per the recommendation.

**Step 2.** Enable SSH on ESXi to run following commands:

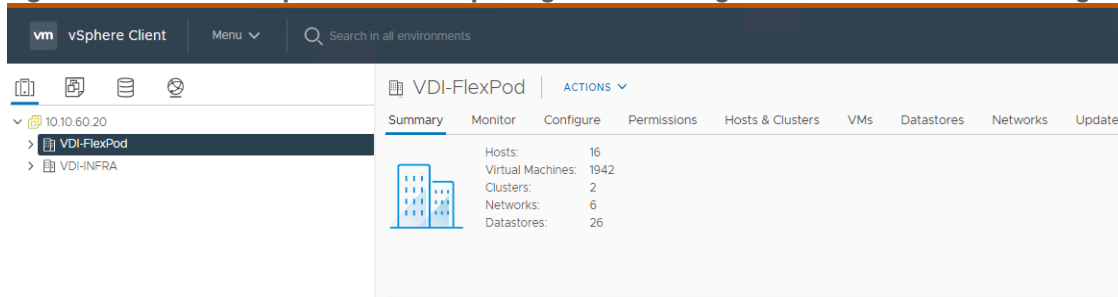
```
esxcli software vib update -d /path/offline-bundle.zip
```

## VMware Clusters

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

- Datacenter: FlexPod - NetApp Storage AFF A400 with Cisco UCS
- Cluster: FlexPod-VDI - Single-session/Multi-session OS VDA workload
- Infrastructure Cluster: Infrastructure virtual machines (vCenter, Active Directory, DNS, DHCP, SQL Server, VMware Horizon Connection Servers and other common services), Login VSI launcher infrastructure were connected using the same set of switches.

**Figure 43. VMware vSphere WebUI Reporting Cluster Configuration for this Validated Design**

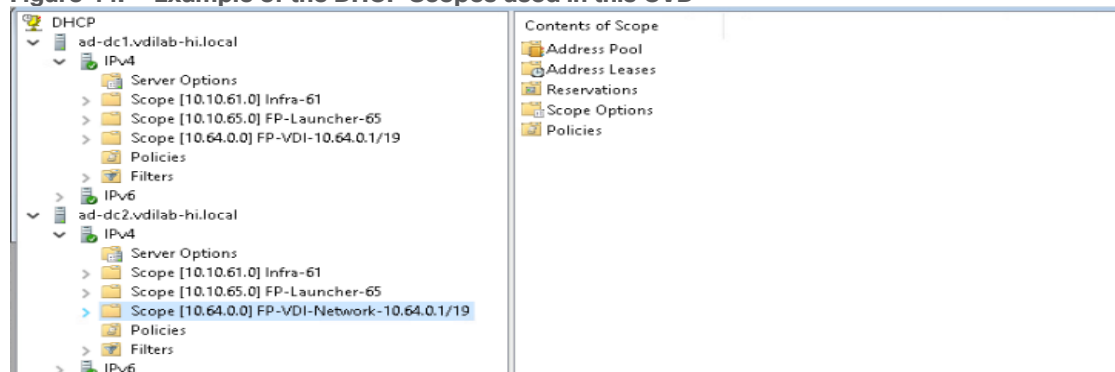


## Build the Virtual Machines and Environment for Workload Testing

### Prerequisites

Create all necessary DHCP scopes for the environment and set the Scope Options.

**Figure 44. Example of the DHCP Scopes used in this CVD**



### Software Infrastructure Configuration

This section explains how to configure the software infrastructure components that comprise this solution.

Install and configure the infrastructure virtual machines by following the process listed in [Table 24](#).

**Table 24. Test Infrastructure Virtual Machine Configuration**

Configuration	Key Management Server	VMware Horizon Replica Server Virtual Machines
Operating system	Microsoft Windows Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	4	4
Memory amount	8 GB	8 GB
Network	VMXNET3 k21-Infra-Mgmt-61	VMXNET3 k21-Infra-Mgmt-61
Disk-1 (OS) size	60 GB	60 GB
Disk-2 size	-	-
Operating system	Microsoft Windows Server 2019	VCSA - SUSE Linux
Virtual CPU amount	4	16
Memory amount	8 GB	32 GB
Network	VMXNET3 k21-Infra-Mgmt-61	VMXNET3 k21-InBand-Mgmt-60
Disk size	60 GB	

Configuration	Microsoft SQL Server Virtual Machine	VMware Horizon Connection Server Virtual Machine
Operating system	Microsoft Windows Server 2019 Microsoft SQL Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	4	4
Memory amount	12GB	12 GB
Network	VMXNET3 k21-Infra-Mgmt-61	VMXNET3 k21-Infra-Mgmt-61
Disk-1 (OS) size	60 GB	60 GB
Disk-2 size	100 GB SQL Databases\Logs	-

**Note:** The additional Horizon Replica servers, Microsoft Key Management Server and additional SQL server for database logs will be configured similar RAM/CPU. The Amount of RAM (Gb) and Virtual CPU for each Infrastructure is subjected to adjust based on amount of load the virtual machine generates

## Horizon 8 Infrastructure Components Installation

This subject contains the following:

- [Install Horizon Connection and Replica Servers](#)
- [Create a Microsoft Management Console Certificate Request](#)
- [Configure the Horizon 8 Environment](#)
- [Configure Event Database](#)
- [Configure Horizon 8 Licenses](#)
- [Configure vCenter](#)
- [Configure Instant Clone Domain Admins](#)

**Note:** The prerequisites for installing the view connection server, replica server(s) and composer server is to have Windows 2012, 2016 or 2019 virtual machines ready.

**Note:** In this study, we used Windows Server 2019 virtual machines for all Horizon infrastructure servers and Other Infrastructure VMs.

**Note:** This following section provides a detailed, systematic installation process for VMware Horizon 2209 or VMware Horizon 8.7.0

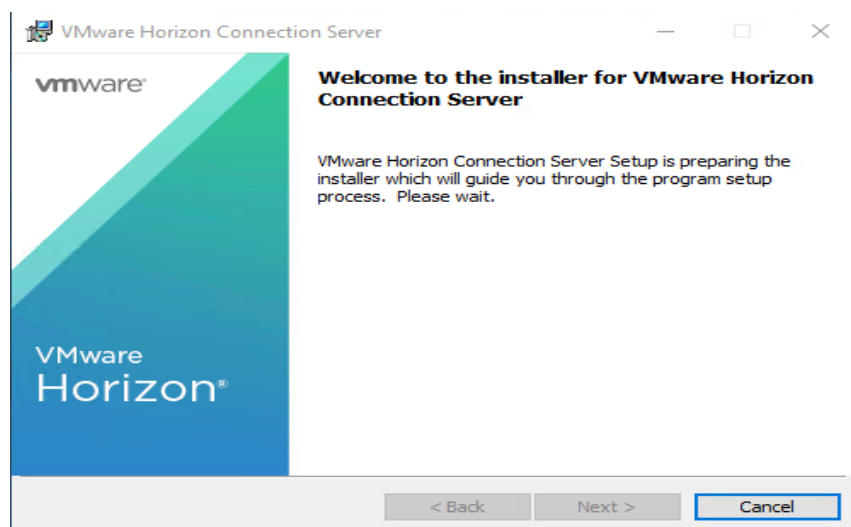
### Procedure 1. Install Horizon Connection and Replica Servers

**Step 1.** Download the VMware Horizon 8 installation package from this link:

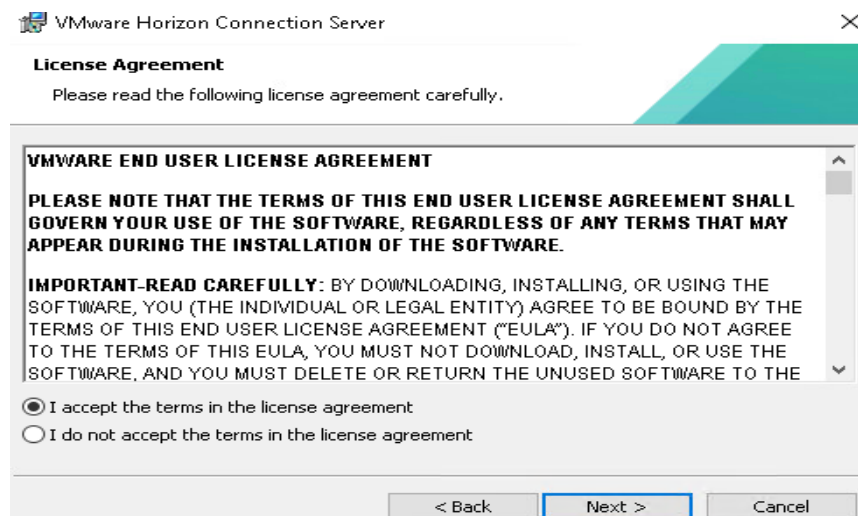
[https://customerconnect.vmware.com/downloads/info/slug/desktop\\_end\\_user\\_computing/vmware\\_horizon/2209](https://customerconnect.vmware.com/downloads/info/slug/desktop_end_user_computing/vmware_horizon/2209)

**Step 2.** Open view connection server installation, VMware-viewconnectionserver-x86\_64-8.7.0-20649599.exe.


**Step 3.** Click Next.



**Step 4.** Accept the EULA then click Next.

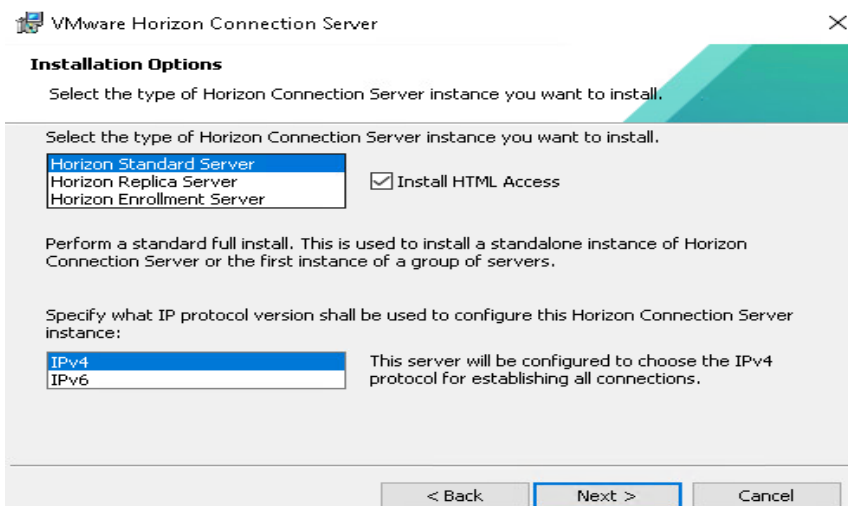


**Step 5.** Keep the default destination folder and click Next.

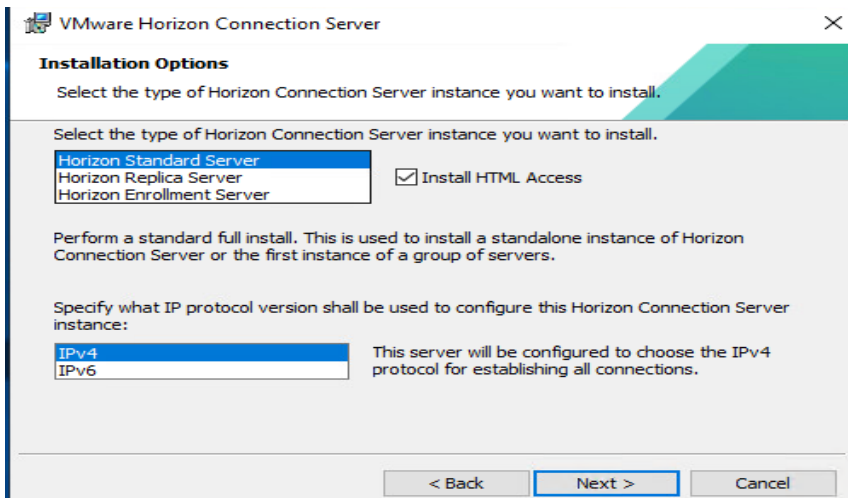


**Step 6.** Select type of instance intended to install.

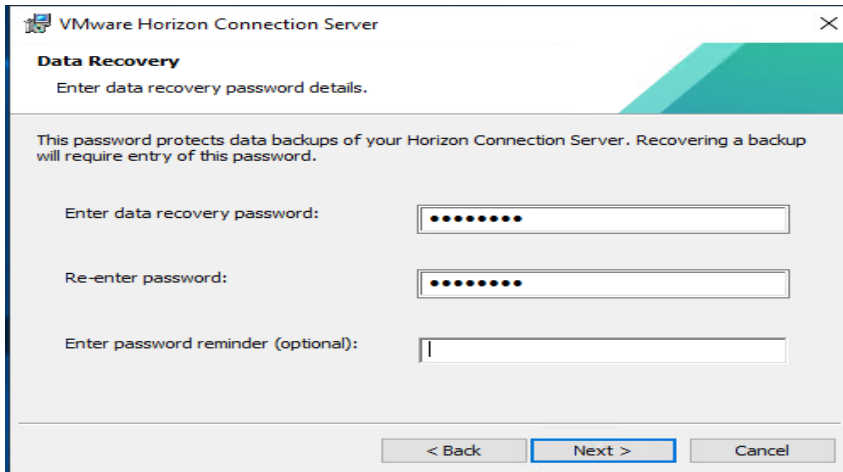
**Step 7.** Select Standard Server instance for primary connection server installation.



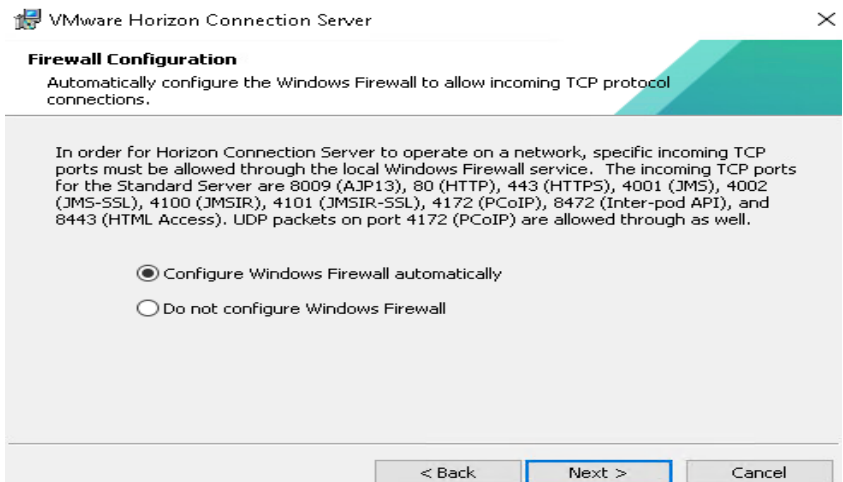
**Step 8.** Select Replica server instance for fault tolerant connection server configuration after completion of Standard Server instance installation.



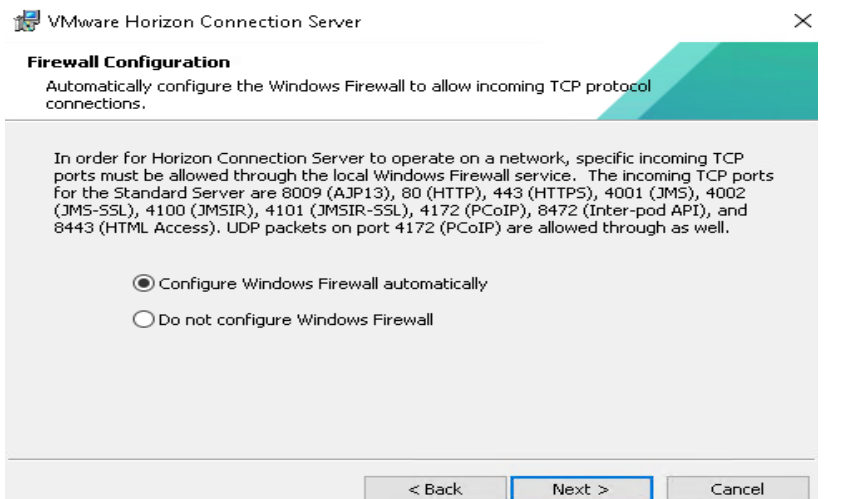
**Step 9.** Enter the Data Recovery Password.



**Step 10.** Click Next.

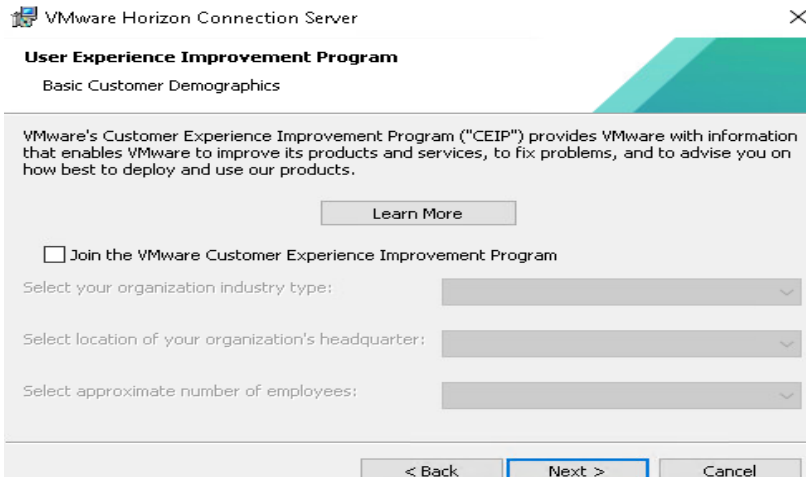


**Step 11.** Select authorized users and group, click Next.

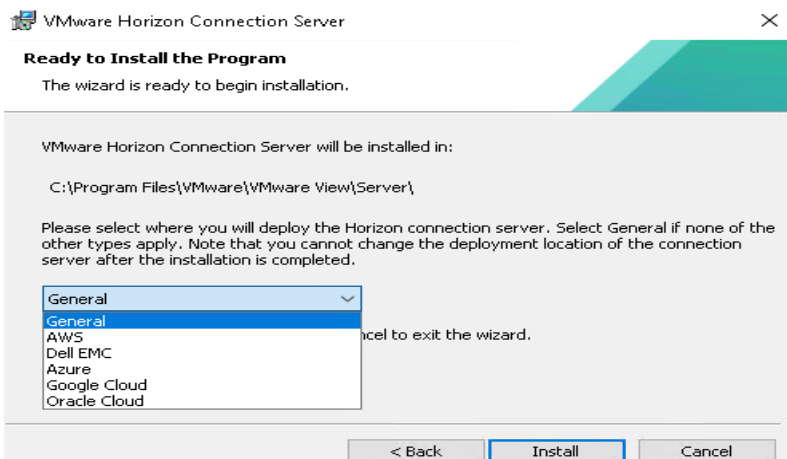


**Step 12.** Enter domain credentials for previously specified domain user/group.

**Step 13.** Opt-in or Opt-out of User Experience Improvement Program. Click Next.



**Step 14. Click Install.**



**Step 15. Click Finish.**



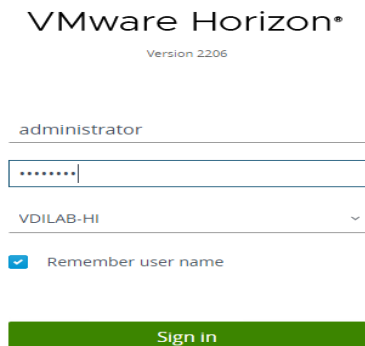
**Procedure 2. Create a Microsoft Management Console Certificate Request**

**Step 1.** To generate a Horizon View SSL certificate request, use the Microsoft Management Console (MMC) Certificates snap-in:  
[https://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=206866](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=206866)



### Procedure 3. Configure the Horizon 8 Environment

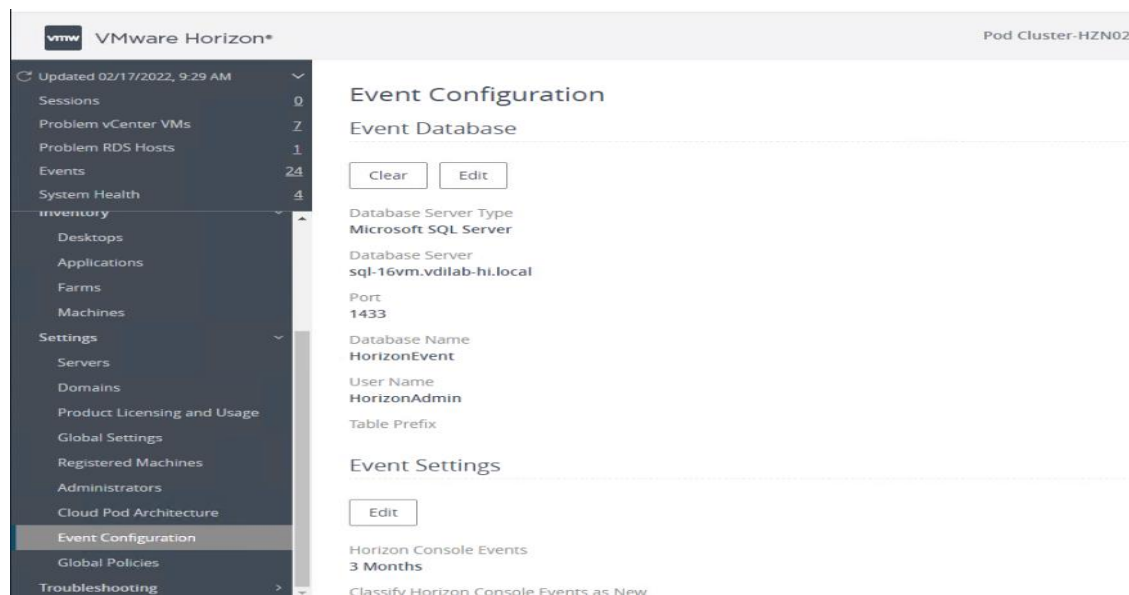
**Step 1.** Open WebUI, Login to [https://<Horizon\\_Connection\\_server\\_Management\\_IP\\_Address>/admin](https://<Horizon_Connection_server_Management_IP_Address>/admin).



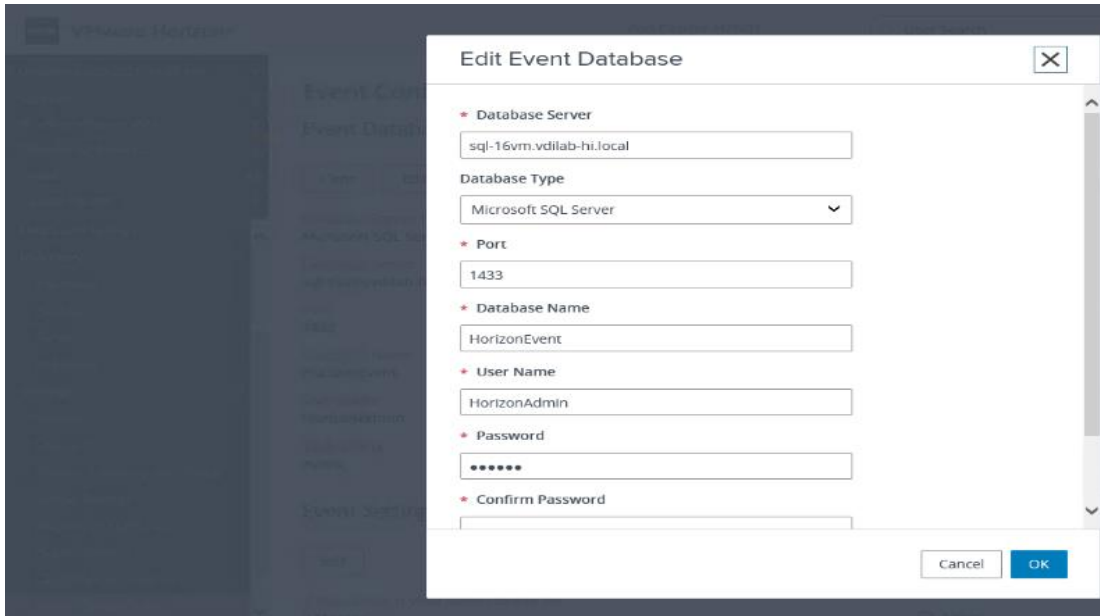
### Procedure 4. Configure Event Database

**Step 1.** Configure the Event Database by adding Database Server, Database name, login credentials and prefix for the table from the Horizon 8 Administrator, View Configuration, Event Configuration node of the Inventory pane.

**Step 2.** Click Edit in the action pane.

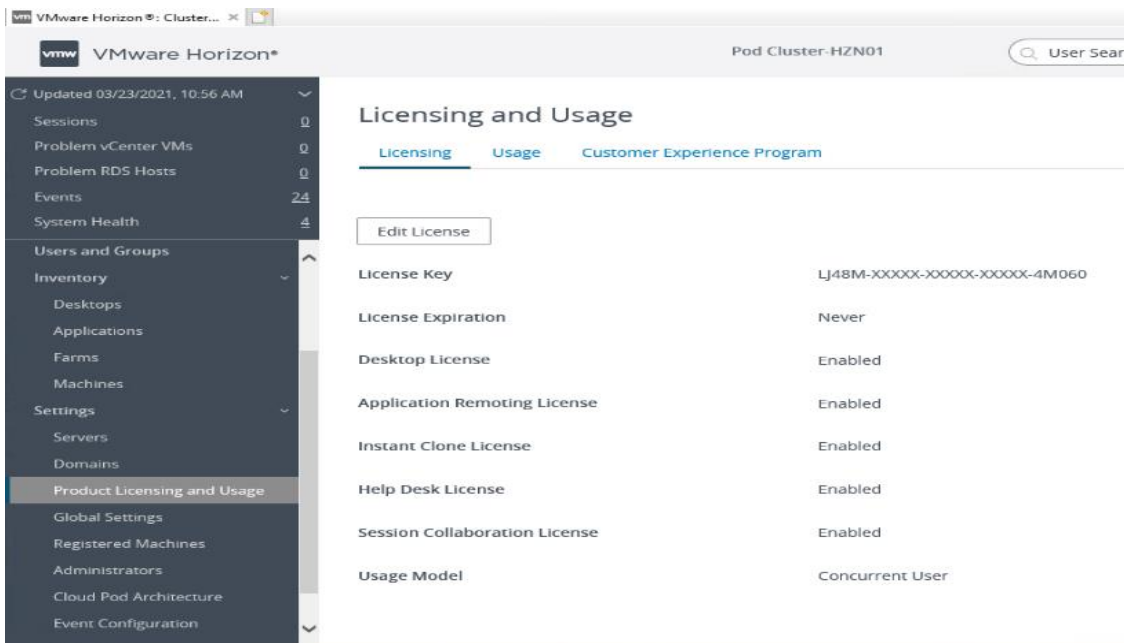


The details are shown below:



## Procedure 5. Configure Horizon 8 Licenses

- Step 1.** Click View Configuration.
- Step 2.** Select Product Licensing and Usage.
- Step 3.** Click Edit License in the action pane.
- Step 4.** Add the License Serial Number.
- Step 5.** Click OK.



## Procedure 6. Configure vCenter

- Step 1.** In View Configuration, Select Servers. Click Add vCenter Server tab.
- Step 2.** Enter vCenter Server IP Address or FQDN, login credentials.

