



FlexPod Datacenter with VMware Horizon and VMware vSphere 7 for up to 2300 Seats

Deployment Configuration Guide for Virtual Desktop Infrastructure built on Cisco UCS B200 M6 with 3rd Generation Intel Xeon Scalable Processors, Cisco UCS Manager 4.2.(1), NetApp Storage for VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions, Windows 10 Desktops, and VMware vSphere 7.0 U2 Hypervisor

Published: June 2022



In partnership with:



Document Organization





























This document is organized into the following chapters:

- [Executive Summary](#)
- [Solution Overview](#)
- [Technology Overview](#)
- [Solution Components](#)
- [Storage Configuration](#)
- [Storage Configuration – Boot LUNs](#)
- [Configuration and Installation](#)
- [Master Image Creation for Tested Horizon Deployment Types](#)
- [Test Setup, Configuration, and Load Recommendation](#)
- [Test Procedure](#)
- [Test Results](#)
- [Summary](#)
- [About the Author](#)
- [References](#)
- [Appendices](#)

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>.

Icons Used in this Document

	Layer 3 Switch (Multilayer switch)		Layer 3 Switch Stack		Layer 3 Routed Link		Internet (untrusted)
	Layer 2 Switch		Layer 2 Switch Stack		Layer 2 Switched Link		Private Network or the remainder of campus network (trusted)
	Router		SD-Access Embedded Wireless		Layer 3 EtherChannel Routed Link		Private WAN Circuit (trusted)
	Router		Layer 3 Switch (Multilayer switch)		Layer 2 EtherChannel Switched Link		Wired Endpoint (802.1X)
	StackWise Virtual (SVL) or Virtual Switching System (VSS)		Firewall		Redundancy Port (WLC)		Wireless Endpoint (802.1X)
	WLC (Wireless LAN Controller)		Cisco DNA Center		Multi-box, single logical unit such as HA Pair, VSS, SVL		Wired Endpoint
	AP (Access Point)		Identity Service Engine		Services (DHCP, DNS, AD, NTP, etc)		Wireless Endpoint

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 data center 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 7.0U2 hypervisor platform.

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 2300 End User Compute users. The solution virtualized on Cisco UCS B200 M6 blade server, booting VMware vSphere 7.0 Update 2 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 (2300), pooled/non-persistent hosted virtual Windows 10 Instant Clones desktops (1700) 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 2300-seat 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: Data Center](#)
- [Physical Topology](#)
- [Configuration Guidelines](#)
- [What is FlexPod?](#)

Introduction

The current industry trend in data center 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 B200 M6 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, 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 B200 M6 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 B200 M6 Blade Server , 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 M6 Servers featuring the 3rd Gen Intel Xeon Scalable processors.

Highlights for this design include:

- Deploying and managing Cisco UCS 5108 chassis equipped with Cisco UCS B200 M6 blade server using UCS (Unified Computing System)
- Support for Cisco UCS B200 M6 blade servers with 3rd Gen Intel Xeon Scalable Family processors and 3200 MHz memory
- Support for the Cisco UCS Manager 4.2
- Validation of Cisco Nexus 9000 with NetApp AFF A400 system
- Validation of Cisco MDS 9000 with NetApp AFF A400 system
- Support for NetApp Storage AFF A400 with ONTAP version 9.10.1P1
- VMware Horizon 2111 Remote Desktop Server Hosted Sessions
- VMware Horizon 2111 Horizon instant clone virtual machines
- VMware Horizon 2111 Horizon persistent full desktops
- Support for VMware vSphere 7.0 U2
- Fully automated solution deployment covering FlexPod infrastructure and vSphere virtualization

FlexPod Cisco Validated Design Advantages for VDI

The data center 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 for pre-designed 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 Data Center (small failure domains)
- Service Provider Data Center (small failure domains)
- Commercial Data Center
- 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 B200 M6 blade servers dual 32-core 2.0 GHz Intel Xeon (Gold 6338) 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 6238 32-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 B200 M6 Cisco UCS 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 2300 user Remote Desktop Hosted (RDSH) Sessions and 1700 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 2300 user Remote Desktop Hosted (RDSH) Sessions and 1700 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 1700 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 7.0U2 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 data center management:** Cisco maintains industry leadership with the new Cisco UCS Manager 4.2(1f) 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.
- **Cisco 40G Fabric:** Our 40G unified fabric story gets additional validation on 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: Data Center

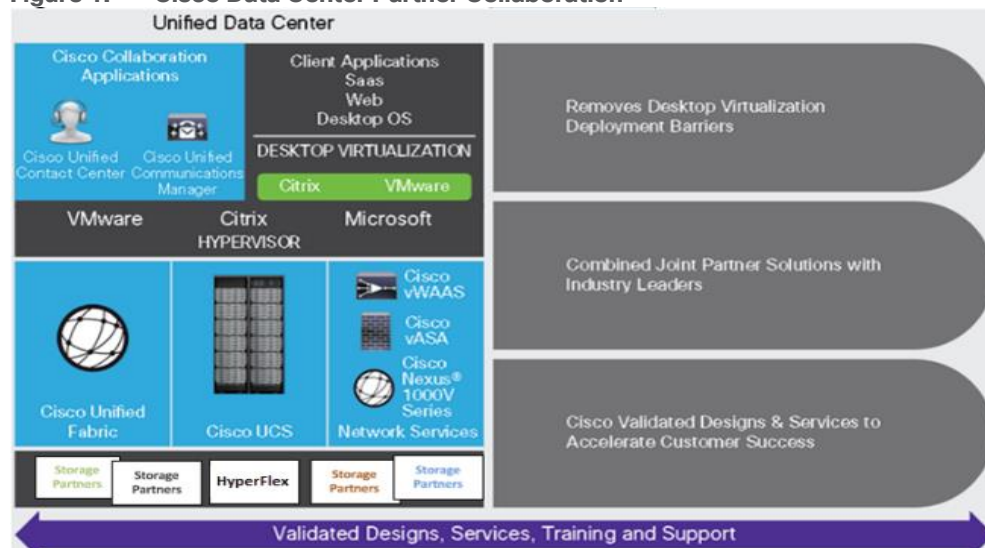
The Evolving Workplace

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 Data Center 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 data center 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 data center 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 data center 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 data center infrastructure (Cisco UCS and Cisco Nexus Family solutions) for desktop virtualization provides strong data center, 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 2300 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, data center applications, and cloud computing.

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

Savings and Success

The simplified, secure, scalable Cisco data center 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

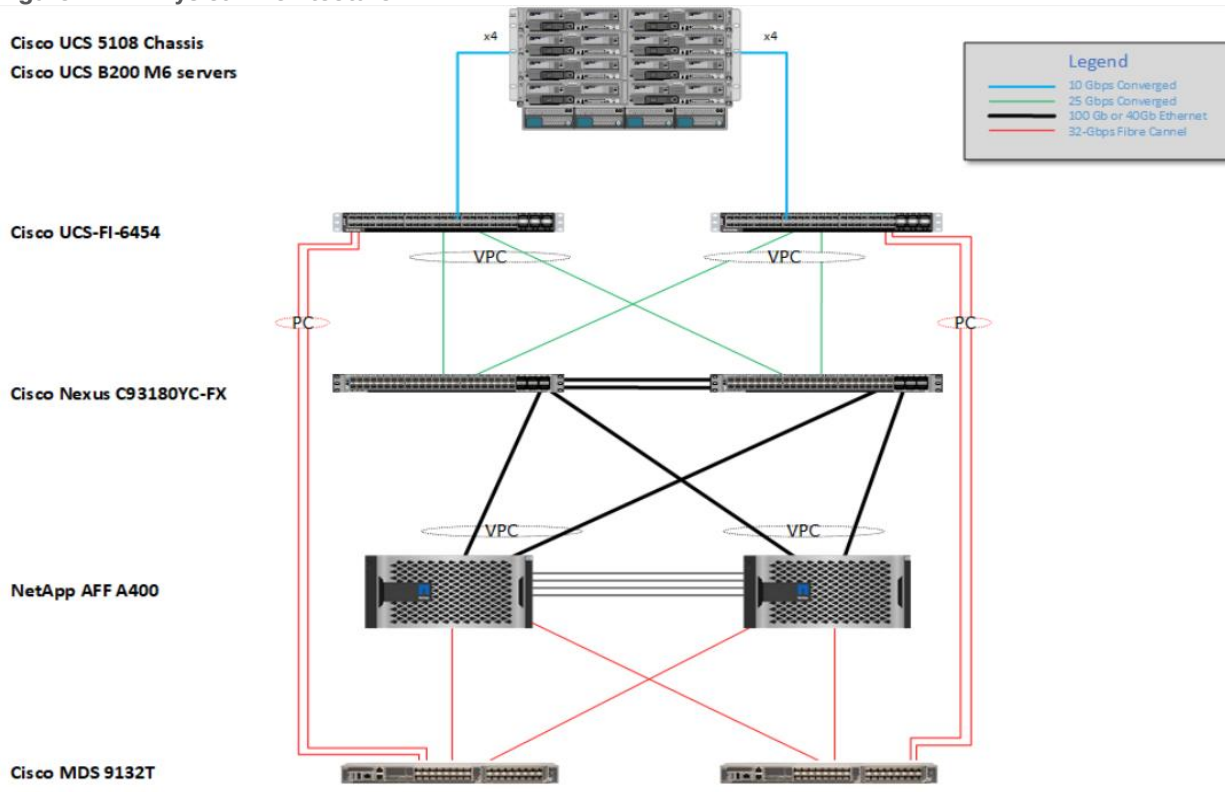
The following are some typical use cases:

- 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
- Eight Cisco UCS B200 M6 Blade Servers (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 7.02.

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:

- 2300 Random Hosted Shared (RDSH) Server 2019 user sessions with Microsoft Office 2016 (Instant Clones)
- 1700 Random Pooled Windows 10 Desktops with Microsoft Office 2016 (Instant Clone virtual machines)
- 1700 Static Full Copy Windows 10 Desktops with Microsoft Office 2016 (Full Clone 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 2300 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 data center rack or be deployed according to a customer's data center 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



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

Technology Overview

This chapter contains the following:

- [Cisco Unified Computing System](#)

Cisco Unified Computing System

This subject contains the following:

- [Cisco UCS Differentiators](#)
- [Cisco UCS Manager](#)
- [Cisco Intersight](#)

Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10-100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform with a unified management domain for managing all resources.

Cisco Unified Computing System consists of the following subsystems:

- **Compute** - The compute piece of the system incorporates servers based on the Second-Generation Intel® Xeon® Scalable processors. Servers are available in blade and rack form factor, managed by Cisco UCS Manager.
- **Network** - The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lower costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.
- **Virtualization** - The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtual environments to support evolving business needs.
- **Storage access** - Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.
- **Management:** The system uniquely integrates compute, network, and storage access subsystems, enabling it to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager increases IT staff productivity by enabling storage, network, and server administrators to collaborate on Service Profiles that define the desired physical configurations and infrastructure policies for applications. Service Profiles increase business agility by enabling IT to automate and provision resources in minutes instead of days.

Cisco UCS Differentiators

Cisco Unified Computing System is revolutionizing the way servers are managed in the datacenter. The following are the unique differentiators of Cisco Unified Computing System and Cisco UCS Manager:

- **Embedded Management** – In Cisco UCS, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating the need for any external physical or virtual devices to manage the servers.
- **Unified Fabric** – In Cisco UCS, from blade server chassis or rack servers to FI, there is a single Ethernet cable used for LAN, SAN, and management traffic. This converged I/O results in reduced cables, SFPs and adapters – reducing capital and operational expenses of the overall solution.
- **Auto Discovery** – By simply inserting the blade server in the chassis or connecting the rack server to the fabric interconnect, discovery and inventory of compute resources occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.
- **Policy Based Resource Classification** – Once a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD showcases the policy-based resource classification of Cisco UCS Manager.
- **Combined Rack and Blade Server Management** – Cisco UCS Manager can manage Cisco UCS B-series blade servers and Cisco UCS C-series rack servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.
- **Model based Management Architecture** – The Cisco UCS Manager architecture and management database is model based, and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.
- **Policies, Pools, Templates** – The management approach in Cisco UCS Manager is based on defining policies, pools, and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network, and storage resources.
- **Loose Referential Integrity** – In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each other. This provides great flexibility where different experts from different domains, such as network, storage, security, server, and virtualization work together to accomplish a complex task.
- **Policy Resolution** – In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organizational hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then the special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.
- **Service Profiles and Stateless Computing** – A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.

- **Built-in Multi-Tenancy Support** – The combination of policies, pools and templates, loose referential integrity, policy resolution in the organizational hierarchy and a service profiles-based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environments typically observed in private and public clouds.
- **Extended Memory** – The enterprise-class Cisco UCS Blade server extends the capabilities of the Cisco Unified Computing System portfolio in a half-width blade form factor. It harnesses the power of the latest Intel® Xeon® Scalable Series processor family CPUs and Intel® Optane DC Persistent Memory (DCPMM) with up to 18 TB of RAM (using 256GB DDR4 DIMMs and 512GB DCPMM).
- **Simplified QoS** – Even though Fibre Channel and Ethernet are converged in the Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

Cisco UCS Manager

Cisco UCS Manager (UCSM) provides unified, integrated management for all software and hardware components in Cisco UCS. Using [Cisco Single Connect](#) technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive graphical user interface (GUI), a command-line interface (CLI), or a through a robust application programming interface (API).

Cisco UCS Manager is embedded into the Cisco UCS Fabric Interconnect and provides a unified management interface that integrates server, network, and storage. Cisco UCS Manager performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers a comprehensive set of XML API for third party integration, exposes thousands of integration points, and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

Cisco UCS™ Manager 4.2 provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) across multiple chassis and Cisco UCS servers. Cisco UCS Manager 4.2 is a unified software release for all supported Cisco UCS hardware platforms. Release 4.2 enables support for Cisco UCS 6454 Fabric Interconnects, Cisco UCS VIC 1400 series adapter cards on Cisco UCS M6 servers and third-Generation Intel® Xeon® Scalable processor refresh and Intel® Optane™ Data Center persistent memory modules on UCS Intel-based M6 servers.

For more information on Cisco UCS Manager Release 4.2, refer to:

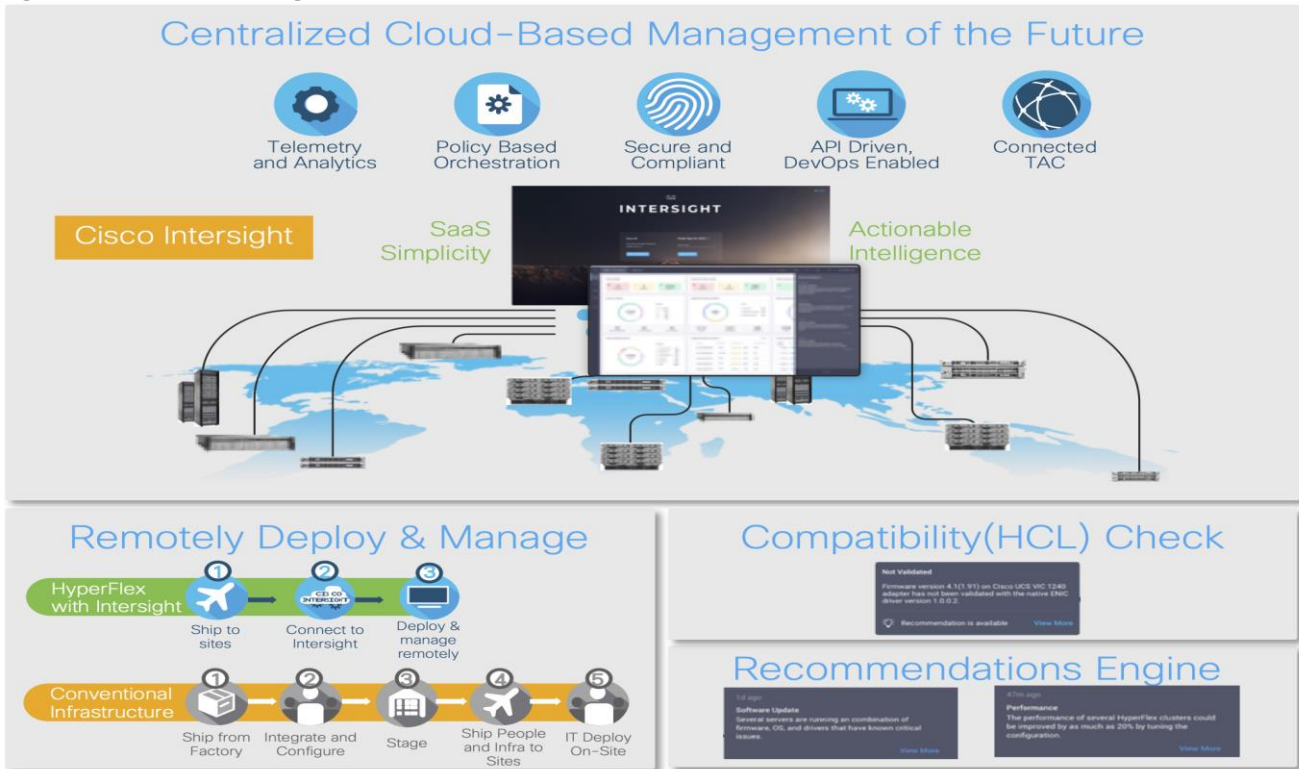
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/release/notes/cisco-ucs-manager-rn-4-2.html

Cisco Intersight

Cisco Intersight™ is a lifecycle management platform for your infrastructure, regardless of where it resides. In your enterprise data center, at the edge, in remote and branch offices, at retail and industrial sites—all these locations present unique management challenges and have typically required separate tools. Cisco Intersight Software as a Service (SaaS) unifies and simplifies your experience of the Cisco Unified Computing System™ (Cisco UCS®)

Intersight software delivers a new level of cloud-powered intelligence that supports lifecycle management with continuous improvement. It is tightly integrated with the Cisco® Technical Assistance Center (TAC). Expertise and information flow seamlessly between Cisco Intersight and IT teams, providing global management of Cisco infrastructure, anywhere. Remediation and problem resolution are supported with automated upload of error logs for rapid root-cause analysis.

Figure 4. Cisco Intersight



- Automate your infrastructure

Cisco has a strong track record for management solutions that deliver policy-based automation to daily operations. Intersight SaaS is a natural evolution of our strategies. Cisco designed Cisco UCS to be 100 percent programmable. Cisco Intersight simply moves the control plane from the network into the cloud. Now you can manage your Cisco UCS infrastructure wherever it resides through a single interface.

- Deploy your way

If you need to control how your management data is handled, comply with data locality regulations, or consolidate the number of outbound connections from servers, you can use the Cisco Intersight Virtual Appliance for an on-premises experience. Cisco Intersight Virtual Appliance is continuously updated just like the SaaS version, so regardless of which approach you implement, you never have to worry about whether your management software is up to date.

- DevOps ready

If you are implementing DevOps practices, you can use the Cisco Intersight API with either the cloud-based or virtual appliance offering. Through the API you can configure and manage infrastructure as code—you are not merely configuring an abstraction layer; you are managing the real thing. Through the API and support of cloud-based RESTful API, Terraform providers, Microsoft PowerShell scripts, or Python software, you can automate the deployment of settings and software for both physical and virtual layers. Using the API, you can simplify infrastructure lifecycle operations and increase the speed of continuous application delivery.

- Pervasive simplicity

- Simplify the user experience by managing your infrastructure regardless of where it is installed.
- Automate updates to Cisco UCS Data Platform software, reducing complexity and manual efforts.

-
- Actionable intelligence
 - Use best practices to enable faster, proactive IT operations.
 - Gain actionable insight for ongoing improvement and problem avoidance.
 - Manage anywhere
 - Deploy in the data center and at the edge with massive scale.
 - Get visibility into the health and inventory detail for your Intersight Managed environment on-the-go with the Cisco Intersight Mobile App.

For more information about Cisco Intersight and the different deployment options, go to: [Cisco Intersight – Manage your systems anywhere.](#)

Solution Components

This chapter contains the following:

- [Cisco UCS Fabric Interconnect](#)
- [Cisco UCS B200 M6 Blade Server](#)
- [Cisco Switches](#)
- [Cisco Intersight](#)
- [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](#)
- [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 UCS Fabric Interconnect

The Cisco UCS Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. Depending on the model chosen, the Cisco UCS Fabric Interconnect offers line-rate, low-latency, lossless 10 Gigabit, 25 Gigabit, 40 Gigabit, or 100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel connectivity. Cisco UCS Fabric Interconnects provide the management and communication backbone for the Cisco UCS C-Series, S-Series and HX-Series Rack-Mount Servers, Cisco UCS B-Series Blade Servers, and Cisco UCS 5100 Series Blade Server Chassis. All servers and chassis, and therefore all blades, attached to the Cisco UCS Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabrics, the Cisco UCS Fabric Interconnects provide both the LAN and SAN connectivity for all servers within its domain.

For networking performance, the Cisco UCS 6454 Series uses a cut-through architecture, supporting deterministic, low latency, line rate 10/25/40/100 Gigabit Ethernet ports, 3.82 Tbps of switching capacity, and 320 Gbps bandwidth per Cisco 5108 blade chassis when connected through the IOM 2408 model. The product family supports Cisco low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over the Ethernet fabric from the servers to the uplinks. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6454 Fabric Interconnect

The Cisco UCS 6454 Fabric Interconnect is a one-rack-unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has eight (8) 10/25-Gbps fixed Ethernet ports, which can optionally be configured as 8/16/32-Gbps FC ports (ports 1 to 8), thirty-six (36) 10/25-Gbps fixed Ethernet ports (ports 9 to 44), four (4) 1/10/25-Gbps Ethernet ports (ports 45 to 48), and finally six (6) 40/100-Gbps Ethernet uplink ports (ports 49 to 54). For more information, refer to the Cisco UCS 6454 Fabric Interconnect spec sheet: (<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/6400-specsheet.pdf>)

Figure 5. Cisco UCS 6454 Fabric Interconnect



Cisco UCS B200 M6 Blade Server

The Cisco UCS B200 M6 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads, including Virtual Desktop Infrastructure (VDI), web infrastructure, distributed databases, converged infrastructure, and enterprise applications such as Oracle and SAP HANA. The Cisco UCS B200 M6 blade server can quickly deploy stateless physical and virtual workloads through programmable, easy-to-use Cisco UCS Manager and Cisco Intersight™ and simplified server access through Cisco® SingleConnect technology. It includes:

- 3rd Gen Intel® Xeon® Scalable and processors with up to 40 cores per socket
- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane™ PMem
- Up to 2 Small Form-Factor (SFF) drives or up to 4 M.2 SATA drives
- Up to 80 Gbps of I/O throughput

Figure 6. Cisco UCS B200 M6 Blade Server



Cisco UCS VIC 1440 mLOM Interface Card

The Cisco UCS VIC 1440 mLOM Interface Card is a quad-port Enhanced Small Form-Factor Pluggable (SFP+) 10/25/40-Gbps Ethernet, and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS C-Series Rack Servers. The Cisco UCS VIC 1440 mLOM is used in conjunction with the Cisco UCS 6454 model Fabric Interconnects. The mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases. The card enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host, each dynamically configured as either a network interface card (NICs) or host bus adapter (HBA). The personality of the interfaces is set programmatically using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and

Worldwide Name [WWN]), failover policy, adapter settings, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all specified using the service profile.

Figure 7. Cisco UCS VIC 1440 mLOM Interface Card



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 data centers.


- 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 data centers 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 8. 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 9](#)) 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 Data Center Network Manager.

Figure 9. Cisco MDS 9132T 32-Gb 32-Port Fabric Channel Switch



Figure 10. 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 troubleshot using Cisco Data Center Network Manager (DCNM), which currently manages the entire suite of Cisco data center 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 DCM-SAN
 - Cisco DCM-SAN can be used to monitor, configure, and analyze Cisco 32Gbps Fibre Channel fabrics and show information about the Cisco Nexus switching fabric. Cisco DCM-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 DCM 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 DCM 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 DCM platform. The device connector provides a secure way for connected Cisco DCM to send and receive information from the Cisco Intersight portal, using a secure Internet connection.

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 11. Cisco Intersight Overview



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

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

Cisco Intersight Virtual Appliance and Private Virtual Appliance

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

Cisco Intersight Assist

Cisco Intersight Assist helps customers add endpoint devices to Cisco Intersight. A data center could have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight, but does not connect 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, go to: https://intersight.com/help/getting_started#licensing_requirements

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 data center, hybrid cloud platforms, and services to securely deploy and manage infrastructure resources across data center 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 12. Example of User-Customizable Cisco Intersight Dashboard for FlexPod UCS Domain

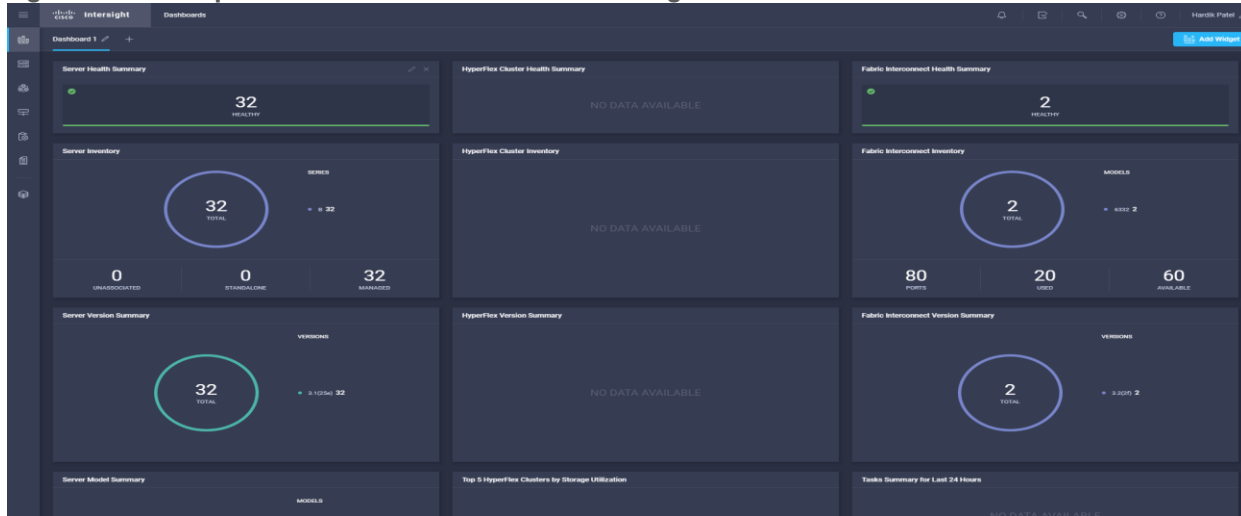
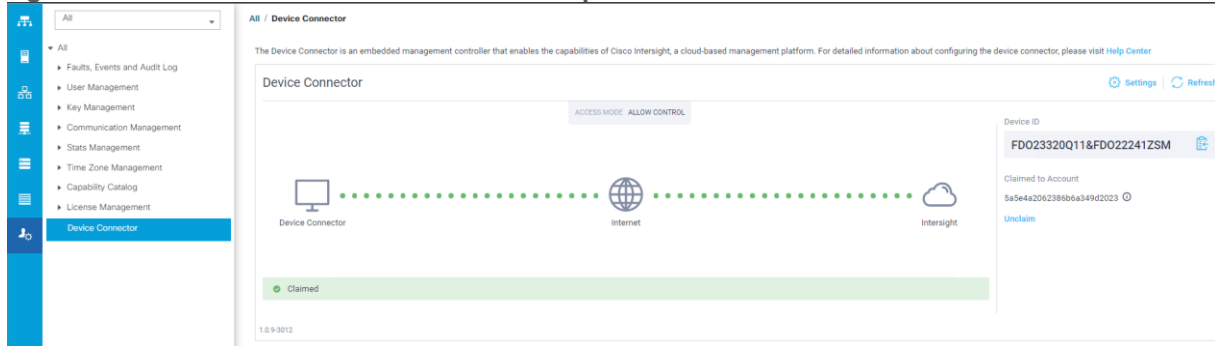


Figure 13. Cisco UCSM Device Connector Example



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: <https://docs.vmware.com/en/VMware-Horizon/2111/published-desktops-applications/GUID-F411EE75-58EB-4DF6-B216-F1889A526FDF.html>

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 data center, 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?platformTypeId=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
FCoE target ports, UTA2	n/a	n/a	n/a	64

Specifications	AFF A250	AFF A400	AFF A800	AFF A900
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 data center 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 15. NetApp AFF A400 Front View



Figure 16. 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.

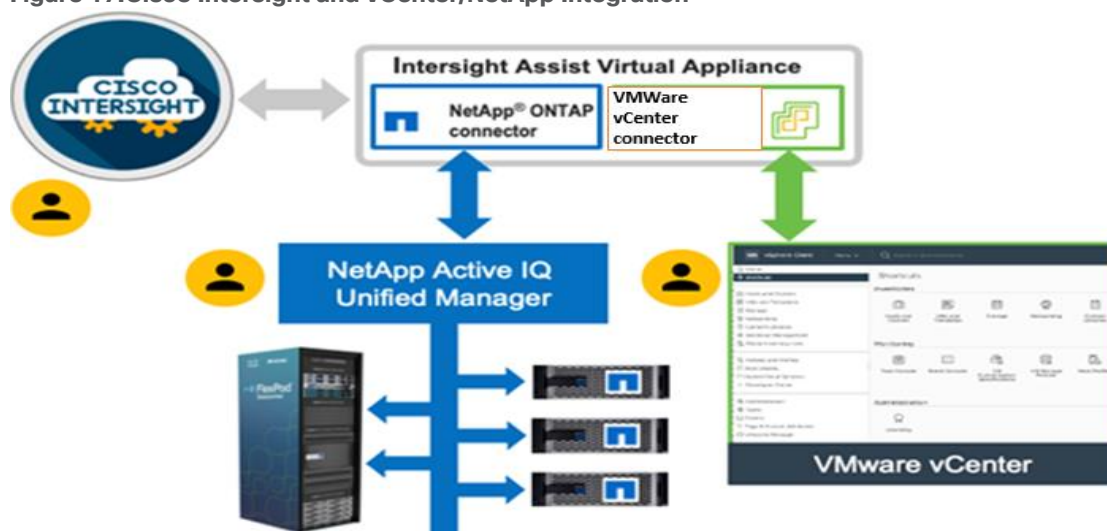
NetApp ONTAP 9

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 data center. ONTAP 9 has a host of features to simplify deployment and data management, accelerate and protect critical data, and make infrastructure future-ready across hybrid-cloud architectures.

NetApp ONTAP 9 is the data management software that is used with the NetApp AFF A400 all-flash storage system in this solution design. ONTAP software offers secure unified storage for applications that read and write data over block- or file-access protocol storage configurations. These storage configurations range from high-speed flash to lower-priced spinning media or cloud-based object storage. ONTAP implementations can run on NetApp engineered 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 17. Cisco Intersight and vCenter/NetApp Integration



:

Table 2. AFF Technical Specifications

NetApp ONTAP 9.10.1P1

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 14](#).

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 14. Storage Efficiency Features

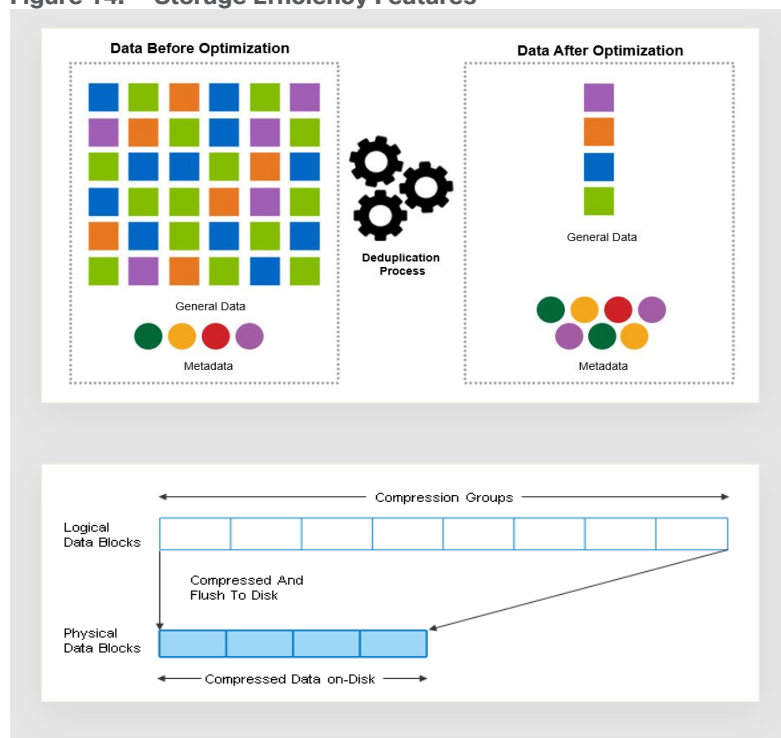
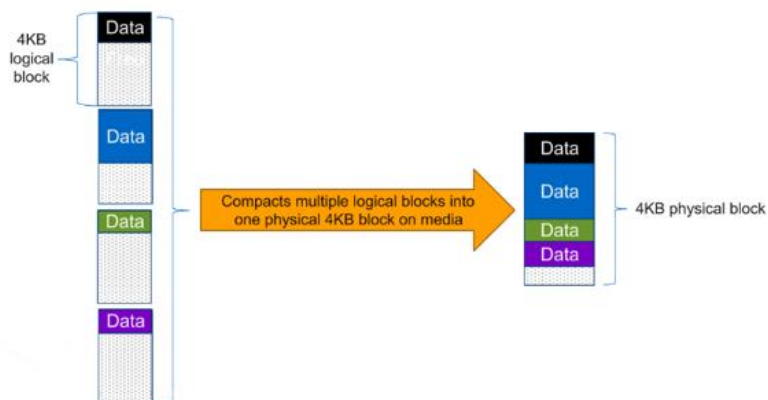


Figure 15. Storage Efficiency



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 3](#) 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 3. Compression and Deduplication Guidelines

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

Workload	Storage Efficiency Guidelines		
		deduplication	<ul style="list-style-type: none"> Inline zero-block deduplication
File Services	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Adaptive background compression Deduplication Inline zero-block deduplication
Mixed Workload	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary workloads, use: <ul style="list-style-type: none"> Deduplication Inline zero-block deduplication For secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Adaptive background compression Deduplication Inline zero-block deduplication

Space Savings

[Table 4](#) 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 4. 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 FlexClone volumes (where applicable) technologies]	Ratio Range
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 linked 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 16. 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 17. 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 18. 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 19. 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 20](#) 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 20. 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

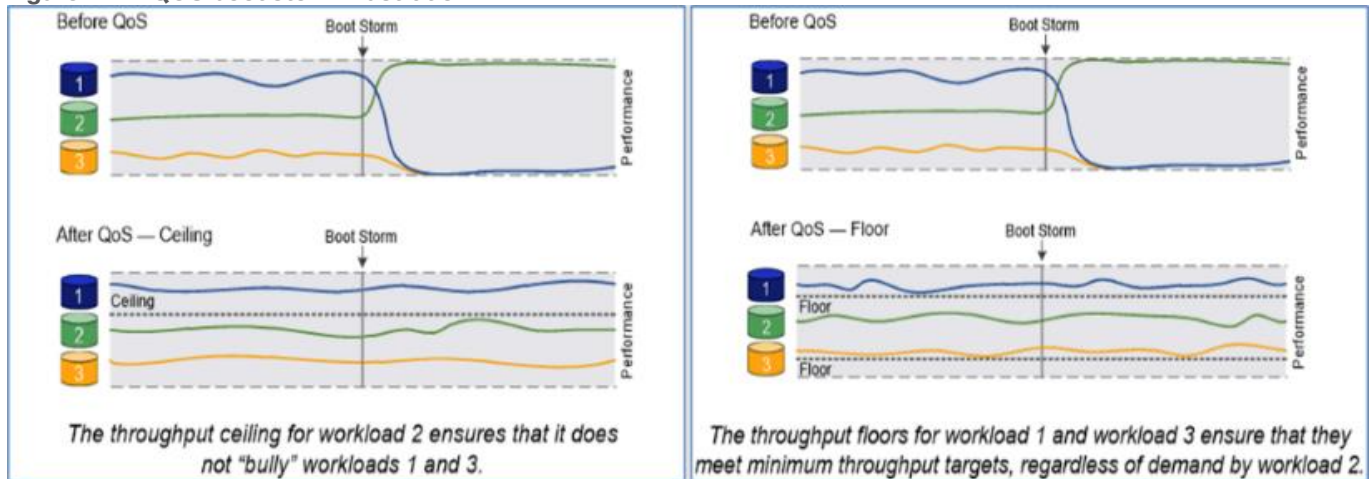
maintaining the ratio of IOPS to TBs/GBs. To assist in managing your QoS, Active IQ unified manager will provide QoS suggestions based on historical performance and usage.

Three default adaptive QoS policy groups are available, as shown in [Table 5](#). You can apply these policy groups directly to a volume.

Table 5. 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 21. 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.

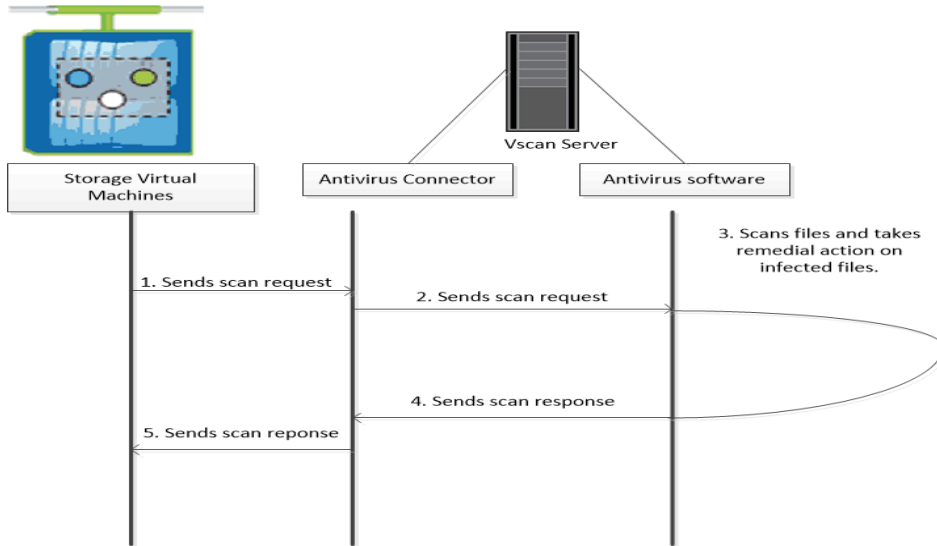
Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. The ONTAP Antivirus Connector, provided by NetApp and installed on the external server, handles communication between the storage system and the antivirus software.

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.

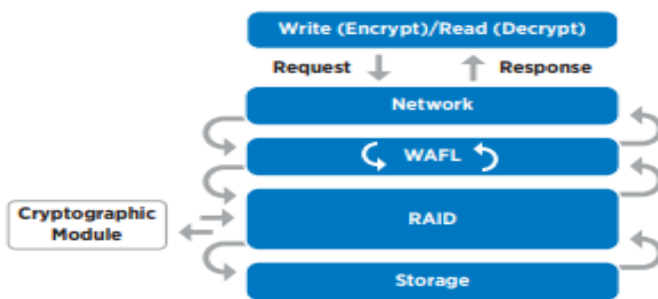


NetApp Volume Encryption(NVE) and NetApp Aggregate Encryption (NAE)

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 22](#).

Figure 22. 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 data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

VMware vSphere 7.0 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

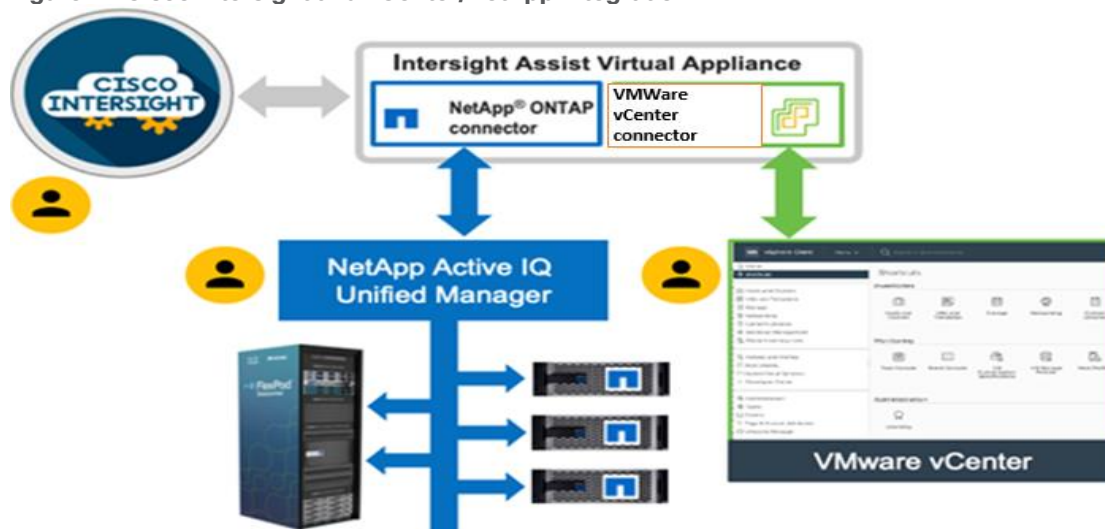
VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest 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 17. Cisco Intersight and vCenter/NetApp Integration



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

Enterprise SAN and NAS workloads can benefit equally from the integrated management solution. The integration architecture enables FlexPod customers to use new management capabilities with no compromise in their existing 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.

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.

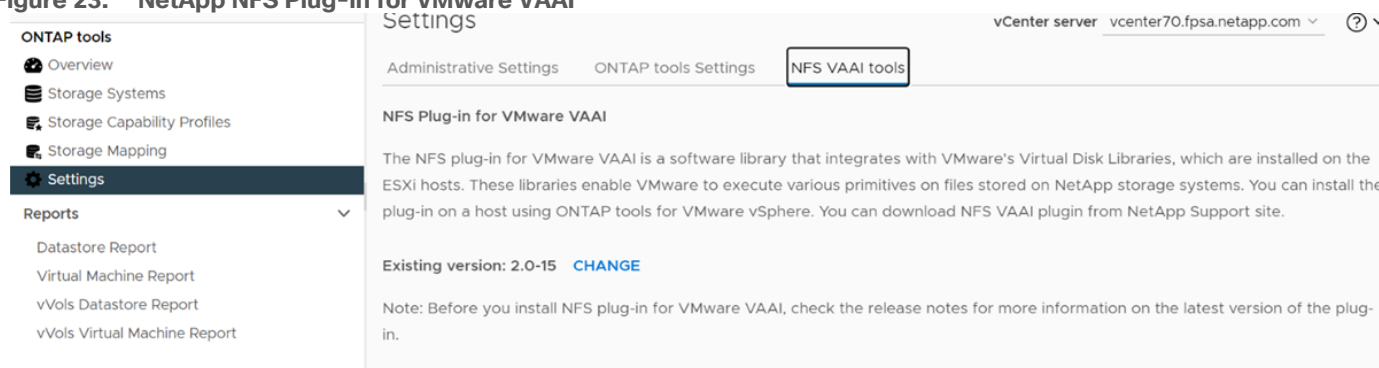
To download ontap tools for vmware vsphere, go to:

<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 23. NetApp NFS Plug-in for VMware VAAI



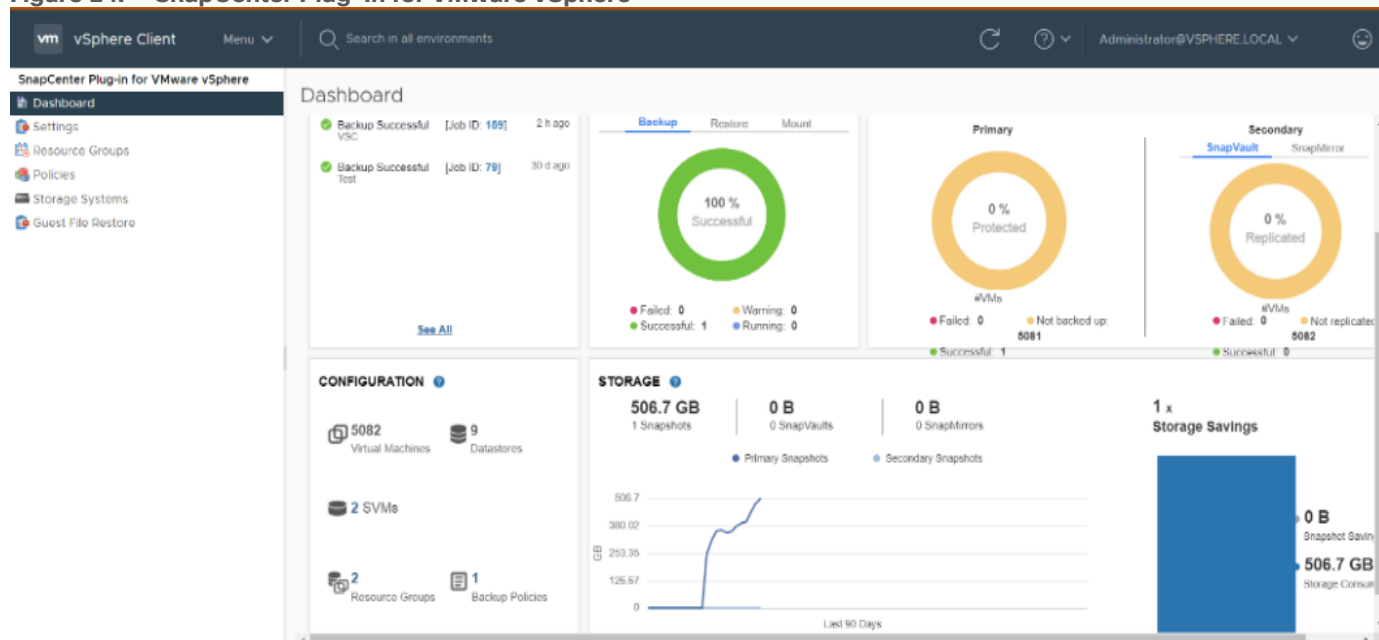
The copy offload feature and space reservation feature improve the performance of VSC operations. The NetApp NFS Plug-in for VAAI is not shipped with VSC, but you can install it by using VSC. You can download the plug-in installation package and obtain the instructions for installing the plug-in from 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 24. 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

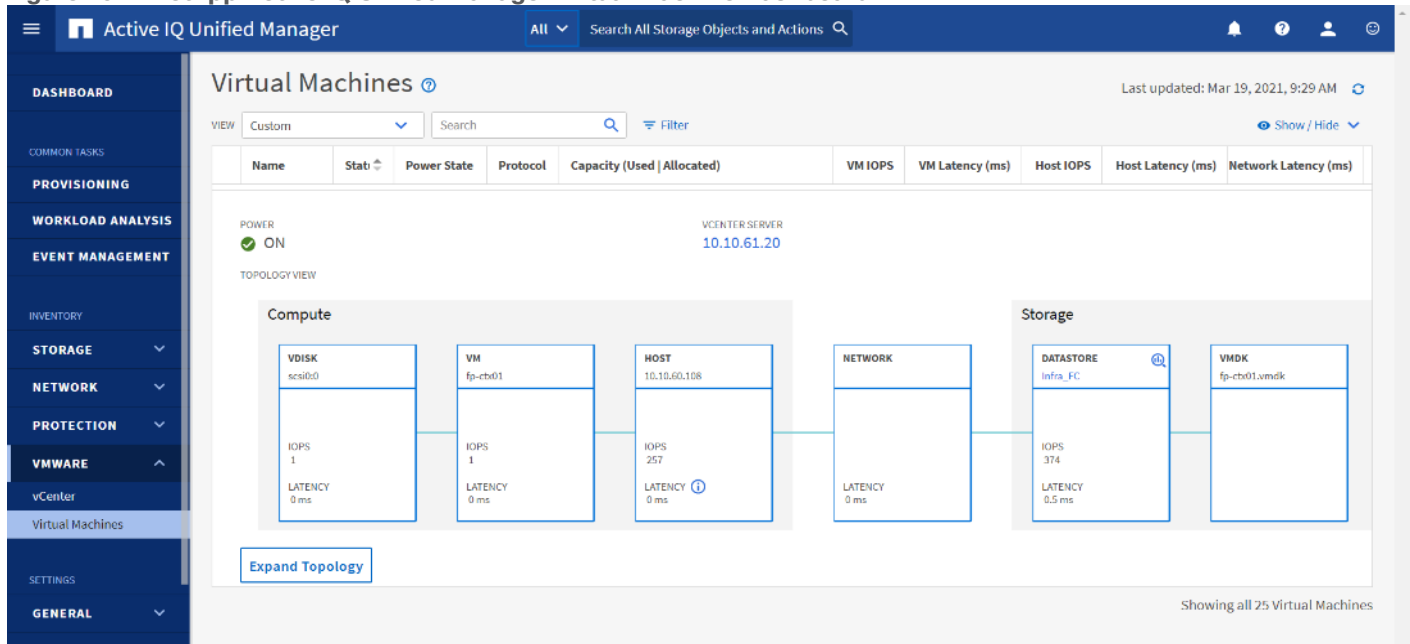
NetApp Active IQ Unified Manager (Unified Manager) is a comprehensive monitoring and proactive management tool for NetApp ONTAP systems to help manage the availability, capacity, protection, and performance risks of your storage systems and virtual infrastructure. You can deploy Unified Manager on a 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 25. 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, 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 data center 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 2.