**Step 3.** Advanced Settings options can be modified to change existing operations limit. Keep the advanced settings options as default.

The screenshot shows a configuration window titled "Edit vCenter Server - VCSA.VDILAB-HI.LOCAL". At the top, a red asterisk (\*) denotes a required field. The fields are as follows:

- Server address:** VCSA.VDILAB-HI.LOCAL
- User Name:** administrator@vsphere.local
- Password:** masked with seven dots
- Description:** VC for FP
- SSL:** checked
- Port:** 443

At the bottom, there is a "Deployment Type" label and two buttons: "Cancel" and "OK".

**Step 4.** Click View certificate. Accept the certificate and click to complete adding vCenter.

**Step 5.** Click Next.

### Procedure 7. Configure Instant Clone Domain Admins

**Step 1.** Under View Configuration, Click on Instant Clone Domain Admins.

**Step 2.** Click Add. Enter credentials for domain user/group.

The screenshot shows a dialog box titled "Edit Domain Admin" overlaid on a dark background. The fields are as follows:

- Full domain name:** vdilab-hi.local
- Username:** administrator
- Password:** masked with seven dots

At the bottom right, there are two buttons: "Cancel" and "OK".

## Master Image Creation for Tested Horizon Deployment Types

This chapter contains the following:

- [Create the Master Image for the tested Horizon Deployment Types](#)
- [Prepare Microsoft Windows 10 and Server 2019 R2 with Microsoft Office 2016](#)
- [Optimize Base Windows 10 or Server 2019 Guest OS](#)
- [Install the Virtual Desktop Agent Software Installation for VMware Horizon](#)
- [Install Additional Software for testing End User Experience](#)
- [Create a Native Snapshot for Automated Desktop Pool Creation](#)
- [Create Customization Specification for Virtual Desktops](#)
- [Create RDSH Farm](#)
- [Create the Horizon 8 Remote Desktop Server Hosted \(RDSH\) Sessions Published Desktop Pool](#)
- [Create VMware Horizon Instant Clone and Full Clone Persistent Windows 10 Desktop Pool](#)
- [VMware Horizon Persistent Windows 10 Desktop Pool Creation](#)
- [Configure FSLogix for VMware Remote Desktop Session Host \(RDSH\) Server Sessions & Windows 10 Virtual Desktops Profiles Profile Container](#)
- [Agent Installation](#)

### Procedure 1. Create the Master Image for the tested Horizon Deployment Types

**Step 1.** Select an ESXi host in an existing infrastructure cluster and create the virtual machines to use as Golden Images with Windows 10 and Office 2016 for Instant Clone, and Full Clone desktops.

**Note:** We used a 64-bit version of Microsoft Operating System and Office for our testing.

**Note:** A third master image has been created using Microsoft Windows Server 2019 for Remote Desktop Server Sessions (RDSH) server session virtual machines.

[Table 25](#) lists the parameters use for Master Image virtual machines.

**Table 25. Golden Image Virtual Machine Parameters**

Attribute	Instant / Non-Persistent Clone	Persistent / Full Clone	RDSH server
Desktop operating system	Microsoft Windows 10 Enterprise (64-bit)	Microsoft Windows 10 Enterprise (64-bit)	Microsoft Windows Server 2019 standard (64-bit)
Hardware	VMware Virtual Hardware Version 13	VMware Virtual Hardware Version 13	VMware Virtual Hardware Version 13
vCPU	2	2	4
Memory	3.5 Gb	3.5 Gb	24GB
Memory reserved	3.5 Gb	3.5 Gb	

Attribute	Instant / Non-Persistent Clone	Persistent / Full Clone	RDSH server
Video RAM	35 MB/Auto detect	35 MB/Auto detect	4MB/Auto detect
3D graphics	Off	Off	Off
NIC	1	1	1
Virtual network adapter 1	VMXNet3 adapter	VMXNet3 adapter	VMXNet3 adapter
Virtual SCSI controller 0	Paravirtual	Paravirtual	Paravirtual
Virtual disk: VMDK 1	40 GB	80 GB	80 GB
Virtual floppy drive 1	Removed	Removed	Removed
Virtual CD/DVD drive 1	-	-	-
Applications	Login VSI 4.1.40 application installation Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Microsoft Internet Explorer Microsoft Office 2016	Login VSI 4.1.40 application installation Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Microsoft Internet Explorer Microsoft Office 2016	Login VSI 4.1.40 application installation Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Microsoft Internet Explorer Microsoft Office 2016
VMware tools	Release 13325	Release 13325	Release 13325
VMware View Agent	Release 8.7.0-20606795 (Version 2209)	Release 8.7.0-20606795 (Version 2209)	Release 8.7.0-20606795 (Version 2209)
Attribute	Instant-clone	Persistent/Full Clone	RDSH Remote Server Sessions

\* For Persistent Desktops, we configured 3.5 GB of RAM as amount of memory allocated is enough to run LoginVSI Knowledge Worker workload. FlexPod ESXi nodes and compute-only node were configured with 1TB of total memory for this performance study.

## Prepare Microsoft Windows 10 and Server 2019 R2 with Microsoft Office 2016

Prepare your master image for one or more of the following use cases:

- VMware Horizon 2209 Remote Desktop Server Sessions (RDSH) Server 2019 Virtual Machines
- VMware Horizon 2209 Instant Clone non-persistent virtual machines
- VMware Horizon 2209 Persistent full clone virtual machines

---

Include Microsoft Office 2016 and other applications used by all pool users in your organization into your master image.

Apply the required Microsoft updates and patches to your master images.

**Note:** For this study, we added Login VSI target software to enable the use the Login VSI Knowledge Worker workload to benchmark end user experience for each use case.

## Optimize Base Windows 10 or Server 2019 Guest OS

This subject contains the following procedures:

- [Install the Virtual Desktop Agent Software Installation for Horizon](#)
- [Install Additional Software](#)
- [Create a Native Snapshot for Automated Desktop Pool Creation](#)
- [Create Customization Specification for Virtual Desktops](#)
- [Create RDSH Farm](#)
- [Create the Horizon 8 Remote Desktop Server Hosted \(RDSH\) Sessions Published Desktop Pool](#)
- [Create VMware Horizon Instant Clone and Full Clone Persistent Windows 10 Desktop Pool](#)
- [VMware Horizon Persistent Windows 10 Desktop Pool Creation](#)
- [Configure FSLogix for VMware Remote Desktop Session Host \(RDSH\) Server Sessions & Windows 10 Virtual Desktops Profiles Profile Container](#)
- [Agent Installation](#)

Click the links below for information about how to optimize windows 10 for VDI deployment:

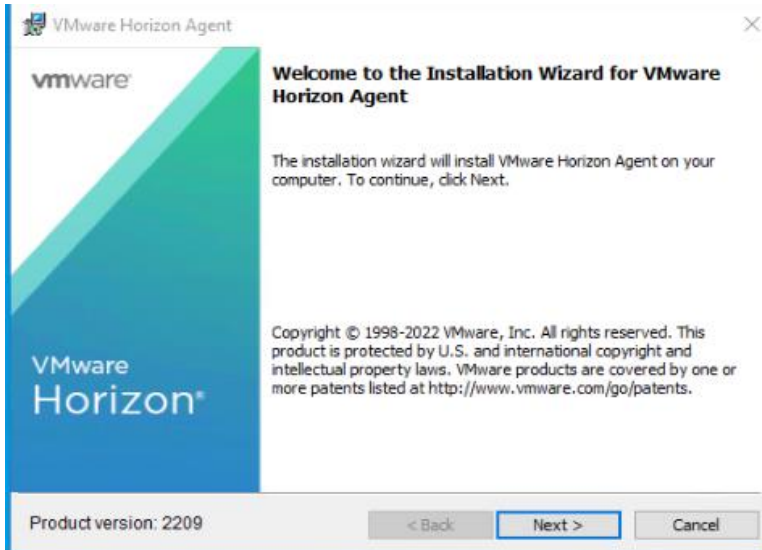
VMware Windows Operating System Optimization Tool Guide:

<http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/vmware-view-optimizationguidewindows7-en-white-paper.pdf>

VMware Optimization Tool for HVD or HSD Deployment: <https://labs.vmware.com/flings/vmware-os-optimization-tool>

### Procedure 1. Install the Virtual Desktop Agent Software Installation for Horizon

**Step 1.** For each master image created, open the Horizon View Agent Installer, VMware-viewagent-2209 or version 8.7.0-20606795.exe Click Next to install.

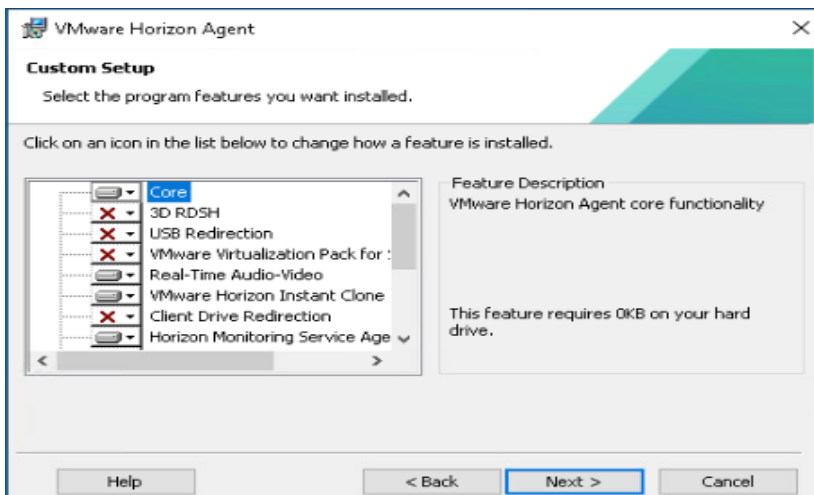


**Step 2.** Review and accept the EULA Agreement. Click Next.

**Step 3.** Select Network protocol configuration, click Next.

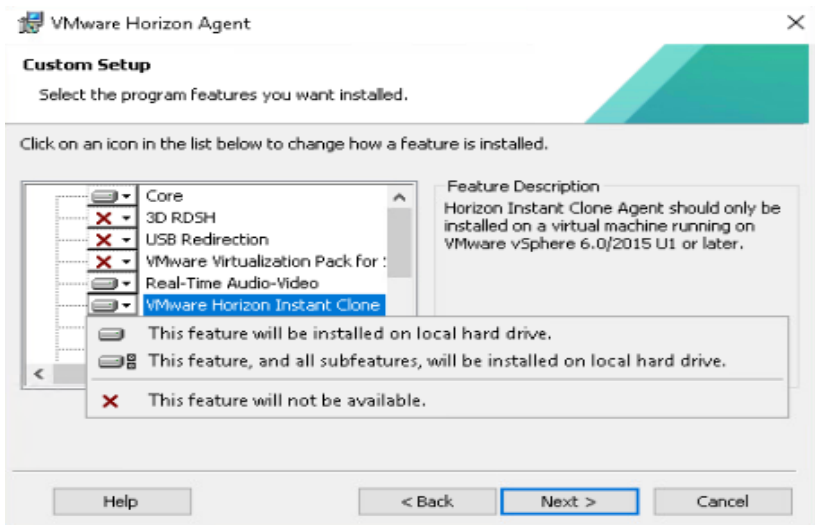
**Step 4.** Based on the Desktop pool you want to create, select either View Composer Agent or Instant Clone Agent installation. Do not install both features on the same master image.

**Step 5.** Enable installation of the VMware Horizon View Instant Agent for Instant -clone VDI virtual machines.



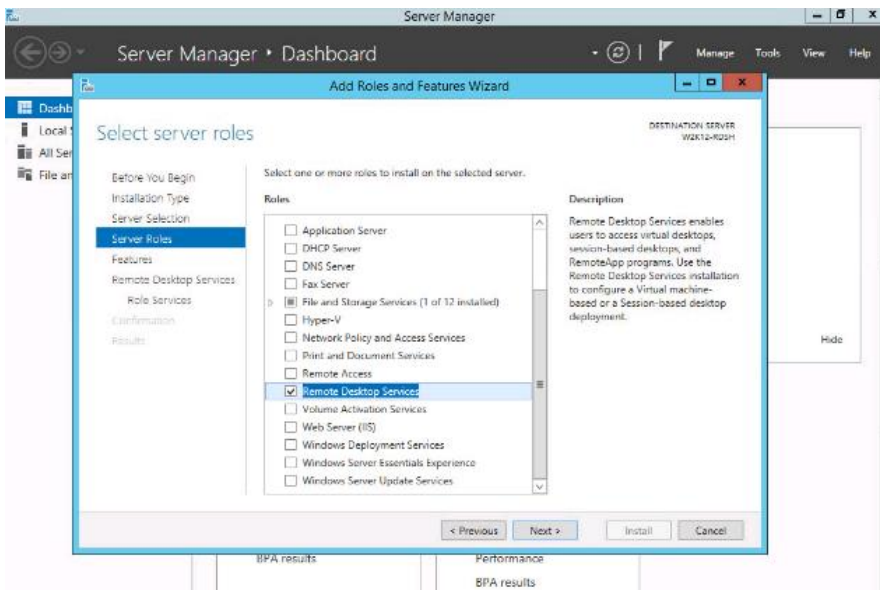
**Step 6.** Disable the Horizon View Composer Agent and enable the Horizon Instant Clone Agent for Instant Clone floating assigned desktop pool creation.

**Note:** The VMware Horizon Composer was not tested for this CVD.



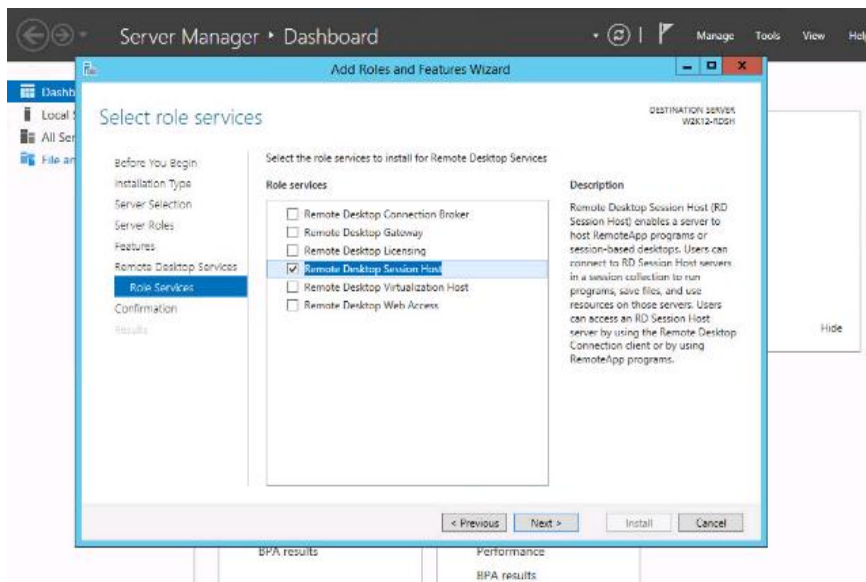
**Note:** Prior to installing the Horizon View Agent on a Microsoft Server 2019 virtual machine, you must add the Remote Desktop Services role and the Remote Desktop Session Host role service.

**Step 7.** To add Remote Desktop Services role on Windows Server OS from the Server Manager, use the Add Roles and Features wizard:

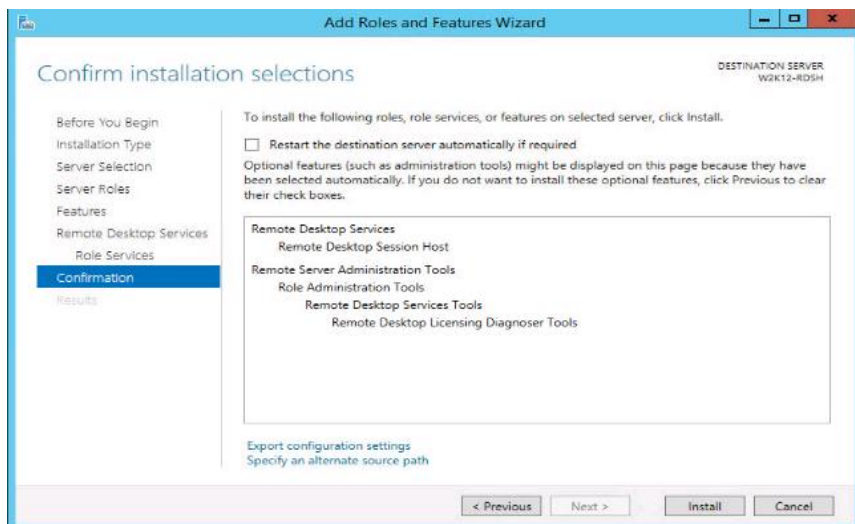


**Step 8.** Add Remote Desktop Session Host services.





**Step 9.** Click Install.

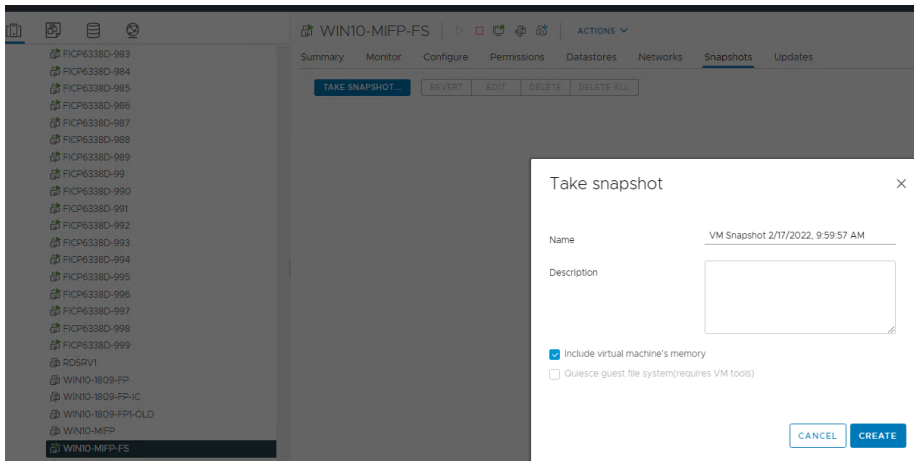


## Procedure 2. Install Additional Software

- Step 1.** For testing, we installed Microsoft Office 2016 64-bit version.
- Step 2.** Log into the VSI Target software package to facilitate workload testing.
- Step 3.** Install service packs and hot fixes required for the additional software components that are being added.
- Step 4.** Reboot or shut down the VM as required.

## Procedure 3. Create a Native Snapshot for Automated Desktop Pool Creation

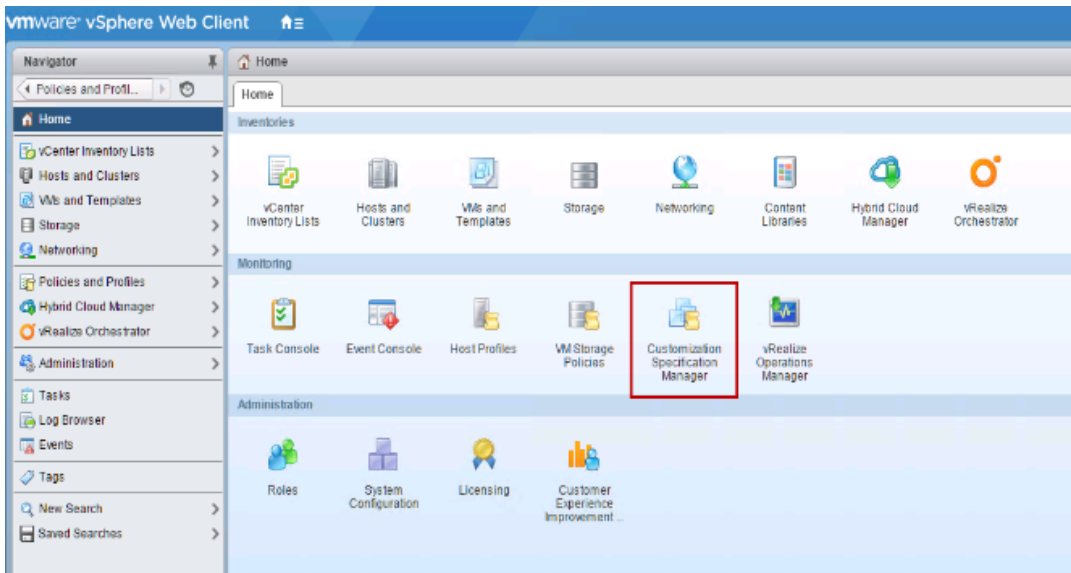
- Step 1.** Log into vCenter WebUI.
- Step 2.** Select the master image for the automated desktop pool creation.
- Step 3.** Right-click and select Master Image X Data Platform > Snapshot Now.



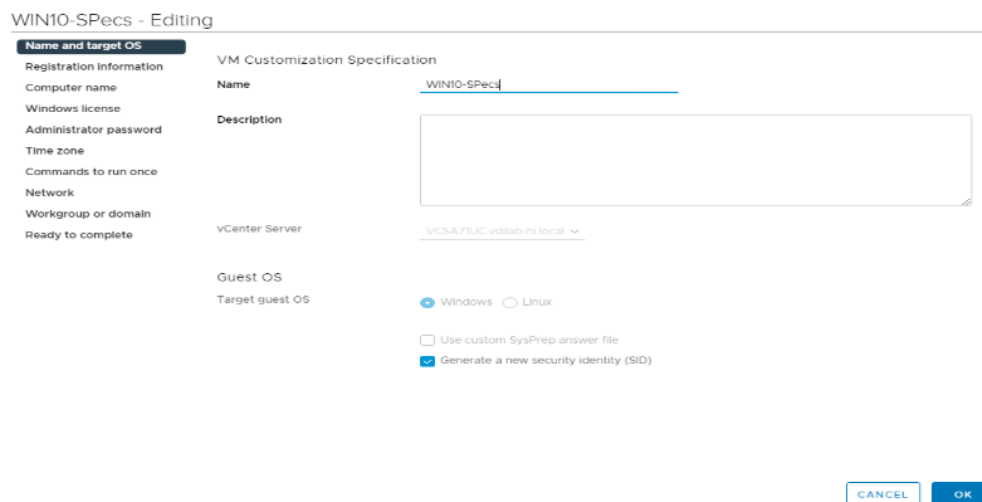
**Step 4.** Enter a name for the WIN 10 Master Image snapshot.

### Procedure 4. Create Customization Specification for Virtual Desktops

**Step 1.** On vCenter WebUI, select Customization Specification Manager.



**Step 2.** Select VM Operating System as Windows for Windows based guest OS optimization. Enter a name.



**Step 3.** Provide name and organization details.

### WIN10-SPEcs - Editing

Name and target OS

**Registration information**

Owner name	Administrator
Computer name	
Windows license	Owner organization
	vdilab-hi.local

Administrator password

Time zone

Commands to run once

Network

Workgroup or domain

Ready to complete

**Step 4.** Provide a computer name. For this solution, we selected Use the virtual machine name.

**Step 5.** Provide the product License key if required.

### WIN10-SPEcs - Editing

Name and target OS

Registration information

**Computer name**

- Use the virtual machine name ⓘ
- Enter a name in the Clone/Deploy wizard
- Enter a name

Windows license

Administrator password

Time zone

Commands to run once

- Append a unique numeric value. ⓘ

Network

- Generate a name using the custom application configured with the vCenter Server

Workgroup or domain

Ready to complete

Argument

**Step 6.** Provide Password credentials.

### WIN10-SPEcs - Editing

Name and target OS

Registration information

Computer name

Windows license

**Administrator password**

Password	.....
Confirm password	.....

Time zone

- Automatically logon as Administrator

Commands to run once

Number of times to logon automatically 1

Network

Workgroup or domain

Ready to complete

**Step 7.** Select the Timezone.

WIN10-SPecs - Editing

Name and target OS

Registration information

Computer name

Windows license

Administrator password

**Time zone**

Commands to run once

Network

Workgroup or domain

Ready to complete

Time zone

- (UTC-12:00) International Date Line West
- (UTC-11:00) Coordinated Universal Time-11
- (UTC-10:00) Aleutian Islands
- (UTC-10:00) Hawaii
- (UTC-09:30) Marquesas Islands
- (UTC-09:00) Alaska
- (UTC-09:00) Coordinated Universal Time-09
- (UTC-08:00) Baja California
- (UTC-08:00) Coordinated Universal Time-08
- (UTC-08:00) Pacific Time (US & Canada)**
- (UTC-07:00) Arizona
- (UTC-07:00) Chihuahua, La Paz, Mazatlan
- (UTC-07:00) Mountain Time (US & Canada)
- (UTC-06:00) Central America
- (UTC-06:00) Central Time (US & Canada)
- (UTC-06:00) Easter Island
- (UTC-06:00) Guadalajara, Mexico City, Monterrey
- (UTC-06:00) Saskatchewan
- (UTC-05:00) Bogota, Lima, Quito, Rio Branco
- (UTC-05:00) Chetumal
- (UTC-05:00) Eastern Time (US & Canada)
- (UTC-05:00) Haiti
- (UTC-05:00) Havana
- (UTC-05:00) Indiana (East)
- (UTC-04:00) Asuncion
- (UTC-04:00) Atlantic Time (Canada)

**Step 8.** Add the commands to run when the first-time user logs in if there are any.

**Step 9.** Provide the network information whether to use the DHCP server to assign IP address, or manual configuration.

WIN10-SPecs - Editing

Name and target OS

Registration information

Computer name

Windows license

Administrator password

Time zone

Commands to run once

**Network**

Workgroup or domain

Ready to complete

Use standard network settings for the guest operating system, including enabling DHCP on all network interfaces  
 Manually select custom settings

ADD    EDIT    DELETE

	Description	IPv4 Address	IPv6 Address
<input type="radio"/>	NIC1	Use DHCP	Not used

**Step 10.** Provide the domain name and user credentials.

WIN10-SPecs - Editing

Name and target OS

Registration information

Computer name

Windows license

Administrator password

Time zone

Commands to run once

Network

**Workgroup or domain**

Ready to complete

Workgroup    WORKGROUP

Windows Server domain    vdlilab-hi.local

Specify a user account that has permission to add a computer to the domain.

Username    Administrator

Password    .....

Confirm password    .....

**Step 11.** Review and click Next to complete creating the Customization Specs.

**Step 12.** Click Finish.

## WIN10-Specs - Editing

Name and target OS		
<b>Registration information</b>	Name	WIN10-Specs
<b>Computer name</b>	Target guest OS	Windows
<b>Windows license</b>	OS options	Generate new security ID
<b>Administrator password</b>	Registration info	Owner name: Administrator Organization: vdi-lab-hi.local
<b>Time zone</b>	Computer name	Use Virtual Machine name
<b>Commands to run once</b>	Product key	No product key specified
<b>Network</b>	Server license mode	Per server (Maximum Connections: 5)
<b>Workgroup or domain</b>	Administrator access	Do not log in automatically as Administrator
<b>Ready to complete</b>	Time zone	(UTC-08:00) Pacific Time (US & Canada)
	Network type	Standard
	Windows Server domain	vdi-lab-hi.local Username: Administrator

CANCEL OK

## Procedure 5. Create RDSH Farm

**Note:** Before you can create an RDSH desktop pool, you must first create a RDSH Farm.

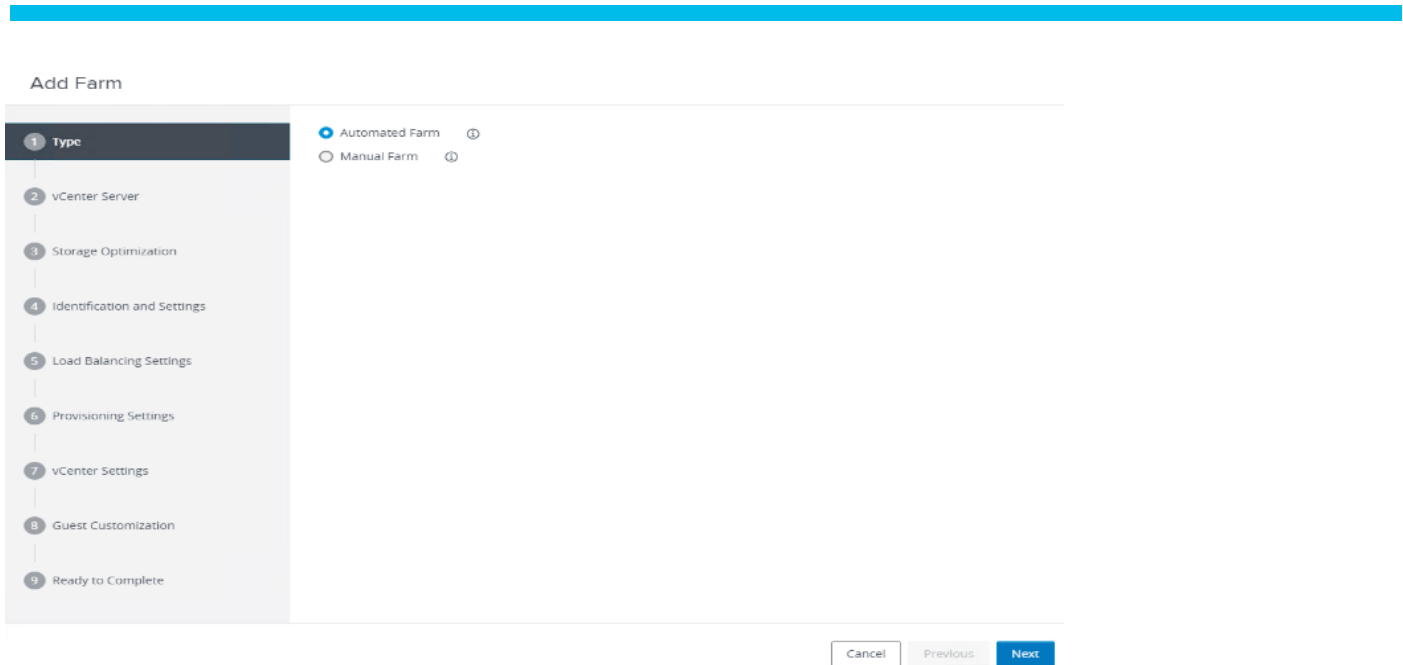
**Step 1.** In the VMware Horizon Administration console, select Farms under the Resource node of the Inventory pane.

**Step 2.** Click Add in the action pane to create a new RDSH Farm.

The screenshot shows the VMware Horizon Administration console interface. The top navigation bar includes the VMware logo, 'VMware Horizon', 'Pod Cluster-11201', a search bar, and user information 'administrator'. The left sidebar contains a navigation menu with 'Farms' selected. The main content area displays the 'Farms' page with a table of RDSH Farms. The table has columns for ID, Type, Source, RDS Hosts, Application Pools, Sessions, and Max Number. One farm is listed with ID 'B0', Type 'Automated', Source 'vCenter (instant clone)', RDS Hosts '40', Application Pools '0', Sessions '0', and Max Number 'Unlimited'. Above the table, there are buttons for 'Delete', 'More Commands', and 'Access Group', along with a filter input field.

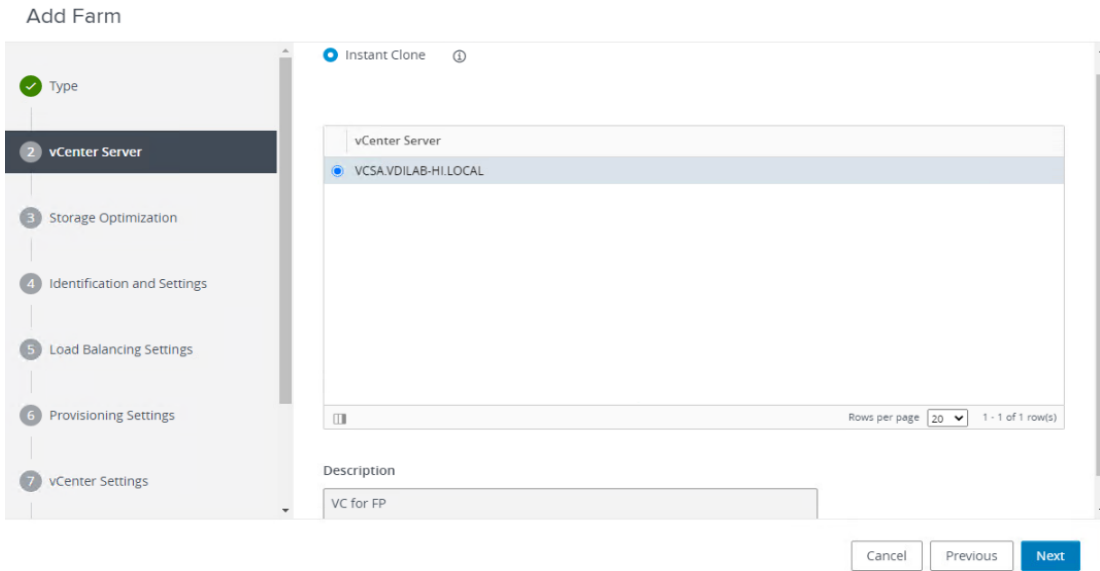
**Step 3.** Select either to create an Automated or Manual Farm. In this solution, we selected Automated Farm.

**Note:** A Manual Farm requires a manual registration of each RDSH server to Horizon Connection or Replica Server instance.



**Step 4.** Select the vCenter Server and Horizon Composer server that you will use to deploy the Horizon RDSH Farm.

**Step 5.** Click Next.



**Step 6.** Enter the RDSH Farm ID, Access group, Default Display Protocol (Blast/PCoIP/RDP).

**Step 7.** Select if users are allowed to change the default display protocol, Session timeout, Logoff Disconnected users, and select the checkbox to Enable HTML access.

**Step 8.** Click Next.

Add Farm - RDS-FARM

- Type
- vCenter Server
- Storage Optimization
- 4 Identification and Settings**
- 5 Load Balancing Settings
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Guest Customization
- 9 Ready to Complete

Asterisk (\*) denotes required field

\* ID  
RDS-FARM

Description  
VMware RDS-FARM for Server 2019 Sessions

Access Group  
/

Farm Settings

Default Display Protocol ⓘ  
Microsoft RDP

Allow Users to Choose Protocol  
Yes

3D Renderer ⓘ  
Manage using vSphere Client

vSphere doesn't support 3D option other than NVIDIA Grid vGPU for Windows Server OS

Pre-launch Session Timeout (Applications Only) ⓘ

Cancel Previous Next

**Step 9.** Select the provisioning settings, naming convention for RDSH server VM to deploy, and the number of VMs to deploy.

**Note:** In this study, we deployed 2600 RDSH virtual machines sessions using 168 RDS VMs spread evenly on 8 node Cisco UCS X-Series Cluster.

**Step 10.** Click Next.

Add Farm - RDS

- Type
- vCenter Server
- Storage Optimization
- Identification and Settings
- Load Balancing Settings
- 6 Provisioning Settings**
- 7 vCenter Settings

Asterisk (\*) denotes required field

Basic

Enable Provisioning ⓘ

Stop Provisioning on Error

Virtual Machine Naming ⓘ

\* Naming Pattern  
RDS

Farm Sizing

\* Maximum Machines  
80

\* Minimum Number of Ready (Provisioned) Machines during Instant Clone Maintenance Operations  
0

Cancel Previous Next

**Step 11.** Click Next.

**Step 12.** Select vCenter settings. For example, Master Image, snapshot, folder, Host or Cluster, resource pool, storage selection.

**Step 13.** Click Next.

## Add Farm - RDS-FARM

- ✔ Type
- ✔ vCenter Server
- ✔ Storage Optimization
- ✔ Identification and Settings
- ✔ Load Balancing Settings
- ✔ Provisioning Settings
- 7 vCenter Settings
- 8 Guest Customization
- 9 Ready to Complete

### Default Image

Asterisk (\*) denotes required field

\* Golden Image in vCenter

\* Snapshot

### Virtual Machine Location

\* VM Folder Location

### Resource Settings

\* Cluster

\* Resource Pool

\* Datastores  
 1 selected

Network

**Step 14.** For step 6 Datastores: Browse and click Data Stores.

**Step 15.** Click OK.

### Select Instant Clone Datastores

Select the instant clone datastores to use for this Automated Farm. Only datastores that can be used by the selected host or cluster can be selected.

Show all datastores (including local datastores) ↻

Datstore	Capacity (GB)	Free Space (...	FS Type	Drive Type	Storage Overcommit
<input checked="" type="checkbox"/> VDI-DS	92,160	91,983.25	NFS		Unbounded
<input type="checkbox"/> VDI-ESX	600	513.84	NFS		

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% Utilization (GB)	Max Recommended (GB)
Instant clones	91,983.25	2,080	2,720	4,320

**Step 16.** Click Next.

**Step 17.** Select the Active Directory Domain, the Active Directory OU into which the RDSH machines will be provisioned, and the Sysprep file created as part of the customization specific configuration performed earlier.

**Note:** If you choose the instant clone pool for the RDSH FARM creation, you may not see the Sys prep guest customization step shown in the screenshot shown below.



**Step 18. Click Next.**

Add Farm - RDS-FARM

Asterisk (\*) denotes required field

Domain: vdlab-hi.local(Administrator)

\* AD Container: OU=Target,OU=Computers,OU=LoginVSI [Browse]

Allow Reuse of Existing Computer Accounts ⓘ

Image Publish Computer Account: [ ] ⓘ

Use ClonePrep

Power-Off Script Name: [ ] ⓘ

Power-Off Script Parameters: [ ] ⓘ  
Example: p1 p2 p3

Post-Synchronization Script Name: [ ] ⓘ

Post-Synchronization Script Parameters: [ ] ⓘ

[Cancel] [Previous] [Next]

**Step 19. Review the pool creation information.**

**Step 20. Click Finish.**

Add Farm - RDS-FARM

ID	RDS-FARM
Description	VMware RDS-FARM for Server 2019 Sessions
Access Group	/
<b>Farm Settings</b>	
Default Display Protocol	Microsoft RDP
Allow Users to Choose Protocol	Yes
3D Renderer	Manage using vSphere Client
Pre-launch Session Timeout (Applications Only)	10 minutes
Empty Session Timeout (Applications Only)	1 minute
When Timeout Occurs	Disconnect
Logoff Disconnected Sessions	Never
Allow Session Collaboration	Disabled
Load Balancing Settings	

[Cancel] [Previous] [Submit]

The VMware Horizon Administration console displays the status of the provisioning task and pool settings:

VMware Horizon® Pod Cluster-HZ01

User Search About Settings Help administrator

Updated 04/19/2021, 12:22 PM

Sessions Problem vCenter VMs Problem RDS Hosts Events System Health

Monitor Dashboard Events Sessions Help Desk Users and Groups Inventory Desktops Applications Farms

### Farms

Add Edit Delete More Commands Access Group

Access Group All Filter

ID	Type	Source	RDS Hosts	Application Pools	Sessions	Max Number
RDFARM-	Automated	vCenter (instant clone)	40	0	0	Unlimited

VMware Horizon® Pod Cluster-HZ02

User Search About Settings Help administrator

Updated 03/25/2022, 4:07 PM

Sessions Problem vCenter VMs Problem RDS Hosts Events System Health

Monitor Dashboard Events Sessions Help Desk Users and Groups Inventory Desktops Applications Farms Machines Settings Servers Domains Product Licensing and Usage Global Settings Registered Machines

### RDFARM-

Summary RDS Hosts RDS Pools Sessions

Recover Remove From Farm More Commands

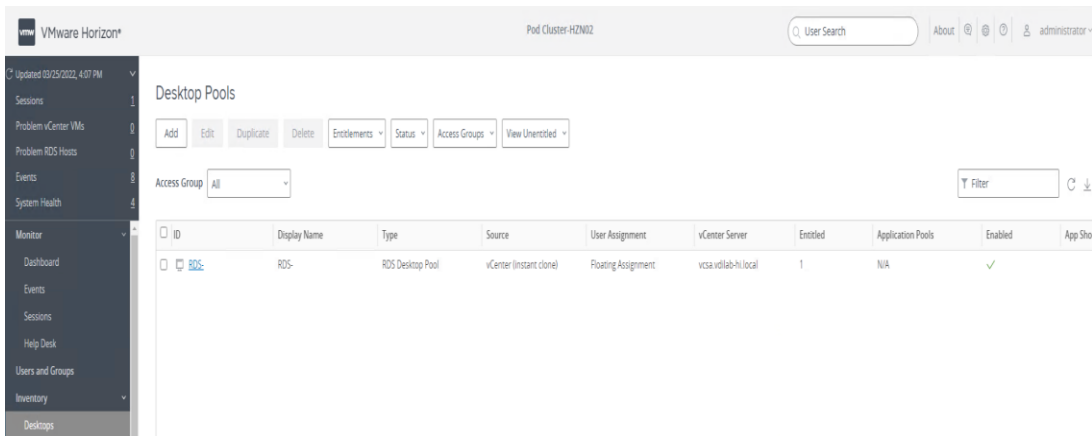
Filter

DNS Name	Type	Image	Pending Image	Task	Max Number of Connections	Agent Version	Enabled	Status
rdfarm-1.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-10.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-11.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-12.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-13.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-14.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-15.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-16.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-17.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-18.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-19.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-2.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-20.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available
rdfarm-21.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✓	Available

## Procedure 6. Create the Horizon 8 RDS Published Desktop Pool

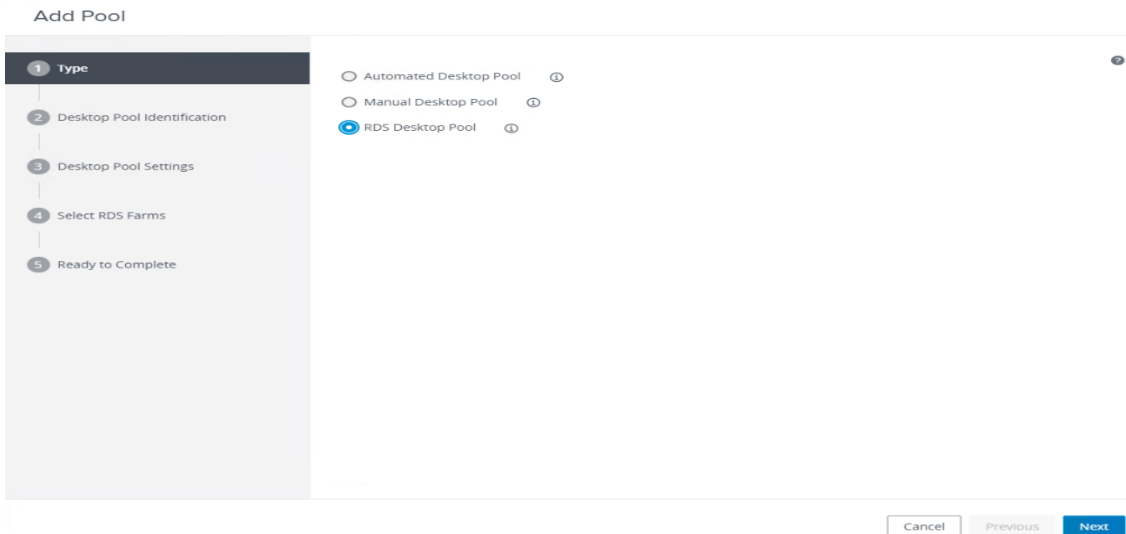
**Step 1.** In the Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.

**Step 2.** Click Add in the action pane.



**Step 3.** Select RDS Desktop pool.

**Step 4.** Click Next.



**Step 5.** Enter Pool ID and Display name.

**Step 6.** Click Next.



**Step 7.** Accept the default settings on Desktop Pool Settings page.

**Step 8.** Click Next.

Add Pool - RDS-POOL

State: Enabled

Connection Server Restrictions: None

Category Folder: None

Client Restrictions:  Enabled

Allow Separate Desktop Sessions from Different Client Devices: No

Buttons: Cancel, Previous, Next

**Step 9.** Click the “Select an RDS farm for this desktop pool” radio button.

**Step 10.** Click the farm created in the previous section or click create a new RDS Farm if not done so.

**Step 11.** Click Next.

Add Pool - RDS-POOL

Create a new RDS farm

Select an RDS farm for this desktop pool

Filter

Farm ID	Description	RDS Hosts	Max Number of Connections	Status
No records available.				

Buttons: Cancel, Previous, Next

**Step 12.** Review the pool settings.

**Step 13.** Select the checkbox “Entitle users after this wizard finishes” to authorize users for the newly create RDSH desktop pool.

**Step 14.** Click Finish.

**Step 15.** Select the Users or Groups checkbox, use the search tools to locate the user or group to be authorized, highlight the user or group in the results box.

**Step 16.** Click OK.

Find User or Group
✕

---

Type

Domain

Name/User Name

Description

**Find**

Users  Groups

Entire Directory

Starts with login

Starts with

**Find**

<input type="checkbox"/>	Name	User Name	Email	Description	In Folder
<input type="checkbox"/>	LoginVSI	LoginVSI/vdilab-hi.local			vdilab-hi.local/Login\

**Step 17.** You now have a functional RDSH Farm and Desktop Pool with users identified who are authorized to utilize Horizon RDSH sessions.

**Procedure 7. Create VMware Horizon Instant Clone and Full Clone Persistent Windows 10 Desktop Pool**

- Step 1.** In Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.
- Step 2.** Click Add in the action pane.
- Step 3.** Select assignment type for pool.
- Step 4.** Click Next.
- Step 5.** Select Floating or Dedicated user assignment.

Add Pool

- Type
- vCenter Server
- 3 User Assignment
- 4 Storage Optimization

Floating ⓘ

Dedicated ⓘ

Enable Automatic Assignment

Enable Multi-User Assignment ⓘ

Automatic assignment is not supported for multi-user assignment pools.

- Step 6.** Select the applicable options.
- Step 7.** Click Next.

Add Pool - WIN10-VDI-

- Type
- vCenter Server
- User Assignment
- Storage Optimization
- 5 Desktop Pool Identification**
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Desktop Pool Settings
- 9 Remote Display Settings

Asterisk (\*) denotes required field

\* ID ⓘ

Display Name ⓘ

Access Group ⓘ

Description

Cancel Previous Next

**Step 8.** Enter pool identification details.

**Step 9.** Click Next.

Add Pool - WIN10-VDI-

- Type
- vCenter Server
- User Assignment
- Storage Optimization
- Desktop Pool Identification
- 6 Provisioning Settings**
- 7 vCenter Settings
- 8 Desktop Pool Settings
- 9 Remote Display Settings

**Basic**

Enable Provisioning ⓘ

Stop Provisioning on Error

---

**Virtual Machine Naming** ⓘ

Specify Names Manually

0 names entered
Enter Names

Use a Naming Pattern ⓘ

\* Naming Pattern



---

**Provision Machines**

Machines on Demand

Min Number of Machines

All Machines Up-Front

---

**Desktop Pool Sizing**

\* Maximum Machines

\* Spare (Powered On) Machines

Cancel Previous Next

**Step 10.** Select Desktop Pool settings.

**Note:** Be sure to scroll down in this dialogue to configure all options.

**Step 11.** Click Next.

**Step 12.** Select Provisioning Settings.

**Step 13.** Click Next.

Add Pool - WIN10-VDI-

Default Image

Asterisk (\*) denotes required field

\* Golden Image in vCenter  
/VDI-DC/vm/WIN10-NEW-0222 Browse

\* Snapshot  
/WIN10-NEW-0325-SS/WIN10-NEW-0222-SS-2GBRV Browse

Virtual Machine Location

\* VM Folder Location  
/VDI-DC/vm/Discovered virtual machine Browse

Resource Settings

\* Cluster  
/VDI-DC/host/HX45 Browse

\* Resource Pool  
/VDI-DC/host/HX45/Resources Browse

\* Datastores  
1 selected Browse

Network

Cancel Previous Next

**Step 14.** Click Next.

**Step 15.** Click Next.

**Step 16.** Select each of the six required vCenter Settings by using the Browse button next to each field.

**Step 17.** For Datastore selection, select the correct datastore and set the Storage Overcommit as “Unbounded.”

**Step 18.** Click OK.

Add Pool - WIN10-VDI-

Asterisk (\*) denotes required field

Domain  
vdlab-hi.local(Administrator)

\* AD Container  
OU=Target,OU=Computers,OU=LoginVSI Browse

Allow Reuse of Existing Computer Accounts ⓘ

Image Publish Computer Account ⓘ

Use ClonePrep

Power-Off Script Name ⓘ

Power-Off Script Parameters  
Example: p1 p2 p3

Post-Synchronization Script Name ⓘ

Post-Synchronization Script Parameters

Cancel Previous Next

**Step 19.** Click Next.

**Step 20.** Set the Advanced Storage Options using the settings shown in the following screenshot.

**Step 21.** Click Next.

**Step 22.** Select Guest optimization settings.

**Step 23.** Select the Active Directory domain, browse to the Active Directory Container where the virtual machines will be provisioned and then select either the QuickPrep or Sysprep option you would like to use. Highlight the Customization Spec previously prepared.

**Step 24.** Click Next.

**Step 25.** Select the checkbox “Entitle users after pool creation wizard completion” if you would like to authorize users as part of this process. Follow instructions provided in the Create Horizon 8 RDS Desktop Pool to authorize users for the Instant Clone Pool.

**Step 26.** Click Finish to complete the Instant Clone Pool creation process.

Add Pool - WIN10-VDI-

<input type="checkbox"/> Entitle Users After Adding Pool	
Type	Automated Desktop Pool
User Assignment	Floating Assignment
vCenter Server	10.10.50.39
Unique ID	WIN10-VDI-
Description	-
Display Name	WIN10-VDI-
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Client Restrictions	Disabled
Log Off After Disconnect	Never
Connection Server Restrictions	None

Cancel Previous Submit

## Procedure 8. VMware Horizon Persistent Windows 10 Desktop Pool Creation

**Step 1.** In Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.

**Step 2.** Click Add in the action pane.

**Step 3.** Select assignment type for pool.

**Step 4.** Click Next.

Add Pool - WIN10-VDI-

Type

Automated Desktop Pool

Manual Desktop Pool

RDS Desktop Pool

**Step 5.** Select the Dedicated radio button.



**Step 6.** Select the Enable automatic assignment checkbox, if desired.

**Step 7.** Click Next.

**Step 8.** Select the Full Virtual Machines radio button and highlight your vCenter and Composer.

**Step 9.** Click Next.

Add Pool - WIN10-VDI-

Type

vCenter Server

3 User Assignment

4 Storage Optimization

5 Desktop Pool Identification

6 Provisioning Settings

7 vCenter Settings

8 Desktop Pool Settings

9 Remote Display Settings

Instant Clone ⓘ

Full Virtual Machines ⓘ

vCenter Server
10.10.50.39

Description

VC 7/UC

Cancel Previous Next

**Step 10.** Enter the pool identification details.

Add Pool - WIN10-VDI-

Type

vCenter Server

3 User Assignment

4 Storage Optimization

5 Desktop Pool Identification

6 Provisioning Settings

Floating ⓘ

Dedicated ⓘ

Enable Automatic Assignment

Enable Multi-User Assignment ⓘ

Automatic assignment is not supported for multi-user assignment pools.