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 hosted shared sessions) or a Windows 10 Desktop OS (for, instant clone and persistent VDI 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 26. VMware Horizon Design Overview

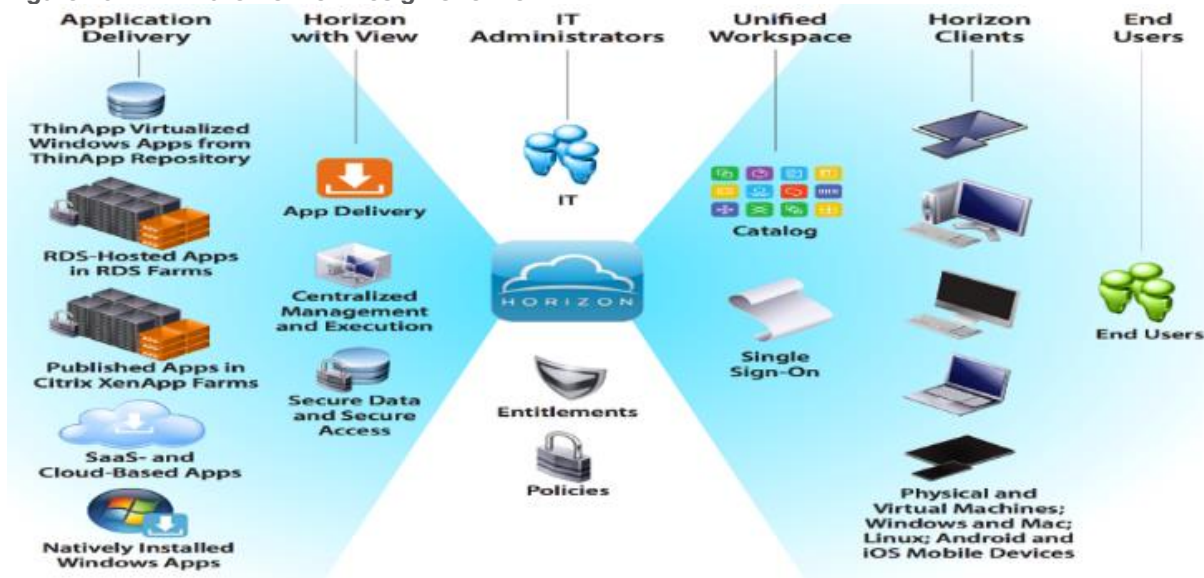
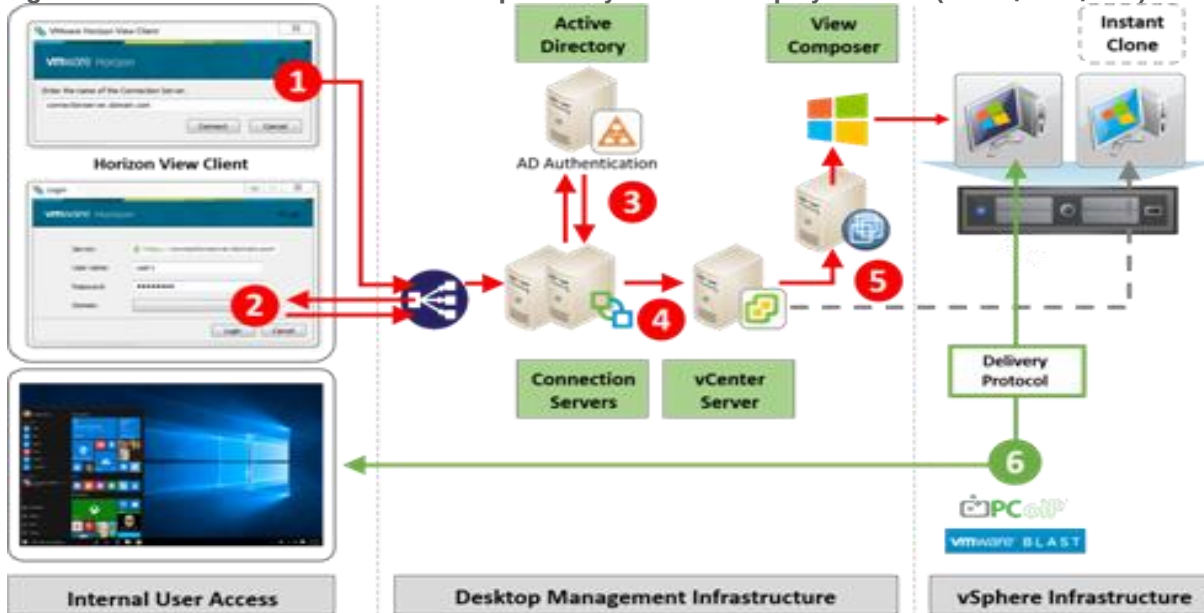


Figure 27. Horizon VDI and RDSH Desktop Delivery Based on Display Protocol (PCoIP/Blast/RDP)

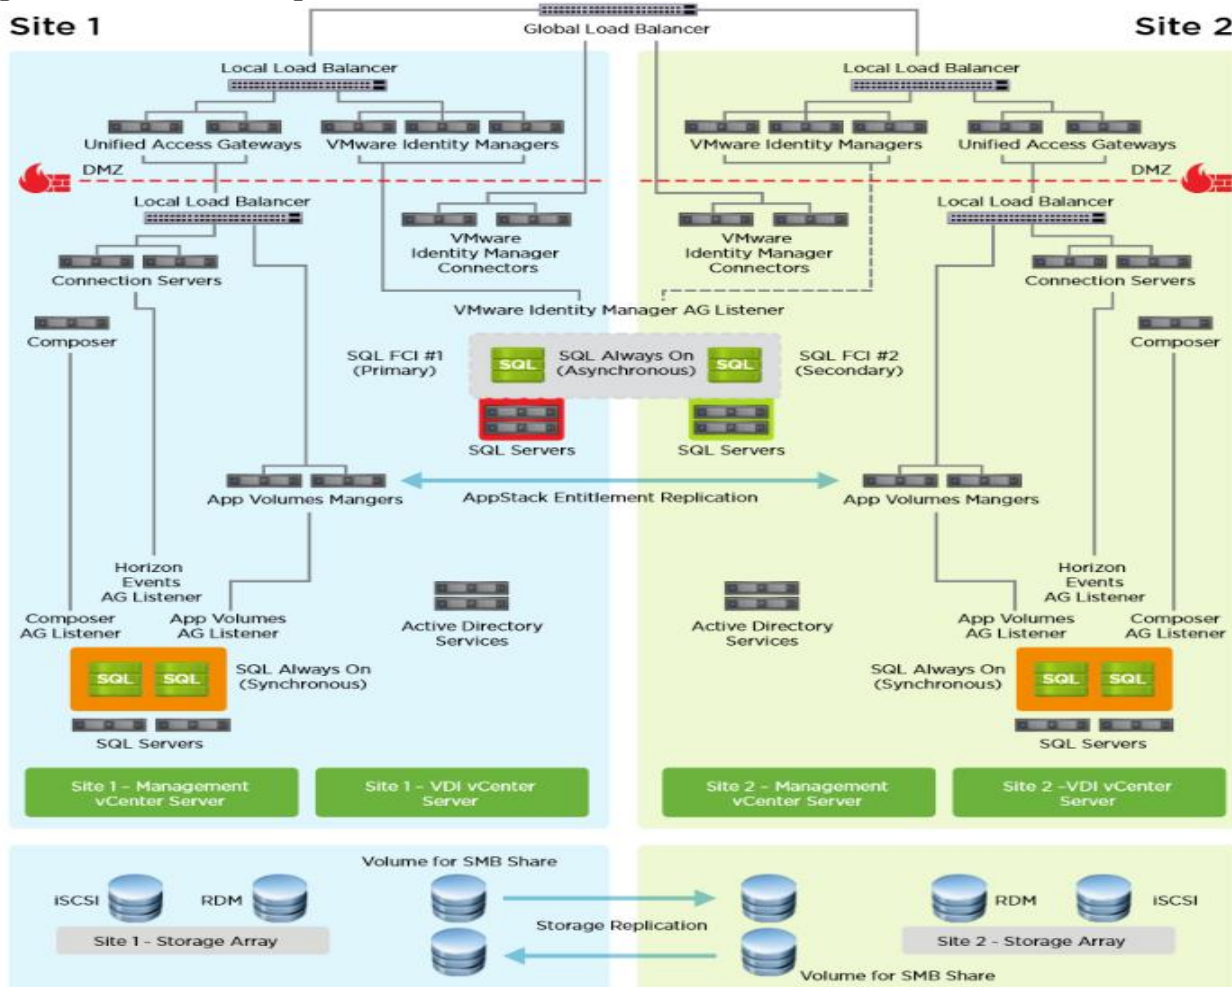


Multiple-Site Configuration

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 28](#) illustrating sites, shows a site created in two data centers. 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 28. Multisite Configuration Overview



Based on the requirement and no of data centers 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: 2300 Windows Server 2019 random pooled desktops were configured and tested
- Single-session OS Solution:
 - VMware Horizon Instant Clones non-persistent virtual machines: 1700 Windows 10 Virtual desktops random pooled were configured and tested
 - VMware Horizon Full Clone persistent virtual machines: 1700 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, Linked- clone automated pool 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://www.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 data center rack, including the access layer network switches.

Products Deployed

This CVD details the deployment of up to 2300 Multi-session OS, 1700 Single-session OS VDI users featuring the following software:

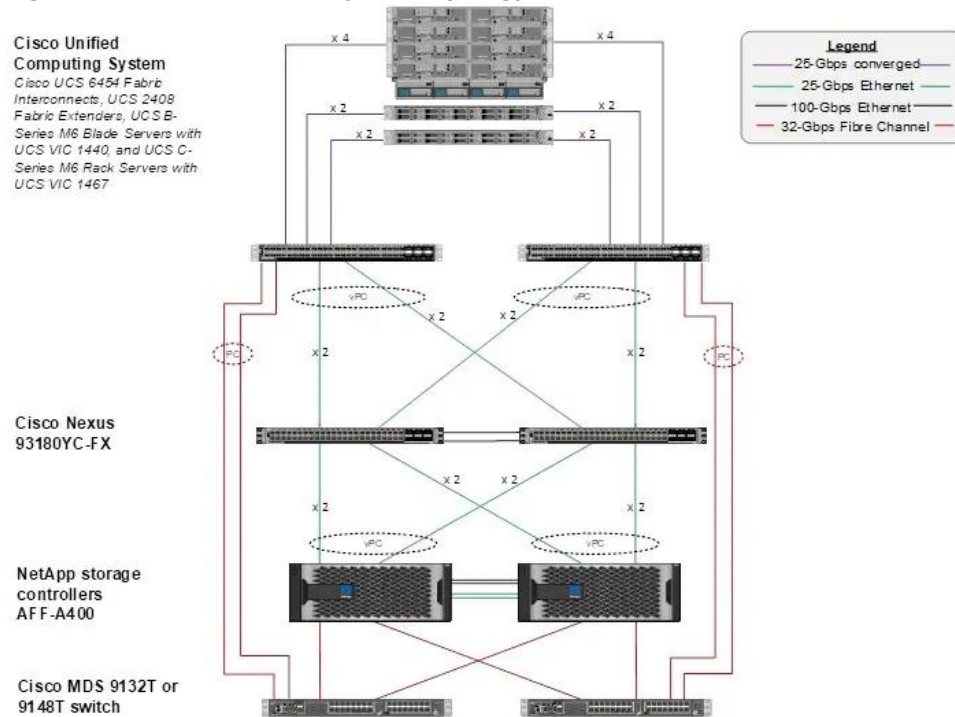
- VMware vSphere ESXi 7 Update 2 Hypervisor
- Microsoft SQL Server 2019 SP1
- Microsoft Windows Server 2019 and Windows 10 64-bit virtual machine Operating Systems
- VMware Horizon 2111 Remote Desktop Server Hosted (RDSH) Sessions provisioned as Linked Clones and stored on the NFS storage
- VMware Horizon 2111 Non-Persistent Win10 Virtual Desktops (VDI) provisioned as Instant Clones and stored on NFS storage
- VMware Horizon 2111 Persistent Win10 Virtual Desktops (VDI) provisioned as Full Clones and stored on NFS storage
- NetApp ONTAP Tools for VMware vSphere 9.10P1
- 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 U2 delivers a Virtual Desktop Infrastructure that is redundant, using the best practices of Cisco and NetApp Storage. The solution includes VMware vSphere 7.0 U2 hypervisor installed on the Cisco UCS M6 blade 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 29. FlexPod VDI - Physical Topology



[Figure 29](#) 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.
- 1 Cisco UCS 5108 Blade Server Chassis with two Cisco UCS-IOM-2408 IO Modules.
- 8 Cisco UCS B200 M6 Blade Servers with Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM, and one Cisco VIC1440 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 6](#) lists the software versions of the primary products installed in the environment.

Table 6. Software and Firmware Versions

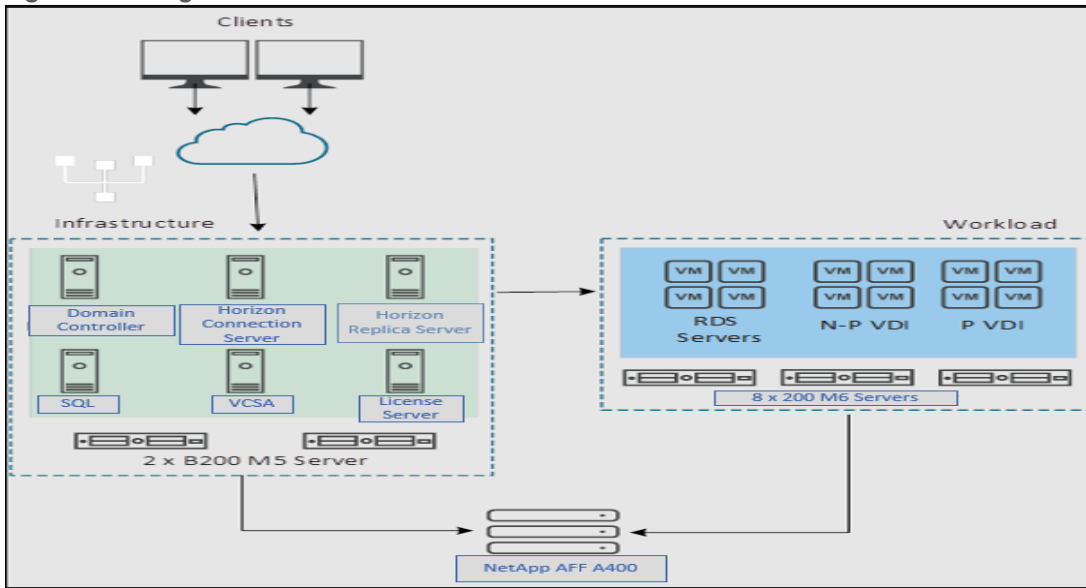
Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	4.2(1f) bundle release

Vendor	Product / Component	Version / Build / Code
Cisco	UCS Manager	4.2(1f) bundle release
Cisco	UCS B200 M6 Blades	4.2(1f) bundle release
Cisco	VIC 1440	4.2(1f) 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	7.0.2.18455184
VMware	vSphere 7. U2	7.0.2.00500
VMware	Horizon Connection Server	8.4.0.18964782
VMware	Horizon Agent	8.4.0.18964730
VMware	Tools	11.2.5.17337674

Logical Architecture

The logical architecture of the validated solution which is designed to support up to 1700 users on a single chassis containing Eight Cisco UCS B200 M6 blade servers, with physical redundancy for the blade servers for each workload type is illustrated in [Figure 30](#).

Figure 30. Logical Architecture Overview



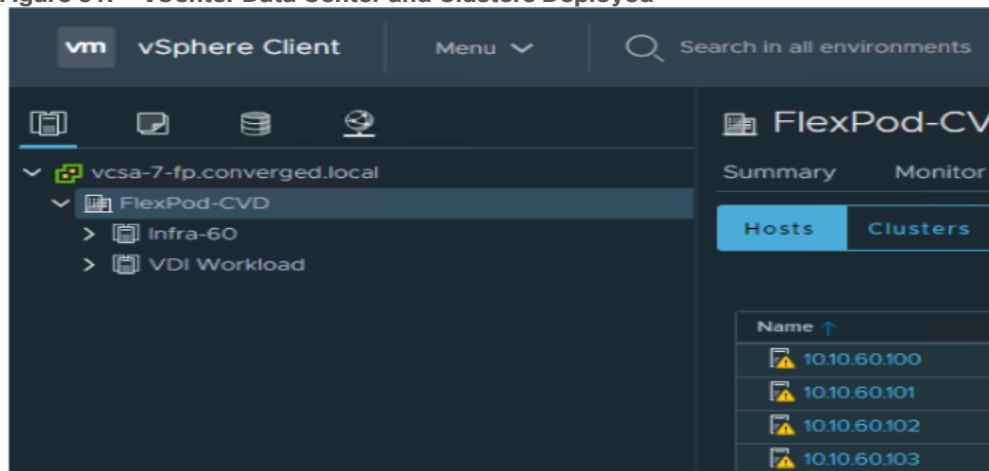
VMware Clusters

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

- VDI Cluster FlexPod Data Center 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 streamed with VMware Horizon, Windows 10 Streamed 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, LVS Web Servers and LVS 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 31. vCenter Data Center 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 2111 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 7](#).

Table 7. VLANS Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
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	132	Out-of-Band management interfaces

VSANS

Two virtual SANs configured for communications and fault tolerance in this design as outlined in [Table 8](#).

Table 8. 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](#)
- [Physical Connectivity](#)
- [FlexPod Cisco Nexus Base](#)
- [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 32](#) 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 32. FlexPod Solution Cabling Diagram

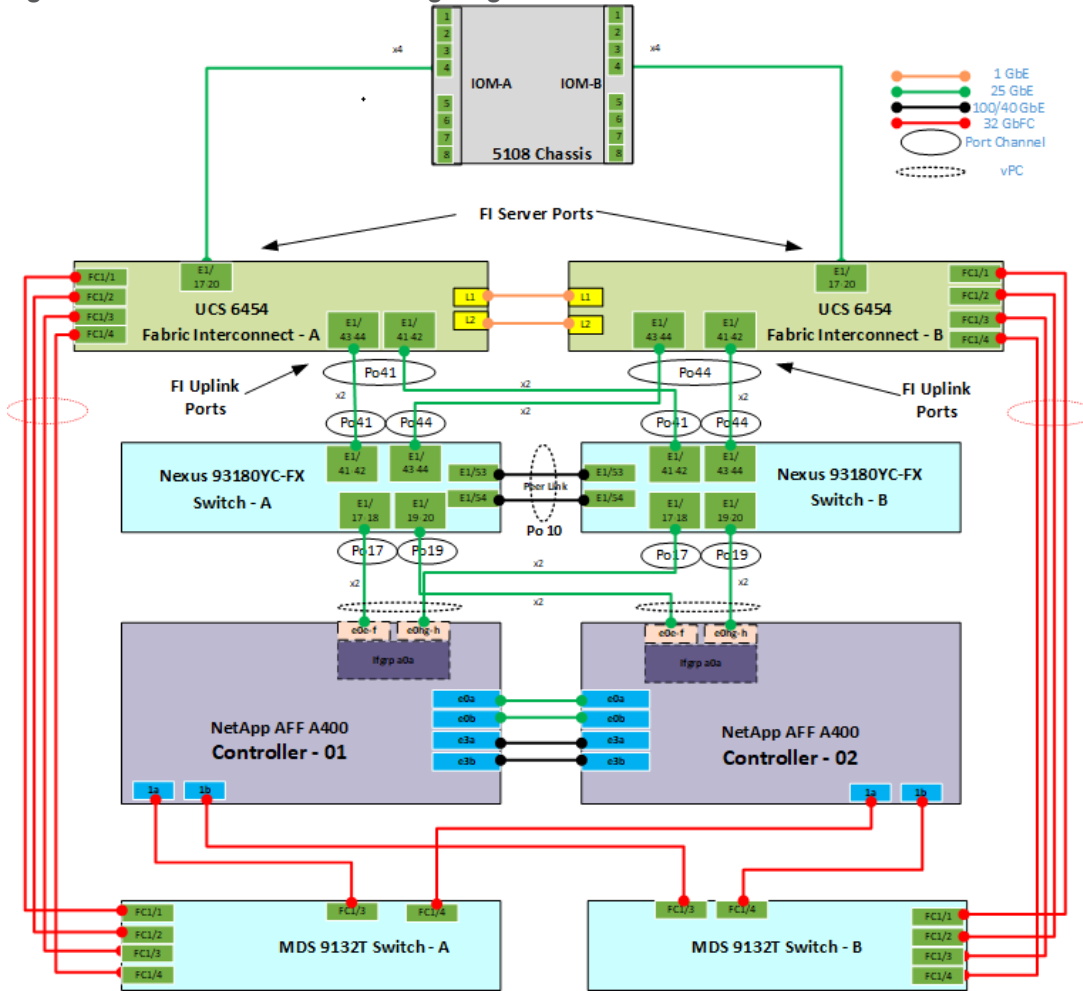


Table 9. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX A	Eth1/19	25GbE	NetApp Controller 2	e0e
	Eth1/20	25GbE	NetApp Controller 2	e0f
	Eth1/17	25GbE	NetApp Controller 1	e0e
	Eth1/18	25GbE	NetApp Controller 1	e0f
	Eth1/41	25GbE	Cisco UCS fabric interconnect A	PO 1/43
	Eth1/42	25GbE	Cisco UCS fabric interconnect A	PO 1/44
	Eth1/43	25GbE	Cisco UCS fabric interconnect B	PO 1/43
	Eth1/44	25GbE	Cisco UCS fabric interconnect B	PO 1/44
	Eth1/53	40GbE	Cisco Nexus 93180YC-FX B	Eth1/53
	Eth1/54	40GbE	Cisco Nexus 93180YC-FX B	Eth1/54

Local Device	Local Port	Connection	Remote Device	Remote Port
	MGMT0	GbE	GbE management switch	Any

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

Table 10. Cisco Nexus 93180YC-FX-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX B	Eth1/17	25GbE	NetApp Controller 1	e0g
	Eth1/18	25GbE	NetApp Controller 1	e0h
	Eth1/19	25GbE	NetApp Controller 2	e0g
	Eth1/20	25GbE	NetApp Controller 2	e0h
	Eth1/41	25GbE	Cisco UCS fabric interconnect A	P0 1/41
	Eth1/42	25GbE	Cisco UCS fabric interconnect A	P0 1/42
	Eth1/43	25GbE	Cisco UCS fabric interconnect B	P0 1/43
	Eth1/44	25GbE	Cisco UCS fabric interconnect B	P0 1/44
	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 11. 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 12. 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 13. Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect	Eth1/41	25GbE	Cisco Nexus 93180YC-FX B	Eth1/41
	Eth1/42	25GbE	Cisco Nexus 93180YC-FX B	Eth1/42

Local Device	Local Port	Connection	Remote Device	Remote Port
A	Eth1/43	25GbE	Cisco Nexus 93180YC-FX A	Eth1/41
	Eth1/44	25GbE	Cisco Nexus 93180YC-FX A	Eth1/42
	FC1/1	32GbE	Cisco MDS 9132T A	IOM 1/1
	FC1/2	32GbE	Cisco MDS 9132T A	IOM 1/2
	FC1/3	32GbE	Cisco MDS 9132T A	IOM 1/3
	FC1/4	32GbE	Cisco MDS 9132T A	IOM 1/4
	MGMT0	GbE	GbE Management switch	
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 14. Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth1/41	25GbE	Cisco Nexus 93180YC-FX B	Eth1/43
	Eth1/42	25GbE	Cisco Nexus 93180YC-FX B	Eth1/44
	Eth1/43	25GbE	Cisco Nexus 93180YC-FX A	Eth1/43
	Eth1/44	25GbE	Cisco Nexus 93180YC-FX A	Eth1/44
	FC1/1	32GbE	Cisco MDS 9132 T B	FC 1/1
	FC1/2	32GbE	Cisco MDS 9132 T B	FC 1/2
	FC1/3	32GbE	Cisco MDS 9132 T B	FC 1/3
	FC1/4	32GbE	Cisco MDS 9132 T B	FC 1/4
	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 Eth1/41
description <ucs-clustername>-a:1/43
udld enable
interface Eth1/42
description <ucs-clustername>-a:1/44
udld enable
```

```
interface Eth1/43
description <ucs-clustername>-b:1/43
udld enable
interface Eth1/44
description <ucs-clustername>-b:1/44
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 Eth1/41
description <ucs-clustername>-a:1/41
udld enable
interface Eth1/42
description <ucs-clustername>-a:1/42
udld enable
interface Eth1/43
description <ucs-clustername>-b:1/41
udld enable
interface Eth1/44
description <ucs-clustername>-b:1/42
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 Eth1/53-54
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 Eth1/41-42
channel-group 121 mode active
no shutdown
interface Po142
description <ucs-clustername>-b
interface Eth1/43-44
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 Po141
inherit port-profile FP-UCS
interface Po142
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 Po141
vpc 121
interface Po142
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 Po141
vpc 121
interface Po142
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](#)
- [Cisco UCS Manual Deployment](#)
- [FlexPod Cisco MDS Switch Configuration](#)

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 \(Infra-SVM\) 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 15](#) lists the information needed to configure two ONTAP nodes. Customize the cluster-detail values with the information applicable to your deployment.

Table 15. 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 16. Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>

Cluster Detail	Cluster Detail Value
Cluster Admin SVM	<cluster-adm-svm>
Infrastructure Data SVM	<infra-data-svm>
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.

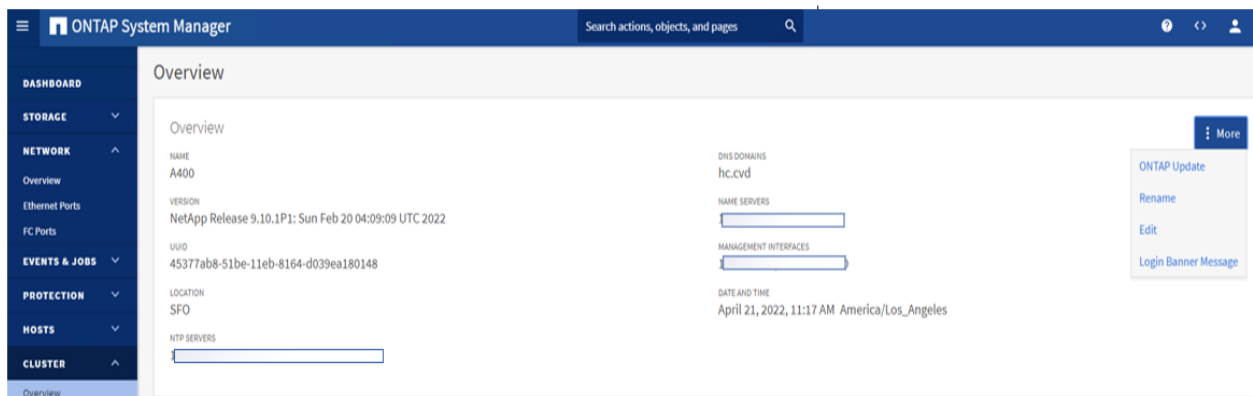
The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

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.

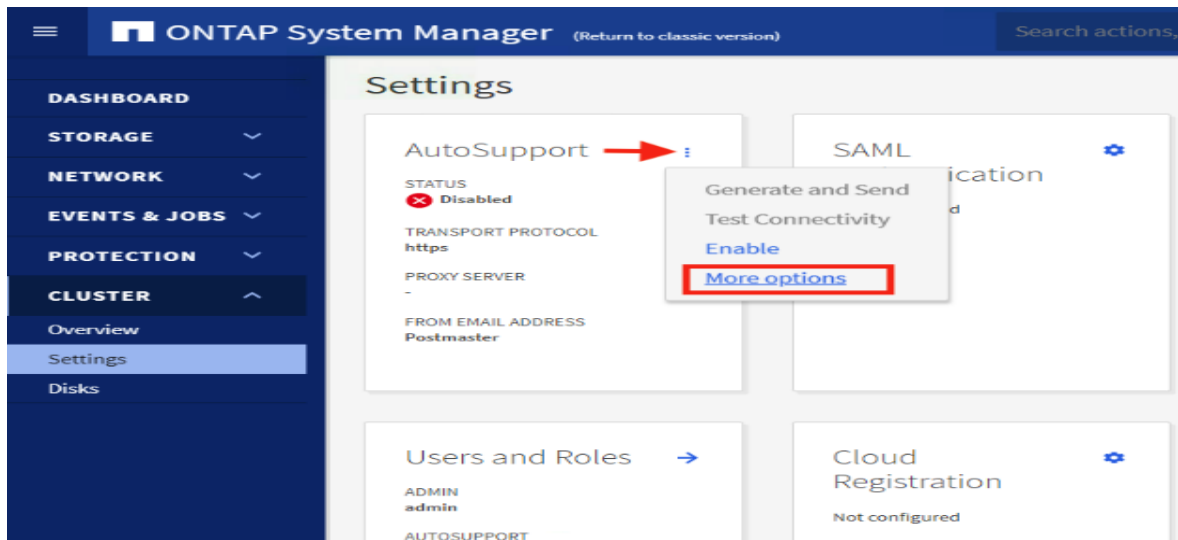


Step 19. Add additional cluster configuration details and click Save to make the changes persistent:

- 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.

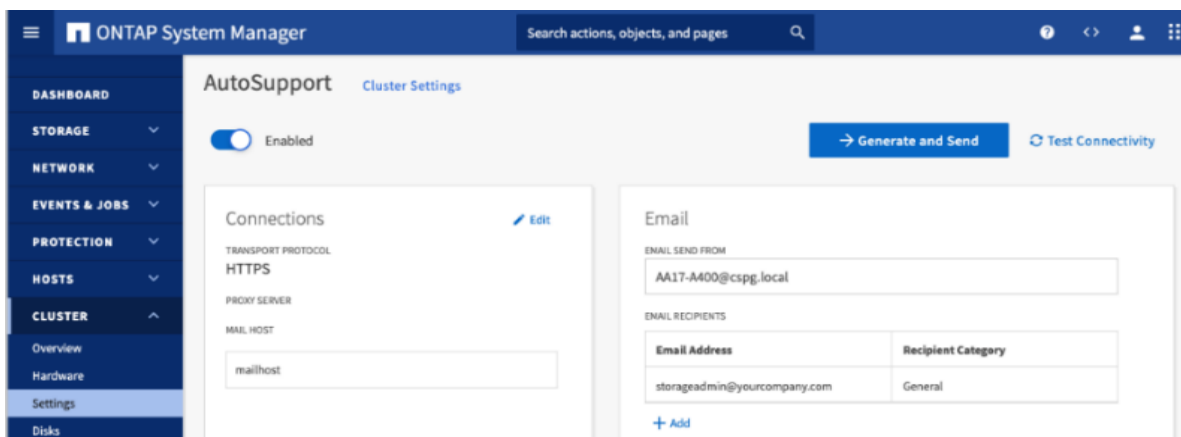
Step 24. Click Edit to change the transport protocol, add a proxy server address and a mail host as needed.

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

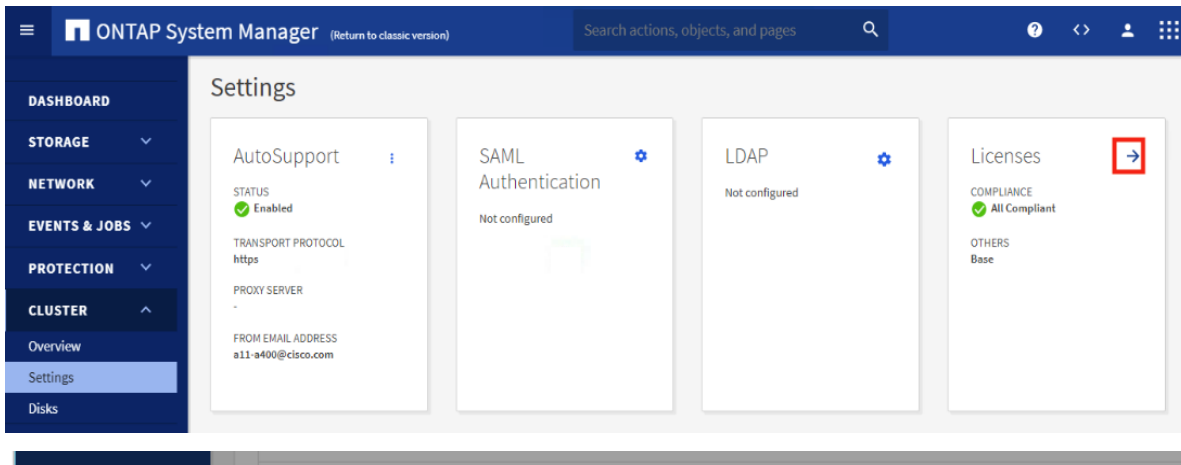
Step 27. Click Save when complete.



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

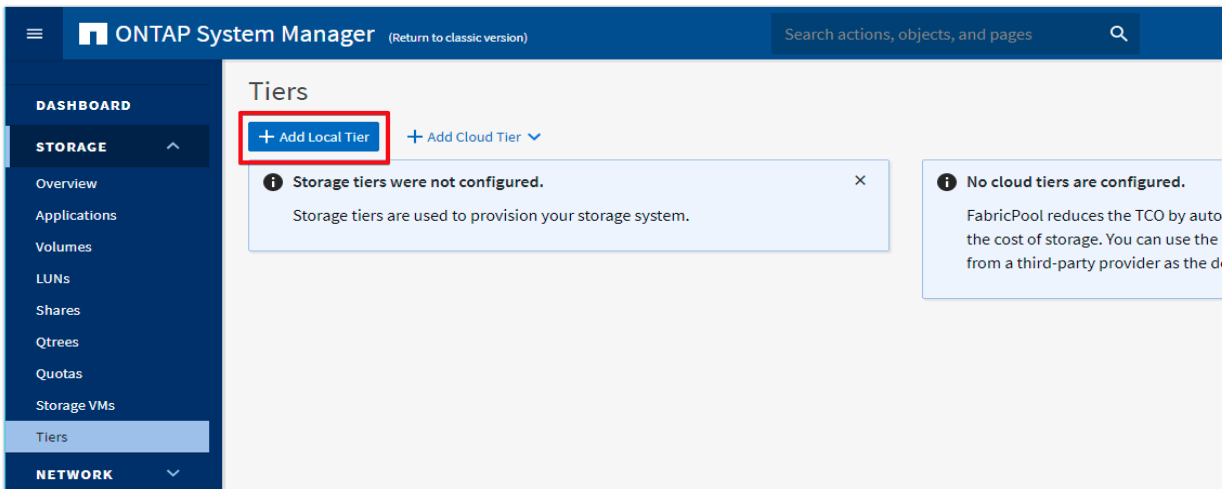
Step 29. Locate the Licenses tile on the right and click the detail arrow.

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

LOCAL TIER 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

..... X

.....

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:

```
AA17-A400::> storage failover show
                                Takeover
Node           Partner           Possible State Description
-----
AA17-A400-01  AA17-A400-02     true      Connected to AA17-A400-02
AA17-A400-02  AA17-A400-01     true      Connected to AA17-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.

```
AA17-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:

```
AA17-A400::> storage failover hwassist show
Node
-----
AA17-A400-01
                Partner: AA17-A400-02
                Hwassist Enabled: true
                Hwassist IP: 192.x.x.84
                Hwassist Port: 162
                Monitor Status: active
                Inactive Reason: -
                Corrective Action: -
                Keep-Alive Status: healthy

AA17-A400-02
                Partner: AA17-A400-01
                Hwassist Enabled: true
                Hwassist IP: 192.x.x.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> -ip-space 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
```

```
AA17-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
AA17-A400-01	e0e	auto	auto	none
AA17-A400-01	e0f	auto	auto	none
AA17-A400-01	e0g	auto	auto	none
AA17-A400-01	e0h	auto	auto	none
AA17-A400-02	e0e	auto	auto	none
AA17-A400-02	e0f	auto	auto	none
AA17-A400-02	e0g	auto	auto	none
AA17-A400-02	e0h	auto	auto	none

```
8 entries were displayed.
```

```
AA17-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
AA17-A400-01	e3a	auto	auto	none
AA17-A400-01	e3b	auto	auto	none
AA17-A400-02	e3a	auto	auto	none
AA17-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:

```
AA17-A400::> network port ifgrp show

      Port      Distribution      Active
Node   IfGrp      Function      MAC Address      Ports  Ports
-----
AA17-A400-01
      a0a      port      d2:39:ea:29:d4:4a  full  e0e, e0f, e0g, e0h
AA17-A400-02
      a0a      port      d2:39:ea:29:ce:d5  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:

```
AA17-A400::> network port vlan show

      Network Network
Node  VLAN Name Port  VLAN ID  MAC Address
```

```

-----
AA17-A400-01
  a0a-60    a0a    60    d2:39:ea:29:d4:4a
  a0a-61
            a0a    61    d2:39:ea:29:d4:4a
  a0a-62
            a0a    62    d2:39:ea:29:d4:4a
  a0a-63
            a0a    63    d2:39:ea:29:d4:4a
AA17-A400-02
  a0a-60    a0a    60    d2:39:ea:29:ce:d5
  a0a-61
            a0a    61    d2:39:ea:29:ce:d5
  a0a-62
            a0a    62    d2:39:ea:29:ce:d5
  a0a-63
            a0a    63    d2:39:ea:29:ce:d5

```

8 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-SVM -rootvolume infra_svm_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-SVM aggregate list for the NetApp ONTAP Tools:

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

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

```
vserver nfs create -vserver Infra-SVM -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:

```
AA17-A400::> vserver nfs show -fields vstorage
vserver    vstorage
-----
Infra-SVM enabled
```

Procedure 26. Configure CIFS Servers

Note: You can enable and configure CIFS servers on storage virtual machines (SVMs) with NetApp FlexVol® 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-SVM -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-SVM -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-SVM -cifs-server Infra-CIFS -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-SVM -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-SVM -volume infra_svm_root_m01 -aggregate <aggr1_node01> -
size 1GB -type DP

volume create -vserver Infra-SVM -volume infra_svm_root_m02 -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-SVM:infra_svm_root -destination-path Infra-
SVM:infra_svm_root_m01 -type LS -schedule 15min
```

```
snapmirror create -source-path Infra-SVM:infra_svm_root -destination-path Infra-
SVM:infra_svm_root_m02 -type LS -schedule 15min
```

Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path Infra-SVM:infra_svm_root
```

Step 4. To verify, run the following:

```
AA17-A400::> snapmirror show -type ls
```

Source Path	Destination Type	Mirror Path	Relationship State	Total Progress	Healthy	Progress Last Updated
AA17-A400://Infra-SVM/Infra_SVM_root	LS	AA17-A400://Infra-SVM/infra_svm_root_m01	Snapmirrored Idle	-	true	-
		AA17-A400://Infra-SVM/infra_svm_root_m02	Snapmirrored Idle	-	true	-

2 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-SVM -status-admin up
```

Step 2. To verify, run the following:

```
AA17-A400::> vserver fcp show
```

Vserver	Target Name	Admin	Status
-----	-----	-----	-----
Infra-SVM	20:00:d0:39:ea:29:ce:d4	up	

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

Procedure 30. Create iSCSI Block Protocol Service - required only for iSCSI configuration

Step 1. Run the following command to create the iSCSI service:

```
vserver iscsi create -vserver <infra-data-svm>
```

Step 2. To verify, run the following:

```
AA17-A400::> vserver iscsi show
```

Vserver	Target Name	Target Alias	Admin	Status
-----	-----	-----	-----	-----
Infra-SVM	iqn.1992-08.com.netapp:sn.63144a05ad1211eb8a7ad039ea29d44a:vs.3	Infra-SVM	up	

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

Procedure 31. Vserver Protocol Verification

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

```
AA17-A400::> vserver show-protocols -vserver Infra-SVM
```

Vserver: Infra-SVM

Protocols: nfs, fcp, iscsi, ndmp, nvme

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

```
vserver add-protocols -vserver <infra-data-svm> -protocols < iscsi or fcp >
```

Procedure 32. 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 -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -
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-SVM 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 -vserver Infra-SVM
```

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 33. 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-SVM -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-SVM -volume infra_svm_root -policy default
```

Procedure 34. 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-SVM -policyname cifs
export-policy rule create -vserver Infra-SVM -policyname cifs -clientmatch <domain_name>
-rorule
```

```
krb5i,krb5p -rwrule krb5i,krb5p
```

Procedure 35. 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-SVM -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-SVM -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-SVM -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-SVM -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-SVM:infra_svm_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 36. 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-SVM -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-SVM -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-SVM -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 37. 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-SVM -volume infra_swap
```

Procedure 38. 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-SVM -share-name <CIFS_share_1> -path
/infra_datastore_01 -share properties oplocks,browsable,continuously-
available,showsnapshot

cifs share create -vserver Infra-SVM -share-name <CIFS_share_2> -path
/infra_datastore_02 -share properties oplocks,browsable,continuously-
available,showsnapshot

cifs share create -vserver Infra-SVM -share-name <CIFS_share_3> -path /cifs_vol_03 -
share properties oplocks,browsable,continuously-available,showsnapshot
```

Procedure 39. Create NFS LIFs

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

```
network interface create -vserver Infra-SVM -lif nfs-lif-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-SVM -lif nfs-lif-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:

```
AA17-A400::> network interface show -vserver Infra-SVM -data-protocol nfs

          Logical      Status      Network          Current      Current Is
Vserver   Interface  Admin/Oper Address/Mask     Node         Port      Home
-----
Infra-SVM
          nfs-lif-01   up/up      192.168.30.1/24  AA17-A400-01 a0a-62
                                               true
          nfs-lif-02   up/up      192.168.30.2/24  AA17-A400-02 a0a-62
                                               true
```

2 entries were displayed.

Procedure 40. Create CIFS LIFs

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

```
network interface create -vserver Infra-SVM -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-SVM -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 41. 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-SVM -lif fcp-lif-01a -role data -data-protocol
fcp -home-node <st-node01> -home-port 1a -status-admin up
network interface create -vserver Infra-SVM -lif fcp-lif-01b -role data -data-protocol
fcp -home-node <st-node01> -home-port 1b -status-admin up
network interface create -vserver Infra-SVM -lif fcp-lif-02a -role data -data-protocol
fcp -home-node <st-node02> -home-port 1a -status-admin up
network interface create -vserver Infra-SVM -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:

```
AA17-A400::> network interface show -vserver Infra-SVM -data-protocol fcp
      Logical      Status      Network      Current      Current Is
Vserver  Interface  Admin/Oper  Address/Mask  Node          Port      Home
-----  -
Infra-SVM
      fcp-lif-01a  up/up      20:01:d0:39:ea:29:ce:d4
                                                AA17-A400-01  1a        true
      fcp-lif-01b  up/up      20:02:d0:39:ea:29:ce:d4
                                                AA17-A400-01  1b        true
      fcp-lif-02a  up/up      20:03:d0:39:ea:29:ce:d4
                                                AA17-A400-02  1a        true
      fcp-lif-02b  up/up      20:04:d0:39:ea:29:ce:d4
                                                AA17-A400-02  1b        true
```

4 entries were displayed.

Procedure 42. Create iSCSI LIFs - required only for iSCSI configuration

Step 1. To create four iSCSI LIFs, run the following commands (two on each node):

```
network interface create -vserver <infra-data-svm> -lif iscsi-lif-01a -role data -data-
protocol iscsi -home-node <st-node01> -home-port a0a-<infra-iscsi-a-vlan-id> -address
<st-node01-infra-iscsi-a-ip> -netmask <infra-iscsi-a-mask> -status-admin up
network interface create -vserver <infra-data-svm> -lif iscsi-lif-01b -role data -data-
protocol iscsi -home-node <st-node01> -home-port a0a-<infra-iscsi-b-vlan-id> -address
<st-node01-infra-iscsi-b-ip> -netmask <infra-iscsi-b-mask> -status-admin up
network interface create -vserver <infra-data-svm> -lif iscsi-lif-02a -role data -data-
protocol iscsi -home-node <st-node02> -home-port a0a-<infra-iscsi-a-vlan-id> -address
<st-node02-infra-iscsi-a-ip> -netmask <infra-iscsi-a-mask> -status-admin up
network interface create -vserver <infra-data-svm> -lif iscsi-lif-02b -role data -data-
protocol iscsi -home-node <st-node02> -home-port a0a-<infra-iscsi-b-vlan-id> -address
<st-node02-infra-iscsi-b-ip> -netmask <infra-iscsi-b-mask> -status-admin up
```

Procedure 43. Configure FC-NVMe Datastore for vSphere 7U2 on existing SVM - Infra-SVM - for FC-NVMe configuration only

Note: To Configure FC-NVMe Datastores for vSphere 7U2, enable the FC-NVMe protocol on an existing SVM or create a separate SVM for FC-NVMe workloads. In this deployment, Infra-SVM was used for FC-NVMe datastore configuration.

Step 1. Verify NVMe Capable adapters are installed in the cluster:

```
network fcp adapter show -data-protocols-supported fc-nvme
```

Step 2. Add the NVMe protocol to the SVM and list it:

```
vserver add-protocols -vserver Infra-SVM -protocols nvme
```

Step 3. To verify, run the following:

```
AA17-A400::> vserver show -vserver Infra-SVM -fields allowed-protocols
vserver  allowed-protocols
-----
Infra-SVM nfs, fcp, iscsi, ndmp, nvme
```

Step 4. Create NVMe service:

```
vserver nvme create -vserver Infra-SVM
```

Step 5. To verify, run the following:

```
AA17-A400::> vserver nvme show -vserver Infra-SVM
Vserver Name: Infra-SVM
Administrative Status: up
```

Step 6. Create NVMe FC LIFs:

```
network interface create -vserver Infra-SVM -lif fc-nvme-lif-01a -role data -data-protocol fc-nvme -home-node <st-node01> -home-port 1a -status-admin up
network interface create -vserver Infra-SVM -lif fc-nvme-lif-01b -role data -data-protocol fc-nvme -home-node <st-node01> -home-port 1b -status-admin up
network interface create -vserver Infra-SVM -lif fcp-nvme-lif-02a -role data -data-protocol fc-nvme -home-node <st-node02> -home-port 1a -status-admin up
network interface create -vserver Infra-SVM -lif fcp-nvme-lif-02b -role data -data-protocol fc-nvme -home-node <st-node02> -home-port 1b -status-admin up
```

Step 7. To verify, run the following:

```
AA17-A400::> network interface show -vserver Infra-SVM -data-protocol fc-nvme
Logical      Status      Network      Current      Current Is
Vserver      Interface   Admin/Oper   Address/Mask Node          Port         Home
-----
Infra-SVM
fc-nvme-lif-01a
up/up        20:06:d0:39:ea:29:ce:d4
AA17-A400-01 1b           true
fc-nvme-lif-01b
up/up        20:08:d0:39:ea:29:ce:d4
AA17-A400-01 1a           true
fc-nvme-lif-02a
up/up        20:07:d0:39:ea:29:ce:d4
AA17-A400-02 1b           true
```

```

fc-nvme-lif-02b
                up/up      20:09:d0:39:ea:29:ce:d4
                                AA17-A400-02  1a      true

```

Note: You can only configure two NVMe LIFs per node on a maximum of four nodes.

Step 8. Create volume:

```

vol create -vserver Infra-SVM -volume NVMe_Datastore_01 -aggregate
AA17_A400_01_NVME_SSD_1 -size 500G -state online -space-guarantee none -percent-
snapshot-space 0

```

Procedure 44. Add Infrastructure SVM Administrator and SVM Administration LIF to In-band Management Network

Step 1. Run the following commands:

```

network interface create -vserver Infra-SVM -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-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-
gateway>

```

Step 3. To verify, run the following:

```

AA17-A400::> network route show -vserver Infra-SVM
Vserver          Destination      Gateway          Metric
-----
Infra-SVM
                0.0.0.0/0       192.168.17.254  20

```

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

```

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

```

```

security login unlock -username vsadmin -vserver Infra-SVM

```

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 45. 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>
```

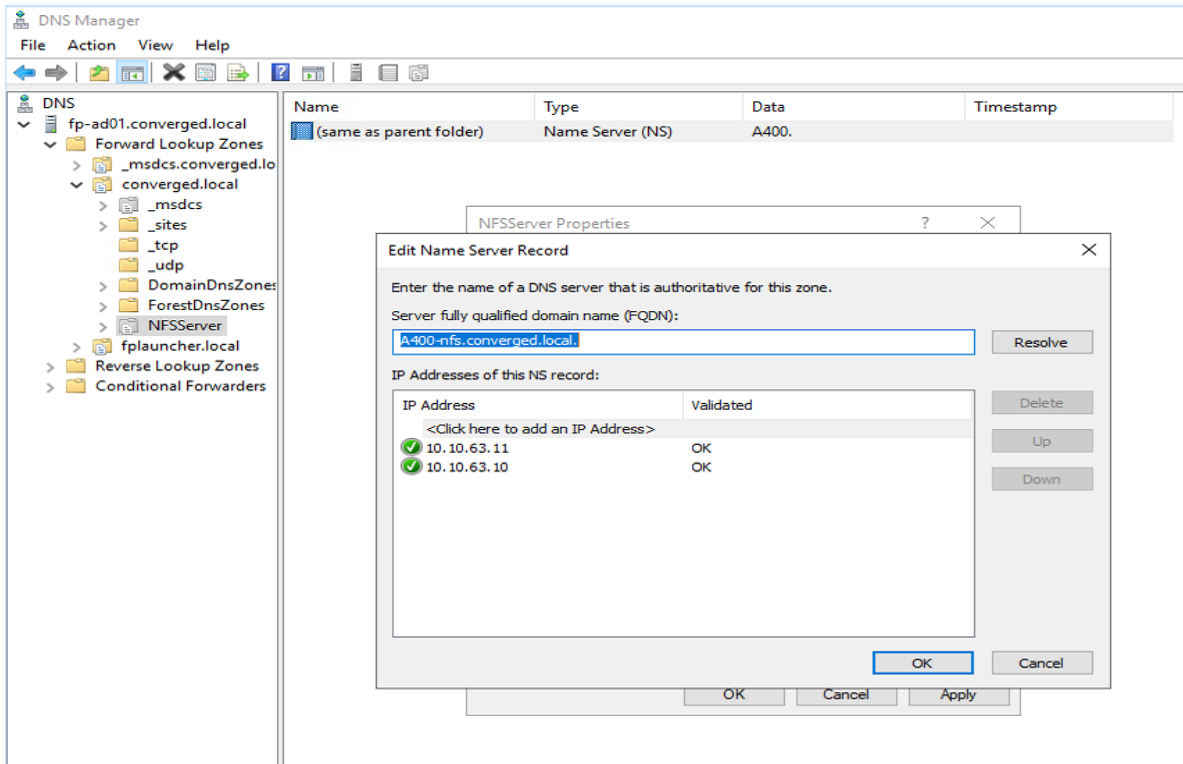
Name	Number of Members	Size	Adaptive QoS Policy	Expected IOPS (2048 * Allocated Space)	Peak IOPS (4096 * Used Space)
VDI	8	30TB (12% used)	performance	61440	14745.6
Data	8	10TB (25% used)	performance	20480	10240

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 UCS Manual Deployment

This subject contains the following procedures:

- [Configure Fabric Interconnects at Console](#)
- [Configure Fabric Interconnects for a Cluster Setup](#)
- [Configure Base Cisco Unified Computing System](#)
- [Synchronize Cisco UCSM to NTP](#)
- [Configure Global Policies](#)
- [Set Fabric Interconnects to Fibre Channel End Host Mode](#)
- [Configure FC SAN Uplink Ports](#)
- [Configure Server Ports](#)
- [Configure Ethernet LAN Uplink Ports](#)
- [Create Uplink Port Channels to Cisco Nexus Switches](#)
- [Configure VLAN](#)
- [Configure VSAN](#)
- [Create New Sub-Organization](#)
- [IP Pool Creation](#)
- [UUID Suffix Pool Creation](#)
- [Server Pool Creation](#)
- [MAC Pool Creation](#)

- [WWNN and WWPN Pool Creation](#)
- [Set Jumbo Frames in both Cisco Fabric Interconnects](#)
- [Create Host Firmware Package](#)
- [Create Server Pool Policy](#)
- [Create Server Pool Policy Qualifications](#)
- [Create a Target Pool and Qualification](#)
- [Create Network Control Policy for Cisco Discovery Protocol](#)
- [Create Power Control Policy](#)
- [Create Server BIOS Policy](#)
- [Configure Maintenance Policy](#)
- [Create vNIC Templates](#)
- [Create vHBA Templates](#)
- [Create Server Boot Policy for SAN Boot](#)
- [Create SAN Policy A](#)
- [Create SAN Policy B](#)
- [Create and Configure a Service Profile Template](#)
- [Clone Service Profile Template](#)
- [Create Service Profiles from Template and Associate to Servers](#)
- [Create First Four Service Profiles from Template](#)

Cisco Unified Computing System Base Configuration

This section details the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power, and installation of the chassis are described in the [Cisco UCS Manager Getting Started Guide](#) and it is beyond the scope of this document. For more information about each step, refer to the following document, [Cisco UCS Manager - Configuration Guides](#).

Cisco UCS Manager Software Version 4.2(1f)

This document assumes you are using Cisco UCS Manager Software version 4.2(1f). To upgrade the Cisco UCS Manager software and the Cisco UCS 6454 Fabric Interconnect software to a higher version of the firmware,) refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Procedure 1. Configure Fabric Interconnects at Console

Step 1. Connect a console cable to the console port on what will become the primary fabric interconnect.

Step 2. If the fabric interconnect was previously deployed and you want to erase it to redeploy, login with the existing user name and password:

```
# connect local-mgmt
# erase config
# yes (to confirm)
```

Step 3. After the fabric interconnect restarts, the out-of-the-box first time installation prompt appears, type “console” and press Enter.

Step 4. Follow the [Initial Configuration](#) steps as outlined in [Cisco UCS Manager Getting Started Guide](#). When configured, log into UCSM IP Address through Web interface to perform base Cisco UCS configuration.

Procedure 2. Configure Fabric Interconnects for a Cluster Setup

Step 1. Verify the following physical connections on the fabric interconnect:

- a. The management Ethernet port (mgmt0) is connected to an external hub, switch, or router.
- b. The L1 ports on both fabric interconnects are directly connected to each other
- c. The L2 ports on both fabric interconnects are directly connected to each other
- d. Connect to the console port on the first Fabric Interconnect.

Step 2. Review the settings on the console. Answer yes to Apply and Save the configuration.

Step 3. Wait for the login prompt to make sure the configuration has been saved to Fabric Interconnect A.

Step 4. Connect the console port on the second Fabric Interconnect, configure secondary FI.

Figure 33. Initial Setup of Cisco UCS Manager on Primary Fabric Interconnect

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

Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
Enforce strong password? (y/n) [y]: n

Enter the password for "admin":
Confirm the password for "admin":

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes
Enter the switch fabric (A/B) []: A
Enter the system name: VCC-AAD17
Physical Switch Mgmt0 IP address : 10.29.164.246
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 10.29.164.1
Cluster IPv4 address : 10.29.164.245
Configure the DNS Server IP address? (yes/no) [n]:
Configure the default domain name? (yes/no) [n]:
Join centralized management environment (UCS Central)? (yes/no) [n]:
Following configurations will be applied:
Switch Fabric=A
System Name=VCC-AAD17
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=10.29.164.246
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.29.164.1
Ipv6 value=0

Cluster Enabled=yes
Cluster IP Address=10.29.164.245
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized.
UCSM will be functional only after peer FI is configured in clustering mode.

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

Configuration file - OK

Cisco UCS 6300 Series Fabric Interconnect
VCC-AAD17-A login: █
```

Figure 34. Initial Setup of Cisco UCS Manager on Secondary Fabric Interconnect

```
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:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.164.246
Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
Cluster IPv4 address      : 10.29.164.245

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

Physical Switch Mgmt0 IP address : 10.29.164.247

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

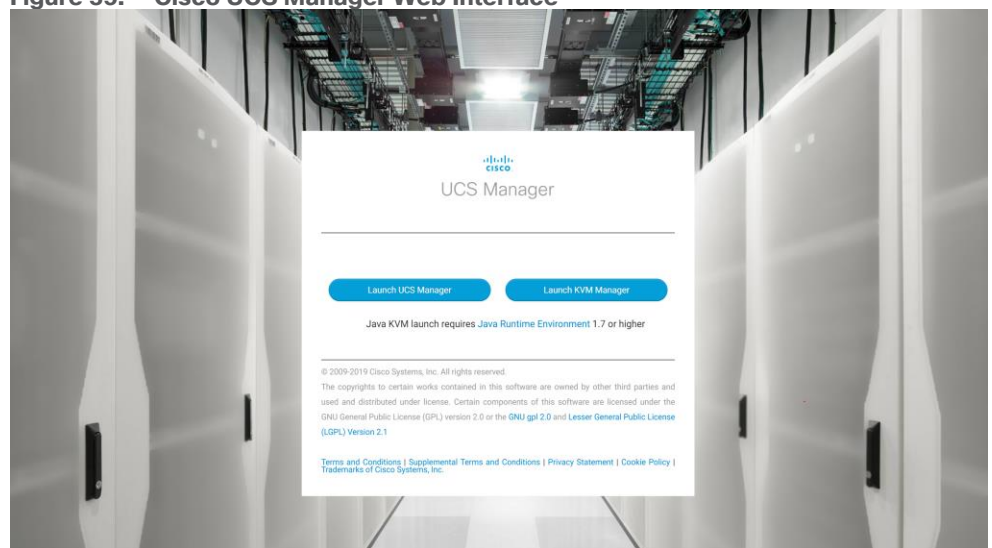
Fri Feb 16 18:53:15 UTC 2018
Configuration file - Ok

Cisco UCS 6300 Series Fabric Interconnect
VCC-AAD17-B login: █
```

Step 5. Log into the Cisco Unified Computing System (Cisco UCS) environment; open a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address configured above.