Cancel Previous Next

**Step 11.** Click Next.

## Add Pool - WIN10-Persistent

Asterisk (\*) denotes required field

- Type
- vCenter Server
- User Assignment
- Storage Optimization
- 5 Desktop Pool Identification**
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Desktop Pool Settings
- 9 Remote Display Settings

\* ID

Display Name

Access Group

Description

Cancel Previous Next

**Step 12.** Select Desktop Pool settings.

**Step 13.** Click Next.

## Add Pool - WIN10-Persistent

Enable Provisioning

Stop Provisioning on Error

Virtual Machine Naming

Specify Names Manually

Enter Names

Start machines in maintenance mode

# Unassigned Machines Kept Powered On

Use a Naming Pattern

\* Naming Pattern

Provision Machines

Machines on Demand

Min Number of Machines

All Machines Up-Front

Desktop Pool Sizing

\* Maximum Machines

Cancel Previous Next

**Step 14.** Select the provisioning settings to meet your requirements.

**Step 15.** Click Next.

**Step 16.** Click Next.

**Step 17.** Select each of the five vCenter Settings.

**Step 18.** Click Next.

Add Pool - WIN10-Persistent

Virtual Machine Template

\* Template  
/VDI-DC/vm/WIN10-NEW-2022-FC

Virtual Machine Location

\* VM Folder Location  
/VDI-DC/vm

Resource Settings

\* Host or Cluster  
/VDI-DC/host/HX45

\* Resource Pool  
/VDI-DC/host/HX45/Resources

\* Datastores  
1 selected

Cancel Previous Next

**Step 19.** For Datastore selection, select the datastore with storage overcommit as “Unbounded.”

**Step 20.** Click OK.

**Step 21.** Select Advance Storage Options and enable the View Storage Accelerator.

**Step 22.** Click Next.

Add Pool - WIN10-Persistent

State  
Enabled

Connection Server Restrictions  
None

Category Folder  
None

Client Restrictions  Enabled

Session Types  
Desktop

Remote Machine Power Policy  
Take no power action

Log Off After Disconnect  
Never

Allow Users to Restart Machines  
No

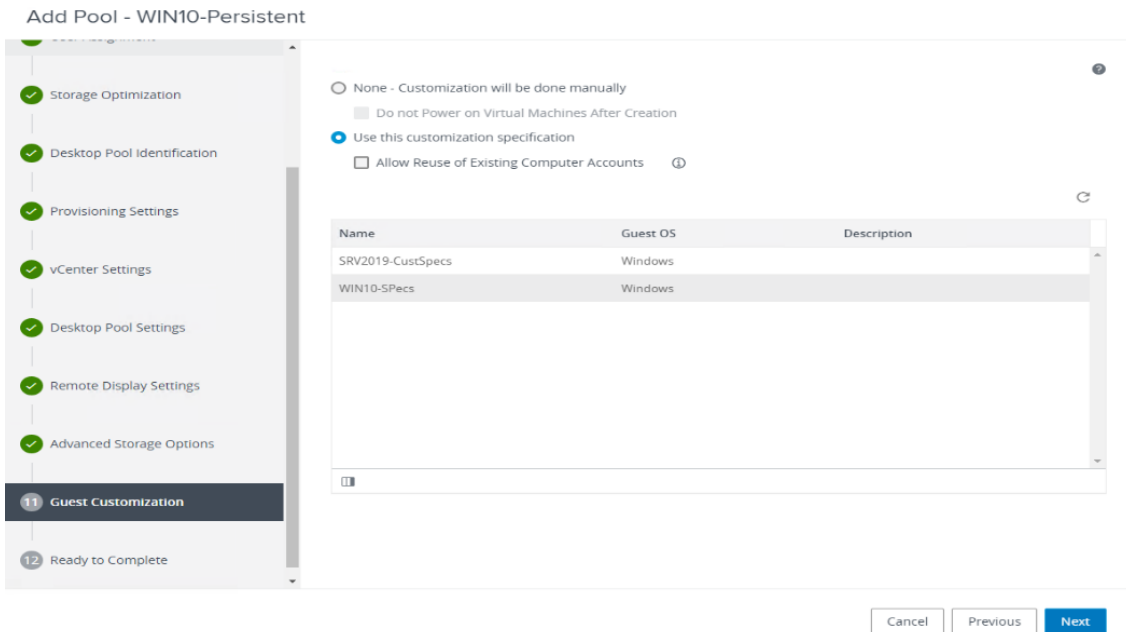
Show Assigned Machine Name

Show Machine Alias Name

Cancel Previous Next

**Step 23.** Select Guest optimization settings.

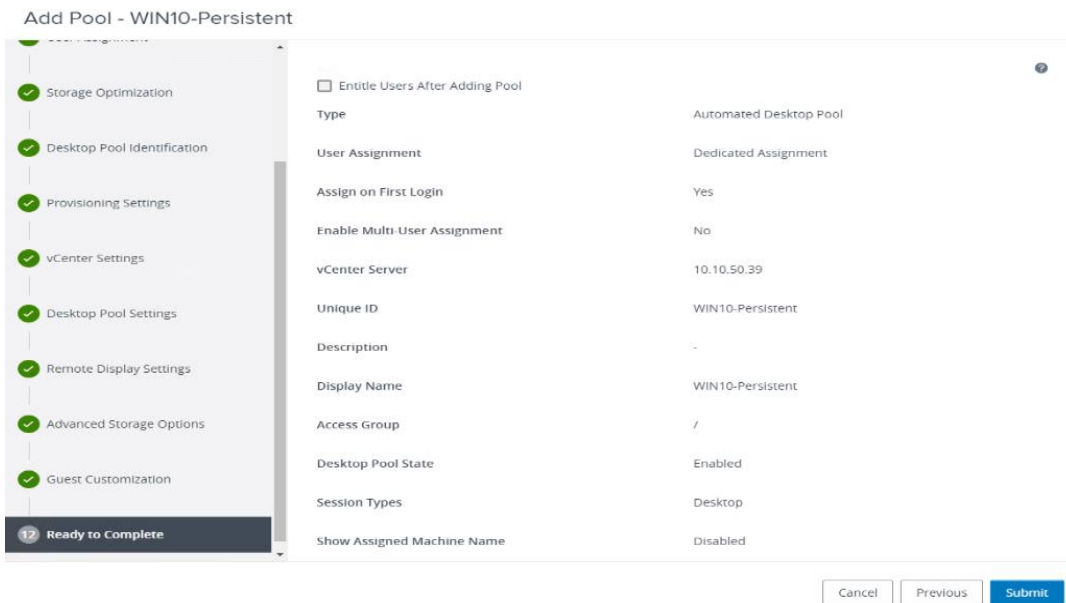
**Step 24.** Click Next.



**Step 25.** Review the summary of the pool you are creating.

**Step 26.** Select the checkbox “Entitle users after pool creation wizard completion” to authorize users for the pool.

**Step 27.** Click Finish.



**Step 28.** Follow the instructions provided in the Create Horizon 8 RDS Desktop Pool to authorize users for the Instant Clone Pool.

### Procedure 9. Configure FSLogix for VMware Remote Desktop Session Host (RDSH) Server Sessions and Windows 10 Virtual Desktops Profiles Profile Container

#### Tech tip

Profile Container is a full remote profile solution for non-persistent environments. Profile Container redirects the entire

user profile to a remote location. Profile Container configuration defines how and where the profile is redirected.

**Note:** Profile Container is inclusive of the benefits found in Office Container.

**Note:** When using Profile Container, both applications and users see the profile as if it's located on the local drive.

## Prerequisites

**Step 1.** Verify that you meet all [entitlement and configuration requirements](#).

**Step 2.** [Download and install FSLogix Software](#)

**Step 3.** Consider the storage and network requirements for your users' profiles (in this CVD, we used the NetApp A400 to store the FSLogix Profile disks).

**Step 4.** Verify that your users have [appropriate storage permissions](#) where profiles will be placed.

**Step 5.** Profile Container is installed and configured after stopping use of other solutions used to manage remote profiles.

**Step 6.** Exclude the VHD(X) files for Profile Containers from Anti-Virus (AV) scanning.

Configure FSLogix Profile Management


**Step 7.** When the FSLogix software is downloaded, copy the 'fslogix.admx and fslogix.adml' to the 'PolicyDefinitions' folder in your domain to manage the settings with Group Policy.

**Step 8.** On your VDI master image, install the FSLogix agent 'FSLogixAppsSetup' and accept all the defaults.

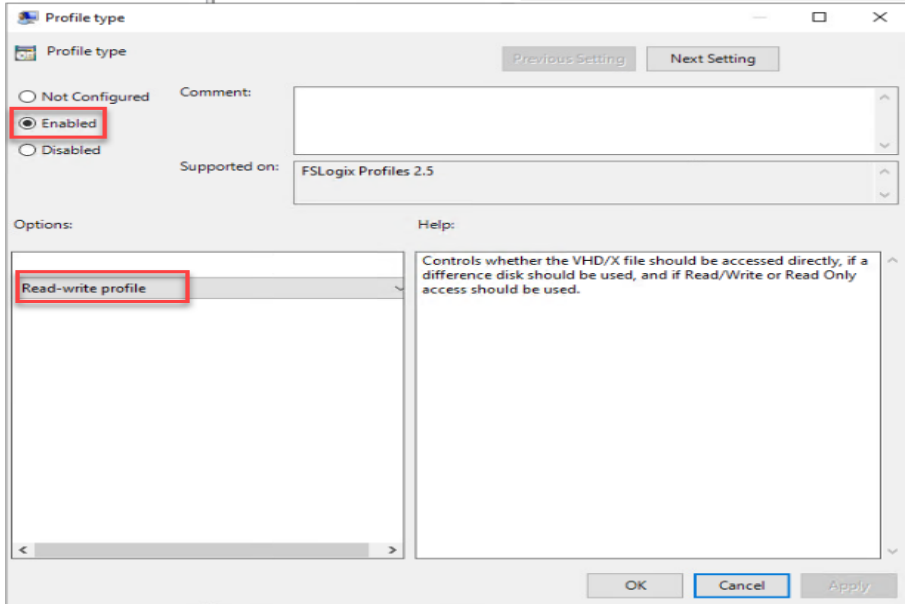
**Step 9.** Create a Group Policy object and link it to the Organizational Unit the VDI computer accounts.

**Step 10.** Right-click the FSLogix GPO policy.


**Step 11.** Enable FSLogix Profile Management.




**Step 12.** Select Profile Type (in this solution, we used Read-Write profiles).




**Step 13.** Enter the location of the Profile location (our solution used a CIFS share on the Netapp Array).



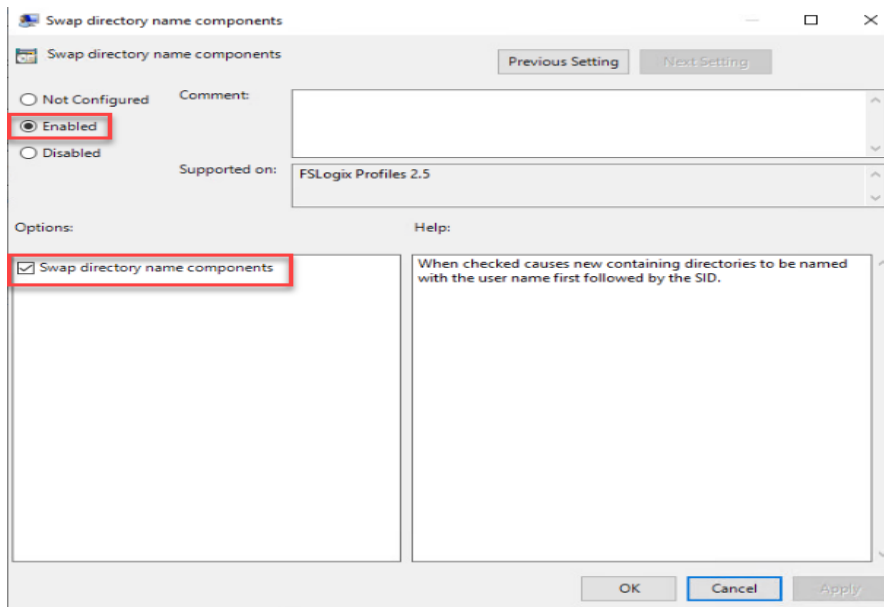
**Note:** We recommend using the Dynamic VHDX setting.



**Note:** VHDX is recommended over VHD.



**Note:** We enabled the 'Swap directory name components' setting for an easier administration but is not necessary for improved performance.



### Tech tip

FSLogix is an outstanding method of controlling the user experience and profile data in a VDI environment. There are many helpful settings and configurations for VDI with FSLogix that were not used in this solution.

A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Profile management in VDI environments is an integral part of the user experience. FSLogix, a Microsoft tool, was used to manage user profiles in this validated design.

FSLogix allows you to:

- Roam user data between remote computing session hosts
- Minimize sign in times for virtual desktop environments
- Optimize file IO between host/client and remote profile store
- Provide a local profile experience, eliminating the need for roaming profiles.
- Simplify the management of applications and 'Gold Images'

More information about the tool can be found [here](#).


### Procedure 10. Agent Installation

**Step 1.** FSLogix download file [here](#).

**Step 2.** Run FSLogixAppSetup.exe on VDI master image (32 bit or 64 bit depending on your environment).

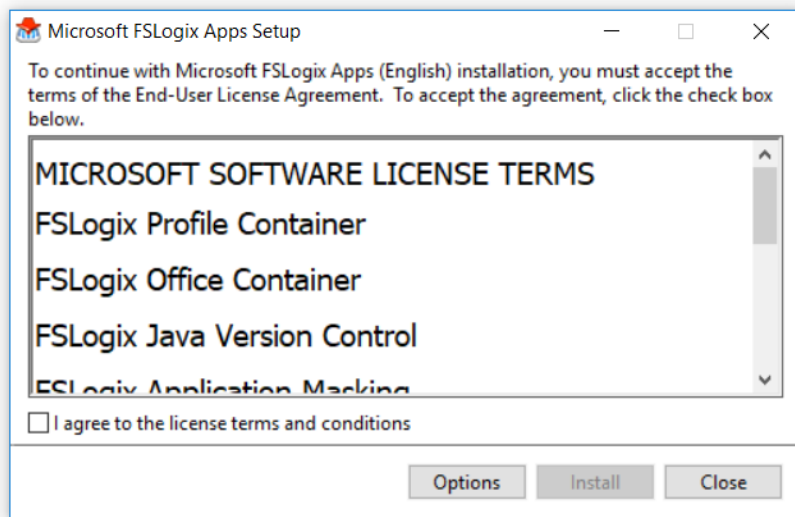
**Step 3.** Click OK to proceed with default installation folder.





**Step 4.** Review and accept the license agreement.

**Step 5.** Click Install.



**Step 6.** Reboot.

**Tech tip**

Consider enabling and configuring FSLogix logging as well as limiting the size of the profiles and excluding additional directories.

Figure 45. Example of FSLogix Policy

FXLogics-H8HZN

Scope Details Settings Delegation

VDI(LAB-HI) Enterprise Admins Edit settings, delete, modify security No

**Computer Configuration (Enabled)** hide

**Policies** hide

**Administrative Templates** hide

Policy definitions (ADMX files) retrieved from the local computer.

**FSLogix/ Profile Containers** hide

Policy	Setting	Comment
Dynamic VHD(X) allocation	Enabled	
Dynamic VHD(X) allocation	Enabled	Enabled
Enabled	Enabled	
Enabled	Enabled	Enabled
Profile type	Enabled	
Profile type	Enabled	Read-write profile
Size in MBs	Enabled	
Size in MBs	Enabled	3000
VHD location	Enabled	
VHD location	Enabled	\\10.10.61.121:fp01\VDIProfiles

**FSLogix/ Profile Containers/ Container and Directory Naming** hide

Policy	Setting	Comment
Swap directory name components	Enabled	
Swap directory name components	Enabled	Enabled
Virtual disk type	Enabled	
Virtual disk type	Enabled	VHDX

Activate Windows  
Go to Settings to activate Windows.

## Test Setup, Configuration, and Load Recommendation

This chapter contains the following:

- [Cisco UCS Test Configuration for Single Blade Scalability](#)
- [Cisco UCS Test Configuration for Full Scale Testing](#)
- [Test Methodology and Success Criteria](#)

We tested a single Cisco UCS X-Series 210C blade to validate against the performance of one and eight Cisco UCS X-Series 210C on a single chassis to illustrate linear scalability for each workload use case studied.

### Cisco UCS Test Configuration for Single Blade Scalability

This test case validates Recommended Maximum Workload per host server using VMware Horizon 8 Remote Desktop Server Hosted (RDSH) Sessions 375 Multi-session OS sessions and 280 Single-session Windows 10 OS sessions for Instant clones and Full clone virtual machines tests.

Figure 46. Test Configuration for Single Server Scalability VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi-session OS sessions

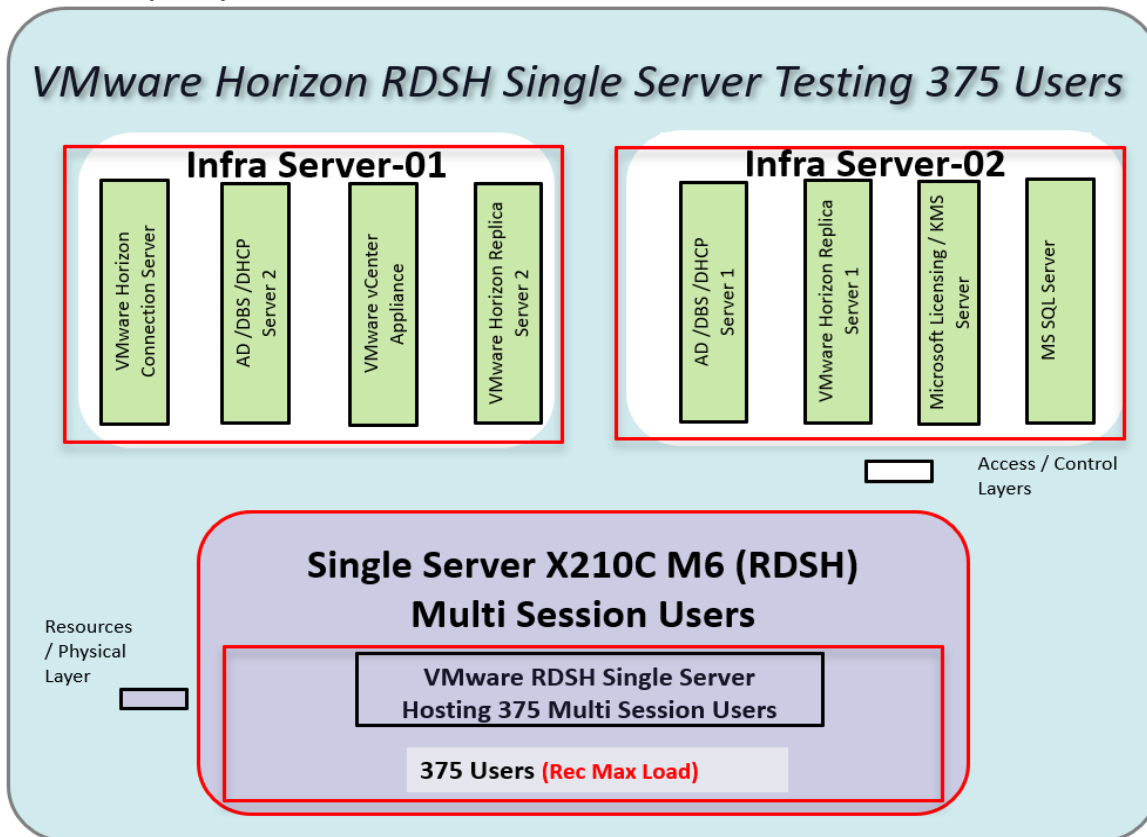


Figure 47. Test Configuration for Single Server Scalability VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi-session OS sessions on ESXi Host

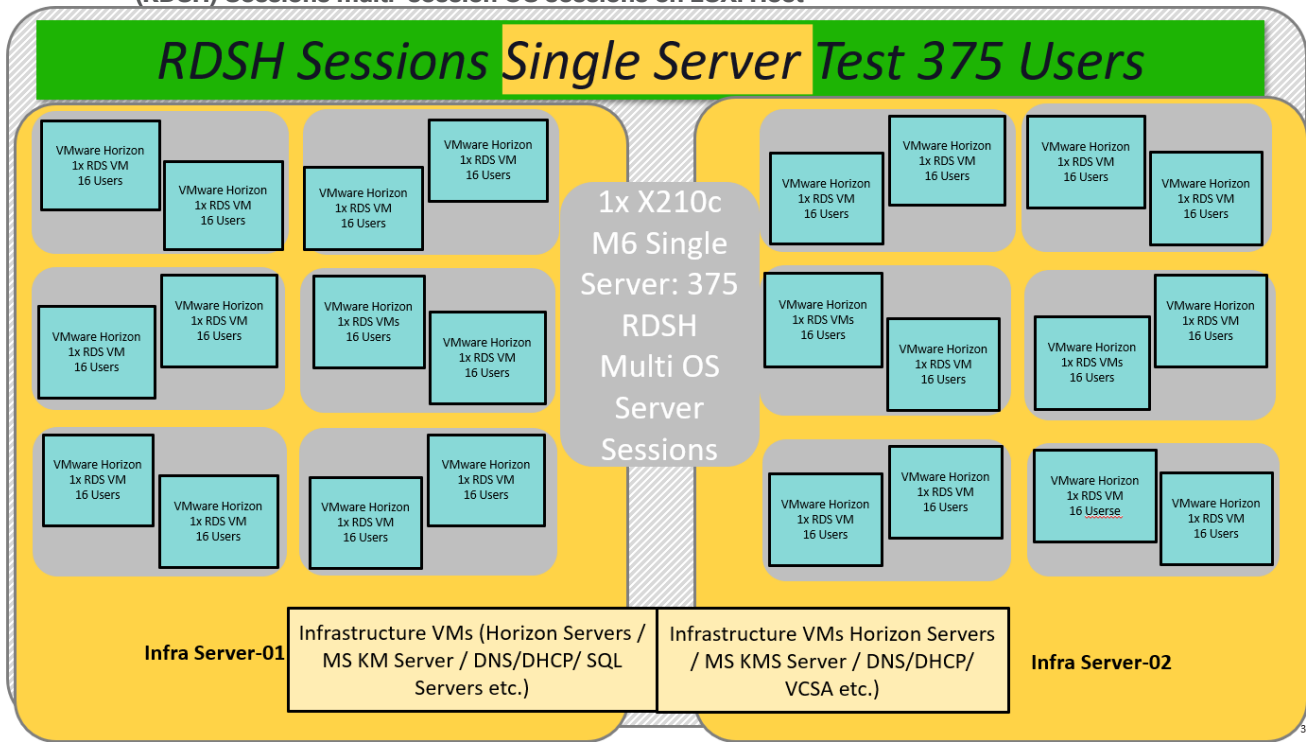
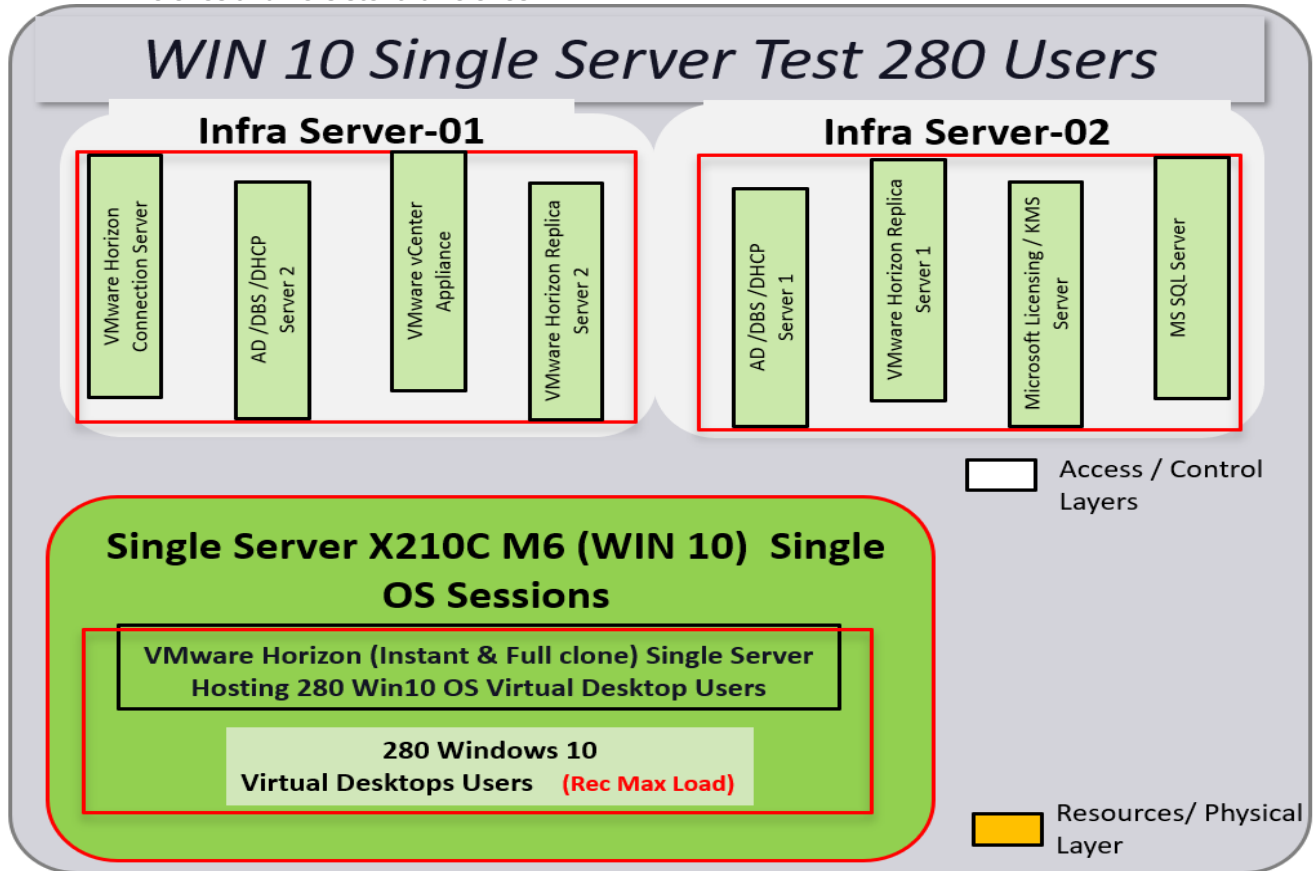


Figure 48. Test configuration for Single Server Scalability VMware Horizon WIN 10 Virtual Desktops for Instant clones and Persistent full clones



---

#### Hardware components:

- Cisco UCS 9508 X-Series Server Chassis
- 2 Cisco UCS 6454 4<sup>th</sup> Gen Fabric Interconnects
- 1 Cisco UCS X-Series 210C Servers with Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM for host blade
- Cisco VIC 14425 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32Gb 32-Port Fibre Channel Switches
- NetApp Storage AFF A400 with dual redundant controllers, with Twenty 1.92TB DirectFlash NVMe drives

#### Software components:

- Cisco UCS firmware 5.0(1b)
- NetApp ONTAP 9.10.1P1
- ESXi 7.0 Update 3 for host blades
- VMware Horizon 2209 RDSH Server Sessions and WIN 10 Virtual machines.
- Microsoft SQL Server 2019
- Microsoft Windows 10 64 bit (1909), 2vCPU, 3.5 GB RAM, 40 GB (Instant clones) / 80 GB (Full clones)
- Microsoft Windows Server 2019 (1809), 4vCPU, 24GB RAM, 80 GB vDisk (master)
- Microsoft Office 2016 32-bit
- FSLogix 2.9.7979.62170
- Login VSI 4.1.40 Knowledge Worker Workload (Benchmark Mode)

### Cisco UCS Test Configuration for Full Scale Testing

These test cases validate eight blades in a cluster hosting three distinct workloads using VMware Horizon RDSH Server Sessions and WIN 10 Virtual Desktops:

- 2600 VMware Horizon Remote Desktop Server Hosts (RDSH) sessions
- 1800 VMware Horizon Instant clone random pooled Windows 10 desktops
- 1800 VMware Horizon Full clone dedicated Windows 10 desktops

**Note:** Server N+1 fault tolerance is factored into this solution for each cluster/workload.

Figure 49. Test Configuration for Full Scale VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions

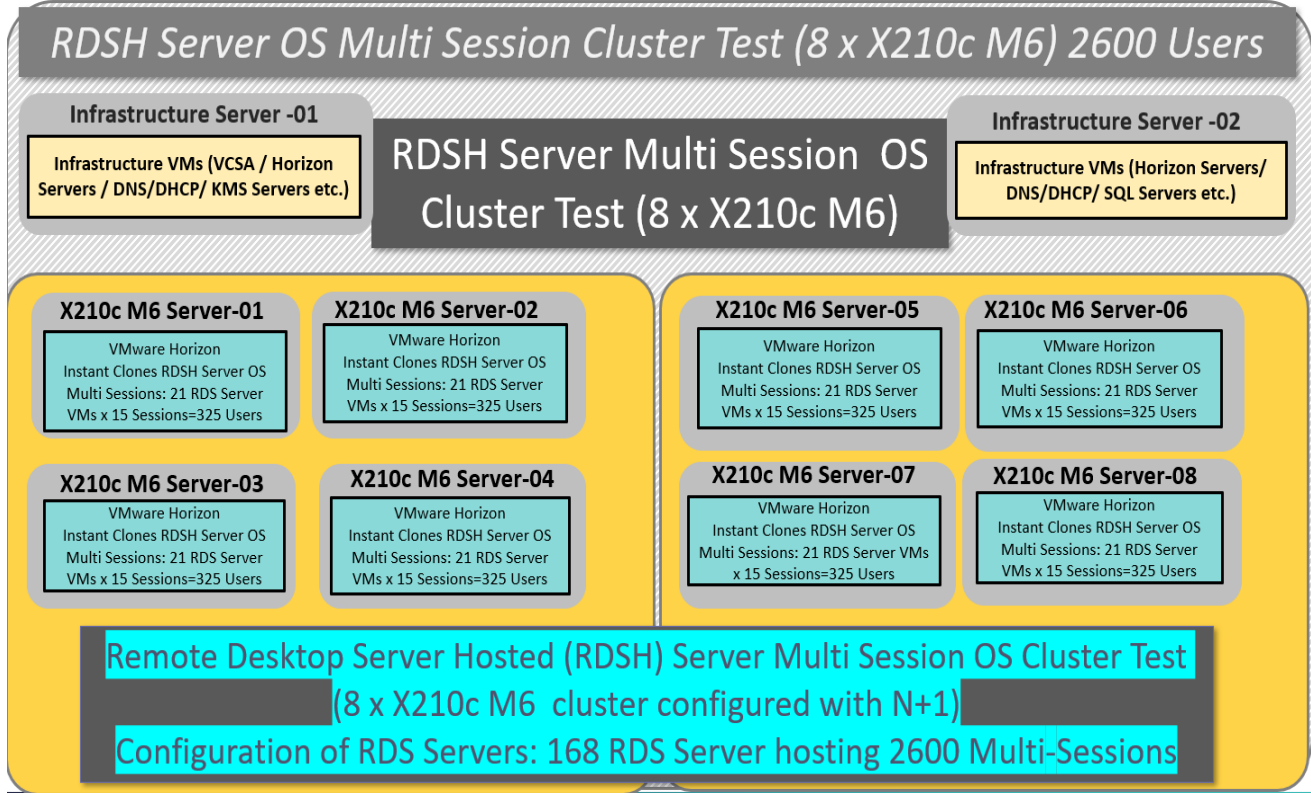


Figure 50. Test Configuration for VMware Instant Clones non-persistent Desktops Single-Session OS