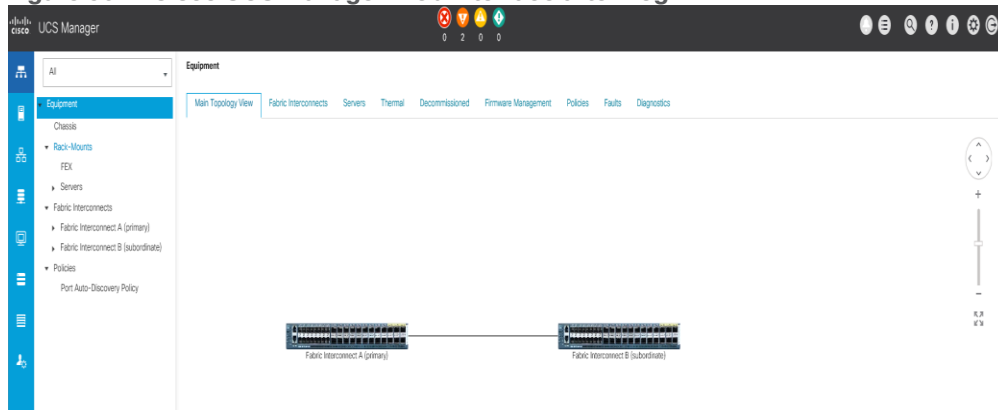
Step 6. Click the Launch UCS Manager link to download the Cisco UCS Manager software. If prompted, accept the security certificates.

Figure 35. Cisco UCS Manager Web Interface



Step 7. When prompted, enter the user name and password enter the password. Click Log In to login to Cisco UCS Manager.

Figure 36. Cisco UCS Manager Web Interface after Login



Configure Base Cisco Unified Computing System

Note: The following are the high-level steps involved for a Cisco UCS configuration.

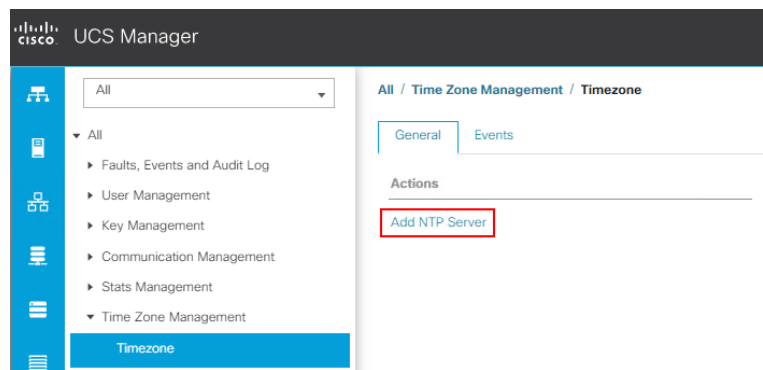
- Configure Fabric Interconnects for a Cluster Setup
- Set Fabric Interconnects to Fibre Channel End Host Mode
- Synchronize Cisco UCS to NTP
- Configure Fabric Interconnects for Chassis and Blade Discovery
- Configure Global Policies
- Configure Server Ports
- Configure LAN and SAN on Cisco UCS Manager
- Configure Ethernet LAN Uplink Ports
- Create Uplink Port Channels to Cisco Nexus Switches
- Configure FC SAN Uplink Ports
- Configure VLAN
- Configure VSAN
- Configure IP, UUID, Server, MAC, WWNN and WWPN Pools
- IP Pool Creation
- UUID Suffix Pool Creation
- Server Pool Creation
- MAC Pool Creation
- WWNN and WWPN Pool Creation
- Set Jumbo Frames in both the Cisco Fabric Interconnect
- Configure Server BIOS Policy
- Create Adapter Policy
- Configure Update Default Maintenance Policy

- Configure vNIC and vHBA Template
- Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

Procedure 1. Synchronize Cisco UCS Manager to NTP

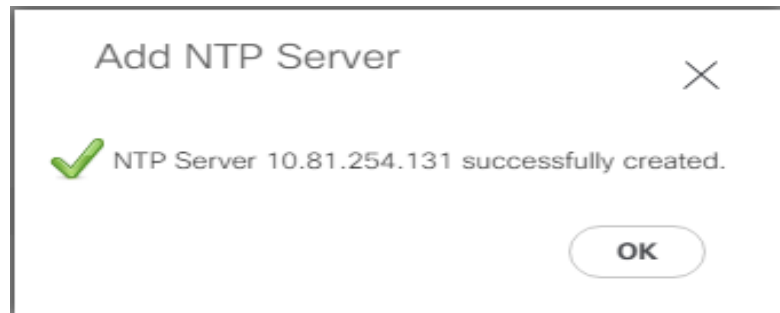
- Step 1.** In Cisco UCS Manager, in the navigation pane, click the Admin tab.
- Step 2.** Select All > Time zone Management.
- Step 3.** In the Properties pane, select the appropriate time zone in the Time zone menu.
- Step 4.** Click Save Changes and then click OK.
- Step 5.** Click Add NTP Server.



- Step 6.** Enter the NTP server IP address and click OK.



- Step 7.** Click OK to finish.



- Step 8.** Repeat steps 1-7 to configure additional NTP servers.
- Step 9.** Click Save Changes.

Configure Fabric Interconnects for Chassis and Blade Discovery

Cisco UCS 6454 Fabric Interconnects are configured for redundancy. It provides resiliency in case of failures. The first step is to establish connectivity between blades and Fabric Interconnects.

Procedure 1. Configure Global Policies

Note: The chassis discovery policy determines how the system reacts when you add a new chassis. We recommend using the platform max value as shown. Using platform max helps ensure that Cisco UCS Manager uses the maximum number of IOM uplinks available.

Step 1. In Cisco UCS Manager, go to Equipment > Policies > Global Policies > Chassis/FEX Discovery Policies. As shown in the screenshot below, for Action select “Platform Max” from the drop-down list and set Link Grouping to Port Channel.

Step 2. Click Save Changes.

Step 3. Click OK.

Figure 37. UCS Global Policy

The screenshot shows the 'Equipment' section of the Cisco UCS Manager interface, specifically the 'Policies' tab. The 'Global Policies' sub-tab is active, and the 'Chassis/FEX Discovery Policy' is selected. The configuration options are as follows:

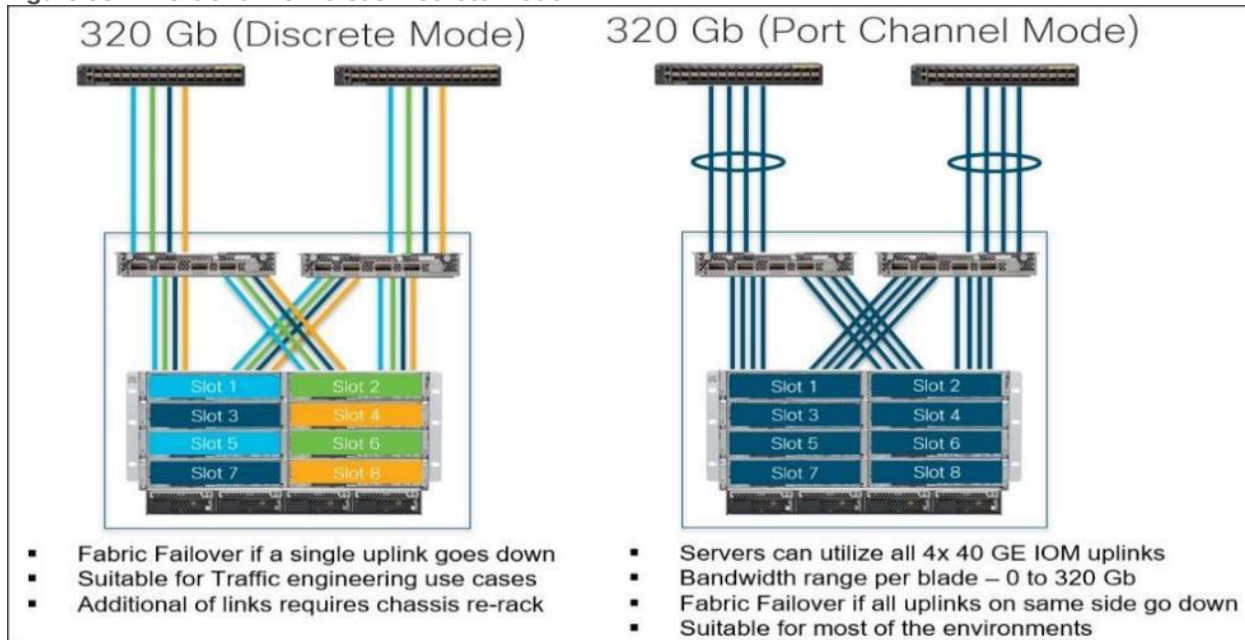
- Chassis/FEX Discovery Policy:**
 - Action: Platform Max (dropdown)
 - Link Grouping Preference: Port Channel (radio button selected)
- Rack Server Discovery Policy:**
 - Action: Immediate (radio button selected)
 - Scrub Policy: <not set> (dropdown)
- Rack Management Connection Policy:**
 - Action: Auto Acknowledged (radio button selected)
- Power Policy:**
 - Redundancy: N+1 (radio button selected)
- MAC Address Table Aging:**
 - Aging Time: Mode Default (radio button selected)
- Global Power Allocation Policy:**
 - Allocation Method: Policy Driven Chassis Group Cap (radio button selected)
- Firmware Auto Sync Server Policy:**
 - Sync State: No Actions (radio button selected)

At the bottom right, there are two buttons: 'Save Changes' and 'Reset Values'.

Fabric Ports: Discrete versus Port Channel Mode

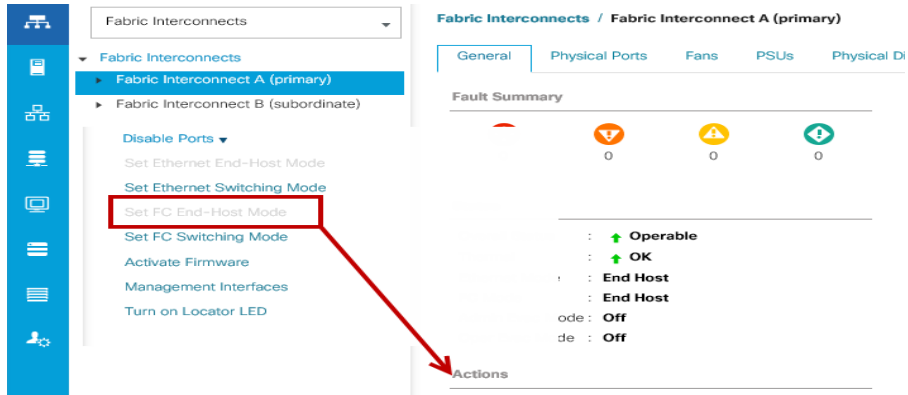
Figure 38 illustrates the advantage of Discrete Vs Port-Channel mode in Cisco UCS Manager.

Figure 38. Port Channel versus Discrete Mode



Procedure 2. Set Fabric Interconnects to Fibre Channel End Host Mode

Step 1. In order to configure the FC Uplink ports connected to the Cisco UCS MDS 9132T 32-Gb FC switch, set the Fabric Interconnects to the Fibre Channel End Host Mode. Verify that the fabric interconnects are operating in “FC End-Host Mode.”



Note: The fabric interconnect automatically reboots if switched to operational mode; perform this task on one FI first, wait for the FI to come up and repeat this process on the second FI.

Procedure 3. Configure FC SAN Uplink Ports

Step 1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > General tab > Actions pane, click Configure Unified Ports.

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate)

General | Physical Ports | Fans | PSUs | Physical Display | FSM | Neighbors | Faults | Events | Statistics

Status

Overall Status : ↑ **Operable**
 Thermal : ↑ **OK**
 Ethernet Mode : **End Host**
 FC Mode : **End Host**
 Admin Evac Mode : **Off**
 Oper Evac Mode : **Off**

Actions

Configure Evacuation
Configure Unified Ports
 Internal Fabric Manager
 LAN Uplinks Manager
 NAS Appliance Manager
 SAN Uplinks Manager
 SAN Storage Manager
 Enable Ports ▾
 Disable Ports ▾
 Set Ethernet End-Host Mode
 Set Ethernet Switching Mode
 Set FC End-Host Mode
 Set FC Switching Mode
 Activate Firmware
 Management Interfaces
 Turn off Locator LED

Properties

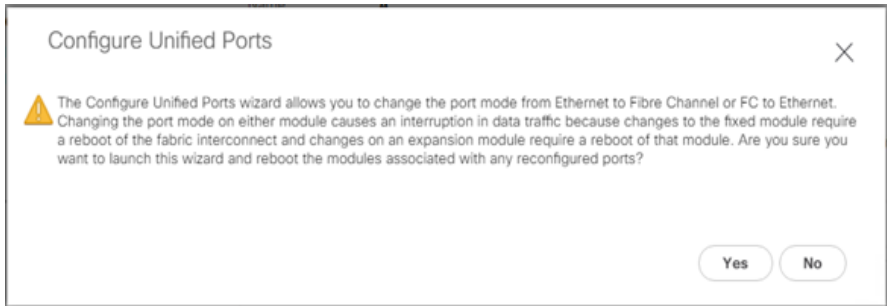
Name : **A**
 Product Name : **Cisco UCS 6454**
 Vendor : **Cisco Systems, Inc.** PID : **UCS-FI-6454**
 Revision : **0** Serial : **FDO22241ZLJ**
 Available Memory : Total Memory : **62.761 (GB)**
 Locator LED : ●

⊕ Part Details
 ⊕ Local Storage Information
 ⊕ Access
 ⊕ High Availability Details
 ⊕ VLAN Port Count
 ⊕ FC Zone Count

Firmware

Boot-loader Version : **v05.40(01/17/2020)**
 Kernel Version : **7.0(3)N2(4.12a)**
 System Version : **7.0(3)N2(4.12a)**
 Service Pack Version : **4.1(2)SP0(Default)**
 Package Version : **4.1(2a)A**
 Startup Kernel Version : **7.0(3)N2(4.12a)**
 Activate Status : **Ready**

Step 2. Click Yes to confirm in the pop-up window.

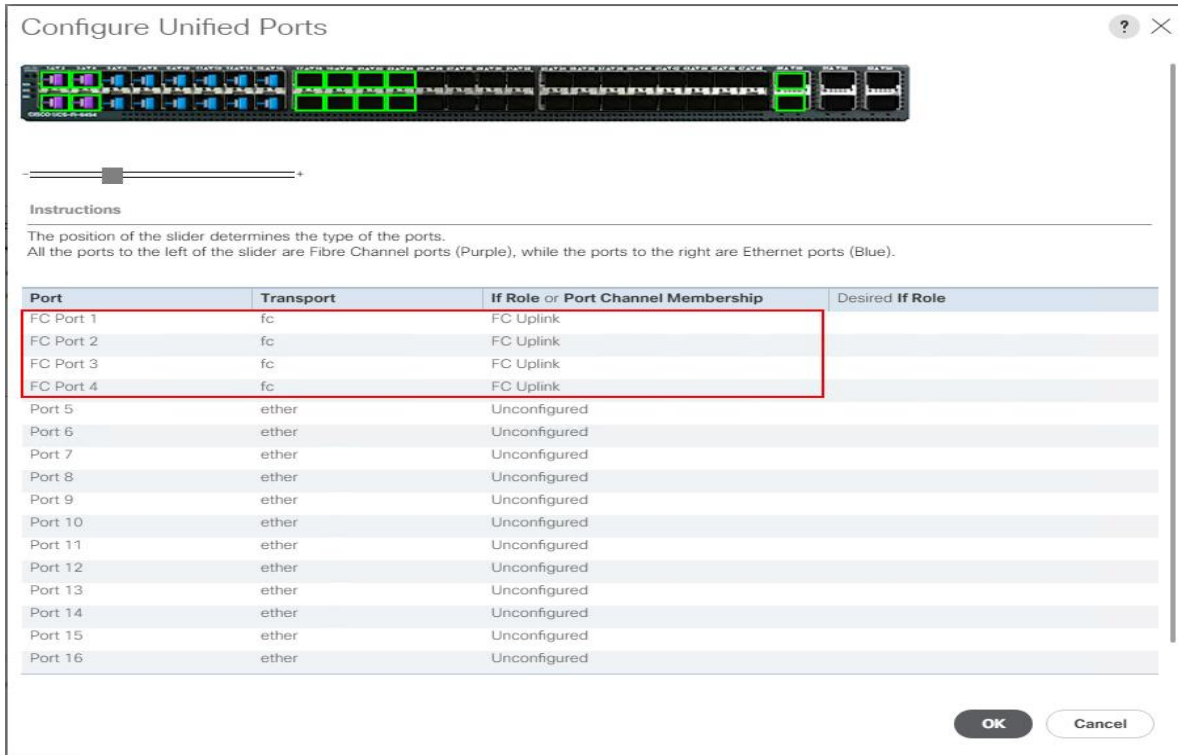


Step 3. Move the slider to the right.

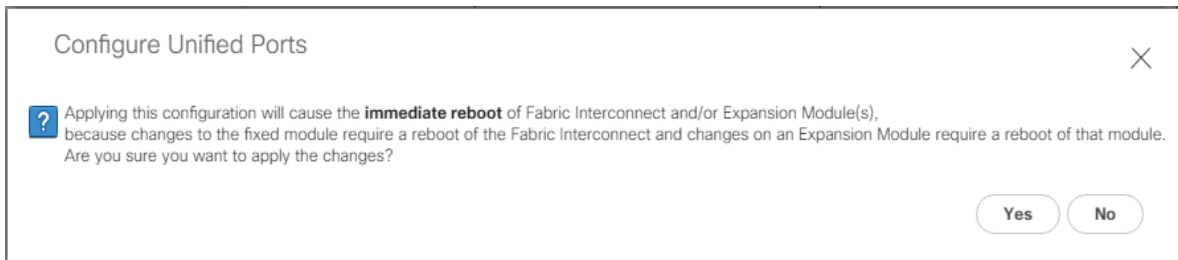
Step 4. Click OK.

Note: Ports to the right of the slider will become FC ports. For our study, we configured the first four ports (Ports are configured in sets of 4 ports) on the FI as FC Uplink ports.


Note: Applying this configuration will cause the immediate reboot of the fabric interconnect and/or the expansion module(s).



Step 5. Click Yes to apply the changes.



Step 6. Click OK to proceed.



After the FI reboot, your FC Ports configuration will look like [Figure 39](#).

Step 7. Repeat steps 1-6 on Fabric Interconnect B.

Figure 39. FC Uplink Ports on Fabric Interconnect A

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module / FC Ports

FC Ports

Advanced Filter Export Print All Unconfigured Network Storage Monitor

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
1	1	20:01:90:3A:9C:0E:33:20	Network	Physical	Up	Enabled
1	2	20:02:00:3A:9C:0E:33:20	Network	Physical	Up	Enabled
1	3	20:03:00:3A:9C:0E:33:20	Network	Physical	Up	Enabled
1	4	20:04:00:3A:9C:0E:33:20	Network	Physical	Up	Enabled

Procedure 4. Configure Server Ports

Configure the server ports to initiate chassis and blade discovery.

Step 1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > Ethernet Ports.

Step 2. Select the ports (for this solution ports are 17-24) which are connected to the Cisco IO Modules of the two B-Series 5108 Chassis.

Step 3. Right-click and select “Configure as Server Port.”

Figure 40. Configure Server Port on Cisco UCS Manager Fabric Interconnect for Chassis/Server Discovery

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module / Ethernet Ports

Ethernet Ports

Advanced Filter Export Print All Unconfigured Network Server FCoE Uplink Unified Uplink Appliance Storage FCoE Storage Unified Storage Monitor

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	17	00:3A:9C:0E:33:38	Unconfigured	Physical	Admin Down	Disabled	
1	0	18	00:3A:9C:0E:33:39	Unconfigured	Physical	Admin Down	Disabled	
1	0	19	00:3A:9C:0E:33:3A	Unconfigured	Physical	Admin Down	Disabled	
1	0	20	00:3A:9C:0E:33:3B	Unconfigured	Physical	Admin Down	Disabled	
1	0	21	00:3A:9C:0E:33:3C	Unconfigured	Physical	Admin Down	Disabled	
1	0	22	00:3A:9C:0E:33:3D	Unconfigured	Physical	Admin Down	Disabled	
1	0	23	00:3A:9C:0E:33:3E	Unconfigured	Physical	Admin Down	Disabled	
1	0	24	00:3A:9C:0E:33:3F	Unconfigured	Physical	Admin Down	Disabled	
1	0	25	00:3A:9C:0E:33:40	Unconfigured	Physical	Stp Not Present	Disabled	
1	0	26	00:3A:9C:0E:33:41	Unconfigured	Physical	Stp Not Present	Disabled	
1	0	27	00:3A:9C:0E:33:42	Unconfigured	Physical	Stp Not Present	Disabled	
1	0	28	00:3A:9C:0E:33:43	Unconfigured	Physical	Stp Not Present	Disabled	
1	0	29	00:3A:9C:0E:33:44	Unconfigured	Physical	Stp Not Present	Disabled	

Step 4. Click Yes to confirm and click OK.

Step 5. Repeat steps 1-4 to configure the Server Port on Fabric Interconnect B.

When configured, the server port will look like [Figure 41](#) on both Fabric Interconnects.

Figure 41. Server Ports on Fabric Interconnect A

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	17	00:3A:9C:0E:33:38	Server	Physical	Up	Enabled	sys/chassis-1/slot-2/fabri...
1	0	18	00:3A:9C:0E:33:39	Server	Physical	Up	Enabled	sys/chassis-1/slot-2/fabri...
1	0	19	00:3A:9C:0E:33:3A	Server	Physical	Up	Enabled	sys/chassis-2/slot-2/fabri...
1	0	20	00:3A:9C:0E:33:3B	Server	Physical	Up	Enabled	sys/chassis-2/slot-2/fabri...
1	0	21	00:3A:9C:0E:33:3C	Server	Physical	Up	Enabled	sys/chassis-3/slot-2/fabri...
1	0	22	00:3A:9C:0E:33:3D	Server	Physical	Up	Enabled	sys/chassis-3/slot-2/fabri...
1	0	23	00:3A:9C:0E:33:3E	Server	Physical	Up	Enabled	sys/chassis-4/slot-2/fabri...
1	0	24	00:3A:9C:0E:33:3F	Server	Physical	Link Up	Enabled	sys/chassis-4/slot-2/fabri...
1	0	25	00:3A:9C:0E:33:40	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	26	00:3A:9C:0E:33:41	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	27	00:3A:9C:0E:33:42	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	28	00:3A:9C:0E:33:43	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	29	00:3A:9C:0E:33:44	Unconfigured	Physical	Sfp Not Present	Disabled	

Step 6. After configuring Server Ports, acknowledge both the Chassis. Go to Equipment > Chassis > Chassis 1 > General > Actions > select “Acknowledge Chassis”. Similarly, acknowledge the chassis 2-4.

Step 7. After acknowledging both the chassis, re-acknowledge all the servers placed in the chassis. Go to Equipment > Chassis 1 > Servers > Server 1 > General > Actions > select Server Maintenance > select option “Re-acknowledge” and click OK. Repeat this process to re-acknowledge all eight Servers.

Step 8. When the acknowledgement of the Servers is completed, verify the Port-channel of Internal LAN. Go to the LAN tab > Internal LAN > Internal Fabric A > Port Channels as shown in [Figure 42](#).

Figure 42. Internal LAN Port Channels

Name	Slot ID	Port ID	Aggr. Port ID	Peer Slot ID	Peer Port ID	Fabric ID	Peer
Eth Interface 1/17	1	17	0	2	1	A	sys/switch-A/access-eth/ep...
Eth Interface 1/18	1	18	0	2	5	A	sys/switch-A/access-eth/ep...

Procedure 5. Configure Ethernet LAN Uplink Ports

- Step 1.** In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
- Step 2.** Select Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
- Step 3.** Expand Ethernet Ports.

Step 4. Select ports (for this solution ports are 49-50) that are connected to the Nexus switches, right-click them, and select Configure as Network Port.

Figure 43. Network Uplink Port Configuration on Fabric Interconnect Configuration

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 35	1	35	00:3A:9C:0E:33:4A	Unconfigured	Physical	Sfp Not Present	Disabled
Port 36	1	36	00:3A:9C:0E:33:4B	Unconfigured	Physical	Sfp Not Present	Disabled
Port 37	1	37	00:3A:9C:0E:33:4C	Unconfigured	Physical	Sfp Not Present	Disabled
Port 38	1	38	00:3A:9C:0E:33:4D	Unconfigured	Physical	Sfp Not Present	Disabled
Port 39	1	39	00:3A:9C:0E:33:4E	Unconfigured	Physical	Sfp Not Present	Disabled
Port 40	1	40	00:3A:9C:0E:33:4F	Unconfigured	Physical	Sfp Not Present	Disabled
Port 41	1	41	00:3A:9C:0E:33:50	Unconfigured	Physical	Sfp Not Present	Disabled
Port 42	1	42	00:3A:9C:0E:33:51	Unconfigured	Physical	Sfp Not Present	Disabled
Port 43	1	43	00:3A:9C:0E:33:52	Unconfigured	Physical	Sfp Not Present	Disabled
Port 44	1	44	00:3A:9C:0E:33:53	Unconfigured	Physical	Sfp Not Present	Disabled
Port 45	1	45	00:3A:9C:0E:33:54	Unconfigured	Physical	Sfp Not Present	Disabled
Port 46	1	46	00:3A:9C:0E:33:55	Unconfigured	Physical	Sfp Not Present	Disabled
Port 47	1	47	00:3A:9C:0E:33:56	Unconfigured	Physical	Sfp Not Present	Disabled
Port 48	1	48	00:3A:9C:0E:33:57	Unconfigured	Physical	Sfp Not Present	Disabled
Port 49	1	49	00:3A:9C:0E:33:58	Unconfigured	Physical	Admin Down	Disabled
Port 50	1	50	00:3A:9C:0E:33:5C	Unconfigured	Physical	Admin Down	Disabled
Port 51	1	51	00:3A:9C:0E:33:60	Unconfigured	Physical	Sfp Not Present	Disabled
Port 52	1	52	00:3A:9C:0E:33:64	Unconfigured	Physical	Sfp Not Present	Disabled
Port 53	1	53	00:3A:9C:0E:33:68	Unconfigured	Physical	Sfp Not Present	Disabled
Port 54	1	54	00:3A:9C:0E:33:6C	Unconfigured	Physical	Sfp Not Present	Disabled

Step 5. Click Yes to confirm ports and click OK.

Step 6. Verify the Ports connected to Cisco Nexus upstream switches are now configured as network ports.

Step 7. Repeat steps 1-6 for Fabric Interconnect B. The screenshot below shows the network uplink ports for Fabric A.

Figure 44. Network Uplink Port on Fabric Interconnect

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 37	1	37	00:3A:9C:0E:33:4C	Unconfigured	Physical	Sfp Not Present	Disabled
Port 38	1	38	00:3A:9C:0E:33:4D	Unconfigured	Physical	Sfp Not Present	Disabled
Port 39	1	39	00:3A:9C:0E:33:4E	Unconfigured	Physical	Sfp Not Present	Disabled
Port 40	1	40	00:3A:9C:0E:33:4F	Unconfigured	Physical	Sfp Not Present	Disabled
Port 41	1	41	00:3A:9C:0E:33:50	Unconfigured	Physical	Sfp Not Present	Disabled
Port 42	1	42	00:3A:9C:0E:33:51	Unconfigured	Physical	Sfp Not Present	Disabled
Port 43	1	43	00:3A:9C:0E:33:52	Unconfigured	Physical	Sfp Not Present	Disabled
Port 44	1	44	00:3A:9C:0E:33:53	Unconfigured	Physical	Sfp Not Present	Disabled
Port 45	1	45	00:3A:9C:0E:33:54	Unconfigured	Physical	Sfp Not Present	Disabled
Port 46	1	46	00:3A:9C:0E:33:55	Unconfigured	Physical	Sfp Not Present	Disabled
Port 47	1	47	00:3A:9C:0E:33:56	Unconfigured	Physical	Sfp Not Present	Disabled
Port 48	1	48	00:3A:9C:0E:33:57	Unconfigured	Physical	Sfp Not Present	Disabled
Port 49	1	49	00:3A:9C:0E:33:58	Network	Physical	Up	Enabled
Port 50	1	50	00:3A:9C:0E:33:5C	Network	Physical	Up	Enabled

You have now created two uplink ports on each Fabric Interconnect as shown above. These ports will be used to create Virtual Port Channels.

Procedure 6. Create Uplink Port Channels to Cisco Nexus Switches

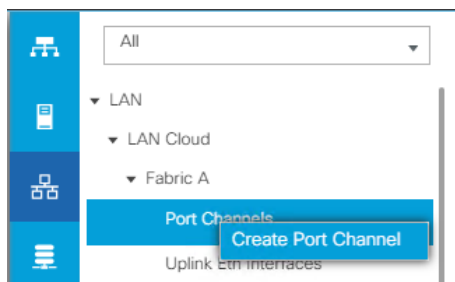
Note: In this procedure, two port channels are created one from Fabric A to both Cisco Nexus 93180YC-FX switches and one from Fabric B to both Cisco Nexus 93180YC-FX switches.

Step 1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

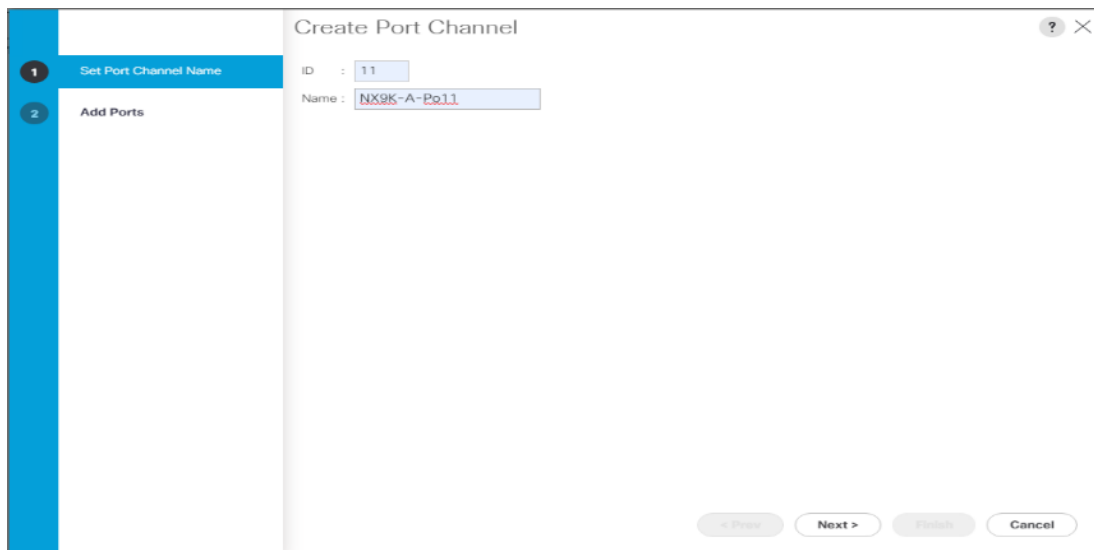
Step 2. Click LAN > LAN Cloud > Fabric A.

Step 3. Right-click Port Channels.

Step 4. Select Create Port Channel.

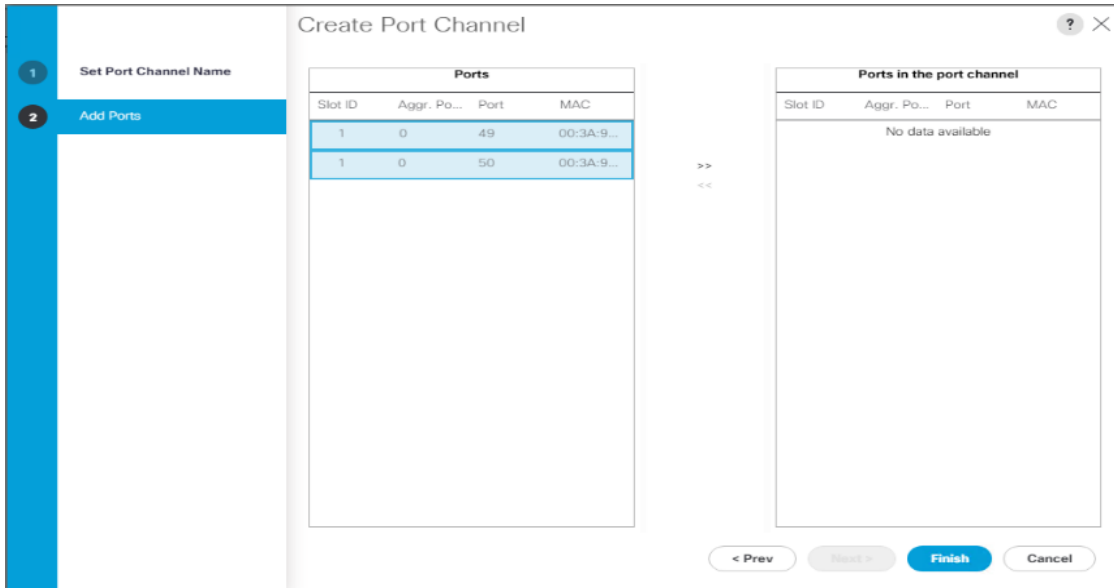


Step 5. Enter 11 as the unique ID of the port channel and name of the port channel.

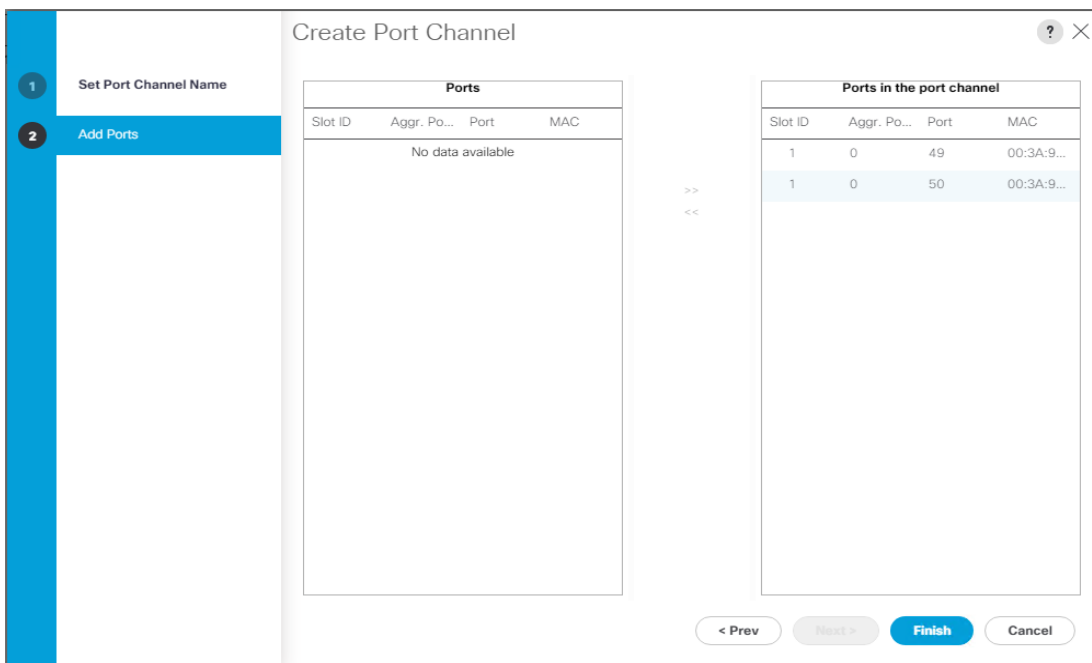


Step 6. Click Next.

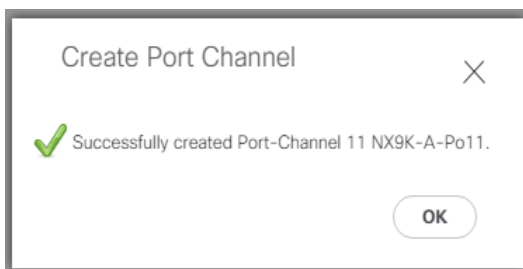
Step 7. Select Ethernet ports 49-50 for the port channel.



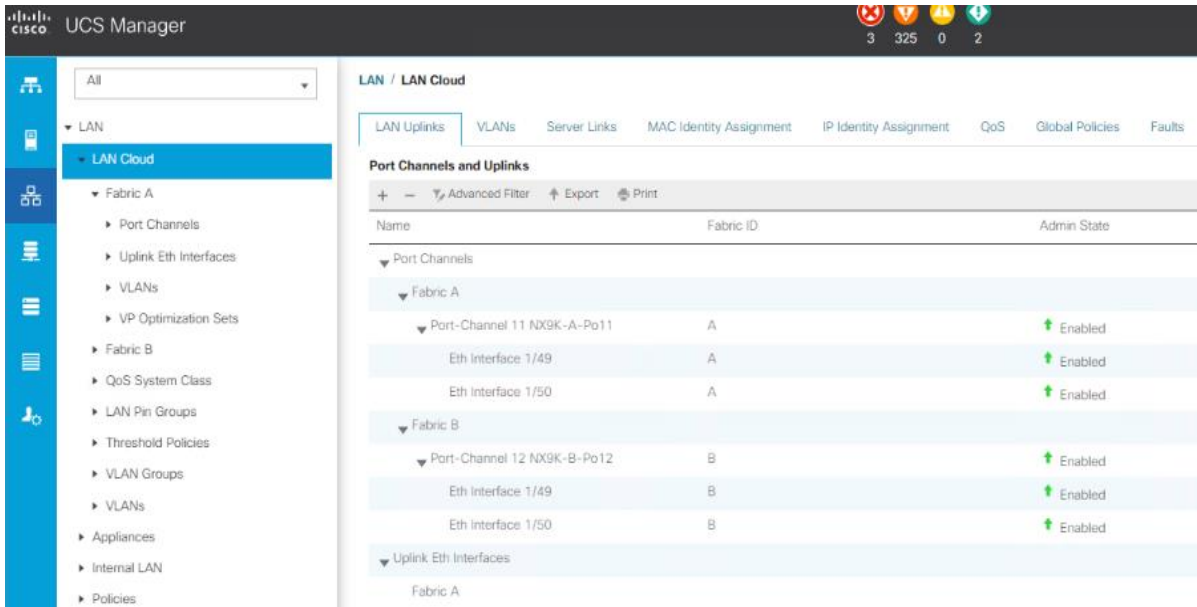
Step 8. Click Finish.



Step 9. Click OK.



Step 10. Repeat steps 1–9 for the Port Channel configuration on FI-B.



Procedure 7. Configure VLAN

- Step 1.** In Cisco UCS Manager, click the LAN tab in the navigation pane.
- Step 2.** Click LAN > LAN Cloud.
- Step 3.** Right-click VLANs.
- Step 4.** Select Create VLANs.
- Step 5.** Enter InBand-Mgmt as the name of the VLAN to be used for Public Network Traffic.
- Step 6.** Keep the Common/Global option selected for the scope of the VLAN.
- Step 7.** Enter 60 as the ID of the VLAN ID.
- Step 8.** Keep the Sharing Type as None.
- Step 9.** Click OK.

Create VLANs ? X

VLAN Name/Prefx :

Multicast Policy Name : [Create Multicast Policy](#)

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics. Enter the range of VLAN IDs.(e.g. " 2009-2019", " 29,35,40-45", " 23", " 23,34-45")

VLAN IDs :

Sharing Type : None Primary Isolated Community

Step 10. Repeat steps 1–9 to create required VLANs. [Figure 45](#) shows the VLANs configured for this solution.

Figure 45. VLANs Configured for this Solution

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN In-Band-Mgmt (60)	60	Lan	Ether	No	None
VLAN Infra-Mgmt (61)	61	Lan	Ether	No	None
VLAN CIFS-VLAN (62)	62	Lan	Ether	No	None
VLAN NFS-Vlan (63)	63	Lan	Ether	No	None
VLAN VDI-VLAN (64)	64	Lan	Ether	No	None
VLAN Launcher-65 (65)	65	Lan	Ether	No	None
VLAN vMotion (66)	66	Lan	Ether	No	None

IMPORTANT! Create both VLANs with global access across both fabric interconnects. This makes sure the VLAN identity is maintained across the fabric interconnects in case of a NIC failover.

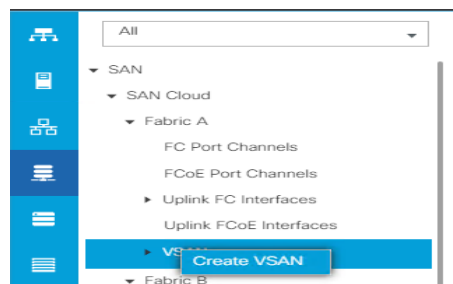
Procedure 8. Configure VSAN

Step 1. In Cisco UCS Manager, click the SAN tab in the navigation pane.

Step 2. Select SAN > SAN Cloud.

Step 3. Under VSANs, right-click VSANs.

Step 4. Select Create VSANs.



Step 5. Enter the name of the VSAN, such as FlexPod-A.

Note: In this solution, we created two VSANs; VSAN FlexPod-A 400 on the Cisco UCS Fabric A and VSAN FlexPod-B 401 on the Cisco UCS Fabric B for SAN Boot and Storage Access.

Step 6. Select Disabled for FC Zoning

Note: In this solution we used two Cisco MDS 9132T 32Gb switches that provide Fibre Channel zoning.

Step 7. Select Fabric A for the scope of the VSAN:

- a. Enter 400 as VSAN ID and FCoE VLAN ID.
- b. Click OK.

Create VSAN



Name : FlexPod- VSAN-A

FC Zoning Settings

FC Zoning : Disabled Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 400

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN : 400

OK

Cancel

Step 8. Repeat steps 1-7 to create the VSANs necessary for this solution.

[Figure 46](#) shows VSAN 400 and 401 configured for this solution.

Figure 46. VSANs Configured for this Solution

Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operational State
VSAN 400-K (400)	400	A	Virtual	Network	Fc	430	OK
VSAN 401 (401)	401	A	Virtual	Network	Fc	421	OK

Procedure 9. Create New Sub-Organization

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Select root > Sub-Organization.
- Step 3.** Right-click Sub-Organization.
- Step 4.** Enter the name of the Sub-Organization.
- Step 5.** Click OK.

Procedure 10. IP Pool Creation

Note: An IP address pool on the out-of-band management network must be created to facilitate KVM access to each compute node in the Cisco UCS domain.

- Step 1.** In Cisco UCS Manager, in the navigation pane, click the LAN tab.
- Step 2.** Click Pools > root > Sub-Organizations > FlexPod-CVD > IP Pools > click Create IP Pool.
- Step 3.** Select the option Sequential to assign IP in sequential order then click Next.

Create IP Pool [?] X

1 Define Name and Description

2 Add IPv4 Blocks

3 Add IPv6 Blocks

Name : FlexPod- KVM Pool

Description :

Assignment Order : Default Sequential

Step 4. Click Add IPv4 Block.

Step 5. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information as shown below.

Create IP Pool [?] X

1 Define Name and Description

2 Add IPv4 Blocks

3 Add IPv6 Blocks

Create Block of IPv4 Addresses [?] X

From : 10.29.164.166 Size : 32

Subnet Mask : 255.255.255.0 Default Gateway : 10.29.164.1

Primary DNS : 0.0.0.0 Secondary DNS : 0.0.0.0

OK Cancel

+ Add - Delete

Procedure 11. UUID Suffix Pool Creation

Step 1. In Cisco UCS Manager, click the Servers tab in the navigation pane.

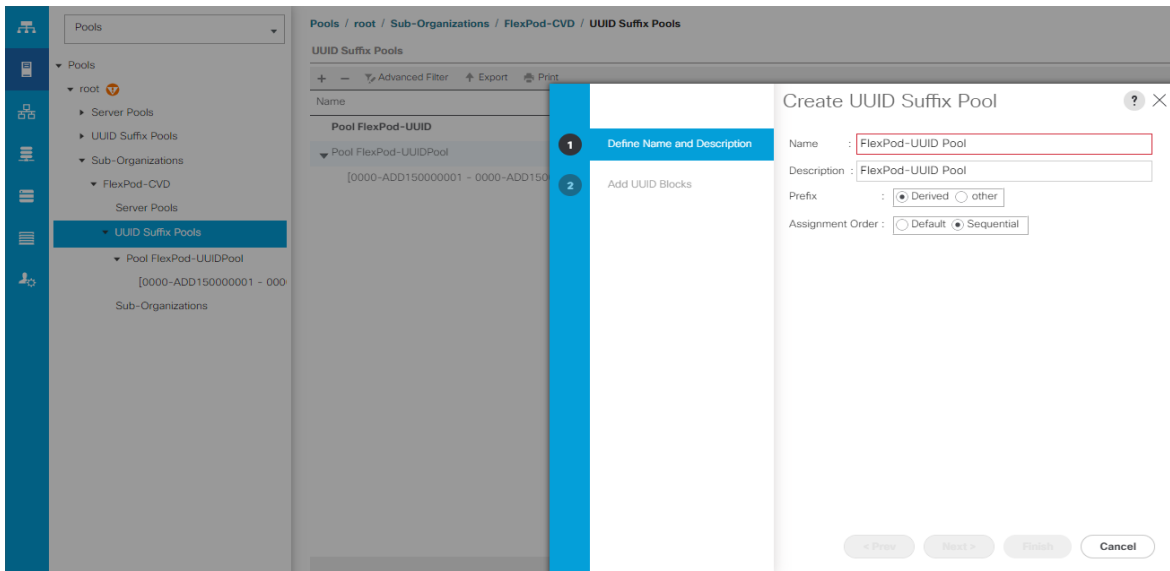
Step 2. Click Pools > root > Sub-Organization > FlexPod-CVD.

Step 3. Right-click UUID Suffix Pools and then select Create UUID Suffix Pool.

Step 4. Enter the name of the UUID name.

Step 5. Optional: Enter a description for the UUID pool.

Step 6. Keep the prefix at the derived option and select Sequential in as Assignment Order then click Next.



Step 7. Click Add to add a block of UUIDs.

Step 8. Create a starting point UUID as per your environment.

Step 9. Specify a size for the UUID block that is sufficient to support the available blade or server resources.



Procedure 12. Server Pool Creation

Tech tip

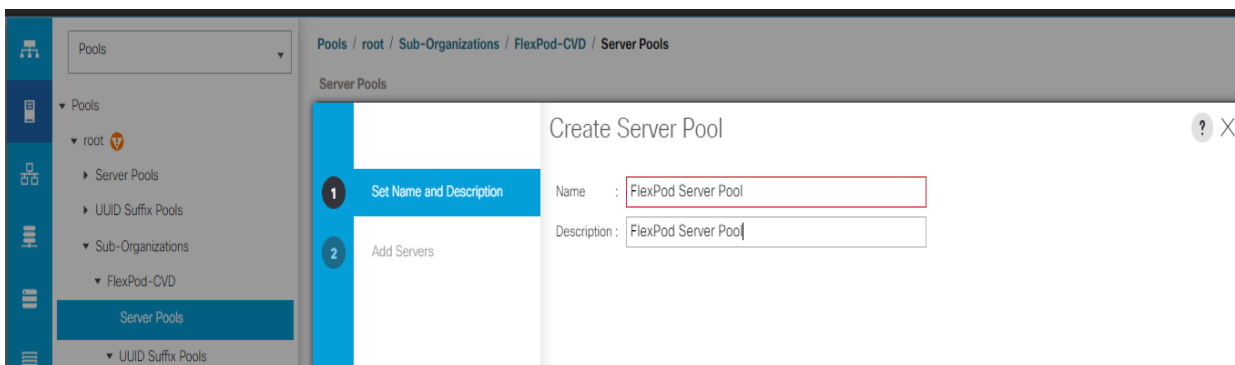
Consider creating unique server pools to achieve the granularity that is required in your environment.

Step 1. In Cisco UCS Manager, click the Servers tab in the navigation pane.

Step 2. Click Pools > root > Sub-Organization > FlexPod-CVD > right-click Server Pools > Select Create Server Pool.

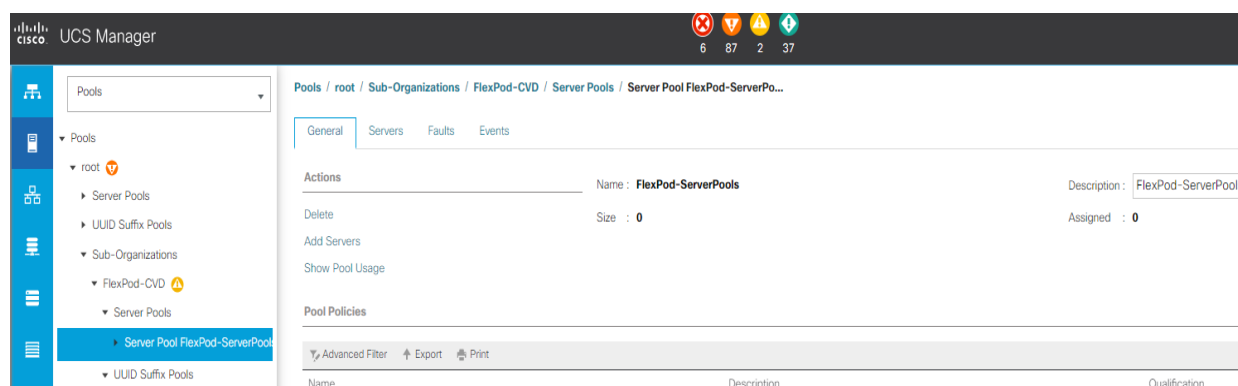
Step 3. Enter name of the server pool.

Step 4. Optional: Enter a description for the server pool then click Next.



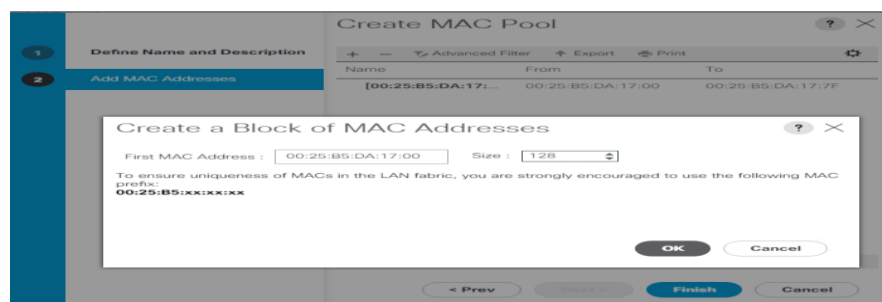
Step 5. Select servers to be used for the deployment and click > to add them to the server pool. In our case we added thirty servers in this server pool.

Step 6. Click Finish and then click OK.

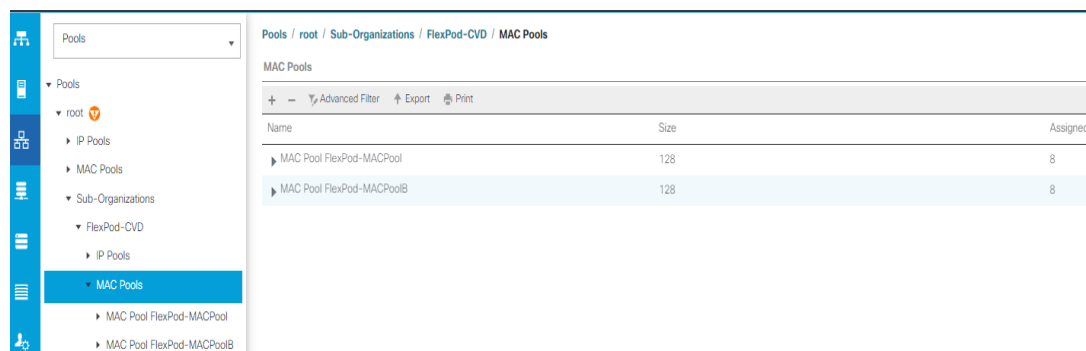


Procedure 13. MAC Pool Creation

- Step 1.** In Cisco UCS Manager, click the LAN tab in the navigation pane.
- Step 2.** Click Pools > root > Sub-Organization > FlexPod > right-click MAC Pools under the root organization.
- Step 3.** Click Create MAC Pool to create the MAC address pool.
- Step 4.** Enter name for MAC pool. Select Assignment Order as “Sequential.”
- Step 5.** Enter the seed MAC address and provide the number of MAC addresses to be provisioned.
- Step 6.** Click OK and then click Finish.
- Step 7.** In the confirmation message, click OK.



Step 8. Create MAC Pool B and assign unique MAC Addresses as shown below.



Procedure 14. WWNN and WWPN Pool Creation

WWNN Pool

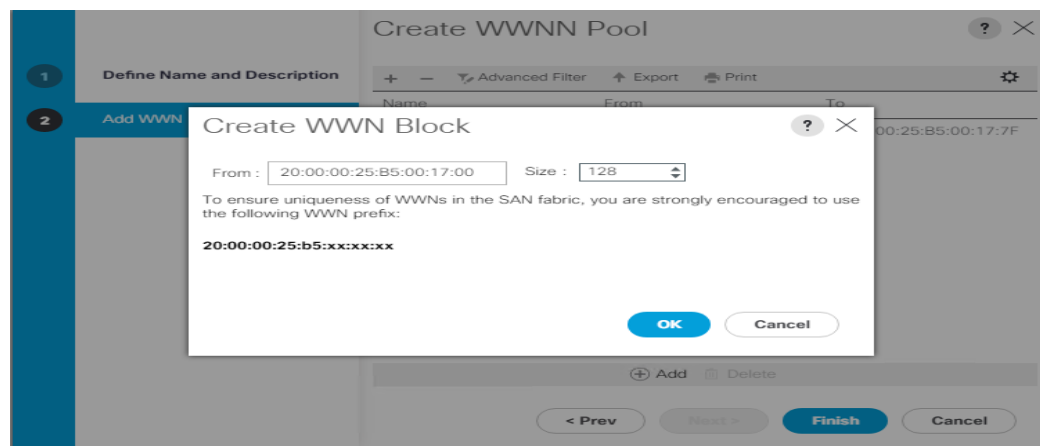
Step 1. In Cisco UCS Manager, click the SAN tab in the navigation pane.

Step 2. Click Pools > Root > Sub-Organization > FlexPod-CVD > WWNN Pools > right-click WWNN Pools > select Create WWNN Pool.

Step 3. Assign name and Assignment Order as sequential.

Step 4. Click Next and then click Add to add block of Ports.

Step 5. Enter Block for WWN and size of WWNN Pool as shown below.



Step 6. Click OK and then click Finish.

WWPN Pool

Note: We created two WWPN as WWPN-A Pool and WWPN-B as Worldwide Port Name as shown below. These WWNN and WWPN entries will be used to access storage through SAN configuration.

Step 1. In Cisco UCS Manager, click the SAN tab in the navigation pane.

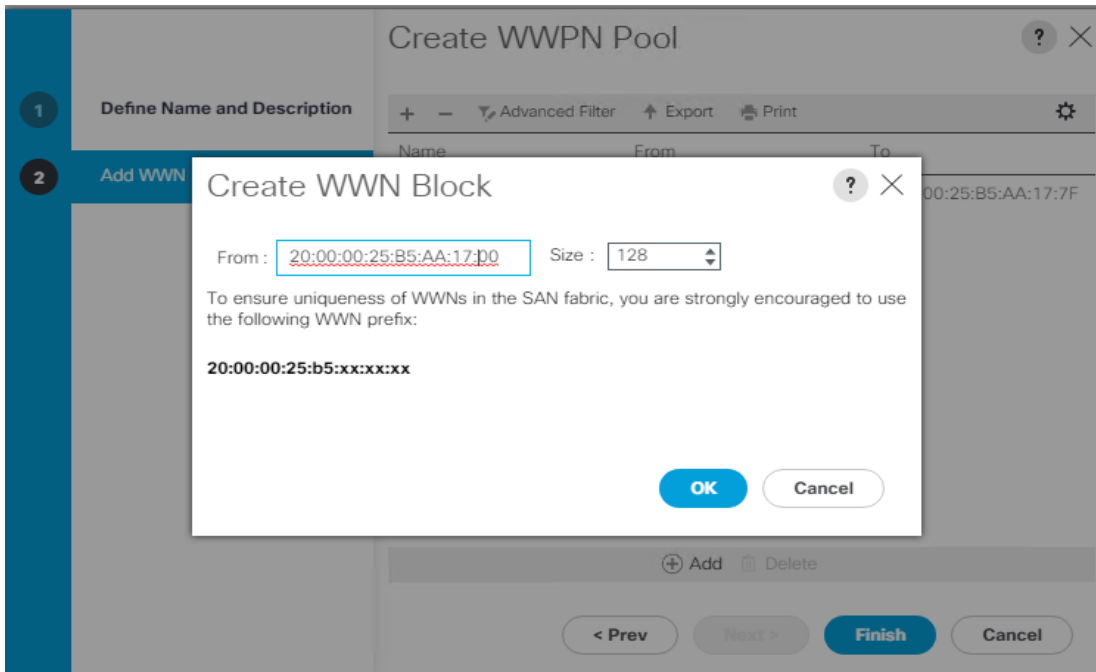
Step 2. Select Pools > Root > WWPN Pools > right-click WWPN Pools > select Create WWPN Pool.

Step 3. Assign name and Assignment Order as sequential.

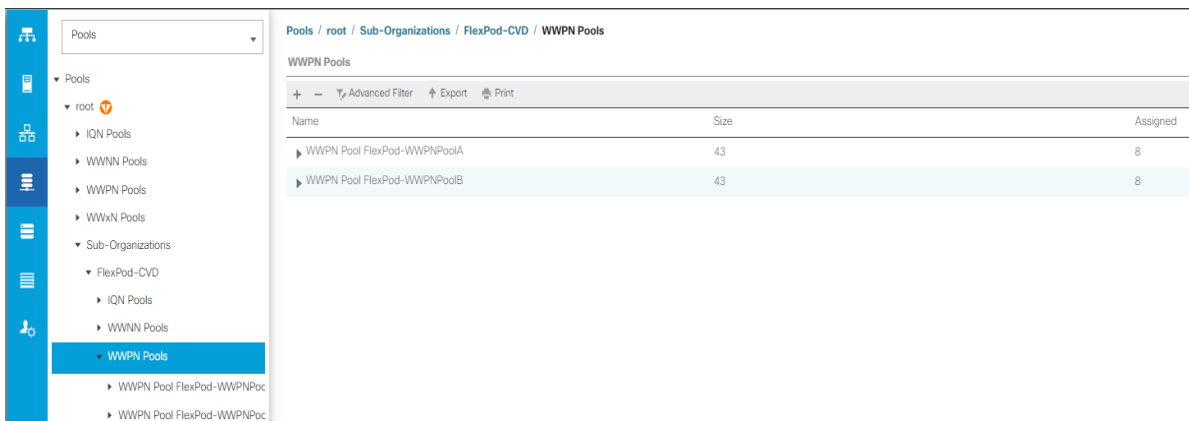
Step 4. Click Next and then click Add to add block of Ports.

Step 5. Enter Block for WWN and size.

Step 6. Click OK and then click Finish.



Step 7. Configure the WWPN-B Pool and assign the unique block IDs as shown below.



Procedure 15. Set Jumbo Frames in both Cisco Fabric Interconnect

Step 1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

Step 2. Select LAN > LAN Cloud > QoS System Class.

Step 3. In the right pane, click the General tab.

Step 4. On the Best Effort row, enter 9216 in the box under the MTU column.

Step 5. Click Save Changes.

Step 6. Click OK.

LAN Cloud / QoS System Class

General Events FSM

Actions Properties

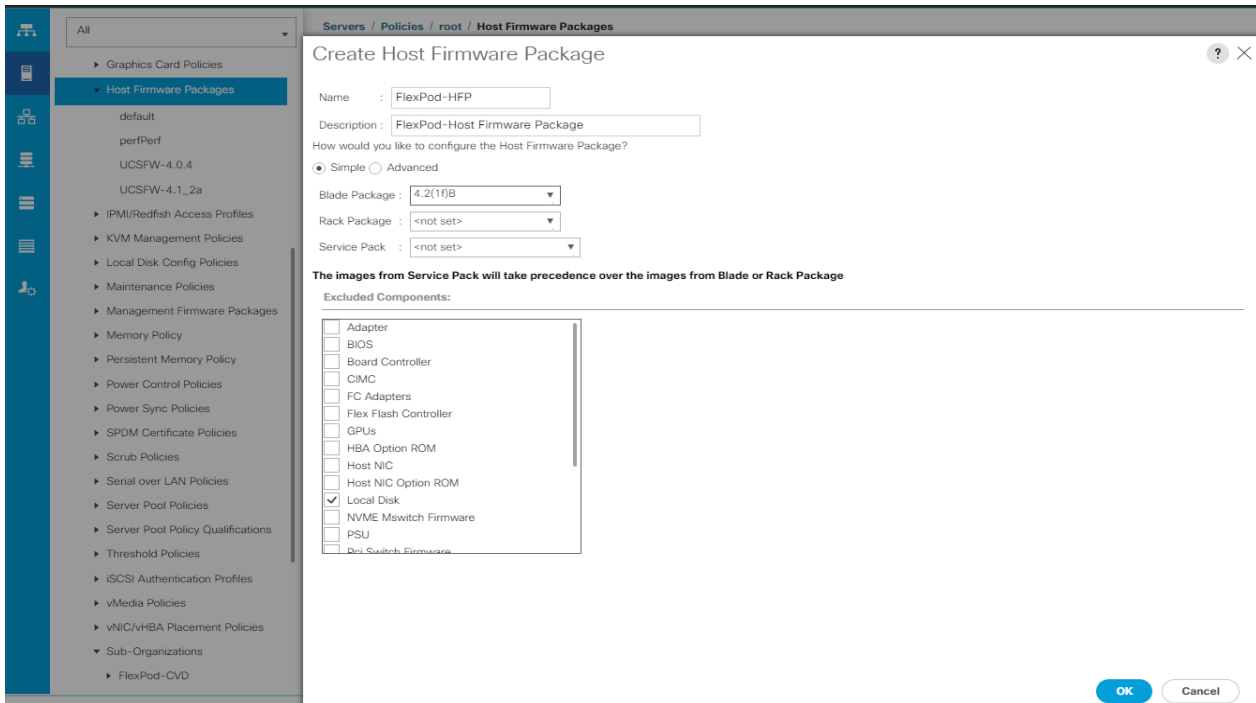
Use Global Owner: Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

Procedure 16. Create Host Firmware Package

Note: Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Select root > Sub-Organization > FlexPod-CVD > Host Firmware Packages.
- Step 3.** Right-click Host Firmware Packages.
- Step 4.** Select Create Host Firmware Package.
- Step 5.** Enter name of the host firmware package.
- Step 6.** Leave Simple selected.
- Step 7.** Select the version 4.2(1f) for both the Blade Package.
- Step 8.** Click OK to create the host firmware package.

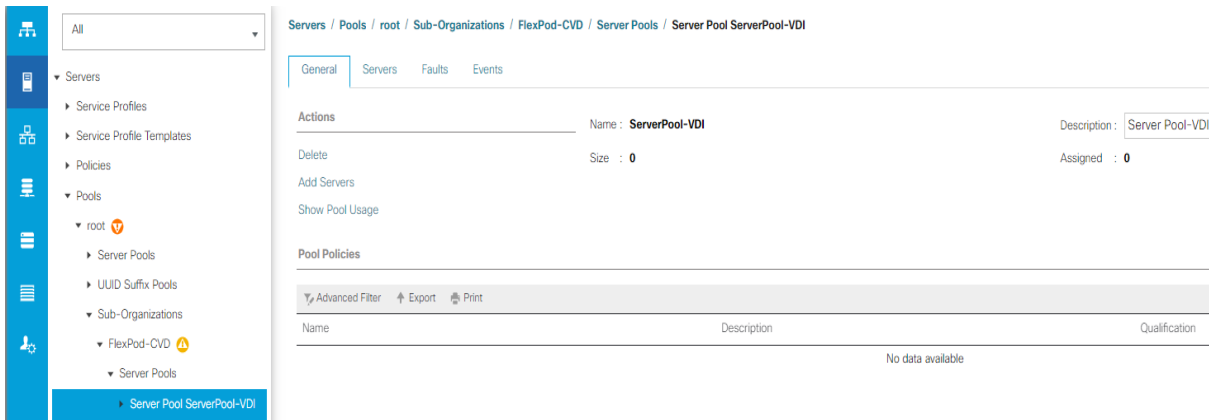


Procedure 17. Create Server Pool Policy

Note: Creating the server pool policy requires you to create the Server Pool Policy and Server Pool Qualification Policy.

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Click Pools > root > Sub-Organization > FlexPod-CVD > Server Pools.
- Step 3.** Right-click Server Pools > select Create Server Pools Policy; Enter Policy name.
- Step 4.** Select server from left pane to add as pooled server.

Note: In our case, we created two server pools policies. For the “VDI-CVD01” policy, we added Servers as Chassis 1 Slot 1-8 and Chassis 3 Slot 1-8 and for the “VDI-CVD02” policy, we added Chassis 2 Slot 1-8 and Chassis 4 Slot 1-8.



Procedure 18. Create Server Pool Policy Qualifications

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.

- Step 2.** Click Pools > root > Sub-Organization > FlexPod-CVD > Server Pool Policy Qualification.
- Step 3.** Right-click Server Pools Select Create Server Pool Policy Qualification; Enter Policy name.
- Step 4.** Select Chassis/Server Qualification from left pane to add in Qualifications.
- Step 5.** Click Add or OK to either Add more servers to existing policy to Finish creation of Policy.

Create Server Pool Policy Qualification [?] [X]

Naming

Name :

Description :

This server pool policy qualification will apply to new or re-discovered servers. Existing servers are not qualified until they are re-discovered

Actions

- Create Adapter Qualifications
- Create Chassis/Server Qualifications**
- Create Memory Qualifications
- Create CPU/Cores Qualifications
- Create Storage Qualifications
- Create Server PID Qualifications
- Create Power Group Qualifications
- Create Rack Qualifications

Qualifications

Name	Max	Model	From	To	Architec...	Speed	Stepping	Power G...
Chassis id range [1 - 1]			1	1				

[+] Add [X] Delete [i] Info

Note: In our case, we created two server pools policies. For the “VDI-CVD01” policy, we added Servers as Chassis 1 Slot 1-8 and Chassis 3 Slot 1-8 and for the “VDI-CVD02” policy, we added Chassis 2 Slot 1-8 and Chassis 4 Slot 1-8.

Server Pool Policy Qualifications

Name	Max	Model	From	To
▼ VCC-CVD01-Qual				
Chassis id range [1 - 1]			1	1
Chassis id range [3 - 3]			3	3
▼ VCC-CVD02-Qual				
Chassis id range [2 - 2]			2	2
Chassis id range [4 - 4]			4	4

Procedure 19. Create a Target Pool and Qualification

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Click Pools > root > Sub-Organization > FlexPod-CVD > Server Pool Policies.
- Step 3.** Right-click Server Pool Policies and Select Create Server Pool Policy; Enter Policy name.
- Step 4.** Select Target Pool and Qualification from the drop-down list.
- Step 5.** Click OK.

Create Server Pool Policy [?] [X]

Name :

Description :

Target Pool :

Qualification :

Note: We created two Server Pool Policies to associate with the Service Profile Templates “VDI-CVD01” and “VDI-CVD02” as described in this section.

Procedure 20. Create Network Control Policy for Cisco Discovery Protocol

- Step 1.** In Cisco UCS Manager, click the LAN tab in the navigation pane.
- Step 2.** Click Policies > root > Sub-Organization > FlexPod-CVD > Network Control Policies.
- Step 3.** Right-click Network Control Policies.
- Step 4.** Click Create Network Control Policy.
- Step 5.** Enter policy name.
- Step 6.** Select the Enabled option for “CDP.”
- Step 7.** Click OK to create the network control policy.

Create Network Control Policy

Name : CDP_Enabled

Description :

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

OK Cancel

Procedure 21. Create Power Control Policy

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Click Policies > root > Sub-Organization > FlexPod-CVD > Power Control Policies.
- Step 3.** Right-click Power Control Policies.
- Step 4.** Click Create Power Control Policy.
- Step 5.** Select Fan Speed Policy as “Max Power.”
- Step 6.** Enter NoPowerCap as the power control policy name.
- Step 7.** Change the power capping setting to No Cap.
- Step 8.** Click OK to create the power control policy.

Create Power Control Policy

Name : NoPowerCap

Description :

Fan Speed Policy : Max Power

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

No Cap cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK Cancel

Procedure 22. Create Server BIOS Policy

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Click Policies > root > Sub-Organization > FlexPod-CVD > BIOS Policies.
- Step 3.** Right-click BIOS Policies.
- Step 4.** Click Create BIOS Policy.
- Step 5.** Enter B200-M6-BIOS as the BIOS policy name.
- Step 6.** Click OK to create policy.

Create BIOS Policy

Name : B200M6 - BIOS

Description : B200M6 - BIOS Policy

Reboot on BIOS Settings Change :

OK Cancel

- Step 7.** Leave all BIOS Settings as “Platform Default.”

Procedure 23. Configure Maintenance Policy

- Step 1.** In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Step 2.** Click Policies > root > Sub-Organization > FlexPod-CVD > Maintenance Policies.
- Step 3.** Right-click Maintenance Policies to create a new policy.
- Step 4.** Enter name for Maintenance Policy
- Step 5.** Change the Reboot Policy to User Ack.
- Step 6.** Click Save Changes.
- Step 7.** Click OK to accept the change.

Policies / root / Sub-Organizations / FlexPod-CVD / Maintenance Policies / FlexPodMaint

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : FlexPodMaint

Description : FlexPod

Owner : Local

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : Immediate User Ack

Reboot Policy : Immediate User Ack Timer Automatic

On Next Boot (Apply pending changes at next reboot.)

Procedure 24. Create vNIC Templates

- Step 1.** In Cisco UCS Manager, click the LAN tab in the navigation pane.
- Step 2.** Click Policies > root > Sub-Organization > FlexPod-CVD > vNIC Template.
- Step 3.** Right-click vNIC Templates.
- Step 4.** Click Create vNIC Template.
- Step 5.** Enter name for vNIC template.

- Step 6.** Keep Fabric A selected. Do not select the Enable Failover checkbox.
- Step 7.** For Redundancy Type, Select “Primary Template.”
- Step 8.** Select Updating Template as the Template Type.
- Step 9.** Under VLANs, select the checkboxes for desired VLANs to add as part of the vNIC Template.
- Step 10.** Set Native-VLAN as the native VLAN.
- Step 11.** For MTU, enter 9000.
- Step 12.** In the MAC Pool list, select MAC Pool configure for Fabric A.
- Step 13.** In the Network Control Policy list, select CDP_Enabled.
- Step 14.** Click OK to create the vNIC template.

Create vNIC Template [?] [X]

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print [Settings]

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	InBand-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Infra-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Launcher	<input type="radio"/>
<input checked="" type="checkbox"/>	VM-Network	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MACPool-A(128/128)

QoS Policy : <not set>

Network Control Policy : CDP_Enabled

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

[OK] [Cancel]

Step 15. Repeat steps 1-14 to create a vNIC Template for Fabric B. For Peer redundancy Template Select “vNIC-Template-A” created in the previous step.

Create vNIC Template ? X

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template :

Target

Adapter

VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	InBand-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Infra-Mgmt	<input type="radio"/>
<input checked="" type="checkbox"/>	Launcher	<input type="radio"/>

Step 16. Verify that vNIC-Template-A Peer Redundancy Template is set to “vNIC-Template-B.”

Procedure 25. Create vHBA Templates

- Step 1.** In Cisco UCS Manager, click the SAN tab in the navigation pane.
- Step 2.** Click Policies > root > Sub-Organization > FlexPod-CVD > vHBA Template.
- Step 3.** Right-click vHBA Templates.
- Step 4.** Click Create vHBA Template.
- Step 5.** Enter vHBA-A as the vHBA template name.
- Step 6.** Keep Fabric A selected.
- Step 7.** Select VSAN created for Fabric A from the drop-down list.
- Step 8.** Change to Updating Template.
- Step 9.** For Max Data Field keep 2048.
- Step 10.** Select WWPN Pool for Fabric A (created earlier) for our WWPN Pool.
- Step 11.** Leave the remaining fields as-is.
- Step 12.** Click OK.

Create vHBA Template



Name :

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size :

WWPN Pool :

QoS Policy :

Pin Group :

Stats Threshold Policy :

Step 13. Repeat steps 1-12 to create a vHBA Template for Fabric B.

Procedure 26. Create Server Boot Policy for SAN Boot

Note: All Cisco UCS B200 M6 Blade Servers for the workload and the two Infrastructure servers were set to boot from SAN for this Cisco Validated Design as part of the Service Profile template. The benefits of booting from SAN are numerous; disaster recovery, lower cooling, and power requirements for each server since a local drive is not required, and better performance, to name just a few.

Tech tip

We strongly recommend using “Boot from SAN” to realize the full benefits of Cisco UCS stateless computing features, such as service profile mobility.

Note: This process applies to a Cisco UCS environment in which the storage SAN ports are configured as explained in the following section.

Note: A Local disk configuration for the Cisco UCS is necessary if the servers in the environments have a local disk.

Step 1. Go to tab Servers > Policies > root > Sub-Organization > FlexPod-CVD > right-click Local Disk Configuration Policy > Enter “SAN-Boot” as the local disk configuration policy name and change the mode to “No Local Storage.”

Step 2. Click OK to create the policy.

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : Disable Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately. Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : Disable Enable

FlexFlash Removable State : Yes No No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily. Please ensure SD cards are not in use before changing the FlexFlash Removable State.

OK Cancel

Note: As shown in the screenshot below, the NetApp Storage AFF have four active FC connections that pair with the Cisco MDS 9132T 32-Gb switches. Two FC ports are connected to Cisco MDS-A and the other Two FC ports are connected to Cisco MDS-B Switches. All FC ports are 32 Gb/s. The SAN Port 1a of NetApp Storage AFF Controller 1 is connected to Cisco MDS Switch A and SAN port 1b is connected to MDS Switch B. The SAN Port .1a of NetApp Storage AFF Controller 2 is connected to Cisco MDS Switch A and SAN port .1b connected to MDS Switch B.

	Node	1a	1b
▼	A400-01	32 GB/s	32 GB/s
▼	A400-02	32 GB/s	32 GB/s

Procedure 27. Create SAN Policy A

Tech tip

The SAN-A boot policy configures the SAN Primary's primary-target to be port 1a of controller 1 on the NetApp Storage cluster and SAN Primary's secondary-target to be port 1a of controller 2 on the NetApp Storage cluster. Similarly, the SAN Secondary's primary-target should be port 1b of controller 2 on the NetApp Storage cluster and SAN Secondary's secondary-target should be port 1b of controller 1 on the NetApp Storage cluster.

Step 1. Log into the storage controller and verify all the port information is correct. This information can be found in the NetApp Storage GUI under System > Connections > Target Ports.

Note: You have to create a SAN Primary (hba0) and a SAN Secondary (hba1) in SAN-A Boot Policy by entering WWPN of NetApp Storage FC Ports as explained in the following section.

Step 2. Go to Cisco UCS Manager and then go to Servers > Policies > root > Sub Organization > FlexPod-CVD > Boot Policies. Right-click and select Create Boot Policy.


Step 3. Enter SAN-A as the name of the boot policy.



Step 4. Expand the Local Devices drop-down list and click Add CD/DVD.

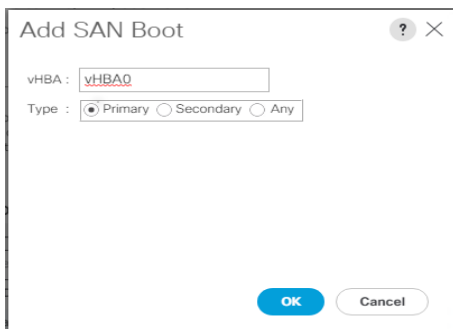


Step 5. Expand the vHBAs drop-down list and click Add SAN Boot.



Note: The SAN boot paths and targets will include primary and secondary options in order to maximize resiliency and number of paths.

Step 6. In the Add SAN Boot dialog box, for Type select “Primary” and name vHBA as “vHBA0”. Click OK to add SAN Boot.



Step 7. Select add SAN Boot Target.

⊖ vHBAs

Add SAN Boot

Add SAN Boot Target

Step 8. Keep **1** as the value for Boot Target LUN. Enter the WWPN for FC port FC0.1a of NetApp Controller 1 and add SAN Boot Primary Target.

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN
CD/DVD	1				
San	2				
SAN Primary		vHBA0	Primary		
SAN Target ...			Primary	0	20:0C:D0:39:EA:18:01:47
SAN Target ...			Secondary	0	20:0E:D0:39:EA:18:01:47
SAN Secondary		vHBA1	Secondary		

Step 9. Add a secondary SAN Boot target into same hba0, enter the boot target LUN as **1** and WWPN for FC port .1a of NetApp Controller 2 , and add SAN Boot Secondary Target.

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN
SAN Primary		vHBA0	Primary		
SAN Target ...			Primary	0	20:0C:D0:39:EA:18:01:47
SAN Target ...			Secondary	0	20:0E:D0:39:EA:18:01:47
SAN Secondary		vHBA1	Secondary		
SAN Target ...			Primary	0	20:0B:D0:39:EA:18:01:47
SAN Target ...			Secondary	0	20:0D:D0:39:EA:18:01:47

Step 10. From the vHBA drop-down list and click Add SAN Boot. In the Add SAN Boot dialog box, enter "vHBA1" in the vHBA field. Click OK to SAN Boot, then click Add SAN Boot Target.

Add SAN Boot ? X

vHBA :

Type : Primary Secondary Any

Step 11. Keep **1** as the value for the Boot Target LUN. Enter the WWPN for FC port .1b of NetApp Controller 2 and add SAN Boot Primary Target.

Add SAN Boot Target
?
✕

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK
Cancel

Step 12. Add a secondary SAN Boot target into same vhb1 and enter the boot target LUN as 1 and WWPN for FC port 1b of NetApp Controller 1 and add SAN Boot Secondary Target.

Add SAN Boot Target
?
✕

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK
Cancel

Step 13. Click Save Changes.

- Sub-Organizations
- FlexPod-CVD
 - Adapter Policies
 - BIOS Policies
 - Boot Policies
 - Boot Policy A300Boot
 - Boot Policy A400
 - Boot Policy a400-BFS**
 - Boot Policy A400Boot
 - Boot Policy Converged
 - Diagnostics Policies
 - Graphics Card Policies
 - Host Firmware Packages
 - IPMI/Redfish Access Profiles
 - KVM Management Policies
 - Local Disk Config Policies
 - Maintenance Policies
 - FlexPodMaint
 - Management Firmware Packages
 - Persistent Memory Policy
 - Power Control Policies
 - Power Sync Policies
 - SPDM Certificate Policies
 - Scrub Policies

Policies / root / Sub-Organizations / FlexPod-CVD / Boot Policies / Boot Policy a400-BFS

General

Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **a400-BFS**

Description :

Owner : **Local**

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Warning

The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/i...	Type	LUN Name	WWN
CD/DVD	1				
San	2				
SAN Primary		vHBA0	Primary		
SAN Target Pri...			Primary	0	20:0B:D0:39:EA:18:01:47
SAN Target S...			Secondary	0	20:00:D0:39:EA:18:01:47
SAN Secondary		vHBA1	Secondary		

Step 14. After creating the FC boot policy, you can view the boot order in the Cisco UCS Manager GUI. To view the boot order, navigate to Servers > Policies > Boot Policies. Click Boot Policy SAN-Boot-A to view the boot order in the right pane of the Cisco UCS Manager as shown below:

© 2022 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

Page 141 of 331

Boot Policies		Events				
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN	Slot Number
▶ Boot Policy A300Boot						
▶ Boot Policy A400						
▼ Boot Policy a400-BFS						
CD/DVD	1					
▼ San	2					
▼ SAN Primary		vHBA0	Primary			
SAN Targe...			Primary	0	20:0B:D0:39:EA:18:01:...	
SAN Targe...			Secondary	0	20:0D:D0:39:EA:18:01:...	
▼ SAN Secondary		vHBA1	Secondary			
SAN Targe...			Primary	0	20:0C:D0:39:EA:18:01:...	
SAN Targe...			Secondary	0	20:0E:D0:39:EA:18:01:47	

Procedure 28. Create SAN Policy B

Tech tip

The SAN-B boot policy configures the SAN Primary's primary-target to be port 1b of Controller 1 on the NetApp Storage cluster and SAN Primary's secondary-target to be port 1b Controller 2 on the NetApp Storage cluster. Similarly, the SAN Secondary's primary-target should be port 1a of Controller 2 on the NetApp Storage cluster and SAN Secondary's secondary-target should be port 1a of Controller 1 on the NetApp Storage cluster.

Step 1. Log into the storage controller and verify all the port information is correct. This information can be found in the NetApp Storage GUI under System > Connections > Target Ports.

Note: You have to create SAN Primary (vHBA1) and SAN Secondary (vHBA0) in SAN-B Boot Policy by entering WWPN of NetApp Storage FC Ports as explained in the following section.

Step 2. Go to Cisco UCS Manager and then go to tab Servers > Policies > root > Sub Organization > FlexPod-CVD > Boot Policies.

Step 3. Right-click and select Create Boot Policy. Enter SAN-B as the name of the boot policy.

Create Boot Policy ? X

Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
No data available									

Step 4. Expand the Local Devices drop-down list and Click Add CD/DVD. Expand the vHBAs drop-down list and click Add SAN Boot.

Note: The SAN boot paths and targets include primary and secondary options in order to maximize resiliency and number of paths.

Step 5. In the Add SAN Boot dialog box, for Type select “Primary” and name vHBA as “vHBA0.” Click OK to add SAN Boot.

Add SAN Boot ? X

vHBA :

Type : Primary Secondary Any

Step 6. Select Add SAN Boot Target to enter WWPN address of storage port. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port 1b of NetApp Controller 1 and add SAN Boot Primary Target.

Step 7. Add the secondary SAN Boot target into the same hba0; enter boot target LUN as 1 and WWPN for FC port 1bof NetApp Controller 2 and add SAN Boot Secondary Target.

Add SAN Boot Target ? X

Boot Target LUN : 1

Boot Target WWPN : 20:0C:D0:39:EA:18:01:47

Type : Primary Secondary

OK Cancel

Step 8. From the vHBA drop-down list, click Add SAN Boot. In the Add SAN Boot dialog box, enter " hba1" in the vHBA field. Click OK to SAN Boot, then click Add SAN Boot Target.

Add SAN Boot ? X

vHBA : vHBA0

Type : Primary Secondary Any

OK Cancel

Step 9. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port 1a of NetApp Controller 2 and Add SAN Boot Primary Target.

Add SAN Boot Target ? X

Boot Target LUN : 1

Boot Target WWPN : 20:0C:D0:39:EA:18:01:47

Type : Primary Secondary

OK Cancel

Step 10. Add secondary SAN Boot target into same hba1 and enter boot target LUN as 1 and WWPN for FC port 1a of NetApp Controller 1 and add SAN Boot Secondary Target.

Step 11. Click OK.

Create Boot Policy

Name : SAN-B

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

+ Local Devices

+ CIMC Mounted vMedia

+ vNICs

- vHBAs

Add SAN Boot

Add SAN Boot Target

+ iSCSI vNICs

+ EFI Shell

Boot Order									
Name	Order	vNIC/...	Type	LUN ...	WWN	Slot N...	Boot ...	Boot ...	Descri...
CD/DVD	1								
San	2								

Move Up Move Down Delete

Set User Boot Parameters

OK Cancel

Step 12. After creating the FC boot policies, you can view the boot order in the Cisco UCS Manager GUI. To view the boot order, navigate to Servers > Policies > Boot Policies. Click Boot Policy SAN-Boot-B to view the boot order in the right pane of the Cisco UCS Manager as shown below:

Policies / root / Sub-Organizations / FlexPod-CVD / Boot Policies

Boot Policies Events

Name	Order	vNIC/vHBA/iSCSI v...	Type	LUN Name	WWN
▶ Boot Policy A300Boot					
▶ Boot Policy A400					
▼ Boot Policy a400-BFS					
CD/DVD	1				
▼ San	2				
▼ SAN Primary		vHBA0	Primary		
SAN Target ...			Primary	0	20:0B:D0:39:EA:18:01:47
SAN Target ...			Secondary	0	20:0D:D0:39:EA:18:01:47
▼ SAN Secondary		vHBA1	Secondary		
SAN Target ...			Primary	0	20:0C:D0:39:EA:18:01:47
SAN Target ...			Secondary	0	20:0E:D0:39:EA:18:01:47

Note: For this solution, we created two Boot Policy as “SAN-A” and “SAN-B.” For 32 Cisco UCS B200 M6 blade servers, you will assign the first 16 Service Profiles with SAN-A to the first 16 servers and the remaining 16 Service Profiles with SAN-B to the remaining 16 servers as explained in the following section.

Procedure 29. Create and Configure a Service Profile Template

Tech tip

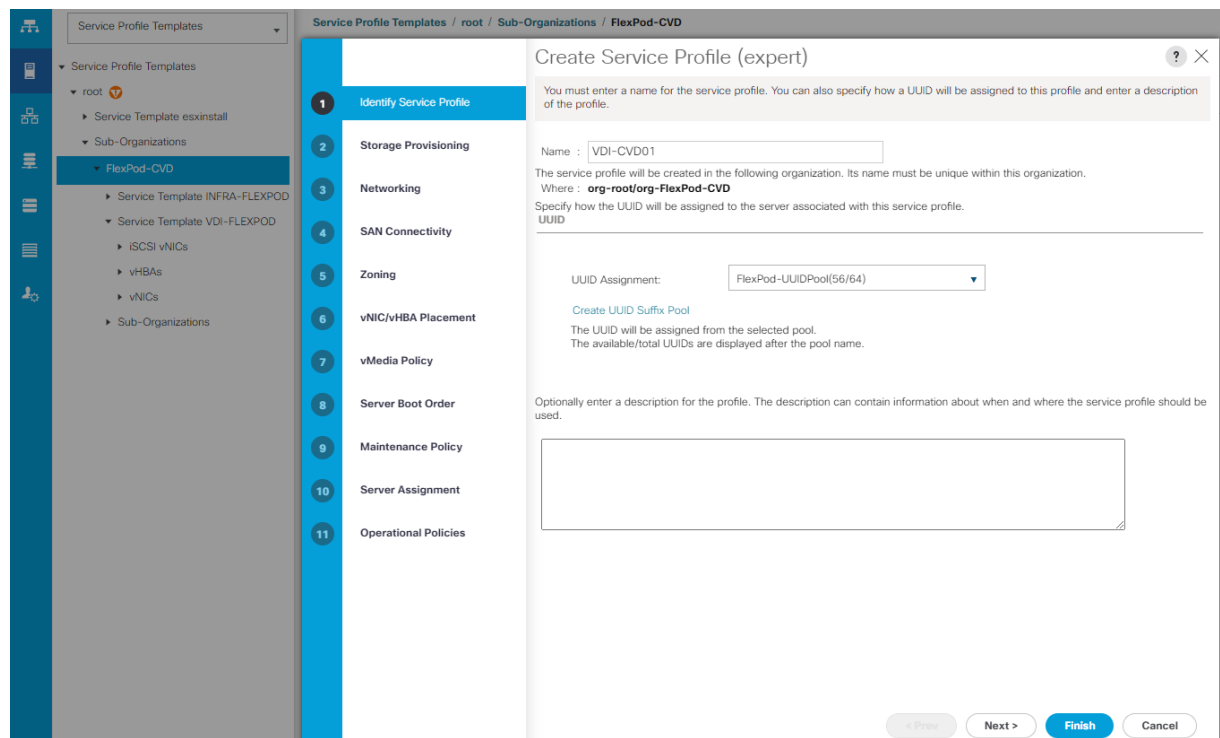
Service profile templates enable policy-based server management that helps ensure consistent server resource provisioning suitable to meet predefined workload needs.

Note: You will create two Service Profile templates; the first Service profile template “VDI-CVD01” uses the boot policy “SAN-A” and the second Service profile template “VDI-CVD02” uses the boot policy “SAN-B” to utilize all the FC ports from NetApp Storage for high-availability in case any FC links go down.

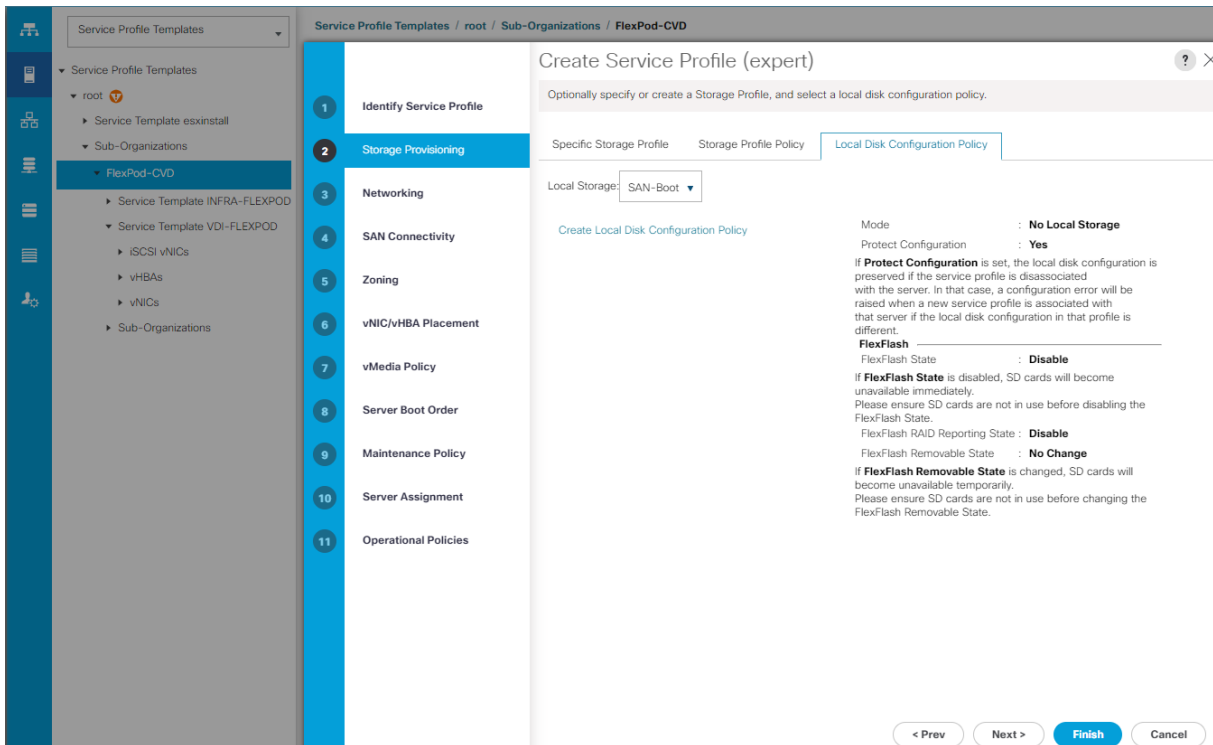
Note: You will create the first VDI-CVD01 as explained in the following section.

Step 1. In the Cisco UCS Manager, go to Servers > Service Profile Templates > root Sub Organization > FlexPod-CVD > and right-click Create Service Profile Template.

Step 2. Enter the Service Profile Template name, select the UUID Pool that was previously created, and click Next.



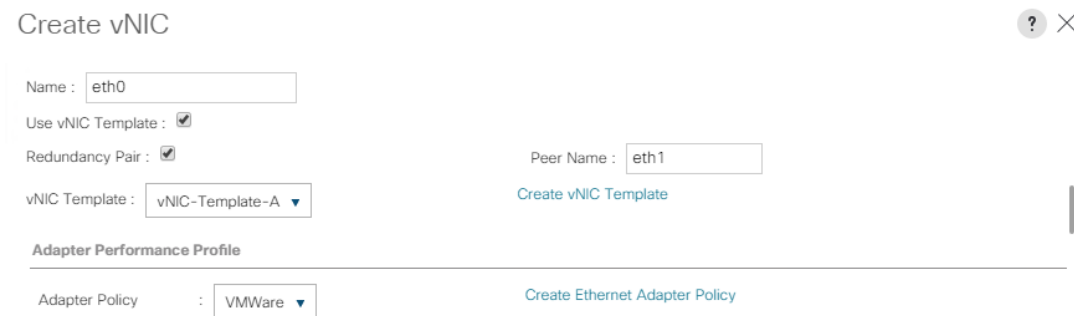
Step 3. Select Local Disk Configuration Policy to SAN-Boot as No Local Storage.



Step 4. In the networking window, select Expert and click Add to create vNICs. Add one or more vNICs that the server should use to connect to the LAN.

Note: Now there are two vNICs in the create vNIC menu; you provided a name for the first vNIC as “eth0” and the second vNIC as “eth1.”

Step 5. Select vNIC-Template-A for the vNIC Template and select VMware for the Adapter Policy as shown below.



Step 6. Select vNIC-Template-B for the vNIC Template, created with the name eth1. Select VMware for the vNIC “eth1” for the Adapter Policy.

Note: eth0 and eth1 vNICs are created so that the servers can connect to the LAN.

Step 7. When the vNICs are created, you need to create vHBAs. Click Next.

Step 8. In the SAN Connectivity menu, select Expert to configure as SAN connectivity. Select WWNN (World Wide Node Name) pool, which you previously created. Click Add to add vHBAs.

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?
 Simple Expert No vHBAs Use Connectivity Policy

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.
 World Wide Node Name

WWNN Assignment:

The WWNN will be assigned from the selected pool.
 The available/total WWNNs are displayed after the pool name.

Name	WWPN
No data available	

The following four HBAs were created:

- vHBA0 using vHBA Template vHBA-A
- vHBA1 using vHBA Template vHBA-B
- vHBA2 using vHBA Template vHBA-A
- vHBA3 using vHBA Template vHBA-B

Figure 47. vHBA0

Create vHBA

Name :

Use vHBA Template :

Redundancy Pair : Peer Name :

vHBA Template : [Create vHBA Template](#)

Adapter Performance Profile

Adapter Policy : [Create Fibre Channel Adapter Policy](#)

Figure 48. vHBA1
Modify vHBA

Name : **vHBA1**

Use vHBA Template :

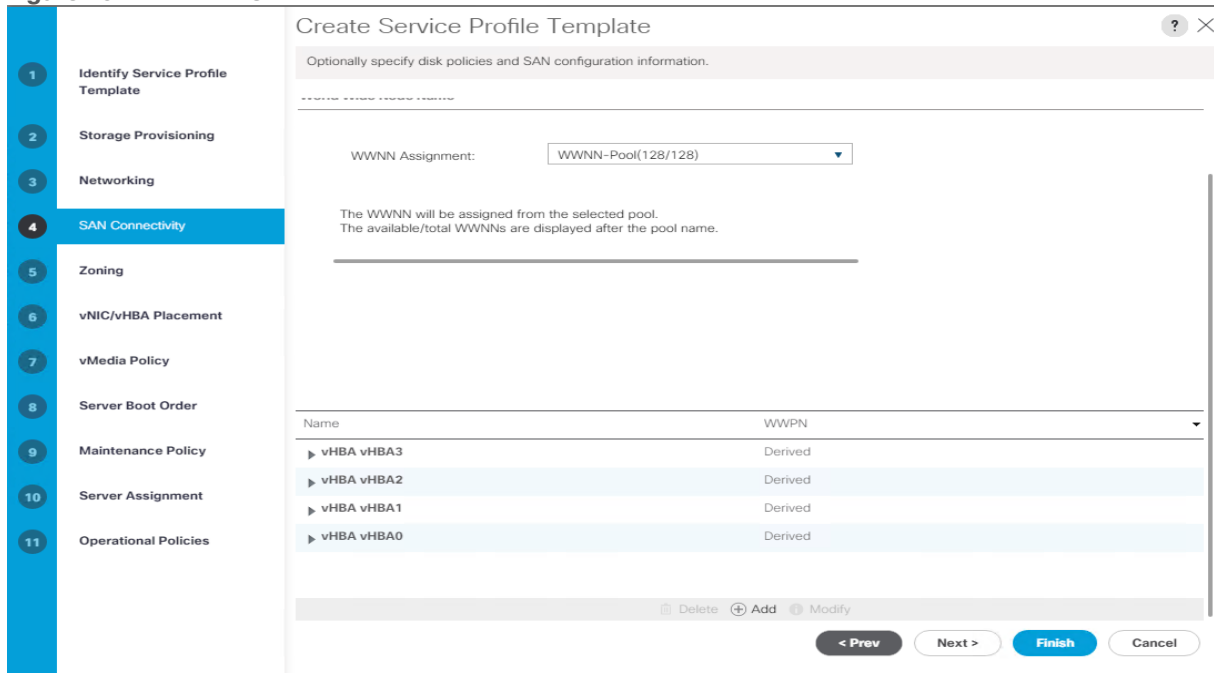
[Create vHBA Template](#)

vHBA Template :

Adapter Performance Profile

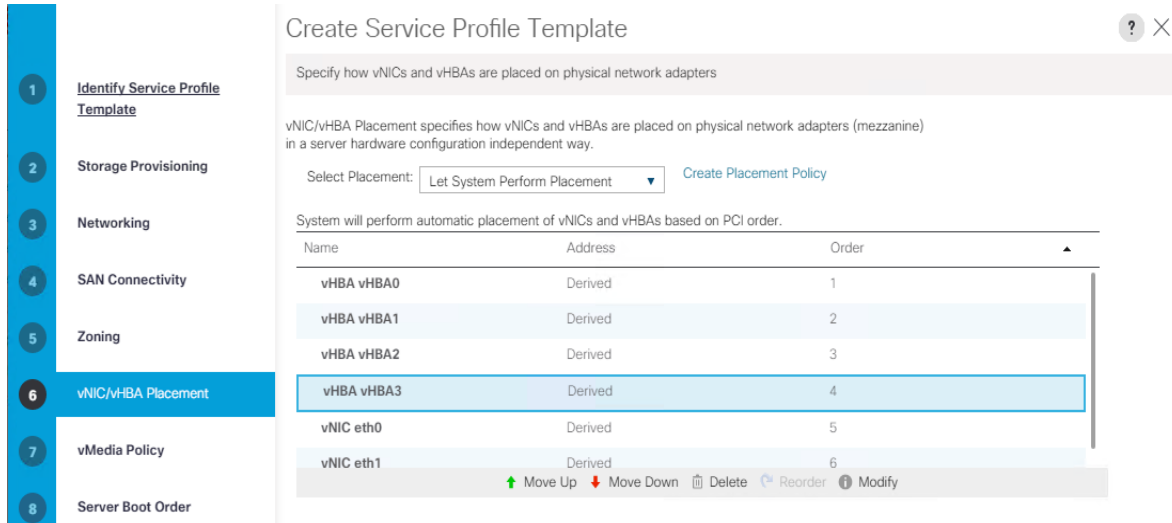
Adapter Policy : [Create Fibre Channel Adapter Policy](#)

Figure 49. All vHBAs



Step 9. Skip zoning. For this FlexPod Configuration, the Cisco MDS 9132T 32-Gb is used for zoning.