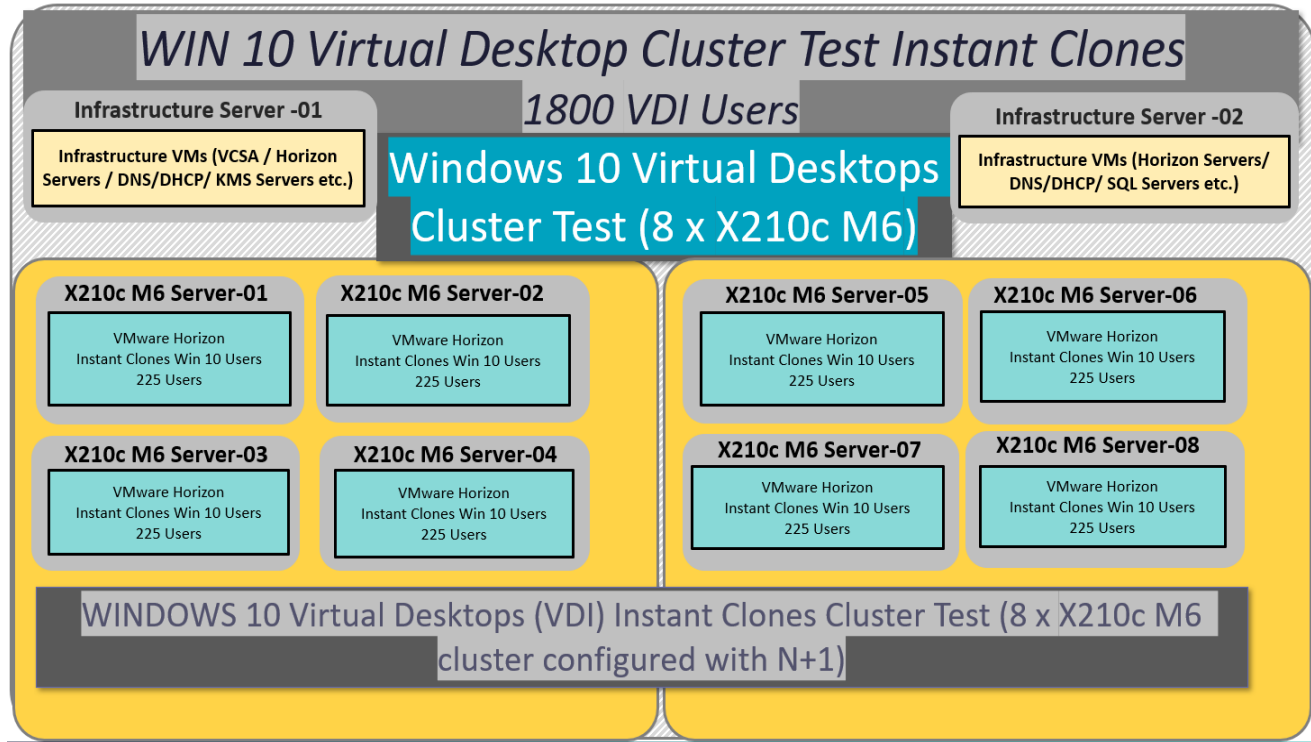
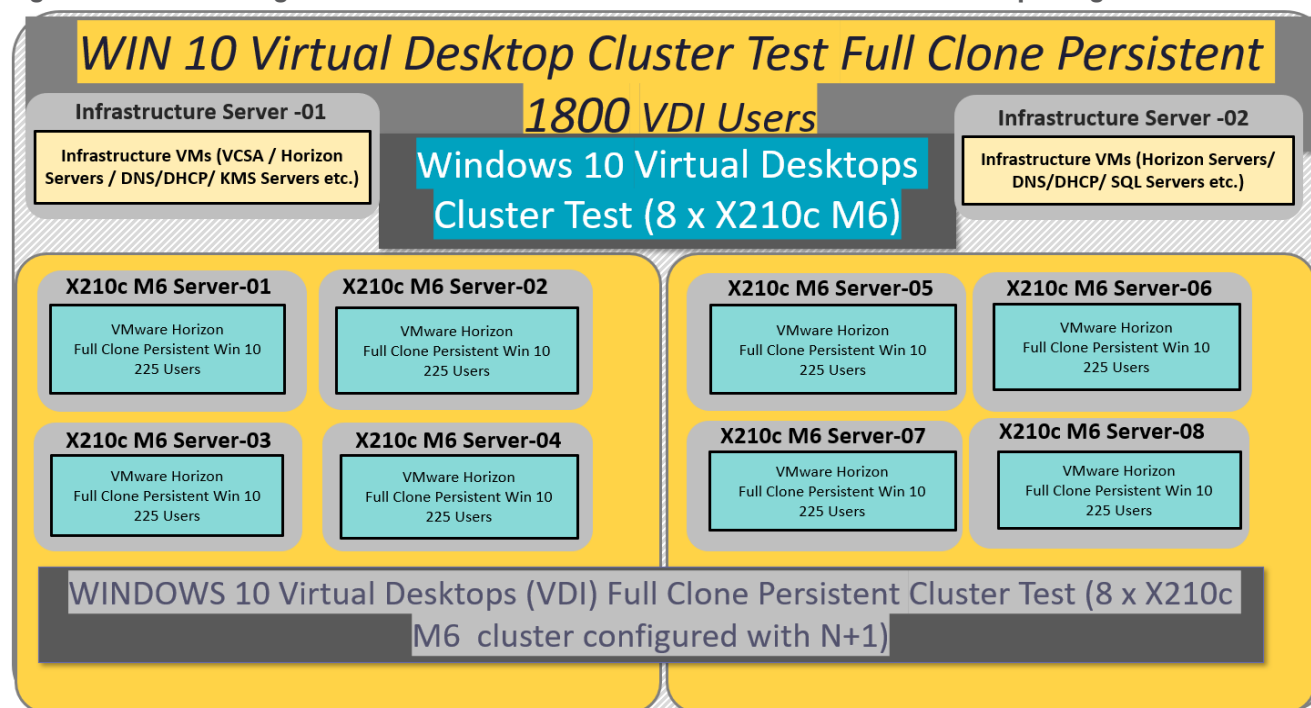


Figure 51. Test Configuration for Full Scale VMware Horizon Full Clone Virtual Desktops Single-Session OS



Hardware components:

- Cisco UCS 9508 Blade Server Chassis
- 2 Cisco UCS 6454 4<sup>th</sup> Gen Fabric Interconnects
- 8 Cisco UCSX-210c Servers with Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM for all host blades
- Cisco VIC 14425 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32Gb, 32-Port Fibre Channel Switches
- 1 NetApp Storage AFF A400 with dual redundant controllers, with Twenty 1.92TB DirectFlash NVMe drives

Software components:

- Cisco UCS firmware 5.0(1b)
- NetApp ONTAP 9.10.1P1
- ESXi 7.0 Update 3 for host blades
- VMware Horizon 2209 Remote Desktop Server Hosted (RDSH) Sessions and Windows 10 virtual desktops
- Microsoft SQL Server 2019
- Microsoft Windows 10 64-bit (1909), 2vCPU, 3.5 GB RAM, 40 GB (Instant clones) / 80 GB (Full clones) HDD (master) for virtual desktop configuration
- Microsoft Windows Server 2019 (1809), 4vCPU, 24GB RAM, 80 GB vDisk (master) for RDS Server VM configuration
- Microsoft Office 2016 32-bit

- 
- FSLogix 2.9.7979.62170
  - Login VSI 4.1.40.1 Knowledge Worker Workload (Benchmark Mode)

## Test Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the RDSH/VDI Session under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from <http://www.loginvsi.com>



---

## Test Procedure

This chapter contains the following:

- [Pre-Test Setup for Single and Multi-Blade Testing](#)
- [Test Run Protocol](#)
- [Success Criteria](#)
- [VSImax 4.1.x Description](#)
- [Server-Side Response Time Measurements](#)
- [Calculating VSImax v4.1.x](#)
- [Single-Server Recommended Maximum Workload](#)

The following protocol was used for each test cycle in this study to ensure consistent results.

### Pre-Test Setup for Single and Multi-Blade Testing

All virtual machines were shut down utilizing the VMware Horizon Administrator Console and vCenter.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

### Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. For testing where the user session count exceeds 1000 users, we will now deem the test run successful with up to 1% session failure rate.

Additionally, Cisco requires that the Login VSI Benchmark method be used for all single server and scale testing. This assures that our tests represent real-world scenarios. For each of the three consecutive runs on single server tests, the same process was followed. To do so, follow these steps:

1. Time 0:00:00 Start PerfMon/Esxstop Logging on the following systems:
  - a. Infrastructure and VDI Host Blades used in the test run
2. vCenter used in the test run
3. All Infrastructure virtual machines used in test run (AD, SQL, brokers, image mgmt., and so on)
4. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
5. Time 0:05: Boot Virtual Desktops/RDS Virtual Machines using View Connection server.
6. The boot rate should be around 10-12 virtual machines per minute per server.
7. Time 0:06 First machines boot.
8. Time 0:30 Single Server or Scale target number of desktop virtual machines booted on 1 or more blades.
9. No more than 30 minutes for boot up of all virtual desktops is allowed.
10. Time 0:35 Single Server or Scale target number of desktop virtual machines desktops available on View Connection Server.

- 
11. Virtual machine settling time.
  12. No more than 60 Minutes of rest time is allowed after the last desktop is registered on the VMware Horizon Console or available in Horizon Connection Server dashboard. Typically, a 30-45-minute rest period is sufficient.
  13. Time 1:35 Start Login VSI 4.1.x Office Worker Benchmark Mode Test, setting auto-logoff time at 15 minutes, with Single Server or Scale target number of desktop virtual machines utilizing sufficient number of Launchers (at 20-25 sessions/ per launcher).
  14. Time 2:23 Single Server or Scale target number of desktop virtual machines desktops launched (48-minute benchmark launch rate).
  15. Time 2:25 All launched sessions must become active. id test run within this window.
  16. Time 2:40 Login VSI Test Ends (based on Auto Logoff 15 minutes period designated above).
    - a. Time 2:55 All active sessions logged off.
  17. Time 2:57 All logging terminated, Test complete.
  18. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows machines.
  19. Time 3:30 Reboot all hypervisor hosts.
  20. Time 3:45 Ready for the new test sequence.

## Success Criteria

Our pass criteria for this testing is as follows:

- Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 4.1.x Office Worker workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The VMware Horizon Console must be monitored throughout the steady state and will make sure of the following:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered, unavailable, or available state at any time during steady state
- Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Stuck sessions define a test failure condition.
- Cisco requires three consecutive runs with results within +/-1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with the proposed white paper.)

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing. FlexPod Datacenter with Cisco UCS and VMware Horizon Remote Desktops and Windows 10 virtual desktops on VMware ESXi 7.0 Update 3 Test Results.

---

The purpose of this testing is to provide the data needed to validate VMware Horizon Remote Desktop Sessions (RDS) and VMware Horizon Virtual Desktop (VDI) instant-clones and VMware Horizon Virtual Desktop (VDI) full-clones models using ESXi and vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2019 sessions on Cisco UCS X-Series Blade Servers using the NetApp Storage AFF A400 storage system.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of VMware products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single blade performance and multi-blade, linear scalability.

### **VSI<sub>max</sub> 4.1.x Description**

The philosophy behind Login VSI is different from conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put: the faster the execution time or the bigger the throughput, the faster the system is according to the benchmark.

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for HSD or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer, and Adobe PDF reader. By gradually increasing the amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what is its true maximum user capacity.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSI<sub>max</sub>. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time, it will be clear the response times escalate at saturation point.

This VSI<sub>max</sub> is the “Virtual Session Index (VSI).” With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

### **Server-Side Response Time Measurements**

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system and are initiated at logon within the simulated user’s desktop session context.

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI, the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solution, for a benchmark like Login VSI.

## Calculating VSImax v4.1.x

The simulated desktop workload is scripted in a 48-minute loop when a simulated Login VSI user is logged on, performing generic Office worker activities. After the loop is finished it will restart automatically. Within each loop, the response times of sixteen specific operations are measured in a regular interval: sixteen times in within each loop. The response times of these five operations are used to determine VSImax.

The operations from which the response times are measured are:

- Notepad File Open (NFO)
- Loading and initiating VSINotepad.exe and opening the openfile dialog. This operation is handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.
- Notepad Start Load (NSLD)
- Loading and initiating VSINotepad.exe and opening a file. This operation is also handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Zip High Compression (ZHC)

This action copy's a random file and compresses it (with 7zip) with high compression enabled. The compression will very briefly spike CPU and disk IO.

- Zip Low Compression (ZLC)


This action copy's a random file and compresses it (with 7zip) with low compression enabled. The compression will very briefly disk IO and creates some load on the CPU.

- CPU


Calculates a large array of random data and spikes the CPU for a short period of time.

These measured operations within Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, and so on. These operations are specifically short by nature. When such operations become consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. If such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

**Figure 52. Sample of a VSI Max Response Time Graph, Representing a Normal Test**



**Figure 53. Sample of a VSI Test Response Time Graph with a Performance Issue**



When the test is finished, VSI<sub>max</sub> can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSI<sub>max</sub> is not reached, and the amount of sessions ran successfully.

The response times are quite different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. These response times of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSI<sub>max</sub> models, this weighting much better represents system performance. All actions have similar weight in the VSI<sub>max</sub> total. The following weighting of the response times is applied.

The following actions are part of the VSI<sub>max</sub> v4.1.x calculation and are weighted as follows (US notation):

- Notepad File Open (NFO): 0.75
- Notepad Start Load (NSLD): 0.2
- Zip High Compression (ZHC): 0.125
- Zip Low Compression (ZLC): 0.2
- CPU: 0.75

This weighting is applied on the baseline and normal Login VSI response times.

With the introduction of Login VSI 4.1.x, we also created a new method to calculate the basephase of an environment. With the new workloads (Taskworker, Powerworker, and so on) enabling 'basephase' for a more reliable baseline has become obsolete. The calculation is explained below. In total the 15 lowest VSI response time samples are taken from the entire test; the lowest 2 samples are removed. and the 13 remaining samples are averaged. The result is the Baseline. To summarize:

- Calculate the Basephase
  - Take the lowest 15 samples of the complete test
  - From those 15 samples remove the lowest 2
  - Average the 13 results that are left is the baseline

The VSI<sub>max</sub> average response time in Login VSI 4.1.x is calculated on the number of active users that are logged on the system.

---

Always a 5 Login VSI response time samples are averaged + 40 percent of the number of “active” sessions. For example, if the active sessions are 60, then latest 5 + 24 (=40 percent of 60) = 31 response time measurement is used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5 percent and bottom 5 percent of the VSI response time samples are removed from the average calculation, with a minimum of 1 top and 1 bottom sample. As a result, with 60 active users, the last 31 VSI response time sample are taken. From those 31 samples, the top 2 samples are removed, and the lowest 2 results are removed (5 percent of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

VSI<sub>max</sub> v4.1.x is reached when the VSI<sub>base</sub> + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSI<sub>max</sub> response time can grow 2 - 3x the baseline average. In end-user computing, a 3x increase in response time in comparison to the baseline is typically regarded as the maximum performance degradation to be considered acceptable.

In VSI<sub>max</sub> v4.1.x this latency threshold is fixed to 1000ms, this allows better and fairer comparisons between two different systems, especially when they have different baseline results. Ultimately, in VSI<sub>max</sub> v4.1.x, the performance of the system is not decided by the total average response time, but by the latency it has under load. For all systems, this is now 1000ms (weighted).

The threshold for the total response time is: average weighted baseline response time + 1000ms.

When the system has a weighted baseline response time average of 1500ms, the maximum average response time may not be greater than 2500ms (1500+1000). If the average baseline is 3000 the maximum average response time may not be greater than 4000ms (3000+1000).

When the threshold is not exceeded by the average VSI response time during the test, VSI<sub>max</sub> is not hit, and the number of sessions ran successfully. This approach is fundamentally different in comparison to previous VSI<sub>max</sub> methods since it is required to saturate the system beyond VSI<sub>max</sub> threshold.

Lastly, VSI<sub>max</sub> v4.1.x is now always reported with the average baseline VSI response time result. For example: “The VSI<sub>max</sub> v4.1.x was 125 with a baseline of 1526ms”. This helps considerably in the comparison of systems and gives a more complete understanding of the system. The baseline performance helps to understand the best performance the system can give to an individual user. VSI<sub>max</sub> indicates what the total user capacity is for the system. These two are not automatically connected and related.

When a server with an extremely fast dual core CPU, running at 3.6 GHz, is compared to a 10 core CPU, running at 2,26 GHz, the dual core machine will give an individual user better performance than the 10-core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

However, the server with the slower 10 core CPU will easily have a larger capacity than the faster dual core system. This is indicated by VSI<sub>max</sub> v4.1.x, and the higher VSI<sub>max</sub> is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSI<sub>max</sub> method is introduced: VSI<sub>max</sub> v4.1.x. This methodology gives much better insight into system performance and scales to extremely large systems.

## Single-Server Recommended Maximum Workload

For both the VMware Remote Desktops Server Hosted (RDSH) Sessions and Windows 10 Virtual Desktops, use cases, a recommended maximum workload was determined by the Login VSI Knowledge Worker Workload in VSI Benchmark Mode end user experience measurements and blade server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to ensure that end user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90–95 percent.

Memory should never be oversubscribed for Desktop Virtualization workloads.

**Table 26. Phases of Test Runs**

Test Phase	Description
Boot	Start all RDS and VDI virtual machines at the same time
Idle	The rest time after the last desktop is registered on the VMware Horizon Console. (Typically, a 30–45 minute, <60 min)
Logon	The Login VSI phase of the test is where sessions are launched and start executing the workload over a 48 minutes duration
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files (typically for the 15-minute duration)
Logoff	Sessions finish executing the Login VSI workload and logoff

---

## Test Results

This chapter contains the following:

- [Single-Server Recommended Maximum Workload Testing for Remote Desktop Server Hosted \(RDSH\) Server Sessions and Win 10 Virtual Desktops](#)
- [Single-Server Recommended Maximum Workload for VMware Horizon Instant Clone Remote Desktop Server Hosted \(RDSH\) Multi Session OS Random Sessions with 375 Users](#)
- [Single-Server Recommended Maximum Workload for VMware Horizon Instant Clone Windows 10 with 280 Users](#)
- [Full Scale Workload Testing](#)
- [Scalability Considerations and Guidelines](#)
- [Scalability of VMware Horizon Remote Desktop Server Hosted \(RDSH\) Sessions and Win 10 Virtual Desktops Configuration](#)

### Single-Server Recommended Maximum Workload Testing for Remote Desktop Server Hosted (RDSH) Server Sessions and Win 10 Virtual Desktops

This section shows the key performance metrics that were captured on the Cisco UCS host blades during the single server testing to determine the Recommended Maximum Workload per host server. The single server testing comprised of following three tests:

- 375 VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions (Random)
- 280 VMware Horizon Instant Clone desktops (Random)
- 280 VMware Horizon Full clone desktops (Dedicated)

### Single-Server Recommended Maximum Workload for Instant Clone Remote Desktop Server Hosted (RDSH) Multi Session OS Random Sessions with 375 Users

The recommended maximum workload for a Cisco UCS X-Series blade server with dual Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM is 375 Instant Clone Remote Desktop Server Hosted (RDSH) Multi Session OS with 4 vCPU and 24 GB RAM. The X-Series server ran Windows Server 2019 RDS Virtual Machines.

Login VSI performance data is as follows:



Figure 54. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | VSI Score

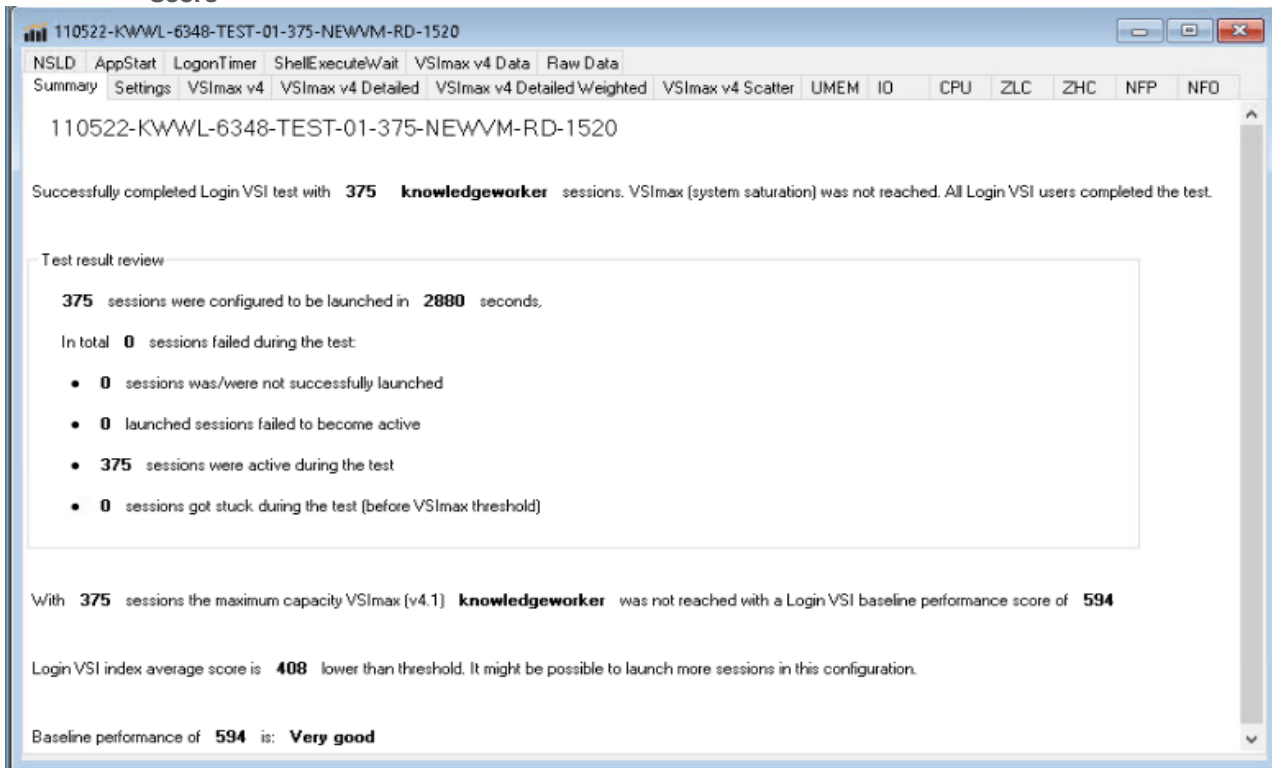




Figure 55. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | VSI Score




Performance data for the server running the workload is as follows:


**Figure 56. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | Host CPU Utilization**




**Figure 57. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | Host Memory Utilization**




**Figure 58. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | Host Network Utilization**



**Figure 59. Processor Time Utilization**



**Figure 60. Processor Time Utilization -Repeat Test**



## AFF A400 Storage Charts for 375 Remote Desktop Server Hosted (RDSH) Session User NFS Data Stores on Two Storage Controllers

Figure 61. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 1 for 375 Remote Desktop Server Hosted (RDSH) Sessions Test

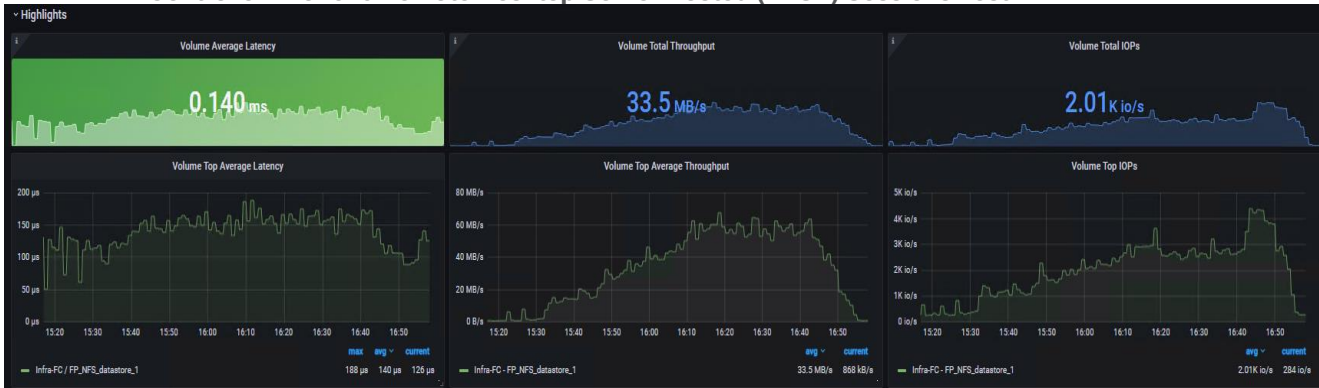
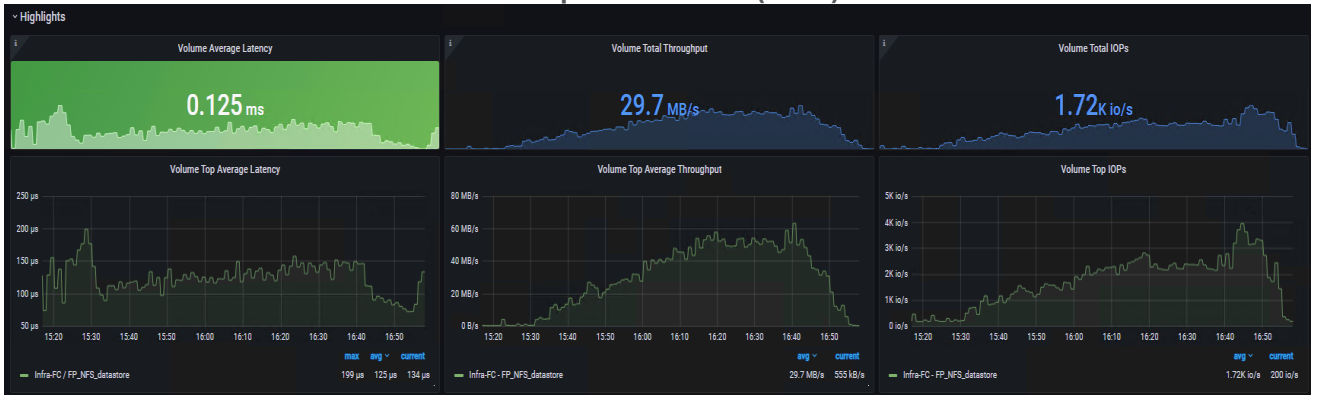


Figure 62. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 2 for 375 Remote Desktop Server Hosted (RDSH) Sessions Test



The recommended maximum workload for a Cisco UCS X-Series blade server with dual Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM is 280 Windows 10 64-bit VDI non-persistent VMware Horizon virtual machines with 2 vCPU and 3.5GB RAM.