Step 10. Select the default option Let System Perform Placement in the Placement Selection menu.



Step 11. For the Server Boot Policy, select SAN-A, which you previously created.

Create Service Profile Template ? X

Optionally specify the boot policy for this service profile template.

Select a boot policy.
 Boot Policy: [Create Boot Policy](#)

Name : **SAN-A**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Legacy**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vH...	Type	WWN	LUN Name	Slot Num...	Boot Name	Boot Path	Description
San	2								
▶ SAN Primary		vHBA0	Primary						
▶ SAN Second...		vHBA1	Secondary						
Remote CD/DVD	1								

[Create iSCSI vNIC](#) [Set iSCSI Boot Parameters](#) [Set UEFI Boot Parameters](#)

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

Create Service Profile (expert) ? X

Optionally specify the boot policy for this service profile.

Select a boot policy.
 Boot Policy: [Create Boot Policy](#)

Name : **A400Boot**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Legacy**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	vNIC/v...	Type	WWN
CD/DVD			
San			
▶ SAN Primary	hba0	Primary	
SAN Targe...		Primary	20:06:D0:39:EA:18:01:47
SAN Targe...		Secondary	20:08:D0:39:EA:18:01:47
▶ SAN Secondary	hba1	Secondary	

[Create iSCSI vNIC](#) [Set iSCSI Boot Parameters](#) [Set UEFI Boot Parameters](#)

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

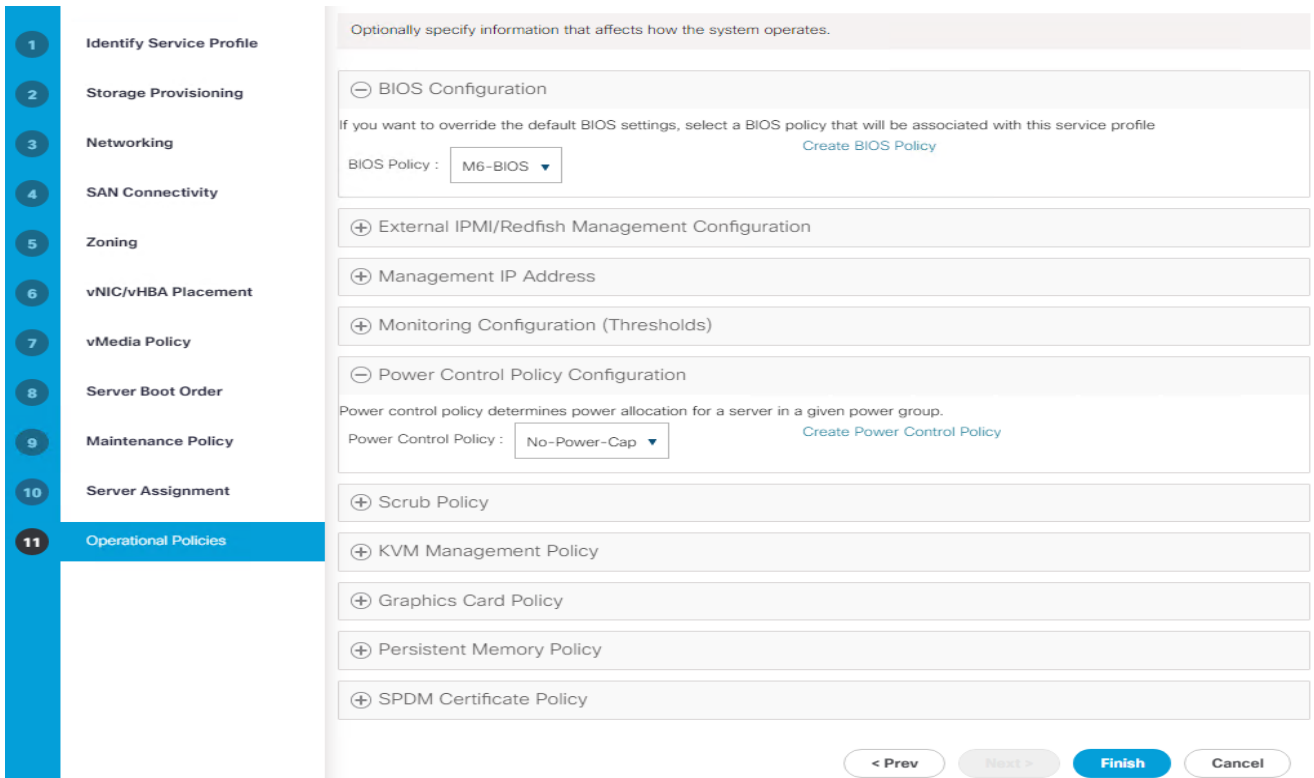
Note: The default setting was retained for the remaining maintenance and assignment policies in the configuration. However, they may vary from site-to-site depending on workloads, best practices, and policies. For example, we created a maintenance policy, BIOS policy, Power Policy, as detailed below.

Step 12. Select UserAck maintenance policy, which requires user acknowledgement prior rebooting server when making changes to policy or pool configuration tied to a service profile.

Step 13. Select Server Pool policy to automatically assign service profile to a server that meets the requirement for server qualification based on the pool configuration.

Step 14. On the same page you can configure “Host firmware Package Policy” which helps to keep the firmware in sync when associated to server.

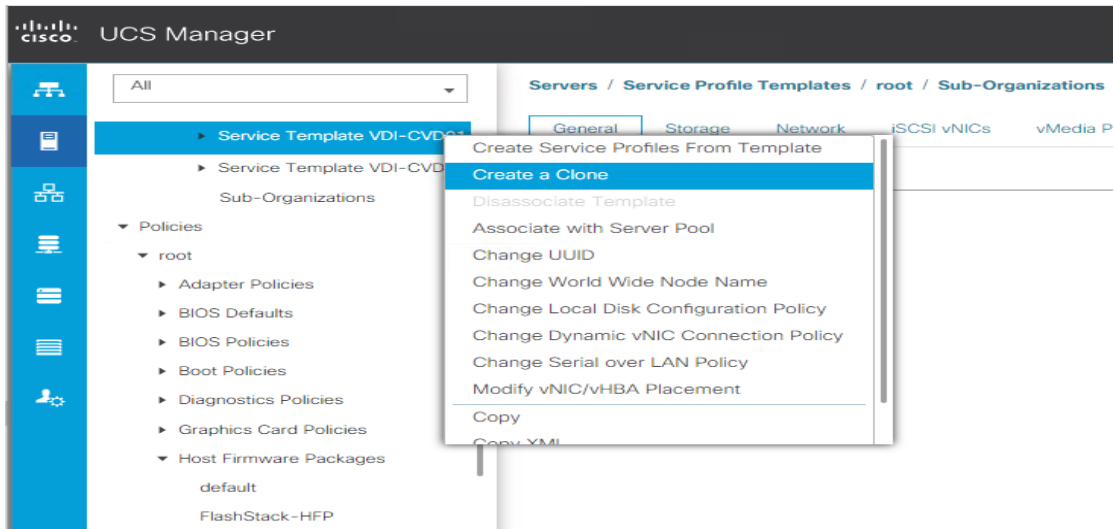
Note: On the Operational Policy page, we configured BIOS policy for B200 M6 blade server, Power Control Policy with “NoPowerCap” for maximum performance and Graphics Card Policy for Cisco UCS B200 M6 blade server configured with Nvidia P6 GPU card.



Step 15. Click Next and then click Finish to create service profile template as “VDI-CVD01.”

Procedure 30. Clone Service Profile Template

Step 1. In the Cisco UCS Manager, go to Servers > Service Profile Templates > root > Sub Organization > FlexPod-CVD > Service Template VDI-CVD01 and right-click Create a Clone as shown below.



Step 2. Enter name to create Clone from existing Service Profile template. Click OK.

Create Clone From VDI-FLEXPOD

Clone Name
: VDI-FP-01

Org
: FlexPod-CVD

OK Cancel Help

Note: This VDI-CVD02 service profile template will be used to create the remaining sixteen service profiles for VDI workload and Infrastructure server02.

Step 3. To change boot order from SAN-A to SAN-B for VDI-CVD02, click Cloned Service Profile template > Select Boot Order tab. Click Modify Boot Policy.

Service Profile Templates / root / Sub-Organizations / FlexPod-CVD / Service Template VDI-FLEXPOD

General Storage Network iSCSI vNICs vMedia Policy **Boot Order** Policies Events FSM

Actions

Modify Boot Policy

Global Boot Policy

Name : **Converged**

Description :

Reboot on Boot Order Change : **No**

Enforce vNIC/vHBA/iSCSI Name : **Yes**

Boot Mode : **Legacy**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/iSCSI...	Type	LUN Name	WWN	Slot Number	Boot Name
CD/DVD	1						
San	2						

Step 4. From the drop-down list, for the Boot Policy, select SAN-B and click OK.

Modify Boot Policy

Boot Policy: SAN-A

Name : **SAN-A**

Description :

Reboot on Boot Order Change : **No**

Enforce vNIC/vHBA/iSCSI Name : **Yes**

Boot Mode : **Legacy**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Note: You have now created the Service Profile template “VDI-CVD01” and “VDI-CVD02” with each having four vHBAs and two vNICs.

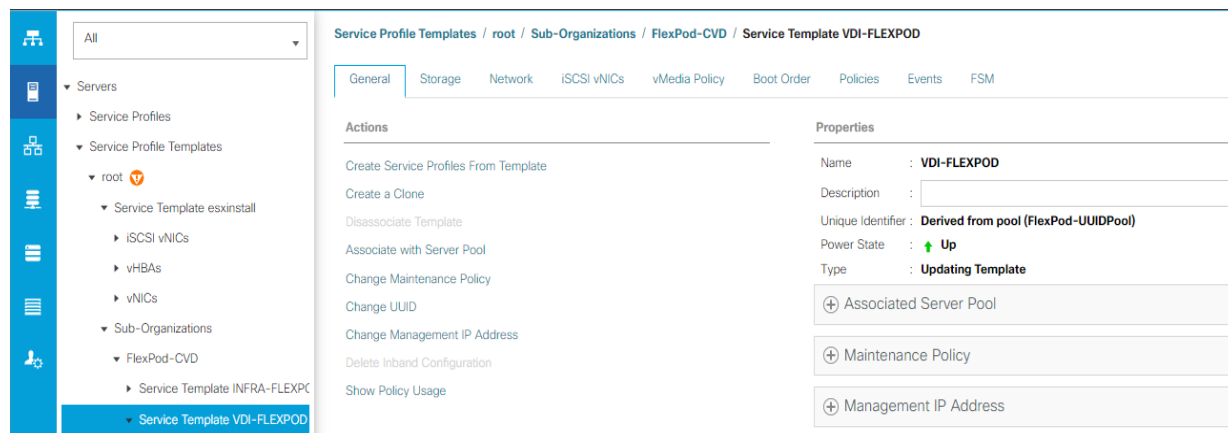
Create Service Profiles from Template and Associate to Servers

Note: You will create 16 service profiles from the VDI-CVD01 template and 16 service profiles from the VDI-CVD02 template as explained in the following sections.

Note: For the first 15 workload nodes and infrastructure node 01, you will create 16 service profiles from the template VDI-CVD01. The remaining 15 workload nodes and infrastructure node 02, will require creating another 16 service profiles from the template VDI-CVD02.”

Procedure 1. Create First Four Service Profiles from Template

Step 1. Go to the Servers tab > Service Profiles > root > Sub-Organization > FlexPod-CVD and right-click Create Service Profiles from Template.



Step 2. Select “VDI-CVD01” for the Service profile template which you created earlier and name the service profile “VDI-HostX.” To create four service profiles, enter 16 for the Number of Instances, as 16 as shown below. This process will create service profiles “VDI-HOST1”, “VDI-HOST2”, and “VDI-HOST16.”

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

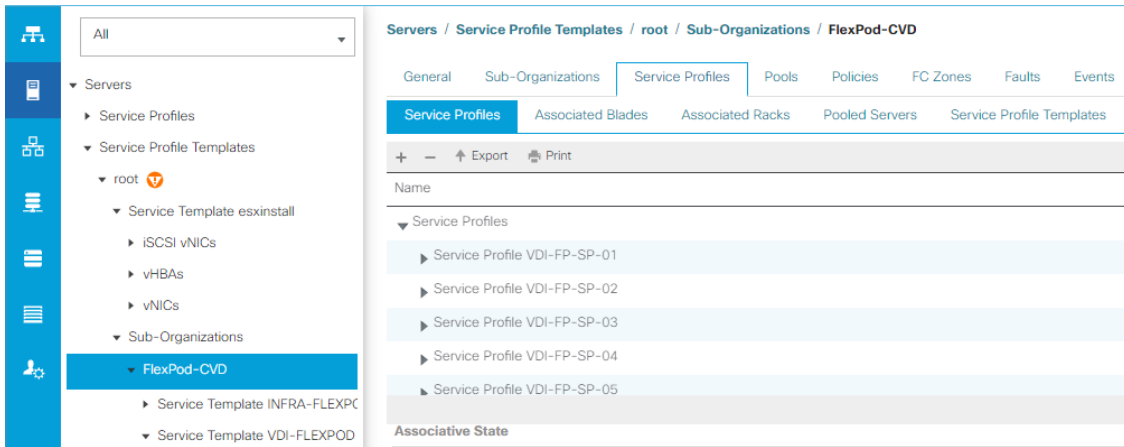
Number of Instances :

Step 3. Create the remaining four Service Profiles “VDI-HOST17”, “VDI-HOST18”, and “VDI-HOST32” from Template “VDI-CVD02.”

Note: When the service profiles are created, the association of Service Profile starts automatically to servers based on the Server Pool Policies.

Step 4. Rename the Service Profiles on Chassis 3/8 as VDI-Infra01 and Service Profile on Chassis 4/8 as VDI-Infra02. Rename rest as necessary to have VDI-Host1 to VDI-Host30.

Step 5. Service Profile association can be verified in Cisco UCS Manager > Servers > Service Profiles. Different tabs can provide details on Service profile association based on Server Pools Policy, Service Profile Template to which Service Profile is tied to, and so on.



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-SVM-fcp-lif-01a pwn <fcp-lif-01a-wwpn>
device-alias name Infra-SVM-fcp-lif-02a pwn <fcp-lif-02a-wwpn>
device-alias name VM-Host-Infra-01-A pwn <vm-host-infra-01-wwpna>
device-alias name VM-Host-Infra-02-A pwn <vm-host-infra-02-wwpna>
device-alias name VM-Host-Infra-03-A pwn <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-SVM-fcp-lif-01b pwnn <fcp-lif-01b-wwpn>
device-alias name Infra-SVM-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-SVM-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-SVM-fcp-lif-01a target
member device-alias Infra-SVM-fcp-lif-02a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member Infra-SVM-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-SVM 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-SVM-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-SVM-fcp-lif-01b target
```

```
member device-alias Infra-SVM-fcp-lif-02b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member Infra-SVM-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:

- [ONTAP Boot Storage Setup](#)
- [Install VMware ESXi 7.0](#)
- [VMware vCenter 7.0](#)

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-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -size 32GB -ostype vmware -space-reserve disabled
```

```
lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-02 -size 32GB -ostype vmware -space-reserve disabled
```

```
lun create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-03 -size 32GB -ostype vmware -space-reserve disabled
```

Procedure 2. Create igroups

Step 2. Create initiator groups (igroups) by entering the following commands from the storage cluster management node Secure Shell (SSH) connection:

```
lun igroup create -vserver Infra-SVM -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-SVM -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-SVM -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-SVM -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 3. 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-SVM -path /vol/esxi_boot/VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-id 0
```

```
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-02 -igroup VM-Host-Infra-02 -lun-id 0
```

```
lun mapping create -vserver Infra-SVM -path /vol/esxi_boot/VM-Host-Infra-03 -igroup VM-Host-Infra-03 -lun-id 0
```

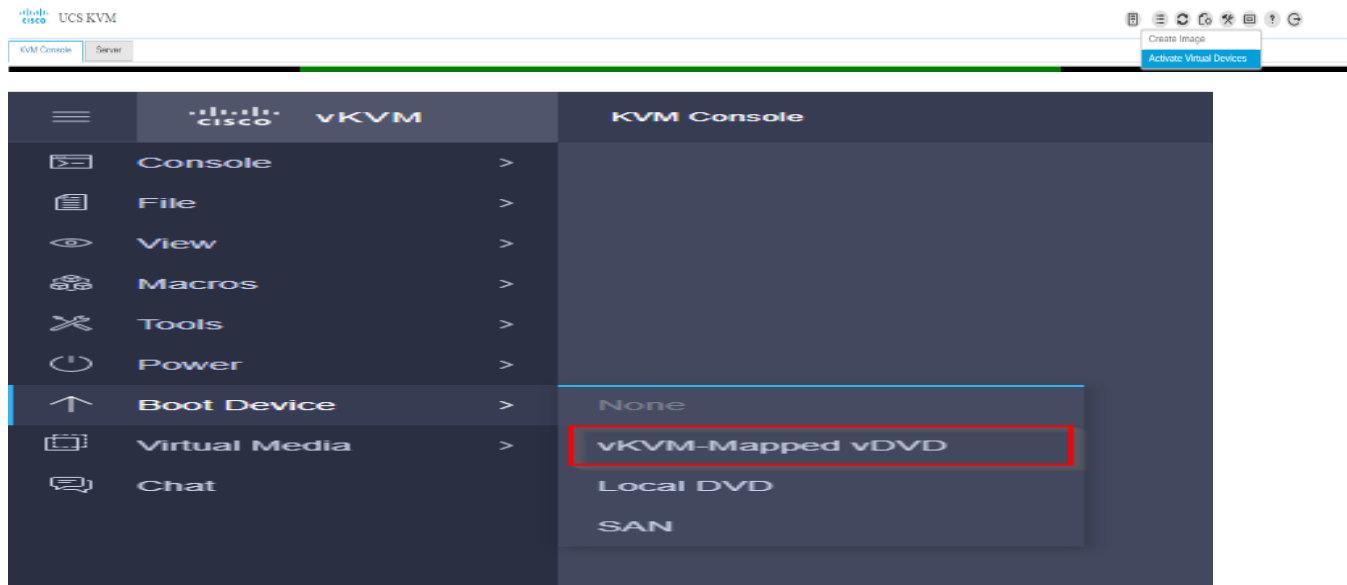
Step 2. Download the Cisco Custom Image for VMware ESXi 7.0 Update 2a, from the [VMware vSphere Hypervisor 7.0 U2](#) page click the “Custom ISOs” tab.

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.

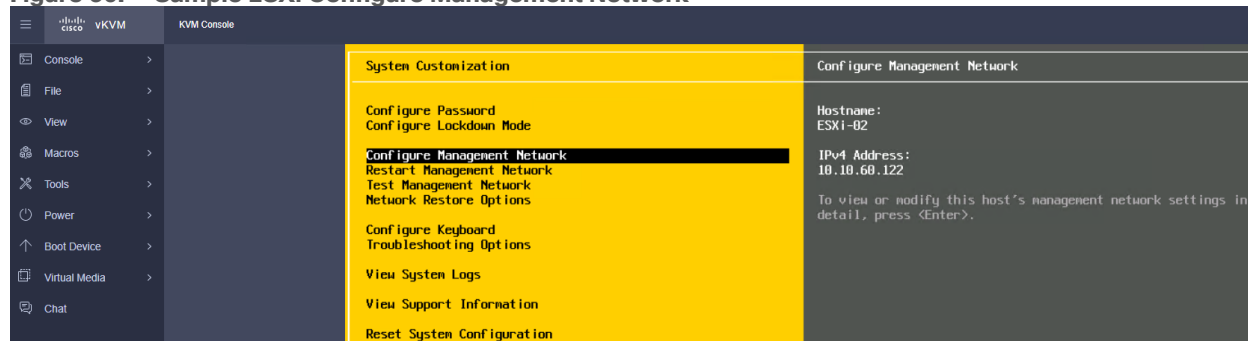
Step 18. Enter the IP address of the primary and secondary DNS server. Enter Hostname

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 50. 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 the 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-ESXI70U2-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 user name 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.U2, 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_1.1.5-10EM.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:

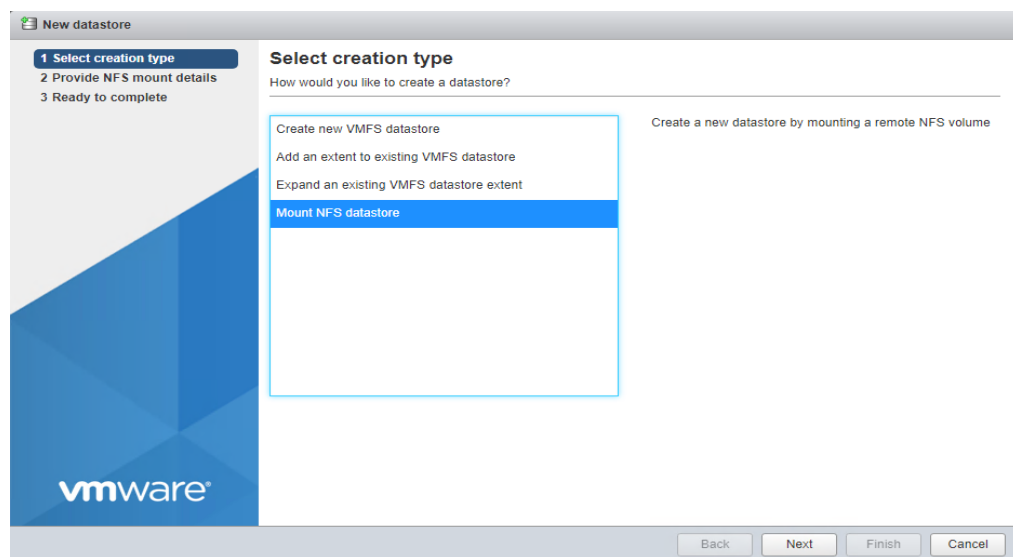
Step 35. Click Networking and the VMkernel NICs tab to confirm configured virtual adapters. The adapters listed should be similar to the following example:

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

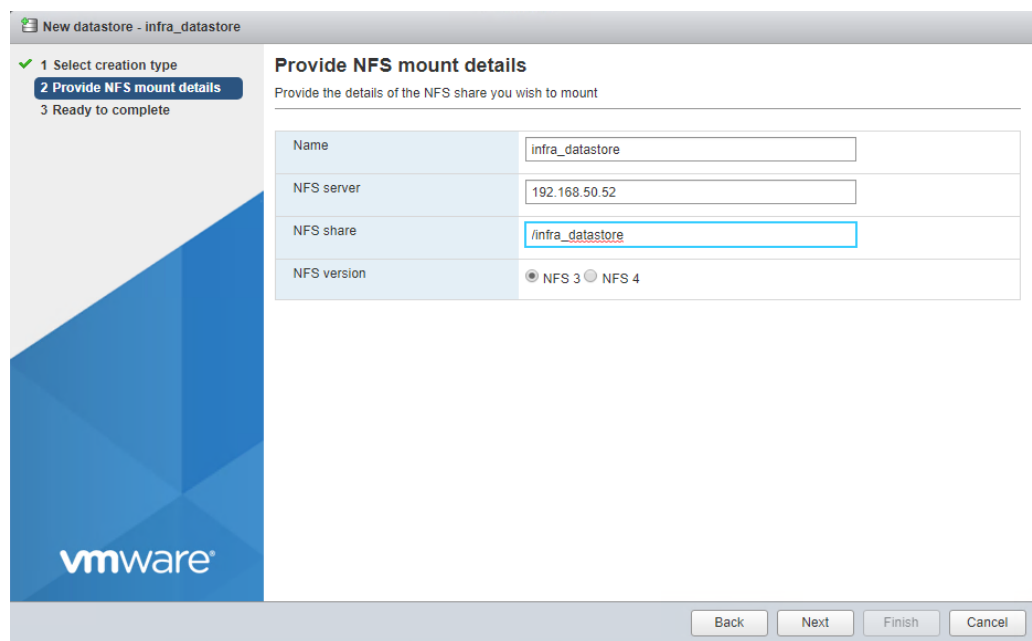
Procedure 11. Mount Required Datastores on ESXi Host VM-Host-Infra-01

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

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

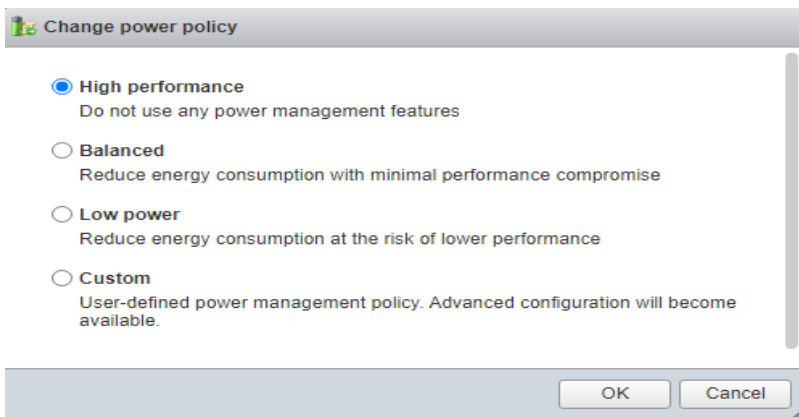
Note: Implementing this policy is recommended in [Performance Tuning Guide for Cisco UCS M5 Servers](#) for maximum VMware ESXi performance. If your organization has specific power policies, please set this policy accordingly.

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.U2-7867351.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 7.0.U2 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 vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appears.

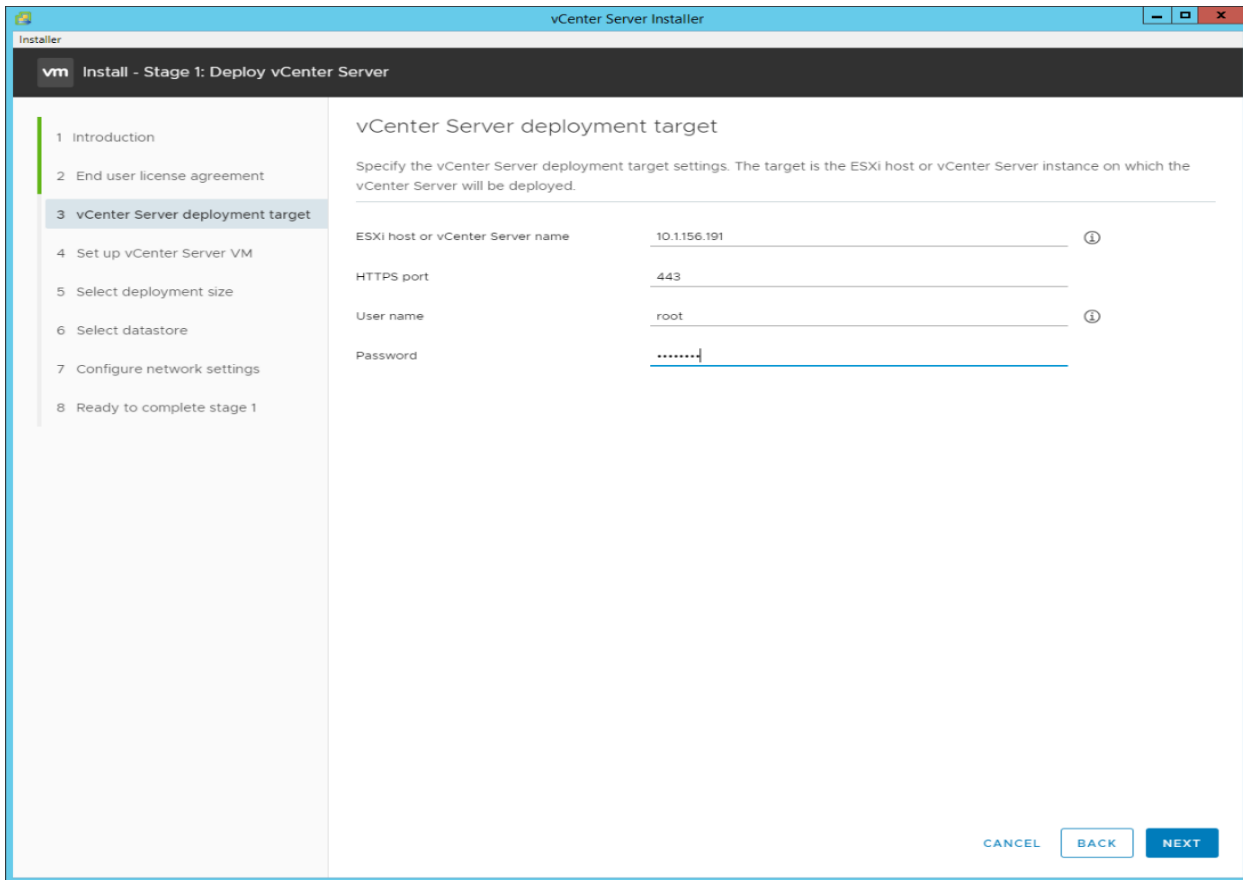


Step 4. Click Install to start the vCenter Server Appliance deployment wizard.

Step 5. Click NEXT in the Introduction section.

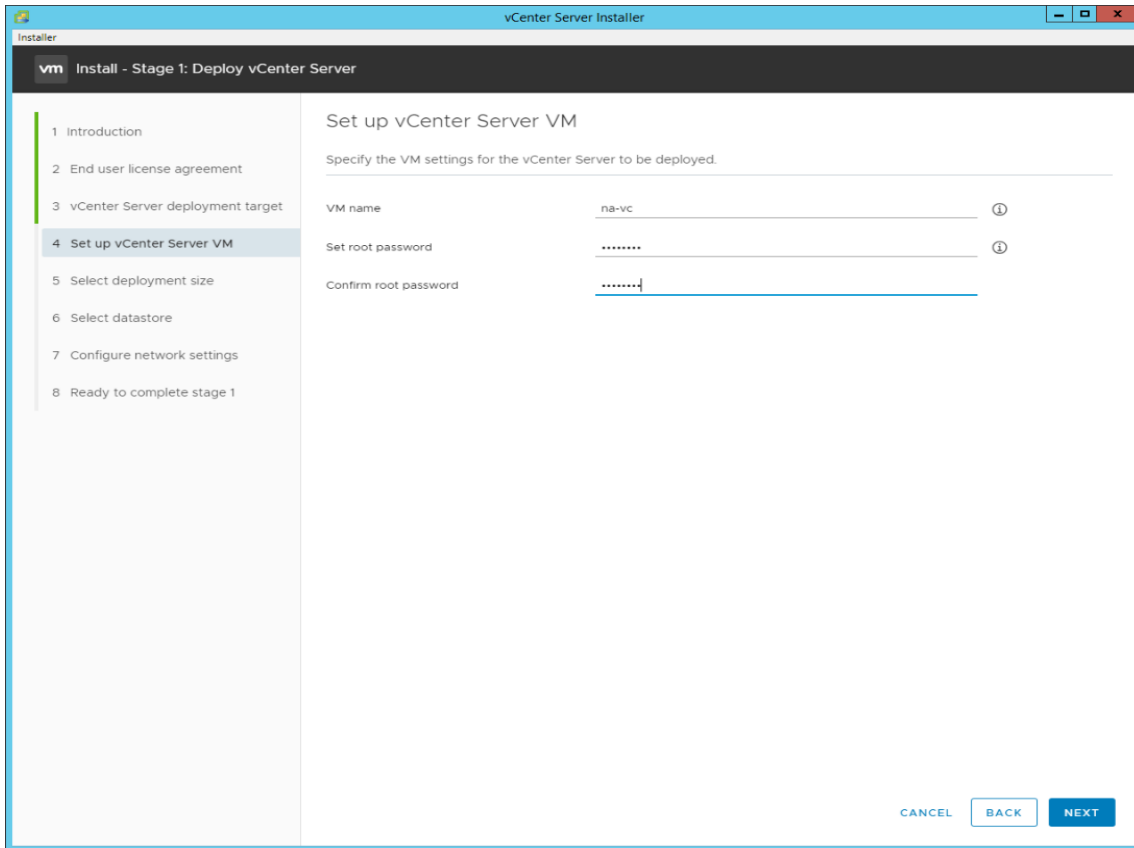
Step 6. Read and accept the license agreement and click NEXT.

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.

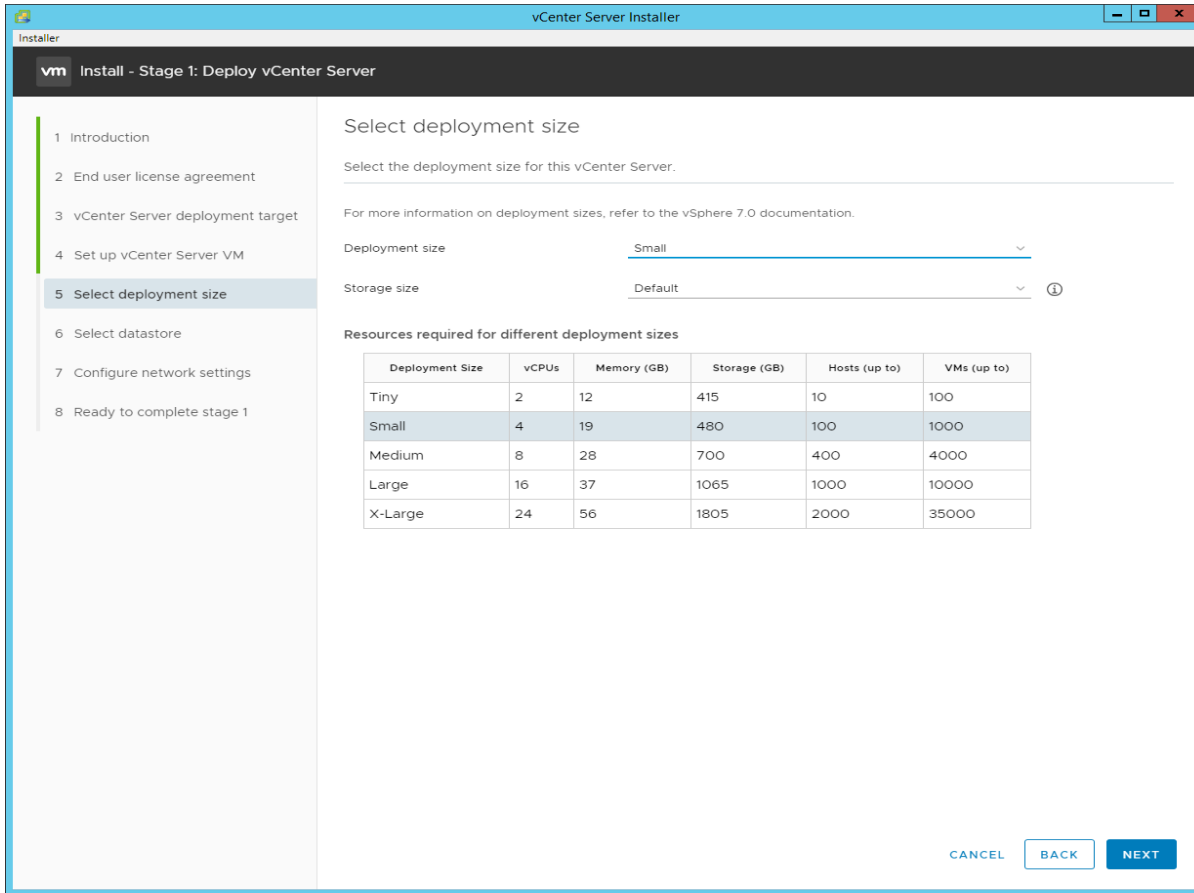


Step 8. Click YES to accept the certificate.

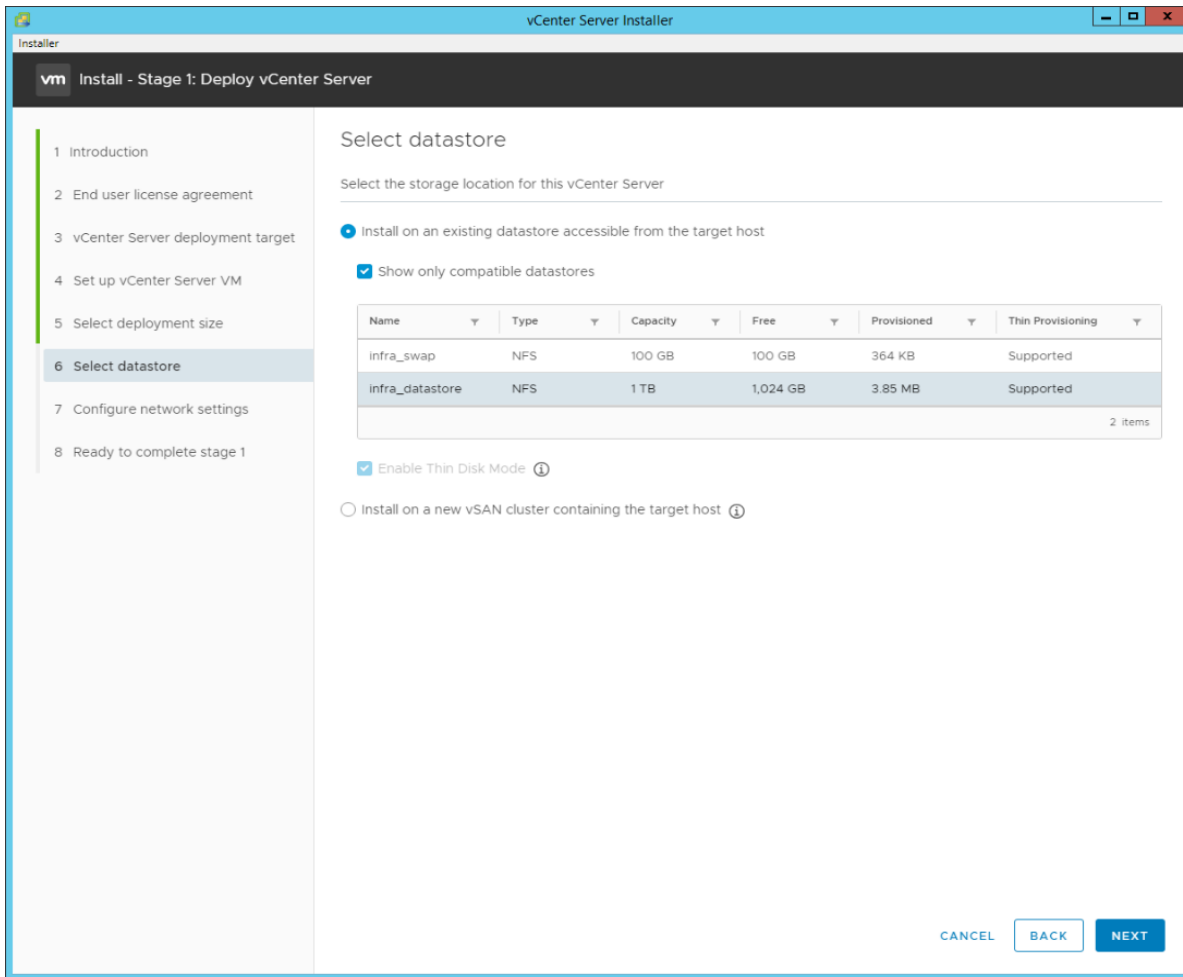
Step 9. Enter the Appliance VM name and password details in the “Set up vCenter Server VM” section. Click NEXT.



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**.

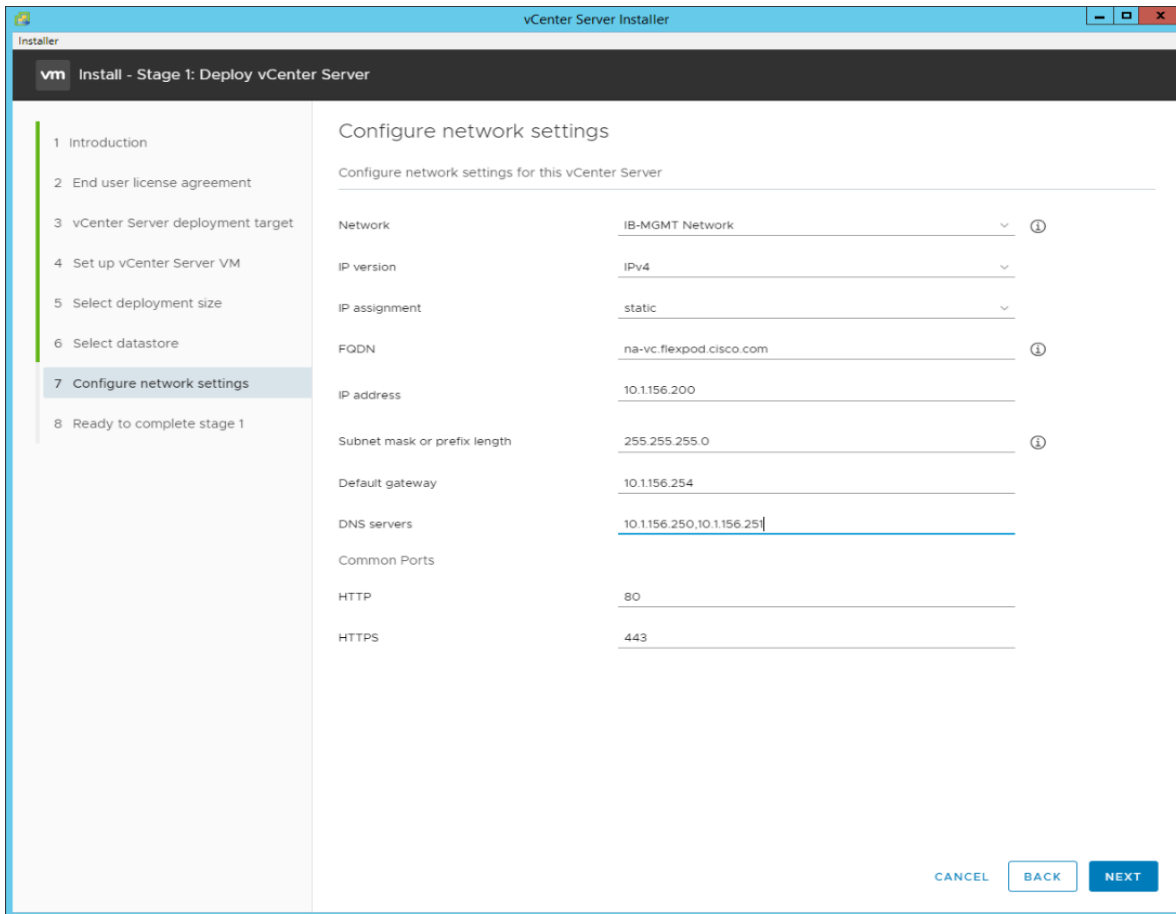


Step 12. In the “Network Settings” section, configure the following settings:

- 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.

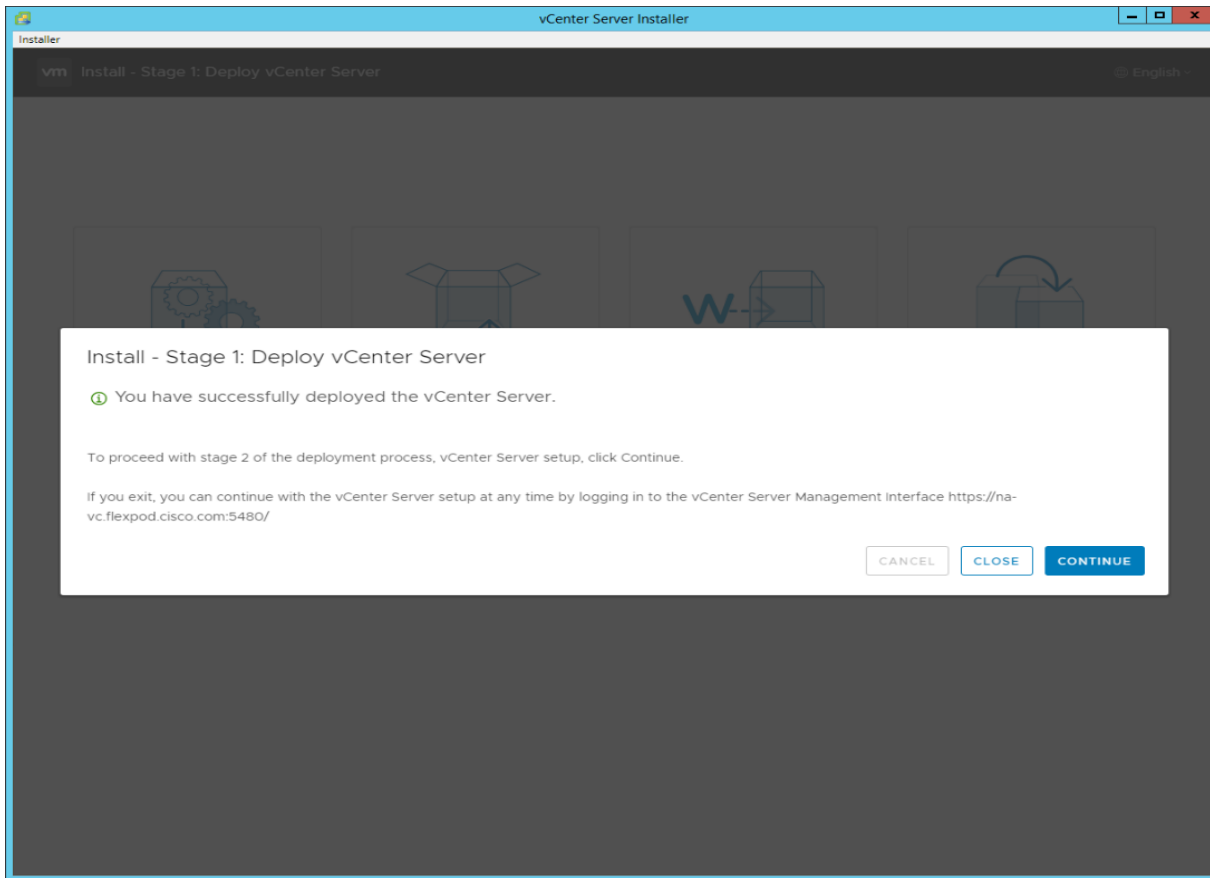
- b. IP version: IPV4
- c. IP assignment: static
- d. FQDN: <vcenter-fqdn>
- e. IP address: <vcenter-ip>
- f. Subnet mask or prefix length: <vcenter-subnet-mask>
- g. Default gateway: <vcenter-gateway>
- h. DNS Servers: <dns-server1>,<dns-server2>



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.

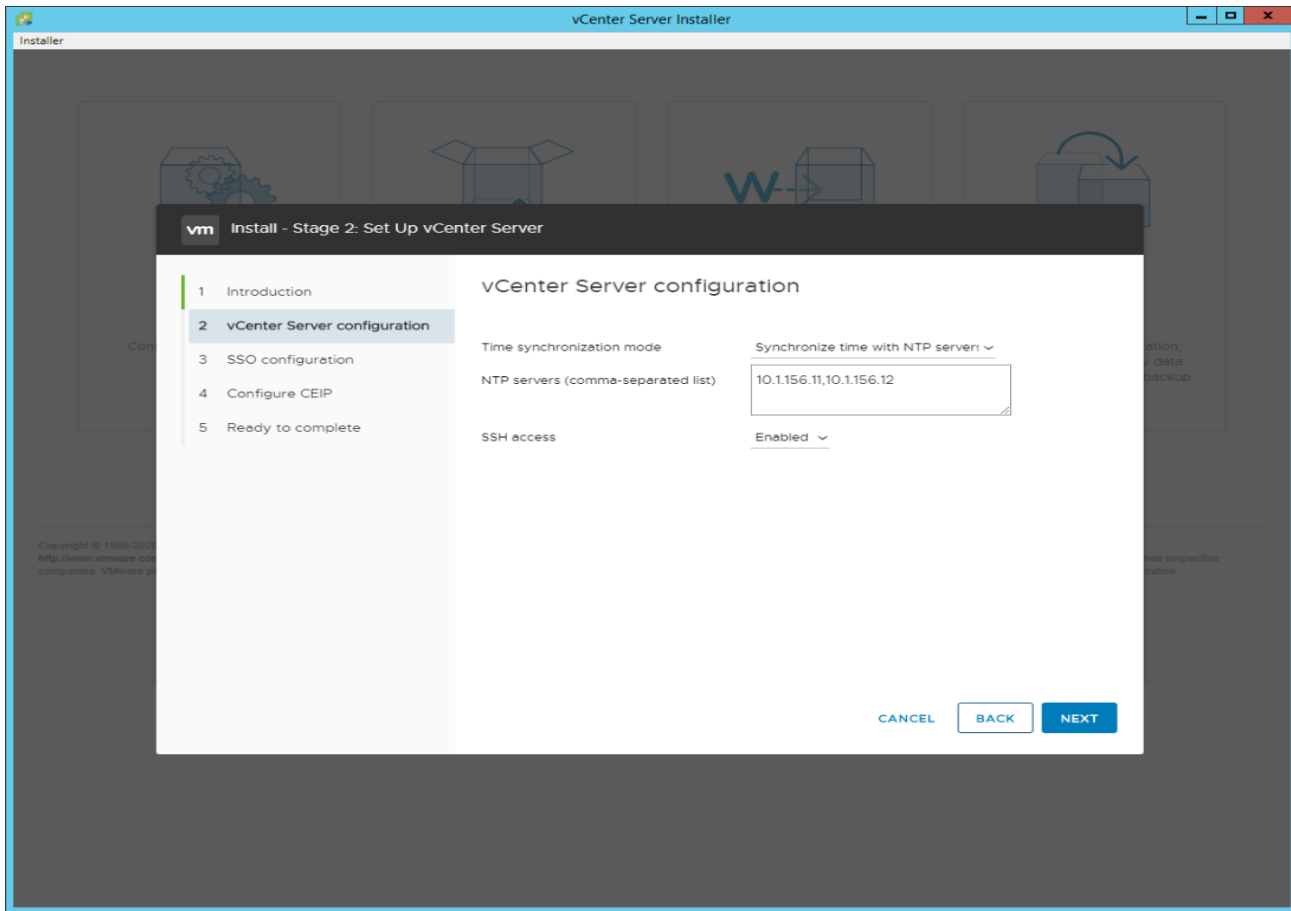


Step 15. Click CONTINUE to proceed with stage 2 configuration.

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.

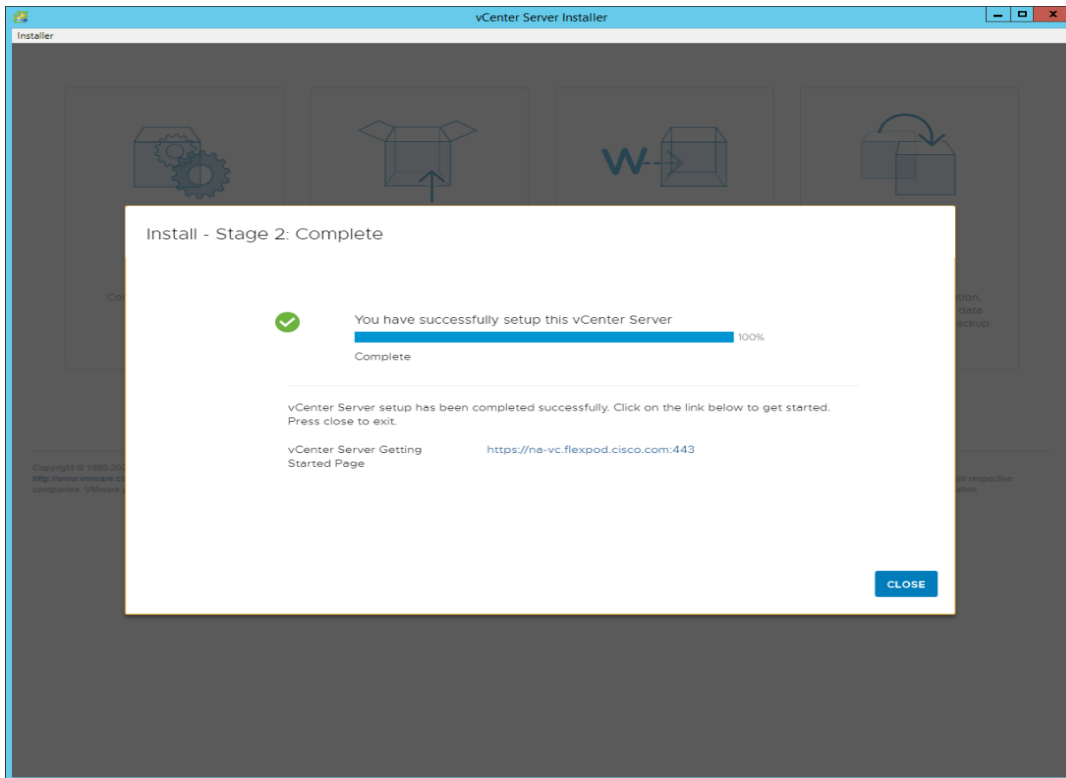
Step 21. Decide whether to join VMware's Customer Experience Improvement Program (CEIP).

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.

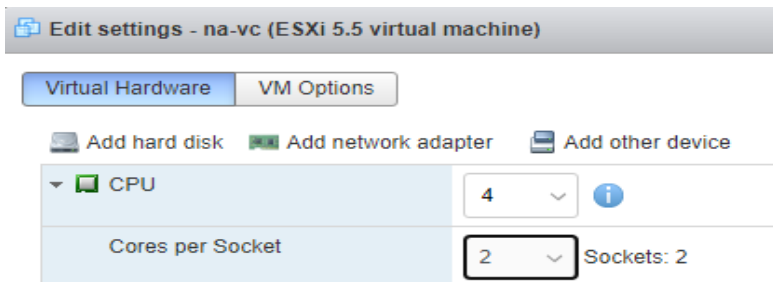
Step 7. In the Edit settings window, expand CPU and check the value of Sockets.



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 7.0.2.00500 was installed.

Step 10. Go to root > Logout to logout of the Appliance Management interface.

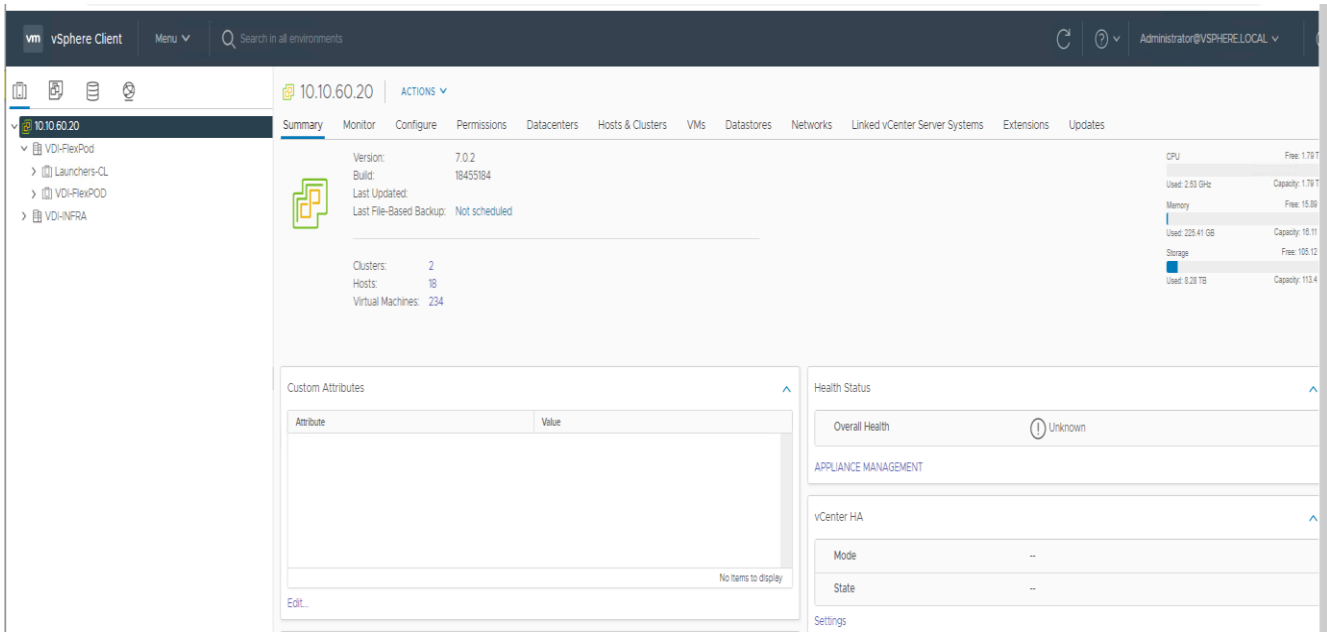
Step 11. Using a web browser, navigate to <https://<vcenter-fqdn>>. You will need to navigate security screens.

Note: With VMware vCenter 7.0.U2 the use of the vCenter FQDN is required.

Step 12. Click LAUNCH VSPHERE CLIENT (HTML5).

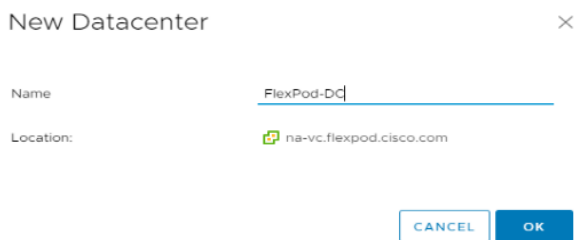
Note: Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is the only option in vSphere 7 and will be used going forward.

Step 13. Log in using the Single Sign-On username (administrator@vsphere.local) 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.



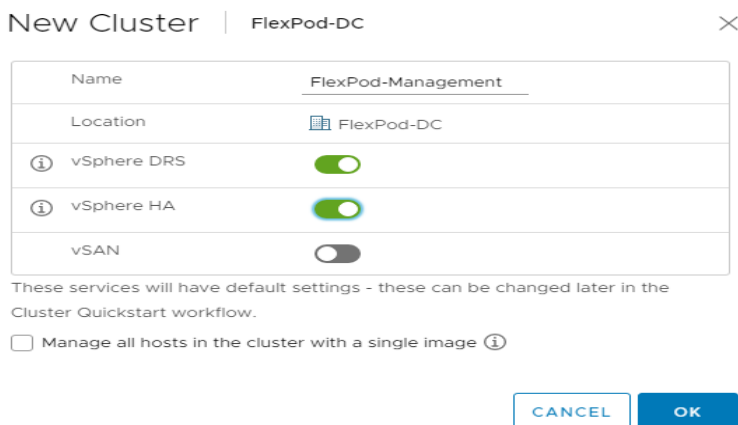
Step 16. Click OK.

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.

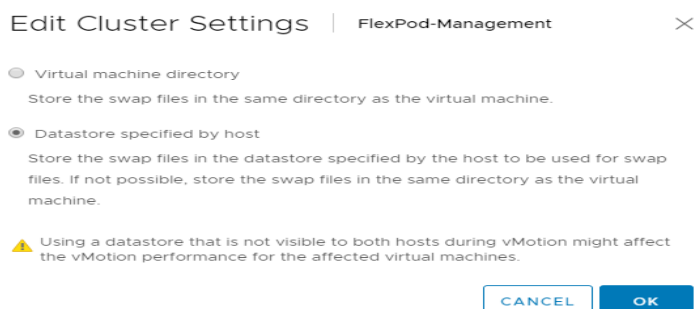


Step 21. Click OK to create the new cluster.

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.

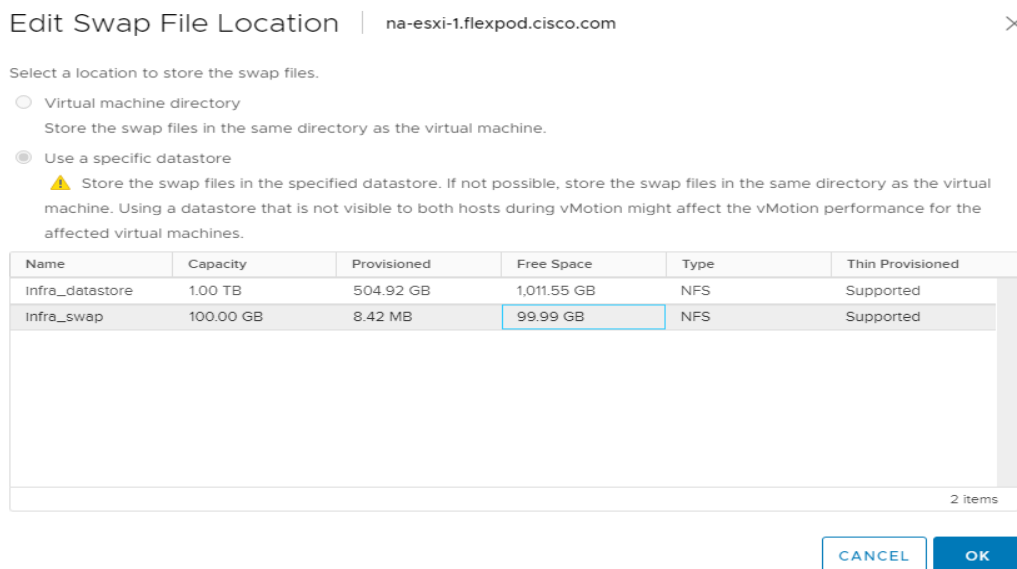
Step 30. The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.

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.



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
 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	Datstore	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](#)

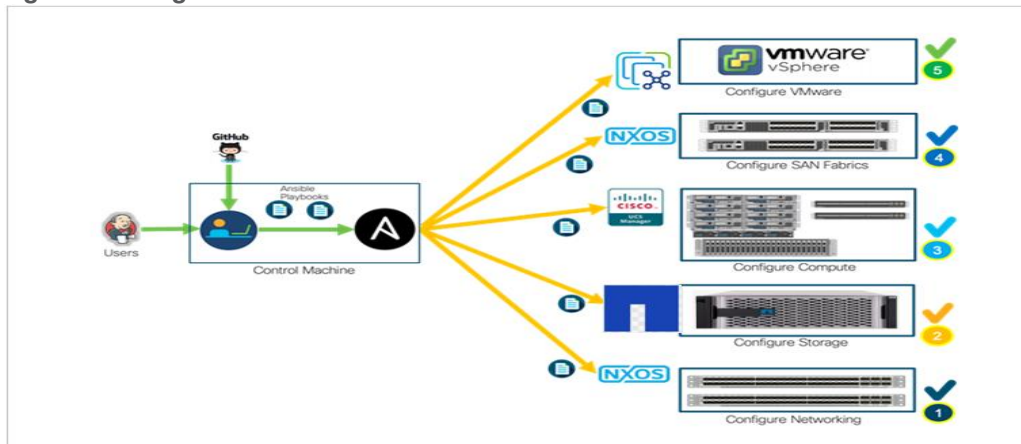
FlexPod Automated Deployment with Ansible

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 U2 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 51. 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 55). 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 ucsm sdk
```