Login VSI performance data is as follows:

# Instant Clones Testing

Figure 63. Single Server | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | VSI Score

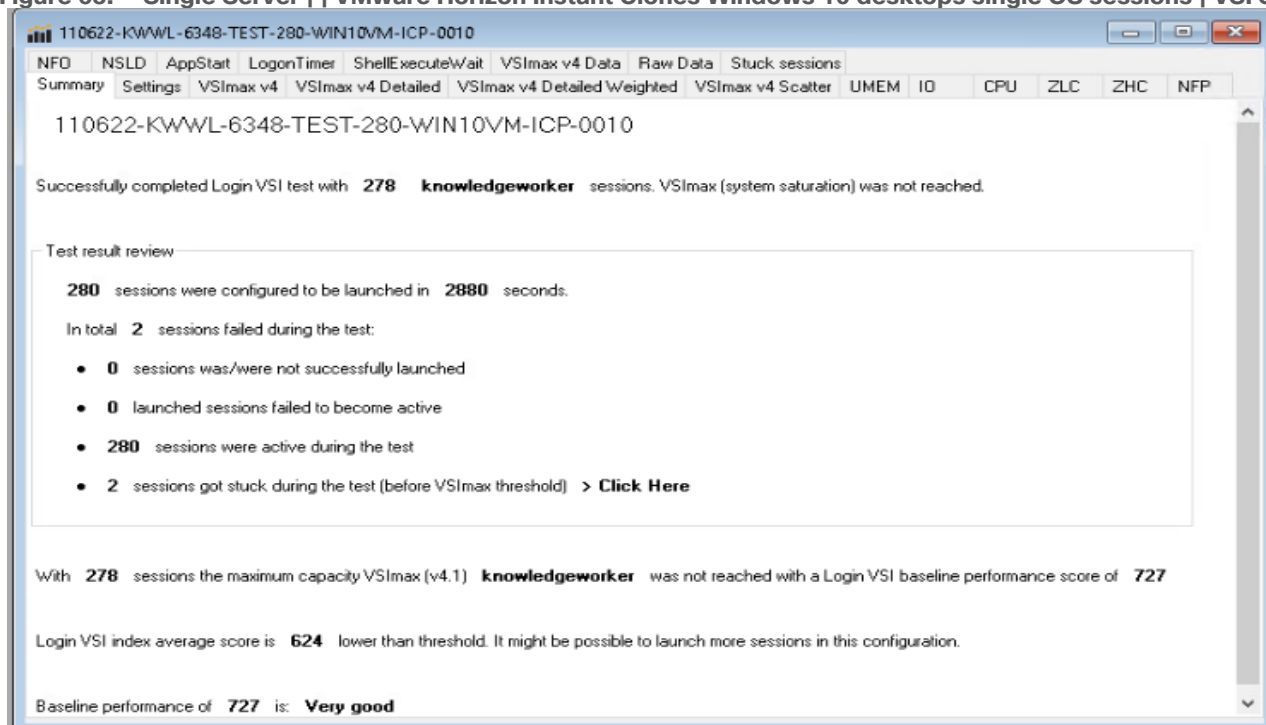




Figure 64. Single Server | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | VSI Score




Performance data for the server running the workload is as follows:


**Figure 65. Single Server Recommended Maximum Workload | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host CPU Utilization**



**Figure 66. Single Server Recommended Maximum Workload | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host Memory Utilization**

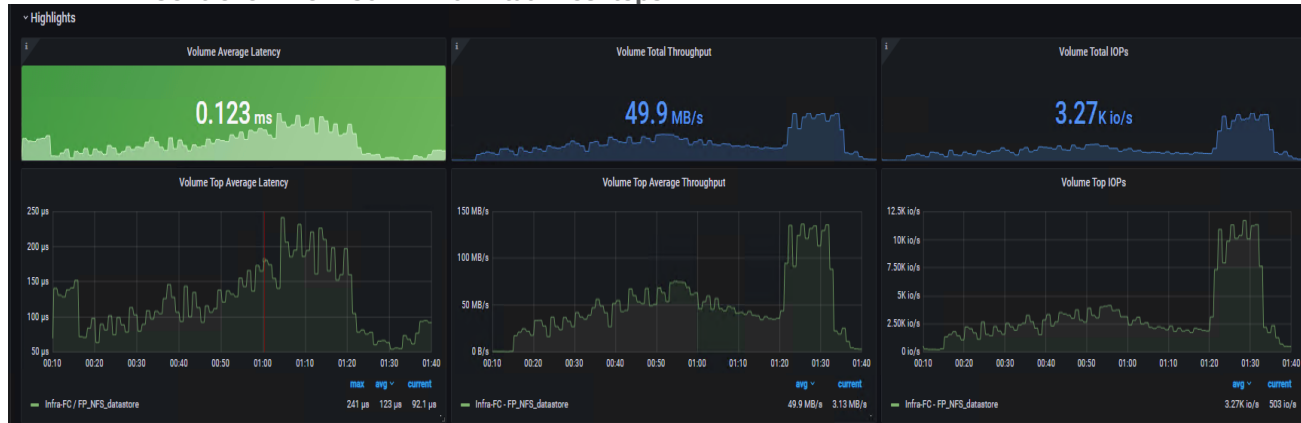


**Figure 67. Single Server | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host Network Utilization**

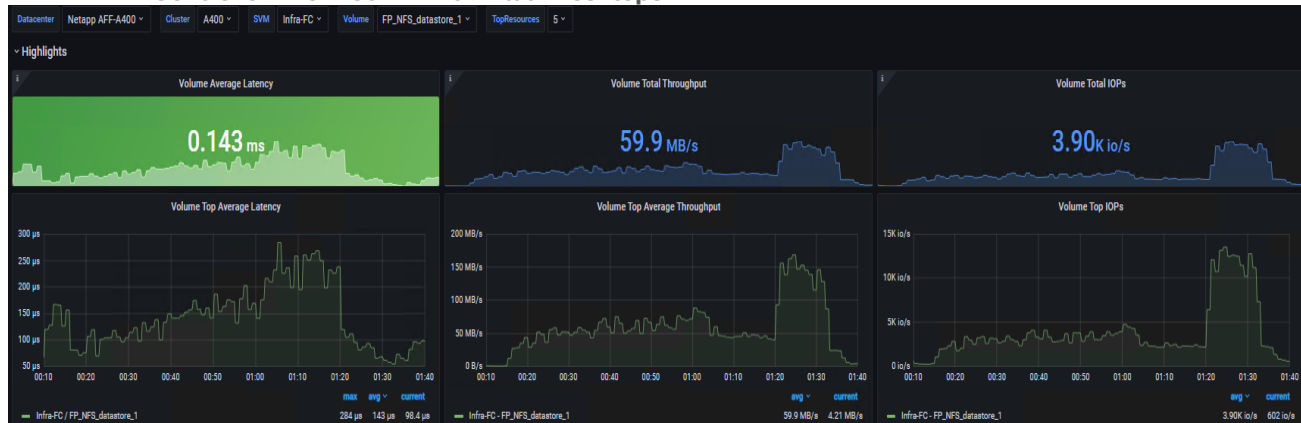


**Storage Charts for 280 Users Win10 Virtual Desktops Instant Clones Testing**

**Figure 68. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 1 for 280 Win 10 Virtual Desktops**



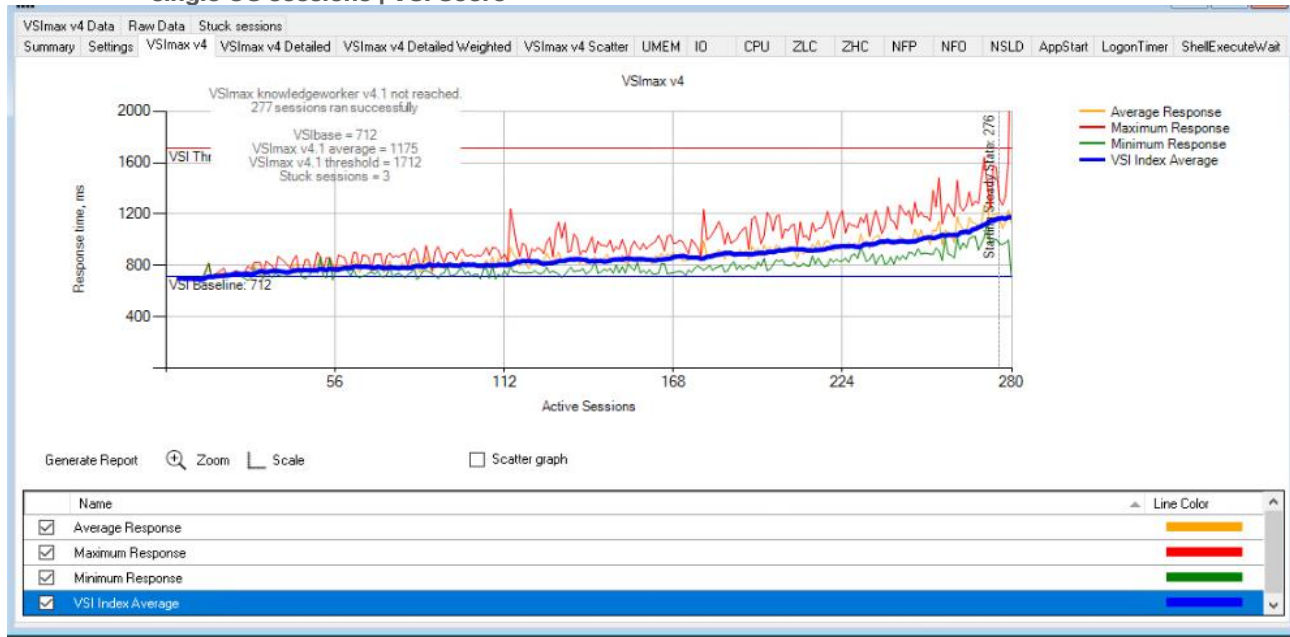
**Figure 69. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 2 for 280 Win 10 Virtual Desktops**



The recommended maximum workload for a Cisco UCS X-Series blade server with dual Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM is 280 Windows 10 64-bit VDI persistent Full Clone VMware Horizon virtual machines with 2 vCPU and 3.5GB RAM.


LoginVSI data is as follows:

**Figure 70. Single Server Recommended Maximum Workload | | VMware Horizon Full Clone Windows 10 desktops single OS sessions | VSI Score**




Performance data for the server running the workload is as follows:

**Figure 71. Single Server Recommended Maximum Workload Single Server | | VMware Horizon Full clone Windows 10 desktops single OS sessions | Host CPU Utilization**






**Figure 72. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 1 for 280 full clone Win 10 Virtual Desktops**



**Figure 73. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 2 for 280 full clone Win 10 Virtual Desktops**



## Full Scale Workload Testing

This section describes the key performance metrics that were captured on the Cisco UCS, during the full-scale testing. Full Scale testing was done with following Workloads using 8 Cisco UCS X210C M6 Blade Servers, configured in a single ESXi Host Pool, and designed to support single Host failure (N+1 Fault tolerance):

- 2600 VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi-session OS sessions
- 1800 VMware Horizon Instant clones Windows 10 Virtual Desktops
- 1800 VMware Horizon Full Clone Windows 10 Virtual Desktops

To achieve the target, sessions were launched against each workload set at a time. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

### Full Scale Recommended Maximum Workload Testing for VMware Horizon Multi Session Remote Desktop Server Hosted (RDSH) Sessions with 2600 Users

This section describes the key performance metrics that were captured on the Cisco UCS and NetApp Storage AFF A400 array during the full-scale testing with 2600 VMware instant Clone Remoted Desktop Server Hosted (RDSH) Sessions using 8 X-Series blades in a single RDS pool.

The workload for the test is 2600 RDSH User sessions and to achieve the target, sessions were launched against all workload hosts concurrently. As per the Cisco Test Protocol for VDI solutions, all sessions were

launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

The configured system efficiently and effectively delivered the following results:

**Figure 74. Full Scale | 2600 Users | Cluster Test | VMware Horizon RDSH Sessions Instant Clones RDSH Server sessions | VSI Score**

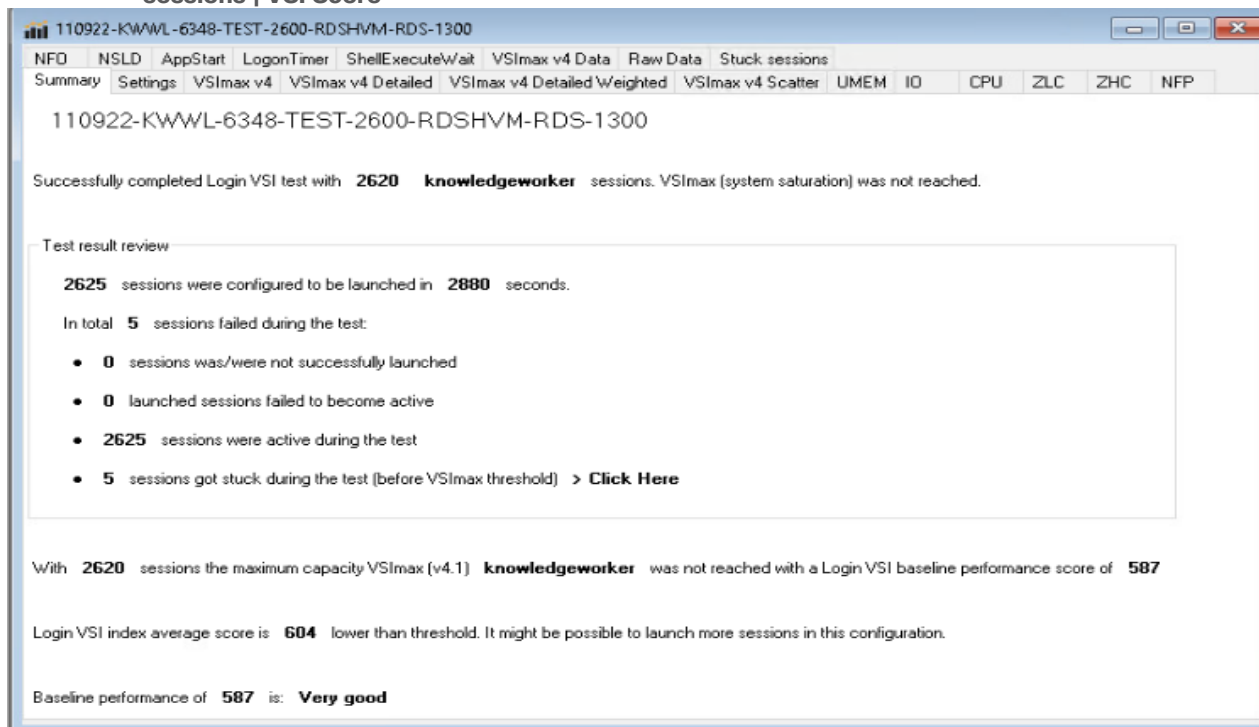


Figure 75. Full Scale | 2600 Users | Cluster Test | VMware Horizon RDSH Server Sessions Instant Clones | VSI Score




Figure 76. ESXTOP CPU Util% for 8 Hosts for 2600 Cluster Test Remote Desktop Server Hosted (RDSH) Sessions Test




Figure 77. Memory Non-Kernal Mbytes for 8 Hosts Cluster Test for 2600 Remote Desktop Server Hosted (RDSH) Sessions Test







Figure 78. Network Util for 8 Hosts Cluster Test for 2600 Remote Desktop Server Hosted (RDSH) Sessions Test




**Figure 79. RDS Server VM Processor Times for 2600 RDSH Cluster Test: From RDS Server 150: Test -01**




**Figure 80. RDS Server VM Processor Times for 2600 RDSH Cluster Test: From RDS Server -50 : Test -01**



**Figure 81. RDS Server VM Processor Times for 2600 RDSH Cluster Test: From RDS Server -168: Test -01**

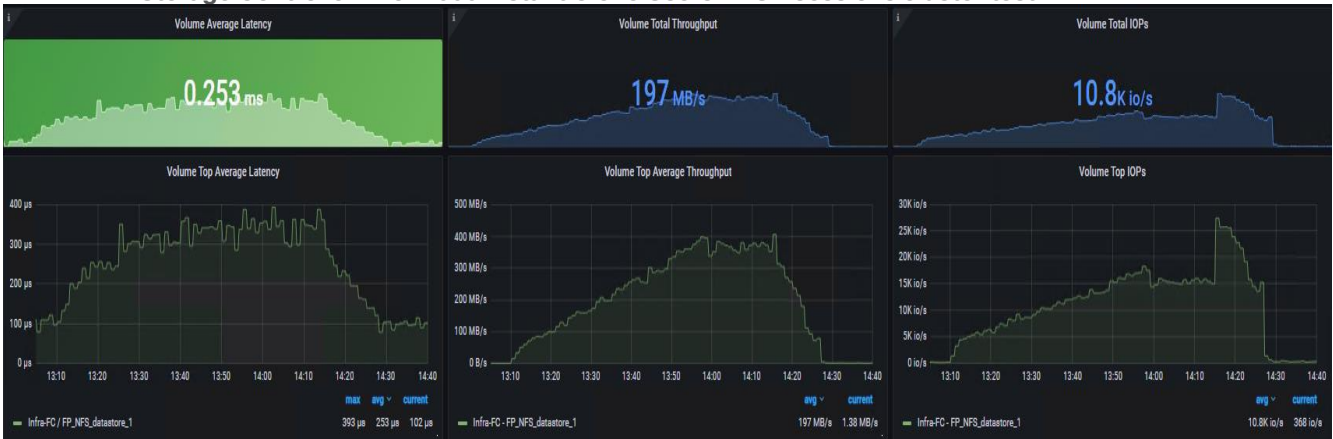


**Figure 82. Login VSI End User Experience Comparison for 3 repeat tests for Full Scale 2600 multi session RDSH Cluster Tests**




# AFF A400 Storage Charts for 2600 Remote Desktop Server Hosted (RDSH) Sessions Cluster Full Scale Test.

**Figure 83. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 1 for 2600 Instant Clone Users RDSH sessions cluster test**



**Figure 84. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 2 for 2600 Instant Clone Users RDSH sessions cluster test**



**Figure 85. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for SMB Share for 2600 Instant Clone Users RDSH sessions cluster test**




Figure 86. Full Scale | 1800 Users | VMware Horizon Win 10 Virtual Desktops Instant Clone single- OS Machines | LoginVSI VSI Score for 1800 Instant Clones Cluster Test.


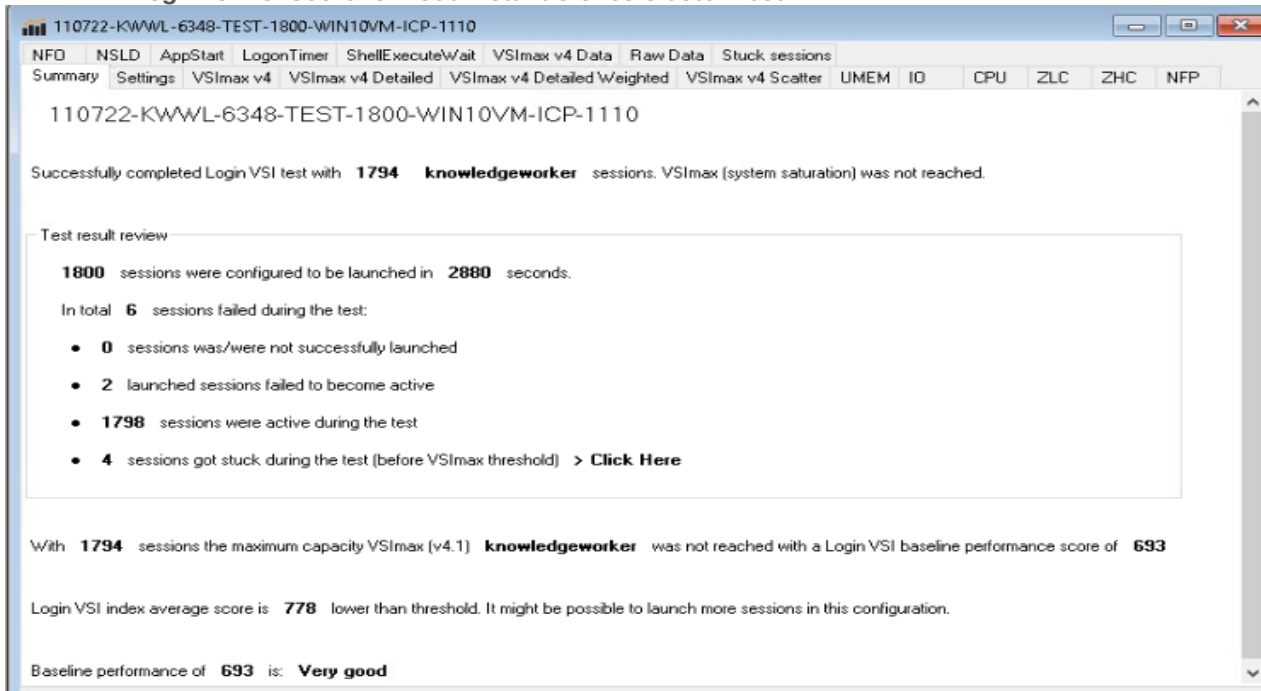



Figure 87. Full Scale | 1800 Users | VMware Horizon Win 10 Virtual Desktops Instant Clone single- OS Machines | LoginVSI VSI Score for 1800 Instant Clones Cluster Test







**Figure 88. Login VSI End User Experience Comparison for 3 repeat tests for FullScale 1800 single session Win 10 Instant Clone Virtual Desktop Cluster Test**




**Figure 89. Full Scale | 1800 Users | VMware Horizon Instant Clone Virtual Desktops Single-session OS machine | ESXTOP CPU Util% for 8 hosts on Cluster**



**Figure 90. Full Scale | 1800 Users | VMware Horizon Instant Clone Virtual Desktops Single-session OS machine | Host Memory Utilization for 8 Hosts on Cluster**



**Figure 91. Full Scale | 1800 Users | VMware Horizon Windows 10 Virtual Desktops | Host Network Utilization for 8 hosts on Cluster**



# AFF A400 Storage Charts for VMware Horizon 1800 Windows 10 Instant Clones Virtual Desktops Cluster Full Scale Test.

**Figure 92. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage Controller 1 for 1800 Instant Clone Users Cluster Test**



**Figure 93. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage Controller 2 for 1800 Instant Clone Users Cluster Test**



**Figure 94. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 SMB share for 1800 Instant Clone Users Cluster Test**



## Full Clone 1800 Users WIN 10 VDI testing: Cluster Test

Figure 95. Full Scale | 1800 Users | VMware Horizon Full clones single VM Login VSI End User Experience VSI Score

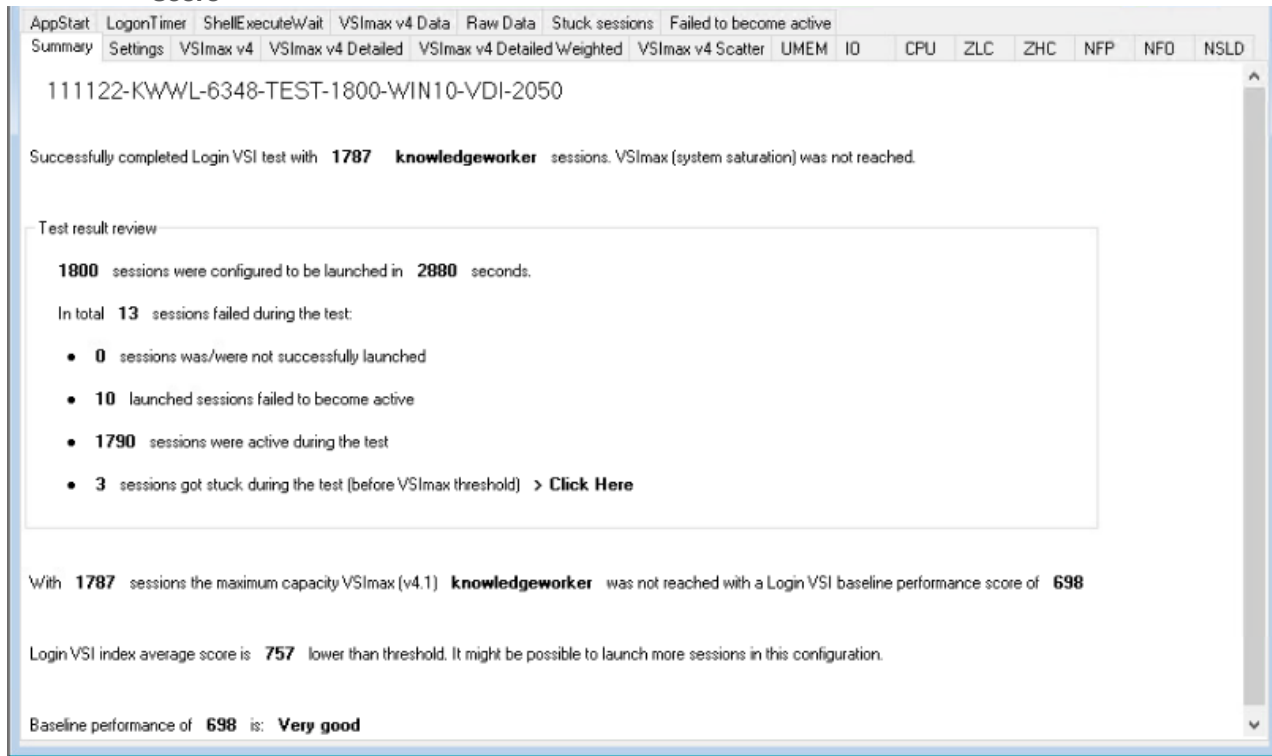




Figure 96. Full Scale | 1800 Users | VMware Horizon Full Clones Single VM Login VSI End User Experience VSI Score



**Figure 97. Login VSI End User Experience for 3 Repeat Tests: Full Scale test 1800 Full Clone Win10 Virtual Desktops Comparison**



**Figure 98. ESXTOP CPU Util% for 1800 Full Clone Cluster Test Win10 Virtual Desktops for 8 hosts on Cluster**




Figure 99. ESXTOP Memory util% for 1800 Full Clone Cluster Test Non-Kernel Mbytes for 8 hosts on Cluster





Figure 100. Full Scale | 1800 Users | VMware Horizon Full Clone Virtual Desktops Single-Session OS Machine | Network Util% for 8 hosts on Cluster



# AFF A400 Storage Charts for VMware Horizon 1800 Windows 10 Full Clones Virtual Desktops Cluster Full Scale Test

Figure 101. Volume Average Latency, Volume Total Throughput and Volume Total IPS AFF A400 Controller 1 for 1800 Full Clone Users Cluster Test




Figure 102. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Controller 2 for 1800 Full Clone Users Cluster Test

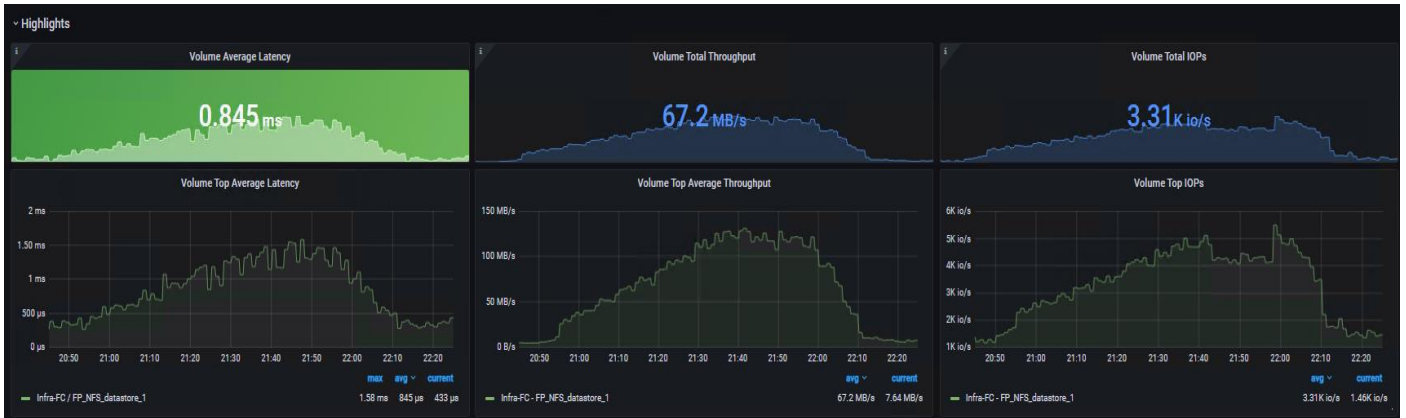


Figure 103. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 SMB Share for 1800 Full Clone Users Cluster Test



---

## Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 2600 users, which this reference architecture has successfully tested. This 2600-seat solution provides a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

### Cisco UCS System Scalability

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested:

- Cisco UCS Manager Software supports up to 20 Cisco UCS chassis within a single Cisco UCS domain with Cisco UCS 6454 Fabric Interconnect. A single UCS domain can grow to 160 blades for an enterprise deployment.
- Cisco UCS Central, the manager of managers, extends UCS domains and vastly increases the reach of the Cisco UCS system. Simplify daily operations by centrally managing and automating routine tasks and expediting problem resolution. Our powerful platform eliminates disparate management environments. Use it to support up to 10,000 Cisco UCS servers (blade, rack, composable, and Mini) and manage multiple Cisco UCS instances or domains across globally-distributed locations.
- As scale grows, the value of the combined UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100 percent of the time.
- To accommodate the Cisco Nexus 9000 upstream connectivity in the way we describe in the network configuration section, two Ethernet uplinks are needed to be configured on the Cisco UCS 6454 Fabric Interconnect.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the NetApp scaling section. Please refer the NetApp section that follows this one for scalability guidelines.