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:

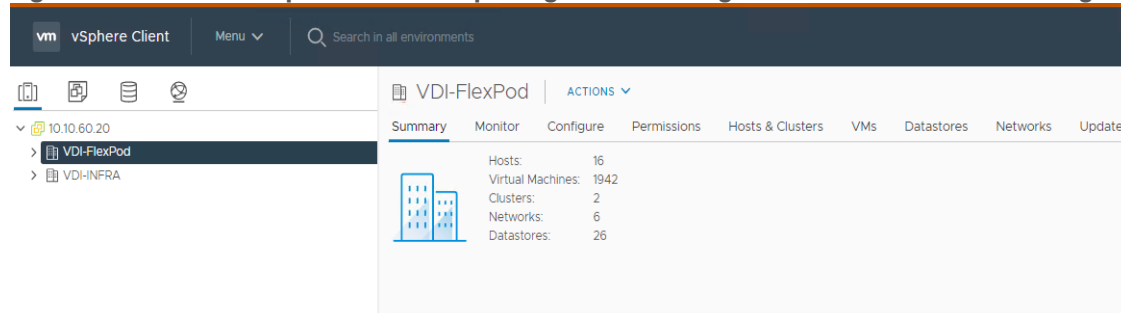
```
esxcli software vib update -d /path/offline-bundle.zip
```

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 52. VMware vSphere WebUI Reporting Cluster Configuration for this Validated Design

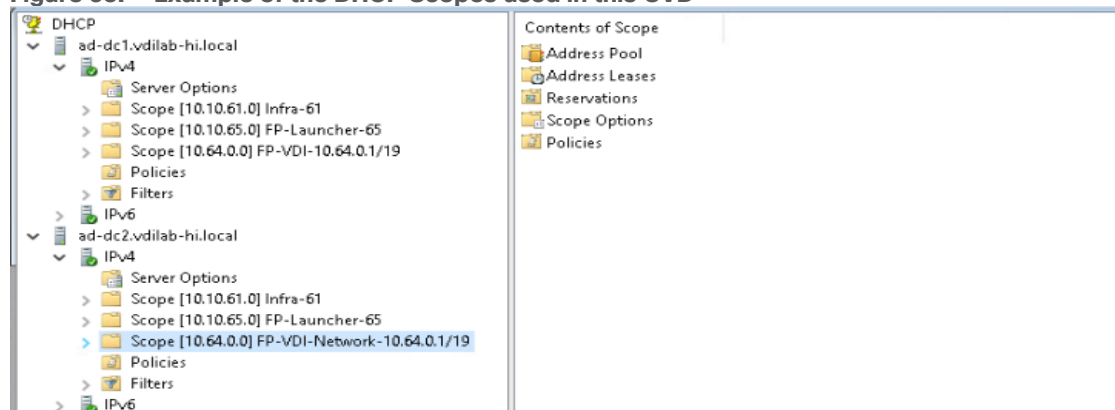


Build the Virtual Machines and Environment for Workload Testing

Prerequisites

Create all necessary DHCP scopes for the environment and set the Scope Options.

Figure 53. 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 17](#).

Table 17. 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	40 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	8 GB
Network	VMXNET3 k21-Infra-Mgmt-61	VMXNET3 k21-Infra-Mgmt-61
Disk-1 (OS) size	40 GB	40 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 2111 or Horizon 8.4

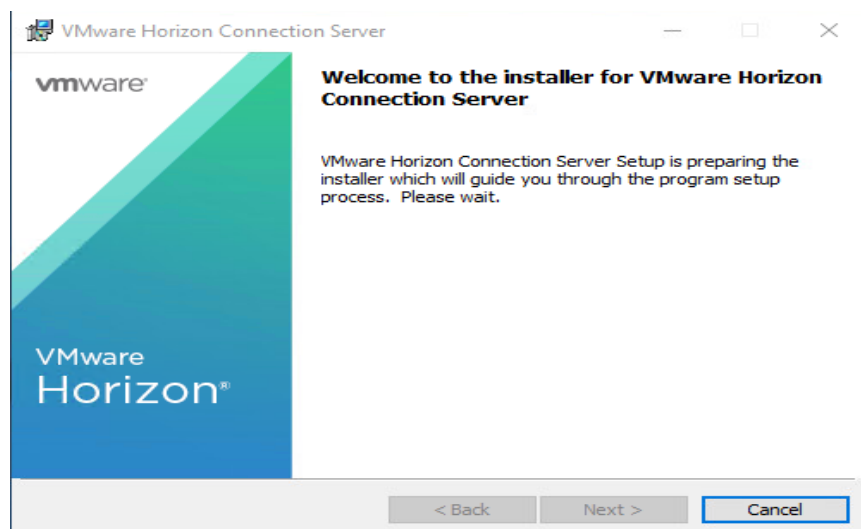
Procedure 1. Install Horizon Connection and Replica Servers

Step 1. Download the VMware Horizon 8 installation package from this link:

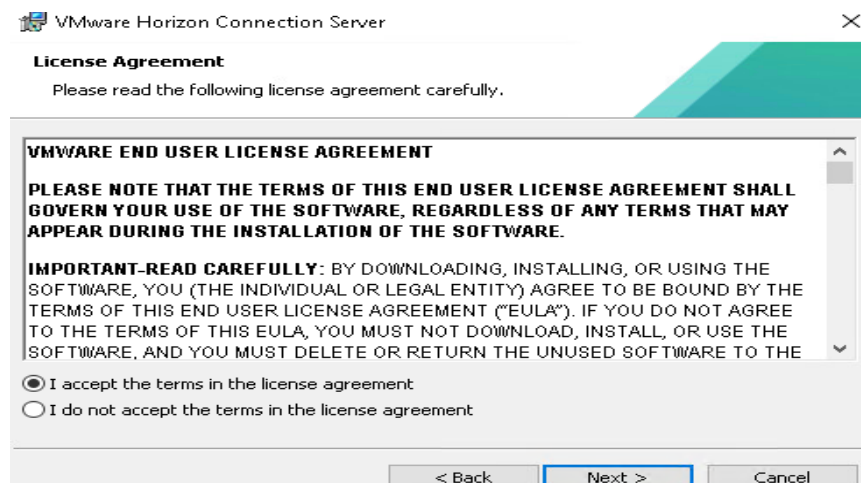
https://customerconnect.vmware.com/downloads/info/slug/desktop_end_user_computing/vmware_horizon/2111

Step 2. Open view connection server installation, VMware-viewconnectionsrvr-x86_64-8.4.0-18964782.exe.

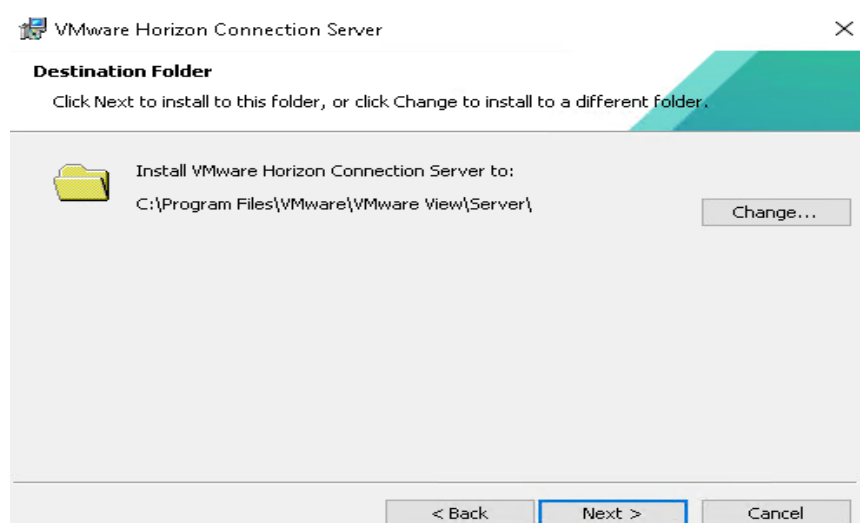
Step 3. Click Next.



Step 4. Accept the EULA then click Next.

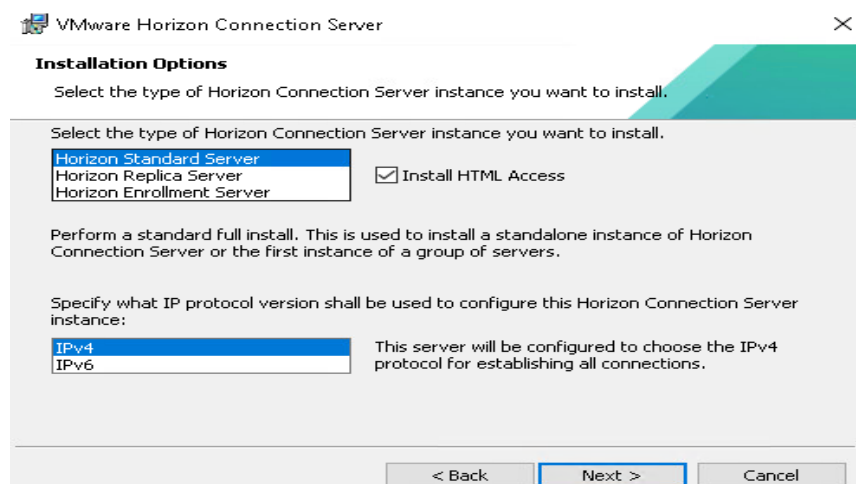


Step 5. Keep the default destination folder and click Next.

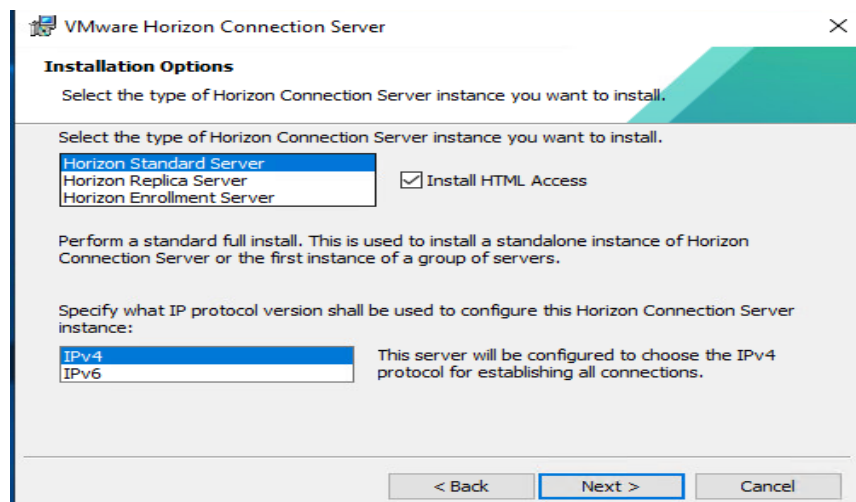


Step 6. Select type of instance intended to install.

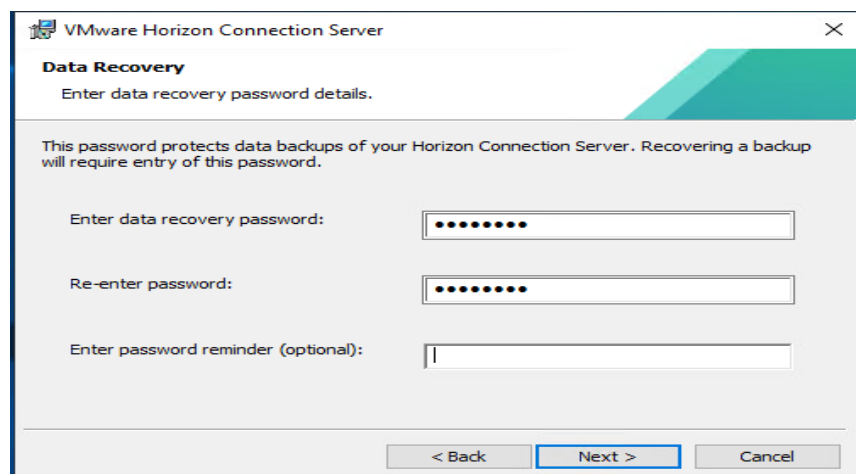
Step 7. Select Standard Server instance for primary connection server installation.



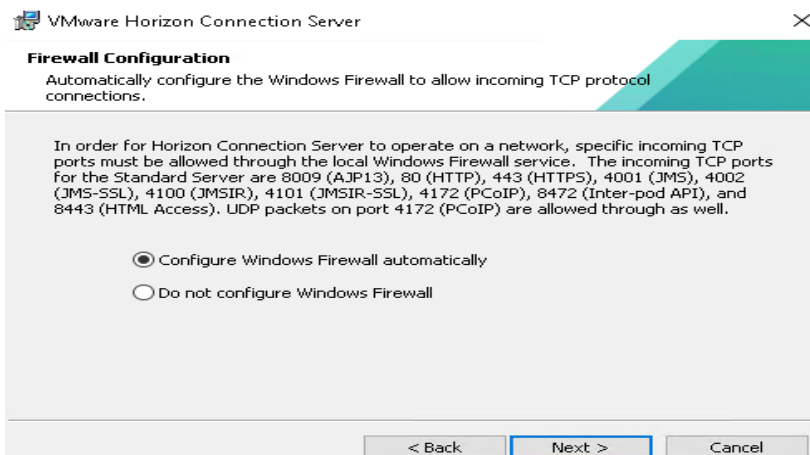
Step 8. Select Replica server instance for fault tolerant connection server configuration after completion of Standard Server instance installation.



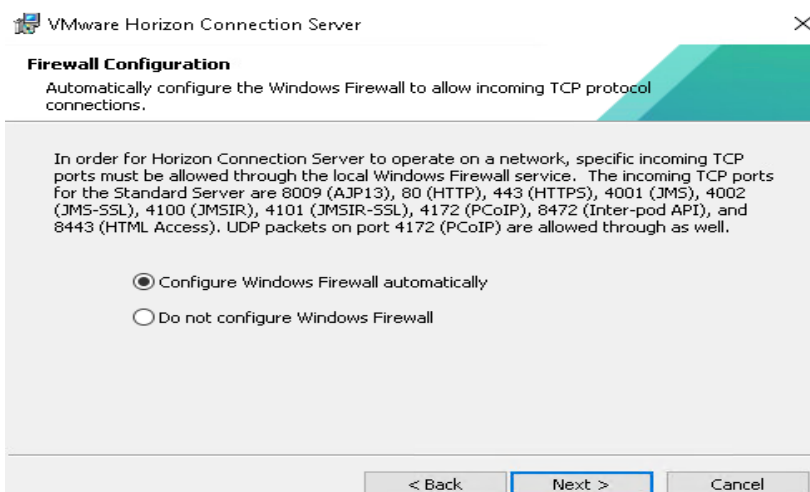
Step 9. Enter the Data Recovery Password.



Step 10. Click Next.

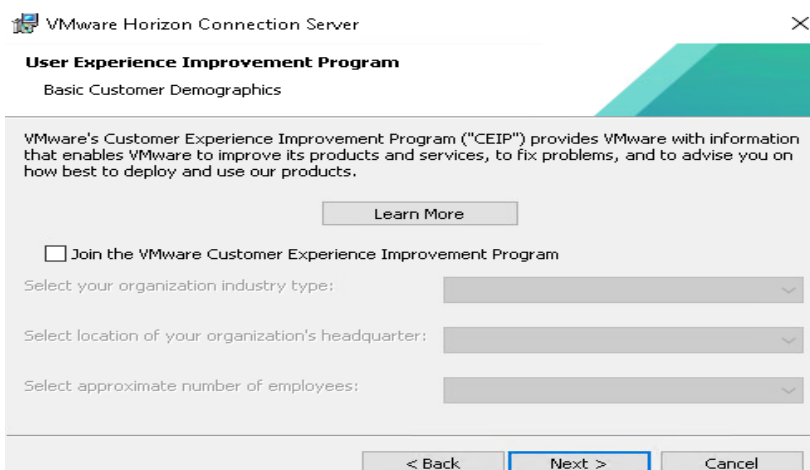


Step 11. Select authorized users and group, click Next.

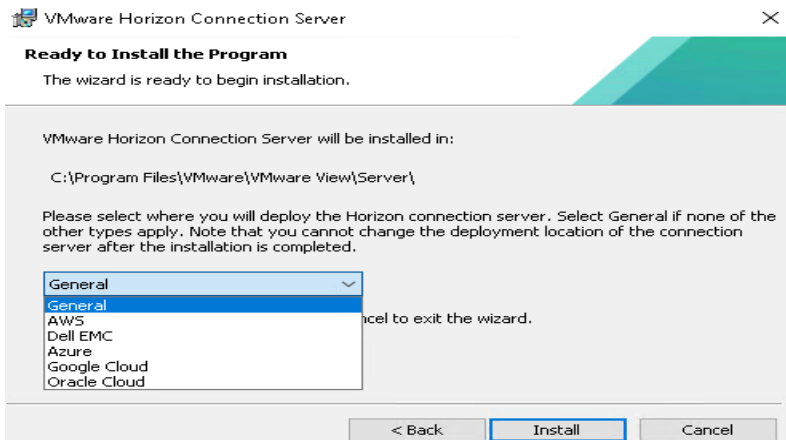


Step 12. Enter domain credentials for previously specified domain user/group.

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=2068666

Procedure 3. Configure the Horizon 8 Environment

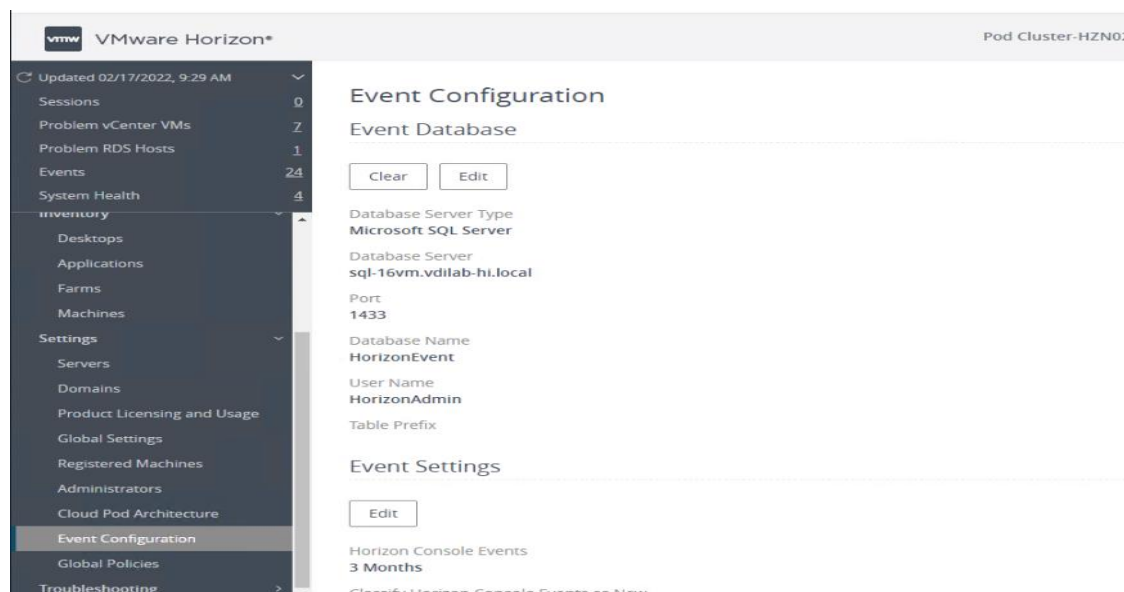
Step 1. Open WebUI, Login to https://<Horizon_Connection_server_Management_IP_Address>/admin.



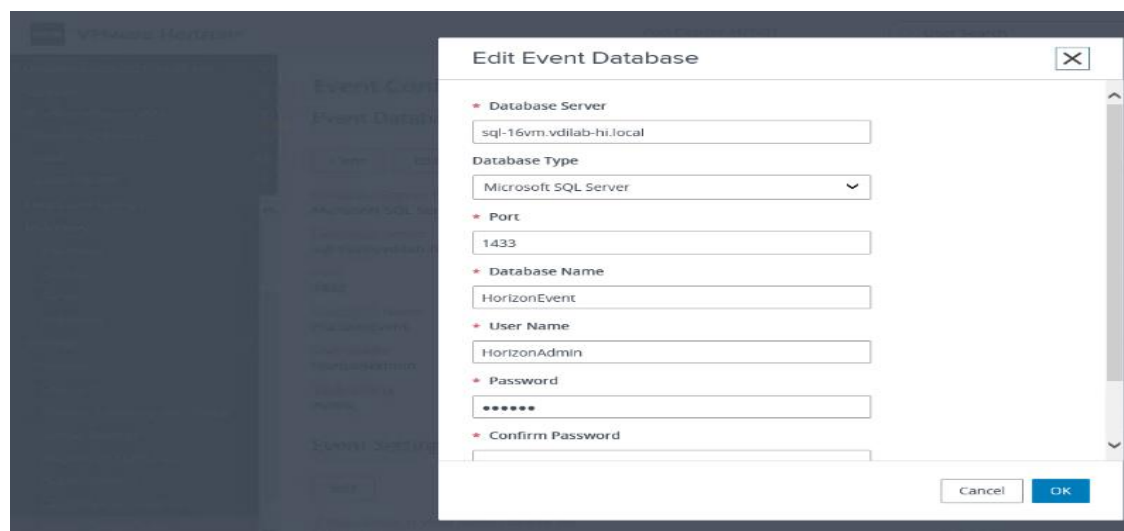
Procedure 4. Configure Event Database

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

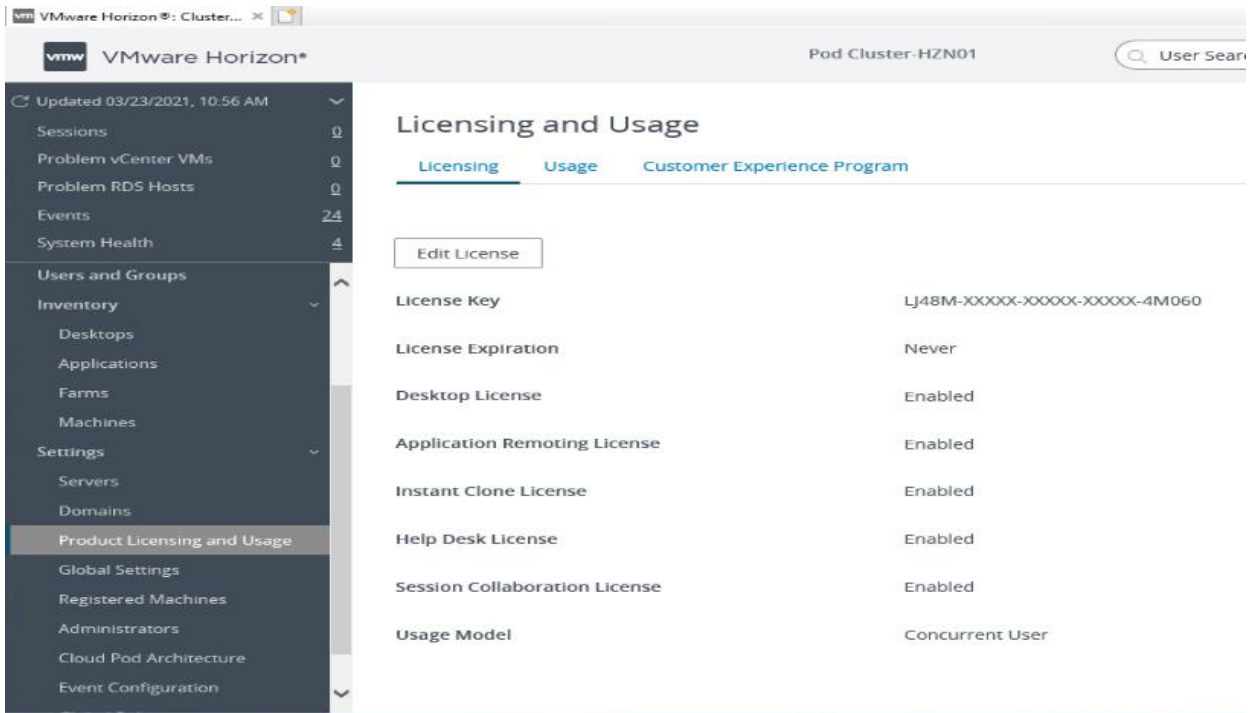
Step 1. Click View Configuration.

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.

Edit vCenter Server - VCSA.VDILAB-HI.LOCAL ✕

Asterisk (*) denotes required field

* Server address ⓘ

VCSA.VDILAB-HI.LOCAL

* User Name

administrator@vsphere.local

* Password

.....

Description

VC for FP

SSL

* Port

443

Deployment Type

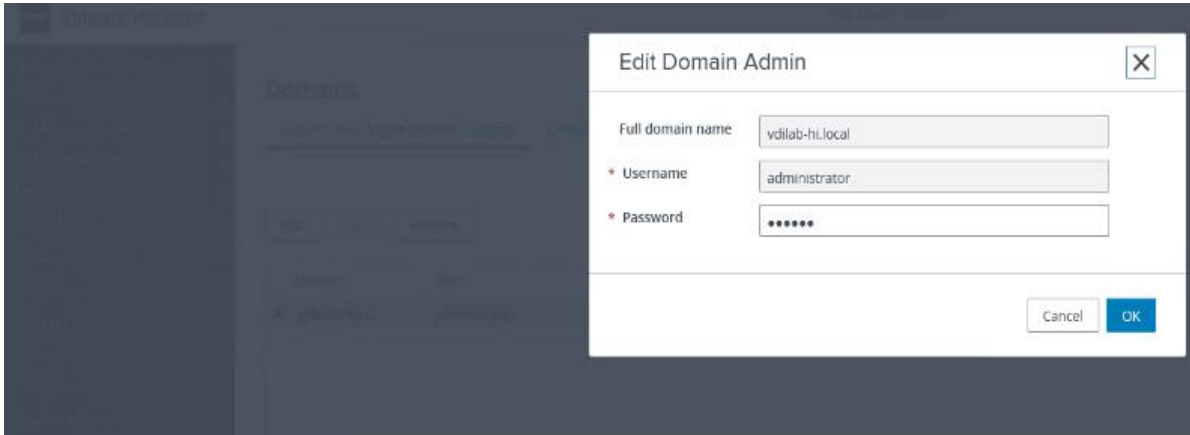
Step 4. Click View certificate. Accept the certificate and click to complete adding vCenter.

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.



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 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](#)

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 18](#) lists the parameters use for Master Image virtual machines.

Table 18. 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	8
Memory	3.5 Gb	3.5 Gb	32768MB
Memory reserved	3.5 Gb	3.5 Gb	32768MB

Attribute	Instant / Non-Persistent Clone	Persistent / Full Clone	RDSH server
Video RAM	35 MB	35 MB	4MB
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 11333	Release 11333	Release 11333
VMware View Agent	Release 8.4.0-18964730 (Version 2111)	Release 8.4.0-18964730 (Version 2111)	Release 8.4.0-18964730 (Version 2111)
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 2111 Remote Desktop Server Sessions (RDSH) Server 2019 Virtual Machines
- VMware Horizon 2111 Instant Clone non-persistent virtual machines
- VMware Horizon 2111 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:

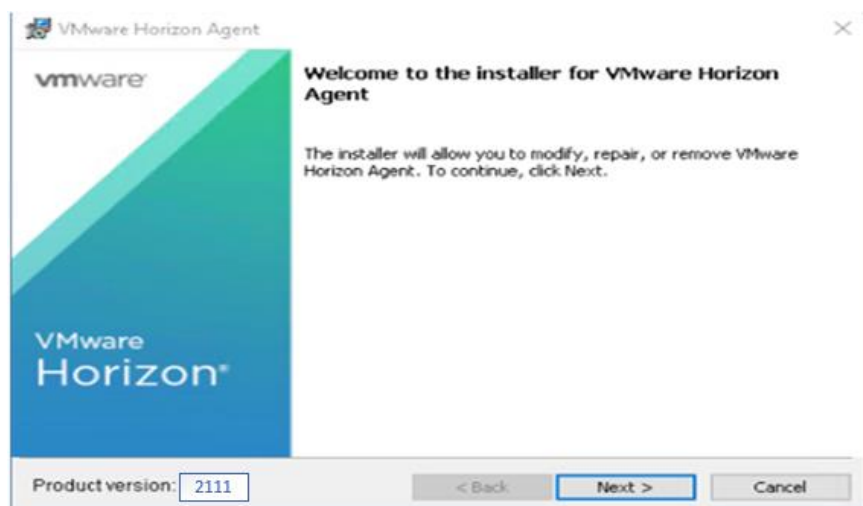
VMware Windows Operating System Optimization Tool Guide:

<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-2111 or version 8.4.0-18964730.exe Click Next to install.

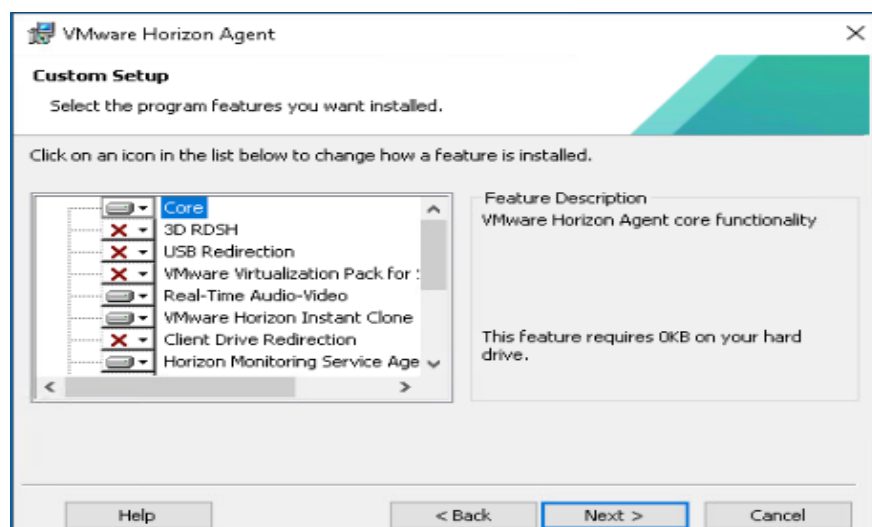


Step 2. Review and accept the EULA Agreement. Click Next.

Step 3. Select Network protocol configuration, click Next.

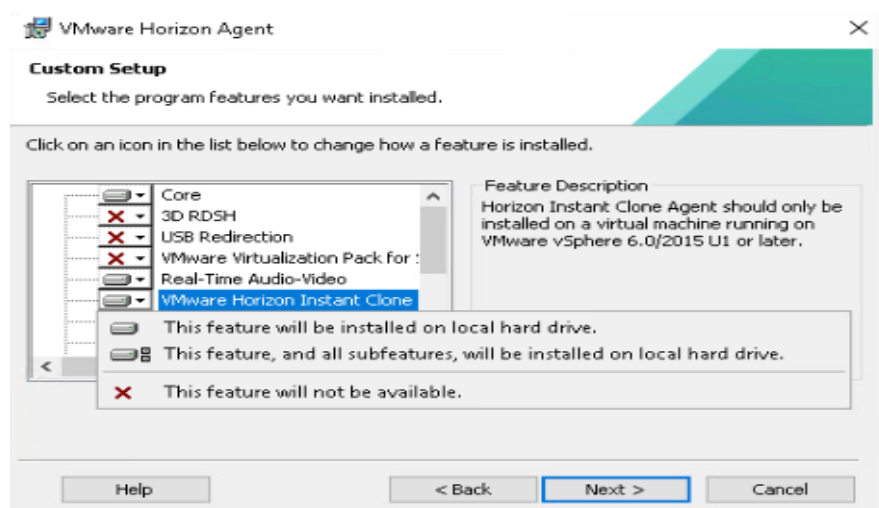
Step 4. Based on the Desktop pool you want to create, select either View Composer Agent or Instant Clone Agent installation. Do not install both features on the same master image.

Step 5. Enable installation of the VMware Horizon View Instant Agent for Instant -clone VDI virtual machines.



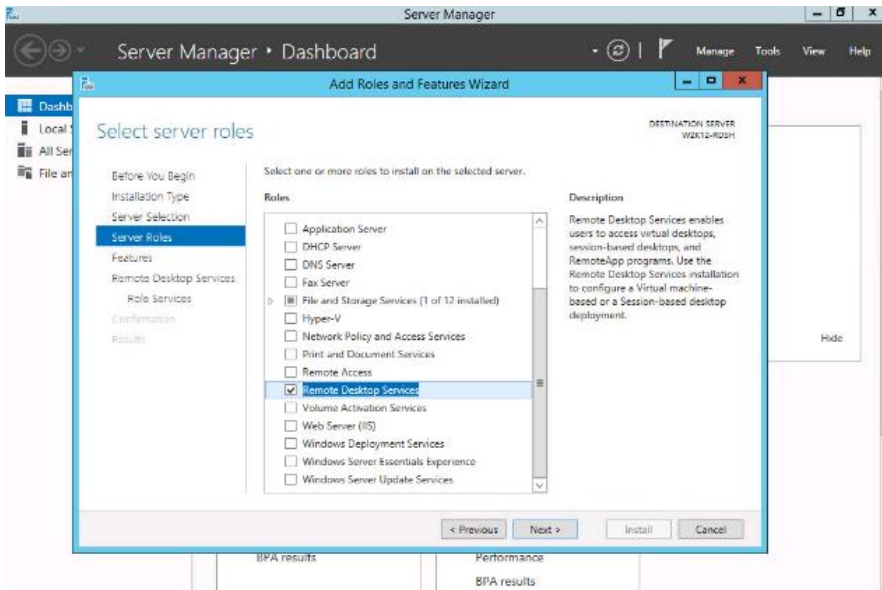
Step 6. Disable the Horizon View Composer Agent and enable the Horizon Instant Clone Agent for Instant Clone floating assigned desktop pool creation.

Note: The VMware Horizon Composer was not tested for this CVD.

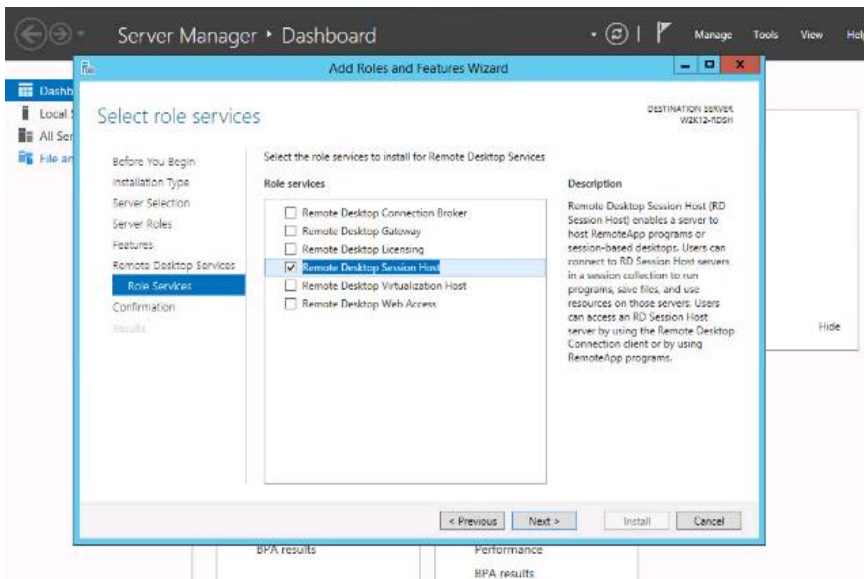


Note: Prior to installing the Horizon View Agent on a Microsoft Server 2019 virtual machine, you must add the Remote Desktop Services role and the Remote Desktop Session Host role service.

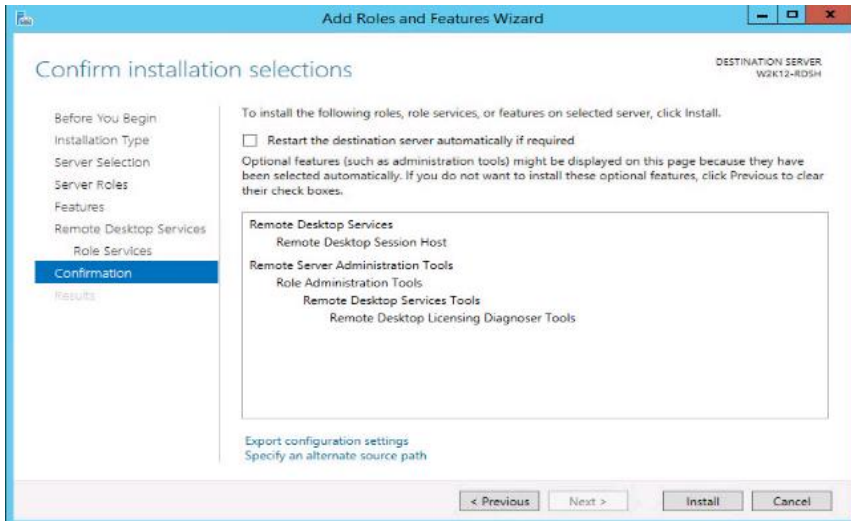
Step 7. To add Remote Desktop Services role on Windows Server OS from the Server Manager, use the Add Roles and Features wizard:



Step 8. Add Remote Desktop Session Host services.



Step 9. Click Install.

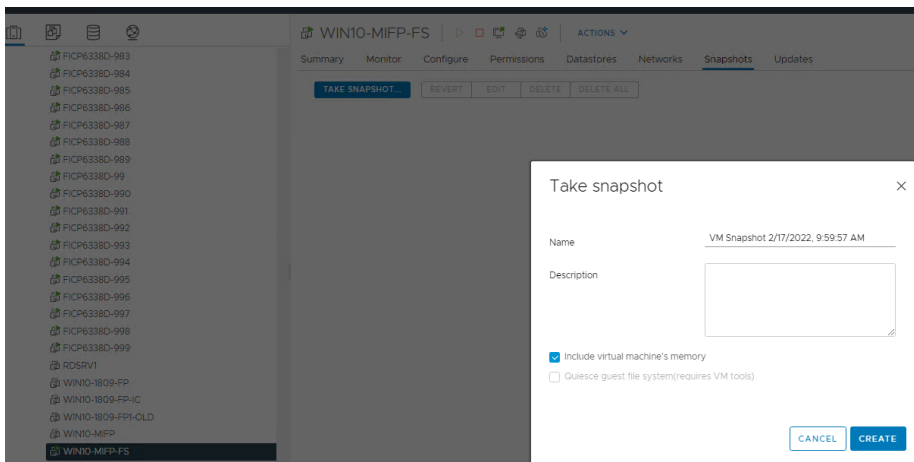


Procedure 2. Install Additional Software

- 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

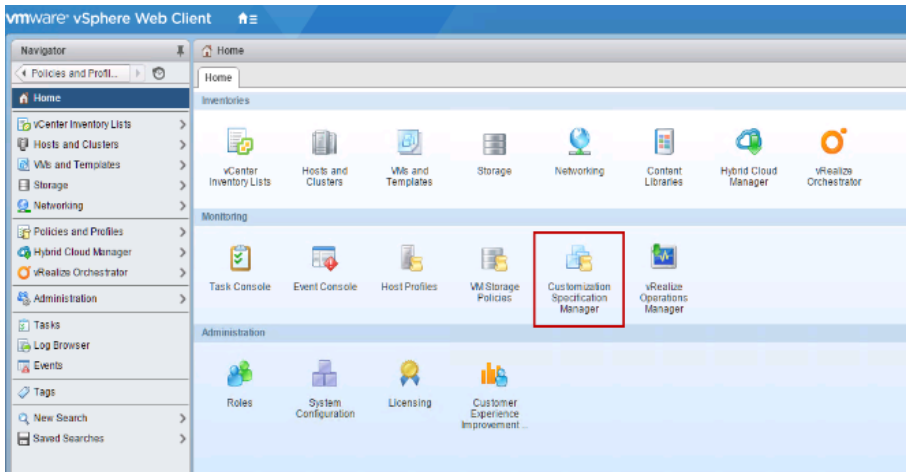
- Step 1.** Log into vCenter WebUI.
- Step 2.** Select the master image for the automated desktop pool creation.
- Step 3.** Right-click and select Master Image X Data Platform > Snapshot Now.



- 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.

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

VM Customization Specification

Name WIN10-Specs

Description

vCenter Server VCSA71UC.vdilab-hi.local

Guest OS

Target guest OS Windows Linux

Use custom SysPrep answer file

Generate a new security identity (SID)

CANCEL OK

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

Windows license

Administrator password

Time zone

Commands to run once

Network

Workgroup or domain

Ready to complete

Use the virtual machine name ⓘ

Enter a name in the Clone/Deploy wizard

Enter a name

Append a unique numeric value. ⓘ

Generate a name using the custom application configured with the vCenter Server

Argument _____

Step 6. Provide Password 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

Password _____

Confirm password _____

Automatically logon as Administrator

Number of times to logon automatically 1

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)

CANCEL OK

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

Use standard network settings for the guest operating system, including enabling DHCP on all network interfaces

Manually select custom settings

Computer name

Windows license

ADD EDIT DELETE

Administrator password

Time zone

Commands to run once

	Description	IPv4 Address	IPv6 Address
<input type="radio"/>	NIC1	Use DHCP	Not used

Network

Workgroup or domain

Ready to complete

Step 10. Provide the domain name and user credentials.

WIN10-Specs - Editing

Name and target OS

Registration information

Workgroup

WORKGROUP

Computer name

Windows license

Windows Server domain

vdllab-hi.local

Administrator password

Specify a user account that has permission to add a computer to the domain.

Time zone

Username

Administrator

Commands to run once

Password

Network

Workgroup or domain

Confirm password

Ready to complete

Step 11. Review and click Next to complete creating the Customization Specs.

Step 12. Click Finish.

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

Name
Target guest OS
OS options
Registration info
Computer name
Product key
Server license mode
Administrator access
Time zone
Network type
Windows Server domain

WIN10-Specs
Windows
Generate new security ID
Owner name: Administrator
Organization: vdlab-hi.local
Use Virtual Machine name
No product key specified
Per server (Maximum Connections: 5)
Do not log in automatically as Administrator
(UTC-08:00) Pacific Time (US & Canada)
Standard
vdlab-hi.local
Username: Administrator

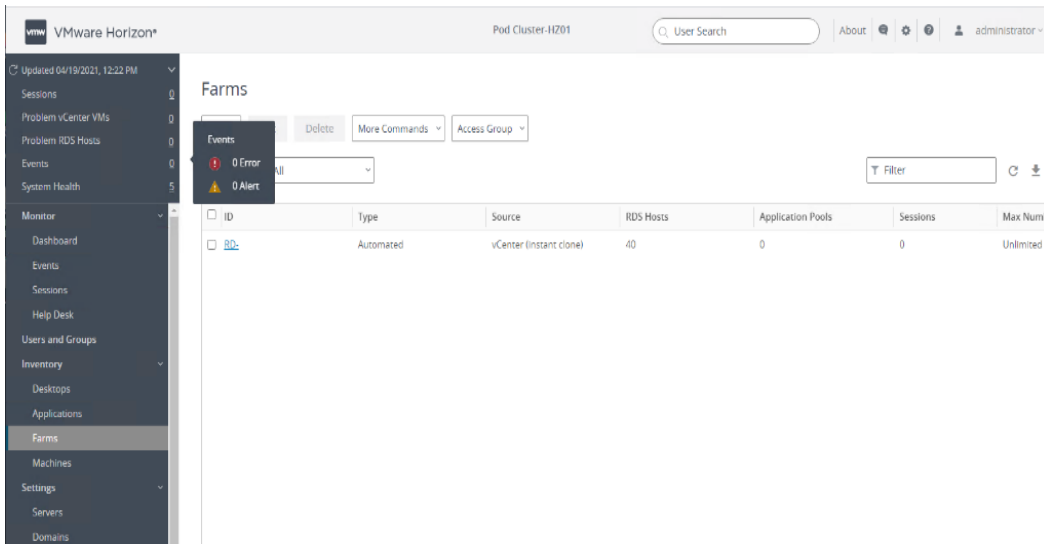
CANCEL OK

Procedure 5. Create RDSH Farm

Note: Before you can create an RDSH desktop pool, you must first create a RDSH Farm.

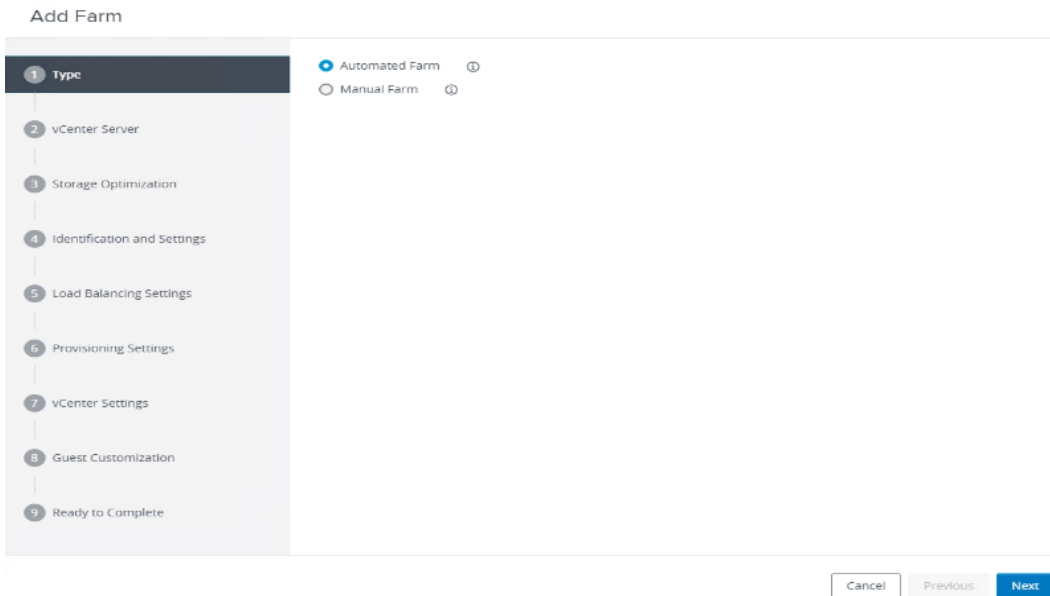
Step 1. In the VMware Horizon Administration console, select Farms under the Resource node of the Inventory pane.

Step 2. Click Add in the action pane to create a new RDSH Farm.



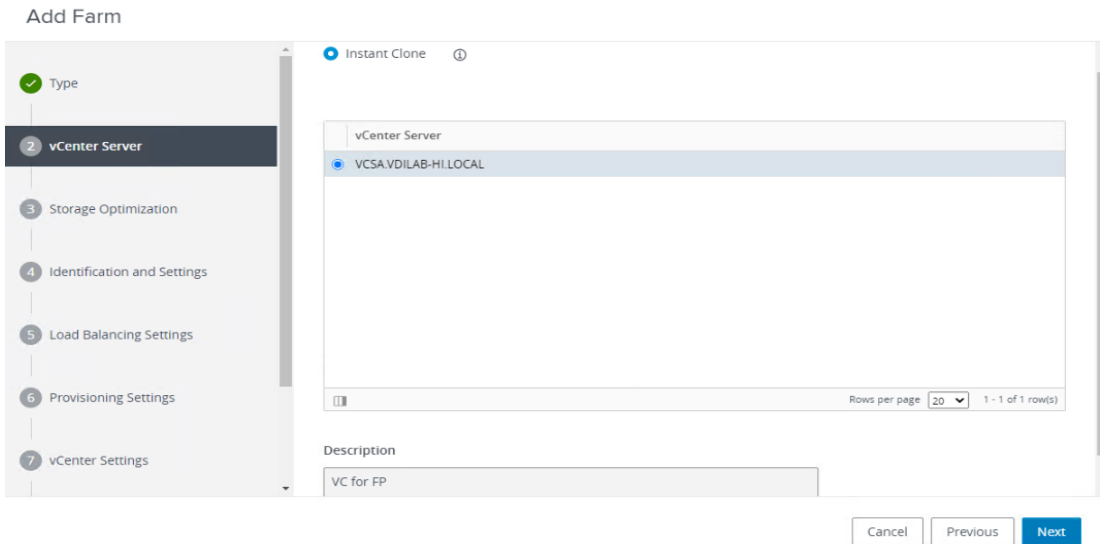
Step 3. Select either to create an Automated or Manual Farm. In this solution, we selected Automated Farm.

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.

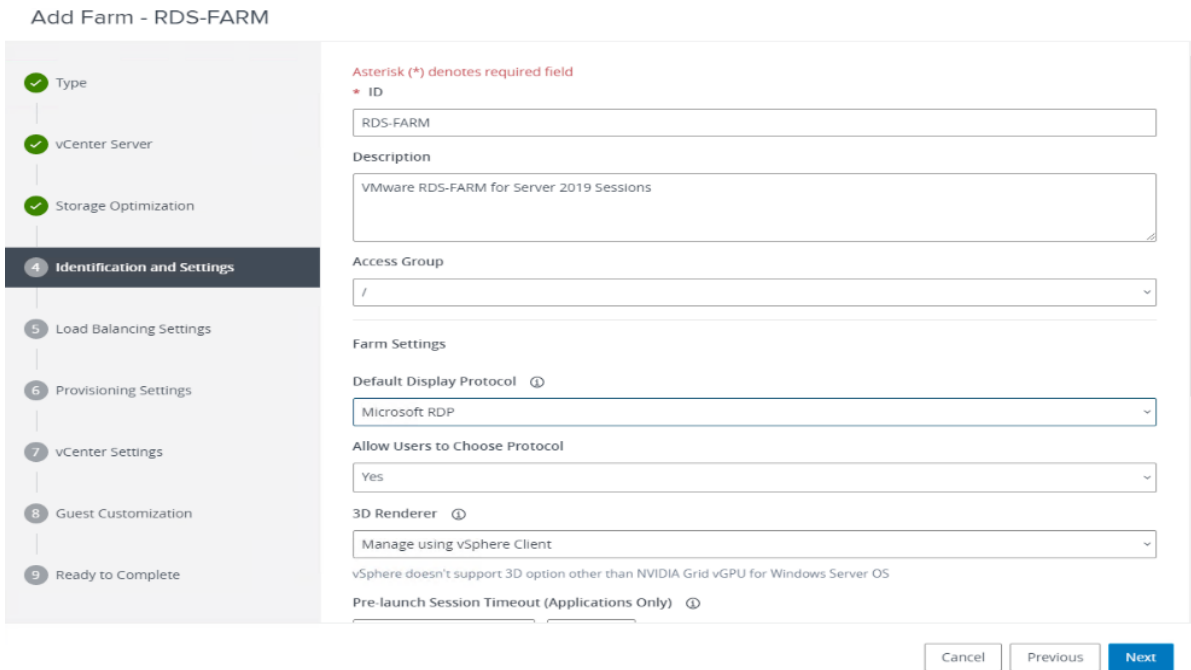


Note: You can choose to create either Instant clones or View Composer linked clones for the RDSH server FARM server VMs. Both have benefits and limitations, but detailing these differences are beyond the scope of this CVD. Please refer to your VMware documentation for more information.

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.



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 2300 RDSH virtual machines across our 8 node Cisco UCS B200 M6 Cluster.

Step 10. Click Next.

Add Farm - RDS

The screenshot shows the 'Add Farm - RDS' configuration page. On the left is a vertical navigation pane with steps 1 through 7. Step 6, 'Provisioning Settings', is highlighted. The main content area is titled 'Basic' and contains the following settings:

- Asterisk (*) denotes required field**
- Basic**
 - Enable Provisioning ⓘ
 - Stop Provisioning on Error
- Virtual Machine Naming** ⓘ
 - * Naming Pattern:
- Farm Sizing**
 - * Maximum Machines:
 - * Minimum Number of Ready (Provisioned) Machines during Instant Clone Maintenance Operations:

At the bottom right, there are three buttons: 'Cancel', 'Previous', and '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

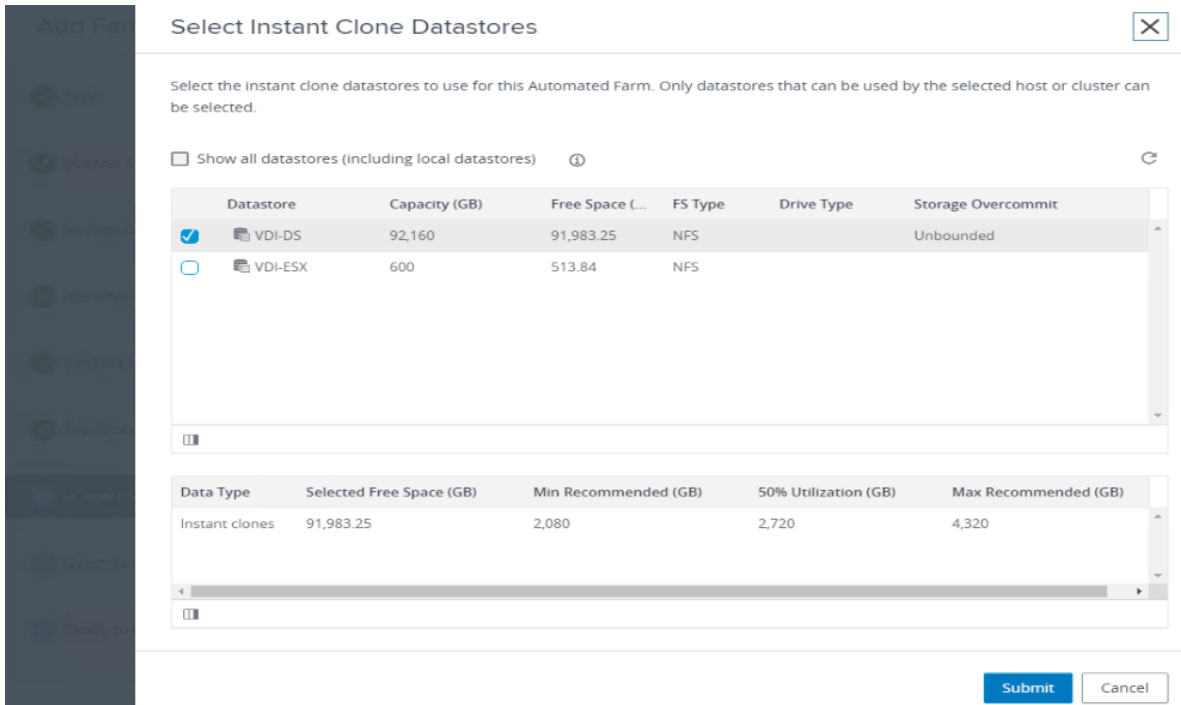
The screenshot shows the 'Add Farm - RDS-FARM' configuration page. On the left is a vertical navigation pane with steps 7 through 9. Step 7, 'vCenter Settings', is highlighted. The main content area is titled 'Default Image' and contains the following settings:

- Asterisk (*) denotes required field**
- Default Image**
 - * Golden Image in vCenter:
 - * Snapshot:
- Virtual Machine Location**
 - * VM Folder Location:
- Resource Settings**
 - * Cluster:
 - * Resource Pool:
 - * Datastores: 1 selected
- Network**

At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

Step 14. For step 6 Datastores: Browse and click Data Stores.

Step 15. Click OK.

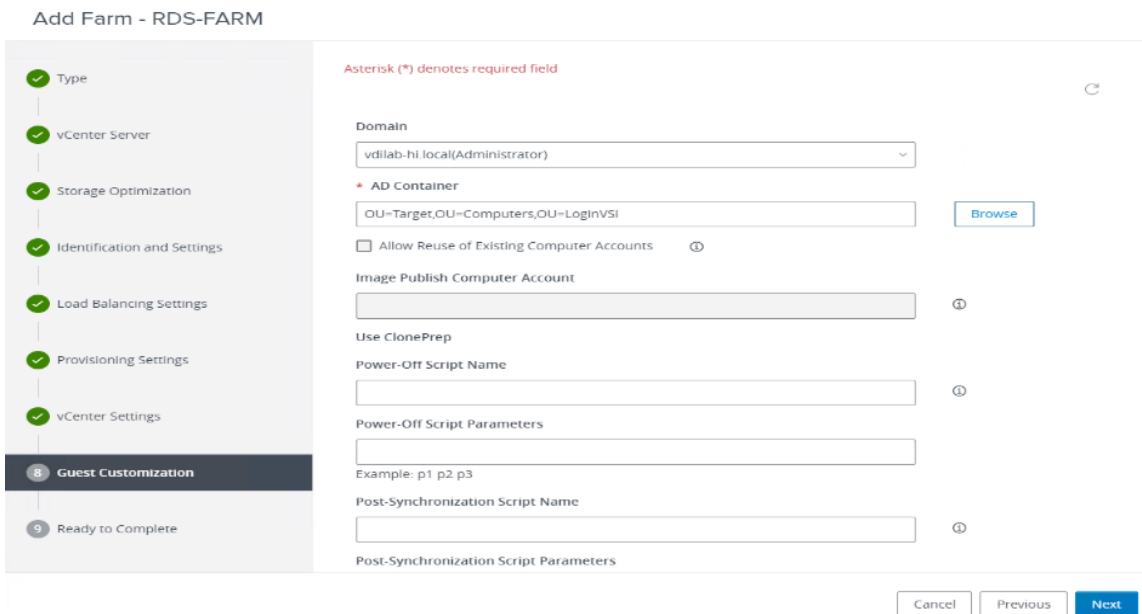


Step 16. Click Next.

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.



Step 19. Review the pool creation information.

Step 20. Click Finish.

Add Farm - RDS-FARM

- ✔ Type
- ✔ vCenter Server
- ✔ Storage Optimization
- ✔ Identification and Settings
- ✔ Load Balancing Settings
- ✔ Provisioning Settings
- ✔ vCenter Settings
- ✔ Guest Customization
- 9 Ready to Complete

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
Pre-launch Session Timeout (Applications Only)	10 minutes
Empty Session Timeout (Applications Only)	1 minute
When Timeout Occurs	Disconnect
Logout 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 ⚙️ 🔒 👤 administrator

- Updated 04/19/2021, 12:22 PM
- Sessions 0
- Problem vCenter VMs 0
- Problem RDS Hosts 0
- Events 0
- System Health 5
- Monitor
- Dashboard
- Events
- Sessions
- Help Desk
- Users and Groups
- Inventory
- Desktops
- Applications
- Farms

Farms

Access Group:
Filter

ID	Type	Source	RDS Hosts	Application Pools	Sessions	Max Number
RD-	Automated	vCenter (instant clone)	40	0	0	Unlimited

VMware Horizon* Pod Cluster: HZN02

Updated: 03/25/2022, 4:07 PM

Sessions: 1
Problem vCenter VMs: 0
Problem RDS Hosts: 0
Events: 8
System Health: 4

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
rtfarm-1.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-10.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-11.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-12.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-13.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-14.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-15.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-16.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-17.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-18.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-19.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-2.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-20.vdliab-hi.local	Windows Server 2016 or above	RDSRV1 - RDSRV1...		None	35	8.4.0-19446757	✔	Available
rtfarm-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.

VMware Horizon* Pod Cluster: HZN02

Updated: 03/25/2022, 4:07 PM

Sessions: 1
Problem vCenter VMs: 0
Problem RDS Hosts: 0
Events: 8
System Health: 4

Monitor

- Dashboard
- Events
- Sessions
- Help Desk
- Users and Groups
- Inventory
- Desktops

Desktop Pools

Add Edit Duplicate Delete Entitlements Status Access Groups View Unentitled

Access Group: All

Filter

ID	Display Name	Type	Source	User Assignment	vCenter Server	Entitled	Application Pools	Enabled	App Shor
RDS-	RDS-	RDS Desktop Pool	vCenter (Instant clone)	Floating Assignment	vcsa.vdliab-hi.local	1	N/A	✔	

Step 3. Select RDS Desktop pool.

Step 4. Click Next.



Add Pool

- 1 Type
- 2 Desktop Pool Identification
- 3 Desktop Pool Settings
- 4 Select RDS Farms
- 5 Ready to Complete

Automated Desktop Pool ⓘ

Manual Desktop Pool ⓘ

RDS Desktop Pool ⓘ

Cancel Previous **Next**

Step 5. Enter Pool ID and Display name.

Step 6. Click Next.

Add Pool - RDS-POOL

- 1 Type
- 2 Desktop Pool Identification
- 3 Desktop Pool Settings
- 4 Select RDS Farms
- 5 Ready to Complete

Asterisk (*) denotes required field

* ID ⓘ

RDS-POOL

Display Name ⓘ

RDS-POOL

Description

RDS-POOL- Server: 2019 RDS Sessions.

Cancel Previous **Next**

Step 7. Accept the default settings on Desktop Pool Settings page.

Step 8. Click Next.



Add Pool - RDS-POOL

- Type
- Desktop Pool Identification
- Desktop Pool Settings**
- Select RDS Farms
- Ready to Complete

State: Enabled

Connection Server Restrictions: None

Category Folder: None

Client Restrictions: Enabled

Allow Separate Desktop Sessions from Different Client Devices: No

- 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

- Type
- Desktop Pool Identification
- Desktop Pool Settings
- Select RDS Farms**
- Ready to Complete

Create a new RDS farm

Select an RDS farm for this desktop pool

Farm ID	Description	RDS Hosts	Max Number of Connections	St
No records available.				

- 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 Users Groups

Domain

Name/User Name

Description

<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.

Add Pool

1 Type

2 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

Automated Desktop Pool ⓘ

Manual Desktop Pool ⓘ

RDS Desktop Pool ⓘ

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 View Composer Linked Clones, highlight your vCenter and View Composer virtual machine.
- 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 Linked Clone Pool.

Step 26. Click Finish to complete the Linked 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-

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-

The screenshot shows the 'vCenter Settings' step of the 'Add Pool - WIN10-VDI-' wizard. On the left, a vertical navigation pane lists steps 1 through 9, with 'vCenter Server' (Step 7) highlighted. The main area contains two radio buttons: 'Instant Clone' (unselected) and 'Full Virtual Machines' (selected). Below the radio buttons is a list box titled 'vCenter Server' containing the IP address '10.10.50.39'. A 'Description' text box below the list box contains the text 'VC 7/UC'. At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

Step 10. Enter the pool identification details.

Add Pool - WIN10-VDI-

The screenshot shows the 'User Assignment' step of the 'Add Pool - WIN10-VDI-' wizard. On the left, the vertical navigation pane highlights 'User Assignment' (Step 8). The main area shows two radio buttons: 'Floating' (unselected) and 'Dedicated' (selected). Under the 'Dedicated' option, there is a checked checkbox for 'Enable Automatic Assignment' and an unchecked checkbox for 'Enable Multi-User Assignment'. Below these checkboxes is the text: 'Automatic assignment is not supported for multi-user assignment pools.' At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

Step 11. Click Next.

Add Pool - WIN10-Persistent

Asterisk (*) denotes required field

Type

- 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.

Add Pool - WIN10-Persistent

Name	Guest OS	Description
SRV2019-CustSpecs	Windows	
WIN10-SPecs	Windows	

Buttons: Cancel, Previous, 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.

<input type="checkbox"/> Entitle Users After Adding Pool	
Type	Automated Desktop Pool
User Assignment	Dedicated Assignment
Assign on First Login	Yes
Enable Multi-User Assignment	No
vCenter Server	10.10.50.39
Unique ID	WIN10-Persistent
Description	-
Display Name	WIN10-Persistent
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Show Assigned Machine Name	Disabled

Buttons: Cancel, Previous, Submit

Step 28. Follow the instructions provided in the Create Horizon 8 RDS Desktop Pool to authorize users for the Linked 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

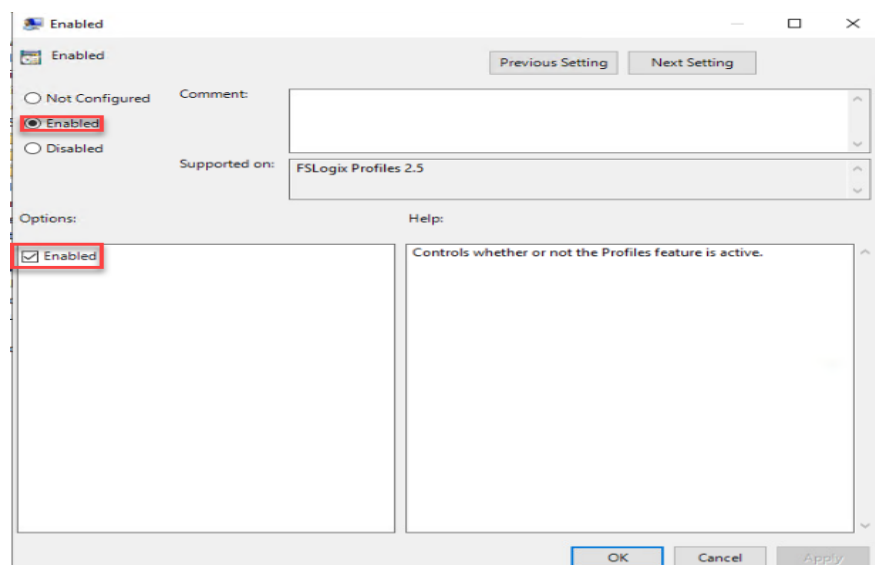
Step 7. When the FSLogix software is downloaded, copy the 'fslogix.admx and fslogix.adml' to the 'PolicyDefinitions' folder in your domain to manage the settings with Group Policy.

Step 8. On your VDI master image, install the FSLogix agent 'FSLogixAppsSetup' and accept all the defaults.

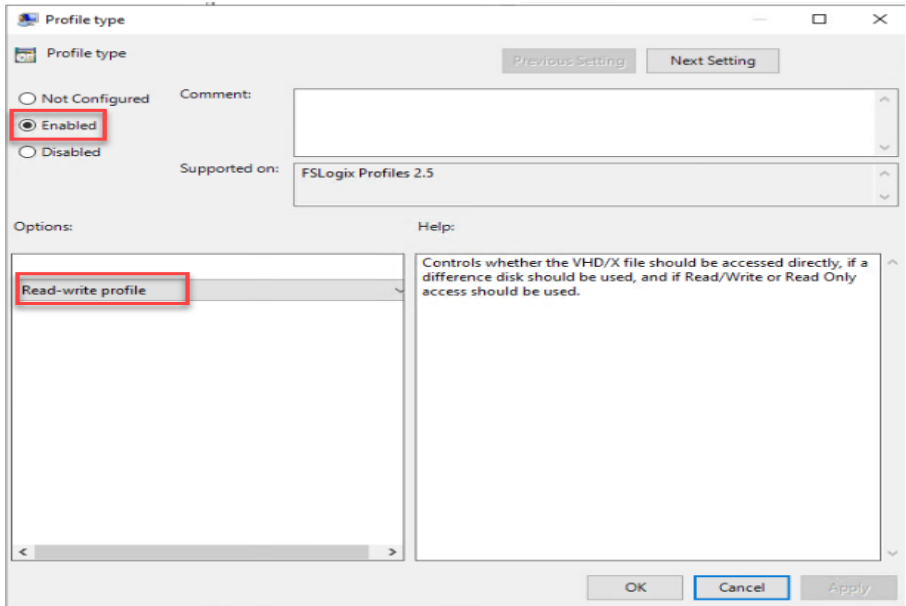
Step 9. Create a Group Policy object and link it to the Organizational Unit the VDI computer accounts.

Step 10. Right-click the FSLogix GPO policy.

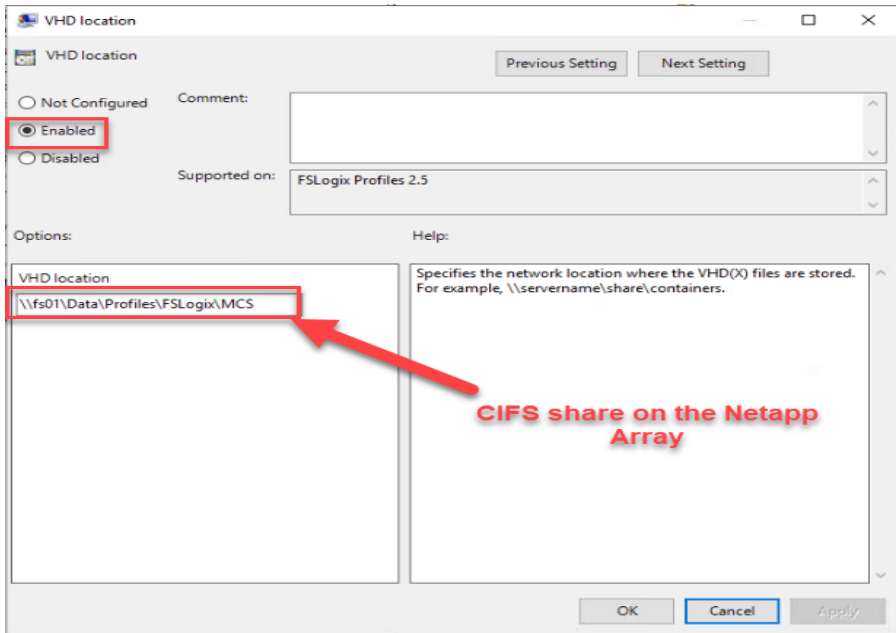
Step 11. Enable FSLogix Profile Management.



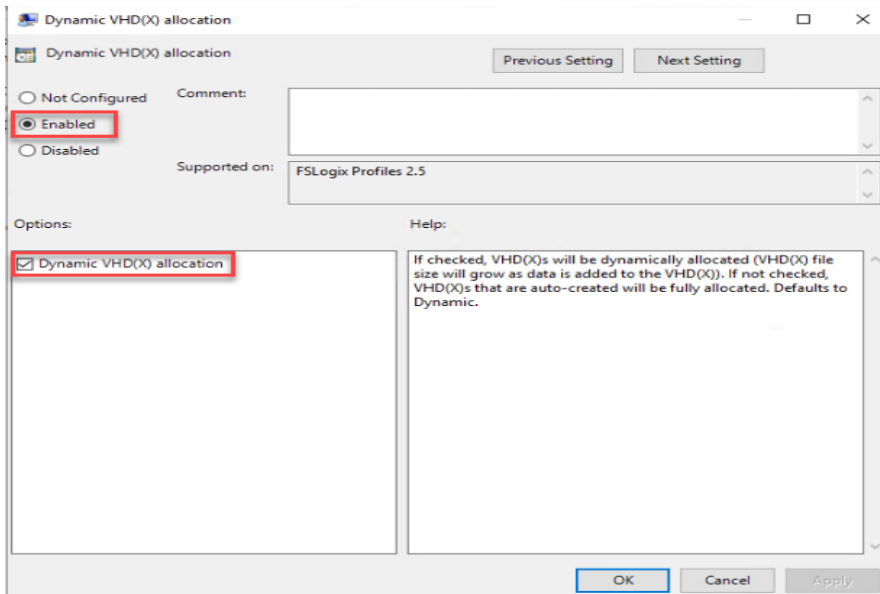
Step 12. Select Profile Type (in this solution, we used Read-Write profiles).



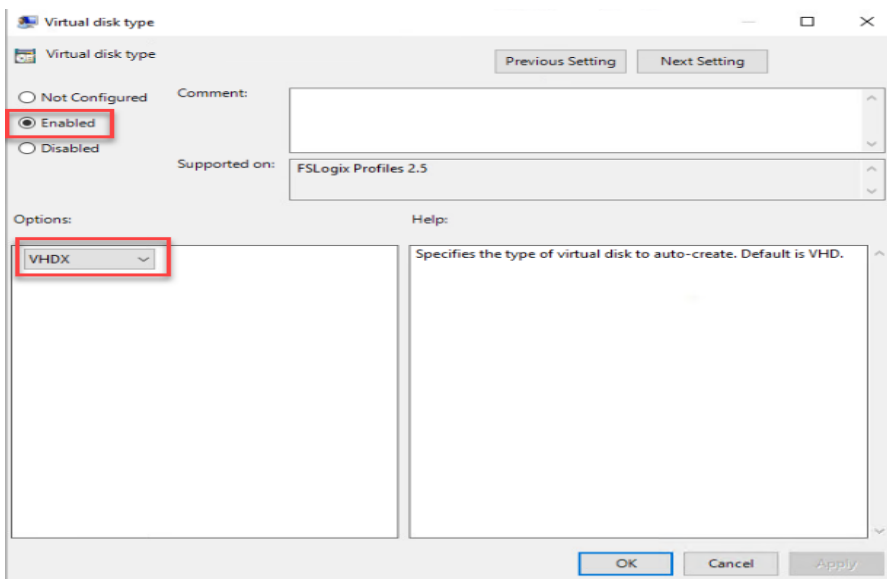
Step 13. Enter the location of the Profile location (our solution used a CIFS share on the Netapp Array).




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'

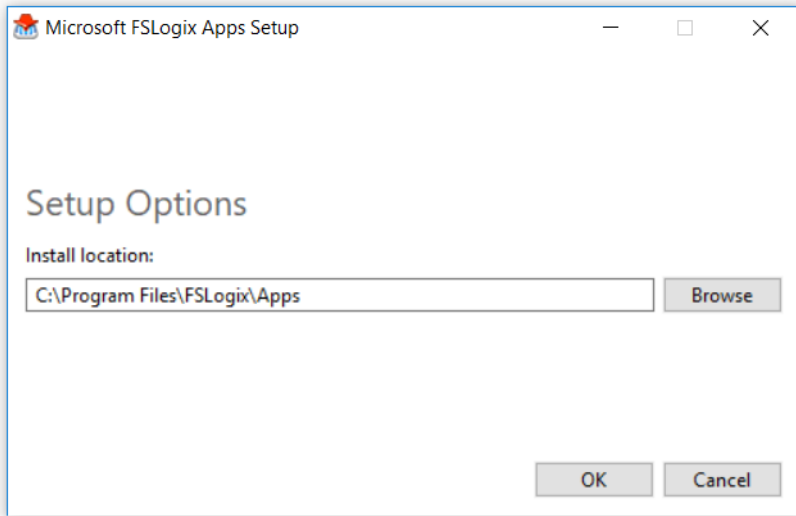
More information about the tool can be found [here](#).

Procedure 10. Agent Installation

Step 1. FSLogix download file [here](#).

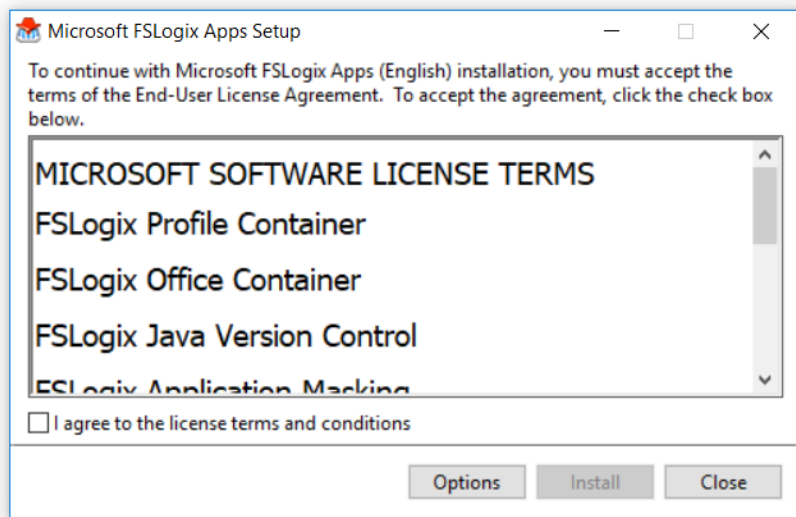
Step 2. Run FSLogixAppSetup.exe on VDI master image (32 bit or 64 bit depending on your environment).

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 54. Example of FSLogix Policy

FXLogics-H8HZN

Scope: Details Settings Delegation

VDILAB-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	
Profile type	Enabled	
	Read-write profile	
Size in MBs	Enabled	
Size in MBs	3000	
VHD location	Enabled	
VHD location	\\10.10.61.12\fp01\VDIProfiles	

FSLogix/Profile Containers/Container and Directory Naming hide

Policy	Setting	Comment
Swap directory name components	Enabled	
Swap directory name components	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 B200 M6 blade to validate against the performance of one and eight B200 M6 blades 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 325 Multi-session OS sessions and 245 Single-session Windows 10 OS sessions for linked clones and Full clone virtual machines tests.

Figure 55. Test Configuration for Single Server Scalability VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi-session OS sessions

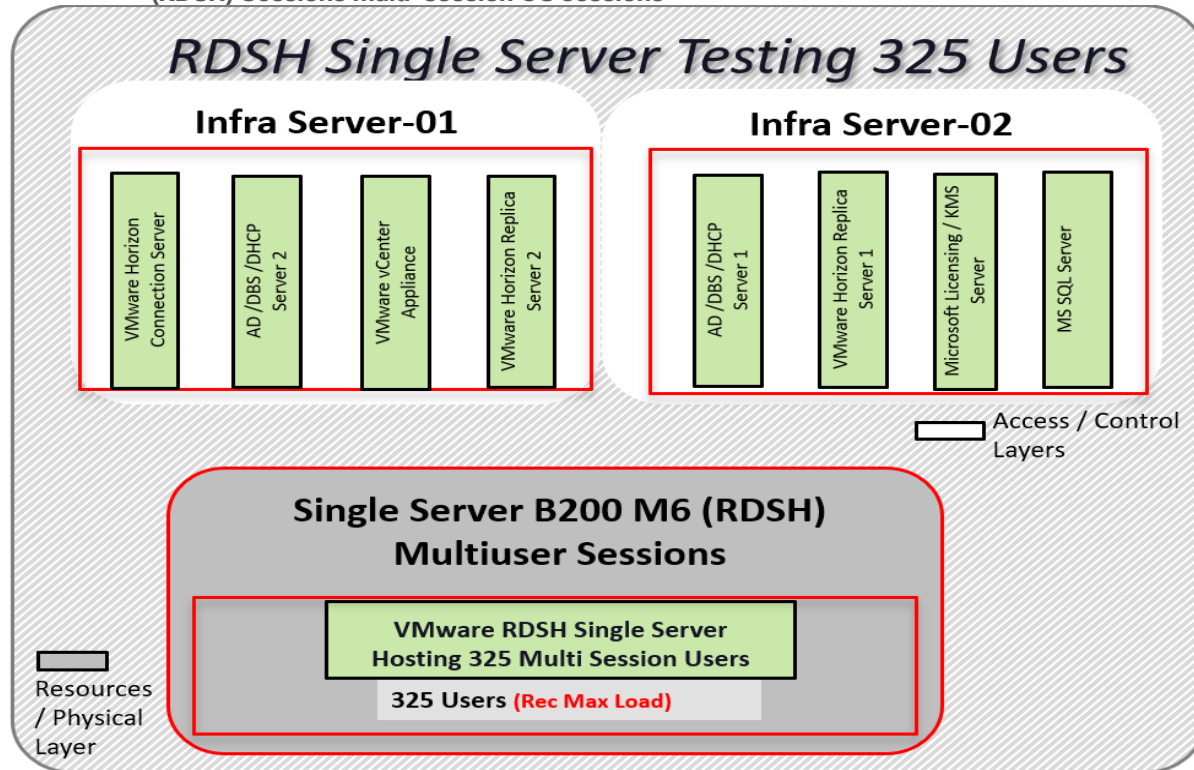


Figure 56. Test Configuration for Single Server Scalability VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi-session OS sessions on ESXi Host

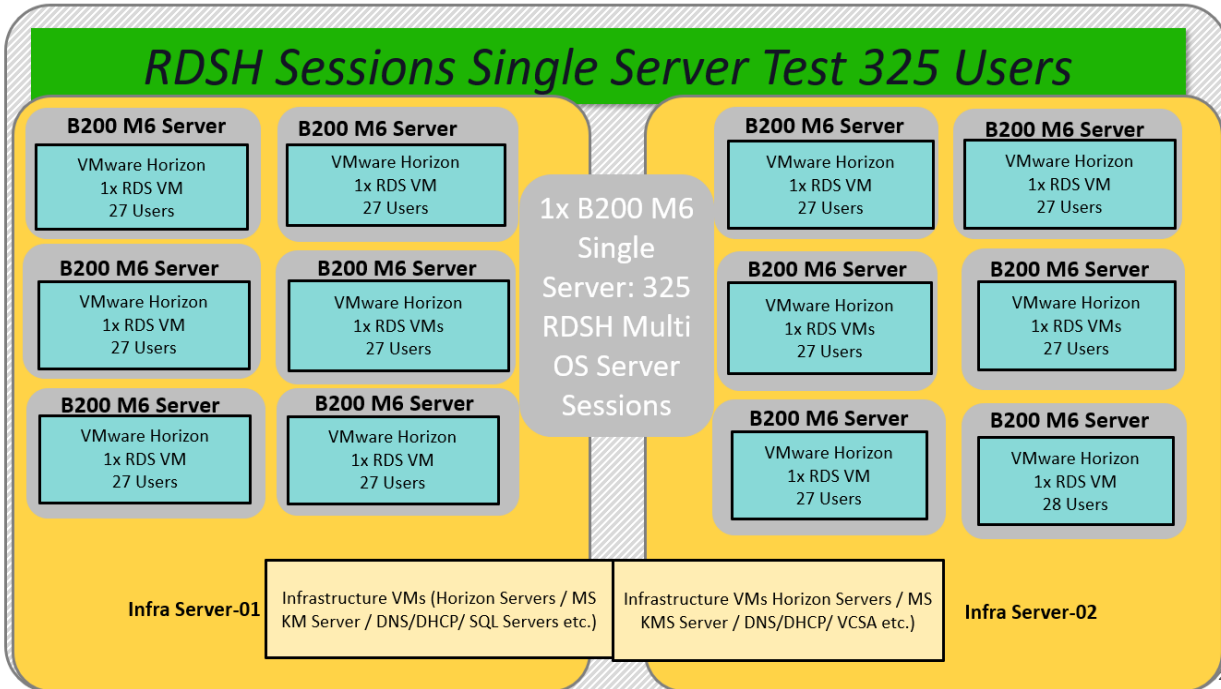
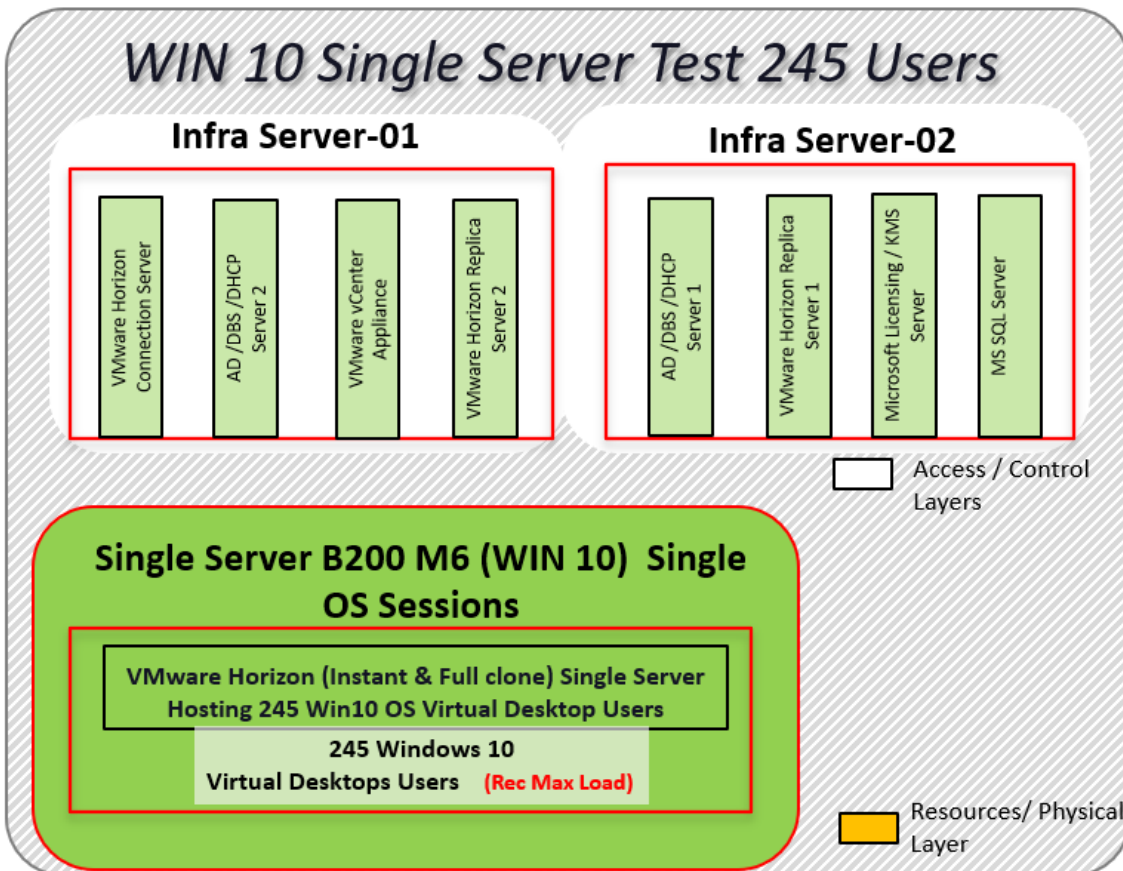


Figure 57. Test configuration for Single Server Scalability VMware Horizon WIN 10 Virtual Desktops for Instant Clones and Persistent full clones



Hardware components:

-
- Cisco UCS 5108 Blade Server Chassis
 - 2 Cisco UCS 6454 4th Gen Fabric Interconnects
 - 1 Cisco UCS B200 M6 Blade Servers with Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM for all host blades
 - Cisco VIC 1440 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 4.2(1f)
- NetApp ONTAP 9.10.1P1
- ESXi 7.0 Update 2a for host blades
- VMware Horizon 2111 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 HDD (master)
- Microsoft Windows Server 2019 (1809), 8vCPU, 32GB RAM, 60 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:

- 2300 VMware Horizon Remote Desktop Server Hosts (RDSH) sessions
- 1700 VMware Instant clone random pooled Windows 10 desktops
- 1700 VMware Full clone dedicated Windows 10 desktops

Note: Server N+1 fault tolerance is factored into this solution for each cluster/workload.