### NetApp FAS Storage Guidelines for Scale Desktop Virtualization Workloads

Storage sizing has three steps:

1. Gathering solution requirements.
2. Estimating storage capacity and performance.
3. Obtaining recommendations for the storage configuration.

### Solution Assessment

Assessment is an important first step. Liquidware Labs Stratusphere FIT, and Lakeside VDI Assessment are recommended to collect network, server, and storage requirements. NetApp has contracted with Liquidware Labs to provide free licenses to NetApp employees and channel partners. For information on how to obtain software and licenses, refer to this [FAQ](#). Liquidware Labs also provides a storage template that fits the NetApp system performance modeler. For guidelines on how to use Stratusphere FIT and the NetApp custom report template, refer to [TR-3902: Guidelines for Virtual Desktop Storage Profiling](#).

Virtual desktop sizing depends on the following:

- The number of the seats
- The VM workload (applications, VM size, and VM OS)
- The connection broker (VMWare Horizon Remote Desktop Server Hosted (RDSH) Sessions & Win 10 Virtual Desktops)



- The hypervisor type (VMware vSphere, Citrix Hypervisor, or Hyper-V)
- The provisioning method (NetApp clone, and Full Clone Provisioning)
- Future storage growth
- Disaster recovery requirements
- User home directories

NetApp has developed a sizing tool called the System Performance Modeler (SPM) that simplifies the process of performance sizing for NetApp systems. It has a step-by-step wizard to support varied workload requirements and provides recommendations for meeting your performance needs.

Storage sizing has two factors: capacity and performance. NetApp recommends using the NetApp SPM tool to size the virtual desktop solution. To use this tool, contact NetApp partners and NetApp sales engineers who have the access to SPM. When using the NetApp SPM to size a solution, NetApp recommends separately sizing the VDI workload (including the write cache and personal vDisk if used), and the CIFS profile and home directory workload. When sizing CIFS, NetApp recommends sizing with a heavy user workload. Eighty percent concurrency was assumed in this solution.

### Performance Considerations

The collection of performance requirements is a critical step. After using Liquidware Labs Stratusphere FIT and Lakeside VDI Assessment to gather I/O requirements, contact the NetApp account team to obtain recommended software and hardware configurations.

Size, the read/write ratio, and random or sequential reads comprise the I/O considerations. We use 90 percent write and 10 percent read for PVS workload. Storage CPU utilization must also be considered. Use [Table 27](#) as guidance for your sizing calculations for a PVS workload when using a LoginVSI heavy workload.

**Table 27. Typical IOPS without RamCache plus Overflow Feature**

	Boot IOPS	Login IOPS	Steady IOPS
Write Cache (NFS)	8-10	9	7.5
vDisk (CIFS SMB 3)	0.5	0	0
Infrastructure (NFS)	2	1.5	0

### Scalability of VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions and Win 10 Virtual Desktops Configuration

Remote Desktop Server Hosted (RDSH) Sessions & Win10 Virtual Desktops environments can scale to large numbers. When implementing Remote Desktop Server Hosted (RDSH) Sessions & Win 10 Virtual Desktops, consider the following in scaling the number of hosted shared and hosted virtual desktops:

- Types of storage in your environment
- Types of desktops that will be deployed
- Data protection requirements

When designing and deploying this CVD environment Cisco and VMware Horizon recommends using N+1 schema for virtualization host servers to accommodate resiliency. In all Reference Architectures (such as this CVD), this recommendation is applied to all host servers.

---

## Summary

FlexPod delivers a platform for Enterprise End User Computing deployments and cloud datacenters using Cisco UCS Blade and Rack Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS 9100 Fibre Channel switches, and NetApp Storage AFF A400 Storage Array. FlexPod is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlexPod provides to customers wishing to deploy enterprise-class VDI.

---

## About the Author

### **Ramesh Guduru, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.**

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

Ramesh is a subject matter expert on Desktop/Server virtualization, Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA/AMD Graphics.

### **Ruchika Lahoti, Technical Marketing Engineer, NetApp.**

Ruchika has more than five years of experience in the IT industry. She focusses on FlexPod Hybrid Cloud Infrastructure, implementation, validation, and automation. Ruchika holds a bachelors' degree in Computer Science.

## Acknowledgements

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

- Jeff Nichols, Leader-Technical Marketing Engineering, Cisco Systems Inc.

---

## References

This section provides links to additional information for each partner's solution component of this document.

### Cisco UCS X-Series Servers

- <https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-x-series-modular-system/series.html>
- <https://www.cisco.com/site/us/en/products/computing/servers-unified-computing-systems/ucs-x-series-modular-systems/index.html>
- [https://intersight.com/help/saas/resources/cisco\\_x\\_series\\_management\\_guide](https://intersight.com/help/saas/resources/cisco_x_series_management_guide)

### Cisco Intersight Configuration Guides

- <https://www.cisco.com/c/en/us/support/servers-unified-computing/intersight/products-installation-guides-list.html>
- [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/Intersight/b\\_Intersight\\_Managed\\_Mode\\_Configuration\\_Guide.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide.html)
- [https://intersight.com/help/saas/supported\\_systems#supported\\_hardware\\_for\\_intersight\\_managed\\_mode](https://intersight.com/help/saas/supported_systems#supported_hardware_for_intersight_managed_mode)

### Cisco Nexus Switching References

- <http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html>
- <http://www.cisco.com/c/en/us/products/switches/nexus-93180YC-FX-switch/index.html>

### Cisco MDS 9000 Service Switch References

- <http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html>
- <http://www.cisco.com/c/en/us/products/storage-networking/product-listing.html>
- <http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/datasheet-listing.html>

### VMWare References

- <https://docs.vmware.com/en/VMware-Horizon/8-2209/rn/vmware-horizon-8-2209-release-notes/index.html>
- <https://docs.vmware.com/en/VMware-Horizon-Client-for-Windows/2209/rn/vmware-horizon-client-for-windows-2209-release-notes/index.html>
- <https://techzone.vmware.com/resource/quick-start-tutorial-vmware-horizon-8#components-and-architecture>
- <https://techzone.vmware.com/resource/best-practices-published-applications-and-desktops-vmware-horizon-and-vmware-horizon-apps>
- <https://techzone.vmware.com/resource/what-vmware-horizon>

### FlexPod

- <https://www.flexpod.com>
- [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_esxi65u1\\_n9fc.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html)

---

## VMware References

- <https://docs.vmware.com/en/VMware-vSphere/index.html>
- <https://labs.vmware.com/flings/vmware-os-optimization-tool>
- <https://pubs.vmware.com/view-51/index.jsp?topic=%2Fcom.vmware.view.planning.doc%2FGUID-6CAFF558-A0AB-4894-A0F4-97CF556784A9.html>

## Microsoft References

- [https://technet.microsoft.com/en-us/library/hh831620\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/hh831620(v=ws.11).aspx)
- [https://technet.microsoft.com/en-us/library/dn281793\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/dn281793(v=ws.11).aspx)
- <https://support.microsoft.com/en-us/kb/2833839>
- [https://technet.microsoft.com/en-us/library/hh831447\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/hh831447(v=ws.11).aspx)

## Login VSI Documentation

- [https://www.loginvsi.com/documentation/Main\\_Page](https://www.loginvsi.com/documentation/Main_Page)
- [https://www.loginvsi.com/documentation/Start\\_your\\_first\\_test](https://www.loginvsi.com/documentation/Start_your_first_test)

## NetApp Reference Documents

- <http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>
- <http://www.netapp.com/us/products/data-management-software/ontap.aspx>
- <https://mysupport.netapp.com/documentation/docweb/index.html?productID=62379&language=en-US>
- <http://www.netapp.com/us/products/management-software/>
- <http://www.netapp.com/us/products/management-software/vsc/>

---

## Appendices

The appendices are as follows:

[Appendix A - Cisco Switch Configuration](#)

[Appendix B - Glossary of Acronyms](#)

[Appendix C - Glossary of Terms](#)

### Appendix A—Cisco Switch Configuration

This chapter contains the following:

- [Network Configuration](#)
- [Fibre Channel Configuration](#)

#### Network Configuration

##### N93180YC-FX -A Configuration

```
!Command: show running-config
version 9.3 (7a)I1(3b)
switchname DV-Pod-2-N9K-A
class-map type network-qos class-platinum
match qos-group 2
class-map type network-qos class-all-flood
match qos-group 2
class-map type network-qos system_nq_policy
match qos-group 2
class-map type network-qos class-ip-multicast
match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-platinum
    mtu 9216
  class type network-qos class-default
    mtu 9216
vdc DV-Pod-2-N9K-A id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
feature telnet
cfs ipv4 distribute
```

```
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock protocol none vdc 1
no password strength-check
username admin password 5 $1$tYYajkfc$7P7nLjWYvfTWAlvFDnwJZ. role network-admin
ip domain-lookup
ip access-list NFS_VLAN63
  10 permit ip 10.10.63.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-A_64
  10 permit ip 10.10.64.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-B_65
  10 permit ip 10.10.65.0 255.255.255.0 any
  20 deny ip any any
class-map type qos match-any class-platinum
  match cos 5
policy-map type qos jumbo
  class class-platinum
    set qos-group 2
  class class-default
    set qos-group 0
system qos
  service-policy type network-qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0xf747567d6cfecf362a9641ac6f3cefc9 priv
0xf747567d6cfecf362a9641ac6f3cefc9 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.81.254.202
vlan 1-2,60-70,102,164
vlan 60
  name In-Band-Mgmt
vlan 61
```

```
name Infra-Mgmt
vlan 62
  name CIFS
vlan 63
  name NFS
vlan 64
  name iSCSI-A
vlan 65
  name iSCSI-B
vlan 66
  name vMotion
vlan 67
  name N1KV
vlan 68
  name LauncherPXE
vlan 69
  name Launcher81
vlan 70
  name other-3
vlan 102
  name VDI
vlan 164
  name Out-Of-Band-Mgmt
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 10.29.164.66 source 10.29.164.65
  delay restore 150
  peer-gateway
  auto-recovery
interface Vlan1
  no ip redirects
  no ipv6 redirects
```



```
interface Vlan2
  description Default native vlan 2
  no ip redirects
  no ipv6 redirects
interface Vlan60
  description Out of Band Management vlan 60
  no shutdown
  no ip redirects
  ip address 10.10.60.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 60
    preempt
    priority 110
    ip 10.10.60.1
interface Vlan61
  description Infrastructure vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.61.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 61
  description NetApp_AFF400_Node-01_port_e0e_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216
  channel-group 18 mode active
interface Ethernet1/4
  description NetApp_AFF400_Node-01_port_e4a_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216
  channel-group 18 mode active
interface Ethernet1/5
  description NetApp_AFF400_Node-02_port_e0f_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216
  channel-group 13 mode active
interface Ethernet1/6
  description NetApp_AFF400_Node-02_port_e4a_CIFS
```

```
switchport mode trunk
switchport trunk allowed vlan 62,64-65
mtu 9216
channel-group 13 mode active
interface Ethernet1/7
description NetApp_AFF400_Node-01_port_e0f_CIFS
switchport mode trunk
switchport trunk allowed vlan 62,64-65
mtu 9216
channel-group 17 mode active
interface Ethernet1/8
description NetApp_AFF400_Node-01_port_e1a_CIFS
switchport mode trunk
switchport trunk allowed vlan 62,64-65
mtu 9216
channel-group 17 mode active
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
description Uplink_from_FI-A_6k
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
mtu 9216
channel-group 11 mode active
interface Ethernet1/18
description Uplink_from_FI-A_6k
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
mtu 9216
channel-group 11 mode active
interface Ethernet1/19
description Uplink_from_FI-B_6k
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
mtu 9216
channel-group 12 mode active
```

```
interface Ethernet1/20
  description Uplink_from_FI-B_6k
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 12 mode active
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
  description Uplink_from_LoginVSI_Launchers_FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 15 mode active
interface Ethernet1/46
  description Uplink_from_LoginVSI_Launchers_FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 16 mode active
```

```
interface Ethernet1/47
interface Ethernet1/48
    description TOR
    switchport access vlan 164
interface Ethernet1/49
    description VPC Peer Link between 9ks
    switchport mode trunk
    switchport trunk allowed vlan 1-2,60-70,102,164
    channel-group 10 mode active
interface Ethernet1/50
    description VPC Peer Link between 9ks
    switchport mode trunk
    switchport trunk allowed vlan 1-2,60-70,102,164
    channel-group 10 mode active
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
    vrf member management
    ip address 10.29.164.65/24
line console
line vty
boot nxos bootflash://sup-1/n9000-dk9.7.0.3.I1.3b.bin
N93180YC-FX -B Configuration
!Command: show running-config
!Time: Fri Nov 7 16:47:01 2022
version 9.3 (7a)I1(3b)
switchname DV-Pod-2-N9K-B
class-map type network-qos class-platinum
match qos-group 2
class-map type network-qos class-all-flood
match qos-group 2
class-map type network-qos system_nq_policy
match qos-group 2
class-map type network-qos class-ip-multicast
match qos-group 2
policy-map type network-qos jumbo
    class type network-qos class-platinum
        mtu 9216
    class type network-qos class-default
        mtu 9216
```

```
vdc DV-Pod-2-N9K-B id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
feature telnet
cfs ipv4 distribute
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock protocol none vdc 1
no password strength-check
username admin password 5 $1$fp3LrGLC$PF8eML85qkPBgdH/bZAKK/ role network-admin
ip domain-lookup
ip access-list NFS_VLAN63
  10 permit ip 10.10.63.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-A_64
  10 permit ip 10.10.64.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-B_65
  10 permit ip 10.10.65.0 255.255.255.0 any
  20 deny ip any any
class-map type qos match-any class-platinum
  match cos 5
policy-map type qos jumbo
  class class-platinum
    set qos-group 2
  class class-default
    set qos-group 0
system qos
service-policy type network-qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0x13ec164cc65d2b9854d70379681039c8 priv
0x13ec164cc65d2b9854d70379681039c8 localizedkey
```

```
ntp master 8
vlan 1-2,60-70,102,164
vlan 60
    name In-Band-Mgmt
vlan 61
    name Infra-Mgmt
vlan 62
    name CIFS
vlan 63
    name NFS
vlan 64
    name iSCSI-A
vlan 65
    name iSCSI-B
vlan 66
    name vMotion
vlan 68
    name LauncherPXE
vlan 69
    name Launcher81
vlan 70
    name other-3
vlan 102
    name VDI
vlan 173
    name Out-Of-Band-Mgmt
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
    ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
vpc domain 10
    peer-switch
    role priority 10
    peer-keepalive destination 10.29.164.65 source 10.29.164.66
    delay restore 150
    peer-gateway
    auto-recovery
```

```
interface Vlan1
  no ip redirects
  no ipv6 redirects
interface Vlan2
  description Default native vlan 2
  no ip redirects
  no ipv6 redirects
interface Vlan60
  description Out of Band Management vlan 60
  no shutdown
  no ip redirects
  ip address 10.10.60.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 60
    preempt
    priority 110
    ip 10.10.60.1
interface Vlan61
  description Infrastructure vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.61.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 61
    preempt
    ip 10.10.61.1
interface Vlan62
  description CIFS vlan 62
  no shutdown
  no ip redirects
  ip address 10.10.62.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 62
    preempt
    priority 110
    ip 10.10.62.1
interface Vlan63
  description NFS vlan 63
  no shutdown
```

```
no ip redirects
ip address 10.10.63.3/24
no ipv6 redirects
hsrp version 2
hsrp 63
    preempt
    ip 10.10.63.1
interface Vlan64
    description iSCSI Fabric A path vlan 64
    no shutdown
    no ip redirects
    ip address 10.10.64.3/24
    no ipv6 redirects
    hsrp version 2
    hsrp 64
        preempt
        priority 110
        ip 10.10.64.1
interface Vlan65
    description iSCSI Fabric B path vlan 65
    no shutdown
    no ip redirects
    ip address 10.10.65.3/24
    no ipv6 redirects
    hsrp version 2
    hsrp 65
        preempt
        ip 10.10.65.1
interface Vlan66
    description vMotion network vlan 66
    no shutdown
    ip address 10.10.66.3/24
    hsrp version 2
    hsrp 66
        preempt
        ip 10.10.66.1
interface Vlan67
    description vlan 67
    no shutdown
    ip address 10.10.67.3/24
    hsrp version 2
    hsrp 67
```



```
preempt
ip 10.10.67.1
interface Vlan68
description LoginVSI Launchers vlan 68
no shutdown
no ip redirects
ip address 10.10.68.3/24
no ipv6 redirects
hsrp version 2
hsrp 68
preempt
ip 10.10.68.1
interface Vlan69
description LoginVSI Launchers 10.10.81-network vlan 69
no shutdown
no ip redirects
ip address 10.10.81.3/24
no ipv6 redirects
hsrp version 2
hsrp 69
preempt
ip 10.10.81.1
interface Vlan102
description VDI vlan 102
no shutdown
no ip redirects
ip address 10.2.0.3/19
no ipv6 redirects
hsrp version 2
hsrp 102
preempt delay minimum 240
priority 110
timers 1 3
ip 10.2.0.1
ip dhcp relay address 10.10.61.30
interface port-channel10
description VPC-PeerLink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
spanning-tree port type network
vpc peer-link
interface port-channel11
```

```
description FI-A_6k_UCS-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
spanning-tree port type edge trunk
mtu 9216
vpc 11
interface port-channel12
description FI-B_6k_UCS-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
spanning-tree port type edge trunk
mtu 9216
vpc 12
interface port-channel13
description NetApp_AFF400_Node_02_CIFS
switchport mode trunk
switchport trunk allowed vlan 62,64-65
spanning-tree port type edge trunk
mtu 9216
vpc 13
interface port-channel14
description NetApp_AFF400_Node_02_NFS
switchport mode trunk
switchport trunk allowed vlan 63
spanning-tree port type edge trunk
mtu 9216
vpc 14
interface port-channel15
description FI-A_6k_Launchers-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
spanning-tree port type edge trunk
mtu 9216
vpc 15
interface port-channel16
description FI-B_6k_Launchers-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102,164
spanning-tree port type edge trunk
mtu 9216
vpc 16
interface port-channel17
```

```
description NetApp_AFF400_Node_01_CIFS
switchport mode trunk
switchport trunk allowed vlan 62,64-65
spanning-tree port type edge trunk
mtu 9216
vpc 17
interface port-channel18
description NetApp_AFF400_Node-01_port_NFS
switchport mode trunk
switchport trunk allowed vlan 63
spanning-tree port type edge trunk
mtu 9216
vpc 18
interface Ethernet1/1
description NetApp_AFF400_Node-02_port_e0g_NFS
switchport mode trunk
switchport trunk allowed vlan 63
mtu 9216
channel-group 14 mode active
interface Ethernet1/2
description NetApp_AFF400_Node-02_port_elb_NFS
switchport mode trunk
switchport trunk allowed vlan 63
mtu 9216
channel-group 14 mode active
interface Ethernet1/3
description NetApp_AFF400_Node-01_port_e0g_NFS
switchport mode trunk
switchport trunk allowed vlan 63
mtu 9216
channel-group 18 mode active
interface Ethernet1/4
description NetApp_AFF400_Node-01_port_e4b_NFS
switchport mode trunk
switchport trunk allowed vlan 63
mtu 9216
channel-group 18 mode active
interface Ethernet1/5
description NetApp_AFF400_Node-02_port_e0h_CIFS
switchport mode trunk
switchport trunk allowed vlan 62,64-65
mtu 9216
```

```
channel-group 13 mode active
interface Ethernet1/6
  description NetApp_AFF400_Node-02_port_e4b_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216
  channel-group 13 mode active
interface Ethernet1/7
  description NetApp_AFF400_Node-01_port_e0h_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216
  channel-group 17 mode active
interface Ethernet1/8
  description NetApp_AFF400_Node-01_port_elb_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216
  channel-group 17 mode active
interface Ethernet1/9
  description Jumphost ToR
  switchport access vlan 60
  spanning-tree port type edge
  speed 1000
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
  description Uplink_from_FI-A_6k
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 11 mode active
interface Ethernet1/18
  description Uplink_from_FI-A_6k
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
```

```
channel-group 11 mode active
interface Ethernet1/19
  description Uplink_from_FI-B_6k
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 12 mode active
interface Ethernet1/20
  description Uplink_from_FI-B_6k
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
channel-group 12 mode active
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
  description Uplink_from_LoginVSI_Launchers_FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
```

```
channel-group 15 mode active
interface Ethernet1/46
  description Uplink_from_LoginVSI_Launchers_FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 16 mode active
interface Ethernet1/47
interface Ethernet1/48
  description TOR
  switchport access vlan 164
interface Ethernet1/49
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  channel-group 10 mode active
interface Ethernet1/50
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  channel-group 10 mode active
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
  vrf member management
  ip address 10.29.164.66/24
line console
line vty
boot nxos bootflash://sup-1/n9000-dk9.7.0.3.I1.3b.bin
```

## Fibre Channel Configuration

### Cisco MDS 9132T - A Configuration

```
!Command: show running-config
!Time: Wed Nov  7 00:49:39 2022
version 8.1(1)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
```

```
rule 5 permit show feature environment
rule 4 permit show feature hardware
rule 3 permit show feature module
rule 2 permit show feature snmp
rule 1 permit show feature system
no password strength-check
username admin password 5 $1$DDq8vFlx$EwCSM003dlXZ4j1Py9ZoC. role network-admin
ip domain-lookup
ip host MDS-A 10.29.164.238
aaa group server radius radius
snmp-server contact jnichols
snmp-server user admin network-admin auth md5 0x2efbf582e573df2038164f1422c231fe
priv 0x2efbf582e573df2038164f1422c231fe localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1163
snmp-server host 10.155.166.14 traps version 2c public udp-port 1163
snmp-server host 10.29.132.18 traps version 2c public udp-port 1163
snmp-server host 10.29.164.130 traps version 2c public udp-port 1163
snmp-server host 10.29.164.250 traps version 2c public udp-port 1164
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community public group network-operator
vsan database
  vsan 400 name "FlexPod-A"
device-alias database
  device-alias name A400_N1P3 pwnn 50:01:73:80:59:16:01:12
  device-alias name A400_N2P1 pwnn 50:01:73:80:59:16:01:20
  device-alias name A400_N2P3 pwnn 50:01:73:80:59:16:01:22
  device-alias name A400_N3P1 pwnn 50:01:73:80:59:16:01:30
  device-alias name A400_N3P3 pwnn 50:01:73:80:59:16:01:32
  device-alias name VDI-1-hba1 pwnn 20:00:00:25:b5:3a:00:3f
  device-alias name VDI-2-hba1 pwnn 20:00:00:25:b5:3a:00:0f
  device-alias name VDI-3-hba1 pwnn 20:00:00:25:b5:3a:00:1f
  device-alias name VDI-4-hba1 pwnn 20:00:00:25:b5:3a:00:4e
  device-alias name VDI-5-hba1 pwnn 20:00:00:25:b5:3a:00:2e
  device-alias name VDI-6-hba1 pwnn 20:00:00:25:b5:3a:00:3e
  device-alias name VDI-7-hba1 pwnn 20:00:00:25:b5:3a:00:0e
  device-alias name VDI-8-hba1 pwnn 20:00:00:25:b5:3a:00:4d
  device-alias name Infra01-8-hba1 pwnn 20:00:00:25:b5:3a:00:4f
  device-alias name Infra02-16-hba1 pwnn 20:00:00:25:b5:3a:00:2f
```

```

device-alias commit
fcdomain fcid database
  vsan 1 wwn 52:4a:93:72:0d:21:6b:11 fcid 0x290000 dynamic
  vsan 1 wwn 52:4a:93:72:0d:21:6b:10 fcid 0x290100 dynamic
  vsan 1 wwn 20:20:00:2a:6a:d3:df:80 fcid 0x290200 dynamic
  vsan 1 wwn 24:01:00:2a:6a:d3:df:80 fcid 0x290400 dynamic
  vsan 1 wwn 52:4a:93:72:0d:21:6b:00 fcid 0x290400 dynamic
  vsan 1 wwn 50:01:73:80:59:16:01:10 fcid 0x290500 dynamic
  vsan 1 wwn 50:01:73:80:59:16:01:20 fcid 0x290600 dynamic
!
  [A400_N2P1]
  vsan 1 wwn 50:01:73:80:59:16:01:30 fcid 0x290700 dynamic
!
  [A400_N3P1]
  vsan 1 wwn 50:01:73:80:59:16:01:12 fcid 0x290800 dynamic
!
  [A400_N1P3]
  vsan 1 wwn 50:01:73:80:59:16:01:22 fcid 0x290900 dynamic
!
  [A400_N2P3]
  vsan 1 wwn 50:01:73:80:59:16:01:32 fcid 0x290a00 dynamic
!
  [A400_N3P3]
  vsan 400 wwn 50:01:73:80:59:16:01:10 fcid 0xa30400 dynamic
  vsan 400 wwn 50:01:73:80:59:16:01:20 fcid 0xa30400 dynamic
!
  [A400_N2P1]
  vsan 400 wwn 50:01:73:80:59:16:01:30 fcid 0xa30500 dynamic
!
  [A400_N3P1]
  vsan 400 wwn 50:01:73:80:59:16:01:12 fcid 0xa30600 dynamic
!
  [A400_N1P3]
  vsan 400 wwn 50:01:73:80:59:16:01:22 fcid 0xa30700 dynamic
!
  [A400_N2P3]
  vsan 400 wwn 50:01:73:80:59:16:01:32 fcid 0xa30800 dynamic
!
  [A400_N3P3]
  vsan 1 wwn 20:4d:54:7f:ee:83:42:00 fcid 0x290b00 dynamic
  vsan 1 wwn 20:4e:54:7f:ee:83:42:00 fcid 0x290c00 dynamic
  vsan 1 wwn 20:4f:54:7f:ee:83:42:00 fcid 0x290d00 dynamic
  vsan 1 wwn 20:50:54:7f:ee:83:42:00 fcid 0x290e00 dynamic
  vsan 400 wwn 50:0a:09:84:80:d3:67:d3 fcid 0x680000 dynamic
  vsan 400 wwn 20:03:00:a0:98:af:bd:e8 fcid 0x680001 dynamic
!
  [A400-02-0g]
  vsan 400 wwn 50:0a:09:84:80:13:41:27 fcid 0x680100 dynamic
  vsan 400 wwn 20:01:00:a0:98:af:bd:e8 fcid 0x680101 dynamic
!
  [A400-01-0g]
  vsan 400 wwn 20:02:00:de:fb:90:a0:80 fcid 0x680200 dynamic
  vsan 400 wwn 20:03:00:de:fb:90:a0:80 fcid 0x680400 dynamic
  vsan 400 wwn 20:04:00:de:fb:90:a0:80 fcid 0x680400 dynamic

```



```
vsan 400 wwn 20:01:00:de:fb:90:a0:80 fcid 0x680500 dynamic
vsan 400 wwn 20:00:00:25:b5:3a:00:49 fcid 0x680308 dynamic
!
[VDI-29-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1a fcid 0x680415 dynamic
!
[VDI-28-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4b fcid 0x680206 dynamic
!
[VDI-19-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0a fcid 0x680508 dynamic
!
[VDI-27-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0c fcid 0x680307 dynamic
!
[VDI-17-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2c fcid 0x680402 dynamic
!
[VDI-15-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3a fcid 0x680210 dynamic
!
[VDI-26-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4a fcid 0x680505 dynamic
!
[VDI-24-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2a fcid 0x680413 dynamic
!
[VDI-25-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1c fcid 0x680207 dynamic
!
[VDI-18-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3c fcid 0x680502 dynamic
!
[VDI-32-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0b fcid 0x68020b dynamic
!
[VDI-22-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4c fcid 0x680208 dynamic
!
[VDI-14-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:39 fcid 0x680306 dynamic
!
[VDI-30-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0d fcid 0x68040d dynamic
!
[VDI-12-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1e fcid 0x680501 dynamic
!
[VDI-31-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2b fcid 0x680202 dynamic
!
[VDI-20-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0e fcid 0x680203 dynamic
!
[VDI-7-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1b fcid 0x680509 dynamic
!
[VDI-23-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2f fcid 0x680401 dynamic
!
[Infra02-16-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4d fcid 0x680302 dynamic
```

```

!           [VDI-9-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1d fcid 0x680507 dynamic
!           [VDI-13-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3d fcid 0x68040e dynamic
!           [VDI-11-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2d fcid 0x680305 dynamic
!           [VDI-10-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3b fcid 0x680303 dynamic
!           [VDI-21-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0f fcid 0x680201 dynamic
!           [VDI-2-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3f fcid 0x680506 dynamic
!           [VDI-1-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3e fcid 0x680304 dynamic
!           [VDI-6-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4f fcid 0x680406 dynamic
!           [Infra01-8-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1f fcid 0x680204 dynamic
!           [VDI-3-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4e fcid 0x680504 dynamic
!           [VDI-4-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2e fcid 0x68050a dynamic
!           [VDI-5-hba1]
vsan 1 wwn 56:c9:ce:90:0d:e8:24:02 fcid 0x290f00 dynamic
!Active Zone Database Section for vsan 400
zone name A400_VDI-1-hba1 vsan 400
  member pwn 20:00:00:25:b5:3a:00:3f
!           [VDI-1-hba1]
  member pwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
zone name A400_VDI-2-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0f
!           [VDI-2-hba1]
zone name A400_VDI-3-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]

```

```
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:1f
!
[VDI-3-hba1]
zone name A400_VDI-4-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4e
!
[VDI-4-hba1]
zone name A400_VDI-5-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2e
!
[VDI-5-hba1]
zone name A400_VDI-6-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:3e
!
[VDI-6-hba1]
zone name A400_VDI-7-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:0e
!
[VDI-7-hba1]
zone name A400_Infra01-8-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4f
!
[Infra01-8-hba1]
zone name A400_VDI-9-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
```

```
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4d
!
[VDI-9-hba1]
zone name A400_VDI-10-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2d
!
[VDI-10-hba1]
zone name A400_VDI-11-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:3d
!
[VDI-11-hba1]
zone name A400_VDI-12-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:0d
!
[VDI-12-hba1]
zone name A400_VDI-13-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:1d
!
[VDI-13-hba1]
zone name A400_VDI-14-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4c
!
[VDI-14-hba1]
zone name A400_VDI-15-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
```

```
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2c
!
[VDI-15-hba1]
zone name A400_Infra02-16-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2f
!
[Infra02-16-hba1]
zone name A400_VDI-17-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:0c
!
[VDI-17-hba1]
zone name A400_VDI-18-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:1c
!
[VDI-18-hba1]
zone name A400_VDI-19-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4b
!
[VDI-19-hba1]
zone name A400_VDI-20-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2b
!
[VDI-20-hba1]
zone name A400_VDI-21-hba1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
```

```
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:3b
!
[VDI-21-hba1]
zone name A400_VDI-22-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:0b
!
[VDI-22-hba1]
zone name A400_VDI-23-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:1b
!
[VDI-23-hba1]
zone name A400_VDI-24-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:4a
!
[VDI-24-hba1]
zone name A400_VDI-25-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:2a
!
[VDI-25-hba1]
zone name A400_VDI-26-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:3a
!
[VDI-26-hba1]
zone name A400_VDI-27-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
```

```
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:0a
!
[VDI-27-hba1]
zone name A400_VDI-28-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:1a
!
[VDI-28-hba1]
zone name A400_VDI-29-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:49
!
[VDI-29-hba1]
zone name A400_VDI-30-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:39
!
[VDI-30-hba1]
zone name A400_VDI-31-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:1e
!
[VDI-31-hba1]
zone name A400_VDI-32-hba1 vsan 400
member pwn 20:01:00:a0:98:af:bd:e8
!
[A400-01-0g]
member pwn 20:03:00:a0:98:af:bd:e8
!
[A400-02-0g]
member pwn 20:00:00:25:b5:3a:00:3c
!
[VDI-32-hba1]
zoneset name FlexPod_FabricA vsan 400
member A400_VDI-1-hba1
member A400_VDI-2-hba1
```

```
member A400_VDI-3-hba1
member A400_VDI-4-hba1
member A400_VDI-5-hba1
member A400_VDI-6-hba1
member A400_VDI-7-hba1
member A400_Infra01-8-hba1
member A400_VDI-9-hba1
member A400_VDI-10-hba1
member A400_VDI-11-hba1
member A400_VDI-12-hba1
member A400_VDI-13-hba1
member A400_VDI-14-hba1
member A400_VDI-15-hba1
member A400_Infra02-16-hba1
member A400_VDI-17-hba1
member A400_VDI-18-hba1
member A400_VDI-19-hba1
member A400_VDI-20-hba1
member A400_VDI-21-hba1
member A400_VDI-22-hba1
member A400_VDI-23-hba1
member A400_VDI-24-hba1
member A400_VDI-25-hba1
member A400_VDI-26-hba1
member A400_VDI-27-hba1
member A400_VDI-28-hba1
member A400_VDI-29-hba1
member A400_VDI-30-hba1
member A400_VDI-31-hba1
member A400_VDI-32-hba1
zoneset activate name FlexPod_FabricA vsan 400
do clear zone database vsan 400
!Full Zone Database Section for vsan 400
zone name A400_VDI-1-hba1 vsan 400
    member pwn 20:00:00:25:b5:3a:00:3f
!
    [VDI-1-hba1]
    member pwn 20:01:00:a0:98:af:bd:e8
!
    [A400-01-0g]
    member pwn 20:03:00:a0:98:af:bd:e8
!
    [A400-02-0g]
zone name A400_VDI-2-hba1 vsan 400
    member pwn 20:01:00:a0:98:af:bd:e8
```



```
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0f
! [VDI-2-hba1]
zone name A400_VDI-3-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1f
! [VDI-3-hba1]
zone name A400_VDI-4-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4e
! [VDI-4-hba1]
zone name A400_VDI-5-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2e
! [VDI-5-hba1]
zone name A400_VDI-6-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3e
! [VDI-6-hba1]
zone name A400_VDI-7-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0e
! [VDI-7-hba1]
zone name A400_Infra01-8-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
```

```
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1e
! [VDI-31-hba1]
zone name A400_VDI-9-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4d
! [VDI-9-hba1]
zone name A400_VDI-10-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2d
! [VDI-10-hba1]
zone name A400_VDI-11-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3d
! [VDI-11-hba1]
zone name A400_VDI-12-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0d
! [VDI-12-hba1]
zone name A400_VDI-13-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1d
! [VDI-13-hba1]
zone name A400_VDI-14-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
```

```
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4c
! [VDI-14-hba1]
zone name A400_VDI-15-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2c
! [VDI-15-hba1]
zone name A400_Infra02-16-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2f
! [Infra02-16-hba1]
zone name A400_VDI-17-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0c
! [VDI-17-hba1]
zone name A400_VDI-18-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1c
! [VDI-18-hba1]
zone name A400_VDI-19-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4b
! [VDI-19-hba1]
zone name A400_VDI-20-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
```

```
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2b
! [VDI-20-hba1]
zone name A400_VDI-21-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3b
! [VDI-21-hba1]
zone name A400_VDI-22-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0b
! [VDI-22-hba1]
zone name A400_VDI-23-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1b
! [VDI-23-hba1]
zone name A400_VDI-24-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4a
! [VDI-24-hba1]
zone name A400_VDI-25-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2a
! [VDI-25-hba1]
zone name A400_VDI-26-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
```

```
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3a
! [VDI-26-hba1]
zone name A400_VDI-27-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0a
! [VDI-27-hba1]
zone name A400_VDI-28-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1a
! [VDI-28-hba1]
zone name A400_VDI-29-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:49
! [VDI-29-hba1]
zone name A400_VDI-30-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:39
! [VDI-30-hba1]
zone name A400_VDI-31-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1e
! [VDI-31-hba1]
zone name A400_VDI-32-hba1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
```

```
! [A400-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
! [A400-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3c
! [VDI-32-hba1]
zoneset name FlexPod_FabricA vsan 400
  member A400_VDI-1-hba1
  member A400_VDI-2-hba1
  member A400_VDI-3-hba1
  member A400_VDI-4-hba1
  member A400_VDI-5-hba1
  member A400_VDI-6-hba1
  member A400_VDI-7-hba1
  member A400_Infra01-8-hba1
  member A400_VDI-9-hba1
  member A400_VDI-10-hba1
  member A400_VDI-11-hba1
  member A400_VDI-12-hba1
  member A400_VDI-13-hba1
  member A400_VDI-14-hba1
  member A400_VDI-15-hba1
  member A400_Infra02-16-hba1
  member A400_VDI-17-hba1
  member A400_VDI-18-hba1
  member A400_VDI-19-hba1
  member A400_VDI-20-hba1
  member A400_VDI-21-hba1
  member A400_VDI-22-hba1
  member A400_VDI-23-hba1
  member A400_VDI-24-hba1
  member A400_VDI-25-hba1
  member A400_VDI-26-hba1
  member A400_VDI-27-hba1
  member A400_VDI-28-hba1
  member A400_VDI-29-hba1
  member A400_VDI-30-hba1
  member A400_VDI-31-hba1
  member A400_VDI-32-hba1

interface mgmt0
  ip address 10.29.164.238 255.255.255.0
interface port-channel1
```

```
channel mode active
switchport rate-mode dedicated
interface port-channel2
channel mode active
switchport rate-mode dedicated
interface port-channel30
switchport rate-mode dedicated
vsan database
vsan 400 interface fc1/37
vsan 400 interface fc1/38
vsan 400 interface fc1/43
vsan 400 interface fc1/44
vsan 400 interface fc1/45
vsan 400 interface fc1/46
switchname MDS-A
no terminal log-all
line console
terminal width 80
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.8.1.1.bin
boot system bootflash:/m9100-s5ek9-mz.8.1.1.bin
interface fc1/13
switchport speed 8000
interface fc1/14
switchport speed 8000
interface fc1/15
switchport speed 8000
interface fc1/16
switchport speed 8000
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/43
interface fc1/44
interface fc1/45
```

---

```
interface fc1/46
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/17
interface fc1/18
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/47
interface fc1/48
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/19
interface fc1/20
interface fc1/21
```



```
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/1
    switchport trunk mode off
    port-license acquire
    no shutdown
interface fc1/2
    switchport trunk mode off
    port-license acquire
    no shutdown
interface fc1/3
    switchport trunk mode off
    port-license acquire
    no shutdown
interface fc1/4
    switchport trunk mode off
    port-license acquire
    no shutdown
interface fc1/5
    port-license acquire
    no shutdown
interface fc1/6
    port-license acquire
    no shutdown
interface fc1/7
    port-license acquire
    no shutdown
interface fc1/8
    port-license acquire
    no shutdown
interface fc1/9
    port-license acquire
interface fc1/10
    port-license acquire
interface fc1/11
    port-license acquire
interface fc1/12
```

```
    port-license acquire
interface fc1/13
    port-license acquire
    no shutdown
interface fc1/14
    port-license acquire
    no shutdown
interface fc1/15
    port-license acquire
    no shutdown

interface fc1/16
    port-license acquire
    no shutdown
interface fc1/17
    port-license acquire
    channel-group 1 force
    no shutdown
interface fc1/18
    port-license acquire
    channel-group 1 force
    no shutdown
interface fc1/19
    switchport description AFFA400 CTRL-A:01
    port-license acquire
    no shutdown
interface fc1/20
    switchport description AFFA400 CTRL-A:05
    port-license acquire
    no shutdown
interface fc1/21
    switchport description Launcher-FIA
    port-license acquire
    no shutdown
interface fc1/22
    switchport description Launcher-FIA
    port-license acquire
    no shutdown
interface fc1/23
    switchport description Launcher-FIA
    port-license acquire
    no shutdown
```

```
interface fc1/24
  switchport description Launcher-FIA
  port-license acquire
  no shutdown
interface fc1/25
  port-license acquire
  no shutdown
interface fc1/26
  port-license acquire
  no shutdown
interface fc1/27
  port-license acquire
  no shutdown
interface fc1/28
  port-license acquire
  no shutdown
interface fc1/29
  port-license acquire
interface fc1/30
  port-license acquire
interface fc1/31
  port-license acquire
interface fc1/32
  port-license acquire
interface fc1/33
  port-license acquire
interface fc1/34
  port-license acquire
interface fc1/35
  port-license acquire
interface fc1/36
  port-license acquire
interface fc1/37
  switchport trunk mode off
  port-license acquire
  no shutdown
interface fc1/38
  switchport trunk mode off
  port-license acquire
  no shutdown
interface fc1/39
  port-license acquire
```

```
no shutdown
interface fc1/40
  port-license acquire
  no shutdown
interface fc1/41
  port-license acquire
  no shutdown
interface fc1/42
  port-license acquire
  no shutdown
interface fc1/43
  port-license acquire
  no shutdown
interface fc1/44
  port-license acquire
  no shutdown
interface fc1/45
  port-license acquire
  no shutdown
interface fc1/46
  port-license acquire
  no shutdown
interface fc1/47
  port-license acquire
  no shutdown
interface fc1/48
  port-license acquire
  no shutdown
ip default-gateway 10.29.164.1
MDS-A#
```

## Appendix B—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 Adaptive 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**—Datacenter

**DHCP**—Dynamic Host Configuration Protocol

**DM**—Dense-Mode (multicast)

**DMVPN**—Dynamic Multipoint Virtual Private Network

**DMZ**—Demilitarized Zone (firewall/networking construct)

**DNA**—Cisco Digital Network Architecture

**DNS**—Domain Name System

**DORA**—Discover, Offer, Request, ACK (DHCP Process)

**DWDM**—Dense Wavelength Division Multiplexing

---

**ECMP**—Equal Cost Multi Path

**EID**—Endpoint Identifier

**EIGRP**—Enhanced Interior Gateway Routing Protocol

**EMI**—Electromagnetic Interference

**ETR**—Egress Tunnel Router (LISP)

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

**FHR**—First-Hop Router (multicast)

**FHRP**—First-Hop Redundancy Protocol

**FMC**—Cisco Firepower Management Center

**FTD**—Cisco Firepower Threat Defense

**GBAC**—Group-Based Access Control

**GbE**—Gigabit Ethernet

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

**GRE**—Generic Routing Encapsulation

**GRT**—Global Routing Table

**HA**—High-Availability

**HQ**—Headquarters

**HSRP**—Cisco Hot-Standby Routing Protocol

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

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

**ICMP**—Internet Control Message Protocol

**IDF**—Intermediate Distribution Frame; essentially a wiring closet.

**IEEE**—Institute of Electrical and Electronics Engineers

**IETF**—Internet Engineering Task Force

**IGP**—Interior Gateway Protocol

**IID**—Instance-ID (LISP)

**IOE**—Internet of Everything

**IoT**—Internet of Things

**IP**—Internet Protocol

**IPAM**—IP Address Management

**IPS**—Intrusion Prevention System

**IPSec**—Internet Protocol Security

---

**ISE**—Cisco Identity Services Engine

**ISR**—Integrated Services Router

**IS-IS**—Intermediate System to Intermediate System routing protocol

**ITR**—Ingress Tunnel Router (LISP)

**LACP**—Link Aggregation Control Protocol

**LAG**—Link Aggregation Group

**LAN**—Local Area Network

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

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

**LHR**—Last-Hop Router (multicast)

**LISP**—Location Identifier Separation Protocol

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

**MAN**—Metro Area Network

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

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

**MnT**—Monitoring and Troubleshooting Node (Cisco ISE persona)

**MOH**—Music on Hold

**MPLS**—Multiprotocol Label Switching

**MR**—Map-resolver (LISP)

**MS**—Map-server (LISP)

**MSDP**—Multicast Source Discovery Protocol (multicast)

**MTU**—Maximum Transmission Unit

**NAC**—Network Access Control

**NAD**—Network Access Device

**NAT**—Network Address Translation

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

**NFV**—Network Functions Virtualization

**NSF**—Non-Stop Forwarding

**OSI**—Open Systems Interconnection model

**OSPF**—Open Shortest Path First routing protocol

**OT**—Operational Technology

**PAgP**—Port Aggregation Protocol

---

**PAN**—Primary Administration Node (Cisco ISE persona)

**PCI DSS**—Payment Card Industry Data Security Standard

**PD**—Powered Devices (PoE)

**PETR**—Proxy-Egress Tunnel Router (LISP)

**PIM**—Protocol-Independent Multicast

**PITR**—Proxy-Ingress Tunnel Router (LISP)

**PnP**—Plug-n-Play

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

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

**PSE**—Power Sourcing Equipment (PoE)

**PSN**—Policy Service Node (Cisco ISE persona)

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

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

**QoS**—Quality of Service

**RADIUS**—Remote Authentication Dial-In User Service

**REST**—Representational State Transfer

**RFC**—Request for Comments Document (IETF)

**RIB**—Routing Information Base

**RLOC**—Routing Locator (LISP)

**RP**—Rendezvous Point (multicast)

**RP**—Redundancy Port (WLC)

**RP**—Route Processor

**RPF**—Reverse Path Forwarding

**RR**—Route Reflector (BGP)

**RTT**—Round-Trip Time

**SA**—Source Active (multicast)

**SAFI**—Subsequent Address Family Identifiers (BGP)

**SD**—Software-Defined

**SDA**—Cisco Software Defined-Access

**SDN**—Software-Defined Networking

**SFP**—Small Form-Factor Pluggable (1 GbE transceiver)

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



---

**SGACL**—Security-Group ACL

**SGT**—Scalable Group Tag, sometimes reference as Security Group Tag

**SM**—Spare-mode (multicast)

**SNMP**—Simple Network Management Protocol

**SSID**—Service Set Identifier (wireless)

**SSM**—Source-Specific Multicast (PIM)

**SSO**—Stateful Switchover

**STP**—Spanning-tree protocol

**SVI**—Switched Virtual Interface

**SVL**—Cisco StackWise Virtual

**SWIM**—Software Image Management

**SXP**—Scalable Group Tag Exchange Protocol

**Syslog**—System Logging Protocol

**TACACS+**—Terminal Access Controller Access-Control System Plus

**TCP**—Transmission Control Protocol (OSI Layer 4)

**UCS**—Cisco Unified Computing System

**UDP**—User Datagram Protocol (OSI Layer 4)

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

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

**URL**—Uniform Resource Locator

**VLAN**—Virtual Local Area Network

**VM**—Virtual Machine

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

**VNI**—Virtual Network Identifier (VXLAN)

**vPC**—virtual Port Channel (Cisco Nexus)

**VPLS**—Virtual Private LAN Service

**VPN**—Virtual Private Network

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

**VPWS**—Virtual Private Wire Service

**VRF**—Virtual Routing and Forwarding

**VSL**—Virtual Switch Link (Cisco VSS component)

**VSS**—Cisco Virtual Switching System

**VXLAN**—Virtual Extensible LAN

**WAN**—Wide-Area Network

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

**WoL**—Wake-on-LAN

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

## Appendix C—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.

<p><b>aaS/XaaS</b> <b>(IT capability provided as a Service)</b></p>	<p>Some IT capability, X, provided as a service (XaaS). Some benefits are:</p> <ul style="list-style-type: none"><li>• The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it.</li><li>• There are 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.</li><li>• 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.</li><li>• Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes.</li></ul> <p>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.</p> <p>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.</p> <p>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.</p>
<p><b>Ansible</b></p>	<p>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).</p> <p><a href="https://www.ansible.com">https://www.ansible.com</a></p>
<p><b>AWS</b> <b>(Amazon Web Services)</b></p>	<p>Provider of IaaS and PaaS.</p> <p><a href="https://aws.amazon.com">https://aws.amazon.com</a></p>
<p><b>Azure</b></p>	<p>Microsoft IaaS and PaaS.</p> <p><a href="https://azure.microsoft.com/en-gb/">https://azure.microsoft.com/en-gb/</a></p>

---

**Co-located datacenter**

“A colocation center (CoLo)...is a type of datacenter 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](https://en.wikipedia.org/wiki/Colocation_centre)

<b>Containers (Docker)</b>	<p>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).</p> <p><a href="https://www.docker.com">https://www.docker.com</a></p> <p><a href="https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html">https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html</a></p>
<b>DevOps</b>	<p>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.</p> <p><a href="https://en.wikipedia.org/wiki/DevOps">https://en.wikipedia.org/wiki/DevOps</a></p> <p><a href="https://en.wikipedia.org/wiki/CI/CD">https://en.wikipedia.org/wiki/CI/CD</a></p>
<b>Edge compute</b>	<p>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.</p> <p>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.</p> <p><a href="https://en.wikipedia.org/wiki/Mobile_edge_computing">https://en.wikipedia.org/wiki/Mobile_edge_computing</a></p>
<b>IaaS (Infrastructure as-a-Service)</b>	<p>Infrastructure components provided aaS, located in datacenters 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).</p>
<b>IaC (Infrastructure as-Code)</b>	<p>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.</p> <p><a href="https://en.wikipedia.org/wiki/Infrastructure_as_code">https://en.wikipedia.org/wiki/Infrastructure_as_code</a></p>
<b>IAM (Identity and Access Management)</b>	<p>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.</p> <p><a href="https://en.wikipedia.org/wiki/Identity_management">https://en.wikipedia.org/wiki/Identity_management</a></p>
<b>IBM (Cloud)</b>	<p>IBM IaaS and PaaS.</p> <p><a href="https://www.ibm.com/cloud">https://www.ibm.com/cloud</a></p>
<b>Intersight</b>	<p>Cisco Intersight™ is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.</p> <p><a href="https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html">https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html</a></p>

<b>GCP</b> <b>(Google Cloud Platform)</b>	Google IaaS and PaaS. <a href="https://cloud.google.com/gcp">https://cloud.google.com/gcp</a>
<b>Kubernetes</b> <b>(K8s)</b>	Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. <a href="https://kubernetes.io">https://kubernetes.io</a>
<b>Microservices</b>	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. <a href="https://en.wikipedia.org/wiki/Microservices">https://en.wikipedia.org/wiki/Microservices</a>
<b>PaaS</b> <b>(Platform-as-a-Service)</b>	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.
<b>Private on-premises datacenter</b>	A datacenter infrastructure housed within an environment owned by a given enterprise is distinguished from other forms of datacenter, with the implication that the private datacenter 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 datacenter, 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.
<b>REST API</b>	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. <a href="https://en.wikipedia.org/wiki/Representational_state_transfer">https://en.wikipedia.org/wiki/Representational_state_transfer</a>
<b>SaaS</b> <b>(Software-as-a-Service)</b>	End-user applications provided “aaS” over the public Internet, with the underlying software systems and infrastructure owned and managed by the provider.
<b>SAML</b> <b>(Security Assertion Markup Language)</b>	Used in the context of Single-Sign-On (SSO) for exchanging authentication and authorization data between an identity provider, typically an IAM system, and a service provider (some form of SaaS). The SAML protocol exchanges XML documents that contain security assertions used by the aaS for access control decisions. <a href="https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language">https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language</a>
<b>Terraform</b>	An open-source IaC software tool for cloud services, based on declarative configuration files. <a href="https://www.terraform.io">https://www.terraform.io</a>

---

## Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](https://cs.co/en-cvds) 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 APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Datacenter Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW\_P3)

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>.

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