Figure 58. Test Configuration for Full Scale VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions

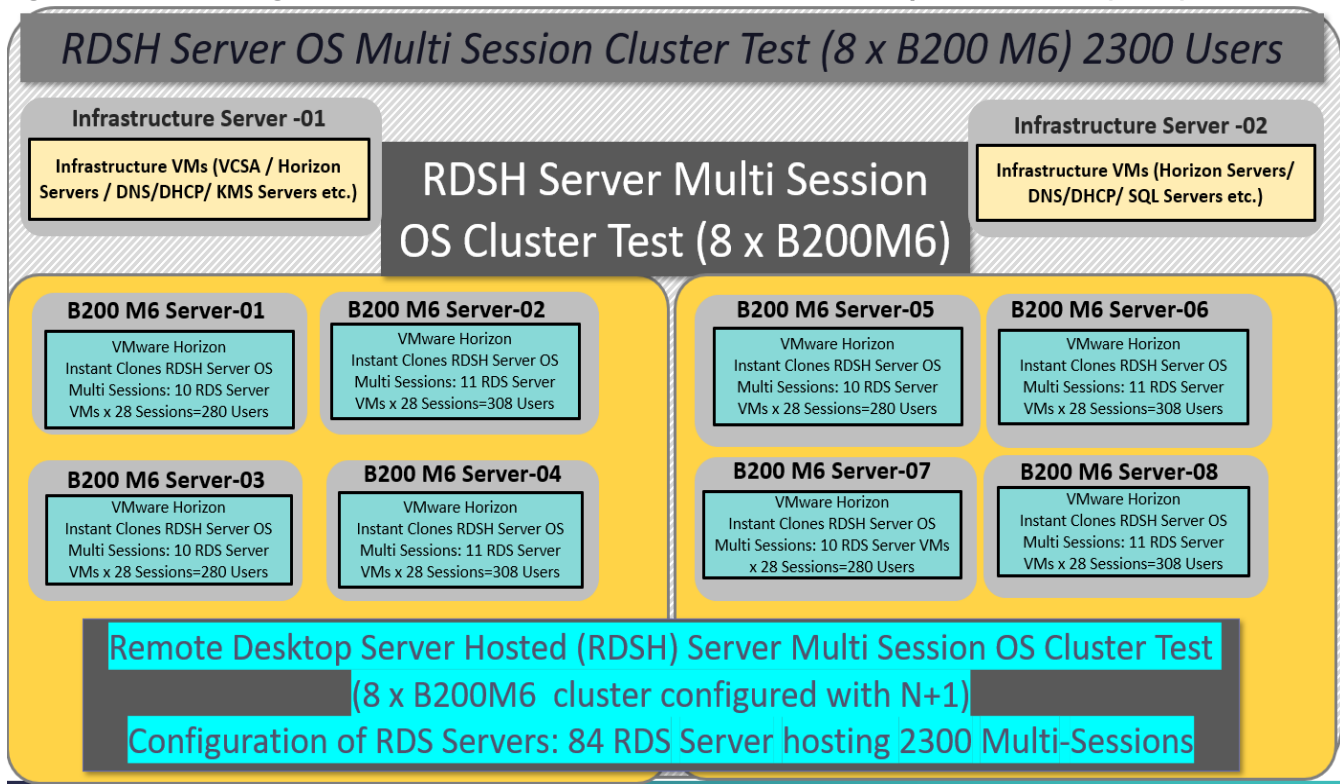


Figure 59. Test Configuration for VMware Instant Clones non-persistent Desktops Single-Session OS

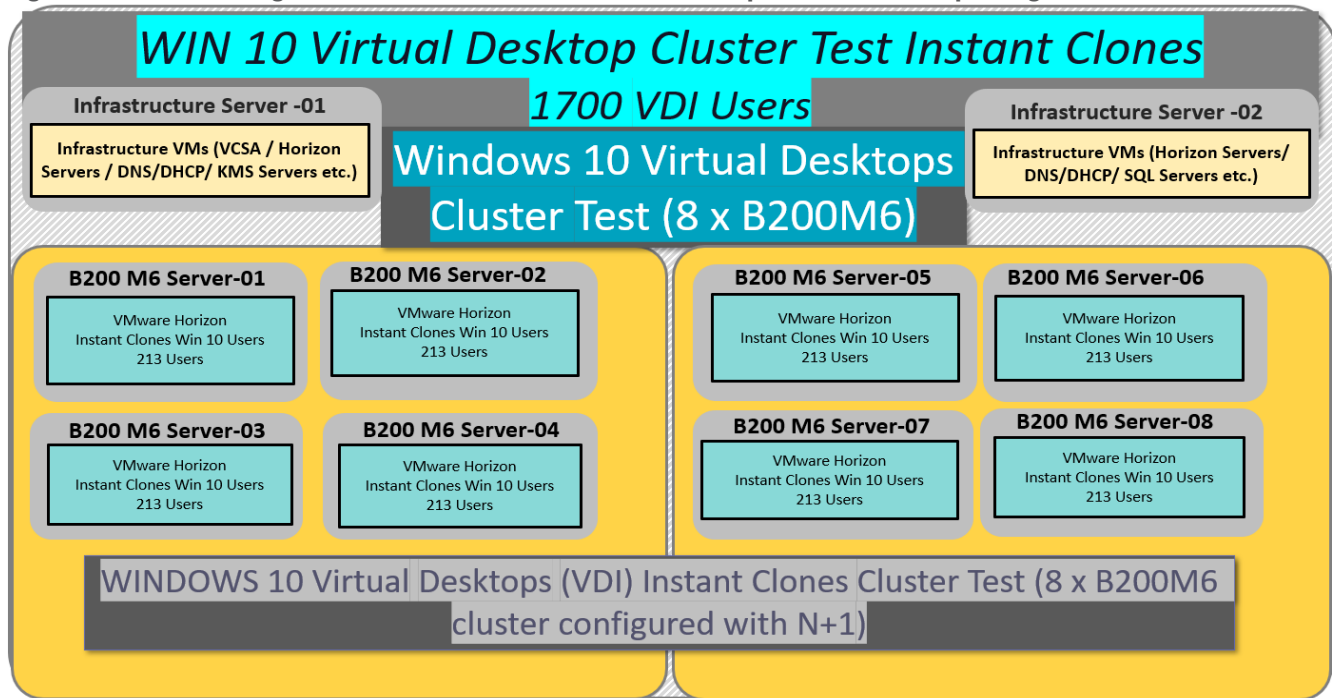
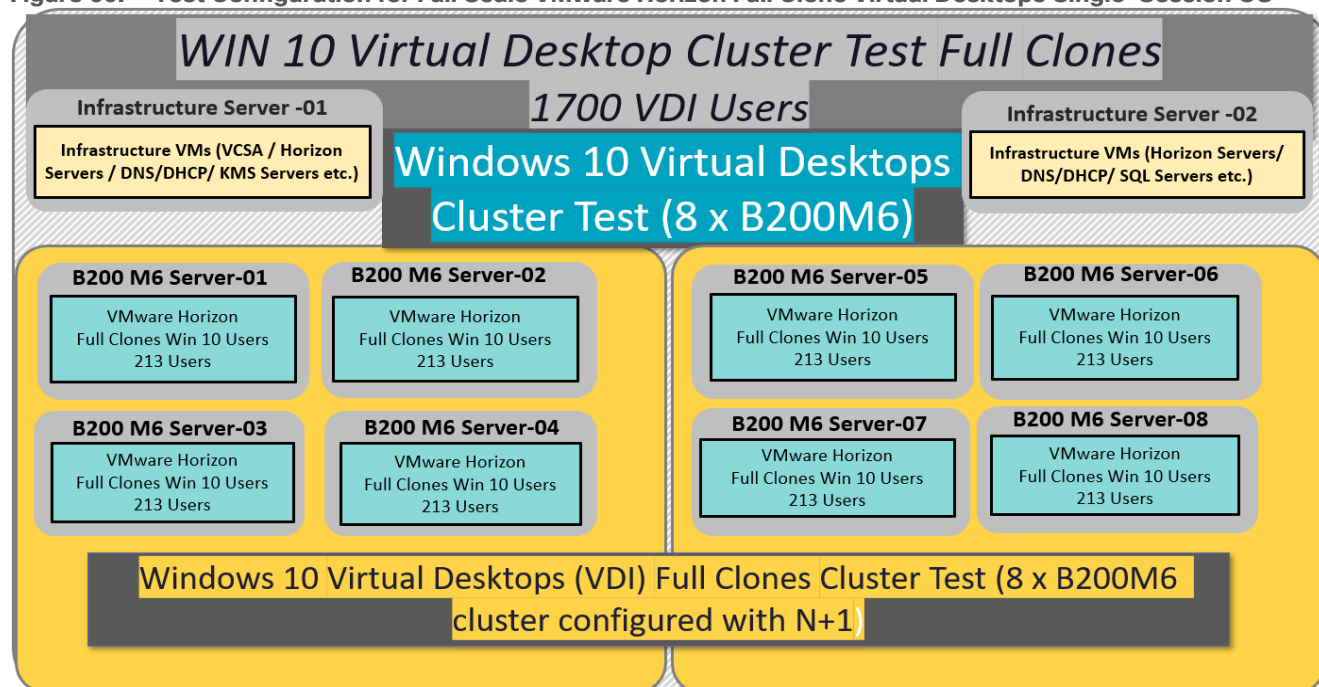


Figure 60. Test Configuration for Full Scale VMware Horizon Full Clone Virtual Desktops Single-Session OS



Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 8 Cisco UCS B200 M6 Blade Servers with Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM for all host blades
- Cisco VIC 1440 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 4.2(1f)
- NetApp ONTAP 9.10.1P1
- ESXi 7.0 Update 2a for host blades
- VMware Horizon 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 HDD (master) for virtual desktop configuration
- Microsoft Windows Server 2019 (1809), 8vCPU, 32GB RAM, 60 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/Esxtop 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 Data Center with Cisco UCS and VMware Horizon Remote Desktops and Windows 10 virtual desktops on VMware ESXi 7.0 Update 2a 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 B200 M6 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_{max} 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_{max}. 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_{max} 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 61. Sample of a VSI Max Response Time Graph, Representing a Normal Test





Figure 62. Sample of a VSI Test Response Time Graph with a Performance Issue



When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax 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 VSImax models, this weighting much better represents system performance. All actions have similar weight in the VSImax total. The following weighting of the response times is applied.

The following actions are part of the VSImax 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 VSImax 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_{max} v4.1.x is reached when the VSI_{base} + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSI_{max} 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_{max} 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_{max} 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_{max} is not hit, and the number of sessions ran successfully. This approach is fundamentally different in comparison to previous VSI_{max} methods since it is required to saturate the system beyond VSI_{max} threshold.

Lastly, VSI_{max} v4.1.x is now always reported with the average baseline VSI response time result. For example: “The VSI_{max} 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_{max} 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_{max} v4.1.x, and the higher VSI_{max} is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSI_{max} method is introduced: VSI_{max} 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 19. 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 Instant Clone Remote Desktop Server Hosted \(RDSH\) Multi Session OS Random Sessions with 325 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:

- 325 VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions (Random)
- 245 VMware Horizon Instant Clone desktops (Random)
- 245 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 325 Users

The recommended maximum workload for a Cisco UCS B200 M6 blade server with dual Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM is 325 Instant Clone Remote Desktop Server Hosted (RDSH) Multi Session OS with 8 vCPU and 32 GB RAM. The B200 M6 server 12 ran Windows Server 2019 Virtual Machines. Each virtual server was configured with 8 vCPUs and 32GB RAM.

Login VSI performance data is as follows:

Figure 63. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | VSI Score


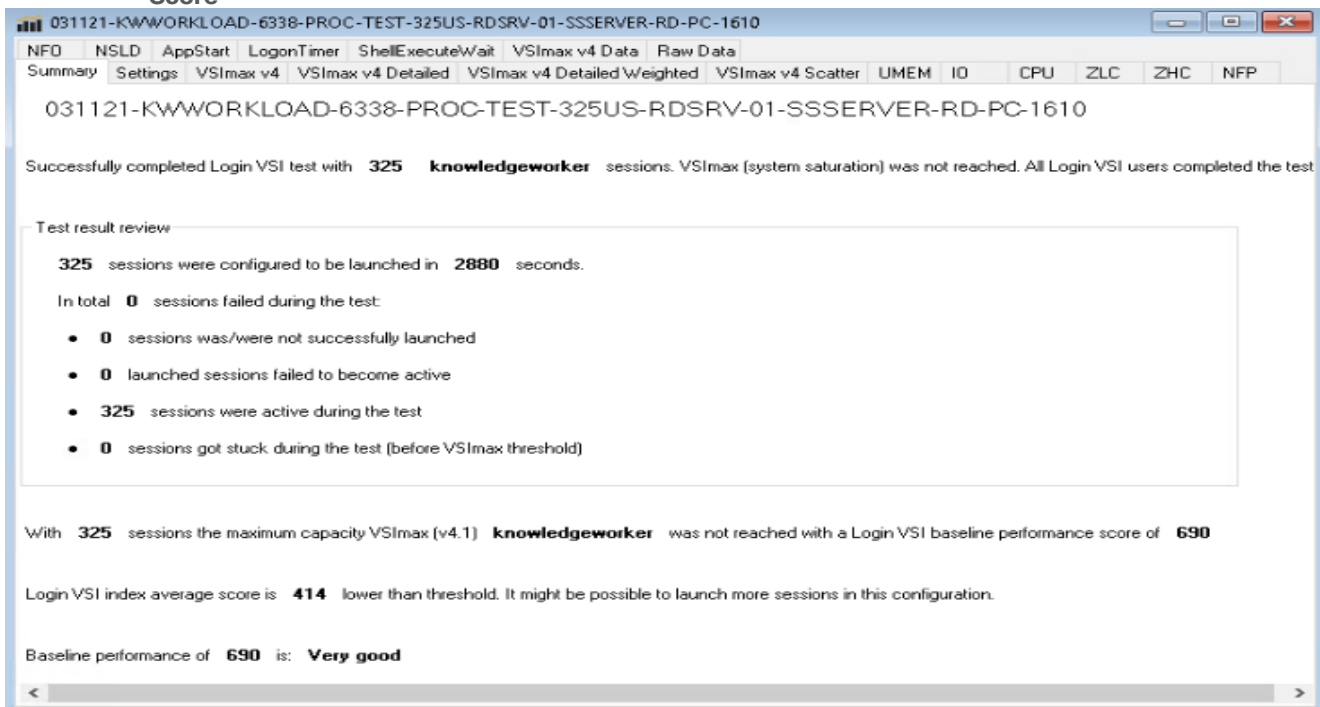


Figure 64. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | VSI Score



Performance data for the server running the workload is as follows:

Figure 65. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | Host CPU Utilization




Figure 66. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | Host Memory Utilization




Figure 67. Single Server | VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi sessions | Host Network Utilization





Figure 68. Processor Time Utilization



AFF A400 Storage Charts for 325 Remote Desktop Server Hosted (RDSH) Session User NFS Data Stores on Two Storage Controllers

Figure 69. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 1 for 325 Remote Desktop Server Hosted (RDSH) Sessions Test

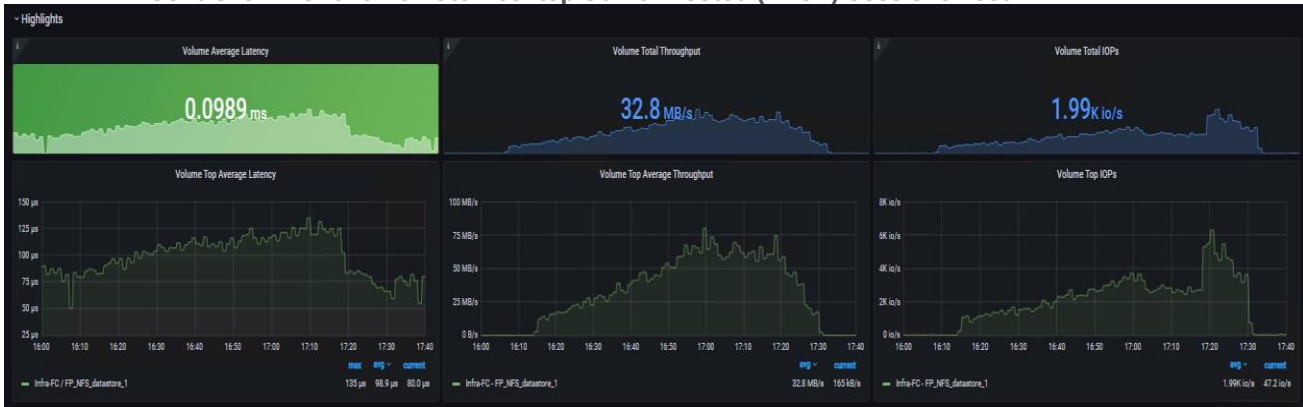


Figure 70. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 2 for 325 Remote Desktop Server Hosted (RDSH) Sessions Test



Figure 71. SVM Average latency, SVM Total Throughput and SVM Volume Total IOPS from AFF A400 Storage for 325 Remote Desktop Server Hosted (RDSH) Sessions Test



NAS Drill-Down for 325 Users Remote Desktop Server Hosted (RDSH) Sessions Test

Figure 72. NAS Drill-Down




Figure 73. FCP Drill-Down for 325 Users Remote Desktop Server Hosted (RDSH) Sessions Test



Figure 74. FCP LIF Send/Receive Throughput for 325 Users Remote Desktop Server Hosted (RDSH) Sessions Test



Figure 75. NFSv3 Front End Drill-Down for 325 Users Remote Desktop Server Hosted (RDSH) Sessions Test



The recommended maximum workload for a Cisco UCS B200 M6 blade server with dual Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM is 245 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:

Figure 76. Single Server | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | VSI Score

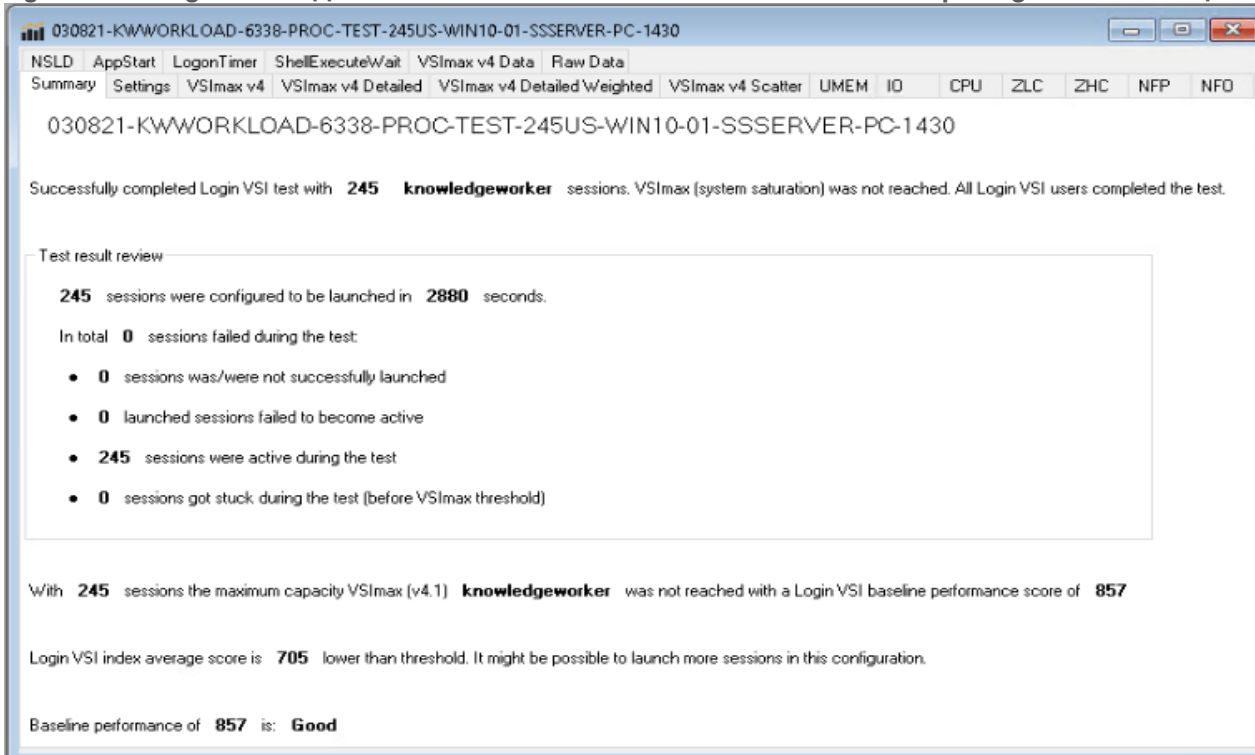



Figure 77. Single Server | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | VSI Score



Performance data for the server running the workload is as follows:

Figure 78. Single Server Recommended Maximum Workload | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host CPU Utilization




Figure 79. Single Server Recommended Maximum Workload | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host Memory Utilization





Figure 80. Single Server | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host Network Utilization



Storage Charts for 245 Users Win10 Virtual Desktops Instant Clones Testing

Figure 81. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 1 for 245 Win 10 Virtual Desktops

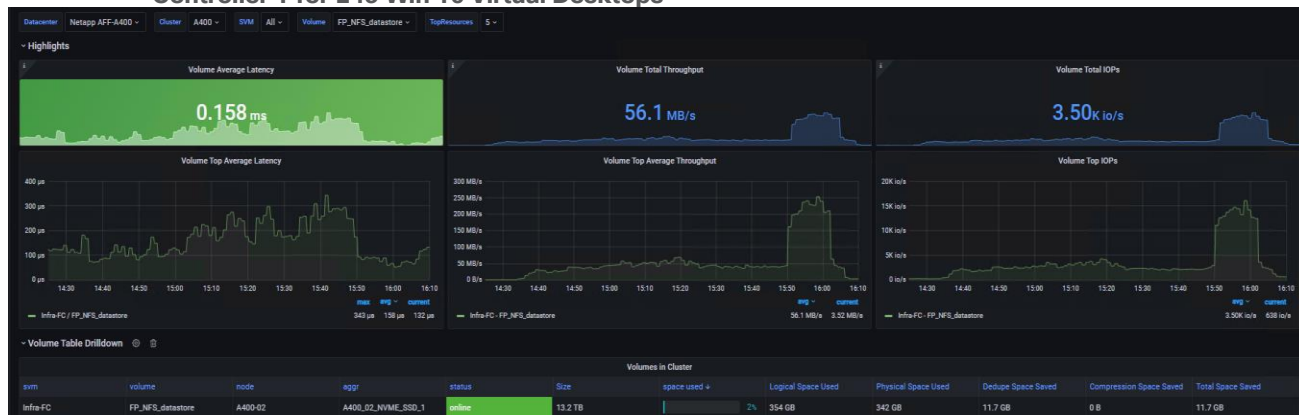


Figure 82. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage Controller 2 for 245 Win 10 Virtual Desktops



Figure 83. Volume Average latency, Volume Total Throughput and Volume Total IOPS from AFF A400 Storage for SMB Share for 245 Win 10 Virtual Desktops



Figure 84. SVM Average latency, SVM Total Thruput and SVM Total IOPS from AFF A400 Storage for 245 Win 10 Virtual Desktops



Figure 85. SVM Average latency, SVM Total Throughput and SVM Total IOPS from AFF A400 Storage for SMB Share for 245 Win 10 Virtual Desktops

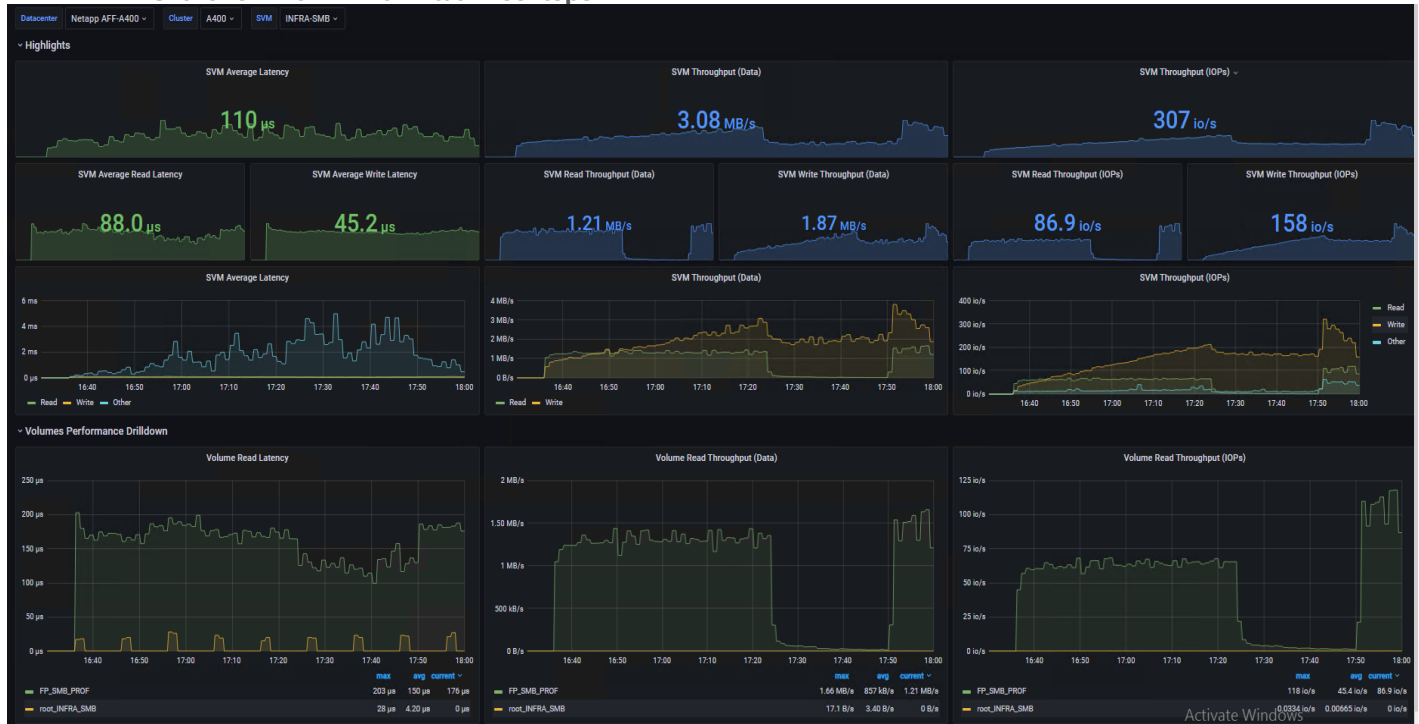



Figure 86. NAS LIF Send Throughput and NAS LIF receive Throughput AFF A400 Storage for SMB Share for 245 Win 10 Virtual Desktops



Figure 87. NFSv3 Avg Latency, NFSv3 Throughput and NFSv3 Throughput IOPS from AFF A400 Storage for SMB Share for 245 Win 10 Virtual Desktops




Figure 88. NAS Send LIF Throughput NAS LIF Receive Throughput AFF A400 Storage for SMB Share for 245 Win 10 Virtual Desktops



Note: The recommended maximum workload for a Cisco UCS B200 M6 blade server with dual Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM is 245 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 89. 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 90. Single Server Recommended Maximum Workload Single Server | | VMware Horizon Full clone Windows 10 desktops single OS sessions | Host CPU Utilization




Figure 91. Single Server Recommended Maximum Workload | | VMware Horizon Full clone Windows 10 desktops single OS sessions | Host Memory Utilization




Figure 92. Single Server | | VMware Horizon Instant Clones Windows 10 desktops single OS sessions | Host Network Utilization




Figure 93. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 1 for 245 full clone Win 10 Virtual Desktops



Figure 94. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 2 for 245 full clone Win 10 Virtual Desktops



Figure 95. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage SMB Share for 245 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 B200M6 Blade Servers, configured in a single ESXi Host Pool, and designed to support single Host failure (N+1 Fault tolerance):

- 2300 VMware Horizon Remote Desktop Server Hosted (RDSH) Sessions multi-session OS sessions
- 1700 VMware Horizon Instant clones Windows 10 Virtual Desktops
- 1700 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 2300 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 2300 VMware instant Clone Remoted Desktop Server Hosted (RDSH) Sessions using 8 B200 M6 blade s in a single pool.

The workload for the test is 1700 Non-Persistent VDI users. 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 96. Full Scale | 2300 Users | Cluster Test | VMware Horizon RDSH Sessions Instant Clones RDSH Server sessions | VSI Score

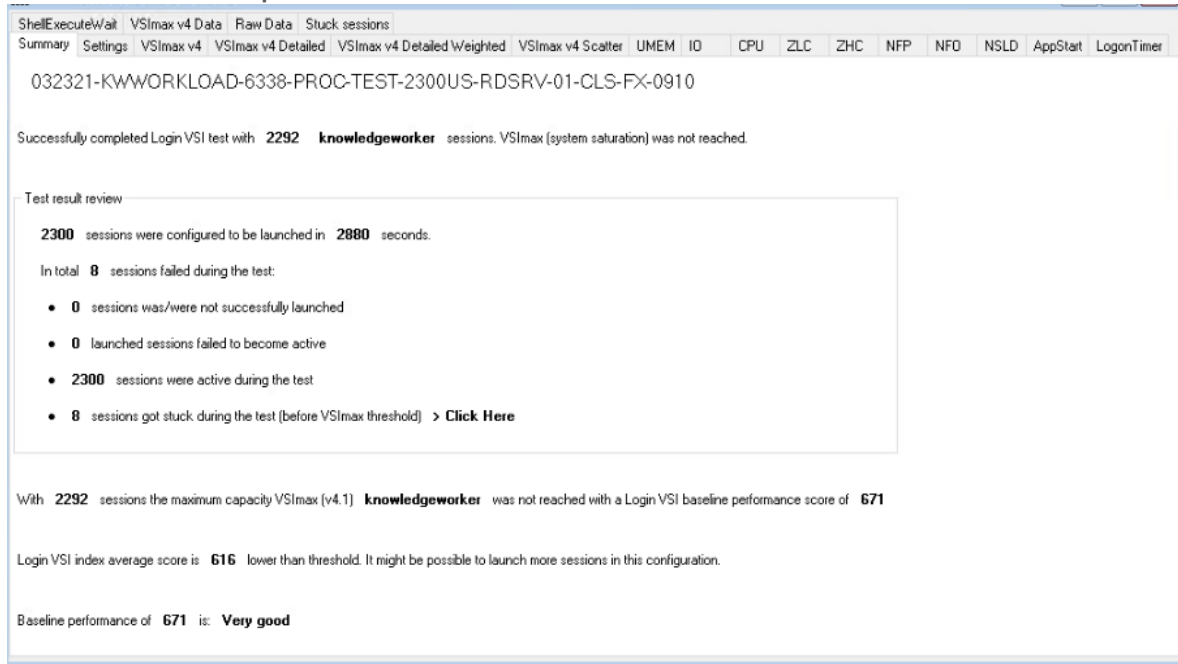


Figure 97. Full Scale | 2300 Users | Cluster Test | VMware Horizon RDSH Server Sessions Instant Clones | VSI Score




Figure 98. ESXTOP CPU Util% for 8 Hosts for 2300 Cluster Test Remote Desktop Server Hosted (RDSH) Sessions Test




Figure 99. Memory Non-Kernal Mbytes for 8 Hosts Cluster Test for 2300 Remote Desktop Server Hosted (RDSH) Sessions Test




Figure 100. Network Util for 8 Hosts Cluster Test for 2300 Remote Desktop Server Hosted (RDSH) Sessions Test




Figure 101. RDS Server VM Processor Times for 2300 RDSH Cluster Test: From RDS Server 10: Test -01





Figure 102. RDS Server VM Processor Times for 2300 RDSH Cluster Test: From RDS Server -50: Test -01



Test-02

Figure 103. RDS Server VM Processor Times for 2300 RDSH Cluster Test: From RDS Server -50 from Test-2




Figure 104. RDS Server VM Processor Times for 2300 RDSH Cluster Test: From RDS Server 80 Test -02





Figure 105. Login VSI End User Experience Comparison for 4 repeat tests for FullScale 2300 multi session RDSH Cluster Test



AFF A400 Storage Charts for 2300 Remote Desktop Server Hosted (RDSH) Sessions Cluster Full Scale Test.

Figure 106. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 1 for 2300 Instant Clone 2300 Users RDSH sessions cluster test



Figure 107. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for Storage Controller 2 for 2300 Instant Clone 2300 Users RDSH sessions cluster test



Figure 108. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage for SMB Share for 2300 Instant Clone 2300 Users RDSH sessions cluster test

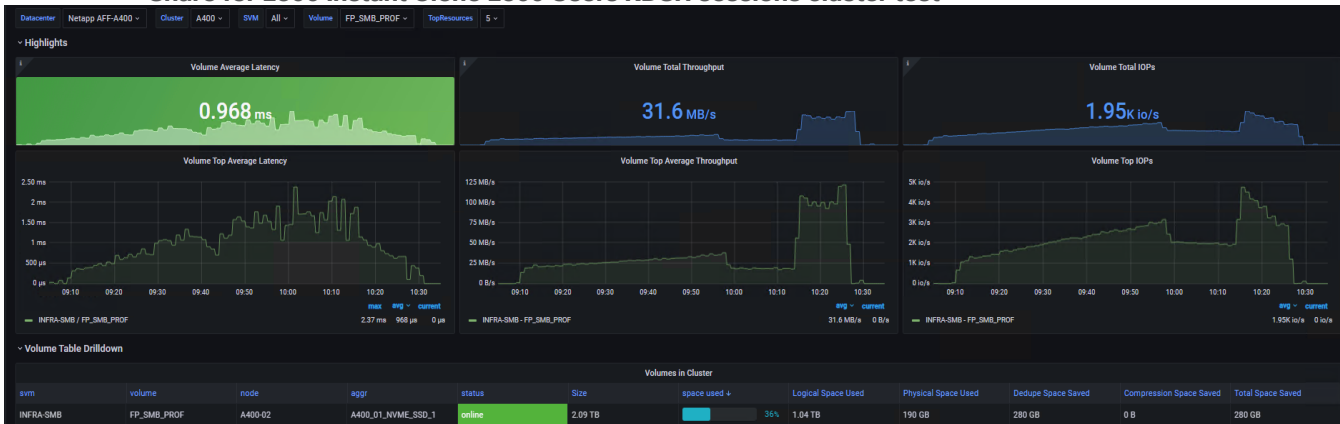


Figure 109. Full Scale | 1700 Users | VMware Horizon Win 10 Virtual Desktops Instant Clone single- OS Machines | LoginVSI VSI Score for 1700 Instant Clones Cluster Test.




Figure 110. Full Scale | 1700 Users | VMware Horizon Win 10 Virtual Desktops Instant Clone single- OS Machines | LoginVSI VSI Score for 1700 Instant Clones Cluster Test

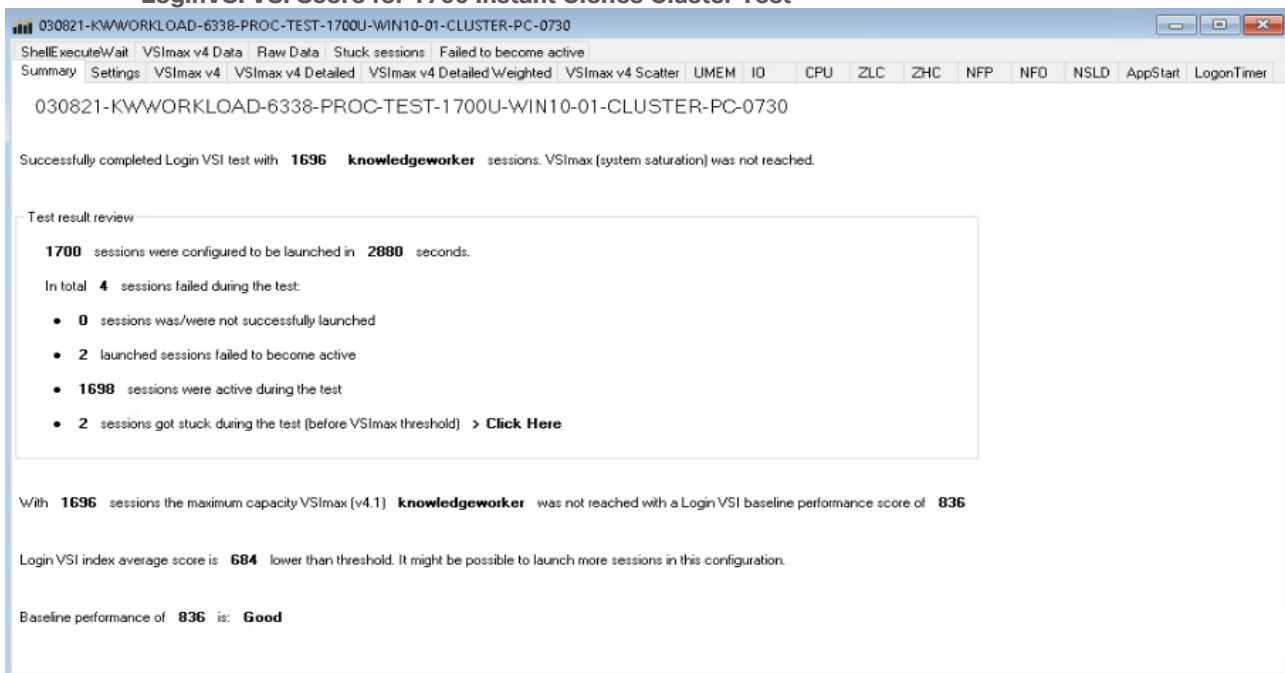


Figure 111. Login VSI End User Experience Comparison for 4 repeat tests for FullScale 1700 single session Win 10 Instant Clone Virtual Desktop Cluster Test



Figure 112. Full Scale | 1700 Users | VMware Horizon Instant Clone Virtual Desktops Single-session OS machine | ESXTOP CPU Util% for 8 hosts on Cluster




Figure 113. Full Scale | 1700 Users | VMware Horizon Instant Clone Virtual Desktops Single-session OS machine | Host Memory Utilization for 8 Hosts on Cluster





Figure 114. Full Scale | 1700 Users | VMware Horizon Windows 10 Virtual Desktops | Host Network Utilization for 8 hosts on Cluster



AFF A400 Storage Charts for VMware Horizon 1700 Windows 10 Instant Clones Virtual Desktops Cluster Full Scale Test.

Figure 115. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage Controller 1 for 1700 Instant Clone Users Cluster Test



Figure 116. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Storage Controller 2 for 1700 Instant Clone Users Cluster Test



Figure 117. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 SMB share for 1700 Instant Clone Users Cluster Test



Figure 118. Full Scale | 1700 Users | VMware Horizon Full clones single VM Login VSI End User Experience VSI Score

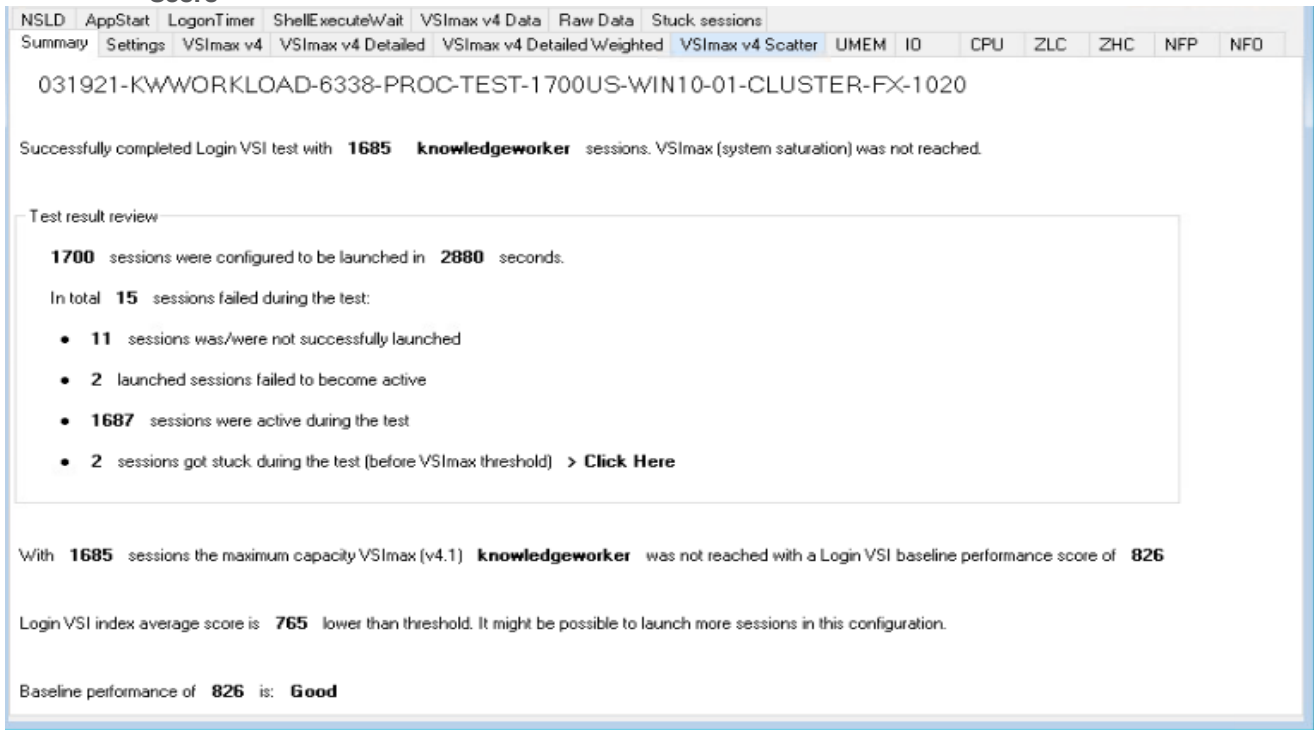


Figure 119. Full Scale | 1700 Users | VMware Horizon Full Clones Single VM Login VSI End User Experience VSI Score




Figure 120. Login VSI End User Experience for 4 Repeat Tests: Full Scale test 1700 Full Clone Win10 Virtual Desktops Comparison



Figure 121. ESXTOP CPU Util% for 1700 Full Clone Cluster Test Win10 Virtual Desktops for 8 hosts on Cluster
 ESXTOP CPU Util% Full Scale test 8 ESXi hosts 1700 Win 10 Users. Full clone cluster test




Figure 122. ESXTOP Memory util% for 1700 Full Clone Cluster Test Non-Kernel Mbytes for 8 hosts on Cluster





Figure 123. Full Scale | 1700 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 1700 Windows 10 Full Clones Virtual Desktops Cluster Full Scale Test.

Figure 124. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Controller 1 for 1700 Instant Clone Users Cluster Test



Figure 125. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 Controller 2 for 1700 Full Clone Users Cluster Test

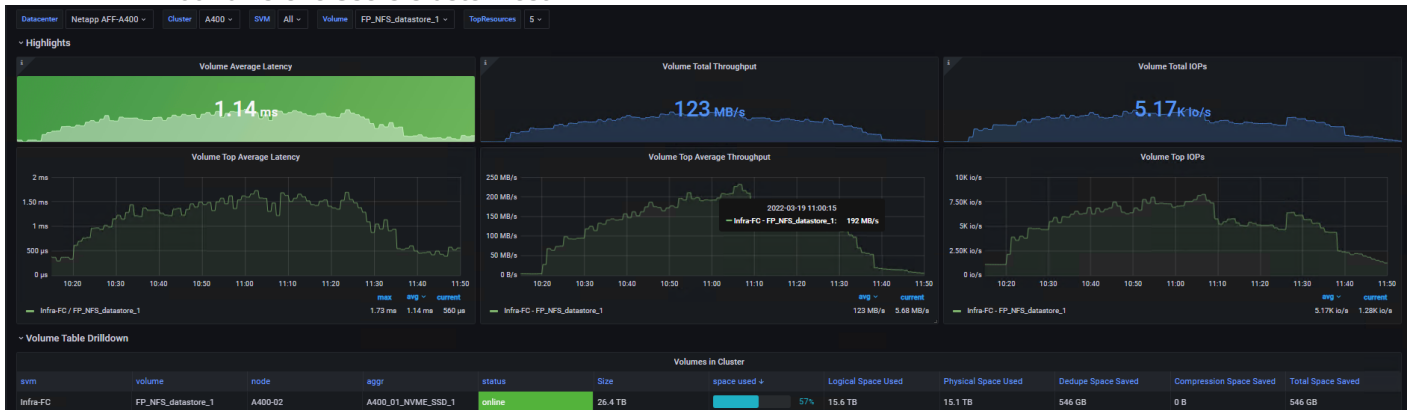


Figure 126. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 SMB Share for 1700 Full Clone Users Cluster Test



Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 2300 users, which this reference architecture has successfully tested. This 2300-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, Linked 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 20](#) as guidance for your sizing calculations for a PVS workload when using a LoginVSI heavy workload.

Table 20. 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 data centers 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's 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.

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:

- Jyh-shing Chen, Technical Marketing Engineer, NetApp
- Kamini Singh, Technical Marketing Engineer, NetApp
- Ruchika Lahoti, Technical Marketing Engineer, NetApp

References

This section provides links to additional information for each partner's solution component of this document.

Cisco UCS B200 M6 Servers

- <https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b200m6-specsheet.pdf>
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/blade-servers/B200M6/B200M5_preface_0100.html
- <https://www.intel.com/content/dam/www/central-libraries/us/en/documents/c45-2372313-00-cisco-ucs-b200-m6-blade-server-aag-v1c.pdf>

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://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%202111/rn/vmware-horizon-8-2111-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://docs.vmware.com/en/VMware-Horizon/2111/horizon-architecture-planning/GUID-C653B8FB-BCB5-4E33-BA0B-44B5FCAA2762.html>
- <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

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-6CAFE558-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/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 Feb 26 16:47:01 2016
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 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 10.29.164.65 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_e1b_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_e1b_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 Feb 7 00:49:39 2018
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$EwCSM003dlXZ4jlPy9ZoC. 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 pwwn 50:01:73:80:59:16:01:12
  device-alias name A400_N2P1 pwwn 50:01:73:80:59:16:01:20
  device-alias name A400_N2P3 pwwn 50:01:73:80:59:16:01:22
  device-alias name A400_N3P1 pwwn 50:01:73:80:59:16:01:30
  device-alias name A400_N3P3 pwwn 50:01:73:80:59:16:01:32
  device-alias name VDI-1-hba1 pwwn 20:00:00:25:b5:3a:00:3f
  device-alias name VDI-2-hba1 pwwn 20:00:00:25:b5:3a:00:0f
  device-alias name VDI-3-hba1 pwwn 20:00:00:25:b5:3a:00:1f
  device-alias name VDI-4-hba1 pwwn 20:00:00:25:b5:3a:00:4e
  device-alias name VDI-5-hba1 pwwn 20:00:00:25:b5:3a:00:2e
  device-alias name VDI-6-hba1 pwwn 20:00:00:25:b5:3a:00:3e
  device-alias name VDI-7-hba1 pwwn 20:00:00:25:b5:3a:00:0e
  device-alias name VDI-8-hba1 pwwn 20:00:00:25:b5:3a:00:4d

```

```

device-alias name Infra01-8-hbal pwnn 20:00:00:25:b5:3a:00:4f
device-alias name Infra02-16-hbal 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 pwwn 20:00:00:25:b5:3a:00:3f
!           [VDI-1-hba1]
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
zone name A400_VDI-2-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 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:00:4f
!
[Infra01-8-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
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 pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:0f
!           [VDI-2-hba1]
zone name A400_VDI-3-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:1f
!           [VDI-3-hba1]
zone name A400_VDI-4-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:4e
!           [VDI-4-hba1]
zone name A400_VDI-5-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:2e
!           [VDI-5-hba1]
zone name A400_VDI-6-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:3e
!           [VDI-6-hba1]
zone name A400_VDI-7-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [A400-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [A400-02-0g]
  member pwwn 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 fcl/24
  switchport description Launcher-FIA
  port-license acquire
  no shutdown
interface fcl/25
  port-license acquire
  no shutdown
interface fcl/26
  port-license acquire
  no shutdown
interface fcl/27
  port-license acquire
  no shutdown
interface fcl/28
  port-license acquire
  no shutdown
interface fcl/29
  port-license acquire
interface fcl/30
  port-license acquire
interface fcl/31
  port-license acquire
interface fcl/32
  port-license acquire
interface fcl/33
  port-license acquire
interface fcl/34
  port-license acquire
interface fcl/35
  port-license acquire
interface fcl/36
  port-license acquire
interface fcl/37
  switchport trunk mode off
  port-license acquire
  no shutdown
interface fcl/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—Data Center

DHCP—Dynamic Host Configuration Protocol

DM—Dense-Mode (multicast)

DMVPN—Dynamic Multipoint Virtual Private Network

DMZ—Demilitarized Zone (firewall/networking construct)

DNA—Cisco Digital Network Architecture

DNS—Domain Name System

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

DWDM—Dense Wavelength Division Multiplexing

ECMP—Equal Cost Multi Path

EID—Endpoint Identifier

EIGRP—Enhanced Interior Gateway Routing Protocol

EMI—Electromagnetic Interference

ETR—Egress Tunnel Router (LISP)

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

FHR—First-Hop Router (multicast)

FHRP—First-Hop Redundancy Protocol

FMC—Cisco Firepower Management Center

FTD—Cisco Firepower Threat Defense

GBAC—Group-Based Access Control

GbE—Gigabit Ethernet

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

GRE—Generic Routing Encapsulation

GRT—Global Routing Table

HA—High-Availability

HQ—Headquarters

HSRP—Cisco Hot-Standby Routing Protocol

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

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

ICMP—Internet Control Message Protocol

IDF—Intermediate Distribution Frame; essentially a wiring closet.

IEEE—Institute of Electrical and Electronics Engineers

IETF—Internet Engineering Task Force

IGP—Interior Gateway Protocol

IID—Instance-ID (LISP)

IOE—Internet of Everything

IoT—Internet of Things

IP—Internet Protocol

IPAM—IP Address Management

IPS—Intrusion Prevention System

IPSec—Internet Protocol Security

ISE—Cisco Identity Services Engine

ISR—Integrated Services Router

IS-IS—Intermediate System to Intermediate System routing protocol

ITR—Ingress Tunnel Router (LISP)

LACP—Link Aggregation Control Protocol

LAG—Link Aggregation Group

LAN—Local Area Network

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

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

LHR—Last-Hop Router (multicast)

LISP—Location Identifier Separation Protocol

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

MAN—Metro Area Network

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

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

MnT—Monitoring and Troubleshooting Node (Cisco ISE persona)

MOH—Music on Hold

MPLS—Multiprotocol Label Switching

MR—Map-resolver (LISP)

MS—Map-server (LISP)

MSDP—Multicast Source Discovery Protocol (multicast)

MTU—Maximum Transmission Unit

NAC—Network Access Control

NAD—Network Access Device

NAT—Network Address Translation

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

NFV—Network Functions Virtualization

NSF—Non-Stop Forwarding

OSI—Open Systems Interconnection model

OSPF—Open Shortest Path First routing protocol

OT—Operational Technology

PAgP—Port Aggregation Protocol

PAN—Primary Administration Node (Cisco ISE persona)

PCI DSS—Payment Card Industry Data Security Standard

PD—Powered Devices (PoE)

PETR—Proxy-Egress Tunnel Router (LISP)

PIM—Protocol-Independent Multicast

PITR—Proxy-Ingress Tunnel Router (LISP)

PnP—Plug-n-Play

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

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

PSE—Power Sourcing Equipment (PoE)

PSN—Policy Service Node (Cisco ISE persona)

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

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

QoS—Quality of Service

RADIUS—Remote Authentication Dial-In User Service

REST—Representational State Transfer

RFC—Request for Comments Document (IETF)

RIB—Routing Information Base

RLOC—Routing Locator (LISP)

RP—Rendezvous Point (multicast)

RP—Redundancy Port (WLC)

RP—Route Processer

RPF—Reverse Path Forwarding

RR—Route Reflector (BGP)

RTT—Round-Trip Time

SA—Source Active (multicast)

SAFI—Subsequent Address Family Identifiers (BGP)

SD—Software-Defined

SDA—Cisco Software Defined-Access

SDN—Software-Defined Networking

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

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

SGACL—Security-Group ACL

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

SM—Spare-mode (multicast)

SNMP—Simple Network Management Protocol

SSID—Service Set Identifier (wireless)

SSM—Source-Specific Multicast (PIM)

SSO—Stateful Switchover

STP—Spanning-tree protocol

SVI—Switched Virtual Interface

SVL—Cisco StackWise Virtual

SWIM—Software Image Management

SXP—Scalable Group Tag Exchange Protocol

Syslog—System Logging Protocol

TACACS+—Terminal Access Controller Access-Control System Plus

TCP—Transmission Control Protocol (OSI Layer 4)

UCS— Cisco Unified Computing System

UDP—User Datagram Protocol (OSI Layer 4)

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

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

URL—Uniform Resource Locator

VLAN—Virtual Local Area Network

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>aaS/XaaS (IT capability provided as a Service)</p>	<p>Some IT capability, X, provided as a service (XaaS). Some benefits are:</p> <ul style="list-style-type: none">• The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it.• 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.• The service charge is an IT OpEx cost (pay-as-you-go), whereas the CapEx and the service infrastructure is the responsibility of the provider.• Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes. <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>Ansible</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>https://www.ansible.com</p>
<p>AWS (Amazon Web Services)</p>	<p>Provider of IaaS and PaaS.</p> <p>https://aws.amazon.com</p>
<p>Azure</p>	<p>Microsoft IaaS and PaaS.</p>



	https://azure.microsoft.com/en-gb/
Co-located data center	<p>“A colocation center (CoLo)...is a type of data center where equipment, space, and bandwidth are available for rental to retail customers. Colocation facilities provide space, power, cooling, and physical security for the server, storage, and networking equipment of other firms and also connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity.”</p> <p>https://en.wikipedia.org/wiki/Colocation_centre</p>

Containers (Docker)	<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>https://www.docker.com</p> <p>https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html</p>
DevOps	<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>https://en.wikipedia.org/wiki/DevOps</p> <p>https://en.wikipedia.org/wiki/CI/CD</p>
Edge compute	<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>https://en.wikipedia.org/wiki/Mobile_edge_computing</p>
IaaS (Infrastructure as-a-Service)	<p>Infrastructure components provided aaS, located in data centers operated by a provider, typically accessed over the public Internet. IaaS provides a base platform for the deployment of workloads, typically with containers and Kubernetes (K8s).</p>
IaC (Infrastructure as-Code)	<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>https://en.wikipedia.org/wiki/Infrastructure_as_code</p>
IAM (Identity and Access Management)	<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>https://en.wikipedia.org/wiki/Identity_management</p>
IBM (Cloud)	<p>IBM IaaS and PaaS.</p> <p>https://www.ibm.com/cloud</p>
Intersight	<p>Cisco Intersight™ is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.</p> <p>https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html</p>

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

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 Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_P5)

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)